

On Interpretability of Deep Learning based Skin Lesion Classifiers using Concept Activation Vectors

Adriano Lucieri
Hochschule Pforzheim
German Research Center for Artificial
Intelligence GmbH (DFKI)
Kaiserslautern, Germany
0000-0003-1473-4745

Muhammad Naseer Bajwa
Technische Universität Kaiserslautern
German Research Center for Artificial
Intelligence GmbH (DFKI)
Kaiserslautern, Germany
0000-0002-4821-1056

Stephan Alexander Braun
Department of Dermatology
Universitätsklinikum Münster
Universitätsklinikum Düsseldorf
Münster, Germany
0000-0002-5600-0195

Muhammad Imran Malik
National University of Science and
Technology (NUST)
National Center of Artificial Intelligence
Islamabad, Pakistan
0000-0002-8079-5119

Andreas Dengel
Technische Universität Kaiserslautern
German Research Center for Artificial
Intelligence GmbH (DFKI)
Kaiserslautern, Germany
0000-0002-6100-8255

Sheraz Ahmed
Smart Data and Knowledge Services
German Research Center for Artificial
Intelligence GmbH (DFKI)
Kaiserslautern, Germany
0000-0002-4239-6520

Abstract—Deep learning based medical image classifiers have shown remarkable prowess in various application areas like ophthalmology, dermatology, pathology, and radiology. However, the acceptance of these Computer-Aided Diagnosis (CAD) systems in real clinical setups is severely limited primarily because their decision-making process remains largely obscure. This work aims at elucidating a deep learning based medical image classifier by verifying that the model learns and utilizes similar disease-related concepts as described and employed by dermatologists. We used a well-trained and high performing neural network developed by REasoning for COMplex Data (RECOD) Lab for classification of three skin tumours, i.e. Melanocytic Naevi, Melanoma and Seborrheic Keratosis and performed a detailed analysis on its latent space. Two well established and publicly available skin disease datasets, PH² and derm7pt, are used for experimentation. These datasets provide annotations for various disease-related concepts in addition to image-wise diagnosis labels. Human understandable concepts are mapped to RECOD image classification model with the help of Concept Activation Vectors (CAVs), introducing a novel training and significance testing paradigm for CAVs. Our results on an independent evaluation set clearly shows that the classifier learns and encodes human understandable concepts in its latent representation. Additionally, TCAV scores (Testing with CAVs) suggest that the neural network indeed makes use of disease-related concepts in the correct way when making predictions. We anticipate that this work can not only increase confidence of medical practitioners on CAD but also serve as a stepping stone for further development of CAV-based neural network interpretation methods.

Index Terms—Skin Lesion Classification, Medical Image Analysis, Computer-Aided Diagnosis, Explainable Artificial Intelligence, Concept Activation Vectors, Convolutional Neural Networks.

This work is funded by National University of Science and Technology (NUST), Pakistan through Prime Minister’s Programme for Development of PhDs in Science and Technology, and BMBF projects ExplAINN (01IS19074) and DeFuseNN (01IW17002).

I. INTRODUCTION

United Nations (UN) has recognised healthcare and well-being as one of the 17 Sustainable Development Goals (SDGs) to create a better future for all by 2030 [1]. However, achieving this goal requires concerted and sustained efforts in utilizing all available means to improve healthcare since many people are needlessly suffering from preventable diseases. In 2017, AI for Good [2], a UN initiative to provide a global platform for researchers, identified great potential of Artificial Intelligence (AI) to achieve these SDGs and to help solve the greatest global challenges.

Numerous remarkable studies have been conducted in the last few years successfully applying deep learning for disease classification using various medical image modalities [3]–[5]. However, the acceptance of such Computer-Aided Diagnosis (CAD) solutions with doctors and patients remains dubious at best due to the fact that the process behind learning and encoding features in latent space by computer models is not very well understood. This lack of transparency in the whole decision-making process cannot be overlooked in various critical application areas including medical diagnosis. Especially after Europe’s General Data Protection Regulations (GDPR) [6] came into effect in 2018, data subjects are entitled to *Right to Explanation* for any automated decision made by computer algorithms. It is, therefore, need of the hour to elucidate the working principle of deep learning based classifiers so that practical applications of AI in medical diagnosis can be realized expeditiously.

Compared to other fields of applications of Deep Neural Networks (DNNs), medical image analysis often presents unique challenges due to inherent complexity of this task. Manual classification of complex diseases involves recogniz-

ing subtle features and high-level concepts that are challenging to grasp without expert knowledge. Even with expert knowledge, doctors’ subjective understanding of disease biomarkers leads to low inter-expert agreement [7], [8]. Therefore, common explanation methods like visualization of saliency maps, which strongly rely on spatial divisibility of concepts, work well on common object detection tasks [9]–[11] that have well distinguishable features, but fail on more complex medical image analysis tasks.

