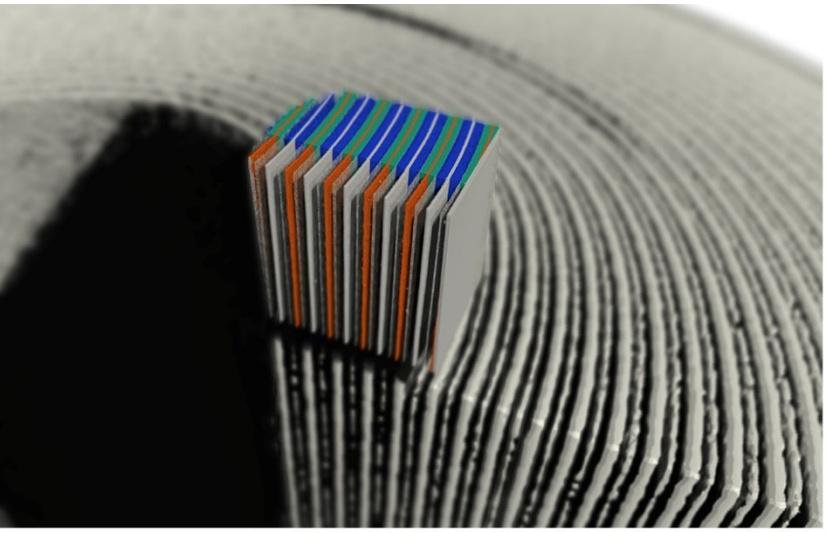

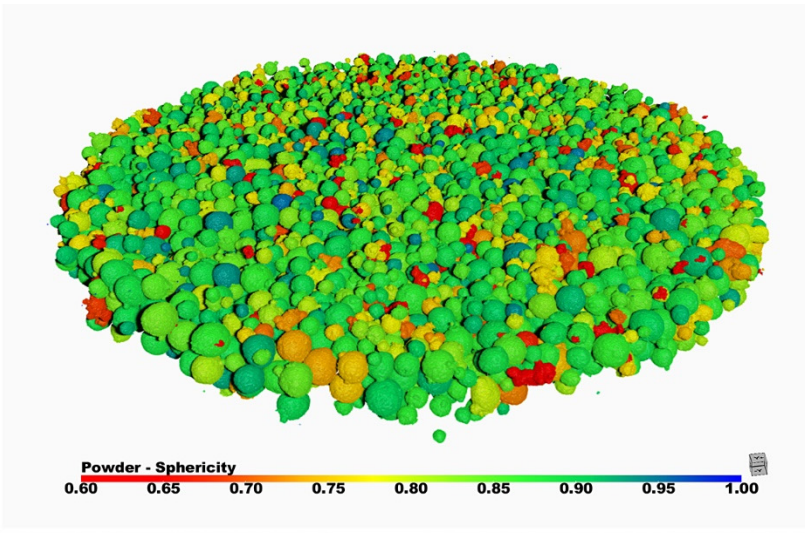

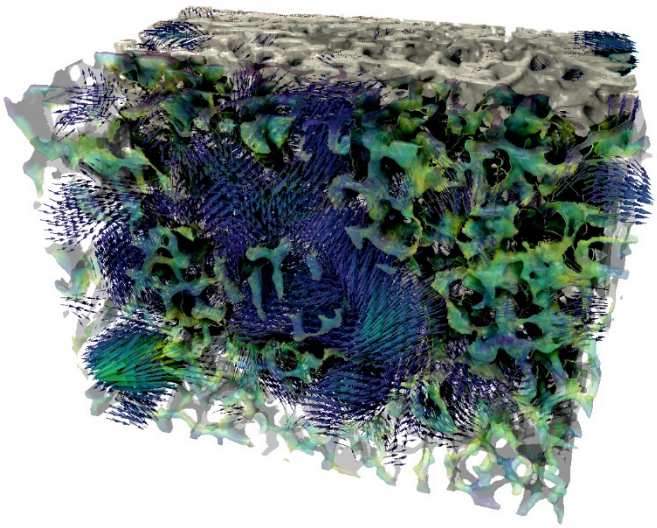

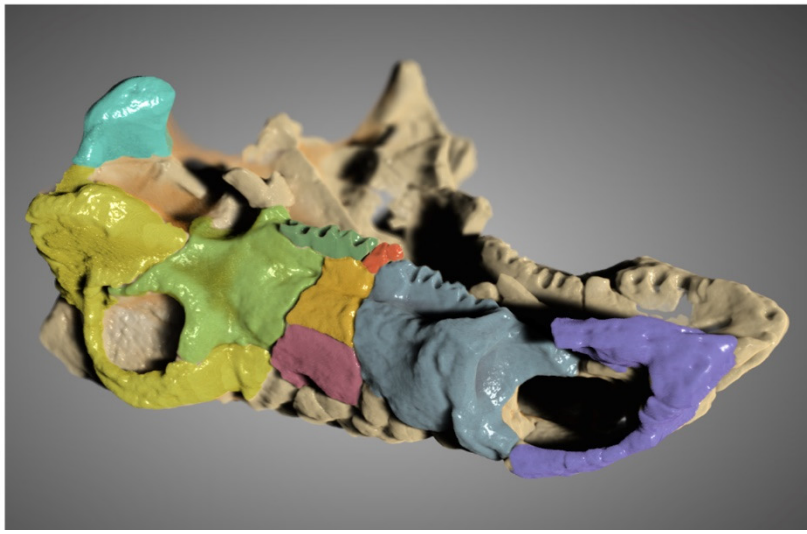

VERSION 2022.2

Learn all about the new features, product enhancements, and other improvements implemented in the Dragonfly and Dragonfly Pro 2022.2 software release. Features of note include **elastic registration** to better understand dataset deformation, **pre-trained deep learning models** to accelerate model training, **normalization of data ranges** prior to model training and inference, and **new measurements** for characterizing material properties and measuring bone mineral density.

Cover Images

Clockwise from the top left: 3D rendering from a correlative workflow (courtesy of ZEISS), additive manufacturing powder characterization, vector field mapping of a 4D deformation study looking at the interface between an additive manufactured spacer and a human vertebrae while undergoing continuous compression (CT data courtesy TESCAN/sample courtesy Rush University), segmented Ankylosaurus jawbone (courtesy of DigiMorph).

Object Research Systems (ORS) Inc.
460, rue Ste-Catherine Ouest, Suite 600
Montréal (Québec), Canada H3B 1A7

www.theobjects.com

Contents

New Features and Product Enhancements

The software release for Dragonfly 2022.2 and Dragonfly Pro 2022.2 marks the first integration of **Feature Analysis** for performing cross-table analysis of feature vectors and other measurements. Other new features of note include **elastic registration** to better understand how datasets were deformed, **pre-trained deep learning models** to accelerate model training, **normalization of data ranges** prior to deep model training or inference, and **new measurements** for characterizing material properties and measuring bone mineral density. Significant enhancements for the Segmentation Wizard, CT Reconstruction, and the Movie Maker, as well as new 3D rendering features and optimizations, are also included in Dragonfly 2022.2.

Feature Analysis

This software release includes an alpha version of Dragonfly's **Feature Analysis** module, which provides data integration for cross-table analysis of feature vectors and other measurements to help classify objects that share specific characteristics.

Choose **Utilities** > **Feature Analysis (alpha)** on the menu bar to open the Feature Analysis dialog, shown below.

Feature Analysis dialog

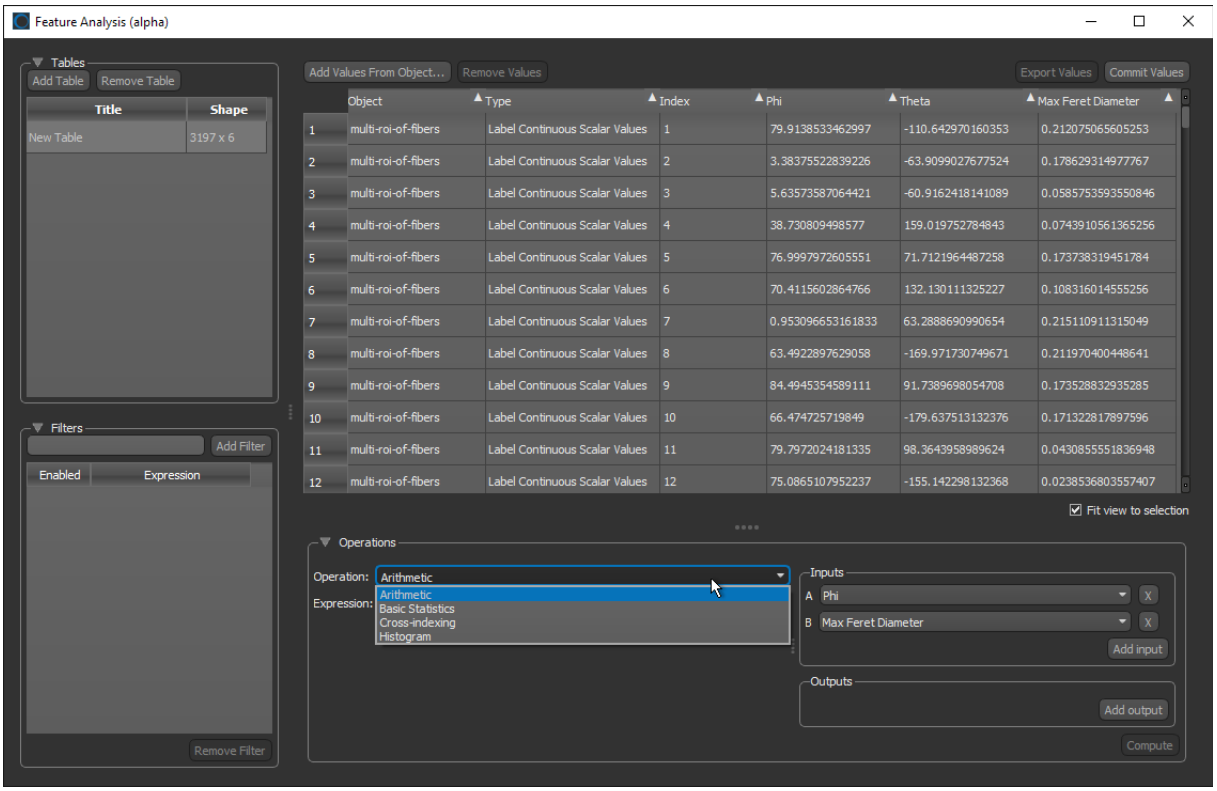

In the Feature Analysis dialog, you can:

- Add scalar values generated from multi-ROIs, meshes, and graphs to multiple tables.
- Compute basic statistics, such as the mean, median, minimum, maximum, standard deviation, and variance of imported scalar values, as well as modify selected values manually or by applying arithmetic expressions.
- Plot selected inputs on 1D histograms or on 2D histograms, where the X-axis represents the values of the first input and the Y-axis the values of the second input.
- Create polar charts that use values and angles to show information as polar coordinates. Rather than using the standard X and Y coordinates, each point on a polar plane is expressed using these two values:
    - Radius (r), which is the distance from the center of the plot.
    - Theta (θ), which is the angle from a reference angle.

  The plane itself is made up of concentric circles expanding outward from the origin, or the pole. Polar plots are often used when the analyzing data that has a cyclical nature.

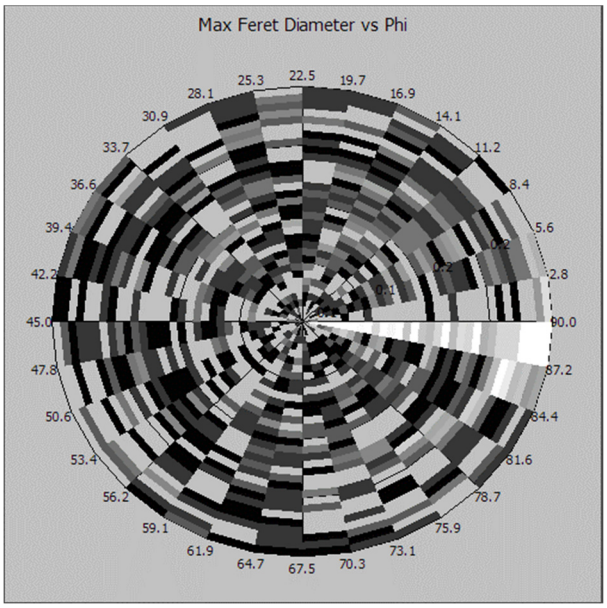

- Create classes by selecting instances in the data that match some criteria. Instances can be added to classes on 1D histograms with the Range Selector tool and they can be selected by painting on 2D histograms and polar plots. You should note that in the case of polar plots, painting is done on annular sectors.
- Select an item in the table to automatically zoom to the selected item in the current 2D/3D views of any data type, including multi-ROIs, graphs, and meshes.

Deep Learning

Updates for Dragonfly's **Deep Learning Tool** include pre-trained models, the option to normalize data ranges prior to training or inference, training on multiple GPUs, and deep model inference on multiple axes. A new deep learning model architecture, custom loss functions and layers, and additional macro blocks for automating deep model training are also included in this release.

Pre-Trained Models

To help you get started quickly with implementing Dragonfly's deep learning capabilities, you can now download a selection of deep models pre-trained by the Dragonfly Team. You can then use *transfer learning* to apply the learning to a specific regression or semantic segmentation problem. Starting with a pre-trained model often provides better and faster results with smaller training sets than using an untrained model. In the examples below, the progress of training with a pre-trained and an untrained deep model are shown as a function of increasingly larger training sets — from 2, 4, 8, 16, to 32 labeled tiles.

Validation Dice coefficients for pre-trained and untrained U-Net dl-5 ifc-64 models using increasingly larger training sets

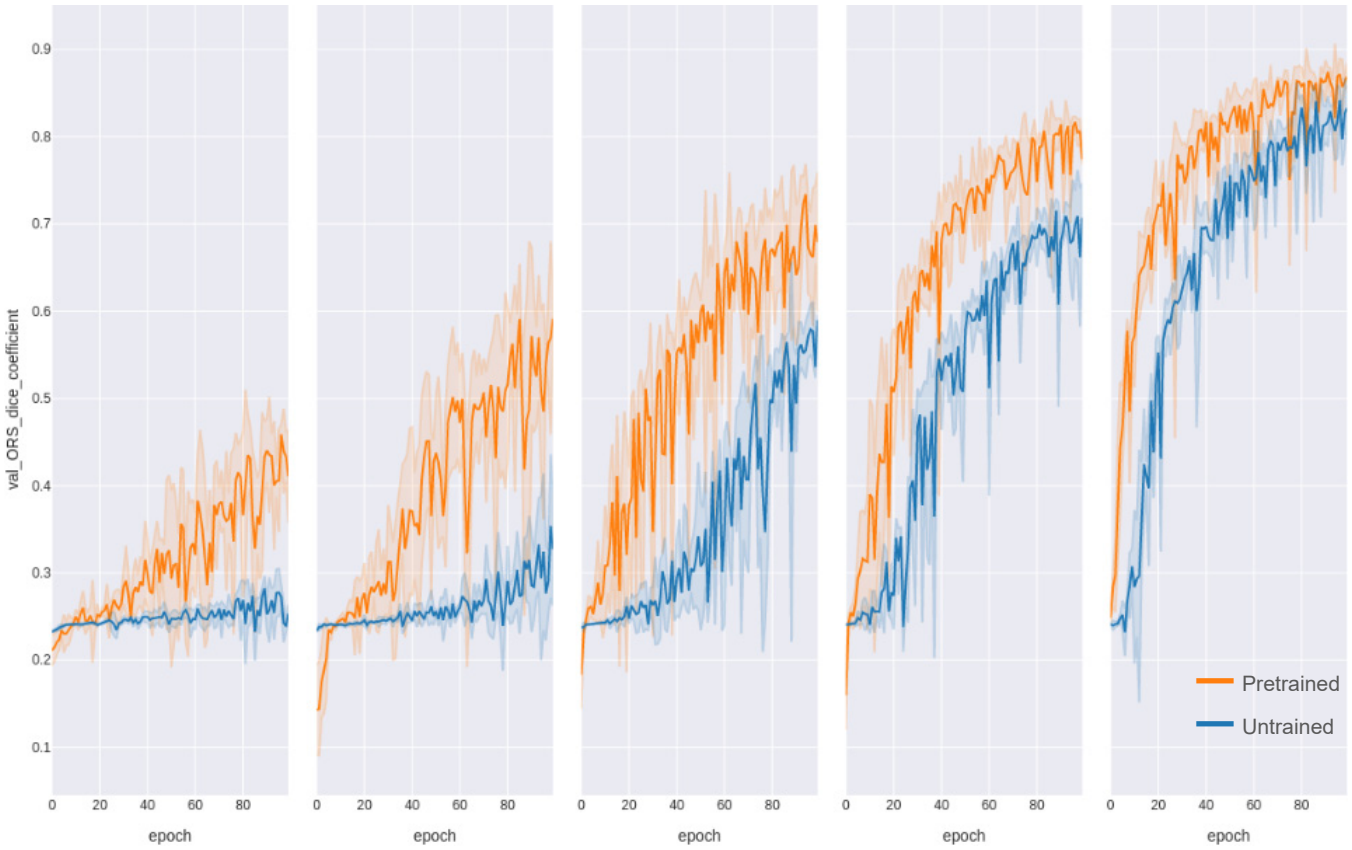

You should note that the objective when training a neural network is to identify the correct weights for the network by multiple forward and backward iterations. In the case using pre-trained models for semantic segmentation tasks, you take advantage of learned feature maps to reduce laborious and time-consuming labeling of training sets. Pre-trained models are generally more label efficient than untrained models as they also use the unlabeled examples for learning.

The following pre-trained models are available for download:

Pre-trained models

Model	Description
Generic U-Net dl-5 ifc-64 v-1.0	**Architecture:** U-Net **Depth level**: 5 **Initial filter count:** 64 **Input dimension:** 2D
Generic U-Net dl-6 ifc-64 v-1.0	**Architecture:** U-Net **Depth level**: 6 **Initial filter count:** 64 **Input dimension:** 2D
Generic U-Net dl-7 ifc-32 v-1.0	**Architecture:** U-Net **Depth level**: 7 **Initial filter count:** 32 **Input dimension:** 2D
Generic U-Net dl-6 ifc-64 slice-3 v-1.0	**Architecture:** U-Net **Depth level**: 6 **Initial filter count:** 64 **Input dimension:** 2.5D, input slices count: 3
Generic U-Net dl-7 ifc-32 slice-3 v-1.0	**Architecture:** U-Net **Depth level**: 7 **Initial filter count:** 32 **Input dimension:** 2.5D, input slices count: 3

NOTE The above models were pre-trained by the Dragonfly team with subsets extracted from a large repository of diverse greyscale images, with uninformative and duplicate data removed.

NOTE Refer to the topic and tutorial *Transfer learning and fine-tuning*, which is available at https://www.tensorflow.org/tutorials/images/transfer_learning for additional information about working with pre-trained models.

Do the following to download and generate a pre-trained model:

1. Open the Deep Learning Tool and then click the **New** button.

   The Model Generator dialog appears.

2. Choose **Pre-trained (by the Dragonfly Team)** as the Architecture.

3. Enter a **Model type** and **Class count** (for segmentation models only), as well as a **Name** and **Description**, optional.

4. Double-click inside the Value cell of the pre-trained model and then choose the required model in the drop-down menu, as shown below.

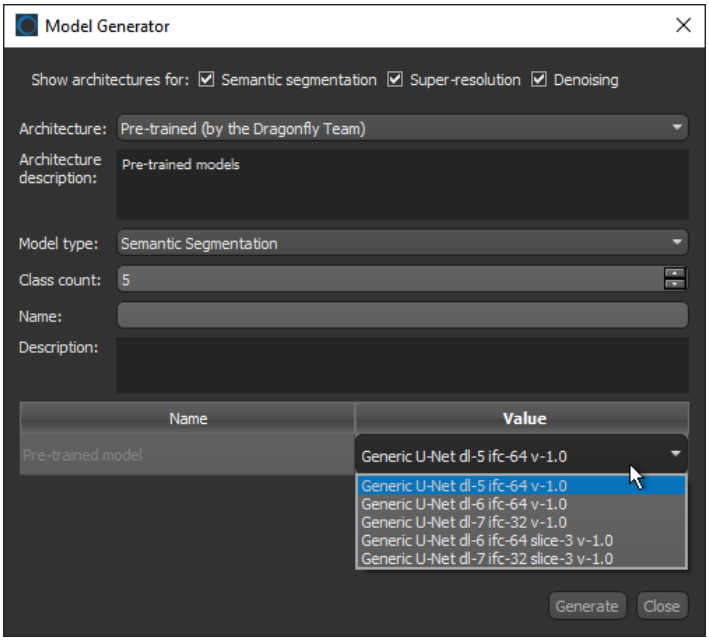

5. Click the **Generate** button.

   The model will be downloaded to a local cache the first time that it is generated and a copy will then be made for the current session. Subsequent models that are generated will be copied from the local cache. You should note that Dragonfly will check for new or updated pre-trained models each time that the application is started, provided that your system is connected to the Internet. You should also note that cached models will be available even when Dragonfly runs offline.

   The local cache is located at:

   `…\AppData\Local\ORS\Dragonfly2022.2\pythonUserExtensions\PythonPluginExtensions\DeepTrainer\PretrainedWeights.`

Normalize Data Range Prior to Training or Inference

Compromised (poor quality) deep model training and inference can occur in cases in which datasets have a limited dynamic range. In these cases, you can now normalize calibrated and uncalibrated datasets prior to training in the Deep Learning Tool dialog, as well as before inference in the Segment with AI dialog and in the Filter with AI dialog.

Click the **Edit Calibration** button, as shown below, to open the Normalization dialog.

Deep Learning Tool dialog

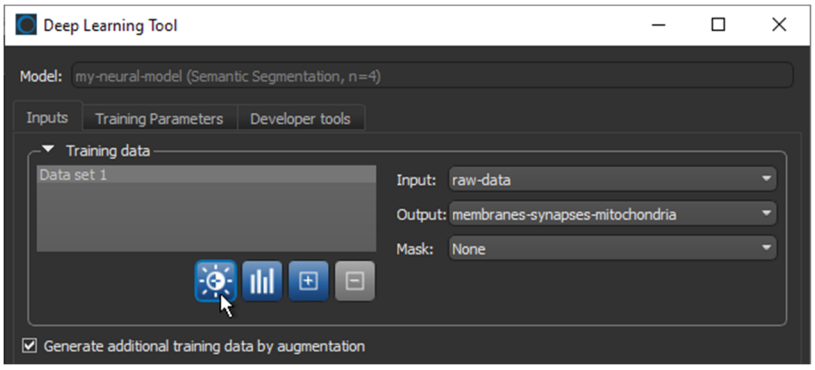

Adjusting the normalization boundary can be useful in cases in which you only want to apply the model on the current input and the input histogram is not spread over the data range. You should also note that the normalization range selected during training is saved with the model and that the saved range is loaded and set as the default at inference. You can edit the normalization range at inference before applying the model on an input, but this adjustment is not saved. The next time that the model is loaded, the normalization range is reset, as shown below.

Normalization dialog

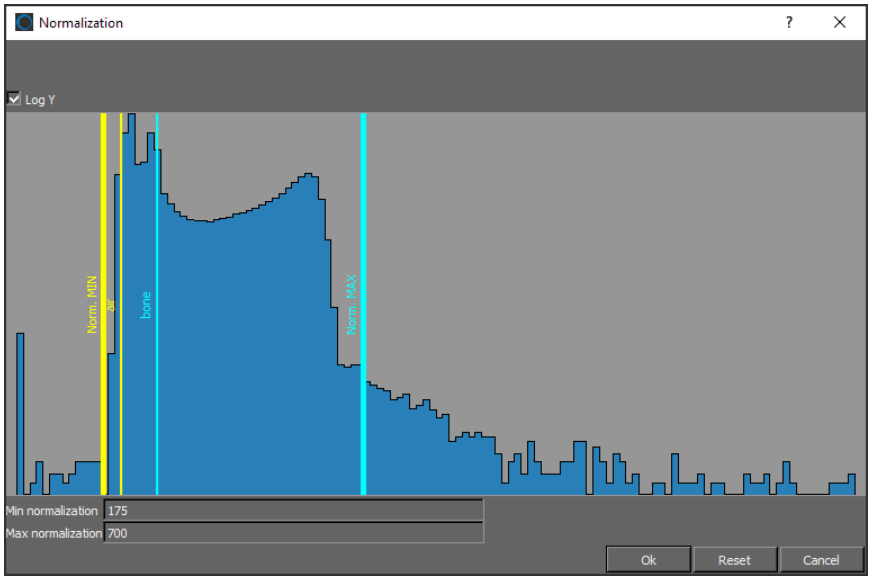

Deep Model Inference on Multiple Axes

As an additional option for deep model interference, you can choose to segment your datasets on multiple axes in the Segment with AI panel. In this case, the deep model will be applied on multiple axes and then the class label will be predicted based on the selected type of vote — **Majority** or **Soft**.

The Segment with AI panel, shown below, is available by default on the left sidebar's Segment tab. You can also right-click the dataset you need to segment and the choose **Segment with AI** in the pop-up menu to open the panel. The options for segmenting on multiple axes are circled on the screen capture of the Segment with AI panel.

Segment with AI panel

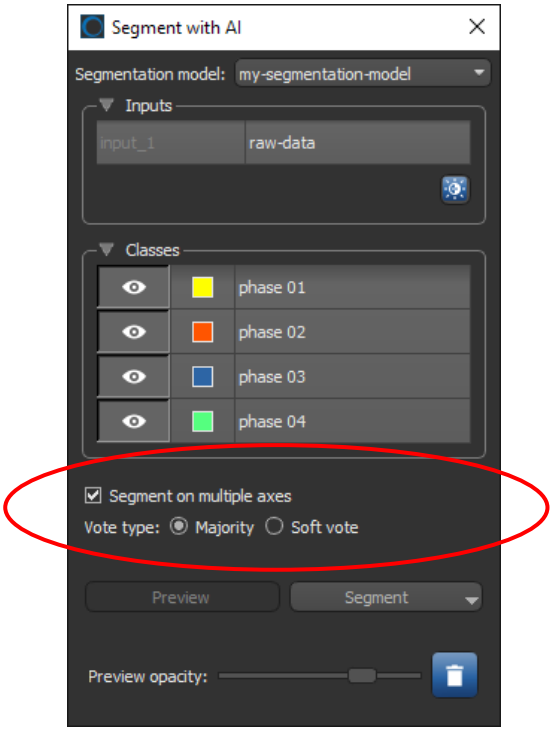

You can select 'Majority vote' or 'Soft vote' for the multiple axes inferences as follows.

