AI Remote TCM Diagnosis Software via Palm Image Analysis

Dawei Wang [1,*] and Jingxiu Zhang [2]

[1] College of Intelligent Science and Engineering, Harbin Engineering University, Harbin 150001, Heilongjiang, China
[2] Hefei Mingchen Electromechanical Technology Co., Ltd., Hefei 230093, Anhui, China

* Corresponding author. Tel.: +86 15181440044; email: daweiwang@hrbeu.edu.cn
Manuscript submitted June 12, 2026; accepted October 13, 2026; published November 25, 2026
doi: 10.17706/jsw.21.2.78-92

Abstract: To address the uneven distribution of Traditional Chinese Medicine (TCM) medical resources in grassroots and remote regions and geographical restrictions of offline pulse-taking and face-inspection diagnosis, this paper designs and implements an Artificial Intelligence (AI)-powered remote TCM diagnosis software prototype based on palm-image analysis. Guided by TCM meridian and holographic theories, a four-layer modular software architecture including data-acquisition layer, storage-management layer, AI-model layer and application-interface layer is constructed, together with cross-layer privacy-protection mechanisms for sensitive medical data. Three independent U-Net semantic-segmentation models are deployed to realize pixel-level segmentation for palm Region of Interest (ROI), subcutaneous cyan veins and cross-mark textures. A weighted Binary Cross-Entropy (BCE)-Dice mixed loss function is adopted to alleviate severe foreground-background class imbalance of tiny pathological targets. Post-processing constraints based on palm ROI masks are applied to eliminate false predictions outside valid palm areas. The model is trained and validated on samples from the public Tongji University Palm Vision (TJU-PV) palmprint dataset; pixel-level medical segmentation masks are manually annotated in-house for this study. The prototype supports image acquisition, feature segmentation, TCM nine-grid holographic zoning syndrome differentiation and structured auxiliary diagnostic report generation. Experiments achieve favorable segmentation performance for palm ROI. Restricted by limited manually-annotated samples and extremely small pixel occupation of pathological features, faint cyan veins and cross-marks still suffer from missed detections. This prototype only provides auxiliary health assessment and cannot replace formal clinical diagnosis by TCM physicians. It offers technical references for non-contact remote intelligent TCM healthcare.

Keywords: traditional Chinese medicine palm diagnosis, remote diagnosis, U-Net, semantic segmentation, palm-print features, intelligent healthcare

1. Introduction

In the digital era, artificial intelligence has profoundly transformed traditional medical service modes. Grassroots and remote regions in China are suffering from shortages of high-quality Traditional Chinese Medicine (TCM) resources. Patients often need long-distance travel to large hospitals, which brings considerable time and economic costs. Remote intelligent diagnosis breaks geographical restrictions and enables end-users to obtain auxiliary health assessment at home, which is of great value for early sub-health screening under the “preventive treatment of disease” philosophy of TCM.

Inspection diagnosis is a core branch of TCM, which includes face, tongue and palm inspection. As a non-contact diagnostic modality, palm diagnosis is highly suitable for mobile-phone-based remote data collection. According to the holographic and meridian theories of TCM, pathological changes of internal zang-fu organs are reflected on palms in the form of color, vein and texture variations, as recorded in Huangdi's Internal Classic Spiritual Pivot. Observing subcutaneous cyan veins, pathological cross-marks and skin color variations can assist the evaluation of physical constitution and sub-health risks.

Existing AI-driven TCM research focuses heavily on tongue-image analysis for segmentation, constitution identification and syndrome differentiation. Comparatively, palm-diagnosis-oriented studies remain insufficient. Conventional palm-feature extraction heavily relies on hand-crafted operators including Canny edge detection, Hough transform and Harris corner detection. Such methods are vulnerable to illumination changes, skin-color variations and complex fine textures on palm surfaces, leading to frequent false detections. Many existing publications only study coarse main palm lines, while ignoring clinically meaningful tiny pathological marks such as subcutaneous cyan veins and cross-marks. Furthermore, most existing works only deliver isolated algorithm prototypes rather than complete end-to-end software covering image acquisition, persistent storage, AI inference and dual-role doctor-patient interaction workflows.

Three main contributions are summarized in this paper:

1. A four-layer modular end-to-end architecture for remote TCM palm diagnosis is presented, with complete patient-doctor workflows and role-based access control for medical-data privacy protection.

2. Three decoupled U-Net models are designed to segment palm- Region of Interest (ROI), subcutaneous cyan-veins and pathological cross-marks. Compensating for the research gap that existing palm segmentation mainly serves biometric identification, we adopt BCE-Dice loss and ROI-mask post-processing to improve segmentation performance on small pathological features.

3. A fully functional AI-assisted palm-diagnosis prototype is implemented. Combined with TCM nine-grid holographic zoning rules, structured diagnostic reports can be automatically produced. Model evaluation, limitation analysis and future improvement directions for intelligent TCM palm diagnosis are also provided.

2. Related Work

In recent years, computer vision and deep learning have greatly promoted non-invasive medical auxiliary diagnosis, bringing new opportunities for intelligent TCM palm diagnosis. Existing palm-image studies can be divided into medical auxiliary diagnosis, palm biometric recognition, and general palmline analysis.

In medical TCM scenarios, researchers have verified the correlation between palm visual phenotypes and physical conditions, including respiratory diseases, visceral lesions, and TCM constitution identification [1–5]. These studies demonstrate the feasibility of using palm images for health screening. Some works have developed palm image acquisition devices and objective evaluation frameworks to standardize TCM palm inspection [6–7]. In the biometric field, most palmprint learning methods focus on contour detection, few-shot recognition, and constrained ROI extraction [8–10]. Nevertheless, these methods mainly serve identity authentication and cannot achieve fine-grained semantic segmentation of pathological palm features.

From the perspective of AI-TCM interdisciplinary research, existing reviews show that current intelligent TCM diagnosis studies predominantly focus on tongue recognition and syndrome classification, while pixel-level feature segmentation for palm diagnosis remains insufficient [11–15]. Most existing palm-learning techniques rely on classification models that output overall disease risk without locating specific pathological regions, making it difficult to explain which palm textures support the diagnostic results.

