



## **DermX: An end-to-end framework for explainable automated dermatological diagnosis**

**Jalaboi, Raluca; Faye, Frederik; Orbes-Arteaga, Mauricio ; Jørgsen, Dan; Winther, Ole; Galimzianova, Alfiia**

*Published in:*  
Medical Image Analysis

*Link to article, DOI:*  
[10.1016/j.media.2022.102647](https://doi.org/10.1016/j.media.2022.102647)

*Publication date:*  
2023

*Document Version*  
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

*Citation (APA):*  
Jalaboi, R., Faye, F., Orbes-Arteaga, M., Jørgsen, D., Winther, O., & Galimzianova, A. (2023). DermX: An end-to-end framework for explainable automated dermatological diagnosis. *Medical Image Analysis*, 83, Article 102647. <https://doi.org/10.1016/j.media.2022.102647>

---

### **General rights**

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



# DermX: An end-to-end framework for explainable automated dermatological diagnosis

Raluca Jalaboi<sup>a,b,\*</sup>, Frederik Faye<sup>b</sup>, Mauricio Orbes-Arteaga<sup>b</sup>, Dan Jørgensen<sup>b</sup>, Ole Winther<sup>a,c,d</sup>, Alfiia Galimzianova<sup>b</sup>

<sup>a</sup> Department of Applied Mathematics and Computer Science at the Technical University of Denmark, Richard Petersens Plads, Building 324, DK-2800 Kongens Lyngby, Denmark

<sup>b</sup> Omhu A/S, Silkegade 8 st, DK-1113 Copenhagen C, Denmark

<sup>c</sup> Bioinformatics Centre, Department of Biology, University of Copenhagen, Copenhagen, Denmark

<sup>d</sup> Center for Genomic Medicine, Rigshospitalet, Copenhagen University Hospital, Copenhagen, Denmark

## ARTICLE INFO

### MSC:

92B20

92C50

68T45

### Keywords:

Dermatology

Explainability

Dataset

Convolutional neural networks

## ABSTRACT

Dermatological diagnosis automation is essential in addressing the high prevalence of skin diseases and critical shortage of dermatologists. Despite approaching expert-level diagnosis performance, convolutional neural network (ConvNet) adoption in clinical practice is impeded by their limited explainability, and by subjective, expensive explainability validations. We introduce DermX, an end-to-end framework for explainable automated dermatological diagnosis. DermX is a clinically-inspired explainable dermatological diagnosis ConvNet, trained using DermXDB, a 554 image dataset annotated by eight dermatologists with diagnoses, supporting explanations, and explanation attention maps. DermX+ extends DermX with guided attention training for explanation attention maps. Both methods achieve near-expert diagnosis performance, with DermX, DermX+, and dermatologist F1 scores of 0.79, 0.79, and 0.87, respectively. We assess the explanation performance in terms of identification and localization by comparing model-selected with dermatologist-selected explanations, and gradient-weighted class-activation maps with dermatologist explanation maps, respectively. DermX obtained an identification F1 score of 0.77, while DermX+ obtained 0.79. The localization F1 score is 0.39 for DermX and 0.35 for DermX+. These results show that explainability does not necessarily come at the expense of predictive power, as our high-performance models provide expert-inspired explanations for their diagnoses without lowering their diagnosis performance.

## 1. Introduction

Skin diseases affect a third of the global population (Hay et al., 2014) and are the fourth leading cause of disability worldwide (Karimkhani et al., 2017). The increasing demand for dermatological care is exacerbated by the low performance of general practitioners when diagnosing skin conditions (Federman et al., 1999), and by the global scarcity of expert dermatologists (Feng et al., 2018; Kringos et al., 2015).

Automation may help alleviate this problem. Convolutional neural networks (ConvNets) have been shown to achieve near expert-level performance in diagnosing dermatological conditions from images of skin lesions (Thomsen et al., 2020; Esteva et al., 2017), and that they are able to assist general practitioners as well as less experienced dermatologists in improving their diagnostic performance (Tschandl et al., 2020; Jain et al., 2021). However, the lack of a good explanation

mechanism (Kelly et al., 2019) for ConvNet decisions is one of the main obstacles to their adoption as automated diagnosis systems (Goodman and Flaxman, 2017; Kelly et al., 2019; Topol, 2019). A good explanation is expected to be both *plausible*, i.e. as similar as possible to a human explanation, and *faithful*, i.e. to accurately represent the inner workings of the network (Jacovi and Goldberg, 2020).

Different mechanisms for explaining ConvNet decisions have been proposed (Simonyan et al., 2014; Selvaraju et al., 2017; Ribeiro et al., 2016). Within the medical imaging literature, the most common explainability methods are saliency-based methods, such as raw saliency maps (Simonyan et al., 2014) and gradient-weighted class-activation attention maps (Grad-CAM) (Singh et al., 2020). While other methods were criticized due to their lack of faithfulness, Grad-CAMs have been shown to perform well (Adebayo et al., 2018). However, there remains a lack of standard metrics for plausibility validation, as the explanations they provide are often incomplete and difficult to quantify (Tschandl

\* Corresponding author at: Department of Applied Mathematics and Computer Science at the Technical University of Denmark, Richard Petersens Plads, Building 324, DK-2800 Kongens Lyngby, Denmark.


E-mail address: [rjal@dtu.dk](mailto:rjal@dtu.dk) (R. Jalaboi).

<https://doi.org/10.1016/j.media.2022.102647>


Received 15 December 2021; Received in revised form 17 August 2022; Accepted 27 September 2022

Available online 7 October 2022

1361-8415/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).



**Fig. 1.** Clinically-inspired convolutional neural network architecture for image diagnosis with explanations in the form of skin lesion characteristics. Given an image, the model is trained to predict the diagnosis together with the supporting characteristics, and to focus its attention on image sections that contain relevant characteristics. The diagnosis is predicted using both the characteristics identified by the model (similar to how dermatologists diagnose cases), and the extracted image features. Using the extracted features alongside the predicted characteristics ensures that no relevant information is lost, e.g. the age or the skin tone. The explanation module offers plausible, faithful explanations to the diagnosis predicted by the model, while also localizing the explanations in the image.



**Fig. 2.** Sample from the DermXDB dataset. A seborrheic dermatitis image from the SD-260 dataset was evaluated by eight dermatologists. Three evaluations are depicted in this figure. One dermatologist correctly diagnosed it as seborrheic dermatitis due to the presence of plaque. Another dermatologist incorrectly diagnosed it as acne due to the presence of open comedones, papules, and pustules, while a third dermatologist diagnosed it as acne due to the presence of pustules.

et al., 2020). More specifically, common ConvNet explainability methods provide no semantic information alongside the explanation, but rather focus on the image section where the network pays attention. In complex domains such as dermatology, this information is not enough to explain the decision mechanisms: knowing that the network focuses on the skin lesion does not explain why it diagnosed a case as acne and not rosacea. Moreover, such complex tasks require that thorough explanation validation be done by domain experts, which is a time consuming and expensive process. Current dermatological datasets focus either solely on disease diagnosis, or on lesion segmentation (DermNetNZ, 2021; Sun et al., 2016; Tschandl et al., 2018). Having access to expert-annotated dermatological diagnosis explanations would improve the validation of explainability methods and allow the training of intrinsically explainable models. However, to the best of our knowledge, no such dataset exists.

Our contributions are twofold. First, to enable a quantitative assessment of the explainability of dermatological diagnosis models, we introduce DermXDB, a dermatological explainability dataset with gold standard diagnostic explanations provided by eight board-certified dermatologists. DermXDB consists of 554 images from DermNetNZ (2021) and SD-260 (Sun et al., 2016) associated with one of six diagnoses and their explanations in the form of skin lesion characteristics, as defined by Nast et al. (2016). This labeling procedure mimics clinical practice, where dermatologists assess the characteristics of skin lesions to derive and support a tentative diagnosis (Oakley, 2017). An annotation example can be seen in Fig. 2.

Second, we introduce DermX – a novel, clinically-inspired ConvNet architecture for skin disease diagnosis and explanations. This architecture is illustrated in Fig. 1. Following the clinical approach of explaining dermatological diagnoses through skin lesion characteristics,

DermX first identifies relevant characteristics in the image (which can also be interpreted as diagnosis explanations), and then relies on them, alongside the image features, to diagnose the case. Using Grad-CAM (Selvaraju et al., 2017), we then localize the predicted characteristics in the image. We validate the plausibility and faithfulness of our explanations using DermXDB as the gold standard for explanations.<sup>1</sup>

### 1.1. Related work

Machine learning-based dermatological diagnosis systems have been widely investigated, achieving results on par with human experts (Esteve et al., 2017; Tschandl et al., 2020; Jain et al., 2021). These advances in the automated diagnosis of skin lesions were made possible in part by the emergence of various dermatological datasets, which contain images diagnosed by medical experts (Tschandl et al., 2018; DermNetNZ, 2021; Sun et al., 2016). The widely used ISIC dataset (Tschandl et al., 2018) also includes lesion segmentations that can partially serve as a basis for objective explanation measurement. However, these segmentations were not collected to explain the diagnosis, but rather to localize the lesions. This shortcoming becomes critical in diseases such as actinic keratosis, where the area surrounding the lesion is just as important for the diagnosis as the lesion itself (Tschandl et al., 2020).


Explainability is an important topic in machine learning in general and in medical imaging in particular. Saliency-based explainability methods, e.g. Grad-CAM (Selvaraju et al., 2017), are often used as a way to investigate if the models learn relevant features (Tschandl et al., 2020; Zhang et al., 2019; Barata et al., 2021). Other explainability methods, such as LIME (Ribeiro et al., 2016), Kernel-SHAP (Lundberg and Lee, 2017), and Sharp-LIME (Graziani et al., 2021) are less commonly used in the medical imaging literature.

