Physio-Anatomical Prior-Guided Probabilistic Field Learning for Ultrasound-Based Neuro-Fascial Interface Landmark Localization

Guowei Dai[1], Duwei Dai[2], Yi Zhang[3], Yulong Ji[4*], and Hu Chen[1*]

[1] College of Computer Science, Sichuan University, China
[2] Second Affiliated Hospital of Xi'an Jiaotong University, China
[3] School of Cyber Science and Engineering, Sichuan University, China
[4] School of Aeronautics and Astronautics, Sichuan University, China

Abstract. Accurate localization of acupoints in ultrasound imaging is critical for safe and effective acupuncture, yet remains challenging due to low tissue contrast, speckle noise, and significant inter-subject anatomical variability. Existing methods often rely on implicit texture learning or rigid landmark detection, failing to account for the depth-dependent physiological nature of the Deqi response. To address these limitations, we propose the Physio-Anatomical Probability Field Network (PAPF-Net), a novel framework that shifts from deterministic point regression to uncertainty-aware distribution modeling. The core technical contribution lies in synergizing two complementary mechanisms. To address anatomical heterogeneity, the Category-Conditional Anatomical Prototype (CCAP) module explicitly incorporates site-specific priors by deforming learnable templates to conform to individual anatomical geometries. Concurrently, to mitigate acoustic ambiguity, the Uncertainty-Aware Decoder transitions from deterministic regression to probabilistic modeling, capturing signal attenuation in deep tissues through distribution-aware Sinkhorn supervision. Validated across a large multi-center cohort, our approach surpasses current standards in both accuracy and robustness. Its focus on neuro-fascial interfaces ensures physiological plausibility, offering an interpretable, standardized basis for ultrasound-guided acupuncture.

Keywords: Ultrasound Navigation · Landmark Detection · Anatomical Priors · Uncertainty Estimation · Deqi Localization

1 Introduction

Acupuncture has demonstrated notable bio-regulatory effects in the rehabilitation of neurological disorders, the management of chronic pain, and the treatment of functional illnesses [6,19]. The therapeutic efficacy of acupuncture is closely linked to inducing the Deqi state, a key determinant of treatment success [12,19]. Physiologically, Deqi manifests as a composite sensation, subjectively perceived

[4] †Corresponding authors: Hu Chen, huchen@scu.edu.cn; Yulong Ji, jyl@scu.edu.cn.

by the patient as soreness, numbness, distension, or heaviness, and objectively felt by the practitioner as a needle grasp from the surrounding tissue [15,18]. Modern anatomical and interventional ultrasound studies have shown that this phenomenon arises from the mechanical interaction between the needle tip and deep anatomical structures, particularly the neuro-fascial interfaces, including deep fascia, neurovascular bundles, and intermuscular septa.

Automating this localization process is critical for standardizing acupuncture therapy but presents unique computer vision challenges [9,18,7]. Unlike typical landmark detection tasks (e.g., facial keypoints or anatomical landmarks) where targets are visually distinct [13,3], the Deqi response zone is functionally defined and visually ambiguous, often lacking clear boundaries in ultrasound images. The anatomical structures surrounding acupoints (e.g., the metacarpal spacing for Hegu versus the forearm muscles for Shousanli) exhibit significant morphological heterogeneity across subjects and sites, rendering standard template matching ineffective [4,5,8]. The Deqi sensation is inherently probabilistic; the optimal stimulation point is not a Dirac delta function but a spatial field whose spread correlates with tissue depth and patient sensitivity [20]. Deterministic regression models fail to capture this aleatoric uncertainty, leading to overconfident but physiologically implausible predictions.

