Improving Skin Lesion Segmentation with Generative Adversarial Networks

Federico Pollastri*, Federico Bolelli*, Roberto Paredes †, Costantino Grana*
*DIEF, Università degli Studi di Modena e Reggio Emilia, Italy
†PRHLT Research Center, Universitat Politécnica de Valencia, Spain

UNIMORE UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

UNIVERSITAT POLITÈCNICA DE VALÈNCIA

AImageLab

PRHLT

Problem Statement

Malignant melanoma is the most dangerous skin cancer, with a substantial death rate. Automated skin lesion segmentation is a fundamental step to help experts in early diagnosis, but requires a huge amount of data to be performed. Unfortunately, manual image segmentation is a very time consuming job that demands the work of a competent specialist.

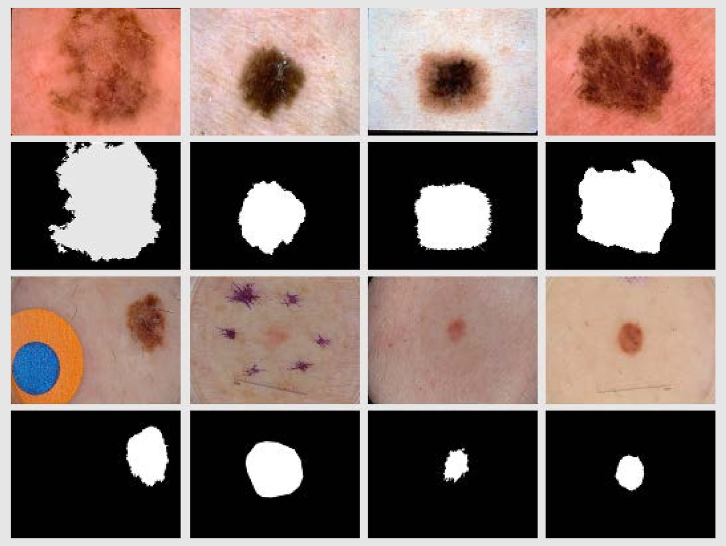

Samples from the ISIC dataset.

Goals:

- **decrease the amount of manually segmented images required by automated analysis;**
- **support the clinical decision making.**

Baseline Architecture

Our model maps the input dermoscopic image into a posterior probability map, exploiting an architecture based on the CDNN which won the International Skin Imaging Collaboration (ISIC) challenge in 2017 *[1]*.

