

# Uncertainty Quantified Deep Learning and Regression Analysis Framework for Image Segmentation of Skin Cancer Lesions

1<sup>st</sup> Elhoucine Elfatimi

*Department of Pathology Laboratory Medicine  
University of California  
Irvine, CA, USA  
eelfatim@uci.edu*

2<sup>nd</sup> Pratik Shah\*

*Department of Pathology Laboratory Medicine  
Biomedical Engineering  
University of California  
Irvine, CA, USA  
pratik.shah@uci.edu*

**Abstract**—Deep learning models (DLMs) frequently achieve accurate segmentation and classification of tumors from medical images. However, DLMs lacking feedback on their image segmentation mechanisms such as Dice coefficients and confidence in their performance face challenges processing previously unseen images in real-world clinical settings. Uncertainty estimates to identify DLM predictions at the cellular or single-pixel levels requiring clinician review can enhance trust, however their deployment requires significant computational resources. This study reports two DLMs, one trained from scratch and another based on transfer learning, with Monte Carlo dropout or Bayes-by-backprop uncertainty estimations to segment lesions from the publicly available The International Skin Imaging Collaboration-19 dermoscopy image database with lesions of cancer. A novel approach to compute pixel-by-pixel uncertainty estimations of DLM segmentation performance in multiple clinical regions from a single dermatoscopy image with corresponding Dice scores is reported for the first time. Image-level uncertainty maps demonstrated correspondence between imperfect DLM segmentation and high uncertainty levels in specific skin tissue regions with or without lesions. Four new linear regression models that can predict the Dice performance of DLM segmentation using constants and uncertainty measures either individually or in combination from lesions, tissue structures, and non-tissue pixels regions critical for clinical diagnosis and prognostication in skin images (Spearman's correlation,  $p < 0.05$ ), are reported for the first time for low-compute uncertainty estimation workflows.

**Index Terms**—Deep learning, Skin cancer, Uncertainty prediction, Segmentation, Statistical modeling, Image processing

## I. INTRODUCTION

Skin cancer is the most common cancer in the United States, affecting about one in five Americans in their lifetime [1]. Differentiating between melanoma, melanocytic nevus (mole), and seborrheic keratosis can be clinically challenging. Seborrheic keratosis appears as discolored, raised patches and is unrelated to sun exposure, while melanoma, an aggressive skin cancer, is caused by UV light damaging skin cell DNA [2]. Melanocytic nevi are benign skin lesions from

melanocytes, generally harmless, making visual inspection for skin cancer difficult. Deep neural networks (DNNs) are employed for image detection, classification, and segmentation in clinical applications. Research has focused on using deep learning algorithms to classify and segment dermoscopic images of skin lesions [3]. However, DNNs often lack mechanisms for risk evaluation and comprehensive uncertainty estimations, which are essential for assessing the reliability and confidence of model predictions [4]. Many DNNs use softmax activation, which can lead to unjustified confidence levels for unknown inputs [5]. Bayesian neural networks that provide output and uncertainty by setting a distribution over model weights, are computationally complex and impractical for high-dimensional inputs [6]. Approximating Bayesian distributions using dropout-trained networks can address these challenges [5].

One key goal of this study was to address clinically relevant challenges of uncertainty estimation in DNNs by employing Bayesian uncertainty methods-inspired approaches of Monte Carlo dropout (MCD) and Bayes-by-backprop (BBP) for segmentation of melanoma, melanocytic nevus, and seborrheic keratosis skin conditions from the The International Skin Imaging Collaboration (ISIC) 2019 challenge images [7]. The most prevalent DNNs models do not to quantify prediction uncertainty at the pixel level, which limits their clinical applicability where prediction confidence is crucial. This study integrates MCD and BBP to provide pixel-level uncertainty estimates, thereby enhancing both the trustworthiness and interpretability of predictions. Another goal was approximating Bayesian probability distributions in a more computationally feasible manner, making them suitable for high-dimensional dermoscopy images with multiple classes. Additionally, deep learning models (DLMs) trained from scratch (TFS) and based on transfer learning (TL) were evaluated to compare their performance with MCD and BBP techniques. A novel fast computing framework that integrates multiple clinical skin region-based pixel-by-pixel uncertainty from a single image

\*Corresponding author: Dr. Pratik Shah Ph.D. (pratik.shah@uci.edu)

to predict DLMs output uncertainties using linear regression models is reported for the first time to visualize, interpret and improve DNN predictions for clinical applications.

### A. Related Work

The relationship between Dice scores and uncertainty in deep learning models (DLMs) for segmenting clinically important regions in skin and medical images is not fully understood, and face challenges such as lack of voxel-level uncertainty feedback and methods for individualized clinical outcomes. A single Dice score is commonly used in organ segmentation studies. For instance, Senapati et al. evaluated liver segmentation using Bayesian neural networks with Dice scores [8], Hoebel et al. used linear regression with uncertainty metrics for thorax segmentation [9], and Do et al. demonstrated a correlation between Dice scores and uncertainty in myocardial imaging with a UNet-like architecture [10]. Most studies on dermoscopic images use DLMs for classification without regression analysis or Dice score evaluations. For example, uncertainty distributions from MCD have been reported without Dice scores [11]. Similarly, Combalia et al. used Monte Carlo sampling without computing Dice scores [12], and Shapna employed a CNN model without pixel-level analysis [13]. In skin lesion segmentation, studies like Ashraf et al. used test-time augmentation and Bayesian learning to evaluate Dice coefficients but did not include linear regression [14]. Similarly, De Angelo et al. used Dice and Jaccard index for multiple regions of interests (ROIs) without linear regression [15]. Another study reported improved U-Net architectures for skin segmentation on the ISIC dataset, but did not focus on uncertainty estimation using training from scratch (TFS) or transfer learning (TL) with MCD and BBP approaches used in our paper [16]. There is a need for a framework that evaluates the relationship between Dice scores, uncertainty, and the performance of DLMs trained with TFS or TL using MCD and BBP, particularly for clinical diagnosis of various skin lesions.