Majority vote… Sums the predictions for each class label and predicts the class label with the most votes.

Soft vote… Sums the predicted probabilities for each class label and predicts the class label with the largest probability.

Deep Model Training on Multiple GPUs

In some cases, you may find it advantageous to train deep models on multiple GPUs instead of a single GPU. If you have multiple GPUs available, then you can set how many GPUs to use for deep model training at the beginning of training. It can be any number from 1 to the total number of available GPUs. You should note that you can select your preferred GPUs in the Deep Learning preference 'GPU card for Deep Learning' (see Deep Learning Preferences on page 81).

The option for selecting the number of GPU cards to use for model training is available on the **Training Parameters** tab in the Deep Learning Tool dialog, as shown below.

Training Parameters tab

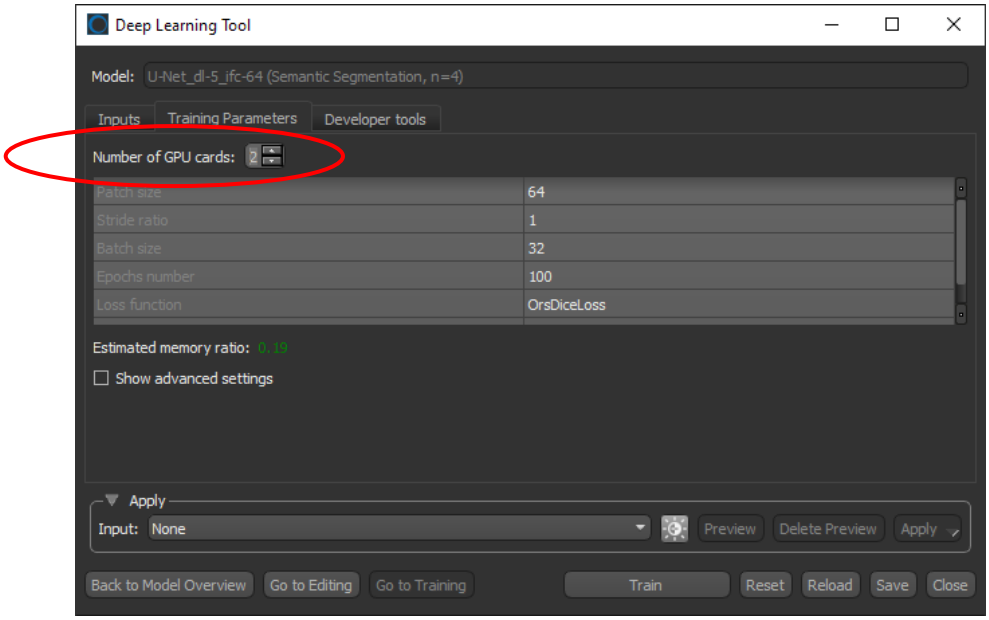

NOTE Selecting this option will add some overhead to model training as models will be replicated for each GPU. In addition, if not all GPUs are the same, then the slowest card may create a bottleneck as patches are split evenly, not proportionally to each GPU.

New Deep Learning Models

A new deep learning architecture — **Multi-Level Wavelet U-Net** — is available in this software release for training regression models. This model is a U-Net style encoder-decoder network with discrete wavelet transform (DWT) encoding and inverse wavelet transform (IWT) decoding. In this modified U-Net architecture, DWT reduces the size of feature maps in the contracting subnetwork, while an additional convolutional layer is used to decrease the channels of feature maps. In the expanding subnetwork, IWT is then deployed to reconstruct high-resolution feature maps.

NOTE Refer to the publication Pengju Liu, Hongzhi Zhang, Wei Lian, Wangmeng Zuo, *Multi-Level Wavelet Convolutional Neural Networks*, IEEE Access, Volume 7, 2019 for more information about this U-Net style encoder. Available at: (https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8732332).

NOTE Information about the implementation of the Multi-Level Wavelet U-Net model, as well as experimental test results, are available at: https://github.com/lpj0/MWCNN.

Default Loss Function

You should note that **Categorical Crossentropy** is now the default loss function for semantic segmentation, instead of **OrsDiceLoss**.

Custom Loss Functions

For training regression tasks, this software release includes a number of new custom loss functions. Double-click the editable **Loss function** field on the Training Parameters tab of the Deep Learning Tool to view the available loss functions, as shown below.

Custom loss functions

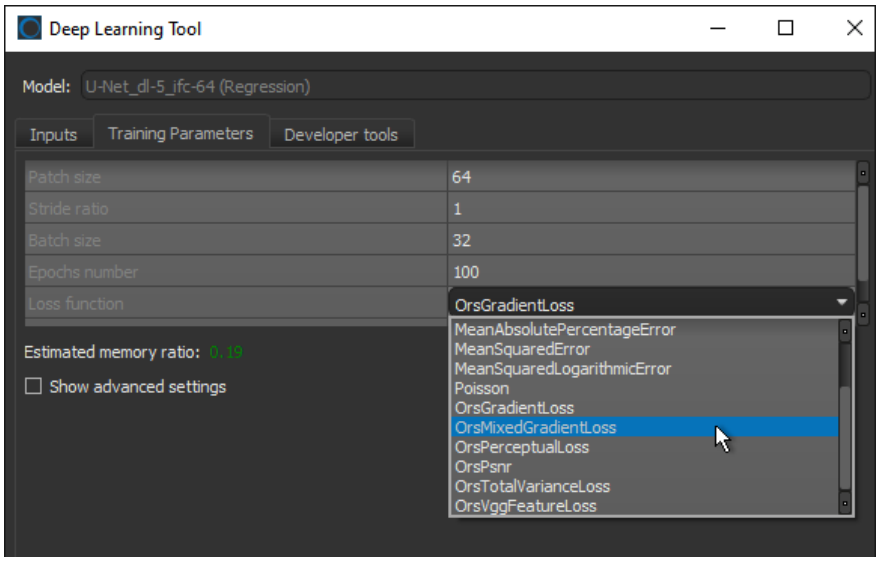

The following custom loss functions are available in Dragonfly 2022.2.

Custom loss functions

Loss Function	Description
OrsGradientLoss	A simple loss function that helps preserve image structures for super-resolution and is the squared error loss between the gradients in X and Y of true and predicted batch. Loss is normalized on twice the batch size to consider the two directions of gradient. **NOTE** The input and output must be well registered or the translation may be learned. In addition, noise has a high gradient and blotchy artifacts can be created in empty areas. **REFERENCE** Abrahamyan et al. *Gradient Variance Loss for Structure-Enhanced Image Super-Resolution*, arXiv, February 2, 2022 (https://arxiv.org/pdf/2202.00997.pdf).
OrsMixGradientLoss	This loss function is similar to OrsGradientLoss but is more forgiving than gradient loss for dataset registration. Includes Mean Square Error in the loss calculation and Sobel edges are used instead of gradient. **NOTE** Uses fixed weight for the gradient and mean square portion of total loss. Noise has a high gradient. **REFERENCE** Zhengyang Lu and Ying Chen, *Single Image Super Resolution Based on a Modified U-Net with Mixed Gradient Loss*, arXiv, November 21, 2019 (https://arxiv.org/pdf/1911.09428.pdf).
OrsTotalVarianceLoss	Total variance defined by sum of integral of absolute value. This loss function tries to minimize the variance in an image while maintaining a fit to the original. Essentially tries to reproduce the image in piecewise constant functions. **NOTE** Weight between variance and image fit automatically adapted based on scale of image features. May be less effective for complicated structures. **REFERENCE** David Strong and Tony Chan, *Edge-Preserving and Scale-Dependent Properties of Total Variation Regularization*, Inverse Problems, Volume 19, Number 6 (https://iopscience.iop.org/article/10.1088/0266-5611/19/6/059).
OrsVGGFeatureLoss	This loss function compares the activation map at a layer of a pretrained VGG16 model between true and predicted batches. VGG style networks have shown good discriminative performance and tend to preserve structures well with less blur in output image. **NOTE** Uses memory to load VGG model. **REFERENCE** Johnson et al. *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*, arXiv, March 27, 2016 (https://arxiv.org/pdf/1603.08155.pdf).
OrsPerceptualLoss	Designed for use with Gan-style models, this loss function includes style loss (4 activation layers of VGG), feature loss (1 activation layer of VGG), TV loss, and MSE loss. This loss function tends to preserve structures well with less blur in output images. **NOTE** Uses memory to load VGG model, weights of losses in total loss is fixed. **REFERENCE** Johnson et al. *Perceptual Losses for Real-Time Style Transfer and Super-Resolution*, arXiv, March 27, 2016 (https://arxiv.org/pdf/1603.08155.pdf).

Custom Layers

In some cases, you may want to edit the network structure of a deep learning model by modifying the combination of pre-defined components (functions) called layers. In addition to all the layers previously defined for the Deep Learning Tool, this software release includes three additional custom layers — OrsDiscreteWaveletTransform ,OrsInverseWaveletTransform, and OrsConvolveBatches.

Click the **Go to Editing** button at the bottom of the Deep Learning Tool dialog to open the Model Editing panel. You can view the available layers in the **Select New Layer Type** drop-down menu, as shown below.

Model Editing panel

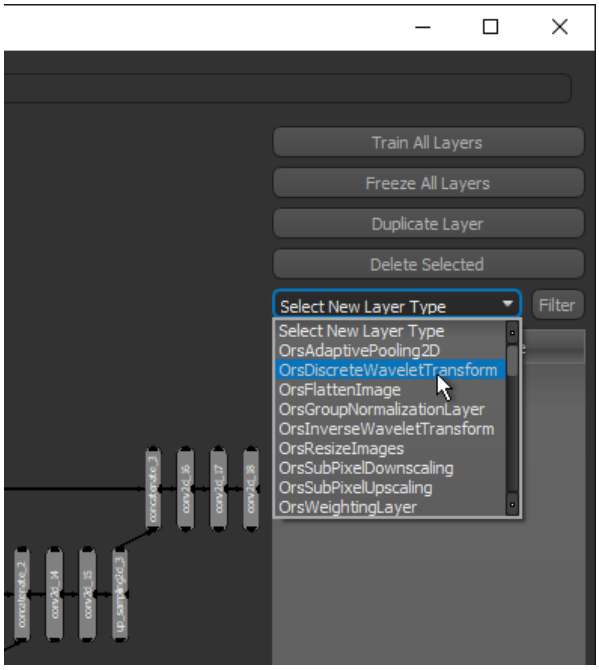

The parameters of each layer can be set by selecting the layer and then setting the values in the side panel on the right. The model can be edited by dragging these layers on the workspace and then defining the connection graph between the layers.

Custom layers

Layer	Description
OrsDiscreteWaveletTransform	This custom layer performs a discrete wavelet transform using Haar wavelets on input images. The output is four quarter-sized images. This layer can be used to replace pooling layers as it enlarges the perceptual field, but unlike pooling is perfectly inversible. **NOTE** Outputs quadruple channel size and may complicate models. **REFERENCE** Liu et al. *Multi-Level Wavelet Convolutional Neural Networks*, IEEE Access, Volume 7, 2019 (https://ieeexplore.ieee.org/abstract/document/8732332).
OrsInverseWaveletTransform	This layer performs an inverse wavelet transform using Haar wavelets on input images. The layer takes four channels as input and outputs one channel with a size four times larger than the inputs. Used as the up-sampling scheme for discrete wavelet transform.

Layer	Description
	Reference Liu et al. *Multi-Level Wavelet Convolutional Neural Networks*, IEEE Access, Volume 7, 2019 (https://ieeexplore.ieee.org/abstract/document/8732332).
OrsConvolveBatches	This layer performs a convolution between images and kernels batch-wise and can be used to handle large ranges of motion during sparse-view interpolation. Kernel size must be specified and the layer does not perform any padding. **Note** Equivalent operation to iterating a kernel over a batch but is very memory intensive. **Reference** Niklaus et al. *Video Frame Interpolation via Adaptive Separable Convolution*, arXiv, August 5, 2017 (https://arxiv.org/pdf/1708.01692.pdf).

Additional Data Augmentation Setting

Data augmentation is the process of increasing the size of a dataset by transforming it in ways that a neural network is unlikely to learn by itself. This software release provides an additional data augmentation setting called 'Jitter'. If selected, a jitter effect will be applied randomly to input and output patches within the set range. You should note that in many cases jitter can help prevent overfitting during training.

Data augmentation settings are available on the Training screen's Inputs tab, as shown below.

Data augmentation settings

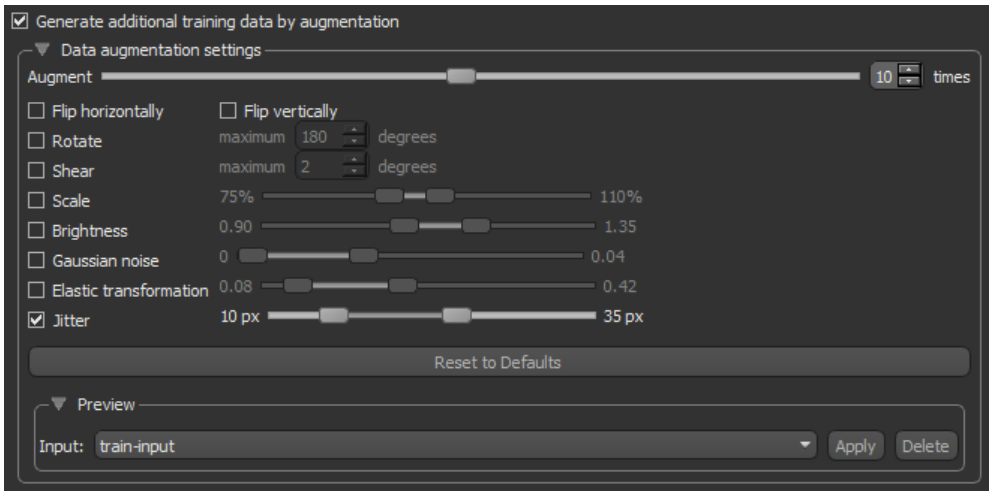

Note This additional data augmentation setting is also available in the Segmentation Wizard on the Data Augmentation tab in the Training Parameters dialog and for custom model architectures on the Train tab.

New Macro Blocks for Deep Learning

A number of new Macro Builder blocks, shown below, are available in Dragonfly 2022.2 for automating deep learning model training with Dragonfly's Macro Builder.

Choose **Utilities** > **Macro Builder** on the menu bar to open Dragonfly's Macro Builder.

New Macro Builder blocks for deep learning model training

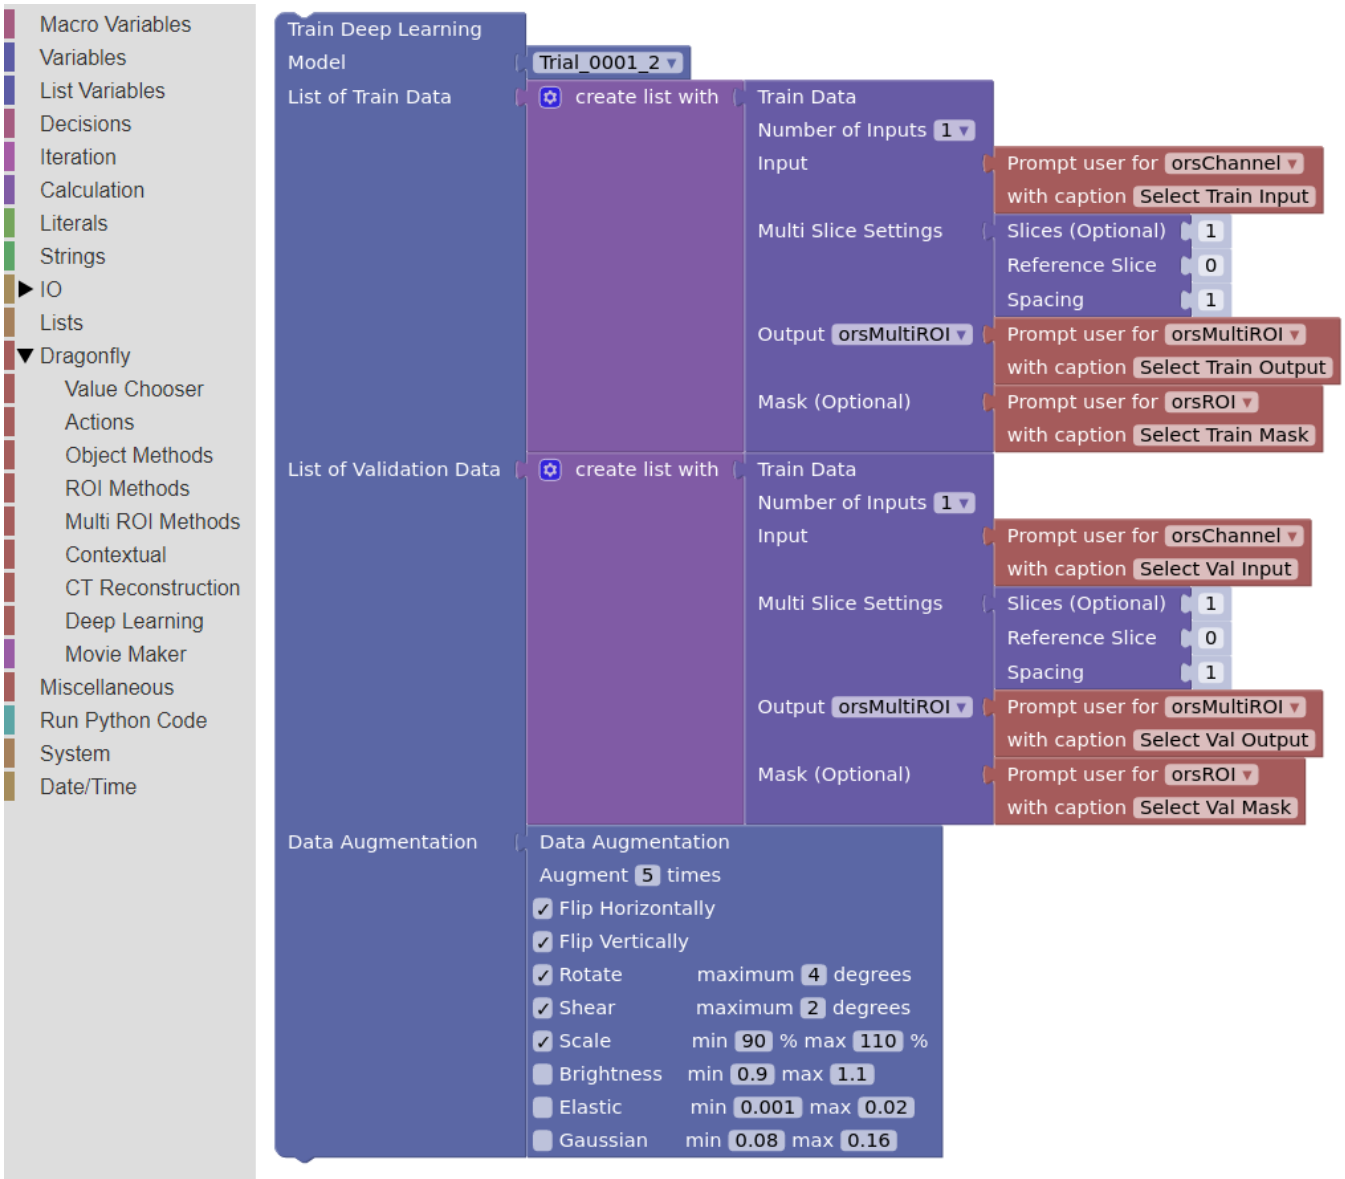

Segmentation Wizard

Pre-trained models, the option to train deep models on multiple GPUs, and excluding image modalities from training are available for Dragonfly's Segmentation Wizard in this software release. Other updates include new auto compute settings and the option to add trained deep models to user-defined strategies.

Pre-Trained Models

You can now download a selection of deep models pre-trained by the Dragonfly Team. In many cases, starting with a pre-trained model provides better and faster results with smaller training sets than using an untrained model. In the case of using pre-trained models for semantic segmentation tasks, you take advantage of learned feature maps to reduce laborious and time-consuming labeling of training sets. Refer to the topic Pre-Trained Models on page 9 for descriptions of the available pre-trained models and additional information about working with pre-trained models.

You can add pre-trained models to a user-defined strategy, as described below. Selected models will be downloaded to a local cache the first time that they are imported and then a copy will be made for the current Segmentation Wizard session. Subsequent models will be copied from the local cache. You should note that Dragonfly will check for new or updated pre-trained models each time that the application is started, provided that your system is connected to the Internet. You should also note that cached models will be available even when Dragonfly runs offline.

Do the following to add a pre-trained model to a user-defined strategy:

1. Click the **Model Generation Strategy** icon on the Models tab.
2. Select or create a user-defined strategy in the Model Generation Strategy dialog.
3. Click the **Add New Model** button.
4. Scroll to and select **Pre-Trained (by the Dragonfly Team)** in the Select drop-down menu, as shown in the following screen capture.

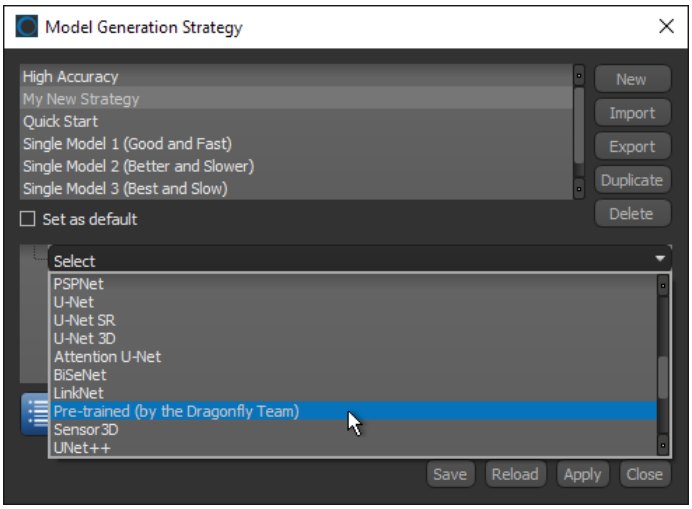

5. Double-click inside the editable Pre-trained model field and then choose the required model in the drop-down menu, as shown below.

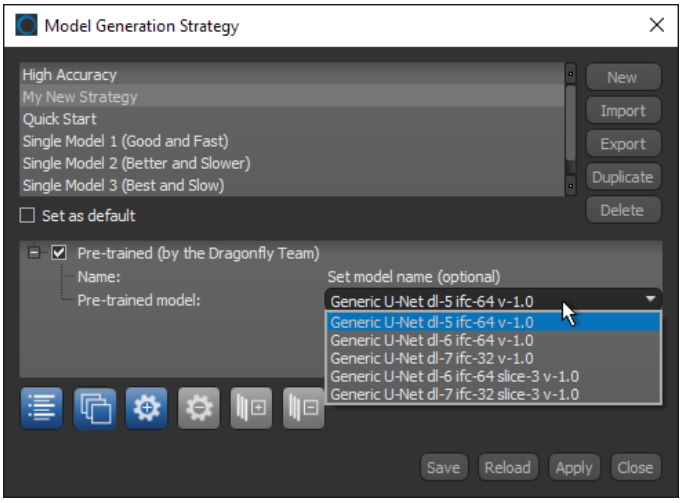

6. Click the **Apply** button to add the selected model to the strategy.

   The selected model will be downloaded, if required, and then copied to the current Segmentation Wizard session.

Deep Model Training on Multiple GPUs

In some cases, you may find it advantageous to train deep models on multiple GPUs instead of a single GPU (see Deep Model Training on Multiple GPUs on page 14).

The option for selecting the number of GPU cards to use for model training is available on the Settings tab, as shown below.

Settings tab

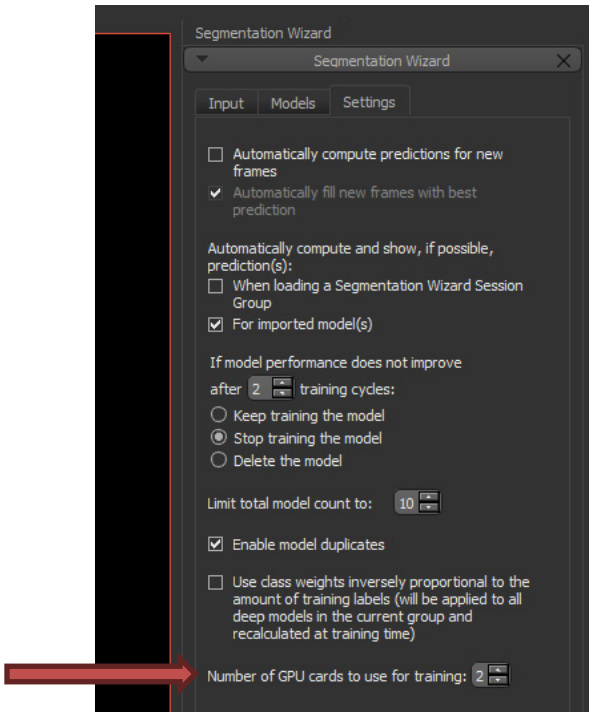

Auto Compute When Loading Session Groups

The option to automatically compute and show model predictions by default when loading a Segmentation Wizard Session Group is available in this software release. This option is available on the Settings tab, as shown below.

