4th International Conference on Electronics, Engineering, Computer Science and Applied Development (EESD 2026)

Article

Development of a Generative Art App for Individuals with Limb Loss: Exploring Mirror Visual Feedback and Tactile Substitution as Complementary Therapeutic Modalities

Qiqi Wu[1,*]

[1] Qingdao Huanghai University, Qingdao, China
* Correspondence: Qiqi Wu, Qingdao Huanghai University, Qingdao, China

Abstract: Individuals with limb loss may experience altered sensorimotor perception, phantom limb sensations, and changes in body image, creating a need for accessible complementary forms of interaction. This study develops a generative art app integrating digitally reconstructed mirror visual feedback with vibrotactile substitution. Seven representative movement states were selected from NinaPro DB3, and ten consistently available surface electromyography channels were used as input. A lightweight one-dimensional convolutional neural network recognized movement states and drove a mirrored digital-limb representation. A constrained generative engine mapped movement class, intensity, and confidence to particle, color, trajectory, and line parameters, while a tactile encoder converted movement events into distinguishable vibration patterns. Subject-independent evaluation compared the proposed classifier with support vector machine and random forest baselines. Prototype testing further examined visual and haptic latency, frame rate, memory consumption, and tactile trigger accuracy across four feedback configurations. The evaluation indicates that movement recognition, mirror animation, generative graphics, and vibration commands can be coordinated within a responsive mobile framework. The findings concern technical feasibility and cross-modal synchronization rather than clinical efficacy. Future studies should recruit individuals with limb loss to assess usability, sustained engagement, body-image experience, and clinically supervised outcomes.

Keywords: limb loss; generative art; mirror visual feedback; tactile substitution; mobile health

Received: 08 July 2026
Revised: 18 August 2026
Accepted: 30 August 2026
Published: 05 September 2026

1. Introduction

1.1. Background

Limb loss can have consequences extending beyond physical impairment, including phantom limb sensations, phantom limb pain, altered body image, and difficulties in adapting to changes in sensorimotor experience [1]. Mirror visual feedback has therefore been explored as a non-invasive approach for presenting the visual illusion that a missing limb is present and moving. In conventional mirror-based practice, movements of the intact limb are reflected symmetrically and visually attributed to the missing side [2]. This visual correspondence may help users engage with imagined bilateral movements, although reported outcomes vary across intervention protocols, participant groups, and assessment periods.

Recent mobile technologies create opportunities to extend this principle beyond a physical mirror box. Smartphone cameras can capture intact-limb movements, real-time graphics can transform movement trajectories into dynamic visual content, and embedded vibration motors can communicate interaction events through tactile cues [3].

Generative art is particularly relevant because movement direction, speed, amplitude, and continuity can be mapped to changes in particles, colors, lines, and textures. These capabilities allow a mobile application to combine mirrored movement representation with personalized visual expression and non-invasive tactile substitution. Such an application should not be positioned as a replacement for clinical assessment, medication, rehabilitation, or psychological care [4,5]. Instead, it can be examined as a low-cost and accessible complementary modality that may support home-based engagement under appropriate professional guidance.

1.2. Research Gaps

Despite these opportunities, three limitations remain. First, many mirror-feedback interfaces rely on relatively fixed visual representations. Repeated exposure to an unchanged mirrored image may provide limited opportunities for creative expression and sustained interaction. Second, visual and tactile feedback are frequently developed as separate functions. Without a shared encoding mechanism, a visual event and its corresponding vibration may be temporally or semantically inconsistent, increasing the user’s perceptual burden. Third, although generative graphics have been used in digital art and well-being applications, their suitability for limb-loss-oriented interaction remains insufficiently evaluated. Highly abstract graphics may conceal the mirrored limb, while computationally intensive generation may introduce latency and weaken action–feedback correspondence.

1.3. Research Questions

Accordingly, this study addresses three research questions:

RQ1: Can an open dataset collected from individuals with limb loss support mobile recognition of representative upper-limb movements?

RQ2: Can generative art enhance visual expression while preserving the recognizability of mirrored movement contours?

RQ3: Can limb movements, visual events, and vibration patterns be mapped synchronously with sufficiently low latency for real-time interaction?

1.4. Contributions

This study makes three contributions. First, it proposes a mobile application architecture integrating mirror visual feedback, movement-responsive generative art, and vibrotactile substitution. Second, it establishes a cross-modal encoding mechanism that maps recognized movement classes and movement intensity to visual parameters and distinguishable vibration patterns. Third, it uses the NinaPro DB3 dataset and prototype-level testing to evaluate movement recognition, visual rendering, and real-time feedback performance. The resulting evidence concerns technical feasibility and interaction consistency rather than clinical efficacy, providing a reproducible foundation for subsequent user-centered and clinically supervised studies.

