SFU SIMON FRASER UNIVERSITY
ENGAGING THE WORLD

Deep learning for skin image analysis

Beyond more data and faster GPUs

Ghassan Hamarneh

hamarneh@sfu.ca

www. **Medical Image Analysis** .com

17 June 2019

ISIC Skin Image Analysis Workshop at CVPR

Acknowledgements

Current and former students

Dr. Jeremy Kawahara

Dr. Hengameh Mirzaalian

Dr. Aicha Bentaieb

Saeid Asgari

Zahra MiriKharaji

Saeed Izadi

Yiqi Yan

Kumar Abhishek

Chris Yoon

Funding: research, infrastructure, computing

Disclosure

Scientific Advisor and Shareholder:
Triage Technologies Inc.

Skin image analysis system

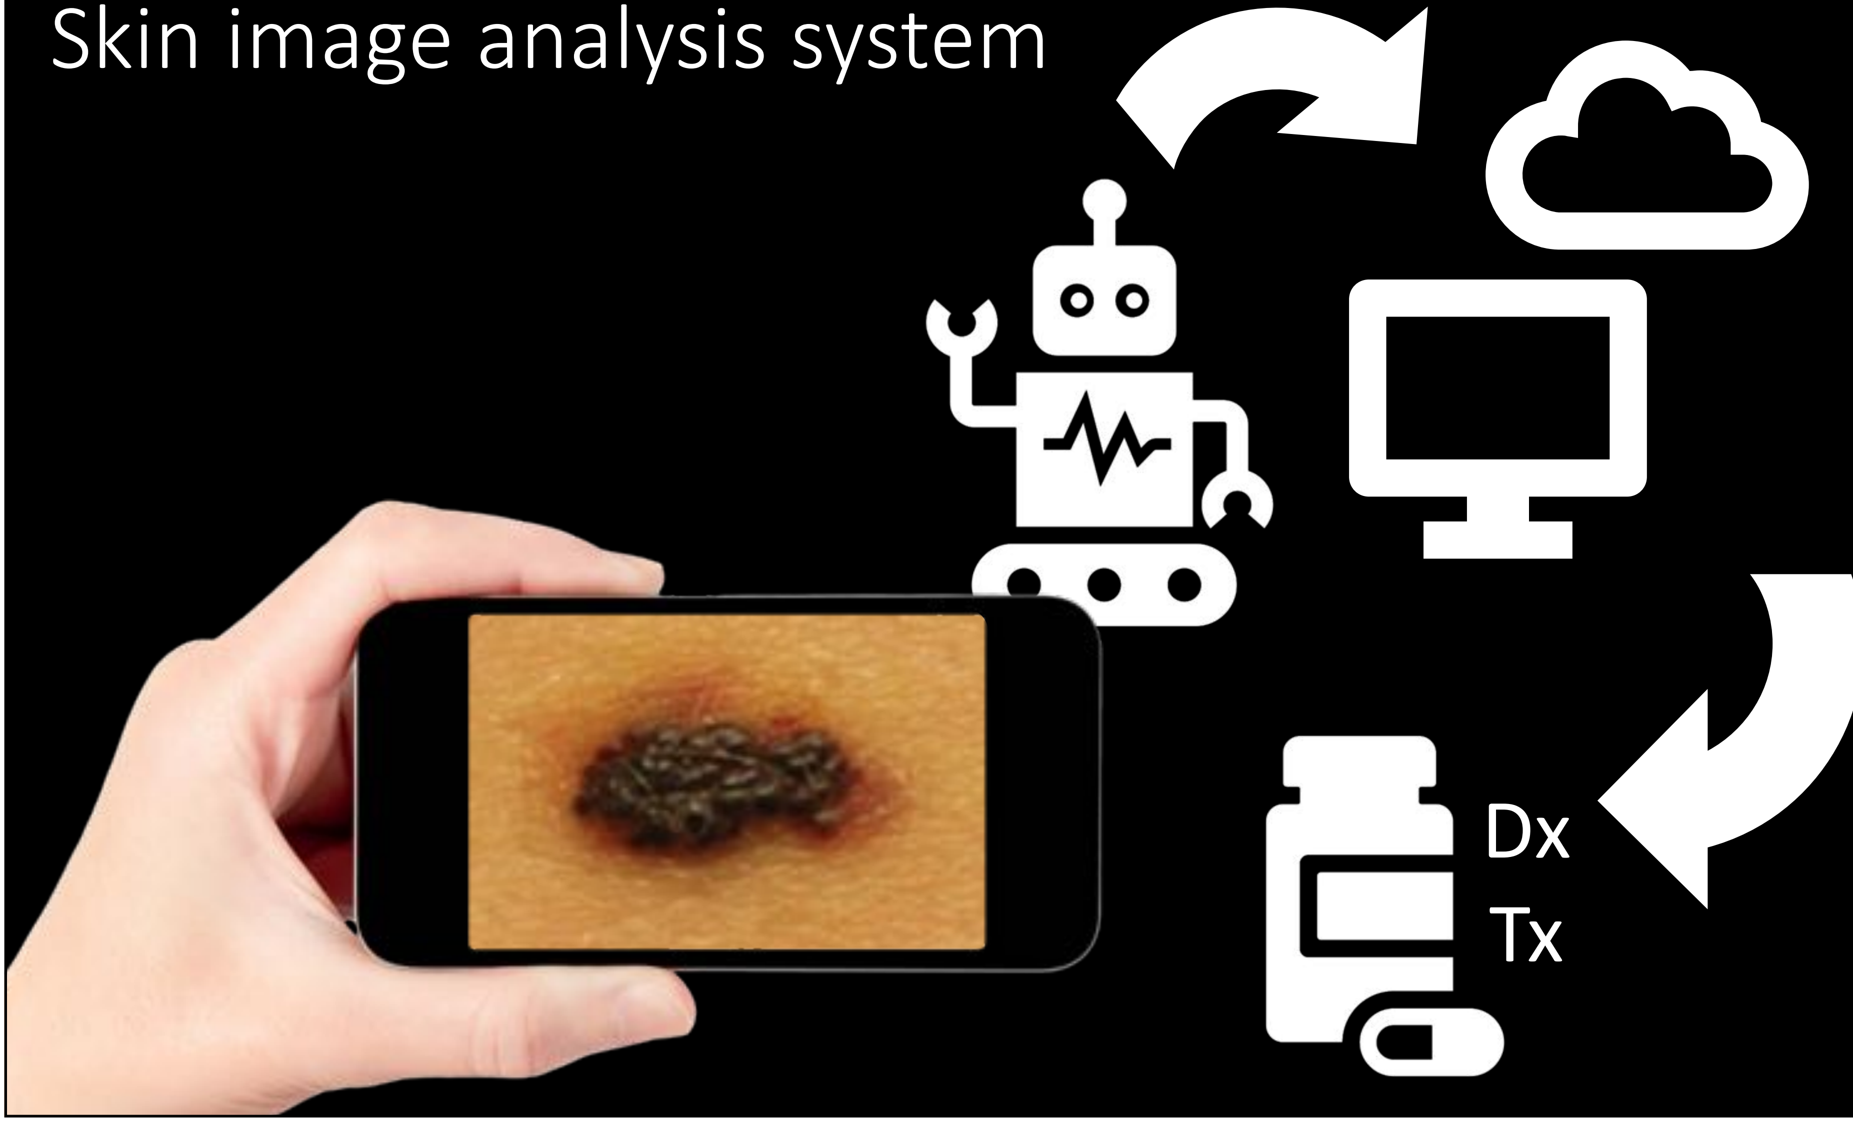

Image analysis via deep learning

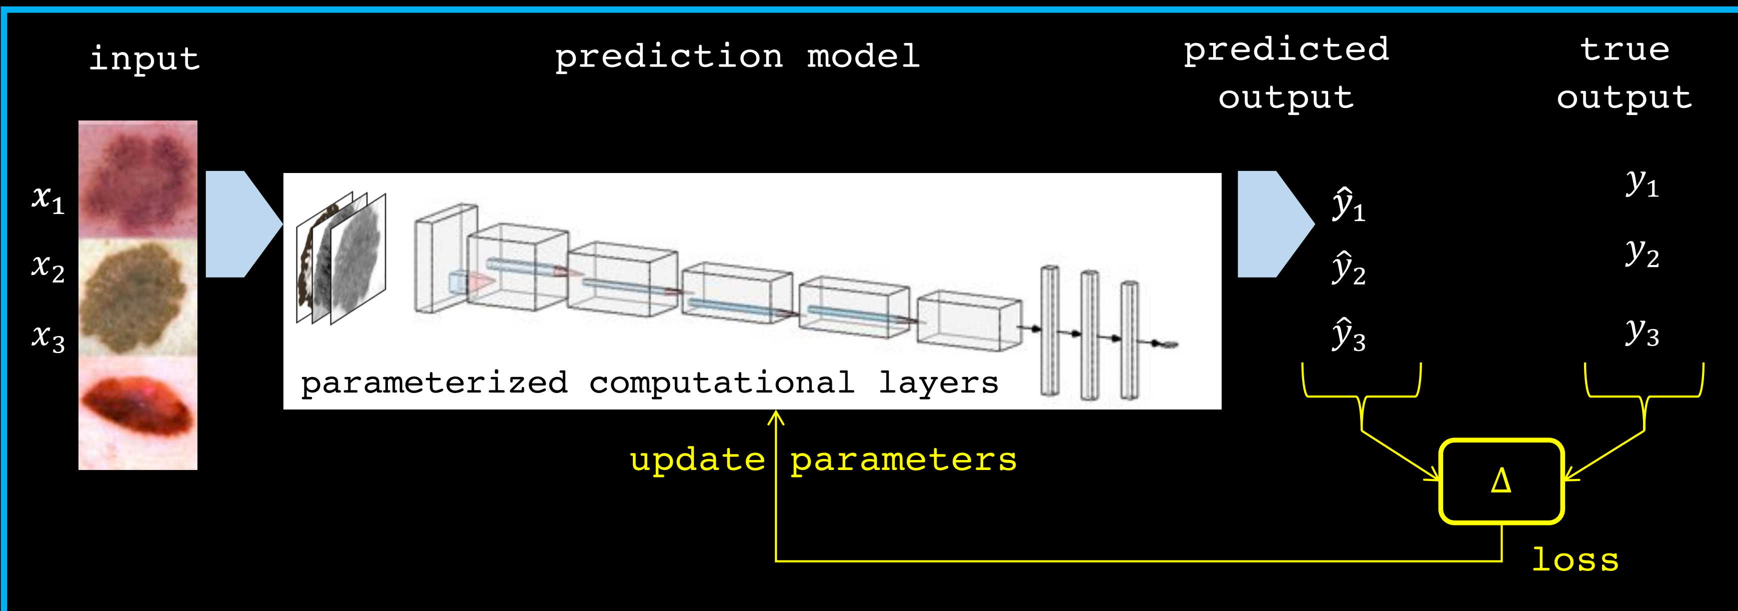

hyperparameter tuning

training, validation, testing (hold out), deployment

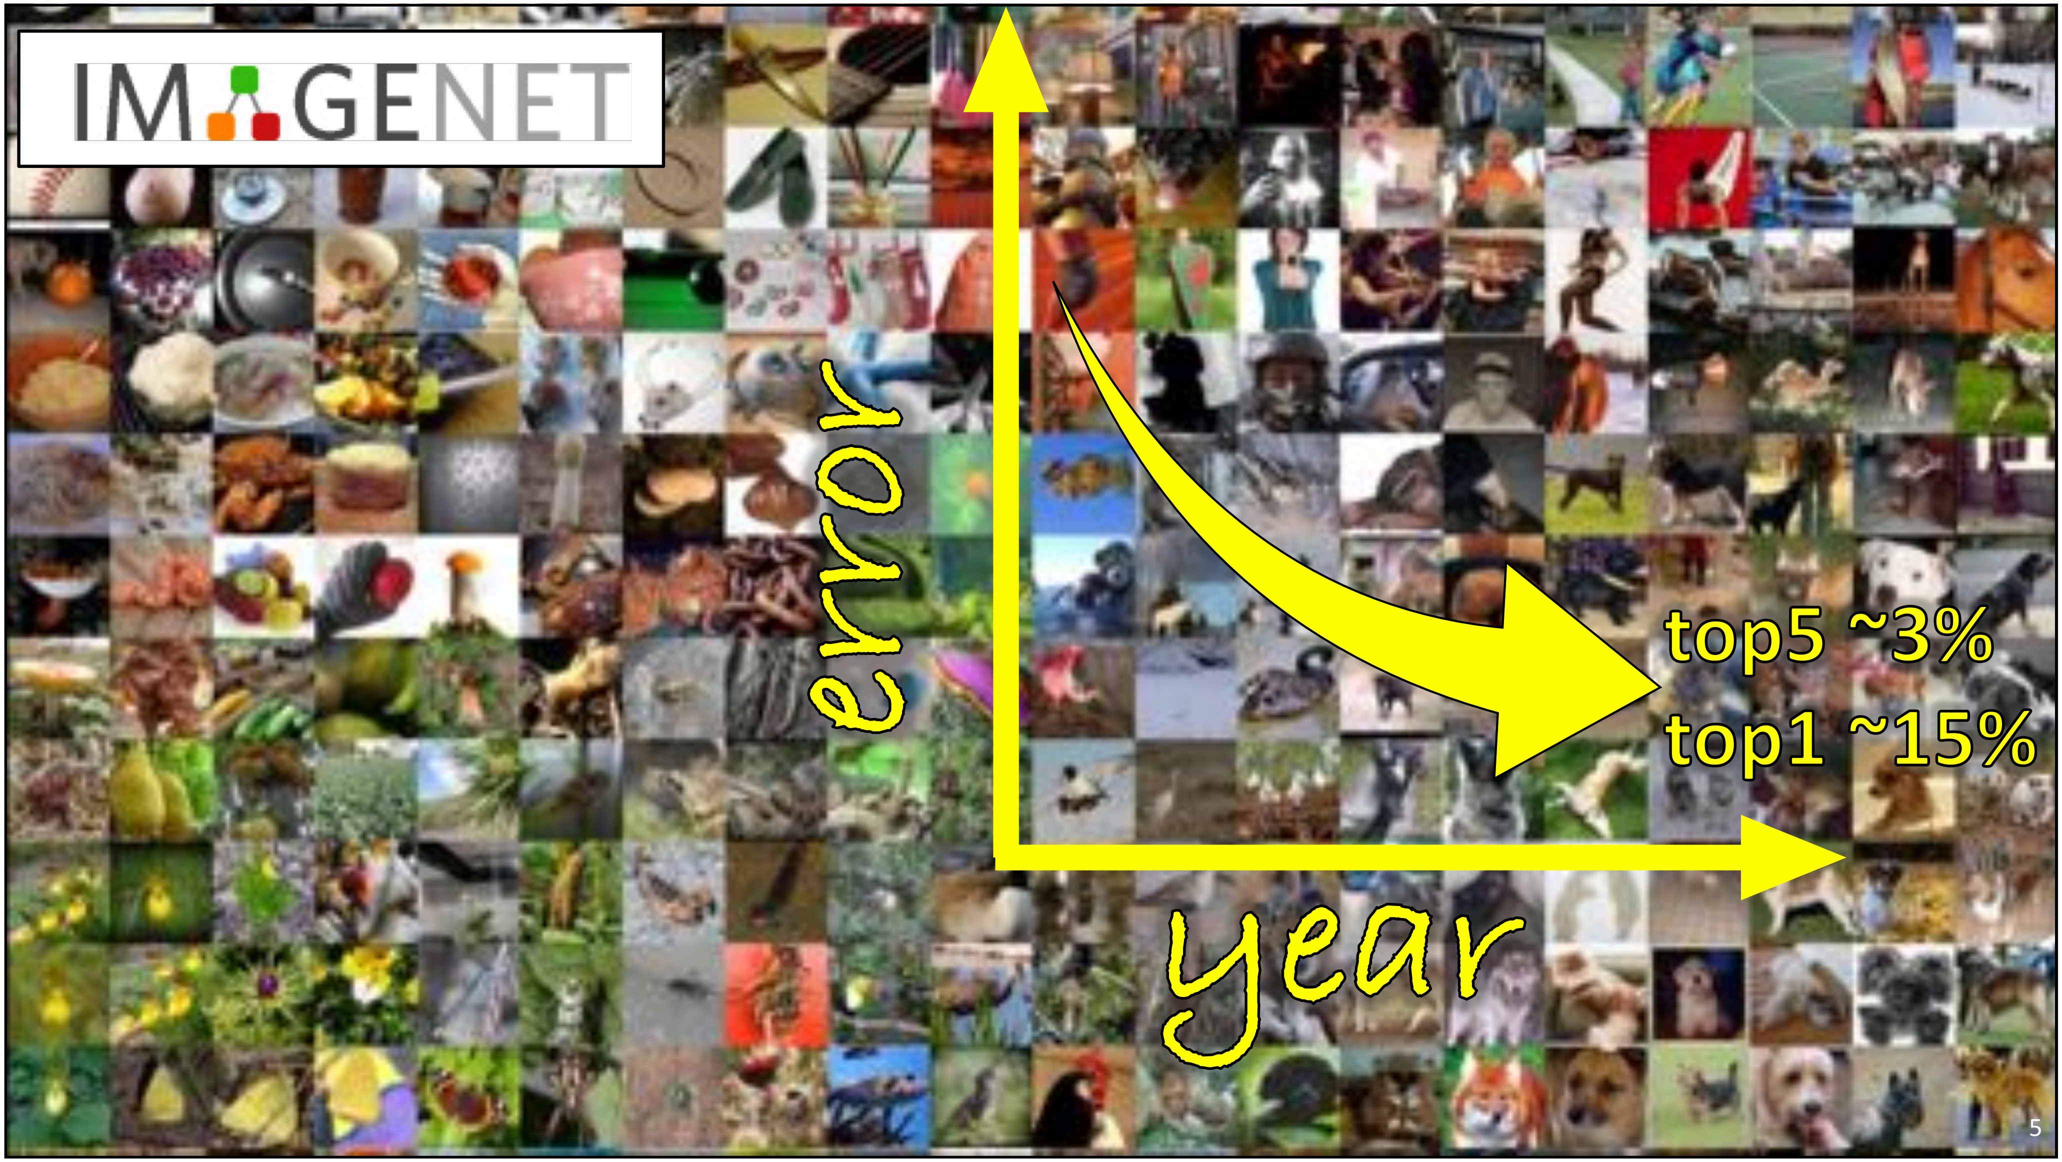
IMAGENET
error
year
top5 ~3%
top1 ~15%

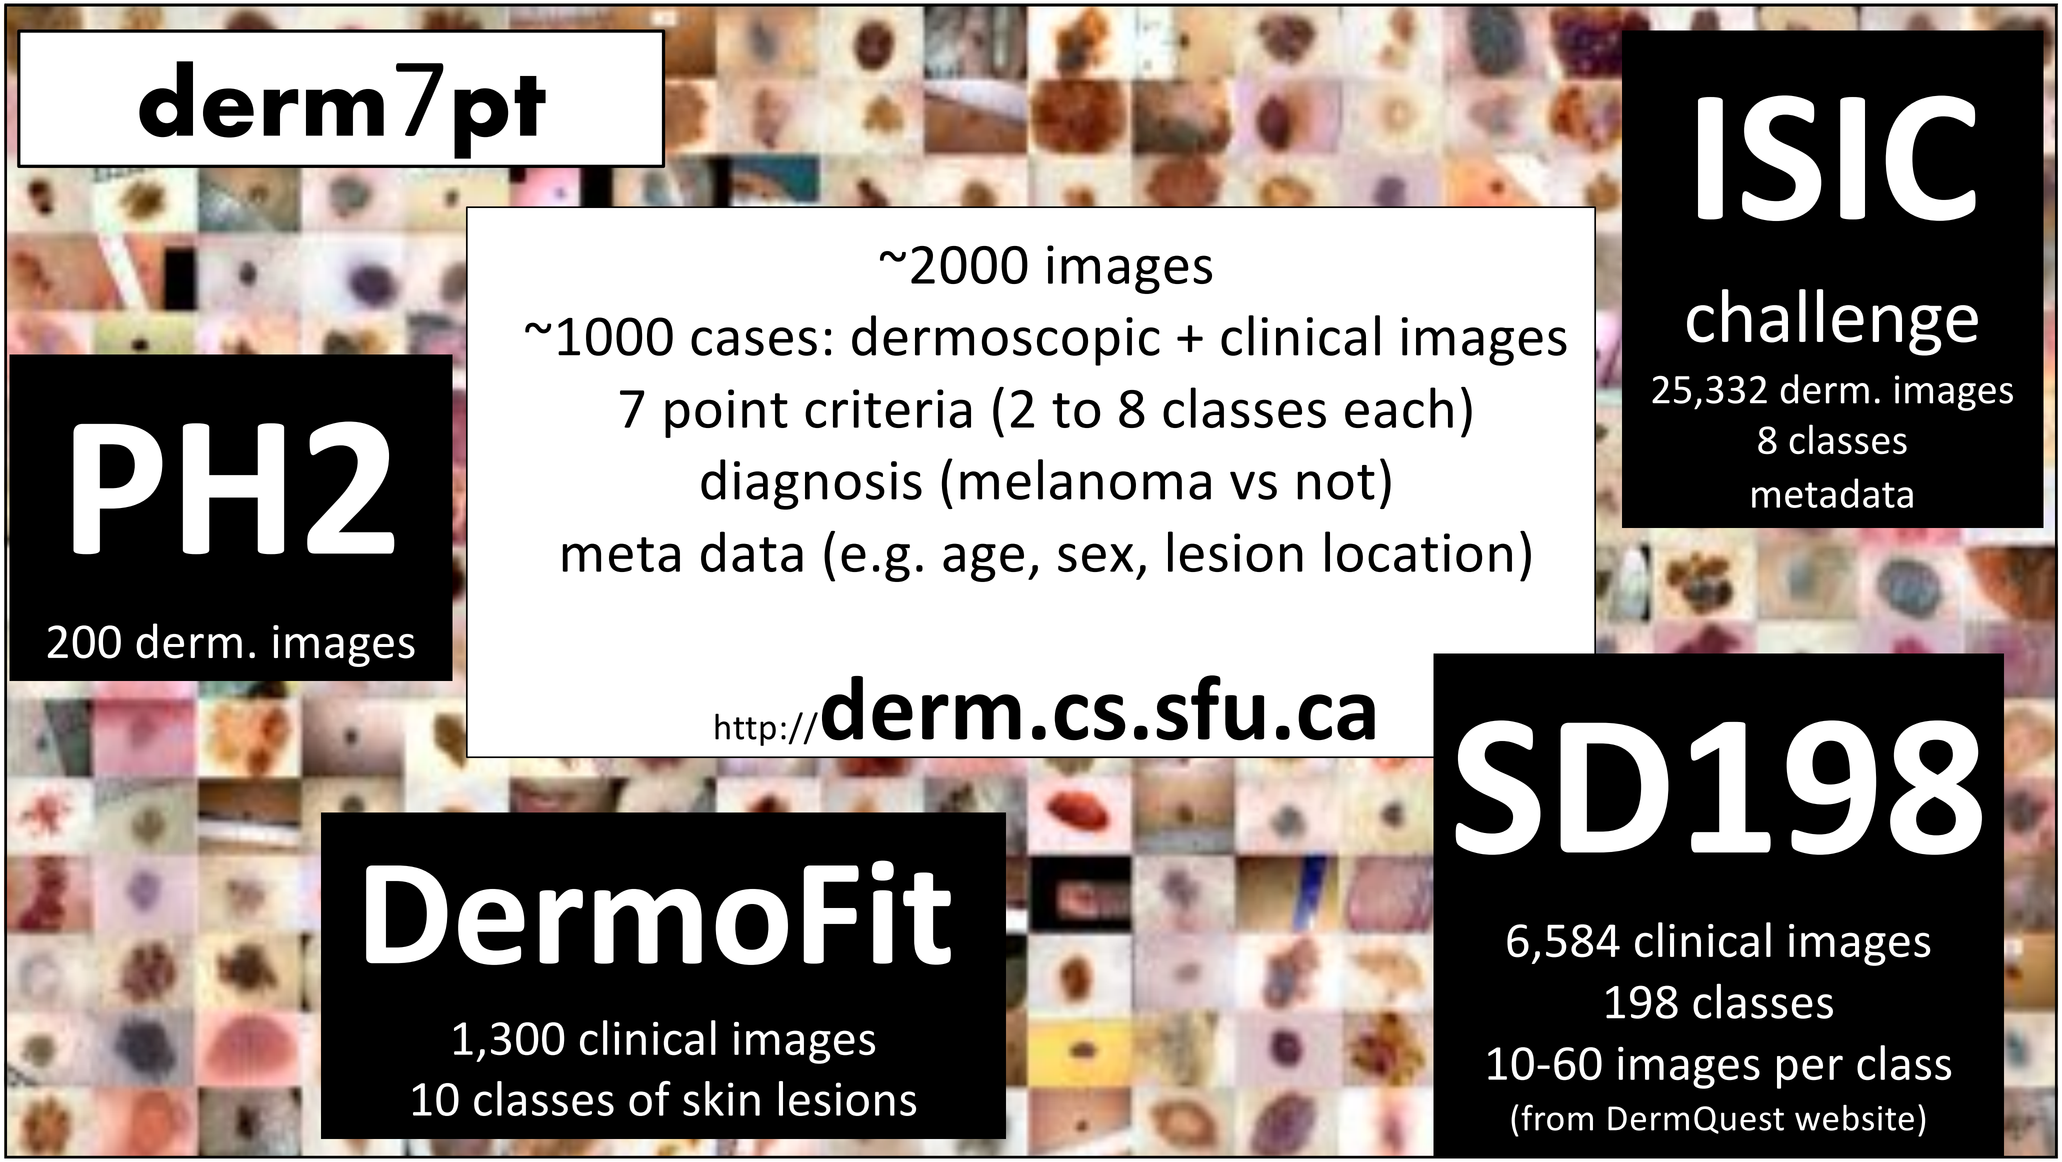
derm7pt
~2000 images
~1000 cases: dermoscopic + clinical images
7 point criteria (2 to 8 classes each)
diagnosis (melanoma vs not)
meta data (e.g. age, sex, lesion location)
http://derm.cs.sfu.ca
ISIC
challenge
25,332 derm. images
8 classes
metadata
PH2
200 derm. images
DermoFit
1,300 clinical images
10 classes of skin lesions
SD198
6,584 clinical images
198 classes
10-60 images per class
(from DermQuest website)

nature

Letter | Published: 25 January 2017

Jan. 2017

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau & Sebastian Thrun

Nature **542**, 115–118 (02 February 2017) | Download Citation

https://www.nature.com/articles/nature21056

GENERAL DERMATOLOGY

Feb. 2019

BJD British Journal of Dermatology

Deep-learning-based, computer-aided classifier developed with a small dataset of clinical images surpasses board-certified dermatologists in skin tumour diagnosis*

Y. Fujisawa,[1] Y. Otomo,[2] Y. Ogata,[3] Y. Nakamura,[1] R. Fujita,[2] Y. Ishitsuka,[1] R. Watanabe,[1] N. Okiyama,[1] K. Ohara[4] and M. Fujimoto[1]

https://onlinelibrary.wiley.com/doi/abs/10.1111/bjd.17470

ANNALS OF ONCOLOGY

"CNN ROC AUC greater than mean ROC area of dermatologists… higher specificity"

Man against machine: diagnostic performance of a deep learning convolutional neural network for dermoscopic melanoma recognition in comparison to 58 dermatologists

H. A. Haenssle[1*,†], C. Fink[1†], R. Schneiderbauer[1], F. Toberer[1], T. Buhl[2], A. Blum[3], A. Kalloo[4], A. Ben Hadj Hassen[5], L. Thomas[6], A. Enk[1] & L. Uhlmann[7]

August 2018

https://academic.oup.com/annonc/article/29/8/1836/5004443

EJC EUROPEAN JOURNAL OF CANCER

May 2019

Deep learning outperformed 136 of 157 dermatologists in a head-to-head dermoscopic melanoma image classification task

Titus J. Brinker[a,b,*], Achim Hekler[a], Alexander H. Enk[b], Joachim Klode[c], Axel Hauschild[d], Carola Berking[e], Bastian Schilling[f], Sebastian Haferkamp[g], Dirk Schadendorf[c], Tim Holland-Letz[h], Jochen S. Utikal[i,j,1], Christof von Kalle[a,1], Collaborators[2]

https://www.sciencedirect.com/science/article/pii/S0959804919302217

human vs machine

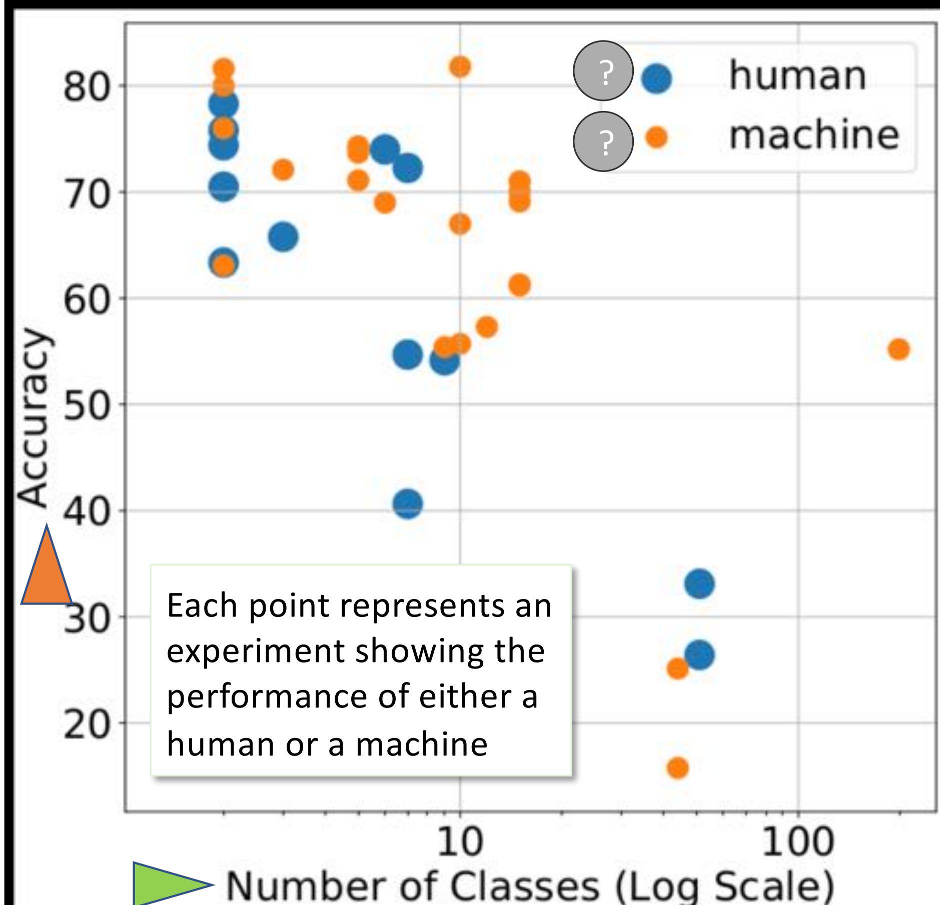

arXiv.org > cs > arXiv:1906.01256

