

Huygens Essential

USER GUIDE FOR VERSION 19.10



Scientific Volume Imaging B.V.

Huygens Essential

USER GUIDE FOR VERSION 19.10



Scientific Volume Imaging B.V.

Copyright ©1995-2019 by Scientific Volume Imaging B.V. for Huygens version 19.10
All right reserved.
This manual was created on November 21, 2019.

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	https://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

1	Introduction	1
2	Installation	3
	Microsoft Windows	3
	Mac OS X	3
	Linux (Debian)	3
	Linux (RPM)	3
	After the Installation	4
	The License String	4
	Updating the Software	6
	Removing the Software	7
	System Requirements for Huygens Essential	7
	Support on Installation	8
	User preferences	8
	Measuring the software usage	10
3	Deconvolution Wizard	11
	The Processing Stages in the Wizard	11
	Loading an Image	12
	Verifying Microscopy Parameters	13
	Starting the Deconvolution Wizard	14
	Using a Measured PSF	16
	The Intelligent Cropper	16
	Stabilizing STED data	17
	The Image Histogram	18

Estimating the Average Background	19
The Deconvolution Stage	20
Finishing or Restarting a Deconvolution Run	22
The Final Stage	22
The Comparison View	22
Multi-channel Images	23
Z-drift Correcting for Time Series	24
Saving the Result	24
4 Deconvolution Express	25
Why Deconvolution Express	25
Verifying the image parameters	25
Starting Deconvolution Express	25
Unsupervised Profiles	26
5 Batch Express	29
Difference between Batch Express and other Huygens deconvolution tools	29
The Batch Express window	30
Working with the Batch Express	30
Batch Express preview	31
6 The PSF Distiller	33
Beads suited for PSF Distillation	33
The PSF Distiller Window	34
The Processing Stages in the Wizard	35
Averaging Stage	35
Distillation Stage	36
Finalizing the result	37
Full-width at half-maximum estimator	37
7 The Batch Processor	39
The Batch Processor	39
Usage	41

Menus	44
8 Huygens STED Deconvolution	47
Introduction	47
STED principle	47
STED parameters	48
Example settings	49
Estimating STED parameters	49
STED Deconvolution	51
Summary: deconvolution procedure	52
9 Huygens Array Detector Deconvolution	53
Array Detector and Zeiss Airyscan images	53
Introduction	53
The Zeiss Airyscan microscope	54
Loading an Array Detector Image	54
Visualizing an Array Detector Image	54
Array Detector image parameters	55
Deconvolving Array Detector images	56
10 Light Sheet Deconvolution	59
LSFM image restoration	59
LSFM Parameters	60
LSFM deconvolution	63
11 Light Sheet Fusion & Deconvolution Wizard (Legacy)	65
Starting the Fusion and Deconvolution Wizard	65
Options Menu	65
Selecting input images	66
Deconvolution settings	67
Transformation and resampling parameters	67
Setting the rotation parameters	68
Channels and Time frames	69

Starting the deconvolution and the fusion	69
12 The Interactive Fuser	71
Opening The Interactive Fuser & Decon Window	72
A global overview of The Interactive Fuser window	73
Preprocessing settings	74
Input view specific settings	75
Interactive alignment	76
Registration & fusion settings	77
Starting the registration and fusion process	81
13 The Twin Slicer	83
Using the Slicer in Basic Mode	84
Using the Slicer in Advanced Mode	86
Measurement	90
14 The Orthogonal Slicer	91
Orthogonal views	92
The Crosshair Cursor	92
Visualization parameters	92
Measurements	93
Auto-Zoom	93
Display Options	93
15 The MIP Renderer	95
Basic Usage	96
Options menu	97
Simple Animations	97
16 The SFP Renderer	99
Basic Usage	100
Advanced Usage	102
Simple Animations	106

17 The Surface Renderer	107
Basic Usage	108
Advanced Usage	109
Simple Animations	111
18 The Movie Maker	113
An Overview	113
Creating and Adjusting Keyframes	114
Using the Storyboard	115
Working with Movie Projects	116
Using the Timeline	116
Advanced Topics	117
19 The Gallery	119
20 Tile Stitching & Deconvolution Wizard	121
Stitching & Deconvolution	121
Vignetting and Shading correction	121
Starting the Wizard	122
21 The Hot & Cold Pixel Remover	125
Why correct for hot/cold pixels?	125
Hot Pixel Remover window	125
Detection Settings	126
Saving the correction	127
Using a Mask	128
Adding hot or cold pixel removal to a deconvolution template	128
22 The Crosstalk Corrector	129
What is crosstalk?	129
Starting the Crosstalk Corrector	129
Adjust the crosstalk coefficient	131
23 Bleaching Corrector	133

What is Bleaching?	133
Starting the Bleaching Corrector	133
Estimate Bleaching	134
Adjusting Bleaching Factors	135
Correct Bleaching	135
Working with templates	135
24 The Chromatic Aberration Corrector	137
Causes of chromatic aberration	137
Starting the Chromatic Aberration Corrector	138
Estimation of the chromatic Aberration	138
Visualization of the chromatic aberration	139
Editing the chromatic shift vectors	140
Working with templates	141
25 The Object Stabilizer	143
The Object Stabilizer Window	143
Preprocessing	143
Stabilization of 3D Time Series	144
Alignment of Slices in 3D Stacks	147
Stabilization Settings	149
Cropping the Result	149
Stabilization Summary	150
26 The Object Tracker	151
Introduction	151
The Object Tracker Wizard	151
The Track Analyzer	154
27 Introduction to the Object Analyzer	159
Starting the Object Analyzer	159
Segmenting the Objects: Setting the Threshold	161
Interaction with the Objects	164

Render Pipes	164
Object Statistics	165
Storing your Results	168
Further Reading	168
28 Object Analyzer Geometry Measurements	169
Iso-surface	169
Principal Axis	169
Length and width	169
Sphericity	170
Aspect Ratio	171
More Parameters and Filtering	171
29 Object Analyzer Component Reference	173
Main window components	174
30 The Colocalization Analyzer	187
Overview of the modes	187
How to use the Colocalization Analyzer	189
Time frame & channel selection and the 2D Histogram	190
Threshold settings	190
Region of interest selection	192
RBNCC	193
Computation and results	194
31 Huygens multi-GPU acceleration	197
Huygens GPU options	197
Huygens and GPU acceleration	198
How to prepare GPU drivers and CUDA® for Huygens	198
GPU Check List	199
GPU status Window	199
Multi-GPU support	200
32 Huygens for Multiple Users	203

Remote Display	203
Floating License(s)	204
Huygens web accessibility with HRM	205
33 Establishing Image Parameters	207
Image Size	207
Signal to Noise Ratio	207
Black Level	209
Sampling Density	209
Computing the Backprojected Pinhole Radius and Distance	210
34 Improving Image Quality	213
Data Acquisition Pitfalls	213
Deconvolution Improvements	216
35 Appendix	217
The Point Spread Function	217
Quality Factor	217
File Series	217
Adjusting the Global Color Scheme	219
Hue Selector	220
Setting the Coverslip Position	221
Excitation Beam Overfill factor	222
Brightfield Images	223
2D Histogram	223
36 Support and Contact Information	225
Contact Information	225
Support	225
License String Details	226

Chapter 1

Introduction

Huygens Essential is an image processing software package tailored for restoration, visualization and analysis of microscopic images. Its wizard driven user interface guides through the process of deconvolving images from light microscopes. Huygens Essential is able to deconvolve a wide variety of images ranging from 2D widefield images to 4D multi-channel multi-photon, multi detector confocal time series. To facilitate comparison of raw and deconvolved data or results from different deconvolution runs Huygens Essential is equipped with a dual 4D slicer tool. Also 3D images and animations can be rendered with its powerful visualization tools. Post-restoration analysis is possible using the interactive analysis tools.

Based on the same image processing engine (the compute engine) as Huygens Professional, Huygens Essential combines the quality and speed of the algorithms available in Huygens Professional with the ease of use of a wizard driven intelligent user interface complemented by a versatile and intuitive batch processor.

Huygens Essential is build on cross-platform technology. It is available on Microsoft Windows 7, Windows 8, and Windows 10 (64 bit) , Linux (64 bit), and Mac OS X Mountain Lion, Mavericks, Yosemite, El Capitan (X11, XQuartz).

Chapter 2

Installation

Huygens Essential can be downloaded from the SVI website¹.

Microsoft Windows

Double click on the Huygens installer executable, e.g. `huygens-19100p0_x86_64.exe`. During installation the directory `C:\Program Files\SVI` will be created by default, and four Huygens icons appear on the desktop.

Mac OS X

Double click the package file, for instance `huygens-19.10.0-p0-Lion-i386-x86_64.pkg`. 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. Since Huygens is only available for Mac as an X11 application, installation of XQuartz² is required.

Linux (Debian)

Debian packages are commonly used by Ubuntu and other Debian-based Linux distributions. Double click the package file, e.g. `huygens_19.10.0-p0_amd64.deb`, and follow the steps in the package manager. To install the package from a terminal, use the command:

```
dpkg -i huygens_19.10.0-p0_amd64.deb
```

Linux (RPM)

RPM (RedHat Package Manager) packages are used natively by RedHat, Fedora, SUSE, and other RPM-based Linux distributions. Double click the package file, e.g. `huygens-19.10.0-p0.x86.rpm`, and follow the steps in the package manager, or install the package from the command line:

```
rpm -ivh --force huygens-19.10.0-p0.x86.rpm
```

¹<https://www.svi.nl/Download>

²<https://www.svi.nl/MacOSX>

After the Installation

After starting Huygens Essential for the first time, a message will be shown informing that no license is available yet. Without a license, Huygens Essential will run in Freeware mode which gives access to the basic renderers but not to the image restoration or analysis functions. The list of free available functionalities can be found on the website ³. A System ID is required for generating a license. The System ID can be obtained by pressing the REQUEST TEST LICENSE button after starting Huygens Essential. It can also be found in the HELP → ABOUT menu. The next section explains how to obtain and install a license string.

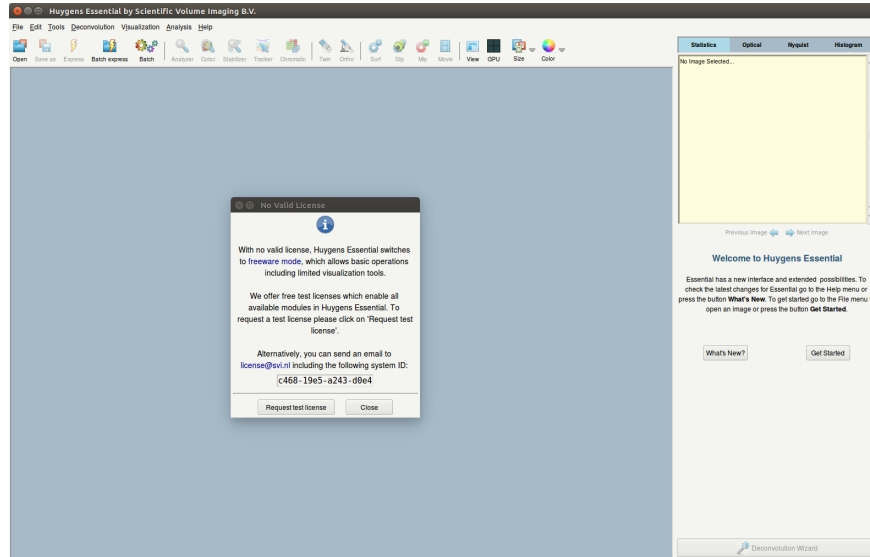


Figure 2.1: The startup window of Huygens Essential. If no license string is installed the software runs in *freeware* mode

The License String

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

```
HuEss-19.10-wcnps-d-tvAC-emnps-eom2019Dec31-e7b7c623393d708e-
user@domain.com-4fce0dbe86e8ca4344dd
```

At startup Huygens Essential searches for a license file `huygensLicense`, which may contain a license. A license string is provided by SVI via e-mail. The procedure for installing a license string is the same for all platforms.

³<https://www.svi.nl/FreeWare>

Obtaining a License String

If Huygens Essential 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). Please send it to sales@svi.nl, a license string will be provided.

To prevent any typing error please use the COPY button to save the ID to the clipboard. Please paste it into an email to sales@svi.nl, e.g., using EDIT → PASTE in your email application.

To see whether there is a new version of Huygens Essential available from the SVI company server, the dialog box also contains a button *Check for updates*.

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 Essential and go to HELP → LICENSE MANAGER: 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 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 start but remain in Freeware mode. When the license is successfully added to the license file the message “*Added license successfully*” will appear.

Restart Huygens Essential to activate the new license.

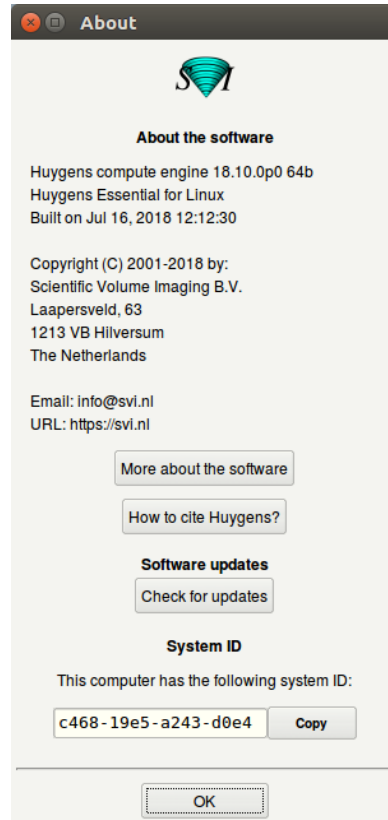


Figure 2.2: The HELP → ABOUT window. The system ID is shown at the bottom.

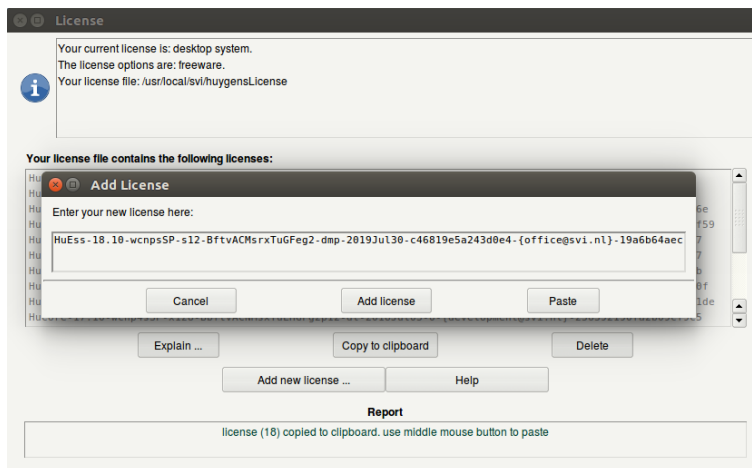


Figure 2.3: The license window allows to add, delete and troubleshoot licenses.

Location of the License File

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

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

Table 2.1: The default installation paths per platform. The path name on Mac OS X depends on where the software is installed. This is a typical example.

On 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

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. To have a license string explained, select it in the license window (HELP → LICENSE OVERVIEW) and press the EXPLAIN LICENSE button. All details for the current license will be listed (Figure 2.4). This information can be used to troubleshoot licensing problems.

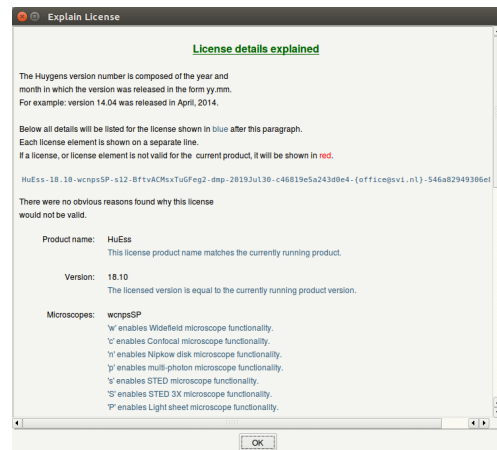


Figure 2.4: The *Explain License* window lists all license details.

Updating the Software

Upon startup, Huygens can check for available updates, if the system is connected to the Internet and the *Check update* option is selected in the GENERAL OPTIONS window accessible via EDIT → PREFERENCES. A pop-up window will appear when a newer version is available. The HELP → ABOUT window offers the option to manually check for updates.

SVI’s website can also be consulted for updates. Twice a year around April and October new releases will become available. Download the new version from the SVI website⁴. Proceed with the installation as explained above.

Do not uninstall the old version as this will delete the license string. The new version will (by default) automatically replace the old Huygens version. However, for the MAC OS X, we advise to first uninstall the software and then reinstall a newer version. Please read the next subsection “Removing the Software” carefully for more information.

⁴<https://www.svi.nl/Download>

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. See “Location of the License File” on page 6 to read where the license string is stored. Table 2.2 on page 7 shows the procedure to uninstall the software for each 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 huygens and uninstall it. This could also be handled with the command line; type <code>dpkg -r huygens</code> to uninstall a Debian package or <code>rpm -e huygens</code> to uninstall an RPM package.
Mac OS X	Drag the installed version to the waste basket.

Table 2.2: The uninstallation procedure per platform.

System Requirements for Huygens Essential

The following tables show the requirements for Windows, Mac OS X, and Linux.

Operating system	Huygens runs on Microsoft Windows 7, Windows 8, and Windows 10 (64bit)
Processor	1 GHz or faster 64-bit processor
Memory	8GB or more.
Graphics card	Any fairly modern card will do. For GPU acceleration a CUDA-enabled NVIDIA® graphics card is needed (see Chapter 31).

Table 2.3: System requirements for Microsoft Windows.

Operating system	Huygens runs on Mac OS X Mountain Lion, Mavericks, Yosemite, El Capitan, (High) Sierra, Mojave (X11, XQuartz) ⁶ .
Processor	Intel 64 bit.
Memory	8GB or more.
Graphics card	Any fairly modern card will do.

Table 2.4: System requirements for Mac OS X.

Operating system	Most popular distributions like Ubuntu, RedHat, Fedora, and SuSE are supported. Linux (64 bit)
Processor	1 GHz or faster 64-bit processor
Memory	8GB or more.
Graphics card	Any fairly modern card will do. For GPU acceleration a CUDA-enabled NVIDIA® graphics card is needed (see Chapter 31).

Table 2.5: System requirements for Linux.

⁵OS X 10.7 or higher with X11 is required for full 64 bit capabilities.

⁶OS X 10.7 or higher with X11 is required for full 64 bit capabilities.

Support on Installation

If any problem is encountered when installing either the program or the licenses and these guidelines fall short to fix it, please search the SVI Wiki⁷ or contact SVI (See "Contact Information" on page 225).

User preferences

Under EDIT → PREFERENCES, preferences can be set to configure the usability of the software (see Figure 2.5). Also, the default directories for the File Open and Save dialogs, and for the *Image Feeder* and *Batch Express* can be defined here. More details can be seen in the tooltips that are shown when hovering over the preferences with the mouse cursor. The preferences are grouped under two main tabs:

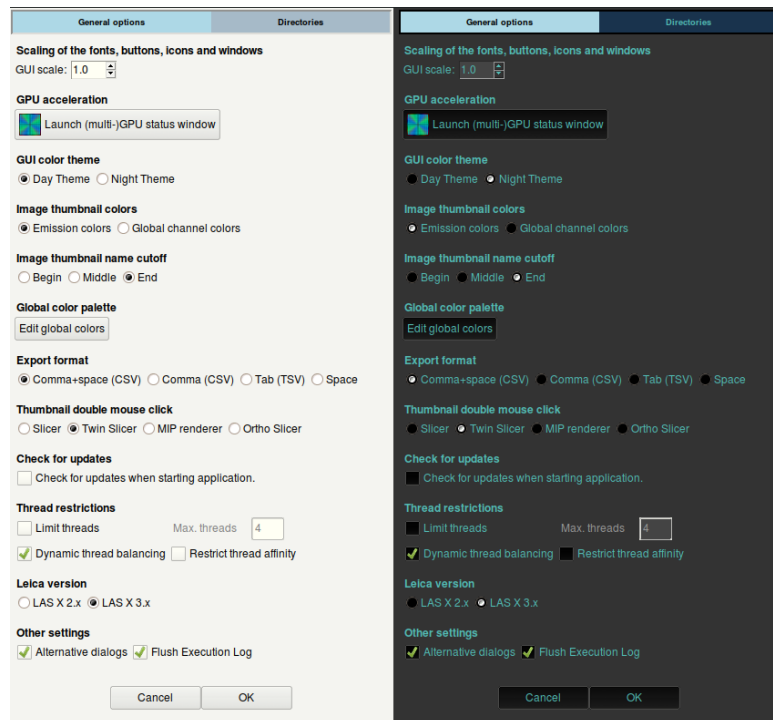


Figure 2.5: The Preference window, available under the EDIT menu, allows to change the default settings at startup. For example, in this window the appearance of all the Huygens windows could be changed from a day (left side) to a night (right side) theme.

General options

- **Scaling for 4K monitors.** Huygens is able to scale fonts, buttons, windows and icons for 4K monitors, which can have vertically more than 3000 pixels. The first entry box located under the GENERAL OPTIONS tab allows to manually adjust the scaling value. After changing the scaling setting, Huygens needs to be restarted for the changes to take effect.
- **GPU acceleration.** Enable/disable GPU processing. Select which GPU cards to use when the system has several cards. Define the default GPU card.
- **GUI color theme.** With this option, the appearance of the Huygens user interface can be changed from a Day to a Night theme, and vice versa. For optimal visibility in this manual, only Day theme versions of Huygens are shown.

⁷https://www.svi.nl/FAQ#Installation_questions

- **Image thumbnail colors and Global color palette.** These settings determine the first color mode of the thumbnail when a new image is opened in Huygens. The global colors can be adjusted with the Color picker tool. See also “Adjusting the Global Color Scheme” on page 219.
- **Image thumbnail name line break.** An image name can be too long to fit properly under a thumbnail and needs to be cropped. Select which part of the name will be discarded.
- **Export format.** This option sets the format for exporting ASCII data.
- **Thumbnail double mouse click.** Option to select which visualization tool to open upon double-clicking an image thumbnail in the main Huygens window.
- **Check for updates.** Upon start, Huygens can check whether an updated version is available. See also See “Updating the Software” on page 6.
- **Thread restrictions.** The number of threads used by Huygens can be customized here. By enabling dynamic thread balancing, Huygens will take threads conservatively. With the option switched off, Huygens will take the maximum amount of threads, resulting in maximum performance but little computing power for other applications.
- **Other settings.** Alternative dialogs other than those specific for the operating system can be shown in Huygens when checked. By flushing the execution log, entries are not kept in a temporary buffer but immediately saved, ensuring that all information from the session is preserved if debugging is needed.
- **Leica version.** If the Huygens license includes the option to allow an exchange of images with Leica LASX, then the proper LASX version needs to be defined.

Directories

- **Program history.** Huygens can keep the history of a certain number of images for a specific number of days. The history is available via the FILE → OPEN RECENT.
- **Disk locations.** Locations of important folders for loading/saving images, templates, etc.
- **The Image Feeder.** The image feeder is a functionality that can automatically open images in Huygens if these images are saved in a specified Image Feeder directory. To have this functionality active, select *Open Images*. It is also possible to open files that are already present within the Image Feeder directory at the time Huygens is starting. With Huygens version 18.04 the new option *Batch Express* was introduced, which can automatically deconvolve all images from a Watch folder using unsupervised profiles, and save the resulting images in an Output folder. The locations of both of these folders at startup can be specified here. It is also possible to change these locations from within the *Batch Express* window (see Chapter 5 on page 29).
- **Default Strategy.** *Batch Express* offers different unsupervised profiles for deconvolving images in the queue. The type of profile that is applied by default can be selected in this window. It is also possible to change the profile from within *Batch Express* window.

Measuring the software usage

The usage of the Huygens basics (e.g. Essential, Professional) software is logged every time the software is launched. In addition, the usage of those tools that require a license-flag in the license string are saved. This usage can be viewed with the Usage Report window, see Figure 2.6

Go to `HELP → USAGE REPORT` and wait for the software to analyze the data. Once the analysis is done the *General information* tab shows the total usage of the software. The *Detailed information* tab shows the usage

per tool. The usage can be filtered by product, user, version and time. Simply select one of the options in the menu-buttons and the data will be re-analyzed. This usage report can be saved through `FILE → SAVE USAGE` as a `.csv` file.

The usage is written to a log file located in the installation directory. When the software is updated and/or uninstalled, the usage logs are kept.

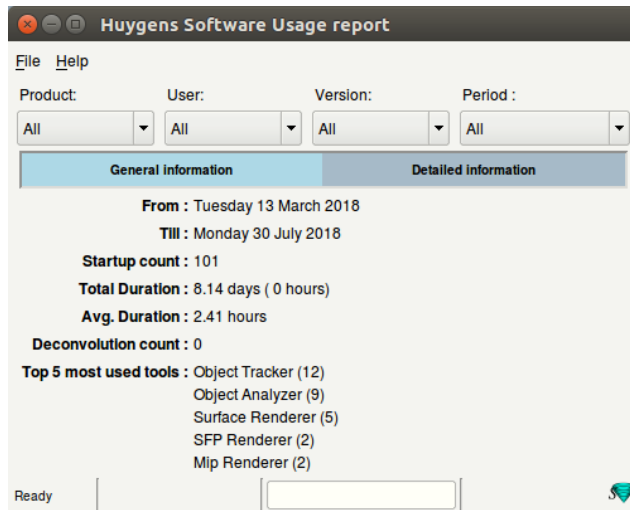


Figure 2.6: The usage report window.

Chapter 3

Deconvolution Wizard

The Processing Stages in the Wizard

The deconvolution wizard guides you through the process of microscopy image deconvolution in several stages. In Huygens Essential one needs to verify the microscopic parameters before continuing with deconvolution. This is done with the Microscopic Parameter Editor.

Each stage is composed of one or more tasks. While proceeding, each stage is briefly described in the bottom-left *Help* window tab. The stage progress is indicated in the *Wizard* pane, below the *Help* tab. Additional information can be found in the online help (HELP → ONLINE HELP) as well as by clicking on the highlighted help questions.

The following steps and stages are to be followed:

- Loading an image.
- Verifying microscopy parameters.
- Start Stage: here, an overview is given of the most important parameters for ensuring optimal deconvolution. You can continue with using the wizard or applying a deconvolution template.
- Loading a measured PSF that has been distilled from bead images.
- Preprocessing Stage: this stage goes through all preprocessing steps.
 - Cropping stage
 - Stabilization (STED data)
 - Select channel
 - Inspecting the image histogram.
 - Background Estimation
 - Selecting Deconvolution algorithm
- Deconvolution Setup:
 - Setting the final deconvolution parameters.
 - Run the deconvolution.
 - Select final result
- Post-processing Stage: correct for Z-drift if any and accept the result.
- End Stage: Option to Restart, Save deconvolution template and Finish.

The next sections of this chapter will explain the wizard stages in detail.

Loading an Image

Select FILE → OPEN or the OPEN icon in the Huygens main window to open the file dialog, browse to the directory where the images are stored, and select the image, for example faba.h5. The demo image (faba.h5) is placed in the Images subdirectory of the installation path (see Table 2.1 on page 6).

Most file formats from microscope vendors are supported, but some of them require a special option in the license to be read. See the SVI support Wiki¹ for updated information.

If the file is opened successfully, select its corresponding thumbnail and check whether the image dimensions are correct by looking under the tab STATISTICS located at the top-right of the main window. Some tools, such as those to convert the image dimensions, are described in more detail in the next subsections.

Converting a Dataset

Before applying deconvolution, it is important that the image dimensions are correctly defined. Dimensions can be changed with the conversion tools listed under the TOOLS menu. For example, a 3D stack can be converted here into a 3D time series to a 3D time series using TOOLS → CONVERT → CONVERT XYZ TO XYZT option, or vice versa, or a 3D stack can be converted into a time series of 2D images (TOOLS → CONVERT → XYZ TO XYT or vice versa). In case of a refractive index mismatch, it is important to set the orientation of the image correctly with respect to the imaging direction and also to the coverslip position (TOOLS → MIRROR). Read more at “Setting the Coverslip Position” on page 221.

Time Series

A time series is a sequence of images recorded at uniform time intervals. Each recorded image is a time frame. Huygens Essential is capable of deconvolving 2D-time or 3D-time data automatically with the *Time option* (see Table 36.2 on page 228). There are some tools that are intended only for time series, such as the widefield *bleaching corrector* or the *z-drift corrector*.

Adapting the image

In the TOOLS menu, you can find *Convert brightfield* image, which is required for the processing of brightfield images (See “Brightfield Images” on page 223.). A crop tool is also available, but its use is recommended only after properly tuning the image parameters. This tool will be explained at a later stage in this manual.

¹<https://www.svi.nl/FileFormats>

Verifying Microscopy Parameters

In Huygens Essential it is mandatory to verify your microscopy parameters before deconvolution. If they are not verified yet and you start the Deconvolution Wizard, you will be prompted to do so using the Parameter Editor. The editor can be found by right-clicking on the image thumbnail and selecting MICROSCOPIC PARAMETERS. See also Figure 3.1 on page 14.

Parameter	Explanation
Microscope type	Select from Widefield, Confocal, Array detector Confocal, Rescan Confocal, Spinning disk, 4Pi, STED, Light sheet, or Unspecified. For a multiphoton microscope, set the excitation photon count higher than 1. Select Widefield if a non-descanned detector was used, and Confocal in case a pinhole of less than a few Airy disc units was used.
Numerical aperture	The NA of the objective lens.
Objective quality	Select from perfect, poor, or something in between. If you are not sure about this setting, leave it to Good.
Coverslip position	The position of the glass interface relative to the nearest slice of the stack. See also page 221
Imaging direction	Select from upward or downward. Upward means that the objective lens is closest to the bottom slice in the stack. Upward stands for an inverted microscope where the objective points upward.
Backprojected pinhole radius	The radius (in nm) of the pinholes in the spinning disk as it appears in the specimen plane. This is the physical pinhole radius divided by the total magnification of the detection system.
Backprojected pinhole spacing	The distance (in μm) between the pinholes in the spinning disk as it appears in the specimen plane. This is the physical pinhole distance divided by the total magnification of the detection system.
Lens refractive index	The RI of the immersion medium for the objective lens.
Medium refractive index	The RI of the specimen embedding medium.
Excitation wavelength	The wavelength (in nm) of the excitation light (usually a laser line).
Emission wavelength	The wavelength (in nm) of the emitted light.
Excitation photon count	Number of photons used in multi-photon microscopy.
Excitation fill factor	The width of the beam relative to the aperture. The default value is 2, meaning that the aperture has a diameter of 2σ , where σ is the standard deviation of the Gaussian distribution in the beam.
STED-specific parameter	See “STED parameters” on page 48.
Light Sheet parameters	See “LSFM Parameters” on page 60.
Deviating microscope type	The deviating microscope type for the current channel.

Table 3.1: Optical parameters explained.

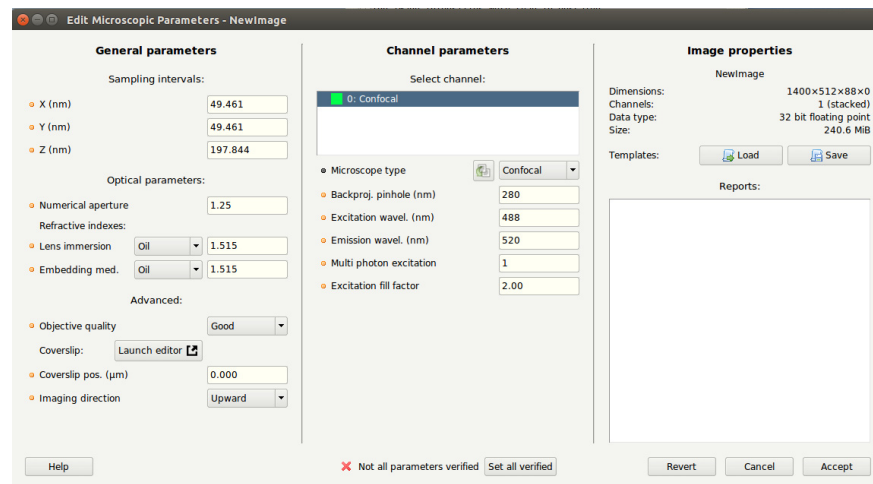


Figure 3.1: The microscopy parameter editor. This window can be opened by right-clicking on the thumbnail and selecting MICROSCOPIC PARAMETERS

Once the proper parameters have been set and verified, they can be saved to a Huygens microscopy template file (.hgsm). Templates can be applied at the start of the Parameter Wizard, in the Parameter Editor, or in the Batch Processor.

The LOAD button allows the selection of a template from a list of template files. The Huygens common templates directory is named Templates, and resides in the Huygens installation directory (See Table 2.1 on page 6). The user's personal templates directory is called SVI/Templates and can be found in the user's home directory. User home directories are usually located in C:\Users on Windows. On Mac OS X they are usually in /Users and on Linux in /home.

Starting the Deconvolution Wizard

When you have verified all the microscopy parameters, you can continue with the *Deconvolution Wizard* by pressing the DECONVOLUTION WIZARD button (on the right bottom corner of the main window), by selecting the wizard icon in the main task bar, or by selecting the wizard from the DECONVOLUTION menu. You can also start the wizard by using the right mouse button on a thumbnail and selecting the DECONVOLUTION WIZARD option from the menu that appears. If the wizard was started already before you verified the microscopy parameters, Huygens will automatically continue with the wizard.

If you prefer to generate a point spread function from measured beads first, select PSF DISTILLER WIZARD in the DECONVOLUTION menu (See Chapter 6 “The PSF Distiller” on page 33). A special license is needed for activating the PSF Distiller option.

The first window of the *Deconvolution Wizard* shows a summary of the most important parameters for ensuring optimal deconvolution. Review them carefully before proceeding if you have not done so already (see Figure 3.2). You can continue with using the deconvolution wizard or you can execute a deconvolution template previously made.

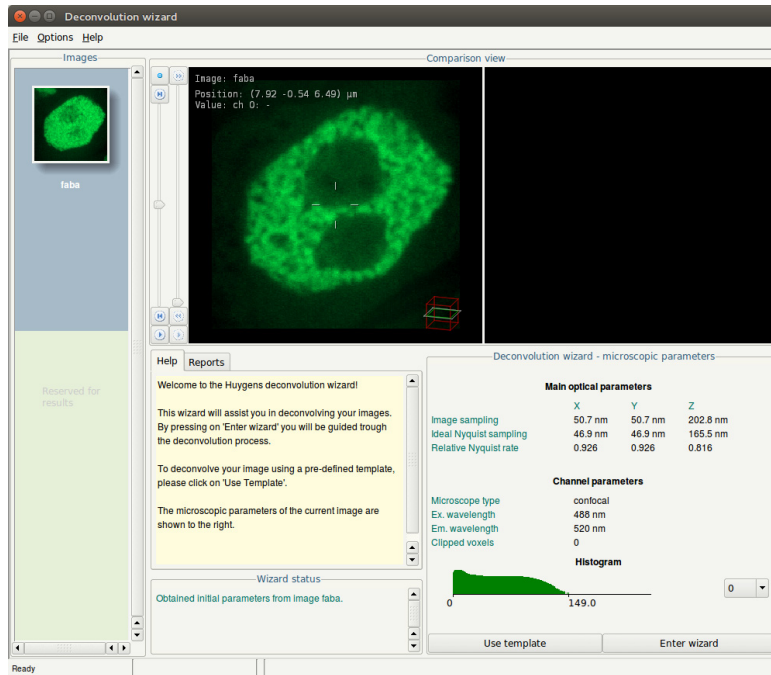


Figure 3.2: Starting window of the *Deconvolution Wizard* with a summary of the most essential parameters.

The latter option will show a list of available templates and a button to load a template file. Below the list of files is shown whether the template considers an experimental or theoretical PSF (Figure 3.3). You can overrule this by using the button **OVERRIDE PSF FROM TEMPLATE**. A next screen, will give you the option then to select an experimental PSF image file.

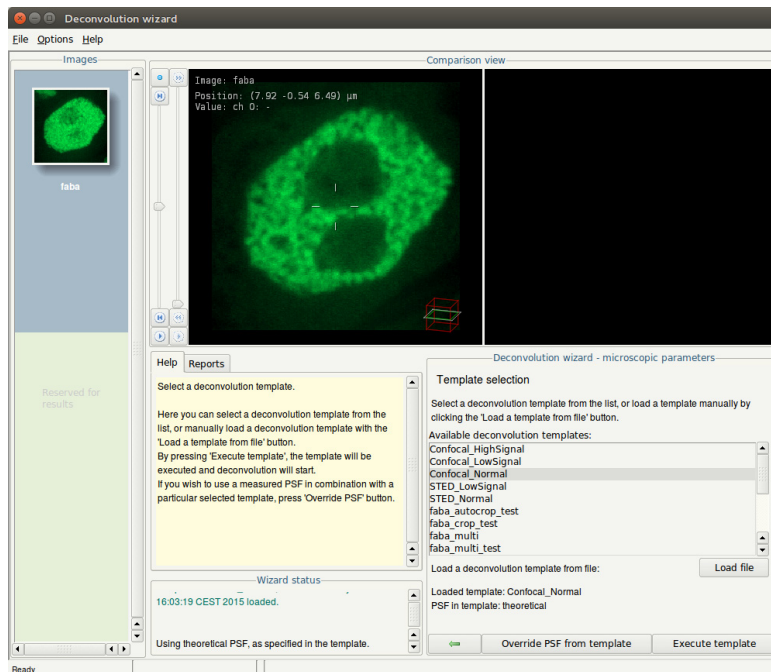


Figure 3.3: Deconvolution templates can be selected from the list or can be loaded. A template specifies if an experimental or theoretical PSF is used, but this can be overruled with the button **OVERRIDE PSF FROM TEMPLATE**.

Using a Measured PSF

Measured PSF's improve the deconvolution results and may also serve as a quality test for the microscope. If the measured PSF contains less channels than the image, a theoretical PSF will be generated for the channels that have no available PSF. See Chapter 6 "The PSF Distiller" on page 33 and "The Point Spread Function" on page 217 for more information.

A previously created measured Point Spread Function (PSF) can be loaded from a file or directly loaded from the Huygens Essential main window. A measured PSF should only be used for deconvolution if the image and the bead(s) were recorded with the same microscope at the same parameter settings. The remaining steps in the wizard will be skipped when selecting the button EXECUTE TEMPLATE.

The Intelligent Cropper

The time needed to deconvolve an image increases more than proportionally with its volume. Therefore, the deconvolution can be accelerated considerably by *cropping* the image.

Huygens Essential and Professional are equipped with an intelligent cropper which automatically surveys the image to find a reasonable proposal for the crop region (See Figure 3.4). In computing this initial proposal, the microscopical parameters are taken into account, making sure that cropping will not have a negative impact on the deconvolution result. Because the survey depends on accurate microscopical parameters it is recommended to use the cropper as a final step in the preprocessing stage (press YES when the wizard asks to launch the cropper), but it can also be launched from outside the wizard through the menu TOOLS → CROP.

Cropping in X, Y, and Z.

The borders of the proposed cropping region are indicated by a colored contour. The initial position is computed from the image content and the microscopic parameters at launch time of the cropper.

The three views shown are *maximum intensity projections* (MIPs) along the main axes. For more information about these projections read Chapter 15 "The MIP Renderer" on page 95. By default the entire volume (including all time frames) is projected. The red, yellow, and blue triangles can be dragged to restrict the projected volume.

The cropper allows manual adjustment of the proposed crop region. To adjust the crop region, drag the corners or sides of the cropbox to the desired position or use the entries in the *Specifications* panel. To crop the original image press the CROP button. To create a new cropped image, press the EXTRACT button. The EXTRACT button is not visible if the cropper is used within the deconvolution wizard. Do not crop the object too tightly, because that would remove blur information relevant for deconvolution, and background area for the efficient estimation of the background level.

Cropping in Time and removing channels

Select the from and to frame under *Specifications* to crop time series. There are entries for the time-selection, if the image is a time series. Only a continuous range can be selected to crop in time series.

All available channels are listed under CHANNELS in the *Operation Window* and can be removed individually. By using the right-click on a channel check box, all channels can be turned on or off, except for the selected channel. In the main window you can separate the image into individual channels using SPLIT CHANNELS, which can be found under the TOOLS menu or as a taskbar icon.




Figure 3.4: The crop tool.

Cropper Customizations and Options

The cropper is extended with extra visualization options available under OPTIONS. LOCK RATIO fixes the aspect ratio of the cropping box for each projection separately. PIXEL PROPERTIES shows the pixel intensities and position of the mouse if you hover over the image. The intensities, however, are only shown when one time frame and one channel is selected. If the PREVIEW box is checked, a small renderer shows the image after cropping.

Stabilizing STED data

STED image acquisition is often subjective to drift. Thus, it is strongly recommended to stabilize the STED images before deconvolving them. This can be easily done within the *Deconvolution Wizard*, after the cropping stage.

The option to stabilize the data along the z direction is present in the *Deconvolution Wizard* with STED data. Whether stabilization will actually be applied depends on the percentage level of STED 3X. Stabilization will automatically correct the raw data for drifts and misalignments along the z direction. Deconvolution will be performed on the corrected, stabilized dataset. Once run, the stabilization cannot be undone.

This stage can be skipped at the Deconvolution Wizard if the image has already been stabilized by other means like for example, the Huygens Object Stabilizer.

To stabilize the image in the Deconvolution Wizard just click on the “AUTO STABILIZE” button at the “STABILIZATION” stage (See Figure 3.5).

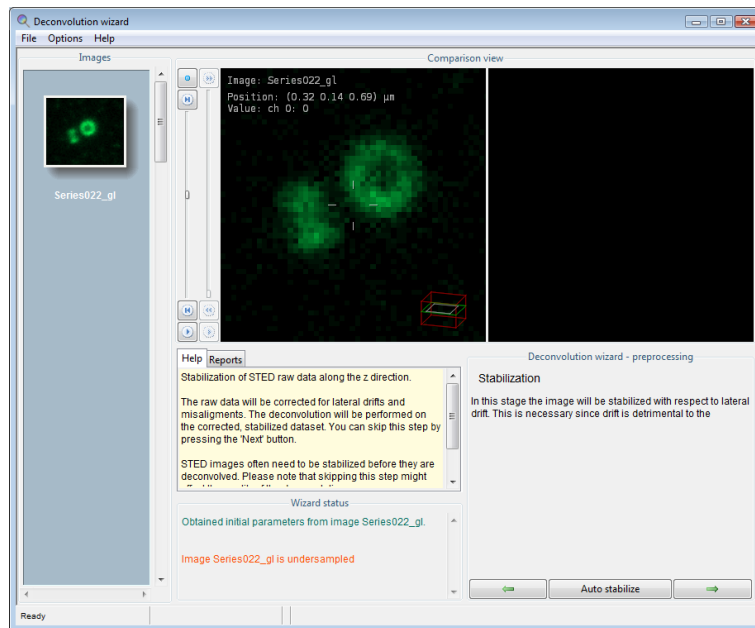


Figure 3.5: Stabilization stage for STED data in the Deconvolution Wizard.

The Image Histogram

The histogram is an important statistical tool for inspecting the image. It is included in the deconvolution wizard to be able to spot problems that might have occurred *during the acquisition*. The histogram shows the number of pixels as a function of the intensity (gray value) or groups of intensities. If the image is an 8 bit image, gray values vary between 0 and 255.

The *x*-axis is the gray value and the *y*-axis is the number of pixels in the image with that gray value. If the image is more than 8 bit, then gray values are collected to form a *bin*. For example, gray values in the range 0-9 are collected in bin 0, values in the range 10-19 in bin 1, etc. The histogram plot now shows the number of pixels in every bin.

The histogram in Figure 3.6 shows that the intensity distribution in the demo image is of reasonable quality. The narrow peak shown at the left represents the background pixels, all with similar values. The height of the peak represents the amount of background pixels (note that the vertical axes uses *logarithmic* scaling).

In this case there is also a small black gap at the left of the histogram. This indicates an electronic offset, often referred to as *black level*, in the signal recording chain of the microscope.

If a peak is visible at the extreme right hand side of the histogram it indicates *saturation* or *clipping*. Clipping is caused by intensities above the maximum allowed digital value. Usually, all

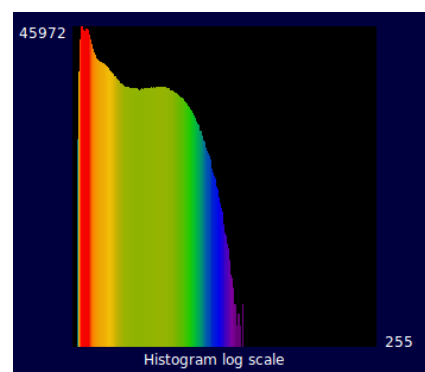


Figure 3.6: The image histogram. The vertical mapping mode can be selected from linear or logarithmic.

values above the maximum value are replaced by the maximum value. On rare occasions they are replaced by zeros. Clipping will have a negative effect on the results of deconvolution, especially with widefield images.

The histogram stage is included in the deconvolution wizard for examining purposes only. It does not affect the deconvolution process that follows².

Estimating the Average Background

In this stage the *average* background in a volume image is estimated. The average background corresponds with the noise-free equivalent of the background in 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.

In this stage the background value of the image (channel) can be automatically estimated or manually entered. For the automatic estimation you can choose the estimation mode and the area radius. The following estimation strategies are available:

- **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 by default set to 0.7 μm .
- **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 the planar region with radius r that has the lowest value.

Press the AUTO button in the wizard to continue with the automatic estimation with the options entered in this stage. Press the MANUAL button, if you want to skip the automatic estimation step and estimate it yourself.

In the next step the estimated value can be adapted either by altering the value in the *Absolute background* field or in the *Relative background* field. Setting the latter to -10, for example, lowers the estimated background by 10%. You can also choose how the background is stored in the template. By default the relative background is stored. If you plan to use this template also for other images that have been acquired with the same settings, you may want to use the absolute background instead.

In case you choose the manual setting, the absolute value is automatically stored in the template.

The last option determines the deconvolution algorithm. The default is *Classical Maximum Likelihood Estimation* (CMLE). In case of widefield images, you are offered the possibility to select *Quick Maximum Likelihood Estimation* (QMLE). The QMLE can typically be applied to images with a high signal-to-noise level. For STED images you

have the option to use either the CMLE or the *Good's roughness Maximum Likelihood Estimation* (GMLE) algorithm. So, the wizard offers a different set of algorithms depending on what microscope type is used. If done press ACCEPT to proceed to the deconvolution stage.

²Learn more about histograms at <https://www.svi.nl/ImageHistogram>

The Deconvolution Stage

Figure 3.7 shows the deconvolution stage where the deconvolution parameters can be set. The following parameters are available for the CMLE and GMLE algorithm. The options for QMLE are almost the same, except that the Quality threshold and the Iteration mode are not available. If you are interested in quantifying signal, we like to refer to our wiki for optimal settings.³

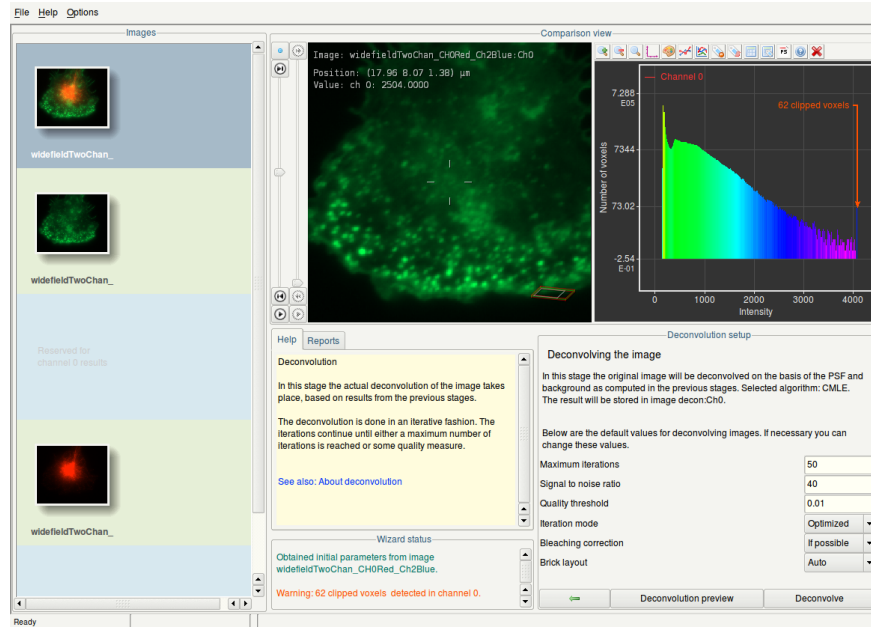


Figure 3.7: The Deconvolution Stage in the wizard with the deconvolution preview option.

1. **Number of iterations.** MLE is an iterative process that never stops if no *stopping criterion* is given. This stopping criterion can simply be the maximum number of iterations. This value can be adjusted, depending on the desired final quality of the image. For an initial run the value can be left at its default. To achieve the best result this value can be increased to e.g. 100. Another stopping criterion is the *Quality threshold* of the process (See Item 3).
2. **Signal to noise ratio.** The SNR is a parameter than controls the sharpness of the restoration 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. See also “Signal to Noise Ratio” on page 207.
3. **Quality threshold.** Beyond a certain amount of iterations, typically below 100, the difference between successive iterations becomes insignificant and the progress grinds to a halt. Therefore it is a good idea to monitor the progress with a quality measure, and to stop the iterations when the change in quality drops below a threshold. At a high setting of this quality threshold, e.g. 0.01, the quality difference between subsequent iterations may drop below the threshold before the indicated maximum number of iterations has been reached. The smaller the threshold the larger the number of iterations that will be reached; the higher the quality of the restoration. Still, the extra quality gain becomes very small at higher iteration counts. The absolute value of the final quality factor much depends on the data, the microscope type, and the background. It is a global value computed over the entire image, so the contribution of a local resolution increase can be small, as such it can only be used in a relative way to compare iterations and should not be seen as a value to compare the quality of different images.

