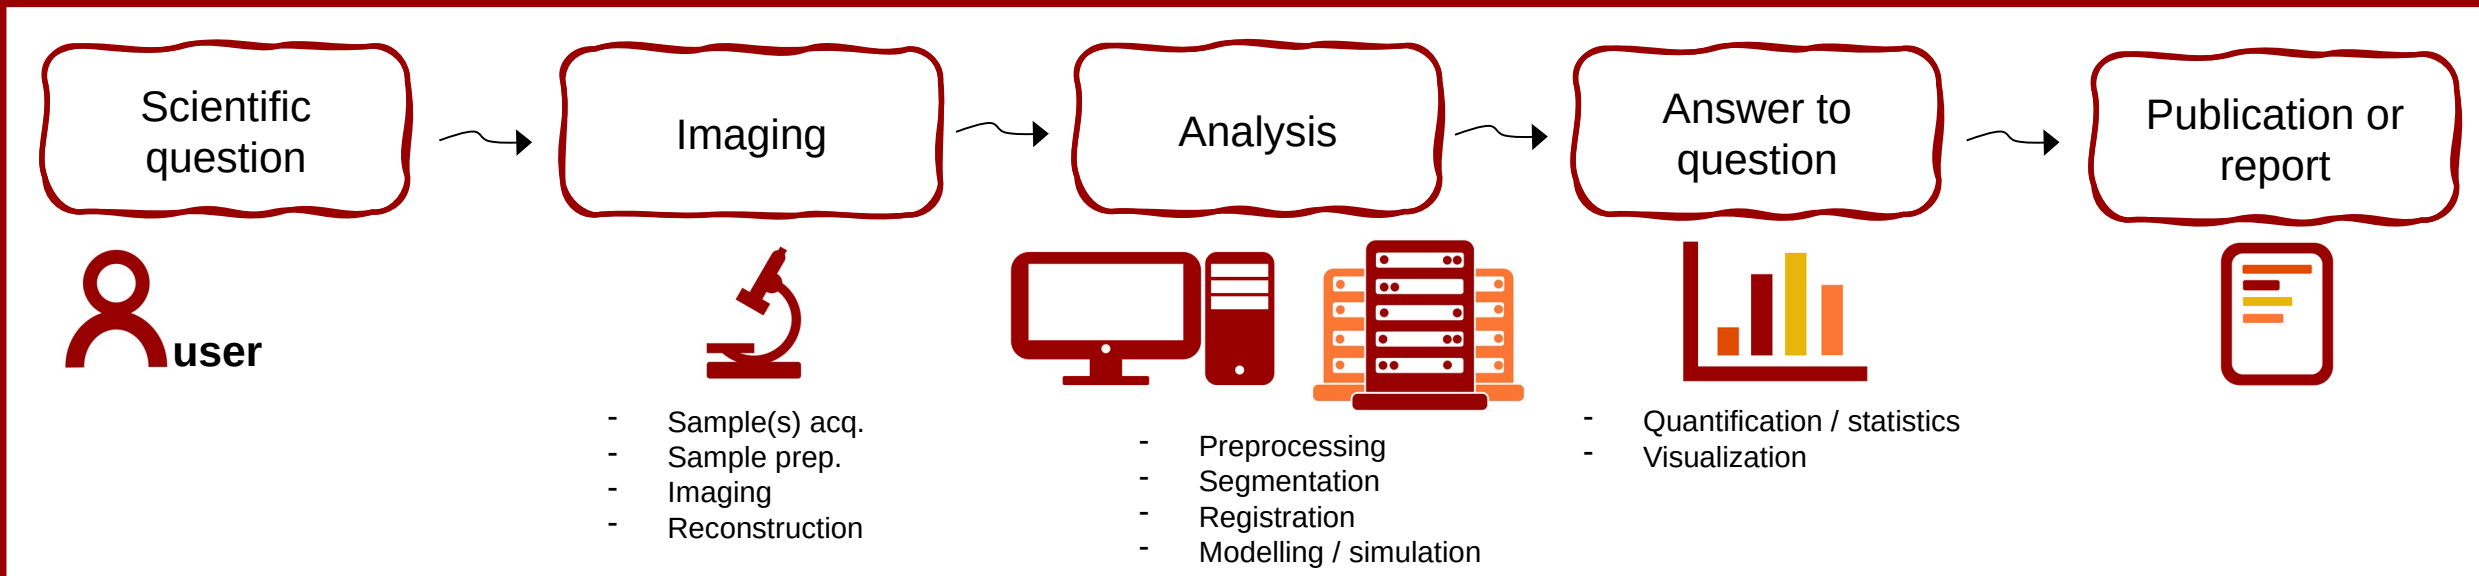



02509: High performance computing for **experimental** 3D image analysis

02510: Deep learning for **experimental** 3D image analysis

Course responsables/contributors:

- Marco Pizzolato (Assistant Professor)
- H. Martin Kjer (Researcher)
- Vedrana Andersen Dahl (Associate Professor)
- Anders Bjorholm Dahl (Professor)
- Patrick Møller Jensen (postdoc)





2 / 5 < >

Section for Visual Computing

The Section for Visual Computing at DTU Compute carries out research in image analysis and computer graphics. The goal of image analysis is to extract information from images whereas the goal of computer graphics is concerned with image synthesis. Despite the fact that the goals are converse, there is a large overlap in methodology between the two fields.

The research is focused on the following five areas:

- > **Medical Image Analysis**
- > **Geometry and Appearance Modelling**
- > **Computer Vision**
- > **Multivariate Statistics for Images**
- > **Analysis of 3D Microstructure**

Within the above areas, our overarching research aim is to make methodological contributions of broad applicability but motivated by applications and often informed by our collaborations with partners in applied domains. Our contributions have often resulted in innovation, and we are proud of the fact that several startups have been founded by former members of the section.

We are always looking for research collaborators and students, and you are welcome to contact us if you have an idea for a research project or need advice within our field of expertise. You can find more information about our section in the description of our research areas.

Analysis of 3D Microstructure

