dragonfly

Release Notes

Version 2024.1

Dragonfly 3D World

Learn all about the new features, product enhancements, and other improvements implemented in the **Dragonfly 3D World** and **Dragonfly Pro** 2024.1 software release.

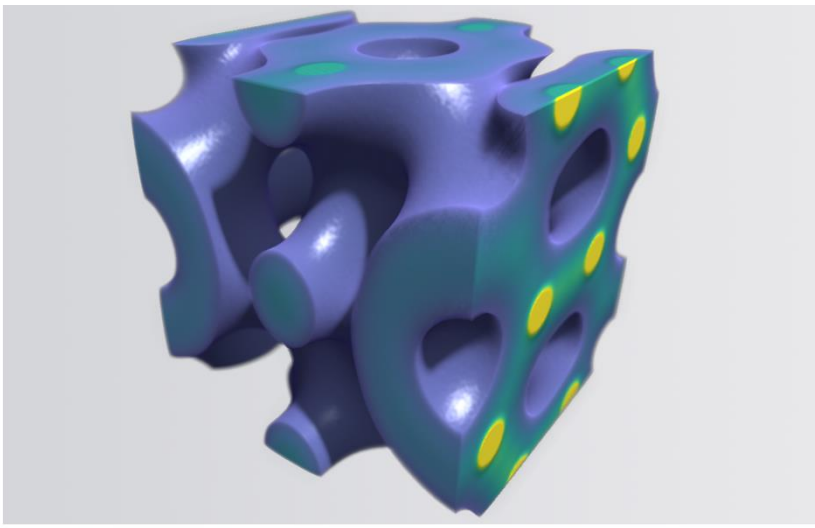

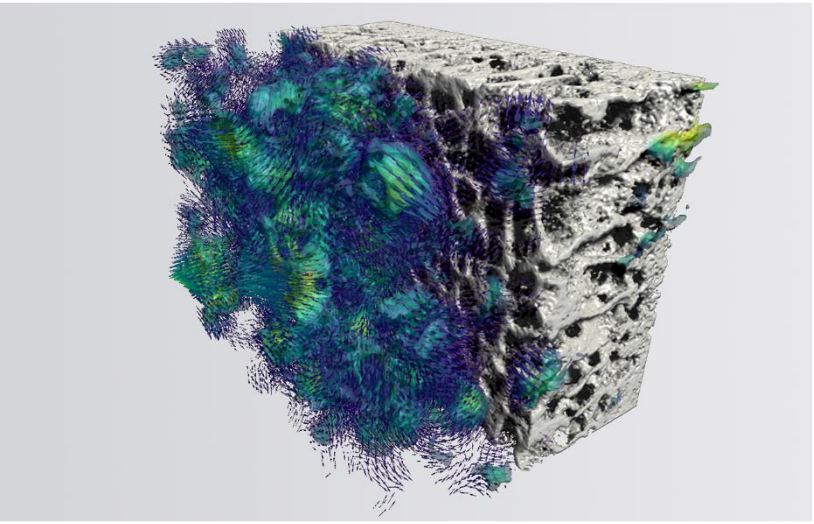

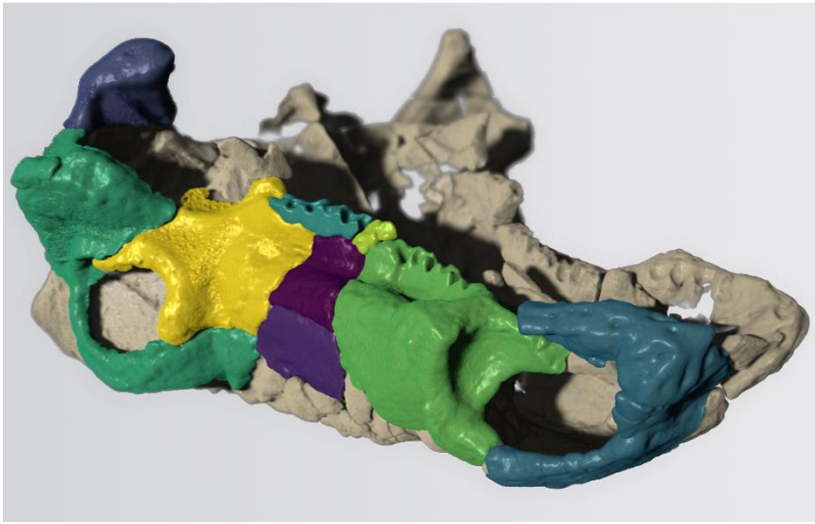

Copyright

© 2024 Comet Technologies Canada Inc. All rights reserved.

The present end-user documentation is confidential and proprietary information of Comet Technologies Canada Inc. (the "Company"). Only licensees of the Company have a right to use the information contained herein. Only licensees have the right to copy and/or transfer the information for internal use, unless otherwise agreed with the Company. Any unauthorized use, disclosure, transfer, or reproduction of this confidential information may give rise to a right in the Company to seek a legal remedy against such use, disclosure, transfer, or reproduction.

Except as expressly provided otherwise in writing, the information provided in this document is provided AS IS, without any condition or warranty, whether written, oral, implied, legal, or statutory. The Company makes no warranty as to its accuracy. Any use of the documentation or the information contained herein is at the risk of the user. Documentation may include technical or other inaccuracies or typographical errors. Information is subject to change without notice.

Trademarks

Dragonfly and the Dragonfly logos are trademarks of Comet Technologies Canada Inc.

Third-Party Trademarks

Python is a trademark or registered trademark of the Python Software Foundation. Adobe, Acrobat, Flash, and Reader are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries. Intel, Pentium, and Pentium 4 are registered trademarks or trademarks of Intel Corporation or its subsidiaries in the United States and other countries. NVIDIA and GeForce are registered trademarks or trademarks of NVIDIA Corporation in the United States and/or other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Advanced Micro Devices, AMD, and ATI are either registered trademarks or trademarks of Advanced Micro Devices Incorporated in the United States and/or other countries. VMware is a registered trademark and VMmark is a trademark of VMware, Inc. VMware VMmark is a product of VMware, an EMC Company. VMmark utilizes SPECjbb©2005 and SPECweb®2005, which are available from the Standard Performance Evaluation Corporation (SPEC). All other brand names, product names, or trademarks belong to their respective holders and should be noted as such.

Notices

Portions of Dragonfly's 3D engine licensed from the University of Münster.

Dragonfly is NOT a medical device.

Cover Images

Clockwise from the top left: 3D rendering of a Venus flytrap (Dionaea muscipula) scanned on a Comet Yxlon CT system, 3D rendering of a clipped boron dataset, segmented Ankylosaurus jawbone (courtesy of DigiMorph), vector field mapping of a 4D deformation study looking at the interface between an additive manufactured spacer and a human vertebra while undergoing continuous compression (CT data courtesy TESCAN/sample courtesy Rush University).

Comet Technologies Canada Inc.
460, rue Ste-Catherine Ouest, Suite 600
Montréal (Québec), Canada H3B 1A7

Dragonfly – a brand of Comet

dragonfly.comet.tech

Contents

New Branding 7

New Features and Product Enhancements 8

Dragonfly Social 9

Registering an Account 9

Connecting to Dragonfly Social 11

Segmentation Wizard 12

Sparse Labeling for Deep Model Training 12

Starting Sessions with a Trained Model 12

Revised Model Generation Strategies 14

Support for Multiple Datasets 15

Deep Learning Tool 16

Ready-To-Use Deep Model Library 17

Ready-To-Use Models 19

Updates for Segment with AI and Filter with AI 20

Additional Pre-Trained Models 21

Sparse Labeling 23

New Architectures 23

More Efficient Patch Blending for Accelerated Inference 24

Compile for Faster Inference 25

New Advanced Tools 26

Patch Inspection 27

Generate Augmented Deep Learning Training Sets 28

Automatic Loading/Unloading of Models 28

Automatic Saving of Models 29

Export Training Sets in a Single File 29

External Processes for Deep Learning Task Processing 30

Select Dataset for Visual Feedback 31

Validation Settings 31

Easy Sharing of Deep Models 32

NGSolve Integration 33

Open Fiber Segmentation 35

Automated Porosity Segmentation 37

Compute Measurements 39
Radiomics 40
Gray Level Co-occurrence Matrix (GLCM) 41
Gray Level Dependence Matrix (GLDM) 42
Gray Level Run Length Matrix (GLRLM) Features 42
Gray Level Size Zone Matrix (GLSZM) Features 43
Neighbouring Gray Tone Difference Matrix (NGTDM) Features 43
Longest Length 44
Phi and Theta Computations 45
Analyze and Classify Measurements 47
Histogram Analysis 50
Meshes 50
3D Mesh Painter 51
Mesh Smoothing with Constraint 52
Remeshing with Poisson Surface Reconstruction 53
Mesh Capping 54
Adjustable Thickness for Outlined Meshes 54
ROI Painter 55
ROI Painting Preferences 57
Thresholding 57
3D Rendering 58
Multiple Lights 58
Lighting and Shadows 60
Tone Mapping 61
Suppression of Moiré Effects 62
Material Presets 63
Settings for ROIs and Multi-ROIs 64
Planes 65
Exploded Views 65
Show Grayscale 66
3D View Controls 67
Automatic View Alignment 67
Constrain Pan Movements 67
Gyroscope Axis Indicators 68
New preferences for the Axis Indicator 69

CT Reconstruction 69
Organizer 70
Image Registration 70
Natural Language Search 71
Importing RAW Data 72
File Export 72
Image Filtering 73
Developer Tools 74
Tutorial Videos 75
Menu Bar Changes 76
Workflows Menu 76
Utilities Menu 76
Help Menu 77
Changes to Contextual Menus 78
Dataset Contextual Menu 78
ROI Contextual Menu 79
Multi-ROI Contextual Menu 80
Mesh Contextual Menu 80
Graph Contextual Menu 81
Vector Field Contextual Menu 81
Views Contextual Menu 81
Preferences Changes 82
Views Preferences 82
Deep Learning Preferences 82
System Requirements 83
Third-Party Packages 83
Publicly Available Source Code 84
TensorFlow 84
Qt 84
Porous Microstructure Analysis (PuMA) 84
Netgen/NGSolve 85

Learning Dragonfly **86**
Dragonfly Workshops 86
Online Training 86
Requesting New Features **86**
Reporting Issues **87**

New Branding

This software release features our new branding for the Dragonfly product line. It introduces the new name — **Dragonfly 3D World** — for this version of Dragonfly and includes updated logos and icons, as seen on the front cover of this publication and on the title bar of the application.

New Dragonfly 3D World icon

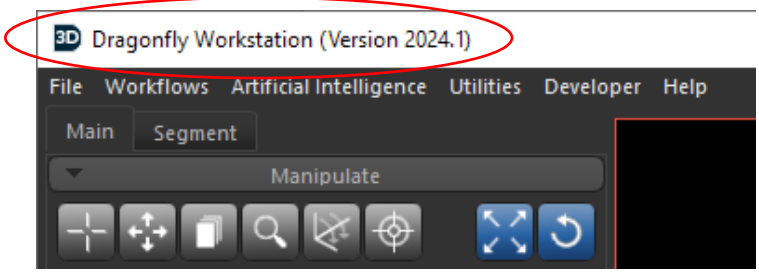

These changes are part of an extensive rebranding initiative to align Dragonfly as a brand of Comet. Comet is a globally leading, innovative technology company based in Switzerland with a focus on plasma control and X-ray technology. Comet's innovative solutions are in demand in the semiconductor and electronics market, the aerospace and automobile industry, and in security inspection. Headquartered in Flamatt, Switzerland, Comet has a presence in all world markets and employs more than 1,700 people worldwide, including about 600 in Switzerland. Besides production facilities in China, Denmark, Germany, Malaysia, Switzerland, and the USA, Comet maintains various other subsidiaries in Canada, China, Japan, Korea, Taiwan, and the USA. Comet (COTN) is listed on the SIX Swiss Exchange. To learn more about Comet, go to https://comet.tech/.

You can expect to hear more from us soon about additions to the Dragonfly product line, such as **Dragonfly 2D World** and **Dragonfly View**, as well as targeted applications that meet the needs of industrial users, such as **Void Inspect** and **Battery Insights**.

New Features and Product Enhancements

The software release for the newly branded **Dragonfly 3D World 2024.1** and **Dragonfly Pro 2024.1** includes new features and performance enhancements that are sure to boost your productivity. Just some of the highlights of this new release include:

- The introduction of **Dragonfly Social**, an online community platform that lets Dragonfly users start discussions, chat, access material, post announcements, as well as share extensions and deep models (see Dragonfly Social on page 9).
- Support for **sparsely-labeled ground truths**, which means that you can prepare training data for deep model training quicker (see Sparse Labeling for Deep Model Training on page 12).
- Access to an expanded remote library of **ready-to-use deep models** for image enhancement and segmentation (see Ready-To-Use Deep Model Library on page 17).
- **Accelerated inference** (see More Efficient Patch Blending for Accelerated Inference on page 24) and **advanced tools** for inspecting patches and generating augmented training sets for deep learning (see New Advanced Tools on page 26).
- Integration of **NGSolve** for quantifying fluid flow and heat dissipation in engineering and research applications (see NGSolve Integration on page 33).
- **Automated segmentation routines**, including extraction of individual fibers from datasets of reinforced composites (see Open Fiber Segmentation on page 35) and automated methods to reliably segment porosity in single-material samples, such as castings (see Automated Porosity Segmentation on page 37).
- **Advanced image analytics** to answer nuanced research questions and **radiomics** texture extraction to identify and analyze image features (see Compute Measurements on page 39 and Radiomics Features on page 40).
- The production version of Dragonfly 3D World's **Analyze and Classify Measurements** module, which lets you import scalar data from multi-ROIs, meshes, graphs, and vector fields for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
- Multiple lights, 'material' definitions, and improved tone mapping to make your **3D renderings** even more sensational (see 3D Rendering on page 58).
- **Extended accessibility**, with the option to install the application on a system/laptop without a dedicated GPU (see System Requirements on page 83).

Dragonfly Social

Dragonfly Social is a new online community platform that lets Dragonfly users start discussions, access material, post announcements, as well as upload and download extensions and deep models for the application. In addition, an integrated chatbot that lets users can ask basic questions about working with Dragonfly and learn more about best practices will soon be added to the platform.

Dragonfly Social

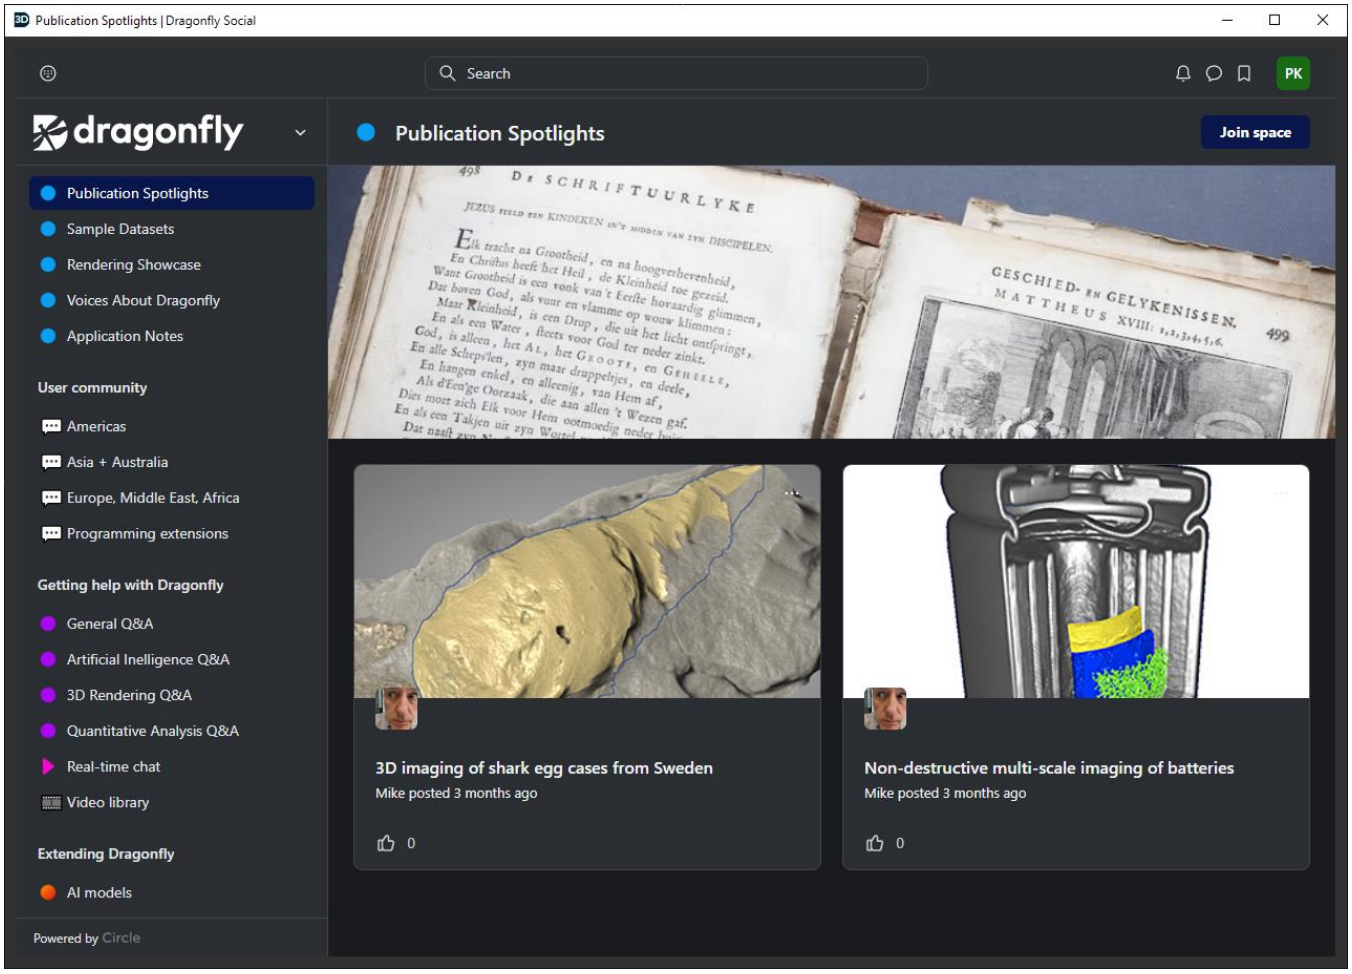

Registering an Account

You can register a new Dragonfly Social account as follows:

1. Do one of the following:
   - Enter your information in the Introducing Dragonfly Social dialog, as shown in the following screenshot, and then continue to Step 3.

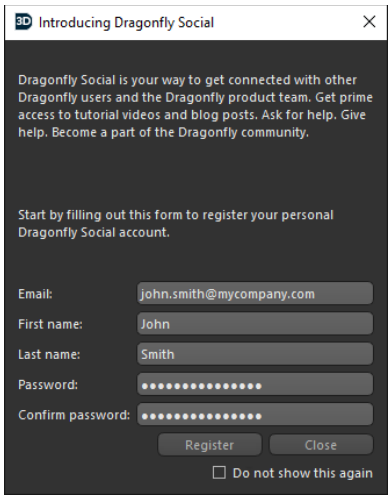

- Choose **Help** > **Manage Dragonfly Social Accounts** on the menu bar.

  The Manage Dragonfly Social Accounts dialog appears.

2. Do the following in the Manage Dragonfly Social Accounts dialog:
   - Click the **Register New Account** button, as shown below.

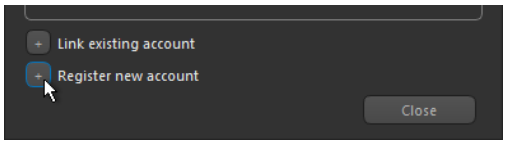

   - Enter your information in the Register New Account dialog.

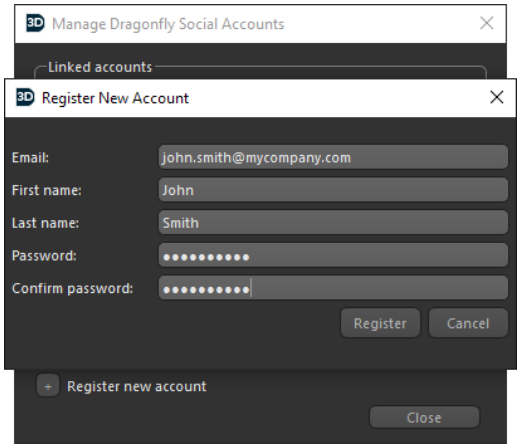

3. Click the **Register** button.

   The End User Licence Agreement dialog appears.

4. Review the terms and conditions and click **Agree** if you accept the terms and conditions.

   The Dragonfly Social Account Created Successfully dialog appears.

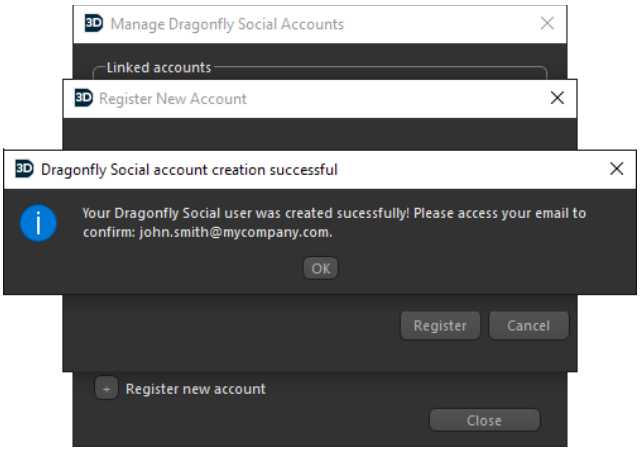

5. Click **OK** in the dialog.

6. Verify your email.

   If verification is successful, you will be able to login with the registered account.

Connecting to Dragonfly Social

After you have registered for an account, you can connect to Dragonfly Social by doing the following. You should note that you need to be connected to the Internet to access Dragonfly Social.

1. Do one of the following:

   - Click **Connect** in the Connect Now dialog, as shown below.

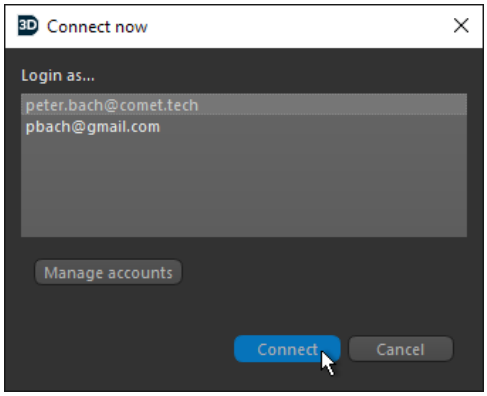

   - Choose **Help** > **Connect to Dragonfly Social As** > **Login As** on the Dragonfly menu bar to connect to the platform.
   - Go to https://dragonflysocial.comet.tech/ from any browser on a system connected to the Internet.

2. Sign into your account.

Segmentation Wizard

Sparse labeling for deep model training, starting sessions with a trained deep model, easy access to our remote library of ready-to-use segmentation models, and support for working with multiple datasets are all available for the Segmentation Wizard in this software release. Other deep learning workflow enhancements also included in the Segmentation Wizard are described in the topic Deep Learning Tool on page 16.

Sparse Labeling for Deep Model Training

One of the significant challenges of training models for semantic segmentation can be the need to densely label ground truth multi-ROIs. To address this issue and improve workflow efficiency, this software release reduces labeling requirements for deep model training by providing support for sparse labeling. In many cases, you should be able to achieve comparable results to those obtained by training with densely labeled ground truths, which can be much more laborious to obtain.

Starting Sessions with a Trained Model

In some cases, you may want to start a Segmentation Wizard session with a deep model that was already trained either by yourself or by the Dragonfly team. This option is now available in the Model Selection dialog, shown below, which appears whenever you launch the Segmentation Wizard.

Model Selection dialog

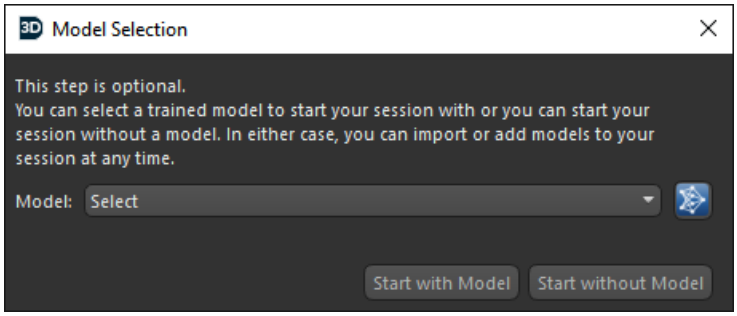

You can either select a trained model in the Model drop-down menu or click the **Remote Library** button to open the Remote Library of Ready-to-Use Segmentation Models dialog, which is shown on the following page.

Remote Library of Ready-to-Use Segmentation Models dialog

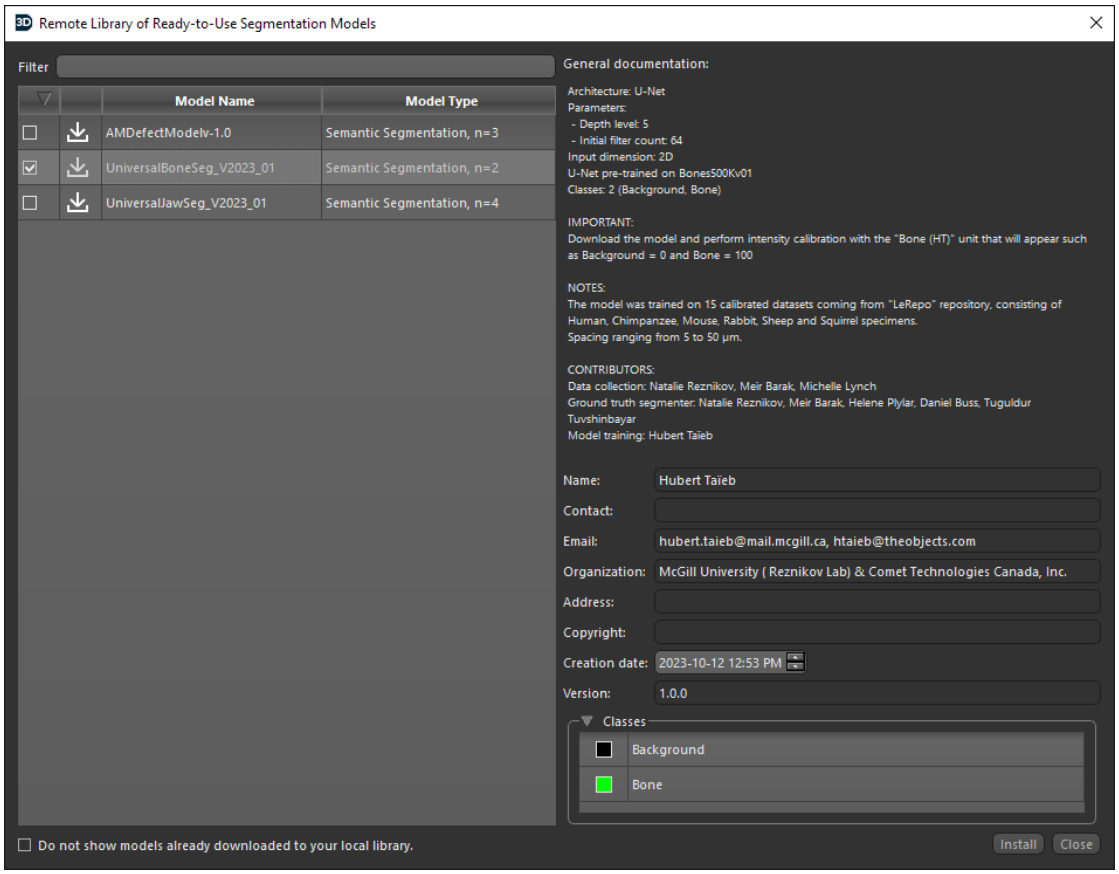

Downloaded ready-to-use models will appear in the Model drop-down menu the next time that you launch Dragonfly 3D World. Refer to the topic Ready-To-Use Deep Model Library on page 17 for additional information about ready-to-use-models.

NOTE You can also open the Remote Library of Ready-to-Use Segmentation Models dialog from the **Remote Library** button on the Models tab after you launch the Segmentation Wizard, as shown below.

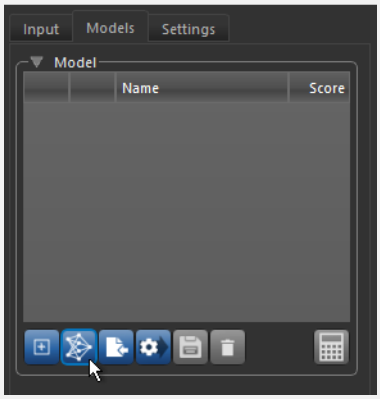

Revised Model Generation Strategies

This software release provides revised model generation strategies based on the availability of pre-trained models and sparse labeling for deep model training. You should note that the Model Generation Strategy dialog, which includes all default and user-defined strategies, appears whenever you click **Train** and the models list is empty. You can also click the **Open** button the Models tab to open the dialog, shown below.

