THE DEEP HEALTH EUROPEAN PROJECT AND THE RESEARCH ON MEDICAL IMAGING @AIMAGELAB

Federico Bolelli

federico.bolelli@unimore.it

Università degli Studi di Modena e Reggio Emilia, DIEF, *Italy*

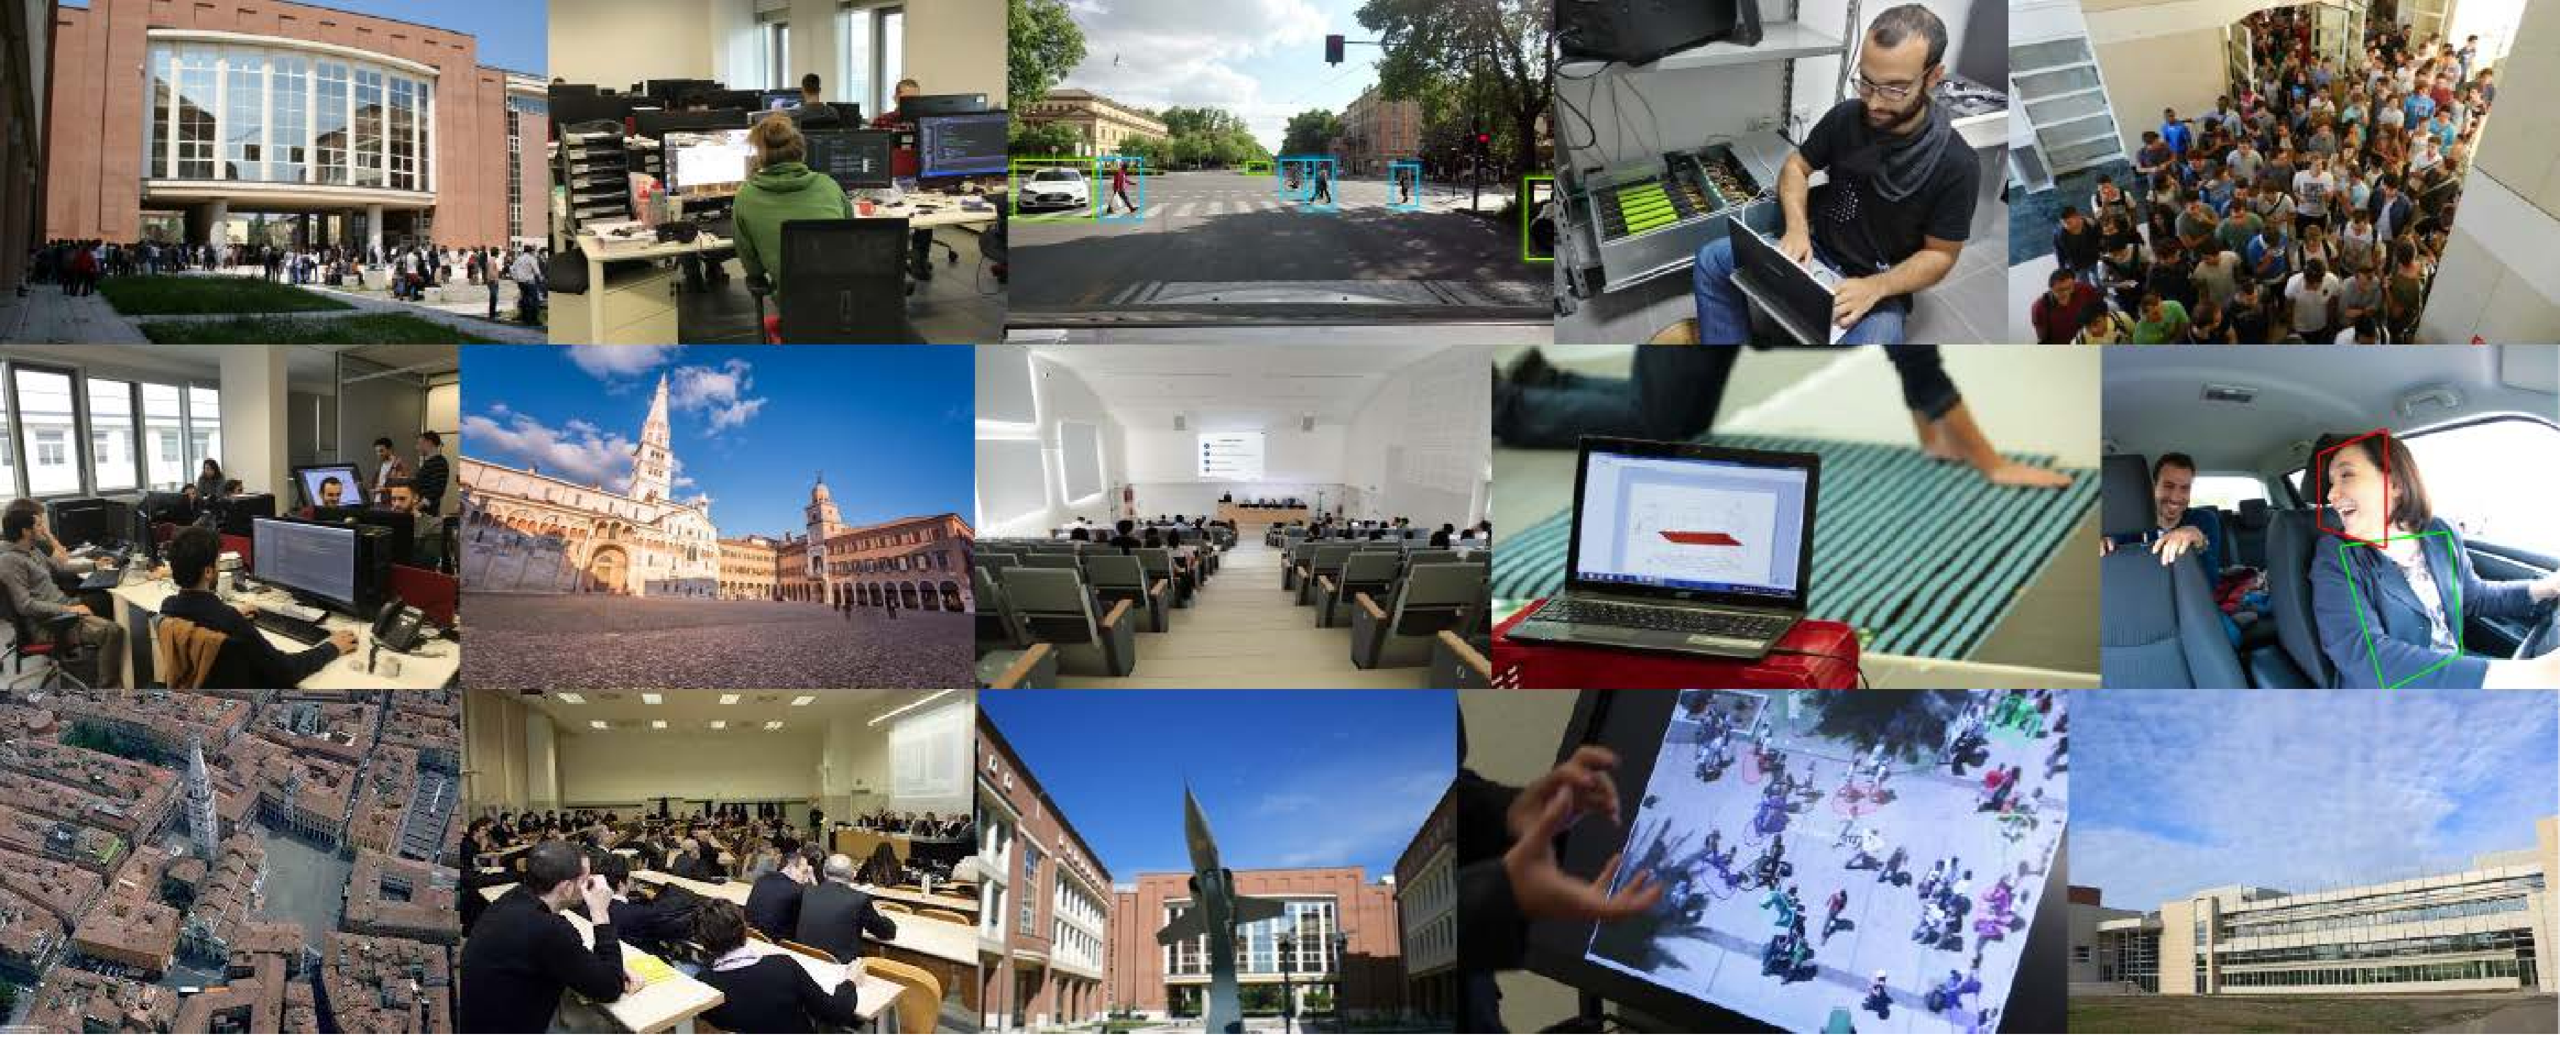

DEEP HEALTH H2020

AImage Lab

UNIMORE UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

DEEPHEALTH: DEEP-LEARNING AND HPC TO BOOST BIOMEDICAL APPLICATIONS FOR HEALTH

AImage Lab

- 21 partners (France, Germany, Greece, Italy, Netherlands, Romania, Spain, Sweden, Switzerland)
- **Duration:** 3 years
- **Period:** 2019 - 2021
- **EU grants:** 12.77 M€
- **Budget:** 14.64 M€
- **Keywords:** High performance computing, Big data, Very large medical image data bases, Scalability, Deep Learning, Artificial Intelligence, Heterogeneous Architectures

DEEPHEALTH

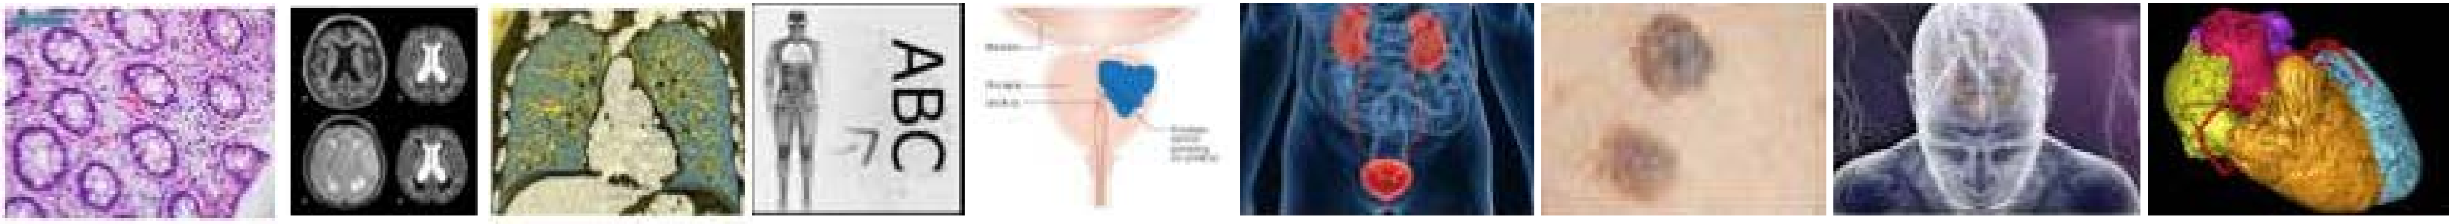

MULTIDISCIPLINARY GROUP OF 21 PARTNERS

AImage Lab

9 Research organizations:

- Universitat Politècnica de València
- Commissariat à l'Énergie Atomique et aux Énergies Alternatives
- Barcelona Supercomputing Center
- Università degli Studi di Torino
- **Università degli Studi di Modena e Reggio Emilia**
- Centro di Ricerca Sviluppo e Studi Superiori in Sardegna
- Otto von Guericke University Magdeburg
- Ecole Polytechnique Federale De Lausanne
- Fundación para el fomento de la Investigación Sanitaria y Biomédica de la Comunitat Valenciana.