2. Related Work and Design Foundations

2.1. Mirror Visual Feedback for Limb Loss

Conventional mirror visual feedback uses a mirror box positioned along the user’s body midline. The intact limb is placed on the reflective side, while the residual limb or missing side remains concealed [6]. When the intact limb moves, its reflection creates the visual impression of symmetrical bilateral movement. This approach has been investigated in relation to phantom limb sensations, phantom limb pain, motor imagery, and perceived phantom-limb control.

Digital mirror systems extend this principle through cameras, computer vision, virtual limbs, augmented reality, or virtual reality [7]. Compared with a physical mirror, these systems can modify viewpoint, movement range, task difficulty, and visual content while recording interaction data and providing guided exercises. However, evidence does not show uniform effects across users. Studies differ in amputation level, pain history, intervention frequency, duration, control conditions, and outcome measures. Small samples and short follow-up periods further limit generalization [8]. Therefore, this study neither assesses pain reduction nor claims equivalence to clinical mirror therapy. It focuses on implementing recognizable and temporally consistent digital mirror feedback.

2.2. Tactile Sensory Substitution

Many visually controlled assistive systems provide limited information about contact, movement completion, or interaction intensity [9]. Users may therefore depend heavily on vision, increasing attention demands when visual information is incomplete. Sensory substitution addresses this limitation by converting information associated with a missing sensory channel into signals perceived through another available body location.

Non-invasive approaches include vibrotactile, electrotactile, and mechanical pressure feedback [10]. Vibrotactile feedback is relatively safe, controllable, and compatible with compact devices. Electrotactile feedback provides distinguishable patterns but requires careful calibration of current intensity and electrode placement. Pressure-based systems may provide more spatially intuitive cues but generally require additional hardware [11]. By contrast, a smartphone vibration motor offers a low-cost means of communicating movement initiation, target completion, or deviation. Here, tactile substitution refers to encoding digital interaction events as vibration patterns perceived by an available skin area. It does not restore physiological touch or reproduce natural sensation in the missing limb [12].

2.3. Generative Art in Digital Health

Generative art uses predefined rules or computational models to produce visual content that responds to input data [13,14]. In a movement-based health application, movement class, direction, amplitude, and confidence can control colors, trajectories, particle density, line width, or texture. This transformation makes repeated interaction less static and allows users to create individualized visual outcomes rather than only follow standardized instructions.

However, visual novelty does not necessarily improve therapeutic relevance or usability. Excessive abstraction may conceal the mirrored limb, weaken action–feedback correspondence, or introduce distracting complexity [15]. Unrestricted image-generation models may also produce unpredictable outputs and delays incompatible with real-time sensorimotor interaction. The proposed application therefore uses constrained generation: the mirrored limb contour remains visible, while generative effects are limited to predefined parameters surrounding or following its trajectory.

2.4. Research Positioning

This study is positioned at the intersection of mobile interaction, digital mirror feedback, generative art, and non-invasive sensory substitution. Its objective is neither to improve movement-classification accuracy as an isolated machine-learning task nor to demonstrate the clinical effectiveness of mirror therapy. Instead, it investigates how movement recognition, mirrored representation, constrained generative graphics, and vibrotactile encoding can be integrated into a functioning mobile prototype. The evaluation consequently emphasizes recognition performance, visual continuity, response latency, and cross-modal synchronization, establishing a technical basis for future studies involving target users and clinically supervised outcomes.

3. App Design and Methodology

3.1. Design Requirements

The application was designed according to four requirements. First, recognizability requires the mirrored limb contour and movement direction to remain visible during interaction. Generative elements are therefore rendered around or behind the digital limb rather than replacing it with a fully abstract image. Second, real-time responsiveness requires low latency between movement input, classification, visual rendering, and vibration activation. Third, cross-modal consistency requires repeated instances of the same movement to produce semantically stable visual and tactile patterns. Colors and particle density may vary within bounded ranges, but their associations with movement classes remain unchanged. Fourth, safety requires the interface to avoid rapid flashing, excessive contrast, and continuous high-intensity vibration. The application also states that it is a complementary interaction prototype rather than a substitute for clinical assessment or treatment.

3.2. Overall App Architecture

Figure 1 presents the five-layer application architecture. The data/input layer receives sEMG signals corresponding to those provided by NinaPro DB3. In practical deployment, these signals would require a compatible external wearable sEMG sensor connected to the smartphone. An optional camera input is reserved for intact-limb motion capture in future live interaction but is not used to train or evaluate the signal-based classifier in this study. The preprocessing layer performs signal filtering, normalization, sliding-window segmentation, and feature construction. The resulting windows enter the movement recognition layer, where a lightweight one-dimensional convolutional neural network (1D-CNN) estimates the movement class, confidence score, and movement intensity.

