Application Note

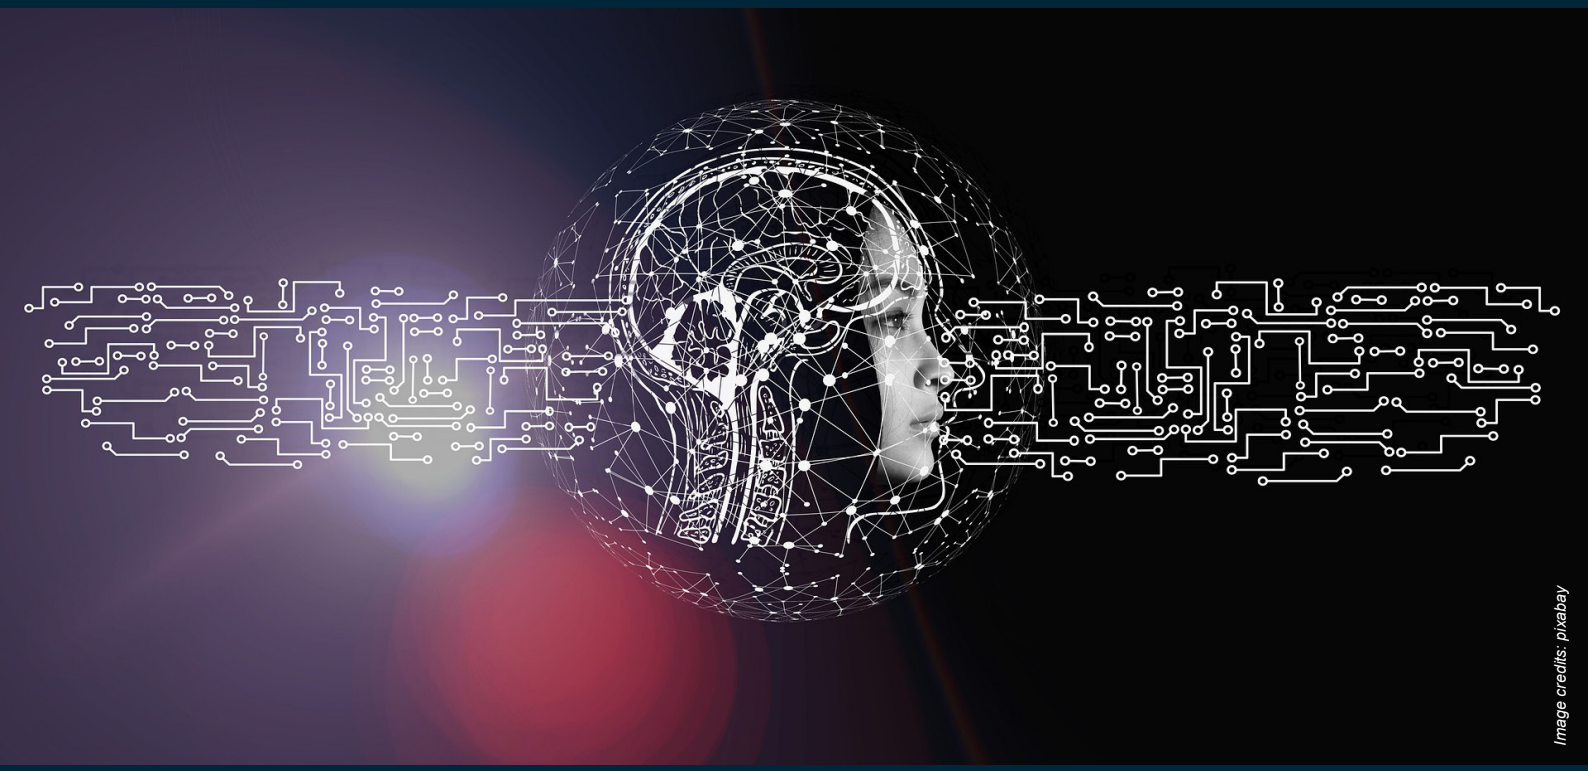

*Image credits: pixabay*

# Everything you always wanted to know about **Dragonfly Deep Learning** but were too afraid to ask

Anton du Plessis

dragonfly

# Dragonfly Deep Learning

**Why use deep learning for Computed Tomography (CT) image data?**

- Deep learning can **make image segmentation possible where it was previously impossible** or near impossible (due to noise, artifacts, lack of traditional segmentation tool capabilities, etc.)
- Deep learning **removes human bias from segmentation tasks** (improving repeatability & reproducibility)
- Deep learning **can facilitate automation** of image processing and analysis tasks due to the above reasons, especially in combination with macro's
- Deep learning methods can be used for **image enhancement**, for subsequent better or easier segmentation and analysis