### B. Summary of Contributions

This study focuses on training multiple DNNs using MCD and BBP techniques with TFS and TL approaches for segmentation of melanoma, melanocytic nevus, and seborrheic keratosis skin conditions from the ISIC-19 images and estimate DLM uncertainty at the pixel level. Another key goal was to enhance the segmentation accuracy and reliability of binary segmentation tasks, particularly in identifying skin or cancer lesions, from multiple clinically important ROI's from dermoscopic Red Green Blue (RGB) images. The performance of these approaches was assessed using five-fold cross-validation and other comprehensive statistical metrics reported in the methods section and summarized in Tables I and II. By leveraging ground-truth skin lesion labels, all images were categorized into two clinical ROIs: pixels with lesions and pixels without skin lesions. Additionally, three individual clinical region-based uncertainties and an overall uncertainty measure for each image were computed specific to each of the three skin lesion

diagnosis classes. A novel algorithm (Algorithm 1) is proposed to calculate region-specific uncertainty, facilitating a more granular evaluation of DLM performance. Four separate linear models were trained and validated to learn the correspondence between the Dice coefficient, calculated by comparing DLM output with ground truth segmentation masks for each image, and the clinical region-specific uncertainties derived from the MCD and BBP approaches (Algorithm 1). Statistical analysis revealed significant correlations ( $p < 0.05$ ) between Dice scores and the predictive performance of the MCD and BBP DLMs. The linear models reported in this study demonstrated high-accuracy in predicting the Dice performance of DLMs using constants and uncertainty measures from two clinical ROIs, either individually or in combination (Equations 3 and 4).

## II. DATA DESCRIPTION

The International Skin Imaging foundation released a data set in 2019 (ISIC 2019) containing 13,786 3-channel RGB images of resolution ranging from  $(7000 \times 4500)$  to  $(1024 \times 720)$  ( $W \times H$ ) with available ground truth (GT) binary mask associated with each image as its label [17]. Each binary mask had white pixels denoting regions of the image where skin moles and or cancer labels were present (class 1). Black pixels in the binary masks represented non-target and or non-disease regions (class 0). Each image also has a clinical class (from a set of 17 classes) associated with it based on the dermoscopy diagnosis. The dataset used comprised 2,935 images with the top three available clinical classes: nevus ( $n=1497$ ), melanoma ( $n=1019$ ), and seborrheic keratosis ( $n=419$ ). These classes are highly prevalent in clinical practice, making them a good fit for research and diagnostic algorithms. This dataset has been widely used to identify the most common types of skin cancer from dermoscopic images [18]. Detailed description on data pre-processing and training of DLMS with five-fold cross validation were performed as previously described [?].

## III. METHODS

### A. Deep Learning Model Training

The ISIC 2019 dataset was split into 80% for training (~2,347 images) and 20% for testing (~588 images). A Bayesian VGG-UNet CNN, enhanced with MCD and BBP techniques, was trained on skin lesion images and their labels. The network used an encoder-decoder structure [19], based on a modified VGG-16 [20]. To address uncertainties in segmentation, uncertainty quantification mechanisms were integrated into the decoder component to mitigate noise during the downscaling phase. The MCD and BBP approximated Bayesian inference, minimizing the approximate posterior within a deep Gaussian process framework [5]. All DLMs were trained for 50 epochs with a batch size of two using the Adam optimizer and a learning rate of ( $\alpha_{lr} = 0.001$ ). Input images were resized to  $608 \times 416$  (width  $\times$  height), with outputs at  $304 \times 208$ . Training was done on an NVIDIA GeForce RTX 4080 Ti with 32GB of GPU memory.

The DLM training, implementation, and image processing workflow, shown in Fig.1, involved training using both TL and TFS. Outputs were generated from DLMs with Monte Carlo Dropout (DLM MCD) and Bayes-by-Backprop (DLM Bc). The output predictions from DLM MCD and DLM Bc were combined and compared with ground truth segmentation masks, and Dice scores were computed. Correlation analysis was used to establish the relationship between Dice scores and uncertainty estimations derived from the DLM by developing separate linear models and algorithms for predictions that incorporated uncertainties. This framework integrates the segmentation performance of multiple DLMs with specific regions of clinical relevance in individual skin images to calculate the overall segmentation performance of the deep learning model.

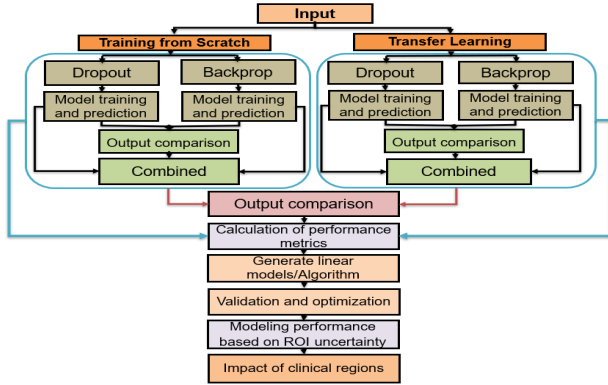


Fig. 1. Schematic diagram of the deep learning framework for interpretable skin lesion image segmentation with uncertainty quantification and linear regression

### B. Deep Learning Model Prediction And Evaluation

For each model prediction, the numbers of True Positive (TP), False Positive (FP), False Negative (FN), and True Negative (TN) pixels were counted over the entire output. A TP instance was when an image lesion pixel was segmented accurately, and TN when a non-lesion or non-tissue pixel indicated in ground truth clinical labels was not segmented by the deep learning models. False negative pixels were those that were denoted as lesion pixels in the ground truth, but were not segmented by the model, and FP pixels were those segmented as a lesion by the model but classified as non-lesion tissue or non-tissue in the ground truth data. Using these definitions, four metrics for model predictions for all the images in the test data - the True Positive Rate (TPR), False Positive Rate (FPR) ( Table I), and the AUROC and Dice scores that evaluated the similarity between the ground truth clinical labels and the DLM generated segmentation were calculated (TableI). The Dice scores were used for subsequent linear model estimations. The overall median Dice, AUROC, TPR and FPR values for all 2,935 test images are reported in Table I. A comprehensive comparative analysis of the uncertainty from TFS and TL backprop vs. dropout DLM was learned with linear regression for predicting the Dice coefficient of skin lesion and melanoma segmentation. For each image, a distribution of pixel-level

