



**BIO-INSPIRED
MATERIALS**

NATIONAL CENTER OF COMPETENCE
IN RESEARCH

Introduction to ImageJ

Session 4: 3D

Dimitri Vanhecke



adolphe merkle institute
excellence in pure and applied nanoscience

ETH

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich




ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE



**UNIVERSITÉ
DE GENÈVE**



SWISS NATIONAL SCIENCE FOUNDATION



nature methods



Article

<https://doi.org/10.1038/s41592-024-02580-4>

Segment Anything for Microscopy

Received: 21 August 2023

Accepted: 26 November 2024

Published online: 12 February 2025

 Check for updates

Anwai Archit¹, Luca Freckmann¹, Sushmita Nair¹, Nabeel Khalid^{1,2,3}, Paul Hilt⁶, Vikas Rajashekar², Marei Freitag¹, Carolin Teuber¹, Genevieve Buckley⁴, Sebastian von Haaren⁵, Sagnik Gupta¹, Andreas Dengel^{2,3}, Sheraz Ahmed² & **Constantin Pape¹**✉

Accurate segmentation of objects in microscopy images remains a bottleneck for many researchers despite the number of tools developed for this purpose. Here, we present Segment Anything for Microscopy (μ SAM), a tool for segmentation and tracking in multidimensional microscopy data. It is based on Segment Anything, a vision foundation model for image segmentation. We extend it by fine-tuning generalist models for light and electron microscopy that clearly improve segmentation quality for a wide range of imaging conditions. We also implement interactive and automatic segmentation in a napari plugin that can speed up diverse segmentation tasks and provides a unified solution for microscopy annotation across different microscopy modalities. Our work constitutes the application of vision foundation models in microscopy, laying the groundwork for solving image analysis tasks in this domain with a small set of powerful deep learning models.



Going digital – what is a digital image?

A digital image is an ordered rectangular array (or grid) of **integers (numbers: 0,1,2,3...)**. Each element (=number) in the grid is also known as a picture element or 'Pixel'

Spectrum

1 dimensional array

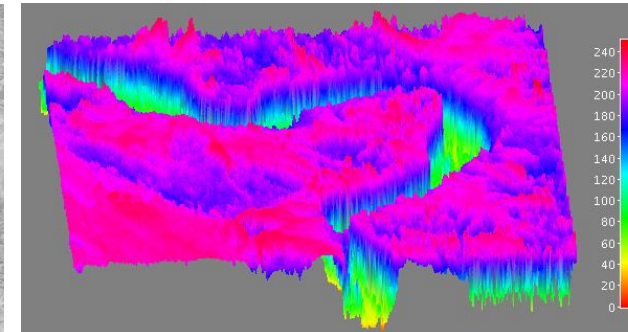
208
209
213
198
191
191
195
184
190
188
191
188
200



Image


2 dimensional array

82	72	78	86	65	41
157	144	167	188	201	191
185	191	195	188	188	191
193	195	195	191	189	171
173	170	181	192	194	191
210	214	206	202	203	201
237	224	221	230	232	221
183	180	190	188	192	181
178	170	159	187	195	181
167	164	170	186	192	181
159	162	164	184	170	161
180	172	165	172	185	171
193	180	196	195	185	171
167	184	182	183	180	171
195	191	182	189	195	181
183	188	184	183	174	161
161	166	165	170	169	161




Stacks

3D array
(= volume stack or video/Timelapse)




Types of 3D data



Filters, point operators, ... and stacks

Upon running a function over a stack, you will often get a question:



Hence, all

- Filters
- Bandpass filters
- Point operators
- Binary functions
- etc...

Are also valid for stacks

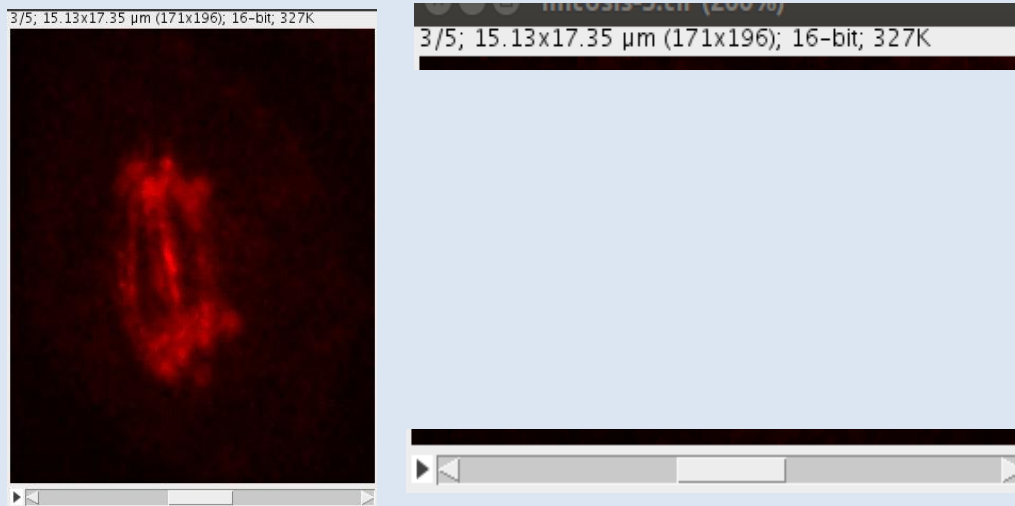


Sobel filter on RGB Lena

Stacks

Prerequisites

- All the slices in a stack must be the same size (X,Y) and bit depth.
- The slice thickness is considered constant (Z)