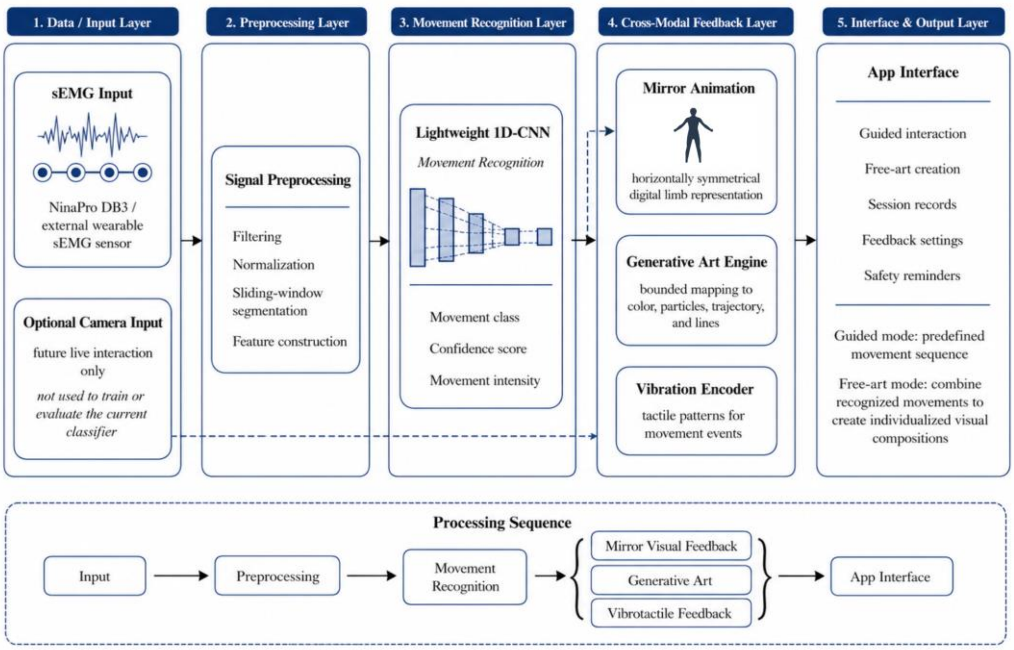

Figure 1. Overall Architecture of the Generative Art App Integrating Mirror Visual Feedback and Tactile Substitution

The recognition outputs are distributed to three components in the cross-modal feedback layer. The mirror animation component generates a horizontally symmetrical digital limb representation. The generative art engine converts movement characteristics into bounded visual parameters, while the vibration encoder assigns tactile patterns to movement events. The interface and output layer provides guided interaction, free-art creation, session records, feedback settings, and safety reminders. Guided mode presents a predefined sequence of movements, whereas free-art mode allows users to combine recognized movements to construct an individualized visual composition.

3.3. Dataset and Movement Selection

NinaPro DB3 contains recordings from 11 individuals with transradial limb loss who attempted up to 49 hand movements in addition to rest. The acquisition used up to 12 Delsys Trigno sEMG electrodes sampled at 2 kHz. The database also provides accelerometer signals, uncalibrated CyberGlove measurements, and force records for specific exercises. These variables were excluded because they were not consistently available or directly comparable across the selected movement states. The experiments therefore used sEMG as the sole model input.

Subjects 7 and 8 were recorded with 10 rather than 12 sEMG electrodes because of limited residual-limb space. Therefore, the ten sEMG channels consistently available across participants were retained as the model input. Accelerometer, CyberGlove, and force signals were excluded from the experiments to maintain a consistent input configuration across participants and movement states. This design also avoided treating unavailable sEMG channels as measured zero activity.

Seven states were selected according to their relevance to simple mirror-based interaction: hand opening, hand closing, wrist flexion, wrist extension, pronation, supination, and rest. The limited movement set reduced interface complexity and provided visually distinguishable actions for generative mapping. Because several participants discontinued parts of the acquisition owing to fatigue or discomfort, the availability of movement repetitions was checked separately in each validation fold (As shown in Table 1).

Table 1. Dataset Characteristics, Selected Movements, and Experimental Settings

Category	Content
Dataset	NinaPro DB3
Participants	11 individuals with transradial limb loss
Acquisition	Up to 49 attempted movements plus rest
Primary input	Ten consistently available sEMG channels
Sampling rate	2 kHz
Selected states	Open, close, flexion, extension, pronation, supination, and rest
Segmentation	200-ms windows with 50% overlap
Primary model	Lightweight 1D-CNN
Baselines	Support vector machine and random forest
Validation	Leave-one-subject-out evaluation
Metrics	Accuracy, Macro-F1, inference latency, and feedback latency
Feedback outputs	Mirrored animation, generative graphics, and vibration

3.4. Signal Preprocessing and Movement Recognition

The sEMG channels were band-pass filtered between 20 and 450 Hz, followed by power-line interference removal. Each channel was standardized using statistics calculated exclusively from the training participants:

$$\hat{x}_{i,t} = \frac{x_{i,t} - \mu_i^{train}}{\sigma_i^{train} + \epsilon}, \quad (1)$$

where $x_{i,t}$ is the value of channel i at time t, μ_i^{train} and σ_i^{train} are the corresponding training-set mean and standard deviation, and ϵ prevents division by zero. The normalized signals were divided into 200-ms windows with 50% overlap. Window-level root-mean-square amplitude was additionally calculated as an indicator of movement intensity.