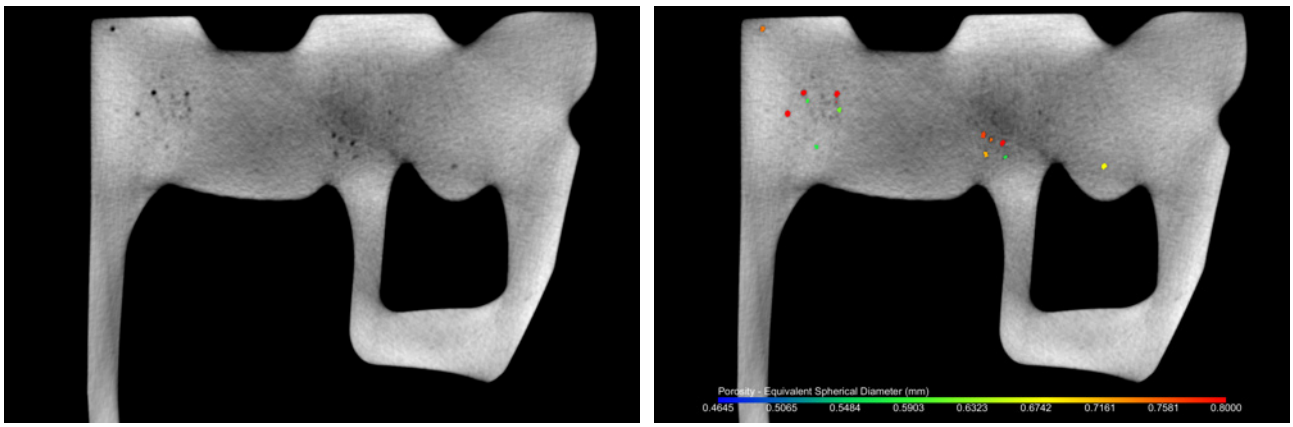


*Image credits: Comet Yxlon Applications Lab, Hamburg*

Image segmentation: In this example a deep learning model was used to segment the casting porosity despite underlying grey value variations (bright and dark regions) and small feature sizes. Regular thresholding methods do not work in this case due to artifacts and partial volume effect making small pores less dark than larger ones.

# Misconceptions

Before we discuss the benefits, it is important to highlight some frequent misconceptions:

- **Misconception 1: “Deep learning is fast”**

Actually this is far from the truth, deep learning is almost never something that is fast and never should be attempted as a “quick test to see if it works”. Deep learning uses lots of computer resources. It takes time to train a model, to check it, and even to apply it to a full dataset, even with today’s biggest graphics cards and best software for the job (Dragonfly). A typical training and inference process on industrial CT data would take about 1-4 hrs, much longer than typically spent on traditional segmentation tasks. As an example, the inference process alone (application of the trained deep learning model) to a full CT dataset takes 23 minutes, using a laptop with an Nvidia 3050 GPU. This is for a dataset with 750 x 560 x 1350 voxels. Clearly this is not fast, but the value addition comes with clever ways of using a trained model as part of an automated workflow or applying the same model on multiple datasets, or using image enhancement or segmentation methods that save time in the long run. If you can solve your image problem with traditional image tools, then this is almost always the better option, especially when those can be automated (e.g. standard image filter de-noising, automatic Otsu threshold). Deep learning is the long-distance racer, not the fastest car available.


*Image credits: pixabay*

- **Misconception 2: “Deep learning can create false or unreliable information”**

This misconception comes from the era of chatbots that are trained to make conversations and is not a concern in image analysis software. Most important to understand is that in deep learning in the context of image segmentation and image enhancement, all the underlying information is clear and unchanged – there is no image manipulation and all results can be checked. Deep learning image segmentation is based on ground truth segmentations provided by the user and follow well defined paths to create and train the models. In Dragonfly, all the deep learning architecture and parameter information is saved with the model (nothing is hidden). All the trained models are also saved together with their training history, providing information on how well the model converged and many other metrics that are open to the user to analyze (and even compare different models to see which is actually better). As with all tools, there are good ways to use them and bad ways to use them, here the software you use is key to ensure you use it optimally and your verification of the results are also important, especially for some test cases. Dragonfly makes this particularly easy for the user.

