deep learning for automated segmentation of tomographic images

Copil BIGMÉCA January 2021, Monday 2021-01-13

João P C Bertoldo

supervisors

Henry Proudhon Etienne Decencière David Ryckelynck

introduction

myself

João Paulo Casagrande Bertoldo

yes, it's a pretty long name... (:

2019-2020 Paris-Dauphine - PSL University

Master IASD: Artificial Intelligence, Systems, Data

2017-2021 MINES ParisTech - PSL University

Executive engineering - Data Science Minor

2013-2020 University of São Paulo (USP)

Mechatronics engineering

<u>contact</u>

e-mails: joaopcbertoldo@gmail.com

joao.bertoldo@mines-paristech.fr

linkedIn: linkedin.com/in/joaopcbertoldo

github: github.com/joaopcbertoldo

project

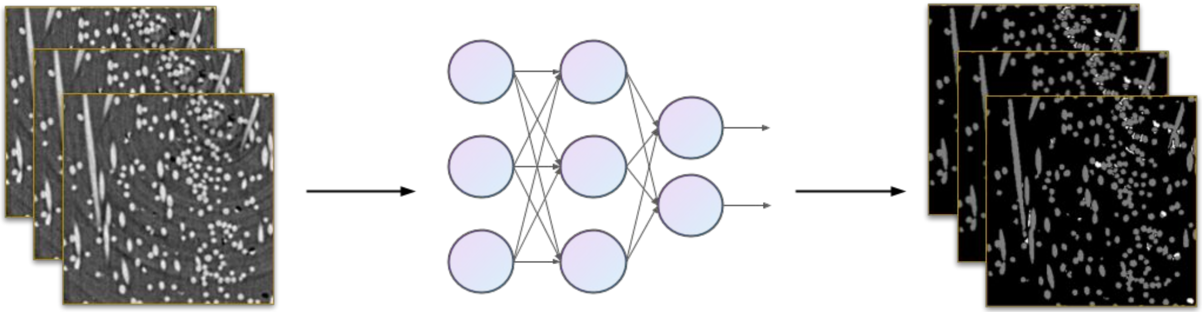

Tomography slices $I \in [0, 1]^{W \times L \times D}$

Neural Network $f : I \rightarrow S$

Segmentation $S \in \{0, 1, 2\}^{W \times L \times D}$

goal segment an x-ray tomography with a neural network

why? because doing it (semi-)manually is time-consuming

data a composite material with fibers and porosities

project repository: github.com/joaopcbertoldo/tomo2seg

presentation

- ➔ **background and context**
  material, u-net architecture
- ➔ **overview**
  project framework
- ➔ **in detail**
  data, architecture, optimization
- ➔ **results**
  hyperparameters, fun videos

background and context

material

we wish to locate separate phases voxel by voxel

glass fiber-reinforced composite

- 3 phases
  - PolyAmide 66 matrix
  - Glass fiber reinforcement
  - Porosities (holes) in the matrix
- specimen size: 2mm x 2mm x 6mm
- voxel size: 1.3 μm^3

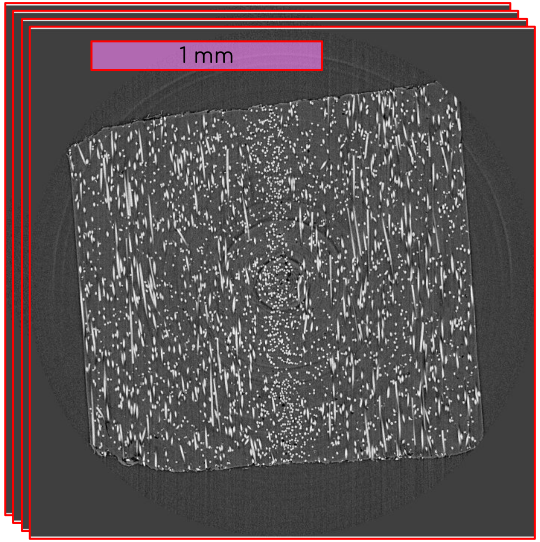

2048 x 2048 x 2048 ≈ 9 * 10^9 voxels

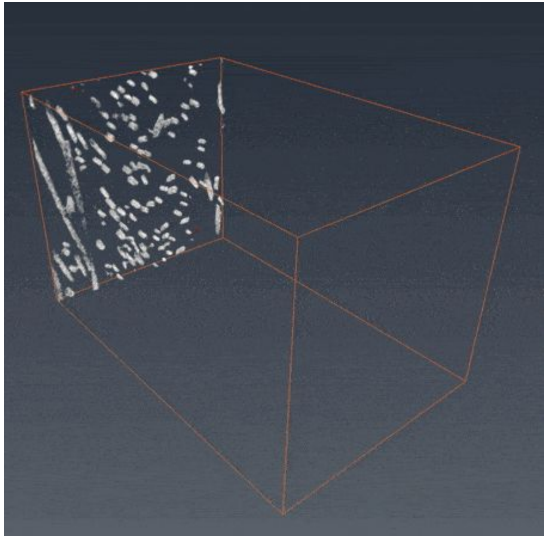

glass fibers (white)
and porosities (red)
link to the video

u-net

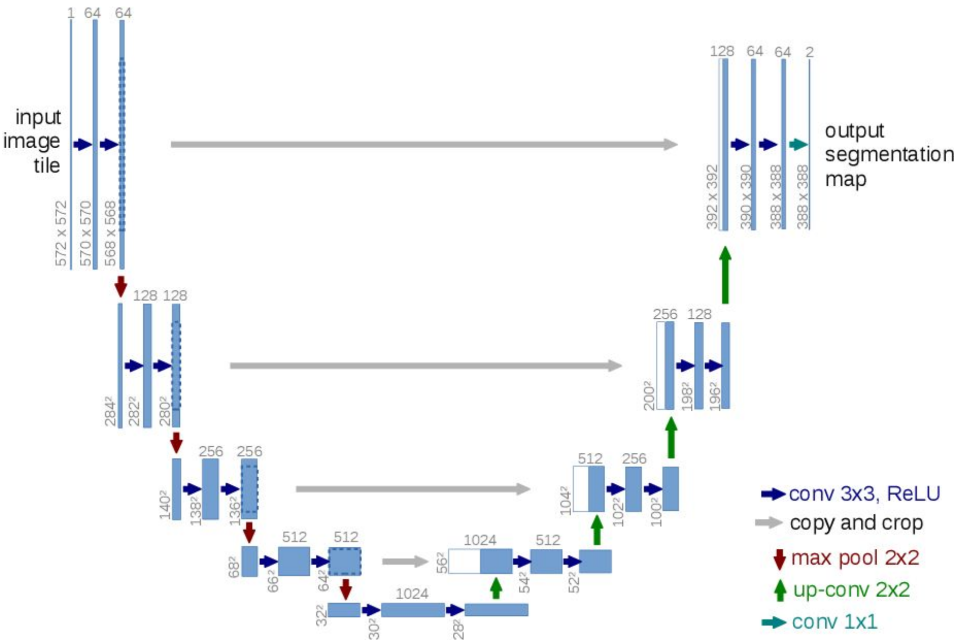

U-Net: Convolutional Networks for Biomedical Image Segmentation.

Ronneberger, Olaf & Fischer, Philipp & Brox, Thomas.

University of Freiburg, Germany.

DOI: [10.1007/978-3-319-24574-4_28]

International Conference on Medical Image Computing and Computer-Assisted Intervention (2015).

cited 18,000 times (January 2021, source: Semantic Scholar)

3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation

Çiçek, Özgün & Abdulkadir, Ahmed & Lienkamp, Soeren & Brox, Thomas & Ronneberger, Olaf.

University of Freiburg, Germany.

DOI: 10.1007/978-3-319-46723-8_49

International Conference on Medical Image Computing and Computer-Assisted Intervention (2016).

cited 1,900 times (January 2021, source: Semantic Scholar)

which one is better?

overview

Annotate, train, evaluate

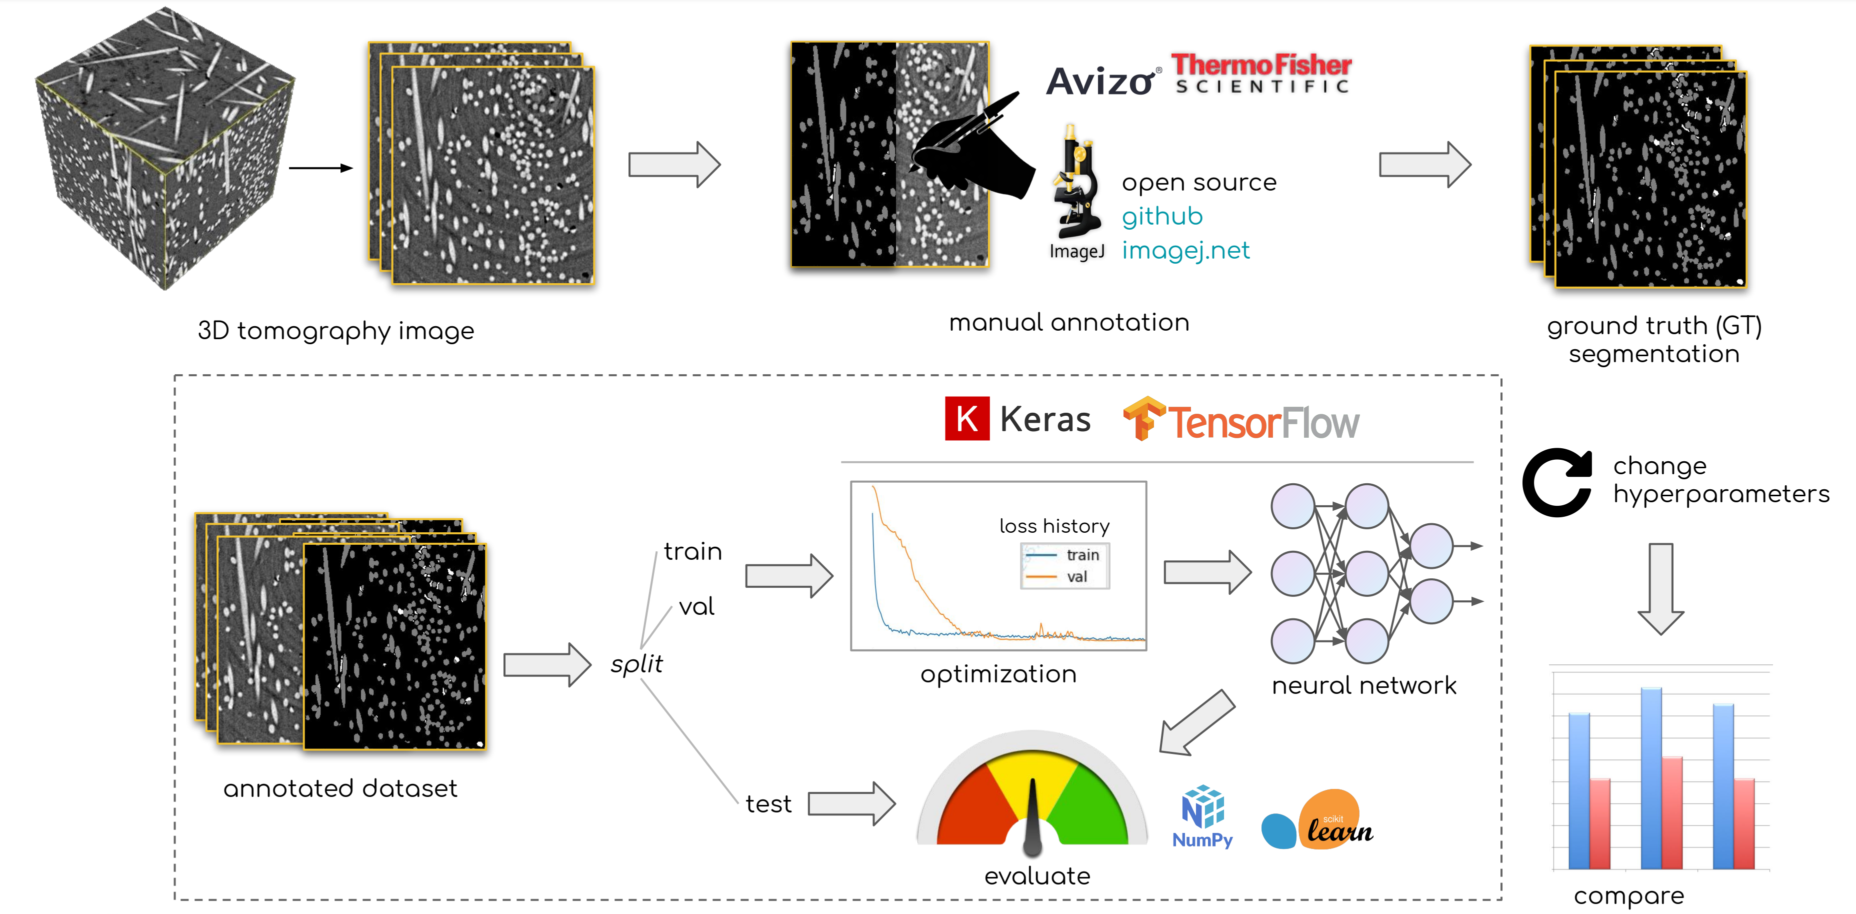

in detail

data annotation

phase 1
double Seeded Region Growing (fiji)