To address these challenges, we propose the Physio-Anatomical Probability Field Network (PAPF-Net), which reformulates acupoint localization from deterministic coordinate regression to uncertainty-aware probability field learning. Targeting the morphological heterogeneity of musculoskeletal structures, we introduce a Category-Conditional Anatomical Prototype (CCAP) module. This component retrieves site-specific structural priors and dynamically registers them to the input via deformable alignment, explicitly injecting geometric constraints to robustify feature learning against speckle noise. Simultaneously, to capture the aleatoric uncertainty arising from acoustic attenuation, we design a Dual-Head Uncertainty Decoder. By modeling the output as a heteroscedastic Gaussian distribution, the network provides depth-adaptive confidence maps that align with the physical reality of ultrasound propagation. We validate our method on Deqi-US-1k, a curated multi-center dataset of 1,000 subjects with labels verified by patient feedback during needle insertion. The dataset provides functional Deqi response annotations for deep ultrasound analysis. Our main contributions are summarized as follows:

- We propose PAPF-Net, a deep learning framework dedicated to localizing functional Deqi response zones in ultrasound, bridging the gap between traditional medicine and modern computer vision.
- We introduce the CCAP module, which leverages deformable alignment to inject site-specific structural priors, effectively handling anatomical heterogeneity.
- We formulate a depth-adaptive uncertainty learning mechanism that models the Deqi response as a probabilistic field, offering physiologically interpretable confidence maps.

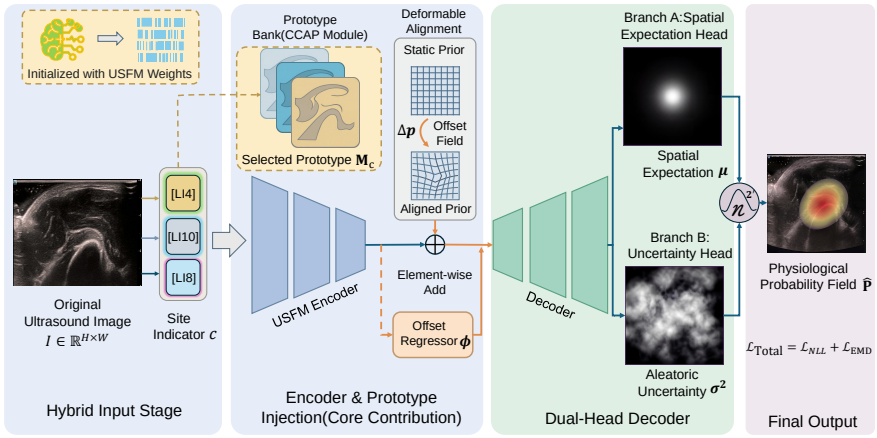

Fig. 1. Overview of PAPF-Net. Given input (I, c), where c is the site indicator, CCAP selects the corresponding learnable prototype $\mathbf{M}_c$ from a bank of K site-specific prototypes and aligns it with offsets $\Delta\mathbf{p}$. A dual-head decoder then regresses $(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$ to generate the final *Deqi* probability field $\hat{\mathbf{P}}$.

2 Method

We propose the Physio-Anatomical Probability Field Network (PAPF-Net) to localize functional Deqi response zones as probabilistic fields rather than deterministic coordinates. As shown in Fig. 1, the architecture synergizes deep acoustic features with explicit anatomical priors through a U-shaped design operating on B-mode images I and a site indicator $c \in \{1, \ldots, K\}$, where K is the number of anatomical categories. To tackle significant morphological heterogeneity across anatomical sites, the bottleneck incorporates a CCAP module. This mechanism retrieves a learnable, site-specific structural template and dynamically aligns it to the input via deformable registration, thereby injecting geometric constraints such as bone curvature and fascial depth into the latent representation. Complementing this structural guidance, a Dual-Head Uncertainty Decoder models observation-dependent aleatoric uncertainty induced by ultrasound attenuation. By jointly regressing a spatial expectation $\boldsymbol{\mu}$ and a depth-adaptive variance map $\boldsymbol{\sigma}^2$, the network constructs a heteroscedastic Gaussian field, ensuring uncertainty-aware predictions in visually ambiguous deep tissue regions.

2.1 From Coordinates to Physiological Fields

Let $\mathcal{D} = \{(I_i, \mathbf{y}_i, c_i)\}_{i=1}^{N}$ denote the dataset, where I_i is the ultrasound image, c_i represents the anatomical site, and $\mathbf{y}_i$ is the ground-truth coordinate verified by functional *Deqi* feedback (Fig. 2). Unlike conventional methods that model landmarks as deterministic Dirac deltas, we postulate the *Deqi* response as a continuous probability density function governed by neuro-fascial mechanics.