Skin cancer is the most common type of cancer in the U.S [12]. According to a recent study [13], skin cancer related death rate forecast for U.S. in 2019 amounted to 11,650 people. These rising rates of skin cancer incidences can not only cost precious lives but also incur huge burden on healthcare systems. It is estimated that approximately 3 million people are treated annually for skin cancer in the U.S. and it costs around 8.1 billion USD [14].

In this study, we choose classification of skin diseases as a use case to understand what DNNs learn and what they rely on for their predictions in medical diagnosis. We attempt to understand if the *concepts* learnt by classifiers in complex Medical Image Analysis (MIA) tasks are similar to those used by dermatologists. We use two publicly available datasets of dermoscopic images i.e. PH² and derm7pt. These datasets are selected because they provide concept annotations in addition to image-wise diagnosis labels. Summarizing our contributions, this study presents;

- A new training and significance testing paradigm for Concept Activation Vectors (CAVs) using identically distributed data.
- Mapping of concepts learnt by a deep model, in its latent space for skin lesion classification, to dermatologically significant human-understandable concepts using CAVs.
- Analysis of contributions of different dermoscopic criteria to the predictions of deep models, revealing agreement between reasoning process of doctors and deep models.

II. RELATED WORK

Understanding the way neural networks learn and explaining their prediction behaviour are active areas of research [15], [16]. To interpret these inherently nonlinear mathematical models, three main types of methods are usually employed.

A. Saliency-Based Neural Network Explanations

Saliency-based methods for neural network explanation were among the first tools towards explainable AI and are still highly in use today. Examples for those methods are GradCAM [11], SmoothGrad [17], Integrated Gradient [18] and Layer-Wise Relevance Propagation (LRP) [19] to name a few. Since these methods create importance maps based on individual input samples, they provide only local interpretations and are unable to explain network’s decisions on a global scale.

To date, there are only very few works focusing on interpreting deep classifiers for dermoscopic images using saliency-based attribution methods [20], [21]. This might partly be due to the innate difficulty in skin lesion classification that

mandates huge amount of expert knowledge to recognize complex and subtle structures. It could also be due to a large variation in fine nuances of these structures that are hard to discern yet can drastically change diagnosis. Moreover, the visual artefacts corresponding to various diseases in skin lesion images sometimes overlap and are usually distributed all over the image, which does not fare well with saliency-based model interpretation. Other domains in which these methods are frequently used for interpretation like object classification [9]–[11] usually show more discriminatory features corresponding to specific parts of an object – for instance tires or headlights of cars.

B. Text-Based Neural Network Explanations

Textual explanation methods for neural networks can be either template-based [22]–[24] that generate justifications from some auxiliary information or rule-based [25]–[28], where a classifier is trained with images as well as additional natural language explanations. Zhang et al. [27] proposed a unified network called MDNet following a rule-based approach that generates diagnostic reports along with corresponding attention maps of input images in order to increase the semantic and visual interpretability of MIA task at hand. Jing et al. [28] proposed a multi-task learning framework that is able to localize abnormal regions in medical images, predict tags and generate their descriptions.

C. Concept-Based Neural Network Explanations

Concept-based explanations address the problem of explaining black-box models by finding human-understandable concepts in neural networks’ latent representation. Kim et al. [29] introduced *Concept Activation Vectors* (CAVs) that are used to map human-understandable concepts to latent representation learnt by models in a supervised way using general human concept patches that were taken from various domains. By calculating these main concept directions and by leveraging directional derivatives, they were able to quantify the influence of a concept to the prediction of single output classes. Zhou et al. [30] followed a similar approach by decomposing neural networks’ activations into semantically meaningful components pre-trained from a large concept corpus. Graziani et al. [31], [32] restated the problem of classification in CAVs and employed regression instead. They applied their so-called *Regression Concept Vectors* (RCVs) on problems from medical domain like binary classification of breast cancer histopathology slides and classification of Retinopathy of Prematurity (ROP) states. As an extension to the work in [29], Ghorbani et al. [33] recently developed a method for unsupervised clustering of object datasets by first applying segmentation of single objects and then clustering activations of object patches into semantically meaningful clusters.

To the best of our knowledge, concept-based explanation methods have not previously been explored for skin lesion classification networks. Due to the nature of this problem,

not all of the previously described methods can be directly applied on this task. Unsupervised clustering as used in [33], for example, is not suitable in skin lesions as there is a huge spatial concept overlap and thus no possibility for distinct part segmentation. RCVs are also not applicable as skin lesion concepts are hardly quantifiable. The method in [30] requires a concept corpus which is not readily available for this specific task. Any type of textual explanation generation is also not applicable, as no diagnostic reports or descriptions of diagnosis are provided with any public dermoscopic skin lesion dataset. The computation of CAVs as seen in [29] requires patches corresponding to general human-understandable concepts. In this work we adopt the TCAV method to the problem of skin lesion classification. Instead of providing general, out-of-distribution concept patches, we train CAVs using samples from identically distributed datasets to map human-understandable concepts to the network’s latent space.