Computer Science > Computer Vision and Pattern Recognition

Visual Diagnosis of Dermatological Disorders: Human and Machine Performance

June 2019

https://arxiv.org/abs/1906.01256

Jeremy Kawahara, Ghassan Hamarneh

	Year	Dataset	N.Images	N.Test	Derm.	Clinic.	Meta	H.vs.M	Classes	Acc.
[89]	2015	Internal	-	65	✓			human	2	63.35
[89]	2015	Internal	273	65	✓			machine	2	63.08
[90]	2017	ISIC-100	-	100	✓			human	2	70.50
[90]	2017	ISIC-100	1000	100	✓			machine	2	76.00
[96]	2018	Internal	-	100	✓			human	2	74.40
[96]	2018	Internal	-	100	✓		✓	human	2	78.30
[96]	2018	Internal		100	✓			machine	2	81.60
[94]	2018	Asan	-	1133		✓		human	2	75.80
[94]	2018	Asan	49,567	1133		✓		machine	2	80.00
[97]	2019	ISIC-100	13737	100	✓			machine	2	84.02
[39]	2019	ISIC-100	-	100	✓			human	2	62.82
[39]	2019	MED-NODE	-	100		✓		human	2	69.40
[41]	2017	Stanford	-	180	✓	✓		human	3	65.78
[41]	2017	Stanford	127,463	127,463	✓	✓		machine	3	72.10
[22]	2013	Dermofit	960	960		✓		machine	5	74.30
[83]	2018	Atlas	2018	395	✓	✓		machine	5	71.10
[83]	2018	Atlas	2018	395	✓	✓	✓	machine	5	73.70
[93]	2017	Internal	348	50	✓			human	6	74.00
[93]	2017	Internal	348	50	✓			machine	6	69.00
[102]	2002	Internal	-	256		✓*	✓*	human	7	40.62
[102]	2002	Internal	-	256	✓		✓	human	7	54.69
[102]	2002	Internal	-	256	✓*	✓*	✓*	human	7	72.27
[41]	2017	Stanford	-	180	✓	✓		human	9	54.15
[41]	2017	Stanford	127,463	127,463	✓	✓		machine	9	55.40
[57]	2015	Dermofit	1300	1300		✓		machine	10	67.00
[66]	2016	Dermofit	1300	1300		✓		machine	10	81.80
[42]	2018	Dermofit	20,689	1300		✓		machine	10	55.70
[42]	2018	Asan	19,389	1,276		✓		machine	12	57.30
[98]	2019	Internal	-	1260		✓		human	14	41.70
[98]	2019	Internal	-	1820		✓		human	14	59.70
[98]	2019	Internal	6009	1142		✓		machine	14	76.50
[45]	2017	MoleMap	40,173	1776	✓	✓		machine	15	69.10
[80]	2017	MoleMap	26584	7975		✓		machine	15	61.20
[80]	2017	MoleMap	26584	7975	✓			machine	15	61.30
[80]	2017	MoleMap	26584	7975	✓	✓		machine	15	70.00
[77]	2017	MoleMap	32,194	8,012	✓	✓		machine	15	71.00
[82]	2014	dermis	2309	1429		✓		machine	44	15.76
[82]	2014	dermis	2309	1429		✓	✓	machine	44	25.12
[100]	2017	Internal	-	2072		✓		human	51	26.40
[100]	2017	Internal	-	2072	✓			human	51	33.10
[32]	2016	SD-198	6,584	3292		✓		machine	198	52.19
[76]	2018	SD-198	-			✓		human	198	49.00
[76]	2018	SD-198	-			✓		human	198	52.00
[76]	2018	SD-198	-			✓		human	198	[illegible]

Skin conditions

1029 skin conditions
excluding melanoma, neoplasms

ICD-10 Version:2016

- XII Diseases of the skin and subcutaneous tissue
  - L00-L08 Infections of the skin and subcutaneous tissue
  - L10-L14 Bullous disorders
  - L20-L30 Dermatitis and eczema
  - L40-L45 Papulosquamous disorders
  - L50-L54 Urticaria and erythema
  - L55-L59 Radiation-related disorders of the skin and subcutaneous tissue
  - L60-L75 Disorders of skin appendages
  - L80-L99 Other disorders of the skin and subcutaneous tissue

https://icd.who.int/browse10/2016/en

- D03 Melanoma in situ
- D04 Carcinoma in situ of skin
- D23 Other benign neoplasms of skin
- C43-C44 Melanoma and other malignant neoplasms of skin

Which classes to predict?

Large number of hierarchical conditions
Which conditions, at what level of granularity?