Model Generation Strategy dialog

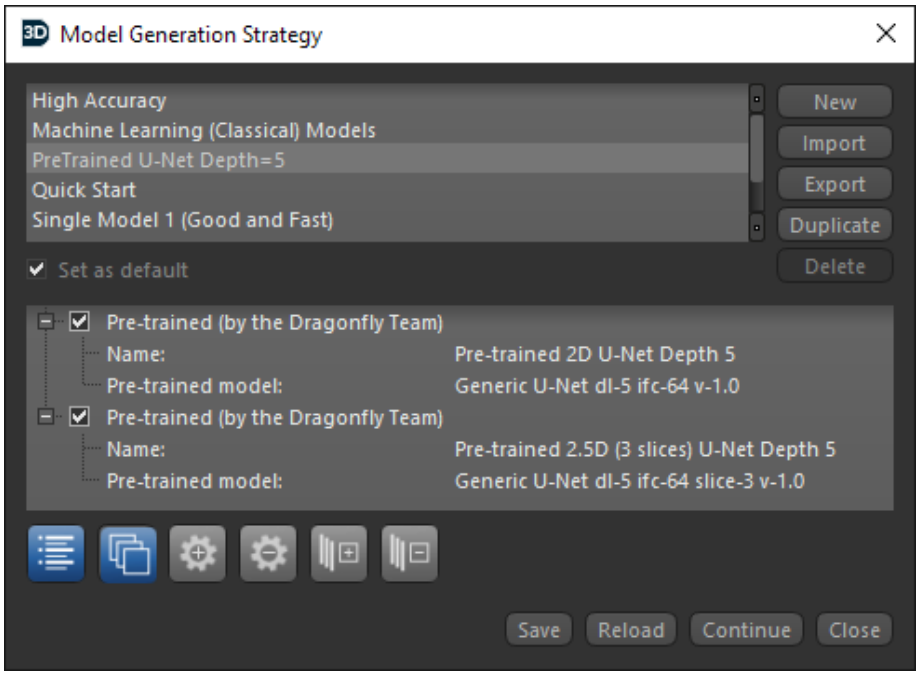

The following table provides a summary of the changes and additions to the default model generation strategies. Any strategy not listed below remains unchanged from previous versions.

Model generation strategies updates

Model	Description
Machine Learning (Classical) Models	Renamed, but still includes the two random forest models available in the original 'Sparse Training Data' strategy.
Pre-Trained U-Net Depth = 5	The models in this new strategy provide both good and fast results with the 2D U-Net pre-trained model, as well as better but slower results with the 2.5D U-Net pre-trained model. You should note that starting with a pre-trained model often provides better and faster results with smaller training sets than using an untrained model. The models in this new strategy include: • Pre-trained 2D U-Net with depth level 5 and initial filters count of 64. • Pre-trained 2.5D (3 slices) U-Net with depth level 5 and initial filters count of 64.

Support for Multiple Datasets

This software release includes support for working with multiple datasets in the Segmentation Wizard. For example, in cases in which you are working with 2D datasets it may not be possible to label an adequate number of pixels for training and validation on a single image. Imported datasets will appear in the Data sets box, as shown below.

Segmentation Wizard panel

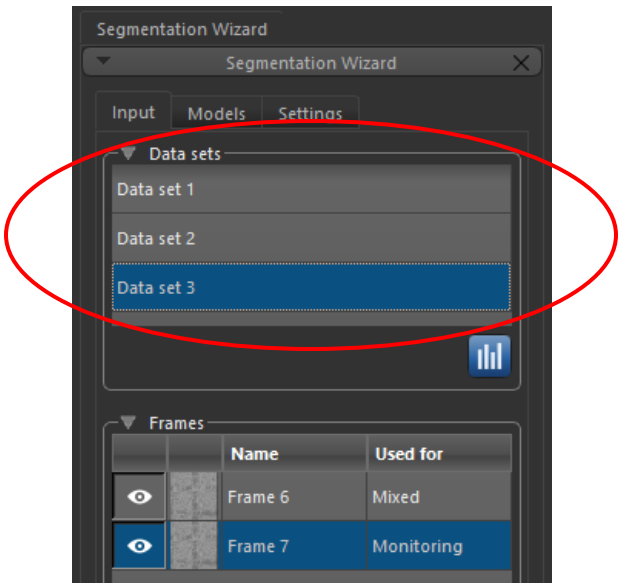

Do the following to load multiple datasets in the Segmentation Wizard:

1. Choose **Artificial Intelligence** > **Segmentation Wizard** on the menu bar.

   The Data Selection dialog appears.

2. Click the **Add Dataset** button to add the required dataset(s), as shown below.

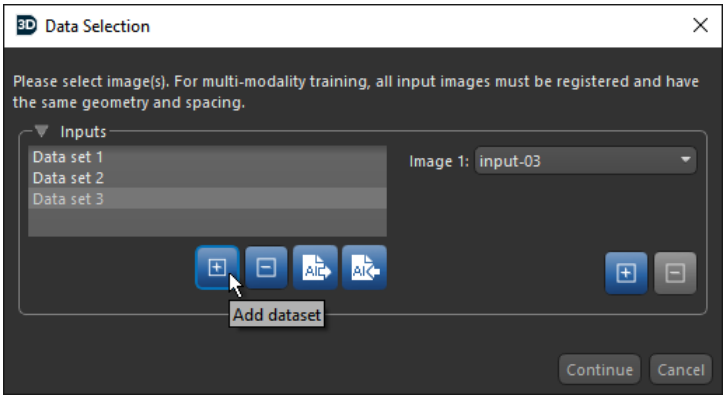

3. Choose an image or images for each input dataset.

Deep Learning Tool

This software release includes new features and innovations for Dragonfly's **Deep Learning Tool** that are sure to boost your productivity. These include support for sparsely labeled ground truths, a dedicated remote library of ready-to-use deep models for image enhancement and segmentation, accelerated inference, advanced tools for inspecting patches and generating augmented training sets, new architectures for training regression and semantic segmentation models, as well as other enhancements.

In addition to updating to **TensorFlow 2.11**, you should note the changes for the options on the Model panel. These changes as described below.

Model panel options

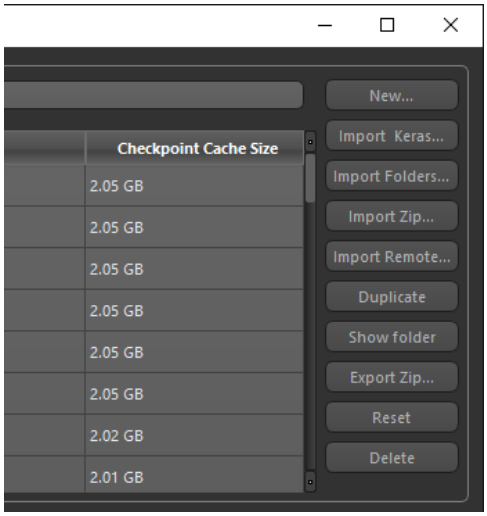

Import/Export Zip... Lets you export and import deep models for sharing with colleagues and the Dragonfly community (see Easy Sharing of Deep Models on page 32).

Import Remote... Lets you browse our library of ready-to-use models and download selected models to your local library (see Ready-To-Use Deep Model Library on page 17).

Reset... This button, which previously appeared at the bottom of the Deep Learning Tool dialog, lets you return models to their untrained state by randomizing their weights. If you do this, your model architecture will be preserved, but you will lose the 'learnt knowledge' of the model.

Load/Unload... These options were deprecated in this version of Dragonfly. In this release, deep models are automatically loaded and unloaded. Refer to Automatic Loading/Unloading of Models on page 29 for additional information about this change.

Ready-To-Use Deep Model Library

Ready-to-use deep models can greatly accelerate workflows for image enhancement and segmentation tasks. This software release includes a new utility that lets you easily browse our library of ready-to-use deep models and download selected models to your local library. Downloaded models can be used immediately or fine-tuned by performing additional training with your own datasets.

In the example below, a ready-to-use deep model for segmenting jawbone, dentin, and enamel was quickly fine-tuned with a small dataset of five slices and then applied to a dataset of a Lutra lutra mandible.

Result of applying the ready-to-use 'UniversalJawSeg' model with classes for enamel, dentin, and bone. Original microCT dataset acquired by Dr. Kornelius Kupczik (Universidad de Chile).

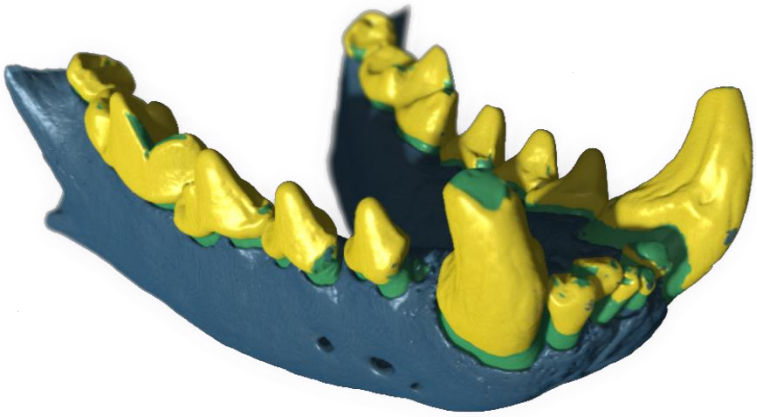

Click the **Import Remote** button in the Deep Learning Tool dialog, as shown below, to open the Remote Library of Ready-to-Use Regression and Segmentation Models dialog.

Deep Learning Tool dialog

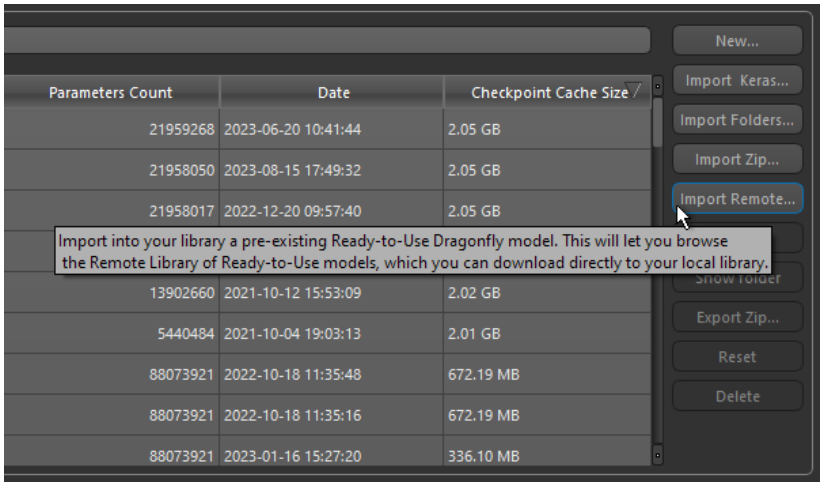

Models are listed as (available to download) and (already installed in local library). In addition, the general documentation provides information about each model's architecture, parameters, input dimensions, the number of classes for semantic segmentation models, as well as usage notes such as calibration requirements.

The Remote Library of Ready-to-Use Regression and Segmentation Models dialog is shown below.

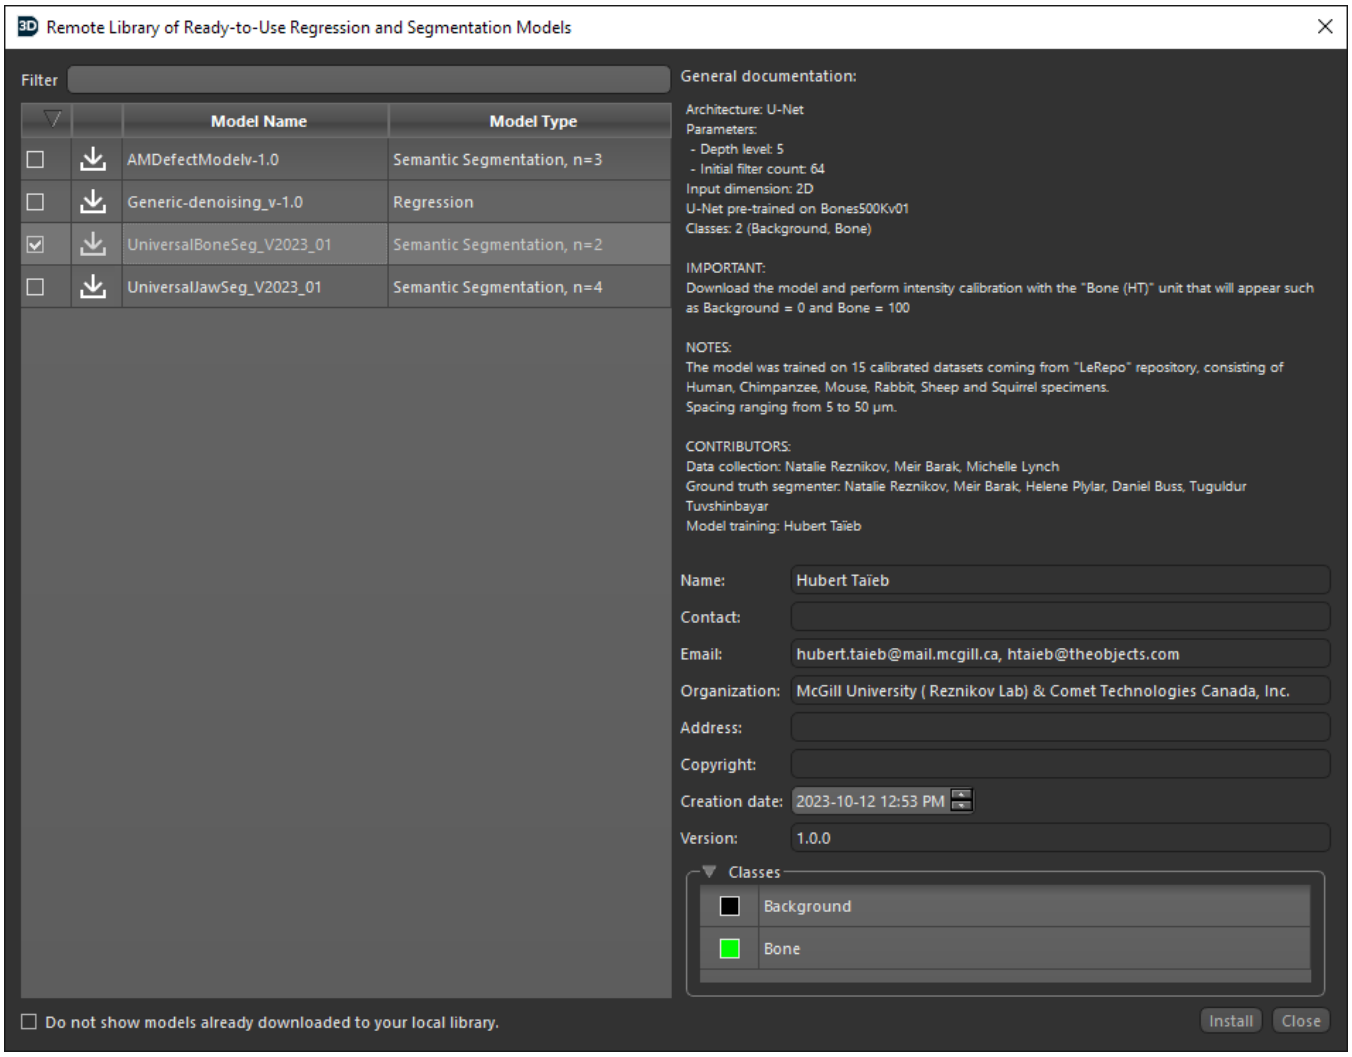

NOTE	Ready-to-use models are filtered by version, that is, only models that can be used by the currently installed version of Dragonfly will appear in the list.

NOTE	You need to be connected to the Internet to access the remote library of ready-to-use deep models.

NOTE	Contact us if you are interested in sharing any of your own ready-to-use deep models. You can also upload trained models to the online platform **Dragonfly Social** (see Dragonfly Social on page 9).

Ready-To-Use Models

The ready-to-use models listed below are now available for download from the Remote Library of Ready-to-Use Regression and Segmentation Models dialog.

Ready-to-use models

Model	Description
AMDefectModelv-1.0	Trained for segmenting pores, material, and background on calibrated datasets of AM metal cubes. **Architecture:** U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2.5D, input slice count: 3 **Classes:** 3 (pores, material, background) Input data must be calibrated such that the background = 0 and material = 100. **Reference** Taute, C., Möller, H., Du Plessis, A., Tshibalanganda, M. and Leary, M., 2021. *Characterization of additively manufactured AlSilOMg cubes with different porosities*. Journal of the Southern African Institute of Mining and Metallurgy, 121(4), pp.143-150.
Generic-denoising_v-1.0	This model was trained for general denoising with the following settings: **Architecture:** U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2D
UniversalBoneSeg	This model can be used for segmenting bone and was pre-trained on Bones500Kv01 and then fine-tuned on 15 calibrated datasets consisting of human, chimpanzee, mouse, rabbit, sheep, and squirrel specimens. Spacing values ranged from 5 to 50 µm. The settings of the model are as follows: **Architecture:** U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2D **Classes:** 2 (background, bone) Input data must be calibrated such that background = 0 and bone = 100.
UniversalJawSeg	This model can be used for segmenting jawbone, dentin, and enamel and was pre-trained on ORS500Kv01 and then fine-tuned on four calibrated datasets from the repository "LeRepo". Spacing values ranged from 15 to 50 µm. The settings of the model are as follows: **Architecture:** U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2D **Classes:** 4 (background, bone, dentin, enamel) Input data must be calibrated such that background = 0 and enamel = 100.

Updates for Segment with AI and Filter with AI

The option to download ready-to-use deep models for semantic segmentation or regression tasks is also available in the Segment with AI and Filter with AI dialogs. Click the **Import Remote** button, as shown below, to open the Remote Library of Ready-to-Use Regression and Segmentation Models dialog.

Segment with AI and Filter with AI dialogs

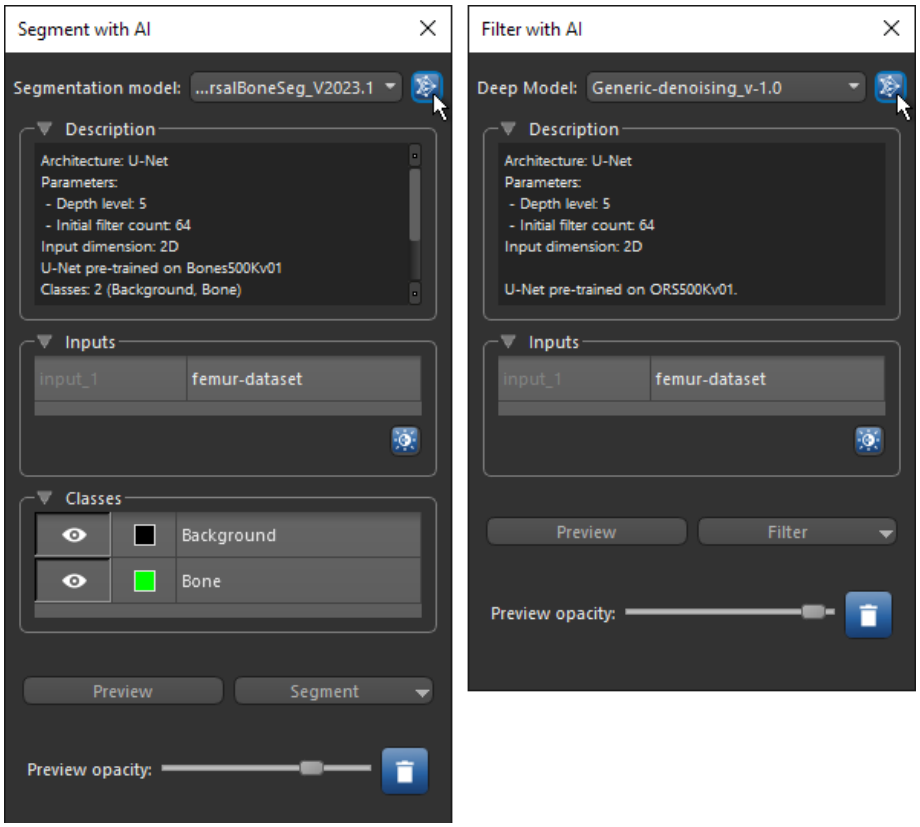

Additional changes to the dialogs include the addition of the description for the selected model.

NOTE Refer to the topic Ready-To-Use Deep Model Library on page 17 for additional information about ready-to-use deep models.

NOTE The 'Segment on Multiple Axes' feature, which lets you apply a semantic segmentation model on multiple axes and predict class labels based on a selected type of type of vote, is not available in this software release. This feature will be reinstated as soon as possible.

Additional Pre-Trained Models

You can now download additional pre-trained deep models to help you get started quickly with implementing Dragonfly's deep learning solutions. You should note that starting with a pre-trained model often provides better and faster results with smaller training sets than using an untrained model.

The following additional pre-trained models are now available for download from the Deep Learning Tool and from the Segmentation Wizard.

New pre-trained models

Model	Description
Generic Trans U-Net dl-5 ifc-32	**Architecture:** Trans U-Net **Parameters:** - **Patch size:** 224 - **Depth level:** 5 - **Initial filter count:** 32 - **Use batch normalization:** True **Input dimension:** 2D
Generic Trans U-Net dl-5 ifc-32 slice-3	**Architecture:** Trans U-Net **Parameters:** - **Patch size:** 224 - **Depth level:** 5 - **Initial filter count:** 32 - **Use batch normalization:** True **Input dimension:** 2,5D, input slices count: 3
Generic Attention U-Net dl-5 ifc-64	**Architecture:** Attention U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2D
Generic INet fc-64	**Architecture:** INet **Parameters:** - **Filter count:** 64 **Input dimension:** 2D
Generic Attention U-Net dl-5 ifc-64 slice-3	**Architecture:** Attention U-Net **Parameters:** - **Depth level:** 5 - **Initial filter count:** 64 **Input dimension:** 2,5D, input slices count: 3
Generic INet fc-64 slice-3	**Architecture:** INet **Parameters:** - **Filter count:** 64 **Input dimension:** 2.5D, input slices count: 3

Pre-trained models are available in the Model Generator dialogs as shown below. Refer to the Dragonfly 3D World Help topic '*Pre-Trained Models*' for additional information about working with pre-trained models.

Deep Learning Tool Model Generator dialog

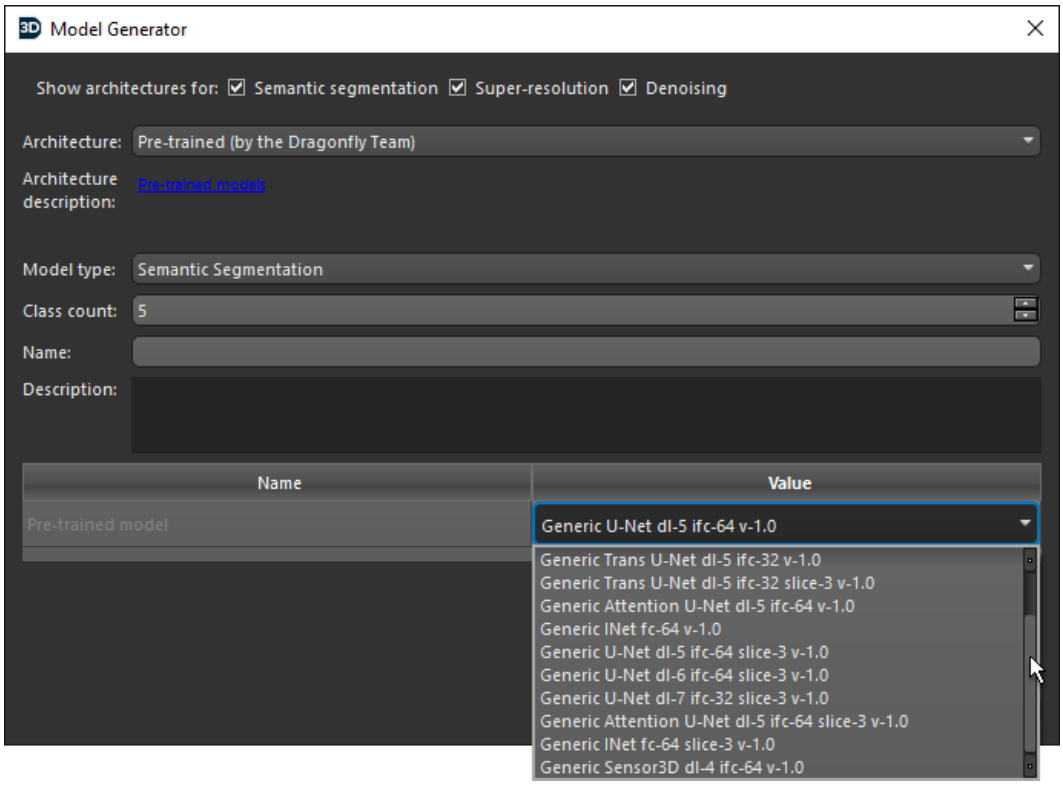

Segmentation Wizard Model Generator dialog

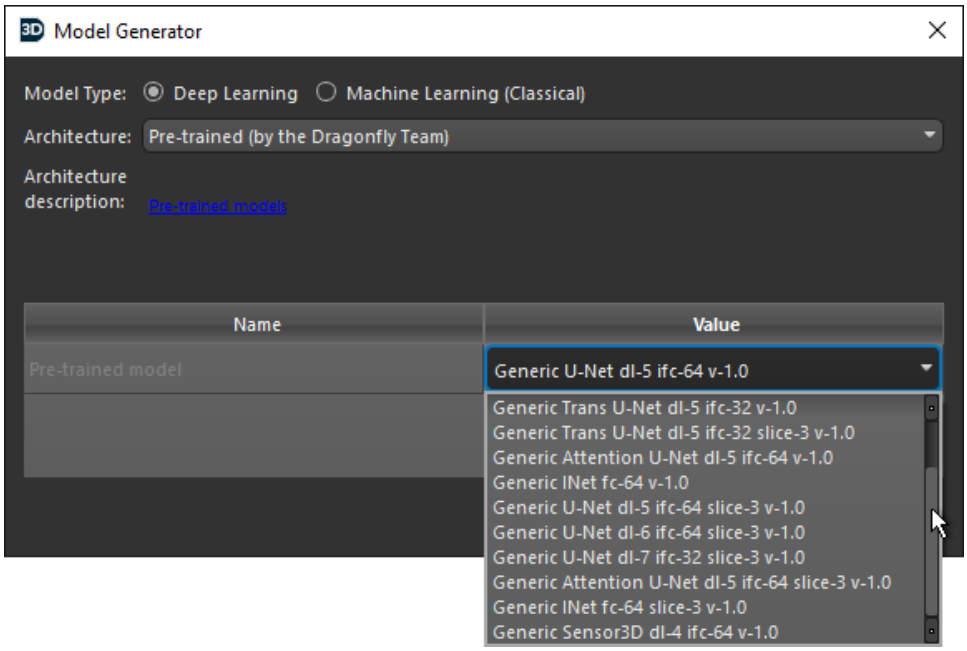

Sparse Labeling

As an option for accelerating workflows for training semantic segmentation models, this software release includes support for sparse labeling. With sparse labeling, you should be able to achieve comparable results to those obtained with densely labeled ground truths, which can be much more laborious to obtain.

New Architectures

Several new deep model architectures are available in this software release for training regression and semantic segmentation models. These newly added architectures are listed below.

FreqNet… A frequency-domain image super-resolution model that can explicitly learn the reconstruction of high-frequency details from low resolution images. Refer to R. Cai et al., *FreqNet: A Frequency-domain Image Super-Resolution Network with Discrete Cosine Transform*, rXiv:2111.10800v1, 2012 (https://arxiv.org/pdf/2111.10800.pdf) to learn more about this model.

Multi-Level Wavelet U-Net… A U-Net style encoder/decoder network with discrete wavelet transform (DWT) encoding and inverse wavelet transform (IWT) decoding for image denoising and single image super-resolution. Haar wavelets are used to perform the transformations. Refer to P. Liu et al., *Multi-level Wavelet Convolutional Neural Networks*, arXiv:1907.03128v1, 2019 (https://arxiv.org/pdf/1907.03128.pdf) for additional information about this model.

Multi-Stage Wavelet Denoise… A multi-stage image denoising model based on a dynamic convolutional block, two cascaded wavelet transform and enhancement blocks, and a residual block. Refer to C. Tian et al., *Multi-stage image denoising with the wavelet transform,* arXiv:2209.12394v3, 2022 (https://arxiv.org/pdf/2209.12394.pdf) for additional information about this model.

Non-local Fourier U-Net… A U-Net style encoder/decoder network using non-local fast Fourier convolution (NL-FCC) layers that has shown promising results in image super-resolution. The NL-FFC layer combines local and global features in spectral domain and spatial domain. Refer to A. K. Sinha et al., *NL-FFC: Non-Local Fast Fourier Convolution for Image Super Resolution*, 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (DOI: 10.1109/CVPRW56347.2022.00062).

Swin UNETR… A U-Net architecture combining a pure Swin transformer encoder and a deconvolutional decoder. Refer to A. Hatamizadeh et al., *Swin UNETR: Swin Transformers*

for Semantic Segmentation of Brain Tumors in MRI Images, arXiv:2201.01266v1, 2022 (https://arxiv.org/pdf/2201.01266.pdf) for additional information about this model.

TransUNet… A U-Net architecture using transformer layers as a bottleneck. The architecture has been optimized with TransBTS hyperparameters and registers. Refer to J. Chen et al., *TransUNet: Transformers Make Strong Encoders for Medical Image Segmentation*, arXiv:2102.04306v1, 2021 (https://arxiv.org/pdf/2102.04306.pdf) for additional information about this model, as well as W. Wang et al., TransBTS: Multimodal Brain Tumor Segmentation Using Transformer, arXiv:2103.04430v2, 2021 (https://arxiv.org/pdf/2103.04430.pdf) and T. Darcet et al., *Vision Transformers Need Registers*, arXiv:2309.16588v1, 2023 (https://arxiv.org/pdf/2309.16588.pdf).