- **Misconception 3: “You need a PhD to do deep learning”**

Actually today it is possible to use easy to follow software tools to train and apply models, without having any knowledge or understanding of the underlying algorithms (and without needing any advanced degrees!). All you need is some logic and to follow some general “good practice” rules, often these are included into software tools to make it easy for you to use default models and reach good solutions. In Dragonfly, there are two main ways to train deep learning models. The easiest is the “Segmentation Wizard” that guides you through the process of defining ground truth, choosing a model and training a model. In this process, you can visually see if the results look reasonable and, if not, you can stop and adjust the ground truth and continue the training again. This interactive process is intuitive and easy to implement, without needing any knowledge of model architectures, depths, patch sizes or other technical aspects. When you have a need to access some of these more advanced tools, you can “open the hood” and modify all these parameters to values different from default. The second method in Dragonfly to train deep learning models is to use the “Deep Learning Tool” that allows you the freedom to modify all parameters and even adjust the architecture of the model manually, meant for expert users. Dragonfly is the only CT software that makes deep learning easy for the user, while also giving full control.

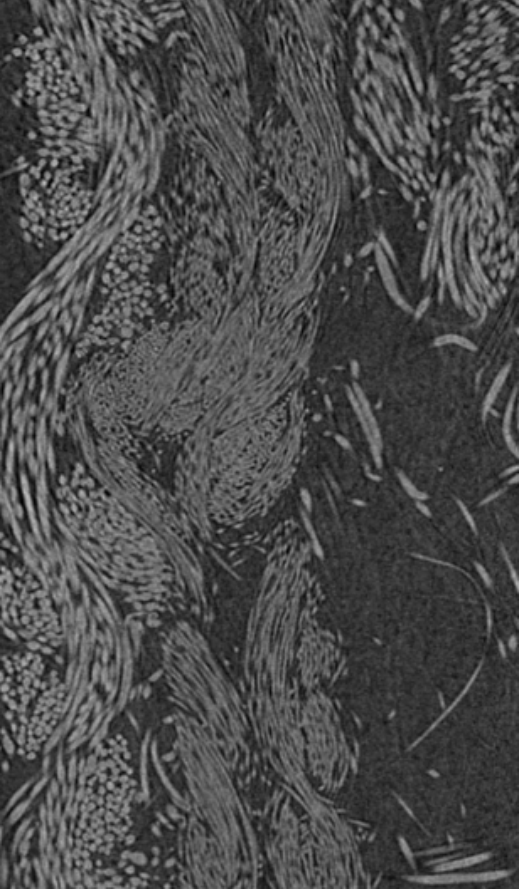
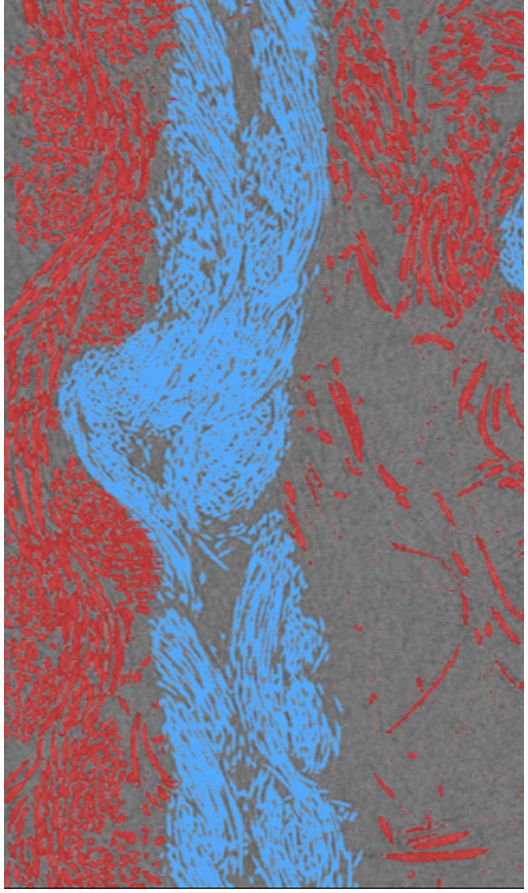
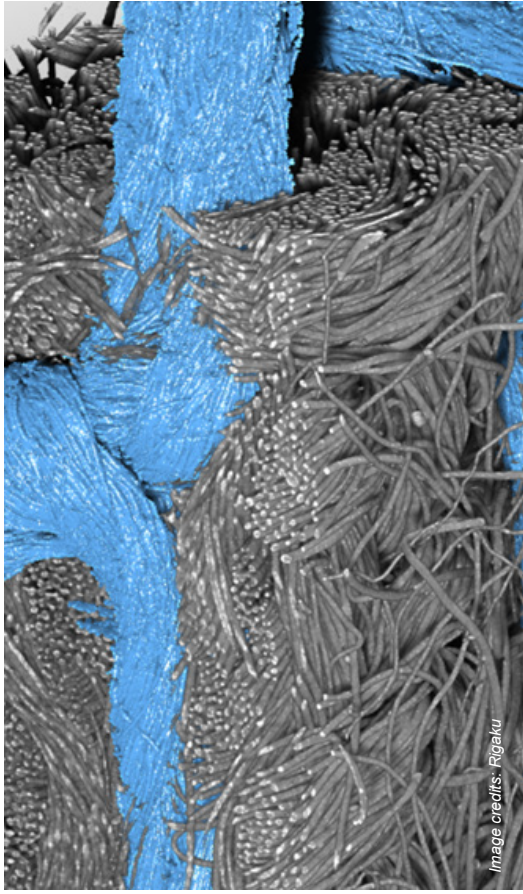