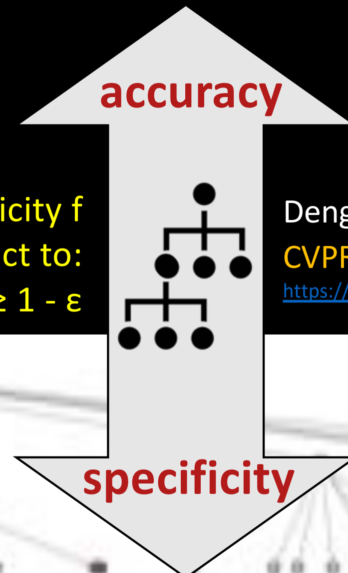

max specificity f
subject to:
accuracy (f) ≥ 1 - ε

Deng, Krause, Berg, Fei-Fei
CVPR 2012 IMAGENET
https://ieeexplore.ieee.org/document/6248086

Romero-Lopez et al.
SOCML 2019

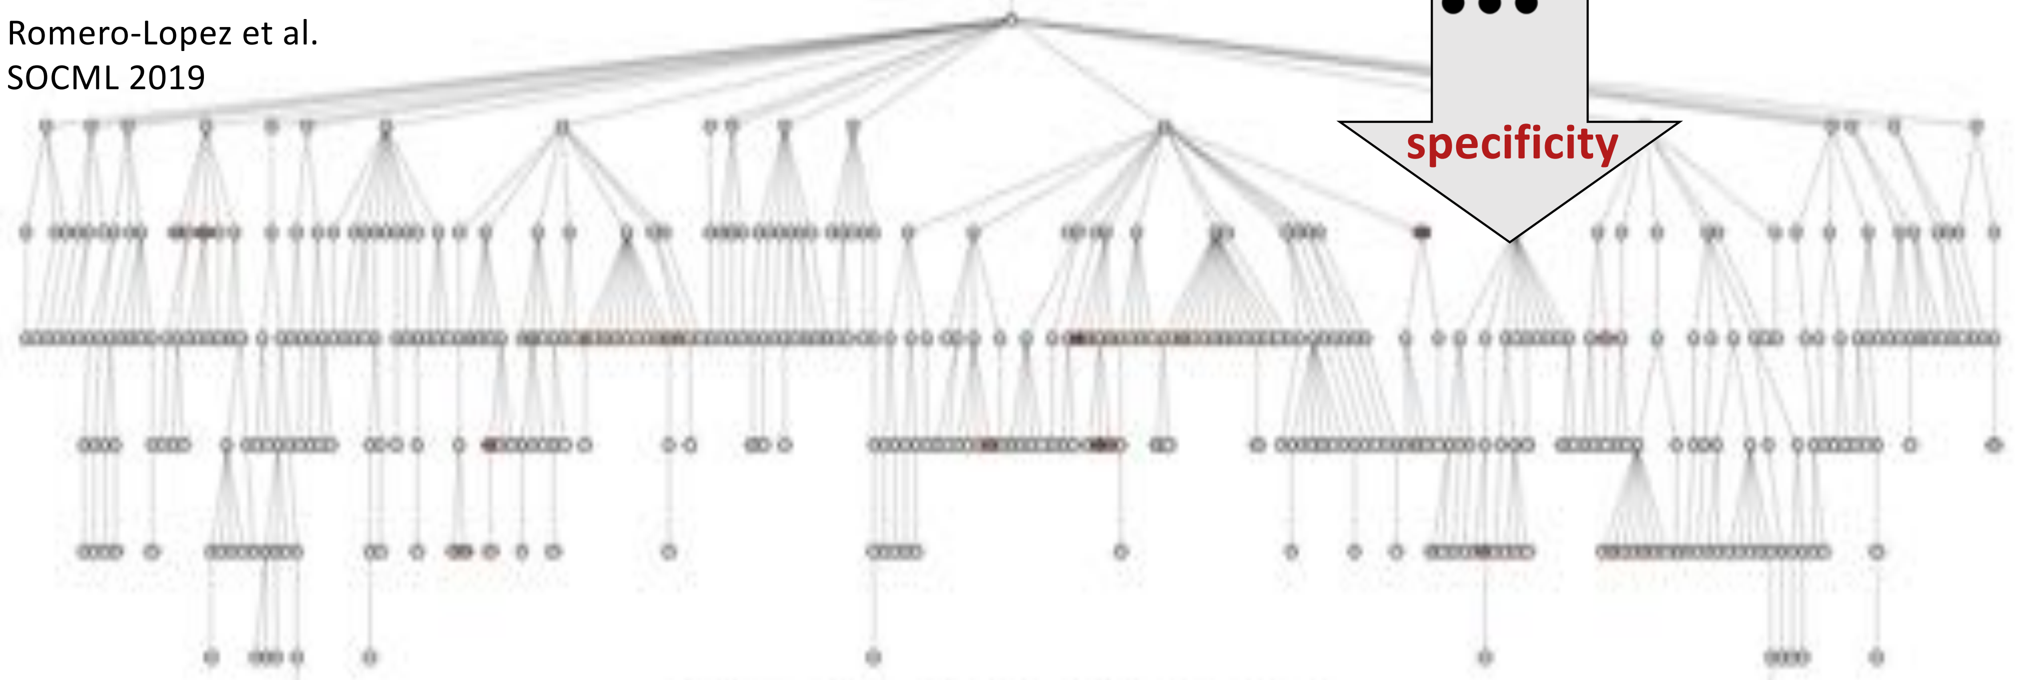

Triage skin conditions ontology

133 diagnosis nodes parents to 588 different skin conditions

Other prediction tasks

localize

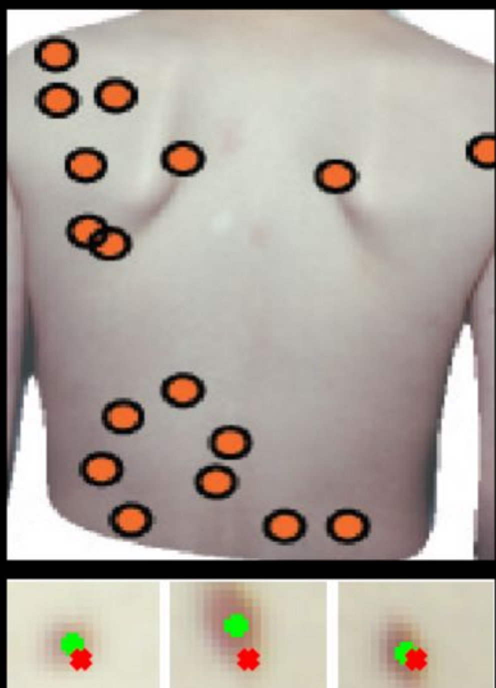

manual predicted

Mirzaalian, Hamarneh, Lee
CVPR 2009
https://doi.ieeecomputersociety.org/10.1109/CVPR.2009.5206725

segment

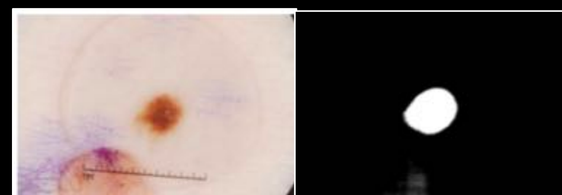

Mirikharaji, Hamarneh
MICCAI 2018
https://link.springer.com/chapter/10.1007/978-3-030-00937-3_84

Izadi, Mirikharaji, Kawahara, Hamarneh
ISBI 2018
https://ieeexplore.ieee.org/abstract/document/8363712

Taskonomy CVPR 2018
Zamir, Savarese, Malik

hair removal

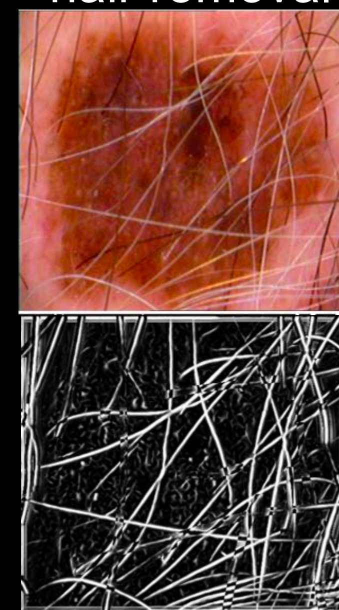

Mirzaalian, Hamarneh, Lee
IEEE TIP 2014
https://ieeexplore.ieee.org/document/6918479/

Other prediction tasks

Joint skin lesion localization and segmentation

Vesal, Patil, Ravikuma, Maier
ISIC 2018
https://link.springer.com/chapter/10.1007/978-3-030-01201-4_31

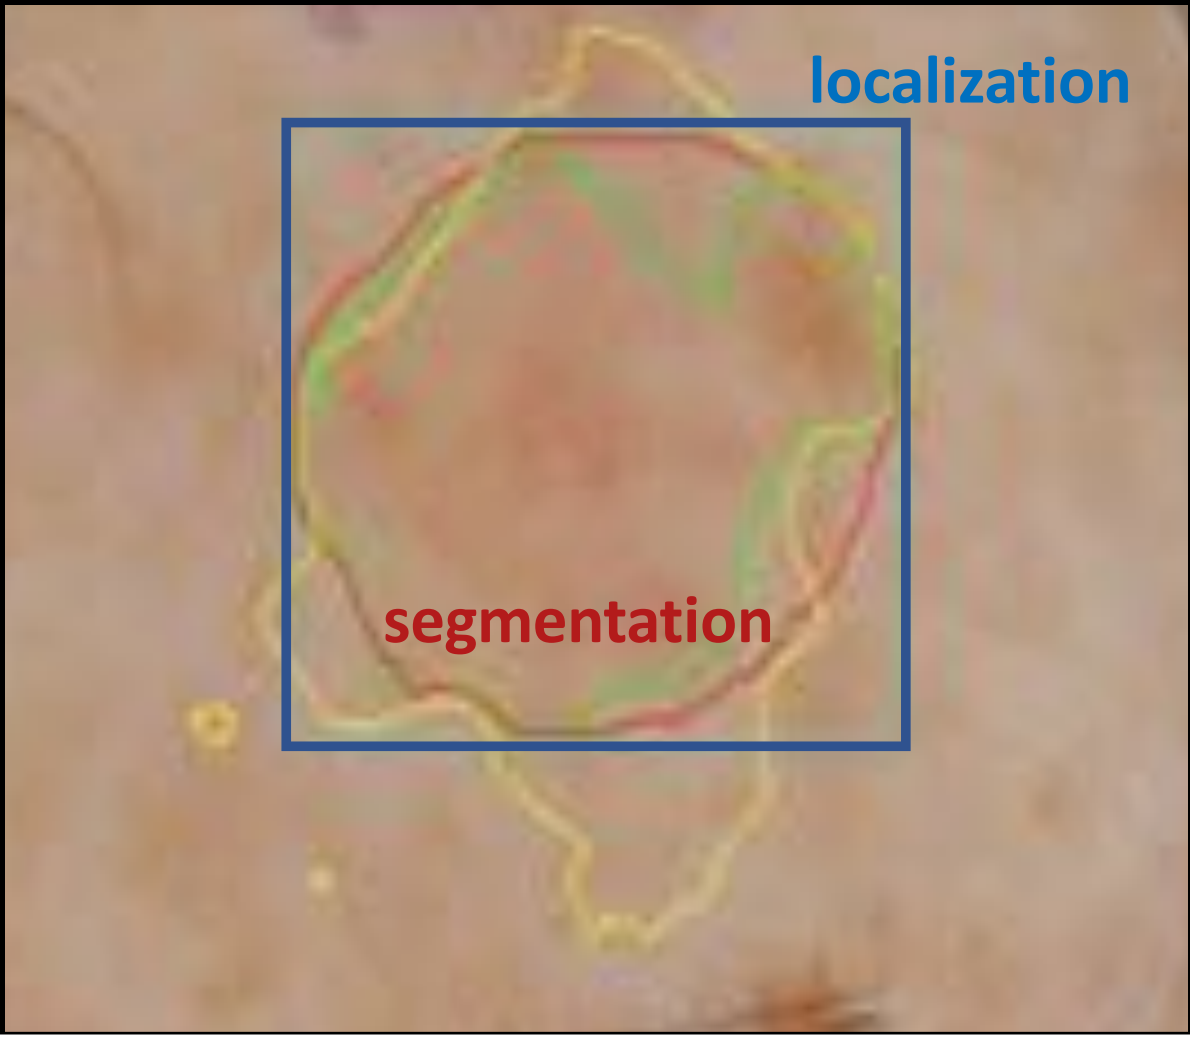

Blue: detected bounding box
Green: GT lesion boundary
Yellow: SkinNet
Red: Faster-RCNN+SkinNet

Real-life skin images

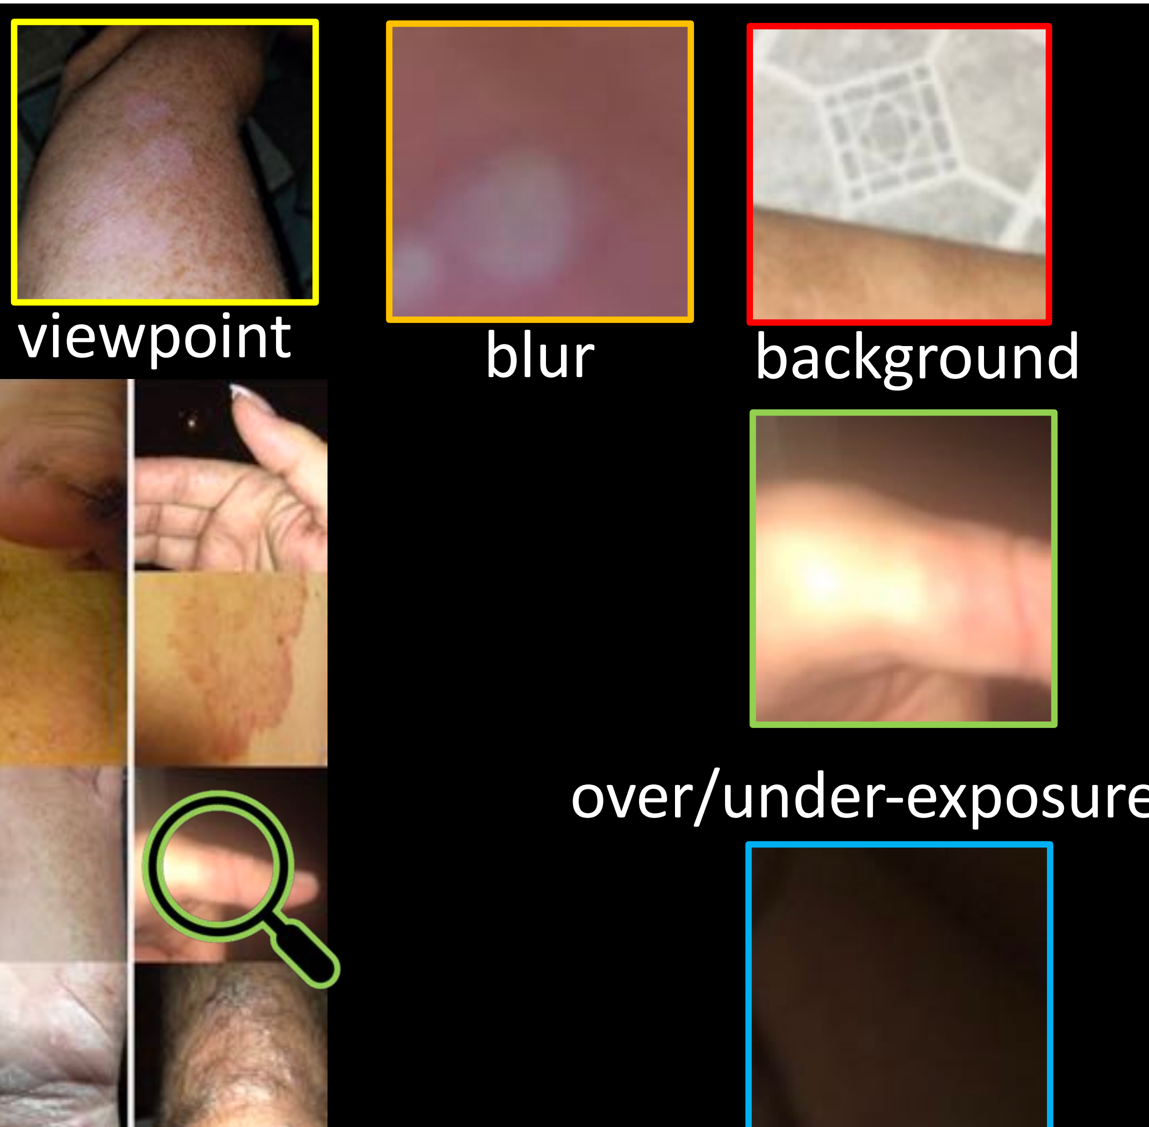

Annotations

data
100 clean
1500 noisy

Fully- weakly- un-supervised

100 images with clean labels

vs

1000 images with noisy labels?