³www.svi.nl/SignalQuantification

4. **Iteration mode.** In *optimized* mode (highly recommended) the iteration steps are bigger than in classical mode. The advantage of classical mode is that the direction of its smaller steps is sure to be in the right direction; this is not always the case in optimized mode. Fortunately, the algorithm detects if the optimized mode hits upon a sub optimal result. If so, it switches back to the classical mode to search for the optimum.
5. **Bleaching correction.** If this option is set to *if possible*, then the data is inspected for bleaching. 3D stacks and time series of widefield images will always be corrected. Confocal images can only be corrected if they are part of a time series, and when the bleaching over time shows exponential behavior.
6. **Brick layout.** When this option is set to *auto*, then Huygens Essential splits the image into bricks in two situations:
 - (a) The *system's memory* is not sufficiently large to allow an image to be deconvolved as a whole.
 - (b) *Spherical aberration* is present, for which the point spread function needs to be adapted to the depth.
7. **Array detector reduce mode.** This option only appears for array detector images and describes how the detectors will be reduced and/or supersampled prior to deconvolution. Subsets of detectors can each be combined into a single image, yielding a new image with fewer detectors. This can greatly improve the speed of the deconvolution as well as reduce the required memory. In some of the reduction modes the image is also supersampled, yielding an image with a greater number of pixels than the original. The reduction modes that can be selected depend on the type of array detector image. All in all the following reduction modes exist:
 - *Auto*: automatically selects the most appropriate reduction mode for the image.
 - *None*: no reduction is applied which means that a PSF of each detector is calculated and all PSFs and images are used simultaneously within the deconvolution procedure. Requires more RAM and is slower than other modes.
 - *All*: all Array Detectors are reduced to one single image which is subsequently deconvolved with one PSF matching with the reduced image.
 - *Super Y*: this mode is meant to be used for images acquired in *Fast mode* where the image is acquired with an elongated laser excitation where 4 lines are scanned simultaneously. All detectors are combined into a single supersampled image. The image is supersampled in the y direction by a factor of 4.
 - *Super XY*: all detector are combined into a single supersampled image. The image is supersampled in the x and y directions. The supersampling factors depend on the type of array detector image. A regular 32 detector *Airyscan* image will be supersampled by a factor of 2 in both the x and y direction, while a *Fast mode* image will be supersampled by a factor of 2 and 8 in the x and y direction respectively.

The button DECONVOLUTION PREVIEW can be used to quickly test a defined set of deconvolution parameters on a selected region of interest (ROI). This ROI is shown as a yellow frame within the left scene and includes, if possible, ten planes in the z dimension from the position of the slice shown. Thus, the position of the ROI depends on the position of the z slider and can be changed in xy with the left mouse button. Any adjustment of the ROI position or deconvolution parameters will immediately start a new and fast deconvolution of the ROI and update the deconvolved image in the scene marked by the ROI.

If you are ready to apply the tested deconvolution parameters to the complete image, simply press the button DECONVOLVE. You can of course skip the deconvolution preview process and immediately proceed with deconvolving the image. Pressing STOP DECONVOLUTION halts the iterations and retrieves the result from the previous iteration. If the first iteration is not yet complete an empty image will be shown.

Finishing or Restarting a Deconvolution Run

When a deconvolution run is finished, the result is shown in the right viewer, and you can choose to **RESTART ALL**, **RESTART CHANNEL**, **RESUME** or **ACCEPT/ALL DONE** the restoration:

- **Restart all** will return to the start stage and all results of all channels are discarded.
- **Restart (channel)** keeps the result so far and will return to the background stage where you can change the background setting and rerun the process with different deconvolution parameters. A new result will be generated to compare with previous results. This can be repeated several times.
- **Resume** keeps the result and returns to the stage where the deconvolution parameters can be entered. The software will ask to continue where it left off, or to start from the raw image again. A new result will be generated to compare with previous results. This can be repeated several times.
- **Accept, to next channel** proceeds to the first channel that has not yet been deconvolved. In case of All done, all channels have at least one deconvolution result or it was not a multi-channel image and it will proceed to the final stage.

The Final Stage

The last deconvolution result of each channel is automatically selected for the final image. By moving image names from the *Available* to the *Selected* list, you can combine channels as desired. You can use the mouse-scroll to scroll the lists. Also, when clicking on a thumbnail in the left *Images* pane, the corresponding name is highlighted in the *Available* list, if it was not already moved to the *Selected* list. Not only the deconvolved results are available, but also you can include an original image in the final result. Continue to the next stage if you are satisfied with the selection. Please note, pressing **RESTART** will discard all the results and will immediately return to the start stage.

What now follows is the possibility to **RESTART** all again (which discards the results), **SAVE TEMPLATE** for deconvolution, or press **DONE** to export the result to the main window of Huygens Essential. You can also use an icon of one of the specific visualization and analysis options to immediately continue with the deconvolved image.

The Comparison View

Detailed image comparisons of all the thumbnails in the left panel of the Wizard can be done using the *Comparison View* displays with loading images via a right-click on their corresponding thumbnails and select “Show left” or “Show right” (See Figure 3.8).

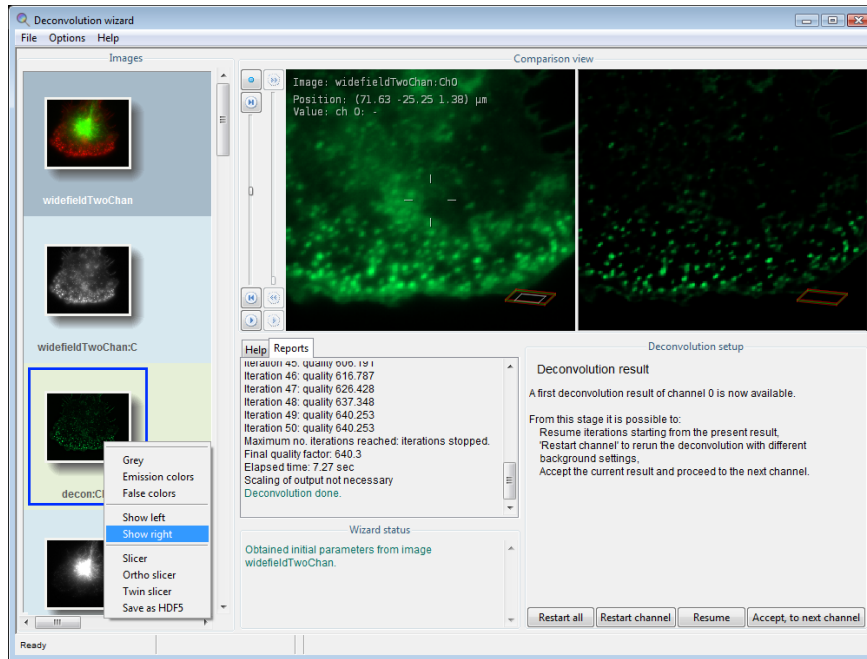


Figure 3.8: Showing an image on the Comparison View: right-click on the thumbnail and select “Show left” or “Show right”.

In order to plot the image intensities along a specific path click on the image and draw a line without releasing the mouse. A plot will be shown on the Comparison View display (left or right) opposite to where the line is drawn. To hide the plot, click a point of the Comparison View off the drawn line. To show the same plot again, click on the drawn line. This operation can be performed on both sides of the Comparison View. The plot will show the intensities of the two images loaded on the displays, for comparing purposes, see Figure 3.9 on page 24. If one of the displays is empty the plot will only show the intensities the image shown.

Multi-channel Images

Multi channel images can be deconvolved in a semi-automatic fashion. Still, the results can be fine-tuned for each individual channel. After the preprocessing stage the multi-channel image is split into single channel images named <imagename>:Ch0, <imagename>:Ch1, etc. The first is automatically selected for deconvolution.

The procedure to deconvolve a multi channel data set is exactly the same as for a single channel image. Therefore multiple reruns on the channel can be done manually, just as with single channel data. Press ACCEPT, TO NEXT CHANNEL in the last stage, to select the following channels for deconvolution. There is the possibility to skip one or more channels.

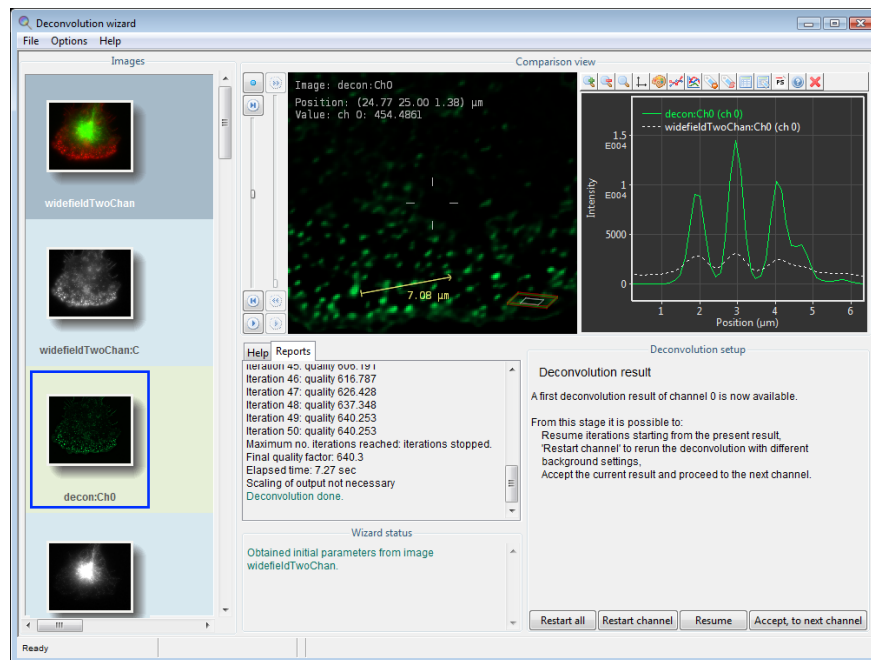


Figure 3.9: Comparing intensity profiles of two images in the Deconvolution Wizard. Notice that the plot prompts at the display opposite to where the line is drawn.

When the last channel has been processed, the wizard offers the option to specify what channels and in what order they should be combined for the final deconvolved image.

Z-drift Correcting for Time Series

For 3D time series the wizard shows an additional step to enable the correction for movement in the z direction (axial) that could have been occurred, for instance, by thermal drift of the microscope table. In case of a multi channel image, the corrector can survey *All channels* and determine the mean z position of the channels, or it can take *One channel* as set by the *Reference channel* parameter.

After determining the z positions per frame, the z positions (not the image) can be filtered using a *median*, *Gaussian* or *Kuwahara* filter of variable width. If the drift is gradual, a Gaussian filter is probably best. In case of a drift with sudden reversals or outliers a median filter is best. In case the z positions show sudden jumps, we recommend the Kuwahara filter.

Saving the Result

After exporting the deconvolution result to the main window, the result can be saved. Select the image to be saved and select FILE → SAVE AS... in the menu bar. The *HDF5* file format preserves all microscopic parameters and applies a lossless compression. To see which other file formats the Huygens software supports, see <https://www.svi.nl/FileFormats>.

Select FILE → SAVE STAGE REPORT to store the information as displayed in the Report tab.

Chapter 4

Deconvolution Express

Why Deconvolution Express

Deconvolution Express is a tool that helps to reveal great image detail, less blur and less noise, with one click of a button. Huygens makes this possible by using unsupervised profiles for finding acceptable parameter values, for which user input is generally needed.

Notice that although the *Deconvolution Wizard* is still the advanced recommended tool for optimizing deconvolution, *Deconvolution Express* can be quite helpful in getting a first impression of the quality attainable on an image.

Verifying the image parameters

The quality of deconvolution depends, among other things, on the image meta data or the lack thereof. In other words, the microscopy parameters describing the image (sampling sizes, wavelengths, etc.) are as important as the deconvolution parameters (number of iterations, quality threshold, etcetera) for the restoration result. Thus, before deconvolving the image please make sure that the image meta data reliably describes the acquisition settings. This is particularly important for the sampling sizes, the refractive indexes and the imaging direction. For more details see “Verifying Microscopy Parameters” on page 13.

If the meta data parameters are correct it is usually helpful to save them to a Huygens Microscopy Parameters template file for later re-use. A template can also be applied to a batch of images.

Starting Deconvolution Express

The *Deconvolution Express* can be started by clicking on one image in the main window, and by selecting in the main task bar DECONVOLUTION → DECONVOLUTION EXPRESS or by clicking on the icon *Express* in the icon task bar (see Figure 4.1 on page 26).

At the top right, the Deconvolution Express window shows a single slice of the selected image. At the left part of the window sliders are shown with which you can change the z plane and time frame of the displayed image. Below these sliders, active channels can be unselected and selected, and the color scheme can be adjusted. Also, brightness and gamma can be adjusted to optimize the visibility of the image data. A *Report window* presents information on how you can proceed and what is relevant for the interpreting the final result.

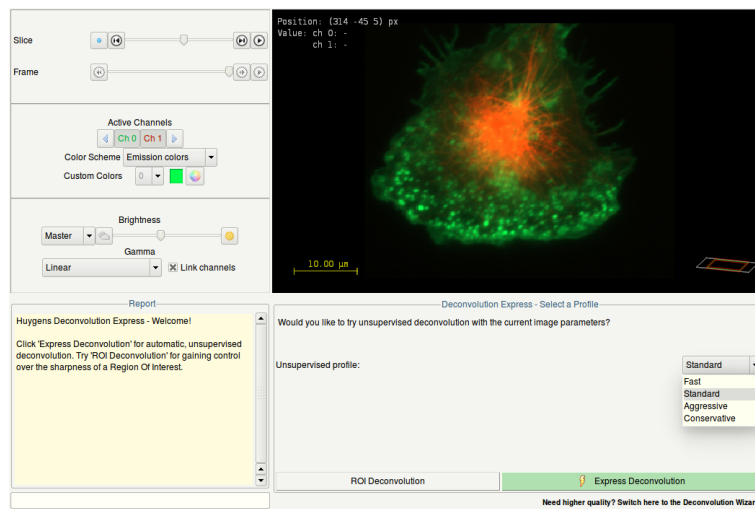


Figure 4.1: Deconvolution Express showing the unsupervised profiles.

The two buttons EXPRESS DECONVOLUTION and ROI DECONVOLUTION, can be used to immediately run an automatic unsupervised deconvolution, or to first control the sharpness settings within a ROI.

By first pressing ROI DECONVOLUTION, you can adjust the sharpness of each channel independently and see instantaneously the effect on the quality of deconvolution within a specific region of interest (ROI). This ROI is defined by the yellow frame projected on the displayed image, and can be moved in the XY plane using the mouse left button. Similarly, a change in the unsupervised profile that is applied from the drop down menu will immediately update the deconvolution result in the ROI.

Once the sharpness setting for every channel has been optimized, you can press the button EXPRESS DECONVOLUTION to apply these settings.

You can, of course, skip the ROI Deconvolution and immediately proceed with EXPRESS DECONVOLUTION. Also then, a different unsupervised profile can be selected. Pressing ABORT halts the iterations and retrieves the result from the previous iteration. After the deconvolution is finished, you can save the Task report as a .txt file, and/or continue with the *Twin Slicer* to compare the deconvolved image with the original data side-by-side. To correct for additional image distortions you can decide to directly switch to the *Object Stabilizer* or *Chromatic Aberration Corrector*, or if you prefer to fine-tune the optimal deconvolution parameters you can continue with using the *Deconvolution Wizard*.

Unsupervised Profiles

Deconvolution Express offers different unsupervised profiles (fast, standard, aggressive, conservative) for finding the optimal deconvolution method for each image (Figure 4.1 on page 26). Each profile surveys the image with different conditions to arrive at the best deconvolution parameters for that image. These profiles are available for user selection and can be generally understood as a speed versus quality gauge:

- **Fast:** in this unsupervised profile Huygens uses faster deconvolution algorithms or fewer iterations than in other profiles. The ratio “speed versus quality gain” is particularly suited for widefield images.
- **Standard (Default choice):** this choice is a good trade-off between speed, resolution gain, and noise reduction.
- **Aggressive:** this is the method of choice for reaching higher resolution. It is recommended for images containing high enough signal. It is less suitable for images with sparse signal such as e.g., STED or low signal confocal images.
- **Conservative:** this profile is more cautious on attempting a very high resolution gain. Instead, it is very effective at reducing image noise. Therefore, it should be a safe mode for avoiding image artifacts.

Based on the selected unsupervised profile, the *Deconvolution Express* tool will decide on the best deconvolution parameters for the input image. The tool will first estimate the image background and *Signal to Noise Ratio (SNR)*. Notice that the more complete and accurate the image parameters are, the more reliable the SNR findings. In turn, this has important implications on the resulting image quality. For example, Leica¹ provides extra information on their LIF images which are saved with the LAS X software (version 3.x or higher) making the SNR estimations more reliable. As previously described, the sharpness of the deconvolved result can be further adjusted using the button ROI DECONVOLUTION.

In addition to estimating the image statistics, the Deconvolution Express tool also makes smart choices on the type of algorithm, the number of iterations, the quality threshold, and the optimal number of bricks, among others.

¹Leica Microsystems GmbH

Chapter 5

Batch Express

With Huygens 18.04, we have introduced the new option *Batch Express*, which is able to deconvolve all your images located within a specified directory fully automatically using the unsupervised profiles that are also available for Deconvolution Express. Please note that the specifications of your Huygens license determines which microscope files can be processed and which file formats can be read¹.

Within *Batch Express* you can specify a specific directory as the “Watch folder”. This folder is constantly scanned by *Batch Express* for new images that are added. For every image it finds, a new deconvolution task is created and placed within the queue in the *Batch Express* window. Images within this queue will be deconvolved one by one from the top down. Every task consists of a fully automated parameter estimation of the image, the actual deconvolution, and the saving of the deconvolved result to the directory specified as “Output folder”.

Difference between Batch Express and other Huygens deconvolution tools

Huygens Essential offers different tools for deconvolving image data, ensuring that every user is served optimally.

For example, the *Deconvolution Wizard* is intended for users than prefer to be guided through the different stages involved in optimizing the parameters for deconvolution. During these stages, image statistics are verified and deconvolution parameters can be adjusted manually. Within Huygens Professional, the *Operation Window* permits you to have more freedom to change parameters and thus demands more user experience. *Deconvolution Express* and *Batch Express*, both present in Essential and Professional, use the same unsupervised profiles and Huygens algorithms to deconvolve images. With Deconvolution Express, you first need to open the image in Huygens, whereas with *Batch Express* you only need to define the location of the images and the destination for saving the resulting deconvolved files.

Batch Express and the *Batch Processor* can deconvolve many images with the difference that the *Batch Processor* requires pre-made templates for the microscopy image parameters and deconvolution parameters. The advantage of using a deconvolution template is that it can include a range of additional pre- and post-processing steps like for example bleaching and chromatic aberration correction, and stabilization. The possibility of applying the same template to multiple images guarantees that these images are processed similarly which may be necessary for comparison and quantification studies.

¹The files that Huygens can read depends on the options in your Huygens license. For more information see <https://www.svi.nl/LicenseStringDetails>.

Please note that *Batch Express* selects the best possible deconvolution parameters for each individual image and gives a first good impression of the quality attainable, with the consequence that different images may be processed differently.

The Batch Express window

Batch Express can be started by clicking in the main task bar DECONVOLUTION → BATCH EXPRESS or by clicking the icon in the top icon task bar.

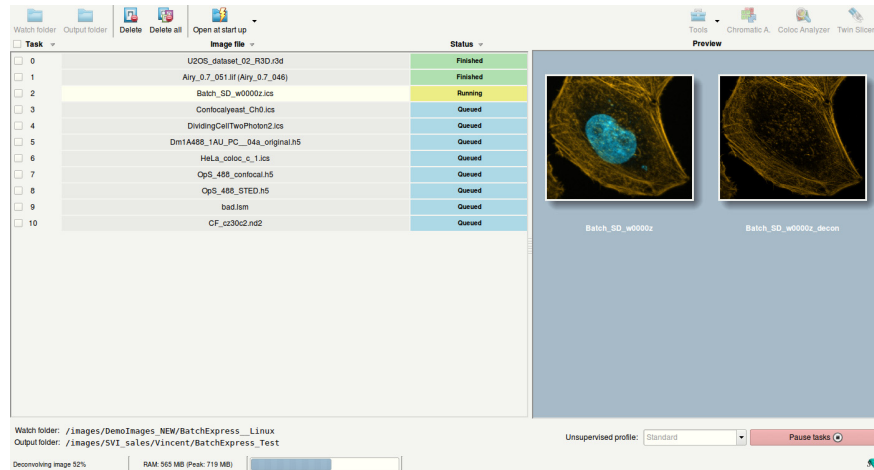


Figure 5.1: The *Batch Express* window shows images of various file formats in a queue. The colored status buttons indicate that two tasks are already finished, task number 2 is still running, and other tasks are in the queue scheduled for deconvolution.

The windows (see Figure 5.1) is split into seven main parts. In the left top corner, you can find a button bar to organize tasks. The main center space is available for the queued tasks and will be empty when you start *Batch Express* for the first time. In the left bottom corner, the current selected directories are shown for the “Watch folder” and “Output folder”. At the top right, a button bar enables the user to interact with the result of the deconvolution tasks. The center space at the right is reserved for displaying the input image and its deconvolved result. Below this space, a pull down menu is present for selecting the deconvolution strategy using a profile, and a play and pause button is present to start and pause the deconvolution of images from the queue. Information located at the bottom of the *Batch Express* window shows the current process, the amount of RAM and a progress indicator.

In the right bottom corner, the unsupervised profiles can be chosen. These profiles are the same as in the *Deconvolution Express* option (see Chapter 4 on page 25). When changing the profile, a pop up window will come up with the question if the user want to re-deconvolve finished tasks with the new profile.

Working with the Batch Express

By clicking on the WATCH FOLDER button at the top left of the window you can select the location of the images that you like to deconvolve. *Batch Express* will automatically find images in this directory and one level below which means images in the sub-directories of the “Watch folder”. Next, you can select with OUTPUT FOLDER button the directory to which the deconvolved results should be saved. The output folder cannot be the same folder or a sub folder of the watch folder as stored images will be detected again by the *Batch Express*. This would result in an infinite loop of deconvolving the same image.

Every image detected in the watched folder will result in a single task in the queue, with the exception of lif files and tiff series. One lif file can contain multiple images and for each image a separate task will be generated. A serie of Tiff files (time sequence or Z stack) can be seen as one image for which a single task is made. Once in the queue, a blue-colored status button behind each task indicates it is *Ready to run*. A task can be set to *Deactivated* by clicking on the status button. This can be undone by clicking again.

Tasks can be selected by clicking on the select box in the first column on the left. With SHIFT-left-mouse click on the select box, a range of tasks can be selected by clicking only on the first and the last box of the range. With CTRL-left-mouse click on the image name, multiple tasks can be selected. All tasks can be selected or deselected by clicking on the most upper select box (located at the left of “Task”). Selected tasks or all tasks can be deleted by clicking the DELETE or the DELETE ALL button, respectively. The status of all the selected tasks can be changed by just clicking on only one of the selected task’s status buttons. Next to the delete all button, you can find the OPEN AT START UP button. When this button is set to *enable*, the *Batch Express* will open itself when Huygens Essential is started.

By clicking on the green RUN TASKS button, which then switches to PAUSE TASKS, Batch Express will start deconvolving all the queued tasks. The status button of the first task will turn yellow and show *Running*, and the status button of all other tasks will show a light blue color and Queued. Upon clicking the PAUSE TASKS button, RUN TASKS will appear again and the task that is currently running is *Stopped* as indicated by the orange status button. If you click on the status button marked with *Stopped* or *Ready to run*, this specific task becomes *Deactivated* and will be skipped if you resume with deconvolving tasks by clicking the RUN TASKS button.

At the moment a task is running, it cannot be deactivated. It first needs to be stopped by clicking on the button RUN TASKS. Also, while running, directories cannot be changed.

After the deconvolution of a task is completed, the image is saved as an ICS2 file in the output folder. ICS2 is a file type that ensures that the full 32bit range is kept, is compatible with many software packages, and which handles metadata very well. The status button of completed tasks become green and show *Finished*. When the deconvolution of the other tasks proceeds, you can restart a finished task by clicking on its status button and confirming its reactivation. The status button will then turn again to light-blue and *Queued*, and the task will be processed when the other tasks in the queue are completed. Instead, if you click on the status button of a finished task when the deconvolution has stopped, the status button changes to blue and shows *Ready to run*.

A status button that is red-colored and marked as *Failed*, indicates that a task has not been completed. The reason for this failure is that the microscopy parameters cannot be read correctly from the image file. It is then advised to verify the microscopy image parameters .

Batch Express is constantly in a waiting mode for new images that are saved in the “Watch folder”. When a new image is detected, it appears as a new task at the bottom of the queue.

Batch Express preview

When a task is running, the *Preview* windows shows two thumbnails with on the left the original image, and on the right the channels that have sofar been deconvolved. When a task is finished, these thumbnails are automatically refreshed with the images of the next task.

After stopping the deconvolution process by clicking the PAUSE TASKS button, you can take a closer look at images of tasks that are marked as *Finished*. Tasks previously marked as *Queued* are now marked as *Ready to run*. They cannot be viewed at this point, since they have not yet been loaded by *Batch Express*. If you click on an image file name in the queue that is marked as *Finished*, you can see the original and deconvolved image side by side. With the TWIN SLICER button in the right top corner, the input and output image will be loaded in the *Twin Slicer* which allows a closer comparison of both (see Chapter 13 “The Twin Slicer” on page 83). Buttons marked CHROMATIC A. and COLOC ANALYZER will open multichannel images in the *Chromatic Aberration Corrector* or *Colocalization Analyzer*, respectively (see Chapter 24 on page 137 and Chapter 30 on page 187). If these two buttons are inactive the image selected contains only one channel.

Under the TOOLS button, you have immediate access to the image file via the file manager of your operating system, and you can open the image within one of the other Huygens tools. These and additional tools can also be accessed by right-clicking on one of the two *Batch Express* thumbnails and using the popup menu. From here, it is also possible to open the image in Huygens main window or to save the deconvolution template. This template includes the parameters applied by *Batch Express*, and can subsequently be used in the *Batch Processor* to deconvolve many images with the same deconvolution parameters. When opening an image with any of the tools, it will also be exported to the main Huygens window. This will take the same amount of time as opening the image directly from the Huygens main window into this same option.

Chapter 6

The PSF Distiller

Huygens Essential is optionally equipped with the *PSF Distiller*. This wizard driven tool guides the user through the process of distilling a *Point Spread Function* (PSF) from 3D bead images. It helps in selecting bead images, and in creating and saving the PSF for further deconvolution runs. The wizard is able to measure a PSF from one or multiple images, each containing one or more fluorescent beads. It is also possible to distill multi-channel PSF's from multi-color bead images, or to assemble a multi-channel PSF from separate single channel bead images PSF's.

Measured PSF's improve deconvolution results and may also serve as a quality test for the microscope¹. The measured PSF acts as a calibration of the microscope in the sense of relating a physical known object (bead) with what the microscope actually measures (bead image). Figure 6.1 shows an example of a theoretical, and a measured PSF.




Figure 6.1: An *xz* cross section of a theoretical PSF (left) versus a measured PSF (right) for the same confocal setup.

Beads suited for PSF Distillation

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. In order to measure a PSF from beads, the diameter of the beads should be in the order of the *half intensity width*² (HIW; also referred to as *full width at half maximum*, FWHM) of the expected PSF. Larger beads will reduce the accuracy of the Distiller, while smaller beads yield insufficient signal for accurate averaging, resulting also in reduced accuracy. Beads ranging from 100 to 200 nm can be used. Typically beads with a diameter of around 170 nm perform very well for many types of microscopy.

Use the default SNR settings. If possible, it is recommended 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 170 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 microscopy parameters that you will use later to image your specimens. Please find more practical information about beads for PSF measurements on the SVI Wiki³.

¹We advise to measure the PSF after a change in the recording setup and certainly after each maintenance job in which the optics or scanning device was serviced.

²<https://www.svi.nl/HalfIntensityWidth>

³<https://www.svi.nl/RecordingBeads>

The PSF Distiller Window

If your license includes the PSF Distiller option, select `DECONVOLUTION` → `PSF DISTILLER WIZARD` to start the Distiller. In the opening window (Figure 6.2) different panes show the input beads field, the Help field, the distiller status field, the report field, the wizard field and the PSFs & Accus field:

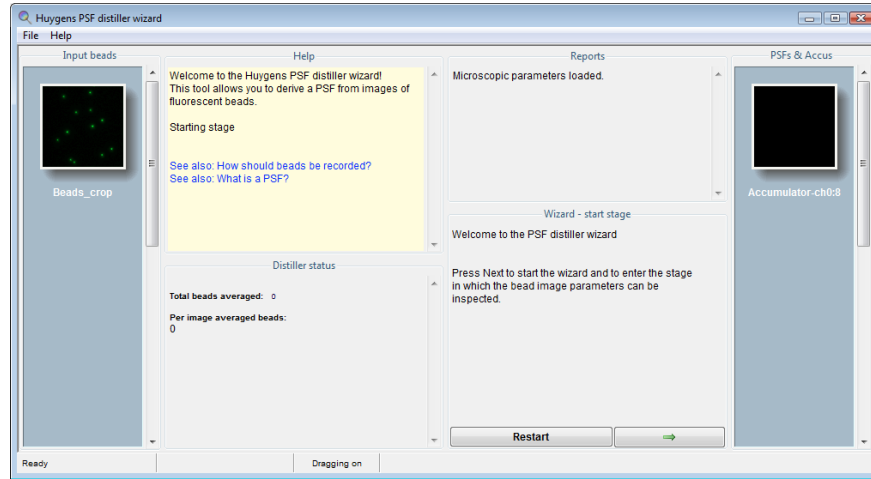


Figure 6.2: The PSF Distiller window. Different panes show the input beads field, the Help field, the distiller status field the report field, the wizard field and the PSF and Accus field.

- *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 grayed 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.
- *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 PSFs & Accus pane*
Displays the intermediate results of the distilling process. The thumbnails can be inspected with the Sliders. 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.

The Processing Stages in the Wizard

The following steps and stages are to be followed:

- Loading an image.
- Start Stage: here the possibility exists to load a microscopic parameter template and check the microscopic parameters.
- Averaging Stage: in this stage (all channels of) the image are searched for beads that meet the selection criteria. After each successful or unsuccessful search there is the possibility to load in additional bead images, or go to the next stage.
- Distillation Stage: in this stage the PSF is measured from the averaged beads, for all available channels.
- Finalizing the result: in case it is desired to combine results from earlier distillations with the current result to obtain a multi channel PSF, an earlier result can be added here. It is also possible to add single or multi channel previous results to a current multi channel result.
- Save the result.

The next sections of this chapter will explain the wizard stages in detail. See “Loading an Image” on page 12 and “Saving the Result” on page 24 for more information on handling image files.

Starting the Distiller

After launching Huygens Essential, open the first bead image via FILE → OPEN. If the license includes the PSF Distiller option, start the PSF distiller wizard via the menu DECONVOLUTIONANALYSIS → PSF DISTILLER WIZARD. When the window is opened, one or more *accumulator images* will be created into which later on the averaged beads will be kept. Now the start stage will be entered.

Verifying Microscopic Parameters

See “Verifying Microscopy Parameters” on page 13 for more information on the microscopic parameters. Next to the optical parameters listed in Table 3.1 on page 13, it is in particular important to check the *sampling densities*.

Do not use undersampled bead images. If any of the entry fields for the sampling density turns orange or red, the data is *unusable* for distilling PSF's.

If there are multiple bead images, then the parameters of bead images loaded at a later stage should match the ones to establish in this stage; a warning of any mismatch will be given.

Averaging Stage

After setting the *bead diameter* and estimating the *Signal to Noise Ratio (SNR)*, the image is searched for beads that meet the following selection criteria:

- A bead should not be *too close to another bead*. If a bead is too close to another bead, their signals will interfere. In widefield bead images this is quite problematic due to the large size of the blur cone. Fortunately, widefield PSF's can be derived from a single bead within an image.

- 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 usable beads are found, an explanation and some advice will be displayed in a pop up window. First, the software will try to find beads with ideal selection criteria. If this does not yield a single bead, it will automatically retry with reduced inter-bead distance criteria.

After having added the last bead recording press DISTILL to enter the *distillation* stage (See Figure 6.3).

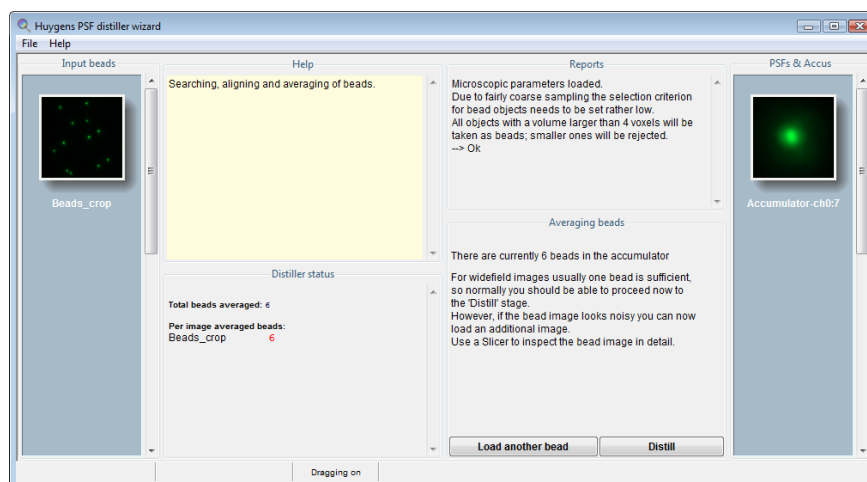


Figure 6.3: The averaging stage in the PSF Distiller wizard.

Confocal and Two Photon Bead Images

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. To load more bead images press LOAD ANOTHER BEAD in the averaging stage and either select an image from the main window or open a new one.

Widefield Bead Images

Images from 160 nm beads should look like smooth fuzzy blobs with no visible noise. Use the default SNR settings. Averaging beads is not necessary for widefield images.

Distillation Stage

The distillation stage usually requires no user intervention, though in some cases a pop up window will be displayed with a question or a warning. All channels will be processed automatically.

Finalizing the result

At this stage, a previously obtained PSF image can be added as a channel before or after the current result. Press **ADD CHANNELS** if this is desirable. The Distiller will compare the microscopic image parameters of the selected PSF image and check its content. In case there are differences, the software will ask to decide between ignoring these and discarding the selected file. Press **DONE**, if you are finished adding channels to the distilled PSF.

Full-width at half-maximum estimator

This stage also offers the possibility to measure the quality of the PSF using the PSF full-width at half-maximum (FWHM) Estimator tool. If you have the PSF distiller option in your license, you can also find this tool in the task bar menu of Huygens Essential under **ANALYSIS**.

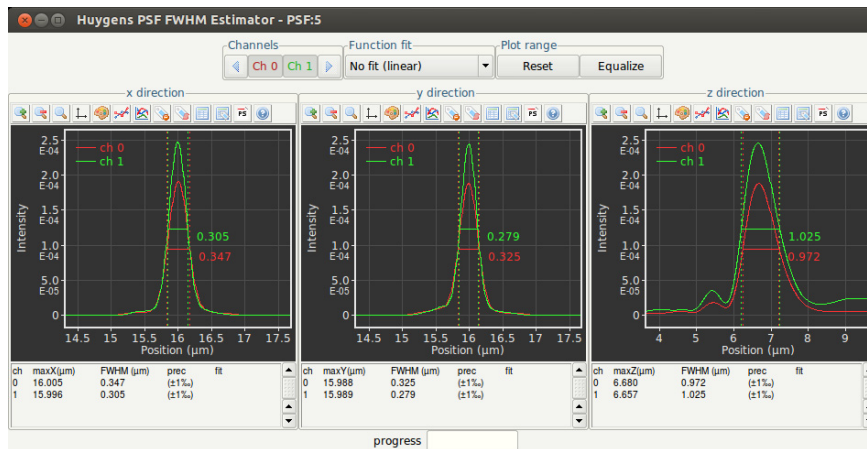


Figure 6.4: The PSF-FWHM overview. No curve fitting was selected.

The PSF FWHM Estimator creates a line intensity profile through the center of the PSF in all dimensions, and measures the X value of the peak maximum and the width of each peak at approximately half of the maximum value using a curve fitting procedure. Curve fitters such as Gaussian, Lorentzian, Voigt and Pearson, are available for this. This is done for all channels. For further background information on the FWHM⁴, PSFs and the PSF Distiller⁵, please refer to the SVI Wiki.

Finally, press **EXPORT & CLOSE**. This will export the distilled PSF to the main window, where the result can be saved or used immediately in a subsequent deconvolution run.

⁴<https://www.svi.nl/HalfIntensityWidth>

⁵<https://www.svi.nl/PsfDistiller>

Chapter 7

The Batch Processor

Once you are familiar with a particular kind of dataset and the restoration parameters are determined, more 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 multi-threading or multi-GPU 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 Essential and Professional.

The Batch Processor

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

The main Batch Processor window can be separated into five different elements from the top to the bottom. It shows initially no tasks (see Figure 7.1 on page 40). The elements are:

- **Upper button** bar shows from left to right three buttons to select a destination directory, create a new destination directory below the currently selected one, and change the output file format of the resulting images. The next two buttons enable opening and saving a task list to and from a batch file stored on the computer. The next four buttons enable editing the task list and starting or stopping a run of tasks. The right-most three buttons control the resources allocated to the tasks when they are run: one button switches between single-GPU mode (use the currently selected default GPU device) and multi-GPU mode (use all selected devices in the GPU status window), the other buttons control the number of concurrent tasks and CPU threads per task that are allocated.

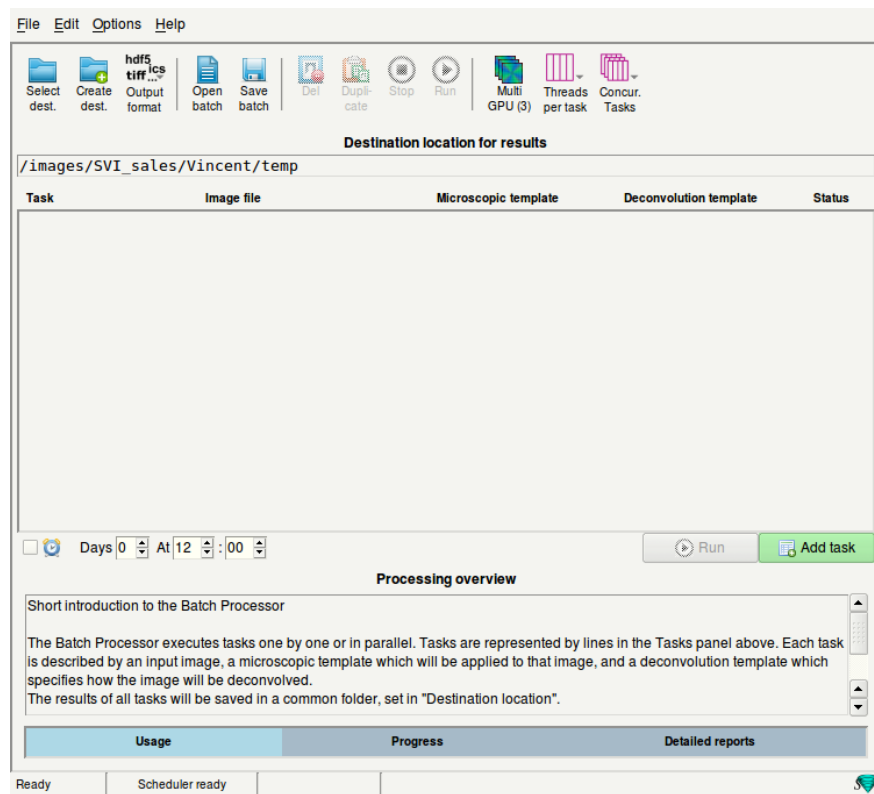


Figure 7.1: The Batch Processor main window.

- **Destination Location for Results.** This is the directory where the resulting images will be placed during the batch run. With the two folder buttons (in the upper button bar) 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** (located below the task area) has a clock at the left side to delay the beginning of the processing, and buttons at the right side to run and add tasks to the list (one by one or many at the same time). These tasks can also be found as buttons in the upper button bar.
- **The Processing overview** in the detailed information is given about the whole process in its different stages.
- **The status bar** at the bottom-end of the batch processor window supplies some status information about Huygens Essential. The leftmost tab shows the state of Huygens Essential, and the tab to the right of it shows information of the batch scheduler. The rightmost of these tabs gives information about the button the mouse is currently pointing at (now empty).

Usage

Before starting a batch processing job, the save location of the results should be defined (*Destination Location for Results* field) using one of the two destination buttons in the upper button bar, and a file format should be selected in which the results should be stored (OUTPUT FORMAT button in upper button bar or via 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 the green ADD TASKS button below the task field a new menu entitled Selected images is expanded at the right (see Figure 7.2).

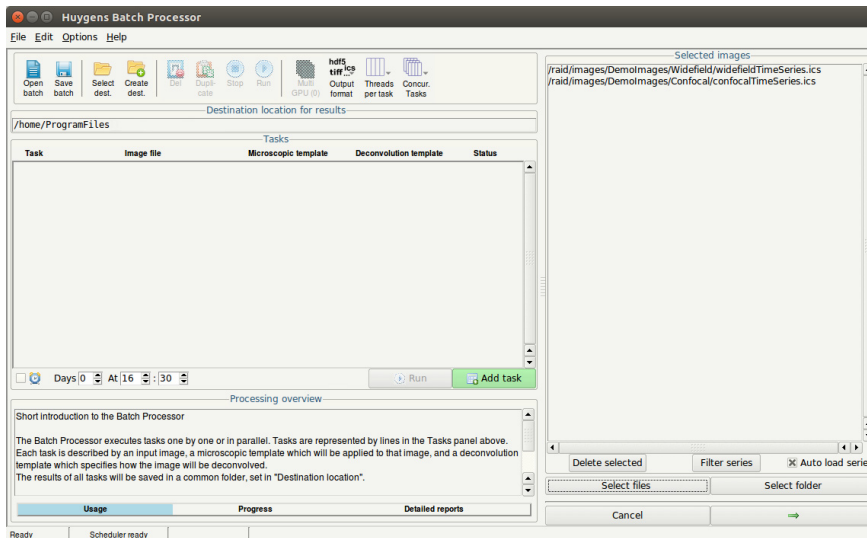


Figure 7.2: The Batch Processor with the *Selected images* field expanded.

Either a complete folder containing images can be selected or folders can be directed to select a single file. If a file is selected containing multiple sub-images (for example 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. Upon checking the *Auto load series* box (on), the Batch processor will try to interpret each listed file as a start of a file series.

If the box is unchecked (off), only the listed files will be deconvolved. If a complete file series is selected and you would like to deconvolve that series as one image, you can click *Filter series* to prune all members of a file series except the first from the list.

If you select an entire folder containing one or more file series and the *Auto load series* box is checked (on), only the first file in a series will be listed. Leave *Auto load series* to unchecked (off) if you do not want this; you can always prune the list with *Filter series*.

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

Microscopic Templates: Describing the Images

To guarantee the quality of the deconvolution, it is very important that the image acquisition conditions be properly described.

The microscopy parameters can be edited and saved to a template using the Edit Microscopic Parameters window. See “Verifying Microscopy Parameters” on page 13. These templates can then be reused at a later time. Parameter templates can be edited at any moment via the EDIT → EDIT MICROSCOPIC PARAMETERS menu. You can also edit a pre-made template within the Batch Processor or create a completely NEW TEMPLATE. In addition to editing the microscopic parameters, Huygens also offers the possibility to deconvolve the image with the microscopy parameters stored in the image metadata. Because different manufacturers save the microscopic data with different confidence levels¹ the Batch Processor allows you to choose whether only reliable metadata should be taken into account for the deconvolution (see Figure 7.3).




Figure 7.3: The Batch Processor metadata selector.

It is therefore important to understand what the different microscopy parameters refer to and know how to establish them (See Table 3.1 on page 13). A typically conflictive one is *the backprojected pinhole radius*. This parameter is not difficult to calculate², especially with the assistance of the online *backprojected pinhole calculator*³.

As mentioned above, the entered parameters can be stored in *templates* for convenience sake, so the same template can be applied to a series of images acquired with the same settings.

Click the NEXT button (⇒) again when finished with the microscopic template.

Deconvolution Templates: Configuring the Restoration Process

Like the microscopy parameters, the deconvolution parameters can be saved in a template. Again, it is important that it be understood what the role of these parameters is (see “The Deconvolution Stage” on page 20 and the *Restoration Parameters* article⁴).

At this stage, you can load one of the existing deconvolution templates from the list or create a NEW TEMPLATE. If you select NEW TEMPLATE, the first tab allows you to LOAD a previously made template or SAVE a new template.

You can adjust the settings for your deconvolution procedure and select the deconvolution ALGORITHM that you like to apply. Here, you can also choose to use a Theoretical PSF or to load and use an Experimental PSF. The experimental PSF file can be selected under the next tab, named PSF. Please note that you need to review the deconvolution parameters for all the channels present in the image file. For more detailed information on the deconvolution parameters look under “The Deconvolution Stage” on page 20.

With the PRE and POST tabs, you will be able to select *Autocrop*, perform *Stabilize along z* for 3D images, *Pre-adjust baseline* and perform *Z-Drift correction* with time series images. Furthermore, chromatic aberration correction is possible via an additional tab by loading a pre-made template (see “24” on page 141).

¹<https://www.svi.nl/ConfidenceLevels>

²<https://www.svi.nl/BackProjectedPinholeRadius>

³<https://www.svi.nl/BackprojectedPinholeCalculator>

⁴<https://www.svi.nl/RestorationParameters>

The **OPTIONS** tab offers you the possibility to adjust the *Timeout* setting for limiting the computational time spent on one task. This prevents the *Batch Processor* from getting stuck on a specific file while there are more tasks queued.

Once all settings are done, click the **DONE** button at the bottom of the screen. You can also have a look at our detailed *Batch Processor Tutorial* on our support wiki⁵.


A typically conflictive point is setting the *signal-to-noise ratio*. Note that this is not a number describing the image, but something that can be tuned to achieve different deconvolution results. Please see “The Deconvolution Stage” on page 20 and the *Set the Signal to Noise Ratio article* in the SVI Wiki⁶.

Location of the Saved Templates

The microscopic and deconvolution templates are by default saved in the **SVI/TEMPLATES** folder in the user’s home directory⁷. 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 clicking the *Done button*, 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 by clicking the **DELETE** button  or selecting **EDIT → DELETE TASK** from the menu.

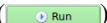
New tasks can still be added to the queue before starting the computations, or start computing right away and add new tasks afterward.

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 will be 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 by clicking the **SAVE BATCH** button or by selecting **FILE → SAVE** in the menu.

Add as many tasks as required, single files of complete folders, the Huygens *Batch Processor* will run them all. By pushing the **Run** 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 according to the progress of the process.

The restored and deconvolved images are saved in the selected destination directory as soon

⁵<https://www.svi.nl/BatchProcessor>

⁶<https://www.svi.nl/FrontPageSetTheSignalToNoiseRatio>

⁷The user home directories are usually located in **C:\Users** on Microsoft Windows Vista, Windows 7, Windows 8, and Windows 10 (64 bit). On Mac OS X Mountain Lion, Mavericks, Yosemite, El Capitan. Sierra (X11, XQuartz) they are usually in **/Users**, and on Linux in **/home**.

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 should not block other activities, 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 when the processing should start. The timer checkbox is then selected automatically; deselect it to disable the timer.

Exiting the batch Processor

If the Batch Processor is closed while running tasks, it will stop all running tasks. The Batch Processor window can be scrolled down while running tasks. Just exit the Batch Processor after all the jobs have been finished.

Menus

Most options in the menu are also represented by a button in the upper task bar. The FILE menu can be used to save and load the tasks list for future reference. In this menu also the information reported during the batch processing can be saved.

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

The OPTIONS menu has four sub menus:

- **OUTPUT FORMAT:** this sub-menu shows several options for the file format to select for saving the restoration result.
- **INPUT CONVERSION:** this sub-menu selects what to do with input images that are stored in byte or 16-bits integer format. The default option in the Batch Processor is to read them as integer images. The alternative is to force conversion to 32-bit float images. Since the input format for a deconvolution may affect the output format, this option may help to prevent unwanted loss of information by setting it to "to float" or reduce the memory requirement by setting it to "to int". In cases where a deconvolution result is different between interactive use and Batch Processor use (and when the same result would be desirable), changing this option is a good candidate to test.
- **THREADS PER TASK:** this sub-menu allows to set the number of processor cores 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 a different value as threads.
- **CONCURRENT TASKS:** if the system has multiple processor cores, it can be chosen 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. In multi-GPU mode, the default number of tasks will be set to the number of participating GPU devices (the number of GPUs with the "selected" check box ticked in the GPU status window). See also Figure 31.1 on page 199. If you run into problems due to a lack of memory for running this many tasks in parallel, the easiest to understand way to change the number of multi-GPU concurrent tasks is to have fewer selected GPU devices in the GPU status window. If you are certain that more than one task will fit on each GPU device, you can use the concurrent tasks option in the Batch Processor to have more than one task per GPU. Be careful to check if everything really works the way you expect it to when you attempt this, however, because (i) this will need really large amounts of system RAM compared to each image's size, and (ii) the

slowdown that will result when some task does not fit on its allotted GPU device is often much worse than the speed-up gained when it works.

The HELP menu can bring you to our online documentation wiki page and gives under the ABOUT option more information on, for example, what current Huygens version you are using.

Chapter 8

Huygens STED Deconvolution

Introduction

Stimulated-Emission Depletion (STED) microscopy is a fluorescence microscopy technique which overcomes and improves the diffraction-limited resolution of regular confocal microscopy techniques up to 4-5 times¹, leading to the so-called sub-diffraction resolution or super resolution in the lateral and, in case of STED 3D, also the axial direction.

The Huygens STED optical option offers support for deconvolving STED images, yielding stunning results in xy and in z . Several other Huygens Essential tools offer an additional STED mode for STED-specific image processing. For example, the PSF distiller (see Chapter 6 on page 33) can generate Point Spread Functions and can estimate the STED microscopic parameters out of STED bead images, and the Object Stabilizer (see Chapter 25 on page 143) is equipped with a special STED stabilization mode. The Deconvolution Wizard (see Chapter 3 on page 11) and Batch Processor (see Chapter 7 on page 39) can automatically stabilize STED raw images, which often contain drifts along the z direction.

