

# When Medical Imaging Met Self-Attention: A Love Story That Didn't Quite Work out

Tristan Piater<sup>a</sup>, Niklas Penzel<sup>b</sup>, Gideon Stein<sup>c</sup> and Joachim Denzler<sup>d</sup>

*Computer Vision Group, Friedrich Schiller University, Jena, Germany*

**Keywords:** Self-Attention Mechanisms, Feature Analysis, Medical Imaging.

**Abstract:** A substantial body of research has focused on developing systems that assist medical professionals during labor-intensive early screening processes, many based on convolutional deep-learning architectures. Recently, multiple studies explored the application of so-called self-attention mechanisms in the vision domain. These studies often report empirical improvements over fully convolutional approaches on various datasets and tasks. To evaluate this trend for medical imaging, we extend two widely adopted convolutional architectures with different self-attention variants on two different medical datasets. With this, we aim to specifically evaluate the possible advantages of additional self-attention. We compare our models with similarly sized convolutional and attention-based baselines and evaluate performance gains statistically. Additionally, we investigate how including such layers changes the features learned by these models during the training. Following a hyperparameter search, and contrary to our expectations, we observe no significant improvement in balanced accuracy over fully convolutional models. We also find that important features, such as dermoscopic structures in skin lesion images, are still not learned by employing self-attention. Finally, analyzing local explanations, we confirm biased feature usage. We conclude that merely incorporating attention is insufficient to surpass the performance of existing fully convolutional methods.

## 1 INTRODUCTION

Computer vision models can aid medical practitioners by providing visual analysis of skin lesions or tumor tissues. In both cases, a correct classification can help save lives since the survival rate of cancer patients increases drastically when their condition is detected early. Malignant melanomata, for example, still result in over 7000 deaths in the US per year (Society, 2022). Hence, a substantial body of research, e.g., (Tschandl et al., 2018; Mishra and Celebi, 2016; Codella et al., 2019; Celebi et al., 2019; Lee and Paeng, 2018; Bera et al., 2019; Khened et al., 2021), has been focused on developing automated systems that examine medical images to assist medical practitioners with the demanding screening process. Usually, those models are built on convolutional neural networks (CNNs), a well-established deep-learning archetype.

Previous work investigated convolutional neural network (CNN) models and their feature usage (Reimers et al., 2021; Penzel et al., 2022) specifically for skin lesion classification and found that these networks learn only a subset of medically relevant features. To be specific, of the dermatological ABCD rule (Nachbar et al., 1994), only **A**symmetry and **B**order irregularity were incorporated consistently in the prediction process of state-of-the-art melanoma models. Further, **C**olor and **D**ermoscopic structures were often not learned. Additionally, modern automatic skin lesion classifiers often overfit on biases contained in the training data (Reimers et al., 2021), e.g., spurious colorful patches (Scope et al., 2016).

Recently, emerging from natural language processing (Vaswani et al., 2017), multiple studies explored the application of so-called self-attention mechanisms in the vision domain (Kolesnikov et al., 2021; Wang et al., 2018; Ramachandran et al., 2019). These studies often report empirical improvements over state-of-the-art approaches on various datasets and tasks. Following the intuition that self-attention mechanisms possibly allow for a better comprehension of global features, e.g., dermoscopic structures

<sup>a</sup> <https://orcid.org/0009-0008-1938-6261>

<sup>b</sup> <https://orcid.org/0000-0001-8002-4130>

<sup>c</sup> <https://orcid.org/0000-0002-2735-1842>

<sup>d</sup> <https://orcid.org/0000-0002-3193-3300>

in skin lesions or number of cells in tumor tissue, previous approaches directly evaluated popular transformer architectures on medical image classification tasks (Yang et al., 2023; He et al., 2022; Krishna et al., 2023). While they typically report an improvement in accuracy, it is not self-evident that the observed improvement is due to the attention mechanisms and not other circumstances, e.g., possibly increased parameter counts or altered input representation (Trockman and Kolter, 2022; Tolstikhin et al., 2021) which often go hand in hand when deploying transformer architectures. Furthermore, we find that the question of whether self-attention mechanisms can help to learn additional medically relevant features is understudied. To answer these questions, we study the possible benefits of self-attention mechanisms on two medical classification tasks: skin lesion classification (ISIC, 2022) and tumor tissue classification (Bandi et al., 2018). Specifically, we extend two widely adopted CNN-backbones: ResNet (He et al., 2016) and EfficientNet (Tan and Le, 2019) with different self-attention variants and evaluate in-distribution and out-of-distribution performance. Further, we compare our models with convolutional and vision transformer baselines. In the literature, two strategies have been proposed for incorporating self-attention into CNNs. The global self-attention approach (Wang et al., 2018) and the local self-attention block (Ramachandran et al., 2019), which we adopt and extend. Furthermore, we also investigate combinations of these approaches. For all our architectures, we perform hyperparameter searches and report the parameter count to enable fair comparisons.