Adaptively handle noisy annotations via modified deep model optimization

1. learn weight map W to control pixel contribution to loss
2. ↑W ⇔ ↑ agreement with clean data
   agreement in loss gradient

Mirikharaji, Yan, Hamarneh, 2019
https://arxiv.org/abs/1906.03815

Dice %

original pretrain on noisy fine-tune on clean	modified
78.6	N/A
76.1	80.7
75.0	80.3
73.0	79.5
70.5	73.6

Annotations

with only ~100 clean annotations, the reweighting outperforms the fully-supervised with 2000 clean annotations

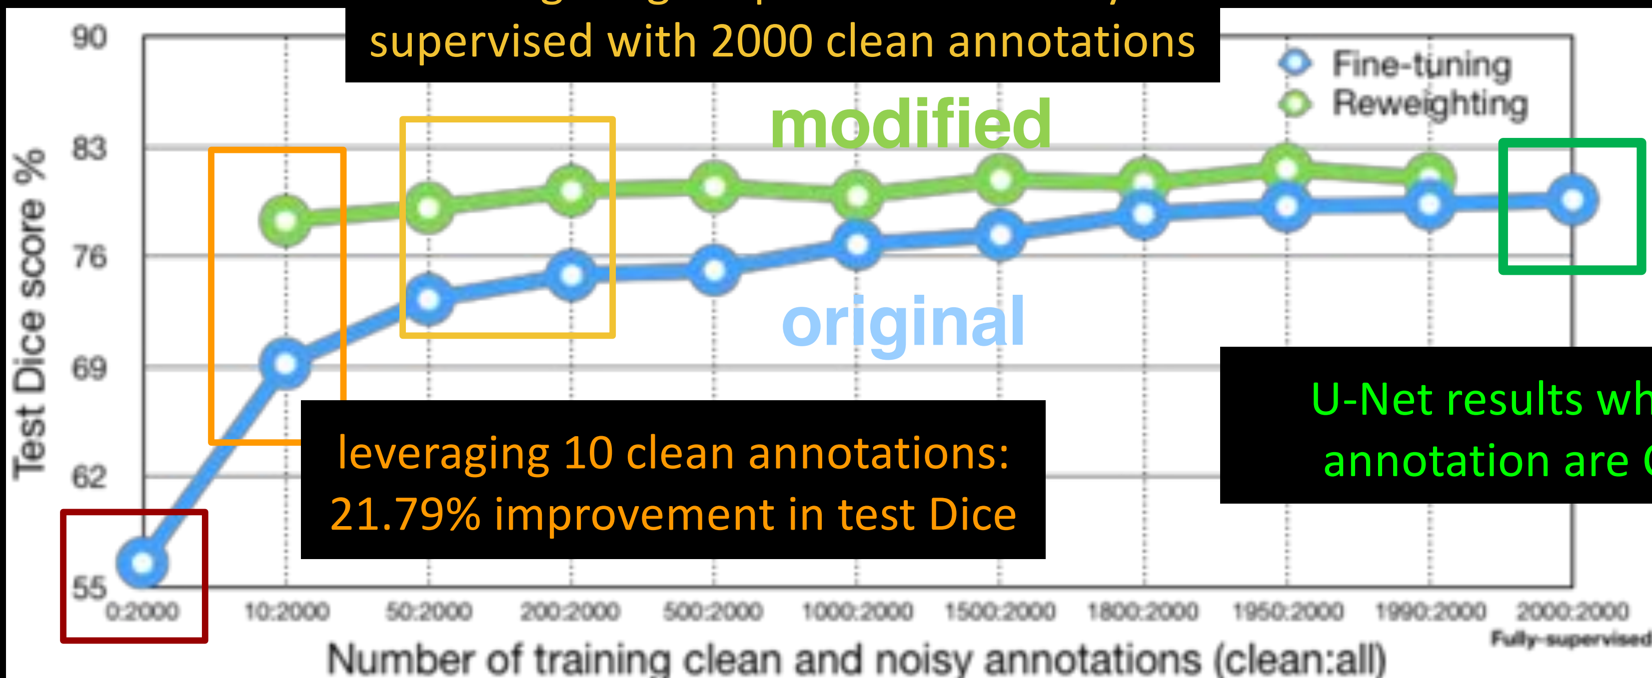

leveraging 10 clean annotations: 21.79% improvement in test Dice

U-Net results when all annotation are Clean

U-Net results when all annotation are noisy

Loss choice

Kawahara, Daneshvar, Argenziano, Hamarneh
IEEE JBHI 2019
https://ieeexplore.ieee.org/document/8333693

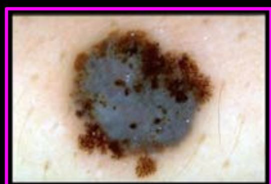

CNN parameters

7 Criteria

weighted cross-entropy

$$L(x, y, z; \theta) = \ell(x, y; \theta) + \sum_{j=1}^{7} \ell(x, z_j; \theta)$$

Diagnosis

7-point Criteria

The j-th criterion from the 7-point Criteria

1. Basal cell carcinoma (BCC)
2. Nevus (multiple) (NEV)
3. Melanoma (multiple) (MEL)
4. Seborrheic keratosis (SK)
5. Miscellaneous (MISC)

1. Pigment Network (PN)
2. Blue Whitish Veil (BWV)
3. Vascular Structures (VS)
4. Pigmentation (PIG)
5. Streaks (STR)
6. Dots and Globules (DaG)
7. Regression Structures (RS)

Weighted cross-entropy loss

Given input data e.g,. image

Number of labels in a class

Standard cross-entropy loss

$$\ell(x, c; \theta) = -\frac{1}{N_b}\sum_{i=1}^{N_b}\sum_{j=1}^{N_c} w(c)_j \cdot \overbrace{c_j^{(i)} \cdot \underbrace{\log\left(\phi(x^{(i)};\theta)_{c,j}\right)}}$$

Given class labels (e.g., diagnosis or 7-point criteria)

Number of samples in a mini-batch

Weighting function

Predicted log-likelihood of the j-th label for the c-th class of the i-th sample

Class imbalance

Context: Melanoma diagnosis; 7 dermoscopic criteria

Ensure each mini-batch includes ≥k (random) samples from each label

1. Pigment Network (PN)
2. Blue Whitish Veil (BWV)
3. Vascular Structures (VS)
4. Pigmentation (PIG)
5. Streaks (STR)
6. Dots and Globules (DaG)
7. Regression Structures (RS)

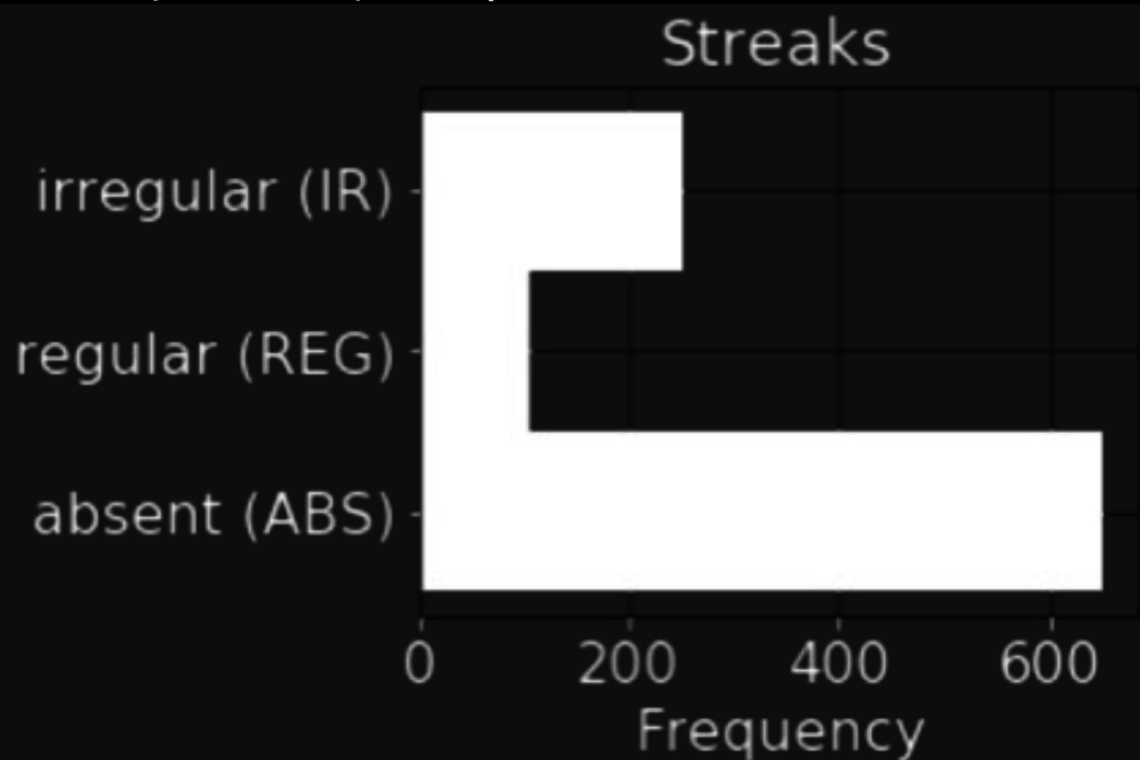

Higher cross entropy weights assigned to infrequent labels in a mini-batch

Data augmentation

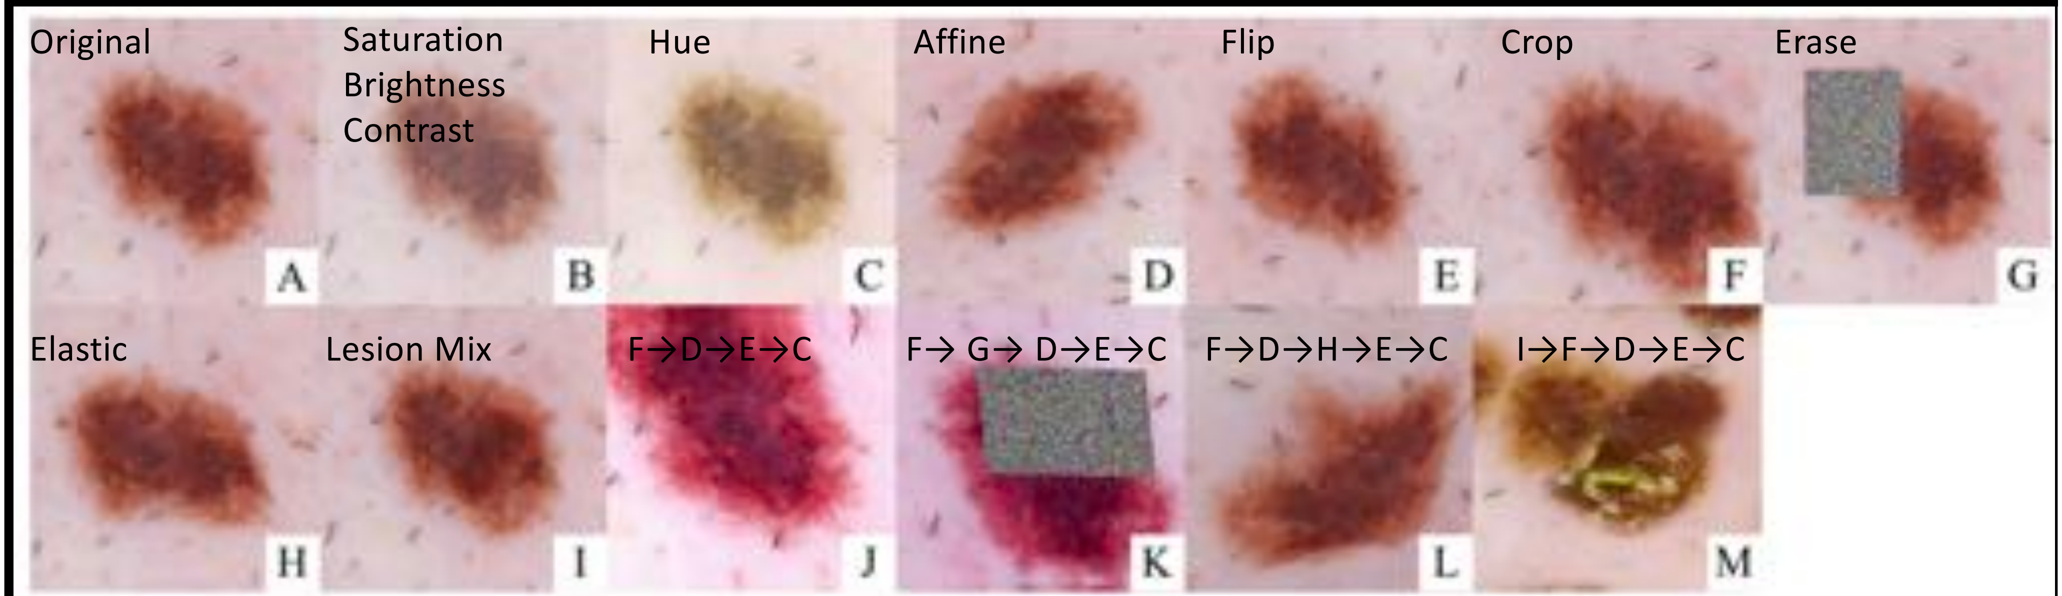

Data Augmentation for Skin Lesion Analysis
Perez, Vasconcelos, Avila, Valle
ISIC 2018
https://link.springer.com/chapter/10.1007/978-3-030-01201-4_33

Data augmentation

Simulation via Physically- and Statistically-based Warps

DeformIt

MICCAI 2008

Hamarneh, Jassi, Tang, Booth

https://link.springer.com/chapter/10.1007/978-3-540-85988-8_55
https://ieeexplore.ieee.org/document/6867974

$$I = \bar{I} + \alpha \mathbf{Pb} + (1-\alpha)\Phi \mathrm{u}$$

variational PCA

Vibrational FEM

$\uparrow$data ➔ $\uparrow\alpha$
rely more on statistical model and less on knowledge-based models

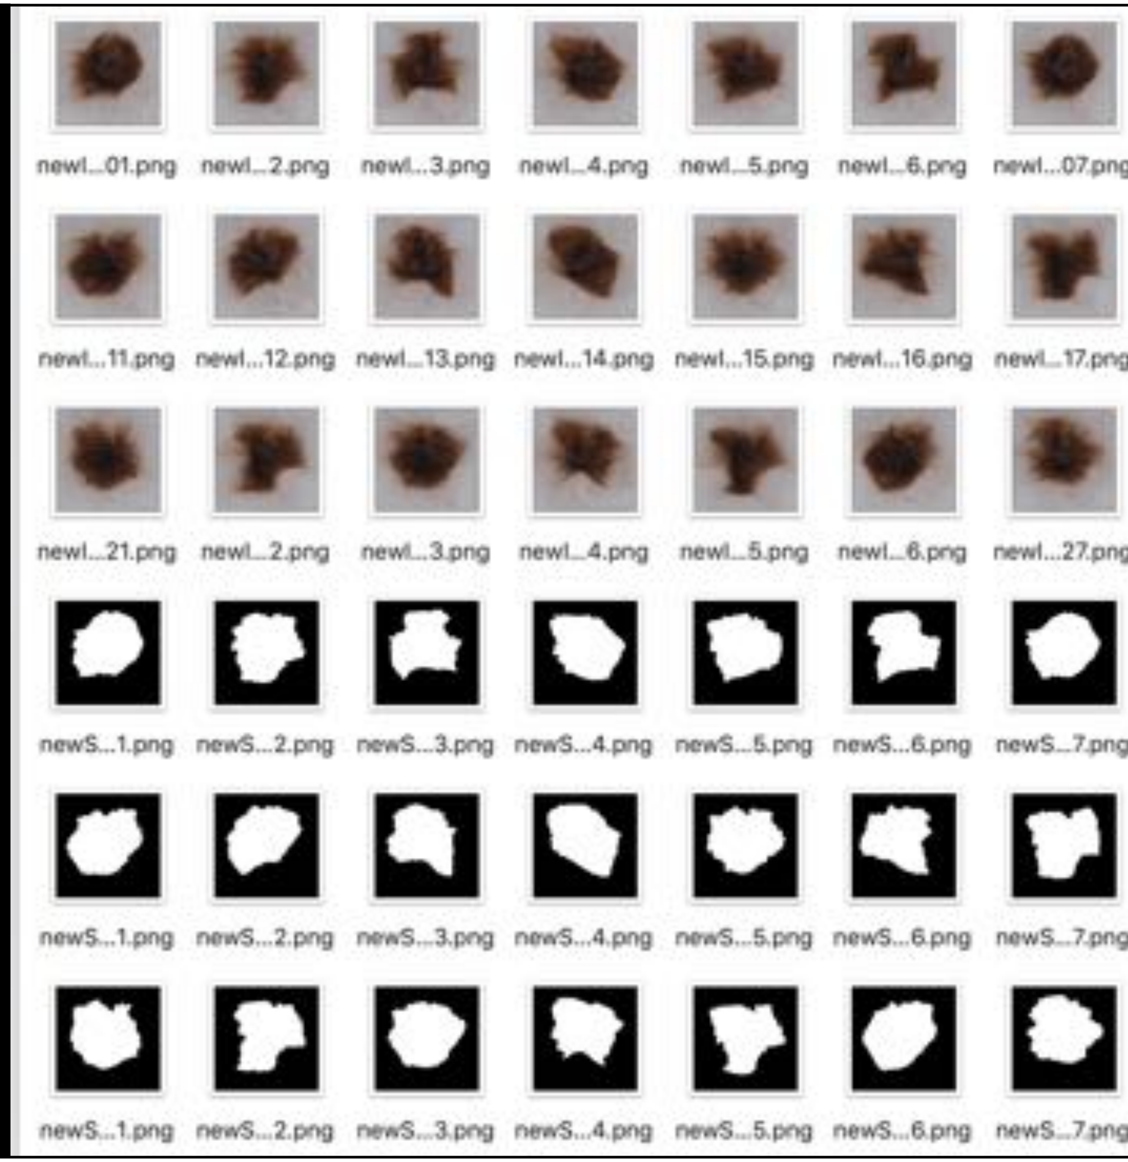

Augmentation

Hair Occlusion Simulator

HairSim

IEEE TIP 2014

Mirzaalian, Lee , Hamarneh

https://ieeexplore.ieee.org/document/6918479

- medial A-B curve synthesizer
- hair-thickening: dilation radius ∝ geodesic distance to A and B

$$r(p) = \min\{T, \alpha\Gamma(p, A), \alpha\Gamma(p, B)\}$$

- New image (H): blending of clean image I with a colored C hair mask M, Hair color C

$$\begin{bmatrix} H_R \\ H_G \\ H_B \end{bmatrix} = I(\mathbf{1} - G_\sigma * \boldsymbol{M}) + \begin{bmatrix} C_R \\ C_G \\ C_B \end{bmatrix} (G_\sigma * M)$$

Augmentation

Generative Adversarial Networks GANS

Generating Highly Realistic Images of Skin Lesions with GANs (MelanoGan)
Baur, Albarqouni, Navab ISIC 2018
https://link.springer.com/chapter/10.1007/978-3-030-01201-4_28
https://arxiv.org/abs/1804.04338

Skin Lesion Synthesis with GANs
Bissoto, Perez, Valle, Avila, ISIC 2018
https://link.springer.com/chapter/10.1007/978-3-030-01201-4_32

Augmenting data with GANs to segment melanoma skin lesions
Pollastri, Bolelli, Paredes, Grana
Multimedia Tools and Applications 2019
https://link.springer.com/article/10.1007/s11042-019-7717-y