Two works, one in natural language processing and the other in dermatological imaging, have a similar approach to explainability as ours. Within natural language processing, Mathew et al. (2021) propose a framework that explains hateful speech identification. Human readers were asked to identify the most important tokens in a sentence for the prediction of hateful speech. Then, the explanation plausibility and faithfulness of the model-generated explanations were quantified by comparing to the human annotations. Within dermatological image analysis, Barata et al. (2021) investigate how hierarchical taxonomies for skin lesion classification can be used to improve ConvNet skin cancer diagnosis capabilities. They train networks to follow the hierarchical classification of diseases in their prediction, and to focus on relevant parts of the image.

<sup>1</sup> The DermXDB dataset and the implementation of DermX and DermX+ are available at <https://github.com/ralucaj/dermx>.

**Table 1**  
Distribution of images over DermNetNZ and SD-260, and over the six possible diagnoses.

	Acne	Actinic keratosis	Psoriasis	Seborrheic dermatitis	Viral warts	Vitiligo	Total
DermNetNZ	58	48	47	15	46	77	291
SD-260	61	43	51	79	20	9	263
Total	119	91	98	94	66	86	554



**Fig. 3.** Localizable characteristics taxonomy. All characteristics were tailored to the six DermXDB diseases using medical resources (Nast et al., 2016; Oakley, 2017), and with the help of two senior dermatologists.

In this work, we combine the two approaches by detecting diagnosis-explaining characteristics, each with its own localization, and train two ConvNets to focus on the relevant part of the image for each characteristic. Both networks are evaluated for the plausibility and faithfulness of their explanations.

## 2. Material and methods

### 2.1. Explainability dataset

To enable explainable modeling, we identified the clinically relevant explanation taxonomy, designed an appropriate annotation protocol, and collected expert-labeled data. This resulted in DermXDB: a novel dermatological explainability dataset designed to enable the training of the proposed end-to-end explainable models and quantitative explainability evaluation. The dataset consists of 554 images that belong to one of the following classes: acne, actinic keratosis, psoriasis, seborrheic dermatitis, viral warts, or vitiligo. Images were sourced from DermNetNZ (2021) and SD-260 (Sun et al., 2016) with written permission from the owners. The distribution over datasets and diseases is described in Table 1. All images were evaluated by eight board-certified dermatologists, with between four and twelve years of clinical experience. Each evaluation consists of a diagnosis and supporting explanations in the form of global tags, localizable characteristics, their segmentations, and additional descriptive terms for basic characteristics.

The development of this dataset included several steps. First, we performed multiple experiments to define the target diseases and the nature of the explanations. Second, we selected the six diagnoses and defined the explanation taxonomy illustrated in Fig. 3. Third, the labelers were allowed a short period of time to get accustomed to the annotation protocol and the labeling tool by evaluating images from an internal dataset. Finally, DermXDB images were selected and sent to the dermatologists for labeling.

**Preliminary investigation.** Nine diseases were initially investigated: psoriasis, rosacea, vitiligo, seborrheic dermatitis, pityriasis rosea, viral warts, actinic keratosis, acne, and impetigo. These diseases were chosen based on prevalence (Lim et al., 2017) and by the expectation that they could be diagnosed using images as the only source of patient information (Oakley, 2017). Dermatologists were asked to diagnose and explain their decision in free-text for over 100 images. During this step, dermatologists could see the original diagnosis of the image, but had the option to disagree with it. This step led to the exclusion of rosacea, impetigo, and pityriasis rosea from future experiments due to

the difficulty in diagnosing them in the absence of the patient medical history. It also led to the introduction of a structured ontology for the diagnosis explanations to avoid manual processing of typos and synonyms.


**Diagnosis and explanation ontology.** Preliminary investigations also highlighted the importance of having a consistent explanation ontology. After analyzing free-text explanations, they were formalized as an extended list of skin lesion characteristics (Nast et al., 2016). The characteristics set was selected to sufficiently explain the six target diseases (Oakley, 2017). With the help of two senior dermatologists, several other relevant characteristics were added.

The resulting set of characteristics was split into non-localizable characteristics (e.g. age or sex), localizable characteristics (e.g. plaque or open comedo), and additional descriptive terms (e.g. red or well-circumscribed), according to the International League of Dermatological Societies' classification (Nast et al., 2016). Fig. 3 illustrates the final DermXDB explanation taxonomy, while more information about the other two types of labels is available in Appendix Figs. A.11 and A.12.

**Annotation protocol.** Dermatologists were first asked to diagnose the image, and then tag it with characteristics that explain their diagnosis. No information about the gold standard diagnosis or the disease distribution was made available. If the dermatologists were unable to evaluate the image due to poor quality, or if the image depicted a different disease than the target conditions, they had the option to discard it.

Dermatologists could then select diagnosis-supporting non-localizable characteristics as global image tags. Afterwards, they could select and outline localizable characteristics. Dermatologists were instructed to highlight all relevant areas for each characteristic, and were only allowed to include irrelevant areas if separating them from the characteristic was too time consuming or difficult. In other words, they were instructed to favor sensitivity over specificity. Finally, basic terms (as defined in Fig. 3) could be enriched with additional descriptive terms when required for the diagnosis explanation. Once all tags and characteristics were added, the image could be marked as complete.

After the taxonomy and annotation protocol were defined, all dermatologists underwent two rounds of on-boarding in Darwin, a browser-based labeling tool (V7-Labs, 2021). A screenshot of the labeling interface is shown in Appendix Fig. A.13. Following this, they were asked to annotate a set of 630 images from the DermNetNZ and SD-260 datasets.




**Fig. 4.** DermX architecture for image diagnosis with explanations in the form of skin lesion characteristics. The model is trained to predict both diagnoses and characteristics. Image features go through a dimensionality reduction linear layer to ensure that the characteristics are not overshadowed by the image features. The explainability module identifies diagnosis explanations in the form of characteristics, and their localization on the image can be detected through Grad-CAMs.

**Data cleaning.** Once annotations were performed, the dataset went through two cleanup steps. First, to avoid ambiguities in the dataset, annotations with diagnoses outside the target conditions were discarded. This resulted in 33 images being removed from the dataset because all eight dermatologists tagged these as ‘other disease’, e.g. acne keloidalis nuchae. The second step was to manually group images from the same patients. For all patients with more than one image, only the first image based on alphabetical order was kept. After cleanup, 554 images were left. Out of all evaluations performed on these images, 150 were discarded due to reports of low image quality, resulting in 4202 individual evaluations.

## 2.2. Explainable models

We propose two inherently explainable models for joint prediction of diagnosis and explanations. First, we design DermX: an end-to-end clinically-inspired architecture for explainable diagnosis, and train it on the reference diagnosis and expert-identified explanation labels. Next, we build an enhanced model that also includes learning of the explanation localization – DermX+. In the following, we provide a detailed description of each of the models.

**DermX model.** We propose a clinically-inspired model trained using the data described above. Following the multi-task learning paradigm, the model learns how to predict a diagnosis and its supporting characteristics at the same time. Using a ConvNet as an image feature extractor, we flatten and pass these features into the two prediction modules. The explainability module passes the features through a dense block, composed of a dropout layer, a linear layer with ReLU activations, and another dropout layer. This output is then passed into a linear layer with ten neurons and a logistic function is applied to each to give the probabilistic multi-label predictions, i.e. multiple characteristics can be predicted at the same time. The diagnosis module processes the image features using a similar dense block, after which they are concatenated to the characteristic logits. For this module, the dense block also doubles as a dimensionality reduction component, allowing the image features and the characteristics to have the same order of magnitude. The concatenated features are then passed through a linear layer with six neurons, followed by a softmax function to give our single-label prediction head for diagnoses. Fig. 4 illustrates the DermX architecture.



**Fig. 5.** DermX+ architecture used to generate explanations using guided attention. In addition to the DermX architecture described in Fig. 4, we introduce an additional loss term for the characteristics attention map. The Grad-CAM attention is computed for each predicted characteristic using the features extracted by the last convolutional layer in the backbone network. Characteristic Grad-CAMs are then compared to the downsized fuzzy fusion maps for each characteristic.

DermX optimizes the loss defined as follows. Let  $y_{i,d} \in \{0, 1\}$  and  $z_{i,c} \in \{0, 1\}$  be the target diagnosis and target characteristics for image  $i \in \{1, \dots, N\}$  in a batch of size  $N$ , where  $d \in \{1, \dots, D\}$  and  $c \in \{1, \dots, C\}$  denote the diagnosis and characteristic class, and let  $\hat{y}_{i,d} \in (0, 1)$  and  $\hat{z}_{i,c} \in (0, 1)$  be the diagnosis and characteristics predictions, respectively. The loss can then be written as

$$L = \lambda_D L_D + \lambda_C L_C, \quad (1)$$

where  $L_D$  is the categorical cross-entropy diagnosis loss defined as

$$L_D = -\frac{1}{ND} \sum_{i=1}^N \sum_{d=1}^D y_{i,d} \log \hat{y}_{i,d}, \quad (2)$$

$L_C$  is the binary cross-entropy characteristics loss defined as

$$L_C = -\frac{1}{NC} \sum_{i=1}^N \sum_{c=1}^C (z_{i,c} \log(\hat{z}_{i,c}) + (1 - z_{i,c}) \log(1 - \hat{z}_{i,c})), \quad (3)$$

and  $\lambda_D$  and  $\lambda_C$  are hyper-parameters for weighing the relative loss contributions.

**DermX+ model.** We build on top of the DermX architecture by introducing a guided attention element (Li et al., 2018). Fig. 5 highlights the difference between DermX and DermX+, namely the addition of a characteristic attention component.

In addition to the two losses optimized by DermX and described in Eq. (1), the DermX+ model also optimizes the attention loss term  $L_A$ :

$$L = \lambda_D L_D + \lambda_C L_C + \lambda_A L_A, \quad (4)$$


where  $L_A$  is the Dice loss for attention

$$L_A = \frac{1}{NC} \sum_{i=1}^N \sum_{c=1}^C \left( 1 - \frac{2A_{i,c}M_{i,c}}{A_{i,c} + M_{i,c}} \right), \quad (5)$$

with  $A_{i,c}$  being the attention map, and  $M_{i,c}$  being the fuzzy localization label, both for image  $i$  and characteristic  $c$ .

