FSSM-ResNet: Hair-Aware Skin Lesion Segmentation with Embedded Fuzzy Structural Suppression

Ayaulym Raikhankyzy*, Amira Toksanbay*, Saltanat Zarkhinova*, Rimma Kubanova*, Nail Fakhrutdinov*, Symbat Bayakhmetova†, Anel Amantayeva†, Adnan Yazici*
*Department of Computer Science, Nazarbayev University, Astana, Kazakhstan
{adnan.yazici, ayaulym.raikhankyzy, amira.toksanbay, saltanat.zarkhinova, rimma.kubanova,nail.fakhrutdinov}@nu.edu.kz
†Department of Biomedical Sciences, Nazarbayev University, Astana, Kazakhstan
{symbat.bayakhmetova, anel.amantayeva}@nu.edu.kz

***Abstract*—Skin lesion segmentation is a key step in computer-aided skin cancer analysis, yet dermoscopic images often contain hair that obscures lesion boundaries and degrades segmentation performance. Most existing methods address hair via a separate preprocessing stage (e.g., hair removal or inpainting), which complicates the pipeline and can introduce additional artifacts. In this paper, we propose FSSM-ResNet, a hair-aware segmentation model that embeds a Fuzzy Structural Suppression Module (FSSM) directly within the network. FSSM exploits simple, interpretable fuzzy logic based on typical hair characteristics, dark, thin, and linear, to generate a hair-likelihood map and selectively suppress hair regions during decoding, eliminating the need for external preprocessing. To rigorously assess robustness to hair, we introduce VHI (Very Hairy Images), a manually annotated benchmark of 936 strongly hair-occluded dermoscopic images collected from multiple public datasets. Experiments on PH2 and ISIC 2016–2018 show that this fuzzy logic–driven structural suppression not only improves performance on standard benchmarks, but also maintains superior reliability under severe hair occlusion on VHI.**

***Index Terms*—skin lesion segmentation, fuzzy logic, ResNet, attention mechanism, U-Net,**

I. Introduction

Skin cancer is among the most prevalent and rapidly increasing malignancies worldwide, affecting all populations [1]. According to Global Burden of Disease estimates, melanoma and non-melanoma skin cancers jointly account for roughly 128,000 deaths per year and contribute substantially to morbidity and disability-adjusted life years [2]. Clinically, skin lesions span a broad spectrum from benign growths to highly malignant and aggressive tumors. Benign lesions typically exhibit slow growth, symmetry, uniform pigmentation, and well-defined borders, whereas malignant lesions arise from uncontrolled cellular proliferation and often display asymmetry, irregular borders, color variegation, and rapid expansion [3]. Although initially confined to the skin, these cancers can metastasize to distant organs, making early detection critical. Timely and accurate diagnosis of malignant skin lesions can improve five-year survival rates by up to 90% and reduce the need for extensive surgery, chemotherapy, and radiotherapy, thereby lowering treatment-related morbidity and overall socio-economic burden [4].

Distinguishing malignant from benign skin lesions is challenging even for experienced dermatologists, given the subtlety and variability of visual patterns. Dermoscopic interpretation and histopathological assessment are both time-consuming and subject to inter-observer variability [3]. In addition, access to these diagnostic modalities and trained specialists is often limited in rural or underserved regions. Integrating automated diagnostic tools into broader smart e-health and remote monitoring frameworks [5] can help bridge this healthcare gap. Together, these factors highlight the need for reliable computer-assisted tools that can support clinicians in evaluating dermoscopic images and detecting patterns indicative of malignancy.

A critical prerequisite for reliable computer-aided diagnosis is accurate lesion segmentation, since boundary- and shape-related cues directly influence clinically interpretable criteria such as asymmetry, border irregularity, and diameter. However, dermoscopic images frequently exhibit occlusion by body hair, which appears as dark, thin, elongated structures crossing lesion borders. These artifacts can mislead segmentation models by introducing uncertainty and discontinuities, and most existing pipelines address the problem via explicit hair removal or inpainting as a separate preprocessing step [6]–[8]. While sometimes effective, such preprocessing adds computational overhead and may introduce irreversible distortions that propagate to downstream analysis.

In this work, our main goal is to develop a *hair-robust and interpretable* skin lesion segmentation pipeline without explicit hair preprocessing. To achieve this, we propose FSSM-ResNet, an encoder–decoder architecture that integrates fuzzy reasoning about hair directly into the segmentation network. Inspired by fuzzy attention mechanisms in medical image segmentation [9], the proposed Fuzzy Structural Suppression Module (FSSM) computes a continuous hair-likelihood map from dermatologically motivated cues and suppresses hair-corrupted regions during decoding. This design preserves end-to-end efficiency, improves robustness under severe hair occlusion, and maintains interpretability through explicitly defined fuzzy memberships and rules.

The main motivation for this work stems from the gap between clinical needs and current technical solutions for dermoscopic analysis. Early and accurate detection of malignant skin lesions critically depends on reliable lesion segmentation, since shape- and border-based criteria directly influence diagnostic decisions. However, dermoscopic images are frequently occluded by body hair, dark, thin, elongated structures crossing lesion boundaries, that can severely degrade both traditional and deep learning–based segmentation methods.

Most existing pipelines treat hair as a nuisance to be removed in a separate preprocessing step (e.g., morphological filtering, inpainting, or GAN-based repair). While sometimes effective, such approaches introduce several drawbacks: additional computational overhead, brittle binary hair masks that fail on thin or low-contrast hairs, and irreversible distortions of fine lesion texture that propagate to downstream classifiers. Recent fuzzy and attention-based segmentation networks improve robustness to generic uncertainty but do not explicitly model hair artifacts in dermoscopy, and fuzzy hair detection is typically confined to standalone preprocessing rather than being integrated into the segmentation model.

