Key tasks for AI in musculoskeletal imaging

ESSR 2023, Bilbao

Jean Feydy
HeKA team, Inria Paris, Inserm, Université Paris-Cité

Thursday, 22 June 2023

Today's talk: a computer's perspective on MSK imaging

1. What is an **image**?

2. Software **bottlenecks** for AI research

3. What you can **expect** going forward

I do not have any conflict of interest to disclose.

What is an image?

What do you see on a medical image? [Zyg]

What do you see on a medical image? [Zyg]

1. Pixels

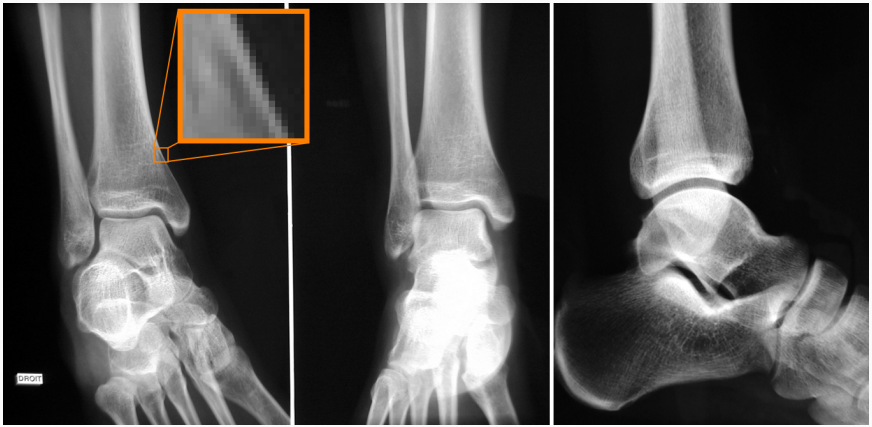

What do you see on a medical image? [Zyg]

1. Pixels
2. Anatomy

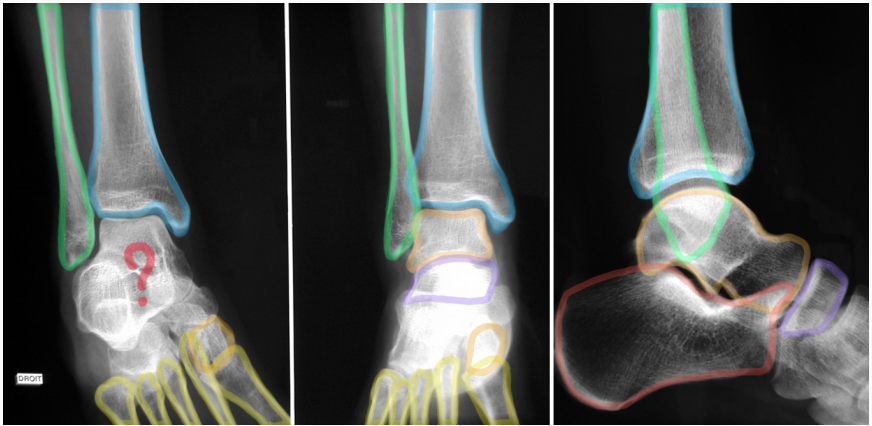

What do you see on a medical image? [Zyg]

1. Pixels
2. Anatomy
3. Function

What do you see on a medical image? [Zyg]

1. Pixels
2. Anatomy
3. Function

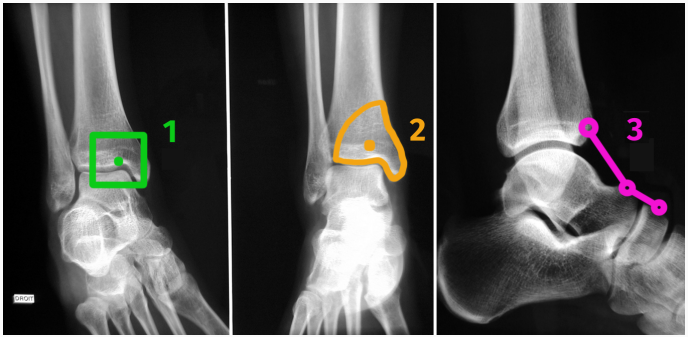

Simplifying a bit, each level of analysis corresponds to a way of **grouping pixels** with their neighbors.