## 2.3. Model training and validation

**Data.** Given the limited size of the dataset, we create a stratified ten-fold cross-validation setup to train explainable models, leading




**Fig. 6.** Differences between dermatologist-labeled attention maps, distributed over the six diseases: acne, actinic keratosis (AK), psoriasis, seborrheic dermatitis (SD), viral warts, and vitiligo. The maps were computed as the union of all characteristics labeled by each of the first three dermatologists. Each color represents a different supporting characteristic.

to approximately 500 training images and 50 test images for each fold. Results presented in this paper are aggregated over all ten folds. For diagnosis prediction we use the gold standard diagnosis label, as defined in the source datasets. A characteristic was marked as relevant for a diagnosis if at least one dermatologist included the characteristic in their decision explanation. Characteristic labels for localization were created as aggregated fuzzy maps, i.e. each pixel value in a mask was generated as a fraction of how many dermatologists included it in their characteristic localization. Only characteristics selected for the correct diagnosis with regard to the gold standard were included both in defining the presence of a characteristic and in the fuzzy map aggregation. This way, we avoid introducing noise due to a mismatch between the diagnosis a dermatologist was explaining and the diagnosis label used to train the network. Additionally, we exclude characteristics that appear in fewer than 30 samples throughout the dataset and characteristics with an inter-rater F1 score below 0.30. We thus focus on closed comedo, dermatoglyph disruption, open comedo, papule, patch, plaque, pustule, scale, scar, and sun damage.

**Implementation details.** In all experiments, we use an EfficientNet-B2 (Tan and Le, 2019) ConvNet pre-trained on the ImageNet image recognition dataset (Deng et al., 2009) for feature extraction, with all layers fine-tuned on the DermXDB data. Both models were trained for 93 epochs using the AdamW optimizer (Loshchilov and Hutter, 2018), the cosine annealing with warm restarts learning rate scheduler (Loshchilov and Hutter, 2016), and a starting learning rate of 0.0005. Within the dense block we use linear layers with 64 neurons, dropout layers with 0.2 probability, and ReLU activations. DermX is trained with  $\lambda_D = 1$ ,  $\lambda_C = 1$ , while DermX+ uses  $\lambda_D = 1$ ,  $\lambda_C = 1$ , and  $\lambda_A = 10$ . Further information about the hyper-parameters used for training and other implementation details can be found in Appendix Table B.12.

#### 2.4. Explainability evaluation

We measure the performance with regard to the image diagnosis of both our dermatologists and our trained models using the F1 score,



**Fig. 7.** Characteristics labeled by a dermatologist for an acne case. Following the instructions, no characteristic was segmented but rather the region where they were present was identified without necessarily following the lesion boundaries. For more difficult characteristics to locate, e.g. scars, dermatologists were instructed to brush over entire areas containing the characteristic.

**Table 2**

Dermatologist inter-rater agreement for the presence or absence of characteristics (mean $\pm$ std). This analysis shows significant variation in the selection and agreement rates. Characteristics commonly considered important for diagnosing one of the diseases (e.g. comedones, plaques) have higher agreement rates, while uncommon characteristics (e.g. dermatoglyph disruption) display low selection and agreement rates.

	F1 score	Average selection
Dermatoglyph disruption	0.36 $\pm$ 0.35	21.50 $\pm$ 14.34
Closed comedo	0.54 $\pm$ 0.14	41.5 $\pm$ 22.20
Open comedo	0.65 $\pm$ 0.10	50.88 $\pm$ 23.28
Papule	0.67 $\pm$ 0.07	138.25 $\pm$ 33.44
Patch	0.76 $\pm$ 0.11	114.00 $\pm$ 43.10
Plaque	0.78 $\pm$ 0.06	205.12 $\pm$ 35.98
Pustule	0.76 $\pm$ 0.06	58.62 $\pm$ 13.52
Scale	0.89 $\pm$ 0.03	188.50 $\pm$ 31.16
Scar	0.46 $\pm$ 0.14	41.75 $\pm$ 26.84
Sun damage	0.46 $\pm$ 0.26	31.62 $\pm$ 14.91
Mean	0.63 $\pm$ 0.16	86.42 $\pm$ 67.05

**Table 3**

Dermatologist inter-rater localization agreement for localizable characteristics (mean $\pm$ std). Overlap measures show a significant variation between raters in outlining characteristics. Sensitivity values are high for characteristics that occupy larger areas and that often display well-circumscribed borders (e.g. plaque, scale), but tend to be lower in smaller characteristics (e.g. comedones, pustules).

	F1 score	Sensitivity	Specificity
Dermatoglyph disruption	0.68 $\pm$ 0.17	0.82 $\pm$ 0.16	0.98 $\pm$ 0.03
Closed comedo	0.20 $\pm$ 0.21	0.46 $\pm$ 0.36	0.89 $\pm$ 0.17
Open comedo	0.20 $\pm$ 0.18	0.44 $\pm$ 0.33	0.91 $\pm$ 0.15
Papule	0.27 $\pm$ 0.25	0.49 $\pm$ 0.34	0.94 $\pm$ 0.12
Patch	0.65 $\pm$ 0.18	0.80 $\pm$ 0.19	0.93 $\pm$ 0.11
Plaque	0.64 $\pm$ 0.21	0.79 $\pm$ 0.21	0.93 $\pm$ 0.10
Pustule	0.26 $\pm$ 0.17	0.50 $\pm$ 0.30	0.98 $\pm$ 0.08
Scale	0.51 $\pm$ 0.23	0.70 $\pm$ 0.25	0.93 $\pm$ 0.10
Scar	0.30 $\pm$ 0.23	0.56 $\pm$ 0.34	0.89 $\pm$ 0.14
Sun damage	0.71 $\pm$ 0.21	0.84 $\pm$ 0.19	0.76 $\pm$ 0.22
Mean	0.42 $\pm$ 0.20	0.62 $\pm$ 0.16	0.91 $\pm$ 0.06

sensitivity, and specificity. The same metrics are used to quantify the inter-rater agreement on image diagnosis and characteristics selection between dermatologists. The model performance on characteristics is measured with regard to the fuzzy fusion label for characteristics using the same three metrics. F1 score (also known as the Dice-Sørensen coefficient for pixel-level segmentation), sensitivity, and specificity are also used to measure the inter-rater agreement for the localizable

**Table 4**

Comparison of model diagnosis performance with regard to the gold standard, presented as the mean F1 score  $\pm$  std. The models compared are the diagnosis-only model (Dx), the clinically-inspired diagnosis and characteristics model (DermX), and the DermX model trained with guided attention (DermX+). Dermatologist scores are summarized as mean  $\pm$  std across the experts. The gold standard is the original image diagnosis as defined by the source dataset.

	Dx	DermX	DermX+	Expert
Acne	0.87 $\pm$ 0.05	0.87 $\pm$ 0.05	0.86 $\pm$ 0.06	0.94 $\pm$ 0.02
Actinic keratosis	0.80 $\pm$ 0.06	0.79 $\pm$ 0.14	0.73 $\pm$ 0.10	0.79 $\pm$ 0.12
Psoriasis	0.77 $\pm$ 0.07	0.73 $\pm$ 0.11	0.80 $\pm$ 0.09	0.87 $\pm$ 0.04
Seborrheic dermatitis	0.77 $\pm$ 0.07	0.74 $\pm$ 0.09	0.74 $\pm$ 0.10	0.75 $\pm$ 0.08
Viral warts	0.76 $\pm$ 0.18	0.76 $\pm$ 0.11	0.76 $\pm$ 0.15	0.92 $\pm$ 0.05
Vitiligo	0.78 $\pm$ 0.10	0.83 $\pm$ 0.10	0.86 $\pm$ 0.08	0.95 $\pm$ 0.02
Mean	0.79 $\pm$ 0.05	0.79 $\pm$ 0.04	0.79 $\pm$ 0.04	0.87 $\pm$ 0.08

**Table 5**

Performance comparison for characteristics identification with regard to dermatologist-generated labels, reported as mean F1 scores. We compare the clinical diagnosis and characteristics model (DermX), the DermX model trained with guided attention (DermX+), and the inter-rater agreement among dermatologists. A characteristic was tagged as present if at least one dermatologist marked it in an image. The F1 score for dermatologists is based on the pairwise inter-rater agreement on characteristics (mean  $\pm$  std).

	DermX	DermX+	Expert	Samples
Closed comedo	0.76 $\pm$ 0.08	0.81 $\pm$ 0.12	0.55 $\pm$ 0.15	96
Dermatoglyph disruption	0.74 $\pm$ 0.20	0.70 $\pm$ 0.20	0.37 $\pm$ 0.35	54
Open comedo	0.80 $\pm$ 0.06	0.80 $\pm$ 0.08	0.66 $\pm$ 0.10	110
Papule	0.79 $\pm$ 0.07	0.80 $\pm$ 0.07	0.68 $\pm$ 0.07	278
Patch	0.76 $\pm$ 0.06	0.79 $\pm$ 0.06	0.78 $\pm$ 0.11	249
Plaque	0.88 $\pm$ 0.03	0.90 $\pm$ 0.03	0.79 $\pm$ 0.06	352
Pustule	0.79 $\pm$ 0.10	0.81 $\pm$ 0.06	0.77 $\pm$ 0.06	106
Scale	0.79 $\pm$ 0.04	0.82 $\pm$ 0.05	0.91 $\pm$ 0.02	275
Scar	0.78 $\pm$ 0.10	0.80 $\pm$ 0.08	0.47 $\pm$ 0.14	115
Sun damage	0.66 $\pm$ 0.11	0.64 $\pm$ 0.15	0.46 $\pm$ 0.27	78
Mean	0.77 $\pm$ 0.03	0.79 $\pm$ 0.03	0.64 $\pm$ 0.16	171.30

**Table 6**

DermX performance for characteristics localization with regard to the fuzzy dermatologist localization maps, reported as mean soft sensitivity, specificity, and F1 score. DermX performance metrics are computed only on samples where both the model and the dermatologists agree on the relevance of a characteristic, in order to decouple localization performance from the identification performance.

