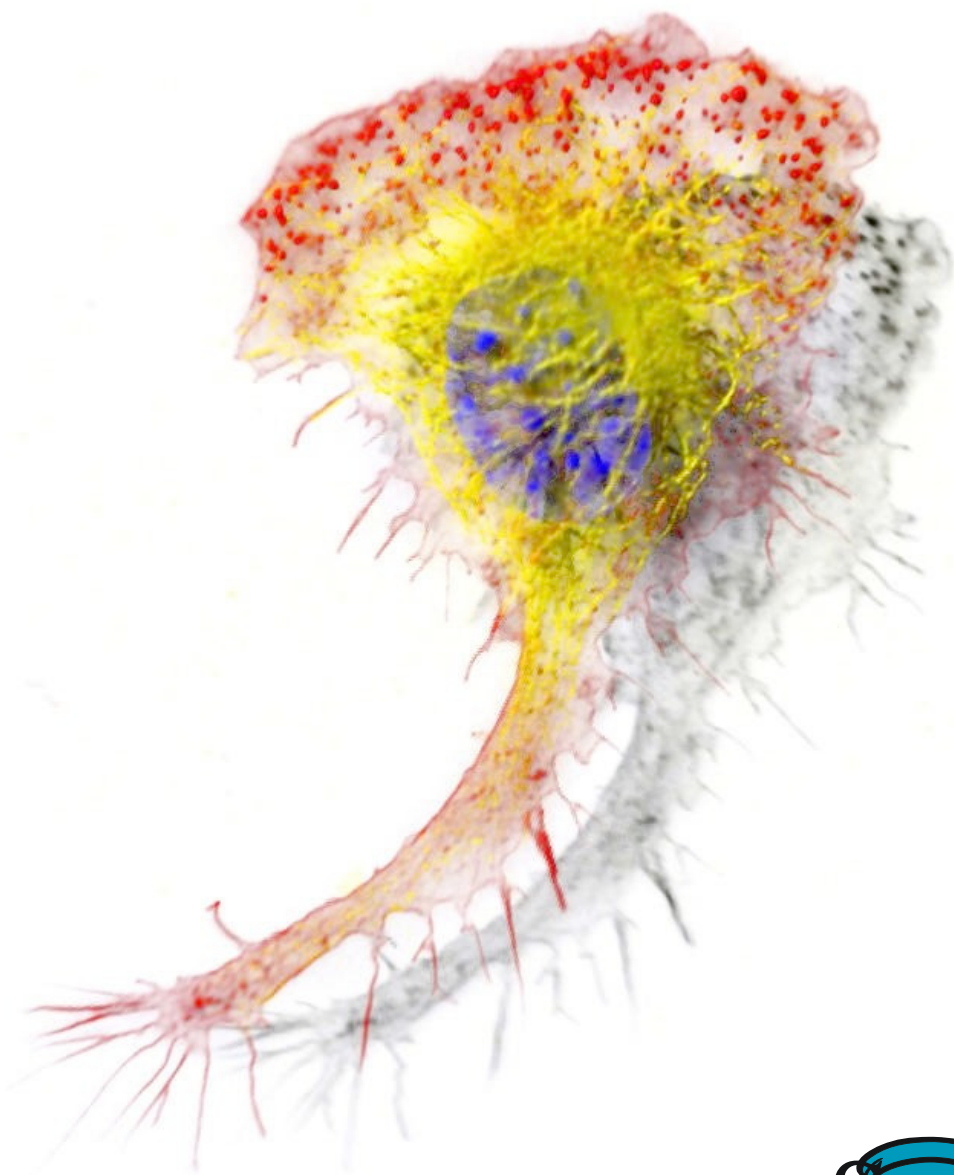

Huygens Professional

User Guide for version 3.7



Scientific Volume Imaging B.V.

Huygens Professional

User Guide for version 3.7

Scientific Volume Imaging B.V.

Copyright © 1995-2011 by Scientific Volume Imaging B.V.
All rights reserved

Mailing Address Scientific Volume Imaging B.V.
Laapersveld 63
1213 VB Hilversum
The Netherlands

Phone +31 35 6421626

Fax +31 35 6837971

E-mail info@svi.nl

URL <http://www.svi.nl/>

Cover illustration: Macrophage recorded by Dr. James Evans (White-head Institute, MIT, Boston MA, USA) using widefield microscopy, as deconvolved with Huygens[®]. Stained for tubulin (yellow/green), actin (red) and the nucleus (DAPI, blue).

Contents

CHAPTER 1	<i>Introduction</i>	1
	What is Huygens Professional?	1
CHAPTER 2	<i>Installation</i>	5
	Microsoft Windows	5
	Microsoft Windows 64 bit Edition	5
	Mac OS X	5
	Linux (Debian)	5
	Linux (RPM)	6
	After the Installation	6
	The License String	6
	Updating the Software	8
	Removing the Software	9
	System Requirements for Huygens Professional	9
	Support on Installation	10
CHAPTER 3	<i>Getting Started</i>	11
	Step 1: Start Huygens Professional	11
	Step 2: Load an Image	12
	Step 3: Inspect your Image	12
	Step 4: Generate a Point Spread Function (PSF)	17
	Step 5: Estimate the Average Background Value	18
	Step 6: Estimate the Signal to Noise Ratio (SNR)	19
	Step 7: Perform a Deconvolution Run	19
	Step 8: Saving your Image	21

CHAPTER 4	<i>The Batch Processor</i>	23
	The Batch Processor Window	23
	Usage	24
	Menus	27
CHAPTER 5	<i>The Slicer</i>	29
	Color	29
	Contrast	30
	Time	30
	Linking	30
CHAPTER 6	<i>The Twin Slicer</i>	31
	Using the Slicer in Basic Mode	32
	Using the Slicer in Advanced Mode	34
	Measurement	36
CHAPTER 7	<i>The Orthogonal Slicer</i>	39
	The Crosshair Cursor	40
	Visualization parameters	40
	Measurements	40
	Auto-Zoom	40
	Display Options	40
CHAPTER 8	<i>The MIP Renderer</i>	41
	Basic Usage	41
	Advanced Usage	43
	Simple Animations	43
CHAPTER 9	<i>The SFP Renderer</i>	45
	Basic Usage	45
	Advanced Usage	47
	Simple Animations	48
CHAPTER 10	<i>The Surface Renderer</i>	49
	Basic Usage	50
	Advanced Usage	51
	Simple Animations	52

CHAPTER 11	<i>The Movie Maker</i>	55
	An Overview	55
	Creating and Adjusting Keyframes	56
	Using the Storyboard.	57
	Working with Movie Projects.	58
	Using the Timeline.	58
	Advanced Topics	59
CHAPTER 12	<i>The Object Stabilizer</i>	61
	Stabilization of 3D Time Series.	61
	Alignment of Slices in 3D Stacks	66
CHAPTER 13	<i>Introduction to the Object Analyzer</i>	69
	Starting the Object Analyzer	69
	Segmenting the Objects: Setting the Threshold	71
	Interaction with the Objects.	73
	Render Pipes	74
	Object Statistics	75
	Storing your Results	79
	Further Reading.	79
CHAPTER 14	<i>Object Analyzer Geometry Measurements</i>	81
	Iso-surface	81
	Principal Axis.	81
	Length and width.	81
	Sphericity	82
	Aspect Ratio.	83
	More Parameters and Filtering.	83
CHAPTER 15	<i>Object Analyzer Component Reference</i>	85
	Main window components	86
CHAPTER 16	<i>The Colocalization Analyzer</i>	101
	How to use the Colocalization Analyzer	101
	Iso-colocalization object analysis.	104
	Backgrounds vs. thresholds in colocalization	104
CHAPTER 17	<i>The Chromatic Shift Corrector</i>	107
	Starting the Chromatic Shift Corrector	107

	Estimation of the chromatic shifts	108
	Visualization of the chromatic shifts	108
	Editing the chromatic shift vectors	109
	Working with templates	110
CHAPTER 18	<i>The PSF Distiller</i>	111
	Acquiring the Beads	111
	The PSF Distiller Window	111
	Using the PSF Distiller Wizard	113
CHAPTER 19	<i>Establishing Image Parameters</i>	115
	Image Size	115
	Signal to Noise Ratio	115
	Black Level	117
	Sampling Density	117
	Computing the Backprojected Pinhole Radius and Distance	118
CHAPTER 20	<i>Improving Image Quality</i>	123
	Data Acquisition Pitfalls	123
	Deconvolution Improvements	126
CHAPTER 21	<i>Appendix</i>	127
	License String Details	127
	The Point Spread Function	128
	Quality Factor	129
	File Series	129
	Adjusting the Global Color Scheme	130
	Hue Selector	131
	Image Statistics	131
	Setting the Coverslip Position	132
	Excitation Beam Overfill Factor	133
	Brightfield Images	134
	Support and Contact Information	134

This guide is useful for the beginner in deconvolution and for the expert that starts with the Huygens Professional toolkit.

In this introductory chapter we explain the structure of the guide and provide general information: installation instructions, license details, system requirements, contact addresses, conventions used throughout the manual and the status of the current release of Huygens Professional.

What is Huygens Professional?

Huygens Professional is an image processing software package tailored for deconvolution of microscopic images. It enables you to deconvolve and analyze a wide variety of images ranging from 2D widefield (WF) images to 4D multi-channel two-photon confocal images or from spinning disk confocal microscopes. Also people who use experimental set-ups like 4pi-microscopes may benefit from the software.

The Huygens Professional toolbox contains the following features:

Image Deconvolution Functions

- Accelerated Maximum Likelihood Estimation restoration algorithm optimized for low light level images.
- Iterative Constrained Tikhonov-Miller restoration algorithm.
- Quick Tikhonov-Miller and Quick MLE restoration algorithms.
- PSF Distiller and Point Spread Function (PSF) measurement tool box to derive a microscopic PSF from finite sized micro bead images, containing:
 - Automatic alignment and averaging procedure to combine the signal from different micro beads in one or more images.
 - PSF reconstruction tool to correct for the finite size of micro beads.
- Generates a theoretical Point Spread Function for widefield, confocal and two-photon microscopes based on electromagnetic diffraction theory.
- Automatic bleaching correction of 3D and 4D widefield images and 4D confocal and multi-photon images.

-
- z-drift corrector tool for time series that enables you to correct for movement in the z (axial) direction that could have occurred for instance by thermal drift of the microscope table.

Basic Image Processing

- Capability to handle multiple images.
- Time series support.
- Multi parameter (multi channel) image elements (stacked or packed).
- Basic data types: unsigned byte, 16 bit signed integer, 32 bit float, 2x32 bit complex.
- Per image undo/redo capabilities.
- An Image History that keeps track of all operations. This history can be saved along with the files or as a Huygens script to automate tasks.
- Scripting and batch processing environment based on Tcl (Tool Command Language).

Core Image Processing Functions

- Create, delete, copy, copy block, convert, split, join, zoom, rotate, iso-sample, shift, replicate image.
- Add/remove border, shift to sub-pixel accuracy, mirror image, swap image octants.
- Arithmetic operations on two image operands, one image operand and a scalar, mathematical functions on one image operand, soft clipping & thresholding.
- 4D Gaussian filter of arbitrary widths, 4D Laplacian filter.
- Generate solid and hollow bandlimited spheres, generate Poisson and Gaussian noise.
- Real and complex 4D Fast Fourier transforms.

Reporting & Display Operations

- Image statistics.
- Report sampling density with respect to Nyquist rate.
- Image histograms of images with up to two channels.
- Plots of energy flux as function of time and axial position.
- Plots of intensity profiles along lines in any orientation.

Image File I/O

- Reads HDF5, ICS, Imaris classic, TIFF (file series with arbitrary numbering), Zeiss LSM, Metamorph STK, Biorad PIC, Olympus TIFF (Fluoview and SIS), Delta Vision IMSubs (r3d), MRC files, IPLab IPL, Image-Pro SEQ, OME XML, and plain text (TXT and CSV).
- Optionally reads Zeiss ZVI and LSM, Leica LIF, Olympus OIF, and BioVision IPM and IPM Tiff.
- Writes HDF5, ICS, Imaris classic, Biorad PIC, TIFF (Leica style and classic numbered TIFF), OME XML, Delta Vision IMSubs (r3d), plain text (TXT and CSV), and AVI (Mjpeg).

Analysis Functions

- Threshold and label 3D image.
- Analyze labelled objects: compute center of mass, volume and integrated intensity.
- Estimate background.
- Measure distance.
- Plot line profiles.
- Compute image ratio.
- Compute colocalization coefficients.
- Compute the co-occurrence matrix of an image.

Visualization Capabilities

- Thumbnail images.
- Multiple Slicers on one or more images. Each Slicer is able to;
 - show x - y , x - z or y - z slices for selectable points in time while optimizing contrast on a global or per-plane basis;
 - besides x - y , x - z and y - z slices, show slices in any arbitrary orientation;
 - display multi channel images in spectral or false color;
 - report individual pixel/voxel positions and values;
 - swing through planes or time;
 - dynamically link the position, orientation, time, and zoom.
- View your data from three sides simultaneously using the Orthogonal Slicer.
- Maximum intensity projections (MIP) of multi channel images in any orientation.
- Volume rendering using the SFP renderer. The Simulated Fluorescence Process (SFP) algorithm computes high quality 3D images based on the physics of light.
- Iso-surface rendering.
- Export of sophisticated 3D animations:
 - combine high quality MIP, SFP, and surface renderings;
 - use predefined animation presets;
 - animate the viewpoint, zoom, and any other visualization parameter;
 - use linear or smooth interpolation in transitions;
 - export to AVI and Tiff series.

Huygens Professional can be downloaded from the SVI website¹.

Microsoft Windows

Double click on the Huygens installer executable, e.g. `huygens-370p0.exe`. Double click its icon to start the installation. During installation the directory `C:\Program files\SVI\` will be created by default. After completion the four Huygens icons appear on the desktop. Double clicking on the Huygens Professional icon starts the program.

*Microsoft Windows
64 bit Edition*

Double click the Huygens installer executable, e.g. `huygens-370p0_x86_64.exe`. Note that the 64 bit Windows version will only run on 64 bit editions of Microsoft Windows 7, Vista and XP. During installation the directory `C:\Program files (x86)\SVI\` will be created by default. Both the 32 and 64 bit Huygens versions will be installed in this directory. After completion the four Huygens icons appear on the desktop.

Mac OS X

Double click the package file, for instance `huygens-3.7.0-p0-Leopard-i386.pkg.tar.gz`. The archive manager expands it to a `.pkg` file, which will be placed in the same directory. Double click this file, and follow the installation wizard.

Linux (Debian)

Debian packages are natively used by Ubuntu and other Debian-based Linux distributions. Double click the package file, e.g. `huygens-3.7.0-p0_i386.deb`, and follow the steps in the package manager. To install the package through the command line:

```
dpkg -i huygens-3.7.0-p0_i386.deb
```

1. <http://www.svi.nl/>

Linux (RPM)

RPM (RedHat Package Manager) packages are natively used by RedHat, Fedora, SUSE, and other RPM-based Linux distributions. Double click the package file, e.g. `huygens-3.7.0-p0.i386.rpm`, and follow the steps in the package manager, or install the package through the command line:

```
rpm -ivh --force huygens-3.7.0-p0.i386.rpm
```

After the Installation

After a first-time installation there is not yet a license available. However, still the software can be started. Without a license it will run in *Freeware mode*. The *System ID*, necessary for generating a license, pops up (See Figure 2.1) when opening Huygens Professional and it can be found in the `HELP→LICENSE` menu. The next section explains how to obtain and install a license string.

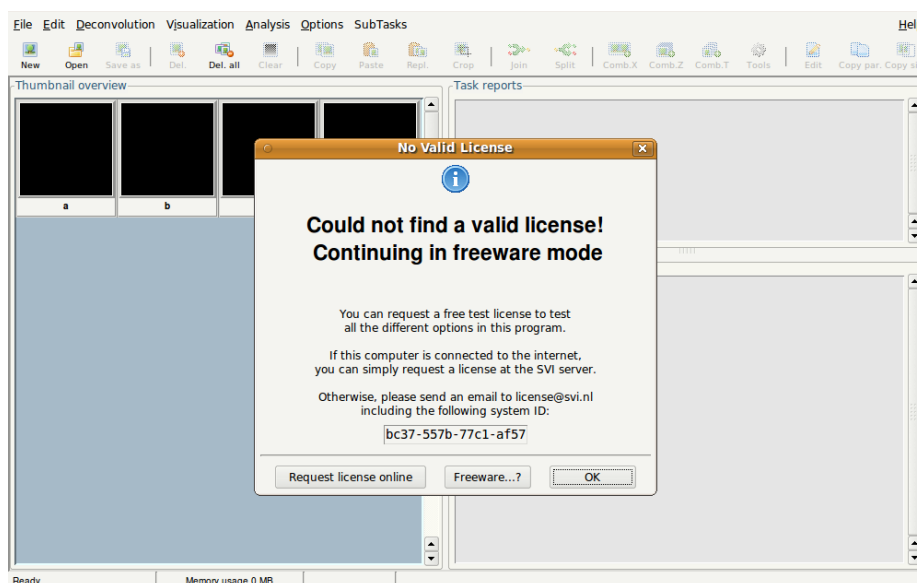


FIGURE 2.1. The startup window of Huygens Professional. If no license string is installed the software runs in *freeware mode*.

The License String

The license key used by all SVI software is a single string per licensed package. It may look as follows:

```
HuPro-3.7-wcnp-d-tvAC-emnps-eom2012Dec31-  
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

At startup Huygens Professional searches for a license file `huygensLicense` which contains a license string. This license string is provided by SVI via e-mail. Installing the license string is the same for all platforms.

Obtaining a License String

If upgrading is not handled from a previous installation it is likely that a license is not yet available. To enable us to generate a license string, we need the *fingerprint* of the computer used, the so called *system ID* number. If Huygens Professional is not already running, please start it. The system ID pops up as long as no valid license is available and is displayed in the HELP→ABOUT dialog (Figure 2.2). Send it to sales@svi.nl, and a license string will be provided. To prevent any typing error use the COPY button to save the ID to the clipboard. It can be printed into the license mail message with the EDIT→PASTE menu item of the mail program.

This dialog box also contains a button to *Check for updates* on the SVI company server.

Installing the License String

Select the license string in the e-mail message and copy it to the clipboard using EDIT→COPY in the mailing program. Start Huygens Professional and go to HELP→LICENSE: a dialog box pops up. Then press the ADD NEW LICENSE button and paste the string into the text field (Figure 2.3). Complete the procedure by pressing ADD LICENSE; this

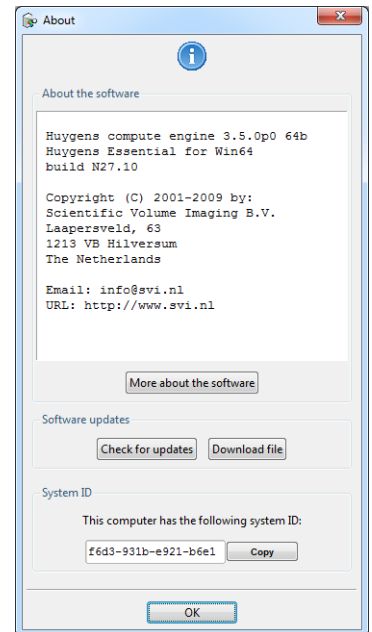


FIGURE 2.2. The HELP→ABOUT window. The system ID is shown at the bottom.

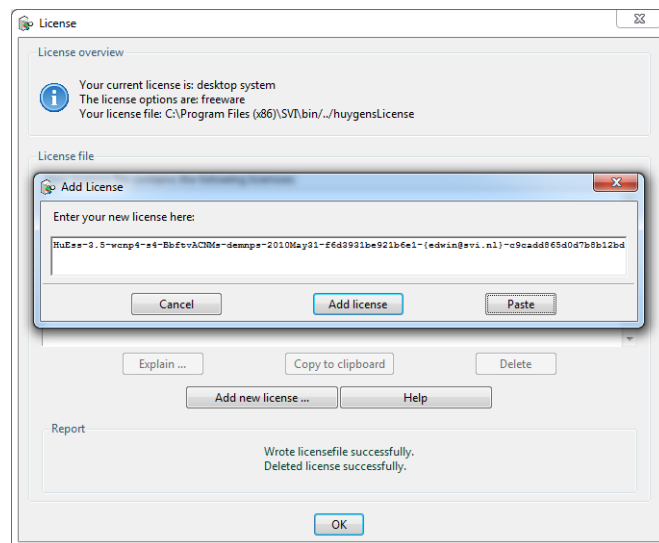


FIGURE 2.3. The license window allows to add, delete and troubleshoot licenses.

will add the string to the huygensLicense file. Please try to avoid typing the license string by hand: any typing error will invalidate the license. With an invalid license, the software will remain in Freeware mode. When the license is correct the message “*Added license successfully*” will appear.

Restart Huygens Professional to activate the new license.

Location of the License File

The license string is added to the file `huygensLicense` in the SVI directory (Table 2.1 on page 6).

TABLE 2.1. The default installation paths per platform.

Platform	Installation path
Windows	C:\Program files\SVI\
Windows 64 bit Edition	C:\Program files (x86)\SVI\
Mac OS X	/Applications/SVI/ ^a
Linux	/usr/local/svi/

- a. The path name on Mac OS X depends on where the software is installed. This is a typical example.

On Irix and Linux and Mac OS X an alternative location is the user's home directory. On OS X this is especially convenient when updating frequently.

Troubleshooting License Strings

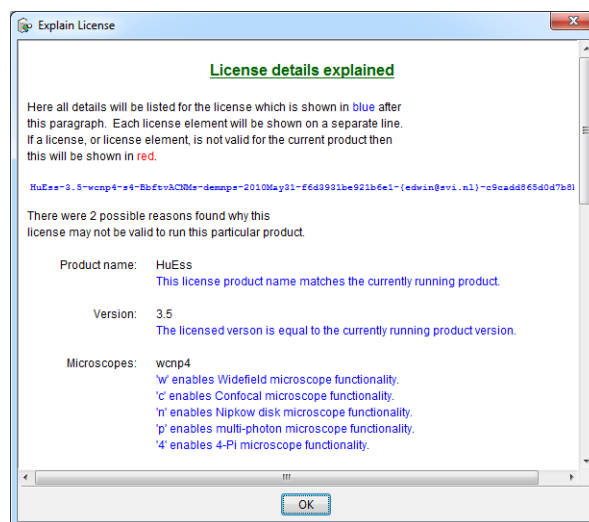


FIGURE 2.4. The *Explain License* window lists all license details.

The license string as used by SVI has the same appearance on all supported platforms. For each product it is required to have a license string installed. Select a license string in the license window (HELP→LICENSE) and press the EXPLAIN LICENSE button. All details for the current license will be listed (Figure 2.4). If running into licensing problems this information can be used to analyze the problem.

Updating the Software

When the system is attached to the internet a pop-up window will appear when a newer version is available. The website can also be consulted for updates. Twice a year (April and October) new releases will become available. During and shortly after this period it is advisable to consult more frequently. Download the new version from the SVI website². Proceed with the installation as explained above.

2. <http://www.svi.nl/>

Do not uninstall the old version as this will delete the license string. The newer version will by default automatically replace the older one. On Mac OS X please make a backup of the license string in a safe place before removing the previous installation.

Removing the Software

Removing the software will also cause the license string to be removed. If it is preferred to uninstall the current version prior to installing a newer one, take care to store the license string in a safe place. Table 2.2 on page 7 shows the uninstallation procedure for each platform.

TABLE 2.2. The uninstallation procedure per platform.

Platform	Procedure
Windows	Open the start menu and select: PROGRAMS→HUYGENS SUITE→UNINSTALL→REMOVE THE HUYGENS SUITE.
Linux	Open the package manager, search for <i>huygens</i> and uninstall it. This could also be handled with the command line; type <code>dpkg -r huygens</code> to install a Debian package or <code>rpm -e huygens</code> to install an RPM package.
Mac OS X	Drag the installed version to the waste basket.

System Requirements for Huygens Professional

Tables Table 2.3, Table 2.4, and Table 2.5 list the requirements for Windows, Mac OS X, and Linux.

TABLE 2.3. System requirements for Microsoft Windows.

Operating system	Huygens runs on Microsoft Windows 2003 Server, XP (32 and 64 bit), Vista (32 and 64 bit), and Windows 7 (32 and 64 bit)
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.
Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

TABLE 2.4. System requirements for Mac OS X

Operating system	Huygens runs on Mac OS X Tiger (32 bit only) and (Snow)Leopard (32 and 64 bit) ^a .
Processor	G5 PowerPC or Intel.
Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

a. OS X 10.5 or higher with X11 is required for full 64 bit capabilities.

TABLE 2.5. System requirements for Linux

Operating system	Most popular distributions like Ubuntu, RedHat, Fedora, and SuSE are supported (32 and 64 bit).
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.

TABLE 2.5. System requirements for Linux

Memory	2 Gb or more.
Graphics card	Any fairly modern card will do.

*Support on
Installation*

If any problem are encountered in installing the program or the licenses which could not be solved with the guidelines here included, please search the support Wiki³ or contact SVI (See “Support and Contact Information” on page 128).

3. <http://www.svi.nl/FrontPage>

This chapter will help you to get through the basic procedures in deconvolving an image and to become familiar with the user interface components of Huygens Professional. Note that the steps 4 and 5 described below are automated in the deconvolution run (step 7), but they will give you some insight in the process of deconvolving an image.

Step 1: Start Huygens Professional

On Linux, start the software by clicking its icon or by typing `huygenspro` into a Unix shell. On Mac OS-X and Windows click its icon. When Huygens Professional is started from a Unix shell, it is a good idea to navigate to your image directory first.

The `huygenspro` command starts a graphical user interface (GUI) that comes up with four empty thumbnail images (Figure 3.1).

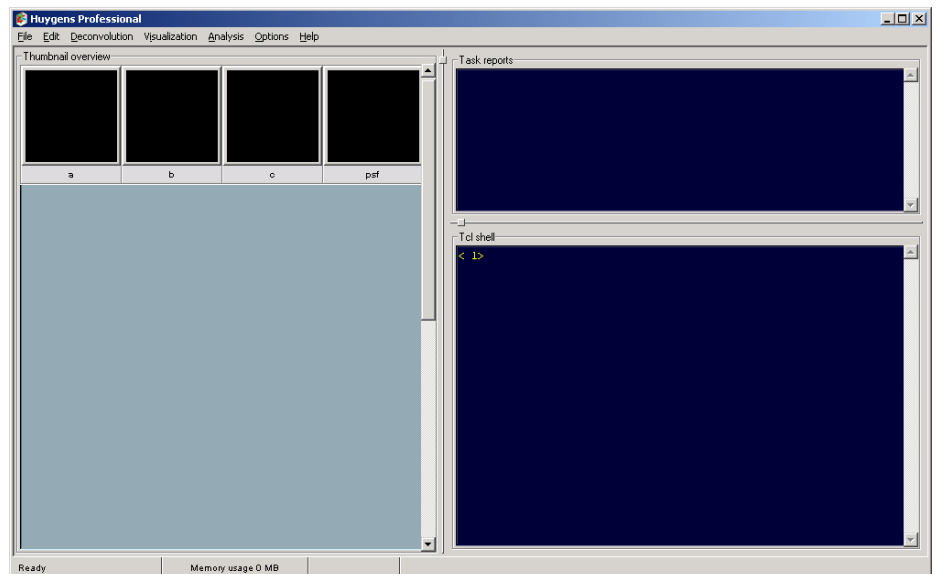


FIGURE 3.1. The main window on Microsoft Windows. Four empty thumbnails are shown by default.

Thumbnail Images

Huygens Professional initially loads four empty images, named `a`, `b`, `c`, and `psf`. These images are tagged with default microscopy parameters, but the image data consists of just zeroes. The fourth image, `psf`, will be used in the following steps as a container for the point spread function.

Step 2: Load an Image

First, an image file has to be opened. Select `FILE`→`OPEN` from the menu in the main window and find the image that you would like to deconvolve, visualize or analyze.

In the distribution you will find the demo image `'faba128'`. The image consists of the file pair `faba128.ids` and `faba128.ics`. The `IDS` file contains the image data and the `ICS` file the meta data, including the microscope parameters. You may click on one of these files to open the image.

In a Unix shell, steps 1 and 2 can be combined in a single command that loads the images automatically after starting Huygens Professional:

```
huygenspro faba128.ics dama.ids&
```

This command loads the files `'faba128'` and `'dama'` and shows their thumbnail images in the main window. The ampersand symbol (`&`) makes sure that the command prompt will return.

If you would like to open a series of Tiff slices, please read “Numbered Tiff Series” on page 129 for the naming convention in order to be able to read a multi-dimensional image as a whole.

Step 3: Inspect your Image

The slicer

Select the `'faba128'` image and open the Slicer via the menu `VISUALIZATION`→`SLICER` in the main window, or by right-clicking the thumbnail and choosing `VIEW SLICER`. The Slicer enables the user to show a single 2D plane extracted from a 3D volume. There are controls to select any plane orientation in space, zoom, and scroll through the available planes. For time series, a separate control is available to scroll through the available time frames. Furthermore, controls are available to set the contrast, brightness and color scheme for the different channels. Color schemes include `'spectral'` colors, i.e. the displayed color of a channel matches the true color of the emission wavelength in so far this is possible.

You can open as many Slicers as you like, on the same image or on different images (Figure 3.2). Multiple Slicers can be linked and `'listen'` to other Slicer windows. When linking is active, right clicking on the main Slicer will change the other view's center to the clicked coordinate, regardless of its orientation. Note that the option `'other'` from the `'linking'` panel has to be selected in order listen to Slicer windows of other images.

More information on the Slicer and other Huygens visualization tools can be read starting from “The Slicer” on page 29 onwards.

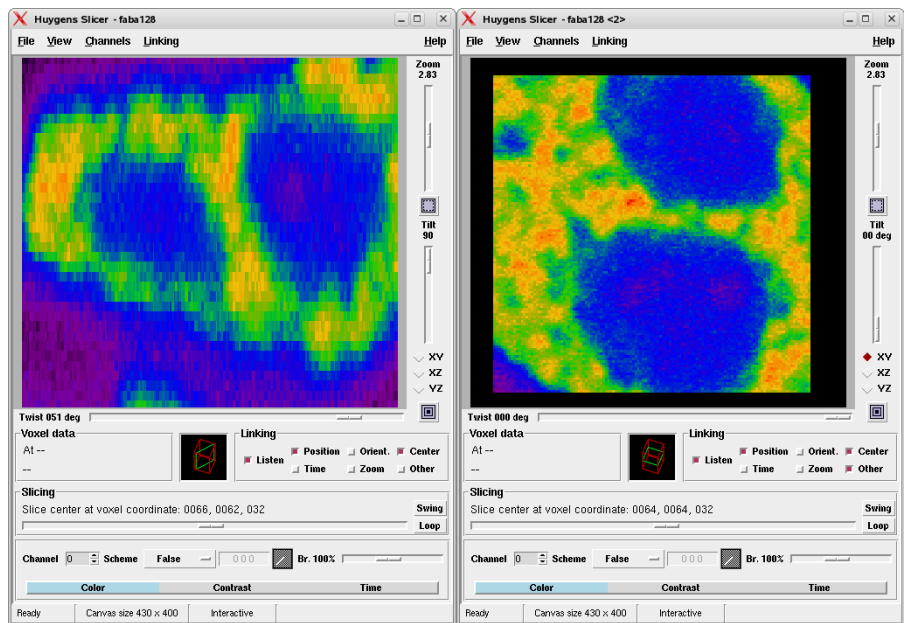


FIGURE 3.2. Two slicers opened from the same image.

Microscopic Parameters

Verify the microscopic parameters (needed for the generation of a PSF) like NA, microscope type, sample size, etc. by selecting EDIT→EDIT MICROSCOPIC PARAMETERS from the main menu (Figure 3.3).

Did you use the correct sampling size during the image acquisition? To find out select EDIT→NYQUIST RATE from the main menu, which displays an information window as shown in Figure 3.4. Read more about this important topic at the SVI-wiki¹. It is also important to make sure that the sample sizes reported by the microscope during the acquisition are correct².

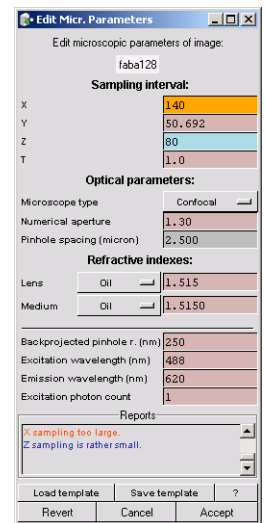


FIGURE 3.3. The microscopy editor. The values in this window should be checked carefully before deconvolution

1. <http://support.svi.nl/wiki/NyquistCalculator>
 2. <http://support.svi.nl/wiki/CalibrateAxialDistances>

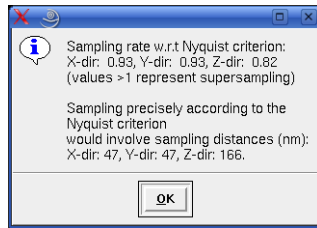


FIGURE 3.4. The Nyquist information window. Values smaller than one indicate undersampling

In case of the 'faba128' image, we see that the image is fairly good sampled in the x and y directions since the Nyquist rate is 0.93, which is near to the value of one. However, the z direction has a Nyquist rate of 0.82, which indicates undersampling (the sampling was too coarse). It is advised to sample in the z direction using a sample distance of 166 nm or less. It can be concluded that if the Nyquist tool leaves you with values smaller than one, the image appears to be undersampled. Your future recordings may benefit from this information. Note that supersampling is not a problem, although your dataset becomes unnecessarily large without providing

extra information.

For now we continue with the z undersampled demo image.

The Intelligent Cropper

The time needed to deconvolve an image increases more than proportional with its volume. Therefore, deconvolution can be accelerated considerably by cropping the image. Huygens Professional is equipped with an intelligent cropper that automatically surveys the image to find a reasonable proposal for the crop region. In computing this initial proposal the microscopic parameters are taken into account, making sure that cropping will not have a negative impact on the deconvolution result.

To start the cropper, select DECONVOLUTION→CROP IMAGE from the menu in the main window. Initially, the cropper suggests that the bottom planes from 'faba128' can be removed without significant loss of information that will decrease the deconvolution quality (Figure 3.5). To adjust this crop region manually, place the cursor inside the red

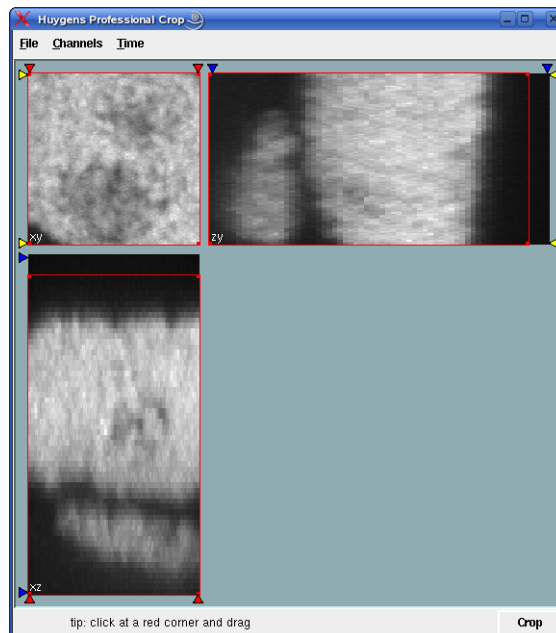


FIGURE 3.5. The cropper window.

boundary, press the left mouse button and drag it to sweep out a volume. Accept the new borders by selecting FILE→CROP. Do not crop the object too tightly, because you would remove blur information relevant for deconvolution.

The three views shown are maximum intensity projections (MIPs) along the main axes. The projections are computed by tracing parallel rays perpendicular to the projection plane through the data volume, each ray ending in a pixel of the projection image. The maximum intensity value found in each ray path is projected. For example, each pixel in the xy projection image corresponds with the maximum value in the vertical column of voxels above it.

By default the projections are over the whole dataset (including all the frames in time series), but this might be confusing sometimes. The small colored triangles can be used to restrict the projections within a specific range of slices. Note that this does not affect the cropping.

Channels from a multi-channel image can be removed by selecting CHANNELS→SELECT CHANNELS from the cropper menu. You also can reduce the number of time frames by selecting TIME →SELECT FRAMES.

The Image Histogram

Select the image `faba128` in the main window and display its histogram via EDIT→IMAGE HISTOGRAM in the main menu. The histogram enables you to visually inspect the intensity distribution of the image visually. The distribution of `faba128`, as seen in the left window in Figure 3.6, is of reasonable quality, i.e. there are no sharp

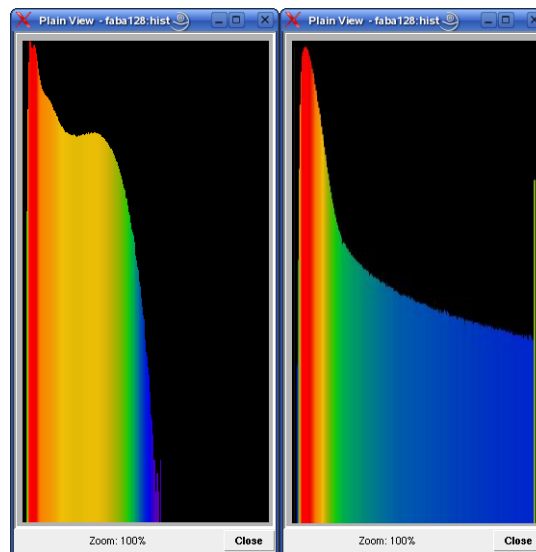


FIGURE 3.6. Two image histogram windows. In the right image, the spike at the right hand side indicates clipping (saturation).

peaks at both sides of the histogram that may indicate clipping. The right window in Figure 3.6 shows an example of the histogram of a clipped image.

Clipping occurs when input signals that are too high, are mapped to the highest value available, or when `negative` input signals are mapped to zeroes by the CCD camera. Clipping at the right hand side of the histogram indicates saturation, as is the case in the right histogram in Figure 3.6. Saturation is caused by overexposure, e.g. by raising the laser intensity above the maximum pixel value available on your microscope. Usually, all values above the maximum value are replaced by the maximum value. On rare occasions they are replaced by zeroes. Clipping will have a negative effect on the results of decon-

volution, especially with widefield images. See “Clipping” on page 124 for more information.

A histogram may also show an offset at the left hand side, which indicates a positive blacklevel. A large positive blacklevel value will reduce the effective dynamic range of your microscope, but will do no harm to the deconvolution with Huygens Professional. For more information, see “Black Level” on page 117 or the SVI wiki page³.

The Operations Window

In the former steps all operations were performed without any ‘operation parameters’, i.e. variables or options applicable to a particular function. For example, if one wants to know the image statistics no extra information is needed but the name of the image. All these types of operations are accessible from the main window. However, most operations, like adding some constant value to an image, need extra input data. This can be done in the ‘operations window’, shown in Figure 3.7. This window can be accessed via

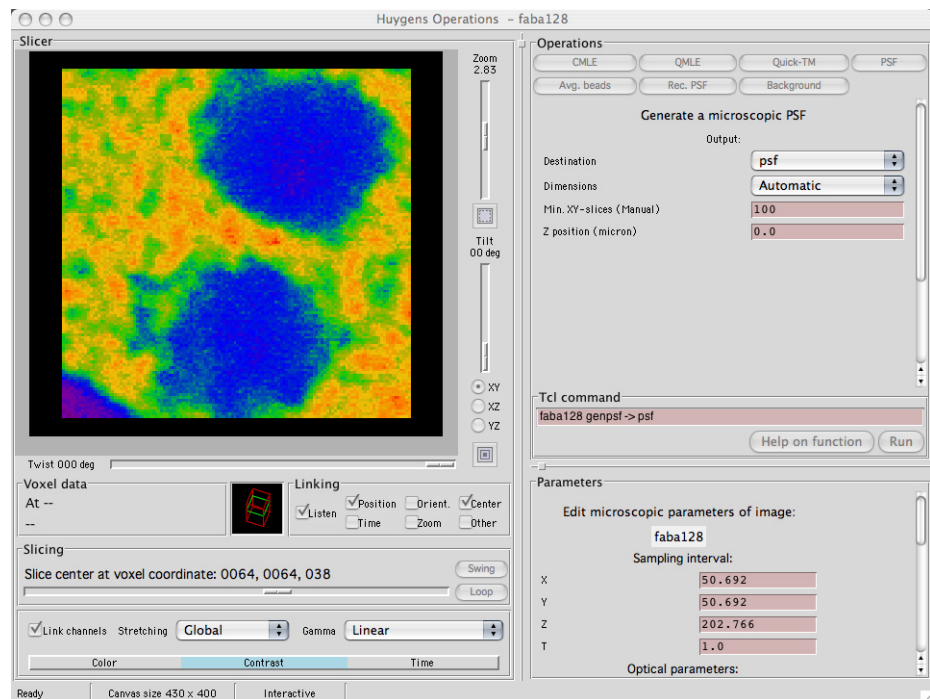


FIGURE 3.7. The operations window as seen on Mac OS X. Various operations can be selected from the menu bar or the quick access buttons (top right).

DECONVOLUTION⇒OPERATIONS WINDOW in the main menu.

In the operations window, an operation may be selected from the menu bar or from the quick access buttons on the ‘operations’ panel at the right hand side of the window. Note that the left side of the operation window is similar to the Slicer. Each function can be controlled by editing the parameters in the ‘operations’ panel. To execute the operation, press the RUN button. Note that the corresponding Tcl command is displayed at the bottom of the panel and, after execution, also at the Tcl shell in the main window.

3. <http://support.svi.nl/wiki/BlackLevel>

Using the operations window, it is also possible to change or convert the dimensions of an image. By selecting a menu item from EDIT→CONVERSIONS, you can e.g. change xyz into xyt or convert a 3D stack into a 4D time series.

Intermezzo: Bleaching Correction

This extra check is given here for the sake of completeness for those using this schematic approach for widefield images or confocal time series.

From the menu in the operations window, select ANALYSIS→PLOT FLUX. A typical example is shown in Figure 3.8. Note that the demo image 'faba128' is a regular confocal image; a bleaching correction test makes no sense here.

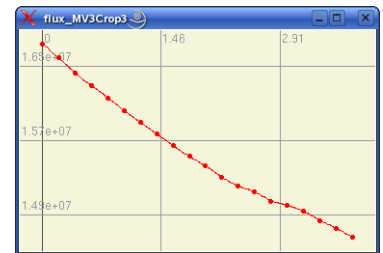


FIGURE 3.8. A flux plot generated by the 'Plot flux' tool. The plot shows a typical example of the intensity distribution along the z axis of a bleached widefield image.

Step 4: Generate a Point Spread Function (PSF)

The point spread function (PSF) is the way in which an imaginary, infinitesimal small light source would be sampled by your microscope. Often this object is no longer a point but blurred and spread out. The aim of the PSF generation is to estimate the amount of blurring along the x , y , and z axes. In the final step of deconvolution the PSF is used to come to a measured deconvolution, in contradiction to the so-called 'blind deconvolution'.

A PSF can be obtained either by *recording small beads* with known bead diameter -- 180 nm beads work fine-- and reconstructing a measured PSF from the bead image, or by *calculating a theoretical PSF* from the information about your microscope settings. Huygens Professional has many tools to handle experimental PSF's, but these are beyond the scope of this user guide.

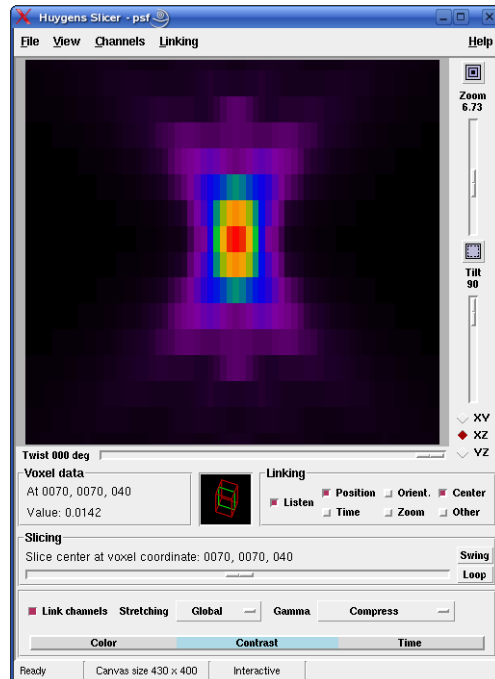


FIGURE 3.9. The PSF in x-z view. False coloring and a compressed contrast highlight the low intensity values showing the typical 'diabolo shape' in lateral view.

A theoretical PSF can be generated by clicking the button PSF from the 'operations' panel in the operations window. This function computes a PSF from the microscopic parameters. By default, the result will be displayed in the image named 'psf', but any available image can be selected as destination. Pressing the RUN button will execute the function. The Slicer can be used to examine the generated PSF in more detail (Figure 3.9).

Note that although this step gives you insight in the PSF, it is actually better letting the deconvolution run generate the PSF on the fly. Especially while using large images that are processed brick-wise (see "Brick-Wise Processing" on page 115), the PSF is prevented from being unnecessary large, it adapts its size to the brick size. In case of a multi channel image, the PSF is generated for one channel at the time instead for all at once.

All images can be cleared using the command EDIT→ZERO IMAGE in the main window. You can do this for the image 'psf' as it will not be used in the next steps.

Step 5: Estimate the Average Background Value

Huygens estimates the average background in a volume image. The average background is thought to correspond with the noise free equivalent of the measured (noisy) image. It is important for the search strategy that the microscopic parameters of the image are correct, especially the sampling distance and the microscope type are of importance.

The average background estimation is invoked by selecting ANALYSIS→ESTIMATE BACKGROUND from the menu in the operations window. The following choices are possible:

- *Lowest value (Default)*

The image is searched for a 3D region with the lowest average value. The axial size of the region is about 0.3 μm , the lateral size is controlled by the 'radius' parameter which is set to 0.5 μm by default.

- *In/near object*

The neighborhood around the voxel with the highest value is searched for a planar region with the lowest average value. The size of the region is controlled by the 'radius' parameter.

- *Widefield*

First the image is searched for a 3D region with the lowest values to ensure that the region with the least amount of blur contributions is found. Subsequently the background is determined by searching this region for a planar region with the lowest average value. Again, the size of the region is controlled by the 'radius' parameter.

In case of the image `faba128', the best choice is `Lowest value', because we are looking for the background outside the object in a confocal image. After pressing the RUN button the calculated value is displayed in the `Task reports' panel of the main window. It is one of the values that should be known in the next steps.

Note that like PSF generation, this step is automated in the deconvolution run (step 7) too.

Step 6: Estimate the Signal to Noise Ratio (SNR)

The signal to noise ratio (SNR, or S/N) is used as a parameter that controls the sharpness of the deconvolution result. It is best estimating these values manually, because automatic SNR estimation is tricky. The easiest way to obtain some SNR is to look at the textures of bright areas in your object image. Figure 19.1 on page 116 shows examples of an image having different noise levels. Inspect your image and decide if your image is noisy (SNR wiki, which is accessible in a fast way by typing the command `help SNR` in the Tcl shell).

Step 7: Perform a Deconvolution Run

In the `operations' panel, there are three deconvolution methods to choose from. These choices are `Classic Maximum Likelihood Estimation' (CMLE), `Quick Maximum Likelihood Estimation' (QMLE), and `Quick Tikhonov-Miller restoration' (Quick-TM). Quick-TM is intended for very special circumstances only. CMLE is the most general restoration method available and should be used in most circumstances. For more information on these restoration methods, visit the SVI wiki⁴.

After clicking the CMLE button, the following parameter values can be adapted (Figure 3.10).

- *PSF (if available)*

From the list of all opened images, an image can be selected that is used as the PSF. This can be either a measured PSF or an earlier generated PSF (using the PSF button from the operations window). If you do not use a measured PSF, it is recommended to select an empty image. In this case, a PSF is `not available' and the software uses the microscopic parameters to generate PSF's on the fly.
- *Destination*

The destination image in which the result will be stored.
- *Signal/Noise per channel*

This is the value that was estimated in step 6.
- *Max. iterations*

The deconvolution process will stop when the maximum number of iterations is reached or when the increase of quality drops below the `Quality change threshold'. Depending on initial quality of your image, a higher number of iterations will result in a better restoration result.
- *Search for background*

The option `Auto' will choose the best estimation mode for you and calculates the background value. In `Manual' mode the background value should be specified in the field below. This value was calculated in step 5.
- *Backgr. per ch. (absolute or %)*

4. <http://www.svi.nl/RestorationMethod>

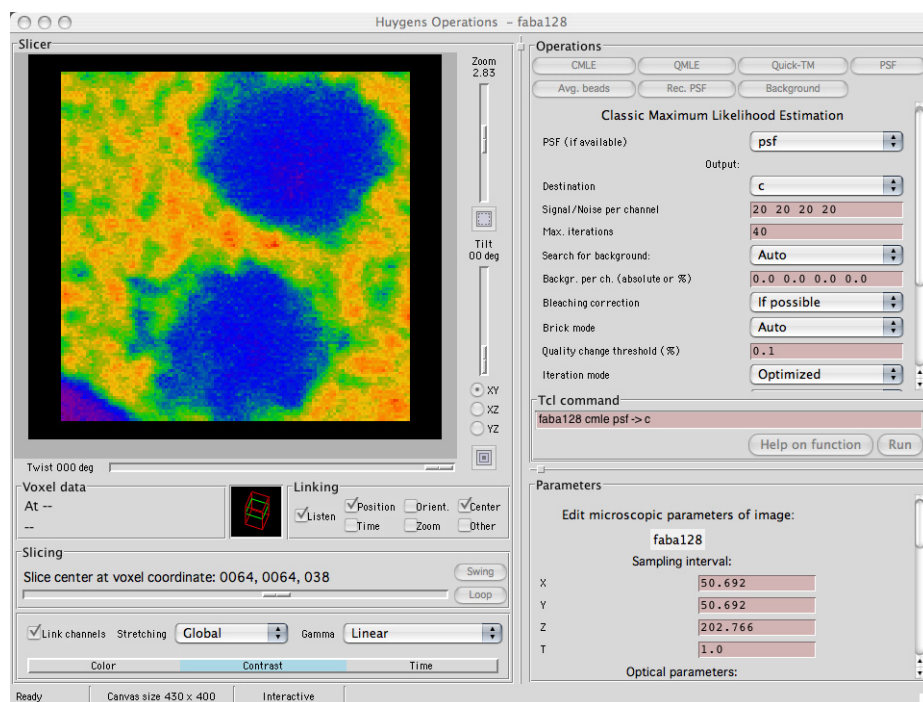


FIGURE 3.10. The operations window as seen on Mac OS X. The right panel shows the parameters for the CMLE deconvolution run.

Multi channel images may have different background values per channel. If the previous parameter was set to 'Manual', this list of values holds the background for each of the channels. Since the backgrounds established by the survey are conservative values (unless 'In/near object' is selected) they can be increased by a percentage. For instance, 10% means all background estimates are increased by 10%. Negative percentages are also valid.

- *Bleaching correction*

If this parameter is set to 'If possible', bleaching correction will be performed for widefield images or confocal time series showing exponential behavior.

- *Quality change threshold*

This number gives the minimum quality increase between subsequent iterations. If the increase of quality drops below the 'Quality change threshold', the deconvolution process will stop. The lower the threshold value and the larger the number of iterations, the higher is the quality of deconvolution.

- Other parameters like 'Brick mode'⁵, 'Iteration mode'⁶, and 'Padding mode'⁷ are best kept in their default modes. More information on these modes can be found in the SVI wiki pages.

Pressing the RUN button will end up with a restored result in the destination image (Figure 3.11). You may compare this result with the original image using two linked Slicers (see step 3 on page 12).

5. <http://support.svi.nl/wiki/BrickSplitting>

6. <http://support.svi.nl/wiki/IterationMode>

7. <http://support.svi.nl/wiki/PaddingMode>

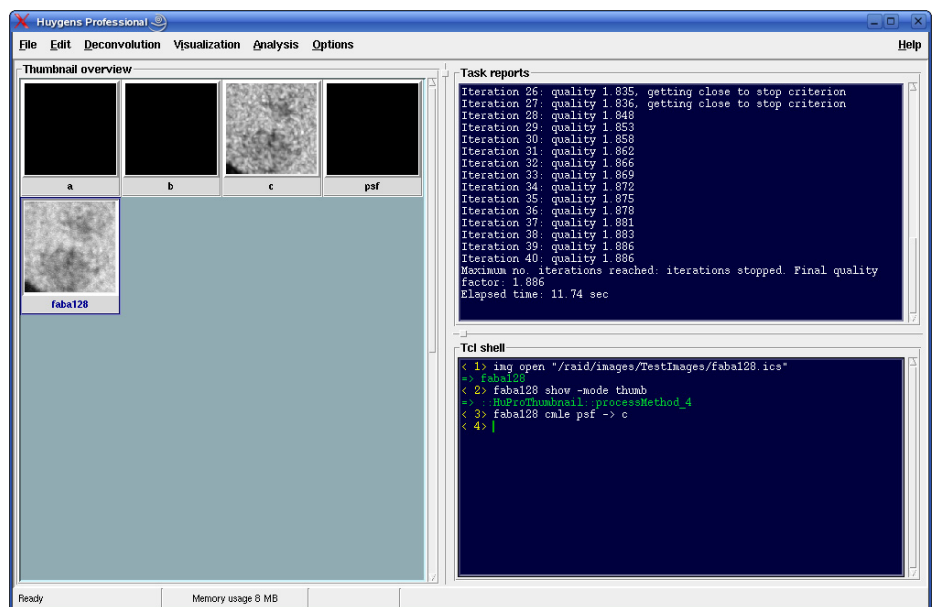


FIGURE 3.11. The main window with a deconvolved image 'c'. Note that the 'Task reports' panel shows information on the deconvolution process.

After inspecting the result you may want to perform another deconvolution run. Maybe you also wish to change one or more parameters. In case the output image destination is not changed, you will be asked if you like to do a completely new run or replace the previous result.

Step 8: Saving your Image

If you are done and wish to save the result, select the thumbnail of the restored image and save the file by selecting FILE→SAVE AS from the menu in the main window. There are different formats to choose from.

Why is automatic deconvolution essential?

If the procedure is not automated, what should you do to deconvolve an n-channel image? You will have to split your parent image into n single-channel images and then generate a PSF, estimate the background and the signal to noise ratio for each of them. Next, you have to run the deconvolution algorithm channel by channel which ends you up with n restored one-channel images. Finally you have to 'join' the separate images to a restored n-channel image. Although all tools are available in Huygens Professional, it is a time consuming task. Not to speak about multi channel time series!

Once you are familiar with a particular kind of dataset and the restoration parameters are determined, a couple or more of similar datasets can be restored automatically. This is called *batch processing*.

A batch process consists of a number of image restoration tasks (one per image) which are executed one by one until all are finished. Depending on the multithreading capabilities of the computer multiple tasks can be executed *in parallel*.

For example batch scripts can be programmed with Huygens Scripting, which makes it possible to run scripts written in Tcl, using the extensive set of Tcl-Huygens image processing commands.

Batch processes can also be configured easily with the interactive Huygens Batch Processor. The Batch Processor is the tool to do large scale deconvolution of multiple images within Huygens Professional.

The Batch Processor Window

To launch the Batch Processor first open Huygens Professional, then click on the menu DECONVOLUTION→BATCH PROCESSOR.

The different elements that form the window (initially with no tasks, see Figure 4.1 on page 20) are:

- **The Save Location.** This is the directory where the resulting images will be placed during the batch run. With the two folder buttons a location can be respectively selected or a new location can be created in the currently selected folder.
- **The Tasks area** shows a list of tasks (empty at start). Tasks are jobs that will be processed by the Batch Processor one by one. Each task line consists of an image, a microscopic template and a deconvolution template. These templates can be updated after a task line is added to the list to tune the values in each particular case. In the *Usage* section this is explained in more detail.
- **The button bar** has at the left side a clock to delay the beginning of the processing, and at the right side buttons to delete, duplicate, stop, start, and add tasks to the list (one by one or many at the same time).

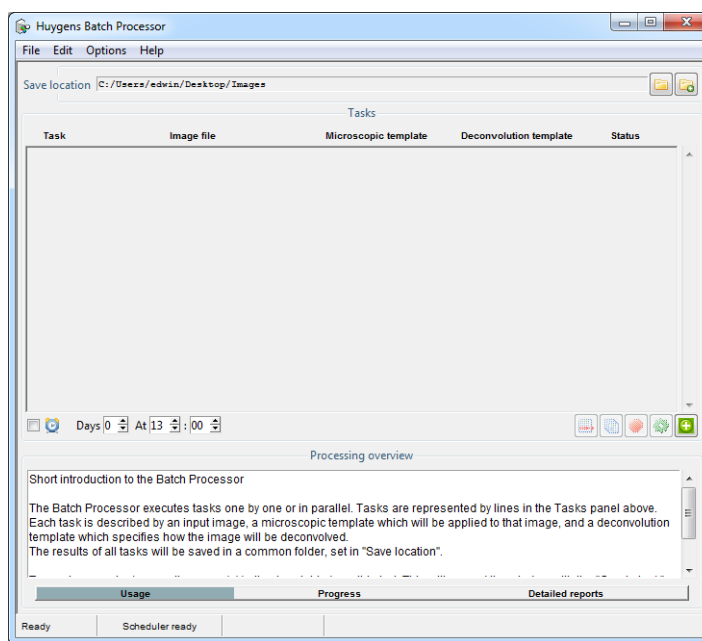


FIGURE 4.1. The Batch Processor main window.

- **The Usage, Progress, and Processing report tabs** in the *Processing overview* area give information about the whole process in its different stages.
- **The status tabs** at the bottom of the batch processor window supply some state information about Huygens Professional. The most left tab shows the state of Huygens Professional and the tab to the right of it shows information of the batch scheduler. The most right of these tabs gives information about the button the mouse is currently pointing at.

Usage

Before to start with creating and running tasks, the places should be defined where the results are saved (*Save location* field), and a file format in which the results should be stored (OPTIONS→OUTPUT FORMAT).

Selecting Input Files

The Batch Processor has a wizard to guide in creating new tasks with only a few clicks. By clicking *add* button (⊕) below the task field a new field titled *Selected images* is expanded at the right (See Figure 4.2 on page 21).

Either a complete folder containing images can be selected or files can be browsed to select single files. If a file is selected containing multiple sub-images (e.g. a Leica LIF file), a secondary menu will pop up to select which sub-image to deconvolve. Each selected sub-image will be added as a new task in the queue.

After selecting the images to restore, the *next* button (⇒) can be clicked to select or create a microscopic template.

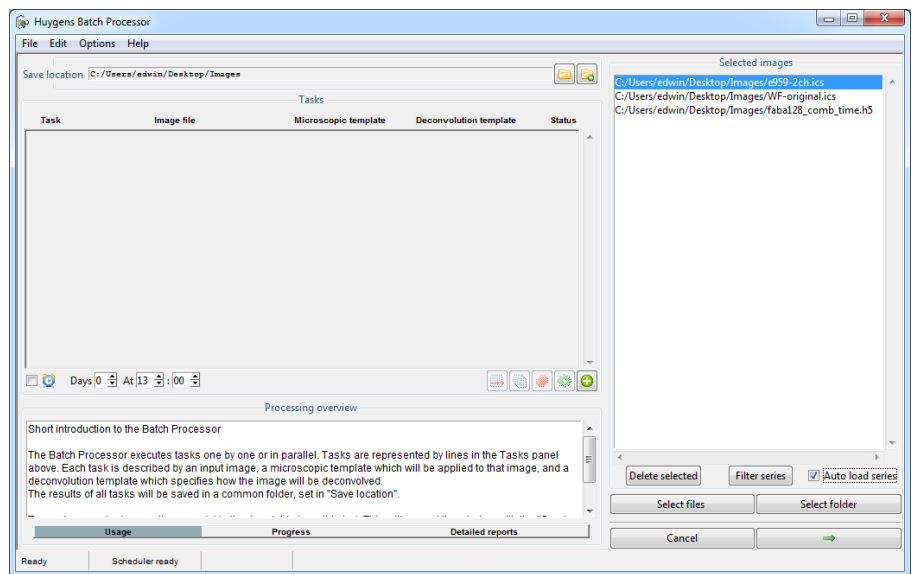


FIGURE 4.2. The Batch Processor with the *Selected images* field expanded.

Microscopic Templates: Describing the Images

To guarantee the quality of the deconvolution it is very important that the image acquisition conditions are properly described. Luckily, the parameters that the Huygens algorithms need are just a few. Because they are so important the Batch Processor does not trust the image metadata included in the file but applies only the entered settings.

This is because different microscope manufacturers report these parameters in different ways and units. It is therefore important to understand what the different microscopic parameters refer to and know how to establish them. A typically conflictive one is the *backprojected pinhole radius*, but once it is understood to what it refers. It is very easy to calculate¹, especially with the assistance of the online *backprojected pinhole calculator*².

The entered parameters can be stored in *templates* for convenience reasons, so the same template can be applied to a series of images acquired with the same physical conditions.

Click the *next* button (→) again when finished with the image description, possibly having reused a previous template.

Deconvolution Templates: Configuring the Restoration Process

Like the microscopic parameters, the deconvolution algorithms are also gathered in a template. Again, it is important that there is understood what these parameters do. Please see “The Deconvolution Stage” on page 15 and the *Restoration Parameters* article³ in the support Wiki.

1. <http://support.svi.nl/wiki/BackProjectedPinholeRadius>

2. <http://support.svi.nl/wiki/BackprojectedPinholeCalculator>

3. <http://support.svi.nl/wiki/RestorationParameters>

A typically conflictive point is setting the *signal-to-noise ratio*. Mind that this is not a number describing the image, but something that can be tuned to achieve different deconvolution results. In practice this is not difficult at all, please see “The Deconvolution Stage” on page 15 and the *Set the Signal to Noise Ratio* article in the support Wiki⁴.

As before, the parameters can be stored in a newly created template, or just a previously saved one can be selected to reuse it.

Saving Templates

The microscopic and deconvolution templates are by default saved in the `SVI/Templates` folder in the user’s home directory⁵, where Huygens can find them. The next time the Batch processor is used, the saved templates will be found in the wizard to set up the batch task.

There are some sample templates in a global location, where the system administrator can also store templates for everybody to use. This global location is a subdirectory `Templates` of the Huygens installation directory (See Table 2.1 on page 6).

Adding the Task

After selecting the deconvolution template the **DONE** button appears. By pushing it the task is created and shown in the *Tasks* field. The task can be deactivated by clicking the **READY TO RUN** button in the status column, setting the status to **DEACTIVATED**. Marked tasks can be deleted by clicking the *delete* button (🗑️) or selecting **EDIT→DELETE TASK** from the menu.

Tasks can still be added to the queue before starting the computations, or start computing right away and add new tasks afterward.

Duplicating Tasks

The *duplicate tasks* button (📄) is very convenient to prepare series of tasks for the same image. Just push the button and the copy of the selected line is ready for modification. For example it might be required to vary one deconvolution parameter to find the optimal value.

Running the Batch Job

When the batch process is configured, its configuration can be saved for future reference by selecting **FILE→SAVE** in the menu.

Add as many tasks as required, single files of complete folders, the Huygens Batch Processor will take care to run them all. By pushing the *start* button (🌀) the Batch processor will start and go over the task list.

The progress of the Batch Processor and the report for each individual task are shown on the tabs in the *Processing overview* area. The status of each task in the task list changes accordingly to the evolution of the process.

4. <http://support.svi.nl/wiki/SetTheSignalToNoiseRatio>

5. The user home directories are usually located in `C:\Users` on Windows 7 and Vista and in `C:\Documents and Settings` on Windows XP. On Mac OS X they are usually in `/Users` and on Linux in `/home`.

The restored images are saved in the selected destination directory as soon as they are ready, along with the image history and an independent task description that can be loaded later in the Batch Processor to re-execute it.

If the computations are very demanding for the system and it should not be blocked, at this moment, the beginning of the queue processing can be delayed by using the timer (🕒). Just adjust the delay in days (zero for today) and set the time of the day in which the processing should start. The timer checkbox is selected automatically; deselect it to disable the timer.

Exiting the Batch Processor

If the batch processor is quit while it is running tasks, the batch processor will stop all running tasks. The Batch processor window can be rolled down while running task. Just exit the batch processor after all the jobs have been finished.

Menus

The FILE menu can be used to store the tasks list for future reference, or in case it should not be started immediately but loaded later for execution. In this menu also the reported information can be saved during the batch processing.

The EDIT menu can be used to duplicate or delete tasks in the list.

The OPTIONS menu has three sub menus:

- **OUTPUT FORMAT:** this sub-menu shows several options for the file format to select for saving the restoration result.
- **THREADS PER TASK:** this sub-menu allows to set the number of processors per job. Typically, in a run where tasks are processed sequentially, the computational work will still be distributed over the available processors, depending on license limitations.

The number of threads Huygens can use in parallel is by default set to AUTO, but in cases where it is required to restrict the computing resources, set the number of threads.

- **CONCURRENT TASKS:** if the system has multiple processors there can be selected to run multiple jobs at the same time. However, it is not necessarily true that concurrent execution of tasks is faster than sequential execution, because in the former case multiple tasks will compete for the available memory (*deconvolution demands a lot of memory*) If the available memory is insufficient, a slowdown will occur.

The Slicer (see Figure 3.2 on page 13) allows you to quickly inspect an image, but also to compare multiple images (e.g. a deconvolution result with the original), or view the same image in multiple orientations simultaneously.

To start the Slicer, select VISUALIZATION→SLICER from the menu in the main window, or double click the image's thumbnail. You can open as many Slicers as you like, on the same image or on different images, and multiple Slicers can be linked and 'listen' to other Slicer windows.

The Slicer enables the user to show a single 2D plane extracted from a 3D volume. There are controls to select any plane orientation in space, zoom, and scroll through the available planes. For time series, a separate control is available to scroll through the available time frames. Pixel intensity values for the cursor position on the image are displayed in the 'Voxel data' panel. You can move the image by clicking the middle mouse button and keeping it pressed while moving the image to the desired position.

Color

For each channel, one out of five color schemes can be chosen:

- *Grey*
Pixel values are assigned different shades of grey ranging from black for the lowest values to white for the highest values.
- *Emission*
The displayed color of a channel matches the true color of the emission wavelength in so far this is possible.