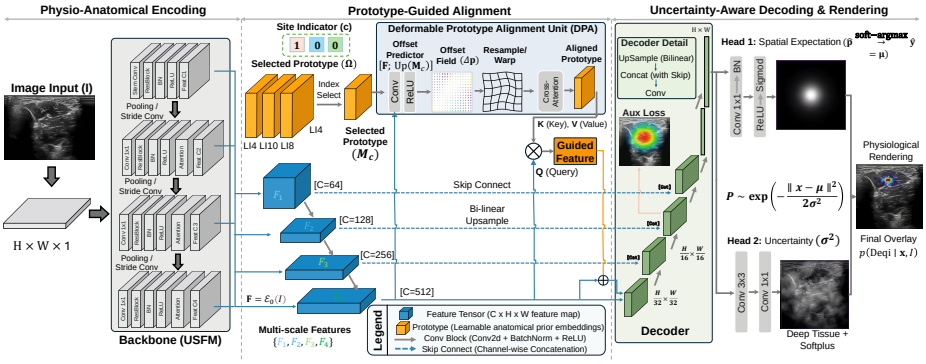

Fig. 2. Mechanisms of PAPF-Net. (a) Prototype alignment to features $\mathbf{F}$ via offsets $\Delta\mathbf{p}$, where $\mathbf{M}_c$ is selected by the site indicator c. (b) Uncertainty: shallow (low σ^2) vs. deep (high σ^2) regions. (c) Sinkhorn supervision minimizes the transport cost between the predicted field $\hat{\mathbf{P}}$ and a normalized Gaussian target heatmap $\mathbf{Q}_i$ centered at $\mathbf{y}_i$; the auxiliary branch applies the same supervision at 1/4 resolution.

Our goal is to learn a mapping $\mathcal{F}_\Theta : (I, c) \mapsto \hat{\mathbf{P}}$ that quantifies this physiological field.

We employ a U-shaped architecture where the encoder $\mathcal{E}_\theta$ extracts hierarchical features $\{\mathbf{F}_\ell\}_{\ell=1}^4$. The bottleneck $\mathbf{F}_4$ is processed by the prototype module (Sec. 2.2), and the decoder $\mathcal{D}_\psi$ progressively aggregates features via skip connections. Crucially, to capture observation-dependent aleatoric uncertainty arising from acoustic attenuation, the decoder terminates in two coupled heads: (1) an *Expectation Head* producing the normalized field $\hat{\mathbf{P}}$, from which the spatial mean is derived via differentiable soft-argmax as $\hat{\mathbf{y}} = \sum_{\mathbf{x}} \mathbf{x}\,\hat{\mathbf{P}}(\mathbf{x})$; and (2) an *Uncertainty Head* predicting a pixel-wise variance map $\hat{\sigma}^2(\mathbf{x})$ via softplus activation. The uncertainty head does not estimate epistemic uncertainty over model weights; instead, it models image-dependent ambiguity caused by attenuation and shadowing, and we use the variance at the predicted location, $\hat{\sigma}_i^2 = \hat{\sigma}^2(\hat{\mathbf{y}}_i)$, in Eq. 1.

$$p(\mathbf{y}_i \mid I_i) = \mathcal{N}\big(\mathbf{y}_i;\, \hat{\mathbf{y}}_i,\, \hat{\sigma}_i^2\big)\,, \tag{1}$$

where $\hat{\sigma}_i^2$ is the sample-wise variance predicted by the uncertainty head. This formulation allows the network to adaptively relax localization precision in regions with high anatomical complexity or shadowing, transforming rigid regression into uncertainty-aware estimation.

2.2 Category-Conditional Anatomical Prototype Injection

Anatomical contexts vary drastically across acupoints (e.g., metacarpal LI4 versus radial-ulnar LI10), rendering single global priors ill-posed. To capture this heterogeneity, we propose a CCAP mechanism.