uncertainty was obtained. The overall mean uncertainty and other uncertainty metrics for all 2,935 test images are reported in Table II. Uncertainty estimates resulting from this evaluation range between a minimum value of 0 for the model being fully certain of its prediction to a maximum of 1 when the model is fully uncertain about its prediction. The analysis was expanded to explore the combined effect of MCD and Bc approaches within the TFS and TL frameworks (Fig. 2). By integrating these techniques, we aimed to harness the synergistic benefits they offer in mitigating uncertainty. To obtain pixel-level uncertainty, 50 MC iterations were taken where the difference between 67<sup>th</sup> and 33<sup>rd</sup> percentiles of their sigmoid values from the DLMs predictions were set as uncertainty estimates and represented as heatmaps (see Fig. 3 and Fig. 4). Two hundred and ninety-four images containing lesion labels from the test dataset constituted the nevus class subset, 224 images from the test data containing lesion label constituted the melanoma class subset, and 87 images that contained seborrheic keratosis lesion labels from the test data set were used. Mathematically, for a given DL MCD model  $M$  and input test image  $X_{test_i}$ , the DLM prediction was calculated by:  $O_{M_p}^j = \text{argmax}(M_p(X_{test_i}))$  for the  $j^{\text{th}}$  MC iteration, where  $M_p$  is the model  $M$  in testing phase. The final model output  $O_{M_i}$  was generated by taking a mean of all 50 MC iteration outputs. To generate the predicted segmentation outputs, a value based thresholding was used to designate the predicted class to which the pixel belonged. Predicted ‘class 0’ (pixel value of ‘0’ in binary scale) or non-lesion pixels were assigned a value of 0 only if the mean value for the corresponding pixel across all MC iterations was  $\leq 0.95$ , and predicted ‘class 1’ (pixel value of ‘1’ in binary scale) or lesion pixels as 1 were assigned only if the mean value from the MC iterations was  $> 0.95$ . Thus,

$$O_{M_{p1}} = 1/\alpha \sum_{j=1}^{\alpha} O_{M_p}^j \quad (1)$$

where  $\alpha = 50$  is the number of MC iterations. And the final model output,  $O_{M_i} = [O_{M_{p1}} \text{ s.t } \forall \text{ pixels, } p \text{ in } O_{M_{p1}}, (|p| > 0.95) \implies (|p| = 1) \text{ and } (|p| \leq 0.95) \implies (|p| = 0)]$ . To obtain the corresponding uncertainty map,  $Unc_{map_i}$ , each MC iteration step produced an uncertainty map as  $Unc_{M_p}^j = \max(M_p(X_{test_i}))$ . Thus,

$$Unc_{map_i} = P_{67}([Unc_{M_p}^j]_{j=1}^{\alpha}) - P_{33}([Unc_{M_p}^j]_{j=1}^{\alpha}) \quad (2)$$

### C. Modelling Performance Based on Uncertainty

Algorithm 1 is proposed to calculate the estimated Dice performance of DLMs using uncertainties for clinical ROIs in skin images used in this study. Equation 3 calculated the Dice of the segmentation model based on a linear combination of lesion-tissue and non-lesion skin pixels uncertainties, and equations 4(i) and 4(ii) calculated the Dice of the DLM segmentation from the pixels of three individual ROIs. Two subsets of test data were used for generating the linear regression models and for obtaining the values of the coefficients for equations 3 to 5. A detailed breakdown of the coefficients and associated

---

**Step 1: DLM Predictions and Uncertainty Mapping**

**Input:** Optimized deep learning model  $M$ , test dataset  $X_{test}$  with  $n$  images, and corresponding labels  $X_{GT}$ .

**Output:** DLM segmentation output,  $O_{M_i}$  and uncertainty map,  $Unc_{map_i}$  (see methods Section III-B) corresponding to  $X_{test_i}$ . Thus,  $\forall i = [1, 2, \dots, n]$  images,

$$X_{Dice_i} = Dice(O_{M_i}, X_{GT_i}), List_{Dice} = [X_{Dice_i}]_{i=1}^n = \\ List_{Unc} = [Unc_{map_i}]_{i=1}^n$$

**Step 2: Generate ROI Uncertainty:**

**Input:** Test image  $X_{test_i}$  and ground-truth mask  $X_{GT_i}$ .

•  $X_{test_i}$ -lesion =  $X_{GT_i}$ , where '0' represents non-lesion pixels, '1' represents lesion pixels.

•  $X_{test_i}$ -non-lesion =  $X_{test_i}$ -lesion<sup>C</sup>, complement of  $X_{test_i}$ -lesion, where '0' represents lesion pixels, '1' represents non-lesion pixels.

**Output:** Uncertainty maps for specific regions computed as:

•  $Unc_{iL} = Unc_{C_{test_i}$ -lesion =  $(X_{test_i}$ -lesion  $\odot$   $Unc_{map_i}$ );  $List_{T_{unc}} = [mean(Unc_{iL})]_{i=1}^n$ .

•  $Unc_{iNL} = Unc_{C_{test_i}$ -non-lesion =  $(X_{test_i}$ -non-lesion  $\odot$   $Unc_{map_i}$ );  $List_{NT_{unc}} = [mean(Unc_{iNL})]_{i=1}^n$ . The symbol  $\odot$  represents the Hadamard product.

**Step 3: Generate linear regression models for Dice Score Prediction:**

**Input:** Dependent variable  $List_{Dice} = Y_{Dice}$  (say) =  $[y_{Dice_i}]_{i=1}^n$  for  $n$  images, and independent variables:

•  $List_{T_{unc}} = X_0$  (say) =  $[x_{0i}]_{i=1}^n$  (overall uncertainty as shown in Fig 3 and 4 column D).

•  $List_{T_{unc}} = X_1$  (say) =  $[x_{1i}]_{i=1}^n$  (lesion tissue uncertainty as shown in Fig 3 and 4 column E)

•  $List_{NT_{unc}} = X_2$  (say) =  $[x_{2i}]_{i=1}^n$  (non-lesion tissue uncertainty as shown in Fig 3 and 4 column F) – that predict  $Y_{Dice}$ .

