

VAST *Lite* 1.2.1
Volume Annotation and Segmentation Tool

User Manual

Daniel R. Berger

November 15th, 2018

Contents

1	Introduction	1
2	Getting Started	5
2.1	System Requirements	5
2.2	Program Setup	5
2.2.1	Try It Out!	6
2.2.2	Preferences	6
2.3	Additional Files Included With VAST	8
3	Working with VAST	9
3.1	Image Stack Importing	9
3.1.1	Importing image stacks: Pattern-based names	10
3.1.2	Lossless and lossy compression	11
3.1.3	Importing 3D volume files	12
3.1.4	Image scale and description	12
3.1.5	Preparing .VSVI files for large image stacks	13
3.1.6	Remote image stacks	16
3.2	Viewing and Navigating an Image Stack	17
3.2.1	The sidebar	17
3.2.2	The RAM (main memory) usage indicator	17
3.2.3	Getting and setting coordinates	18
3.2.4	Layers	18
3.3	Painting	20
3.3.1	Multi-scale painting	21
3.3.2	Automatic Z-filling	22
3.3.3	Masking	22
3.3.4	Semi-automatic segmentation from boundary maps	23
3.3.5	Semi-automatic segmentation from an imperfect automatic segmentation	23
3.4	Filling	24
3.4.1	Filling segments in the same layer	25
3.4.2	Filling segments from a source layer to a target layer	25
3.5	Splitting	26
3.6	Segments	27
3.6.1	Picking segments	27
3.6.2	The segment hierarchy	27
3.6.3	Re-ordering and moving segments in the tree	27
3.6.4	Collapsing and expanding tree branches	28
3.6.5	Adding new segments	28
3.6.6	Using anchor points and other coordinate triplets	28

3.6.7	Helper functions for arranging segments	29
3.6.8	Select recently selected segments	29
3.6.9	Global operations: Deleting and welding segment subtrees	29
3.6.10	Global operations: Reevaluating bounding boxes	29
3.6.11	Global operations: Segmentation cleaning	30
3.6.12	Segment tags	30
3.6.13	Editing the color of a segment	30
3.6.14	Exporting segment metadata	31
3.6.15	Segment information	31
3.6.16	Searching for a segment with a given name or ID	31
3.6.17	The 'Collect' tool	31
3.7	Saving Segmentations	32
3.7.1	Save Segmentation As Special	32
3.8	Segmentation Merging	32
3.9	Importing Segmentations From Image Stacks	33
3.10	Exporting Image Stacks	35
3.11	The 3D Viewer	37
3.11.1	Using the 3D viewer to set the 2D view location	37
4	The VastTools Matlab Toolbox	39
4.1	Getting started with the VastTools Matlab Toolbox	39
4.2	Exporting 3D Models	40
4.3	Export Particle Clouds (3D Object Instancing)	43
4.4	Export 3D Boxes, Planes and Scale Bars	43
4.5	Exporting Projection Images	44
4.6	Measuring	46
4.6.1	Measure Segment Volumes	46
4.6.2	Measure Segment Lengths	46
4.6.3	Measure Segment Surface Area	46
4.6.4	Euclidian Distance Measurement Tool	46
4.7	Simple Navigator Images	47
4.8	Target Lists	47
A	FAQ and Trouble Shooting	49
A.1	Frequently Asked Questions	49
A.2	Typical Use Cases	51
A.3	Some Performance Tips	52
A.4	Setting up VAST with a Wacom screen	53
A.5	Keyboard Shortcuts in VAST	54
A.6	Terms of Usage and Privacy Statement	55
B	Technical Information	57
B.1	Size limitations	57
B.2	Supported file formats for importing / exporting	57
B.3	API Function Reference	58

Chapter 1

Introduction

VAST [7] is a utility program for manual annotation and segmentation of large volumetric image (voxel) data sets. It enables users to work with data sets in the Terabyte or even Petabyte range at interactive speeds, to explore them visually and to label structures of interest by voxel painting.

Voxel painting has a number of advantages over other manual segmentation approaches like bread crumbing (placing a number of labeled points inside each object) and skeletonization (placing labeled points and connecting them with edges to form a 'skeleton'). It reveals the true shape of objects and makes visualizations of the data more comprehensive. It also allows measurement of volume and surface shape properties of the labeled objects. When working on a dense segmentation, the fact that voxel painting labels areas rather than single points or lines in the cross-sections of objects makes it a lot easier to spot objects that have not been labeled yet (visual pop-out). It also enables additional functionality, for example reliably determining neighbor regions, using voxel overlap of several segmentations for merging regions or determining synaptic connectivity, and voxel-by-voxel masking, which offers powerful capabilities to use (imperfect) automatic segmentations to speed up manual painting. Some of these functions are also difficult to implement when using outline labeling (drawing vectorized lines around process cross-sections in 2D slices). In addition, outline labeling can introduce problems in identifying corresponding outlines in different slices and can cause unintended region overlap. Also, many machine learning algorithms for automatic segmentation rely on voxelized labelings for training, and the output of many such algorithms are labeled voxelized regions. VAST can be used to generate volumetric training data sets in their native format, and can also be used for importing, proof-reading and correcting results of segmentation algorithms.

The disadvantages of voxel painting which are mentioned most often are that it is slower than for example skeletonization or bread crumbing, and that it needs more storage space than alternative methods, especially when working with very large data sets. VAST tries to alleviate both disadvantages. For voxel painting, the time needed to label an object is largely determined by the number of outlines that have to be drawn. VAST can do automatic convex Z-filling to reduce the number of outlines by a factor of up to 8 if accurate boundary tracing is not needed. It can also take advantage of (even imperfect) automatic segmentation to define object boundaries automatically during painting. Even whole supervoxels can be copied from one layer to another using the fill tool. Painting is also optimized to be highly responsive, and the user interface provides quick access to functions which are used often (navigating through the image stack, changing the tooltip size, color picking, and switching between paint and delete mode). To reduce the amount of required memory and to enable interactive painting speed with any tooltip size, VAST implements multi-scale painting.

Considering that alternative labeling methods are likely to be more error-prone and will probably require more time until an acceptably low error rate is reached, voxel painting might be the overall faster alternative for fully manual labeling. If an object skeleton is needed (for example to compute length of a dendrite or number of spines etc.), it can be computed from the voxelized segmentation,

whereas the inverse operation (computing accurate volumes from skeletons) is much harder. In a preliminary test we found that using VAST, a very experienced user can produce a dense segmentation of well-stained and well-aligned cortex neuropil data (6x6x30 nm voxel size) at a speed of about 4 cubic microns per hour. If cross-sections are labeled with a color dot in the middle rather than accurately outlining and filling the cross-sections, the labeling speed can be increased by a factor of 3 at the expense of accuracy (12 cubic microns per hour). If the filling function is used with a reasonable automatic segmentation, tracing processes can reach speeds exceeding one millimeter per hour and is essentially limited only by the speed of data streaming and human perception.

VAST is also very portable and light-weight. It consists of a single Windows executable file which is independent of third-party libraries. It does not require to be installed and can be easily copied and for example run from a memory stick.

Many tools for labeling voxel data exist, but such tools usually have drawbacks when it comes to painting in large data sets. Most tools do not support voxel painting but instead do bread crumbing, skeletonizing or storing object outlines as vector graphics (for example splines). Also most of them require the data to be loaded in RAM completely. This is not possible for the data sets in the Terabyte to Petabyte range which are currently produced by serial-section electron microscopy of biological samples. Many of these tools are also developed as cross-platform applications, so that they can be run on Windows, Mac and Linux systems. This usually means that a non-native GUI system has to be used (e.g. Qt) which makes the program a lot more bulky and more difficult to install and maintain. VAST is a Windows-only program and uses native Windows GUI and graphics functions.

VAST is currently being used in several scientific projects to label cells in electron-microscopic and light microscopic image stacks. It was also used for the recently published studies [3], [4], [5] [6], and [8]. VAST itself was recently also published [7]; please use this reference for citations.

The key concepts of VAST are:

- Image stacks can be imported into VASTs .vsv file format, where they are stored as dices. This, together with pre-computation of mip maps, allows for fast panning and zooming through the data when it is opened in VAST.
- .vsv files support lossless and lossy compression to reduce the resulting file size
- Image data and segmentations are stored in single files, which makes it easy to copy them from one place to another
- Alternatively, which is particularly useful for very large volumes, VAST can load image data directly from image tiles if mip-maps have been precomputed, using a .vsvi descriptor file
- Image stacks can also be loaded from a web server over HTTP (openconnecto.me, neurodata.io, butterfly and Google brainmaps formats)
- Dynamic multi-threaded cacheing in RAM with pre-loading for low-latency display update
- Layering: Several image and segmentation stacks can be opened and displayed together with a number of blending and tinting options
- Multi-scale painting (in VAST, painting always happens at the currently displayed resolution (mip level))
- Automatic convex Z-filling during painting to speed up coarse labeling
- Automatic 2D-filling of closed contours
- 3D flood filling tool

- Painting and filling can be constrained by masks defined by a separate image or segmentation layer
- Label color patterns create a larger number of distinguishable label colors
- Label hierarchies allow for fast and reversible grouping of labeled segments in a tree structure
- Segments have anchor points to quickly find them in the volume
- Exporting of segmentations, EM stacks and mixed image stacks ('screenshots') in multiple formats
- Importing of segmentations from image stacks
- File-modification free editing. Image files are not changed (except if you explicitly update information in them). Segmentation files are only changed if you save the segmentation back to the same file.
- TCP/IP-based API through which external programs can directly communicate with VAST, and supporting Matlab script 'VastTools' which provides supplementary functions like measuring labeled segments and exporting of 3D surface models.

Current limitations:

- Currently only 16-bit segmentations are supported
- There is no 'Undo' function
- Exporting of 3D models of segmented objects is currently only supported externally (using the VastTools Matlab script)
- Only rudimentary 3D model display
- No multi-user support
- Not an open-source project, but the VAST *Lite* executable can be copied and used freely (see section A.6 for details)

Please note that VAST is still under development, and is subject to change as new features are added and bugs are fixed. For bug reports or helpful suggestions, please contact me at: danielberger@fas.harvard.edu.

Chapter 2

Getting Started

2.1 System Requirements

VAST currently runs on 64-bit Windows computers that support DirectX 11. These are Windows Vista, 7, 8 and 10, with DirectX 11 or later available. Windows XP and older versions will not work. Currently only 64 bit versions of these operating systems are supported, because 32 bit programs are limited in the amount of RAM they can handle (4 GB max. theoretically, less realistically). The computer also has to have a DirectX 11 compatible graphics card, which should however be available in all modern PCs.

Recommended system configuration:

- Windows PC with 64 bit Windows 7 or later
- 16 GB of RAM (the more the better)
- DirectX 11 compatible graphics card
- 2 TB of disk space (depends on the size of the data you work with)
- Wacom Cintiq 13HD or other pen touch screen with configurable two-button pen

Minimal system configuration:

- Windows PC with 64 bit Windows 7
- 2 GB of RAM (at least 4 GB recommended)
- DirectX 11 compatible on-board graphics card
- Standard screen and mouse

2.2 Program Setup

To use VAST, simply copy the executable program into a folder where you have read/write access, and set up links on your desktop, start menu and/or taskbar if desired. It is important that VAST has read and write access to the folder where the executable is because it will write a configuration file (`vast_preferences.dat`) into the same folder to store your settings. Then start the executable.

2.2.1 Try It Out!

The quickest way to try out VAST is to use an online data set. Several online data sets are included in the .ZIP package of supplementary files as .vsvr files. You can save this package from the executable by clicking 'Yes' in the 'First Start' pop-up window when you start VAST for the first time, or by choosing 'Save Documentation .ZIP To Disk...' from the 'Info' menu.

Unzip the package (or drag&drop the contents to a folder outside the .ZIP using Windows Explorer for example).¹ Then, in VAST, go to 'Open EM...' and select one of the .vsvr files in the VAST_package/Online Datasets/ folder, for example 'kasthuri11_neurodata.vsvr'.² This will load images of a big EM stack from the Johns Hopkins neurodata.io server. Your computer has to have internet access for this to work.³

Click and drag the EM slice to pan. Use the mouse wheel to zoom. Use UP/DOWN arrow keys or A/Z to scroll through the stack.⁴

Click on the little pencil icon in the toolbar to switch to 'Paint' mode. Choose 'Yes' in the popup window. Click and drag over the image to paint. You can select different paint colors in the 'Segment Colors' tool window. To erase, hold down the 'Delete' key while painting (or click and hold the right and left mouse buttons together). Select 'Keyboard Shortcuts' from the 'Window' menu for a list of available keyboard functions.

2.2.2 Preferences

VAST will set up the preferences for you when you run it for the first time on a computer (whenever it cannot find the preferences file `vast_preferences.dat`). You can edit these preferences if you want to. In the main menu of VAST, go to 'File / Preferences ...' to open the Preferences dialog window. You should at least check once whether the folder in which VAST puts its temporary disk cache is on a hard drive with lots of free space. Depending on what you do, temporary disk cache files can get similarly large as the segmentation files you are working with, in particular if you use global segmentation editing functions like merging, segment deleting, or segment welding. Each VAST instance will write a cache file (`vastsegcache*.vss`) to this folder and delete it when VAST is closed. In case VAST crashed (which should not happen) or power was lost while VAST was running, a cache file may be left there which takes up disk space, so you may want to check for and delete old cache files once in a while if VAST was terminated abnormally.

Memory and Cacheing

On the left side of the Preferences dialog you can set how much cache memory VAST will use maximally for voxel images and for segmentations. VAST chooses initial values which are reasonable for your system. The rule of thumb is: If you can afford it, leave 1-2 GB for the system, and split the rest 1/3 each for image cache, segmentation cache, and general usage of VAST (don't assign). On a system with 8 GB RAM, this means to give 2000 MB to the image RAM cache and 2000 MB to the segmentation RAM cache. If you are only viewing images and not using segmentations, you can increase the size of the image cache and reduce the size of the segmentation cache accordingly. If you plan to use other programs at the same time or run two instances of VAST, please reduce these values as needed. VAST will not immediately use all of the allotted memory, but it will stop reserving new memory for cache blocks and re-use old blocks when it reaches the limit.

¹Extracting the package contents to actual files is important! `vasttools.m` for example will not work properly when started from within the compressed .zip file.

²Alternatively, you can simply drag and drop the .vsvr file from the Windows Explorer window onto the main window of VAST to open it.

³Remote EM images are loaded from a server using a cutout service over HTTP. The .vsvr file is just a text file which defines the web address of the dataset to load and its dimensions.

⁴Please be patient for the images to load, the neurodata.io server is very slow.

In general, do not allow VAST to allocate more memory than the system has. This can result in severe performance issues. There is a memory usage indicator in the upper right corner of the VAST window which shows you how much memory is currently used. The blue frame indicates the maximum amount of RAM which VAST uses for image and segmentation cacheing. If the memory indicator becomes red and your system slows down, try to REDUCE the cache limits to allow Windows and VAST to use more RAM for other data.

Some of the segmentation cache is used for holding the currently displayed part of the segmentation in memory. When you exit the preferences dialog, VAST will tell you how much of the segmentation cache it needs for the current display settings and whether the cache size is sufficient.

'Disk Cache Directory': Here you can specify the folder where VAST stores its temporary disk cache. Click the '[...]' button to browse. Set this to a folder where you have lots of free space (more than the size of the largest segmentation file you will be working with, since for certain functions VAST has to duplicate the segmentation data).

Painting

The segmentation bit depth is currently fixed at 16 bits. This allows for a maximum of 65535 labels; however, since each segment is represented in the tree view of the 'Segment Colors' window in VAST, memory limitations in the Windows system might prevent VAST from using that many labels.⁵

'Tablet Mode (Pen Paints, Finger Moves)': On some pen-enabled tablet computers VAST can distinguish finger and pen input. If this mode is enabled, the pen should paint and the finger should move the view when in Paint Mode. This may not work on some newer systems.

'Generate New Colors From:': Allows you to select which color space to use for assigning random colors to new segments. Segment colors can of course also be changed manually (see section 3.6.13).

'Adjust Scroll Speed With Mouse Wheel While Z-Scrolling:': If this option is enabled the z-scroll speed (the rate at which you go through slices when holding down A, Z, S, X, Up, Down, PageUp, PageDown) is decoupled from the key repeat speed and can be controlled with the mouse wheel while one of those keys is held down. Roll the mouse wheel towards you to 'pull down' the scroll speed. Roll away from you to increase the scroll speed until you reach the maximum, which is the window update speed. The mouse wheel will still adjust the zoom level while you are not z-scrolling.

Display Properties

'Maximum Window Width' specifies the width (or height, whichever is greater) of the largest window you will be using, in pixels. This value is used to determine how many textured tiles are needed to fill the entire window at all zoom levels. Setting this value smaller reduces memory consumption and increases cacheing speed, but if the value is too small, the image texture might not reach all the way to the sides of the window at all zoom levels.

'Target Resolution Smaller Than' lets you specify the effective resolution of the displayed textures on the screen in screen pixels per texture pixel. This affects at what zoom levels which mipmaps are used. '2' is a good general setting; '1' makes it more detailed but slower (and more memory-consuming), and '4' makes it faster but blurry.

'Texture size m (texture is m^2):' defines how large the texture tiles will be which are used for displaying image and segmentation textures. Depending on the graphics card some texture sizes might be faster than others. I recommend to leave this setting at 128.

'Texture Smoothing': You can set here whether you want to use texture interpolation. This reduces aliasing effects but can result in a slightly blurred appearance of the textures. The most

⁵It runs fine with more than 36000 labels in one of our data sets.

natural setting of this is, in my opinion, 'All except Mip 0', which will show pixels with sharp boundaries only if you zoom in more than the native resolution of the image data.

'Link Tool Windows': When enabled, VAST's floating tool windows will automatically move to their default locations with respect of the main window when the main window is moved or rescaled.

The remaining options are self-explanatory. Opacity values have to be set between 0 (fully transparent) and 255 (fully opaque).

Other

'Enable API Remote Connections at Startup': It can be useful to enable this option if you often use the VAST API to link to external programs. The API remote connection functionality can also be enabled and disabled in the 'Remote Control API Server' dialog under 'Window' in the VAST main menu.

Press OK after you're done configuring the preferences.

2.3 Additional Files Included With VAST

Under 'Info / Save Documentation .ZIP To Disk ...' in the main menu you can save out additional files which are packaged into the VAST executable, as a ZIP file. Select a target location and save, then unzip the ZIP file.

Currently this includes a set of .vsvr files to access some large EM data sets remotely (see section 3.1.6), some Matlab scripts which can be useful for analyzing VAST data in Matlab, including the VastTools toolbox which can communicate directly with VAST through the API and provides additional functionality (see chapter 4), and this documentation as a PDF file.

Chapter 3

Working with VAST

VAST uses its own file format `.vsv`¹ to store image data. It can open `.vsv` files immediately and navigate in them quickly. If your data is a stack of for example `.png` images, you can either import it into a `.vsv` data file (recommended for image stacks of a few hundred Gigabytes or smaller), or prepare tiles, mipmaps and a `.vsvi` file to let VAST load the image data directly from image files (see section 3.1.5). After opening a `.vsv` or `.vsvi` image file in VAST, you can create a segmentation by painting on top of the images, or you can open an associated segmentation file (`.vss`)² and view image and segmentation together. `.vss` files tend to get big quickly, but can be packed efficiently, for example in a ZIP file. You can view segmentations, modify and save them. You can export segmentations as image stacks for using them in other analysis programs or to render the segmented objects in a 3D animation program like 3D Studio MAX. You can also import segmentation image stacks that were generated externally.

3.1 Image Stack Importing

Typically volumetric image data is stored as a series of 2D images, or as a serial 3D block of data, which is not suitable for fast interactive viewing. When you import such a data set into VAST as a `.vsv` file, it converts the images to a diced data structure, and computes and includes mipmaps for the images, all stored within the `.vsv` file.³ Using diced data does not only speed up loading of parts of images, but also enables fast loading of volumetric sub-regions or 2D sections at other orientations through the image data.

VAST does currently not include image alignment and stitching functions. If you are starting with an unaligned stack of images, you will first have to align the images with a different program (Fiji or Photoshop, for example) and then save a stack of aligned images which all have the same dimensions and are named and numbered in a consistent way (for example `img000.png`, `img001.png`, ...). Ideally all images should be in the same folder.