Recently, advanced palm-image segmentation methods have further improved fine-grained palm-feature

extraction [16–25]. Current technical improvements include multi-scale convolution, Transformer enhancement, context fusion, and multi-level feature fusion strategies for palm-vein and palm-line segmentation. Data augmentation and lightweight segmentation pipelines have also been explored to alleviate annotation scarcity. However, almost all existing palm-segmentation tasks are designed for biometric recognition. Existing studies only segment palm outlines or general palm veins, without targeting TCM-specific pathological features, such as subcutaneous cyan veins and pathological cross-marks.

In summary, current medical palm-diagnosis studies mainly adopt image-level classification, lacking precise pixel-wise parsing of local pathological markers. Meanwhile, existing palm-segmentation methods and public datasets are oriented toward biometric identification and lack medical annotations for TCM diagnostic features. To address this research gap, this paper proposes a multi-branch U-Net segmentation framework for palm core ROI, subcutaneous vein, and cross-mark segmentation. Based on the extracted TCM pathological features, we further construct an intelligent remote palm-diagnosis system to support objective and explainable auxiliary diagnosis of TCM.

3. System Overall Design

3.1. Four-layer Overall Architecture

The proposed AI-powered remote TCM palm-diagnosis software adopts a layered modular architecture consisting of four layers from bottom to top: Data Acquisition Layer, Storage-Management Layer, AI-Model Layer, Application-Interface Layer. A cross-layer security-and-privacy module protects sensitive palm images and diagnostic records throughout all workflows. The overall architecture is shown in Fig. 1.

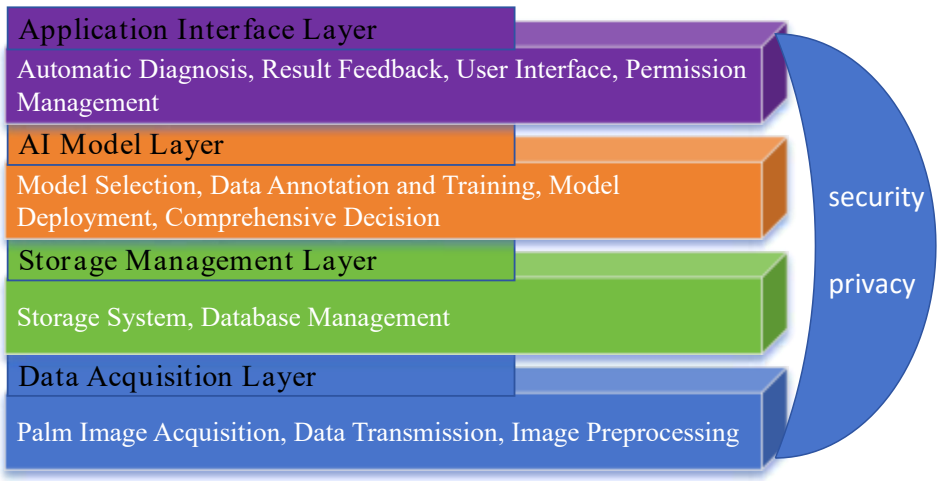


Fig. 1. Schematic diagram of four-layer overall system architecture.

3.1.1. Data acquisition layer

This bottom layer obtains raw palm images, including palm-image capture, encrypted transmission and preliminary image pre-processing. Users capture palm photographs via mobile cameras; image enhancement and noise reduction are performed. Built-in quality-control rules reject invalid samples suffering insufficient illumination or distorted hand posture before delivering valid images upward.

3.1.2. Storage-Management Layer

Responsible for persistent data storage, containing file-storage subsystem and relational database. Original palm photographs, pre-processed images, patient profiles, AI diagnostic records and doctor-review feedback are stored separately. Decoupling image files and business metadata improves system scalability for future function expansion.

3.1.3. AI-Model Layer

As the core computing unit, it manages dataset annotation, model training, model deployment and hybrid TCM rule-based decision-making. Three independent U-Net semantic-segmentation models are deployed for palm-ROI segmentation, subcutaneous cyan-vein segmentation and cross-mark segmentation. Extracted segmentation masks are fed into the TCM nine-grid holographic zoning module for feature statistics to generate preliminary auxiliary diagnostic results for upper-layer delivery.

3.1.4. Application-Interface Layer

Two separated portals are provided for end-users: patient side and doctor side. Functions cover automatic diagnostic-report generation, bidirectional result feedback, graphical user interface and role-based permission management. Patients can upload palm images and view sub-health conditioning suggestions; TCM physicians can inspect AI outputs, manually review and revise diagnostic conclusions.

3.1.5. Cross-layer security-and-privacy guarantee

Role-based access isolation differentiates permissions for ordinary patients, physicians and administrators to prevent unauthorized access and leakage of sensitive medical data.

3.2. Business Workflow for Remote Palm-image Intelligent Diagnosis

The complete remote palm-health-recognition workflow is illustrated in Fig. 2. Pre-diagnosis reminders are displayed before formal testing: four contraindications for palm inspection (do not diagnose after alcohol intake, violent emotional outbursts, sexual activity, or under poor lighting). Three pre-query items collect age, gender and medication history. Standardized palm images are acquired afterwards.