Learnable Prototype Bank. We construct a bank $\mathcal{M} = \{\mathbf{M}_k\}_{k=1}^K$ where K is the number of site categories and each prototype $\mathbf{M}_k$ encodes a site-specific

latent structural template (Fig. 2). The site label $c \in \{1, \ldots, K\}$ is represented as a one-hot indicator in implementation, and the corresponding prototype $\mathbf{M}_c$ is selected for each sample. In our implementation, each prototype is a randomly initialized learnable tensor with the same shape as the bottleneck feature map and is optimized end-to-end with the network. To ensure distinctiveness and prevent mode collapse, we enforce an orthogonality regularization on the prototype embeddings:

$$\mathcal{L}_{\text{ortho}} = \sum_{k \neq k'} \langle \text{vec}(\mathbf{M}_k), \text{vec}(\mathbf{M}_{k'}) \rangle^2 . \tag{2}$$

Deformable Alignment and Injection. Direct fusion of canonical prototypes with instance features $\mathbf{F}$ is suboptimal due to variations in patient pose and probe orientation. We address this via a Deformable Prototype Alignment (DPA) module. A lightweight network $\mathcal{G}_\phi$ first estimates a dense offset field $\Delta\mathbf{p}$ based on the correlation between the feature map and the upsampled prototype:

$$\Delta\mathbf{p} = \mathcal{G}_\phi([\mathbf{F};\ \text{Up}(\mathbf{M}_c)]). \tag{3}$$

The prototype is then spatially warped to match the instance anatomy via a differentiable sampling operation $\mathcal{W}$, yielding $\widetilde{\mathbf{M}}_c = \mathcal{W}(\mathbf{M}_c, \Delta\mathbf{p})$. Finally, to rigorously constrain the feature learning, the aligned prototype is injected into the visual stream via a gated cross-attention mechanism:

$$\mathbf{F}_{\text{guided}} = \mathbf{F} + \alpha \cdot \text{CrossAttn}(\mathbf{F}, \widetilde{\mathbf{M}}_c, \widetilde{\mathbf{M}}_c), \tag{4}$$

where α is a zero-initialized learnable scalar. This design allows the network to selectively attend to relevant structural priors (e.g., bone boundaries) while gradually adapting to patient-specific morphology during training.

2.3 Training Objective and Notation

The full training objective is written explicitly to avoid ambiguity. For each sample, the landmark coordinate $\mathbf{y}_i$ is converted into a normalized Gaussian target heatmap $\mathbf{Q}_i$ at the supervision resolution:

$$\mathbf{Q}_i(\mathbf{x}) = \frac{\exp\left(-\|\mathbf{x} - \mathbf{y}_i\|_2^2 / 2\tau^2\right)}{\sum_{\mathbf{x}'} \exp(-\|\mathbf{x}' - \mathbf{y}_i\|_2^2 / 2\tau^2)}. \tag{5}$$

The Sinkhorn term is computed between the predicted probability field and $\mathbf{Q}_i$, and the auxiliary branch applies the same supervision at $1/4$ resolution:

$$\mathcal{L}_{\text{Sinkhorn}} = \text{Sinkhorn}(\hat{\mathbf{P}}_i, \mathbf{Q}_i) + \text{Sinkhorn}(\hat{\mathbf{P}}_i^{1/4}, \mathbf{Q}_i^{1/4}). \tag{6}$$

The heteroscedastic Gaussian negative log-likelihood is applied to the coordinate prediction:

$$\mathcal{L}_{\text{NLL}} = -\log \mathcal{N}\left(\mathbf{y}_i;\ \hat{\mathbf{y}}_i,\ \hat{\sigma}_i^2\right), \tag{7}$$

where $\hat{\mathbf{y}}_i = \sum_{\mathbf{x}} \mathbf{x}\, \hat{\mathbf{P}}_i(\mathbf{x})$. The final objective is

$$\mathcal{L}_{\text{total}} = \lambda_{\text{nll}} \mathcal{L}_{\text{NLL}} + \lambda_{\text{sink}} \mathcal{L}_{\text{Sinkhorn}} + \lambda_{\text{reg}} \mathcal{L}_{\text{ortho}}. \tag{8}$$

Here, λ_{nll}, λ_{sink}, and λ_{reg} are the only optimization weights reported in the implementation details; the auxiliary branch is included inside $\mathcal{L}_{\text{Sinkhorn}}$ rather than treated as an additional independent term.

