



UNIVERSITÀ
DEGLI STUDI
FIRENZE

FLORE

Repository istituzionale dell'Università degli Studi di Firenze

Skin Cancer Classification Using Inception Network and Transfer Learning

Questa è la Versione finale referata (Post print/Accepted manuscript) della seguente pubblicazione:

Original Citation:

Skin Cancer Classification Using Inception Network and Transfer Learning / Priscilla Benedetti; Damiano Perri; Marco Simonetti; Osvaldo Gervasi; Gianluca Reali; Mauro Femminella. - ELETTRONICO. - LNTCS 12249:(2020), pp. 536-545. (Intervento presentato al convegno International Conference on Computational Science and Its Applications tenutosi a Cagliari nel 01/07/2020 - 04/07/2020) [10.1007

Availability:

The webpage <https://hdl.handle.net/2158/1293224> of the repository was last updated on 2022-12-09T15:57:55Z

Publisher:

SPRINGER INTERNATIONAL PUBLISHING AG

Published version:

DOI: 10.1007/978-3-030-58799-4_39

Terms of use:

Open Access

La pubblicazione è resa disponibile sotto le norme e i termini della licenza di deposito, secondo quanto stabilito dalla Policy per l'accesso aperto dell'Università degli Studi di Firenze (<https://www.sba.unifi.it/upload/policy-oa-2016-1.pdf>)

Publisher copyright claim:

Conformità alle politiche dell'editore / Compliance to publisher's policies

Questa versione della pubblicazione è conforme a quanto richiesto dalle politiche dell'editore in materia di copyright.

This version of the publication conforms to the publisher's copyright policies.

La data sopra indicata si riferisce all'ultimo aggiornamento della scheda del Repository FloRe - The above-mentioned date refers to the last update of the record in the Institutional Repository FloRe

(Article begins on next page)

Skin Cancer Classification using Inception Network and Transfer Learning

Priscilla Benedetti¹ *ORCID:0000-0002-3029-0681*
Damiano Perri² *ORCID:0000-0001-6815-6659*
Marco Simonetti² *ORCID:0000-0003-2923-5519*
Osvaldo Gervasi³ *ORCID:0000-0003-4327-520X*
Gianluca Reali¹ *ORCID:0000-0002-8567-5917*
Mauro Femminella¹ *ORCID:0000-0002-6695-5956*

¹ University of Perugia, Dept. of Engineering, Perugia, Italy

² University of Florence, Dept. of Mathematics and Computer Science, Florence, Italy

³ University of Perugia, Dept. of Mathematics and Computer Science, Perugia, Italy

Abstract. Medical data classification is typically a challenging task due to imbalance between classes. In this paper, we propose an approach to classify dermatoscopic images from HAM10000 (Human Against Machine with 10000 training images) dataset, consisting of seven imbalanced types of skin lesions, with good precision and low resources requirements. Classification is done by using a pretrained convolutional neural network. We evaluate the accuracy and performance of the proposal and illustrate possible extensions.

Keywords: Machine Learning, Convolutional Neural Network, Keras, TensorFlow

1 Introduction

Training of neural networks for automated diagnosis of pigmented skin lesions can be a difficult process due to the small size and lack of diversity of available datasets of dermatoscopic images. The HAM10000 (“Human Against Machine with 10000 training images”) dataset is a collection of dermatoscopic images from different populations acquired and stored by different modalities. We used the benchmark dataset, with a small number of images and a strong imbalance among the 7 different types of lesions, to prove the validity of our approach, which is characterized by good results and light usage of resources.

Exploiting a highly engineered convolutional neural network with transfer learning, customized data augmentation and a non-adaptive optimization algorithm, we show the possibility of obtaining a final model able to precisely recognize multiple categories, although scarcely represented in the dataset. The whole training process has a limited impact on computational resources, requiring no more than 20 GB of RAM space. The rest of paper is structured as follows:

Section 2 describes the related work in the field of medical images processing. Section 3 illustrates the dataset of interest. Section 4 gives an overview of the model architecture. Section 5 includes the training process and shows experimental results. Finally, some final comments and future research directions are reported in Section 6.