Settings tab

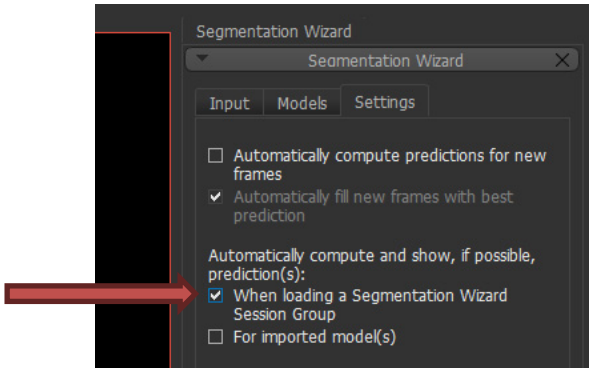

Adding Models to a User Strategy

You can now add trained deep models to a user-defined strategy from the current Segmentation Wizard session. This can be particularly useful if you often generate the same model(s) with edited training parameters.

Right-click a trained deep model on the Models tab and then choose **Add to Strategy** to add the model to either a new strategy or to an existing user-defined strategy.

Models tab

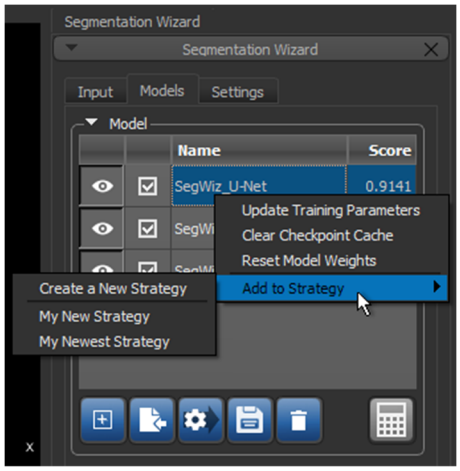

Create a New Strategy... Select this option to add the selected model to a new strategy.

Current strategies list... Select this option to add the selected model to a listed user-defined strategy.

You can also add trained deep models to a user-defined strategy in the Model Generation Strategy dialog. To do so, click the **Add New Model** button on the dialog and then choose a trained model in the Pretrained models list.

NOTE You can edit the name and/or description of the model that was added to a user-defined model generation strategy, as shown below.

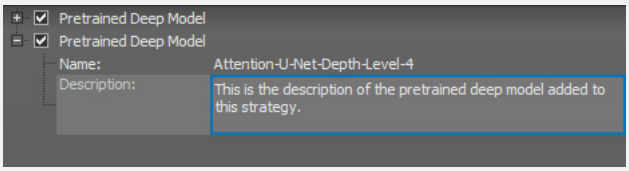

Exclude Image Modalities from Training

In some cases, you might want to exclude an image modality from model training. For example, in cases in which a modality in a multi-modality training set will only be used for labeling voxels. The new option to include or exclude an input modality from model training is available in the Image modalities list, as shown below.

Image modalities

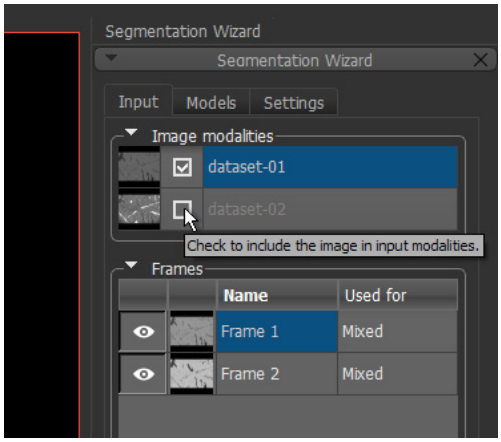

Checked… If checked, the dataset will be included with input modalities during model training.

Unchecked… If not checked, the dataset will be excluded from model training and will only used for labeling. Check to include the image in the input modalities.

Image Quality Metrics

You can now evaluate and compare the image quality of previews and other single-slice datasets with a number of different metrics, such as Sharpness, Shannon Entropy, SNR, and FactorQ, in the newly implemented Image Quality Metrics dialog.

Right-click the data that you want to include in your evaluation and then choose **Open Image Quality Metrics** in the pop-up menu to open the dialog. You can then choose the objects to include in the comparison, apply a mask to limit computations, and select the metrics for comparison. You can also export the results in the comma-separated values file (*.csv extension) format.

Image Quality Metrics dialog

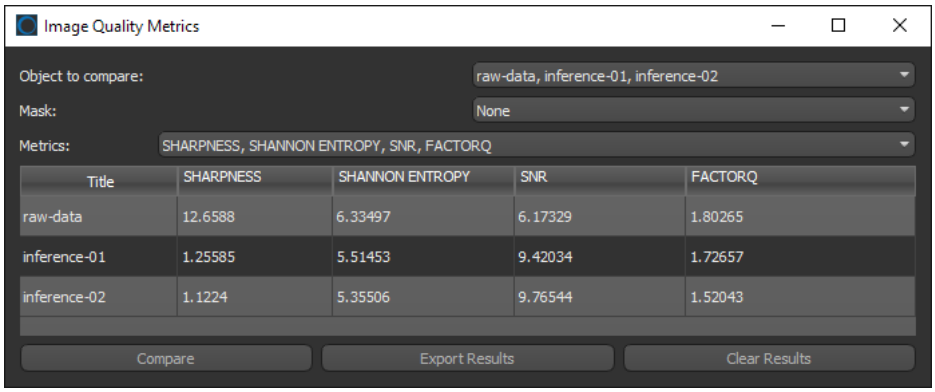

Image quality metrics

Metric	Description
Sharpness	Is a measure of how accurately a sample is represented and is related to the edge contrast of an image. **NOTE** This metric is calculated using the function numpy.gradient (https://numpy.org/doc/stable/reference/generated/numpy.gradient.html).
Shannon Entropy	Is a measure of the uncertainty in predicting voxel values. Decreasing values would mean less uncertainty and better image quality. **NOTE** This metric is calculated using the function skimage.measure.shannon_entropy (https://scikit-image.org/docs/0.13.x/api/skimage.measure.html#skimage.measure.shannon_entropy).
SNR	Is the signal-to-noise ratio, which is defined as the ratio of average and standard deviation of a signal or measurement. The higher the ratio, the better the image quality. **NOTE** This definition is only useful for variables that are always non-negative.
FactorQ	Is a measure of the degree of separation of two material classes in the analyzed image and is calculated on the base of a grey value histogram. **NOTE** FactorQ is only meaningful for datasets with two classes.

Intensity Scale Calibration

This software release provides several updates for calibrating the intensity scale of images to a set of calibration standards. By setting-up a standard calibration, the intensity values of different materials or tissue will be consistent, regardless of the sample acquired. In most cases, you can re-use the same deep model for the inference of different datasets that are calibrated to the same standard.

Right-click the data you need to calibrate and then choose **Calibrate Intensity Scale** in the pop-up menu to open the Intensity Calibration dialog, shown below.

Intensity Calibration dialog

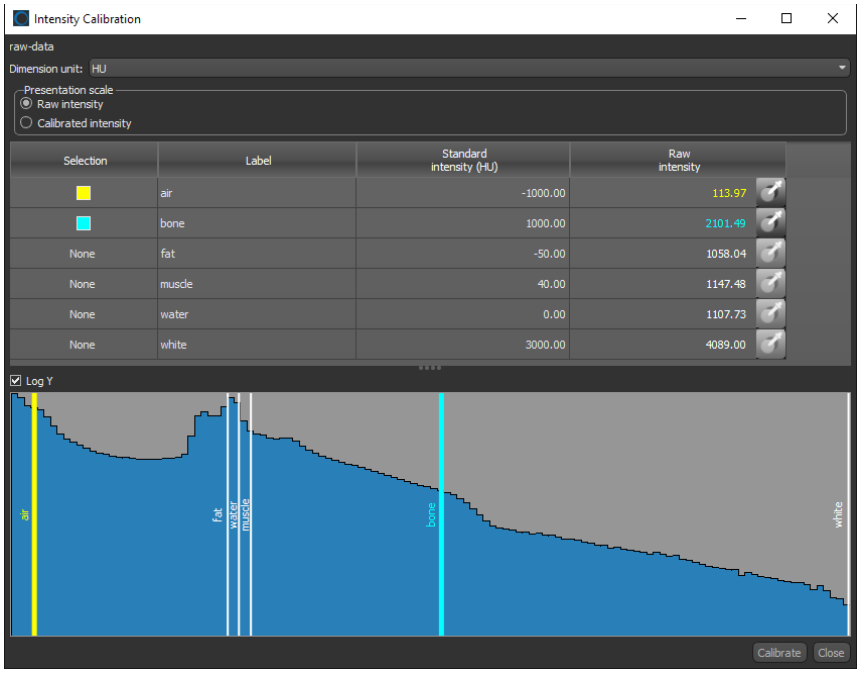

NOTE An intensity scale calibration does not change the data distribution; it only computes a slope/offset transformation so that the selected standard intensity values will be obtained from the raw intensity values. This is the same as manually modifying the slope/offset values in the Image Properties panel.

The following updates are available in the Intensity Calibration dialog:

- You can use the integrated Probe tools to extract the '*mode*' of the raw intensity values of each selection, which is extracted from the histogram and the size of the spherical brush. Mode is the intensity value that corresponds to the center of the bin having the maximum element count. You should note that the histogram is built with 1024 bins spread over the range of data of the dataset.

  You should also note that it is normal to have a different value for mode than for the mean. The advantage of using mode instead of the mean is to help identify

the intensity of a given material that is mixed with other materials, without being influenced by those other materials.

- You can adjust your selections directly on the histogram.
- For calibrated datasets, you can toggle between Raw intensity and Calibrated intensity by selecting a presentation scale.
- You can remove calibrations from raw data by clicking the **Reset Calibration** button. When reset, the offset and slope values will be set to 0 and 1 respectively. A reset button is also available on the Image Properties panel, as shown below.

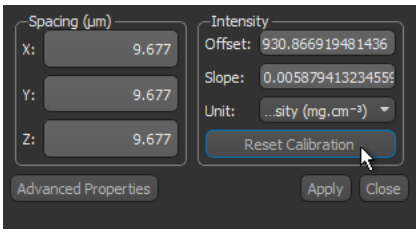

Configurable Actions for Intensity Scale Calibrations

The following actions can be configured for intensity scale calibrations.

Configurable actions for intensity scale calibrations

Action	Default Key
Decrease probe size (in Intensity Calibration)	Mouse wheel down
Increase probe size (in Intensity Calibration)	Mouse wheel up
Probe pre-calibrated value in Intensity Calibration	Left mouse

Preset Calibration Standards

Preset calibration standards for **Hounsfield units (HU)** and **Bone Mineral Density** are now available for calibrating raw data. You can review the set labels and intensity values for these presets in the Dimension Unit Manager.

Choose **Utilities** > **Dimension Unit Manager** on the menu bar to open the Dimension Unit Manager dialog, shown below.

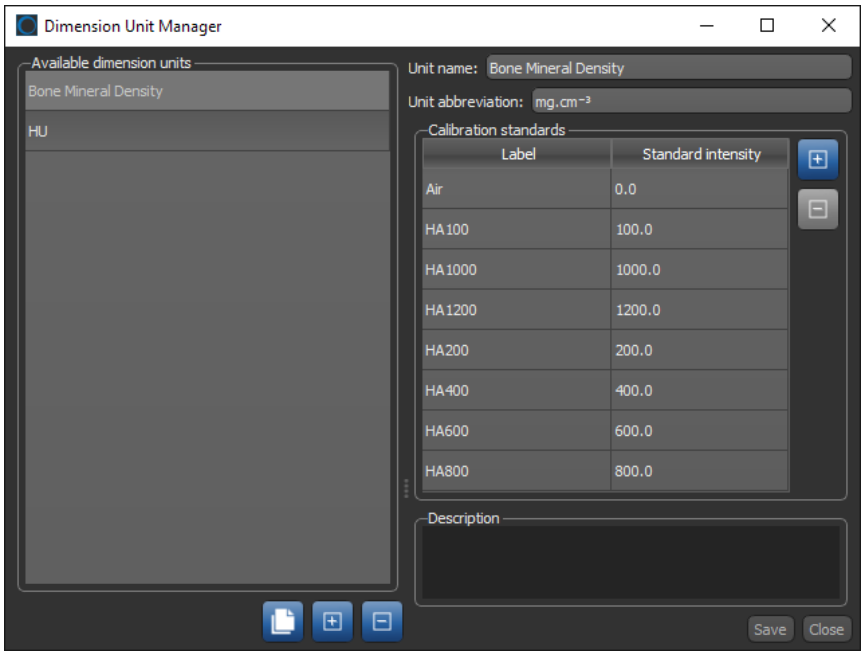

NOTE Refer to the topic Intensity Scale Calibration on page 26 for information about the updates available for calibrating the intensity scale of images. Refer to the topic Updates for Bone Analysis on page 29 for information about bone mineral density measurements.

Elastic Registration

To better understand how a dataset was deformed, this software release provides *non-rigid* or elastic registration for mapping points from one image to the corresponding points in a reference image.

Right-click the fixed image that you want to use as the reference for the registration and then choose **Image Registration** in the pop-up menu to open the Image Registration dialog, shown below. You can then choose **Non-rigid** as the registration type, as well as the mobile image for which the deformable registration is computed and the resulting transform is used to resample it. The output settings include the registered image, displacement modulus, displacement vector fields and generated X, Y, and Z displacements.

Image Registration dialog

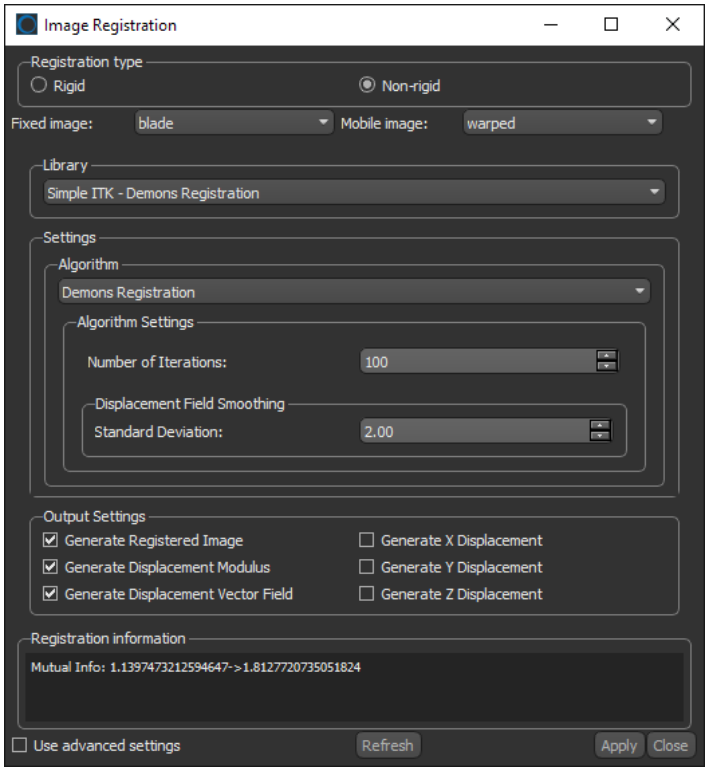

NOTE Rigid and non-rigid registration of images in Dragonfly is based on *elastix* and *SimpleITK*, which are open-source software packages based on the well-known Insight Segmentation and Registration Toolkit (ITK)

You can also apply multiple transforms, assign weights to similarity and penalties metrics, as well as choose settings for the selected optimizer and interpolator, by selecting the **Use advanced settings** option.

The advanced options for *Simple Elastix – Elastic Registration* are shown in the following screen capture.

Advanced settings

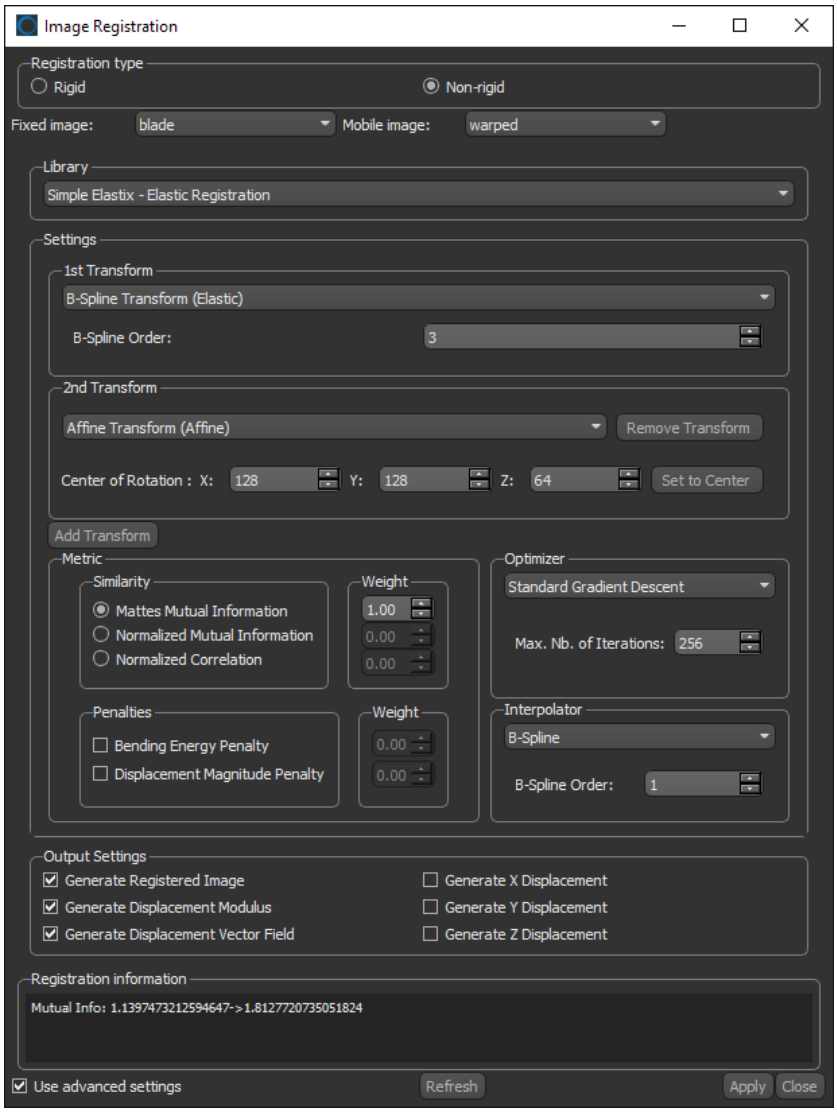

3D Data Augmentation

For some processing and analysis purposes, you may want to create deformed or *warped* datasets. For example, to test invert problems such a non-rigid registration or to create additional data for deep model training. For these purposes, Dragonfly includes a 3D data augmentation generator that lets you create warped datasets from an original dataset.

Right-click the original dataset, region of interest, or multi-ROI and then choose **3D Data Augmentation** in the pop-up menu to open the 3D Data Augmentation dialog, shown below.

3D Data Augmentation dialog

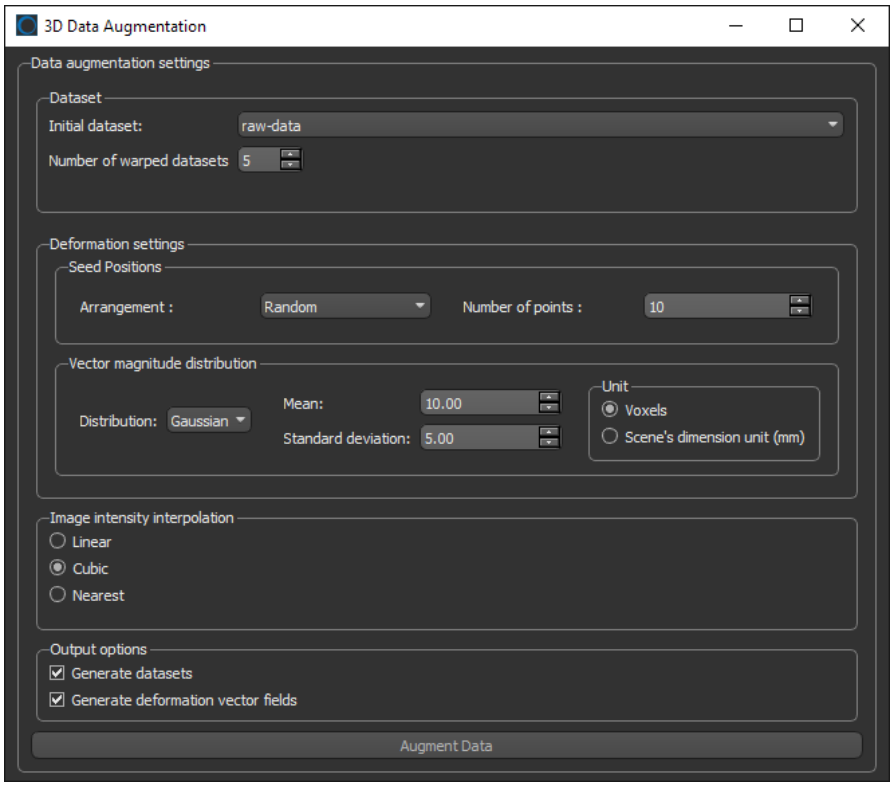

In the 3D Data Augmentation dialog, you can:

- Select the number of outputs required.
- Set the seed positions and vector magnitude distribution, as well as an image intensity interpolation.
- Choose to generate deformation vector fields, in addition to the generated datasets.

An example of a generated warped dataset with its associated deformation vector field is shown below.

Warped dataset with deformation vector field

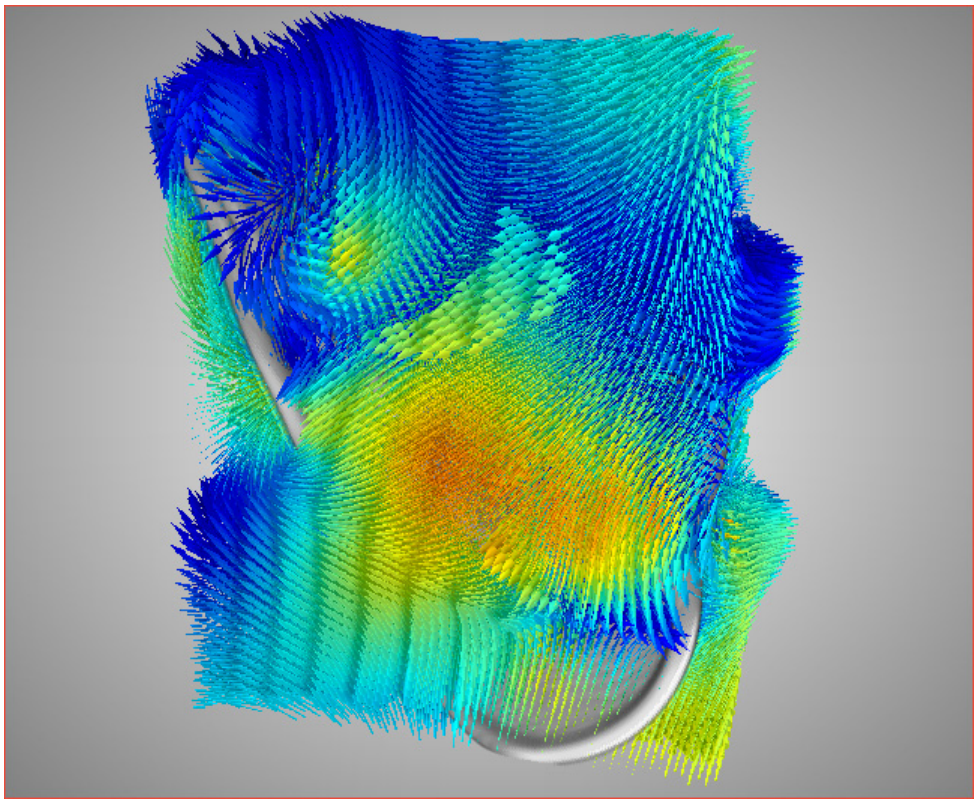

Bone Analysis

This software release includes a number of additional measurements for quantifying bone micro-architectures and measuring bone mineral density. These measurements are available in the Global Measurements dialog whenever you perform a detailed analysis with Dragonfly's Bone Analysis wizard.

You should note that you can now filter the available measurements by input. For example, only measurements related to trabecular bone will appear if you only input a trabecular bone region of interest. Likewise, if you add the filled bone ROI, then the measurements related to those inputs will appear in the dialog, as shown below.

Global Measurements dialog

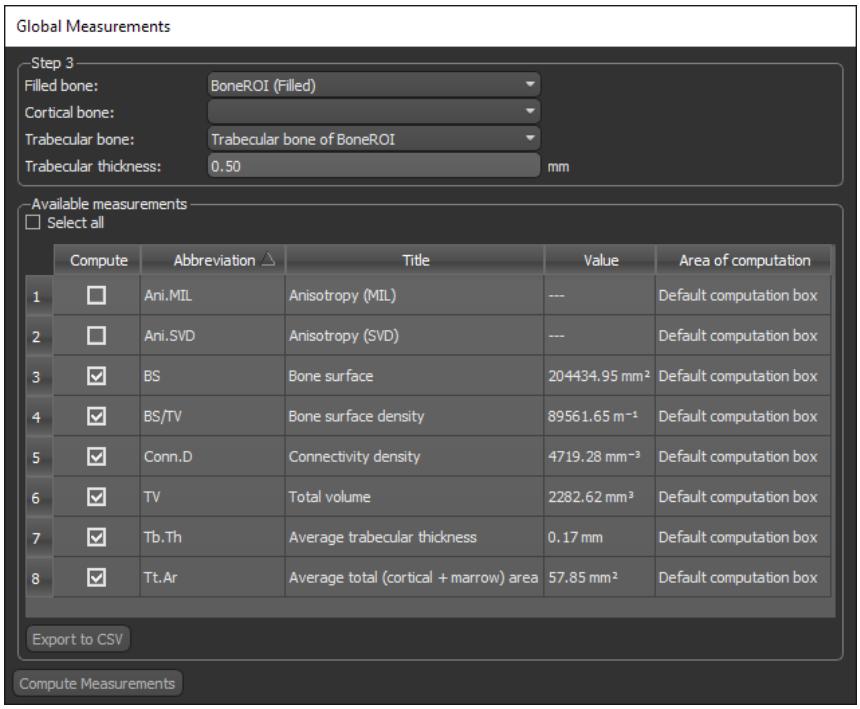

New Global Measurements

The following new measurements are available for providing quantitative descriptions of bone micro-architecture and bone mineral density from segmented bone regions.

