Will Artificial Intelligence Replace Dermatologists?

Susan M. Swetter, MD

Professor of Dermatology

Director, Pigmented Lesion & Melanoma Program

Physician Leader, Cutaneous Oncology Program

Stanford University Medical Center and Cancer Institute

Stanford Artificial Intelligence

Stanford | MEDICINE

2018 AAD Annual Meeting
San Diego, California
February 16-20, 2018
SAN DIEGO CONVENTION CENTER

DISCLOSURE OF RELATIONSHIPS WITH INDUSTRY

Susan M. Swetter, MD

S027 Managing Patients with Melanoma or Other Melanocytic Neoplasms

DISCLOSURES

I do not have any relevant relationships with industry.

How do we know what we know?

- Rules-based systems
- Unconscious pattern matching

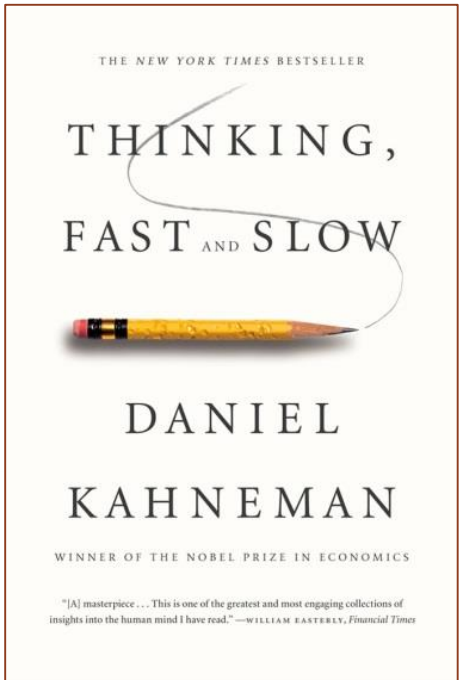

Systems drive the way we think:

System 1 – fast, intuitive, emotional

System 2 – slower, more deliberate and logical

Stanford MEDICINE

Dual-process theory applied to Dermatologists

System 1 (intuitive) – rapid, primarily visual, operates below level of perceptible consciousness. i.e. “gut feeling”

- focus on pattern recognition: “blink, think” and “10-second rule” (*Giuseppe Argenziano, Naples*)

- why “ugly duckling” rule discriminates melanoma more accurately than ABCD clinical warning signs (*Gaudy-Marquest C et al. JAMA Dermatol 2017*)

System 2 (analytical) – deliberate judgment, based on conscious applications of rules acquired through learning

Transition from intuitive to analytical reasoning can hinder clinical reasoning and increase diagnostic error

Stanford Artificial Intelligence

Pelaccia T et al. *Med Educ Online* 2011. Stolper E et al. *BMC Fam Pract* 2009. Norman GR, Eva KW. *Med Educ* 2010.

Stanford Medicine

Skin Cancer Facts

- Skin cancer - most common cancer in the US
- 1 in 5 Americans will develop skin cancer in their lifetime
- Latest estimate (2012): >5.4 million cases BCC/SCC treated in >3.3 million persons in the US (*Rogers HW et al. Arch Dermatol 2015*)
- In 2017 –
  - estimated >91K new cases of invasive melanoma and >87K melanomas in situ in the US
  - nearly 10,000 melanoma-associated deaths
- Survival rate for melanoma is >95% if detected early

Stanford Artificial Intelligence

Siegal RL et al. Cancer Statistics, 2018. *CA Cancer J Clin* 2018

Stanford Medicine

“MELANOMA WRITES ITS MESSAGE IN THE SKIN WITH ITS OWN INK AND FOR ALL OF US TO SEE”

-Neville Davis, Queensland, Australia

…so why is early detection so hard?

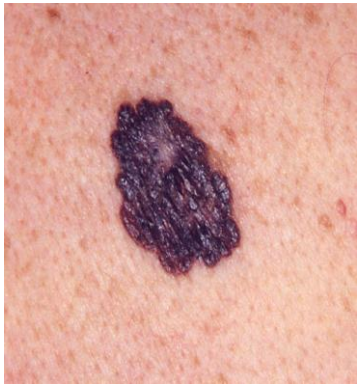

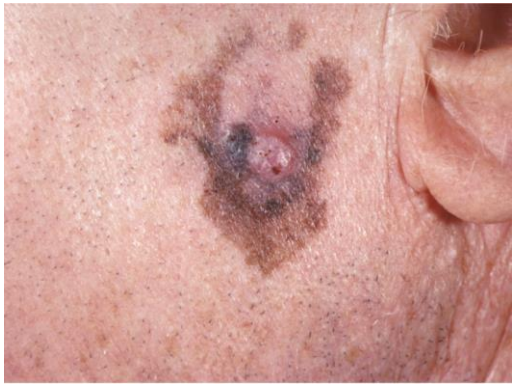

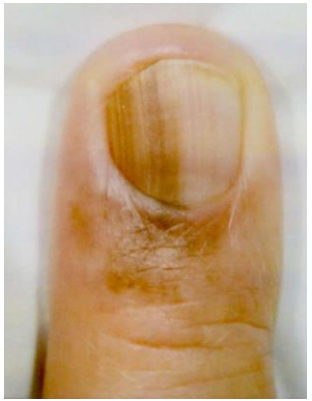

Stanford Artificial Intelligence

Stanford Medicine

Global Health Care Accessibility

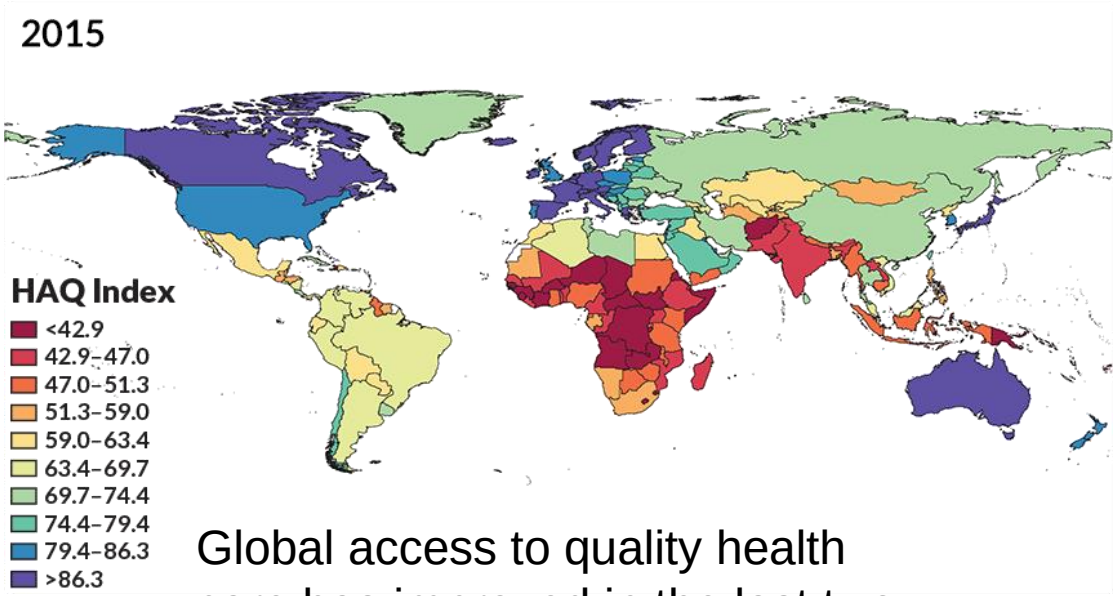

Global access to quality health care has improved in the last two decades, but the gap between countries with the most and least effective treatments has grown

Barber R et al. *The Lancet* 2017

Stanford Artificial Intelligence

Stanford Medicine

How do we democratize health care access?

- Data collection at scale
- Diagnostics at scale
- Diminish health disparities
- **6.3 billion smartphones globally by 2021**

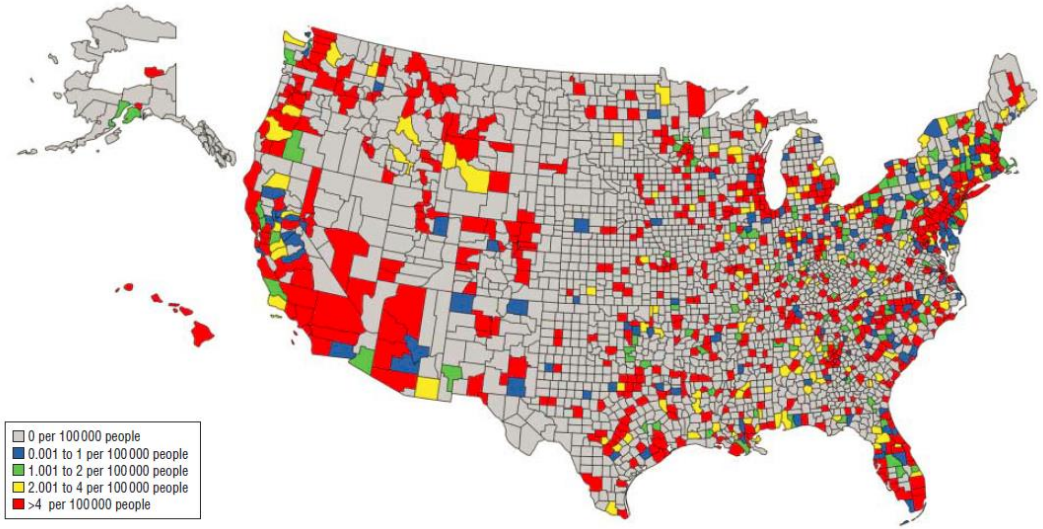

Can we use AI to expand access to dermatologists?