III. BACKGROUND

This section briefly describes CAVs and the method of calculating TCAV scores used in this work to quantify the contribution of a concept to DNN’s prediction. Moreover, dermoscopic concepts explaining the classifier’s decisions are also introduced.

A. Concept Activation Vectors

In order to achieve human-centered interpretability of DNNs, Kim et al. [29] introduced *Concept Activation Vectors*. A CAV, \vec{v}_c , is a vector in the embedding space of a neural network pointing into the direction that encodes the given concept c . CAVs can be calculated by training a binary concept classifier dividing samples containing a given concept from samples where the concept is absent. The CAV is then defined as the normal to the hyperplane separating the two classes.

a) TCAV Score: The metric introduced in [29] to estimate the influence of a CAV on a class of input images is the TCAV score. It makes use of directional derivatives $S_{C,k,l}(x)$ to measure the contextual sensitivity of a concept towards an entire input class, therefore providing global explanations. The TCAV score is given by:

$$TCAV_{Q_{C,k,l}} = \frac{|x \in X_k : S_{C,k,l}(x) > 0|}{|X_k|}, \quad (1)$$

where X_k denotes all inputs, k represents the class labels and $S_{C,k,l}(x)$ the directional derivative of a sample’s activation x from layer l with respect to class k and concept C . The TCAV score effectively measures the ratio of class k ’s inputs, that are positively affected by concept C without taking any magnitude into account. As compared to saliency maps or other per-feature metrics, the TCAV score allows for quantitative evaluation of concepts on whole input classes.

B. Dermoscopic Concepts used for Analysis

The concepts used in this work to interpret the deep classifier are briefly defined below in accordance with standardized

terminology agreed upon by expert dermatologist in 3rd Consensus Conference of the International Society of Dermoscopy (IDS) [34]. Fig. 1 depicts examples of some concepts listed below.

1) Pigment Networks: Pigment Networks consist of interconnected pigmented lines forming a gridlike pattern. Depending on the subtype of Pigment Network, it can either have minimal variability in colour, thickness and spacing of the lines, forming a symmetric grid (Typical Pigment Network) or have greater variability in colour, thickness and spacing of the lines, forming an asymmetric grid (Atypical Pigment Network). Apart from those two general types, more subtypes are also defined in literature. Atypical Pigment Networks can be a clue for Melanoma (although many dysplastic naevi also have atypical networks) whereas typical Pigment Networks normally indicate benign melanocytic lesions (Naevi).

2) Streaks: Streaks describe an abnormality of the lesion that can either have the form of pure straight radial extensions, radial extensions with bulbous and often kinked projections on their ends, or a widening of broken lines with incomplete connections. Streaks are referred to as irregular if they are irregularly distributed along the edge of the lesion and are brown-black in colour [35]. Regular Streaks indicate benign lesions and Irregular Streaks are clues for malignant Melanoma.

3) Regression Structures: Regression Structures are characterized by the appearance of either areas of fine, grey-blue dots, or areas of skin whiter than the surrounding normal-looking skin without blood vessels or shiny-white structures. Its presence is highly indicative of melanoma [35].

4) Dots and Globules: Dots are small structures of pigmented areas clustered in any distribution on or around the lesion. Dots clustered in center regions or on the network lines are referred to as regular, otherwise irregular. Globules are round, oval or polygonal structures larger than dots that can have high variability in colour, size and shape along with asymmetric distribution (Irregular Globules) or minimal variability along with symmetric distribution (Regular Globules). Regular Dots and Globules are indicators for benign melanocytic lesions and irregular Dots and Globules indicate melanoma [35].

5) Blue-Whitish Veils: Blue-Whitish Veils describe an irregularly shaped, structureless blotch on the lesion area that is characterized by a blue hue with an overlying whitish ground-glass haze. In [36] it is rated as the most useful single diagnostic indicator for melanoma.

6) Asymmetry: Asymmetry is the most important factor in malignant melanoma identification using ABCD rule [37]. In our work, asymmetry refers to an asymmetrical lesion contour as well as asymmetrical distributions of structures and colours within a lesion [38]. The asymmetry concept is further divided into symmetric and asymmetric in one or two axes.

7) Colour: This concept refers to colour present within the lesion area. As the appearance of single colours is not yet indicative of any diagnosis, a combined concept of *three or more colours* is used in the analysis. The presence of




Fig. 1: Some exemplary cases of skin lesion concepts from derm7pt dataset

three or more colours increases the probability of melanoma drastically [35].

The intricate explanations of concepts given above along with the concepts' innate variability offer much room for interpretation, implying the complexity of the problem itself. This is evident by the fact that even doctors tend to have notable disagreements when it comes to diagnosis, localization or identification of concept [7], [8].