This chapter describes the STED parameters and the restoration procedure to achieve optimal results with STED images. A short introduction to the STED principle is also included. A step-by-step summary of the restoration process for STED images is listed at the end of the chapter.

STED principle

The STED microscope overcomes the diffraction-limited resolution of the conventional fluorescence and confocal microscopes, by exploiting a strong non-linear effect in the depletion of excited fluorophores. In the STED microscope two laser beams are focussed on the same location. The first laser beam excites the fluorophore molecules located in the imaged volume, in the same way as the confocal microscope. The second laser, also referred to as depletion beam or STED beam, goes through a shape changing phase filter resulting in a doughnut-shaped focus. In the outer ring of this focus, where the intensity is high, excited fluorophore molecules are forced out of the excited state. This depletion effect is very non-linear, so that in effect above a certain intensity depletion rises quickly.

This results in a narrow region around the optical axis being hardly depleted, whereas beyond this region depletion increases steeply. The diameter of the depletion-free region can be as small as 25-50nm, easily resulting in a four fold increase in resolution over good quality confocal resolution.

¹“Two-photon excitation STED microscopy”. Gael Moneron and Stefan W.Hell. Optics Express 17 ; 17 (2009).

STED parameters

To achieve optimum results with the Huygens STED deconvolution it is recommended to verify the microscopic parameters of the STED image.

If the image metadata contains information about the STED parameters Huygens Essential will incorporate them into the image. Otherwise, parameter defaults will be loaded instead. Therefore, it is advisable to review the parameter values of the STED datasets.

To edit the microscopic parameters of an image right-click on the image and select “MICROSCOPIC PARAMETERS”, as explained in Chapter 3 on page 11. The specific STED parameters will be presented as in Figure 8.1:

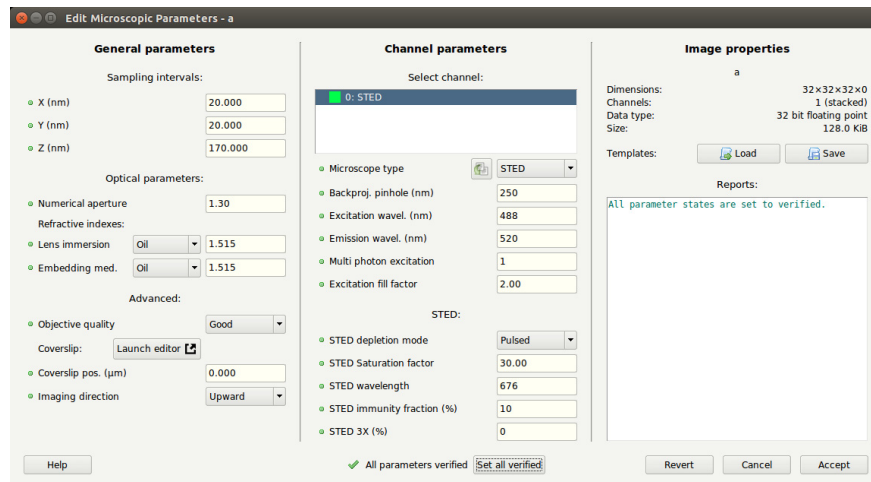


Figure 8.1: The specific STED parameters can be reviewed and, if needed, edited in the Parameter Editor.

- **STED depletion mode:** this mode determines the type of STED depletion that is being used. You can set this to Pulsed, CW (continuous wave) non-gated detection, or CW gated detection.
- **STED saturation factor:** this parameter describes how much the fluorescence is suppressed by the STED beam. The higher this factor the more fluorescence suppression off the optical axis, the more resolution. At values below 1, hardly any resolution is gained. Typical values are in the range of 10 to 50². A very large value of the saturation factor usually also implies a large resolution improvement. Because this parameter can be difficult to quantify it is recommended to use the Huygens PSF distiller (see below) to get an automatic estimation. Also, please notice that in practice the fluorescence is never completely suppressed in the whole depletion region as the fluorophore molecules can be partly ‘immune’ to depletion.
- **STED wavelength:** the wavelength of the STED depletion laser beam (nm). The STED wavelength must be a value within the range of the fluorophore emission spectrum.

²“Resolution scaling in STED microscopy”. Harke B, Keller J, Ullal CK, Westphal V, Schönle A, Hell SW. Optics Express 16 ; 6 (2008).

- **STED immunity fraction:** the fraction of fluorophore molecules not susceptible, ‘immune’, to the depletion beam. This parameter is specified as a percentage; 100% meaning that all the fluorophore molecules are immune to the depletion beam, 0% meaning that no fluorophore molecules are immune to the depletion beam. The value that should be entered here is usually between 0% and 10%. Because this parameter can be difficult to quantify it is advised to use the Huygens PSF distiller (see below) to get an automatic estimation.
- **STED 3X:** the percentage of power used in the z depletion beam. The remaining power is used for the vortex beam path. This value is read from the image metadata

The STED saturation factor can in principle be determined experimentally, see the footnote on page 48. An alternative is to use the Huygens PSF distiller which can estimate these parameters automatically from a suitable bead image.

Example settings

In Table 8.1 you find the (ranges of) normal values for the STED microscopic parameters for some STED systems.

Parameters	Microscope types			
	TCS STED pulsed STED	TCS STED CW STED CW	TCS SP8 gSTED	STED (3X) STED 3X
specified resolution (FWHM)*	70	80	50	**
typical resolution (FWHM)	60	70	40	**
excitation fill factor	0.5-1.0 (0.8)	1.2	1.2	1.2
STED saturation factor	< 60 (40)	< 40 (25)	< 40 (25)	< 40 (25)
STED wavelength (nm)	750-780	592	592	592/660/775
STED immunity fraction	<25% (10)	<25% (10)	<10% (5)	<10% (5)
Imaging direction	upwards	upwards	upwards	upwards

Table 8.1: * measured on 40nm beads; ** depending on laser and 3D depletion

Estimating STED parameters

Estimating the STED parameters is an automated process guided by the PSF distiller. Open an image of STED beads in the PSF distiller and proceed to distill a PSF as explained in Chapter 6 “The PSF Distiller” on page 33.

The PSF distiller will recognize that the image has been recorded with a STED microscope and will offer the possibility to estimate the STED parameters after distilling the PSF (see Figure 8.2).

The possibility to generate a theoretical PSF with the estimated STED parameters is offered in the PSF distiller. This can be of interest to compare the distilled PSF with the theoretical PSF.

Figure 8.3 shows the PSF distiller after estimating the STED parameters of a beads image. In this example, the estimated STED saturation factor is 6 and the (very high) immunity factor of 29%.

Estimated STED parameters will be attached to the new distilled PSF, which will be shown as a thumbnail in the main window of Huygens Essential. The parameters are

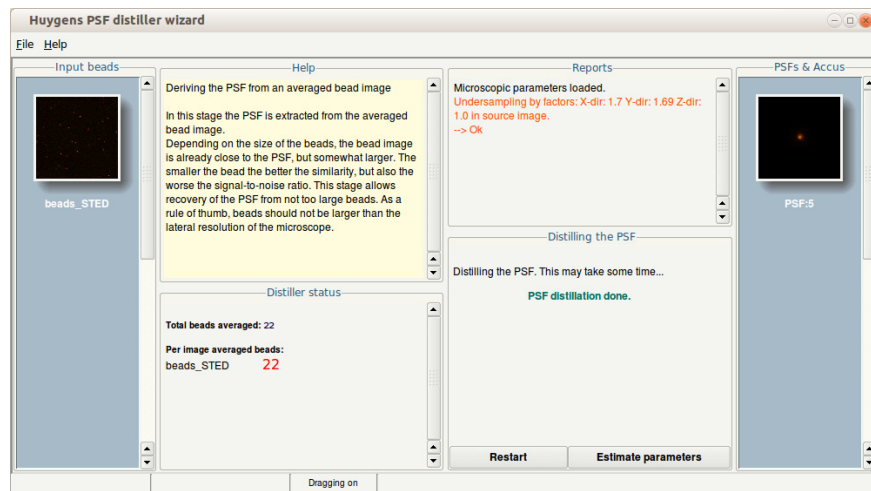


Figure 8.2: Estimating the STED microscopic parameters of a beads image automatically with the Huygens PSF distiller.

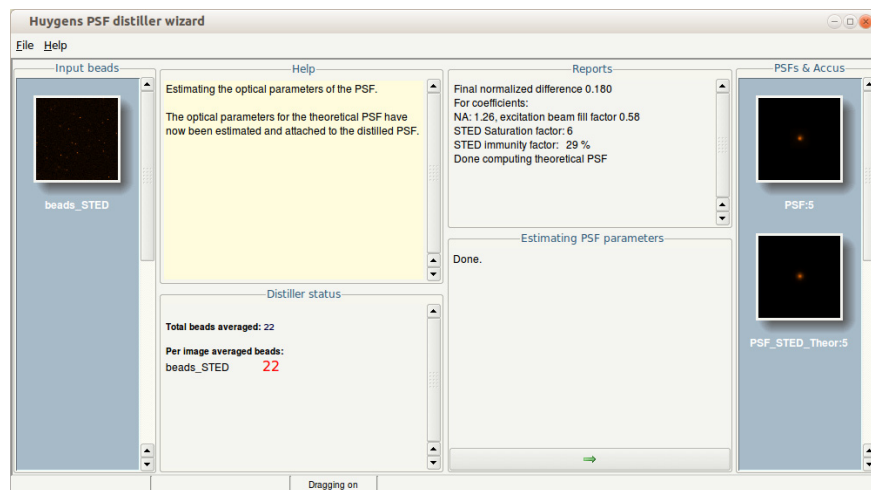


Figure 8.3: The PSF distiller showing the STED microscopic parameters automatically estimated from a beads image.

reported by the PSF distiller (see Figure 8.3) and can be reviewed by opening the MICROSCOPIC PARAMETERS by right-clicking on the PSF thumbnail (see Figure 8.4).

The microscopic parameters of the PSF can be exported as a template file within the Edit Microscopic Parameters window (see Figure 8.4). The microscopic parameter template file can, later-onwards, be applied to other STED images that have been recorded under the exact same conditions. This makes future parameter editing much easier and reliable.

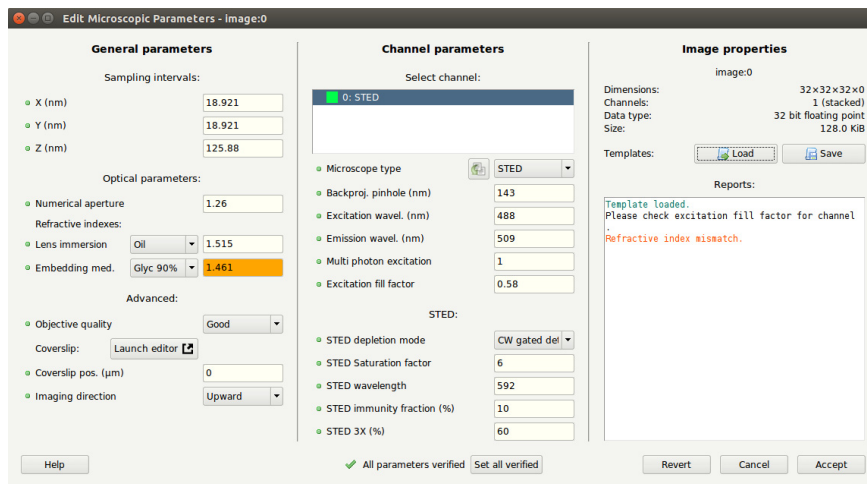


Figure 8.4: The Microscopic Parameters of the distilled PSF now contain the values automatically estimated by Huygens Essential.

STED Deconvolution

To deconvolve a STED image, open first the Edit Microscopic Parameter window (Via the main menu EDIT → EDIT MICROSCOPIC PARAMETERS). Verify the microscopic parameters or load the microscopic parameter template file that was saved earlier. The latter option will import the PSF microscopic parameters to the STED raw dataset. Both the PSF and the raw image will now have the same estimates of the STED parameters. Next, open the data set in the Deconvolution Wizard (select the image → DECONVOLUTION → DECONVOLUTION WIZARD). Select the PSF at the first stage of the Deconvolution Wizard. If no PSF is provided, the Deconvolution Wizard will create a theoretical PSF from the microscopic parameters of the image.

Continue to the next Wizard stage and proceed step by step as explained in Chapter 3 “Deconvolution Wizard” on page 11. The Deconvolution Wizard will recognize a STED dataset and can offer a stabilization option because STED image acquisition is often subjective to drift. The severity of drift in STED image data depends on how much of the STED power is used for vortex and for z depletion. Stabilization is recommended for pure vortex-based STED images. With increased z depletion in STED 3X data, drift will be more difficult to estimate but also less of an issue with respect to image quality. Huygens automatically decides whether the option to stabilize the image is needed, by taking the aspect ratio of the STED PSF into account.

To stabilize the image in the Deconvolution Wizard just click on the AUTO STABILIZE button at the “STABILIZATION” stage (See Figure 8.5). If Huygens applies stabilization a remark will be shown in the REPORT window. Note that the same level of stabilization will also be applied to the confocal channel, if present.

The stabilization stage can be skipped at the Deconvolution Wizard if the image has already been stabilized by the Huygens Object Stabilizer (See Chapter 25 on page 143.)

After this stage, complete the remaining Deconvolution Wizard stages as usual.

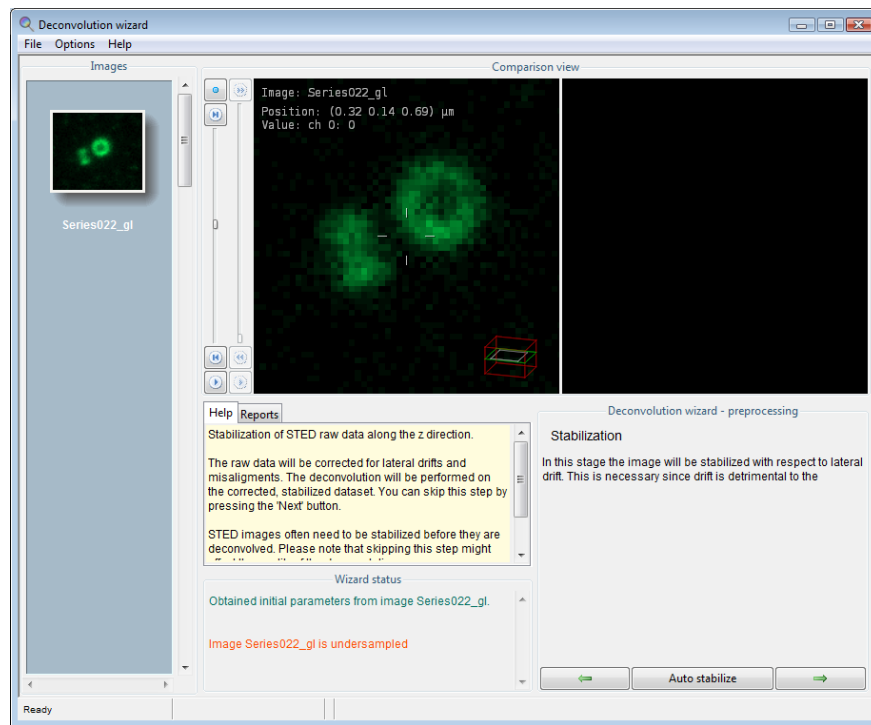


Figure 8.5: Stabilization stage for STED data in the Deconvolution Wizard.

Summary: deconvolution procedure

The STED deconvolution procedure can be summed up in the following steps:

- Load a STED beads image.
- Distill a PSF from the beads image with the PSF distiller. Allow it to estimate the STED parameters automatically.
- Export the microscopic parameters of the new PSF to a template file (via the main menu EDIT → EDIT MICROSCOPIC PARAMETERS).
- Load the STED raw dataset.
- Import the PSF microscopic parameters template file to the STED raw data (via).
- Open the raw dataset in the Deconvolution Wizard.
- Specify the PSF image in the Deconvolution Wizard and continue to the subsequent Wizard steps.
- At the “Stabilization” stage select “Auto Stabilize” if the STED raw data has not yet been stabilized.
- Complete the remaining stages of the Deconvolution Wizard as for any other image.

Chapter 9

Huygens Array Detector Deconvolution

Array Detector and Zeiss Airyscan images

Introduction

From version 18.10 of the Huygens software it is possible to deconvolve images from confocal microscopes equipped with an array detector and from Zeiss LSM microscopes with an Airyscan detector. Introduced by ZEISS¹, the Zeiss Airyscan uses a 32-element array detector organized in a honeycomb-grid. Next to supporting the various Zeiss Airyscan modes, the Huygens Array Detector Confocal Optical Option also supports other multi-detector layouts. This includes nanoSPAD detector layouts, or other generic multi-detector layouts².

In conventional confocal microscopy, the emission light is detected by a single detector. Light passing the physical pinhole is integrated to a single value and recorded in a pixel of the final image corresponding to the position of the scanner. When the scanner moves to a new position, the emission light passing the pinhole is detected and integrated, written to the next pixel, and so on, until all pixels in the image are filled. In this way, there is a one-to-one relationship between scanner and pixel positions. Replacing the single pinhole by an array of pinholes, each collecting the emitted light from a slightly different location in the sample, breaks this relationship. Now each scan position yields multiple samples which may be added to pixels of a suitably densely sampled destination image. Another way to look at this is to interpret the result of an array detector scan as an array of images, one per detector. Each of these represents an image as in the case of a single detector scan, albeit with a shifted detector position. For example, in the case of the Airyscan microscope, the raw dataset is then a set of 32 slightly shifted images. Comparing these images shows shifts which correspond to the positions of the detectors in the array.

Array detector microscopes open up new possibilities in deconvolution of the data. The Huygens Array Detector Optical Option is the solution to obtain unprecedented high resolution and contrast in confocal imaging.

¹Carl Zeiss AG

²Technology Note: The Airyscan Detector from ZEISS

The Zeiss Airyscan microscope

The Zeiss Airyscan microscope can be used in four modes. For each of these, Huygens offers high quality deconvolution:

- *Standard mode*: this is the standard confocal mode of the Airyscan system. Images acquired with this mode can be deconvolved in Huygens using the Confocal Optical Option, setting the pinhole size to the size of the pinhole used during acquisition.
- *Virtual pinhole mode*: in this mode a circular shaped subset of the detectors in the array are summed to produce a similar result as a confocal microscope equipped with a single circular pinhole of the same size. This has the advantage that a trade-off between signal and resolution can be made and revised in post-processing. Since the resulting images have mimic confocal images they can be deconvolved in Huygens using the Confocal Optical Option, setting the pinhole size to the size of the pinhole selected in post-processing.
- *Super Resolution mode*: the datasets generated in this mode contain the raw data of each of the 32 detectors. Depending on signal strength and sampling density, these can be deconvolved in various ways using the Huygens Array Detector Confocal Optical Option.
- *Fast mode* recorded with 16 wider spaced detectors organized in a column along the y direction. With these, four x-lines lines are acquired with one scan along x. The Huygens Airyscan Optical Option allows deconvolution of these images at x and y sampling densities equal to the microscope's x sampling. In this way, each scan position is converted to four pixels. In favorable conditions, this can be increased even further, yielding 16 reconstructed pixels per scan position.

The following sections of this chapter will explain how to deconvolve the images acquired in *Super Resolution mode* and *Fast mode* using the Array Detector Confocal Optical Option in Huygens.

Loading an Array Detector Image

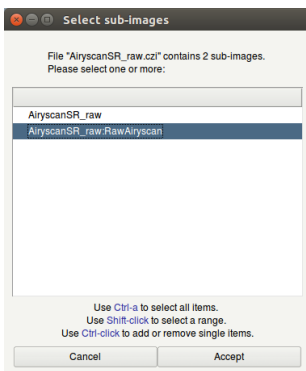


Figure 9.1: Select sub-image window

Select FILE → OPEN or the OPEN icon in the Huygens Essential main window to open the file dialog, browse to the directory where the images are stored, and select the raw-data image. The Huygens processing of Array Detector images requires the raw data as input, and therefore the complete multi-detector dataset saved from the microscope after image acquisition. In the case of a CZI file, often multiple sub-images are present in the data set. In these cases, an additional selection window will pop up showing the sub-images. Select the sub image labelled 'raw Airyscan' (see figure 9.1).

After loading the dataset, a quick survey of the data will be made to establish metadata about the detectors. This may take a brief moment.

Visualizing an Array Detector Image

All individual detector components of an array detector dataset can be visualized using the *Twin Slicer* (see also “The Twin Slicer” on page 83). To view a particular detector, select GO TO ADVANCED MODE and move to the DETECTOR tab. Use the *Detector* slider to scroll through the raw detector images.




Figure 9.2: Inspecting the multiple raw detector images in the DETECTOR tab of the *Twin Slicer*.

Array Detector image parameters

As with any other microscopy image, we recommend to first check the Microscopy parameters (see also “Verifying Microscopy Parameters” on page 13). With array detector data, the microscope type should be indicated as *Array det. Confocal* with in-between brackets the specific detector model. This detector model can be either nanoSPAD, Zeiss Airyscan, Zeiss Fast Airyscan, or generic. In case of Zeiss Airyscan data, the parameters should be correctly reported in the CZI file, but it is good practice to check them before proceeding with deconvolution.

The detector model can be changed within the Editor Window launched with the Launch Editor button located behind the parameter title *Array Detector parameters* (see figure 9.3).

This window shows additional detector parameters such as *Detector spacing*, *Aspect ratio (x : y)*, *Rotation*, and *Mirroring*. Typically, the *Aspect ratio* is around 1.0, the *Rotation* zero. *Mirroring* is often 'X'.

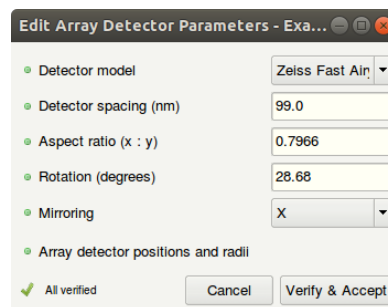


Figure 9.3: Array Detector parameters

Huygens estimates their actual value from the multiple detectors and specifies these here, and if needed, they can be adjusted. After verifying and accepting these values, the other microscopy parameters can be reviewed. For example, it is recommend to review the wavelength properties of each channel. Also, set the Embedding medium parameter to the type or refractive index value used for the specimen and adjust the Coverslip position and image direction if needed. By correctly defining these parameters, Huygens is able to compensate for spherical aberration with a depth-dependent PSF in the deconvolution procedure (See “Spherical Aberration Correction” on page 216). The back-projected pinhole value can be disregarded for array detector images. Next, all microscopy parameters can be set as verified, saved as a template, and accepted to close the parameter window.

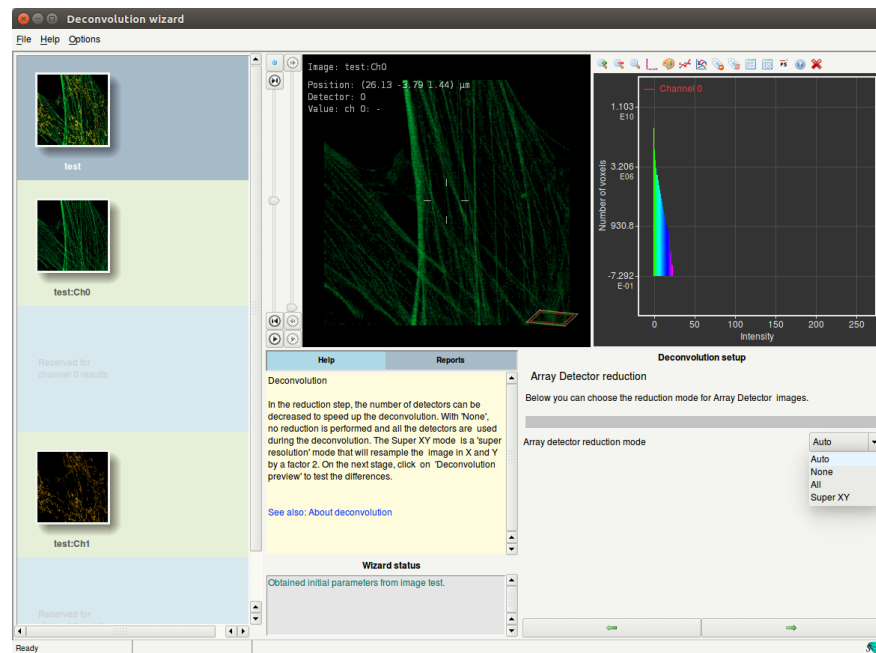


Figure 9.4: The Array Detector reduction stage in the Huygens Deconvolution wizard

Deconvolving Array Detector images

With the image parameters verified, the image can be deconvolution using Deconvolution Express, Deconvolution Wizard, Batch Processor or the Operation Window. Deconvolution Express will automatically estimate the background and signal-to-noise ratio (SNR) from the central detector image of the array since in most cases this contains the strongest signal. This SNR value is then extrapolated to all other detectors in the array. For obtaining the most optimal deconvolution results, we recommend using the Deconvolution Wizard or Operation Window and test different SNR values to obtain the best possible result.

Apart from one specific parameter for Array Detector data, the so-called “Array Detector Reduction mode”, all deconvolution parameters are similar to those for any other microscopy image. The Array Detector Reduction mode appears in the Deconvolution Wizard after the background estimation stage and is also listed in the next “Deconvolution Setup” stage. Reduction mode refers to the way detectors may or not may be combined by the deconvolution algorithm, or to the way the detector outputs are assigned to different pixels. Combining detectors may greatly reduce the time and memory requirements of the deconvolution algorithm. Assigning the detector outputs to different, smaller pixels may offer superior resolution at the cost of an increase in computing time and memory requirements. In addition, the type of Huygens Array Detector Reduction modes offered in the Deconvolution Wizard depends on whether the image is recorded in *Superresolution* Airyscan or *Fast mode* Airyscan.

To deconvolve Array Detector images with the Operation Window, select the CMLE algorithm. The “Array Detector Reduction mode” parameter is presented as the last choice in the right panel.

Depending on the image, the following Array Detector Reduction modes are available:

- Auto: Huygens selects and applies the most appropriate mode for the image.
- None: No reduction is applied which means that all detectors are treated individually. In this mode, different PSFs are calculated for each detector, and all PSFs and detector images are used simultaneously within the deconvolution procedure. This mode requires much more memory and is also much slower than other modes.

- All: All Array Detectors are reduced to one single image which is subsequently deconvolved with one PSF corresponding with the properties reduced image.
- Super XY: The detector outputs are assigned to different, smaller pixels and combined into a single supersampled image. The image is supersampled in both x and y directions.

The supersampling factors depend on the type of array detector image. A regular 32 detector *Superresolution* Airyscan image will be supersampled by a factor of 2 in both the x and y direction, while a *Fast mode* image will be supersampled by a factor of 2 and 8 in the x and y direction respectively.

- Super Y: This mode only applicable to images acquired in *Fast mode*. In this mode the fluorescence excitation distribution is elongated in y to allow the excitation of multiple x -lines in one sweep of the scanner. In the *Fast mode* Airyscan 4 lines are scanned simultaneously. In Super Y mode Huygens combines the detector outputs into a single supersampled image with voxels which are square in x and y . The resulting supersampling factor is 4 in the y direction.

Chapter 10

Light Sheet Deconvolution

Light Sheet Fluorescence Microscopy (LSFM), also known as Selective Plane Illumination Microscopy (SPIM), is a technique which combines fast 3D image acquisition with optical-sectioning by focusing with an excitation objective a thin light-sheet into the specimen. This light sheet is perpendicularly positioned with respect to the objective and detector. The detector collects the emitted fluorescent signal as a 2D image.




Figure 10.1: Principle of Light sheet fluorescence microscopy. The excitation light is shaped by a cylindrical lens in order to illuminate a thin plane of the sample. The objective is oriented perpendicular to the illumination plane to capture an image of a thin section of the sample

A three dimensional stack can be imaged by moving the specimen along the optical axis. Multiple stacks can be acquired from different angles and aligned to account for possible light loss and shading effects. The Light Sheet technique is very well suited for imaging living specimens since the illumination is restricted to the focal plane which minimizes photo-damage and improves contrast. Also no point-scanning is needed shortening the acquisition time significantly. For simplicity we will continue using only the term LSFM or 'Light Sheet' and not SPIM.

LSFM image restoration

The LSFM deconvolution option is available for Huygens Essential, Professional, and Core. Deconvolution of light sheet data has been tuned extensively by calculating the correct point-spread-function and by optimally dividing data sets in pieces: either to deal with large data sets, or to deal with the changes of the PSF over the light sheet. In this way, the

PSF is accurately calculated by theoretical modeling. Huygens accounts for the variation of the thickness of the light sheet over the field, which can be substantial for large specimens. It is also possible to use a measured PSF, which can be distilled from fluorescent bead images using the *PSF Distiller* (see Chapter 6 on page 33). Our latest special LSFM version also offers the option to correct for light scattering, which can be a serious issue in thick samples with light sheet imaging. Deconvolution and scattering correction can also be combined with fusion. See “Light Sheet Fusion & Deconvolution Wizard (Legacy)” on page 65.

Huygens LSFM deconvolution uses GPU acceleration optimally, and can fully exploit multiprocessor systems. The algorithm itself is also optimized for speed.

This chapter describes the specific microscopy parameters of LSFM images that are relevant for achieving an optimal deconvolution result.

LSFM Parameters

To get the most out of Huygens LSFM deconvolution, it is recommended to verify the microscopy parameters of the image data first. If the image meta-data contains information about the LSFM parameters Huygens will incorporate them into the image. Otherwise, parameter defaults will be loaded instead.

To review the parameters of the LSFM datasets, select the image thumbnail within the main window of Huygens. Then, right-click on the image and select “MICROSCOPIC PARAMETERS”, as explained in Chapter 3 on page 11. If “MICROSCOPIC PARAMETERS” has been selected, the microscopy parameters are displayed as in Figure 10.2.

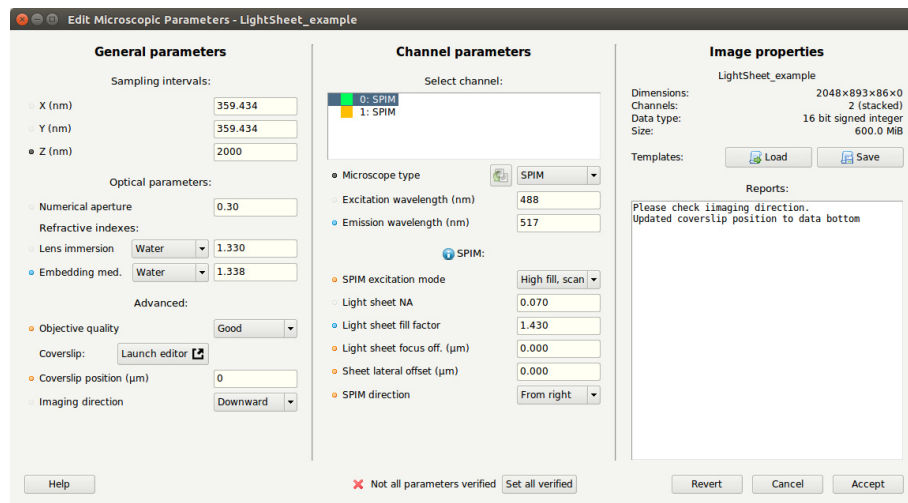


Figure 10.2: The general and LSFM-specific microscopy parameters can be reviewed and edited in the Edit Microscopic Parameters window.

There are various types of LSFM setups. These differ mainly in the way the light sheet is generated. The Huygens LSFM optical option currently supports the following light sheet types:¹

1. Light sheets with a Gaussian profile, generated by illumination from one side, either by a cylindrical lens, or by scanning a beam.
2. Light sheets with a Gaussian profile, generated by simultaneous illumination from two opposing sides.

¹All LSFM can be used as well in multi-photon mode. In those cases you need the multi-photon option as well.

3. Light sheets generated by excitation lenses that are overfilled at their entry pupils, either by using a cylindrical lens, or by scanning a beam².
4. Light sheets generated with a scanning Bessel beam, most used for samples which have a thickness below 50 micron due to its diminished blur.
5. Scanning Bessel lattice (lattice light sheet)³.

The difference between these types of excitation are important for deconvolution, because they strongly influence the shape of the PSF. Light sheets that are generated by lenses that are overfilled tend to be thinner, but at the cost of a more complex profile along the optical detection axis, and a sheet thickness that is less uniform over the image. For the most up-to-date recommendations concerning the settings for LSFM systems such as Zeiss Z1, Leica DLS, LaVision, and MuVi-LSFM (Luxendo), we like to refer you to our wiki page on LSFM deconvolution⁴

To review and edit the microscopy parameters of an image right-click on the image and select "MICROSCOPIC PARAMETERS", and fill in the correct values. More information on the LSFM-specific parameters related to the various types is explained below.

- **Light sheet excitation mode:** this mode defines how the light sheet is produced and offers the following choices:
 - Gaussian light sheet** Appropriate for LSFM system that use a cylindrical lens or a scanning beam where the fill factor at the entry pupil of the excitation is low.
 - Gaussian MultiView light sheet.** As "Gaussian light sheet", but with simultaneous illumination from opposing sides.
 - High fill factor, scanning beam.** Appropriate for a LSFM system where a scanning beam is used to form the light sheet, and where the entry pupil of the excitation lens is overfilled.
 - Scanning Bessel beam.** Use this option for a LSFM system where a scanning Bessel beam is used to form the light sheet. If the multi-photon excitation parameter is adjusted, Huygens is able to account for multi-photon illumination as well.
 - Scanning Bessel lattice.** Use this option for a LSFM system where a scanning Bessel lattice is used to form the light sheet. If the multi-photon excitation parameter is adjusted, Huygens is able to account for multi-photon illumination as well.
 - High fill factor, cylinder.** Appropriate for a LSFM system where a cylindrical lens is used to form the light sheet, using a high fill factor at the entry pupil.

The following parameter appears only in the first two modes ("Gaussian light sheet" and "Gaussian MultiView light sheet").

- **Width of the Gaussian sheet (microns):** this parameter specifies the width (thickness) of the sheet in micrometer. The width of the Gaussian profile is defined as the distance between the two points where the value is equal to $2/e^2$ of the peak value.

The following parameter appears in the last four excitation modes ("High fill factor, scanning beam", "Scanning Bessel beam", "Scanning Bessel lattice" and "High fill factor, cylinder"):

- **Light sheet NA:** the numerical aperture of the excitation lens. Note that the NA of the detection lens is specified under the (Main) Optical Parameters. If the NA of the excitation lens or its fill factor are unknown, but the effective NA is known, proceed as follows:
 - Set the light sheet NA to 4 times the effective NA

²<https://www.svi.nl/ExcitationFillFactor>

³For deskewed LLS you need to use the Object Stabilizer prior to using the light sheet and fusion bundle.

⁴<https://www.svi.nl/SPIMDeconvolution>

- Set the fill factor to 0.25

The following parameter is only applicable for the "High fill factor, scanning beam" and "High fill factor, cylinder" setting.

- **Light sheet fill factor:** this value indicates the Fill Factor of the excitation (LSFM) lens. The Fill Factor is the ratio between the beam width and the diameter of the objective pupil. Huygens default value is "0.5", meaning that the illumination beam is half as wide as the objective pupil of the LSFM lens.

The following parameters are active for all modes:

- **Light sheet focus offset (microns):** this value defines how far the light sheet is located below (negative value) or above (positive value) the focal point of the detection lens. This parameter cannot be read from the image file. Default value is "0", which is the optimal value in a well-aligned system.
- **Sheet lateral offset (microns):** specifies the distance between the focal point of the excitation lens and the optical axis of the detection lens. A negative value indicates that the "middle" of the light sheet is shifted towards the excitation lens. A positive value indicates a shift away from the excitation lens. Generally, this value is not read from the image file. The value is equal to zero in a well-aligned system. However, this may be different in cases where the image is cropped, or where the optical setup allows you to change this, or where it is simply not well aligned.
- **Light sheet direction:** the parameter defines where the excitation objective is positioned with respect to the detection lens.

Scatter parameters

From version 17.04 onwards, Huygens is able to correct for scattering in your sample by adjusting the PSF according to the scattering parameters. You can fine-tune the scattering correction via the Scatter parameter launch button. you can adjust:

- Scatter model (uniform exponential, uniform Gaussian, and 1D X-direction exponential).
- The length of the free path (in micron) of the emission light, and the percentage of the scattered light with respect to the non-scattered direct light.

General parameter settings for LSFM

In the `General` parameter column, it is assumed that the detection lens has an "Upward" position with respect to the image (stack). You can check whether this is correct by using the `LAUNCH EDITOR` button in the `MICROSCOPIC PARAMETER` window. If necessary, the position of the detection lens with respect to the image can be adjusted here. See also "Setting the Coverslip Position" on page 221.

In many LSFM experiments, the refractive indices of the lens and sample medium match. If you set these parameters correctly in Huygens, the coverslip position will be ignored and as it should, no spherical aberration correction will be applied.

LSFM deconvolution

After verifying the general and LSFM-specific microscopy parameters, the image is ready to be deconvolved as any other image via either the *Deconvolution Wizard* or *Batch Processor*. As most LSFM images are multiview these can also be deconvolved within the Interactive 3D Fusion & Deconvolution, where the deconvolution is combined with fusion of the multiple views. See “The Interactive Fuser” on page 71. You can find the Interactive 3D Fusion & Deconvolution under DECONVOLUTION → INTERACTIVE 3D FUSION.

Chapter 11

Light Sheet Fusion & Deconvolution Wizard (Legacy)

In many Light Sheet microscopy and SPIM setups the option exists to rotate the sample to image it from different directions. By combining these views into a single image stack, a superior image can be obtained. Since Huygens 19.10, the Huygens *Light Sheet Fusion and Decon Wizard* is superseded by The Interactive Fuser. The Huygens *Light Sheet Fusion and Decon Wizard* offers a solution for the deconvolution and fusion of multiple views of the same specimen into a single superior image. By integrating the deconvolution into the fusion process, improved results can be obtained, while significantly simplifying the post-acquisition work-flow. The Huygens *Light Sheet Fusion and Decon Wizard* offers support for all commonly known Light Sheets setups.

Starting the Fusion and Deconvolution Wizard

The Wizard opens as a separate window in Huygens when selecting TOOLS → LEGACY → LIGHT SHEET FUSION & DECON WIZARD. Follow the wizard to select the images to fuse, to set the deconvolution and fusion parameters, and finally to perform the fusion operation. First however, you may want to change options within the OPTIONS menu in the task bar of the Wizard. These options may help you to save memory, increase processing speed, and to limit the resulting image file size. Note that these settings also may change the quality of the fusion and the resulting image.

Options Menu

The fusion of the input images is done by deconvolving each input image, transforming it by rotating and shifting, and then fusing it into the final result. Normally, the intermediate deconvolved and transformed images are not kept to save memory. To save the intermediate images, select the "Keep intermediate images" entry in the OPTIONS menu.

When fusing images the result can take up substantially more memory. By default, the output image is cropped to the size of the input images to save memory because the size of the result image may need to be increased to fully contain the input images after rotation and shifting. To retain the full-size result, de-select the "CROP OUTPUT" entry in the OPTIONS menu.

By the default the OUTPUT DATA TYPE of the result image is set to signed 16 bits, even though the results are internally calculated in a higher 32-bit floating point precision. If you prefer to use unsigned 16 bit, you have the option to select it in this menu. To keep the full precision, select the "Floating point" entry. To save memory, you can also select the "8-bit" entry in the same sub-menu.

Fusion settings menu

To obtain the final fusion result, the transformed inputs are combined voxel-by-voxel into a single 3D image. The FUSION SETTINGS sub-menu under the OPTIONS menu lets the user select how the voxels are combined according to one of two modes:

- **Weight pixels with the selection criterion value:** The value of each voxel is a weighted sum of the input values.
- **Take pixel with the maximum selection criterion value:** The output value of each voxel is the maximum of the input values. The input values may be the intensity values of the voxels that are combined, which leads to a simple weighted sum of the inputs, or to the selection of the voxel with the maximal intensity.

To include more local structural information, the weighted sum or maximum operation can instead be done using a local information criterion which is derived from the intensities of the voxels and of the voxels in its immediate vicinity. The Huygens Fusion Wizard implements two methods for calculating the local information criterion which are local variance and local entropy. The selection criterion can be changed within the FUSION SETTINGS menu:

- **Pixel values:** uses no local information, but uses the voxel intensities directly.
- **Local Variance:** uses the local variance criterion. In each input image, the variance of the intensities of the voxels in the immediate vicinity of each voxel is calculated.
- **Local entropy:** uses the local entropy criterion. In each input image, the entropy of the voxels in the immediate vicinity of each voxel is calculated. This approach is theoretically more accurate than local variance, but in practice the local variance reaches a similar level of quality using much less computational resources.

Selecting input images

The input images for the fusion process can be added by pressing either the ADD IMAGES button, or the ADD IMAGE FILES button. The ADD IMAGES button presents a list of the images that are already loaded in Huygens and that can be selected as inputs.

Instead of using images that are already loaded in Huygens, the ADD IMAGE FILES button can be used to select images that reside on disk. These files will be loaded on demand, so they are only loaded when they are needed. This is advantageous if multiple large files need to be processed that may not fit in memory simultaneously.

The REMOVE INPUTS button can be used to remove one or more images from the inputs.

The inputs appear as thumbnails in a column at the left side within the *Fusion and Deconvolution Wizard* window. They are shown in the order they have been added, and they will be processed also in this order. The inputs can be reordered by selecting the thumbnails with the mouse and pressing the arrow buttons below the *Images* panel with the thumbnails. (Use Shift-click to select a range of images, and Control-click to select or unselect individual images.) Selected inputs can also be removed using the thumbnail context menu (Figure 11.1 on page 67).

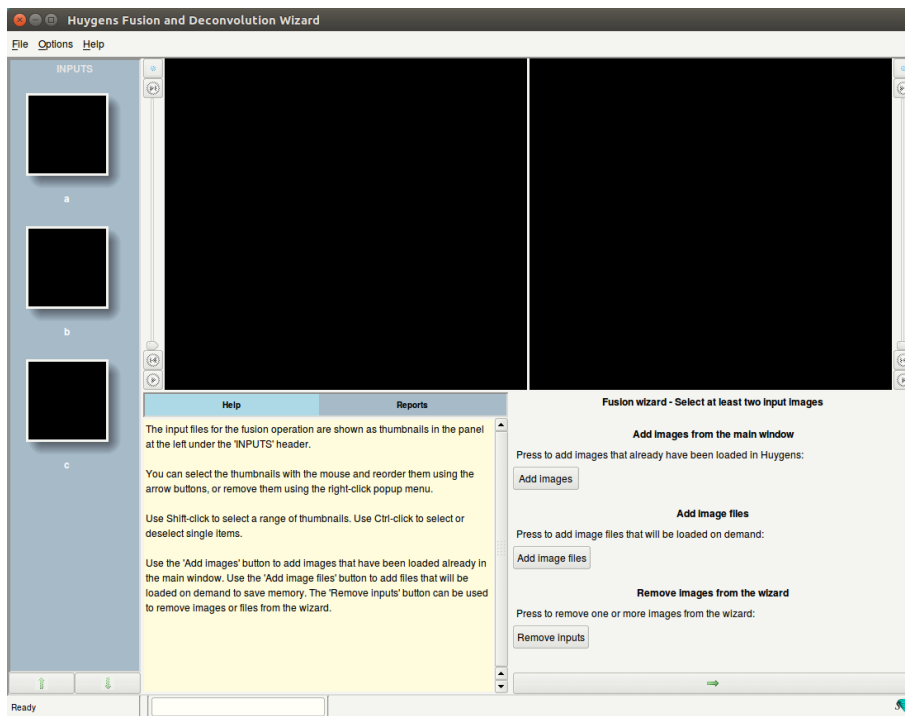


Figure 11.1: The *Fusion and Deconvolution Wizard* main window after having selected thumbnails a, b, and c with the ADD IMAGES button.

Deconvolution settings

After entering the inputs, the wizard continues with the deconvolution settings stage. The deconvolution parameters can be set by pressing the SET DECONVOLUTION PARAMETERS button. This opens the *Huygens deconvolution template editor*, which allows the user to set the parameters for the deconvolution or load a deconvolution template. Each input image will be deconvolved using these parameters prior to the fusion. If no deconvolution parameters are set, the deconvolution step will be skipped. If a template is selected a DISABLE button will appear for canceling the use of the selected template. The microscopy parameters of the input images can be edited at this stage by pressing the button SET MICROSCOPE PARAMETERS. This will open the *Huygens Microscopic Parameter Editor* (“Verifying Microscopy Parameters” on page 13).

Transformation and resampling parameters

Setting the type of the transformation

The next step in the wizard is used to determine how the input images should be transformed before they are fused into a single image. Three types of transformation can be selected:

- **Shift:** The images are aligned by shifting in the x , y and z direction. This option should typically be selected if no rotation is involved and the position of the structure in the image is just shifted. This option allows to overlay/fuse views from the left and from the right in microscopes that have two excitation light sheets but no rotation of the detection objective lens with respect to the sample.
- **Rotate/shift:** The images are rotated around either the x or y axis and aligned by shifting in the x , y and z directions. This option should typically be chosen if the sample rotates with respect to the detection objective lens between the different views in the acquisition.

- **None:** No rotation or shift transformations are applied. This option should be selected if the input images are already aligned to a satisfactory level and only the overlaying/fusion step of combining them into one image is desired.

Setting resampling options

Microscopy images are frequently sampled differently in the z direction: the voxels have a different size in the z direction than in the x and y directions. In the individual images, the larger voxel size in z is usually acceptable because the optical resolution of microscopy images is generally poorer in that direction. However, in the case of the fusion of rotated images, the final fused image may have an improved z resolution and it may be more appropriate to have an equal voxel size in all directions. The *Fusion and Deconvolution Wizard* offers three options to equalize the sampling:

- **Do not equalize:** The sampling of the input is not changed.
- **Scale the z dimension to match xy sampling:** The 3D image is resampled along the z direction to make the sampling equal to the xy sampling. This usually leads to an image that is sampled on a finer grid along z . This is generally the option that should be selected if the fusion improves the z resolution of the final image compared to the inputs.
- **Scale the x and y dimensions to match z sampling:** The 3D image is resampled in xy to make the sampling equal to the sampling along z . This usually increases the sampling distance in xy , which is generally not desired. However, it leads to smaller images that may take less computational resources to fuse.

Optionally, a calibration factor can be entered for the sampling along the z direction. This is useful if the sampling along the optical axis (z), as it is set in the parameter editor, is slightly off. This would adversely affect the result because the rotated images would be scaled incorrectly in the fusion process. To compensate for this, select “Apply a calibration factor to the sampling rate in z ” and enter a value to multiply the sampling rate with. This will only affect the sampling rate in z in the result image, not its final size.

The correct value for this calibration factor needs to be estimated empirically by visualizing the result. It is also possible to run a simple check on the calibration by selecting CHECK z CALIBRATION in the TOOLS menu. This will print an estimate of the necessary calibration value in the report window. If this value differs from 1 (by one percent or more), this indicates that another run with a different calibration value may be necessary. The printed value can be used as an approximate estimate for the calibration value.

To facility quick processing it is also possible to downsample the 3D inputs in the xy planes, or along the z direction: select the desired downsampling option, and enter the downsampling factor to reduce the number of voxels in the xy , or z direction.

Setting the rotation parameters

If the Rotation transformation is selected, the next stage of the Wizard is used to set the parameters needed to rotate the input images to the correct orientation. Rotation is assumed to occur around one main axis, which should be selected by the user. The amount of rotation that is needed can be determined in several ways:

- **Estimate from the data:** The angle of rotation is estimated from the data itself. This requires no further prior knowledge, but it is the most time-consuming option.
- **Increase with a fixed step:** With this option, it is assumed that each input image is rotated with a fixed amount compared to the previous input. For instance, four images evenly acquired over one full rotation correspond with a step of 90 degrees. The

rotation step must be entered, but if it is not known exactly, an uncertainty can also be entered. For instance, entering 90 +/- 5 degrees indicates that the step is between 85 and 95 degrees. The exact angle within this range will be determined automatically during the fusion. Note that if this option is used, the input images should be ordered accordingly.

- **Enter manually:** With this option the angles for all images can be entered manually. For instance, four views acquired over 360 degrees at equal steps correspond to "0 90 180 270". As in the previous option, an uncertainty in the rotation angle can also be entered.

In the last two cases, it is necessary to know in which direction the rotation occurs. Select "clockwise", or "counterclockwise" to set the rotation direction for your instrument accordingly.

Finally, because the true axis of rotation may not be exactly aligned with the selected axis (x , y or z), the orientation of the image can be fine-tuned in 3D. To enable this, check the FINE-TUNE THE 3D ROTATION ANGLES option and enter the range over which the axis orientation should be optimized.

Channels and Time frames

Selecting the channels used by the registration

If the input images are multi-channel images, the next stage of the wizard shows a channel selector. The user can select the channels that will be used by the *Fusion and Deconvolution Wizard* to calculate the exact alignment parameters that are needed for the fusion. The user can use this option to select one or more channels that provide the best information. For instance, sparse images with many features are more likely to help in finding the correct alignment parameters than smooth images with few features. Selecting channels that are rich in features may improve the fusion result.

Selecting the time frames used for the registration

If the inputs are time-series the next stage of the *Fusion and Deconvolution Wizard* will allow the user to choose if each time-frame should be aligned separately, or if the alignment parameters can be derived from a single time-frame and applied to all of them. By default the *Fusion and Deconvolution Wizard* aligns each frame separately, but if the alignment parameters remain the same over time, this is not necessary. In this case, unchecking the REGISTER EACH FRAME INDIVIDUALLY option may speed up the fusion process significantly.

Starting the deconvolution and the fusion

After the parameters have been set, the fusion process is started by pressing START THE FUSION. The inputs are processed sequentially, loading them on the fly if necessary, applying the deconvolution and image transforms, and finally fusing them into a single image. This process can be time-consuming and can be interrupted using the STOP THE FUSION... button.

After the images are fused, and maybe also deconvolved, you can click on buttons to restart or quit the wizard, or you can click on DONE, which will place the image in the main Huygens window with suffix `_fusion` or `_decon_fusion`.