Stanford Artificial Intelligence

Stanford Medicine

Aneja S et al. *Arch Dermatol* 2012

AI is changing our world

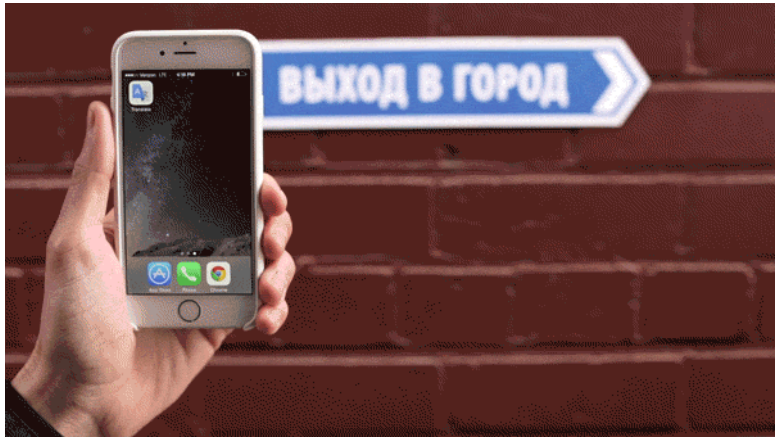

- Driverless cars
- Translation capabilities
- Mortgage lending
- Financial markets

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Neural Networks

- Modeled on human brain/neurons
- Inputs cause neuron to “fire” if over threshold
- AI: inputs fire again and again to get the right output

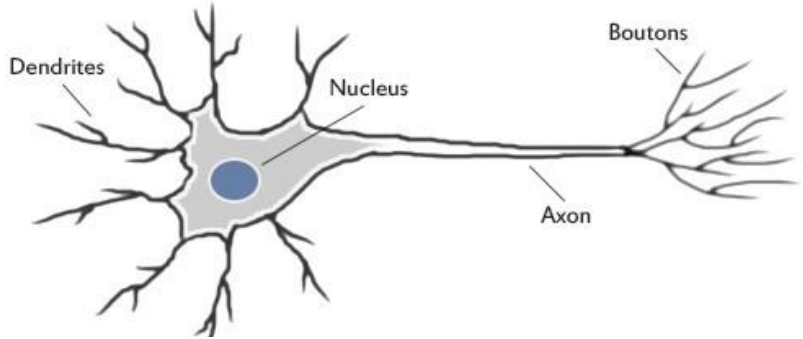

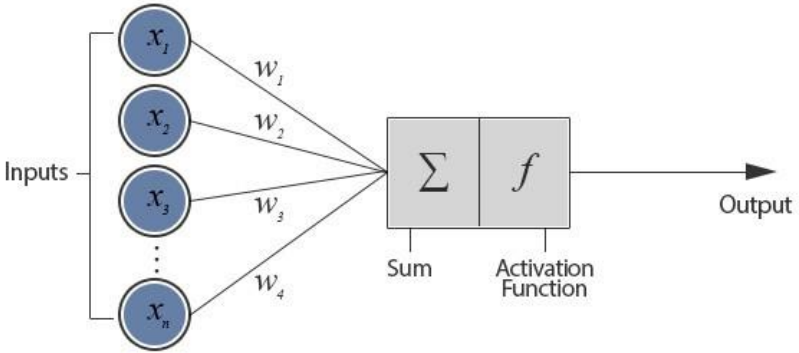

Harsh Pokharna, 2016. Introduction to Neural Networks. Medium.com

Multi-layer Neural Networks

Teach system to preserve the correct output/answer

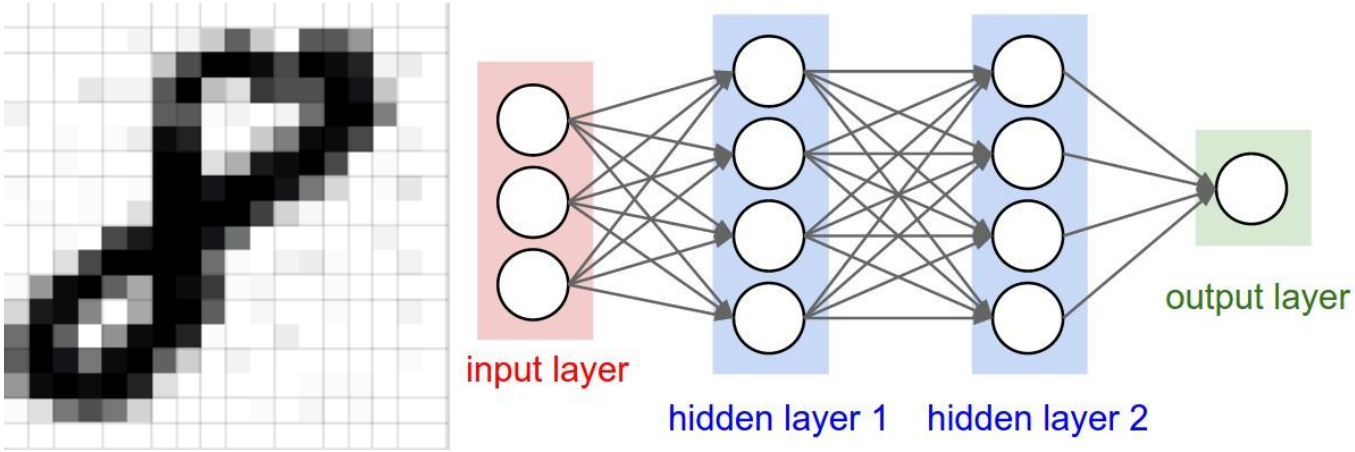

https://medium.com/@ageitgey/machine-learning-is-fun-part-3-deep-learning-and-convolutional-neural-networks-f40359318721

http://cs231n.github.io/convolutional-networks/

Multi-layer Neural Networks

Generate an output

Supervised learning

- Output value is compared with the known value of the image- (“8” dog, cat, etc.)

Backpropagation→ run it back through, adjusting weights in order to reduce error

- Does this across thousands or millions of images

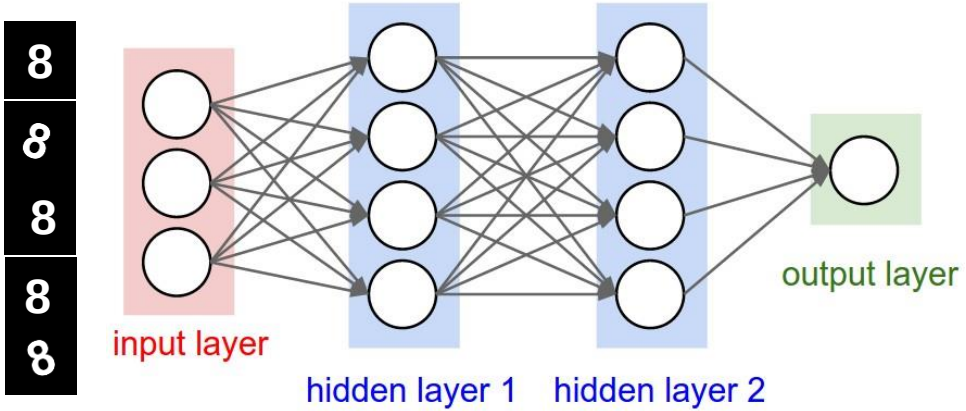

Convolutional Neural Networks

Involves local connections, shared weights, pooling, and many layers

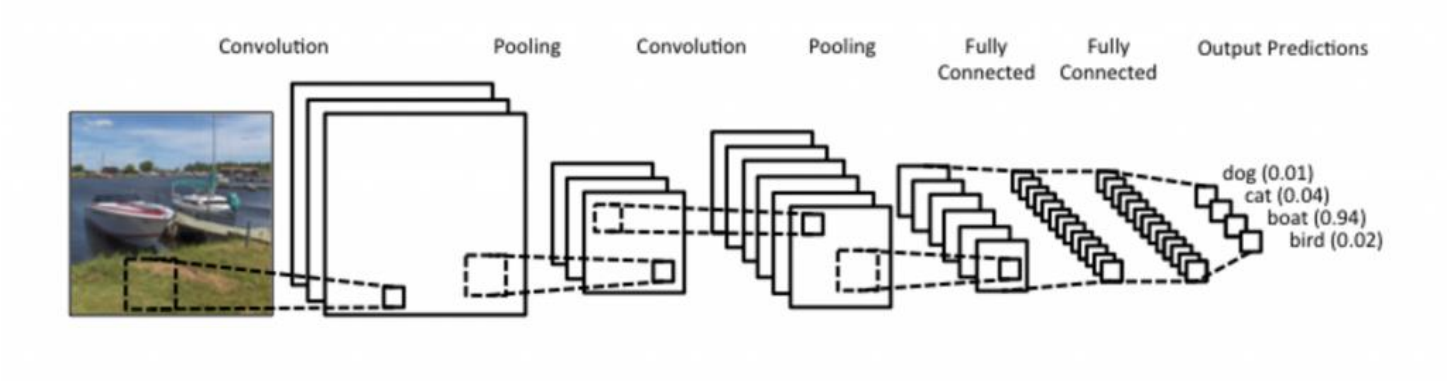

- Capture relationship of pixels to each other using filter as a matrix
- Runs algorithms over and over until errors are minimized across all images
- Allows for profound level of pattern recognition beyond human brain capability

Stanford Artificial Intelligence

https://ujjwalkarn.me/2016/08/11/intuitive-explanation-convnets/

Stanford Medicine

Advances in AI

Breakthroughs in artificial intelligence over past 5 years

Convolutional neural networks+large databases+processing power= **deep learning**