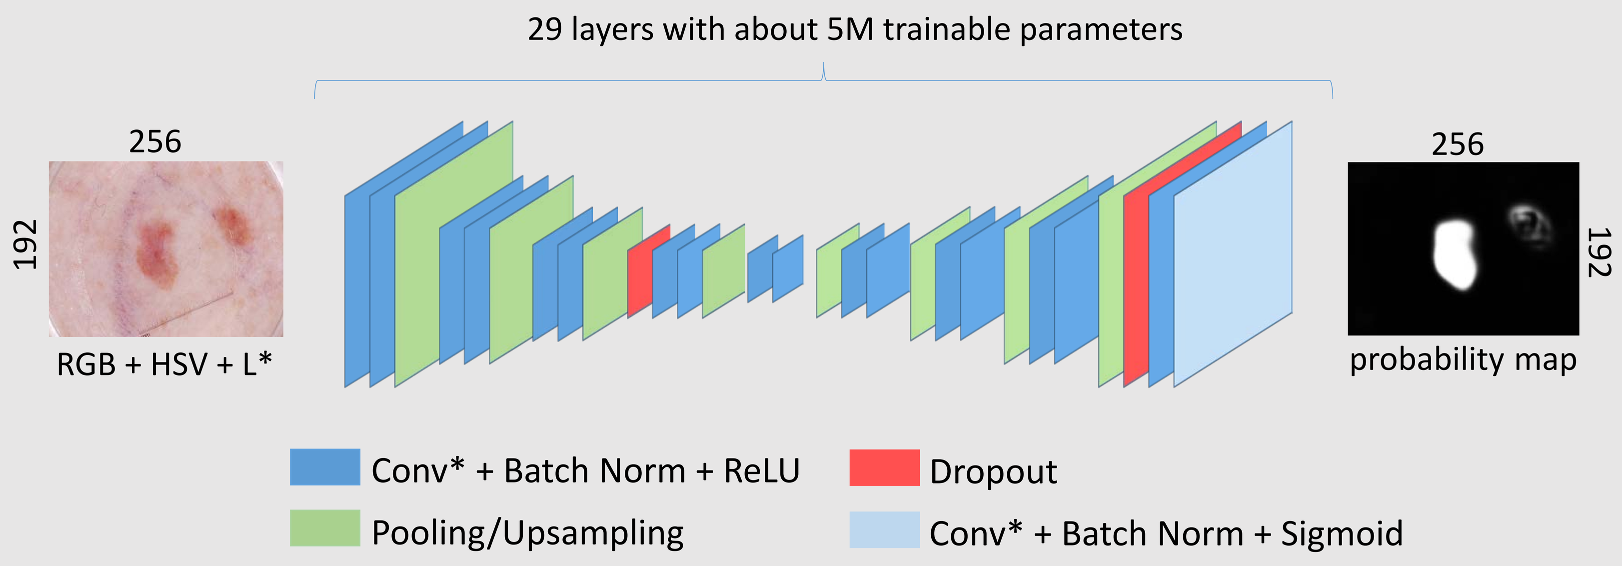

* fixed stride of 1 pixel.

[1] Yuan, Yading, Ming Chao, and Yeh-Chi Lo. "Automatic skin lesion segmentation with fully convolutional-deconvolutional networks." arXiv preprint arXiv:1703.05165 (2017).

Hyperparameters Analysis

The baseline CDNN is mainly affected by three hyperparameters, which are stressed in our analysis:

1) *Input image size*:
   - 96x128
   - 192x256
   - 384x512
2) *Image channels*:
   - RGB
   - RGB + HSV + L*
   - RGB + HSV + CIELAB
3) *Loss function:*

$$L = 1 - \frac{\sum_{i,j} t_{ij} p_{ij}}{\sum_{i,j} t_{ij}^2 + \sum_{ij} p_{ij}^2 - \sum_{i,j} t_{ij} p_{ij}} \qquad (Eq.\ 1)$$

$$d_j = 1 - \frac{\sum_{i,j} t_{ij} p_{ij}}{\sum_{i,j} t_{ij} + \sum_{ij} p_{ij} - \sum_{i,j} t_{ij} p_{ij}} \qquad (Eq.3)$$