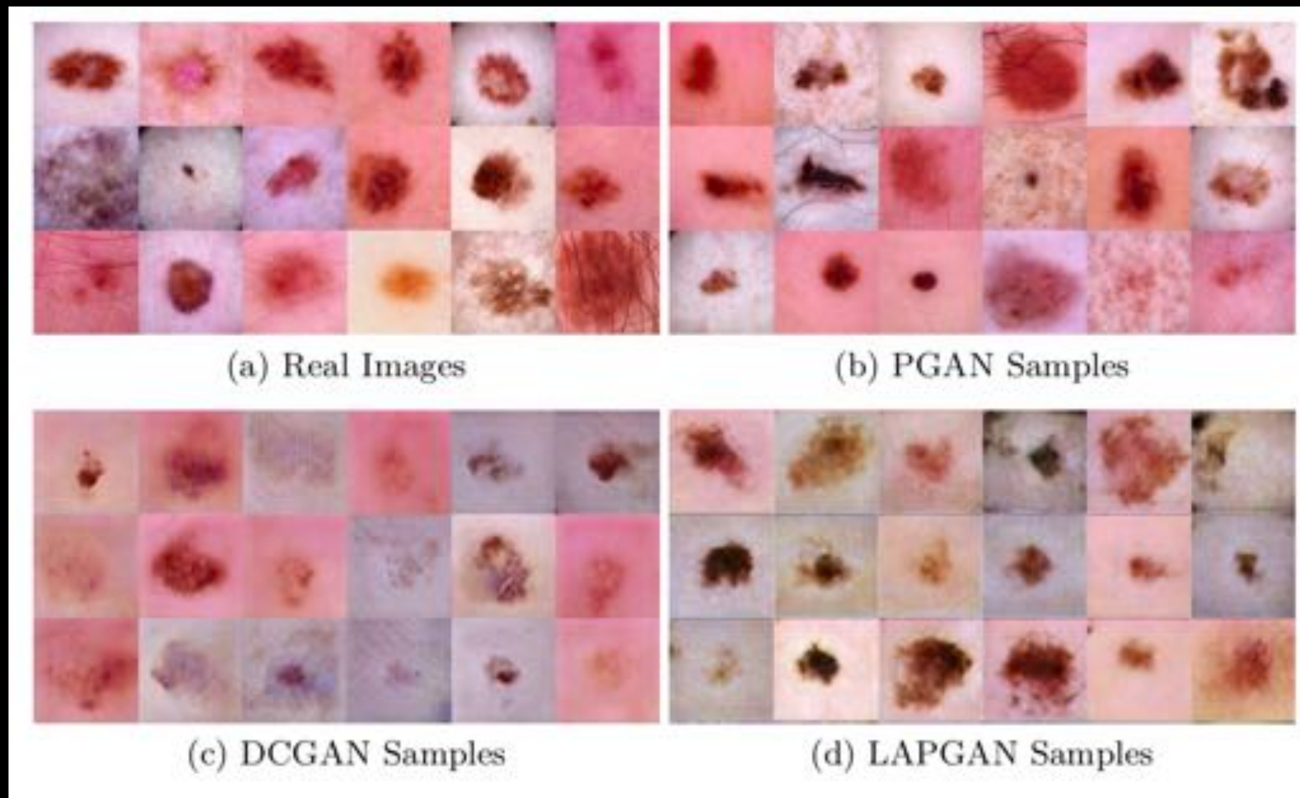

Augmentation

GAN-based **Mask2Lesion** translation

Abhishek, Hamarneh. 2019 https://arxiv.org/abs/1906.05845

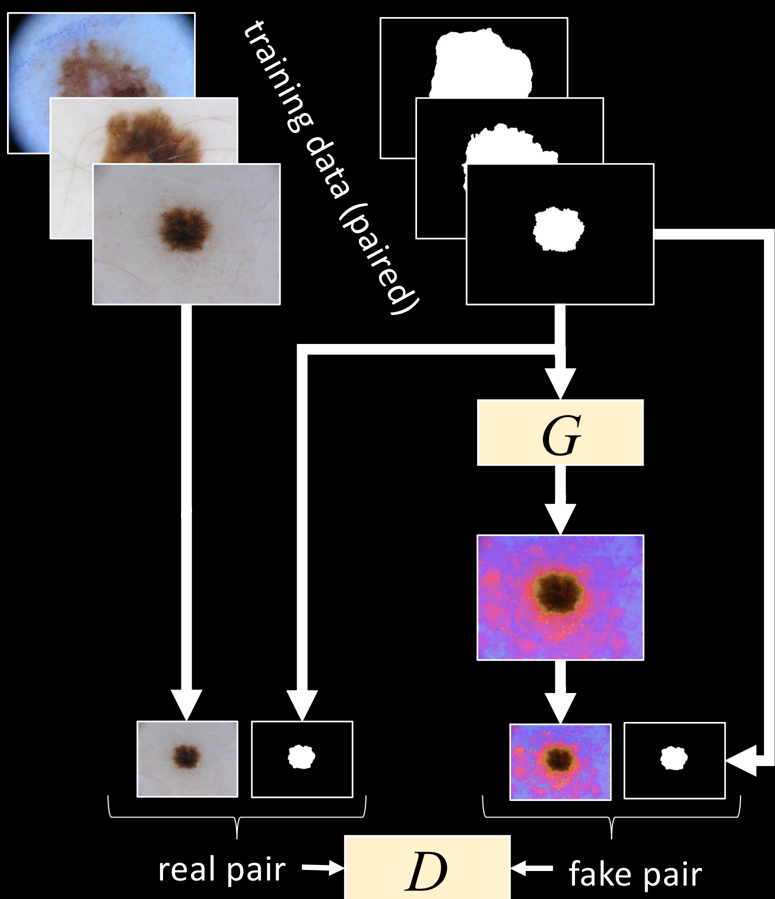

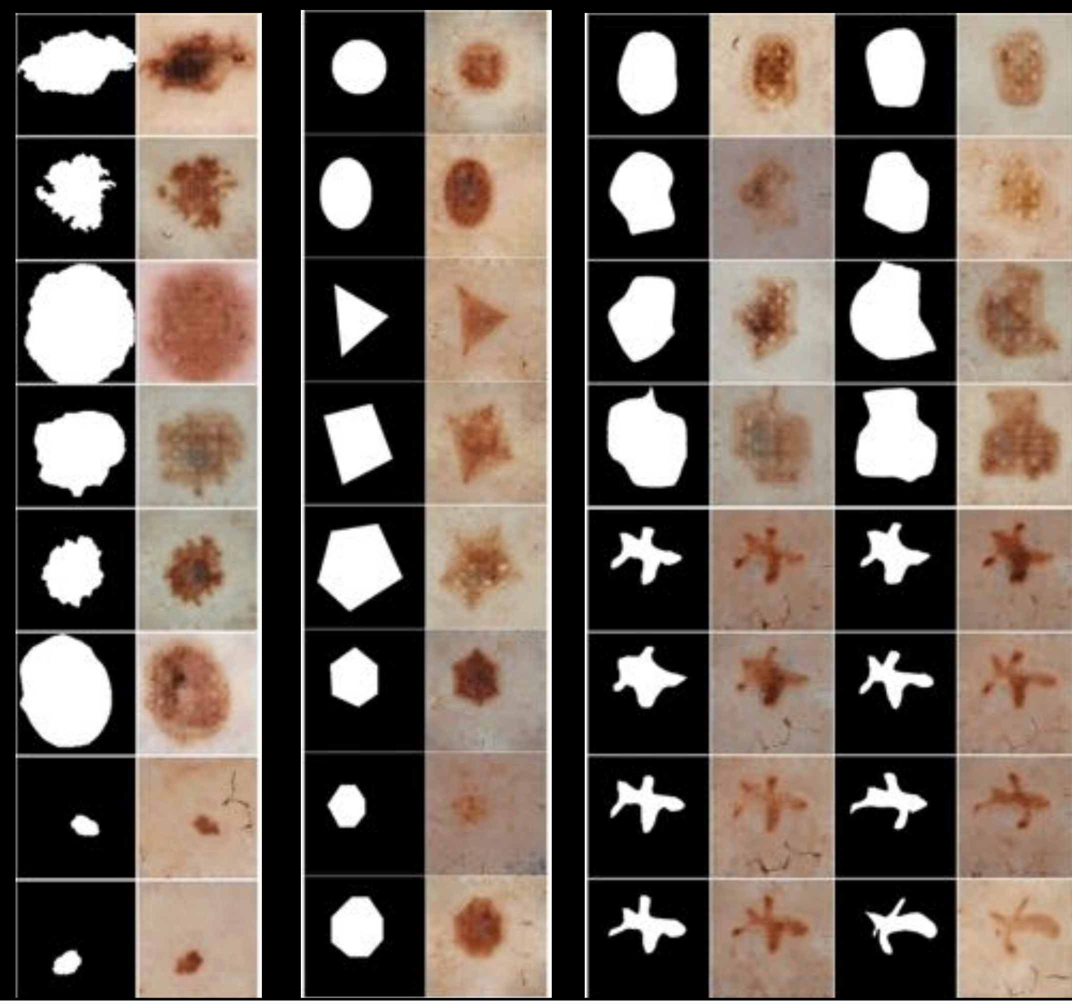

Network architecture

Deep auto-context FCN for skin lesion segmentation
Mirikharaji, Izadi, Kawahara, Hamarneh ISBI2018
https://ieeexplore.ieee.org/document/8363711

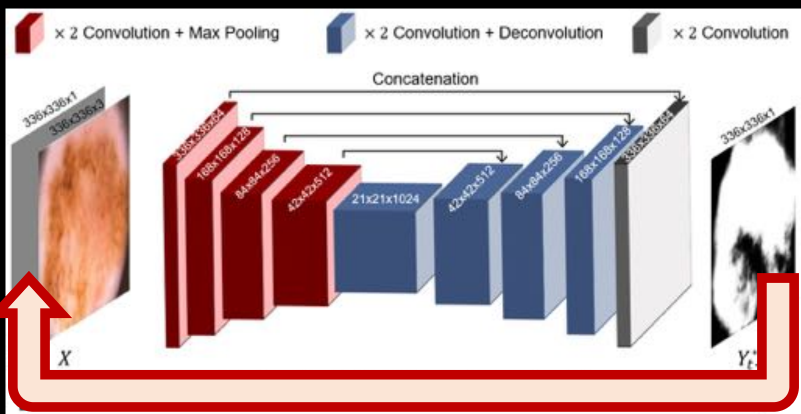

Generative adversarial networks to segment skin lesions
Izadi, Mirikharaji, Kawahara, Hamarneh ISBI 2018
https://ieeexplore.ieee.org/abstract/document/8363712

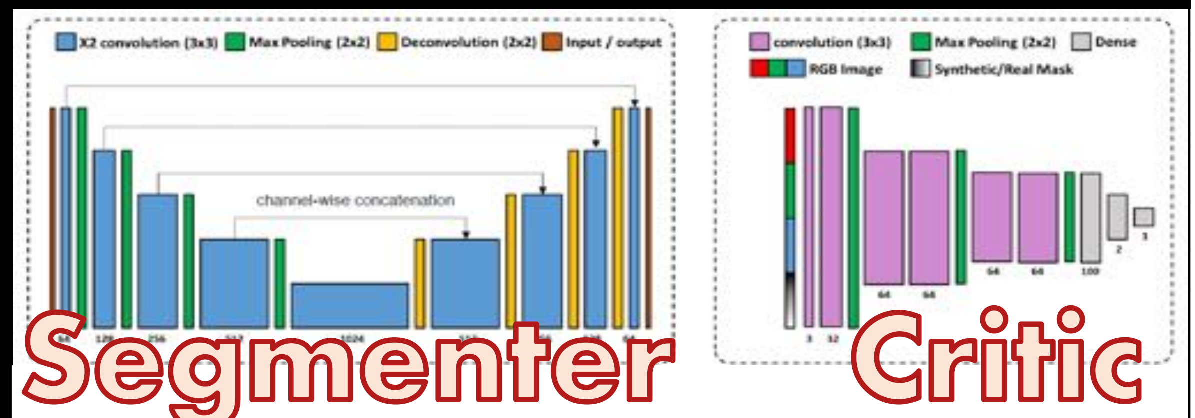

diagnosis &
7-point derm.
criteria

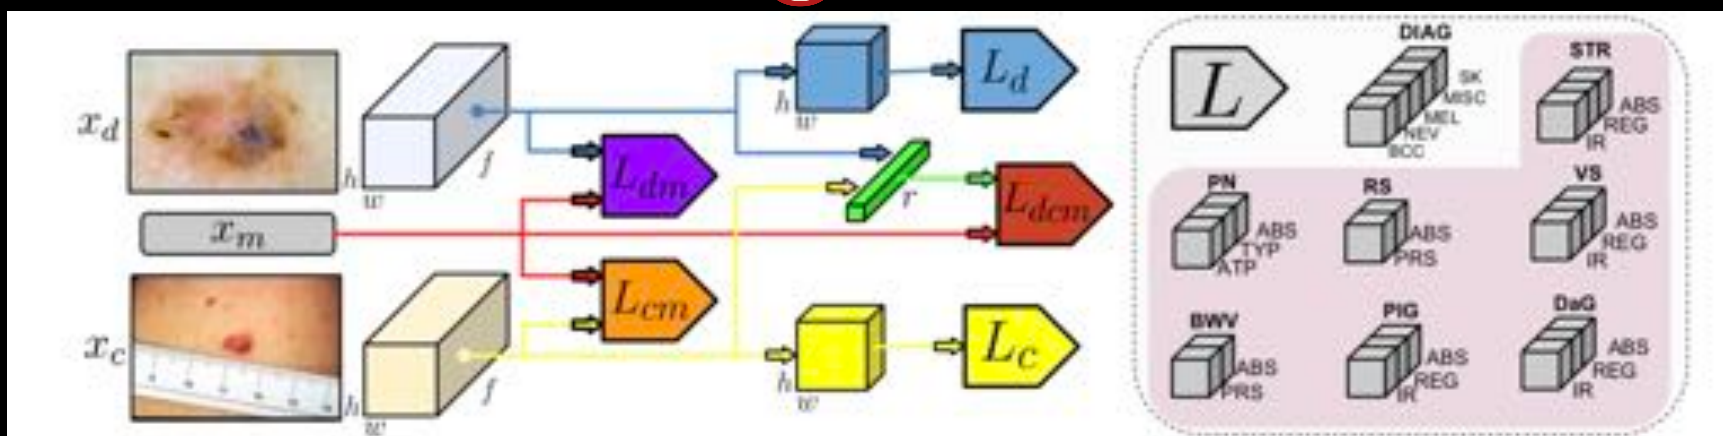

Network architecture search (NAS)

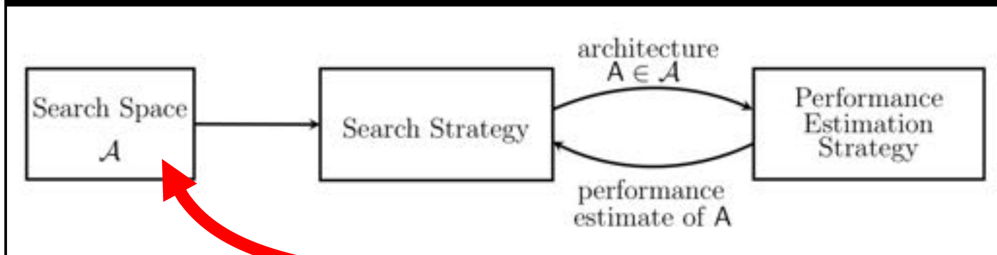

Neural Architecture Search: A Survey
Thomas Elsken, Jan Hendrik Metzen, Frank Hutter; 20(55):1–21, 2019
https://arxiv.org/abs/1808.05377

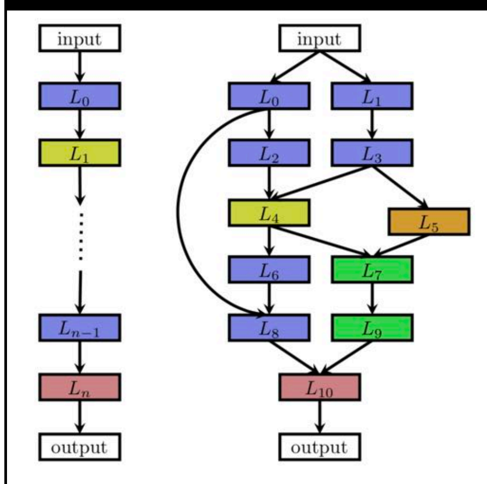

Arrangement of layers

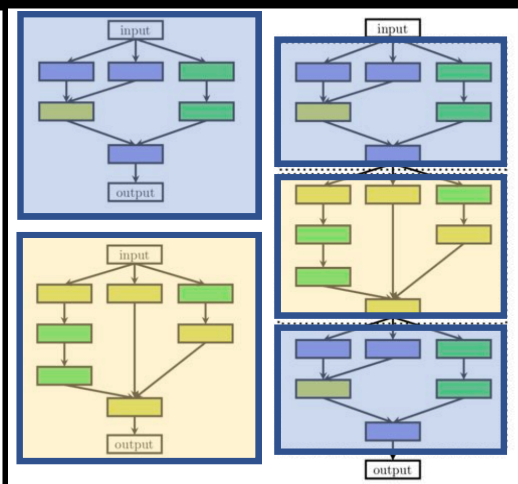

or blocks of layers

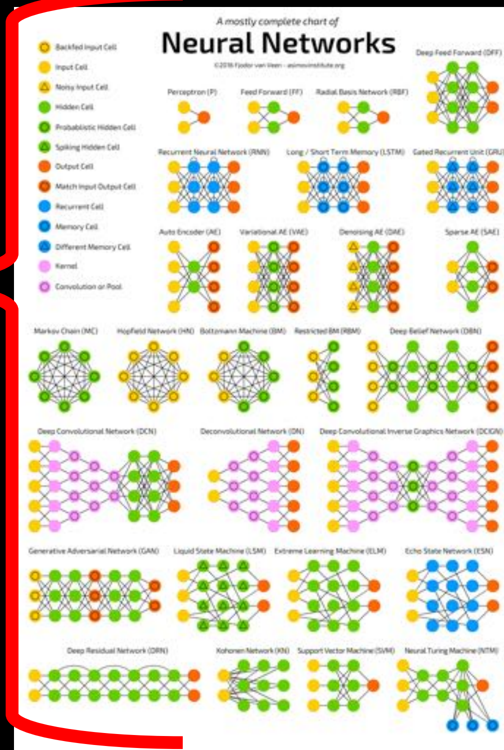

https://www.asimovinstitute.org/author/fjodorvanveen

Network architecture search (NAS)

