Deep Learning Models for Segmentation of Mobile-Acquired Dermatological Images

Catarina Andrade^{1(✉)}, Luís F. Teixeira², Maria João M. Vasconcelos¹,
and Luís Rosado¹

¹ Fraunhofer Portugal AICOS, 4200-135 Porto, Portugal
{[catarina.andrade](mailto:catarina.andrade@fraunhofer.pt),[maria.vasconcelos](mailto:maria.vasconcelos@fraunhofer.pt),[luis.rosado](mailto:luis.rosado@fraunhofer.pt)}@fraunhofer.pt

² FEUP/INESC, 4200-465 Porto, Portugal
luisft@fe.up.pt

Abstract. With the ever-increasing occurrence of skin cancer, timely and accurate skin cancer detection has become clinically more imperative. A clinical mobile-based deep learning approach is a possible solution for this challenge. Nevertheless, there is a major impediment in the development of such a model: the scarce availability of labelled data acquired with mobile devices, namely macroscopic images. In this work, we present two experiments to assemble a robust deep learning model for macroscopic skin lesion segmentation and to capitalize on the sizable dermoscopic databases. In the first experiment two groups of deep learning models, U-Net based and DeepLab based, were created and tested exclusively in the available macroscopic images. In the second experiment, the possibility of transferring knowledge between the domains was tested. To accomplish this, the selected model was retrained in the dermoscopic images and, subsequently, fine-tuned with the macroscopic images. The best model implemented in the first experiment was a DeepLab based model with a MobileNetV2 as feature extractor with a width multiplier of 0.35 and optimized with the soft Dice loss. This model comprehended 0.4 million parameters and obtained a thresholded Jaccard coefficient of 72.97% and 78.51% in the Dermofit and SMARTSKINS databases, respectively. In the second experiment, with the usage of transfer learning, the performance of this model was significantly improved in the first database to 75.46% and slightly decreased to 78.04% in the second.

Keywords: Skin lesion segmentation · Macroscopic images · Dermoscopic images · Convolution neural networks · Knowledge transfer

1 Introduction

Skin cancer is the most prevalent malignancy worldwide [5]. Among the carcinogenic skin lesions, malignant melanoma is less common yet the most lethal, with a worldwide morbidity of 60.7 thousand people in 2018 [5]. Timely and accurate

skin cancer detection is clinically highly relevant since the estimated 5-year survival rate for malignant melanoma drops from over 98% to 23% if detected when the metastases are distant from the origin point [20]. Nonetheless, this effective diagnosis raises another paradigm: the clinical presentation of most common cutaneous cancers is, every so often, identical to benign skin lesions.

The mobile technological advancement and the ubiquitous adoption of smart-phones associated with the high performance of deep learning algorithms have the potential to improve skin cancer triage with the creation of an algorithm which can match or outperform the visual assessment of skin cancer. Convolutional neural networks have been the staple method used in the skin lesion segmentation challenge. Most methods are based on modifications of encoder-decoder architecture of the U-Net [15]. From small changes, as modifying the number of input channels and loss optimization [11] to the addition of recurrent layers and residual units [1], Sarker et al. [18] developed a U-Net with an encoder path consisting of four pre-trained dilated residual networks and a pyramid pooling block.

Nevertheless, the scarce availability of labelled data acquired with mobile devices, namely macroscopic images may prove to be a major impediment for the creation of such a method. Habitually, the cutaneous skin lesions are diagnosed by skin lesion surface microscopy (dermoscopy) which allows for the visualization of the subsurface skin structures which are usually not visible to the naked eye. This compelled the creation of several dermoscopic databases of substantial size. Withall, the direct inference between the macroscopic and dermoscopic domain is not advisable due to their paradoxical characteristics and challenges namely, the acquisition of images with the dermoscope generates several structures, colours and artefacts which are not detectable in the macroscopic image. The polarized light that permits the visualization of these characteristics eliminates the surface glare of the skin, which is abundantly common in the macroscopic setting. Additionally, structures clearly visible in dermoscopic images like pigmented network, streaks, dots, globules, blue-whitish veil or vascular patterns are usually less noticeable or even imperceptible in macroscopic images. Furthermore, the flat outward aspect in the dermoscopic images, caused by the direct contact of the dermoscope with the skin is paradoxical with the visual depth normally present in the macroscopic images. In fact, even for the diagnosis, there are rules and methods specific for each domain [9].