# Image segmentation

Any image segmentation task can be achieved by deep learning, by understanding these basic concepts:

- You have to provide ground truth data, typically by manually segmenting one or more slices in the image.
- The model is only as good as the ground truth data – that means you need to give it examples of everything that you expect in the images to be segmented (if inclusions are present in the new data but was never in the training data, the model will not know what to do, for example)
- You can teach a model to segment in any way you like, for example you can specify fibers to be larger than they actually are, by manually painting the ground truth like that. This can be useful in some special cases, e.g. when you want an easier visualization

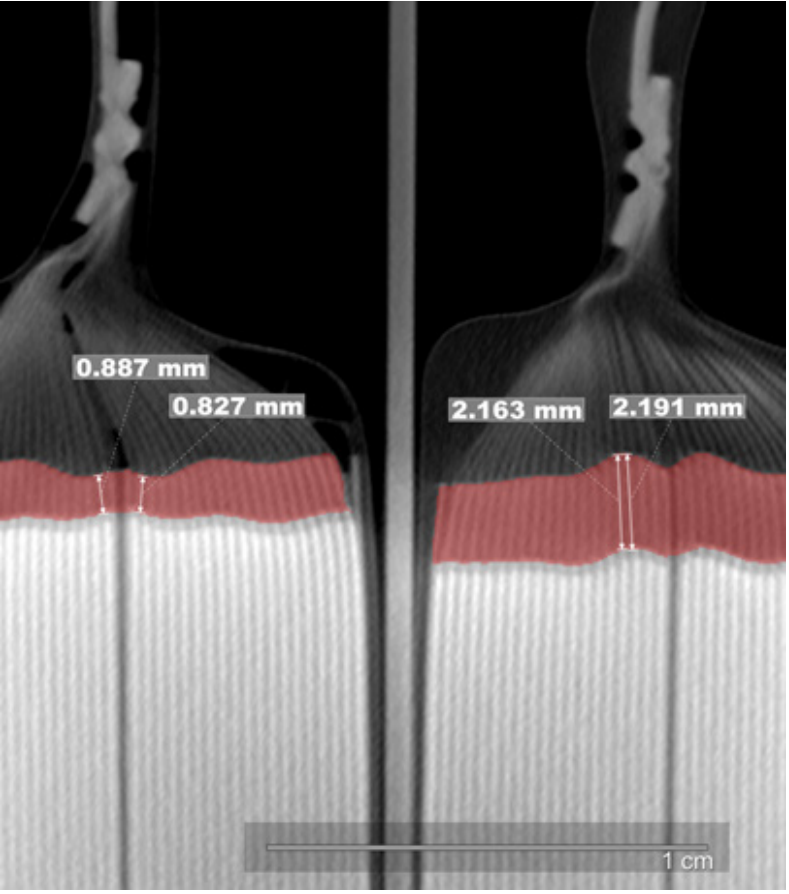


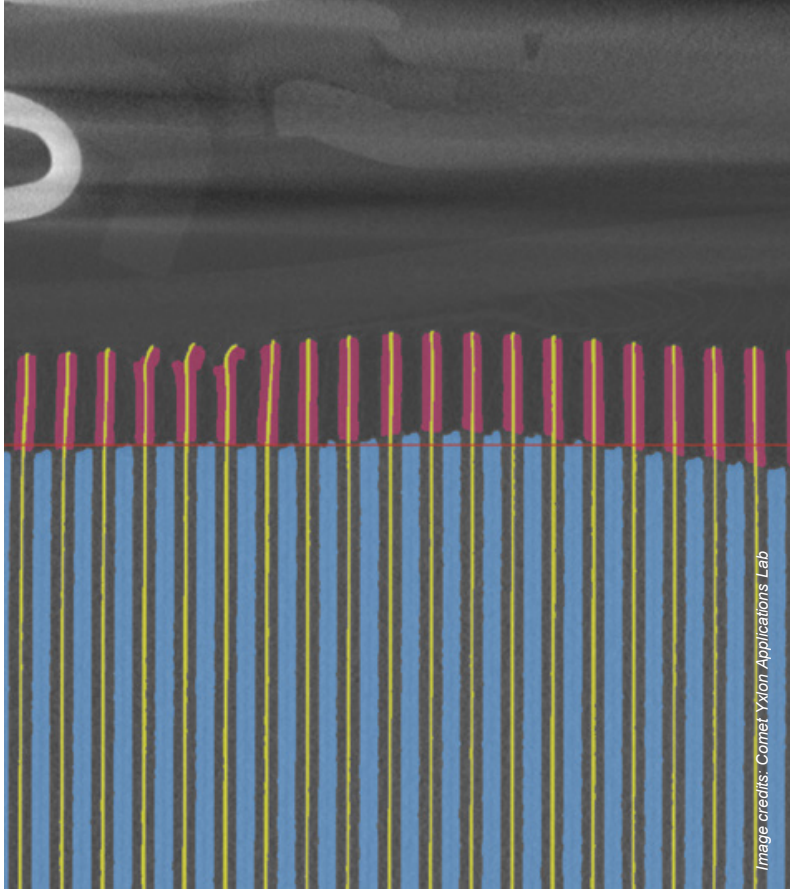


# Image segmentation

The examples shown above are two cases of battery overhang segmentation, in both cases performed by deep learning. In the case on the left, the entire overhang region is segmented. This is done intentionally due to lack of image resolution and the benefit of obtaining 3D local thickness mapping information easily with such a segmentation. On the image on the right, three different materials are individually segmented – in this case for detailed individual overhang measurements. These examples illustrate the flexibility in using the deep learning tools to segment in different ways, depending on the need.