2 Related work

Processing of biomedical images has always been a field strongly beaten by CNN pioneers. The first related papers date back to 1991 [1], with a strong impulse in the following years in the search for methods for automating the classification of pathologies and related diagnosis [2, 3].

Nowadays, almost thirty years later, reliability of networks reached a rather high level, as well as intrinsic complexity. This reliability allowed a wide diffusion of the approach of subjecting diagnostic images to automatic classification systems, from evolutionary algorithms [4–7] to deep networks [8–11], being them either convolutive or not. Even in the medical sector of dermatology, automatic image recognition and classification was used for decades to detect tumor skin lesions [12, 13].

Recent and promising research has highlighted the possibility that properly trained machines can exceed the human recognition and classification capability to recognize skin cancers. The scores obtained are very encouraging [14] and we are confident that in the near future the recognition capacity of these forms of pathologies will become almost total.

Today CNNs are used for image feature extraction. Features are used for image classification [15, 16].

3 The Dataset

Dermatoscopy is often used to get better diagnoses of pigmented skin lesions, either benign or malignant. With dermatoscopic images is also possible to train artificial neural networks to recognize pigmented skin lesions automatically. Nevertheless, training requires the usage of a large number of samples, although the number of high quality images with reliable labels is either limited or restricted to only a few classes of diseases, often unbalanced.

Due to these limitations, some previous research activities focused on melanocytic lesions (in order to differentiate between a benign and malignant sample) and disregarded non-melanocytic pigmented lesions, even if very common. In order to boost research on automated diagnosis of dermatoscopic images, HAM10000 has been providing the participant of the ISIC 2018 classification challenge, hosted by the annual MICCAI conference in Granada, Spain [17], specific images.

The set of 10015 8-bit RGB color images were collected in 20 years from populations from two different sites, specifically the Department of Dermatology of the Medical University of Vienna, and the skin cancer practice of Cliff

Rosendahl in Queensland. Relevant cases include a representative collection of all important diagnostic categories of pigmented lesions [17]:

- **akiec**: Actinic Keratoses and Intraepithelial Carcinoma, common noninvasive variants of squamous cell carcinoma that can be treated locally without surgery. [327 images]
- **bcc**: Basal cell carcinoma, a cancer that rarely metastasizes but grows destructively if untreated. [514 images]
- **bkl**: Generic label that includes seborrheic keratoses, solar lentigo and lichenplanus like keratoses (LPLK), which corresponds to a seborrheic keratosis or a solar lentigo with inflammation and regression, often mistaken for melanoma. [1099 images]
- **df**: Dermatofibroma, a benign skin lesion. [115 images]
- **nv**: Melanocytic nevi are benign neoplasms of melanocytes. [6705 images]
- **mel**: Melanoma, if diagnosed in an early stage, it can be cured by simple surgical excision. [1113 images]
- **vasc**: Vascular skin lesions in the dataset range from cherry angiomas to angiokeratomas and pyogenic granulomas. [142 images]

More than 50% of lesions are confirmed through histopathology (histo), the ground truth for the rest of the cases is either follow-up examination (followup), expert consensus (consensus), or confirmation by in-vivo confocal microscopy (confocal). Other features in the individual dataset include age, gender and body-site of lesion (localization) [17].