3 Experiments

3.1 Experimental Setup

Dataset and Protocols We validate our method on Deqi-US-1k, a large-scale ultrasound dataset from a multi-center randomized controlled trial coordinated by a collaborating hospital. The cohort comprises 1,000 subjects with peripheral facial paralysis. High-resolution B-mode sequences were acquired using Mindray Resona 7/9 systems. The dataset targets three acupoints with distinct anatomical contexts: *Hegu* (LI4, $N = 330$), *Shousanli* (LI10, $N = 330$), and *Xialian* (LI8, $N = 340$). A rigorous function-driven protocol was enforced: ground truth coordinates $\mathbf{y}_i$ were annotated at the needle tip position only upon patient confirmation of the Deqi sensation, ensuring labels reflect physiological response zones rather than surface approximations.

Implementation Details. The proposed framework is implemented in PyTorch and trained on 4× NVIDIA A100 GPUs (80GB). We use the AdamW optimizer with a base learning rate of 5×10^{-4} and weight decay of 1×10^{-4}. The learning rate follows a cosine annealing schedule with 5 warm-up epochs. Input images are resized to 384×384. Data augmentation includes random rotation ($\pm 15°$), scaling ($0.8 - 1.2\times$), and intensity jittering. Loss weights are set to $\lambda_{nll} = 1.0$, $\lambda_{sink} = 0.5$, and $\lambda_{reg} = 0.1$; the auxiliary 1/4-resolution branch uses the same Sinkhorn supervision and is absorbed into λ_{sink}. The prototype bank contains one randomly initialized tensor per site, and the backbone is initialized with pre-trained weights from the Universal Ultrasound Foundation Model (USFM) [10]. Training takes approximately 4 hours for 100 epochs with a total batch size of 32.

Evaluation Metrics We employ standard Mean Radial Error (MRE) and Successful Detection Rate (SDR) with thresholds $\delta \in \{0.5, 1.0, 1.5\}$ mm to evaluate geometric accuracy. To rigorously validate physiological plausibility, we propose the Anatomical Layer Hit Rate (ALHR). Unlike geometric distance, ALHR measures whether the predicted peak $\hat{\mathbf{y}}_i$ falls within the correct neuro-fascial tissue layer (e.g., deep fascia) defined by expert-annotated masks $\mathcal{S}_i$. It is a task-specific layer-consistency metric rather than a generic segmentation precision score. It is formulated as:

$$\text{ALHR} = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}\left(\hat{\mathbf{y}}_i \in \mathcal{S}_i\right) \times 100\% \tag{9}$$

where N is the number of test samples and $\mathbb{I}(\cdot)$ is the indicator function returning 1 if the condition holds, and 0 otherwise.

Table 1. Performance comparison of different landmark localization methods. Lower MRE is better, while higher SDR and ALHR are better. The best results are highlighted in **bold**, and the second-best results are underlined. The reported values represent the mean ± standard deviation across a 5-fold cross-validation scheme.