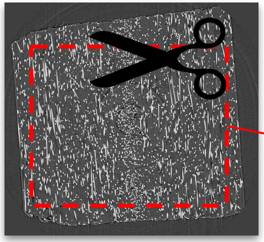

original volume
2048 x 2048 x 2048

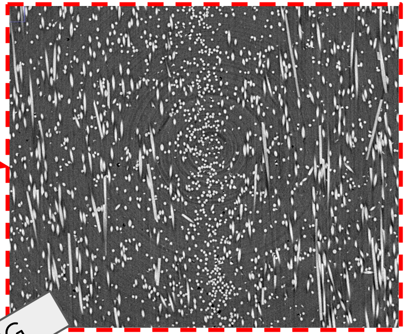

1300 x 1040 x 1900

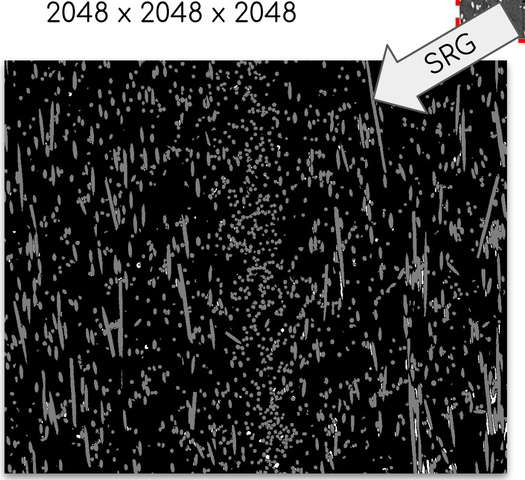

ground truth
$v_{i,j,k} \in \{0,1,2\}$

phase 2
artifacts correction (avizo)

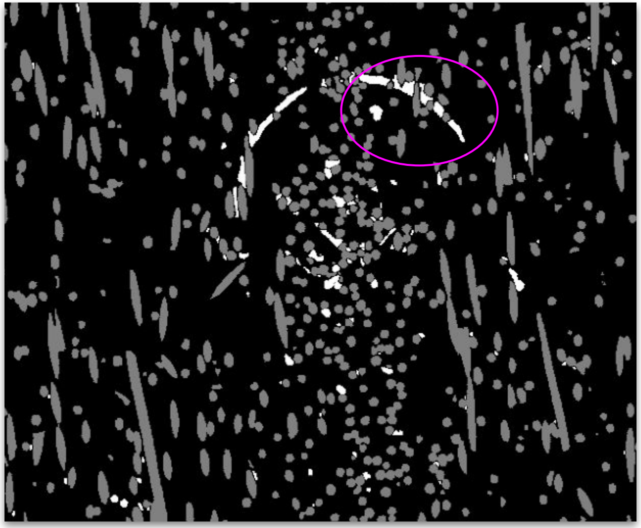

issue	mitigation strategy
• ring artefacts	• find large, long 2d blobs*
• ill-defined porosity boundary	• manually alter the voxels

** "blob": connected set*

data split & class imbalance

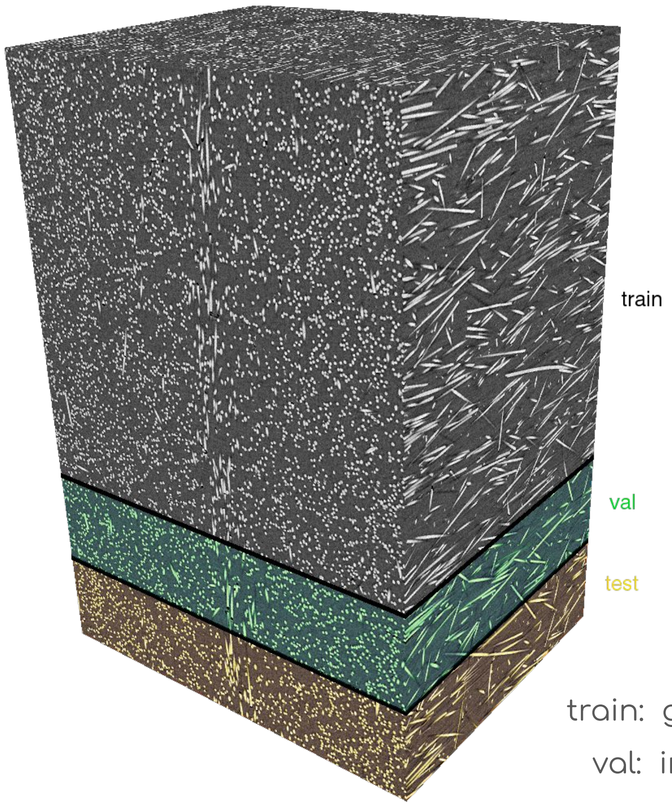

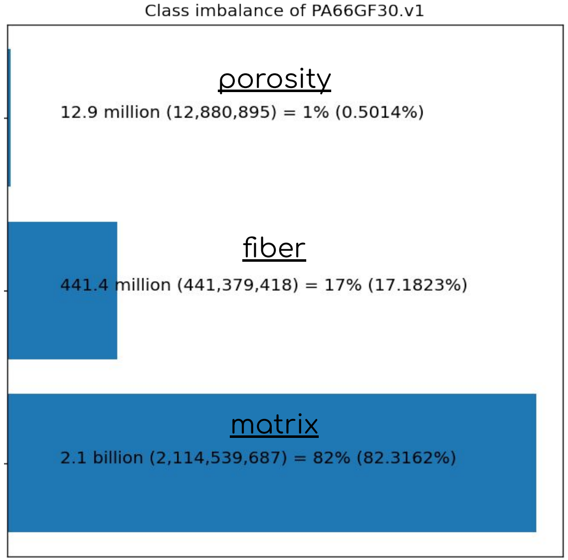

train: gradient descent

val: in-the-loop model selection

test: evaluation

set	nb. layers (1300x1040)	nb. voxels (millions)	proportion
train	1300	1.800	68%
val	128	170	9%
test	300	400	16%

data augmentation

1. random 3D crop
2. geometric transformation
3. value shift

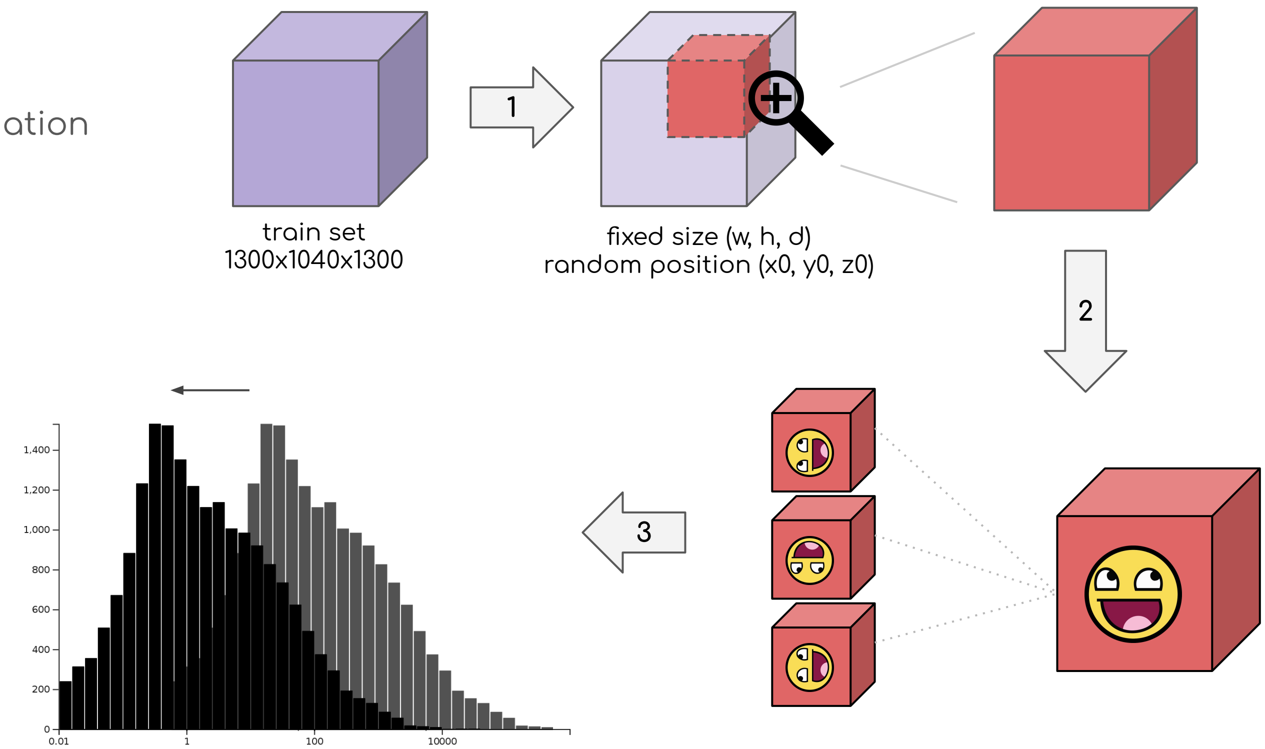

value shift
add a constant to every voxel in the crop

random transformation
(rotate 90°, flip, transpose, combinations...)

