
Doctoral Dissertations

Student Theses and Dissertations

Summer 2024

Improvements in Biomedical Image Analysis with Computational Intelligence and Data Fusion Techniques

Akanksha Maurya

Missouri University of Science and Technology

Follow this and additional works at: https://scholarsmine.mst.edu/doctoral_dissertations



Part of the [Computer Engineering Commons](#)

Department: **Electrical and Computer Engineering**

Recommended Citation

Maurya, Akanksha, "Improvements in Biomedical Image Analysis with Computational Intelligence and Data Fusion Techniques" (2024). *Doctoral Dissertations*. 3314.

https://scholarsmine.mst.edu/doctoral_dissertations/3314

This thesis is brought to you by Scholars' Mine, a service of the Missouri S&T Library and Learning Resources. This work is protected by U. S. Copyright Law. Unauthorized use including reproduction for redistribution requires the permission of the copyright holder. For more information, please contact scholarsmine@mst.edu.

IMPROVEMENTS IN BIOMEDICAL IMAGE ANALYSIS WITH
COMPUTATIONAL INTELLIGENCE AND DATA FUSION TECHNIQUES

by

AKANKSHA MAURYA

A DISSERTATION

Presented to the Graduate Faculty of the

MISSOURI UNIVERSITY OF SCIENCE AND TECHNOLOGY

In Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

in

ELECTRICAL ENGINEERING

2023

Approved by:

R. Joe Stanley, Advisor
William V. Stoecker
Kurt Kosbar
Donald Wunsch
Randy H. Moss

© 2023

Akanksha Maurya

All Rights Reserved

PUBLICATION DISSERTATION OPTION

This dissertation consists of the following three articles:

Paper I, found on pages 6–20, has been published in *Skin Research and Technology*

Paper II, found on pages 21–50, has been under review in the *Journal of Digital Imaging*.

Paper III, found on pages 51–83, has been under review in the *Journal of Digital Imaging*.

ABSTRACT

An estimated 2 million new cases of basal cell carcinoma (BCC) are diagnosed each year in the United States, making it one of the most common skin cancers. Earlier detection of these cancers enables less invasive biopsies. Clinical detection consists of a preliminary visual observation of these skin lesions by an experienced dermatologist making it a specialized task highly dependent on their time, availability, and resources. Hence, there is a need for automating this process that can assist healthcare staff. In recent years, deep learning (DL) has been used extensively and successfully to diagnose different cancers in dermoscopic images. Telangiectasia or narrow blood vessels that typically appear serpiginous or arborizing, are a critical indicator of basal cell carcinoma (BCC), aiding dermatologists in BCC diagnosis. Most DL approaches lack such clinical inputs that could aid in higher accuracy and explainability. Hence, in this research, we exploit the following computational and data fusion techniques for BCC feature detection and diagnosis: 1. Automate the segmentation of telangiectasia with the application of image processing techniques and a semantic deep learning model. 2. Apply ensemble learning on a combination of Deep learning features and handcrafted features from semantically segmented telangiectasia masks for BCC diagnosis. 3. Explore topological data analysis (TDA) techniques to create a DL-TDA based hybrid classification model. Through this research we achieve state-of-the-art results in BCC diagnosis and provide pathways for automating diagnosis/classification for similar datasets and problem statements.

ACKNOWLEDGMENTS

I am grateful to my advisor, Dr. R. Joe Stanley, for his support, encouragement, and kindness throughout my PhD. His feedback and constant motivation lifted my spirits and helped me tremendously in this research. I consider myself extremely fortunate to have had the opportunity to work with Dr. William V. Stoecker. His guidance throughout has been another highlight of my PhD. I would like to thank Dr Randy Moss, Dr. Kurt Kosbar and Dr Donald Wunsch for their insightful comments and feedback.

I would like to thank my life partner, Corey Harrington for being my pillar of strength throughout my Ph.D. His presence and unwavering support made every hurdle easier. A very special thanks to all my colleagues: Anand Nambisan, Jason Hagerty and Norsang Lama for their help and friendship during my time in our lab. I again, consider myself lucky to have found such inspiring and amazing friends and co-workers.

I am indebted to my family in India for their faith in me and my vision, especially my elder brother Aadarsh who nurtured me and encouraged me every step of the way. My mother Vidya Maurya and my father L. C. Maurya are extraordinary parents and words fall short in describing my love for them and the support they have given me throughout my life.

Lastly, I want to thank all the great artists of the world, especially musicians for creating the beautiful art that accompanied me almost every day of my PhD.

TABLE OF CONTENTS

	Page
PUBLICATION DISSERTATION OPTION	iii
ABSTRACT.....	iv
ACKNOWLEDGMENTS	v
LIST OF ILLUSTRATIONS.....	x
LIST OF TABLES	xiii
NOMENCLATURE	xiv
 SECTION	
1. INTRODUCTION.....	1
1.1. OVERVIEW	2
1.2. PROBLEM DESCRIPTION	2
1.3. SUMMARY OF CONTRIBUTIONS	4
1.3.1. A Deep Learning Approach to Detect Telangiectasia in Basal Cell Carcinoma.....	4
1.3.2. Basal Cell Carcinoma Diagnosis with Fusion of Deep Learning and Telangiectasia Features.....	4
1.3.3. Hybrid Topological Data Analysis and Deep Learning for Basal Cell Carcinoma Diagnosis.....	5
 PAPER	
I. A DEEP LEARNING APPROACH TO DETECT BLOOD VESSELS IN BASAL CELL CARCINOMA.....	6
ABSTRACT.....	6
1. INTRODUCTION	7

2. IMAGE DATASETS AND PRE-PROCESSING	8
2.1. IMAGE DATASETS	8
2.2. IMAGE PRE-PROCESSING.....	9
2.3. IMAGE SEGMENTATION	10
3. DEEP LEARNING NETWORK	10
3.1. NETWORK ARCHITECTURE	10
4. EXPERIMENTS PERFORMED	12
5. EXPERIMENTAL RESULTS.....	13
5.1. VESSEL DETECTION PERFORMANCE FOR DIFFERENT LOSS FUNCTIONS	14
6. DISCUSSION	15
REFERENCES.....	18
II. BASAL CELL CARCINOMA DIAGNOSIS WITH FUSION OF DEEP LEARNING AND TELANGIECTASIA FEATURES	21
ABSTRACT.....	21
1. INTRODUCTION	22
2. MATERIALS AND METHODS.....	24
2.1. IMAGE DATASETS	24
2.2. DATA AUGMENTATION	26
2.3. PROPOSED METHODOLOGY	28
2.4. TRAINING DETAILS	35
3. EXPERIMENTAL RESULTS.....	36
3.1. VESSEL SEGMENTATION TEST RESULTS FOR BCC AND NON-BCC IMAGES	36
3.2. DEEP LEARNING TRAINING RESULTS AFTER FINE-TUNING.....	37

3.3. FEATURE IMPORTANCE WITH RANDOM FOREST.....	39
3.4. FINAL CLASSIFICATION WITH DEEP LEARNING AND RANDOM FOREST CLASSIFIER	40
3.5. PERFORMANCE COMPARISON WITH EXISTING METHODS	42
4. DISCUSSION	43
5. CONCLUSION AND FUTURE WORK.....	45
ACKNOWLEDGEMENT	46
REFERENCES.....	46
III. HYBRID TOPOLOGICAL DATA ANALYSIS AND DEEP LEARNING FOR BASAL CELL CARCINOMA DIAGNOSIS.....	51
ABSTRACT.....	51
1. INTRODUCTION	52
2. MATERIALS AND METHODS.....	54
2.1. IMAGE DATASETS	54
2.2. PRE-PROCESSING.....	55
2.3 PROPOSED METHODOLOGY	56
2.3.1. U-Net for Telangiectasia Masks.....	57
2.3.2. Topological Data Analysis (TDA) and Persistent Homology	57
2.3.3. Persistence Statistics for Topological Features	65
2.3.4. Transfer Learning with EfficientNet-B5 for Feature Extraction	68
2.3.5. Class Probabilities with Majority Voting	70
2.4. TRAINING DETAILS.....	71
3. EXPERIMENTAL RESULTS.....	71
3.1. U-NET TELANGIECTASIA SEGMENTATION RESULTS	72

3.2. BCC DIAGNOSIS WITH PERSISTENCE STATISTICS AND RANDOM FOREST	73
3.3. DEEP LEARNING RESULTS WITH EFFICIENTNET-B5	74
3.4. BCC VS NON-BCC WITH HYBRID TDA-DL MODEL	75
3.5. PERFORMANCE COMPARISON WITH EXISTING METHODS	76
4. DISCUSSION	77
5. CONCLUSION	78
REFERENCES	79
SECTION	
2. SUMMARY AND CONCLUSIONS	84
BIBLIOGRAPHY	85
VITA	91

LIST OF ILLUSTRATIONS

PAPER I	Page
Figure 1. Vessels in BCC. Arborizing and serpiginous telangiectasia vs. non-specific sun-damage telangiectasia.....	9
Figure 2. Color augmentations.....	10
Figure 3. U-Net architecture	12
Figure 4. Predicted binary masks and overlays	14
Figure 5. Training and validation loss curves for combo loss	15
Figure 6. Mask annotations for the same image by different team members	16
Figure 7. Example of disagreement with manual mask. True positives shown by green, false positives shown by blue and false negatives shown by yellow	17
PAPER II	
Figure 1. Left to right, first row: first two images are BCC from the HAM10000 (ISIC2018) dataset; last two images are BCC from the NIH study dataset. Second row presents telangiectasia overlays for the images in the first row	25
Figure 2. From left, first row: benign keratosis, nevus, actinic keratosis; from left second row: dermatofibroma and vascular lesion	26
Figure 3. Number of images from each category used in the non-BCC dataset.....	26
Figure 4. Different geometric augmentations	27
Figure 5. Left: Color augmentations Right: vessel masks dilation.....	28
Figure 6. Proposed architecture employing a fusion of deep learning and handcrafted features from vessels for BCC classification	29
Figure 7. Object surround masks: Top row from left: Example lesion image 1 with telangiectasia, its ground truth mask, its surround mask for object 1 (the contiguous connected vessel area), its surround mask for object 2 (the vessel on the bottom right). Bottom row from left: Example lesion image 2 with telangiectasia, its ground truth mask, its surround mask for object 1 (the contiguous connected vessel area), its surround mask for object 2 (the vessel on the bottom right)	30

Figure 8. MBConv1 & MBConv6 blocks.....	32
Figure 9. Squeeze & Excitation block	34
Figure 10. From left: Column 1 shows non-BCC images, column 2 shows their predicted masks, column 3 shows BCC images, and column 4 shows predicted BCC masks.....	37
Figure 11. Loss curves before and after fine-tuning the EfficientNet-B5 based model ...	38
Figure 12. Accuracy curves before and after fine-tuning the EfficientNet-B5 based model.....	38
Figure 13. Importance scores for deep learning features	39
Figure 14. Importance scores for handcrafted features.....	39
Figure 15. ROC curve for our final fusion classification model.....	42
PAPER III	
Figure 1. From left: actinic keratosis, benign keratosis, dermatofibroma, basal cell carcinoma, nevus and vascular lesion	55
Figure 2. Pre-processing flowchart	56
Figure 3. Pipeline investigated employing a hybrid TDA-DL method for BCC classification.....	57
Figure 4. Flow of the persistent homology algorithm to generate persistence diagrams .	58
Figure 5. Simplicial Complex containing a 3D simplex (tetrahedron) and a 2D simplex (triangle)	60
Figure 6. Persistent homology filtration process leading to formation of birth-death pairs in the persistence diagram	64
Figure 7. Top row, from left: BCC skin lesion; Its corresponding P0 persistence diagram and P1 Persistence diagram for the red plane; Bottom row, from left: non-BCC (actinic keratosis) skin lesion; Its corresponding P0 persistence diagram and P1 Persistence diagram for the red plane.....	65
Figure 8. Generation of persistence statistics-based features	66

Figure 9. From left: non-BCC image with vessel mask prediction, BCC image with vessel mask prediction	72
Figure 10. Training and validation accuracy and loss curves for fine tuning EfficientNet-B5 on BCC vs non-BCC dataset	75

LIST OF TABLES

PAPER 1	Page
Table 1. Different hyperparameters used for training U-Net.....	13
Table 2. Evaluation metrics for different loss functions with U-Net.....	14
 PAPER II	
Table 1. Handcrafted Feature Description	31
Table 2. Efficientnet-B5 Based DL Model Description.....	35
Table 3. Performance Comparison of Different Fusion Models.....	41
Table 4. Performance Comparison with Baseline DL Model.....	41
Table 5. Performance Comparison with Other Methods	42
 PAPER III	
Table 1. Persistence Statistics calculated for our methodology.....	68
Table 2. Our EfficientNet-B5 based deep learning model.....	70
Table 3. Hyperparameters for the EfficientNet-B5 model.....	71
Table 4. Some example PS features calculated from vessel masks	73
Table 5. Random Forest classification with Persistence Statistics for R, G, B, V, Z and vessels	74
Table 6. Metric improvements with Subsets of PS Features	74
Table 7. Performance comparison of different deep learning models	75
Table 8. Performance comparison of DL model, PS model and hybrid TDA-DL model	76
Table 9. Performance comparison with other studies.....	76

NOMENCLATURE

Symbol	Description
AUC	Area under curve
BCC	Basal Cell Carcinoma
TDA	Topological Data Analysis
DL	Deep Learning
TP	True Positive
FP	False Positive
TN	True Negative
FN	False Negative
y	True Value
\hat{y}	Predicted Value
DSC	Dice Coefficient
JAC	Jaccard Index
GLCM	Grey Level Co-occurrence Matrix
MLP	Multi-Layer Perceptron
PH	Persistent Homology
PS	Persistence Statistics

1. INTRODUCTION

Basal Cell Carcinoma or BCC accounts for 80% of all skin cancer cases in the United States. As mentioned by the American Cancer Society “According to one estimate, about 5.4 million basal and squamous cell skin cancers are diagnosed each year in the US (occurring in about 3.3 million Americans, as some people have more than one). About 8 out of 10 of these are basal cell cancers” [1][2].

BCC usually appears on exposed areas of the skin such as arms, face, and neck. A BCC lesion may appear raised, bumpy, or scaly patch which can be skin colored or brown and may even sometimes resemble a sore. Due to this, it sometimes may not be timely treated. Even though the cancer is not considered deadly, if left untreated, it is prone to become locally invasive may lead to disfigurement and other complications. These complications can include ulceration, bleeding, infection, recurrence and sometimes metastasis. Hence, it is important to diagnose and treat BCC early to prevent it from spreading and causing further damage. Early detection improves the chances of a favorable outcome.[3][4]

BCC is usually diagnosed by a dermatologist and the screening involves visual examination, where the dermatologist is looking for any suspicious lesions. They may use a dermascope for magnified versions of the lesion. If there is suspicion, they may perform a biopsy and take a small sample of the tissue. To prevent invasive biopsies and associated costs, early detection through other methods of diagnosis is being developed to assist dermatologists and healthcare staff, which may also benefit patients in remote areas.[5][6]

1.1. OVERVIEW

Basal Cell Carcinoma (BCC) is one of the most widespread forms of skin cancer worldwide. To diagnose this cancer, dermatologists perform visual inspection of the affected skin lesion area which involves looking for certain cancer indicating physical features, such as Telangiectasia or thin blood vessels. Digital images of the lesion are a zoomed in version of the skin lesion. To assist dermatologists in accessing these lesions for biomarkers such as Telangiectasia, we explore several automation techniques for BCC diagnosis as well as clinical feature segmentation. Such techniques would help them avoid invasive biopsies and have extra diagnostic assistance.

1.2. PROBLEM DESCRIPTION

In recent years, biomedical image analysis with the help of a plethora of computational methods has gained fruitful momentum. Computational methods such as deep learning applied with the goal of classification of medical conditions as well as segmentation/extracting features have resulted in a lot of success and understandably, increased clinical and research interests.[3-8]

Computational and quantitative methods especially for digital medical image analysis strive to find the underlying distinguishing patterns that may or not be observable by the human eye. As such Deep Learning (DL) techniques using dermoscopic images have recently shown diagnostic accuracy exceeding that of dermatologists. Due to the black box nature of deep learning models, these results may sometimes lack interpretability, especially from a clinical perspective [9-13]

Telangiectasis or thin arborizing blood vessels present within the skin lesion are an important biomarker for BCC. Detecting and segmenting these vessels either by visual inspection or computational methods can help dermatologists proceed with a more probable and confident BCC diagnosis. These blood vessels usually appear a little darker than the surrounding skin and hence a pixel rule-based color drop vessel detection methods have been used in the past [14-16]. Stacked sparse autoencoders, independent component analysis, k-means clustering, and shape filters to detect vessels and other vascular structures are some of the other methods that have been used. There are very few image processing studies that utilize these vascular structures to diagnose BCC. An adaptive critic design approach to discriminate vessels from competing structures, and features extracted from these vascular structures to classify BCC using a random forest classifier are some of the other works. Features learned from sparse autoencoders, combined with patient profile information, are some meta-data fusion diagnosis methods used in the past [17-22]. Recently, used clustering-based color features and GLCM-based texture features to train VGG16 and MLP models to extract deep learning features to subsequently train another MLP model was applied [22-35].

In this dissertation we explore a combination of different data fusion techniques and a hybrid of TDA and DL methods to diagnose BCC. We dive into first integrating image processing to segment Telangiectasia and then utilize it for BCC diagnosis. We then fuse handcrafted features from telangiectasia in a DL based ensemble learning model [35-46]. Lastly, we apply TDA to create a hybrid model for BCC diagnosis with less computational cost.

1.3. SUMMARY OF CONTRIBUTIONS

This dissertation is comprised of 3 journal papers as listed in the publications list.

The unique contributions from each article can be summarized as follows:

- First study of pixel and DL-based telangiectasia segmentation in skin lesions
- A unique dataset consisting of thousand binary vessel masks for telangiectasia publicly released on Zenodo
- An improved fusion-DL model for BCC diagnosis with handcrafted telangiectasia features
- A hybrid TDA-DL model with state-of-the-art BCC diagnosis

1.3.1. A Deep Learning Approach to Detect Telangiectasia in Basal Cell

Carcinoma. Telangiectasia or thin blood vessels is the first step in identifying BCC. In this study, we automate the detection and segmentation of telangiectasia by using a pixel based Deep Learning approach. A unique binary vessel mask dataset is curated for applying a U-Net based segmentation model that outputs vessel mask in candidate images. To overcome the similarity of vessel and lesion pixels a preprocessing histogram equalization method is used on the BCC images. We optimize the performance by using a combination loss function to manage class imbalance of images and pixel imbalance of skin versus vessel pixels. We achieve Jaccard scores within the variation of human observers and analyze the results for issues in medical segmentation such as inter-observer variability.

1.3.2. Basal Cell Carcinoma Diagnosis with Fusion of Deep Learning and

Telangiectasia Features. We utilize the segmented telangiectasia masks from the study above to implement a novel BCC diagnosis technique. We apply a range of image

processing techniques to yield different telangiectasia and lesion shape and color features. On a hold-out dataset of 395 lesion images (195 BCC and 200 non-BCC) we achieve a binary classification accuracy of 97.2% and an AUC of 0.99. We demonstrate metric improvements in three stages: 1) the addition of handcrafted telangiectasia features to deep learning features, 2) including areas near telangiectasia (surround areas), 3) discarding the noisy lower-importance features. Examining the surround areas of telangiectasia and calculating respective features, we offer another novel approach to feature finding with weak annotations. Our experimental results show state-of-the-art accuracy and precision in the diagnosis of BCC, compared to three benchmark techniques.