The primary classifier contained three one-dimensional convolutional blocks with 32, 64, and 128 filters. Each block used batch normalization, rectified linear unit activation, and temporal pooling. Global average pooling was followed by a fully connected output layer, with dropout applied during training. The probability of movement class k was calculated as

$$p_{t,k} = \frac{\exp(z_{t,k})}{\sum_{j=1}^{K} \exp(z_{t,j})}, c_t = \arg\max_k p_{t,k}, q_t = \max_k p_{t,k}, \quad (2)$$

where $z_{t,k}$ is the class logit, K is the number of states, c_t is the predicted class, and q_t is its confidence. A support vector machine and random forest were implemented as baselines using the same input channels, movement classes, and participant partitions.

A leave-one-subject-out protocol was adopted. In each fold, one participant was retained for testing, while the remaining participants were used for model development. Validation data were selected only from the development participants. Thus, no signal windows from the test participant appeared in training, reducing participant-identity leakage.

3.5. Mirror Visual Feedback

The predicted movement was applied to an articulated digital limb model and horizontally reflected toward the side corresponding to the missing limb. Joint relationships, movement direction, and the external limb contour were retained to maintain visual recognizability. The predicted sequence was temporally smoothed to prevent rapid transitions between adjacent classes.

If q_t fell below the confidence threshold, the application briefly retained the most recent stable state and then returned the model to rest. Low-confidence predictions therefore did not generate persistent incorrect movements. A future camera-supported mode could capture landmarks from the intact limb and generate a direct visual reflection without sEMG acquisition. However, the present prototype uses sEMG-based movement recognition to drive the digital limb, and camera-derived data are not included in the reported evaluation.

3.6. Constrained Generative Art Mapping

The generative art engine mapped each recognized movement to a bounded visual parameter vector:

$$a_t = g(c_t, v_t, r_t, q_t; \Theta_a), \quad (3)$$

where c_t denotes movement class, v_t is movement velocity, r_t is movement intensity derived from normalized sEMG amplitude, q_t is classification confidence, and Θ_a contains the predefined visual constraints. The output a_t includes color, particle number, diffusion direction, trajectory curvature, and line width.

Each movement class was assigned a distinguishable visual rule. Hand opening produced outward particle diffusion, while closing caused particles to aggregate toward the mirrored hand. Flexion generated inward-curving lines, whereas extension produced outward light trails. Pronation generated clockwise color gradients, while supination generated counterclockwise gradients. These directional differences allowed the two rotational movements to remain visually distinguishable.

Movement intensity controlled particle density, line width, and transformation amplitude, whereas classification confidence limited the opacity of the generated content. Low-confidence predictions therefore produced weaker visual changes rather than abrupt transformations. The mirrored digital limb remained in the foreground, while particles, trajectories, and color gradients were restricted to the surrounding or background area. This constrained composition preserved the recognizability of movement direction and limb contour while allowing each interaction to generate a variable artistic output.

3.7. Vibrotactile Substitution

The smartphone vibration motor encoded movement and interaction events according to

$$h_t = f_h(c_t, e_t, q_t; \Theta_h), \quad (4)$$

where $h_t = (d_t^h, I_t^h, n_t^h, \Delta_t^h)$ represents pulse duration, intensity, pulse number, and inter-pulse interval. The variable e_t denotes movement initiation, maintenance, completion, or deviation, while Θ_h contains the predefined safety-constrained encoding rules.

Movement initiation generated a short pulse, successful completion generated two pulses, sustained movement generated intermittent low-intensity feedback, and unstable input generated a distinguishable low-frequency pattern. A tactile pattern was activated only when classification confidence exceeded the threshold and the minimum inter-trigger interval had elapsed. These conditions prevented overlapping windows from causing repeated vibration. Users could select reduced, standard, or enhanced vibration levels or disable the function.

Vibration intensity encoded digital interaction events rather than pain severity. Accordingly, the component provided tactile substitution but did not claim to restore physiological touch in the missing limb. The mirrored animation, constrained generative graphics, and vibration patterns were thus generated from a shared movement estimate while remaining independently adjustable.