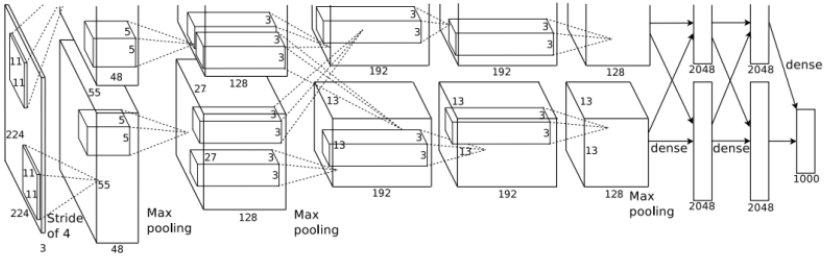

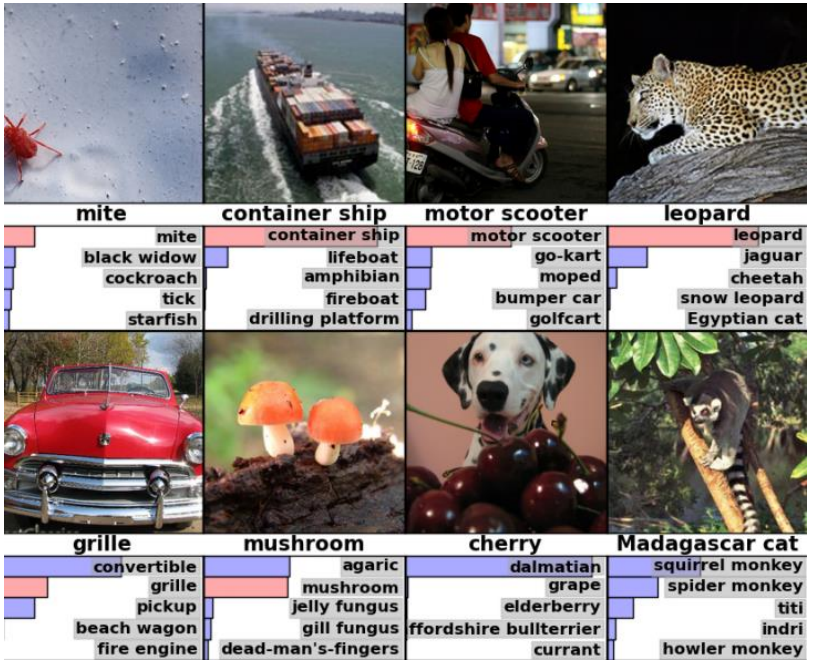

Stanford Artificial Intelligence

Krizhevsky, Sutskever, Hinton et al. ImageNet Classification with Deep Convolutional Neural Networks, NIPS 2012

Algorithms train themselves to recognize what is important and what is not, without human intervention

Fig. 1a: Malamute

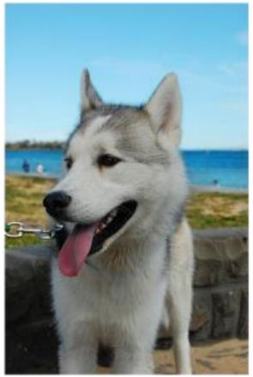

Fig. 1b: Eskimo Dog

Fig. 1c: Husky

- ImageNet: large, deep CNNs trained to classify 1.4 mm high-resolution images into 1000 different classes with low error rates
- If given sufficient examples of huskies, chihuahuas, basal cell carcinomas or melanomas, algorithms learn the relevant patterns of these categories

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Genesis of Stanford AI/Dermatology Interface

> “If AI can differentiate between hundreds of dog breeds, it could make a great contribution to dermatology.”

Stanford Dermatologist/ Dermatopathologist Roberto Novoa, MD
January 27, 2015

Stanford ARTIFICIAL INTELLIGENCE

Stanford MEDICINE

Stanford AI in Dermatology Team

Andre Esteva, MS

Brett Kuprel, MS

Sebastian Thrun, PhD (Udacity)

Helen Blau, PhD

Roberto Novoa, MD

Justin Ko, MD, MBA

Susan Swetter, MD

Clinical and Dermoscopic Images Spanning Breadth of Disease

Constructing the dataset:

Nearly 130,000 images:
- clinician-labeled and/or biopsy-proven from 18 different, open-access online repositories

- clinical data from Stanford Dermatology clinics

Biopsy proven images:
- University of Edinburgh library
- ISDIS ISBI challenge (dermoscopic)

Images comprised more than 2,000 dermatologic diseases

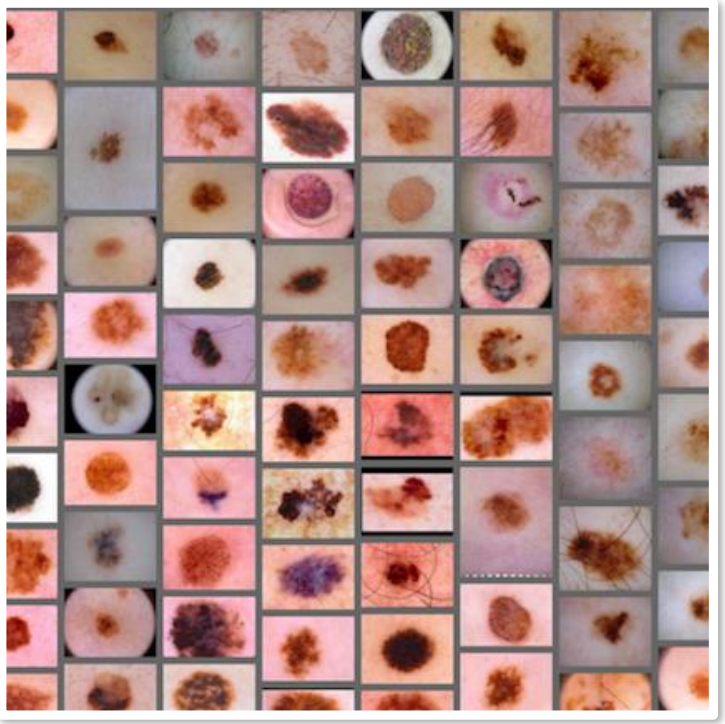

Stanford Artificial Intelligence

Stanford Medicine

Visual Taxonomy

129K images
2K diseases

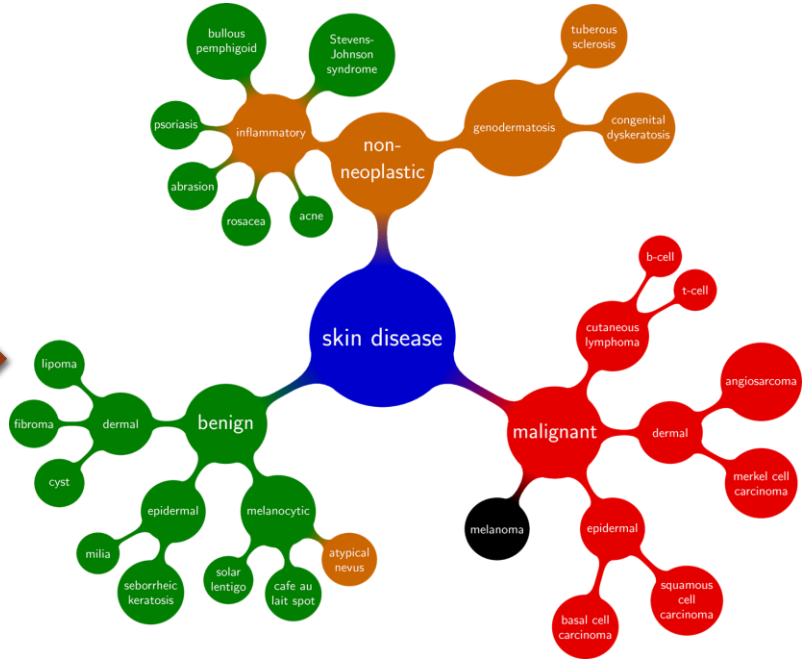

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Our Objectives

Evaluate performance of deep learning algorithms (namely, a single CNN) on classification of cutaneous malignancies

Compare to dermatologists

2 main questions:

- Lesion benign or malignant?
- Biopsy/treat or reassure?

Stanford Artificial Intelligence

Stanford Medicine

Methods

Recruited 21 board-certified dermatologists

- Stanford University
- University of Pennsylvania
- Massachusetts General Hospital/Harvard
- University of Iowa

Three 100+ question tests

a. Epidermal tumors (BCC/SCC vs SKs)
b. Clinical images of melanocytic lesions (melanoma vs benign nevi, excluded SKs)
c. Dermoscopic images of melanocytic lesions (separate dataset)