Contrary to our expectations, we found that the tested self-attention mechanisms do not significantly improve the performance of medical image classifiers, nor do vision transformers when kept at a similar parameter count as the convolutional baselines. In contrast, we sometimes measure a significant decrease in performance.

With these results in mind, we derive global and local explanations of our models' behavior. To be specific, for the skin lesion models using the global explainability method described in (Reimers et al., 2020). Here, we directly follow the approaches described in (Reimers et al., 2021; Penzel et al., 2022). We find that adding self-attention does generally not help in learning medical-relevant features. While sometimes additional medically relevant features are learned, this is often accompanied by relying on more biases.

Furthermore, we derive local explanations for a selection of images using both Grad-CAM (Selvaraju et al., 2017) as well as visualizing the learned global

attention maps where applicable. While self-attention natively provides local explanations that can lead to the discovery of structural biases, we find that out-of-the-box methods, we used Grad-CAM (Selvaraju et al., 2017), can provide very similar insights. Also, observing these insights, we find that the global self-attention maps in our experiments often reveal unwanted attention spikes and artifacts in the background. Similar behavior was also observed for other transformer models, e.g., in (Darcet et al., 2023; Xiao et al., 2023).

To conclude: our empirical analysis suggests that merely incorporating attention is insufficient to surpass the performance of existing fully convolutional methods.

## 2 APPROACH

To explain our setup, we first describe both, global (Wang et al., 2018) and local (Ramachandran et al., 2019), approaches to extend CNNs with self-attention methods. Further, we propose another local self-attention implementation for a smoother integration into pre-trained networks. Afterward, we detail how we include those attention methods into two selected CNNs: ResNet18 (He et al., 2016) and EfficientNet-B0 (Tan and Le, 2019).

### 2.1 Global Self-Attention

We implement the global self-attention (GA) method following the details in (Wang et al., 2018). To add further details, it is a global operation where each of the output values depends on all of the input vectors. Therefore, it should increase the capability of capturing long-range dependencies, which are especially helpful in the medical area. An example here would be global features like skin lesion asymmetry or dermoscopic structures. Behind the GA operation, a zero-initialized  $1 \times 1 \times 1$   $GA_w$  convolution is added. This convolution, together with a residual connection, results in an identity mapping at initialization (Wang et al., 2018). Hence, inserting it into a network keeps the pre-trained behavior but somewhat increases the parameter count. The computational graph can be seen in Figure 1.

### 2.2 Local Self-Attention

Secondly, we investigate a local self-attention (LA) mechanism based on (Ramachandran et al., 2019). This version does not operate on the whole feature




Figure 1: Computational graph of the global self-attention.

map for each input pixel, but rather on the local neighborhood  $N_k(i, j) = \{a, b \in \mathbb{Z} : |a - i| < \frac{k}{2} \wedge |b - j| < \frac{k}{2}\}$ , where  $k$  determines the neighborhood size. The authors introduce this approach by replacing the final convolutional layers of a CNN. The implementation is based on three  $1 \times 1$  convolutions, which require fewer parameters compared to a  $k \times k$  ( $k > 1$ ) convolution. Another advantage is its content-based computation. However, it is not capable of imitating the computation of the convolutional layer it replaces. Therefore, adding it to a CNN destroys the pre-trained behavior.

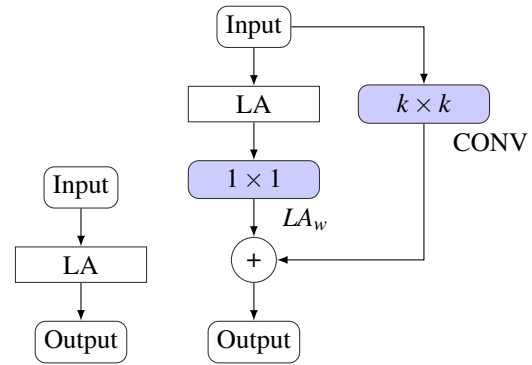
To overcome this issue, we propose an alternative implementation, which we call embedded local self-attention, henceforth “ELA”. The implementation is inspired by the design of GA described above. Specifically, we add a residual connection around the introduced LA block, containing the original convolutions and pre-trained weights. Similarly to GA, we add a zero-initialized  $1 \times 1 \times 1$  convolution after the LA block to conserve pre-trained behavior. Hence, it is content-based and keeps pre-trained behavior at initialization, at the cost of adding four  $1 \times 1$  convolutions. Both operations are visualized in Figure 2 as computational graphs, showing the similarity of ELA to GA.

