AI: from cats to medical imaging

Roska Tamás Doctoral School of Sciences and Technology

Ákos Kovács
akos.kovacs@mediso.com

Mediso Worldwide

Mediso is a medical imaging company developing hardware and software of CT, MR, PET and SPECT modalities.

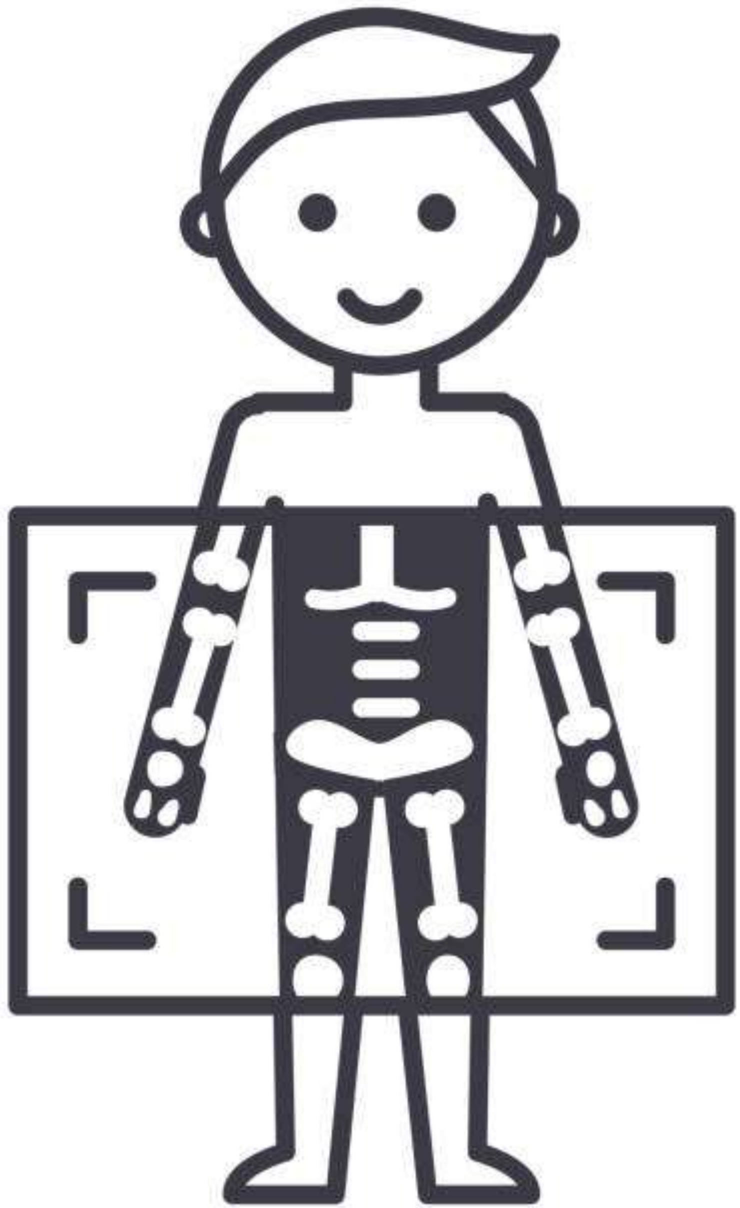

We are focusing on artificial intelligence as much as about any new technology.

We are using GPUs (CUDA, OpenCL) for reconstructions (e.g. Monte Carlo simulations) and FPGA-s for signal processing

When should one use a CT or an MR

If we need anatomical images

Structures are important

We would like to know about a region how it looks like

Iterative CT reconstruction

- Measured data: 20000-80000 projections, 896x16 pixels (1-4 GB)
- Reconstructed data: 30-300 slices, 512x512 pixels (30-300 MB)
- With iterative CT reconstruction
  - Quality of the reconstructed image can be improved
  - Radiation dose can be reduced
- Mapping projection data into volume data using rays
  - Parallelized with GPU (CUDA)
  - Texture memory for interpolation
- One iteration step is 2s with Nvidia Geforce 2080Ti (100-2000 total iterations required)

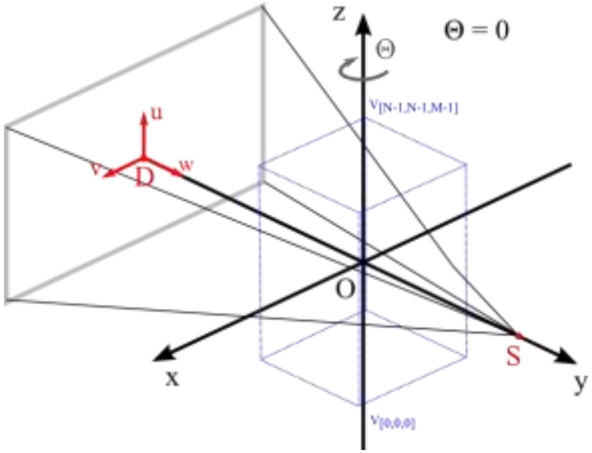

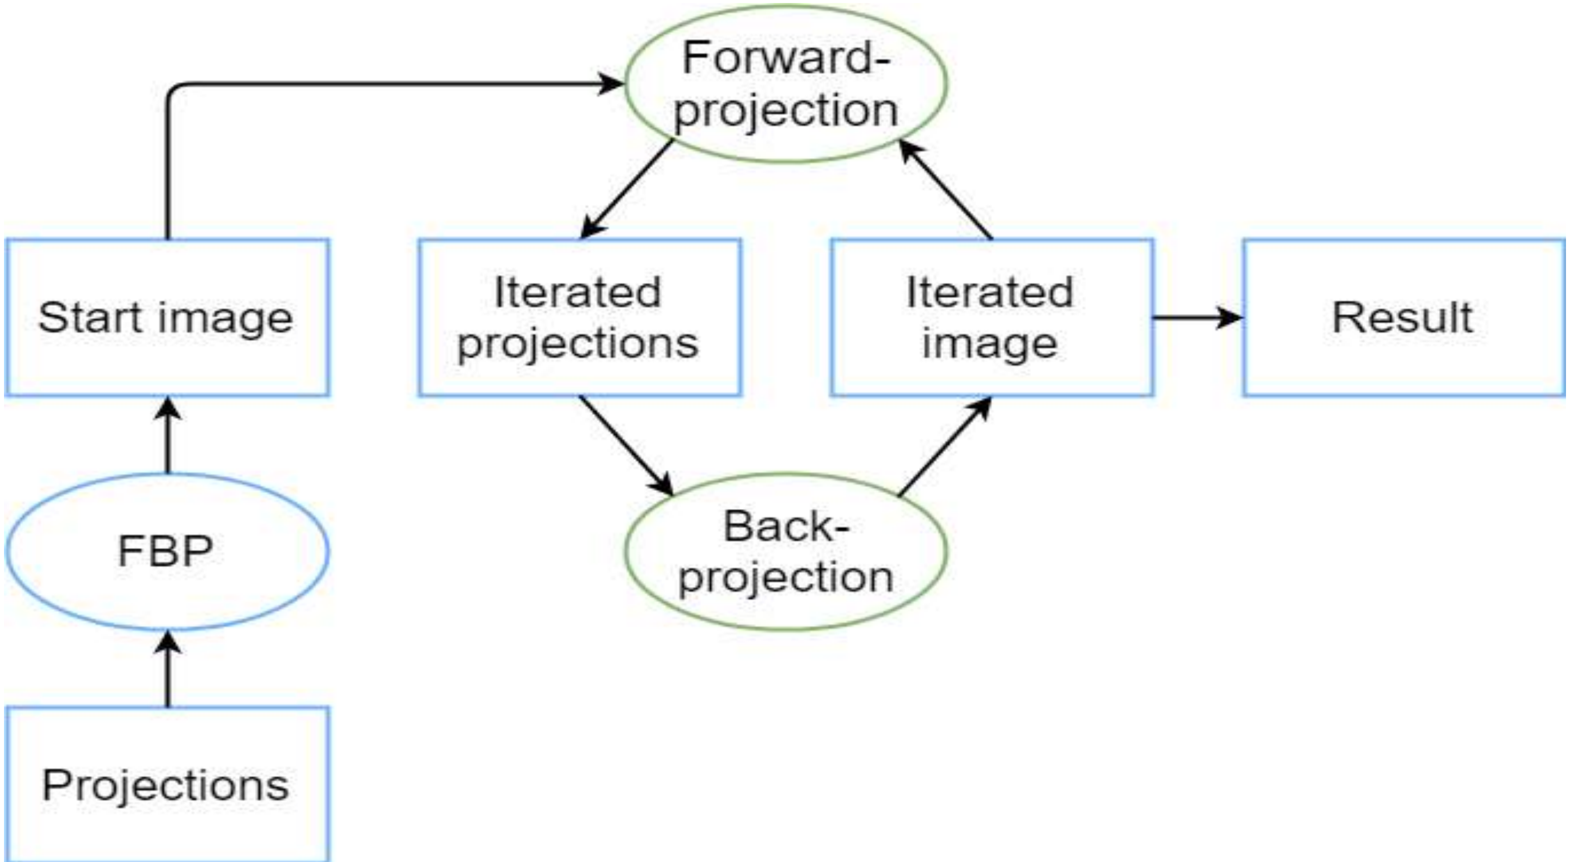
Forward-
projection
Start image
Iterated
projections
Iterated
image
Result
FBP
Back-
projection
Projections