Stanford ARTIFICIAL INTELLIGENCE

Stanford MEDICINE

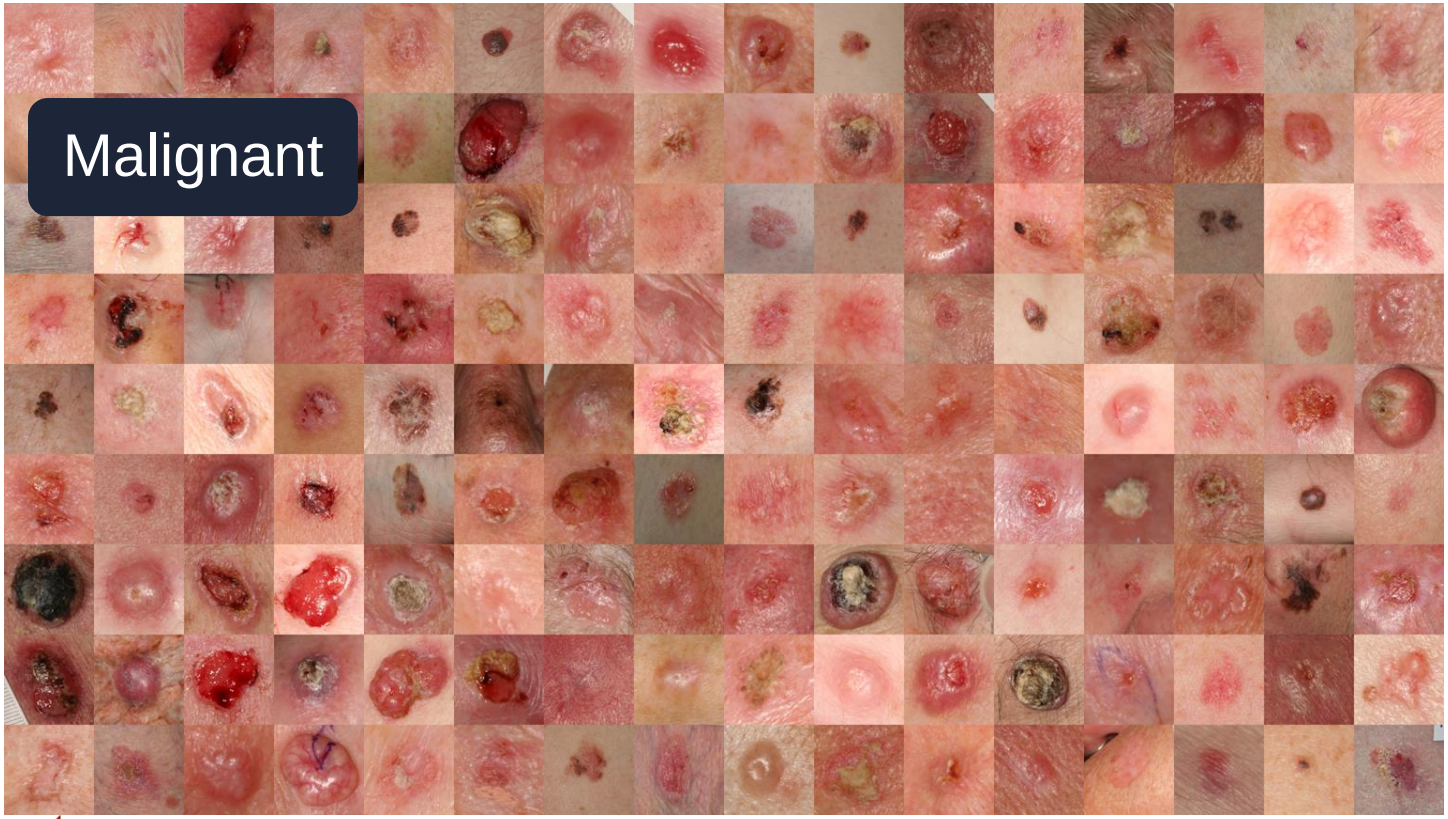

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

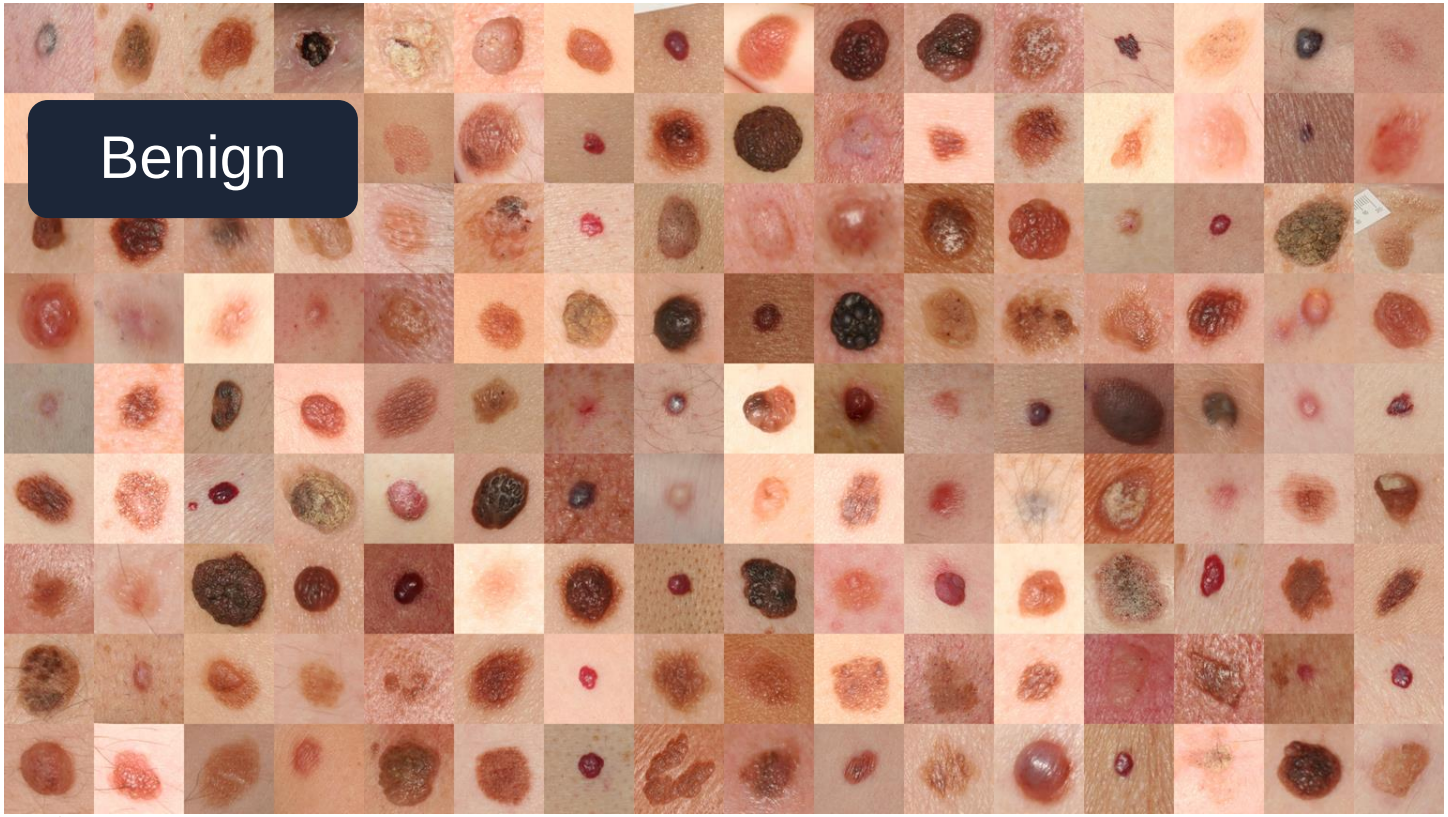

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Training

Goal: to build one system that could accommodate significant variation inherent in photographic images—lighting, zoom, angle, etc. with no pre-processing or lesion segmentation

Essentially to be able to feed any captured skin image directly into the system and output a classification

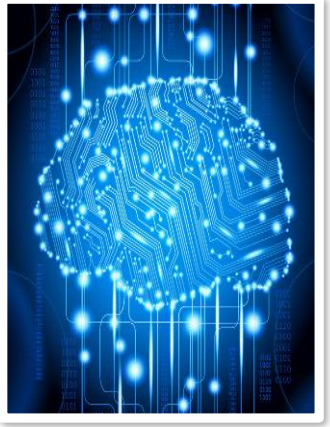

Stanford Artificial Intelligence

Stanford Medicine

Deep Convolutional Neural Network (CNN)

Skin lesion image

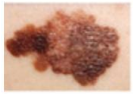

GoogleNet Inception v3 CNN architecture pretrained on ImageNet;
- the net on our dataset was then trained and fine-tuned using transfer learning;
- results in probability distribution over clinical classes of skin
- training classes defined by applying a partitioning algorithm to our taxonomy

Stanford Artificial Intelligence

Stanford Medicine

Sample Test Images

Utilized a new dermatologist-labeled dataset of >129K clinical images, including >3300 separate dermoscopic images