## 2.3 Implementation Details

We base our analysis on two commonly used CNN architectures, ResNet (He et al., 2016) and EfficientNet (Tan and Le, 2019). For all models, we rely on commonly used ImageNet (Russakovsky et al., 2015) pre-trained weights as initialization. In the following, we describe our specific architecture choices.

### 2.3.1 ResNet (He et al., 2016)

Regarding ResNet architectures, we opt for ResNet18, a well-established and frequently used



(a) LA

(b) ELA

Figure 2: Computational graphs for both local approaches: local self-attention (LA) and embedded local self-attention (ELA).

variant. We insert the GA block after the first and second blocks of ResNet18, as recommended in (Wang et al., 2018). Regarding local attention, for both LA and ELA, we replace all convolutional layers of the last block. A similar setup is discussed in (Ramachandran et al., 2019). Note that all ResNet architectures feature four blocks. Hence, our inclusion of attention mechanisms could be implemented similarly for other variants.

### 2.3.2 EfficientNet (Tan and Le, 2019)

Again, we choose the smallest variant of this model family, EfficientNet-B0, featuring seven blocks. To insert the GA component, we again follow the recommendations of (Wang et al., 2018) and include it after the second and third architecture blocks. Similarly, we add LA and ELA by replacing convolutions in the last block. Note that EfficientNet-B0 utilizes so-called MBConv blocks (Sandler et al., 2018). Hence, minimizing the number of parameters and increasing efficiency. An MBConv block contains a convolutional layer using groups, amongst other techniques, to reduce parameters. Hence, naively replacing this grouped convolution significantly increases the parameter count. To avoid this issue, we replace the whole MBConv block instead, leading to a similar parameter reduction as observed for the ResNet models described above. We argue that this parameter reduction aligns better with the reasoning stated in (Ramachandran et al., 2019) and increases comparability between our selected CNN architectures.

### 2.3.3 Model Selection

Table 1 summarizes the different architectures we investigate in this study to analyze the influence of self-attention on skin lesion and tissue classification.

Table 1: Parameter counts of all architectures (in million) and relative change compared to the corresponding baselines.

	ResNet18	EfficientNet-B0
Base	11.18 $\pm$ 0.00%	4.01 $\pm$ 0.00%
+ GA	11.22 $\pm$ 0.37%	4.02 $\pm$ 0.12%
+ LA	5.67 $\pm$ 49.25%	3.48 $\pm$ 13.29%
+ GA + LA	5.71 $\pm$ 48.87%	3.48 $\pm$ 13.17%
+ ELA	14.98 $\pm$ 34.02%	4.30 $\pm$ 7.16%
+ GA + ELA	15.02 $\pm$ 34.40%	4.30 $\pm$ 7.27%
ViT	5.49 $\pm$ 0.00%	

Additionally, we report the parameter counts to ensure that none of our architecture changes widely skew the model capacities. For comparison, we also list our selected baselines: ResNet18 (He et al., 2016), EfficientNet-B0 (Tan and Le, 2019), and two attention-only ViT models (Kolesnikov et al., 2021). Here, we select the ViT tiny variant with a parameter count of 5.49 million parameters which is in between both selected CNN architectures. We investigate a version pre-trained on the commonly used ImageNet 1K challenge (Russakovsky et al., 2015) to stay comparable to our other experiments and baselines, as well as one ViT pre-trained on the larger ImageNet 21K dataset. We add the latter model, given empirical observations that larger pre-training datasets increase the downstream performance of transformer models (Kolesnikov et al., 2021).

### 3 EXPERIMENTS

We start by describing the general experimental setup we employed before discussing the achieved results. We then analyze the feature usage of the best-performing skin lesion models. Finally, we visualize the learned attention maps and compare them to local Grad-CAM (Selvaraju et al., 2017) explanations to gather additional insights.