These limitations motivate a shift from "hair removal" to *hair-aware segmentation*: embedding interpretable fuzzy hair modeling directly into the network so that hair is treated as a structured cue rather than mere noise. By learning a continuous hair-likelihood map and using it to modulate decoder features, the model can suppress hair responses while preserving diagnostically important lesion boundaries. At the same time, there is a clear need for a dedicated benchmark of strongly hair-occluded images (VHI) to rigorously stress-test hair robustness beyond standard, relatively clean datasets.

The main contributions of this paper are summarized as follows:

- We introduce FSSM-ResNet, a segmentation architecture that embeds fuzzy hair suppression directly into the network, eliminating the need for separate hair-removal or inpainting stages.
- We design a Mamdani-style fuzzy inference module with learnable Gaussian memberships and compare single-rule vs. multi-rule variants, enabling an explicit interpretability–performance trade-off.
- We construct VHI (Very Hairy Images), a manually annotated set of 936 severely hair-occluded dermoscopic images from multiple public sources, and use it as an external stress test for hair robustness.

The remainder of this paper is organized as follows. Section II reviews related work on hair handling and fuzzy/attention-based medical image segmentation. Section III describes the proposed FSSM-ResNet architecture and the Fuzzy Structural Suppression Module in detail. Section IV presents the datasets, the new VHI benchmark, implementation details, and experimental results, including ablation studies and comparisons with state-of-the-art methods. Finally, Section V summarizes our findings and outlines directions for future research.

II. Related Work

Hair detection remains a critical preprocessing bottleneck in automated dermoscopic image analysis, directly impacting the accuracy of lesion segmentation and the preservation of clinically relevant textural patterns. Traditional approaches rely on morphological filtering [10], edge-based operators [11], intensity-thresholding heuristics [12], and curvilinear structure modeling [11]. Although effective in relatively constrained settings, these methods struggle in challenging scenarios such as thin, overlapping, or low-contrast hairs, particularly when hair structures resemble lesion textures, thereby interfering with accurate lesion boundary segmentation and feature extraction [13], [14].

Sequential pipelines typically follow hair detection with post-processing or repair strategies such as interpolation, inpainting, or patch-based synthesis [15]–[17]. Although deep learning–based reconstruction methods can produce visually plausible hair-free images, they may smooth or blur high-frequency lesion details and generate textures that differ from the underlying anatomy, potentially affecting subsequent segmentation performance [12]. Although reconstruction techniques can generate visually plausible hair-free images, they often produce blurring and artifacts when filling large masked regions, indicating potential loss of fine structural details [17]. In addition, comprehensive surveys of computer-aided melanoma analysis note that AI-based diagnostic pipelines are often computationally intensive and face challenges in preserving detailed feature information during preprocessing and segmentation [18].

Recent deep learning approaches aim to mitigate these limitations by integrating hair-related cues directly into the segmentation pipeline or by learning robust feature representations that are less sensitive to hair artifacts. Encoder–decoder architectures such as ChimeraNet [13] employ multi-scale feature fusion to learn detailed hair morphology, while segmentation networks such as ADF-Net [19] leverage adaptive attention mechanisms to implicitly suppress artifact-corrupted features during lesion segmentation, thereby improving robustness and segmentation performance without relying on explicit hair mask generation. Similarly, fuzzy attention–based models like FA-SegNet [20] employ adaptive fuzzy weighting to emphasize important regions, enhancing robustness in low-contrast or ambiguous areas.

Fuzzy logic provides a natural framework to address these challenges. By assigning continuous membership values, fuzzy systems avoid hard binary decisions, better capturing low-contrast or partially occluded structures [12]. In medical imaging, fuzzy clustering techniques such as Maximum Variance Fuzzy Clustering (MVFC) have been used to identify hair-like regions by integrating color, intensity, and structural cues into graded memberships [12]. Yet, such methods remain preprocessing steps, feeding memberships into classical repair algorithms rather than directly informing segmentation networks.

Emerging fuzzy attention networks and fuzzy-integrated architectures demonstrate how these principles can be embedded into deep learning models to modulate feature responses according to uncertainty. FuzzyNet [21] uses fuzzy attention for polyp segmentation, dynamically weighting features based on local ambiguity. FuzzyTransNet [22] fuses fuzzy transformer layers for rectal polyp and skin lesion segmentation, effectively capturing uncertain boundaries. FCAU-Net [23] employs frequency-channel fuzzy attention to highlight informative regions while suppressing noise, improving robustness in low-contrast and texture-variable areas. Complementary work in fuzzy attention models, such as FA-SegNet [20], demonstrates how differentiable fuzzy weighting can enhance medical image segmentation by adaptively emphasizing relevant features in low-contrast or ambiguous regions.

Despite these advances, no existing work explicitly addresses hair-aware fuzzy segmentation in dermoscopy. In this context, hair artifacts can be modeled as a continuous fuzzy set H, where $\mu_H(x) \in [0, 1]$ quantifies the degree to which a pixel x belongs to hair, informed by local intensity, gradient magnitude, directional curvature, and contextual lesion probability. Some recent studies integrate hair detection and lesion segmentation within a unified deep architecture [24], but these models typically rely on separate processing streams for hair and lesion boundaries. We argue that injecting fuzzy membership information directly into the segmentation decoder, e.g., through fuzzy attention mechanisms or fuzzy-weighted loss functions, would support joint optimization of hair representation and lesion contours, thereby obviating the need for separate repair modules. Such designs could better preserve high-frequency textural details while reducing computational overhead. This direction represents a natural evolution from hair removal to hair-aware segmentation, aligns with uncertainty-aware modeling principles, and promotes improved interpretability and diagnostic fidelity in clinical dermoscopy.

III. Methodology

This section presents the proposed FSSM-ResNet framework. We first describe the overall encoder–decoder architecture and baseline backbone, then introduce the Fuzzy Structural Suppression Module (FSSM), including its hair-likelihood modeling and feature suppression strategy. Finally, we outline the training setup, loss functions, and implementation details used to optimize the network.