Chapter 12

The Interactive Fuser

The interactive Huygens Fuser is an option that can be used for the fusion and restoration of multi-view images. Such images are typically acquired using light sheet fluorescent microscopy (LSFM), although in principle any multi-view data set can be fused. Light Sheet microscopes often acquire images by illuminating an object with one - or sometimes more - sheets of light. The light is captured using objective lenses in the direction perpendicular to the sheet of light. LSFM data frequently suffers from stripes and shading issues due to photon scattering and sample-dependent photon absorption. These problems are commonly addressed by rotating and viewing the sample from different directions. The different views can then be fused into a single superior image where all parts of the specimen are imaged optimally.

Like images from all other optical microscopes, LSFM data is subjected to blurring induced by the diffraction-limited optics of the microscope. Deconvolution can correct for this distortion and is critically dependent on the inherent resolution of the data and on the signal-to-noise ratio. Because light-sheet data provides a good compromise in this regard, it is well suited for deconvolution. Thus, deconvolution combined with multi-view reconstruction, allows a significant improvement/restoration of LSFM data. The INTERACTIVE FUSER & DECON offers these two options in one user-friendly tool.

If you plan to perform both deconvolution and fusion, then the deconvolution should be done on the raw input images before they are fused. The reason for this is that the model for the LS point spread function has been designed for situations where a 3D z -stack is acquired using a detection objective and a specimen whose relative orientation is fixed.

In this chapter you can read about the following topics

- Starting and opening images in the INTERACTIVE FUSER & DECON Window.
- Checking the image orientation and applying preprocessing steps.
- Manually align the multiple views.
- Select the options for optimal automatic registration of all views.
- Optional customization of the method for fusing the aligned views.

Furthermore the extra possibilities for handling multi-channel and time series images are shown.

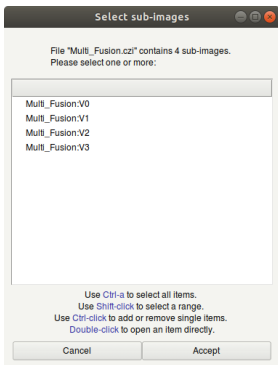
Opening The Interactive Fuser & Decon Window

This section explains how the Interactive Fuser & Decon Window handles images during the opening stage. Light sheet imaging involves a large field-of-view volume, which results in large files. This is exacerbated by having multiple acquisitions of each specimen. To prevent having all input data loaded into the system's memory simultaneously, the INTERACTIVE FUSER & DECON Window has two modes of opening images.

1. Images that are already open in Huygens

This first mode is similar to how the rest of Huygens handles images and is useful if there is sufficient memory available to perform deconvolution (if desired) as well as fuse the images. When all images that are to be fused are available from the Huygens main window, select their thumbnails simultaneously (by using the Ctrl or Shift key while clicking on them with the left mouse button) and then right click on one of them. In the pop-up menu, select the option INTERACTIVE FUSER & DECON. This option is also shown in the main menu bar in Huygens under DECONVOLUTION → INTERACTIVE FUSER & DECON.

2. Files from disk



This second mode works much like the HUYGENS STITCHING & DECON WIZARD and is useful if memory-use is a constraint. Start by opening the Interactive Fuser & Decon window, without selecting any thumbnails. Selecting the option from the main menu bar in Huygens under DECONVOLUTION → INTERACTIVE FUSER & DECON opens the new window together with a file open dialog window, where you select the file names of the input data on disk.

Image files that contain multiple views as sub-images are also supported, for example .czi or .lif files. Opening them shows a sub-image dialog screen as shown in figure 12.1. Select the views you would like to fuse and press Accept. After finishing, the window starts and down-sampled thumbnail and overlay images become visible. These images are used during the interactive alignment of the different views.

Figure 12.1: Dialog for selecting sub-images of multi-view image files.

A global overview of The Interactive Fuser window

This section gives a global overview of the parts that constitute the window. The next section describes the components in detail.

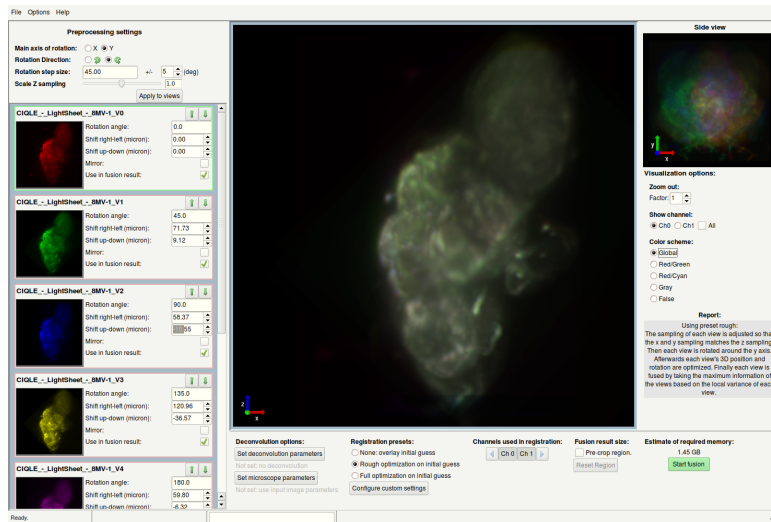


Figure 12.2: The initial state of the window after opening the data set.

The top-left in figure 12.2 shows the preprocessing settings. Start here to indicate the rotation angles of the multi-view images and also possibly diagnose z-sampling issues with the slider.

Directly underneath the preprocessing options, there is a list of all the views. Each view has its own set of 3-dimensional transformation parameters to move and rotate the view to manually overlap in the main window on the right. More about the optimization of this alignment is discussed in Registration & fusion settings on page 77.

In the center of the window, a large preview image of the fusion result is shown in the form of a maximum intensity projection (MIP) along the main axis of rotation. This image is the main image to use for interactively overlaying the different views. Each view is rotated and shifted views after with the transformation parameters as set with each view. Left-clicking and dragging the mouse cursor displaces the last selected view. Shift+left-clicking and dragging will rotate the last selected view around its center.

On the right, there is a smaller MIP along the z-direction. This image remains static during the interactive overlaying of each view. During the automatic registration this image is updated and shows the transformations applied to each view.

Beneath this smaller image are the visualization options. Change which channels are shown and change the color scheme from the default red/green to red/cyan or gray. These settings do not influence the fusion result, they are merely for visualizing the views. It is also possible to zoom out if there is not enough room to fit everything. Be aware that zooming out does add extra volume to the fusion result image. To help with memory management, crop away sections of the image before starting the fusion process by checking PRE-CROP REGION.

Finally, at the bottom there are several options pertaining to the deconvolution and automatic registration and fusion steps. From left to right, there are the microscopic and deconvolution parameters (for information on light sheet deconvolution see Chapter 10). Next, there are registration presets, a pre-cropping option to indicate how large the result image should be, and finally the estimated required memory with the START FUSION button.

The next sections contain a more in-depth description for each of these aforementioned segments and provides a step-by-step description to perform optimal fusion with your images.

Preprocessing settings

As mentioned in the previous section on page 73, the **Preprocessing settings** are used to describe the specific configuration of your measurement. These settings effect all views and can be found on the top-left in the Huygens Fuser. Set these before starting the visual interactive part of overlaying the views.

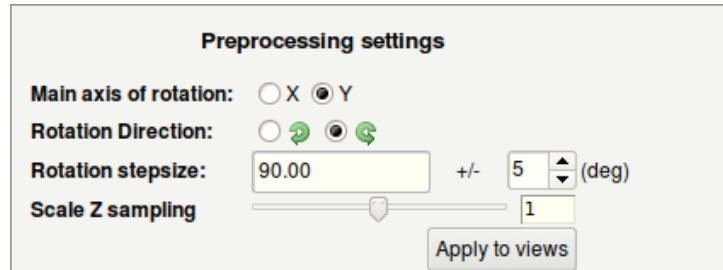


Figure 12.3: The preprocessing settings detailing the rotation and orientation of each image.

The first option **Main axis of rotation** determines around which axis the preview images are rotated. Note that the z -axis is not shown, as this is interpreted as the direction of the optical path, which is perpendicular to the light sheet. Rotating around the z -axis would not result in more information in the fused image. Most of the times the axis of rotation is the y -axis of an image. When the main axis of rotation is not known (e.g. when the Huygens user is using this window for the first time with data from an unknown type of microscope) this setting is one of the first things to determine to get the multiple views overlaid.

The second option **Rotation direction** denotes the direction in which the views are rotated. This has effect on both the small previews and the large canvas in the middle of the window.

The third option **Rotation step size** denotes the degrees rotated in between each acquisition. This setting defaults to the 360 degrees divided by the number of acquisitions loaded in the window. The extra “+/- ... (deg)” addition to the step size is used in the automatic registration that performs fine-tuning on the transformation to limit the deviation from the given value.¹

The fourth option is the **Scale Z sampling** with a slider and a real-valued number. The default is to scale by a factor of 1.0, which takes the aspect ratio of the voxels *as-is* from the input images. If the raw image does not contain correct meta data about the sample sizes, then each of the maximum intensity projections are likely compressed or stretched out in one direction. This setting can help to diagnose this problem, and temporarily fix it. The preferred solution is to get the correct meta data beforehand and then start fusion. This can be done by either saving in a different file format with the acquisition software to store the images, or by using a microscopic parameter template with the correct settings.²

¹If the used light sheet microscope does not perform rotations, but only images the biological sample from left and right, then put both the rotation and maximum deviation to 0°.

²Some of the registration methods can also fine-tune the aspect ratio, but it needs to be correct to within 10% to have a good enough initial overlap.

Input view specific settings

At the middle- and bottom-left of the window are previews of the individual input images with settings that can be changed for each of them (see Figure 12.4). Above each image is the name of the image in the original data set. Use the up and down arrows to change the order of the fusion. If the first preprocessing settings are set correctly, they should all have the same orientation but not necessarily the correct left-right and up-down shifts.

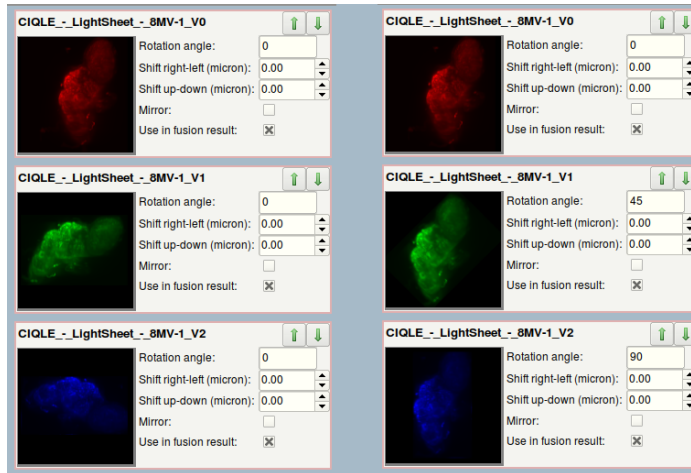


Figure 12.4: Example data set with 8 views, showing the first 3. Left: The input views before rotation angle is set. Right: After the rotation is set. Note that there are still relative horizontal and vertical shifts.

Each of these images has a check box **Use in fusion result**. Checking this box shows the thumbnail image in color and the image is included later on when the fusion process starts. The images that are not (yet) included show up in a gray color scheme. By default only the first two views are active. The rest is not for ease of manipulating the relative shifts. For more information on the color schemes and the interactive overlay, see section Interactive alignment on page 76).

Each of the input images has a **Rotation angle** setting, the initial value of which is determined from the rotation step size from the preprocessing settings (see subsection Preprocessing settings on page 74). The rotation angles can be set individually, though. Note that pressing the APPLY button in the preprocessing settings would overwrite any manually set rotation angles again.

Use the **Shift right-left** and **Shift up-down** entries or the left mouse button to change the position of one of the input images with respect to the others in the overlay preview. The recommended way to accomplish an approximate overlay, is by dragging with the mouse cursor (see subsection Interactive alignment on page 76).

The **mirror** checkbox is used to mirror the individual input images in the plane of the thumbnail before they are used in subsequent steps. This transformation cannot otherwise be described using angles and changing the main axis of rotation. It should normally not be needed when the microscope has a single light sheet and a single camera, but it can become useful in some cases.

Now that each view option is explained, we can move on to the large canvas section containing the overlay of all active views. In the next section we show how to intuitively move around the different views in order to manually and then automatically fine tune the overlap between the images.

Interactive alignment




Figure 12.5: Visualization options.

The large overlay preview image is intended to allow the different views to be shifted perpendicular to the main axis of rotation in such a way that all the views are overlapping when the manipulation is done.

Important to note is that the different views that are included are distinguishable by color. The different color schemes can be selected in the Visualization options on the right of the window (see Figure 12.5). The **Global** color scheme uses the same color encodings for the different views as the Global setting in visualization tools for multi-channel images in Huygens, i.e. red for the first, green for the second, blue for the third channel, etc.

The default second color scheme **Red/Green** displays the currently selected channel in green and all other participating channels in red. Compared to Global with just two channels, the brightness and contrast are somewhat enhanced as well. This is the recommended color scheme to start, and then selecting the images two at a time. A **Red/Cyan** option is also included for people who have trouble distinguishing red and green colors.

The color scheme options **Grey** and **False** do the same thing for input views in Huygens Interactive Fuser as they do for visualizing channels elsewhere in Huygens.

One of the input images is the currently selected, or *active* view. The active view is shown in the list of views with a green border around its transformation settings. In the Red/Green color scheme mode, the MIP is shown in green. Activate a view by clicking the thumbnail, or by adding it to the fusion result with the checkbox. Only the active view moves with respect to the others by dragging the mouse cursor with the left mouse button in the large preview window. The goal is to let each added view overlap as much as possible with the already positioned views.

The recommended sequence to manually align the different input images is:

1. Set the preprocessing settings to their correct values, such that the rotation and the sizes of the object match up. We deal with possible shifts in the next step.
2. Start by aligning (shifting) the second (green) to the first (red) view in Red/Green mode. Drag the green image in the large preview until it overlaps with the red.
3. Temporarily deselect the first image (by un-checking the “Use in fusion result” setting) and select the third image (either by checking “Use in fusion result” or clicking on its thumbnail image).
4. Repeat these two steps until all images have been aligned two at a time.
5. Now select the “Global” color scheme and select *all* input images by checking their “Use in fusion result” setting or clicking on the gray images’ thumbnails. Check that the overall alignment is good. If not, figure out which image is misaligned. Click on the thumbnail to make it the active view. Then realign it.

Registration & fusion settings

The registration and fusion settings describe what happens to each input image after the orientation and manual alignment are fixed. The following settings decide what happens once the START FUSION is pressed and the processing starts.




Figure 12.6: Registration and fusion options for deconvolution and presets for automatically aligning and fusing the views.

Deconvolution

There are two buttons for loading templates, labeled “Set deconvolution parameters” and “Set microscope parameters”. Either they should not be used at all, in which case no deconvolution is performed prior to registration and fusion, or they should both be used. Like processing multiple images in the Batch Processor, the deconvolution and microscopic templates (stored on disk as `.hgscd` and `.hgsm` files respectively) enable the same deconvolution procedure to be applied to all of the input images. If templates are specified, then each input image is first deconvolved according to the settings and the deconvolution result is used in subsequent steps. See “Verifying Microscopy Parameters” on page 13 for more information on how to use microscopic parameter templates. See “The Deconvolution Stage” on page 20 and the *Restoration Parameters* article³ for more information on deconvolution templates.

Optimization

Huygens Interactive Fuser supports three presets for registration:

- None: Overlay initial guess
- Rough optimization on initial guess
- Full optimization on initial guess

The **None** method, specifies that no fine-tuning on the manual alignment should be attempted at all. The input images, optionally deconvolved, are all fused with the transformation parameters as set in the window. The voxel grid of each image is adjusted to an isotropic grid that matches the largest sampling size. This is a fast method that requires little memory. It can be used to quickly check if the manual alignment in the lateral direction is done correctly. Any necessary shifts along the rotation axis are not applied with this method. This is necessary in some cases. Then, it is advised to use either the second or third method.

The second and third method perform **Optimization on initial guess** and employ an iterative optimization algorithm that requires an initial alignment that is reasonably close to the optimum value. This method can optimize multiple transformation parameters at the same time. If this method is used, then be sure to first align the different views approximately in the large preview window if they are misaligned. The optimization starts by testing the shift of the images along the main axis of rotation, a parameter that you cannot set graphically in the window. Then, by default, the algorithm uses the following 7 parameters to optimize the alignment of the views in 3D:

³<https://www.svi.nl/RestorationParameters>

- 3 angle parameters to describe an arbitrary rotation in 3D
- 3 shift parameters to describe a shift along the x , y , and z axes of the rotated image.
- a compression or extension factor to slightly change the axial/lateral aspect ratio of the sampling sizes

The algorithm iteratively tries to find a better set of parameters that maximize the degree of 3D overlap between the to-be-fused view and the views that have been fused so far. The large image preview and the small image preview are updated to show the current estimate during the optimization process, and the status bar of the window displays the current degree of overlap. The ranges for these parameters can be tuned. See Customizing registration & fusion options on page 79.

The “Channels used during registration” control is a list of channels that are present in the individual input images. They can each be individually enabled or disabled (“all” and “none” being shortcuts to having all or none of them selected). The fusion and the result still use all of the available channels (if you do not want that, you can split the raw data and only keep the relevant channels). The registration uses only the selected channels, which can be used to base the degree of overlap that is maximized on only those channels that contain certain biological features, or perhaps fiducial markers. If the channels resulted from the simultaneous or sequential acquisition of the same specimen where there should be the same transformation between the views in each of the channels, the choice “use all of the available channels” is usually optimal, and therefore it is the default setting.

Customizing registration & fusion options

By opening the customizable settings with the CUSTOM SETTINGS button, the differences between the presets are shown. For example, change the optimization to FULL OPTIMIZATION ON INITIAL GUESS by clicking on the button on the left-hand side in the custom settings window. This changes, among other things, the voxel sizes to a finer grid so that the Z sampling is adjusted to the XY sampling values. Also from the extra features, a multi-resolution scale space fit is applied during the registration, as well as performing a background subtraction to optimally overlap the biological object. These options both do require more memory.

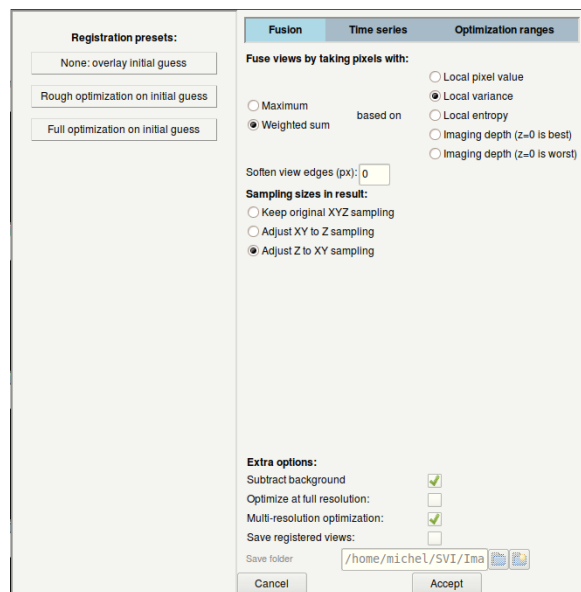


Figure 12.7: The custom fusion settings with presets on the left. Determine how the voxel information from all views are combined. Choose extra methods for the automatic registration phase.

Fusion settings

When it comes to combining the information of the aligned views, the Huygens Interactive Fuser & Decon window supports the same two methods of fusion as its predecessor, the Fusion and Deconvolution Wizard. As is also described in its chapter on page 65, these two options are to take either the **Maximum** or a **Weighted sum** of voxels in both views, which can be combined with the types local information around each voxel as described in Table 12.1

In the case we use **Maximum based on Local Pixel Value**, then at each individual voxel location the fusion algorithm checks whether the already fused image or the new to-be-fused view has a higher voxel value at that location. It then uses the value that is maximum out of the two. If the shading effect makes each input image dimmer at the side that is furthest away from where the light sheet is coming from, then selecting the maximum of the local pixel value already mostly selects the locally “best” image.

However, if the image does not become less bright, but just contains less detail (more blur) towards the side that is furthest away from the light sheet, then it can be beneficial to choose the view that provides the highest local variance or highest local entropy at that location. The entropy is computationally more expensive, but theoretically better, while the variance is less intensive to compute and gives almost equal results in practice (both still require extra memory to store the extra variance or entropy information of the images).

Where the “Maximum” method chooses either one or the other voxel value from the two inputs (the fusion result so far and the to-be-fused image), the “Weighted sum” image gives a

Option	Description
Local Pixel value	Simply use the voxel value as is, when determining the maximum or weighted sum of the two views.
Local Variance	Calculate the variance around the voxel value and either select the voxel from the view with the maximum variance (most information), or weight the voxel values in both views by their respective local variance.
Local Entropy	Calculate the entropy in the immediate vicinity of the voxel. Either select the voxel from the view with the maximum entropy (most information), or weight the voxel values in both views by their respective local entropy.
Imaging depth	A linear gradient is applied to each image as a weight based on the imaging depth. The further a light sheet moves inside an object, the less light reaches that part. Thus, weighing the information with a depth factor can really improve some results. ($Z=0$ best means that a non-rotated view has the clearest information on the bottom).

Table 12.1: Choices for what type of local information to use from the views when fusing.

result where all of the inputs have a contribution. The “Weighted sum” uses the voxel value, the variance, or the entropy as weight factor and then compute a weighted summation.

Extra features

On the bottom of the Customizable settings window, there are some extra features. Checking these options can improve the results of your fusion, but beware that they might use extra memory.

Subtract background is used to estimate the image’s background. Background subtraction is done prior to further processing. In case the images are deconvolved as well, this option will have little effect as deconvolution will perform this action as well.

Optimize at full resolution will add an additional registration step, which uses the result of the regular registration as its initial guess, and performs the registration at the full image resolution, while the default procedure will do it at reduced resolution. In theory, this might improve the resolution of the result, but it does so at much increased memory and computation time. In practice, the value of this step may be limited since the default procedure is accurate to a sub-voxel level at the reduced scale and other factors than the voxel size of the registration mostly dominate the accuracy of the result. Therefore, this option is disabled by default in all presets.

Multi-resolution optimization is a useful feature that is enabled in the last preset. It starts by smoothing the views to be registered. It reduces the amount of smoothing as the optimization process progresses and thereby finer and finer details of the original images are taken into account. This option can help to prevent the process from getting stuck as opposed to finding the optimal overlap. It also makes it less important that the initial guess be very precise.

Save registered views and **save folder** will save the images after registration, but before fusion, to a designated folder on disk. This allows inspection of the overlaid images individually in case of problems. It is recommended to open the saved images in the main window separately and then inspect them either in the Twin Slicer or in linked Slicer windows.

Time series

When tracking objects over time that were acquired with multiple views, it becomes necessary to perform fusion on time series. Open the custom settings and select the TIME SERIES tab to see the possible options for handling time series.

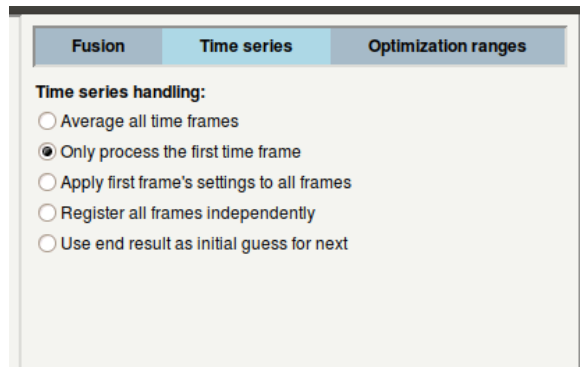


Figure 12.8: The options for handling optimization parameters in time series. The default option to only fuse the first frame is shown here.

By default only the first time frame is processed during fusion. The INTERACTIVE FUSER & DECON window makes it possible to apply the transformation of each time frame identically with the option “Apply first frame’s setting to all frames”. However, sometimes the objects move during the acquisition. To address this, there are two other options. One registers all frames independently by starting from the initial guess, as manually aligned. The other, final, option uses the transformation parameters after optimization of the previous frame as a starting point for the next frame.

Optimization ranges

To prevent the optimization from moving the views too far apart from the optimal solution, the Fuser has a set of ranges on how much the shifts and rotation angles are allowed to deviate from their respective initial guess values. Find these ranges under the Optimization ranges tab in the “Custom settings” window.

The re-scaling values are shown on top. The value for z is not shown in this window, as it was already dealt with in the Preprocessing steps.

The change in angles around which the images will rotate is reported in degrees. The shifts between views is reported in pixels. Alter these values if you notice that the optimization does not completely reach the desired overlap in the end.

By default the scaling (or stretching) in the x and y direction is turned off. This is done by setting these values to 0 as most of the time sampling in x and y is performed accurately. The z -direction is allowed to make small changes. These factors are reported in a *logarithmic* scale, or decibel scale (dB for short). This means that a value of 3 in this box allows the image to be expanded or shrunken down by a factors of 2 in size.

Starting the registration and fusion process

The START FUSION button is shown in green if at least two input images and settings are available that would enable the deconvolution, registration and fusion process to be started. Press this button to mark the end of the interactive phase of using the window where the user can change the parameters, and start the computation of the result (which can be time-consuming).

Computing the fusion of large empty regions would be wasteful of both time and memory

Fusion	Time series	Optimization ranges
Optimization rescaling values		
x scale	<input type="text" value="1.0"/>	
y scale	<input type="text" value="1.0"/>	
Optimization parameter ranges:		
x scale (dB)	<input type="text" value="0"/>	angle phi (deg) <input type="text" value="5"/>
		shift in x (px) <input type="text" value="25"/>
y scale (dB)	<input type="text" value="0"/>	angle theta (deg) <input type="text" value="5"/>
		shift in y (px) <input type="text" value="25"/>
z scale (dB)	<input type="text" value="1"/>	angle gamma (deg) <input type="text" value="5"/>
		shift in z (px) <input type="text" value="25"/>

Figure 12.9: The x - and y -scaling values, together with the ranges for the optimization process.

in this case. That is where the **Pre-crop region** checkbox can be used. If checked, a box is drawn inside the large preview as well as in the side view (“xy overlap view”). This box can be manipulated by dragging the lines or the corners, similar to the Cropper tool (see section 3 on page 16). This can drastically reduce the computation time and the size of the output image. As long as the “Pre-crop region” box is checked, the other functionality of the large preview window, such as manipulating the relative shifts, is disabled. Click **RESET CROP** to start over with a selection. Note that the side preview has two lines as well to restrict the result image size in the axial direction, but they may be less visible if dragged to the extreme edges of the image.

During the computation the **START FUSION** button changes to a **ABORT FUSION** button which can be used to discard the deconvolution and fusion run that is currently running, and return to changing the parameters.

If at any point you are happy with the current state of the automatic registration, it is possible to move on to the next view by clicking the **NEXT VIEW...** button that appears during this stage.

When the fusion is successful, then the fused image is added to the Huygens main window thumbnail panel and a pop-up window appears to notify the user that the fusion has completed. Each view will also have been saved at the location as specified in the custom settings.

Chapter 13

The Twin Slicer

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 Twin Slicer window, click the image name in the drop-down menu above the left or right view port (See Figure 13.1).




Figure 13.1: 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 13.1 on page 84:





Options	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).
SCALE BAR	show or hide the scale bar.
SVI LOGO	Show or hide the SVI logo in the lower right of the view port.
RESET VIEW TO DEFAULT	Sets all viewing settings to default.

Table 13.1: The options in the Twin Slicer’s VIEW menu.

Panning

Click and hold the *middle mouse button* on the slice to move it around. If the middle mouse button is missing, click and hold the *left mouse button* while holding the Ctrl button. 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 *Go to Basic Mode* and *Go to Advanced Mode*. In basic mode, all controls are visible in the panels below the view ports (See Figure 13.1). These controls apply to both slicers. This is in contrast to the advanced mode, which allows independent control of the left and right slicers (See “Using the Slicer in Advanced Mode” on page 86). At the right of the button to switch to advanced mode (so only under basic mode), there is a button *Add Slicer*. This button can be used to open a third slicer.

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.

Projection mode

You can select here *Slice* to activate the slider for slicing through the z-stack. It is also possible to show a *Maximum Intensity Projection (MIP)* or *Sum* of all the slices.

Orientation

Make a selection in the *Orientation* panel to change the plane that is displayed.

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.
- *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 Essential main menu: EDIT → PREFERENCES... → EDIT GLOBAL COLORS.
- *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.

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 or Strong compression*: 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 85).

- *Widefield or Strong 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 right *mouse button is clicked*, the Twin Slicer shows a context menu. The option “ZOOM IN BRIGHTEST SPOT” will automatically center and zoom in on the brightest spot in a 3D neighborhood around the cursor.

Using the Slicer in Advanced Mode

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. All controls are available in twofold and accessible through the tabs at 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.

Changing Detector

For Array Detector images a specific slider in the *Detector* tab allows for slicing through the detectors.

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* (🔄).

Setting the animation speed

Both the time slider and the slice sliders below each view port have a play mode to animate the slicing. With the slider between the slicers one can control the speed of this animation. For slow animation, move the slider towards the snail (🐌). This will add delays between the rendering of the scenes. When the slider is set all the way to the hummingbird icon (🐝) there is no delay between the renderings and the limitation of the visualization is simply how fast your computer can render the scenes. When decreasing the sizes of the view port, the rendering can go even faster, since small scenes are rendered more quickly than large scenes.

Rotation


The three radio buttons in the *Orientation* 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.

Projection mode

You can select here *Slice* to activate the slider for slicing through the z stack. It is also possible to show a Maximum Intensity Projection (MIP) or Sum of all the slices.

Changing Display Colors

Click the *Channels & Colors* tab key to view the color settings panel. The *Active channels* buttons (or drop down menu, if the image contains more than ten channels) can be used to enable or disable channels. The button colors correspond to the channel colors.

In addition to the *color schemes* that are available in basic mode (“Changing Display Colors” on page 85), 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 (master).

The second 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 85 for an overview). If the *Link* box is checked, the way of mapping data values to pixel intensities is the same for all channels; if not, the range is automatically adjusted according to the minimum and maximum intensity in each channel.

For more advanced gamma options, click on the *Contrast Editor* cartoon which opens a new interactive window. Here, you can manually change the image contrast for all channels at once or for each channel individually. It is possible to change the curves by adding new points, and dragging or deleting points. Preset contrast curves can be found in the drop down menu under the label *Compression*. After selecting a new compression, the corresponding curve and settings will be adjusted. By changing any of the preset curves, a custom curve is created and can be selected in the drop down menu as well. In the background, the histogram of each channel is shown.

Channels can be switched on or off by clicking on the check mark at the upper right corner. This enables you to change the contrast for the selected channel. Histograms will only be shown for the channels which are selected.

Contrast curves can be saved in two ways, as points or complete curves. Saving them as contrast curves will generate for each channel two columns, the x axis and y axis. On the x axis, the image intensity range is shown and on the y axis the contrast scaling factor. When the contrast curves are saved as points, the points can be loaded at a later moment into the contrast editor (not yet implemented).

When selecting the compression mode *Error function*, a manually soft threshold can be generated which suppresses the lower image intensities. The *Error function* is defined as the cumulative distribution function of the normal (or Gaussian) distribution¹. The pivot and the sigma correspond to the mean and sigma of the normal distribution, respectively. Increasing the pivot value will increase the threshold. The sigma corresponds to the range of the error function. Setting sigma to a small value (0.001) will result in a hard threshold, clipping the image intensities.

The third drop down menu under the *Contrast* panel allows you to either set the contrast stretch to globally (*Global*; the complete dataset will be taken into account), or for each screen view (*Per Screen*).

¹https://en.wikipedia.org/wiki/Normal_distribution

Linking, View & Plot

In the *Linking, View & Plot* panel are the most used options, which you can also find in the file menus. In one view you can easily see and set the links, view and plot properties. More options are available in the menus. For more information about the linking options, see the Linking Controls below.

Animate

The *Animate* tab allows you to set the frame rate in frames per second making it possible to create animations through time or through the *z* stack.

Shortcut Information

In the *Shortcut Info* panel all the available shortcuts are listed which are explained in the previous section. If your mouse or laptop lacks a middle mouse button, one can use the left mouse button while holding the **Ctrl** key to simulate the middle mouse button press.

Linking Controls

The **LINKING** file menu can be used to change the way in which both slicers communicate. The options in this menu are listed in Table 13.2. Note that settings get synchronized once the controls are being used.

Options	Description
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.

Table 13.2: The options in the Twin Slicer's **LINKING** menu (accessible in advanced mode).

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 the same piece of data is displayed for comparison.
- *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.
 - Go to LINKING and link the slice position, time frame, zoom level, and panning. Unlink the rotation.
 - Select the *Orientation* tab at the bottom of the window and select xz and xy .

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) beside 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.
 - Go to LINKING (either the menu, or the panel *Linking, View & Plot*) and link the slice position, time frame, and rotation. Unlink the zoom level and panning.
 - Drag the line separating the two scenes to the left to make the left slicer 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 13.2).




Figure 13.2: The Twin Slicer in *advanced mode*, with all controls but zoom and panning linked.

Measurement

Markers

Double click in one of the images or right click and select SET MARKER 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 beside 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 is shown in μm . Press and hold the Shift key while creating a new ruler and the ruler will be parallel to one of the axis, depending on your mouse motion.

Click and drag the end points of the ruler to make adjustments. Press and hold the Shift 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 Shift 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

If a ruler in the left slicer is selected, the right slicer will be replaced by a plot window and vice versa. The intensity profiles for both the left and right image are shown in the same plot. Unselect PLOT → PLOT BOTH SLICERS from the menu to only show the plot for one image. Graphs for the left slicer will have solid lines, while the graphs for the right are dashed (See Figure 13.3 and our online SVI Wiki²).




Figure 13.3: 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).

²<https://www.svi.nl/DataPlotter>

Chapter 14

The Orthogonal Slicer

The Orthogonal Slicer is designed to show the same point in 3D space from three orthogonal directions (see Figure 14.1).

It can be started by clicking on an image thumbnail and selecting ORTHO SLICER from the main VISUALIZATION menu, or using the 2D VISUALIZATION menu present under the right-button of your mouse.

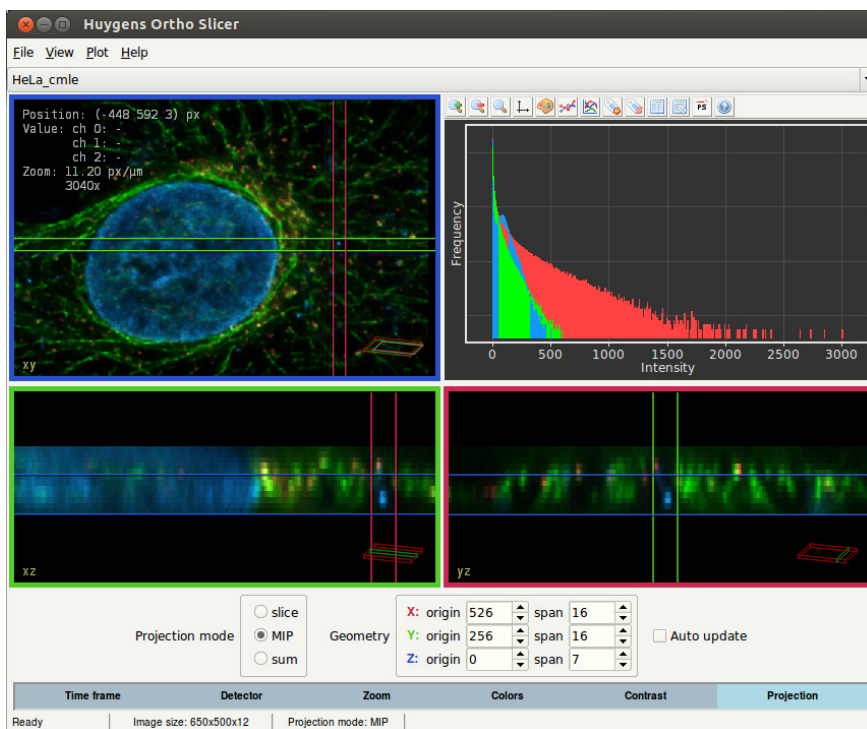


Figure 14.1: The Huygens *Orthogonal Slicer*.

Orthogonal views

The Orthogonal Slicer tool shows 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).

Each view is enclosed by a colored frame whose color indicates to which slices in the other views it corresponds.

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 behavior makes the *Ortho Slicer* a useful tool to study small objects in 3D.

The histogram (top right) shows the number of pixels in the complete image plotted against the intensity value. More information on the *Ortho Slicer* can be found on the wiki¹.

The Crosshair Cursor

The position of your mouse is projected as a cross-hairs pointer on all views. 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 the real pointer is more towards you (in front of your screen). If you click with your left mouse button on a position in one of the views, the slide-borders (indicated with colored lines) will be centered to that specific position.

Visualization parameters

Changing the visualization parameters in the Orthogonal Slicer is similar to the The Twin Slicer (on page 83). There are tools to:

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

Panning can be achieved by clicking and dragging the middle mouse button. To center the slice, use the menu than can be activated by placing the cursor on the view and clicking the right mouse buttons or by pressing 'c'.

¹<https://www.svi.nl/OrthoSlicerIntro>

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.

If a ruler is drawn, the histogram panel 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 capabilities².

Auto-Zoom

When you click the right mouse button, several options appear. The first can be used to 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, graphics 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.

²<https://www.svi.nl/DataPlotter>

Chapter 15

The MIP Renderer

The Maximum Intensity Projection (MIP) Renderer produces 3D to 2D projections of the highest intensities in an object. The result is similar to the MIP view in the *Twin Slicer* (see chapter 13) and *Orthogonal Slicer* (see chapter 14) tools, but in the *MIP Renderer* it can be obtained for any given viewpoint.

The *MIP Renderer* projects image voxels onto the screen by tracing rays from a given viewpoint through the data. For each ray the maximum encountered intensity is determined and displayed (See Figure 15.1). To start the *MIP Renderer*, right-click on a thumbnail and select 3D VISUALIZATION → MIP RENDERER from the pop-up menu, or choose VISUALIZATION → MIP RENDERER from the main menu. Finally, you can start the *MIP Renderer* by clicking its icon in the *Render Panel*.

Starting from version 19.04, the *MIP Renderer* supports GPU acceleration. This greatly reduces render times. Note that the rendered results are identical between CPU and GPU, which enables dynamic switching. Due to the nature of the ray tracing algorithm, rendered MIP images always show a certain amount of perspective. This means that objects farther away from the viewpoint appear smaller and that parallel lines converge.




Figure 15.1: A schematic overview of MIP rendering. The maximum intensities along rays originating in the viewpoint are projected.

Basic Usage

Orientation, Zoom, and Panning

Adjust the viewing angle by moving the *Tilt* and *Twist* sliders (See Figure 15.2), or by left-clicking and dragging the mouse pointer across the scene. The magnification can be adjusted by using the *Zoom slider* or the scroll wheel. Use the center mouse button to *pan* the center of the projection.

Note that the scene is rendered in a lower resolution while the scene is being changed. Once the scene is no longer being changed, a full resolution version will automatically be rendered. This automatic rendering can be deactivated by pressing the SWITCH OFF rendering button (🛑). When either the FAST MODE render button (🌿) or the HIGH QUALITY render button (🌳) are pressed the automatic full resolution rendering is reactivated.

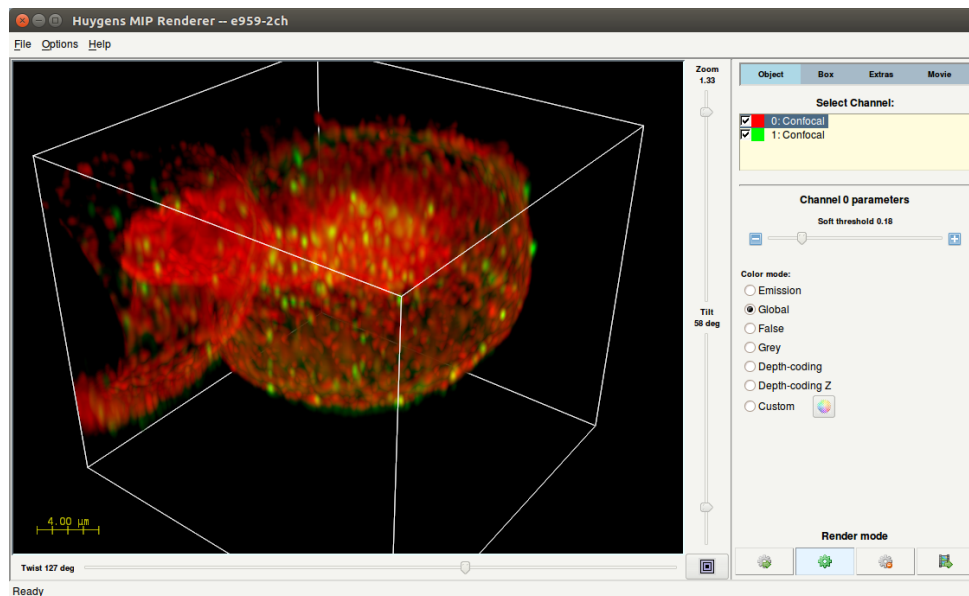


Figure 15.2: The MIP Renderer window.

Tabs for changing scene and movie settings

- **Object:** This tab contains settings affecting how the different channels are rendered. With the checkboxes in front of the listed channels, you can pick which channels are used. The *Soft threshold* slider in the *Channel parameters* panel sets the threshold level of the selected channel. Applying a threshold *reduces the background* in the image, so 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 discarded), a soft threshold uses a *smooth transition* around the threshold value. The color scheme for each channel can also be set. Note that *Global colors* are defined in the EDIT → PREFERENCES window.
- **Box:** This tab contains settings related to the bounding box, scale bar, and hue bar. Each of these elements can be toggled and their colors can be set. Furthermore for the bounding box the transparency can be set and for the scale bar the length can be set.
- **Extras:** Set the resolution of the scene. By selecting *Canvas* the scene will be rendered at a resolution that fits the window.
- **Movie:** See “Simple Animations” on page 97.

Options menu

Table 15.1 gives an overview of the render options available through the OPTIONS menu.

Option	Description
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, use the right mouse button to pick up and move the scene canvas.
COLOR MODE	Choose between EMISSION, GLOBAL (See “Adjusting the Global Color Scheme” on page 219), FALSE, GREY, DEPTH-CODING, DEPTH-CODING Z, or CUSTOM.
SHOW SVI LOGO	Hide or show the SVI logo in the bottom right corner.

Table 15.1: Render options for the *MIP Renderer*.

Depth-coding color modes

The *MIP Renderer* has two special color modes that do not appear in the other renderers: DEPTH-CODING and DEPTH-CODING Z. In these modes each point is given a hue based on the depth at which the maximum intensity was found. This depth can either be the distance to the viewpoint (in case of DEPTH-CODING, viewing-angle dependent), or simply the true z depth in the sample (in case of DEPTH-CODING Z, not dependent on viewing angle). The hues are automatically spread out over the complete range of depths visible in the current view. By adding a *Hue bar*, the relation between color and depth is shown.

Saving scenes

Choose FILE → SAVE SCENE... to save the rendered scene as a Tiff file. A scene template can also be saved so that it can be re-used or applied to other images.

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 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 113.) allows you to create sophisticated animations using the *MIP*, *SFP*, and *Surface Renderer*.

However, the *MIP Renderer* can also create simple animations on its own by transitioning between two custom keyframes. Set the render parameters for the first frame and click the MOVIE tab, SET VIEW → FIRST SCENE. Now adjust the parameters for the final frame, and click SET VIEW → LAST SCENE. The *frame count* and *frame rate* can also be set in the MOVIE tab. Finally press the *animate* button (⏮) to render the movie as an AVI or as a series of Tiff files.

The exported AVI files use the MJPEG¹ codec and can be played by most video players, including Windows Media Player and Apple Quicktime.

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

Chapter 16

The SFP Renderer

The *SFP Renderer* generates *realistic 3D scenes*, based on a 3D microscopy image, which is interpreted as a distribution of fluorescent material. The computational work is done by the *Simulated Fluorescence Process* (SFP) algorithm¹, simulating what happens when the material is excited and how the subsequently emitted light travels to the observer (See Figure 16.1).

The properties of this algorithm allow rendering of the object at different depths, unveiling layers beneath the object's surface. The SFP algorithm is not limited by boundaries or sharp gradients and is exclusively suited for rendering 3D microscopy data. Starting from version 19.04, the *SFP Renderer* supports GPU acceleration, which greatly reduces render times.

The *SFP Renderer* can be used in basic mode or in advanced mode. In basic mode a limited number of render parameters are available, and cutting planes and surfaces can not be added. The advanced mode is enabled by a special flag in the Huygens license string (The advanced visualization flag). See Render Parameters on page 103 for more info on which render parameters are available in basic mode.

To start the *SFP Renderer*, right-click on a thumbnail and select 3D VISUALIZATION → SFP VOLUME RENDERER from the pop-up menu, or choose VISUALIZATION → SFP VOLUME RENDERER from the main menu. It can also be launched directly from the *Render Panel* after selecting an image thumbnail in the main window.




Figure 16.1: In the *SFP Renderer* excitation of and subsequent emission of light by 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 excitation and emission light controls to what extent the viewer can peer inside the object.

¹<https://www.svi.nl/SFP>

Basic Usage

Orientation, Zoom and Panning

Adjust the viewing angle by moving the *Tilt* and *Twist* sliders (See Figure 16.2), or by left-clicking and dragging the mouse pointer across the scene. The magnification can be adjusted by using the *Zoom slider* or the scroll wheel. Use the center mouse button to *pan* the center of the projection.




Figure 16.2: The *SFP Renderer* window.

Render Quality

The *SFP renderer* can render the image with three different levels of quality. These can be set with buttons in the bottom-right of the SFP window:

- **Low Quality.** This is the lowest quality level and produces renderings with reduced resolution. This mode is only used while the scene parameters are being changed, if necessary.
- **Fast Quality.** Produces good quality renderings at full resolution.
- **High Quality.** Produces very high quality renderings while being more computationally expensive. In most cases the difference between fast and high quality are negligible, but it can be beneficial when trying to render scenes with very sharp gradients / surfaces.

The render quality can be set to fast (🔧), high (⚙️), or low (🔧) by using the corresponding buttons in the bottom right corner of the window. Note that regardless of the quality level, the scene will generally be rendered in low quality while the scene is being edited to allow for quick and smooth feedback.

Penetration depth

The top of the *Object* panel shows the *penetration depth*, which determines the general "hardness" of the objects. When light travels through the object, it is attenuated according to the properties and local density of the material. The length scale on which this attenuation takes place is called the *penetration depth*. By moving this slider the excitation, emission, and shadow transparencies are changed at the same time for all channels. These parameters can later be tweaked individually using the sliders below. On startup a penetration depth is estimated based on the image properties and contents.

Select Channel

The *Select Channel* panel provides a list of image channels with their colors and microscope types. The checkboxes in front of the channels are used to toggle which channels show up in the scene. By selecting a channel its render parameters can be edited independently.

Excitation transparency

The *Excitation transparency* determines the transparency of a certain channel material for excitation light. When excitation light is absorbed by an object, it casts a shadow on itself and all the objects beneath it (relative to the lighting direction) for all channels.

Emission transparency

The *Emission transparency* determines the transparency of a certain channel material for emission light. When emission light is absorbed by an object, it blocks the view of objects behind it (relative to the viewer) for all channels.

Shadow transparency

The *Shadow transparency* determines the "hardness" of the shadow on the table. For a low transparency the shadow will be very dark and sharp, and for a high transparency it will be bright and smooth. By default the *Shadow transparency* is set equal to the *Excitation transparency*, as this looks the most natural.

Object brightness

The *Object brightness* can be used to set brightness of a specific channel in the rendered scene. When set to 0%, the rendered channel will show up as black. When set to 100%, the rendered channel will show up as its corresponding color.

Soft threshold

The *Soft threshold* slider in the *Object* panel sets the threshold level of the selected channel. Applying a threshold *reduces the background* in the image, so 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 discarded), a soft threshold uses a *smooth transition* around the threshold value.

Color mode

Set the color for each channel by choosing between *Global* (See “Adjusting the Global Color Scheme” on page 219)), *Emission*, and or *Custom*. When *Custom* is selected, the color for the channel can be set using the hue and saturation sliders beneath the color mode selector.

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 own characteristic excitation, emission, and shadow absorption. The absorption properties can be controlled by the user (See the *transparencies* in Table 16.1 on page 104).

To excite the fluorescent material 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, or not 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* is placed below the data volume on which the cast shadows are visible.

After excitation, the fluorescent matter will emit light at a longer wavelength. Since this emitted light has a different wavelength it can not excite the same fluorescent matter: multiple scattering does not occur. Thus only 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 channels reaching the viewpoint.

Excitation light direction

An important consideration when making a rendering or animation using the SFP renderer is where the excitation light should be emanating from. The lighting direction azimuth and zenith angles are found under the *Environment* tab. The lighting direction will affect which parts of your object are clearly visible and which are shadowed. It also affects the shape and orientation of the shadow on the table. In general you want the light to be coming roughly from behind the camera. To aid with this, there is a *Camera tracking mode* that locks the azimuth angle of the lighting direction to the twist angle of the viewing direction, so that virtual light source appears to be locked to the camera. This means that when rotating the object, different parts of the object will be illuminated and the shadow on the table will change. By default, the *Camera tracking* is switched off, meaning that the shadows are baked into the object and the table and don't change when changing the viewing angle.

Adding an iso-surface

By using the settings in the *Surface* tab, an iso-surface can be added to the SFP scene. This surface is similar to the surfaces found in the *SFP Renderer*. However, in the SFP scene, the surface has more advanced interactions with the excitation and emission light, creating a much greater sense of space and depth. The surface is effectively interpreted as a completely impenetrable barrier for both excitation and emission light, and thus casts hard shadows on the SFP material beneath it and on the table. It is also partially obscured by SFP material which might be in front of it. Finally, the surface can also cast shadows on itself, so interesting renderings and animations can be made even without any fluorescent material.

Cutting planes

By using the cutting plane settings in the *Cuts & box* tab, up to six cutting planes can be added to the SFP scene with custom orientations and offsets. These planes completely remove all fluorescent material and surfaces on one side of them. This allows you to peer inside the object or to highlight objects that would otherwise not be visible. By using a combinations of cutting planes arbitrary regions can be selected. In image 16.3 the effect of adding multiple cutting planes to a scene is shown.