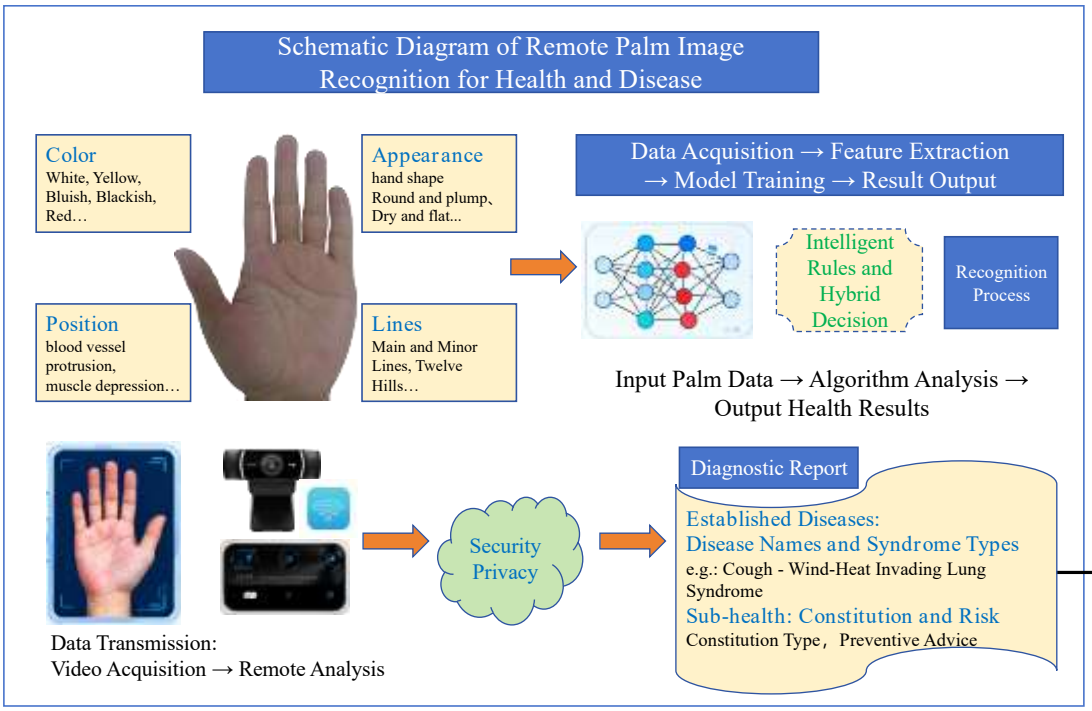


Fig. 2. Workflow diagram for remote palm-image-based diagnosis.

Four major categories of palm-health features are extracted:

1. Color features: palm skin tone, including white, yellow, cyan-purple and other tonal manifestations;

2. Appearance features: overall hand morphology, such as plump-versus-gaunt palm shapes;

3. Position-related pathological features: local lesions including protruding subcutaneous cyan veins and palm-surface muscular depressions;

4. Line features: primary palm lines, secondary fine wrinkles, pathological marks (cross-marks and grid-shaped textures) distributed over twelve palm mounds.

Images are transmitted to backend services under privacy protection following the pipeline: Data Acquisition → Feature Extraction → Model Inference → Result Output. Three U-Net models produce segmentation masks for palm ROI, cyan veins and cross marks. The palm-ROI binary mask multiplies vein and cross-mark prediction outputs to zero out false detections outside valid palm regions. The extracted features are imported into the hybrid intelligent-rule decision-making module for syndrome reasoning according to TCM nine-grid zang-fu organ mapping.

The final structured diagnostic report contains two parts:

1. Confirmed-disease section: disease names and corresponding TCM syndrome types;

2. Sub-health-risk section: constitution evaluation, health-risk warnings and targeted TCM conditioning recommendations.

General data-flow summary: Palm-image input → Algorithm analysis → Health-diagnosis-result output. Physicians can review reports and initiate remote video consultation if required.

3.3. Prior Mapping Between Palm-Print Pathological Features and Diseases

Based on the TCM palm-diagnosis rules summarized in Tables 1 and 2, subcutaneous veins, palm cross-shaped marks and palm creases correspond to different physical conditions and pathological signs. Table 1 elaborates the correlations between the positions of subcutaneous blue veins on the palm and their pathological manifestations, while Table 2 classifies palm marks by shape and interprets their pathological implications. In particular, palm cross-shaped marks can reflect potential lesions and support preventive diagnosis for "undiseased states". This feature-pathology mapping provides key diagnostic evidence for both the AI segmentation module and clinicians. However, most existing palmprint-related studies focus on biometric identification, and few studies incorporate such TCM pathological prior knowledge for disease diagnosis, which limits medical-oriented practical applications.

Table 1. Correlation between palm blue-vein positions and pathological manifestations

Blue-vein Location on Palm	Pathological Implications from Traditional Palm-Diagnosis
Finger region	Children: indigestion of gastrointestinal tract; Adults: gastrointestinal disorder, insufficient cerebral blood supply, dizziness, early-warning risk of stroke
Dorsum of hand	Lumbar-back stagnation, sub-health signals related to lumbar muscle strain
Thenar eminence of palm	Lumbocrural rheumatism, joint impediment-pain
Wrist crease area	Gynecological disorders, irregular menstruation, leukorrhea problems
Lifeline trajectory	Early-warning signals of hepatobiliary diseases
Tiger-mouth region	Mastalgia during menstrual period, liver-qi stagnation and blood stasis
Palm-finger junction	Intestinal dysfunction, constipation
Inner palm area	Gastrointestinal diseases, risk of hypertension-hyperlipidemia-hyperglycemia, acidic constitution, dizziness and physical weakness

Table 2. Correlation between palm mark shapes and pathological implications

Sketch Icon	Palm Mark Type	Pathological Implications from Traditional Palm-Diagnosis
+	Cross-shaped mark	Pathological condition has occurred or is going to occur
△	Triangular mark	Pathological condition is occurring
□	Quadrilateral mark	Old-standing/chronic disease
#	Grid-shaped mark	Chronic old disease with high recurrence risk

3.4. Software-engineering Characteristics Analysis

Different from many existing intelligent-medical software products that only deploy isolated AI models, our system is specially customized for TCM palm-diagnosis scenarios. The layered-modular architecture

decouples AI-algorithm logic from business logic, and several key software-engineering properties are analyzed as follows.

1. Scalability: Modules are loosely coupled; new segmentation models or updated TCM diagnostic rules can be plugged-in without system-wide reconstruction.

2. Reliability: Pre-acquisition quality-control filters low-quality input images and reduces useless inference computation overhead.

3. Security: Role-based permission control isolates patient-side and doctor-side data views to prevent unauthorized access to sensitive palm medical data.

4. Real-time performance: Model inference can run on ordinary consumer-grade Central Processing Units (CPUs) for prototype verification; the layered architecture supports future migration toward cloud-server deployment for online real-time service.

3.5. Prototype Software Implementation

Based on the four-layer system architecture proposed above, a desktop graphical prototype is implemented using Python and built-in Tkinter library under VSCode development environment, as shown in Fig. 3. The prototype realizes core business modules including palm-image upload, algorithm-model selection, AI inference simulation, and diagnostic-report generation with physician suggestions.



Fig. 3. A Prototype of TCM palm-print tele-diagnosis system.

The graphical user interface contains four functional zones:

1. Image upload and preview area: Supports local palm-image file selection and preview display. Pre-diagnosis tips of TCM palm inspection ("four contraindications and three inquiries") are prompted before analysis.

2. Model & algorithm selection panel: Drop-down list for selecting three independent U-Net segmentation models: Palm-ROI segmentation, Subcutaneous cyan-vein segmentation, Cross-mark segmentation.