A. Overall pipeline

Our proposed solution Fig.1 focuses on hair-robust lesion segmentation. Given a dermoscopic RGB image x, we predict a binary lesion mask $\hat{m}$ which can be further used for classification tasks. Hair artifacts are a major obstacle in this step because hairs can occlude diagnostically relevant lesion boundaries and internal structures [6]–[8]. Our key idea is to avoid external hair-removal preprocessing and instead make the segmentation network *hair-aware by design*.

B. Skin lesion segmentation

FSSM-ResNet adopts an encoder–decoder architecture in which a ResNet backbone (ResNet34/ResNet50) extracts multi-scale features, and a U-Net–style decoder progressively upsamples them with skip connections for precise localization [25], [26]. Unlike a standard U-Net decoder, each decoder block in FSSM-ResNet incorporates a lightweight Fuzzy Self-Attention (FSA) module and is modulated by the predicted hair-likelihood map $H(x)$ produced by FSSM.

C. Relation to F^{2}CAU-Net (baseline)

Our baseline is F^{2}CAU-Net [9], which introduces fuzzy modeling to improve robustness to *uncertainty* in medical images. In particular, the authors emphasize fuzzy components to handle ambiguous boundaries and noise, stating that they integrate a fuzzy convolution module and a fuzzy attention mechanism to improve robustness [9]. However, F^{2}CAU-Net is not designed for *dermoscopic hair artifacts* specifically, which are among the most disruptive artifacts for segmentation and structure detection [7]. We therefore retain the general motivation (uncertainty-aware modeling) but add a *task-specific fuzzy hair module* and redesign attention for efficiency on high-resolution dermoscopy.

D. Fuzzy Structural Suppression Module (FSSM)

Classical artifact-handling pipelines typically apply explicit hair removal *before* segmentation, e.g., DullRazor [6] or SharpRazor [7]. While effective, such preprocessing increases runtime and can introduce inpainting errors that propagate into segmentation and downstream diagnosis. In contrast, FSSM is an internal module that produces a *continuous* hair-likelihood map and suppresses hair responses directly inside the network, enabling end-to-end training.

a) Dermatological prior.: Hairs in dermoscopy appear prototypically *dark, thin, and elongated*. This is consistent with early hair-removal literature as they report that dark thick hairs confused their program, resulting in unsatisfactory segmentation [6] and later work emphasizing the criticality of hair removal for lesion analysis [7], [8]. We encode this prior using a Mamdani-style fuzzy inference rule with interpretable antecedents [27].

b) Feature maps.: From the input RGB image $x \in [0, 1]^{3\times H\times W}$ we compute three normalized maps: (i) grayscale intensity $I(x)$, (ii) gradient magnitude $G(x)$ using fixed Sobel kernels, and (iii) a learned line-response map $L(x)$ using a lightweight ConvNet.

c) Gaussian membership functions.: We define three linguistic sets `Dark`, `Edge-like`, and `Line-strong` using Gaussian membership functions

$$\mu(z; c, \sigma) = \exp\left(-\frac{(z - c)^2}{2\sigma^2}\right), \qquad (1)$$

with learnable (c, σ) initialized to plausible values (e.g., lower c for `Dark`).

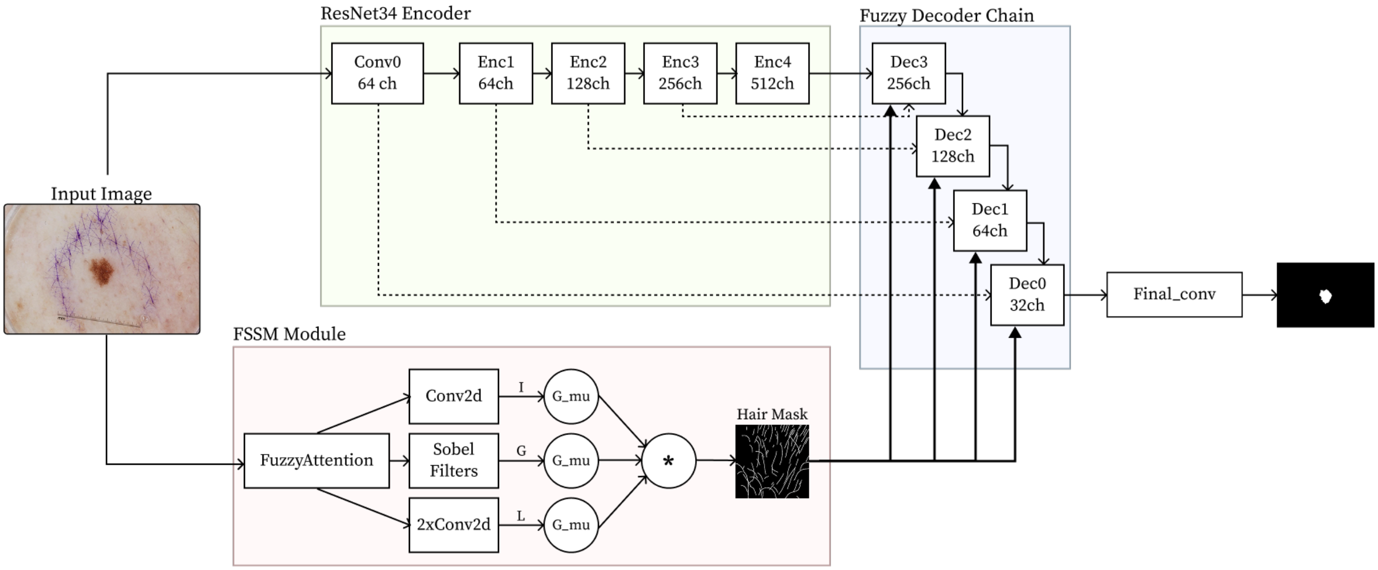

Fig. 1. The proposed FSSM-ResNet34 with 1 rule architecture diagram. The model consists of a ResNet34 Encoder, a parallel Fuzzy Hair Module for mask generation, and a Fuzzy Decoder that integrates both to predict the lesion mask.

d) Single-rule (1-rule) inference.: Our simplest and most interpretable variant uses a single rule:

R_1: IF intensity is Dark AND gradient is Edge-like AND line response is Line-strong THEN hair-membership is High.