IV. MATERIALS & METHOD

A. Model

The model used in this work as the basis for our exploration and experimentation is developed by the University of Campinas in Brazil. Their RECOD Lab (REasoning for COMplex Data) made their submission [39] to the IEEE International Symposium on Biomedical Imaging (ISBI) 2017 challenge and is publicly available on github¹. By applying a transfer learning approach combined with extensive ensembling using an SVM meta-layer on top of seven base models trained on different data subsets, they achieved best Area Under Receiver Operating Characteristic (ROC) Curve (AUC) for Melanoma (MEL) classification (87.4%), 3rd best AUC for Seborrheic Keratosis (SK) classification (94.3%), and 3rd best combined/mean AUC (90.8%) in part 3 of 2017 challenge. In this work, we intentionally refrained from training our own skin lesion classification model as our primary objective was explainability of these models instead of their classification performance. Thus, for our experimentation we only focused on a single module from RECOD's well-trained architecture. We used one of the base models² with Inception v4 [40] architecture, subsequently referred to as *model* or *network*. This base model was trained on RECOD's "deploy" set of 9,640 images using per-image normalization.

B. Datasets

The datasets used for concept training are PH² dataset [38] and Seven-Point Checklist Dermatology dataset abbreviated as derm7pt [41].

The PH² dataset consists of 200 dermoscopic images of melanocytic lesions, including 80 common naevi, 80 atypical naevi, and 40 melanomas. Along with the images, colour and lesion segmentation masks are provided as well as extensive well-curated annotations as seen in Table I. The derm7pt dataset consists of 1,011 clinical and dermoscopic images. Each sample is assigned to either a miscellaneous class or one of 4 diagnosis classes. Two of these diagnosis classes i.e. Melanoma and Naevi (NV) are further divided into 13 subclasses. From this dataset, only MEL and NV samples have been considered, resulting in 823 images. SK samples have been discounted due to their low count of only 45 samples. Table I provides an overview of number of samples for each concept class.

For evaluation purposes, the original ISBI 2017 challenge dataset [42] is used. The train set of ISBI 2017 challenge contains 1372 samples of NV, 374 samples of MEL and 254 samples of SK whereas test set contains 393 images of NV,

TABLE I: Distribution of image samples into different concept classes in PH² and derm7pt datasets. Note that PH² dataset does not distinguish between regular and irregular streaks.

Concepts	Presentation	Abbreviation	PH ² [38]	derm7pt [41]
Pigment Network	Typical	PN_T	84	335
	Atypical	PN_AT	116	216
		PN	N/A	551
Streaks	Regular	ST_R	N/A	96
	Irregular	ST_IR	N/A	237
		ST	30	333
Regression Structures		RS	25	233
Dots & Globules	Regular	DG_R	54	300
	Irregular	DR_IR	59	390
		DG	113	690
Blue-Whitish Veils		BWV	36	182
Asymmetry	1-Axis	Asym_1	31	N/A
	2-Axis	Asym_2	52	N/A
		Sym	117	N/A
Colours	3 or more	C_3	39	N/A
Total Samples			200	823

¹<https://github.com/learningtitans/isbi2017-part3>

²*checkpoint.rc25 of RECOD model*

117 images of MEL and 90 images of SK.

In order to verify statistical significance of our results, we trained random CAVs to compare against our concept CAVs. For this purpose, random concept labels are assigned to a subset of the ISIC archive³ images, excluding MEL and NV classes, resulting in 2,870 samples. The idea behind leaving out those two classes is that remaining samples hardly contain concepts similar to the ones used for concept training.

C. Experimental Setup

As previously described, all experiments have been conducted on one of the Inception v4 base models from [39]. For each concept, binary classifiers are trained on network’s activations to find the concepts’ directions in the embedding space. The training and evaluation scheme is depicted in Fig. 2. First, activations are extracted from *mixed_6h* layer of the Inception v4 model using PH² and *derm7pt* datasets. A clustering-based under-sampling technique as well as stratified splitting is applied to ensure evenly balanced train and validation sets for each binary concept training. Second, TCAV score is used to evaluate a concept’s importance to a specific target class. To account for differences in pre-processing and classifier initialization, each classifier training is repeated 20 times on a randomly sampled dataset split, resulting in different CAVs and different TCAV scores.

To check statistical significance of learned concepts, we trained additional 50 random CAVs per layer. The random datasets are produced by repeatedly sampling 1,000 random images from ISIC archive subset described in section IV-B and assigning random binary labels. The distribution of random

concept TCAV scores and actual concept TCAV scores is then compared by conducting a two-sided *t*-test with $\alpha = 0.05$ to assure significance of the found CAVs. In the results section, statistical insignificance is represented by red asterisks on top of the plotted bars.

V. RESULTS AND ANALYSIS

The lack of quantifiability in most applicable explanation methods does not allow for proper comparison with previous approaches. Hence, we will focus on quantitative evaluation of concept classifier’s accuracies and TCAV scores as well as a qualitative analysis of the resulting CAVs.