VAST can import single-tile image stacks, multi-tile image stacks, and 3D volume files. In a single-tile image stack, each slice of the stack consists of a single image file. In a multi-tile image stack, each slice is composed of several tiles in a XY grid, and each tile is stored in a separate image file. A 3D volume file stores all slice images in a single file (currently supported are multi-image

¹The file name extension `.vsvol` can be used instead, and may be preferable, because Microsoft thinks `.vsv` is a 'Microsoft Visio' file.

²Equivalently, you can use `.vsseg` instead of `.vss` as an alternative file name extension for VAST segmentation files.

³A *mipmap* is a downsampled version of an image. VAST uses power-of-two (2D, XY) mipmaps. For example, for an original image of 1024x1024 pixels, it will compute mipmaps of 512x512 and 256x256 pixels. It does this for every slice image in a stack.

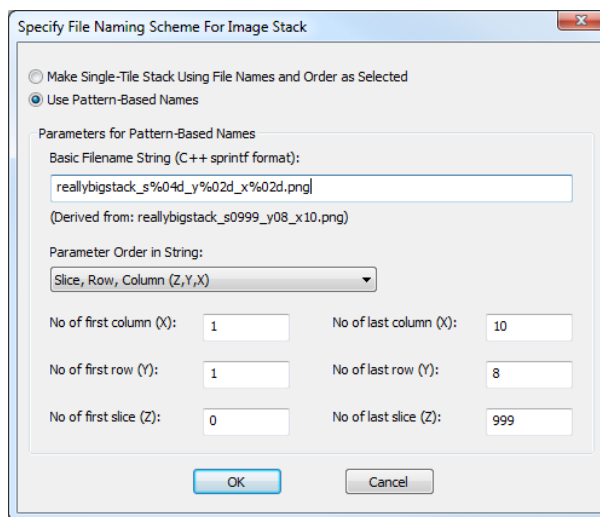


Figure 3.1: First dialog for importing EM image stacks: Specification of pattern-based names

.TIFF and NiftI .nii files). VAST will convert image data to either 8-bit graylevel or 24-bit RGB when importing.

For importing and dicing, VAST will use the RAM cache which is normally used for cacheing EM image data during viewing and painting. Having lots of cache memory available will make importing somewhat faster, because images have to be re-loaded less often. You can set the size of the EM image cache in the Preferences (see section 2.2.2).

3.1.1 Importing image stacks: Pattern-based names

In the main menu of VAST, go to 'File / Import / Import EM Stack from Images to .VSV file ...'. VAST will show a file browser dialog in which you can select one or several image files. For importing 3D NiftI or multi-image TIFF files, please select only one file. If you import a single-tile stack and do not want to use pattern-based names, select all slice images in the correct order, because images will be stacked in the same order in which they appear in the system's list of selected files. The order is usually correct if you select the last image first, then shift-click (hold the SHIFT key down and click left with the mouse) the first image to select the whole range. You can also 'Select All' by pressing CTRL-A if the folder only contains the image files you want to import. If you are worried about the order of the images and want more precise control, you can use pattern-based names. If you make use of pattern-based names to import single- or multi-tile image stacks, it is sufficient to select one file, but even better to select the first and the last file in your set of images. Then click 'Open'.

After selecting one or more image files (except for NiftI and multi-image TIFF files), VAST will display the dialog shown in Figure 3.1. To import without pattern-based names, select 'Make Single-Tile Stack Using File Names and Order as Selected' and press OK.

If you select the second option, 'Use Pattern-Based Names', the parameters in the lower part of the dialog window will be enabled. With pattern-based names, you specify a template string for the file names which contains placeholders for numbers, and ranges for these numbers. With this you can also import image stacks in which each slice is stored in several image files (multi-tile image stacks). This is useful for data sets in which a single slice is so large that it can not be stored in a single image file, but is stored in a set of tiles which form a regular grid. Please note that these tiles should *not* be unstitched image tiles as they come off a microscope, but they have to fit seamlessly.

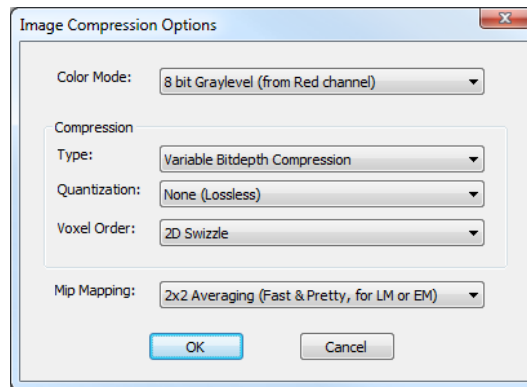


Figure 3.2: Second dialog for importing EM image stacks: Image compression options

If you have a set of raw microscopic images which are not yet stitched and aligned, please use an external program to generate a stitched and aligned image stack first, and store each slice as a single image or a set of image tiles. You can then import those images into VAST.

VAST will use the file(s) you selected in the previous dialog to determine the source directory where the image files are and to generate a basic template for the file name. It assumes that all images of the stack are in the same folder⁴ (the one you picked an image from), and are named consistently with numbers for slices, rows and columns. In this dialog, you specify these ranges as well as a schema to derive the filename for a given slice/row/column coordinate. Not all image tiles have to exist, but VAST will warn you if image tiles can not be found.

Let's say, for example, you have a data set called 'bigstack', which has 1000 slices, numbered from 0 to 999, and each slice has 10x8 tiles, numbered from 1 to 10 and 1 to 8. Assume the image in the upper left corner of the first slice is called 'bigstack_s0000_x01_y01.png', the image tile right of it is called 'bigstack_s0000_x02_y01.png', and so on. The last image in the lower right corner of the last slice would be called 'bigstack_s0999_x10_y08.png'.

First, make sure that the file name in the edit box at the top contains the correct *C++ format string* (as it is used by `printf()`). In general, numbers which specify the slice, column and row coordinates have to be replaced by codes like '%d' (integer number) or '%04d' (integer number with zero-padding to 4 digits). VAST will then fill in those numbers for each image. The correct string for the example above would be: 'bigstack_s%04d_x%02d_y%02d.png'. For more information about format strings, please refer to a C++ manual or ask the internet.

You can also wildcards ('*') in the file name to allow for part of the name to change arbitrarily, for example if the wafer number or date or other information is included in the filename.

In the combo box below, select which coordinates are used in the file names, and in which order they appear. The edit boxes below let you specify the range of (integer) numbers for the three coordinates. Figure 3.1 shows how to set up these parameters for the example discussed above. After you entered all parameters, press OK.

3.1.2 Lossless and lossy compression

Next, VAST will show a dialog where you can specify the color mode and image compression options (Figure 3.2).

Under 'Color Mode', please select if you want to import the images as 8-bit graylevel or 24-bit

⁴There is a workaround if images are stored in separate folders per section – You can specify a relative path within the filename template and use the slice parameter to index the folder. Remove an additional slice counter in the filename by using a wildcard (*).

color images, and for graylevel which source color channel to use. When importing from graylevel images, please select the first option ('from RED channel').

Under 'Compression' you can specify compression options. Under 'Type' you can select the compression method - Uncompressed, Variable Bitdepth Compression, zlib Compression, and Spectral Compression. The three different compression algorithms are by themselves lossless, but might produce slightly smaller or larger file size depending on your data. Variable Bitdepth Compression should be fastest when reading from the compressed files.

'Quantization' specifies whether the compression should be lossy or lossless. Lossyness is achieved by quantizing, meaning throwing away bits. For example, if you set Quantization to -2 bits, graylevel images will have only 6 bits resolution (64 different gray levels) rather than 8 bits (256 different gray levels). Throwing bits away reduces both file size and image quality.

'Voxel Order' defines in which order the pixels in the images will be stored. This can have an effect on compressed file size. '2D Swizzle' stores pixels in 2D Z-order. I usually get best results using '2D Swizzle', but '3D Swizzle' might be superior for very well aligned data (e.g. FIB-SEM).

'Mip Mapping' lets you select whether you would like to average or subsample pixels for the generation of mip maps. Averaging gives superior image quality for EM and LM image stacks, but changes the pixel values. Therefore, if pixel values should be preserved, for example if the image stack represents an automatic segmentation result (see section 3.3.5), subsampling should be used instead of averaging.

Next VAST will ask you to specify a target location and file name for the resulting .vsv image volume file. Use '.vsv' as extension for the file name. Choose a location where you have enough storage space for the file. The file will not only contain the original image data, but also the mip maps. For example, if you import 1024 images of 1024x1024 pixels each and store uncompressed, the .vsv file will be approximately $1024 \cdot 1024 \cdot 1024 + 512 \cdot 512 \cdot 1024 + 256 \cdot 256 \cdot 1024 = 1409286144$ Bytes (≈ 1.3 Gigabytes) large. Lossless compression will reduce the file size, and lossy compression even more, but by how much depends strongly on your data and the compression method used.

Then, the images will be read, diced, and put into the target file. After that, VAST will compute the mipmaps and put those in the target file too. Depending on the size of the data, this process can take several hours. For example, a big data set of 350 GB takes about 5 hours to import on a recent desktop machine. The limiting factor is usually the transfer speed of data from and to the hard drives.

You can cancel the importing, but the target file will then be incomplete / corrupted and can not be used with VAST.

3.1.3 Importing 3D volume files

Importing a 3D volume file is easier than importing an image stack. VAST currently supports multi-image TIFF and NiftI files. VAST will ask you to specify the name of the source file and the name of the target .vsv file. For NiftI, files are loaded into memory completely for importing, so this only works for smaller volumes, and they are expected to store 8-bit data.

3.1.4 Image scale and description

After importing, VAST will show a dialog where you can set the voxel size in nanometers of your data in the file. Press 'Save to file' in the dialog to save the information you entered to the VSV file. The voxel size entered here is used for the scale bar which you can enable in the main menu under 'Info / View Scale Bar', for the 3D viewer, and for scaling models and measurements in VastTools. You can access and change these numbers for the selected image layer under 'Info / Volume properties ...' in the main menu. This dialog also displays how large your image stack is in voxels.

Under 'Info / EM File Information ...' in the main menu you can enter and view text which will be stored in your VSV file as well. This can contain a description of the data, copyright information, or other.

3.1.5 Preparing .VSVI files for large image stacks

For image stacks which are larger than a few hundred Gigabytes, importing to a .vsv file may become unfeasible. Importing might take weeks or months, and the resulting file might become too large for the file system or for storage on a single drive. For these cases VAST now supports reading image stacks directly from sets of image tiles. However, the images have to be prepared in a particular way and a .vsvi file has to be created which tells VAST how to read different parts of the stack from the images. The following conditions have to be met:

- All image tiles should be stored in the same folder tree and named in a regular way so that VAST can generate file names to load all tiles needed for a particular region. The rules of tile naming are described in the .vsvi file for the data set (see below).
- Powers-of-two mipmaps (2D, XY) have to be precomputed and stored with the original-resolution tiles.
- Image tiles should be small enough so that jumping to a new location will not result in a long delay before images show up, and the overhead for loading a region of data remains reasonable. We have made good experiences with tiles which are 1024x1024 pixels large. Of course there is a tradeoff between the size of a single tile and the number of tiles which have to be stored. Split into tiles of 1024x1024 pixels, a 100x100x100 micrometer dataset at a resolution of 4x4x30 nanometers, which is roughly 2 Terabytes large, results in more than 2.8 million image files.

Storing image data as a set of files also has other advantages: Parts of the image stack can easily be read or exchanged by other programs, and the image stack can also be extended easily.⁵ Also, since different resolutions (mip maps) of the stack are stored separately, it is possible to only copy lower-resolution mip maps to a target drive and thus generate a much smaller lower-resolution version of the same data set.

.vsvi files are text files in a JSON-like format. Below are the contents of a .vsvi file which describes a tiled image stack with mip maps:

```
{
  "Comment": "IARPA R0 100mu normalized",
  "ServerType": "imagetiles",

  "SourceFileNameTemplate": ".\\mip0\\%04d*_montaged\\%04d*_tr%d-tc%d.png",
  "SourceParamSequence": "ssrc",
  "SourceMinS": 1,
  "SourceMaxS": 3394,
  "SourceMinR": 1,
  "SourceMaxR": 25,
  "SourceMinC": 1,
  "SourceMaxC": 25,

  "MipMapFileNameTemplate": ".\\mip%d\\slice%04d\\%04d_tr%d-tc%d.png",
  "MipMapParamSequence": "mssrc",
  "SourceMinM": 1,
  "SourceMaxM": 7,

  "SourceTileSizeX": 1024,
```

⁵The downside of this is that other dependent data sets, like segmentation layers, will not fit to this image stack any more if its size or alignment is changed.

```

"SourceTileSizeY": 1024,
"SourceBytesPerPixel": 1,
"MissingImagePolicy": "nearest",

"TargetDataSizeX": 25600,
"TargetDataSizeY": 25600,
"TargetDataSizeZ": 3394,
"OffsetX": 0,
"OffsetY": 0,
"OffsetZ": 0,
"OffsetMip": 0,
"TargetVoxelSizeXnm": 4,
"TargetVoxelSizeYnm": 4,
"TargetVoxelSizeZnm": 30,
"TargetLayerName": "IARPA R0 100mu normalized"
}

```

The different parameters in the file have the following meaning:

- "Comment": This text will be displayed in the user selects 'Info / EM File Information ...' from the main menu in VAST
- "ServerType": This currently has to be "imagetiles"
- "SourceFileNameTemplate": The filename template (with relative path from the location of this file) for accessing the full-resolution image tiles, in printf() format. This may include wildcards (*).
- "SourceParamSequence": The parameter order of number parameters in the above template. "ssrc" for example means Slice, Slice, Row, Column
- "SourceMinS": The minimal value used for the Slice (Z) parameter in the stack. Slices count from the top to the bottom of the stack (front to back in 3D, assuming they are sections which are cut from a block).
- "SourceMaxS": The maximal value used for the Slice (Z) parameter in the stack
- "SourceMinR": The minimal value used for the Row (Y) parameter in the full-resolution stack. Rows count from the top to the bottom of the screen
- "SourceMaxR": The maximal value used for the Row (Y) parameter in the full-resolution stack
- "SourceMinC": The minimal value used for the Column (X) parameter in the full-resolution stack. Columns count from left to right on the screen
- "SourceMaxC": The maximal value used for the Column (X) parameter in the full-resolution stack
- "MipMapFileNameTemplate": The filename template (with relative path from the location of this file) for accessing the mip-map image tiles, in printf() format. This may include wildcards (*).
- "MipMapParamSequence": The parameter order of number parameters in the above template. "mssrc" for example means Miplevel, Slice, Slice, Row, Column.
- "SourceMinM": The minimal value used for the mip levels (M) parameter in the stack. This should always be 1. In VAST, mip level m has a reduction factor of $1/2^m$, meaning that mip level 0 is full resolution, 1 is half resolution (in both X and Y), 2 is quarter resolution and so on.

- "SourceMaxM": The maximal value used for the mip levels (M) parameter in the stack
- "SourceTileSizeX": This is the horizontal size of a single image tile, in pixels
- "SourceTileSizeY": This is the vertical size of a single image tile, in pixels
- "SourceBytesPerPixel": This value is 1 for graylevel images and 3 for RGB images
- "MissingImagePolicy": This option describes what VAST will do if an image tile file can not be found. It can be set to "black" (show black tile), "white" (show white tile), or "nearest" (search up or down in the local 16-slice window for the nearest image tile which exists and use that instead; if none is found, leave black)
- "TargetDataSizeX": X extent of the whole stack, in pixels at full resolution
- "TargetDataSizeY": Y extent of the whole stack, in pixels at full resolution
- "TargetDataSizeZ": Z extent of the whole stack, in pixels at full resolution
- "OffsetX": X placement offset of image stack. Please leave this 0 at the moment.
- "OffsetY": Y placement offset of image stack. Please leave this 0 at the moment.
- "OffsetZ": Z placement offset of image stack. Please leave this 0 at the moment.
- "OffsetMip": Resolution shift in number of mipmaps
- "TargetVoxelSizeXnm": X size of one voxel at full resolution, in nanometers
- "TargetVoxelSizeYnm": Y size of one voxel at full resolution, in nanometers
- "TargetVoxelSizeZnm": Z size of one voxel at full resolution, in nanometers
- "TargetLayerName": Name of the layer when loaded into VAST. (Currently unused; instead, the file name is displayed as layer name)

Optional parameters:

- "SourceSectionOrder": For example "SourceSectionOrder": "36:68, 72, 70, 84:166"; optional direct specification of the numbers to be used for the source sections, in Matlab notation. If this parameter is included, SourceMinS, SourceMaxS, and OffsetZ still have to be included but are ignored.

In this example the main folder contains the .vsvi file and subfolders mip0, mip1, ..., mip7. mip0 contains the full-resolution tiles, in a separate subfolder per slice. The subfolder for slice 0 files is named 0001_W01_Sec001_montaged for example, and the subfolder for slice 3394 is named 3394_W13_Sec271_montaged. As you can see the name components describing the wafer number W and section number Sec are changing but are not used for indexing the slice. They are therefore replaced by a wildcard (*) in the filename template in the .vsvi file.

VAST has a function to help generating a .vsvi file for an image data set. To run this, select 'File / Import / Prepare .VSVI Image Stack ...' in the main menu of VAST. First, VAST will ask you to specify a target file name for the .vsvi file and a location where to save it. Ideally the .vsvi file should be stored in the main folder for the data set and all the image tiles should be stored in subfolders, so that they can be referenced by a relative path. Next, VAST will ask you to select one or more example full-resolution image tiles. Then it will show you a dialog to specify the parameters of the .vsvi file. For this example the parameters should be set as shown in Figure 3.3.

Once you click the 'Save .VSVI' button, the information you entered will be saved to the .vsvi file you specified beforehand. VAST will not immediately open that file but you can load it in the

Prepare .VSVI Descriptor File for EM Image File Stack

Name: IARPA R10 100mu normalized

Description: IARPA R10 100mu normalized

Full-res (Mip0) File Name Template (sprintf() Format, Wildcards Allowed):

Full-res (Mip0) Parameter Order: ssrc (Use s (slice), r (row), and c (column))

Mip0 Min Column (c): 1 Source Tile Size X (pixels): 1024

Mip0 Max Column (c): 25 Source Tile Size Y (pixels): 1024

Mip0 Min Row (r): 1 Missing Image Policy: Copy Nearest

Mip0 Max Row (r): 25 Mip0 Voxel size X (nm): 4.000000

Mip0 Min Slice (s): 1 Mip0 Voxel Size Y (nm): 4.000000

Mip0 Max Slice (s): 3394 Voxel Size Z (nm): 30.000000

Mip-Map 1-m File Name Template (sprintf() Format, Wildcards Allowed):

Mip-Map 1 - m Parameter Order: msssrc (Use m (mip), s (slice), r (row), and c (col))

Figure 3.3: Dialog for the preparation of a .vsvi file, filled out for the example in this section

same way as any other .vsv or .vsvr file. Since the .vsvi file is a text file, you can also view and edit it with normal text editors.

The .vsvi preparation tool will compute the total extent of your image stack from the tile size and number of tiles which you provided. It will also compute how many mip levels will be expected by VAST and store that value in the SourceMaxM parameter. The lowest-resolution mip map which VAST uses is the first one in which both width and height of the complete mip map slice image are less than 512 pixels. You can check which value VAST wrote to the .vsvi file to make sure you provide all mip maps needed.