Analysis of 3D microstructure is one of the core research areas within the **Visual Computing section** at DTU Compute. We develop new methods for fast and accurate quantification of 3D microstructure from large 3D volumes. Our research is conducted in close collaboration with scientists developing and using 3D imaging. Together, we promote 3D imaging as a reliable tool for measuring microstructure.

Thanks to the construction of the large-scale facilities **MAX IV** and **ESS** in Lund, Sweden, there is now a fantastic opportunity for 3D imaging at extremely high spatial and temporal resolutions close to DTU. An obstacle in fully utilizing this opportunity is a time-consuming image analysis, which is challenged by the size and complexity of the acquired data sets. There is therefore a need for expertise and new analysis tools for 3D imaging. Developing this expertise and associated tools is our research focus.

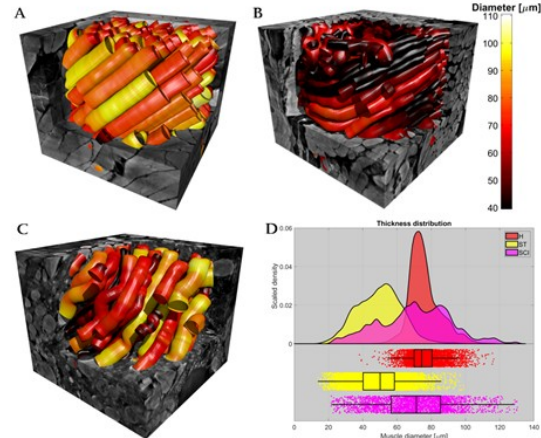
We have established the **QIM: The Center for Quantification of Imaging Data** in a collaboration with University of Copenhagen and Lund University. QIM is also a part of the **DTU 3D Imaging Center - 3DIM**, which is our X-ray μ CT laboratory that is also tied closely to the **DANMAX** beamline at MAX IV.

Below, you find our research highlights:

- Segmentation based on machine learning
- Dictionary-based segmentation
- Graph-based segmentation

Research core

Quantitative characterization of 3D images involves measuring the geometric and intensity characteristics of imaged samples. Often, the geometric features are of interest, because they determine the properties of a sample. This puts focus on methods for efficient and easy segmentation.