This work aims to evaluate the possibility of designing a deep learning algorithm for segmenting the lesion in macroscopic images which would operate fully in the mobile environment. This involves creating a fast and lightweight algorithm with expert-level accuracy to be integrated into the mobile environment. To assemble such a model, we explored the capitalization of the sizable dermoscopic databases and designed two separate experiments.

2 Methodology

2.1 Databases and Problem Definition

As there were several databases available which provided matching binary segmentation masks, it was possible to assemble two distinct datasets: the

dermoscopic (set D) and the macroscopic (set M). The set D was constituted by the combined images of all ISIC Challenges (2016 [7], 2017 [4] and 2018 [3,21]) and the PH2 database [13] and set M was comprised of the Dermofit image Library [12] and the SMARTSKINS database [22].

For the PH2 and Dermofit image Library, an 80/20 split was used for the creation of training/validation and test sets. In the case of the SMARTSKINS database, a 50/50 split was used due to the database the small size. Considering the three ISIC challenges had duplicate images in the different years, the training datasets of ISIC 2016, ISIC 2017 and ISIC 2018 and the validation of ISIC 2017 were combined and the duplicates removed. Subsequently, the image instances of the test dataset of ISIC 2017 were also eliminated from the combined dataset and reserved as a test subset. When the databases were classified, the division was structurally made to maintain an equal percentage of each class in the training/validation and test sets. The characterization and the slitting of the databases in each dataset are summarised in Table 1.

Table 1. Overview of available segmentation databases and separation into train/validation and test subsets.

Set	Database	No. images (type)	No. train/val.	No. test
Set M	Dermofit	1300 (MD)	1036	264
	SMARTSKINS	80 (MP)	39	41
Set D	ISIC	2594 (De)	1994	600
	PH2	200 (De)	160	40

No. - Number; De - Dermoscopic images; MD - Macroscopic images acquired with a digital camera; MP - Macroscopic images acquired with a mobile phone;

As it can be observed in Table 1, the size of set D is almost double of the size of set M, which lead to the creation of two experiments. In the first experiment, a comparative study, using exclusively the set M, was performed with two major groups of deep Learning models, U-Net and DeepLab based. From this study, a model was to be chosen to be used in the following experiment. The second experiment tested the possibility of transferring the knowledge from the dermoscopic to the macroscopic domain was tested. This was accomplished by re-training the chosen model in the first experiment with the set D and subsequent fine-tuning of the model with set M.

2.2 Deep Learning Models for Semantic Segmentation

For implementation, we adopt the Tensorflow API r1.15 in Python 3.7.3 on three NVIDIA Tesla V100 PCIe GPU module, two with 16GB and the other with 32GB. Initially, the standard image resizing (512×512 pixels) and standardization were performed as a preprocessing stage. As for the training protocol, we employ a batch size of 4, a 90/10 partition for the training/validation subsets and the Adam optimizer to perform stochastic optimization with a cyclic

learning rate (CLR) [19]. The cycle used was the cosine annealing variation with periodic restarts [8] associated with early stopping and model checkpointing.

U-Net Based Models. As a baseline model, we implemented a classical U-Net [15] with the addition of dropout layers (0.2) and zero-padding. The second model was an Attention U-Net [14] (AttU-Net) which main modification lies in the addition of an Attention Gate (AG) at the skip connection of each level. The third U-Net based model trained was the R2U-Net [1] which is a recurrent residual convolutional neural network based on U-Net. The last model implemented was a combination of the two aforementioned models (AttR2U-Net) [10]. The optimal number of initial feature channels was also analysed. This model parameter can be used to decrease the model complexity, however, it can also downgrade the performance of the model. This hypothesis was tested, using the values 16, 32 and 64, to ascertain the fidelity of this technique.

DeepLab Based Models. The second proposed approach was the state-of-the-art DeepLabv3+ model [2]. Initially, the original modified Aligned Xception encoder was used [2]. However, due to its considerable size which is not suited to the mobile environment, the MobileNetV2 as in [17] was also implemented. Primarily, all models were tested with randomly initialized weights and then two different sets of pre-trained weights were used: one pre-trained in Cityscapes and the other on the Pascal VOC 2012. Furthermore, some encoder specific experiments were also performed. To the modified Aligned Xception encoder, two output strides (OS), which refer to the spatial resolution ratio between the input and the output images, were tested: 8 and 16. In the case of the MobileNetV2, the variation of width multiplier (α) was analyzed. This hyperparameter allows the manipulation of the input width of a layer, which can lead to the reduction or augmentation of the models by a ratio roughly of α^2 .