#### 3.1 Experimental Setup

To investigate the influence of additional self-attention blocks, we analyze the performance of our hybrid approaches in three distinct ways. First, we trained the models on skin-lesion and tumor tissue classification, comparing the balanced accuracies on in-distribution and out-of-distribution test sets. Concerning skin lesions, we use data from the ISIC archive (ISIC, 2022) as training data and rely on the PAD-UFES-20 dataset (Pacheco et al., 2020) as out-of-distribution (OOD) test data. To align these tasks, we simplified the classification problem to the three

classes benign nevi, melanoma, and others. Regarding tumor tissue classification, we use the Camelyon17 dataset (Bandi et al., 2018), which already provides an out-of-distribution split in addition to the training data. The domain shift in Camelyon17 is introduced by sampling the OOD set from a different hospital.

To increase reliability, we performed a small hyperparameter search, considering the learning rate and weight decay. To reduce the search space, we fixed our batch size to 32. For data augmentation, we followed the recommendations from (Perez et al., 2018) and chose their best-performing augmentation scheme. This scheme randomly applies cropping, affine transformations, flipping, and changing the hue, brightness, contrast, and saturation of the input. In our experiments, we use SGD as an optimizer. To determine the performance of a specific hyperparameter combination, we perform the search with 2 random data splits.

After determining the optimal learning rate and weight decay for each model, which is shown in Table 5, we train each architecture 10 times on different data splits to determine the mean and standard deviation of the achieved performances. Additionally, we test for statistical significance by performing Welch’s T-Test (Welch, 1947) on these samples. Here, we employ the widely used significance level of  $p < 0.05$ .

#### 3.2 Performance Analysis

Table 2 summarizes the balanced accuracy for all examined models and corresponding baselines on both datasets and the corresponding OOD sets.

On the ISIC archive data, the ResNet18 variants with added attention mechanisms generally underperformed compared to the convolutional baseline. Further, while the EfficientNet-B0 variants primarily exhibit improvements, on average, over the baseline, these improvements are not statistically significant. On the Camelyon17 dataset, we observed similar performances between the baselines and our hybrid models, with the only significant result being the EfficientNet + ELA, which performs worse on average. Hence, overall, we found no reliable upward trend in performance when including self-attention mechanisms in a model.

We, however, noticed smaller differences between the specific attention versions. Specifically, we found that local approaches show, on average, increased performance in combination with GA. Nevertheless, the overall performance still decreases compared to the baselines.

While the significant reduction of performance for

Table 2: Balanced Accuracies on each Dataset and each Model. Values inside the parentheses show the balanced accuracies on the OOD-dataset. Bold values indicate **statistically significant** differences to the respective base model.

	ISIC Dataset		Camelyon17	
	ResNet	EfficientNet	ResNet	EfficientNet
Base	73.9 $\pm$ 1.74 (57.4 $\pm$ 8.87)	75.4 $\pm$ 1.75 (57.6 $\pm$ 15.18)	98.4 $\pm$ 0.18 (94.4 $\pm$ 1.66)	98.5 $\pm$ 0.16 (94.5 $\pm$ 2.33)
+GA	<b>72.1</b> $\pm$ 1.31 (56.3 $\pm$ 12.92)	76.6 $\pm$ 1.83 (55.8 $\pm$ 12.67)	98.4 $\pm$ 0.07 (94.0 $\pm$ 1.73)	98.6 $\pm$ 0.17 (95.4 $\pm$ 0.24)
+LA	<b>70.8</b> $\pm$ 2.96 (56.3 $\pm$ 5.83)	75.5 $\pm$ 1.58 (55.5 $\pm$ 7.77)	98.2 $\pm$ 0.36 ( <b>92.9</b> $\pm$ 3.15)	98.4 $\pm$ 0.21 (94.3 $\pm$ 2.01)
+GA+LA	<b>71.2</b> $\pm$ 1.44 (55.5 $\pm$ 4.64)	75.8 $\pm$ 1.14 (55.4 $\pm$ 12.48)	98.4 $\pm$ 0.16 ( <b>92.2</b> $\pm$ 5.90)	98.6 $\pm$ 0.26 (94.3 $\pm$ 0.63)
+ELA	<b>72.3</b> $\pm$ 1.45 (58.4 $\pm$ 7.05)	73.8 $\pm$ 2.29 (57.9 $\pm$ 8.43)	98.2 $\pm$ 0.23 ( <b>91.4</b> $\pm$ 8.22)	<b>98.3</b> $\pm$ 0.07 (94.1 $\pm$ 2.09)
+GA+ELA	73.5 $\pm$ 1.35 (58.8 $\pm$ 10.41)	75.9 $\pm$ 1.77 (54.7 $\pm$ 11.84)	98.4 $\pm$ 0.14 (95.1 $\pm$ 0.44)	98.5 $\pm$ 0.40 (94.3 $\pm$ 3.01)
	1k	21k	1k	21k
ViT	66.0 $\pm$ 17.36 (52.6 $\pm$ 113.45)	75.0 $\pm$ 0.90 (56.7 $\pm$ 6.94)	98.3 $\pm$ 0.43 (94.8 $\pm$ 0.91)	98.3 $\pm$ 0.11 (94.5 $\pm$ 0.96)

some models could be attributed to the reduced parameter count (LA models), we also observe similar behavior in models with increased parameter count (ELA). This is further corroborated by the observed higher performance of the EfficientNet models in the skin lesion task despite a lower parameter count compared to the ResNet variants.