$$L = \frac{1}{n} \sum_{i,j} (t_{ij} - p_{ij})^2 \qquad (MSE)$$

GAN Augmented Data

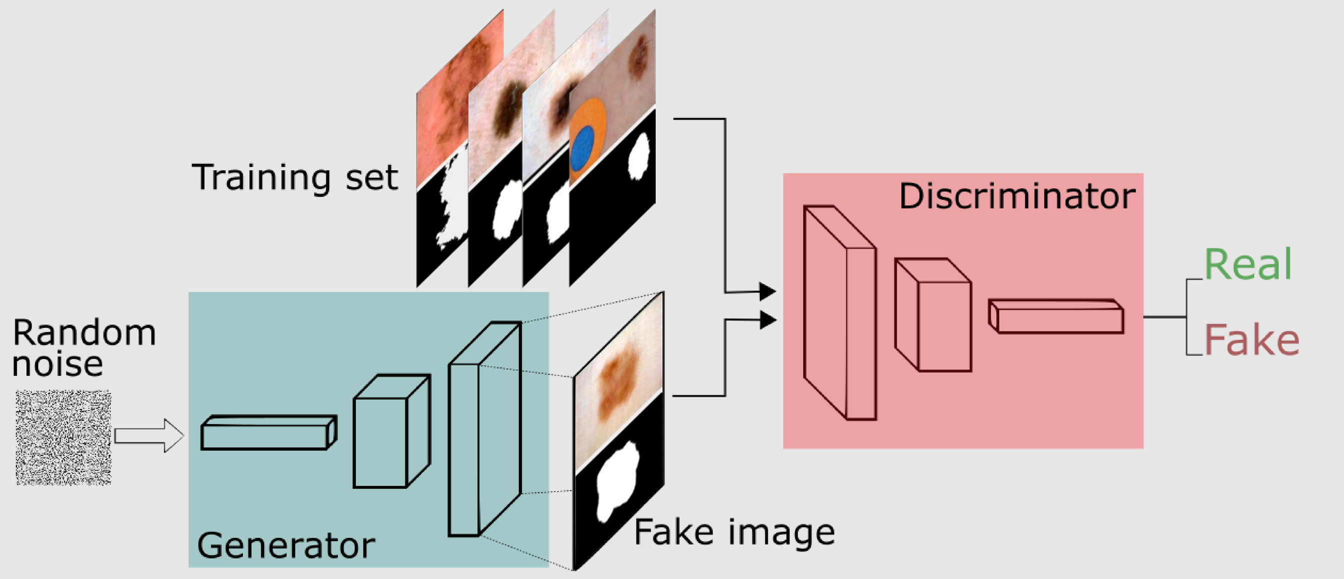

A DCGAN is employed to generate both the skin lesion image and its segmentation mask, improving the data augmentation process.

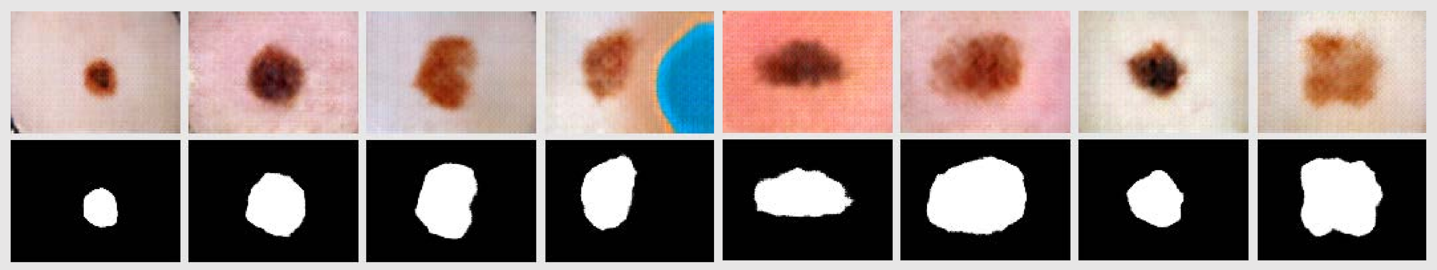

Samples from the generated dataset.

Experimental Results

Neural Network	Input Size	Number of Channels	Loss Function	Without Augmentation	With Augmentation
$CDNN_0$	192x256	7	Eq. 1	0.731	0.743
$CDNN_1$	192x256	3	Eq. 1	0.732	0.753
$CDNN_2$	192x256	9	Eq. 1	0.734	0.743
$CDNN_3$	96x128	7	Eq. 1	0.735	0.750
$CDNN_4$	384x512	7	Eq. 1	0.700	-
$CDNN_5$	192x256	7	Eq.3	0.738	0.738
$CDNN_6$	192x256	7	MSE	0.738	0.739
				Ensemble:	**0.781**