Fig. 3a shows all mean validation accuracies achieved by individual binary concept classifiers along with their standard deviation trained on *derm7pt* embeddings from *mixed_6h* layer. Mean baseline results from training on 50 random concept subsets are depicted by horizontal red line along with light red surrounding area marking standard deviation. It is evident from the figure that all concept classifiers achieved significantly higher validation accuracies than random baseline. The overall accuracies achieved might not seem very high, however it has to be mentioned here that computation of CAVs requires the use of linear classifiers to calculate normal vector to decision hyperplane. The results are clear evidence that network’s latent space is structured in a way that allows activation’s separation with respect to similar concepts.

Fig. 3b shows the classifiers’ validation accuracy trained on PH² dataset embeddings from *mixed_6h* layer. It is notable that many concepts achieved relatively mediocre accuracies near the random baseline. This can be explained by very small number of positive concept samples available in PH² dataset.

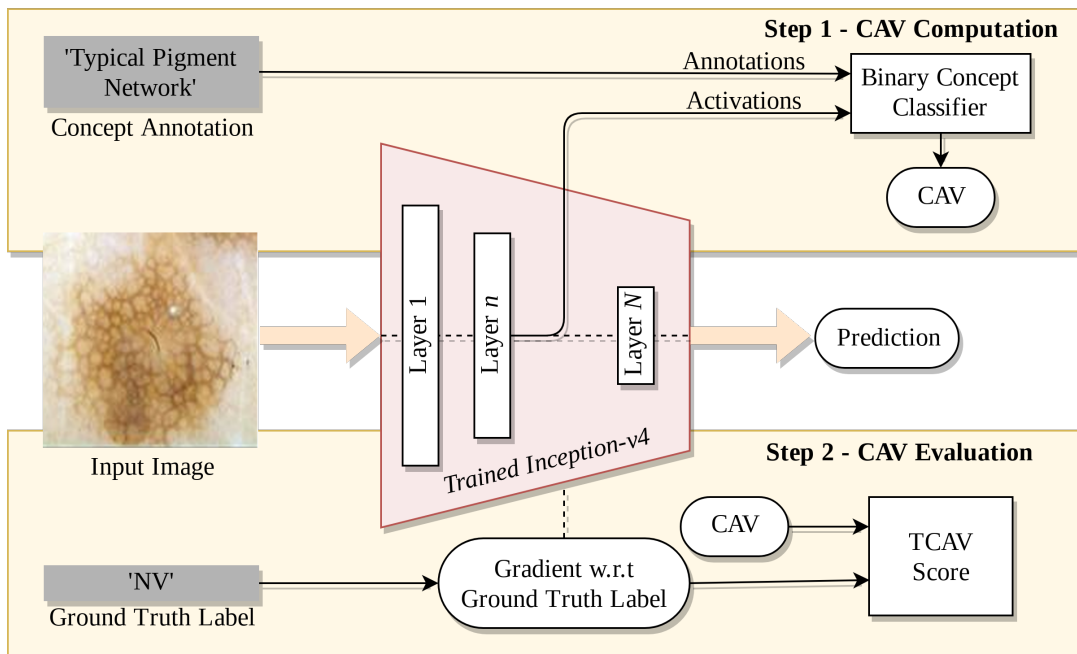




Fig. 2: Overview of training concept classifiers and calculating CAV and TCAV scores

³<https://isic-archive.com/> retrieved in November 2019



(a) Derm7pt dataset



(b) PH² dataset

Fig. 3: Validation accuracies of all concept classifiers trained and tested individually on derm7pt and PH². Random baseline is denoted by horizontal red line along with light red area marking standard deviation. Insignificant classifiers are marked with a red asterisk.

The TCAV score quantifies positive or negative influence of a given concept towards a specific target class. Values above 0.5 indicate positive influence of the concept to the prediction and lower values indicate negative influence. Figures 4 and 5 show the TCAV scores achieved by evaluating 20 CAVs per

concept on the *mixed_6h* layer separately trained on both datasets. Average baseline scores of all 50 random concepts are again depicted by red horizontal lines along with their standard deviation in light red. Statistically insignificant results are marked by red asterisks.




Fig. 4: TCAV scores of each concept for derm7pt with respect to each target class on *mixed_6h* layer of RECOD model.




Fig. 5: TCAV scores of each concept for PH^2 with respect to each target class on *mixed_6h* layer of RECOD model.