4 Health organizations

- Karolinska Institutet
- Azienda Ospedaliera Citta Della Salute E Della Scienza Di Torino
- Centre Hospitalier Universitaire Vaudois, Lausanne
- Spitalul Clinic "Prof. Dr. Th. Burghele"

4 Large industrial partners

- Everis Spain SLU
- Philips Medical Systems Nederland B.V.
- THALES SIX GTS FRANCE SAS
- SIVECO Romania SA

4 SME industrial partners

- WINGS ICT Solutions Information & Communication Technologies Ike
- Tree Technology S.A.
- PRO DESIGN Electronic GmbH
- STELAR SECURITY TECHNOLOGY LAW RESEARCH UG

MAIN OBJECTIVES

AImage Lab

- DeepHealth will provide **HPC** computing power at the service of biomedical applications and apply **Deep Learning** to support new and more efficient ways of diagnosis, monitoring and treatment of diseases.

- DeepHealth will provide enhanced productivity of expert-users (technical staff in hospital, not machine learning experts) working within health care institutions.

MAIN OBJECTIVES

AImage Lab

- The *Deep*Health framework will be based on **two new core technology libraries** to be developed within the project and integrated in existing biomedical software platforms: the **European Distributed Deep Learning Library** (EDDLL) and the **European Computer Vision Library** (ECVL).

- These libraries will take full advantage of the current and coming development of HPC systems deployed over Europe, and **will provide a transparent use of heterogeneous hardware accelerators** to optimize the training of predictive models, while considering performance and accuracy trade-offs.

- *The* ***Deep*****Health** *framework will also* ***include a front-end to simplify the use of the core technology libraries***, both for the manipulation of images and for the training of Deep Neural Networks (DNNs).

MAIN OBJECTIVES

- The purpose of *Deep*Health is to offer a **unified framework** completely adapted to exploit underlaying **heterogeneous HPC and Big Data architectures** and assembled with state-of-the-art techniques in **Deep Learning and Computer Vision** for facilitating the daily work of expert users.

- Expert users need to manage large image datasets and train predictive models.

- *Deep*Health will make it **easy and transparent** to do so in a **distributed way** on hybrid and heterogeneous HPC + Big Data architectures.

TECHNICAL OBJECTIVES

AImage Lab

- **Adapt** existing Machine Learning (ML) and Deep Learning (DL) algorithms to run on HPC systems;
- **Develop efficient** versions of those algorithms that require more computing power to take advantage of hardware accelerators such as GPUs and FPGAs;
- **Make the use** of heterogeneous hardware as **transparent** as possible to expert users, during and beyond the duration of the project;
- Facilitate the **integration** of new techniques and algorithms into the core libraries so that they can easily take advantage of **future hardware designs**;
- Implement an asynchronous protocol to efficiently train DNNs on **distributed** architectures;
- **Integrate** the core technology libraries into existing end-user software platforms to improve them;
- Provide a free-software and open-source toolkit, *i.e.* the ***Deep*****Health** toolkit, for bio-imaging analytics in order to **support** medical **diagnoses** and **studies** on the evolution of diseases.

PROJECT CONTRIBUTIONS – TOOLS AND LIBRARIES

AImage Lab

TL1: A **European Distributed Deep Learning Library (EDDLL)**, free-software and open-source, ready to be integrated into end-user software platforms or applications, and ready to run DL algorithms on Hybrid HPC + Big Data architectures with heterogeneous hardware, including CPU, GPU, FPGA IBM Power AI, etc.

TL2: A **European Computer Vision Library (ECVL)**, free-software and open-source, to facilitate the integration and exchange of data between existing state-of-the-art Computer Vision (CV) and Image Processing libraries (IP).

TL3: A **front-end** for facilitating the use of the functionalities provided by EDDLL and ECVL (the core technology libraries). The front-end will include the executable programs for designing, training and testing predictive models, for image processing, and will have a graphical user interface via web. Front-end + EDDLL + ECVL = the ***Deep*****Health** toolkit.

PROJECT CONTRIBUTIONS – TOOLS AND LIBRARIES

AImage Lab

https://github.com/deephealthproject

DEEPHEALTH

DeepHealth

Repositories 6 | Packages | People 1 | Projects

eddl
European Distributed Deep Learning (EDDL) Library
C · MIT · 4 · 14 · 1 · 0 · Updated 4 minutes ago

docker-libs
Makefile · 2 · 0 · 0 · 0 · Updated 4 hours ago

ecvl
European Computer Vision Library
C++ · MIT · 1 · 4 · 1 · 0 · Updated 3 days ago

pyeddl
Python wrapper for the EDDL library
C++ · MIT · 2 · 2 · 1 · 0 · Updated 5 days ago

pyecvl
Python wrapper for the ECVL
C++ · MIT · 1 · 0 · 0 · 0 · Updated 14 days ago

ceddl
Plain C wrapper for the EDDL library
C · MIT · 0 · 3 · 0 · 0 · Updated on 15 Jan

EUROPEAN COMPUTER VISION LIBRARY (ECVL)

- Already provides 3D viewer tool and all the basic functionalities for pre- and post-processing operations, including data augmentation.