	Sensitivity	Specificity	F1 score	Samples
Closed comedo	0.69 $\pm$ 0.09	0.69 $\pm$ 0.05	0.40 $\pm$ 0.04	75
Dermatoglyph disruption	0.69 $\pm$ 0.09	0.69 $\pm$ 0.05	0.28 $\pm$ 0.06	36
Open comedo	0.69 $\pm$ 0.06	0.68 $\pm$ 0.04	0.36 $\pm$ 0.05	90
Papule	0.63 $\pm$ 0.08	0.72 $\pm$ 0.05	0.34 $\pm$ 0.04	219
Patch	0.57 $\pm$ 0.07	0.78 $\pm$ 0.04	0.43 $\pm$ 0.05	188
Plaque	0.65 $\pm$ 0.06	0.75 $\pm$ 0.04	0.43 $\pm$ 0.03	314
Pustule	0.69 $\pm$ 0.07	0.69 $\pm$ 0.05	0.24 $\pm$ 0.05	88
Scale	0.65 $\pm$ 0.07	0.76 $\pm$ 0.04	0.41 $\pm$ 0.03	222
Scar	0.64 $\pm$ 0.06	0.72 $\pm$ 0.05	0.46 $\pm$ 0.05	90
Sun damage	0.44 $\pm$ 0.08	0.87 $\pm$ 0.04	0.56 $\pm$ 0.06	50
Mean	0.64 $\pm$ 0.02	0.74 $\pm$ 0.01	0.39 $\pm$ 0.02	137.20

characteristics region outlining overlap. All values are reported as the mean and the standard deviation (std) over the 10 folds.

We define the explainability of our models as having two components: plausibility and faithfulness. For plausibility, we focus on both the identification and the localization of characteristics. First, we measure the F1 score, sensitivity, and specificity per characteristic to measure the models' ability to correctly identify the right explanations. Similar to Mathew et al. (2021), we compare the Grad-CAM activations per characteristic with the fuzzy attention maps for each characteristic, and measure their similarity using the F1 score, sensitivity, and specificity. All pixel-based metrics are implemented using fuzzy logic, as follows:

$$F1 = \frac{2 \sum_{p \in P} \min(A_p, M_p)}{\sum_{p \in P} (A_p) + \sum_{p \in P} (M_p)} \quad (6)$$

**Table 7**

DermX+ performance for characteristics localization with regard to the fuzzy dermatologist localization maps, reported as mean soft sensitivity, specificity, and F1 score. DermX+ values are computed only on samples where both the model and the dermatologists agree on the relevance of a characteristic, in order to decouple localization performance from the identification performance.

	Sensitivity	Specificity	F1 score	Samples
Closed comedo	0.11 $\pm$ 0.08	0.96 $\pm$ 0.02	0.10 $\pm$ 0.09	63
Dermatoglyph disruption	0.60 $\pm$ 0.15	0.97 $\pm$ 0.02	0.60 $\pm$ 0.12	34
Open comedo	0.06 $\pm$ 0.05	0.97 $\pm$ 0.02	0.06 $\pm$ 0.05	93
Papule	0.19 $\pm$ 0.06	0.97 $\pm$ 0.02	0.18 $\pm$ 0.04	232
Patch	0.56 $\pm$ 0.06	0.89 $\pm$ 0.03	0.53 $\pm$ 0.05	191
Plaque	0.65 $\pm$ 0.06	0.91 $\pm$ 0.01	0.61 $\pm$ 0.05	312
Pustule	0.04 $\pm$ 0.05	0.98 $\pm$ 0.01	0.03 $\pm$ 0.03	91
Scale	0.58 $\pm$ 0.11	0.89 $\pm$ 0.03	0.49 $\pm$ 0.07	224
Scar	0.35 $\pm$ 0.15	0.92 $\pm$ 0.04	0.35 $\pm$ 0.12	93
Sun damage	0.53 $\pm$ 0.19	0.90 $\pm$ 0.10	0.58 $\pm$ 0.20	47
Mean	0.37 $\pm$ 0.01	0.94 $\pm$ 0.00	0.36 $\pm$ 0.01	138.00

$$\text{Sensitivity} = \frac{\sum_{p \in P} \min(A_p, M_p)}{\sum_{p \in P} (M_p)} \quad (7)$$

$$\text{Specificity} = \frac{\sum_{p \in P} \min(1 - A_p, 1 - M_p)}{\sum_{p \in P} (1 - M_p)} \quad (8)$$

where  $P$  represents the pixels included in the analysis,  $A$  defines the class activations, and  $M$  represents the fuzzy label maps.

Following the comprehensiveness evaluation described by DeYoung et al. (2020), we measure the faithfulness of our models through the use of contrastive examples. Given a model  $m$ , an input image  $x$ , a set of explanation outlines  $e$ , a contrastive image  $x_e$  where all areas marked as an explanation for the image  $x$  were occluded, and the class probability output  $m(x)$  for the predicted class on the original input  $x$  we measure the faithfulness  $F$  as

$$F = m(x) - m(x_e) \quad (9)$$


In other words, the faithfulness describes what impact removing the explanations  $e$  from the image would have on the decision of model  $m$ . We decided not to include the sufficiency metric as it would lead to out-of-distribution images, such as a blank background with a plaque or a couple of pustules.

Finally, given the intrinsic disagreement between experts within medical fields, we postulate that explainable models should be able to properly argue their decisions, regardless of whether it matches the gold standard or not. Similar to how dermatologists may debate the correct diagnosis for a case by highlighting different explanations that support their decision, we expect an explainable model to do the same. However, as we do not always have the gold standard explanation for a wrong diagnosis, we need to define a basic set of explanations for any disease. To this end, we define the expected explanation as the prevalence of each characteristic within the dermatologists explanations for a diagnosis (Appendix Table A.11). Then, for the wrongly predicted diagnoses we compare the set of characteristics associated with that prediction with the expected explanation for the predicted diagnosis. For example, a case incorrectly classified as psoriasis is expected to be explained using one or several of papule, plaque, and scale, which are commonly used by dermatologists in their explanations of psoriasis. We evaluate how the model explanation for wrong diagnoses by computing the precision of the model's explanations with regard to the expected explanation for a diagnosis.


### 3. Results

#### 3.1. DermXDB analysis


We first analyzed the data focusing on dermatologist performance with regard to the gold standard diagnosis and their inter-rater agreement on both diagnoses and supporting characteristics. A total of 554



**Fig. 8.** Characteristic attention maps for a correctly classified psoriasis case, for the two identified characteristics: plaque (first row) and scale (second row). The first column shows the dermatologist-derived fuzzy attention map, the second one illustrates the Grad-CAM for each characteristic generated by DermX, while the last column shows the DermX+ Grad-CAM maps. DermX+ displays much closer results to the gold standard, while DermX maps include more irrelevant information, such as the finger knuckles.



**Fig. 9.** Characteristic attention maps for a psoriasis case wrongly classified as seborrheic dermatitis by both DermX and DermX+. While both models detect plaque and scale in the image, their detection of patch, a characteristic indicative of seborrheic dermatitis, leads them to misdiagnose the image.



**Fig. 10.** Contrastive samples on a DermXDB psoriasis image for DermX and DermX+. DermX, a more sensitive and less specific model, occludes large parts of the image, while DermX+, a more specific and less sensitive model, occludes only the lesion. When evaluating the contrastive sample for DermX+, the model has the possibility to use other diagnosis hints in the image (defined in DermXDB as non-localizable characteristics) that are occluded in the DermX sample.

images were included in this analysis, each with eight evaluations performed by board-certified dermatologists. The dermatologist diagnostic

performance in terms of mean F1 score with regard to the gold standard varies between 0.75 for seborrheic dermatitis and 0.95 for vitiligo. Aggregated F1 scores can be seen in Table 4. A full description of the dermatologist performance with regard to the gold standard is available in Appendix Table A.8.

Inter-rater agreement on characteristics, as described in Table 2, varies significantly more, partially due to the lower number of selections per class. Most basic terms display high levels of agreement, with F1 scores between 0.67 and 0.89. The exceptions are macule with an F1 score of 0.12 and nodule with an F1 score of 0.17, both also displaying low selection rates. Several additional terms, such as open and closed comedones, display levels of agreement similar to the basic terms. Fig. 6 illustrates an example of disagreement between three dermatologists on the location of supporting characteristics on one random case for each disease, while Fig. 7 highlights how the protocol was followed by a dermatologist in an acne case. Additional metrics for the full set of characteristics are described in Appendix Table A.9.

Outlining characteristics is a more difficult task, as confirmed by the low inter-rater F1 scores reported in Table 3. The lower F1 values can also be explained by how difficult outlining small or poorly circumscribed characteristics is. In terms of sensitivity, we notice the same trend as in binary agreement: dermatologists tend to agree more on the basic terms. Metrics for the full set of localizable characteristics are presented in Appendix Table A.10.

### 3.2. Explainable model

We trained a clinically-inspired model from Fig. 1 (DermX), and the same model architecture trained with guided attention (DermX+) for characteristics localization. We also train a diagnosis-only model (Dx) to check whether adding explanations impacts the diagnosis performance of DermX and DermX+.

Table 4 compares the diagnostic performance between all three models and the dermatologists with regard to the gold standard diagnosis. More information about their diagnostic performance is presented in Appendix B. For comparison, we trained a diagnosis-only model with a ResNet50 (He et al., 2016) base to validate the choice of architecture, and a diagnosis-only model trained with proportional class weights. The ResNet-based model achieved a macro F1-score of  $0.79 \pm 0.06$ , while the weighted class model showed a similar macro F1-score of  $0.78 \pm 0.05$ . More information about these two models is available in Appendix Table B.13. Additionally, we trained four interpretable models on the characteristics data for diagnosis prediction: a logistic regression model, a decision tree, a k-nearest neighbor with five neighbors, and a categorical naive Bayes models. These models obtained macro F1 scores of  $0.86 \pm 0.04$ ,  $0.85 \pm 0.05$ ,  $0.80 \pm 0.05$ , and  $0.86 \pm 0.05$ , respectively.