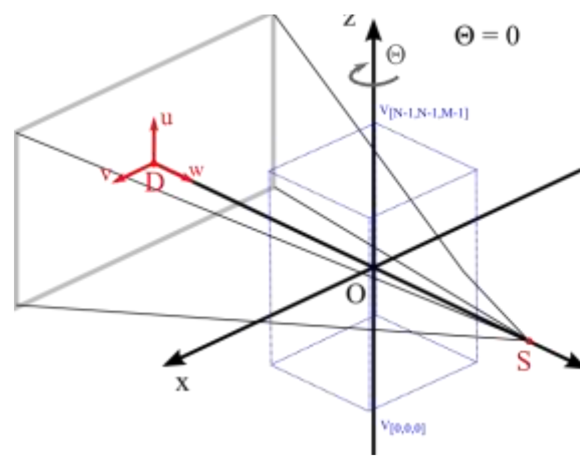
z
Θ
Θ = 0
u
v
w
D
O
S
x
y

When should one use a SPECT or PET

If we need functional images

We would like to answer these questions:

- Where and how much pharmaceutical is gathered up by cells?
- Where are regions with outlier activity? For example cancer cells?

Iterative PET reconstruction

GPGPU is perfect for PET reconstruction because …

- Input data: up to 300-700 million float32 values
- Output data: up to 50-150 million float32 values

Steps of reconstruction:

- Projections (geometric or Monte-Carlo)
- Corrections (attenuation, scatter, random, deadtime, detector-model, …)
- Filters (mostly volumetric)

There are analytical (2D-FBP) and iterative (EM) recons

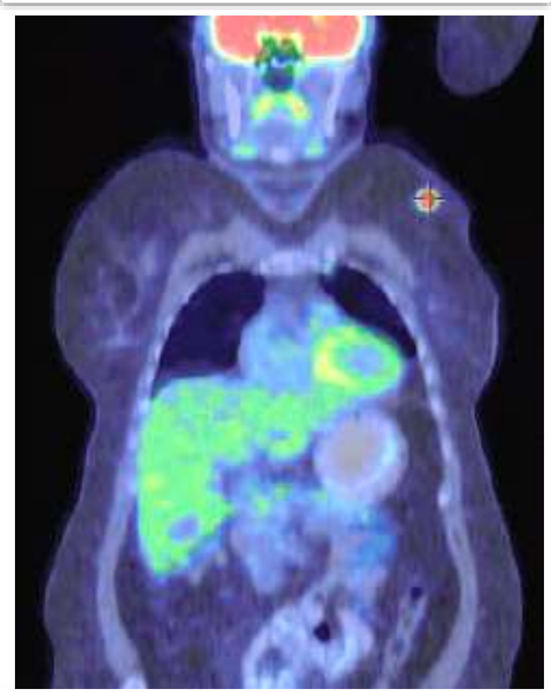

Iterative SPECT reconstruction

Goal: Determine the source distribution of the gamma emitting radionuclides in the body

- Follow the path of photons one by one
- Precise modeling of physical processes
- Simulation of random processes („Monte Carlo”)