- *False*
Pixels values are assigned different colors ranging from black/dark purple for the lowest values to bright red for the highest value.
- *Global*
In a multi-channel image, colors can be assigned to the different channels by using the 'global color palette'.
- *Custom*

A custom color can be chosen for each of the channels. As with RGB, pixel values are assigned different shades of a particular color ranging from black for the lowest values to the brightest possible shade for the highest values.

Contrast

For the whole view, three basic gamma contrast settings can be chosen:

- *Linear*
In this mode the pixel values are mapped to screen buffer color intensities in a linear fashion. Note that the actual translation of the screen buffer values to the actual brightness of a screen pixel is usually quite non-linear.
- *Compress*
When an image contains a few very bright spots and some larger darker structures, using the Linear mode will result in poor visibility of the darker structures. Restoration of such images is likely to further increase the dynamic range resulting in the large structures becoming even dimmer. In such cases use, the compress display mode to increase the contrast of the low valued regions and reduce the contrast of the high-valued regions. Another way to improve the visibility of dark structures is the usage of false colors, as mentioned before.
- *Strong compression*
Same as above, with stronger effect.
- *Widefield*
In restoring widefield images it sometimes happens that blur removal is not perfect, for instance when one is forced to use a theoretical point spread function in sub optimal optical conditions. In such cases the visibility of blur remnants can be effectively suppressed.
- *Strong Widefield*
Same as above, with stronger effect.

Time

Slices of a time series can be selected by entering a frame number or using the slider.

Linking

When you open two or more Slicers on the same image, then the Slicers are automatically linked. This means the position and orientation of both views are linked by default. A right-click on one of the images will change the center in the other Slicers.

Besides centering and orientation, the 'Time', 'Zoom', and 'Twist' and 'Tilt' sliders can be linked by setting the appropriate option in the 'Linking' panel.

If you would like to link the Slicers of two different images, e.g. a deconvolved image and its original, select 'Other' in order to let the current Slicer listen to other Slicers.

The Twin Slicer allows to synchronize views of two images, measure distances, plot line profiles, etc. In *basic mode*, which is also available without a license, image comparison is intuitive and easy, while the *advanced mode* gives the user the freedom to rotate the cutting plane to any arbitrary orientation, link (synchronize) or unlink viewing parameters between the two images, and more.

To launch the Huygens Twin Slicer, select an image and select VISUALIZATION→TWIN SLICER from the main menu. To view another image in an existing slicer, click the image name in the drop-down menu above the left or right view port (See Figure 5.1).

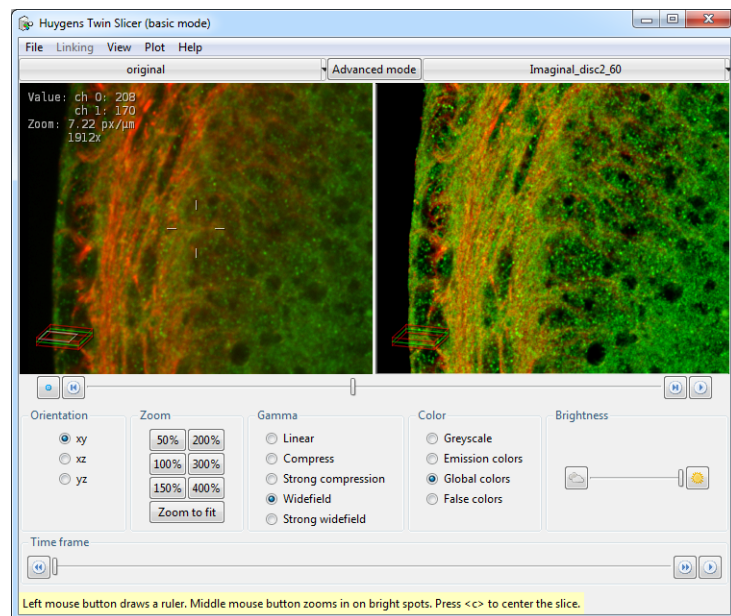


FIGURE 5.1. The Twin Slicer in *basic mode*, showing an original and deconvolved image side-by-side.

The View Menu

Use the VIEW menu to show or hide image properties and guides. These are listed in Table 5.1:

TABLE 5.1. The options in the Twin Slicer’s VIEW menu.

Option	Description
POINTER COORDINATES	Display the position of the mouse pointer in μm or in voxel coordinates.
TIME	Display the time for the current slice in seconds or frame numbers.
INTENSITY VALUES	Display the intensity values for all channels on the current pointer location.
ZOOM	Display the zoom value in screen-pixels per micron. A magnification factor is displayed as well; using the pixel density for the monitor, this value gives an estimation for the absolute magnification.
ROTATION ANGLES	Display the tilt and twist angles in degrees.
DROP SHADOWS	Enhance the contrast for the overlaid lines and text by showing drop shadows.
SLICE BOUNDARIES	Draws the slice boundaries for the left image in the right one and vice versa. This is helpful when both slicers are used.
WIREFRAME BOX	Show or hide the wireframe box, which gives visual feedback on the position and orientation of the cutting plane (green), and the displayed slice (gray) in the data volume (red).
SVI LOGO	Show or hide the SVI logo in the lower right of the view port.

Panning

Click and hold the *right mouse button* on the slice to move it around. Clicking the *center button* (•) or pressing the ‘c’ key centers the slice.

Slicing

Drag the *slider below the view ports* to move the cutting plane back and forth. This can also be achieved using the buttons adjacent to the slider (⏪ and ⏩), the up/down arrow keys on the keyboard, or by placing the mouse pointer over the slider and using the scroll wheel. The play button (▶) moves the cutting plane through the data volume. The pointer coordinates can be displayed through the VIEW menu. Note that it is possible to move the cutting plane out of the volume. Pressing the *center button* (•) or pressing the ‘c’ key centers the plane again.

Using the Slicer in Basic Mode

The button centered at the top of the window enables switching between *basic* and *advanced* mode. In basic mode, all controls are visible in the panels below the view ports (See Figure 5.1). In contrast to the advanced mode, which allows independent control of the left and right slicer (See “Using the Slicer in Advanced Mode” on page 28), the basic mode shows a single set of controls that apply to both slicers.

Changing Time Frames

Drag the slider in the lower *Time frame* panel to change the time frame or press the play button (▶) to animate the time series. The time frame can be displayed through the VIEW menu.

Orientation

Make a selection in the most left *Orientation* panel to change the plane used to display the image.

Zooming

Click the buttons in the *Zoom* panel or use the scroll wheel to zoom in or out on the location of the mouse pointer.

Changing Display Colors

Click an option in the *Color* panel to select a color scheme:

- **Greyscale:** the image is displayed in gray tints. For single-channel images, this gives a higher contrast than the emission or global colors.
- **Emission colors:** if the emission wavelengths are set correctly, this gives the most intuitive view.
- **False colors:** a false color is given to each intensity value. This view gives a high contrast and makes it easy to spot areas of homogeneous intensity.
- **Global colors:** the colors as defined in the *global color scheme*. The global color scheme applies to all visualization tools and can be modified via the Huygens Professional main menu: OPTIONS→PREFERENCES...→EDIT GLOBAL COLORS.
- **Custom**

Tuning the Brightness and Contrast

The brightness can be adjusted in the most right *Brightness* panel using the buttons (☁ and ☀), dragging the slider, or putting the mouse pointer over the slider and using the scroll wheel. The *Gamma* panel provides a linear and some nonlinear ways of mapping data values to pixel intensities. These are:

- **Linear** (default): pixel values are mapped to screen buffer color intensities in a linear fashion. Note that the actual translation of the screen buffer values to the actual brightness of a screen pixel is usually quite nonlinear.
- **Compress:** where an image contains a few very bright spots and some larger darker structures using linear mode will result in poor visibility of the darker structures. Restoration of such images is likely to further increase the dynamic range resulting in the large structures becoming even dimmer. In such cases use the compress display mode to increase the contrast of the low valued regions and reduce the contrast of the high-valued regions. Another way to improve the visibility of dark structures is the usage of *false colors* (See “Changing Display Colors” on page 27).
- **Widefield:** in restoring widefield images it sometimes happens that blur removal is not perfect, for instance when one is forced to use a theoretical point spread function in sub optimal optical conditions. In such cases the visibility of blur remnants can be effectively suppressed.

Automatic Panning, Slicing and Zooming

When the *middle mouse button is clicked*, the Twin Slicer will automatically center and zoom in on the brightest spot in a 3D neighborhood around the cursor.

The button centered at the top of the window offers switching between *basic* and *advanced* mode. The advanced mode allows independent control of the left and right slicer. In this mode, all controls are available in twofold and accessible through the tabs in the bottom of the window.

Changing Time Frames

Drag the slider in the *Time frame* tab to change the time frame or press the play button (▶) to animate the time series. The time frame can be displayed through the VIEW menu.

Zooming

Use the scroll wheel to zoom in or out on the location of the mouse pointer, or access the *Zoom* tab. The four buttons in this tab respectively *zoom out* (◀), *zoom in* (▶), *zoom 1:1* (🔍), (the x-sample distance matches 1 pixel), and *view all* (🌐).

Rotation

The three radio buttons in the *Rotate* tab can be used to switch between axial (*xy*), frontal (*xz*), and transverse (*yz*) orientations. The *Twist* slider rotates the cutting plane around a *z*-axis, while the *Tilt* button rotates the cutting plane around an axis in the *xy* plane. The tilt and twist angles can be displayed through the VIEW menu. Note that the wireframe box in the bottom left of each view port gives visual feedback about the position and orientation of the slice.

Changing Display Colors

Click the *Colors* tab key to view the color settings panel. The *Active channels* drop down menu can be used to enable or disable channels.

In addition to the *color schemes* that are available in basic mode (“Changing Display Colors” on page 27), the advanced mode allows the use of *custom colors*. Use the color picker (🎨) to manually select a color for each channel.

Tuning the Brightness and Contrast

The brightness and contrast controls are accessible in the *Contrast* panel. The brightness can be changed per channel, or for all channels at once. The *Gamma* drop down menu provides a linear and some non-linear ways of mapping data values to pixel intensities (See “Tuning the Brightness and Contrast” on page 27 for an overview).

If the *Link channels* box is checked, this means that the way of mapping data values to pixel intensities is the same for all channels; if not, the range is automatically adjusted for to minimum and maximum in each channel.

Linking Controls

The LINKING menu can be used to change the way in which both slicers communicate. The options in this menu are listed in Table 5.2. Note that settings get synchronized once

TABLE 5.2. The options in the Twin Slicer's LINKING menu (accessible in advanced mode).

Option	Description
ADVANCED LINKING	Enables the user to change the linking of the slice position, panning, and rotation. Doing so may lead to complex situations regarding orientation.
POINTER LOCATION	Shows the position of the mouse pointer in the other slicer.
SLICE POSITION	Makes sure that the cutting plane for the right slicer crosses the center of the left slice, and vice versa.
TIME FRAME	Synchronize the time.
ZOOM LEVEL	Synchronize the level of magnification.
PANNING	This does not affect position of the cutting plane, but it shifts the right slice such that the projection of the center of the left slice is in the center of the right slice, and vice versa.
ROTATION	Makes sure that the rotation angles for both cutting planes are the same.
ACTIVE CHANNELS	The left and right slicer will have the same channels enabled and disabled.
COLOR SCHEME	Makes sure that the left and right slicer use the same colors scheme.
CUSTOM COLORS	Use the same custom color scheme for both slicers.
BRIGHTNESS	Synchronize the brightness.
GAMMA	Synchronize the gamma setting.

the controls are being used.

Some useful ways of linking the controls are:

- **Comparison mode:** to configure the Huygens Twin Slicer to compare two images, e.g. original and deconvolved, it is best to link all orientation parameters, i.e. slice position, time frame, zoom level, panning and rotation. This ensures that there is always looked at the same piece of data.
- **Orthogonal mode:** to view a part of an image in two orthogonal directions, for instance axial (xy) and frontal (xz), do the following:
 - Select the same image for both the left and right slicer.
 - Tick ADVANCED LINKING and link the slice position, time frame, zoom level, and panning. Unlink the rotation.
 - Select the *Rotate* tab at the bottom of the window and select the xz and xz orientation.

Now it is possible to zoom, pan, and slice while the centers of the left and right slice are always aligned. Note that when the cutting planes are not the same, the projected mouse pointer will show a distance (in μm) besides it. If this number is positive, it means that real pointer is more towards the observer (in front of the screen).

- **Overview mode:** An easy overview mode can be configured as follows:
 - Select the same image for both the left and right slicer.

- Tick **ADVANCED LINKING** and link the slice position, time frame, and rotation. Unlink the zoom level and panning.
- Drag the sash to the left to make the left slicer a bit smaller.
- Select the *Zoom* tab at the bottom and click the *view all* button (🏠).

Now the right slicer can be used to zoom in on the data, while the left slicer shows the position in the image (See Figure 5.2).

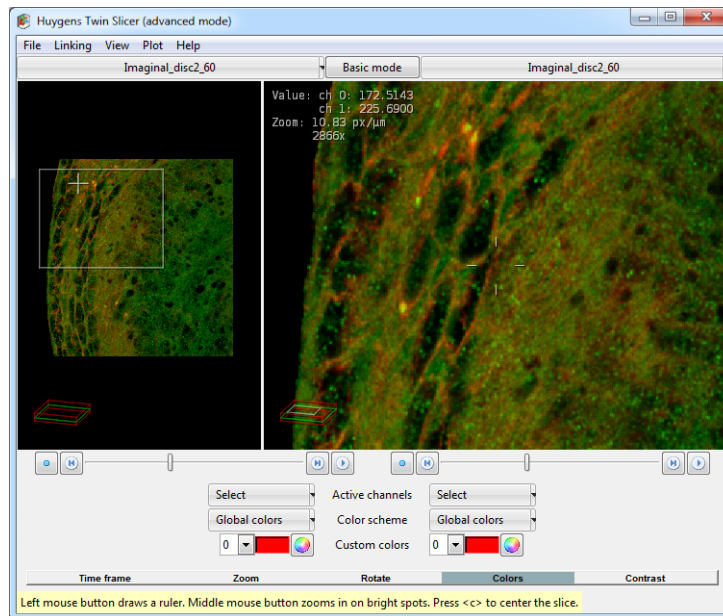


FIGURE 5.2. The Twin Slicer in *advanced mode*, with all controls but zoom and panning linked.

Measurement

Markers

Double click in one of the images to place a marker at the position of the mouse pointer. As configured in the VIEW menu, the marker shows the coordinates and intensity values besides it. To remove the marker, click it and press the `Delete` key.

Rulers

To overlay a ruler on the image, *hold the left mouse button and drag*. The length of the line in μm is displayed beside it. Click and drag the end points of the ruler to make adjustments. Press and hold the `Ctrl` key while dragging an end point to change length without changing direction. Click and drag the middle of the ruler to move it in its entirety, without changing length or direction. Press and hold the `Ctrl` key while dragging the ruler to move it perpendicular to its direction. To remove the ruler, click it and press the `Delete` key.

Intensity Profiles

When a ruler in the left slicer is selected, the right slicer will be replaced by a plot window and vice versa. See the online SVI support Wiki¹ for more information on the data plotter.

Select PLOT→PLOT BOTH SLICERS from the menu to show the intensity profiles for both the left and right image in the same plot. The graphs for the left slicer will have solid lines, while the graphs for the right one are dashed (See Figure 5.3).

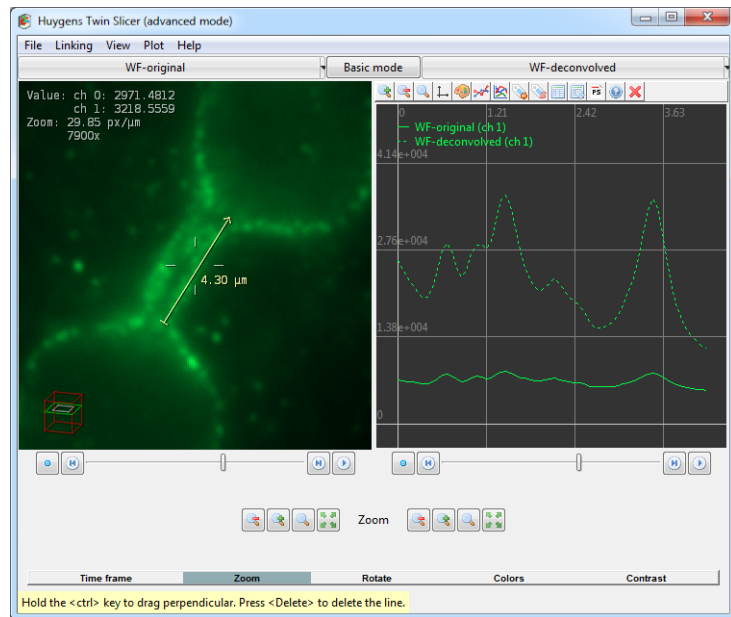


FIGURE 5.3. Measuring the intensity profile along a line. The plot can be configured such that it shows the profile of both images (left solid, right dashed).

1. <http://support.svi.nl/wiki/DataPlotter>



The Huygens Orthogonal Slicer, shown in Figure 6.1, is designed to show the same point in 3D space from three orthogonal directions;

- axial or xy (top left);
- frontal or xz (bottom left);
- transverse or yz (bottom right).

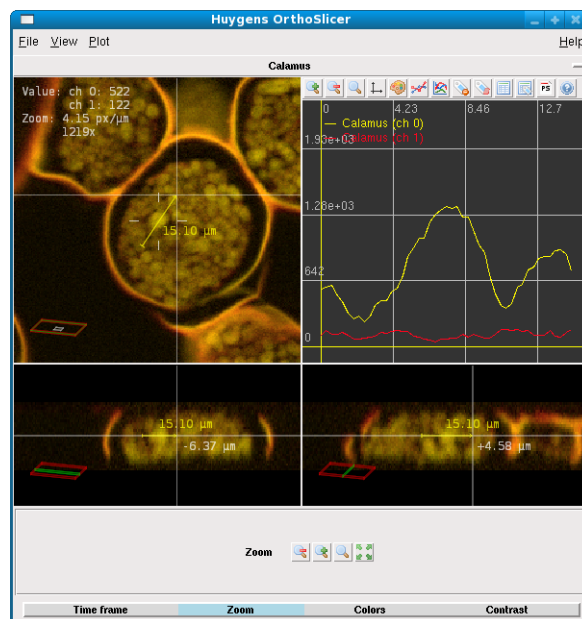


FIGURE 6.1. The Huygens Orthogonal Slicer

If you move one of the slices, the others will follow to make sure that the center of each of the slices intersects in the same point in space. This behaviour makes the Ortho Slicer a useful tool to study small objects in 3D.

The Crosshair Cursor

The position of your mouse is projected as a cross-hairs pointer on all slices. The value besides the center of the cross-hairs gives the distance of the mouse position to this projection. If this number is positive, it means that real pointer is more towards you (in front of your screen).

Visualization parameters

Changing the visualization parameters in the Orthogonal Slicer is similar to the Huygens Twin Slicer (on page 25). There are tools to;

- change time frames
- zoom in, out, fit or zoom 1:1
- change display colors
- tune the brightness and contrast

Panning can be achieved by right-clicking and dragging an image. To center the slice, press the Center button (the blue dot) at the lower left of the image or press 'c'.

Measurements

To overlay a ruler on the image, hold the left mouse button and drag. The length of the line in microns is displayed beside it. Left-click and drag the end points of the ruler to make adjustments. Note that the other orthogonal directions show a projection of this ruler. Press and hold `Ctrl` while dragging an end point to change length without changing direction. Left-click and drag the middle of the ruler to move it in its entirety, without changing length or direction. Press and hold the `Ctrl` key while dragging the ruler to move it perpendicular to its direction. To remove the ruler, left-click somewhere else on the image.

When a ruler is drawn, the help pane will be replaced by a plot that shows the intensity profile along it. See the SVI wiki for more information about the data plotter's capabilities.¹

Auto-Zoom

When you click the middle mouse button, the Orthogonal Slicer will automatically center and zoom in on the brightest spot in a 3D neighborhood around the mouse pointer.

Display Options

The VIEW menu allows you to show or hide information and guides within the image overlay, including pointer coordinates, time, intensity, zoom, rotation, and the wire-frame.

The 'Global value range' option in the PLOT menu uses the maximum and minimum value of the image(s) to determine the visible range of the plot, otherwise it uses the maximum and minimum values of the plot data.

1. <http://support.svi.nl/DataPlotter>

The *Maximum Intensity Projection* (MIP) Renderer enables the possibility to obtain an orthogonal projection of 3D data from any given viewpoint.

The renderer projects the voxels with maximum intensity that fall in the way of parallel rays traced from inside the image volume to the screen (See Figure 7.1). Notice that this implies that two MIP renderings from opposite viewpoints show symmetrical images.

To start the MIP Renderer, right-click on a thumbnail and select VIEW→MIP RENDERER from the pop-up menu.

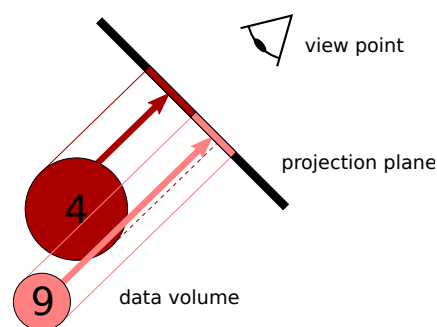


FIGURE 7.1. A schematic overview of MIP rendering. The maximum intensities on rays perpendicular to the screen are projected.

Basic Usage

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 7.2), or by dragging the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel. Use right mouse button to *pan* the center of the projection.

Note that the thumbnail preview (the top right) reflects changes in the configuration instantly, while the large view should be updated manually. To update the large view, press the *fast mode* button (🌿) or the *high quality* button (🌿).

Threshold

The *Soft threshold* slider in the *Channel parameters* panel at the right affects the threshold level. The application of a threshold is a preprocessing step that *reduces the background* in the image, i.e. voxels with intensity values below the threshold value become transparent. Contrary to a standard threshold, which is ‘all or nothing’ (values above the

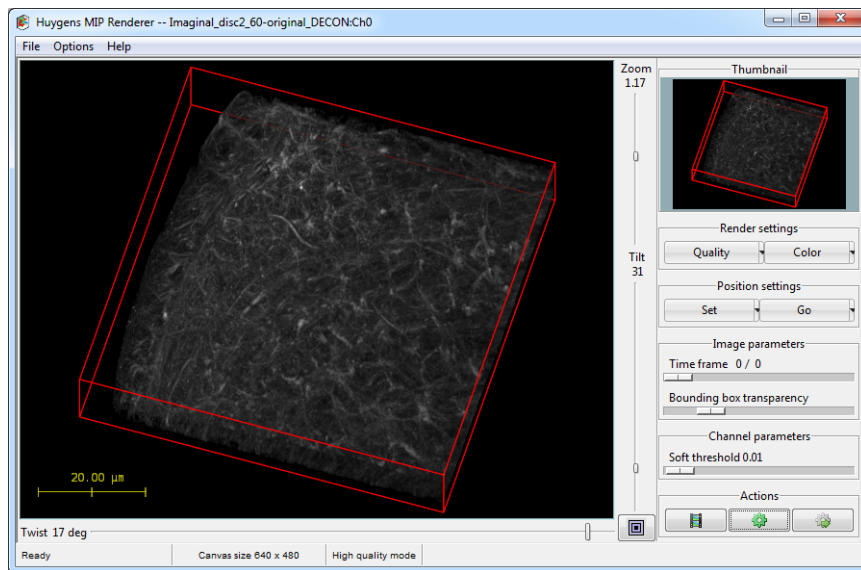


FIGURE 7.2. The MIP Renderer window.

threshold are kept, values below it are deleted), the soft threshold function handles images in a different way. It makes a *smooth transition* between the original and the deleted value.

Saving Scenes

Choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

*Advanced Usage***Render Options**

Table 7.1 gives an overview of the different render options that are available through the

TABLE 7.1. Render options for the MIP Renderer.

Option	Description
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the <i>middle mouse button</i> to pick up and move the rendered image.
RENDER QUALITY	Set the default quality (FAST or HIGH QUALITY). This setting will be used for rendering animations.
COLOR MODE	Choose between GREY, EMISSION COLORS, GLOBAL PALETTE (See “Adjusting the Global Color Scheme” on page 124), or FALSE COLOR.
BOUNDING BOX	Enable or disable the bounding box, or adjust the line color.
SHOW SCALE BAR	Enable or disable the scale bar.
SOFT THRESHOLD MODE	Adjust the smoothness of the soft threshold (See “Threshold” on page 35).
SHOW SVI LOGO	Hide or show the SVI logo at the bottom right.
CENTER SCENE	Undo both the panning of the projection center (right mouse button) and the rendered image itself (middle mouse button).


OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension `.hgsv` and they can be applied to any image that is loaded in the MIP Renderer.

Simple Animations

The Huygens *Movie Maker* (See “The Movie Maker” on page 49) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the MIP Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *animate* button (), and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG¹ codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

-
1. <http://en.wikipedia.org/wiki/Mjpeg>

The SFP Renderer generates *realistic 3D scenes*, based on taking the 3D microscopy image as a distribution of fluorescent material. The computational work is done by the *Simulated Fluorescence Process (SFP) algorithm*¹, simulating what happens if the material is excited and how the subsequently emitted light travels to the observer (See Figure 8.1). The unique properties of this algorithm enable it to create depth cue rich images from unprocessed data. Because it does not rely on boundaries or sharp gradients, it is eminently suited to render 3D microscopic data sets. Since the SFP algorithm is based on *ray tracing* that runs efficiently on multi-core computers, it does not require a special graphics card.

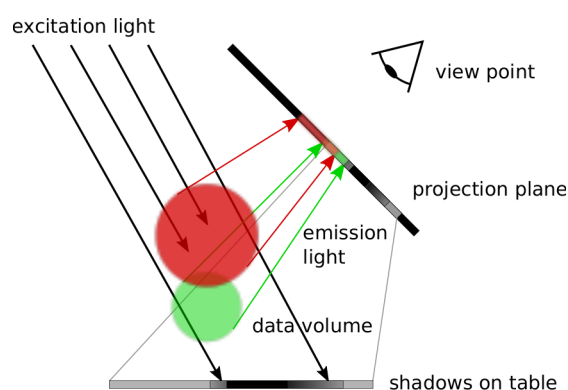


FIGURE 8.1. In the SFP renderer excitation and subsequent emission of light of fluorescent materials is simulated. Each subsequent voxel in the light beam is affected by shadowing from its predecessors. The transparency of the object for the emission light controls to what extent the viewer can peer inside the object.

To start the SFP Renderer, right-click on a thumbnail and select VIEW→SFP RENDERER from the pop-up menu, or choose VISUALIZATION→SFP RENDERER from the main menu.

Basic Usage

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 8.2) or by dragging the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel.

1. <http://support.svi.nl/wiki/SFP>

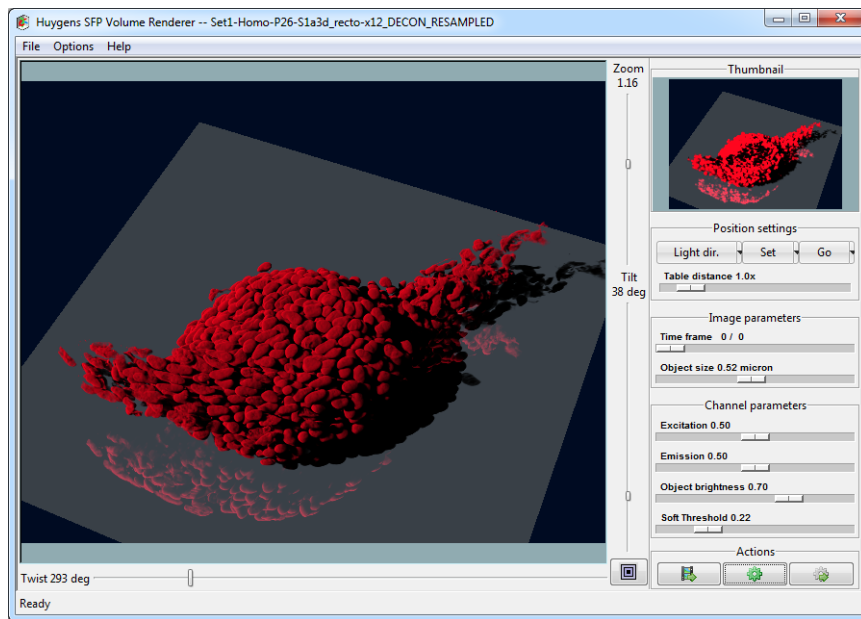


FIGURE 8.2. The SFP Renderer window.

Note that the thumbnail preview (the top right) reflects changes in the configuration instantly, while the large view should be updated manually. To update the large view, press the *fast mode* button (🌿) or the *high quality* button (🌿).

Threshold

The *Soft threshold* slider in the *Channel parameters* panel at the right affects the threshold level. The application of a threshold is a preprocessing step that *reduces the background* in the image, i.e. voxels with intensity values below the threshold value become transparent. Contrary to a standard threshold, which is 'all or nothing' (values above the threshold are kept, values below it are deleted), the soft threshold function handles images in a different way. It makes a *smooth transition* between the original and the deleted value.

Object Size

The characteristic object size can be set by the *Object size slider* in the *Image parameters* panel at the right. This parameter affects both the excitation and the emission *transparency*. While traveling through the object, the light intensity is attenuated to some degree. This enables us to define some definition for penetration depth at which the light intensity is decreased to some extent, for instance 10 % of its initial value. This penetration depth should be in line with the object size. A transparent object is small with respect to the penetration depth. Thus for the same physical properties of the light one object can be transparent while the other is oblique due to its size. To find a reasonable range in transparencies the object size may be altered. The initial object size is computed from the microscopic sampling sizes and number of pixels the image is composed of. If the microscopic sampling sizes of the image are incorrect, then the object size is set according to some default parameters. and may not be related to the actual object size.

Saving Scenes

Choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

Advanced Usage

SFP Fundamentals

The voxel values in the image are taken as the *density of a fluorescent material*. In case of a multi channel image, each channel is handled as a different fluorescent dye. Each dye has its specific excitation and emission wavelength with corresponding distinct absorption properties. The absorption properties can be controlled by the user (See the *transparencies* in Table 8.1 on page 41). The different emission wavelengths give each dye its specific color.

To excite the fluorescent matter light must traverse other matter. The resulting attenuation of the excitation light will cause objects, which are hidden from the light source by other objects, to be weakly illuminated, if at all. The attenuation of the excitation light will be visible as shadows on other objects. To optimally use the depth perception cues generated by these shadows, a flat *table* below the data volume is placed on which the cast shadows become clearly visible. In Figure 8.2 the table is rendered as a mirror.

After excitation the fluorescent matter will emit light at a longer wavelength. Since this emitted light has changed wavelength it is not capable to re-excite the same fluorescent matter: multiple scattering does not occur. Thus only the light emitted in the direction of the viewer, either directly or by way of the semi reflecting table is of importance. By simulating the propagation of the emitted light through the matter, the algorithm computes the final intensities of all wavelengths (the spectrum) of the light reaching the viewpoint.

The properties of the interaction between object and light (transparency), both for excitation and emission, can be adapted interactively by the user to produce different sceneries.

Render Parameters

Table 8.1 gives an overview of all render parameters in the SFP Renderer.

TABLE 8.1. SFP render parameters

Parameter	Description
Light direction	Change the angle for the excitation illumination.
Table distance	Adjust the distance between the object and the table.
Time frame	Set the time frame (in case of a time series).
Object size	Adjust the total transparency of the rendered object. See “Object Size” on page 40
Excitation	Adjust the excitation transparency for the matter in the selected channel.
Emission	Adjust the emission transparency for the matter in the selected channel.
Object brightness	Set the intensity level for the excitation light source for the selected channel.
Soft threshold	Adjust the threshold level for the selected channel. See “Threshold” on page 40.

Render Options

Table 8.2 gives an overview of the different render options that are available through the

TABLE 8.2. Render options for the SFP Renderer.

Option	Description
COMPUTE SHADOW	Enable or disable the computation of shadows.
SHOW SVI LOGO	Show or hide the SVI logo at the bottom right.
SMALL THUMBNAIL	Reduce the size of the preview thumbnail. This enhances the interactivity on slower computers.
TABLE STYLE	The table style can be set to OFF (hidden), NORMAL (grey), and MIRROR (reflective).
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the <i>middle mouse button</i> to pick up and move the rendered image.
RENDER QUALITY	Set the default quality (FAST, NORMAL, or BEST). This setting will be used for rendering animations.
COLOR MODE	Choose between EMISSION COLORS or GLOBAL PALETTE (See See “Adjusting the Global Color Scheme” on page 124).


OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension .hgsv and they can be applied to any image that is loaded in the SFP Renderer, but keep in mind that the sampling sizes of the data affect the transparency.

Simple Animations

The Huygens *Movie Maker* (See “The Movie Maker” on page 49) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the SFP Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *animate* button (), and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG² codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

2. <http://en.wikipedia.org/wiki/Mjpeg>

An iso-surface is a 3D surface representation of points with *equal intensities* in a 3D stack; it is the 3D equivalent of a contour line (See Figure 9.3). The Huygens Surface Renderer is a powerful 3D visualization tool that enables the visualization of these surfaces and thus the representation of the data in a convenient way to clearly see *separated volumes*. Shading enhances the perception of 3D shapes and texture (See Figure 9.1).

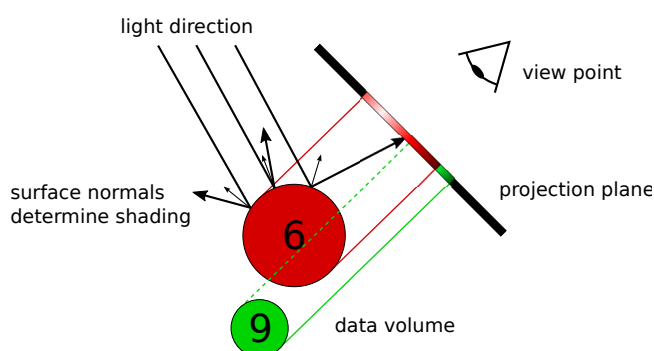


FIGURE 9.1. A schematic overview of surface rendering. The orientation of the surface determines the shading.

Besides iso-intensity surfaces, this renderer is able to generate MIP projections which are blended with the surfaces to be used as a reference to the original microscopic data (See Chapter 7 on page 35).

Because the Surface Renderer is based on *fast ray tracing* that runs efficiently on multi-core computers, there is no need for any special graphics card as would be necessary for conventional polygon based techniques.

To start the Surface Renderer, right-click on a thumbnail and select VIEW→SURFACE RENDERER from the pop-up menu, or choose VISUALIZATION→SURFACE RENDERER from the main menu.

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 9.2) or by dragging

