Software bottlenecks for 3D AI

Sciences @ PSC seminar

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

Thursday, 7 November 2024

HeKA – on the 3rd floor: a translational research team for public health

Hôpitaux

Inria

Inserm

Universités

New methods for healthcare data?

Machine **learning**:

- **Numerous** patients.
- **Lots** of data per patient.
- Numbers, **images**, **texts**…

Healthcare data:

- Varying sampling **times**.
- **Missing** data.
- Multiple **centers**.

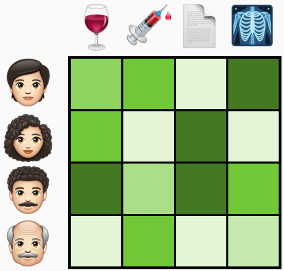

New methods for healthcare data?

Machine **learning**:

- **Numerous** patients.
- **Lots** of data per patient.
- Numbers, **images**, **texts**…

Healthcare data:

- Varying sampling **times**.
- **Missing** data.
- Multiple **centers**.

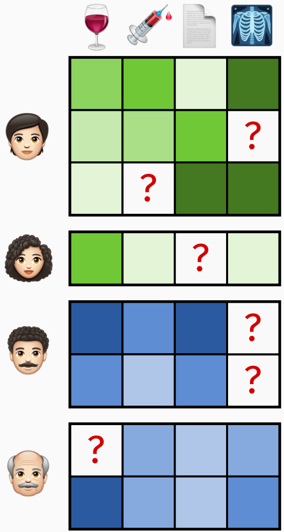

A key research area: encoding patient data

Some **topics**:

- Natural **language** processing for medical documents.
- Visualization and **generation** of complex data.
- Digital **anatomy** for 3D medical imaging.

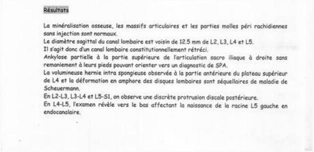

Résultats

La minéralisation osseuse, les massifs articulaires et les parties molles péri rachidiennes sans injection sont normaux.
Le diamètre sagittal du canal lombaire est voisin de 12.5 mm de L2, L3, L4 et L5.
Il s'agit donc d'un canal lombaire constitutionnellement rétréci.
Ankylose partielle à la partie supérieure de l'articulation sacro iliaque à droite sans remaniement à leurs pieds pouvant orienter vers un diagnostic de SPA.
La volumineuse hernie intra spongieuse observée à la partie antérieure du plateau supérieur de L4 et la déformation en amphore des disques lombaires sont séquellaires de maladie de Scheuermann.
En L2-L3, L3-L4 et L5-S1, on observe une discrète protrusion discale postérieure.
En L4-L5, l'examen révèle vers le bas affectant la naissance de la racine L5 gauche en endocanalaire.

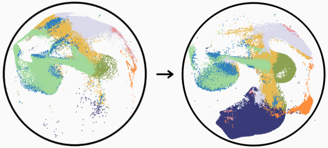

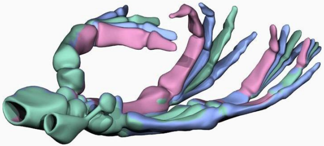

Maths @ PSC ⇔ **clinical work** @ AP-HP

Today's talk: a computer's perspective on 3D shapes

1. What is an **image**?

2. Software **bottlenecks** for AI research

3. Current **research** focus

What is an image?