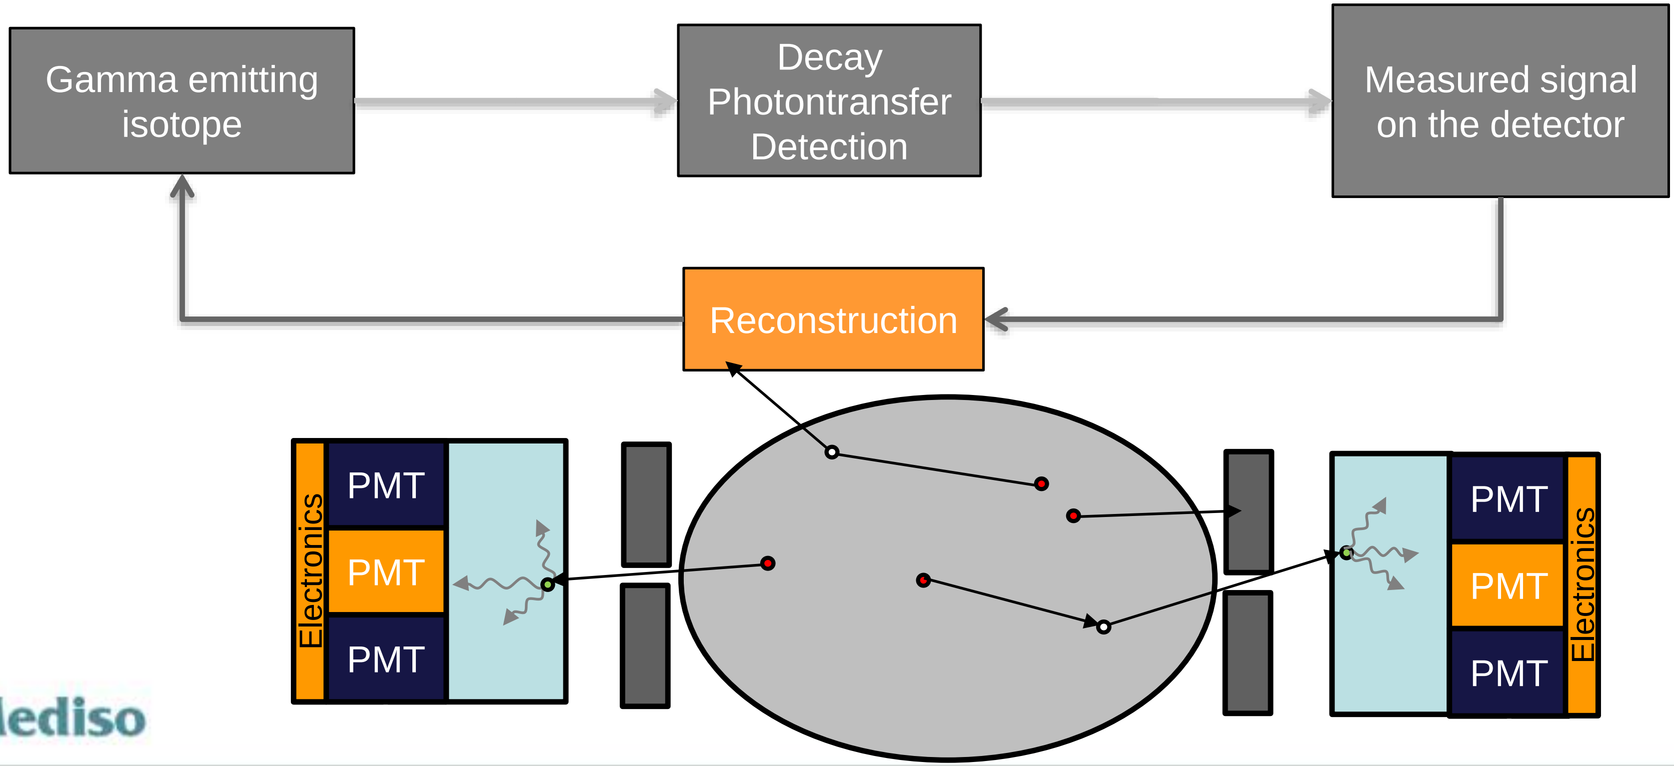

CT

MR

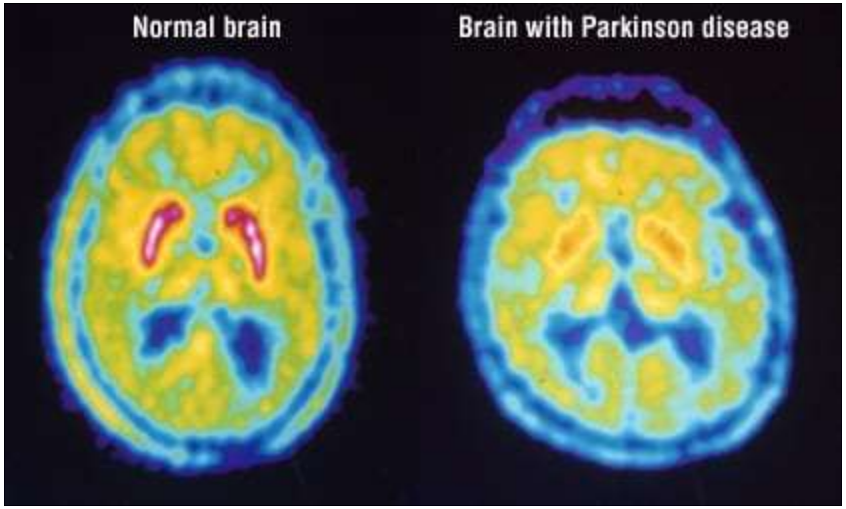

PET/SPECT

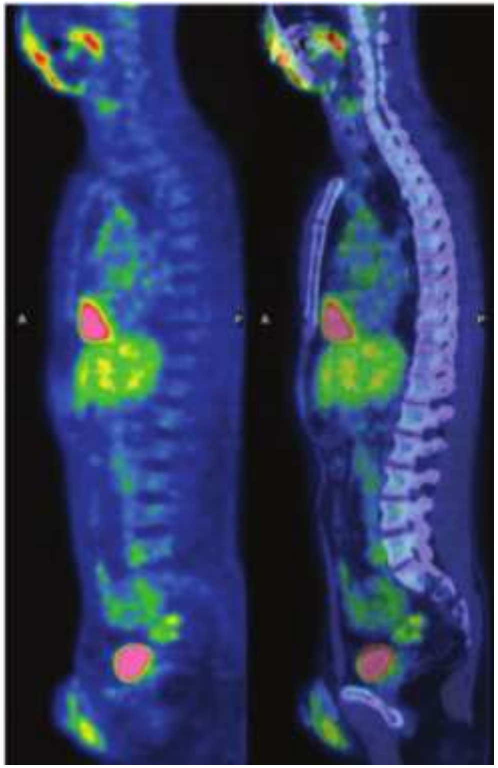

38.4%

https://www.cancer.gov/about-cancer/understanding/statistics

AI: from cats to medical imaging

Literature review showed us that neural networks can perform exceptionally well

- But the articles focus on technology
- They do not detail real-life use cases

Apart from developing neural networks, the other necessary steps:

- Prove the validity of the algorithms
- Test the methods
- Attack the networks to recognize their limits and capabilities

Datasets

MNIST

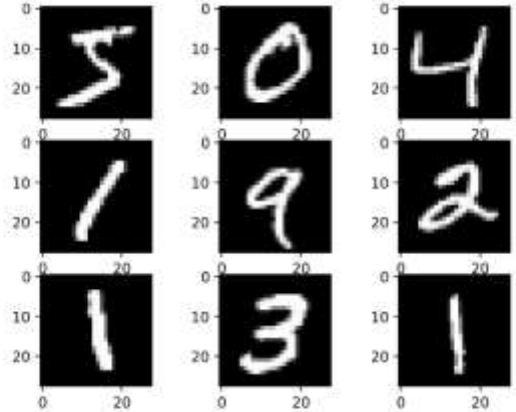

CIFAR10

Microsoft COCO
Cityscapes
Imdb_reviews
Iris

Instance segmentation


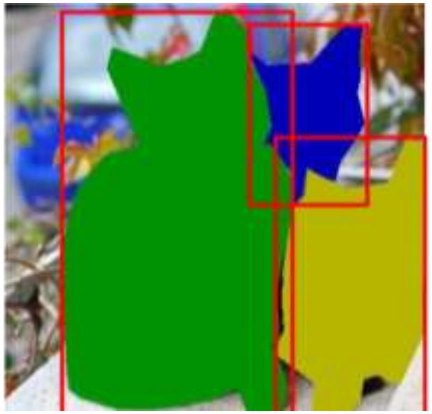


In medical imaging:

segment organs → for evaluation

→ motion correction

segment all lesions (cancer) seperately

https://medium.com/@danielmechea/what-is-panoptic-segmentation-and-why-you-should-care-7f6c953d2a6a

Semantic segmentation

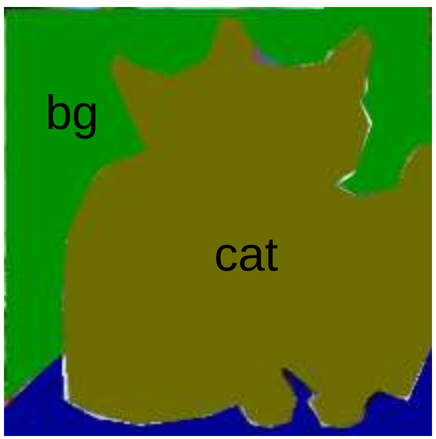

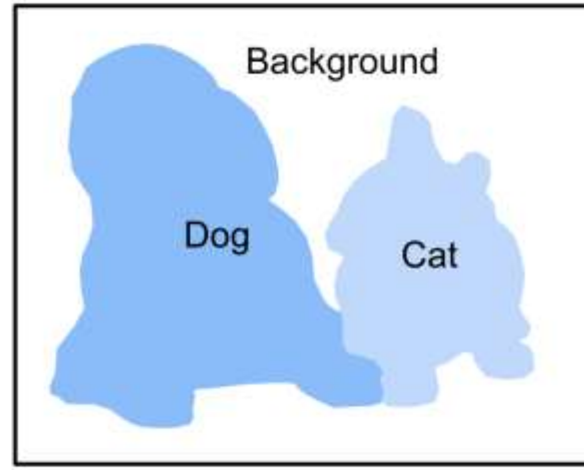

In medical imaging:

segment tissue types → for attenuation correction

background – foreground separation

https://www.d2l.ai/chapter_computer-vision/semantic-segmentation-and-dataset.html

https://medium.com/@danielmechea/what-is-panoptic-segmentation-and-why-you-should-care-7f6c953d2a6a

