Huygens Essential

User Guide for version 17.10

Scientific Volume Imaging B.V.

Huygens Essential

User Guide for version 17.10

Scientific Volume Imaging B.V.

Mailing Address	Scientific Volume Imaging B.V. Laapersveld 63 1213 VB Hilversum The Netherlands
Phone	+31 35 6421626
Fax	+31 35 6837971
E-mail	info@svi.nl
URL	http://www.svi.nl/

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

Contents

CHAPTER 1 *Introduction* 1

CHAPTER 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
4K Monitors 8
Measuring the software usage 8

CHAPTER 3 *Titan* 9

Titan as Freeware 9
Titan does not edit or move your valuable image data 10
Starting Titan 10
Searching images 12
Annotations and Projects 12
Viewing and forwarding images 13

CHAPTER 4 *Deconvolution Wizard* 15

The Processing Stages in the Wizard 15
Loading an Image 16

Verifying Microscopy Parameters 17
Starting the Deconvolution Wizard 19
Using a Measured PSF 20
The Intelligent Cropper 20
Stabilizing STED data 22
The Image Histogram 22
Estimating the Average Background 23
The Deconvolution Stage 24
Finishing or Restarting a Deconvolution Run 26
The Final Stage 26
The Comparison View 26
Multi-channel Images 27
Z-drift Correcting for Time Series 28
Saving the Result 28

CHAPTER 5 *Deconvolution Express* 29

Why Deconvolution Express 29
Verifying the image parameters 29
Starting Deconvolution Express 29
Unsupervised Profiles 30

CHAPTER 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 36
Distillation Stage 37
Finalizing the result 37
Full-width at half-maximum estimator 37

CHAPTER 7 *The Batch Processor* 39

The Batch Processor Window 39
Usage 40
Menus 44

CHAPTER 8 *Huygens STED Deconvolution* 45

Introduction 45
STED principle 45
STED parameters 46
Example settings 47
Estimating STED parameters 47

STED Deconvolution 49
Summary: deconvolution procedure 50

CHAPTER 9 *SPIM/Light Sheet Deconvolution 51*

SPIM image restoration 51
SPIM parameters 51
SPIM deconvolution 54

CHAPTER 10 *SPIM/Light Sheet Fusion & Deconvolution Wizard 55*

Starting the Fusion and Deconvolution Wizard 55
Options Menu 55
Selecting input images 56
Deconvolution settings 57
Transformation and resampling parameters 57
Setting the rotation parameters 58
Channels and Time frames 59
Starting the fusion 59

CHAPTER 11 *The Twin Slicer 61*

Using the Slicer in Basic Mode 62
Using the Slicer in Advanced Mode 64
Measurement 67

CHAPTER 12 *The Orthogonal Slicer 69*

The Crosshair Cursor 70
Visualization parameters 70
Measurements 70
Auto-Zoom 71
Display Options 71

CHAPTER 13 *The MIP Renderer 73*

Basic Usage 73
Options menu 75
Simple Animations 75

CHAPTER 14 *The SFP Renderer 77*

Basic Usage 78
Advanced Usage 80
Simple Animations 81

CHAPTER 15 *The Surface Renderer* **83**
Basic Usage **84**
Advanced Usage **85**
Simple Animations **86**

CHAPTER 16 *The Movie Maker* **89**
An Overview **89**
Creating and Adjusting Keyframes **90**
Using the Storyboard **91**
Working with Movie Projects **92**
Using the Timeline **92**
Advanced Topics **93**

CHAPTER 17 *The Gallery* **95**

CHAPTER 18 *Stitching & Deconvolution Wizard* **97**
Stitching & Deconvolution **97**
Vignetting and Shading correction **97**
Starting the Wizard **97**

CHAPTER 19 *The Hot Pixel Remover* **101**
Why correct for hot pixels? **101**
Hot Pixel Remover window **101**
Hot Pixel Detection **102**
Correcting and finalizing **103**
Using a Mask **104**
Batch Hot Pixel correction **104**

CHAPTER 20 *The Crosstalk Corrector* **107**
What is crosstalk? **107**
Starting the Crosstalk Corrector **107**
Adjust the crosstalk coefficients **109**

CHAPTER 21 *Bleaching Corrector* **111**
What is bleaching? **111**
Starting the Bleaching Corrector **112**
Estimate Bleaching **112**
Adjusting Bleaching Factors **113**
Correct Bleaching **114**

CHAPTER 22 *The Chromatic Aberration Corrector* **115**

Causes of chromatic aberration **115**
Starting the Chromatic Aberration Corrector **116**
Estimation of the chromatic Aberration **116**
Visualization of the chromatic aberration **117**
Editing the chromatic shift vectors **118**
Working with templates **119**

CHAPTER 23 *The Object Stabilizer* **121**

Stabilization of 3D Time Series **121**
Alignment of Slices in 3D Stacks **126**

CHAPTER 24 *The Object Tracker* **127**

Introduction **127**
The Object Tracker Wizard **127**
The Track Analyzer **131**

CHAPTER 25 *Introduction to the Object Analyzer* **135**

Starting the Object Analyzer **135**
Segmenting the Objects: Setting the Threshold **137**
Interaction with the Objects **140**
Render Pipes **140**
Object Statistics **141**
Storing your Results **145**
Further Reading **145**

CHAPTER 26 *Object Analyzer Geometry Measurements* **147**

Iso-surface **147**
Principal Axis **147**
Length and width **147**
Sphericity **148**
Aspect Ratio **149**
More Parameters and Filtering **149**

CHAPTER 27 *Object Analyzer Component Reference* **151**

Main window components **152**

CHAPTER 28 *The Colocalization Analyzer* **167**

How to use the Colocalization Analyzer **167**

The colocalization map 169
Iso-colocalization object analysis 171
Threshold Estimators 171
Threshold Options 172
Selecting a ROI image 173
RBNCC 173

CHAPTER 29 *Huygens multi-GPU acceleration* 175

Huygens GPU options 175
Huygens and GPU acceleration 176
How to prepare GPU drivers and CUDA for Huygens 176
GPU Check list 176
GPU status Window 177
Multi GPU Support 178

CHAPTER 30 *Huygens for Multiple Users* 181

Remote Display 181
Floating License(s) 182
Huygens web accessibility with HRM 183

CHAPTER 31 *Establishing Image Parameters* 185

Image Size 185
Signal to Noise Ratio 185
Black Level 187
Sampling Density 187
Computing the Backprojected Pinhole Radius and Distance 189

CHAPTER 32 *Improving Image Quality* 193

Data Acquisition Pitfalls 193
Deconvolution Improvements 196

CHAPTER 33 *Appendix* 197

The Point Spread Function 197
Quality Factor 197
File Series 198
Image Feeder 199
Adjusting the Global Color Scheme 199
Hue Selector 200
Setting the Coverslip Position 200
Excitation Beam Overfill Factor 202

Brightfield Images 203
2D Histogram 204

CHAPTER 34 *Support and Contact Information* *207*

Contact Information 207
Support 207
License String Details 208

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 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 fortified with a versatile and intuitive batch processor.

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

CHAPTER 2

Installation

Huygens Essential can be downloaded from the SVI website[1].

Microsoft Windows

Double click on the Huygens installer executable, e.g. `huygens-17101p0_i386.exe`. Double click its icon to start the installation. During installation the directory `C:\Program files\SVI\` will be created by default, and five Huygens icons appear on the desktop. Double clicking on the Huygens Essential icon starts the program.

Mac OS X

Double click the package file, for instance `huygens-17.10.1-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 a X11 application, you need to install X11[2].

Linux (Debian)

Debian packages are natively used by Ubuntu and other Debian-based Linux distributions. Double click the package file, e.g. `huygens_17.10.1-p0_amd64_olderGLIBC.deb`, and follow the steps in the package manager. To install the package through the command line:

```
dpkg -i huygens_17.10.1-p0_amd64_olderGLIBC.deb
```

Linux (RPM)

RPM (RedHat Package Manager) packages are natively used by RedHat, Fedora, SUSE, and other RPM-based Linux distributions. Double click the package file, e.g. `huygens-17.10.1-p0.x86_64.rpm`, and follow the steps in the package manager, or install the package through the command line:

1. `http://www.svi.nl/Download`
2. `http://www.svi.nl/MacOSX`

```
rpm -ivh --force huygens-17.10.1-p0.x86_64.rpm
```

After the Installation

After a first-time installation there is not yet a license available. However, the software can still be started. Without a license it will run in *Freeware mode* . The *System ID*, necessary for generating a license, is obtained by pressing the REQUEST TEST LICENSE button is obtained by pressing the GET A LICENSE button (See Figure 2.1) when opening . It can also be found in the HELP→LICENSE OVERVIEW menu. The next section explains how to obtain and install a license string.

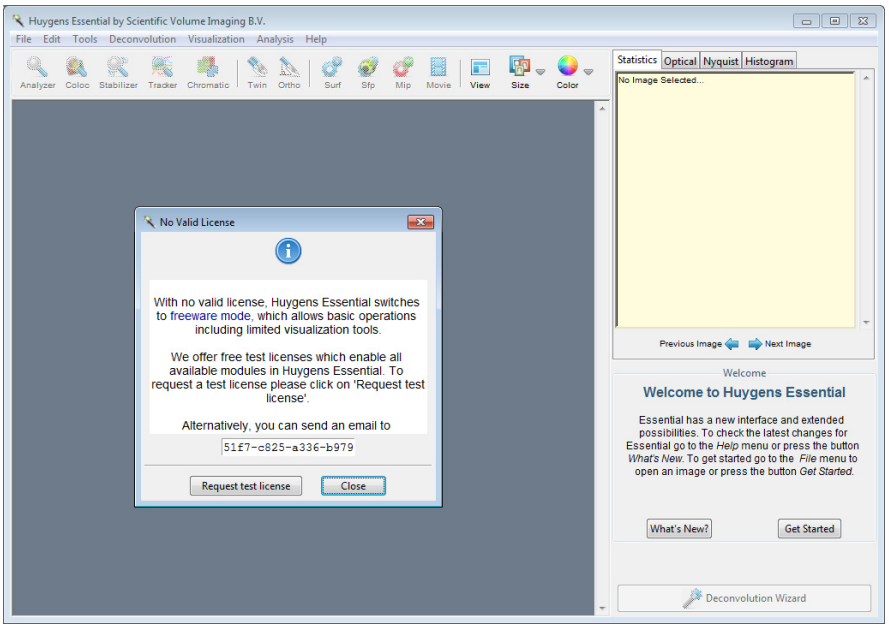

FIGURE 2.1. The startup window of Huygens Essential. If no license string is installed the software runs in *freeware mode*.