We compute the firing strength using a product T-norm:

$$\alpha_{R_1}(p) = \mu_{\text{dark}}(I(p))^{\alpha_I} \cdot \mu_{\text{edge}}(G(p))^{\alpha_G} \cdot \mu_{\text{line}}(L(p))^{\alpha_L}, \quad (2)$$

where $\alpha_I, \alpha_G, \alpha_L$ are scalar weights (we upweight line evidence, since hair is strongly linear). The output $H(x) \in [0,1]^{1\times H\times W}$ is a continuous hair-likelihood map.

e) Multi-rule (MR) extension.: To address broader hair appearance variation (e.g., thick vs. thin, dark vs. light) and explicitly reject confusing artifacts, we implement a multi-rule variant. Unlike the single-rule baseline, this approach aggregates evidence from six distinct fuzzy prototypes: two positive rules to maximize recall and four negative rules to ensure precision by suppressing specific artifacts. The rule base is defined as follows:

- *Positive Rules (Hair Evidence):*

R_{h1}: IF intensity is *Dark* AND line-response is *Line-like* AND width is *Thin* AND orientation is *Coherent* THEN hair-membership is *High*. (Targets standard dark hairs).

R_{h2}: IF line-response is *Line-like* AND width is *Thin* AND orientation is *Coherent* THEN hair-membership is *High*. (Targets light-colored hairs by relaxing the intensity constraint).

- *Negative Rules (Artifact Suppression):*

R_{n1}: IF gradient is *Strong-Edge* THEN hair-membership is *Low*. (Suppresses sharp skin folds).

R_{n2}: IF width is *Wide* THEN hair-membership is *Low*. (Suppresses thick vessels or lesions).

R_{n3}: IF structure is *Blob-like* THEN hair-membership is *Low*. (Suppresses bubbles and globules).

R_{n4}: IF orientation is *Isotropic* THEN hair-membership is *Low*. (Suppresses non-directional noise).

Here, *orientation* refers to the dominant local direction of an image structure. It represents *orientation coherence*, defined as

$$\text{coh} = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \varepsilon}, \quad (3)$$

where λ_1 and λ_2 are the eigenvalues of the local structure tensor. High coherence indicates a strongly directional, anisotropic pattern, as expected for hair-like fibers, whereas low coherence corresponds to *isotropic* structures that exhibit similar properties in all directions and therefore lack a dominant orientation. In addition, *blob-like* structures denote compact, rounded, non-elongated regions. Within the multi-rule fuzzy system, coherent orientation supports positive hair rules, while isotropic and blob-like responses activate negative rules to suppress non-hair artifacts such as bubbles, globules, and non-directional noise.

This configuration increases coverage but reduces simplicity, as interpretability is shared across several rules. In experiments, we therefore report both the 1-rule (maximal interpretability) and multi-rule (maximal coverage) variants.

f) Hair-aware suppression inside the decoder.: Given a decoder feature map F and a resized hair map $H_{\downarrow}$, we suppress hair regions by

$$F \leftarrow F \odot (1 - H_{\downarrow}), \quad (4)$$

so the network learns to rely less on corrupted evidence when forming lesion boundaries.

E. Fuzzy Self-Attention (FSA): efficient fuzzy fusion

F^2CAU-Net [9] presents fuzzy attention as a mechanism to reduce redundancy and improve robustness under uncertainty.

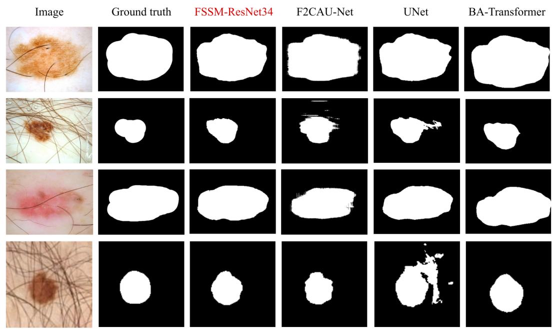

Fig. 2. Qualitative comparison of lesion segmentation predictions on the VHI dataset. Our FSSM-ResNet produces cleaner lesion boundaries and fewer hair-induced artifacts.

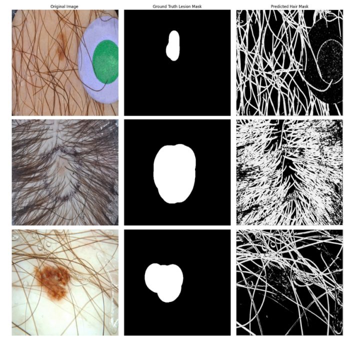

Fig. 3. Qualitative comparison of hair mask predictions of the FSSM on the VHI dataset. FSSM produces clean and accurate hair masks.

We keep the fuzzy motivation but redesign attention to be *fast and spatially native* for dermoscopy.

Given $X \in \mathbb{R}^{C\times H\times W}$, FSA applies: *(1) channel attention* (SE-style) [28], *(2) spatial attention* (CBAM-style) [29], and *(3) fuzzy residual fusion* with a *single learnable scalar* $w \in (0, 1)$.

After channel attention A_c and spatial attention A_s, we obtain the refined features

$$X_s = (X \odot A_c) \odot A_s, \tag{5}$$

which has the same shape as X. We introduce a single learnable scalar $s \in \mathbb{R}$ and convert it to a mixing weight $w = \sigma(s) \in (0, 1)$, then fuse by

$$X' = (1 - w)X + wX_s. \tag{6}$$

This scalar w can be interpreted as a *membership degree* controlling how much to trust the refined features. Compared with heavier attention designs, this is parameter-efficient (+1 scalar per block) and has linear complexity $O(CHW)$, which is important for high-resolution dermoscopy.