8 possibilities in 2d
50 possibilities in 3d

modular u-net

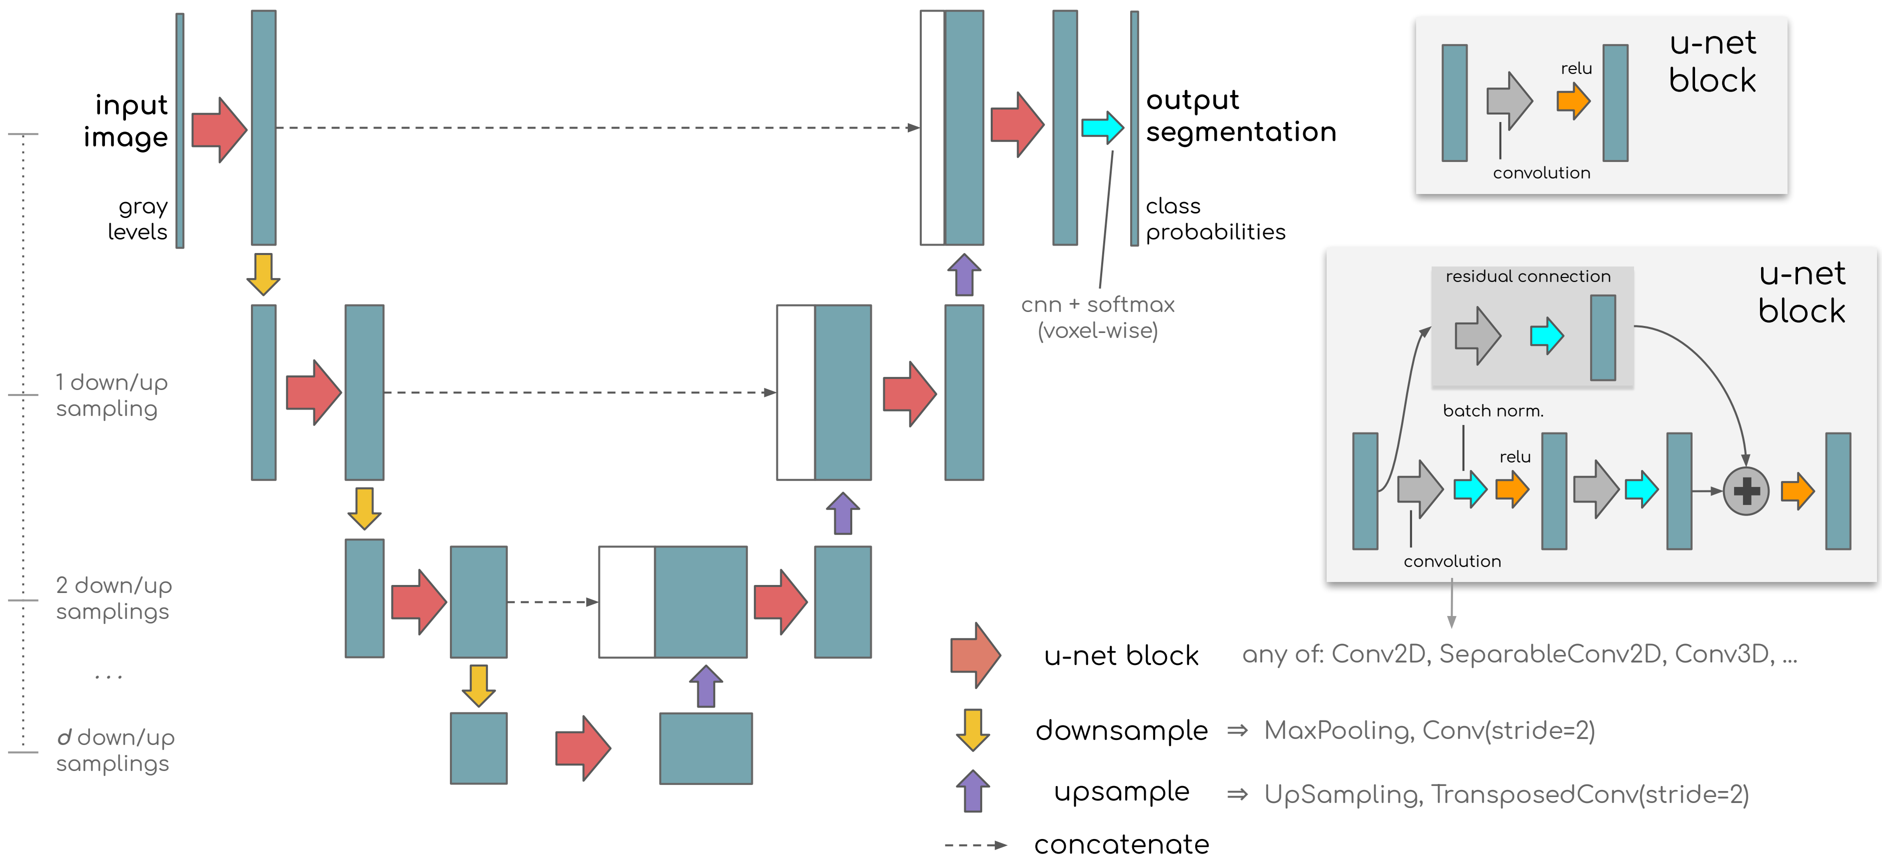

model variations

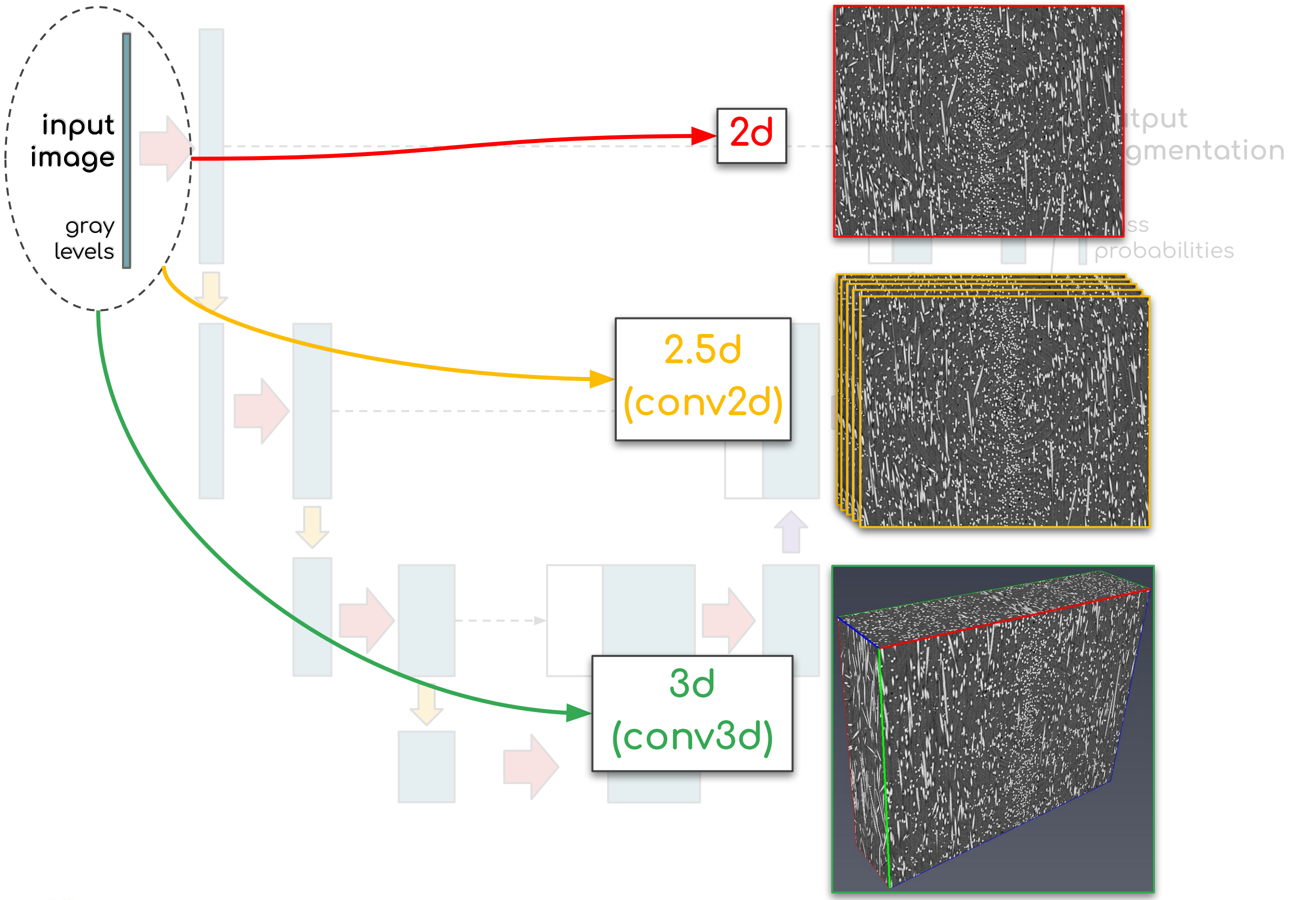

optimization

computational resources:

- 1 machine: 2 x Quadro P4000 (2 x 8Gb) *...thanks Centre des Matériaux!*
- 3 machines: 1 x Quadro P2000 (5Gb)

optimizer: adam (Keras's implementation)

loss function: (adapted) jaccard index

$$J = \frac{|A \cap B|}{|A \cup B|}$$

ref: all voxels classified as matrix ⇒ L = 32%

learning rate: a triangular-like shaped schedule

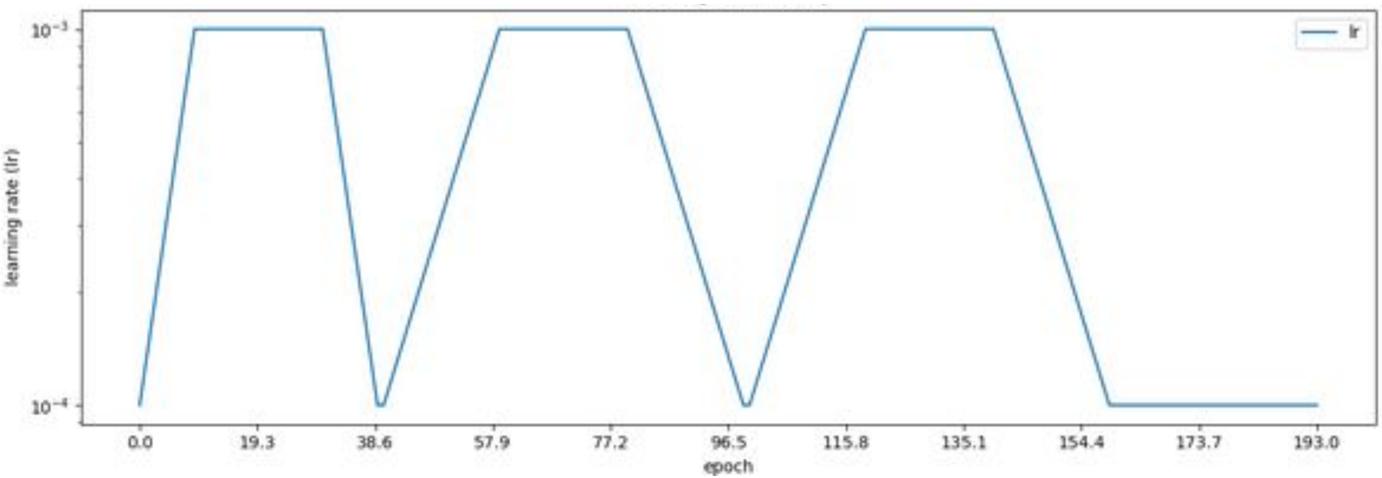

optimization

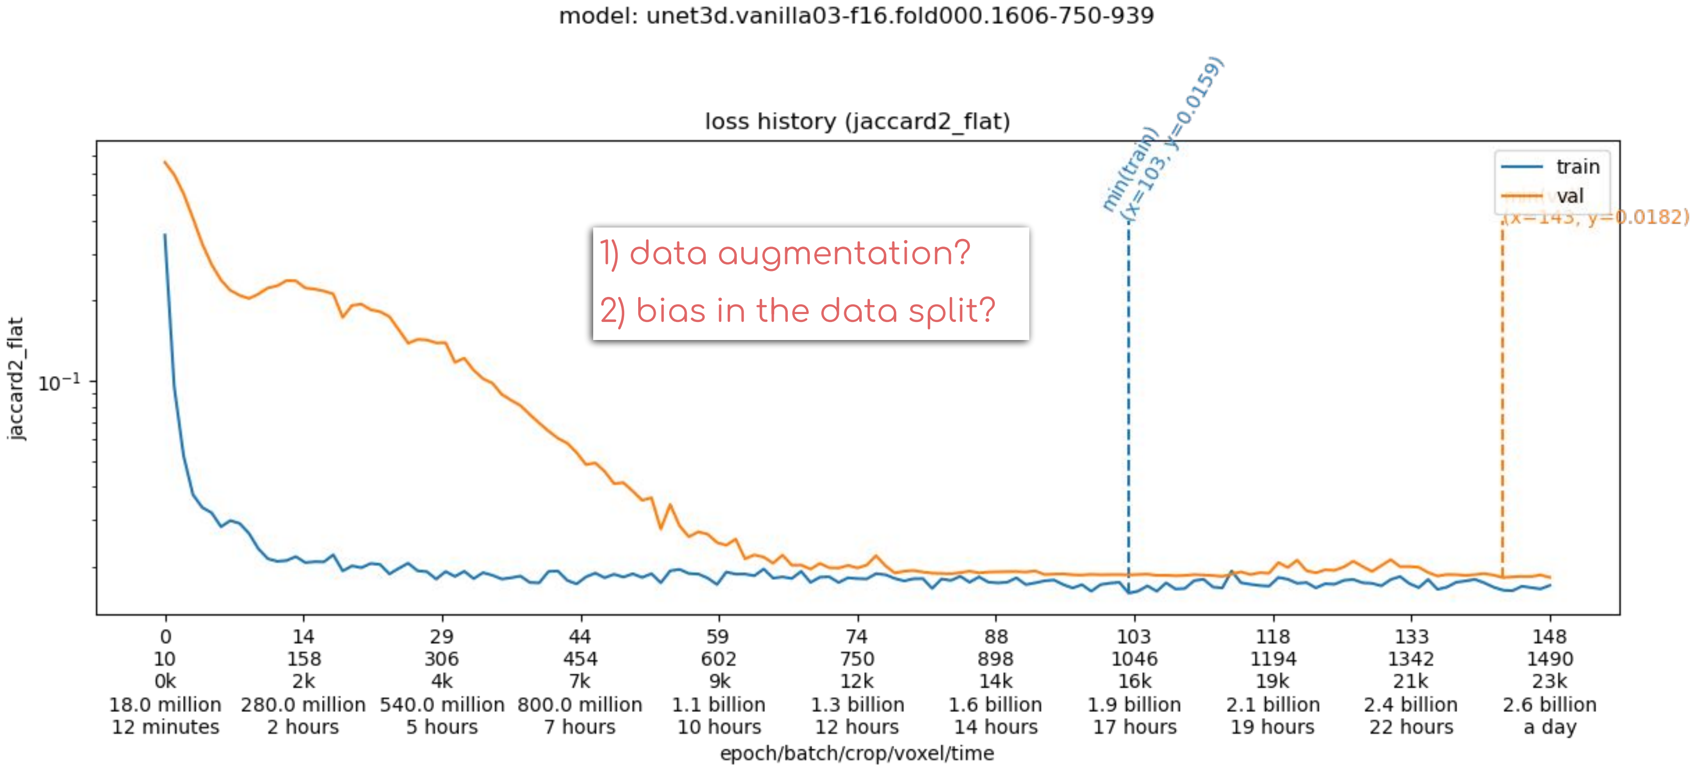

training time: 8 ~ 24h

results

blue: glass fibers

yellow: porosities

link to the video

test set (300 slices, 400M voxels)

processing time: 2~7 min

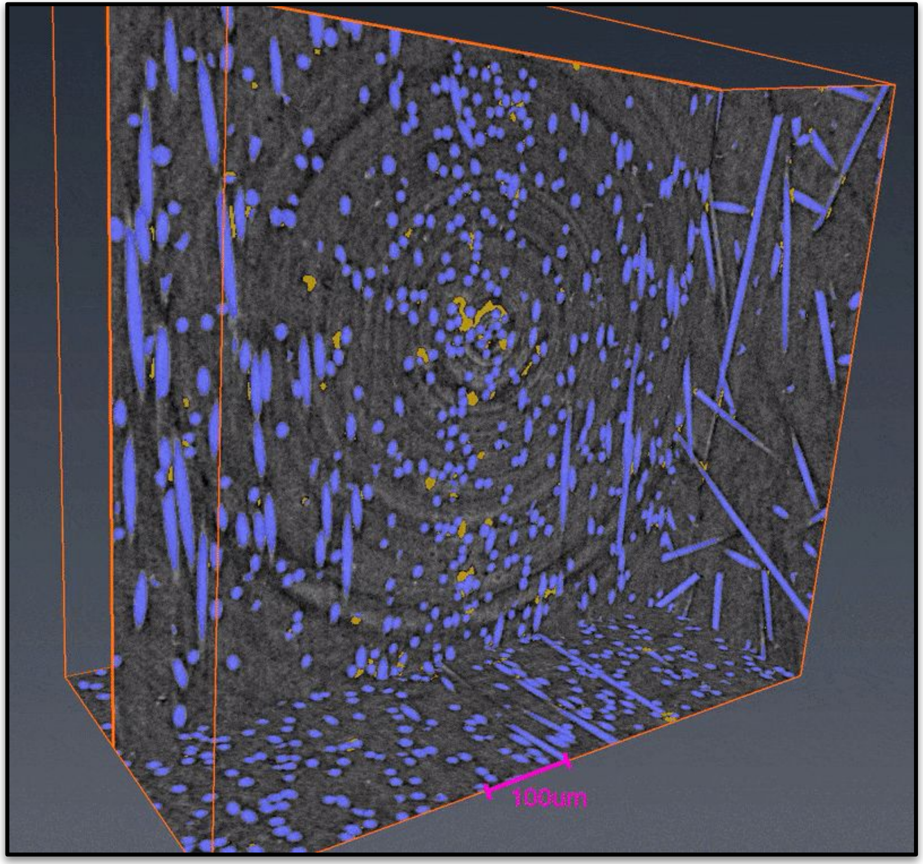

blue: glass fibers

yellow: porosities

test set (300 slices, 400M voxels)