More Efficient Patch Blending for Accelerated Inference

This software release includes more efficient patch blending for accelerated inference. The following changes were implemented to increase inference speeds:

- Removal of the maximum patch size constraint and optimization of patch size computation to reduce patch blending.
- Model predictions and patch blending are now done in parallel.

Refer to the following graphs for a comparison of inference speeds recorded for versions 2022.2 and 2024.1 for a semantic segmentation task on high-performance and typical GPUs. Test datasets varied in slice sizes from 500×500 pixels to 2500×2500 pixels with a constant depth of 782 pixels.

Inference times for high-memory and low-memory GPUs for Dragonfly versions 2022.2 and 2024.1

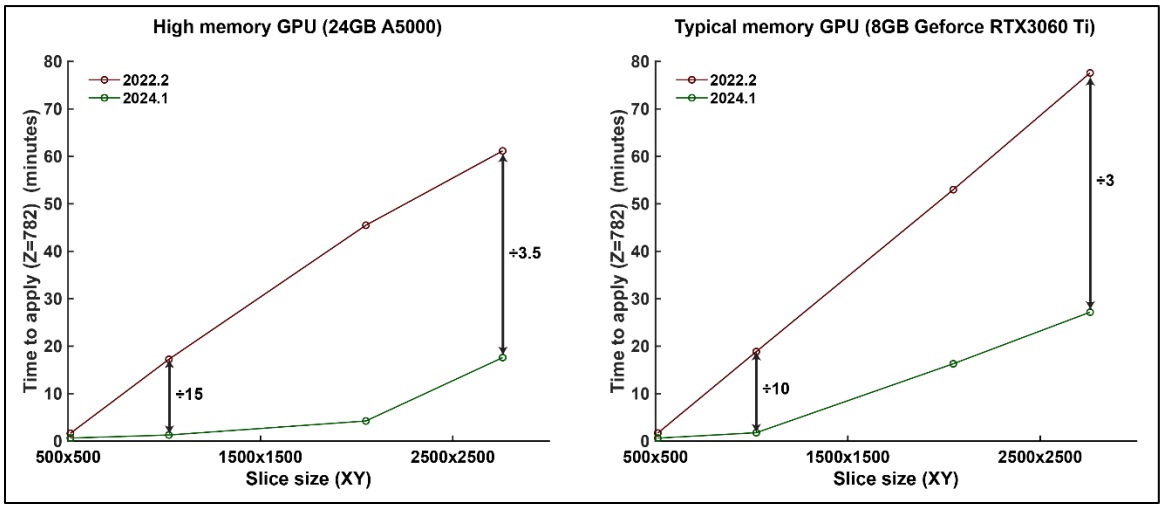

Compile for Faster Inference

This software release includes the opportunity to increase inference efficiency for lower end GPUs with TensorFlow-TensorRT (TF-TRT), a deep-learning compiler for TensorFlow that optimizes models for inference. This process includes five types of optimizations — precision calibration, layers and tensor fusion, kernel auto-tuning, dynamic tensor memory, and multiple stream execution. Inference efficiency can be a concern when deploying deep models on systems with limited GPUs because of latency, memory utilization, and power consumption.

The option to compile deep models for faster inference is available in the Deep Learning preferences, as shown below.

Deep Learning preferences

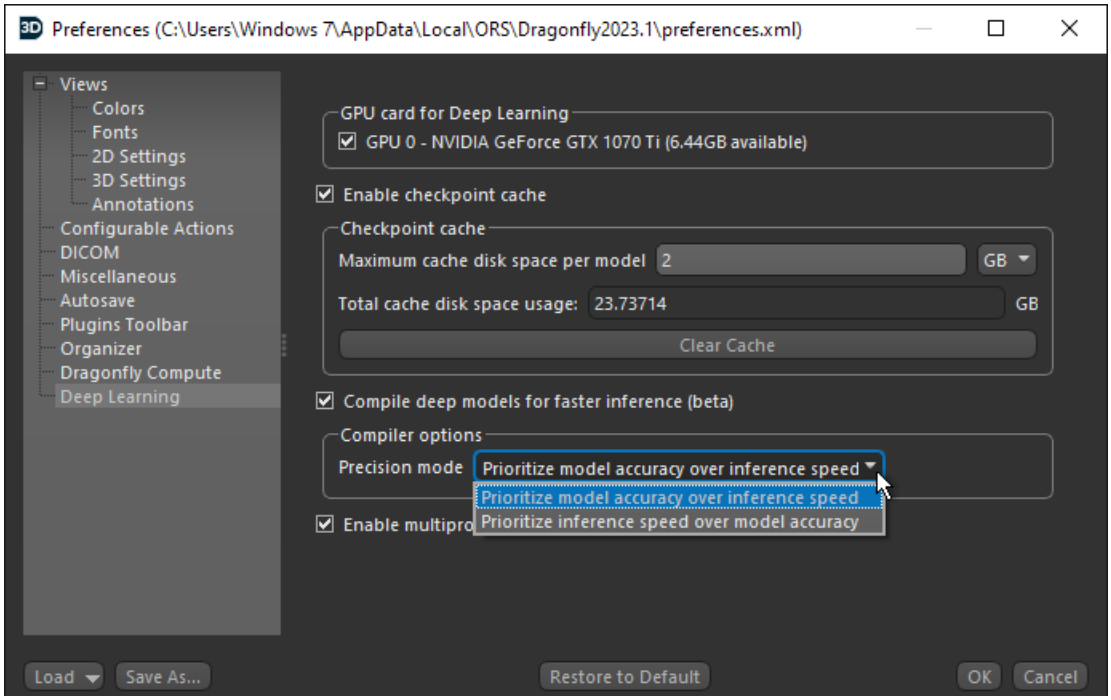

NOTE TensorFlow-TensorRT (TF-TRT) is a deep-learning compiler for TensorFlow that optimizes TF models for inference on NVIDIA devices. TF-TRT is the TensorFlow integration for NVIDIA's TensorRT (TRT) High-Performance Deep-Learning Inference SDK. It focuses specifically on running an already-trained network quickly and efficiently on NVIDIA hardware.
Refer to https://docs.nvidia.com/deeplearning/frameworks/tf-trt-user-guide/index.html for more information about TF-TRT.

The following options are available for deep model compilation:

Compile deep models for faster inference (beta)... If selected, deep models will be compiled for faster inference. If not selected, deep models will be compiled with the default settings as in previous versions of Dragonfly.

Precision mode... You can select from two precision modes to compile a deep model for inference — 'Prioritize inference speed over model accuracy' or 'Prioritize model accuracy over inference speed'.

As shown in the Segmentation Comparator dialog, prioritizing inference speed over model accuracy may adversely affect accuracy.

Comparison of Precision modes with a ground truth

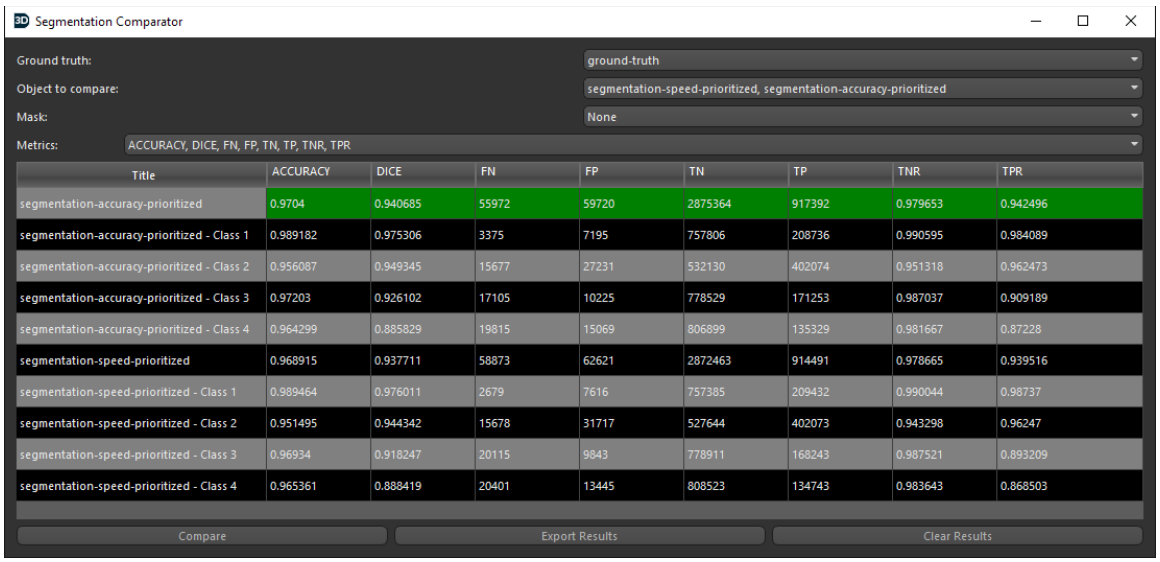

Title	ACCURACY	DICE	FN	FP	TN	TP	TNR	TPR
segmentation-accuracy-prioritized	0.9704	0.940685	55972	59720	2875364	917392	0.979653	0.942496
segmentation-accuracy-prioritized - Class 1	0.989182	0.975306	3375	7195	757806	208736	0.990595	0.984089
segmentation-accuracy-prioritized - Class 2	0.956087	0.949345	15677	27231	532130	402074	0.951318	0.962473
segmentation-accuracy-prioritized - Class 3	0.97203	0.926102	17105	10225	778529	171253	0.987037	0.909189
segmentation-accuracy-prioritized - Class 4	0.964299	0.885829	19815	15069	806899	135329	0.981667	0.87228
segmentation-speed-prioritized	0.968915	0.937711	58873	62621	2872463	914491	0.978665	0.939516
segmentation-speed-prioritized - Class 1	0.989464	0.976011	2679	7616	757385	209432	0.990044	0.98737
segmentation-speed-prioritized - Class 2	0.951495	0.944342	15678	31717	527644	402073	0.943298	0.96247
segmentation-speed-prioritized - Class 3	0.96934	0.918247	20115	9843	778911	168243	0.987521	0.893209
segmentation-speed-prioritized - Class 4	0.965361	0.888419	20401	13445	808523	134743	0.983643	0.868503

NOTE Optimized inference usually works best for simple models, such as U-Net. It might not work well or it may not be used to compile more complex models, such as Sensor3D. There are also only limited benefits for optimizing inference for high-end GPUs.

New Advanced Tools

Two new advanced tools — **Inspect Patches** and **Generate Augmented Trainset** — are available on the Advanced Tools tab on the Model Training panel in this release.

Advanced Tools tab

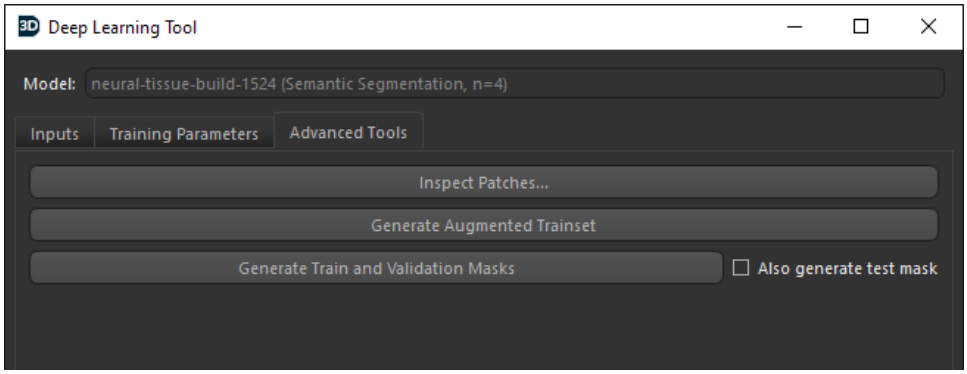

Patch Inspection

Patch size can be a critical factor when training deep models for semantic segmentation and must be selected in consideration of the specific characteristics of the input data and the selected model architecture. Generally, patch sizes should strike a balance between computational efficiency, contextual information, and performance for the specific semantic segmentation task at hand.

To review patches, click the **Advanced Tools** tab on the Model Training panel and then click the **Inspect Patches** button. Patches will be extracted based on the current patch size, data augmentation settings, and the selected input count and dimension. In the Patches Review dialog, you can review training and validation patches, visually inspect each patch, and extract additional information about all patches.

Patches Review dialog

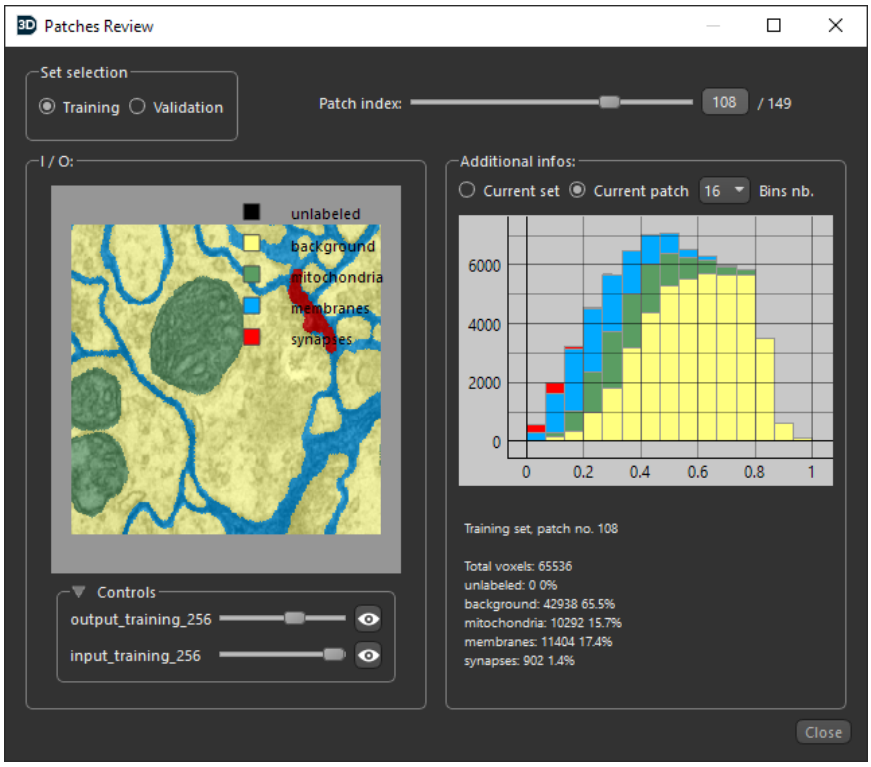

Generate Augmented Deep Learning Training Sets

In some cases, you may want to generate an augmented training set to increase training speed. Augmented training sets, which include training and validation sets, are generated based on the current patch size, data augmentation settings, and the selected input count and dimension.

To generate an augmented training set, open the **Advanced Tools** tab on the Model Training panel and then click **Generate Augmented Trainset**. Training and validation sets, as well as masks, are added to the Data Properties and Settings panel as shown.

Extracted augmented training set

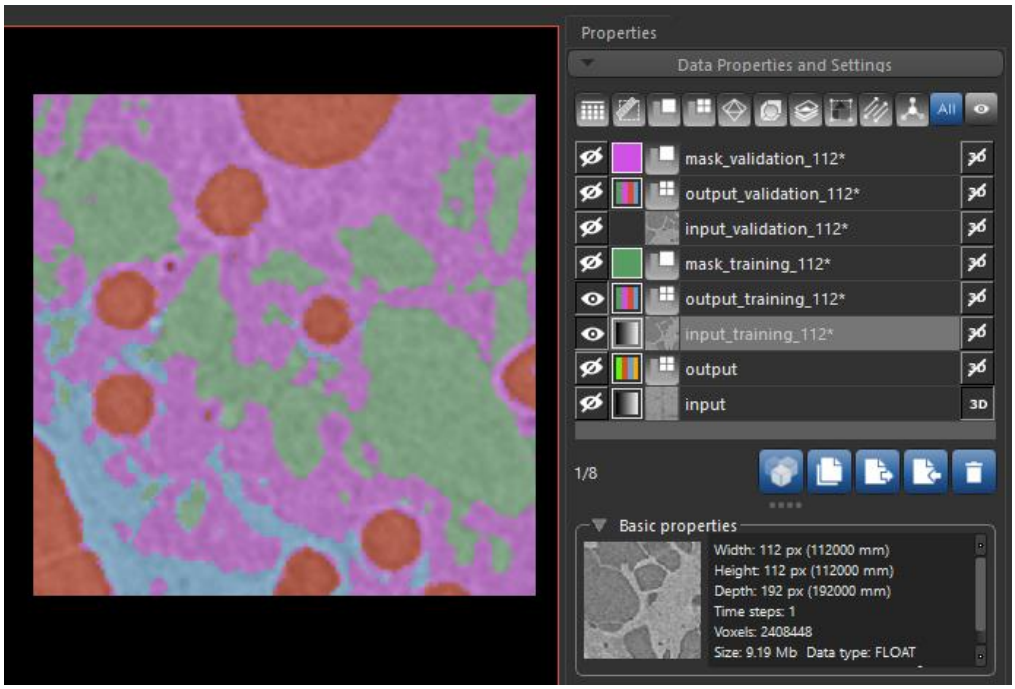

NOTE You can train models with augmented training sets without using data augmentation. This may improve training speed but will have less diversity in augmentation.

Automatic Loading/Unloading of Models

With the goal of providing a better and more uniform experience when using the Deep Learning Tool, this software release includes automatic loading and unloading of deep models. This means that there is no longer any need to load or unload models manually prior to training, generating previews, or applying a model. This implemented change will also ensure that there is never more than one model loaded at any given time.

NOTE The **Load** and **Unload** buttons were removed from the Deep Learning Tool dialog and the model states **Not Loaded** and **Ready** were deprecated in this software version.

Automatic Saving of Models

You should note that models are now saved automatically after training is completed or stopped. This change and others, can be seen in the updates to the options available at the bottom of the Deep Learning Tool dialog, as shown below.

Comparison between previous versions (top) and version 2024.1 (bottom)

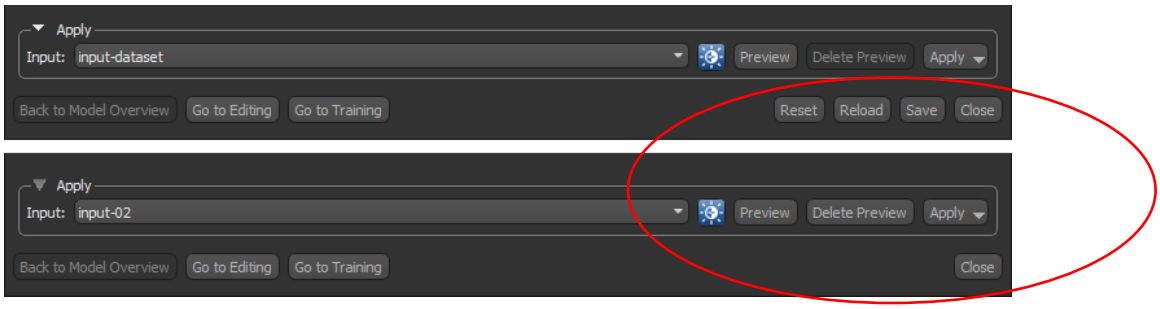

Reset… This option, which lets you return models to their untrained state by randomizing their weights, was moved to the Model box on the Model Overview panel.

Reload… This option, which lets you revert a model back to its last saved version, was renamed as **Revert** and is now available only on the Model Editing panel.

Save… This option, which lets you save changes to a model's architecture, is now only available on the Model Editing panel.

Export Training Sets in a Single File

New options in the Deep Learning Tool, circled below, let you export all your training data, as well as import training data that was exported.

Export as AI and Import as AI buttons

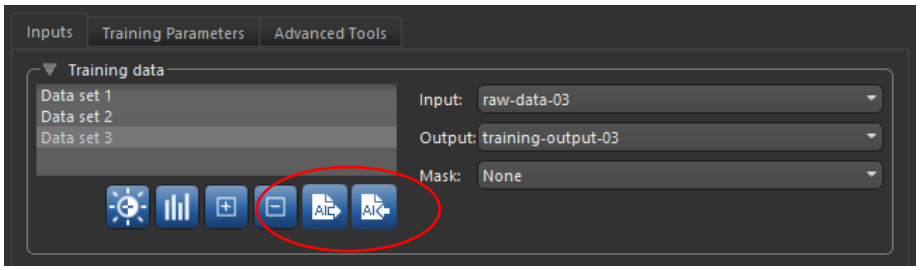

Export as AI… Saves all training data inputs, outputs, and masks as slices in the ORSObject (*.ORSObject extension) file format. Saving only data that is used for training can significantly reduce the size of training sets.

Import as AI… Imports all training data that was exported as AI. All saved inputs, outputs, and masks for each training set will be added to the dialog when the data is imported, minimizing any additional manual entry.

External Processes for Deep Learning Task Processing

You should note that TensorFlow functionality, such as generating models, model training, and inference, can be run as an external process in Dragonfly 3D World and Dragonfly Pro for Windows installations. The main benefit of this is that memory can be released for subsequent processes after training or inference is completed. In previous versions, memory used by TensorFlow could only be released by exiting Dragonfly.

You need to check the setting 'Execute Deep Learning tasks in an external process' to run separate processes for TensorFlow functionality, as circled below.

Deep Learning preferences

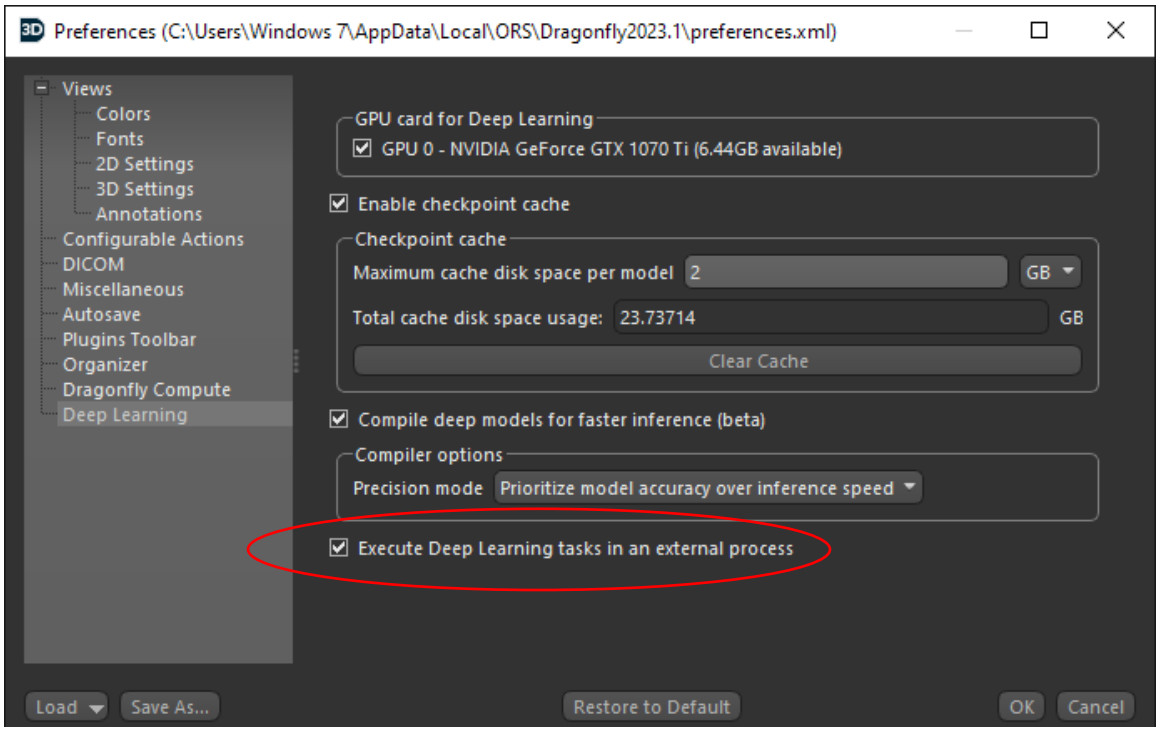

NOTE	Running TensorFlow functionality as an external process is not available for Linux installations. In this case, the 'Execute Deep Learning tasks in an external process' option is disabled by default.

NOTE	This feature is only available as beta in this software release.

Select Dataset for Visual Feedback

You can now select any dataset for visual feedback. For example, in cases in which you are using multiple training sets to train a model or if you want to monitor training progress on a dataset that is not part of your training data. The option for selecting a dataset for visual feedback is available in the Deep Learning Tool dialog on the Inputs tab as shown below.

Input options for Visual feedback

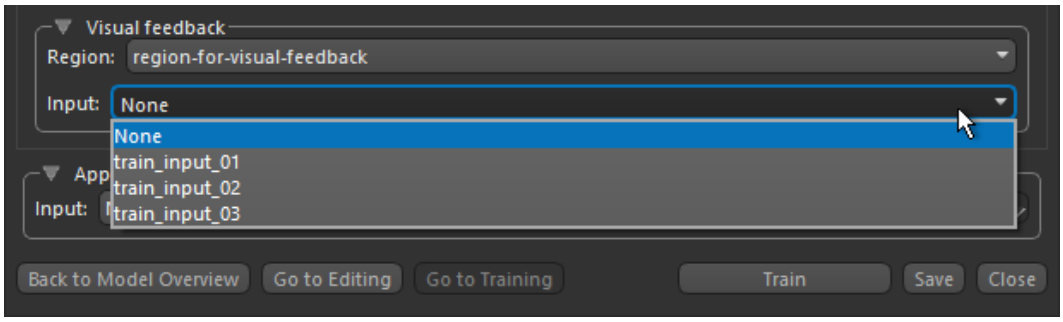

Validation Settings

New validation settings — **Use voxel-wise sampling** and **Use patch-wise sampling** — are available in this software release for training segmentation models. These settings, circled below, are available on the Training panel's Inputs tab. You should note that in either case, the maximum number of patches will be extracted from the training set. The difference is that in the case of voxel-wise sampling, patches will contain voxels designated as training, validation, and test, while in the case of patch-wise sampling, voxels in sampled patches will be exclusively for training, validation, or test.

Validation settings

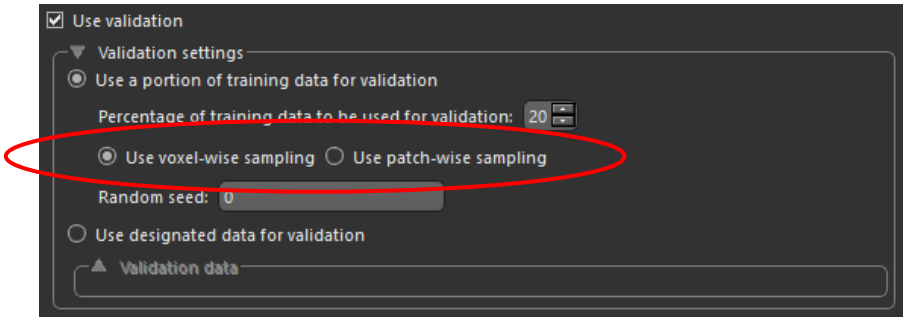

NOTE In the Segmentation Wizard, sampling is now done exclusively voxel-wise when data is split for training, validation, and test. This setting cannot be changed.

Easy Sharing of Deep Models

To enable easy sharing of deep models, this software release includes the option to export models to a ZIP file and to import models from ZIP files. The **Import Zip** and **Export ZIP** options are circled on the screenshot.

Import Zip and Export Zip buttons

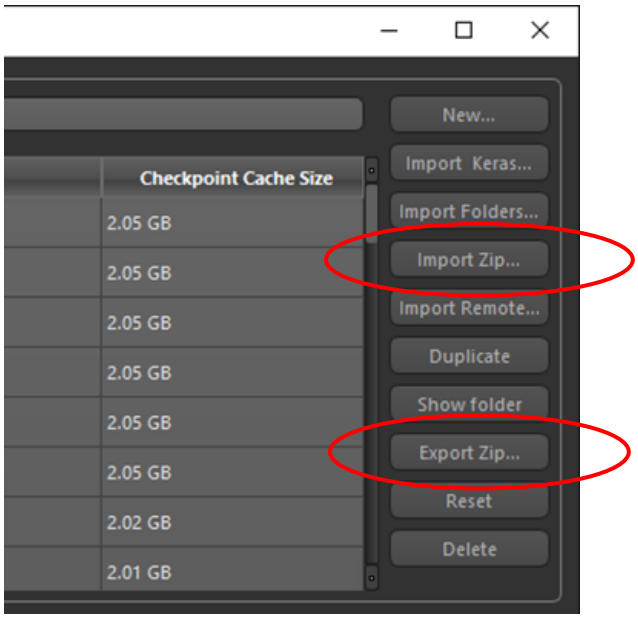

NGSolve Integration

For quantifying fluid flow and heat dissipation in engineering and research applications, this software release includes the integration of Netgen/NGSolve. You should note that Netgen is a mesh generation tool while NGSolve is a widely used high-performance finite element analysis software. This implementation in Dragonfly includes a Navier-Stokes equation for solving transient fluid mechanics, a Stokes model to solve steady-state fluid flow, and a heat equation for steady-state conduction. In addition, new partial differential equations and solution algorithms can be implemented easily thanks to Dragonfly's flexible Python interface.

Below is an example of a fluid mechanics simulation done with NGSolve in Dragonfly, thanks to the vector field generation capability of Dragonfly.