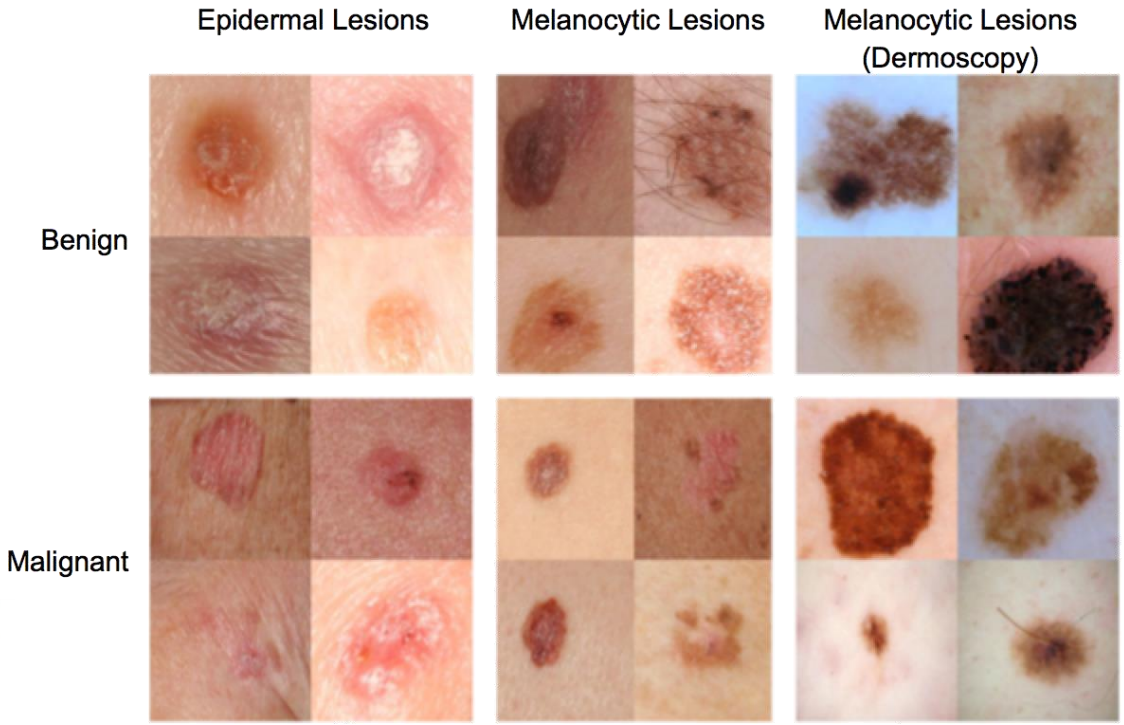

Stanford Artificial Intelligence

Stanford Medicine

Skin cancer: through the machine's eye

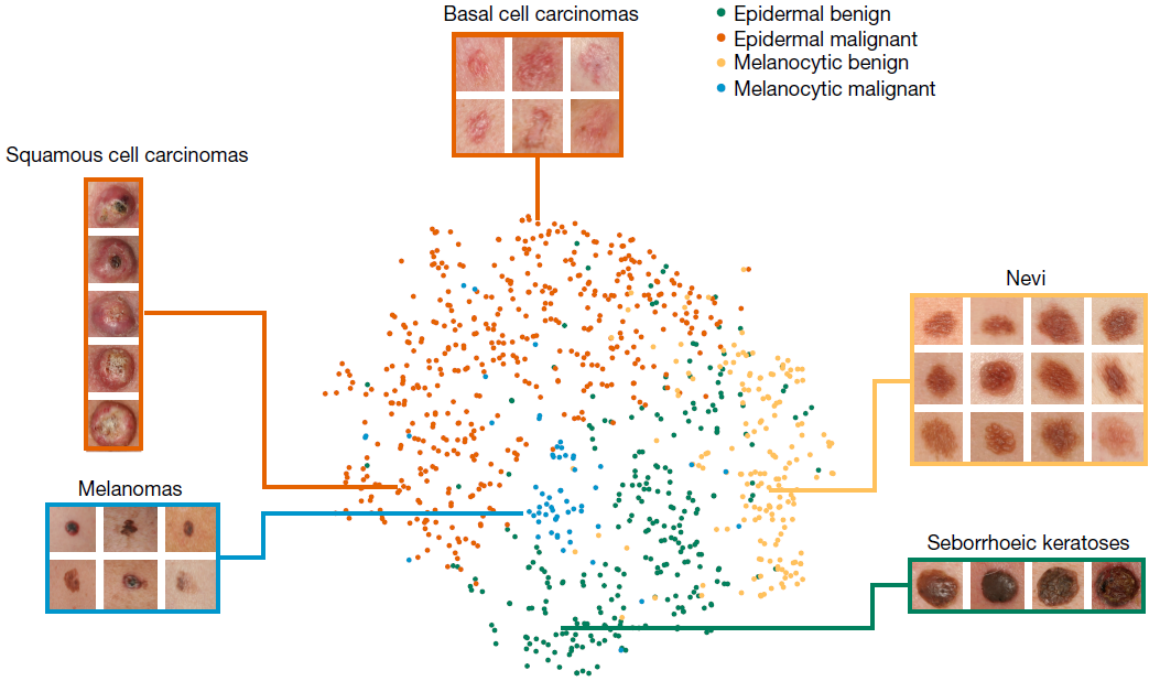

Stanford Artificial Intelligence

Stanford Medicine

Algorithm ready for testing within 1 year

CNN outputs a malignancy probability for each image

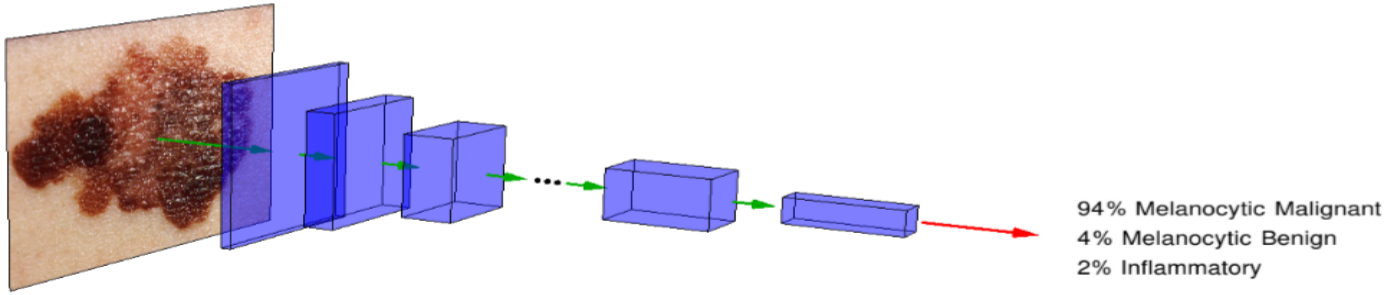

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Evaluation

Algorithm validated in 2 ways using:
i) 3-class disease partition of 1st level nodes in our taxonomy (non-neoplastic, benign neoplastic, and malignant neoplastic) – with 72% overall accuracy
ii) class disease partition with 2nd level nodes – with 55% overall accuracy

Then only biopsy-labeled images used to conclusively validate the algorithm using the same CNN for all 3 tasks:

- -keratinocyte carcinoma vs SK
- -melanoma vs nevus
- -melanoma vs nevus on dermoscopy

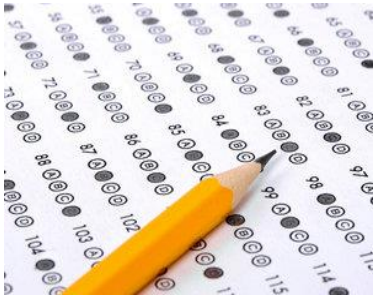

Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017 2;542:115-118.

Stanford MEDICINE

Results for CNN

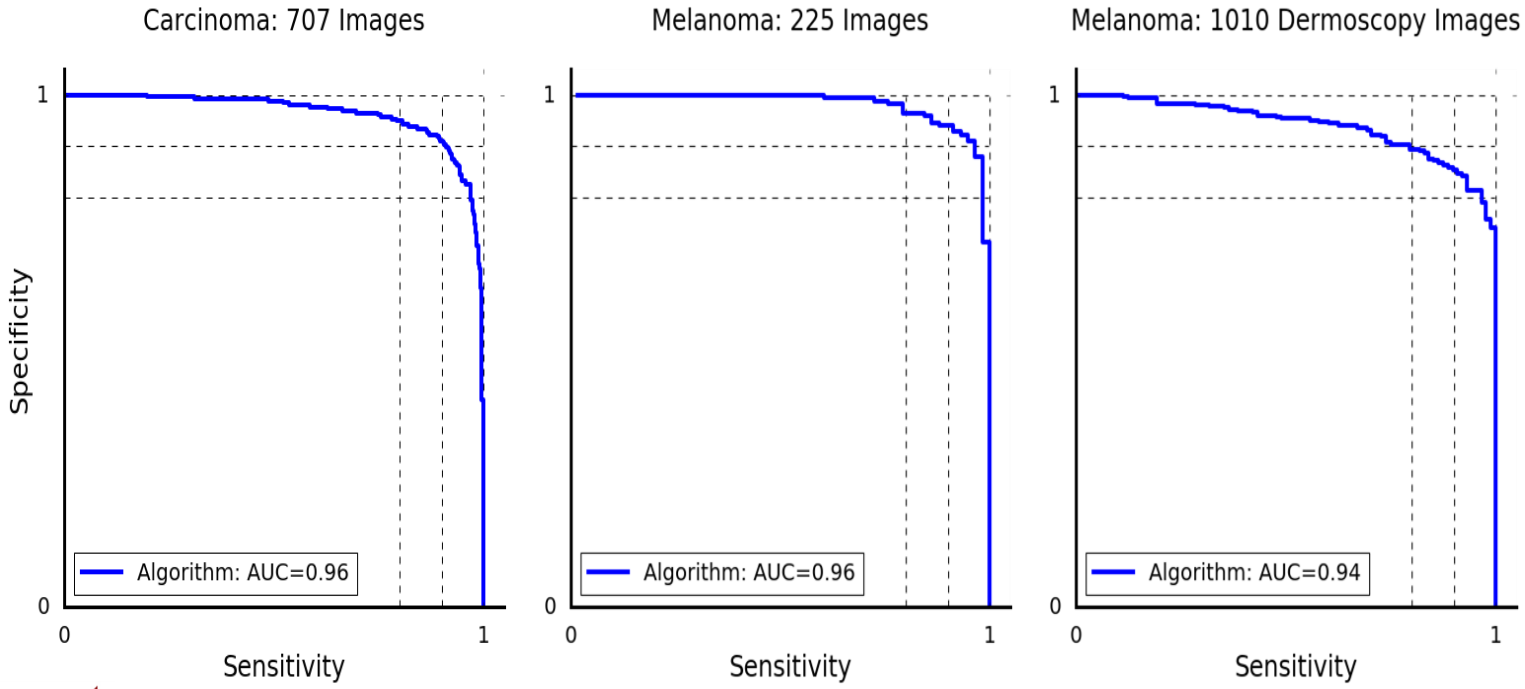

Stanford Artificial Intelligence

Esteva A, Kuprel B, Novoa RA, Ko J, Swetter SM, Blau HM, Thrun S. Dermatologist-level classification of skin cancer with deep neural networks. *Nature*. 2017 2;542:115-118.