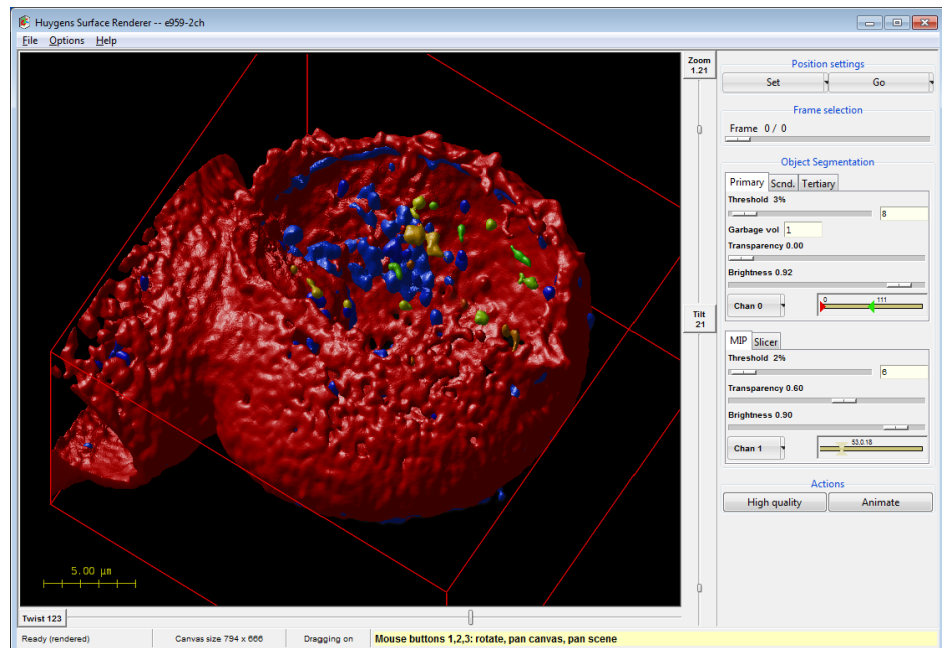


FIGURE 9.2. The Surface Renderer window.

the mouse pointer on the large view. The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel. Use right mouse button to *pan* the center of the projection

Threshold

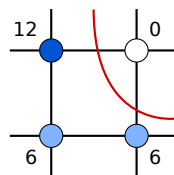


FIGURE 9.3. A contour line for an interpolated value of 5. Because 5 is much closer to 6 than to 12, the distance of the contour to the voxel with value 12 is larger than the distance to the bottom-right voxel with value 6

Use the *Threshold* slider in the *Render pipes* panel to apply different thresholds to the data channels (See Figure 9.3). Voxels that are spatially connected and have intensities above this threshold define *closed volumes*. These volumes are represented by the 3D (iso-intensity) surfaces containing them, each object having a different surface color.

The three *render pipes*, in the *Object Segmentation* frame, referred to as *primary*, *secondary*, and *tertiary*, allow us to define three threshold levels that can be applied to the same or to different data channels. The data channel can be selected using the menu button in the *Object Segmentation* panel. The color range in which the different objects inside a render pipe will be displayed can be adjusted with the hue selector next to it.

Saving Scenes

Press the HIGH QUALITY button in *Actions* panel to apply full scene *anti aliasing* to the rendering and choose FILE→SAVE SCENE... to save the rendered scene as a Tiff file.

*Advanced Usage***Adding a Maximum Intensity Projection**

Besides the surface pipes there are additional rendering pipes to redirect data to the scene. The *MIP* pipe works projecting the voxels with maximum intensity laying in the path of the rays traced along the viewing direction (See Chapter 7 on page 35). In combination with the surface pipes, very clear representations can be obtained of the different objects in the image. The MIP of a channel can be a good spatial reference for the objects in other channels.

Adding a Slice

The *Slicer* pipe is available to represent a single slice of the 3D dataset in its corresponding location.

Render Parameters

Table 9.1 gives an overview of all render parameters in the Surface Renderer.

TABLE 9.1. Surface render parameters

Parameter	Description
Frame	Set the time frame (in case of a time series).
Threshold (surface)	Adjust the threshold level for the selected pipe, i.e. the intensity for which the iso-intensity surfaces are defined. See “Threshold” on page 44
Threshold (MIP)	Set the soft threshold level for the MIP pipe. See “Threshold” on page 35
Garbage volume	Volumes that contain less voxels than defined by the <i>garbage volume</i> parameter will not be rendered. This is useful for rendering only significant objects in noisy images.
Transparency	Set the level of transparency to other pipes.
Brightness	Adjust the brightness for the selected pipe.
Slice Z-position	Set the position of the slice in the <i>Slicer</i> pipe.

Render Options

Table 9.2 gives an overview of the different render options that are available through the

TABLE 9.2. Render options for the Surface Renderer.

Option	Description
ANIMATION FRAME COUNT	Set the number of frames that will be rendered in a movie. 180 frames with a frame rate of 24 fps result in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate; a rate of 24 frames per second is fine for smooth movies.
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the <i>middle mouse button</i> to pick up and move the rendered image.
TRANSPARENCY DEPTH	This option defines how different surfaces are seen through the others: <i>Simple</i> : see through one surface, the surface closest to the viewer. Quite often this is sufficient. <i>Normal</i> : see through two surfaces. <i>Deep</i> : consider many more screening levels, making the final rendering computationally more complex.
BOUNDING BOX	Enable or disable the bounding box, or adjust the line color.
SCALE BAR	Enable or disable the scale bar.
SHOW SVI LOGO	Hide or show the SVI logo at the bottom right.
High quality MIP	Render the MIP pipe in high quality mode.
CENTER SCENE	Undo both the panning of the projection center (right mouse button) and the rendered image itself (middle mouse button).

OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

Templates

All scene settings, i.e. both the render options and all parameters, can be exported to a template file via FILE→SAVE SCENE TEMPLATE. The template files have the extension `.hgsv` and they can be applied to any image that is loaded in the Surface Renderer.

Simple Animations

The Huygens *Movie Maker* (See “The Movie Maker” on page 49) allows to create easily sophisticated animations using the MIP, SFP, and Surface Renderer.

Without the Movie Maker the Surface Renderer has the option to make simple animations of the image, changing the view point in different frames. Set the render parameters for the first frame and click SET→HOME in the *Position settings* panel at the right. Now adjust the viewpoint for the final frame, and click SET→END. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu may be adjusted. Finally press the *Animate* button in the *Actions* pane, and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG¹ codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

1. <http://en.wikipedia.org/wiki/Mjpeg>

The Movie Maker is a tool that allows to create easily sophisticated animations of multi-channel 3D images using the powerful Huygens visualization tools. Animations from the *MIP Renderer* (See Chapter 7 on page 35), the *SFP Renderer* (See Chapter 8 on page 39), and the *Surface Renderer* (See Chapter 9 on page 43) can be combined in a single movie.

The Movie Maker assists the user in creating the *key frames* that define the main scenes, and the smooth *transitions* between them. Interactive manipulation of the scenes is possible using the interfaces of the renderers or by dragging nodes in the *Timeline*.

The movies can be exported to AVI files or to TIFF series that can be combined with other software. Movie projects can be saved for later editing or for usage with other 3D datasets.

An introductory tutorial can be found in the HELP menu at the top right. This interactive tutorial guides the user step by step through the process of creating a simple movie.

An Overview

Figure 10.1 shows the Movie Maker's user interface. The numbered areas are:

1. The *storyboard*: this filmstrip shows the main elements of the movie, which are the keyframes and the transitions between them.
2. The *preview area*: this mini movie player quickly creates a low-resolution version of the movie.
3. The *timeline*: this interactive plot shows how render parameters change over time. Use the menu-button below this timeline to the render parameter to be visualized. Green nodes, representing render parameters at each keyframe, can be dragged vertically.

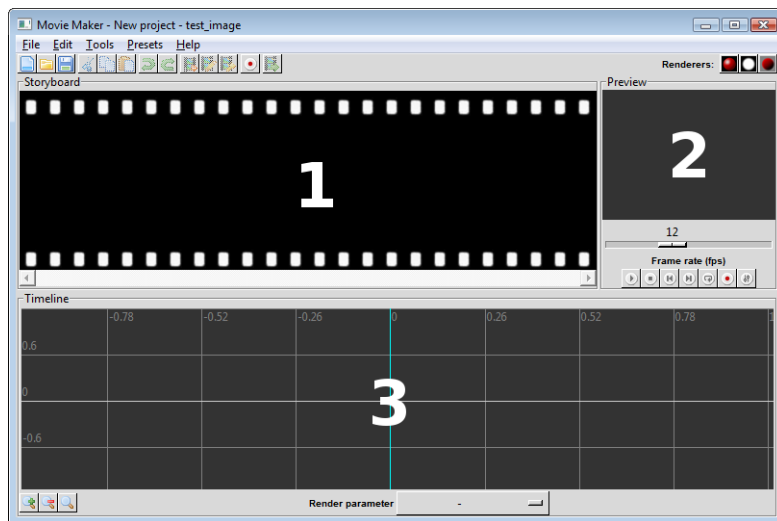


FIGURE 10.1. An overview of the Movie Maker user interface, showing the storyboard (1), the preview area (2), and the timeline

Creating and Adjusting Keyframes

About Keyframes

A keyframe defines a *control point* within a transition. This can be either a start point, end point, or an intermediate point in time. The appearance of the 3D rendered image is *fixed* in these frames. In between keyframes, the Huygens Movie Maker calculates a smooth or linear transition (a technique called *tweening*).

Inserting New Keyframes

To add the first keyframe to the storyboard, one of the renderers should be launched from the Movie Maker window by pressing the corresponding button (■, □, or ■). In the renderer a scene can be defined using the available controls; see Chapter 7, Chapter 8, and Chapter 9 for more information on these windows. Once the looks of the first frame are satisfactory, press the *add keyframe* button (■) to capture this configuration and add the keyframe to the storyboard. All render parameters are captured and can most of them be smoothly animated.

Because a movie needs at least two keyframes (the start and end point of a *transition*), the same renderer should be used to define a second configuration. Once the looks of the second frame are satisfactory, press the *add keyframe* button (■) again. The storyboard will now show two keyframes with an arrow in between. This arrow indicates transition from one keyframe to another.

The Huygens Movie Maker allows the user to mix keyframes from different renderers, but transitions can only be made between keyframes from the same renderer, as shown in Figure 10.2.

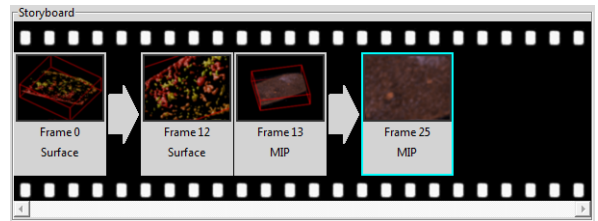



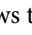
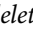
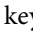
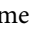
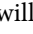
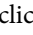
FIGURE 10.2. The Movie Maker's storyboard showing two surface renderer keyframes, followed by two MIP renderer keyframes.

Editing Keyframes

To edit an existing keyframe, double click it or select the frame and choose **EDIT**→**EDIT KEYFRAME** from the menu. This will load the keyframe's settings in the corresponding renderer. The renderer's controls can now be used to adjust the 3D scene. To submit the changes to the Movie Maker, press the *add keyframe* button () again. Because the original keyframe is still selected, the Movie Maker will ask if the original frame should be replaced.

Using the Storyboard

Rearranging Keyframes

The storyboard allows the user to copy, delete, and rearrange keyframes. The *cut* (), *copy* (), *delete* (), and *edit* () buttons are activated whenever a keyframe is selected. When there is no keyframe selected, the *paste* button () will append the copied or cut keyframe to the last frame in the storyboard. If one of the keyframes is selected, the Movie Maker will ask where the frame should be inserted. All these operations can be undone by clicking the *undo* button ().

Changing Transitions

Transitions can be changed by double clicking the arrow that joins two keyframes. This will pop-up a window in which the duration of transition (in frames or seconds) can be changed, as well as the transition type of the parameters that are animated (See Figure 10.3). For most parameters, the Huygens Movie Maker can calculate linear or smooth transitions. An exception to this is the twist, which also requires a direction of rotation (clockwise or counter clockwise) and a number of rotations.

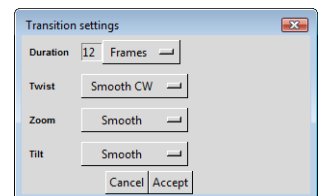
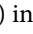
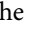




FIGURE 10.3. The transition settings dialog is shown by double clicking a transition arrow in the storyboard.

Playing a Preview Movie

To preview the movie, just press the *play* button () in the *Preview* area (See Figure 10.4). The Movie Maker quickly creates a low-resolution movie and displays it in the preview area. If *loop mode* () is on, the movie will be repeated until the *stop* button () is pressed.

To change the animation settings (*aspect ratio*, *size*, *frame rate*, etc.), press the *render settings* button () in the *Preview* area. The *High quality* setting and the *AVI quality* are not reflected in the movie preview, but only noticeable in the final result.

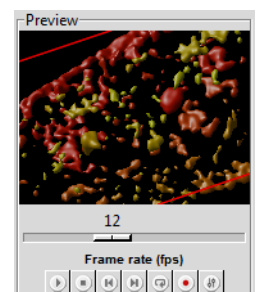


FIGURE 10.4. The preview area.

Export to AVI or Tiff Series

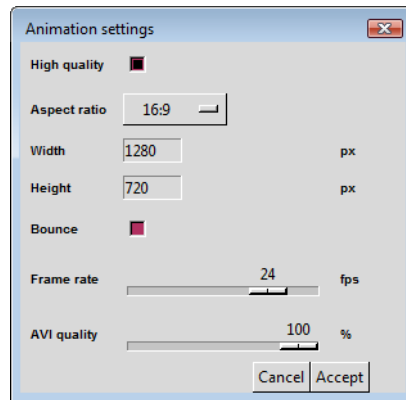


FIGURE 10.5. The animation settings dialog.

If the result is satisfactory, press the record button (●) below the preview (See Figure 10.4) to render the final movie and export it to AVI or a Tiff series. Before the save dialog appears, the Huygens Movie Maker will show the animation settings dialog (See Figure 10.5), where e.g. the AVI quality can be set. Note that large movies will take several minutes to render.

Once the movie has been exported to AVI, it can be opened in a movie player like Apple QuickTime or Windows Movie Player. To quickly open the last saved movie in the operating system's default movie player, press the button labeled "Open AVI in external movie player" (📂).

Working with Movie Projects

Saving Projects

The collection of keyframes, transitions, and animation settings is called a *project*. The image itself does not belong to the project. To save the current project to disk, press the *save* button (💾) or choose FILE→SAVE PROJECT in the menu. The Movie Maker project files have the extension `.hgsa` (Huygens animation template).

The Movie Maker will ask if the project should be saved when one attempts to close the Movie Maker while there are unsaved changes to the project.

Reloading and Appending Projects

If a saved project is reloaded in a Movie Maker that has the same image attached, then the final movie will be exactly the same. However, the Movie Maker allows the user to apply saved projects to different images or append saved projects to the current storyboard. To load or append a project from disk, select FILE→OPEN PROJECT... or FILE→APPEND PROJECT... from the menu respectively.

Besides projects from disk, the Movie Maker has some *presets* that can be appended to the movie. These can be found in the PRESETS menu.

Using the Timeline

Visual Feedback

The timeline is an interactive plot which shows the frame number on the horizontal axis and the value of a render parameter on the vertical axis (See Figure 10.6). It gives a more detailed visual feedback on values of each of the animated parameters. A mouse click somewhere in the timeline area will select the corresponding frame and display a preview of the frame in the preview area. The left and right arrow keys can be used to navigate through the frames.

To zoom in on the timeline, click near the frame of interest and use the scrollwheel or the *magnifying glass* buttons below the timeline to change the zoom level.

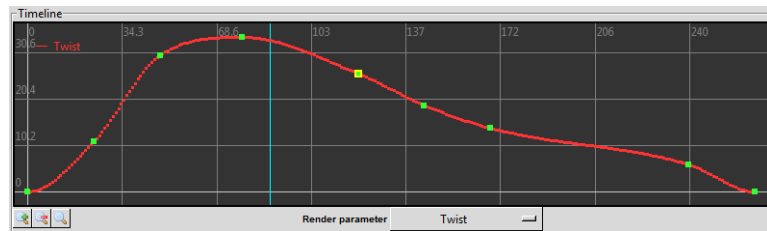


FIGURE 10.6. The timeline area gives a detailed visual feedback on the animated parameters.

Changing Render Parameters

The keyframe nodes are displayed in green and can be dragged vertically to adapt the value of the parameter. The (interpolated) transition frames are displayed as smaller red dots (See Figure 10.6). If a render parameter has been changed, the Movie Maker will recalculate the transitions, update the thumbnails in the storyboard, and update the still of this frame in the preview area.

To select the render parameter which value should be shown in the graph, open the drop-down menu below the timeline and choose the preferred parameter. Because the amount of parameters that can be animated is huge, the menu only lists the ones that change during the movie. If “Other render parameter...” is selected, a dialog window will pop up that shows a list of all available parameters.

In order to change one of the parameters in a frame which is not a keyframe, the frame can be converted into a keyframe. To do this, select the frame in the timeline and choose EDIT→CONVERT TO KEYFRAME from the menu.

Advanced Topics

Stretching Movie Length

The number of frames in a transition can be changed by double clicking the transition arrow. However, if a movie gets complex it is easier to use the *stretch tool* to e.g. double the number of frames in all transitions. This tool can be found in the menu TOOLS→STRETCH MOVIE.

The stretch tool shows the *frame rate*, the *number of frames*, and the *duration* in seconds (See Figure 10.7). The two sliders can be used to change the frame rate and the number of frames; changing these will affect the duration of the movie. The stretch tool tries to redistribute the total number of frames over the whole movie in such a way that the relative length of each transition does not change.

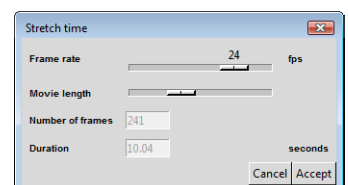


FIGURE 10.7. The stretch dialog helps increasing or decreasing the number of frames in a complex movie.

Synchronizing Transitions in Time and Slice Plane Transitions

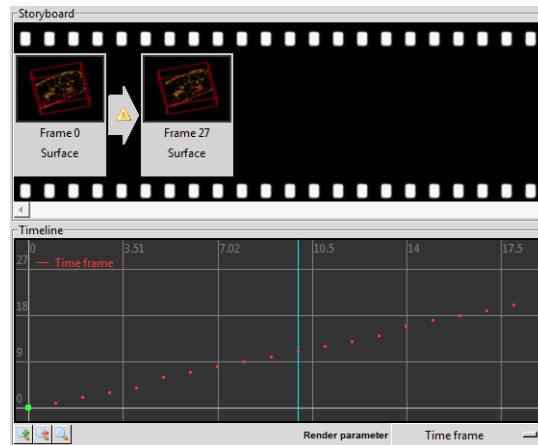


FIGURE 10.8. A transition in which the *time frame* parameter is out of sync with the number of frames. In this example, the transition counts 27 frames, while the *time frame* parameter increases linearly from 0 to 30. Note the irregular change of this parameter in the plot.

Most render parameters, like the *tilt*, *twist*, and *zoom*, can be set to non-integer values. The *time frame* (in case of a time series), and the *slicer z-position*, however, are fixed to integer values. When the transition of such a parameter is not linear, or when the change in value does not match the number of frames, then this parameter is *out of sync*. In that case the Movie Maker will show a warning symbol (▲) on the transition arrow (See Figure 10.8). In the final result, these asynchronous transitions may show irregularities.

To restore the synchronization, right click on the transition arrow and choose SYNCHRONIZE TIME FRAMES... from the pop-up menu. The Movie

Maker will set the transition type to linear and add or remove some frames from the transition to achieve a 1:1, 2:1, 1:2, etc. synchronization.

Creating Loopable and Bouncing Movies

In two simple steps, a movie can be made *loopable*, i.e. it can be played seamlessly in repeat mode:

1. Copy the first keyframe and paste it to the end.
2. Right-click on the final keyframe that has just been pasted, and click SKIP THIS FRAME in the pop-up menu.

The result is the best when all transitions are set to *smooth*.

To create a *bouncing* animation, i.e. an animation that is played in reverse when the last frame is reached, mark the *bounce* option in the animation settings dialog. Doing so will not insert extra keyframes on the storyboard, but the frames are appended in reverse to the final AVI file or Tiff series.

About Movie Quality

In the Huygens Movie Maker, two types of quality can be set:

1. The *High quality* check box determines if the quality of the renderers should be set to the highest possible setting when rendering the final movie.
2. The *AVI quality* scale bar determines the compression level of the AVI file. Set to 100 %, the quality is the best, but then the file size will be large.

These quality settings are not reflected in the movie preview, but only noticeable in the final result.

The Huygens Object Stabilizer measures and corrects for cell motion, thermal drift, shaking, and other types of movement (i.e. translation in x , y , and z , and axial rotation). Both the measurement and subsequent stabilization are done in 3D and at sub-pixel level. The Stabilizer not only stabilizes 2D or 3D time series, but it also allows the alignment of slices within a 3D stack.

To launch the stabilizer, select an image and select DECONVOLUTION→OBJECT STABILIZER from the main menu.

Stabilization of 3D Time Series

Stabilization over time works best if the 3D time series has been deconvolved first. Deconvolution enhances resolution and reduces noise, which helps to analyze motion.

For time series there are four stabilization methods available, which are explained in the next paragraphs:

- *Cross correlation*
- *Model-based correlation*
- *Multi object tracking*
- *Center of mass alignment*

If a 2D time series was read in as a 3D stack, then the Huygens Tcl command `convertZ2T` can be used to convert move z dimension to the time dimension. In Huygens Essential, this function is available from the main menu TOOLS→CONVERT XYZ TO XYT.

Pre-processing

The first step in the wizard is the same for all methods. In this screen there is a button to launch the cropping tool (See “The Intelligent Cropper” on page 14), and a selector for disabling the channels that are not relevant. The channel selector is hidden for single-channel images.

Which step will follow next in the wizard depends on the stabilization method chosen.

The Cross Correlation Method

This can be considered an *all-round method*. It can correct for both x - y - z translation and axial rotation. Adjacent time frames are compared and the program tries to find the best alignment by maximizing structural overlap. A spherical region of interest can be defined to stabilize a particular part of the image.

After the pre-processing step, the *correlation settings* step is shown. This screen allows to enable or disable rotation detection (rotation detection takes more time), and it is possible to select a region of interest in the image, as shown in Figure 11.1. The NEXT button

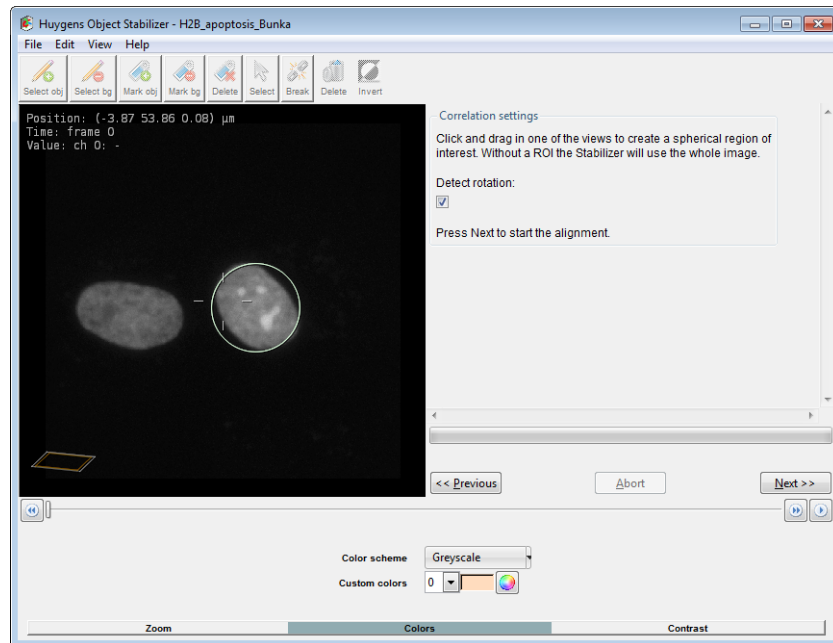


FIGURE 11.1. The *Correlation settings* screen in the Object Stabilizer.

will start the alignment process.

When the Object Stabilizer has finished measuring the displacements, the *stabilization* screen is shown. Continue reading “Stabilization Settings” on page 59.

The Model-Based Correlation Method

When the geometry of the imaged object did not change much during the acquisition, then the time series can be stabilized using a model of the object. The stabilizer creates the model automatically.

Pushing the NEXT button on the pre-processing screen immediately starts the alignment process, because this method does not require any additional user input. When the Object Stabilizer has finished measuring the displacements, the *stabilization* screen is shown. Continue reading “Stabilization Settings” on page 59.

The Multi Object Tracking Method

When your image contains well-defined objects, i.e. nuclei or small particles, then the image may be stabilized using object tracking. Objects are tracked over time and their average movement is used to stabilize the time series.

After the pre-processing step, the *detection settings* step is shown (Figure 11.2). In this screen the detection algorithm can be configured to *filter on brightness only*. This is a little faster and usually sufficient for tracking bright (or dark) spots in images with a high contrast between the spots and the background.

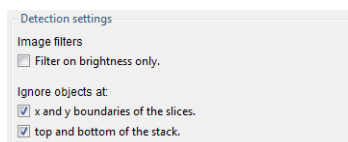


FIGURE 11.2. The *Detection settings* screen.

When the detection algorithm finds many false objects at the image borders (this may happen for very small object sizes), then the Object Stabilizer can be configured to *ignore objects* at *x* and *y* boundaries, and at the top and bottom of the stack.

Next, the user is asked to use the *Select object* tool (🔍) and *Select background* tool (🗑️) to mark a few objects and background areas in the first time frame. Figure 11.3 shows an example of selected objects (green selection) and background (red selection). It is important to make sure that the size of the object selections roughly matches the real size of the objects, because this is a parameter that is used to tune the detection filters.



FIGURE 11.3. An example of object and background selections.

Note that the selections are three-dimensional, i.e. spheres instead of circles. The *x-z* and *y-z* slicers can be used to view and modify the position of the selections along the optical (*z*) axis.

When the NEXT button is pressed, the Object Stabilizer analyzes the selections and uses that information to detect objects in the first time frame. The detected objects are marked by blue spheres as shown in Figure 11.4. When the detection fails, then the selection tools can be used to select new objects and background, mark detected objects as object (🔍) or as background (🗑️), and delete selections. When the selections are modified, pressing NEXT will recompute the detection instead of proceeding to the next screen.

The histogram and sliders in the *optimizing object detection* screen (See Figure 11.4) can be used to apply thresholds on the *number of objects*, *probability*, *width*, and *brightness*. Objects that fall outside of one of these thresholds will be ignored. The *probability* is a statistic that reflects the certainty of the detection algorithm, i.e. objects with a low probability are probably noise. When the thresholds are modified, pressing NEXT will recompute the detection instead of proceeding to the next screen.

In the next screen the tracking parameters can be modified. The detection algorithm can be configured to automatically adapt its settings per time frame (to correct for bleaching, for example). The tracking algorithm can be configured to use the geometry and orientation of the detected objects –besides the position, brightness, and filter values– for connecting the tracks (default).

The NEXT button will start the tracking process. When this is finished the *edit tracks* screen is shown (See Figure 11.5 on page 59). In this stage the user can select (🖱️), break (🔪), and delete (🗑️) the detected tracks. The Object Stabilizer averages the remaining tracks to correct for the average displacement of the objects.

When the NEXT button is pressed again, the *stabilization* screen is shown. Continue reading “Stabilization Settings” on page 59.

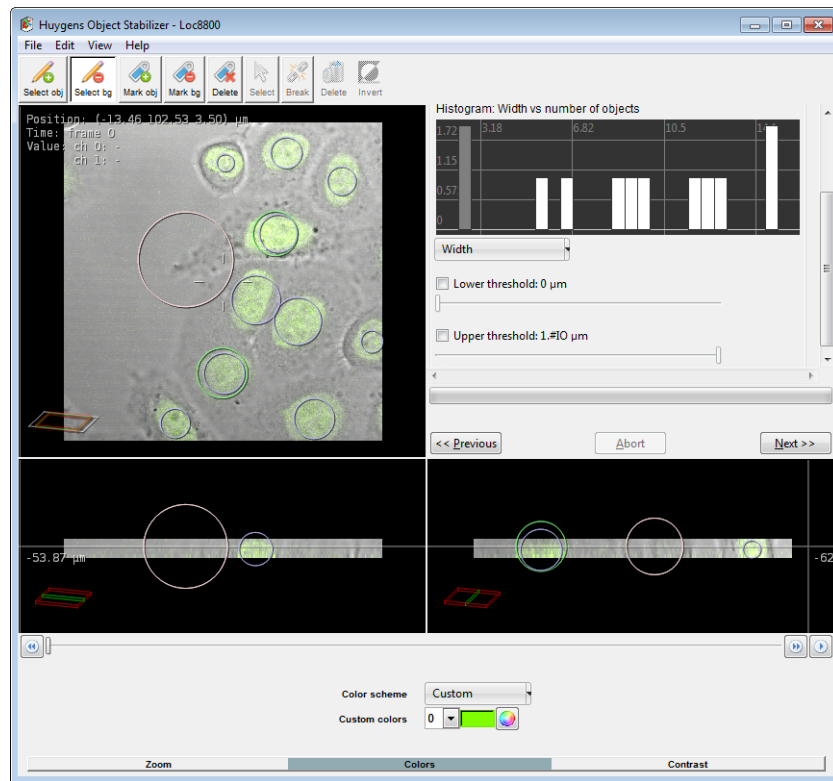


FIGURE 11.4. The *Optimizing object detection* screen in the Object Stabilizer.

The Center of Mass Alignment Method

This method works best if the image contains a single large object. No objects should cross the image borders, and the contrast between object and background should be high.

Pushing the NEXT button on the pre-processing screen immediately starts the alignment process, because this method does not require any additional user input. When the Object Stabilizer has finished measuring the displacements, the *stabilization* screen is shown. Continue reading “Stabilization Settings” on page 59.

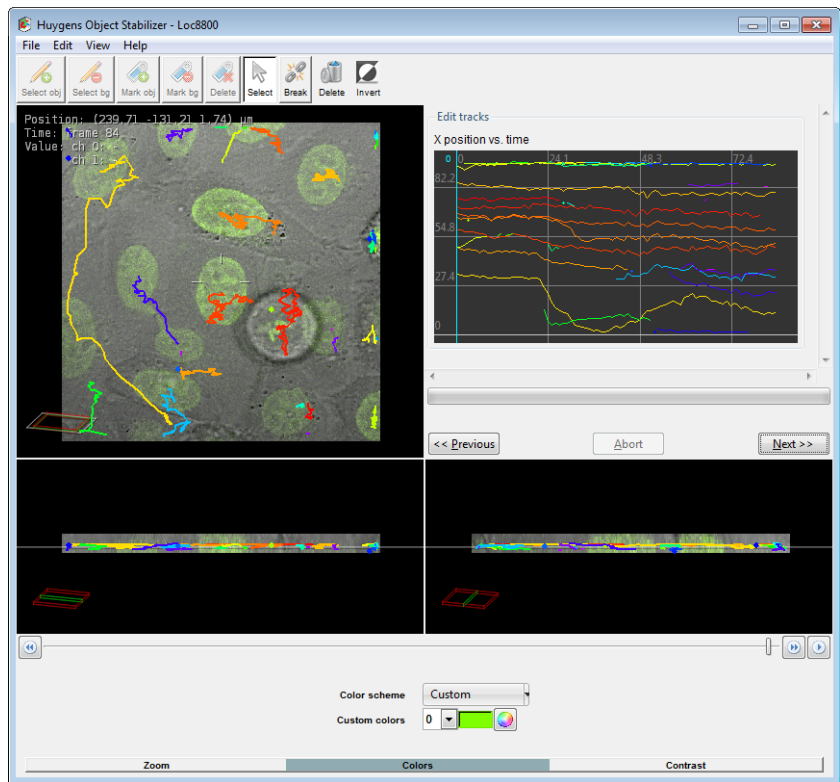


FIGURE 11.5. The *Edit tracks* screen in the Object Stabilizer.

Stabilization Settings

The *stabilization settings* screen shows the measured displacements in x -, y -, and z -direction, and the rotation in four plots. The three sliders above the plots can be used to filter outliers, drift, and noise from the displacement curves. When the displacement curves shows sudden steps, then the outlier filter can be used to remove those steps. The drift and noise filters can be used to respectively remove the drift and keep the random motion, or keep the drift and remove the random motion. The effect of those filters is shown in Figure 11.6.

At the bottom of the *stabilization settings* screen there is the option to select the quality of the interpolation filter that is used to resample the images. The *high quality* mode is not available when correcting for rotation, because it uses Fourier transforms to shift the image.

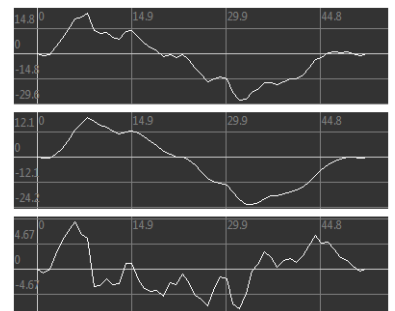


FIGURE 11.6. Effect of filters in the *Stabilization settings* screen. The top shows the measured displacement, the curve in the middle is noise filtered, and the bottom is drift filtered.

Push the NEXT button to apply the stabilization and proceed to the cropping screen.

Cropping the Result

Because there is some extra space needed for the translated and rotated frames, the stabilized image will be larger than the original one. The *cropping* screen (See Figure 11.7)

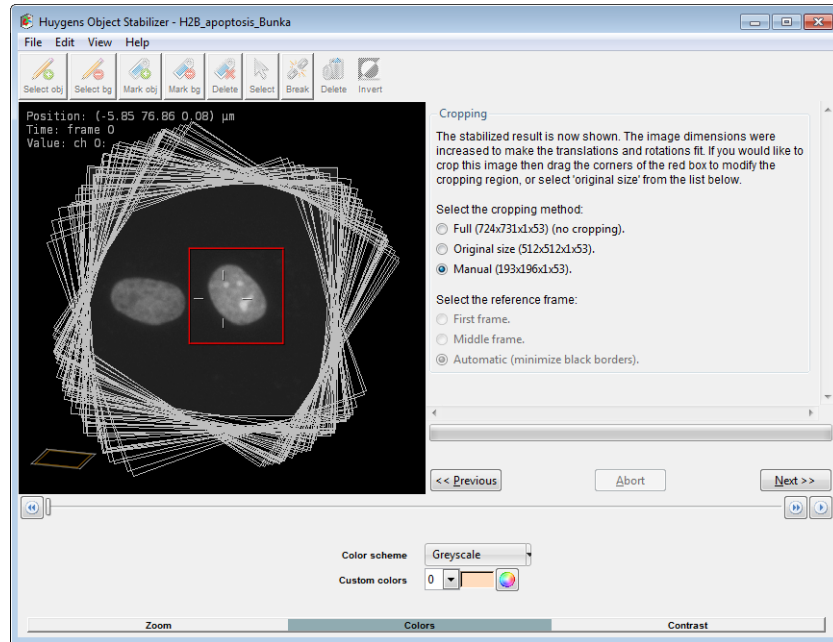


FIGURE 11.7. The stabilization result can be cropped manually to cut out the objects of interest. The original image borders are shown as grey rectangles.

allows the user to manually cut off the black borders, crop the image to the original size, or keep the full size.

Alignment of Slices in 3D Stacks

3D alignment can be applied prior to deconvolution if the stack suffers from *severe* misalignment between adjacent slices. In those cases the Object Stabilizer may improve the deconvolution of 3D stacks and can be applied beforehand. In all other case applying deconvolution first is preferable.

The 3D alignment method always uses the cross correlation method for comparing adjacent frames, and it can correct for x and y translation and rotation. The steps in the wizard are the same as for the stabilization of time series, as explained in “The Cross Correlation Method” on page 56. A comparison between a misaligned and aligned z -stack is shown in Figure 11.7.

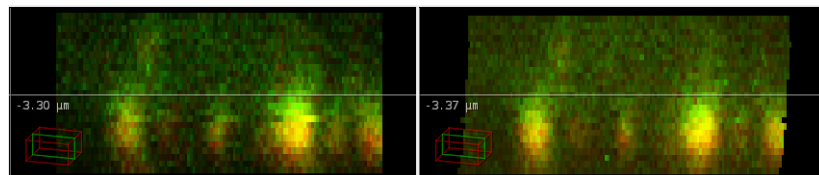


FIGURE 11.8. An x - z slice of a misaligned z -stack (left), and the alignment result (right). The chromatic shift in stack like this one can be correcting using the Chromatic Shift Corrector.

Note that 3D alignment deforms the volume. It is advisable to check the alignment result carefully in using the x - z and y - z views of the Twin Slicer (See “Using the Slicer in Basic Mode” on page 26).

The interactive Object Analyzer (OA) tool allows you to obtain *statistics of individual objects* by clicking on them, or analyzing all objects with a single button press.

In this context, an object is a distinct group of interesting voxels that are spatially connected one to another. Interesting voxels are distinguished from the background by using a seed and threshold criterion. Therefore, defining objects in an image implies:

1. *Segmentation*: Separating interesting voxels from the background according to a given criterion;
2. *Labeling*: Grouping them together and assigning them a distinct name or label.

This is done interactively by the Object Analyzer. To remove too small objects in an early stage from the analysis, a garbage level can be set below which objects are discarded. After that, detected objects are automatically labeled and sent to a continuous iso-surface renderer (See Chapter 9 on page 43).

The Object Analyzer is an extended optional tool, and is enabled by a special flag in the Huygens license string.

This chapter is written in the form of a step-by-step, introductory tutorial to the basic functions of the Object Analyzer. A reference guide that describes all the components of this tool can be found on page 79.

Starting the Object Analyzer

- Launch Huygens Essential or Huygens Professional.
- Load an image you want to analyze. To explore all the OA possibilities, better use a multi channel image.
- Select the image thumbnail, and in the top menu go to ANALYSIS→OBJECT ANALYZER ADVANCED.

You can find this introductory tutorial on-line in the SVI support Wiki¹, from where you can also download the test image we will use in the following explanations.

1. <http://support.svi.nl/wiki/ObjectAnalyzerTutorial>