Fluid mechanics simulation: Navier-Stokes flow around an obstacle with square cross-section

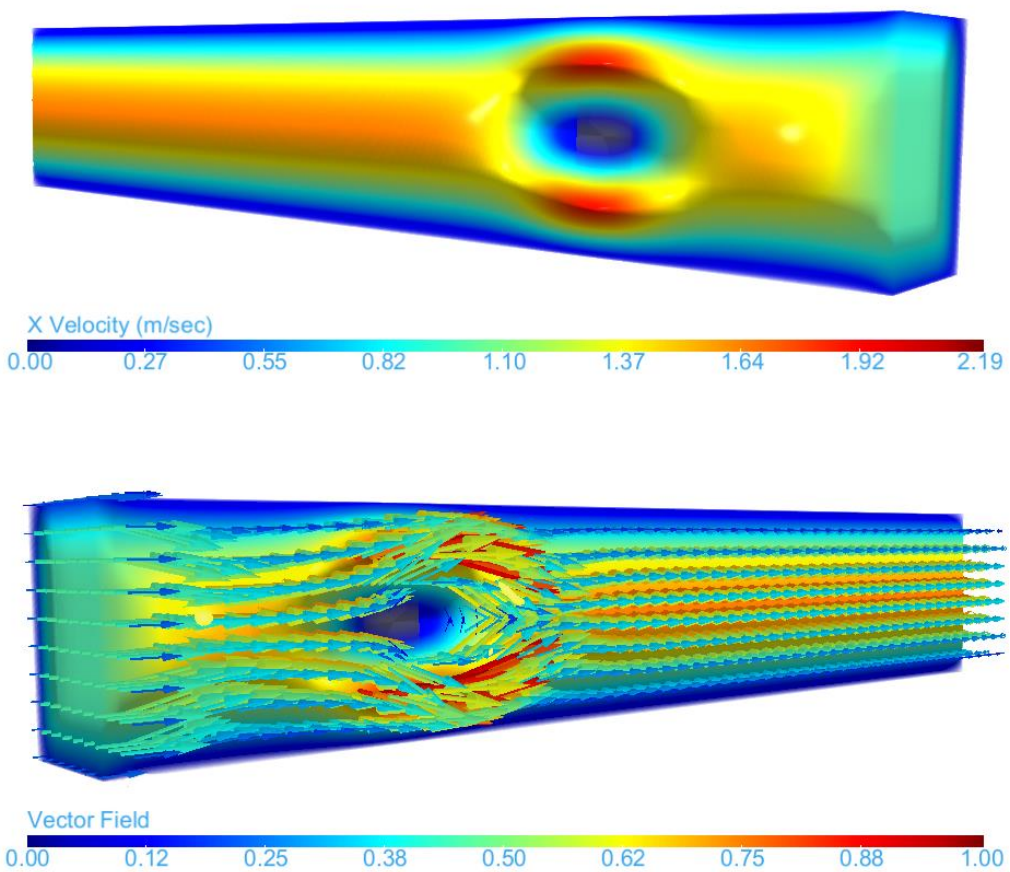

You should note that Dragonfly provides the option to assign boundary conditions using a mesh painter (see 3D Mesh Painter on page 51).

Choose **Workflows** > **NGSolve** on the menu bar to open the NGSolve Computation dialog. Shown below is the **Meshing** tab, on which you can generate finite element meshes from regions of interest and multi-ROIs, and the **Solver** tab, on which you can choose a solver, the simulation parameters, and the required outputs.

NGSolve Computation dialog

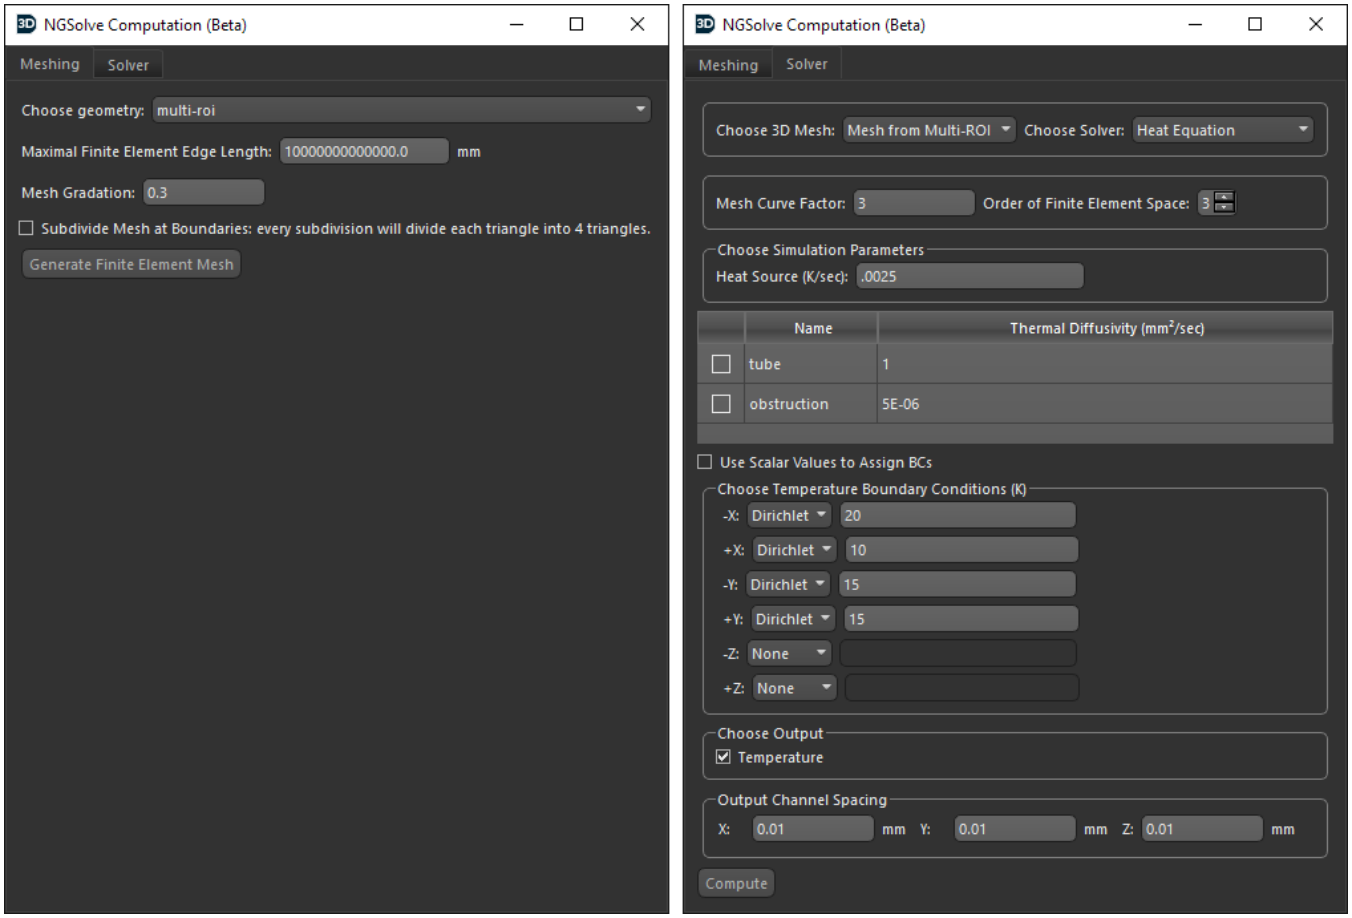

Tutorials and implementation examples for NGSolve and the mesh generator Netgen are available on the NGSolve YouTube channel (https://www.youtube.com/@NGSolve). Documentation about Netgen/NGSolve is available at https://docu.ngsolve.org/latest/.

NOTE You can set boundary conditions for simulations for solving heat equations and other simulations with the 3D Mesh Painter (see 3D Mesh Painter on page 51).

You can export generated meshes from Dragonfly for use in other finite element software.

Open Fiber Segmentation

This software release includes an innovative segmentation tool for automatically extracting individual fibers from datasets of short fiber reinforced composites. Based on the Open Fiber Segmentation algorithm developed by Facundo Sosa-Rey et al. (https://github.com/lm2-poly/OpenFiberSeg), the input is a region of interest of segmented fibers and the output is a multi-ROI in which each label is an instance of the segmented fibers. The method only requires minimal input — selection of the fiber parameters of minimum length, smoothing diameter, and maximum stitching distance, as well as the convexity check parameters. The output allows for the analysis of correlated length, orientation distribution, and other parameters.

Choose **Workflows** > **Open Fiber Segmentation** to open the Open Fiber Segmentation panel, shown below.

Open Fiber Segmentation panel and segmented fibers in a 3D view

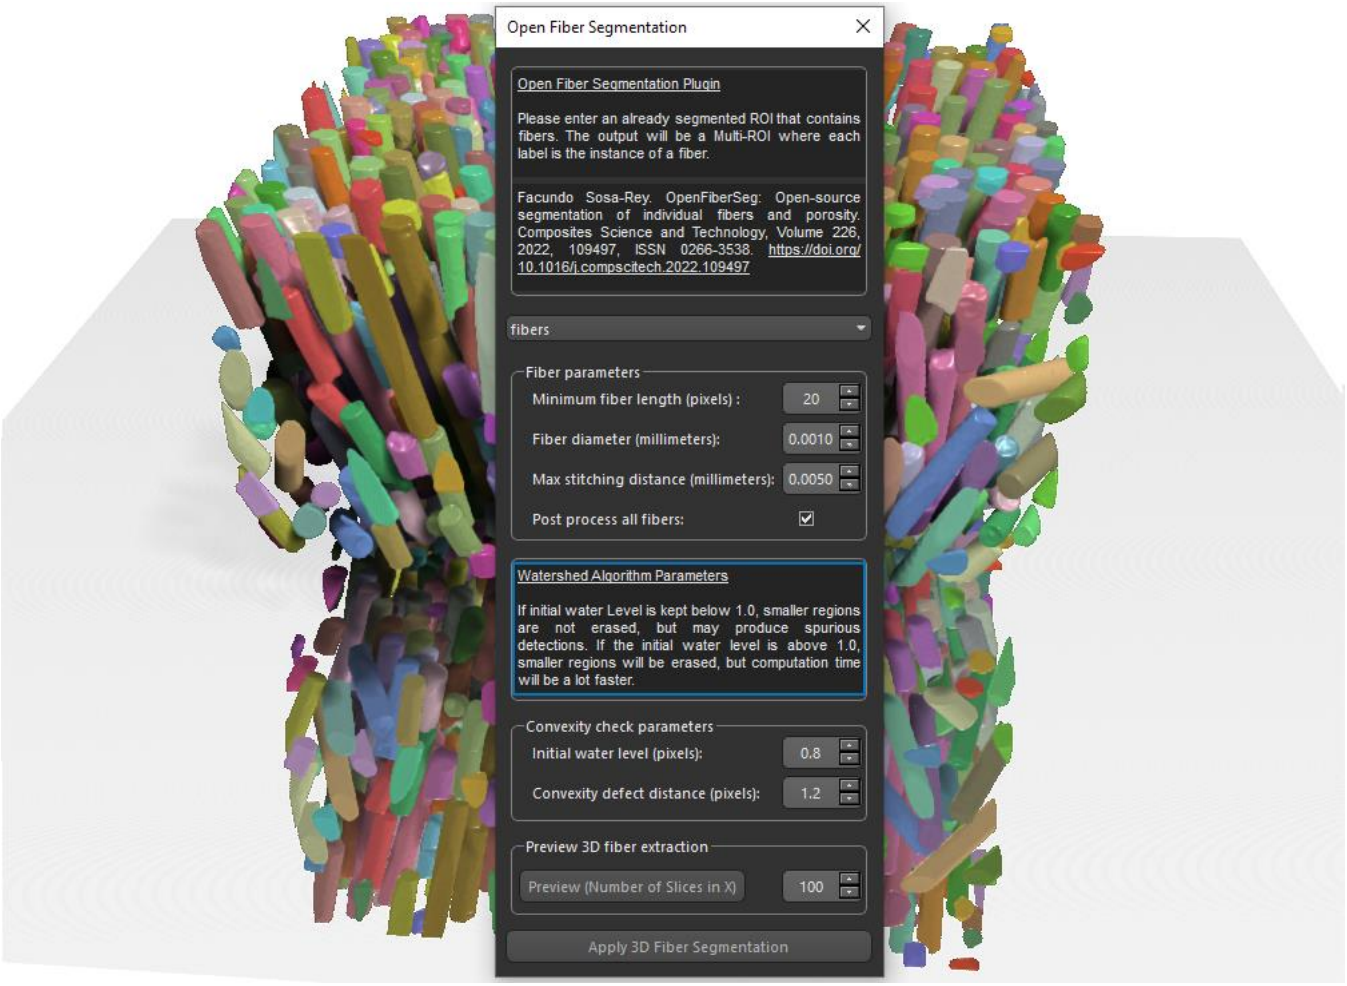

The following parameters and settings are available for the instance segmentation of fibers with the Open Fiber Segmentation plugin.

Parameters for the instance segmentation of fibers

Parameter	Description
Fiber parameters	
Minimum fiber length (pixels)	Is the minimum value of fiber length, in pixels, that the algorithm will accept as a fiber.
Fiber diameter (default unit)	Is used mostly for representation, will smooth the fibers with the selected diameter value.
Max stitching distance (default unit)	Is the maximum distance between two fibers that will be considered to stitch both fibers as one.
Post process all fibers	If checked, all fibers will go through the post processing algorithm for smoothing, gap filling.
Convexity check parameters	
Initial water level (pixels)	Is the distance parameter to find all the convex regions using a 'cv.watershed' function. **NOTE** If the initial water level is kept below 1.0, smaller regions will not be erased but spurious detections may occur. If the initial water level is above 1.0, then smaller regions will be erased but computation time will be faster.
Convexity defect distance (pixels)	Is the distance parameter, which is defined as a convexity defect size to find regions that are not convex, to verify if two fibers are touching.
Preview 3D fiber extraction	
Preview (number of slices in X)	Will compute the specified number of slices in X as a preview.

NOTE Refer to the publication Facundo Sosa-Rey et al. *OpenFiberSeg: Open-source segmentation of individual fibers and porosity*. Composites Science and Technology, Volume 226, 2022, 109497, ISSN 0266-3538. DOI: https://doi.org/10.1016/j.compscitech.2022.109497.

Automated Porosity Segmentation

The Dragonfly 3D World 2024.1 release provides automated thresholding methods to reliably segment porosity in single-material samples with closed pores, such as castings. This lets you quickly and accurately prepare data for analyzing pores with parameters such as volume fraction percentage, pore size distribution, sphericity, and others.

Right-click the image data that you need to segment and then choose **Automated Porosity Segmentation** in the pop-up menu. You can then choose a method to segment pores. The result of an automated porosity segmentation is shown below.

Automated porosity segmentation (colored red) of an automotive part

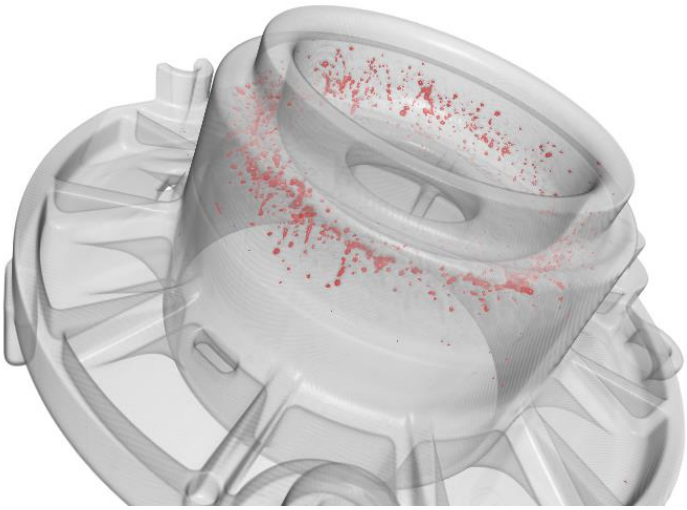

The following options are available for applying an automated porosity segmentation.

Basic Otsu… This is the fastest available method and typically captures the largest pores in the sample but may not detect smaller ones. The Basic Otsu method includes creating a region of interest for the air and void space based on the lower Otsu threshold for the full 3D dataset. This ROI is then refined to remove segmented exterior air using the 'Process Islands' tool to remove the largest component within the ROI.

Segmentation of pores according to the Basic Otsu method for a casting dataset

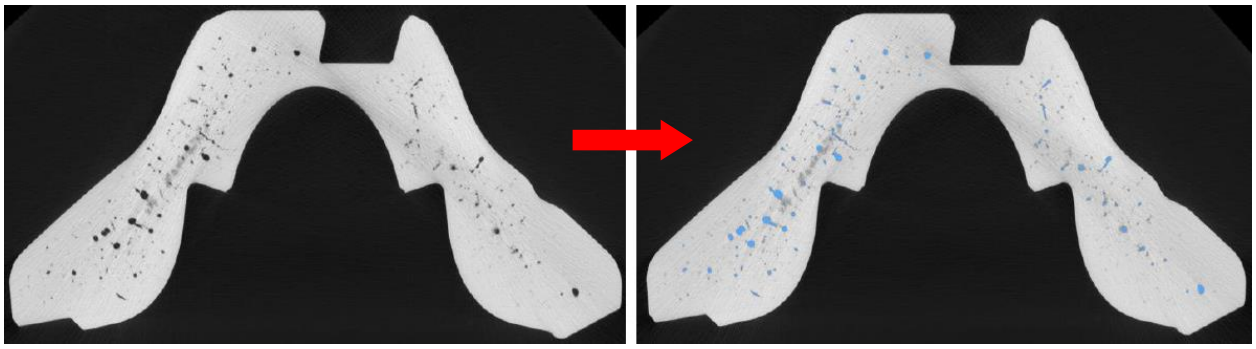

Advanced Otsu… This fast and reliable method typically captures smaller pores than the Basic Otsu method and creates two regions of interest that describe the pores and material in the sample respectively.

In this method, a region of interest is first defined for the material using the upper Otsu threshold. Internal pores are then filled and a dilate function is applied with a 3D kernel of size 5 to increase the ROI size beyond the edge of the segmented material. The ROI is then split at the Otsu threshold to effectively separate the pores from the segmented material. The layer around the material in the resulting 'Porosity ROI' is then removed with the 'Process Islands' tool.

Segmentation of pores according to the Advanced Otsu method for a casting dataset

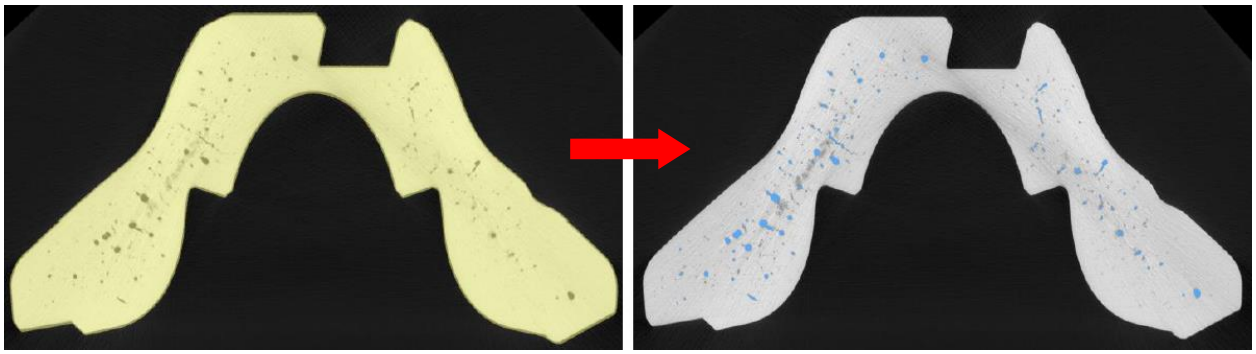

Sobel Threshold… This advanced method can find the smallest pores in a sample and can be applied to challenging data. For example, in cases in which brightness variations are present across the image and in cases of low signal to noise.

For the segmentation step, this method uses Sobel image filtering without any pre-smoothing to highlight the edges of pores. This initial ROI is then 'closed' and inner areas are subsequently filled, while the edge that defines the material is removed.

Segmentation of pores according to Sobel method for a casting dataset shown

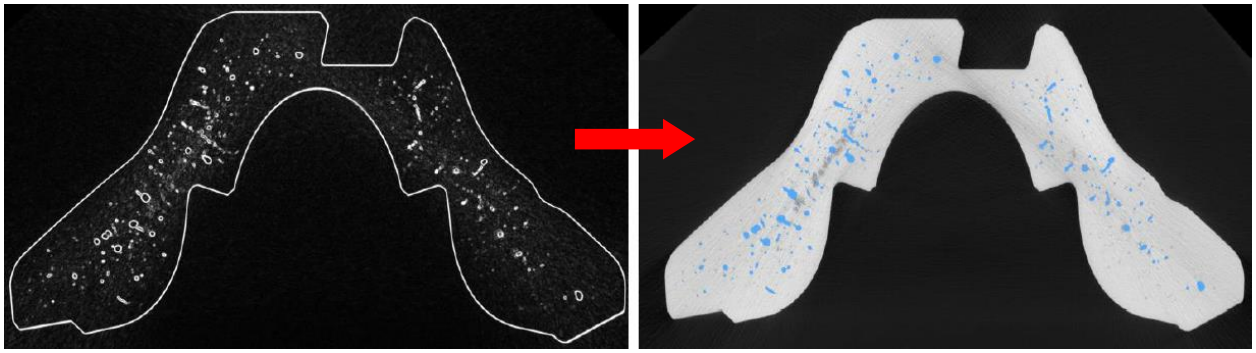

Compute Measurements

The option to compute measurements from multi-ROIs, meshes, graphs, and vector fields is now available in the **Compute Measurements** dialog. You should note that this dialog simply replaces the Scalar Generator dialog that was available in previous versions of Dragonfly. In addition, new metrics — 'Radiomics with Dataset' and 'Longest Length' and — are available in this software release.

Right-click a multi-ROI, mesh, graph, or vector field and then choose **Compute Measurements** to open the dialog shown below.

Compute Measurements dialog

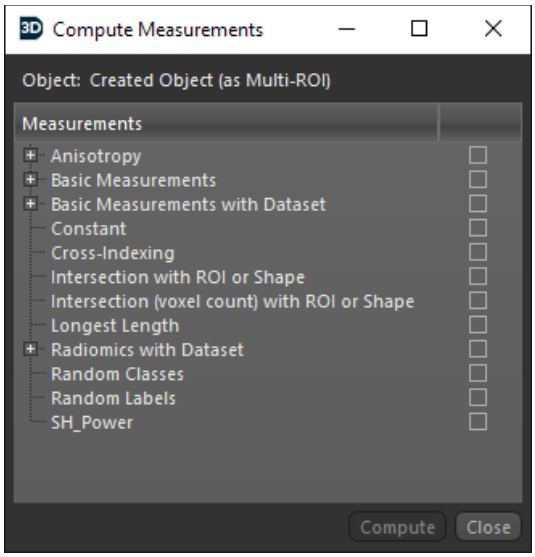

NOTE The option to compute measurements and statistical properties in the **Object Analysis** dialog was deprecated in this software version. Measurements can only be computed in the Compute Measurements dialog as of this software release.
Refer to the screenshots below for a comparison of the Object Analysis dialog in version 2022.2 (on the left) and in version 2024.1 (on the right).

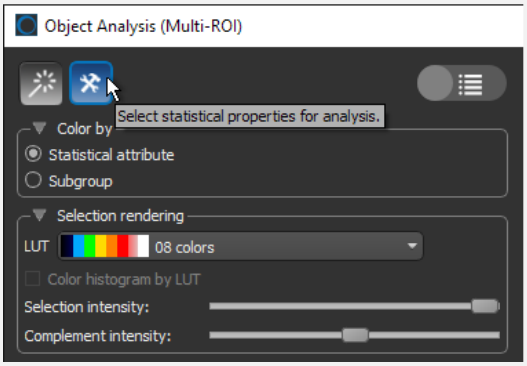

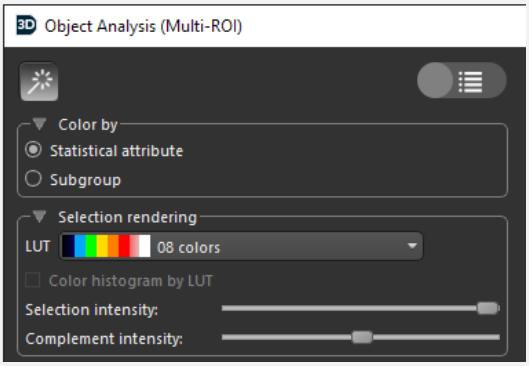

Radiomics

A new set of measurements — **Radiomics with Dataset** — is available in this software release for extracting quantitative features from images using data-characterization algorithms. Once features are extracted, computational algorithms and statistical models can be applied to analyze and interpret the extracted data.

The radiomic features now available in Dragonfly's Compute Measurements dialog can be divided into groups that include descriptors of the relationships between image voxels, as well as local texture variations and gradients. Right-click a multi-ROI and then choose **Compute Measurements** to open the dialog and access the Radiomics with Dataset set of measurements.

Compute Measurements dialog

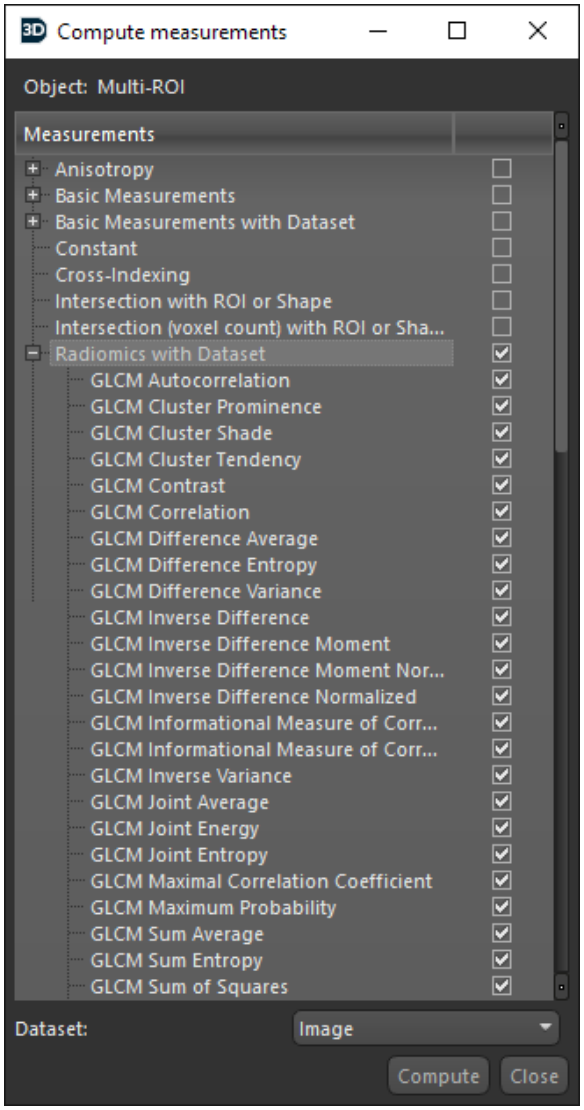

The various features that can be extracted using radiomics, on either an original or derived image, can be subdivided into the following classes:

Gray Level Co-occurrence Matrix (GLCM)… These 24 features provide for quantitative and objective analyzes of image textures.

Gray Level Dependence Matrix (GLDM)… The 14 features of the GLDM provide information about the coarser-scale texture properties of an image compared to GLCM features. The strength and direction of gray-level dependencies captured can be particularly useful for analyzing patterns characterized by variations in gray-level values.

Gray Level Run Length Matrix (GLRLM)… These 16 features capture information about the spatial arrangement and distribution of pixel runs in an image. Providing insights into the coarseness, regularity, and complexity of underlying textures, these features can identify regions with specific texture patterns for image segmentation.

Gray Level Size Zone Matrix (GLSZM)… The 16 features of the GLSZM provide information about the distribution and characteristics of connected regions within an image by capturing the heterogeneity, fragmentation, and spatial organization of texture patterns. Theses features can identify regions of interest with specific size and texture characteristics for image segmentation.

Neighbouring Gray Tone Difference Matrix (NGTDM)… The five features of the NGTDM provide information about the local texture variations and gradients in an image by capturing the differences in gray-level values between adjacent pixels to reveal fine-scale details of texture patterns. They can also identify regions of interest with specific local texture variations for image segmentation.

Gray Level Co-occurrence Matrix (GLCM)

The Gray Level Co-occurrence Matrix (GLCM) is a widely used method in radiomics for extracting texture-based features. GLCM quantifies the spatial relationships between pairs of pixels based on their gray-level intensities within an image and is constructed by counting the number of occurrences of pixel pairs with specific intensity values and their spatial relationships at a selected offset. This process is repeated for different offsets or directions to capture various spatial dependencies within the image, as described below.

Contrast… Measures the local intensity variations between neighboring pixels.

Energy (or Angular Second Moment)… Represents uniformity or image homogeneity.

Homogeneity… Measures the closeness of distribution of pixel pairs to the diagonal elements of the GLCM.

Entropy... Quantifies the randomness or uncertainty of the pixel-pair distribution.

Correlation... Measures the linear dependency between gray levels of neighboring pixels.

NOTE Refer to https://pyradiomics.readthedocs.io/en/latest/features.html#module-radiomics.glcm for more information about GLCM.