Results for NV and MEL classes for concepts trained using *derm7pt* look very much as expected. Although, the score for *PN* turned out to be insignificant in one experiment, features indicating benign melanocytic lesions like *PN_T*, *ST_R* and *DG_R* all contributed positively towards NV class. On the other hand, strong signs for malignant melanoma like *PN_AT*, *ST_IR*, *RS*, *DG_IR* and *BWV* show strong negative influence. Also, it is notable that the presence of Streaks in general (*ST*) has a stronger negative influence as compared to the presence of particularly regular Streaks (*ST_R*). Results for MEL class show the exact opposite behaviour, which is perfectly aligned with the descriptions in medical literature. It is again notable that the presence of Dots and Globules (*DG*) and the presence of Streaks (*ST*) show higher positive impact on MEL class as compared to their regular variants, for example, regular Streaks (*ST_R*). Results for SK class show similar concept influence as for MEL, except for Pigment Networks (*PN*) exhibiting negative influence. In [43] the appearance of network-like structures in Seborrheic Keratosis has been confirmed. The model might have encoded those structures in the atypical Pigment Network (*PN_AT*) concept, as their appearance slightly differs from the classical Pigment Networks definition. In the same study, evidence for Dots and blue-gray areas in SK lesions have been found as well.

Fig. 5 shows resulting TCAV scores for CAVs trained on PH^2 dataset. All concepts achieving less than 55% validation accuracy have not been considered. Again, TCAV scores for NV and MEL show expected behaviour. Only typical Pigment Networks (*PN_T*), regular Dots and Globules (*DG_R*) and Symmetry (*Sym*) contribute positively towards Naevi class. For melanoma, the exact opposite holds again which can be

confirmed by the concept descriptions in Section III-B. Additionally, from the results it appears that asymmetric lesions (*Asym_2*) and lesions containing more than three colours (*C_3*) tend to be classified as melanoma. For SK we can again observe low influence of typical Pigment Networks (*PN_T*) as well as high influence for all other concepts including asymmetry (*Asym_2*) and colour diversity (*C_3*).

To further validate that the model has comprehensively learnt these disease-related concepts instead of learning some random concepts, we had our model sort all the test images with respect to degree of visibility of a certain concept in each image. In other words, the model ordered all 300 test images starting from those that presented very obvious existence of a concept and ending with those which had least evidence of that concept. This ordering is performed based on euclidean distance in a CAV's direction. Fig. 6 through Fig. 8 show first five and last five images from the sorted test set with respect to different concepts. The first row of each figure shows positive examples, where the concept is most clearly visible, and the second row shows negative examples, where the concept is virtually absent. It is evident from these figures that the proposed method for explaining skin disease classifiers does not only provide justification of classifier's decision on global dataset scale but also sensibly identifies reasons for per-image predictions.

VI. CONCLUSION

Concept-based methods for network explanation offer great potential especially for complex classification tasks in critical application areas like MIA. This work strives to leverage these methods to verify the ability of DNNs to learn and utilize




Fig. 6: First row shows the test images that are the most similar to Typical Pigment Network (PN_T) concept according to euclidean distance in CAV direction. The second row shows the least similar images to PN_T .




Fig. 7: First row shows the test images that are the most similar to Irregular Streaks (ST_{IR}) concept according to euclidean distance in CAV direction. The second row shows the least similar images to ST_{IR} .




Fig. 8: First row shows the test images that are the most similar to Regression Structure (RS) concept according to euclidean distance in CAV direction. The second row shows the least similar images to RS .

human understandable concepts for skin lesion classification. Our results show strong correlation between DNN’s learnt representation of various concepts and those routinely used by dermatologists. This work corroborates that deep learning based CAD systems are able to learn and utilize similar disease-related concepts for prediction as used by dermatologists. We hope that physicians would be able to confer higher confidence to such CAD systems that are able to justify their prediction by listing the concepts which influenced positively or negatively towards a certain output. It has also been shown that Testing with CAVs (TCAV) is applicable using complete identically distributed images instead of general concept patches. However, this work can further be improved by using more granular labelling of diseases indicative concepts to get deeper insight into model’s classification processes as well as further validation of its decisions.

Due to the complexity of the problem, possibly subjective annotations of training set by various expert, and small number of concept training samples, not all human-defined disease-related concepts were thoroughly analysed. Standardizing the annotation according to one *school of thought* in dermatology community, for example following [34], can decrease inter-observer disagreement but it would require an enormous amount of time and effort by dermatology experts. Nevertheless, the obtained results are remarkably aligned with common diagnostic criteria. In order to allow for a more comprehensive interpretation of the TCAV scores for this specific task, it would be desirable to curate a high-quality dataset with reliable fine-grained labels of concepts that are known to be highly indicative of specific diagnoses.

Supervised concept learning is highly dependent on high quality and precisely annotated human concept examples. Therefore, more focus should be placed on generating clean datasets of high-quality concept annotations that can be used for explaining models in medical imaging applications to speed up their deployments in clinical routines. As supervised concept classification from network activations has already been proven to be effective, the following logical and highly desirable extension of unsupervised concept discovery should be considered. Not only would this be another improvement towards simplifying the interpretability of networks by eliminating the necessity for laborious expert annotations but it could also allow insights in a network’s own concepts, potentially revealing new knowledge for domain experts or unexpected biases in the network.