4. Experimental Design and Prototype Evaluation

4.1. Experimental Protocol

Movement recognition followed a leave-one-subject-out protocol using seven states from NinaPro DB3. In each of the 11 folds, one participant was retained for testing, and the remainder were used for model development, with 20% of their repetitions assigned to validation. No participant appeared in both the training and test partitions, preventing identity leakage.

The support vector machine (SVM), random forest, and 1D-CNN used the same ten sEMG channels, filtering procedure, movement labels, and 200-ms windows with 50% overlap. The SVM used a radial basis function kernel, while the random forest contained 200 trees. The 1D-CNN was trained for up to 120 epochs using Adam, a batch size of 64, an initial learning rate of 10^{-3}, and early stopping with a patience of 12 epochs. Experiments were repeated with five random seeds, and results are reported as mean ± standard deviation across the 11 folds and five runs.

Model development used a workstation with an Intel Core i7-12700K processor, 32 GB RAM, and an NVIDIA RTX 3060 GPU. Prototype inference was evaluated in an Android Studio environment with four processor cores, 4 GB memory, and Android 14 using TensorFlow Lite. Haptic latency measured the interval from movement recognition to vibration-command dispatch, not the motor's physical onset.

For statistical comparisons, the five runs within each fold were averaged to obtain one result per test participant. Two-sided paired t-tests were used when paired differences satisfied normality; otherwise, Wilcoxon signed-rank tests were applied. Multiple comparisons used Holm correction, with significance defined as adjusted $p<0.05$.

4.2. Evaluation Metrics

Movement recognition was evaluated using accuracy, Macro-F1, and per-class recall. Accuracy measured the proportion of correctly classified windows, whereas Macro-F1 assigned equal weight to every movement class. Per-class recall was used to identify movements that were frequently missed or confused.

App performance was measured using inference time per window, visual generation latency, average frame rate, and peak memory usage. Visual latency was defined as the interval between the availability of a movement prediction and presentation of the corresponding rendered frame. Frame rate was averaged over a five-minute prototype session.

Cross-modal synchronization was evaluated using movement-to-visual latency, movement-to-vibration-command latency, tactile trigger accuracy, and false-trigger rate. A tactile command was considered correct when its encoded event matched the recognized movement and was dispatched within 100 ms of classification. The false-trigger rate represented the percentage of vibration commands activated during low-confidence or unstable predictions.

4.3. Movement Recognition Results

As shown in Table 2, the 1D-CNN achieved the highest accuracy of 83.1±5.4% and Macro-F1 of 81.7±5.8%. The random forest obtained 76.4±6.1% accuracy and 74.8±6.5% Macro-F1, while the SVM achieved 72.8±6.7% and 70.9±7.1%, respectively. The 1D-CNN outperformed the random forest in accuracy (adjusted $p=0.006$) and Macro-F1 (adjusted $p=0.004$). Its corresponding improvements over the SVM were also significant for accuracy (adjusted $p<0.001$) and Macro-F1 (adjusted $p<0.001$).

Table 2. Movement Recognition, System Performance, and Ablation Results.

Part A. Movement-recognition performance			
Method	**Accuracy (%)**	**Macro-F1 (%)**	**Inference Time (ms)**
SVM	72.8 ± 6.7	70.9 ± 7.1	4.8 ± 0.6
Random forest	76.4 ± 6.1	74.8 ± 6.5	6.9 ± 0.8
1D-CNN	83.1 ± 5.4	81.7 ± 5.8	8.7 ± 1.1

Part B. Prototype ablation results				
App Variant	**Visual Latency (ms)**	**Haptic Latency (ms)**	**Frame Rate (fps)**	**Trigger Accuracy (%)**
A: Mirror only	24.7 ± 2.9	N/A	58.9 ± 1.1	N/A
B: Mirror + art	38.6 ± 4.5	N/A	55.8 ± 2.6	N/A
C: Mirror + haptic	26.2 ± 3.1	31.8 ± 3.7	58.1 ± 1.5	96.2 ± 1.8
D: Complete system	42.9 ± 5.1	35.6 ± 4.2	53.7 ± 3.1	94.8 ± 2.3

Figure 2 presents the per-class recall of the 1D-CNN. Rest was the most recognizable state, with a recall of 94.0%. Opening and closing achieved recalls of 86.9% and 84.7%, respectively. Lower recall was observed for flexion (80.2%), extension (77.8%), pronation (79.6%), and supination (78.4%). Flexion and extension produced the lowest combined recall among the paired movements, while pronation and supination showed similarly limited performance. These results correspond to the two principal confusion pairs observed in the classification outputs. The errors can be attributed to similarities in residual-muscle activation and inter-participant variation in residual-limb anatomy and electrode placement.