Gray Level Dependence Matrix (GLDM)

The Gray Level Dependence Matrix (GLDM) is another method used in radiomics to extract texture-based features. GLDM characterizes the spatial dependencies and relationships between gray-level values in an image. The GLDM is constructed by counting the number of occurrences of pixel pairs with specific gray-level differences and their corresponding spatial relationships at the selected distance. The statistical measures to quantify the texture properties of an image include:

Contrast... Measures the local intensity differences between neighboring pixel pairs.

Dissimilarity... Quantifies the average difference in gray-level values between neighboring pixel pairs.

Homogeneity... Measures the closeness of distribution of pixel pairs to the diagonal elements of the GLDM.

Angular Second Moment... Represents the uniformity or homogeneity of the image.

Entropy... Measures the randomness or uncertainty of the pixel pair distribution.

Correlation... Measures the linear dependency between gray levels of neighboring pixels.

NOTE Refer to https://pyradiomics.readthedocs.io/en/latest/features.html#module-radiomics.gldm for more information about GLDM.

Gray Level Run Length Matrix (GLRLM) Features

The Gray Level Run Length Matrix (GLRLM) is a radiomics method used characterize the lengths and frequencies of consecutive runs of pixels with the same gray-level value in an image. The GLRLM is constructed by counting the number of occurrences of consecutive runs of pixels with specific gray-level values and their corresponding lengths in a chosen direction. GLRLM features include:

Short Run Emphasis (SRE)... Measures the distribution of short runs, indicating the presence of fine textures.

Long Run Emphasis (LRE)... Represents the distribution of long runs, indicating the presence of coarse textures.

Gray-Level Nonuniformity (GLN)... Quantifies the variation in gray-level values within the runs, reflecting the heterogeneity of texture.

Run Length Nonuniformity (RLN)... Measures the variation in the lengths of the runs, indicating the variability of texture patterns.

Run Percentage (RP)... Represents the proportion of the image covered by runs of a specific gray-level value.

NOTE Refer to https://pyradiomics.readthedocs.io/en/latest/features.html#module-radiomics.glrlm for more information about GLRLM.

Gray Level Size Zone Matrix (GLSZM) Features

The Gray Level Size Zone Matrix (GLSZM) is a radiomics method used to extract texture-based features from images by characterizing the size and spatial distribution of connected regions with the same gray-level value. The GLSZM is constructed by counting the number of occurrences of connected regions with specific sizes and their corresponding gray-level values. Once the GLSZM is constructed, various statistical measures can be derived from it to describe the texture properties of the image.

Small Area Emphasis (SAE)... Emphasizes the contribution of small, connected regions to the texture analysis.

Large Area Emphasis (LAE)... Emphasizes the contribution of large, connected regions to the texture analysis.

Gray-Level Nonuniformity (GLN)... Quantifies the variation in gray-level values within the connected regions.

Zone Percentage (ZP)... Represents the proportion of the image covered by connected regions of a specific gray-level value and size.

NOTE Refer to https://pyradiomics.readthedocs.io/en/latest/features.html#module-radiomics.glszm for more information about GSZM.

Neighbouring Gray Tone Difference Matrix (NGTDM) Features

The Neighbouring Gray Tone Difference Matrix (NGTDM) characterizes the local variations in gray-level values between neighboring pixels in an image. To compute NGTDM features, a specific distance or neighborhood is defined, and the gray-level differences between the central pixel and its neighbors within this distance are calculated. The NGTDM is constructed by counting the number of occurrences of specific gray-level differences and their corresponding frequencies.

Once the NGTDM is constructed, various statistical measures can be derived from it to describe the texture properties of the image. Common NGTDM features include:

Contrast... Measures the local variations in gray-level differences between the central pixel and its neighbors.

Complexity... Quantifies the complexity or randomness of the gray-level differences.

Strength... Represents the overall magnitude or strength of the gray-level differences.

Busyness... Reflects the level of local variations or activity in the image.

NOTE Refer to https://pyradiomics.readthedocs.io/en/latest/features.html#module-radiomics.ngtdm for more information about NGTDM.

Longest Length

In some cases, you may want to compute the longest length of each labeled object in a multi-ROI. For example, as part of the analysis of a fibers dataset that was segmented with the new Open Fiber Segmentation plugin (see Open Fiber Segmentation on page 35). This option is available in the Compute Measurements dialog for multi-ROIs, as shown in the following screenshot.

Compute Measurements dialog

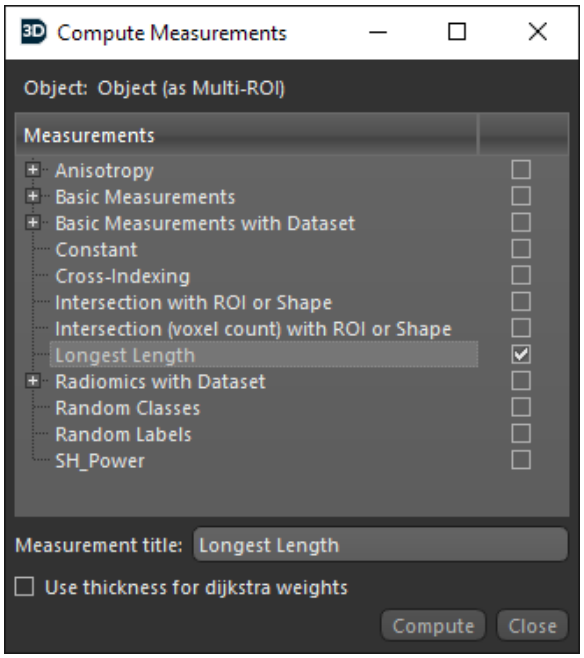

NOTE If required, you can select the 'Use thickness for Dijkstra weights' option for the computation.

When the computation is complete, select '**Longest Length**' in Measurement drop-down menu in the Scalar information box, as shown below, to review the measurements.

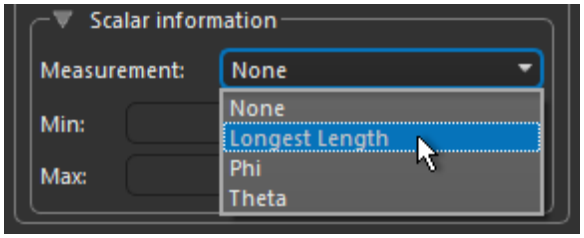

Phi and Theta Computations

This software release includes the option to compute Phi and Theta in both the world coordinate and image coordinate systems, as well as to compute Phi and Theta of the minimum, middle, and maximum Eigen vectors.

Results of Phi computations in image and world coordinates

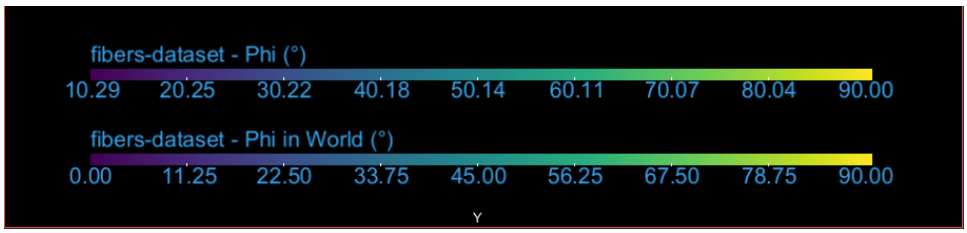

NOTE The option to compute Phi and Theta in the world coordinate system lets you directly compare the Phi and Theta of datasets that are spatially oriented differently.

These new computations are available in the Basic Measurements section of the Compute Measurements dialog, as shown in the following screenshot.

Compute Measurements dialog

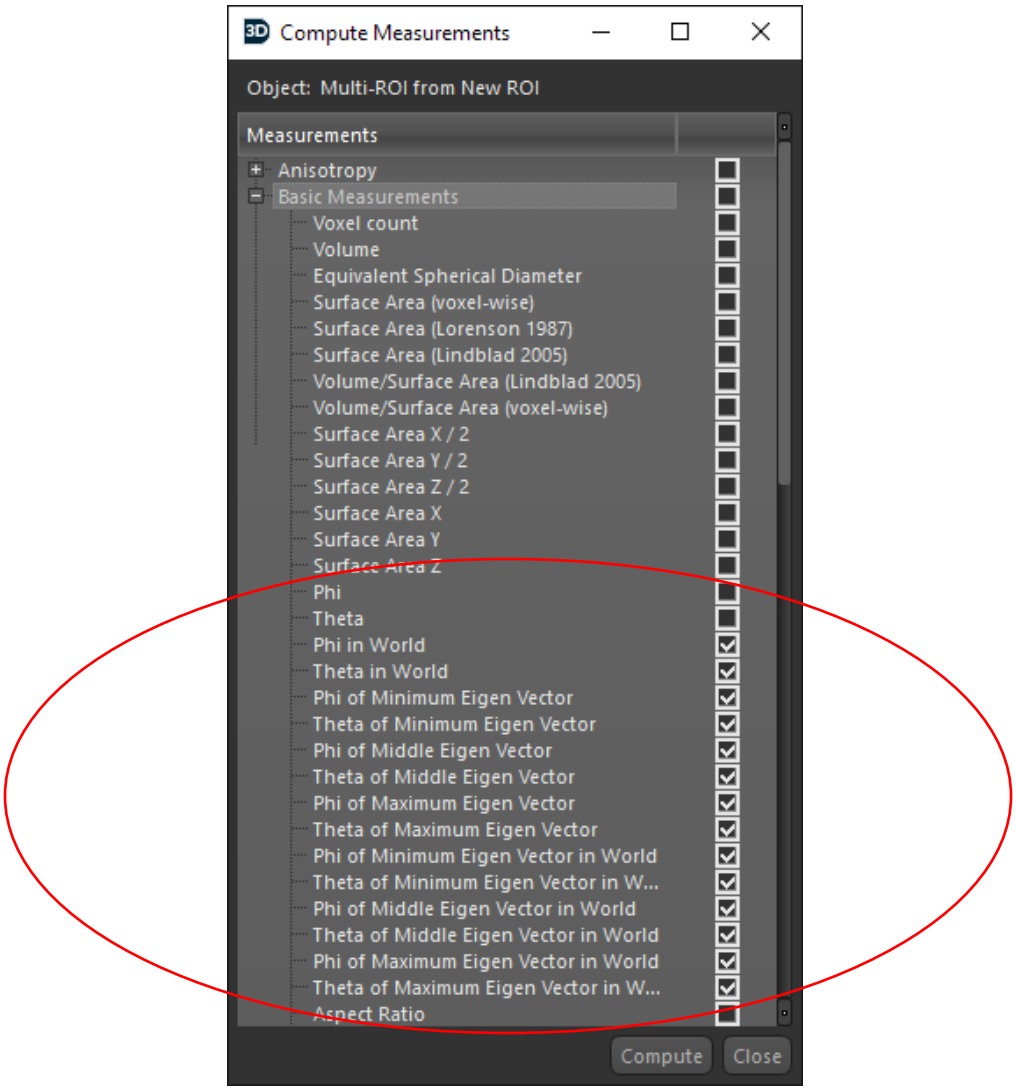

NOTE	Computations of Phi and Theta are equivalent to computations of Phi of Minimum Eigen Vector and Theta of Minimum Eigen.

Analyze and Classify Measurements

This software release includes the production version of Dragonfly 3D World's **Analyze and Classify Measurements** module, which lets you import scalar data from multi-ROIs, meshes, graphs, and vector fields for the cross-table analysis of feature vectors. You can apply arithmetic operations to the data, add cross-reference measurements, plot histograms to gain insights into the distribution and characteristics of the data, categorize objects according to different criteria, calculate statistics or ratios based on classifications, and more.

Right-click a multi-ROI, mesh, graph, or vector field that contains scalar values and then choose **Analyze and Classify Measurements** in the pop-up menu to open the Analyze and Classify Measurements dialog, shown below.

Analyze and Classify Measurements dialog

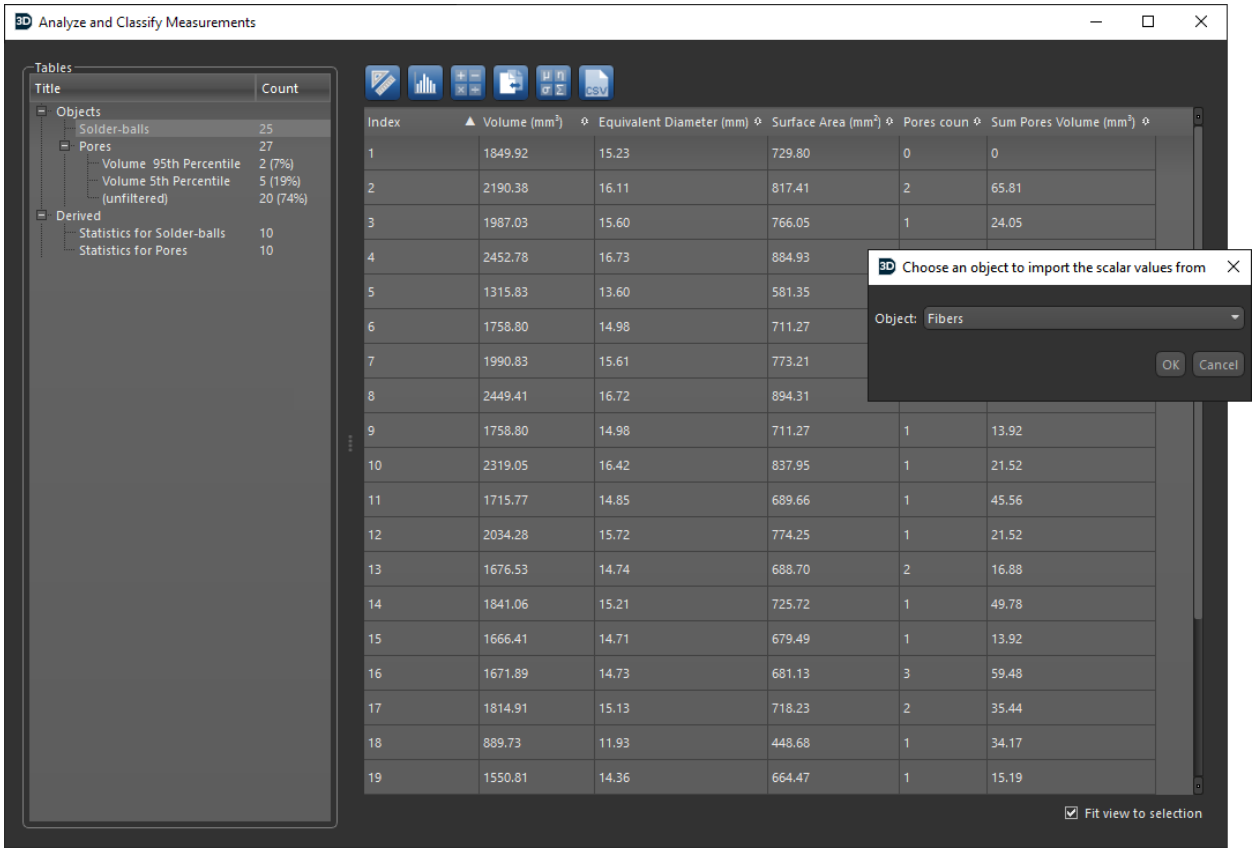

You can choose the object(s) you need to analyze in the Choose an Object to Import Scalar Values From dialog. You can add data from additional objects to the dialog by right-clicking 'Objects' on the left sidebar and then choosing **Add** in the pop-up menu.

Tables panel… The Tables panel lists all the tables that were created by importing the scalar data of objects and all derived tables in a tree menu. In addition, you can apply

filtering to create sub-groups that match specific criteria. For example, data can be filtered within a range, by percentile, or by logical expressions such as less than, equal to, greater than, and so on.

Operations... The tools at the top of the dialog let you to compute new measurements, start a histogram analysis, apply arithmetic operations, add cross-reference measurements, compute basic statistics, and export data in the CSV file format.

Operations

Operation	Icon	Description
New Measurement		Opens the Compute Measurements dialog, shown below, in which you can compute new measurements for the object(s) that you are analyzing. 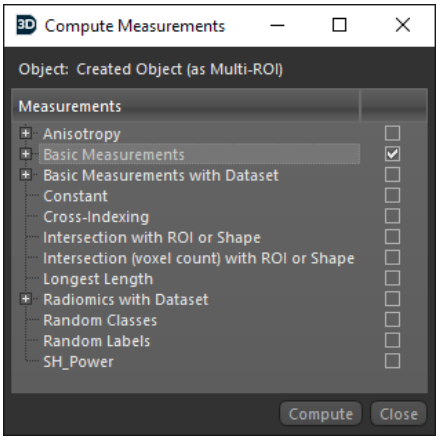 See Compute Measurements on page 37 for information about computing measurements for multi-ROIs, meshes, graphs, and vector fields.
Histogram Analysis		Opens the Histogram Analysis dialog, in which you can plot and analyze the distribution of scalar data contained in an object, compute basic statistics, as well as create classes by selecting instances in the data that match some criteria. See Histogram Analysis on page 50 for information about updates to the dialog.
Arithmetic Operations		Opens the Arithmetic Operations dialog, shown below, in which you can apply arithmetic expressions to one or more inputs to modify measurements or to compute new ones. 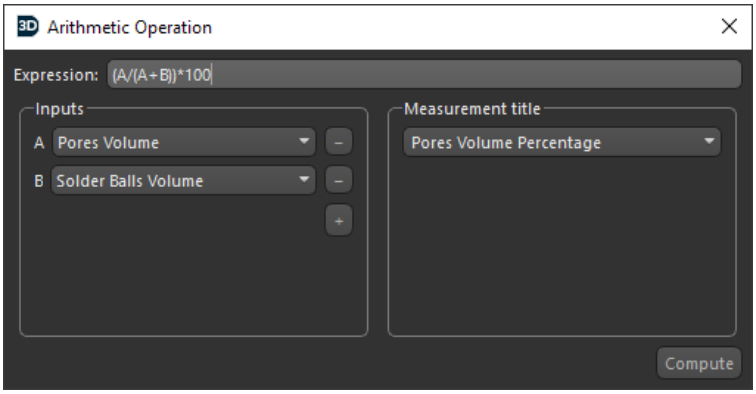 **Note** The convention for identifying each input within an expression is simply 'A' for the first input, 'B' for the second input, 'C' for the third, and so on. For example, to get the sum of the first input and the second input, the expression would be 'A+B'. More complex

Operation	Icon	Description
		expressions can also be written, such as ‘D*[(A+B)+(2*A-1.8*B)]/C’. You can also include NumPy functionalities in an expression, such as ‘np.sin’ and ‘np.cos’.
Cross-Reference Measurement		Opens the Cross-Reference Measurement dialog, shown in the following screenshot. You can add cross-reference measurements to a table in this dialog. 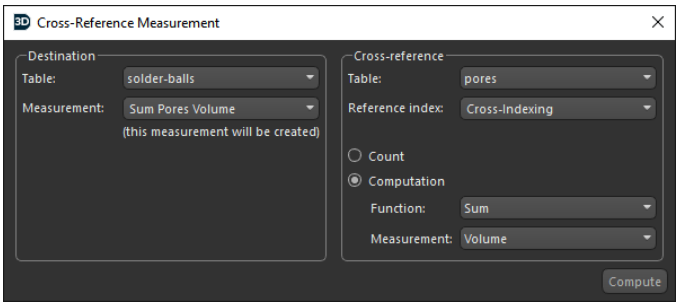 As shown in the following example, the measurements of 'Pores Count' and 'Pores Volume' were added to a Solder-balls table from the pores labels that were cross-referenced to the solder-balls. 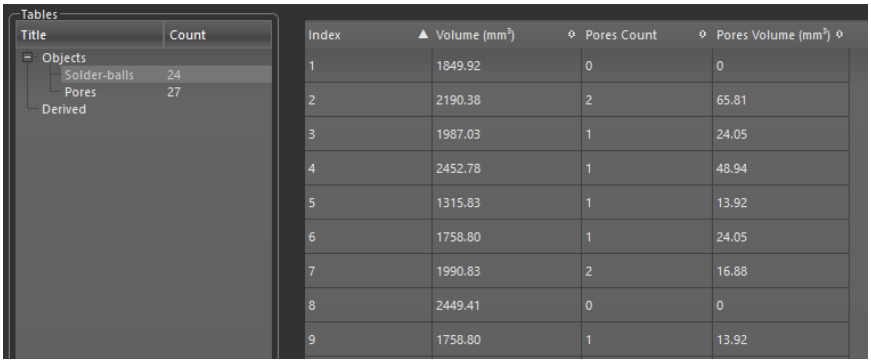 Tables Title / Count Objects Solder-balls 24 Pores 27 Derived Index / Volume (mm³) / Pores Count / Pores Volume (mm³) 1 1849.92 0 0 2 2190.38 2 65.81 3 1987.03 1 24.05 4 2452.78 1 48.94 5 1315.83 1 13.92 6 1758.80 1 24.05 7 1990.83 2 16.88 8 2449.41 0 0 9 1758.80 1 13.92
Basic Statistics		Opens the Basic Statistics dialog, in which you can compute basic statistics, such as the mean, median, minimum, maximum, and standard deviation of computed measurements. These statistics can be added to the original table as metadata or as a new derived table. These output options are available in the Basic Statistics dialog, as shown below. 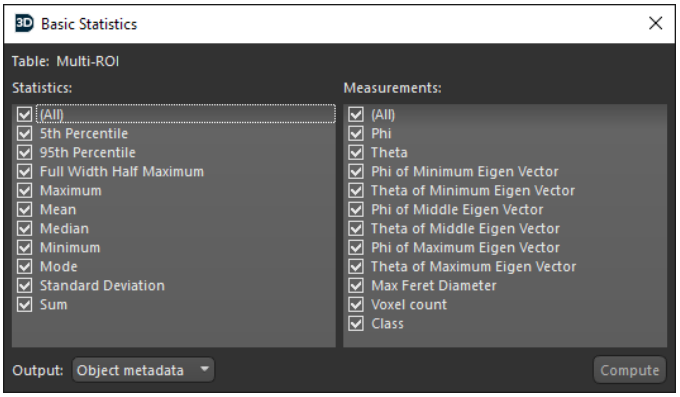
Export as CSV		Exports the current table or selected columns in the comma-separated values (*.csv extension) file format.

Tables… All imported, filtered, and derived scalar data is organized in tables, with columns representing features and rows representing components or instances within

the scalar data, such as the labeled objects in a multi-ROI, the faces or vertices of a mesh or graph, or the vectors of a vector field. Features can be length, area, volume, angle, circularity, gray-level intensity value, or anything else that is useful for the application at hand, and every row can be interpreted as a feature vector.

Histogram Analysis

Updates for Histogram Analysis include the option to plot 2D histograms using Cartesian or polar coordinates, as well as a Classification tab on which you can 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.

Click the **Histogram** tool on the Data Properties and Settings panel or click the **Histogram Analysis** button on the Analyze and Classify Measurements dialog to open the Histogram Analysis dialog, as shown in the following screen captures.

Histogram Analysis dialog

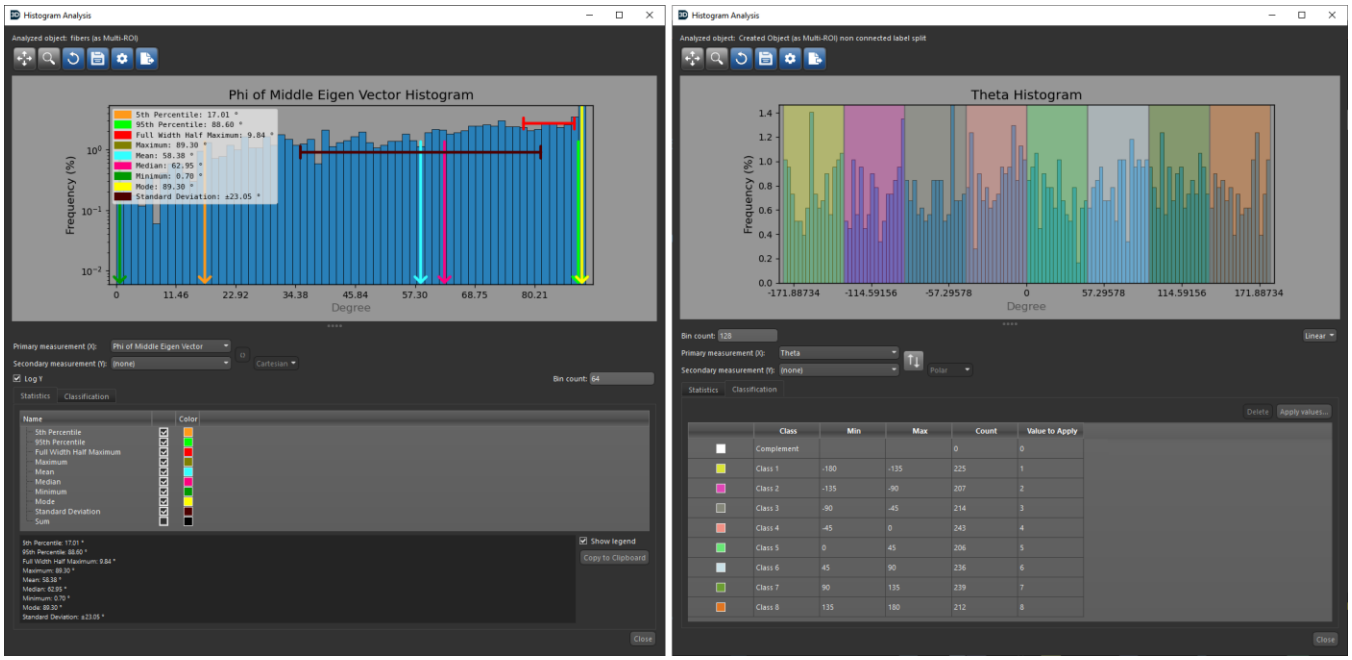

Meshes

New options are available in this software release for working with meshes. These include a 3D Mesh Painter, with which you can edit a mesh's scalar values or create a

new scalar slot for painting, the option to smooth meshes with a constraint, adjustable thickness for outlined meshes, and mesh capping on all faces.

3D Mesh Painter

For some 3D modeling applications, you may need to edit the scalar values of a mesh or create a new scalar slot for painting values. These options, which are now available in Dragonfly's 3D Mesh Painter, let you:

- Set boundary conditions for simulations, such as for solving heat equations with NGSolve (see NGSolve Integration on page 33).
- Correct any inaccurate values contained within a scalar mesh.
- Define where an operation is applied, such as mesh smoothing with a constraint (see Mesh Smoothing with Constraint on page 52).

Choose **Utilities** > **3D Mesh Painter** to open the 3D Mesh Painter panel, shown below.

3D Mesh Painter panel and painted mesh

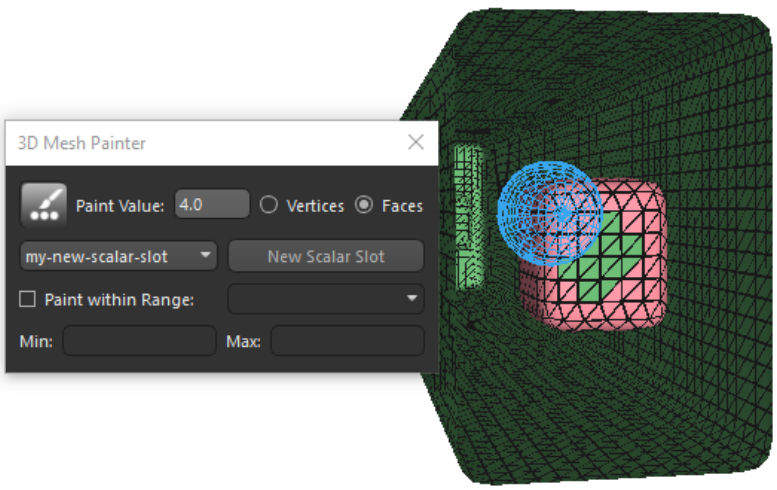

You can do the following with the 3D Mesh Painter:

- Paint on vertices or faces.
- Paint on an existing scalar slot or create a new slot for painting.
- Paint within the range of another scalar slot. For example, when constraining painting to the boundary faces of a solid mesh. You should note that the initial value set for boundary faces is -1.

The following configured actions are available for working with the 3D Mesh Painter.

Configured actions for the 3D Mesh Painter

Action	Default Setting

Action	Default Setting
Decrease brush size (in 3D Mesh Painter)	Mouse wheel down
Increase brush size (in 3D Mesh Painter)	Mouse wheel up
Paint vertex or face (in 3D Mesh Painter)	Left Ctrl + Left mouse

Mesh Smoothing with Constraint

This software release includes the opportunity to smooth meshes using a selected scalar slot as the per-vertices Laplacian relaxation parameter, as well as a selected degree of smoothing. As shown in the example below, smoothing with a constraint lets you limit smoothing to selected areas of a mesh.

Original mesh (left) and the original mesh smoothed with constraint (right)

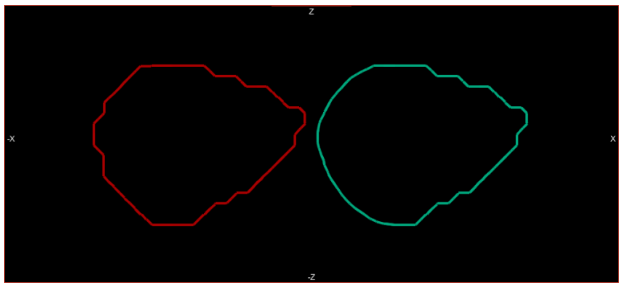

NOTE The scalar slot selected for smoothing with constraint should have vertex values between 0 (no relaxation) and less than 2. You should note that a value to 1 corresponds to 'normal' smoothing. In most cases, it is best to create a scalar slot specifically for smoothing with the 3D Mesh Painter (refer to the topic 3D Mesh Painter on page 51 for information about mesh painting).