4 Model Architecture

Since their first appearance in Le Cun *et al.* publication [18], Convolutional Neural Networks (CNN) have been widely applied to data that have a known, grid-like structure. Possible set of interests are time-series data, which can be modeled as a 1D grid taking samples at regular time intervals, and image data, which can be thought of as a 2D grid of pixels. The foundational layer of a convolutional network consists of three stages. In the first stage, the layer performs several parallel convolutions to produce a set of linear activations. In the second stage, each convolution output is run through a nonlinear activation function, such as the rectified linear activation function (ReLU). In the third stage, a pooling function is used to modify the output of the layer further [19]. The max pooling operation, used in this work, reports the maximum output within a rectangular neighborhood. Since the objective images of skin lesions present great variation in size, we decided to use a network made by inception modules, which make use of filters of different size operating at the same level. This usage of wide modules with multiple cheap convolutional operations entails a reduced computational complexity with respect to a deep network with large convolutional layers. In the specific model we used, based on Inception-ResNet-v2, another point of speed improvement is the introduction of residual connections, which replace pooling operations within the main inception modules. However, the previously

mentioned max pooling operations are still present in the reduction blocks. The structure of the network used in this work is shown on Fig.1

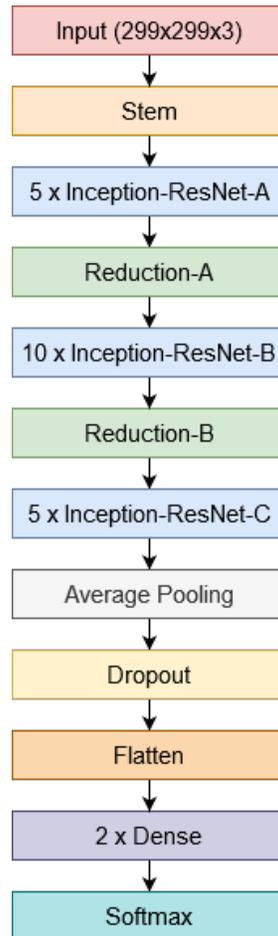


Fig. 1. General schema for the used network, an Inception-ResNet-v2 architecture with the addition of a flattening layer, 2 fully-connected layers and a final softmax activation.

The original Inception-ResNet-v2 architecture [20] has a stem block consisting of the concatenation of multiple convolutional and pooling layers, while Inception-ResNet blocks (A, B and C) contain a set of convolutional filters with an average pooling layer. As previously mentioned, reduction blocks (A, B) replace the average pooling operation with a max pooling one. This structure has been extended with a final module consisting of a flattening step, two fully-

connected layers of 64 units each, and the softmax classifier. The overall module is trainable on a single GPU with reduced memory consumption.

5 Training process and experimental results

This work consists of two training rounds, after a step of data processing in order to deal with the strong imbalance of the dataset:

- A first classification training process using class weights.
- Rollback of previous obtained best model to improve classification performance with a second training phase.

5.1 Data Processing

In the first stage of data processing, after the creation of a new column with a more readable definition of labels, each class was translated into a numerical code using *pandas.Categorical.codes*. Afterwards, missing values in "age" column was filled with column mean value. Fig.2 and Fig.3 show the HAM10000 data distribution.

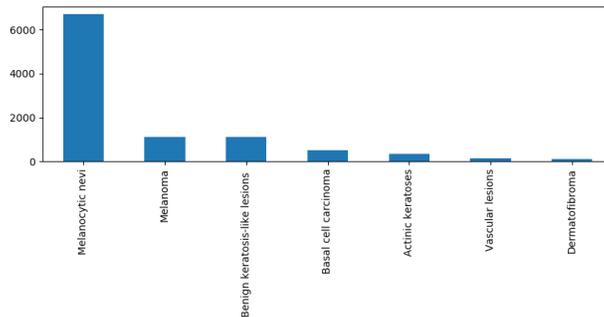


Fig. 2. Plotting the frequency of each class, the imbalance between Melanocytic Nevi and the rest of the possible categories is manifest.