IV. Experiments & Results

This section evaluates FSSM-ResNet on public benchmark datasets and the proposed VHI dataset. We first describe the datasets, evaluation metrics, and baseline methods, then present quantitative segmentation results, robustness under heavy hair occlusion, and ablation studies assessing the contribution of each module. We also provide qualitative visual results to illustrate hair suppression behavior and discuss the implications of these findings.

A. Experimental Setup

We evaluate segmentation robustness under hair occlusion by reporting performance on both standard dermoscopy benchmarks and our hair-stress benchmark.

1) Datasets:

- Training data: All segmentation models are trained on HAM10000 (10,015 dermoscopic images from multiple acquisition sources) [30]. We use an 80/20 train/validation split and do *not* use any VHI images for training or model selection.
- VHI (Very Hairy Images): We curate VHI, a custom benchmark of 936 dermoscopic images selected to exhibit *strong hair occlusion* (dense, dark, line-like artifacts crossing the lesion boundary). Images are sourced from multiple public dermoscopy collections, including ISIC 2019 [31], ISIC 2020 [32], ISIC 2024 [33], and the MILK10k benchmark [34]. Every VHI image was manually annotated (or quality-checked when an initial mask existed) to produce a binary lesion mask using CVAT [35]. VHI is used *only* for evaluation (no train/validation/test split, no hyperparameter tuning) and we verified no overlap between VHI and HAM10000.
- Hair mask dataset: The fuzzy hair module (FSSM) is trained separately on a dedicated hair-mask dataset containing 500 dermoscopic images with manually refined hair segmentation masks [36]. This provides direct pixel-level supervision for learning hair-likelihood maps without altering the lesion segmentation labels.
- Additional evaluation benchmarks: To contextualize performance beyond hair-occluded images, we evaluate on PH2 [37] and ISIC2016–2018 [38], [39], [40] lesion segmentation benchmarks. For ISIC2016, we follow the official challenge setting with expert-derived ground truth.

2) Metrics: We report Intersection over Union (IoU) and Dice, which are overlap-based metrics commonly recommended for medical image segmentation evaluation [42]. IoU, also called the Jaccard Index, measures how much the predicted region overlaps with the true region relative to their union. Dice coefficient, also called the F1 score for segmentation, measures overlap with slightly different weighting. IoU penalizes under-/over-segmentation more strongly, while Dice is more forgiving when boundaries are uncertain (typical in dermoscopy under occlusion).

3) Training Details: All models were trained for 100 epochs with 320×320 input shape and a Binary Cross-Entropy (BCE) with Dice loss function [9]. The FSSM is either (i) frozen after standalone training or (ii) made trainable end-to-end (*tunable*) to test whether joint optimization improves hair robustness.

TABLE I
QUANTITATIVE COMPARISON OF SEGMENTATION MODELS ACROSS FIVE BENCHMARK DATASETS. THE BEST-PERFORMING RESULT FOR EACH METRIC WITHIN A DATASET IS SHOWN IN **BOLD**, WHILE THE SECOND-BEST RESULT IS INDICATED BY <u>UNDERLINING</u>.

Models	VHI		PH2		ISIC 2016		ISIC 2017		ISIC 2018	
	IoU↑	Dice↑	IoU↑	Dice↑	IoU↑	Dice↑	IoU↑	Dice↑	IoU↑	Dice↑
U-Net [25]	0.4838	0.6188	0.7551	0.8458	0.7672	0.8462	0.6007	0.7098	0.6883	0.7966
BA-Transformer (ResNet-50 backbone) [41]	**0.5961**	**0.6763**	0.8272	0.8862	0.8175	0.8617	0.6669	0.7321	0.7781	0.8434
F^2CAU-Net (baseline) [9]	0.5159	0.6406	0.8328	0.9011	0.8008	0.8678	0.6217	0.7260	0.7636	0.8518
Ours FSSM-ResNet50 (FSSM frozen, multi-rule)	0.4248	0.5712	**0.8730**	**0.9275**	0.8605	0.9204	**0.7391**	**0.8349**	**0.8072**	**0.8849**
Ours FSSM-ResNet50 (FSSM frozen, 1-rule)	0.5055	0.6483	0.8611	0.9214	<u>0.8621</u>	<u>0.9210</u>	0.7250	0.8248	0.8027	0.8822
Ours FSSM-ResNet50 (no FSSM)	0.4775	0.6238	<u>0.8663</u>	<u>0.9253</u>	**0.8629**	**0.9218**	<u>0.7317</u>	<u>0.8308</u>	<u>0.8031</u>	<u>0.8835</u>
Ours FSSM-ResNet34 (FSSM frozen, multi-rule)	0.5051	0.6310	0.8377	0.9049	0.8023	0.8679	0.6393	0.7462	0.7639	0.8547
Ours FSSM-ResNet34 (FSSM frozen, 1-rule)	<u>0.5354</u>	<u>0.6647</u>	0.8470	0.9083	0.8072	0.8724	0.6737	0.7697	0.7800	0.8637
Ours FSSM-ResNet34 (no FSSM)	0.5159	0.6406	0.8328	0.9011	0.8008	0.8678	0.6217	0.7260	0.7636	0.8518
Ours FSSM-ResNet34 (tunable, 1-rule)	0.5128	0.6372	0.8436	0.9077	0.8219	0.8846	0.6507	0.7537	0.7708	0.8578

TABLE II
COMPARISON OF THE MODELS IN TERMS OF PARAMETER COUNT AND INFERENCE THROUGHPUT.

Model	**Parameters**	**Throughput (samples/s)**
UNet	7.85M	363.57
F2CAUNet	8.55M	155.00
BA-Transformers	46.74M	126.22
Ours: FSSM-ResNet34	24.48M	345.00
Ours: FSSM-ResNet50	44.00M	200.00

B. *Segmentation Results and Analysis*

Table I reports the segmentation results. Overall, hair robustness and clean-image performance do not necessarily correlate.