Regarding OOD performance, we see that there are substantial differences between the datasets. For the Camelyon17 dataset, the degradation of performance is minimal for all models, and they still achieve high balanced accuracies. We observe higher average OOD performance for the EfficientNet + GA hybrid model. However, this increase is not statistically significant. Further, the only statistically reliable results regarding the OOD performance are the observed decreases for some of the ResNet hybrid models. In the skin lesion classification task, the drop in performance is much more severe, also leading to higher standard deviations. Hence, we did not observe any significant improvements of the self-attention hybrid models over the fully convolutional baselines in this task.

The performance of the ViT heavily depends on the pretraining data size in the skin lesion task. While pretraining, the model with the 21k classes version of ImageNet (Russakovsky et al., 2015) outperforms the ResNet baseline in the skin lesion classification task, pretraining with default ImageNet 1K does perform significantly worse. This result confirms previous observations, e.g., (Kolesnikov et al., 2021), that larger pretraining datasets can increase downstream performance. In our experiment, this specifically meant that two of our ten ViT 1K models diverged, resulting in random guessing capabilities and explaining the huge observed standard deviation.

### 3.3 Global Explanations: Feature Usage

To further investigate the model adaptations described in Sec. 2.3.3, we investigated our hybrid models and

baselines on a feature level in the skin lesion classification task. Here we employ the feature attribution method described in (Reimers et al., 2020). Previous work (Reimers et al., 2021; Penzel et al., 2022) also applied this method to skin lesion classification and investigated the usage of features related to the ABCD rule (Nachbar et al., 1994) and known biases in skin lesion data (Mishra and Celebi, 2016).

Here, we specifically follow (Reimers et al., 2021; Penzel et al., 2022) and extract the same four features related to the ABCD rule, namely **A**symmetry, **B**order irregularity, **C**olor, and **D**ermoscopic structures, as well as the four bias features age of a patient, sex of the patient, skin color, and the occurrence of large colorful patches (Scope et al., 2016). For a detailed description of these features, we refer the reader to the original work (Reimers et al., 2021). Furthermore, we follow the hyperparameter settings described in (Penzel et al., 2022) and select three conditional independence tests, namely cHSIC (Fukumizu et al., 2007), RCoT (Strobl et al., 2019), and CMiknn (Runge, 2018). We report the majority decision and use a significance level of  $p < 0.01$  (Penzel et al., 2022).

Table 3 contains the results of our feature analysis. Regarding the bias features, the age and skin color of the patients are both learned by all analyzed models. The colorful patches and patient's sex features are also often learned, however, predominantly by the ResNet variants. The EfficientNet models seem to be more robust in that regard. Looking at the meaningful ABCD rule features, we observe that the models often learn the asymmetry or border irregularity while the color feature is rarely, and the dermoscopic structures feature is never learned. In contrast, both ViT models learn to incorporate the color of the lesion into their decision process. In general, GA seems to produce slightly better results, given that all extended models that utilize the color feature also include GA. However, this observation is inconsistent as other combinations of GA do not result in this behavior. Some

Table 3: Feature usage of our best-performing skin lesion models according to balanced accuracy. We abbreviate Asymmetry, Border irregularity, Color, and Dermoscopic structures with the associated ABCD rule letter. We use significance level  $p < 0.01$  (Reimers et al., 2021).

	Model	A	B	C	D	Age	Sex	Skin color	Colorful patches
ResNet	Base	✓	✓	✗	✗	✓	✓	✓	✓
	+GA	✓	✓	✓	✗	✓	✓	✓	✓
	+LA	✗	✓	✗	✓	✓	✓	✓	✓
	+GA+LA	✓	✓	✓	✗	✓	✗	✓	✓
	+ELA	✓	✓	✗	✗	✓	✓	✓	✓
	+GA+ELA	✓	✓	✗	✗	✓	✓	✓	✓
EfficientNet	Base	✓	✓	✗	✗	✓	✗	✓	✗
	+GA	✗	✓	✗	✗	✓	✗	✓	✓
	+LA	✓	✓	✗	✗	✓	✓	✓	✓
	+GA+LA	✓	✓	✓	✗	✓	✗	✓	✓
	+ELA	✓	✓	✗	✗	✓	✓	✓	✗
	+GA+ELA	✗	✓	✗	✗	✓	✓	✓	✓
ViT	1K	✓	✓	✓	✗	✓	✗	✓	✓
	21K	✓	✓	✓	✗	✓	✓	✓	✓

models containing GA regress and do not learn to utilize the skin lesion asymmetry, which is consistently learned by the CNN baselines and which has also been reported previously (Reimers et al., 2021).