Model Optimization. The selection of a suitable loss function for the challenge at hand is pivotal to reach the appropriate capacity of the model. In total, five losses were tested. Initially, cross-entropy (CE) was chosen as the standard loss. Then four losses were tested: the soft Dice coefficient (DI) loss ($1 - DI$) and soft Jaccard coefficient (JA) loss ($1 - JA$), and the logarithmic combinations with the cross-entropy ($CE - \log JA$, $CE - \log DI$). These losses use a soft variant of the DI or JA, which uses the predicted probabilities instead of a binary mask, to decrease the effect of the class imbalance amidst each sample.

Performance Assessment. In the first experiment all the image segmentation models were trained on set M. Here, all the model configurations were tested and selected based on the results on the validation dataset. The measures used to evaluate the performance of the models were the ones used in the ISIC challenge of 2018 [3]: threshold JA (TJA), JA, DI, accuracy (AC), sensitivity (SE) and specificity (SP). The decision of the best model was a balance between model complexity and TJA, as it was the scoring metric of the ISIC 2018 challenge. The threshold on which these metrics were taken was inferred with the result of a JA analysis performed on the validation subset. Considering the sigmoid nonlinearity used in the last layer, the binarization threshold was estimated

with the intent of maximizing the resultant JA. For this purpose, the JA was evaluated in 50 thresholds within the $[0.5, 1]$ range. In the second experiment, the model was evaluated in the test subset and compared with the results of the selected model of the first experiment.

3 Experimental Results

3.1 First Experiment

The best performing configurations for each model architecture of the U-Net based models and each encoder of the Deeplab based models are summarized in Table 2.

Regarding the best performing U-Net based models, the addition of the AG proved to be quite disadvantageous. The models with this extension reach underperforming results leading to a decrease in 2% in JA and TJA when compared with the classical U-Net structure. Besides, the requirement of a higher number of ifC, 64 instead of 32 needed in the baseline U-Net elucidates to the inefficiency of the networks with the AG in learning representative features. On the other hand, the addition of the recurrent residual unit leads to 3% improvement in JA and 4% in TJA when compared to the Classical U-Net. Contrarily to the AttU-Net, the R2U-Net reached higher performance with the lowest ifC tested, 16. However, the number of parameters of this network increases significantly when compared with the other networks. Consequentially, the R2U-Net with 16 ifC and the classical U-Net with 32 ifC have a similar number parameter, around 7M. When adding the recurrent residual units with the AG the results are entirely consistent with the aforementioned conclusions. Essentially, the AttR2U-Net result improves in comparison to the classical U-Net however it decreases 1% the JA when comparing with the R2U-Net.

Table 2. Performance metrics for each model architecture used in the U-Net approach best configuration and for each encoder used in the DeepLab approach, evaluated in the validation subset of set M.

U-Net based models										
Model architecture	L	ifC	P	T	TJA	JA	DI	AC	SE	SP
U-Net	$1 - JA$	32	8 M	0.52	78.44	80.80	89.24	93.16	90.42	94.65
AttU-Net	$1 - DI$	64	32 M	0.50	76.55	78.78	87.91	93.27	88.27	95.69
R2U-Net	$1 - JA$	16	6 M	0.50	83.02	83.02	90.60	94.39	90.05	96.48
AttR2U-Net	$CE - \log JA$	64	96 M	0.50	82.08	82.08	90.00	93.98	89.38	96.18
Deeplab based models										
Encoder	L	PreT	P	T	TJA	JA	DI	AC	SE	SP
Xception, OS = 16	$1 - JA$	Pascal	41 M	0.50	85.37	85.37	92.00	95.54	91.46	97.05
MobileNetV2, $\alpha = 1.0$	$1 - JA$	Pascal	2 M	0.55	84.80	84.80	91.70	95.18	91.60	96.55
MobileNetV2, $\alpha = 0.35$	$1 - DI$	None	0.4 M	0.50	82.85	82.85	90.52	94.65	88.31	97.27

L - Loss Function; ifC - initial feature channels; P - Model Parameters; T - binarization threshold used in the output probability map; PreT - Pretraining of weights;

Concerning the DeepLab approach, an overall conclusion can be drawn from the robustness of the combination of the soft JA loss function ($1 - JA$) and the pre-trained model in Pascal VOC 2012 which lead to the top-performing results for each of the encoders. Concerning the output stride of $OS = 16$ surpasses in terms of performance the $OS = 8$, meaning a denser feature extraction in the last layers of the model decoder is not suitable for skin lesion segmentation. Not surprisingly, the addition of the inverted residual depthwise separable convolution of the MobileNetV2 encoder leads to a dramatical reduction of model complexity. In fact, there is almost a reduction of approximately nineteen times fewer parameters and with the loss of less than 1% in all of the performance metrics. This result prompted a second study, which focuses only on the effects on the reduction of the MobileNetV2. Thus, several models with various α and no pre-training of weights were implemented and optimized with the five designed loss functions. The best result from this study is summarised in Table 2 (Deeplab based models, row 3).