Object Detection

Tutorials from http://medium.com or http://towardsdatascience.com

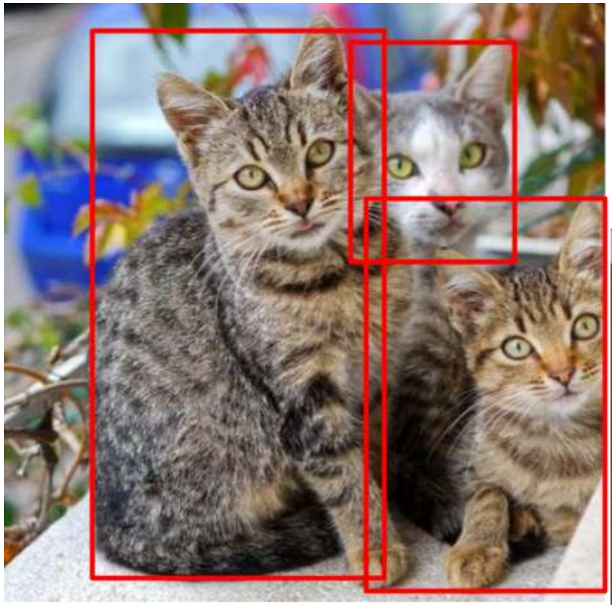

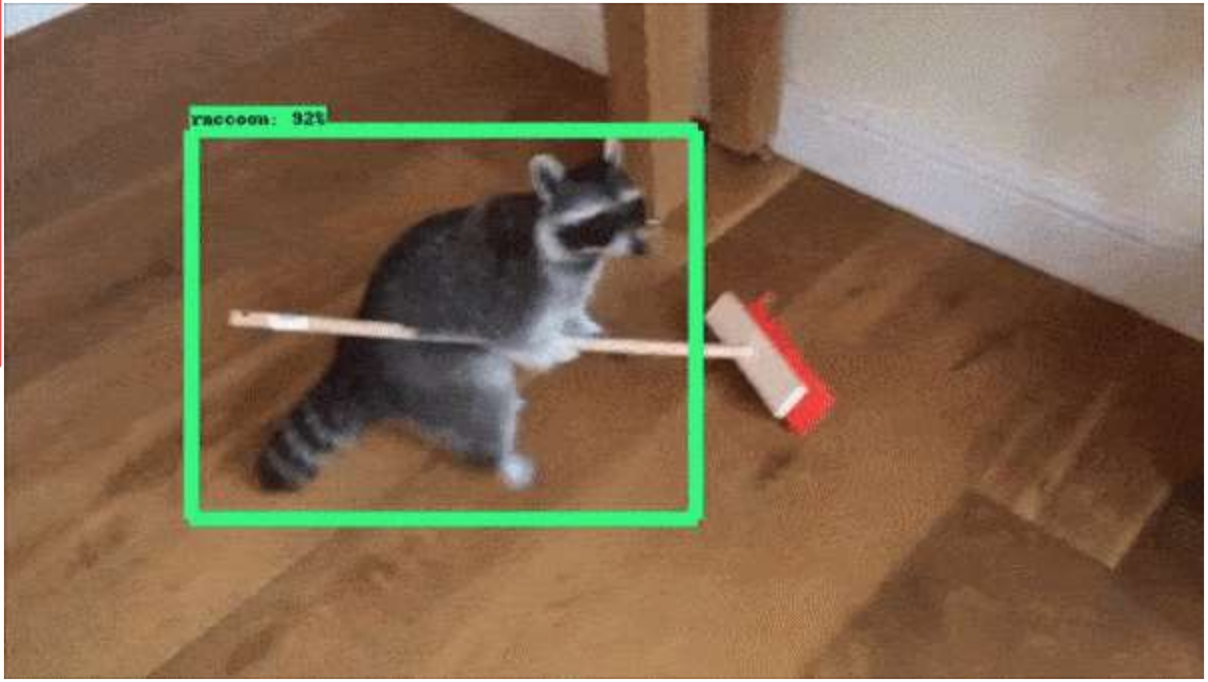

https://towardsdatascience.com/how-to-train-your-own-object-detector-with-tensorflows-object-detector-api-bec72ecfe1d9

Object Detection: lesions

Detect cancer! Where are the lesions?

Doctors annotate only ~100 measurements

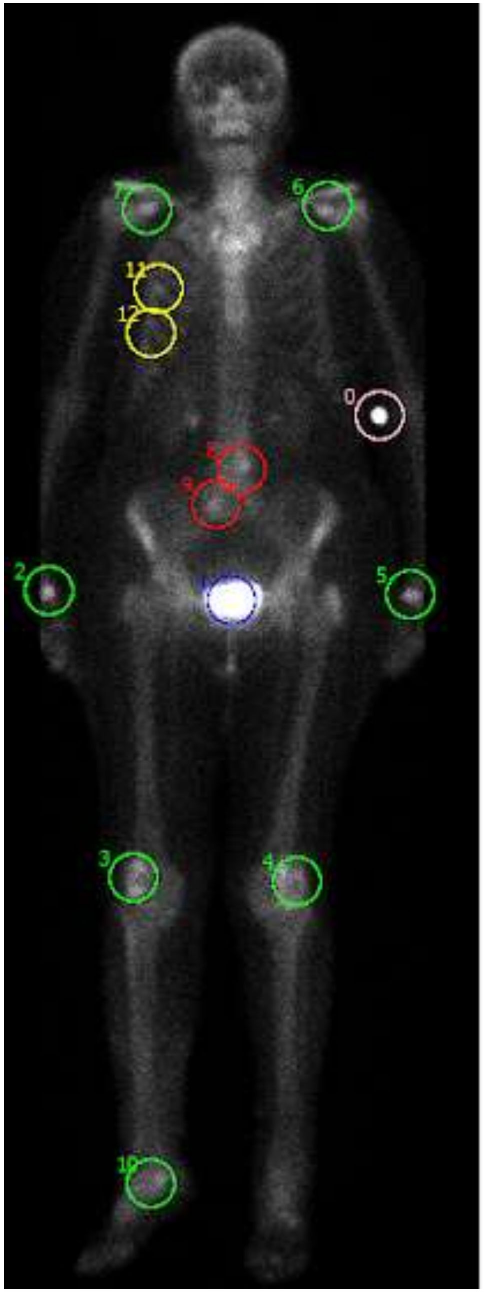

Object Detection

We created a solution to **select lesion candidates** and a separate network to **classify** them.

For lesion candidate proposal network

- we created a **synthetic lesion generator**
  - It can simulate the characteristics of gamma planar camera and generate realistic lesions

The classification network:

- It can falsely predict lesions to low risk or high risk, because we only have annotated data about ~100 measurements.
- And some measurement come from healthy patients, without lesions.

Object Detection

We created a solution to **select lesion candidates** and a separate network to **classify** them.

For lesion candidate proposal network

- we created a **synthetic lesion generator**
  - It can simulate the characteristics of gamma planar camera and generate realistic lesions

The **lesion candidate proposal network** has a good performance. **Solved!**

The classification network:

- It can falsely predict lesions to low risk or high risk, because we only have annotated data about ~100 measurements.
- And some measurement come from healthy patients, without lesions.

The classification is not robust. **More Data!**

Segmentation

Pilot study by a student

Bone segmenter for gamma planar images:

- only bones
- leave out soft tissues (kidneys, bladder emit a large signal in this measurements)

It uses a UNET like architecture

Future plans: change some of the Mediso's traditional segmentation methods to AI driven solutions, based on this pilot study

Image enhancing

We don’t need annotations by doctors:

- Much easier task at the beginning than Object Detection or Classification

Serious problems arise here

Two direction: **superresolution** and **noise filtering**

Original	1/8 statistics
	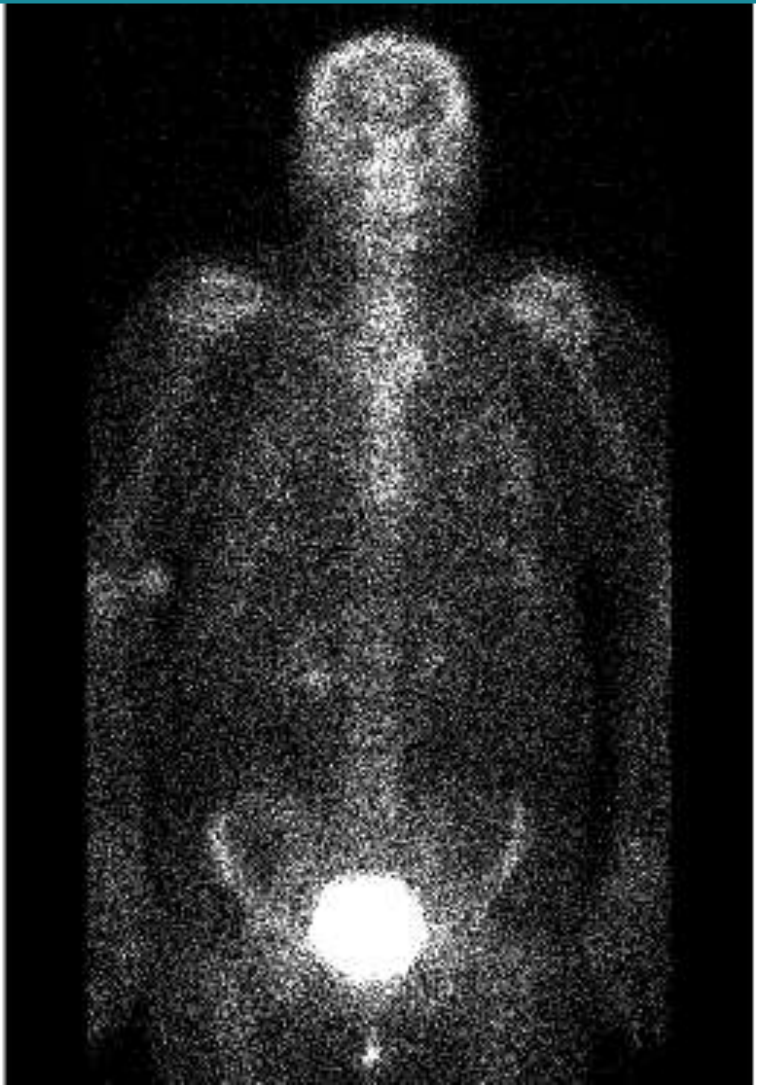

Noise filtering results

Superresolution

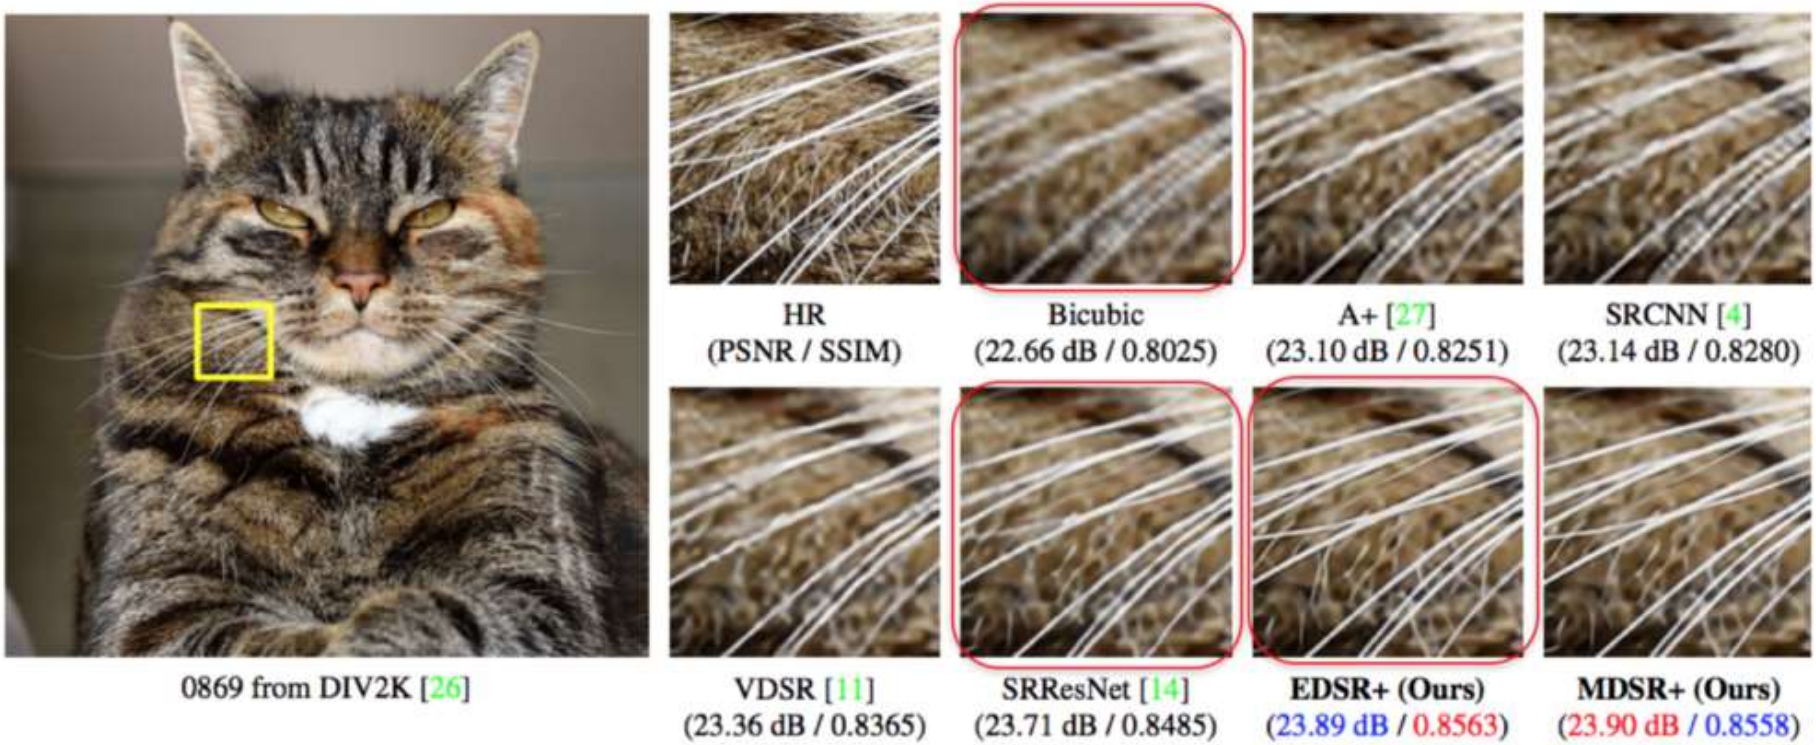

https://medium.com/@hyponymous/part-of-the-series-a-month-of-machine-learning-paper-summaries-fde1f72888c0

Superresolution (with a GAN)

16x64 (input)	256x1024 (predicted)	256x1024, (ground truth)