Figure 16.3: SFP rendering of a mouse blastocyst. In the image on the left, no cutting planes are used. In the middle image a single cutting plane is used with a plane normal roughly pointing towards the camera. In the right-hand image an additional cutting plane is used, with a plane normal pointing roughly upward. Image courtesy of Dr. Marc Duque Ramirez, Dr. Ritsuya Niwayama (Hiiragi group), and Dr. Stefan Terjung (ALMF) from the EMBL in Heidelberg, Germany.

Render Parameters

Table 16.1 gives an overview of all render parameters in the *SFP Renderer*. The options which are only available in advanced mode are followed by an asterisk.

Parameter	Description
Twist, tilt, zoom & pan	Adjust the viewing angle and camera position.
Render mode	Set whether the scene must be rendered in fast mode, in high quality mode, rendered in a movie or not rendered.
Time frame	Select the time frame for time series images.
Detector	Select the detector for multi-detector images.
Penetration depth	Adjust the overall transparency of the objects.
Select Channel	Select channels for display and for adjusting parameters.
Excitation transparency*	Adjust excitation transparency for the selected channel.
Emission transparency*	Adjust emission transparency for the selected channel.
Shadow transparency*	Adjust shadow transparency for the selected channel.
Object brightness*	Set the intensity level of the emission light for the selected channel.
Soft threshold	Adjust the threshold level for the selected channel.
Color mode	Adjust the color of the selected channel. Custom color mode not available in basic mode.
Select Channel (surface)*	Select the channel for the iso-surface. When set to <i>Off</i> , no surface is shown.
Threshold (surface)*	Set the threshold for the iso-surface.
Garbage volume (surface)*	Set the garbage volume for the iso-surface. Segmented regions smaller than this volume are not shown.
Brightness (surface)*	Set the overall brightness of the iso-surface.

Saturation (surface)*	Set the color saturation of the surface. High saturation leads to bright colors whereas a low saturation will lead to a gray (or white) surface.
Color range (surface)*	Range of colors used for separating the regions segmented by the iso-surface.
Camera tracking*	Let the excitation lighting direction follow the camera.
Light twist & tilt*	Adjust the excitation lighting direction (azimuth and zenith).
Table	Choose whether or not the table underneath the object should be included in the scene.
Table distance*	Adjust the distance between the object and the table.
Table reflection*	Adjust the reflectivity of the table. When 0, no reflections are visible.
Table size*	Adjust the size of the table.
Table color & brightness*	Adjust the color of the table through hue, saturation, and brightness.
Background color & brightness*	Adjust the color of the background through hue, saturation, and brightness.
Cut plane selection*	Activate or de-activate cutting planes. An easy way to set and activate a cutting plane is by using the <i>Standard cut planes</i> buttons, which produce orthogonal cutting planes going through the center of the image.
Cut plane twist*	Set the twist of the surface normal of the selected cutting plane.
Cut plane tilt*	Set the tilt of the surface normal of the selected cutting plane.
Cut plane offset*	Move the cutting plane along its plane normal to specify its position.
Show box	Show a bounding box around the image volume.
Box color	Set the color of the bounding box.
Show rulers	Show sets of rulers on the sides of the bounding box to determine real-life distances.
Ruler color	Set the color of the ruler lines.
Hide foreground rulers	When active, only show rulers behind the object so as not to obscure the object.
Ruler delta X, Y, and Z	Set the distance between subsequent ruler lines on the bounding box. Automatically snaps to round numbers.
First scene	Set the render parameters for the first scene of the movie.
Last scene	Set the render parameters for the last scene of the movie.
Set number of movie frames	Set the number of frames that will be included in the movie.
Set movie frame rate	Set the framerate (fps) of the movie.
Render size*	Set the size (in pixels) of the SFP scene. If set to "Canvas", fills the window. Note that the render size is locked in basic mode.

Table 16.1: SFP render parameters. Options followed by an asterisk (*) are only available in advanced mode.

Options Menu

Table 16.2 gives an overview of the different render options that are available through the OPTIONS menu. The options which are only available in advanced mode are followed by an asterisk.

Option	Description
SHOW SVI LOGO*	Show or hide the SVI logo at the bottom right of the scene.
COLOR MODE	Choose between GLOBAL (See “Adjusting the Global Color Scheme” on page 219), EMISSION or CUSTOM. Note that custom colors are not available in basic mode.
VIRTUAL RENDER SIZE*	Adjust the size of the rendered image. If the render size exceeds the display area, you can use the <i>right mouse button</i> to move the rendering canvas. Note that the render size is locked in basic mode.

Table 16.2: Menu options for the *SFP Renderer*. Options followed by an asterisk (*) are only available in advanced mode.

Saving Scenes


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

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 can be applied to any image that is loaded in the *SFP Renderer*. 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 113.) allows you to create sophisticated animations using the *MIP*, *SFP*, and *Surface Renderer*.

However, the *SFP Renderer* can also create simple animations on its own by transitioning between two custom keyframes. Set the render parameters for the first frame and click the MOVIE tab, SET VIEW → FIRST SCENE. Now adjust the parameters for the final frame, and click SET VIEW → LAST SCENE. The *frame count* and *frame rate* can also be set in the MOVIE tab. Finally press the *animate* button () to render the movie as an AVI or as a series of Tiff files.

The exported AVI files use the MJPEG² codec and can be played by most video players, including Windows Media Player and Apple Quicktime.

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

Chapter 17

The Surface Renderer

The Huygens *Surface Renderer* is a powerful 3D visualization tool that enables the visualization of iso-surfaces of volumes. The *Surface Renderer* is available as an option, and is enabled by a special flag in the Huygens license string (The advanced visualization flag).

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 17.1). The iso-surface envelops voxels with intensities above a certain threshold. If those voxels are distributed in groups that are not spatially connected, they will be considered and labeled as independent objects.

Shading enhances the perception of 3D shapes and texture (See Figure 17.2). Alongside the iso-surfaces, the *Surface Renderer* can also render a MIP projection (See Chapter 15 on page 95) and cross-sectional slices.

Starting from version 19.04, the *Surface Renderer* supports GPU acceleration. This greatly reduces render times. Note that the rendered results are identical between CPU and GPU, which enables dynamic switching. Due to the nature of the ray tracing algorithm, rendered surface images always show a certain amount of perspective. This means that objects farther away from the viewpoint appear smaller and that parallel lines converge.

To start the *Surface Renderer*, right-click on a thumbnail and select 3D VISUALIZATION → SURFACE RENDERER from the pop-up menu. Alternatively, choose the corresponding icon from the taskbar or select the VISUALIZATION → SURFACE RENDERER from the main menu. Finally, the *Surface Renderer* can be launched from the *Render Panel* after selecting an image thumbnail in the main window.




Figure 17.1: 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.




Figure 17.2: A schematic overview of surface rendering. The orientation of the surface determines the shading.

Basic Usage

Orientation, Zoom, and Panning

Adjust the viewing angle by moving the *Tilt* and *Twist* sliders (See Figure 17.3), or by left-clicking and dragging the mouse pointer across the scene. The magnification can be adjusted by using the *Zoom slider* or the scroll wheel. Use the center mouse button to *pan* the center of the projection.

Quality

Using the *quality* buttons in the bottom right the render quality can be set to *fast* (🌿) or *high* (🌳). In high quality mode, anti-aliasing is performed to smooth out sharp edges in the surfaces. For most images, the difference between *high* and *fast* quality is quite small.




Figure 17.3: The *Surface Renderer* window.

Threshold and Surfaces

Use the *Threshold* slider in the *Object Segmentation* panel to apply different thresholds to the data channels. 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 you to define three threshold levels that can be applied to the same or to different data channels. The data channels can be selected using the selector in the *Object Segmentation* panel. The color range in which the different objects in a render pipe will be displayed can be adjusted with the hue selector next to it. To increase segmentation of objects, watershedding can be applied. For more information see “Using the watershed segmentation” on page 162.

Advanced Usage

Adding a Maximum Intensity Projection

Alongside the surfaces a MIP projection of a specific channel can be added (See Chapter 15 on page 95) by selecting the *MIP* pipe. When the MIP pipe is used in combination with the surface pipes, a clearer representation can be obtained of the different objects in the image. In other words, the MIP of a channel can be used as a spatial reference for the objects present in other channels.

Adding a Slice

A *slicer* can be added to the surfaces and MIP by selecting the *Slicer* pipe. The slicer shows a single slice for a single channel of the 3D dataset in its corresponding location. Slicing can be performed along each of the orthogonal directions and the position of the slice can be set using the slider. Using a slicer actually cuts away the surface so you can look inside. The MIP projection is not cut by the slicer.

Render Parameters

Table 17.1 gives an overview of all render parameters in the *Surface Renderer*.

Parameter	Description
Frame	Set the time frame (in case of a time series).
Detector	Set the detector (in case of a multi-detector image).
Use watershed (surface)	Toggles watershed segmentation.
Sigma (surface)	Sets the watershed sigma (see “Using the watershed segmentation” on page 162)
Threshold (surface)	Adjust the threshold level for the selected surface, i.e. the threshold for the iso-surface. See “Threshold and Surfaces” on page 108
Seed (surface)	Only objects with an intensity higher than the seed level remain, while the rest are discarded
Garbage volume (surface)	Objects that contain fewer voxels than defined by the <i>garbage volume</i> parameter will not be rendered. This is useful for filtering out tiny objects in noisy images.
Transparency (surface)	Adjust the transparency of the selected surface.
Brightness (surface)	Adjust the brightness of the selected surface.
Channel (surface)	Set the surface channel for this pipe. Up to three channels can be rendered using the different surface pipes, or the same channel surface can be rendered with different thresholds.
Hue (surface)	Using the hue range slider the range of hues used for the surface for the selected pipe can be set.
Threshold (MIP)	Set the soft threshold level of the MIP.
Transparency (MIP)	Adjust the transparency of the MIP.
Brightness (MIP)	Adjust the brightness of the MIP.
Channel (MIP)	Set the MIP channel. Only a single channel MIP can be added.
Hue & saturation (MIP)	Using the hue & saturation slider the color of the MIP can be set.

Transparency (Slicer)	Adjust the transparency of the slice.
Brightness (Slicer)	Adjust the brightness of the slice.
Orientation (Slicer)	Set the orientation of the slice. For now only orthogonal orientations are allowed.
Position (Slicer)	Set the position of the slicer perpendicular to the slicing surface.
Show (Slicer)	Toggle the slicer. For now only the first channel can be sliced.
First scene	Set the render parameters for the first scene of the movie.
Last scene	Set the render parameters for the last scene of the movie.
Set number of movie frames	Set the number of frames that will be included in the movie.
Set movie frame rate	Set the number of frames per second of the movie.

Table 17.1: Surface render parameters

Options Menu

Table 17.2 gives an overview of the different render options that are available through the OPTIONS menu.

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 results in a movie with a duration of 7.5 seconds.
ANIMATION FRAME RATE	Adjust the frame rate of the movie.
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, the <i>right mouse button</i> can be used to move the render canvas.
TRANSPARENCY DEPTH	This option determines how many layers of a surface are rendered: <i>Simple:</i> Only render the outermost layer of each surface. <i>Normal:</i> Render two surface layers. The backside of objects can now be seen. <i>Deep:</i> Render many layers. Allows objects to be seen through other objects in the same pipe.
BOUNDING BOX	Enable or disable the bounding box.
SCALE BAR	Enable or disable the scale bar.
SHOW SVI LOGO	Hide or show the SVI logo in the bottom right corner.
HIGH QUALITY MIP	Render the MIP pipe in high quality mode.
CENTER SCENE	Undo both the panning of the projection center (middle mouse button) and the rendering canvas (right mouse button).

Table 17.2: Surface render parameters available in the OPTIONS menu.

Saving Scenes


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

Templates

All scene settings, i.e. both the render options and the 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 113.) allows you to create sophisticated animations using the *MIP*, *SFP*, and *Surface Renderer*.

However, the *Surface Renderer* can also create simple animations on its own by transitioning between two custom keyframes. Set the render parameters for the first frame and click the MOVIE tab, SET VIEW → FIRST SCENE. Now adjust the parameters for the final frame, and click SET VIEW → LAST SCENE. The *frame count* and *frame rate* can also be set in the MOVIE tab. Finally press the *animate* button () to render the movie as an AVI or as a series of Tiff files.

The exported AVI files use the MJPEG¹ codec and can be played by most video players, including Windows Media Player and Apple Quicktime.

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

Chapter 18

The Movie Maker

The *Movie Maker* is a tool that allows the user to easily create sophisticated animations of multi-channel 3D images using the powerful Huygens visualization tools. Animations from the *MIP Renderer* (See Chapter 15 on page 95), the *SFP Renderer* (See Chapter 16 on page 99), and the *Surface Renderer* (See Chapter 17 on page 107) can be combined in a single movie.

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

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 in the top right. This interactive tutorial guides the user step by step through the process of creating a simple movie.

An Overview

Figure 18.1 shows the *Movie Maker*'s 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 select the render parameter to be visualized. Green nodes, representing render parameters at each keyframe, can be dragged vertically.

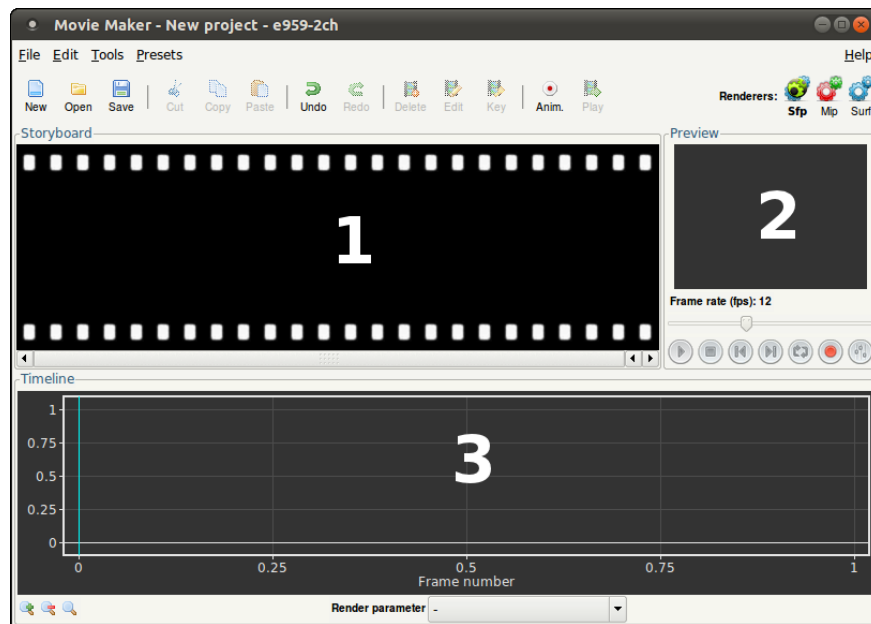


Figure 18.1: An overview of the *Movie Maker* interface, showing the *Storyboard* (1), the *Preview* (2), and the *Timeline*. (3).

Creating and Adjusting Keyframes

About Keyframes

A keyframe defines a *control point* within a transition. This can be either a starting 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 (Sfp, Mip, or Surf). Within the renderer, a scene can be defined using the available controls; see Chapter 15, Chapter 16, and Chapter 17 for more information. If the view of the first frame is satisfactory, press the *Add keyframe* button (K) to capture this configuration and to add the keyframe to the storyboard. All render parameters are captured and most of them can be smoothly animated.

A movie needs at least two keyframes from the same renderer to set the start and end point of a *transition*. If the second frame view is satisfactory, press the *Add keyframe* button (K) again. The storyboard will now show two keyframes with an arrow in between. This arrow indicates the transition from one keyframe to another.




Figure 18.2: The *Movie Maker*'s storyboard showing two surface-rendered keyframes, followed by two MIP-renderer keyframes.

The *Huygens Movie Maker* accepts a mix of keyframes from different renderers, but transitions can only be made between keyframes from the same renderer, as shown in Figure 18.2.

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. If 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 with 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 will be animated (See Figure 18.3). For most parameters, the Huygens *Movie Maker* uses linear or smooth transitions. An exception to this is the twist, which also requires a direction of rotation (clockwise or counter clockwise) and a value for the number of rotations.




Figure 18.3: 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 18.4). 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 applied to the movie preview, but only to the final result.




Figure 18.4: The Preview area.

Export to AVI or Tiff Series

Press the *record* button (Ⓜ) (See Figure 18.4) to render the final movie and to 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 18.5), where 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*" (📂).

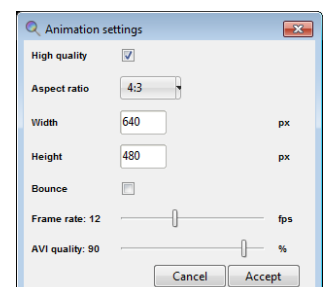


Figure 18.5: The animation settings dialog.

Working with Movie Projects

Saving Projects

A 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, press the *Save* button (📁) or choose FILE → SAVE PROJECT in the menu. *Movie Maker* project files have the extension .hgsa (Huygens animation template).

When closing the *Movie Maker* window, you are asked to save the project if unsaved changes have been made.

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 identical. It is also possible to apply saved projects to other images, or to the current storyboard. To load or append a project from disk, select FILE → OPEN PROJECT... or FILE → APPEND PROJECT... from the menu, respectively.

The *Movie Maker* has some example projects available under the menu PRESETS that can be used.

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 selected render parameter on the vertical axis (See Figure 18.6). It gives a more detailed visual feedback on values of each of the animated parameters. A mouse click in the *Timeline* area will select the nearest frame and displays a preview of this 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 scroll-wheel or the *magnifying glass* buttons below the timeline.




Figure 18.6: The *Timeline* area gives detailed visual feedback on the animated parameters

Changing Render Parameters

The keyframe nodes are displayed in green and can be dragged vertically to tune the value of the parameter. The (interpolated) transition frames are displayed as smaller red dots (See Figure 18.6). 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 to be shown in the graph, open the drop-down menu below the timeline. Because the number of parameters that can be animated is extensive, the menu only lists the ones that change during this movie. If “*Other render parameter...*” is selected, a dialog window will pop up that shows a list of all available parameters.

To change parameters in a frame which is not a keyframe, the frame first needs to 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*. 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 18.7). 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 complete movie in such a way that the relative length of each transition does not change.




Figure 18.7: The stretch dialog helps to increase or decrease the number of frames in a complex movie.

Synchronizing Transitions in Time and Slice Plane Transitions

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 18.8). In the final result, these asynchronous transitions may show irregularities.

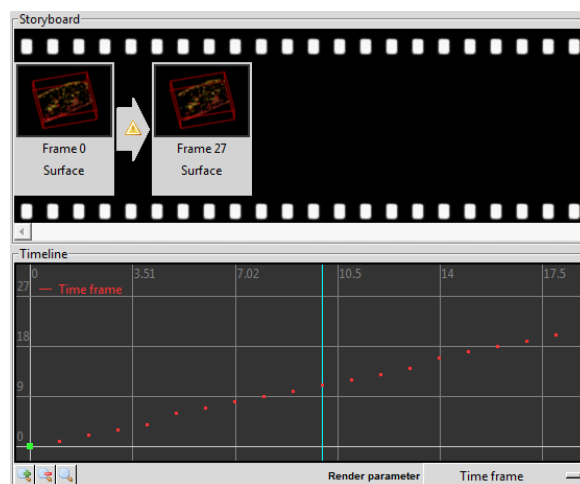


Figure 18.8: A transition in which the time frame parameter is out of sync with the number of frames. In this example, the transition counts 27 frames, while the time frame parameter increases linearly from 0 to 30. Note the irregular change of this parameter in the plot.

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 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 best, but then the file size will be large.

These quality settings are not applied to the movie preview, but only to the final result.

Chapter 19

The Gallery

The Gallery can be found under the menu **VISUALIZATION** and gives you a quick overview of your microscopic image. Instead of slicing through your image using a slider, the *Gallery* tool displays the individual slices of your z-stack or time series next to each other (See Figure 19.1). As a user, you can control which z-slices or time frames are displayed

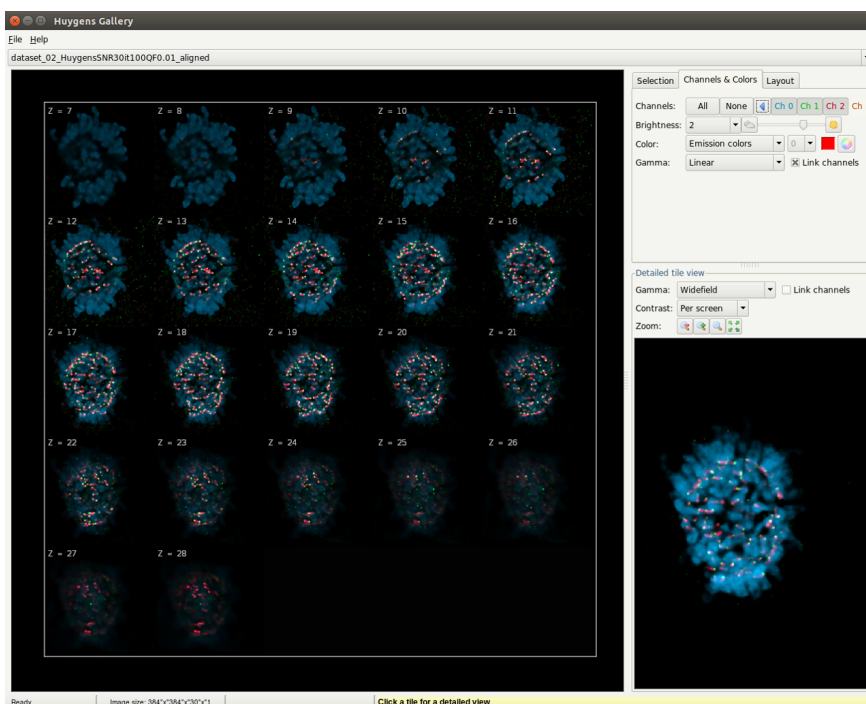


Figure 19.1: The Gallery tool showing the individual slices of a Z stack image.

in the gallery. Display settings for the overview can be adjusted using the options listed under the **CHANNELS & COLORS**. Tab The **LAYOUT** tab offers the possibility to change the layout and decoration of the overview. The overview can be exported as a tiff file via the **FILE** menu. Upon selecting a tile with the left mouse button, a detailed view appears in the window *Detailed tile view*. **GAMMA**, **CONTRAST** and **ZOOM** controls allow optimal and closer viewing of a selected tile in this window. For more details on the use of these controls see “Using the Slicer in Advanced Mode” on page 86.

Chapter 20

Tile Stitching & Deconvolution Wizard

Stitching & Deconvolution

When imaging large objects at high resolution, the microscope field of view can become too small. This problem is typically solved by subdividing the region to be imaged into multiple smaller images (tiles), which after acquisition are combined to a large overview. Common issues with stitched images are the appearance of shading artifacts due to vignetting and the increasing file size.

Improving images with deconvolution can be a challenging task as extensive computer resources for allocating both raw and processed images may be needed. The Huygens *Stitching & Deconvolution Wizard* offers a unique solution for this by combining high-quality stitching with deconvolution. This integrated approach offers the advantage of minimizing the resources needed for obtaining restored contrast-rich and high-resolution images of large objects. Furthermore, the post-acquisition work-flow is considerably simplified for the user.

Vignetting and Shading correction

Individual building blocks (tiles) of a stitched mosaic often suffer from vignetting and shading because of uneven field illumination. Such distortions can be seen in individual tiles, but are more apparent when tiles are stitched. To correct for these artifacts, the Huygens *Tile Stitching & Deconvolution Wizard* is equipped with an automatic vignetting/shading correction. As an alternative, there is a “manual” vignetting and shading correction option for which flat-field and dark-frame images need to be loaded.

To summarize, the Huygens *Tile Stitching & Deconvolution Wizard* is an easy-to-use wizard that allows for:

- optimal stitching of tiles.
- automatic or manual vignetting and shading correction.
- integrated deconvolution.
- Stitching of 2D, 3D multichannel and time series images.
- Stitching of unlimited files sizes and numbers of tiles.

Starting the Wizard

If Zeiss CZI, Leica LIF, or DeltaVision images with tiles, or a XML file with accompanying Tiff tiles are opened via the file menu, the tile content will be automatically recognized and opened in the *Stitching & Deconvolution Wizard*. The wizard can also be started in Huygens from the DECONVOLUTION menu. Next, click the LOAD MULTITILE IMAGE button in the *Welcome* screen, to select tiles that should be loaded into the wizard. This opens the window *Select sub-images* (See Figure 20.1). Use CTRL-CLICK to select or de-select individual tiles, SHIFT-CLICK to select a range, and CTRL+A to select all tiles.

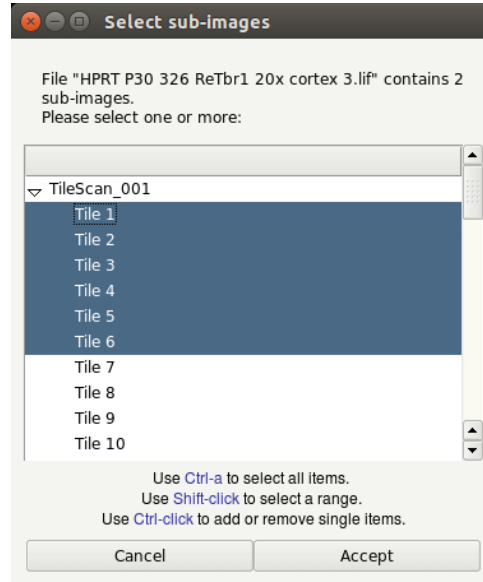


Figure 20.1: Selection of tiles contained in a dataset.

Single channel Tiff tile series of which the XML file with the position information is missing, but which individual tiles still contain the position information, can be stitched with Huygens by defining adjusting the custom stitching pattern in the *Tile Stitching & Deconvolution Wizard*.

Upon opening the first tile of such a Tiff series via the FILE → OPEN menu, or from within the wizard, Huygens will recognize the file as part of a file series and open the *File Series Tool*. See “File Series” on page 217. With this tool you can select the *Tile* dimension for the correct index in the file name, and click on LOAD SELECTION. This will automatically start the *Tile Stitching & Deconvolution Wizard*.

Pattern selection

The next window concerns the *Pattern selection* (see Figure 20.2). It shows the tiles' initial positions read from the metadata. If the ordering of the tiles is not specified within the metadata, as is the case with Tiff tile series lacking an XML, Huygens will fit the initial tile position. The *Pattern selection* window can be used to customize the stitching pattern. Parameters such as acquisition pattern, starting point, percentage of overlap between tiles, and number of rows can be defined at this stage. Also buttons are present to flip tiles in x, in y, and in both directions. The buttons SHOW TILES and SHOW MIPS can be used for viewing the actual content of the tiles. A slider, present under the *Help and Report* windows allows you to slice through the Z dimension if the tiles are 3D stacks. Another slider under the *Time frame* tab, which is located at the bottom, will allow you to visualize other time frames. Several acquisition pattern parameters can be adjusted here and the tile overview will be immediately updated. The Huygens *Stitching & Deconvolution Wizard* allows a minimum value of "0" percent as acquisition overlap.

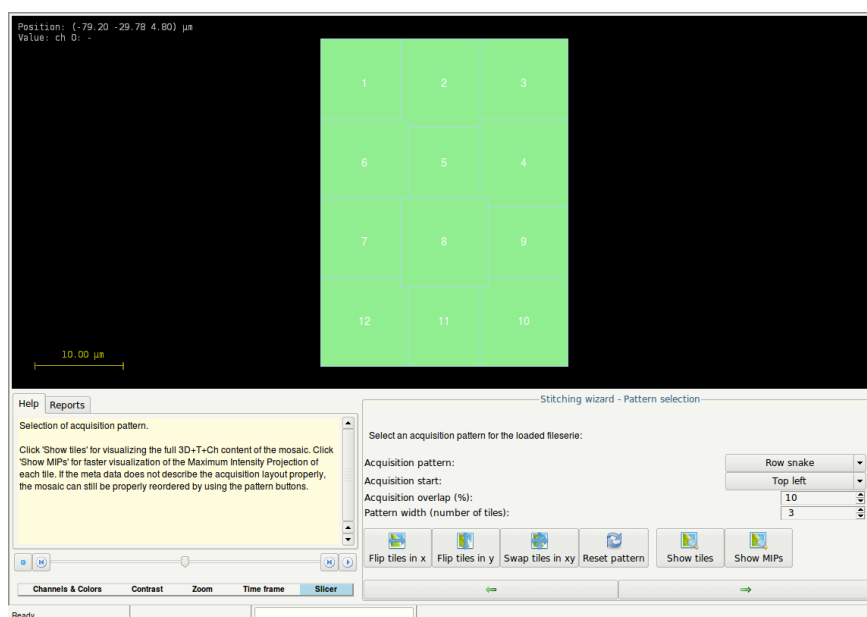


Figure 20.2: The Stitching and Deconvolution Wizard window showing the tiles' initial positions.

Tile Selection

Also within the next *Tile selection* window, you can show the actual content of the tiles. By default, all tiles of the mosaic are selected. Use the *mouse left click* plus the CTRL/SHIFT keys to select a specific region of tiles, or to select and de-select tiles.

Vignetting and shading correction

On the next stage, you can select whether automatic vignetting and shading correction must be applied, or if this should be manually done by loading a flat-field and a dark frame image. When automatic is selected, you need to select which channels will be corrected, and what model for the vignetting needs to be applied. This can be either a default, circular, or ellipsoidal model. The severity of vignetting can be adjusted with the slider. The steepness of the vignetting model at the tile edges can also be adjusted with the presented slider. The optimal settings for these parameters may need to be determined experimentally. An example of a stitched image with and without automatic vignetting/shading correction is shown in the *Twin Slicer* window in Figure 20.3.

Optimization of transitions

Before the actual stitching is executed, The *Tile Stitching & Deconvolution Wizard* will calculate the optimal transitions for each pair of tiles. At this stage, you can select what channels need to be included for this optimization step. If the overlap between the tiles is not optimal, you may want to go back to the step *Pattern selection*, and adjust the acquisition overlap percentage.

Deconvolution of tiles

The *Tile Stitching & Deconvolution Wizard* offers the user the unique possibility to deconvolve each tile automatically before stitching. Every single tile is first deconvolved before being stitched into the mosaic. By doing so, Huygens not only simplifies the work-flow but also minimizes the processing workload that would be needed for deconvolving a stitched mosaic image, or for stitching tiles that have increased in size because they have already been deconvolved. To deconvolve the tiles before stitching, you just need to select a microscopic template and deconvolution template (“Verifying Microscopy Parameters” on page 13, and See “The Final Stage” on page 22.). These templates should be created or edited first before starting the *Stitching & Deconvolution Wizard* by using the options EDIT → EDIT DECONVOLUTION TEMPLATE (for details on the parameters See “The Deconvolution Stage” on page 20.) or EDIT → EDIT MICROSCOPIC PARAMETERS (See “Verifying Microscopy Parameters” on page 13)




Figure 20.3: Huygens stitched tiles of widefield fluorescent Leica LIF data. The same data is shown in the Twin Slicer without (left panel) and with (right panel) Huygens automatic vignetting correction. Note that the overlapping regions are less apparent after vignetting correction. Image represents a developing mouse cortex (P30) stained for Tbr1, reelin and an unspecified molecule.

Finally, start the complete stitching process with the optional vignetting correction and deconvolution, using the **STITCH TILES !** button. It is recommended to take a closer look at the *Report* window for any specific messages. For example, if the image data is not containing sufficient information for vignetting correction, the Wizard may decide to omit this step.

Chapter 21

The Hot & Cold Pixel Remover

Why correct for hot/cold pixels?

Hot pixels appear as bright pixels within your image. They do not originate from the specimen or noise, but are caused by individual pixels of the CCD camera with higher than normal rates of charge leakage. In fact, all pixels on a CCD detector have some charge leakage. This is why the amount of hot pixels will increase when you increase the exposure time. The charge leakage gets worse at higher temperatures.

Cold pixels are the opposite: pixels with a much lower than expected intensity. They are caused by individual sensors in the CCD or CMOS camera with lower than normal sensitivity (even as low as zero, so called 'dead' pixels).

Hot and cold pixels can negatively affect the contrast range in an image. As a consequence, deconvolution, image stabilization, visualization, and data analysis can be negatively affected. Hot pixels can also be a product of deconvolution. If a tiny object in your original image fits the PSF perfectly, deconvolution may result in a very bright spot in your image. This can disturb the contrast range and therefore the visualization of the image. We have come across this situation a few times. So, if the image seems black after deconvolution, try to remove hot pixels first and see if this brings back the original contrast of the image.

Hot Pixel Remover window

The *Hot & Cold Pixel Remover* can be started from the DECONVOLUTION menu in Huygens Essential. DECONVOLUTION → “HOT & COLD PIXEL REMOVER”

The *Hot & Cold Pixel Remover* interface initially shows an xz maximum intensity projection of the original image (top left) and a corrected image (top right). See Figure 21.1.

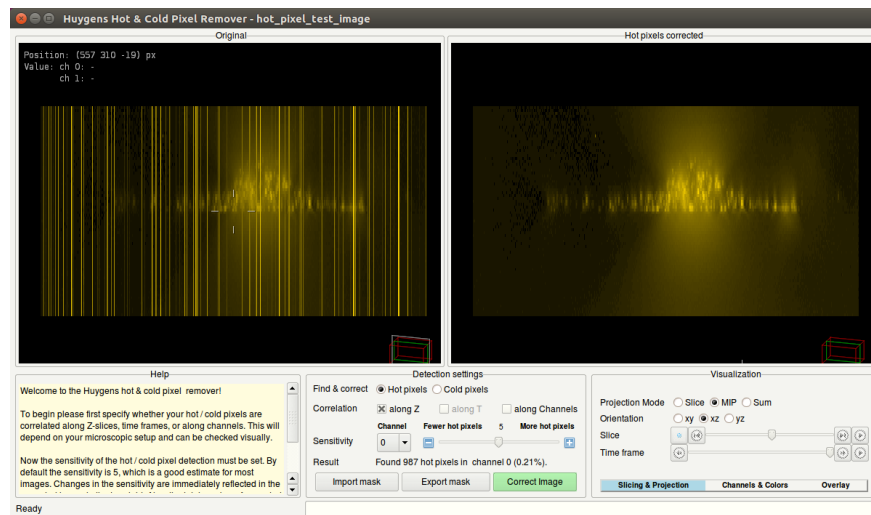


Figure 21.1: The start-up window of the Hot Pixel Remover.

By default a MIP projection is shown in the xz direction as this makes it easier to recognize hot pixels by eye. For cold pixel removal, a sum projection is used. At first, only a single channel is shown because the correction needs to be performed for each channel separately. At the very bottom a status bar is shown with a progress bar. In the bottom left a help box explains the general features. The bottom center shows the *Detection settings* box with the detection parameters and different correction and mask options. The bottom right shows a standard visualization settings box:

- *Slicing*: contains three slicers which allow you to slice through the xyz stack (only enabled when the projection mode is slice), time frames, and detectors.
- *Projection*: here you can switch between a *Slice*, *MIP* (Maximum Intensity Projection), or *Sum* projection. The orientation of the projection can also be set here.
- *Channels & Colors*: five different color schemes are available in the first drop down menu. If you are interested in viewing other channels, they can be selected or deselected. By default only one channel is shown. The additional gamma drop down menu can be used to optimize the contrast of the image which can be useful if the hot pixels have a large range of intensities.
- *Overlay*: within the display menu you can choose between viewing only the image, a hot/cold pixel overlay, or both. The transparency and color of the hot/cold pixel overlay can be adjusted with the slider and color selector, respectively. These tools are helpful in verifying whether the identified hot pixels are correct.

Detection Settings

The most important setting is the top one: *Find & correct*. Here, you can specify whether you want to detect and remove hot pixels, or cold pixels. Only one type of pixel aberration can be corrected at a time. In most cases, an image will only suffer from hot pixels, so this is the standard mode. In case of an image with both hot and cold pixels (rare), it is recommended to first remove the hot pixels, and then to remove the cold pixels from the resulting image.

The locations that are flagged as hot/cold pixels will depend on the pixel detection settings. Whenever these settings are changed the shown correction changes accordingly. Also next to *Result* you will see the total number of unique hot/cold pixels that have been corrected. The settings are split up into two groups: The *Correlation* settings and the *Sensitivity*. The correlation settings allow you to specify whether there are some additional correlations between the hot/cold pixel locations, which the detection algorithm can use. If you are not sure if any correlations exist, you can simply disable all correlations. Next we will explain why and when these correlations can occur.

Correlation along Z

When a camera defect is the cause of hot/cold pixels, this means that certain pixels in the 2D sensor grid will often produce a hot/cold pixel. When the image is recorded slice-by-slice (in the z direction), this means that each z slice has a high likelihood of having a hot/cold pixel at this exact location. This leads to “hot columns” and “cold columns” respectively, which show up as vertical lines in an xz projection (like a barcode). Since this is almost always the case, the *along Z* correlation is enabled by default.

Correlation along T

For the same reason as above, when a time series is recorded with a single camera (as is usually the case in a time series), the hot/cold pixels will show up at the same location in each time frame. In this case the correlation *along T* option should be enabled. Note that if the hot/cold pixels in a sensor change rapidly over time, this will no longer be true.

Correlation along channels

For the same reason as above, when the channels of a multi-channel image are recorded one by one, with the same camera, the hot/cold pixels will show up in the same locations in each channel. In this case the correlation *along Channels* should be enabled. Note that if the hot/cold pixels in a sensor change rapidly over time, this will no longer be true. When *along Channels* is checked, the image windows will show all channels at once.

After specifying the repeating parameters the correction will generally already be good. However, the sensitivity of the tool can be fine-tuned for each channel using the sensitivity slider. When a channel is selected next to the sensitivity slider it is automatically shown in the image windows. When the sensitivity is lowered, fewer hot/cold pixels will be found, and vice versa. A higher sensitivity increases the odds of flagging regular pixels as hot/cold pixels, while a lower sensitivity increases the chance of flagging hot/cold pixels as regular pixels. Ideally you want to lower the sensitivity as much as possible before hot/cold pixels show up in the correction.

Saving the correction

When the corrections have been verified and (optional) the sensitivities have been fine-tuned for each channel, the correction can be saved by clicking the **CORRECT IMAGE** button. This will load the correction directly into the main Huygens window, from which it can be saved or edited further.

Using a Mask

The locations of the hot/cold pixels can be saved by clicking the EXPORT MASK button. This will save the pixel mask as a .h5 file. This mask file can then be imported back into the *Hot & Cold Pixel Remover* for correcting another image, or can be added to a deconvolution template to correct a whole series of images with identical hot or cold pixels. A previously determined mask can be imported from file by clicking the IMPORT MASK button. This mask will then override the currently defined correction settings. It is recommended to verify the correction settings that will be applied with this mask. Note that when the detection settings are changed, these will take priority again and the imported mask will no longer be used.

Adding hot or cold pixel removal to a deconvolution template

A hot or cold pixel removal step can be added to a deconvolution template using either the *Deconvolution Template Editor* from the EDIT menu or using the internal deconvolution template editor in the *Batch Processor*. In either case you must navigate to the pre-processing options (*Pre* tab). Adding a hot or cold pixel removal step simply means importing a previously created mask from file. When the specified mask is valid, it will light up green, and a pixel correction step is added to the deconvolution template. If the path is invalid, it will light up red and the pixel correction step is skipped. Whenever this occurs you will be notified of the reason for this invalidity.

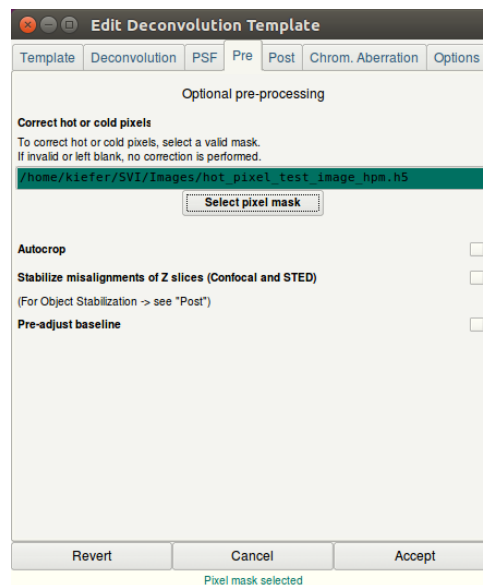


Figure 21.2: Tool for editing the deconvolution parameters and including hot pixel correction.

Chapter 22

The Crosstalk Corrector

The Crosstalk Corrector detects and corrects for (linear) crosstalk between all channels in multi-channel images.

What is crosstalk?

In fluorescence microscopes, crosstalk (also frequently referred to as bleedthrough) can occur when acquiring a multi-channel image. In that case, the emission radiation of a given emission wavelength is detected by the wrong detector because part of the photons go through the wrong optical path inside the microscope (e.g. because the filters efficiency is not 100%). Therefore, some signal is actually recorded as coming from a certain dye when it really comes from a different one. The amount of crosstalk into the other channel depends on the signal intensity and it shows a linear dependency.

To avoid crosstalk, microscopes usually excite each dye alternatively, making sure that all the detected radiation comes from a single dye type. But some experiments (like e.g. Fluorescence Resonance Energy Transfer - FRET) require simultaneous acquisition of signal from all the present dyes, with the possible risk of crosstalk. Also the separation of the excitation and emission spectra of different fluorescences may overlap and result in crosstalk.

Crosstalk within multi-channel images dramatically affects almost any type of data analysis, including colocalization analysis.

In line with our aim to improve microscopic image quality and measurements, we have implemented a Crosstalk Corrector tool in Huygens Essential to correct for this imaging artifact.

Starting the Crosstalk Corrector

To start the Crosstalk Corrector in Huygens Essential, select an image and go to DECONVOLUTION → CROSSTALK CORRECTOR. This tool can handle multi-channel time serie images up to 32 channels.

The tool does not automatically estimate the crosstalk at startup. This is to avoid slow loading time, especially if the image contains many channels. The more channels there are in the image, the longer the estimation will take.




Figure 22.1: The startup window of the Crosstalk Corrector with a large 2D histogram of the original data (left) and a preview of the 2D histogram and MIP of the corrected data (right).

The Crosstalk corrector shows a large 2D histogram in the left panel of the startup window, see Figure 22.1. This histogram is from the original raw data set and is useful for the detection and visualization of the crosstalk coefficients.




Figure 22.2: A 2D histogram, showing the intensity range of the first channel on the x-axis and the intensity of the second channel on the y-axis. The triangle shows that for every intensity in the first channel there is always a fraction of it in the second channel which is typical for crosstalk.

Assume, for example, that in a 2-channel image the first channel bleeds into the second channel. Then at least a fraction of every signal (or intensity) in the first channel exists in the second channel. And so, for every signal in the first channel no intensity combination exists with low signal (smaller than a fraction of the first channel) in the second channel. This is visualized by an empty (except for some noise) triangular shape, from the origin along the axis of the first channel, see Figure 22.2.

Note that the colors within the 2D histogram have nothing to do with the colors used to represent the channels. The colors show the density or count of intensity combinations within the two channels.¹

The slope of the diagonal of the triangle is the crosstalk coefficient of, in this case, channel 0 in channel 1. For each crosstalk or bleedthrough in each channel there is a crosstalk coefficient. For a 2-channel image there are two crosstalk coefficients, for a 3-channel images there are six, for a 4-channel image there are twelve.

In general, for an n -channel image there are $n^2 - n$ crosstalk coefficients. These coefficients can be displayed in a matrix (see box entitled ‘View’) with element (i,j) the crosstalk coefficient of channel i in channel j . The elements (i,i) can be ignored, since comparing channels with itself will result in a crosstalk of 1.0. Note that the crosstalk coefficient matrix is not symmetric.

¹Read more about 2D histograms at <https://www.svi.nl/2Dhistogram>

The crosstalk corrector can handle multi-channel images up to 32 channels, but the histogram can only show two channels at a time. That is why you can select two channels to visualize the 2D histogram, but the correction will take place on all the channels simultaneously.

Next to the large 2D histogram, within the preview frame, the 2D histogram and MIP view of the corrected data are shown. This way you can quickly validate the crosstalk correction.

With the options at the bottom of the window you can adjust all of the available display settings: set the two channels for the histogram views, toggle between histogram view and matrix view, toggle between a MIP preview of the corrected and original data, select for the MIP preview all channels (shown in gray) or only the selected channels (shown in their corresponding channel selector color) and to center the MIP preview. Use the mouse left button to activate a magnification tool for inspecting all the displays.

With clicking the REDUCE NOISE button, you will apply a noise correction to the image which will facilitate the correct estimation of the crosstalk coefficient from the 2D histogram. This can be undone by clicking on the same button which is now marked ORIGINAL IMAGE. With the ESTIMATE button you can estimate all the crosstalk coefficients. With the CORRECT IMAGE button you apply the final crosstalk coefficient matrix to the originally loaded image and the result image with suffix etc, will be shown in the main window of Huygens Essentiall.

Adjust the crosstalk coefficient

There are three ways to set the crosstalk between channels. The simplest is to press the ESTIMATE button and let the tool determine the crosstalk coefficients. If the estimation is not satisfactory the crosstalk coefficients can be adjusted manually or you can try to see whether the button REDUCE NOISE will produce a 2D histogram from which the coefficients can more easily be estimated.

The crosstalk coefficients can be manually adjusted by moving the handles within the 2D histogram, which are the small circles at the end of the crosstalk lines. When the handles are moved, the crosstalk line changes, and thus the crosstalk coefficient.

Click and drag on the middle of a crosstalk line, to move this line completely. This does not change the crosstalk coefficient, because the slope does not change. It is intended for verification only.

The crosstalk coefficient can also be adjusted directly by switching to matrix view and changing the element in the matrix which corresponds to the correct crosstalk coefficient. For example, if the crosstalk coefficient of channel 2 in channel 3 needs to be adjusted, edit element (2,3) (row, column) in the crosstalk coefficient matrix.

Chapter 23

Bleaching Corrector

The *Bleaching Corrector* is a tool in Huygens that can estimate and correct for bleaching effects and lamp jitter within a 3D and/or time series image.

What is Bleaching?

In fluorescence microscopy, the result is a fading of the emission intensity with respect to the Z and/or T dimension of the image. Bleaching is a change in composition or state of the dye molecule, making it unable to fluoresce.¹

Large differences in integrated intensity between the first and last plane/frame of the image is not only an imaging artifact in the raw data, but may also, in combination with the use of discrete Fourier transforms during deconvolution, cause incorrect intensity values at the edges, so-called wrap-around effects. In addition, intensity instability throughout the image also results in unwanted artifacts in the restoration result.

Starting the Bleaching Corrector

To start the Bleaching Corrector in Huygens Professional or Essential, select an image from the main window and select from the main taskbar DECONVOLUTION → BLEACHING CORRECTOR.

The main window of the *Bleaching Corrector* tool shows on the left side a single slice of the image up for correction (see Figure 23.1). More information regarding the scene and additional visualization options can be found in the *View* menu in the taskbar at the top of the window. At the bottom of the window, tabs are present whose actions are coupled to the rendered image on the left. These four tabs labeled *Time*, *Channels & Colors*, *Slicer* and *Contrast* offer the possibility to change the rendered image properties.

The *Time* tab shows a slider that can be used to change the frame that is shown. The tab *Channels & Colors* lets you choose which channels you want to show and with what type of color scheme. The *Slicer* tab lets you change the plane that is shown, and the *Contrast* tab can be used to change the brightness and gamma of the image.

On the right side of the displayed image an empty data plot is displayed under each of the three different tabs. The plot remains empty until an estimation is performed.

¹Vicente et al., Journal of Physics: Conference Series 90 (2007); Song et al., Biophysical Journal (1995) 68: 2588-2600.

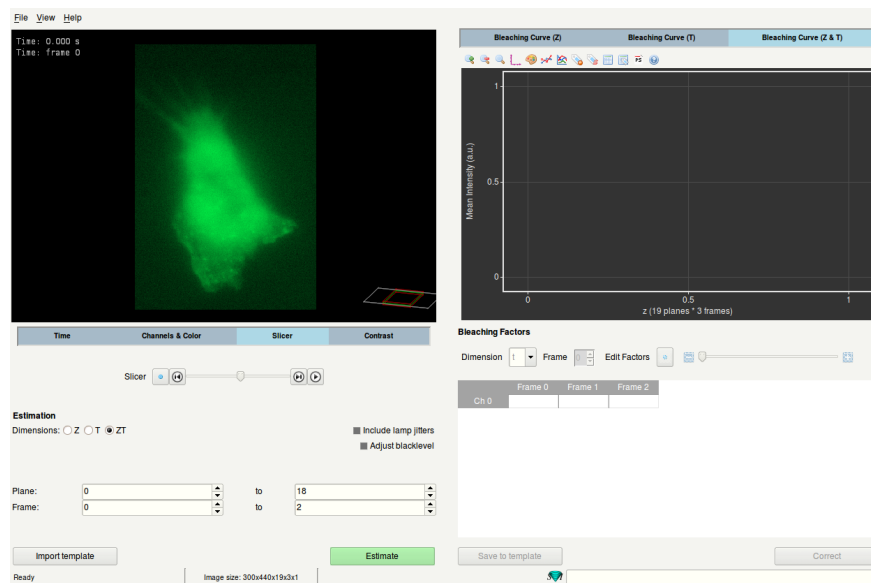


Figure 23.1: Startup window of the Bleaching Corrector

Estimate Bleaching

Under the header *Estimation* on the left side of the window, it is possible to set the initial parameters for the estimation of the bleaching. Depending on the image dimensions, you can choose from which dimension (Z, T or ZT) you want to estimate the bleaching by clicking on the corresponding buttons. You will have the option to check the box **INCLUDE LAMP JITTERS** to include a correction for illumination instability. This option is only possible for Widefield image data. In addition, there is an option to adjust for the presence of a black level. Next, choose the range of planes and/or frames in which you want to check for bleaching effects by scrolling through the spin boxes. Press the **ESTIMATE** button to estimate the intensity correction. The data plot on the right will be filled with the original (solid line) and corrected (dashed line) intensities along time and Z. See Figure 23.2.