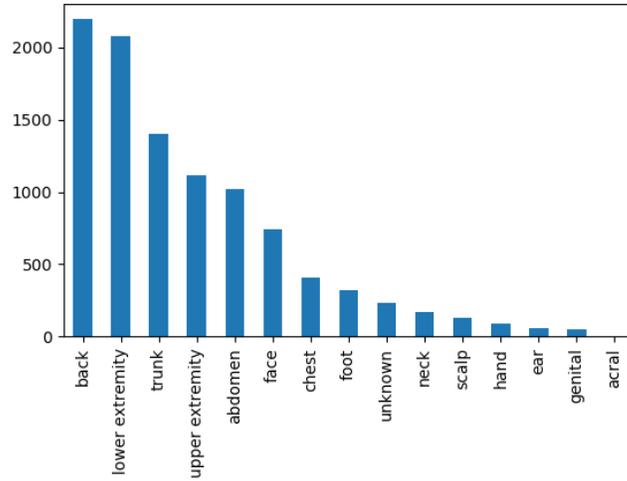


Fig. 3. Plot of samples' body locations.

Finally, images are loaded and resized from 450×600 to 299×299 in order to be correctly processed by the network. After a normalization step on RGB arrays, we split the dataset into a training and validation set with 80:20 ratio.



Fig. 4. Some HAM10000 images resized with OpenCV.

In order to re-balance the dataset, we chose to shrink the amount of images for each class to an equal maximum dimension of 450 samples. This significant decrease of available images is then mitigated by applying a step of data augmentation. Training set expansion is made by altering images with small transformations to reproduce some variations, such as horizontal flips, vertical flips, translations, rotations and shearing transformations.

5.2 Baseline

Due to the limited number of samples for the training process, we decided to take advantage of transfer learning, utilizing Inception-ResNet-v2 pre-trained on ImageNet [21] and Tensorflow, a deep learning framework developed by Google, for fine-tuning of the last 40 layers. Keras library offers a wide range of optimizers: Adaptive optimization methods such as AdaGrad, RMSProp, and Adam are widely used for deep neural networks training due to their fast convergence times. However, as described in [22], when the number of parameters exceeds the number of data points these optimizers often determine a worse generalization capability compared with non-adaptive methods. In this work we used a stochastic gradient descent optimizer (SGD), with learning rate set to 0.0006 and usage of momentum and Nesterov Accelerated Gradient in order to adapt updates to the slope of the loss function (categorical crossentropy) and speed up the training process. The total number of epochs was set to 100, using a small batch size of 10. A set of class weight was introduced in the training process to get more emphasis on minority class recognition. A maximum patience of 15 epochs was set to the early stopping callback in order to mitigate the overfitting visible in Fig.5, which shows the history of training and validation process.

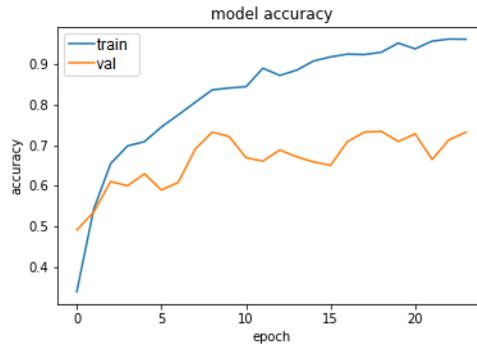


Fig. 5. Accuracy and Loss for each epoch.

Finally, the model achieves an accuracy of 73.4% on the validation set, using weights from the best epoch. Fig.6 shows the confusion matrix for the model on the validation set.

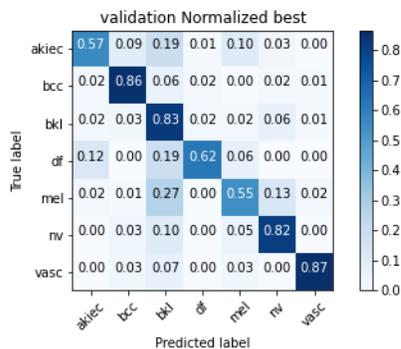


Fig. 6. Two of the minority classes, Actinic Keratoses (akiec) and Dermatofibroma (df), are not properly recognized. Melanoma (mel) is often mistaken for generic keratoses (bkl), as already mentioned in 3

5.3 Resuming training from the best epoch:

In order to improve classification performance, specially on minority classes, we loaded the best model obtained in the first round to extend the training phase and explore other potential local minimum points of the loss function, by using an additional amount of 20 epochs. This second step led to an enhancement in overall predictions, reaching the maximum accuracy value of 78.9%.