Contact



Hans Martin Kjer
Researcher
DTU Compute
+45 45 25 30 35
hmkj@dtu.dk



Vedrana Andersen Dahl
Associate Professor
DTU Compute
+45 27 35 98 81
vand@dtu.dk



Marco Pizzolato
Assistant Professor
DTU Compute
mapiz@dtu.dk



Anders Bjørholm Dahl
Professor, Head of Section
DTU Compute
+45 45 25 39 07
abda@dtu.dk

Medical Image Analysis


Medical image analysis takes place at the intersection of image acquisition, processing, analysis and interpretation. While our main focus is on the modelling and analysis of data, this can only be done through a detailed understanding of the image data and the underlying anatomical context. Our vision is to push the frontiers of integrated image acquisition, modelling and its interpretation. Our research is grounded in industrial, clinical and societal needs.

Below, you find the following highlights:

- **Magnetic resonance imaging (MRI) of the brain**
- **Explainability, fairness and uncertainty quantification**
- **Topology aware learning for medical imaging**
- **Shape and appearance modelling**
- **Geometric machine learning**
- **Cross-disciplinary applications**
- **Synchrotron imaging for validation and prediction models of diffusion MRI**
- **Light-sheet Fluorescence microscopy is crosslinked with MRI**

Magnetic resonance imaging of the brain

Magnetic resonance imaging (MRI) is an extremely versatile modality that allows us to acquire 3D images of the **brain** with many different "contrasts" leading to the collection of 4D or sometimes N-D volumes. **Modeling** these images means applying and developing **signal/image processing, biophysical models, and mathematical procedures** to report summary parameters that can reveal the presence of pathology like stroke, edema, trauma, neurodegeneration, demyelination, etc. There is a variety of MRI techniques like functional and structural MRI (relaxometry, susceptibility, etc.).



Individual muscle fibers segmented from the 3D X-ray synchrotron data. Segmentation allows for quantification of various features - here local thickness. A) A sample from a Healthy Control (H). B) A sample from a stroke patient (ST). C) A sample from a Spinal Cord Injury patient (SCI). D) The statistics of the three samples (Pingel et al. 2022).

Contact



Aasa Feragen
Professor
DTU Compute
+45 26 22 04 98
afar@dtu.dk



Anders Nymark Christensen
Associate Professor
DTU Compute
+45 45 25 52 58
anym@dtu.dk



Rasmus Reinhold Paulsen
Professor
DTU Compute
+45 45 25 34 23
rapa@dtu.dk



Tim Bjørn Dyrby
Professor
DTU Compute
+45 45 25 34 24
tbdy@dtu.dk



Marco Pizzolato
Assistant Professor
DTU Compute
mapiz@dtu.dk

Background of courses 02509 and 02510

3DIM




Rajmund Mokso

Senior Researcher,
Department of Physics

- UltraSound and Biomechanics
- Neutrons and X-rays for Materials Physics
- Center for Fast Ultrasound Imaging

**DTU 3D imaging center
April 11th (facility visit)**



- ▶ Week 1: Introduction
- ▶ Week 2: image processing fundamentals
- ▶ Week 3: x-ray and tomographic reconstruction
- ▶ Week 4: edges, gradients, and smoothing
- ▶ Week 5: 3D scale space and filtering
- ▶ Week 6: Porosity analysis
- ▶ Week 7: Image registration
- ▶ Week 8: MAX IV visit
- ▶ Week 9: Advanced segmentation
- ▶ Week 10: Surfaces and orientation

02509

- ▶ Week 1: Introduction
- ▶ Week 2: Convolutions
- ▶ Week 3: Setting up DL-models
- ▶ Week 4: Entering the 3D dimension
- ▶ Week 5: Fast and efficient training
- ▶ Week 6: Tomography
- ▶ Week 7: Annotations
- ▶ Week 8: Preprocessing and alignment
- ▶ Week 9: Sparse labels
- ▶ Week 10: Quantification
- ▶ Week 11-13: Project work



02510

DL-accelerated segmentation for Histopathology

Comparison between 2D and 3D CNNs for volumetric data semantic segmentation

GOAL

Compare cardiac tissue vessel semantic segmentation based on partially annotated volume using 2D CNN U-Net with stitching into 3D volume and direct 3D volume segmentation using V-Net.