The License String

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

```
HuEss-17.10-wcnps-d-tvAC-emnps-eom2015Dec31-
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

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

Obtaining a License String

If upgrading is not handled from a previous installation it is likely that a license is not yet available. To enable us to generate a license string, we need the *fingerprint* of the computer used, the so-called *system ID* number.

If Huygens 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). Send it to `sales@svi.nl`, and a license string will be provided.

To prevent any typing error use the COPY button to save the ID to the clipboard. It can be printed into the license mail message with the EDIT→PASTE menu item of the mail program.

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

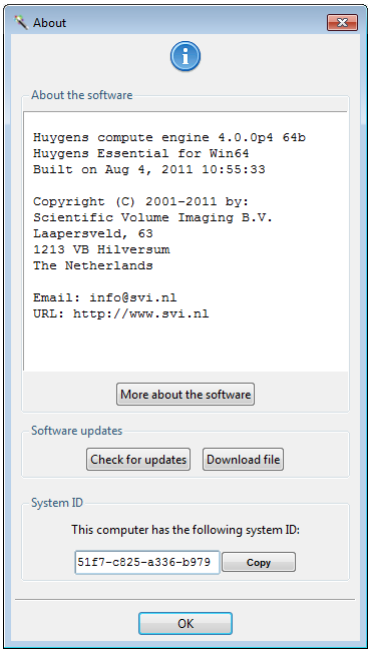

FIGURE 2.2. The HELP→ABOUT window. The system ID is shown at the bottom.

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: 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 remain in Freeware mode. When the license is correct the message "*Added license successfully*" will appear.

Restart Huygens Essential to activate the new license!

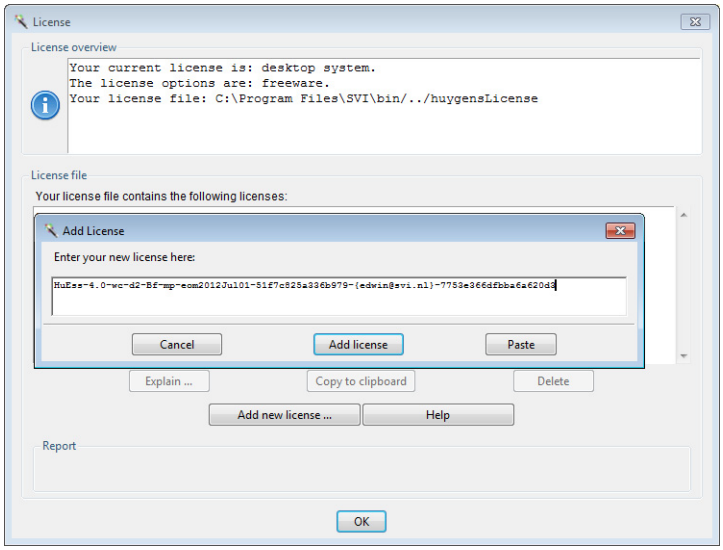

FIGURE 2.3. The license window allows to add, delete and troubleshoot licenses.

Location of the License File

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

TABLE 2.1. The default installation paths per platform.

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

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

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

Troubleshooting License Strings

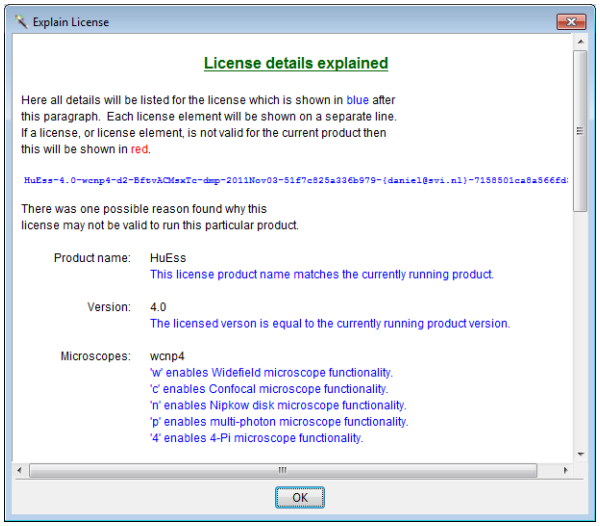

FIGURE 2.4. The *Explain License* window lists all license details.

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

Updating the Software

If the system is attached to the Internet a pop-up window will appear when a newer version is available. The website can also be consulted for updates. Twice a year around April and October new releases will become available. Download the new version from the SVI website[3]. 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,

3. `http://www.svi.nl/Download`

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.

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 to license string is stored. Table 2.2 on page 7 shows the uninstall procedure for each platform.

TABLE 2.2. The uninstallation procedure per platform.

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

System Requirements for Huygens Essential

The tables shown below the requirements for Windows, Mac OS X, and Linux.

TABLE 2.3. System requirements for Microsoft Windows.

Operating system	Huygens runs on Microsoft Windows Vista, Windows 7, Windows 8 , and Windows 10 (64 bit)
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.
Memory	4 Gb or more.
Graphics card	Any fairly modern card will do.

TABLE 2.4. System requirements for Mac OS X

Operating system	Huygens runs on Mac OS X Mountain Lion, Mavericks, Yosemite, El Capitan (X11, XQuartz)[a].
Processor	G5 PowerPC or Intel.
Memory	4 Gb or more.
Graphics card	Any fairly modern card will do.

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

TABLE 2.5. System requirements for Linux

Operating system	Most popular distributions like Ubuntu, RedHat, Fedora, and SuSE are supported. Linux (64 bit)
Processor	AMD Athlon 64 or Intel Pentium 4 and higher.
Memory	4 Gb or more.
Graphics card	Any fairly modern card will do.

Support on Installation

If any problem is encountered in installing the program or the licenses which could not be solved with the guidelines here included, please search the SVI Wiki[4] or contact SVI (See "Contact Information" on page 207).

4K Monitors

Since version 17.04, Huygens is able to scale fonts, buttons, windows and icons which is useful when using Huygens on a 4K monitor. 4K monitors are monitor with vertically more than 3000 pixels. For Windows and Linux, Huygens is able to detects 4K monitors automatically, and will scale accordingly. Settings regarding scaling can be found in the *Preferences* window (EDIT> PREFERENCE> SCALING). With the entry box, you can manually select the scaling value. A check box allows Huygens to automatically detect the right scaling setting (checkbox enbled) or to use the scaling setting as defined in the Preference file that was saved with the last Huygens instance (checkbox disabled). After changing the scaling setting, Huygens needs to be restarted to have a fully scaled Huygens.

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 your license string is saved. This usage can be viewed with the Usage Report window, see Figure 2.5.

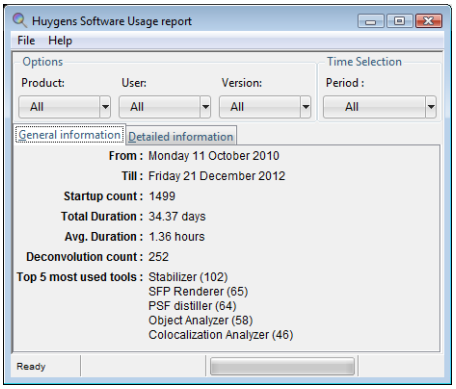

FIGURE 2.5. The usage report window.

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 is 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 in the directory `UsageLogs` and all users are able to write to it. Each year a new log-file will be created and the Usage Report tool analyzes all usage logs in this directory. When the software is updated and/or uninstalled, the usage logs are kept.

4. `http://www.svi.nl/FAQ#Installation_questions`

CHAPTER 3

Titan

How to deal with my growing amount of image data and how to find my images fast without too much effort. These are questions that most microscopists will face at some point.

Related questions like: How can I add keywords and experimental information to images without modifying them; How to get a quick overview of the acquisition parameters; Can I group images into projects without actually editing or moving the raw data; I am not looking forward to re-organize my data and use a database approach, what alternative do I have.

If you recognize yourself in these questions, you may find *Huygens Titan* useful.

Titan as Freeware

Titan is not modifying or moving your valuable microscopy data. It indexes your files and gives you a quick overview of all your images with thumbnails, metadata, histograms, statistics, and annotations (Figure 3.1). Titan reads the common microscope file

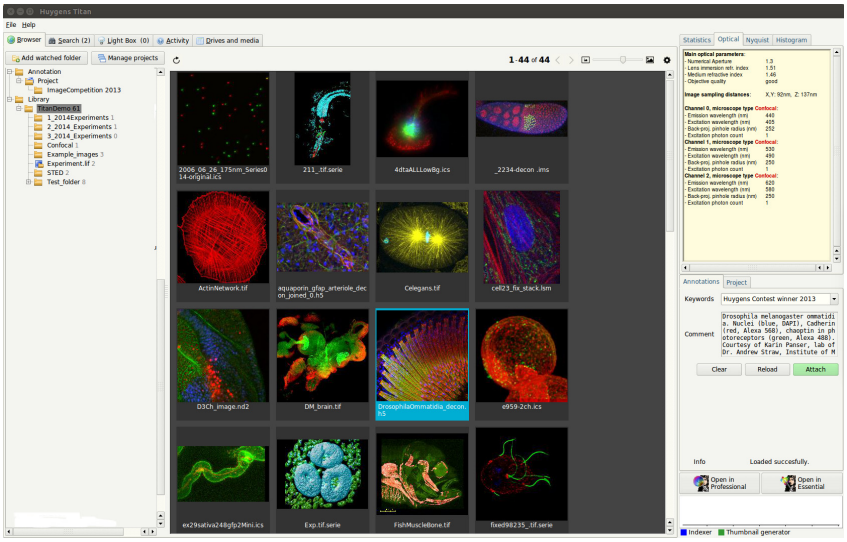

FIGURE 3.1. Huygens Titan shows microscopy data in *Freeware mode*.

formats, and shows embedded images from container formats such as Leica LIF and Zeiss CZI (for a list of formats See “Image File I/O” on page 2). Titan has no problem with the location of the images. Microscopy data may have been saved centrally, located on external hard disks or USB drives, and spread over multiple file storage locations. The ***Freeware version*** of Titan has a capacity of 200 Gigabyte image data. If you need more than that, you can contact us for a test license (see “Contact Information” on page 207).

Titan does not edit or move your valuable image data

There is no need to go through all your existing data if you want to organize your images with Titan. Titan simply does that for you! Just start by specifying in Titan the location of your images. This can be even multiple locations. Titan just visits the images and indexes them automatically, smartly focusing on images that you want to see first. You can quickly search for specific images across all indexed storage locations with search terms including file or folder names, Nyquist values, and/or metadata properties such as channel number, excitation/emission wavelength values, and microscope type (like for example Confocal, STED, Widefield). Keywords and comments can be added to the images, and it is possible to group images into projects. Again, this can all be done without editing the raw images and their location. Selected images can be opened in the *Light Box*, which is a viewer that allows you to inspect multi-dimensional image data (2D-5D). Exporting images to Essential and Professional for further processing is easily done with their corresponding launch buttons.

Starting Titan

Click on the Huygens Titan desktop icon to get started. First, you will see a “Welcome to Titan” screen which gives you a quick and thorough tour through Titan. See Figure 3.2.

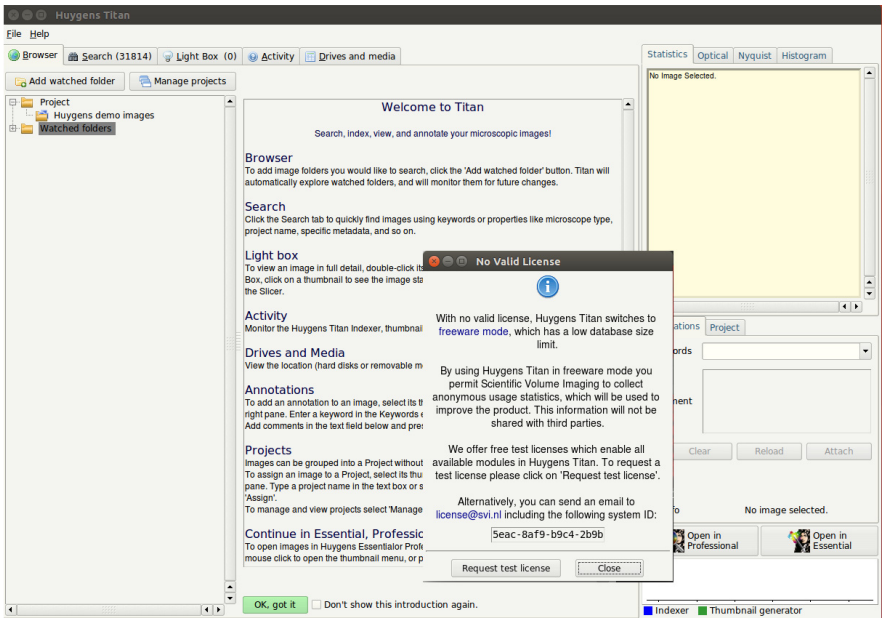

FIGURE 3.2. Upon closing the “No valid license” window, Titan can be used in *Freeware mode* and index up to 200 Gigabyte of data. With a valid test or purchased license, you will be able to index terabytes of data. The “Welcome to Titan” screen will help you start organizing your image data within a few minutes.

Also, a "No valid License" window appears if you do not have a valid Titan test or purchased license. With the help of this window you can request a test license, which we will be happy to send to you. Such a test license will enable you to index terabytes of images. If you only close this window, you can use Titan in *FreeWare* mode and index up to 200 Gigabyte of data.

Define the location of your images

Use within the main window the ADD WATCHED FOLDER button at the top-left, or look under FILE, to open the *Preferences* window for adding image storage locations (folders) to be indexed. The indexing priority can be adjusted by clicking on the respective field. Under the Tab *"Indexer"*, you can also define at what times Titan should leave enough computer resources available for other activities, and what should be the maximum file size to be indexed. The number of CPU's available for Titan can be set, as well as the storage location of where Titan should store the image thumbnails it generates.

Upon hitting the Apply button, Titan immediately starts indexing, as can be seen in the Titan activity graph at the bottom-right. The *Watched folder* name with the number of images it contains, is displayed in the left main window. If a location, such as for example a USB stick or external hard disk, is no longer available, a red cross will appears on a *Watched folder* name. Re-connecting the file storage location will remove this cross again.

Display options for colors and how the thumbnails should be sorted can be set with the settings (wheel) button located on the right above the thumbnail window. The slider next to it enables you to adjust the size of the thumbnails. If the number of images exceed one hundred they will be listed over multiple pages. Scrolling through these pages can be done with the arrow buttons.

At the right side within the main Titan window, there are several tabs located with more information of the selected thumbnail, and on what annotations and which project is attached to the image. (see "Annotations and Projects" on page 12).

Thumbnail messages

There are several messages that can appear on the created thumbnails. They show the indexing status of this specific image. An overview of these messages and the actions that can be taken, can be seen on this wiki page[1]. If, for example, the size of the image is larger than scheduled for indexing (see next paragraph), you may want to force the image to be indexed. This can be done by right-clicking on the thumbnail and selecting *Force Refresh image.*

Scheduling index activity

Indexing can be a computational and time-consuming process, so if there is many data to be indexed you may need to leave Titan active for some time. Once your images are indexed you may appreciate how fast Titan shows their thumbnails and metadata. By default, Titan runs in background mode between 7:00 and 19:00 hours on working days. Under the INDEXER tab within the *Preferences* window (FILE→PREFERENCES), you can manage these settings to save computer resources.

1. `http://www.svi.nl/TitanHelp`

Searching images

Titan allows you to easily search through your images using keywords and multiple metadata properties. Quick searches are available to select for images that have been acquired with correct Nyquist sampling, and for those that are multichannel (suitable for colocalization analysis), and have multiple frames (stabilization). Images can also be searched by file type, and with what microscope type they were acquired. Combinations of search terms can also be applied, and search parameters can be cleared with the cross button above the thumbnail window (Figure 3.3 on page 12). To see the directory of the image shown under the search tab, you can select the image and click on *view containing directory in the titan browser* under the right mouse button menu.

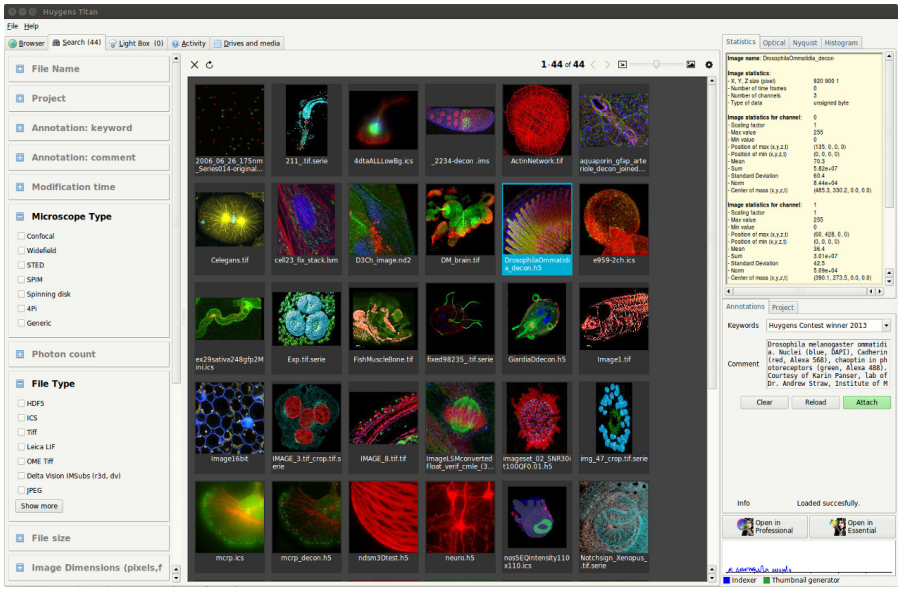

FIGURE 3.3. Titan offers several parameters that can be set to narrow down your search for images.

Annotations and Projects

You can tag images with *Keywords* and *Comments* in their corresponding fields under the ANNOTATIONS tab in the right panel. By clicking on the green ATTACH button, the entry is saved. If you want to clear the entry, use the CLEAR button and then click on the ATTACH button to confirm the empty fields.

To assign images to a project, you can select a thumbnail, click on the PROJECT tab and write a new project name in the *Project* field, or select an existing project name from the drop-down menu. If the project name is new, it will be created automatically. This project will not be shown under the *Project* tree under the BROWSER tab (top-left in the main Projects can window), unless you click the SHOW button. If you want to assign multiple images to a project, you can select the images by choosing a directory under the BROWSER TAB or by searching for images under the SEARCH tab. The assign all button will give these images the specified project name.

Projects can be managed using the MANAGE PROJECTS button or via FILE→MANAGE PROJECTS. Projects can be added or removed here, and project names can be selected for display under the *Project* tree view.

Viewing and forwarding images

At any stage, you can double-click on a thumbnail to add the image to the *Light-box* in which another double-click allows detailed multi-dimensional viewing (Figure 3.4). This can also be done via the menu that appears when clicking the right mouse button on the thumbnail. The same menu can be used to highlight the image in the directory it is located in (use for this purpose *open containing directory*).

Furthermore, you can use this menu to open selected images for further processing in Huygens Essential or Professional. Multiple images can be selected by holding the Shift or CTRL-key. You can also use the HUYGENS ESSENTIAL and HUYGENS PROFESSIONAL icons in Titan for this purpose. Please note that the display of channel colors may change if the image is opened in Essential or Professional. This is due to a different color setting in the software which can be changed easily (see “Adjusting the Global Color Scheme” on page 199).

For more information and a tutorial video we like to refer you to the wiki on Huygens Titan[2].

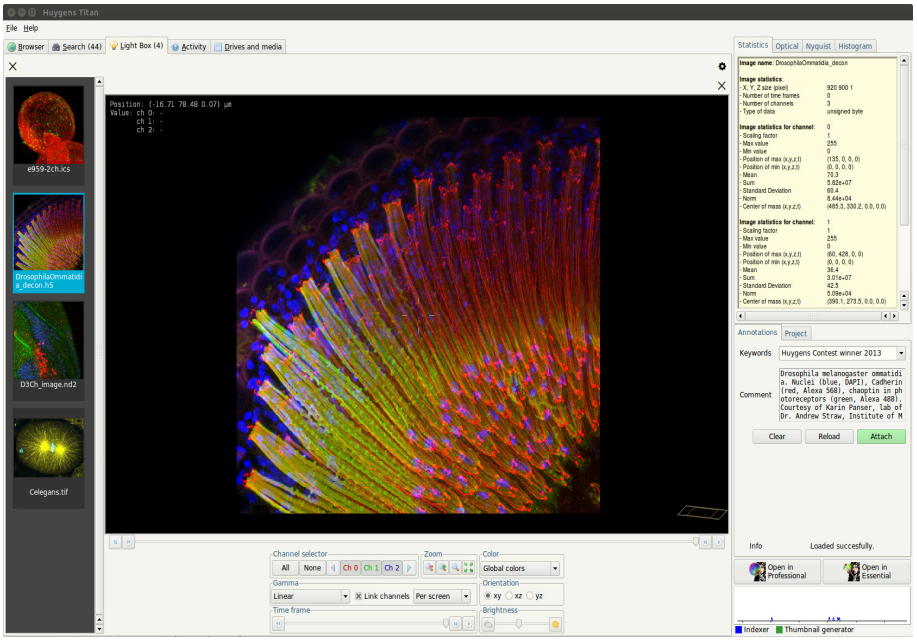

FIGURE 3.4. Slicing through 2D-3D time series and multichannel images is easy in Titan *Light box*.

2. http://www.svi.nl/HuygensTitan

CHAPTER 4

Deconvolution Wizard

The Processing Stages in the Wizard

The deconvolution wizard guides you through the process of microscopic image deconvolution in several stages. In Huygens Essential one needs to verify the microscopic parameters before continuing with deconvolution. This is done with either the Microscopic Parameter Editor or with the Parameter wizard, which in turn guides you through the process of parameter verification. The parameter wizard is especially useful for inexperienced users.

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 status* 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). For searching your image, you may find Huygens Titan a very useful option as it shows your images as thumbnails and offers advanced search tools (see Chapter 3 on page 9).

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 Wiki[1] 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 TOOL menu. For example, a 3D stack can be converted here into a 3D time series (TOOLS→CONVERT->CONVERT XYZ TO XYZT) 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 ALONG Z). The MIRROR tab in the converter tool can also be used for that purpose. Read more at "Setting the Coverslip Position" on page 200.

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 34.2 on page 209). 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 *Invert image,* which is required for the processing of brightfield images (See "Brightfield Images" on page 203.). 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. http://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 either the *Parameter Wizard* or the *Parameter Editor*. The editor and wizard can be found by right-clicking on the image thumbnail and selecting EDITOR or PARAMETER WIZARD. See also Figure 4.1 on page 18and Figure 4.2 on page 18).

TABLE 4.1. Optical parameters explained.

Parameter	Explanation
Microscope type	Select from *Widefield*, C*onfocal*, S*pinning disk*, *4Pi or STED, SPIM, or unspecified*. For a multiphoton microscope insert an *excitation photon count* higher than 1. Select *Widefield* if a nondescanned 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 200
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 46.
SPIM/LSFM parameters	See “SPIM parameters” on page 51.
Deviating microscope type	The deviating microscope type for the current channel.

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

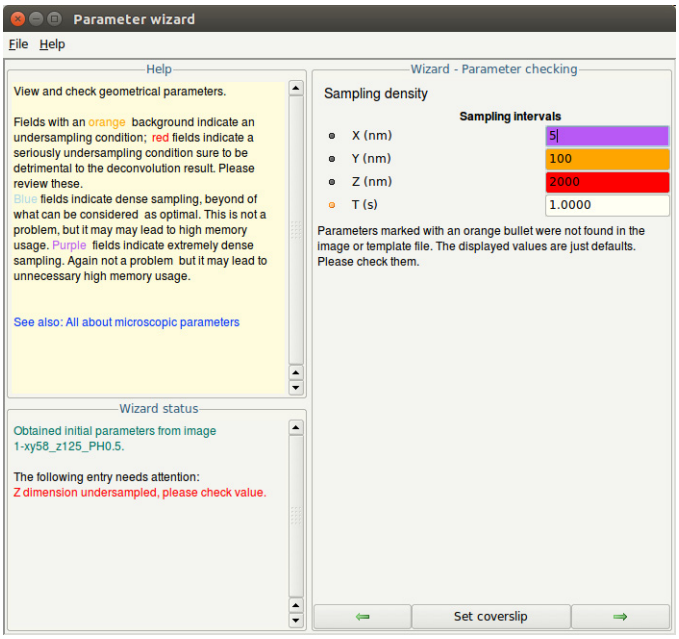

FIGURE 4.1. Parameter wizard: check sampling density. Red coloring indicates a suspicious value, and orange a non-optimal value..EDIT PARAMETERS

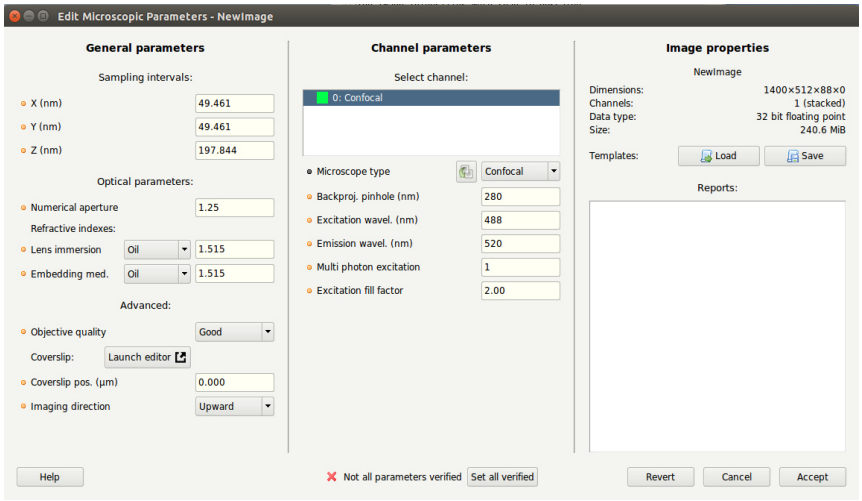

FIGURE 4.2. The microscopy parameter editor. This window can be opened by right-clicking on the thumbnail and selecting EDIT PARAMETERS.

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 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 4.3). You can continue with using the deconvolution wizard or you can execute a deconvolution template previously made.

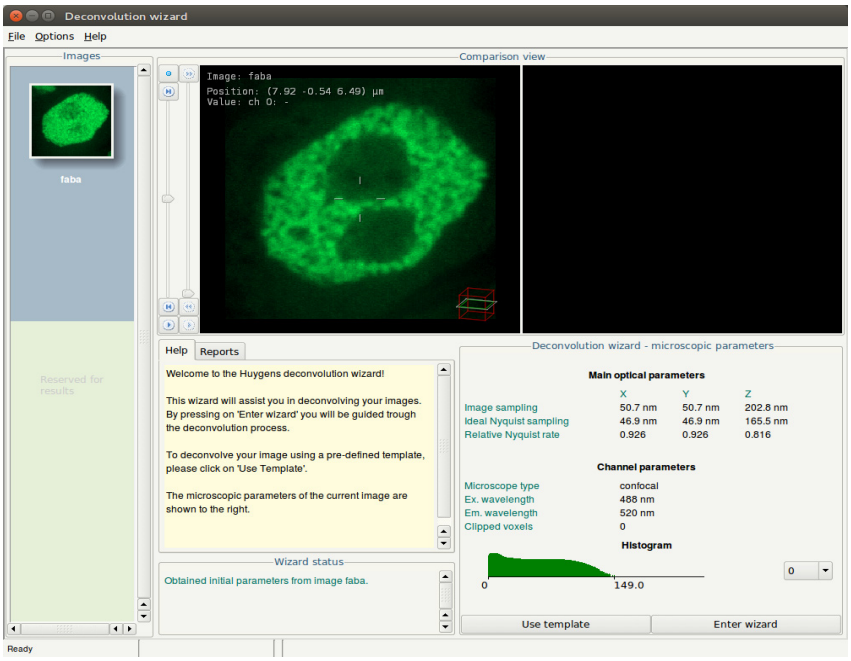

FIGURE 4.3. Starting window of the *Deconvolution Wizard* with a summary of the most essential parameters.

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 4.4). 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.

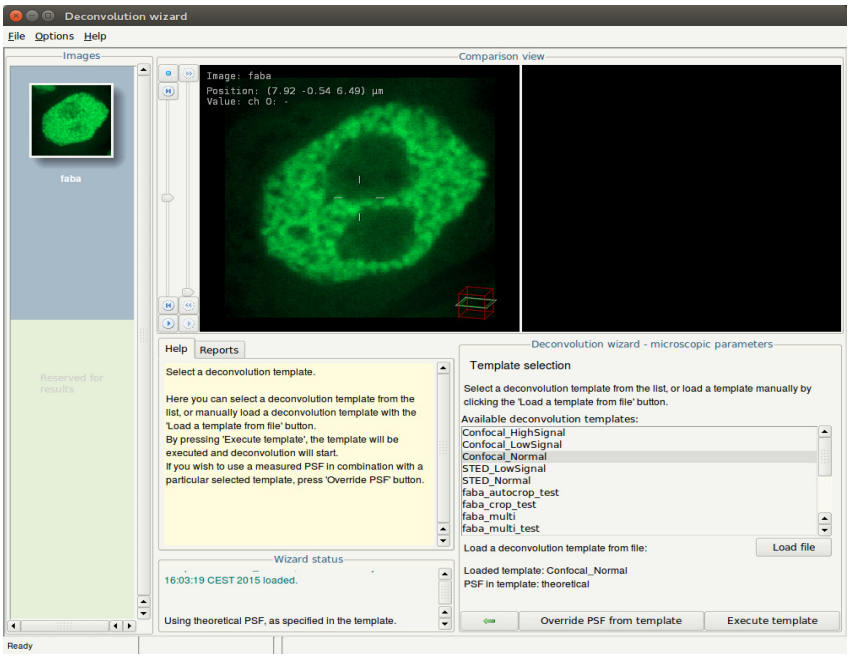

FIGURE 4.4. Deconvolution templates can be selected from the list or can be loaded. A template specifies if an experimental or theoretical PSF is used, but this can be overruled with the button OVERRIDE PSF FROM TEMPLATE.

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 197 for more information.

A previously created measured Point Spread Function (PSF) can be loaded from a file or directly loaded from the main Essential 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 4.5). 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 TOOL→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 13 “The MIP Renderer” on page 73. 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, which can be found under the TOOLS menu or as a taskbar icon.

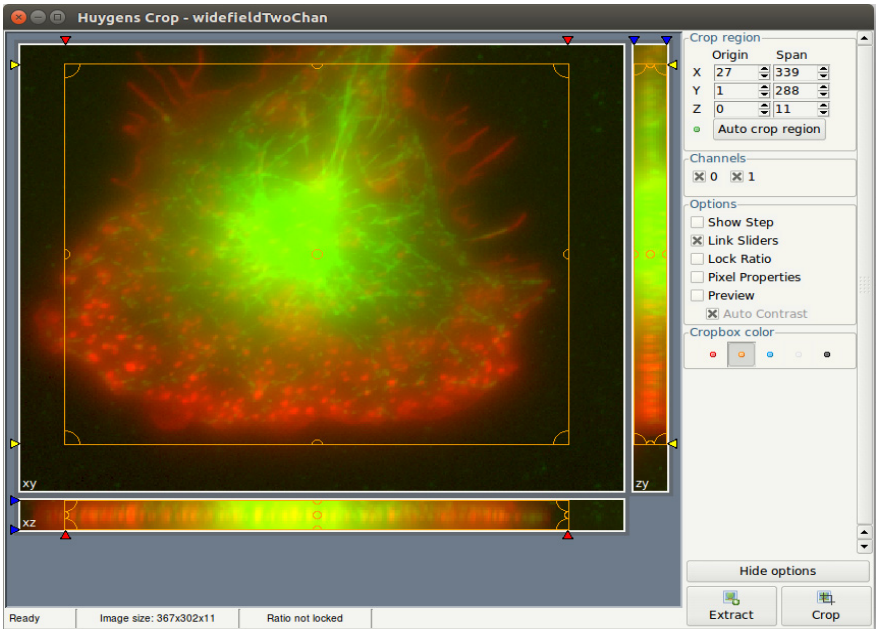

FIGURE 4.5. The crop tool.

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 4.6).

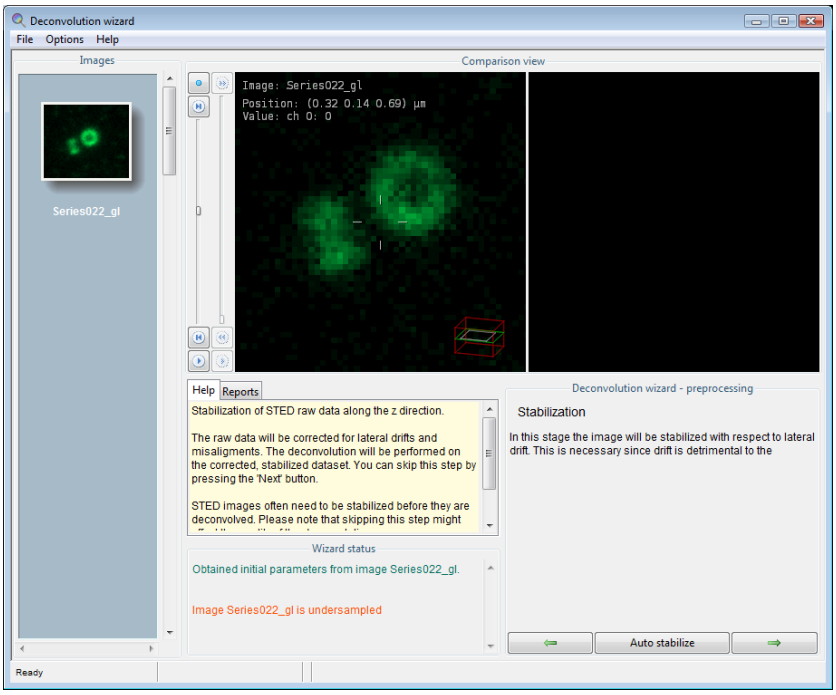

FIGURE 4.6. Stabilization stage for STED data in the Deconvolution Wizard.

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 4.7 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).

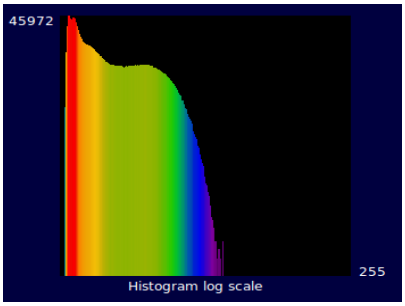

FIGURE 4.7. The image histogram. The vertical mapping mode can be selected from linear or logarithmic.

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 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 follows[2].

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. Subse-

2. Learn more about histograms at `http://www.svi.nl/ImageHistogram`

quently 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 Likelyhood 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.

The Deconvolution Stage

Figure 4.8 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]

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 185.
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

3. www.svi.nl/SignalQuantification

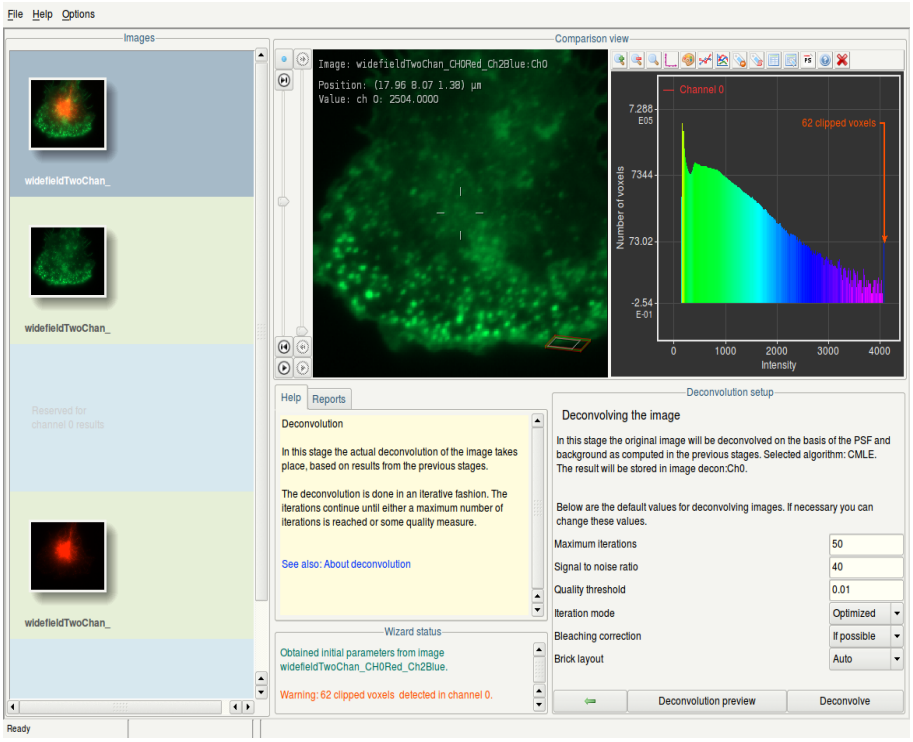

FIGURE 4.8. The Deconvolution Stage in the wizard with the deconvolution preview option.

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.

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.

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 4.9).

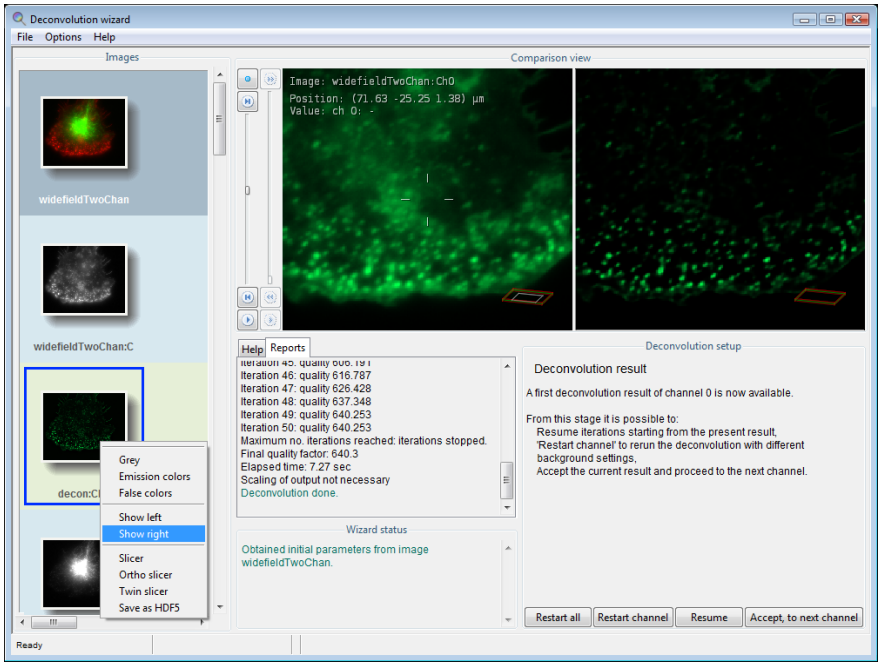

FIGURE 4.9. Showing an image on the Comparison View: right-click on the thumbnail and select "Show left" or "Show right".

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 4.10 on page 28. 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.

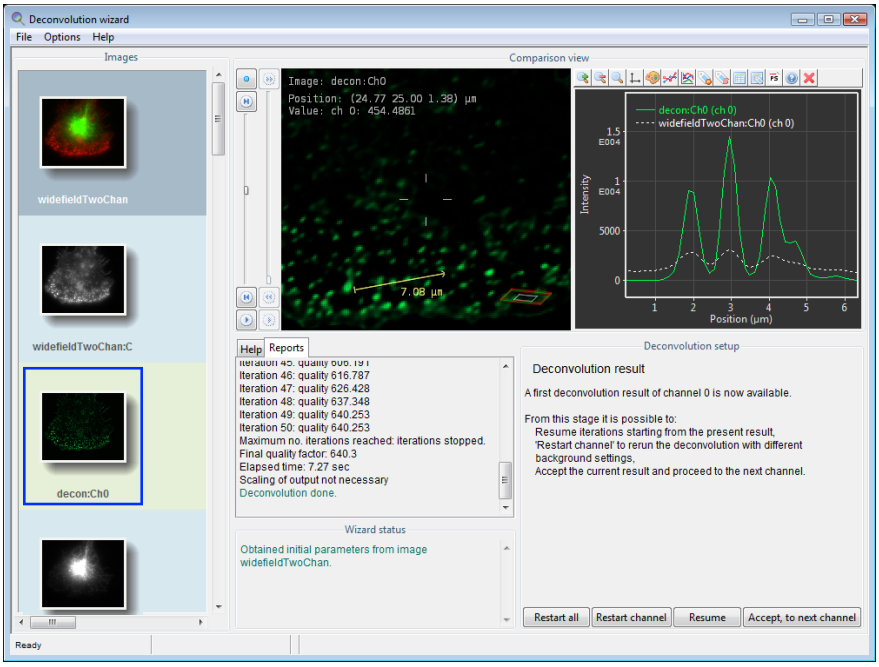

FIGURE 4.10. Comparing intensity profiles of two images in the Deconvolution Wizard. Notice that the plot prompts at the display opposite to where the line is drawn.

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 `http://www.svi.nl/FileFormats`.

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

CHAPTER 5

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 17.

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 5.1 on page 30).

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

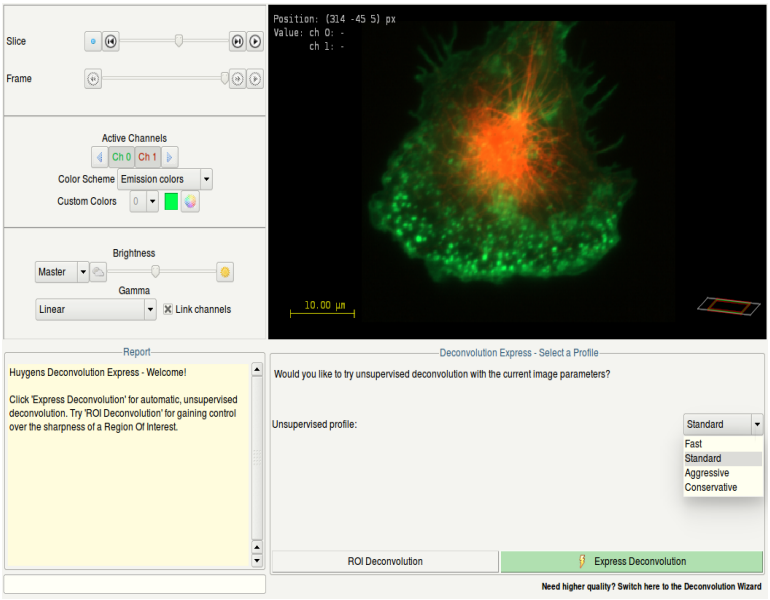

FIGURE 5.1. Deconvolution Express showing the unsupervised profiles.

presents information on how you can proceed and what is relevant for the interpreting the final result.

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 5.1 on page 30). 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, *LEICA*[1]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.

1. Leica Microsystems GmbH

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 microscope[1]. 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.

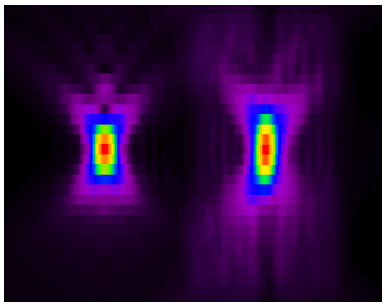

FIGURE 6.1. An *xz* cross section of a theoretical PSF (left) versus a measured PSF (right) for the same confocal setup.

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

1. We 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. `http://www.svi.nl/HalfIntensityWidth`

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 Wiki[3].

The PSF Distiller Window

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

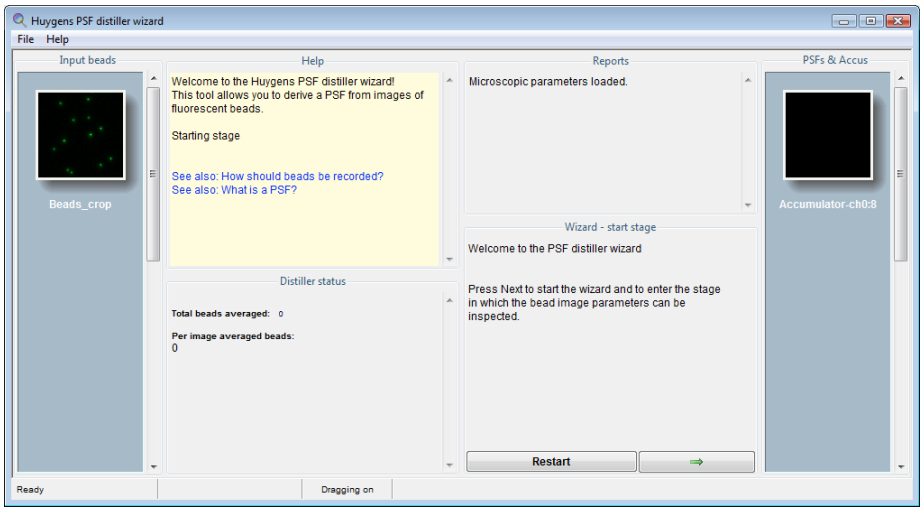

FIGURE 6.2. The PSF Distiller window. Different panes show the input beads field, the Help field, the report field, the wizard fields and the PSF and Accus field.

- *The Input beads pane*

  Shows the selected file with the bead images for the PSF distiller process. You can import and use multiple files with beads for the distillation of a PSF. With the cursor you can hover over the thumbnail and with a right mouse click you can roll down a menu to open the image in one of the slicers. The edit parameters function is 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.

3. http://www.svi.nl/RecordingBeads

- *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 16 and "Saving the Result" on page 28 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 via the menu DECONVOLUTION -> PSF DISTILLER. 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 17 for more information on the microscopic parameters. Next to the optical parameters listed in Table 4.1 on page 17, 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).

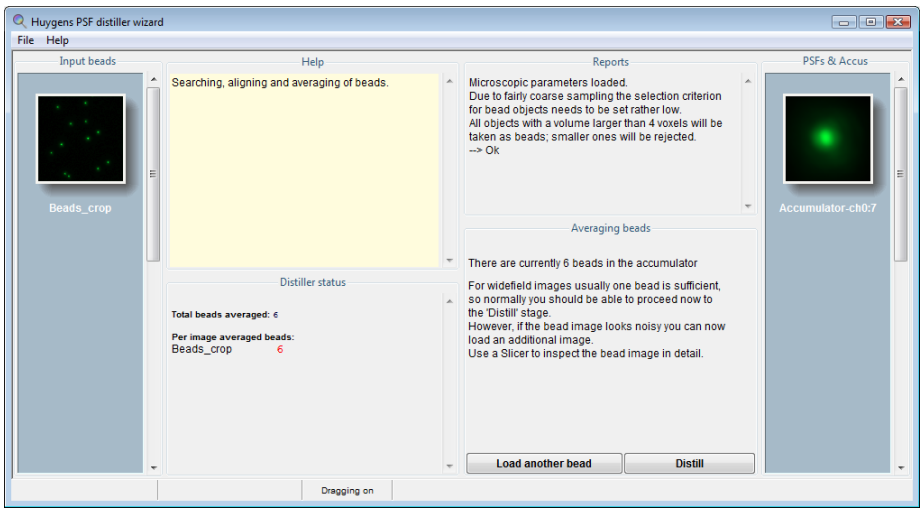

FIGURE 6.3. The averaging stage in the PSF Distiller wizard.

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 Professional under ANALYSIS.

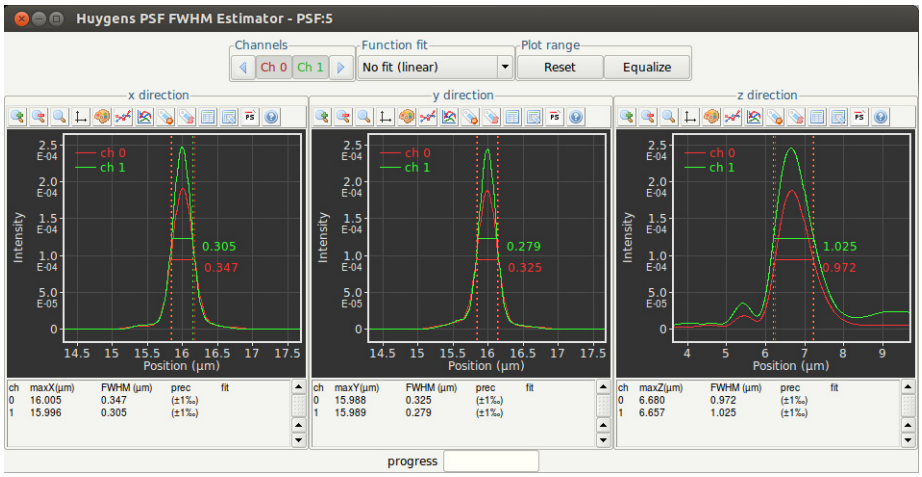

FIGURE 6.4. The PSF-FWHM overview. No curve fitting was selected.

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, Voight and Pearson, are available for this. This is done for all channels. For further background information on the FWHM[4], PSFs and the PSF Distiller[5], please refer to the SVI Wiki.

4. http://www.svi.nl/HalfIntensityWidth
5. http://www.svi.nl/PsfDistiller

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.

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 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 Window

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 the left to right two buttons to load and save batch files into the tasks, two buttons to select and create the destination for the results of the batch run, four buttons to edit the tasks that have been entered in the task area, and four buttons for: using multi-GPU, selecting the output file format, selecting the number of CPU threads that should be used per task, and selecting how many concurrent tasks should be running.
- ***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.

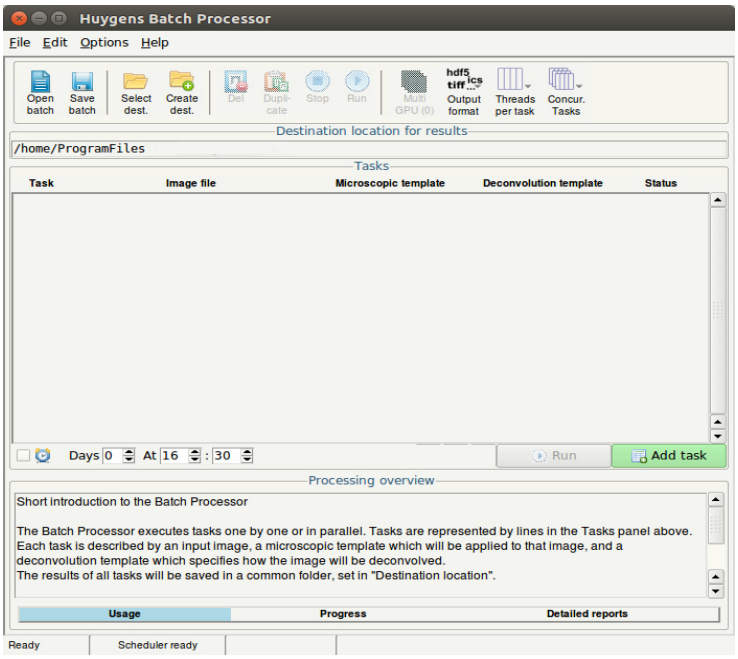

FIGURE 7.1. The Batch Processor main window.

- **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 delete, duplicate, stop, 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 TASK button below the task field a new menu entitled *Selected images* is expanded at the right (see Figure 7.2).