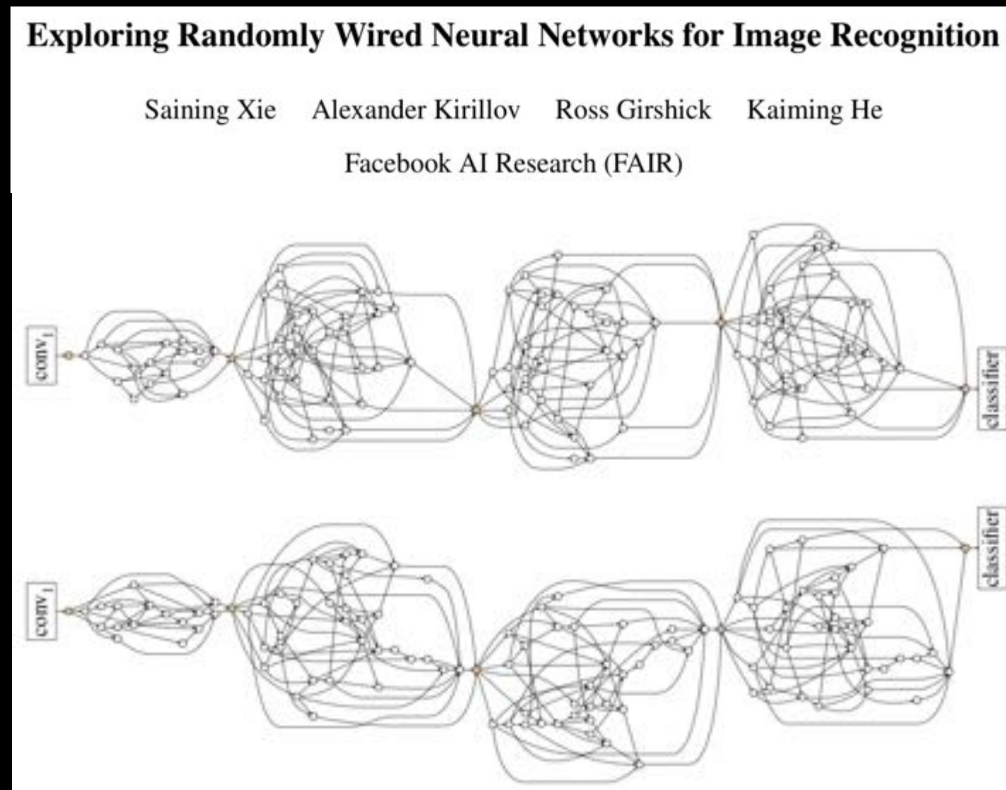
Exploring Randomly Wired Neural Networks for Image Recognition

Saining Xie Alexander Kirillov Ross Girshick Kaiming He

Facebook AI Research (FAIR)

Exploring Randomly Wired Neural Networks for Image Recognition
Xie, Kirillov, Girshick, He 2019
https://arxiv.org/abs/1904.01569

Layer design

Layers: [de]conv, fully conn., [un]pool, sequence eg LSTM, activation eg RELU

Radial Basis Function RBF layer with learnable width β, center c, transformation Ψ, bias b

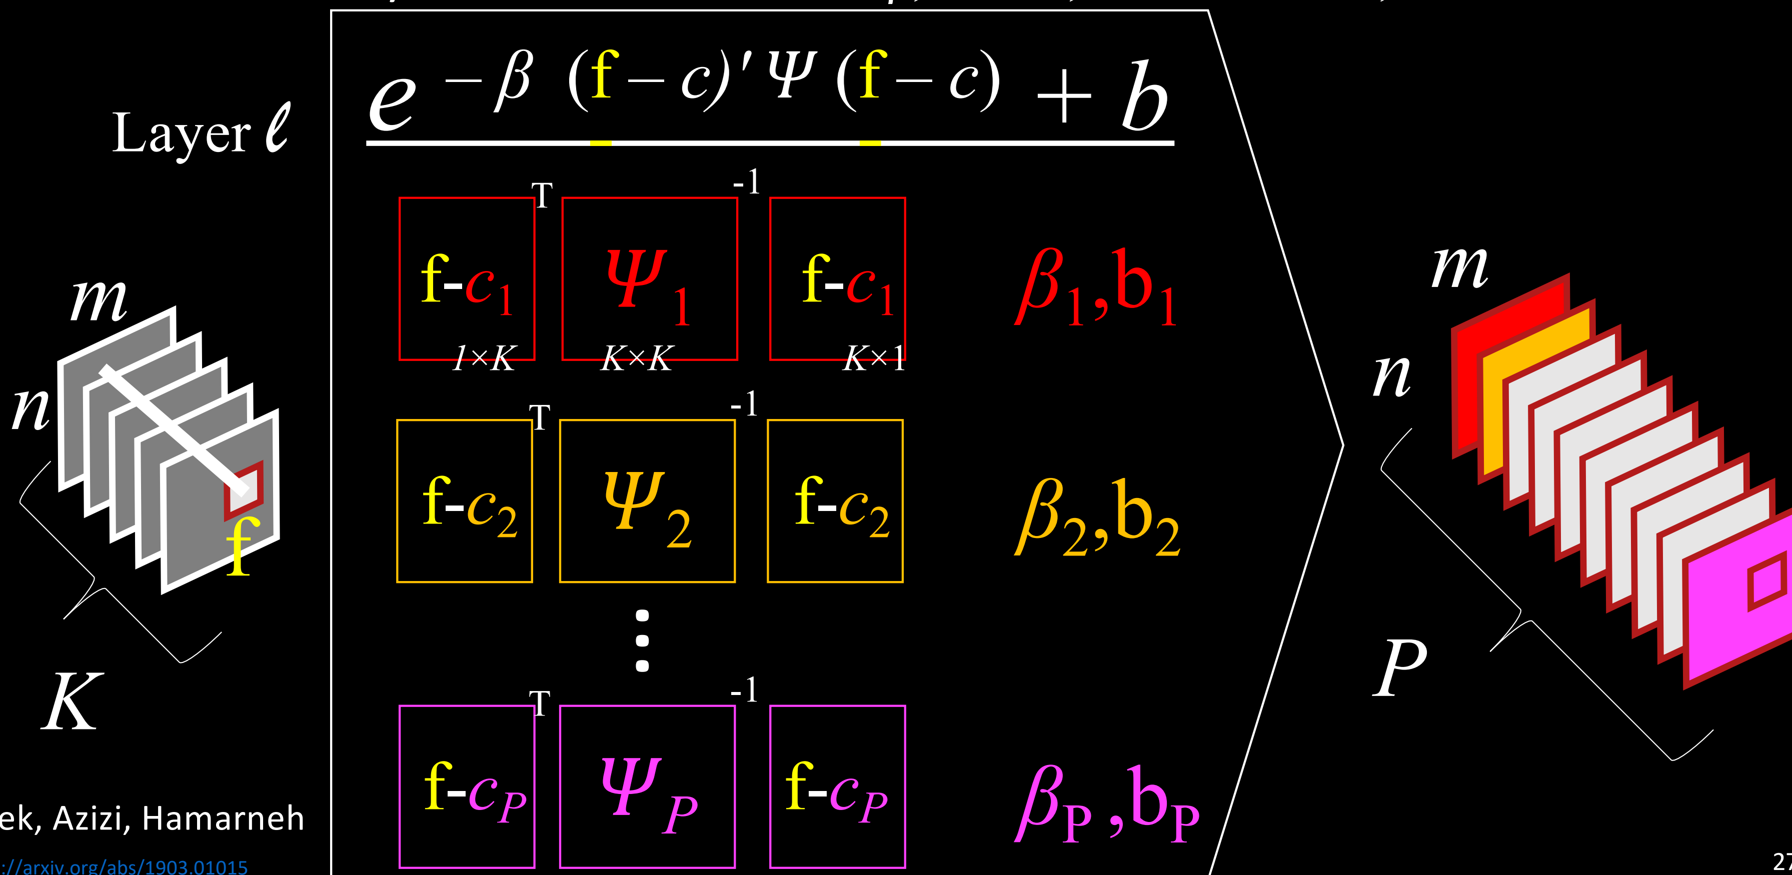

Asgari, Abhishek, Azizi, Hamarneh
CVPR 2019 https://arxiv.org/abs/1903.01015

Layer design

Radial Basis Function RBF layer with learnable width β, center c, transformation Ψ, bias b

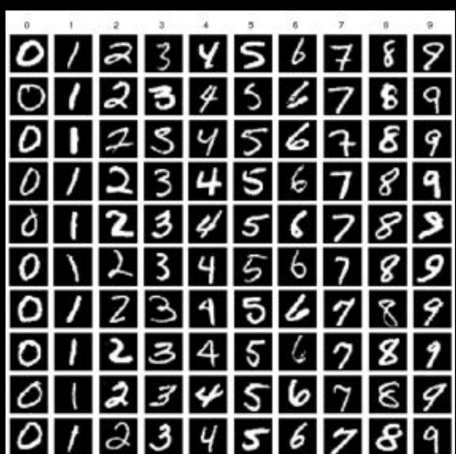

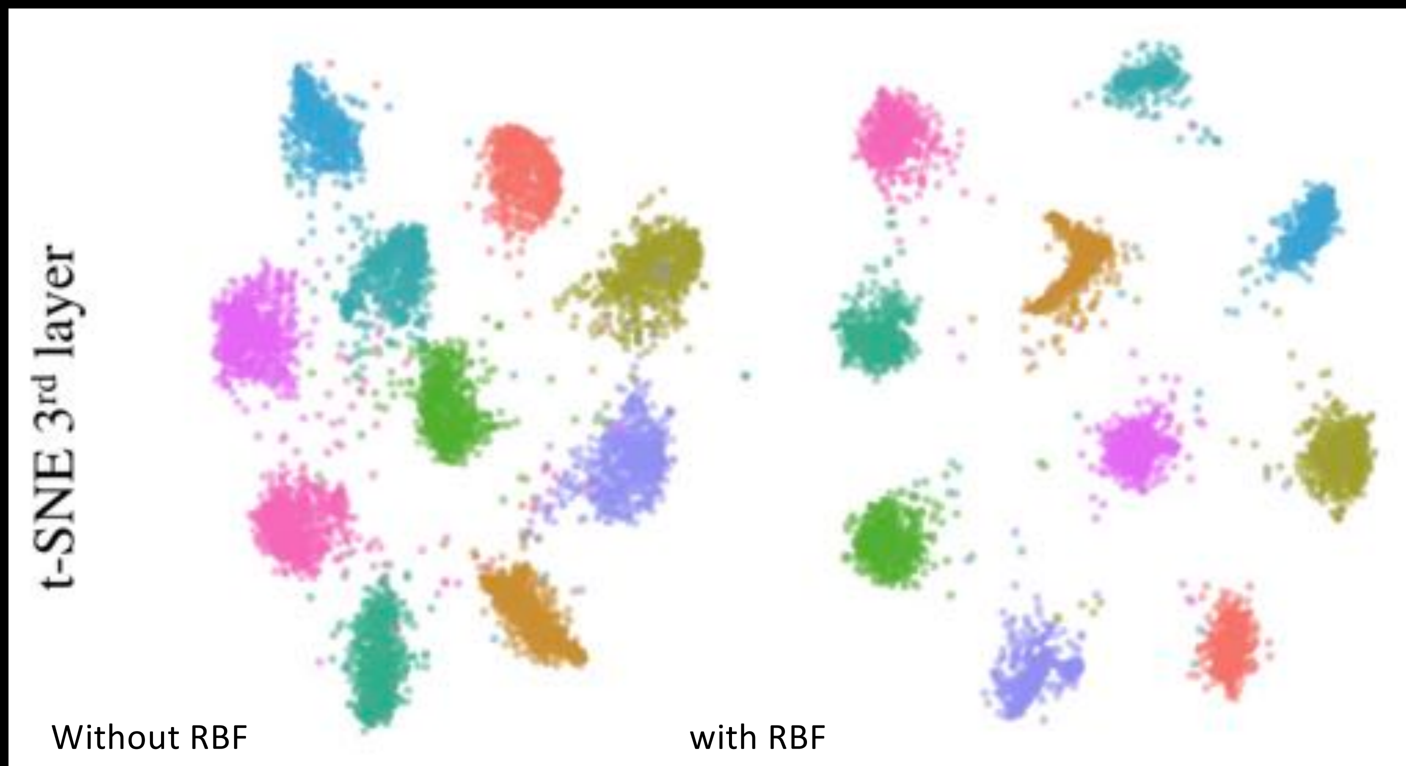

Asgari, Abhishek, Azizi, Hamarneh
CVPR 2019 https://arxiv.org/abs/1903.01015

Adversarial attacks

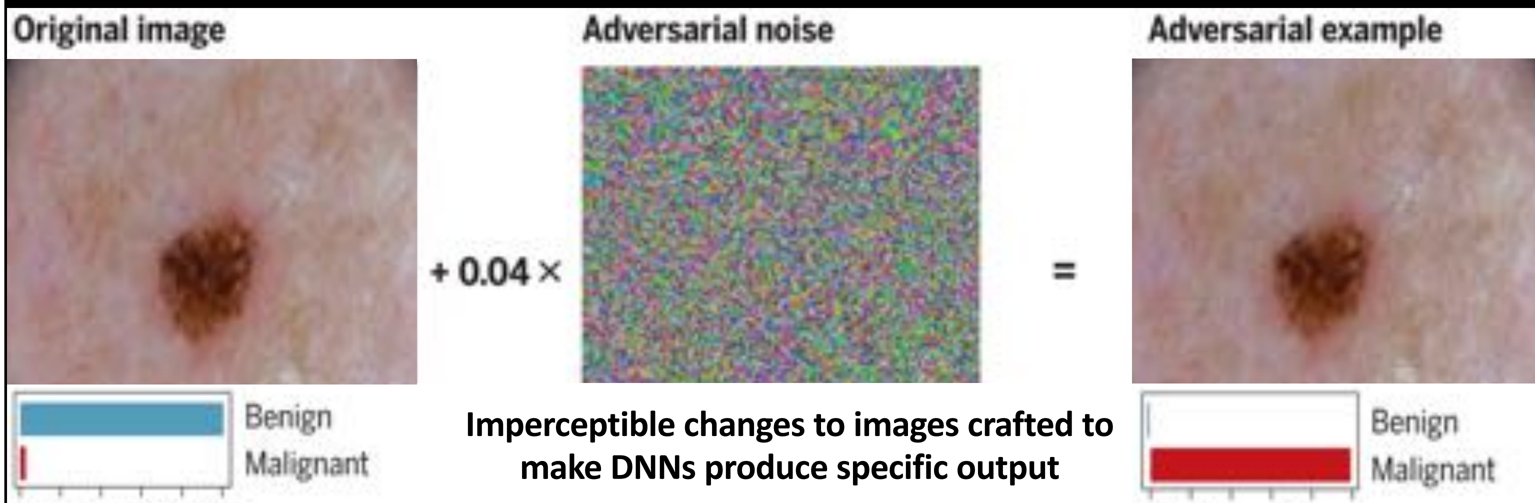

Imperceptible changes to images crafted to make DNNs produce specific output

Finlayson, Bowers, Ito, Zittrain, Beam, Kohane
Science 2019
https://science.sciencemag.org/content/363/6433/1287

Asgari, Abhishek, Azizi, Hamarneh
CVPR 2019
https://arxiv.org/abs/1903.01015

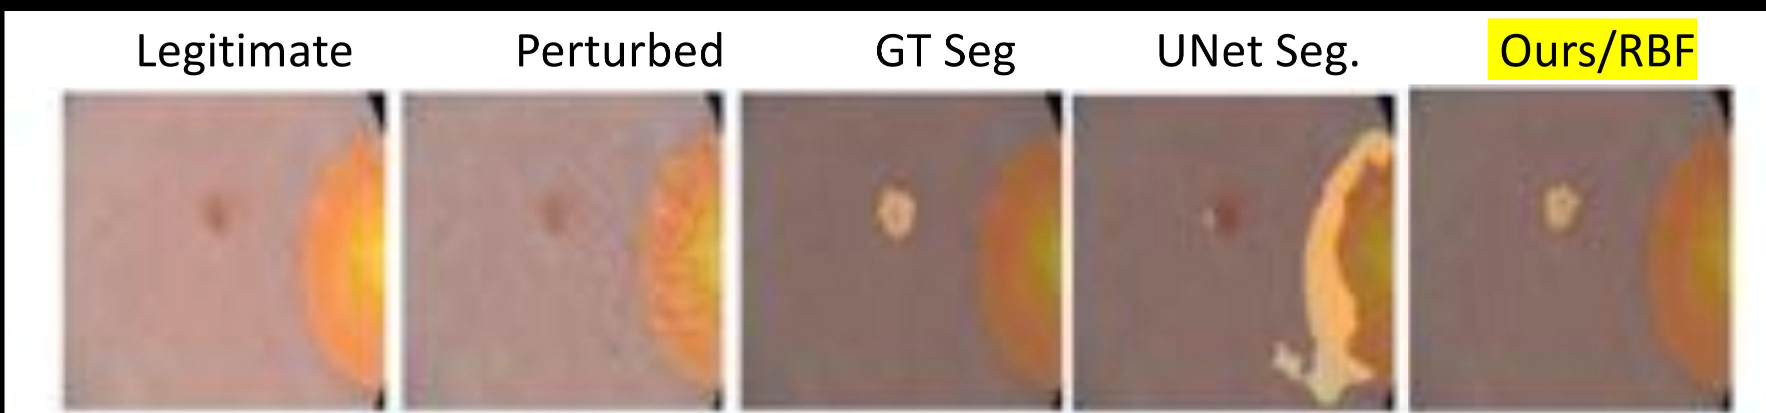

Dataset shift

Bias and fairness

Machine Learning and Health Care Disparities in Dermatology

Unfortunately, most ML programs are largely learning on light skin. For example, in the International Skin Imaging Collaboration: Melanoma Project, which is one of the largest and often-used, open-source, public-access archives of pigmented lesions, much of the patient data are heavily collected from fair-skinned populations in the United States, Europe, and Australia.[3] Thus, no matter how advanced the ML algorithm, it may underperform on images of lesions in skin of color.

JAMA Dermatology 2018;154(11)
Adamson, Smith
https://jamanetwork.com/journals/jamadermatology/article-abstract/2688587

ORIGINAL ARTICLE

Classification of the Clinical Images for Benign and Malignant Cutaneous Tumors Using a Deep Learning Algorithm

Seung Seog Han[1,7], Myoung Shin Kim[2,7], Woohyung Lim[3], Gyeong Hun Park[4], Ilwoo Park[5] and Sung Eun Chang[6]

Because of the different patient demographics in the three validation datasets we tested with our algorithm, the sensitivity and specificity of these datasets were analyzed over a change in threshold from 0.0000 to 1.0000 (Figure 4). The sensitivities of the Asan and Hallym test dataset over this threshold were similar. However, the specificities for BCC, squamous cell carcinoma, and melanoma between the Asan test dataset and Edinburgh dataset showed substantial differences, which may have been due to malignancy subtypes and the skin colors around the lesions. It may be necessary, therefore, to choose different thresholds or generate different models for different ethnic groups.