3. AI inference output panel: Displays segmentation feedback, feature statistics from TCM nine-grid holographic zoning. Real model inference interfaces are reserved; current version simulates output for prototype demonstration.

4. Diagnostic-report display area: Automatically generates structured reports including confirmed-disease items, sub-health risk evaluation and physician-provided conditioning suggestions. Reports can be exported and saved as local TXT files.

Note that this desktop program serves as system prototype verification. After completing full model training, the simulated inference block can be replaced by loading trained.pth weight files of three U-Net networks for real pixel-level segmentation.

4. Algorithm and Experimental Design

4.1. Overall Algorithm Scheme

Traditional hand-crafted image operators suffer from interference caused by uneven illumination, skin-color diversity and dense fine palm textures when extracting palm-ROI, subcutaneous cyan-vein and cross-mark features. We adopt U-Net semantic-segmentation networks for pixel-wise feature extraction.

Cyan veins and cross-marks belong to tiny foreground targets occupying very few pixels and bring severe foreground-background class imbalance. Moreover, cyan veins and cross-marks share similar visual appearances and easily cause feature confusion. Instead of multi-task joint training with shared encoder weights, we adopt three completely independent U-Net models (palm-ROI model, subcutaneous-cyan-vein model, cross-mark model). Each model has separate weights, independent annotated datasets and decoupled training procedures to avoid mutual interference between different target classes. The algorithm pipeline includes three phases: image pre-processing, U-Net inference, and ROI-mask-based post-processing filtering.

4.2. U-Net Network Architecture

Classic U-Net serves as segmentation backbone, consisting of encoder (contracting down-sampling path) and decoder (expansive up-sampling path). Skip connections transfer shallow detailed encoder features toward the decoder to preserve fine spatial information of palm edges and subtle textures alongside high-level semantic representations.

The DoubleConv basic block consists of two successive 3 × 3 convolutions, batch normalization and ReLU activation. Max-pooling realizes down-sampling; transposed convolution completes resolution restoration during up-sampling, followed by feature concatenation with corresponding encoder skip-connection feature maps. A final 1 × 1 convolution outputs single-channel logit probability maps.

Network input: Red Green Blue (RGB) three-channel image resized to 256 × 256 pixels.

Network output: single-channel segmentation mask with pixel values {0, 1}; value 1 denotes foreground target and 0 denotes background.

4.3. Loss Function and Optimization Strategy

Because foreground pixels of cyan veins and cross-marks occupy a tiny fraction of total image pixels, vanilla binary cross-entropy loss biases network predictions toward background and produces missed

small-target detections. We employ weighted BCE-Dice mixed loss to mitigate class imbalance.

$$L_{\text{BCE}} = -\frac{1}{N}\sum_{i=1}^{N}[\, y_i \log(\hat{y}_i) + (1 - y_i)\log(1 - \hat{y}_i)] \tag{1}$$

$$L_{\text{Dice}} = 1 - \frac{2\sum_{i=1}^{N} y_i \hat{y}_i}{\sum_{i=1}^{N} y_i + \sum_{i=1}^{N} \hat{y}_i} \tag{2}$$

$$L_{\text{total}} = \lambda_1 \cdot L_{\text{BCE}} + \lambda_2 \cdot L_{\text{Dice}} \tag{3}$$

where y_i denotes ground-truth mask label, $\hat{y}_i$ denotes network sigmoid predicted probability, N is total pixel count of one image. Hyper-parameters are set as $\lambda_1 = 0.5$, $\lambda_2 = 0.5$.

Adam optimizer is utilized, initial learning rate = 1×10^{-4}, batch size = 2. For the ROI segmentation task and cross-vein (cross-mark) segmentation task, the maximum training epoch was set to 50, while the maximum epoch for palmar vein segmentation was set to 100. Validation loss monitoring with early-stopping strategy saves weights achieving minimal validation loss to prevent over-fitting from final epoch checkpoints.

4.4. Dataset Construction and Partitioning

Palm photographs with standard hand posture and appropriate illumination are screened from the public TJU-PV palmprint dataset [26]. Pixel-level manual annotations targeting palm-ROI, subcutaneous cyan veins and cross-mark textures are completed using Labelme software. Three independent mask label sets are generated for different segmentation tasks. The raw palm images are shared across three tasks, while mask files are completely independent without overlapping labels. The samples are randomly split into training set and validation set following an 8:2 ratio, yielding 200 training samples and 50 validation samples. All input images are resized to 256 × 256 and normalized into floating-point range [0, 1].

4.5. Experimental Hardware and Software Environment

The hardware and software environment used in the experiment is listed in Table 3.

Table 3. Experimental hardware-software environment

Item	Configuration
Operating System	Windows 11
CPU	Intel-Core-i7-12700H
Programming Language	Python 3.10
Deep-learning Framework	PyTorch 2.1
Image-processing Libraries	OpenCV-Python 4.8, Pillow, scikit-image

4.6. Evaluation metrics

Four standard semantic-segmentation quantitative metrics are adopted: Precision, Recall, F1-Score, mean Intersection-over-Union (mIoU). Single-image CPU inference latency is recorded to evaluate processing efficiency.

$$\text{Precision} = \frac{\text{TP}}{\text{TP}+FP} \tag{4}$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{5}$$

$$\text{F1-Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{6}$$

$$mIoU = \frac{1}{k+1}\sum_{i=0}^{k}\frac{TP}{TP_i+FP_i+FN_i} \quad (7)$$

where TP_i, FP_i and FN_i denote the true positive, false positive and false negative pixels of the i-th class, and k denotes the number of segmentation classes.

4.7. Experimental results and analysis

4.7.1. Loss convergence analysis

Three U-Net models for palm ROI, subcutaneous vein and palm cross-mark segmentation are trained iteratively, and the convergence and fitting performance of each model are analyzed based on training and validation loss curves. Fig. 4 shows the variation of the training loss and validation loss of the palm ROI segmentation model during 50 training epochs. Fig. 5 presents the corresponding loss curves of the subcutaneous cyan-vein segmentation model over 100 training epochs, and Fig. 6 reports those of the cross-mark segmentation model over 50 training epochs.