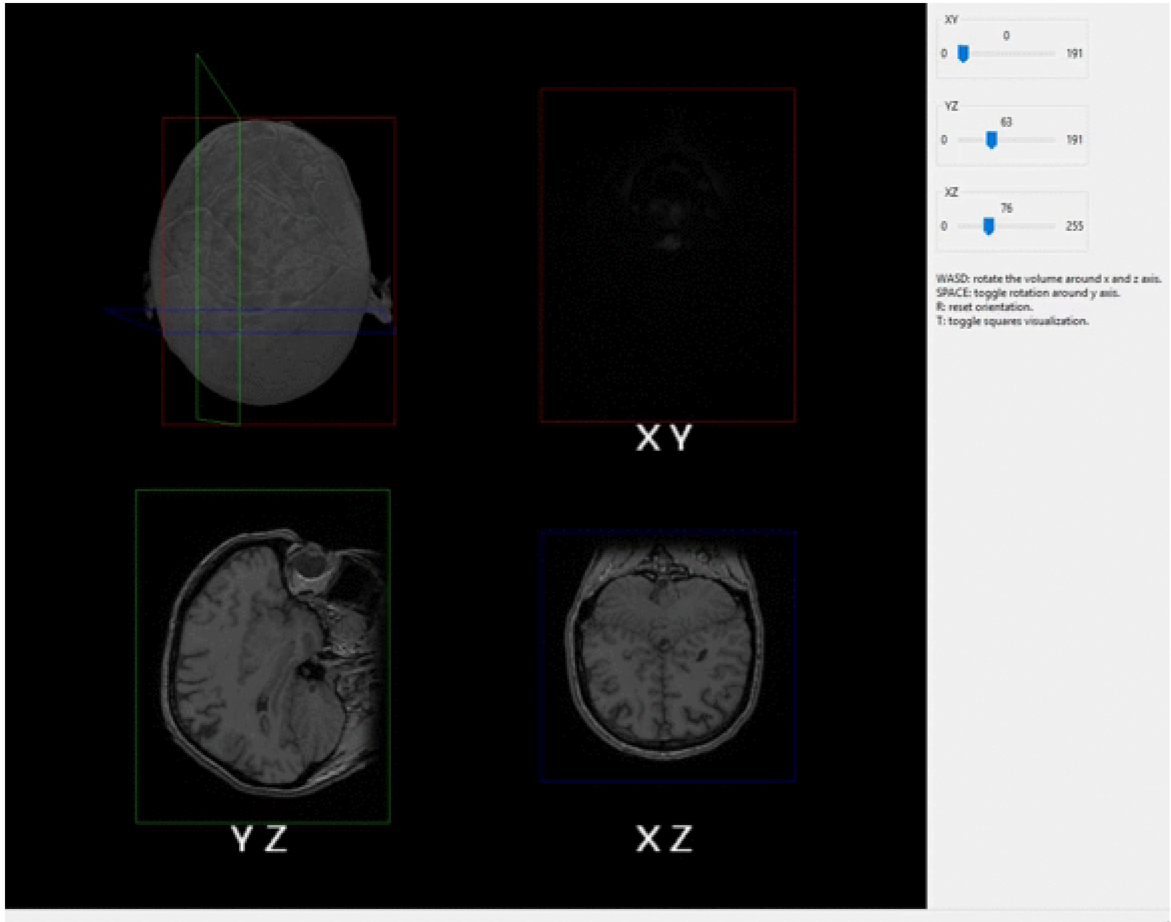

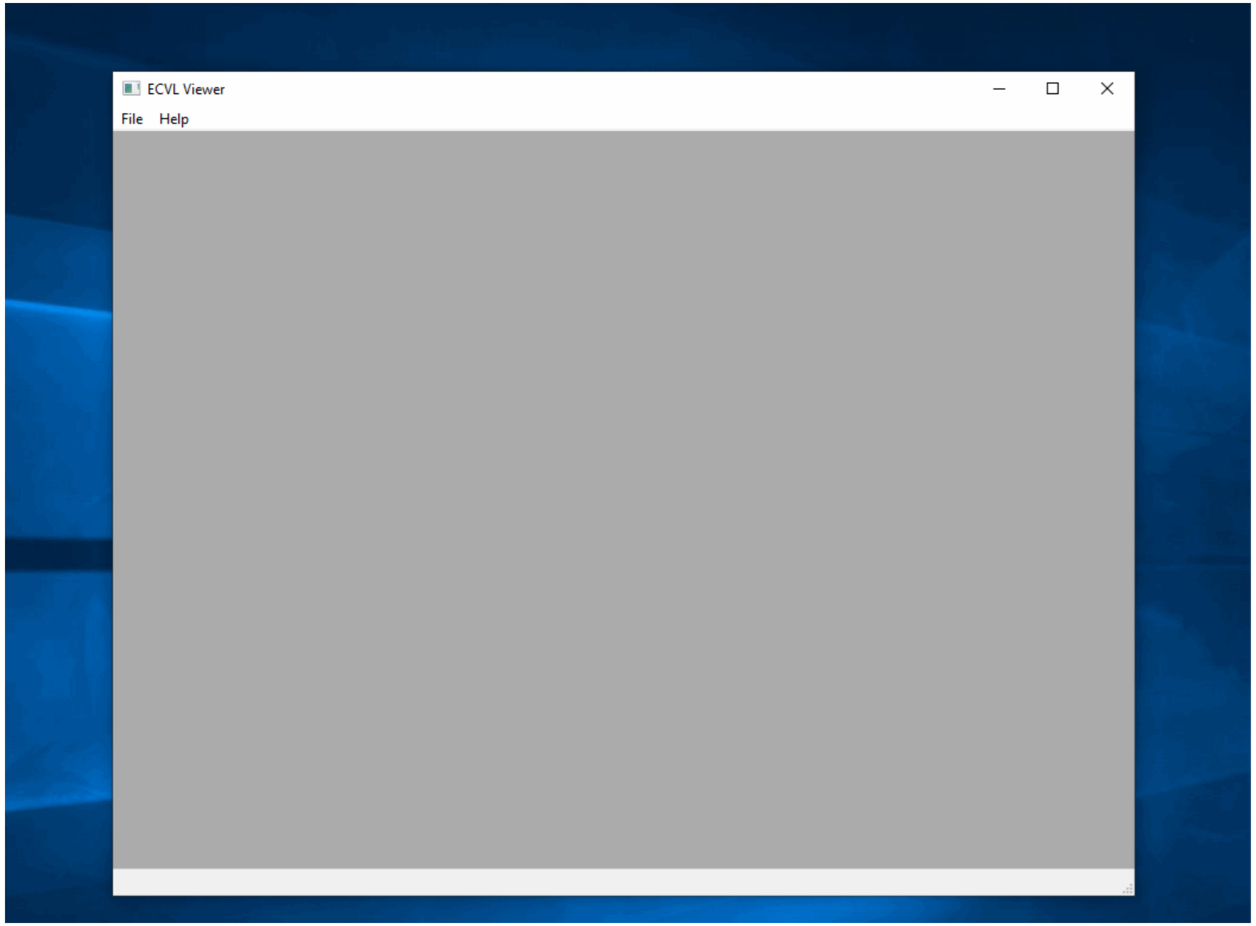

DEMONSTRATORS BASED UPON USE CASES

AImage Lab

- UC1: Migraine & Seizures prediction (WINGS)
- UC2: Path dataset and use-case (UNITO)
- UC3: Brain dataset and use-case (UNITO)
- UC4: Chest dataset and use case (SALUTE)
- UC5: Deep Image Annotation (UNITO)
- UC6: Promort (KAROLINSKA and CRS4)
- UC7: Major Depression (OVGU-MAGDEBURG)
- UC8: Dementia (OVGU-MAGDEBURG)
- UC9: Study of structural changes in lumbar spine pathology (FISABIO)
- UC10: Predictive Model for Alzheimer's Disease using Structural Neuroimaging (FISABIO)
- UC11: Image Analysis and prediction for Urology (SIVECO + SCTH)
- **UC12: Skin cancer melanoma detection (UNIMORE)**
- UC13: Epileptic seizure detection (CHUV NeuroTech)
- UC14: Prediction of the multiple sclerosis patients outcome (CHUV NeuroTech)

KPI	Description
TOTM	Time of Training Models
TOPPI	Time of Pre-Processing Images
TTMIP	Time to Model in Production
Speed-Up	T(1) / T(nt) - where **nt** is the number of execution threads
Efficiency of Parallelism	T(1) / (nt * T(nt))

AMBITION

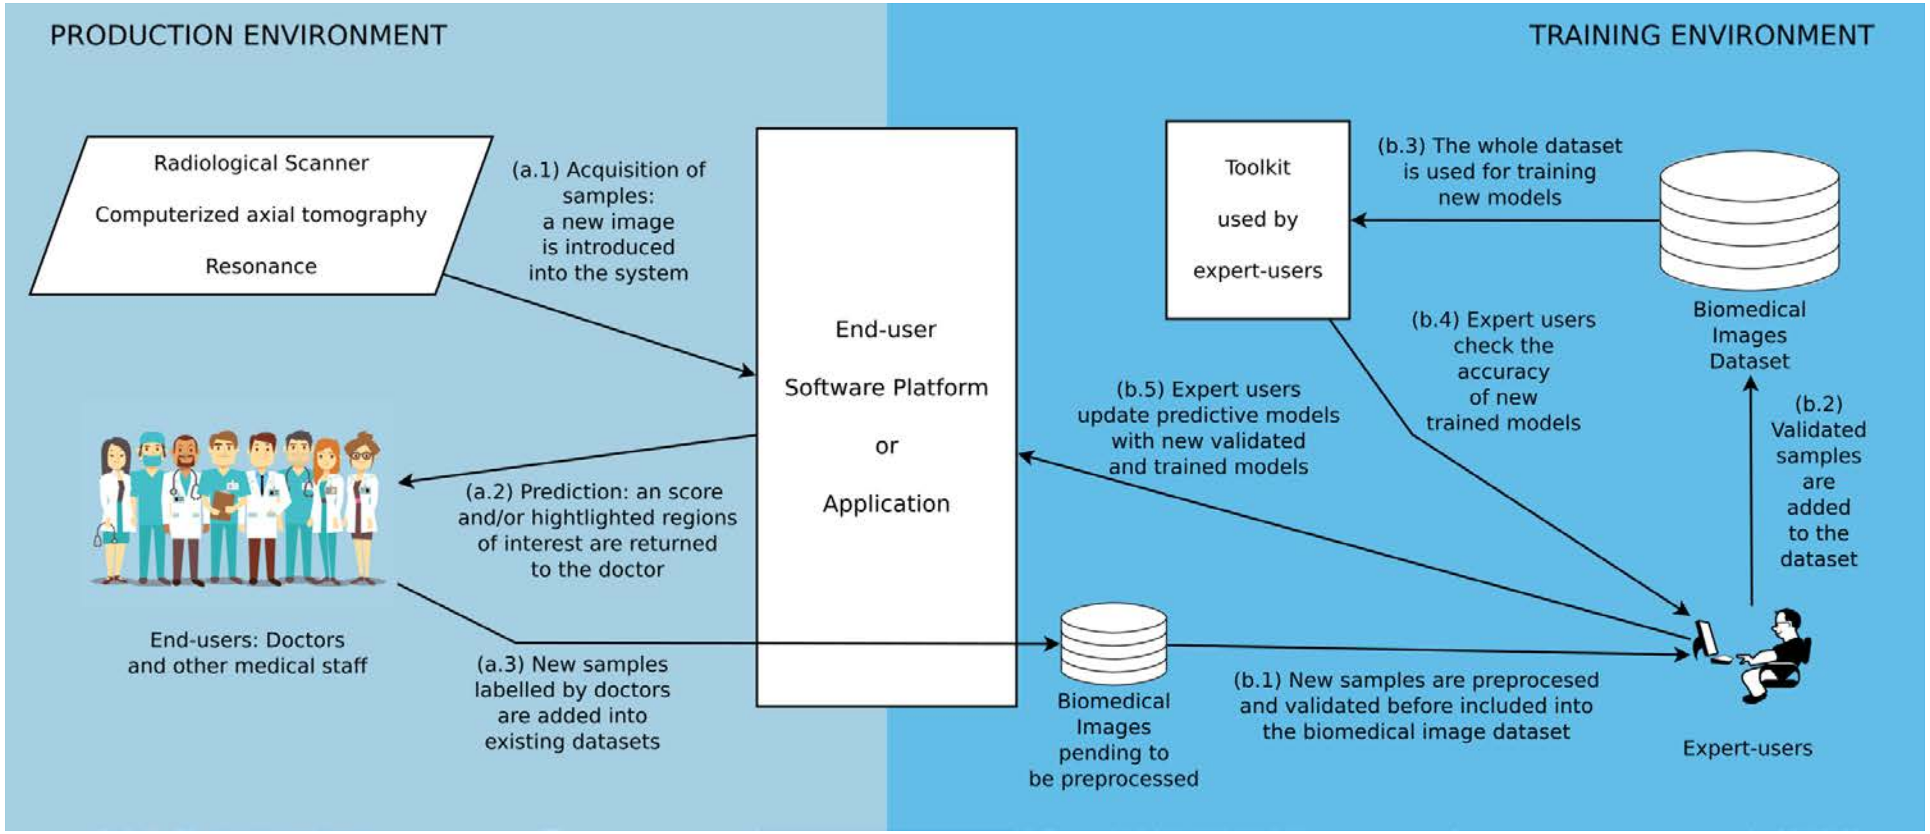

- One of the main ambitions of *Deep*Health is to **increase the productivity of expert users**.
- The technical goal of *Deep*Health is to **cover the needs of expert users**, regardless of whether they provide ad hoc solutions to end users or make use of existing software platforms.

WHY DO WE NEED NEW LIBRARIES?

AImage Lab

- Very well-designed libraries in both Computer Vision/Image Processing and Machine Learning/Deep Learning fields exists

OpenCV

ImageJ

TensorFlow

VTK Kitware

PYTORCH

WHY DO WE NEED NEW LIBRARIES?

- **At the beginning** of the project we will not implement new libraries from scratch, but **we will integrate existing ones**.
- We want to **cover all the needs of expert users,** so the libraries must provide all functionalities and data-formats used in medical fields. State-of-the-art software is often very specific or, when flexible, usually requires computer science expertise to combine modules and obtained the desired result.
- The two libraries must be **completely interoperable** each other **and** must **support processing on** HPC **heterogeneous hardware** (CPU, GPU, FPGA and so on);
- **EU payed us!**