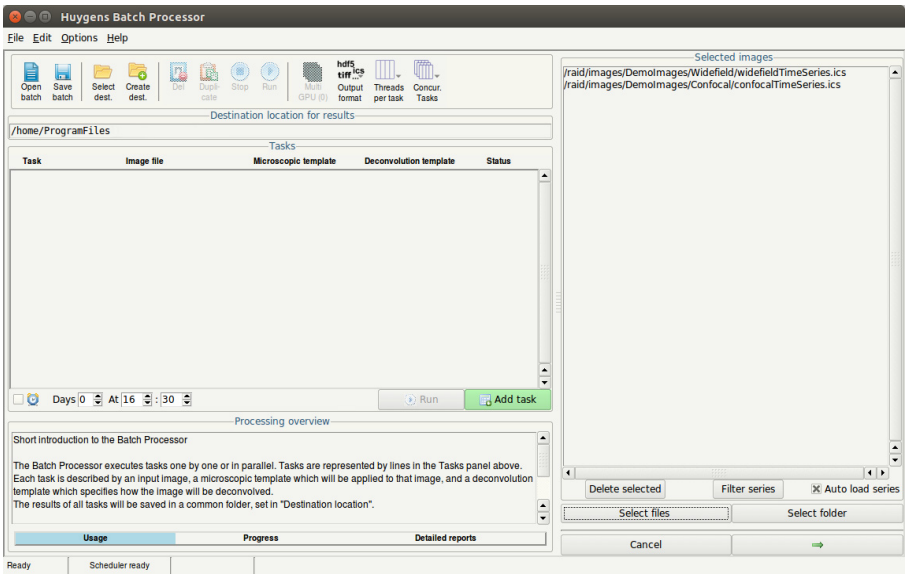

FIGURE 7.2. The Batch Processor with the *Selected images* field expanded.

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

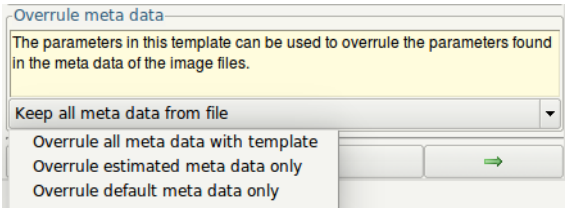

FIGURE 7.3. The Batch Processor metadata selector.

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

The microscopy parameters can be edited and saved to a template using the Parameter Editor or Wizard. See “Microscopy Parameters” on page 13. and “Verifying

Microscopy Parameters" on page 17. These templates can then be reused in the future. 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 levels[1] the Batch Processor allows you to choose whether only reliable metadata should be taken into account for the deconvolution (see Figure 7.3).

It is therefore important to understand what the different microscopy parameters refer to and know how to establish them (See Table 4.1 on page 17). A typically conflictive one is the *backprojected pinhole radius*. This parameter is not difficult to calculate[2], 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 is understood what the role of these parameters is (see "The Deconvolution Stage" on page 24 and the *Restoration Parameters* article[4]).

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 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 24.

With the PRE and POST taps, 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 "Working with templates" on page 119).

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

1. http://www.svi.nl/ConfidenceLevels
2. http://www.svi.nl/BackProjectedPinholeRadius
3. http://www.svi.nl/BackprojectedPinholeCalculator
4. http://www.svi.nl/RestorationParameters

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 wiki[5].

A typically conflictive point is setting the *signal-to-noise ratio*. Mind that this is not a number describing the image, but something that can be tuned to achieve different deconvolution results. Please see "The Deconvolution Stage" on page 24 and the *Set the Signal to Noise Ratio* article in the SVI Wiki[6].

Location of the Saved Templates

The microscopic and deconvolution templates are by default saved in the `SVI/Templates` folder in the user's home directory[7]. 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 deleted 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 (Run) the Batch Processor will start and go over the task list.

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

5. `http://www.svi.nl/BatchProcessor`
6. `http://www.svi.nl/SetTheSignalToNoiseRatio`
7. The 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 (X11, XQuartz) they are usually in `/Users`, and on Linux in `/home`.

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

If the computations are very demanding for the system and 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 three sub menus:

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

  The number of threads Huygens can use in parallel is by default set to AUTO, but in cases where it is required to restrict the computing resources, set a different value as threads.
- CONCURRENT TASKS: if the system has multiple processors 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.

The HELP menu can bring you to our online documention 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-**E**mission **D**epletion (STED) microscopy is a fluorescence microscopy technique which overcomes and improves the diffraction-limited resolution of regular confocal microscopy techniques up to 4-5 times[1], 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 23 on page 121) is equipped with a special STED stabilization mode. The Deconvolution Wizard (see Chapter 4 on page 15) 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.

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

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 25-50m, easily resulting in a four fold increase in resolution over good quality confocal resolution.

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 "EDIT PARAMETERS" or "PARAMETER WIZARD", as explained in Chapter 4 on page 15. If "EDIT PARAMETERS" has been selected, the specific STED parameters will be presented as in Figure 8.1:

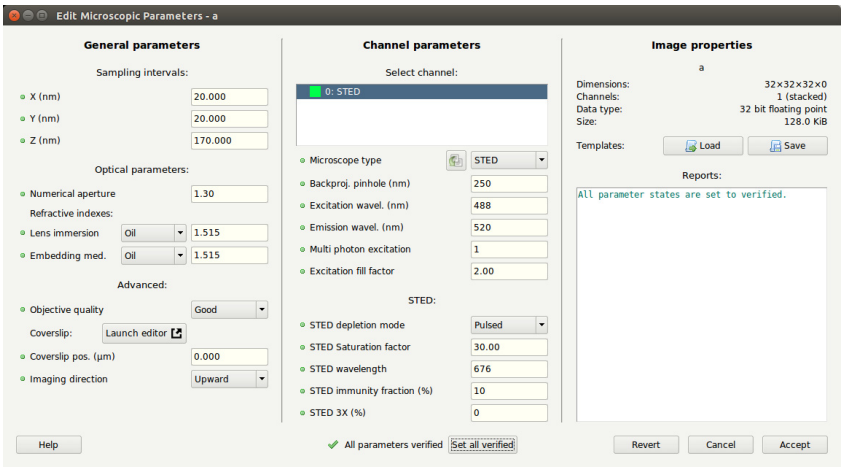

FIGURE 8.1. The specific STED parameters can be revewed and, if neede, edited in the Parameter Editor.

- **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 50[1]. 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.

1. "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 46. 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.

TABLE 8.1.

Parameter	Microscope types			
	TCS STED	TCS STED CW/TCS SP8 STED (3X)		
	pulsed STED	STED CW	gSTED	STED 3X
specified resolution (FWHM)*	70	80	50	n.d.**
typical resolution (FWHM)	60	70	40	n.d.**
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

* measured on 40nm beads; ** not determined yet

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

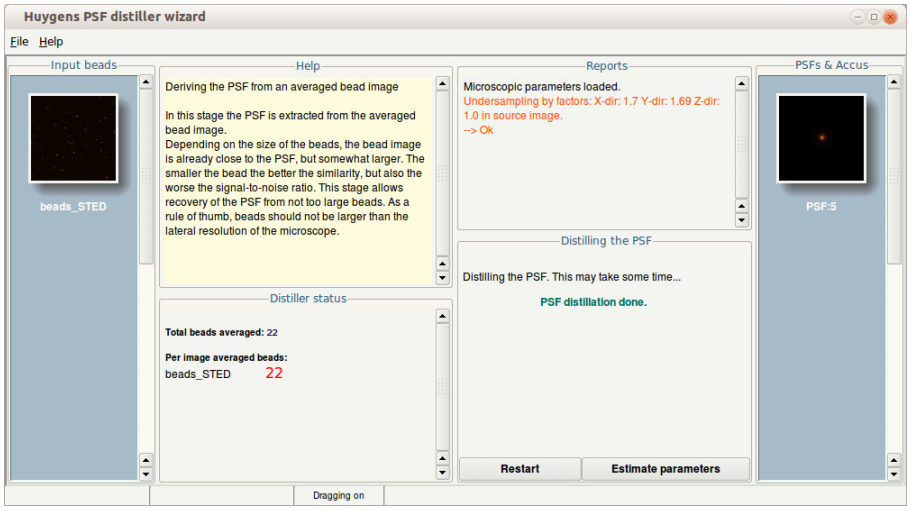

FIGURE 8.2. Estimating the STED microscopic parameters of a beads image automatically with the Huygens PSF distiller.

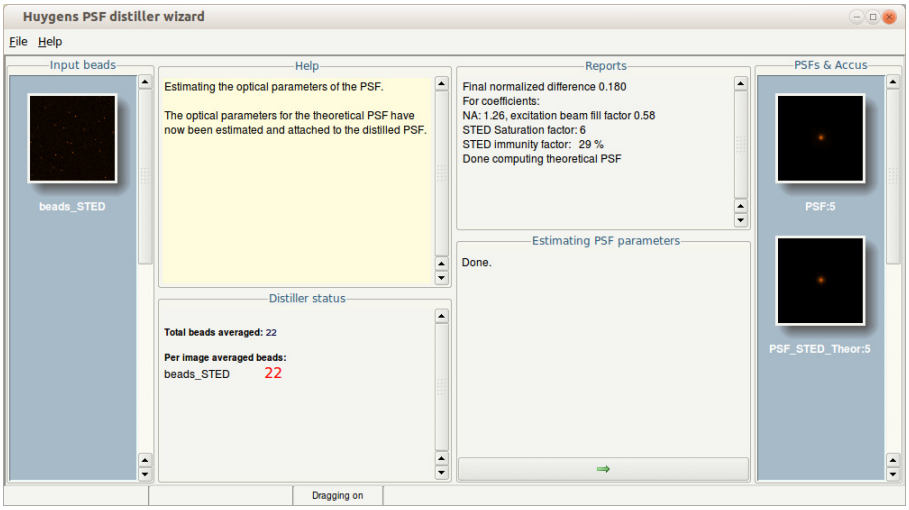

FIGURE 8.3. The PSF distiller showing the STED microscopic parameters automatically estimated from a beads image.

reported by the PSF distiller (see Figure 8.3) and can be reviewed by, for example, opening the Parameter Editor 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 Parameter Wizard or the Parameter Editor (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.

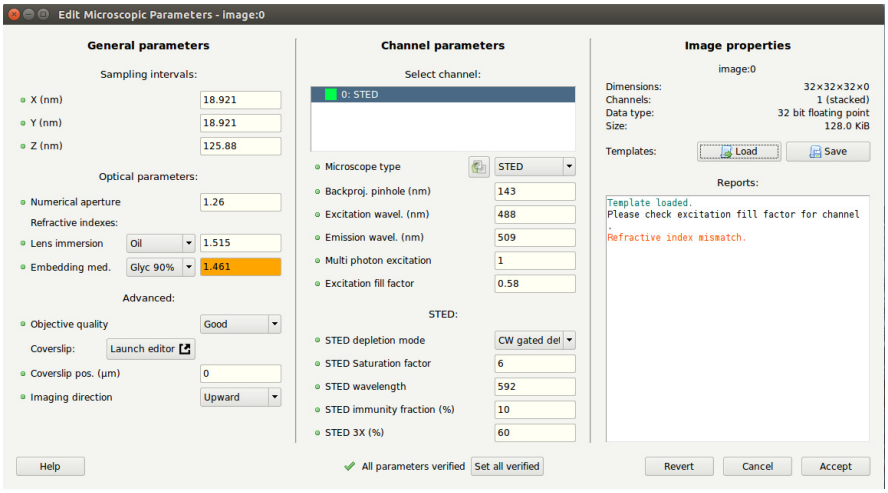

FIGURE 8.4. The Microscopic Parameters of the distilled PSF now contain the values automatically estimated by Huygens Essential.

STED Deconvolution

To deconvolve a STED image, open first the Parameter Editor or Parameter Wizard. 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 4 “Deconvolution Wizard” on page 15. 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 23 on page 121.)**

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

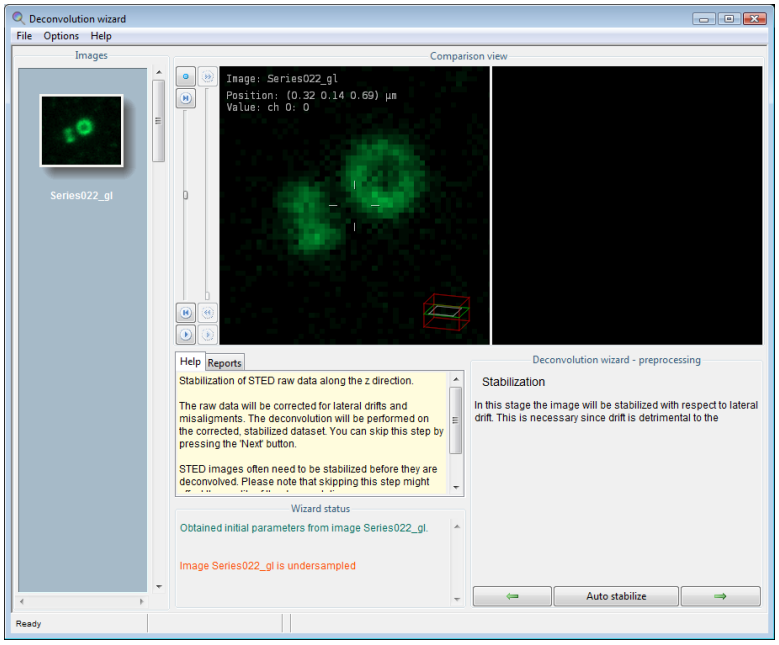

FIGURE 8.5. Stabilization stage for STED data in the Deconvolution Wizard.

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 Parameter Editor or Parameter Wizard).
- Load the STED raw dataset.
- Import the PSF microscopic parameters template file to the STED raw data (via Parameter Editor or Parameter Wizard).
- 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

SPIM/Light Sheet Deconvolution

Selective **P**lane **I**llumination **M**icroscopy (SPIM), also know as **L**ight **S**heet **F**luorescence **M**icroscopy (LSFM), 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. 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 SPIM/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 SPIM and not LSFM.

SPIM image restoration

The SPIM 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 SPIM version (in beta) also offers the option to correct for light scattering, which can be a serious issue in thick samples with light-sheet imaging.

Huygens SPIM deconvolution uses GPU acceleration optimally, and can fully exploit multi-processor systems. The algorithm itself is also optimized for speed.

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

SPIM parameters

To get the most out of Huygens SPIM deconvolution, it is recommended to verify the microscopy parameters of the image data first. If the image meta-data contains informa-

tion about the SPIM parameters Huygens will incorporate them into the image. Otherwise, parameter defaults will be loaded instead.

To review the parameters of the SPIM datasets, select the image thumbnail within the main window of Huygens. Then, right-click on the image and select "EDIT PARAMETERS" or "PARAMETER WIZARD", as explained in Chapter 4 on page 15. If as been selected, the microscopy parameters are displayed as in Figure 9.1."EDIT PARAMETERS".

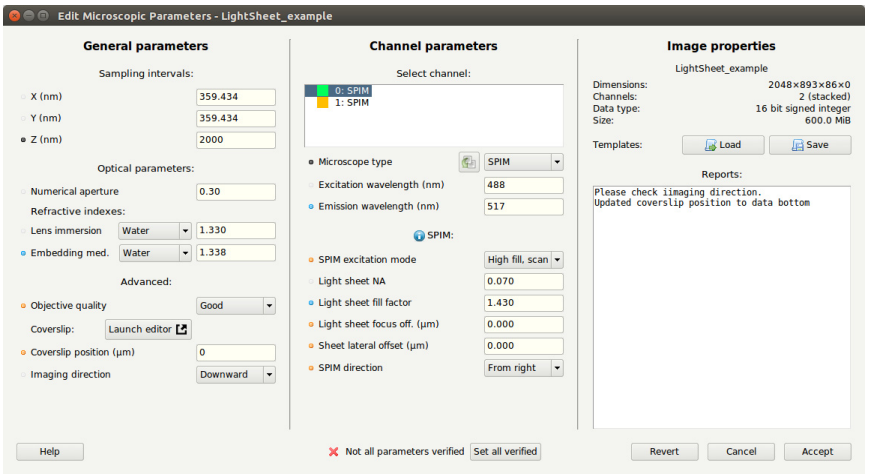

FIGURE 9.1. The general and SPIM-specific microscopy parameters can be reviewed and edited in the Parameter Editor window.

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

1. Light sheets with a Gaussian profile, generated by illumination from one side, either by a cylindrical lens, or by scanning a beam.
2. Lightsheets with a Gaussian profile, generated by simultaneous illumination from two opposing sides.
3. Lightsheets generated by excitation lenses that are overfilled at their entry pupils, either by using a cylindrical lens, or by scanning a beam[1].
4. Lights sheets generated with a scanning Bessel beam

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 SPIM systems such as Zeiss Z1, Leica DLS, LaVision, and MuVi-SPIM (Luxendo), we like to refer you to our wiki page on SPIM deconvolution.[2]

To review and edit the microscopy parameters of an image right-click on the image and select "PARAMETER EDITOR""EDIT PARAMETERS" or "PARAMETER WIZARD", as explained in Chapter 4 on page 15. If "PARAMETER EDITOR""EDIT PARAMETERS" has been selected,

1. www.svi.nl/ExcitationFillFactor
2. www.svi.nl/SPIMDeconvolution

the microscopy parameters for a SPIM image will be presented as in Figure 9.1. More information on the SPIM-specific parameters are explained below.

- **SPIM excitation mode:** this mode defines how the light sheet is produced and offers the following choices:

  Gaussian light sheet Appropriate for SPIM 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 SPIM 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 SPIM system where a scanning Bessel beam is used to form the light sheet. If the multiphoton excitation parameter is adjusted, Huygens is able to account for multi-photon illumination as well.

  High fill factor, cylinder. Appropriate for a SPIM 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 ligh 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 $1/e^2$ of the peak value.

The following parameter appears in the last three excitation modes ("High fill factor, scanning beam", "Scanning bessel beam", 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
  - 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 (SPIM) 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 SPIM 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 simple not well aligned.

- **SPIM 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 by adjusting the:

- 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 SPIM

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 'PARAMETER EDITOR window. In the PARAMETER WIZARD you will find this option under the button SET COVERSLIP. 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 200.

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

SPIM deconvolution

After verifying the general and SPIM-specific microscopy parameters, the image is ready to be deconvolved as any other image via either the *Deconvolution Wizard* or *Batch Processor.* Multiview SPIM images can also be deconvolved within the Fusion & Deconvolution Wizard, where the deconvolution is combined with fusion of the multiple views. See “SPIM/Light Sheet Fusion & Deconvolution Wizard” on page 55.

CHAPTER 10

SPIM/Light Sheet Fusion & Deconvolution Wizard

In many SPIM or light (sheet) microscopy 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. The Huygens SPIM *Fusion & Deconvolution 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.

Starting the Fusion and Deconvolution Wizard

The *Wizard* opens as a separate window in Huygens when selecting via the menu in the main window DECONVOLUTION→FUSION & DECONVOLUTION. 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 helps 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 that are rotated, the final fusion result can be larger in the dimensions perpendicular to the rotation axis, and 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. To retain the full-size result, de-select the "Crop output" entry in the OPTIONS menu.

By the default the data type of the result image is set to 16 bits, even though the results are internally calculated in a higher 32-bit floating point precision. To keep the full preci-

sion, select the "Floating point" entry in the OUTPUT DATA TYPE sub-menu in the OPTIONS menu. To save memory, the user 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***: The value of each voxel is a weighted sum of the input voxels. The weights are derived from the local environments of the input voxels.
- ***Maximum***: The output value of each voxel is the maximum of the input values, or by calculating a local quality criterion and selecting the value of the voxel with the highest quality value.

To determine how the inputs are weighted (Weight mode), or which input is selected (Maximum mode), a local information criterion is used: for each input voxel a value is derived from the values of the voxel and its immediate neighbors using one of the three following options, that can be selected in this FUSION SETTINGS menu:

- ***Pixel values***: The average of the inputs is used (Weight mode), or the input with the maximum value is used (Maximum mode).
- ***Variance***: The input values are weighted by the variance of the voxel values in the immediate neighborhood (Weight mode), or the input with the highest local variance is selected (Maximum mode).
- ***Local entropy***: As in the Variance option, but instead of the local variance of the input voxels, the local entropy is used.

Selecting input images

The input images for the fusion process can be added by pressing either the ADD INPUTS button, or the ADD FILES button. The ADD INPUTS 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 FILES button can be used to select images that reside on disk. Next to the ADD FILES button, a LOAD ON DEMAND check button can be used to indicate that these files should be loaded only 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 & 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 10.1 on page 57).

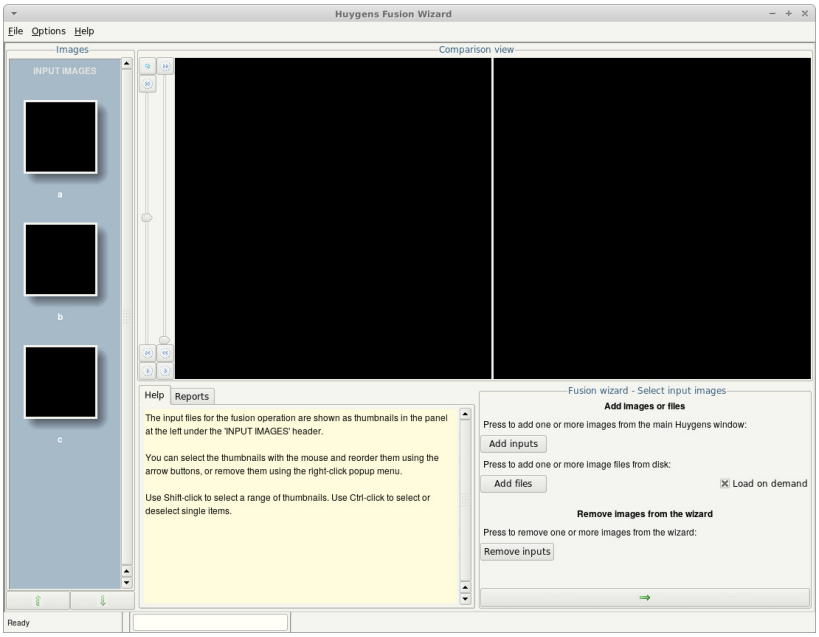

FIGURE 10.1. The *Fusion & Deconvolution Wizard* main window after having selected thumbnails a, b, and c with the ADD INPUTS button.

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 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 red-colored icon 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 Microscopic Parameter Editor ("Verifying Microscopy Parameters" on page 17).

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 slightly shifted.
- ***Rotation***: The images are rotated around one axis and aligned by shifting in the X, Y and Z directions.
- ***None***: No rotation or shift transformations are applied. This option should be selected if the position of the structure is similar in both images.

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 & Deconvolution Wizard* offers three options to equalize the sampling:

- ***Do not equalize***: The sampling of the input is not changed.
- ***Adjust Z sampling to match XY***: 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.
- ***Adjust the XY sampling to match Z***: The 3D image is resampled in XY to make the sampling equal to the sampling along Z. This usually reduces the sampling 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 multiplication factor. This will only affect the sampling rate in Z in the result image, not its final size. The precise value for this calibration factor needs to be estimated empirically.

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.

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 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 THE3D 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 & 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 & Deconvolution Wizard* will allow the user to select if each time-frame should be aligned separately, or if the alignment parameters can be derived from a single time-frame. By default the *Fusion & 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 fusion

After the parameters have been set, the fusion process is started by pressing START 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 11

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 slicer, click the image name in the drop-down menu above the left or right view port (See Figure 11.1).

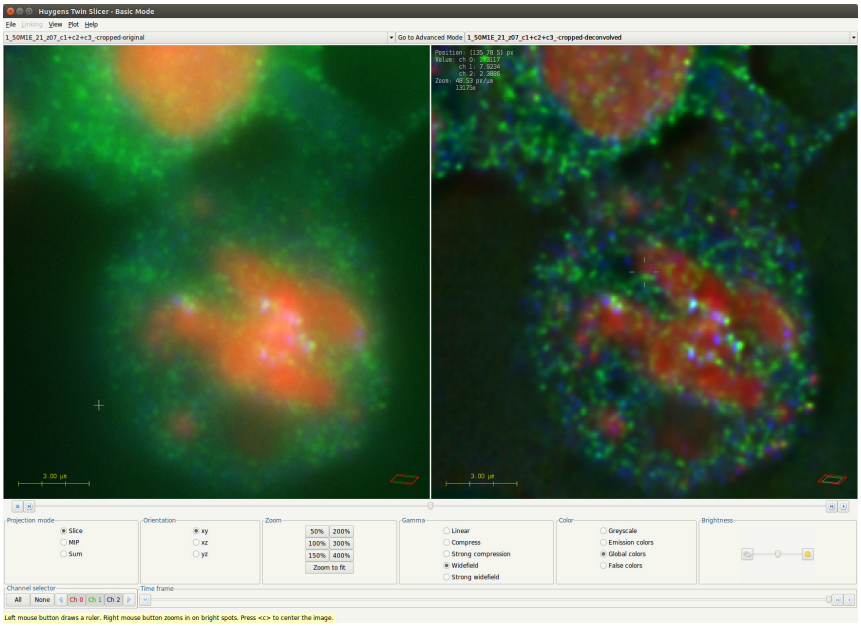

FIGURE 11.1. The Twin Slicer in basic mode, showing an original and deconvolved image side-by-side.

The View Menu

Use the VIEW menu to show or hide image properties and guides. These are listed in Table 11.1:

TABLE 11.1. The options in the Twin Slicer's VIEW menu.

Option	Description
POINTER COORDINATES	Display the position of the mouse pointer in µm or in voxel coordinates.
TIME	Display the time for the current slice in seconds or frame numbers.
INTENSITY VALUES	Display the intensity values for all channels on the current pointer location.
ZOOM	Display the zoom value in screen-pixels per micron. A magnification factor is displayed as well; using the pixel density for the monitor, this value gives an estimation for the absolute magnification.
ROTATION ANGLES	Display the tilt and twist angles in degrees.
DROP SHADOWS	Enhance the contrast for the overlayed 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.

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 *basic* and *advanced* mode. In basic mode, all controls are visible in the panels below the view ports (See Figure 11.1). In contrast to the advanced mode, which allows independent control

of the left and right slicer (See "Using the Slicer in Advanced Mode" on page 64), the basic mode shows a single set of controls that apply to both slicers.

Changing Time Frames

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

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: TOOLS→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 63).

- *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 ON 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.

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 63), 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 63 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) distribution[1]. 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*).

1.https://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 11.2. Note that settings get synchronized once the controls are being used.

TABLE 11.2. The options in the Twin Slicer's LINKING menu (accessible in advanced mode).

Option	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.

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 11.2).

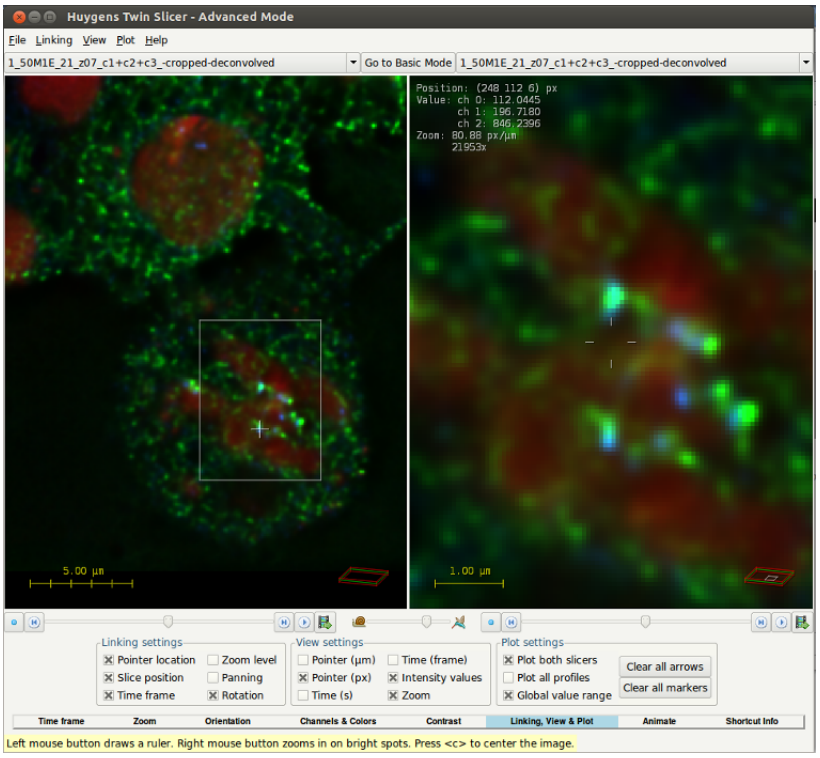

FIGURE 11.2. The Twin Slicer in *advanced mode*, with all controls but zoom and panning linked.

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 11.3 and our online SVI Wiki[2].

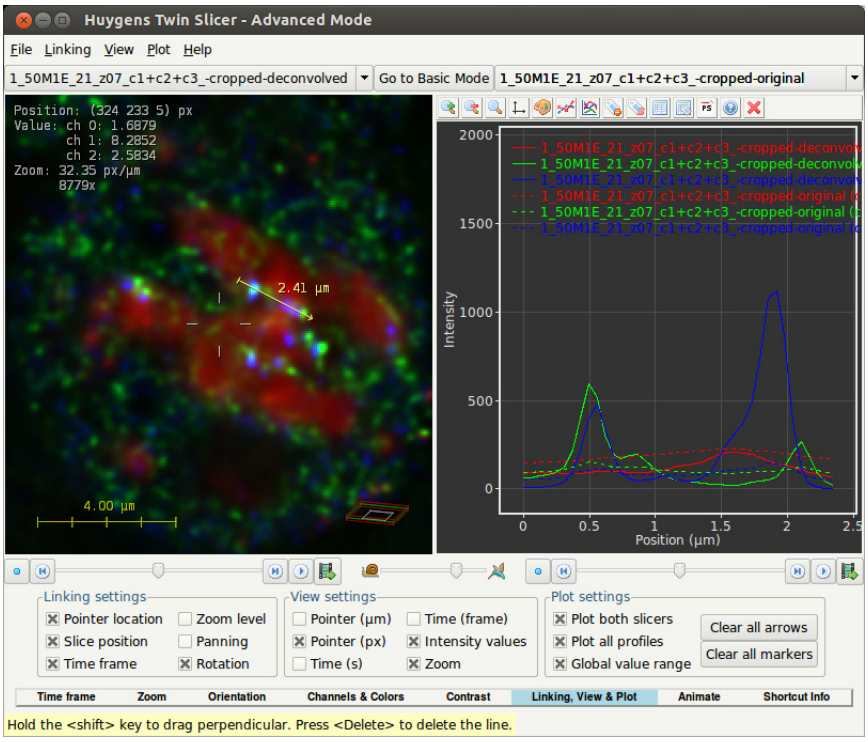

FIGURE 11.3. Measuring the intensity profile along a line. The plot can be configured such that it shows the profile of both images (left solid, right dashed).

2. http://www.svi.nl/DataPlotter

CHAPTER 12

The Orthogonal Slicer

The Huygens Orthogonal Slicer tool is shown in Figure 12.1. It displays the image that is selected from the main window. The Ortho Slicer is designed to show the same point in 3D space from three orthogonal directions;

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

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

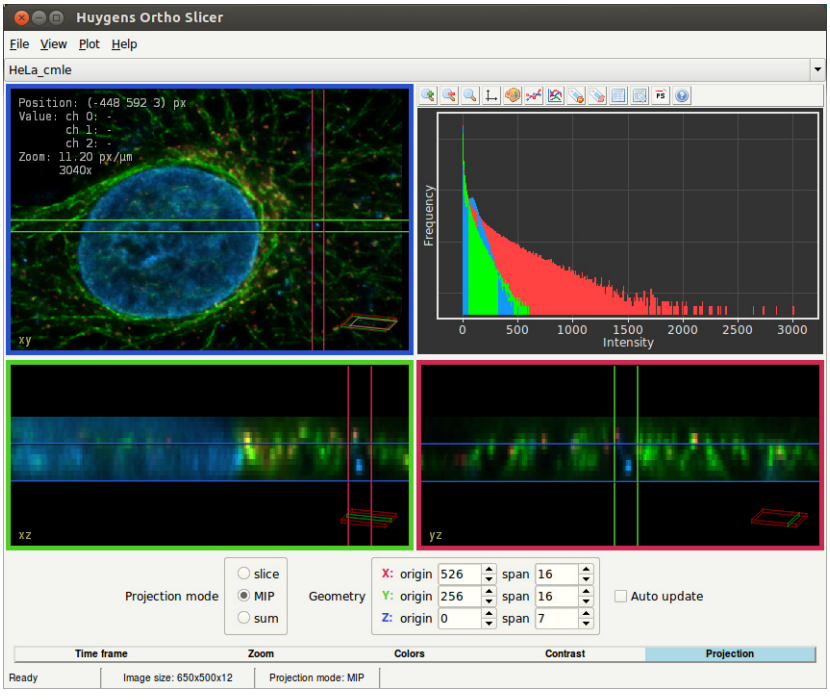

FIGURE 12.1. The Huygens Orthogonal Slicer.

If you move one of the slices, the others will follow to make sure that the center of each of the slices intersects in the same point in space. This 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 orthoslicer can be found on the wiki[1].

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 Huygens Twin Slicer (on page 61). There are tools to;

- change time frames
- 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'.

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 capabilities.[2]

1. http://www.svi.nl/OrthoSlicerIntro
2. http://www.svi.nl/DataPlotter

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 wireframe.

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.

CHAPTER 13

The MIP Renderer

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

The MIP renderer projects the image voxels on the screen by tracing near-parallel rays from a viewpoint far away through the data volume. Along each ray the maximum intensity encountered is taken for the rendered image (See Figure 13.1). Notice that this implies that two MIP renderings from opposite viewpoints show symmetrical images.

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

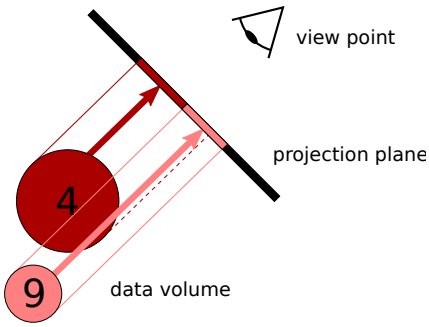

FIGURE 13.1. A schematic overview of MIP rendering. The maximum intensities on rays parallel to the screen are projected.

Basic Usage

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders (See Figure 13.2), or by 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.

Notice that the scene is rendered with low quality while being edited, which allows for fast scene edition. The scene is automatically rendered with good quality when the edition is finished and the scene released. This automatic rendering can be deactivated by pressing the SWITCH OFF rendering button ( ). However, when either the FAST MODE render button ( ) or the HIGH QUALITY render button ( ) are pressed the automatic rendering occurs again upon releasing the MIP scene.

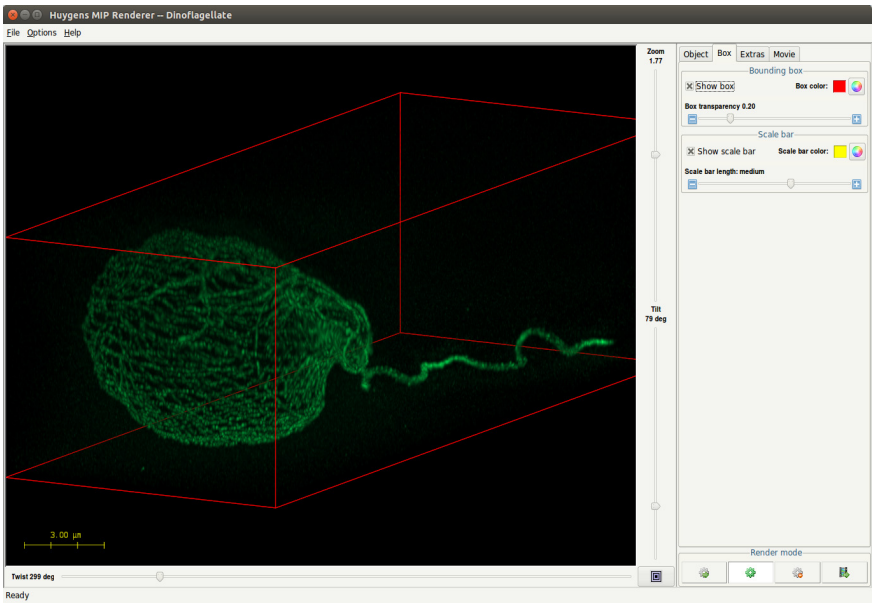

FIGURE 13.2. The MIP Renderer window showing a deconvolved STED 3X images of 30x11x10 micron (xyz) of a *Dinoflagellate*.

Tabs for changing image display and movie settings

- Object: This tab can be used to adjust the display of the different channels. With check boxes in front of the listed channels, you can select or unselect channels to be included for display. The *Soft threshold* slider in the *Channel parameters* panel at the right affects the threshold level of the channel that is selected. The application of a threshold is a preprocessing step that *reduces the background* in the image, i.e., voxels with intensity values below the threshold value become transparent. Contrary to a standard threshold, which is 'all or nothing' (values above the threshold are kept, values below it are deleted), the soft threshold function handles images in a different way. It makes a *smooth transition* between the original and the deleted value. The mode of the colors of the channels can be changed whereby Global colors are defined within the TOOLS->PREFERENCES window.
- Box: The bounding box and scale bar can be switched on/off, and their colors can be changed here. The transparancy of the box can be adjusted with a slider if it interferes with the image display, and the scale bar size can be adjusted with another slider.
- Extras: The virtual rendering size of the image can be changed and customized by selecting an option under this tab.
- Movie: For details see "Simple Animations" on page 75

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.

Options menu

Table 13.1 gives an overview of the render options available through the OPTIONS menu.

TABLE 13.1. Render options for the MIP Renderer.

Option	Description
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the *right mouse button* to pick up and move the scene canvas.
COLOR MODE	Choose between GREY, EMISSION COLORS, GLOBAL PALETTE (See "Adjusting the Global Color Scheme" on page 199) or FALSE COLOR.
SHOW SVI LOGO	Hide or show the SVI logo at the bottom right.

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

Without the Movie Maker the MIP Renderer has the option to make simple animations of the image, changing the view point in different movie frames. Set the render parameters for the first frame and click MOVIE tab, SET →FIRST SCENE. Now adjust the viewpoint for the final frame, and click SET→LAST SCENE. Set also the *frame count* and *frame rate*, and adjust any other general options such as colors using the additional tabs. Finally press the *animate* button ( ), and select a directory to save the AVI movie or the TIFF frames to.

The exported AVI files use the MJPEG[1] codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime. The TIFF frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

1. `http://en.wikipedia.org/wiki/Mjpeg`

CHAPTER 14

The SFP Renderer

The SFP Renderer generates *realistic 3D scenes*, based on the 3D microscopy image, which is taken as a distribution of fluorescent material. The computational work is done by the *Simulated Fluorescence Process* (SFP) algorithm[1], simulating what happens if the material is excited and how the subsequently emitted light travels to the observer (See Figure 14.1).

The properties of this algorithm allow to render the object at different depths, unveiling layers under the object's surface. The SFP algorithm is not limited by boundaries or sharp gradients and is exclusively suited to render 3D microscopy data. Since the SFP algorithm is based on *ray-tracing* techniques it does not require special graphic cards.

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