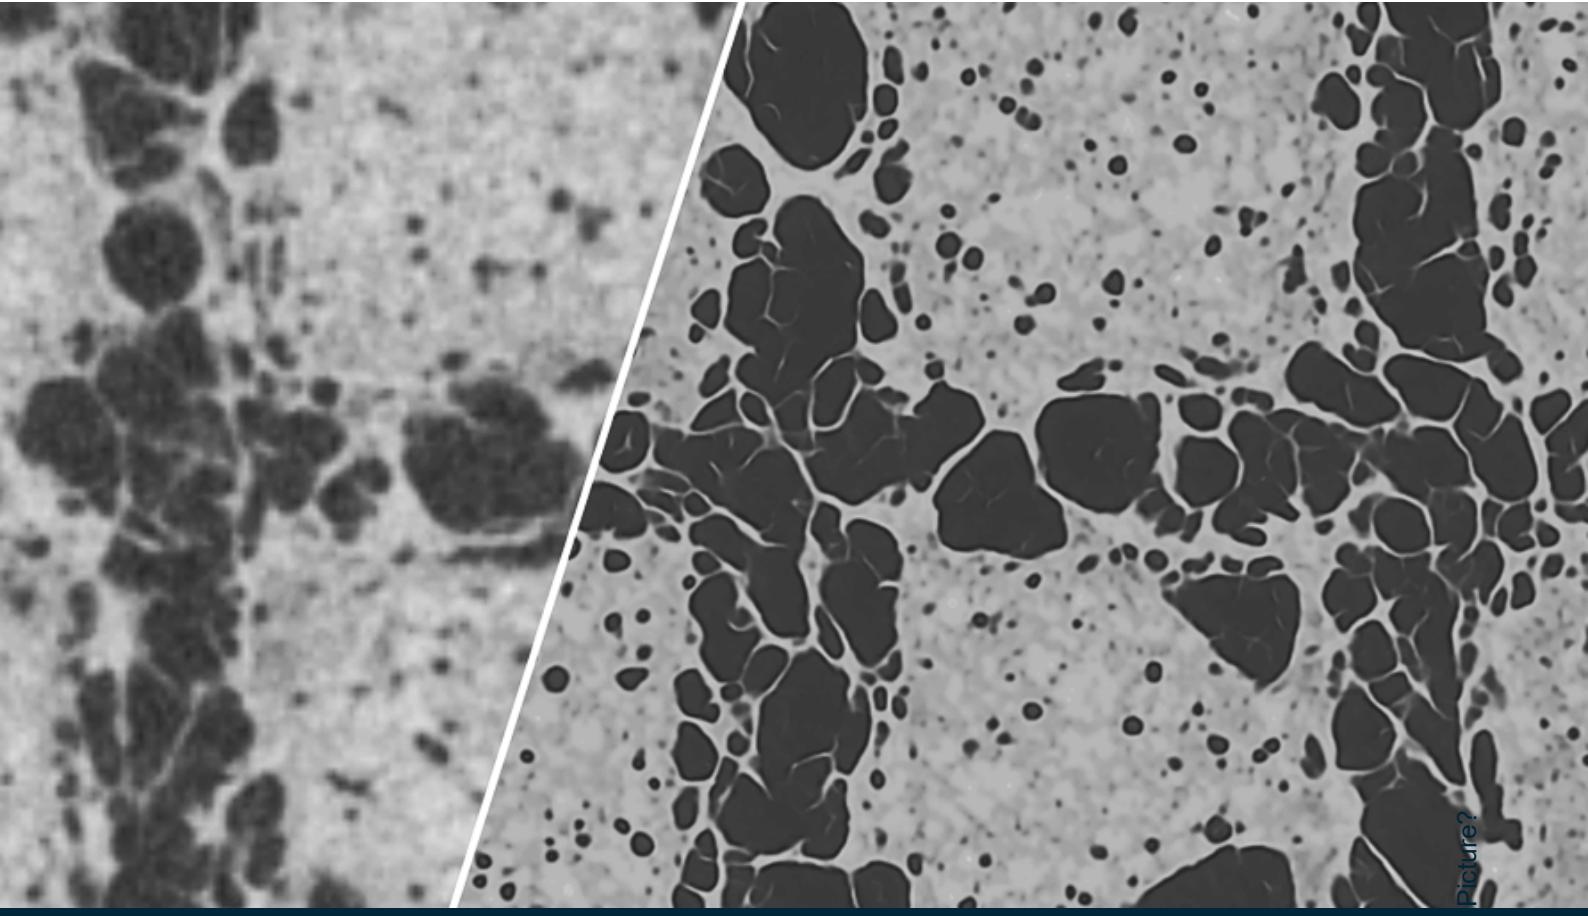


# Image enhancement

Deep learning image enhancement methods allow to improve the quality of images, by taking noisy or poor data and using deep learning methods to improve the quality of these images. Even if it looks and sounds like magic, it is based on solid physics. No new image features are created, only existing visible features are enhanced in contrast, noise removed and overall the image quality improved. Image enhancement methods using deep learning include variations of “super resolution” and Noise2Void, amongst others.

# Super resolution

In the case of super-resolution (one of the methods for deep learning image enhancement), you use two scans to train the model. One is a poor resolution scan (e.g. 16 µm scan on left of image), and one is a high resolution scan (4 µm scan on right of image). These two scans must cover the same field of view which is used to teach a deep learning model how to move from the poor data to the good data. Basically, a deep learning model is trained to reach something “that looks like” the high resolution scan. The effect is shown in the middle in the example below. A new image with less noise and better contrast is the result.

Other variations of super resolution involve long high quality scans and fast noisy scans – this allows to create a model that converts noisy, fast scans to better quality, saving lots of scan time when applied to many similar samples. This is sometimes referred to as super-contrast