1st level: a pixel grid

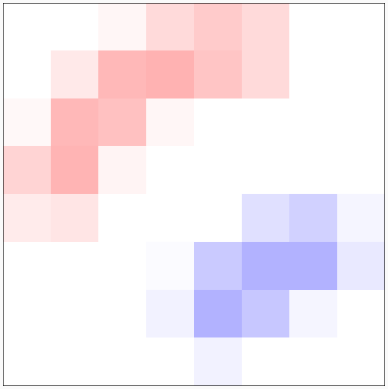

$N_x \times N_y \times N_z$ array of pixels.

Bitmap images and volumes:

- `.bmp`, `.png`, `.jpg`
- Standard in **radiology**.

\+ Ordered memory structure.
\+ Explicit neighborhoods.
\+ Fast **local** filters.

→ **Texture** analysis.
→ Organ **segmentation**.
→ Pattern **detection**.

2nd level: point clouds and 3D surfaces

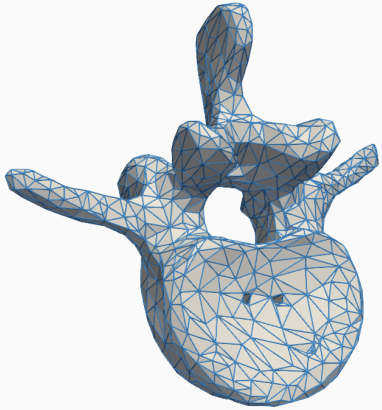

$N_{points} \times 3$ array of (x, y, z) coordinates.

Clouds of points ($\pm$ triangles):

- `.svg`
- Standard for **video games**.

$+$ Compact representation.
$+$ High precision geometry.
$+$ **Easy to deform**.

$\rightarrow$ **3D visualization**.
$\rightarrow$ Anatomical **atlas**.
$\rightarrow$ **Shape** analysis.

3rd level: biomechanical and/or physiological model [Zyg]

Volumetric mesh,
graph of interactions.

Mechanical/biological model:

- Finite elements, networks.
- Standard for **CAD**.

\+ Prior **knowledge**.
\+ **Robust** to noise.
\+ **Realistic** behaviour.

→ **Physiological** interpretation.
→ **Infer** what cannot be seen (stress).
→ **Simulate** a surgery.

Strengths and weaknesses of these image formats

Looking for the **neighbors** of a point in 3D space?

- On a **grid** : read adjacent memory cells.
- With N **points** (x, y, z) : computation of N distances.

Want to **rotate** a bone by 10°?

- On a **grid** : artifacts, loss of details, transfers between memory cells.
- With N **points** (x, y, z) : simple arithmetics on the coordinates.

Computational **speed** $\iff$ Training on **large datasets**.

To summarize

AI = **statistical regression** method + relevant **computational model**.

In medical imaging, we represent patient data as:

1. A 2D or 3D **pixel grid**.
2. An array of (x, y, z) **coordinates**.
3. A **web** of complex interactions.
4. All three at once!

In most cases, we define a large **structured formula**:

$$\text{image} \xrightarrow{F} F(\text{image}) \simeq \text{diagnostic}$$

F is a parametric computing **architecture**
$\simeq$ **model** to fit $\simeq$ **network** to train.

Software bottlenecks for AI research

The AI revolution is driven by gaming computers

Digital images and machine **learning** have been studied for **decades**.
Breakthrough in 2010-15 : using **PlayStations** to do **science** became **easy**.

Research effort at all levels towards:

- Increasingly powerful **computers**.
- Increasingly convenient **software toolkits**.
- Increasingly relevant **models**.

Spectacular results in a few applications

$\Longrightarrow$ massive **investments**, industry + governments.