FSSM-ResNet50 (frozen, multi-rule) achieves the highest IoU and Dice scores on PH2 and ISIC2017–2018, making it the best-performing model on these benchmarks. In contrast, FSSM-ResNet50 (no FSSM) attains the top result on ISIC 2016, indicating that the effect of the proposed hair suppression mechanism may vary across datasets. This suggests that dataset-specific image characteristics influence the extent to which explicit hair-aware processing is beneficial. When combined with a sufficiently expressive encoder, the proposed FSSM module appears to act as a regularizing filter, suppressing dark structures unrelated to the lesion and improving boundary consistency on typical dermoscopic images.

A different trend is observed on VHI, which serves as a challenging stress test for severe hair occlusion. On this dataset, BA-Transformer achieves the best segmentation performance, while FSSM-ResNet34 ranks second. This result is reasonable given the substantially higher representational capacity of BA-Transformer as shown in Table II, which uses 46.74M parameters and exhibits the lowest inference throughput among the compared models (126.22 samples/s). The throughput was measured on NVIDIA RTX 4090. Its stronger performance on VHI therefore comes at a considerable computational cost. By contrast, FSSM-ResNet34 attains competitive segmentation accuracy while operating with nearly half the parameter count (24.48M) and much higher throughput (345 samples/s). This indicates that the best scores of BA-Transformer comes at the cost of efficiency, while the proposed FSSM-ResNet34 provides a substantially better efficiency–performance balance, making them more attractive for real-world clinical systems where both segmentation quality and inference speed are critical.

The multi-rule and ResNet50 configurations degrade on VHI benchmark. This indicates a practical trade-off: (i) stronger backbones may overfit to non-hair cues prevalent in standard datasets, and (ii) multi-rule fuzzy detectors can become *overly aggressive* under severe occlusion, suppressing not only hair but also genuine dark lesion regions, thereby reducing IoU. In other words, the hair-occlusion stress test favors simpler fuzzy hair logic.

The *tunable* variant does not surpass the frozen 1-rule model on VHI. A plausible explanation is that segmentation loss alone does not explicitly reward accurate hair localization. In the absence of a dedicated hair loss, gradients can drive the hair mask toward shortcut behaviors (e.g., suppressing any dark region), which in turn degrades the boundaries of the lesion in heavily occluded images.

When the hair module is disabled, FSSM-ResNet50 (no FSSM) has lower scores than FSSM-ResNet50 (FSSM frozen, 1-rule). This emphasizes the importance of FSSM for hairy images. Furthermore, FSSM-ResNet34 (no FSSM) matches the baseline F^2CAU-Net, since the resulting architecture is essentially the similar decoder-attention design without hair gating. This highlights that performance gains are specifically attributable to the FSSM component.

This interpretation is further supported by the qualitative results shown in Fig. 3, which illustrate the hair segmentation outputs produced by the proposed FSSM module. It can be seen that FSSM generates clean and well-localized masks that are subsequently used within the decoder to suppress occluding structures. Importantly, the module also responds to bubble artifacts commonly introduced during dermoscopic imaging. Since air bubbles are also recognized as a source of interference in dermoscopic image analysis, their suppression may provide an additional benefit to the segmentation of the lesion [43]. This behavior offers a plausible explanation for why the proposed models remain competitive not only on heavily occluded images, but also on standard public benchmarks, where such artifacts can also degrade segmentation quality.

Taken together, the results show that integrating explicit hair reasoning improves robustness. As illustrated in Fig. 2, FSSM-ResNet produces cleaner and more complete lesion masks under hair occlusion, with fewer boundary breaks and reduced leakage into hair structures than competing methods. Quantitatively, the multi-rule, large-backbone model performs best on standard benchmarks, whereas the compact, single-rule variant generalizes better under extreme hair occlusion, revealing an accuracy–robustness trade-off across operating conditions. At the same time, the throughput analysis shows that these gains are not achieved at the expense of excessive computational cost, further supporting the practicality of the proposed approach.

C. Ablation Study

To assess the contribution of the main components, we conducted an ablation study by varying the backbone architecture, the complexity of the fuzzy rules, and the training configuration of FSSM. The quantitative results are reported in Table I.

1) Impact of Backbone Architecture: We compared ResNet-34 and ResNet-50 as backbone feature extractors. Overall, ResNet-50 performs better on most datasets. Under the "FSSM frozen, 1-rule" setting, it achieves higher IoU on PH2 (0.8611 vs. 0.8470), ISIC 2016 (0.8621 vs. 0.8072), and ISIC 2017 (0.7250 vs. 0.6737). Although ResNet-34 remains competitive on VHI, the deeper ResNet-50 generally provides stronger feature representations for lesion segmentation.

2) Single Rule vs. Multi-Rule: We also compared a simple 1-rule fuzzy configuration with a more complex multi-rule variant. The 1-rule version uses a compact fuzzy expression based on intensity, gradient, and line responses, while the multi-rule version uses six rules, including both positive hair-detection rules and negative rules to suppress non-hair artifacts.

As shown in Table I, *FSSM-ResNet50 (FSSM frozen, multi-rule)* improves performance on PH2 and the ISIC datasets compared to *FSSM-ResNet50 (no FSSM)*, but performs worse on VHI. This suggests that the effectiveness of the multi-rule design depends on dataset characteristics such as artifact type and lesion appearance. In contrast, the 1-rule variant shows more consistent performance across both ResNet-34 and ResNet-50, indicating that a simpler fuzzy design may generalize better. The stronger VHI results of ResNet-34 with 1-rule further suggest that performance is influenced not only by backbone depth, but also by dataset-specific characteristics and backbone–module interactions.

3) FSSM Configuration:: Finally, we examined the training strategy for the FSSM module by comparing three configurations: "No FSSM" (baseline), "Tunable" (learnable fuzzy parameters), and "Frozen" (fixed fuzzy parameters). Introducing FSSM improves performance over the baseline. However, the FSSM "Frozen" setting consistently outperforms the tunable variant. For instance, on the ISIC 2016 dataset, the frozen 1-rule model attains an IoU of 0.8621, whereas the tunable 1-rule model reaches only 0.8219. This indicates that fixing the fuzzy parameters helps avoid overfitting and allows the network to concentrate on learning better feature representations, thereby improving generalization.