**Output:** Model for  $Y_{Dice}$  in terms of  $X_1$  and  $X_2$ :

$$Y_{Dice} = \alpha_0 + \alpha_1 * X_1 + \alpha_2 * X_2 \quad (3)$$

for individual ROIs of each image:

$$(i)Y_{Dice} = \beta_0 + \beta_1 * X_1, \quad (ii)Y_{Dice} = \gamma_0 + \gamma_1 * X_2 \quad (4)$$

for the overall uncertainty,  $X_0$ :

$$Y_{Dice} = \theta_a + \theta_b * X_0 \quad (5)$$

---

**Algorithm 1: Uncertainty Estimations of Clinical Regions of Interest (ROI)** : Predicting deep learning model (DLM) performance with uncertainty estimations of clinical class-based image regions

values that explain the uncertainty within specific ROIs for three classes are in Table III. The predictive capabilities of the proposed Algorithm 1 and linear models are further described in subsection IV-A4.

## IV. RESULTS

### A. Deep Learning Model Performance and Uncertainty Analysis

This study focused on developing DLMS for skin lesion segmentation, leveraging uncertainty estimates to enhance

trustworthiness and predict Dice coefficients for outputs using region based linear regression models. The utilization of MCD and BBP techniques for training two models, TFS and another based on TL, allowed for a comprehensive assessment of model uncertainty at the pixel level. Subsequently, the analysis was expanded to explore the combined effect of dropout and Bayes-by-backprop techniques within the TFS and TL frameworks. Five-fold cross-validation was implemented to segment skin lesions and cancer tissue to estimate segmentation uncertainties in the trained DLM models reported in this study. The confidence intervals (CI) were calculated using the empirical bootstrap method with  $n = 5,000$  simulations. Table I shows that for the TFS model with dropout (tfs\_d), the reported Dice coefficient was 0.8809 (95% CI: 0.8705 - 0.8802). For the TFS model with backpropagation (tfs\_bc), the reported Dice coefficient was 0.8162 (95% CI: 0.8005 - 0.8250). When utilizing the TL model, the reported Dice coefficient for TL with dropout (tl\_d) were 0.8313 (95% CI: 0.8205 - 0.8471), and for TL with backpropagation (tl\_bc), the Dice coefficient was 0.7702 (95% CI: 0.7605 - 0.7824). In addition to the results from the entire test dataset, we also calculated deep learning model performance on three clinical classes individually. The models demonstrated distinct mean uncertainty values, for e.g. in melanoma tfs\_d had the highest mean uncertainty at 0.061448, while tl\_bc had the lowest at 0.03723. In the seborrheic class, tfs\_d also showed the highest mean uncertainty at 0.060842, with tl\_bc with the lowest at 0.037722. For nevus, the trend continued with tfs\_d resulting in the highest mean uncertainty at 0.052898 and tl\_bc the lowest at 0.026754. The results indicated that both the TFS and TL models achieved high performance (Dice) in skin lesion and cancer segmentation using MCD and BBP (Table I). The results also indicated that the tfs\_d model showed higher uncertainty across all metrics followed by the tl\_d model (Table II), but also achieved higher segmentation performance (Table I). This heightened uncertainty can be attributed to the stochastic behavior of dropout, where units are randomly omitted during training introducing variability in predictions, resulting in higher uncertainty values. Training with backpropagation layers (tfs\_bc and tl\_bc) generally exhibited lower uncertainty compared to their dropout counterparts (Table II), with corresponding lower segmentation performance (Table I). We hypothesize that the dropout layers contribute to higher uncertainty metrics compared to backpropagation layers in both TL and TFS models, and therefore tfs\_d and tl\_d layers exhibit the higher uncertainty metrics compared with tfs\_bc and tl\_bc models as shown in Table (II). Therefore, the choice between dropout and backpropagation in DLMS can be determined on the need for balanced performance and uncertainty management. The dropout method may be suitable for scenarios requiring both accuracy and awareness of uncertainty. Conversely, backpropagation provides stable predictions with lower uncertainty but may result in lower segmentation accuracy compared to dropout.

Analyzing 2935 test images revealed that most of the uncertainty with the tl\_bc approach ( $n=2759$  images) was

TABLE I  
MEDIAN PERFORMANCE METRICS FOR TRAINED FROM SCRATCH (TFS) AND TRANSFER LEARNING (TL) DEEP LEARNING MODELS WITH BAYES BACKPROPAGATION (BC) OR MONTE CARLO DROPOUT (D) LAYERS BASED ON DICE SCORE, AUROC, TRUE POSITIVE RATE (TPR), AND FALSE POSITIVE RATE (FPR).

Model	Performance			
	Dice	AUROC	TPR	FPR
tfs_bc	0.8162	0.8032	0.7829	0.0194
tfs_d	0.8809	0.8662	0.7730	0.0008
tl_bc	0.7702	0.7197	0.7068	0.0278
tl_d	0.8313	0.8168	0.8353	0.0162

TABLE II  
UNCERTAINTY METRICS FOR DEEP LEARNING MODELS TRAINED FROM SCRATCH (TFS) AND TRANSFER LEARNING (TL) WITH BAYES BACKPROPAGATION (BC) OR MONTE CARLO DROPOUT (D).

Model	Uncertainty			
	Maximum	Minimum	Mean	Std. dev.
tfs_bc	0.1094	0.0031	0.0128	0.0147
tfs_d	0.1406	0.0063	0.0250	0.0420
tl_bc	0.0976	0.0029	0.0115	0.0145
tl_d	0.1305	0.0061	0.0159	0.0130

primarily attributed to the skin lesion boundary, while 176 images exhibited uncertainty inside the lesion region (Fig 2). Similarly, using the tfs\_bc approach on the same set of test images, only 253 images demonstrated uncertainty within the lesion region, with the majority (2682 images) showing uncertainty localized on the boundaries (Fig 2). These findings indicate that the MCD Bc (tl\_bc, tfs\_bc) used in this study prioritizes the detection of uncertainty in boundary segmentation. Conversely, with the tl\_d approach, 2924 images exhibited segmentation uncertainty inside the lesion region, and only 11 images had boundary pixels with uncertainty. Similarly, with the tfs\_d approach, uncertainty was centered in the interior of the lesion regions in 2927 images, with only eight images emphasizing the boundary. Thus, the MCD approach with both TFS and TL in DLMs shows the highest uncertainty inside skin lesions and lower uncertainty in other regions of the RGB images.