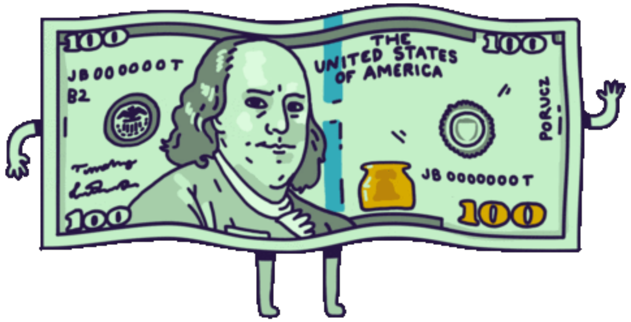

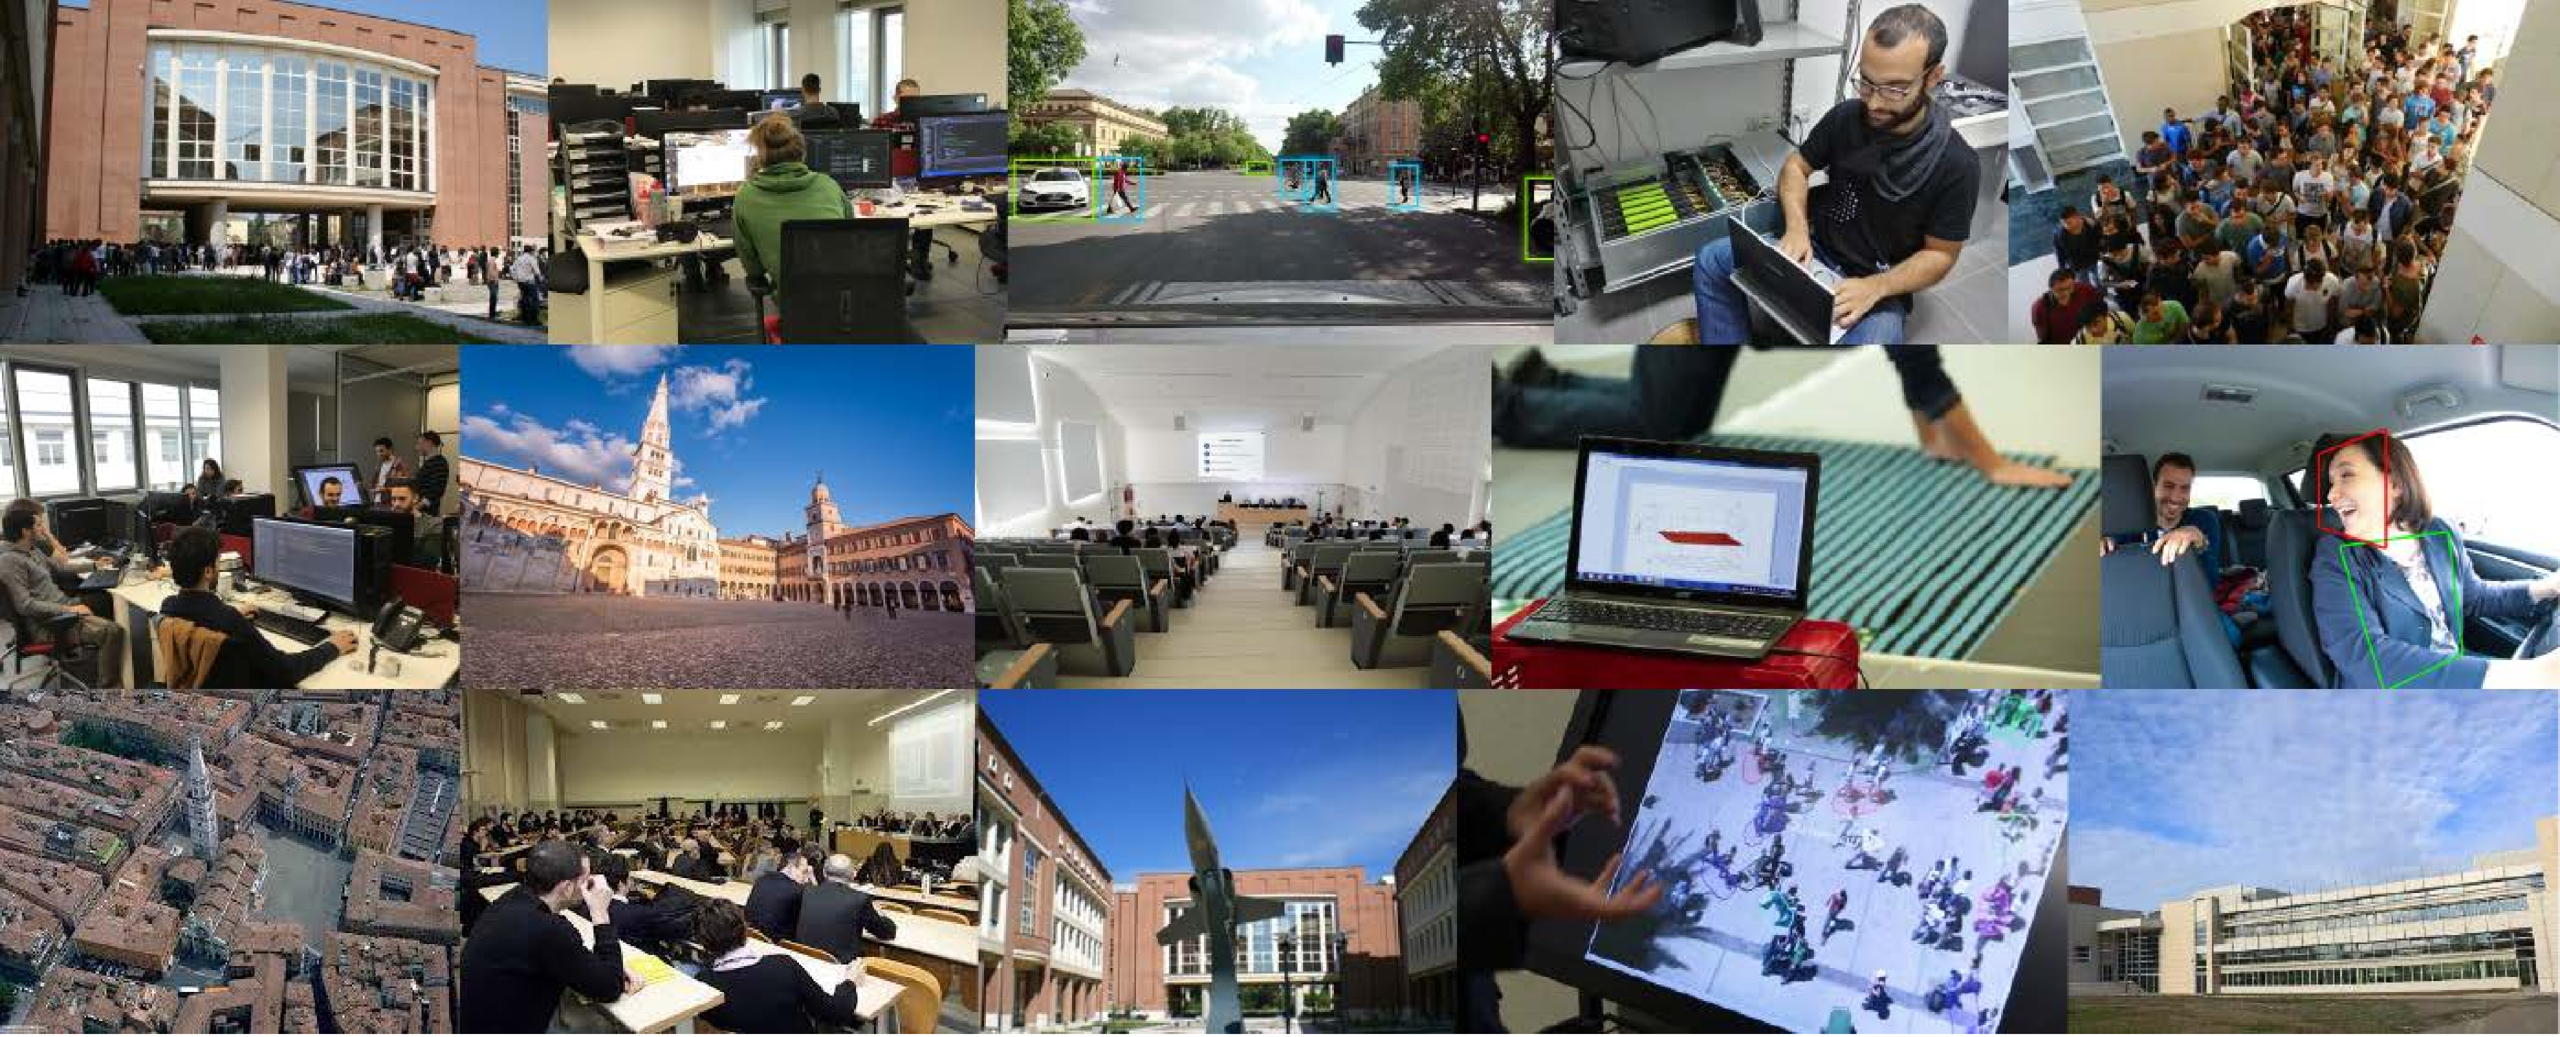

DERMOSCOPY IMAGE ANALYSIS

AImage Lab

UNIMORE UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

WARNING

The following slides contain graphic images that some readers may find disturbing