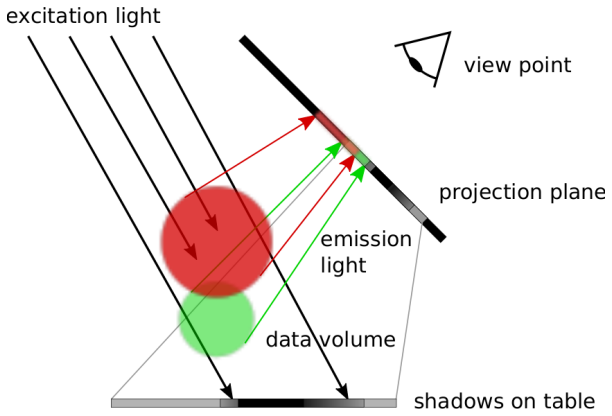

FIGURE 14.1. In the SFP renderer excitation and subsequent emission of light of fluorescent materials is simulated. Each subsequent voxel in the light beam is affected by shadowing from its predecessors. The transparency of the object for the emission light controls to what extent the viewer can peer inside the object.

1. http://www.svi.nl/SFP

Orientation and Zoom

Adjust the viewpoint by moving the *Tilt* and *Twist* sliders or by clicking and dragging the mouse pointer across the scene. The magnification can be adjusted by using the *Zoom* slider or the scroll wheel. See Figure 14.2

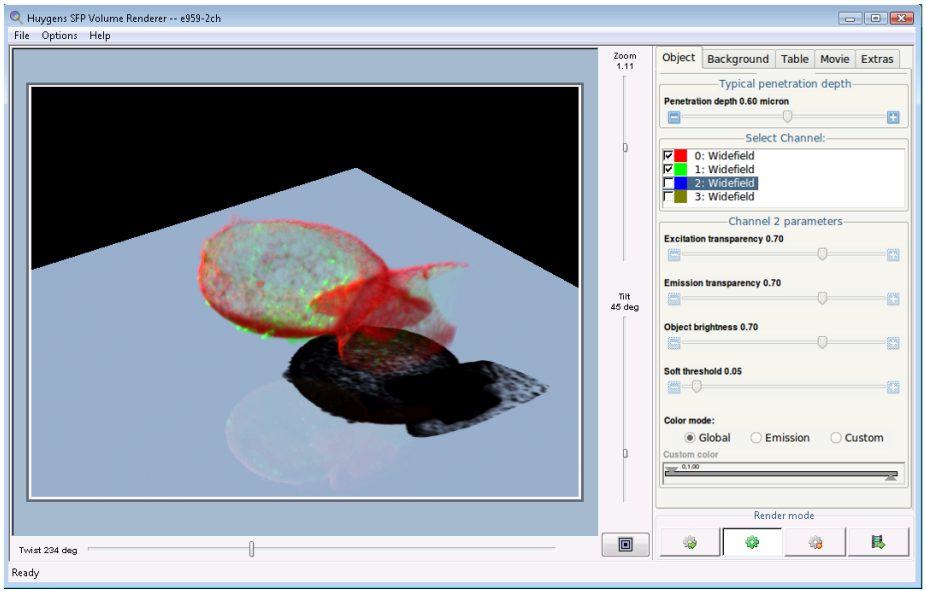

FIGURE 14.2. The SFP Renderer window.

Render quality

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

- Low Quality. This level is used when the scene is being moved or edited in any other way (zoomed in and out, adjusted with new parameter values, etc). This quality level is also used after editing the image if the Fast Mode and High Quality levels are switched off.
- Fast Mode. This quality level improves the scene greatly with respect to the Low Quality level. The Fast Mode level is used automatically when the renderer detects that the user stopped editing the scene (rotating, zooming, adjusting parameter values, etc) provided that the FAST MODE render button ( ) is activated.
- High Quality. This mode can deliver even better quality than the Fast Mode although it can take more time to render the scene. The High Quality level is used automatically when the renderer detects that the user stopped editing the scene (moving, zooming, adapting parameters) provided that the HIGH QUALITY render button ( ) is activated.

When you edit the scene, it is rendered at low quality to allow for fast scene edition. The scene is automatically rendered with good quality (Fast Mode or High Quality) when the edition is finished and the scene released. 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 rendering occurs again upon releasing the SFP scene.

When using poorly sampled images the FAST MODE and HIGH QUALITY modes allow for more limited zooming than the low quality mode. In that case a subtle zoom jump might be noticeable when switching between the low quality and the better quality modes.

Select Channel

The *Select Channel* panel shows for each channel what color (indicated by the colored squares) is used for its display in the scene, and what microscope type is defined (within the microscopy parameters). If needed, the channel colors can be changed under *Color Mode*. When clicking on one of the channels in this panel, you select it for adjusting the channel parameters for rendering. The checkboxes can be used to select and un-select channels for display.

Soft threshold

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

Penetration depth

The characteristic object size can be set by the *Penetration depth* slider at the right side. This parameter affects both the excitation and the emission transparency. While traveling through the object, the light intensity is attenuated to some degree. This enables us to define some definition for penetration depth at which the light intensity is decreased to some extent, for instance 10% of its initial value. This penetration depth should be in line with the object size. A transparent object is small with respect to the penetration depth. Thus for the same physical properties of the light one object can be transparent while the other is opaque due to its size. To find a reasonable range in transparencies the Penetration depth may be altered. The initial object size is computed from the microscopic sampling sizes and number of pixels the image is composed of. If the microscopic sampling sizes of the image are incorrect, then the penetration depth is set according to some default parameters.

Excitation and Emission transparency

The characteristic wavelength-dependent absorption properties of the object can be set for each channel by adjusting the transparencies for the excitation and emission wavelengths individually. For the same physical properties (such as size) an object can be transparent at one wavelength and opaque at another due to the different interaction of light with matter.

Saving Scenes

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

SFP Fundamentals

The voxel values in the image are taken as the *density of a fluorescent material*. In case of a multi channel image, each channel is handled as a different fluorescent dye. Each dye has its specific excitation and emission wavelength with corresponding distinct absorption properties. The absorption properties can be controlled by the user (See the *transparencies* in Table 14.1 on page 80). The different emission wavelengths give each dye its specific color.

To excite the fluorescent matter light must traverse other matter. The resulting attenuation of the excitation light will cause objects, which are hidden from the light source by other objects, to be weakly illuminated, if at all. The attenuation of the excitation light will be visible as shadows on other objects. To optimally use the depth perception cues generated by these shadows, a flat *table* below the data volume is placed on which the cast shadows become clearly visible. In Figure 14.2 the table is rendered as a mirror.

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

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

Render Parameters

Table 14.1 gives an overview of all render parameters in the SFP Renderer.

TABLE 14.1. SFP render parameters

Parameter	Description
Time frame	A selection of the time frame for time series.
Penetration depth	Adjusts the total transparency of the rendered object.
Select Channel	Select channels for display and for adjusting parameters.
Excitation transparency	Adjust excitation transparency for the matter in selected channel.
Emission transparency	Adjust emission transparency for the matter in selected channel.
Object brightness	Set the intensity level for the excitation light source for the selected channel.
Soft threshold	Adjust the threshold level for the selected channel.
Color mode	Adjust the color of the selected channel.
Background color	Adjust the color of the scene background.
Compute shadow	Choose whether or not the object's shadow should be included.
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 degree of reflection of the table.
Table size	Adjust the size of the underlying table.

TABLE 14.1. SFP render parameters

Parameter	Description
Table color	Adjust the color of the table
First scene	Set the viewpoint of the first scene of the movie.
Last scene	Set the viewpoint of 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.
Light direction	Set the position of the light source.
Light intensity	Set the intensity of the light source.
Render size	Set the size (in pixels) of the SFP scene.
Render mode	Set whether the scene must be rendered in fast mode, in high quality mode, rendered in a movie or not rendered.

Render Options

Table 14.2 gives an overview of the different render options that are available through the OPTIONS menu.

TABLE 14.2. Render options for the SFP Renderer.

Option	Description
SHOW SVI LOGO	Show or hide the SVI logo at the bottom right of the scene.
COLOR MODEV	Choose between EMISSION COLORS or GLOBAL PALETTE (See "Adjusting the Global Color Scheme" on page 199.).
VIRTUAL RENDER SIZE	Adjust the size of the rendered image. When the render size exceeds the display area, then use the *middle mouse button* to drag and drop the rendered image.

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 images 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 89.) allows to easily create sophisticated animations using the MIP, SFP, and Surface Renderer.

SFP Renderer has its own option to make simple animations. Set the render parameters for the first frame by clicking the MOVIE tab, SET →FIRST SCENE. Now adjust the viewpoint for the final frame, and click SET→LAST SCENE. Set also the *frame count*, *frame rate*, and other general options such as colors and light properties. Finally press the *animate* button ( ), and select a directory to save the AVI movie to.The exported AVI files use the MJPEG[2] codec and can be loaded in most movie players, including Windows Movie Player and Apple Quicktime.

2. `http://en.wikipedia.org/wiki/Mjpeg`

CHAPTER 15

The Surface Renderer

Huygens Surface Renderer is a powerful 3D visualization tool that enables the visualization of iso-surfaces of volumes.

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. 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.

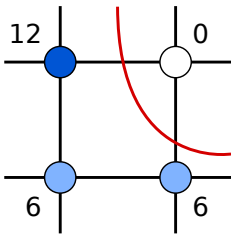

FIGURE 15.1. A contour line for an interpolated value of 5. Because 5 is much closer to 6 than to 12, the distance of the contour to the voxel with value 12 is larger than the distance to the bottom-right voxel with value 6.

The iso-surface is a smooth surface through all points between the voxel locations which have an interpolated value equal to the threshold value (see Figure 15.1).

Shading enhances the perception of 3D shapes and texture (See Figure 15.2). Besides iso-intensity surfaces, this renderer is able to generate MIP projections which are blended with the surfaces to be used as a reference to the original microscopic data (See Chapter 13 on page 73).

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

To start the Surface Renderer, right-click on a thumbnail and select VIEW→SURFACE RENDERER from the pop-up

menu. Alternatively, choose the corresponding icon from the taskbar or select the VISUALIZATION→SURFACE RENDERER from the main menu.

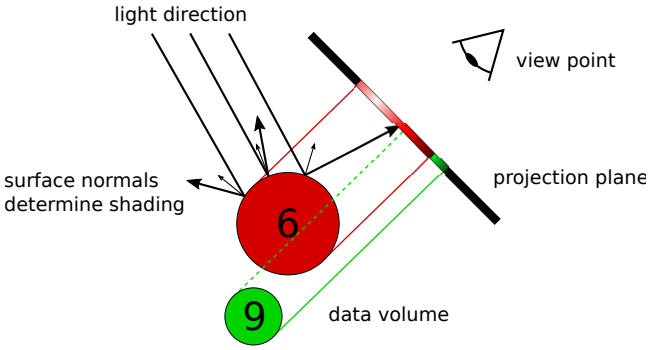

FIGURE 15.2. A schematic overview of surface rendering. The orientation of the surface determines the shading.

Basic Usage

Orientation and Zoom

Within the Surface Renderer window, you can adjust the viewpoint by moving the *Tilt* and *Twist* sliders or by dragging the mouse pointer on the large view (See Figure 15.3). The magnification can be adjusted using the *Zoom* slider or by using the scroll wheel. Use the middle mouse button to *pan* the center of the projection.

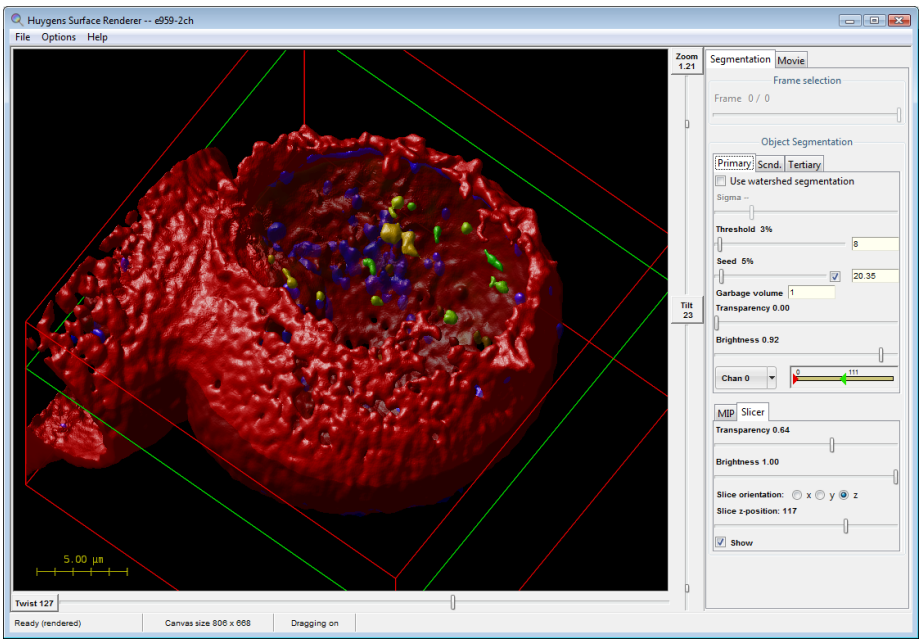

FIGURE 15.3. The Surface Renderer window.

Threshold

Use the *Threshold* slider in the *Render pipes* 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 us to define three threshold levels that can be applied to the same or to different data channels. The data channel can be selected using the menu button in the *Object Segmentation* panel. The color range in which the different objects inside a render pipe will be displayed can be adjusted with the hue selector next to it. To facilitate the segmentation of the objects, a watershedding can be applied. For more information see "Using the watershed segmentation" on page 139.

Saving Scenes

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

Advanced Usage

Adding a Maximum Intensity Projection

Besides the surface pipes there are additional rendering pipes to redirect data to the scene. The *MIP* pipe works by projecting the voxels with maximum intensity laying in the path of the rays traced along the viewing direction (See Chapter 13 on page 73). 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 used as a spatial reference for the objects present in other channels.

Adding a Slice

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

Render Parameters

Table 15.1 gives an overview of all render parameters in the Surface Renderer.

TABLE 15.1. Surface render parameters

Parameter	Description
Frame	Set the time frame (in case of a time series).
Threshold (surface)	Adjust the threshold level for the selected pipe, i.e. the intensity for which the iso-intensity surfaces are defined. See "Threshold" on page 85
Threshold (MIP)	Set the soft threshold level for the MIP pipe.
Seed	Only objects with an intensity higher than the seed level remain, while the rest are discarded
Garbage volume	Volumes that contain less voxels than defined by the *garbage volume* parameter will not be rendered. This is useful for rendering only significant objects in noisy images.
Transparency	Set the level of transparency to other pipes.

TABLE 15.1. Surface render parameters

Parameter	Description
Brightness	Adjust the brightness for the selected pipe.
Slice Z-position	Set the position of the slice in the *Slicer* pipe.

Render Options

Table 15.2 gives an overview of the different render options that are available through the OPTIONS menu. The ANIMATION FRAME COUNT, ANIMATION FRAME RATE and RENDER QUALITY apply to the rendering of simple movies as explained in the next section.

TABLE 15.2. Render options for the Surface Renderer.

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

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 Surface Renderer has the option to make simple animations of the image. When selecting the movie-tab, one can set the render parameters for the first scene and click the FIRST SCENE button in the *Position settings* panel of the *Set view* settings. Now adjust the viewpoint for the last scene, and click the LAST SCENE button of the *Set view* settings. Also the *frame count*, *frame rate*, or other render options in the OPTIONS menu can be adjusted.

If you want to view or adjust the first or last scene, press the corresponding button next to *Show*. This will show you the stored scenes. You can adjust these settings and again store the scenes via the *Set view* buttons.

Finally press the *Animate* button in the *Actions* pane, and select a directory for storing the AVI movie or the Tiff frames.

The exported AVI files use the MJPEG[1] codec and can be played in most movie players, including Windows Movie Player and Apple Quicktime. Tiff frames are useful to combine multiple animations or edit the movie in e.g. Windows Movie Maker.

For more sophisticated movie making options, and for creating animations using combinations of MIP, SFP, and Surface rendered images, we like to refer you to the Huygens *Movie Maker* (See "The Movie Maker" on page 89.)

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

CHAPTER 16

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 13 on page 73), the *SFP Renderer* (See Chapter 14 on page 77), and the *Surface Renderer* (See Chapter 15 on page 83) 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 e *Timeline* or by using the interfaces of the renderers.

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

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

An Overview

Figure 16.1 shows the Movie Maker's user interface. The numbered areas are:

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

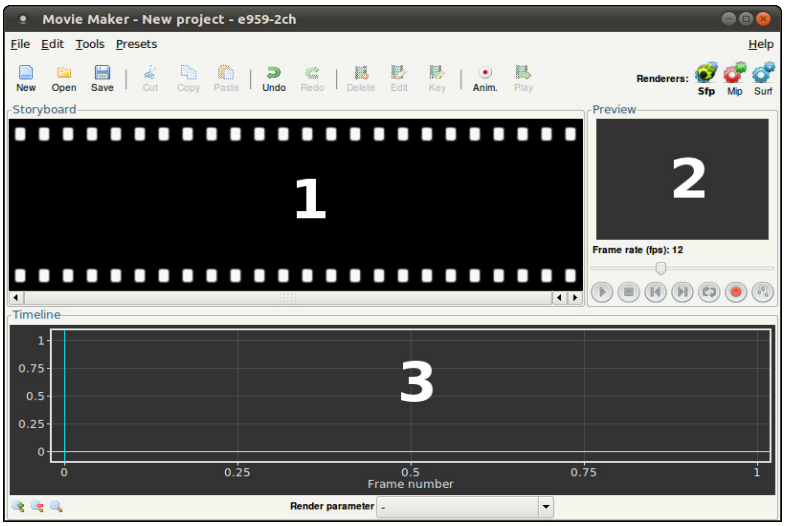

FIGURE 16.1. An overview of the Movie Maker user interface, showing the *Storyboard* (1), the *Preview* (2), and the *Timeline.* (3).

Creating and Adjusting Keyframes

About Keyframes

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

Inserting New Keyframes

To add the first keyframe to the storyboard, one of the renderers should be launched from the Movie Maker window by pressing the corresponding button (, , or). Within the renderer, a scene can be defined using the available controls; see Chapter 13, Chapter 14, and Chapter 15 for more information. If the view of the first frame is satisfactory, press the *Add keyframe* button ( ) 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 ( ) again. The storyboard will now show two keyframes with an arrow in between. This arrow indicates the transition from one keyframe to another.

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 16.2.

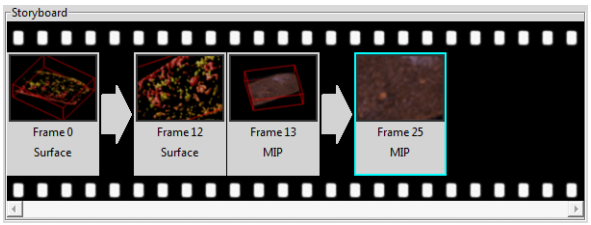

FIGURE 16.2. The Movie Maker's storyboard showing two surface-renderered keyframes, followed by two MIP renderer keyframes.

Editing Keyframes

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

Using the Storyboard

Rearranging Keyframes

The storyboard allows the user to copy, delete, and rearrange keyframes. The *Cut* ( ), *Copy* ( ), *Delete* ( ), and *Edit* ( ) buttons are activated whenever a keyframe is selected. 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 16.3). 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.

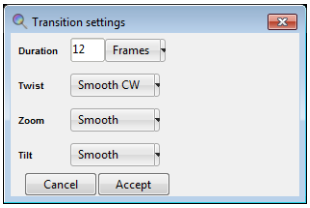

FIGURE 16.3. The transition settings dialog is shown by double clicking a transition arrow in the storyboard.

Playing a Preview Movie

To preview the movie, just press the *play* button ( ) in the *Preview* area (See Figure 16.4). 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.

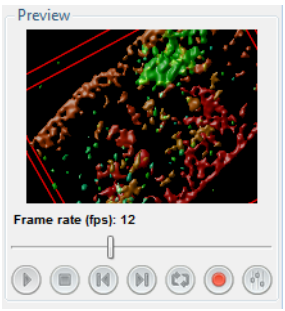

FIGURE 16.4. The Preview area.

Export to AVI or Tiff Series

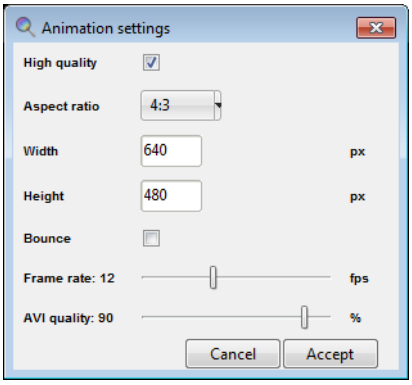

FIGURE 16.5. The animation settings dialog.

Press the record button ( ) (See Figure 16.4) 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 16.5), 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*" ( ).

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 an identical copy. 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 16.6). 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.

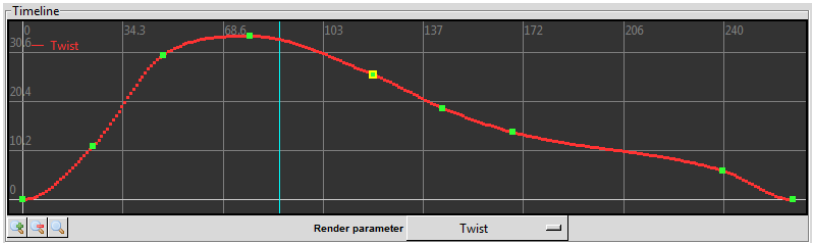

FIGURE 16.6. The *Timeline* area gives detailed visual feedback on the animated parameters.

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 16.6). 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 16.7). 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.

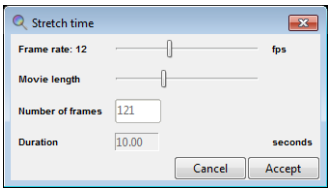

FIGURE 16.7. The stretch dialog helps increasing or decreasing the number of frames in a complex movie.

Synchronizing Transitions in Time and Slice Plane Transitions

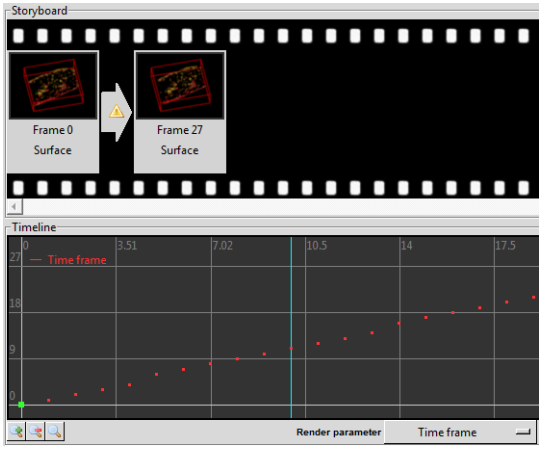

FIGURE 16.8. A transition in which the *time frame* parameter is out of sync with the number of frames. In this example, the transition counts 27 frames, while the *time frame* parameter increases linearly from 0 to 30. Note the irregular change of this parameter in the plot.

Most render parameters, like the *Tilt*, *Twist*, and *Zoom*, can be set to non-integer values. The *Time frame* (in case of a time series), and the *Slicer z-position*, however, are fixed to integer values. When the transition of such a parameter is not linear, or when the change in value does not match the number of frames, then this parameter is *out of sync*. In that case the Movie Maker will show a warning symbol (⚠) on the transition arrow (See Figure 16.8). In the final result, these asynchronous transitions may show irregularities.

To restore the synchronization, right click on the transition arrow and choose SYNCHRONIZE TIME FRAMES... from the pop-up menu. The Movie Maker will set the transition type to linear and add or remove some frames from the transition to achieve a 1:1, 2:1, 1:2, etc. synchronization.

Creating Loopable and Bouncing Movies

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

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

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

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

About Movie Quality

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

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

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

CHAPTER 17

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 17.1). As a user, you can control which z-slices or time frames are displayed

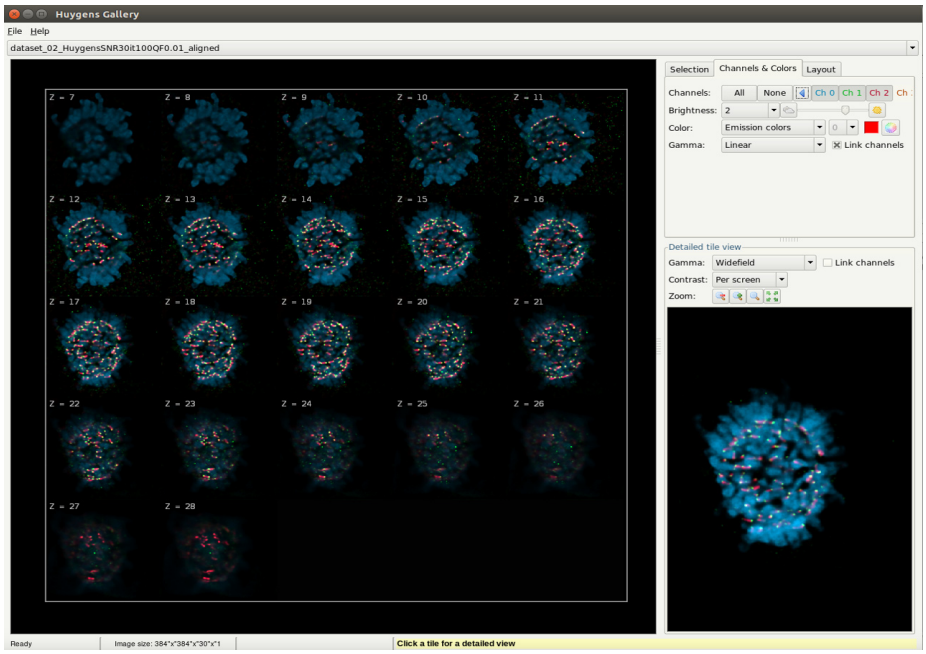

FIGURE 17.1. The Gallery tool showing the individual slices of a Z stack image.

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 slice 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 slice in this window. For more details on the use of these controls see "Using the Slicer in Advanced Mode" on page 64.

CHAPTER 18

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.

Improving such stitched images with deconvolution is a challenging task as extensive computer resources for allocating both raw and processed images are 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 *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 image need to be loaded.

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

- optimal stitching of tiles.
- automatic 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 Leica LIF or Zeiss CZI 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 MULTITILE IMAGE – OPEN button in the *Welcome* screen, to select tiles that should be loaded into the *Stitching & Deconvolution Wizard*. This opens the window *Select sub-images* (See Figure 18.1). Use CTRL-click to select or de-select individual tiles, SHIFT-CLICK to select a range, and CTRL+ A to selects all tiles.

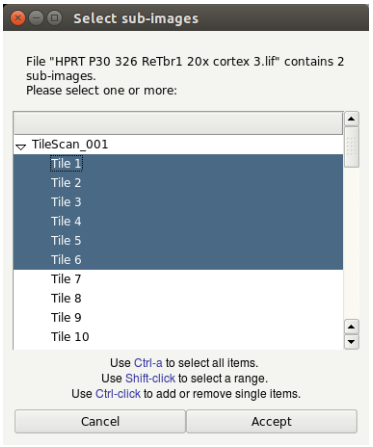

FIGURE 18.1. Selection of tiles contained in a dataset.

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 using custom stitching patterns in the *Stitching & Deconvolution Wizard*. Upon opening the first tile of the 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 198.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 *Stitching & Deconvolution Wizard*.

Pattern selection

If the position of the tiles is known from the metadata, the *Stitching & Deconvolution Wizard* will continue with *Tile selection* (see below). However, if the ordering of the tiles is not specified within the metadata, as is the case with Tiff tile series lacking an XML, a *Stitcher* window will be presented with options to customize the stitching pattern. Parameters such as acquisition pattern, starting point, percentage of overlap between tiles, and number of rows can be defined in at this stage. Also buttons are present to flip tiles in x, in y, and in both directions. By showing the content of the tiles with the button TILES – SHOW CONTENTS, you may find it more easy to select the correct pattern settings, but keep in mind that this will use more computer RAM memory.

Pattern selection

The next window concerns the *Pattern selection* (see Figure 18.2). It shows the tiles' initial positions read from the metadata or as defined within the Pattern Selection window. 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 will allow you to visualize other time frames. Several acquisition pattern parameters can be

selected here if this needs adjusting. The Huygens *Stitching & Deconvolution Wizard* allows a minimum value of "0" percent as acquisition overlap.

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..

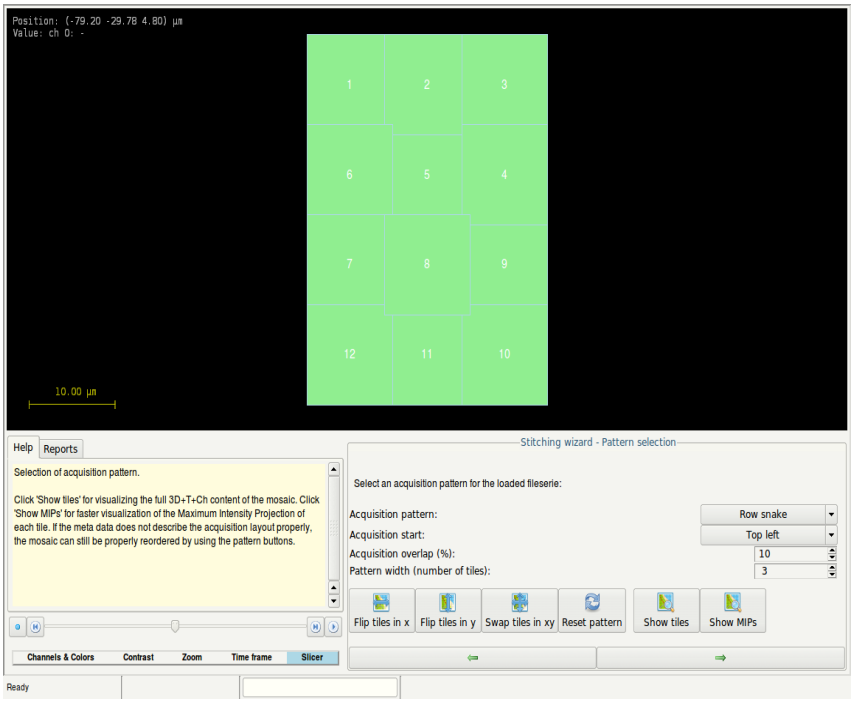

FIGURE 18.2. The *Stitching and Deconvolution Wizard* window showing the tiles' initial positions.

Vignetting and shading correction

Continue and follow the instructions of the wizard to 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 18.3.

Optimization of transitions

Before the actual stitching is executed, The *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 *Stitching & Deconvolution Wizard* offers the user the unique possibility to deconvolve each tile automatically before stitching. Every single tile is first deconvolved before being stiched 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 17, and See "The Final Stage" on page 26.). 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 24.) or EDIT→EDIT MICROSCOPIC PARAMETERS *(*See "Verifying Microscopy Parameters" on page 17.)

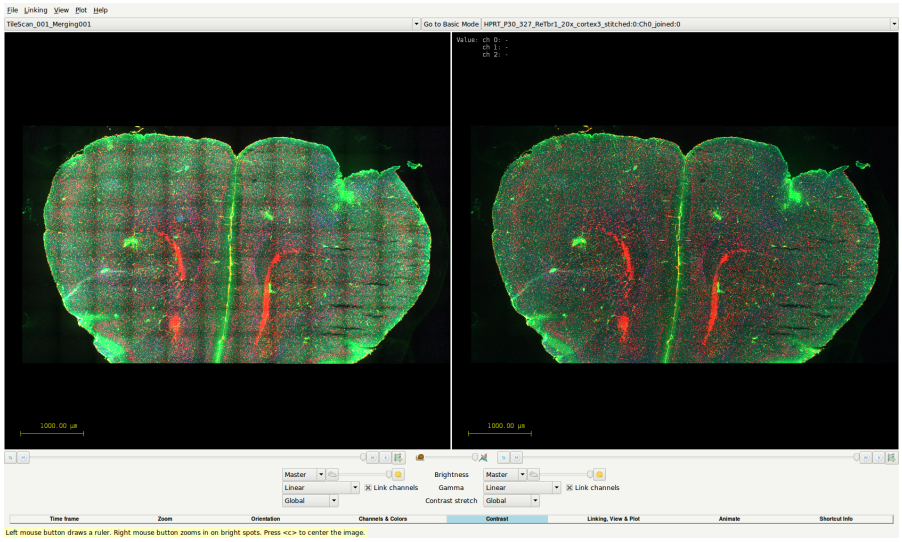

FIGURE 18.3. Huygens stitched tiles of widefield fluorescent Leica LIF data. The same data is shown in the Twin Slicer without (left panel) and with (right panel) Huygens automatic vignetting correction. Note that the overlapping regions are less apparent after vignetting correction. Image represents 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 TILE 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 19

The Hot Pixel Remover

The *Hot Pixel Remover* can detect and correct for hot pixels automatically.

Why correct for hot 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.

The removal of hot pixels is useful for several reasons:

- For deconvolution; hot pixels affect the deconvolution result, because they disturb the normal signal distribution in your image. Especially when those hot pixels are not due to strong noise, but caused by deficiencies of the camera.
- For image restoration; hot pixels may affect the performance of restoration tools such as the Object Stabilizer
- To avoid false colocalization; when hot pixels occur at the same position in each channel, this will affect the colocalization analysis.
- For visualization; hot pixels disturb the contrast range and could force the contrast of the image to be very low. This may result in that objects of interest with lower intensity values become almost invisible.

Hot pixels may also be a product of deconvolution. If a tiny object in your original image fits the PSF perfectly, deconvolution may result in one very bright spot in your image. This can disturb the contrast range and as such 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 Pixel Remover* can detect and correct hot pixels in an image. The tool can be started from the DECONVOLUTION menu in Huygens Essential. TOOLS → "REMOVE HOT PIXELS...".