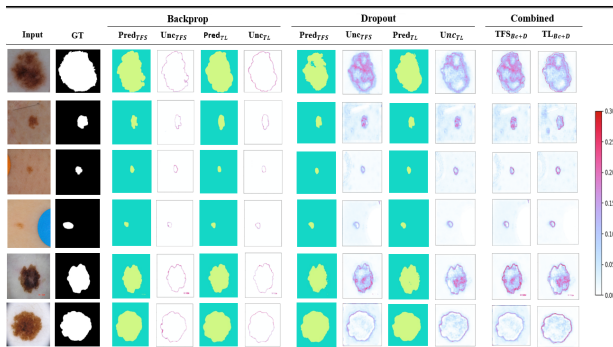


Fig. 2. Segmentation and uncertainty maps obtained from Bayes-by-backprop (Bc) and Monte Carlo dropout (MCD) augmented transfer learning (TL) and trained from scratch (TFS) deep learning models. GT - Ground Truth labels, Pred - Model prediction, Unc-Uncertainty.

This also suggests that the MCD approach exhibits lower performance when predicting within lesion tissue regions but performs well in other regions.

1) *Performance On Prediction Of Nevus Skin Tissue:* From lesion containing images  $\alpha_0 = 0.8449$ ,  $\alpha_1 = -1.1932$  and  $\alpha_2 = -1.5108$  underscore the predictive behavior with low RMSE values ( $\leq 0.16$ ) indicating accurate predictions across all four linear regression models. The mean Dice value of 0.837642 further indicated the DLM and linear regression model efficacy in delineating nevus skin tissue pixels. After evaluating the coefficients and ROI-based uncertainties in Table III column 3, the regression model exhibited varying levels of certainty in different regions of interest. These insights can assist with the interpretation of predictive performance and the reliability of DLMs in segmenting nevus skin lesions.

2) *Performance On Prediction Of Melanoma Skin Lesions:* Table III column 4 shows results for melanoma skin lesions detailing the coefficients and values associated with ROI-based uncertainty. From these images  $\alpha_0 = 0.9644$ ,  $\alpha_1 = -8.6780$  and  $\alpha_2 = -3.0881$  were calculated. The RMSE values for all four linear regression models were  $\leq 0.16$  and the mean Dice value were 0.822958.

3) *Performance On Prediction Of Seborrheic Keratosis Skin Tissue:* Table III column 5 shows a detailed breakdown of the coefficients and associated values that explain the uncertainty within specific ROIs for seborrheic keratosis skin tissue pixels. From these lesion containing images,  $\alpha_0 = 0.8835$ ,  $\alpha_1 = -8.4215$  and  $\alpha_2 = -4.7513$  were calculated. The low RMSE values ( $\leq 0.19$ ) obtained across all four linear models indicated that the models reported high segmentation performance in predicting the lesion presence in these images. Additionally, the mean Dice coefficient of 0.845437 reflects the efficacy of the models in delineating seborrheic keratosis skin tissue pixels.

TABLE III  
ANALYSIS OF COEFFICIENTS AND CORRELATION IN SKIN TISSUE CLASSES. ROI: CLINICAL REGIONS OF INTEREST IN IMAGES. THE VALUES OF  $\beta_0$ ,  $\beta_1$ ,  $\gamma_0$ ,  $\gamma_1$ ,  $\theta_a$ , AND  $\theta_b$  ARE LINEAR REGRESSION COEFFICIENTS.  $\rho$  REPRESENTS THE SPEARMAN'S CORRELATION COEFFICIENT.

ROI-based Uncertainty	Coefficients	Nevus	Melanoma	Seborrheic
Lesion uncertainty	$\beta_0$	0.8350	0.9135	0.8038
	$\beta_1$	-1.5644	-7.7102	-6.3832
	$\rho$	-0.0243	-0.486	-0.3042
Non-lesion uncertainty	$\gamma_0$	0.8433	0.8311	0.8049
	$\gamma_1$	-1.5579	-0.5570	-3.1506
	$\rho$	0.1275	0.0008	-0.1538
Overall uncertainty	$\theta_a$	0.8449	0.9643	0.8769
	$\theta_b$	-1.4596	5.7393	-5.6249
	$\rho$	-0.8496	-0.5556	-0.4164

4) *Predictive capabilities of the proposed linear models:* Eq. 3 was further expanded with the class information for an input image as categorical variable. Thus, Eq. 3 can be rewritten for this case as  $Y_{Dice} = (\alpha_0 + \alpha_1 * X_1 + \alpha_2 * X_2 + \phi_1 * C_1 + \phi_2 * C_2 + \phi_3 * C_3)$ , where  $C_1$ ,  $C_2$  and  $C_3$  (that can take either a value of 0 or 1 based on presence or absence of the class information) correspond to the melanoma, nevus and seborrheic keratosis diagnosis classes. Solving for the values of these coefficients based on the entire test set of 2,935 images,  $\alpha_0 = 0.8852$ ,  $\alpha_1 = -7.5594$ ,  $\alpha_2 = -2.6110$ ,  $\phi_1 = 0.06$ ,

$\phi_2 = -0.0229$  and  $\phi_3 = -0.0372$  were calculated. A low RMSE of 0.1543 was calculated for this linear regression model. A positive coefficient for all classes indicated high segmentation performance. For all derived coefficients in each of the three clinical classes, negative coefficient ( $\alpha_1, \alpha_2, \beta_1, \gamma_1$  and  $\theta_b$ ) values indicated that uncertainty in linear regression models predictions decreased its performance. From the coefficients derived for Eq. 3, we further note the importance of each clinical region as well as the degree to which a clinical diagnosis class affects the trained model's overall performance – a lesser (more negative) value of the coefficient implying a greater contribution in reducing Dice. For the images belonging to the nevus class, the non-lesion region uncertainties had the highest coefficient value (weight) relative to the other region. These results also indicated that linear regression models performance was most affected (reduced) by the uncertainty of predictions of non-lesion region, followed by the lesion-tissue regions. For the melanoma and seborrhoeic keratosis images, lesion tissue region uncertainties contributed more substantially than pixels without any lesion tissue.