DECREASE THE AMOUNT OF MANUAL ANNOTATION REQUIRED USING GAN

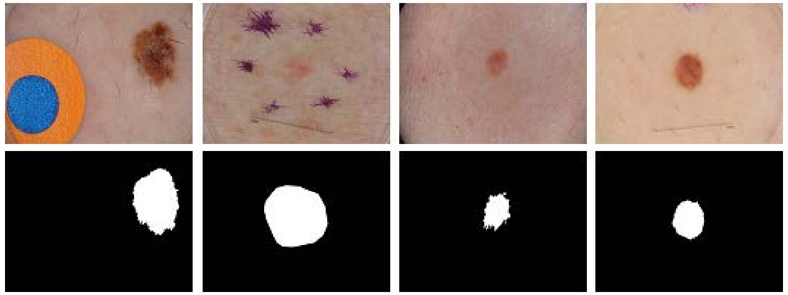

Original skin lesion images

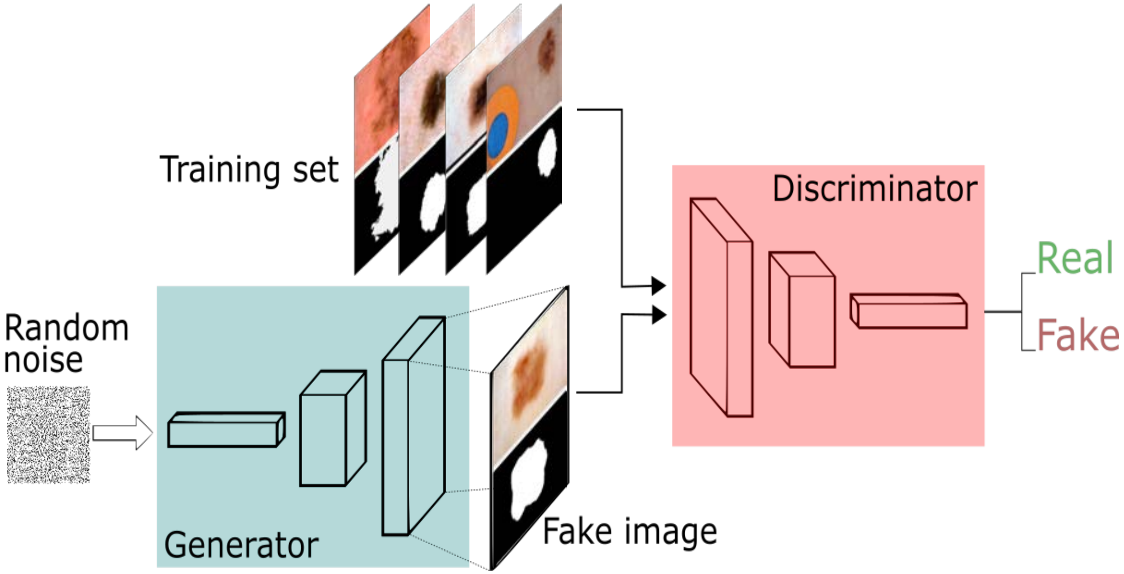

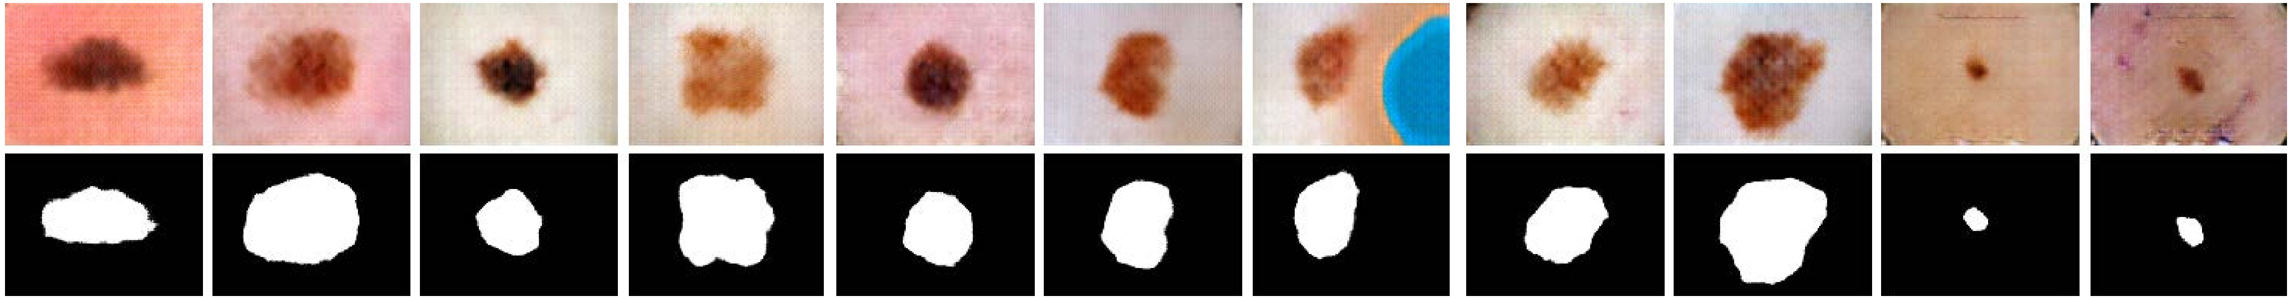

Generated skin lesion images

SAMPLES FROM DIFFERENT KINDS OF GAN

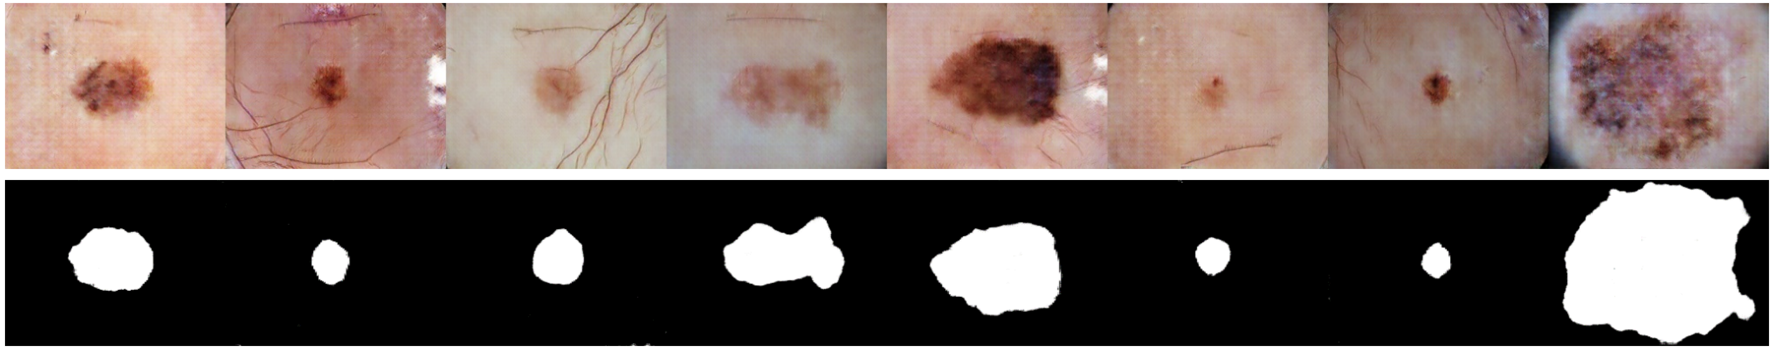

Deep Convolutional GAN [1*]

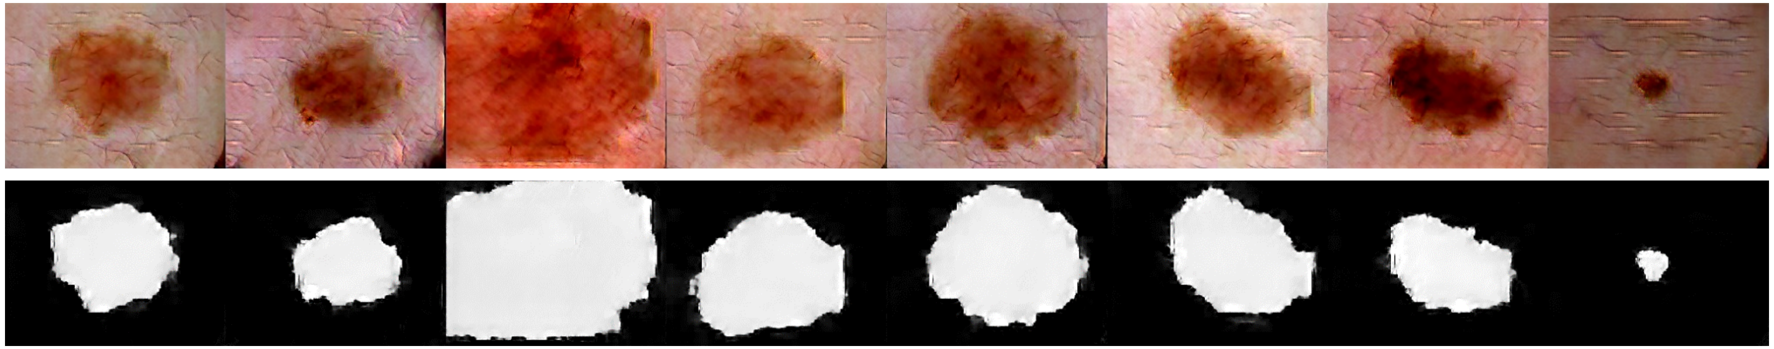

Laplacian Pyramid GAN

[1] Radford, A., Metz, L., & Chintala, S. (2015). Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434*

[2]* Denton, E. L., Chintala, S., & Fergus, R. (2015). Deep generative image models using a laplacian pyramid of adversarial networks. In *Advances in neural information processing systems* (pp. 1486-1494).

WHAT ABOUT SCALABILITY?