processing time: 2~7 min

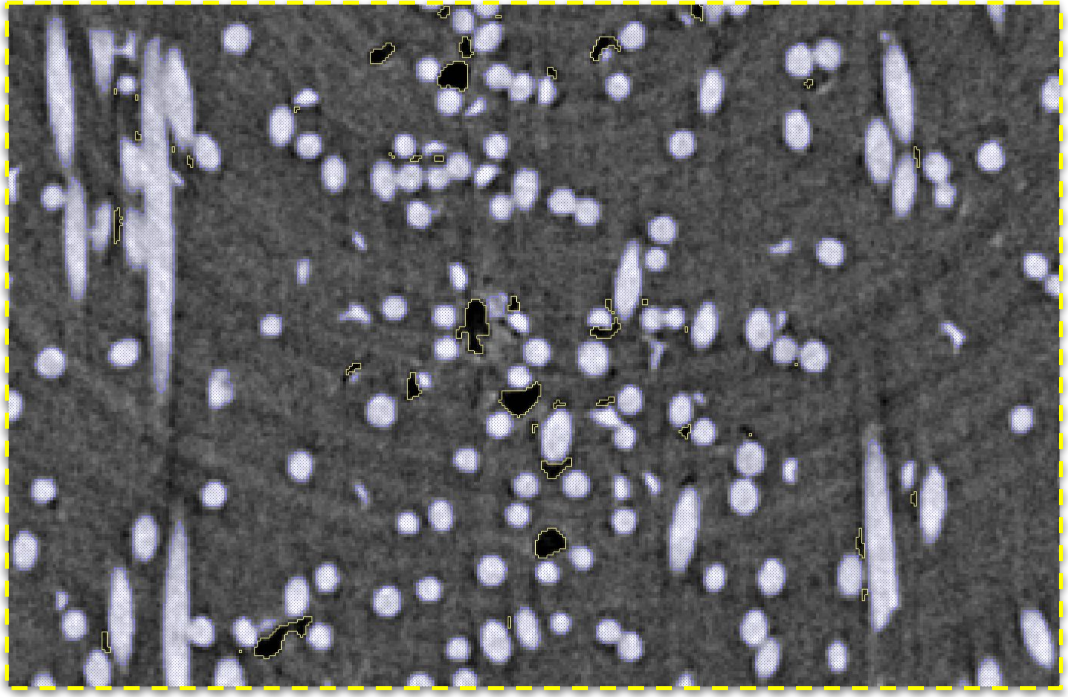

results

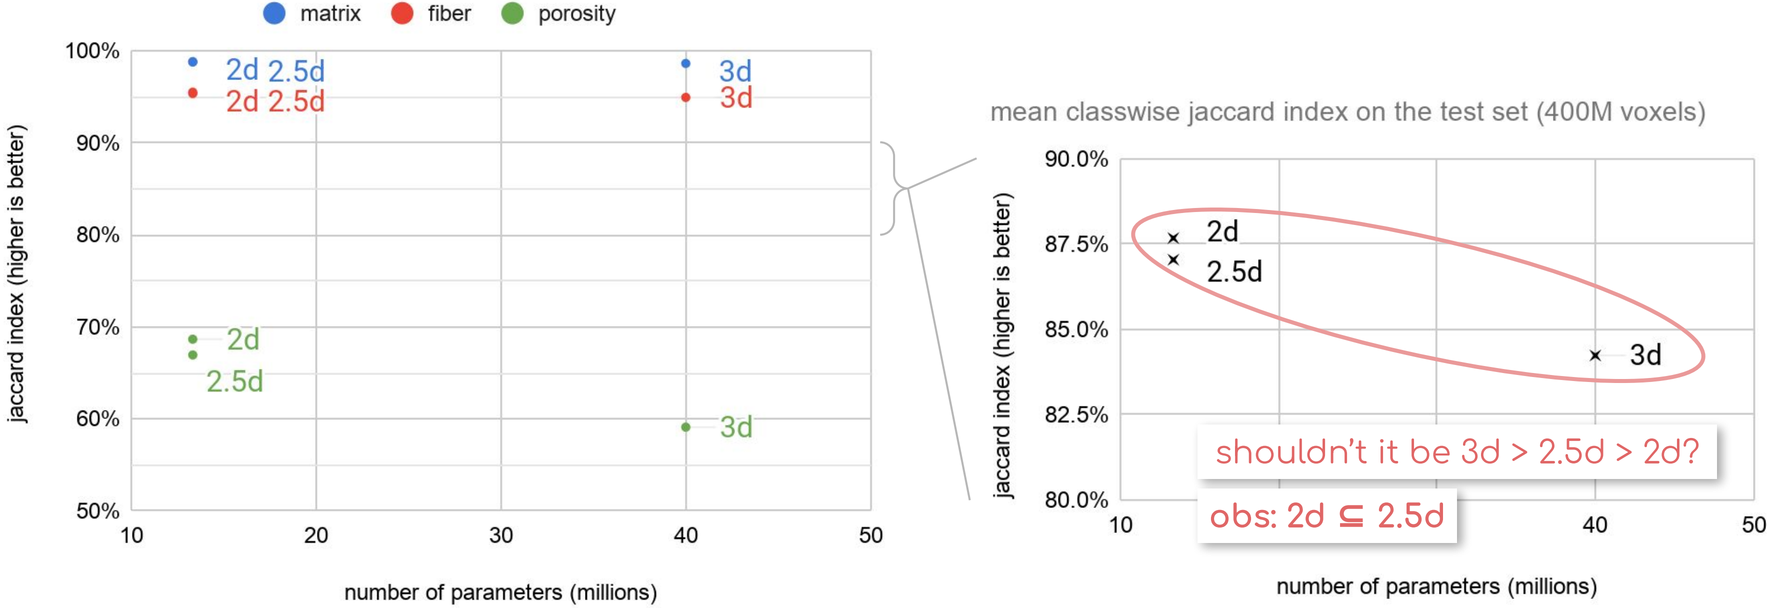

link to the chart

link to the chart

	disk size
2d/2.5d	150 Mb
3d	460 Mb

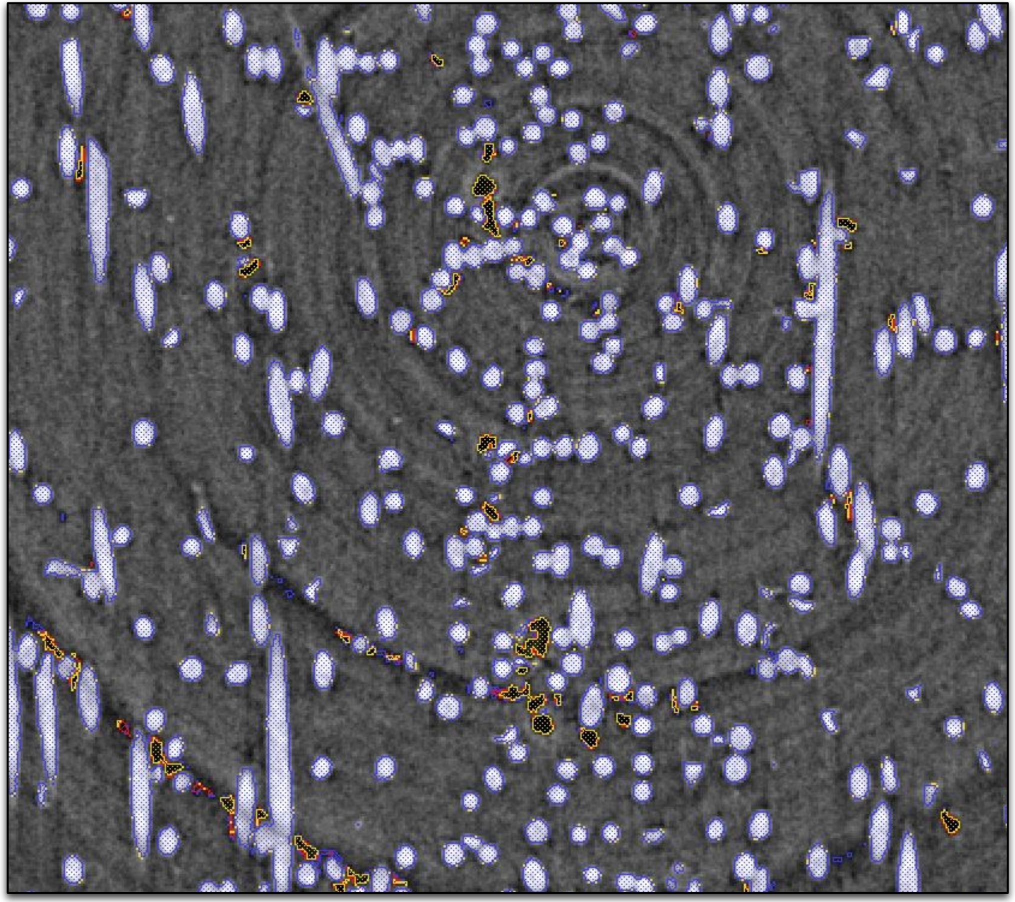

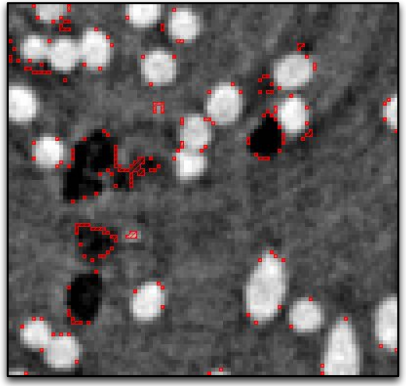

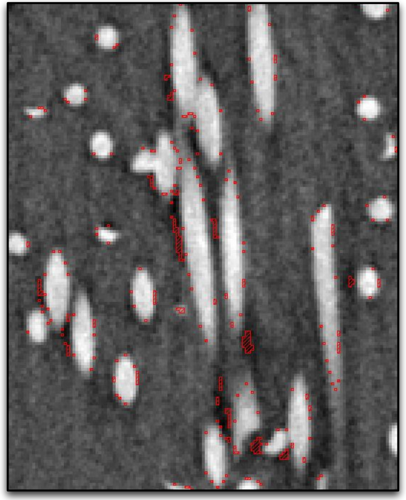

blue: glass fibers

yellow: porosities

red: errors

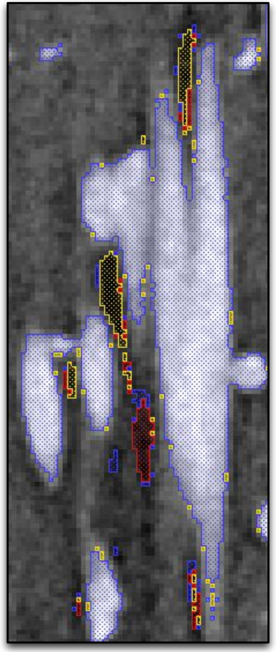

other volumes

	dimension	≈ nb. voxels	processing time
'sister'	2048 x 2048 x 2048	9B	75'
bending test	1842 x 1766 x 1946	6B	40'
biaxial test	1579 x 1845 x 2002	6B	32'
original (manual)	*1300 x 1040 x 1900*	*3B*	*12 + 20 person-hour*

