


Dual Stream Network with Selective Optimization for Skin Disease Recognition in Consumer Grade Images

Introduction

- Skin disease localisation and classification on consumer-grade images is more challenging compared to that on dermoscopic imaging.
- Such images, in addition to having the skin condition of interest in a very small area of the image, has other noisy non-clinical details introduced due to the lighting conditions and the distance of the hand held device from the anatomy at the time of acquisition.
- We propose a novel deep network architecture & a new optimization strategy for classification with implicit localisation of skin diseases from consumer grade images.
- Such a strategy resulted in a 5% increase of accuracy over the current state-of-the-art methods on SD-198 dataset, which is publicly available. The proposed algorithm is also validated on a new dataset containing over 12,000 images across 75 different skin conditions.

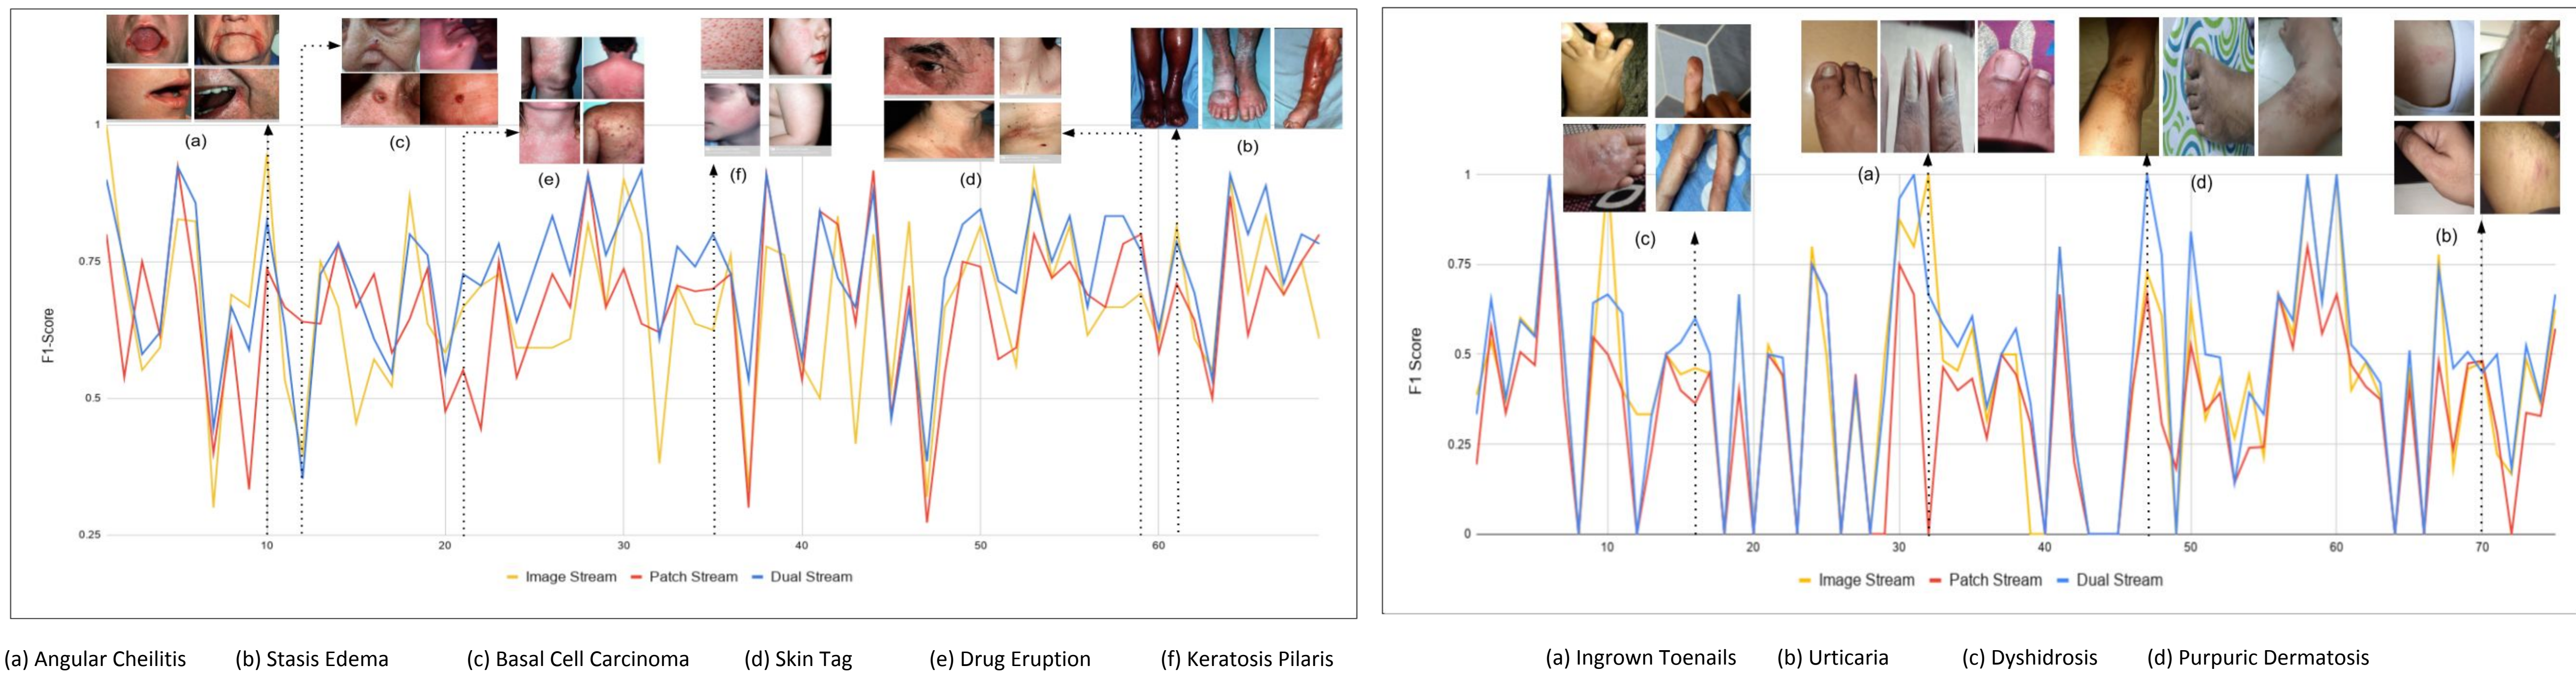


Methodology



- A weakly supervised segmentation algorithm is first employed to extract Region of Interests (RoI) from the image.
- The RoI and the original image form the two input streams of the proposed architecture. Each stream of the architecture learns high level and low level features from the original image and the RoI, respectively.
- The two streams are independently optimised until the loss stops decreasing after which both the streams are optimised collectively with the help of a third combiner sub-network.

Results



Method	Train-Test Split	Precision	Recall	F1	Accuracy	G-Mean	M-AUC
SIFT+CN+SVM	50:50	52.87	N/R	N/R	N/R	N/R	N/R
Derma Criteria	50:50	57.87	N/R	N/R	N/R	N/R	N/R
Dual Stream (Ours)	50:50	61.1 ± 1.2	60.2 ± 1.3	60.6 ± 1.2	60.9 ± 1.3	38.8 ± 3.3	60.7 ± 0.9
SPBL	80:20	71.4 ± 1.7	65.7 ± 1.6	66.2 ± 1.6	67.8 ± 1.8	42.8 ± 4	68.5 ± 1.6
Dual Stream (Ours)	80:20	73.1 ± 1.4	69.2 ± 1.1	70.9 ± 1.2	71.4 ± 1.1	46.8 ± 3.7	71.1 ± 1.2