1.3.3. Hybrid Topological Data Analysis and Deep Learning for Basal Cell Carcinoma Diagnosis. Topological Data Analysis is a growing field in applied mathematics and is rapidly being explored for deep learning applications. In this study, we apply persistence homology to generate topological features known as persistence statistics and produce results comparable and even better than a few DL models. We use a hybrid TDA-DL model to ultimately improve diagnostic accuracy of BCC. We show improvements in two stages: first by adding telangiectasia features to a TDA classification model and second: by adding TDA features to an EfficientNet based BCC classification model. We achieve state-of-the-art accuracy of 97.4% and an AUC of 0.994 on a hold-out test set of 395 images. Our results demonstrate telangiectasia improves BCC diagnosis and TDA can potentially improve DL results.

PAPER

I. A DEEP LEARNING APPROACH TO DETECT BLOOD VESSELS IN BASAL CELL CARCINOMA

Akanksha Maurya¹, R. Joe Stanley¹, Norsang Lama¹, Sadhika Jagannathan², Daniyal Saeed³, Samantha Swinfard¹, Jason R. Hagerty⁴, William V. Stoecker⁴

¹Missouri University of Science & Technology, Rolla MO 65209 USA

²University of Missouri, Kansas City MO USA

³St Louis University, MO USA

⁴S&A Technology, Rolla, MO, 65401 USA

ABSTRACT

Purpose: Blood vessels called telangiectasia are visible in skin lesions with the aid of dermoscopy. Telangiectasia are a pivotal identifying feature of basal cell carcinoma. These vessels appear thready, serpiginous and may also appear arborizing, i.e., wide vessels branch into successively thinner vessels. Due to these intricacies, their detection is not an easy task, neither with manual annotation nor with computerized techniques. In this study, we automate the segmentation of telangiectasia in dermoscopic images with a deep learning U-Net approach.

Methods: We apply a combination of image processing techniques and a deep learning-based U-Net approach to detect telangiectasia in digital basal cell carcinoma skin cancer images. We compare loss functions and optimize the performance by using a combination loss function to manage class imbalance of skin vs vessel pixels.

Results: We establish a baseline method for pixel-based telangiectasia detection in skin cancer lesion images. An analysis and comparison for human observer variability in annotation is also presented.

Conclusion: Our approach yields Jaccard score within the variation of human observers as it addresses a new aspect of the rapidly evolving field of deep learning: automatic identification of cancer-specific structures. Further application of DL techniques to detect dermoscopic structures and handle noisy labels is warranted.

KEYWORDS: Blood vessels, telangiectasia, dermoscopy, deep learning, skin cancer, basal cell carcinoma.

1. INTRODUCTION

An estimated two million new cases of basal cell carcinoma (BCC), the most common type of skin cancer, are diagnosed each year in the USA [1]. Earlier detection of these cancers enables less invasive treatment [2,3].

Deep Learning (DL) techniques using dermoscopic images have recently shown diagnostic accuracy exceeding that of dermatologists [4-6]. Recent studies have shown improved results for skin cancer diagnoses by fusing ensembles, in some cases handcrafted and DL techniques [7-11]. However, these studies have not employed DL at the dermoscopic level, e.g., to detect blood vessels, a critical sign for BCC (Figure 1).

In previous studies, Cheng et al. [12] used a local pixel color drop technique to identify candidate vessel pixels. Cheng et al. [13] used an adaptive critic design approach to better discriminate vessels from competing structures. Kharazmi et al. [14] used independent component analysis, k-means clustering, and shape filters to detect

vessels and other vascular structures. Kharazmi et al. [15] used a stacked sparse auto-encoders (SSAE) DL approach to detect vessel patches in BCC.

The detection of these intricate cancer-signaling vessels is not an easy task as the data and annotations for these images are limited. Also, these vessels may be blurry and share color similarity with surrounding skin. In this study we explore U-Net, a deep learning-based neural network for vessel segmentation. Our approach is a pixel-based method that captures structures that can elude patch-based methods. Inter-observer variability that is reflected in object labelling is also a widespread issue among medical image datasets and recent studies have shown this can affect model training significantly [16]. Hence, we analyze our annotations for the vessel data and consider this variability for comparison of metrics.

The techniques described in this paper will allow those new to the field, including the growing number of mid-level providers, to automatically identify these critical structures for early cancer detection. It will also benefit researchers seeking to precisely capture features needed for classification and diagnosis of BCC combining other deep learning methods.

2. IMAGE DATASETS AND PRE-PROCESSING

2.1. IMAGE DATASETS

We use two datasets of dermoscopic images of BCC. The first is the HAM10000 [17] dataset of Tschandl et al., a publicly available dermoscopic image dataset with images of size 450x600. The second data set was taken from NIH studies R43

CA153927-01 and CA101639-02A2 with images of size 768x1024. We processed 690 images—413 from the NIH study and 217 from HAM10000 data.

2.2. IMAGE PRE-PROCESSING

Since the two datasets of BCC images differ significantly in image characteristics such as image size and color range, we use stratified training, validation and test sets. From the 690 images, we randomly chose 445 images for training, 112 images for validation, and 133 images for the test set. We cropped all non-square images to square dimensions and resized to 512 x 512. We performed these augmentations:

1. Geometric augmentations: random rotation, horizontal flip, and vertical flip.
2. Color augmentations: To overcome the similarity in red pigmented skin and vessels, we apply histogram stretching to each color channel followed by contrast limited adaptive histogram equalization, normalization, and brightness enhancement.

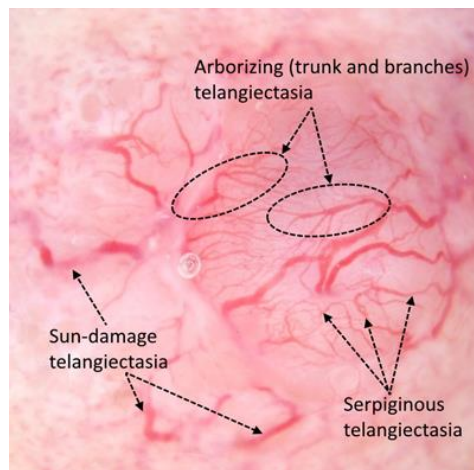


Figure 1. Vessels in BCC. Arborizing and serpiginous telangiectasia vs. non-specific sun-damage telangiectasia

After the augmentations, we used 3114 images for training and 784 images for validation.




Figure 2. Color augmentations

2.3. IMAGE SEGMENTATION

We confined the analysis to the BCC by automatic segmentation using U-Net [18] with details below.

3. DEEP LEARNING NETWORK

3.1. NETWORK ARCHITECTURE

Biomedical image segmentation often employs U-Net [18] due to its ability to perform accurate pixel-based classification. The network includes contractive and expansive paths where the contractive (encoder) path follows architecture similar to a convolutional neural network and the expansive (decoder) path uses transposed 2D convolutional layers. The encoder is downsampled 4 times and the decoder is up-sampled

4 times to restore the high-level semantic feature map produced by the encoder to the original size of the image. (Figure 3).

Each convolutional layer in the encoder part is followed by a maxpool down-sampling operation for the network to encode the input image into feature representations at multiple levels. The decoder includes up-sampling and concatenation, succeeded by convolution operations. The decoder projects the lower dimensional discriminative features learned by the encoder into a higher-resolution space. It upsamples the feature map while simultaneously concatenating it with its higher resolution feature map from the encoder part. The final layer does a 1x1 convolution to map the last feature map to the respective classes.

Since vessels only constitute about 2-10% of the image, it is essential to use a loss function that addresses this severe class imbalance. We use a combination loss function, a weighted sum of Dice loss and binary cross entropy, defined as:

$$DL(y, \hat{y}) = 1 - (2y\hat{y} + 1)/(y + \hat{y} + 1)$$

$$L_{W-bce} = -\frac{1}{N} \sum_i \beta (y - \log(y)) + (1 - \beta)(1 - y) \log(1 - \hat{y})$$

$$m = \alpha L_{W-bce} - (1 - \alpha) DL(y, \hat{y})$$

Where DL is Dice loss and LW-bce is weighted binary cross entropy [19].

We employ U-Net [18] with our modifications for vessel segmentation due to its accuracy in pixel-based classification, as detailed in Section 4.

4. EXPERIMENTS PERFORMED

Our vessel U-Net model uses varied input sizes for each of the RGB image channels of 32, 64, 128, 256 and 512. We use Exponential Linear Units activations instead of the traditional U-Net Relu activations as they tend to converge faster and produce more accurate results [21].

$$F(x) = \begin{cases} z & z > 0 \\ \alpha \cdot (e^z - 1) & z \leq 0 \end{cases}$$

The weights are initialized from truncated normal distribution centered on zero with standard deviation = $\sqrt{2/\text{fan_in}}$ where fan_in is the number of input units in the weight tensor.

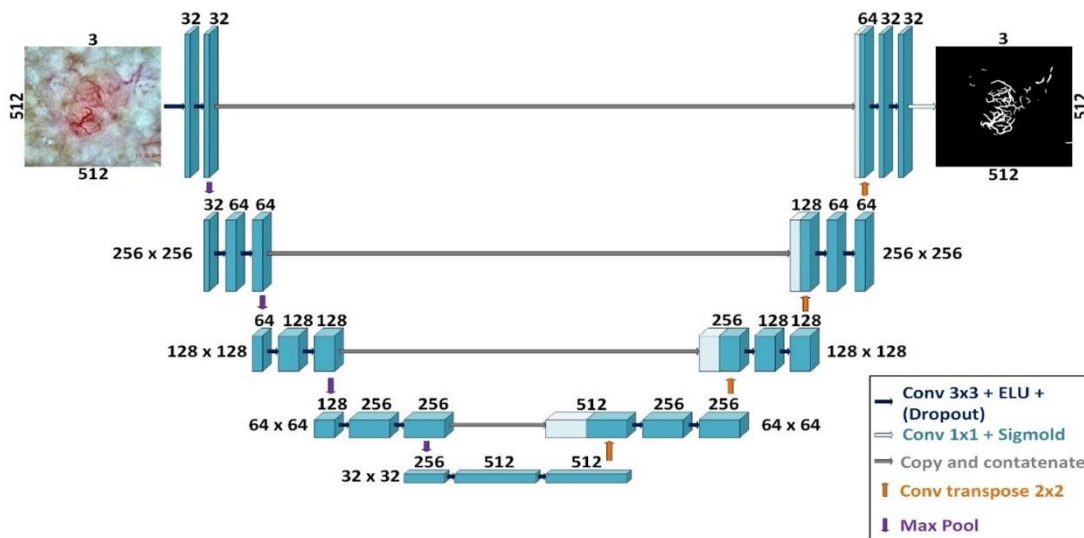


Figure 3. U-Net architecture

To introduce regularization, we use dropout layers with probability 0.1 for both the encoder and decoder. To prevent overfitting, we use early stopping with a patience value of 5 and save the best model. The hyperparameters are listed in Table 1.

Table 1. Different hyperparameters used for training U-Net

Hyperparameters	Value
One weight	0.89
Zero weight	0.11
Learning rate	0.0001
Epochs	20
Metrics	Jaccard loss
Batch Size	8
Alpha for combo loss	0.7

Automatic lesion borders were determined by U-Net trained on the 2594-image ISIC 2018 Task 1 Lesion Segmentation dataset [19]. Images were resized to 320x320 using bilinear interpolation. We randomly split the images into training and validation set of 80:20. Hyperparameters were similar to Table 1 except the model was trained up to 100 epochs with a batch size of 10 and an Adam optimizer and Dice loss function.

5. EXPERIMENTAL RESULTS

We evaluate our model with the Jaccard (intersection over union) metric. This avoids the over-representation of negative pixels in sparse features.

$$\text{Jaccard Index} = \frac{TP}{TP+FP+FN} \text{ (for binary classification)}$$

Our U-Net model achieves a Jaccard score of 37.8% on the test set, which exceeds the mean Jaccard score among our 5 observers who created vessel masks. Results of our model are shown in Figure 4 as green overlays, to compare with the predicted masks, white overlays.




Figure 4. Predicted binary masks and overlays

5.1. VESSEL DETECTION PERFORMANCE FOR DIFFERENT LOSS FUNCTIONS

Table 2. Evaluation metrics for different loss functions with U-Net

Loss Functions	Accuracy	Jaccard	Precision	Recall
Weighed BCN	0.98	0.311	0.351	0.734
Tversky Loss	0.985	0.351	0.442	0.629
Focal Tversky Loss	0.985	0.354	0.435	0.657
Dice Loss	0.989	0.364	0.565	0.506
Combo Loss	0.987	0.376	0.574	0.521

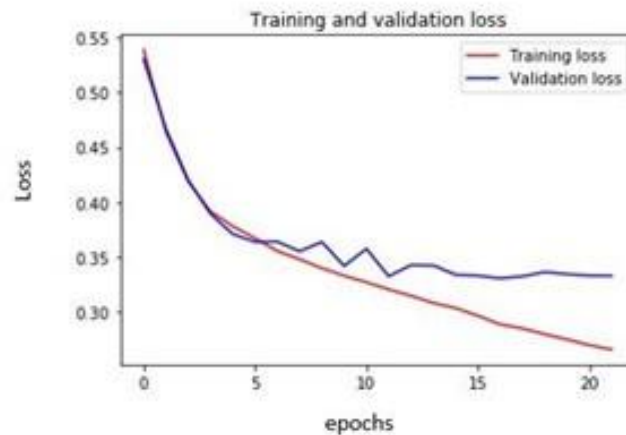


Figure 5. Training and validation loss curves for combo loss

In this study, 20 epochs yield a stable validation loss (Figure 5).

6. DISCUSSION

We used the basic U-Net structure of Ronneberger et al. [18] with modifications as noted in network architecture. We used different loss functions [19] (table 2). The mean accuracy, precision, recall and Jaccard index, precision and recall on the test set for the combination loss function were 0.987, 0.38, 0.574 and 0.521. This compares with mean accuracy, precision and recall for Kharazmi et al. [15] of 0.954, 0.947 and 0.917. The latter two scores are higher than we obtained. However, the two studies are not comparable because we score presence or absence of vessels on a pixel-by-pixel basis and Kharazmi scores presence or absence of a vessel within a patch of 32 x 32 pixels. Additionally, the Kharazmi masks include vascular structures other than vessels (figures 5-7) [15]. Other studies [12][13][14] lack pixel-by-pixel scoring.

To understand our results better, we took samples of 10 random images from each person's mask set and had all 5 observers mark those masks, to create 5 masks for

each of the 50 images. One such example of an image and the corresponding mask done by each person is shown below. We can see from the images that there is some difference in the way each person views and draws the vessels.




Figure 6. Mask annotations for the same image by different team members

For the example images shown in Figure 6, the mean Jaccard calculated for each pair of observers on a pixel-by-pixel basis for is only 0.285, even though they appear rather similar.

Over the entire set the median pairwise Jaccard is 0.271. Some of the possible reasons for low Jaccard values are as follows:

1. Inter-rater variability: The most common variation we observed is the extent of vessel pixel covering, due to inexact fading of the edges of the vessels, and variable covering in the mask. The second most common variation was the varying inclusion of vague vessel structures.

2. Software differences: Two observers used Paint.net and three used Photoshop.
3. Tool variation: For freehand drawing of a mask, variations can also arise from the use of a stylus with Photoshop and a mouse with Paint.net.

Most false positives we observed were the result of blurry or thin vessels missed during mask creation, Figure 7.




Figure 7. Example of disagreement with manual mask. True positives shown by green, false positives shown by blue and false negatives shown by yellow

This research enables optimal detection of a critical dermoscopy structure in early basal cell carcinoma: thready blood vessels called telangiectasia. We accomplish this vessel segmentation task by combining image pre-processing with U-Net using our hyperparameters. The study also compares different loss functions and manages class imbalance by using composite loss functions.

The subjective nature of structure identification by different observers has not received sufficient attention in the literature. Accordingly, we present an analysis of differences in vessel detection by different observers. We show that these differences are not a deterrent to accurate detection of these structures. We are able to achieve deep learning results that are more in agreement with each observer than the observers are with each other. In this study, we establish a path to detection of other cancer-critical signs for earlier cancer detection. In the future, we would like to further explore differences in machine and manual annotation to develop more sophisticated models and different U-Net approaches [20][21].

REFERENCES

1. Rogers HW, Weinstock MA, Feldman SR, Coldiron, BM. Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the US population, 2012. *JAMA Dermatology*. 2012;151(10):1081-1086.
2. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks, *Nature* 2017;542(7639):115-118.
3. Rigel DS, Torres AM, Ely HJ. Imiquimod 5% cream following curettage without electrodesiccation for basal cell carcinoma: preliminary report. *J Drugs Dermatol*.2008;7(1suppl 1):15-16.
4. Marchetti MA, Codella NCF, Dusza SW, Gutman DA, Helba B, Kalloo A, Mishra N, Carrera C, Celebi ME, DeFazio JL, Jaimes N, Marghoob AA, Quigley E, Scope A, Yélamos O, Halpern AC; International Skin Imaging Collaboration. Results of the 2016 International Skin Imaging Collaboration International Symposium on Biomedical Imaging challenge: Comparison of the accuracy of computer algorithms to dermatologists for the diagnosis of melanoma from dermoscopic images. *J Am Acad Dermatol*. 2018;78(2):270-277.

5. Haenssle HA, Fink C, Toberer F, Winkler J, Stolz W, Deinlein T, Hofmann-Wellenhof R, Lallas A, Emmert S, Buhl T, Zutt M, Blum A, Abassi MS, Thomas L, Tromme I, Tschandl P, Enk A, Rosenberger A; Reader Study Level I and Level II Groups. Man against machine reloaded: performance of a market-approved convolutional neural network in classifying a broad spectrum of skin lesions in comparison with 96 dermatologists working under less artificial conditions. 2020; 31(1):137-143.
6. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017;542(7639):115-118
7. T. Majtner, S. Yildirim-Yayilgan, and J. Y. Hardeberg, Combining deep learning and hand-crafted features for skin lesion classification, 6th Int Conf Image Process Theory Tools Appl. (IPTA) Dec 12 2016. Oulu, Finland.
8. Codella NCF, Cai J, Abedini M, Garnavi, R, Halpern A, Smith JR. Deep learning, sparse coding, and SVM for melanoma recognition in dermoscopy images. International Workshop on Machine Learning in Medical Imaging (MLMI) 2015. Oct 5, Munich Germany.
9. Codella NCF, Nguyen Q-T, Pankanti S, Gutman D, Helba B, Halpern A, Smith JF. Deep learning ensembles for melanoma recognition in dermoscopy images. *IBM Res Dev*. 2017;61(4/5):1-5.
10. González-Díaz I, DermaKNet: Incorporating the knowledge of dermatologists to convolutional neural networks for skin lesion diagnosis. *IEEE J Biomed Health Inform*. 2019;23(2):547-559.
11. Hagerty JR, Stanley RJ, Almubarak HA, Lama N, Kasmi R, Guo P, Drugge RJ, Rabinovitz HS, Oliviero M, Stoecker WV. Deep learning and handcrafted method fusion: Higher diagnostic accuracy for melanoma dermoscopy images. *IEEE J Biomed Health Inform*. 2019;23(4):1385-1391.
12. Cheng B, Erdos D, Stanley RJ, Stoecker WV, Calcara, DA, Gomez DD. Automatic detection of basal cell carcinoma using telangiectasia analysis in dermoscopy skin lesion images, *Skin Res Technol*. 2011;17(3):278-87.