Method	MRE (mm) ↓		SDR (%) ↑			ALHR (%) ↑
	L_1	L_2	0.5 mm	1.0 mm	1.5 mm	
U-Net [11]	0.51 ± 0.034	0.68 ± 0.052	64.80 ± 2.81	92.17 ± 0.72	97.40 ± 1.37	60.5
DM-ResNet [16]	0.47 ± 0.031	0.62 ± 0.045	71.25 ± 2.14	94.30 ± 0.65	98.12 ± 0.98	66.8
HRNet [14]	0.45 ± 0.028	0.59 ± 0.041	74.60 ± 1.95	95.55 ± 0.58	98.80 ± 0.82	69.4
TransUNet [2]	0.44 ± 0.026	0.57 ± 0.038	76.12 ± 1.88	96.02 ± 0.51	99.05 ± 0.65	71.2
FARNet [1]	0.43 ± 0.024	0.56 ± 0.033	78.35 ± 1.76	96.88 ± 0.45	99.20 ± 0.55	73.5
TGCN-ICF [7]	0.41 ± 0.021	0.53 ± 0.029	81.50 ± 1.62	97.45 ± 0.40	99.45 ± 0.42	76.8
EDA [17]	0.40 ± 0.019	0.52 ± 0.027	82.85 ± 1.55	97.90 ± 0.38	99.60 ± 0.35	78.1
IT-UNet [8]	0.39 ± 0.017	0.49 ± 0.022	84.15 ± 1.48	98.10 ± 0.35	99.68 ± 0.31	79.8
CalibratedSL [3]	0.38 ± 0.013	0.50 ± 0.025	84.90 ± 1.50	98.80 ± 0.32	99.75 ± 0.28	80.5
DGMLD [4]	0.39 ± 0.015	0.51 ± 0.024	85.60 ± 1.40	98.65 ± 0.28	99.80 ± 0.22	81.2
PAPF-Net (Ours)	$\mathbf{0.36 \pm 0.011}$	$\mathbf{0.46 \pm 0.018}$	$\mathbf{88.45 \pm 1.25}$	$\mathbf{99.15 \pm 0.20}$	$\mathbf{99.92 \pm 0.10}$	**83.6**

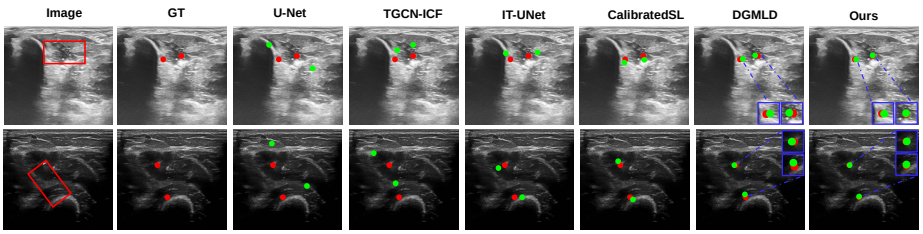

Fig. 3. Qualitative comparison showing PAPF-Net robust localization of neuro-fascial interfaces under severe acoustic shadowing and ambiguous multi-layer textures, surpassing competing methods limited by local speckle noise.

3.2 Comparison with State-of-the-Art

Table 1 compares our PAPF-Net against ten leading methods, spanning generic medical segmentation and specialized landmark detection. Our method consistently achieves the best performance across all metrics. Specifically, it reduces the Mean Radial Error (L_2) to 0.46 mm, outperforming the second-best DGMLD by a significant margin of 9.8%. For the strict 0.5 mm threshold, PAPF-Net attains an SDR of 88.45%, demonstrating superior fine-grained localization in speckle-heavy ultrasound images. Notably, while generic CNNs struggle with low-contrast fascial interfaces (MRE $>$ 0.6 mm), specialized models like HR-Net improve geometric consistency but lack explicit structural guidance. The gains of PAPF-Net mainly come from CCAP and uncertainty-aware modeling. The 83.6% ALHR surpasses the nearest competitor by 2.4%. This metric indicates that our probabilistic field approach is not only geometrically precise but also physiologically plausible, correctly targeting the deep neuro-fascial response zones rather than superficial tissues.

3.3 Qualitative Analysis

Fig. 3 visualizes the localization performance on challenging samples characterized by severe acoustic shadowing and ambiguous multi-layer fascial textures. Generic methods (e.g., U-Net, TGCN-ICF) exhibit significant spatial deviations, often gravitating towards high-contrast bone surfaces or incorrect muscle layers due to their reliance on local texture features, which are easily corrupted by speckle noise. Even specialized baselines like DGMLD show minor misalignments in layer selection. PAPF-Net aligns more closely with the ground truth in the zoomed-in regions. The CCAP module imposes global structural constraints that reduce local anatomical ambiguity, and the uncertainty-aware mechanism suppresses false positives in noisy regions. As a result, the predictions target the physiologically correct neuro-fascial interface rather than merely salient image features.