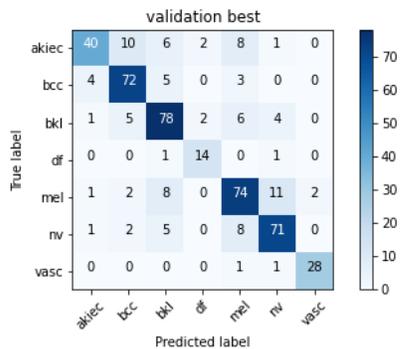


Fig. 7. Final confusion matrix.

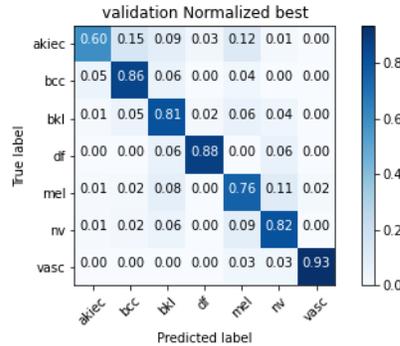


Fig. 8. Normalized final confusion matrix.

Fig.7 shows the normalized confusion matrix on the validation set for the final fine-tuned model. In this case, 6 out of 7 categories are classified with a total ratio of True Positives higher than 75%, even in presence of extremely limited sample set, as vascular lesions (vasc), 30 samples, and dermatofibroma (df), 16 samples. The whole process of training has required less than four hours on Google Colab cloud’s GPU, for an overall RAM utilization below 20 GB.

6 Conclusions and future works

In conclusion, in this paper we investigate the possibility of obtaining improved performances in the classification of 7 significantly unbalanced different types of skin diseases, with a small amount of available images. With use of a fine-tuned deep inception network, data augmentation and class weights, the model can achieve a good final diagnostic accuracy. The described training process has a light resource usage, requiring less than 20 GB of RAM space, and it can be executed in a Google Colab notebook. For future improvements larger datasets of dermatoscopic images are needed. The model shown in this paper can be regarded as a starting point to implement a lightweight diagnostic support system for dermatologists, for example in the Web as well as through a mobile application.

References

1. Wei Zhang and Akira Hasegawa and Kazuyoshi Itoh and Yoshiki Ichioka, Image processing of human corneal endothelium based on a learning network, 4211–4217 (1991)
2. Zhang, Wei and Doi, Kunio and Giger, Maryellen L. and Wu, Yuzheng and Nishikawa, Robert M. and Schmidt, Robert A., Computerized detection of clustered microcalcifications in digital mammograms using a shift-invariant artificial neural network (1994)