Pertaining to the loss function the results are quite consistent. For each model architecture, the stochastic optimization performed by a loss function, which takes into account the soft variation of JA and DI, leads to improved results. The soft DI and soft JA losses yielded the best results in all the models except the AttR2U-Net. Therefore, the use of a loss function which takes into consideration the measure of overlap between two samples is an effective way of reducing the class imbalance between the surrounding skin and the lesion pixels.

Based on the aforementioned approaches and experiments, one model was chosen to be evaluated in the test subset of set M. The main rationale behind this selection was choosing the model which offers the best a balance between two desirable but usually incompatible aspects features: performance and model complexity. The selected model was the reduced Mobile DeepLab with $\alpha = 0.35$ and optimized with soft DI loss function mainly due to its reduced size and its JA and TJA values above the 80% threshold, which is above the interobserver agreement and the visual correctness threshold [4].

3.2 Second Experiment

After the selection of the reduced Mobile DeepLab, the model was retrained with set D. The same training procedure and network parameters were used. Subsequently, the model trained with set D was fine-tuned with set M. The obtained results of both experiments, evaluated in the test subset of the SMARTSKINS and Dermofit, are presented in Table 3. From this table, it is possible to infer that the fine-tuning, with the macroscopic data, of the pre-trained model, on set D, leads to a 2.49% TJA improvement in digital-acquired images (Dermofit) and a slight decrease of less than 0.48% in the mobile-acquired images (SMARTSKINS).

For the SMARTSKINS database, there is no standard used for the slitting of the database into train-validation-test subset. Therefore, the comparison with the models in the literature might not be as equitable as desired. Nevertheless, the reduced Mobile DeepLab attains in the first experiment the performance of

Table 3. Comparison of the performance metrics of the reduced Mobile DeepLab from the two proposed experiments, evaluated in the test subset of the SMARTSKINS and Dermofit.

Test subset	Experiment	T	TJA	JA	DI	AC	SE	SP
SMARTSKINS	Exp 1	0.50	78.51	82.64	90.14	98.96	95.40	99.15
	Exp 2	0.56	78.04	82.21	89.89	98.90	96.05	99.09
Dermofit	Exp 1	0.50	72.97	80.26	88.26	93.51	87.56	96.86
	Exp 2	0.56	75.46	81.03	88.79	93.78	89.68	96.13

T - binarization threshold used in the output probability map;

82.64% in JA, 90.14% in DI and 99.15% in AC. These values set a new state-of-the-art performance in the SMARTSKINS database which previously was of 81.58% in JA [16], 83.36% in DI [6] and 97.38% in AC [16].

Figure 1 presents several examples of the predicted segmentation mask of the model trained in each experiment compared with the ground truth label (GT). The model in both experiments shows highly satisfactory results when the lesion is pigmentedated with high contracts with the skin (Fig. 1, row 1). The presence of lesions with dysplastic form and uneven pigmentation can lead to the underperformance of the model of the second experiment (Fig. 1, row 2). The model of the second experiment outperforms the other in the presence of dark hair and lesions with other moles near (Fig. 1, row 3). Both experiments seem to underperform when the lesion presents red regions amidst the normal skin and vascularizations near the lesion border (Fig. 1, row 4).

Fig. 1. Examples of successful and failed segmentation results on the SMARTSKINS (left) and Dermofit (right) test subset. In the comparison images: yellow - true positives; red - false positives; green - false negatives; black - true negatives; (Color figure online)

4 Conclusion and Future Work

The yielded results show considerable potential in the use of models with decreased complexity and size. Altogether, the selected network had less than half a million parameters and a decrease in performance of TJA of 3% when compared with a model with approximately 41 M parameters.