3.1.6 Remote image stacks

In addition to using an image stack in VAST which has been imported into a local .vsv file, or .vsvi image stacks, you can also open and access image stacks which are hosted online. VAST supports the 'Open Connectome Project Cutout Service' from <http://www.openconnectome.me>⁶, the newer neurodata.io data servers, Harvard's *Butterfly* cutout service, and Google's 'Brainmaps' framework. To access a remote image stack you need a .vsvr file which specifies the parameters of the data set. .vsvr files are text files in a JSON-like format and share many parameters with the .vsvi files as described in section 3.1.5. Here is the content of the openconnectome_kasthuri11.vsvr file:

```
{
  "Comment": "Source: http://openconnectome.me/ocp/ca/kasthuri11/info/",
  "ServerType": "openconnectome",
  "ServerName": "openconnectome.me",
  "ServerFolder": "/ocp/ca/kasthuri11",
  "SourceDataSizeX": 21504,
  "SourceDataSizeY": 26624,
  "SourceDataSizeZ": 1850,
  "TargetDataSizeX": 10747,
  "TargetDataSizeY": 12895,
```

⁶See for example <https://arxiv.org/abs/1306.3543>. Essentially VAST requests zipped [128x128x16] pixel blocks of the data set from the data server with URLs which specify the requested region, like: <http://openconnectome.me/ocp/ca/kasthuri11/zip/6/1,129/1,129/1,17/>. The received file is then unzipped to extract the image data.

```

"TargetDataSizeZ": 1850,
"OffsetX": 0,
"OffsetY": 0,
"OffsetZ": 0,
"OffsetMip": 1,
"TargetVoxelSizeXnm": 6,
"TargetVoxelSizeYnm": 6,
"TargetVoxelSizeZnm": 30,
"TargetLayerName": "Kasthuri1110OpenConnectome"
}

```

VAST comes with several pre-defined .vsvr files which you can use to open and view some example data sets.

3.2 Viewing and Navigating an Image Stack

You can open .vsv, .vsvol, .vsvi, or .vsvr data sets by using 'File / Open EM Stack ...' from the main menu. Opened files will be added to a list under 'File / Open Recent EM Stack', from where you can quickly access them again. The list contains the 16 most recent image stack files.

You can also open .vsv (and .vsvol, .vsvi, .vsvr, .vss, .vsseg) files by drag-and-drop from a file browser (Windows Explorer) onto the VAST window.

VAST currently has 'Move', 'Paint', 'Collect', 'Pick', 'Fill' and 'Split' modes, which you can set by clicking the tool buttons in the toolbar. The cross of arrows icon selects 'Move' mode and the little pencil selects 'Paint' mode. In this section we will explain how to use the 'Move' mode. For an explanation of the other modes please refer to section 3.3.

The easiest way to navigate in the image stack is by using the mouse in 'Move' mode. You can pan (move the image sideways) by left clicking and dragging it. You can use the mouse wheel to zoom in and out. Alternatively, you can zoom using the N and M buttons, or the sidebar (see below). Use the UP and DOWN arrow keys or A and Z to scroll through the slices of the stack, or the sidebar to scroll more quickly.

3.2.1 The sidebar

VAST provides a sidebar for quickly zooming and moving through the stack. The sidebar is a region close to the left and the right edge of the main window. When you move the mouse cursor to the left or right edge of the window you will see it appear as a transparent white overlay strip.⁷ Clicking into the sidebar and dragging the mouse up or down will scroll through the slices of the stack (left mouse button) or zoom (right mouse button). If you move the mouse cursor too far away from the side of the window, the view will 'jump back' to the previous view. If you move the mouse cursor very close to the top or bottom of the window while scrolling (not zooming), VAST will start to scroll continuously, with a speed depending on mouse cursor position. You can use this function to quickly scroll through a very large image stack.

3.2.2 The RAM (main memory) usage indicator

At the right side of the toolbar you can see a little field named 'RAM:' which shows the current RAM (random access memory) usage in your computer. The blue frame indicates how much RAM VAST will use maximally for segmentation- and image cache combined. Make sure that this frame is not dedicating more than 2/3 of your total RAM (you can adjust these settings in the Preferences, see section 2.2.2). The solid blue block shows how much RAM VAST has currently allocated for

⁷You can set the opacity of the sidebar in the Preferences, under 'Side Bar Opacity'

segmentation and image cache. The light blue area shows how much memory is allocated by VAST for other purposes. The green area shows how much RAM the Windows system and other programs are using. The colors will change to yellow if the total memory usage goes above 90%, and red if they go above 96%. Running out of available RAM can slow down your system significantly. However, in some cases Windows uses large amounts of the available RAM for disk caching and can free those instantly if more memory is needed by programs without affecting the system performance.

3.2.3 Getting and setting coordinates

VAST uses a coordinate system with a zero point in the upper left corner of the first slice, with positive X to the right and positive Y down in the slice, and Z marking the slice number (front-to-back).⁸ Coordinates are given in pixels at full resolution (the coordinates are independent of the mip map displayed). The coordinates displayed in the upper left corner of the main window show the current location of the center of the main window. You can switch the displayed coordinates on and off by using 'Info / View Coordinates' from the main menu. Zooming in or out will not move the center point of the window and therefore also not change its coordinates. Getting or setting coordinates will also use the coordinates of the center of the screen, as do the 'anchor points' of segments (see section 3.3). While you drag the slice with the mouse VAST displays a semi-transparent cross which indicates the location of the center.⁹

Once you load an image stack, a tool window labeled 'Coordinates' will appear in the upper right corner of the main window. If the tool window is not displayed you can open it using 'Window / Coordinates' from the main menu. It shows you the current center coordinates and allows you to read and set these values. The edit field in the tool window is updated as you navigate through the stack. To save the current location, simply copy the coordinates from that text field (mark with the mouse and press CTRL-C), then paste it into the text editor of your choice. You can also set the coordinates by entering or pasting numbers here and pressing Enter. VAST will then jump to the new coordinates. The exact format of the string does not matter; VAST simply looks for the first three numbers in the string. VAST does not mind whether there are commas or brackets or other non-numerical characters.

This function is quite useful if you want to store coordinates of interesting points in an external text file or spread sheet. Please keep in mind that the coordinate denotes the center of the current view. The center is indicated by semi-transparent crosshairs when you pan the view. You can also center any point by right-clicking that point with the mouse in 'Move' mode and selecting 'Center' from the context menu.

The dropdown-listbox in the Coordinates tool window lists the up to 64 most recent locations you visited. A new entry is added every time you pan the view (but currently not if you scroll through Z). You can go back to previous locations by selecting the coordinates from this list.

3.2.4 Layers

VAST can open several image and segmentation stacks at the same time, provided that they have the same stack size. Each stack is listed as a 'Layer' in the 'Layers' tool window. The order in the list defines the order of the layers. Layers BELOW in the list are 'in front'. The segmentation layers are always rendered on top of all image stack layers. You can change the order of the layers by drag-and-drop in the list. If you can not see all layers in the list, increase the size of the tool window by dragging a corner.

⁸The logic behind the front-to-back order of the image stack is that the images represent sections which were cut from the front surface of a tissue block, and therefore the first section should be closest to the observer for correct geometry. This assumption is used for 3D viewing and model export.

⁹You can set the opacity of the center cross in the Preferences, under 'Center Cross Opacity'

One image layer, and, if available, one segmentation layer is the selected layer of its kind, and its name is displayed in bold text in the 'Layers' tool window. One of these, the one that was clicked last, is also highlighted, and is the layer for which the 'Layer Properties' are shown in the 'Layers' tool window, below the list of layers:

- 'Solo': If this function is enabled, only the currently selected layer will be displayed.
- 'Editable': Uncheck this for segmentation layers you do not want to accidentally paint into. Disabled for image layers
- 'Visible' (image stacks only): Transparency value for this layer. Switch off to hide layer.
- 'Bright' (image stacks only): Image Brightness; switch on to enable brightness control with the slider
- 'Contrast' (image stacks only): Image Contrast; switch on to enable contrast control with the slider
- 'Alpha' (segmentation stacks only): Opacity value for this layer. Switch off to hide layer. Duplicated in toolbar.
- 'SelAlpha' (segmentation stacks only): Opacity value for selected segments (and children). Duplicated in toolbar.
- 'Pattern' (segmentation stacks only): Contrast of segment patterns. Duplicated in toolbar.

VAST can blend image layers with different transparency modes. Click on the button 'Menu' to access more layer options. Under 'Blend Mode', you can select either 'Blend' for alpha-blending, 'Additive' for additive blending, or 'Subtractive' for subtractive blending. The different settings for the transparency computation are:

- 'Flat': All pixels in the image share the same transparency [Default]
- 'Dark Transparent': The darker a pixel $((R+G+B)/3)$, the more transparent it is
- 'Bright Transparent': The brighter a pixel $((R+G+B)/3)$, the more transparent it is
- 'Max(RGB) Dark Transparent': The darker a pixel $(\text{Max}(R,G,B))$, the more transparent it is
- 'Max(RGB) Bright Transparent': The brighter a pixel $(\text{Max}(R,G,B))$, the more transparent it is

For RGB image stacks, 'RED Channel Target Color', 'GREEN Channel Target Color' and 'BLUE Channel Target Color' let you swap and disable different color channels individually, and even assign an arbitrary mixed color to each channel for more sophisticated color space transformations. This is particularly useful for mapping colors in fluorescence microscopy image stacks.

'Color Filter ...' will open a color selection dialog where you can choose a color by which the layer images should be (subtractively) filtered for display. To not filter the images, choose white (255,255,255) [Default].

For image layers, if you enable 'Show Color Clamping' in the Layers window menu, VAST will show which pixels reach saturation while the 'Bright' and 'Contrast' sliders are moved, by inverting the brightness (by color channel) for saturated pixels. This can be helpful for example when part of the image should become fully transparent in the 3D viewer by setting those pixels to saturated black or white.

Segmentation layers can also be used as masks to make part of the displayed image stack black or white. This can be useful for example to show only the interior organelles within a segmented neuron in the 3D viewer. To enable masking, select the segmentation layer in the Layers window, and select the desired masking mode in the Layers window menu under 'Blend Mode'.

3.3 Painting

The main function currently provided by VAST is painting of segmentations as a colored overlay of the image data. When a stack of EM images is loaded, you can enter 'Paint Mode' by clicking the little pencil icon in the toolbar. When you start a new segmentation like this, VAST will ask you if you want to add 16 segments (label colors) to your segment list. Also, two floating tool windows will appear at the right side. The upper one, 'Drawing Properties' (Figure 3.4), provides options for drawing, whereas the lower, 'Segment Colors', lets you select and organize the segment labels and their colors in the segmentation.

To add a new separate segmentation layer, click 'Menu' in the 'Layers' tool window and choose 'Add New Segmentation Layer'. Each segmentation layer holds a separate set of segment colors and is (or will be, when saved) associated with its own segmentation file.

When in paint mode, you can paint on top of the currently displayed EM image. Select a color (label number) from the 'Segment Colors' window at the right by clicking on it. Then click the left mouse button where you want to paint in the image.¹⁰ You will see the outline of your current tooltip as a circle. By clicking and dragging the mouse you can paint larger regions. All painting happens in an overlay plane, the segmentation layer, which is blended over the EM image (the EM image itself will not be changed). You can use the 'Alpha:' checkbox in the toolbar to switch the selected segmentation layer on and off, and the slider right of it to set its opacity.¹¹ Most colors are not solid colors, but have patterns. Use the 'Pattern:' checkbox in the toolbar to switch patterns on and off, and use the slider right of it to manipulate the contrast of the patterns. If you enable the 'SelAlpha:' checkbox, the opacity of the selected segment and its children will be controlled separately by the SelAlpha slider. You can use this to highlight a particular segment or set of segments. You can also switch the EM image layer on and off, by clicking on the 'EM' checkbox in the toolbar. This is sometimes useful if you want to inspect just the segmentation. These controls are duplicated in the 'Layers' tool window for the selected layer.

You can change the size of the pen tooltip. The easiest way is, if you are using a pen tablet and VAST is properly configured, to hold down one of the pen buttons and to move the pen up or down on the screen. When using a mouse, you can hold down either the 'Tab' key or the '\|' key and click and drag up and down instead. You can also use the - and + keys on the keyboard. The current pen diameter (in pixels) is displayed in the Drawing Properties tool window. The third way of changing the tooltip size is to edit the 'Pen Diam.' text field in the tool window. You can also lock the current tooltip size if you don't want it to be changed accidentally, for example if the size of the tooltip is important for your data analysis, by switching on 'Lock'.

The checkbox 'Fill' next to it switches automatic filling of closed contours on and off. If enabled, after each paint stroke VAST checks a rectangular area with the approximate extent of the stroke for empty closed contours of paint color and fills them with the paint color.

Below you can choose from 'Paint All', 'Background' and 'Parent'. This determines which voxels in the current segmentation are paintable. If you select 'Paint All', you will paint over or delete anything, no matter if it was painted before or not. If it is set to 'Background', previous paints will not change, but your paint will only be applied to voxels which have not yet been painted. If you erase, only the current paint color will be erased to empty (background). This is the most useful painting mode.

Instead of only affecting background pixels, 'Parent' mode will affect only pixels which have the color of the immediate parent of the current paint color (see below for a description of segment hierarchies). When erasing, voxels with the current paint color will be changed back to the parent color. This mode is only useful in special cases, in particular when re-labeling a previously painted area to a new color, for example when splitting a labeled object into a hierarchy of parts.

¹⁰Even though you can use VAST with a mouse, it is designed to be used with a pen tablet.

¹¹You can also temporarily hide all segmentation layers by holding down the H key.

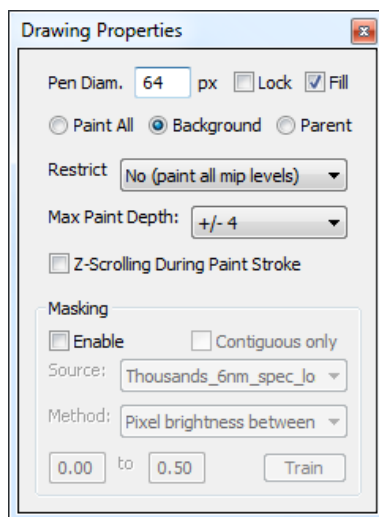


Figure 3.4: The Drawing Properties tool window

3.3.1 Multi-scale painting

A specialty of VAST is that it allows you to paint at different resolutions. In fact, VAST limits you to always paint at the currently displayed resolution. The advantage of this is that the amount of data that has to be manipulated when you paint a stroke is limited by the window size and screen resolution. Otherwise, painting in very large volumes could easily lead to a situation in which the amount of data that has to be written for a paint stroke is much larger than what can be loaded in RAM at one time, which would cause all sorts of problems, including very slow painting. Also it does not make sense to paint at a resolution which is much higher than the screen resolution because mouse (or pen) precision is also limited. Finally, allowing low-resolution painting can save a lot of memory, if large objects are painted coarsely.

In VAST, images are stored as a pyramid of mipmaps with reduced resolution using powers-of-two factors. Painting always happens at the resolution of the currently displayed mipmap. This means that you can change the resolution at which you are painting by zooming. A single segmentation can be composed of parts at different resolutions. For example it is possible to draw a rough outline of an object at a low resolution, and then to zoom in and correct the object's shape at a high resolution. VAST will automatically upscale and downscale the displayed segmentation as you zoom, but zooming will not change the painted segmentation. The segmentation is stored at the resolution at which it was painted. If you paint at a low resolution first and then correct at a high resolution, part of the low-resolution segmentation will be replaced by a high-resolution version. If you paint at a high resolution first and then correct at a low resolution, part of the high-resolution segmentation will be replaced by a low-resolution version, including pixels in the vicinity.¹²

Sometimes you might want to make sure that a painted segmentation has a certain resolution. You can enforce painting at only one resolution by restricting painting to a particular mipmap ('Restrict' in the Drawing Properties tool window). VAST will then enable painting only when the image stack is zoomed to display the selected mipmap.

¹²I have to do this because at the time of painting at a low resolution, not all higher-resolution images may be available in RAM (they may even be too large to be loaded in RAM).

3.3.2 Automatic Z-filling

The time that has to be spent to manually paint a segment in VAST depends largely on the number of 2D outlines that have to be drawn. Especially if you follow a process that runs vertically through the volume, you have to paint (almost) the same outline over and over again, for every slice. If you want to just get a rough outline of an object and you're not interested in a high precision of the boundary, you could increase the painting speed by a factor of n if you paint the outline only in every n -th slice, or paint n slices at a time. In the first case, you get gaps of $n - 1$ slices between the painted outlines, in the second case it is hard to determine what you are actually painting because you can't see where your color goes in most of the slices.

VAST uses a third method. It supports automatic Z-filling of intermediate slices where the regions of the lower and the upper painted region overlap. It turns out that in most cases neuronal objects are *locally convex*. Exceptions are branches, for example when a spine neck runs very close to the dendritic shaft. Automatic z-filling will only fill in the volumes of the overlap between the specified painted regions, and in most cases (for convex objects) the filled regions will stay inside the segmented object.

Z-filling makes sense across a few slices only, because there will be no overlap if your object moves too much from slice to slice (runs at an oblique angle). Also, VAST has to load multiple slices in RAM to be able to fill in those slices. The maximal distance across which VAST lets you fill in depends on the size of the image cubes used. Currently the cubes are set to be 16^3 voxels large, and VAST allows you to fill in up to ± 8 slices (because it makes sure that two layers of cubes are loaded at the time of painting). In the data sets we are using this is approximately as far as z-filling makes sense, and it speeds up painting by a factor of up to 8.

You can set how far the z-filling will reach by setting 'Max Paint Depth' in the 'Drawing Properties' tool window. This value controls both the distance at which Z-filling occurs, and the stepping distance for navigating with S, X or PageUp, PageDown keys, to ensure gap-free painting.

Automatic Z-filling is only applied while painting, not when erasing. This makes it easier to correct what has been filled in in the case of non-convex neighborhoods. This also means that the best strategy to draw an object coarse-to-fine is to try to paint conservatively (try to stay within the object boundaries), and correct by adding paint rather than removing paint, because deleting has to be done in every slice individually.

'Z-Scrolling During Paint Stroke' is by default disabled to prevent painting errors when accidentally switching to the next slice before the paint stroke is finished. However, if you enable it, you can very quickly coarsely label a long neurite running through your stack vertically by scrolling through the stack while following the neurite with your pen – provided that loading of the image stack keeps up with the update rate of the screen. You can adjust the scroll speed with the mouse wheel if you enable 'Adjust Scroll Speed With Mouse Wheel While Z-Scrolling' in the Preferences (see section 2.2.2).

3.3.3 Masking

The last section of the 'Drawing Properties' tool window handles the settings for 'masking'. If you switch on masking by clicking the check box 'Enable', only pixels will be painted for which the image in a selectable source layer fulfills certain criteria. You can choose from several methods which determine the paintable pixels depending on whether the brightness or color of the image pixel in the layer set under 'Source' is in a certain range, which you can set. The value range for minimum and maximum brightness is $[0..1]$. For absolute brightness range modes and 'Trained Color' (see below), the source layer is restricted to be an image layer (not a segmentation layer).

If 'Contiguous only' is checked, VAST will restrict painting to pixels which fulfill the 'masking' criteria and are topologically connected to the place where the pen is put down, by performing a masked flood filling operation. This is for example important for semi-automatic segmentation with

boundary maps (see below).

The method 'Use Trained Color' lets you interactively set which color or brightness range to paint from examples. If you click the 'Train' button, a small tool window will appear which lets you set and clear positive and negative examples from painted segments. This currently only uses the brightness / color statistics of single pixels and can not learn patterns or textures. To use this function, paint over some image area which you want to have labeled in one segment color, select that segment color and click 'Add Positive from Selected Segment'. Then use a different segment color to label some image area which should be excluded from painting, make sure that segment color is selected, and click 'Add Negative from Selected Segment'.¹³ Then set the 'Method' to 'Use trained color'. Subsequently VAST will only paint over pixels which have brightness or color similar to the positive example pixels. If you add positive or negative examples several times, the new examples will be added to the set of already defined examples, until you clear the training data.¹⁴

3.3.4 Semi-automatic segmentation from boundary maps

Masked painting with a fixed brightness range as described above can be used to constrain painted areas by using a boundary map, so that the outline of painted objects is traced automatically.

A boundary map is a gray level image stack which, for each pixel, encodes the likelihood of the pixel being on a boundary as a brightness value. Such maps can be computed for example by using trained convolutional networks on the original EM stack and then loaded into VAST as a separate image layer. To use the boundary map for painting, masking is enabled with the 'Contiguous only' option, and the boundary map is selected as source layer. The 'Method' is set to 'EM pixel between' and the brightness range is set so that boundary pixels are excluded from painting. Then the tooltip size is increased to cover the cross-section of the to-be-segmented object, and then a single painting click on the cross-section should label the object in that section with correct outline. If the boundary map is imperfect, spills can be easily corrected by manual deletion.

The boundary map layer does not have to be visible for this to work. You can hide it and just show the EM image layer, and still use the boundary map layer as source for masking.