V. Conclusion

In this paper, we presented FSSM-ResNet, a hair-aware skin lesion segmentation architecture designed to address hair occlusion in dermoscopic images. By integrating the Fuzzy Structural Suppression Module (FSSM) directly into a ResNet–U-Net backbone, the proposed approach removes the need for separate hair-removal or inpainting stages. The Mamdani-style fuzzy inference module encodes simple dermatological priors by modeling hair as dark, thin, and linear structures, producing an interpretable hair-likelihood map that suppresses hair-corrupted regions during decoding. To rigorously evaluate hair robustness, we also introduced the VHI (Very Hairy Images) benchmark, a manually annotated dataset focused on severe hair artifacts. Experiments on PH2, ISIC 2016–2018, and VHI show that FSSM-ResNet improves segmentation quality on standard benchmarks and provides stronger robustness under heavy hair occlusion, while revealing a trade-off between backbone capacity, fuzzy rule complexity, and robustness. Overall, the results show that simple and interpretable fuzzy modeling can effectively complement deep learning for artifact-aware segmentation. A limitation of the proposed approach is the possible over-suppression of dark lesion regions under severe occlusion.

As future work, we plan to extend the proposed segmentation framework to broader computer-aided diagnosis tasks, including lesion classification, subtype characterization, and automated extraction of ABCD criteria and shape-based biomarkers. We also aim to explore multi-task learning for joint segmentation, attribute detection, and risk scoring, potentially with clinical metadata. In addition, we intend to adapt FSSM-ResNet to large-scale total-body imaging and real-time dermoscopy workflows to further assess its clinical utility.

Acknowledgment

Funding for this research was provided by the Ministry of Science and Higher Education of the Republic of Kazakhstan under the grant titled "Smart-Care: Innovative Multi-Sensor Technology for Elderly and Disabled Health Management" (AP23487613) for the 2024-2026 academic years.

References

[1] R. Wang, Y. Chen, X. Shao, T. Chen, J. Zhong, Y. Ou, and J. Chen, "Burden of skin cancer in older adults from 1990 to 2021 and modelled projection to 2050," *JAMA Dermatology*, vol. 161, pp. 715–722, 07 2025.

[2] M. Arnold, D. Singh, M. Laversanne, J. Vignat, S. Vaccarella, F. Meheus, A. E. Cust, E. de Vries, D. C. Whiteman, and F. Bray, “Global burden of cutaneous melanoma in 2020 and projections to 2040,” *JAMA Dermatol.*, vol. 158, pp. 495–503, May 2022.

[3] J. Y. Chen, K. Fernandez, R. P. Fadadu, R. Reddy, M.-O. Kim, J. Tan, and M. L. Wei, “Skin cancer diagnosis by lesion, physician, and examination type: A systematic review and meta-analysis,” *JAMA Dermatology*, vol. 161, pp. 135–146, 08 2024.

[4] M. Perez, J. A. Abisaad, K. D. Rojas, M. A. Marchetti, and N. Jaimes, “Skin cancer: Primary, secondary, and tertiary prevention. part I,” *J. Am. Acad. Dermatol.*, vol. 87, pp. 255–268, Aug. 2022.

[5] A. Yazici, D. Zhumabekova, A. Nurakhmetova, Z. Yergaliyev, H. Y. Yatbaz, Z. Makisheva, M. Lewis, and E. Ever, “A smart e-health framework for monitoring the health of the elderly and disabled,” *Internet of Things*, vol. 24, p. 100971, 2023.

[6] T. Lee, V. Ng, R. Gallagher, A. Coldman, and D. McLean, “Dullrazor: a software approach to hair removal from images,” *Computers in Biology and Medicine*, 1997. PubMed: 9437554.

[7] R. Kasmi *et al.*, “Sharprazor: Automatic removal of hair and ruler marks from dermoscopy images,” *Skin Research and Technology*, 2023. PMCID: PMC10234178.

[8] L. Talavera-Martínez, P. Bibiloni, and M. Gonzalez-Hidalgo, “Hair removal in dermoscopic images using deep multitask learning,” in *CVPR Workshops*, 2022.

[9] T. Zhou, H. Wang, S. Geng, H. Ju, J. Huang, F. Fu, and W. Ding, “F^2CAU-Net: A dual fuzzy medical image segmentation cascade method based on fuzzy feature learning,” *Applied Soft Computing*, vol. 184, p. 113692, 2025.

[10] S. Chatterjee, D. Dey, and S. Munshi, *Preprocessing and segmentation of skin lesion images*, pp. 25–52. 01 2022.

[11] P. Choudhary, S. Jyoti, and J. Yadav, “Curvelet and fast marching method-based technique for efficient artifact detection and removal in dermoscopic images,” *International Journal of Imaging Systems and Technology*, vol. 31, 07 2021.

[12] X. Song, S. Guo, L. Han, L. Wang, W. Yang, G. Wang, and A. B. Çekderi, “Research on hair removal algorithm of dermatoscopic images based on maximum variance fuzzy clustering and optimization criminisi algorithm,” *Biomedical Signal Processing and Control*, vol. 78, p. 103967, 09 2022.

[13] N. Lama, R. Kasmi, J. R. Hagerty, R. J. Stanley, R. Young, J. Miinch, J. Nepal, A. Nambisan, and W. V. Stoecker, “ChimeraNet: U-Net for hair detection in dermoscopic skin lesion images,” *Journal of Digital Imaging*, vol. 36, pp. 526–535, Apr. 2023.

[14] W. Li, A. N. Joseph Raj, T. Tjahjadi, and Z. Zhuang, “Digital hair removal by deep learning for skin lesion segmentation,” *Pattern Recognition*, vol. 117, p. 107994, Sept. 2021.

[15] N. Bansal and S. Sridhar, “Hexa-gan: Skin lesion image inpainting via hexagonal sampling based generative adversarial network,” *Biomedical Signal Processing and Control*, vol. 89, p. 105603, 03 2024.