When comparing the two experiments it can be inferred that the knowledge transfer between the dermoscopic and macroscopic domains still resulted in an overall improvement of the model. Despite the slight decrease in performance on the SMARTSKINS dataset, the improvement in the Dermofit dataset is significantly larger. It should be noted that the Dermofit dataset has more variety of skin lesion classes, including non-pigmented lesions that are not present in the SMARTSKINS dataset, thus we can assume that the fine-tuning procedure brought an overall model improvement. Nevertheless, there's still room for improvements, namely further experiments should be done in order to effectively take advantage of the sizable dermoscopic datasets.

Acknowledgements. This work was done under the scope of project “DERM.AI: Usage of Artificial Intelligence to Power Tele dermatological Screening” and supported by national funds through ‘FCT–Foundation for Science and Technology, I.P.’, with reference DSAIPA/AI/0031/2018.

References

1. Alom, M.Z., Hasan, M., Yakopcic, C., Taha, T.M., Asari, V.K.: Recurrent residual convolutional neural network based on U-Net (R2U-Net) for medical image segmentation (2018)
2. Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) ECCV 2018. LNCS, vol. 11211, pp. 833–851. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01234-2_49
3. Codella, N., et al.: Skin lesion analysis toward melanoma detection 2018: a challenge hosted by the international skin imaging collaboration (ISIC) (2019)
4. Codella, N.C.F., 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), April 2018
5. Ferlay, J., et al.: Estimating the global cancer incidence and mortality in 2018: globocan sources and methods. *Int. J. Cancer* **144**(8), 1941–1953 (2019)
6. Fernandes, K., Cruz, R., Cardoso, J.S.: Deep image segmentation by quality inference. In: 2018 International Joint Conference on Neural Networks (IJCNN), pp. 1–8. IEEE (2018)
7. Gutman, D., et al.: Skin lesion analysis toward melanoma detection: a challenge at the international symposium on biomedical imaging (ISBI) 2016, hosted by the international skin imaging collaboration (ISIC) (2016)
8. Huang, G., Li, Y., Pleiss, G., Liu, Z., Hopcroft, J.E., Weinberger, K.Q.: Snapshot ensembles: train 1, get M for free. arXiv preprint [arXiv:1704.00109](https://arxiv.org/abs/1704.00109) (2017)

9. Korotkov, K., Garcia, R.: Computerized analysis of pigmented skin lesions: a review. *Artif. Intell. Med.* **56**(2), 69–90 (2012)
10. LeeJunHyun: Pytorch implementation of U-Net, R2U-Net, attention U-Net, attention R2U-Net (2019). <https://github.com/LeeJunHyun/Image.Segmentation>. Accessed 1 Jan 2020
11. Lin, B.S., Michael, K., Kalra, S., Tizhoosh, H.R.: Skin lesion segmentation: U-nets versus clustering. In: 2017 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1–7. IEEE (2017)
12. Ltd., E.I.: Dermofit image library - edinburgh innovations (2019). <https://licensing.eri.ed.ac.uk/i/software/dermofit-image-library.html>. Accessed 11 June 2019
13. Mendonça, T., Ferreira, P.M., Marques, J.S., Marcal, A.R., Rozeira, J.: PH²-a dermoscopic image database for research and benchmarking. In: 2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 5437–5440. IEEE (2013)
14. Oktay, O., B., et al.: Attention U-Net: learning where to look for the pancreas. arXiv preprint [arXiv:1804.03999](https://arxiv.org/abs/1804.03999) (2018)
15. Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4_28
16. Rosado, L., Vasconcelos, M.: Automatic segmentation methodology for dermatological images acquired via mobile devices. In: Proceedings of the International Joint Conference on Biomedical Engineering Systems and Technologies, vol. 5, pp. 246–251 (2015)
17. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV 2: inverted residuals and linear bottlenecks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, June 2018
18. Sarker, M.M.K., et al.: SLSDeep: skin lesion segmentation based on dilated residual and pyramid pooling networks. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) MICCAI 2018. LNCS, vol. 11071, pp. 21–29. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00934-2_3
19. Smith, L.N.: Cyclical learning rates for training neural networks. In: 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 464–472. IEEE (2017)
20. American Cancer Society: Cancer facts and figures 2019 (2019)
21. Tschandl, P., Rosendahl, C., Kittler, H.: The HAM10000 dataset, a large collection of multi-source dermoscopic images of common pigmented skin lesions. *Sci. Data* **5**, 180161 (2018)
22. Vasconcelos, M.J.M., Rosado, L., Ferreira, M.: Principal axes-based asymmetry assessment methodology for skin lesion image analysis. In: Bebis, G., et al. (eds.) ISVC 2014. LNCS, vol. 8888, pp. 21–31. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-14364-4_3

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.