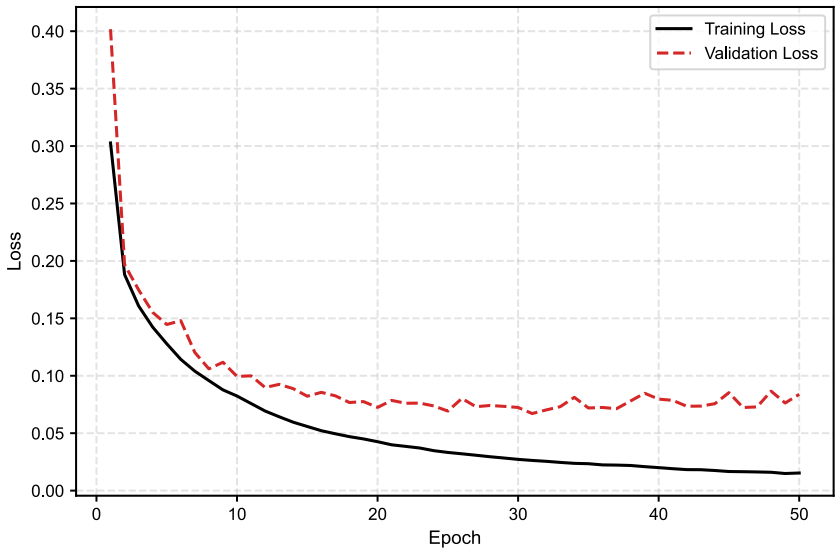


Fig. 4. Training and validation loss curves for the palm ROI segmentation model (50 epochs).

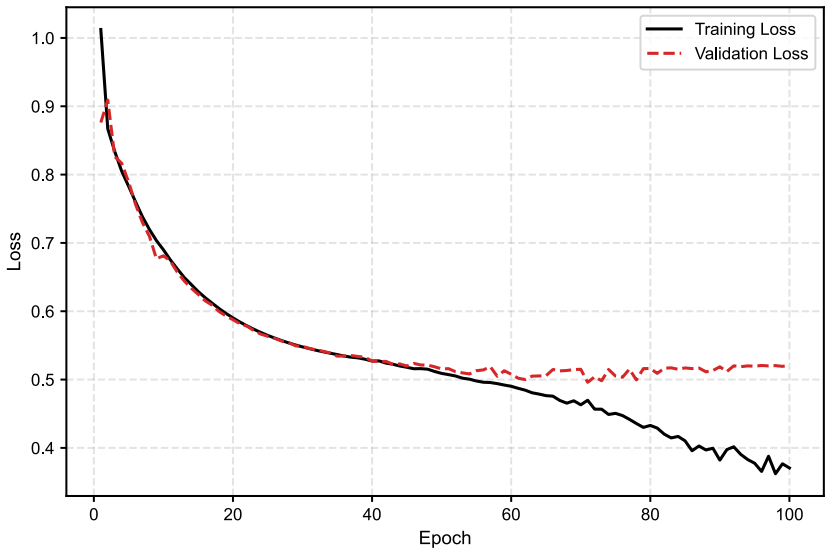


Fig. 5. Training and validation loss curves for the subcutaneous cyan-vein segmentation model (100 epochs).

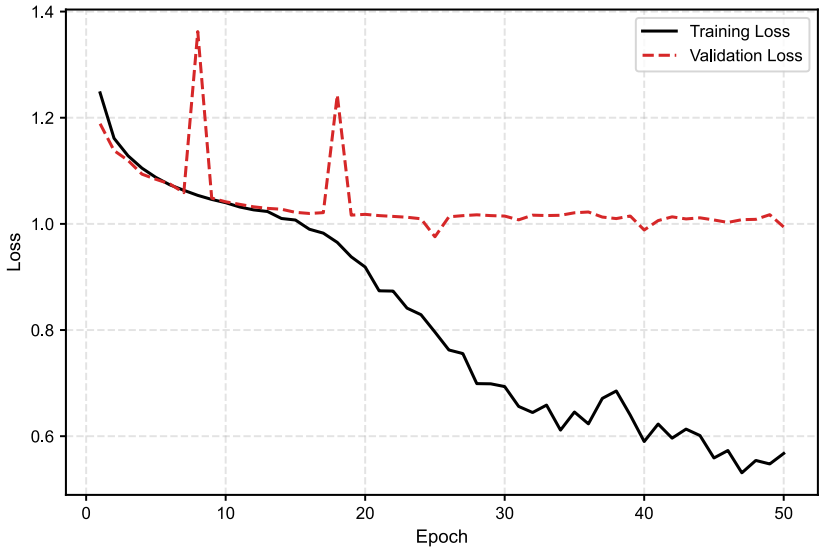


Fig. 6. Training and validation loss curves for the cross-mark segmentation model (50 epochs).

The palm ROI segmentation model converges rapidly. In the first 20 epochs, the training loss decreases from 0.3027 to 0.0426 and the validation loss drops from 0.4018 to 0.0724. The final training loss converges to approximately 0.015, while the validation loss stabilizes between 0.07 and 0.085 with slight fluctuations. The model presents good fitting ability and stable feature extraction with minor overfitting, which satisfies the segmentation requirements of key palm regions.

Subcutaneous vein segmentation is more difficult due to fine textures and large inter-sample differences. The model shows a slow loss decline in the early training stage and gradually converges after 70 epochs. The final training loss is about 0.37 and the validation loss stabilizes in the range of 0.50 to 0.52. The stable loss platform indicates that the model has fully learned vein features and reached the optimal convergence state.

Palm cross-marks are extremely fine textures with sparse labeled samples and severe sample imbalance, leading to the highest segmentation difficulty. Within 50 training epochs, the initial training loss of 1.2469 finally decreases to approximately 0.56. The model converges steadily in the first 7 epochs, while two obvious validation loss spikes occur at the 8th and 18th epochs, causing generalization fluctuations. In the middle and late stages, the training loss continues to decrease, but the validation loss remains stable at 1.00~1.02. The expanding gap between training and validation loss indicates significant model overfitting.

To avoid accuracy degradation caused by overfitting, validation loss is monitored in real time and only the optimal weights are retained. The model at the 25th epoch with the minimum validation loss of 0.9757 is selected for inference. All three models achieve effective convergence and can stably extract palm texture features, providing reliable feature support for intelligent TCM palm diagnosis.