2x Quadro P4000

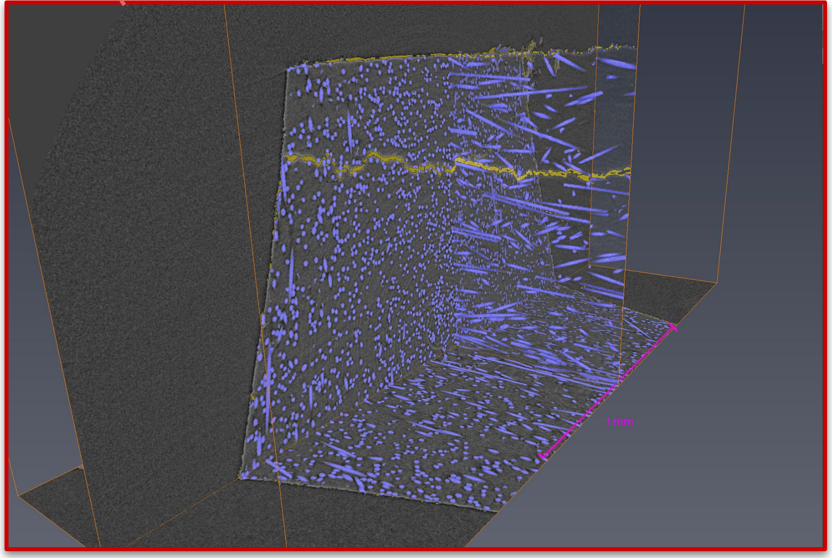

blue: glass fibers

yellow: ~~porosities~~ crack

link to the video

*the outside and the crack are **undefined** phases*

learning curve

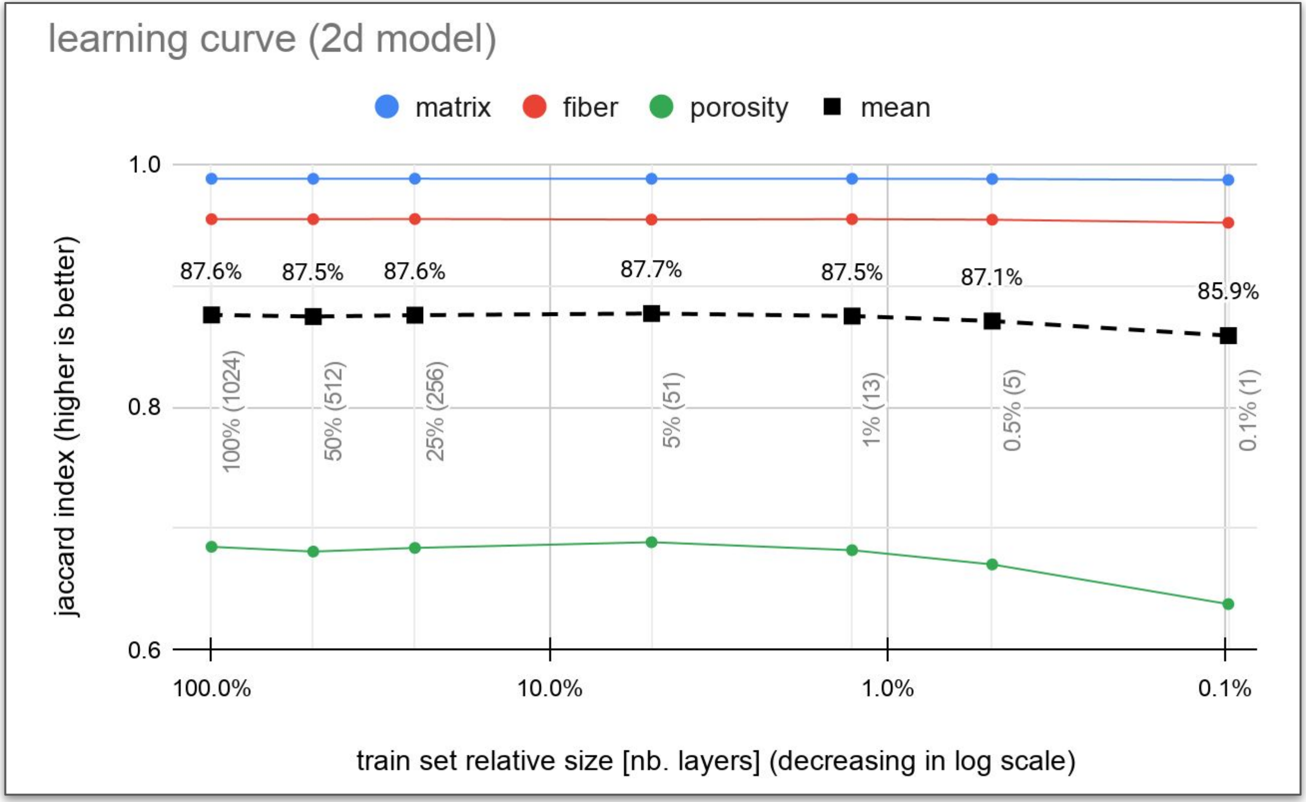

link to the chart

thanks to the data augmentation?

"In many biomedical applications, only very few images are required to train a network that generalizes reasonably well."

Ö. Çiçek, A. Abdulkadir, S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation," presented at the Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2016, doi: 10.1007/978-3-319-46723-8_49.

conclusion

takeaways

⇒ satisfactory results
human-like or under-resolution errors

⇒ training
8h ~ 24h
1 ~ 10 annotated layers (~1000 x 1000 pixels)

⇒ processing
~ 30 layers / minute (2048 x 2048 images)

⇒ model selection takeout
2d > 2.5d > 3d

next

⇒ application
segment harder cases like fractures

⇒ leverage the 2d model to build the 2.5d
construct a 2.5d model using the weights of a pre-trained 2d model

⇒ funny val loss curve
investigate why the validation is “late”

thank you for your attention!

Copil BIGMÉCA January 2021, Monday 2021-01-13

João P C Bertoldo

supervisors

Henry Proudhon Etienne Decencière David Ryckelynck

extras

jaccard2

jaccard index

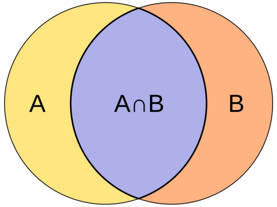

$$J = \frac{|A \cap B|}{|A \cup B|} = \frac{|A \cap B|}{|A| + |B| - |A \cap B|}$$

$$J \in [0, 1]$$

issue: it doesn't fit

for each voxel:

	matrix k=1	fiber k=2	porosity k=3
ground truth (y)	0	0	1
model's output (ỹ)	.10	.10	.80

solution: generalization of J for K classes

$$I = \sum_{k=1}^{K} y_k \, \tilde{y}_k \qquad J_2 = \frac{I}{\sum_{k=1}^{K} y_k^2 + \sum_{k=1}^{K} \tilde{y}_k^2 - I}$$

Shouldn't it be a <u>loss</u>?

$$\text{loss} = 1 - J_2$$

reference values

theoretical model	**jaccard2**
all probabilities 1/3	67%
all voxels classified as matrix	32%

2.5d variation: shared encoder-decoder

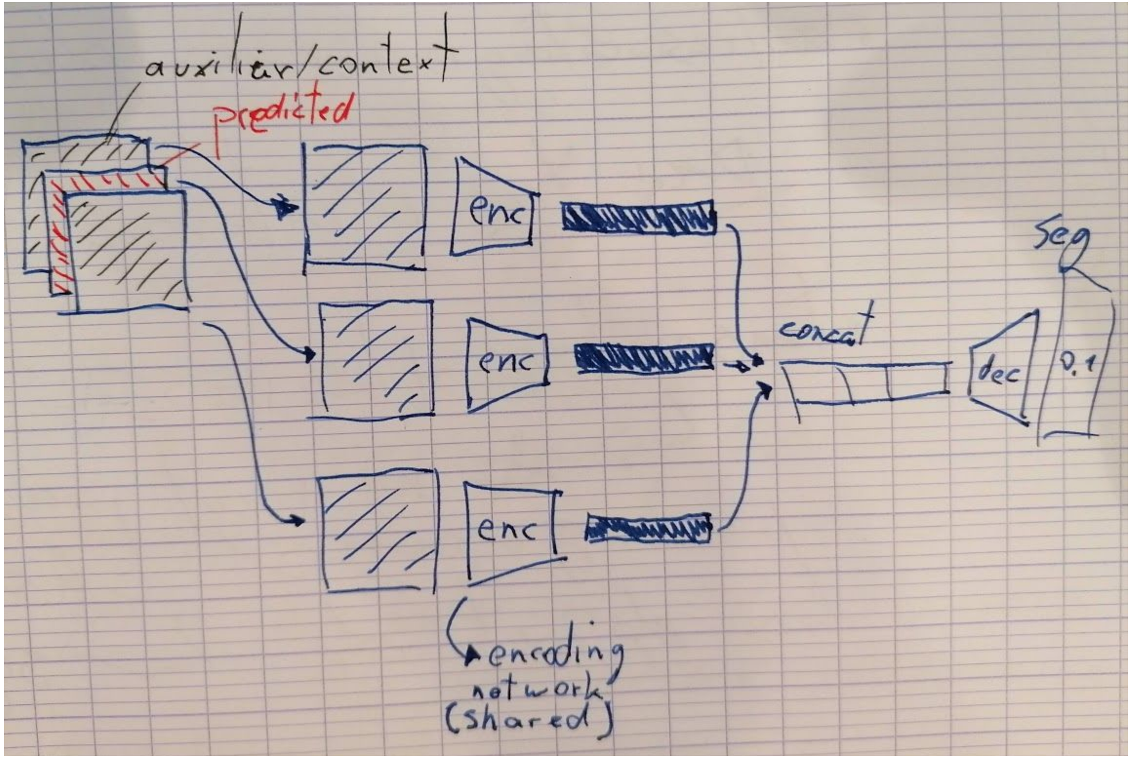

model evolution (model 2d)

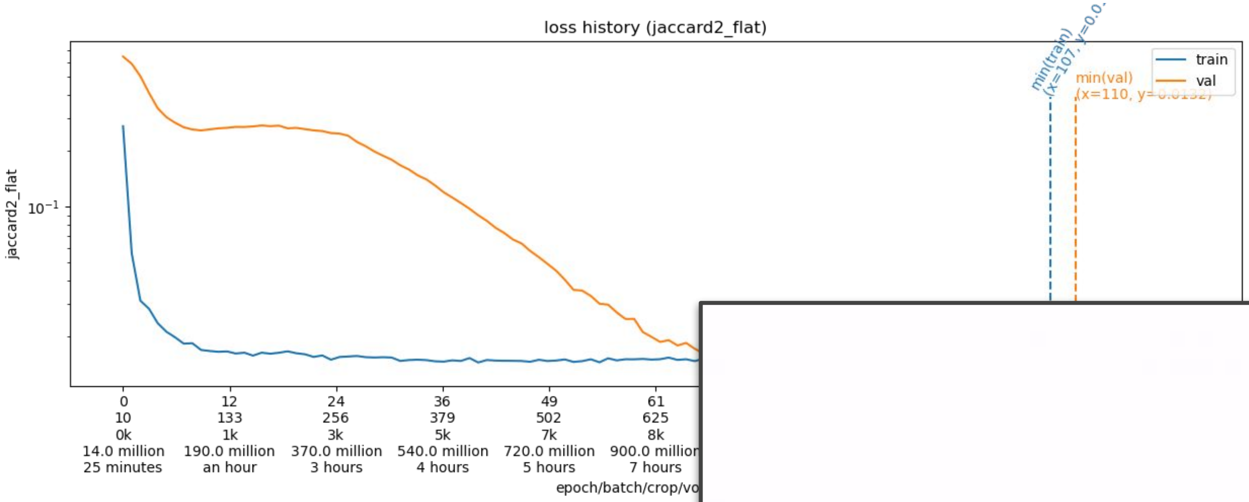

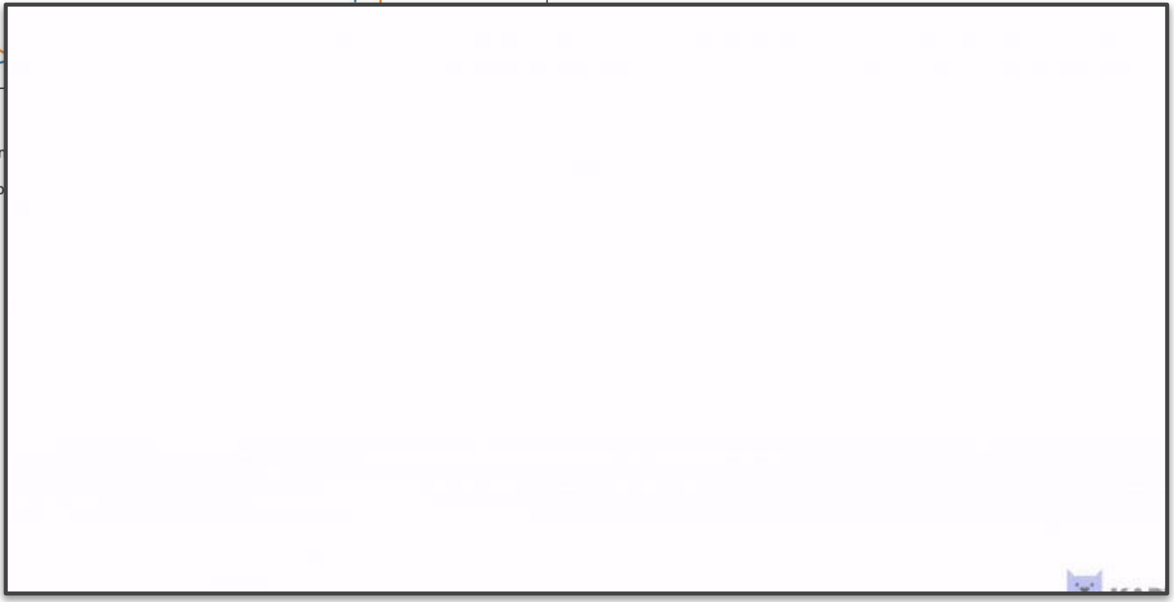

black: matrix

gray: fibers

white: porosities

link to the video

crack volume

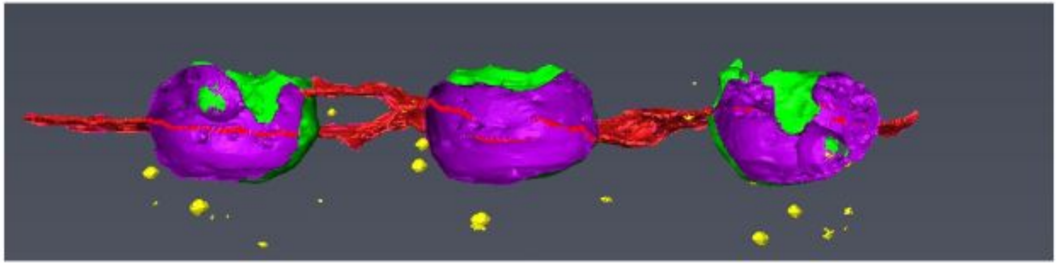

(a)