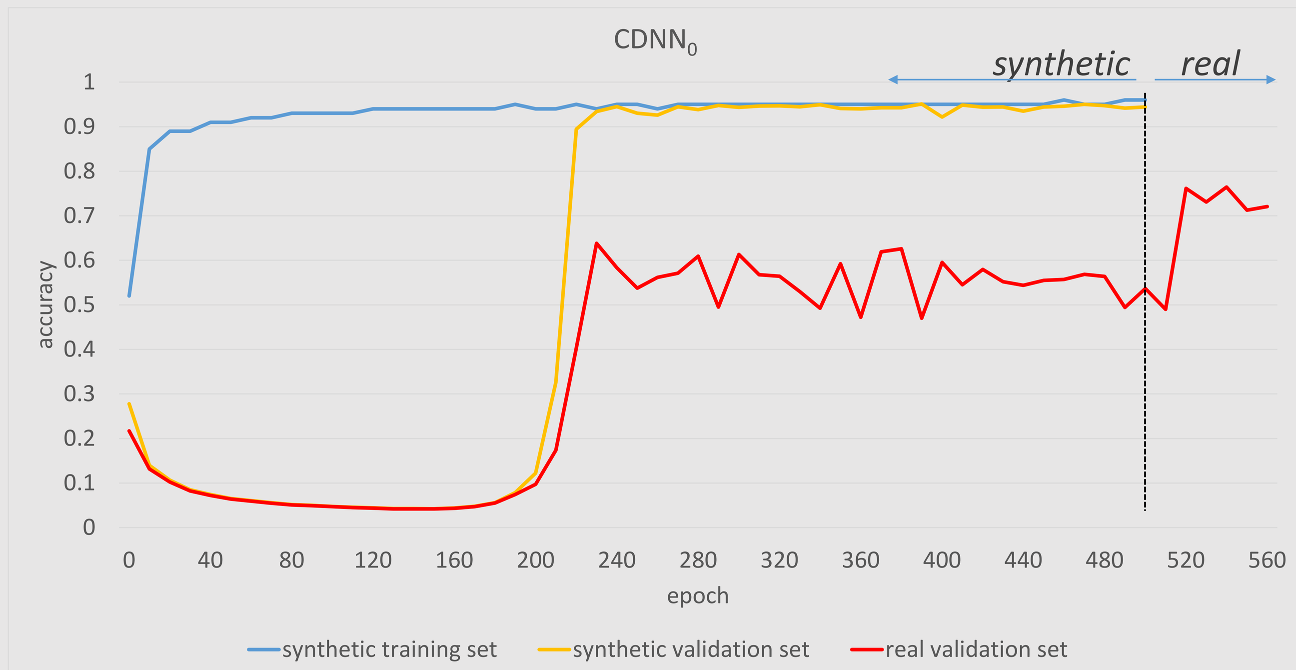

Visualization of the training process: synthetic data are exploited to performed an "initialization" of the CDNN, which is then fine-tuned using the real data.