All models display similar F1 scores on all six diseases. The best results are obtained for vitiligo and acne, two disease classes where dermatologists also display high F1 score values. Seborrheic dermatitis on the other hand seems to be a difficult disease class for both dermatologists and models. For the rest of the results section we will focus on DermX and DermX+.

In terms of explanation plausibility, we look at both the identification of explanations, defined as the ability to detect the same characteristics as a dermatologist, and at their localization in the image. A comparison of F1 scores is described in Table 5. The two models perform well for explanation identification, with DermX+ obtaining slightly better results on most characteristics. Compared to dermatologists, the models perform within standard deviation bounds of the inter-rater agreement. Additional metrics are reported in Appendix Tables B.16 and B.17.

The localization plausibility of the models' explanation is quantified in Tables 6 and 7, with more statistics being presented in Appendix Tables B.18, B.19, B.20, and B.21. DermX performs adequately well on all characteristics. DermX+ is better at localizing large characteristics, e.g. patches or scales, but performs poorly on smaller characteristics, e.g. open and closed comedones. Dermatologists F1 scores indicate




Fig. A.11. Non-localizable characteristics taxonomy. These characteristics were added to the International League of Dermatological Societies' classification as global image tags after being flagged as relevant by our senior dermatologists.




Fig. A.12. Additional descriptive terms for localizable characteristics. All terms were tailored for the six diseases from medical resources (Nast et al., 2016; Oakley, 2017), and with the help of two senior dermatologists.




Fig. A.13. Labeling tool interface, exemplified for a psoriasis case from the SD-260 dataset. In the global tag search box (area 1, bottom right), dermatologists can select the disease, relevant demographics information, and lesion distribution. The brush selection menu (area 2, top left) allows them to select and mark localizable characteristics on the image. The full annotation menu (area 3, top right) is used to select of additional descriptive terms for the localized basic terms.

**Table A.8**  
Dermatologist diagnosis performance with regard to the gold standard (mean±std).

	F1 score	Sensitivity	Specificity	Average selection
Acne	0.94 ± 0.02	0.90 ± 0.04	0.99 ± 0.00	102.75 ± 10.00
Actinic keratosis	0.79 ± 0.12	0.68 ± 0.16	1.00 ± 0.00	57.62 ± 16.54
Psoriasis	0.87 ± 0.04	0.88 ± 0.04	0.97 ± 0.02	98.00 ± 8.09
Seborrheic dermatitis	0.75 ± 0.08	0.64 ± 0.11	0.99 ± 0.01	60.38 ± 9.23
Viral warts	0.92 ± 0.05	0.85 ± 0.08	1.00 ± 0.00	55.38 ± 5.05
Vitiligo	0.95 ± 0.02	0.90 ± 0.05	1.00 ± 0.00	76.25 ± 4.82
Mean	0.87 ± 0.08	0.81 ± 0.11	0.99 ± 0.01	75.06 ± 19.15

that the two models are, in some characteristics, within standard deviation of the inter-rater agreement. For other characteristics, such as dermatoglyph disruption for DermX and pustule for DermX+, the model performance is below the expert inter-rater agreement. Fig. 8 illustrates the explanations given for a correctly predicted psoriasis case by DermX and DermX+, respectively, while Fig. 9 shows the explanations given by the two ConvNets for a misclassified psoriasis case.

Explanation precision scores for the correct diagnosis prediction were computed with regard to the dermatologist labels. The resulting values are  $0.88 \pm 0.03$  for DermX and  $0.90 \pm 0.03$  for DermX+. On the wrong diagnosis prediction, DermX precision is  $0.85 \pm 0.06$ , while DermX+ precision is  $0.86 \pm 0.04$ . Mean faithfulness results are  $0.42 \pm 0.06$  for DermX and  $0.27 \pm 0.06$  for DermX+.

**Table A.9**

Dermatologist inter-rater agreement for the presence or absence of characteristics (mean $\pm$ std). This analysis shows significant variation in the selection and agreement rates. Characteristics commonly considered important for diagnosing one of the diseases (e.g. comedones, plaques) have higher agreement rates, while uncommon characteristics (e.g. leukotrichia, telangiectasia) display low selection and agreement rates.

	F1 score	Sensitivity	Specificity	Cohen's kappa	Average selection
<b>Basic terms</b>					
Macule	0.12 $\pm$ 0.15	0.18 $\pm$ 0.23	0.95 $\pm$ 0.07	0.09 $\pm$ 0.15	25.25 $\pm$ 30.37
Nodule	0.17 $\pm$ 0.16	0.26 $\pm$ 0.27	0.97 $\pm$ 0.04	0.16 $\pm$ 0.15	17.88 $\pm$ 17.42
Papule	0.67 $\pm$ 0.07	0.69 $\pm$ 0.13	0.86 $\pm$ 0.08	0.52 $\pm$ 0.09	138.25 $\pm$ 33.44
Patch	0.76 $\pm$ 0.11	0.79 $\pm$ 0.18	0.92 $\pm$ 0.08	0.68 $\pm$ 0.15	114.00 $\pm$ 43.10
Plaque	0.78 $\pm$ 0.06	0.80 $\pm$ 0.11	0.84 $\pm$ 0.08	0.62 $\pm$ 0.09	205.12 $\pm$ 35.98
Pustule	0.76 $\pm$ 0.06	0.78 $\pm$ 0.11	0.96 $\pm$ 0.02	0.73 $\pm$ 0.06	58.62 $\pm$ 13.52
Scale	0.89 $\pm$ 0.03	0.90 $\pm$ 0.06	0.93 $\pm$ 0.04	0.82 $\pm$ 0.04	188.50 $\pm$ 31.16
<b>Additional terms</b>					
Closed comedo	0.54 $\pm$ 0.14	0.64 $\pm$ 0.25	0.96 $\pm$ 0.04	0.51 $\pm$ 0.14	41.5 $\pm$ 22.20
Cyst	0.20 $\pm$ 0.16	0.31 $\pm$ 0.34	0.99 $\pm$ 0.01	0.19 $\pm$ 0.16	6.25 $\pm$ 6.70
Dermatoglyph disruption	0.36 $\pm$ 0.35	0.39 $\pm$ 0.38	0.97 $\pm$ 0.03	0.37 $\pm$ 0.34	21.50 $\pm$ 14.34
Leukotrichia	0.43 $\pm$ 0.41	0.45 $\pm$ 0.43	1.00 $\pm$ 0.01	0.45 $\pm$ 0.41	4.62 $\pm$ 2.69
Open comedo	0.65 $\pm$ 0.10	0.71 $\pm$ 0.21	0.96 $\pm$ 0.04	0.61 $\pm$ 0.11	50.88 $\pm$ 23.28
Scar	0.46 $\pm$ 0.14	0.57 $\pm$ 0.26	0.95 $\pm$ 0.05	0.42 $\pm$ 0.14	41.75 $\pm$ 26.84
Sun damage	0.46 $\pm$ 0.26	0.53 $\pm$ 0.29	0.97 $\pm$ 0.02	0.43 $\pm$ 0.25	31.62 $\pm$ 14.91
Telangiectasia	0.17 $\pm$ 0.25	0.19 $\pm$ 0.28	0.99 $\pm$ 0.01	0.19 $\pm$ 0.25	6.12 $\pm$ 5.60
Thrombosed capillaries	0.36 $\pm$ 0.27	0.45 $\pm$ 0.37	0.98 $\pm$ 0.02	0.35 $\pm$ 0.26	15.88 $\pm$ 11.78
Mean	0.49 $\pm$ 0.24	0.54 $\pm$ 0.22	0.95 $\pm$ 0.04	0.45 $\pm$ 0.21	60.48 $\pm$ 62.97

**Table A.10**

Dermatologist inter-rater localization agreement for localizable characteristics (mean $\pm$ std). Overlap measures show a significant variation between raters in outlining characteristics. Sensitivity values are high for characteristics that occupy larger areas and that often display well-circumscribed borders (e.g. plaque, scale), but tend to be lower in smaller characteristics (e.g. comedones, pustules).

	F1 score	Sensitivity	Specificity
<b>Basic terms</b>			
Macule	0.21 $\pm$ 0.16	0.44 $\pm$ 0.3	0.95 $\pm$ 0.11
Nodule	0.31 $\pm$ 0.24	0.55 $\pm$ 0.33	0.96 $\pm$ 0.09
Papule	0.27 $\pm$ 0.25	0.49 $\pm$ 0.34	0.94 $\pm$ 0.12
Patch	0.65 $\pm$ 0.18	0.80 $\pm$ 0.19	0.93 $\pm$ 0.11
Plaque	0.64 $\pm$ 0.21	0.79 $\pm$ 0.21	0.93 $\pm$ 0.10
Pustule	0.26 $\pm$ 0.17	0.50 $\pm$ 0.30	0.98 $\pm$ 0.08
Scale	0.51 $\pm$ 0.23	0.70 $\pm$ 0.25	0.93 $\pm$ 0.10
<b>Additional terms</b>			
Closed comedo	0.20 $\pm$ 0.21	0.46 $\pm$ 0.36	0.89 $\pm$ 0.17
Cyst	0.39 $\pm$ 0.27	0.59 $\pm$ 0.31	0.98 $\pm$ 0.07
Dermatoglyph disruption	0.68 $\pm$ 0.17	0.82 $\pm$ 0.16	0.98 $\pm$ 0.03
Leukotrichia	0.50 $\pm$ 0.14	0.70 $\pm$ 0.21	0.96 $\pm$ 0.07
Open comedo	0.20 $\pm$ 0.18	0.44 $\pm$ 0.33	0.91 $\pm$ 0.15
Scar	0.30 $\pm$ 0.23	0.56 $\pm$ 0.34	0.89 $\pm$ 0.14
Sun damage	0.71 $\pm$ 0.21	0.84 $\pm$ 0.19	0.76 $\pm$ 0.22
Telangiectasia	0.29 $\pm$ 0.14	0.53 $\pm$ 0.28	0.93 $\pm$ 0.12
Thrombosed capillaries	0.42 $\pm$ 0.22	0.65 $\pm$ 0.28	0.99 $\pm$ 0.02
Mean	0.41 $\pm$ 0.18	0.62 $\pm$ 0.14	0.93 $\pm$ 0.05

## 4. Discussion

To the best of our knowledge, DermX is the first end-to-end framework created for the purpose of explaining automated dermatological diagnoses. The two ConvNets we introduce, DermX and DermX+, mimic the dermatological approach to diagnosing skin conditions: first they recognize supporting characteristics, then they use these characteristics as well as other high level information to arrive at a diagnosis. In addition to identifying supporting characteristics as explanations to a diagnosis, DermX+ also learns the localization of the explanations via the guided attention loss. The decision to use an attention mechanism for localization rather than a semantic segmentation approach was guided by the design of the annotation protocol. Because dermatologists were instructed to highlight explanation regions in an image with a focus on sensitivity instead of specificity, the resulting outlines are not well suited as segmentation masks. For this work to be possible, we collected diagnoses and supporting characteristics for 554 images from eight board-certified dermatologists.