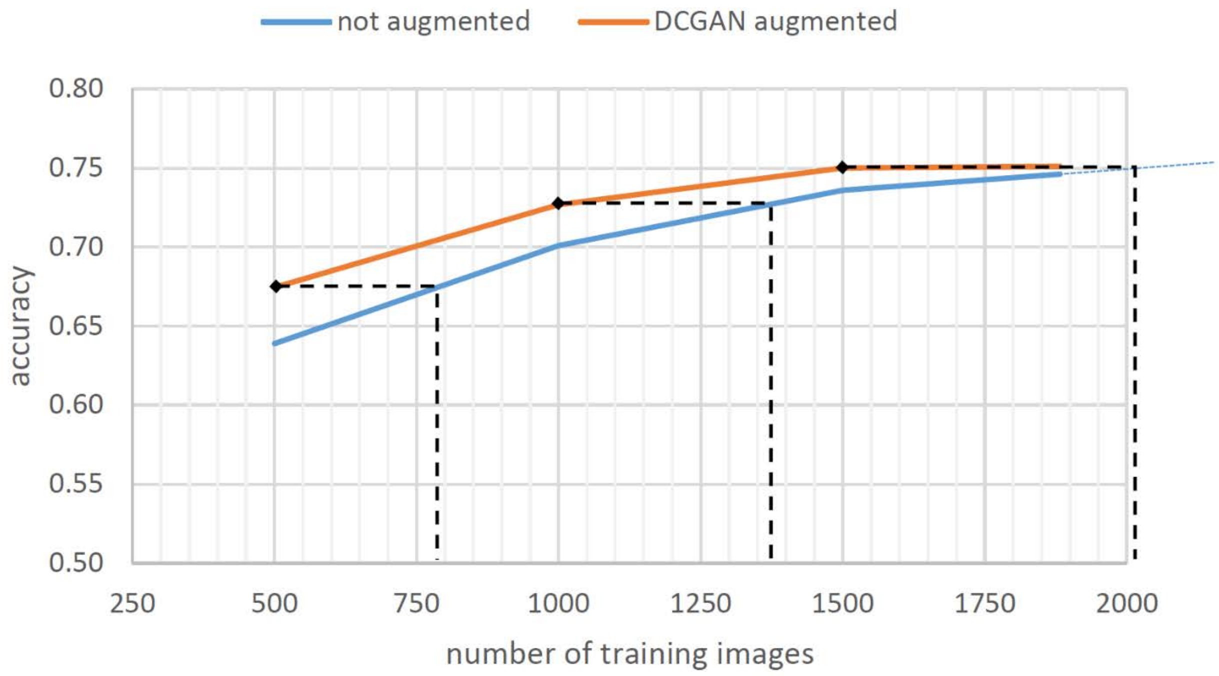

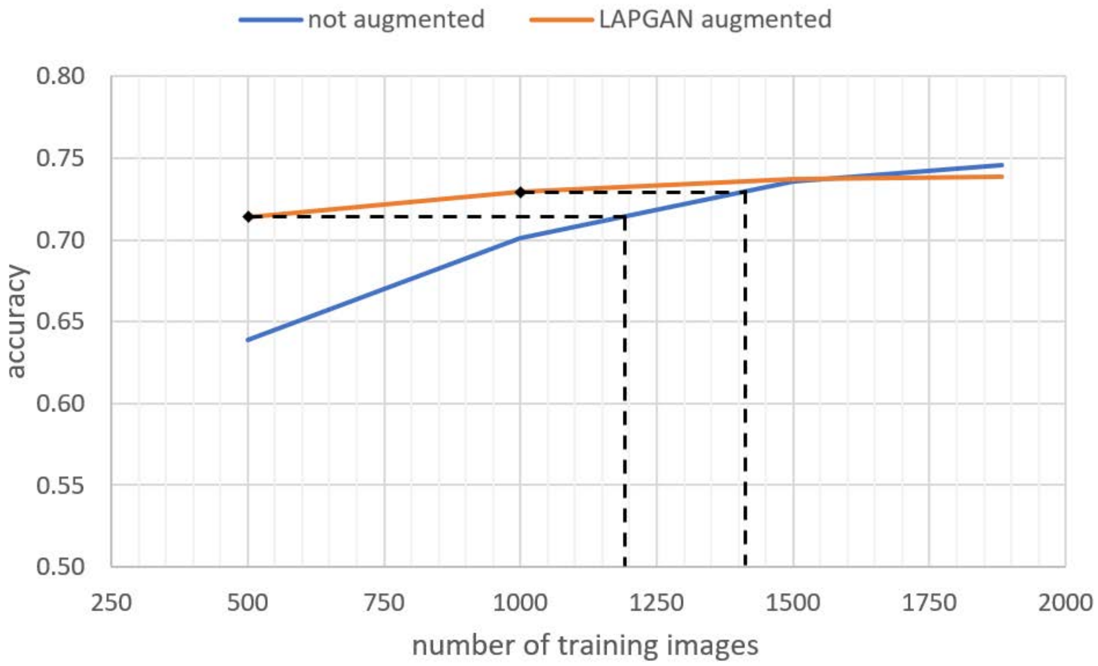

SAMPLES FROM DIFFERENT KINDS OF GAN

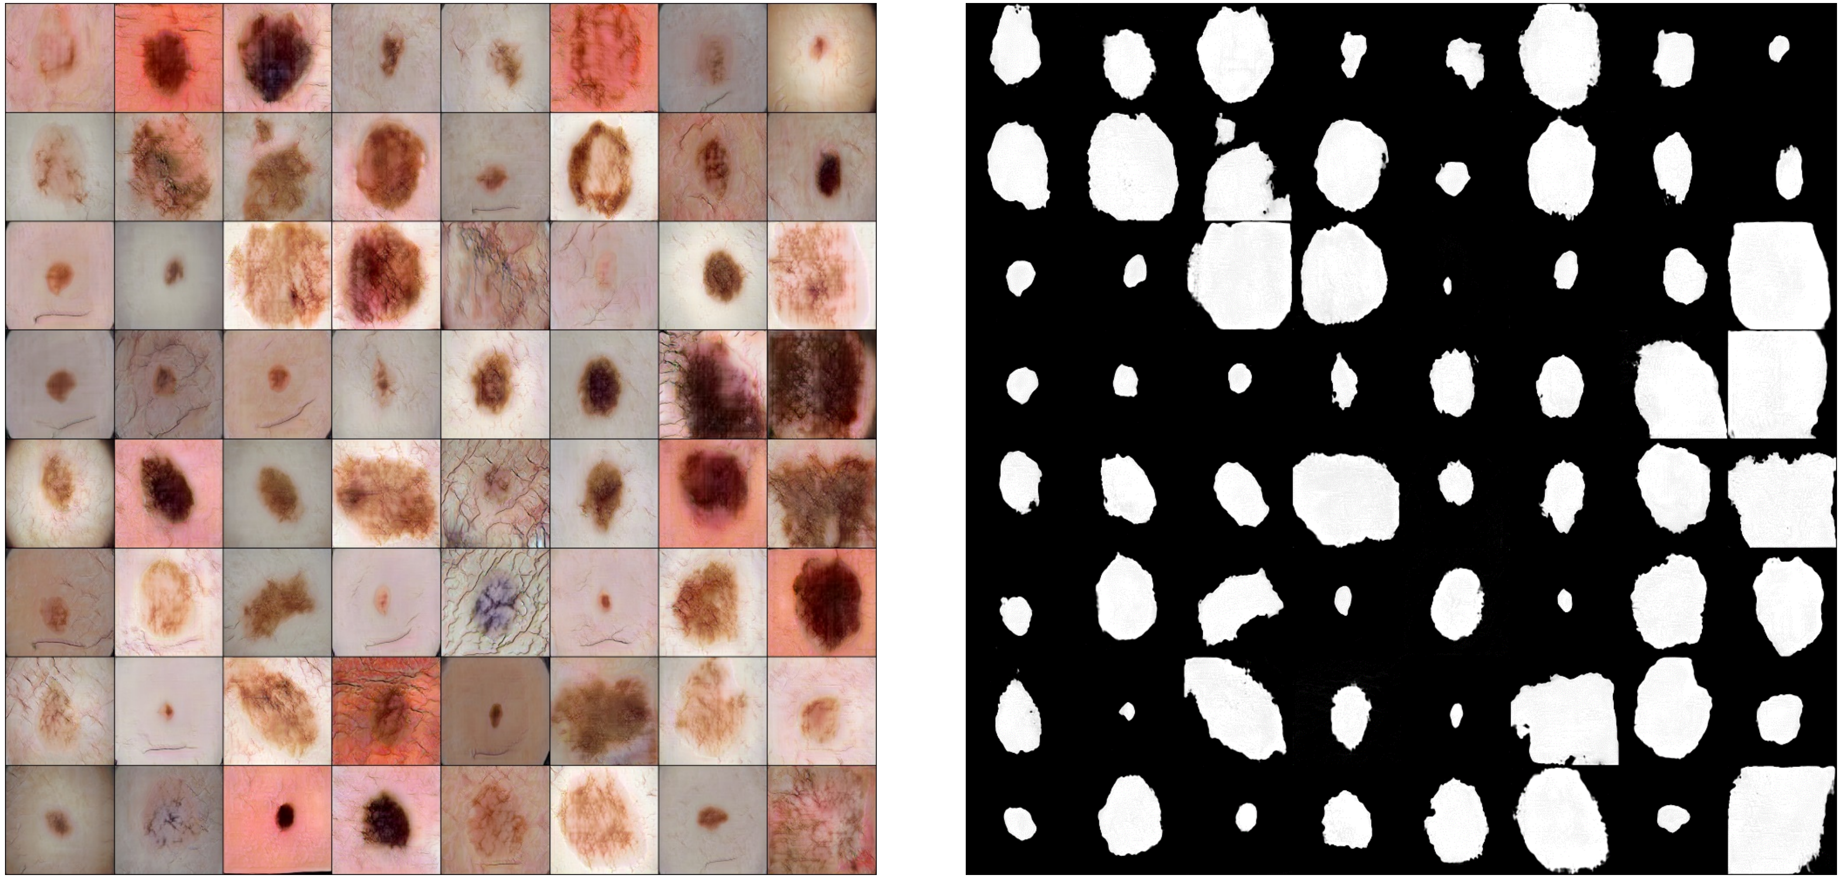

Progressive Growing of GAN

[1] Karras, T., Aila, T., Laine, S., & Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196.*

SUPPORT THE CLINICAL DECISIONS

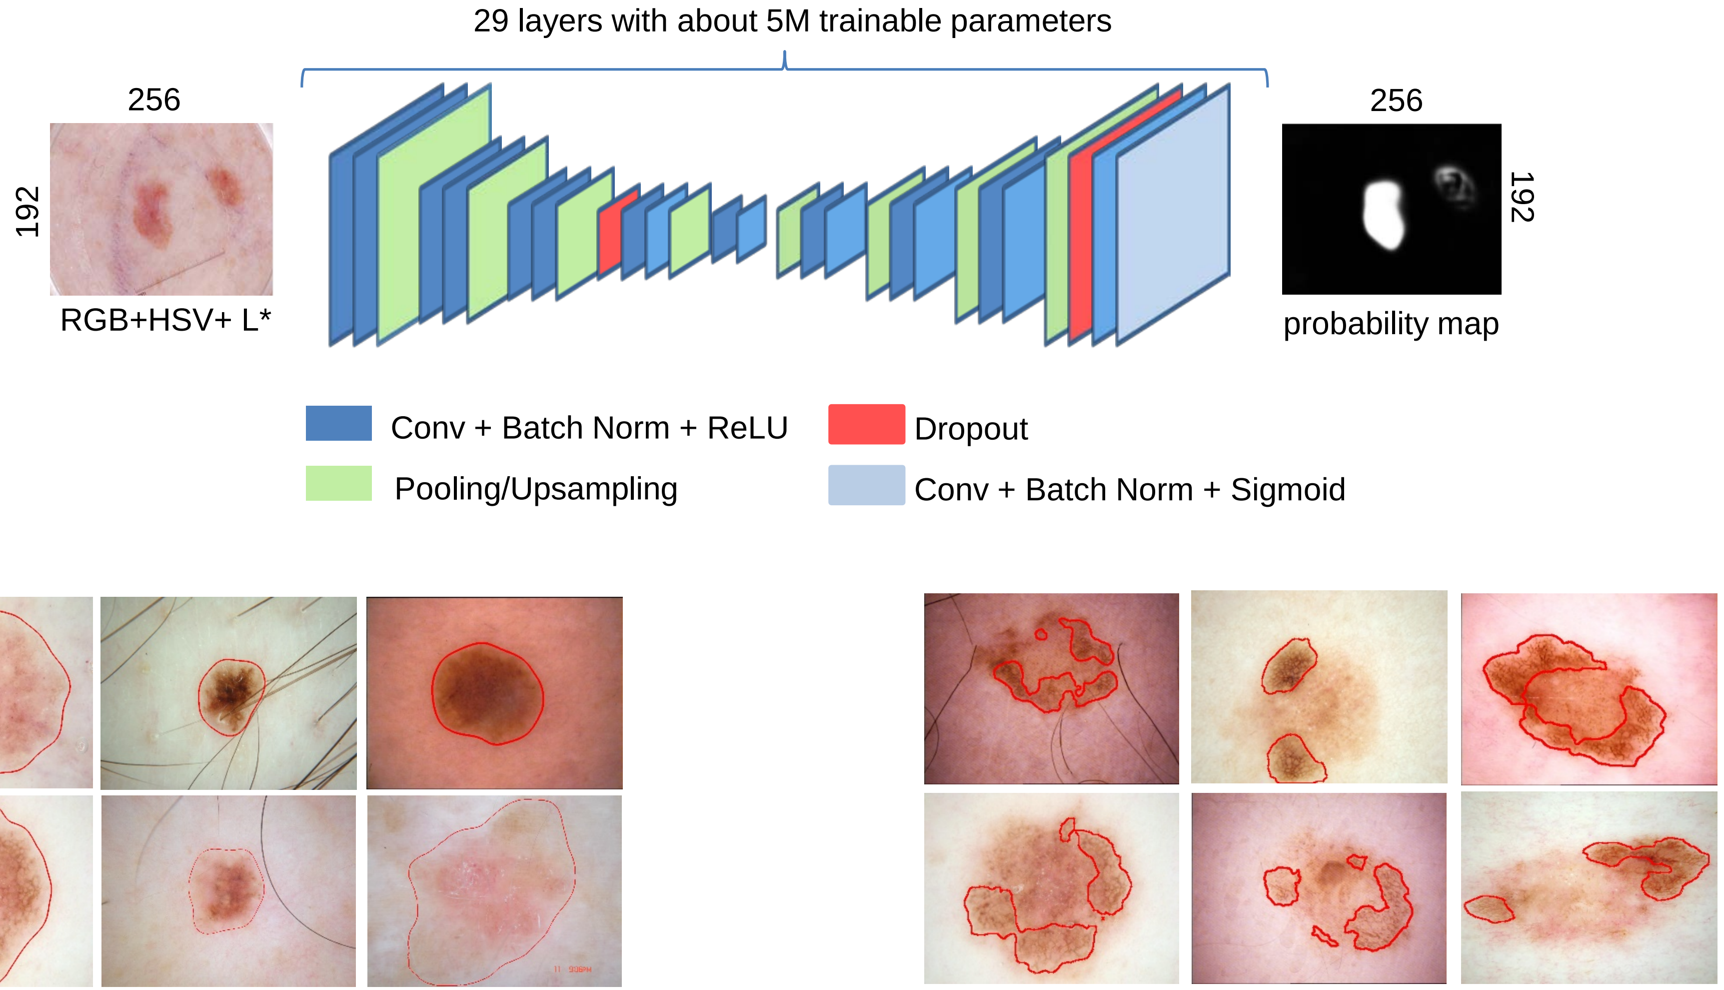

Lesion boundary segmentation

Lesion attribute detection

DETECTION-SEGMENTATION PIPELINE

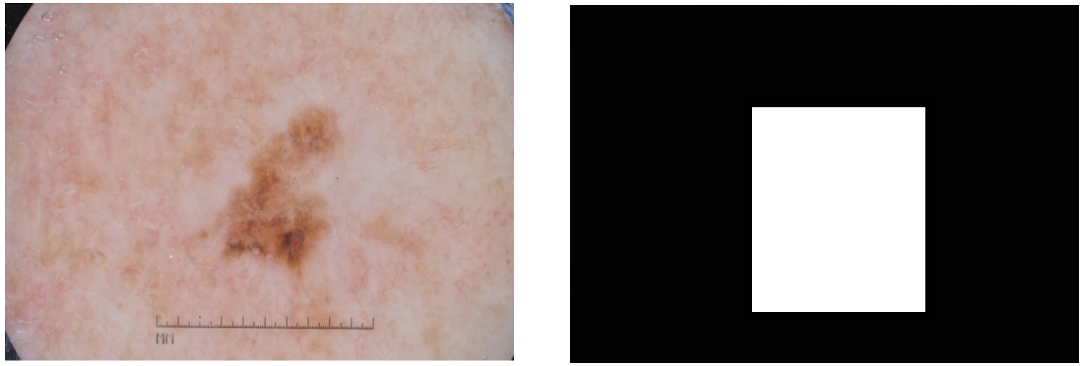

- Skin lesions are coarsely detected (MASK-RCNN)
- Then cropped
- And finally segmented