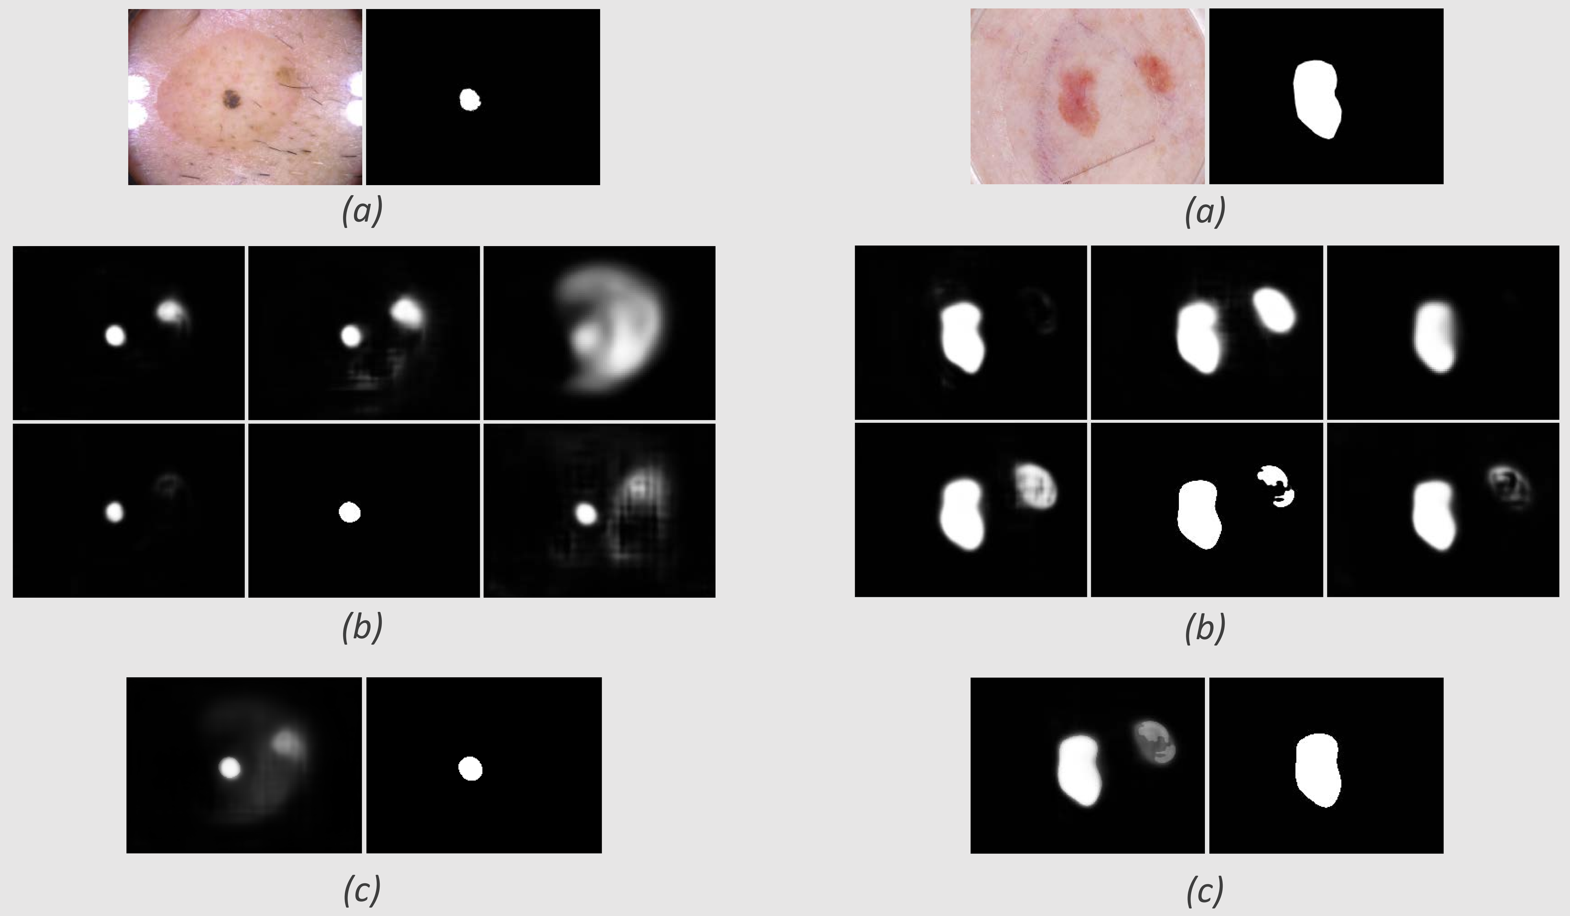

(a) Input image and ground truth. (b) Left-to-right then top-to-bottom, output prediction of $CDNN_1$, $CDNN_2$, $CDNN_3$, $CDNN_0$, $CDNN_5$ and $CDNN_6$. (c) Outputs ensemble before and after binarization.