4.7.2. Quantitative segmentation results

Quantitative results on validation set are shown in Table 4. As can be observed from the quantitative metrics on the validation set, the palm-ROI segmentation achieves excellent performance with Precision, Recall, F1-Score and mIoU all above 0.95, which enables accurate localization of the palm region of interest. In contrast, the segmentation accuracy of subcutaneous cyan vein and cross-mark is relatively low, with mIoU values of 0.001 and 0.036 respectively. This is mainly caused by extremely thin vein textures, tiny cross-mark targets and severe sample imbalance. The inference latency for a single image is approximately 1.8 s for all three tasks, which basically meets the requirement of offline analysis.

Table 4. Quantitative segmentation results on validation dataset

Segmentation Target	Precision	Recall	F1-Score	mIoU	Inference latency (s)
Palm ROI	0.975	0.983	0.979	0.959	1.890
Subcutaneous cyan vein	0.001	0.250	0.002	0.001	1.831
Cross-mark	0.043	0.097	0.058	0.036	1.759

4.7.3. Visualization and post-processing effect

In Fig. 7, the left sub-figure shows the original palmprint image, while the right one presents the segmentation visualization result. The blue contour denotes the boundary of palm Region-of-Interest (ROI). The red semi-transparent area marks subcutaneous veins and the yellow semi-transparent area marks cross textures. The visualization result intuitively verifies that the proposed model can simultaneously segment and extract multiple features including palm-ROI, veins and cross textures.

Compared with conventional OpenCV-based hand-crafted pipelines, U-Net segmentation exhibits stronger robustness against fine palm noise and uneven lighting. Multiplying vein/cross-mark prediction masks by the palm-ROI binary mask clears spurious predicted pixels outside valid palm regions and suppresses background false-positive blobs.

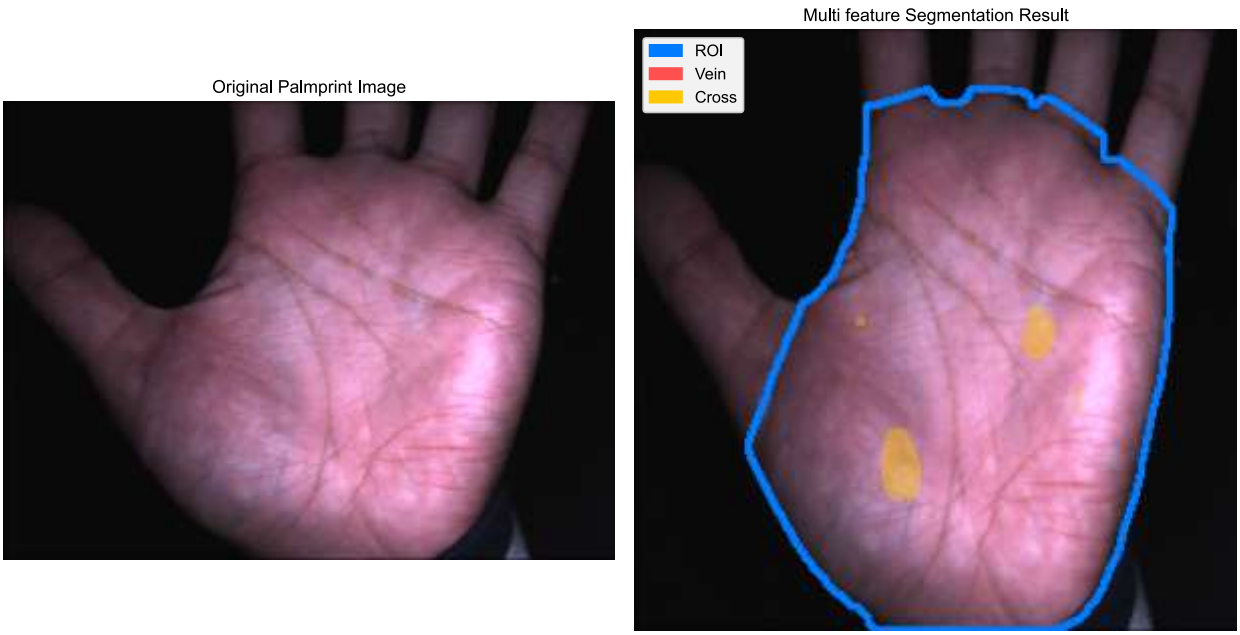


Fig. 7. Visual comparison of multi-feature segmentation for palmprint.

4.7.4. Analysis of Reasons for advantages for TCM palm-diagnosis scenarios

To interpret the better segmentation performance for palm-core ROI, subcutaneous vein and cross-mark tasks, three ROI-extraction schemes (Fig. 8) are compared as visual evidence.

Instead of feeding full-palm images into the network, our U-Net-based ROI segmentation (blue contour, Fig.8-left) accurately extracts the diagnostic-effective palm-core region. The MediaPipe baseline (green contour) produces an undersized ROI missing lateral diagnostic features, whereas the morphology-based method (red contour) incorrectly covers the thumb. Our ROI output suppresses finger-crease noise and improves subsequent segmentation for veins and cross-marks.

Furthermore, existing palm-vein studies mostly serve biometric recognition rather than TCM pathological detection, and little work has explored automatic cross-mark extraction for disease screening. According to Table 2, palm cross-marks indicate potential pathological changes and support preventive diagnosis. By

jointly segmenting palm-core ROI, subcutaneous veins and cross-marks, our method possesses unique advantages for remote TCM palm diagnosis.