Please explore the image with the Twin Slicer (See Chapter 5 on page 25) to get an impression of it. This is a deconvolved image, to reduce noise and blur artifacts. It is always a good idea to perform object analysis with deconvolved datasets.

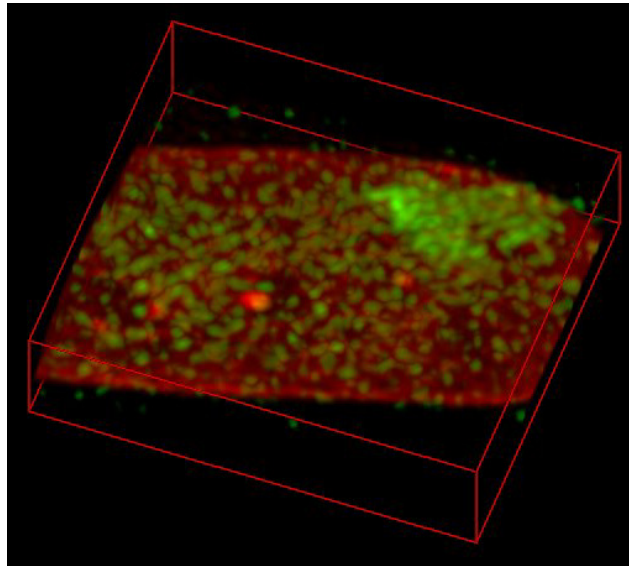


FIGURE 12.1. Maximum intensity projection of the test image. Cell nucleus FISH-stained, recorded at the Nuclear organization Group, SILS, University of Amsterdam (head: Prof. Roel van Driel), under the 3D-Genome research project.

The image in Figure 12.1 is a MIP projection of the test 3D dataset. Notice that there are a few bright objects in the red channel against a more or less homogeneous background (it is actually a quite flat cell nucleus), and lots of scattered objects of different sizes and intensities in the green channel. The red channel is in general very dim with the exception of the inner bright objects and some increase in intensity in its periphery, making something like a border.

When the image is opened it in the OA this is what we first see what is shown in Figure 12.2.

The bright objects in the red channel are recognizable, now from a top view. The image was automatically rendered with some default settings. The intensity range was explored to set a convenient threshold to segment the objects in the first (red) channel, that are shown as iso-surfaces. Every independent object gets a different color, ranging from red to green.

How were these objects separated from the background, isolated from the rest of the intensities in the image so that they could be represented on the rendering as independent entities? They are said to be *segmented*.

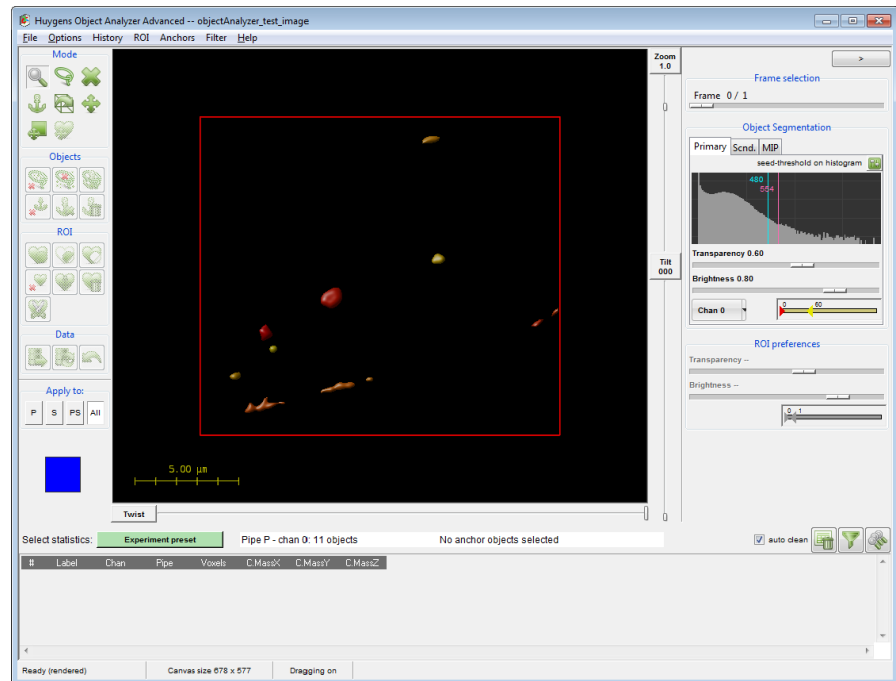



FIGURE 12.2. Start up screen of the Object Analyzer. The objects in the image's first channel are automatically segmented with some default parameters calculated based on the intensity distribution.

Segmenting the Objects: Setting the Threshold

The segmentation method currently available in the Object Analyzer is more flexible than simply setting a threshold: it uses a combined seed-and-threshold method. You can ignore the seed if you want, for starters. It acts as a secondary threshold level, so that objects that do not reach it in intensity (in at least one voxel) are discarded. In this introductory tutorial we will basically ignore the seed, and leave it at the automatic linked value, slightly larger than the threshold itself. For more details on how it works, you may refer to the expert on-line tutorial².

Let's try different segmentation parameters and see what happens. First we lower the threshold value from the automatically calculated value down to something around 360 (the mentioned numerical values refer to this particular test image, of course). You can drag the blue line in the channel histogram and shift it to lower values, or click on the blue-font label showing the threshold value to enter any number. The threshold line on the histogram can be found at the right of the window, in a pane labeled *Object Segmentation*. Next to it, in magenta, there is the seed level. As it is linked to the threshold by default, when you shift the blue line the other one goes with it.

(You can also switch to a percentage representation of threshold and seed by clicking the small button  at the top-right of the histogram. In this alternative slider view, the only one before Huygens 3.5, you can also control whether the seed is linked to the threshold or it remains independent).

2. <http://support.svi.nl/wiki/ObjectAnalyzerExpertTutorial>

We see many things happening when we shift the threshold to lower values. First, the objects we had defined grow in size because more voxels around them get attached to them. Some objects that were separated before grow so much that they now get connected, and they define new single objects. We also get more and more objects of all sizes. To appreciate the details better, you can increase the zoom factor up to 1.30 to see what is shown in Figure 12.3.

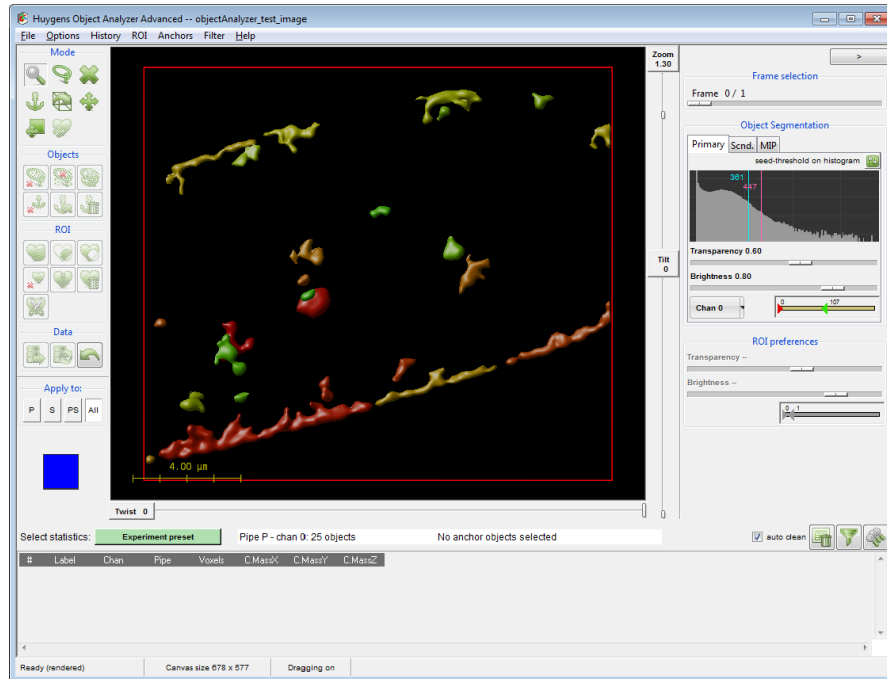



FIGURE 12.3. Different segmentation conditions. Lowering the threshold increases the objects size, and adds more objects to the scene.

In this new segmentation, after lowering the threshold, we have made the original objects larger, but also included many new ones that may be in the way. If you are not interested in all those new objects now, you can try to remove them. Here is where the seed plays an interesting role, but we decided to forget it by now. We can also filter the objects based on their measurable properties, or just use a garbage voxel level, to discard objects that are too small.

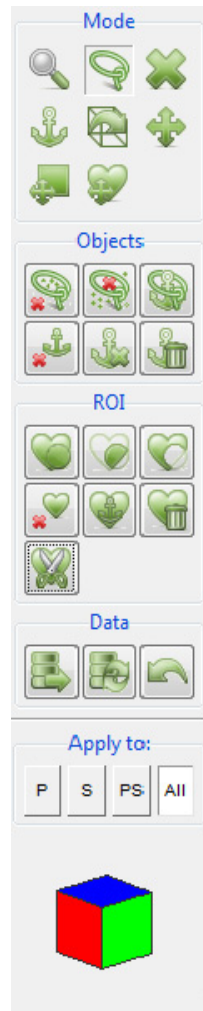
Setting a Garbage Volume Level

A quick way of removing disturbing objects is the GARBAGE VOLUME. You can find this entry in the alternative slider view of the threshold, to which you can switch by clicking the small button () at the top-right of the histogram. The garbage entry will be shown right below the seed slider. Objects with a number of voxels below the garbage level are discarded. This means that when you set it to 1, no segmented object is discarded, but if you set it to e.g. 100, any object with a volume smaller than 100 voxels will be removed.

So far for the segmentation. You can apply some post-segmentation filtering in complex ways, the details are explained in the expert tutorial. In the top menu you can also find some predefined Filters that you can use for quick access to discard objects based on some basic geometrical properties or on the way they relate to other objects.

By now, these two tools (threshold and garbage level) are already powerful enough to continue our exploration of other Object Analyzer features.

Interaction with the Objects



On the left of the window you can find a column of buttons. Many of the buttons are deactivated at this moment, they will be activated when it makes sense later. But almost all in the first group are always available. They control the *Mouse mode*, i.e. what the mouse does when you left-click on the image. When you hover with the mouse over these buttons you get a tip on what they are for. If they have a keyboard shortcut to activate them, it is also shown here. For example, the first mouse mode is intended to Analyze objects, and you can always activate it from within the rendering view by pressing the keyboard key '1'.


These are the currently available mouse modes (from left to right, and from top to bottom):

- **Analyze object** lets you click on different defined objects and obtain the local statistics.
- **Select area** lets you define a 2D region on the current view of the image so you can do different things with it: analyze or discard objects below it, anchor them as references, or interactively define regions of interest (ROI). We will see what all this is useful for.
- **Discard object** lets you discard irrelevant objects one by one. Just select this mode and click on the disturbing objects.
- **Select object** (as anchor) lets you select and deselect objects to be 'anchors', for example to act as references to measure distances from other objects when asking for local statistics. When you set an object as a reference anchor it will 'light up' and change color on the screen to indicate its new status. It is possible to select a group of anchor objects, and you can operate with them through the Anchors menu.
- **Rotate scene** interacts with the full image to rotate it in the space, by dragging the mouse pointer on the rendering view. That can also be achieved by moving the Tilt and Twist sliders along the rendering.
- **Pan scene** interacts with the full image to move it in space *laterally*. This means that you can pan the scene in the 2D plane of your screen, not along the third dimension, along your line of sight.
- **Pan canvas** is similar to 'pan scene' but not exactly the same. It allows you to explore the canvas by not re-rendering the scene. This only makes sense when you have a canvas larger than your rendering window, of course. See `OPTIONS→VIRTUAL RENDER SIZE` in the top menu.
- **Shift the ROI**. This is the only mouse mode that is not always enabled: you need to have defined a *region of interest* (ROI) before moving it around.

The buttons in the second, third and fourth groups are not mouse modes: they do not set new behaviors for your mouse when interacting with the objects view but execute operations on previously defined conditions, for example deleting the objects under a selected

area, or discarding every object that has not been selected as anchor. We will not consider them in this basic tutorial, as you don't need them to perform basic measurements! Please see the "Object Analyzer Component Reference" on page 79 for a detailed description of these toolbox buttons.

At the end of the buttons column there's a colored reference cube that will help you in orienting in space when you rotate the dataset, specially with large zoom factors that do not let you see the surrounding box frame in the rendered image. The initial view of this cube is the blue top face corresponding to $z = 1$. Hovering the mouse over the cube faces brings a tooltip with the face label (x , y and z with values 0 or 1).

The *Analyze object* mouse mode  is the default one. Click on an object and see what happens. You probably already did by now and noticed that some parameters were reported on the table at the bottom of the window. We will see how to report even more in a minute.

First, we need to briefly explain what a pipe is, just to know how to select what channels from the image you want to analyze.

Render Pipes

When we opened the analyzer we got the first channel of the image directly shown on the screen, but we can change that and select the other channel from the multi-channel original image to explore it. For that we simply need to select another channel to be shown in the PRIMARY PIPE: where it reads *Chan 0* you just select *Chan 1* (See Figure 12.4).

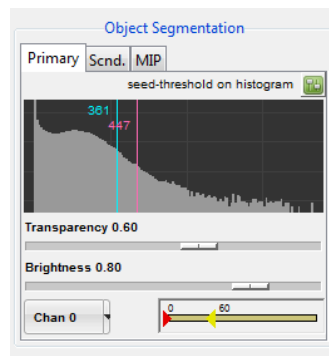


FIGURE 12.4. Rendering pipes control panel.

But what if you want to explore both channels at the same time?

The term *pipe* suggests that data go from your original image to the final rendering in the Object Analyzer through a 'computational duct' in which some processing occurs. You have two of these pipes in the OA to redirect data through.

So we have a secondary pipe too. You can activate its control pane by clicking on the tab that reads SCND. 'Secondary' does not mean here 'less important', it is just that we have the first pipe (*Primary*) and the second pipe (*Secondary*), abbreviated P and S in some places. We could have called them also *Red* and *Green* pipes, independently of their real colors, like it is common in conventional colocalization analysis (and the name may so be changed in future versions of the software). We keep numbers to refer to image chan-

nels, which is something different, because in each of these pipes you can put any channel you want, in any order. You can even segment the same channel twice, with different conditions in each of the pipes.

All the different objects in a pipe are colored differently to be able to distinguish them. The range of colors assigned to each pipe can be controlled with a HUE SELECTOR (See “Hue Selector” on page 125). You can collapse this range completely if you want that all objects in a pipe get the same color.

There is also a maximum intensity projection (MIP) pipe that doesn't interfere in the analysis but that can be used to set a spatial reference to our eyes.

Before continuing our exploration, let's put channel 0 again in the primary pipe with a threshold of 40% (488 in absolute terms), nothing (Off) in the secondary pipe, and channel 0 again in the MIP, so we see something like Figure 12.5.

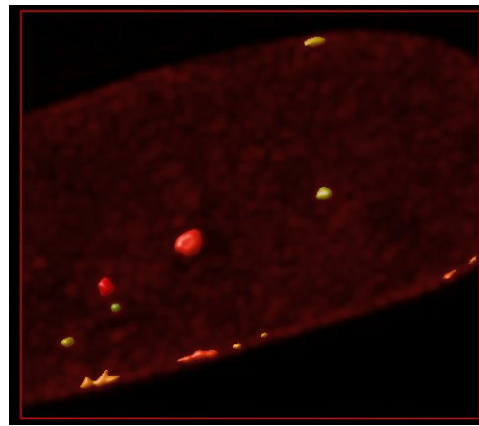



FIGURE 12.5. Objects and MIP. Objects in a surface pipe rendered together with a MIP pipe for spatial reference. The data channel is the same one in both pipes.

Object Statistics

The Object Analyzer can measure a lot of different parameters on the segmented objects, but only a few are reported by default. Otherwise the statistic table would be too saturated with information and you will not be able to find your way in it. To obtain object information you can click on each object interactively while you are in the *Analyze object* mouse mode, or press the ANALYZE ALL button  on the top-right of the table to automatically process all the segmented objects in all pipes. (When you have an area selected on the screen, enclosing a few objects, this button analyzes only these objects).

There are many parameters you can measure for each object, not only about its 'internal' information but also about its relationship with other objects and regions in the image. Let's see now how to retrieve all this information.

Configuring the Reported Parameters

To simplify the usability of the OA there is a big button next to *Select Statistics* on the top left of the table that reads EXPERIMENT PRESET at start-up. Click on it and you will get a preset selection dialog like the one in Figure 12.6.

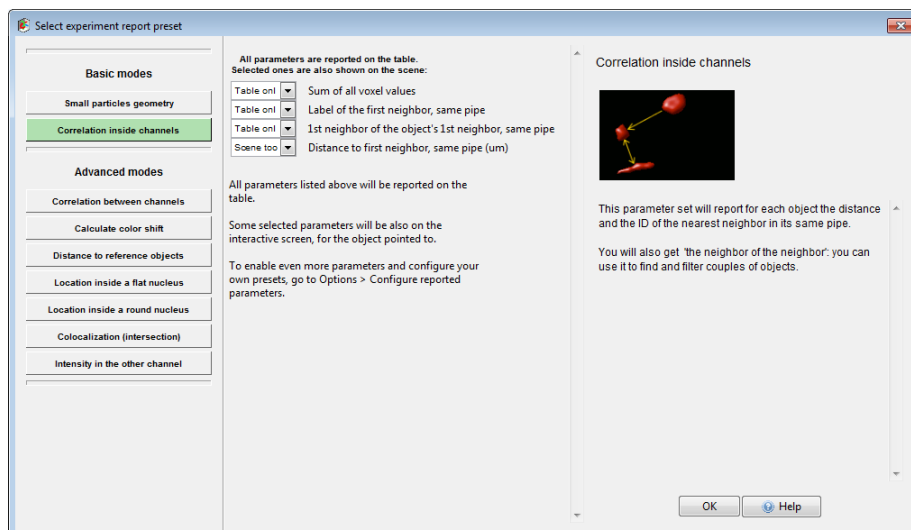


FIGURE 12.6. The Experiment Presets dialog. Collections of parameters can be selected to be reported on the table and interactively on the scene, depending on the experimental needs.

On its left column a series of different experimental needs are listed. When you click on each of them, a new list of parameters is listed in the middle column, and a description is shown on the right column. Even more, when you hover with your mouse over the listed parameters you get a tooltip text explaining each parameter with more detail. Please read the descriptions briefly to get an impression of them. Let's select now the preset called CORRELATION INSIDE CHANNELS and inspect it in detail. Please read the descriptive text.


"This parameter set will report for each object the distance and the ID of the nearest neighbor in its same pipe."

For each object we are interested in, we are going to find also the closest object.

Notice that from the listed parameters in this preset only one is selected by default to be reported on the scene. All of them will be calculated and reported on the table, but only the selected ones will be shown on the rendering canvas as well, when you interact with the objects. In this case it is smart to report only the distance there.

Once you have selected this CORRELATION INSIDE CHANNELS preset, press OK to continue.

Measuring the Objects

Make sure you have the *Analyze object*  mouse mode selected, click on a nice object on the screen, and notice the new columns that appear on the statistics table.

Before showing the actual object details, a row in the table informs you about the segmentation conditions for this pipe. Then comes the data itself. When you move your mouse over the column titles at the top of the table you get a description of each parameter at the very bottom of the window. **Distance to first neighbor, same pipe (um)** As many parameters are available, sometimes measuring things complex to describe, there is only room for a cryptic label in the column title. You can always find out what each parameter is by looking at its tooltip. The description of each parameter will be also stored in your file when you export the table later.

The tooltip of the column header also shows the sum and the average of all the cells below it, that you can optionally restrict to include data of a given pipe only.

Click on another object to add its data to the table. You will see that the conditions are not reported again, because they have not changed. A checkbox option at the top-right of the table clears it automatically when the segmentation or report conditions change, so you always have an organized table. Deselect this option to simply accumulate rows on the table, so you can save it later and process it with another program. Another button next to that checkbox clears the table manually whenever you want.

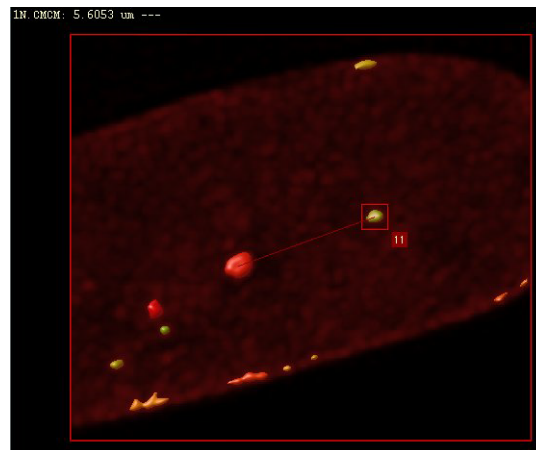


FIGURE 12.7. Interaction with the objects. When the distance to the nearest neighbor is computed, it is also displayed on the screen for the object aimed at.

You may also have noticed that something else happens when you interact with the rendered objects. In the example of Figure 12.7 the user clicked on object 11. A line joins the center of mass (CM) of this object with the center of mass of the nearest neighbor, and this distance is reported at the top of the window. (This is the only parameter reported on the canvas because that is the only one selected on the preset selection dialog, but all the interesting information is now included in the table too).

Other Measuring Parameters


There are two basic parameter presets: this one we have used here to explore the nearest neighbors, and another one to retrieve morphological parameters about objects, called SMALL PARTICLES GEOMETRY.

This other parameter preset can report object information like *length*, *width*, *aspect ratio*, and *sphericity*. Details on how these parameters are defined can be found on page 75.

The presets are organized in basic and advanced modes. The basic modes refer to presets on which you don't have to do much before starting measuring: you click on the objects and you get the desired information. The advanced modes require that you define some conditions first: to define a ROI, for example. They are not intrinsically more complex, but they demand from the user to know how to set these reference conditions first.

For more details about any parameter preset just click the Help button you can find at the lower-right part of the preset selection dialog, and follow the on-screen tooltips during the interaction with the module.

Exploring the Table

Let's try another thing: click on the ANALYZE ALL button  at the top-right of the table so the data of all objects are gathered. When the computations are finished (it should be quite fast in this example) move your mouse over the table rows and see what happens. The object corresponding to the current row will be highlighted on the canvas, and the distance to its nearest neighbor will be shown.

47	36	1	S	14	102.79	92.715	13.935	0	2636.93	2	1.7649	-8.1495	6.2257	-8.1499
48	37	1	S	7	144.58	163.71	14	0	1316.35	8	3.0245	27.968	-35.824	-2.0866
49	38	1	S	6	217.81	172.84	14	0	1130.53	8	4.2358	-45.366	-44.955	-2.0866
50	39	1	S	11	32.902	78.999	15	0	1983.05	9	1.272	15.083	-10.844	-1.7127
51	40	1	S	40	77.396	94.902	15.224	0	8238.29	2	2.2204	17.244	4.0366	-9.4394

Ready (rendered) Canvas size 678 x 577 Dragging on 38/S 1NP.CMCM: 4.2358 um

FIGURE 12.8. Exploring the table. Hovering the cursor over the table rows also highlights the corresponding object on the scene.

A good way to find an object in a very long table is by clicking on it on the rendering canvas while the *Analyze object* mode is active: the table will be shifted to show its corresponding row, and it will be highlighted.

By moving the mouse over the column titles you get an expanded description of the parameter plus some basic statistics of all the cells in that column. You can also right-click on a column and select STATS for more detailed statistics.

To plot a histogram of the distribution of values in a column, select the column (or a subset), click with the right mouse button, and select HISTOGRAM from the pop-up menu.

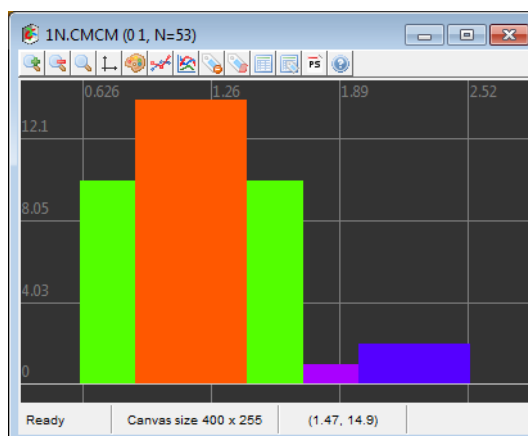


FIGURE 12.9. Histogram representation of one column values. After computing all objects sphericity parameters, for example, and having them reported on the table, we can plot a histogram of the particle sphericity distribution.

Storing your Results

In the top menu you can do FILE→SAVE OBJECT STATISTICS to export the table you see on the screen to a file that you can import elsewhere. The file can be read directly, or imported in conventional spreadsheet programs or data plotters. The table can be stored as a plain text file, a csv-file, that can be imported in e.g. Microsoft Excel, OpenOffice Calc or GNUplot, and as an m-file, that can be imported in Matlab, to do further analysis and/or calculations.

You can also save the current scene as you see it on the rendering canvas to a TIFF file on disk with FILE→SAVE SCENE. In the OPTIONS menu you can find different options that affect how the scene is rendered. You can set the MIP pipe to high quality, for example, or render each analyzed object together with its numeric ID label, so that you can link them with the exported data visually.

In the HISTORY menu, you can save your analysis history as a template, to reproduce it using the same or other dataset.

Further Reading

This tutorial has covered the very basic features to the Huygens Object Analyzer. If you want to learn much more, consider following the expert on-line tutorial³, where you will be told on many other powerful things you can do with this versatile analyzer.

A reference description of the Object Analyzer components can be found in Chapter 12 on page 63.

3. <http://support.svi.nl/wiki/ObjectAnalyzerExpertTutorial>

Object Analyzer Geometry Measurements

Iso-surface

In the Object Analyzer, the threshold set for the segmentation is also used to define an iso-surface around the object.

An iso-surface is a 3D surface representation of points with equal values in a 3D data distribution. Is the 3D equivalent of a 2D contour line (See Figure 9.3 on page 44).

Based on this, one can measure volume and surface in *high resolution*, by fine polygonization at a sub-pixel level.

Principal Axis

Segmented objects are geometrically analyzed in terms of their principal moments of inertia. (In this sense, the recorded light intensity registered in the image is used as density: the pixel 'values' are interpreted as local mass, so brighter regions weight more).

The principal axes of an object establish a natural system of reference based on its mass distribution. When you rotate an object around one of its principal axes, the angular momentum is parallel to it. This does not happen in general, and is what makes these axes so special.

Around these axes the principal moments of inertia of the object are defined. For one of these axes, the rotation inertia of the object is minimal (around this axis the object would rotate with the least effort). This axis usually lies along the length of the object. The other two axes are orthogonal to it, and orthogonal to each other.

Length and width

One can easily define a box, with dimensions L , $pBoxW0$, $pBoxW1$ in the system of reference of the principal axes, that encloses the object completely. The sides of this box are in general not parallel to the main planes of the image, because the principal axes do not coincide with the image x , y , or z -axes in general. It is as if the principal box is rotated with respect to the image in order to properly enclose the object, which may not be aligned with any of the image axes.

The length of the object is the largest distance measured along the three principal axes, it coincides with the largest dimension of the principal box L .

One could use the other two dimensions of the principal box as width and thickness of the object, but for some practical uses this may be too simplistic.

The width of the objects is actually computed with a search algorithm that acts as a virtual caliper held perpendicular to the length axis. To find the largest width of an object one would rotate the caliper around the object and repeat this procedure while sliding the caliper along the length axis. However, because microscopic data, even when it is deconvolved, often shows orientation dependent imaging due to the lower axial resolution, structures are often elongated in the axial direction. Moreover, in most cases, the voxels themselves are much higher than they are wide, causing all small objects to be elongated.

Clearly, without correction, the anisotropy in resolution would result in an overestimation of the width. To avoid this problem the rotation angles at which the caliper is held are divided in axial directions and lateral directions.

To measure in the axial directions several slices are taken out of the object and analyzed one by one. Each of these slices is parallel to the 'caliper plane', perpendicular to the length axis of the object. In any slice there are directions more oriented towards the optical axis (axial directions) than others. The largest axial width of each slice is obtained by holding the caliper in these directions and searching for the largest among them.

After all slices along the length axis have been examined, the largest axial width of the object is reported as *WiAx*.

The lateral directions in the caliper plane are the directions closer to the *xy*-plane. The caliper measures now the width of each slice in directions near-perpendicular to the optical axis. Taking the largest figure among these lateral directions might again introduce a bias due to elongation, so now two values are computed per slice: the largest and the smallest width along the lateral directions.

After exploring all slices some global relevant figures are reported. The largest lateral width obtained while sliding the caliper plane along the length axis is reported as *WiLat*. The largest of the smallest widths is reported as *WiLatC*. In case of small objects and moderate to high ratios between the axial and lateral resolution this last value is likely to suffer least from the orientation dependent imaging.

The waist (the smallest of the smallest widths) in the lateral directions is reported as *WaistLat*.

What about objects with a vertical length axis? In that case the 'caliper plane', perpendicular to the length axis, will be horizontal, parallel to the *xy*-plane. In a horizontal caliper plane all directions in it are perpendicular to the vertical *z*-axis. As a result there is no 'most axial direction' in that plane. The software then orients the lateral width towards the *x*-axis and the axial direction towards the *y*-axis.

Sphericity

The sphericity is reported in two ways in the Object Analyzer:

1. The axial sphericity is defined in general as the ratio of the volume of an ellipsoid with axes length *L*, width *W* and thickness *T* to the volume of a sphere circumscribed around the segmented object, defined by its length.

Because it is based on three axes, it gives an idea of the 3D aspect ratio of the objects.

Depending on what of the previously reported parameters we choose for W and T we have different practical definitions of axial sphericity.

Probably the most intuitive one is the axial sphericity of the principal box $axSphPB$: the ratio of the volume of an ellipsoid with axes L , $pBoxW0$, $pBoxW1$ to the volume of a sphere circumscribed around the principal box (see above) using the largest side (the length of the object L) as diameter (Equation 1).

$$AxSphPB = \left(\frac{pBoxW0 \cdot pBoxW1 \cdot L}{L^3} \right)^{1/3} \quad (\text{EQ 1})$$

2. Another definition (reported as $AxSph$) involves the lateral and axial widths discussed above, obtained with the virtual caliper algorithm (Equation 2).

$$AxSph = \left(\frac{L \cdot WiAx \cdot WiLatC}{L^3} \right)^{1/3} \quad (\text{EQ 2})$$

The roughness sphericity characterizes the roughness of the iso-surface, it measures how close the volume-to-surface ratio is to the one of an ideal sphere. This is conventionally defined as:

$$SurfSph = \frac{\pi^{1/3} \cdot (6V_i)^{2/3}}{A_i} \quad (\text{EQ 3})$$

where V_i is the iso-volume and A_i is the iso-surface of the segmented object. Both sphericity values become 1 for an ideal sphere. The $SurfSph$ is the inverse of the 'surface factor' fs used in Goetze et al.¹

Aspect Ratio

Again, the aspect ratio of an object can be defined in terms of different dimensions:

- The axial aspect ratio $AxRatio$ is the ratio of the object length to its axial width $WiAx$.
- Similarly, the lateral aspect ratio $LatRatio$ is the length divided by $WiLatC$.

More Parameters and Filtering

Many more parameters, geometrical and of other kinds, can be calculated, reported, and used to filter the data by the Object Analyzer. Please refer to the on-line tutorials to learn how to use the different parameters and Experiment presets.

Parameters are available that report:

- *Correlation* inside and between channels, by analyzing neighbor objects
- Location of objects with respect to *reference objects* (anchors).
- Location of objects inside *regions of interest* (for example bodies inside a cell nucleus)
- *Colocalization* of objects (by computing the volume and the intensity overlap of segmented objects in different pipes).

Any of the calculated parameters can be used to filter out objects and further segment your image in elaborated ways.

1. Goetze et al., *Molecular Cellular Biology* 27, p. 4475-4487 (2007)



This section describes the components of the Object Analyzer (OA) interface.

This section is intended as a quick reference. See “Introduction to the Object Analyzer” on page 63 to learn how to use the basic components in context.

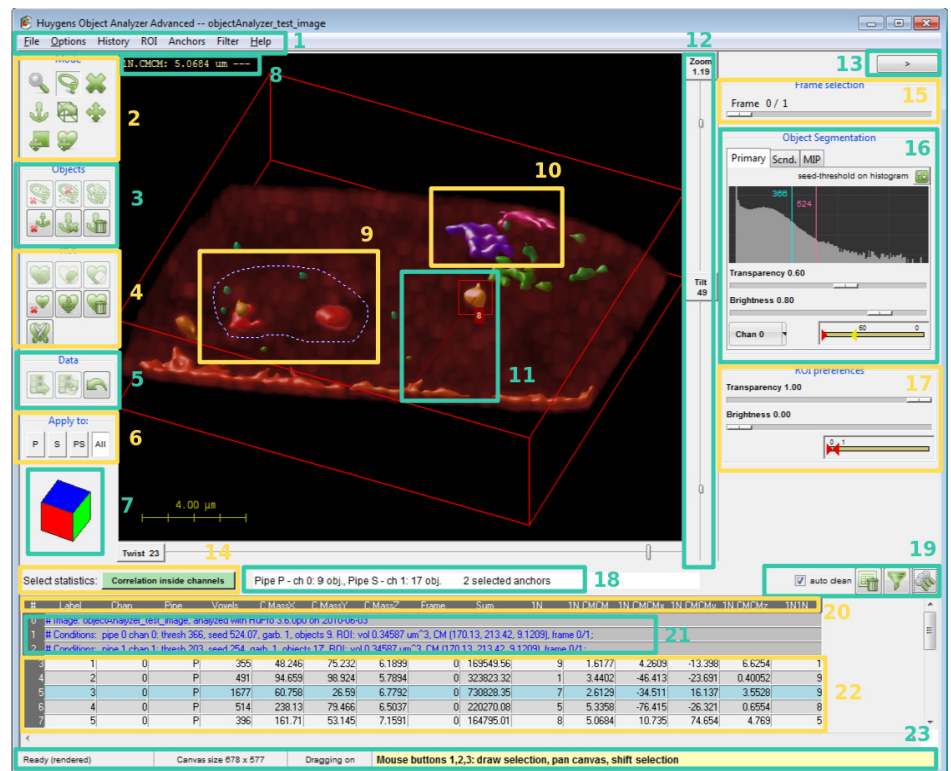


FIGURE 14.1. The Object Analyzer interface. Different regions are enumerated to describe them in sections of this chapter.

Main window components

The OA main window is shown in Figure 14.1 with most of its components in an active state. Different regions of the interface are enumerated to describe them conveniently along this chapter.

When the OA is launched, not all the buttons in the toolbox are enabled. Most buttons are automatically enabled when they are usable, depending on conditions set by the user.

Main Menu (1)

FILE: Entries to save the rendered image, export or clear the object statistics, and analyze all time frames in a series.

OPTIONS: Rendering options and advanced statistics configuration:

- Virtual render size: sets the size of the canvas on which the Ray Tracing algorithm renders the image. This canvas can be larger than the OA window or even the screen (that's why you can pan the canvas to inspect other regions). Like that you can render and save high-resolution TIFF images.