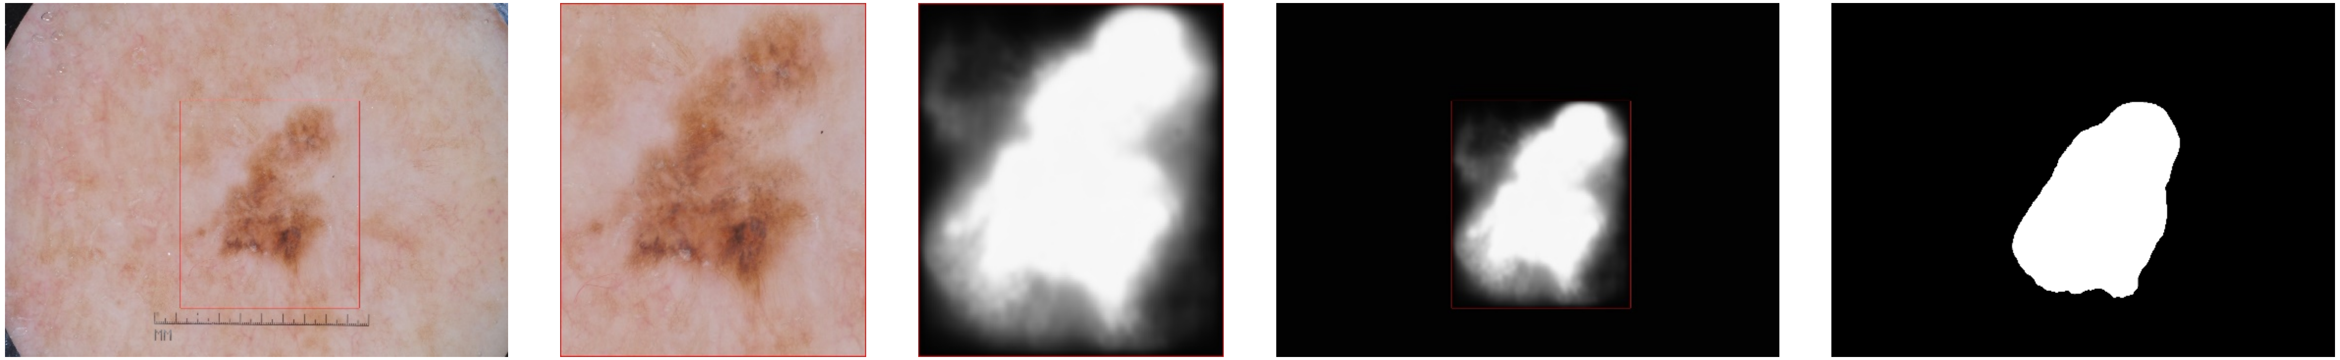

DEEP LEARNING BASED VS OLD STYLE APPROACHES

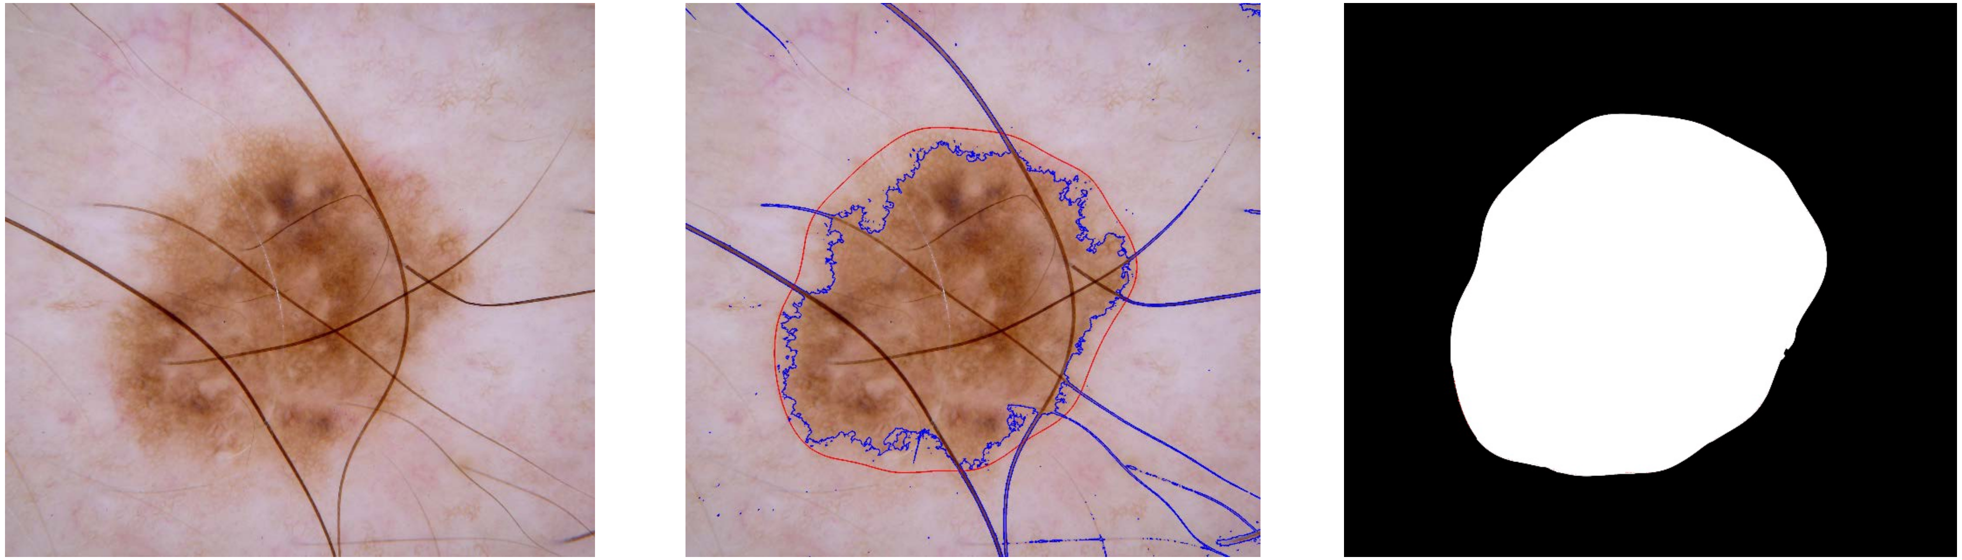

original image

old style

deep learning

ground truth

EXPERIMENTAL RESULTS

AImage Lab

Method	Test IoU	Test TIoU
Ours (ensemble)	**0.850**	0.827
SegAN [35]	0.785	—
GAN Augmented [27]	0.781	—
DCL-PSI [2]	0.777	—
DeepLabv3+* [7]	0.769	—
(RE)-DS-U-ResnetFCN34 [22]	0.772	—
SegNet* [1]	0.767	—
Challenge winners [37]	0.765	—
Tiramisu* [18]	0.765	—
U-Net* [29]	0.740	—