### B. Comparative Qualitative Evaluation of Deep Learning Models

Monte Carlo dropout consistently outperformed backpropagation DLMs in skin lesion and cancer segmentation (Table I). Consequently, the qualitative evaluation focused on MCD in TFS and TL models. Visualizations of segmentation masks and uncertainty maps (Figs. 3 and 4) revealed MCD model performance across clinical scenarios. The study also examined the impact of training approach (TFS vs. TL) and MCD utilization on model performance. Figures 3 and 4 display RGB skin dermoscopy images with model output segmentation masks and uncertainty heatmaps. Model predictions and overall uncertainty maps are shown in columns C and D, while region-based uncertainty maps,  $Unc_{iL}$  and  $Unc_{iNL}$ , are in columns E and F (Figs. 3 and 4).

1) *Qualitative Evaluation of the Monte Carlo Dropout Trained From Scratch Deep Learning Model:* **Figure 3** shows six (two from each clinical class) representative skin dermoscopy images from the test data and MCD model output segmentation masks and corresponding uncertainty heatmaps from trained TFS model (tfs\_d). Images 1 and 2 in Fig. 3 belong to the melanoma class, images 3 and 4 to the nevus class and images 5 and 6 to seborrhoeic keratosis. Predictions for images 2, 4, 5 and 6 show almost no false negative and no or very low false positive regions indicating good segmentation accuracy. Image 1 exhibited the highest false negative while image 3 showed the highest false positive regions. For image 1, the false negative regions were prominent where the tfs\_d model showed the highest uncertainties – this could have been a result of the model being unable to differentiate between lighter colored pixels tissues which are part of the lesion ground-truth and the background that contains non-lesion skin tissue which are of similar coloration. Such lighter pixels color in skin lesion and tumor signatures were not effectively captured by the tfs\_d possibly due to the sparsity

of such examples in the data set and is indicated by the high uncertainty values for the region (see Fig. 3 image 1D and 1E).

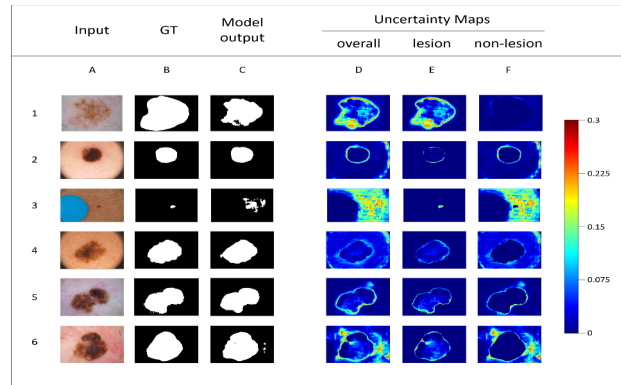


Fig. 3. Visualization of skin cancer segmentation and uncertainty by a Monte Carlo dropout (MCD) model trained from scratch on dermoscopy images. **Columns left to right:** A Input RGB image; B Clinical ground-truth binary mask; C MCD segmentation binary mask; D MCD uncertainty map for segmentation; E MCD uncertainty map for lesion regions; F MCD uncertainty map for non-lesion regions. The color bar shows model uncertainty, with red indicating high uncertainty and blue indicating low uncertainty (0). 'GT' stands for clinical ground-truth.

Figure 3 image 1F shows low uncertainty for pixels in non-lesion regions, indicating superior performance. The uncertainty maps (column D) and lesion uncertainty maps (column E) demonstrate that the tfs\_d model effectively learned the lesion regions and exhibited less uncertainty in non-lesion regions (column F). Most detected uncertainties were within the non-lesion region, at the lesion and non-lesion boundary, or at skin tissue and non-tissue boundaries (images 2 and 4). This indicates that the tfs\_d model accurately learned lesion regions but showed uncertainty at non-lesion skin tissue, non-tissue junctions, and tissue intersections.

In images 3 and 6, uncertainties in non-lesion regions suggest that the tfs\_d model may inaccurately segment darker non-lesion skin tones with lesions, especially in image 3 where the lesion color closely matches the non-lesion color. Additionally, in image 3, the model effectively ignored a blue artifact, as indicated by no uncertainty in that region. Visualizing these ROI-based uncertainty maps (Fig. 3) helps identify high-uncertainty regions and make informed decisions on model output. High uncertainties in the lesion region of image 1 suggest potential lower segmentation performance within lesion regions during inference, while similar uncertainties in image 6 for both lesion and non-lesion regions indicate acceptable predictions when differential diagnosis depends on disease pixels.

2) *Qualitative evaluation of the Monte Carlo dropout transfer learning deep learning model:* **Figure 4** shows six representative skin dermoscopy images from the test data along with the output segmentation masks and corresponding uncertainty heatmaps from the trained Monte Carlo dropout transfer learning deep learning (tl\_d) model. Images 1 and 2

in Fig. 4 belong to the melanoma class, images 3 and 4 to the nevus class, and images 5 and 6 to the seborrheic keratosis. Similar to the tfs\_d evaluation, we observed that predictions for images 1, 2, 4, 5, and 6 exhibited almost no false negatives and no or very low false positive regions, indicating good segmentation accuracy. However, there were differences in the uncertainty patterns and performance analysis compared to the tfs\_d. For instance, in image 1 (Fig. 3 and 4) the tfs\_d model exhibited higher uncertainties in the lesion region compared to the tl\_d model, indicating potential challenges in accurately segmenting melanoma lesions during inference. Furthermore images 3 and 6 (Fig. 3 and 4) showed lower uncertainties in the non-lesion regions in tl\_d model compared to the tfs\_d, suggesting improved performance of the tl\_d model in distinguishing between lesion and non-lesion skin tissue in these cases.