13. Cheng B, Stanley RJ, Stoecker WV, Hinton K. Automatic telangiectasia analysis in dermoscopy images using adaptive critic design, *Skin Res Technol.* 2012;18(4):389-96
14. Kharazmi P, AlJasser MI, H. Lui H, Wang ZJ, Lee TK. Automated detection and segmentation of vascular structures of skin lesions wseen in dermoscopy, with an application to basal cell carcinoma classification, *IEEE J Biomed Health Inform.* 2017;21(6):1675-1684
15. Kharazmi P, Zheng J, Lui H, Wang ZJ, Lee TK. A computer-aided decision support system for detection and locatization of cutaneous vasculature in dermoscopy images via deep feature learning., *J Med Syst,* 2018;42(2):33.
16. Karimi D, Dou H, Warfield S K, Gholipour A: Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis, 2019, pp. 1–17, arXiv:1912.02911. [Online]. Available: <http://arxiv.org/abs/1912.02911>
17. Tschandl P, Rosendahl C, Kittler H. , Data Descriptor: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. 2018;Sci Data 5, 180161.
18. O. Ronneberger, P. Fischer, and T. Brox, U-Net: Convolutional networks for biomedical image segmentation. [Online]. Available: <http://lmb.informatik.uni-freiburg.de/>. Accessed 17 11 2020
19. Jadon, Shruti, A survey of loss functions for semantic segmentation, 2020, [Online], Available: <https://arxiv.org/abs/2006.14822> Available: <https://arxiv.org/abs/2006.14822> Accessed 17 11 2021.
20. Codella NCF, Gutman DA, Celebi ME, Helba B. Skin lesion analysis toward melanoma detection 2018: A challenge hosted by the International Skin Imaging Collaboration (ISIC). arXiv:1902.03368 [cs.CV] Available: <https://arxiv.org/abs/1902.03368> 2019. Accessed 17 11 2021.
21. Oskal KRJ, Risdal, M., Janssen EAM, Undersrud ES, Gulsrud TO. . et al. A U-net based approach to epidermal tissue segmentation in whole slide histopathological images. *SN Appl. Sci* 2019; 1:672.

II. BASAL CELL CARCINOMA DIAGNOSIS WITH FUSION OF DEEP LEARNING AND TELANGIECTASIA FEATURES

Akanksha Maurya¹, Member, IEEE, R. Joe Stanley¹, Senior Member, IEEE, Hemanth Y. Aradhyula², Norsang Lama¹, Anand K. Nambisan¹, Gehana Patel³, Daniyal Saaed³, Samantha Swinfard¹, Colin Smith⁴, Sadhika Jagannathan⁵, Jason R. Hagerty⁶, William V. Stoecker⁶

¹Missouri University of Science & Technology, Rolla MO, USA

²Ford Motor Company, MI, USA

³University of Missouri, Columbia MO USA

⁴A.T. Still University of Health Sciences, MO USA

⁵University of Missouri, Kansas City, MO USA

⁶S&A Technology, Rolla, MO, 65401 USA

ABSTRACT

Telangiectasia, narrow blood vessels that typically appear serpiginous or arborizing, are a critical indicator of basal cell carcinoma (BCC), aiding dermatologists in BCC diagnosis. In recent years, deep learning has been used extensively and successfully to diagnose different cancers in dermoscopic images. However, most approaches lack clinical inputs that could aid in higher accuracy and explainability. We demonstrate a novel BCC diagnosis technique by applying ensemble learning on a combination of EfficientNet-B5 based DL features and handcrafted features from semantically segmented telangiectasia masks (U-Net-based). On a hold-out dataset of 395 lesion images (195 BCC and 200 non-BCC) we achieve a binary classification accuracy of 97.2% and an AUC of 0.99. We demonstrate metric improvements in three stages: 1) the addition of handcrafted telangiectasia features to deep learning features, 2) including areas near telangiectasia (surround areas), 3) discarding the noisy lower-importance features. Examining the surround areas of telangiectasia and calculating respective

features, we offer another novel approach to feature finding with weak annotations. Our experimental results show state-of-the-art accuracy and precision in the diagnosis of BCC, compared to three benchmark techniques. Further exploration of deep learning techniques for individual dermoscopy feature detection is warranted.

Index Terms— basal cell carcinoma, deep learning, fusion, telangiectasia, transfer learning, dermoscopy

1. INTRODUCTION

Basal cell carcinoma is one of the two most common types of skin cancer in the USA, with over two million new cases diagnosed yearly. [1] Dermatologists usually diagnose BCC by visual inspection. However, certain benign lesions can be confused with BCC and lead to a unnecessary biopsy. Automating this diagnosis and ensuring early detection will reduce the burden on patients and healthcare professionals and produce more accurate results [2,3].

Deep learning methodologies applied to dermoscopy images have yielded high diagnostic accuracy, now exceeding that of dermatologists [4-6]. Skin cancer diagnosis from images has advanced by implementing fusion ensembles, metadata, and some handcrafted features [7-11].

Telangiectasia or thin arborizing vessels within the lesion is a crucial factor for dermatologists when looking for BCC. Their detection, either by visual inspection or any computational method can provide pathways to make BCC diagnosis more accurate.

Cheng et al. [12] investigated a local pixel color drop technique to identify vessel pixels. Kharazmi et al. [13] applied independent component analysis, k-means clustering,

and shape for detecting vessels and other vascular structures. Kharazmi et al. [14] detected vessel patches by using stacked sparse autoencoders (SSAE) as their deep learning model. Maurya et al. [15] used a U-Net-based deep-learning (DL) model to perform semantic segmentation of these blood vessels. Semantically segmented precise binary masks provide means to effectively quantify the vessel feature information.

There are very few image processing studies that utilize these vascular structures to diagnose BCC. Cheng et al. [16] used an adaptive critic design approach to discriminate vessels from competing structures, enabling BCC classification. Kharazmi et al. [17] utilize features extracted from these vascular structures to classify BCC using a random forest classifier. Kharazmi et al. [18] learned from sparse autoencoders, combined them with patient profile information, and fed them to a Softmax classifier for BCC diagnosis. Recently, Serrano et.al [19] used clustering-based color features and GLCM-based texture features to train VGG16 and MLP models to extract deep learning features that they use to train another MLP model. The final MLP classifies lesions with either presence or absence of one of seven BCC patterns, providing BCC classification with high accuracy.

In this study, we achieve state-of-the-art accuracy in BCC classification by making the following unique contributions: 1. Clinically inspired and explainable BCC diagnosis with deep learning-based telangiectasia mask generation as an intermediate step. 2. An ensemble learning classifier utilizing a hybrid input feature set consisting of object, shape, and color telangiectasia features integrated with deep learning features improving overall BCC classification.

The remainder of the paper is organized as follows. Section 2 explains the image datasets, our proposed methodology, and training details. Section 3 presents and explains our experimental results. Section 4 provides a discussion on our approaches and results. Section 5 gives the conclusion and possible future work.

2. MATERIALS AND METHODS

2.1. IMAGE DATASETS

The skin lesion images used in this study come from 3 datasets: the HAM10000 dataset (ISIC 2018) of Tschandl et al. [20], a publicly available skin lesion dermoscopy dataset containing over 10,000 skin images for seven diagnostic categories, the ISIC 2019 dataset [20-22], and datasets R43 from NIH studies CA153927-01 and CA101639-02A2 [23]. For training the U-Net model, that generates telangiectasia masks, 127 images were selected from the HAM10000 dataset, 90 images from the ISIC 2019 dataset, and 783 images from the NIH study dataset, leading to a total of 1000 BCC images. The ISIC 2019 dataset included a few repeat images, omitted from our BCC dataset. The ground truth binary vessel masks were manually drawn by our team and verified by a dermatologist (WVS). The BCC dataset and the ground truth masks are shown in Figure 1.

For the non-BCC dataset, we selected 1000 images from the HAM10000 dataset from five benign categories: benign keratosis, nevus, actinic keratosis, dermatofibroma, and vascular lesion, shown in Figure 2. The distribution for each of these categories is shown in Figure. 3. The images are 8-bit RGB of size 450x600 from the HAM10000

dataset and 1024x768 from the NIH study dataset. All the images are resized to 448x448 before training.




Figure 1. Left to right, first row: first two images are BCC from the HAM10000 (ISIC2018) dataset; last two images are BCC from the NIH study dataset. Second row presents telangiectasia overlays for the images in the first row

For both U-Net and EfficientNet-based models' training, standard train-test splits of 80-20 were used. The training set was further split 80-20 to create a validation set. For BCC classification, the images were randomly selected, leading to subsets of 1288 for training, 324 for validation, and 395 for testing. For the U-Net model, the training set consisted of 650 images, the validation set comprised 162 images, and the test set included 195 images. These 195 BCC images combined with 200 non-BCC images make up the holdout test set for the BCC classification model.



Figure 2. From left, first row: benign keratosis, nevus, actinic keratosis; from left second row: dermatofibroma and vascular lesion




Figure 3. Number of images from each category used in the non-BCC dataset

2.2. DATA AUGMENTATION

For medical datasets with relatively few examples and a lack of variation, data augmentation helps create more training samples. Augmentation provides deep learning models the ability to generalize and hence provides regularization without overfitting. For both the U-Net and EfficientNet-based BCC classification models, the geometric augmentations include rotation ranging $+30^\circ$ to -30° in reflect mode, not to distort the

vessels, horizontal and vertical flip, width shift with a range of $(-0.2, +0.2)$, height shift with a range of $(-0.2, +0.2)$, and shear with a range $(-0.2, +0.2)$. These geometric augmentations are shown in Figure 4. For the U-Net model, to overcome the similarity in red-pigmented skin and vessels, we perform color augmentations as well. We apply histogram stretching to each color channel followed by contrast-limited adaptive histogram equalization, normalization, and brightness enhancement [15], as shown in Figure 5.



Figure 4. Different geometric augmentations

Due to the narrowness of the vessels, all the vessel masks are dilated with a 3×3 structuring element and closed with a 2×2 structuring element. This dilation prevents the vessels in the masks from being broken when augmented and covers some boundary pixels, as shown in Figure 5. Since the goal is to identify vessels within the lesion, the vessel masks are multiplied with U-Net-generated lesion masks to yield vessels only within the lesion. For both models, the images were square-cropped and then resized to 448×448 .




Figure 5. Left: Color augmentations Right: vessel masks dilation

2.3. PROPOSED METHODOLOGY

Our proposed methodology integrates clinically relevant handcrafted features of telangiectasia with high-level features extracted from a pre-trained deep learning model according to a feature importance score determined by the average Gini impurity decrease, calculated from a random forest structure. We then utilize the higher predictive ability of ensemble learning methods to feed this hybrid feature set to a random forest classifier and create our novel fusion BCC diagnosis technique. Our method includes four main components: 1. Semantically segmenting telangiectasia with a U-Net-based model to yield a binary vessel mask. 2. Applying image processing and statistical methods to calculate descriptive vessel features based on the objects in the vessel mask. 3. Extracting high-level deep learning features from fine-tuning a pre-trained EfficientNet-based model. 4. Calculating the feature importance score for all the features and selecting only the top-most features above a threshold. 5. Classifying the skin lesions into BCC or non-BCC using a random forest classifier trained on the hybrid feature set. Our methodology is illustrated in Figure 6.

In recent years, U-Net [24] based semantic segmentation models have been the go-to for biomedical segmentation. Our vessel detection deep learning model is based on the U-Net model [15]. As vessels occupy only 2-10% of the skin lesion image, a combination loss function addresses this severe class imbalance [19].

Traditional image processing techniques provide several tools to calculate telangiectasia features explicitly. Features are generated using objects in the vessel mask. Table 1 explains the handcrafted features we generated.

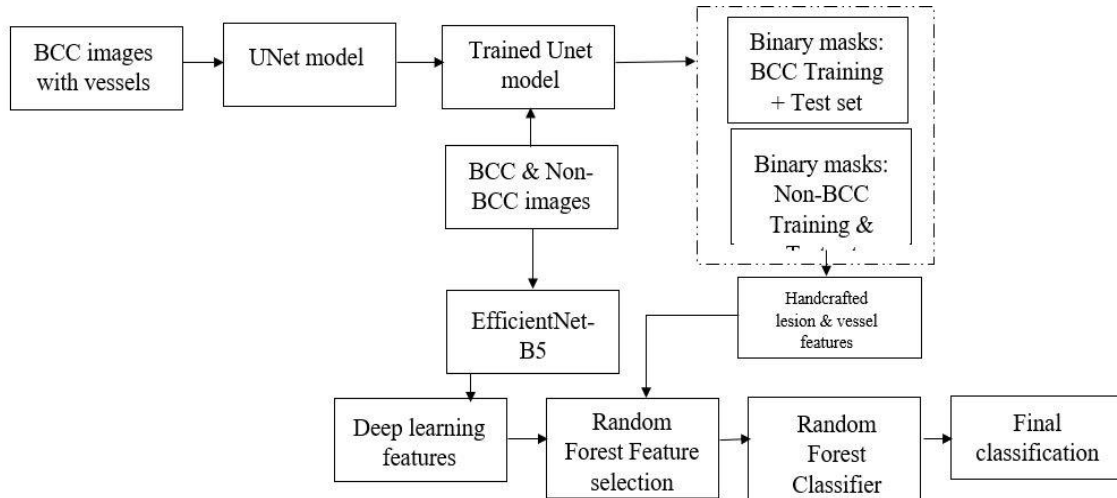


Figure 6. Proposed architecture employing a fusion of deep learning and handcrafted features from vessels for BCC classification

Features 1 to 8 are general vessel descriptors [12]. Features 1-4 represent BCC's narrower, longer, and more numerous vessels. Eccentricity features 5- 8 are calculated to account for straighter BCC vessel structures. From Maurya et al. [15], annotations for telangiectasia suffer from interobserver variability, fine or blurry telangiectasia, and poor contrast in images. There may be missed vessels along the boundary of annotations that

need to be included. Hence, regions surrounding the vessels can also be helpful in differentiating basal cell carcinoma.

To include these probable missed features, we generate a surround mask for the vessel objects. Every object from the vessel mask is extracted and dilated with a disk structuring element of radius 12 (d1) and radius 5 (d2), resulting in two dilation variants of the object. Removing d2 from d1 gives the object surround mask, as shown in Figure 7. For the example images shown, 347 and 240 vessel objects are detected. The figure shows surround masks generated for two such vessel objects. Features 5 to 8 are calculated for the surround masks and make up features 9 to 12. Features 13 to 22 include the number of objects calculated after morphologically eroding the final vessel mask with a circular structuring element of radius from 1 to 10. Features 23 to 32 include the area of objects calculated after morphologically eroding the final vessel mask with a circular structuring element of radius from 1 to 10. Features 33 to 44 are color features for vessels. Features 45 to 56 are color features in the HSV plane and features 57 to 80 are features 33 to 56 (RGB and HSV features) applied to the surround of vessel objects.




Figure 7. Object surround masks: Top row from left: Example lesion image 1 with telangiectasia, its ground truth mask, its surround mask for object 1 (the contiguous connected vessel area), its surround mask for object 2 (the vessel on the bottom right). Bottom row from left: Example lesion image 2 with telangiectasia, its ground truth mask, its surround mask for object 1 (the contiguous connected vessel area), its surround mask for object 2 (the vessel on the bottom right)

Table 1. Handcrafted Feature Description

No.	Abbreviation	Measure	Description	Meaning
1	no_vessels	Number of vessels	Number of vessels in the final vessel mask	BCC has more vessels
2	avg_len	Average length	Average length for all vessels within a lesion	BCC vessels are longer
3	avg_wid	Average width	Average width for all vessels within a lesion	BCC vessels are narrower
4	avg_area	Average area	Average area for all vessels within a lesion	BCC vessels are larger
5	max_ecc	Maximum eccentricity	Maximum ratio of the distance between the foci of the ellipse enclosing the vessels and its major axis length	BCC vessels are straighter
6	min_ecc	Minimum eccentricity	Minimum ratio of the distance between the foci of the ellipse enclosing the vessels and its major axis length	BCC vessels are straighter
7	avg_ecc	Average eccentricity	Average eccentricity of all the vessels per vessel mask	BCC vessels are straighter
8	STD_ecc	Standard deviation of eccentricity	Average standard deviation of eccentricity of all the vessels per vessel mask	BCC vessels are straighter and more uniform
9 to 12	1 to 8 with prefix 'sur'	Surround features	Same as features 5 to 8 for the surrounding objects in the vessel mask	Regions around the vessel may contain distinguish-ing information
13 to 22	no_objN; N: 1 to 10	Eroded vessel objects	Number of objects after the vessel mask is eroded with a disk structure of radius 1 to 10	BCC objects are fewer after given number of erosions
23 to 32	areaN; N: 1 to 10	Eroded vessel area	Area of objects after the vessel mask is eroded with a disk structure of radius 1 to 10	BCC object areas are smaller after given number of erosions
33 to 35	max_R, max_G, max_B	Color features	Maximum red, green, and blue value of every vessel; then averaged over total number of vessels per image	BCC vessels appear darker than lesion
36 to 38	min_R, min_G, min_B	Color features	Minimum red, green, and blue value of every vessel averaged over total number of vessels per image	BCC vessels appear darker than lesion
39 to 41	avg_R, avg_G, avg_B	Color features	Average red, green, and blue value of every vessel averaged over total number of vessels per image	BCC vessels appear darker than lesion
42 to 44	STD_R, STD_G, STD_B	Color features	Standard Deviation of red, green, and blue value of every vessel averaged over total number of vessels per image	BCC vessels appear darker than lesion
45 to 56	33 to 44 with prefix "HSV"	Color features	Features 33 to 44 applied in HSV plane	HSV plane is more robust to lighting and shadow variations
57 to 80	33 to 56 with prefix 'sur'	Surround features	Features 33 to 56 applied on the surround of objects in the vessels	Regions around the vessel may contain distinguish-ing information

To calculate deep learning features, we use a pretrained EfficientNet model. EfficientNets are a series of convolutional neural network models introduced by Tan et al. [25] that uniformly scale all network depth, width, or resolution dimensions by a compounding coefficient. These models achieved state-of-art top-1 accuracy on the ImageNet [26] dataset with fewer parameters. Their primary building block is a mobile inverted bottleneck called MBConv. The family of EfficientNet networks has different numbers of these MBConv blocks. The EfficientNetB5 model consists of two main blocks: MBConv1 and MBConv6. The detailed structures of these blocks are shown in Figure 8.

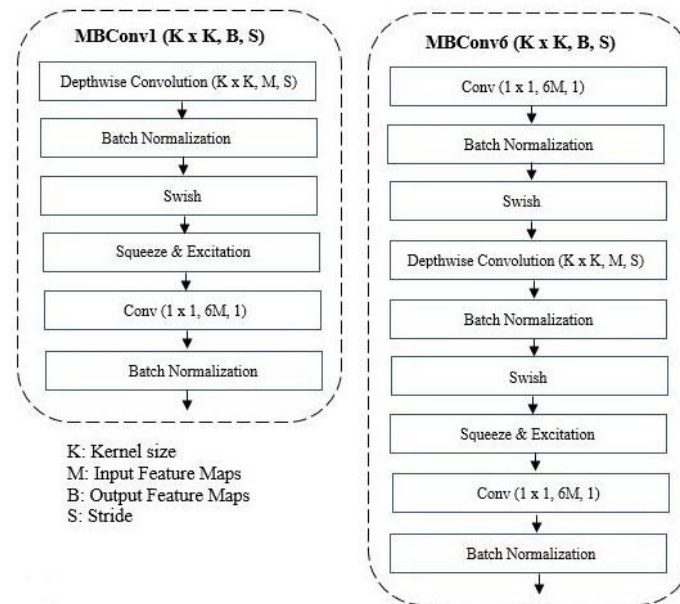


Figure 8. MBConv1 & MBConv6 blocks

The MBConv block or the inverted residual block improved the MobileNet [27] or residual blocks and followed a narrow-> wide->narrow approach, which means that the connections in the MBConv blocks move from one bottleneck to another, using a

residual connection. The basic MBConv block implements the following operations: 1x1 convolution that expands the dimensionality from the narrow channels to wider channels, a 3x3 or 5x5 channel-wise or depth-wise convolution operation to get output features, ultimately followed by another 1x1 convolution that downsamples the number of channels to the initial value. Since this output block and the initial block have the same dimensionality, they are added together.