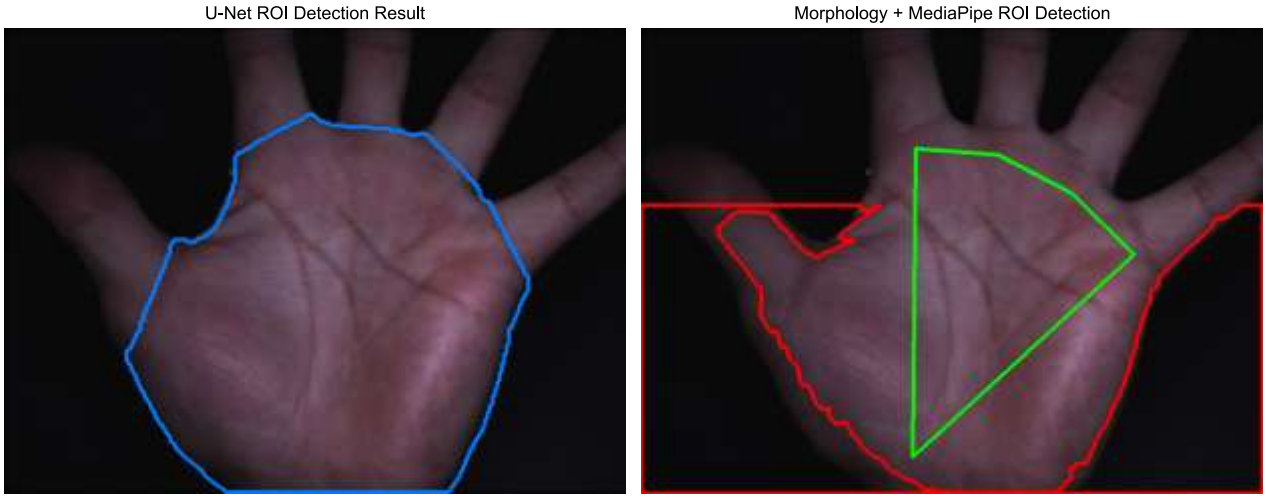


Fig. 8. Comparison of results from three palm-core ROI extraction schemes. Left: multi-feature segmentation results obtained by the proposed U-Net (Blue-ROI); Right: results of two baseline methods (Red: morphology-based method, Green: MediaPipe landmark-based method).

4.8. Algorithm limitations

Cyan veins and cross-marks are faint small-size targets with low pixel proportions, which easily lead to missed detections and over-fitting risks caused by limited annotated samples. Model performance is sensitive to illumination variations, skin-color diversity and shooting angles. Degradation occurs when test samples deviate from the distribution of training data. Although pre-processing including image normalization is adopted, the current model lacks dedicated illumination-elimination modules. Further practical deployment-related limitations are discussed in Section 5.2.

5. Discussion

5.1. Comparative Analysis with Existing Works

Many existing palm-image-based TCM studies mainly focus on extracting coarse-grained main palm lines such as the life-line and head-line, while few researches pay attention to clinically-concerned subtle pathological features including subcutaneous cyan veins and pathological cross-marks. Traditional hand-crafted pipelines combining Canny edge detection, Hough transform and Harris corner detection can produce acceptable outputs only under ideal lighting conditions, but yield high false-positive rates under complex real-world illumination and skin-color interference. Different from multi-task learning schemes that adopt shared encoder weights, this work employs three completely independent U-Net networks for separate segmentation tasks, which avoids mutual feature interference between heterogeneous tiny targets. Furthermore, the ROI-mask-based post-processing strategy effectively suppresses background false-positive blobs outside the valid palm region.

From the software-engineering perspective, most existing published literature only report isolated algorithm results without complete supporting application systems. In contrast, this work embeds segmentation algorithms into an end-to-end remote-diagnosis software prototype covering palm-image acquisition, persistent storage, AI inference, doctor-patient role separation and automatic report generation. Compared with general intelligent medical software that simply encapsulates pre-trained models, our

four-layer loosely-coupled architecture decouples AI algorithm logic from upper-layer business logic. New segmentation models or updated TCM diagnostic rules can be plugged-in without system-wide reconstruction, which improves system scalability. Pre-diagnosis image quality-control modules filter low-quality input samples and reduce useless inference overhead for better reliability. Role-based cross-layer access control strictly isolates patient-side and doctor-side data views to enhance medical-data security. The prototype can run on ordinary consumer-grade CPUs and supports future migration to cloud-server deployment, balancing real-time performance for offline prototype verification.

5.2. Study Limitations

The images adopted for model training are acquired under controlled dark-background environments with dedicated auxiliary lighting equipment from the TJU-PV dataset, a public biometric palmprint dataset rather than a medically-oriented TCM palm-diagnosis dataset; all medical segmentation masks are manually annotated in-house. Significant fluctuation of ambient illumination in practical usage will alter palm texture and subcutaneous vein visual manifestation, further degrading segmentation and recognition accuracy. As a feasible improvement direction, we may standardize the hardware-driven image-acquisition workflow, referring to the non-contact palm-image capture scheme described in literature [26]. Standardized lighting parameters, shooting distance and background constraints following this reference can greatly enhance the consistency of acquired palm images and mitigate illumination-induced performance drop in real-world deployment scenarios.

5.3. Future-work Outlook

For future improvements, we plan to expand expert-annotated palm datasets by collecting samples covering diverse skin tones, illumination conditions and age groups. Online data-augmentation strategies, class-balanced loss functions and few-shot meta-learning will be adopted to mitigate segmentation difficulties for faint small pathological targets. We will introduce illumination-normalization and color-calibration pre-processing modules to reduce interference brought by shooting-device and environmental variations. Cooperation with TCM hospitals will be arranged to carry out prospective clinical validation and iterative optimization of both algorithm modules and application workflows. Moreover, we intend to explore federated-learning frameworks, which allow model iterative updating without centralizing raw sensitive palm-image data, to further strengthen patient-data privacy protection for real-world remote-diagnosis scenarios.

6. Conclusion

Targeting practical requirements of remote TCM palm diagnosis, this paper implements a complete AI-powered remote TCM diagnostic software prototype based on palm-image analysis. A four-layer layered-modular architecture is constructed to realize full-link workflows for data acquisition, storage management, AI computation and doctor-patient interaction. Three independent U-Net networks are adopted for pixel-level segmentation of palm ROI, subcutaneous cyan veins and cross-marks. The BCE-Dice mixed loss alleviates foreground-background class imbalance for small targets. Post-processing with palm-ROI masks suppresses false detections outside palm boundaries. The prototype supports palm-image analysis and generates structured auxiliary diagnostic reports complying with TCM nine-grid zoning rules.

Experiments show satisfactory segmentation performance for palm ROI. Constrained by limited annotated samples, missed detections still exist for faint cyan veins and cross-marks. The presented system is an auxiliary-evaluation prototype and cannot replace professional TCM clinical diagnosis. Future work will expand annotated datasets, optimize small-target segmentation algorithms and promote hospital-co-

operated clinical verification to advance digitization of TCM inspection diagnosis and remote intelligent-medical applications.

Conflict of Interest

The authors declare no conflict of interest.

Author Contributions

D. W. conducted the research, completed the program, wrote the paper; J. Z. conducted the research, analyzed the data results; all authors had approved the final version.