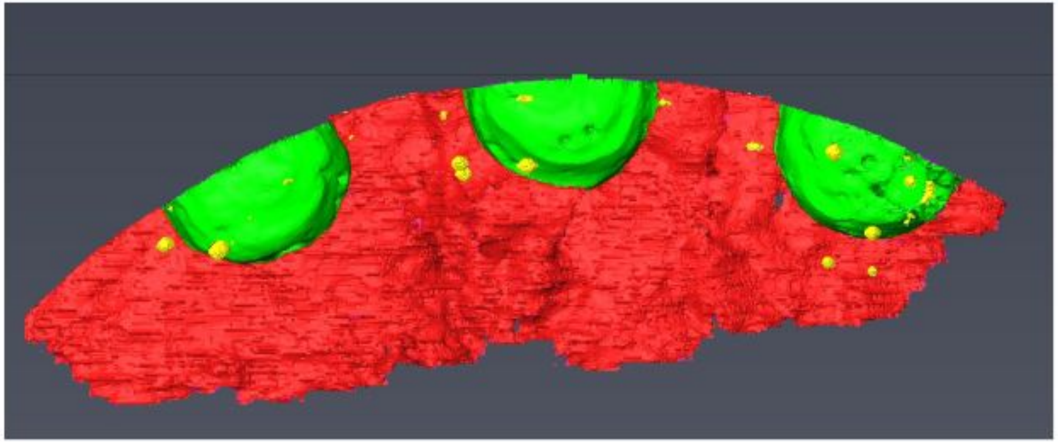

(b)

J. Ajuelos, L. Marcin, C. Montebello, V. Maurel, and H. Proudhon, “Caractérisation et modélisation du rôle des défauts microstructuraux sur les propriétés en fatigue de superalliage base nickel élaboré par fabrication additive,” p. 71.

crack volume

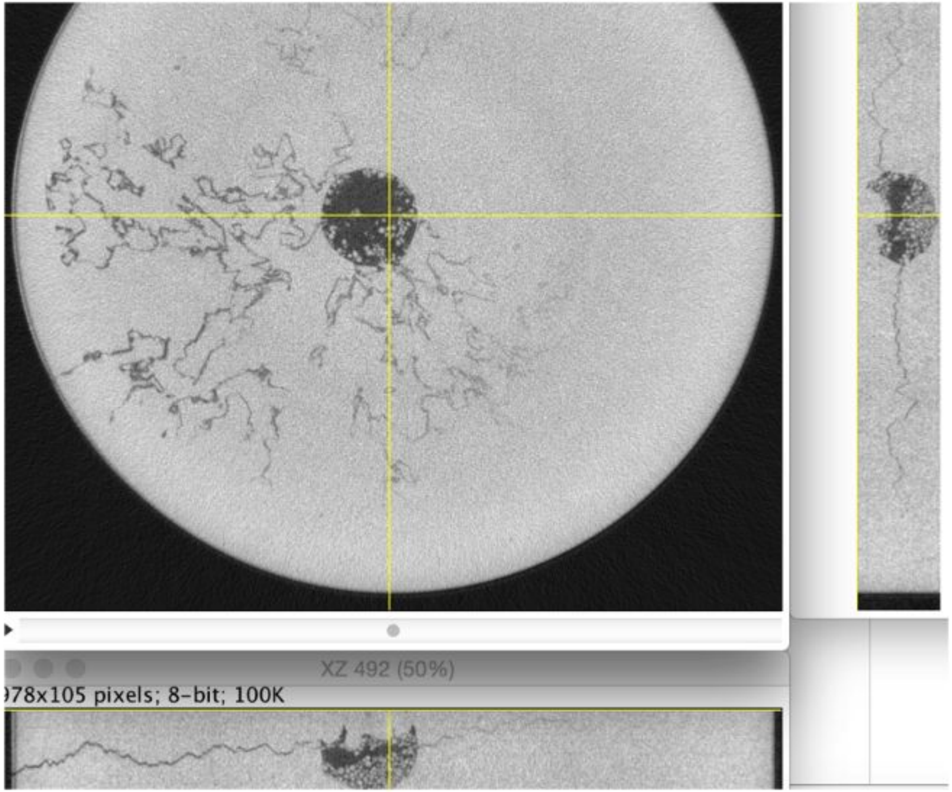

J. Ajuelos, L. Marcin, C. Montebello, V. Maurel, and H. Proudhon, “Caractérisation et modélisation du rôle des défauts microstructuraux sur les propriétés en fatigue de superalliage base nickel élaboré par fabrication additive,” p. 71.

inside the model

crack volume

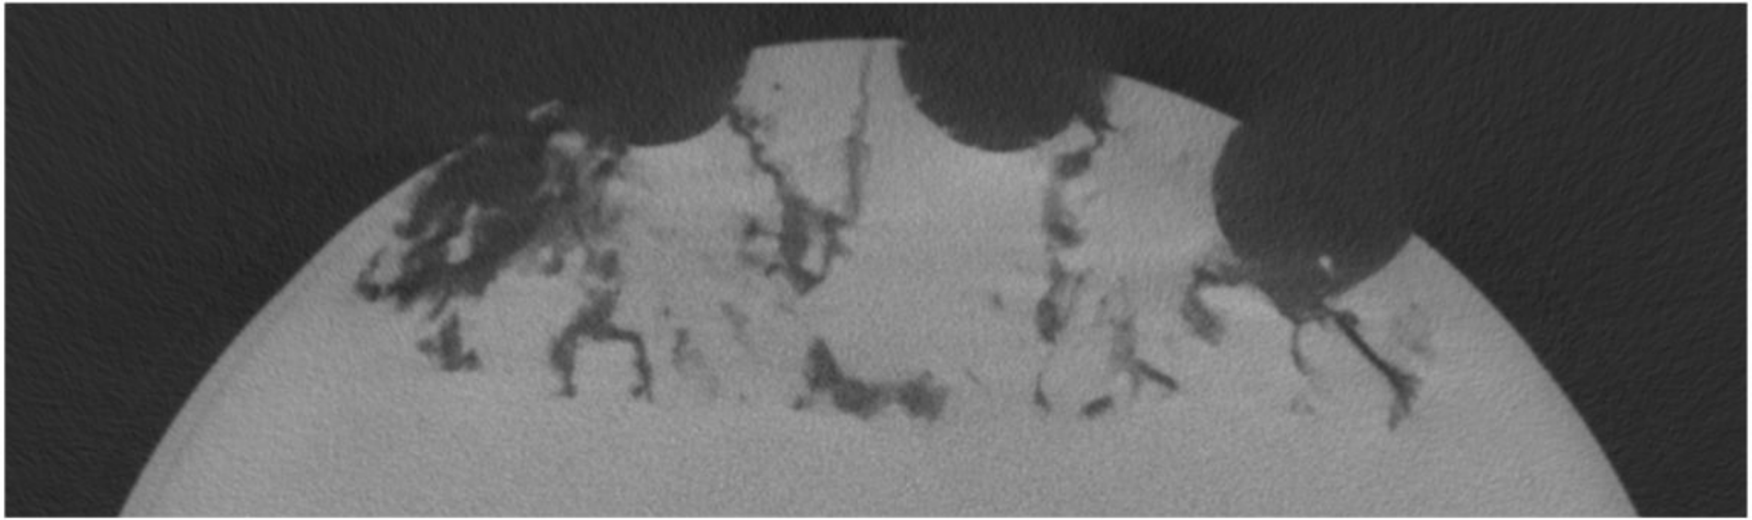

J. Ajuelos, L. Marcin, C. Montebello, V. Maurel, and H. Proudhon, “Caractérisation et modélisation du rôle des défauts microstructuraux sur les propriétés en fatigue de superalliage base nickel élaboré par fabrication additive,” p. 71.

ablation

aux

model: (a) modular u-net

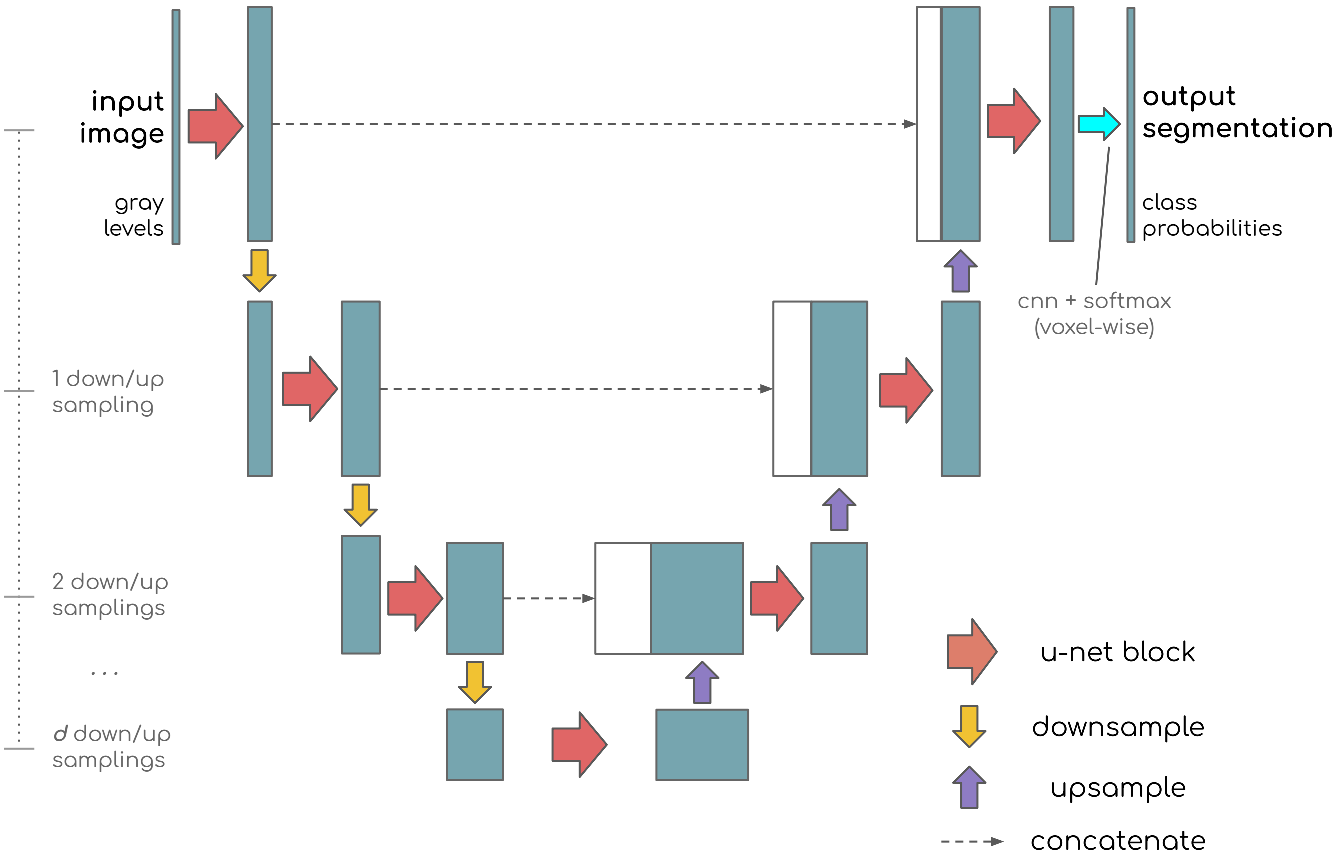

https://www.researchgate.net/publication/258554113_Utilisation_de_la_Correlation_d'Images_Numeriques_et_de_la_Methode_de_l'Ecart_a_l'eQuilibre_pour_la_caracterisation_mecanique_de_tubes_obtenus_par_enroulement_filamentaire

Crouzeix, Laurent & Périé, Jean-Noël & Torres, Mauricio & Douchin, Bernard & Collombet, Francis & Hernández-Moreno, Hilario. (2009). Utilisation de la Corrélation d'Images Numériques et de la Méthode de l'Ecart à l'éQuilibre pour la caractérisation mécanique de tubes obtenus par enroulement filamentaire.