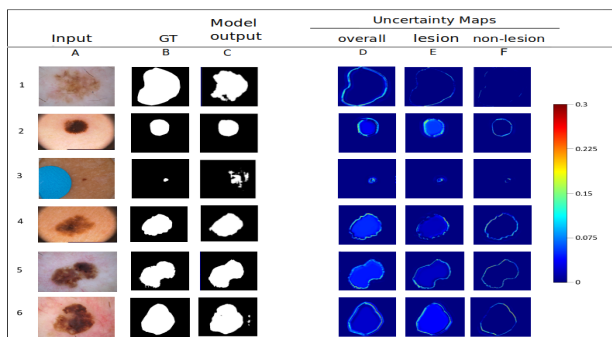


Fig. 4. Visualization of skin cancer segmentation and uncertainty by a Monte Carlo dropout (MCD) model trained with transfer learning on dermoscopy images. **Left to right columns:** **A** Input RGB image; **B** Clinical ground-truth binary mask; **C** MCD segmentation binary mask; **D** MCD uncertainty map for segmentation; **E** MCD uncertainty map for lesion regions; **F** MCD uncertainty map for non-lesion regions. The color bar shows model uncertainty, with red indicating high uncertainty and blue indicating low uncertainty (0). ‘GT’ stands for clinical ground-truth.

Visualizing the ROI-based uncertainty maps (Fig. 4) helped identify regions of high uncertainties, informing decisions about the model’s output. In image 1, high uncertainties were observed at the boundaries, indicating potential improvement areas for the tl\_d model. Unlike the tfs\_d model, the tl\_d model showed minimal uncertainties in both lesion and non-lesion regions, suggesting more reliable predictions for differential diagnosis reliant on disease pixels. Overall uncertainty maps (column D) and lesion uncertainty maps (column E) indicated that the tl\_d model learned lesion regions well, exhibiting less uncertainty compared to tfs\_d (column F) (Fig. 3 and 4). Most detected uncertainties were within the lesion and non-lesion boundary regions or at the junctions between skin tissue and non-tissue. This demonstrated that the tl\_d model had effectively learned lesion regions and showed less uncertainty at non-lesion skin tissue and tissue intersections. Both models showed strengths in segmentation accuracy and learning lesion regions but faced challenges with specific lesion types. The tl\_d model was better in distinguishing between lesion and non-lesion tissues, providing more reliable

predictions for disease pixels compared to tfs\_d (Fig. 3 and 4). For the tl\_d model, uncertainties were prominent around the edges where lesion tissue pixels meet non-lesion or background tissue pixels, while for the tfs\_d model, uncertainties were prominent within lesion tissues. The tfs\_d model was less optimal in predicting non-lesion tissue regions, possibly due to insufficient examples with clear boundary demarcations and differential features in the dataset.

## V. DISCUSSION

In this study, the performance of deep learning models (DLMs) using TFS and TL with Bayes-by-Backprop (Bc) and Monte Carlo Dropout (MCD) techniques for skin cancer segmentation was examined (Fig. 2). The models employing TFS or TL with Bayes-by-Backprop (tfs\_bc, tl\_bc) showed higher segmentation uncertainty concentrated on the boundaries of the skin lesion, indicating lower uncertainty within the lesion region. These boundaries typically represent areas where skin tumor tissue pixels intersect with non-lesion or background tissue pixels. In contrast, TFS or TL with Monte Carlo Dropout (tfs\_d, tl\_d) showed more prevalent uncertainties within the tumor tissues compared to the DLM Bc model (see Fig. 2).

The combination of DLMs Bc (tfs\_bc, tl\_bc) and DLMs MCD (tfs\_d, tl\_d) techniques involved integrating predictions from each method. This approach leveraged the strengths of both techniques to improve overall performance and generalization. These results show that combining DLM Bc and DLM MCD predictions within the TFS and TL frameworks reduced uncertainty and provided more robust predictions, improving segmentation accuracy with low uncertainties (Table I and II).

This study also developed a novel method to estimate uncertainty in DLM predictions for skin lesion segmentation. By correlating uncertainty with Dice scores using linear regression, the proposed algorithm enhances performance estimation and interpretability through the generation of uncertainty maps.

Quantitative analysis through five-fold cross-validation revealed significant correlations between predictive performance, uncertainty measures, and ground-truth segmentation masks. Qualitative evaluation validated the effectiveness of the trained models, reporting accurate segmentation with minimal false positive and false negative pixels. The study also offered specific guidelines for choosing between TFS and TL approaches in skin lesion and cancer image segmentation tasks by DLMs. High performance metrics in AUROC, Dice coefficient, and TPR were achieved by employing TFS with dropout layers, delivering superior performance and robust uncertainty estimates (Table I and II). Using TL DLM with dropout layers balanced performance and uncertainty, making it suitable for scenarios requiring both accuracy and uncertainty awareness. While TFS with backpropagation layers provided stable predictions with lower uncertainty estimations, albeit with potential trade-offs in lower segmentation performance (Table I and II). Thus, integrating both dropout and backpropagation techniques within TFS and TL frameworks results in a comprehensive uncertainty reduction strategy for

enhancing model robustness and precision for skin lesion segmentation (Fig. 2).

The trained deep learning models from this study accurately align with ground-truth data and can be adapted and fine-tuned using Algorithm 1 for small amounts of unlabelled, out-of-distribution (OoD) data to minimize false positives and maintain high performance levels. When tested with OoD data lacking true ground-truth, the study employs historical performance metrics such as Dice scores and AUROC (Table I and II) to evaluate model effectiveness, utilizing Equation 5 where appropriate. Thus, Algorithm 1 is versatile and promising for working with other skin imaging datasets for improving segmentation tasks.

a) *Limitations:* While the performance metrics of the framework, approach, and models in this study are promising, it is important to acknowledge that the calculated performance could vary for other practitioners. This variability can be attributed to several factors unique to their study, including the complexity of the model architecture, data preprocessing techniques, refinement of feature selection processes, and the choice of pipeline.