NOTE Bone mineral density measurements are only available in Dragonfly Pro. Contact Object Research Systems for information about the availability of Dragonfly Pro.

New global measurements for Bone Analysis

Abbr.	Measurement	Description	Unit
BMD.Me[1]	Bone mineral density (trabecular bone + medullary cavity)	Is the mean intensity of bone mineral density (BMD) computed for calibrated bone data that corresponds to trabecular bone and the medullary cavity. **NOTE** This parameter relates to a mixed bone-soft tissue region. Refer to **TMD.BV** and **TMD.Co** for measurements of the mineral density of bone itself. **NOTE** See Intensity Calibration for BMD and TMD Computations on page 35 for information about calibrating bone data.	$mg.cm^{-3}$
BS	Bone surface	Is the surface of the trabecular bone and is computed from the input trabecular bone ROI. **NOTE** This parameter can be computed with the marching cube or weighted voxel estimation methods.	(default unit)[2]
BS/BV	Specific bone surface	Is the ratio of the bone surface (trabecular) to the bone volume (trabecular + cortical) and is computed from the input bone ROIs. **NOTE** Surface to volume ratio or "specific bone surface" can be a useful parameter for characterising the complexity of structures.	$(default\ unit)^{-1}$
BS/TV	Bone surface density	Is the ratio of the trabecular bone surface to the total volume of the mixed bone-soft tissue region (total volume).	$(default\ unit)^{-1}$
Conn.D	Connectivity density	Is a measure of the degree of connectivity of trabeculae normalized by the volume of the trabecular bone region of interest.	$(default\ unit)^{-3}$
Ct.Po	Cortical porosity	Is the ratio of pores to the total volume in the input cortical bone region and is computed as: [total pore volume (Po.V)/total volume of cortical bone (Ct.V)].	none
SMI[2]	Structure model index	Is an indicator of the relative structure of the trabeculae and involves measurements of surface convex curvature. Ideal plates, cylinders, and spheres have SMI values of 0, 3 and 4 respectively. Conversely, cylindrical and spherical cavities have SMI values of -3 and -4 respectively. **NOTE** This parameter is of importance in osteoporotic degradation of trabecular bone, which is characterised by a transition from a plate-like to rod-like architecture.	none
Tb.N	Trabecular number	Is a measure of the average number of trabeculae per unit length.	$(default\ unit)^{-1}$
TMD.BV[3]	Tissue mineral density (Bone volume)	Is the mean intensity of tissue mineral density (TMD) computed for calibrated bone data that corresponds to the input cortical and trabecular bone regions. **NOTE** See Intensity Calibration for BMD and TMD Computations on page 35 for information about calibrating bone data.	$mg.cm^{-3}$
TMD.Co[3]	Tissue mineral density (Cortical bone)	Is the mean intensity of tissue mineral density (TMD) computed for calibrated bone data that corresponds to the cortical bone ROI. **NOTE** See Intensity Calibration for BMD and TMD Computations on page 35 for information about calibrating bone data.	$mg.cm^{-3}$

[1] Bone mineral density (BMD) is defined as the volumetric density of calcium hydroxyapatite (CaHA) in a biological tissue in terms of $mg.cm^{-3}$. It is calibrated by means of phantoms with known density values of CaHA.

[2] Refer to Hildebrand T, Ruegsegger P. *Quantification of Bone Microarchitecture with the Structure Model Index,* Comput Methods Biomech Biomed Engin. 1997; 1: 15–23 for more information about SMI.

[3] In contrast to BMD, tissue mineral density (TMD) provides information about the material density of the bone itself and ignores surrounding soft tissue.

Intensity Calibration for BMD and TMD Computations

Measurements of bone mineral density (BMD) and tissue mineral density (TMD) require calibrated bone data, which is calibrated by means of phantoms with a known density of calcium hydroxyapatite (CaHA). Calibration can be done beforehand and is also made available on demand in the Bone Analysis wizard by clicking the **Calibrate Intensity Now** button in the Global Measurements dialog, as shown below.

Global Measurements dialog

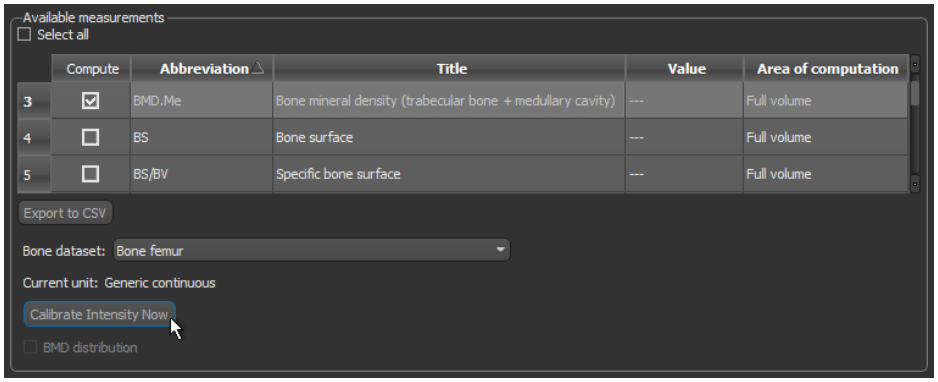

You can then calibrate the data to standard intensity values in the Intensity Calibration dialog, which is shown below.

Intensity Calibration dialog

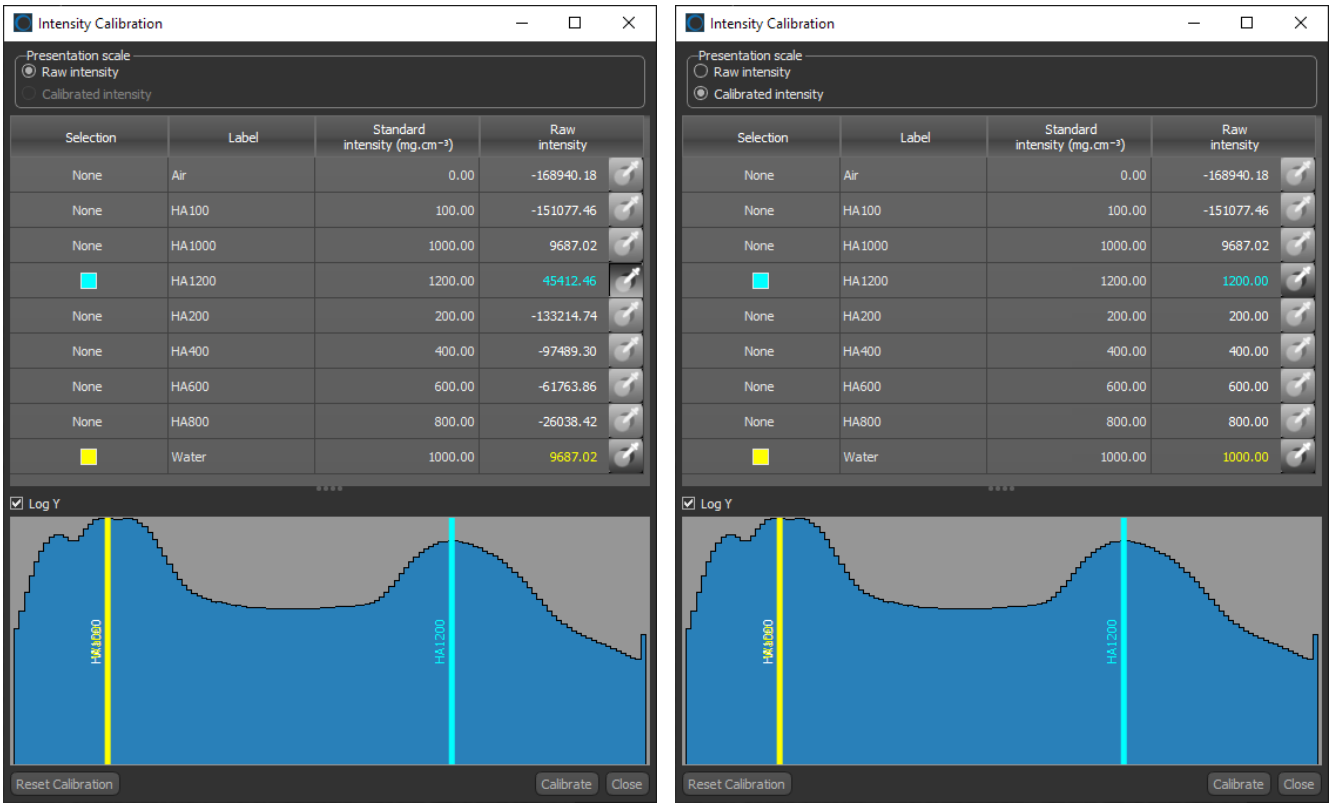

In the Intensity Calibration dialog, you can:

- Select the required calibration points. For example, water and your CaHA phantom.
- Probe the bone data to extract the mode of the raw intensity values.
- Adjust the probed values with the indicators on the histogram.
- Toggle between raw intensity values and calibrated intensity values.

NOTE Refer to the topic Intensity Scale Calibration on page 26 for more information about intensity scale calibrations.

Bone Mineral Density Distributions

This software release includes the option to add histograms of bone mineral density (BMD.Me) and tissue mineral density (TMD.BV and TMD.Co) distributions to Bone Analysis reports. Plotted histograms include labels for the 5th percentile, 95th percentile, mode, mean, and full width half maximum (FWHM) computed for calibrated bone data, as shown below.

Histogram of bone mineral density distribution

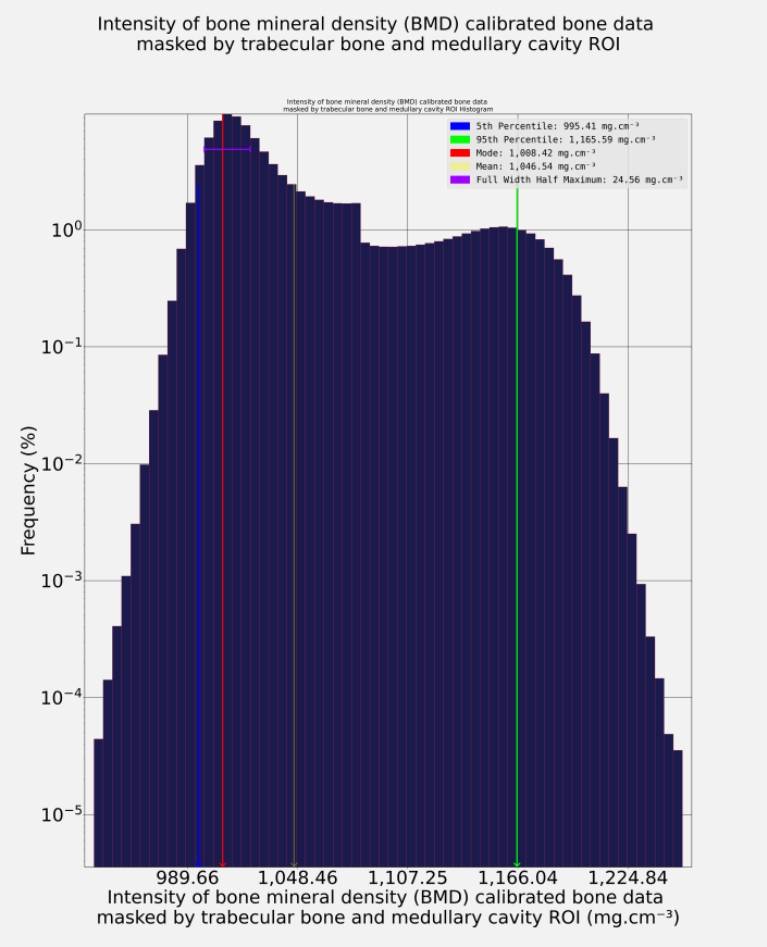

To add a histogram to a report, select and compute the required BMD and/or TMD measurements and then check the BMD and/or TMD distribution options, as circled on the example shown below.

Global Measurements dialog

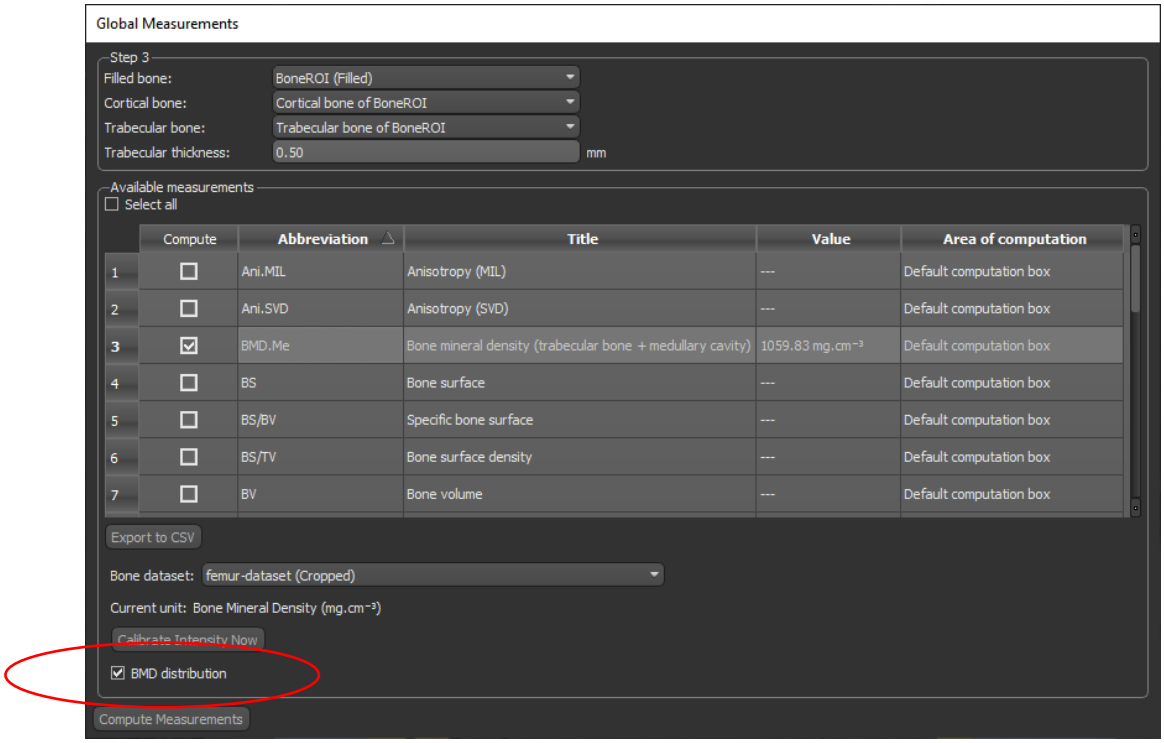

NOTE Bone mineral density measurements and distributions are only available in Dragonfly Pro. Contact Object Research Systems for information about the availability of Dragonfly Pro.

Scalar Generator

A number of new measurements are available in the Scalar Generator for characterizing and measuring object properties, as well as for computing locations as discrete pixel indices.

Spherical Harmonics

In mathematics and physical science, spherical harmonics (SH) are special functions that are defined on the surface of spheres. In Dragonfly 2022.2, you can add scalar values of spherical harmonics to the classes of a multi-ROI to help classify clusters of spherical and quasi-spherical objects and to identify outliers. In some cases, spherical harmonics may provide a more robust measurement than 'roughness' for particle analysis and can be used as a feature vector for machine learning classification tasks.

Right-click the multi-ROI to which you want to add scalar values of spherical harmonics and then choose **Scalar Generator** in the pop-up menu. Spherical harmonics are available in the Scalar Generator dialog, as shown below.

Scalar Generator

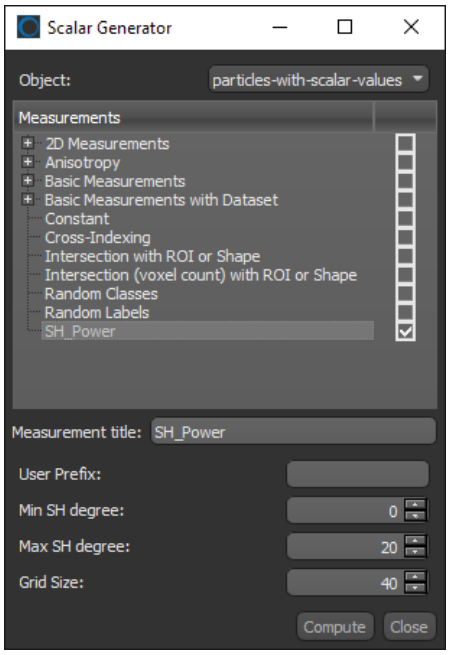

In the Scalar Generator dialog, you can:

- Add a user prefix to scalar values of spherical harmonics, optional.
- Select the minimum and maximum power spectrum that will be computed.
- Select a 'grid size', which determines the number of rays projected to make the projection map.

2D histogram of plotted SH powers in the Object Analysis dialog

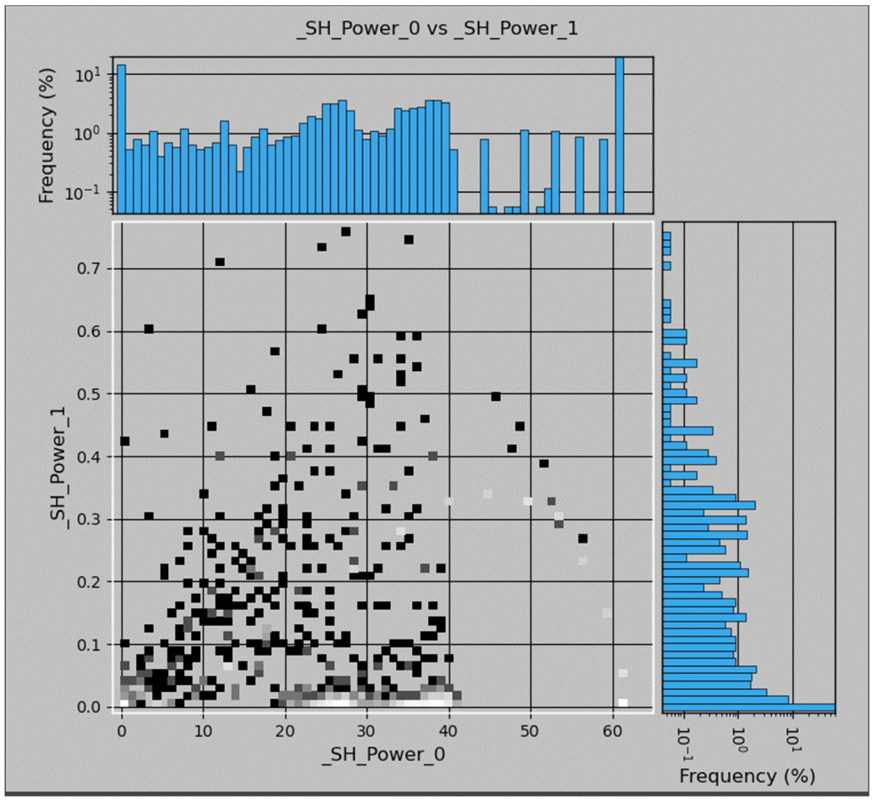

NOTE For simulation and experimentation purposes, you can generate meshes from spherical harmonics (see Generate Meshes from Spherical Harmonics on page 72).

References

Refer to the following for more information about spherical harmonics and the implemented Python software package used for spherical harmonic transforms.

- Overview of spherical harmonics: (https://mathworld.wolfram.com/SphericalHarmonic.html).
- Integration package of Spherical Harmonic Tools (pyshtools 4.10): (https://shtools.github.io/SHTOOLS/index.html).
- Tutorial that demonstrates how to analyse global data on the sphere using spherical harmonic functions: (https://nbviewer.org/github/SHTOOLS/SHTOOLS/blob/master/examples/notebooks/low-level-spherical-harmonic-analyses.ipynb).

Anisotropy

This software release provides the opportunity compute anisotropy for the classes within a multi-ROI using the mean intercept length (MIL) or star volume distribution (SVD) method. The settings for computing the degree of anisotropy are available in the Scalar Generator dialog, as shown below.

Scalar Generator dialog

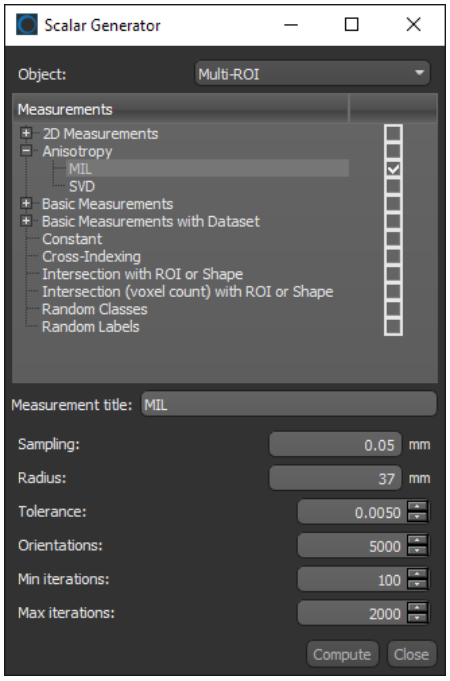

NOTE Anisotropy is a measure of how highly oriented substructures are within a volume. For an isotropic (perfectly oriented) system, the degree of anisotropy (DA) is equal to 0. As the system becomes more anisotropic (less well-oriented), the DA increases to some value less than 1.

Refer to the table below for a description of the settings applicable to anisotropy.

Settings for computing anisotropy

Parameter	Description
Sampling	Is the resolution, or distance between subsequent samples along each vector. The entered value can be equal to the voxel size of the input multi-ROI.
Radius	Is the radius of the sampling sphere, which determines the length of each sampling vector.
Tolerance	Is the coefficient of variation. Sampling new random points will continue until either a coefficient of variation equal to the tolerance is reached or the maximum number of iterations is completed.
Orientations	Is the number of lines to analyze per sampling sphere.
Min iterations	Is the minimum number of random points in the sample that will be analyzed.
Max iterations	Is the maximum number of random points in the sample that will be analyzed. Fitting will stop automatically after this number is completed or if the coefficient of variation (tolerance) is reached.

Basic Measurements

The following basic measurements are now available for multi-ROIs in this software release.

Basic measurements

Measurement	Description
Center of Mass Index X/Y/Z	Indicates the X, Y, or Z voxel index of the center of mass, which is calculated from the centroid of the object.
Min Index X/Y/Z	Indicates the minimum X, Y, or Z voxel index of an object.
Max Index X/Y/Z	Indicates the maximum X, Y, or Z voxel index of an object.

Basic Measurements with Datasets

The following new basic measurements with datasets are available for multi-ROIs in this software release.

Basic measurements with datasets

Measurement	Description
Weighted Center of Mass X/Y/Z Index	Indicates the X, Y, or Z voxel index of the weighted center of mass.
Position of Minimum Value X/Y/Z	Indicates the X, Y, or Z position (in default units) of the minimum value within each object.
Position of Maximum Value X/Y/Z	Indicates the X, Y, or Z position (in default units) of the maximum value within each object.
Index of Minimum Value X/Y/Z	Indicates the X, Y, or Z voxel index of the minimum value within each object.
Index of Maximum Value X/Y/Z	Indicates the X, Y, or Z voxel index of the maximum value within each object.

Artificial Data Generator

This software release features an integration of the algorithms used in Pore Network Modeling (OpenPNM) and Porous Microstructure Analysis (PuMA) for generating artificial image data and graphs.

Choose **Utilities** > **Artificial Data Generator** on the menu bar to open the Create an Object dialog, shown below.

Create an Object dialog

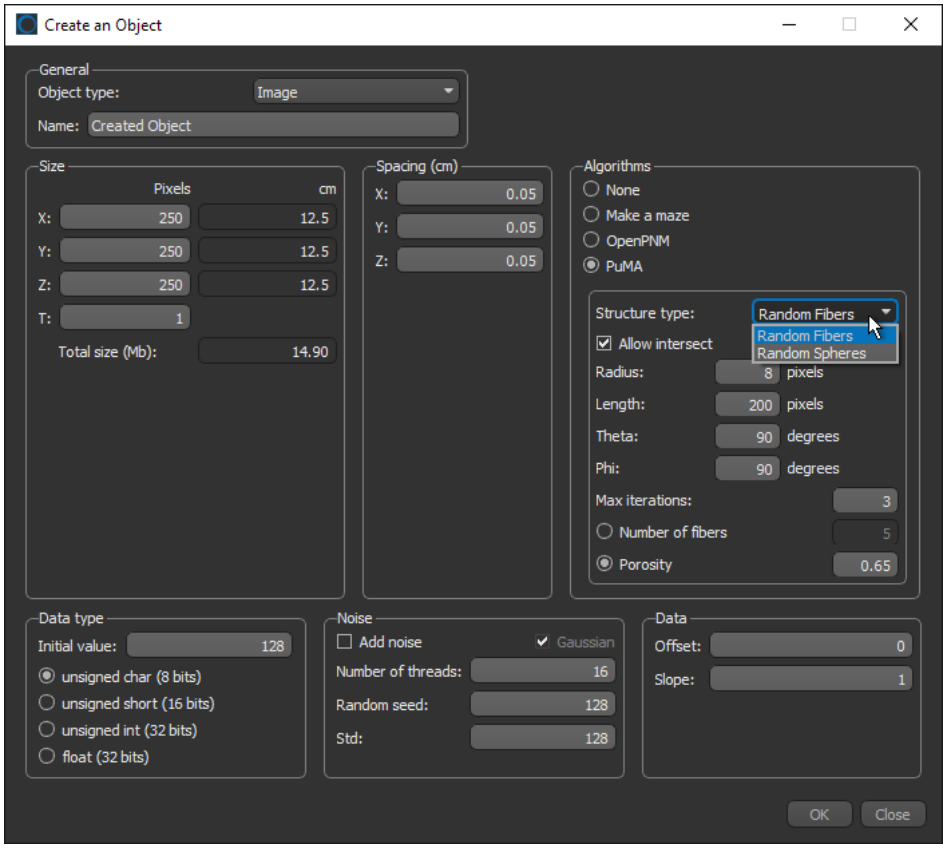

OpenPNM

The options for generating artificial image data for pore network simulations, which include blobs, cylinders, lattice spheres, overlapping spheres, and Voronoi edges, are available in the Create an Object dialog whenever **OpenPNM** is selected as the algorithm for image data. These structure types are available in the Structure type drop-down menu.