DATASET

CT- scan from Cardiac tissue from patients who succumbed to Covid-19.

Single 3D volume [1024, 1024, 1024], a total of 1.1e09 voxels. A mask with only 0.14% annotated voxels is provided for training the model, with 2 level encoding:

- 1 - Blood Vessel.
- 2- Background tissue.



METHODOLOGY

The project has been carried out following the standard Image Analysis pipeline:

Data Inspection: Understanding the data distribution and features. CT

Pre-Processing: The trainset is composed of sub-volumes of the original image containing annotated voxels in the center of the volume, random rotations along the axis are performed.

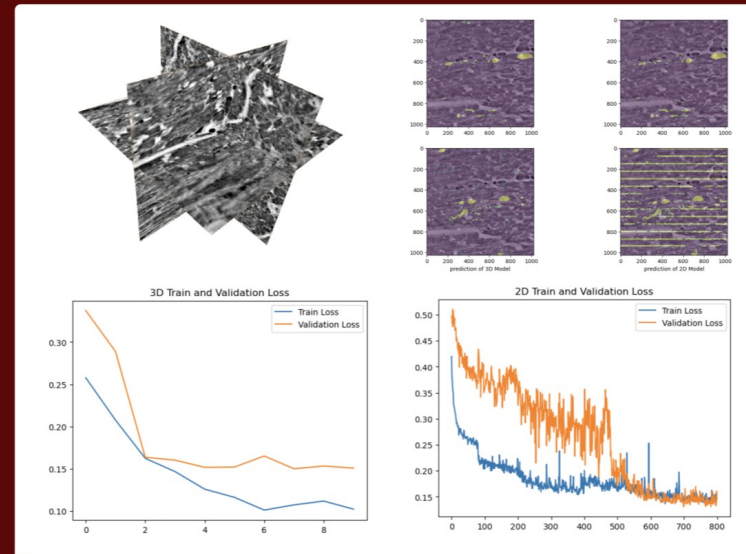
Training: Networks are trained in the standard way, using Dice-Loss and Adaptive Moment Estimation optimizer (ADAM). Encoder-Decoder structure.

- Algorithm for finding orientation labels in the 2D approach: we write a function to find labeled 2D crops. For each 3D sub-volume, only one of sectional planes is saved.

Validation: A mask containing only annotated voxels is used to analyze the performance over the whole annotated set. Dice loss.

Post-Processing: Stitching back together the volume, using Monai's SlidingWindowInference class.

Quantification: Result comparison. Dice Score.



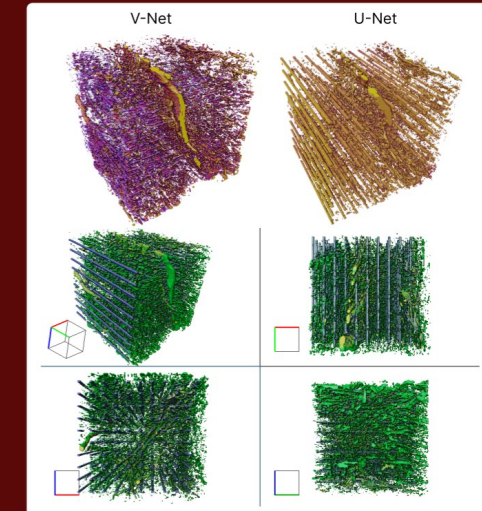
RESULTS

Using V-Net's output mask as ground truth, the inference of the 2D U-Net can be compared.

- There is clear artifacts repeated along the sagittal plane of the volume.

- Performance comparison:
Computing time: Training and result generation
Memory usage

	V-Net	U-Net	Dice Similarity
Dice Score	0.005459	0.007614	35,4%
Num. Params.	76,850,113	25.970.809	



FUTURE WORK

Pre-processing: Data Augmentation to avoid overfitting

Other Architectures: Med-DANet, nnFormer, KIU-Net

Improving ground truth: to improve quantification and vessel thickness analysis.


Get the brain out!

A random forest classifier segmentation study, using a DL-based method as a reference




GOAL

Extract a mask of the brain and quantify volume/performance. Then repeat for cerebellum only.