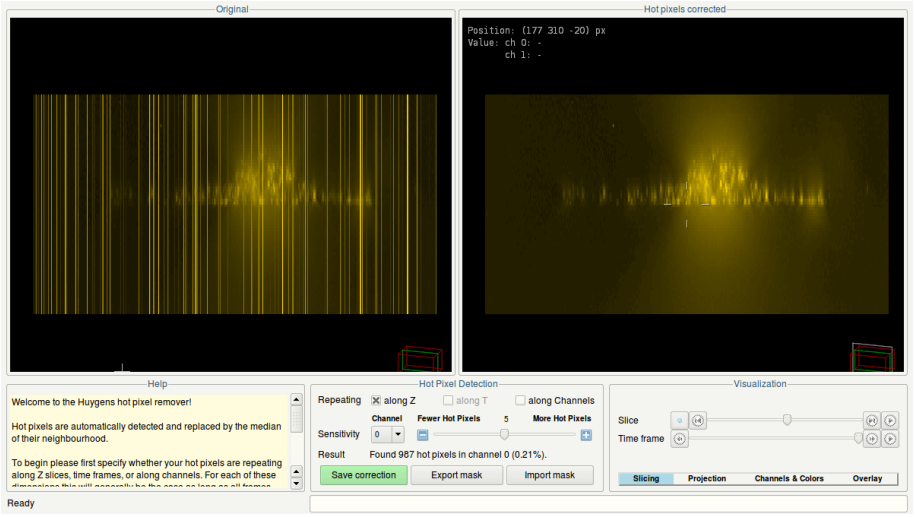

FIGURE 19.1. The start-up window of the Hot Pixel Remover.

The *Hot Pixel Remover* interface shows at start a XZ maximum intensity projection of the original image (top left) and corrected image (top right). Such projections are shown by default as they facilitate the recognition of hot pixels by eye. Note that only the first channel is shown, and the correction needs to be performed on every channel separately. At the very bottom a status bar is shown with a progress bar.

At the left side of the window, a *Help* box explains the general features of the *Hot Pixel Remover*. The *Hot Pixel Detection* box in the middle shows the detection parameters and different correction and mask options (see also below), and a *Visualization* settings box enables you to inspect the image closer with the offered tabs (see Figure 19.1):

- *Slicing*: simply slice through the XYZ stack. The slicer is only enabled when you have selected Slice under the *Projection* tab. The slicer below allows you to slice through time frames..
- *Projection*: here you can switch between Slice, Maximum Intensity Projection (MIP), or Sum projection. The orientation of the views can also be selected here.
- *Channels & Colors*: five different color schemes are available under the first drop down menu (See “Color” on page 75.) 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 show a large variety of intensities.
- *Overlay*: within the display menu you can select between viewing only the image, the hot pixels, or both. The transparency and color of the hot pixels can be adjusted with the slider and color selector, respectively. These tools are helpful in verifying whether the identified hot pixels are correct.

Hot Pixel Detection

The locations that are flagged as hot pixels will depend on the hot 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 pixels that will be corrected, and the percentage with respect to the total number of pixels in a single XY slice. The settings are split up into two groups: The repeating settings and the sensitivity. The

repeating settings allow you to specify whether there is some additional correlation between the hot pixel locations, which the detection algorithm can use. If you are not sure if any correlations exist, you can simply disable all repeating settings. Next we will explain why and when these correlations can occur.

Repeating along Z

If the hot pixels originate from an abnormal charge leakage or defect in the CCD/CMOS, this means that certain pixels in the 2D sensor grid will always produce a hot pixel. Thus, a recorded slice-by-slice (in the Z direction) image will have a hot pixel at these exact locations. This leads to "hot columns" which show up as vertical lines in the XZ MIP projection (like a bar code). Since this is almost always the case, the "along Z" flag is true by default, and false positives will most likely be eliminated.

If you want to remove hot pixels which are a product of deconvolution, all the repeating options should be turned off. Because in that case, you are searching for a single hot pixel which is not conserved in any dimension.

Repeating 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 pixels will show up at the same location in each time frame. In this case the repeating along T option should be checked. Note that if the hot pixels in a sensor change rapidly over time, this will no longer be true.

Repeating along channels

Similarly, if channels of a multi-channel image are recorded one by one, with the same camera, the hot pixels will show up at the same locations in each channel. In such cases, the repeating along Channels option should be checked. Note that if the hot 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 properly specifying the repeating parameters the correction will very likely give an optimal result. However, the sensitivity for detecting hot pixels can be fine-tuned manually for each channel using the sensitivity slider. A channel can be selected with the drop down menu located at the left side of the sensitivity slider, this will also shown in the selected channel in the image windows. With lowering the sensitivity, fewer hot pixels will be found and vice versa. A higher sensitivity setting increases the odds of flagging regular pixels as hot pixels, while a lower sensitivity setting increases the chance of flagging hot pixels as regular pixels. Ideally you want to lower the sensitivity as much as possible and still minimize the number of hot pixels that show up in the corrected image.

Correcting and finalizing

Once you are satisfied with the hot pixel detection, you can correct your image. The hot pixels will be replaced by the median of their 3x3 neighborhood. Simply press CORRECT IMAGE and a new image with suffix 'hpc' (hot pixel corrected) is created in the main window of Huygens Essential, from which it can be saved or edited further.

Using a Mask

Locations of hot pixels can be saved by clicking the EXPORT MASK button. This will save the hot pixel mask as a .h5 file. This mask file can then be imported back into the hot pixel remover for correcting another image, or it can be added to a deconvolution template to correct other images that have identical hot pixel locations.

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.

Batch Hot Pixel correction

Hot pixel correction can be included in the deconvolution template using the "Edit Deconvolution Template" tool (Figure 19.2), which can be found under the EDIT menu in the main wondow. Alternatively, you can use the same template editor in the batch processor. In either case, you must navigate to the pre-processing options ("Pre" tab). Adding a hot pixel removal step simply means importing a previously created mask from file. If the specified mask is valid, it will light up green, and a hot pixel correction step is added to the deconvolution template. If the path is invalid, it will light up red and the hot pixel removal step is skipped. Whenever this occurs you will be notified of the reason for this incompatibility.

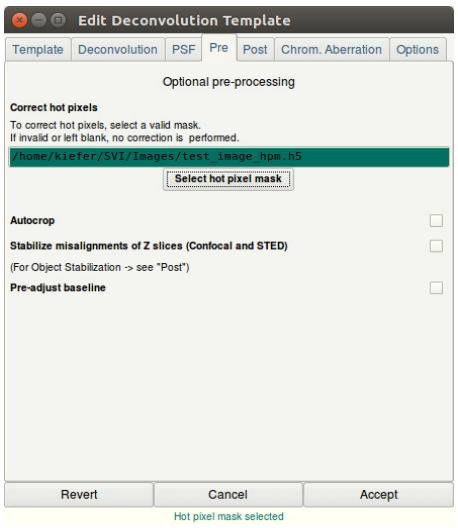

FIGURE 19.2. Tool for editing the deconvolution parameters and including hot pixel correction.

CHAPTER 20

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 (or 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.

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.

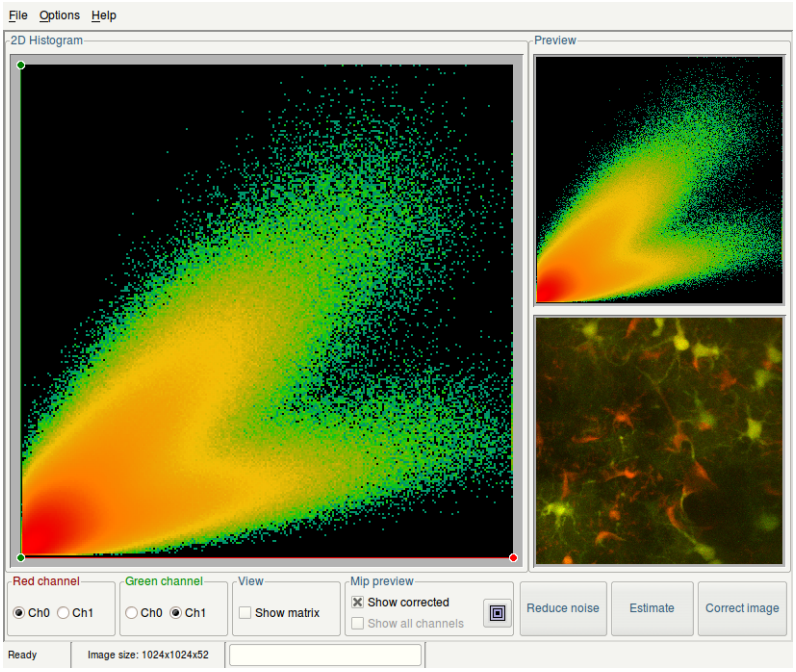

FIGURE 20.1. The startup window of the Crosstalk Corrector with a large 2D histogram of the original data (left) and a preview of the 2D histogram and MIP of the corrected data (right).

The Crosstalk corrector shows a large 2D histogram in the left panel of the startup window, see Figure 20.1. This histogram is from the original raw data set and is useful for the detection and visualization of the crosstalk coefficients.

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 20.2.

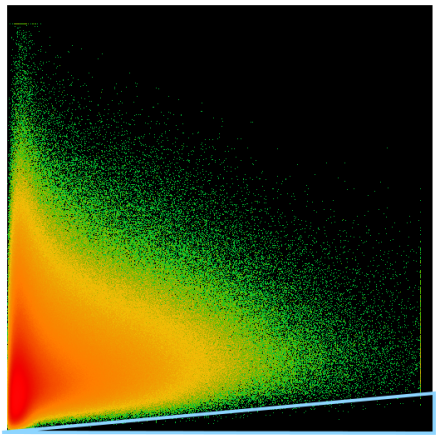

FIGURE 20.2. A 2D histogram, showing the intensity range of the first channel on the x-axis and the intensity of the second channel on the y-axis. The triangle shows that for every intensity in the first channel there is always a fraction of it in the second channel which is typical for crosstalk.

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,

1. Read more about 2D histograms at `http://www.svi.nl/2Dhistogram`

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.

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 grey) 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 button you apply the final crosstalk coefficient matrix to the image and the result image with suffix .ctc, will be shown in the main window of Huygens Essential.

Adjust the crosstalk coefficients

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 21

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 microscopes, bleaching effects can occur when acquiring the image. Bleaching is the fading of the emission intensity with respect to the Z and/or T dimension of the image. This is the result of 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.

In general, bleaching can only occur if the signal in the image is derived from fluorescence and if the image contains two or more time frames. An exception is made for Widefield 3D images which also can exhibit bleaching in the z direction. Here is a summary of the types of images that are supported by Huygens and which can show bleaching effects:

- Widefield 3D
- Widefield time series
- Confocal time series
- STED time series
- Spinning Disc time series

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

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. See FIGURE 21.1 .

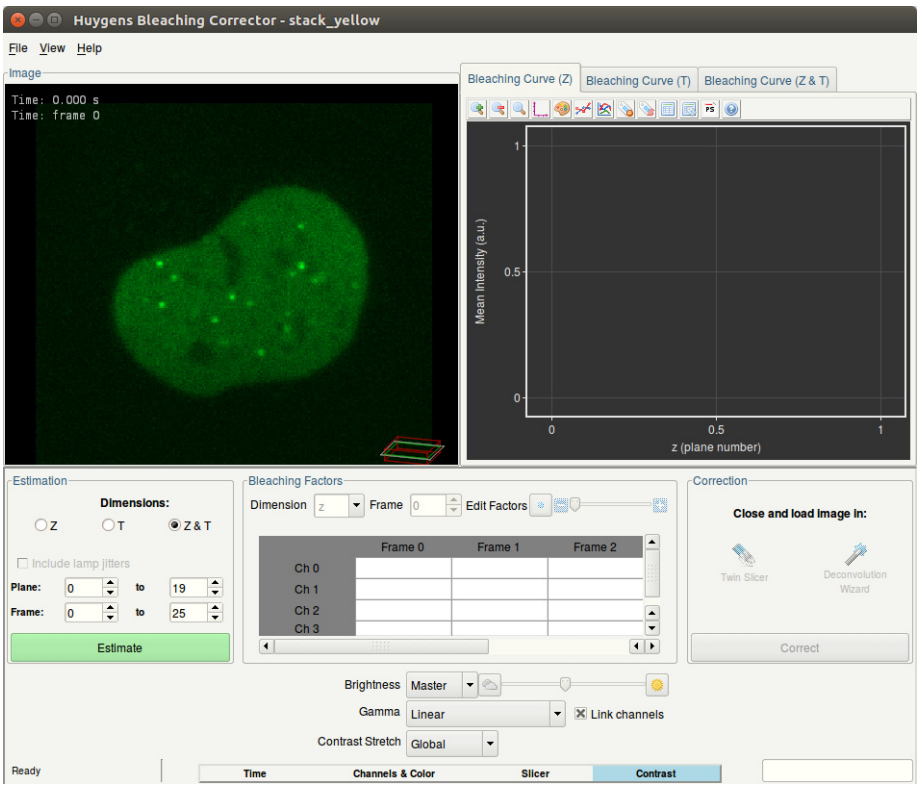

FIGURE 21.1. Startup window of the bleaching Corrector

The main window of the *Bleaching Corrector* tool shows on the left side a single slice of the image up for correction. More information regarding the scene and view options can be selected and unselected from under the View menu in the taskbar at the top of the window. At the bottom of the window, tabs are present which 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 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, contrast 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.

Estimate Bleaching

In the frame *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 Z&T) 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. 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 (straight line) and corrected (dashed line) intensities along time and Z. See Figure 21.2.

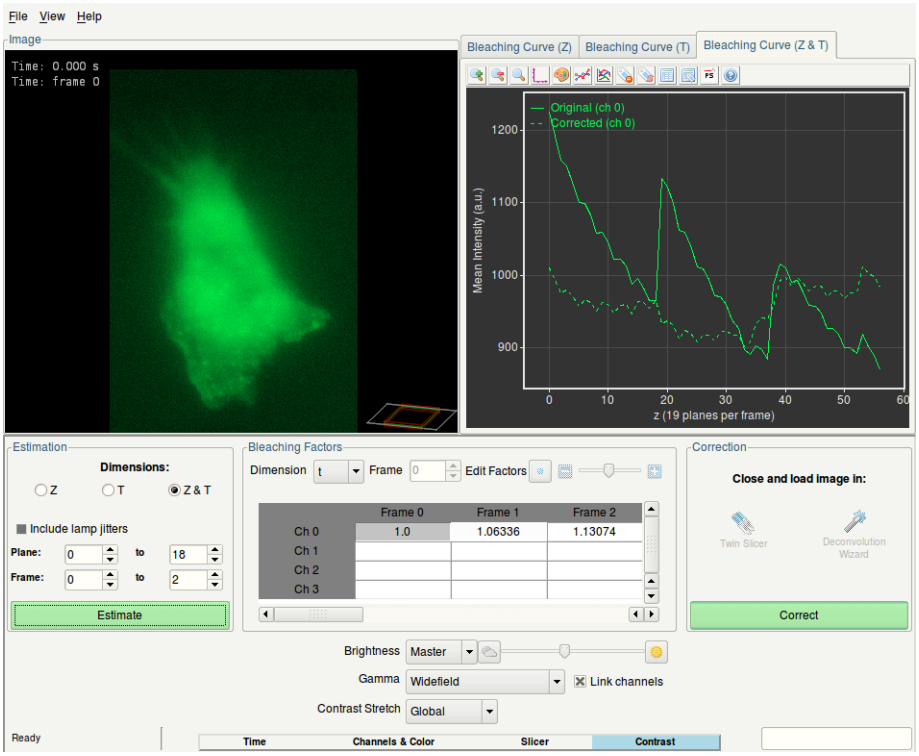

FIGURE 21.2. The window of the *Bleaching Corrector* after estimation of the bleaching.

Adjusting Bleaching Factors

After estimation, the table shown in the middle 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 by pressing the CORRECT button. This will create and export a corrected image to the Huygens main window. After correction, it is possible to continue immediately with the image in the *Twin Slicer* or *Deconvolution Wizard* Chapter 11 on page FirstChapter 4 on page First. You can also decide to directly close the *Bleaching Corrector* window and simultaneously with other Huygens tools. You can either compare the original and corrected image with the Twin Slicer or use the corrected image to start up the Deconvolution Wizard.

CHAPTER 22

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 69.). Below the orthogonal slicer 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 22.1.

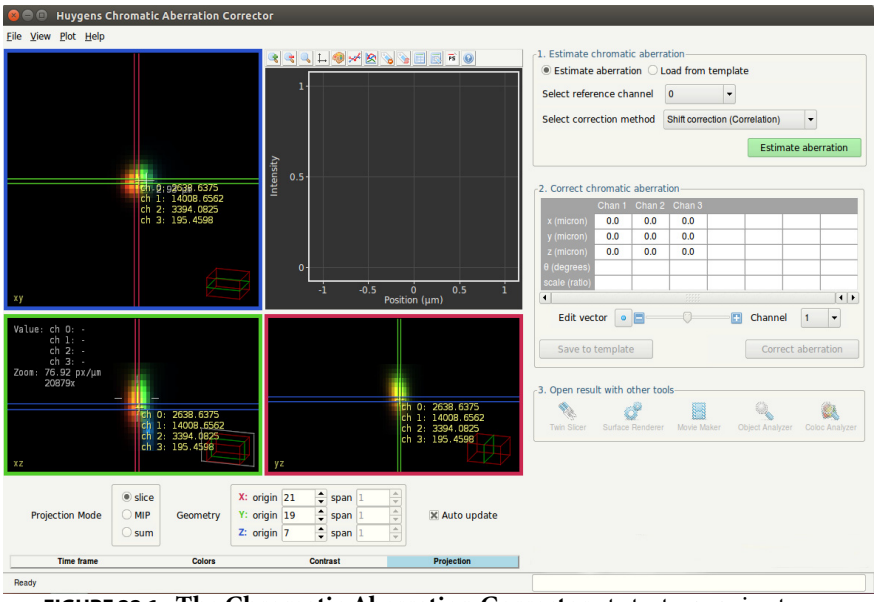

FIGURE 22.1. The Chromatic Aberration Corrector at start-up, prior to estimating the chromatic aberration.

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 22.2. 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.

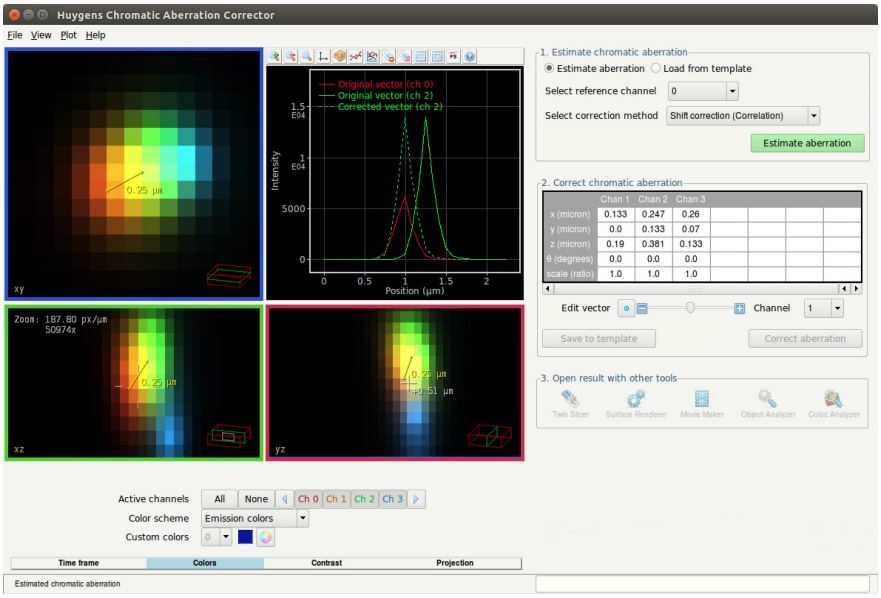

FIGURE 22.2. The Chromatic Aberration Corrector after estimating the shifts. A plot shows the gap between channel 2 and reference channel 0. The estimated shift is drawn as a vector in the orthogonal slicer scenes.

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 22.3. 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.

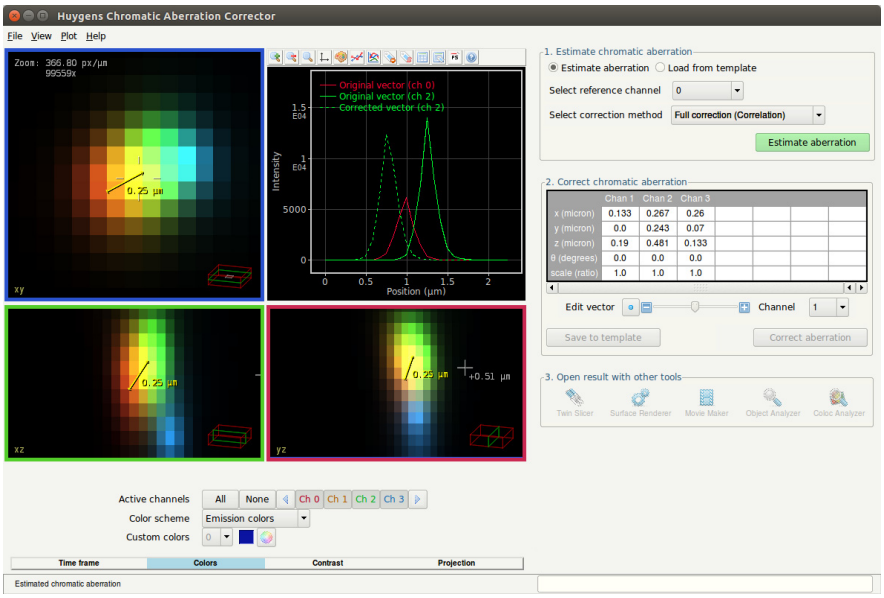

FIGURE 22.3. Using the Edit Channel tool of the Chromatic Aberration Corrector. The estimated vector has been manually lengthened, and as a result the dashed line plot is slightly moved to the left.

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 23

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). Both the 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 stabilizes 2D or 3D time series, and also aligns slices within a 3D stack. Within the opening screen, you can select whether the correction needs to be applied on time series or 3D stack data.

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 next paragraphs:

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

→CONVERTIn Huygens Essential, this function is available from the operations windows main menu EDIT→CONVERSIONS→ XY-Z TO XY-T.

Pre-processing

After selecting the stabilization method, a *Pre-Processing* step is shown where there is the possibility to launch the cropping tool. If the image concerns multichannel data, you can also select here the channels you plan to use for calculating the stabilization. Addition-

ally, you can select here the quality of the interpolation that is being used to resample the images.

The Cross Correlation Method

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

After the pre-processing step, the *correlation settings* screen 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 23.1. The NEXT button will start the alignment process.

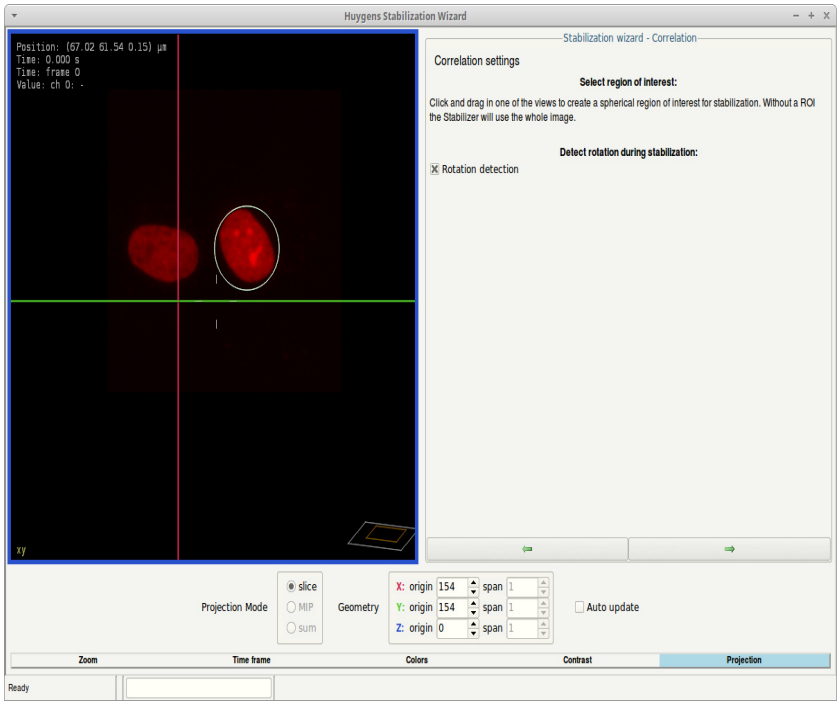

FIGURE 23.1. The *Correlation settings* screen in the Object Stabilizer.

When the Object Stabilizer has finished measuring the displacements, the *Stabilization* screen is shown. Continue reading “Stabilization Settings” on page 125.

The Model-Based Correlation Method

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

Pushing the NEXT button on the pre-processing screen immediately starts the alignment process, because this method does not require any additional user input. When the Object Stabilizer has finished measuring the displacements, the *Stabilization* screen is shown. Continue reading “Stabilization Settings” on page 125.

The Multi Object Tracking Method

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

After the pre-processing step, the *detection settings* step is shown (Figure 23.2). In this screen the detection algorithm can be configured to *filter on brightness only*. This is a little faster and usually sufficient for tracking bright (or dark) spots in images with a high contrast between the spots and the background.

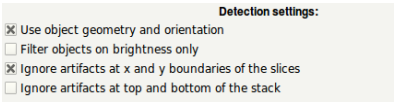

FIGURE 23.2. The *Detection settings* screen.

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

Next, the user is asked to use the *Select object* tool ( ) and *Select background* tool ( ) to mark a few objects and background areas in the first time frame. Figure 23.3 shows an example of selected objects (green selection) and background (red selection). It is important to make sure that the size of the object selections roughly matches the real size of the objects, because this is a parameter that is used to tune the detection filters.

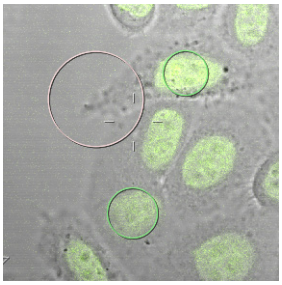

FIGURE 23.3. An example of object and background selections.

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

When the NEXT button is pressed, the Object Stabilizer analyzes the selections and uses that information to detect objects in the first time frame. The detected objects are marked by blue spheres as shown in Figure 23.4. When the detection fails, then the selection tools can be used to select new objects and background, mark detected objects as object ( ) or as background ( ), and delete selections. When the selections are modified, pressing NEXT will recompute the detection instead of proceeding to the next screen.

The histogram and sliders in the *optimizing object detection* screen (See Figure 23.4) 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 screen.

In the next screen 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

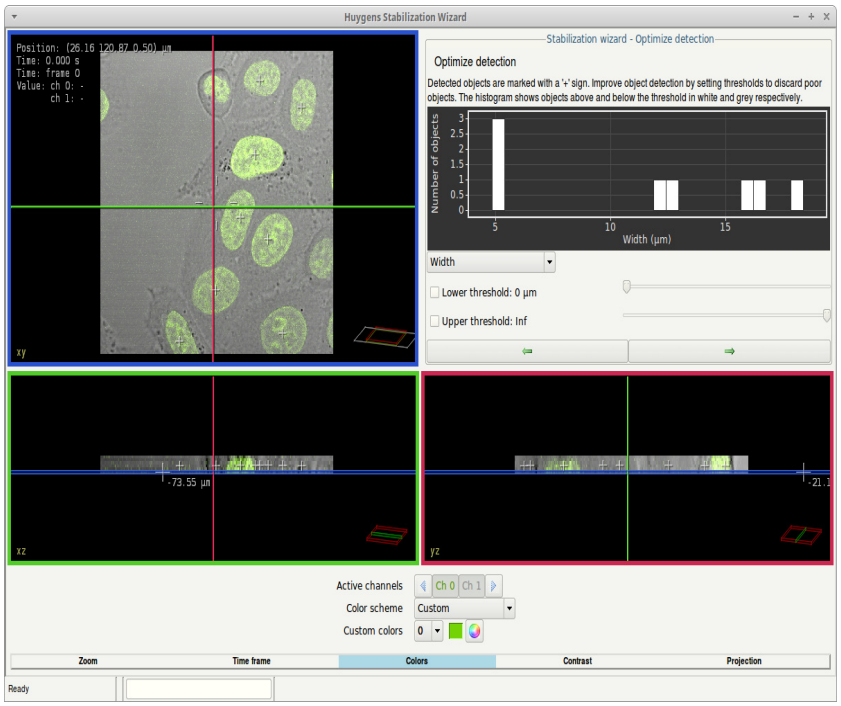

FIGURE 23.4. The *Optimizing object detection* screen in the Object Stabilizer.

to use the geometry and orientation of the detected objects –besides the position, brightness, and filter values– for connecting the tracks (default).

The NEXT button will start the tracking process. When this is finished the *edit tracks* screen is shown (See Figure 23.5 on page 125). In this stage the user can select ( ), break ( ), and delete ( ) the detected tracks. The Object Stabilizer averages the remaining tracks to correct for the average displacement of the objects.

When the NEXT button is pressed again, the *Stabilization* screen is shown. Continue reading “Stabilization Settings” on page 125.

The Center of Mass Alignment Method

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

Pushing the NEXT button on the pre-processing screen immediately starts the alignment process, because this method does not require any additional user input. When the Object Stabilizer has finished measuring the displacements, the *Stabilization* screen is shown. Continue reading “Stabilization Settings” on page 125.

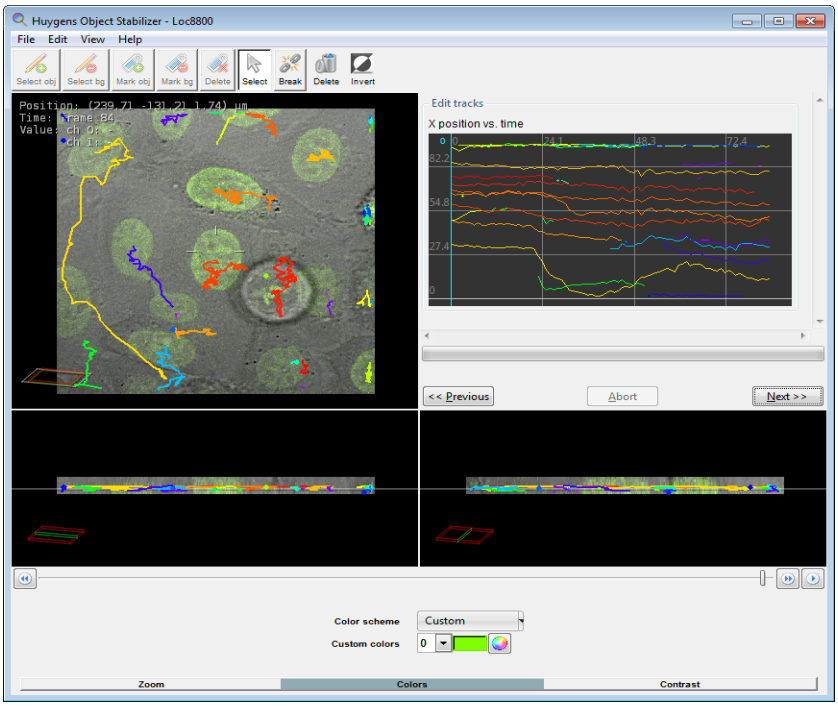

FIGURE 23.5. The *Edit tracks* screen in the Object Stabilizer.

Stabilization Settings

The *Stabilization settings* screen shows the measured displacements in *x*-, *y*-, and *z*-direction, and the rotation in plots under four different tabs. If the axial displacement is small no correlation will be applied, unless the checkbox is ticked. The three sliders above the plots can be used to filter outliers, drift, and noise from the displacement curves. When the displacement curves shows sudden steps, then the outlier filter can be used to remove those steps. The drift and noise filters can be used to respectively remove the drift and keep the random motion, or keep the drift and remove the random motion. The effect of those filters is shown in Figure 23.6.

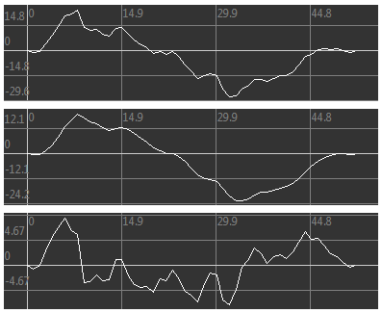

FIGURE 23.6. Effect of filters in the *Stabilization settings* screen. The top shows the measured displacement, the curve in the middle is noise filtered, and the bottom is drift filtered.

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

Cropping the Result

Because there is some extra space needed for the translated and rotated frames, the stabilized image will be larger than the original one. The *Cropping* screen (See Figure 23.7) allows the user to manually cut off the black borders, crop the image to the original size, or keep the full size. When selecting *Original size*, the reference frame needs to be selected.

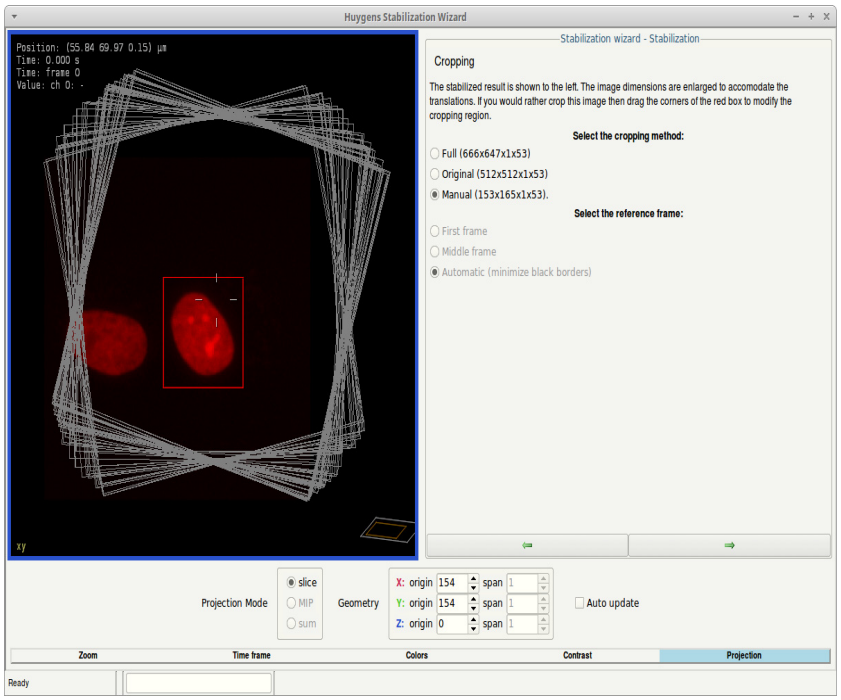

FIGURE 23.7. The stabilization result can be cropped manually to cut out the objects of interest. The original image borders are shown as grey rectangles.

Alignment of Slices in 3D Stacks

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

The 3D alignment method always uses the cross correlation method for comparing adjacent frames, and it can correct for *x* and *y* translation and rotation. The steps in the wizard are the same as for the stabilization of time series, as explained in "The Cross Correlation Method" on page 122. A comparison between a misaligned and aligned *z*-stack is shown in Figure 23.8.

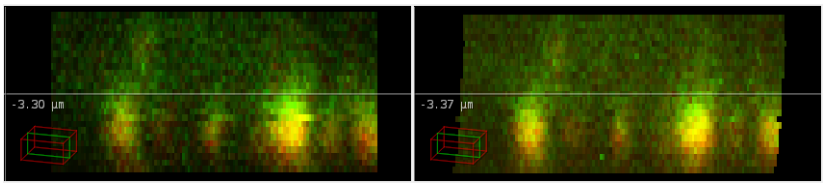

FIGURE 23.8. An *x*-*z* slice of a misaligned *z*-stack (left), and the alignment result (right). The chromatic shift in stack like this one can be correcting using the Chromatic Shift Corrector.

Note that 3D alignment deforms the volume. It is advisable to check the alignment result carefully in using the *x*-*z* and *y*-*z* views of the Twin Slicer (See "Using the Slicer in Basic Mode" on page 62.).

CHAPTER 24

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 24.1). 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 4.2 on page 18) 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 20.). 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.

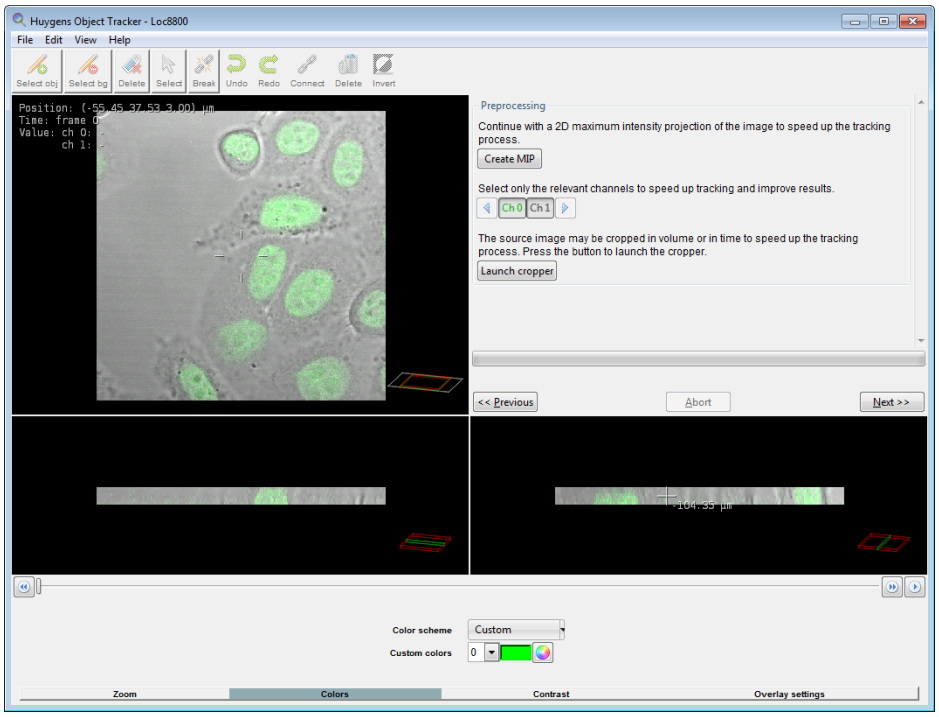

FIGURE 24.1. The preprocessing stage in the Huygens Object Tracker.

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.

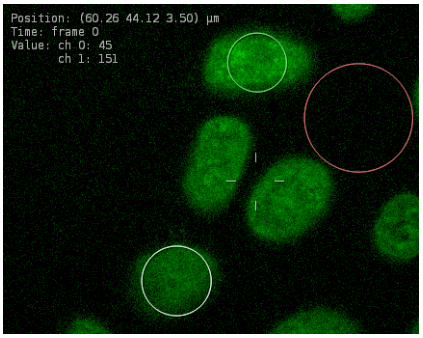

FIGURE 24.2. Two object selections (green) and a background selection (red).

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 24.2).

Like the Orthogonal Slicer (page 69), 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 139.), 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 24.3).

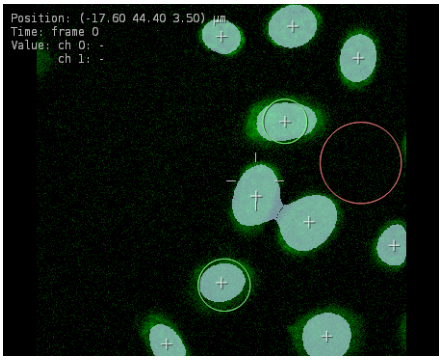

FIGURE 24.3. Detected objects.

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.

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 24.4).

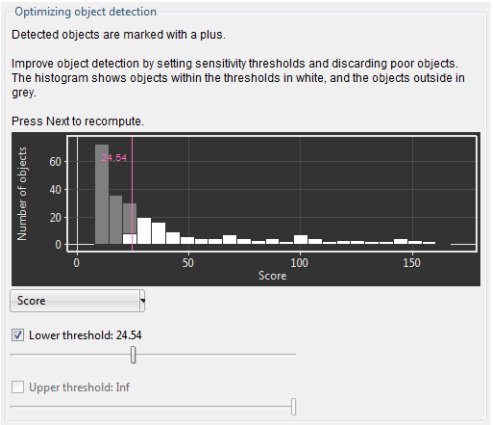

FIGURE 24.4. Increase the threshold on the score to reject more objects.

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.

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 deleted 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 24.5). The projection angle can be changed using the slider below the graph.

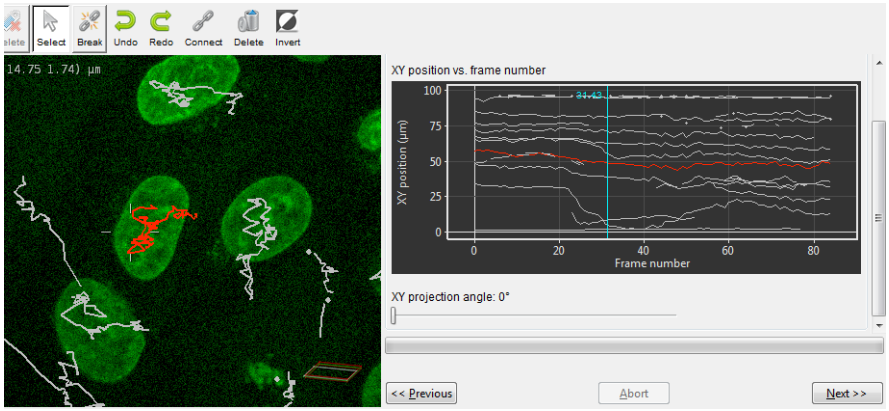

FIGURE 24.5. The *quick edit tracks* stage allows to connect, break, and delete tracks.

An angle of 0° projects on the *x*-axis while 90° projects on the *y*-axis.

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 24.6) can select a group of tracks based on common properties.

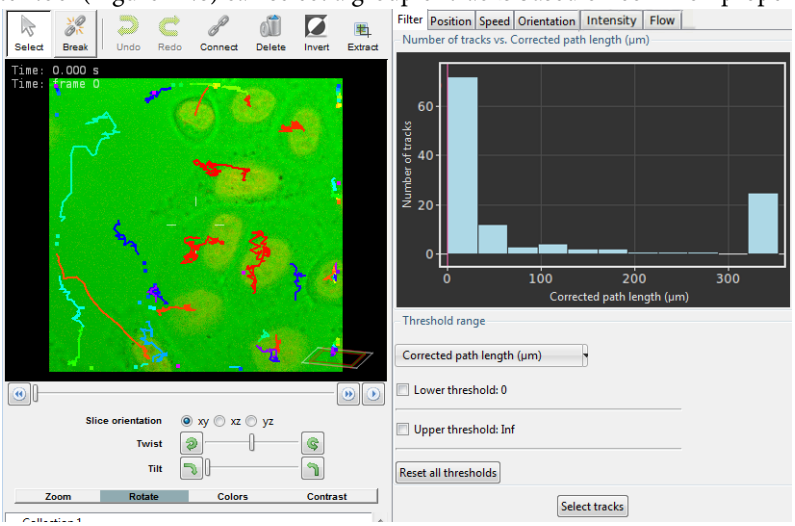

FIGURE 24.6. The filter tab in the Track Analyzer.

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.

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 24.7, gives the number of objects versus the distance to their track origins for a specific duration. Example; if the time frame slider is

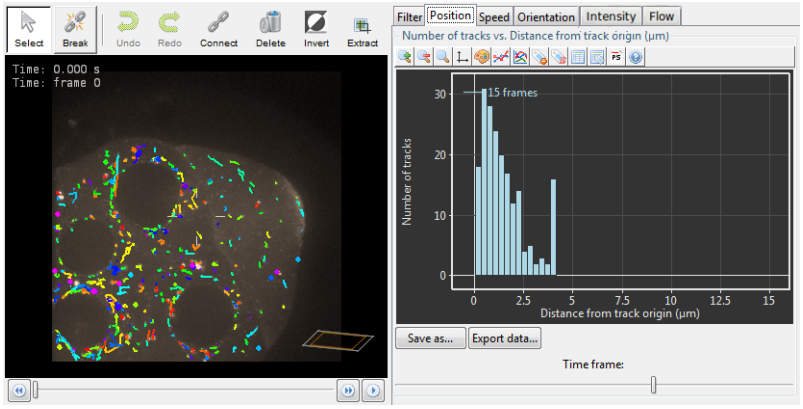

FIGURE 24.7. In case of a large number of objects that show Brownian motion, the shape of the histogram in the *position* tab will be a half-normal distribution.

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.

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 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.

$$MD(t) = \frac{1}{N}\sum_{i=1}^{N}\left|\bar{x}_{i,\,t} - \bar{x}_{i,\,0}\right| \qquad \textbf{(EQ 1)}$$

The third plot shows the mean squared displacement, or *MSD*, as given in Equation 2. Note that the x-axis of this plot shows the duration t in seconds instead of frames.

$$MSD(t) = \frac{1}{N}\sum_{i=1}^{N}\left|\bar{x}_{i,\,t} - \bar{x}_{i,\,0}\right|^2 \qquad \textbf{(EQ 2)}$$

There is also a close relation between the *MSD* and the histogram in the first plot (Figure 24.7). 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 24.8)..

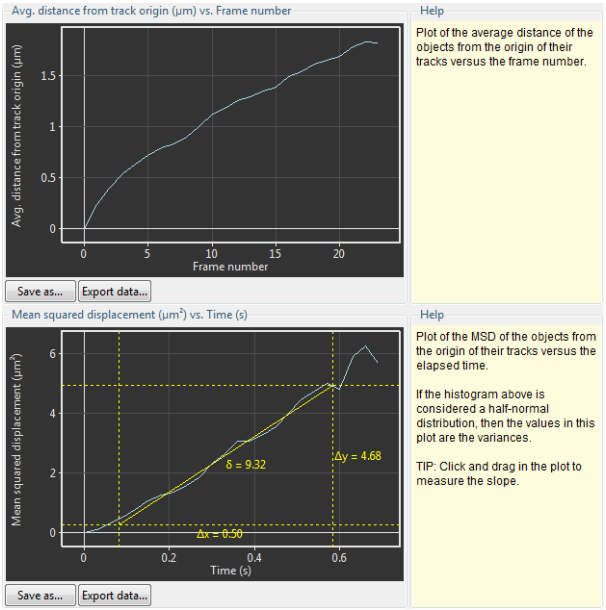

FIGURE 24.8. The slope of the mean squared displacement plot indicates a diffusion coefficient of 9.32 µm²/s.

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 24.9.

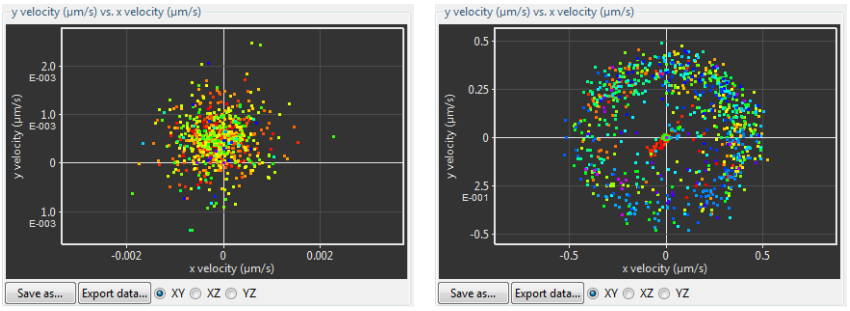

FIGURE 24.9. Left: An *xy* velocity scatter plot showing a drift in the positive *y* direction. Right: An *xy* velocity scatter plot of objects moving with a constant speed of about 0.35 µm/s in random directions.

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 numers).

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 24.10).

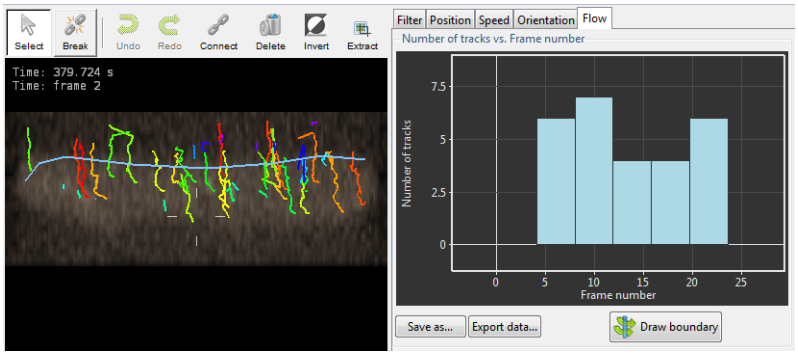

FIGURE 24.10. The flow histogram shows that objects are crossing the boundary between time frame 4 and 24.

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 25

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 15 on page 83).

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 151.

Starting the Object Analyzer

- Launch 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 Wiki[1], 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 11 on page 61) 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.

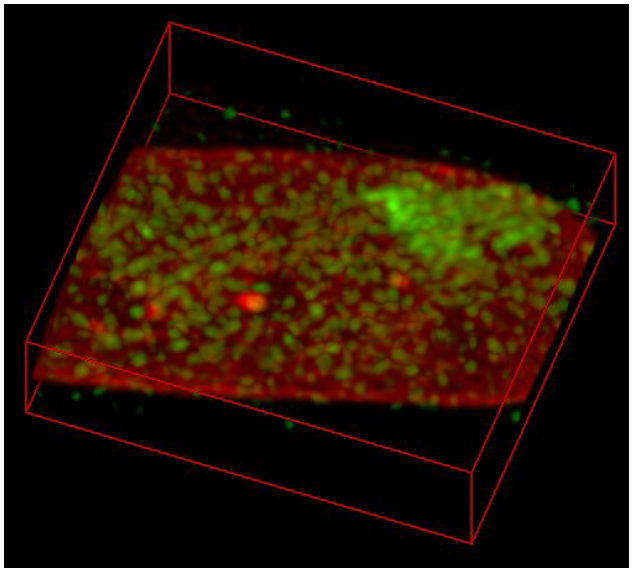

FIGURE 25.1. Maximum intensity projection of the test image. Cell nucleus FISH-stained, recorded at the Nuclear organization Group, SILS, University of Amsterdam (head: Prof. Roel van Driel), under the 3D-Genome research project.

The image in Figure 25.1 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 25.2). 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 and saturation sliders.

1. http://www.svi.nl/ObjectAnalyzerTutorial

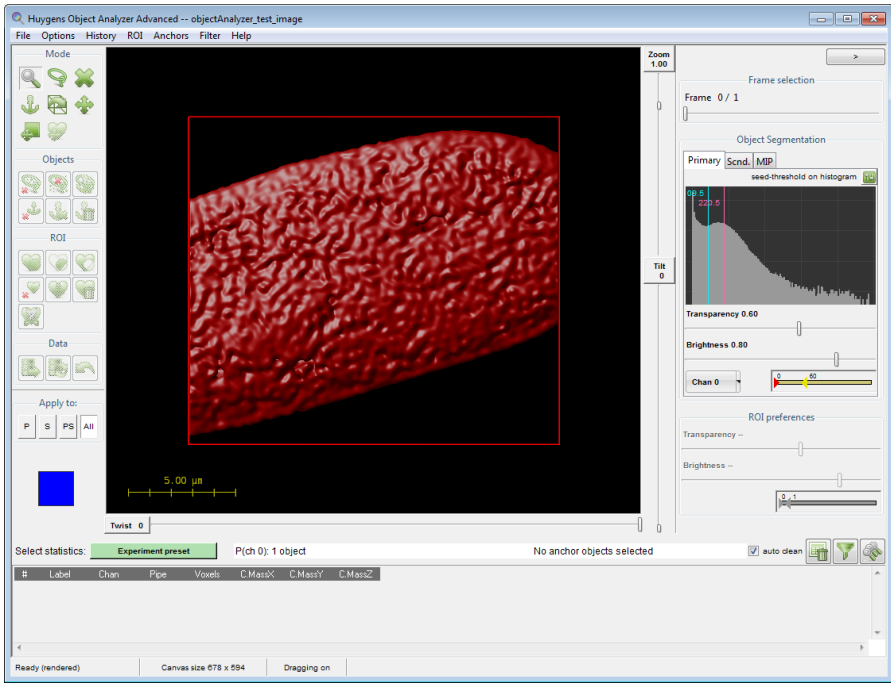

FIGURE 25.2. Start up screen of the Object Analyzer. The object in the image's first channel is automatically segmented with initial parameters calculated based on the intensity distribution (Otsu's segmentation).

Segmenting the Objects: Setting the Threshold

The segmentation method in the Object Analyzer uses a combined seed-and-threshold method, see Figure 25.3. 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 tutorial[2].

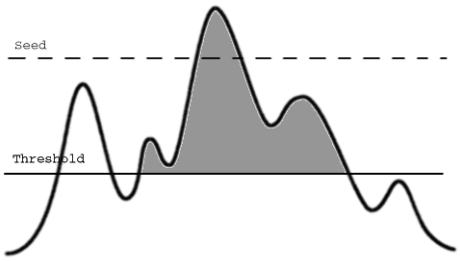

FIGURE 25.3. Visual representation of the seed-and-threshold segmentation method.

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.

2. http://www.svi.nl/ObjectAnalyzerExpertTutorial

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.

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 25.4.

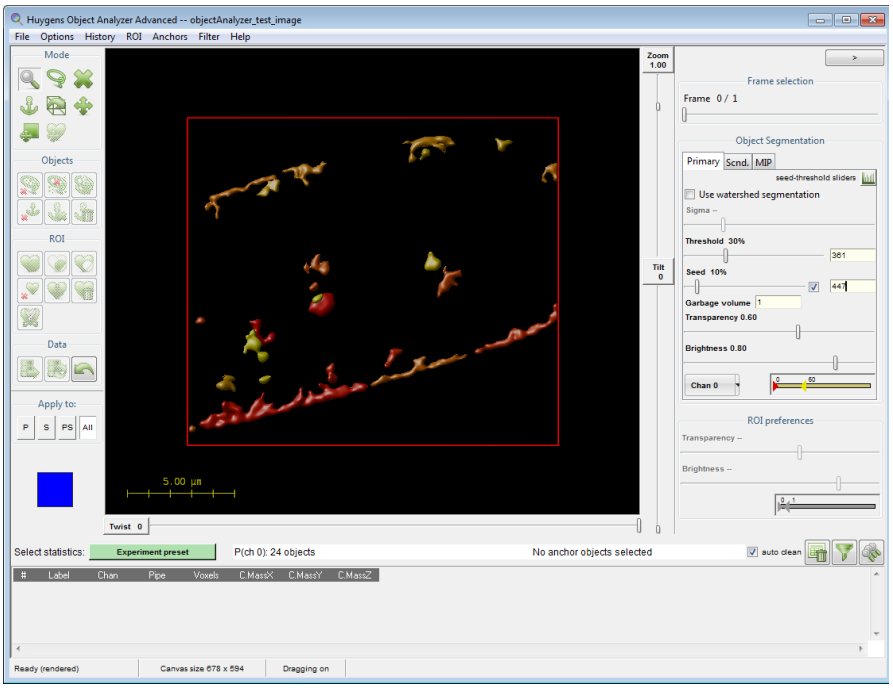

FIGURE 25.4. Different segmentation conditions. Increasing the threshold decreases the objects size, and adds more objects to the scene.

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 25.5.

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.

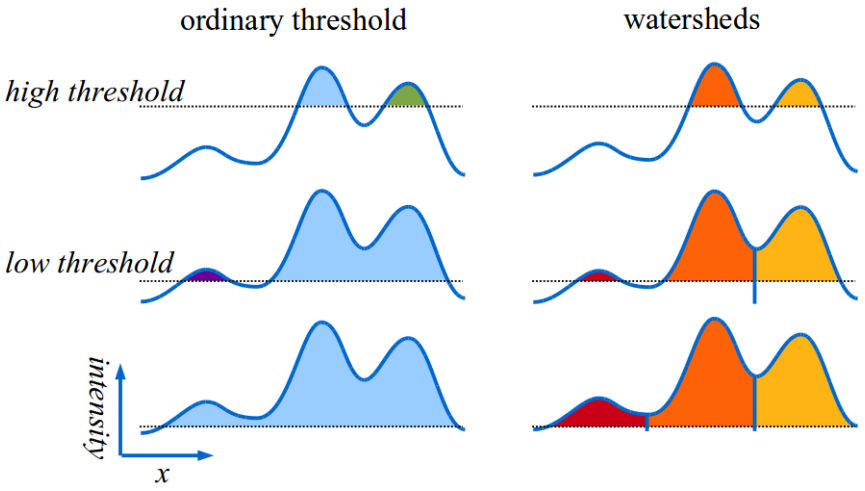

FIGURE 25.5. Difference between the seed-and-threshold and watershed segmentation methods.

Interaction with the Objects

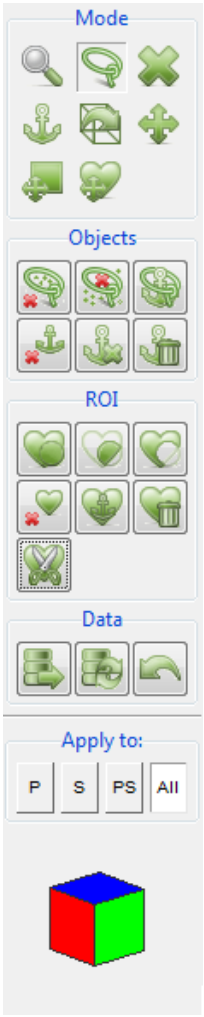