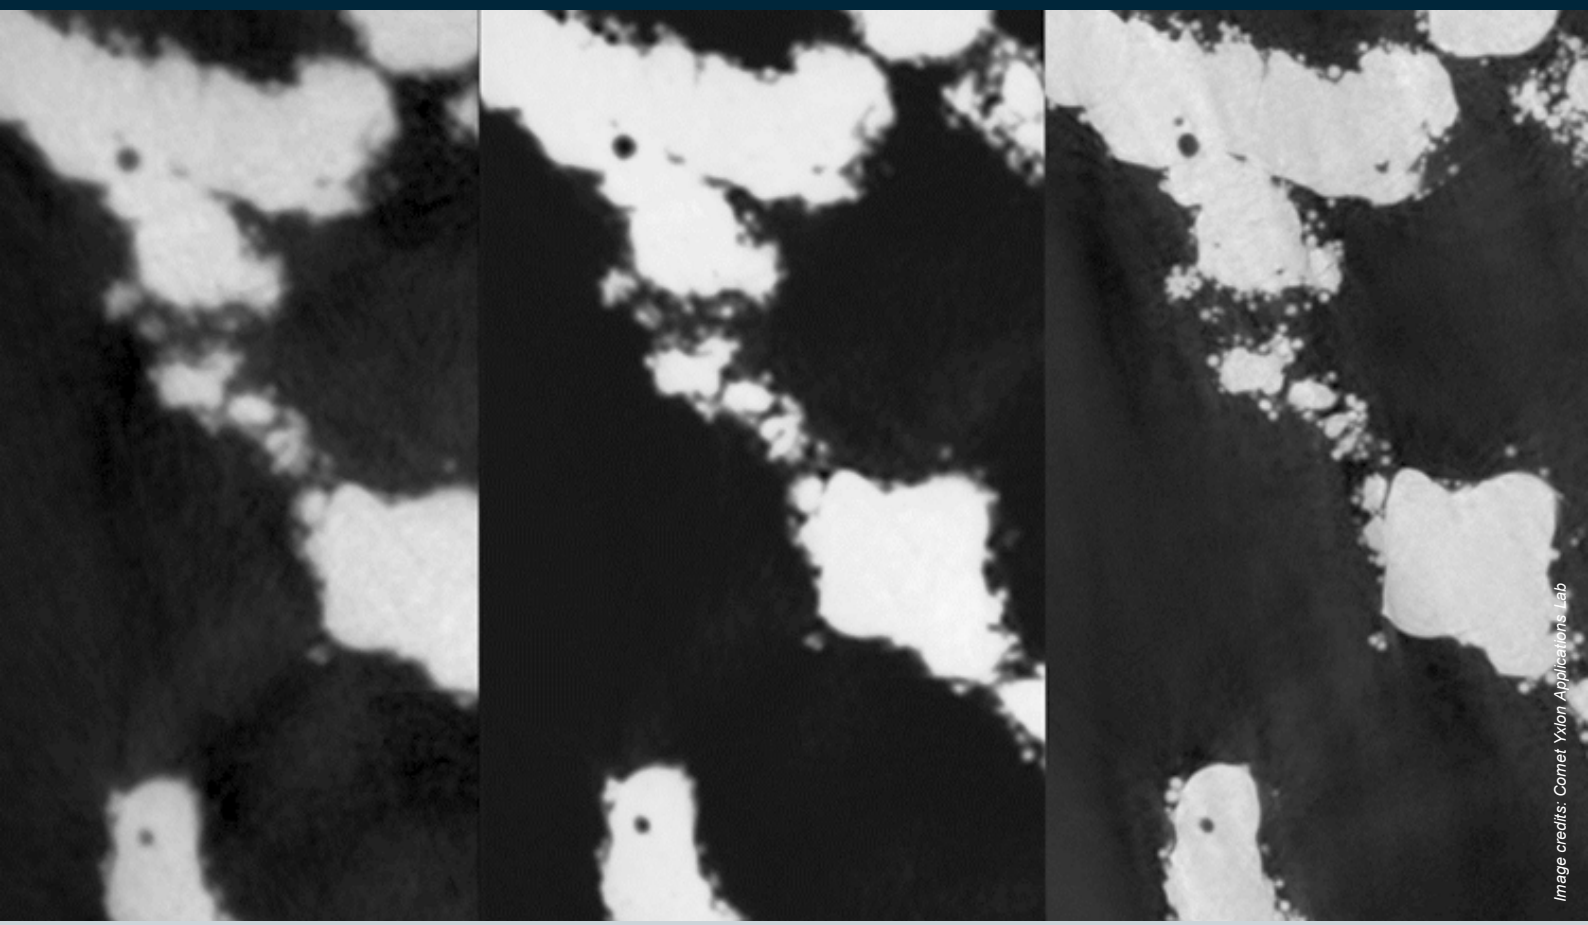

*Image credits: Comet Yxlon Applications Lab*

16 µm | Super resolution | 4 µm

# Some ways to use deep learning

**Image segmentation of challenging images**

In today's industrial CT applications, the features of interest are often very small or have poor contrast. This is an ideal case for deep learning. Until today, if the segmentation task failed on such data, the operator would go back to the acquisition and attempt higher quality scans, taking more time and possibly never reaching good enough image quality. Even in the face of noise or poor contrast, when the feature of interest can be identified by eye (i.e. when you can see what you need to segment), a model can be trained to segment it. This opens a world of possibilities to the user: examples include segmenting datasets with artifacts (streaks, beam hardening, rings, etc.), excessive noise or poor contrast, variations across the image, etc., can be performed by deep learning.

**Applying segmentation models to multiple datasets**

Since deep learning is based on greyscale images, variations between scans may cause some challenges. Dragonfly includes the possibility to calibrate the intensity scale prior to training a model, which makes the model more robust to such changes between scans. This means that you can use your specially developed model on all similar scans of the same object type.

**Adding data to existing models to refine them**

Transfer learning is the approach of taking an existing model and adding more training data to it, in order to make it more robust and better suited to a different dataset type. An example is a generic model included in Dragonfly for automated porosity segmentation in additively manufactured parts. Since the model was trained on 50 datasets of specific size and percentage porosity, with specific dataset types, it is not expected to work perfectly for all scans of additively manufactured parts. The idea is to use this model as a start point, and add data of your type to it, to fit it to the task at hand.

# Terminology

**Artificial intelligence**

All forms of problem solving using machines, using some form of human-like intelligence. This is the original concept developed from which machine learning and deep learning methods were developed.

**Machine Learning**

Machine learning involves all systems that learn and adapt from experience, without being explicitly programmed. This means these ML models get better with more training. They are a subset of AI, based on correlations and (often linear) regressions.

**Deep learning**

Deep learning is a subset of AI and machine learning that uses artificial neural networks to mimic the learning process of the human brain. This is the main driver of the explosion in AI in all fields at the moment. Deep learning makes use of multiple layers to learn hierarchical models, which makes more complex, non-linear correlations possible, and are more robust than traditional ML methods.