Generated artificial pore network

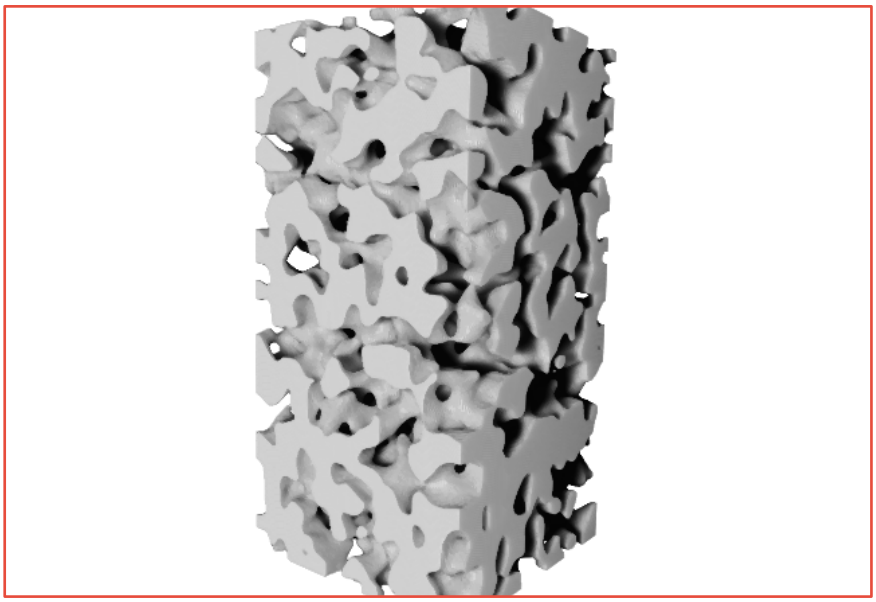

You also generate graphs for pore network simulations with selectable geometry and connectivity by selecting **Graph** as the Object type.

Generated graph

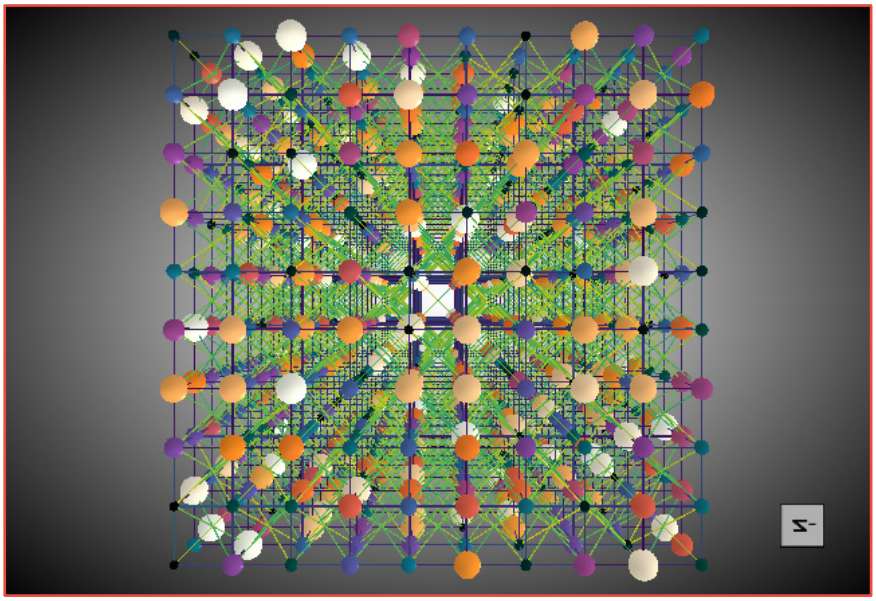

PuMA

You can now generate random fiber and random sphere specimens with PuMA's Sample Generator whenever **PuMA** is selected as the algorithm. You can select a size for the specimen and effective porosity, as well as specify a number of additional parameters such as radius, length, orientation, and intersection for fiber structure types. The parameters for sphere structure types include diameter and intersection.

Random fiber specimen

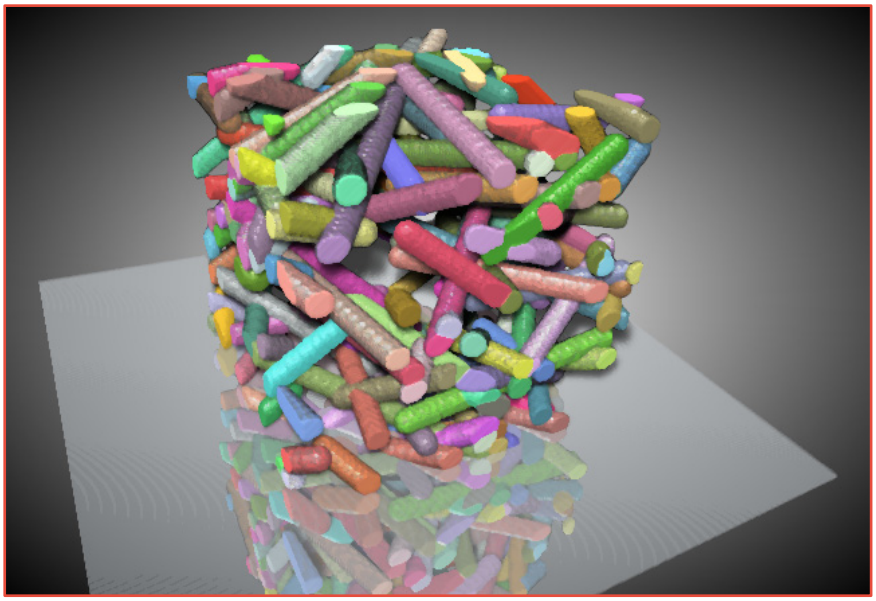

Porous Microstructure Analysis (PuMA)

A number of updates for performing material response simulations on microstructures of porous media with Dragonfly's **Porous Microstructure Analysis (PuMA)** module are available in this software release. These updates are described below.

Elasticity Computation Type

This software release includes the addition of **Elasticity** as a computation type for investigating effective material properties and for performing material response simulations on microstructures of porous media. You should note that the required inputs for elasticity computations is a multi-ROI, the parameters for the labels, and a direction or directions for solving.

PuMA dialog with results for elasticity

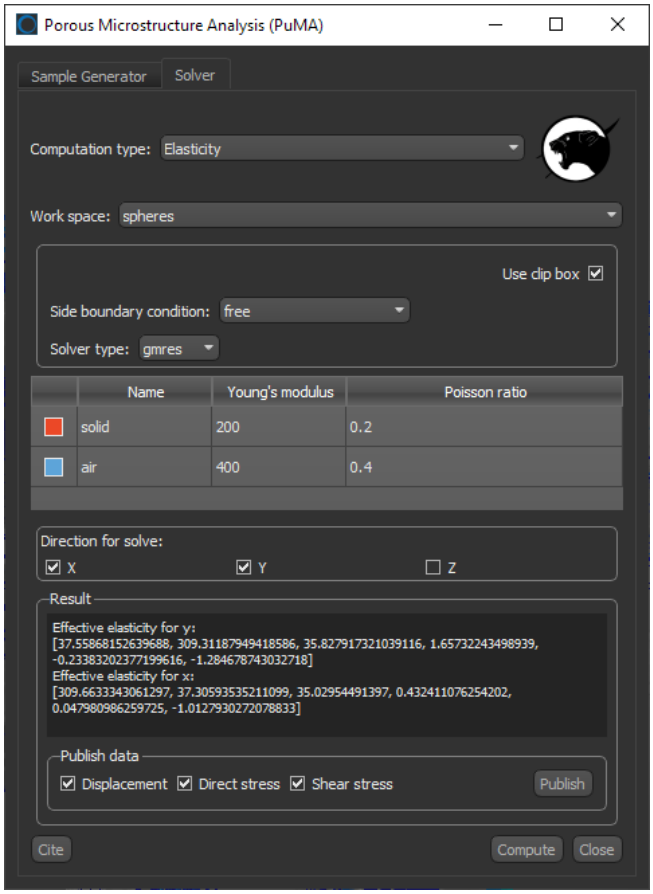

You can also generate vector fields to visualize published data of the displacement, direct stress, and shear stress coefficients computed for effective elasticity, as shown below.

Vector field of direct stress

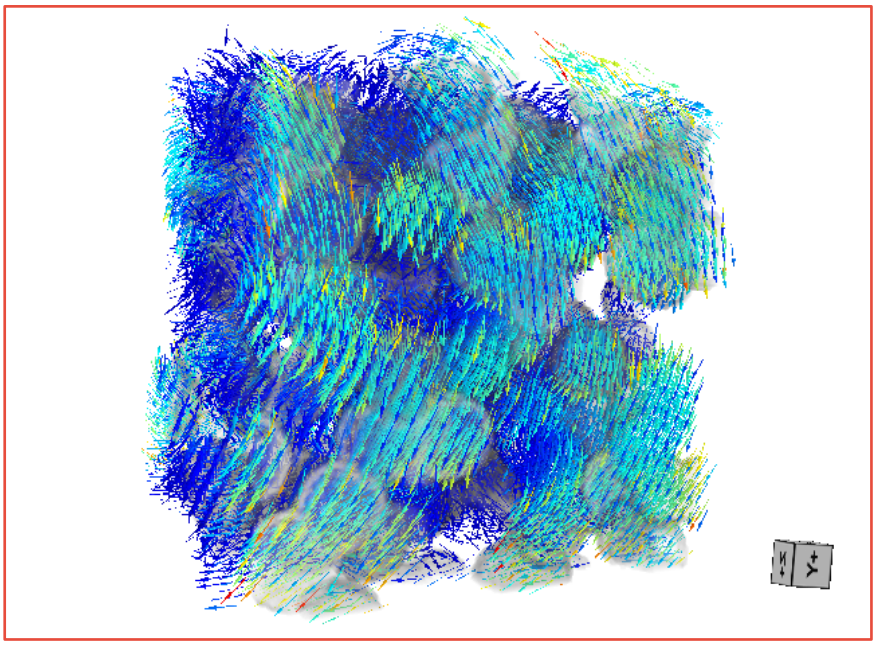

Limit Computations to Data Subsets

Computations of continuum tortuosity, elasticity, orientation detection, permeability, and thermal conductivity can now be limited to a subset of the data contained within a user-defined box or within a clip box applied to the workspace. These options are available in the computation settings box, as circled below.

PuMA dialog

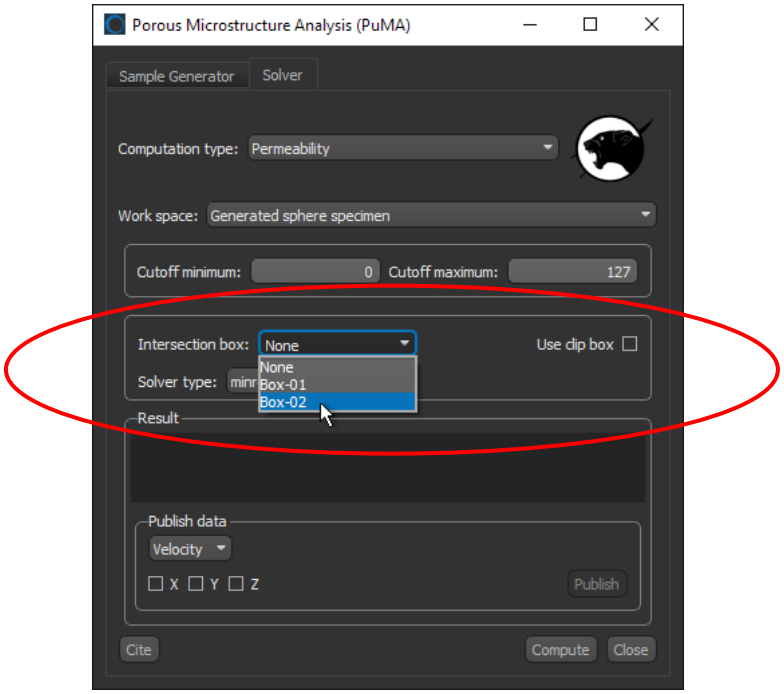

CT Reconstruction

This software release features a refactoring of the CT Reconstruction dialog, as well as the integration of RTK for the reconstruction of parallel-beam projections and the addition of SART for the reconstruction of cone-beam projections. You should note that RTK is an open-source software package based on the Insight Toolkit (ITK).

Choose **Workflows** > **CT Reconstruction** on the menu bar to open the CT Reconstruction dialog, shown below.

CT Reconstruction dialog

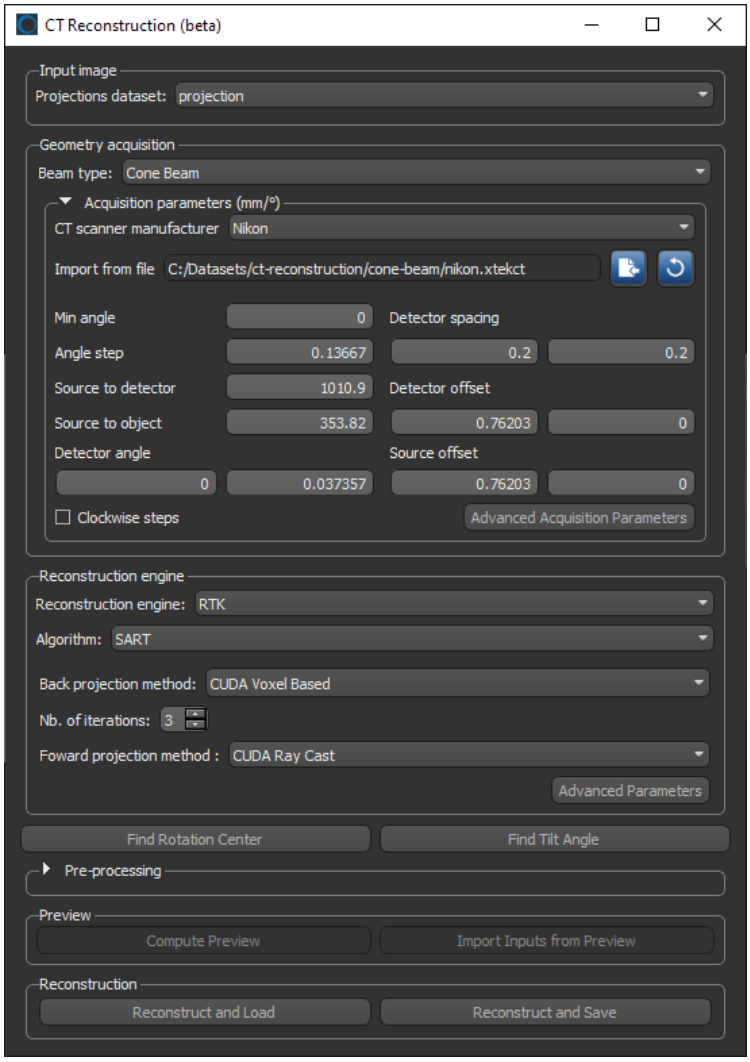

Tilt Angle Finder

In cases in which projection data was acquired at a tilt angle, reconstructions may include artifacts and other anomalies. To help optimize the quality of your reconstructions, this software release includes a CT Tilt Angle Finder.

Click the **Find Title Angle** button in the CT Reconstruction dialog to open the CT Tilt Angle Finder dialog, shown below.

CT Tilt Angle Finder

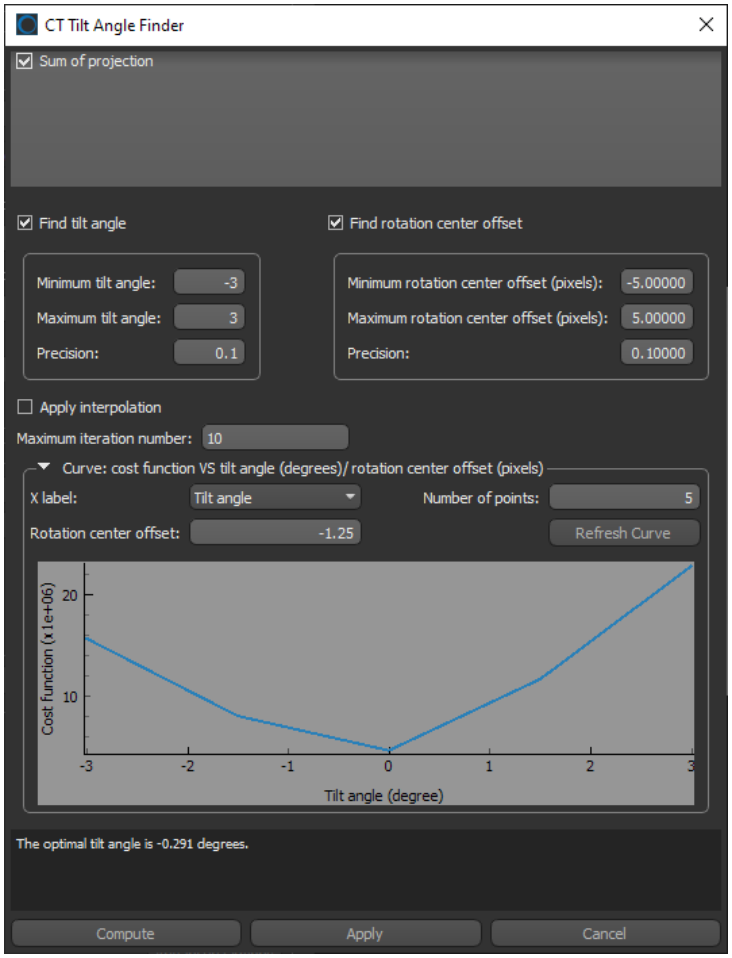

NOTE The CT Tilt Angle Finder dialog also includes an alternate method to find the rotation center offset. In this case, it is best to find the tilt angle first and then the rotation center offset.

Macro Builder Blocks for CT Reconstruction

With new Macro Builder blocks for CT Reconstruction, you can now fully configure and automate the reconstruction of cone beam and parallel beam projections in Dragonfly's Macro Builder.

Choose **Utilities** > **Macro Builder** on the menu bar to open the Macro Builder.

Fully configured macro for reconstructing cone beam projections

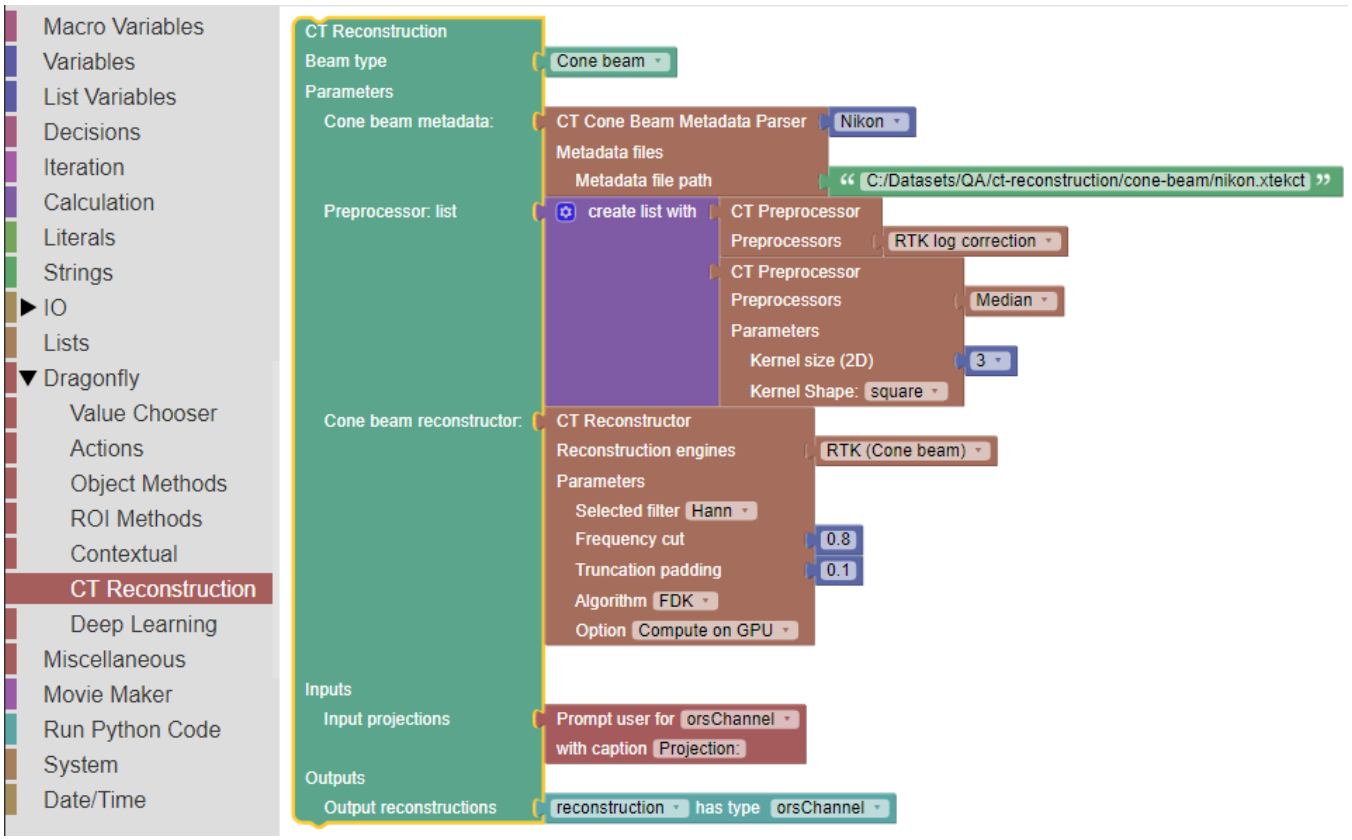

Radiographic Dendrochronology

Dragonfly's new **Radiographic Dendrochronology wizard** provides a guided workflow for analyzing wood samples or planks from furniture or works of art, such as old panel paintings, that were acquired with non-destructive testing techniques. The results from an analysis, which include the width of each tree ring ordered in the direction of growth, can visually inspected and then exported in the Heidelberg format.

Dendrochronology (or tree-ring dating) is the scientific method of dating tree rings (also called growth rings) to the year they were formed and derives from Ancient Greek dendron meaning "tree", and khronos, meaning "time".

Choose **Workflows** > **RAD-DENDRO** on the menu bar to open the Radiographic Dendrochronology Wizard, shown below.

Radiographic Dendrochronology wizard

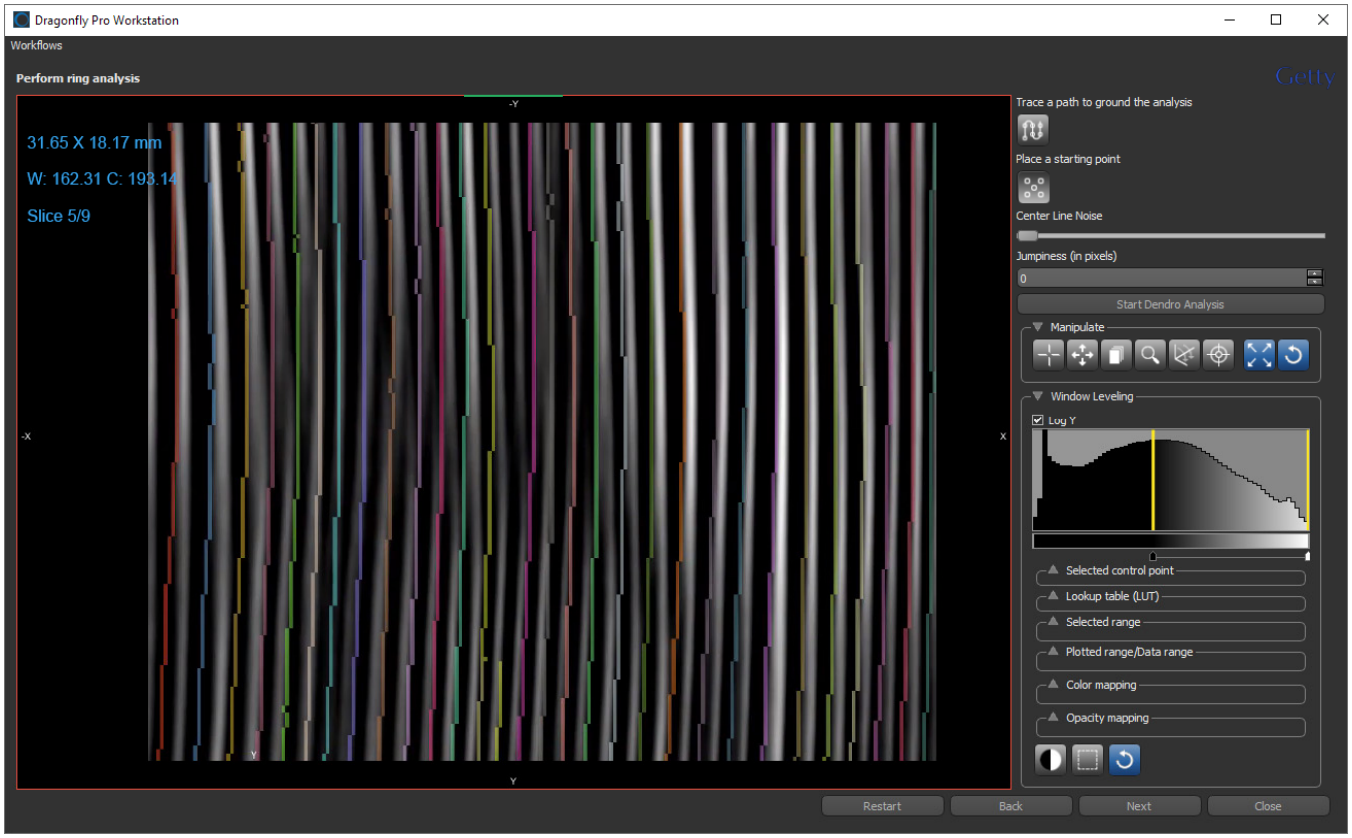

NOTE The Radiographic Dendrochronology wizard was developed by Blake C. Sharp, Arlen Heginbotham, Pascale Fraiture, Andrea Seim from the J. Paul Getty Museum in Los Angeles in conjunction with the Dragonfly team.