Table 2. Component-wise ablation study on the Deqi-US-1k dataset. **Proto**: Static Anatomical Prototype Injection; **DPA**: Deformable Prototype Alignment; **Uncert**: Uncertainty-Aware Head; **EMD**: Sinkhorn supervision between the predicted field and the normalized Gaussian target heatmap.

Model ID	Components				MRE	SDR (%) ↑			ALHR
	Proto	DPA	Uncert	EMD	(mm) ↓	0.5 mm	1.0 mm	1.5 mm	(%) ↑
Baseline	–	–	–	–	0.58 ± 0.04	64.2	92.5	97.1	68.4
Static	✓	–	–	–	0.52 ± 0.03	71.8	94.8	98.2	74.2
Deform	✓	✓	–	–	0.45 ± 0.02	79.5	97.2	99.0	79.6
Prob.	✓	✓	✓	–	0.40 ± 0.02	84.1	98.4	99.5	81.8
PAPF-Net	✓	✓	✓	✓	**0.36 ± 0.01**	**88.5**	**99.2**	**99.9**	**83.6**

3.4 Ablation Study

Table 2 dissects the contribution of each proposed module. The Baseline model, relying solely on visual features, yields a suboptimal MRE of 0.58 mm and a low ALHR of 68.4%, indicating that implicit texture learning is insufficient to overcome the inherent speckle noise and low contrast of fascial interfaces. Injecting a Static Prototype provides global anatomical guidance, reducing the MRE to 0.52 mm. However, the most significant performance leap occurs with the DPA, which boosts the SDR@0.5mm by 7.7% and ALHR by 5.4%. This confirms that rigid priors are inadequate for handling inter-subject morphological heterogeneity. The DPA module effectively warps the canonical template to match patient-specific bone curvature, ensuring the features are strictly constrained within the valid anatomical manifold. Incorporating the Uncertainty-Aware Head further reduces the MRE to 0.40 mm. By modeling observation-dependent aleatoric uncertainty, the network learns to distrust ambiguous signals in deep tissues and to focus on high-confidence structural landmarks. Finally, introducing the EMD

achieves the best overall performance. Unlike pixel-wise regression, which penalizes local coordinate shifts rigidly, EMD optimizes the global geometry of the predicted probability field against the Gaussian target heatmap, ensuring that the Deqi response distribution aligns with the ground-truth neuro-fascial pathway.

4 Conclusion

In this paper, we present PAPF-Net, a novel framework that redefines ultrasound acupoint localization by integrating explicit anatomical priors with uncertainty-aware modeling. Unlike conventional regression methods limited by speckle noise, our approach leverages deformable prototypes to disentangle complex inter-subject anatomical variations and utilizes depth-adaptive probability fields to capture the physiological nature of the Deqi response. Extensive validation on a multi-center dataset demonstrates that PAPF-Net achieves superior geometric accuracy and robustly targets the correct neuro-fascial interfaces. By bridging the gap between data-driven vision and physiological reality, this work provides a reliable, interpretable foundation for safe ultrasound-guided automated acupuncture. The current evaluation covers three clinically distinct acupoint categories; broader anatomical coverage and prospective external validation remain future work.

Acknowledgments This work was supported in part by the Sichuan Natural Science Foundation under Grant 2024NSFSC0725.

Disclosure of Interests The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ao, Y., Wu, H.: Feature Aggregation and Refinement Network for 2D Anatomical Landmark Detection. Journal of Digital Imaging **36**(2), 547–561 (Nov 2022)
2. Chen, J.: TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation. In: Proceedings of the International Conference on Machine Learning Workshop on Interpretable Machine Learning in Healthcare. pp. 1–13 (2021)
3. Feng, Y., Yang, J., Tang, L.e.a.: Uncertainty Quantification and Quality Control for Heatmap-Based Landmark Detection Models. IEEE Transactions on Medical Imaging **44**(8), 3451–3463 (Aug 2025)
4. Gong, H., Lu, Y., Wan, X.e.a.: Domain Generalized Medical Landmark Detection via Robust Boundary-Aware Pre-Training. Proceedings of the AAAI Conference on Artificial Intelligence **39**(3), 3140–3148 (Apr 2025)
5. Gong, H., Wan, B., Kang, L.e.a.: Boundary as the Bridge: Toward Heterogeneous Partially-Labeled Medical Image Segmentation and Landmark Detection. IEEE Transactions on Medical Imaging **44**(7), 2747–2756 (Jul 2025)