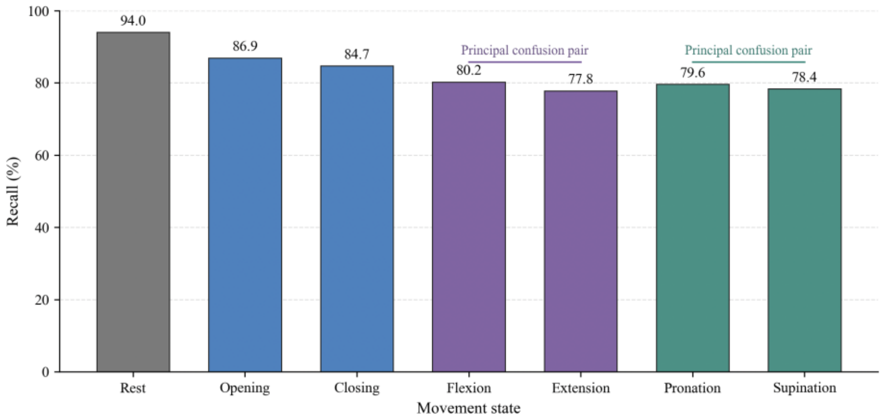

Figure 2. Per-Class Recall of the 1D-CNN for Seven Movement States. Flexion–extension and Pronation–supination Represent the Two Principal Confusion Pairs.

Although the 1D-CNN required more inference time than the SVM and random forest, its average latency remained below 10 ms per window. It therefore provided the most favorable balance between movement-recognition performance and prototype responsiveness. The subject-independent protocol also imposed a more demanding evaluation condition than random window-level splitting because the model could not rely on signal characteristics specific to the test participant.

4.4. Real-Time Feedback and Ablation Results

Four application variants were evaluated. Variant A contained mirror animation only. Variant B combined mirror animation with generative art. Variant C integrated mirror animation and vibrotactile feedback. Variant D contained all three feedback components.

Variant A achieved a visual latency of 24.7 ± 2.9 ms and a frame rate of 58.9 ± 1.1 fps. Adding generative art increased latency to 38.6 ± 4.5 ms and reduced the frame rate to 55.8 ± 2.6 fps. The latency increase was significant after Holm correction (adjusted p <

0.001), primarily because of particle updating, trajectory smoothing, and layered rendering.

Variant C achieved a visual latency of 26.2 ± 3.1 ms, haptic latency of 31.8 ± 3.7 ms, and frame rate of 58.1 ± 1.5 fps. Its visual latency did not differ significantly from Variant A (adjusted $p = 0.118$), indicating limited rendering overhead from vibration dispatch. Tactile trigger accuracy reached 96.2 ± 1.8%, with a false-trigger rate of 2.1 ± 0.8%.

Variant D recorded a visual latency of 42.9 ± 5.1 ms, haptic latency of 35.6 ± 4.2 ms, and frame rate of 53.7 ± 3.1 fps. Visual latency was higher than in Variant B (adjusted $p = 0.013$), while haptic latency exceeded that of Variant C (adjusted $p = 0.021$). Nevertheless, both remained below the 100-ms interaction threshold, and rendering remained above 50 fps.

Peak memory usage was 146, 214, 158, and 231 MB for Variants A–D, respectively, showing that generative rendering accounted for most additional memory use. Variant D achieved 94.8 ± 2.3% trigger accuracy, slightly below Variant C (adjusted $p = 0.047$), while its false-trigger rate increased to 3.4 ± 1.0%. Rapid transitions occasionally activated adjacent generative and tactile events, although confidence gating and the minimum inter-trigger interval prevented most repeated vibration.

Overall, generative art imposed a greater computational cost than vibrotactile encoding. Despite the highest latency and memory use, the complete system maintained responsive rendering and tactile-command dispatch, supporting the technical compatibility of the three feedback components without establishing clinical effectiveness.

4.5. Visual and Synchronization Analysis

Figure 3(a) compares the accuracy and Macro-F1 of the three recognition models. The 1D-CNN maintained a similar advantage for both metrics, indicating that its performance was not limited to the more frequently represented states. The error bars also show substantial between-participant variation, supporting the use of subject-independent rather than randomly mixed window-level evaluation.