Journal of Investigative Dermatology 2018; 138(7)
Han et al.
https://jamanetwork.com/journals/jamadermatology/article-abstract/2688587

Dataset shift

Bias and fairness

Test dataset of European population:
10 classes - 1300 images

Train and test on same dataset

Deep features to classify skin lesions
Kawahara, BenTaieb, Hamarneh
ISBI 2016
https://ieeexplore.ieee.org/document/7493528

Train on Asian, test on European

Classification of the Clinical Images for Benign and Malignant Cutaneous Tumors Using a Deep Learning Algorithms
Han, Kim, Lim, Park, Park, Chang
Journal of Investigative Dermatology
https://www.jidonline.org/article/S0022-202X(18)30111-8/

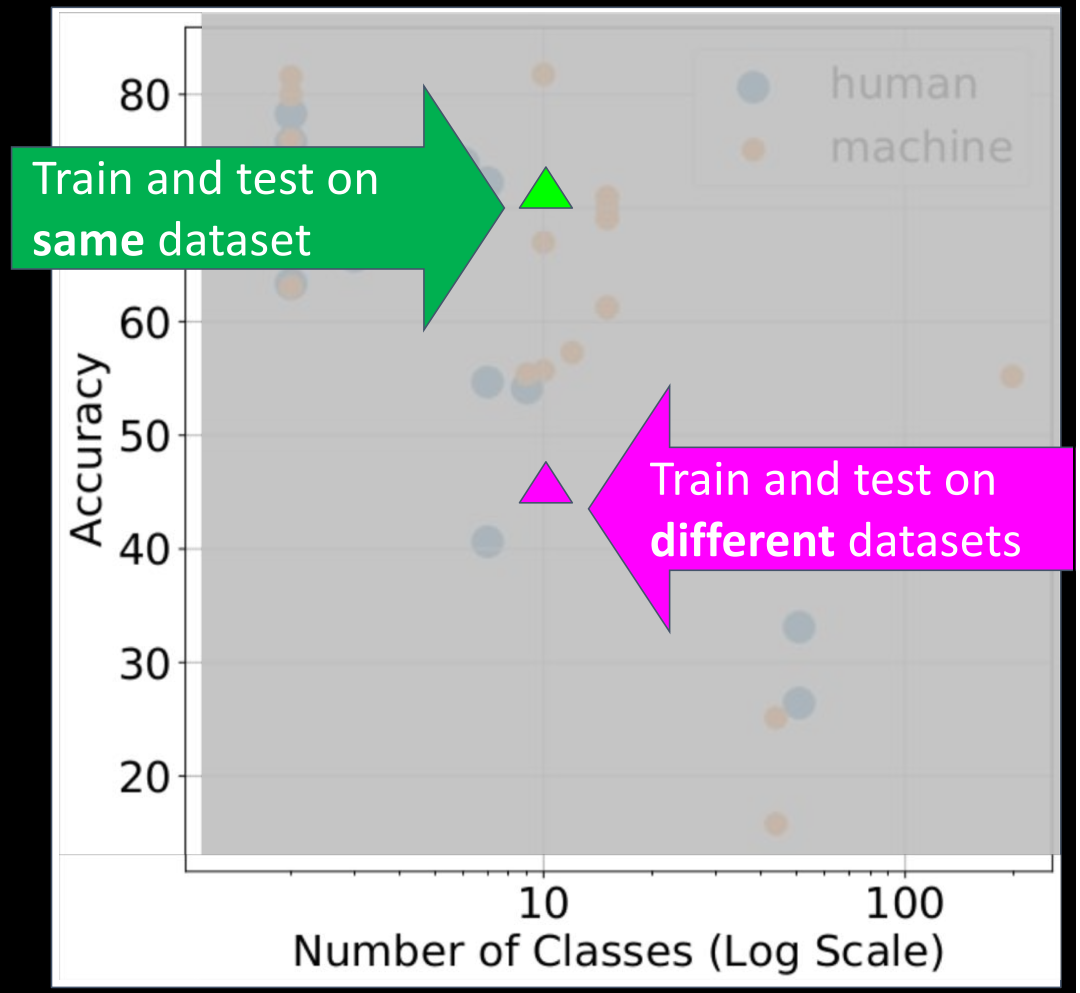

Dataset shift

MICCAI 2019 Yoon, Hamarneh, Garbi

7 Domains:
1 primary: HAM10000
6 secondary: Dermofit+MSK+UDA+ONIC+Derm7pt+PH2
n_s samples/class

CCSA loss: classification & contrastive semantic alignment [Motiian ICCV 2017] CE loss + feature alignment/separation losses

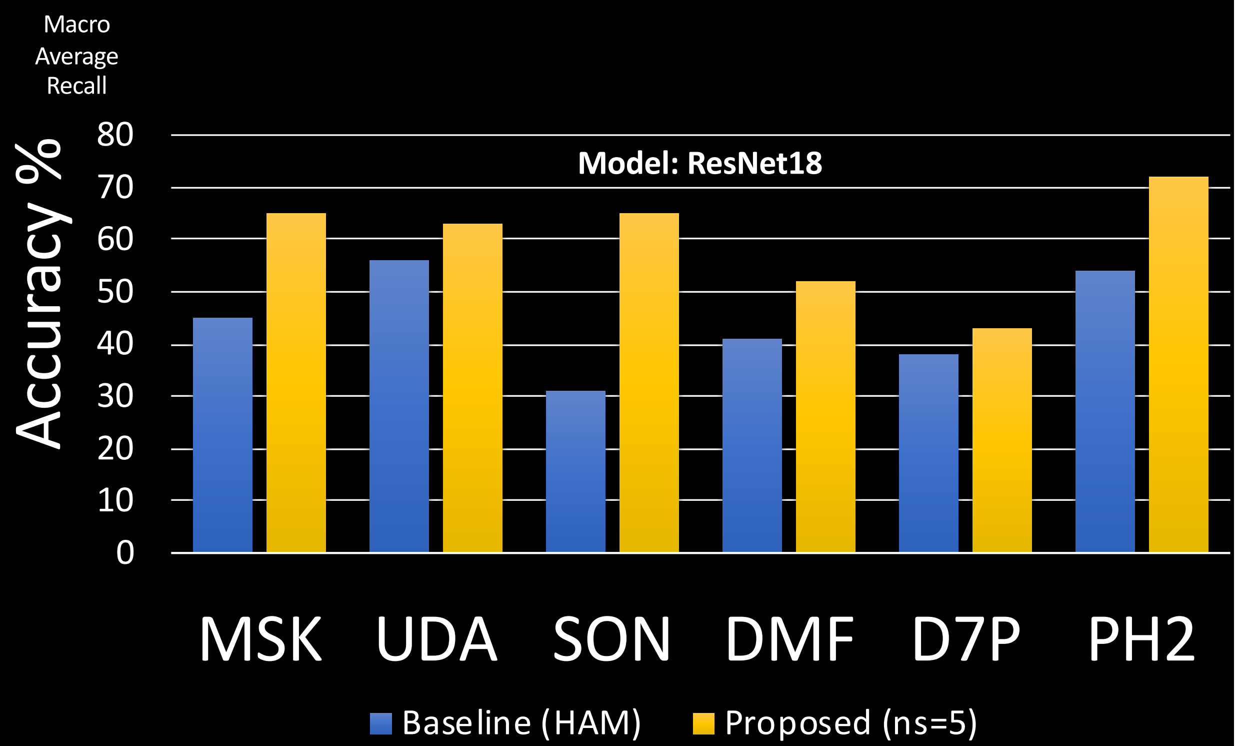

Class imbalance:
- Intra-domain P(nevus) >> P(melanoma)
- Inter-domain dermatofibroma ∉ Domain2

Dynamic sampling

two image-label pairs across domain: $(x_1, y_1), (x_2, y_2)$

Adaptive weighting

of CCSA loss based on $P(y = c_i)$ and $P(y_1 = c_i, y_2 = c_j)$

Interpretability / explainability

Activation and attention maps

Selvaraju, Cogswell, Das, Vedantam, Parikh, Batra. Grad-CAM. ICCV 2017

Melanoma Recognition via Visual Attention
Yan, Kawahara, Hamarneh. IPMI 2019
https://link.springer.com/chapter/10.1007/978-3-030-20351-1_62

Guide (add prior to) the attention maps to ROIs known to discriminatory:

lesion segmentation & dermoscopic features

$$\mathcal{L}_D(\mathcal{A}, \bar{\mathcal{A}}) = 1 - D(\mathcal{A}, \bar{\mathcal{A}})$$

$$\mathcal{L} = \mathcal{L}_{focal} + \lambda_1 \mathcal{L}_D(\mathcal{A}^{(3)}, \bar{\mathcal{A}}^{(3)}) + \lambda_2 \mathcal{L}_D(\mathcal{A}^{(4)}, \bar{\mathcal{A}}^{(4)})$$

Prior

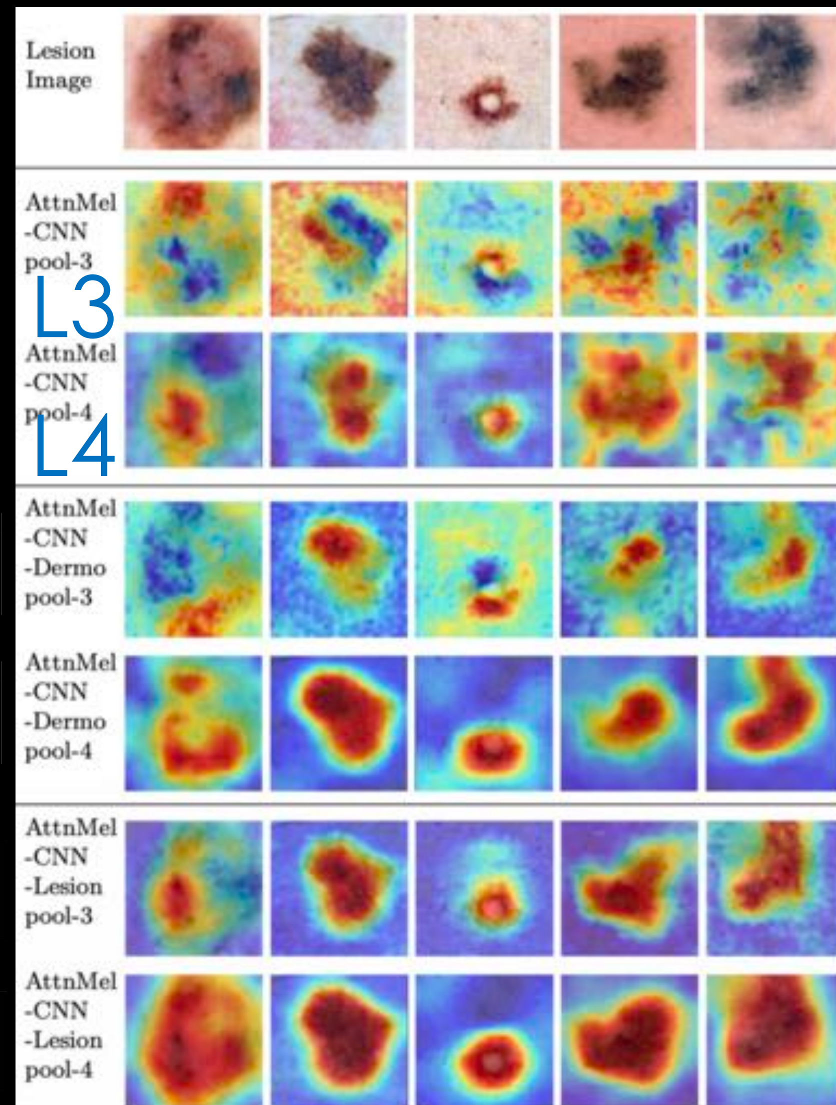

Longitudinal tracking

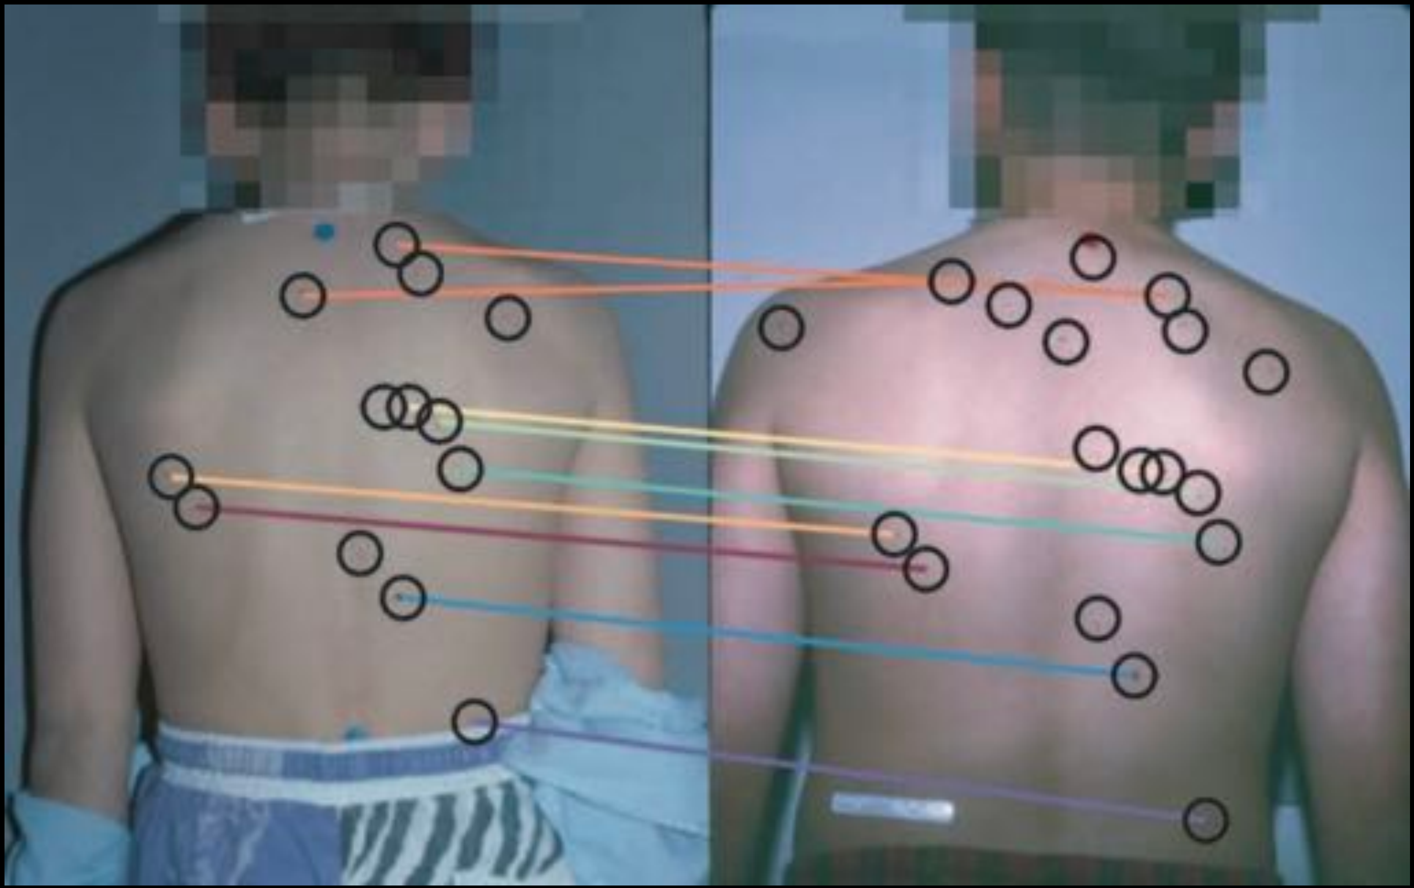

Mirzaalian, Lee, Hamarneh
CVPR 2009, MICCAI 2012, JBHI 2013, MedIA 2015, ISBI2015
https://ieeexplore.ieee.org/abstract/document/5206725 https://ieeexplore.ieee.org/document/6681908/
https://www.sciencedirect.com/science/article/abs/pii/S1361841515000353 https://ieeexplore.ieee.org/document/7164139