## VI. CONCLUSION

This study analyzes the relationship between prediction uncertainty and the Dice score of whole slide image segmentation. It demonstrates that clinical region-based linear models can predict DLM performance (Dice) using individual or combined uncertainty estimations (Equations: 3- 5). Unlike other studies that focus only on overall uncertainty measures, this study decouples clinical and region-based uncertainty estimates for skin lesions and cancer signatures. Significant correlations ( $p < 0.05$ ) between overall and ROI-based uncertainties with Dice scores were established using Spearman's rank correlation. High negative values indicated a strong negative correlation between uncertainty and model performance, supported by linear regression coefficients, most of which were negative. Low RMSE indicated a good fit for these results. Equation 3 predicts performance based on all regions simultaneously, identifying regions contributing to greater uncertainty and poorer performance. Algorithm 1 estimates model performance unsupervised, generating uncertainty maps and mean uncertainty values. Clinical ROI-based uncertainty estimates and predicted Dice scores can inform real-world clinical decisions. Future research will evaluate models across diverse skin and other organs for multi-class segmentation, integrate domain knowledge, and simulate aleatoric and epistemic noise during training for further refinement in clinical and machine learning outcomes.

## VII. DATA, CODE AND MODEL AVAILABILITY

All resources, results, data, code, and models, are deposited in a GitHub repository [https://github.com/Prof-Pratik-Shah-Lab-UCI/skin\\_project\\_2024](https://github.com/Prof-Pratik-Shah-Lab-UCI/skin_project_2024), and will be made available upon request to the corresponding author.

## REFERENCES

- [1] Skin Cancer Foundation., "Skin cancer facts page," <https://www.skincancer.org/skin-cancer-information/skin-cancer-facts/>, 2024, accessed February 2024.
- [2] American Academy of dermatology Association., "Skin cancer facts page," <https://www.aad.org/media/stats-skin-cancer>, 2022, accessed 22-April-2022.
- [3] M. Goyal, T. Knackstedt, S. Yan, and S. Hassanpour, "Artificial intelligence-based image classification methods for diagnosis of skin cancer: Challenges and opportunities," *Computers in biology and medicine*, vol. 127, p. 104065, 2020.
- [4] G. Choy, O. Khalilzadeh, M. Michalski, S. Do, A. E. Samir, O. S. Pinykh, J. R. Geis, P. V. Pandharipande, J. A. Brink, and K. J. Dreyer, "Current applications and future impact of machine learning in radiology," *Radiology*, vol. 288, no. 2, pp. 318–328, 2018.
- [5] Y. Gal and Z. Ghahramani., "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [6] C. M. Bishop., "Bayesian neural networks," *Journal of the Brazilian Computer Society*, vol. 4, pp. 61–68, 1997.
- [7] 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, no. 1, pp. 1–9, 2018.
- [8] J. Senapati, A. G. Roy, S. Pölsterl, D. Gutmann, S. Gatidis, C. Schlett, A. Peters, F. Bamberg, and C. Wachinger, "Bayesian neural networks for uncertainty estimation of imaging biomarkers," in *Machine Learning in Medical Imaging: 11th International Workshop, MLMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 11*. Springer, 2020, pp. 270–280.
- [9] K. Hoebel, V. Andrearczyk, A. Beers, J. Patel, K. Chang, A. De-peursinge, H. Müller, and J. Kalpathy-Cramer, "An exploration of uncertainty information for segmentation quality assessment," in *Medical Imaging 2020: Image Processing*, vol. 11313. SPIE, 2020, pp. 381–390.
- [10] H. P. Do, Y. Guo, A. J. Yoon, and K. S. Nayak, "Accuracy, uncertainty, and adaptability of automatic myocardial asl segmentation using deep cnn," *Magnetic resonance in medicine*, vol. 83, no. 5, pp. 1863–1874, 2020.
- [11] P. Van Molle, T. Verbelen, C. De Boom, B. Vankeirsbilck, J. De Vylder, B. Diricx, T. Kimpe, P. Simoons, and B. Dhoedt, "Quantifying uncertainty of deep neural networks in skin lesion classification," in *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures*. Springer, 2019, pp. 52–61.
- [12] M. Combalia, F. Hueto, S. Puig, J. Malvehy, and V. Vilaplana, "Uncertainty estimation in deep neural networks for dermoscopic image classification," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 2020, pp. 744–745.
- [13] M. Shapna Akter, H. Shahriar, S. Sneha, and A. Cuzzocrea, "Multi-class skin cancer classification architecture based on deep convolutional neural network," *arXiv e-prints*, pp. arXiv–2303, 2023.
- [14] H. Ashraf, A. Waris, M. F. Ghafoor, S. O. Gilani, and I. K. Niazi, "Melanoma segmentation using deep learning with test-time augmentations and conditional random fields," *Scientific Reports*, vol. 12, no. 1, p. 3948, 2022.
- [15] G. G. De Angelo, A. G. Pacheco, and R. A. Krohling, "Skin lesion segmentation using deep learning for images acquired from smartphones," in *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2019, pp. 1–8.
- [16] M. Nawaz, T. Nazir, M. Masood, F. Ali, M. A. Khan, U. Tariq, N. Sahar, and R. Damaševičius, "Melanoma segmentation: A framework of improved densenet77 and unet convolutional neural network," *International Journal of Imaging Systems and Technology*, vol. 32, no. 6, pp. 2137–2153, 2022.
- [17] "ISIC Archive homepage," 2019, <https://www.isic-archive.com/#!/topWithHeader/wideContentTop/main>.
- [18] M. Combalia, N. C. Codella, V. Rotemberg, B. Helba, V. Vilaplana, O. Reiter, C. Carrera, A. Barreiro, A. C. Halpern, S. Puig et al., "Bcn20000: Dermoscopic lesions in the wild," *arXiv preprint arXiv:1908.02288*, 2019.
- [19] S. Ghosal and P. Shah, "A deep-learning toolkit for visualization and interpretation of segmented medical images," *Cell Reports Methods*, vol. 1, no. 7, 2021.
- [20] K. Simonyan and A. Zisserman., "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.