**Neural network**

Deep learning uses neural networks to encode and decode information and make correlations that are hard to find. Different network types exist such as the Unet, Sensor 3D, Attention Unet, and many more. These are all optional architectures in Dragonfly, with parameters that can be varied. The default is Unet, as this is most widely used and works very well for image segmentation tasks.

**Segmentation Wizard**

This is the name of Dragonfly's wizard-based method of training AI models, including machine learning and deep learning models. It is a user-friendly interface that allows you to define ground truth segmentations in one or more selected regions or frames, selecting a frame for visual feedback and selecting a model or models to train. The training process gives visual information on the success of the segmentation, allowing to stop the training and add data and continue training.

# Terminology

**Inference**

This refers to the application of a model to selected data to make the segmentation or “infer” the segmentation based on the model.

**Classes**

Each material type to be segmented is referred to as a class in a multiROI. The output segmentation is also a multiROI with classes for each material type or phase.

**Ground truth**

This refers to the input of information to teach the model. The user provides segmentations in one or more frames or slices, to provide the input information for training the model. This is done using traditional 2D segmentation tools and should be done carefully.

**Data augmentation**

This refers to Dragonfly’s internal method of taking ground truth segmentation and creating additional versions of this ground truth, to increase the training data. Rotated, stretched and other forms of augmentation are used in this process and can be increased by the user.

**Monitoring frame**

This refers to the visual feedback – the live segmentation of a specific region of the image by the current model being trained. This helps the user to understand if the training is working or improving or not, and when a good time is to stop the training process.

# Some frequently asked questions

**What is the difference between deep learning and machine learning?**

Many software tools today offer machine learning methods to reach reasonable results. In some cases this is sufficient, but often images are more complex or variations between images of different samples require more robust solutions. Deep learning is a specialized form of machine learning that uses neural networks to encode and decode image information, reaching much more robust solutions. This means deep learning is much better suited to the varying conditions expected for CT images. Dragonfly is by far the leader in providing deep learning tools that can be trained by the user and is user friendly.

**Is deep learning a black box?**

Many software tools do make their AI tools into black boxes, but in Dragonfly the training history is recorded, making it possible to evaluate the convergence reached, and many other metrics of the model are recorded with the model, for comparison to other models directly. All parameters and details of the model architecture are also recorded. In Dragonfly, the entire process is open to the user with all information saved.

**It sounds great, what are the down-sides?**

Since deep learning is based on image grey values, sometimes a model seems perfect in one slice but does not make the needed segmentation in another region of the image. In this case it might be necessary to go back and add more training data and continue training, making the whole process time consuming. This can also sometimes mean the segmentation needs manual cleanup, in some cases. All segmentations need to be checked and the main problems occur when the training was not sufficient. It is an art form and the more you do it the better your models become.

# More questions

**What makes Dragonfly deep learning the best?**

It is the only commercial CT software that allows the user to train real deep learning models for image segmentation and image enhancement in an easy to use wizard-based process, with all parameters open to the user to fine tune the model as needed. It is the only commercial tool with such a wide variety of model architectures included, allows any number of classes or phases to be included into the model, allows very deep models, allows variations of 2D, 2.5D and fully 3D models, and performs variable image augmentation methods to increase the ground truth data size in clever ways. It allows also multi-modality inputs, that means you can provide different input datasets (e.g. different scans of the same object) to better teach the model how to recognize specific features.

What all of that means is that Dragonfly is the most versatile, open to the user, user friendly and has the highest level of complexity that is on the market today for CT data.

**What are the computer requirements for deep learning?**

You need a good Nvidia graphics card in your machine, something with at least 8 Gb of onboard memory (and more is better), detailed specifications can be found here: https://dragonfly.comet.tech/gettingstarted

**Is it possible for Dragonfly team to train a model for me?**

Yes this is possible and even included in the specialized Battery Insights product. Typically at least 3 datasets are needed and the process is explained clearly with some training included.

**How do I get started?**

It is easy to get started in Dragonfly. Either get in touch with your local sales contact, or go to https://dragonfly.comet.tech/getfreetrial to get your 30 day trial of the software. Self-learning materials are available on our website (https://dragonfly.comet.tech/gettingstarted) and youtube channel (https://www.youtube.com/channel/UCuFI2zHcyStR2RJpMXbi6ow). It is also possible to purchase online training sessions to help you to hit the ground running.

dragonfly

Dragonfly - a brand of Comet
Comet Technologies Canada Inc.
460, rue Ste-Catherine Ouest
Suite 600
Montréal (Québec)
Canada  H3B 1A7
dragonfly.comet.tech