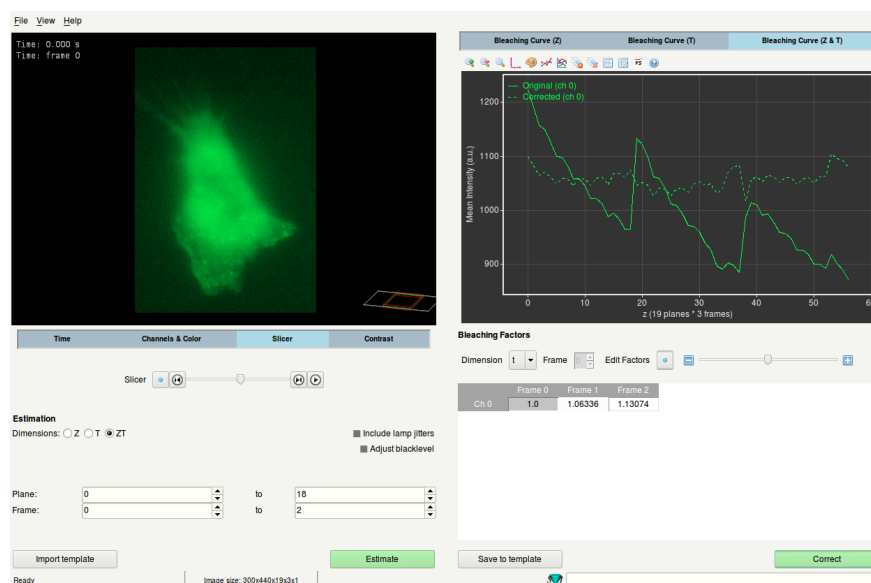


Figure 23.2: The window of the Bleaching Corrector after estimation of the bleaching.

Adjusting Bleaching Factors

After estimation, the table shown on the right-hand side of the window will be filled with bleaching correction factors. Each represents the value of the corrected intensity divided by the original intensity. The bleaching factors in the table correspond to the intensity curve shown in the data plot. It is possible to switch to different dimensions of bleaching by toggling between the different data plot tabs, or by changing the dimension in the combo box located above the table itself. Also, it is possible here to select a different frame from the corresponding field, or by clicking on the respective field in the table. For large numbers of channels, frames or planes, it is possible to use the scroll bar located below the table. Click on a cell to select a bleaching factor or double click on a channel row title to select a complete row. By using the slider or keyboard the user can change the bleaching factors and adjust the corrected intensity accordingly. The blue bullet button sets the selected cells to their initial estimated bleaching factor.

Note that over-correcting, i.e. correcting for effects that are not bleaching related can result in unwanted loss of intensities.

Correct Bleaching

After the bleaching factors have been estimated (and maybe edited), the image can be corrected with the **CORRECT** button. This will create and export a corrected image to the Huygens main window from which you can continue working with the corrected image. For example, you can compare the original and corrected image with the Twin Slicer or use the corrected image to start the Deconvolution Wizard.

Working with templates

The estimated bleaching factors can be saved to a template using the **SAVE TO TEMPLATE** button in the *Bleaching Factors* frame. Saved templates can be imported and applied to other images. A list with saved templates is shown in the *Load Template* frame. Use the **IMPORT TEMPLATE** button if the preferred template is located in a different folder than the one shown in the template list. The template values will be loaded and shown in the bleaching factors table. Furthermore, the plot will show the corresponding correction as if the **ESTIMATE** button was pressed. Once the template is loaded, the user can continue with optionally adjusting the bleaching factors and correcting the image as is described in the previous paragraphs. It is recommended to apply templates to images that were acquired with the same microscope settings.

Chapter 24

The Chromatic Aberration Corrector

Since Huygens version 15.05, the *Chromatic Shift Corrector* option has been renamed *Chromatic Aberration Corrector* as, besides removing misalignments between channels, it is now also capable of automatically estimating and correcting differences in scaling and rotation between channels.

The support for templates in this option allows chromatic correction to be applied from one image to other images. This is particularly interesting when the estimation is carried out on a multi-channel bead image. The use of templates for chromatic aberration correction is also available in the batch processor.

Causes of chromatic aberration

The chromatic shift, usually being the largest component of the chromatic aberration, is a misalignment across the channels of a multichannel image which can be corrected by simply shifting the channels.

There are several circumstances that can lead to chromatic shifts in the microscopic images:

- Chromatic aberrations in the microscope optics.
- Misaligned beam splitters.
- Misaligned excitation lasers.
- Misaligned color filters.
- Faulty color interpolation in color cameras.
- Any other internal misalignments in the microscope.

Therefore, chromatic aberration can be a rather frequent - though correctable - imperfection in multichannel images.

Starting the Chromatic Aberration Corrector

- Launch Huygens Essential.
- Load a multichannel image to be corrected for chromatic aberration.
- Select the image thumbnail and in the top menu go to DECONVOLUTION → CHROMATIC ABERRATION CORRECTOR.

The *Chromatic Aberration Corrector* will open and show the image on an orthogonal slicer where the existing chromatic shifts can be seen in a 3D view (For operating the slicer See “The Orthogonal Slicer” on page 91.). Below the orthogonal slicer is a z-slicer, a time slicer, and other visualization tools such as contrast, color 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 Aberration Corrector* at start-up with a loaded two-channel bead image is shown in Figure 24.1.

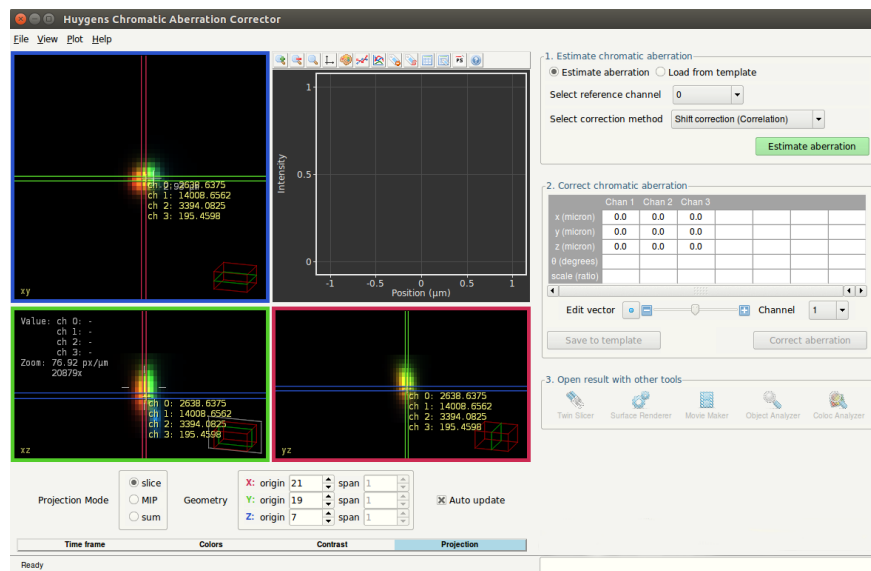


Figure 24.1: The Chromatic Aberration Corrector at start-up, prior to estimating the chromatic aberration.

Estimation of the chromatic Aberration

The following three methods can be chosen for the automatic estimation of chromatic aberration:

- Shift Correction (Cross correlation). This can be considered an ‘all-round’ method. The software searches for the best alignment across channels by maximizing the overlap.
- Shift Correction (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.
- Full Correction (Cross Correlation). This approach corrects for shifts, scaling, and rotation between channels. The software searches for the best alignment across channels by maximizing the overlap. Both scaling and rotation correction are performed with respect to the image center

Chromatic shifts will be quantified by a three dimensional vector, indicating how much a channel is shifted with respect to the given reference channel.

The channel to act as reference can be selected by the user via the Reference Channel selection box. Because this channel will have no chromatic aberration with itself, it will not be reported.

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

Visualization of the chromatic aberration

The estimated chromatic aberration will be reported within the table. The aberration of each channel but the reference will be listed. The length unit of the three dimensional shift vectors is set to micrometers. Rotation will be presented in degrees, and a difference in scaling will be represented as a ratio value.

The user can select any channel for editing except for the one used as a reference. This can be done by using the EDIT CHANNEL selection box under the vector table.

The shift vector selected for editing will be drawn on the orthogonal slicer, so that the estimated shift can be easily assessed. Notice that the vector components are projected onto each orthogonal view correspondingly. The total length of the vector is displayed next to each vector projection. This length is a measure of the estimated chromatic shift.

At the same time, a plot shows the intensity profiles along the direction of the shift vector estimated for the reference channel and the channel selected for editing. 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, and allows us to see to what extent the intensity of the 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 aberration estimation is shown in Figure 24.2. 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 should show no remaining shift, having similar shapes and peaks roughly located at the same x positions.

The Chromatic Aberration Corrector will return accurate and reliable estimations of the existing chromatic aberration. Still, a possibility to edit and customize the estimated shifts is offered so that the user can reach more precision if necessary.

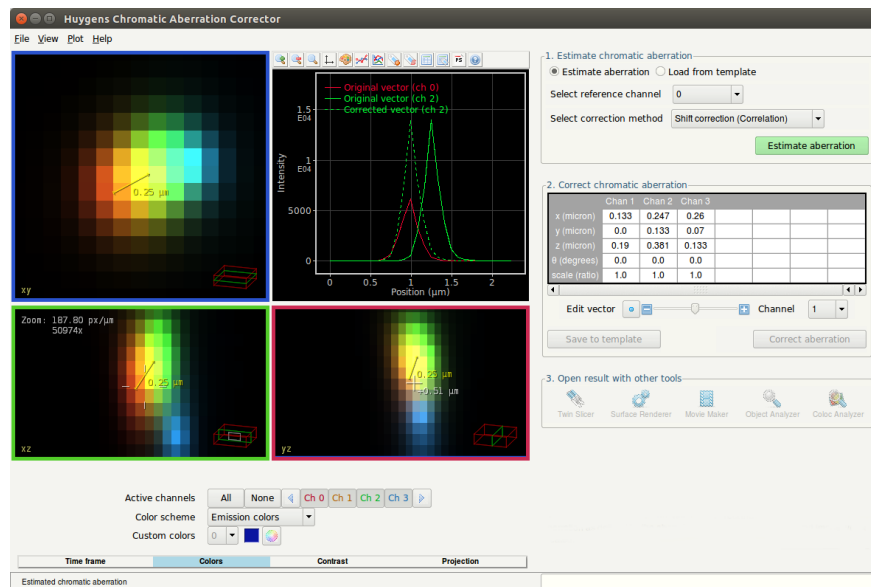


Figure 24.2: The Chromatic Aberration Corrector after estimating the shifts. A plot shows the gap between channel 2 and reference channel 0. The estimated shift is drawn as a vector in the orthogonal slicer scenes.

Editing the chromatic shift vectors

The user can also edit the shifts that have been estimated automatically by the Chromatic Aberration Corrector. The components of the selected shift vector can be modified by using the EDIT VECTOR tool, which allows to shorten and lengthen the estimated shift.

While the shift vector 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 24.3. Channels can be edited one by one and their shifts modified with again the EDIT VECTOR tool.

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

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

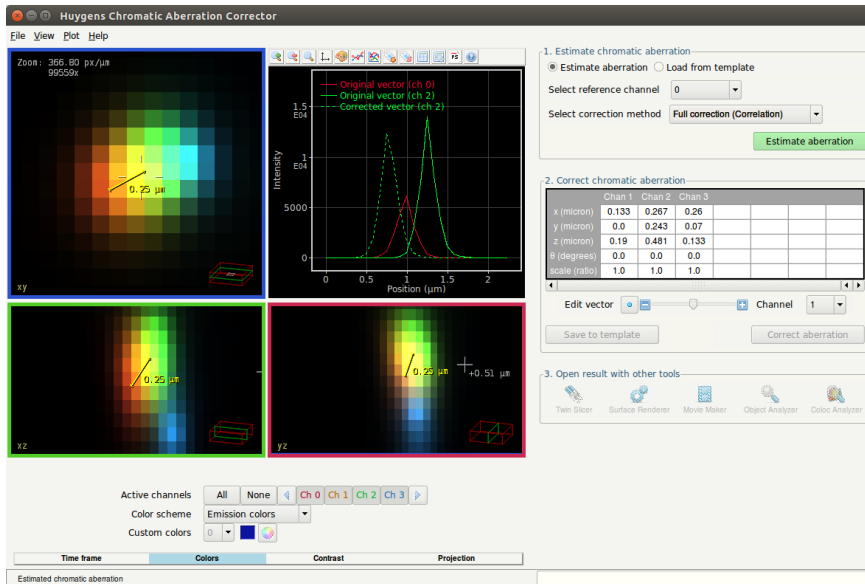


Figure 24.3: Using the Edit Channel tool of the Chromatic Aberration Corrector. The estimated vector has been manually lengthened, and as a result the dashed line plot is slightly moved to the left.

Working with templates

The estimated aberration correction can be saved to a template by using the template tool of the Chromatic Aberration Corrector.

Saved templates can be imported and applied to other images. The template values will be loaded and listed in the table. Furthermore, the plot will show the corresponding shift in the intensity profiles as if the ESTIMATE ABERRATION button had been pressed.

Subsequently, the image can be corrected by clicking on the CORRECT ABERRATION button, which will create a new corrected image in the main window. This will also activate buttons with which the corrected image can be immediately opened in Huygens tools for visualization and analysis. Some of these tools are available as options.

It is recommended to apply templates to images that have the same emission and excitation wavelengths as the image with which the template was created.

Chapter 25

The Object Stabilizer

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). Stabilization can be performed over time (aligning subsequent 2D or 3D time frames), or over z (aligning the slices in a z stack). Both the displacement measurement and subsequent stabilization are done in 3D and at sub-pixel level.

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

The Object Stabilizer Window

Upon opening the *Object Stabilizer* you are presented with three orthogonal projections of the image. The menu spanning the bottom of the window houses the visualization settings just as in the *Ortho Slicer*. In the upper right quadrant a wizard is embedded which will guide you through the stabilization process. First you will need to specify whether to stabilize your image over time, or along z slices, and then you need to select your stabilization method. Since the stabilization methods and subsequent stages depend strongly on this choice we will describe these modes separately.

Preprocessing

The preprocessing stage is common to both time and z stabilization and allows you to crop your image, either automatically or manually using the *Cropper* tool. Furthermore the channels that are taken into account during stabilization can be set. Unselecting a channel here means that it does not contribute to the calculation of the inter-frame or inter-slice displacements. Note however that the result image will always contain all of the original channels.

Next the interpolation quality can be set. This determines how the slices/time frames are resampled during stabilization. The *Lanczos* method delivers the highest quality result and can take sub-pixel shifts into account. This mode should always be used unless there are special reasons not to. The *standard* mode can also take sub-pixel shifts into account but is a bit faster. Finally the *low* method does not take sub-pixels shifts into account and will only perform integer shifts. The most significant advantage of this method is that it allows the data to keep its original data type, whereas the *standard* and *Lanczos* methods will automatically produce a floating point result image. Therefore this mode should be used if you want to stabilize a large byte or integer image and you are low on free memory.

Stabilization of 3D Time Series

Stabilization over time works best if the 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 following 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 tool under TOOLS → CONVERT can be used to convert the z dimension to the time dimension.

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 stabilizer 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* stage appears in which rotation detection can be enabled or disabled (rotation detection takes more time), and where it is possible to select a region of interest in the image, as shown in Figure 25.1. The NEXT button will start the alignment process.

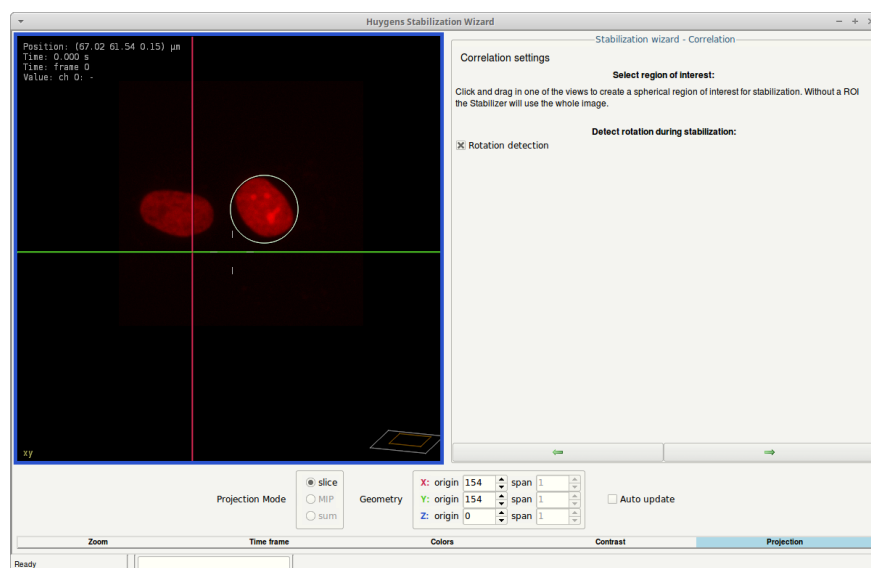


Figure 25.1: The *Correlation* stage in the *Object Stabilizer*.

When the *Object Stabilizer* has finished measuring the displacements, the *Stabilization* stage is shown. Continue reading “Stabilization Settings” on page 149.

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 *Object Stabilizer* creates this model automatically.

Pushing the NEXT button on the pre-processing stage 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* stage is shown. Continue reading “Stabilization Settings” on page 149.

The Multi Object Tracking Method

When your image contains well-defined objects, i.e. nuclei or small particles, 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 25.2). In this stage 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.

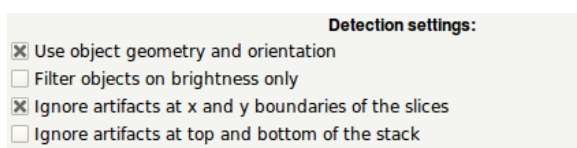


Figure 25.2: The *Detection settings* stage.

When the detection algorithm finds many false objects at the image borders (this may happen for very small object sizes), 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 25.3 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.

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 25.4.

The histogram and sliders in the *optimizing object detection* stage (See Figure 25.4) can be used to apply thresholds on the *number of objects*, *score*, *width*, and *brightness*. Objects that fall outside of one of these thresholds will be ignored. The score is a statistic that reflects the certainty of the detection algorithm, i.e. objects with a low score are probably noise. When the thresholds are modified, pressing NEXT will recompute the detection instead of proceeding to the next stage.

In the next stage the tracking parameters can be modified. A limit can be set to restrict the maximum distance over which an object is allowed to move in between two time frames. This means that the tracker will not connect two objects in subsequent time frames if their mutual distance is larger than this value.

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).



Figure 25.3: An example of object and background selections.

The NEXT button will start the tracking process. When this is finished the *Edit tracks* stage is shown (See Figure 25.5 on page 146). 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* stage is shown. Continue reading “Stabilization Settings” on page 149.

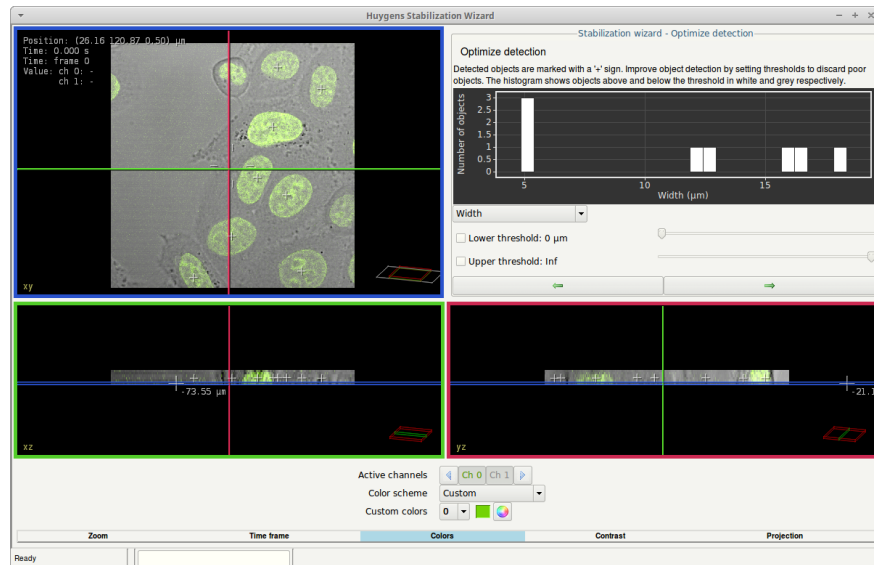


Figure 25.4: The *Optimizing object detection* stage in the *Object Stabilizer*.

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 stage 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* stage is shown. Continue reading “Stabilization Settings” on page 149.

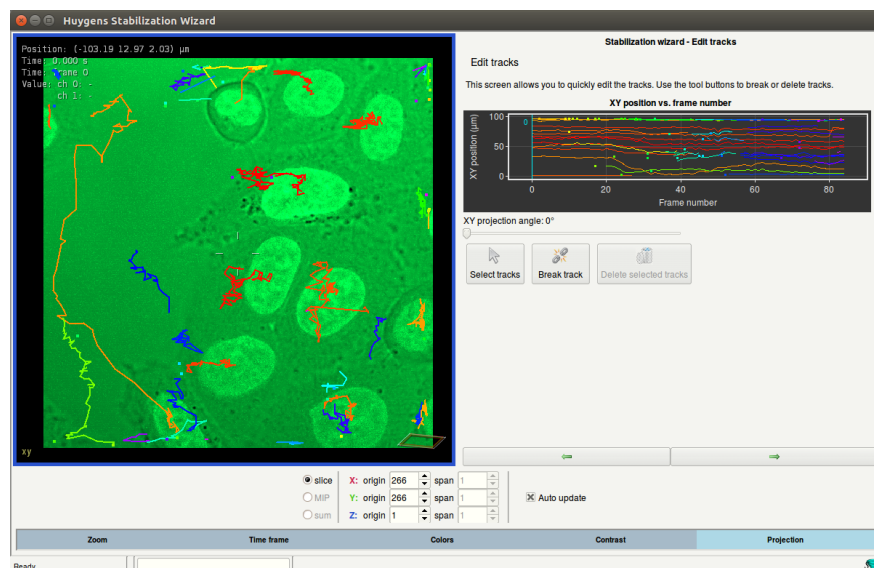


Figure 25.5: The *Edit tracks* stage in the *Object Stabilizer*.

Alignment of Slices in 3D Stacks

Misalignment of the slices in your image can severely affect the quality of deconvolution. Therefore, in contrast to stabilization over time, stabilization along z should always be performed before deconvolution. There are three different stabilization methods for stabilization along z , which will be explained in the following paragraphs:

- *Cross correlation*
- *STED thermal drift correction*
- *Light sheet shearing correction*

In the case of an image which contains both z slices and time frames, displacement measurement and stabilization will be performed for each time frame individually. In case both z and time stabilization are required, z stabilization should be performed first. Then the image can optionally be deconvolved before finally stabilizing over time.

A comparison between a misaligned and aligned z stack is shown in Figure 25.6.




Figure 25.6: An x - z slice of a misaligned z stack (left), and the alignment result (right). The chromatic shift in a stack like this one can be corrected using the *Chromatic Aberration Corrector*.

Note that 3D alignment deforms the volume. It is advisable to check the alignment result carefully by using the x - z and y - z views of the *Twin Slicer* (See “Using the Slicer in Basic Mode” on page 84).

Cross Correlation

This is the most general method and can be used on all images. Cross correlation along z works just like cross correlation over time, as explained in “The Cross Correlation Method” on page 144. In the *Correlation settings* stage, you can choose whether or not to consider rotations and whether or not to use *Iterative filtering*. This last method attempts to smooth out the inter-slice displacements and is recommended for very noisy images and images with empty slices as the top and/or bottom of the stack. After clicking NEXT you will be taken to the *Stabilization* stage.

STED Thermal Drift Correction

This method is optimized for correcting thermal drift in STED imaging systems. No rotation is taken into account and no further input is required. After going through the preprocessing stage you will be taken directly to the *Stabilization* stage.

Light Sheet Shearing Correction

In certain light sheet systems the objective is oriented at an angle relative to the sample stage. This means that the stage moves in a direction that is not parallel to the optical axis, which causes a shear in the resulting image. By selecting *Light sheet shearing correction* you will be taken to an additional stage that allows you to specify the parameters of your imaging setup. This can be seen in Figure 25.7.

The *objective angle* is the angle between the optical axis and the movement direction of the sample stage. For a regular microscope this angle would be zero degrees and there would be no shear. The stage movement direction indicates the direction in which the stage moves as measured in the coordinate system of the image. If you are unsure about this setting you can just check the direction in which your image appears to move when moving up through the z slices. Finally the z *sampling interpretation* refers to the true meaning of the z sampling in the image. It can either be absolute movement of the stage in between z slices or the movement as projected onto the optical axis. In case of the former an additional correction of the z sampling is performed after stabilization. After clicking NEXT you will be taken to the *Stabilization* stage.

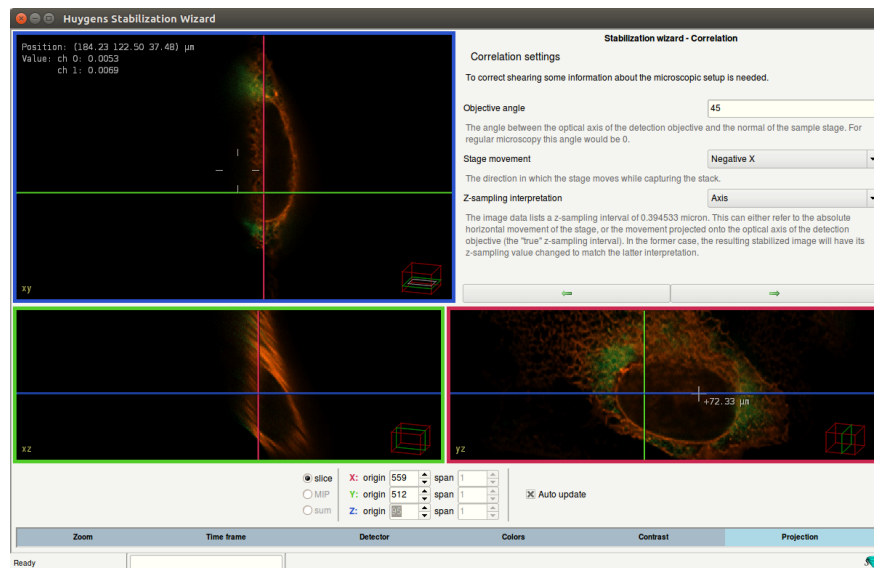


Figure 25.7: To perform a light sheet shearing correction some additional parameters are necessary¹.

¹Image is of a breast carcinoma cell acquired using a dithered, multi Bessel beam Lattice Light Sheet microscope from 3i. Image courtesy of Tobias C. Walther and Shane D. Elliott, Department of Genetics and Complex Diseases, Harvard University.

Stabilization Settings

The *Stabilization* stage shows the found displacements and allows you to fine-tune these displacements, depending on the stabilization direction and method. The displacements are shown in two ways: in the slicers and as a set of graphs in the wizard stage. The displacement graphs allow you to see the total displacement in the x , y , and z directions separately as a function of the slice/frame. Furthermore if applicable the total rotation is also graphed. In the slicers bounding boxes show the final positions of the frames/slices if the image were to be stabilized according to the current displacements.

Depending on the stabilization method, a reference slice/frame can be selected. This slice/frame is guaranteed to keep its original position and orientation while the other slices/frames are moved around it. Furthermore, in the case of z stabilization with an image with time frames, the currently shown time frame can also be selected.

The three sliders can be used to remove outliers, noise, and drift. Removing outliers removes the most extreme frame to frame or slice to slice displacements. Filtering noise means suppressing short term changes in the displacement, only keeping the long-term structure of the displacements, while filtering drift does the opposite. The effect of the noise and drift filters can be seen in Figure 25.8.

Clicking NEXT will apply the stabilization and send you to the *Cropping* stage.

Cropping the Result

Because extra space is required for the translated and rotated frames/slices, the stabilized result will be larger than the original image. The *Cropping* stage (See Figure 25.9) lets you determine how the result will be cropped. Depending on the stabilization method there can be up to four different methods:

- *Tight*: Use the largest possible rectangular area that does not contain any empty regions.
- *Full*: Use the full size of the stabilized image. This makes sure no information from the original image is removed, but may lead to empty regions in the result.
- *Original*: Keep the result the same size as the original. The reference frame/slice selected in the *Stabilization* stage will be in exactly the same position and orientation as in the original image.
- *Manual*: A cropping region can be set manually by dragging the red bounding boxes in the slicers. See Figure 25.9.

When you are done cropping you can click NEXT and you will be taken to the final *Stabilization summary* stage.




Figure 25.8: Effect of filters in the *Stabilization* stage. The top shows the measured displacement, the curve in the middle is noise filtered, and the bottom is drift filtered.

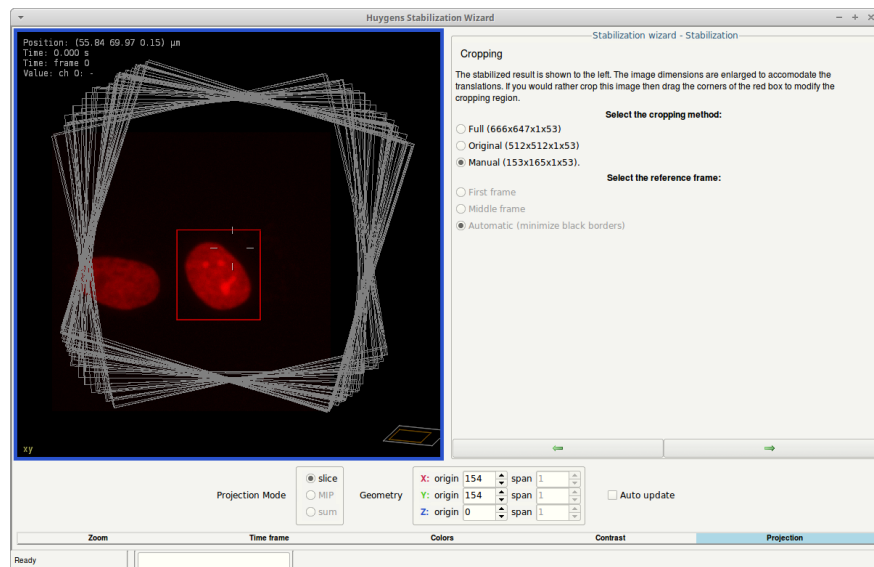


Figure 25.9: The stabilization result can be cropped manually to cut out the objects of interest. The original image borders are shown as grey rectangles.

Stabilization Summary

Here the stabilized result is shown in the slicers. If you are satisfied with the result you can click DONE and a stabilized version of the original image will appear in the main Huygens window. If you are not satisfied with the result you can either press RESTART to completely restart the wizard or QUIT to simply exit the wizard.

Chapter 26

The Object Tracker

Introduction

The Huygens Object Tracker wizard and Track Analyzer can be used to study 3D motion of cells and smaller particles in time series images.

By manually selecting just a few objects and background regions, the Object Tracker is trained to differentiate between object and background volumes, and thus to detect new objects. In this way not only can bright spot tracking be done, but also tracking of more complex scenes. Because the Object Tracker is designed to detect and track objects with a well-defined center, it works best with spherical, disk-like (flat) or slightly elongated objects.

After the Object Tracker wizard has finished processing the time-lapse data, the Track Analyzer is launched. The Track Analyzer can also be launched directly if tracks were stored from a previous session of the Object Tracker. The Track Analyzer helps to filter, edit, and analyze the results. It presents the properties of the measured tracks in informative graphs and allows to export track data to a file.

The Object Tracker Wizard

The Preprocessing Stage

If the Object Tracker is launched –through the main menu ANALYSIS → OBJECT TRACKER or the thumbnail pop-up menu– it will start the *preprocessing* stage (Figure 26.1). The wizard checks if the sampling distances and time interval are reliable (reported from file or verified) and will pop-up a parameter editor (Figure 3.1 on page 14) if those parameters need attention.

Depending on the image dimensions (2D/3D, single/multi channel) the preprocessing stage will give the option to create an axial maximum intensity projection (MIP), select relevant channels, and/or crop the image (See “The Intelligent Cropper” on page 16). These preprocessing operations help to speed up and enhance the tracking process.

Detection Settings

In this stage the parameters that control the object detection can be adjusted. The first screen of this stage shows the choice to use a wizard or to load a template of detection settings saved in a previous session.

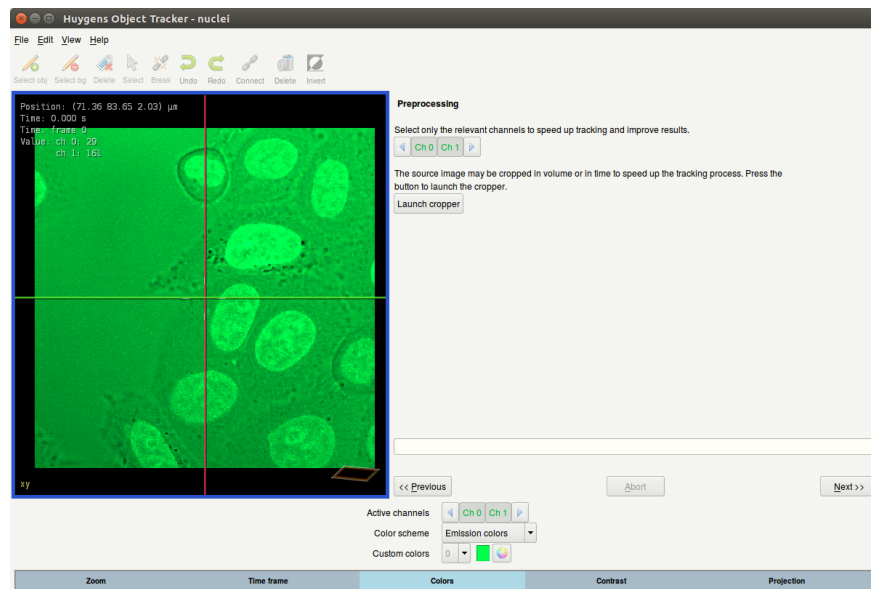


Figure 26.1: The preprocessing stage in the Huygens Object Tracker.

The next screen gives the option to *filter on brightness only*. This disables other image filters, and is faster and usually good enough for bright (or dark) spot tracking.

The *ignore artifacts* choices give the option to discard objects that are detected close to the image borders, as image borders can be sensitive to false positives.

Object Detection Stage

In the *object detection* stage the user is requested to select a few objects and background areas. In this stage one can also create a region of interest (ROI) such that only objects within this ROI are analyzed. By pressing the DEFINE A ROI, a new window will be opened in which an ROI can be created. When done, press the CLOSE AND EXPORT ROI button. The ROI will be visualized in the tracker as a new overlay image. The search for objects is then limited to object within the ROI.




Figure 26.2: Two object selections (green) and a background selection (red).

To start the training, use the *object selection tool* (👉) to draw green outlined spheres to select objects. Use the *background selection tool* (👉) to draw red outlined spheres to select regions containing mostly background voxels (Figure 26.2).

Like the Orthogonal Slicer (page 91), the Object Tracker displays the same point in space (the image center) from three orthogonal directions (xy , xz , and yz). The object and background selections are visible and can be drawn and modified in either of these projections.

It is important to make sure that the diameter of the object selections is about the same as the diameter of the real objects, because the Object Tracker uses that information to tune the detection filters. The size of the background selection does not affect the detection.

When at least one object selection and one background selection is created, pressing the NEXT button will initialize the object detection and advance to the next step in the wizard.

Note that you can choose any slice for training. The chosen slice can be modified by dragging the colored slicer handles in the views.

During training, Linear Discriminant Analysis is used to find the best features to distinguish between object and background. In combination with the watershed segmentation method (See “Using the watershed segmentation” on page 162.), objects are defined. This process is done for each slice.

Optimizing Object Detection

After pressing the NEXT button in the object detection stage, the wizard will highlight the detected objects with a light blue overlay. A 'plus' symbol marks the center of each detected object (Figure 26.3).

In this screen it is still possible to add, delete, or modify the selections. In case the selections are altered, the NEXT button will reload this screen instead of advancing to the next stage.




Figure 26.3: Detected objects.

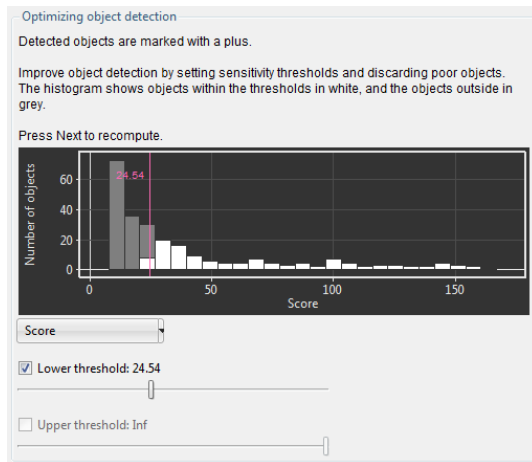


Figure 26.4: Increase the threshold on the score to reject more objects.

The right side of the wizard shows a histogram of the number of objects versus their *score*. The score is a qualitative measure for the probability that the detected object should indeed be classified as 'object'. By default all objects with a positive score are taken into account, but this threshold can be changed in order to make the Object Tracker less or more sensitive. To change the threshold, tick the *lower threshold* box and drag the slider (Figure 26.4).

Besides the score, thresholds can be applied to the *number of objects*, and the *width* and *brightness* of the objects. The width and brightness can be limited by both a lower and upper threshold. There is by default a threshold of 100 objects on the number of objects. An important thing to note is that *all* thresholds are applied, not only the currently selected one.

Besides the score, thresholds can be applied to the *number of objects*, and the *width* and *brightness* of the objects.

The width and brightness can be limited by both a lower and upper threshold.

When one of the thresholds is modified, the NEXT button will reload this screen instead of advancing to the next. When the detection is satisfactory, press NEXT again to advance to the next step.

Tracking Parameters

This is the final stage before the automated tracking starts. The parameters that can be modified in this screen affect the linking of the detected objects between subsequent time frames.

The first parameter limits the distance over which an object is allowed to move between time frames. This means that the tracker will not connect two objects in subsequent time frames if their mutual distance is larger than this value.

The second option enables or disables the use of object geometry and orientation for linking objects. Enabling this option makes sense if the objects of interest have distinct sizes.

Quick Edit Tracks

When all time frames are processed and the detected objects have been linked into tracks, the wizard advances to the *quick edit tracks* stage. In this stage broken tracks can be connected, and wrong tracks can be or broken down into smaller pieces that can be re-connected.

The graph shows the position of the objects in the xy plane versus the time frame number (Figure 26.5). The projection angle can be changed using the slider below the graph. An angle of 0° projects on the x axis while 90° projects on the y axis.

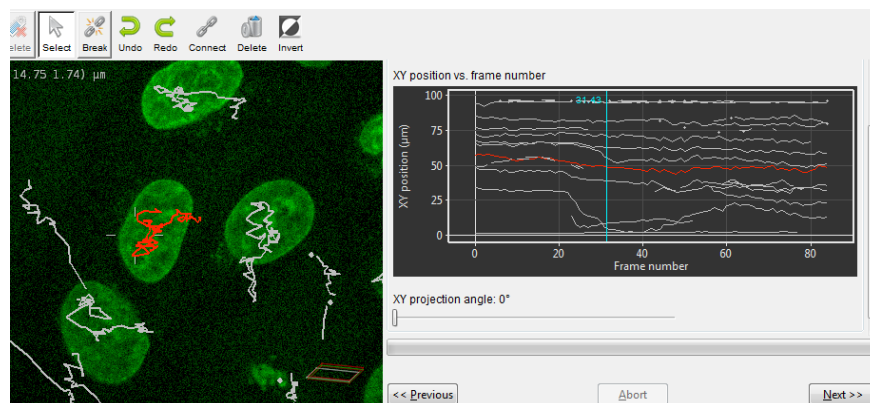


Figure 26.5: The *quick edit tracks* stage allows to connect, break, and delete tracks.

To break tracks down into smaller pieces, simply select the *break* tool (✂) and click between the nodes that should be disconnected. Multiple tracks can be selected using the CTRL key. If the nodes of the selected tracks do not overlap in time, then those tracks can be joined by pressing the *connect* button (🔗).

Pressing NEXT will show the last stage where you can save the tracks to a file, and/or save the tracker settings in a template. The tracks will be saved in a .tracks file and this file can be imported within the Track Analyzer. The tracker template can be used to quickly track similar images. The tracks can also be saved in the Track Analyzer, but the tracker template can only be saved at this stage.

Pressing NEXT again will close the wizard and load the tracks in the Track Analyzer.

The Track Analyzer

The Track Analyzer is automatically launched after the Object Tracker has finished processing the data, but can be started individually as well. Select the image which corresponds to the .tracks file to be analyzed and go to the main menu ANALYSIS → TRACK ANALYZER. The .tracks file can be imported through FILE → OPEN TRACKS....

The Filter Tab

The Track Analyzer can be used to analyze the tracks that have been generated by the Object Tracker. The scene on the top left displays these tracks, which can be saved via FILE → SAVE TRACKS. Viewing options can be selected under VIEW in the menubar. Tracks can still be edited using the break tool (✂) and connect tool (🔗), or the delete button. The window at the bottom left shows the collection of tracks and the objects for each track. Clicking on the track or object number highlights the concerned track in the image. Vice versa, the track number is highlighted when clicking on a specific track in the image.

The filter tool (Figure 26.6) can select a group of tracks based on common properties. These properties can be the number of *detected objects* (nodes) per track, *start time*, *duration*, the *path length*, *average speed* and the *average angular speed*. The upper and lower threshold sliders can be used to define a range within those properties. When the SELECT TRACKS button is pressed, all objects within the threshold ranges are selected. The INVERT button in the taskbar inverts the selection to the other tracks.

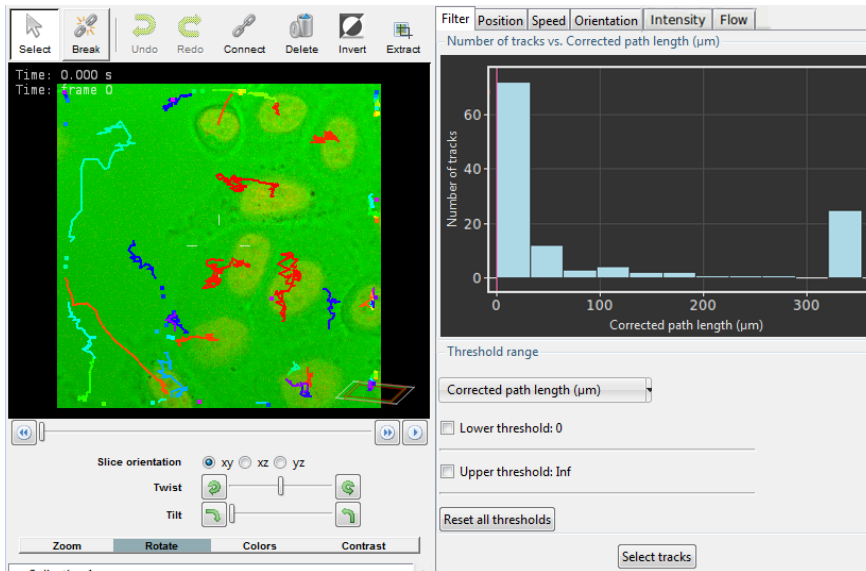


Figure 26.6: The filter tab in the Track Analyzer.

An important thing to note is that all thresholds are applied when SELECT TRACKS is pressed. Press the RESET ALL THRESHOLDS button to disable all thresholds.

To remove tracks with less than 4 objects, for example, select the *detected object per track* property, tick the *upper threshold* box, drag the upper threshold slider to a value just above 4, and click SELECT TRACKS. Now those tracks are selected. These tracks can be removed by pressing the DELETE button.

The Position Tab

This tab provides information on the position and displacement of the tracked objects. The histogram at the top, as shown in Figure 26.7, gives the number of objects versus the distance to their track origins for a specific duration. Example; if the time frame slider is set on 3 frames, and the height of the bar at a distance of 1.2 μm is 9, then there are 9 objects that travelled a net distance of 1.2 μm in 3 time frames.

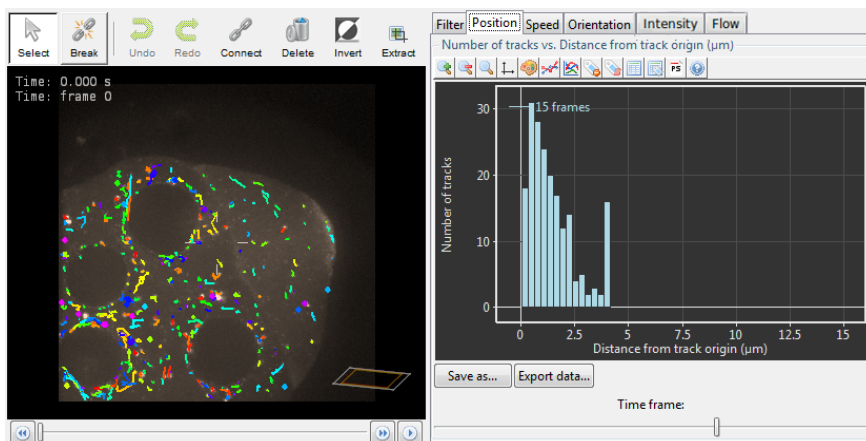


Figure 26.7: In case of a large number of objects that show Brownian motion, the shape of the histogram in the position tab will be a half-normal distribution.

The second and third plot in this tab are closely related. This second plot shows the average displacement for a certain duration, as given in Equation 26.1. In this equation t is the duration in frames and N is the number of objects. $\bar{x}_{i,t}$ is the position vector of object i at time frame t .

$$MD(t) = \frac{1}{N} \sum_{i=1}^N | \bar{x}_{i,t} - \bar{x}_{i,0} | \quad (26.1)$$

The third plot shows the mean squared displacement, or MSD , as given in Equation 26.2. Note that the x -axis of this plot shows the duration t in seconds instead of frames.

$$MSD(t) = \frac{1}{N} \sum_{i=1}^N | \bar{x}_{i,t} - \bar{x}_{i,0} |^2 \quad (26.2)$$

There is also a close relation between the MSD and the histogram in the first plot (Figure 26.7). In case of pure Brownian motion and a large number of objects, the histogram will show a half-normal distribution. In this special case, the variance of that distribution is given by $MSD(t)$. The rate of change of this variance is known as the diffusion coefficient D , which thus can be estimated by measuring the slope of the MSD plot. Just click and drag in any of the plots to measure distances and slopes (Figure 26.8).

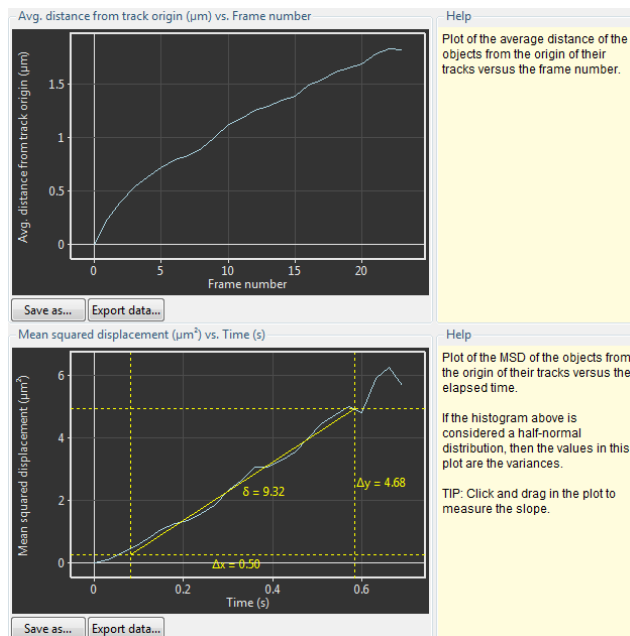


Figure 26.8: The slope of the mean squared displacement plot indicates a diffusion coefficient of $9.32 \mu m^2/s$.

The Speed Tab

This tab displays information on the velocity of the tracked objects. The first plot is a histogram that shows the distribution of the speed of the objects, while the second plot shows the average speed over time.

The third plot informs about the velocity distribution of the objects. This not only gives information on the speed (a scalar value), but also on the direction of the velocity, as shown in Figure 26.9.




Figure 26.9: Left: An *xy* velocity scatter plot showing a drift in the positive *y* direction. Right: An *xy* velocity scatter plot of objects moving with a constant speed of about 0.35 $\mu\text{m/s}$ in random directions.

The Orientation Tab

The first graph under this tab shows the axial rotation of objects in time, while the second graph show the 3D angular speed versus time (indicated as frame numbers).

The Intensity Tab

Under this tab, you can see the intensity of the voxel, which represents the center of mass of the object that is under the selected track, as a function over time. If multiple tracks are selected the average intensity of all the center-of-mass voxels of these tracks are plotted.

The Flow Tab

This tab can be used to gain insight into the amount of objects that cross a certain boundary over time. Boundary surfaces can be defined as a 2D poly line in the *xy* slice image. These lines are extruded to planes in the axial direction. To start drawing one, click the DRAW BOUNDARY button (📏), and then click on the first point of the boundary. A single click adds a new line segment and a double click finishes the poly line.

When a boundary has been defined, then the plot in the Flow tab shows the number of objects that cross this boundary at a certain time frame (Figure 26.10).

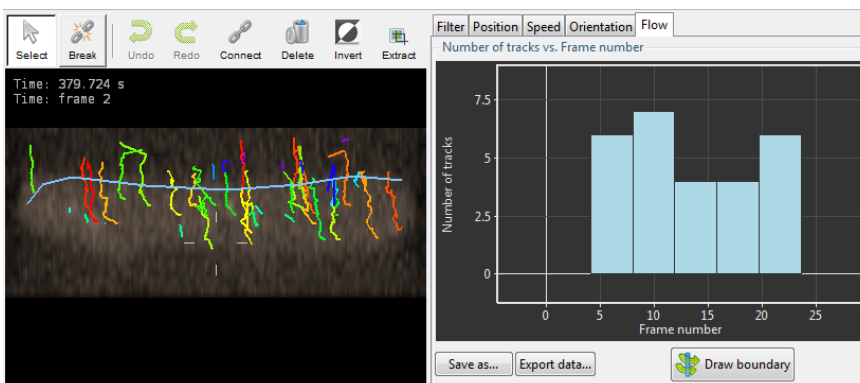


Figure 26.10: The flow histogram shows that objects are crossing the boundary between time frame 4 and 24.