On the left side of the Object Analyzer window you can find a column of buttons. Many of the buttons are deactivated at this moment, they will be activated when it makes sense later. But almost all icons in the first group are always available. The first group of icons control the *Mouse mode*, i.e. 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 154 .

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 151 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

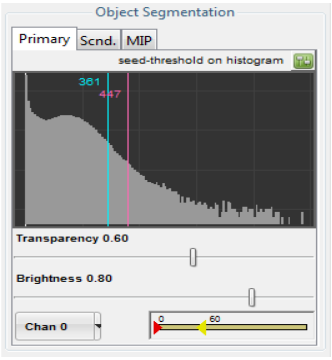

FIGURE 25.6. Rendering pipes control panel.

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 25.6).

With this approach, you can also explore two channels at the same time?

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 import-

ant', 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 200.). 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 25.7.

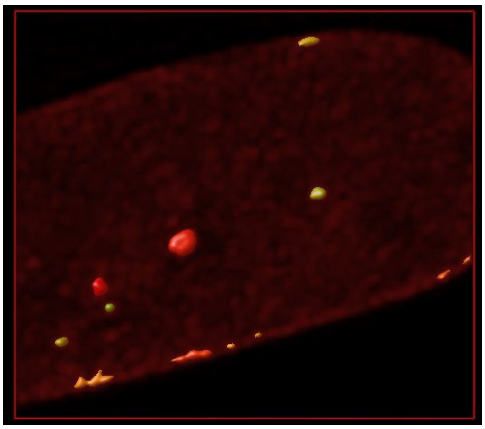

FIGURE 25.7. Objects and MIP. Objects in a surface pipe rendered together with a MIP pipe for spatial reference. The data channel is the same one in both pipes.

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 25.8.

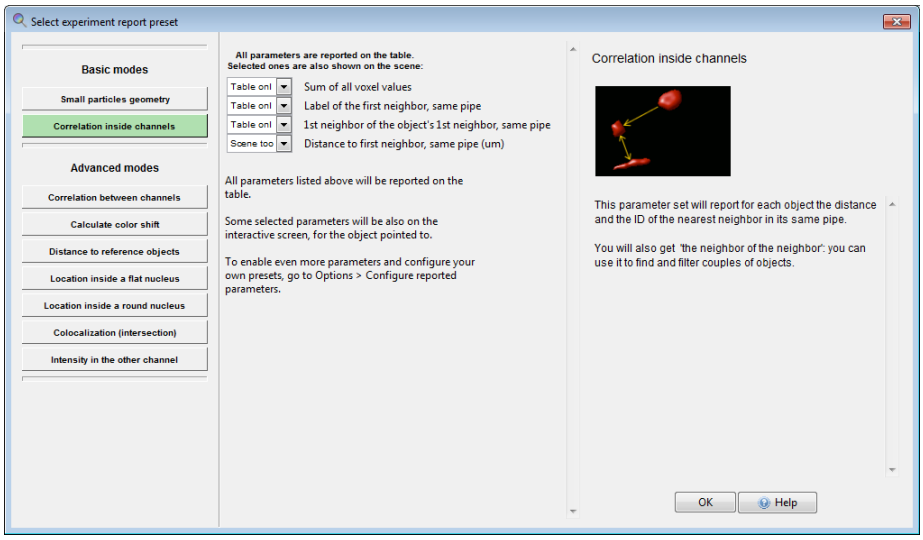

FIGURE 25.8. The Experiment Presets dialog. Collections of parameters can be selected to be reported on the table and the scene, depending on your specific experimental needs.

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.

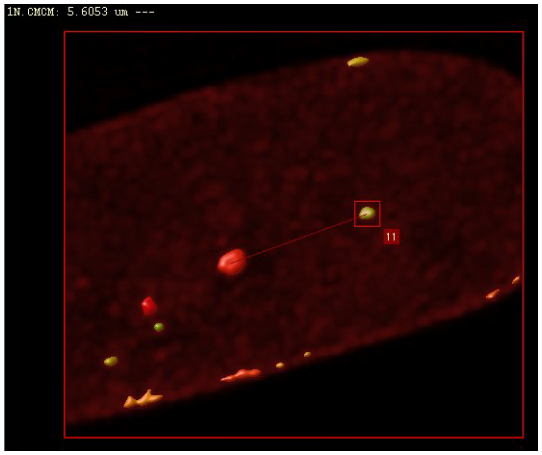

FIGURE 25.9. Interaction with the objects. When the distance to the nearest neighbor is computed, the distance line is also displayed on the screen for the object aimed at.

You may have noticed that something else happens when you interact with the rendered objects. In the example of Figure 25.9, 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 147.

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 25.10). Furthermore, the object corresponding to the current row will be highlighted on the canvas, and the distance to its nearest neighbor will be shown.

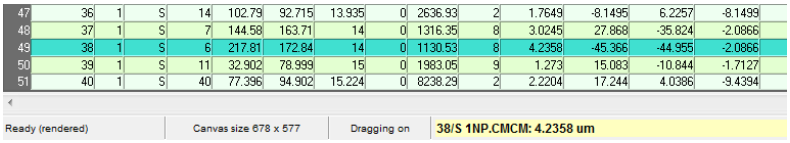

FIGURE 25.10. Exploring the table.

You can find an object in 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 25.11, for example, shows a plot of the partile sphericity distribution.

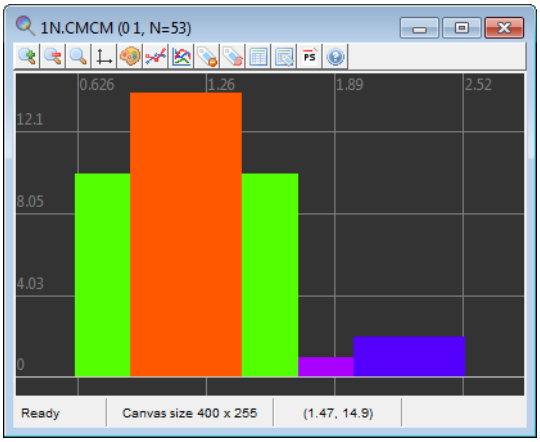

FIGURE 25.11. Histogram representation of one column values.

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 tutorial[3] 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 27 on page 151.

3. http://www.svi.nl/ObjectAnalyzerExpertTutorial

CHAPTER 26

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 15.1 on page 83).

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 *WaistLat*.

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

Sphericity

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

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

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

Depending on 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 1).

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

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

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

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

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

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

Aspect Ratio

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

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

More Parameters and Filtering

Many 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.

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

CHAPTER 27

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 135 to learn how to use the basic components in context.

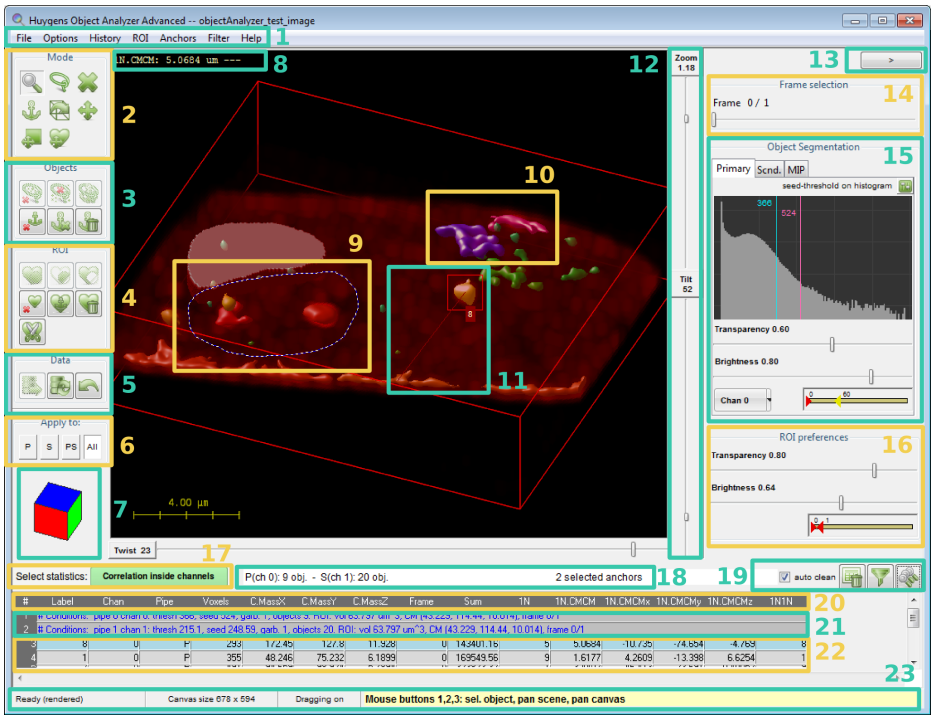

FIGURE 27.1. The Object Analyzer interface. Different regions are enumerated to describe them in sections of this chapter.

Main window components

The OA main window is shown in Figure 27.1 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 157).
- 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 160).

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 156).
  - 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 161, 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 161.

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. Hoovering 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 27.1, 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 157) 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 opperations:

- 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 deleted data was not yet exported to overwrite the original, you can always reload the original dataset with this button.
-  Undo last operation.

Active Pipe Mode (6)

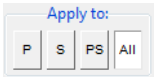

These radio buttons control, on which of the active data pipes (*Primary*, *Secondary* or both) and channels (*Selected pipe*, or *All*) they operate. Pipes are explained in "Render Pipes (15-16)" on page 159.

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)

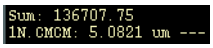

The currently selected Experiment preset (See "Experiment Presets (17)" on page 160) selects the statistics to report many parameters to the table (See Figure 27.1, item 20, 21, and 22) A few of these parameters can be also reported on the screen for the current object for easy reading. The magnitude of the distance that is plotted on the screen is indicated by a triple dash '---'

The Selected Area (9)

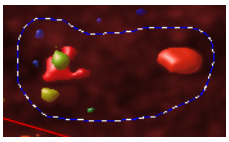

The *Select area* mouse mode allows you to define a region in the image. So for example you can: analyze objects 'inside' the ROI, anchor them as references, or use the selected volume to define a region of interest (ROI).

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.

1. http://www.svi.nl/RayTracing

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 155) 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)

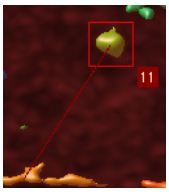

Objects under your cursor are highlighted. Once a certain object is analyzed, the object will also show a red box surrounding the object.To analyze an object just click on the object with the *Analyze object* mouse mode, or press the ANALYZE ALL button (See Figure 27.1, item 19).

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 27.1, 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 156), 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 selector (14)

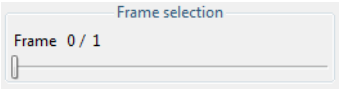

When time series are loaded in the OA, this slider controls which time frame is take for visualization and analysis.

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

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

Render Pipes (15-16)

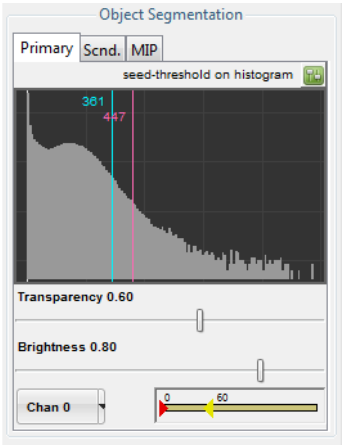

The Huygens Object Analyzer has two surface pipes (named *Primary* and *Secondary* pipes) for you to put image data through for object analysis and another pipe to simultaneously visualize a MIP of one of the channels (See Figure 27.1, item 16).

'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. One can even segment the same channel twice with different conditions in each of the pipes if necessary!

This pane controls which data channel goes through which pipe, how its objects are segmented (with a seed and a threshold[2]), 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 vailable. This alternative representation also show their values as a percentage of the channel maximum and allows you to link the seed with the threshold.

2. http://www.svi.nl/SeedAndThreshold

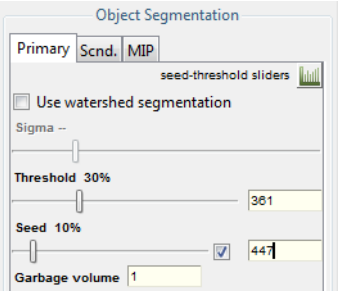

The garbage volume control is also in the slider pane, not in the histogram. To switch between the histogram and the sliders control panes you have to click on the small button ( ) at the right of the pane title.

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 that 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 rage), 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.

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 200.). You can collapse this range completely if you want so all objects in a pipe get the same color.

Experiment Presets (17)

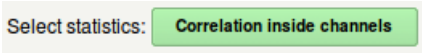

This button opens a preset selection dialog that allows to select an experiment preset, a collection of parameters that make sense to be reported together in the context of certain experimental needs (See Figure 25.8 on page 142).

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

In this dialog you can 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 161). The check boxes allow a few of these parameters to be also reported on the screen (See Figure 27.1, 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:

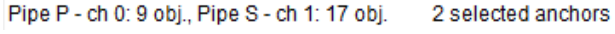

- left: the number of objects currently segmented in each surface pipe.
- right: the total number of objects selected as anchors.

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
- 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 27.1, item 6). When a 2D area has been selected (See Figure 27.1, 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.

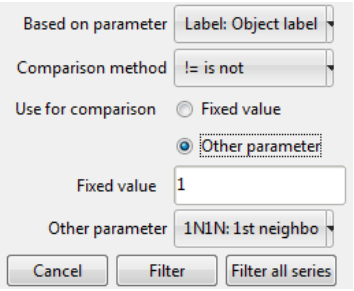

Because the filter is based on the reported parameters, you may need to configure the reported parameters or select other statistics first in order to filter based on the desired property.

For interesting usages of this filter, see for example the neighbors article in the SVI Wiki[3].

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.CMCM	1N.CMCMx	1N.CMCMy
0	# Image: objectAnalyzer_test_image, analyzed with HuPro 3.6.0p0 on 2010-06-03												
1	# Conditions: pipe 0 chan 0: thresh 366, seed 524.07, 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	355	48.246	75.232	6.1899	0	169549.56	9	1.6177	4.2609	-13.398
4	2	0	P	491	94.659	98.924	5.7894	0	323823.32	1	3.4402	-46.413	-23.691
5	3	0	P	1677	60.758	26.59	6.7792	0	730828.35	7	2.6129	-34.511	16.137
6	4	0	P	514	238.13	79.466	6.5037	0	220270.08	5	5.3358	-76.415	-26.321
7	5	0	P	396	161.71	53.145	7.1591	0	164795.01	8	5.0684	10.735	74.654

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

The Table Columns and Their Headers (20)

Distance to first neighbor, same pipe (um)

When you move your mouse over the column titles at the top of the table a description of each parameter will be shown at the very bottom of the window (See Figure 27.1, 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 27.1, 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

3. http://www.svi.nl/ObjectAnalyzerNeighbors

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

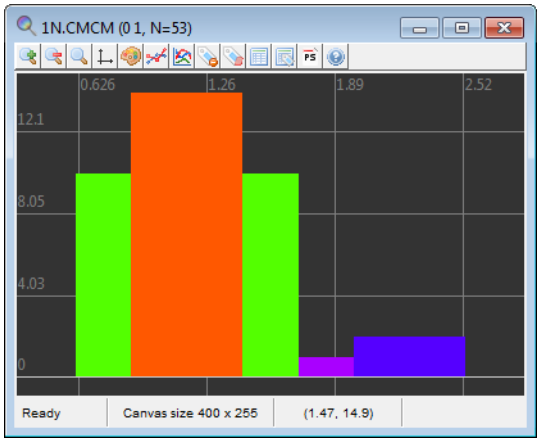

Clicking on a column title selects the whole column for you to copy. Multiple columns can be selected by holding the `Ctrl` key. To plot a histogram of the distribution of values in a column, select the column (or a subset) and select Histogram from the pop-up menu.

Selected columns also act as 'special parameter selector'. The value of the parameter for selected columns will be reported on the status bar (See Figure 27.1, item 23) when you explore cell values for given objects. Like this, you can quickly compare different parameter values for the same object, by looking at the report in the status bar while you move your mouse over the table cells.

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

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

Conditions Report (21)

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

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

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

There is another way to calculate the CM of the ROI, not considering it homogeneous but taking into account the real image intensities in envelops. 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[4].

4. http://www.svi.nl/ObjectAnalyzerROI

The Table Rows, How to Explore Them (22)

Each row in the table is either a *Conditions report* (See "Conditions Report (21)" on page 163) 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 156). 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 27.1, 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 27.1, item 23). This, combined with the selection of rows (see below) or columns (See Figure 27.1, 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 um`. 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 alsoo shown: 4.23 microns.

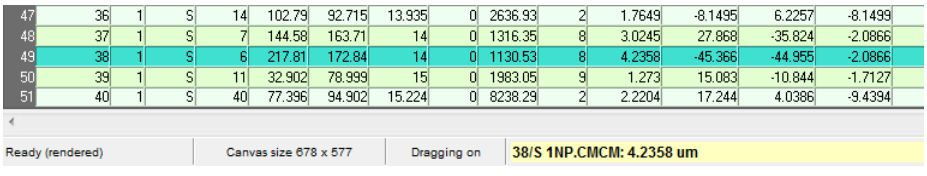

That tooltip region can also show other values that can be used as reference. Try this: while keeping the `Ctrl` key pressed on the keyboard, select a couple of rows by clicking on the row number at the very left of the table. The selected rows will turn green. If you now hover the mouse over a cell 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 seperated in the table.

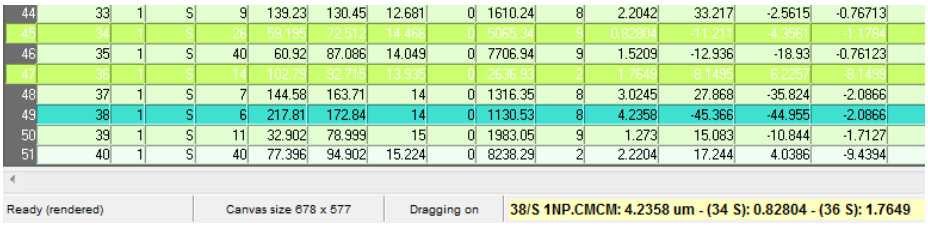

If you select columns instead of rows in the table the tooltip will display the corresponding parameter values for the same object. It is possible to highlight columns and rows at the same time but this is confusing to interpret, and thus not recommended.

The Status Bar and Tooltip (23)

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

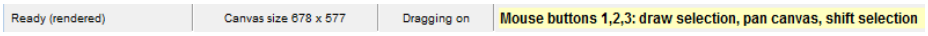

The left side reports the current status of the renderer and analyzer. You can see whether a long computation is running or if the analyzer is ready for further interaction.

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 27.1, 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 28

The Colocalization Analyzer

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

This overlap can be defined in various ways; Huygens supports the *colocalization coefficients* most commonly used in literature: (*Object*) *Pearson, (Object) Spearman*, *Overlap*, *Intersection*, and *Manders M* and *K*. The SVI Wiki provides extensive information on these coefficients[1,2].

The Colocalization Analyzer works in two steps: First, it searches for colocalizing regions in the complete image and calculates the coefficients values. Then in the second step it segments these regions and treats them as objects required for further visualization, and for the analysis of the object-based coefficients.

Note that the Object Analyzer module (See page 135) also provides a form of colocalization measurement which complements the functionality of the Colocolation Analyzer. Whereas the Colocalization Analyzer takes the whole image into account, the Object Analyzer works at the object level; it first segments the 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 208.).

How to use the Colocalization Analyzer

To *start* the Colocalization Analyzer in Huygens Essential, right-click an image thumbnail to open the contextual menu, then select ANALYZE → COLOCALIZATION ANALYZER. Alternatively, select a thumbnail and select in the menu bar ANALYSIS → COLOCALIZATION ANALYZER. The image must be a multi-channel image (See "Multi-channel Images" on page 27.) as colocalization is based on the overlap of different channel intensities. Wait for the analyzer to initialize and to compute the 2D histogram with the default settings.

1. http://www.svi.nl/ColocalizationTheory
2. http://www.svi.nl/ColocalizationCoefficients

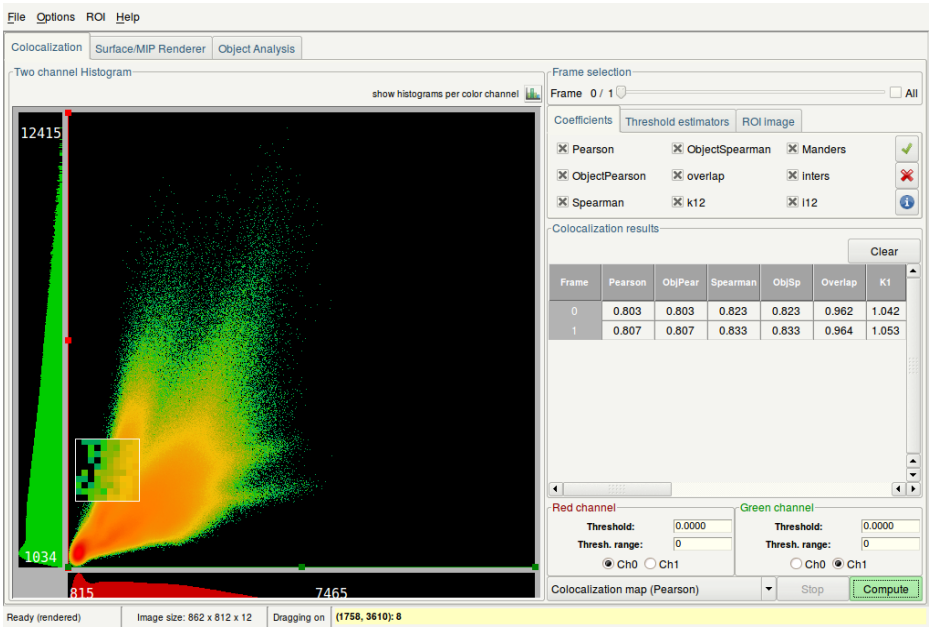

FIGURE 28.1. The first tab of the Colocalization Analyzer showing the colocalization coefficients and histograms.

When the colocalization analyzer is opened you will see three main tabs (See Figure 28.1), of which the first is visible. In this tab the colocalization coefficients can be calculated. You can choose the time frame, select colocalization coefficients, set the threshold settings (see "Threshold Estimators" on page 171, select a ROI image ("Selecting a ROI image" on page 173), and choose a colocalization map (see "The colocalization map" on page 169). On startup, the thresholds are initialized to 0 and the ROI image is disabled, so you can immediately press COMPUTE.

First we select the data to analyze. For time series, the *Frame* slider selects the time coordinate. In the lower part of the window we can select which data channels from our image we want to use. For the channels we follow the usual naming convention in colocalization theory: *Red* (R) for the first channel, *Green* (G) for the second channel.

Whenever we change the Red or the Green channel selection a two-dimensional *histogram* is calculated. This 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 would produce a 2D histogram with values only on the coordinate axes. The histogram can give an indication of some image defects such as clipping, hot pixels, and dynamic range differences. For more information on the 2D histogram see the SVI Wiki[3].

Along the x and y line of the 2D histogram, there is a separate 1D histogram for each color channel. Clicking the SHOW HISTOGRAM PER COLOR CHANNEL button near the top-right corner of the histogram shows enlarged versions of these histograms.

By moving the colored *treshold* lines in either the 2D-histogram, the enlarged 1D-histograms, or by changing the numeric values in the input fields, these values are subtracted

3. http://www.svi.nl/TwoChannelHistogram

from the voxels intensities when calculating the coefficients (if the result is negative, it is understood as a zero). Generally the colocalization coefficients depend much on correct estimation of the image background and resolution. For these reasons we strongly recommend to compute colocalization coefficients only on deconvolved images.

It is possible to set a threshold value and/or range prior to calculation of the colocalization coefficients. To see how these values influence the calculation of the coefficients, we like to refer you to our wiki pages.[4]

The colocalization map

Next we select what *colocalization map* we want to calculate: (*Object*) *Pearson, (Object) Spearman*, *Overlap*, *Manders M* or *K* or *Intersection*. You can find the selector at the bottom of tab 1 next to the COMPUTE button.

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

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

Your choice of the colocalization map is needed for the functionality in the second tab (See Figure 28.2). In this tab an iso-colocalization surface is generated based on the colocalization map. This iso-surface rendering is only possible if a colocalization map exists. Therefore, at startup the iso-colocalization surface sliders are deactivated, as we have to calculate a colocalization map first, by pressing COMPUTE.

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

The obtained colocalization map is represented in the renderer window by iso-colocalization surfaces. These iso-surfaces represent points which all have the same colocalization value, thus regions in which the degree of colocalization exceeds a certain value become objects. This "certain value" can be controlled by the "*Isosurface threshold*" slider in the Iso-colocalization surface parameters. The transparency and the brightness of this surface pipe can be controlled with the corresponding sliders. The color range in which the objects are displayed can be modified using a hue selector (page 200). One can also switch the surface pipe off.

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

4. `https://svi.nl/ColocalizationTheory`.

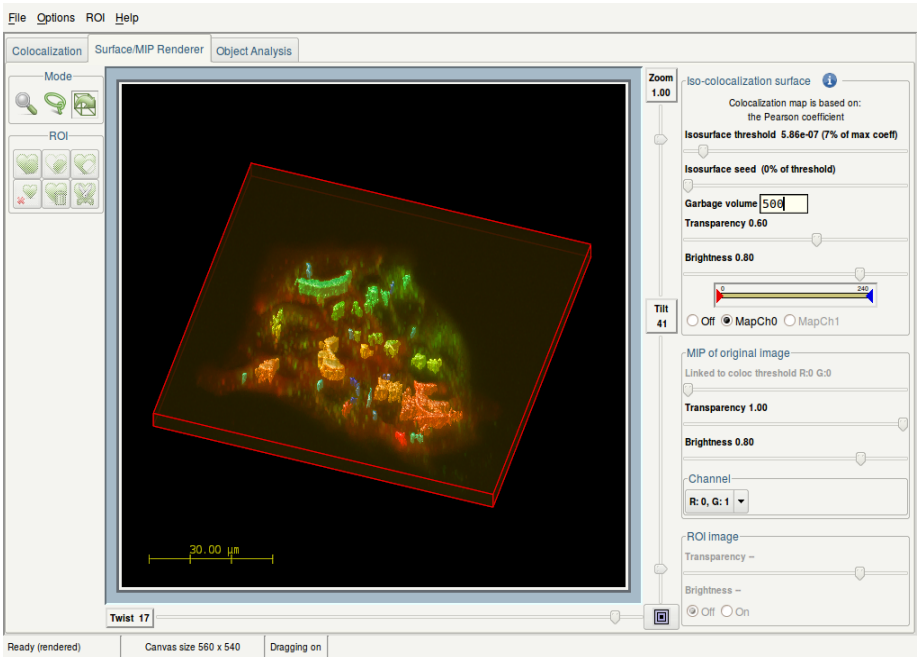

FIGURE 28.2. Tab two of the Colocalization Analyzer shows surfaces for regions of overlapping channels

nel to be used for the rendering can be selected using the MapCh0 / MapCh1 radiobuttons.

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

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

On the left side of the Surface/MIP Renderer tab, different tools are available for object analysis and ROI selection. The 'lasso' tool can be used to hand-draw a 2D area in the surface rendering, which can be used to create an ROI by clicking on the left-top icon under the section 'ROI'. In fact multiple selections can be draws and added to the ROI by repeating these steps. As a result, the ROI is enabled in the *Colocalization* tab under the sub-tab "ROI image" and applied to the original image, after which the 2D histogram is redrawn and the COMPUTE button can be used to recalculate the selected coefficients for the ROI. Using the menubar, the ROI can be exported to the main window or stored (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.

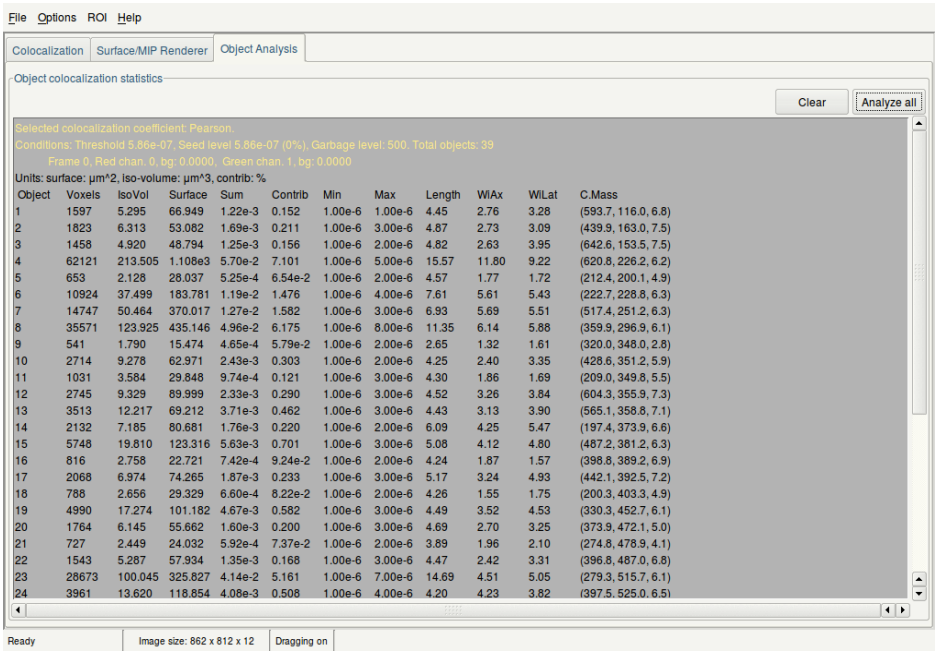

FIGURE 28.3. Tab three of the Colocalization Analyzer if for object analysis for the colocalized regions of the colocalization map.

Iso-colocalization object analysis

One of the features of the colocalization analyzer is iso-colocalization *object analysis* (Figure 28.3 on page 171). It allows to determine quickly the properties of the different colocalization regions in the data. This is realized by visualizing the colocalization map as iso-colocalization surfaces. In this way regions in which the degree of colocalization exceeds a certain value become objects. By clicking on the objects local colocalization parameters are computed and reported. To relate the iso-colocalization objects to the original data the surface objects can be blended with a MIP projection of the data (See "The MIP Renderer" on page 73.). The color range in which these objects are displayed can be modified using a hue selector (page 200).

Threshold Estimators

The Colocalization Analyzer contains a number of threshold estimators. These can be used to separate object-voxels from non-object voxels. For more information please visit our online wiki[5].

The threshold estimators are:

- Gaussian minimum
- Costes method
- Optimized search

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..

5.https://svi.nl/ColocalizationTheory

Costes Method

The Costes method is based on the estimation explained in the article "*Automatic and Quantitative Measurement of Protein-Protein Colocalization in Live Cells*" by Sylvain V. Costes, Dirk Daelemans, Edward H. Cho, Zachary Dobbin, George Pavlakis and Stephan Lockett, Biophysical Journal, volume 86, June 2004 (3993-4003).

In their approach, a linear regression line is calculated first for the 2D histogram data, using a least squares fitting with perpendicular offsets. Next, the highest point on the regression line is taken to set upper intensity limits for both channels, and Pearson's coefficient is calculated over the voxels with intensities below these limits. This process is repeated moving down on the regression line, until the Pearson's coefficient closest to 0 has been found. The intensitiy limits where Pearson's coefficient is closest to 0 are then chosen as the threshold values. Please be aware that for larger datasets, this method can take a significant amount of time.

Optimized search

This method is an extension of the Costes method, but without the assumption that the ideal background combination lies on the regression line. This method uses an iterative process to search the whole 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 input dataset. For both the Costes and the Optimized search it is assumed that the noise in the image is Poisson distributed, because the Pearson coefficient of two Poisson distributions is zero.

Threshold Options

Besides the threshold estimators, there are some additional options to set the threshold settings, especially 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 cancel* the process and it will return the threshold settings, which were found before you pressed Stop.
- *Delete*: This will delete all the threshold settings for all time frames at once and sets 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 *exact* threshold settings to the other time frames.
- *Reset*: This will reset any manually changed threshold of the current time frame, back to the estimated threshold setting of the selected method.
- *Recalc*: To redo the estimation process, which could be useful for the Costes and the Optimized method.

Selecting a ROI image

A ROI image can be selected to calculate the colocalization in a region of interest (ROI). The ROI option can be enabled by ticking the "Enable" checkbox. When the ROI option is enabled, it is required to select an image that meets the dimensions of the original source image.

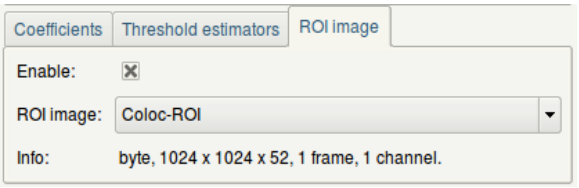

FIGURE 28.4. The mask option enabled with information about the mask image parameters.

When you hover over and click on the image name, a different image can be selected. The ROI image should already be opened in the main window of Huygens Essential.

The ROI image may only contain 1 channel. Whether the ROI image meets the criteria is reported at *Info*. You can hover over the hightlighted image parameters to read the tooltip in the statusbar to see what is wrong.

Once all the criteria are met, simply press the COMPUTE button to only calculate the colocalization in the area denoted in the ROI image. The ROI image used for the analysis is also reported in the table.

RBNCC