Both variations of the MBConv blocks contain the Squeeze and Excitation sub-block [25]. The primary purpose of the Squeeze operation is to extract global information from each of the channels of an image. Each block starts with a feature transformation on an image X to get features U , which are then squeezed to a single value using global average pooling [25]. This output is then fed to a fully connected layer followed by a ReLU function to add nonlinearity and reduce complexity. From here, another fully connected layer followed by a sigmoid function performs the excitation operation to get per-channel weights. The final output is achieved by rescaling these feature maps U with these activations. The detailed structure of the Squeeze and Excitation block is shown in Figure 9.

Since the original EfficientNet-B5 model was built for ten class classification, we remove the top layers to add a global average pooling layer, a dropout layer, and a final dense layer for binary classification. For our model, the initial input image size is 448x448x3. Our model contains 14 different stages and is used first for classification, followed by feature extraction with the trained model.

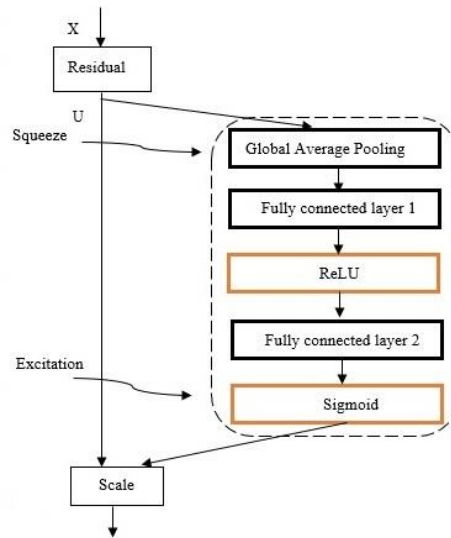


Figure 9. Squeeze & Excitation block

For the classification stage, we start with a 3×3 filter convolution, batch normalization and swish activation function reducing the image dimensions in half from 448 to 224 and increasing the channels from 3 to 48. Hence, the feature map dimensions are $224 \times 224 \times 48$. Stage 2 consists of 3 layers of an MBCConv1 block with a 3×3 filter that maintains the previous stage resolution, but decreases the number of channels, resulting in a feature map of size $224 \times 224 \times 24$. Stages 3 (5 layers), 4 (5 layers) and 5 (7 layers) use 3 MBCConv6 blocks each of kernel size 5×5 continuously reducing the resolution but increasing the feature map size to $28 \times 28 \times 128$ (end of stage 5). Stages 6 (7 layers), 7 (9 layers) and 8 (3 layers) apply 3 more MBCConv6 blocks each, with kernel sizes 3×3 , 5×5 and 3×3 , producing a feature map of size $14 \times 14 \times 2048$. At stage 9, a 1×1 convolution with 2048 filters results in a feature map of size $14 \times 14 \times 2048$. Stages 10 and 11 apply batch normalization and Softmax activation retaining the feature size as the previous layer. Stage 12 uses global average pooling to bring the resolution to 2048, followed by a

dropout (stage 13) and dense layer (stage 14) leading to the final classification. The model is fine-tuned after the 200th layer and the best model is saved. We perform feature extraction at stage 12 after the global average pooling layer, thereby resulting in a 2048 size feature vector for the training, validation, and test sets. The stages, operations, resolutions, channels, and layers are shown in Table 2.

Table 2. Efficientnet-B5 Based DI Model Description

Stage	Operator/Block	Resolution	Channels	Layers
1	Conv 3x3 + BN + Swish	224x224	48	1
2	MBCConv1, k3x3	224x224	24	3
3	MBCConv6, k5x5	112x112	40	5
4	MBCConv6, k5x5	56x56	64	5
5	MBCConv6, k5x5	28x28	128	7
6	MBCConv6, k3x3	28x28	176	7
7	MBCConv6, k5x5	14x14	304	9
8	MBCConv6, k3x3	14x14	512	3
9	Conv 1x1	14x14	2048	1
10	BN	14x14	2048	1
11	Activation	14x14	2048	1
12	Global Average Pooling	2048	1	1
13	Dropout	2048	1	1
14	Dense	1	1	1

2.4. TRAINING DETAILS

All models were built using Keras with a Tensorflow backend in Python 3.7 and trained using a single 32GB Nvidia V100 graphics card. The training and network parameters for the U-Net model generating the vessel masks are taken from [15]. The EfficientNet-B5 model is fine-tuned after the 200th layer. The model is trained for 120 epochs with a learning rate of 0.0001 and a batch size of 20. To prevent overfitting, we use an early stopping criterion with a patience of 5. The loss function used is binary cross entropy with RMSprop optimizer. The dropout rate is 0.2. For the random forest classifier, 1000 estimators are used with the Gini index criterion. The minimum samples

per split are 2 with bootstrapping, and the maximum number of features considered per decision is the square root of the number of features.

3. EXPERIMENTAL RESULTS

We present test results for each of the four stages of the proposed architecture. The results are evaluated on the holdout test set that consists of 195 BCC and 200 non-BCC images. The evaluation metrics used are Accuracy, Sensitivity, Specificity, and Precision which are defined as:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

$$Recall = Sensitivity = \frac{TP}{TP + FN}$$

$$Specificity = \frac{TN}{TN + FP}$$

$$Precision = \frac{TP}{TP + FP}$$

where TP stands for True Positives, TN stands for True Negatives, FP stands for False Positives, and FN stands for False negatives. All variables denote pixel counts.

3.1. VESSEL SEGMENTATION TEST RESULTS FOR BCC AND NON-BCC IMAGES

Figure 10 shows an example of vessel masks generated from the U-Net model for BCC and non-BCC images. From left, columns 1 and 3 show non-BCC and BCC images, respectively, whereas columns 2 and 4 show the predicted masks from the U-Net model. The second image in column 1 has some vessels outside the lesion captured in the predicted masks. Compared to the telangiectasia vessels in the BCC images, these vessels

appear disconnected and thinner. Handcrafted masks can capture such distinctive properties.




Figure 10. From left: Column 1 shows non-BCC images, column 2 shows their predicted masks, column 3 shows BCC images, and column 4 shows predicted BCC masks

3.2. DEEP LEARNING TRAINING RESULTS AFTER FINE-TUNING

We achieved the best deep learning results from the EfficientNet-B5 baseline model. The weights after layer 200 were unfrozen, and the model was fine-tuned to our binary BCC vs. non-BCC classification. The preliminary transfer learning model converged at 100 epochs and took nine more epochs to converge after fine-tuning.



Figure 11. Loss curves before and after fine-tuning the EfficientNet-B5 based model



Figure 12. Accuracy curves before and after fine-tuning the EfficientNet-B5 based model

Figure 11 and 12 show the loss and accuracy curves for the training and validation sets. The green vertical line denotes the point at which fine tuning starts. After a jump in accuracy and a dip in loss, the curves flatten, and the model converges. The validation and test set accuracies were 95.9% and 95.2%, respectively. After the model is trained, the 2048 length feature vector is extracted for the training, validation, and test sets.

3.3. FEATURE IMPORTANCE WITH RANDOM FOREST

Feature importance score is calculated for deep learning and handcrafted features using the random forest classifier. The selection of key features results in models requiring optimal computational complexity while ensuring reduced generalization error due to noise introduced by less important features. Figure 13 and 14 show the selected features with their importance scores.

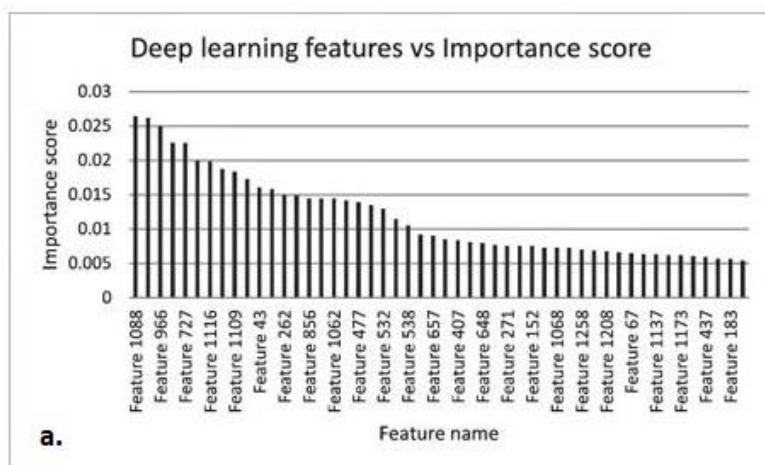


Figure 13. Importance scores for deep learning features

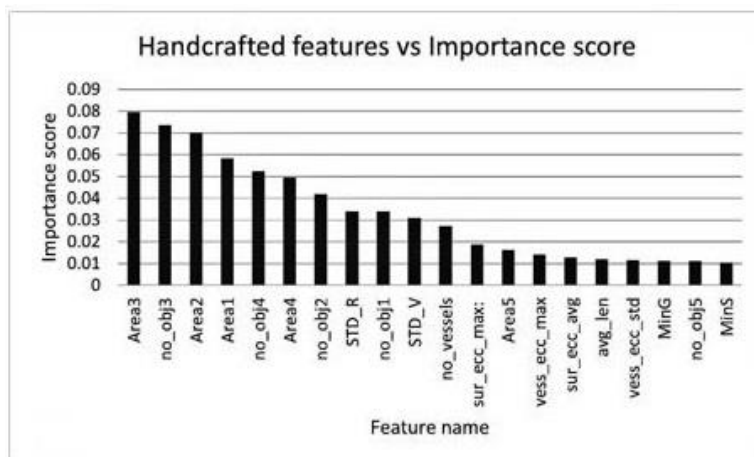


Figure 14. Importance scores for handcrafted features

The top 50 deep learning features with importance score greater than 0.005, and the top 23 handcrafted features with importance score greater than 0.01 are selected. From Figure 14, we observe that the most important handcrafted features generally include area of vessel objects, number of vessel objects, vessel eccentricity features, and color values of the vessel objects and the surrounding area.

3.4. FINAL CLASSIFICATION WITH DEEP LEARNING AND RANDOM FOREST CLASSIFIER

Table 3 shows six different fusion models that we tested. The different models are based on the following two selection criteria:

1. Pretrained model used for extracting deep learning features:
EfficientNetB5, EfficientNetB0, and InceptionV3
2. Feature set size:
 - A. Fusion 2: Select deep learning and handcrafted features with the highest importance score: total 73 features: 50 for deep learning and 23 for handcrafted
 - B. Fusion 1: All deep learning and handcrafted features
 - a. EfficientNet-B5: total 2128 features; 2048 for deep learning and 80 for handcrafted
 - b. EfficientNet-B0: total 1360 features; 1280 for deep learning and 80 for handcrafted
 - c. InceptionV3: total 2128 features; 2048 for deep learning and 80 for handcrafted

We achieve the best scores across all metrics with our Fusion 2 model that uses select critical features extracted from the fine-tuned EfficientNet-B5 model and

handcrafted features extracted from vessel masks that feed a random forest classifier to yield a final classification result. We achieve an accuracy of 0.972, sensitivity of 0.979, specificity of 0.965, and precision of 0.965. We achieve an AUC of 0.995 as shown in Figure 15. Table 4 shows the improvements in metrics as we move from the fine-tuned EfficientNet-B5 model to our fusion model. The accuracy, sensitivity, specificity, and precision improve 1.3%, 3.7%, 1.5% and 1.5%, respectively, suggesting the importance of adding handcrafted features. Omitting the surround features from the list of handcrafted features drops the AUC slightly, from 0.995 to 0.993.

Table 3. Performance Comparison of Different Fusion Models

Model	Feature set size	Acc	Sens	Spec	Prec
EfficientNet-B5-FT-Fusion2	73: 50 EfficientNet-B5-FT + 23 handcrafted	0.972	0.979	0.965	0.965
EfficientNet-B0-FT-Fusion2	73: 50 EfficientNet-B0-FT + 23 handcrafted	0.967	0.979	0.955	0.960
InceptionV3-FT-Fusion2	73: 50 InceptionV3-FT + 23 handcrafted	0.955	0.965	0.95	0.95
EfficientNet-B5-FT-Fusion1	2128: 2048 EfficientNet-B5-FT + 80 handcrafted	0.964	0.967	0.955	0.95
EfficientNet-B0-FT-Fusion1	1360: 1280 EfficientNet-B0-FT + 80 handcrafted	0.934	0.945	0.94	0.933
InceptionV3-FT-Fusion1	2128: 2048 InceptionV3-FT + 80 handcrafted	0.93	0.942	0.934	0.93

Notes for Table: FT: Fine-tuned; Acc: Accuracy; Sens: Sensitivity; Spec: Specificity; Prec: Precision

Table 4. Performance Comparison with Baseline DL Model

Model	Feature set size	Acc	Sens	Spec	Prec
EfficientNet-B5-FT-Fusion2	73: 50 EfficientNet-B5-FT + 23 handcrafted	0.972	0.979	0.965	0.965
EfficientNet-B5-FT	2048 EfficientNet-B5	0.952	0.942	0.950	0.950

Notes for Table: FT: Fine-tuned; Acc: Accuracy; Sens: Sensitivity; Spec: Specificity; Prec: Precision

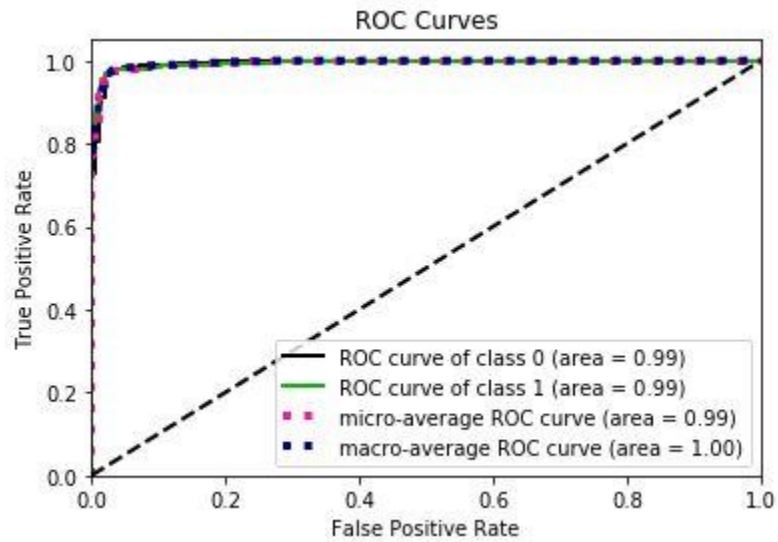


Figure 15. ROC curve for our final fusion classification model

Table 5. Performance Comparison with Other Methods

Manu	Dataset	Feature categories	Final Classifier	Acc	Sens	Prec
Kharazmi et al. 2017	659; 299 BCC and 360 non-BCC	Vascular features	Random Forest	0.965	0.904	0.952
Kharazmi et al. 2018	1199; 599 BCC and 600 non-BCC	Patient profile information & SAE feature learning	Softmax	0.911	0.853	0.877
Serrano et al. 2022	692 BCC and 671 non-BCC	Color and texture features	MLP	0.970	0.993	0.953
Proposed method	2000; 1000 BCC and 1000 non-BCC	EfficientNet-B5 & localized vessel handcrafted color and shape features	Random Forest	0.972	0.979	0.965

Notes for Table: Manu: Manuscript; Acc: Accuracy; Sens: Sensitivity; Spec: Specificity; Prec: Precision

3.5. PERFORMANCE COMPARISON WITH EXISTING METHODS

We compared the performance of our proposed fusion model with 3 other published results [17,18,19]. Table 5 lists the datasets, features, classifiers, and scoring metrics for the models. We achieve better accuracy and precision values with our method

than the existing best values. All the methods listed except ours use some type of color and texture features; only ours semantically segments telangiectasia as an intermediate step.

4. DISCUSSION

A crucial clue for the clinical diagnosis of basal cell carcinoma is the presence of telangiectasia within the lesion. Classical image processing methods to detect these vessels in this study used statistical measures to quantify telangiectasia features. These measures included characteristics such as color values relative to the surrounding lesion area [12] or independent component analysis of melanin and hemoglobin components followed by thresholding and clustering. [18]. From these vessel masks, different color, texture, and shape features are calculated. Our group used a deep learning-based U-Net model to detect these vessel masks with high accuracy vs. ground truth, obtaining a mean Jaccard score within the variation of human observers [15].

Recently, deep learning methods achieved superior results for detecting features such as hair [28] and globules [29]. From these deep learning-generated masks, classical features are calculated with the assumption that if the masks are more accurate, the features will be as well. Developing a single model for diagnosis without extracting individual features from whole images using pre-trained deep learning models also has been used extensively. However, it is impossible to know which features the deep learning model deems more important, contributing to its black box nature. Our structure-based detection model partially remedies this shortcoming by detecting specific features.

Moreover, this report shows improved diagnostic accuracy for BCC vs. non-BCC classification by combining deep learning and classical features with ensemble learning.

Figure 10. displays the advantages of the method presented here. Benign lesions have fewer vessels, and the total vessel area is less, as shown in columns 2 and 4 of Figure 10. The number of vessels, their morphological features and statistical properties show up as the most discriminatory features obtained from these masks, as shown in Figure 14. Eccentricity and color of objects found, all critical features of telangiectasia, are also crucial handcrafted features derived from these masks.

We achieve state-of-the-art results with this approach that are better than deep learning or traditional image processing results, indicating promise for our structure-based detection model. We also achieve clinically explainable results, opening similar pathways to solve other diagnostic challenges. Our results also confirm the superiority of ensemble learning methods for selecting a robust feature combination that improve the model's accuracy. As seen from Table 1, all metrics improve when the features with higher importance scores are used. Another observation concerns the recent study by Serrano et al. [19]. The authors used different BCC features to annotate images with the presence or absence of features. We achieve a similar AUC but slightly better accuracy with our proposed model, using only a single automatically segmented BCC structure: telangiectasia. Due to this added local pixel information, our results achieve state-of-the-art accuracy automatically without observing and annotating every single pattern that may or may not be present. In previous work, we determined significant interobserver variability in vessel annotation [15]. However, DL can learn to detect structures with more consistency than those providing the masks for DL training. Thus, DL appears to be

able to generalize from limited and inexact data sets and can detect vessel-like structures in different kinds of skin lesions, not just BCC. Figure 10 shows that our U-Net-based vessel detection model can identify these structures even in non-BCC images. Once the telangiectasia are detected, distinguishing qualities in BCC vessels (thinness, arborizing) are captured when we calculate the handcrafted features as indicated in Table 1. Using handcrafted features also helps us distinguish between vessels present outside the lesion, as they do not contribute to BCC diagnosis.

To account for the missed vessels due to blurry boundaries [15], our introduction of surround area features also leads to an overall improvement in the AUC value. For segmentation problems, surround area detection by boundary expansion is a novel solution to feature finding that can contribute to better classification.

There are several limitations of this work. The vessel mask marking was supervised by a single dermatologist (WVS). Only one team observer (one of AM, DS, SS, or WVS) annotated each mask. The final sensitivity was less than for the study by Serrano et al., however, the overall accuracy was higher.