The Radiographic Dendrochronology wizard provides the following workflow steps:

- Import and register images that were acquired with high resolution X-ray film or that were reconstructed from CT tomography.
- Crop images to a subset of the input dataset.
- Apply filtering to remove artifacts, such as dust on digitized X-ray films.
- Automatically extract the direction of growth.
- Define a path to automatically analyze and labelize each ring from a chosen starting point.
- Visually inspect the results and export outputs in the Heidelberg format.

Converting Image Data Types

In some cases, you may need to convert the data type of your images. For example, from *unsigned short* to *unsigned char*. In this software release, you can now choose to preserve the original data calibration, in which case the slope and offset values will be modified to match the original data, or to set the values within the usual data range for the selected data type.

Right-click the image data you want to convert and then choose **Modify and Transform > Convert** in the pop-up menu to open the Image Converter panel. The option to preserve the calibration or to set the intensity values within the normalized range is circled below.

Output options for converting image data

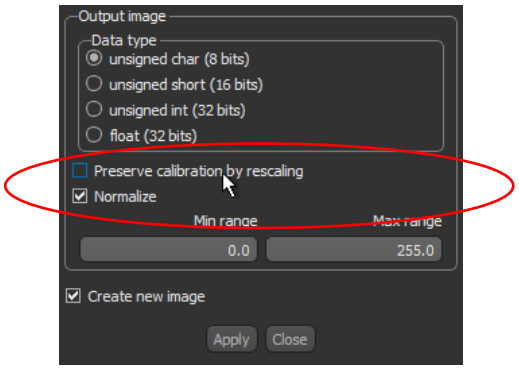

Preserve calibration by rescaling... If checked, the original data calibration will be preserved by transforming the offset and slope values. For example, if you convert an image saved as unsigned short to unsigned char, the intensity values will still range from 0 to 65,000. If unchecked, the offset will be set to 0 and the slope will be set to 1. In this case, intensity values will be within the usual data range of the selected data type. For example, if you convert an image saved as unsigned short to unsigned char, the intensity values will range from 0 to 255.

Normalize... If selected, the data spread will be normalized within the selected Min and Max range values.

Histogram Analysis

The Histogram Analysis dialog now includes the option to compute additional measurements, such as the 5th and 95th percentile, full width half maximum, mode, and sum. The dialog also includes tick marks to indicate measurement locations, as well as a legend that can be superimposed on the plotted values.

Click the **Histogram** tool on the Data Properties and Settings panel to plot a dataset's intensity distribution, as shown in the following screen capture. You can also use the Histogram tool to plot a dataset's intensity distribution within a region of interest and the scalar values within a multi-ROI, mesh, or other object.

Histogram Analysis dialog

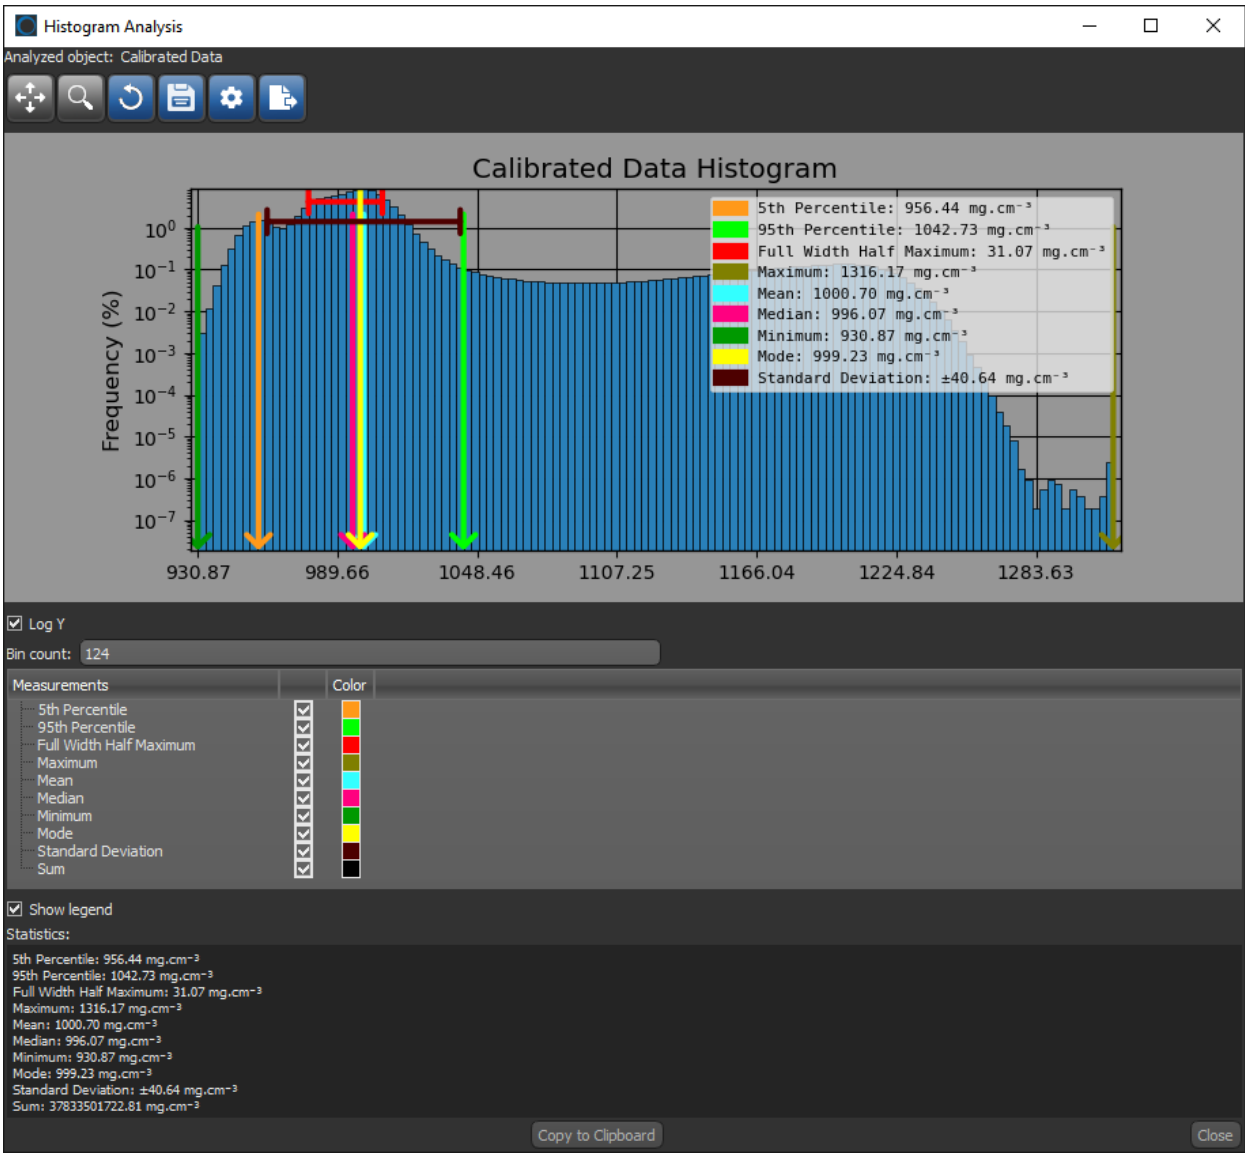

The following measurements are available in the Histogram Analysis dialog.

Histogram Analysis measurements

Measurement	Description
5th Percentile	Is the value below which 5% of the values in the distribution may be found.
95th Percentile	Is the value below which 95% of the values in the distribution may be found.
Full Width Half Maximum	Is the width of the curve as measured between points on the Y-axis which are at the peak half maximum level.
Maximum	Indicates the maximum intensity value.
Mean	Indicates the mean intensity value.
Median	Indicates the median intensity value.
Minimum	Indicates the minimum intensity value.
Mode	Is the most commonly occurring value.
Standard Deviation	Indicates the standard deviation of intensity values. Shown as the range from the mean value on the histogram.
Sum	Is the integral of all intensities of all of the pixels.

Image Filtering

The following new and updated filters are available in Dragonfly 2022.2 to improve image quality and yield information that otherwise would have been missed.

Manual 3D RBF Filter

This software release includes a Manual 3D RBF filter that uses a radial basis function to correct 3D non-uniform illumination and beam hardening effects. You should note that the Manual 3D RBF filter requires at least three(3) seed points to be placed in representative regions on image slices. In contrast to the Manual RBF filter, you can place the seed points on different image slices with the Points tool, which is circled below.

Options for Manual 3D RBF Filter

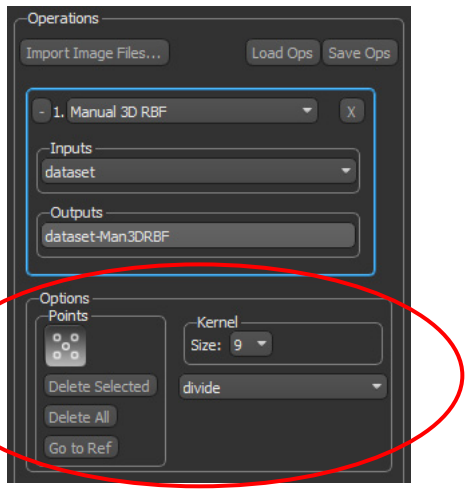

A good indication that a dataset requires shading correction can be provided by the intensity profile of rulers drawn across the image, as shown below.

Intensity profile in Z axis before (in yellow) and after correction (in blue)

Ring Artifact Removal Filter

In cases in which ring artifacts do not originate at the center of a dataset, you can now pick the center of the rings for filtering. You can pick the center of the ring with the Points tool, which is circled below.

Options for the ring artifact removal filter

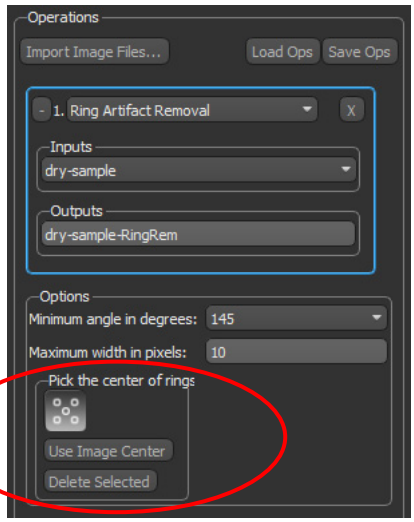

Configurable Actions

The following configurable actions are available for working with the Ring Artifact Removal filter and the Manual 3D RBF filter.

Configurable actions for filtering

Action	Default Key
Create point for ring artifact removal filter	Left mouse
Edit point for ring artifact removal filter	Left mouse
Create points for 3D RBF filter	Left mouse
Edit points for 3D RBF filter	Left mouse

Lossless Compression with ZSTD

This software release features the implementation of **zstd** for saving objects, such as image data, ROIs, meshes, and so on, in the ORS Object file format (*.ORSObject extension), as well as session files (*.ORSSession extension), with lossless compression. Lossless compression is a class of data compression algorithms that allows the original data to be perfectly reconstructed from the compressed data.

NOTE The Zstandard, commonly known by the name of its reference implementation zstd, is a lossless data compression algorithm developed by Yann Collet at Facebook and was released as open-source software in 2016.

The option for saving ORSObject and ORSSession files with lossless compression is available in the Preferences dialog, as shown below.

Miscellaneous Preferences

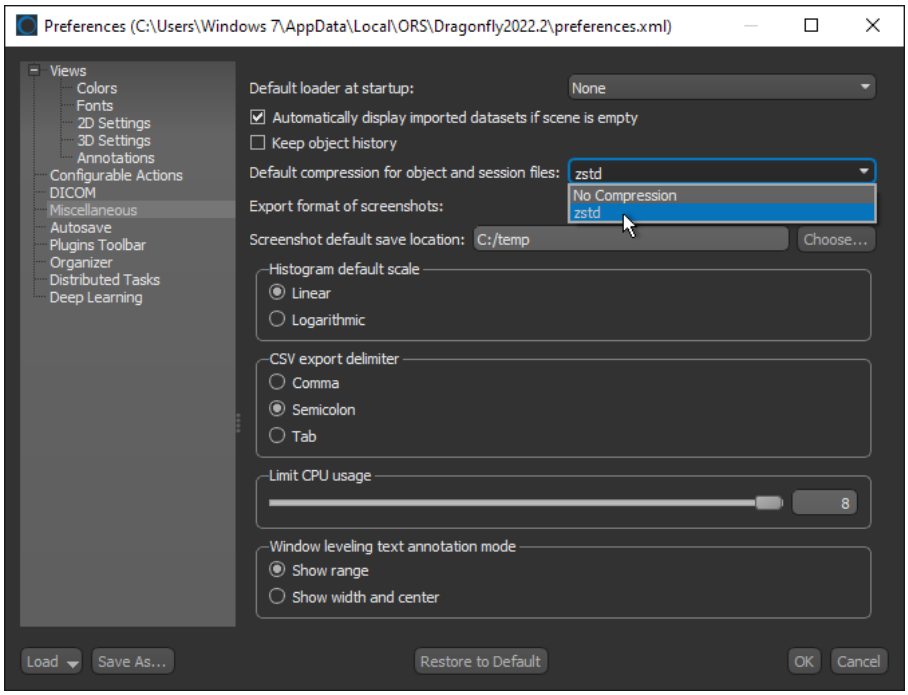

NOTE Compressed object files are NOT compatible with Dragonfly 2022.1 or prior versions. This means that objects saved with compression cannot be loaded in earlier versions of Dragonfly.

3D Renderings and Settings

This software release includes new rendering options and optimizations that should result in improved performance and interactivity, particularly for users who typically deploy **correlative workflows** and work with large datasets.

Additional improvements for 3D visualizations include automatic detection of the appropriate focal length by collision detection with what is in the scene for the Autofocus option, the suppression of moiré effects, and the elimination of darkening in cases in which air surrounds a scanned object.

Multi-volume rendering in a correlative workflow

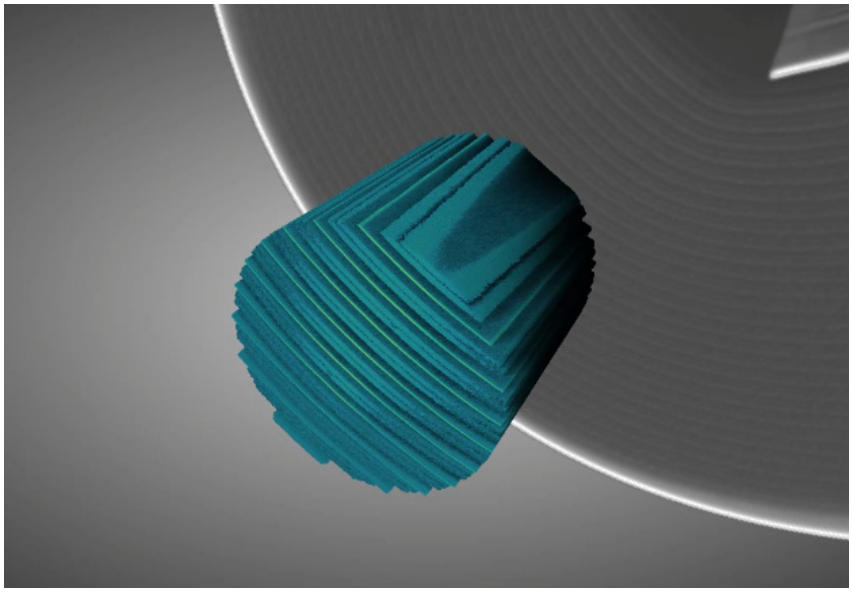

Autofocus feature

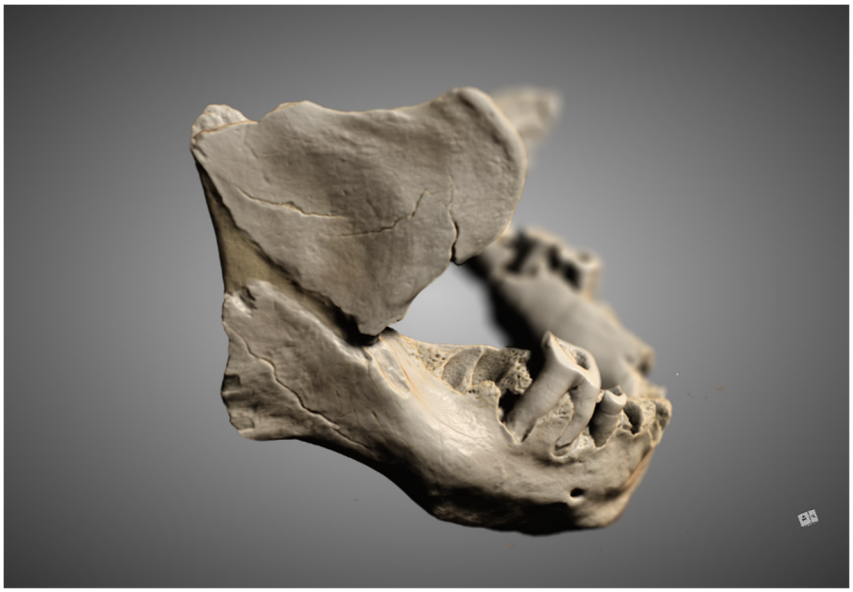

Suppression of moiré effects on the boundaries of a rendered dataset

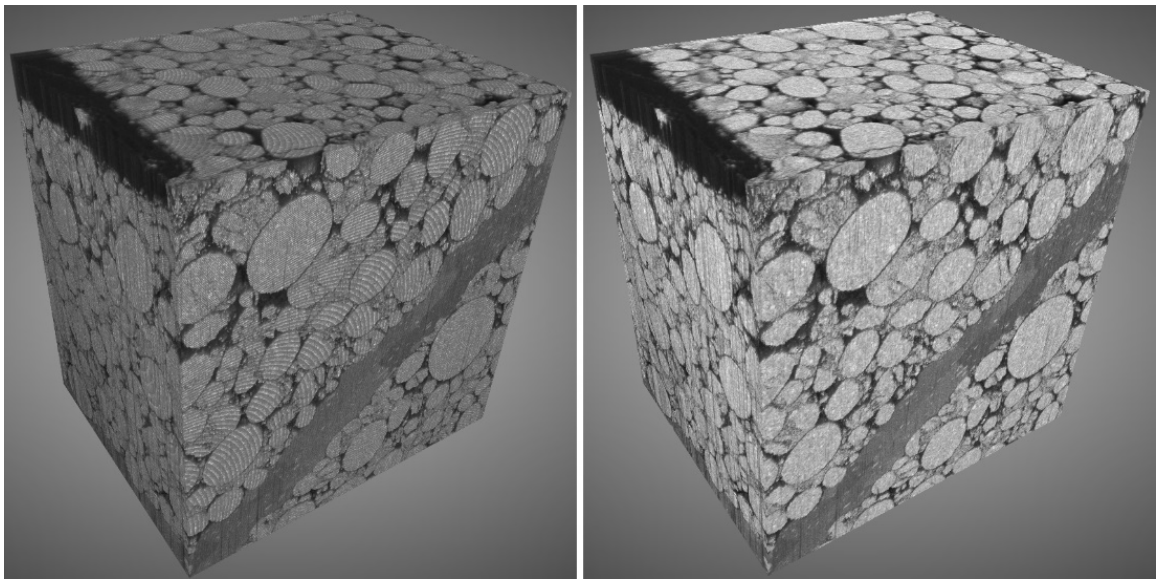

NOTE With improved sampling, and the exclusion of voxel fragments greater than an accumulated threshold at low cost, the **Maximize quality** option has been removed from the Scene's Views Properties panel.

NOTE The option **Flip ROI lighting** was also deprecated in this software release. This is now applied automatically whenever a region of interest or multi-ROI is not facing the camera.

Field of View Control

A new slider on the Scene Views Properties panel, shown below, lets you zoom in and out on objects in a 3D view without moving the camera. In this case, clipping does not occur when zooming in very closely and lighting does not change.

Field of View slider

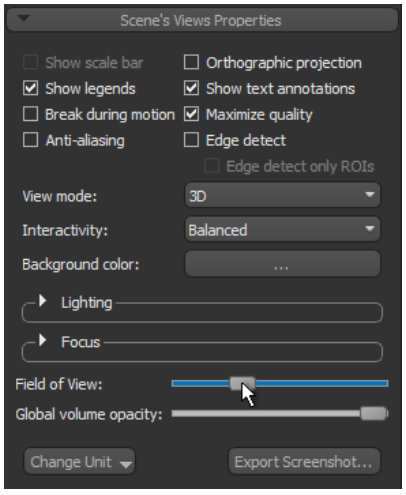

Far and close field of view settings

Plane Settings

In some cases, you may want to make a visual plane solid in order to view the full area covered by the plane. This option is now available in this software release.

Visual plane shown as solid

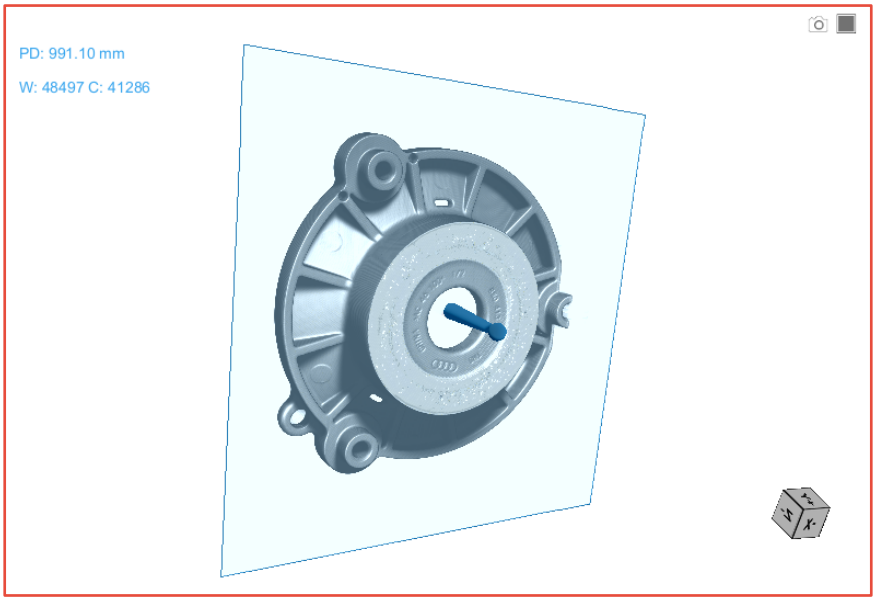

The option to view a visual plane as solid is available in the Plane settings box, as shown below.

Plane settings

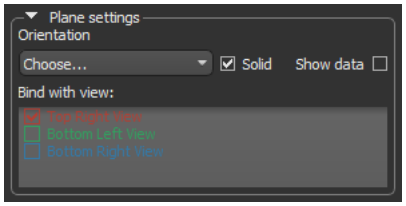

Lighting Controls for ROIs and Multi-ROIs

A number of additional lighting controls – **Diffuse** and **Specular** – are now available for fine-tuning the visualization of regions of interest and multi-ROIs when they are shown without a dataset.

Lighting applied to an ROI without a dataset

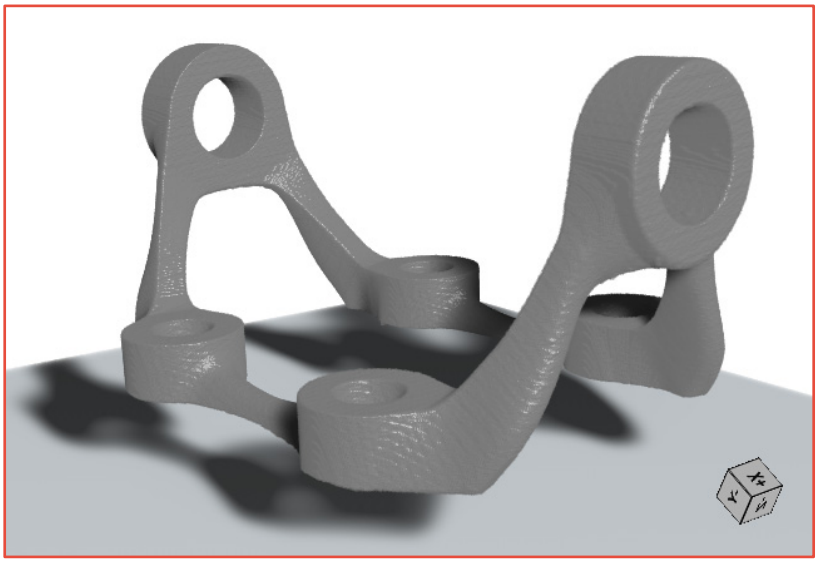

These controls are available on the Data Properties and Settings panel in the Opacity and color box, as shown below.

Opacity and color box

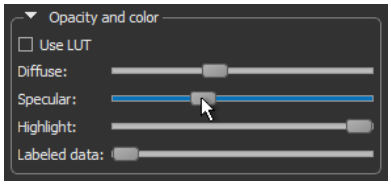

Diffuse slider… Lets you vary the strength of diffuse light reflected from surfaces, which is a characteristic of light absorbed by surfaces. Increasing the diffuse lighting with the Diffuse slider will result in surfaces being more defined.

Specular slider… Lets you vary the strength of specular reflections originating from the surfaces of an object and is a characteristic of light reflected from shiny surfaces. Increasing this property with the Specular slider will result in increased reflections.

NOTE These new options can be better than converting an ROI or multi-ROI to a mesh for visualization purposes.