During the process of collecting the DermXDB data, we found that dermatologists often focus on different characteristics when diagnosing a case. While most explanations for diseases display a set of common characteristics, such as scales, plaques, and papules for psoriasis, there is also a long tail of relevant characteristics that are not always selected. In addition, we found that inter-rater agreement was low for characteristics localization. This may be caused by the difficulty in outlining characteristics with poorly defined boundaries, such as patch, but also by dermatologists differing in their approach to outlining smaller characteristics, such as open and closed comedones.

The contrast between high agreement on diagnoses and low agreement on supporting characteristics illustrates how different experts perceive explanations in different ways. Although they generally agree on the diagnosis, dermatologists focus on different characteristics to explain their decision. To properly evaluate a model's explanations, we must therefore consider the opinions of multiple experts. Moreover, this intrinsic variability in how experts approach explanations lends more urgency to the need for quantifiable explanation methods.

From a modeling perspective, our results contradict the common adage that there must be a trade-off between predictive power and explainability. DermX and DermX+ both report the same diagnosis performance as a standard diagnosis-only ConvNet, while also offering plausible explanations for their decisions. Even in cases where they predict the wrong diagnosis, both models provide arguments that make sense for their prediction. Most explanations given by both models are within standard deviation of the inter-rater agreement on characteristics, suggesting that either model may function as a second opinion with realistic decision explanations.

When compared to interpretable models trained on the characteristics data, both DermX and DermX+ obtain a diagnosis performance within standard deviation of the models using manually labeled features. None of the models we trained obtains a diagnosis performance as high as that of experts. We postulate that this is due to the difficulty of the dataset, as shown by the inter-rater agreement in Table 4, and due to the limited amount of training data. On the other hand, our results are on par with the diagnosis accuracy reported by other research groups using dermatological clinical photography, which varies between 56.7% on 134 classes (Han et al., 2020) and 86.53% on four classes (Burlina et al., 2019).

Our localization results for both models are lower than the inter-rater agreement on expert-derived maps for most characteristics. This may in part be due to the low inter-rater agreement on the localization data, and in part due to the small scale at which the maps were

**Table A.11**  
Characteristics prevalence per disease.

	Acne	Actinic keratosis	Psoriasis	Seborrheic dermatitis	Viral warts	Vitiligo
Closed comedo	0.40	0.00	0.00	0.00	0.00	0.00
Dermatoglyph disruption	0.00	0.00	0.00	0.00	0.38	0.00
Open comedo	0.49	0.00	0.00	0.01	0.00	0.00
Papule	0.70	0.13	0.18	0.04	0.65	0.00
Patch	0.02	0.32	0.02	0.24	0.00	0.97
Plaque	0.04	0.59	0.96	0.71	0.46	0.01
Pustule	0.56	0.00	0.01	0.00	0.00	0.00
Scale	0.01	0.81	0.89	0.78	0.05	0.00
Scar	0.38	0.03	0.00	0.00	0.00	0.00
Sun damage	0.00	0.47	0.01	0.01	0.00	0.01

**Table B.12**  
Training hyper-parameters common to all trained models.

Hyperparameter	Value
Batch size	32
Rotation	10
Zoom	0.15
Brightness	0.35
Contrast	0.20
Saturation	0.20
Scale	(0.85, 1.15)
Translate	(0.15, 0.15)
Hue	0.15
Dropout	0.20

computed (nine by nine pixels for the EfficientNet-B2 architecture). However, the high sensitivity values show that these maps are often good enough to give a visual hint as to the location of the characteristic in an image. Such a hint would be useful in cases where an expert using DermX or DermX+ as a second opinion did not notice that characteristic. Comparing the two models, DermX+ displays lower overall F1 scores than DermX, while showcasing higher overall specificity and high sensitivity on large characteristics. This may be explained by its training target: dermatologist attention maps were linearly scaled down to the size of the feature maps, which may have reduced the target attention map of small characteristics to an almost empty mask. Another possible explanation is DermX's reliance on the sometimes noisy localization data. In particular, for characteristics smaller than 1 cm (closed comedo, open comedo, papule, and pustule), DermX+ is clearly outperformed by DermX due to the lower specificity and higher sensitivity of DermX. In the future, we plan on investigating different ways of downscaling the masks, and to increase the feature map size to take advantage of the high resolution gold standard attention maps.

Mean faithfulness scores above zero for both models prove that the characteristic localizations are indeed explanations about the diagnosis decision mechanisms of the models. DermX+, a more specific model in terms of characteristic localization, has lower faithfulness scores than DermX, which tends to include adjacent regions in its localizations. Fig. 10 showcases the impact a model's specificity and sensitivity have on the contrastive samples, and therefore on the faithfulness metric. In this example, the contrastive sample created by DermX+ still displays image-level non-localizable characteristics, information which is occluded in the DermX contrastive sample. This further confirms the importance of image-level tags in skin lesion diagnosis, e.g. by noticing that acne is predominantly located on the face or upper trunk, or that actinic keratosis most commonly affects elderly people.

This work opens many new research avenues in the domain of medical image diagnosis explainability. From a dermatological data perspective, we plan on adding more diseases and supporting characteristics to DermXDB. The annotation protocol developed as part of DermXDB can serve as an inspiration not only for explaining other dermatological diseases, but also for different radiology and pathology investigations. In radiology, imaging findings are routinely recorded, including the supporting characteristics for the diagnosis. For example,

the malignancy of a lesion seen on a mammogram could be supported by localized characteristics such as calcifications and dense tissue (Sickles et al., 2013). The network architectures we proposed could also be applied to learning the supporting radiological findings as explanations to diagnoses provided appropriately labeled datasets. From a modeling perspective, we will focus on leveraging the full potential of DermXDB by adding image-level explanations to the diagnosis models, and by incorporating the additional descriptive terms into the explanation setup. More work can be done in improving the characteristic localization. We will be focusing in particular on introducing the adversarial loss described in Li et al. (2018) for semi-supervised attention guidance. Another approach we will be to train object detection networks (Tan et al., 2020; Redmon and Farhadi, 2018) to detect the supporting characteristics alongside the diagnosis. Once the localization reaches a higher performance, a true test of the DermX architecture would be to set up a clinical trial where its predictions would be used as a second opinion for health care professionals of various levels of expertise.

## 5. Conclusions

In this work, we introduce DermX – a novel, clinically-inspired explainable ConvNet architecture for skin lesion diagnosis. We also introduce a variation named DermX+ that adds a guided attention loss such that the localization of lesion characteristics becomes a part of the supervised training. We quantify the explanation quality by comparing it with explanations given by board-certified dermatologists with different levels of clinical experience. To facilitate future work, we release this explainability dataset to the public, and describe the annotation protocol used for its creation.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Raluca Jalaboi reports financial support was provided by Innovation Fund Denmark. Ole Winther reports financial support was provided by Novo Nordisk Foundation. Ole Winther reports financial support was provided by Danish National Research Foundation. Frederik Faye, Mauricio Orbes-Arteaga, Dan Jørgensen, and Alfiia Galimzianova have no conflicts of interest.

## Data availability

Data will be made available on request.

## Acknowledgments

Raluca Jalaboi's work was supported in part by the Danish Innovation Fund, Denmark under Grant 0153-00154 A. Ole Winther's work was funded in part by the Novo Nordisk Foundation, Denmark through the Center for Basic Machine Learning Research in Life Science (NNF200C0062606). Ole Winther acknowledges supporting the Pioneer Centre for AI, DNRF, Denmark grant number P1.

**Table B.13**

Comparison of model diagnosis performance with regard to the gold standard, presented as the mean F1 score  $\pm$ std. The models compared are the diagnosis-only model (Dx), diagnosis-only with a ResNet50 base (DxRN), a diagnosis-only model trained with class weights, the clinically-inspired diagnosis and characteristics model (DermX), and the DermX model trained with guided attention (DermX+). Dermatologist scores are summarized as mean  $\pm$ std across the experts. The gold standard is the original image diagnosis as defined by the source dataset.