REFERENCES

- [1] United Nations General Assembly, “Transforming our World: The 2030 Agenda for Sustainable Development,” available at <https://sustainabledevelopment.un.org/post2015/>, October 2015, a/RES/70/1.
- [2] H. Zhao and M. Shingles, “AI for Good: Global Summit,” June 2017.
- [3] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun, “Dermatologist-level classification of skin cancer with deep neural networks,” *Nature*, vol. 542, no. 7639, p. 115, 2017.
- [4] M. N. Bajwa, M. I. Malik, S. A. Siddiqui, A. Dengel, F. Shafait, W. Neumeier, and S. Ahmed, “Two-stage framework for optic disc localization and glaucoma classification in retinal fundus images using deep learning,” *BMC medical informatics and decision making*, vol. 19, no. 1, p. 136, 2019.
- [5] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. Van Der Laak, B. Van Ginneken, and C. I. Sánchez, “A survey on deep learning in medical image analysis,” *Medical image analysis*, vol. 42, pp. 60–88, 2017.
- [6] Council of the European Union, “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation),” available at <http://data.europa.eu/eli/reg/2016/679/2016-05-04>, April 2016.
- [7] A. B. Fortina, E. Peserico, A. Silletti, and E. Zattra, “Where’s the naevus? inter-operator variability in the localization of melanocytic lesion border,” *Skin Research and Technology*, vol. 18, no. 3, pp. 311–315, 2012.
- [8] R. Corona, A. Mele, M. Amini, G. De Rosa, G. Coppola, P. Piccardi, M. Fucci, P. Pasquini, and T. Faraggiana, “Interobserver variability on the histopathologic diagnosis of cutaneous melanoma and other pigmented skin lesions,” *Journal of clinical Oncology*, vol. 14, no. 4, pp. 1218–1223, 1996.
- [9] A. Ghorbani, D. Ouyang, A. Abid, B. He, J. H. Chen, R. A. Harrington, D. H. Liang, E. A. Ashley, and J. Y. Zou, “Deep learning interpretation of echocardiograms,” *bioRxiv*, p. 681676, 2019.
- [10] S. Jolly, B. K. Iwana, R. Kuroki, and S. Uchida, “How do convolutional neural networks learn design?” in *2018 24th International Conference on Pattern Recognition (ICPR)*. IEEE, 2018, pp. 1085–1090.
- [11] R. R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra, “Grad-cam: Why did you say that?” *arXiv preprint arXiv:1611.07450*, 2016.
- [12] American Cancer Society, “Cancer facts & figures 2017,” 2017, <https://www.cancer.org/content/dam/cancer-org/research/cancer-facts-and-statistics/annual-cancer-facts-and-figures/2017/cancer-facts-and-figures-2017.pdf>.
- [13] R. L. Siegel, K. D. Miller, and A. Jemal, “Cancer statistics, 2019,” *CA: a cancer journal for clinicians*, vol. 69, no. 1, pp. 7–34, 2019.
- [14] E. S. Ruiz, F. C. Morgan, C. M. Zigler, R. J. Besaw, and C. D. Schmults, “Analysis of national skin cancer expenditures in the united states medicare population, 2013,” *Journal of the American Academy of Dermatology*, vol. 80, no. 1, pp. 275–278, 2019.
- [15] S. Palacio, J. Folz, J. Hees, F. Raue, D. Borth, and A. Dengel, “What do deep networks like to see?” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3108–3117.
- [16] W. Samek, T. Wiegand, and K.-R. Müller, “Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models,” *arXiv preprint arXiv:1708.08296*, 2017.
- [17] D. Smilkov, N. Thorat, B. Kim, F. Viégas, and M. Wattenberg, “Smoothgrad: removing noise by adding noise,” *arXiv preprint arXiv:1706.03825*, 2017.
- [18] M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks,” in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 3319–3328.
- [19] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation,” *PLoS one*, vol. 10, no. 7, p. e0130140, 2015.
- [20] K. Young, G. Booth, B. Simpson, R. Dutton, and S. Shrapnel, “Deep neural network or dermatologist?” *arXiv preprint arXiv:1908.06612*, 2019.
- [21] J. Wu, X. Li, E. Z. Chen, H. Jiang, X. Dong, and R. Rong, “What evidence does deep learning model use to classify skin lesions?” *arXiv preprint arXiv:1811.01051*, 2018.
- [22] M. Munir, S. A. Siddiqui, F. Küsters, D. Mercier, A. Dengel, and S. Ahmed, “Tsxplain: Demystification of dnn decisions for time-series using natural language and statistical features,” *arXiv preprint arXiv:1905.06175*, 2019.
- [23] P. Guo, C. Anderson, K. Pearson, and R. Farrell, “Neural network interpretation via fine grained textual summarization,” *arXiv preprint arXiv:1805.08969*, 2018.
- [24] L. Anne Hendricks, R. Hu, T. Darrell, and Z. Akata, “Grounding visual explanations,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 264–279.
- [25] B. Hancock, M. Bringmann, P. Varma, P. Liang, S. Wang, and C. Ré, “Training classifiers with natural language explanations,” in *Proceedings*