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

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

1. Pixels

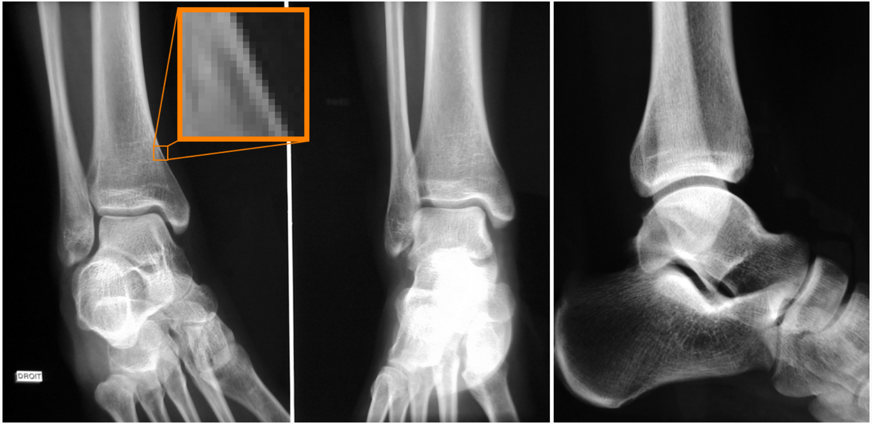

What do you see on a medical image? [Zyg]

1. Pixels
2. Anatomy

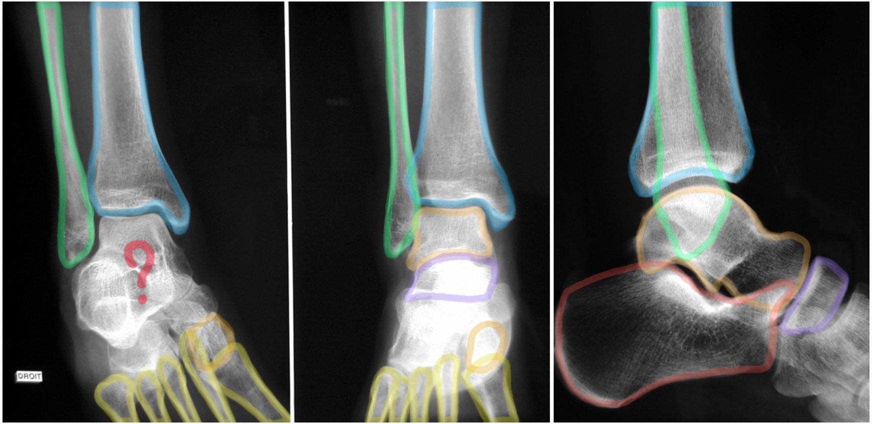

What do you see on a medical image? [Zyg]

1. Pixels 2. Anatomy 3. Function

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

1. Pixels
2. Anatomy
3. Function

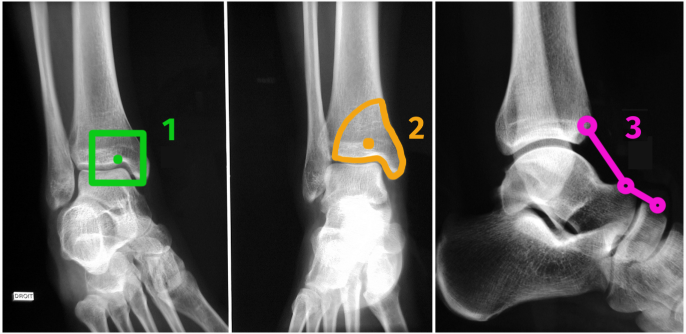

Simplifying a bit, each level of analysis corresponds to a way of **grouping pixels** with their neighbors.

1st level: a pixel grid

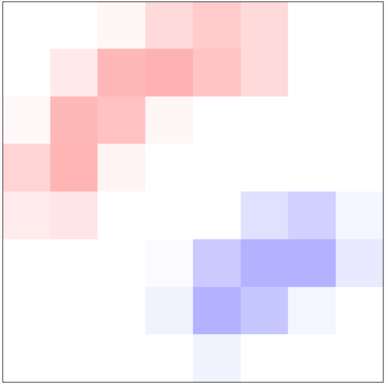

$N_x \times N_y \times N_z$ array of pixels.