	Dx	DxRN	DxW	DermX	DermX+	Expert
Acne	0.87 $\pm$ 0.05	0.90 $\pm$ 0.06	0.85 $\pm$ 0.05	0.87 $\pm$ 0.05	0.86 $\pm$ 0.06	0.94 $\pm$ 0.02
Actinic keratosis	0.80 $\pm$ 0.06	0.74 $\pm$ 0.16	0.77 $\pm$ 0.15	0.79 $\pm$ 0.14	0.73 $\pm$ 0.10	0.79 $\pm$ 0.12
Psoriasis	0.77 $\pm$ 0.07	0.76 $\pm$ 0.07	0.77 $\pm$ 0.06	0.73 $\pm$ 0.11	0.80 $\pm$ 0.09	0.87 $\pm$ 0.04
Seborrheic dermatitis	0.77 $\pm$ 0.07	0.72 $\pm$ 0.11	0.75 $\pm$ 0.14	0.74 $\pm$ 0.09	0.74 $\pm$ 0.10	0.75 $\pm$ 0.08
Viral warts	0.76 $\pm$ 0.18	0.74 $\pm$ 0.16	0.73 $\pm$ 0.20	0.76 $\pm$ 0.11	0.76 $\pm$ 0.15	0.92 $\pm$ 0.05
Vitiligo	0.78 $\pm$ 0.10	0.85 $\pm$ 0.06	0.82 $\pm$ 0.07	0.83 $\pm$ 0.10	0.86 $\pm$ 0.08	0.95 $\pm$ 0.02
Mean	0.79 $\pm$ 0.05	0.79 $\pm$ 0.06	0.78 $\pm$ 0.05	0.79 $\pm$ 0.04	0.79 $\pm$ 0.04	0.87 $\pm$ 0.08

**Table B.14**

DermX diagnostic performance with regard to the gold standard.

	F1 score	Sensitivity	Specificity
Acne	0.87 $\pm$ 0.05	0.89 $\pm$ 0.07	0.96 $\pm$ 0.01
Actinic keratosis	0.79 $\pm$ 0.14	0.74 $\pm$ 0.14	0.97 $\pm$ 0.03
Psoriasis	0.73 $\pm$ 0.11	0.77 $\pm$ 0.11	0.92 $\pm$ 0.05
Seborrheic dermatitis	0.74 $\pm$ 0.09	0.76 $\pm$ 0.10	0.93 $\pm$ 0.04
Viral warts	0.76 $\pm$ 0.11	0.68 $\pm$ 0.15	0.99 $\pm$ 0.01
Vitiligo	0.83 $\pm$ 0.10	0.82 $\pm$ 0.13	0.97 $\pm$ 0.03
Mean	0.79 $\pm$ 0.04	0.78 $\pm$ 0.04	0.96 $\pm$ 0.01

**Table B.15**

DermX+ diagnosis performance with regard to the gold standard.

	F1 score	Sensitivity	Specificity
Acne	0.86 $\pm$ 0.06	0.92 $\pm$ 0.06	0.94 $\pm$ 0.04
Actinic keratosis	0.73 $\pm$ 0.10	0.69 $\pm$ 0.15	0.96 $\pm$ 0.02
Psoriasis	0.80 $\pm$ 0.09	0.83 $\pm$ 0.08	0.95 $\pm$ 0.04
Seborrheic dermatitis	0.74 $\pm$ 0.10	0.76 $\pm$ 0.11	0.94 $\pm$ 0.03
Viral warts	0.76 $\pm$ 0.15	0.70 $\pm$ 0.18	0.98 $\pm$ 0.01
Vitiligo	0.86 $\pm$ 0.08	0.83 $\pm$ 0.11	0.98 $\pm$ 0.02
Mean	0.79 $\pm$ 0.04	0.79 $\pm$ 0.04	0.96 $\pm$ 0.01

**Table B.16**

DermX performance on the presence or absence of characteristics with regard to the dermatologist-generated labels.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.76 $\pm$ 0.08	0.79 $\pm$ 0.01	0.95 $\pm$ 0.03	96
Dermatoglyph disruption	0.74 $\pm$ 0.20	0.68 $\pm$ 0.21	0.99 $\pm$ 0.02	54
Open comedo	0.80 $\pm$ 0.06	0.82 $\pm$ 0.09	0.95 $\pm$ 0.02	110
Papule	0.79 $\pm$ 0.07	0.79 $\pm$ 0.10	0.80 $\pm$ 0.08	278
Patch	0.76 $\pm$ 0.06	0.76 $\pm$ 0.11	0.82 $\pm$ 0.06	249
Plaque	0.88 $\pm$ 0.03	0.89 $\pm$ 0.03	0.74 $\pm$ 0.11	352
Pustule	0.79 $\pm$ 0.10	0.83 $\pm$ 0.14	0.94 $\pm$ 0.02	106
Scale	0.79 $\pm$ 0.04	0.81 $\pm$ 0.07	0.77 $\pm$ 0.07	275
Scar	0.78 $\pm$ 0.10	0.80 $\pm$ 0.14	0.94 $\pm$ 0.04	115
Sun damage	0.66 $\pm$ 0.11	0.64 $\pm$ 0.15	0.96 $\pm$ 0.03	78
Mean	0.77 $\pm$ 0.03	0.78 $\pm$ 0.04	0.88 $\pm$ 0.02	171.30

**Table B.17**

DermX+ performance on the presence or absence of characteristics with regard to the dermatologist-generated labels.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.81 $\pm$ 0.12	0.87 $\pm$ 0.13	0.94 $\pm$ 0.04	96
Dermatoglyph disruption	0.70 $\pm$ 0.2	0.64 $\pm$ 0.23	0.99 $\pm$ 0.01	54
Open comedo	0.80 $\pm$ 0.08	0.85 $\pm$ 0.09	0.93 $\pm$ 0.05	110
Papule	0.80 $\pm$ 0.07	0.83 $\pm$ 0.07	0.76 $\pm$ 0.11	278
Patch	0.79 $\pm$ 0.06	0.77 $\pm$ 0.11	0.87 $\pm$ 0.08	249
Plaque	0.90 $\pm$ 0.03	0.89 $\pm$ 0.05	0.85 $\pm$ 0.06	352
Pustule	0.81 $\pm$ 0.06	0.86 $\pm$ 0.08	0.94 $\pm$ 0.03	106
Scale	0.82 $\pm$ 0.05	0.82 $\pm$ 0.07	0.83 $\pm$ 0.07	275
Scar	0.80 $\pm$ 0.08	0.82 $\pm$ 0.14	0.94 $\pm$ 0.03	115
Sun damage	0.64 $\pm$ 0.15	0.60 $\pm$ 0.17	0.96 $\pm$ 0.03	78
Mean	0.79 $\pm$ 0.03	0.80 $\pm$ 0.04	0.90 $\pm$ 0.02	171.30

**Table B.18**

DermX localization performance for localizable characteristics (mean $\pm$ std) with regard to the fuzzy dermatologist attention maps.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.40 $\pm$ 0.04	0.69 $\pm$ 0.09	0.69 $\pm$ 0.05	75
Dermatoglyph disruption	0.28 $\pm$ 0.06	0.69 $\pm$ 0.09	0.69 $\pm$ 0.05	36
Open comedo	0.36 $\pm$ 0.05	0.69 $\pm$ 0.06	0.68 $\pm$ 0.04	90
Papule	0.34 $\pm$ 0.04	0.63 $\pm$ 0.08	0.72 $\pm$ 0.05	219
Patch	0.43 $\pm$ 0.05	0.57 $\pm$ 0.07	0.78 $\pm$ 0.04	188
Plaque	0.43 $\pm$ 0.03	0.65 $\pm$ 0.06	0.75 $\pm$ 0.04	314
Pustule	0.24 $\pm$ 0.05	0.69 $\pm$ 0.07	0.69 $\pm$ 0.05	88
Scale	0.41 $\pm$ 0.03	0.65 $\pm$ 0.07	0.76 $\pm$ 0.04	222
Scar	0.46 $\pm$ 0.05	0.64 $\pm$ 0.06	0.72 $\pm$ 0.05	90
Sun damage	0.56 $\pm$ 0.06	0.44 $\pm$ 0.08	0.87 $\pm$ 0.04	50
Mean	0.39 $\pm$ 0.02	0.64 $\pm$ 0.02	0.74 $\pm$ 0.01	137.20

**Table B.19**

DermX characteristics localization performance with regard to fuzzy dermatologist attention maps. The results include the localization performance of characteristics identified by the dermatologists but not by the model.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.31 $\pm$ 0.04	0.55 $\pm$ 0.11	0.76 $\pm$ 0.05	96
Dermatoglyph disruption	0.19 $\pm$ 0.06	0.48 $\pm$ 0.17	0.79 $\pm$ 0.07	54
Open comedo	0.30 $\pm$ 0.05	0.57 $\pm$ 0.08	0.73 $\pm$ 0.05	110
Papule	0.27 $\pm$ 0.05	0.50 $\pm$ 0.08	0.78 $\pm$ 0.05	278
Patch	0.33 $\pm$ 0.05	0.44 $\pm$ 0.08	0.83 $\pm$ 0.04	249
Plaque	0.38 $\pm$ 0.02	0.58 $\pm$ 0.05	0.77 $\pm$ 0.04	352
Pustule	0.20 $\pm$ 0.05	0.58 $\pm$ 0.14	0.74 $\pm$ 0.07	106
Scale	0.33 $\pm$ 0.03	0.53 $\pm$ 0.08	0.81 $\pm$ 0.04	275
Scar	0.36 $\pm$ 0.06	0.50 $\pm$ 0.08	0.77 $\pm$ 0.06	115
Sun damage	0.36 $\pm$ 0.09	0.28 $\pm$ 0.07	0.92 $\pm$ 0.03	78
Mean	0.30 $\pm$ 0.01	0.50 $\pm$ 0.01	0.79 $\pm$ 0.00	171.30

**Table B.20**

DermX+ localization performance for localizable characteristics (mean $\pm$ std) with regard to the fuzzy dermatologist attention maps.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.10 $\pm$ 0.09	0.11 $\pm$ 0.08	0.96 $\pm$ 0.02	83
Dermatoglyph disruption	0.60 $\pm$ 0.12	0.60 $\pm$ 0.15	0.97 $\pm$ 0.02	34
Open comedo	0.06 $\pm$ 0.05	0.06 $\pm$ 0.05	0.97 $\pm$ 0.02	93
Papule	0.18 $\pm$ 0.04	0.19 $\pm$ 0.06	0.97 $\pm$ 0.02	232
Patch	0.53 $\pm$ 0.05	0.56 $\pm$ 0.06	0.89 $\pm$ 0.03	191
Plaque	0.61 $\pm$ 0.05	0.65 $\pm$ 0.06	0.91 $\pm$ 0.01	312
Pustule	0.03 $\pm$ 0.03	0.04 $\pm$ 0.05	0.98 $\pm$ 0.01	91
Scale	0.49 $\pm$ 0.07	0.58 $\pm$ 0.11	0.89 $\pm$ 0.03	224
Scar	0.35 $\pm$ 0.12	0.35 $\pm$ 0.15	0.92 $\pm$ 0.04	93
Sun damage	0.58 $\pm$ 0.20	0.53 $\pm$ 0.19	0.90 $\pm$ 0.10	47
Mean	0.36 $\pm$ 0.01	0.37 $\pm$ 0.01	0.94 $\pm$ 0.00	138.00

## Appendix A. Additional dataset information

See Figs. A.11–A.13 and Tables A.8–A.11.