5. CONCLUSION AND FUTURE WORK

This study proposes a telangiectasia-based fusion model approach for classifying BCC vs. non-BCC lesion images. To train our vessel identification model’s deep learning (U-Net) arm, we developed telangiectasia masks for 1000 BCC images, available here. No such telangiectasia overlay database for BCC currently exists. Using the results from [15], we calculate the color and texture features from telangiectasia vessel masks and deep learning features learned from the EfficientNet-B5 model to yield a final

classification result. Using a random forest model to combine features of each model provides a framework for fusion models.

Our fusion model outperforms past BCC classification models in precision and accuracy, over a larger dataset than in previous studies, one that is publicly available. Our state-of-the-art accuracy demonstrates the effectiveness of the proposed fusion techniques for this medical dataset. Our results produce more explainable results than whole-image deep learning results as we target clinically observable and relevant telangiectasia features. The current study is the only one, to the best of our knowledge, which uses semantically generated telangiectasia vessel features for BCC diagnosis.

In the future, we would like to continue this research by including more clinical features for our fusion model and employing additional statistical techniques. For medical datasets limited in number, fusion techniques can help establish state-of-the-art diagnostic models.

ACKNOWLEDGEMENT

The contents of this paper are solely the responsibility of the authors and do not necessarily represent the official views of NIH.

REFERENCES

1. H. W. Rogers, M. A. Weinstock, S. R. Feldman, and B. M. Coldiron, "Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the us population, 2012," *JAMA Dermatol*, vol. 151, no. 10, pp. 1081–1086, 2015, doi: 10.1001/jamadermatol.2015.1187.

2. R. L. Siegel, K. D. Miller, H. E. Fuchs, and A. Jemal, "Cancer statistics, 2021," *CA Cancer J Clin*, vol. 71, no. 1, pp. 7–33, 2021.
3. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
4. M. A. Marchetti, N. C. F. Codella, S. W. Dusza, D. A. Gutman, B. Helba, A. Kalloo, N. Mishra, C. Carrera, M. E. Celebi, J. L. DeFazio, N. Jaimes, A. A. Marghoob, E. Quigley, A. Scope, O. Yélamos, A. C. Halpern, & International Skin Imaging Collaboration (2018). "Results of the 2016 International Skin Imaging Collaboration International Symposium on Biomedical Imaging challenge: Comparison of the accuracy of computer algorithms to dermatologists for the diagnosis of melanoma from dermoscopic images." *J Am Acad Dermatol*, vol.78, no.2, pp.270-277, doi: 10.1016/j.jaad.2017.08.016
5. H. A. Haenssle, C. Fink, F. Toberer, J. Winkler, W. Stolz, T. Deinlein, R. Hofmann-Wellenhof, A. Lallas, S. Emmert, T. Buhl, M. Zutt, A. Blum, M. S. Abassi, L. Thomas, I. Tromme, P. Tschandl, A. Enk, A. Rosenberger, & Reader Study Level I and Level II Groups, "Man against machine reloaded: performance of a market-approved convolutional neural network in classifying a broad spectrum of skin lesions in comparison with 96 dermatologists working under less artificial conditions," *Annals of Oncology*, vol. 31, no. 1, pp. 137–143, Jan. 2020, doi: <https://doi.org/10.1016/j.annonc.2019.10.013>.
6. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
7. T. Majtner, S. Yildirim-Yayilgan, and J. Y. Hardeberg, "Combining deep learning and hand-crafted features for skin lesion classification," 2016 *6th International Conference on Image Processing Theory, Tools and Applications*, IPTA 2016, 2017, doi: 10.1109/IPTA.2016.7821017.
8. N. Codella, J. Cai, M. Abedini, R. Garnavi, A. Halpern, and J. R. Smith, "Deep Learning, Sparse Coding, and SVM for Melanoma Recognition in Dermoscopy Images. In L. Zhou, L. Wang, Q. Wang, Y. Shi (eds) *Machine Learning in Medical Imaging*," *MLMI 2015*, pp. 118–126, doi: 10.1007/JRD.2017.2708299.

9. N. C. F. Codella et al., “Deep Learning Ensembles for Melanoma Recognition in Dermoscopy Images,” *IBM J. Res. Dev.*, vol. 61, no. 4–5, pp. 5:1–5:15, Jul. 2017, doi: 10.1147/978-3-319-24888-2_15.
10. González-Díaz, “DermaKNet: Incorporating the Knowledge of Dermatologists to Convolutional Neural Networks for Skin Lesion Diagnosis,” *IEEE J Biomed Health Inform*, vol. 23, no. 2, pp. 547–559, 2019, doi: 10.1109/JBHI.2018.2806962.
11. J. R. Hagerty et al., “Deep Learning and Handcrafted Method Fusion: Higher Diagnostic Accuracy for Melanoma Dermoscopy Images,” *IEEE J Biomed Health Inform*, vol. 23, no. 4, pp. 1385–1391, 2019, doi: 10.1109/JBHI.2019.2891049.
12. Cheng, D. Erdos, R. J. Stanley, W. V. Stoecker, D. A. Calcara, and D. D. Gómez, “Automatic detection of basal cell carcinoma using telangiectasia analysis in dermoscopy skin lesion images,” *Skin Research and Technology*, vol. 17, no. 3, pp. 278–287, Mar. 2011, doi: <https://doi.org/10.1111/j.1600-0846.2010.00494.x>.
13. P. Kharazmi, M. I. AlJasser, H. Lui, Z. J. Wang, and T. K. Lee, “Automated Detection and Segmentation of Vascular Structures of Skin Lesions Seen in Dermoscopy, With an Application to Basal Cell Carcinoma Classification,” *IEEE J Biomed Health Inform*, vol. 21, no. 6, pp. 1675–1684, Nov. 2017, doi: <https://doi.org/10.1109/JBHI.2016.2637342>.
14. P. Kharazmi, J. Zheng, H. Lui, Z. Jane Wang, and T. K. Lee, “A Computer-Aided Decision Support System for Detection and Localization of Cutaneous Vasculature in Dermoscopy Images Via Deep Feature Learning,” *Journal of Medical Systems*, vol. 42, no. 2, p. 33, Jan. 2018, doi: <https://doi.org/10.1007/s10916-017-0885-2>.
15. Maurya et al., “A deep learning approach to detect blood vessels in basal cell carcinoma,” *Skin Research and technology*, vol. 28, no. 4, pp. 571–576, Jul. 2022, doi: <https://doi.org/10.1111/srt.13150>.
16. Cheng, R. J. Stanley, W. V. Stoecker, and K. Hinton, “Automatic telangiectasia analysis in dermoscopy images using adaptive critic design,” *Skin Research and Technology*, vol. 18, no. 4, pp. 389–396, Nov. 2012, doi: <https://doi.org/10.1111/j.1600-0846.2011.00584>.

17. P. Kharazmi, M. I. AlJasser, H. Lui, Z. J. Wang, and T. K. Lee, "Automated Detection and Segmentation of Vascular Structures of Skin Lesions Seen in Dermoscopy, With an Application to Basal Cell Carcinoma Classification," *IEEE journal of Biomedical and Health Informatics*, vol. 21, no. 6, pp. 1675–1684, Nov. 2017, doi: <https://doi.org/10.1109/JBHI.2016.2637342>.
18. P. Kharazmi, S. Kalia, H. Lui, Z. J. Wang, and T. K. Lee, "A feature fusion system for basal cell carcinoma detection through data-driven feature learning and patient profile," *Skin Research and Technology*, vol. 24, no. 2, pp. 256–264, May 2018, doi: <https://doi.org/10.1111/srt.12422>.
19. Serrano, M. Lazo, A. Serrano, T. Toledo-Pastrana, R. Barros-Tornay, and B. Acha, "Clinically Inspired Skin Lesion Classification through the Detection of Dermoscopic Criteria for Basal Cell Carcinoma," *Journal of Imaging*, vol. 8, no. 7, p. 197, Jul. 2022, doi: <https://doi.org/10.3390/jimaging8070197>.
20. P. Tschandl, C. Rosendahl, and H. Kittler, "Data Descriptor: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions Background & Summary," *Nature Publishing Group*, 2018, doi: 10.1038/sdata.2018.161.
21. N. Codella et al., "Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC)." *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Washington, DC, USA, 2018, pp. 168-172, doi: 10.1109/ISBI.2018.8363547.
22. M. Combalia et al., "BCN20000: Dermoscopic Lesions in the Wild," arXiv:1908.02288 [cs, eess], Aug. 2019, Available: <https://arxiv.org/abs/1908.02288>
23. W.V. Stoecker, Kapil Gupta, B. Shrestha, M. Wronkiewicz, R. Chowdhury, R.J. Stanley, J. Xu, R. H Moss, M. E. Celebi, H. S. Rabinovitz, M. Oliviero, J. M. Malters, I. Kolm, "Detection of basal cell carcinoma using color and histogram measures of semitranslucent areas," *Skin Research and Technology*, 2009, vol. 15, no. 3, pp.283-7. doi: 10.1111/j.1600-0846.2009.00354.x.
24. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation." [Online]. Available: <http://lmb.informatik.uni-freiburg.de/>

25. M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning*, 2019, pp. 6105–6114. doi: 10.48550/arXiv.1905.11946
26. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
27. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., & Chen, L. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4510-4520.
28. N. Lama et al., "ChimeraNet: U-Net for Hair Detection in Dermoscopic Skin Lesion Images," *Journal of Digital Imaging*, Nov. 2022, doi: <https://doi.org/10.1007/s10278-022-00740-6>.
29. K. Nambisan et al., "Deep learning-based dot and globule segmentation with pixel and blob-based metrics for evaluation," *Intelligent Systems with Applications*, vol. 16, p. 200126, Nov. 2022, doi: <https://doi.org/10.1016/j.iswa.2022.200126>.

III. HYBRID TOPOLOGICAL DATA ANALYSIS AND DEEP LEARNING FOR BASAL CELL CARCINOMA DIAGNOSIS

Akanksha Maurya¹, Member, IEEE, R. Joe Stanley¹, Senior Member, IEEE, Norsang Lama¹, Anand K. Nambisan¹, Gehana Patel², Daniyal Saeed², Samantha Swinfard¹, Colin Smith³, Sadhika Jagannathan⁴, Jason R. Hagerty⁵, William V. Stoecker⁵

¹Missouri University of Science & Technology, Rolla MO, USA

²University of Missouri, Columbia MO USA

³A.T. Still University of Health Sciences, MO USA

⁴University of Missouri, Kansas City, MO USA

⁵S&A Technology, Rolla, MO, 65401 USA

ABSTRACT

A critical clinical indicator for basal cell carcinoma (BCC) is the presence of telangiectasia (narrow blood vessels) within the skin lesions. Many skin cancer imaging processes today exploit deep learning (DL) models for diagnosis, segmentation of features and feature analysis. Hence, integration of deep learning models with telangiectasia features could help improve BCC diagnosis. To extend automated diagnosis, recent computational intelligence research has explored the field of Topological Data Analysis (TDA). Persistent homology is a TDA method to identify and quantify topological features of a data set, such as clusters, voids, and tunnels. TDA quantifies these features. In this study we exploit all three aspects namely, telangiectasia, deep learning and TDA to achieve our final BCC classification model. First, we exploit persistent homology-based statistics to implement a color and telangiectasia driven BCC classification model. Second, we use fine-tuning with EfficientNet-B5 based model to achieve robust deep learning features. Finally, we combine the first two models to build our final hybrid TDA-DL model which achieves state of the art accuracy of 97.4% and

AUC of 0.995 on a holdout test of 395 skin lesions for BCC diagnosis. Our hybrid diagnostic model shows that telangiectasia features improve BCC diagnosis and TDA improves DL performance.

Keywords— basal cell carcinoma, TDA, persistent homology, deep learning, fusion, telangiectasia, transfer learning, dermoscopy

1. INTRODUCTION

Over two million cases of basal cell carcinoma (BCC) diagnosed yearly in the US [1]. The initial diagnosis of BCC includes a visual inspection by a dermatologist or mid-level practitioner (nurse practitioner or physician assistant), often with a dermatoscope. If the diagnosis is unclear or if confirmation is needed, an invasive procedure such as a biopsy is performed. Recent research has aimed to improve diagnostic accuracy and minimize the number of biopsies through automatic image processing. In some cases, deep learning (DL) methods applied in dermoscopy have outperformed dermatologists [2-6]. Skin cancer diagnosis from images has advanced by implementing DL and, in some cases, fusion ensembles employing DL, metadata, and handcrafted features [7-12].

Telangiectasia or thin narrow blood vessels within the skin lesions are a critical clinical indicator of BCC. [12-14] Studies have detected these blood vessels through various handcrafted pixel-based techniques or patch-based techniques. Cheng et al. [13] investigated a local pixel color drop technique to identify vessel pixels. Kharazmi et al. [14] applied independent component analysis, k-means clustering, and shape for detecting vessels and other vascular structures. Kharazmi et al. [15] detected vessel patches by using a stacked sparse autoencoder (SSAE) as their DL model. Maurya et. al

[16] employed DL to segment these vessels semantically, a dermoscopic feature-driven approach also used by Nambisan et al to detect dots and globules [17].

Cheng et al. [18] used an adaptive critic design approach in the past to detect and use these vessels for BCC classification. Kharazmi et al. [14] used a random forest-based classifier to diagnose BCC with color and texture features. Kharazmi et al. [19] used a combination of SSAE and patient meta data for BCC diagnosis. Serrano et.al [20] used clustering-based color features and GLCM-based texture features to train VGG16 and MLP models for DL-based BCC classification.

Topology is a branch of mathematics concerned with the properties of geometric objects that are preserved when the object is stretched, bent, or otherwise deformed. Topological Data Analysis (TDA) is an area of mathematics and data analysis that uses tools from topology to study the shape of data. It is a relatively newer research field that is now increasingly used for image classification, feature extraction and image analysis [21-25]. The main idea behind TDA is that the shape of the “point cloud” or clusters of data points can reveal important data properties that may not be immediately apparent from other types of analysis. For example, TDA can be used to identify clusters or groups of data points, detect patterns or trends in the data, and to extract features or characteristics that persist along multiple higher dimensional scales. Hu et.al [24] used TDA based methods for skin lesion segmentation and classification. Bendich et al [25] employed TDA based persistence diagrams to find metadata correlations to the brain artery trees, establishing a correlation between age and brain artery tree topology

This study explores TDA’s ability to extract features from telangiectasia and color-spaces to improve EfficientNet-B5 pre-trained model performance.

This study contributes the following to the existing literature on automatic BCC diagnosis:

- Integrating a clinically observable physical feature: telangiectasia with a DL-TDA model to improve performance.
- Demonstrating an alternative and less computationally intensive TDA model for medical image diagnosis.

The remainder of the paper is organized as follows. Section 2 explains the image datasets, our proposed methodology, and training details. Section 3 presents and explains our experimental results. Section 4 provides a discussion on our approaches and results. Section 5 gives the conclusion and possible future work.

2. MATERIALS AND METHODS

2.1 IMAGE DATASETS

This study uses BCC and benign dermoscopic skin cancer images derived from 3 datasets: the HAM10000 dataset (ISIC 2018) of Tschandl et al. [26], a publicly available skin lesion dermoscopy dataset containing over 10,000 skin images for seven diagnostic categories, the ISIC 2019 dataset [26-28], and datasets R43 from NIH studies CA153927-01 and CA101639-02A2 [29]. The U-Net model is trained on a total of 1000 BCC images, 127 of which come from the HAM10000 dataset, 90 from ISIC 2019 and 783 from the NIH study dataset. We use 1000 non-BCC images from the HAM10000 dataset for our DL-BCC diagnostic model.

The 1000 non-BCC lesions along with their distribution in the dataset are:

- Benign Keratosis: 400

- Nevus: 400
- Actinic Keratosis: 67
- Dermatofibroma: 67
- Vascular Lesion: 66

The 1000 BCC images in the dataset are the same as the U-Net model. All the images are 8-bit RGB of size 450x600 from the HAM10000 dataset and 1024x768 from the NIH study dataset. Example images of these skin lesions are shown in Figure 1.



Figure 1. From left: actinic keratosis, benign keratosis, dermatofibroma, basal cell carcinoma, nevus and vascular lesion

2.2. PRE-PROCESSING

All the images are square cropped centering on the lesion area and resized to 448x448. For the U-Net model, there is an extra step where the images are processed with histogram stretching, contrast limited adaptive histogram equalization (CLAHE), normalization, and brightness enhancement (to make vessels brighter and distinguishable) [16]. The ground truth vessel masks are dilated with a 3x3 structuring element and closed with a 2x2 structuring element. We perform geometric augmentations: rotation of $+30^\circ$ to -30° in reflect mode (to preserve vessel continuity), horizontal and vertical flip, width shift with a range of $(-0.2, +0.2)$, height shift with a range of $(-0.2, +0.2)$, and shear with a range $(-0.2, +0.2)$. Figure 2 shows the steps of pre-processing.




Figure 2. Pre-processing flowchart

2.3. PROPOSED METHODOLOGY

Figure 3 presents our five-component hybrid TDA and deep learning model pipeline investigated in this study, including: 1. A U-Net model that semantically segments telangiectasia in both BCC and non-BCC skin lesion images. 2. A TDA based framework that calculates Persistence Statistics (PS) from vessel masks and -selected color spaces of the images. 3. A deep learning (DL) model based on EfficientNet-B5 for feature extraction. 4. Generating class probabilities from two separate random forest classifiers, one for DL and one for TDA. 5. Majority voting between the probabilities generated from DL and TDA features to yield a final classification. The following subsections describe each of the components in detail.

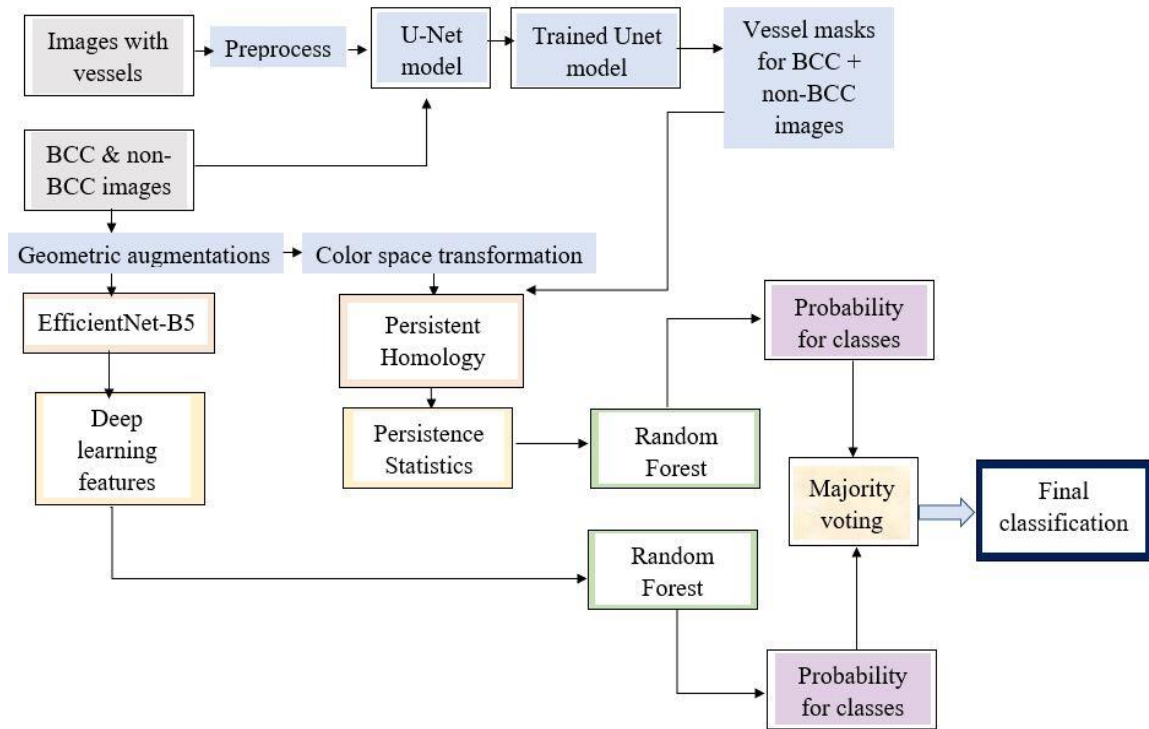


Figure 3. Pipeline investigated employing a hybrid TDA-DL method for BCC classification

2.3.1. U-Net for Telangiectasia Masks. U-Net based segmentation models are widely used in medical image segmentation [30]. The U-Net model and its hyperparameters are taken from [16] for this study. The model produces binary vessel masks for BCC and non-BCC lesion images. A TDA framework used these binary vessel masks (as explained in detail in the subsequent sections) to generate topological features.