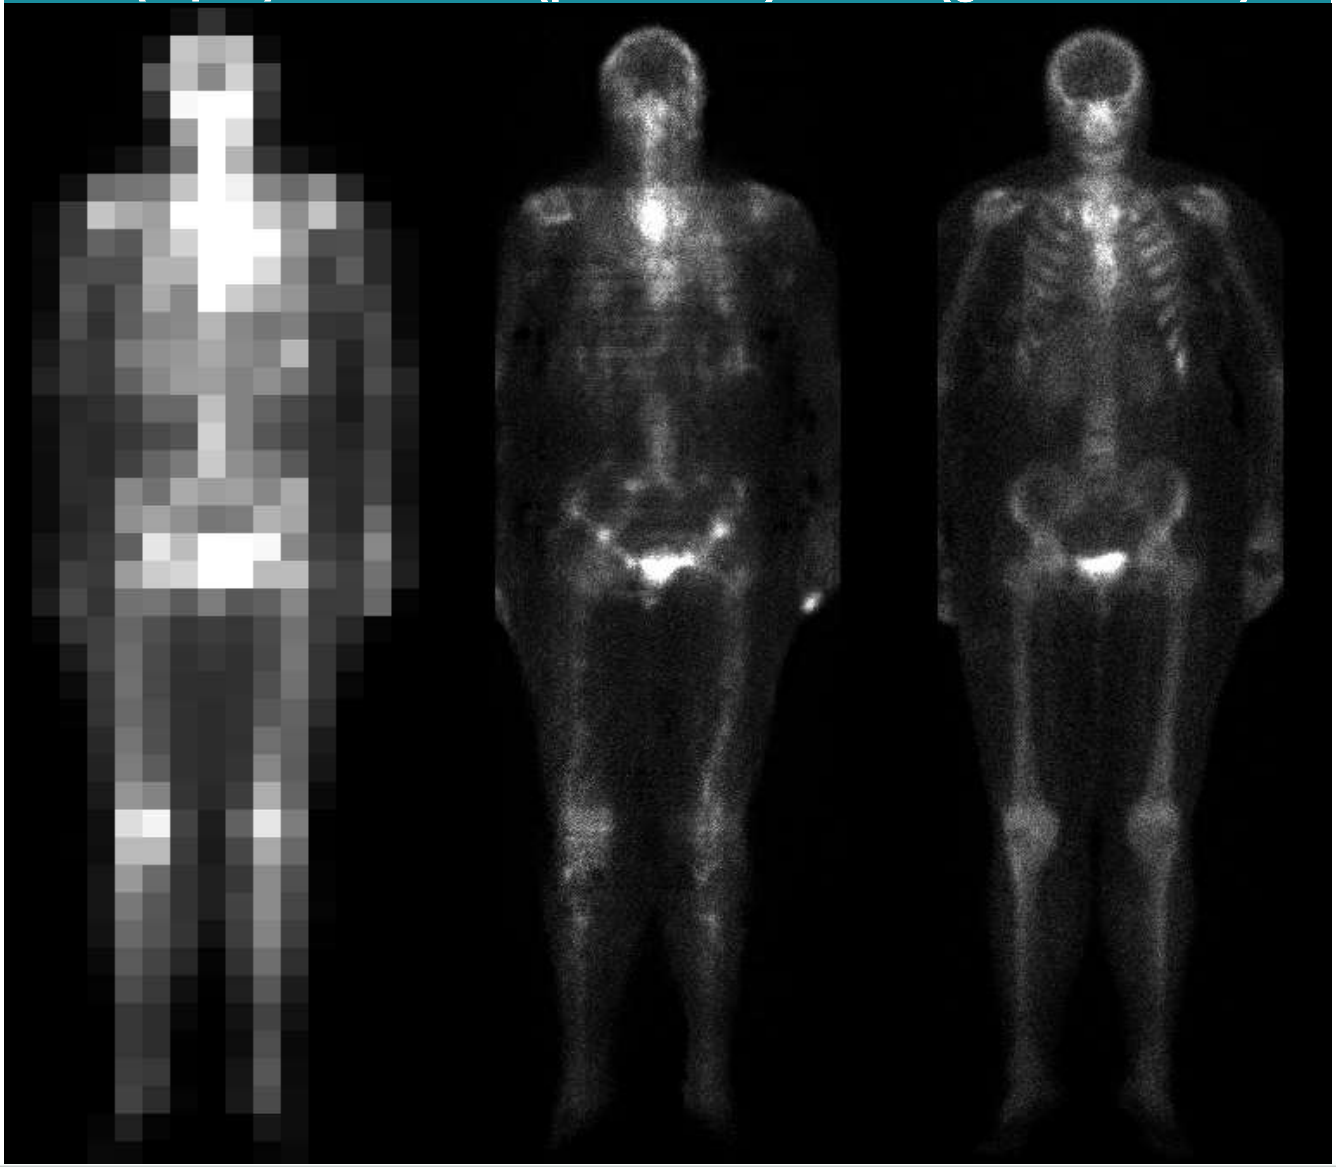

IMAGING FOR SCIENCE

Superresolution (with a GAN)

16x64 (input)	256x1024 (predicted)	256x1024, (ground truth)