Stanford MEDICINE

Results for CNN vs Dermatologists

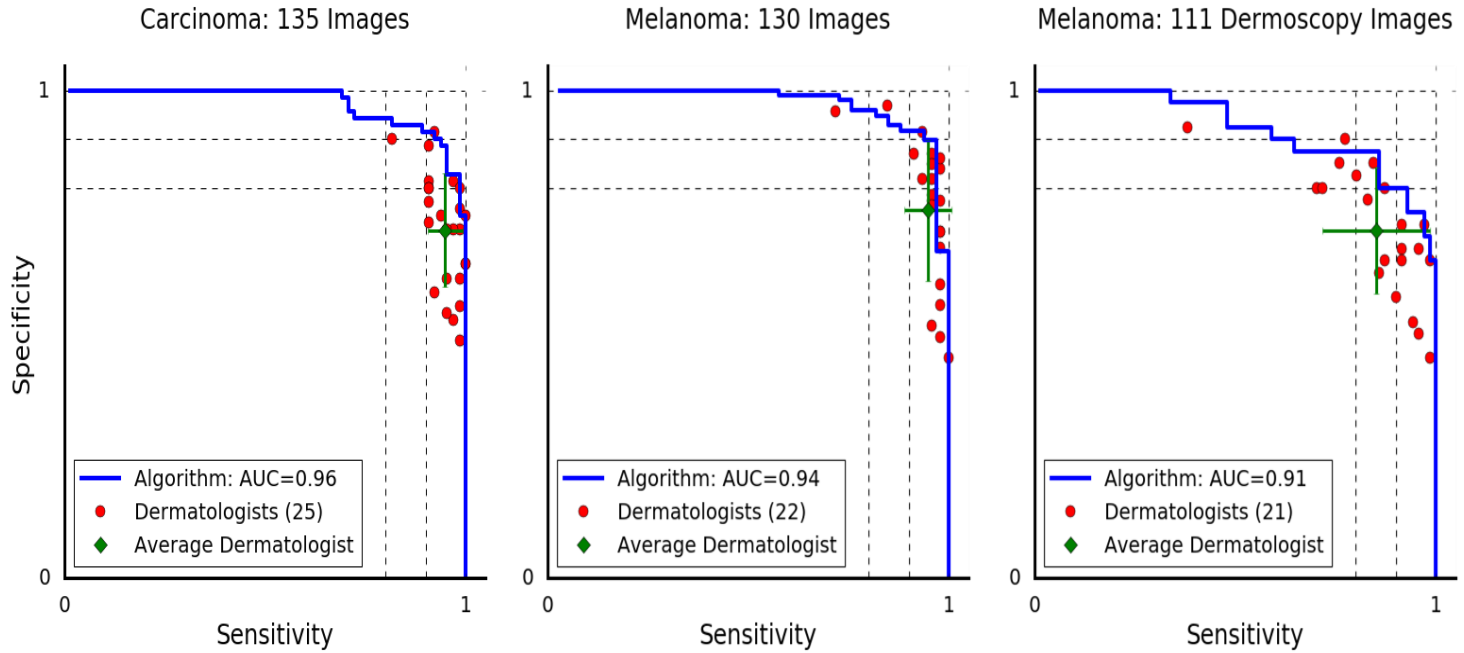

CNN performed at least as well as dermatologists as a whole

(Limitation of dermoscopy test as most dermatologists were not experts in pigmented lesions or dermoscopy)

Stanford Artificial Intelligence

Stanford Medicine

Esteva A et a. *Nature*. 2017 2;542:115-118.

LETTER

doi:10.1038/nature21056

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva[1]*, Brett Kuprel[1]*, Roberto A. Novoa[2,3], Justin Ko[2], Susan M. Swetter[2,4], Helen M. Blau[5] & Sebastian Thrun[6]

Skin cancer, the most common human malignancy[1–3], is primarily diagnosed visually, beginning with an initial clinical screening and followed potentially by dermoscopic analysis, a biopsy and histopathological examination. Automated classification of skin lesions using images is a challenging task owing to the fine-grained variability in the appearance of skin lesions. Deep convolutional neural networks (CNNs)[4,5] show potential for general and highly variable tasks across many fine-grained object categories[6–11]. Here we demonstrate classification of skin lesions using a single CNN, trained end-to-end from images directly, using only pixels and disease labels as inputs. We train a CNN using a dataset of 129,450 clinical images—two orders of magnitude larger than previous datasets[12]—consisting of 2,032 different diseases. We test its performance against 21 board-certified dermatologists on biopsy-proven clinical images with two critical binary classification use cases: keratinocyte carcinomas versus benign seborrheic keratoses; and malignant melanomas versus benign nevi. The first case represents the identification of the most common cancers, the second represents the identification of the deadliest skin cancer. The CNN achieves performance on par with all tested experts across both tasks, demonstrating an artificial intelligence capable of classifying skin cancer with a level of competence comparable to dermatologists. Outfitted with deep neural networks, mobile devices can potentially extend the reach of dermatologists outside of the clinic. It is projected that 6.3 billion smartphone subscriptions will exist by the year 2021 (ref. 13) and can therefore potentially provide low-cost universal access to vital diagnostic care.

There are 5.4 million new cases of skin cancer in the United States[2] every year. One in five Americans will be diagnosed with a cutaneous malignancy in their lifetime. Although melanomas represent fewer than 5% of all skin cancers in the United States, they account for approximately 75% of all skin-cancer-related deaths, and are responsible for over 10,000 deaths annually in the United States alone. Early detection is critical, as the estimated 5-year survival rate for melanoma drops from over 99% if detected in its earliest stages to about 14% if detected in its latest stages. We developed a computational method which may allow medical practitioners and patients to proactively track skin lesions and detect cancer earlier. By creating a novel disease taxonomy, and a disease-partitioning algorithm that maps individual diseases into training classes, we are able to build a deep learning system for automated dermatology.

Previous work in dermatological computer-aided classification[12,14,15] has lacked the generalization capability of medical practitioners owing to insufficient data and a focus on standardized tasks such as dermoscopy[16–18] and histological image classification[19–22]. Dermoscopy images are acquired via a specialized instrument and histological images are acquired via invasive biopsy and microscopy; whereby both modalities yield highly standardized images. Photographic images (for example, smartphone images) exhibit variability in factors such as zoom, angle and lighting, making classification substantially more challenging[23,24]. We overcome this challenge by using a data-driven approach—1.41 million pre-training and training images make classification robust to photographic variability. Many previous techniques require extensive preprocessing, lesion segmentation and extraction of domain-specific visual features before classification. By contrast, our system requires no hand-crafted features; it is trained end-to-end directly from image labels and raw pixels, with a single network for both photographic and dermoscopic images. The existing body of work uses small datasets of typically less than a thousand images of skin lesions[16,18,19], which, as a result, do not generalize well to new images. We demonstrate generalizable classification with a new dermatologist-labelled dataset of 129,450 clinical images, including 3,374 dermoscopy images.

Deep learning algorithms, powered by advances in computation and very large datasets[25], have recently been shown to exceed human performance in visual tasks such as playing Atari games[26], strategic board games like Go[27] and object recognition[6]. In this paper we outline the development of a CNN that matches the performance of dermatologists at three key diagnostic tasks: melanoma classification, melanoma classification using dermoscopy and carcinoma classification. We restrict the comparisons to image-based classification.

We utilize a GoogleNet Inception v3 CNN architecture[9] that was pre-trained on approximately 1.28 million images (1,000 object categories) from the 2014 ImageNet Large Scale Visual Recognition Challenge[6], and train it on our dataset using transfer learning[28]. Figure 1 shows the working system. The CNN is trained using 757 disease classes. Our dataset is composed of dermatologist-labelled images organized in a tree-structured taxonomy of 2,032 diseases, in which the individual diseases form the leaf nodes. The images come from 18 different clinician-curated, open-access online repositories, as well as from clinical data from Stanford University Medical Center. Figure 2a shows a subset of the full taxonomy, which has been organized clinically and visually by medical experts. We split our dataset into 127,463 training and validation images and 1,942 biopsy-labelled test images.

To take advantage of fine-grained information contained within the taxonomy structure, we develop an algorithm (Extended Data Table 1) to partition diseases into fine-grained training classes (for example, amelanotic melanoma and acrolentiginous melanoma). During inference, the CNN outputs a probability distribution over these fine classes. To recover the probabilities for coarser-level classes of interest (for example, melanoma) we sum the probabilities of their descendants (see Methods and Extended Data Fig. 1 for more details).

We validate the effectiveness of the algorithm in two ways, using nine-fold cross-validation. First, we validate the algorithm using a three-class disease partition—the first-level nodes of the taxonomy, which represent benign lesions, malignant lesions and non-neoplastic

[1]Department of Electrical Engineering, Stanford University, Stanford, California, USA. [2]Department of Dermatology, Stanford University, Stanford, California, USA. [3]Department of Pathology, Stanford University, Stanford, California, USA. [4]Dermatology Service, Veterans Affairs Palo Alto Health Care System, Palo Alto, California, USA. [5]Baxter Laboratory for Stem Cell Biology, Department of Microbiology and Immunology, Institute for Stem Cell Biology and Regenerative Medicine, Stanford University, Stanford, California, USA. [6]Department of Computer Science, Stanford University, Stanford, California, USA.
*These authors contributed equally to this work.

Limitations of our study

- Retrospective
- Potential spectrum bias
- Study design→ 2 disease categories
- Differences between in-person exam and telederm
  - › We “blink, think, and compare” and use dermoscopy to help with diagnosis
- Transparency on pathology labels/ dermatopathology accuracy/ wide variability for melanocytic neoplasms (*Elmore JG et al. BMJ 2017*)
- Opportunities for bias/confounding
  - › Need to ensure extensive representation by varied skin types
  - › Extensive additional work is needed
    - Prospective clinical validation studies

Stanford ARTIFICIAL INTELLIGENCE

Stanford MEDICINE

Limitations of Deep Learning

- System is opaque; we don't know why it calls an image benign or malignant
- Dots, rulers, marks, etc. may introduce bias into dataset
- Investigators may not be aware of them or extent
- Impact on MD cognition/learning?