To conclude, we do not find a systematic improvement of the extended models incorporating attention mechanisms over the fully convolutional baselines. There is a small benefit in some attention-based models with respect to the color feature, especially for the ViT architecture. Nevertheless, these results are similar to the empirical evaluation described above, where some of the extended models achieved a performance increase, which was not statistically significant.

### 3.4 Local Explanations: Visualizations

Introducing self-attention into CNNs allows for a direct visualization of important image areas, which can help to gain further insights about the model decision process. We leverage the method from (Wang et al., 2018), which creates a map showing how a certain pixel pays attention to all others and calculates the average over all of these maps. We compare this attention map with Grad-CAM (Selvaraju et al., 2017) and the visualization method of the default ViT (Kolesnikov et al., 2021). Table 4 contains these visualizations generated using five images (all classes for both datasets) and additionally displays the mean of the attention maps over 500 images. To get visually comparable results, we normalized the values between 0 and 255. To stay comparable and only generate one visualization per image, we select the Grad-

CAM for the predicted class for all examples. We find that while models focus on different parts of the images, we observe no visible improvements or biases. Noteworthy, the highlighted area of Grad-CAM is typically larger than the one for the GA visualization due to the way Grad-CAM is implemented in (Selvaraju et al., 2017), which we directly follow. Grad-CAM highlights mostly the whole object of interest. This can be seen in the mean image for the ISIC task because the center, where most of the skin lesions are placed, gets highlighted the most. On the other hand, the attention map visualization highlights more specific features like the border of the skin lesion, framing both interpretability techniques as uniquely different. Interestingly GA, as well as the ViT, put a lot of attention on the corner of the image. While this is initially quite unintuitive, we suggest that this is similar behavior as reported in (Darcet et al., 2023; Xiao et al., 2023). Here, transformers often produce attention spikes in the image background, as visible in our mean images, to store information for further calculations. Unfortunately, this behavior cannot be further unraveled by looking at the attention maps, therefore, revealing a possible disadvantage of these visualizations in comparison to Grad-CAM.

Since we do not carry extensive medical knowledge, we refrain from more interpretations into the visualization here. We, however, note that the attention maps do not provide extensive visual explanations that could not be provided by Grad-CAM approaches in a similar fashion. Additionally, we do not observe specific structural biases since models seem

Table 4: Qualitative examples of explanations generated using either Grad-CAM (Selvaraju et al., 2017) or by visualizing the learned attention maps. We choose one image per class per dataset and the best performing models of either the baselines or adapted with GA or ELA (here abbreviated with G and E).

Model	ISIC				Camelyon17					
	Melanoma	Nevus	Other	Mean	Tumor	No Tumor	Mean			
Grad-CAM	ResNet	-								
		+G								
		+E								
		+GE								
	EfficientNet	-								
		+G								
		+E								
		+GE								
	AttentionMaps	ResNet	+G							
			+GE							
		EfficientNet	+G							
			+GE							
ViT										

to set a reasonable focus concerning the input data typically. Therefore, we attribute no systematic benefit to attention maps neither concerning local feature usage nor interpretability.

## 4 CONCLUSIONS

This work investigated the usage of self-attention mechanisms for medical imaging, specifically classification. Toward this goal, we extended two widely used convolutional architectures with self-attention mechanisms and empirically compared these extended models against fully convolutional and attention-only baselines. We conducted this comparison on a dataset constructed from the well-known ISIC archive (ISIC, 2022) and on the Camelyon17 dataset (Bandi et al., 2018). Additionally, we analyze the OOD generalization on domain-shifted test sets.

While we did observe some minor improvements, none of the attention-featuring architectures led to a statistically significant increase in performance. Sometimes, we even observe a significant decrease in balanced accuracy. In some instances, we find the reduction of performance on the OOD test data is decreased when using self-attention. However, we observe that this behavior is somewhat sensitive to the used backbone model, dataset, and in-distribution performance. Our investigation suggests that self-attention alone does not result in a direct increase in performance or generalization of medical image classification models.