The Colocalization Analyzer can be extended with the Replicate-Based Noise Corrected Correlation (RBNCC) method. Noise causes difficulties around colocalization analysis, since noise also introduces 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 one can correct for noise that contributes to the underestimation of the colocalization coefficients, Pearson and Spearman coefficients. The RBNCC method has been developed by J. Adler, S.N. Pagakis and I. Parmryd from Uppsala University; Journal of Microscopy 230, 121-133; 2008. More information can be found at the SVI Wiki[6].

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. This option can then be enabled via OPTIONS→RBNCC.

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, see Figure 28.5 . The reference image should already be opened in the main window of Huygens Essential and should have the same dimensions as the original. Whether

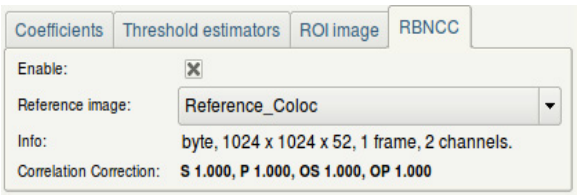

FIGURE 28.5. The RBNCC option is enabled and a reference image can be selected.

6. http://www.svi.nl/RBNCC

the reference image meets the criteria is shown at Info and hovering over any highlighted parameters shows a tool-tip in the status bar 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 the abbreviations stand for a specific coefficient:

- S: Spearman coefficient
- P: Pearson coefficient
- OS: Object Spearman coefficient
- OP: Object Pearson coefficient

CHAPTER 29

Huygens multi-GPU acceleration

Huygens can make full and optimal use of all the computer central processing units (CPUs). With Huygens version 15.05 we introduced the possibility to use the power of graphics processing unit (GPU) acceleration to the full extent, and with version 16.10.0p8 we implemented multi-GPU support. With Huygens GPU acceleration, a section of data that requires significant computation time is run on the GPU, while the rest of the data is processed on the CPU. The GPU consists of hundreds or thousands of separate processing units (cores), thereby making it extremely efficient and fast to process your data.

Huygens GPU options

Huygens GPU acceleration is available free of charge for a CUDA-enabled NVIDIA[1] graphics cards containing a maximum of 2 GB of video RAM and 1024 CUDA cores/ shading units. A few cards with higher specifications are also included (see below). In general, we distinguish the following GPU options:

- Small GPU option (Free of charge): supports up to 1024 CUDA cores and up to 2 GB of RAM
- Medium GPU option: supports up to 3072 CUDA cores and up to 4 GB of RAM
- Large GPU option: supports up to 8192 CUDA cores and up to 24 GB of RAM

More exact details on which NVIDIA CUDA-enabled graphics cards are supported by what GPU option can be found on our website[2].

Many CUDA graphics cards have already been tested for Huygens performance by us and many of our customers. The results of these tests are published on our wiki pages[3]. The performance of Huygens can be up to 30 times faster with using GPU as opposed to only using CPU, but continues to also rely on the performance of the CPU.

1. NVIDIA is a registered trademark of NVIDIA Corporation (USA)
2. https://www.svi.nl/GPUcards
3. https://svi.nl/GPU_benchmarks; https://svi.nl/HuygensGPU

Huygens and GPU acceleration

GPU accelerated Huygens deconvolution is available for CMLE, QMLE and GMLE algorithms, and is suited for Windows and GNU/Linux (Ubuntu / Debian / Red Hat / Fedora / CentOS). Specific stages in the stitching process of tiles are also benefiting from GPU acceleration. We are currently extending the GPU functionality even further to other Huygens algorithms and options. For a recent overview of the Huygens tools and options that benefit from GPU acceleration, we like to refer you to www.svi.nl/WhatsNew.

How to prepare GPU drivers and CUDA for Huygens

The graphic card installation for Windows

- Install the NVIDIA drivers for the graphics card[4].
- Download the CUDA 9.0 Network installer for Windows[5].
- Install CUDA 9.0 using the Network installer.
- Install the latest Huygens version[6].
- 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 de-activate 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 the GPU Checklist [7].

The graphic card installation for Linux

- Download the CUDA 9.0 Network Installer for your Linux operating system[8]
- Follow the instructions as specified
- Install the Huygens GPU version for Linux: sudo dpkg -i huygens_17.04.10.1-p5_amd64.deb
- Install your license or request a test license by sending an email to license@svi.nl
- GPU acceleration is active. You can de-activate 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 chapter.

GPU Check list

For latest details on the GPU checklist please visit this wiki page[9]

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 3.0.
2. **Have you installed the Nvidia driver and CUDA library?** You can install the Nvidia driver and CUDA as described under "The graphic card installation for Windows" or "The graphic card installation for Linux" on page 176

4.http://www.nvidia.com/Download/index.aspx
5.https://developer.nvidia.com/cuda-downloads#win
6.https://svi.nl/Download
7.https://www.svi.nl/GPUChecklist
8.https://developer.nvidia.com/cuda-downloads#linux
9.https://www.svi.nl/GPUChecklist

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 your GPU is listed as a Small GPU card[10]. For cards that fall within the Medium and Large class of cards, you need to purchase an additional Huygens GPU option. 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 177

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 (Figure [figure number]) has three sections (See Figure 29.1).

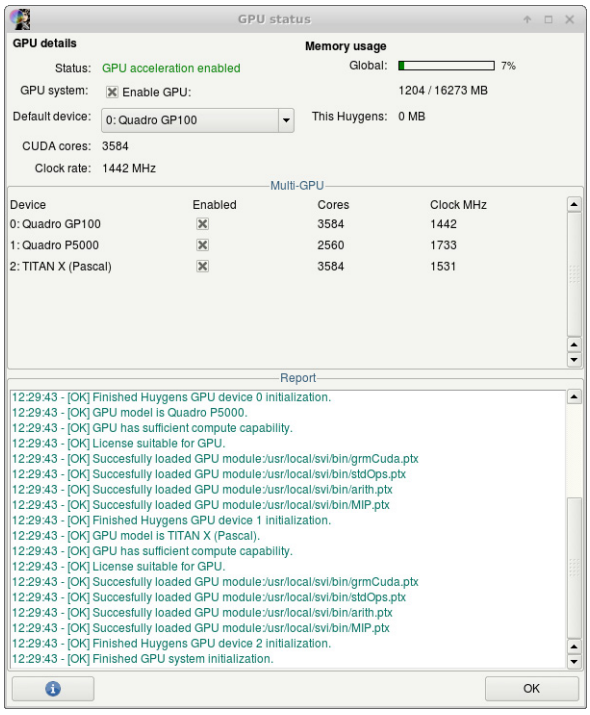

FIGURE 29.1. The GPU status window shows that the installed Quandro GP100 and P5000, and the Titan X cards are enabled in Huygens.

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 drop-down options list the GPU devices that can be selected to become the default GPU device for the current Huygens. 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

10.https://www.svi.nl/HuygensGPU

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 next column ("Enabled") with check boxes is used to enable all or a sub-set of the available devices. The selection of devices that are marked "enabled" are the devices that are used in multi-GPU operations, such as the Batch Processor and multi-GPU deconvolution commands. The rightmost two columns (Cores and Clock MHz) list the number of CUDA cores and clock speed for the devices.

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 selection of enabled devices is modified, 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 website.

The updated multi-GPU version of the *Batch Processor* has a new button bar with a new "multi-GPU" button. A prerequisite for using multiple GPUs is that the license should allow this. If this is the case, the "multi-GPU" button 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 is used like in previous releases of Huygens. If the multi-GPU mode is selected, then the button will display the number of used GPU devices in parentheses. In multi-GPU mode, by default the number of "concurrent jobs" in the batch processor is overridden by the number of (enabled) GPU devices, in order to have one concurrent job per GPU device. In case there is a long queue with many small images to be deconvolved and it is known beforehand that all (or most) images can easily be deconvolved in less than half of the GPU memory, it is possible to set a different number of concurrent jobs manually. For example, if there are two GPU devices and the number of concurrent jobs is set to four, then there will be two simultaneous jobs per GPU. In general, the GPU devices will be assigned to the number of concurrent jobs starting at the device that is designated device "0" in the GPU status window, and go through the list of devices from there, restarting at device "0" again when there are more concurrent jobs than devices.

In case there is a problem with enabling the multi-GPU functionality, the tooltip 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:

- It is greyed out, because the whole GPU subsystem is not available (see the GPU troubleshooting page on the SVI website)
- It is greyed out, because the whole GPU subsystem is not enabled (enable it in the GPU status window)
- It is greyed out, because there are fewer than 2 GPU devices (there is no multi-GPU if you have zero or one GPUs)
- It is greyed 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 multiGpu 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.

CHAPTER 30

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 183.*

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 the 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 wiki[1]

Floating License(s)

The floating license(s) option is available from Huygens 15.10 version onwards. It 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.

1. https://www.svi.nl/RemoteDisplay

- The contact-person receives an active single floating license, which can be send to the users with computers within the IP range.
- The latest Huygens version available on our download page[2] 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 licenses[3].

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 easy 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.

2. https://www.svi.nl/Download
3. https://svi.nl/HuygensFloatingLicenses

- 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.

Set-up:

For the manual: http://huygens-remote-manager.readthedocs.org/latest

For the download: http://sourceforge.net/projects/hrm/files/latest/download

For an overview: https://svi.nl/HuygensRemoteManager

For more detailed information you can read the Huygens Core programmer Guide and the Huygens Remote Manager[4].

4. https://www.svi.nl/Manuals

CHAPTER 31

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 Wiki[1].

Signal to Noise Ratio

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

1. http://www.svi.nl/BrickSplitting

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.

Using a too large SNR value might be risky when restoring noisy originals, because the noise could just being 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 31.1 shows some examples of recordings where different noise levels

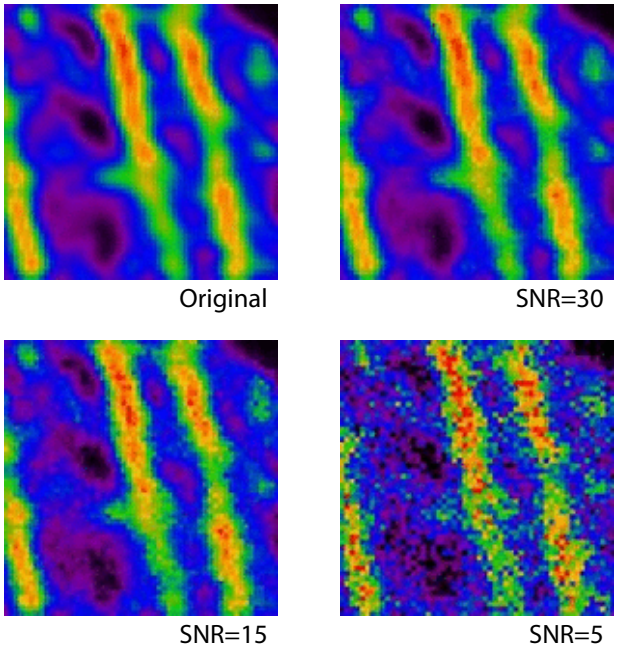

FIGURE 31.1. Examples of different SNR values. Same image, different noise levels.

were added to an original (restored) image.

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

$$SNR = \sqrt{\frac{M}{s}} \qquad \textbf{(EQ 1)}$$

If a significant blacklevel is present, it should be subtracted from *M* and *s*. See "Black Level" on page 187 and the SVI Wiki[2] 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

2. http://www.svi.nl/SNR

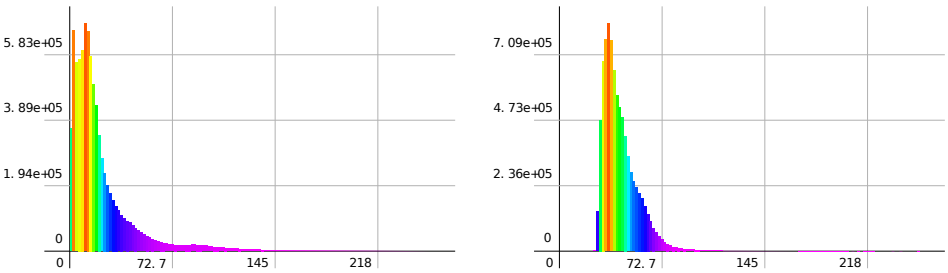

FIGURE 31.2. Left: The histogram of an image without black level. Right: The histogram of a similar image with a significant black level.

very important for the restoration method, and a rough estimation based on the appearance of the image is usually enough (See Figure 31.1)

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.

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.

A positive black level (See Figure 31.2) 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 194.) 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 Wiki[3].

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 tool[4] that computes the Nyquist rate for any widefield, confocal, spinning disc, SPIM/Light Sheet, STED, and 4-Pi microscope.

3. http://www.svi.nl/BlackLevel

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.

Figure 31.3 shows the dependency of this critical sampling distance on the numerical

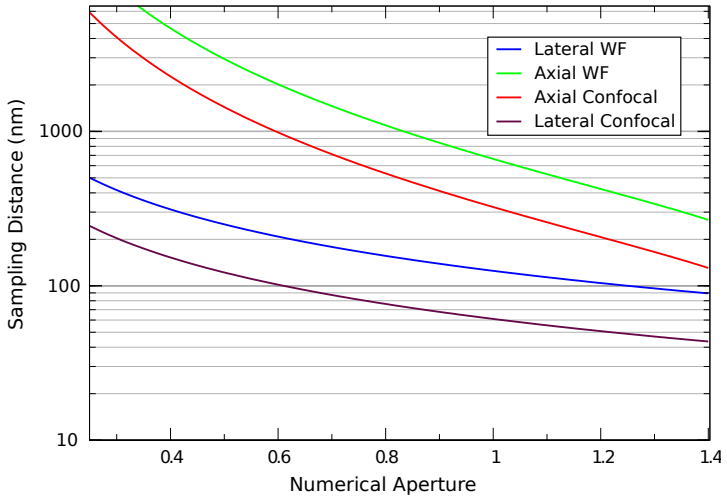

FIGURE 31.3. Critical sampling distance vs. NA. The curves above show the critical sampling distance in axial and lateral directions for wide-field and confocal microscopes. The emission wavelength in both cases is 500 nm.

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

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

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

More information can be found on the SVI Wiki[5].

4. http://www.svi.nl/NyquistCalculator
5. http://www.svi.nl/NyquistRate

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}:

$$r_b = \frac{r_{phy}}{m_{obj} m_{sys}} \quad \textbf{(EQ 2)}$$

The SVI Wiki has a calculator[6] 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 191.).

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:

$$r_b = \frac{0.61 \lambda_{ex} N_{Ad}}{NA} \quad \textbf{(EQ 3)}$$

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

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

6. http://www.svi.nl/BackprojectedPinholeCalculator
7. http://www.svi.nl/AiryDisk

Square Pinholes and the Shape Factor

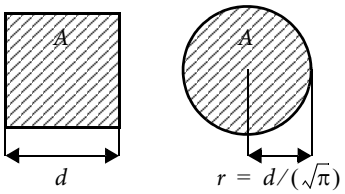

FIGURE 31.4. A square and a disk with an equal area *A*.

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 31.4), the shape factor *c* for a square pinhole becomes $1/(\sqrt{\pi}) = 0.564$. For circular pinholes $c = 0.5$ to just convert the diameter to a radius.

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

$$r_b = \frac{0.69\lambda_{ex}N_{Ad}}{NA} \qquad \textbf{(EQ 4)}$$

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 5 can be used to translate that parameter into a backprojected pinhole radius.

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

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

The backprojected Pinhole Spacing for Spinning Disks

As is the case for the backprojected pinhole diameter, the distances *between* the pinholes in spinning disks must be divided by the system magnification. For the frequently used Yokogawa disks[8], for example, the physical spacing is about 253 µm. This can be checked by imaging a stopped disk. Using a 100× 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×. The pinhole size reported in the Olympus software is the side length of this square pinhole. Therefore, $c = 1/(\sqrt{\pi}) = 0.564$, and

8. http://www.svi.nl/YokogawaDisk

$m_{obj} = 3.82$. Using a 100× objective lens and a reported pinhole size d of 150 μm the backprojected pinhole radius r_b in nm is:

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

The on-line calculator[9] 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 Pawley[10], $53 \times m_{tube}$, with m_{tube} between 1.0 and 1.56 (factor 1.25 for the fluorescence attachment multiplied by factor 1.25 for the DIC attachment). The factor of 53 includes the 8× eyepiece just below the scan head, but does not include that variability in magnification due to the variations in tube-length that are result from the aligning the system.

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

In a Bio-Rad MRC600 with a NA 1.3 60× objective, the Airy disk has a diameter of around 2 to 2.5 mm at the pinhole plane. The diameter of the first Airy zero ring is 7.6 lateral optical units[11] (o.u.), using Equation 7 to express a distance r in dimensionless o.u.:

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

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

9. http://www.svi.nl/BackprojectedPinholeCalculator

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

11. http://www.svi.nl/OpticalUnits

CHAPTER 32

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 Wiki[1].

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 rays[3]. 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

1. http://www.svi.nl/
2. http://www.svi.nl/FishtankEffect
3. http://www.svi.nl/SphericalAbberation

of already a few microns. This is especially harmful for the deconvolution of widefield 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.

- **Total internal reflection**: When the *numerical aperture of the objective lens is larger than the medium refractive index*, total internal reflection[4] 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 193.).

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 Wiki[6].

Undersampling

One of the rules of measurement that is often overlooked is sampling according to the *Nyquist rate* (See "Sampling Density" on page 187). 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 tool[7] 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

4. http://www.svi.nl/TotalInternalReflection
5. http://www.svi.nl/DynamicRange
6. http://www.svi.nl/ClippedImages
7. http://www.svi.nl/NyquistCalculator

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.

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 distored 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 21 "Bleaching Corrector" on page 1110.

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 object[8]. 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 aberration[9] 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 Wiki[10] for detailed steps in configuring the restoration process.

8. http://www.svi.nl/PointSpreadFunction
9. http://www.svi.nl/SphericalAberration
10. http://www.svi.nl/DeconvolutionProcedure

CHAPTER 33

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 diffraction[1] 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 193) 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.

1. http://www.svi.nl/ImageFormation

File Series

There are many ways in which Tiff files or other file series are named. These files can have multiple counters (referring to *slices*, *time frames*, or *channels*), and these counters can have arbitrary prefixes and ordering.

Numbered Tiff Series

If a series is simply numbered like: `slice001.tif`, `slice002.tif`, `...`, `slice0nn.tif`, then Huygens 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

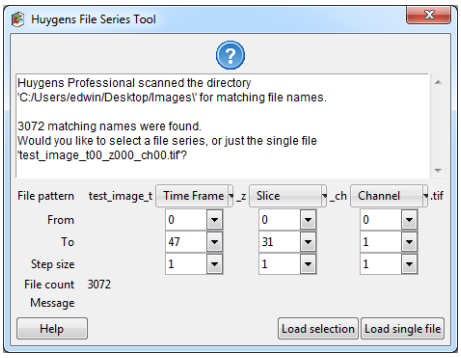

FIGURE 33.1. The Huygens File Series Tool automatically scans a directory for a file series.

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 33.1). This tool enables the user to select a subset of a file series, and select a dimension for each the indices in the file name, so that each image is assigned to the correct *z*-plane, time frame, and channel. 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.

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 18 "Stitching & Deconvolution Wizard" on page 97.
- **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.

Image Feeder

The *Image Feeder* is a functionality that can automatically open images in Huygens if these images are saved or already present in a specified *Image Feeder* directory. This directory can be defined under the DIRECTORIES tab within the Preferences window (EDITSeeFigure 33.2). You need to enable the feeder functionality in this window to make it active. The same window also allows you to open files present within the *Image Feeder* directory during Huygens startup.

Since Huygens version 17.10, files in the *Image Feeder* directory can be automatically deconvolved with the *Deconvolution Express* when Huygens is started and if this option is enabled under this *Preferences* window. Also, deconvolved results can be automatically saved to a destination directory.

Adjusting the Global Color Scheme

Huygens Essential uses a global scheme for coloring the different channels in multi-channel images. These colors can be adjusted through the *Preferences* window via the button under EDIT→PREFERENCES...→EDIT GLOBAL COLORS(See Figure 33.2).

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 colorblinds and non-colorblinds.[2]

2. Published by Okabe and Ito, "*How to make figures and presentations that are friendly to Color-blind people*", J*FLY, 2002

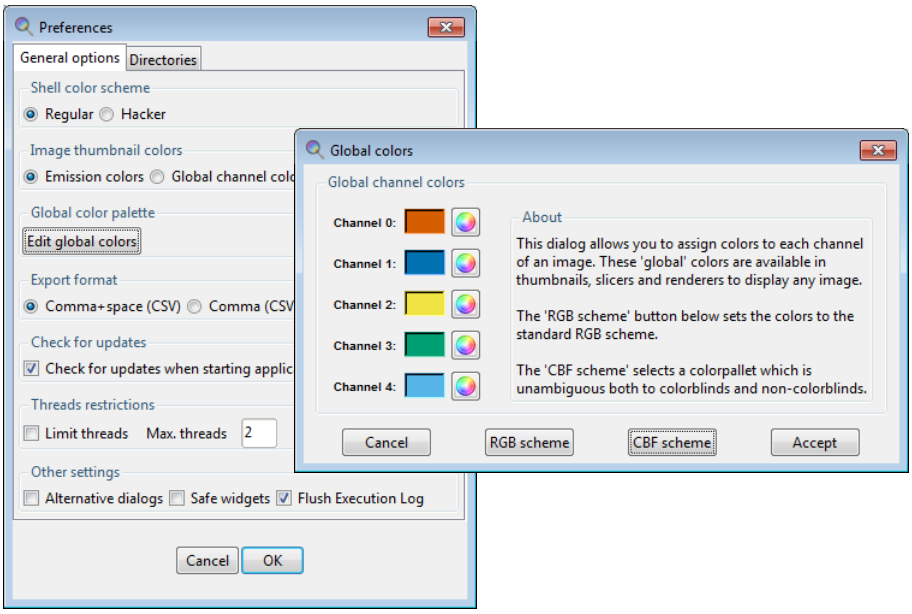

FIGURE 33.2. The global color scheme can be modified through the *Preferences* window.

Hue Selector

The *hue selector* is a component that allows adjustment of the color range in which objects are displayed (See "The Surface Renderer" on page 83 and "The Colocalization Analyzer" on page 167). 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

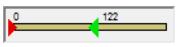

This selector allows the adjustment of a *hue* range. The objects on which this selector acts will get a color that lies within this range. The assignment of colors is based on the position of an object or on another parameter.

Hue Range and Saturation

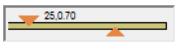

This selector allows the adjustment of a single *hue* value and a *saturation*. The upper triangle defines the color, while the lower triangle sets the saturation for this color; left is white, right is fully saturated.

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 193). 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 33.3).

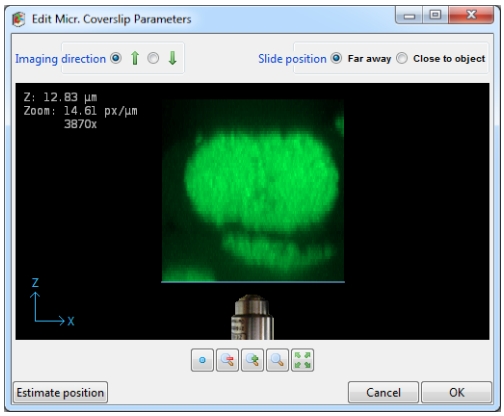

FIGURE 33.3. The coverslip position editor showing an *xy* MIP of the data along they *y*-direction. The coverslip position can be adjusted by dragging the blue line. The imaging direction, here *upwards*, is indicated by the position of the objective relative to the data as shown. The *z*-position shown 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 to100 µm, outside of the image. In this case, or in the case there is no slide, select *Far away* in the top-right selector.

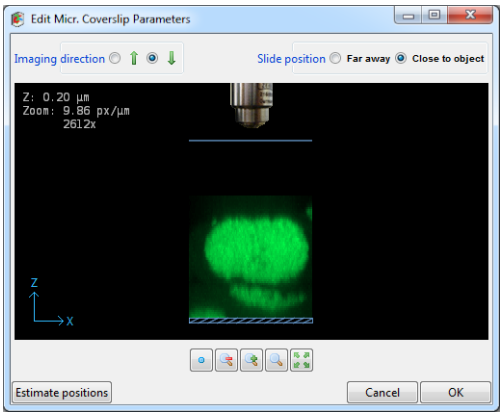

FIGURE 33.4. The Coverslip editor with the slide position set to *Close to object*.

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 profile[3]. 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 33.5) 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.

3. http://en.wikipedia.org/wiki/Gaussian_beam

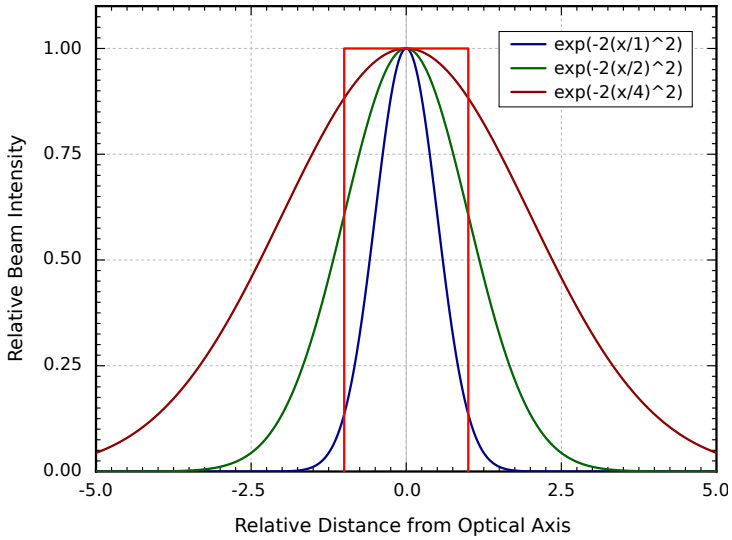

FIGURE 33.5. Lens entry pupil (*red*), beam profile with overfill factor 1 (*blue*), 2 (*green*), and 4(*dark red*). At overfill factor 1, the beam intensity is 14 % of the maximum, at overfill factor 2 the edge intensity is 61 % of the maximum.

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→INVERT IMAGE) and processing them further as incoherent fluorescence wide-field 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.

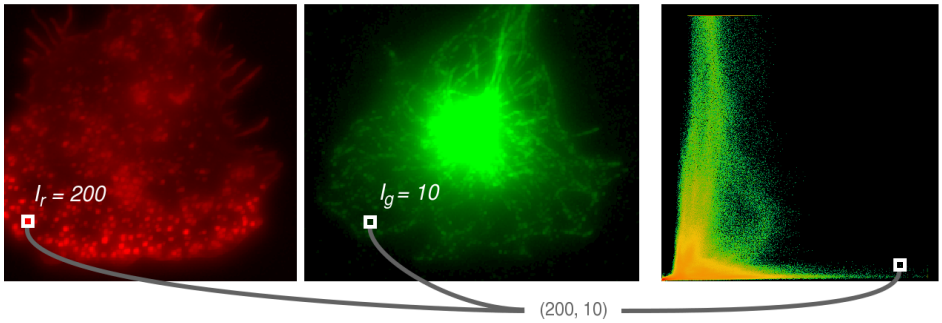

FIGURE 33.6. Example how a 2D histogram is computed, showing the first and second channel and their corresponding 2D histogram. At each position within the channels, the corresponding intensities of both channels are combined to form a coordinate within the 2D histogram. The count at this coordinate of the 2D histogram is then increased by one.

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.

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 occurence of *combinations* of intensities. To compute a 2D histogram the images need to be equal in size. See the example in Figure 33.6; 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 33.6 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 wiki[4].

4. http://www.svi.nl/TwoChannelHistogram.

CHAPTER 34

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	http://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 site[1].

Support

SVI Support Wiki

The SVI Wiki[2] is a rapidly expanding public knowledge resource on 3D microscopy and deconvolution. Based on the WikiWikiWeb principle, it is open to contributions from every visitor. In addition it serves as a support medium for SVI customers and relations to discuss different aspects of the Huygens software.

This is a list of useful starting points in the SVI Wiki to learn more about the Huygens software and microscopical imaging in general:

- Information on the parameters describing the imaging conditions (sampling, numerical aperture, pinholes, etc.):
  http://www.svi.nl/MicroscopicParameters

1. http://www.svi.nl/distributors/
2. http://www.svi.nl/FrontPage

- Information on the restoration parameters (signal to noise ratio, background, quality criteria, etc.) used by the deconvolution algorithms:
  `http://www.svi.nl/RestorationParameters`
- A step by step example on how to tune these parameters to achieve the desired restoration results:
  `http://www.svi.nl/DeconvolutionProcedure`
- Important issues regarding image acquisition and restoration (sampling, clipping, etc.):
  `http://www.svi.nl/ImportantFactors`
- Typical acquisition pitfalls (spherical aberration, undersampling, bleaching, etc.):
  `http://www.svi.nl/AcquisitionPitfalls`
- Information on recording beads to measure a PSF:
  `http://www.svi.nl/RecordingBeads`
- Tutorials and detailed information on using the different aspects of the Huygens software (restoration, visualization, analysis, programming, etc.):
  `http://www.svi.nl/Tutorials`
- Uploading images to SVI:
  `http://www.svi.nl/SendImagesToSvi`

License String Details

Detailed information about the installed license strings can be displayed via HELP→LICENSE. Select the license string of interest and click EXPLAIN LICENSE.

A Huygens license string consists of a set of substrings separated by dashes (-). These substrings describe e.g. the product, version number, options, etc. The checksum at the end of the string should match with all other substrings. A complete string looks like this:

```
HuEss-17.10-wcnps-d-tvAC-emnps-eom2015Dec31-
e7b7c623393d708e-{user@domain.com}-4fce0dbe86e8ca4344dd
```

Table 34.1 lists the building blocks from which this string is composed.

TABLE 34.1. The building blocks of the Huygens license string.

Substring	Description
Product	The product to which the license string applies. This can be `HuEss`, `HuPro`, `HuScript`, `HuCore, 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`' (multi-photon), '`s`' (STED), '`S`' (STED 3X), 'P' (SPIM), and '`4`' (4-Pi experimental microscopes).
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)

TABLE 34.1. The building blocks of the Huygens license string.

Substring	Description
Option flags	This is a set of characters that list the enabled optional modules. An overview of these modules is given in Table 34.2.
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 34.2. Identifiers for the optional modules.

Character	Module
A	Advanced Object Analyzer
b	Small file reader bundle
B	Complete file reader bundle
C	Colocalization Analyzer
f	PSF Distiller
g	Determines what GPU card can be used[a]. By default a card up to 1024 GPU cores and 2GB of VRAM can be used. 'g1' (Medium GPU option) allows the use of cards up to 3072 GPU cores and 4GB RAM, 'g2' (Large GPU option) allows the use of a card up to 8,192 GPU cores and 24GB RAM. The use of multiple cards are specified by the 'Server flag' (see Table 34.1)
G	Stitcher
H	Leica HyVolution LASX wizard
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
v	Surface Renderer visualization
x	Chromatic Aberration Corrector
c	RBNCC option for colocalization
T	Object Tracker
u	Crosstalk Corrector

a. www.svi.nl/GPUcards

Index

Numerics
4K 8
4K Monitors 8

A
Address iv, 207
Anchor 158
Animation
  frame count 86
  frame rate 86
  Movie Maker 89
  SFP 81
  Surface Renderer 86
AVI files 92

B
Background 23
Background estimators 171
Backprojected
  pinhole radius 17
  pinhole spacing 17
Batch Processor 39
Beads 33, 36
Black level 23
Bleaching 195
Bleaching correction 25
Bleaching Corrector 111
bleedthrough 107
Bounce 94
Brick layout 25
Brightfield images 203
Brightness 65

C
Center scene 86
Channels 27
Chromatic aberration 115
Chromatic Aberration Corrector 115
Chromatic shift 115
Client computers 181
Clipping 23, 194
CMLE 24
Colocalization 149, 167
  coefficients 169
  map 169
Color mode 81
Colors 199
Concurrent tasks 44
Concurrent usage 182
Contact SVI 207
Contrast 65
Contrast Editor 65
Convert 16
Correlation 149
Costes background estimation 172
Coverslip position 17, 200
Cropper 20
Cross correlation 116, 122
Crosstalk 107
Crosstalk Corrector 107
CUDA graphics cards 175
CW 46

D
Deconvolution
  templates 42
  wizard 15
Deconvolution Express 29
Deconvolution Wizard 15
Depletion beam 45
Distiller 33
Distributors 207
Drift 28

E
E-mail iv, 207
Emission
  transparency 80
  wavelength 17
Error function 65

Excitation
fill factor 17, 202
photon count 17
transparency 80
wavelength 17
Experiment preset 142, 160

F
Fax number iv, 207
File
Formats 16
Open 16
Save 28
Series 198
Fill factor 17
Filter Objects 161
Filter tracks 131
Floating license 182
Flow 134
Freeware mode 4
Fusion and Deconvolution Wizard 55
Fusion settings 56
Fusion Wizard 55

G
Gallery 95
gamma 65
Garbage volume 85, 138
Gaussian minimum 171
Geometrical distortion 193
Global color scheme 199
GPU acceleration 175
GPU Check list 176

H
hgst file 17
hgsv file 81, 86
Histogram 22, 144
Hot Pixel Remover 101
Huygens Core 183
Huygens Remote Manager (HRM) 183
Huygens Titan 9

I
Image Feeder 199
Imaging direction 17
Installation
Linux 3
Mac OS X 3
path 6
Uninstall 7
Windows 3
Intensity Tab 134
Iso-surface 83
Iteration mode 25
Iterations 24

K
Keyframes 90

L
Leica numbering 198
Lens refractive index 17
Library 11
License string 4, 208
Light Sheet Fluorescence Microscopy 51
Light Sheet Fusion & Deconvolution Wizard 55
Light-box 13
Linux 3
LSFM 51

M
Mac OS X 3
Manders 169
mask image 173
Mean squared displacement 132
Measure
Aspect Ratio 149
Colocalization 149, 167
Correlation 149
Length 147
Principal Axis 147
Sphericity 148
Medium refractive index 17
Memory 7
Microscope type 17
Microscopic parameters 17
Misalignment correction 115
mosaic 97
Mouse mode 140, 154
Movie Maker 89
MSD 132
Multi-channel images 27
Multiple Users 181

N
Number of iterations 24
Numerical aperture 17
Nyquist rate 194

O
Object Analyzer 135
Object size 79
Object Stabilizer 121
Object Tracker 127
Object Tracker Wizard 127
Objective quality 17
Optimized background estimation 172
Optimized search 172
Orthogonal Slicer 69
Overfill factor 202

P
Parameters 17
Pearson 169
Phone number iv, 207
Photodamage 194
Pinhole
radius 17
spacing 17

Pipe
Mode 156
Render 85, 140, 159
Point spread function 33, 197
Presets 92, 142, 160
Principal axis 147
Processor 7
PSF 197
Asymmetric 193
Distiller 33
Experimental 196

Q
QMLE 24
Quality factor 197
Quality threshold 24
Queue manager 183

R
Refractive index
Lens 17
Medium 17
Mismatch 193
Region of interest 152
Remote Desktop Protocol 182
Remote Display 181
Render
pipes 85, 140, 159
SFP 77
size 81, 86
Surface 83
resampling parameters 57
Resolution scaling in STED microscopy 46
ROI 152
ROI Deconvolution 30
rotation parameters 58

S
Saturation 23
Scatter parameters 54
scattering 54
Scheduling index activity 11
Searching images 12
Segmentation 137
seed-and-threshold 137
Selective Plane Illumination Microscopy 51
SFP 77
Shading 97
sharpness 30
Signal to noise ratio 24
Slide position 201
SNR 24
Spearma 169
Spherical aberration 193, 196, 200
Sphericity 148
SPIM 51
SPIM Fusion & Deconvolution Wizard 55
SPIM/Light Sheet Fusion Wizard 55
Stabilizer 121
Statistics 141
STED 45
deconvolution 49
parameter estimation 47
parameters 46, 51
Stitcher 97
Stitcher wizard 97
stitching patterns 98
Storyboard 91
Support 207
Surface Renderer 83
System ID 4, 209
System requirements 7

T
Task 43
Templates
Deconvolution 42
Microscopic 41
Visualization 81, 86
Threads per task 44
Threshold 79, 85, 137
Thumbnail messages 11
Tiff tile series 98
tiles 97
Time series 16
Timeline 92
Titan 9
Total internal reflection 194
Track Analyzer 131
Transformation 57
Transitions 91
Transparency 85
Transparency depth 86

U
Undersampling 194
Uninstall 7
unsupervised profiles 30
Updates 5, 6
Usage report 8

V
Velocity 133
Vignetting 97

W
Watch Directory 11
Watershed segmentation 139
Wavelength
Emission 17
Excitation 17
STED depletion laser 46
Web interface 183
Windows 3
Wizard 15, 35