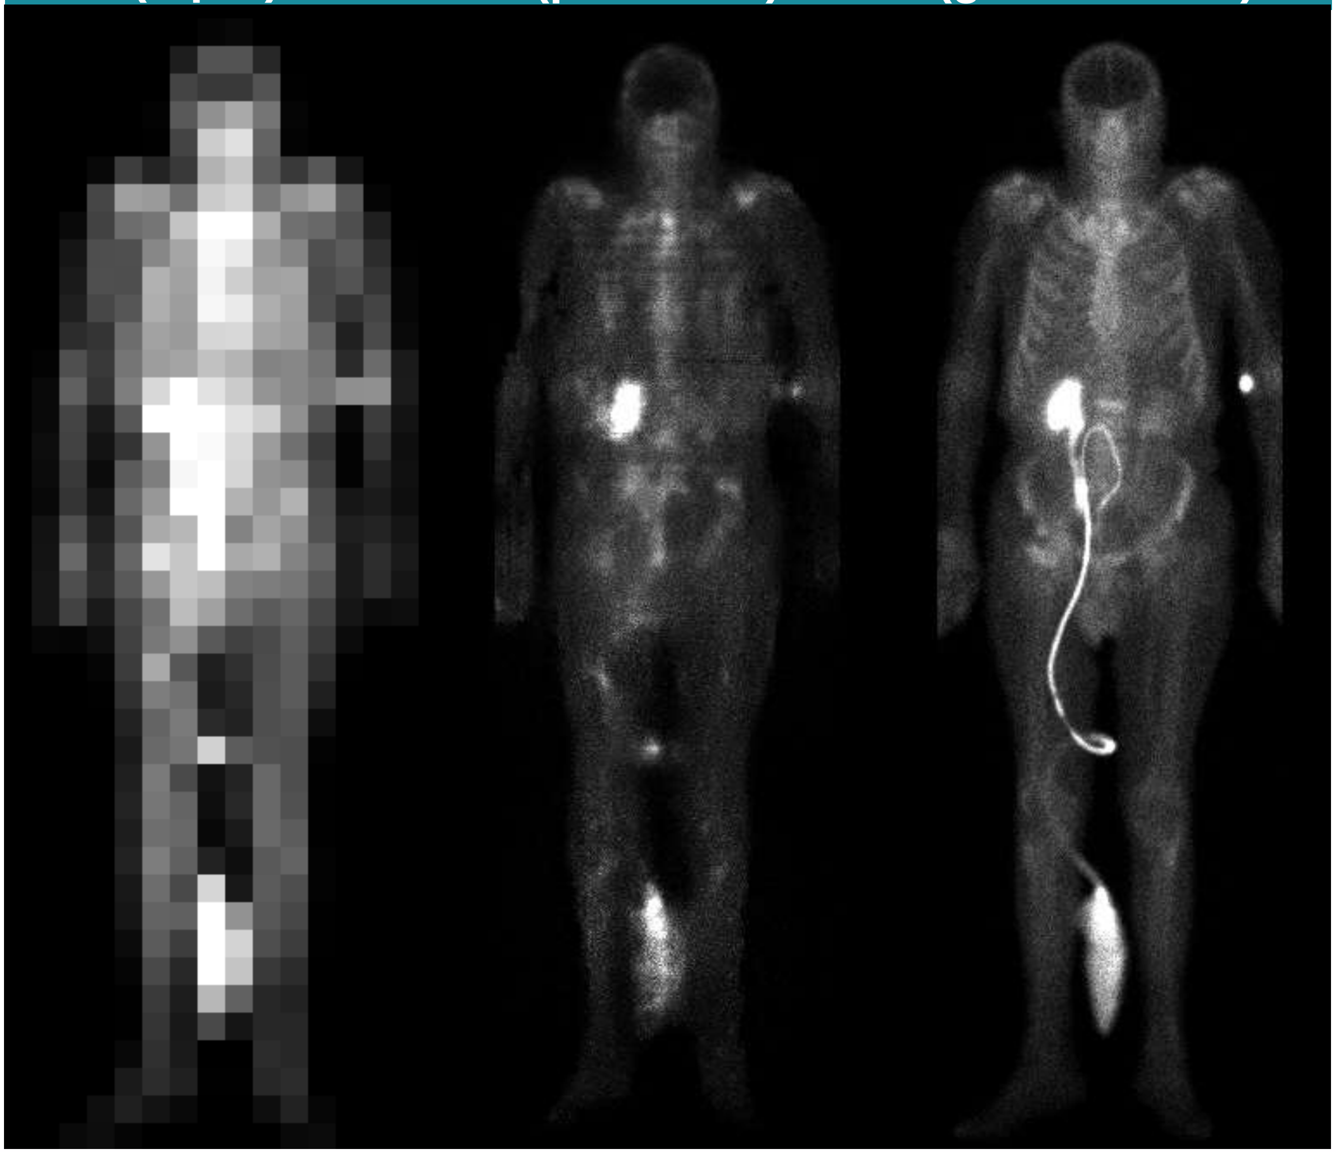

Bluffing

https://phillipi.github.io/pix2pix/

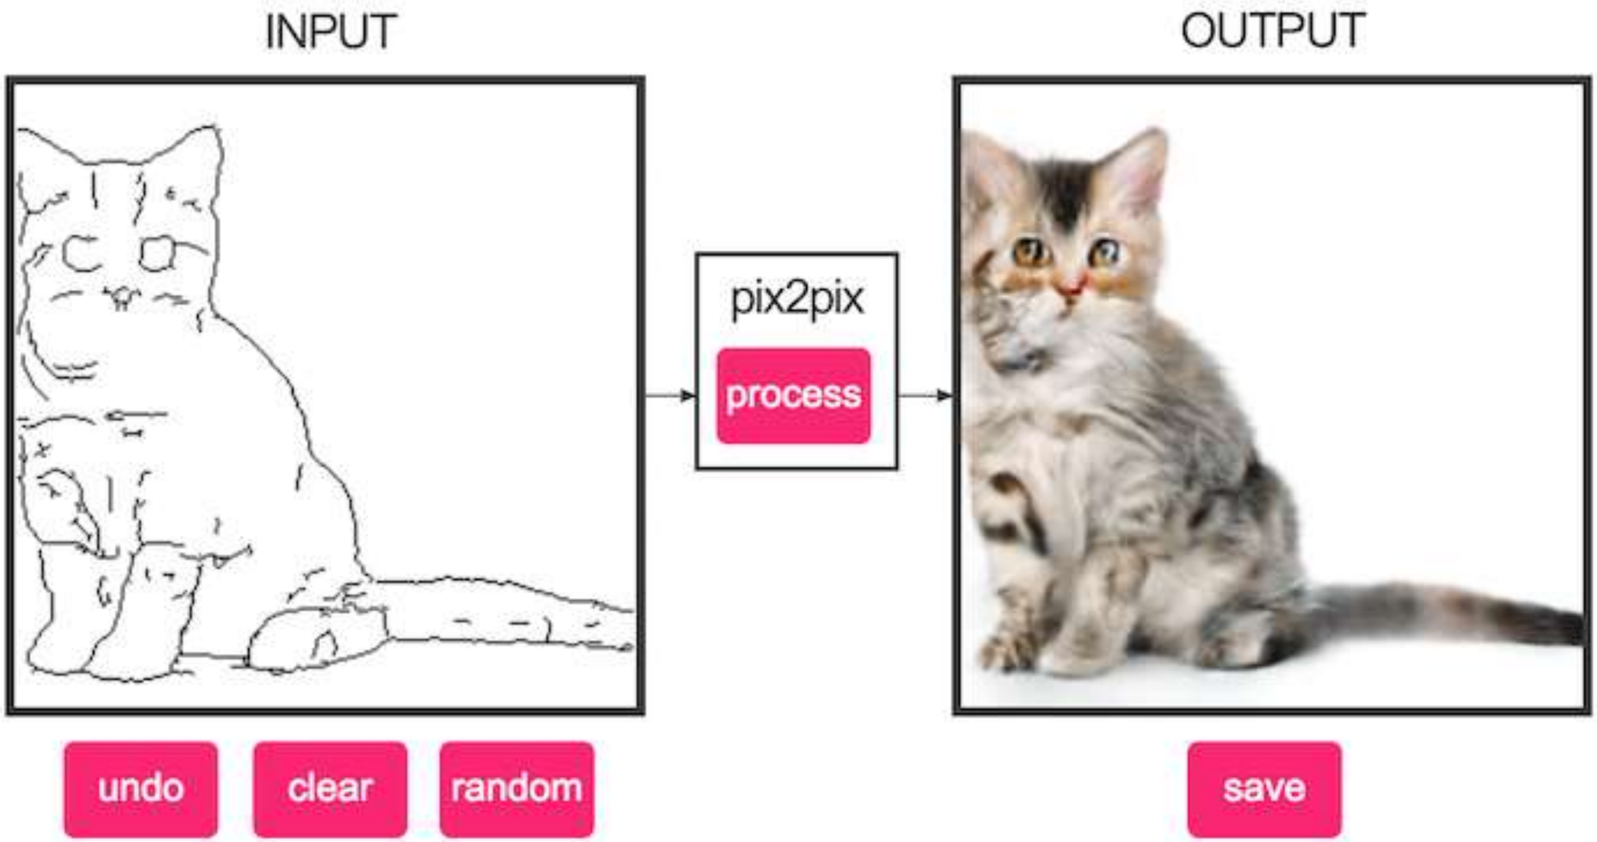

Bluffing

https://phillipi.github.io/pix2pix/

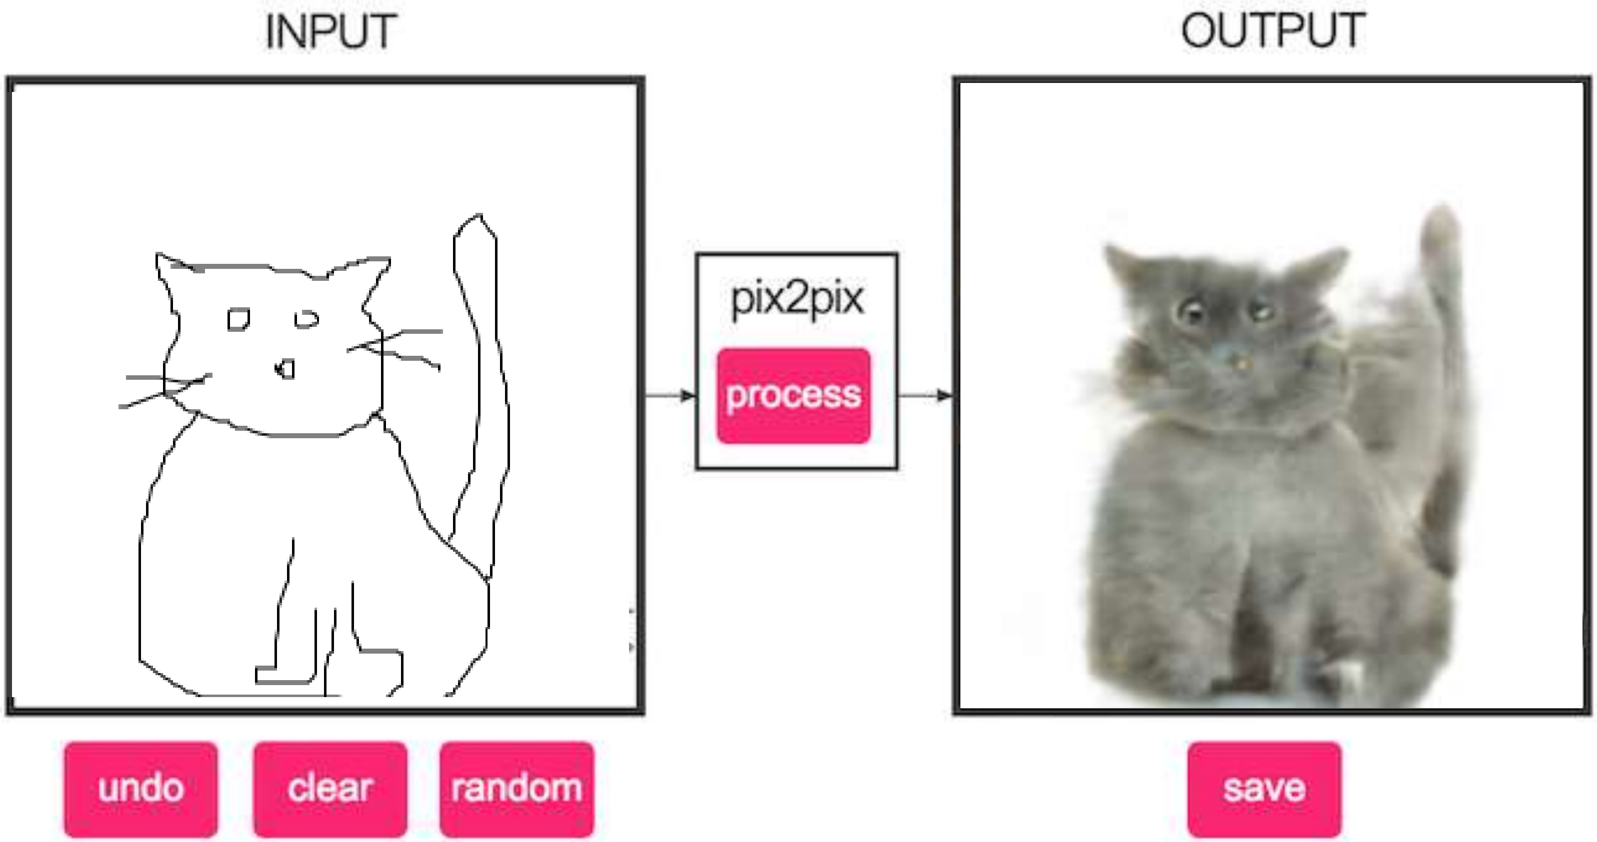


It's exactly what I was afraid of.
What?
Skeletons.