Exporting Data

The plot graphics, as well data in the plots and the track data can be exported to a file. The buttons below each of the plots, `SAVE AS...` and `EXPORT DATA...`, save the plot as an Encapsulated PostScript file (EPS) and export the data as comma separated text (CSV), respectively. The track data of the selected tracks can be exported to a CSV file using the menu `FILE → EXPORT → SELECTED TRACKS AS CSV`.

The EPS format can be imported by most office programs, including Microsoft Office and OpenOffice. The CSV format is plain text and can be parsed by most spreadsheet programs and other software like Matlab and Mathematica.

Chapter 27

Introduction to the Object Analyzer

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 exclude objects that are too small before starting the analysis, a garbage level can be used. Any object smaller than this garbage volume will be discarded. After removing these small objects, all remaining objects are automatically labeled and sent to a continuous iso-surface renderer (See Chapter 17 on page 107).

The Object Analyzer is available as an option, 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 user interface components of this option that can be found on page 173.

Starting the Object Analyzer

- Launch Huygens Essential
- 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 on the SVI Wiki¹, from where you can also download the test image we will use in the following steps. You can explore the image with the Twin Slicer (See Chapter 13 on page 83) for an initial impression of the data. The demo image is a deconvolved image. It is always recommended to perform deconvolution before analysis. Deconvolution will reduce noise and blurring significantly and will thereby improve the analysis quality.




Figure 27.1: 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 27.1 is a MIP projection of the example 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 dim with the exception of the inner bright objects and some increase in intensity in the peripheral area, making something like a border.

When the image is opened in the OA, a first view and segmentation is presented (Figure 27.2). The initial segmentation is based on Otsu's segmentation.

A large bright object is visible in red and viewed from the top. The image was automatically rendered with default settings. If more than one object is present, different colors will be assigned to each object. The specific color range can be adjusted with the hue slider.

¹<https://www.svi.nl/ObjectAnalyzerTutorial>




Figure 27.2: Start up screen of the Object Analyzer. The object in the image’s first channel is automatically segmented with initial parameters calculated based on the intensity distribution (Otsu’s segmentation).

Segmenting the Objects: Setting the Threshold

The segmentation method in the Object Analyzer uses a combined seed-and-threshold method, see Figure 27.3. The seed acts as a secondary threshold level, so that objects that do not reach this intensity (in at least one voxel) are discarded. More details on how the threshold-seed segmentation works can be found in the expert online tutorial².





Figure 27.3: Visual representation of the seed-and-threshold segmentation method.

Let’s try different segmentation parameters and see what happens. First we increase the threshold value from the automatically calculated value. You can drag the blue line in the channel histogram and shift it to higher values, or click on the blue-font label showing the threshold value to enter any number. The threshold lines on the histogram are found at the right of the window, in a pane labeled *Object Segmentation*.

The magenta-colored line is the seed level. When you shift the blue line, the magenta line also moves, as by default the two are linked.

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, you can also control whether the seed is linked to the threshold value or not, by enabling or disabling the check box next to the *Seed* slider.

²<https://www.svi.nl/ObjectAnalyzerExpertTutorial>

When we shift the threshold to higher values the rendering scene will change. First, the object we had defined will shrink in size and we see many smaller objects starting to appear. Some objects that were connected earlier shrink so much that they now get separated, and they define new objects. To visualize the object details better, you can increase the zoom factor up to 1.30 to see what is shown in Figure 27.4.




Figure 27.4: Different segmentation conditions. Increasing the threshold decreases the objects size, and adds more objects to the scene.

The increased threshold breaks the single large object into multiple smaller objects. The very small objects can be easily removed by setting a garbage volume.

Setting a Garbage Volume level

A quick way of removing small objects, is the *Garbage Volume* option. 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 is shown below the seed slider. Objects with a voxel volume below the garbage level are discarded. This means that when you set it to 1, no object is discarded, but if you set it to for example 100, any object with a volume smaller than 100 voxels will be removed from the scene and from further analysis.

Additionally, you can apply some post-segmentation filtering. Details on this are explained in the expert tutorial. In the top menu you can also find some predefined filters for discarding objects. These filters are based on the geometrical properties of individual objects or on how objects are related to other objects.

Using the watershed segmentation

When using the seed and threshold method, it is possible that objects merge at a certain threshold level. To further tune segmentation, you can apply a watershed segmentation.

The watershed segmentation builds “watersheds” at local minimum values to separate merged objects. This can also be explained as a flooding algorithm where the inverse intensity profile is flooded with water. A watershed is build in order to separate the different water-reservoirs. The difference between threshold and watershed segmentation is illustrated in Figure 27.5.

The *Seed* setting has the same meaning for the watershed as it has for the threshold segmentation method, however, the garbage volume filtering uses a slightly different approach. The garbage filter removes objects with a volume below the given value, but if this concerns an object that was separated from an other object by the watershed algorithm (a neighbor) it will first be merged with this neighbor, after which the volume of the new merged object is again reviewed. This merging process will be repeated with new neighbors until it has no direct neighbors. If the volume is it still below garbage level and it has no neighbors anymore, it will be removed, otherwise the complete object is kept.

The watershed segmentation method uses the sigma setting as an extra input. This setting is used to avoid over-segmentation of the data by smoothening intensities with a gaussian filter. Especially with very noise images, local minimum value are frequently present, resulting in many small objects. The Object Analyzer can handle at most $32768 (= 2^{15})$ objects.

If you want to use the watershed segmentation, we recommend to first use the simple seed-and-threshold segmentation. After setting the threshold correctly, the watershed can be activated. Due to its complexity, the watershed segmentation may take some time, especially in comparison with the more simple seed-and-threshold segmentation.

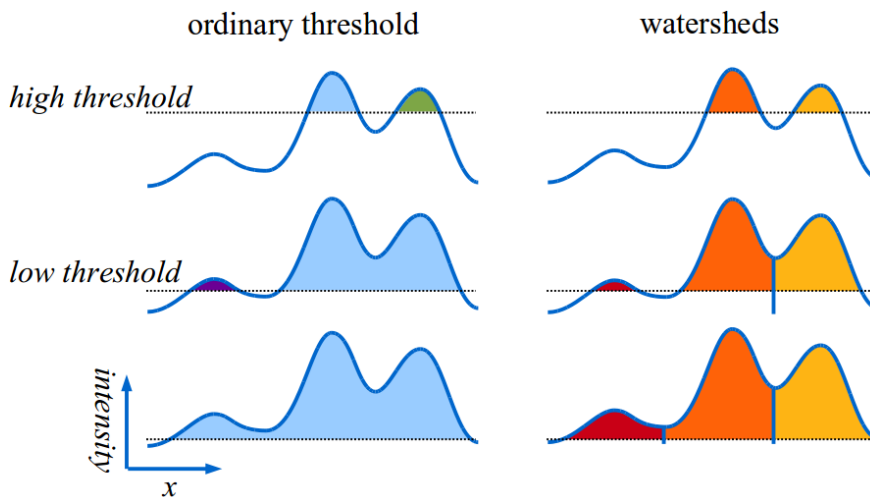



Figure 27.5: Difference between the seed-and-threshold and watershed segmentation methods.




Interaction with the Objects

On the left side of the Object Analyzer 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 icons in the first group are always available. The first group of icons control the *Mouse mode*, i.e. the action of the mouse when you left-click in the scene. A more detailed description of these icons are described in “Mouse Modes (2)” on page 176.

The buttons in the other groups (Objects, ROI and Data) 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 173 for a detailed description of these toolbox buttons.

At the end of the buttons column there is a colored reference cube that will help you to orientate in space when you rotate the dataset, especially 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 mode. Click on an object and see what happens. Notice that some parameters are reported on the table at the bottom of the window. Later in this tutorial we will learn how to analyze more parameters and more objects at once.

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

Render Pipes

When we opened the analyzer, only the first channel of the image is shown on the screen. However, we can change this by simply selecting another channel from the drop-down menu in PRIMARY PIPE: where it reads *Chan 0* you just select *Chan 1* (See Figure 27.6).

The term *pipe* suggests that data flows 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 Object Analyzer to direct data through.

You can activate the control pane of the secondary pipe 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*), also abbreviated as P and S. The primary and secondary pipe should not be confused with the numbering of the image channels. For example; you can even segment the same channel twice, but with different conditions in each of the two 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 220.). If you want all objects in a pipe have the same color, you can collapse this range completely by moving the two triangular sliders together.

There is also a maximum intensity projection (MIP) pipe. This pipe does not interfere with the analysis, but can be used to set a visual reference of the (unsegmented) data.

Before continuing our exploration, let’s enable 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 27.7.




Figure 27.6: Rendering pipes control panel.





Figure 27.7: 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 is able to measure many different parameters, but only a few are computed and reported under the default settings. To obtain information about a specific object, click on this object while you are in the *Analyze object* mouse mode. To automatically process all the segmented objects in all pipes, simply press the ANALYZE ALL button  on the top-right of the table. (If you have a region of interest (ROI) selected on the screen which encloses a few objects, this button analyzes only the objects inside this ROI).

There are many additional useful parameters that you can choose from. Let's have a look at how to retrieve all this information.

Configuring the Reported Parameters

To simplify the usability of the Object Analyzer, 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 selection dialog as shown in Figure 27.8.

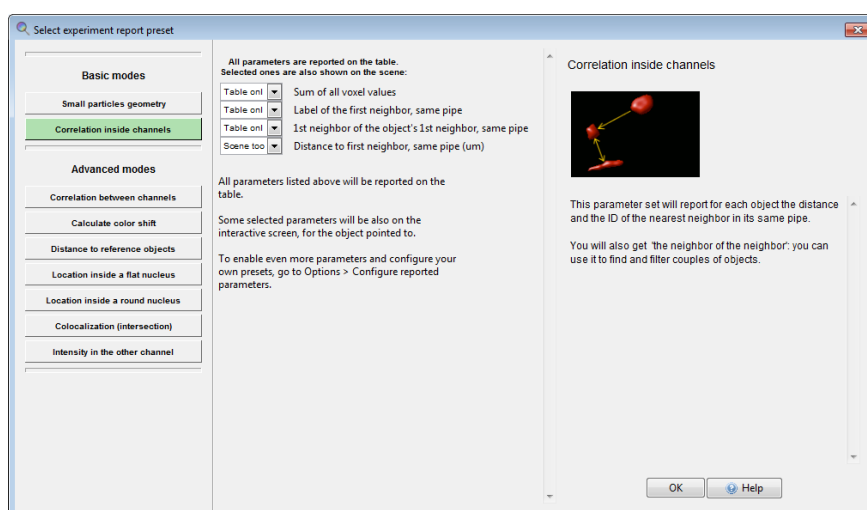


Figure 27.8: The Experiment Presets dialog. Collections of parameters can be selected to be reported on the table and the scene, depending on your specific experimental needs.

The left column shows a series of different experiments. If you click on them, a new list of parameters is listed in the middle column, and a description is shown on the right column. When you hover with your mouse over the listed parameters, a tooltip explains each parameter. Please read the descriptions briefly to get an impression of them. Let's select the preset called CORRELATION INSIDE CHANNELS and inspect it in more 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 pipe.”

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

Notice that among the listed parameters (center column) of this preset, only one of the parameters is selected to be reported in the scene. All parameters will be calculated and reported in the results table, but only the selected ones will be shown on the rendering canvas as well when you interact with the objects. In this parameter set, the distance is a useful parameter to report in the scene, as it provides a visual reference of the distance measurement between neighboring objects.

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 an object in the scene and notice the new columns that appear in the statistics table.

Before showing the actual object details, the first row in the table informs you about the segmentation conditions for this pipe. If you move your mouse over the column header, a description of the corresponding parameter will be reported at the very bottom of the window. You can always find out what each parameter abbreviation stands for by looking at its tooltip. The description of each parameter will be also stored in your file when you export the results table later.

The tooltip of the column header also shows the sum and the average value of all the cells in the column. It is possible to include data of two pipes, or a single given pipe only.

Click on another object to apply a measurement and add the results to the table. You will see that the segmentation conditions are not reported again, because they have not changed. A checkbox option at the top-right of the table can be used to automatically clear the table when the segmentation or report conditions change, so you always have an organized table. When you deselect this option, all the results (despite different segmentation conditions) will accumulate in the table. The button located next to the checkbox allows you to clear the table at any time.




Figure 27.9: Interaction with the objects. When the distance to the nearest neighbor is computed, the distance line is also displayed on the screen for the object aimed at.

You may have noticed that something else happens when you interact with the rendered objects. In the example of Figure 27.9, 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 it is the only that was selected from the preset selection dialog to be shown in the scene.

Other Measuring Parameters


The presets are organized in **basic** and **advanced** modes.

We can distinguish two basic parameter presets: the one we have used here to explore the nearest neighbors, and another one to retrieve morphological parameters about objects, called SMALL PARTICLES GEOMETRY. This parameter preset can report object information like *length*, *width*, *aspect ratio*, and *sphericity*. Details on how these parameters are defined can be found on “Object Analyzer Geometry Measurements” on page 169.

The advanced modes require that you define some conditions first. For example, that you first define a ROI. The advanced modes are not intrinsically more complex, but they demand from the user some knowledge on how to set these reference conditions.

For more details about any parameter preset just click the HELP button 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 calculations are done (it should be quite fast in this example) move your mouse over the table rows and see that the cursor highlights the table rows (see Figure 27.10). Furthermore, 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	27.868	-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 878 x 577 Dragging on 38/S 1NP.CMCM: 4.2358 um

Figure 27.10: Exploring the table.

You can find an object in the results table by clicking on the object in 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.

You can 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 27.11, for example, shows a plot of the particle sphericity distribution.




Figure 27.11: Histogram representation of one column values.

Storing your Results

In the top menu you can select **FILE** → **SAVE OBJECT STATISTICS** to export the table to a file for further analysis and/or calculations. The table can be stored as a plain text file, a csv-file, that can be imported in e.g., Microsoft Excel, LibreOffice Calc, GNUplot, and as a m-file, that can be imported in Matlab.

The current scene as it is currently visible on the rendering canvas can be saved to a TIFF file using **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. Such a template can be used on the same dataset to reproduce the analysis, or on other datasets.

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 and the accompanying wiki pages, where you will learn many other powerful things you can do with this versatile analyzer.

A reference description of the Object Analyzer components can be found in Chapter 29 on page 173.

³<https://www.svi.nl/ObjectAnalyzerExpertTutorial>

Chapter 28

Object Analyzer Geometry Measurements

Iso-surface

In the Object Analyzer, the threshold for the segmentation is 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 17.1 on page 107).

Based on the iso-surface, 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 $Waist-Lat$.

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 which 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 28.1).

$$AxSphPB = \left(\frac{pBoxW0 \cdot pBoxW1 \cdot L}{L^3} \right)^{1/3} \quad (28.1)$$

2. Another definition (reported as *AxSph*) involves the lateral and axial widths discussed above, obtained with the virtual caliper algorithm (Equation 28.2).

$$AxSph = \left(\frac{L \cdot WiAx \cdot WiLatC}{L^3} \right)^{1/3} \quad (28.2)$$

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:

$$SurfSph = \frac{\pi^{1/3} \cdot (6V_i)^{2/3}}{A_i} \quad (28.3)$$

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 additional 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.

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

Chapter 29

Object Analyzer Component Reference

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 159 to learn how to use the basic components in context.

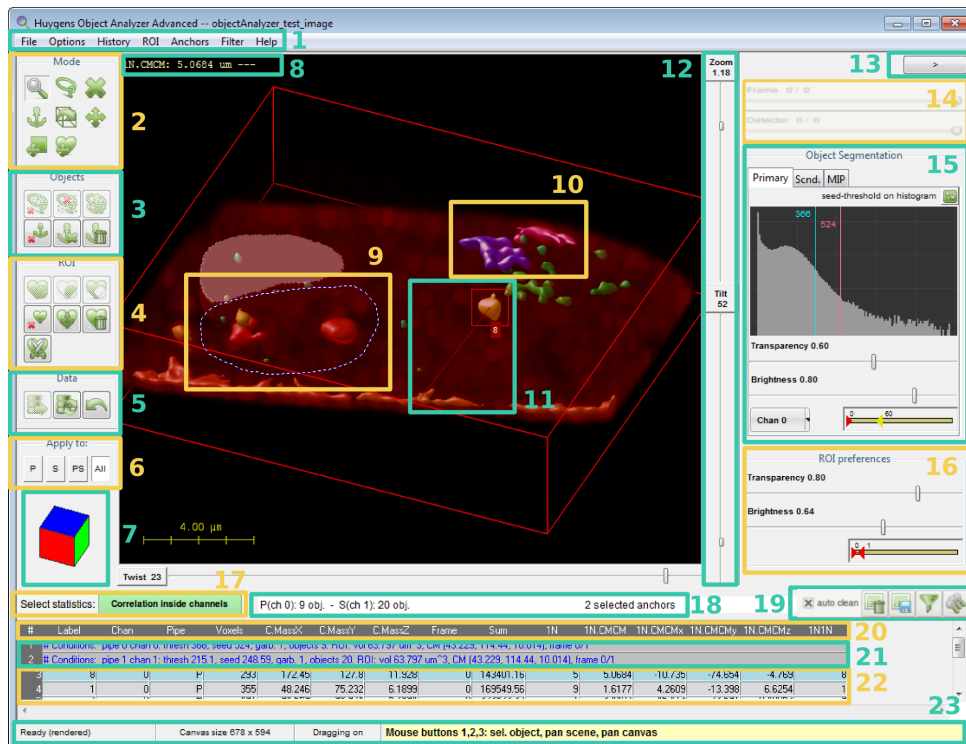


Figure 29.1: The Object Analyzer interface. Different regions are enumerated to describe them in sections of this chapter.

Main window components

The OA main window is shown in Figure 29.1 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). This way 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 179).
- 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 (17)" on page 182).

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. The 2D area is projected along the line of view.
 - 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 178).
 - 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 in one channel intersect with the other. A ROI can defined, and applied to discard objects and parts of objects outside the region, so that only the colocalizing intersecting 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: sets the ROI to the complete dataset volume.
- 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:
 - Export current ROI to the main window, to have the ROI made available for other tools in Huygens.
 - 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

- Anchor objects under selection.
- Select all objects as anchors.
- Set anchors by filtering. This opens a filter dialog as explained in “Table and Analysis Shortcuts (19)” on page 182, but allowing you to select or deselect anchors instead of discarding objects.
- Invert current anchor set.
- Deselect all anchors.
- 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 182.







HELP: Shows on-line help and tutorials.

Mouse Modes (2)

These buttons control the Mouse mode, this defines the mouse functionality when you left-click on the image. Hovering your mouse over these buttons allows you to get a tip on what they are for. If they have a keyboard shortcut, this is also shown in the tooltip. For example, the first mouse mode is intended to *analyze objects*, and you can always activate this mode from within the rendering view by pressing the keyboard key ‘1’.

From left to right, and from top to bottom, the mouse modes are:




-  **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. This allows you to 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 uninteresting 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 with respect to the anchor object. When you set an object as a reference anchor it will ‘light up’ and change color on the screen to indicate the anchor 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.
-  **Pan canvas** This mode allows you to explore the canvas while not re-rendering the scene. This only makes sense when you have a canvas larger than your rendering window, 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’. 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. When a mouse mode is active and the cursor is inside the canvas, the tooltip in the status bar (See Figure 29.1, item 23) shows you what each mouse button can do.

Selection Interactive Operations (3)

The rest of the buttons in the toolbar are not mouse modes, but they 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 179) 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 all objects under the selected area as anchors.

Some objects may partially cross the border of the defined ROI. These can be considered to be either in- or outside the ROI. You can change this definition 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 a lower intensity to clearly indicate which objects are selected. When you have selected certain objects as anchors, you can apply further operations:





-  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 these anchor objects. But you can also use the selected anchors to delete some other objects, or define a 3D ROI based on the anchor object shape. When you have defined a ROI, you can add/subtract new ROI shapes, or perform additional operations with the ROI Interactive Operations buttons:


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.
-  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 this button to crop the dataset and delete (set to zero) all voxel values outside the ROI.
-  Export the current data (as cropped by the ROI, or after deleting objects interactively) to the original image.
-  If 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, on which of the active data pipes (*Primary*, *Secondary* or both) and channels (*Selected pipe*, or *All*) they operate. Pipes are explained in “Render Pipes (15-16)” on page 181.

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 you can notice a colored reference cube that will help you to orientate 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 tracing 1 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 these parameters are taken into account by the renderer, which generates the scene and displays it in the canvas on the screen. Notice that depending on the render size (OPTIONS → VIRTUAL RENDER SIZE) the canvas can be made larger than your screen.


The On-Screen Reported Parameters and Tooltips (8)

The currently selected Experiment preset (See “Experiment Presets (17)” on page 182) selects the statistics to report many parameters to the table (See Figure 29.1, 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 indicated by a triple dash ‘—’

```
Sum: 136707.75
IN. CMC: 5.0821 um ---
```

The Selected Area (9)

The *Select* area  mouse mode allows you to define a region in the image. So for example you can: analyze objects ‘inside’ the ROI, anchor them as references, or use the selected volume to define a region of interest (ROI).




The *selection* is not yet a 3D ROI, but 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’ the current selection. It allows quick and simple interaction with the objects, specially on flat images here the objects remain more or less in a plane.

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 the other objects.

There are different ways to set objects as anchors, interactively (See “Selection Interactive Operations (3)” on page 177) or by using filtering operations (top ANCHORS menu). You can also visit the on-line article about anchor objects for more details on their utility.



Aim (11)



Objects under your cursor are highlighted. Once a certain object is analyzed, the object will also show a red box surrounding the object. To analyze an object just click on the object with the *Analyze object* mouse mode, or press the ANALYZE ALL button (See Figure 29.1, 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 29.1, 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 178), 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.

Time frame and detector slicers (14)



When time series or array detector images are loaded in the OA, these sliders control which time frame of detector is used for visualization and analysis.

Changes to the current frame and detector (like discarding objects or selecting anchors) are remembered when you change the time frame or detector. 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, and likewise for detector series through ANALYZE DETECTOR SERIES.


Render Pipes (15-16)

The Huygens Object Analyzer has two surface 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 29.1, item 16). ‘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*. 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.

This pane controls which data channel goes through which pipe, how its objects are segmented (with a seed and a threshold¹), and how the data is rendered (transparency and brightness). A garbage volume can also be set to remove little spurious objects: objects with a volume smaller than this garbage volume will be discarded and not rendered.

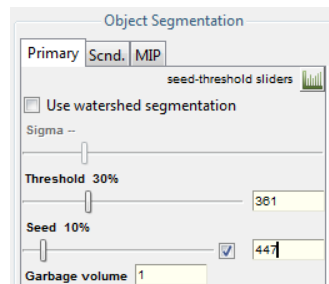

You can see the seed as a secondary threshold. The first threshold segments the data and makes independent objects, but then only objects with intensity that reaches above the seed level will remain, while the rest are discarded.

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 is also available. This alternative representation 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 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.



¹<https://www.svi.nl/SeedAndThreshold>

All the different objects in a pipe are colored differently to distinguish them. The range of colors assigned to each pipe can be controlled with a *Hue Selector* (See “Hue Selector” on page 220.). You can collapse this range completely if you want so all objects in a pipe get the same color.

Experiment Presets (17)

Select statistics: **Correlation inside channels**

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 27.8 on page 165).

In this dialog you can notice 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, a tooltip will appear, 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 183). The check boxes allow a few of these parameters to be also reported on the screen (See Figure 29.1, 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.

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 provide 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

-  Save statistics table: manually save 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 29.1, item 6). When a 2D area has been selected (See Figure 29.1, item 9) this buttons analyzes only the objects under the selection.

Filtering Objects

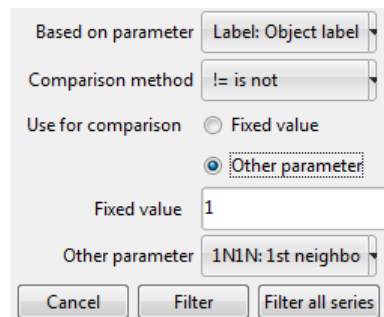
The Filter button  opens a pop up dialog that allows you to discard objects:

- 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 Wiki².

For time series, an additional button allows 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	Voxels	C.MassX	C.MassY	C.MassZ	Frame	Sum	1N	1N.DCMD	1N.DCMDx	1N.DCMDy
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, garb. 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, garb. 1, objects 17. ROI: vol 0.34587 um ³ . CM (170.13, 213.42, 9.1209), frame 0/1:												
3	1	0	P	365	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 on your disk by using the FILE menu.

²<https://www.svi.nl/ObjectAnalyzerNeighbors>

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 a description of each parameter will be shown at the very bottom of the window (See Figure 29.1, item 20). 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 29.1, 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 the 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 (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 29.1, 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 pop up 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 um3, 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 ROI ³.

The Table Rows, How to Explore Them (22)

Each row in the table is either a *Conditions report* (See “Conditions Report (21)” on page 184) 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 178). 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 29.1, item 11). Objects in the Primary pipe are highlighted in red, and objects in the Secondary pipe in green.

A quick way to find the parameters of an object in a very long results 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, which will be highlighted.

When you move over the table cells the current value is shown on the tooltip bar at the bottom (See Figure 29.1, item 23). This, combined with the selection of rows (see below) or columns (See Figure 29.1, 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 μm . 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. The value of the distance itself is also shown: 4.23 microns.

³<https://www.svi.nl/ObjectAnalyzerROI>

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

That tooltip region can also show other values that can be used 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 well separated 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	28	58.195	72.512	14.466	0	5065.34	8	0.88804	-11.211	-4.395	-1.1794
46	35	1	S	40	60.92	87.086	14.049	0	7706.94	9	1.5203	-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. It is possible to highlight columns and rows at the same time but this is confusing to interpret, and thus not recommended.

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 the scene. In slower 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 more detailed description of the reported parameters when you hover over the table headers (See Figure 29.1, 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.

Chapter 30

The Colocalization Analyzer

The *Colocalization Analyzer* provides visual and quantitative information about the degree of overlap between structures in different data channels, for 3D images and time series. The Colocalization Analyzer makes use of GPU acceleration for floating point datatype images. This significantly increases the speed at which colocalization analysis is executed.

Huygens supports a wide range of colocalization coefficients reported in literature: (*Object*) *Pearson*, (*Object*) *Spearman*, *Overlap*, *Intersection*, *Manders M* and *K*, *Li's ICQ*, and *van Steensel's CCF*. The SVI Wiki provides extensive information on these coefficients ^{1 2}. In addition to that, the Huygens Colocalization Analyzer calculates a 3D colocalization map for a selected coefficient. The voxel values in this 3D distribution represent the colocalization locally by showing the contribution of each voxel to the colocalization coefficient.

Note that the Object Analyzer module (See page 159) also provides a form of colocalization measurement that is different from the functionality of the Colocalization Analyzer. The Colocalization Analyzer determines colocalization on a complete image, then it segments the colocalizing regions of the 3D colocalization map into objects for further analysis. The Object Analyzer *starts* at the object level; it first segments the original image into objects and then determines if these objects colocalize.

The Colocalization Analyzer is an optional module, and is enabled by the C flag in the license string (See “License String Details” on page 226.).

Overview of the modes

In Huygens Essential 18.10, there are three options for the Colocalization Analyzer: The Wizard, Advanced, and Compact mode. All three modes consist of a 2D histogram of the image intensities in the top left. Next, the top right shows an interactive viewer panel for rotating and viewing the *Maximum Intensity Projection* (MIP) of the image in 3D (See Figures 30.1 - 30.3). Finally, all three modes contain menus for selecting the coefficients, thresholding options, and drawing a region of interest (ROI).

Wizard

Learn about the possibilities of the Huygens Colocalization Analyzer by starting with the Colocalization Wizard. The stages in the wizard are: Frame/Channel selection, Threshold settings, ROI selection, Coefficient selection and Coefficient results with object analysis.

¹<https://www.svi.nl/ColocalizationTheory>

²<https://www.svi.nl/ColocalizationCoefficients>

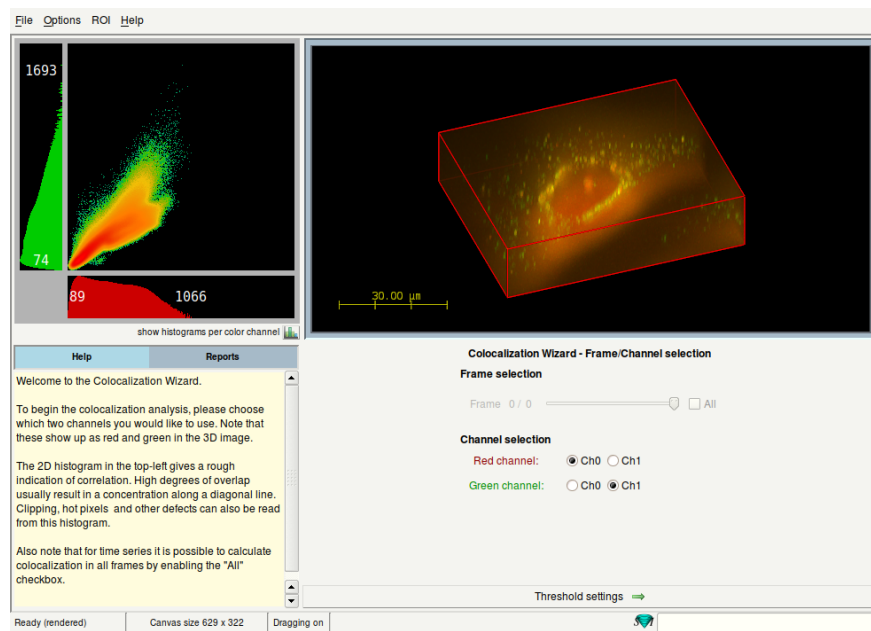


Figure 30.1: Colocalization Wizard with separate stages for each step in colocalization analysis.

Advanced

The Colocalization Analyzer Advanced provides a large overview of all the options. This reduces the number of times clicking back and forth between tabs/wizard stages. The left side shows all the available setup options that you can perform before computing the colocalization including intensity thresholding, region of interest (ROI) selection, and coefficient and colocalization map selection. The top right provides a 3D view of the image including render settings for visualizing the image and colocalization. Finally, the bottom right holds all the coefficient results and a section for object-based colocalization analysis. These objects are created from segmenting the colocalizing regions in the colocalization map.

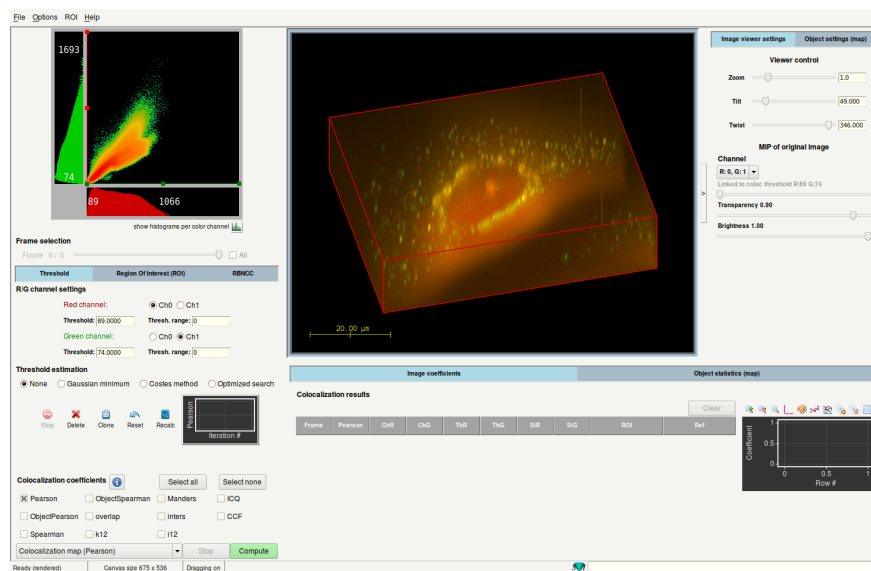


Figure 30.2: The Colocalization Analyzer Advanced consists of three areas: setup options on the left, visualization options on the top right, and analysis of results on the bottom right.

Compact

The Colocalization Analyzer Compact features a tabbed window with the setup options and result section in the control panel. This allows the window to remain compact, while still having the same features as the Advanced mode.




Figure 30.3: The Colocalization Analyzer Compact with tabs for each step in the colocalization analysis.

How to use the Colocalization Analyzer

Colocalization is as much about co-occurrence as it is about correlation of intensities from two different fluorophores. To obtain reproducible results, it is important that the correct sample size is used during image acquisition, as undersampled images can lead to an over-estimation of the degree of overlap. Furthermore, crosstalk and chromatic aberration will also influence the colocalization measurements, and need to be addressed first with the respective options (See page 129 and 137).

To *start* with colocalization analysis in Huygens Essential, right-click an image thumbnail to open the contextual menu, then move to **ANALYZE IMAGE** and select one of the modes. Alternatively, left-click a thumbnail and go to **ANALYSIS** → **COLOCALIZATION ANALYZER** in the top menu bar and select the mode you want from there. The image must be a multi-channel image (See “Multi-channel Images” on page 23) as colocalization is based on the overlap of different channel intensities.

In order to avoid over- or underestimation of your colocalization and get reproducible results it is good practice to only measure in the area of the image that contains your sample. This can be achieved by either using a threshold, which ignores voxels that are below a certain value (i.e. are part of the background or unimportant empty regions of the image). Another way to only measure over the interesting regions is to draw a region of interest (ROI). Both of these methods are available in the Huygens Colocalization Analyzer and are discussed later in section “Threshold settings” on page 190 and “Region of interest selection” on page 192.

The colocalization map and the selected coefficients are calculated by pressing the green **Compute** button. The results table contains the selected coefficients and the viewer will overlay the segmented objects in the colocalization map. The following sections provide details on how to use the available setup options and how to perform object-based colocalization analysis in the Colocalization Analyzer.

Time frame & channel selection and the 2D Histogram

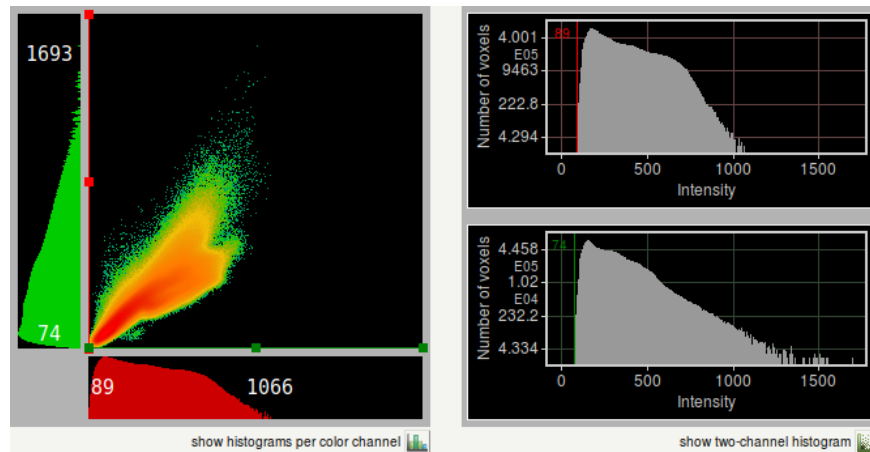


Figure 30.4: The two different histogram layouts. Left: the 2D histogram; right: the two 1D histograms shown separately.

This section of the tool allows for the selection of the time frame and channel indices of interest. For the channels the usual naming convention in colocalization theory is followed: *Red* (R) for the first channel, and *Green* (G) for the second channel. Whenever the time frame or the combination of Red and Green channel selection is changed, the 2D histogram is recalculated and displayed in the histogram panel. 1D representatives of both channels are displayed at x and y axis of this 2D histogram. The 2D histogram provides a rough indication of the degree of overlap between the selected channels. For two channels with a high degree of overlap, the histogram pixels concentrate along the diagonal $y = x$ line. In contrast, total absence of overlap will produce a 2D histogram with values only along the horizontal and vertical axes. The histogram also shows image defects such as clipping, hot pixels, and dynamic range differences. For more information on interpreting the 2D histogram see the examples on the SVI Wiki³. In addition, separate 1D histograms are generated for each channel. They show up after clicking the SHOW HISTOGRAM PER COLOR CHANNEL toggle-button near the bottom right corner of the histogram (See 30.4).

By checking the “All” checkbox next to the frame selection slider, the colocalization coefficients are set to be computed for all frames in the time series.

Where to find it: In the Advanced and Compact modes the histogram and frame/channel selection are on the top left below the 2D histogram. In the Wizard, these options appear in the very first stage after start up.

Threshold settings

Thresholding is a useful method for separating object voxels from uninteresting parts of the image. The *threshold* values are subtracted from all image voxel values. In case negative voxels occur, their values are interpreted as zero. These threshold values can be set manually by moving the colored threshold lines in either the 2D-histogram, the two 1D-histograms, or by changing the numeric values in the input fields.

³<https://www.svi.nl/TwoChannelHistogram>

The *threshold range* is used to set an intensity region around the threshold value in which a soft thresholding is applied through a linear weighting function. To see how these values influence the calculation of the coefficients, please refer to the SVI Wiki⁴.

Huygens supports several estimation methods to automatically set the threshold values. The different methods are: Gaussian minimum, Costes method, and Optimized search.

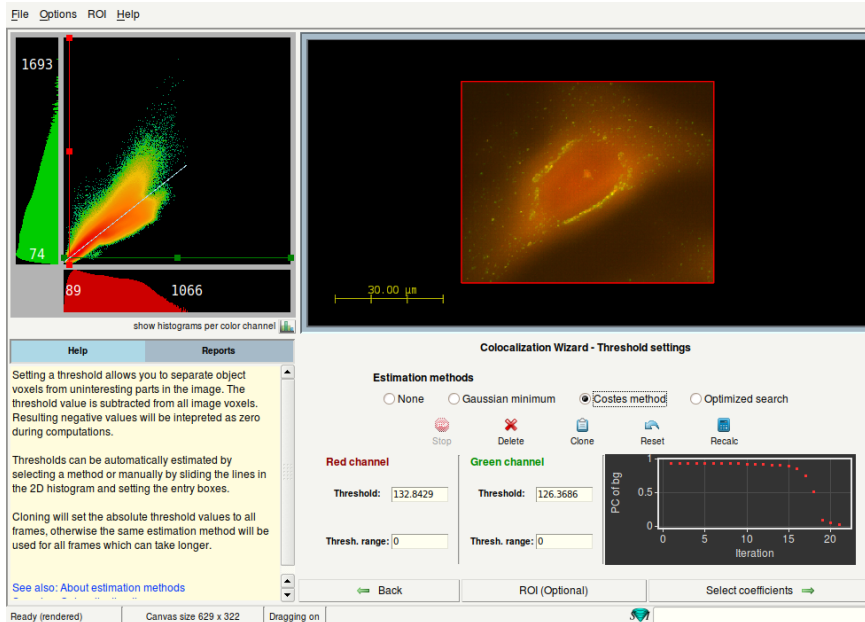


Figure 30.5: *Threshold settings* stage as it appears in the Colocalization Analyzer Wizard. The Costes method is selected, and a regression line is visible in the 2D histogram.

Gaussian minimum

The *Gaussian minimum* is the simplest threshold estimation method. The threshold setting is the minimum intensity value after applying a Gaussian filter on the image. This method is relatively fast and works well for images with low noise levels.

Costes method

The *Costes method* is based on the estimation method explained in the article by Sylvain et al.⁵. This approach consists of three steps that are repeated iteratively to reach the threshold values. First, a linear regression line is calculated that best describes a positive correlation in the 2D histogram data. This is done with a least squares fit with perpendicular offsets. Second, the highest point on the regression line is taken to set upper intensity limits for both channels. Third, the Pearson coefficient is calculated over the voxels with intensities *below* these limits.

This process is repeated moving down on the regression line. After some iterations the upper limits reach the point where the Pearson coefficient is only being calculated over lower signals such as noise. Assuming these signals are Poisson distributed, we expect no correlation. In other words, the threshold values are found when the Pearson coefficient is closest to zero and only voxels with intensities above the set threshold values are used. The plotting window will show the Pearson coefficients found at each iteration (see Figure 30.5).

⁴<https://www.svi.nl/ColocalizationCoefficients>

⁵Sylvain et al. V. Costes, Dirk Daelemans, Edward H. Cho, Zachary Dobbin, George Pavlakis and Stephan Lockett; Automatic and Quantitative Measurement of Protein-Protein Colocalization in Live Cells”, Biophysical Journal (2004) 86:3993-4003).

Please be aware that for large data sets, this method can take a substantial amount of time.

Optimized search

The *Optimized search* is an extension of the *Costes method*, but without the assumption that the ideal threshold combination lies on the regression line. An iterative process is used to search the entire 2D histogram to find the point where the Pearson coefficient of the area below the threshold settings is zero. This iterative method uses a searching algorithm (the Nelder-Mead/amoeba method) that converges to the desired point. While it is, in general, more precise than the *Costes method*, the time required to reach convergence is highly dependent on the size of the input data set. For both the *Costes method* and the *Optimized search* it is assumed that the noise in the image is Poisson distributed, because the Pearson coefficient for two Poisson distributions is zero. Similar to the *Costes method*, the obtained coefficients for each iteration are plotted in the plotting panel upon completion of the search.

Additional buttons

Besides the threshold estimators, there are some additional options available to set the threshold settings, in particular for different time frames.

- **Stop:** This button is disabled by default, but when an estimation method is started, the process can be stopped by pressing this button. This will not invalidate the estimation process and it will return the threshold settings which were found before the Stop button was pressed.
- **Delete:** This will delete all the threshold settings for all time frames at once and set them at the minimum values of the time frame.
- **Clone:** In case the threshold is correctly set for a time frame, this will copy the current threshold values to the other time frames.
- **Reset:** This will reset any manually changed thresholds of the current time frame back to the estimated threshold setting of the selected method.
- **Recalc:** To redo the estimation process.

Where to find it: In the Advanced mode, the threshold menu is located in the left frame which contain all the initialization and setup menus. The Colocalization Wizard and Compact modes have a separate Thresholds stage and tab respectively which contains these threshold options. Note that in the Wizard, adjusting the histogram threshold lines is only possible in this specific stage.

Region of interest selection

By separating objects with threshold settings, you may also miss out on low but relevant voxel values inside your sample region. Region of interest (ROI) selection allows the measurements of colocalization inside the entire relevant regions of your image. Colocalization is then no longer calculated over the entire image, but only with the voxels inside this region. It is also possible to use a combination of both the drawn region of interest in combination with threshold values.

Enable the ROI option by ticking the *Enable* checkbox (see Figure 30.6). Then either select a pre-made ROI image from the Huygens Essential main window via the drop-down menu or draw one by hand.

When selecting a pre-made ROI image, the dimensions of the selected ROI image have to match the dimensions of the original source image. However, the ROI image may only contain 1 channel. Whether the ROI image meets the criteria is reported in Info. Hover

over the highlighted image parameters and read the tooltip in the statusbar to see why the ROI image is invalid. Once all the criteria are met, continue with the rest of the setup to calculate the colocalization in the area selected by the ROI image. The ROI image used for the analysis is also reported in the final results table.

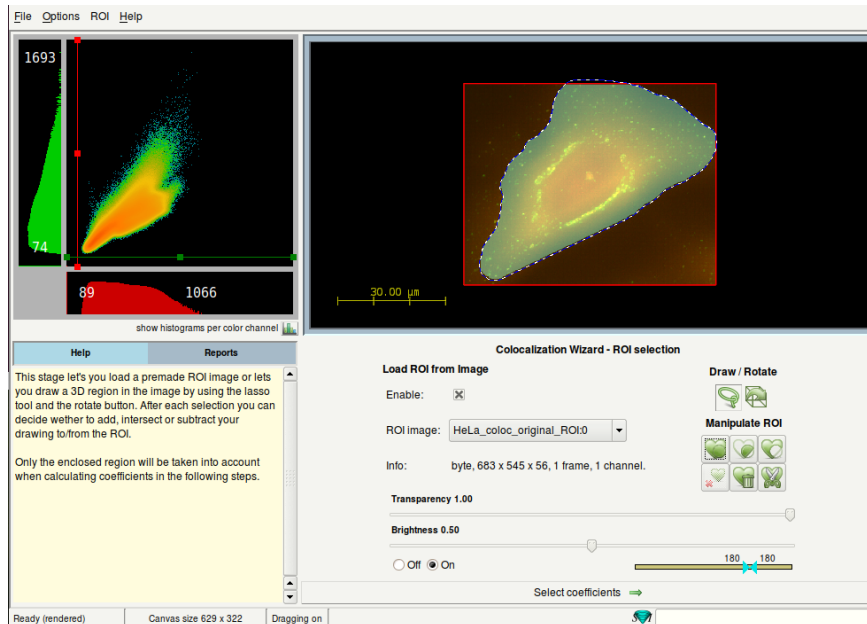


Figure 30.6: ROI image stage of the Colocalization Analyzer Wizard. The tool buttons on the bottom right hand side can be used to draw and customize the region in the viewer panel.

Different buttons are available when drawing an ROI by hand. They are located below the *lasso* button. These buttons allow addition, subtraction, inverting and complete removal of the drawn regions. Multiple selections can be drawn and added to the ROI image by repeating these steps. After adding a selection to the ROI image, the region is enabled and applied to the original image. This is visible by the slight color change that occurs when appending a region to the ROI image. Next, the 2D histogram is redrawn with only intensities of the remaining voxels inside this region.

Tweaking the visualization of the region is available with options such as transparency, brightness and color. These are located below the buttons and the image selection dropdown menu.

Use the menu bar to export the ROI to the main window or store it on disk (see ROI → STORAGE). In addition, both the rendering scene and the colocalization coefficients can be saved to external data (image or text file) using the FILE menu.

Where to find it: In the Advanced and Compact modes the ROI menu is located under the ROI Tab on the left side, and bottom control panel respectively. The Colocalization Wizard has a separate ROI stage, available from the Threshold stage.

RBNCC

The Colocalization Analyzer Advanced and Compact can be extended with the Replicate-Based Noise Corrected Correlation (RBNCC) method.⁶ Noise causes difficulties around colocalization analysis, since noise can also introduce colocalization. This can be handled by using a threshold, which is explained in previous subsections, but when the noise is quite severe a threshold may not be sufficient. With the RBNCC algorithm it is possible to correct for noise that contributes to the underestimation of the Pearson and Spearman

⁶J. Adler, S.N. Pagakis and I. Parmryd from Uppsala University; Journal of Microscopy 230, 121-133; 2008

coefficients. More information can be found at the SVI Wiki⁷. Colocalization coefficients are calculated and directly corrected with RBNCC when this option is activated. To use the RBNCC method a license with the RBNCC option is required.

When RBNCC is activated in the license, it can be enabled using the *Enable* tick-box in the RBNCC tab. With RBNCC enabled, a reference image can be selected by hovering over and clicking on the visible image name. The reference image should already be opened in the main window of Huygens Essential and should have the same dimensions as the original image. Whether the reference image meets the criteria is shown at *Info*, and hovering over any highlighted parameters shows a tooltip in the status bar explaining what is wrong. When the reference image is selected, simply press COMPUTE and the coefficients will be corrected automatically. The correction factors applied to the coefficient are shown in the RBNCC frame, where each abbreviation stands for a specific coefficient:

- S: Spearman coefficient
- P: Pearson coefficient
- OS: Object Spearman coefficient
- OP: Object Pearson coefficient

Where to find it: The Colocalization Advanced and Compact modes have a separate tab located on the top left and on the bottom respectively.

Computation and results

Press the green **Compute** button after the setup is complete to obtain the colocalization results. The histogram will automatically update to show the active time frame in the case that the coefficients are computed for multiple time frames. If the van Steensel's CCF coefficient is selected in the Advanced or Compact mode, the plot on the right hand side of the results table will show the curve of the cross correlation.

Results table

The results table shows the information of each time frame with the selected coefficients, selected channel indices, threshold values, threshold range, region of interest, and optionally the images used for RBNCC (See Figure 30.7).

Frame	Pearson	Spearman	Overlap	ICQ	ChR	ChG	ThR	ThG	SIR	SiG	ROI
0	0.679	0.875	0.790	0.399	0	1	0.00000	0.00000	0.00000	0.00000	test_image_1
1	0.690	0.881	0.791	0.410	0	1	0.00000	0.00000	0.00000	0.00000	test_image_1
0	0.574	0.812	0.646	0.387	0	1	15.000	20.000	0.00000	0.00000	-
1	0.581	0.820	0.649	0.398	0	1	15.000	20.000	0.00000	0.00000	-
0	0.706	0.853	0.786	0.424	0	1	0.00000	0.00000	0.00000	0.00000	-
1	0.714	0.866	0.788	0.432	0	1	0.00000	0.00000	0.00000	0.00000	-

Figure 30.7: The coefficients table with the time frame and channel combination, threshold settings, and name of selected ROI.

Save the table by navigating the top menubar and selecting FILE → SAVE RESULTS → SAVE COLOCALIZATION COEFFICIENTS. Or by pressing the SAVE COEFFICIENTS TABLE button in the final stage of the wizard.

Where to find it: The results table is located behind the *Coefficients* tab in the Compact mode, below the viewer panel on the right hand side in the Advanced mode, and in the final stage of the Wizard.

⁷<https://www.svi.nl/RBNCC>

Colocalization map

In addition to the coefficients that show up in the table, surfaces will also appear on top of the 3D MIP rendered image in the top right viewer (See Figure 30.8). These are isosurfaces that represent voxels which all exceed the *isosurface threshold*. Regions where these voxel values are the same, will connect and become objects. In other words, for regions in the colocalization map to become visible as objects, they have to at least contribute that specific amount to the colocalization coefficient. By default this threshold is obtained with *Otsu's method*⁸. However, the isosurface threshold value can be altered in the render settings, which are available in the Advanced and Compact mode (See Section Render and Object statistics settings below). More information about the seed-and-threshold options can be found in Section Segmenting the Objects: Setting the Threshold on page 161.




Figure 30.8: The Colocalization Analyzer Compact with a surface rendering of the colocalization map. Settings are shown in the *Surface/MIP Renderer* tab selected in the Control panel.

If Li's ICQ coefficient has been selected as the colocalization map, two pop-up windows are generated, showing the related intensity correlation analysis plots.

Render and Object statistics settings

There is one *Maximum Intensity Projection* (MIP) pipe available to overlay a MIP rendering of one or both channels. This acts as a good spatial reference for the objects from the colocalization map. These MIP render options and the colocalization map's isosurface settings are available from the the render options in the Colocalization Analyzer Advanced and Compact modes (See Figure 30.8).