6. Han, R., Hu, J.: Acupuncture: An Overview on Its Functions, Meridian Pathways and Molecular Mechanisms. The American Journal of Chinese Medicine **52**(05), 1215–1244 (Jan 2024)
7. Huang, T., Shi, J., Jin, G., Li, J., Wang, J., Du, J., Shi, J.: Topological GCN for Improving Detection of Hip Landmarks from B-Mode Ultrasound Images. In: Linguraru, M.G., Dou, Q., Feragen, A.e.a. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, vol. 15005, pp. 692–701. Springer Nature Switzerland, Cham (2024)
8. Huang, T., Shi, J., Li, J.e.a.: Involution Transformer Based U-Net for Landmark Detection in Ultrasound Images for Diagnosis of Infantile DDH. IEEE Journal of Biomedical and Health Informatics **28**(8), 4797–4809 (Aug 2024)
9. Kim, J.H., Noh, Y., Lee, H.e.a.: Toward automated detection of microbleeds with anatomical scale localization using deep learning. Medical Image Analysis **101**, 103415 (Apr 2025)
10. Ma, C., Jiao, J., Liang, S., Fu, J., Wang, Q., Li, Z., Wang, Y., Guo, Y.: Tinyusfm: Towards compact and efficient ultrasound foundation models. IEEE Journal of Biomedical and Health Informatics (2026)
11. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W.M.e.a. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, vol. 9351, pp. 234–241. Springer International Publishing, Cham (2015)
12. Sun, W.Z., Chen, C.C., Lin, J.G.: When sng meets acupuncture – a paradigm-shift biomarker for translational research. Journal of Traditional and Complementary Medicine **15**(1), 1–5 (Jan 2025)
13. Wang, A., Elbatel, M., Liu, K.e.a.: Geometric-Guided Few-Shot Dental Landmark Detection with Human-Centric Foundation Model. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2025, vol. 15964, pp. 197–207. Springer Nature Switzerland, Cham (2026)
14. Wang, J., Sun, K., Cheng, T.e.a.: Deep High-Resolution Representation Learning for Visual Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence **43**(10), 3349–3364 (Oct 2021)
15. Wu, J., Zhang, J., Gao, B.e.a.: Temporal Characteristics of Microstate Changes Representing the Neuromodulation Effect of Acupuncture at PC6 With Twirling Manipulation. IEEE Transactions on Computational Social Systems pp. 1–13 (2025)
16. Xu, J., Xie, H., Liu, C.e.a.: Hip Landmark Detection With Dependency Mining in Ultrasound Image. IEEE Transactions on Medical Imaging **40**(12), 3762–3774 (Dec 2021)
17. Yang, F., Zhou, S.K.: EDA: Enhanced Domain-Adversarial Training for Anatomical Landmark Detection. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE, Houston, TX, USA (Apr 2025)
18. Yang, S., Li, Y., Zou, H.e.a.: Exploring an Innovative Deep Learning Solution for Acupuncture Point Localization on the Weak Feature Body Surface of the Human Back. IEEE Journal of Biomedical and Health Informatics **29**(7), 4599–4611 (Jul 2025)
19. Yu, H., Lin, Z., Li, F.e.a.: Spatiospectral Representation and Neural Decoding of Somatic Perception of Acupuncture Stimulations. IEEE Journal of Biomedical and Health Informatics pp. 1–14 (2025)
20. Zhong, Z., Yao, L., Liu, Y.Z., et al.: Objectivization study of acupuncture Deqi and brain modulation mechanisms: a review. Frontiers in Neuroscience **18**, 1386108 (May 2024)