DATASETS

- 4 datasets, MRI scans in nii.gz file format.
- Data type: float64 bit format having values: 0 to 1400
- Channel Number: 1 channel, grayscale
- Volume size: [224, 216, 288]




Mask of brain 1 obtained through HD-Bet.

REFERENCES

Isensee F, Schell M, Tursunova I, Brugnara G, Bonekamp D, Neuberger U, Wick A, Schlemmer HP, Heiland S, Wick W, Bendszus M, Maier-Hein KH, Kickingereder P. Automated brain extraction of multi-sequence MRI using artificial neural networks. *Hum Brain Mapp.* 2019; 1–13. <https://doi.org/10.1002/hbm.24750>

RESULTS

Feature analysis




Intensity has a huge impact for small scales, that decreases with bigger scales, while edges is generally less important. Texture was inactivated after some trial-and-error tests.

Performance evaluation

We used the Intersection over Union to compare the mask produced by the classifier to HD-BET mask. Maximum overlap achieved of 93%. Brain 1 volume: 1.4 L


$$IoU = \frac{\text{Area of Overlap}}{\text{Area of Union}}$$

Brain 3 (left) and 4 (right) obtained through the random forest classifier



No need for pre-processing, but could use dilation and closing. Regionprops might be useful to validate our results - future work: for cerebellum, use centroid location, eccentricity, bounding box.

Cerebellum segmentation



Further segmentation was attempted to extract the cerebellum from the brain. Using META's Segment Anything (SAM) model to label the area based on bounding box and point object recognition..

By repeating the classification task using the same framework, the results were far from ideal maxing out at a 61% IoU.

The cerebellum from brain 1 had a volum of 148.1 cm³. The average volume for the sets was: 150.36 cm³ indicating a slight overestmition

METHODOLOGY

1. We ran the deep learning tool HD-BET on the 4 datasets and obtained a brain mask for each of them. We visually evaluated its performance, concluding that it is a relevant tool.
2. We used these masks to create labels to be used in a random forest classifier.
3. We quantified the importance of features in the random forest classifier.
4. We identified relevant pre- and post-processing methods.
5. We quantified brain volume of each dataset.
6. We compared the performance with HD-BET.
7. We used SAM to segment the cerebellum of 2 datasets and use trainable multi-scale features to segment cerebellum in the remaining 2, while comparing the result with the "manual" segmentation from SAM.
8. We quantified cerebellar volume across datasets in order to perform optimal parameter selection


FUTURE WORK

- Explore alternative methods to perform semantic segmentation, such as active contours (snakes) and layered surfaces.
- Exploit features based on Blob recognition and region properties, to provide geometric information of the relevant areas.
- Alternative classifiers.
- Segmentation could only happen on bigmem nodes and took some minutes. This time could be reduced with more accurate feature analysis.



Goal for the data

- Visualizations and quantitative descriptions of the vasculature system



Data presentation

The supplied dataset consists of five CT Rat Kidney images of resolution (1000x1024x1014) voxels. The image files are changed to the correct format (.nii). Rat: 30, 32, and 33 are loaded into the ITK-Snap software (visualization tool) and rotated using 'Rat 37' as a reference. This ensures that the images are comparable in terms of their rotation. Rat no. 31 and 37 were discarded, as their data seemed incomplete and did not yield a useful image (reasons unknown).

The images of rats no. 30, 32, and 33 were then resliced for further analysis in python. As it was decided that three different thresholding methods, were to be used, only rat no. 30 was analyzed.

Methods

To analyze the vasculature of the rat kidney, the first step is to find a threshold value, to separate the vasculature from the rest of the structures in the kidney. The threshold value is extracted using three different methods, listed below.

Otsu's method for thresholding

- Built-in function localizing the threshold value to be 49.12
- Creating a binary image where all voxels with intensity above 49.12 = 1 and the rest to be 0

Choosing threshold based on histogram

- Creating a histogram of the intensity levels. Choosing the value of intensity 30
- Creating a binary image where all voxels with intensity above 30 = 1 and the rest to be 0

Frangi filtering

- Specific parameters have been chosen for the frangi filter. By looking at a histogram, an 8G6 frangi-filter output can be thresholded, and then overlaid on rat no. 30, eventually laying grounds for a 3D visualization.