Change the viewpoint of the viewer panel by moving the Zoom, Tilt, and Twist sliders. The MIP render quality can be set to either Faster rendering or High quality. The channel selector allows the MIP pipe to display one channel at a time. When a single channel is selected, the first slider can be used to select which voxels are considered for the MIP rendering, depending on their intensities. Note that this viewer threshold is linked to the Red/Green Channel colocalization threshold settings as discussed earlier (See Section Threshold settings on page 190). If both channels R and G are selected to be rendered, their corresponding thresholds as selected in the histogram will be used as projection thresholds. As with the surface pipes, the transparency and the brightness of the MIP can also be controlled.

As soon as the colocalization map has been calculated, a surface is overlaid on top of

⁸https://en.wikipedia.org/wiki/Otsu's_method

the image's MIP rendering. The isosurface threshold settings become available after being grayed out before. Setting the isosurface threshold and seed changes the criteria for the objects to become visible for analysis. Setting a garbage volume for the isosurfaces allows filtering out small objects. The transparency and the brightness of the isosurface can be controlled with the corresponding sliders in the visualization panel. The color range in which the objects are displayed can be modified using the hue selector.

Some coefficients 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, or maps for the minimum and maximum coefficients in case of van Steensel's CCF. In these cases, the iso-colocalization surface parameters offer the possibility of rendering any of the two channels, and thus the isosurface threshold refers only to the active one. The channel to be used for the rendering is selected using the radio-buttons. The surfaces can also simply be switched off.

Where to find it: The render settings can be found in the *Surface/MIP renderer* tab in the Compact mode and on the right-hand side of the viewer panel in the Advanced mode. The Colocalization Analyzer Advanced mode has separate tabs for the MIP and isosurfaces (map) settings.

Object statistics

After the render settings have been applied, clicking the *Object Statistics* tab allows for the analysis of a wide range of statistics for each of the objects as segmented in the viewer panel. Using the *magnifying glass* icon, local object colocalization parameters can be computed and reported by clicking on the individual objects in the surface renderer. This allows for quick determination of the properties of different colocalization regions in the data. Alternatively, it is possible to obtain the statistics of all objects by clicking the *Analyze all* button (See Figure 30.9). These object statistics can be exported through the main menu by selecting FILE → SAVE RESULTS → SAVE OBJECT ANALYSIS. In the wizard the SAVE OBJECT ANALYSIS button in the final stage becomes available as soon as the analysis is performed on the objects.

Object colocalization statistics

Selected colocalization map coefficient: Pearson.
 Conditions: Threshold 2.67e-06, Seed level 2.67e-06 (0%), Garbage level: 2. Total objects: 363
 Frame 0, Red chan. 0, bg: 0.0000, Green chan. 1, bg: 0.0000

Units: surface: μm^2 , iso-volume: μm^3 , contrib: %

Object	Voxels	IsoVol	Surface	Sum	Contrib	Min	Max	Length	WIAx	WILat	C.Mass
1	8	2.86e-3	0.226	2.90e-5	4.27e-3	3.00e-6	4.00e-6	0.33	0.21	0.24	(2.6, 10.3, 7.0)
2	1265	1.029	20.323	5.60e-3	0.825	3.00e-6	1.40e-5	4.84	1.33	1.09	(36.6, 21.5, 7.0)
3	2	3.80e-5	7.92e-2	6.00e-6	8.84e-4	3.00e-6	3.00e-6	0.40	6.60e-2	9.02e-2	(2.0, 15.0, 7.5)
4	2	2.80e-5	6.33e-2	6.00e-6	8.84e-4	3.00e-6	3.00e-6	0.20	0.13	0.12	(16.5, 21.0, 6.0)
5	87	5.89e-2	1.810	3.41e-4	5.02e-2	3.00e-6	7.00e-6	1.05	0.48	0.60	(6.8, 30.9, 5.8)
6	12	4.60e-3	0.385	3.90e-5	5.75e-3	3.00e-6	4.00e-6	0.59	0.30	0.30	(13.5, 26.8, 8.9)
7	1017	0.802	18.151	4.20e-3	0.619	3.00e-6	1.20e-5	4.36	1.18	1.20	(98.1, 38.7, 6.6)
8	623	0.517	8.314	2.99e-3	0.441	3.00e-6	1.40e-5	2.23	1.38	1.40	(27.8, 30.3, 11.0)

Figure 30.9: The Object statistics window after computing the information of all visible objects. The information about the size, contribution to the coefficient, and center of mass is displayed.

Chapter 31

Huygens multi-GPU acceleration

Huygens can make full and optimal use of all the computer's central processing units (CPUs). Huygens version 15.05 introduced the possibility to use the power of graphics processing unit (GPU) acceleration. The GPU contains hundreds or thousands of separate processing units (cores), making it extremely efficient and fast to process your data, provided that the required operations can be done in parallel. Since version 16.10.0p8 multi-GPU support is available and the *Batch Processor* and *Huygens Core* can divided the queued deconvolution tasks over the available GPU's.

Huygens GPU options

GPU size options

Huygens GPU acceleration is available for CUDA-enabled NVIDIA®¹ graphics cards that are covered by the Small GPU option free of charge. Cards with higher specifications may also be included by adding a GPU option to your license.

In general, we distinguish the following GPU options:

- Small GPU option (free of charge)
- Medium GPU option
- Large GPU option
- Extreme GPU option

Up to date details on which NVIDIA® CUDA-enabled graphics cards are supported by which GPU option can be found on our website².

Many CUDA graphics cards have already been tested regarding their performance for Huygens pby SVI and many of SVI's customers. The results of these tests are published on our wiki pages³ ⁴. The performance of Huygens can be up to 30 times faster by using a large GPU as opposed to using only one CPU. However, note that some operations are not yet implemented to run on the GPU, or may not be suitable to be run in parallel. Therefore, the CPU performance should not be ignored when decisions about what hardware to run Huygens on are made.

¹NVIDIA is a registered trademark of NVIDIA Corporation (USA)

²<https://www.svi.nl/GPUcards>

³https://www.svi.nl/GPU_benchmarks

⁴<https://www.svi.nl/HuygensGPU>

Huygens Performance options

Besides the GPU options for Huygens, SVI also offers Performance options. For the use of multi-GPU acceleration the Performance Plus, Mega or Extreme option is needed.

In general, we distinguish the following Performance options:

- Performance Option: Allows the use of up to 6 physical CPU cores (12 logical with hyper-threading) and 1 Small GPU card. This is included in Huygens free of charge.
- Performance Plus: Allows the use of up to 12 physical CPU cores (24 logical with hyper-threading) and 2 Large GPU cards.
- Performance Mega: Allows the use of up to 24 physical CPU cores (48 logical with hyper-threading) and 4 Large GPU cards.
- Performance Extreme: Allows the use of up to 48 physical CPU cores (96 logical with hyper-threading) and 8 Large GPU cards.

Huygens and GPU acceleration

Starting with the CMLE, QMLE and GMLE deconvolution algorithms, an increasing number of operations has gained GPU acceleration support over time. See the SVI web site⁵ for an overview of which parts have been accelerated for single and multi-GPU use.

How to prepare GPU drivers and CUDA® for Huygens

The GPU support installation on Windows

- Install the NVIDIA® drivers for the graphics card⁶.
- Download the latest CUDA version Network installer for Windows⁷.
- Install the latest CUDA version using the Network installer.
- Install Huygens version 19.10⁸.
- Install your Huygens license. If you do not have a license, you can request a test license by sending an email to `license@svi.nl`.
- GPU acceleration is active. You can deactivate it in the Huygens *Preferences* window (under EDIT → PREFERENCES) or via the GPU task bar icon. If the GPU option is grayed-out, please check out the “GPU Check List” on page 199.

The GPU support installation on Linux

- Download the latest CUDA version Network Installer for your Linux operating system⁹.
- Follow the instructions as specified.
- Install your Huygens license. If you do not have a license, you can request a test license by sending an email to `license@svi.nl`.
- GPU acceleration is active. You can deactivate it in the Huygens Preferences window (under EDIT → PREFERENCES) or via the GPU task bar icon. If the GPU option is grayed-out please read the following section.

⁵<https://www.svi.nl/GPUcards>

⁶<https://www.nvidia.com/Download/index.aspx>

⁷<https://developer.nvidia.com/cuda-downloads#win>

⁸<https://www.svi.nl/Download>

⁹<https://developer.nvidia.com/cuda-downloads#linux>

GPU Check List

For the latest details on the GPU checklist please visit this wiki page¹⁰.

1. **Do you have a NVIDIA® CUDA-enabled GPU?** To use Huygens GPU acceleration you need a CUDA-enabled NVIDIA® graphics card with a minimum *compute capability* of 3.0¹¹
2. **Have you installed the NVIDIA® driver and CUDA library?** You can install the NVIDIA® driver and CUDA as described under “The GPU support installation on Windows” or “The GPU support installation on Linux” on page 198
3. **Have you installed Huygens with GPU acceleration?** GPU acceleration is available from Huygens version 15.10 and higher.
4. **Do you have a valid license for the GPU acceleration?** Usage of small GPUs by Huygens is free if a GPU is listed as a Small GPU card¹². For cards that fall within the other classes of cards, an additional Huygens GPU option needs to be purchased. Please contact license@svi.nl for a test license or contact us if you have purchased a GPU option.
5. **Have you activated the GPU acceleration?** Please see “GPU status Window” on page 199

GPU status Window

The GPU acceleration in Huygens can be managed from the GPU status window. The GPU status window can be opened by the GPU STATUS button in the main Huygens button bar, or from the “GPU acceleration” section under EDIT → PREFERENCES.

The GPU status window has three sections (See Figure 31.1).

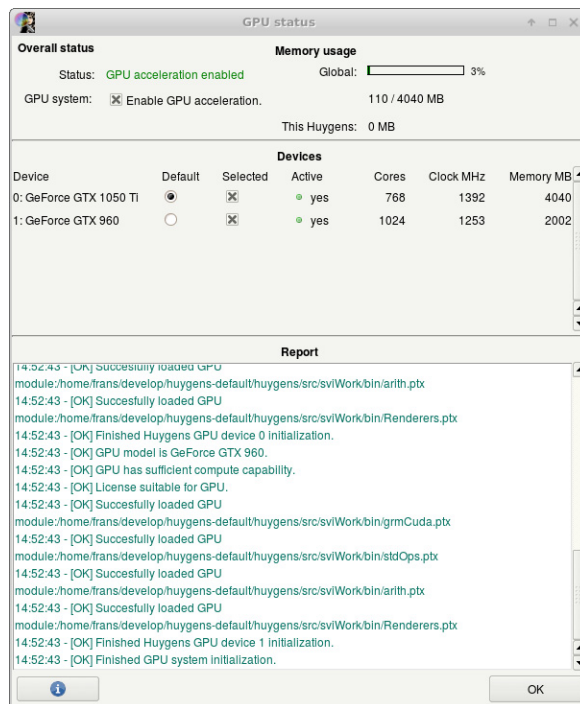


Figure 31.1: The GPU status window shows that the installed GeForce® GTX 1050 Ti and GTX 960 cards are enabled in Huygens.

¹⁰<https://www.svi.nl/GPUchecklist>

¹¹<https://en.wikipedia.org/wiki/CUDA>

¹²<https://www.svi.nl/GPUcards>

In the top section, the GPU acceleration status is indicated. The GPU system check box can be used to enable or disable the GPU acceleration as a whole. The default device is the GPU card that will be used by all GPU operations that are not using multiple GPUs. The rest of the first section displays some information about the currently selected default GPU device, such as its number of CUDA cores, clock rate, total memory, memory currently in use as reported by CUDA, and memory in use by the current Huygens instance.

The second section of the GPU status window lists all CUDA-capable devices that meet the hardware and license requirements. The left column lists the Huygens GPU card number and the card's name as reported by CUDA. The second column ("Default") contains radio buttons to select a new default GPU device for the currently running Huygens. The third column ("Selected") contains check-boxes to select which GPU devices should participate in multi-GPU operations, such as the Batch Processor and multi-GPU deconvolution commands. The fourth column "Active" contains indicators to show which devices would be participating in multi-GPU operations for clarity: they will be turned off either by unchecking the corresponding device-specific box in the "Selected" column or by disabling the GPU system as a whole. The three remaining columns "Cores", "Clock MHz" and "Memory MB" list the number of CUDA cores, nominal clock frequency in MHz and amount of video RAM in megabytes as reported by the NVIDIA® graphics driver.

The third section of the GPU status window is the GPU Report. This is where initialization and status changes are listed, as well as GPU-related error messages. When there are GPU-related problems, please check the contents of this part of the window for error messages. Note that it may happen that the Report contains information about other concurrently running Huygens instances apart from the current one. When the set of selected devices is modified, please note that:

- only enabled devices can be selected to become the new default device;
- when all devices are disabled, the GPU system as a whole is disabled as well;
- the GPU system as a whole can only be enabled if at least one device is enabled;
- if a device is the first device to become enabled, that device automatically becomes the new default device;
- the number of devices that can be enabled simultaneously can be limited by the license.

Multi-GPU support

With Huygens version 16.10.0p8, multi-GPU support has been introduced in the *Batch Processor* and *Huygens Core* to perform deconvolution of a queue of images on multiple GPU devices simultaneously. For the most recent information on GPU developments in Huygens we advise you to visit our website¹³.

A prerequisite for using multiple GPUs is that the license should allow this (See "Huygens Performance options" on page 198). If this is the case, the MULTI-GPU button in the button bar of the *Batch Processor* can switch between using one or the selected set of available GPU devices. If one GPU is used, then the GPU device that is selected in the GPU Status window (see HELP → GPU STATUS). If the multi-GPU mode is selected, then the button will display the number of used GPU devices between parentheses. In multi-GPU mode, by default the number of "concurrent jobs" in the batch processor is overridden by the number of (selected) GPU devices, in order to have one concurrent job per GPU device.

In case there is a problem with enabling the multi-GPU functionality, the tool-tip help that is displayed at the bottom right of the Batch Processor window when the mouse cursor is on the multi-GPU button, can give a short description of the status or problem. Possible states of the multi-GPU button and tool tip include:

¹³<https://www.svi.nl/HuygensGPU>

- It is grayed out, because the whole GPU subsystem is not available (see the GPU troubleshooting page on the SVI website¹⁴)
- it is grayed out, because the whole GPU subsystem is not enabled (enable it in the Preferences or GPU status window);
- it is grayed out, because there are fewer than 2 GPU devices (there is no multi-GPU if you have zero or one GPUs);
- it is grayed out, because the license does not allow the use of more than one GPU device (contact SVI for a trial license);
- it is set to single-GPU and the user may click on the button to set it to multi-GPU;
- it is set to multi-GPU and the user may click on the button to set it to single-GPU.

In case the license either does not allow to use the multi-GPU feature at all, or the license does allow it but for fewer than the available number of devices, a warning is written to the task report.

¹⁴<https://www.svi.nl/gpuCheckList>

Chapter 32

Huygens for Multiple Users

If the demand for Huygens use within one department or spread over campus increases, it becomes a challenge to maintain easy accessibility to Huygens for all users. This chapter will focus on the solutions that we offer within Huygens to facilitate the work of many different users simultaneously.

In situations where the number of processing jobs becomes a challenge, we kindly refer you to more information on our schedulers: *the Batch Processor* (Chapter 7 on page 39) and to the “*Huygens web accessibility with HRM*” on page 205.

We offer three different solution to facilitate access for *Multiple Users*. All are possible with a single or multiple license(s):

1. Remote display
2. Floating licenses
3. Web-based solution via Huygens Remote manager

Remote Display

Huygens offers *Remote Display*, which allows Huygens to be displayed on several *client* computers at the same time, from only one so-called server computer (this server can be a dedicated server, a simple workstation, or even just a laptop).

Features

- Current Huygens Suites already include the remote Display capability
- If the *server* is powerful Huygens will perform accordingly
- Hardware and central storage can be shared efficiently
- Adding more *clients* is easy and flexible
- Users do not have to go elsewhere, but can start jobs in Huygens from their own workstation
- Both *server* and *clients* can run on Windows, Linux or Mac OS. All combinations are possible with multiple clients
- Both *server* and *clients* need to be part of the same network.

How to set up Remote Display

The *server* requires the Remote Desktop Protocol (RDP). All three operating systems (Windows, Linux, and Mac) have their RDP. In general, Windows restricts access to a single user unless a server edition of Windows is used. Even then, the number of simultaneous users may be limited. Linux and Mac allow, in theory, an unlimited number of *client* users simultaneously.

Clients from a Linux or Mac server can access the server via ssh or the X11 command typed in a shell. *Clients* from a Windows server need different protocols per *client* OS. For more information and a full list of currently available free-ware and commercial solutions in this RDP field we refer to our wiki¹.

Floating License(s)

The floating license offers concurrent usage of Huygens on multiple independent computers that do not need to be part of a single network. So, different from the remote display option, no central server or computer network is needed. Only one Huygens license can be set up as a floating license and will already allow the use of Huygens on, for example, the microscope computer for first checks to optimize acquisition parameters, and on a dedicated workstation for extensive image processing jobs.

Unlike most floating license systems, Huygens is much more flexible because it allows even usage beyond the maximum number of concurrent users.

Main Features

- Concurrent usage above the agreed upon number of concurrent licenses is allowed up to 3,000 minutes per year.
- The range of IP addresses on which the Huygens floating licenses can be used is determined by the contact-person, and can be adjusted at any time in consultation with SVI.
- Three different servers distributed over the world ensure a continuity in IP checks so that Huygens functionality is preserved.
- Huygens will continue to work for 72 hours if systems are taken outside the institutes network and effectively go off-line (*The buffer*). This is ideal for users who want to finish work at home or are at traveling to a conference.
- *The buffer* is an extra protection if all three servers are not responsive.
- Users are free to run Huygens on either Windows, Linux or Mac.
- The contact-person receives monthly or weekly overviews of actual usage, so that there is ample time to implement adaptations in IP range and number of floating licenses, if needed.

How to set up Floating licenses

- The floating license functionality is an option that can be added to a Huygens license.
- The contact-person sends a list of IP addresses to SVI.
- The contact-person receives an active single floating license, which can be sent to the users with computers within the IP range.

¹<https://www.svi.nl/RemoteDisplay>

- The latest Huygens version available on our download page² and floating license can be installed on every computer with the agreed IP range.
- If a new version of Huygens is available and the Maintenance and Upgrade contract is still valid, a new single floating license will be issued.
- To ensure that license and software are of the same version, the contact person can best send both the new license as well as a link to the corresponding latest Huygens version to the users within a single email.

Floating licenses are an option and require an uplift per floating license configuration. For more information see our website on floating licenses³.

Huygens web accessibility with HRM

The *Huygens Remote Manager (HRM)* empowered by the *Huygens Core* offers Huygens users a full web-based solution with the huge advantage that users can work from any physical location around the Globe, as long as they have Internet access. The number of simultaneous users can run effectively into hundreds.

HRM is an open-source project in which innovative ICT ideas in microscopy can be implemented. It works as an interface for the *Huygens Core*, which is one of the four Huygens basics optimized for multi-user batch-processing.

Because the interaction with *Huygens Core* occurs via its web interface HRM and not via the Huygens interface directly, as is the case with Essential and Professional, the HRM-HuCore is set up differently. In the HRM-HuCore the user selects images and is guided along a pre-defined path in which the microscopy settings are automatically presented. Obviously, the user can adjust these values if needed. The parameters will be used to automatically generate a Point-Spread Function (PSF) during the deconvolution process. Microscopy parameters can be saved in an image template for future use. The next step is setting the deconvolution and (colocalization) analysis parameters, and submitting the jobs to the server. A link with the resulting files is sent to the user who can view the results interactively and download the files.

Main Features

- HRM/HuCore is based on scalable architecture: it can be installed on one or more servers.
- A single installation on the (Linux) server suffices.
- Contact person will have access to user statistics via the HRM web-interface at all times.
- Contact person can easily manage users and groups via HRM, or externally through LDAP/Active Directory.
- Thanks to the built-in HRM queue manager optimal use of the hardware resources is ensured.
- Users can upload images via HRM, and receive an email notification with a link to the download-able results.
- A file interface in HRM allows users to collect datasets to be deconvolved from a file server or OMERO database server and push back the deconvolution results.

²<https://www.svi.nl/Download>

³<https://www.svi.nl/HuygensFloatingLicenses>

Set-up:

For the manual:

<https://www.svi.nl/Manuals>

For the download:

<https://sourceforge.net/projects/hrm/files/latest/download>

For an overview:

<https://www.svi.nl/HuygensRemoteManager>

For more detailed information you can read the Huygens Core programmer Guide and the Huygens Remote Manager⁴.

⁴<https://www.svi.nl/Manuals>

Chapter 33

Establishing Image Parameters

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 Essential 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 Essential 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 Essential. 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 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 should not be considered as a parameter describing the original image, but more as a tunable parameter that controls the deconvolved result. Since Huygens version 17.10, the *Deconvolution Wizard* is equipped with an option to quickly access the optimal SNR value using deconvolution on a region of interest.

¹<https://www.svi.nl/BrickSplitting>

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 33.1 shows some examples of recordings where different noise levels were added to an original (restored) image.

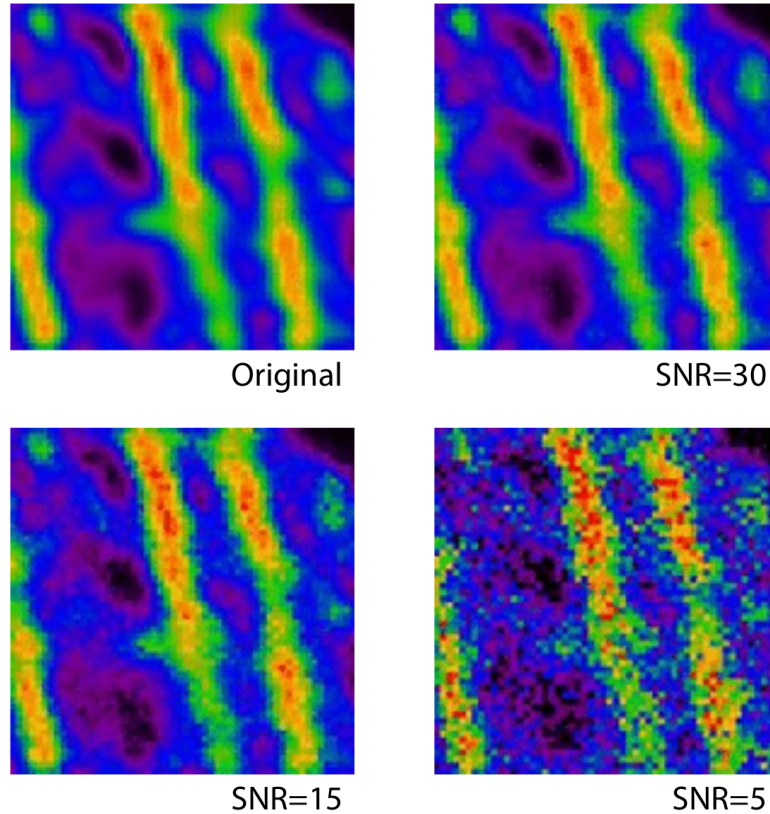


Figure 33.1: Examples of different SNR values. Same image, different noise levels.

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:

$$SNR = \sqrt{\frac{M}{s}} \quad (33.1)$$

If a significant blacklevel is present, it should be subtracted from M and s . See “Black Level” on page 209 and the SVI 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 33.1)

It is recommended to perform the deconvolution with different SNR values around the estimated value and to optimize this SNR value further in order to get the best possible outcome.

²<https://www.svi.nl/SNR>

Black Level

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




Figure 33.2: Left: The histogram of an image without black level. Right: The histogram of a similar image with a significant black level.

A positive black level (See Figure 33.2) will do no harm to the deconvolution since it is 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 Essential 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 214.) 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 Wiki has an on-line tool⁴ that computes the Nyquist rate for any widefield, confocal, spinning disc, SPIM/Light Sheet, STED, 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<https://www.svi.nl/BlackLevel>

⁴<https://www.svi.nl/NyquistCalculator>




Figure 33.3: 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.

Figure 33.3 shows the dependency of this critical sampling distance on the numerical aperture for a wavelength of 500 nm. To apply this plot to another wavelength, simply scale the vertical axis by that wavelength. For example, if a widefield microscope with NideA 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 $(570/50) \times 95 = 108nm$.

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 where 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 Wiki⁵.

Computing the Backprojected Pinhole Radius and Distance

Throughout Huygens Essential pinhole sizes of confocal systems are specified as the *backprojected radius* in nanometer (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} :

$$r_b = \frac{r_{phy}}{m_{obj}m_{sys}} \quad (33.2)$$

⁵<https://www.svi.nl/NyquistRate>

The SVI 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 212.).

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:


$$r_b = \frac{0.61\lambda_{ex}N_{Ad}}{NA} \tag{33.3}$$

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 33.4), the shape factor c for a square pinhole becomes $2/(\sqrt{\pi}) = 0.654$. 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 33.3 and the shape factor c for a square pinhole. Taking into account that Equation 33.3 already converts a diameter into a radius, the relation becomes:

Figure 33.4: A square and a disk with an equal area A .

$$r_b = \frac{0.69\lambda_{ex}N_{Ad}}{NA} \tag{33.4}$$

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 33.5 can be used to translate that parameter into a backprojected pinhole radius.

$$r_b = 10^3 c \frac{(p_8/255)(s_{max} - s_{min}) + s_{min}}{m_{obj}m_{sys}} \tag{33.5}$$

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.

⁶<https://www.svi.nl/BackprojectedPinholeCalculator>

⁷<https://www.svi.nl/AiryDisk>

The backprojected Pinhole Spacing for Spinning Disk

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 200 \times objective lens and a reported pinhole size d of 150 μm the backprojected pinhole radius r_b in nm is:

$$r_b = c \frac{d}{m_{obj} m_{sys}} = 0.564 \frac{150 \times 10^3}{100 \times 3.82} = 221 \quad (33.6)$$

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, $53 \times m_{tube}$, 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 8x 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 \times 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 33.7 to express a distance r in dimensionless o.u.:

$$v = r \frac{2\pi}{\lambda} NA \quad (33.7)$$

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 \times , the system magnification $4918/60 = 82 \pm 9\%$. This value corresponds well with the largest possible system magnification for the MRC600 ($53 \times 1.25 \times 1.25 = 83$).

⁸<https://www.svi.nl/YokogawaDisk>

⁹<https://www.svi.nl/BackprojectedPinholeCalculator>

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

¹¹<https://www.svi.nl/OpticalUnits>

Chapter 34

Improving Image Quality

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 may help in 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 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 fish tank 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 fish tank effect by a different step size. Therefore objects will appear elongated or shortened in the image data.

Huygens Essential 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.

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.

¹<https://www.svi.nl/>

²<https://www.svi.nl/FishtankEffect>

³<https://www.svi.nl/SphericalAbberation>

- **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 213.).

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 Wiki⁶.

Undersampling

One of the rules of measurement that is often overlooked is sampling according to the *Nyquist rate* (See “Sampling Density” on page 209). 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 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. Although this results in an apparently noisier image, there is actually *more information* in the data because sampling density is higher. Deconvolution with Huygens Essential 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.

⁴<https://www.svi.nl/TotalInternalReflection>

⁵<https://www.svi.nl/DynamicRange>

⁶<https://www.svi.nl/ClippedImages>

⁷<https://www.svi.nl/NyquistCalculator>

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 of single z stacks from confocal or two photon systems. But in widefield deconvolution bleaching is more of a problem. Fortunately, the bleaching in widefield images can usually be corrected quite easily. For time series, images from different modalities can be distorted by bleaching. If the bleaching is very severe the correction might not be perfect, resulting in lower quality deconvolution results. For more information on how to correct bleaching see Chapter 23 “Bleaching Corrector” on page 133.

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 Essential 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 , it 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.

Deconvolution Improvements

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 Chapter 6 “The PSF Distiller” on page 33).

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 Wiki¹⁰ for detailed steps in configuring the restoration process.

⁸<https://www.svi.nl/PointSpreadFunction>

⁹<https://www.svi.nl/SphericalAberration>

¹⁰<https://www.svi.nl/DeconvolutionProcedure>

Chapter 35

Appendix

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 Essential, a PSF can be obtained in two different ways:

1. *Generating a theoretical PSF*: When a measured PSF is not available, Huygens Essential 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 Essential does the calculation in the background without any notice.
Images affected by spherical aberration (See “Refractive Index Mismatch” on page 213) 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 33.). Measured PSF’s improve deconvolution results and may also serve as a quality test for the microscope

Quality Factor

Deconvolution as it is done in Huygens Essential 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 Essential 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

¹<https://www.svi.nl/ImageFormation>

arbitrary prefixes and ordering.

Numbered Tiff Series

If a series is simply numbered like: `slice001.tif`, `slice002.tif`, ..., `slice0nn.tif`, then Huygens Essential 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 Essential 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 filename:

```
im_t02_z003_c01.tif
```

The File Series Tool

Although Huygens Essential 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 Essential shows the *File Series Tool* dialog (Figure 35.1). This tool enables the user to select a subset of a file series, and select a dimension for each of the indices in the file name, so that each image is assigned to the correct *z*-plane, time frame, and channel. Only select the first file of a series. If you select more, Huygens will attempt to see every selected file as a start of a new series. Consequently, many windows of the File Series Tool will be opened.

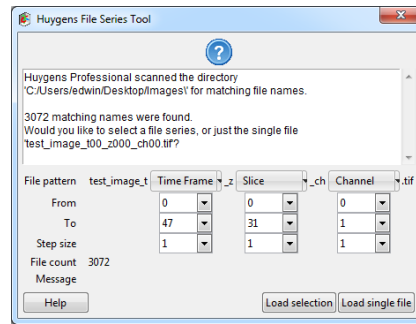


Figure 35.1: The Huygens File Series Tool automatically scans a directory for a file series.

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.
- **Tiles:** The range of this counter becomes the tile dimension and will open the Stitcher. See Chapter 20 “Tile Stitching & Deconvolution Wizard” on page 121.
- **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 Essential uses a global scheme for coloring the different channels in multichannel images. These colors can be adjusted through the Preferences window via the **EDIT GLOBAL COLORS** button under **EDIT** → **PREFERENCES** (See Figure 35.2).

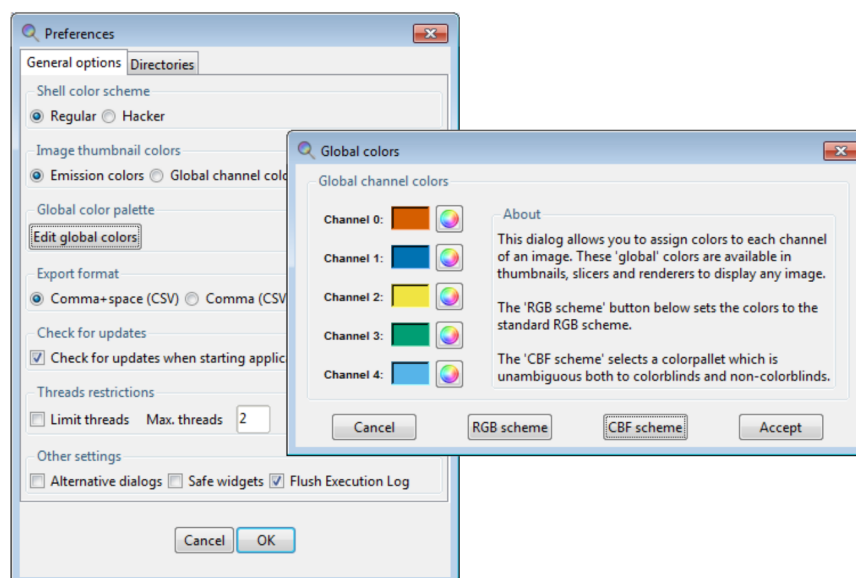



Figure 35.2: The global color scheme can be modified through the *Preferences* window.

There are two color scheme available which are the RGB and CBF schemes. The RGB scheme (Red Green Blue scheme) is the color scheme that starts with the red, green and blue colors for the first 3 channels. The CBF scheme (Color Blind Friendly scheme) are contrasting colors that are unambiguous both to color blind and non-color blind.²

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 107 and “The Colocalization Analyzer” on page 187). 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.

²Published by Okabe and Ito, “How to make figures and presentations that are friendly to Colorblind people”, J*FLY, 2002

Setting the Coverslip Position

If 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 213). As deeper layers in the specimen are imaged, moving away from the coverslip, this distortion will progressively worsen. To compute the level of spherical aberration at a specific depth within the specimen, 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. Next to direct numerical input, the coverslip position and imaging direction can be set using a visual editor that can be started using the LAUNCH EDITOR button located in the parameter editor window (Figure 35.3).

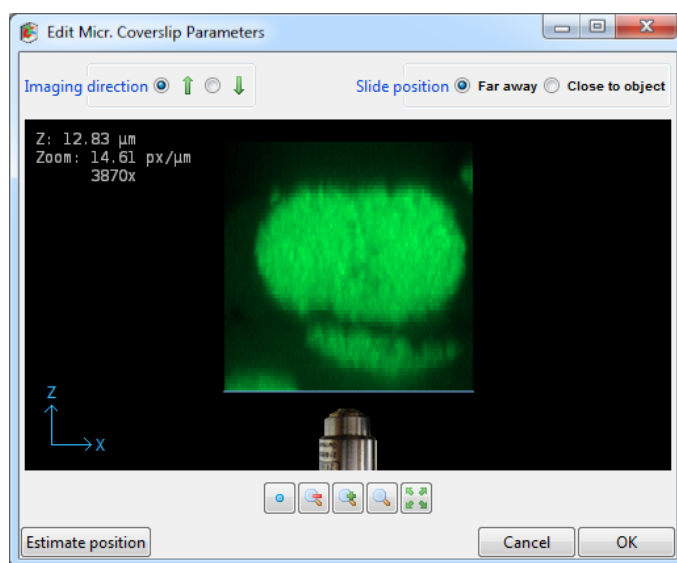


Figure 35.3: The coverslip position editor showing an xy MIP of the data along the y -direction. The coverslip position can be adjusted by dragging the blue line. The imaging direction, here *upwards*, is indicated by the position of the objective relative to the data as shown. The z -position shown top-left in the image indicates the distance in μm of the coverslip to the first data plane.

Inverted Microscope

In addition to the coverslip position, the visual editor shows the 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. 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.

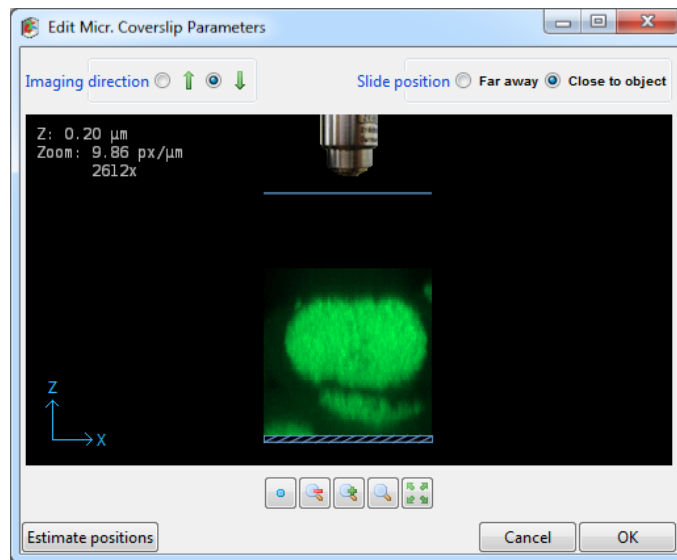


Figure 35.4: The Coverslip editor with the slide position set to *Close to object*.

When the specimen is close to or mounted on the slide, select *Close to object* (upper right corner). 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.

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 profile³. 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 35.5) and is typically in the range from 2 to 4. The overfill factor can be set as a microscopy parameter in Huygens Essential, and is taken into account when computing the point spread function.

³https://en.wikipedia.org/wiki/Gaussian_beam




Figure 35.5: 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.

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 → CONVERT BRIGHTFIELD 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.

2D Histogram

In image processing a 2D histogram shows the relationship of intensities between two images. The 2D histogram is mostly used to compare two channels, where the x-axis represents the intensities of the first channel and the y-axis the intensities of the second channel.




Figure 35.6: Example how a 2D histogram is computed, showing the first and second channel and their corresponding 2D histogram. At each position within the channels, the corresponding intensities of both channels are combined to form a coordinate within the 2D histogram. The count at this coordinate of the 2D histogram is then increased by one.

As a comparison, a 1D histogram is nothing more than counting how many voxels with a particular intensity occur in the image. The intensity range of the image is divided in bins. A voxel then belongs to the bin if its intensity is included within the range the bin represents.

The 2D histogram is the same as the 1D histogram with the difference that it counts the occurrence of *combinations* of intensities. To compute a 2D histogram the images need to be equal in size. See the example in Figure 35.6 ; at position (30,20) the first channel has an intensity of 200 and the second image has an intensity of 10. Then this will add one to the count in the 2D histogram at position (200,10).

The difficulty with a 2D histogram is how to show the actual count per intensity combination. For a 1D histogram the height of the bars represent the count, but this height for a 2D histogram requires a third dimension which is difficult to visualize. Instead, to visualize the count of the combinations, colors are used. In the histogram in Figure 35.6 colors range from green to red. Note that these colors have nothing to do with the colors used to represent the channels. Within the 2D histogram, green represents low counts while red represent high counts.

2D histograms show interesting image properties and are therefore very useful:

- Offset
- Intensity factor
- Colocalization
- Clipping
- Crosstalk
- Hot pixels

More information and examples are in the SVI wiki⁴.

⁴<https://www.svi.nl/TwoChannelHistogram>.

Chapter 36

Support and Contact Information

Contact Information

Addresses and Phone Numbers

Mailing Address	Scientific Volume Imaging B.V. Laapersveld 63 (Entrance C) 1213 VB Hilversum The Netherlands
Phone	+31 35 6421626
Fax	+31 35 6837971
E-mail	info@svi.nl
URL	https://www.svi.nl/

We are directly reachable by phone during office hours (CET) or by e-mail 24/7.

Distributors

An up-to-date list of distributors can be found on our web site¹.

Support

SVI Support Wiki

The SVI Wiki² is a public knowledge resource on 3D microscopy, image restoration (deconvolution), visualization and analysis. 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 Wiki to learn more about the Huygens software and microscopic imaging in general:

- Information on the parameters describing the imaging conditions (sampling, numerical aperture, pinholes, etc.):
<https://www.svi.nl/MicroscopicParameters>

¹<https://www.svi.nl/distributors/>

²<https://www.svi.nl/HomePage>

- Information on the restoration parameters (signal to noise ratio, background, quality criteria, etc.) used by the deconvolution algorithms:
<https://www.svi.nl/RestorationParameters>
- A step by step example on how to tune these parameters to achieve the desired restoration results:
<https://www.svi.nl/DeconvolutionProcedure>
- Important issues regarding image acquisition and restoration (sampling, clipping, etc.):
<https://www.svi.nl/ImportantFactors>
- Typical acquisition pitfalls (spherical aberration, undersampling, bleaching, etc.):
<https://www.svi.nl/AcquisitionPitfalls>
- Information on recording beads to measure a PSF:
<https://www.svi.nl/RecordingBeads>
- Tutorials and detailed information on using the different aspects of the Huygens software (restoration, visualization, analysis, programming, etc.):
<https://www.svi.nl/Tutorials>
- Uploading images to SVI:
<https://www.svi.nl/upload>

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**

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:

```
HuEss-19.10-wcnps-d-tvAC-emnps-eom2019Dec31-e7b7c623393d708e-  
user@domain.com-4fce0dbe86e8ca4344dd
```

Table 36.1 lists the building blocks from which this string is composed.

Table 36.1: The building blocks of the Huygens license string.

Substring	Descriptions
Product	The product to which the license string applies. This can be HuEss, HuPro, HuScript, HuCore, HuLoc, and HuTitan.
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' (multiphoton), 's' (STED), 'S' (STED 3X), 'P' (SPIM/Light sheet), '4' (4Pi experimental microscopes), 'r' (Rescan Confocal) and 'a' (Array Detector Confocal).
Server flag	Determines the number of cores that are enabled for multi-threading, and the use of one or multiple GPU cards. A hyper-threaded core is counted as a single core. It can be 'd' (desktop; maximum of 6 CPU cores, and 1 GPU card), 's' (small server; maximum of 12 CPU cores and 2 GPU cards), 'm' (medium server; maximum of 24 CPU cores and 4 GPU cards), 'l' (larger server; maximum of 48 CPU cores and 8 GPU cards)
Option flags	This is a set of characters that list the enabled optional modules. An overview of these modules is given in Table 36.2.
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).
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 36.2: Identifiers for the optional modules.

Character	Module
a	Object Analyzer Basic
A	Advanced Object Analyzer
b	Small file reader bundle
B	Complete file reader bundle
c	RBNCC option for colocalization
C	Colocalization Analyzer
e	Deconvolution Express / Batch Express
f	PSF Distiller
F	Fusion Wizard
g	Determines what GPU card can be used ⁴ . By default a small GPU card can be used. 'g1' (Medium GPU option) allows the use of a medium GPU card, 'g2' (Large GPU option) allows the use of a large card, 'g3' (Extreme GPU option) allows the use of a Extreme GPU card. The use of multiple cards are specified by the 'Server flag' (see Table 36.1)
G	Stitcher
H	Leica HyVolution LASX wizard
I	Abberior bundle
l	Localization Microscopy
L	Leica LAS AF - Huygens data exchange
M	Movie Maker
N	Enable new, pre-released features
P	Byte limit for image indexer (Titan)
s	Object Stabilizer
t	Time Series
T	Object Tracker
u	Crosstalk Corrector
v	Advanced visualization (Surface Renderer and advanced SFP Renderer)
x	Chromatic Aberration Corrector

³<https://www.svi.nl/GPUcards>

⁴<https://www.svi.nl/GPUcards>

Index

Symbols

4K monitors 8

A

Acquisition Pitfalls 213
Address 225
Airyscan 53
 Load image 54
 modes 54
 Visualize 54
Alignment of Slices 147
Anchor 179
Animation
 Frame count 117
 Frame rate 117
 MIP 97
 Movie Maker 113
 SFP 106
 Surface 111
Array Detector 53
Aspect Ratio 171
AVI files 115

B

Background 19
Backprojected 210
 Pinhole radius 13
 Pinhole spacing 13
Base line 209
Batch Express 29
Batch Processor 39
Bead Images
 Confocal 36
 Two Photon 36
 Widefield 36
Beads 33, 36
Black level 209
Bleaching 133, 215
Bleaching correction 21
Bleaching Corrector 133
Bleaching factors 135
Bleedthrough 129
Bounce 118
Brick layout 21
Brick-Wise 207
Brightfield images 223
Brightness 87

C

Center scene 110
Channels 23
Charge leakage 125
Chromatic Aberration 137
Chromatic Aberration Corrector 137
Chromatic Shift 137
Clipping 18, 214
CMLE 19
Cold pixel 125
Colocalization 187
 Coefficients 187
 Colocalization map 187
Colocalization Analyzer 187
 Isosurface threshold 195
Color 219
Colors 85
Comparison View 22
Concurrent tasks 44
Contrast 87
Convert
 Dataset 12
Corrector
 Bleaching 133
 Chromatic Aberration 137
 Crosstalk 129
 Hot & cold pixels 125
Correlation 127
Coverslip 221
Coverslip position 13
Cropper 16
Cropping
 in x, y and z 16
 in Time 17
Cross correlation 138
Cross Correlation Method 144
Crosshair Cursor 92
Crosstalk 129
Crosstalk coefficients 131
Crosstalk Corrector 129
CUDA graphics cards 197

D

Dead pixels 125
Deconvolution
 Express 25
 Improvements 216
 STED 47
 Tiles 124
 Wizard 11
deconvolution

Airyscan	53
Deconvolution Express	25
Deconvolution Templates	42
Deconvolution Wizard	11
Depletion beam	47
Depth-coding	97
Deskewed LLS	61
Detection lens	61
Distiller	33
Drift	24

E

E-mail	225
Electronic Shift	209
Emission	
Radiation	129
Transparencies	102
Wavelength	13
Emission transparency	101
Equal intensities	107
Error function	87
Establishing Parameters	207
Estimate	
Bleaching	134
Excitation	
Beam	222
Fill factor	13
Photon count	13
Wavelength	13
Excitation lens	61
Excitation mode	61
Excitation transparency	101
Experiment preset	165
Export format	
ASCII data	9
Express	
Batch	29
Deconvolution	25

F

Fading	133
Fax number	225
File	
Open	12
Save	24
Series	217
Fill factor	13
Filter	
Object	183
Tracks	154
Floating License	204
Flow	157
Freeware mode	4
Fusion & Deconvolution	71
Fusion & Deconvolution Wizard	65
Fusion settings	66, 77

G

Gallery	119
Gamma	87
Garbage Volume	162
Gaussian light sheet	61
Gaussian profile	60

Geometrical distortion	213
GMLE	19
GPU	
Acceleration	197
Check list	199
NVIDIA® drivers	198
Status window	199
Support	200
troubleshooting	199
GUI color theme	
Day theme	8
Night theme	8

H

High fill factor	
Cylinder	61
Scanning beam	61
Histogram	18, 223
Hot & Cold Pixel Remover	125
Hot pixel	125
Hot/cold pixel mask	128
HRM	205
Hue selector	220
Huygens Core	205
Huygens Remote Manager	205

I

Image	
Histogram	18
Load	12
Multi-channel	23
Quality	213
Image Feeder	9
Imaging direction	13
Input conversion	44
Instability	
Illumination	215
Mechanical	215
Installation	
Linux	
Debian	3
RPM	3
Mac OS X	3
Support	8
Windows	3
Intensity	157
Interactive Fuser	71
Internal Reflection	215
Internal reflection	214
Iso-surface	169
Iteration mode	21
Iterations	20

J

Jitter	133
--------------	-----

K

Keyframes	114
-----------------	-----

L

Lamp jitter	133
-------------------	-----

lattice light sheet	61
Leica numbering	218
Lens refractive index	13
License	
License File	6
String	4, 226
Troubleshooting	6
License Details	226
Light Scattering	60
Light Sheet	65, 71
Light Sheet Fluorescence Microscopy	59
Linking	88
LSFM	59

M

Mailing Address	225
Maximum Intensity Projection	95
Mean squared displacement	156
Measure	
Aspect ratio	171
Colocalization	171
Correlation	171
Length	169
Principal Axis	169
Sphericity	170
Usage	10
Measured PSF	16
Medium refractive index	13
Memory	7
Microscope type	13
Microscopic Templates	42
Microscopy Parameters	13
MIP Renderer	95
Mouse mode	164
Movie Maker	113
Movie Quality	118
Multi-channel Images	23
Multi-GPU	197
Multiple Users	203

N

Number of iterations	20
Numerical aperture	13
Nyquist rate	209

O

Object	
Analyzer	159
Filtering	183
Measuring	166
Stabilizer	143
Statistics	165
Tracker	151
Object brightness	101
Objective quality	13
Orthogonal Slicer	91
Orthogonal views	92
Output folder	30
Output format	44
Overfill factor	222

P

Parameters	
------------	--

Array Detector	55
Establishing	207
LSFM	60
Microscopy	13
MIP	96
Resampling	67
Rotation	68, 75
Scatter	62
SFP	103
STED	48
Surface	109
Tracking	153
Penetration depth	101
Performance options	198
Phone number	225
Photodamage	214
Pinhole	
Distance	210
Radius	13, 210
Spacing	13
Pipe Mode	178
Pitfalls	213
Pixel	
Cold	125
Hot	125
Point Spread Function	217
Preferences	8
Preview area	113
Principal Axis	169
Projection	85
PSF	
Distiller	33
Measured	16

Q

QMLE	19
Quality	213
Factor	217
Render quality	96, 100, 108
Threshold	20
Queue manager	205

R

RBNCC	193
Reduce mode	21
Reduction	21
Reflection	215
Refractive	
Lens	13
Medium	13
Refractive Index	213
Region of interest	152
Remote desktop protocol	204
Remote Display	203
Render Pipes	164
Render size	97, 103, 109, 174
Renderer	
MIP	95
SFP	99
Surface	107
Resampling parameters	67
ROI	152, 175
ROI Deconvolution	26
Rotation Parameters	68, 75

S	
Sampling Density	209
Saturation	18, 220
Save scene	97, 106, 111
Scanning Bessel beam	61
Scatter Parameters	62
Scattering	60, 62
Segmenting	161
SFP Renderer	99
Shading correction	121, 123
Shadow transparency	101
Sharpness	207
Signal to Noise Ratio	20, 207
Simulated Fluorescence Process	99
Slicer	
Orthogonal	91
Twin	83
SNR	20, 207
Soft threshold	101, 109
Spherical aberration	213
Sphericity	170
SPIM	59
SPIM Fusion & Deconvolution Wizard	65
SPIM Interactive Fuser	71
SPIM/Light Sheet Deconvolution	59
SPIM/Light Sheet Fusion Wizard	65
Stabilizer	143
STED	47
Deconvolution	51
Estimating parameters	49
Parameters	48
Principle	47
Wavelength	48
Stitcher	121
Stitching & Deconvolution Wizard	121
Storyboard	115
Supersampling	21
Support	225
Surface Renderer	107
System Id	4
System Requirements	7
T	
Task	43
Templates	
Bleaching	135
Deconvolution	42
Hot/cold pixel mask	128
Location	43
Microscopic	42
MIP	97
SFP	106
Surface	111
Thermal Effects	215
Thread restrictions	9
Threads per task	44
Threshold	
MIP	96
Quality	20
SFP	101
Surface	108
Threshold Estimation	
Costes method	191
Gaussian minimum	191
Optimized search	192
Tiff Series	115
Tiff series	218
Time frame	180
Time Series	12
Z-drift	24
Timeline	116
Track Analyzer	154
Tracker	151
Transformation	67
Transitions	115
Transparency	
Emission	101
Excitation	101
Shadow	101
Transparency depth	174
Twin Slicer	83
U	
Undersampling	214
Uninstall	7
Unsupervised Profiles	26
Updates	6
Usage Report	10
V	
Velocity	156
Vignetting	121
Vortex	51
W	
Watch folder	30
Watershed segmentation	162
Wavelength	
Emission	13
Excitation	13
STED	48
Website	225
Width of the Gaussian sheet	61
Wizard	11
Z	
Z depletion	51
Z-drift	24