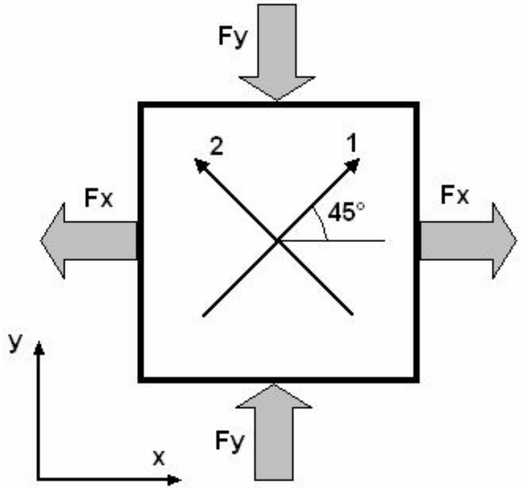

https://www.tec-science.com/material-science/material-testing/bending-flexural-test/

By tec-science -07/13

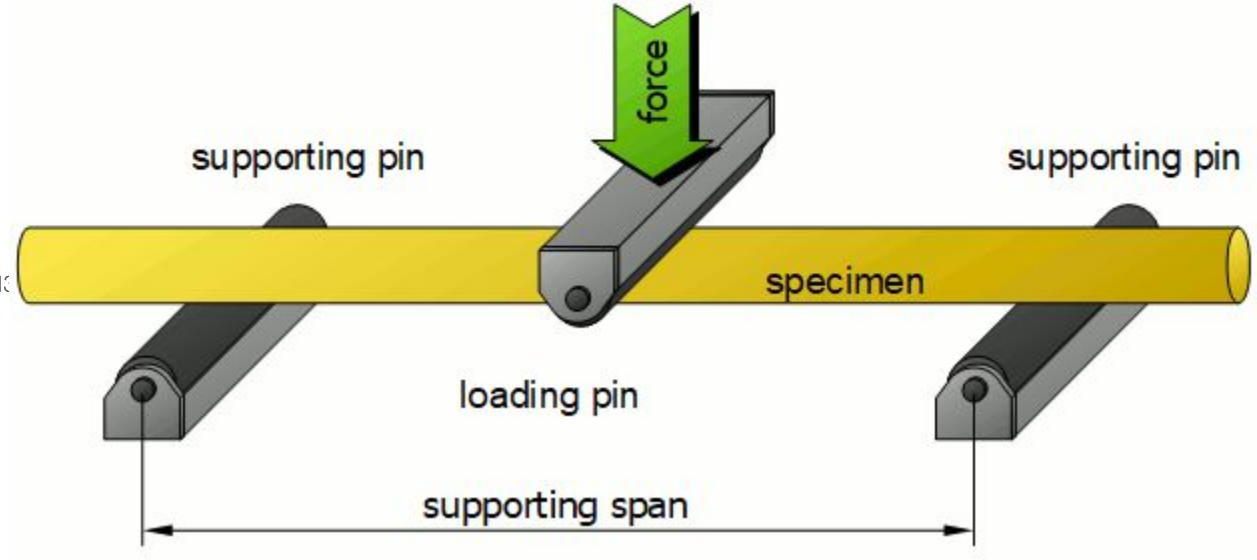

conv2d

Figures ref: [A Comprehensive Introduction to Different Types of Convolutions in Deep Learning](), [Kunlun Bai]()
Tnx, Kunlun!

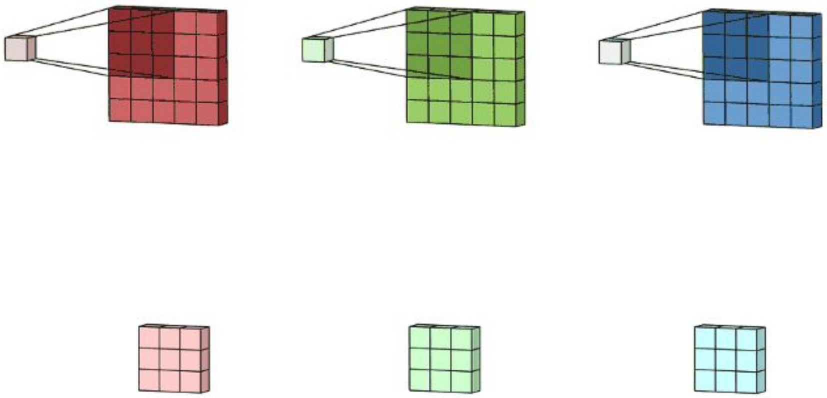

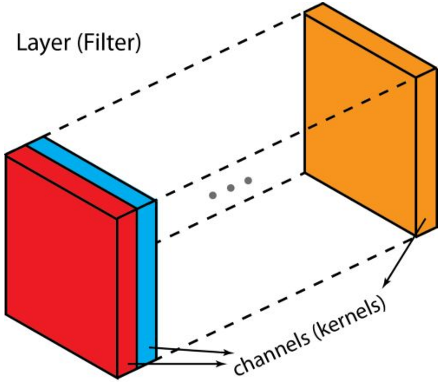

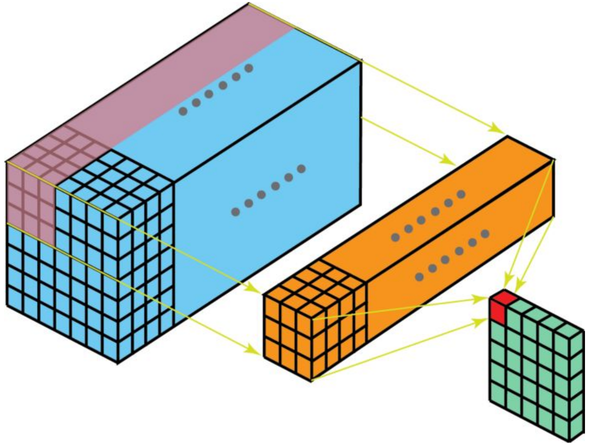

conv3d

Figures ref: A Comprehensive Introduction to Different Types of Convolutions in Deep Learning, Kunlun Bai
Tnx, Kunlun!

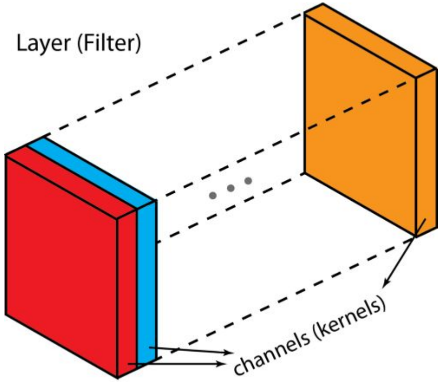

the 4th dimension now!

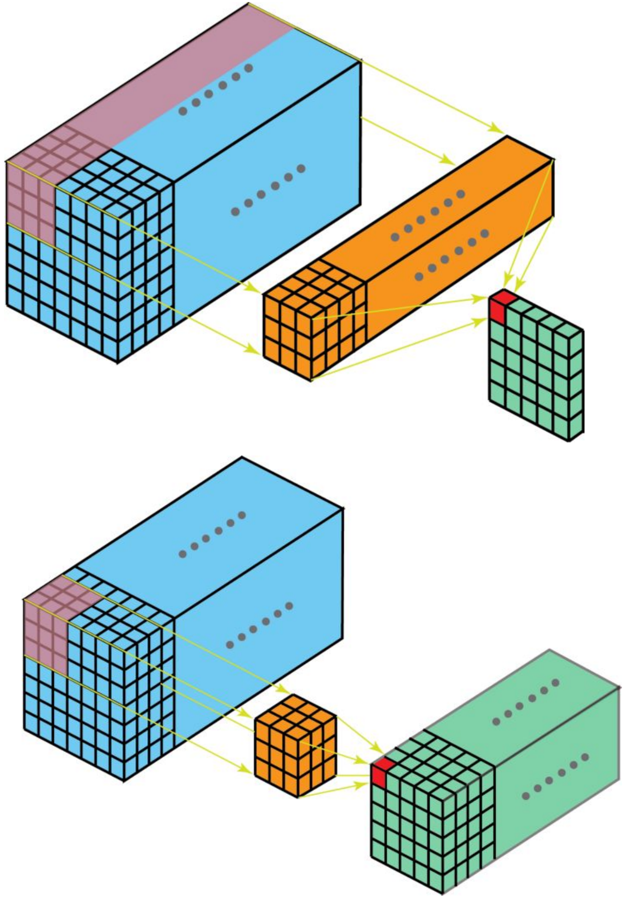

separable conv2d

Figures ref: A Comprehensive Introduction to Different Types of Convolutions in Deep Learning, Kunlun Bai

Tnx, Kunlun!

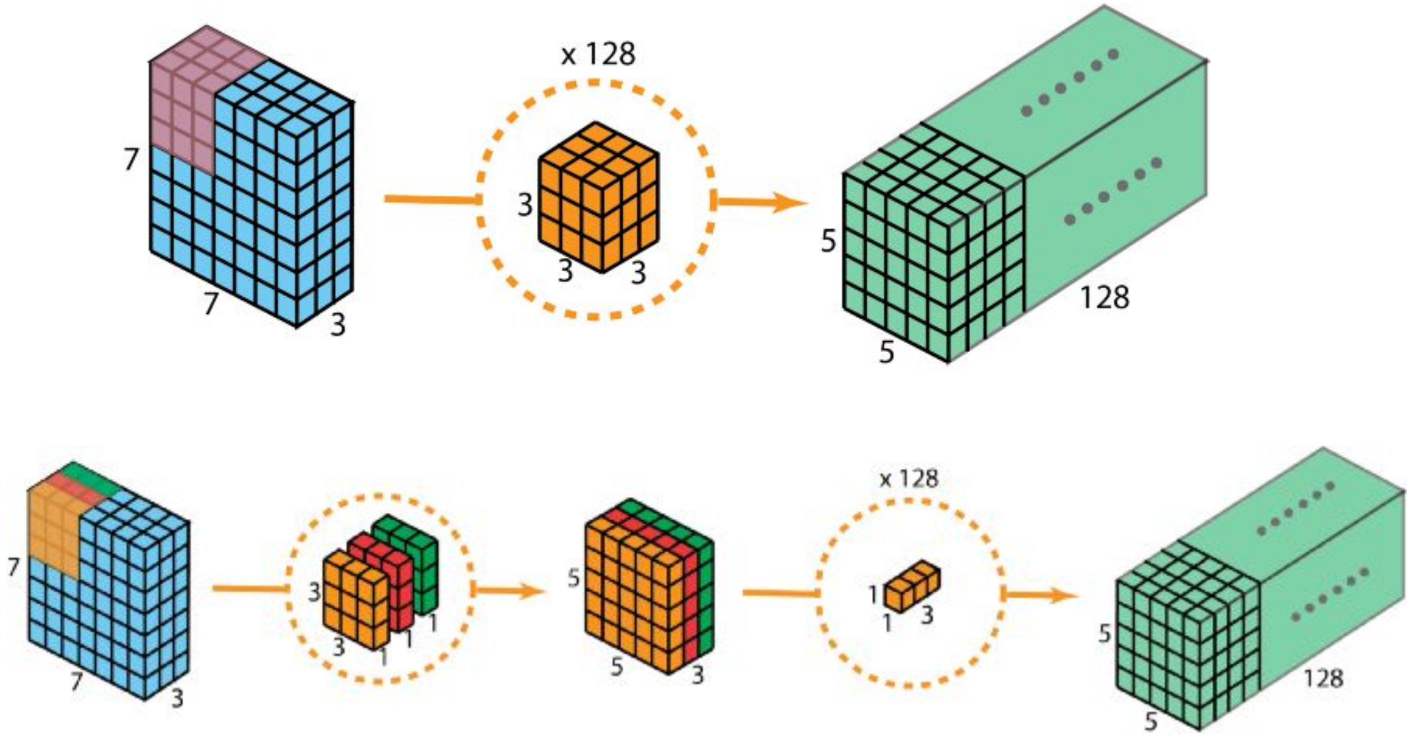

1x1 convolution

Figures ref: [A Comprehensive Introduction to Different Types of Convolutions in Deep Learning](), [Kunlun Bai]()
Tnx, Kunlun!