2.3.2. Topological Data Analysis (TDA) and Persistent Homology. Topological Data Analysis (TDA) applies the concepts and methods of topology for the analysis and visualization of complex data. Persistent homology (PH), a statistical tool of TDA, can detect topological features of the data that persist over larger scales and long intervals of time. PH accounts for the topological features i.e., connected components in dimension 0,

loops in dimension 1 and voids in dimension 2 by creating persistence diagrams [21-25].

The persistent homology algorithm follows the steps shown in Figure 4.

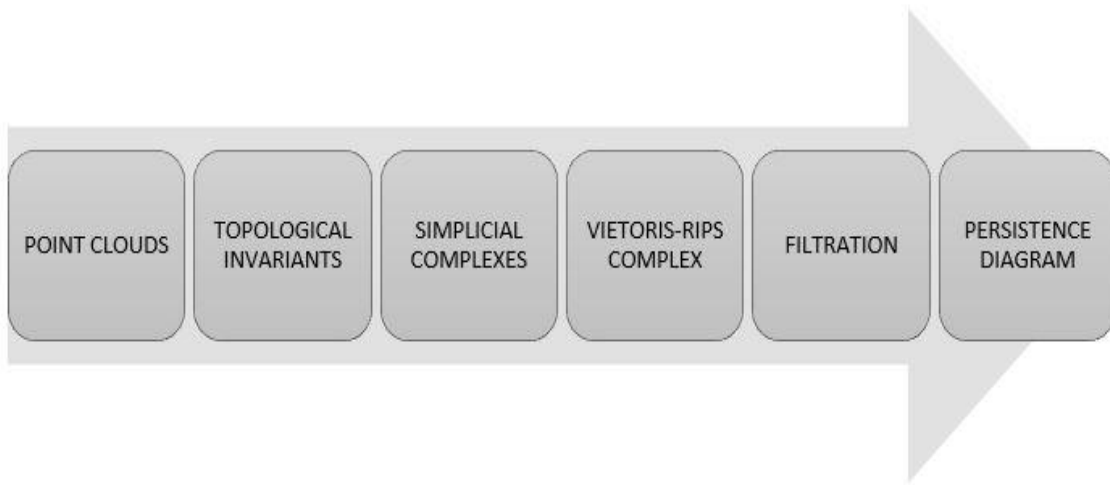


Figure 4. Flow of the persistent homology algorithm to generate persistence diagrams

The main steps of the process are described below:

A. Point clouds: Point clouds are collections of data points that reflect the geometry and spatial relationships of a real-world object or environment in a high-dimensional space, most often a three-dimensional (3D) space. The x , y , and z coordinates of each point in a point cloud, as well as any other qualities like color or intensity, are used to identify each point's location in the space. Point clouds serve as the pixel intensity values for our problem.

B. Topological invariants: Topological invariants are topological space-related mathematical numbers or qualities that are true even if the space is altered in some way. These invariants offer a mechanism to categorize and separate various topological spaces according to their inherent characteristics. Betti numbers are also a type of topological

invariant representing the total number of holes in a space of various sizes. Higher Betti numbers count higher-dimensional holes. The first Betti number counts independent loops, while the zeroth Betti number counts connected components. Topological invariants are frequently generated from algebraic structures known as homology groups or cohomology groups in the setting of algebraic topology, which examines the algebraic features of topological spaces. Chains or cochains, formal combinations of simplices or cells in a topological space, are used to create these groups.

C. Simplicial complex: In the study of combinatorial topology and geometry, a simplex is a fundamental geometric object. It is an extension of the 2-dimensional simplex idea of a triangle to higher dimensions. The convex hull of $(n+1)$ affinely independent points in Euclidean space is formally referred to as an n -dimensional simplex. A simplex is the "simplest" conceivable polytope in n -dimensional space, equivalently, it is a geometric object [21-25]. Here are a few instances:

- A vertex of a zero-dimensional simplex is represented by a single point
- A line segment joining two points is referred to as a one-dimensional simplex
- A triangle having three vertices and three edges is referred to as a two-dimensional simplex
- A tetrahedron with four vertices, six edges, and four triangular faces is referred to as a three-dimensional simplex.




Figure 5. Simplicial Complex containing a 3D simplex (tetrahedron) and a 2D simplex (triangle)

Topological invariants can be computed from the simplicial complex by counting the number of simplexes of different dimensions that make up the complex. Let V be a set of vertices. A subset S of V is called a simplex of dimension n if it contains $n+1$ elements that are affinely independent, meaning that the points do not lie in a lower-dimensional hyperplane. The elements of S are called the vertices of the simplex. A simplicial complex K is a collection of simplexes in V that satisfies the following conditions [21-25]:

- Any face of a simplex in K is also in K , meaning that if S is a simplex in K , then every subset of S that is a simplex is also in K .
- The intersection of any two simplexes in K is either empty or a face of both.

Figure 5 shows a simplicial complex that includes a tetrahedron and a triangle.

D. Vietoris-Rips Complex: We utilize the Vietoris-Rips Complex to produce Simplicial Complexes from our image dataset. It is built by joining together spatial pairs of points that are relatively close to one another, and then joining together higher-

dimensional simplexes while considering the connectivity of the lower-dimensional simplexes.

Formally, for a set of points P of dimension d , where P is a subset of \mathbb{R}^d , then the Vietoris-Rips (VR) complex $V_\epsilon(P)$ at scale ϵ (the VR complex over the point cloud P with parameter ϵ) is defined as [21-25]:

$$V_\epsilon(P) = \{\sigma \subseteq P \mid d(u,v) \leq \epsilon, \forall u \neq v \in \sigma\}$$

Hence, for a set of data points in P , we include a simplex σ (a subset of P) if the points in σ are all within a distance of ϵ from each other. As a result, we obtain a collection of subsets of P that are all simplices, or a simplicial complex of P .

E. Filtration: By varying the values of ϵ to different levels, we discover what appears to produce a significant VR complex. If ϵ is set too small, the complex might just include the initial point cloud or a sparse number of edges connecting the points. On the other hand, the point cloud will just merge into one enormous ultra-dimensional simplex if ϵ is set too large. In order to truly find patterns in a simplicial complex, we must repeatedly change the parameter ϵ (and generate new complexes) from 0 to a maximum that yields a single huge simplex. Then the diagram illustrates what topological features are created and destroyed as ϵ keeps rising. We assume features that persist over a long period of time are significant and vice versa. This process is called filtration.

F. Persistence Diagrams: A persistence diagram is a graphical representation of this process, which consists of a collection of points in a two-dimensional plane. Each point in the diagram represents a topological feature and its corresponding lifespan or persistence, defined as the difference between the scale at which the feature was born and the scale at which it died out. The diagram's horizontal axis represents the birth values of

the topological features, while the vertical axis represents their death values. The diagonal line in the diagram represents features with the same birth and death values and is called the diagonal or the "line of equality".

0D persistent homology and 1D persistent homology refer to the analysis of topological features in different dimensions using the persistent homology framework. 0D persistent homology focuses on analyzing connected components or clusters in a data set. It captures the evolution of these connected components as a parameter, typically related to distance or scale, varies. By systematically increasing or decreasing the parameter, 0D persistent homology tracks the birth and death of connected components. In 0D persistent homology, the filtration complex is constructed by associating each data point in the set with a 0-dimensional simplex. Initially, each data point is a separate connected component. As the parameter increases or decreases, connected components may merge or disappear, resulting in changes in the topology of the data set. The persistence intervals, or barcode intervals, represent the lifespan of the connected components, indicating when they are born and when they die.

1D persistent homology focuses on analyzing loops or cycles in a data set. It captures the evolution of these loops as the filtration parameter varies. By systematically changing the parameter, 1D persistent homology tracks the birth and death of loops. In 1D persistent homology, the filtration complex is constructed by considering both the data points and the edges connecting them. Initially, each data point is a 0-dimensional simplex, and each edge is a 1-dimensional simplex. As the parameter increases or decreases, edges may form loops or cycles, merge with existing loops, or disappear. The

persistence intervals represent the lifespan of the loops, indicating when they are born and when they die.

From this point onwards, we refer to the persistence diagrams corresponding to 0D and 1D persistent homology as P_0 and P_1 .

One approach to understanding this filtration process involves creating a sequence of growing spheres centered on each point and connecting those with overlapping spheres with edges or triangles. Figure 5 illustrates this process.

A. We start with a collection of data points (point clouds) in 2D space. At this point the value of ϵ or the radius of the spheres is 0. Hence the connected components are born at $x=0$. Since there has been no death or “overlap”, there is no corresponding y value.

B. As the concentric spheres around the datapoints increase in size/radii (ϵ increases), the first connected components die or overlap, giving us the first death. Hence we see the first birth-death pair point on the corresponding persistence diagram with birth at $x=0$ and death at $y>0$, with x and y both simply corresponding to the radius ϵ of the spheres.

C. At this stage, with the radius or ϵ increasing, more deaths or overlaps happen leading to more deaths and larger values of y , but we also see the emergence of a loop, hence a birth value for 1D homology. This loop finally disappears in the second substage. Hence, we arrive at a birth and death value for x and y , both greater than zero and accounted for by the orange point in the corresponding persistence diagram.

Therefore, the persistence diagram provides a global summary of the topological features of a dataset, capturing both their presence and their persistence over different scales.

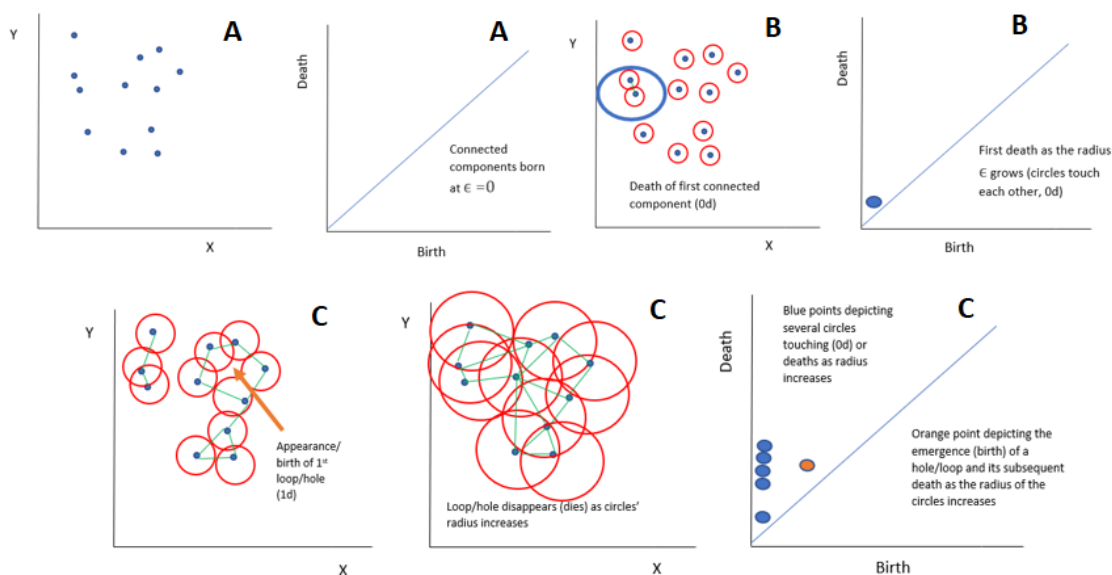


Figure 6. Persistent homology filtration process leading to formation of birth-death pairs in the persistence diagram

For image classification, we use one channel at a time from a 3-channel color space; for example: red color plane from the RGB color space, grayscale image or binary image. PH is used for image analysis by treating image pixels as point clouds, where point clouds are a collection of data points in a high dimensional space. the shape of the point cloud can reveal important data properties that can be used to identify patterns in images, such as textures or shapes, and to measure the similarity between different images. Figure 6 shows persistence diagrams P_0 and P_1 for a BCC and non-BCC image for the red color channel from the RGB color space. We can notice even by visual observation that the birth-death pairs for both images seem distinguishable.

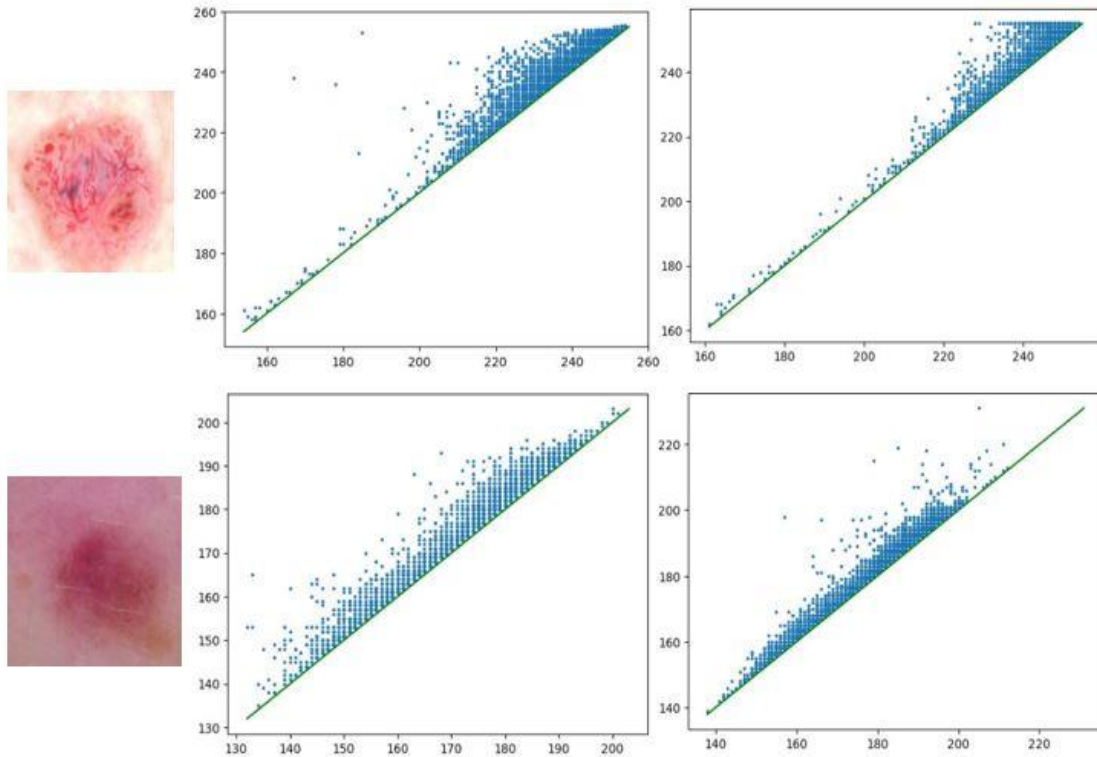


Figure 7. Top row, from left: BCC skin lesion; Its corresponding P0 persistence diagram and P1 Persistence diagram for the red plane; Bottom row, from left: non-BCC (actinic keratosis) skin lesion; Its corresponding P0 persistence diagram and P1 Persistence diagram for the red plane

2.3.3. Persistence Statistics for Topological Features. In the context of a digital image, a point cloud is a set of points in a high-dimensional space representing each pixel's position and color information. Each point in the point cloud corresponds to a single pixel in the image, and its position in the space is determined by its x and y coordinates, while additional dimensions or attributes can represent its color. For our dataset, we treat each channel of a 3-channel image (example RGB) as a grayscale image with pixel intensity values ranging from 0 to 255 [24]. This forms the initial point cloud for the subsequent persistent homology process. We extract five channels from 3

different color spaces namely, R, G, B from the RGB color space, V from the HSV color space and Z from the XYZ color space as these channels performed the best. We also include the predicted telangiectasia mask for the images as another binary image for feature calculation through persistence diagrams. For all the channels we generate both P0 and P1. The process of calculating persistence statistics from P0 is shown in Figure 8.




Figure 8. Generation of persistence statistics-based features

As shown in Figure 7, the persistence diagram P , contains collections of pairs of points that represent the birth and death values of topological features. Our persistence statistics include three quantities that summarize this information in persistent diagrams: total persistence, mid-life coordinates and normalized lifespan [24]. If birth is denoted by b and death is denoted by d , $d-b$ is the lifespan of the topological feature. It represents the

length of time that the corresponding feature persisted during the filtration process. The total persistence is then defined as the sum of the persistence values over all points in the diagram. Mathematically, this can be expressed as:

$$L_i = \sum_{(b,d) \in P_i} d - b$$

where $i = 0,1$ corresponding to P_0 and P_1 . The total persistence provides a global measure of the complexity or richness of a dataset's topological structure, by considering all the topological features and their persistence over different scales. Another statistic is midlife coordinates, expressed mathematically as:

$$M_i = (b + d)/2$$

The third measure is normalized lifespan. It measures the relative persistence or robustness of topological features in a dataset, considering their lifespans and the overall complexity of the persistence diagram. We calculate the normalized lifespan p_i for each point in the diagram as its persistence divided by the total persistence:

$$p_i = (d-b)/L_i$$

The normalized lifespan p_i is a measure of the relative persistence or robustness of a topological feature in comparison to the overall complexity of the persistence diagram [22]. It indicates the proportion of the total persistence contributed by the corresponding feature and provides insight into how long-lasting and persistent the feature is. M_i and p_i are empirical distributions [24] and we apply standard statistical measurements on these distributions, to calculate our feature vector. Table 1 shows the 32 topological features we calculated for each image in our dataset.

Table 1. Persistence Statistics calculated for our methodology

Feature number	Feature name (i =0,1)	Description
1 to 4	Pi_meanmidlife Pi_mean_normalized_lifespan	Means of M_i and p_i
5 to 8	Pi_std_midlife Pi_std_normalized_lifespan,	Standard deviation of M_i and p_i
9 to 12	Pi_skew_midlife Pi_skew_normalized_lifespan	Skewness of M_i and p_i
13 to 16	Pi_kurtosis_midlife Pi_kurtosis_normalized_lifespan	Kurtosis of M_i and p_i
17 to 20	Pi_median_midlife Pi_median_normalized_lifespan	Medians of M_i and p_i
21 to 24	Pi_perc25_midlife Pi_perc25_normalized_lifespan	25 th percentile of M_i and p_i
25 to 28	Pi_perc75_midlife Pi_perc75_normalized_lifespan	75 th percentile of M_i and p_i
29 to 32	Pi_interquart_midlife Pi_interquart_normalized_lifespan	Interquartile ranges of M_i and p_i

2.3.4. Transfer Learning with EfficientNet-B5 for Feature Extraction. A

family of convolutional neural network (CNN) models called EfficientNet has attained cutting-edge performance on a variety of computer vision applications while retaining a manageable number of parameters [31-33]. By properly scaling the network in several dimensions, EfficientNet's major goal is to address the trade-off between model size and accuracy. In the past, scaling a model meant individually expanding its depth, width, or resolution. EfficientNet, on the other hand, suggests a compound scaling technique that takes depth, width, and resolution into account all at once. The compound scaling technique also ensures that the model can be efficiently fine-tuned on smaller datasets without overfitting. Lama et al [34-35] successfully employed EfficientNet-based DL models for lesion segmentation and hair detection. Hence, we chose an EfficientNet based model, specifically EfficientNet-B5 for extracting our deep learning features for the classification model.