3.3.5 Semi-automatic segmentation from an imperfect automatic segmentation

The 'Masking' section of the the Drawing Properties tool window also provides a mode in which an (imperfect) automatic segmentation can be used to help manual drawing. For this you have to open an automatic segmentation as a separate segmentation or image layer. In the 'Method:' drop-down list, select 'Autopick source color', and then select the automatic segmentation layer in the 'Source:' drop-down list. Then, every time you start a pen stroke, VAST will sample the color in the source layer (segment id or pixel color or gray level, depending on source layer type) where you put down the pen (the center of the pen tooltip) and subsequently restrict painting in the active target layer to pixels of exactly that color in the source layer. Enable 'Contiguous only' if the same source color occurs in several segments and you want to constrain painting to the clicked one.

So for example if the cross-section of the object which is painted is correctly segmented in the source layer, that cross-section will be copied from the source to the target section perfectly if it is painted with a large enough tooltip which covers the whole cross-section, or the outline is followed roughly with the pen.

The option 'Autopick nonzero source' in the 'Method:' drop-down list does the same as above, but excludes the background (zero) from being picked as a masking color.

¹³VAST only scans the pixels in the selected segment color which are in the currently displayed slice and region.

¹⁴Internally VAST uses a 64³ HSV cube for example storage and color space mip maps for generalization.

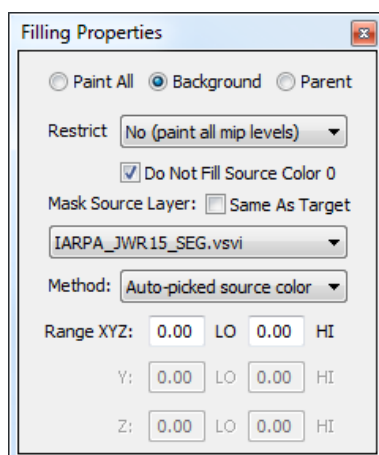


Figure 3.5: The Filling Properties tool window

If 'Z-Scrolling During Paint Stroke' is enabled, you can very quickly draw vertically running processes with perfect outline if you follow the process with a large enough pen while scrolling through the slices. Splits in the source segmentation can be easily corrected by simply lifting the pen and putting it down again within the region of the second source color. Mergers can be corrected by using different target colors for different parts of a source object and accurately drawing the boundary between the objects with the pen.

Using this mode for correcting an automatic segmentation or parts of it has the additional advantage that everything which is painted ends up in a separate layer, and one can be sure that everything in that layer has been looked at. Also the fact that this method works with source segmentations loaded as image layers allows for automatic segmentations to be used which have more than 64k segments. Arbitrary many segments can be converted to 24 bits by throwing away bits and loaded as an RGB image layer. 24 bits would make it very unlikely that two segments with originally different segment numbers, which were now converted to the same segment number, would be located very close to each other. In the rare case that this happens it could be treated in the same way as other mergers.

3.4 Filling

After a segmentation layer is loaded or generated in VAST, the Fill tool will be enabled. To use it, click the paint bucket icon in the tool bar ('Fill Mode'). The 'Filling Properties' tool window will appear.

The fill tool performs constrained three-dimensional filling operations on your data set from a seed point. This can for example be used for recoloring already segmented objects, for splitting merged objects, for supervoxel agglomeration from a source to a target layer, and for filling out complete objects on a reliable boundary map (see section 3.3.4 and below).

VAST loads the parts of the data set which are reached by the filling operation block by block using the caching system, which means that this operation should also work on very large data sets with limited amounts of memory usage.¹⁵

¹⁵There is some overhead since the sides of active blocks are stored in memory for continuing the fill operation into adjacent blocks, but at no time does the whole dataset or object have to be stored in memory at once.

3.4.1 Filling segments in the same layer

If the 'Source Layer' option 'Same As Target' is checked, filling operations will be performed within the selected segmentation layer. The object which is clicked will be flood-filled with the paint color in 3D from the clicked seed point, *at the current mip level*. Just like painting, also filling operations are performed at the resolution of the currently viewed mip level. You can restrict filling to a particular mip level using the drop down menu next to 'Restrict'. This field is linked to the equivalent field in the 'Painting Properties' tool window. The 'Paint all', 'Background' and 'Parent' options are ignored in this mode, since 'Background' and 'Parent' modes would not work. Filling will always change the region of the color clicked to the paint color.

The filling operation uses a 6-neighborhood (fillable voxels at the six sides of a filled voxel will also be filled). If a source object consists of disconnected regions, only the region connected to the seed point will be filled. This can be used to split off parts of an object, for example if two dendrites are accidentally merged. To do this the two parts have to be disconnected topologically (for example by erasing the boundary between them in Paint mode, or by over-painting the to-be-split-off object with the target color, for example using Parent mode), and then filling one side with the target color.

The 'Background' color can also be used as a paint color, to erase objects.

'Do Not Fill Source Color 0' prevents from accidentally filling empty areas.

3.4.2 Filling segments from a source layer to a target layer

Similar to the semi-automatic segmentation functions described in sections 3.3.4 and 3.3.5, the fill tool can also be used to semi-automatically generate a proof-read segmentation in a target layer, by constraining filling by a separate source layer. This adds powerful functions to VAST which can considerably speed up segmentation.

To use a separate source layer, uncheck the 'Source Layer' option 'Same As Target' in the 'Filling Properties' tool window. This will enable additional filling options. First of all, the 'Paint all', 'Background' and 'Parent' options are now accessible, which constrain painting to the target layer and can protect previously painted segments (see section 3.3). These are linked to the equivalent options in the 'Painting Properties' tool window.

Below 'Source Layer' you can choose which layer should be used as source for a mask to constrain filling. 'Method' defines how the mask is derived from the source layer. The following methods are available:

- 'Pixel brightness between': Only pixels which have a brightness in the range defined in 'Range XYZ' in the source layer are fillable. The range values have to be between 0 and 1. This method can be used to perform filling which is constrained by a boundary map (see section 3.3.4).
- 'Pixel brightness XYZ': Same as 'Pixel brightness between', but different value ranges can be defined depending on fill direction (see below).
- 'Autopick source color': Every time the fill tool is applied, the color in the source layer at the seed point will be sampled. If the HI and LO values are 0, filling will be constrained to pixels which have the exact picked color in the source layer. Nonzero values of HI and LO allow filling of pixels with a source layer color in a range around the picked color. This method can for example be used to 'copy' segments from an (imperfect) segmentation layer to a proofread target segmentation (see section 3.3.5), or to copy segments from one segmentation layer to another.
- 'Autopick source XYZ': Same as 'Autopick source color', but different value ranges can be defined depending on fill direction (see below).

- 'Hybrid recolor / autopick': In this mode, if an existing segment in the target layer is clicked, it will be recolored as explained in section 3.4.1, and if an empty area is clicked it will be filled as in mode 'Autopick source color' above.

The XYZ options above can help in cases where a boundary map is used with an image stack which is imperfectly aligned or the Z resolution is low. If consecutive Z sections are shifted more than the width of the boundary, the fill color can 'leak out' into a neighboring segment from one section to the next. To prevent this, boundary maps can be blurred in each XY slice, and the pixel brightness range can be set much more conservatively for filling in Z than for filling in XY. This would let the fill color spread to the next section only where the pixels are far enough away from boundaries, so that leaks are prevented, but in each layer that is reached the fill color could still spread in XY much further to the actual boundary.

3.5 Splitting

The split function in VAST allows to add planar split boundaries to existing segmented objects. For example, spines on a dendrite can be made separate objects by using this function to introduce a straight, minimal-cross-section split boundary in a child color between the dendrite and the spine, and then use the child color to fill the spine.

The split tool can be enabled by clicking the corresponding icon in the tool bar (box with diagonal line and arrow pointing to the right). Enabling it will open the 'Splitting Properties' tool window, in which some options can be set. The split function can currently work in two modes, selectable in the drop-down menu. Either the selected segment is used to paint on its immediate parent ('Use Selected Segment On Parent'), or each click adds a new child segment to the parent clicked, and uses it immediately for the split boundary ('Append numbered child'). In the second case a child segment name and appended number counter can be specified.

To use the splitting tool, click-and-drag at the location where you want to introduce a split within the parent segment, in the 2D view. Dragging will let you define the orientation of the split boundary in the 2D cross-section. When you release the mouse or pen, the split will be applied. At this point VAST will search through planes which intersect the line you specified in 2D, to find the angle at which the plane has the minimal (fully-contained) intersection area with the parent object, and then paste the child color into the parent segment along this plane to separate the two sides. One side can then be filled with the child color to make it a separate sub-object.¹⁶

While you are dragging the split line in 2D, you can hold down the 'Delete' key (typically mapped to a pen button) and release the mouse button or pen to cancel the split. The split line indicator will disappear when you do that. You can also hold down the 'Tab' or '\|' key to translate rather than rotate the split boundary (typically mapped to the other pen button). When in translation mode, the split indicator will appear in green.

A few caveats:

- As with other functions in VAST, the split boundary is drawn at the currently displayed mip resolution. Zoom to the desired resolution first, then introduce the split.
- The search area for the split boundary has a limited size, so this function will fail for very large cross-sections.

¹⁶Be careful when doing the filling – For example, if you fill a spine that you separated with a split plane, but the spine touches another spine, then the fill can fill one spine up, jump to the neighbor spine, fill that spine back down and end up filling the whole dendrite!

3.6 Segments

Segments (segment colors) in VAST are the components of a segmentation. Each segment has a unique ID number which is used to mark voxels as part of this segment in the segmentation images. It also has a number of additional parameters, in particular its color and pattern (see section 3.6.13), its name or label, its position in the segment hierarchy (see section 3.6.2), its anchor point (see section 3.6.6) and its bounding box. The segments which are defined in the currently selected segment layer are listed in the 'Segment Colors' tool window. Segments which have been used are marked with a * at the beginning of their name.

3.6.1 Picking segments

You can select the segment color to paint or fill with in the 'Segment Colors' tool window by clicking on it in the tree view. You can also pick any color you see in the segmentation layer by using the pipette tool. To do this hold down the SHIFT key and click on the segment you wish to select. This makes it very easy to switch between segment colors while painting. Alternatively you can use the 'Pick Segment Mode' which you can select in the main toolbar. If you have several segmentation layers, you can pick colors in a different segmentation layer as well, and VAST will automatically make that segmentation layer active. If segment colors in several layers overlap, VAST will cycle through the layers if you pick several times.

If you hover over segment colors in the main window when in picking mode, VAST will display the name of the segment as a tooltip. If the segment is in a different segmentation layer, the name will be shown in brackets together with the layer name.

3.6.2 The segment hierarchy

VAST can arrange segments in a tree-like hierarchy. This means that each segment can have other segments as children, which can themselves have children, and so on. VAST also allows you to collapse and expand parts of the tree dynamically, which switches between a visualization which shows a whole branch of the tree in the same color or individual sub-branches in individual colors (see section 3.6.4). For example, if all spines of a spiny dendrite are labeled as sub-objects (children) of the dendritic shaft, one can instantly flip between a display in which the whole dendrite has the same color, or each spine has a different color, by opening and closing the dendritic shaft folder. Segments can also be used as folders to group segments, for example to classify labeled objects. You can use tags to designate certain segments as folders, to help external analysis (see section 3.6.12). The grouping can also be applied when segmentations are exported.

Segment hierarchies are visualized and edited in the 'Segment Colors' tool window. This window uses a 'tree view control', similar to the navigation pane of a Windows Explorer window, which makes usage very intuitive. More functions related to the segments can be found in the tool window's menu, which opens either by clicking the 'Menu' button or right-clicking a segment in the tree.

3.6.3 Re-ordering and moving segments in the tree

To re-order the segments, simply drag and drop them with the mouse. You can only select and drag one segment at a time,¹⁷ but if the segment has children the whole branch will be moved (including all children). Please note that to make a segment the first child of another segment, you have to drag it to the right side of the tool window, right of the new parent segment. The new parent segment will then be highlighted in blue, instead of the black line indicating the target space between two segments. You can move any segment, with exception of the 'Background' segment. The Background segment can also not have any children.

¹⁷This is a limitation of the Windows Tree-view control which is used here.

Because it can be cumbersome to move hundreds of items from one folder to another one by one, VAST has two functions 'Make all siblings children' and 'Make all children siblings' which can help in certain situations (see section 3.6.7).

3.6.4 Collapsing and expanding tree branches

You can collapse and expand tree branches, which are displayed in the same way as folders and subfolders are in the Windows explorer, by clicking on the little '+' or '-' sign left of parent segments. When you collapse a folder, all its children will be displayed in the same color as the parent. If you pick a segment color from the segmentation layer by shift-clicking, and the selected segment is in a collapsed folder, the folder will be automatically expanded to show the native color of the segment you selected. In addition the context menu has two functions to set the collapse state of all subfolders of the selected segment at once, under 'Expand / Collapse Child Folders'. If the 'Background' segment is selected, this will be applied to the complete segmentation layer. Also see section 3.6.12.

3.6.5 Adding new segments

The context menu of the 'Segment Colors' tool window provides several functions to add more segments to the selected segment layer. You can add a segment as next sibling or as a child of the selected segment. 'Add 10 Segments' will add 10 segments immediately after the selected segment. 'Add Skeleton Segments' adds a set of child segments to the selected segment which can be used for rudimentary skeletonization. I have not found this function particularly useful though.

The most sophisticated way to add segments is 'Add Named Segments ...', which lets you specify a naming scheme and add multiple named segments at the desired target location in the segment tree. VAST will attempt to guess a naming scheme from the name of the currently selected segment.

You can change the name (label) of any segment in the same way as file names are changed in the Windows Explorer – click a selected segment name a second time, then rename it. Segment names can be maximally 255 characters long.

3.6.6 Using anchor points and other coordinate triplets

Each segment has an 'Anchor Point' stored with it. This is an XYZ coordinate vector which indicates the location of the segment in the image stack. Initially the anchor point is set to the location at which the segment is painted first. You can jump to the anchor point by right-clicking on a used segment in the 'Segment Colors' tool window and selecting 'Go to Anchor Point' from the context menu. You can quickly jump to the anchor point of the selected segment by pressing the 'Home' key or 'G' (for 'go to').¹⁸ You can also set the anchor point of the selected segment to the current view location (as indicated by the center cross) by selecting 'Set Anchor Point' from the context menu. You will have to confirm this action in a pop-up window to prevent accidental setting of anchor points.

Additional coordinate triplets can be stored in the label of each segment. To store the current view coordinates in the label of the selected segment, select 'Add Current Coordinates to Label' in the 'Segment Colors' tool window context menu. You can also copy-paste coordinates as text from the 'Coordinates' tool window or other sources. If those coordinate triplets are framed with round brackets () in the segment label, you can jump to them easily using the context menu. Right-click the segment and choose 'Go to Coordinates in Label', then select the coordinates you want to jump to. VAST supports up to 16 coordinate triplets per segment label, as long as the segment label does not exceed the maximal length of 255 characters.

¹⁸If you pressed the 'Home' key accidentally and want to go back to where you were, you can select the previous location from the drop-down menu in the 'Coordinates' tool window.

3.6.7 Helper functions for arranging segments

Under 'Arrange' in the context menu you can find two useful functions to move many segments at once. 'Make All Siblings Children' will move all siblings of the selected segment into its folder (make them children of the selected segment). 'Make All Children Siblings' moves all children of the selected segment out of its folder and makes them siblings. Be careful with these functions because currently there is no 'Undo'.

3.6.8 Select recently selected segments

Under 'Select Recently Selected' in the 'Segment Colors' tool window's context menu you can find a list of the segments you had recently selected. You can click on one of the listed segments to select it again.

3.6.9 Global operations: Deleting and welding segment subtrees

The context menu of the 'Segment Colors' tool window provides two functions which change segment numbers in the whole segmentation layer, one to delete segments from the segmentation layer ('Delete Segment + Subtree') and one to weld all collapsed folders in a subtree to a single segment each (making them all the folder color, while removing the child segment numbers), 'Weld Collapsed Segments in Subtree'. If the Background segment is selected, *all* collapsed folders in the segmentation layer will be welded to single segments.

Deleting and welding segments is actually more difficult than it seems because it involves traversing the whole segmentation data set and inspecting every single painted voxel. When you choose to delete the selected segment and its children, VAST will actually have to not only set all voxels with those segment numbers to 0, but also renumber all the other voxels so that the used segment numbers will have no 'gaps' (all segment numbers between 0 and n are used).

'Welding' will merge a set of collapsed folders in the subtree of the selected segment into single segments. For this VAST needs to renumber all voxels with segment numbers of children of collapsed folders in the selected segment subtree to the segment number of their collapsed folder, and also, similar to when deleting, renumber the other segment voxels so that the resulting segmentation is free of segment number gaps.

Because these functions change more or less the complete segmentation and VAST is designed so that the opened segmentation file is not changed, it basically has to copy almost all segmentation data to the temporary cache file. Depending on what segmentation you are working on this can take a lot of time, memory and disk space. Also it will change the internal ID numbers of the segments, so please think twice before using these functions in case you are relying on absolute segment IDs in your analysis. These functions are also not well tested; please report any bugs you might encounter.

3.6.10 Global operations: Reevaluating bounding boxes

For each segment VAST keeps a bounding box, which stores the minimal and maximal coordinates of pixels painted with this segment. This can be used to define a cutout region which contains the complete segment. VAST extends this bounding box automatically as a segment grows. However, if a segment shrinks, for example by manual erasing parts of it, the bounding box will not shrink because VAST does not know where else the segment color may be used in the segmentation. This means that bounding boxes can be significantly larger than the object itself, for example if the segment was painted accidentally in the wrong place. For this VAST provides a function to reevaluate the bounding boxes. To use this, zoom the 2D view to the mip level you want to do the evaluation at and select the segment or segment subtree you want to reevaluate (segment and all its children). Then select the function 'Reevaluate Bounding Boxes in Subtree' in the 'Segment Colors' tool window context menu. The higher the resolution which you are using, the longer this process

will take, since VAST has to go through the complete segmentation data at that resolution and find all voxels which have any of the selected segment colors.

3.6.11 Global operations: Segmentation cleaning

When painting segmentations manually, there can be left-over imperfections like holes and stray pixels (for example if part of a segment was incompletely erased). The same can occur in automatic segmentations. You can try to find these for example by using the 3D viewer, but correcting everything manually can be a lot of work. For this purpose VAST provides a function to perform segmentation cleaning. To use it, select **Clean Segments in Subtree ...** in the 'Segment Colors' tool window context menu. A parameter dialog window will open in which you can select the mip level at which to perform the cleaning, which segments to clean, and the cleaning region. Below you can specify the parameters for the cleaning operation itself. Cleaning is performed by loading overlapping parts of the segmentation at a time (a 'search window') and filling holes / removing debris within each windowed region.

'Search Window Diameter (pixels)' specifies the size of the search window. If the window is too small, large pieces of debris and large holes will be missed if they never completely fall within a window. Making the window large will increase its memory consumption.

'Remove Debris up to Volume (pixels):' lets you constrain the size of separate segment regions to be removed, to separate debris from objects which should be kept. If you uncheck this option, debris will not be removed.

'Fill Holes' can also be enabled/disabled, and a volume threshold for hole filling can be set as well.

Please note that with the current implementation of this function there may still remain some pixels and holes in higher-resolution mip levels, which are too small to show up in the mip level used for cleaning. However, if you export images or models at the cleaned resolution, debris and holes should have been properly removed.

3.6.12 Segment tags

Each segment can have a 'Tag' which you can select in the 'Segment Colors' context menu under 'Tags'. A tag is a number between 0 and 15 which can indicate the type of the segment. The tag value is exported with the segmentation metadata text file and can be used to help external analysis. By default the tag of all segments is 0. VAST uses tag 1 to indicate that the segment is a 'Folder Segment', which is not a segmented object by itself but rather a folder which contains other folders and objects. This information can be used to collapse all objects, but expand all folders – select **Expand Only Segments Tagged as Folder** under **Expand / Collapse Child Folders** in the 'Segment Colors' tool window context menu.

Segments which have a tag that is not 0 will have an icon with a colored frame in the Segment Colors tree window. Folder segments have a gray frame.