Select the required mesh and then click **Apply with Constraint** in the Smooth mesh box on the Data Properties and Settings panel.

Smooth mesh box

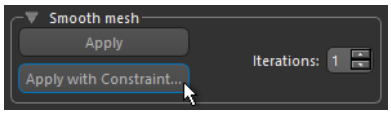

You can then choose the required scalar slot in the Choose a Scalar Values Slot dialog.

Choose a Scalar Values Slot dialog

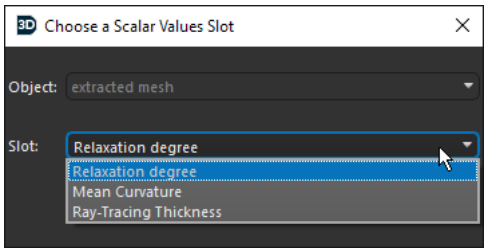

Remeshing with Poisson Surface Reconstruction

This release includes an implementation of a Poisson surface reconstruction algorithm for reconstructing triangle meshes from sets of oriented 3D points. This utility (*vtkPoissonReconstruction*) was adapted from the original implementation by Kazhdan, Bolitho, and Hugues (2006). You should note that Poisson surface reconstruction offers several advantages for remeshing — it is a global solution that considers all the data at once without resorting to heuristic partitioning or blending, and it can create very smooth surfaces that robustly approximate noisy data. An example of remeshing with Poisson surface reconstruction is shown below.

Original mesh (on left) and remeshed with Poisson surface reconstruction (on right)

Right-click the mesh you want to reconstruct and then choose **Remesh with Poisson Surface Reconstruction** to open the Mesh Smoothing dialog, shown below. You can choose the parameters for creating the new mesh in the dialog.

Mesh Smoothing dialog

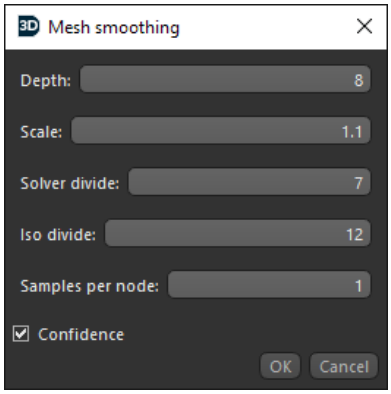

NOTE Refer to the publication Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe, *Poisson Surface Reconstruction*, Eurographics Symposium on Geometry Processing (2006), Konrad Polthier, Alla Sheffer (Editors) for additional information about Poisson surface reconstruction. This publication is available at: https://hhoppe.com/poissonrecon.pdf.

Mesh Capping

In previous software versions, it was only possible to cap clipped meshes on the active Clip box face. In this software version, mesh capping is applied to all clipping faces of the Clip box whenever **Cap mesh on clipping** is selected as a 3D setting.

Cap mesh on clipping

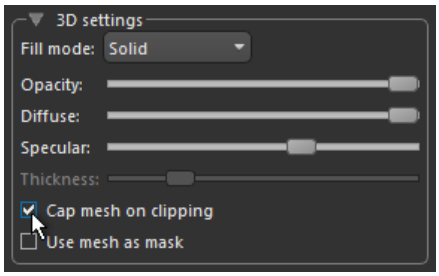

In the example below, a thickness mesh was clipped on two faces and then capped.

Capping applied to a clipped mesh

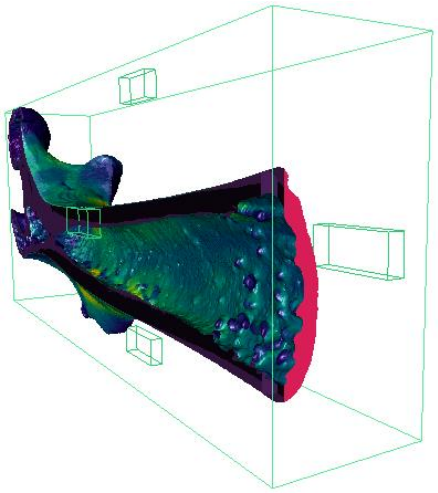

Adjustable Thickness for Outlined Meshes

The Thickness setting in the 3D settings box on the Data Properties and Settings panel, shown below, is now applicable to meshes shown in **Outlined** mode. In previous versions, this setting was only available for meshes shown in Wireframe mode.

Thickness slider

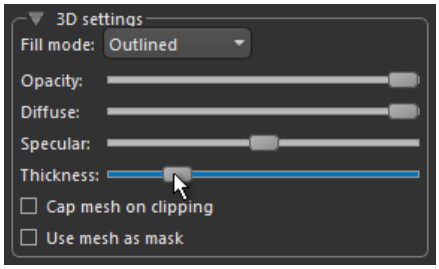

ROI Painter

This software release introduces 'color modes' for labeling regions of interest and multi-ROIs with the ROI Painter tools. Examples of the two available color modes — *Single-color mode* and *Two-color mode* — are shown in the screen captures below.

ROI Painter color modes

Single-color mode… In this mode, which appears whenever one ROI or one multi-ROI class is selected, you can use the Brush tools to label any voxel by holding down *Left Ctrl* (or your configured 'Add with key'). Holding down *Left Shift* (or your configured 'Remove with key') lets you remove voxels from the selected ROI or multi-ROI classes. You should note that in the case of multi-ROIs, this a change from previous versions. Previously, voxels were only removed from the selected class when painting in Single-color mode. In version 2024.1, voxels will be removed from all classes.

Two-color mode… In this mode, which appears whenever multiple ROIs or multiple multi-ROI classes are selected, only the labeled voxels belonging to the selected ROIs (if *Exclusive mode* is selected) or to the selected classes of a multi-ROI will be affected when painting. Holding down *Left Ctrl* (or your configured 'Add with key') will relabel voxels from 'B' and the shown intermediary colors to 'A'. Holding down *Left Shift* (or your

configured 'Remove with key') will relabel voxels from 'A' and the shown intermediary colors to 'B'. The identification of selected ROIs and selected multi-ROI classes as A, B, and C is shown in the following screenshots.

Multiple ROIs selected (left) and multiple multi-ROI classes selected (right)

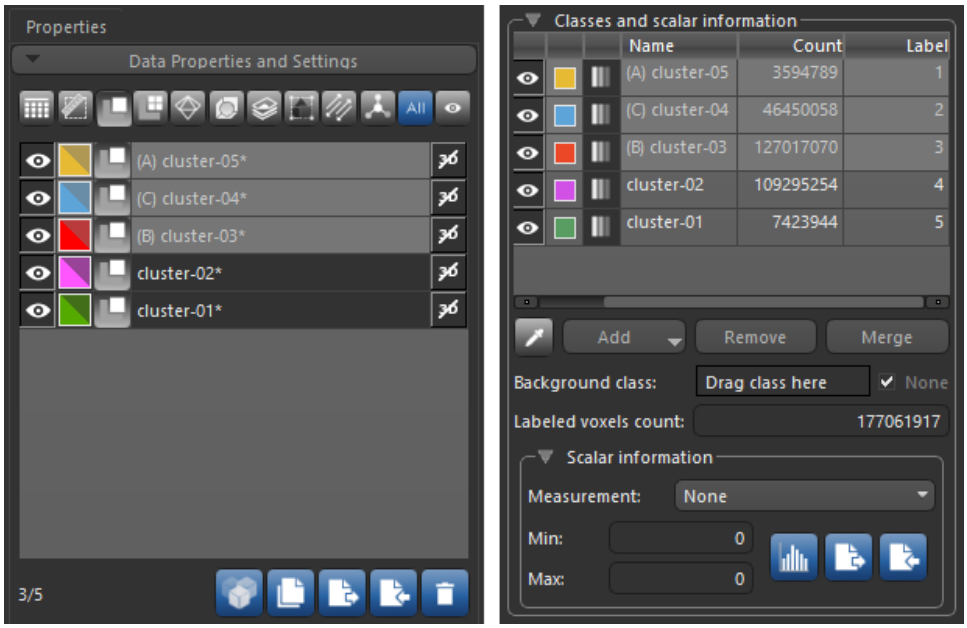

Two-color mode can be particularly useful whenever the results of an automated segmentation may need manual editing to assign labeled voxels from one or more multi-ROI classes to another class or from one ROI to another. As shown below, only the labeled voxels highlighted in yellow are re-assigned to those highlighted in blue when painting in Two-color mode.

Two-color mode painting

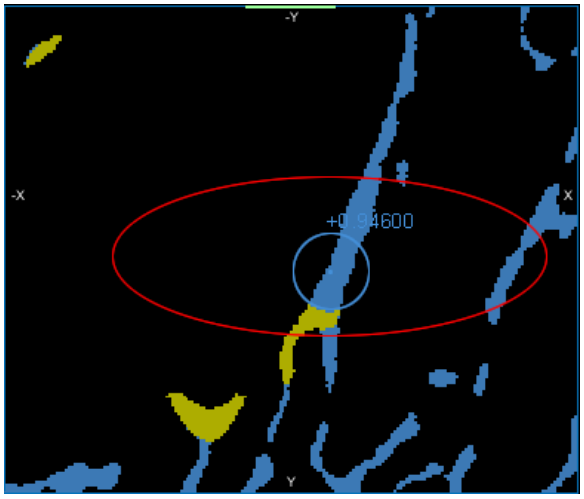

ROI Painting Preferences

A new set of preferences is available for controlling painting behavior in 3D views. **Choose Views > 3D Settings** in the Preferences dialog to view the 'ROI Painting in 3D' preferences, shown below.

3D Settings preferences

Draw only after a stroke is finished… If selected, updates to the region of interest will only appear after the brush is released.

Continuously draw while painting… If selected, updates to the region of interest will appear in real time.

Thresholding

In this software release, the selected range on the ROI Tools panel defines not only the minimum and maximum threshold values, but now also serves as a proxy for calculating the lower and upper Otsu ranges. This means that the Otsu value will be computed within the selected range whenever the **Lower Otsu** or **Upper Otsu** button is clicked. In previous versions, the Otsu value was computed from the full data range, regardless of the currently selected range.

Range box

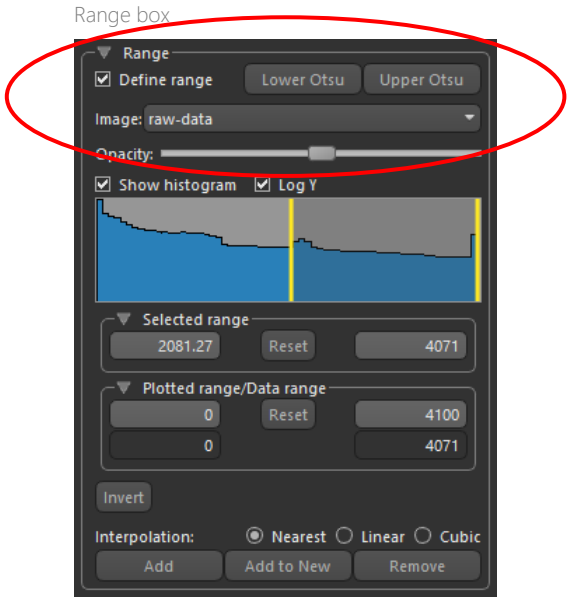

3D Rendering

This software release includes multiple light sources and material presets to make your 3D renderings even more sensational. In addition, optimizations implemented in version 2024.1 are sure to result in improved performance and a reduced memory footprint when rendering multiple volumes, particularly for users who typically deploy **correlative workflows** and work with **large datasets**.

Multiple Lights

Dragonfly 3D World includes multiple lights for achieving a higher degree of photo realism and better shadow detail for 3D renderings. No longer limited to one-point lighting setups, you can now setup a combination of light sources to either draw attention to a special part of the scene or to represent the natural properties of a scene.

As shown in the example below, multiple light sources can help create a sense of depth, contrast, and color to enhance the shape, texture, and details of 3D renderings.

Multiple lights applied to a scene with a virtual floor

Multiple lights are available on the Lighting panel, as shown below.

Lighting panel

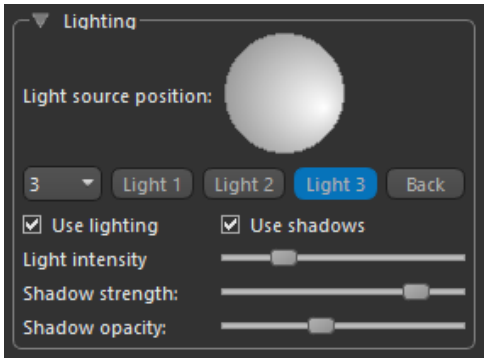

You should note that the lighting source position and light intensity are particular to each light, which means that each can have its own source position and intensity setting. However, you should also note that shadow strength and shadow opacity are universal. Changing these settings for one light will reset the setting(s) for all lights.

NOTE You can switch a light source position to 'Back', as shown below, to apply back lighting with one or more of the selected lights.

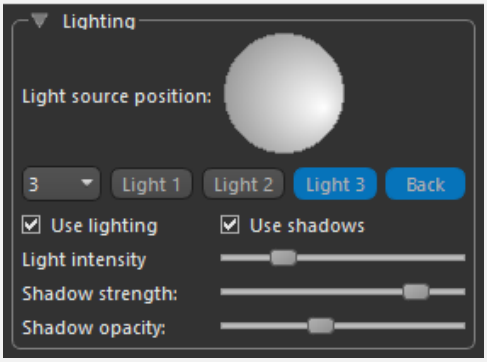

Lighting and Shadows

For this software release, the options 'Use lighting' and 'Use shadows' options were separated in the Lighting panel. This lets you apply lighting without any shadows for a specific purpose and can also help with 3D rendering in cases in which you are using a low-end GPU that might struggle to render shadows.

Dataset without lighting (left), with lighting (center), and with lighting and shadows (right)

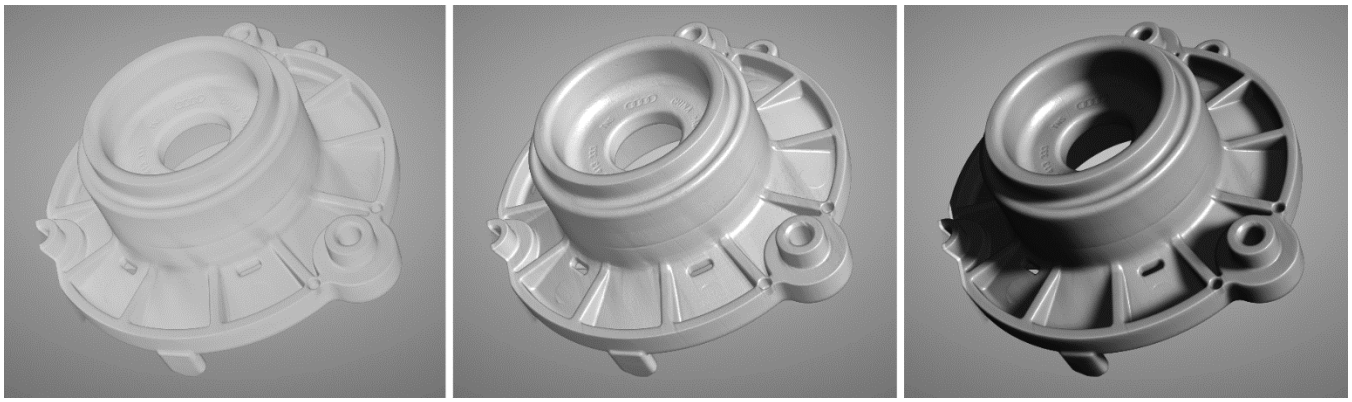

Several additional improvements for shadows are also included in this software release. These include:

- Applying shadows in correlative workflows, as shown below.

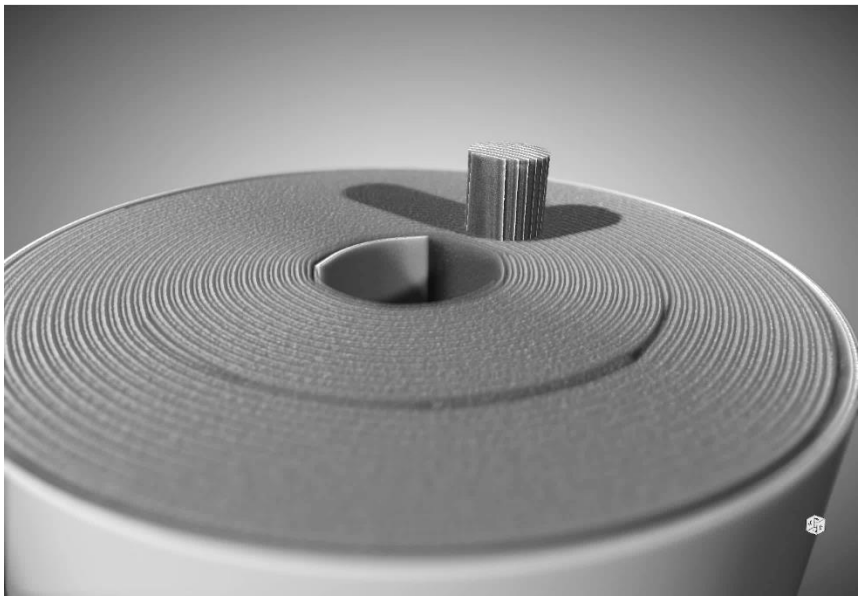

- All visuals can now cast shadows, for example, volumes on meshes, graphs on datasets, and so on.
- Using shadows with reflective floors is now more accurate.

Tone Mapping

Careful mapping of high dynamic range (HDR) values to low dynamic range (LDR) values can be an important part of a 3D rendering pipeline. In this software release, this is now done using ACES (Academy Color Encoding System) filmic tone mapping to avoid saturating shadows and highlights. As shown in the example below, tone mapping can help you reveal the full details of your renderings and give them a hyper realistic look.

Original rendering (left) and with tone mapping applied (right)

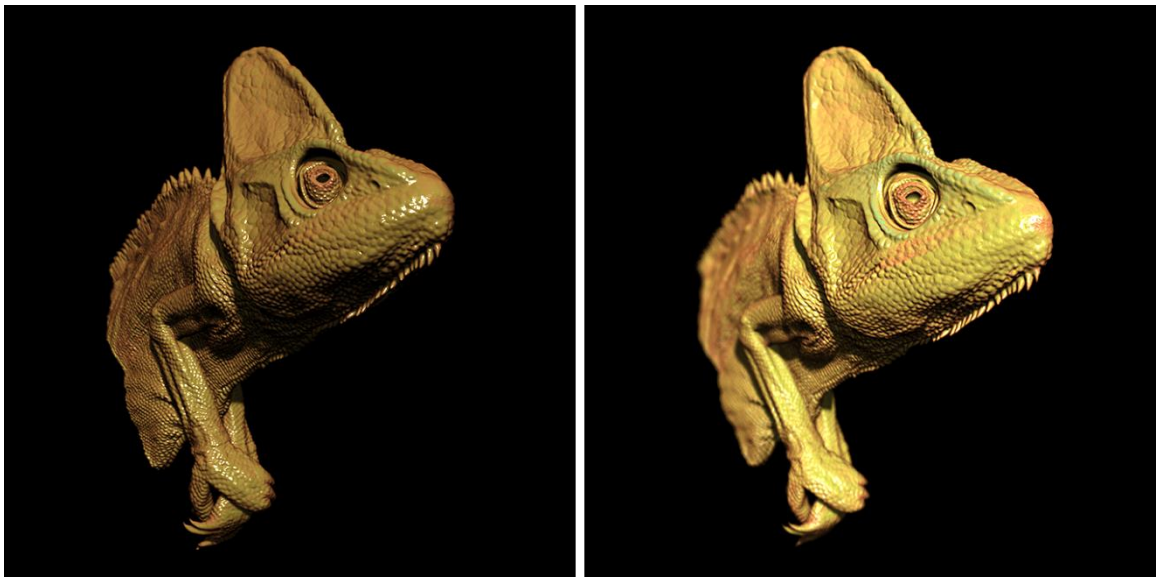

NOTE Tone-mapping consists of adjusting the tonal values of an image with a high dynamic range by scaling the dynamic range down while attempting to preserve the full details of the original image.

The option to apply tone mapping is available in the 3D settings box on the Data Properties and Settings panel, as shown below.

3D settings box

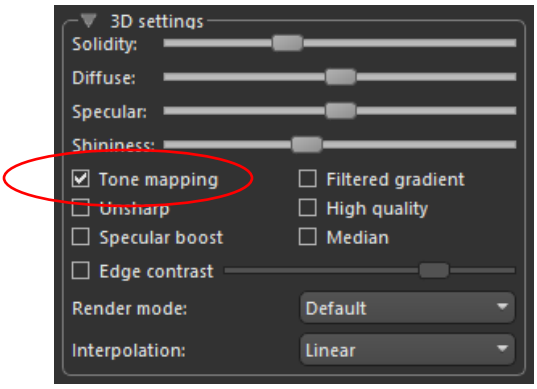

Suppression of Moiré Effects

This software release includes improved filtering to remove moiré effects from both low resolution and high-resolution datasets. This 2D filter has no cost to apply and should work well with most graphic cards.

Below are examples of the suppression of moiré effects on low-resolution and high-resolution datasets. Images on the left are from previous versions of Dragonfly, while images on the right are from the 2024.1 release.

Suppression of moiré effects on a low-resolution dataset

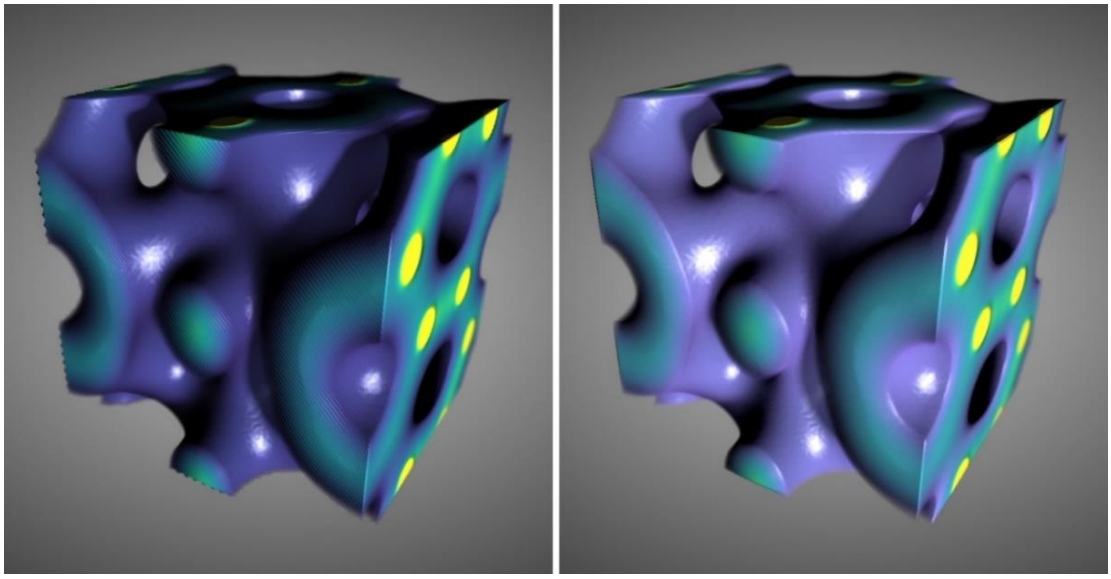

Suppression of moiré effects on a high-resolution dataset

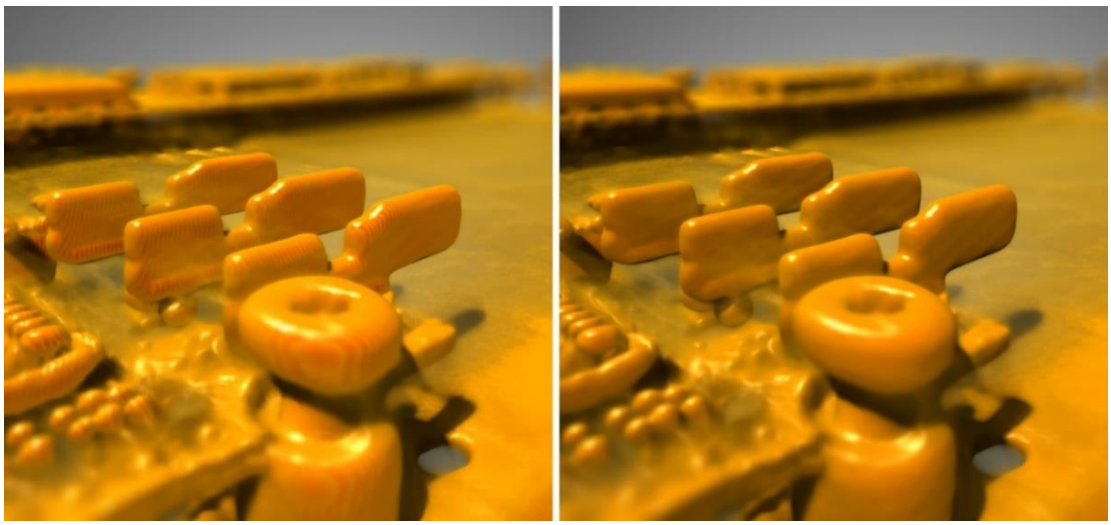

Material Presets

This software release includes phase functions so that you can assign a material, such as aluminum, bronze, cobalt, copper, concrete, and gold, to a dataset or region of interest.

Copper preset applied to a copper lattice (left) and multiple material presets applied to a series of ROIs (right)

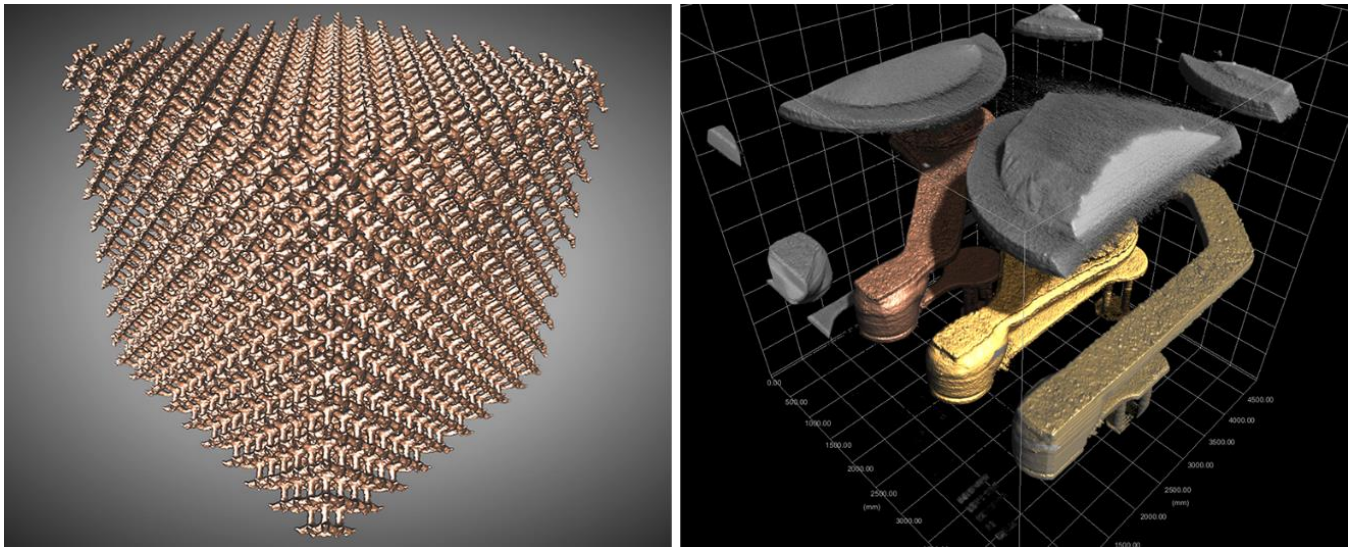

Material settings for images — Aluminum, Bronze, Cobalt, Concrete, Copper, Gold, Plaster, and Sand —are available in the **Presets** drop-down menu in the 3D presets box, circled below. The option to apply material settings to regions of interest is available in the Opacity and color box.

3D presets for images (left) and Opacity and color settings for ROIs (right)

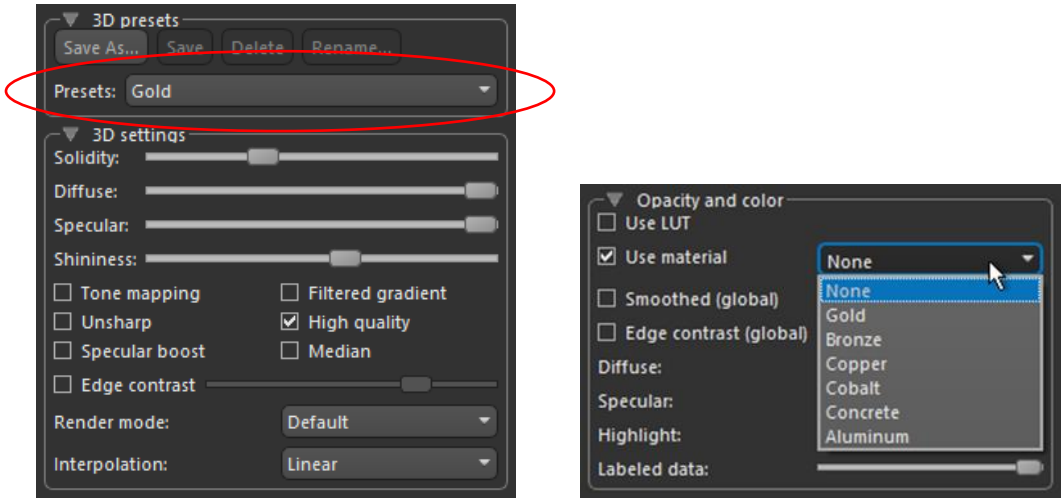

NOTE Additional Presets available in version 2024.1 include Bones ES, Lights x2, Monkey, Reptile, Sand 3, and Sand 4. You should also note that some previously available presets have been updated.