Edge Contrast Factor

In cases in which you want to highlight a segmented feature of interest within an object while still showing the object's structure, a new slider lets you precisely control the edge contrast of images.

Image with edge contrast applied

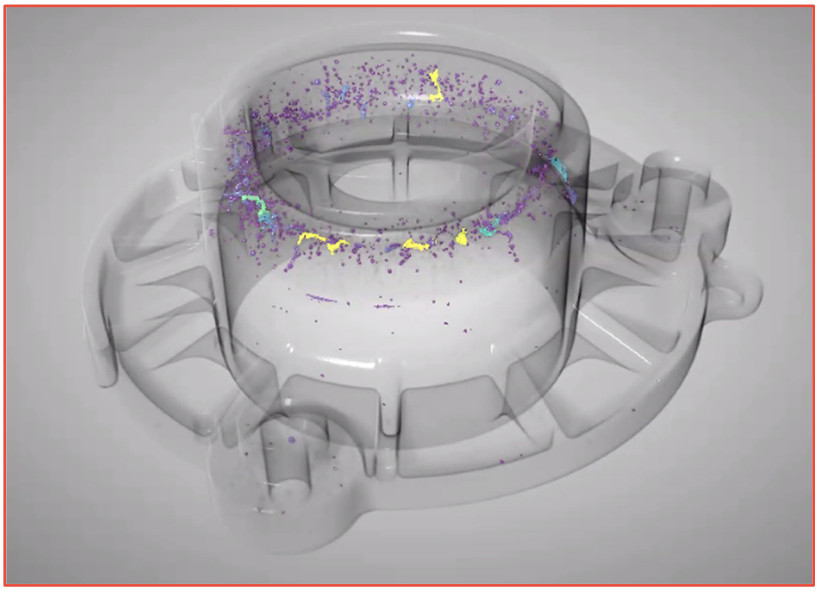

Edge Contrast slider

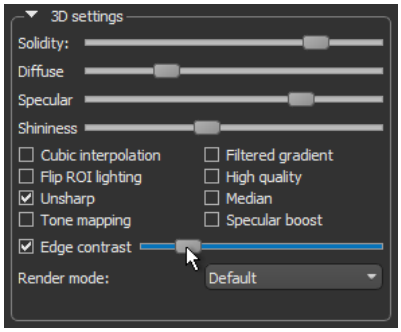

NOTE This control, which can be used in animated sequences, was previously available on the Window leveling panel.

Interpolation

There is now a drop-down menu for the 3D filtering, with options for Nearest, Linear, and Cubic interpolation. These Interpolation options replace the deprecated Cubic checkbox. The addition of the Nearest filter mode in 3D lets you examine datasets with no interpolation. This can particularly useful when viewing datasets exported from PuMA, so that the physical values conveyed are not altered by interpolation.

The additional interpolation options are available in the 3D settings box, as shown below.

3D settings

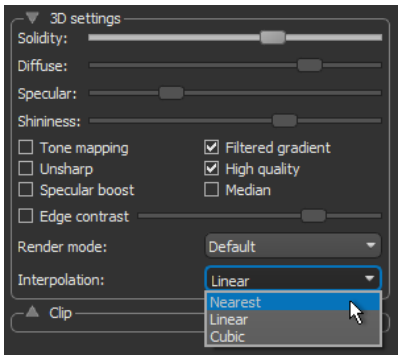

Render Mode

You should note that the Render Mode settings, which can be applied to images in 3D views, has been moved to the Data Properties and Settings panel on the right sidebar from the Scene Views Properties panel on the left sidebar.

3D settings

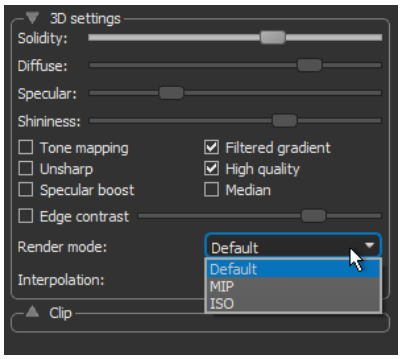

With this change, each dataset in the 3D view can be rendering in a different mode — Default, MIP, or ISO.

Virtual Floor Size Adjustment

The Virtual Floor panel now includes the option to adjust the size of a virtual floor.

Virtual Floor panel

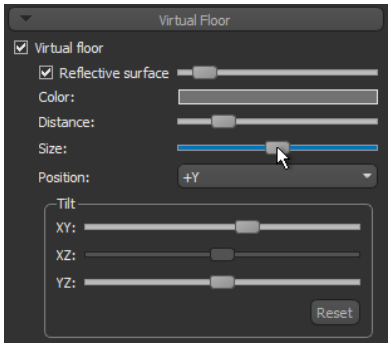

Movie Maker

The following enhancements, described below, were implemented for the Movie Maker in version 2022.2. You should also note that rebuilding movies is now much faster and more robust, particularly in cases in which objects were created and added to timeline after the Movie Maker was launched.

Save and Load Templates

This Dragonfly release includes the option to save and load Movie Maker timelines – either in the Movie File (*.mvf extension) format or in the generic Template File (*.tpl extension) format. In either case, only the camera manipulations, clipping, and other property changes will be saved in the file, which is more efficient that saving all objects and the timeline in a session file.

The options to save and load movie templates are available on the Movie Maker panel, as shown below.

Save Movie Timeline to file

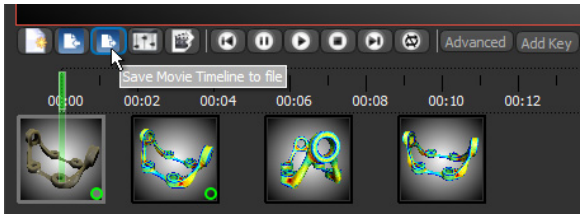

Load Movie Timeline from file

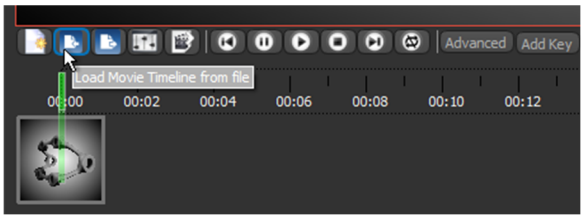

NOTE After loading a Movie Template, you can click the **Update Thumbnails** button to refresh the appearance of the thumbnails on the timeline.

Movie File format... This format is recommended in cases in which you want to save a particular animated sequence and then reproduce the sequence with the same or similarly sized objects of the same data type. You should note that the number of available objects and the name of each object on the Data Properties and Settings panel must be same to load a Movie File template.

Template File format... This format is recommended in cases in which you to save a 'generic' animated sequence that could be re-used for other objects. In this case, the size, position, and spacing of the objects in the animated sequence will be adapted to fit

the movie template stage and properties. You should note that the number of available objects on the Data Properties and Settings panel must be same to load an animated sequence saved as a Template File.

Visual Effects of Shapes

You can now interpolate the visual effects — such as clipping, window leveling, opacity, and others — that are applied between keyframes to shapes, such as boxes, capsules, and cylinders.

Clipping and visual effects applied to a keyframe in an animated sequence

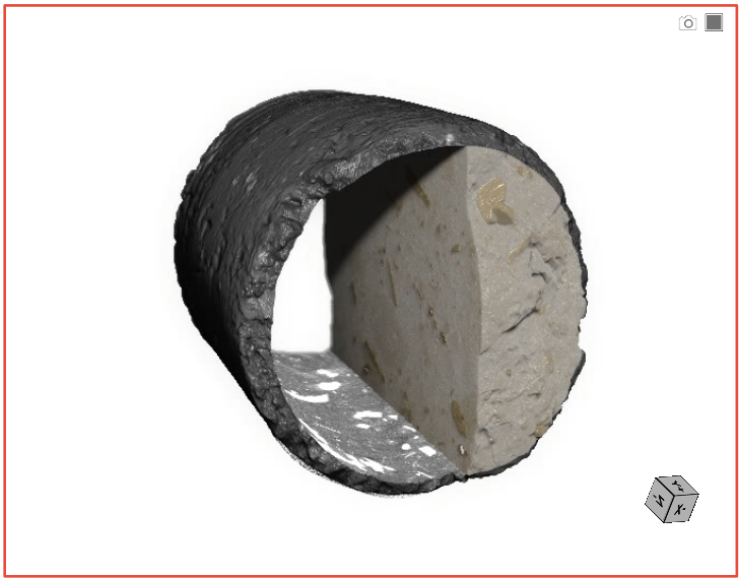

Graph Clipping

You should note that clipping of graphs is now supported in the Movie Maker, as shown in the timeline below.

Movie Maker timeline

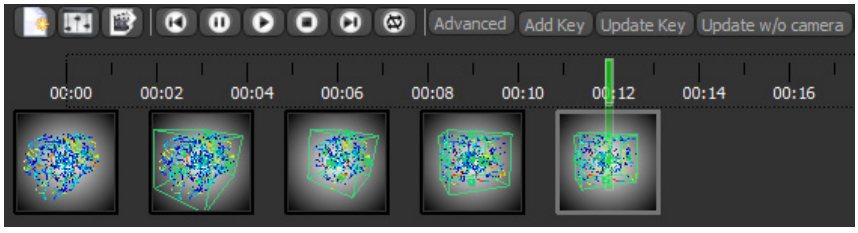

NOTE When clipping graphs, nodes are considered to be either in or out, but edges will be clipped by the boundary box.

Overlays

You can now change the position and opacity of text and image overlays between keyframes, also well as smooth transitions by applying interpolation. These options are available in the Advanced Animations Options dialog, as shown below.

Advanced Animations Options dialog

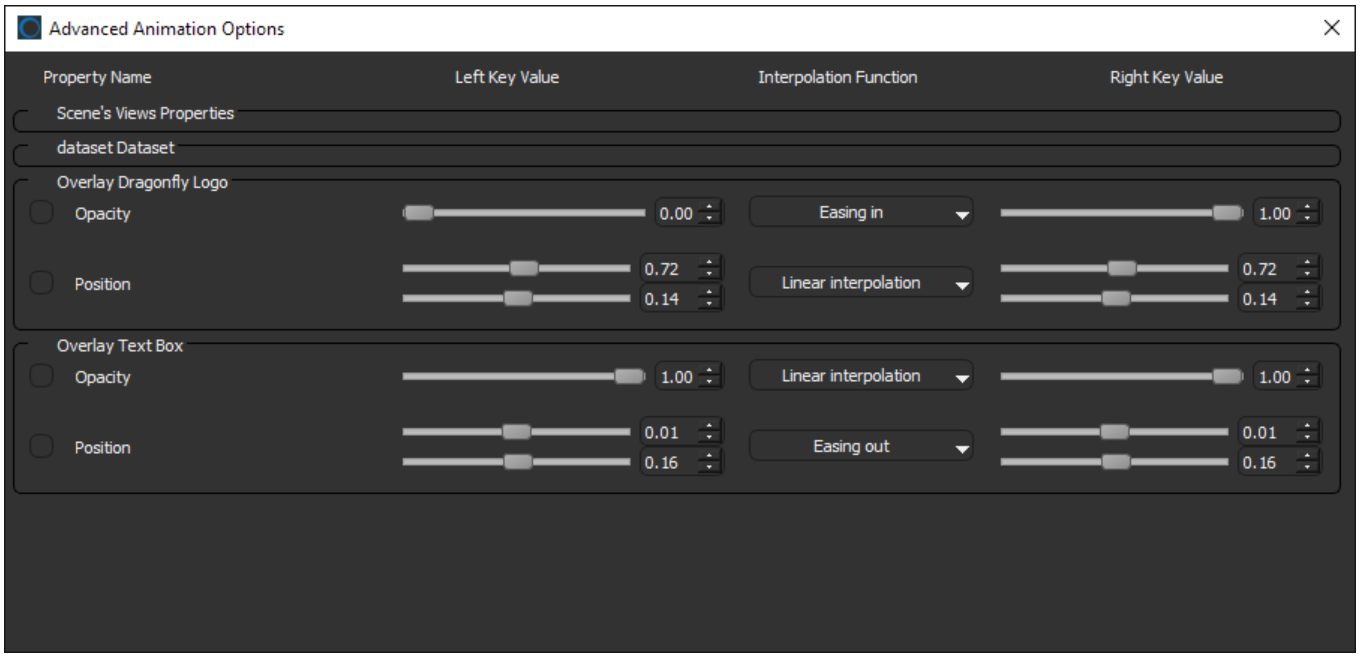

Window Leveling

For fine-tuning the tonal range of 3D renderings, a new **Gain and Bias** Window Leveling mode is available in version 2022.2.

3D renderings of Gila monster head

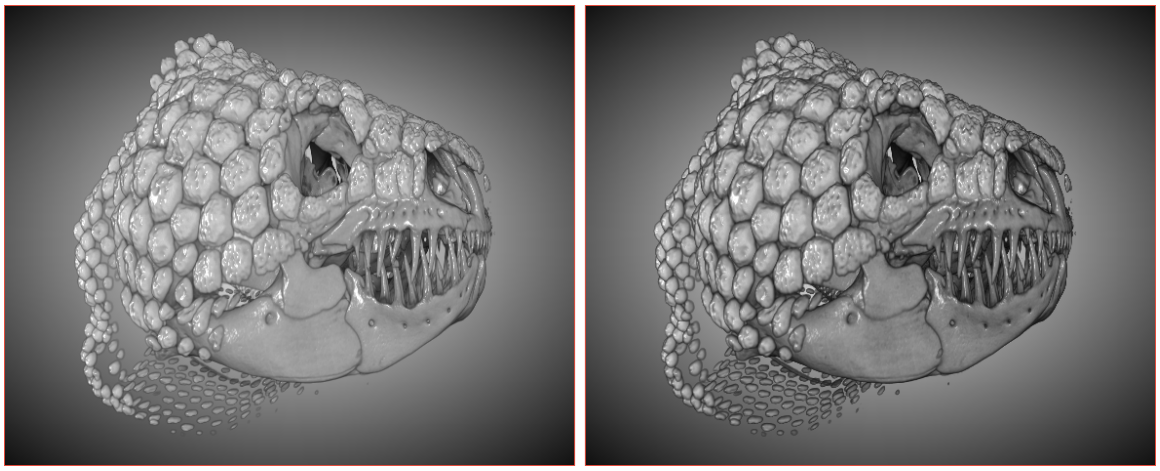

The Gain and Bias window leveling mode, as well as the Gain slider, is available on the Window Leveling panel, as shown below.

Window Leveling panel

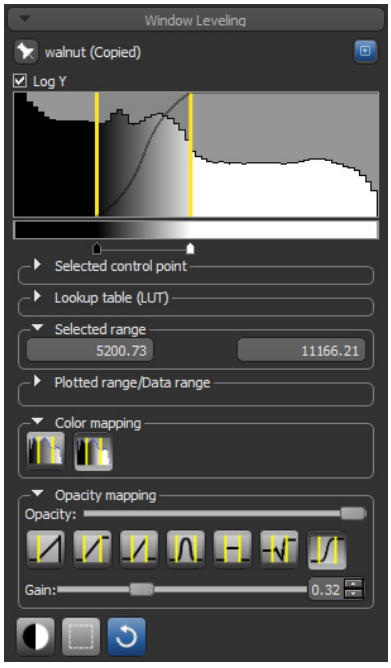

You should note that an image's tonality is initially represented as a straight diagonal line on the histogram. You can then adjust the shape of the curve with the Gain slider to adjust its tonal range.

Regions of Interest

A number of new options available in this software release for working with regions of interest. These new options are described below.

Adaptive Watershed

For cases in which you cannot get good segmentation results using a global threshold, this software release provides an adaptive watershed that separates single components in a manner consistent with the integrity of the connected particles.

Right-click the required ROI and then choose **Adaptive Watershed** in the pop-up menu to compute an adaptive watershed on an initial segmentation. The output is a multi-ROI with all particles indexed and identified.

Input ROI and related dataset (left) and the resulting multi-ROI (right)

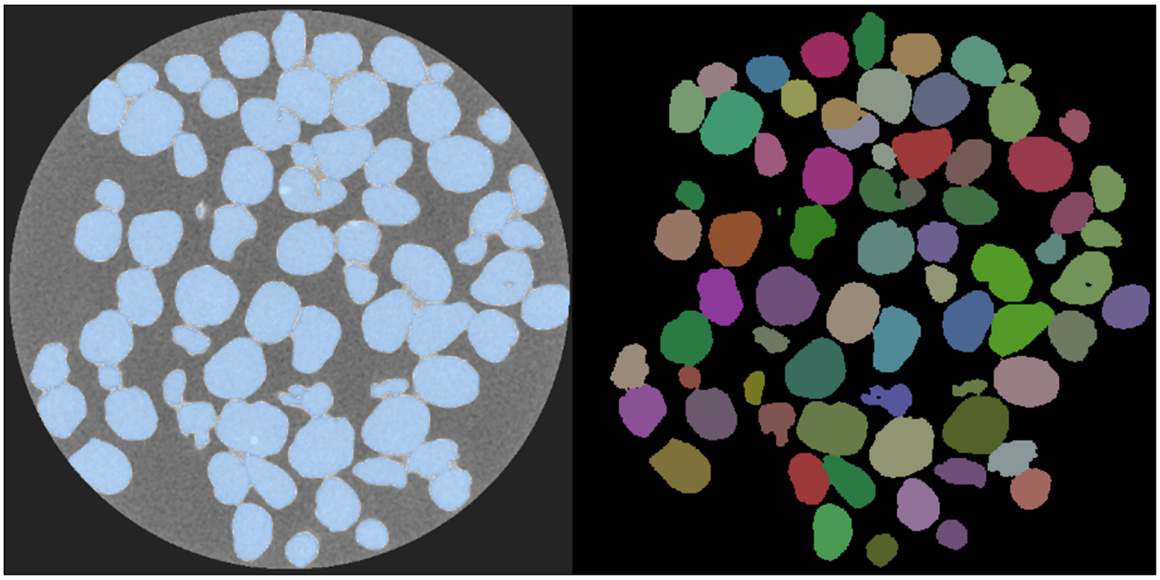

The adaptive watershed can segment a particle image into different regions by treating its inverse distance map as a landscape and the local minima as markers. By labeling each segmented region with a unique index, different particles can be separated, identified, and subsequently analyzed. However, you should note that the performance of the algorithm may be compromised if the particles are irregularly shaped, overlapped, or are overly connected.

NOTE Refer to Sun, H.Q., and Y.J. Luo, Adaptive watershed segmentation of binary particle image, Journal of microscopy 233.2(2009): 326 – 330 for more information about the implemented method. Available at: https://onlinelibrary.wiley.com/doi/10.1111/j.1365-2818.2009.03125.x.

Act on All Timesteps

In cases in which you create regions of interest on 4D image data, you can now choose to have applicable operations — such as clear, invert, and undo, as well as defined ranges, morphological operations, and image operations — applied to all timesteps automatically. The option **Act on all timesteps**, which is circled below, is available whenever you create a new ROI.

ROI Tools panel

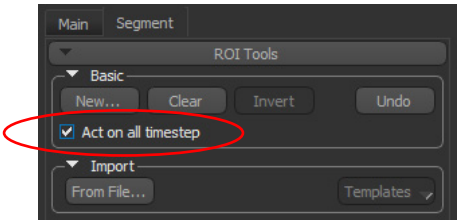

Selection Text Annotation

Whenever you define a range for a threshold segmentation, a new text annotation will appear in the 2D views of the current scene. This annotation, shown below, lets you adjust the range by dragging.

Selection text annotation

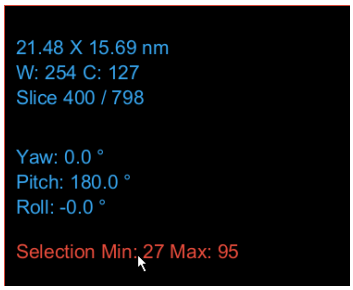

Do the following to adjust the defined range with the Selection text annotation:

- Drag left and right to decrease or increase the selected range (window width)
- Drag up and down to decrease or increase the center value.

New Binary Image (8 bit) from ROI

This item, which is available in the pop-up menu for regions of interest, lets you create a new binary image from an ROI in which all labeled voxels are assigned a selected value and all unlabeled voxels are assigned a value of 0. You can choose a value for the labeled voxels in the Value for Labeled Voxels dialog, as shown below.

Value for Labeled Voxels dialog

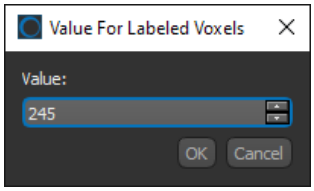

Multi-ROIs

A number of additional options for working with multi-ROIs are available in this software release. These new options are described below.

Gridded Multi-ROIs

In some cases, you may want to make a 'gridded' multi-ROI from a region of interest or another object. For example, to analyze a subsection of a large region of interest or multi-ROI, as shown below.

Gridded multi-ROI of a large crack

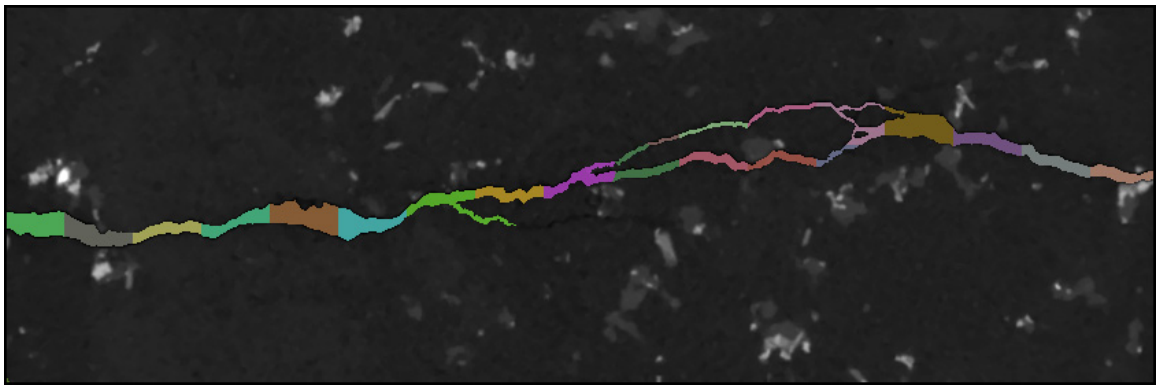

Gridded multi-ROIs can be created from image data, regions of interest, and multi-ROIs. Select the required option, as described below, in the pop-up menu to open the Grid Cell Size dialog, and then select the required sampling size for the generated grid.

Make as a Gridded Multi-ROI Intersection... Creates a gridded multi-ROI that includes only voxels that intersect with the selected region of interest.

Make as a Gridded Multi-ROI of Same Shape... Creates a gridded multi-ROI in the same shape as the selected image data, region of interest, or multi-ROI.

Grow Labels to Feret Box With Otsu

This software release includes the option to grow labels to their Feret box within the local lower or upper Otsu of a selected dataset. Right-click the required multi-ROI and then choose **Grow Labels to Feret Box with Otsu** in the pop-up to open the dialog shown below.

Grow Labels to Feret Box with Otsu dialog

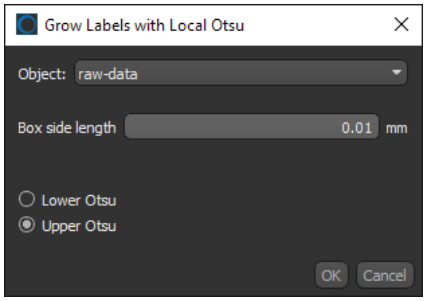

In the dialog, you can:

- Select the object to extract to the local lower or upper Otsu values.
- Grow the Feret box by a specified amount.
- Select local lower or upper Otsu.

NOTE All labels will be processed in order and voxels that intersect with the initial labels within the Feret box of each label will be added.

Meshes

A number of new options for generating meshes and using meshes as a mask, as well as optimizations for computing mesh deviation maps, are available in this software release for 3D modeling.

Generate Meshes from Spherical Harmonics

For simulation and experimentation purposes, this software release provides the option to generate meshes from spherical harmonics.

Choose **Utilities** > **Generate a Mesh from SH** on the menu bar to open a series of dialogs that let you enter the order and degree for the harmonic function. Meshes can be generated with a number of different settings, as shown below.

Meshes generated from spherical harmonics

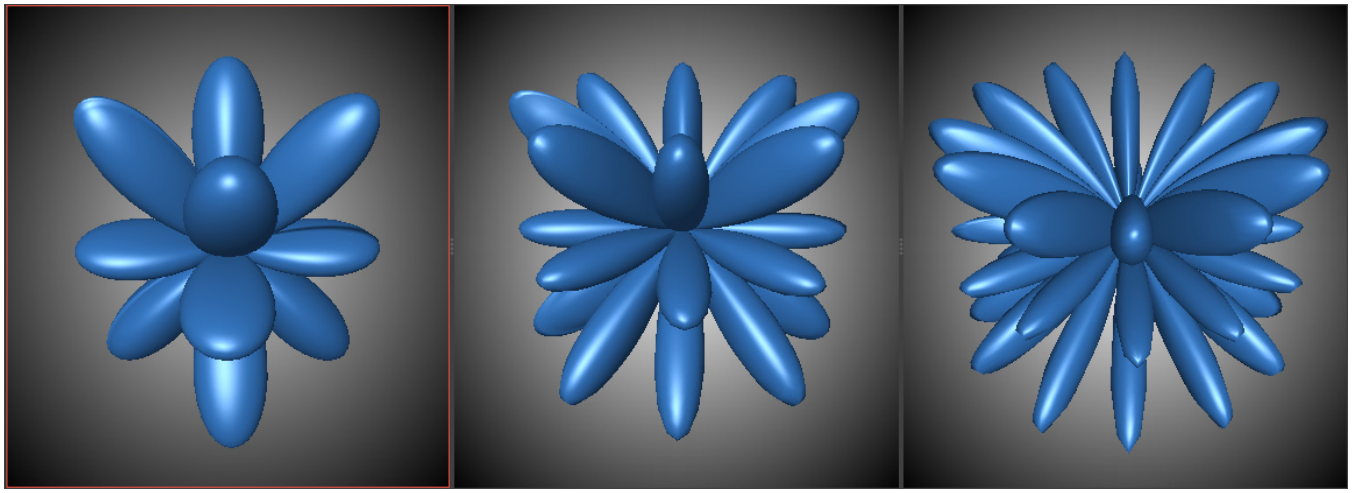

Settings for the above meshes were as follows: L=4, M=2 (left), L=6, M=4 (middle), L=8, M=6 (right).