3.6.13 Editing the color of a segment

The color and pattern of any segment can be changed. Right-click on a segment in the list and go to the sub-menu **Colors** of the context menu, then choose the desired option. Each segment has a primary color, a secondary color, and a pattern that is used to blend between them. You can also randomize the colors of all segments of the selected subtree (the selected segment and all its children), or set the primary or secondary color or pattern of all segments of the selected subtree the same.¹⁹

¹⁹If the Background segment is selected, these functions will apply to all segments in the segment layer (except the Background segment).

Please remember that by collapsing segment folders you can quickly and reversibly switch the displayed color of a segment to the color of the collapsed parent, without the need to change the color of all children individually.

3.6.14 Exporting segment metadata

The entry 'Save Segmentation Metadata ...' in the 'Segment Colors' tool window context menu lets you export the metadata of the segments to a text file, which you can then for example parse with MATLAB to extract colors, hierarchies, names, anchor points etc. of the segments for analysis. A Matlab script which can parse these files is included in the supplementary package (`scanvastcolorfile.m`). The format of the text file is described at the beginning of the text file itself:

1	ID number of the segment
2	Flags field of the segment as 32-bit value
3-6	Primary color as red, green, blue, pattern1
7-10	Secondary color as red, green, blue, pattern2 (unused)
11-13	XYZ coordinates of the segment's anchor point (in voxels)
14-17	IDs of parent, child, previous and next segment (0 if none)
18	If the segment is collapsed into a folder, this is the folder ID
19-24	Segment bounding box (may be incorrect if voxels were deleted)
25	Segment name as text in " "

Table 3.1: Columns of the segment metadata text file

The 'Flags' field uses the following bits: 0: isused, 8: isexpanded, 16,17: isselected (1: selected, 2: in child group of selected)

3.6.15 Segment information

'Segment Information ...' in the context menu opens a window which shows you the internal information (metadata) associated with the selected segment. You can use this information to count children of the segment, get its internal ID or other parameters. Getting the segment info of the 'Background' segment returns statistics for the whole segmentation layer. The text can be copy-pasted if needed.

3.6.16 Searching for a segment with a given name or ID

At the top of the 'Segment Colors' tool window there is an edit field which can be used to find segments. As you type or paste a text string into this field, VAST will select the next segment (after the selected segment in depth-first search order) the name of which contains the typed sub-string. The edit field is case-sensitive. If there is more than one segment which contains this sub-string, you can click on the magnifying glass right of the text edit field to go to the next segment that matches your text. The F3 key has the same function, provided that the segment tree sub-window of the 'Segment Colors' tool window is active.

You can also search for a segment with a particular ID. To do this, type an opening bracket [into the find edit field, followed by the ID number (internal segment number) you are looking for.

3.6.17 The 'Collect' tool

The 'Collect' tool ('Collect Segment Mode') can be selected in the toolbar by clicking the icon with an arrow pointing at a folder. When in Collect mode, segments you click will be moved in the

segment tree to become children of the currently selected segment. This can be useful to quickly classify different objects into different types. Simply make a folder segment for each type, select it, and click on the objects in the image that are of that type with the 'Collect' tool. You can also use this to collect parts of objects in an automatic segmentation into one folder each and weld them (see section 3.6.9) to correct splits.

Using this tool is a bit dangerous because there is no Undo. If you click the wrong object it might be difficult to remember where it came from and there is currently no easy way to 'move it back'. Secondly, when a segment is moved, all its child segments are moved with it, but not its parent(s). So if you are dealing with objects which consist of several parts, make sure that you move the parent segment of the object and not only a side branch of its tree.

3.7 Saving Segmentations

Important: VAST DOES NOT SAVE AUTOMATICALLY while you paint. Your tracings will be held in RAM and/or a cache file on disk until you explicitly save them. If you open a segmentation from a .VSS file and work on it, the file will not be changed unless you explicitly tell VAST to save the changes you made to the opened file by selecting 'Save Segmentation' from the main menu. If you want to keep the previous version and save to a new file, use 'Save Segmentation As ...' instead. VAST will then take all data from the opened file, the RAM cache, and the segmentation cache file, and combine them into a new file on disk.

We have had cases in which people had VAST open for several days without saving and lost a lot of work when the computer crashed. Please save your work once in a while.

3.7.1 Save Segmentation As Special

'Save Segmentation As Special ...' provides you with two functions to save your current segmentation to a new file in a modified way. First, you can choose to save only the selected segment or subtree of the current segmentation to a new file. Alternatively, you can change the resolution of your segmentation and adjust the canvas size on which the saved segmentation will open (this is currently only supported when saving all segments).

To save only the selected segment or the selected segment and its child tree to a new segmentation, select that option from the drop-down list and press 'Save'. Please be aware that the internal IDs of the segments in the new segmentation file will be changed to maintain a gap-free list of IDs from 1 to n for n segments.

The settings in the lower part of the dialog provide limited functionality to adjust the resolution of the saved segmentation and target canvas size.

Normally VSS files only open on top of an image stack of the same size in voxels. You can use this function to save an existing segmentation so that you can open it on a stack which has slightly different size and/or is scaled up or down by a power of two.

Currently the saved segmentation will stay aligned to the upper, left, top corner of the stack. Also, only powers-of-two (X,Y) scaling is possible. If you want to translate the segmentation to a different location or scale in the target stack, you'll have to export the segmentation as an image stack, modify the images accordingly, and re-import (see section 3.9).

3.8 Segmentation Merging

If you have two or more (for example partial) segmentations (.VSS files) of the same image stack you can merge them into a single segmentation. For this, first open one of the segmentations you want to merge in VAST, and then select 'File / Import / Merge .VSS Segmentation Files in ...' from the main menu, and choose the .VSS file(s) you would like to merge in with the opened

segmentation. During merging, VAST will add the selected .VSS files to the current segmentation. VAST lets you choose whether you want to change segment IDs of the merged-in data so that they do not overlap with existing segments, or keep segment IDs of the imported file (then only segments with new IDs will be added to the segment list, and segments with existing IDs will extend pre-existing regions). You can also choose whether to put new segments into a separate folder, and whether or not to overwrite pre-existing voxels or only write into empty voxels. You can select several .VSS files at once for merging, which will then be added one-by-one.

VAST does not save the merged segmentation automatically nor does it change any of the segmentation files. It will generate the merged segmentation in the segmentation cache, and if it does not fit in RAM it will end up in the temporary segmentation file. Please save the resulting merged segmentation if you want to keep it, for example via 'File / Save Segmentation As ...' from the main menu. Please make sure that there is enough space on the drive used for the disk cache.

3.9 Importing Segmentations From Image Stacks

VAST can import segmentations from image stacks, either generating a new segmentation or merging with a previously loaded segmentation. The imported segmentation will be added to the selected segmentation layer, or a new layer will be generated if no segmentation layer exists.

Just like the image files generated during segmentation exporting (see section 3.10 below), the RGB values of the image pixels will be interpreted as segment numbers (the least significant 8 bits (0..7) are in the blue channel, bits 8..15 in the green channel, and bits 16..23 in the red channel). Please keep in mind that VAST can currently only handle 16-bit segmentations, so the red channel should always be 0. However, because of the requirement to have a gap-free list of segment numbers in segmentations, VAST will scan all images for existing segment numbers and renumber them as necessary if it finds gaps or if color channels are used differently than expected. This means that you can use any color mapping for the segments (as long as you don't exceed the 64k limit), and VAST will sort it out for you during importing, but the resulting segment IDs within VAST might differ from what you intended.

To import, select 'File / Import / Import Segmentation from Images ...' from the main menu. VAST will ask you to select one or several image files of the image stack you want to import. Then it will open a dialog where you can specify the parameters for segmentation importing.

File name(s):

You have to specify a name template for all image files in the stack. This has to be a *C++ format string* (as it is used by `printf()`). This means that numbers which specify the slice, column and row coordinates in the file name have to be replaced by codes like '%d' (integer number) or '%04d' (integer number with zero-padding to 4 digits). VAST will then fill in those numbers for each image. For more information about format strings, please refer to a C++ manual or ask the internet. You also have to specify the order and the value limits for the parameters of your file name template, to define the range of names of all the images or image tiles in your stack.

For example, assume you have a single-tile image stack of 100 images in a single folder, and the images are named 'slice_s0000.png' to 'slice_s0099.png'. The file name template should then be set to 'slice_s%04d.png', because the slice counter is an integer number which is zero-padded to four digits. Also, in this case, you would set the 'Parameter order in string:' to 'Slice (Z)', since the only parameter in the template is the slice counter. To get the correct values for the slice counter set 'No of first slice (Z):' to 0 and 'No of last slice (Z):' to 99.

Figure 3.6 shows a more complicated example. Here, segmentation images are tiled and each section is stored in a separate subfolder. The dataset consists of 3400 sections, and each section is stored as 25x25 image tiles of 1024x1024 pixels each. The first tile of the first section

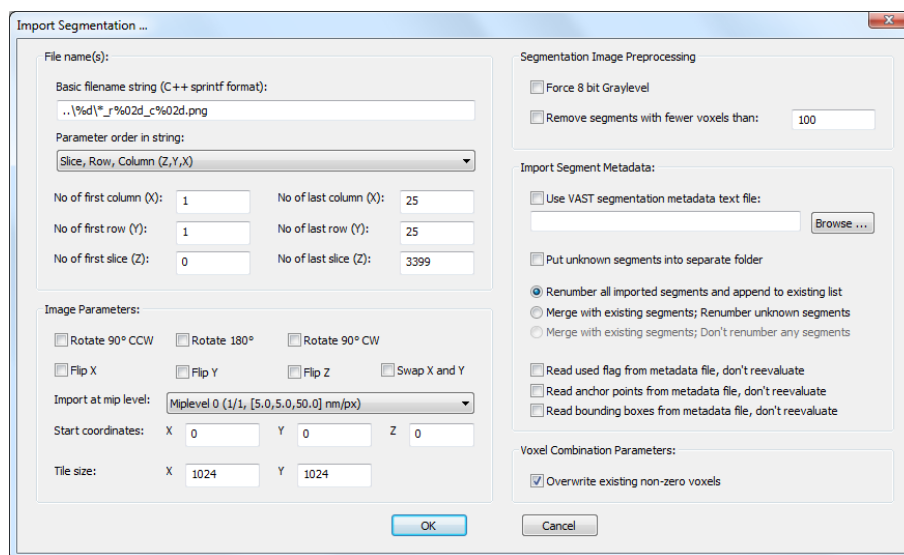


Figure 3.6: Import segmentation options

is named as and located in `.\0\slice1_w01_r01_c01.png`, the last tile of the last section is `.\3399\slice3400_w27_r25_c25.png`. If you select a single image as reference, VAST will assume that directory as start directory, but you can use a relative path in the file name template to go to different directories. `'..\%d\'` in the file name template will go to the correct directory for each slice, and we use the slice counter to fill in this parameter. Next, the beginning of the file name repeats the section counter with an offset of 1, and also has a wafer parameter which we don't need for indexing. We can omit this part of the file name by using a wildcard (`*`). Lastly, the row and column parameters are replaced by `'%02d'` since they are zero-padded to two digits.

Image Parameters:

Here you can rotate and flip the images if necessary, and tell VAST where to put them into the currently opened volume. The 'Tile Size' is extracted from the image file you selected, but you can adjust it here too. The 'Start coordinates' are currently not allowed to be negative. Please crop your images prior to importing if necessary.

Segmentation Image Preprocessing:

The option 'Force 8 bit Graylevel' tells VAST to interpret the images as single-channel 8 bit images, by only using the least-significant 8 bits of the image data. This can help in cases where the loaded images are supposed to be 8 bit graylevel images but are encoded or interpreted as RGB or RGBA color images.

If you enable 'Remove segments with fewer voxels than:', you can clean up messy segmentations upon import by removing small debris segments. All segments with fewer voxels than the given limit will not be imported (and other segments will be renumbered to fill the gaps). If you use this function you currently cannot use a segmentation metadata text file since the size threshold may remove segments which would be needed for the segment hierarchy in the metadata file.

Import Segment Metadata:

If you have a segment metadata file for the imported segmentation stack (same format as the file written in section 3.6.14) you can provide it here. The options below specify how VAST should deal

with the segments in the imported stack. If you are importing into a previously loaded segmentation, you can either keep the imported segments separate by renumbering and appending them to the existing segments, or merge the imported segments which have already existing segment IDs with the existing ones. New segments will be appended. If you use the option 'Merge with existing segments; Renumber unknown segments' together with a metadata file, the hierarchy in the metadata file will be ignored. The three options below let you choose whether you want to import the usage flags, anchor points, and bounding boxes as specified in the metadata file, or reevaluate them from the imported image data.

Voxel Combination Parameters:

If 'Overwrite existing non-zero voxels' is checked, imported non-zero voxels will overwrite existing non-zero voxels. If it is not checked, VAST will keep existing non-zero voxels and ignore imported non-zero voxels at the same place.

Similar to segmentation merging, VAST will not save the segmentation automatically after importing. Please use File / Save Segmentation As ...' from the main menu to save the result.

3.10 Exporting Image Stacks

Exporting EM volumes and/or segmentations as image stacks can be useful to generate slice animations, to transfer data to other programs like Fiji (ImageJ), or to process your results further externally.

To export (parts of) the image stack and/or segmentation stack as a stack of image files, select 'File / Export ...' in the main menu. The dialog shown in Figure 3.7 will pop up. You can export screenshot image stacks, EM image stacks, and segmentation ID image stacks. You can specify a region of the data set to be exported, and a resolution (mip level) for the images. Data regions which are too large for storing in a single image per slice can be exported as a set of image tiles.

Export as:

Choose here if you want to export a single-tile image stack (one tile per slice) or a multi-tile stack (a grid of image tiles per slice). For the second option you can define the tile size to be used.

Region to export (Specify coordinates at full resolution):

This defines the region of the stack that should be exported, for all three targets (Screenshot, EM (image) data, Segmentation). By default it is set to the full stack. You can restrict the export region here. X is the horizontal axis in the slice, pointing to the right; Y is the vertical axis in the slice, pointing downwards on the screen, and Slice (Z) defines the range of slices that should be exported. The first and second columns of the edit fields let you define start and end of the region for each axis, the third column defines the size of the exported region (edit fields change each other to stay consistent).

The 'Full' button will set the values back to the full extent of your stack.

The 'BBox' button sets the region to the bounding box of the currently selected segment. You can use this function to define a cut-out region from a painted segment²⁰, or use a new segment and just paint the upper left top and lower right bottom corners of the cutout region, select it and use the 'BBox' button to set the export region.

The 'Center' button will center the exported region to the current coordinates in the main window. You can for example use this in the following way: First move the image in the main

²⁰Note that the bounding box is not always correct, in particular if you delete parts of what was painted before, the bounding box will not shrink.

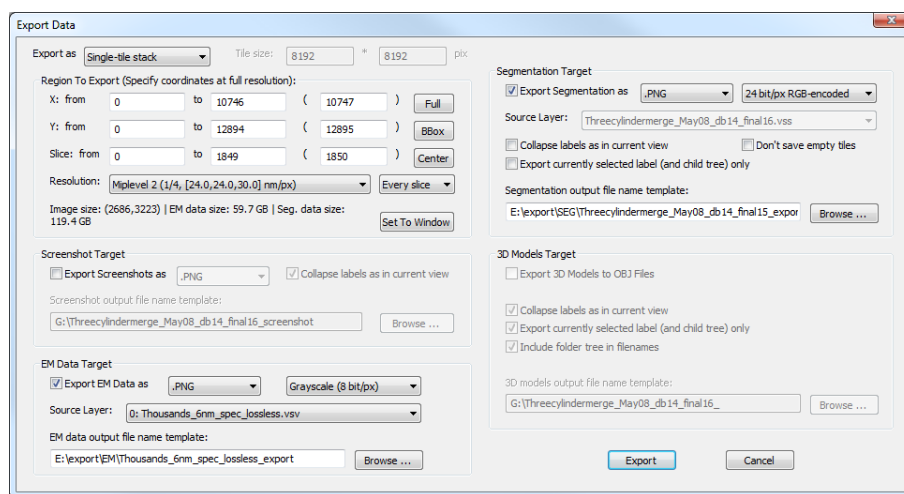


Figure 3.7: Export dialog

window to the location you want to export. Then open the export dialog and specify the size of the image you want to export in pixels, in the third column of the 'Region to Export' section, for example to 1024, 1024, 1 for one image with size (1024,1024). Then click 'Center' to place the export region at the currently viewed location.

The 'Set To Window' button will set the export region to what is currently shown in the main 2D window of VAST.

Under 'Resolution:' you can select at which mip level you want to export. VAST does not support arbitrary scaling, but can export image stacks at its native mipmap scales (which are powers of two). You can also subsample the stack by slices (every n th slice). Below you can see the image size resulting from your settings and an estimate of the (raw) data size that will be exported. Compressed image formats like .PNG can however produce much smaller file sizes, depending on the image content.

Screenshot target:

If you want to export a stack of 'screenshots' how the images look in the main window of VAST, enable the checkbox 'Export Screenshot as'. VAST will reproduce the pattern, blending and tinting settings as they are currently set in the main window in the exported 'screenshots'. Select a target image format and filename prefix / location.

EM data target:

This saves a stack of (EM) image data from the selected layer. You can specify the target format, filename prefix and location.

Segmentation target:

This saves the segment IDs of the currently selected segmentation layer or part of it as an image stack. The segment ID of each pixel (a 16-bit number) will be encoded in the color of the pixel in the exported image. Bits 0-7 will end up in the blue channel and bits 8-15 will end up in the green channel. The red channel will currently stay 0.

3D models target:

This function is currently disabled because it is not fully implemented yet. For now please use Vast-

Tools to export 3D models of segmentations (see section 4.2).

When you are done setting up the parameters for the export, press 'Export' and VAST will start exporting the image stack or stacks (VAST can export more than one target at the same time).

3.11 The 3D Viewer

VAST contains a simple 3D viewer which lets you view a region of your image stack in 3D. To open the 3D viewer, select 'Window / 3D Viewer' in the main menu of VAST.

The 3D viewer uses the image content of the 2D view. Once the image in the 2D window is set up, select 'View / Update' in the menu of the 3D Viewer window to load the data into graphics card memory. The block shown will be more or less centered around where the current 2D view is located (the center cross location), and use the current 2D resolution (mip level). The transparency in the 3D view is derived from the pixel brightness. You can choose to either 'Show bright', 'Show dark' or apply full opacity ('Opaque') by selecting the corresponding option under 'Transparency' in the window menu. To see the effect of the selection, update the data by selecting 'View / Update' in the menu of the 3D Viewer window.

The 3D volume is rendered as a volumetric texture with transparency, rather than surface meshes of objects. This allows for volumetric rendering of all image data, not only segmented objects. Under 'Size' in the menu you can choose from six different sizes for the volumetric texture. The larger the texture is, the more detailed will the 3D model be, but it will also require more memory on the graphics card and update and render slower.

Assume for example you have an EM stack and a segmentation loaded in VAST, and you want to see just one segmented object with its children in 3D. For this, adjust the 2D view so that only that segmented object is shown and everything else is black: Switch off the EM layer and 'Alpha', enable and maximize 'SelAlpha', and select the parent segment of the segments you want to show up in 3D. Center the 2D view to the segment of interest and zoom to the desired resolution. Then, in the 3D Viewer window, make sure that 'Show bright' is selected under 'Transparency' in the window menu, and select 'View / Update' to load just your segmented object(s) into the 3D viewer.

You can have the 3D viewer remember the zoom level and color/transparency settings and decouple them from the 2D window view. The stored settings will be used whenever you update the 3D viewer data to a new location or selection of segments. To hold the current settings, check 'View / Lock Blend Params' and/or 'View / Lock Mip Level' in the menu of the 3D viewer window.

You can zoom the 3D view in and out with the mouse wheel and rotate it in all three axes by dragging with the mouse (the mouse cursor shows which rotation axes will be affected, depending on where the mouse cursor is in the window). Choose 'View / Reset Orientation' from the 3D viewer window menu to set the view back to the orientation of the 2D view.

To save the current 3D scene to an image, select 'View / Save Screenshot ...' from the main menu. You can save the screenshot as PNG, (uncompressed) TIFF or BMP image file.

Further options available in the menu of the 3D viewer include changing the background color and displaying a line cube around the data block.