Which one is the better neural network?

We can find an excellent ground truth for cats.

But for a SPECT measurement we don’t have this convenience

For SPECT we can only use phantoms.

- They are worthless (NNs trained on patients)
- And it is prohibited to train on phantoms

We can’t measure contrast to noise ratio: because we don’t have ground truth noise free regions. And traditional image metrics are useless.

Currently our best bet is to generate synthetic lesions, and measure the performance of the noise filter on them.

We would like to diagnose, not just generate beatiful images.

Bugs and problems

No traditional software bugs

But there is:

- „not expected behaviour”
- Bluffing
- Overfitting
- Mode collapse

And there is another problem with the enhanced, noise filtered images:

Doctors say: **not enough noise!**

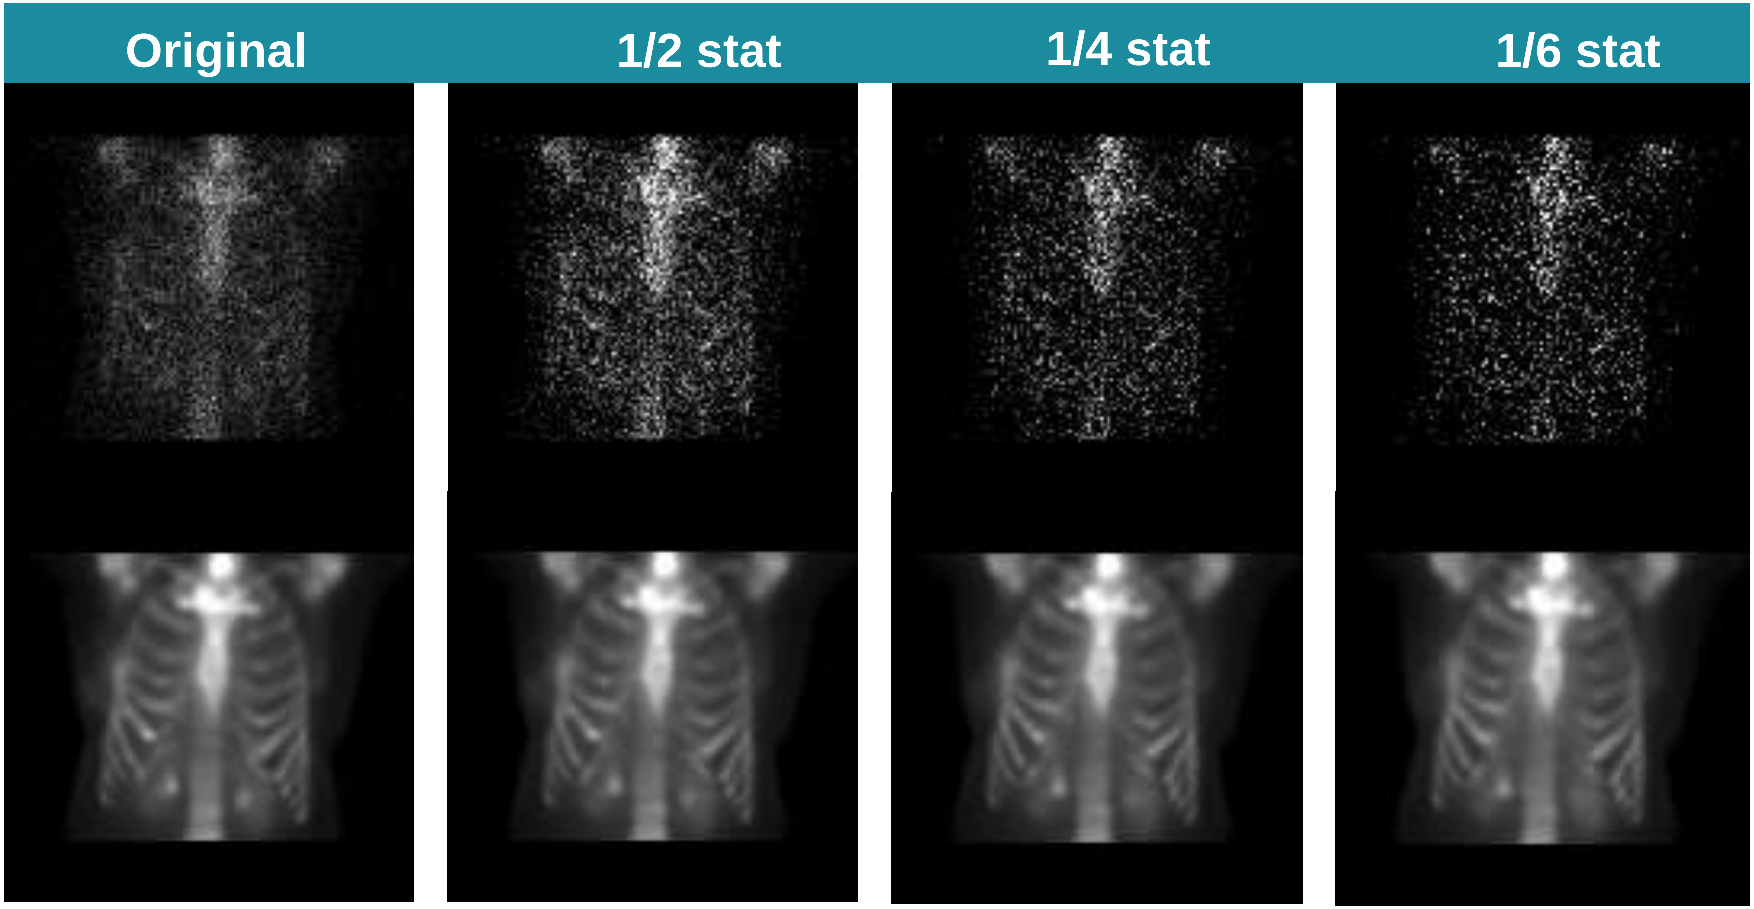

The doctors gain a **confidence factor based on the quality of the image**

We generate noise-free, beatiful images from any statistics. The doctors always think: the network is bluffing.

For clinical validation we have to create a solution to renoise the filtered images

Original → 1/8 stat → AI filtered of 1/8 stat → Renoised AI filtered of 1/8 stat

Clinical prevalidation

We created a set of images:

1. Original measurements
2. Degraded measurements (for example with 1/8 statistics)
3. Filtered original measurements → +add noise to them: renoise
4. Filtered degraded measurements → +renoised

We shuffled the original measurements more times to the test set. We would like to measure the „robustness" of the doctor.

Prevalidation results: the AI filtered images are better on the same statistics. And it seems like, even the ¼ stat measurement with filtering can provide the a same diagnostic quality as a standard measurement.

Why is this work important?

With improving the quality of bone scintigraphy and SPECT measurements:

We get better images:

better, easier diagnostics

The patient can get half or much less radiation dose.

(Radiactive pharmacons in large quantities can be dangerous.)

The measurement time can be reduced:

Instead of 10-45 minutes, some minutes can be enough

Take home message:

We have to devote much more time for validation than developing. We have to **test neural networks and recognize their limits and real capabilities.**

<u>**We are hiring:**</u>

- AI developers (Tensorflow, PyTorch, Keras, MongoDB)
- SPECT,CT, PET imaging developers (C++, **CUDA**)
- <u>C++ developers (Qt)</u>
- <u>C++ developer for visualisation (</u>OpenGL, Vulkan)
- Java <u>developer</u>
- PET physicist
- DevOps expert
- Software or hadware testers

Student's are welcome :)

All open positions: http://mediso.com/jobs.php

Ákos Kovács

akos.kovacs@mediso.com