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

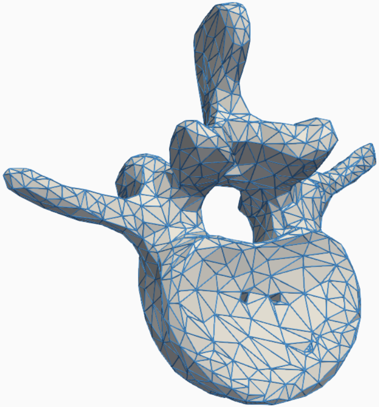

$N_{\text{points}} \times 3$ array of (x, y, z) coordinates.

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.

$\rightarrow$ **Physiological** interpretation.
$\rightarrow$ **Infer** what cannot be seen (stress).
$\rightarrow$ **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** $\Longleftrightarrow$ Training on **large datasets**.

To summarize

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

In biomedical imaging, we represent 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**:

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

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.

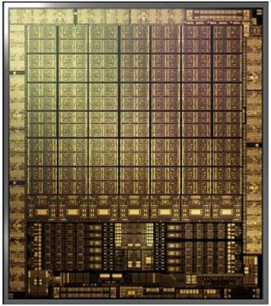

10,000 cores on a GPU.

For grid images: a mature ecosystem

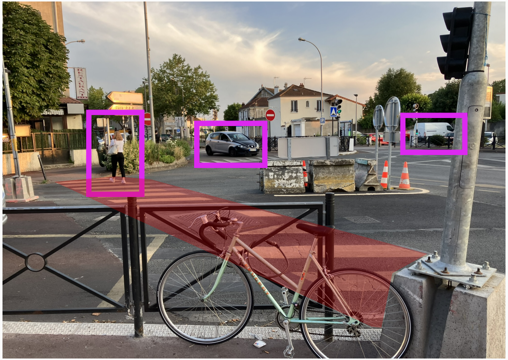

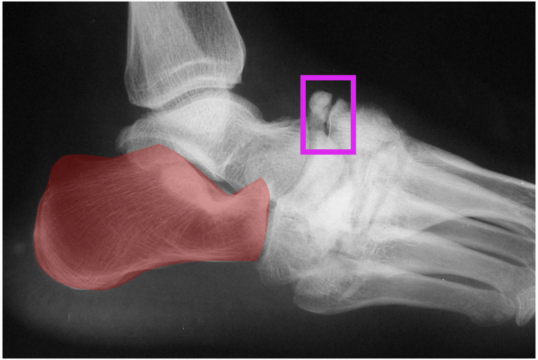

Main motivation for AI in 2012-2022: **self-driving cars**.

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).

For point clouds and graphs: work in progress

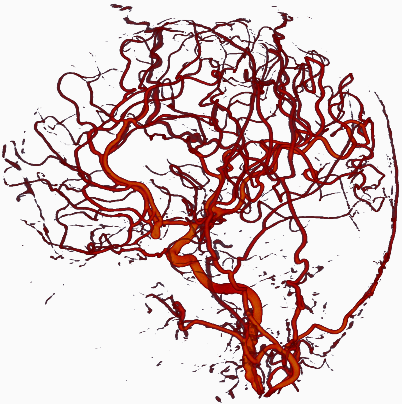

Brain arterial network.
How do we **process this object**?

An ecosystem under construction:

- **KeOps** : since 2017
  - Fast learning with **point clouds**.
- **PyG** : since 2018
  - Fast learning with **graphs**.
- **Warp**, **FEniCSx** and **Taichi** : since 2018
  - Fast learning with **physics**.
- **PyVista** and **Vedo** : since 2019
  - **3D visualisation.**
- **scikit-shapes**: in 2025
  - Easy **morphometrics.**

Untangling brain arterial networks

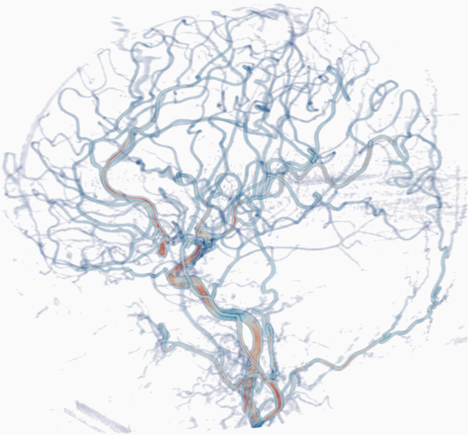

Raw 3D data.