3.11.1 Using the 3D viewer to set the 2D view location

The 3D viewer can be used to set the 2D view to the same location that was clicked in the 3D scene. To do this, hold down the SHIFT key and click on an object in the 3D scene. VAST will compute the coordinates of the closest point of the 3D scene at the clicked location which is more opaque than a threshold, and set the view in the 2D window to those coordinates. The 2D view will update

as soon as it is selected as active window. The threshold can be set to an opacity of 0.1, 0.5 or 0.9 (0 meaning transparent and 1 opaque) under 'Picking' in the menu of the 3D window.

Chapter 4

The VastTools Matlab Toolbox

VAST includes an API through which it can communicate directly with client programs, using a TCP/IP connection. The Matlab script VastTools, which is included with VAST Lite, uses this API to provide users with a number of additional tools for VAST, including target lists, 3D surface exporting and measurement functions. It can be found in the `vast_package.zip` file. Since VastTools is a Matlab script, users can extend the interface with their own functions. Documentation of the VAST API is provided in Appendix B.3.

4.1 Getting started with the VastTools Matlab Toolbox

To start the toolbox, open `vasttools.m` in Matlab and run it. If Matlab asks, change the current directory. A small window should pop up with a menu, message field and cancel button.

Before you can use any of the functions of the toolbox you have to connect VastTools to an instance of VAST which is currently running. First, in VAST, you must enable the Remote Control API Server. Select 'Window / Remote Control API Server' in the main menu of VAST. In the tool window which then opens, enable a TCP/IP port for communication by clicking the 'Enable' check box in the upper left corner. If you are running VAST and VastTools on the same computer, you can use the standard settings (IP 127.0.0.1 and port 22081 on both the VAST and Matlab side). If you are using VastTools a lot and do not want to enable the API every time manually, you can set a flag in the Preferences to have it enabled automatically when VAST starts (see section 2.2.2).

Then, in VastTools, connect to VAST by selecting 'Connect / Connect to VAST' in the main menu of the VastTools window. If the menu item 'Connect / Connect to VAST' changes to 'Disconnect / Disconnect from VAST' you are connected. The message log in the VAST Remote Control API Server window will also show when a remote connection has been accepted. In case you are using a different IP or port, for example because VAST and VastTools run on separate computers, you can set the IP address and port to use under 'Connect / Connection Options' in the VastTools main menu.

Once connected, you can use the different functions of the VastTools toolbox by selecting from the main menu of the VastTools window in Matlab.

Currently VastTools is not locking the selected layer, view, appearance etc. in VAST during exporting, but many exporting functions use these settings. This means that if you change settings in VAST while exporting through VastTools you may get unintended results. To prevent that it is better not to continue working in the source instance of VAST while VastTools is exporting.

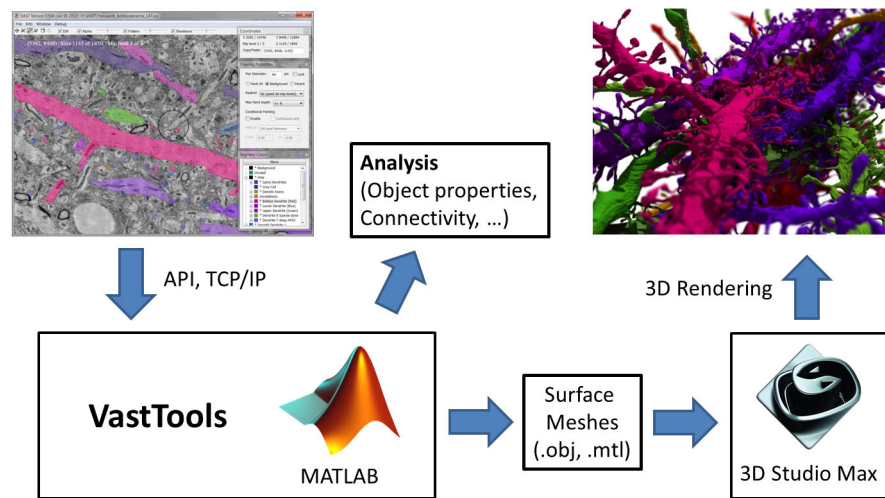


Figure 4.1: Exporting from VAST using VastTools

4.2 Exporting 3D Models

3D model exporting generates surface meshes of the painted voxelized segmentation in VAST, and saves the resulting meshes in Wavefront OBJ files which are widely supported by 3D rendering and animation programs. VastTools uses Matlab's `isosurface` function to generate the meshes. Typically, it will generate surface meshes in parts of the volume to limit memory usage, and then glue together the mesh pieces to generate the final objects. The exporter also saves a .MTL file for each OBJ file which defines the material properties (color) as defined in VAST. You can also measure the surface area of all exported objects and save the results to a text file.

To use this function you have to first run `vasttools.m` in Matlab and connect to VAST, as described above. Then select 'Export / Export 3D Surfaces as OBJ Files ...' from the main menu in VastTools. A dialog window will pop up in which you can specify all the parameters for the export. Once you click 'OK' the exporting will start and run until finished or until you press 'Cancel' in the main VastTools window. This will export objects of the *selected segmentation layer*.

The following parameters for the export can be set:

Render at resolution:

Here you can select the mip level at which the surfaces should be generated. Lower mip levels will result in more detailed models, but also generate larger OBJ files with more triangles and take longer to process. Models and surface computations will be automatically scaled to compensate.

Use every nth slice:

If you want to reduce the resolution in Z you can do so by skipping slices. If you set this value to 2, every second slice will be used, if you set it to 3, every third, and so on. Models and surface computations will be automatically scaled to compensate.

Render from area:

This specifies the region of the current volume in which you want to generate surface models. By default the region is set to cover the whole volume, but you can change the values to cover a smaller region if needed. These values are always specified at full resolution, no matter which mip level is set in the 'Render at resolution' selector.

'Set to full':

Push this button to set the area boundaries to the full extent of the data set opened in VAST.

'Set to selected bbox':

This sets the render area to the combined bounding box of the segment selected in VAST and all its children. CAUTION: The segment bounding boxes are currently not always correct. For example, if you erase voxels you previously painted in VAST, the bounding box will not shrink.

'Set to current voxel':

This button sets the render area to a single voxel, the current center voxel set in VAST. Use together with 'Extend to current voxel' to define a target area by its corner points.

'Extend to current voxel':

This button extends the boundaries of the render area to include the current center voxel set in VAST. Use together with 'Set to current voxel' to define a target area by its corner points.

Voxel size (full res):

These fields will be filled in automatically by the values provided by VAST. The voxel size is used to scale your models correctly. You can override the voxel size manually by entering different values into the edit fields.

Scale models by:

Additional scaling factors for your models. By default these factors are set to 0.001 (1/1000) in all directions to convert the units from nanometers to micrometers.

Model output offset:

By default the models will be placed so that the upper left corner of the data set loaded in VAST is at the origin (0,0,0) of the 3D model space. This means that even if you change the export resolution or restrict the export to a smaller region under 'Render from area', the models will still be placed correctly within the coordinate system of the data set, and models exported from different regions of the same data set will be placed correctly with respect to each other. You can however provide a constant offset here if you want to move the models somewhere else.

Processing block size:

As mentioned above, models are exported in smaller blocks and then glued together; this specifies the block size. A smaller block size will reduce memory consumption and may increase the speed of isosurface computation, but will increase the processing time needed for glueing model parts together. Changing the block size should only have an effect on memory consumption and processing speed, not on the exported models.

Export what:

Here you can select whether you want to export models of all segments, or of a selected branch, and you can choose whether to export all segments individually or glue them together as they are currently displayed in VAST (by selecting a parent segment, and collapsing and expanding folders in VAST). This information will be read from VAST when you press OK, so you can make adjustments in VAST while the VastTools export dialog is open.

The options 'R, G, B isosurfaces from screenshots, at 50%' and 'Isosurfaces from screenshots, 16 brightness levels' (and 32, 64 brightness levels) can be used to generate surface meshes from image stacks directly, without having to segment objects first. You can use the layer mixing and

brightness/contrast settings in the VAST 'Layers' tool window to adjust the result. This can for example be used to export onion-like shells of fluorescence signals from confocal image stacks, to render realistic-looking 3D fluorescence clouds by using surface-normal-dependent transparency and additive blending, as was done in some of the figures in [8].

File name prefix:

All OBJ file names, and also the names of the exported objects, will start with this prefix string.

Object Colors:

Here you can select between 'Object colors from VAST' (the default) and 'Object volumes as JET colormap'. The latter only works if you previously computed object volumes (under 'Measure / Measure Segment Volumes ...').

Target folder:

All generated OBJ and MTL files will be stored to this target folder. Use the 'Browse...' button to select a different folder.

Include Vast folder names in file names:

If enabled, the VAST folder names of the segment hierarchy will be added to all OBJ file names, and also the names of the exported objects. This makes it easy to select and process subsets of the objects based on their names. However, please keep in mind that the length of file names can be limited, depending on your file system, and names which are too long can lead to problems.

Invert Z axis:

If enabled, the models will be mirrored so that they lie below the $z = 0$ plane. This reflects the actual shape of the objects in the tissue, if slices are counted up as they are cut off the surface of a block.

Write 3dsMax bulk loader script to folder:

If enabled, a small file called 'loadallobj_here.ms' will be saved to the target directory. This is a 3dsMax script which, if executed in 3dsMax, will batch-load all objects in the same directory. This is very useful if you are working with many object files.

Close surface sides:

If enabled, meshes will be closed at the sides where they exit the extraction region.

Skip model file generation:

Enable this option if you just want to measure model surface area and not generate OBJ files.

Save surface statistics to file:

If this is enabled, the export script will also compute the surface area of each exported object (by summing up the triangle surface areas) and save the result to the text file with the provided name. This file will always be stored in the same folder as the OBJ files (the target folder), so please give only a file name in the text field, not a file name with target path. Computing the surface area will take additional time, so only enable it if you need it.

4.3 Export Particle Clouds (3D Object Instancing)

This function generates compound 3D mesh models from a reference 3D object (for example, a model of a vesicle) which is placed at different locations, as defined by a segmentation in VAST. Source locations can be either anchor points of objects or the center points of separated painted regions (either handled in each XY section separately, or volumetrically in XYZ).

With this function it is for example possible to make vesicle clouds with spherical vesicle meshes, based on manually annotated vesicles where each vesicle is labeled by a separate dot of paint. The export function places a copy (instance) of the vesicle object in the center of each painted dot. This function can also be used to indicate 3D locations of particular importance with different 3D object glyphs.

Many of the parameters in the 'Export Particle Clouds (3D Object Instancing) ...' dialog are equivalent to 3D object exporting (see above). The following parameters are specific to this export function:

Coords from:

This option defines which segment colors in the selected segmentation layer should be used to define target coordinates.

Coords mode:

Select here whether for each segment, you want to extract the centroid locations of all painted regions in each XY section separately (2D, 4-connectivity), use volumetrically connected region centroids (3D, 6-connectivity), or simply use the anchor point locations of the segments. In the latter case the source resolution, nth slice, extraction area, processing block size, and max region diameter parameters are ignored.

Max region diameter at selected mip level:

If connected regions are extracted from the segmentation voxel data, these values tell VastTools how many additional rows and columns of voxels have to be loaded at the sides of the processing blocks to make sure that complete regions are found. Ideally, the painted markers for the target locations should be small so that no large overhead is needed.

Source model OBJ:

Click 'Browse' to open a file browser where you can select a source 3D model (.obj file) to be placed at the target locations. A 40nm vesicle model (at micron scale) is included with VastTools.

If 'Recenter Model' is checked, the model will be translated so that its center, defined by the average of its vertex coordinates, is at (0,0,0). 'Rescale source model' lets you specify an isometric scaling factor for the source model.

Save coordinates:

If this option is enabled, the locations of the placed 3D objects and the volumes of their regions will be saved to the specified text file. This information can be used to count the objects and to analyze their locations.

4.4 Export 3D Boxes, Planes and Scale Bars

VastTools can export 3D box models in dataset coordinates, which you can use to show regions of the data set in 3D renderings. If you select 'Export / Export 3D Box ...' from the VastTools menu, a dialog window will open with many of the same settings as in the object exporting dialog (see above for a detailed description).

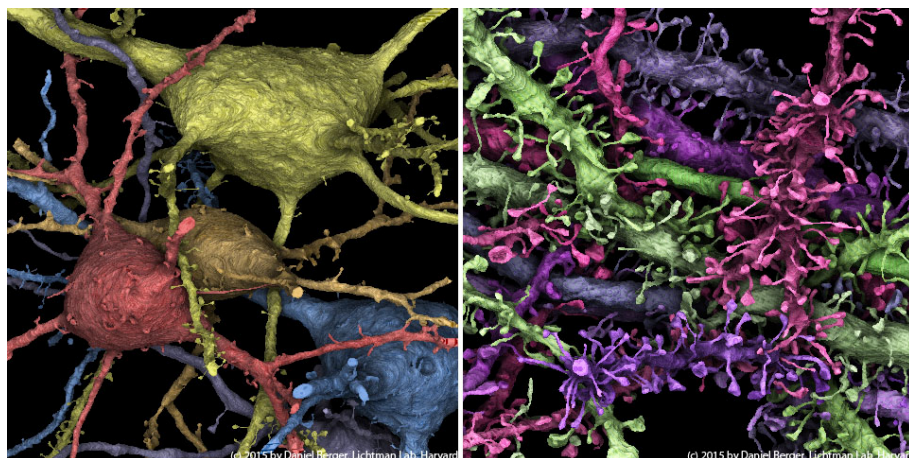


Figure 4.2: Example stack projection images with simulated shadows, generated with VastTools in Matlab

Three box styles are supported:

- 'Wireframe': This will output a 3D box frame mesh with box color and wireframe width as defined. Click the 'Set' button to change the wireframe color.
- 'Solid-color surfaces': This is used to generate a six-sided box, colored in a single color (set under 'Box Color')
- 'Textured surfaces (VAST screenshot images)': Use this to generate a six-sided box with correct textures on its sides. This function will output three files, an .OBJ file specifying the box coordinates, a .MTL file to specify its material with reference to the surface texture, and the surface texture itself as a .PNG file. The surface texture image contains the box surface arranged in a way so that you can print it, cut it out, and make a paper box model of your dataset; however, you'll have to make sure yourself that the texture resolution is isometric.

If the thickness of the box in one dimension is zero (same min and max value under 'Box coordinates'), VAST will export an axis-parallel plane rather than a box. This works for solid-color and textured surface modes.

Scale bar 3D models can be exported under 'Export / Export 3D Scale Bar ...'. You can specify a scale bar length, width, orientation, and color. This function (as well as all other model export functions) uses the data set scale as it is set under 'Info / Volume properties ...' in the VAST main menu. Please make sure that the scaling is correct before exporting scale bars and other models.

4.5 Exporting Projection Images

The 'Export Projection Image' tool allows you to generate 2D projection images in which your microscopic image stacks and/or segmentations are projected along a cardinal axis. This is for example useful to generate 'renderings' of your segmentation or a Z-projection of part of a confocal light microscopic image stack. You can also use these projection images as 'Simple Navigator Images', as described in section 4.7. The following parameters can be set in the export dialog:

Render at resolution:

Sets the mip level at which the projection image is rendered. This affects the resolution of the source images in X and Y, but not Z.

Use every nth slice:

To reduce the resolution in Z (and speed up rendering of very large stacks) you can skip slices if you set this to an integer value larger than 1.

Render from area:

This determines the XYZ block from which the image will be generated. You can use the buttons 'Set to current voxel' and 'Extend to current voxel' together with moving the center cross in VAST to different locations to help you set up a source area. 'Set to full' sets the area to the complete volume, and 'Set to selected bbox' sets it to the combined bounding box of the selected segment and its children. CAUTION: The bounding boxes associated with each segment in VAST are currently not always correct. For example, if you erase voxels you previously painted in VAST, the bounding box will not shrink.

Projection axis, stretching:

Select here along which cardinal axis you would like to project and from what side (which side should be in front in the image), and whether the resulting image should be stretched in Z (for projections along the X or Y axis) if the slice thickness is different from the pixel size. The stretching will be determined by the voxel size of your image stack as set in VAST (see section 3.1.4).

File name, Target folder:

Enter the target file name and location here. The projection image will be saved there after generation if 'Save to file' is enabled.

Segmentation preprocessing:

Select here how you want to use the segmentation for the projection image. This determines which parts of the segmentation are used, and for the 'Screenshots' image source whether the segments should be colored as displayed in VAST (collapsed) or in their individual native colors (uncollapsed).

Expand segments by n pixels:

Makes your segmented regions larger. Useful when using the segmentation as a mask.

Blur edges by n pixels:

Blurs the edges of the segmentation to smooth out the projection image.

Image source:

Select here what you want to use as source images for the projection image. Select 'Screenshots' if you want to use images as they are shown in VAST (using all brightness and contrast settings, layer blending and tinting, segmentation transparency, patterns, etc). A faster option is to just use the primary colors of either the selected or all segmentation layers.

Background color:

Color shown behind the image stack where the stack is transparent

Opacity source:

Determines which parts of each image plane should be rendered opaque for the projection image. Choose to render only segmented areas, only unsegmented areas, or all (Constant). You can also use the pixel brightness of screenshots to define opacity, and either render bright or dark pixels opaque.

Object opacity [0..1]:

Multiplier for the opacity of areas selected in 'Opacity source'. Use a value between 0 (transparent) to 1 (fully opaque).

Blending mode:

Select here whether you want to alpha-blend images with the selected opacity map, add all images, or do a maximum or minimum value projection.

Use shadows, Shadow cone angle:

Enable the check box if you want to include (fake) shadows in the rendering. The shadow cone angle determines how much the shadow spreads out (by means of an image blur filter) from slice to slice.

Depth attenuation (far brightness) [0..1]:

If you set this to values below 1, slices which are further away in the stack will be darkened. This can support the impression of depth in the resulting projection image.

Normalize projection image:

If enabled, the brightness of the final projection image will be adapted so that it uses the whole brightness range (without changing color hue).

4.6 Measuring

VastTools provides several functions to measure objects labeled in VAST.

4.6.1 Measure Segment Volumes

This function can count the number of voxels of different segmented objects in a source area, and save the results to a text file. Parameters in the setup dialog of this function are similar to those of exporting 3D models and measuring surface sizes (see section 4.2 above). Please make sure that the voxel size you use is correct and that you cover the complete area you want to measure to get correct results.

4.6.2 Measure Segment Lengths

This function is not yet implemented.

4.6.3 Measure Segment Surface Area

Since this function relies on a surface mesh to estimate the surface area, it is part of the 'Export 3D Surfaces as OBJ Files' function (see section 4.2). You do not have to generate model files if you just want to measure surface area (check 'Skip model file generation' in the parameter setup dialog). Please make sure that the voxel size you use is correct and that you cover the complete area you want to measure to get correct results.

4.6.4 Euclidian Distance Measurement Tool

This will open a dialog in which you can make simple 3D distance measurements between points in VAST.

First, go to a location in VAST (use the center cross). Then, in VastTools, click the 'Get' button next to the first coordinate to read the current location from VAST. Then go to the other location

and click 'Get' next to the second coordinate. The distance between the two locations will show up as a distance in voxels, and in nanometers.

You can easily jump back to the first and second coordinate in VAST by pressing the 'GO!' button.

Make sure you are using the right voxel size to get accurate nanometer measurements. Click 'Update' to read the current voxel size from VAST. To change the voxel size, please set the voxel size in the .vsv file used in VAST. To do this go to 'Info / Volume Properties ...' in VAST and change the voxel size (see also section 3.1.4). Then click 'Update' in the VastTools Euclidian Distance Measurement Tool.

You can copy the values in the edit fields in the Euclidian Distance Measurement Tool to paste them into different programs, but editing these values directly will currently not work.

4.7 Simple Navigator Images

Simple Navigator Images provide a means to navigate in an image stack using a projection image. After you rendered a projection image in VastTools using 'Export / Export Projection Image ...' (see section 4.5), you can generate a clickable Simple Navigator Image from it. Select 'Navigate / New Simple Navigator Image From Last Projection Image ...'. A new window should pop up which shows the last projection image together with a menu and toolbar. Use the arrow tool to click on a segment anywhere in the projection image to have VAST move to that location in your data set. The magnifying glass tool and the hand tool can be used to zoom and pan the image respectively.