Local Thickness analysis

After creating the binary images, the images were used to visualize and compute the vasculature sizes

- Using built in function 'localthickness'
- Visualize the vessel sizes using histograms
- Adapt parameters of alpha, beta and sigma

Discussion

The methods practiced to analyze the data in this mini-project, has shown promising results in terms of segmenting the blood vessels of the rat kidney (no. 30). The slice using Otsu's threshold has as it seems not captured all the blood vessels, as more of the core of the kidney is visible. The slice using threshold 30 has the most highlighted areas, most likely due to noise and the fact that "too much" tissue was considered/not filtered in the dataset. As expected, the result using Frangi threshold has what it seems to be the best capture of the vascular system. This was as expected, as the Frangi filter highlights tubular structures and is often used in vascular micro-CT. However, due to the computational construction of the Frangi filter, an exceedingly long construction time (8h) can be expected in comparison to the other methods. The Frangi filter also requires more parameter-tuning.


To further substantiate the above mentioned points, regarding the Frangi filter being advantageous to blood-vessel extraction, it can be noticed on the histograms, that the "Self-chosen" and the Otsu's method thresholding methods, has a tendency to favorize smaller blood vessels. The frangi filter, has a broader histogram in comparison to the others which could be interpreted, as it "counting" and differentiating between smaller and larger vessels better than the two other methods.

Results

The results shown are based on rat 30, slice 513 (z-axis).

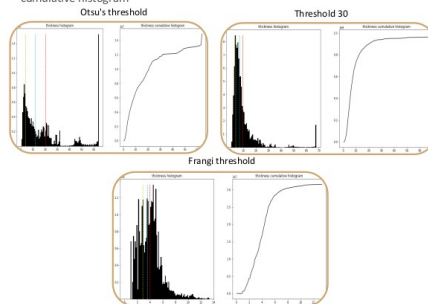
Thresholding

Images showing a 2D slice and 3D slice of the three different methods of localizing the vasculature system of the kidney. The 2D slices are taken from slice [:::513].



Local Thickness Analysis

Three images showing the thickness histogram and thickness cumulative histogram.



Further work


Looking at the results, we could see that the Frangi filter showed the best result. Therefore, for future work, we would focus on optimizing the parameters of the Frangi filter and do further analysis of the size and amount of vessels in the rat kidney.

Wormborer Tunnels

Visualization and quantitative descriptions of wormborer tunnels in wood from μ CT scans

Preprocessing

The image data were normalized, and the bit-depth was reduced from 32 to 8 bits to reduce computational complexity. The images were then thresholded according to pixel intensity to generate binary slices, where only the regions of interest were kept. The background was removed/cropped using contour detection enabling easier segmentation and ensuring only the actual sample was taken into consideration when quantizing. It also reduced the amount of data enabling faster computation.



Segmentation

To segment out the wormholes, different methods were applied. For local thickness analysis, a threshold was found from a voxel intensity histogram, and a voxel intensity-based segmentation was carried out. The local thickness analysis was carried out after smoothening the image using a Gaussian kernel of sigma 0.5. A voxel intensity threshold was then chosen to create a binary volume. Morphological closing of holes was performed using a disk shape of radius 2.


For 3D visualization, contour detection using the algorithm of Suzuki and Abe [1] was carried out on each threshold-binarized slice through the sample. From this, connected component segmentation was carried out. Using region properties of the connected components, regions in a slice with an area less than 270 px² and an eccentricity greater than 0.98 are excluded. This filters out small fragments and cracks, as these tend to be rather linear than circular.



3D- reconstructions of the sample (left) and the segmented tunnels (right).

Local Thickness Analysis

Local thickness analysis was performed on the segmented image. Watershed transform was eventually performed on a modified distance field to label the image features. The local thickness analysis was then carried out and eventually combined with the connected component analysis to relate each label to its corresponding radius.



Result of local thickness analysis. Each color represents a corresponding radius. Purple being small and bright yellow being large.

Tunnel lengths

Using connected component labeling, we identified individual tunnels in the wood, creating a mask that highlighted these tunnels. We estimated the length of each tunnel by counting the number of slices in which it appeared.