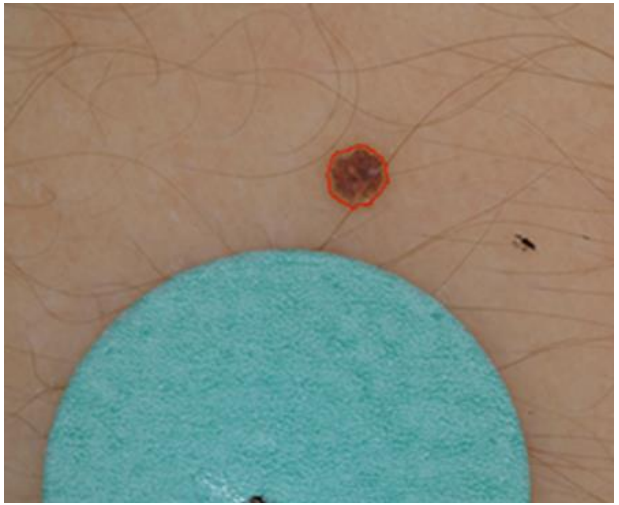

Stanford Artificial Intelligence

Stanford Medicine

Strengths

Computer can assess image data imperceptible to human eye

- Only looking at specific lesions, not whole-body or sequential change detection

Broadly applicable across disciplines

- Algorithms will improve with transfer learning
- Edges, shapes remain important
- Show a skin cancer app cat pictures and it will improve at classifying skin cancer

Capable of running on a smartphone device

- App created to study this, prospective “real-time” validation in clinical setting

Stanford
ARTIFICIAL
INTELLIGENCE

Further directions

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

What the network is “looking at”

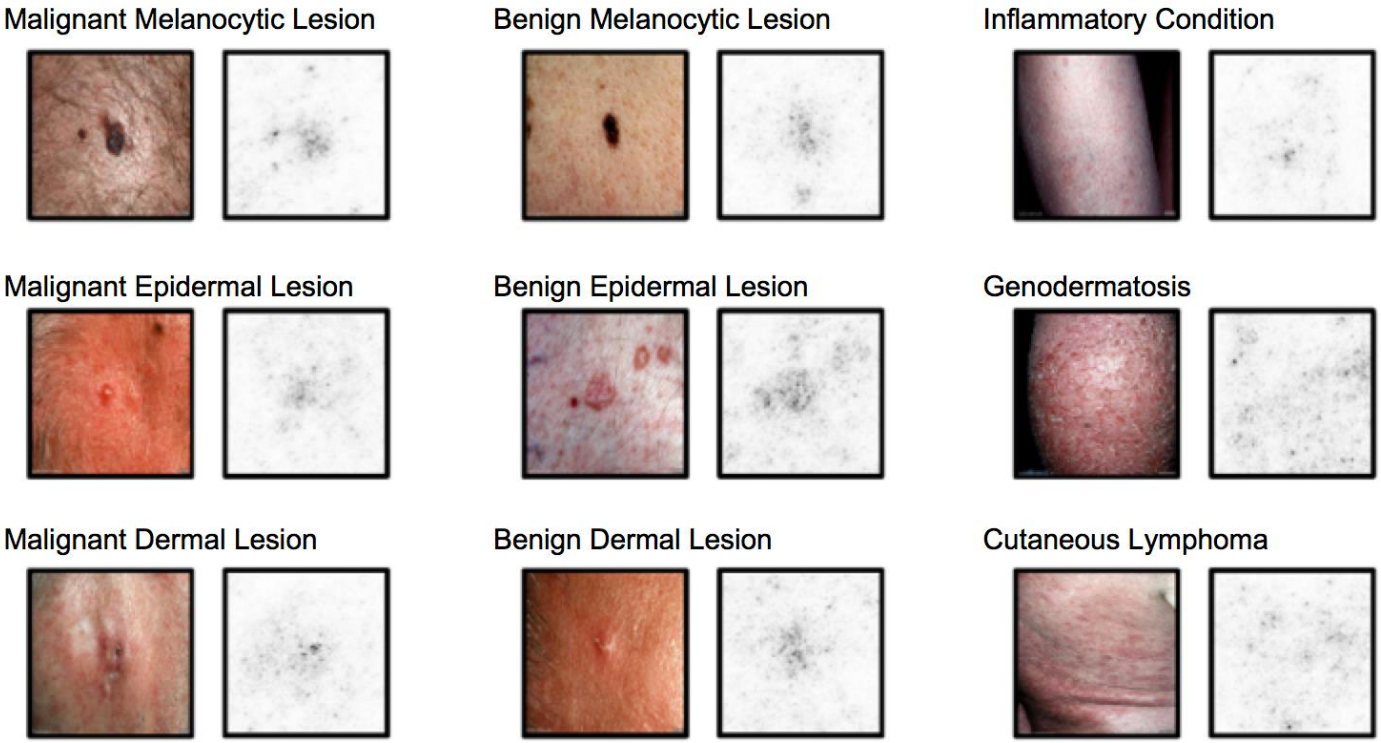

Saliency maps for other skin conditions

Stanford Artificial Intelligence

Stanford Medicine

Can AI be used to quantify and monitor skin disease severity?

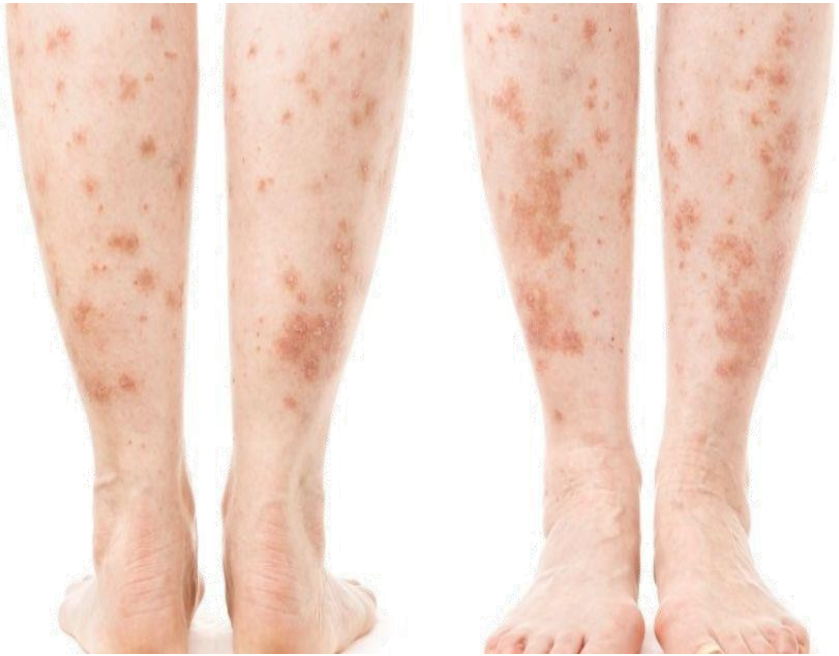

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Context matters

A dermatologist's clinical impression is based on contextual factors beyond visual and dermoscopic examination of a lesion in isolation

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Are we comfortable with AI as black box?

A trained neural net does not necessarily mimic the decision-making approach of humans--rather it identifies its own criteria for informative patterns associated with a disease

Stanford Artificial Intelligence

Decisions need to be made based on machine output of probability, and the pt needs to managed

Stanford Medicine

Is AI a Pandora's box?

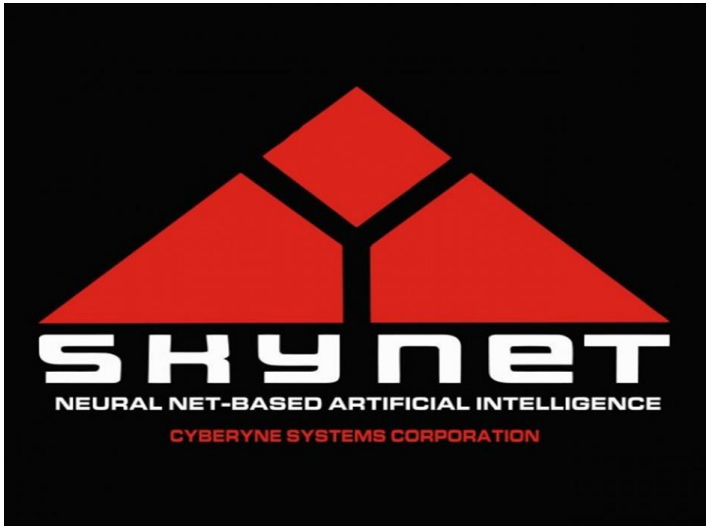

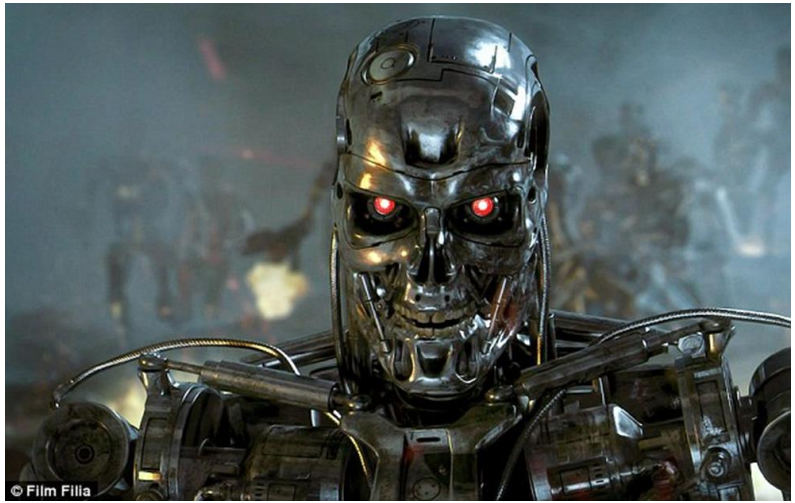

- **Probably not, and it won't replace dermatologists!**
- Roles in health care and education
- Real vs imaginary dangers:
- "AI is a fundamental existential risk for human civilization." *(Elon Musk, 2017)*
- *"*AI software will help us understand biology, understand how to intervene and improve lives very dramatically." *(Bill Gates, 2018)*
- **Regulation and liability need to be addressed proactively**