References

[1] Yu, K., & Matsuki, H. (2022). Diagnostic aid of palm image relating to asthma based on convolution neural network. *Journal of Advances in Artificial Life Robotics*, *3(3)*, 151–154.

[2] Wang, Q. D., Zhao, C. Q., Qiang, Y., *et al.* (2021). Multitask interactive attention learning model based on hand images for assisting Chinese medicine in predicting myocardial infarction. *Computational and Mathematical Methods in Medicine*, *2021*, 6046184.

[3] Fan, Y., Li, J. X., Song, S. Y., *et al.* (2022). Palmprint phenotype feature extraction and classification based on deep learning. *Phenomics*, *2(4)*, 219–229.

[4] Zhang, X. W., Lou, Y. M., Zhou, W. X., *et al.* (2024). Exploration of a machine learning based body constitution model for identifying the twenty-five people of Yin and Yang through hand diagnosis. *Journal of Nanjing University of Traditional Chinese Medicine*, *40(12)*, 1323–1330.

[5] Wang, Q. D., Ji, L. W., Qiang, Y., *et al*. (2022). Multi-branch deep feature fusion method for traditional Chinese medicine intervened human cerebral stroke aided diagnosis. *Journal of Image and Graphics*, *27(3)*, 0935–0947.

[6] Nakhua, H., Motiwala, M., & Katre, N. (2024). Palmistry with deep learning approaches using CNN and OpenCV. In *Proceedings of the 2024 International Conference on Artificial Intelligence and Emerging Technology (Global AI Summit)* (pp. 335–340).

[7] Zhou, W. X., Xu, W. M., Zhang, X. Y., *et al.* (2025). Development of objective multimodal information acquisition equipment for traditional Chinese medicine hand diagnosis. *Journal China Journal of Traditional Chinese Medicine and Pharmacy*, *40(1)*, 83–86.

[8] Shao, H. K., Zhong, D. X., Du, X. F., Du, S. Y., & Veldhuis, R. N. J. (2021). Few-shot learning for palmprint recognition via Meta-Siamese network. *IEEE Transactions on Instrumentation and Measurement*, *70*, 5009812.

[9] Lin, C., Chen, Y., Zou, X., Deng, X., Dai, F., You, J., *et al.* (2024). An unconstrained palmprint region of interest extraction method based on lightweight networks. *PLoS ONE*, *19(8)*, e0307822.

[10] Xu, A. Y., Yang, T., Wang, T. S., *et al.* (2024). Research ideas and methods of morphological feature extraction of traditional Chinese medicine inspection. *Journal of Nanjing University of Traditional Chinese Medicine*, *40(12)*, 1315–1322.

[11] Tian, Z. K., Wang, D. J., Sun, X., *et al.* (2023). Current status and trends of artificial intelligence research on the four traditional Chinese medicine diagnostic methods: A scientometric study. *Annals of Translational Medicine*, *11(3)*, 145.

[12] Patil, S. (2025). Palmistry using machine learning. *arXiv Preprint*, arXiv:2509.02248.

[13] Wu, L., Xu, Y., Cui, Z., Zuo, Y., Zhao, S., & Fei, L. (2021). Triple-type feature extraction for palmprint recognition. *Sensors*, *21(14)*, 4896.

[14] Wang, Y. H., Yang, X., Yang, Y., *et al.* (2026). Research progress and challenges in artificial intelligence-

driven intelligent traditional Chinese medicine diagnosis and treatment. *Shanghai Journal of Traditional Chinese Medicine*, *60(1),* 1–11.

[15] Jongjiamdee, K., Pornwonglert, P., Na Bangchang, N., *et al.* (2025). Artificial intelligence in traditional medicine: Evidence, barriers, and a research roadmap for personalized care. *Frontiers in Artificial Intelligence*, *8*, 1659338.

[16] Sheng, W. S., Zheng, Z. L., & Zhu, H. Z. (2025). A novel approach to palm vein image segmentation combining multi-scale convolution and swin-transformer networks. *Scientific Reports*, *15(1)*, 17539. doi: 10.1038/s41598-025-02757-7

[17] Van, T. P., Nguyen, S. T., Doan, L. B., *et al.* (2021). Efficient palm-line segmentation with U-Net context fusion module. *arXiv Preprint*, arXiv:2102.12127.

[18] Sun, X., Zhou, K., & Tan, X. (2026). Palm image segmentation method based on multi-level feature fusion. *Available at SSRN 6296110.* doi: 10.2139/ssrn.6296110

[19] Wu, Y., & Hu, J. (2024). Deep and shallow feature fusion in feature score level for palmprint recognition. *IET Biometrics*, *2024*, 5683547. doi: 10.1049/2024/5683547

[20] Zhang, M., Zhang, X., Yang, T., *et al.* (2025). Intelligent inspection of traditional Chinese medicine: A brief review. In *Proceedings of the 2025 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)* (pp. 5318–5323. doi: 10.1109/BIBM66473.2025.11356347

[21] Zhang, S., Zhang, S. Y., Zhang, C., *et al.* (2026). Sparse image registration-based marker-free hand acupoint localization using TCM template queries. *Sensors*, *26(17)*, 5552. doi: 10.3390/s26175552

[22] Szymkowski, M., & Saeed, K. (2025). Trimodal machine learning based biometrics system. *Scientific Reports*, *15(1)*, 20949. doi: 10.1038/s41598-025-06288-z

[23] Gao, C. R., Yang, Z. Y., Jia, W., *et al.* (2025). Deep learning in palmprint recognition—A comprehensive survey. *arXiv Preprint*, arXiv:2501.01166.

[24] Zhao, K., Shen, L., Zhang, Y. Y., *et al.* (2022). BézierPalm: A free lunch for palmprint recognition. In *European Conference on Computer Vision Workshops* (pp. 20–34). doi: 10.1007/978-3-031-19778-9_2

[25] Marattukalam, F., Abdulla, W., Cole, D., & Gulati, P. (2023). Deep learning-based wrist vascular biometric recognition. *Sensors*, *23(6)*, 3132. doi: 10.3390/s23063132

[26] Zhang, L., Li, L., *et al.* (2017). Towards contactless palmprint recognition: A novel device, a new benchmark, and a collaborative representation based identification approach. In *IEEE International Conference on Signal and Image Processing* (pp. 199–212).

Copyright © 2026 by the authors. This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited (CC BY 4.0).