Attribute	Measurement	Max value
Length	3.56 mm (mean)	25.57 mm
Radius	1.49 mm (mode)	3.48 mm
Density	88%	

Results

The average tunnel length was found to be 3.56 mm, with the longest tunnel extending to 25.57 mm and the shortest tunnel being just 0.9953 mm in length. Furthermore, the standard deviation of tunnel lengths was calculated to be 4.26 mm. The total volume encompassed by all tunnels amounted to 370.83 cubic millimeters. From the local thickness analysis, a mode radius was determined to be 1.49 mm, but as the tunnels vary in size, tunnel radii up to 3.48 mm were measured. When including the natural cracks in the wood, the material density was approximately 88%.

Discussion

Isolating and segmenting wormholes in wood, and conducting local thickness analysis, enabled quantitative measurements. Worm holes and natural cracks are then easily distinguishable. The quantitative analysis of tunnel lengths and volumes, along with the material density calculations, provides valuable information that could help improve wood treatment techniques, pest control strategies, and woodwork processes.

The 3D visualization could be cleaned up by removing features at the volume border, or by cropping the volume additionally. Also features inside the tunnels were found, and it is debatable whether these should be removed just as with the natural cracks.

Sources

- https://github.com/vedranaa/teaching-notebooks/blob/bb08da792108764211788907c2590cd521d9/Cement_porosity_analysis.ipynb
- [1] Satoshi Suzuki and others. Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing*, 30(1):32–46, 1985.

Contact us




Segmentation of stonewool

Problem

Segmenting fibers in two different types of stonewool

New scan: 2 μm Old scan: ? μm




Methods

Binarization is done by following image analysis methods: Global thresholding, Otsu, Frangi-filtering, local thickness


Best fit binarization is used in training

Model



Channels: 64, 128, 512, (1024)
Input channels: 1
Output channels: 3

Results



Computation



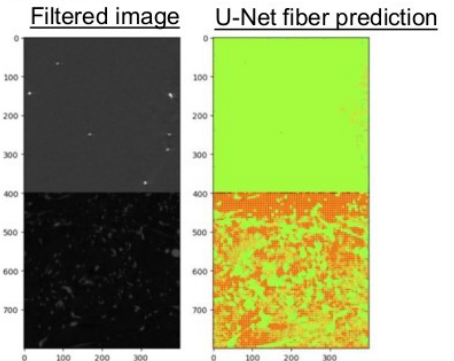


Image size: 400x400x400 voxels
Random Patch size: 96x96x96 voxels
64 patches in training set
16 patches in validation set
16 patches in test set

Label 0: Background
Label 1: New scan
Label 2: Old scan

Raw image
Validation loss: 0.63

With frangi filter
Validation loss: 0.08



Computed with connected components between slices

Augmentation:
Random crops, affine,
flip, rotate

Loss function:
MaskedDICE

Epochs:
130

Optimizer:
Adam

Batch size:
16


Scaler:
Cuda amp GradScaler

Learning rate:
0.0003

References

- [1] Yang, Chen-Hong et al. Deep Hybrid Convolutional Neural Network for Segmentation of Melanoma Skin Lesion. Computational Intelligence and Neuroscience, vol. 2021.
- [2] Ian Goodfellow and Yoshua Bengio and Aaron Courville. Deep Learning. MIT Press. 2016.
- [3] Frangi, Ro F. et al. Multiscale Vessel Enhancement Filtering. Lecture Notes in Computer Science. 2000.

DTU 3D imaging center → April 11th (facility visit)



3D Imaging Center aids in the search for perfect mozzarella

torsdag 22 jun 23 Researchers at 3DIM agree that nothing tastes better than a good pizza after a long day of work, which is one of the many reasons why the center's facilities are used to investigate the microstructure of mozzarella cheese.




3D scans will uncover the secrets of gold treasure

onsdag 08 nov 23 One of the biggest collections of gold treasure in Danish history has been 3D scanned at DTU. Researchers are now trying to unfold the gold pieces digitally to uncover new knowledge about the power dynasties of the Iron Age.

The hunt for the origin of dogs

torsdag 02 nov 23 Population genomics will be combined with the generated 3D skull models of ancient and modern canids to test the identity of controversial Paleolithic specimens



See you on April 11th at DTU