Use 'File / Save Simple Navigator Image ...' to save this Simple Navigator Image to a file which you can open again later. You can have several Simple Navigator Images open at the same time, for example projection images along several axes, and navigate using all of them.

Please note that this functionality has been superseded by the 3D Viewer, which can now also be used to navigate in the image stack (see section 3.11.1).

4.8 Target Lists

Target lists can store the current view coordinates, zoom level and selected segment together with comments. Click the 'Add current VAST location' button to add the current view in VAST to the target list. To move the view in VAST back to a stored location, click on the 'GO!' field of the row in the table. You'll see the change when you make VAST the active window. If you want to update the coordinates, zoom level and/or selected segment number of one of the rows in the target list to the current state in VAST, select it and click the 'Update First Selected' button.

The 'Update Zoom' and 'Update Segnr' checkboxes specify whether zoom level and selected segment number should be transmitted to VAST when you press 'GO!', and also whether they will be updated in the list if you press the 'Update First Selected' button.

Functions to delete, rearrange and add separator lines to the list are provided in the Edit menu. You can save target lists to a file and load them back later. You can have several target lists open at once and cut/copy/paste between them. You can also select the list, or part of the list, and copy (Ctrl-C) and paste (Ctrl-V) the contents into other programs, like a text editor or Excel.

There is now also a function to import a series of coordinates (and zoom levels) into a target list from a Matlab matrix, under 'Edit / Import Coordinates From Matlab Matrix'.

Target lists are stored as simple .mat Matlab files, so you can generate them yourself if you fill a .mat file with the appropriate variables. Simply load a target list file into Matlab to see which variables it contains.

Appendix A

FAQ and Trouble Shooting

A.1 Frequently Asked Questions

My image stack is not aligned. How do I get it aligned into VAST?

VAST does not have any stack alignment (nor stitching) functionality. You'll have to use other programs to render an aligned image stack (for example Adobe Photoshop or plugins in Fiji [1], [2]), and then import that aligned stack into VAST; or you'll have to work on an unaligned image stack.

Can I analyze multi-channel optical image stacks in VAST?

Yes. You can load several image stacks at once, provided they are aligned and the same size, and each one can be RGB or graylevel. You can tint different image stacks in different colors to distinguish different channels and blend them together in different ways (see section 3.2.4).

Does VAST support 4-dimensional data sets (for example time-lapse data of a 3D structure)?

No.

How do I open a .VSS file without a matching .VSV file in VAST?

This is possible with a work-around. If you load any .vsv and then open the .vss in question, and the .vss file has a different size, VAST will tell you its dimensions (size in pixels) in the error message window. Note down these numbers. Now create a dummy .vsvr with these dimensions (vsvr files are just text (.txt) files with a different file name extension). Set the field 'ServerName' to an empty string to make it a dummy layer. Then restart VAST, open the dummy .vsvr, and open your .vss (which should have the same size) on top. Here is an 'empty dummy' .vsvr example:

```
{
  "Comment": "Empty Dummy Layer",
  "ServerType": "openconnectome",
  "ServerName": "",
  "ServerFolder": "",
  "SourceDataSizeX": 49152,
  "SourceDataSizeY": 32768,
  "SourceDataSizeZ": 255,
  "TargetDataSizeX": 49152,
  "TargetDataSizeY": 32768,
  "TargetDataSizeZ": 255,
  "OffsetX": 0,
  "OffsetY": 0,
  "OffsetZ": 0,
  "OffsetMip": 0,
```

```

"TargetVoxelSizeXnm": 6,
"TargetVoxelSizeYnm": 6,
"TargetVoxelSizeZnm": 30,
"TargetLayerName": "Empty Dummy Layer"
}

```

An empty dummy .vsvr file (emptydummy.vsvr) for you to edit is also provided in the support package in the folder 'Online Datasets' (see section 2.3).

How do I deal with self-touching objects?

If you need to be able to recover the true shape of an object, for example for correct skeletonization or computation of the surface area, places where there are self-touches (for example, a dendritic spine touching the dendritic shaft) can be problematic. The easiest solution is to leave a gap between the two sides, but that is not always feasible, in particular in the Z direction. One way to get around this is by using sub-objects and glue. Just like a plastic model which is constructed from parts, you would make the spine a child of the parent and add 'glue' – a different segment which you treat specially in the analysis – to the interface where parent and child object are actually connected.

How do I make shiny 3D pictures and animations from the segmentations in VAST?

VAST does currently not have 3D rendering capabilities apart from the 3D viewer. You can use VastTools to extract .OBJ (Wavefront OBJ) model files of objects segmented in VAST (section 4.2). These models can then be loaded into 3D rendering programs (I use Autodesk 3dsMax).

I accidentally pressed 'Home'. How do I get back to where I was just painting?

The 'Coordinates' tool window keeps a temporally ordered list of recently visited places. Select the second entry from the bottom in the drop-down menu list to jump back.

I exported my segmentation as an image stack, but the images are all black. Why?

When exporting segmentations (not screen shots!) as image stacks, the RGB or graylevel values of the pixels encode the internal segment IDs of your segmentation, starting from 1, as explained in section 3.10. Small ID numbers will show up as very dark gray or blue. Load an image into Photoshop or Fiji and normalize the contrast to check the segments.

Suddenly all internal segment IDs changed – What happened?

VAST currently keeps a contiguous list of segment IDs between 1 and n for n segments at all times. This means that if segments are removed from a segmentation, the other segments will 'move up' to keep the list contiguous. This happens when you use the functions 'Delete Segment + Subtree', 'Weld Segment Subtree' or save a new .vss file with a subset of the current segments using the 'Save Segmentation As Special ...' function. If you are using the internal IDs to identify particular segments and do not want them to change, avoid those functions. For example you can put deleted segments into a 'Deleted' folder and/or re-use them instead of actually deleting them. If you use 'Merge Segmentations in...' or 'Import Segmentation ...' with certain settings, the internal segment IDs of the imported segmentation will also be changed so that they don't overlap with existing IDs. You can specify that in the import dialog.

Why is it called 'VAST Lite' and not just 'VAST'?

The 'Lite' in the name emphasizes that this is not supposed to be the final version of the software. It is currently a usable tool with a limited set of capabilities, which is provided to the scientific community without restrictions. Development of VAST continues, and there may at some point be a released version with more features. Also, 'VAST Lite' is a better name for Google searches than 'VAST'.

A.2 Typical Use Cases

Analyzing synaptic connectivity in a cortex EM stack

It is possible to define the location of synapses and their synaptic partners by painting the synaptic membrane. We do this in two steps. First, we paint the axons and dendrites with individual colors. Then we generate a second data set (a second segmentation layer) in which just the synapse membranes are labeled, using a pen with fixed size (for example 16 pixels diameter) which is wide enough to overlap with pre- and postsynaptic labels. Make sure that the 3D region painted for each synapse is a single connected component and that different synapses are separate connected components. We then export both data sets as image stacks and use a Matlab script to find each synapse by doing connected-component analysis. For each synapse we then find the axon and dendrite which occupy the same voxels as the synapse in the other data set, which gives us the connectivity information. If axons and dendrites are classified as such either via name or segment hierarchy, we can also extract which side is presynaptic and which side is postsynaptic.

Counting and classifying objects by painting

Just as for the synapses in the previous example, you can use connected-component analysis in Matlab to count other objects in the stack, for example neuron cell bodies. If you paint all neurons of each type in the same color and use different colors for different types, connected-component analysis can be used to separate the different cell bodies for each type, given that they are separated in space. The connected-component analysis will also give you the number of objects of each type, and their volume, if you count the painted voxels for each connected component.

For small objects like vesicles you can use the Particle Cloud exporter in VastTools to do the connected-component analysis and export a text file with the center locations and volumes of all connected components; see section 4.3.

Masking out a single cell from a confocal light microscopy image stack

VAST can also load image stacks acquired in light microscopy. If a subset of cells with overlapping branches is fluorescently labeled in such an image stack, VAST can be used to generate a Z-projection image which shows only one cell. For this, first generate a 3D mask by painting over all parts of the cell you want to show. Switch off 'Alpha' for the segmentation layer in VAST so the segmentation is invisible. Then use the VastTools function 'Export / Export Projection Image' to generate a Z-projection image of only the segmented regions of the image stack. For this use 'Screenshots' as image source, 'Segmented areas' as opacity source, and 'Additive' as blending mode. You can optimize brightness/contrast and blending of several layers in VAST to tune the Z-projection image.

Tracking objects in a video

If you translate a video into a sequence of images, you can of course import this image sequence as a stack into VAST (even in color). In the same way as you can label three-dimensional structures in VAST, you can label objects or regions or fiducial points as they move through the video. You can then export the labelings as an image stack and analyze locations in the image and movement, or use the Particle Cloud exporter in VastTools to export a text file with fiducial coordinates; see section 4.3.

Defining fiducial points in an unaligned image stack for manually aided alignment

Some EM image stacks are difficult to align with automatic methods, for example if the image quality is low, there is high-contrast background, or the tissue slices have folds. Manually defined fiducial points which should end up in the same place from slice to slice can help improve the alignment. You can load an unaligned stack into VAST and use manual painting to define fiducial points. Use a different color for each feature you are tracking through the stack, so that in your alignment script

you can tell which points belong together. Again you can use the Particle Cloud exporter in VastTools to export a text file with fiducial coordinates to help the analysis; see section 4.3.

Proof-reading an automatic segmentation

VAST can load automatic segmentations in two ways, either as image stack or as editable segmentation stack. The second option currently requires that the segmentation has less than 2^{16} (65535 or less) segments.

If the segmentation is loaded as an image stack, a proof-read segmentation can be generated in a separate segmentation layer, using masked painting (see section 3.3.5) and filling operations (section 3.4.2). Boundary maps can also be loaded as image stacks and used to generate a proof-read segmentation (see sections 3.3.4 and 3.4.2).

If the segmentation is loaded as a segmentation layer, it can be edited directly. Split segments can be merged by using the Collect tool (section 3.6.17) and subsequent welding (section 3.6.9). Mergers can be corrected by erasing or recoloring the boundary (for example by adding a color for the to-be-split-off object as a child, and using parent mode (section 3.3) to recolor it where it touches the other object), then use filling to completely split it off (section 3.4.1).¹ Of course the user has full freedom to edit the segmentation simply by painting as well.

A.3 Some Performance Tips

- If file access is very slow (when you move through the stack and it takes time until the images appear) consider storing the data locally and/or on SSD drives. In particular, especially when using non-SSD drives, put files which you use together on physically separate drives. The problem is often that two files (for example an EM layer and the segmentation layer) are loaded at the same time from the same non-SSD hard drive, which causes the read/write head to jump forth and back between two locations at high speed. This slows down file access a lot.

If images are loaded from image files (.VSVI), loading can be slow if the image tiles are very large, since a stack of whole image tiles has to be loaded even if only a small region is displayed. Consider making the image tiles smaller (1024x1024 or less) to improve performance.

- If you experience a low frame rate (the mouse cursor is jumping rather than moving smoothly), try to reduce the size of the 'Maximal Window Width' in the Preferences to something like 1280. On very large screens you can set the 'Target resolution smaller than' to 4 (Default: 2) to help.
- If VAST slows down considerably after using it for a while, check if the RAM of your computer is full (check the RAM usage indicator in the upper right corner of the VAST main window, see section 3.2.2). Once memory is full, Windows might swap parts of the data to the hard drive, which can slow down processing a lot. To fix that problem tell VAST to use less memory by REDUCING the maximal RAM cache memory sizes in the Preferences dialog. This should not cause problems even if you are working with large image stacks since VAST does not need to load all data in RAM at once. The only effect which reducing the cache memory size should have is that VAST may have to load parts of your data set from disk more often, which is slower than reading from RAM.

¹This requires however that the split-off branch is a single connected component.

A.4 Setting up VAST with a Wacom screen

When it comes to fast and accurate digital painting, tablets and in particular tablet screens can improve productivity significantly. We are using various Wacom Cintiq screen tablets. While the larger Cintiqs have a better screen, they can be expensive and bulky. I personally prefer the Cintiq 13HD, which can be laid flat onto a desk and close to the user.

An alternative could be tablet laptops, if they fulfill the system requirements for VAST. We tried the Asus EEE Slate EP121; it works but it is a bit slow, and pen button presses tend to be unreliable. Not all tablets are ideal for VAST, for example if pen input is unreliable or if the pen has no buttons or the buttons cannot be customized.

Wacom tablets come with driver software which lets you configure the pen buttons for each program. For optimal workflow in VAST I find it most useful to have 'erase' on one pen button and 'change tooltip size' on the other. For this it is easiest to set one pen button to 'right click' and the other one to 'middle click'. If that does not work for your system (sometimes Windows uses the Right Click event in a special way, for example), you can configure the pen buttons to simulate equivalent key presses (see Appendix A.5).

If you see brief circular animations when you use the pen and drawing is delayed, you should go to 'Pen and Touch' in the Windows Control Panel and switch off 'Flicks'. On the tab 'Flicks', uncheck 'Use flicks to perform common actions quickly and easily' and click 'Apply'.

In Windows 10, find the checkbox 'Use Windows Ink' in the Wacom driver ('Wacom Tablet Properties' in the Control Panel) and uncheck it. This solves most of the lag problems.

A.5 Keyboard Shortcuts in VAST

You can open a window which lists all the keyboard shortcuts in VAST from the main menu under Window / Keyboard Shortcuts. Here is a summary.

SHIFT or ENTER	Pick segment color by mouse click
CTRL or INSERT	Temporarily go to 'move' mode
TAB or \	Move mode: Click left and move pen up/down to zoom; Paint mode: Click left and move pen up/down to change pen diameter
'~ or DELETE	Paint mode: Erase mode
H or L	Hides segmentation while held down

Table A.1: Mode Modifiers (hold down)

UP or A	One slice up
DOWN or Z	One slice down
PAGEUP or S	MAXPAINTDEPTH slices up
PAGEDOWN or X	MAXPAINTDEPTH slices down

Table A.2: Slice Navigation

MAXPAINTDEPTH is the value set under 'Max Paint Depth' in the 'Drawing Properties' dialog.

-_	Decrease tooltip diameter
=+	Increase tooltip diameter
N or Keypad /	Zoom out
M or Keypad *	Zoom in
F	Flash selected segment
I, O, P	Set paint mode to Paint all, Background, Parent
HOME or G	Go to anchor point of selected segment
,< / .>	Select previous / next segment in recently selected list

Table A.3: Other Controls

A.6 Terms of Usage and Privacy Statement

This version of VAST ('the software') is free of charge and may be distributed freely, but not sold. Commercial usage is allowed.

You are using this software at your own risk. Even though it has been tested extensively, it is not free of bugs. Please keep backup copies of your data.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

VAST does not collect usage statistics or other data. In particular, it does not transmit any of your image or segmentation data anywhere.²

This software uses `easyzlib.c`, which is based on the `zlib` library by Jean-loup Gailly and Mark Adler, and `libtiff` with Copyright (c) 1988-1997 Sam Leffler and Copyright (c) 1991-1997 Silicon Graphics, Inc. VastTools uses `jtcp.m` by Kevin Bartlett and `loadawobj.m` by W.S. Harwin, University Reading (Matlab BSD license).

²Except if you set that up explicitly using the Remote Control API Server of course. You can make it transmit segmentation data through the API to VastTools for example, but that is under your control.

Appendix B

Technical Information

B.1 Size limitations

Maximal file size for EM images and segmentations: Limited by maximal file size on disk; theoretical maximum 2^{64} bytes.

Largest EM stack that has been imported into VAST so far: ~ 5 Terabytes

Largest EM stack opened remotely (via .VSVR): ~ 1.3 Petabytes

Maximal number of labels supported: currently $2^{16} - 1$ (but needs lots of RAM; 600 bytes per label).

Largest image supported (at full resolution): $(2^{31} - 1) \cdot (2^{31} - 1)$

B.2 Supported file formats for importing / exporting

Importing of EM images

Currently 8 bit graylevel and 24-bit RGB image stacks are supported.

Stacks and tiled stacks: .JPG, .PNG, .TIF, .BMP

3D Volume files: .NII (NIFTI); will be converted to 8 bits when imported

Versions 1.1 and later of VAST Lite use libtiff for importing from .TIF images, which solves the previous problems caused by the Windows GDI+ TIFF routines.

Exporting of EM images

Stacks and tiled stacks: .PNG 8-bit indexed and 24-bit RGB, .TIF uncompressed, .RAW

Importing of segmentations

Segmentations can be imported from RGB .TIF and .PNG image stacks. The segment number for each pixel is encoded in the RGB value of the image as follows: Bits 0-7 of the label number are expected in the blue channel, bits 8-15 in the green channel, and bits 16-23 in the red channel. This is the same format used for exporting segmentations to image stacks (see below). Please be aware that VAST can currently only handle 16-bit segmentations.

Exporting of segmentations

When exporting segmentations, the available file formats depend on the range of segment numbers used. For example, if the highest segment number is greater than 255, 8 bit indexed file formats will not be available. In that case the label numbers will be encoded into the color channels (for example for RGB, bits 0-7 of the label number will be put into the blue channel, bits 8-15 into the green channel, and bits 16-23 into the red channel).

Stacks and tiled stacks: .PNG, .TIF uncompressed, .RAW
 3D Volume files: .NII (Nifti) 8 bit only (currently unsupported)

Exporting of screen shots

Stacks and tiled stacks: .PNG, .TIF uncompressed, .RAW (all 24 bit RGB)
 3D Volume files: .NII (Nifti) 8 bit only (currently unsupported)

B.3 API Function Reference

VAST includes an API through which external programs can communicate with VAST. The communication is done through a TCP/IP connection. This makes it possible to write client programs in any programming language which supports the TCP/IP protocol. It also allows client programs to run on separate computers and communicate through the network.

VAST comes with a client-side implementation in Matlab, which is used by the VastTools toolbox and included in the `vast_package.zip` file (see section 2.3).

In Matlab the VAST API is implemented as a class, `VASTControlClass.m`. Internally it uses the `jtcp` library for TCP/IP communication in Matlab. The different API functions are implemented as class methods.

Using the VAST API is very simple. With `VASTControlClass.m` and its helper functions available (in the path), simply make an instance of the class and use its functions. Here is a short example which connects to VAST, reads out some basic information, and disconnects:

```
vast=VASTControlClass();
res=vast.connect('127.0.0.1',22081,1000);
if (res==0)
    warndlg('Connecting to VAST at 127.0.0.1, port 22081 failed.','Error');
else
    vinfo=vast.getinfo();
    disp(vinfo);
    vast.disconnect();
end;
```

You can also use VastTools to handle the connection. Simply run VastTools and connect, then run your own script. VastTools uses a global variable 'vdata'. To access the API functions, simply enable access to the global variable, then use the VAST API like so:

```
global vdata;
vinfo=vdata.vast.getinfo();
disp(vinfo);
```

The following API functions are currently available (Version 3, as returned by `getapiversion`):

res = connect(host, port, timeout)

Tries to connect to VAST via TCP/IP. 'host' is given as a text string, 'port' as a number and 'timeout' as a number (in milliseconds). Returns 1 if connected, 0 if connection failed.

res = disconnect()

Disconnects the TCP/IP connection to VAST. Always returns 1. If the request fails, a Java Error will stop the script, so a return of 0 is currently not possible.

errornumber = getlasterror()

Retrieves the error code of the last call of an API function. Error codes have the following meaning:

- 0 No error
- 1 Unknown error
- 2 Unexpected data received from VAST - API mismatch?
- 3 VAST received invalid data. Command ignored.
- 4 VAST internal data read failure
- 5 Internal VAST Error
- 6 Could not complete command because modifying the view in VAST is disabled
- 7 Could not complete command because modifying the segmentation in VAST is disabled
- 10 Coordinates out of bounds
- 11 Mip level out of bounds
- 12 Data size overflow
- 13 Data size mismatch - Specified coordinates and data block size do not match
- 20 Segment number out of bounds
- 21 No segmentation available
- 22 RLE overflow - RLE-encoding makes the data larger than raw; please request as raw
- 30 Layer number out of bounds
- 31 Invalid layer type
- 50 sourcearray and targetarray must have the same length

Table B.1: VAST API Error Codes as returned by getlasterror()

[apiversion, res] = getapiversion()

Returns the version of the API provided by the currently connected VAST executable. The API version described here is version 3.