Settings for ROIs and Multi-ROIs

New and updated settings — 'Smoothed (global)' and 'Edge contrast (global)' — are available in this software release for visualizing regions of interest and multi-ROIs in 3D. These settings are available in the Opacity and color box, as shown below.

Opacity and color box

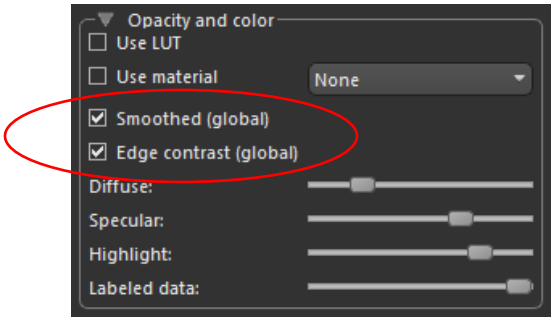

Smoothed (global)... This global setting was updated for version 2024.1 to use reflected light and an increased kernel size for improved real-time smoothing.

Original ROI (left) and smoothed ROI (right)

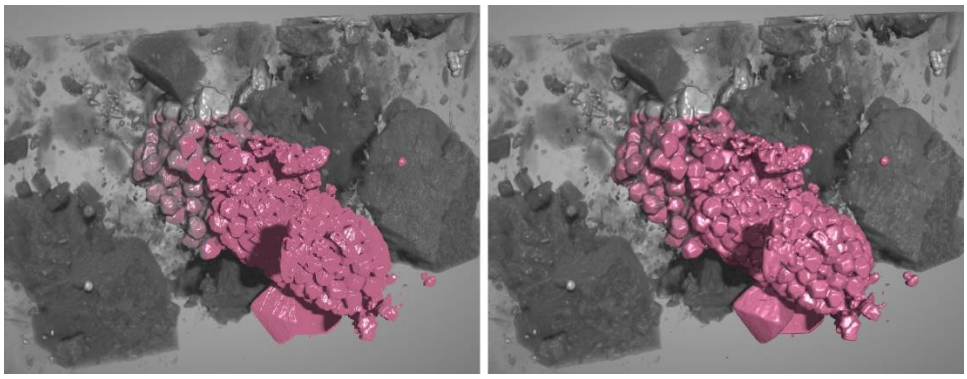

Edge contrast (global)... This global setting produces an image in which the edges are emphasized, and other features are rendered semi-transparent.

Original ROIs (on left) and with edge contrast applied (on right)

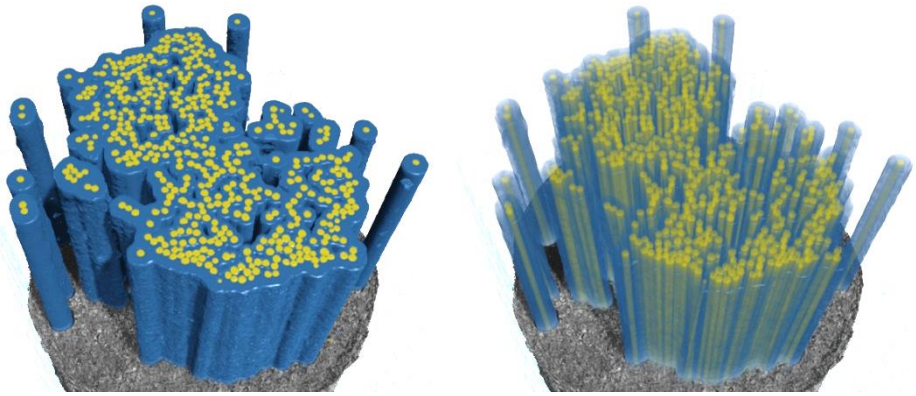

Planes

Exploded views and additional settings are available in this software release for working with planes.

Exploded Views

The new explosion options for visual planes lets you sectionalize 3D renderings to reveal the interior of a sample. You should note that you can include exploded views in animated sequences and generate high-resolution images with this visual effect. An exploded 3D rendering is shown in the illustrations below.

Dataset exploded by 'Translation' on the left and by 'Rotation (screw)' on the right

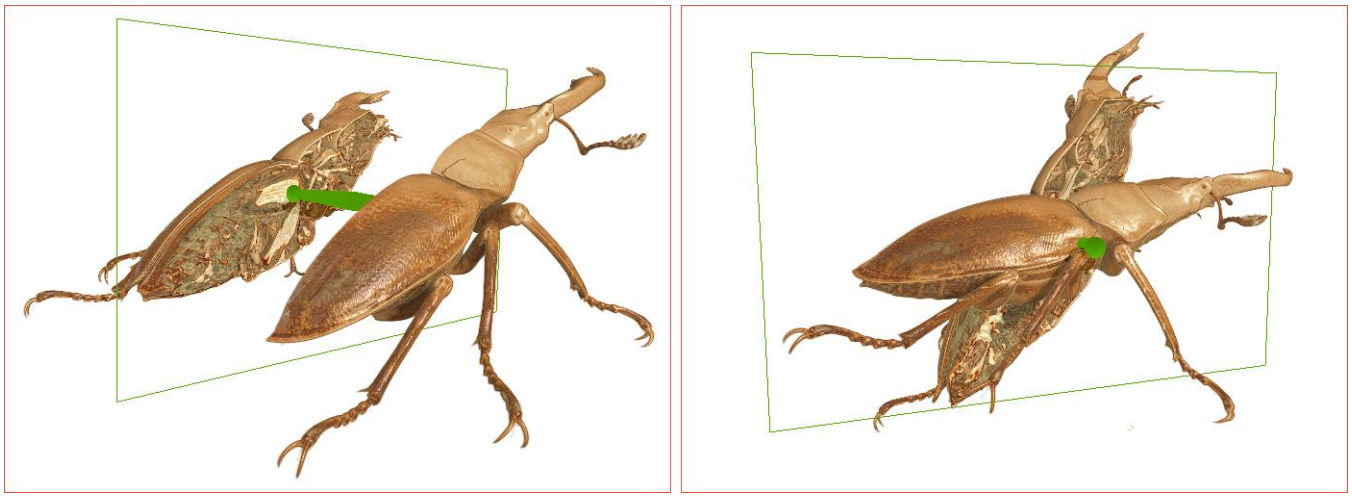

The Explosion panel, shown below, is available for selected planes in the Data Properties and Settings panel. You can choose the type of explosion — **Translation**, **Rotation (hinge)**, or **Rotation (screw)** — in the drop-down menu, as well as the degree of translation or rotation with the Explosion slider. You can also invert rotation explosions.

Plane settings

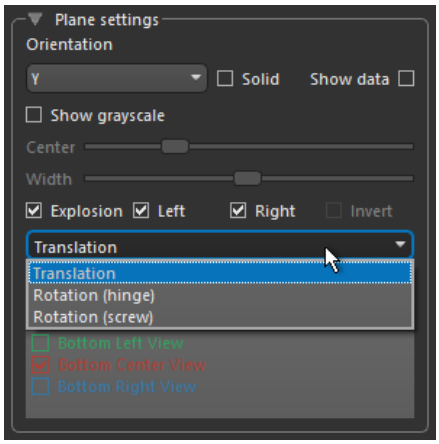

Show Grayscale

This software release includes a new 'Show grayscale' method that can be applied to clipped planes. In some cases, this may provide better results than the 'Show data' method that is also available in previous and current versions of Dragonfly.

A comparison of the 'Show data' and 'Show grayscale' options is illustrated below.

Plane with Show data option applied (left) and with Show grayscale (right)

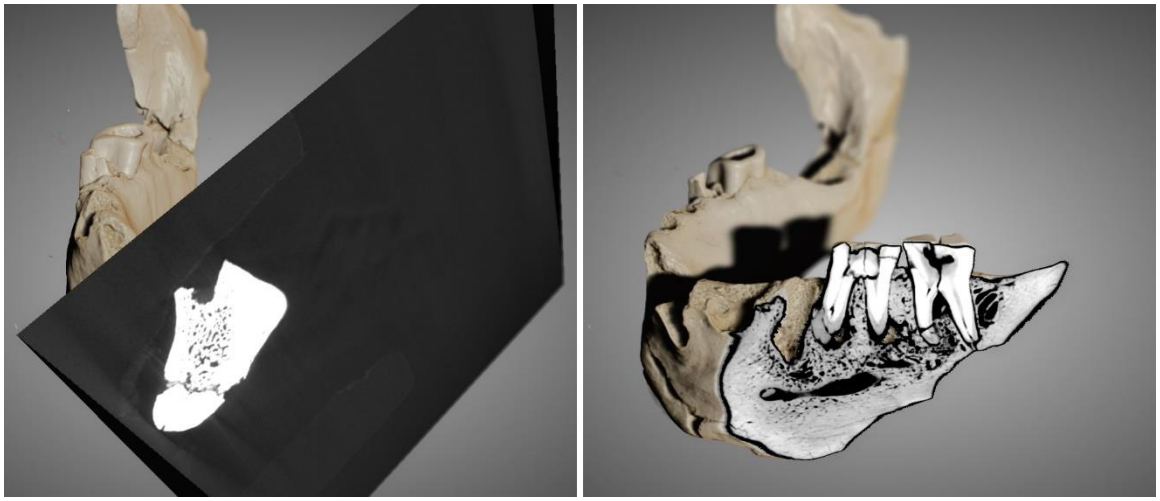

The 'Show grayscale' option and associated window center and width controls are circled below.

Plane settings box

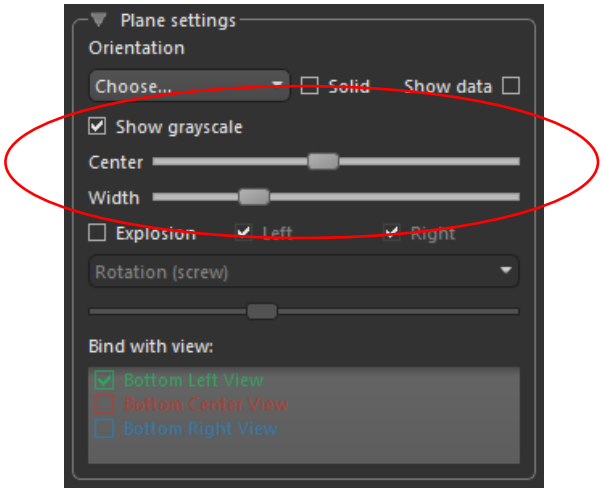

You can adjust the center and width values independently of the window leveling values applied to the dataset.

3D View Controls

New Axis indicator modes and controls are available in version 2024.1 for manipulating 3D views.

Automatic View Alignment

This software release includes the option to automatically align 3D views with the X, Y, or Z axis by double-clicking on the Axis indicator, as shown below.

Axis indicator

NOTE You can also align 3D views by double-clicking the Axis indicator whenever it is shown in Arrow, Gyroscope World, or Gyroscope View mode.

Constrain Pan Movements

In some cases, you may want to constrain pan movements to a single axis. For example, when creating an animated sequence. This software version includes a new set of configured actions that let you constrain pan movements to the X, Y, or Z axis.

After changing the axis indicator to 'Arrows', you can then activate any of the directional arrows and then pan within a constrained direction.

Pan movement constrained to the X axis

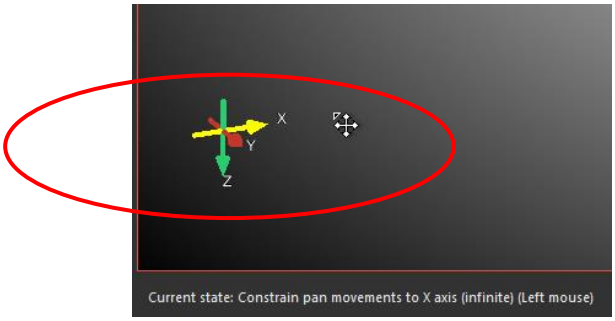

The following configured actions and their default settings are available for constraining pan movements.

Configured actions for constraining Pan movements

Action	Default Setting
Constrain Pan movements to X axis (Infinite mode)	Left mouse
Constrain Pan movements to Y axis (Infinite mode)	Left mouse
Constrain Pan movements to Z axis (Infinite mode)	Left mouse

Gyroscope Axis Indicators

Additional Axis indicator modes — 'Gyroscope World' and 'Gyroscope View' — are available in this software release. You can choose an Axis indicator mode in the 3D view contextual menu, as shown below.

3D view contextual menu

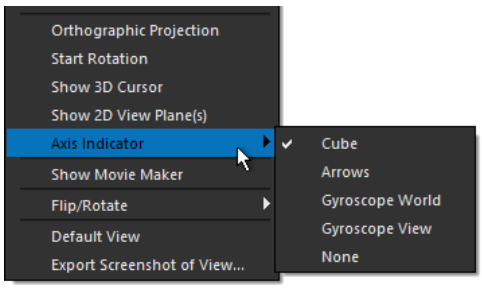

You can constrain Track movements to a selected axis, as shown below. You should note that when an axis is selected, it turns yellow.

Track movement constrained to rotation in the X axis

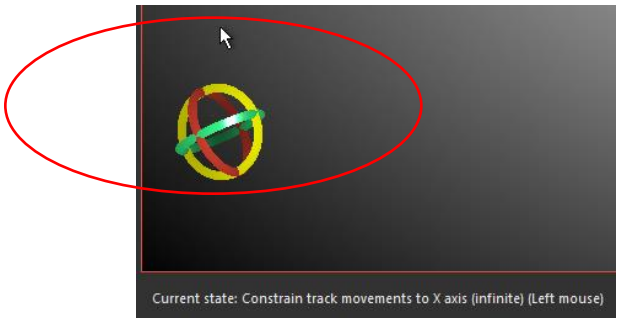

The following configured actions and their default settings are available for constraining Track movements.

Configured actions for constraining Track movements

Action	Default Setting
Constrain Track movements to X axis (Infinite mode)	Left mouse
Constrain Track movements to Y axis (Infinite mode)	Left mouse
Constrain Track movements to Z axis (Infinite mode)	Left mouse

New preferences for the Axis Indicator

A set of new 3D preferences, shown below, are now available for selecting 'Gyroscope (World)' or 'Gyroscope (View)' as the default appearance of the Axis indicator.

CT Reconstruction

An additional pre-processing filter is available in this software release for reconstructing cone-beam and parallel-beam projections in the CT Reconstruction module. A variant of adaptive histogram equalization, contrast limited adaptive histogram equalization (CLAHE) is now available for improving local contrast and enhancing the definition of edges in each region of reconstructed images.

Choose **Workflows** > **CT Reconstruction** on the menu bar to open the CT Reconstruction dialog. The CLAHE pre-processing option is circled below.

Pre-processing options in the CT Reconstruction dialog

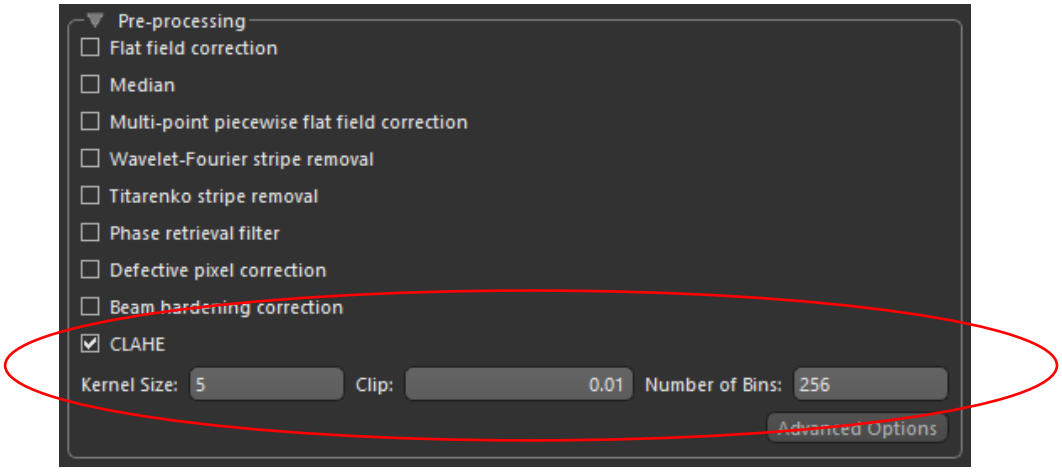

Organizer

You should note that the Organizer now boasts enhanced compatibility with Dragonfly enterprise solutions, such as VizServer, thanks to changes to the backend. The backend is now managed by Neo4j, which is a graph database management system developed by Neo4j, Inc. Go to https://neo4j.com/ to learn more about this administrator tool.

Image Registration

As an update for image registration, this software release provides an initial score for rigid registrations, which is computed prior to applying any registration. As shown below, the starting point for a rigid registration task is indicated in the Registration information box.

Image Registration dialog

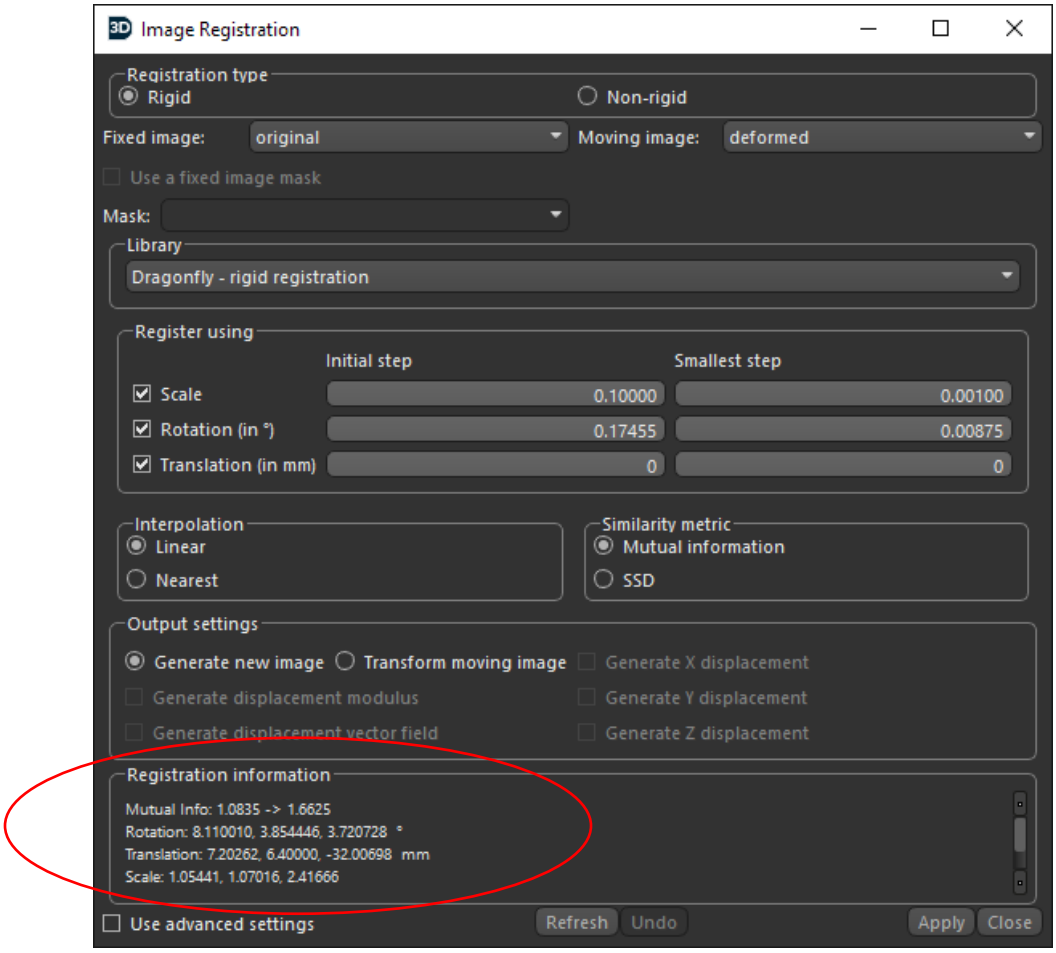

Natural Language Search

This software release includes an implementation of *natural language search (NLS)*, which lets you query the Dragonfly technical documentation using everyday language. Instead of entering keywords in the Search bar you can pose a question using regular spoken language that describes the information you are looking for. For example, if you want to know how to compute a thickness mesh, a keyword-based search query might be 'thickness mesh compute'. But if you are searching using natural language, you might phrase the query as 'How can I compute a thickness mesh?'.

An example of a basic question entered in the Search bar is shown below.

Search bar

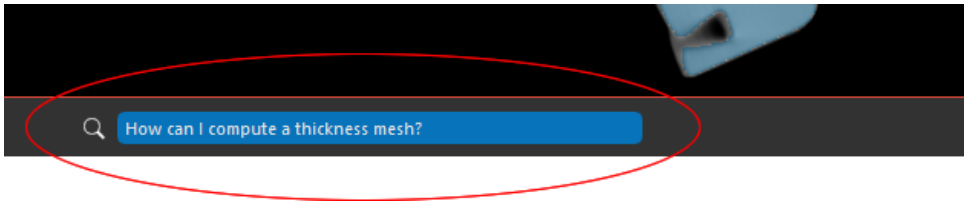

You should note that entered search queries get translated for the search engine using an advanced technique called *natural language processing* (NLP). Results are returned in a pop-up window with links to applicable topics, as shown below.

Search results

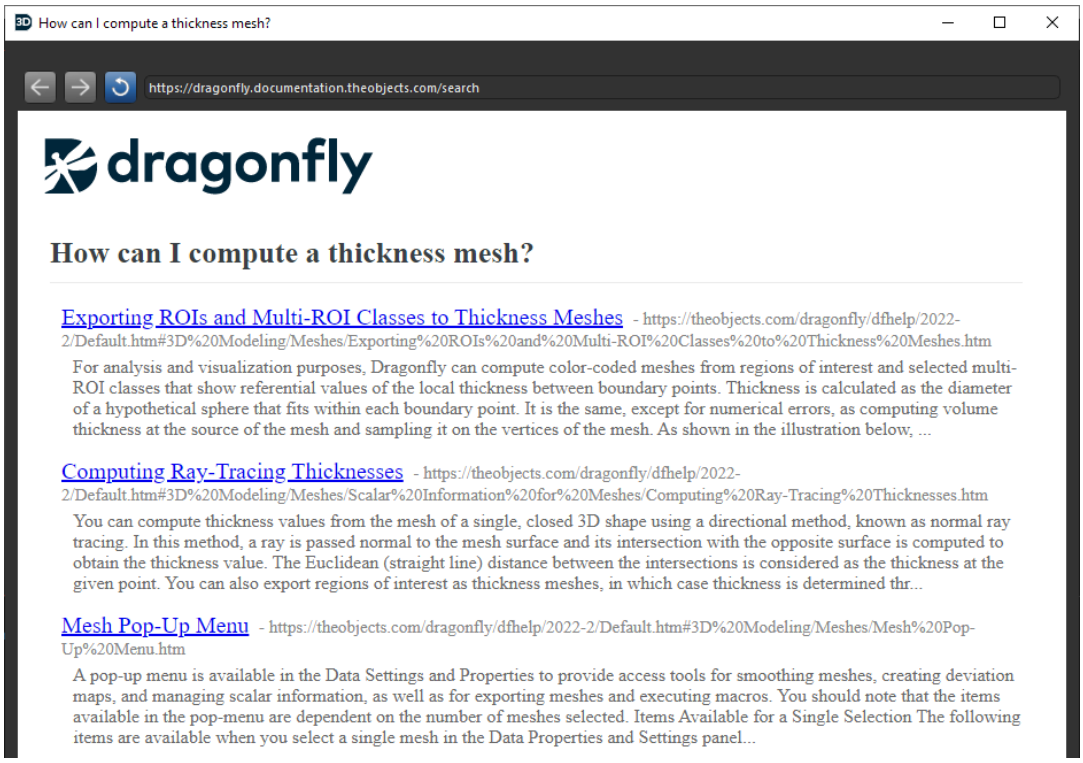

Importing RAW Data

To streamline workflows, you can now sequentially load RAW files without manually re-entering settings for size, data type, and image spacing. In other words, the Image Loader, shown below, will 'remember' your selections when you import RAW files of the same size.

Settings for importing RAW data

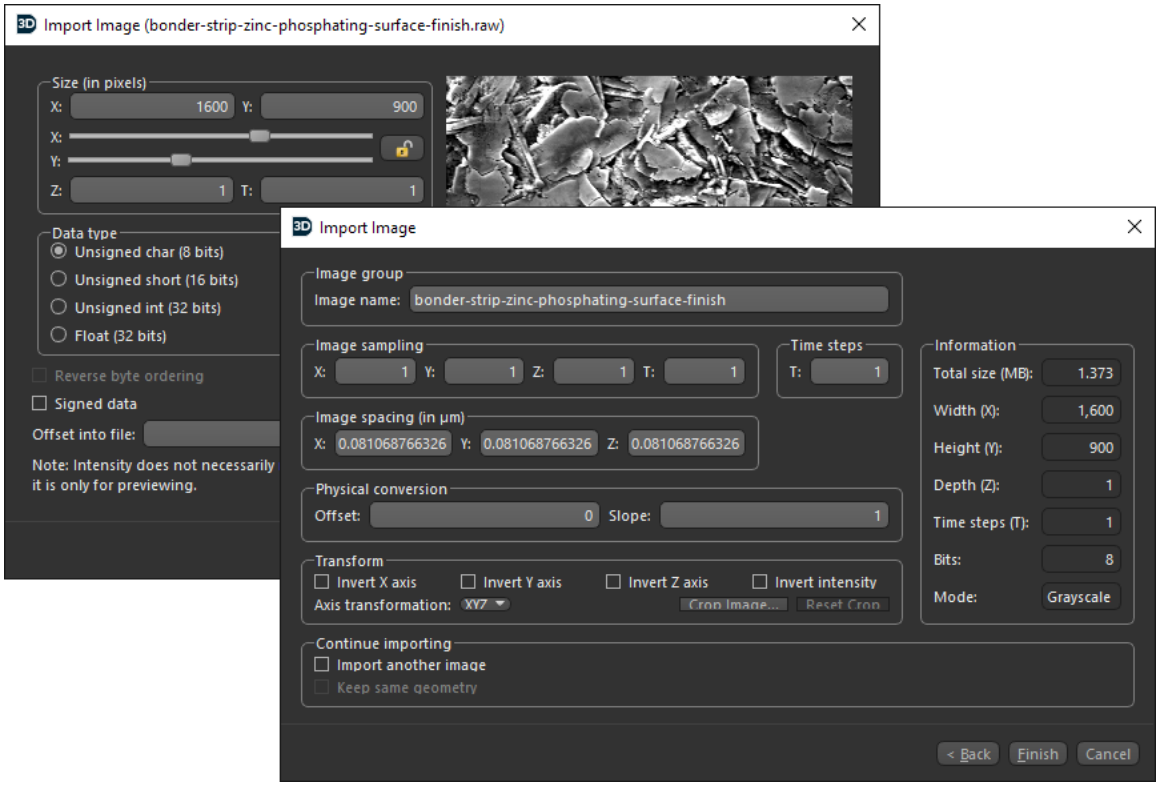

File Export

This software release provides the option to override the set **Miscellaneous** preference for compressing files — 'No compression' or 'zstd' — whenever an ORS Object or Session file is exported.

This option is available in the Export to File dialog, circled on the following screenshot, as well as in the Save Session File dialog.