Longitudinal tracking

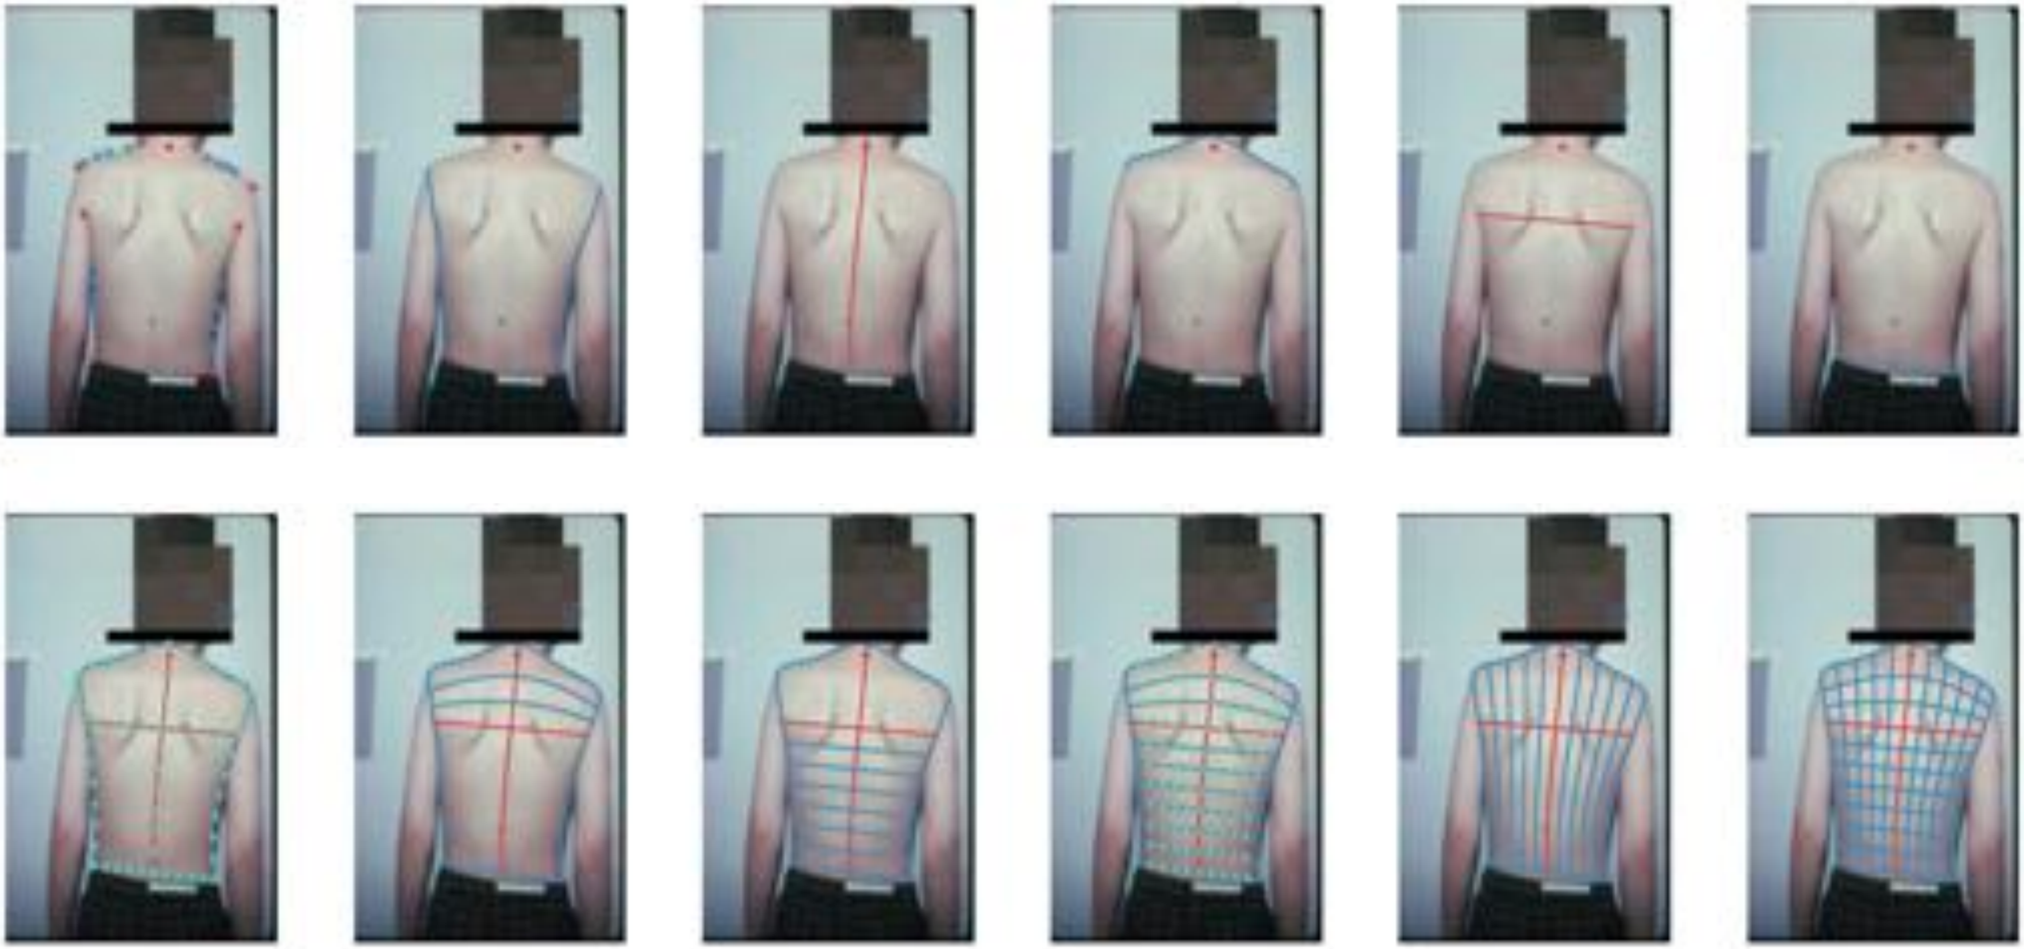

Longitudinal tracking

Visual communication

Beyond reporting predicted class and probabilities

Image Content-Based Navigation of Skin Conditions
Kawahara, Hamarneh. WCD 2015
https://www.cs.sfu.ca/~hamarneh/ecopy/wcd2015a.pdf

Graph Geodesics to Find Progressively Similar Skin Lesion Images
Kawahara, Moriarty, Hamarneh. MICCAI GRAIL 2017
https://link.springer.com/chapter/10.1007/978-3-319-67675-3_4

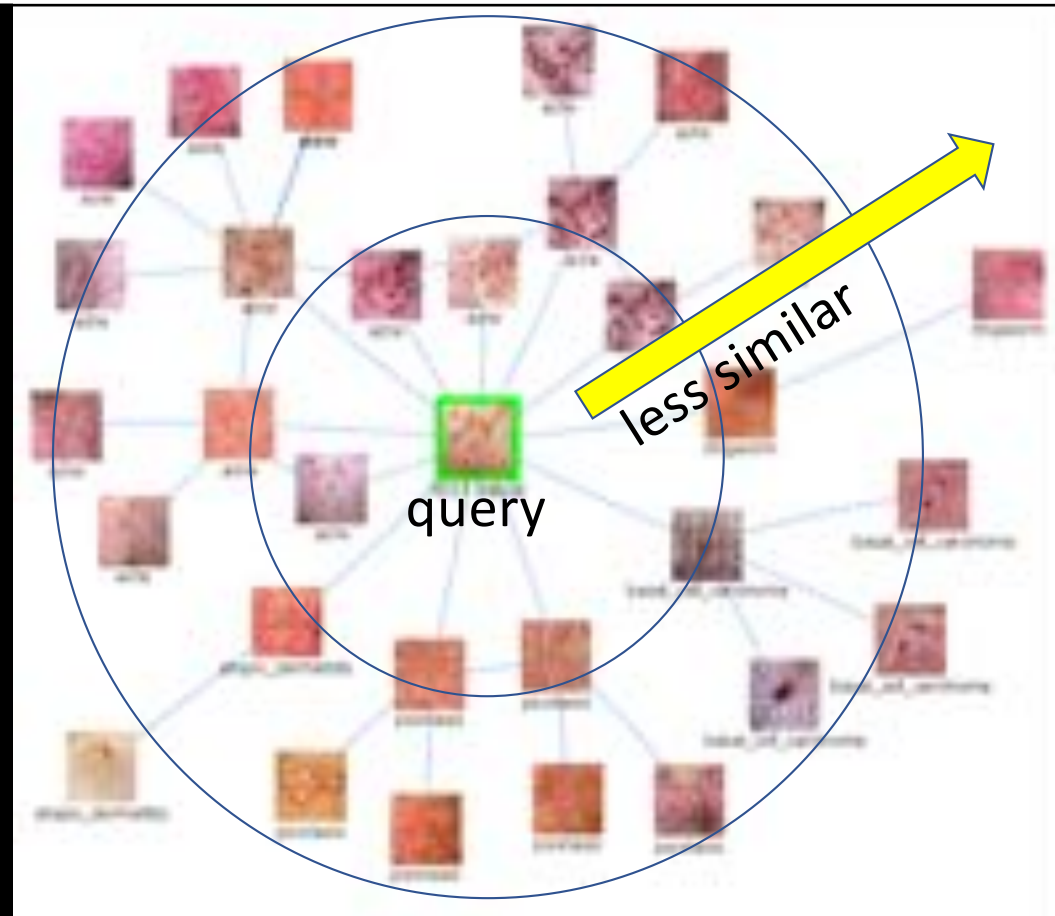

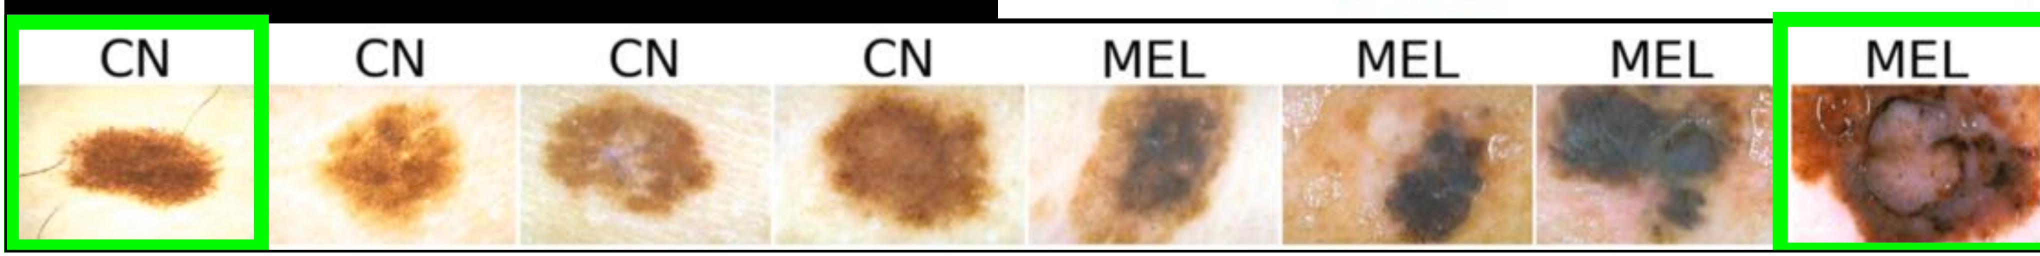

Multi-modal input

Clinical images
Dermoscopic images
Meta-data

$$L(x, y, z; \theta) = \ell(x, y; \theta) + \sum_{j=1}^{7} \ell(x, z_j; \theta)$$

$$\begin{aligned}\mathcal{L}(x_d, x_c, x_m, y, z; \theta) = &\; L((x_d, x_c, x_m), y, z; \theta_{dcm}) \\ &+ L(x_d, y, z; \theta_d) + L((x_d, x_m), y, z; \theta_{dm}) \\ &+ L(x_c, y, z; \theta_c) + L((x_c, x_m), y, z; \theta_{cm})\end{aligned}$$

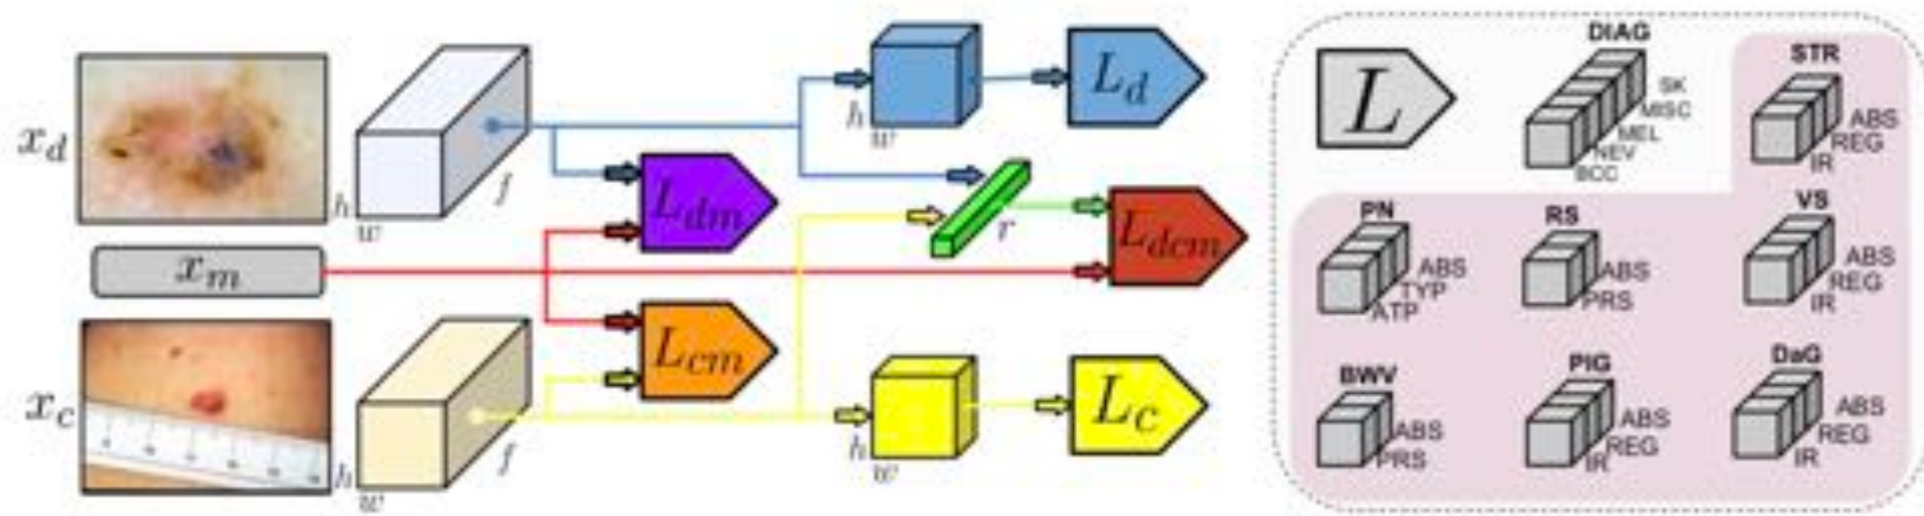

Multi-modal input

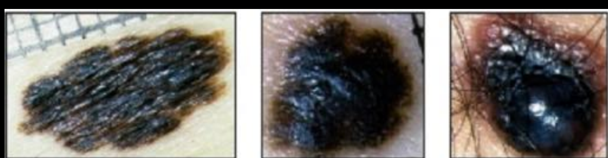

Lesion metadata
body location, roughness / elevation (flat, palpable, nodular)

Patient data: age, gender, race, history

ARE NEURAL NETWORKS EFFECTIVE IN DETECTING MELANOMA USING GENOMIC DATA?

Abder-Rahman Ali,[1] Sally Jane O'Shea,[2,3] Jingpeng Li,[1]
1. Faculty of Natural Sciences, Computing Science and Mathematics, University of Stirling, UK.
2. Dermatology Department, Mater Private Hospital Cork, Ireland.
3. Faculty of Medicine, University College Cork, Ireland.

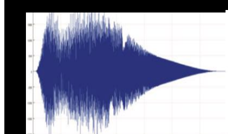

sonification

EBioMedicine
Published by THE LANCET
EBioMedicine 43 (2019) 107–113

Skin cancer detection by deep learning and sound analysis algorithms: A prospective clinical study of an elementary dermoscope

A. Dascalu [a,*], E.O. David [b]
[a] Department of Physiology and Pharmacology, Sackler School of Medicine, Tel Aviv University, Tel Aviv, Israel
[b] Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

European Journal of Cancer 115 (2019) 79–83

Pathologist-level classification of histopathological melanoma images with deep neural networks

Achim Hekler [a], Jochen Sven Utikal [b,c], Alexander H. Enk [d], Carola Berking [e], Joachim Klode [f], Dirk Schadendorf [f], Philipp Jansen [f], Cindy Franklin [g], Tim Holland-Letz [h], Dieter Krahl [i], Christof von Kalle [a], Stefan Fröhling [a], Titus Josef Brinker [a,d,*]

Clinical, dermoscopic, confocal microscopy, optical coherence tomography, histopathology

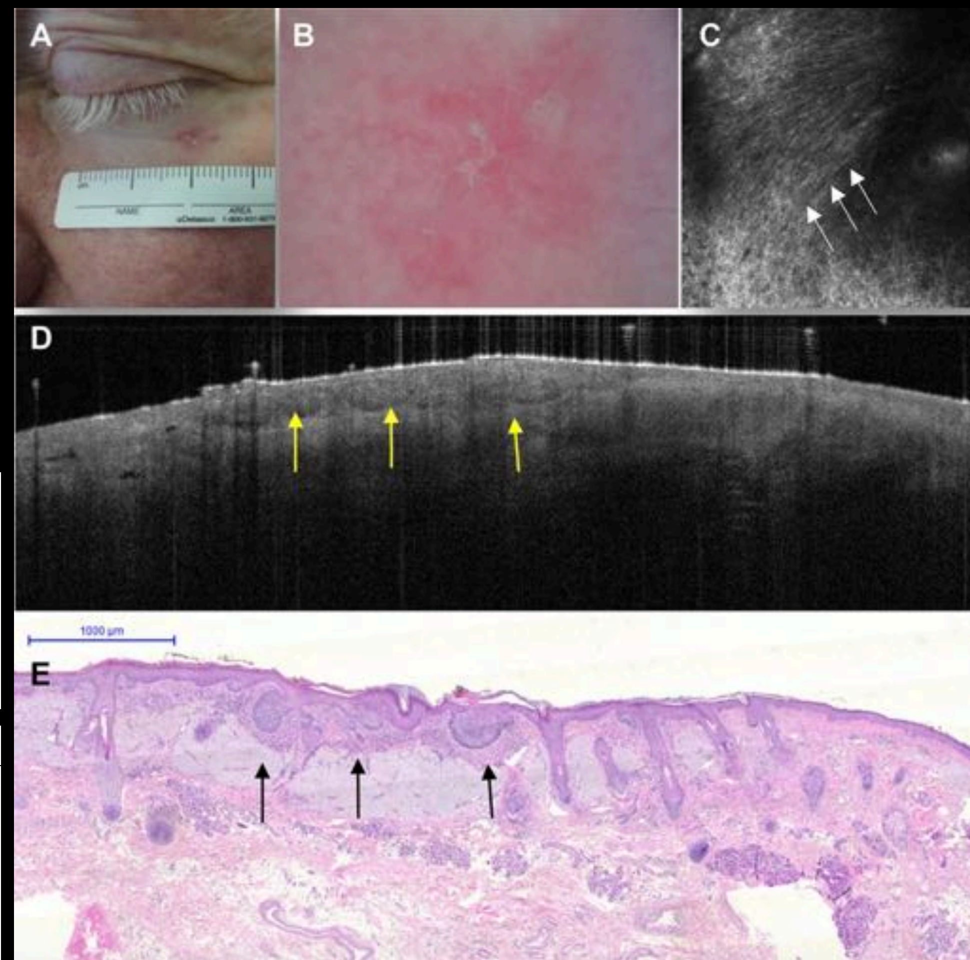

What’s next?

• Disease classes:	<10	→ 1000s
• Datasets:	100s/1000s	→ millions of images
• Training:	full-supervision	→ leveraging weak/no supervision
• Data sources:	homogenous controlled	→ highly heterogenous sources → real-world
• Dimensions:	2D + static	→ 3D + longitudinal/dynamic
• Modalities:	unimodal	→ multi-modal
• Deep modes:	hand-crafted data-driven black-box susceptible	→ automatic → hybrid knowledge- & data-driven models → interpretable → resilient to adversarial attacks
• Beyond technical	communities	→ tighter computational-clinical collaboration
		legal, ethical, societal, economic challenges

Thank you!

hamarneh@sfu.ca