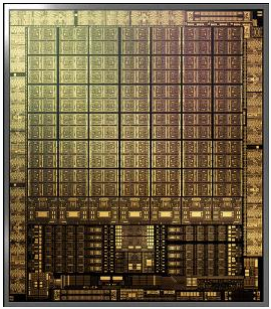

7,000 cores on a GPU.

For grid images: a mature ecosystem

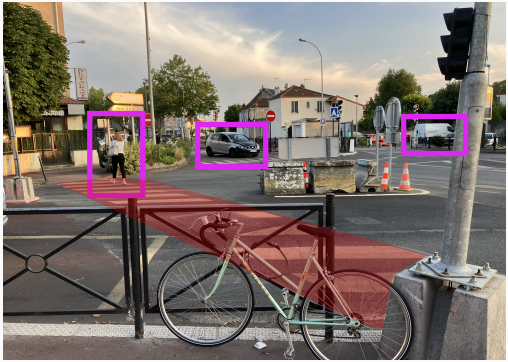

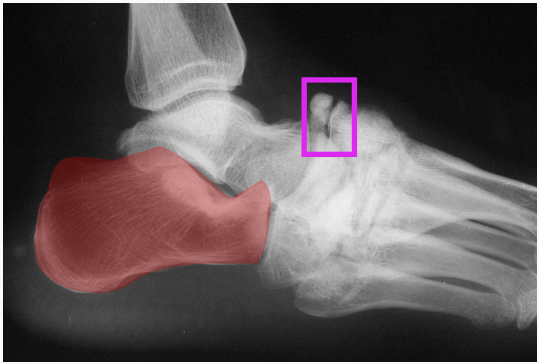

Main motivation for AI in 2012-2022: **self-driving cars**.
Key challenges: **segment** the environment, **detect** other actors.

Two full software suites to manipulate **images as grids of pixels**:
TensorFlow (Google) and PyTorch (Facebook-Meta).

To go beyond prototypes, AI engineers need a full software suite

Graphics: Printer + Driver + **Photoshop** $\Longrightarrow$ Illustrations

Tabular data: GPU + **cuBLAS** + **PyTorch** / TensorFlow $\Longrightarrow$ **"Classical"** neural networks

Pixel grids: GPU + **cuDNN** + PyTorch / TensorFlow $\Longrightarrow$ **Convolutional** neural networks

Point clouds and graphs : GPU + **CUDA** + **??** $\Longrightarrow$ **Geometric** neural networks

For point clouds and graphs: work in progress

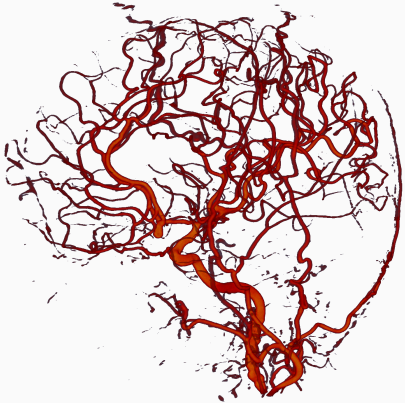

Brain arterial network.

How do we **process this object**?

An ecosystem under construction:

- **KeOps** : since 2017
  - Fast learning with **point clouds**.
- **PyG** : since 2018
  - Fast learning with **graphs**.
- **FEniCSx** and **PhiFlow** : since 2018
  - Fast learning with **physics**.
- **PyVista** and **Vedo** : since 2019
  - **3D visualisation.**
- **scikit-shapes**: released soon
  - Easy **morphometrics.**

Conclusion

- **Gaming computers** (GPUs) are the workhorses of AI.
  A **full software suite** is required to rein in these machines.
- Since 2015, **medical imaging** rides a wave of investment
  from the **FAANG** for **natural** image processing.

  Breakthroughs: **segmentation**, **texture** analysis and lesion **detection**.
- What about **surgical** planning, **morphometrics**, **vascular** analysis… ?

  An **investment in the numerical foundations** of the field is under way.

References

References i

Zygote.

Solid 3d human foot and ankle model.

https://www.zygote.com/cad-models/solid-3d-human-anatomy/cad-human-foot-ankle-model.