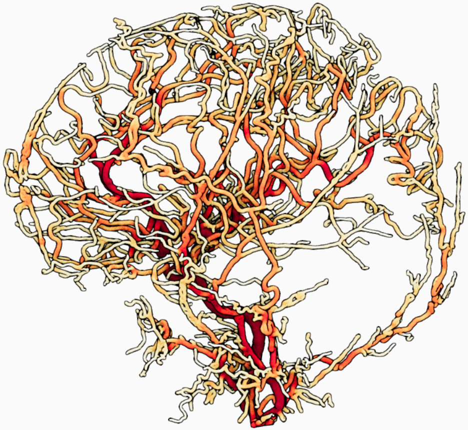

Segmented data.

Untangling brain arterial networks

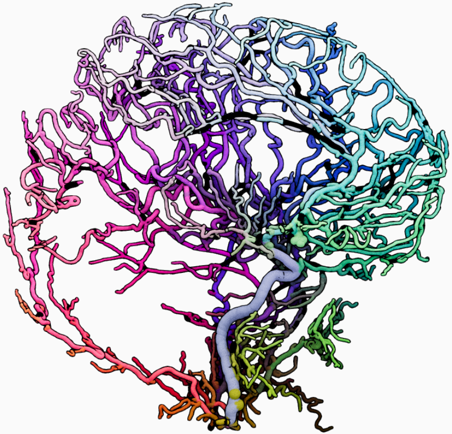

3D view.

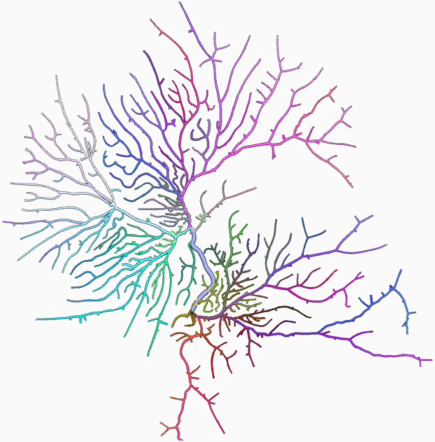

Unfolded 2D view.

Going further

We are currently working on:

- **Population** analyses of arterial networks.
- **Fracture** reconstruction.
- **Software** packaging.

$\Longrightarrow$ Hard open problems are related to **topology changes**.

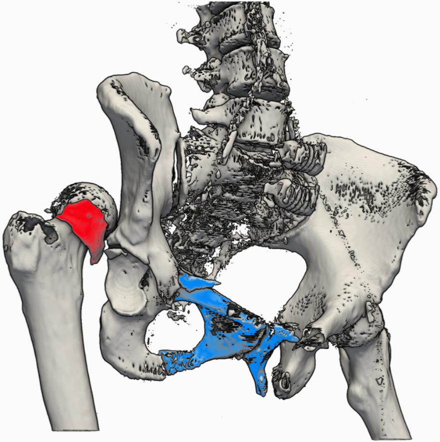

A multiple pelvic fracture.

Conclusion

- **Gaming computers** (GPUs) are the workhorses of AI.
  A **full software suite** is required to rein in these machines.
- Since 2015, **biomedical 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.
  Tradeoffs between ease of use, versatility, speed, portability, etc.

References

References i

Freyr Sverrisson, Jean Feydy, Bruno E. Correia, and Michael M. Bronstein.

Fast end-to-end learning on protein surfaces.

bioRxiv, 2020.

Zygote.

Solid 3d human foot and ankle model.

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