To explain our findings further, we conducted two related analyses, taking a closer look at explanations and feature usage of our architectures. First, we performed a feature analysis of the best-performing skin lesion classification models using the attribution method described in (Reimers et al., 2020). We observed no systematic improvements in medically relevant features or bias features over the fully convolutional baselines, only slight deviations. Important global features, such as dermoscopic structures, are still not learned by employing self-attention.

Second, we performed a qualitative analysis of the learned attention maps. While including attention increases interpretability, e.g., revealing biased feature usage, we find that out-of-the-box explanation methods, we used Grad-CAM (Selvaraju et al., 2017), perform similarly for that purpose. To conclude: our work indicates that merely including self-attention does not directly lead to benefits. Of course, more work is required to ultimately conclude the suitability of self-attention mechanisms for medical image analysis. In any case, we hope to inspire future work to in-

vestigate new architectural changes more thoroughly than merely comparing performance scores since we believe this can provide valuable insights.

## REFERENCES

- Bandi, P., Geessink, O., Manson, Q., Van Dijk, M., Balkenhol, M., Hermesen, M., Bejnordi, B. E., Lee, B., Paeng, K., Zhong, A., et al. (2018). From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE Transactions on Medical Imaging*.
- Bera, K., Schalper, K. A., Rimm, D. L., Velcheti, V., and Madabhushi, A. (2019). Artificial intelligence in digital pathology — new tools for diagnosis and precision oncology. *Nature Reviews Clinical Oncology*, pages 1–13.
- Celebi, M. E., Codella, N., and Halpern, A. (2019). Dermoscopy image analysis: overview and future directions. *IEEE journal of biomedical and health informatics*, 23(2):474–478.
- Codella, N., Rotemberg, V., Tschandl, P., Celebi, M. E., Dusza, S., Gutman, D., Helba, B., Kalloo, A., Liopyris, K., Marchetti, M., Kittler, H., and Halpern, A. (2019). Skin Lesion Analysis Toward Melanoma Detection 2018: A Challenge Hosted by the International Skin Imaging Collaboration (ISIC). *arXiv:1902.03368 [cs]*. arXiv: 1902.03368.
- Darcey, T., Oquab, M., Mairal, J., and Bojanowski, P. (2023). Vision transformers need registers. *arXiv preprint arXiv:2309.16588*.
- Fukumizu, K., Gretton, A., Sun, X., and Schölkopf, B. (2007). Kernel measures of conditional dependence. *Advances in neural information processing systems*, 20.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- He, X., Tan, E.-L., Bi, H., Zhang, X., Zhao, S., and Lei, B. (2022). Fully transformer network for skin lesion analysis. *Medical Image Analysis*, 77:102357.
- ISIC (2022). Isic archive home page. Last accessed 23 July 2023.
- Khened, M., Kori, A., Rajkumar, H., Krishnamurthi, G., and Srinivasan, B. (2021). A generalized deep learning framework for whole-slide image segmentation and analysis. *Scientific Reports*, 11:11579.
- Kolesnikov, A., Dosovitskiy, A., Weissenborn, D., Heigold, G., Uszkoreit, J., Beyer, L., Minderer, M., Dehghani, M., Houlsby, N., Gelly, S., Unterthiner, T., and Zhai, X. (2021). An image is worth 16x16 words: Transformers for image recognition at scale.
- Krishna, G., Supriya, K., K, M., and Sorgile, M. (2023). Lesionaid: Vision transformers-based skin lesion generation and classification.
- Lee, B. and Paeng, K. (2018). A robust and effective approach towards accurate metastasis detection and p-stage classification in breast cancer. In Frangi, A. F.,