Spherical harmonic functions are defined by an order (L), which specifies the number of polynomial terms that a harmonic function contains, and a degree (M), which describes the number of basis functions to be computed for each order. Relatively smooth functions can be obtained using low order spherical harmonics, with increasing order corresponding to a greater number of frequencies captured by the spherical harmonics.

Generate 3D Tetrahedral Meshes

This software release includes the option to generate 3D tetrahedral meshes from selected classes of a multi-ROI. 3D tetrahedral meshes, which are boundary meshes filled with tetrahedrals that can fit arbitrary shaped geometries, are often used for computational fluid dynamics and other simulations.

Clipped 3D tetrahedral mesh (orange) inside a clipped surface mesh (blue)

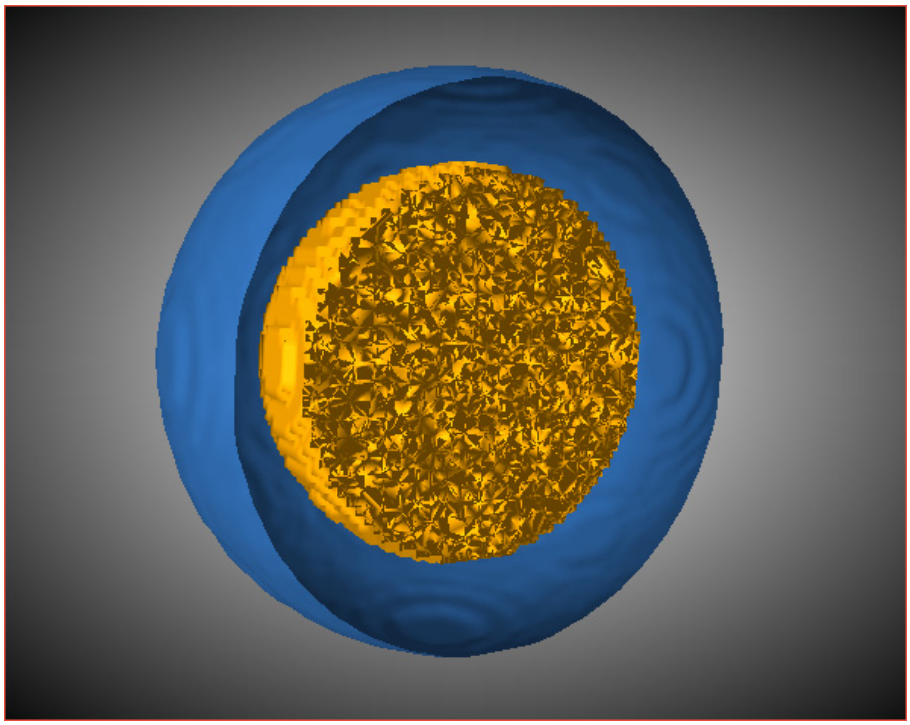

Do the following to generate a 3D tetrahedral mesh:

- Select the classes in a multi-ROI for which you want to generate a 3D tetrahedral mesh.
- Right-click the multi-ROI and then choose **Generate 3D Tetrahedral Mesh** in the pop-up menu.
- Enter the required values in the Generate 3D Tetrahedral Mesh dialog, shown below.

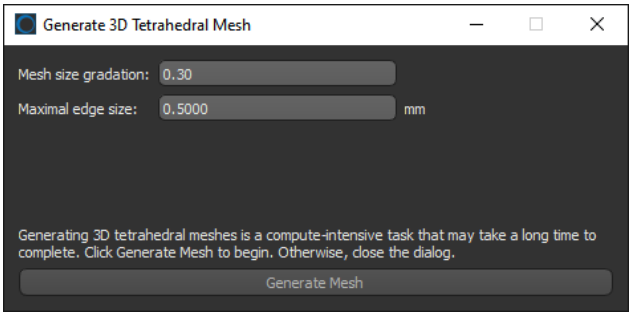

NOTE Tetrahedral elements are commonly constructed as equilateral, such as in systems with circular curvature, or they could be constructed as isosceles tetrahedra when a system has asymmetry. Tetrahedral elements could also be totally unstructured and adapted to arbitrary geometries with high accuracy. The accuracy can be much higher than a simple cubic grid applied in 3D or a square grid applied in 2D.

Masking with Meshes

This software release provides the option to use meshes as a mask. For example, to mask a dataset, as shown below.

3D rendering with mesh (left) and 3D rendering masked with mesh (right)

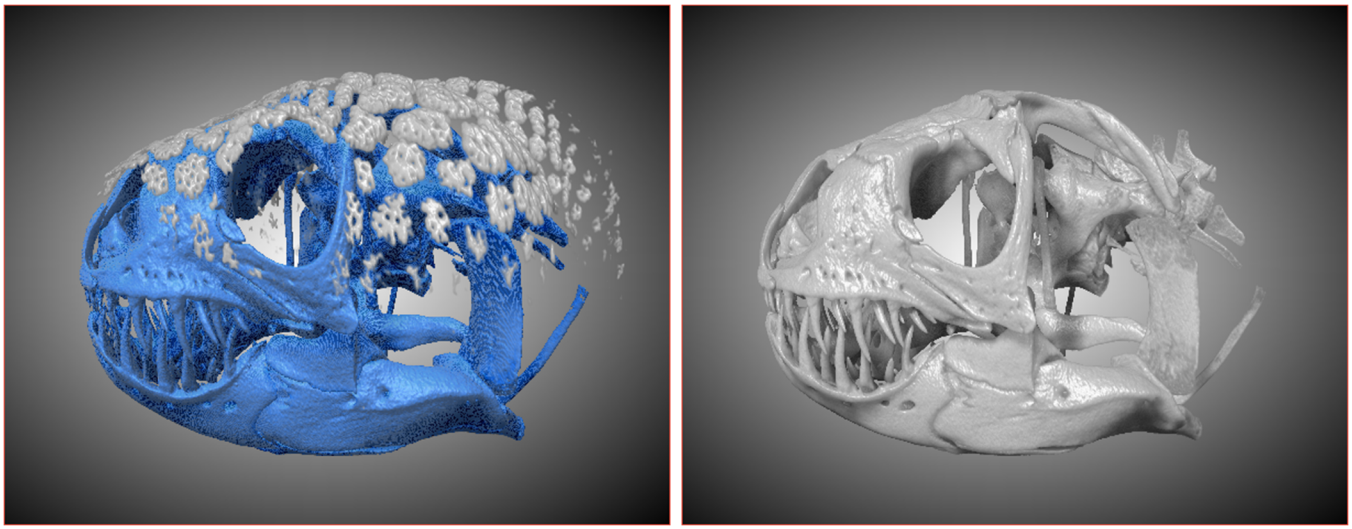

The option to use a mesh as a mask is available in the 3D settings box, as shown below.

3D settings for meshes

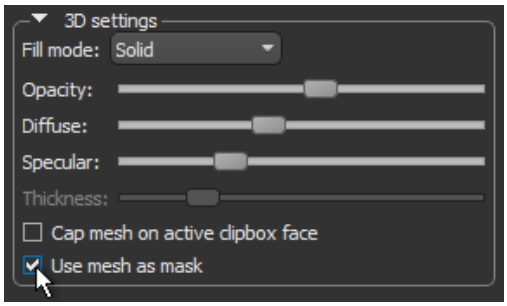

Distance Between Meshes

For optimizing computations of true distances between meshes, this software release provides a Search Distance dialog for specifying the maximum search distance when computing distance maps.

Signed distance map

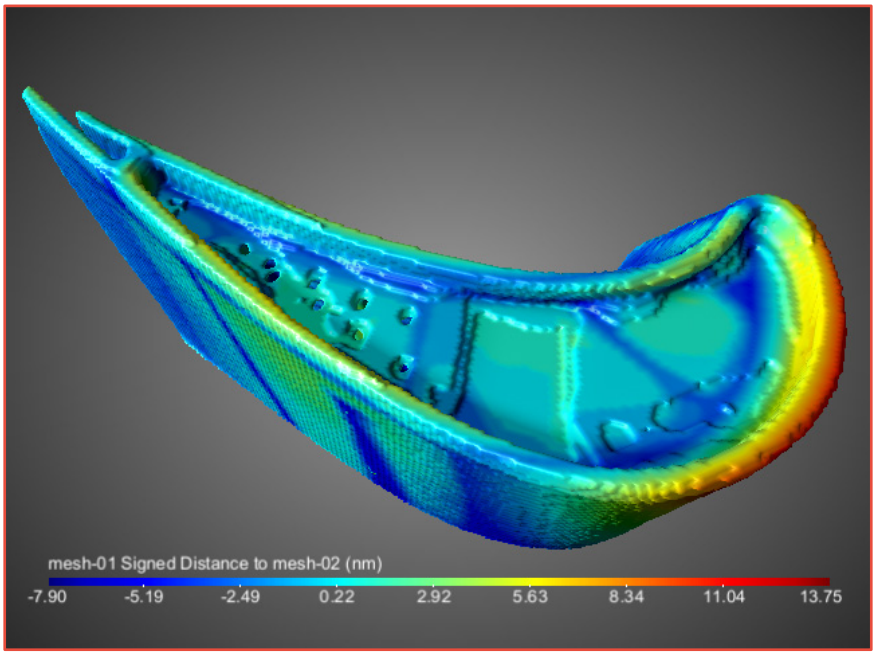

Right-click the required mesh and then choose **Signed Distance Map** or **Deviation Map** in the pop-up menu. You can then choose a reference mesh in the Choose a Mesh as a Reference dialog and a search distance in the Search Distance dialog, shown below.

Search Distance dialog

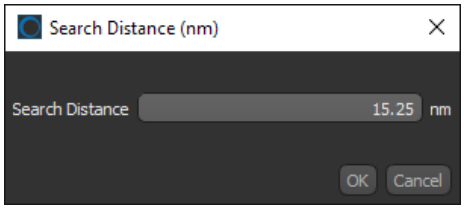

Vector Fields

A new setting for vector fields — **Normalized modulus** — is available in this software release for visualizing vector fields. If applied, the length of each directed line segment will be equal, regardless its magnitude.

Close-up of vector without normalized modulus (left) and with normalized modulus (right)

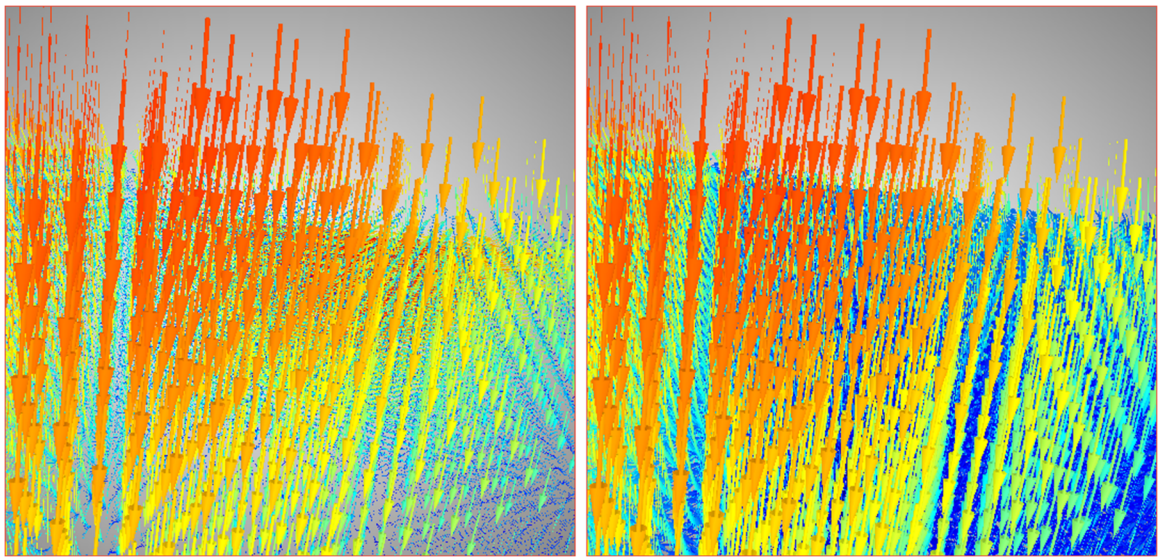

The option to apply the Normalized modulus option is available in the Vector field settings box, as shown below.

Vector field settings

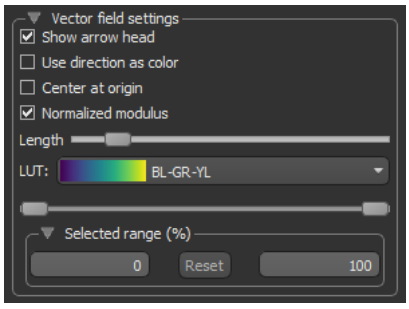

Floating Captions for Annotations

Captions for rulers, paths, angles, and regions can now be floated and then repositioned manually anywhere within a view to avoid overlapping. The option to float a caption is available in the Properties box, as circled below.

Properties for annotations

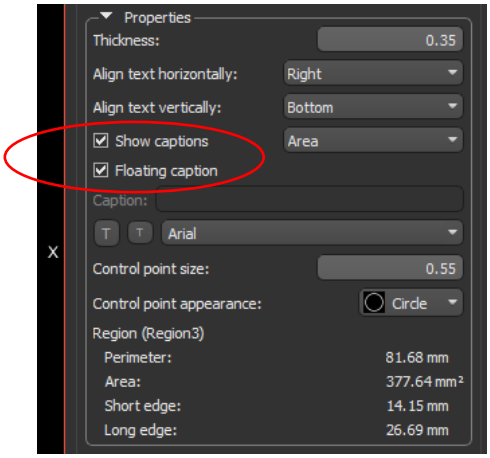

Below is an example of using floating captions for a series of rulers.

Floating captions

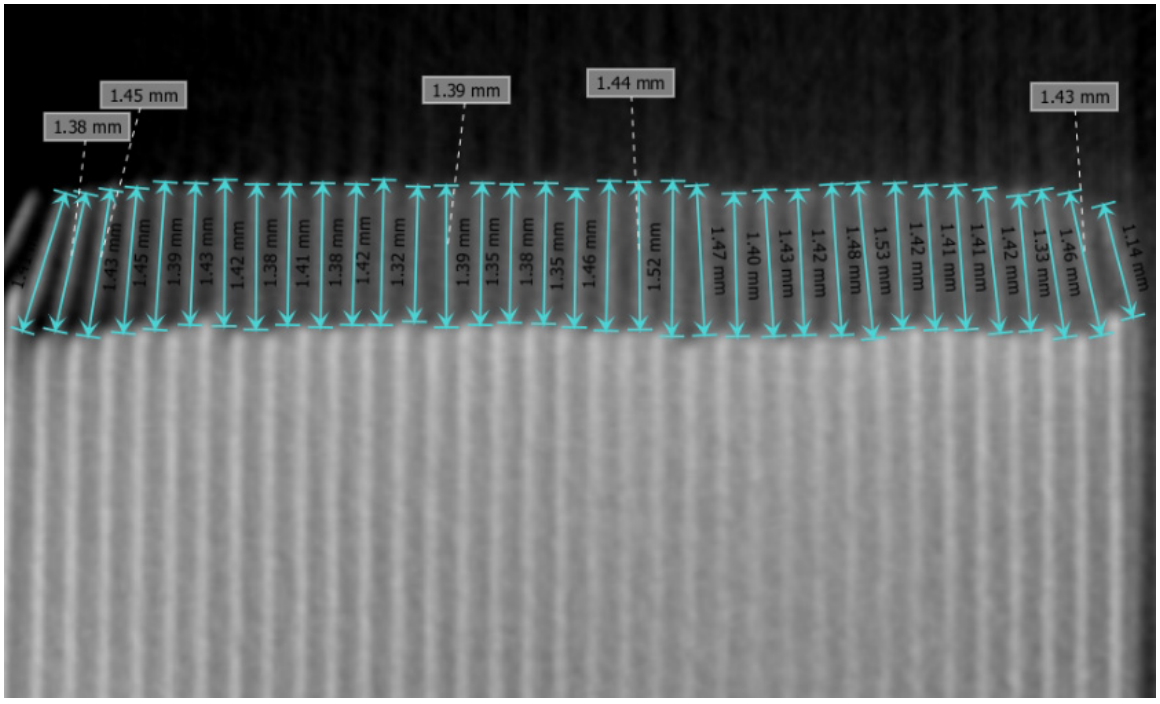

NOTE Selected annotation properties are now recorded by Dragonfly's Macro Player, which can include position, font size, control point appearance, and others. However, the setting 'Automatically update captions' for Points set is not recorded.

Menu Bar Changes

The following new items are available on the menu bar in Dragonfly 2022.2. Other changes are also listed here.

Workflows Menu

A number of changes, listed below, were implemented in the **Workflows** menu.

Workflows menu changes

Measurement	Description
RAD-DENDRO	Opens the Radiographic Dendrochronology Wizard, which provides a guided workflow for analyzing wood samples or planks (see Radiographic Dendrochronology on page 50).

Utilities Menu

A number of changes, listed below, were implemented in the **Utilities** menu.

Utilities menu changes

Measurement	Description
Feature Analysis (alpha)	Opens the Feature Analysis dialog, in which you can perform cross-table analysis of feature vectors and other measurements (see Feature Analysis on page 7).
Generate a Mesh from SH	Lets you generate meshes from spherical harmonics for simulation and experimentation purposes (see Generate Meshes from Spherical Harmonics on page 72).

Changes to Contextual Menus

The following changes for the contextual menus available in the Data Properties and Settings panel are implemented in Dragonfly 2022.2.

Dataset Contextual Menu

A number of new items, listed below, are available in the **Dataset** pop-up menu.

Dataset pop-up menu changes

Item	Description
Derive New from Current View, with Isotropic Spacing	Lets you create a new 3D volume with isotropic spacing that is aligned to the selected viewing angle of the current 2D view.
Make as a Gridded Multi-ROI of Same Shape	Creates a gridded multi-ROI in the same shape as the selected dataset (see Gridded Multi-ROIs on page 70).
3D Data Augmentation	Lets you generate a series of warped datasets (see 3D Data Augmentation on page 31).
Open Image Quality Metrics	Lets you evaluate and compare the image quality of previews and other single-slice datasets with a number of different metrics (see Image Quality Metrics on page 25).

ROI Contextual Menu

A number of new items, listed below, are available in the **Region of Interest** pop-up menu.

ROI pop-up menu changes

Item	Description
Adaptive Watershed	Applies an adaptive watershed that separates single components in a manner consistent with the integrity of the connected particles (see Adaptive Watershed on page 68).
Derive New from Current View, with Isotropic Spacing	Lets you create a new 3D volume with isotropic spacing that is aligned to the selected viewing angle of the current 2D view.
Make as a Gridded Multi-ROI Intersection	Creates a gridded multi-ROI that includes only voxels that intersect with the selected region of interest (see Gridded Multi-ROIs on page 70).
Make as a Gridded Multi-ROI of Same Shape	Creates a gridded multi-ROI in the same shape as the selected region of interest (see Gridded Multi-ROIs on page 70).
New Binary Image (8 bit) from ROI	Lets you create a new binary from an ROI in which all labeled voxels are assigned a selected value and all unlabeled voxels are assigned a value of 0 (see New Binary Image (8 bit) from ROI on page 70).
3D Data Augmentation	Lets you generate a series of warped datasets (see 3D Data Augmentation on page 31).

Multi-ROI Contextual Menu

A number of new items, listed below, are available in the **Multi-ROI** pop-up menu.

Multi-ROI pop-up menu changes

Item	Description
Derive New from Current View, with Isotropic Spacing	Lets you create a new 3D volume with isotropic spacing that is aligned to the selected viewing angle of the current 2D view.
Fill Multi-ROI with New Class	Automatically adds a new complimentary class to the multi-ROI in which all previously unlabeled voxels are labeled.
Make as a Gridded Multi-ROI of Same Shape	Creates a gridded multi-ROI in the same shape as the selected region of interest (see Gridded Multi-ROIs on page 70).
3D Tetrahedral Mesh Generator	Lets you generate 3D tetrahedral meshes from selected classes of a multi-ROI (see Generate 3D Tetrahedral Meshes on page 72).
3D Data Augmentation	Lets you generate a series of warped datasets (see 3D Data Augmentation on page 31).
Grow Labels to Feret box with Otsu	Lets you grow labels to their Feret box within the local lower or upper Otsu of a selected dataset (see Grow Labels to Feret Box With Otsu on page 71).

Preferences Changes

A number of new items, listed below, and other changes have been implemented for setting the application preferences.

Click **Preferences** on the Status bar or choose **File** > **Preferences** on the Menu bar to open the Preferences dialog.

Views Preferences

Changes for the **Views** preferences are indicated in the following table.

Views preferences changes

Item	Description
3D Settings	
Show 2D view plane(s) by default	If selected, 2D view planes will appear in the 3D view by default.

Miscellaneous Preferences

Changes for the **Miscellaneous** preferences are indicated in the following table.

Miscellaneous preferences changes

Item	Description
Default compression for object and session files	Lets you choose to either compress object and session files with zstd when saving or to not apply compression (see Lossless Compression with ZSTD on page 57).

Deep Learning Preferences

Changes for the **Deep Learning** preferences are indicated in the following table

Deep Learning preferences changes

Item	Description
GPU card for Deep Learning	If you have multiple GPUs available, you can select your preferred GPU(s) for training in the Deep Learning Tool and Segmentation Wizard, as shown below (see Deep Model Training on Multiple GPUs on page 14). 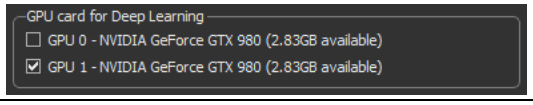

System Requirements

This release supports the base system requirements for Windows and Linux.

NOTE As of Dragonfly version 2022.2, Windows 10 or 11 is required for Deep Learning.

Third-Party Packages

You should note the changes for the third-party packages listed below.

Third-party packages

Item	Version 2022.2	Previous version (2022.1)
TensorFlow	2.4.1	2.4.1
Qt	5.12.9	5.12.9
Python	3.8.12	3.8.6

Publicly Available Source Code

Users can freely download and examine the source code of the following.

TensorFlow

A number of Dragonfly's artificial intelligence applications, such as the Deep Learning Tool, the Segmentation Wizard, and Segment with AI, use TensorFlow APIs to facilitate training and inference of deep neural networks. Created by the Google Brain team, TensorFlow is a free and open-source library widely used for deep learning applications and traditional machine learning.

Users can freely download and examine the TensorFlow source code from: https://github.com/tensorflow/tensorflow/releases/tag/v2.4.1.

Qt

Dragonfly is designed and developed with Qt, which is a free and open-source toolkit for creating graphical user interfaces and cross-platform applications. Qt is developed by The Qt Company, a publicly listed company, and is available under both commercial licenses and open-source GPL 2.0, GPL 3.0, and LGPL 3.0 licenses.

Users can freely download and examine the Qt source code from: https://download.qt.io/official_releases/qt/5.12/5.12.9/.

Porous Microstructure Analysis (PuMA)

The Porous Microstructure Analysis (PuMA) software integrated in Dragonfly was developed by researchers at NASA to compute effective material properties and perform material response simulations on digitized microstructures of porous media. PuMA is able to generate artificial microstructures that mimic real materials and to compute simple morphological properties such as porosity, volume fractions, pore diameter, and specific surface area. Additional capabilities include the determination of effective thermal and electrical conductivity (both radiative and solid conduction - including the ability to simulate local anisotropy for the latter); effective diffusivity and tortuosity from the continuum to the rarefied regime; techniques to determine the local material orientation, as well as the mechanical properties (elasticity coefficient), and the permeability of a material.

PuMA is available under the terms of the NASA Open Source Agreement (NOSA) and is a copyright of the United States Government as represented by the Administrator of the National Aeronautics and Space Administration.

Stable releases of PuMA can be found on the official NASA GitHub repository at: https://github.com/nasa/puma.

The latest development code can be found on the development GitLab repository at: https://gitlab.com/jcfergus/puma-dev.

Extending Dragonfly

Dragonfly's Infinite Toolbox, an open exchange platform for the Dragonfly user community, lets you submit, browse, and download extensions that implement new features and workflows for all your image processing or analysis needs.

Choose **Utilities** > **Infinite Toolbox** on the menu bar to open Dragonfly's Infinite Toolbox.

Requesting New Features

Let us know if you have a suggestion for a new feature, or an idea for an improvement to an existing workflow.

Learning Dragonfly

You'll find some great free resources, such as our latest instructional videos, recorded webinars, and user forums, at https://www.theobjects.com/dragonfly/. You can also take advantage of the course offerings listed below to further benefit from our training team's in-depth product knowledge.

Dragonfly Workshops

Our hands-on training workshops are a great way to get up and running quickly. Combining in-depth instruction with class exercises, our product specialists will quickly teach you how to use key Dragonfly features and functions to optimize your image visualization and analysis tasks.

You can check for updates at https://www.theobjects.com/dragonfly/workshops.html.

Online Training

Get connected with an online training session and follow along with a Dragonfly product specialist as they demonstrate the features and benefits of Dragonfly and walk you through specialized modules. Web-based training, which can be adapted to individual needs and interests, offers an interactive approach that lets you practice new techniques on numerous case studies.

Online training is available upon request. Contact support@theobjects.com for information about availability.

Reporting Issues

Dragonfly includes an integrated reporting module that lets you include comments and other information along with generated DMP and log files whenever you encounter an issue that results in a crash or application freeze.

You should note that you can also report an issue at any time by choosing **Help > Report an Issue** on the menu bar. For example, when an unknown error occurs or an unexpected result is generated.

Report an Issue dialog