3. Vella, F., Neri, I., Gervasi, O., Tasso, S. A simulation framework for scheduling performance evaluation on CPU-GPU heterogeneous system, *Lecture Notes in Computer Science*, 7336, ICCSA 2012, pp. 457–469, Springer, DOI: 10.1007/978-3-642-31128-4_34, (2012).
4. Jong-Chen Chen and Chun-Ming Yeh and Jeh-En Tzeng, Pattern differentiation of glandular cancerous cells and normal cells with cellular automata and evolutionary learning (2008)
5. Guo, Pei Fang and Bhattacharya, Prabir, An Evolutionary Approach to Feature Function Generation in Application to Biomedical Image Patterns (2009)
6. Gervasi, O., Russo, D. and Vella, F.: The AES Implantation Based on OpenCL for Multi/many Core Architecture, 2010 International Conference on Computational Science and Its Applications, Fukuoka, ICCSA 2010, pp. 129–134, Washington, DC, USA, IEEE Computer Society, DOI: 10.1109/ICCSA.2010.44, 2010.
7. Mariotti, M., Gervasi, O., Vella, F., Cuzzocrea, A. and Costantini, A.: Strategies and systems towards grids and clouds integration: A DBMS-based solution, *Future Generation Computer Systems*, vol. 88, pp. 718–729, <http://dx.doi.org/10.1016/j.future.2017.02.047>, 2018.
8. Pang, Shuchao and Yu, Zhezhou and Orgun, Mehmet A., A novel end-to-end classifier using domain transferred deep convolutional neural networks for biomedical images (2017)
9. Zhou, Zongwei and Shin, Jae and Zhang, Lei and Gurudu, Suryakanth and Gotway, Michael and Liang, Jianming, Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally (2017)
10. Gervasi, Osvaldo, Magni, Riccardo and Ferri, Matteo: A Method for Predicting Words by Interpreting Labial Movements, *Lecture Notes in Computer Science*, vol. 9787, pages 450–464, ICCSA 2016, Beijing (China), Springer, DOI: 10.1007/978-3-319-42108-7_34, 2016.
11. Riganelli M., Franzoni V., Gervasi O., and Tasso S., EmEx, a Tool for Automated Emotive Face Recognition Using Convolutional Neural Networks, ICCSA 2017, Workshop on Emotion Recognition, *Lecture Notes in Computer Science*, vol. 10406, pp. 692–704, DOI: 10.1007/978-3-319-62398-6_49, 2017.
12. Lau, Ho Tak and Al-Jumaily, Adel, Automatically early detection of skin cancer: Study based on neural network classification (2009)
13. Dorj, Ulzii-Orshikh and Lee, Keun-Kwang and Choi, Jae-Young and Lee, Malrey, The skin cancer classification using deep convolutional neural network (2018)
14. Maron, Roman C and Weichenthal, Michael and Utikal, Jochen S and Hekler, Achim and Berking, Carola and Hauschild, Axel and Enk, Alexander H and Haferkamp, Sebastian and Klode, Joachim and Schadendorf, Dirk and others, Systematic outperformance of 112 dermatologists in multiclass skin cancer image classification by convolutional neural networks (2019)
15. Biondi Giulio, Franzoni Valentina, Gervasi Osvaldo, Perri Damiano, An Approach for Improving Automatic Mouth Emotion Recognition, LNCS 11619, pp. 649–664, doi: https://doi.org/10.1007/978-3-030-24289-3_48, ISBN=978-3-030-24289-3 (2019)
16. Perri, D., Sylos Labini, P., Gervasi, O., Tasso, S., Vella, F.: Towards a Learning-Based Performance Modeling for Accelerating Deep Neural Networks, LNCS 11619, pp. 665–676, doi: https://doi.org/10.1007/978-3-030-24289-3_49, ISBN:978-3-030-24289-3 (2019)
17. Tschandl, P., Rosendahl, C. and Kittler, H., The HAM10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions, *Sci Data* 5, DOI: <https://doi.org/10.1038/sdata.2018.161>, 2018.

18. Lecun, Y., Bottou, L., Bengio, Y., Haffner, P., Gradient-based learning applied to document recognition. Proceedings of the IEEE, vol. 86, no. 11, pp. 2278–2324, doi: 10.1109/5.726791. (1998)
19. Goodfellow, I., Bengio, Y. and Courville, A., Deep Learning, MIT press, 2016.
20. Szegedy, C., Ioffe, S., Vanhoucke, V. and Alemi, A., Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning, 2016
21. Alex Krizhevsky, Ilya Sutskever and Geoff Hinton: Imagenet classification with deep convolutional neural networks, 25th International Conference on Advance in Neural Information Processing System, pag.1106–1114, 2012.
22. Wilson, A., Roelofs, R., Stern, M., Srebro, N. and Rech, B., The Marginal Value of Adaptive Gradient Methods in Machine Learning, NIPS 2017, 2017.