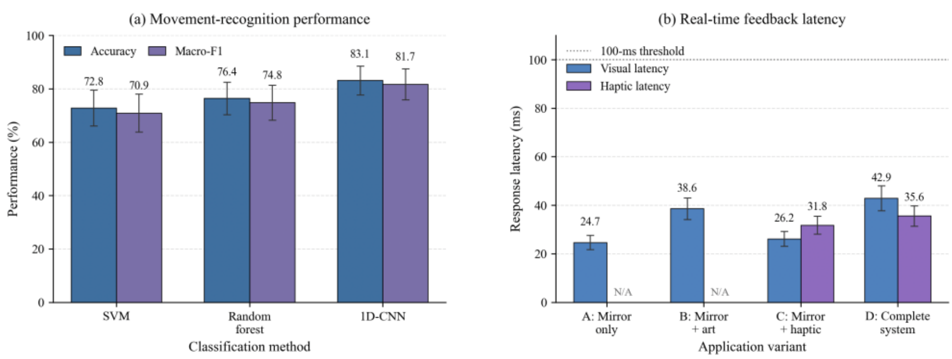

Figure 3. Movement Recognition and Real-Time Feedback Performance of the Proposed App. (A) Movement Recognition Performance, Measured Using Accuracy and Macro-F1 under Subject-Independent Evaluation. (B) Visual and Haptic Response Latency Across the Four Application Variants. Error Bars Represent Standard Deviations; Haptic Latency Is Not Applicable to Variants without Vibration Feedback.

Figure 3(b) compares visual and haptic response latency across the four application variants. Generative rendering produced a larger latency increase than vibration-command dispatch. Nevertheless, the complete system remained below the 100-ms threshold and above 50 fps, indicating that the three feedback channels could be coordinated without interrupting visual continuity.

Overall, the evaluation supports the technical compatibility of movement recognition, mirror animation, constrained generative graphics, and vibration encoding. These findings concern recognition performance and feedback synchronization and do not establish pain reduction, sensory restoration, or clinical effectiveness.

5. Discussion

5.1. Main Findings

The findings address the three research questions from complementary technical perspectives. For RQ1, the NinaPro DB3 data supported subject-independent recognition of seven representative movement states. The 1D-CNN achieved higher accuracy and Macro-F1 than the SVM and random forest while maintaining an inference time below 10 ms. However, flexion–extension and pronation–supination remained the principal confusion pairs, indicating that differences in residual-muscle activation across individuals continue to affect cross-participant recognition.

For RQ2, constrained generative art enriched the visual output without replacing the mirrored limb representation. Retaining the sEMG-driven digital limb as a stable foreground reference preserved the predicted movement direction and contour, while particle diffusion, aggregation, light trails, and directional color gradients provided variable artistic responses. This implementation represents digitally reconstructed mirror feedback rather than direct camera-based reflection of the intact limb. Generative rendering increased latency and memory use, but the complete system maintained more than 50 fps and remained within the predefined response threshold.

For RQ3, vibration commands were synchronized with movement and visual events with relatively low latency. The complete system achieved a tactile trigger accuracy of 94.8%, although rapid state transitions caused a small increase in false triggers. Confidence gating and a minimum inter-trigger interval reduced repeated vibration during unstable predictions.

5.2. Design Implications

Generative art should be treated as an interaction layer that enriches mirror feedback rather than as a replacement for the mirrored limb contour. Similarly, vibrotactile feedback communicates event type, movement completion, or interaction instability; it does not restore natural touch in the missing limb. The smartphone integrates visual rendering, processing, interface control, and vibration output within a portable platform. However, the present sEMG-based configuration still requires a compatible external wearable sensor, whereas a future camera-based mode could reduce this additional hardware requirement. Nevertheless, personalization remains important. Users should be able to adjust color palettes, visual complexity, artistic style, animation intensity, and vibration strength to accommodate perceptual preferences and avoid sensory overload.

5.3. Limitations

NinaPro DB3 contains only 11 participants and primarily represents individuals with transradial upper-limb loss. Offline recordings cannot fully reproduce variations in electrode placement, posture, fatigue, device handling, background applications, and hardware performance during home use. The dataset does not contain pre- and post-intervention mirror-therapy outcomes, and this study did not evaluate phantom limb pain, quality of life, body image, or sustained engagement. Smartphone vibration also cannot reproduce physiological sensation originating from the missing limb. Future longitudinal studies should therefore recruit target users under ethical approval and professional supervision.

6. Conclusion

This study developed a generative art application framework for individuals with limb loss by integrating movement recognition, digital mirror visual feedback, and vibrotactile substitution. First, a lightweight 1D-CNN was used to recognize seven representative movement states from NinaPro DB3 and drive a horizontally mirrored digital-limb representation. The current implementation relies on external sEMG input, while direct camera-based mirroring remains a future extension. Second, a cross-modal mapping mechanism connected movement class, intensity, and confidence with bounded visual parameters and distinguishable vibration patterns. Third, open-data evaluation and prototype testing examined recognition performance, rendering latency, frame rate, memory use, and tactile-command synchronization. The results indicate that the three feedback components can operate together within the response requirements of a mobile prototype.

The system provides an implementable framework for a generative-art-supported digital complementary intervention. Its contribution lies in technical integration and interaction design rather than demonstrated therapeutic efficacy. It should not be considered a replacement for clinical assessment, rehabilitation, medication, or psychological support.