We remove the top layers from the original EfficientNet-B5 model because it was designed for 10 class classification instead of binary classification. We replace them with a global average pooling layer, a dropout layer, and a final dense layer. The initial input image size for our model is $448 \times 448 \times 3$. Our model, which has 14 phases, is first utilized for classification, then feature extraction using the trained model. We begin with a 3×3 filter convolution, batch normalization, and swish activation function for the classification stage, cutting the image dimensions in half from 448 to 224 and raising the number of channels from 3 to 48. As a result, the feature map's measurements are $224 \times 224 \times 48$. Stage 2 is composed of three layers of an MBConv1 block with a 3×3 filter, which reduces the number of channels while maintaining the resolution of Stage 1 to produce a feature map with dimensions of $224 \times 224 \times 24$. Stages 3 (five layers), 4 (five layers), and 5 (seven layers) employ three MBConv6 blocks, each with a kernel size of 5×5 , to gradually decrease the resolution while expanding the size of the feature map to $28 \times 28 \times 128$ (the stage's finish). Stages 6 (7 layers), 7 (9 layers), and 8 (3 layers) each apply three more MBConv6 blocks with kernel sizes of 3×3 , 5×5 , and 3×3 to create a feature map with a final dimension of $14 \times 14 \times 512$. A feature map with the dimensions $14 \times 14 \times 2048$ is produced at stage 9 using a 1×1 convolution with 2048 filters. Stages 10 and 11 maintain the feature size from the preceding layer while applying batch normalization and Softmax activation. Stage 12 uses global average pooling to increase the resolution to 2048, followed by stages 13 and 14 leading to the final classification: a dropout and dense layer. After the 200th layer, the model is fine-tuned, and the best model is saved. After the global average pooling layer, at stage 12, we carry out feature

extraction, producing a 2048-dimensional feature vector for the training, validation, and test sets. Table 2 displays the phases, procedures, resolutions, channels, and layers.

Table 2. Our EfficientNet-B5 based deep learning model

Stage	Operator	Resolution	Channels	Layers
1	Conv 3x3 + BN + Swish	224x224	48	1
2	MBCConv1, k3x3	224x224	24	3
3	MBCConv6, k5x5	112x112	40	5
4	MBCConv6, k5x5	56x56	64	5
5	MBCConv6, k5x5	28x28	128	7
6	MBCConv6, k3x3	28x28	176	7
7	MBCConv6, k5x5	14x14	304	9
8	MBCConv6, k3x3	14x14	512	3
9	Conv 1x1	14x14	2048	1
10	BN	14x14	2048	1
11	Activation	14x14	2048	1
12	Global Average Pooling	2048	1	1
13	Dropout	2048	1	1
14	Dense	1	1	1

2.3.5. Class Probabilities with Majority Voting. The 2048-dimensional feature vector from the EfficientNet-B5 model and the 192-dimensional TDA-PS feature vector are both used as inputs for two different random forest classifiers. These random forest ensemble learners generate probabilities for each class (BCC and non-BCC) resulting in 4 probability values:

- DL_prob_1: probability of a lesion being BCC based on DL features
- DL_prob_0: probability of a lesion being non-BCC based on DL features
- TDA_prob_1: probability of a lesion being BCC based on TDA features
- TDA_prob_0: probability of a lesion being non-BCC based on TDA

features

For each image, the probability for each class is calculated and are compared and the class with the highest probability is chosen as the final class.

2.4. TRAINING DETAILS

Both deep learning models: U-Net and EfficientNet-B5 were built using Keras with a with a Tensorflow backend in Python 3.7 and trained using a single 32GB Nvidia V100 graphics card. Hyperparameters for the U-Net model are the same as for Maurya et al [16]. The hyperparameters for the EfficientNet-B5 model are listed in Table 3. For the random forest classifier, 1000 estimators are used with the Gini index criterion. The minimum samples per split are 2 with bootstrapping.

Table 3. Hyperparameters for the EfficientNet-B5 model

Hyperparameter	Values
Fine tuning layer	200
Epochs	120
Learning rate	0.0001
Batch size	20
Loss function	Binary cross-entropy
Optimizer	Adam
Early stopping criteria	Validation loss
Patience	5
Dropout rate	0.2

3. EXPERIMENTAL RESULTS

In this section we discuss the results of our experiments. All results listed were evaluated on the holdout BCC vs non-BCC test set of 395 skin lesion images (195 BCC and 200 no-BCC). The evaluation metrics used are Accuracy, Sensitivity, Specificity, and Precision which are defined as:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN}$$

$$Recall = Sensitivity = \frac{TP}{TP + FN}$$

$$Specificity = \frac{TN}{TN + FP}$$

$$Precision = \frac{TP}{TP + FP}$$

where TP stands for True Positives, TN stands for True Negatives, FP stands for False Positives, and FN stands for False negatives. All variables denote pixel counts.

3.1. U-NET TELANGIECTASIA SEGMENTATION RESULTS

Figure 6 shows example non-BCC and BCC image with their corresponding predicted vessel masks. We can see that the U-net model can segment vessels in both types of lesions even though the vessels are distinguishable. As seen in the next section persistence statistics exploit this discriminative feature and improve classification.




Figure 9. From left: non-BCC image with vessel mask prediction, BCC image with vessel mask prediction

3.2. BCC DIAGNOSIS WITH PERSISTENCE STATISTICS AND RANDOM FOREST

We present classification results using persistence statistics calculated from the R, G, B, V and Z planes and the binary vessel masks. First, Table 4 shows a subset of the persistence statistics calculated from binary vessel masks for 10 BCC image examples.

Table 4. Some example PS features calculated from vessel masks

Image	P0meanmidlifeVes	P0mean_normalized_lifespanVes	P0std_midlifeVes	P0skew_midlifeVes	P0kurtosis_midlifeVes
1	217.496	0.008	86.0857	-2.149	2.6620
2	224.4636	0.0181	78.3486	-2.577	4.8211
3	212.6923	0.0042	90.8462	-1.908	1.6717
4	202.7542	0.0055	99.7026	-1.548	0.4013
5	201.4303	0.0081	100.057	-1.511	0.3109
6	219.0749	0.0032	84.6016	-2.206	2.8897
7	209.8139	0.0037	93.5049	-1.792	1.2348
8	203.8088	0.0036	97.6522	-1.578	0.5371
9	216.4583	0.0009	84.5038	-2.060	2.4011
10	218.7165	0.0044	83.7334	-2.191	2.8978

Table 5 shows the classification results of a random forest classifier trained on PS features derived from different subsets of color spaces. As we can see, there is a 2.3% jump in accuracy, 3.8% jump in sensitivity, 0.5% jump in specificity and 5% jump in precision after adding vessel features, hence signaling the importance of telangiectasia in improving BCC diagnosis.

Table 5. Random Forest classification with PS for R, G, B, V, Z and vessels

Model	Feature set size	Accuracy	Sensitivity	Specificity	Precision
R, G, B	32x3 = 96	0.900	0.875	0.928	0.875
R, G, B, V	32x4 = 128	0.916	0.889	0.934	0.889
R, G, B, V, Z	32x5 = 160	0.920	0.900	0.945	0.900
R, G, B, V, Z, Vessels	32x6 =192	0.943	0.938	0.950	0.950

To ensure features were not redundant we, run a feature importance test with a random forest classifier with our final feature model (R, G, B, V, Z, vessels) and test metrics by taking subsets of the persistence statistics calculated. Table 6 shows that metrics improve considerably after continuously adding PS features. We observe that all 192 features are needed for high diagnostic accuracy.

Table 6. Metric improvements with Subsets of PS Features

Feature set size	Accuracy	Sensitivity	Specificity	Precision
First 50	0.867	0.855	0.866	0.867
First 90	0.885	0.867	0.889	0.883
First 130	0.902	0.890	0.913	0.900
First 170	0.911	0.905	0.925	0.910
All 192	0.943	0.938	0.950	0.950

3.3. DEEP LEARNING RESULTS WITH EFFICIENTNET-B5

Table 7 lists the three best models we trained for fine tuning on the BCC vs non-BCC dataset. We achieve the best results with the EfficientNet-B5 model and hence chose it for feature extraction for the hybrid model. Table 5 and 7 show that the TDA based random forest model performs better than EfficientNet-B0 and InceptionNetV3 but

slightly worse than EfficientNet-B5. Figure 10 shows the loss and accuracy plots for the EfficientNet-B5 model before and after fine-tuning.

Table 7. Performance comparison of different deep learning models

Model	Feature set size	Accuracy	Sensitivity	Specificity	Precision
InceptionV3-FT	2048	0.920	0.910	0.942	0.934
EfficientNet-B0-FT	1280	0.936	0.925	0.947	0.920
EfficientNet-B5-FT	2048	0.959	0.942	0.950	0.950




Figure 10. Training and validation accuracy and loss curves for fine tuning EfficientNet-B5 on BCC vs non-BCC dataset

3.4. BCC VS NON-BCC WITH HYBRID TDA-DL MODEL

The results of our final hybrid TDA-DL model are shown in Table 8. Adding the TDA -based persistence statistics (PS) features improve the deep learning results in two folds. We can see that without adding the PS features from vessels, the color channel-based PS features still improve the DL results as accuracy is increased by 0.6%, sensitivity by 2.1%, specificity by 2% and precision by 2.1%. However, when we consider the full PS features, including the vessel features, the hybrid model's accuracy

rises by 1.5%, sensitivity jumps by 3%, specificity jumps by 2.8% and precision jumps by 2.9%.

Table 8. Performance comparison of DL model, PS model and hybrid TDA-DL model

Model	Feature set size	Accuracy	Sensitivity	Specificity	Precision
DL-EfficientNet-B5	2048	0.959	0.942	0.950	0.950
TDA (PS based)	192	0.943	0.938	0.950	0.950
DL-TDA Hybrid without vessels	2208	0.965	0.963	0.970	0.971
DL-TDA Hybrid with vessels	2240	0.974	0.972	0.978	0.979

3.5 PERFORMANCE COMPARISON WITH EXISTING METHODS

Table 9 compares the performance of our TDA-DL hybrid model with other published studies on the automation of BCC diagnosis [14] [19] [20]. Kharazmi et al [14] used vascular features from vessels whereas in [19], they used patient meta-data along with DL-based auto-encoder features. Serrano et. al [20] used annotated features accounting for the presence of several clinical biomarkers. Our hybrid model achieves higher accuracy and precision overall and produces segmentation telangiectasia as a sub-step. Ours is the only study (to the best of our knowledge) exploring TDA approaches and focusing on improvement due to clinical features.

Table 9. Performance comparison with other studies

Manuscript	Dataset	Feature categories	Accuracy	Sensitivity	Precision
Kharazmi et al 2017	659; 299 BCC and 360 non-BCC	Vascular features	0.965	0.904	0.952
Kharazmi et al 2018	1199; 599 BCC and 600 non-BCC	Patient profile information & SAE feature learning	0.911	0.853	0.877
Serrano et al 2022	692 BCC and 671 non-BCC	Color and texture features	0.970	0.993	0.953
DL-TDA Hybrid	2000; 1000 BCC and 1000 non-BCC	EfficientNet-B5 & localized vessel and global TDA features	0.977	0.977	0.979

4. DISCUSSION

The inclusion of biomarker-driven features for automation of cancer diagnosis is a rapidly growing field. The automation of telangiectasia detection is an important step in the diagnosis of BCC. Studies on this task include ones based on traditional rule-based image processing techniques such as color drop vessel detection [13] and independent component analysis of melanin and hemoglobin components followed by thresholding and clustering. [19]. Deep learning [16] performed this task at a pixel level by a U-Net segmentation model that obtains a Jaccard score within the variation of human observers. In this study we demonstrate the significance of segmenting telangiectasia and adding features derived from them to improve BCC diagnosis.

TDA has been used extensively and successfully in many applications in medical image analysis [24], biology [36] neurology [25]. TDA can be applied for data with limited or noisy information, since it can work with incomplete or partial data as it can capture the multiscale structure of the data. Our initial random forest classification model based solely on persistence statistics derived from the red, green blue, V channel of HSV color space and Z channel of XYZ color space cannot outperform deep learning models. However, after the adding persistence statistics (PS) features derived from telangiectasia masks, we observe that the random forest classifier outperforms Inception-V3 and EfficientNet-B0 models, indicating the importance of this clinical feature in diagnosis. This observation is also significant as deep learning models learn the abstract data with the help of ground truth labels provided to them whereas TDA-based methods perform feature extraction without ground truth labels i.e., unsupervised learning. The PS-based TDA model accuracy result is slightly lower than that of the EfficientNet-B5.

TDA features can predict the BCC class more accurately for some test cases missed by deep learning. In our final hybrid DL-TDA model, we take advantage of majority voting of the DL and TDA probabilities. The accuracy improves by almost 2% on the holdout test set.

The computational cost of calculating the persistence statistics features is significantly lower than for the deep learning features, i.e., they can be calculated without a high-performance GPU. Handcrafted features have also been used for improving diagnosis with fusion [7-12] [18] but usually feature calculation is feature- and problem-dependent. With TDA analysis, we can bypass those limitations.

Even with the recent improvements in the automation of skin cancer diagnosis, we acknowledge that raising the sensitivity and specificity of these models is an ongoing challenge. We recommend the proposed methodology for further exploration on other medical challenges and datasets.

5. CONCLUSION

This study proposes a deep learning and TDA hybrid approach for classifying BCC vs. non-BCC dermoscopic lesion images. We exploit color space information to calculate persistence homology topological features for our skin lesion images and also include topological features from a clinical biomarker for BCC, telangiectasia. For our deep learning model, we choose a state of the art pretrained model: EfficientNet-B5. Combining the DL and TDA features, our hybrid DL-TDA model outperforms EfficientNet-B5 as well as other convolution neural network-based pretrained models. We achieve state-of-the-art accuracy and precision, over a larger dataset, publicly

available at [37] than in previous studies. With the inclusion of the telangiectasia features and the subsequent improvements in the final classification result we also demonstrate a clinically explainable aspect of the diagnosis that can be extended to other biomarkers.

REFERENCES

1. H. W. Rogers, M. A. Weinstock, S. R. Feldman, and B. M. Coldiron, “Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the us population, 2012,” *JAMA Dermatol*, vol. 151, no. 10, pp. 1081–1086, 2015, doi: 10.1001/jamadermatol.2015.1187.
2. R. L. Siegel, K. D. Miller, H. E. Fuchs, and A. Jemal, “Cancer statistics, 2021,” *CA Cancer J Clin*, vol. 71, no. 1, pp. 7–33, 2021.
3. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
4. M. A. Marchetti, N. C. F. Codella, S. W. Dusza, D. A. Gutman, B. Helba, A. Kalloo, N. Mishra, C. Carrera, M. E. Celebi, J. L. DeFazio, N. Jaimes, A. A. Marghoob, E. Quigley, A. Scope, O. Yélamos, A. C. Halpern, & International Skin Imaging Collaboration (2018). “Results of the 2016 International Skin Imaging Collaboration International Symposium on Biomedical Imaging challenge: Comparison of the accuracy of computer algorithms to dermatologists for the diagnosis of melanoma from dermoscopic images.” *J Am Acad Dermatol*, vol.78, no.2, pp.270-277, doi: 10.1016/j.jaad.2017.08.016
5. H. A. Haenssle, C. Fink, F. Toberer, J. Winkler, W. Stolz, T. Deinlein, R. Hofmann-Wellenhof, A. Lallas, S. Emmert, T. Buhl, M. Zutt, A. Blum, M. S. Abassi, L. Thomas, I. Tromme, P. Tschandl, A. Enk, A. Rosenberger, & Reader Study Level I and Level II Groups, “Man against machine reloaded: performance of a market-approved convolutional neural network in classifying a broad spectrum of skin lesions in comparison with 96 dermatologists working under less artificial conditions,” *Annals of Oncology*, vol. 31, no. 1, pp. 137–143, Jan. 2020, doi: <https://doi.org/10.1016/j.annonc.2019.10.013>.

6. Esteva *et al.*, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
7. T. Majtner, S. Yildirim-Yayilgan, and J. Y. Hardeberg, “Combining deep learning and hand-crafted features for skin lesion classification,” 2016 *6th International Conference on Image Processing Theory, Tools and Applications, IPTA 2016*, 2017, doi: 10.1109/IPTA.2016.7821017.
8. N. Codella, J. Cai, M. Abedini, R. Garnavi, A. Halpern, and J. R. Smith, “Deep Learning, Sparse Coding, and SVM for Melanoma Recognition in Dermoscopy Images. In L. Zhou, L. Wang, Q. Wang, Y. Shi (eds) *Machine Learning in Medical Imaging*,” *MLMI 2015*, pp. 118–126, doi: 10.1007/JRD.2017.2708299.
9. N. C. F. Codella *et al.*, “Deep Learning Ensembles for Melanoma Recognition in Dermoscopy Images,” *IBM J. Res. Dev.*, vol. 61, no. 4–5, pp. 5:1–5:15, Jul. 2017, doi: 10.1147/978-3-319-24888-2_15.
10. González-Díaz, “DermaKNet: Incorporating the Knowledge of Dermatologists to Convolutional Neural Networks for Skin Lesion Diagnosis,” *IEEE J Biomed Health Inform*, vol. 23, no. 2, pp. 547–559, 2019, doi: 10.1109/JBHI.2018.2806962.
11. J. R. Hagerty *et al.*, “Deep Learning and Handcrafted Method Fusion: Higher Diagnostic Accuracy for Melanoma Dermoscopy Images,” *IEEE J Biomed Health Inform*, vol. 23, no. 4, pp. 1385–1391, 2019, doi: 10.1109/JBHI.2019.2891049.
12. Nambisan, A., Maurya, A., Lama, N., Phan, T., Patel, G., Miller, K., Lama, B., Hagerty, J., Stanley, R., & Stoecker, W. V. “Improving Automatic Melanoma Diagnosis Using Deep Learning-Based Segmentation of Irregular Networks”. *Cancers* (2023), 15(4), 1259. <https://doi.org/10.3390/cancers15041259>
13. Cheng, D. Erdos, R. J. Stanley, W. V. Stoecker, D. A. Calcara, and D. D. Gómez, “Automatic detection of basal cell carcinoma using telangiectasia analysis in dermoscopy skin lesion images,” *Skin Research and Technology*, vol. 17, no. 3, pp. 278–287, Mar. 2011, doi: <https://doi.org/10.1111/j.1600-0846.2010.00494.x>