Stanford ARTIFICIAL INTELLIGENCE

Stanford MEDICINE

Clinical validation of app in progress

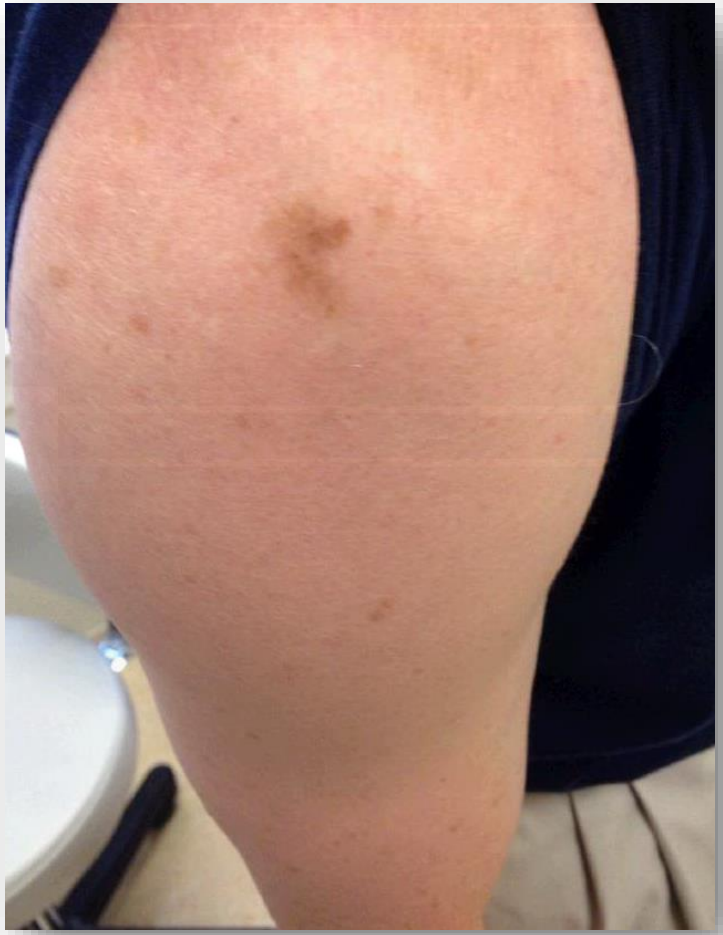

Stanford
ARTIFICIAL
INTELLIGENCE

Stanford
MEDICINE

Can AI help dermatologists?

-35 Y female with atypical mole syndrome, seven cutaneous melanomas (3 melanomas in situ, 3 T1a invasive melanomas and 1 T2a SLN-positive melanoma) and eight severely dysplastic nevi diagnosed since 2016 (negative for p16 mutation)

-5 additional biopsies in 6/2017 showed 2 severely dysplastic nevi and an atypical compound Spitz tumor

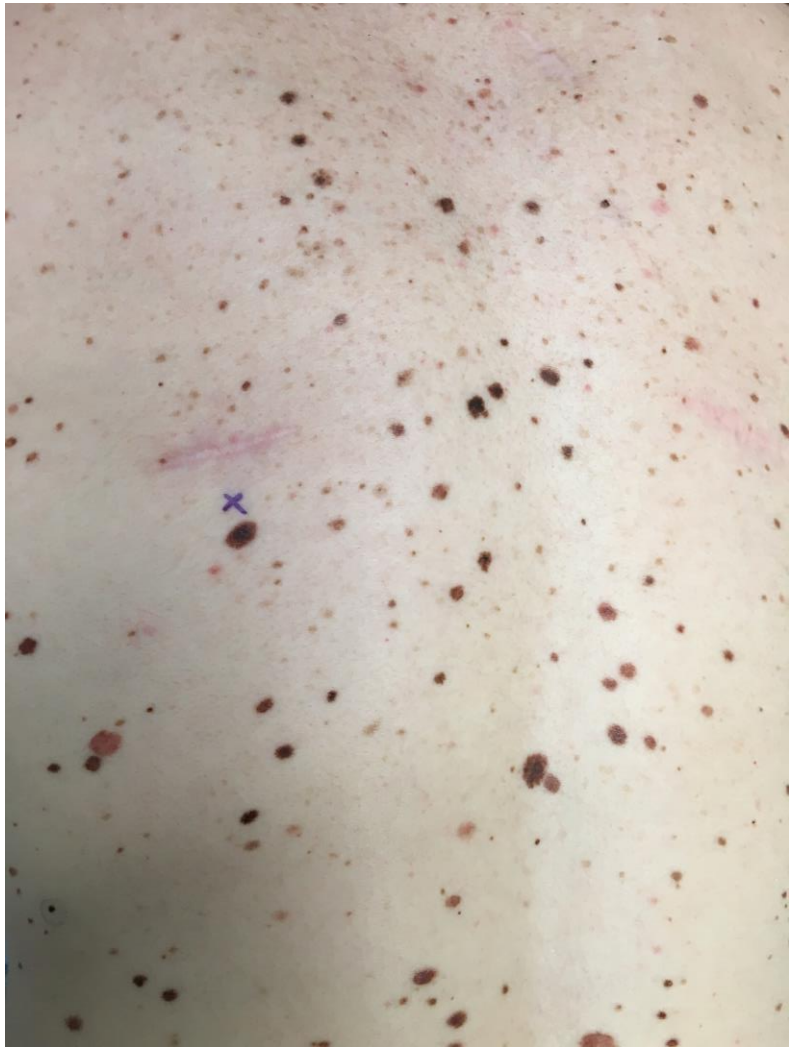

Stanford
ARTIFICIAL
INTELLIGENCE

A tool for surveillance? – NOT YET

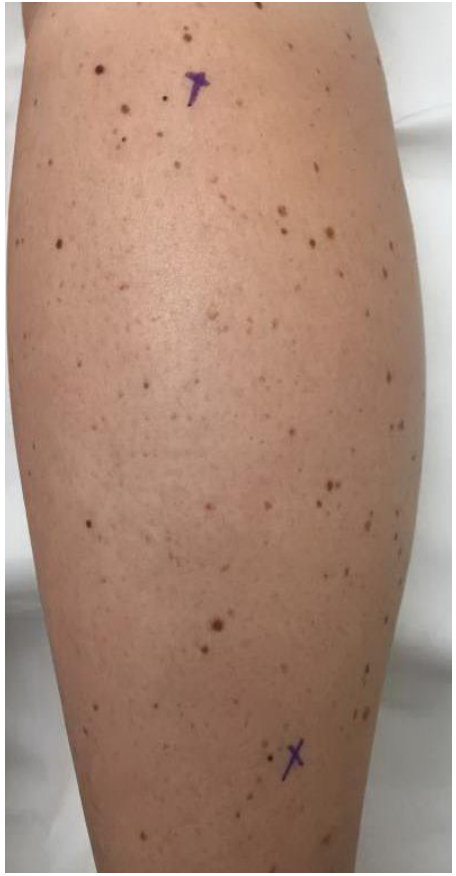

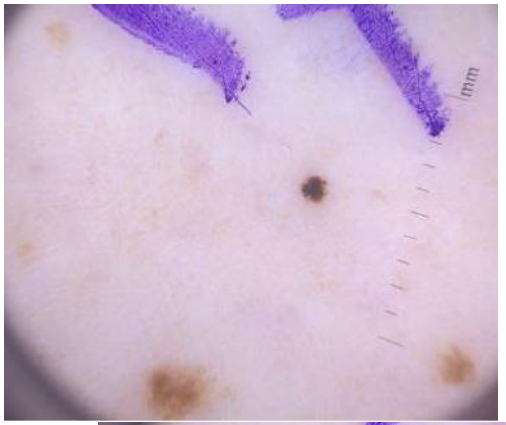

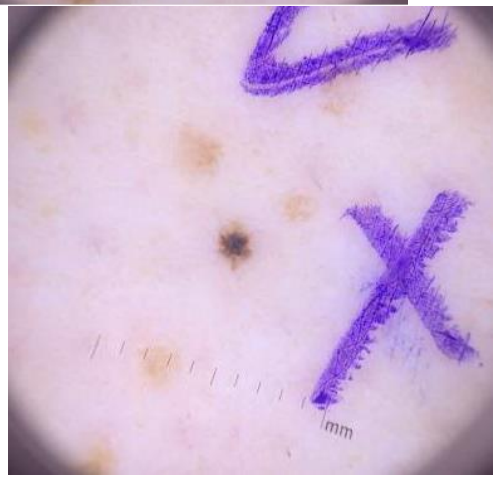

New moderately dysplastic nevi with focal severe atypia

Stanford ARTIFICIAL INTELLIGENCE

Stanford MEDICINE

Bottom line: We still need dermatologists!

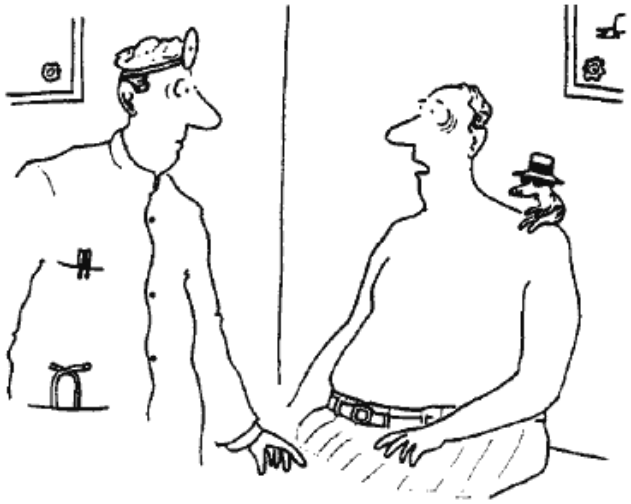

"Doctor, I have a suspicious looking mole on my shoulder."

Stanford Artificial Intelligence

Stanford Medicine