- Schnabel, J. A., Davatzikos, C., Alberola-López, C., and Fichtinger, G., editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 841–850, Cham. Springer International Publishing.
- Mishra, N. K. and Celebi, M. E. (2016). An Overview of Melanoma Detection in Dermoscopy Images Using Image Processing and Machine Learning. *arXiv:1601.07843 [cs, stat]*. arXiv: 1601.07843.
- Nachbar, F., Stolz, W., Merkle, T., Cagnetta, A. B., Vogt, T., Landthaler, M., Bilek, P., Braun-Falco, O., and Plewig, G. (1994). The ABCD rule of dermatoscopy. High prospective value in the diagnosis of doubtful melanocytic skin lesions. *Journal of the American Academy of Dermatology*, 30(4):551–559.
- Pacheco, A. G., Lima, G. R., Salomão, A. S., Krohling, B., Biral, I. P., de Angelo, G. G., Alves Jr, F. C., Es-gario, J. G., Simora, A. C., Castro, P. B., Rodrigues, F. B., Frasson, P. H., Krohling, R. A., Knidel, H., Santos, M. C., do Espírito Santo, R. B., Macedo, T. L., Canuto, T. R., and de Barros, L. F. (2020). Pad-ufes-20: A skin lesion dataset composed of patient data and clinical images collected from smartphones. *Data in Brief*, 32:106221.
- Penzel, N., Reimers, C., Bodesheim, P., and Denzler, J. (2022). Investigating neural network training on a feature level using conditional independence. In *ECCV Workshop on Causality in Vision (ECCV-WS)*, pages 383–399, Cham. Springer Nature Switzerland.
- Perez, F., Vasconcelos, C., Avila, S., and Valle, E. (2018). Data augmentation for skin lesion analysis.
- Ramachandran, P., Parmar, N., Vaswani, A., Bello, I., Levskaya, A., and Shlens, J. (2019). Stand-alone self-attention in vision models. *Advances in Neural Information Processing Systems*, 32.
- Reimers, C., Penzel, N., Bodesheim, P., Runge, J., and Denzler, J. (2021). Conditional dependence tests reveal the usage of abcd rule features and bias variables in automatic skin lesion classification. In *CVPR ISIC Skin Image Analysis Workshop (CVPR-ISIC)*, pages 1810–1819.
- Reimers, C., Runge, J., and Denzler, J. (2020). Determining the relevance of features for deep neural networks. In *European Conference on Computer Vision*, pages 330–346. Springer.
- Runge, J. (2018). Conditional independence testing based on a nearest-neighbor estimator of conditional mutual information. In *International Conference on Artificial Intelligence and Statistics*. PMLR.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. (2015). Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., and Chen, L.-C. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520.
- Scope, A., Marchetti, M. A., Marghoob, A. A., Dusza, S. W., Geller, A. C., Satagopan, J. M., Weinstock, M. A., Berwick, M., and Halpern, A. C. (2016). The study of nevi in children: Principles learned and implications for melanoma diagnosis. *Journal of the American Academy of Dermatology*, 75(4):813–823.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626.
- Society, A. C. (2022). Cancer facts & figures 2022. Last accessed 02 August 2022.
- Strobl, E. V., Zhang, K., and Visweswaran, S. (2019). Approximate kernel-based conditional independence tests for fast non-parametric causal discovery. *Journal of Causal Inference*.
- Tan, M. and Le, Q. (2019). Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR.
- Tolstikhin, I. O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., Lucic, M., and Dosovitskiy, A. (2021). Mlp-mixer: An all-mlp architecture for vision. *CoRR*, abs/2105.01601.
- Trockman, A. and Kolter, J. Z. (2022). Patches are all you need? *CoRR*, abs/2201.09792.
- Tschandl, P., Rosendahl, C., and Kittler, H. (2018). The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Sci. Data*, 5(1):180161.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, X., Girshick, R., Gupta, A., and He, K. (2018). Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803.
- Welch, B. L. (1947). The generalization of ‘student’s’ problem when several different population variances are involved. *Biometrika*, 34(1/2):28–35.
- Xiao, G., Tian, Y., Chen, B., Han, S., and Lewis, M. (2023). Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*.
- Yang, G., Luo, S., and Greer, P. (2023). A novel vision transformer model for skin cancer classification. *Neural Processing Letters*, pages 1–17.

## APPENDIX

Table 5 contains the best hyperparameters, we determined for our models.

Table 5: Best determined hyperparameters.

	Model	ISIC		Camelyon17	
		LR	WD	LR	WD
ResNet	Base	0.001	0.0001	0.001	0.001
	+GA	0.001	0	0.001	0.001
	+LA	0.001	0.0001	0.0001	0.01
	+GA+LA	0.001	0	0.001	0.0001
	+ELA	0.001	0	0.01	0
	+GA+ELA	0.001	0	0.001	0.0001
	EfficientNet	Base	0.01	0	0.001
	+GA	0.01	0	0.001	0.0001
	+LA	0.01	0	0.001	0
	+GA+LA	0.01	0	0.001	0.0001
	+ELA	0.001	0	0.001	0.001
	+GA+ELA	0.01	0	0.01	0
ViT	1K	0.001	0	0.001	0
	21K	0.0001	0.0001	0.001	0.0001

SCITEPRESS  
SCIENCE AND TECHNOLOGY PUBLICATIONS