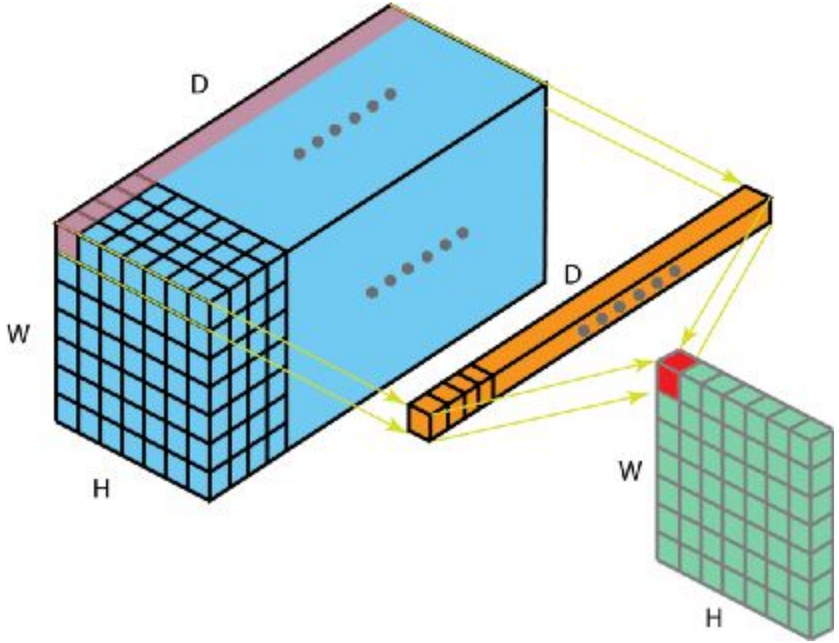

transposed conv2d

Figures ref: [A Comprehensive Introduction to Different Types of Convolutions in Deep Learning](), [Kunlun Bai]()
Tnx, Kunlun!

conv 2d

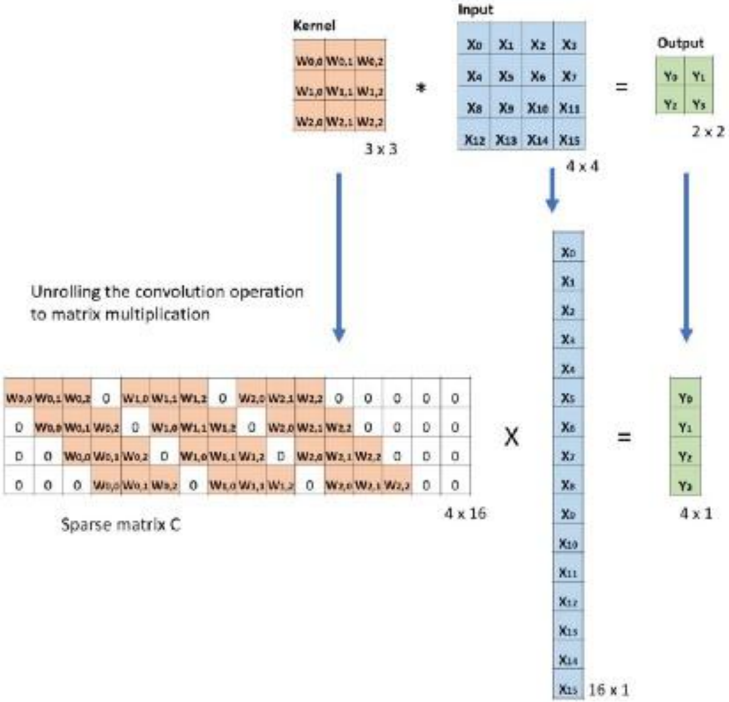

transposed conv 2d

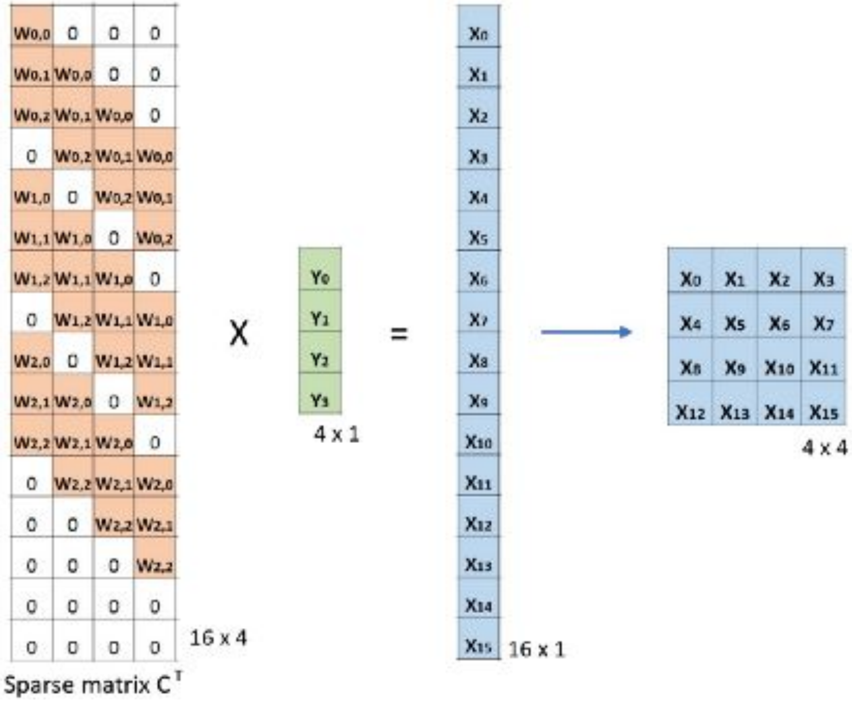

material

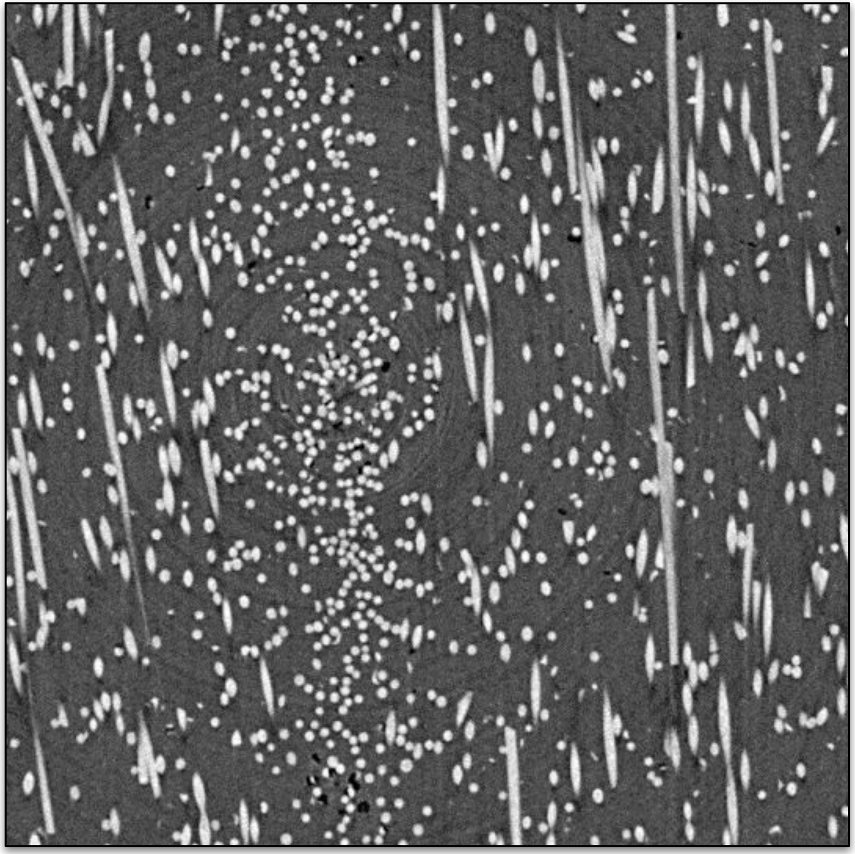

PA66GF30
PolyAmide 66
reinforced with glass fibers
(zoom)

link to the video

A few things to notice:

- *ring artifacts*
- *fibers in different directions*
- *blurred fiber-matrix interfaces*
- *porosities close to fibers*

data annotation

phase 1
double Seeded Region Growing (fiji)

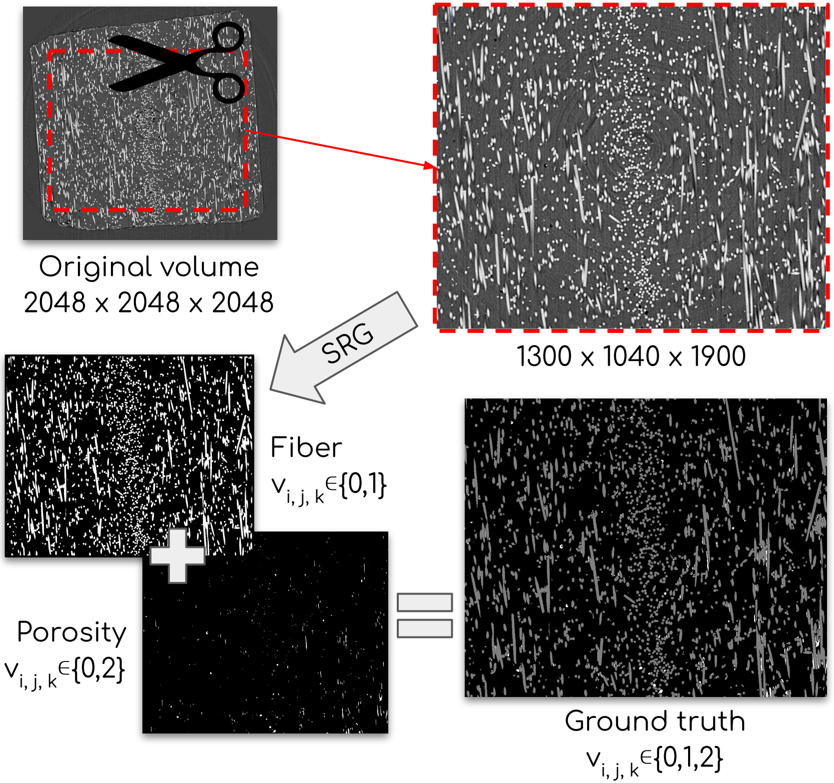

phase 2
artifacts correction (avizo)

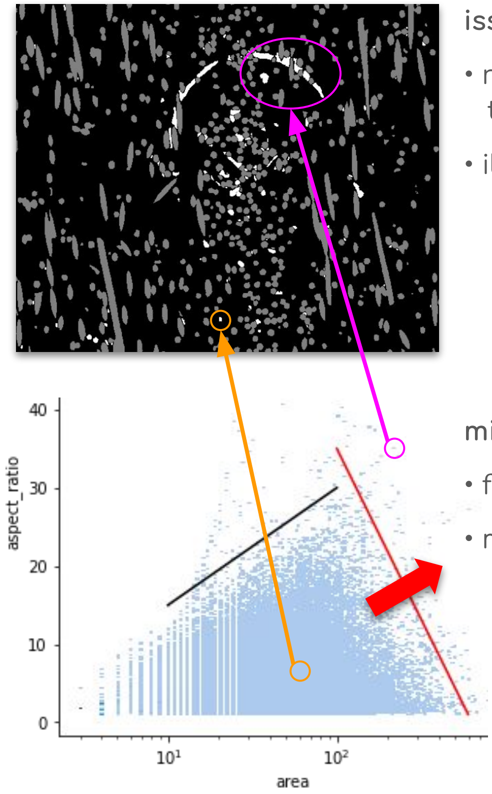

issue

- ring artefacts from the tomography reconstruction
- ill-defined porosity boundary

mitigation strategy

- find large, long 2d blobs*
- manually alter the voxels

* "blob": connected set

optimization

computational resources:

- 1 machine: 2 x Quadro P4000 (2 x 8Gb)
- 3 machines: 1 x Quadro P2000 (5Gb)

...thanks Centre des Matériaux!

optimizer: adam (Keras's implementation)

loss function: (adapted) jaccard index

$$J = \frac{|A \cap B|}{|A \cup B|}$$

ref: all voxels classified as matrix ⇒ L = 32%

epoch size: 10 batches

batch size: 2 ~ 16 samples

learning rate: a triangular-like shaped schedule

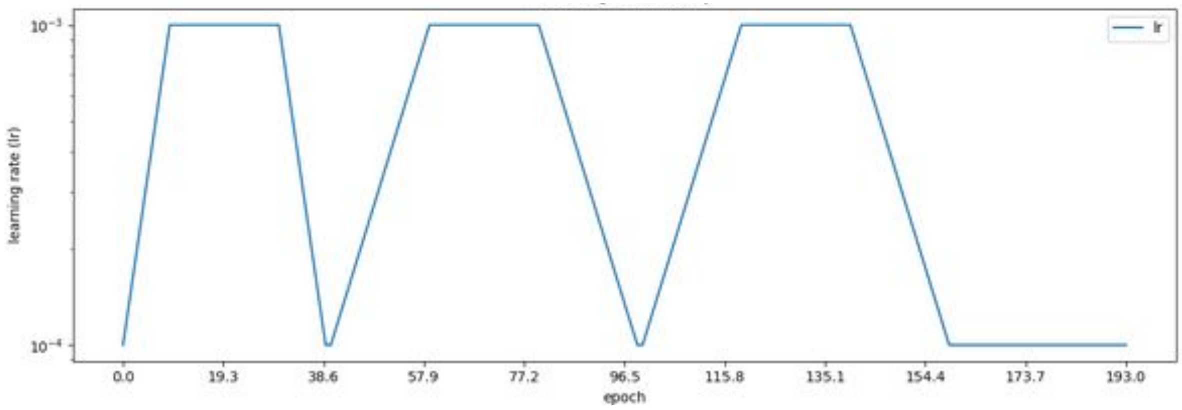

abstract

X-ray Computed Tomography (XCT) generates non-invasive 3D images, which give material scientists a means of quantitatively analyzing a material's internal structure. However, processing tomography images often demands expertise and is a tedious, time-consuming task, creating a bottleneck to scale the analysis of large volumes of data - 3D images can weight several Gigabytes. Classic mathematical morphology-based techniques can help to generate phase segmentation but depend on some level of human intervention. This presentation will show a Deep Learning approach developed at the Centre des Matériaux MINES ParisTech and its viability as an alternative for tomographic image segmentation. We compare 2D and 3D versions of a U-net-based model, respectively, using 2D and 3D convolutions, finding that, unlike expected, the 2D model is more suitable. We also evaluate the trade-off between performance and model size of several architectural variations. Qualitative results show that our models can process, in only 30 minutes, a 6-billion-voxel tomography with human-like quality. Finally, we show that this architecture can achieve such results only using five tomography layers.