Export to File dialog

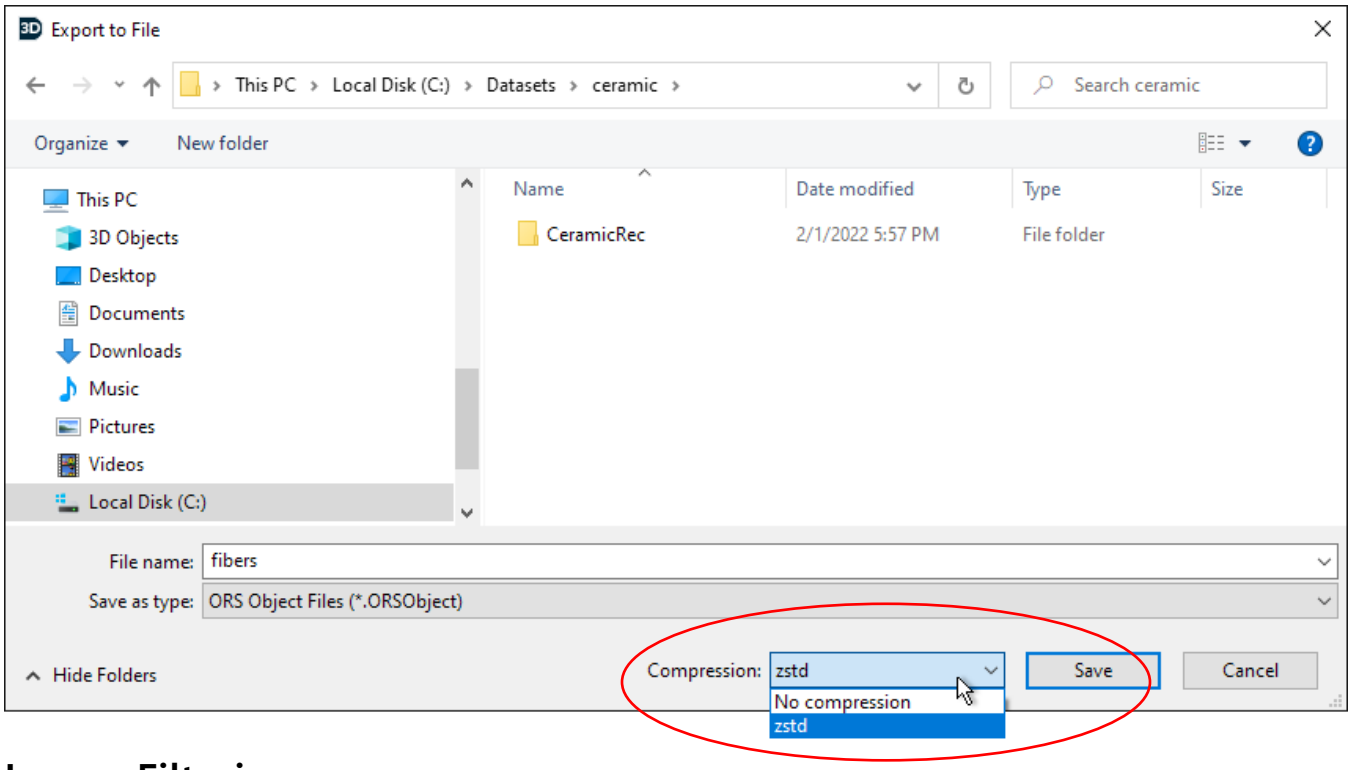

Image Filtering

You should note that some of the set parameter boundaries for image filters were removed in this software release. This provides more freedom to select values according to the nature of any dataset that needs to be processed.

As shown in the following example, the previous upper limit of '21' for the 'window_size' parameter for the Sauvola adaptive thresholding filter was removed.

Sauvola thresholding filter options

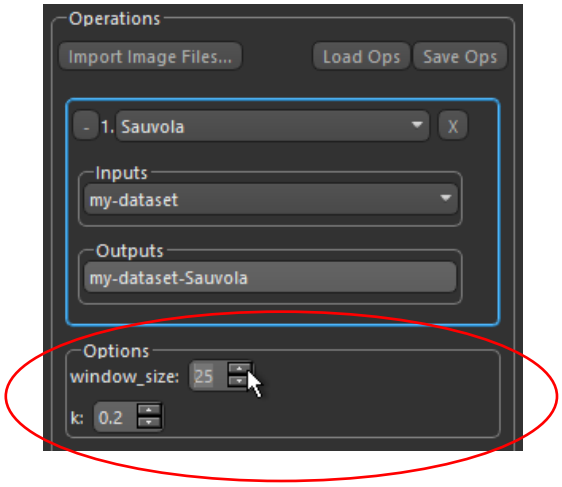

Developer Tools

A number of updates are available in this software release for the **Widget Observer**, which is a dedicated debugging tool for investigating the properties of objects that are part of the Dragonfly UI. New features include tools that let you dynamically highlight selections, select highlight colors and opacity for probed objects, extract information about spacers and layouts, as well as the option to probe elements within a graph.

Choose **Developer Tools** > **Widget Observer** to open the Graph Observer dialog. The new Highlight/Probe tools are circled below.

Graph Observer dialog

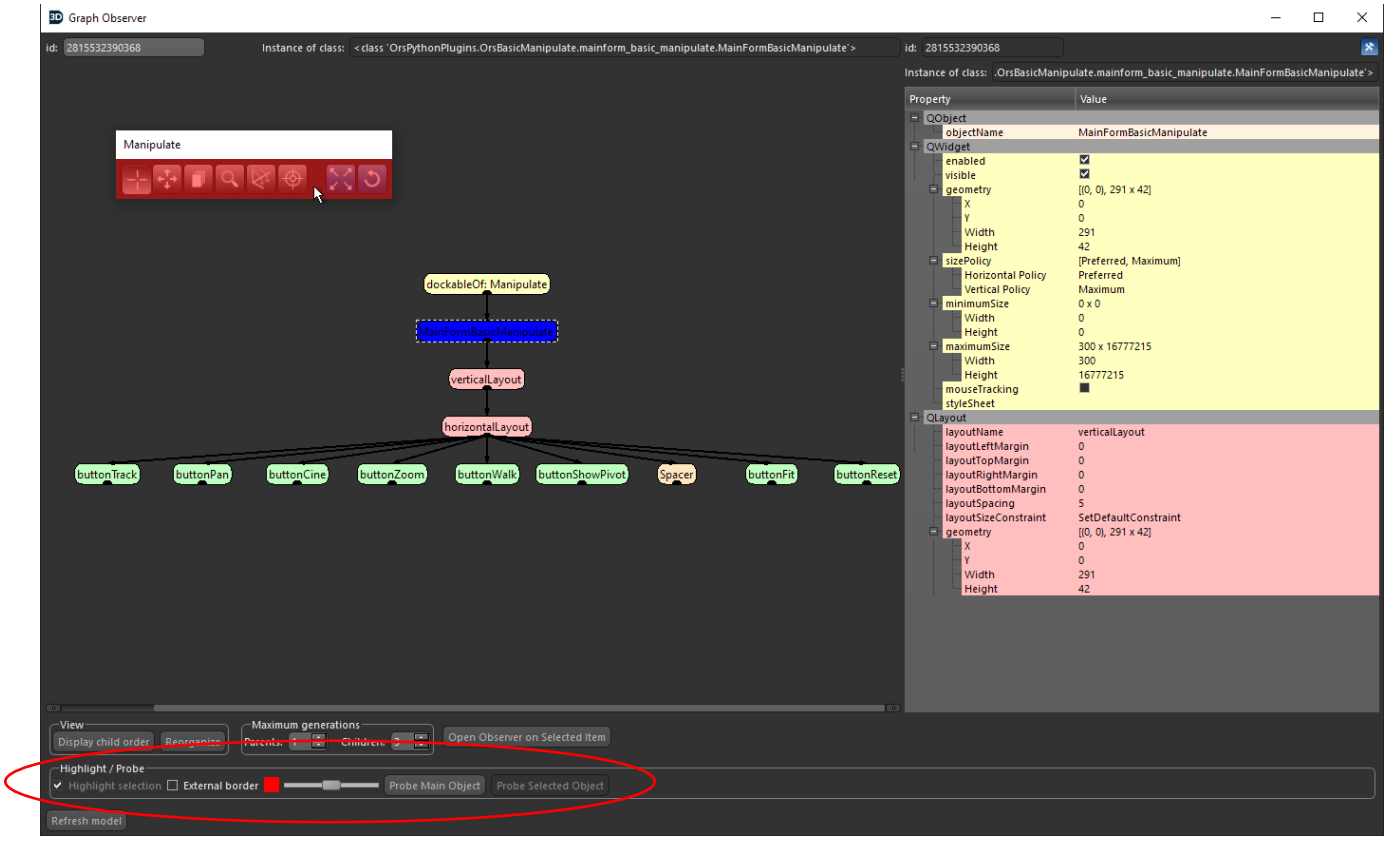

NOTE The Single Widget Observer dialog was deprecated in this software release.

You should also note that a Stylesheet Editor, as shown below, is now available for modifying the properties of a selected object.

Stylesheet Editor (on left)

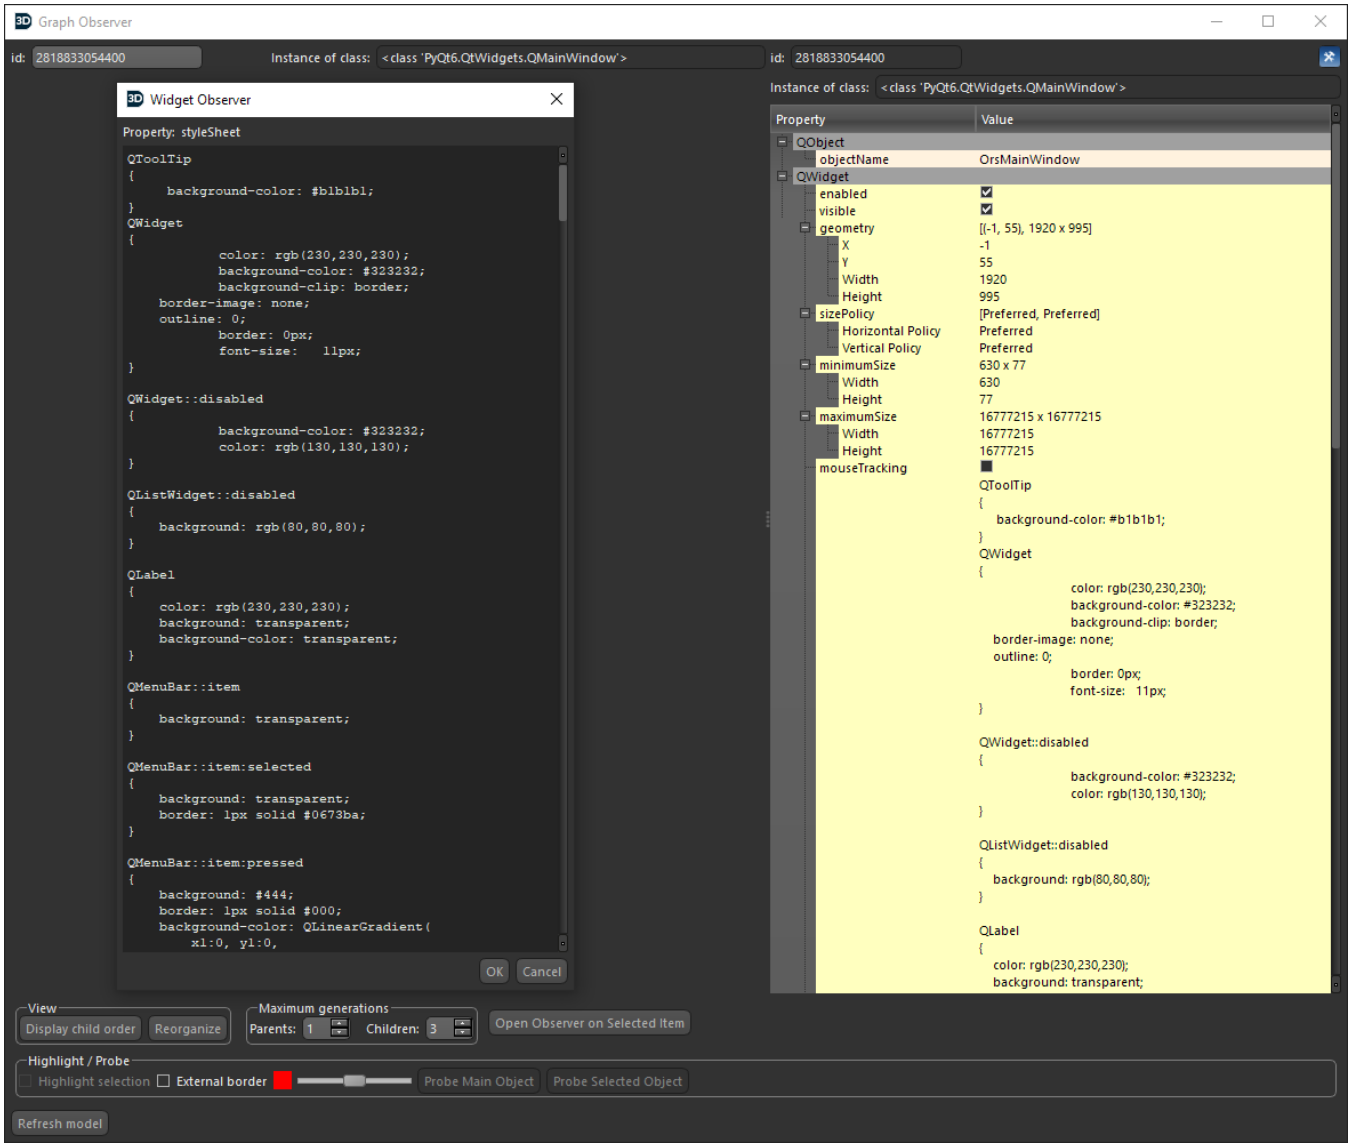

Tutorial Videos

You should note that our YouTube series of videos aimed at developers, which include tutorials about the Object Model Observer, Reimporter, Callback Observer, Context Observer, Docstring Analyzer, and others, was updated for this software release. You can view the full playlist of developer videos at:
www.youtube.com/playlist?list=PLbYyniU4wPOuPZmmjKrWYRT3g8wEmXEUz

Menu Bar Changes

The following new items are available on the menu bar in Dragonfly 2024.1. Other changes are also listed here.

Workflows Menu

The following changes, listed below, were implemented in the **Workflows** menu.

Workflows menu changes

Item	Description
Pore Network Modeling	This item is no longer available.
NGSolve	Opens the NGSolve Computation dialog, in which you can quantify fluid flow and heat dissipation in engineering and research applications (see NGSolve Integration on page 33).
Open Fiber Segmentation	Opens the Open Fiber Segmentation dialog, in which you can perform instance segmentation of the fibers in a region of interest (see Open Fiber Segmentation on page 35).

Utilities Menu

The following changes, listed below, were implemented in the **Utilities** menu.

Utilities menu changes

Item	Description
Analyze and Classify Measurements	Opens the Analyze and Classify Measurements dialog, in which you can import scalar data from multi-ROIs, meshes, graphs, and vector fields for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
Generate a Mesh From Spherical Harmonics	This item is no longer available.
3D Mesh Painter	Lets you edit the scalar values of a mesh (see 3D Mesh Painter on page 51).

Help Menu

The following changes, listed below, were implemented in the **Help** menu.

Help menu changes

Item	Description
Connect to Dragonfly Social As	Lets you connect to the online platform Dragonfly Social (see Dragonfly Social on page 9).
Manage Dragonfly Social Accounts	Lets you create and manage accounts to access Dragonfly Social (see Dragonfly Social on page 9).
Dragonfly Forums (Online)	This item and the online forums are no longer available. You can now use Dragonfly Social to ask questions, find answers, and share your experiences with other Dragonfly users (see Dragonfly Social on page 9).
Request Write Access to Forums	This item is no longer available.
Recover Forums Write Access	This item is no longer available.

Changes to Contextual Menus

The following changes for the contextual menus available in the Data Properties and Settings panel are implemented in Dragonfly 2024.1.

Dataset Contextual Menu

The items listed below are now available in the **Dataset** pop-up menu.

Dataset pop-up menu changes

Item	Description
Automated Porosity Segmentation	Provides automated methods — **Basic Otsu**, **Advanced Otsu**, and **Sobel Threshold** — to reliably segment porosity (see Automated Porosity Segmentation on page 37).
New ROI from Marked Slices	Creates a new region of interest in the same size and shape as the selected dataset and in which each marked slice is fully segmented.
Extract Slab Around Center	Extracts a specified number of rotated datasets (slabs) along the object's central axis. You can a select the number of slabs that will be extracted and the number of slices per slab in the dialogs shown below. 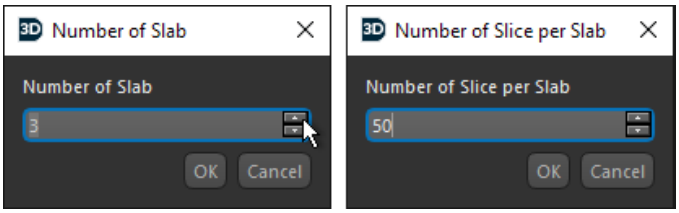 An example of 5 extracted slabs is shown below 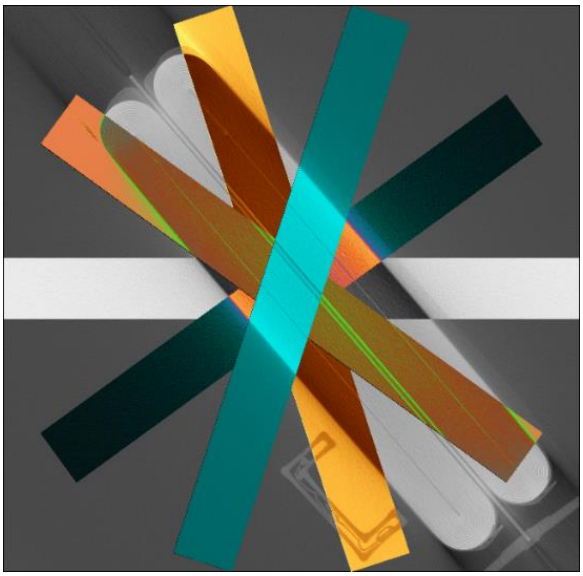
Generate a Multi-ROI from Multiple Otsu Splits	Lets you separate the selected image into several different classes in a multi-ROI, with each one obtained according to the intensity of the gray levels within the image. The number of classes, which determine how many thresholds are calculated, is selectable in

Item	Description
	the Classes dialog, as shown below. 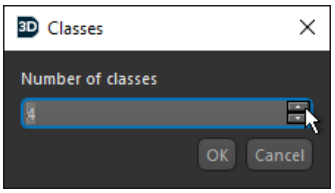 **NOTE** If the number of classes is four, then three threshold values will be calculated to generate the multi-ROI.
Create Local Maximum Map Image	Lets you create an image that shows the coordinates of local peaks. You can enter a threshold for computing the local maximum map in the Threshold dialog, shown below. 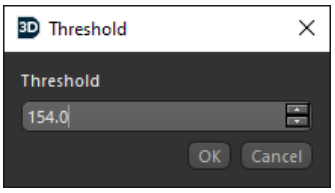
Create Local Maximum Map ROI	Lets you create a region of interest in which only local peaks are labeled. You can enter a threshold for computing the local maximum map in the Threshold dialog.
Create Ridge Image	Lets you create a ridge image.
Create Ridge ROI	Lets you create a region of interest in which detected ridges are labeled.
Open Image Quality Metrics	Now available for multi-slice datasets. Was previously only available for single-slice (2D) datasets.
Generate Phi Theta Orientation from Vector Field in X, Y, Z Images	Lets you generate images with scalar values of Phi and Theta orientations. You need to select at least two images to access this option and then choose the input datasets in the Choose Vector Field Component dialog, shown below. 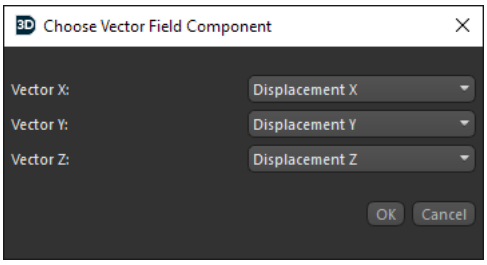 **NOTE** You can use PuMA to compute the required X, Y, and Z images by selecting 'Orientation detection' as the computation type on the Solver tab.

ROI Contextual Menu

The items listed below are now available in the **Region of Interest** pop-up menu.

Region of interest pop-up menu changes

Item	Description
Extract Slab Around Center	Extracts a specified number of rotated ROIs (slabs) along the object's central axis. You can a select the number of slabs that will be extracted and the number of slices per slab.

Multi-ROI Contextual Menu

The items listed below are now available in the **Multi-ROI** pop-up menu.

Multi-ROI pop-up menu changes

Item	Description
Analyze and Classify Measurements	Opens the Analyze and Classify Measurements dialog, in which you can import scalar data from multi-ROIs and other objects for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
Compute Measurements	Replaces the menu item 'Scalar Generator' and opens the Compute Measurements dialog, which has the same functionality of the previously named Scalar Generator dialog (see Compute Measurements on page 37).
Create an ROI from Touching Label Voxels	Lets you create a region of interest that only includes the labeled class voxels that touch each other. The example below shows the ROI created from a multi-ROI with three classes colored yellow, blue, and green. 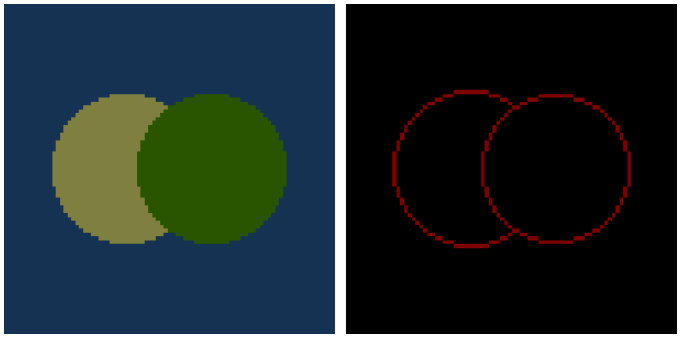 NOTE Processing is done in the same order as the classes are listed in the Classes and scalar information box on the Data Properties and Settings panel.
Extract Slab Around Center	Extracts a set number of rotated multi-ROIs (slabs) along the object's central axis. You can a select the number of slabs that will be extracted and the number of slices per slab.
Grow Selected Labels with Local Otsu	Lets you grow the selected labels of a multi-ROI within the lower or upper local Otsu threshold and by a specified amount.

Mesh Contextual Menu

The items listed below are now available in the **Mesh** pop-up menu.

Mesh pop-up menu changes

Item	Description
Analyze and Classify Measurements	Opens the Analyze and Classify Measurements dialog, in which you can import scalar data from meshes and other objects for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
Compute Measurements	Replaces the menu item 'Scalar Generator' and opens the Compute Measurements dialog, which has the same functionality of the previously named Scalar Generator dialog (see Compute Measurements on page 37).
Remesh with Poisson Surface Reconstruction	Lets you reconstruct triangle meshes from sets of oriented 3D points (see Remeshing with Poisson Surface Reconstruction on page 53).

Item	Description
Merge Meshes	This item, which is available whenever two or more meshes are selected, lets you merge the selected meshes into one mesh.

Graph Contextual Menu

The items listed below are now available in the **Graph** pop-up menu.

Graph pop-up menu changes

Item	Description
Analyze and Classify Measurements	Opens the Analyze and Classify Measurements dialog, in which you can import scalar data from graphs and other objects for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
Compute Measurements	Replaces the menu item 'Scalar Generator' and opens the Compute Measurements dialog, which has the same functionality of the previously named Scalar Generator dialog (see Compute Measurements on page 37).
Map Edge Scalar Values From	Lets you map edge scalar values contained within another object. You can choose the object that includes the scalar values you want to map from in the Choose Scalar Value Slots dialog, as well as a default scalar value.
Export to CSV	Lets you export the vertex coordinates in a comma separated values (*.CSV extension) file.

Vector Field Contextual Menu

The items listed below are now available in the **Vector Field** pop-up menu.

Vector field pop-up menu changes

Item	Description
Analyze and Classify Measurements	Opens the Analyze and Classify Measurements dialog, in which you can import scalar data from vector fields and other objects for the cross-table analysis of feature vectors (see Analyze and Classify Measurements on page 47).
Compute Measurements	Replaces the menu item 'Scalar Generator' and opens the Compute Measurements dialog, which has the same functionality of the previously named Scalar Generator dialog (see Compute Measurements on page 37).

Views Contextual Menu

The items listed below are now available in the **Views** pop-up menu.

Views pop-up menu changes

Item	Description
Export Screenshot of All Views	Lets you save a screenshot of all views in the workspace.

Preferences Changes

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	
Axis indicator	Two new items are available for setting the Axis indicator preferences — Gyroscope (World) and Gyroscope (View). See Gyroscope Axis Indicators on page 68 for more information about these Axis indicator styles and functions.
Default W/L range	Lets you choose the default range that will be applied for window leveling in 3D views — either computed within the upper Otsu range (**Threshold at upper Otsu)** or applied to the full data range (**Full data range**). In the example below of the initial window leveling range applied, the 'Threshold at upper Otsu' option was selected as the default W/L range. 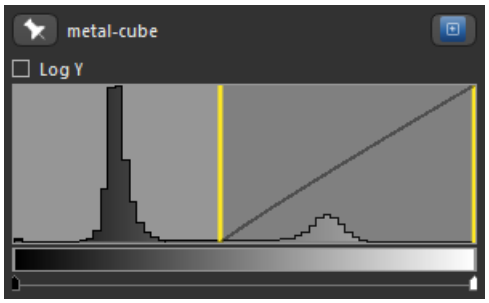
Use lighting by default in 3D	If selected, lighting will be applied by default to new 3D views.
ROI painting in 3D	Lets you choose when updates to regions of interest will appear — after the brush is released or continuously (see ROI Painting Preferences on page 57).

Deep Learning Preferences

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

Deep Learning preferences changes

Item	Description
Enable TensorRT model compilation	Lets you choose to optimize deep models for inference (see Compile for Faster Inference on page 25).
Execute Deep Learning tasks in an external process	If selected, TensorFlow functionality will run as an external process in Windows installations (see External Processes for Deep Model Task Processing on page 29).

System Requirements

As of version 2024.1, Dragonfly 3D World can be installed on systems/laptops without a dedicated graphics processing unit (GPU). The minimum requirement is now an Intel integrated graphics processor (IGP) with UHD Graphics.

NOTE Go to https://en.wikipedia.org/wiki/Intel_Graphics_Technology for a list of Intel IGPs with UHD Graphics.

You should note that performance will likely not be comparable to running Dragonfly 3D World on a system with a dedicated GPU. In addition, some advanced features are not supported on systems/laptops without a dedicated GPU. These include viewing multiple datasets simultaneously, combining images and meshes, using back lighting, and adding reflections to virtual floors.

NOTE Go to https://www.theobjects.com/dragonfly/system-requirements.html for system recommendations to fully experience all the benefits and capabilities of Dragonfly 3D World.

Third-Party Packages

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

Third-party packages

Item	Version 2024.1	Previous version (2022.2)
TensorFlow	2.11	2.4.1
Qt	6.4.2	5.12.9
Python	3.10.12	3.8.12

Publicly Available Source Code

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

TensorFlow

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.

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/.

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 can generate artificial microstructures that mimic real materials and to compute simple morphological properties such as porosity, volume fraction, 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.

Netgen/NGSolve

Netgen/NGSolve is a high-performance multi-physics finite element software. It is widely used to analyze models from solid mechanics, fluid dynamics, and electromagnetics. Users can freely download and examine the NGSolve source code from: https://github.com/NGSolve/ngsolve/releases.

Learning Dragonfly

You'll find our latest instructional videos and recorded webinars on our YouTube channel (https://www.youtube.com/channel/UCuFl2zHcyStR2RJpMXbi6ow). In addition, the new **Dragonfly Social** online community platform lets Dragonfly users start discussions, chat, access material, post announcements, as well as upload and download extensions and deep models for the application.

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.

Contact sales.dragonfly@comet.tech for more information and availability.

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.

Contact sales.dragonfly@comet.tech for information about availability.

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.

Reporting Issues

Dragonfly 3D World includes an integrated reporting module that lets you report any issue that results in a crash. Whenever you choose to send a report, try to include as much information as possible to help the support team to understand why the issue occurred. You can enter details about the issue, such as how to reproduce the problem, in the Send Report to Support Team dialog.

Send Report to Support Team dialog

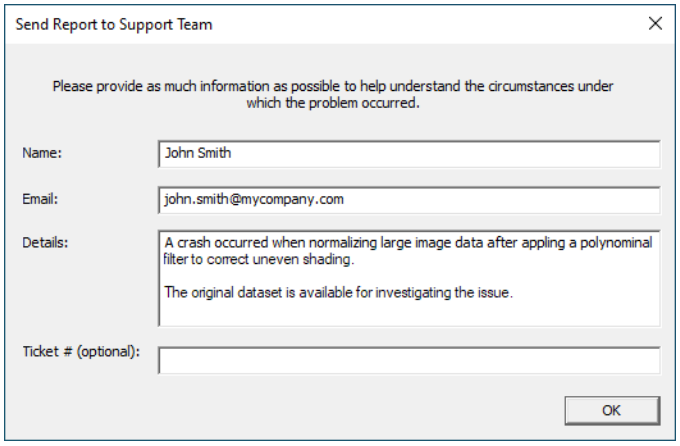

You should note that you can also report an issue at any time by choosing **Help > Report an Issue** on the menu bar. For example, when an unknown error occurs or an unexpected result is generated.

Report an Issue dialog

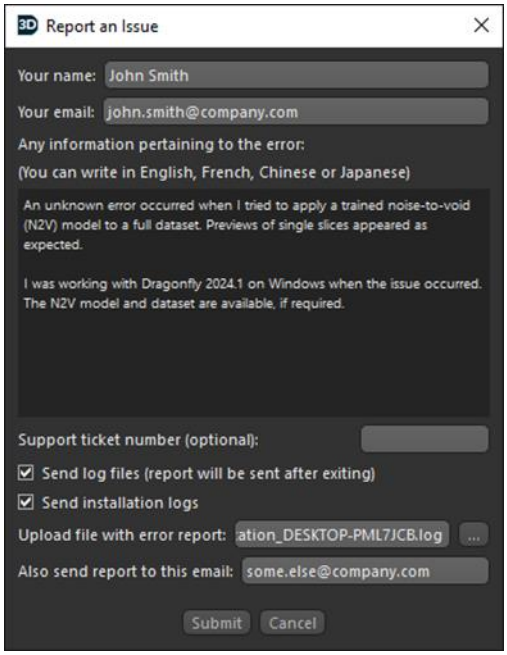

If you decide to report an anomaly, please try to include the following information:

- A description of the issue and when it first occurred.
- The full error message and the steps required to reproduce the problem.
- The installed version of Dragonfly 3D World.
- The operating system of any affected machine.
- If the problem affects a particular dataset or image, you may be asked to provide a copy of the dataset.