- Transparency depth: controls the number of surfaces considered by the renderer in order to show inner cavities and objects inside objects.
- Bounding box: shows or hides the reference 3D bounding box
- Scale bar: shows a scale bar on the scene. The distances are calculated based on the voxel sizes in the original image microscopic parameters.
- Show ID labels: render the scene showing the numerical ID label of each of the analyzed objects.
- Show SVI logo.
- High quality MIP: enables or disables the high quality rendering mode for the MIP pipe
- Show reference cube (See "Reference Cube (7)" on page 85).
- Show on-screen tooltips for interactive actions.
- Relaxed selection: when active, objects partially outside the selected 2D area are also considered.
- Center scene: moves the point of view to show the center of the dataset.
- Configure statistics report: shows all available parameters to let you select which ones are calculated and reported on the table. This is intended for advanced users, beginners should better use the *Experiment presets* (See "Experiment Presets (14)" on page 87).

HISTORY

- Undo the last operation, or Redo it again.
- Reload original data after cropping the image or discarding objects.
- View the whole operations history up to the current point.
- Load and save analysis templates, so that the current view and parameters can be stored and retrieved, or a whole operations history re-executed with other data. This is also useful to store your analysis steps and reproduce them.
- Set analyzer as in any other open instance of the tool.

ROI: operations to define a region of interest in complex ways, and to modify and save the currently defined ROI.

- Set:
 - Set to extruded selected area: uses the interactively defined 2D area to set the ROI to the volume below it.
 - Set using MIP threshold: use the threshold and data channel of the MIP pipe to set a 3D ROI
 - Make coincide with objects: use the currently segmented objects to define a 3D ROI. Objects from the Primary, Secondary or both pipes can be used depending on the active pipes and the pipe mode (See “Active Pipe Mode (6)” on page 84).
 - Make coincide with anchors: use the currently selected anchors to define a 3D ROI. This may leave holes in the interior of the ROI if the anchor is not a solid object.
 - Make coincide with intersection: this is interesting to do object analysis with colocalizing volumes only. When you have two pipes active and some objects colocalize (so that the intersection volume of the two pipes is not empty) a ROI can be defined with this operation, and applied it to discard objects and parts of objects outside it, so that only the colocalizing regions remain.
 - Envelop anchors: use the currently selected anchors to define a 3D ROI, so that also holes inside the objects are 'filled in' and in the ROI.
 - Enclose anchor in a box: define a prism that is an envelop to the selected anchors.
 - Make a spherical ROI of a given diameter, centered in the image. You can shift it later with the shift ROI mouse mode.
 - Set to all volume: maybe a good starting point for further interactive modifications.
- Modify:
 - Fill inner cavities: a ROI defined by using a threshold may contain inner cavities (visible when selecting the deep Transparency Depth). This operation fills them in automatically.
 - Fill inner and cutoff cavities: A cutoff cavity is a hole in the surface of a ROI that touches the image limits.
 - Grow/shrink: the currently defined ROI can be enlarged or reduced in a certain number of VoXels, independently in the xy -plane or in the z -direction. A 3D (xyz) reduction is also possible: here the number of voxels in the xy -plane will be partially adapted to the entered voxels along z to, considering the voxel anisotropy, grow/shrink the volume proportionally.
 - Outer shell: re-define the ROI considering only an outer shell of given thickness.
 - Invert the ROI volume.
- Storage
 - Save current ROI to file
 - Load ROI from file
 - Add ROI from file
 - Intersect with ROI from file
 - Subtract ROI from file
- Center ROI on the anchor CM: align the Center Of Mass (CM) of the currently defined ROI with the CM of the selected anchors.
- Clear ROI
- Keep only objects inside the ROI, discarding anything else. The relaxed selection mode in the options also affects how objects partially inside the ROI are handled.

- Analyze ROI volume computes and reports information on the table about the ROI itself.
- Analyze all objects inside the ROI reports in the table information about objects inside the ROI, or partially outside it, depending on the relaxed selection option.
- Help on Regions of Interest.

ANCHORS

- Select all objects as anchors.
- Deselect all anchors.
- Invert current anchor set.
- Set anchors by filtering. This opens a filter dialog as explained in “Table and Analysis Shortcuts (19)” on page 89, but allowing you to select or deselect anchors instead of discarding objects.
- Keep anchor objects, discarding anything else
- Discard all anchor objects.
- Analyze only objects select as anchors

FILTER: some useful predefined filters to remove objects based on their features, and access to a full control filter tool and to reload the original data.




- Quickly remove objects that are touching the borders of the image, as they are surely incomplete.
- Quickly remove objects based on size or sphericity
- Quickly remove non-colocalizing objects
- Quickly find pairs of objects, inside the same pipe or by combining the two pipes.
- Advanced filtering shows the same filter dialog explained in “Table and Analysis Shortcuts (19)” on page 89.






HELP: shows on-line help and tutorials

Mouse Modes (2)

These buttons control the Mouse mode, what the mouse does when you left-click on the image. When you hover with the mouse over these buttons you get a tip on what they are for. If they have a keyboard shortcut to activate them, it is also shown here. For example, the first mouse mode is intended to *analyze objects*, and you can always activate it from within the rendering view by pressing the keyboard key ‘1’.

From left to right, and from top to bottom:

-  **Analyze object** lets you click on different defined objects and obtain the local statistics.
-  **Select area** lets you define a 2D region on the current view of the image so you can do different things with it: analyze or discard objects below it, anchor them as references, or interactively define regions of interest (ROI). We will see what all this is useful for.
-  **Discard object** lets you discard irrelevant objects one by one. Just select this mode and click on the disturbing objects.




-  **Select object** (as anchor) lets you select and deselect objects to be 'anchors', for example to act as references to measure distances from other objects when asking for local statistics. When you set an object as a reference anchor it will 'light up' and change color on the screen to indicate its new status. It is possible to select a group of anchor objects, and you can operate with them through the Anchors menu.
-  **Rotate scene** interacts with the full image to rotate it in the space, by dragging the mouse pointer on the rendering view. That can also be achieved by moving the Tilt and Twist sliders along the rendering.
-  **Pan scene** interacts with the full image to move it in space laterally. This means that you can pan the scene in the 2D plane of your screen, not along the third dimension, along your line of sight.
-  **Pan canvas** is similar to 'pan scene' but not exactly the same. It allows you to explore the canvas by not re-rendering the scene. This only makes sense when you have a canvas larger than your rendering window, of course. See `OPTIONS→VIRTUAL RENDER SIZE` in the top menu.
-  **Shift the ROI.** This is the only mouse mode that is not always enabled: you need to have defined a region of interest (ROI) before moving it around.

For the advanced users: some of these mouse modes have 'shortcuts' in other mouse modes. Most of the times you can 'pan scene' independently of the selected mouse mode if you use your mouse right button instead of the left one. Similarly, you can 'pan canvas' using the middle button at any time. You will learn this with practice if you need it, but you don't have to care about it right now. When a mouse mode is active and the cursor is inside the canvas, the tooltip in the status bar (See Figure 14.1, item 23) tells you what each mouse button can do.

Selection Interactive Operations (3)

The rest of the buttons in the toolbar are not mouse modes: they do not set new behaviors for your mouse when interacting with the objects view but execute operations on previously defined conditions, for example deleting the objects inside a defined ROI. They are all disabled until these conditions have been set (in the example, until you define a ROI to operate with).



After having drawn a 2D selection (See “The Selected Area (9)” on page 85) in the *Select area* mouse mode, you can click on:

-  Keep all objects under the selected area, discarding anything else.
-  Discard all objects under the selected area
-  Set as anchors all objects under the selected area

The way the drawn 2D area considers objects in its limit can be controlled with an option in the top Options menu: the area can consider only objects fully under the selection, or also objects partially outside it (relaxed selection mode). In any case, the objects not affected by the selection are rendered with dimmer intensity to clearly indicate what objects are selected.







When you have selected certain objects as anchors, you can further operate with them:

-  Deselect all anchors.





-  Keep the selected anchors, discarding anything else.
-  Discard the selected anchors.

More operations are available at the top ANCHORS menu. The main purpose of an anchor is to serve as analysis reference in some experiment presets, for example to measure distances relative to them. But you can also use the selected anchors to delete some objects, or to define a 3D ROI based on them.

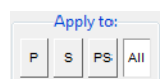
ROI Interactive Operations (4)

-  Reduce the current ROI to the volume under the selected 2D area (intersects the ROI with the selection)
-  Add the volume under the selection to the ROI volume (union of the ROI with the selection).
-  Remove the volume under the selection from the defined ROI (difference of the ROI with the selection).
-  Set all objects inside the ROI as anchors. This requires a ROI to be defined, of course.
-  Keep objects inside the ROI, discarding anything else.
-  Discard the defined ROI, reverting to the whole image.

Data Cropping and Restore Operations (5)

-  Once a ROI is defined, you can use these scissors to crop the dataset and delete (set to zero) all voxels outside the ROI.
-  Export the current data (as cropped by the ROI, or after deleting objects interactively) to the original image in order to save it.
-  If deleted data was not yet exported to overwrite the original, you can always reload the original dataset with this button.
-  Undo last operation.

Active Pipe Mode (6)



These radio buttons control, for most of the interactive, ROI and selection operations, on which of the active data pipes (*Primary*, *Secondary* or both) and channels (*Selected pipe*, or *All*) they act. Pipes are explained in “Render Pipes (16-17)” on page 88.

Examples of operations controlled by these buttons are:

- Interactively clicking on objects. When you click on the scene, only objects in the active pipe are considered.
- Analyze all / Analyze selection
- Set ROI to the visible objects
- Sum all column values (when hovering over a column header in the table)
- Apply the ROI to crop the dataset

There are only a couple of practical differences between the PS and All modes:

When cropping data by applying a ROI, you can crop the channel in the Primary pipe (P), the channel on the Secondary pipe (S), on both (PS), or all channels in the image (All), even if they are not shown in any active pipe.

When summing cell values in a column of the table, you can include in the summation cells for the channel in the Primary pipe (P), the channel on the Secondary pipe (S), both (PS), or all cells (All), including those reporting about the ROI.

For all the other operations the PS and All modes are equivalent.

Reference Cube (7)



At the end of the buttons column there is a colored reference cube that will help you in orienting in space when you rotate the dataset, specially with large zoom factors that do not let you see the surrounding box frame in the rendered image. The initial view of this cube is the blue top face corresponding to $z = 1$. Hovering the mouse over the cube faces brings a tooltip with the face label (x , y and z with values 0 or 1).

The Interactive Rendering Canvas (8-11)


The canvas shows the scene, the result of the ray tracing¹ algorithm rendering the segmented objects.

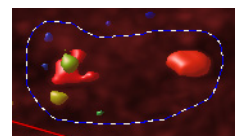
The scene is determined by the objects orientation (*tilt* and *twist*), the *zoom*, the *brightness* of the pipes, what point is centered on the view and so on. All that is taken into account by the renderer, that generates the scene, and puts it in the canvas on the screen for you to see it. Notice that depending on the render size (OPTIONS→VIRTUAL RENDER SIZE) the canvas can be larger than your screen.

The On-Screen Reported Parameters and Tooltips (8)

The currently selected Experiment preset (See “Experiment Presets (14)” on page 87) selects the statistics to report many parameters to the table (See Figure 14.1, item 20, 21, and 22) A few of these parameters can be also reported on the screen for the current object for easy reading. The magnitude of the distance that is plotted on the screen is followed by a triple dash ‘---’

The Selected Area (9)

The *Select area*  mouse mode lets you define a region in the image so you can do different things with it: analyze objects ‘inside’ it, anchor them as references, or use the selected volume to define a region of interest (ROI).



Therefore the *selection* is not yet a 3D ROI, but simply a 2D area. That is why ‘inside’ is quoted in the previous paragraph: objects are inside the selection only from the current point of view, so it would more appropriate to say ‘below’ it. It allows quick and simple interaction with the objects, specially on flat images here the objects remain more or less

1. <http://support.svi.nl/wiki/RayTracing>

in a plane, but only by defining a 3D ROI you have full control on what objects are removed when you want to discard some of them, or when you want to crop the dataset.

In the Select area mouse mode you can use the right mouse button to shift the defined selection around and reuse it multiple times in different locations.

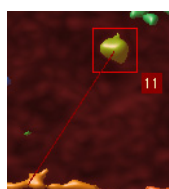
Anchor Objects (10)



Objects that act as references (anchors) are shown in the rendering with magenta or violet colors, depending on the pipe they belong to, so that they are clearly distinct from other regular objects.

There are different ways to set objects as anchors, interactively (See “Selection Interactive Operations (3)” on page 83) or by using filtering operations (top ANCHORS menu). See the on-line article about anchor objects for more details on their utility.

Aim (11)



Objects under your cursor are shown highlighted. Objects that have been already analyzed are also shown enclosed by a box when the mouse moves over them, or over the correspondent table row. To analyze an object just click on it while in the *Analyze object* mouse mode, or press the ANALYZE ALL button (See Figure 14.1, item 19).

Objects in the Primary pipe are framed with a red box, and objects in the Secondary pipe with a green one.

A small label showing the object's number ID is also shown when pointing at it. The background color also indicates if it belongs to the primary (red) or the secondary (green) pipe. Yellow labels are shown whenever the pointer has two objects below it, from different pipes.

If a distance is configured to be reported on-screen (See Figure 14.1, item 8), it is also plotted when pointing to an object.

If you hover over an object's row in the table and point to a cell containing a distance parameter, it will also be plotted on the rendering canvas. Like this you can interactively explore many reported distances.

Notice that, depending on the active pipe mode (See “Active Pipe Mode (6)” on page 84), the interaction with the scene may highlight and affect objects in one pipe only.

Scene Control Sliders (12)



Three sliders run along the canvas (vertically, on its right, and horizontally, on its bottom) to control the point of view of the scene:

- **Zoom:** the number is just indicative, 1 meaning that the whole dataset is shown in the canvas.
- **Tilt:** the angle of rotation (in degrees) around the canvas x -axis.
- **Twist:** the angle of rotation (in degrees) around the image z -axis.

By clicking on the labels you can enter numerical values manually to quickly switch to the desired scene.

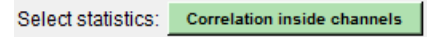
Hide Pane Button (13)

This button at the top right of the window collapses the control pane to make more room for the rendered scene. Once you have defined the segmentation conditions for all pipes, you don't need to interact with those controls anymore, but with the objects, so you can hide the pane to focus on the scene.



Experiment Presets (14)

This button opens a preset selection dialog that allows to select an experiment preset, a collection of parameters that make sense to be reported together in the context of certain experimental needs (See Figure 12.6 on page 70).



Users are very much welcome to send their own suggestions to implement new presets. We will gather all the feedback in different wiki articles that will explain what parameters are best for certain experiments and how to interpret them. Please feel free to tell us about your experiences!

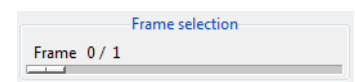
In this dialog you can see three columns. On its left column a series of different experimental needs are listed. When you click on each of them, a different list of parameters is listed in the middle column, and a description is shown on the right column. Even more, when you hover with your mouse over the listed parameters you get a tooltip text explaining each parameter with more detail.

By selecting a particular preset all the listed parameters (apart from the basic ones) will be reported and calculated. Moreover, all the listed parameters will also be available for filtering the objects (See "Filtering Objects" on page 89). The check boxes allow a few of these parameters to be also reported on the screen (See Figure 14.1, item 8) for the current object. The magnitude of the distance that is plotted on the screen is followed by a triple dash '---'. Only one distance can be plotted at a time by clicking on an object, but many can be reported on the table and explored interactively there by simply moving the cursor over the cells.

Only one preset can be selected at a time, and all its parameters will be reported. Advanced and more flexible configuration of the parameters is always possible through `OPTIONS`→`CONFIGURE STATISTICS REPORT` in the menu. In this configuration dialog users can also store any set of reported parameters as a new preset.

Time frame selector (15)

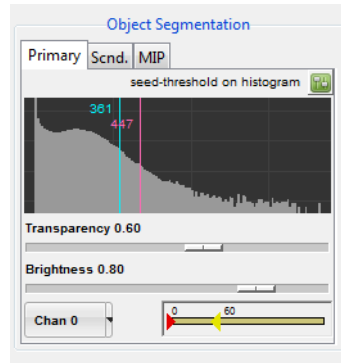
When time series are loaded in the OA, this slider controls which time frame is take for visualization and analysis.



Changes to the current frame (like discarding objects or selecting anchors) are remembered when you change the time frame. But voxel editions (using the ROI to crop the dataset) are lost unless you export them first with *Export the applied ROI to original* (📁).

You can analyze all frames in a time series, accumulating the data in the table, through the `FILE`→`ANALYZE TIME SERIES` menu entry.

Render Pipes (16-17)



The Huygens Object Analyzer has two pipes (named *Primary* and *Secondary* pipes) for you to put image data through for object analysis and another pipe to simultaneously visualize a MIP of one of the channels (See Figure 14.1, item 17).

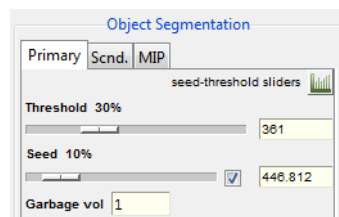
'Secondary' does not mean here 'less important', it is just that you can use the first pipe (*Primary*) and the second pipe (*Secondary*), abbreviated *P* and *S* in some places. We keep numbers to refer to image channels, which is something different, because in each of these pipes one can put any image channel, in any order.

One can even segment the same channel twice with


different conditions in each of the pipes if necessary!

This pane controls what data channel goes through each pipe, how its objects are segmented (with a seed and a threshold²), and how the data is rendered (transparency and brightness). A garbage volume in voxels can also be set to get rid of little spurious objects, with a number of voxels lower than the garbage level.

You can see the seed as a secondary threshold. The first threshold segments the data and makes independent objects, but then only objects with intensity that goes above the seed level remain, while the rest are discarded.



At start-up, the threshold and seed levels are represented by vertical blue and magenta lines, respectively, on top of a histogram of the channel in the pipe. An alternative representation of these levels as sliders is available, that also show their values as a percentage of the channel maximum and allows you to link the seed with the threshold. The garbage volume control is also in the slider pane, not in the histogram.

To switch between the histogram and the sliders control panes you have to click on the small button () at the right of the pane title.

The threshold ranges between the minimum and the maximum values in the channel intensities, considering all the time frames. Its percentage representation refers to the maximum value.

By default, the seed is linked to the threshold value, so when you vary the latter the seed absolute value also changes, in such a way that its relative value remains constant. In its linked mode, the seed is set referred to the span between the actual threshold and the maximum, and ranges from the threshold value itself (0%) to the image maximum (100%). This is because the seed, being an 'upper threshold', can never be lower than the threshold. It is also useful to be able to set the seed to 0% so that it is not used at all and you retrieve the classical threshold-only segmentation. Still you may find convenient to express the seed relative value in the same terms you use for the threshold (relative to the image intensity range), or to be able to set the seed to a fixed value independent of the threshold. For that, deselect the seed checkbox so that 0% also represents the same value as the 0% threshold (the image minimum, or zero), and the two sliders are unlinked. The

2. <http://support.svi.nl/wiki/SeedAndThreshold>

linking mode of the seed does not really affect the segmentation, it is just a matter of convenient representation of relative values: what is applied to the image as threshold and seed are always the absolute numeric values shown on the entry widgets and next to the histogram lines, that you can edit directly by clicking on them.

All the different objects in a pipe are colored differently to be able to distinguish them. The range of colors assigned to each pipe can be controlled with a *Hue Selector* (See “Hue Selector” on page 125). You can collapse this range completely if you want so all objects in a pipe get the same color.

General Object Information (18)

This little bar reports:

Pipe P - ch 0: 9 obj., Pipe S - ch 1: 17 obj. 2 selected anchors

- left: the number of objects currently segmented in each surface pipe.
- right: the total number of objects selected as anchors.

Table and Analysis Shortcuts (19)

These widgets give quick access to some table operations.



- Auto clean checkbox: this option makes the table to be cleaned whenever the segmentation or analysis conditions change, to always have a fresh start. Deselect this option to keep all data and keep accumulating rows in the table.
- Clear statistics table: manually delete the table contents
- Filter: opens a dialog that allows you to filter objects out based on the reported parameters (see below).
- Analyze all / Analyze sel.: This button runs the analysis procedure on all the currently segmented objects, for objects on the pipes selected by the Pipe mode radio-buttons (See Figure 14.1, item 6). When a 2D area has been selected (See Figure 14.1, item 9) this buttons analyzes only the objects under the selection.

Filtering Objects

The Filter button opens a pop up dialog that allows you to discard objects:

- Based on any of the reported parameters...
- Using a certain arithmetic operator...
- To compare the parameter with either a fixed value or with another reported parameter

The pipes that are filtered can be controlled with the active pipe radio-buttons.

You have to select one of the two options (fixed value or another parameter) and enter the value or select the parameter you want to compare with.

Because the filter is based on the reported parameters, you may need to configure the reported parameters or select other statistics first in order to filter based on the desired property.

For interesting usages of this filter, see for example the neighbors article in the SVI-wiki³.

For time series, another button shows to allow filtering all time frames.

Statistics Table (20-22)

The statistics table is the place where all the objects parameters are reported after the analysis. You can explore the table values in interesting ways directly on the Object Analyzer (see below), or export the contents for further analysis in other program.

When you move your mouse over the table rows, the objects they refer to are highlighted on the rendering canvas. When you point to a cell reporting a distance, this distance is also plotted on the screen.

#	Label	Chan	Pipe	Vowels	C.MassX	C.MassY	C.MassZ	Frame	Sum	1N	1N.CMCM	1N.CMCMx	1N.CMCMy
0	# Image: objectAnalyzer_test_image, analyzed with HuPro 3.6.0p0 on 2010-06-03												
1	# Conditions: pipe 0 chan 0, thresh 366, seed 524.07, gaib. 1, objects 9, ROI: vol 0.34587 um ³ , CM (170.13, 213.42, 9.1209), frame 0/1:												
2	# Conditions: pipe 1 chan 1, thresh 203, seed 254, gaib. 1, objects 17, ROI: vol 0.34587 um ³ , CM (170.13, 213.42, 9.1209), frame 0/1:												
3	1	0	P	355	48.246	75.232	6.1899	0	169549.56	9	1.6177	4.2609	-13.398
4	2	0	P	491	94.659	98.924	5.7894	0	323823.32	1	3.4402	46.413	-23.691
5	3	0	P	1677	60.758	26.59	6.7792	0	730828.35	7	2.6129	-34.511	16.137
6	4	0	P	514	238.13	79.466	6.5037	0	220270.08	5	5.3358	-76.415	-26.321
7	5	0	P	396	161.71	53.145	7.1591	0	164795.01	8	5.0684	10.735	74.654

The contents of the table can be copied to your clipboard, or stored to a file in disk by using the FILE menu.

The Table Columns and Their Headers (20)

Distance to first neighbor, same pipe (um) When you move your mouse over the column titles at the top of the table you get a description of each parameter, at the very bottom of the window (See Figure 14.1, item 20). There are many parameters that measure complex things, so in the column title there is only room for a cryptic label. You can always find out what each parameter is by looking at its tooltip. The description of each parameter will be also stored in your file when you export the table later.

This brief parameter description, plus longer explanations, are also given in the selection and configuration dialogs that selects or configures the reported parameters.

For columns with numeric values, basic descriptive statistics of all values in the column are also reported along with the parameter description. You can select whether the sum runs for both pipes (ALL or PS) or only for one of the pipes (P or S) depending on the selected pipe mode (See Figure 14.1, item 6). This provides a quick way of finding the total number of voxels in the object list, or those that are colocalizing, for example. Many interesting questions can be answered by this summations and ratios between them.

You can right click on a column to pop-up a contextual menu, from which you can retrieve more detailed statistics (STATS) for that column. The descriptive statistics will analyze pipes separately and together, and also include ROI information if present. The reported values are the maximum (MAX), the median (MED), the minimum (MIN), the number of items considered (N), the summation of the values (SUM), the average value

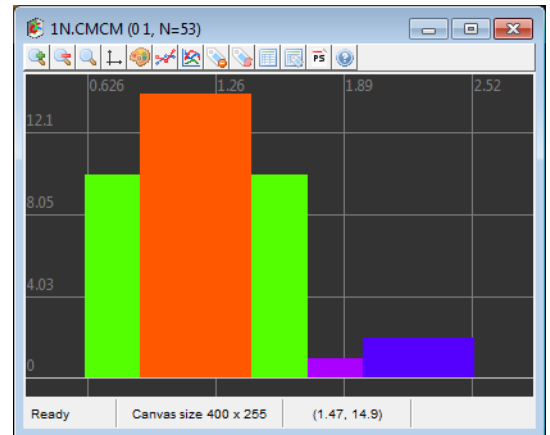
3. <http://support.svi.nl/wiki/ObjectAnalyzerNeighbors>

(AVG) and the standard deviation for the N items (SD- N), reported also as a percentage relative to the average value.

Clicking on a column title selects the whole column for you to copy. Multiple columns can be selected by holding the `Ctrl` key. To plot a histogram of the distribution of values in a column, select the column (or a subset) and select Histogram from the pop-up menu.

Selected columns also act as 'special parameter selector'. The value of the parameter for selected columns will be reported on the status bar (See Figure 14.1, item 23) when you explore cell values for given objects.

Like this, you can quickly compare different parameter values for the same object, by looking at the report in the status bar while you move your mouse over the table cells.



A similar thing can be also done by selecting rows, see below.

When multiple columns are selected, the STATS popup dialog will report descriptive statistics for all them.

Conditions Report (21)

When new segmentation or analysis conditions are set, they are reported when you analyze one or all objects. One row is added to describe the image and the time of the analysis, and another row per active pipe is added that reports the image channel in that pipe, the segmentation conditions, the number of objects in that pipe, and the volume and geometrical center of mass of the ROI:

```
# Conditions: pipe 0 chan 0: thresh 504.8, seed 540.52,
garb. 1, objects 12. ROI: vol 1103 um^3, CM (130, 115,
10)
```

The parameters about the ROI (volume and CM) are the same for both pipes, in case two are active. This is because there is only one ROI for all pipes and channels, and the Center Of Mass (CM) here reported is calculated considering the ROI as a uniform, solid object.

There is another way to calculate the CM of the ROI, not considering it homogeneous but taking into account the real image intensities in envelopes. These intensities, being different per pipe, yield to different ROI's CM per pipe. This is not reported in the conditions rows, but as object rows in the table when you select Analyze all. In this case, the ROI itself is treated as a new object in each pipe and more detailed information is reported in separate rows in the table. See the wiki article about ROI⁴.

4. <http://support.svi.nl/wiki/ObjectAnalyzerROI>

The Table Rows, How to Explore Them (22)

Each row in the table is either a *Conditions report* (See “Conditions Report (21)” on page 91) or information about one of the segmented objects.

In the second case, the object the row refers to is identified by some mandatory parameters (parameters that are always reported):

- *Label*: an integer index that is unique, per object, inside its pipe.
- *Chan*: the image channel this object belongs to
- *Surf*: the surface pipe in the current analysis the object belongs to (*P* is *Primary*, *S* is *Secondary*).
- *Voxels*: the number of voxels in that object
- *C.Mass*: three columns (*x*, *y*, *z*) for the center of mass location of the object, in the image coordinate system, with voxels as units.

The second letter in the *Surf.* column informs whether the object you clicked was in the front (*F*) or the back (*B*) position. This is only relevant in the following situation: it may happen that two objects in different pipes overlap in space (or they apparently do from the current point of view) and when you click somewhere on the screen you are actually selecting both of them, if the pipe interaction mode is PS or ALL (See “Active Pipe Mode (6)” on page 84). In that case, two rows are added to the table, and this second letter lets you know which of the two was in front of the other, from the current point of view.

By moving the mouse over the table rows the corresponding object is highlighted on the canvas (See Figure 14.1, item 11). Objects in the Primary pipe are highlighted in red, and objects in the Secondary pipe in green.

A good way to find an object in a very long table is by clicking on it on the rendering canvas while the *Analyze object* mode is active: the table will be shifted to show its corresponding row, and it will be highlighted.

When you move over the table cells the current value is shown on the tooltip bar at the bottom (See Figure 14.1, item 23). This, combined with the selection of rows (see below) or columns (See Figure 14.1, item 20), allows a quick exploration of the table and the comparison of different parameter values.

In the following example, when the cursor is moved over one cell that contains the distance to the first neighbor, the following is reported in the tooltip: 38 / S 1NP . CMCM : 4 . 2358 μ m. The first part is the label of the object in the current row: label 38 on the Secondary pipe. Then the parameter tag 1NP.CMCM refers to the CM-to-CM distance between this object and the nearest object in the other pipe. Then comes the distance itself: 4.23 microns.

47	36	1	S	14	102.79	92.715	13.935	0	2636.93	2	1.7649	-8.1495	6.2257	-8.1499
48	37	1	S	7	144.58	163.71	14	0	1316.35	8	3.0245	27.868	-35.824	-2.0866
49	38	1	S	6	217.81	172.84	14	0	1130.53	8	4.2358	-45.366	-44.955	-2.0866
50	39	1	S	11	32.902	78.999	15	0	1983.05	9	1.273	15.083	-10.844	-1.7127
51	40	1	S	40	77.396	94.902	15.224	0	8238.29	2	2.2204	17.244	4.0386	-9.4394

Ready (rendered) Canvas size 678 x 577 Dragging on 38/S 1NP.CMCM: 4.2358 μ m

That tooltip region can show not only the information of the cell pointed by the mouse, which would not be really useful, but also other values that can be set as reference. Try this: while keeping the `Ctrl` key pressed on the keyboard, select a couple of rows by clicking on the row number at the very left of the table. The selected rows will turn

green. If you now hover the mouse over a cell on any other row, you will get the value not only of that cell, but also those in the corresponding cells of the selected objects (rows). This is a quick way to compare results for different objects that can be distant in the table.

44	33	1	S	9	139.23	130.45	12.681	0	1610.24	8	2.2042	33.217	-2.5615	-0.76713
45	34	1	S	26	59.195	72.512	14.466	0	5065.34	9	0.82804	-11.211	-4.3561	-1.1784
46	35	1	S	40	60.92	87.086	14.049	0	7706.94	9	1.5209	-12.936	-18.93	-0.76123
47	36	1	S	14	102.79	92.715	13.935	0	2636.93	2	1.7649	-8.1495	6.2257	-8.1499
48	37	1	S	7	144.58	163.71	14	0	1316.35	8	3.0245	27.868	-35.824	-2.0866
49	38	1	S	6	217.81	172.84	14	0	1130.53	8	4.2358	-45.366	-44.955	-2.0866
50	39	1	S	11	32.902	78.999	15	0	1983.05	9	1.273	15.083	-10.844	-1.7127
51	40	1	S	40	77.396	94.902	15.224	0	8238.29	2	2.2204	17.244	4.0386	-9.4394

Ready (rendered) Canvas size 678 x 577 Dragging on **38/S 1NP.CMCM: 4.2358 um - (34 S): 0.82804 - (36 S): 1.7649**

If you select columns instead of rows in the table the tooltip will display the corresponding parameter values for the same object. Do not highlight columns and rows at the same time by now: it may be too confusing to interpret!

The Status Bar and Tooltip (23)

The bottom part of the window is a status bar that also shows an contextual tooltip.

Ready (rendered)	Canvas size 678 x 577	Dragging on	Mouse buttons 1,2,3: draw selection, pan canvas, shift selection
------------------	-----------------------	-------------	---

The left side reports the current status of the renderer and analyzer. You can see whether a long computation is running or if the analyzer is ready for further interaction.

Then the current size of the canvas is reported. When you first start the OA, the canvas size is adapted to the exact room left for the scene rendering, but it can be larger or smaller at wish (OPTIONS→VIRTUAL RENDER SIZE).

The 'Dragging' status refers to the automatic rendering of the scene while the user interacts with it. In very slow systems, the dragging is turned off automatically and the rendering only happens after the user released the mouse buttons or finished moving the segmentation sliders.

The tooltip (the region with light yellow background) shows contextual information:

- A longer description of the reported parameters when you hover over the table headers (See Figure 14.1, item 20) and a sum of the cell below them.
- The value in the cell under the cursor when you point at table rows, plus selected reference values.
- The different actions bound to the mouse buttons, when the cursor is inside the rendering canvas.

The *Colocalization Analyzer* provides information about the amount of spatial overlap between structures in different data channels, for 3D images and time series.

As this overlapping can be defined in many ways, Huygens gives the *colocalization coefficients* most commonly used in literature: (*Object*) *Pearson*, *Spearman*, *Overlap*, *Intersection*, and *Manders M* and *K*. More information on these coefficients can be found in the SVI support Wiki^{1,2}.

The *Colocalization Analyzer* is an extended optional tool, and is enabled by a C flag in the license string (See “License String Details” on page 121).

Notice that the *Object Analyzer* (See page 63) also provides colocalization measurements at the object level. The *Colocalization Analyzer* works more at the level of the whole image, despite local statistics of the colocalizing regions can be easily retrieved. Both analyzers work, in a sense, in complementary ways.

The *Object Analyzer* allows to define objects (segmentation) and observe how much they overlap, in volume or intensity. Objects defined like this can overlap with other objects, or not.

The *Colocalization Analyzer* explores the whole image to search for colocalizing regions based on the usual colocalization coefficients. These regions are then segmented and treated as objects to analyze. These objects are therefore always volumes of intersection.

How to use the Colocalization Analyzer

To *start* the *Colocalization Analyzer* in Huygens Professional, right-click on an image's thumbnail to open the contextual menu, then select COLOCALIZATION ANALYZER. Alternatively, an image's thumbnail can be selected, then in the menu bar select VISUALIZATION → COLOCALIZATION ANALYZER. The image must be multi channel (See “Multi-channel Images” on page 16) as the colocalization is based on the overlapping of different channel intensities. Wait for the analyzer to initialize and to compute the first 2D histogram with the default settings.

-
1. <http://support.svi.nl/wiki/ColocalizationTheory>
 2. <http://support.svi.nl/wiki/ColocalizationCoefficients>

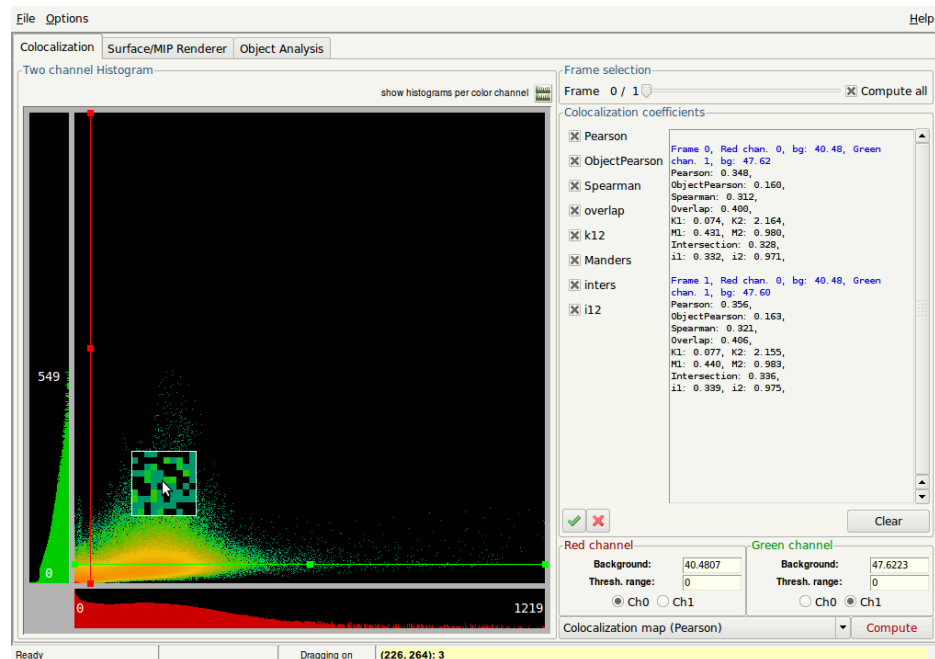


FIGURE 15.1. The first tab of the Colocalization Analyzer showing the colocalization coefficients and histograms.

When the colocalization analyzer is opened you will see three tabs (See Figure 15.1), of which the first is visible. In this tab the colocalization coefficients can be calculated. You can choose the timeframe, select colocalization coefficients, set the background settings and choose a colocalization map. In default start-up, all these settings are set for you, such that you can immediately press COMPUTE.

First we select the data to analyze. For time series, the *Frame* slider selects the time coordinate.

We follow the usual naming convention in colocalization theory for the two compared channels: *Red* (R) for the first channel, *Green* (G) for the second channel. We can select, in the lower part of the window, which data channels from our image are the Red and Green channels to be compared.

A two-channel *histogram* is calculated by default, and updated whenever we change the Red or Green channel selection. This histogram is already an indication of the degree of overlapping between the selected channels: for two channels with a high degree of overlapping, the histogram pixels trend to concentrate along the diagonal $y = x$ line. In contrast, total absence of overlapping would produce a 2D histogram with values only on the coordinate axes¹.

Along the x and y line of the 2D histogram, there are 1D histograms for each color channel separately. The enlarged versions of these 1D histograms are shown, when clicking on the SHOW HISTOGRAM PER COLOR CHANNEL button at the top right of the histogram. By moving the colored *background* lines in either the 2D-histogram, the enlarged 1D-histograms, or by changing the numeric values in the input fields, you can specify the backgrounds for the two selected channels. These values are subtracted from the voxels intensities when calculating the coefficients (if the result is negative, it is understood as a zero). Generally the colocalization coefficients depend much on correct estimation of

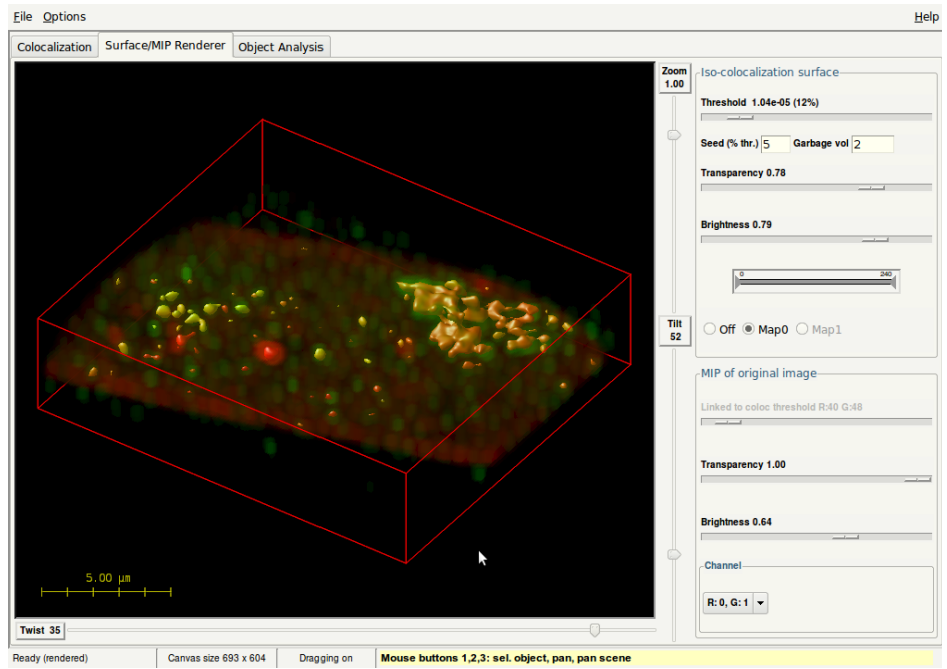


FIGURE 15.2. Tab two of the Colocalization Analyzer shows surfaces for regions of overlapping channels.

the image background and resolution. For these reasons we strongly recommend to compute colocalization coefficients only on deconvolved images.

One sets a background value prior to the calculation to remove signal that would otherwise lead to spurious colocalization, . This is intended for minor tuning, or just in case it is really needed to calculate colocalization in raw images that still have all the measured background. See “Backgrounds vs. thresholds in colocalization” on page 98.

Next we select what *colocalization map* we want to calculate: (*Object*) *Pearson*, *Spearman*, *Overlap*, *Manders M* or *K* or *Intersection*.

Notice the difference between *maps* and *coefficients*: the *colocalization coefficients* parametrize the degree of colocalization of the full image, while a *colocalization map* parametrizes the colocalization locally. In a map, a single colocalization value is calculated per voxel creating a 3D distribution that is represented in a 3D image by iso-surfaces. In other words, the colocalization map shows the contribution of each voxel to the colocalization coefficient.

The Colocalization Analyzer computes only the map selected by the user, but it always computes the selected coefficients. When the COMPUTE button is pressed, the pane *colocalization coefficients* will show the selected colocalization coefficients for the selected time frame. If you have checked the COMPUTE ALL option next to the timeframe slider, the coefficients will be calculated for all time frames sequentially.

Your choice of the colocalization map is needed for the functionality in the second tab (See Figure 15.2). In this tab an iso-colocalization surface is generated based on the colocalization map. This iso-surface rendering is only possible if a colocalization map exists. So, at startup the iso-colocalization surface sliders are deactivated, as we have to calculate a colocalization map first, by pressing COMPUTE.

Together with a surface renderer a MIP renderer is available. The viewpoint of the MIP and surface can be selected by moving the Tilt and Twist slide. Also try changing the zoom.

The obtained colocalization map is represented in the renderer window by iso-colocalization surfaces. These iso-surfaces represents points which all have the same colocalization value, thus regions in which the degree of colocalization exceeds a certain value become objects. This “certain value” can be controlled by the *threshold* slider in the iso-colocalization surface parameters. The transparency and the brightness of this surface pipe can be controlled with the correspondent sliders. The color range in which the objects are displayed can be modified using a hue selector (page 125). One can also switch the surface pipe off.

Some modes generate two-channel *colocalization maps*: colocalization of red with respect to green, and vice versa, e.g. in case of the *Manders M1* and *M2* coefficients. In these cases, the iso-colocalization surface parameters will offer the possibility of rendering any of the two channels, and thus the threshold is referred to the active one.

By clicking on the rendered objects local colocalization parameters are computed and reported, which will be discussed in more detail in “Iso-colocalization object analysis” below.

There is one *Maximum Intensity Projection* (MIP) pipe available to redirect the data channels to. The *MIP rendering* of one channel (maybe one different from those used for colocalization) or the two channels under analysis can be a good *spatial reference* for the objects from the colocalization map. When an original channel is selected, the *threshold slider* can be used to select what voxels are considered for the MIP rendering, depending on their intensities. Notice that here the threshold is simply used for *representation*. If both R and G channels are selected to be rendered, their correspondent backgrounds as selected in the histogram will be used as projection thresholds. As with the surface pipes, the *transparency* and the *brightness* of this MIP can also be controlled.

All the obtained information can be saved to external data (text or image) files through the FILE menu.

Iso-colocalization object analysis

One of the features of the colocalization analyzer is *iso-colocalization object analysis* (Figure 15.3 on page 99). It allows to determine quickly the properties of the different colocalization regions in the data. This is realized by visualizing the colocalization map as iso-colocalization surfaces. In this way regions in which the degree of colocalization exceeds a certain value become objects. By clicking on the objects local colocalization parameters are computed and reported. To relate the iso-colocalization objects to the original data the surface objects can be blended with a MIP projection of the data (See “The MIP Renderer” on page 35). The color range in which these objects are displayed can be modified using a hue selector (page 125).

Backgrounds vs. thresholds in colocalization

Backgrounds are for removing signal prior to the calculation. In an ideally restored image that would not be necessary, because all the signal present in the image gives a valid representation of it. The background would have been removed during the image restoration.

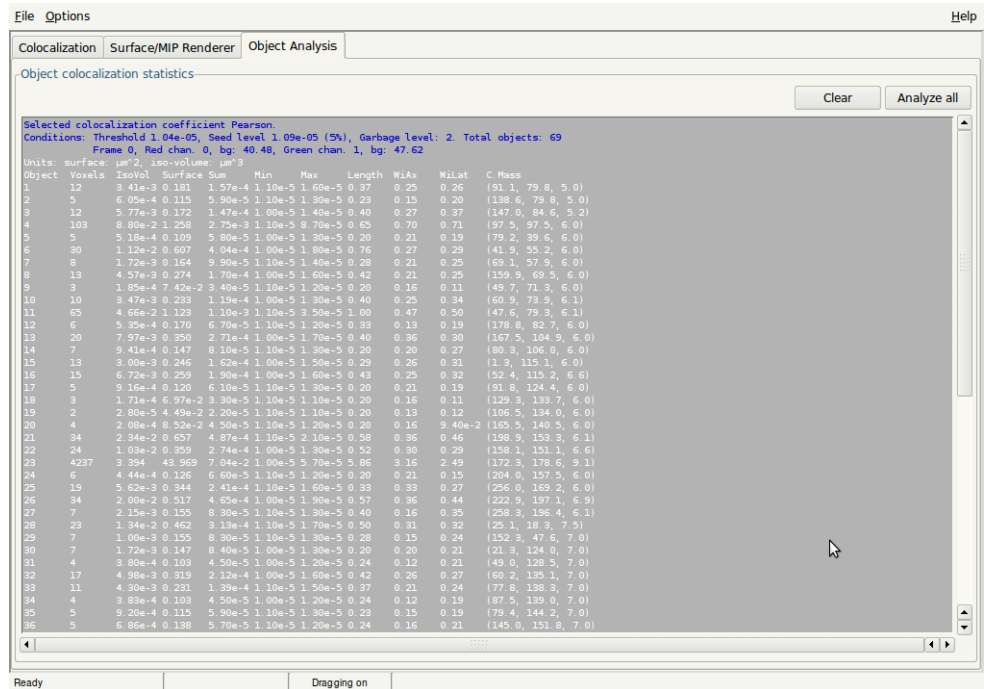


FIGURE 15.3. Tab three of the Colocalization Analyzer if for object analysis for the colocalized regions of the colocalization map.

Thresholds are used to split the colocalization maps in two regions: what are interesting objects, and what are not. Local colocalization values are calculated for every image voxel, but the zero value would be very rarely achieved. In most of the cases there are some non-zero colocalization everywhere, the interest probably is in regions where colocalization exceeds typically a value of 20 % of the maximum.

With the background settings, voxels are in- or excluded based on their voxel intensities: intensities higher than the background are kept, while lower intensities are disregarded for the calculation. Note, the intensity value of voxels do not relate to the level of colocalization. In fact, it may well be that high colocalization levels occur in regions with low voxel intensities.

Therefore, background settings should not be used to remove signal that can still have some colocalization level. Ideally work with deconvolved images and consider all the signal, then study the colocalization levels using thresholds to split regions of high colocalization from regions with low colocalization, something that can not be know beforehand! Setting the background is used for removing really constant background signals.

Thresholds do not affect colocalization, but only the way colocalization maps are represented on the screen and objects are defined. Depending on the locally calculated colocalization objects may merge or split. What affects colocalization, as explained in the colocalization theory at the SVI Wiki, are the backgrounds. In the computation of *Manders* coefficients the background values act like in this example: to the computation of *M1* only pixels in *R* contribute when their corresponding pixel in *G* is above the background.

More information can be found at the SVI Wiki³.

3. <http://support.svi.nl/wiki/ColocalizationAnalyzer>

The Chromatic Shift Corrector is a post-deconvolution tool that can estimate and correct for chromatic shifts, removing the existing misalignments across different channels. The result of this correction is a channel-aligned image free of chromatic shifts.

The support for templates included in this tool allows to apply the chromatic shift correction of one image to other images. This is particularly interesting when the estimation carried out on a beads image is suitable for other sets of images.

The Chromatic Shift Corrector shows its Help content dynamically to guide the user through the process of estimating and correcting for chromatic shifts.

Starting the Chromatic Shift Corrector

- Launch Huygens Essential or Huygens Professional.
- Load a multichannel image to be corrected for chromatic shifts.
- Select the image thumbnail and in the top menu go to DECONVOLUTION->CHROMATIC SHIFT CORRECTOR.

The Chromatic Shift Corrector will open and show the image on an orthogonal slicer where the existing chromatic shifts can be seen in a 3D view. Below the orthogonal slicer a Z slicer, a time slicer, and other visualization tools such as contrast, colour scheme, channel selection and zoom tools can be found. These tools are useful to enhance the view of the image for a better visualization of the chromatic shifts.

A view of the Chromatic Shift Corrector at start-up with a loaded two-channel bead image is shown in Figure 16.1.

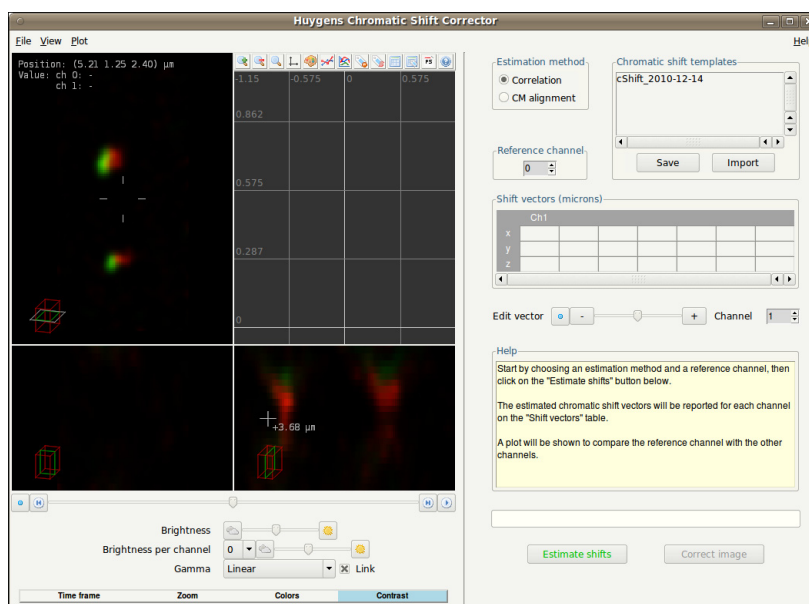


FIGURE 16.1. The Chromatic Shift Corrector at start-up, prior to estimating the chromatic shifts.

Estimation of the chromatic shifts

For the automatic estimation of the chromatic shifts the following two methods can be chosen:

- Cross correlation. This can be considered an 'all-round' method. The software searches for the best alignment across channels by maximizing the overlap.
- Center of mass alignment. This method works best if the image contains a single object. The object should not touch the image borders, and the contrast between object and background should be high.

The chromatic shifts will be quantified by vectors, indicating how much a channel is shifted with respect to a given reference channel.

The channel to act as reference (no shift) can be selected by the user via the Reference Channel selection box. Because this channel will have no chromatic shift its shift will not be reported.

If an estimation method and a reference channel have been selected, the chromatic shifts will be estimated and reported upon clicking on the ESTIMATE SHIFTS button.

Visualization of the chromatic shifts

The estimated chromatic shifts will be reported on the shift vectors table. The shifts of each channel but the reference will be listed. The length unit of the shift vectors is set to micrometers.

The user can select any channel but the reference for editing. This can be done by using the EDIT CHANNEL selection box under the vector table.

A vector will be drawn on the orthogonal slicer showing the estimated shift vector of the edited channel, so that the estimated shift can be easily assessed. At the same time, a plot

will be displayed showing the intensity profiles of the edited channel and the reference channel along the direction of the estimated shift vector.

The plot also shows the intensity profile of the edited channel as if it were corrected with the existing estimated shift (dashed line).

Therefore, the plot serves as a comparison between the reference channel, the edited channel and the corrected edited channel.

This comparison allows us to see to what extent the intensity of edited channel is shifted with respect to the reference channel. At the same time, it shows in advance whether the estimated shift vectors will correct for the existing chromatic shifts properly.

The result of a chromatic shift estimation is shown in Figure 16.2. The shift between the intensity profiles of the reference channel and the edited channel is visible in the embedded plot. Additionally, the dashed line in the plot shows the intensity profile of the edited channel as if it were corrected with the estimated shift.

Ideally, the intensities of the corrected channel (dashed line) and the reference channel would show no gap or shift, having similar shapes and peaks roughly located at the same x positions.

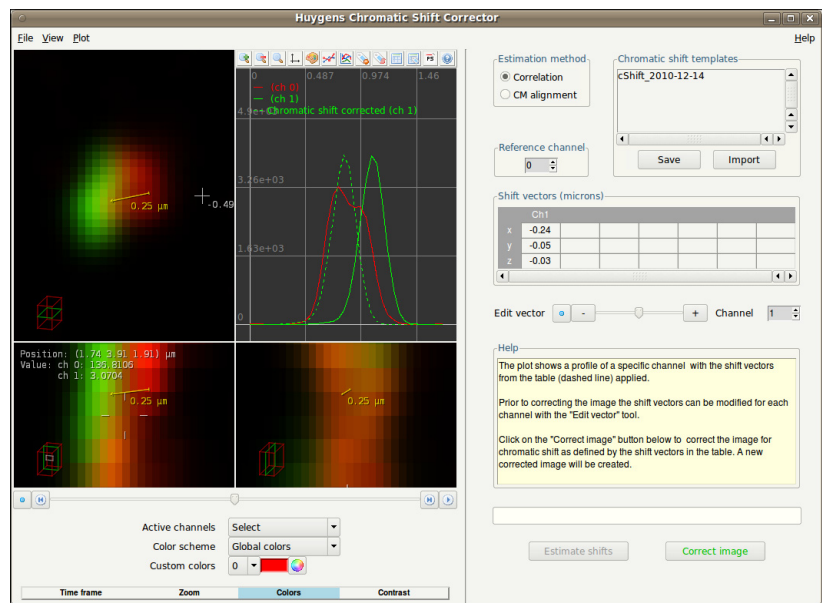


FIGURE 16.2. The Chromatic Shift Corrector after estimating the shifts. A plot shows the gap between two channels. The estimated shift is drawn as a vector on the slicer.

The Chromatic Shift Corrector will return accurate and reliable estimations of the existing chromatic shifts. Nevertheless, the possibility to edit and customize the estimated shifts exists so that the user can reach more precision if necessary.

Editing the chromatic shift vectors

The user can also edit the shifts estimated automatically by the Chromatic Shift Corrector. To apply a different customized correction a channel has to be selected for edition. The components of the selected shift vector can be modified by using the edit tool, which allows to shorten and lengthen the estimated shift.

While the shift of the edited channel is shortened or lengthened the plot updates itself to show how the gap between the reference channel and the edited channel is increased or decreased.

In this way, the contents of the estimated shifts can be modified while checking in advance whether the image will be corrected appropriately.

The chromatic shift estimation of a four-channel image is shown in Figure 16.3. The channels can be edited one by one and their shifts modified with the EDIT CHANNEL tool.

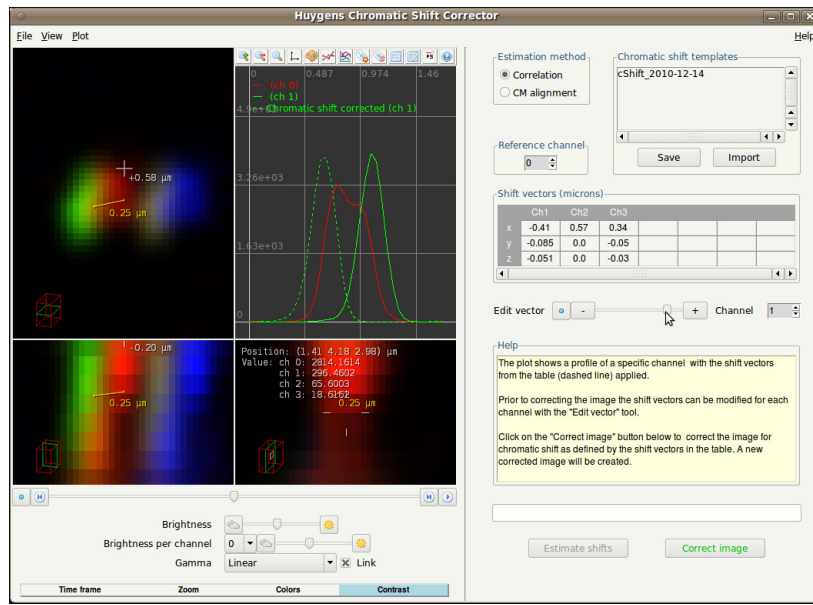


FIGURE 16.3. Using the Edit Channel tool of the Chromatic Shift Corrector. The estimated vector has been lengthened, and the dashed line plot is consequently slightly moved to the left.

Alternatively, any customized shifts can be typed in the vector table for each channel.

Upon clicking on the CORRECT IMAGE button the image will be corrected with the chromatic shifts listed in the vector table. A new corrected image will be created.

Working with templates

The estimated shift vectors can be saved to a template by using the template tool of the Chromatic Shift Corrector.

The saved templates can be imported again to apply its shift vectors to other images. The template vectors will be loaded and listed in the vector table, showing the intensity profiles as if the ESTIMATE SHIFTS button had been pressed.

The image can then be corrected by clicking on the CORRECT IMAGE button, which will create a new corrected image.

In order to be able to work with the Chromatic Shift templates properly it is recommended to apply these templates to images that have the same emission and excitation wavelengths as the image of the template.

The Huygens PSF distiller facilitates the creation of a PSF from bead images. A wizard guides you in selecting your bead images creating and saving the PSF for further use in deconvolution runs. The wizard is able to measure a PSF from one or more images of fluorescent beads, each containing one or more beads. It is also able to distill multi-channel PSF's from information collected from multi wavelength beads, or assemble a multi-channel PSF from single channel PSF's.

Acquiring the Beads

The PSF distiller works by inspecting average images from small beads that are almost sub resolution in size, so they contain much of the PSF information. The diameter of the beads should preferably be in the order of the half intensity width of the expected PSF. Larger beads will reduce the accuracy of the Distiller; smaller beads yield insufficient signal for accurate stacking in the averaging procedure, resulting also in reduced accuracy. Typically beads with a diameter of 160 nm perform very well for many types of microscopy.

Confocal and two photon bead images from 160 nm beads should look like smooth fuzzy blobs with hardly visible noise. Use the default SNR settings. If available, it is a good idea to average 2 to 5 beads. Two photon bead images may look slightly noisy. If so, set the SNR to 20 and average 4 to 10 beads.

Widefield images taken from 160 nm beads should look like smooth fuzzy blobs with no visible noise. The default SNR settings can be used. It is not necessary to average any more beads.

Beads should be recorded with the same microscopic parameters that you will use later to image your specimens. More information on bead images and the PSF Distiller is available on the SVI wiki pages.¹

The PSF Distiller Window

If your license includes the PSF Distiller option, select DECONVOLUTION→PSF DISTILLER to start the Distiller. In the opening window (Figure 18.1), different panes show the input beads field, the Help field, the report field, the wizard fields and the PSF and Accus field:

1. <http://support.svi.nl/wikiRecordingBeads>

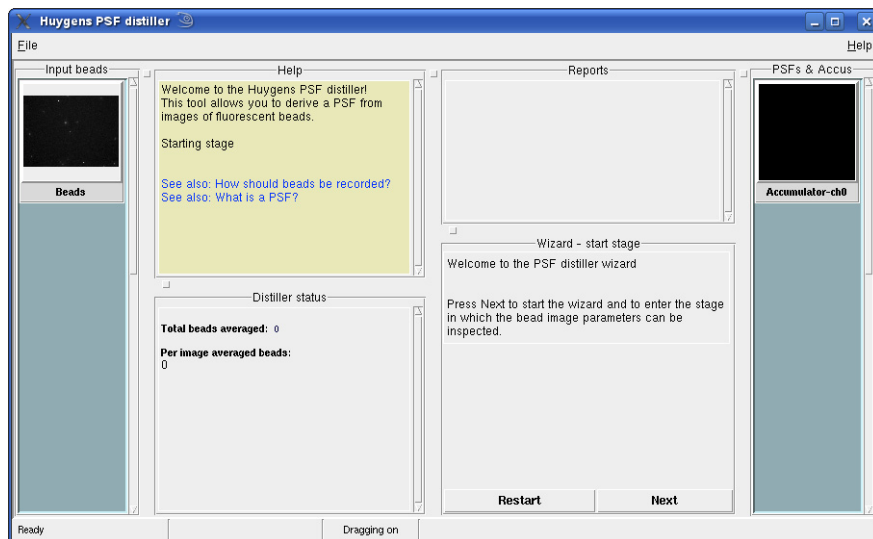


FIGURE 18.1. The PSF Distiller window. Different panes show the input beads field, the Help field, the report field, the wizard fields and the PSF and Accus field.

- *The Input beads pane*
Shows the selected file with the bead images for the PSF distiller process. You can import and use multiple files with beads for the distillation of a PSF. With the cursor you can hover over the thumbnail and with a right mouse click you can roll down a menu to open the image in one of the slicers. The edit parameters function is greyed out here, you can only modify the microscopic parameters in the main microscopic parameter window or the PSF wizard.
- *The Help pane*
Explains the different steps and displays links to relevant wiki pages to explain every step in detail.
- *The Distiller status pane*
Keeps track of all the steps during the distiller run and reports the progress.
- *The Reports pane*
Displays the progress report of the distiller process and the status of the distilling process.
- *The PSF and Accus pane*
Displays the intermediate results of the distilling process. The thumbnails can be inspected with the Slices. The slider on the right side can be used to scroll down if many files are displayed.
- *The Wizard pane*
Shows the steps that will take you from checking image parameters to distilling and averaging the beads from your images.

Using the PSF Distiller Wizard

If you push the NEXT button, the wizard displays the Parameter checking pane (Figure 18.2). In this stage you can again alter the parameters of the selected bead image. The APPLY button is required to bring these changes into effect.

If any of the entry fields for the sampling density turn orange or red, the data is unusable. Undersampled bead images can not be used to measure PSF's.

After pushing the NEXT button again, the wizard will ask you to confirm the bead size and set the signal to noise parameters. Now the distiller process starts with the automated search for beads that are suitable for the distiller, and it will average the selected beads in the Accumulator.

The image is searched for beads which meet all selection criteria:

- A bead should not be too close to another bead. If it is, their images will interfere with each other. In widefield bead images this is quite problematic due to the large size of the blur cone. Fortunately, widefield PSF can be derived from single bead images.
- A bead should not be too close to an image edge. After all, another bead might be located just over the image edge.
- The intensity of a bead should not deviate too much from the median intensity of all beads. If it is brighter, then it may be a cluster of two or more beads. If it is dimmer, then it is not likely to be a bead. In both cases the object geometry is unknown, so they are unusable. If for some reason no beads are found, an explanation and some advice will be displayed in the help window. The software will try to find beads first with ideal selection criteria. If this does not yield a single bead then it will automatically retry with reduced inter bead distance criteria.

When the searching is finished, you may start the distillation or load another bead image to increase the number of beads. Pushing the LOAD ANOTHER BEAD button will open a pane with the file browser to select a second or more file with bead image(s) and display the thumbnail(s) in the input beads pane.

If you have averaged all the beads you can push the DISTILL button. Depending on the size of the file and the type of the microscope the distilling can take some time. Note the Help pane which explains details about the process.

For multichannel images you need multichannel PSF's. To add channels, push the ADD CHANNELS button to repeat the distiller process (See Figure 18.3). When there is enough data to estimate a chromatic shift between the channels, the shift will be reported in the Report pane.

If you are finished you can EXPORT & CLOSE the Distiller window. When returning to the main window you will find the exported PSF ready to save it and use it. The distilled ("measured") PSF can be inspected with the slicers and used to restore images recorded with the same conditions that have been used to record the bead images.

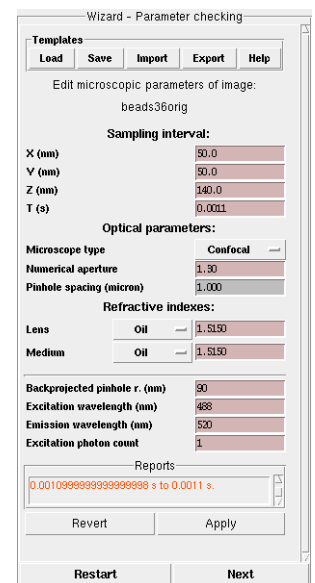


FIGURE 18.2. The Parameter checking pane. This pane gives you the tools to edit the microscopic parameters of your bead image.

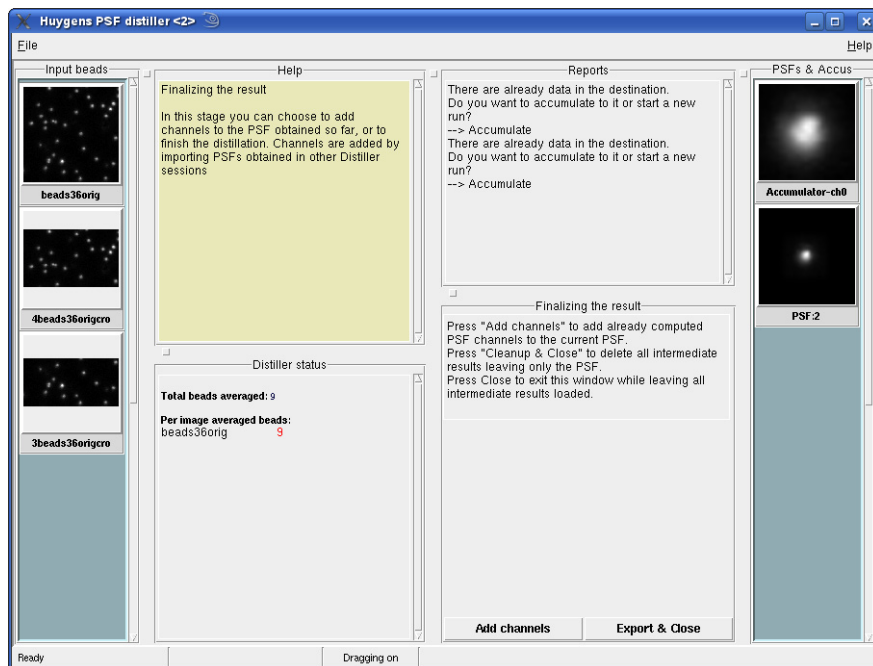


FIGURE 18.3. The PSF Distiller window after distillation. You can choose to add channels or export the result.

The deconvolution algorithm needs to know some of the parameters describing the image acquisition. These are not too many, but careful determination may significantly enhance the deconvolution results.

Image Size

The amount of computing time involved in deconvolving images is more than proportional to the image size. It is therefore sensible to limit the data size as much as possible. Regarding widefield images we recommend to not record planes below and above the object which only contain blur. Huygens Professional does not need these planes to restore the object. Since the blur in these planes might be affected by hard to correct bleaching they might even reduce the quality of the deconvolution result. In any case, never crop the objects of interest. As a rule of thumb, leave about *one extra μm* above and below the objects.

Brick-Wise Processing

Deconvolving images requires much computer memory than the image size, because all computations are done in 32 bit floating point format, and because several extra (hidden) images are needed to store intermediate results. To reduce the memory requirements Huygens Professional will *split the images* into bricks, deconvolve the bricks sequentially, and fit the bricks together in a seamless fashion. Brick-wise processing is an automatic feature of Huygens Professional. To find out the best number of bricks, let the software run in automatic mode for splitting. It will consider many options and go for the most optimal one.

More information can be found on the SVI support Wiki¹.

Signal to Noise Ratio

The Signal to Noise Ratio (referred to as SNR or S/N) is in the Huygens Software used as a regularization parameter, i.e. as a parameter that controls the *sharpness* of the restoration result. The higher this value, the sharper the restored image will be. Therefore it

1. <http://support.svi.nl/wiki/BrickSplitting>

should not be considered as a parameter describing the original image, but more as a tunable parameter that controls the deconvolved result.

Using a too large SNR value might be risky when restoring noisy originals, because the noise could just be enhanced. A noise-free widefield image usually has SNR values higher than 50. A noisy confocal image can have values lower than 20.

A good starting SNR value can be estimated simply by visual inspection from the quality of the image. Figure 18.1 shows some examples of recordings where different noise levels

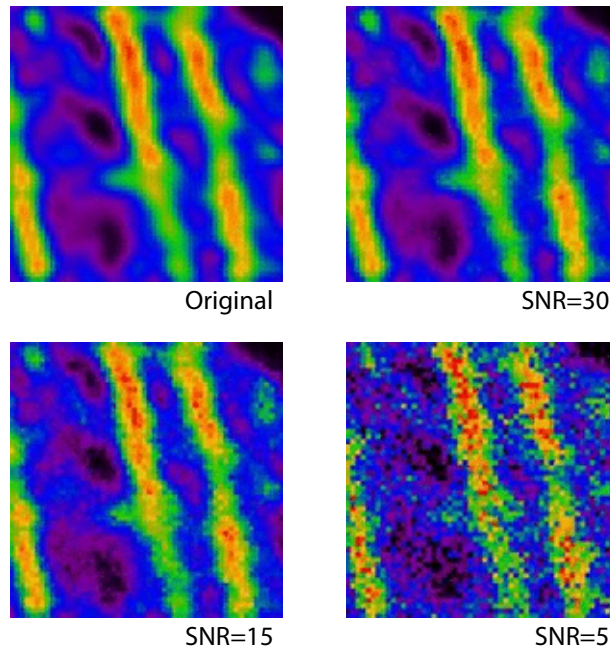


FIGURE 18.1. Examples of different SNR values. Same image, different noise levels.

were added to an original (restored) image.

Estimating the SNR in noisy images is fairly easy. It is based on the idea of establishing the voxel intensity value s corresponding to a single photon hit by looking for such an event in a dark areas of the image. Knowing the intensity value M of the brightest voxel in the image, one can now calculate how many photons are involved in it. The SNR is now defined as:

$$SNR = \sqrt{\frac{M}{s}} \quad (\text{EQ 1})$$

If a significant blacklevel is present, it should be subtracted from M and s . See “Black Level” on page 111 and the SVI support wiki² for more details.

In low noise images this is much more difficult, as single photon events are no longer easily observed. Fortunately, in such cases the establishment of the precise SNR is not very important for the restoration method, and a rough estimation based on the appearance of the image is usually enough (See Figure 18.1).

2. <http://support.svi.nl/wiki/SNR>

Black Level

The blacklevel (also called *base line* or *electronic shift*) is the output of the photo-multiplier if no light is coming through.

A positive black level (See Figure 18.2) will do no harm to the deconvolution since it is

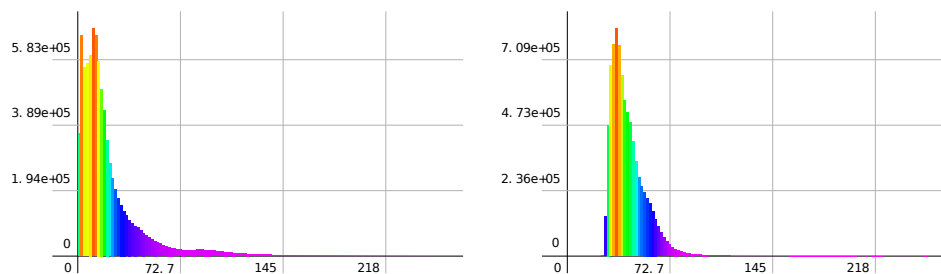


FIGURE 18.2. Left: The histogram of an image without black level. Right: The histogram of a similar image with a significant black level.

automatically accounted for in the background estimation stage. A large black level value, however, will reduce the effective dynamic range of the microscope. Besides that, a large black level will prevent the bleaching decay analyzer to do its job correctly, and the bleaching correction may turn out to be poor or even impossible. To prevent this Huygens Professional has a tool to adjust the base line of an image, and the Batch Processor templates show the option to remove the black level as a preprocessing step.

It is also possible that the black level is *negative*. In the image histogram this will show as a spike on the left. This causes clipping (See “Clipping” on page 118) in the lower intensity limit, and it is impossible to correct: clipped images should be reacquired.

More information can be found on the SVI support Wiki³.

Sampling Density

The sampling density is the number of recorded *samples per unit area volume (3D) or area (2D)*. It is a microscopic parameter that describes the conditions of the image acquisition, established by the way the microscope is configured (usually by the zoom factor). The ideal sampling density depends on the system optics and is determined by the *Nyquist rate*. It is recommended to sample as close to the Nyquist rate as possible. Note that the *actual* sampling distances from the acquisition must be used in the deconvolution.

The SVI support Wiki has an on-line tool⁴ that computes the Nyquist rate for any wide-field, confocal, spinning disc, and 4-Pi microscope.

Sampling according to the Nyquist rate makes sure that *all* information generated by the optics of the microscope is captured in digital form. It can be shown that if the sampling distance is smaller than the so-called *critical sampling distance*, no new information about the object is captured. Apart from practical problems like bleaching, acquisition time and data size there is no objection at all against using a smaller sampling distance than the critical distance, on the contrary.

3. <http://support.svi.nl/wiki/BlackLevel>

4. <http://support.svi.nl/wiki/NyquistCalculator>

Figure 18.3 shows the dependency of this critical sampling distance on the numerical

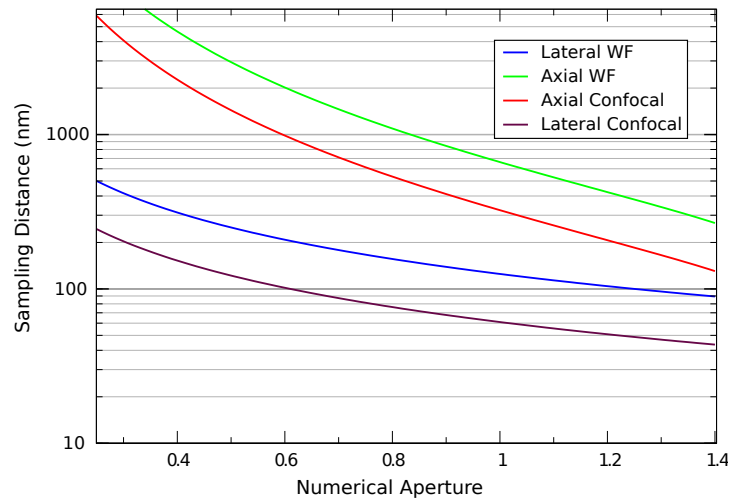


FIGURE 18.3. Critical sampling distance vs. NA. The curves above show the critical sampling distance in axial and lateral directions for wide-field and confocal microscopes. The emission wavelength in both cases is 500 nm.

aperture for a wavelength of 500 nm. To apply this plot of to another wavelength, simply scale the vertical axis by that wavelength. For example, if a widefield microscope with NA 1.3 is used, there can be noticed from the plot that the critical lateral Nyquist sampling distance at 500 nm emission is 95 nm. For an emission wavelength of 570 nm, this becomes $(500/570) \times 95 = 108$ nm .

In the confocal case it is the excitation wavelength which determines the Nyquist sample distance. In theory the pinhole plays no role, but larger pinholes strongly attenuate fine structures at the resolution limit. Therefore, as a rule of thumb, with a common pinhole diameter of 1 Airy disk the lateral critical sampling distance may be increased by 50 % with negligible loss of information. In cases were the pinhole is much larger, the lateral imaging properties much resemble those of a widefield system and the sampling distance can be set accordingly. We do not recommend to increase the axial sampling distance appreciably beyond the critical distance.

In a multi photon excitation microscope, it is the excitation wavelength divided by the photon count which determines the sampling.

More information can be found on the SVI support Wiki⁵.

Computing the Backprojected Pinhole Radius and Distance

Throughout Huygens Professional pinhole sizes of confocal systems are specified as the *backprojected radius* in nm. *Backprojected* means the size of the pinhole *as it appears in the specimen plane*, i.e. the physical pinhole radius r_b divided by the total magnification of the detection system. This total magnification is the product of the (variable) objective magnification m_{obj} multiplied by a fixed internal magnification m_{sys} :

5. <http://support.svi.nl/wiki/NyquistRate>

$$r_b = \frac{r_{phy}}{m_{obj} m_{sys}} \tag{EQ 2}$$

The SVI support Wiki has a calculator⁶ to automatically compute the backprojected pinhole radius for specific microscope models. Note that the Bio-Rad MRC500, 600, and 1024 microscopes have a very high magnification in the detection system (See “Checking the Bio-Rad System Magnification” on page 114).

The equations that can be found in the next pages are intended to orientate the user in finding out the backprojected value for different types of microscopes, but the idea is always the same: given a diameter d of the real pinhole, we might need to multiply it by a factor for unit conversion (to obtain the radius in nm), and divide the result by some other factor that takes account of the magnification of the microscope. These include both the objective and the intrinsic system magnification. In some microscopes with pinholes that are not circular, a geometrical correction will also be needed.

The Airy Disk as Unit for The Backprojected Pinhole

Some confocal microscopes report their pinhole size with the *diameter of the Airy disk*⁷ as unit. shows how to compute the backprojected pinhole radius:

$$r_b = \frac{0.61 \lambda_{ex} N_{Ad}}{NA} \tag{EQ 3}$$

with NA the numerical aperture of the lens, N_{Ad} the number of Airy disks, and λ_{ex} the excitation wavelength. In principle using λ_{ex} is not correct because the Airy diffraction pattern is formed by the *emitted* light. However, we suspect microscope manufacturers prefer to use the excitation wavelength because it is better defined and does not depend on settings of devices like adjustable band filters.

Note that this relation bypasses the need to know internal system and lens magnifications.

Square Pinholes and the Shape Factor

The shape correction from a square to a circular pinhole is based on equalizing the area for both pinholes.

Because the area $A = d^2 = \pi r^2$ (See Figure 18.4), the shape factor c for a square pinhole becomes

$1/(\sqrt{\pi}) = 0.564$. For circular pinholes $c = 0.5$ to just convert the diameter to a radius.

If the size of a square pinhole is given in Airy disk units, then the backprojected radius is a combination of Equation 3 and the shape factor c for a square pinhole. Taking into account that Equation 3 already converts a diameter into a radius, the relation becomes:

$$r_b = \frac{0.69 \lambda_{ex} N_{Ad}}{NA} \tag{EQ 4}$$

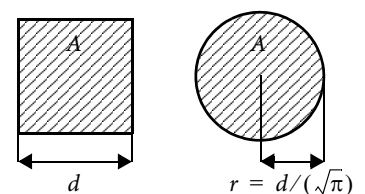


FIGURE 18.4. A square and a disk with an equal area A .

6. <http://support.svi.nl/wiki/BackprojectedPinholeCalculator>
 7. <http://support.svi.nl/wiki/AiryDisk>

Converting from an Integer Parameter

Unfortunately, quite a few microscopes do not report the physical pinhole size or the Airy disk size. Instead, often an integer size parameter p_8 in the range [0,255] (8 bit) is specified. If p_8 maps to a physical size in linear fashion, then Equation 5 can be used to translate that parameter into a backprojected pinhole radius.

$$r_b = 10^3 c \frac{(p_8/255)(s_{max} - s_{min}) + s_{min}}{m_{obj}m_{sys}} \quad \text{(EQ 5)}$$

Here, c is the shape factor and s_{min} and s_{max} are the minimal and maximal pinhole size in μm . The factor 10^3 converts from μm to nm.

The backprojected Pinhole Spacing for Spinning Disks

As is the case for the backprojected pinhole diameter, the distances *between* the pinholes in spinning disks must be divided by the system magnification. For the frequently used Yokogawa disks⁸, for example, the physical spacing is about 253 μm . This can be checked by imaging a stopped disk. Using a 100 \times lens for example, the backprojected distance is about 2.53 μm . If an extra zoom lens is placed between the disk and the sample, its magnification must be also considered.

A Supplied Calibration Curve

If a calibration curve was supplied with the microscope best use that curve to convert the displayed setting to a physical size and from there convert to the backprojected radius.

An Example: the Olympus FV1000

As reported by Olympus engineers, the FV1000 confocal microscope has a square pinhole and an internal magnification of 3.82 \times . The pinhole size reported in the Olympus software is the side length of this square pinhole. Therefore, $c = 1/(\sqrt{\pi}) = 0.564$, and $m_{obj} = 3.82$. Using a 100 \times objective lens and a reported pinhole size d of 150 μm the backprojected pinhole radius r_b in nm is:

$$r_b = c \frac{d}{m_{obj}m_{sys}} = 0.564 \frac{150 \cdot 10^3}{100 \cdot 3.82} = 221 \quad \text{(EQ 6)}$$

The on-line calculator⁹ knows the system magnifications and pinhole shapes for most popular microscopes.

Checking the Bio-Rad System Magnification

The Bio-Rad MRC 500, 600, and 1024 microscopes have a very high magnification in the detection system. The fixed system magnification is, according to Pawley¹⁰,

8. <http://support.svi.nl/wiki/YokogawaDisk>

9. <http://support.svi.nl/wiki/BackprojectedPinholeCalculator>

10. Pawley, J. B., *Handbook of Biological Confocal Microscopy*, 2nd edition, 1995. Plenum Press, New York and London. ISBN 0-306-448262. Page 30.

$53 \times m_{tube}$, with m_{tube} between 1.0 and 1.56 (factor 1.25 for the fluorescence attachment multiplied by factor 1.25 for the DIC attachment). The factor of 53 includes the 8× eyepiece just below the scan head, but does not include that variability in magnification due to the variations in tube-length that are result from the aligning the system.

The high system magnification allows the viewing of the diffraction pattern (Airy disk) at the pinhole plane directly by eye. To enable the correctness verification of the values for the system magnification used in the on-line calculator, the way the system magnification was derived is now explained.

In a Bio-Rad MRC600 with a NA 1.3 60× objective, the Airy disk has a diameter of around 2 to 2.5 mm at the pinhole plane. The diameter of the first Airy zero ring is 7.6 lateral optical units¹¹ (o.u.), using Equation 7 to express a distance r in dimensionless o.u.:

$$v = r \frac{2\pi}{\lambda} NA \quad (\text{EQ 7})$$

In the system described here, an o.u. is 0.3 ± 0.033 mm. At the specimen plane (backprojected) a lateral o.u. is in this case around 61 nm. The total magnification is in that case 4918×, the system magnification $4918/60 = 82 \pm 9$ ×. This value corresponds well with the largest possible system magnification for the MRC600 ($53 \times 1.25 \times 1.25 = 83$).

11. <http://support.svi.nl/wiki/OpticalUnits>

In this chapter discusses basic suggestions on how to acquire better microscope images. These are based on common problems that we find frequently in data provided by users. The recommendations go on the direction of obtaining the highest quality images from the microscope, from the point of view of acquiring as much information as clean as possible. This alone is worth the effort, but it will also be very valuable for the deconvolution afterwards.

Some basic guidelines to improve the deconvolution results are also listed. More information can be found in the SVI support Wiki¹.

Data Acquisition Pitfalls

Refractive Index Mismatch

A mismatch between the refractive index of the lens immersion medium and specimen embedding medium can cause several serious problems:

- **Geometrical distortion:** Frequently referred to as *the fishtank effect*². The axial sampling distance that is recorded in the image file is the step size by which the objective lens moves along the *z*-axis. The focal point inside the sample, however, shifts due to the fishtank effect by a different step size. Therefore objects will appear elongated or shortened in the image data.

Huygens Professional will automatically adapt the PSF to this situation, but it will *not* modify the image geometry. *After deconvolution* the geometric distortion can be corrected by multiplying the *z*-sampling distance by the medium refractive index divided by the immersion refractive index.

- **Spherical aberration:** this phenomenon causes the oblique rays to be focused in a different location than the central rays³. The distance in this focal shift is dependent on the depth of the focus in the specimen. If the mismatch is large, e.g. when going from *oil immersion into a watery medium*, the PSF will become *asymmetric* at depths of already a few microns. This is especially harmful for the deconvolution of wide-field images.

1. <http://support.svi.nl/wiki/>

2. <http://support.svi.nl/wiki/FishtankEffect>

3. <http://support.svi.nl/wiki/SphericalAbberation>

A *workaround* for this problem is to keep the z-range of the data (the number of slices) as small as possible. The *solution* is to use a water immersion lens instead.

- **Total internal reflection:** When the *numerical aperture of the objective lens is larger than the medium refractive index*, total internal reflection⁴ will occur. This is causing excitation light at high angles to be bounced back into the lens and therefore limiting the effective NA.

If spherical aberration is unavoidable, the image can still be improved during restoration using an adaptive point spread function (See “Refractive Index Mismatch” on page 117).

Clipping

The intensity of the light emitted by the microscopic sample is converted to electrical signals that pass an adjustable amplifier. These electrical signals must be converted to numbers processed by the computer. This conversion is done by the CCD camera. Most scientific CCD cameras have a 12 bit converter limiting the output numbers to the range [0,4095] . Negative input signals are usually converted to 0 while positive input values exceeding the *dynamic range*⁵ are all converted to 4095. This phenomenon is called *clipping*; information in the clipped samples is lost.

In practice: be suspicious if the data contains intensity values at the extremes of the numerical range. These ranges are [0,255] for 8 bit data, [0,4095] for 12 bit data, and [0,65535] for 16 bit data.

Read more on the SVI support Wiki⁶.

Undersampling

One of the rules of measurement that is often overlooked is sampling according to the *Nyquist rate* (See “Sampling Density” on page 111). Especially the sampling distance along the optical axis is frequently too large; too few *xy* slices are imaged. This leaves as result a 3D stack in which there is hardly any relation between the adjacent slices.

It is important to know how the sampling conditions should be established in order to recover an image from the sampled values. How the objects should be sampled depends on the microscope type (widefield, confocal, etc.) and on the microscopic parameters used, like the numerical aperture and wavelength.

The SVI support Wiki has an on-line tool⁷ that computes the ideal (Nyquist) sampling distances for any widefield, confocal, spinning disc, and 4-Pi microscope.

Do Not Undersample to Limit Photodamage

Sometimes undersampling is done to limit photodamage to live cells. However, if photodamage plays a role it is actually better to *limit the number of photons per sample* than to limit the number of samples. Having less photons per sample means that the overall photon dose can remain largely constant; as a result bleaching does not need to get worse.

4. <http://support.svi.nl/wiki/TotalInternalReflection>

5. <http://support.svi.nl/wiki/DynamicRange>

6. <http://support.svi.nl/wiki/ClippedImages>

7. <http://support.svi.nl/wiki/NyquistCalculator>

Although this results in an apparently noisier image, there is actually *more information* in the data because sampling density is higher. Deconvolution with Huygens Professional removes the noise and may enhance the image quality dramatically. Of course there are limits, but a fair trade-off can be often found.

Concluding, it is better to record 10 separate noisy slices 100 nm apart than two slices 1000 nm apart and averaged 5 times in order to reduce noise.

Bleaching

Bleaching is a practically unavoidable phenomenon in fluorescence microscopy. Because the image planes are acquired sequentially, bleaching will vary along the z direction. Assuming it is not strong it will not affect deconvolution results on confocal or two photon images. But in widefield deconvolution bleaching is more of a problem. Fortunately, the bleaching in widefield images can usually be corrected quite easily. Huygens Professional will do so automatically. However, if the bleaching is strong the correction might not be perfect, resulting in lower quality deconvolution results.

Illumination Instability

Some widefield systems are equipped with unstable arc lamps which can cause amongst others jitter. These are irregular deviations from the average intensity in time. Huygens Professional will correct this instability, but when the instability is severe it cannot do so sufficiently.

Mechanical Instability

Mechanical instability can take many shapes, for example:

- Vibrations sometimes seen in confocal images. They may seriously hamper deconvolution.
- The z -stage moves irregular or with sudden jumps. This deforms the data along the z -axis and is fatal for widefield and confocal deconvolution.
- The specimen moves. If in widefield data the object can clearly be seen moving when slicing along over a few μm in z . This will cause problems for the deconvolution. Best cause of action, apart from speeding up acquisition, is limiting the z -range of the data as much as possible. Confocal data of moving specimen causes less problems.

Thermal Effects

Thermal effects are known to affect calibration of the z -stage, especially if piezo actuators without feedback control are used. In particular harmful for widefield data. In time series the effect can be seen as a drift of the z -position, or even a periodic movement induced by e.g. an air-conditioning system switching on and off. The z -drift corrector is able to correct this in most situations though.

Internal Reflection

At high NA the angle of incidence of the most oblique rays can be close to 70 degrees. When a ray has to cross the cover-glass to medium interface at such an angle total reflection may occur. To be precise, total reflection occurs when the NA of the lens is higher than the refractive index of the embedding medium. This will reduce the effective NA of the lens.

Acquire an Experimental PSF

A *point spread function* or *PSF* is the image of a single point object⁸. The degree of spreading (blurring) in the image of this point object is a measure for the quality of an optical system. The imaging in a fluorescent microscope is completely described by its PSF. Although in many cases a theoretically calculated PSF very well matches the real one, ideal theoretical calculations can not predict actual misalignments or other problems inside the optical path. Therefore it is always recommendable to measure an experimental PSF and, if it is very different from the ideal one, use it for deconvolution instead of the theoretical one.

The experimental PSF can be measured by acquiring the image of a small bead. When the size of the bead is known, then the PSF Distiller can distill the real shape of the PSF from the recording (See “The PSF Distiller” on page 111).

Spherical Aberration Correction

When there is a refractive index mismatch, then the Huygens software automatically correct for spherical aberration⁹ by adapting the theoretical PSF to the sample depth.

In case the image suffers from severe spherical aberration, it might be better to use a theoretical PSF with this depth-dependent correction than an experimental one.

Improve the Deconvolution Parameters

Some deconvolution parameters, for example the SNR and the background level, can be fine-tuned to get the best out of the restoration process. See the SVI support Wiki¹⁰ for detailed steps in configuring the restoration process.

8. <http://support.svi.nl/wiki/PointSpreadFunction>

9. <http://support.svi.nl/wiki/SphericalAberration>

10. <http://support.svi.nl/wiki/DeconvolutionProcedure>

License String Details

Detailed information about the installed license strings can be displayed via HELP→LICENSE. Select the license string of interest and click EXPLAIN LICENSE.

A Huygens license string consists of a set of substrings separated by dashes (-). These substrings describe e.g. the product, version number, options, etc. The checksum at the end of the string should match with all other substrings. A complete string looks like this:

```
HuPro-3.7-wcnp-d-tvAC-emnps-eom2012Dec31-
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

Table 20.1 lists the building blocks from which this string is composed.

TABLE 20.1. The building blocks of the Huygens license string.

Substring	Description
Product	The product to which the license string applies. This can be HuEss, HuPro, HuScript, and HuCore.
Version	The version number of the product.
Microscope types	This substring consists of one or more characters representing the microscope types for which the deconvolution is enabled. These are 'w' (widefield), 'c' (confocal), 'n' (Spinning disk), 'p' (multi-photon), and '4' (4-Pi experimental microscopes).
Server flag	Determines the number of cores that are enabled for multi-threading. A hyperthreaded core is counted as a single core. It can be 'd' (desktop; 2 cores), 's' (small server; 4 cores), 'm' (medium server; 8 cores), 'l' (larger server; 16 cores), and 'x' (extreme server; 512 cores).
Option flags	This is a set of characters that list the enabled optional modules. An overview of these modules is given in Table 20.2.
Locking policy	A set of characters that indicate to which properties the license is locked. These can be 'd' (expiry date), 'e' (e-mail address), 'm' (system ID), 'n' (number of cores), 'p' (processor type), and 's' (processor details).

TABLE 20.1. The building blocks of the Huygens license string.

Substring	Description
Expiry date	The date on which the license or maintenance ends. When this substring starts with 'eom', then only the maintenance expires; the license remains valid.
System ID	A 16 character hexadecimal string containing hardware identification numbers.
E-mail address	The customer e-mail address.
Checksum	A 20 character hexadecimal checksum on the previous substrings.

TABLE 20.2. Identifiers for the optional modules.

Character	Module
a	Object Analyzer
A	Advanced Object Analyzer
b	Small file reader bundle
B	Complete file reader bundle
C	Colocalization Analyzer
f	PSF Distiller
M	Movie Maker
N	Enable new, pre-released features
s	Object Stabilizer
t	Time series
v	Surface Renderer visualization

The Point Spread Function

One of the basic concepts in image deconvolution is the *point spread function* (PSF). The PSF of the microscope is the image which results from imaging a point object in the microscope. Because of wave diffraction¹ a point object is imaged (spread out) into a fuzzy spot: the point spread function. In fluorescence imaging the PSF completely determines the image formation. In other words: *all microscopic imaging properties are packed into this 3D function*. In Huygens Professional, a PSF can be obtained in two different ways:

1. *Generating a theoretical PSF*: When a measured PSF is not available, Huygens Professional automatically uses a theoretical PSF. The PSF is computed from the microscopic parameters attached to the data. Because a theoretical PSF can be generated without any user intervention Huygens Professional does the calculation in the background without any notice.
Images affected by spherical aberration (See “Refractive Index Mismatch” on page 117) are better restored using a theoretical depth-dependent PSF.
2. *Measuring a PSF*: By using the PSF Distiller a measured PSF can be derived from images of small fluorescent beads (See “The PSF Distiller” on page 111). Measured PSF's improve deconvolution results and may also serve as a quality test for the microscope

1. <http://support.svi.nl/wiki/ImageFormation>

Quality Factor

Deconvolution as it is done in Huygens Professional is based on the idea of finding the best estimate of the object that is imaged by the microscope. To assess the quality of an estimate, Huygens Professional simulates the microscopic imaging of each estimate (the estimated is convolved with the PSF) and compares the simulation with the measured image. From the difference a quality factor is computed. The difference is also used to compute a correction factor to modify the estimate in such a way that the corrected estimate will yield a better quality factor. The quality factor as reported by the software is a measure relative to the first estimate and therefore a number greater than or equal to 1. If the increase in quality drops below the *quality threshold* the iterations are stopped.

File Series

There are many ways in which Tiff files or other file series are named. These files can have multiple counters (referring to *slices*, *time frames*, or *channels*), and these counters can have arbitrary prefixes and ordering.

Numbered Tiff Series

If a series is simply numbered like: `slice001.tif`, `slice002.tif`, ..., `slice0nn.tif`, then Huygens Professional will read the series into a single 3D image. Because Tiff files usually carry no additional microscopic information, check the parameters carefully.

Leica Numbering

Huygens Professional natively supports both reading and *writing* Tiff series with Leica style numbering, if there is more than one channel, slice, or time frame. A single channel 2D time series would be numbered according to the scheme:

```
im_tNN.tif
```

Here, NN is replaced by the time index for each frame. A more complex, multi-channel 3D time series has this pattern:

```
im_tNN_zNNN_cNN.tif
```

In this series, the second channel of the fourth slice of the third time frame has the file-name:

```
im_t02_z003_c01.tif
```

The File Series Tool

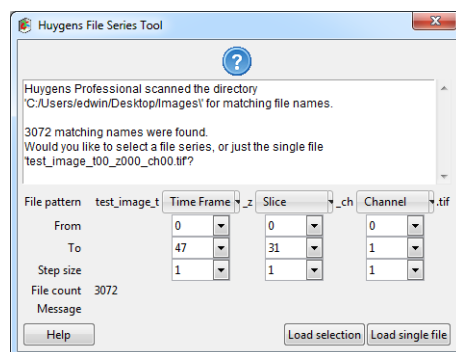


FIGURE 20.1. The Huygens File Series Tool automatically scans a directory for a file series.

Although Huygens Professional uses Leica style numbering for writing files, the software attempts to detect *any type* of file series for reading. Whenever a file is opened that appears to be part of a file series, Huygens Professional shows the *File Series Tool* dialog (Figure 20.1). This tool enables the user to select a subset of a file series, and select a dimension for each the indices in the file name, so that each image is assigned to the correct *z*-plane, time frame, and channel.

The file pattern is shown in the first row in the dialog. The counters in the file name are replaced by menu buttons for selecting

the appropriate dimension for each counter. The options are:

- **Slice:** The range of this counter becomes the *z*-dimension.
- **Time Frame:** The range of this counter becomes the time dimension.
- **Channel:** The range of this counter becomes the channel dimension.
- **Ignore:** the variable is ignored. This is useful to omit e.g. the value of time stamps.
- **The value of the counter** in the selected file: the value of this counter has to match the value in the selected file.

Note that the selection has to be unique, i.e. it is impossible to have ignored variables without having a *Slice*, *Time Frame*, or *Channel* counter.

In the second, third, and fourth row, the range for each of the counters can be defined. A range from 0 to 9 with step size 2 will load the files 0, 2, 4, 6, and 8. Note that the time (in seconds) and *z*-sampling intervals (in nm) are not adapted to the step sizes.

Press the LOAD SELECTION button to load all files in the series into a single image. Before the dialog is closed, the tool will check if all files in the selection are really present in the directory.

Adjusting the Global Color Scheme

Huygens Professional uses a global scheme for coloring the different channels in multi-channel images. These colors can be adjusted through the *Preferences* window via OPTIONS→PREFERENCES...→EDIT GLOBAL COLORS (See Figure 20.2).

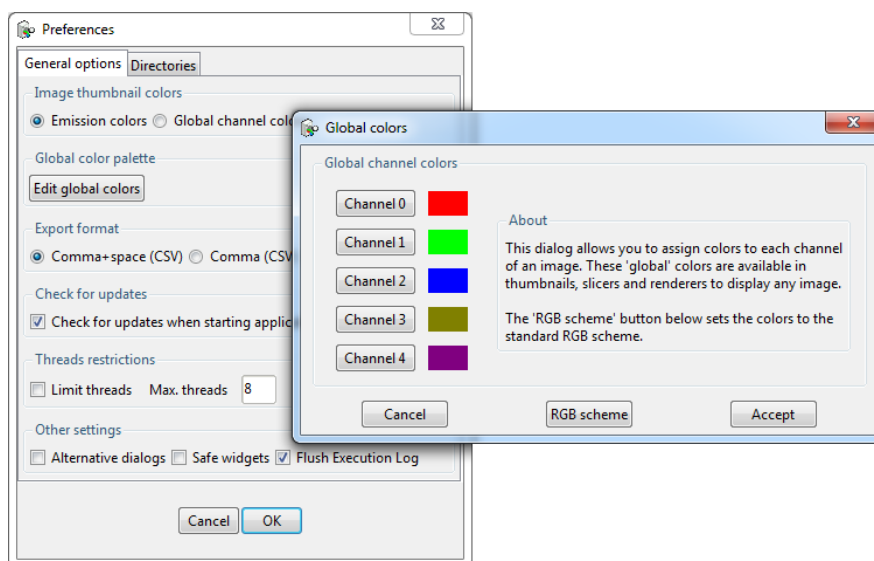


FIGURE 20.2. The global color scheme can be modified through the *Preferences* window.

Hue Selector

The *hue selector* is a component that allows adjustment of the color range in which objects are displayed (See “The Surface Renderer” on page 43 and “The Colocalization Analyzer” on page 95). Objects belonging to different channels can be represented in different hue ranges to make them clearly distinct. The gradual differences inside the selected range make independent objects distinguishable. Also a range can be collapsed to have all objects in a channel displayed with exactly the same color. In Huygens the *hue selector* does appear in two flavors.

Hue Range

This selector allows the adjustment of a *hue* range. The objects on which this selector acts will get a color that lies within this range. The assignment of colors is based on the position of an object or on another parameter.



Hue Range and Saturation

This selector allows the adjustment of a single *hue* value and a *saturation*. The upper triangle defines the color, while the lower triangle sets the saturation for this color; left is white, right is fully saturated.

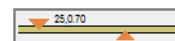


Image Statistics

Right-click on a thumbnail image and select SHOW PARAMETERS from the pop-up menu. This window shows, besides the parameter settings, statistical information of the particular image. Amongst them are the mean, sum, standard deviation, norm, and position of the center of mass.

Setting the Coverslip Position

When there is a mismatch between the refractive index for which the microscope's objective is designed and the actual refractive index of the embedding medium, the shape of the point spread function (PSF) will be distorted due to spherical aberration (See "Refractive Index Mismatch" on page 117). As deeper layers in the specimen are imaged, moving away from the coverslip, this distortion will progressively worsen. To compute the spherical aberration it is necessary to know the distance from the coverslip. Because in many cases the coverslip position does not coincide with the first plane in the data, this position can be set in the microscopic parameter editor. To our knowledge none of the existing microscopic image files record the coverslip position in the meta data.

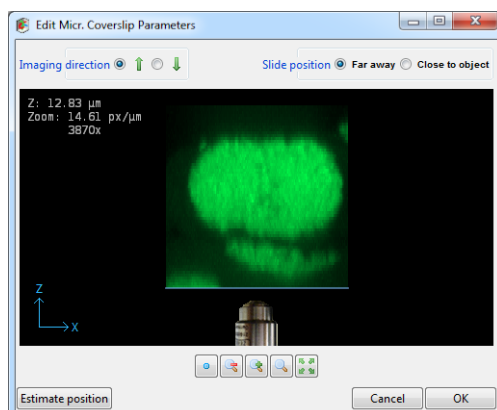



FIGURE 20.3. The coverslip position editor showing an xy MIP of the data along the y -direction. The coverslip position can be adjusted by dragging the blue line. The imaging direction, here *upwards*, is indicated by the position of the objective relative to the data as shown. The z -position shown top-left in the image indicates the distance in μm of the coverslip to the first data plane.

Next to direct numerical input, the coverslip position and imaging direction can be set using a visual editor (Figure 20.3), reachable from the parameter editor by clicking the wrench button ()

Inverted Microscope

The editor shows the coverslip position and imaging direction relative to the data as read from the microscopic file.

In an inverted microscope, with the objective physically below the specimen it is likely that the first xy -plane in the data, corresponding with the lowest location in the xz maximum intensity projection (MIP) on the screen, corresponds with the xy -plane scanned closest to the objective. This situation is shown in REF TO FIG. However, since scan directions and data planes might have been reordered, this match is not

guaranteed. Fortunately, it is often easy to spot the flat side of the object where it adheres to the glass, so the orientation can be verified.

Upright Microscope

In an upright microscope, and a z -scan starting away from the coverslip, the first plane is also likely to be physically the lowest plane. In that case, the imaging direction should be set to downwards and the coverslip position in the top part of the xz MIP projection. However, if the scan started close to the coverslip while storing these first planes first in the data set, the MIP projection will show the data upside down. Consequently, the coverslip position will be in the lower part of the MIP, and the imaging direction is upward.

Slide Position

When the specimen is mounted on the coverslip, the distance from the object to the slide is probably in the range from 50 to 100 μm , outside of the image. In this case, or in the case there is no slide, select *Far away* in the top-right selector.

When the specimen is close to or mounted on the slide, select *Close to object* (). Drag the coverslip to its proper location. When this location is at some distance from the data it might be necessary to zoom out. The image can be dragged by holding down the right mouse button. In terms of imaging quality, when there is a refractive index mismatch between embedding medium and immersion medium, this is not an ideal situation since the light from and to the objective must travel hundreds of wavelengths through the embedding medium, possibly resulting in strong spherical aberration induced bloating of the PSF.

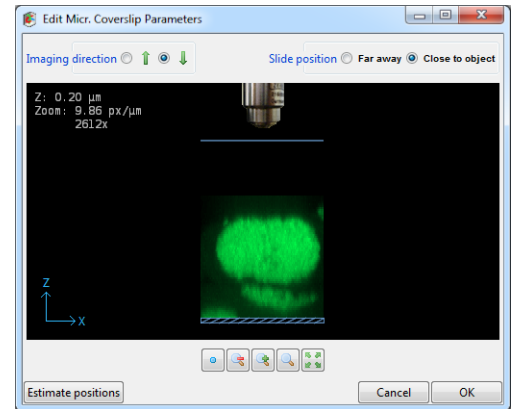


FIGURE 20.4. The Coverslip editor with the slide position set to *Close to object*.

Excitation Beam Overfill Factor

In confocal microscopes, the entry pupil of the microscope objective is illuminated by a laser beam. Usually, laser beams have a Gaussian intensity profile². As a result, the illumination intensity is not constant over the pupil but will decrease towards the edges. Lower edge intensities will lower the effective NA and therefore negatively affect resolution. In most confocal microscopes this is remedied by using a beam width which is significantly larger than the entry pupil, at the cost of loss of excitation power. The ratio between the beam width and the pupil diameter is the excitation beam overfill factor (See Figure 20.5) and is typically in the range from 2 to 4.

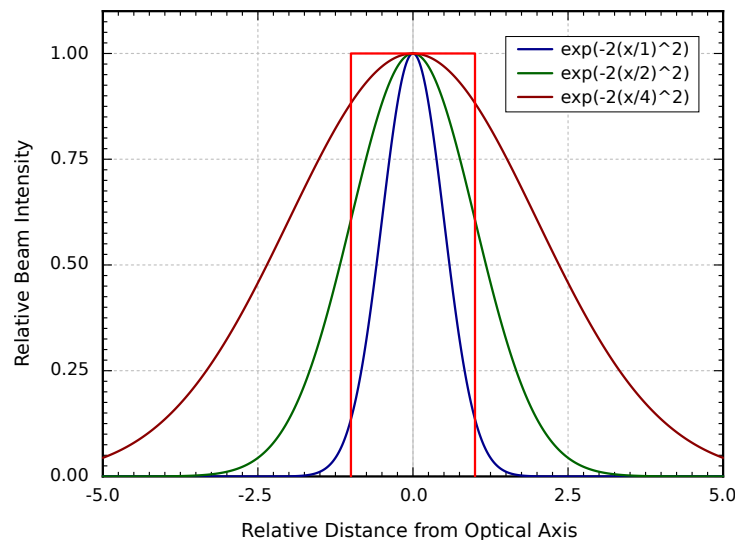


FIGURE 20.5. Lens entry pupil (red), beam profile with overfill factor 1 (blue), 2 (green), and 4 (dark red). At overfill factor 1, the beam intensity is 14 % of the maximum, at overfill factor 2 the edge intensity is 61 % of the maximum.

2. http://en.wikipedia.org/wiki/Gaussian_beam

The overfill factor can be set as a microscopic parameter in Huygens Professional, and is taken into account when computing the point spread function.

Brightfield Images

Brightfield imaging is not a *linear imaging* process. In a linear imaging process the image formation can be described as the linear convolution of the object distribution and the point spread function, hence the name deconvolution for the reverse process. So in principle one cannot apply deconvolution based on linear imaging to non linear imaging modes like brightfield and reflection. One could state that the image formation in these cases *is* linear because it is governed by linear superposition of amplitudes. However, microscopes do not measure light amplitudes but rather intensities, i.e. the absolute squared values of the amplitudes. Taking the absolute square destroys all phase information one would need to effectively apply deconvolution. Fortunately, in the brightfield case the detected light is to a significant degree incoherent. Because in that case there are few phase relations the image formation is largely governed by the addition of intensities, especially if one is dealing with a high contrast image.

In practice one goes about deconvolving brightfield images by inverting them (using TOOLS→INVERT IMAGE image) and processing them further as incoherent fluorescence widefield images. The Tikhonov Miller algorithm was proven to work excellently for brightfield data. This algorithm is available in the Huygens Professional only. With the MLE algorithm one should watch out sharply for interference like patterns (periodic rings and fringes around objects) in the measured image. As a rule these become pronounced in low contrast images. After the deconvolution run a reverse to the original contrast setting is possible.

Support and Contact Information

Addresses and Phone Numbers

Mailing Address	Scientific Volume Imaging B.V. Laapersveld 63 1213 VB Hilversum The Netherlands
Phone	+31 35 6421626
Fax	+31 35 6837971
E-mail	info@svi.nl
URL	http://www.svi.nl

We are directly reachable by phone during office hours (CET) or by e-mail 24/7.

Distributors

A up-to-date list of distributors can be found on our web site³.

SVI Support Wiki

The SVI-wiki⁴ is a rapidly expanding public knowledge resource on 3D microscopy and deconvolution. Based on the WikiWikiWeb principle, it is open to contributions from

3. <http://www.svi.nl/distributors/>

4. <http://support.svi.nl/wiki/>

every visitor. In addition it serves as a support medium for SVI customers and relations to discuss different aspects of the Huygens software.

This is a list of useful starting points in the SVI Support Wiki to learn more about the Huygens software and microscopical imaging in general:

- Information on the parameters describing the imaging conditions (sampling, numerical aperture, pinholes, etc.):
<http://support.svi.nl/wiki/MicroscopicParameters>
- Information on the restoration parameters (signal to noise ratio, background, quality criteria, etc.) used by the deconvolution algorithms:
<http://support.svi.nl/wiki/RestorationParameters>
- A step by step example on how to tune these parameters to achieve the desired restoration results:
<http://support.svi.nl/wiki/DeconvolutionProcedure>
- Important issues regarding image acquisition and restoration (sampling, clipping, etc.):
<http://support.svi.nl/wiki/ImportantFactors>
- Typical acquisition pitfalls (spherical aberration, undersampling, bleaching, etc.):
<http://support.svi.nl/wiki/AcquisitionPitfalls>
- Information on recording beads to measure a PSF:
<http://support.svi.nl/wiki/RecordingBeads>
- Tutorials and detailed information on using the different aspects of the Huygens software (restoration, visualization, analysis, programming, etc.):
<http://support.svi.nl/wiki/Tutorials>
- Uploading images to SVI:
<http://support.svi.nl/wiki/SendImagesToSvi>

Index

A

Address iv, 134
Airy disk 119
Analyze object 73
Anchor 92
Animation
 frame count 43, 48, 52
 frame rate 43, 48, 52
 MIP 43
 Movie Maker 55
 SFP 48
 Surface Renderer 52
AVI files 58

B

Background 18, 104
Backprojected
 pinhole radius 118
 pinhole spacing 118
Base line 117
Batch Processor 23
Beads 111
Black level 16, 117
Bleaching 20, 125
Bounce 60
Bricks 20, 115
Brightfield images 134
Brightness 33

C

Center scene 43, 52
Channel
 Select 15
Clipping 15, 124
Colocalization 83, 101
 coefficients 103
 map 103
Color mode 29, 43, 48
Colors 33, 34, 130
Compress 33

Concurrent tasks 27
Contact SVI 134
Contrast 30, 33
Correlation 83
Coverslip position 132
Cropper 14

D

Deconvolution
 templates 25
Discard object 73
Distributors 134
Dynamic range 117

E

E-mail iv, 134
Emission
 transparency 47
Excitation
 fill factor 133
 transparency 47
Experiment preset 76, 93

F

Fax number iv, 134
File
 Formats 2
 Series 129
Filter Objects 95
Freeware mode 6

G

Gamma 33
Garbage volume 51, 72
Geometrical distortion 123
Global color scheme 130

H

hgsv file 43, 48, 52
Histogram 15, 16, 78

I

- Image
 - Convert 17
- Installation 5
 - Linux 5
 - Mac OS X 5
 - path 8
 - Uninstall 9
 - Windows 5
- Intensity profile 36
- Iso-surface 49
- Iterations 19

K

- Keyframes 56

L

- Leica numbering 129
- License string 6, 127
- Linking 30, 35
- Linux 5, 11

M

- Mac OS X 11
- Mac OS X 5
- Markers 36
- Maximum Intensity Projection 41
- Measure
 - Aspect Ratio 83
 - Colocalization 83, 101
 - Correlation 83
 - Intensity profile 36
 - Length 36, 81
 - Principal Axis 81
 - Sphericity 82
- Memory 9
- MIP 15, 41
- Mouse mode 73, 88
- Movie Maker 55

N

- Nyquist rate 13, 14, 117, 124

O

- Object Analyzer 69, 107
- Object size 46
- Operations Window 16
- Orthogonal Slicer 39
- Overfill factor 133

P

- Parameters
 - Edit 13, 16
 - Microscopic 13
 - Restoration 19
- Phone number iv, 134
- Photodamage 124
- Pinhole
 - radius 118
 - spacing 118
 - Spinning disk 120
- Square 119
- Pipe
 - Mode 90
 - Render 50, 74, 94
- Point spread function 17, 128
- Presets 58, 76, 93
- Principal axis 81
- Processor 9
- PSF 128
 - Asymmetric 123
 - Distiller 111, 113
 - Experimental 126
 - Theoretical 18

Q

- Quality factor 20, 129

R

- Refractive index
 - Mismatch 123
- Region of interest 86
- Render
 - MIP 41
 - pipes 50, 74, 94
 - quality 43, 48
 - SFP 45
 - size 43, 48, 52
 - Surface 49
- Restoration methods 19
- ROI 86

S

- Sampling density 117
- Segmentation 71
- Select area 73
- Select object 73
- SFP 45
- Shape factor 119
- Signal to noise ratio 115
- Slicer 12
- Slide position 132
- SNR 19, 115
- Spherical aberration 123, 126, 132
- Sphericity 82
- Statistics 75, 131
- Storyboard 57
- Support 134
- Surface Renderer 49
- System ID 6, 128
- System requirements 9

T

- Table style 48
- Task 26
- Templates
 - Deconvolution 25
 - Microscopic 25
 - Visualization 43, 48, 52
- Threads per task 27
- Threshold 41, 46, 50, 71
- Timeline 58

Total internal reflection 124
Transitions 57
Transparency 51
Transparency depth 52
Twin Slicer 31

U

Undersampling 124
Uninstall 9
Updates 7, 8

W

Windows 5, 11