If the call succeeded, res will be 1, otherwise 0. Use getlasterror() to retrieve the error code.

[version, res] = getcontrolclassversion()

Returns the version of the VASTControlClass.m script (not part of the API). The current version is 3, introduced in VAST 1.2.

[info, res] = getinfo()

Reads out general information from VAST.

If the call succeeded, res will be 1, otherwise 0. Use getlasterror() to retrieve the error code.

Returns a struct with the following fields if successful, or an empty struct [] if failed:

<code>info.datasizex</code>	X (horizontal) size of the data volume in voxels at full resolution
<code>info.datasizey</code>	Y (vertical) size of the data volume in voxels at full resolution
<code>info.datasizez</code>	Z (number of slices) size of the data volume in voxels
<code>info.voxelsizex</code>	X size of one voxel (in nm)
<code>info.voxelsizey</code>	Y size of one voxel (in nm)
<code>info.voxelsizez</code>	Z size of one voxel (in nm)
<code>info.cubesizex</code>	X size of the internal cubes used in VAST in voxels; always 16
<code>info.cubesizey</code>	Y size of the internal cubes used in VAST in voxels; always 16
<code>info.cubesizez</code>	Z size of the internal cubes used in VAST in voxels; always 16
<code>info.currentviewx</code>	Current view X coord in VAST in voxels at full res (window center)
<code>info.currentviewy</code>	Current view Y coord in VAST in voxels at full res (window center)
<code>info.currentviewz</code>	Current view Z coord in VAST in voxels (slice number)
<code>info.nrofmipllevels</code>	Number of mip levels of the current data set

Table B.2: `getinfo()` struct**[nr, res] = getnumberofsegments()**

Returns the number of segments of the active segmentation layer or [] if failed.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[data, res] = getsegmentdata(id)

Reads out information of the segment with number `id` in the active segmentation layer. This is the same information that gets written to a text file by 'Save Segmentation Metadata ...' in VAST (see section 3.6.14).

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

Returns a struct with the following fields if successful, or an empty struct [] if failed:

<code>data.id</code>	ID of the requested segment
<code>data.flags</code>	Flags field of the segment as 32-bit value
<code>data.col1</code>	Primary color as 32-bit value
<code>data.col2</code>	Secondary color as 32-bit value
<code>data.anchorpoint</code>	XYZ coordinates of the segment's anchor point (in voxels)
<code>data.hierarchy</code>	IDs of parent, child, previous and next segment (0 if none)
<code>data.collapsednr</code>	If the segment is collapsed into a folder, this is the folder ID
<code>data.boundingBox</code>	Segment bounding box (may be incorrect if voxels were deleted)

Table B.3: `getsegmentdata()` struct**[name, res] = getsegmentname(id)**

Returns the name of the segment with the given ID in the active segmentation layer. Returns [] if the ID is out of range.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

res = setanchorpoint(id, x, y, z)

Sets the anchor point of the segment with the given ID in the active segmentation layer to the given coordinates (in voxels at full resolution). Only non-negative values are allowed for x, y and z.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`res = setsegmentname(id, name)`

Sets the name of the segment with the given ID in the active segmentation layer.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`res = setsegmentcolor8(id, r1, g1, b1, p1, r2, g2, b2, p2)`

Sets the primary and secondary colors (and pattern) of the segment with the given ID in the active segmentation layer. Values for `r1`, `g1`, `b1`, `r2`, `g2` and `b2` have to be between 0 and 255. `p1` defines the pattern and has to be between 0 and 15. `p2` is currently unused.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`res = setsegmentcolor32(id, col1, col2)`

Same as `setsegmentcolor8`, but the two colors for the segment are given as two 32-bit values.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`[x, y, z, res] = getviewcoordinates()`

Returns the current view coordinates in VAST in voxels at full resolution.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`[zoom, res] = getviewzoom()`

Returns the current zoom value in VAST.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`res = setviewcoordinates(x, y, z)`

Sets the current `x`, `y`, `z` coordinates (in voxels at full resolution) of the view (center of window) in VAST.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`res = setviewzoom(zoom)`

Sets the current zoom value. Zoom values are currently integer values and can be negative. The higher the value, the more magnified the view. A zoom value of 0 sets the pixel size of the full resolution image to exactly the 'Target Resolution Smaller Than' value specified in the VAST Preferences. A value of -8 zooms out by a factor 2 (and shows the next lower mip level texture).

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[nroflayers, res] = getnroflayers()

Returns the number of layers currently loaded in VAST.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[linfo, res] = getlayerinfo(layernr)

Returns information about the layer with the given number. layernr can be a number between 0 and `getnroflayers()-1`. If the function succeeds, it will return a struct `layerinfo` with fields as described below and res will be 1. If it fails it will return an empty struct and res will be 0. This returns the values which are also visible in the VAST 'Layers' tool window.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

<code>linfo.type</code>	0: EM layer (VSV); 1: Segmentation layer; 3: remote EM layer
<code>linfo.editable</code>	1 if the layer is editable, 0 else
<code>linfo.visible</code>	1 if the layer is visible, 0 else ('Visible' checkbox)
<code>linfo.brightness</code>	'Bright' (brightness) or 'Pattern' checkbox
<code>linfo.contrast</code>	'Contrast' or 'Sel Alpha' checkbox
<code>linfo.opacitylevel</code>	Slider next to 'Visible', which defines layer opacity
<code>linfo.brightnesslevel</code>	Slider next to 'Bright' or 'Pattern'
<code>linfo.contrastlevel</code>	Slider next to 'Contrast' or 'Sel Alpha'
<code>linfo.blendmode</code>	The current blend mode (see below)
<code>linfo.blendoradd</code>	0: alpha blending, 1: additive blending
<code>linfo.tintcolor</code>	Tint color as 32-bit RGBA value. Default: white (0xffffffff)
<code>linfo.name</code>	Name of the layer
<code>linfo.redtargetcolor</code>	Red channel target color (RGBA). Default: red (0xff0000ff)
<code>linfo.greentargetcolor</code>	Green channel target color (RGBA). Default: green (0x00ff00ff)
<code>linfo.bluetargetcolor</code>	Blue channel target color (RGBA). Default: blue (0x0000ffff)

Table B.4: `getlayerinfo()` struct

Blend modes are: 0: no transparency ('Flat'), 1: the darker $\text{mean}(r, g, b)$ the more transparent, 2: the brighter $\text{mean}(r, g, b)$ the more transparent, 3: the darker $\text{max}(r, g, b)$, the more transparent, 4: the brighter $\text{max}(r, g, b)$, the more transparent.

[segdata, res] = getallsegmentdata()

This function is similar to `data = getsegmentdata(id)` above, but retrieves the data for all segments in the active segmentation layer at once and returns a cell array of structs instead of a single struct.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[segdatamatrix, res] = getallsegmentdatamatrix()

Same as `data = getallsegmentdata(id)` above, but returns the data as a matrix with one row per segment and one column per data value. The columns are the same as when exporting segment metadata to a file (see section 3.6.14), with the exception of the segment names.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[segname, res] = getallsegmentnames()

This function is similar to `name = getsegmentname(id)` above, but retrieves the names of all segments in the active segmentation layer at once and returns them in a cell array. If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

res = setselectedsegmentnr(segmentnr)

Sets the selected segment in the active segmentation layer in VAST to the segment with the given ID number. If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[selectedsegmentnr, res] = getselectedsegmentnr()

Returns the ID number of the currently selected segment in the active segmentation layer, or -1 if an error occurred. If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

res = setselectedlayernr(obj)

Sets the selected layer in VAST to the layer with the given ID number (counting from 0 as the first layer in the list). If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[selectedlayernr, selecteddemlayernr, selectedsegmentlayernr, res] = getselectedlayernr()

Returns the numbers of the currently selected layer, EM layer, and segment layer, or -1 if there is no such layer. The 'selected layer' will always be either the 'selected EM layer' or the 'selected segment layer', whichever is actually highlighted in the 'Layers' tool window in VAST. The other number indicates the most recently selected layer of the other type.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[segimage, res] = getsegimageraw(miplevel, minx, maxx, miny, maxy, minz, maxz, flipflag)

Reads out the segmentation as a voxel image for a given mip level (resolution) and area, as defined by minimum and maximum values for the ranges in x, y and z. This function transmits the segmentation image as raw data (a one-dimensional array of bytes); consider using `getsegimageRLEdecoded()` below for a version with faster transmission. If `flipflag` is 1, X and Y axis will be swapped to make the image appear the same as in VAST. If `flipflag` is unspecified or 0, `segimage` will appear with X and Y axis swapped.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

`minx`, `maxx`, `miny`, `maxy`, `minz` and `maxz` are given in voxels *at the resolution defined by miplevel*. You can retrieve the size of the volume at full resolution (mip level 0) using the function `getinfo()`. To compute the size of the volume in X and Y at a different mip level `miplevel`, use the following code:

```

xmin=bitshift(xmin,-miplevel);
xmax=bitshift(xmax,-miplevel)-1;
ymin=bitshift(ymin,-miplevel);
ymax=bitshift(ymax,-miplevel)-1;

```

Note that the size of the volume in Z does not change with mip level (this is the number of slices).

[segimageRLE, res] = getsegimageRLE(miplevel, minx, maxx, miny, maxy, minz, maxz, surfonlyflag)

Reads the requested volume region and returns it as an runlength-encoded (RLE) string. This function is useful if you want to directly process the RLE string; otherwise use one of the functions below.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. See `getsegimage()` above for information on volume coordinates.

The runlength-encoded image is an array of pairs of unsigned 16 bit values. The first value in a pair is the segment ID which follows, and the second value is the number of voxels with that ID to follow. The encoding order of voxels in the volume is row by row (X from left to right) concatenated for one plane (Y from top to bottom), and planes are concatenated from the first to the last slice of the volume. Please note that sequences of the same value defined by one pair can loop around at the end of a row and/or extend into the next plane. Run lengths longer than $2^{16} - 1$ (65535) are encoded in several pairs.

[segimage, res] = getsegimageRLEdecoded(miplevel, minx, maxx, miny, maxy, minz, maxz, surfonlyflag, flipflag)

Same as `getsegimage()` above, except that the data is transmitted run-length encoded (RLE) from VAST to Matlab, and then decoded on the Matlab side. This can improve speed because it reduces the size of the transmitted data. If `surfonlyflag` is 1, all interior voxels (any nonzero voxels for which all six direct neighbors have the same ID) are set to 0. If `flipflag` is 1, X and Y axis will be swapped to make the image appear the same as in VAST. If `flipflag` is unspecified or 0, `segimage` will appear with X and Y axis swapped. See `getsegimageraw()` above for information on volume coordinates.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[values, numbers, res] = getRLEcountunique(miplevel, minx, maxx, miny, maxy, minz, maxz, surfonlyflag)

This function returns the segment IDs (values) and number of voxels of each ID (numbers) in the requested subvolume. See `getsegimageraw()` above for information on volume coordinates.

If the call succeeded, res will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[segimage, values, numbers, res] = getsegimageRLEdecodedcountunique(miplevel, minx, maxx, miny, maxy, minz, maxz, surfonlyflag, flipflag)

Combines the functions `getsegimageRLEdecoded()` and `getRLEcountunique()` above, and returns both the decoded subvolume and a list of segment IDs in the volume and number of voxels for each ID. If `flipflag` is 1, X and Y axis will be swapped to make the image appear the same as in VAST. If `flipflag` is unspecified or 0, `segimage` will appear with X and Y axis swapped. See `getsegimageraw()` above for information on volume coordinates.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[segimage, values, numbers, bboxes, res] = getsegimageRLEdecodedbboxes(miplevel, minx, maxx, miny, maxy, minz, maxz, surfonlyflag, flipflag)

Same as `getsegimageRLEdecodedcountunique()` above, but also returns the bounding boxes for all IDs. The coordinate order for the bounding boxes in `bboxes` is: `[minx,miny,minz,maxx,maxy,maxz]`, one row per segment.¹ `bboxes` is not changed by `flipflag`. See `getsegimageraw()` above for information on volume coordinates.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

res = setsegtranslation(sourcearray, targetarray)

Tells VAST how to translate the segmentation volume before transmitting it through `getsegimage*` and `getRLE*` functions. `sourcearray` and `targetarray` are arrays of unsigned (positive) 16-bit values and have to have the same length. VAST will convert all voxels with IDs in `sourcearray` to the corresponding values in `targetarray`. Voxels with other IDs will be removed (set to 0). To disable the segmentation translation, call `setsegtranslation([], [])`.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[emimage, res] = getemimageraw(layernr, miplevel, minx, maxx, miny, maxy, minz, maxz)

Reads out the EM image of the specified layer as a voxel image for a given mip level (resolution) and area. The area is defined by minimum and maximum values for the ranges in x, y and z. This function transmits the segmentation image as raw data (a one-dimensional array of bytes) with either one or three bytes per voxel, depending on whether the layer has one or three color channels. Use `getemimage` to request a correctly reshaped two-, three- or four-dimensional image.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[emimage,res] = getemimage(layernr, miplevel, minx, maxx, miny, maxy, minz, maxz)

Same as `getemimageraw`, but reshapes the one-dimensional array to a matrix of the requested dimensions. The order of dimensions in the matrix is Y,X,Z,C (C for color, R G B, if the layer has three color channels).

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[screenshotimage,res] = getscreenshotimageraw(miplevel, minx, maxx, miny, maxy, minz, maxz, collapseseg)

Reads out the EM and segmentation stack as currently displayed in VAST as a voxel image for a given mip level (resolution) and area. The area is defined by minimum and maximum values for the ranges in x, y and z. If `collapseseg` is 1, the color of the segmentation will appear as displayed in VAST, if it is 0, each segment will be colored in its native (uncollapsed) color. This function transmits the segmentation image as raw data (a one-dimensional array of bytes). Use `getscreenshotimage` to request a correctly reshaped two-, three- or four-dimensional image. Screenshot images are always returned in RGB format (three bytes per voxel).

¹Sorry that the order is different than in most other places.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

[screenshotimage,res] = getscreenshotimage(miplevel, minx, maxx, miny, maxy, minz, maxz, collapseseg)

Same as `getscreenshotimageraw`, but reshapes the one-dimensional array to a matrix of the requested dimensions. The order of dimensions in the matrix is Y,X,Z,C (C for color, R G B).

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code.

res = setsegimageraw(miplevel, minx, maxx, miny, maxy, minz, maxz, segimage)

Writes the (2D or 3D) image in `segimage` to the active segmentation layer in VAST, at the mip level `miplevel` and target coordinates given in `minx`, `maxx`, `miny`, `maxy`, `minz`, `maxz`. These coordinates have to be specified at the resolution of `miplevel` (see explanation in section `getsegimageraw` above).

VAST will not accept segmentation images which contain segment numbers which exceed the segment list. For existing segments, the segment metadata will be updated to include the imported segmentation image.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 2.

res = setsegimageRLE(miplevel, minx, maxx, miny, maxy, minz, maxz, segimage)

Same as `setsegimageraw`, but this version encodes the segmentation image using RLE and transmits the RLE-encoded data to VAST. This reduces the amount of data that has to be transmitted, but RLE encoding in Matlab is slow, so `setsegimageraw` may be the faster alternative for Matlab. If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 2.

res = setsegmentbbox(id, minx, maxx, miny, maxy, minz, maxz)

Sets the bounding box of the segment with number `id` in the active segmentation layer.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 2.

[firstsegmentnr, res]=getfirstsegmentnr()

Returns the segment number (`id`) of the first segment (after Background) in the active segmentation layer.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 2.

[info, res] = gethardwareinfo()

Returns the hardware properties of the computer on which VAST is running in the struct `info`. The fields of the struct are as follows:

<code>info.computername</code>	text	Name of the computer
<code>info.processorname</code>	text	Name of the processor
<code>info.processorspeed_ghz</code>	double	Clock speed of the processor in GHz
<code>info.nrofprocessorcores</code>	integer	Number of processor cores
<code>info.tickspeedmhz</code>	double	Tick speed of the VAST computer in MHz
<code>info.mmxssecapabilities</code>	text	Text describing the processor capabilities
<code>info.totalmemorygb</code>	double	Memory installed, in GB
<code>info.freememorygb</code>	double	Free memory, in GB
<code>info.graphicscardname</code>	text	Graphics card name
<code>info.graphicsdedicatedvideomemgb</code>	double	Dedicated video memory in GB
<code>info.graphicsdedicatedsystememgb</code>	double	System memory dedicated for graphics, in GB
<code>info.graphicssharedsystememgb</code>	double	System memory shared for graphics, in GB
<code>info.graphicsrasterizerused</code>	text	'Undefined', 'Hardware', 'Warp' or 'Reference'

Table B.5: `gethardwareinfo()` struct

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 2.

[id, res] = addsegment(refid, nextorchild, name)

Adds one segment with the given name to the selected segmentation layer. If `nextorchild` is 0, the new segment will be added directly behind the segment with ID `refid`. If `nextorchild` is 1, the new segment will be added as last child of the segment with ID `refid`.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 3.

res = moveselement(id, refid, nextorchild)

Moves the segment with the specified `id` in the selected segmentation layer to the location indicated by `refid` and `nextorchild`. If `nextorchild` is 0, the segment will be removed from its current location and added directly behind the segment with ID `refid`. If `nextorchild` is 1, the segment will be removed from its current location and added as last child of the segment with ID `refid`.

If the call succeeded, `res` will be 1, otherwise 0. Use `getlasterror()` to retrieve the error code. Introduced in API version 3.

Bibliography

- [1] CARDONA A., SAALFELD S., SCHINDELIN J., ARGANDA-CARRERAS I., PREIBISCH S., ET AL.: *TrakEM2 Software for Neural Circuit Reconstruction*, PLoS ONE, **7**(6):e38011, (2012), doi:10.1371/journal.pone.0038011.
- [2] SAALFELD, S., FETTER, R., CARDONA, A., AND TOMANCAK, P.: *Elastic volume reconstruction from series of ultra-thin microscopy sections*, Nature Methods, **9**(7), (2012), 717-720.
- [3] SRUBEK TOMASSY, G., BERGER, D., CHEN, H., KASTHURI, N., HAYWORTH, K., VERCELLI, A., SEUNG, S., LICHTMAN, J., AND ARLOTTA, P.: *Distinct Profiles of Myelin Distribution Along Single Axons of Pyramidal Neurons in the Neocortex*, Science, **344**(6181), (2014) 319-324, doi:10.1126/science.1249766.
- [4] KASTHURI, N., HAYWORTH, K., BERGER, D., SCHALEK, R., CONCHELLO, J., KNOWLES-BARLEY, S., LEE, D., VAZQUEZ-REINA, A., KAYNIG, V., JONES, T., ROBERTS, M., MORGAN, J., TAPIA, J., SEUNG, H.S., GRAY RONCAL, W., VOGELSTEIN, J., BURNS, R., SUSSMAN, D., PRIEBE, C., PFISTER, H., AND LICHTMAN J.: *Saturated Reconstruction Of A Volume Of Neocortex*, Cell, **162**(3), (2015) 648-661.
- [5] MORGAN, J., BERGER, D., AND LICHTMAN J.: *The Fuzzy Logic of Network Connectivity in Mouse Visual Thalamus*, Cell, **165**, (2016) 192-206.
- [6] QUADRATO, G., NGUYEN, T., MACOSKO, E., SHERWOOD, J., YANG, S., BERGER, D., MARIA, N., SCHOLVIN, J., GOLDMAN, M., KINNEY, J., BOYDEN, E., LICHTMAN, J., WILLIAMS, Z., MCCARROLL, S., AND ARLOTTA, P.: *Cell diversity and network dynamics in photosensitive human brain organoids*, Nature, **545**, (2017), 48-53, doi:10.1038/nature22047.
- [7] BERGER, D., SEUNG, S., AND LICHTMAN J.: *VAST (Volume Annotation and Segmentation Tool): Efficient Manual and Semi-Automatic Labeling of Large 3D Image Stacks*, Frontiers in Neural Circuits, **12**, (2018), Article 88, doi:10.3389/fncir.2018.00088.
- [8] FANG, T., LU, X., BERGER, D., GMEINER, C., CHO, J., SCHALEK, R., PLOEGH, H., AND LICHTMAN, J.: *Nanobody immunostaining for correlated light and electron microscopy with preservation of ultrastructure*, Nature Methods, (2018), doi:10.1038/s41592-018-0177-x.