Future work should recruit individuals with limb loss for ethically approved longitudinal evaluation, including usability, engagement, phantom limb pain, body-image, and quality-of-life measures. Further development should incorporate individualized tactile calibration, compare different generative art styles, and examine whether the interaction framework can be adapted to lower-limb loss. These studies are required to determine its user-level acceptability and potential value within supervised complementary care.

References

1. A. Aternali, H. Lumsden-Ruegg, L. Appel, S. L. Hitzig, A. L. Mayo, and J. Katz, "Phantom limb telescoping in individuals with limb loss: links to anxiety, depression, and pain-related measures," *Frontiers in Pain Research*, vol. 7, p. 1755884, 2026.
2. M. Langeveld, R. Bosman, C. A. Hundepool, L. S. Duraku, C. McGhee, J. M. Zuidam, et al., "Phantom limb pain and painful neuroma after dysvascular lower-extremity amputation: a systematic review and meta-analysis," *Vascular and Endovascular Surgery*, vol. 58, no. 2, pp. 142–150, 2024.
3. E. H. Beisheim-Ryan, G. E. Hicks, R. T. Pohlig, J. Medina, and J. M. Sions, "Body image and perception among adults with and without phantom limb pain," *PM&R*, vol. 15, no. 3, pp. 278–290, 2023.
4. X. Huo, P. Huang, H. Di, T. Ma, S. Jiang, J. Yao, and L. Huang, "Risk factors analysis of phantom limb pain in amputees with malignant tumors," *Journal of Pain Research*, pp. 3979–3992, 2023.
5. C. B. Satala, A. Onofrei, O. Vrînceanu, and D. Mihalache, "From Amputation to Persistent Pain: A Review of Molecular and Cellular Processes in Phantom Limb Pain," *International Journal of Molecular Sciences*, vol. 27, no. 5, p. 2107, 2026.
6. M. Diers, X. Fuchs, R. Bekrater-Bodmann, and H. Flor, "Prevalence of phantom phenomena in congenital and early-life amputees," *The Journal of Pain*, vol. 24, no. 3, pp. 502–508, 2023.
7. E. R. Lontis, K. Yoshida, and W. Jensen, "Non-Invasive Sensory Input Results in Changes in Non-Painful and Painful Sensations in Two Upper-Limb Amputees," *Prosthesis*, vol. 6, no. 1, pp. 1–23, 2023.
8. S. Ishigami and C. Boctor, "Epidemiology and risk factors for phantom limb pain," *Frontiers in Pain Research*, vol. 5, p. 1425544, 2024.
9. K. Szczypiór-Piasecka, "Mirror therapy in the treatment of phantom limb pain following traumatic upper limb amputation–a case report," *Chir Narzadow Ruchu Ortop Pol*, vol. 90, no. 3, pp. 169–173, 2025.
10. G. Marullo, C. Innocente, L. Ulrich, A. Lo Faro, A. Porcelli, R. Ruggieri, et al., "Home-based mirror therapy in phantom limb pain treatment: The augmented humans framework," *Multimedia Tools and Applications*, vol. 84, no. 28, pp. 34145–34177, 2025.
11. C. Prahm, K. Eckstein, M. Bressler, Z. Wang, X. Li, T. Suzuki, et al., "PhantomAR: gamified mixed reality system for alleviating phantom limb pain in upper limb amputees—design, implementation, and clinical usability evaluation," *Journal of Neuroengineering and Rehabilitation*, vol. 22, no. 1, p. 21, 2025.
12. A. S. Eldaly, F. R. Avila, R. A. Torres-Guzman, K. C. Maita, J. P. Garcia, L. P. Serrano, et al., "Virtual and augmented reality in management of phantom limb pain: a systematic review," *Hand*, vol. 19, no. 4, pp. 545–554, 2024.
13. Q. Niu, L. Shi, Y. Niu, K. Jia, G. Fan, R. Gui, and L. Wang, "Motion intention recognition of the affected hand based on the sEMG and improved DenseNet network," *Heliyon*, vol. 10, no. 5, 2024.
14. B. M. Gunterstockman, A. D. Knight, C. E. Mahon, W. L. Childers, T. Cagle, B. D. Hendershot, and S. Farrokhi, "Relationship between phantom limb pain, function, and psychosocial health in individuals with lower-limb loss," *Prosthetics and Orthotics International*, vol. 47, no. 2, pp. 181–188, 2023.
15. A. S. Ivani, M. G. Catalano, G. Grioli, M. Bianchi, Y. Visell, and A. Bicchi, "Tactile perception in upper limb prostheses: mechanical characterization, human experiments, and computational findings," *IEEE Transactions on Haptics*, vol. 17, no. 4, pp. 817–829, 2024.