14. P. Kharazmi, M. I. AlJasser, H. Lui, Z. J. Wang, and T. K. Lee, “Automated Detection and Segmentation of Vascular Structures of Skin Lesions Seen in Dermoscopy, With an Application to Basal Cell Carcinoma Classification,” *IEEE J Biomed Health Inform*, vol. 21, no. 6, pp. 1675–1684, Nov. 2017, doi: <https://doi.org/10.1109/JBHI.2016.2637342>.
15. P. Kharazmi, J. Zheng, H. Lui, Z. Jane Wang, and T. K. Lee, “A Computer-Aided Decision Support System for Detection and Localization of Cutaneous Vasculature in Dermoscopy Images Via Deep Feature Learning,” *Journal of Medical Systems*, vol. 42, no. 2, p. 33, Jan. 2018, doi: <https://doi.org/10.1007/s10916-017-0885-2>.
16. Maurya et al., “A deep learning approach to detect blood vessels in basal cell carcinoma,” *Skin Research and technology*, vol. 28, no. 4, pp. 571–576, Jul. 2022, doi: <https://doi.org/10.1111/srt.13150>.
17. Nambisan, A.K.; Lama, N.; Phan, T.; Swinfard, S.; Lama, B.; Smith, C.; Rajeh, A.; Patel, G.; Hagerty, J.; Stoecker, W.V.; et al. Deep Learning-Based Dot and Globule Segmentation with Pixel and Blob-Based Metrics for Evaluation. *Intell. Syst. Appl.* 2022, 16, 200126.
18. Cheng, R. J. Stanley, W. V. Stoecker, and K. Hinton, “Automatic telangiectasia analysis in dermoscopy images using adaptive critic design,” *Skin Research and Technology*, vol. 18, no. 4, pp. 389–396, Nov. 2012, doi: <https://doi.org/10.1111/j.1600-0846.2011.00584.x>.
19. P. Kharazmi, S. Kalia, H. Lui, Z. J. Wang, and T. K. Lee, “A feature fusion system for basal cell carcinoma detection through data-driven feature learning and patient profile,” *Skin Research and Technology*, vol. 24, no. 2, pp. 256–264, May 2018, doi: <https://doi.org/10.1111/srt.12422>.
20. Serrano, M. Lazo, A. Serrano, T. Toledo-Pastrana, R. Barros-Tornay, and B. Acha, “Clinically Inspired Skin Lesion Classification through the Detection of Dermoscopic Criteria for Basal Cell Carcinoma,” *Journal of Imaging*, vol. 8, no. 7, p. 197, Jul. 2022, doi: <https://doi.org/10.3390/jimaging8070197>.
21. Otter, N., Porter, M. A., Tillmann, U., Grindrod, P., & Harrington, H. A. (2017, August 9). A roadmap for the computation of persistent homology - EPJ Data Science. SpringerOpen. <https://doi.org/10.1140/epjds/s13688-017-0109-5>

22. Carlsson, G. (2014). Topological pattern recognition for point cloud data. *Acta Numerica*, 23, 289-368. doi:10.1017/S0962492914000051
23. E., L., & Z. (2002, November 1). Topological Persistence and Simplification - Discrete & Computational Geometry. SpringerLink. <https://doi.org/10.1007/s00454-002-2885-2>
24. Hu, C. S., Lawson, A., Chen, J. S., Chung, Y. M., Smyth, C., & Yang, S. M. (2021, November 17). TopoResNet: A Hybrid Deep Learning Architecture and Its Application to Skin Lesion Classification. MDPI. <https://doi.org/10.3390/math9222924>
25. Bendich, P., Marron, J. S., Miller, E., Pieloch, A., & Skwerer, S. (2016, March 25). Persistent Homology Analysis of Brain Artery Trees. PubMed Central (PMC). <https://doi.org/10.1214/15-AOAS886>
26. P. Tschandl, C. Rosendahl, and H. Kittler, "Data Descriptor: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions Background & Summary," *Nature Publishing Group*, 2018, doi: 10.1038/sdata.2018.161.
27. N. Codella et al., "Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC)." *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Washington, DC, USA, 2018, pp. 168-172, doi: 10.1109/ISBI.2018.8363547.
28. M. Combalia et al., "BCN20000: Dermoscopic Lesions in the Wild," arXiv:1908.02288 [cs, eess], Aug. 2019, Available: <https://arxiv.org/abs/1908.02288>
29. W.V. Stoecker, Kapil Gupta, B. Shrestha, M. Wronkiewicz, R. Chowdhury, R.J. Stanley, J. Xu, R. H Moss, M. E. Celebi, H. S. Rabinovitz, M. Oliviero, J. M. Malters, I. Kolm, "Detection of basal cell carcinoma using color and histogram measures of semitranslucent areas," *Skin Research and Technology*, 2009, vol. 15, no. 3, pp.283-7. doi: 10.1111/j.1600-0846.2009.00354.x.
30. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation." [Online]. Available: <http://lmb.informatik.uni-freiburg.de/>

31. M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in *International Conference on Machine Learning*, 2019, pp. 6105–6114. doi: 10.48550/arXiv.1905.11946
32. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
33. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., & Chen, L. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4510-4520.
34. Lama, N., J., Hagerty, Nambisan, A., Stanley, R. J., & Stoecker, W. V. "Skin Lesion Segmentation in Dermoscopic Images with Noisy Data". *Journal of Digital Imaging* (2023). <https://doi.org/10.1007/s10278-023-00819-8>
35. Lama, N., Kasmi, R., Hagerty, J., Stanley, J., Young, R., Miinch, J., Nepal, J., Nambisan, A., Stoecker, W. V. "ChimeraNet: U-Net for Hair Detection in Dermoscopic Skin Lesion Images". *Journal of Digital Imaging* (2022). <https://doi.org/10.1007/s10278-022-00740-6>
36. Anand, D.V., Meng, Z., Xia, K. et al. Weighted persistent homology for osmolyte molecular aggregation and hydrogen-bonding network analysis. *Sci Rep* 10, 9685 (2020). <https://doi.org/10.1038/s41598-020-66710-6>
37. Akanksha Maurya, Ronald J Stanley, Hemanth Y Aradhyula, Norsang Lama, Anand K Nambisan, Gehana Patel, Daniyal Saaed, Samantha Swinfard, Colin Smith, Sadhika Jagannathan, Jason Hagerty, & William V Stoecker. (2023). Basal cell carcinoma diagnosis with fusion of deep learning and telangiectasia features [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.7709824>

SECTION

2. SUMMARY AND CONCLUSIONS

This dissertation proposes data fusion techniques used together with deep learning methodologies for skin cancer image analysis. With the first work, the automation of a clinical biomarker for basal cell carcinoma (BCC), i.e., telangiectasia is performed using deep learning-based U-Net model. A unique database of 1000 telangiectasia vessel masks was created as part of this study. The second work utilizes these masks to extract image processing-based features that ultimately help improve a deep learning model for BCC diagnosis. Through this fusion technique, state-of-the-art metrics in BCC are achieved thus establishing the importance of fusion methods as well as introducing more explainability in the framework. The third study also extends data fusion work by exploiting topological features present in the skin lesion images as well as the telangiectasia vessels. This study also improves deep learning results for BCC diagnosis and produces similar results as the previous study. Topological techniques are computationally less extensive and less complex than image processing and deep learning methods and hence provide a pathway for improvements especially for smaller datasets.

To conclude, deep learning has proven to surpass many state-of-the art results in previous and ongoing research but data fusion methods with different computational techniques, can help improve deep learning results further.

BIBLIOGRAPHY

1. R. L. Siegel, K. D. Miller, H. E. Fuchs, and A. Jemal, "Cancer statistics, 2021," *CA Cancer J Clin*, vol. 71, no. 1, pp. 7–33, 2021.
2. Rogers HW, Weinstock MA, Feldman SR, Coldiron, BM. Incidence estimate of nonmelanoma skin cancer (keratinocyte carcinomas) in the US population, 2012. *JAMA Dermatology*. 2012;151(10):1081-1086.
3. Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks, *Nature* 2017;542(7639):115-118.
4. Rigel DS, Torres AM, Ely HJ. Imiquimod 5% cream following curettage without electrodesiccation for basal cell carcinoma: preliminary report. *J Drugs Dermatol*.2008;7(1suppl 1):15-16.
5. Marchetti MA, Codella NCF, Dusza SW, Gutman DA, Helba B, Kalloo A, Mishra N, Carrera C, Celebi ME, DeFazio JL, Jaimes N, Marghoob AA, Quigley E, Scope A, Yélamos O, Halpern AC; International Skin Imaging Collaboration. Results of the 2016 International Skin Imaging Collaboration International Symposium on Biomedical Imaging challenge: Comparison of the accuracy of computer algorithms to dermatologists for the diagnosis of melanoma from dermoscopic images. *J Am Acad Dermatol*. 2018;78(2):270-277.
6. Haenssle HA, Fink C, Toberer F, Winkler J, Stolz W, Deinlein T, Hofmann-Wellenhof R, Lallas A, Emmert S, Buhl T, Zutt M, Blum A, Abassi MS, Thomas L, Tromme I, Tschandl P, Enk A, Rosenberger A; Reader Study Level I and Level II Groups. Man against machine reloaded: performance of a market-approved convolutional neural network in classifying a broad spectrum of skin lesions in comparison with 96 dermatologists working under less artificial conditions. 2020; 31(1):137-143.
7. T. Majtner, S. Yildirim-Yayilgan, and J. Y. Hardeberg, Combining deep learning and hand-crafted features for skin lesion classification, 6th Int Conf Image Process Theory Tools Appl. (IPTA) Dec 12 2016. Oulu, Finland.

8. Codella NCF, Cai J, Abedini M, Garnavi, R, Halpern A, Smith JR. Deep learning, sparse coding, and SVM for melanoma recognition in dermoscopy images. International Workshop on Machine Learning in Medical Imaging (MLMI) 2015. Oct 5, Munich Germany.
9. Codella NCF, Nguyen Q-T, Pankanti S, Gutman D, Helba B, Halpern A, Smith JF. Deep learning ensembles for melanoma recognition in dermoscopy images. IBM J Res Dev. 2017;61(4/5):1-5.
10. González-Díaz I, DermaKNet: Incorporating the knowledge of dermatologists to convolutional neural networks for skin lesion diagnosis. IEEE J Biomed Health Inform. 2019;23(2):547-559.
11. Hagerty JR, Stanley RJ, Almubarak HA, Lama N, Kasmi R, Guo P, Drugge RJ, Rabinovitz HS, Oliviero M, Stoecker WV. Deep learning and handcrafted method fusion: Higher diagnostic accuracy for melanoma dermoscopy images. IEEE J Biomed Health Inform. 2019;23(4):1385-1391
12. Cheng B, Erdos D, Stanley RJ, Stoecker WV, Calcara, DA, Gomez DD. Automatic detection of basal cell carcinoma using telangiectasia analysis in dermoscopy skin lesion images, Skin Res Technol. 2011;17(3):278-87.
13. Cheng B, Stanley RJ, Stoecker WV, Hinton K. Automatic telangiectasia analysis in dermoscopy images using adaptive critic design, Skin Res Technol. 2012;18(4):389-96
14. P. Kharazmi, J. Zheng, H. Lui, Z. Jane Wang, and T. K. Lee, "A Computer-Aided Decision Support System for Detection and Localization of Cutaneous Vasculature in Dermoscopy Images Via Deep Feature Learning," Journal of Medical Systems, vol. 42, no. 2, p. 33, Jan. 2018, doi: <https://doi.org/10.1007/s10916-017-0885-2>.
15. Karimi D, Dou H, Warfield S K, Gholipour A: Deep learning with noisy labels: Exploring techniques and remedies in medical image analysis, 2019, pp. 1–17, arXiv:1912.02911. [Online]. Available: <http://arxiv.org/abs/1912.02911>
16. P. Tschandl, C. Rosendahl, and H. Kittler, "Data Descriptor: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions Background & Summary," *Nature Publishing Group*, 2018, doi: 10.1038/sdata.2018.161.

17. N. Codella et al., “Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC).” *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Washington, DC, USA, 2018, pp. 168-172, doi: 10.1109/ISBI.2018.8363547.
18. M. Combalia et al., “BCN20000: Dermoscopic Lesions in the Wild,” arXiv:1908.02288 [cs, eess], Aug. 2019, Available: <https://arxiv.org/abs/1908.02288>
19. O. Ronneberger, P. Fischer, and T. Brox, U-Net: Convolutional networks for biomedical image segmentation. [Online]. Available: <http://lmb.informatik.uni-freiburg.de/>. Accessed 17 11 2020
20. Jadon, Shruti, A survey of loss functions for semantic segmentation, 2020, [Online], Available: <https://arxiv.org/abs/2006.14822> Available: <https://arxiv.org/abs/2006.14822> Accessed 17 11 2021.
21. N. Codella et al., “Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC).” *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, Washington, DC, USA, 2018, pp. 168-172, doi: 10.1109/ISBI.2018.8363547.
22. Oskal KRJ, Risdal, M., Janssen EAM, Undersrud ES, Gulsrud TO. . et al. A U-net based approach to epidermal tissue segmentation in whole slide histopathological images. *SN Appl. Sci* 2019; 1:672.
23. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
24. P. Kharazmi, M. I. AlJasser, H. Lui, Z. J. Wang, and T. K. Lee, “Automated Detection and Segmentation of Vascular Structures of Skin Lesions Seen in Dermoscopy, With an Application to Basal Cell Carcinoma Classification,” *IEEE J Biomed Health Inform*, vol. 21, no. 6, pp. 1675–1684, Nov. 2017, doi: <https://doi.org/10.1109/JBHI.2016.2637342>.
25. Maurya et al., “A deep learning approach to detect blood vessels in basal cell carcinoma,” *Skin Research and technology*, vol. 28, no. 4, pp. 571–576, Jul. 2022, doi: <https://doi.org/10.1111/srt.13150>.

26. P. Kharazmi, S. Kalia, H. Lui, Z. J. Wang, and T. K. Lee, “A feature fusion system for basal cell carcinoma detection through data-driven feature learning and patient profile,” *Skin Research and Technology*, vol. 24, no. 2, pp. 256–264, May 2018, doi: <https://doi.org/10.1111/srt.12422>.
27. Serrano, M. Lazo, A. Serrano, T. Toledo-Pastrana, R. Barros-Tornay, and B. Acha, “Clinically Inspired Skin Lesion Classification through the Detection of Dermoscopic Criteria for Basal Cell Carcinoma,” *Journal of Imaging*, vol. 8, no. 7, p. 197, Jul. 2022, doi: <https://doi.org/10.3390/jimaging8070197>.
28. W.V. Stoecker, Kapil Gupta, B. Shrestha, M. Wronkiewicz, R. Chowdhury, R.J. Stanley, J. Xu, R. H Moss, M. E. Celebi, H. S. Rabinovitz, M. Oliviero, J. M. Malters, I. Kolm, “Detection of basal cell carcinoma using color and histogram measures of semitranslucent areas,” *Skin Research and Technology*, 2009, vol. 15, no. 3, pp.283-7. doi: 10.1111/j.1600-0846.2009.00354
29. M. Tan and Q. Le, “Efficientnet: Rethinking model scaling for convolutional neural networks,” in *International Conference on Machine Learning*, 2019, pp. 6105–6114. doi: 10.48550/arXiv.1905.11946
30. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248–255.
31. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., & Chen, L. (2018). MobileNetV2: Inverted Residuals and Linear Bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4510-4520.
32. Lama, N., Kasmi, R., Hagerty, J., Stanley, J., Young, R., Miinch, J., Nepal, J., Nambisan, A., Stoecker, W. V. “ChimeraNet: U-Net for Hair Detection in Dermoscopic Skin Lesion Images”. *Journal of Digital Imaging* (2022). <https://doi.org/10.1007/s10278-022-00740-6>
33. K. Nambisan et al., “Deep learning-based dot and globule segmentation with pixel and blob-based metrics for evaluation,” *Intelligent Systems with Applications*, vol. 16, p. 200126, Nov. 2022, doi: <https://doi.org/10.1016/j.iswa.2022.200126>.

34. Esteva et al., “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, pp. 115–118, 2017, doi: 10.1038/nature21056.
35. Nambisan, A., Maurya, A., Lama, N., Phan, T., Patel, G., Miller, K., Lama, B., Hagerty, J., Stanley, R., & Stoecker, W. V. “Improving Automatic Melanoma Diagnosis Using Deep Learning-Based Segmentation of Irregular Networks”. *Cancers* (2023), 15(4), 1259. <https://doi.org/10.3390/cancers15041259>
36. P. Kharazmi, J. Zheng, H. Lui, Z. Jane Wang, and T. K. Lee, “A Computer-Aided Decision Support System for Detection and Localization of Cutaneous Vasculature in Dermoscopy Images Via Deep Feature Learning,” *Journal of Medical Systems*, vol. 42, no. 2, p. 33, Jan. 2018, doi: <https://doi.org/10.1007/s10916-017-0885-2>.
37. Otter, N., Porter, M. A., Tillmann, U., Grindrod, P., & Harrington, H. A. (2017, August 9). A roadmap for the computation of persistent homology - EPJ Data Science. SpringerOpen. <https://doi.org/10.1140/epjds/s13688-017-0109-5>
38. Carlsson, G. (2014). Topological pattern recognition for point cloud data. *Acta Numerica*, 23, 289-368. doi:10.1017/S0962492914000051
39. E., L., & Z. (2002, November 1). Topological Persistence and Simplification - Discrete & Computational Geometry. SpringerLink. <https://doi.org/10.1007/s00454-002-2885-2>
40. Hu, C. S., Lawson, A., Chen, J. S., Chung, Y. M., Smyth, C., & Yang, S. M. (2021, November 17). TopoResNet: A Hybrid Deep Learning Architecture and Its Application to Skin Lesion Classification. MDPI. <https://doi.org/10.3390/math9222924>
41. Bendich, P., Marron, J. S., Miller, E., Pieloch, A., & Skwerer, S. (2016, March 25). Persistent Homology Analysis of Brain Artery Trees. PubMed Central (PMC). <https://doi.org/10.1214/15-AOAS886>
42. P. Tschandl, C. Rosendahl, and H. Kittler, “Data Descriptor: The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions Background & Summary,” *Nature Publishing Group*, 2018, doi: 10.1038/sdata.2018.161.

43. M. Combalia et al., “BCN20000: Dermoscopic Lesions in the Wild,” arXiv:1908.02288 [cs, eess], Aug. 2019, Available: <https://arxiv.org/abs/1908.02288>
44. Lama, N., J., Hagerty, Nambisan, A., Stanley, R. J., & Stoecker, W. V. “Skin Lesion Segmentation in Dermoscopic Images with Noisy Data”. *Journal of Digital Imaging* (2023). <https://doi.org/10.1007/s10278-023-00819-8>
45. Anand, D.V., Meng, Z., Xia, K. et al. Weighted persistent homology for osmolyte molecular aggregation and hydrogen-bonding network analysis. *Sci Rep* 10, 9685 (2020). <https://doi.org/10.1038/s41598-020-66710-6>
46. Akanksha Maurya, Ronald J Stanley, Hemanth Y Aradhyula, Norsang Lama, Anand K Nambisan, Gehana Patel, Daniyal Saaed, Samantha Swinfard, Colin Smith, Sadhika Jagannathan, Jason Hagerty, & William V Stoecker. (2023). Basal cell carcinoma diagnosis with fusion of deep learning and telangiectasia features [Data set]. Zenodo. <https://doi.org/10.5281/zenodo.7709824>

VITA

Akanksha Maurya was born in New Delhi, India. Before attending Missouri University of Science and Technology, Rolla, USA, she attended Gautam Buddha Technical University, India, where she earned a Bachelor of Technology degree in Electrical Engineering in 2012. She then worked as a Senior Systems Engineer at Infosys Ltd. in Pune, India from 2013 to 2016. In 2016, she started her Master of Science in Electrical Engineering at Missouri University of Science and Technology. She received her Ph.D. degree in Electrical Engineering from Missouri University of Science and Technology in July 2023. Her academic experience included working as a Graduate Research Assistant and Graduate Teaching Assistant. Her research and teaching interests were in the fields of biomedical image analysis, deep learning, applications of artificial intelligence, image processing, topological data analysis and statistical modelling.