## Appendix B. Model training and extended performance

See Tables B.12–B.21.

**Table B.21**

DermX+ characteristics localization performance with regard to fuzzy dermatologist attention maps. The results include the localization performance of characteristics identified by the dermatologists but not by the model.

	F1 score	Sensitivity	Specificity	Samples
Closed comedo	0.09 ± 0.08	0.09 ± 0.07	0.97 ± 0.02	96
Dermatoglyph disruption	0.39 ± 0.17	0.4 ± 0.21	0.98 ± 0.02	54
Open comedo	0.05 ± 0.04	0.05 ± 0.04	0.98 ± 0.02	110
Papule	0.15 ± 0.04	0.16 ± 0.06	0.97 ± 0.02	278
Patch	0.41 ± 0.07	0.43 ± 0.07	0.92 ± 0.02	249
Plaque	0.54 ± 0.06	0.58 ± 0.07	0.92 ± 0.01	352
Pustule	0.03 ± 0.03	0.03 ± 0.04	0.98 ± 0.01	106
Scale	0.41 ± 0.07	0.48 ± 0.09	0.91 ± 0.02	275
Scar	0.29 ± 0.1	0.29 ± 0.13	0.93 ± 0.04	115
Sun damage	0.34 ± 0.14	0.31 ± 0.15	0.94 ± 0.08	78
Mean	0.27 ± 0.01	0.28 ± 0.01	0.95 ± 0.00	171.30

## References

Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B., 2018. Sanity checks for saliency maps. In: Proceedings of the 32nd International Conference on Neural Information Processing Systems. pp. 9525–9536.

Barata, C., Celebi, M.E., Marques, J.S., 2021. Explainable skin lesion diagnosis using taxonomies. *Pattern Recognit.* 110, 107413.

Burlina, P.M., Joshi, N.J., Ng, E., Billings, S.D., Rebman, A.W., Aucott, J.N., 2019. Automated detection of erythema migrans and other confounding skin lesions via deep learning. *Comput. Biol. Med.* 105, 151–156.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. Ieee, pp. 248–255.

DermNetNZ, 2021. <https://dermnetnz.org/>, Accessed: 2021-04-01.

DeYoung, J., Jain, S., Rajani, N.F., Lehman, E., Xiong, C., Socher, R., Wallace, B.C., 2020. ERASER: A benchmark to evaluate rationalized NLP models. *Trans. Assoc. Comput. Linguist.*

Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S., 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542 (7639), 115–118.

Federman, D.G., Concato, J., Kirsner, R.S., 1999. Comparison of dermatologic diagnoses by primary care practitioners and dermatologists: a review of the literature. *Arch. Family Med.* 8 (2), 170.

Feng, H., Berk-Krauss, J., Feng, P.W., Stein, J.A., 2018. Comparison of dermatologist density between urban and rural counties in the United States. *JAMA Dermatol.* 154 (11), 1265–1271.

Goodman, B., Flaxman, S., 2017. European union regulations on algorithmic decision-making and a “right to explanation”. *AI Mag.* 38 (3), 50–57.

Graziani, M., Palatnik de Sousa, I., Vellasco, M.M., Costa da Silva, E., Müller, H., Andrearczyk, V., 2021. Sharpening local interpretable model-agnostic explanations for histopathology: Improved understandability and reliability. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 540–549.

Han, S.S., Park, I., Chang, S.E., Lim, W., Kim, M.S., Park, G.H., Chae, J.B., Huh, C.H., Na, J.-I., 2020. Augmented intelligence dermatology: deep neural networks empower medical professionals in diagnosing skin cancer and predicting treatment options for 134 skin disorders. *J. Invest. Dermatol.* 140 (9), 1753–1761.

Hay, R.J., Johns, N.E., Williams, H.C., Bolliger, I.W., Dellavalle, R.P., Margolis, D.J., Marks, R., Naldi, L., Weinstock, M.A., Wulf, S.K., et al., 2014. The global burden of skin disease in 2010: an analysis of the prevalence and impact of skin conditions. *J. Invest. Dermatol.* 134 (6), 1527–1534.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778.

Jacovi, A., Goldberg, Y., 2020. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 4198–4205.

Jain, A., Way, D., Gupta, V., Gao, Y., de Oliveira Marinho, G., Hartford, J., Sayres, R., Kanada, K., Eng, C., Nagpal, K., et al., 2021. Development and assessment of an artificial intelligence-based tool for skin condition diagnosis by primary care physicians and nurse practitioners in teledermatology practices. *JAMA Netw. Open* 4 (4), e217249.

Karimkhani, C., Dellavalle, R.P., Coffeng, L.E., Flohr, C., Hay, R.J., Langan, S.M., Nsoesie, E.O., Ferrari, A.J., Erskine, H.E., Silverberg, J.I., Vos, T., Naghavi, M., 2017. Global skin disease morbidity and mortality: An update from the global burden of disease study 2013. *JAMA Dermatol.* 153 (5), 406–412. <http://dx.doi.org/10.1001/jamadermatol.2016.5538>.

Kelly, C.J., Karthikesalingam, A., Suleyman, M., Corrado, G., King, D., 2019. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med.* 17 (1), 1–9.

Kringos, D.S., Boerma, W.G., Hutchinson, A., Saltman, R.B., Organization, W.H., et al., 2015. Building Primary Care in a Changing Europe. World Health Organization. Regional Office for Europe.

Li, K., Wu, Z., Peng, K.-C., Ernst, J., Fu, Y., 2018. Tell me where to look: Guided attention inference network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9215–9223.

Lim, H.W., Collins, S.A., Resneck Jr., J.S., Bologna, J.L., Hodge, J.A., Rohrer, T.A., Van Beek, M.J., Margolis, D.J., Sober, A.J., Weinstock, M.A., et al., 2017. The burden of skin disease in the United States. *J. Am. Acad. Dermatol.* 76 (5), 958–972.

Loshchilov, I., Hutter, F., 2016. SGDR: Stochastic gradient descent with warm restarts.

Loshchilov, I., Hutter, F., 2018. Decoupled weight decay regularization. In: International Conference on Learning Representations.

Lundberg, S.M., Lee, S.-I., 2017. A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*. Vol. 30, Curran Associates, Inc..

Mathew, B., Saha, P., Yimam, S.M., Biemann, C., Goyal, P., Mukherjee, A., 2021. Hatexplain: A benchmark dataset for explainable hate speech detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 35, (17), pp. 14867–14875.

Nast, A., Griffiths, C.E.M., Hay, R., Sterry, W., Bologna, J.L., 2016. The 2016 international league of dermatological societies’ revised glossary for the description of cutaneous lesions. *Br. J. Dermatol.* 174 (6), 1351–1358.

Oakley, A., 2017. *Dermatology Made Easy*. Scion Publishing Ltd, The Old Hayloft, Vantage Business Park, Bloxham Road, Banbury OX16 9UX, UK.

Redmon, J., Farhadi, A., 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*.

Ribeiro, M.T., Singh, S., Guestrin, C., 2016. “Why should I trust you?” explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1135–1144.

Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D., 2017. Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 618–626.

Sickles, E., D’Orsi, C., Bassett, L., et al., 2013. ACR BI-RADS mammography. In: *ACR BI-RADS® Atlas, Breast Imaging Reporting and Data System*. American College of Radiology.

Simonyan, K., Vedaldi, A., Zisserman, A., 2014. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv:1312.6034*.

Singh, A., Sengupta, S., Lakshminarayanan, V., 2020. Explainable deep learning models in medical image analysis. *J. Imaging* 6 (6), 52.

Sun, X., Yang, J., Sun, M., Wang, K., 2016. A benchmark for automatic visual classification of clinical skin disease images. In: *European Conference on Computer Vision*. Springer, pp. 206–222.

Tan, M., Le, Q., 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning*. PMLR, pp. 6105–6114.

Tan, M., Pang, R., Le, Q.V., 2020. Efficientdet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10781–10790.

Thomsen, K., Iversen, L., Titlestad, T.L., Winther, O., 2020. Systematic review of machine learning for diagnosis and prognosis in dermatology. *J. Dermatol. Treatment* 31 (5), 496–510.

Topol, E.J., 2019. High-performance medicine: the convergence of human and artificial intelligence. 25 (1), pp. 44–56.

Tschandl, P., Rinner, C., Apalla, Z., Argenziano, G., Codella, N., Halpern, A., Janda, M., Lallas, A., Longo, C., Malvehy, J., et al., 2020. Human-computer collaboration for skin cancer recognition. *Nat. Med.* 26 (8), 1229–1234.

Tschandl, P., Rosendahl, C., Kittler, H., 2018. The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci. Data* 5 (1), 1–9.

V7-Labs, 2021. Darwin V7-labs. <https://darwin.v7labs.com>, Accessed: 2021-05-01.

Zhang, J., Xie, Y., Xia, Y., Shen, C., 2019. Attention residual learning for skin lesion classification. *IEEE Trans. Med. Imaging* 38 (9), 2092–2103.