SKIN LESION CLASSIFICATION – 2019 ISIC CHALLENGE

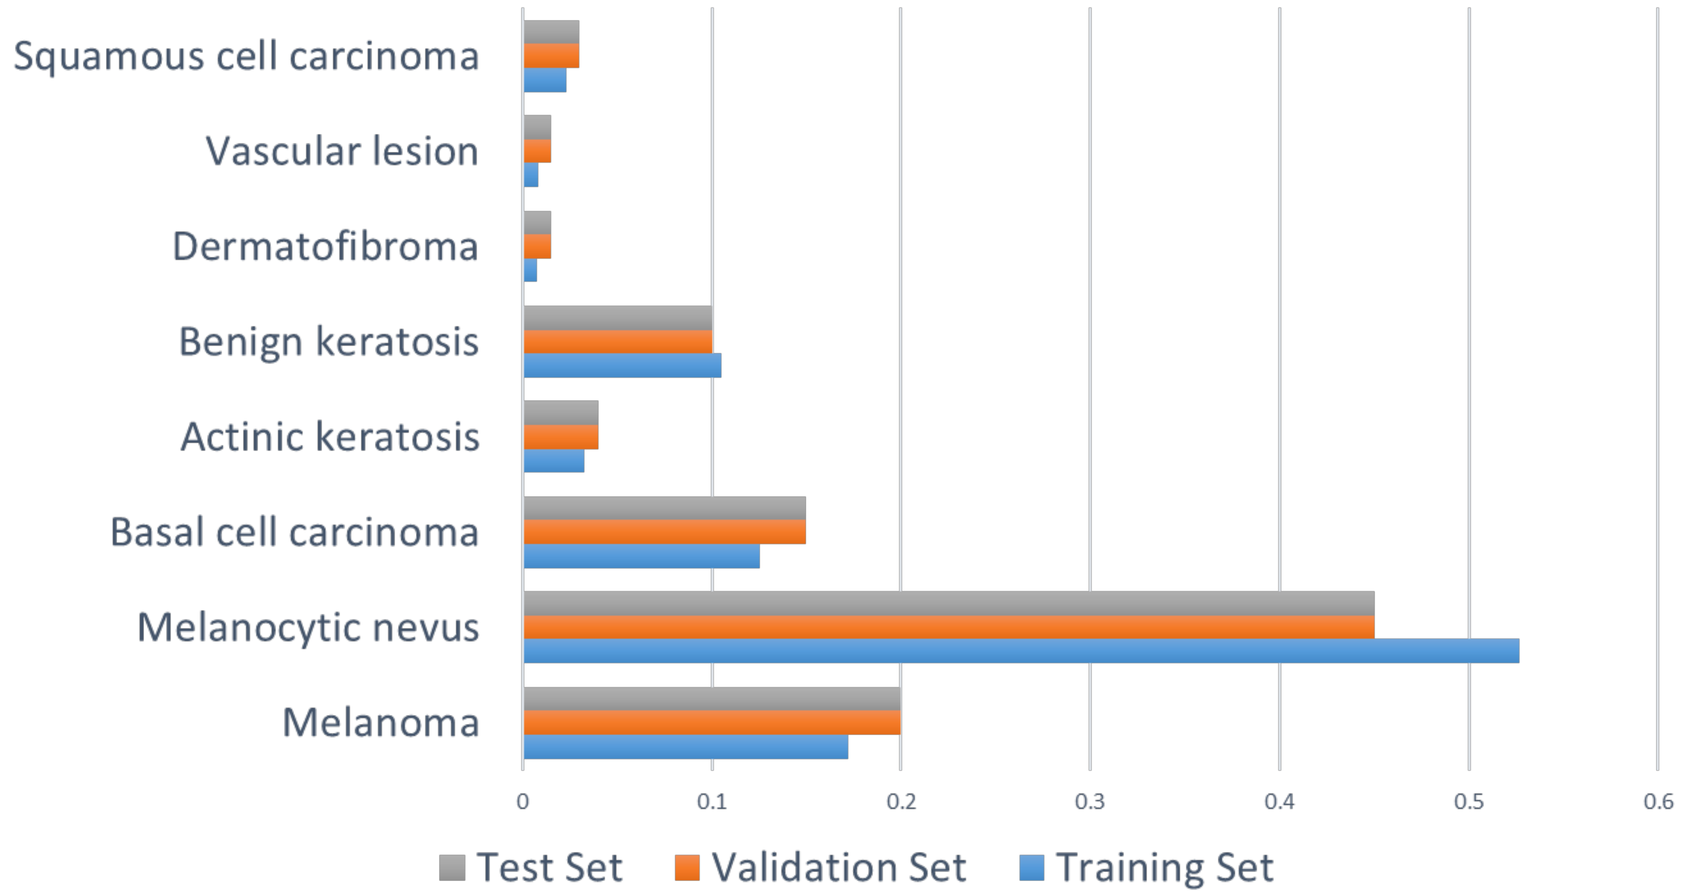

SKIN LESION CLASSIFICATION – 2019 ISIC CHALLENGE

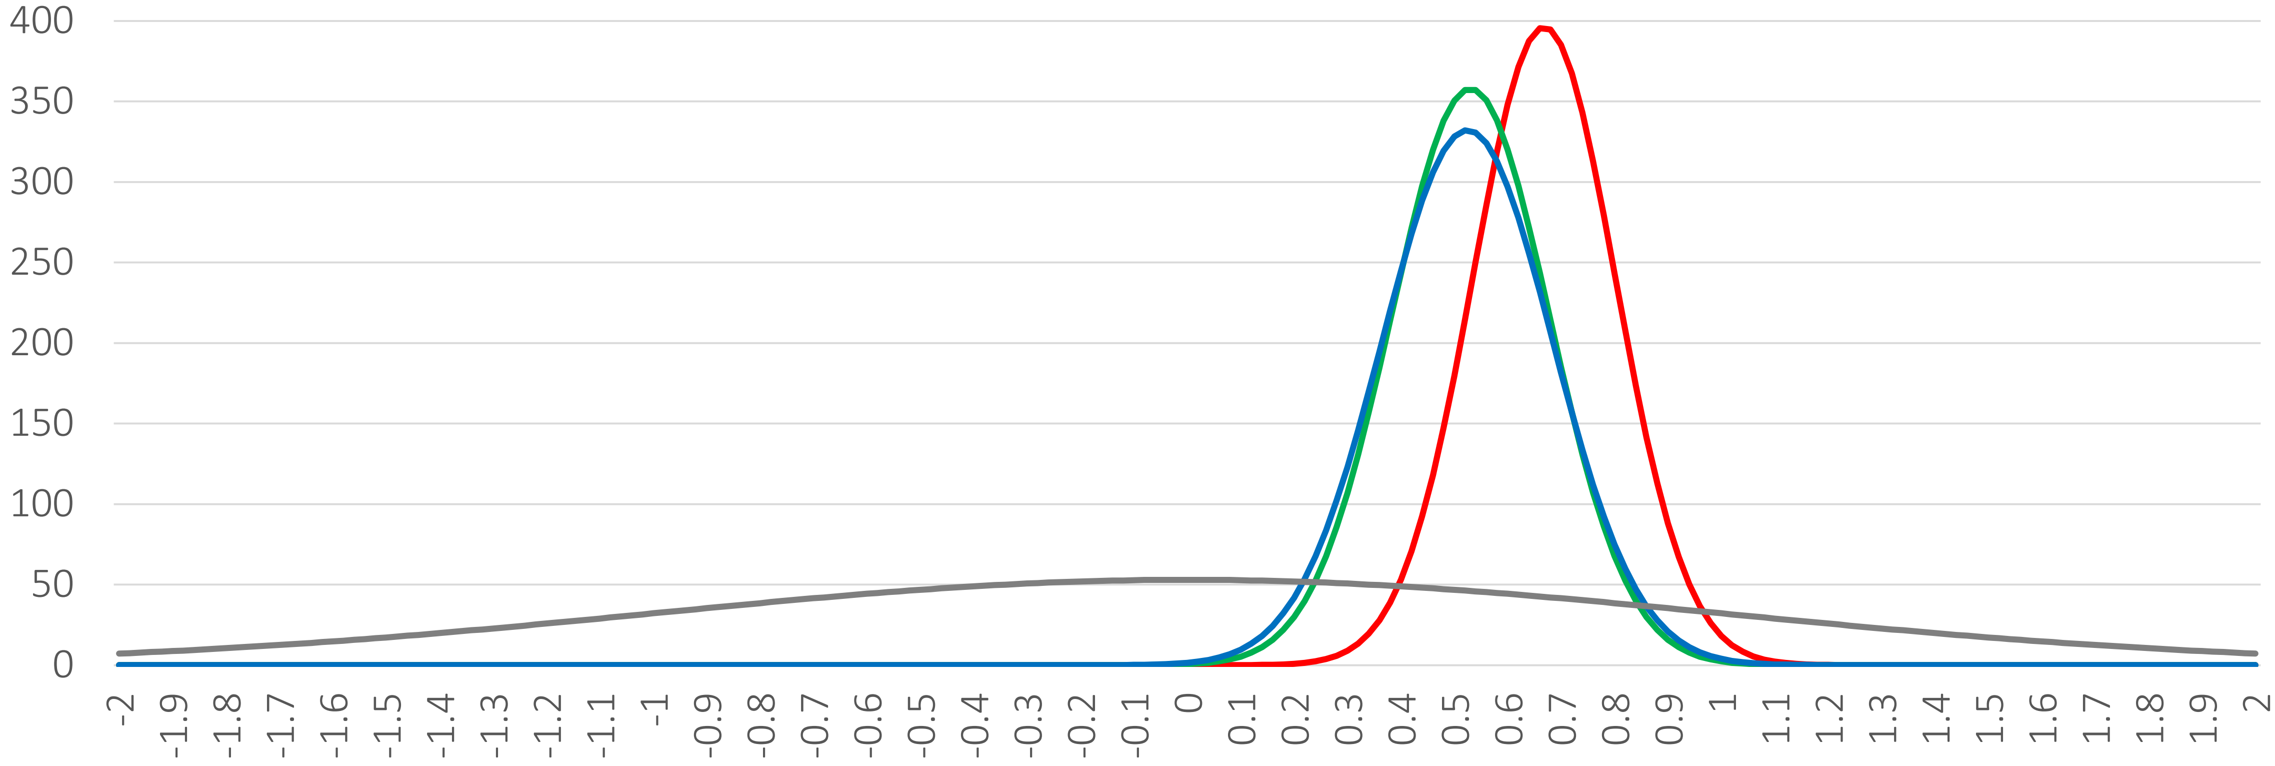

SKIN LESION CLASSIFICATION – 2019 ISIC CHALLENGE

Calibration

Neural nets are overconfident, outputs can be rescaled to obtain a calibrated prediction

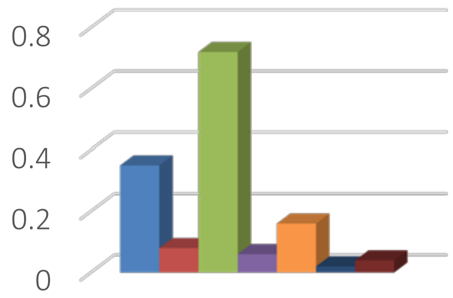

Data Augmentation

During inference, different versions of every image are fed to the model

Averaging Models

- 10 neural networks
- 3 checkpoints for network

The final output is obtained by averaging the prediction of every network

🎉 Third-place (out of 64 research groups) at the 2019 international ISIC challenge 🎉

EXPLAINABILITY AND INTERPRETABILITY

Where do Neural networks look?

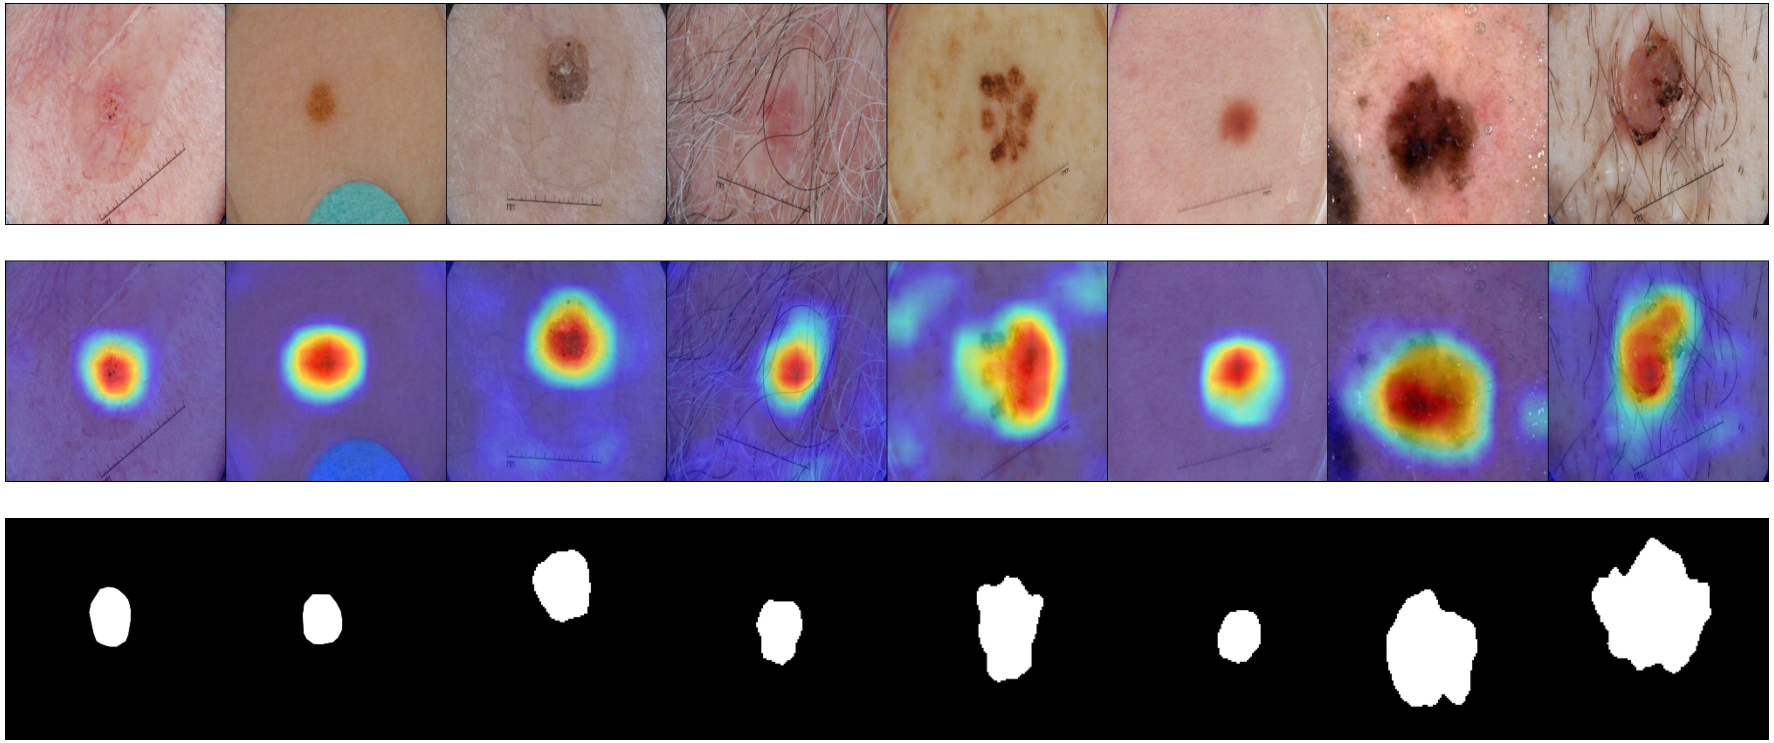

[1*] Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision (pp. 618-626).

UNIMORE

UNIVERSITÀ DEGLI STUDI DI MODENA E REGGIO EMILIA

AImage Lab

THANK YOU!

Federico Bolelli

federico.bolelli@unimore.it

Università degli Studi di Modena e Reggio Emilia, DIEF, *Italy*