[16] D. Bardou, L. Lv, Y. Medjadba, T. Zhang, O. Chouhal, M. Bounezra, S. Karima, Y. Bezza, and A. Snani, “Gad-vae: generative adversarial disentanglement with variational autoencoders for hair removal in dermoscopy images,” *Network Modeling Analysis in Health Informatics and Bioinformatics*, vol. 13, 06 2024.

[17] B. Wang, J. Peng, H. Li, and H. Wang, “Generating high-quality texture via panoramic feature aggregation for large mask inpainting,” *Knowledge-Based Systems*, vol. 286, p. 111382, Feb. 2024.

[18] U. Sharma, P. Aggarwal, and A. Mittal, “Computer-Aided classification of melanoma: A comprehensive survey,” *Archives of Computational Methods in Engineering*, vol. 31, pp. 4893–4927, Dec. 2024.

[19] Z. Huang, H. Deng, S. Yin, T. Zhang, W. Tang, and Q. Wang, “Adf-net: A novel adaptive dual-stream encoding and focal attention decoding network for skin lesion segmentation,” *Biomedical Signal Processing and Control*, vol. 91, p. 105895, 2024.

[20] R. Yang, J. Yu, J. Yin, K. Liu, and S. Xu, “An fa-segnet image segmentation model based on fuzzy attention and its application in cardiac mri segmentation,” *International Journal of Computational Intelligence Systems*, vol. 15, p. 24, 04 2022.

[21] K. B. Patel, F. Li, and G. Wang, “Fuzzynet: A fuzzy attention module for polyp segmentation,” in *NeurIPS ’22 Workshop on All Things Attention: Bridging Different Perspectives on Attention*, 2022.

[22] R. Liu, S. Duan, L. Xu, L. Liu, J. Li, and Y. Zou, “A fuzzy transformer fusion network (fuzzytransnet) for medical image segmentation: The case of rectal polyps and skin lesions,” *Applied Sciences*, vol. 13, no. 16, 2023.

[23] C. Tao, H. Chen, R. Wu, H. Zhi, X. Yan, H. Liu, C. Xu, and M. Jian, “Fcau-net: A frequency channel attention convolutional neural network for medical image segmentation,” in *2023 IEEE Smart World Congress (SWC)*, pp. 1–6, 2023.

[24] J. Qin, D. Pei, Q. Guo, X. Cai, L. Xie, and W. Zhang, “Intersection-union dual-stream cross-attention lova-swinunet for skin cancer hair segmentation and image repair,” *Computers in Biology and Medicine*, vol. 180, p. 108931, 2024.

[25] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015* (N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, eds.), (Cham), pp. 234–241, Springer International Publishing, 2015.

[26] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016.

[27] E. Mamdani and S. Assilian, “An experiment in linguistic synthesis with a fuzzy logic controller,” *International Journal of Man-Machine Studies*, vol. 7, no. 1, pp. 1–13, 1975.

[28] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7132–7141, 2018.

[29] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, “CBAM: Convolutional block attention module,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018.

[30] P. Tschandl, C. Rosendahl, and H. Kittler, “The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions,” *Scientific Data*, vol. 5, p. 180161, 2018.

[31] International Skin Imaging Collaboration (ISIC), “Isic 2019 challenge.” Online, 2019. Accessed 2026-01-31.

[32] V. Rotemberg, N. Kurtansky, B. Betz-Stablein, L. Caffery, E. Chousakos, N. Codella, M. Combalia, S. Dusza, *et al.*, “A patient-centric dataset of images and metadata for identifying melanomas using clinical context,” *Scientific Data*, vol. 8, p. 34, 2021.

[33] N. R. Kurtansky *et al.*, “The SLICE-3D dataset: 400,000 skin lesion image crops extracted from 3d total body photographs for skin cancer detection,” *Scientific Data*, 2024.

[34] International Skin Imaging Collaboration (ISIC), “Milk10k.” Online dataset descriptor (ISIC Archive DOI landing page), 2025. Accessed 2026-01-31.

[35] CVAT.ai Corporation, “Computer vision annotation tool (cvat).” Software (Zenodo), 2024. Accessed 2026-01-31.

[36] S. I. Hossain, S. S. Roy, J. De Goër De Herve, R. E. Mercer, and E. M. Nguifo, “A skin lesion hair mask dataset with fine-grained annotations.” Mendeley Data, Version 2, 2023.

[37] T. Mendonca, P. M. Ferreira, J. S. Marques, A. R. S. Marcal, and J. Rozeira, “PH^2 - a dermoscopic image database for research and benchmarking,” in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, IEEE, July 2013.

[38] D. Gutman, N. C. F. Codella, M. E. Celebi, B. Helba, M. Marchetti, N. Mishra, and A. Halpern, “Skin lesion analysis toward melanoma detection: A challenge at the international symposium on biomedical imaging (isbi) 2016, hosted by the international skin imaging collaboration (isic),” 2016.

[39] M. Berseth, “Isic 2017 - skin lesion analysis towards melanoma detection,” 2017.

[40] N. Codella, V. Rotemberg, P. Tschandl, M. E. Celebi, S. Dusza, D. Gutman, B. Helba, A. Kalloo, K. Liopyris, M. Marchetti, H. Kittler, and A. Halpern, “Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the international skin imaging collaboration (isic),” 2019.

[41] J. Wang, L. Wei, L. Wang, Q. Zhou, L. Zhu, and J. Qin, “Boundary-aware transformers for skin lesion segmentation,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 206–216, Springer, 2021.

[42] D. Müller, I. Soto-Rey, and F. Kramer, “Towards a guideline for evaluation metrics in medical image segmentation,” 2022.

[43] S. W. I. H. O. H. C. X. Emre Celebi M, Alp Aslandogan Y, “Unsupervised border detection in dermoscopy images.,” *Skin Research and Technology*, vol. 13, no. 4, pp. 454–462, 2007.