- of the conference. Association for Computational Linguistics. Meeting, vol. 2018. NIH Public Access, 2018, p. 1884.
- [26] S. Srivastava, I. Labutov, and T. Mitchell, "Joint concept learning and semantic parsing from natural language explanations," in *Proceedings of the 2017 conference on empirical methods in natural language processing*, 2017, pp. 1527–1536.
- [27] Z. Zhang, Y. Xie, F. Xing, M. McGough, and L. Yang, "Mdnet: A semantically and visually interpretable medical image diagnosis network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 6428–6436.
- [28] B. Jing, P. Xie, and E. Xing, "On the automatic generation of medical imaging reports," *arXiv preprint arXiv:1711.08195*, 2017.
- [29] B. Kim, M. Wattenberg, J. Gilmer, C. Cai, J. Wexler, F. Viegas, and R. Sayres, "Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav)," *arXiv preprint arXiv:1711.11279*, 2017.
- [30] B. Zhou, Y. Sun, D. Bau, and A. Torralba, "Interpretable basis decomposition for visual explanation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 119–134.
- [31] M. Graziani, V. Andrearczyk, and H. Müller, "Regression concept vectors for bidirectional explanations in histopathology," in *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer, 2018, pp. 124–132.
- [32] M. Graziani, J. M. Brown, V. Andrearczyk, V. Yildiz, J. P. Campbell, D. Erdogmus, S. Ioannidis, M. F. Chiang, J. Kalpathy-Cramer, and H. Müller, "Improved interpretability for computer-aided severity assessment of retinopathy of prematurity," in *Medical Imaging 2019: Computer-Aided Diagnosis*, vol. 10950. International Society for Optics and Photonics, 2019, p. 109501R.
- [33] A. Ghorbani, J. Wexler, and B. Kim, "Automating interpretability: Discovering and testing visual concepts learned by neural networks," *arXiv preprint arXiv:1902.03129*, 2019.
- [34] H. Kittler, A. A. Marghoob, G. Argenziano, C. Carrera, C. Curiel-Lewandrowski, R. Hofmann-Wellenhof, J. Malvehy, S. Menzies, S. Puig, H. Rabinovitz *et al.*, "Standardization of terminology in dermoscopy/dermatology: Results of the third consensus conference of the international society of dermoscopy," *Journal of the American Academy of Dermatology*, vol. 74, no. 6, pp. 1093–1106, 2016.
- [35] G. Argenziano, H. P. Soyer, S. Chimenti, R. Talamini, R. Corona, F. Sera, M. Binder, L. Cerroni, G. De Rosa, G. Ferrara *et al.*, "Dermoscopy of pigmented skin lesions: results of a consensus meeting via the internet," *Journal of the American Academy of Dermatology*, vol. 48, no. 5, pp. 679–693, 2003.
- [36] S. Menzies, C. Ingvar, and W. McCarthy, "A sensitivity and specificity analysis of the surface microscopy features of invasive melanoma." *Melanoma research*, vol. 6, no. 1, pp. 55–62, 1996.
- [37] R. J. Friedman, D. S. Rigel, and A. W. Kopf, "Early detection of malignant melanoma: the role of physician examination and self-examination of the skin," *CA: a cancer journal for clinicians*, vol. 35, no. 3, pp. 130–151, 1985.
- [38] T. Mendonça, P. M. Ferreira, J. S. Marques, A. R. Marcal, and J. Rozeira, "Ph 2-a dermoscopic image database for research and benchmarking," in *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*. IEEE, 2013, pp. 5437–5440.
- [39] A. Menegola, J. Tavares, M. Fornaciali, L. T. Li, S. Avila, and E. Valle, "Recod titans at isic challenge 2017," *arXiv preprint arXiv:1703.04819*, 2017.
- [40] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [41] J. Kawahara, S. Daneshvar, G. Argenziano, and G. Hamarneh, "Seven-point checklist and skin lesion classification using multitask multimodal neural nets," *IEEE Journal of Biomedical and Health Informatics*, vol. 23, no. 2, pp. 538–546, mar 2019.
- [42] N. C. Codella, D. Gutman, M. E. Celebi, B. Helba, M. A. Marchetti, S. W. Dusza, A. Kalloo, K. Liopyris, N. Mishra, H. Kittler *et al.*, "Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic)," in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*. IEEE, 2018, pp. 168–172.
- [43] R. P. Braun, H. S. Rabinovitz, J. Krischer, J. Kreusch, M. Oliviero, L. Naldi, A. W. Kopf, and J. H. Saurat, "Dermoscopy of pigmented seborrheic keratosis: a morphological study," *Archives of dermatology*, vol. 138, no. 12, pp. 1556–1560, 2002.