Type of stacks

1. 2D images with **encoded Z information** (e.g. AFM)
2. **Channels (or layers)** are multiple images stored within one file. Typically, they contain different color absorption functions of the same object
3. **Z layers** are images recorded at different depth positions through the object
4. **Time lapse** are images recorded at a different time point
5. **Custom multi-dimensional datasets** ($n > 3$). E.g. Hyperstacks, Raman maps and hyperspectral maps

Stacks

3D array
(= volume stack or video)




Height maps

Gwyddion


Open SPM (AFM, SNOM/NSOM, STM, MFM, ...) data analysis software

e.g. AFM height maps


- 2D image
- Pixel value = height = height map



7.87 nm



3D Map (XY view, transformed height map) XZ view (10X)





10 nm

2. Channels

Pseudo-color


= a single channel (grayscale) equipped with a LUT



Count: 247200 Min: 0
Mean: 125.901 Max: 255
StdDev: 73.247 Mode: 201 (3164)

RGB images (24 bit=3x8bit)



3 layers, reflecting the natural red, green and blue colors (or HSL, CMYK, HSV, ...)



Count: 64000 Min: 0
Mean: 71.655 Max: 248
StdDev: 58.251 Mode: 8 (2137)



Composite images (flexible: e.g. 5x16bit)

n layers, separated. For example LSM multi channel data



Count: 1310720 Min: 472
Mean: 1038.686 Max: 8583
StdDev: 730.173 Mode: 567.051 (188628)
Bins: 256 Bin Width: 31.684


2. Channels: Pseudo-color (Lookup tables (LUTs))




- Image = 2D array of numbers
- "Grayscale" Lookup table = a means to giving a graphical meaning to these numbers.
- But "Grayscale" is just one of these Lookup tables!
- Image > Look-up table (not for RGB images)

Look-up tables = dictionary in Python


MPL-inferno




Physics




Edges





brgbcmyw



2. Channels: Lookup tables (LUTs)



2. Channels: composite images



Composite image

Channels tool

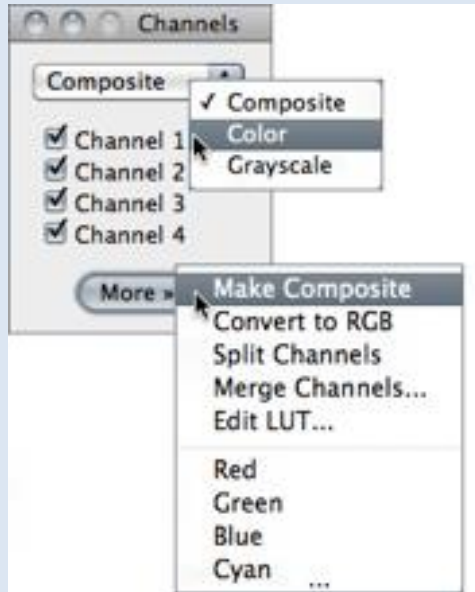


Image > Color > Channels tool

- Composite:** overlaying the layers of choice (also for RGB images)
Color: showing only one layer, with LUT. Change the LUT of the selected layer
Grayscale: showing only one layer, in grayscale LUT
(Clicking on the channel selector = use the channel scrollbar below the image)
- Make composite:** splits the color image in its layers
Convert to RGB: joins the layers into a 2D RGB image (you will end up with 1 window)
Split channels: makes n windows of each channel
Merge channels: Tool to put n single channels together into a composite stack

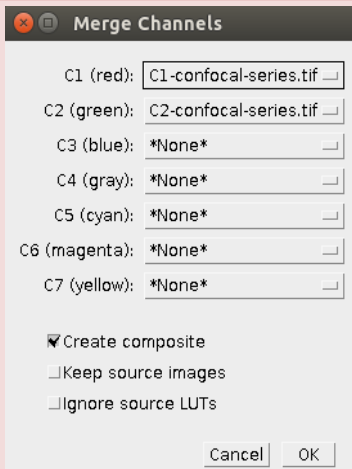


Image > Color > Merge Channels

Combines n images into a composite image

- Prerequisite: all images have the same size (width, height and bitdepth)
- Choose the LUT (color)
- Once merged: check the "Arrange" menu entry (Image > color > Arrange...)

Channels: split, arrange, and merge

EXERCISE

Open Example 1

Convert to Composite

Convert a color image to a composite image (Image > color > channels tool: More > make composite)

Split a composite dataset in its grayscale components

Split the three channels (Channels tool: More > split channels)

Optional: change the LUT of each of the grayscale components

Change LUTs if required (Image Lookup tables)

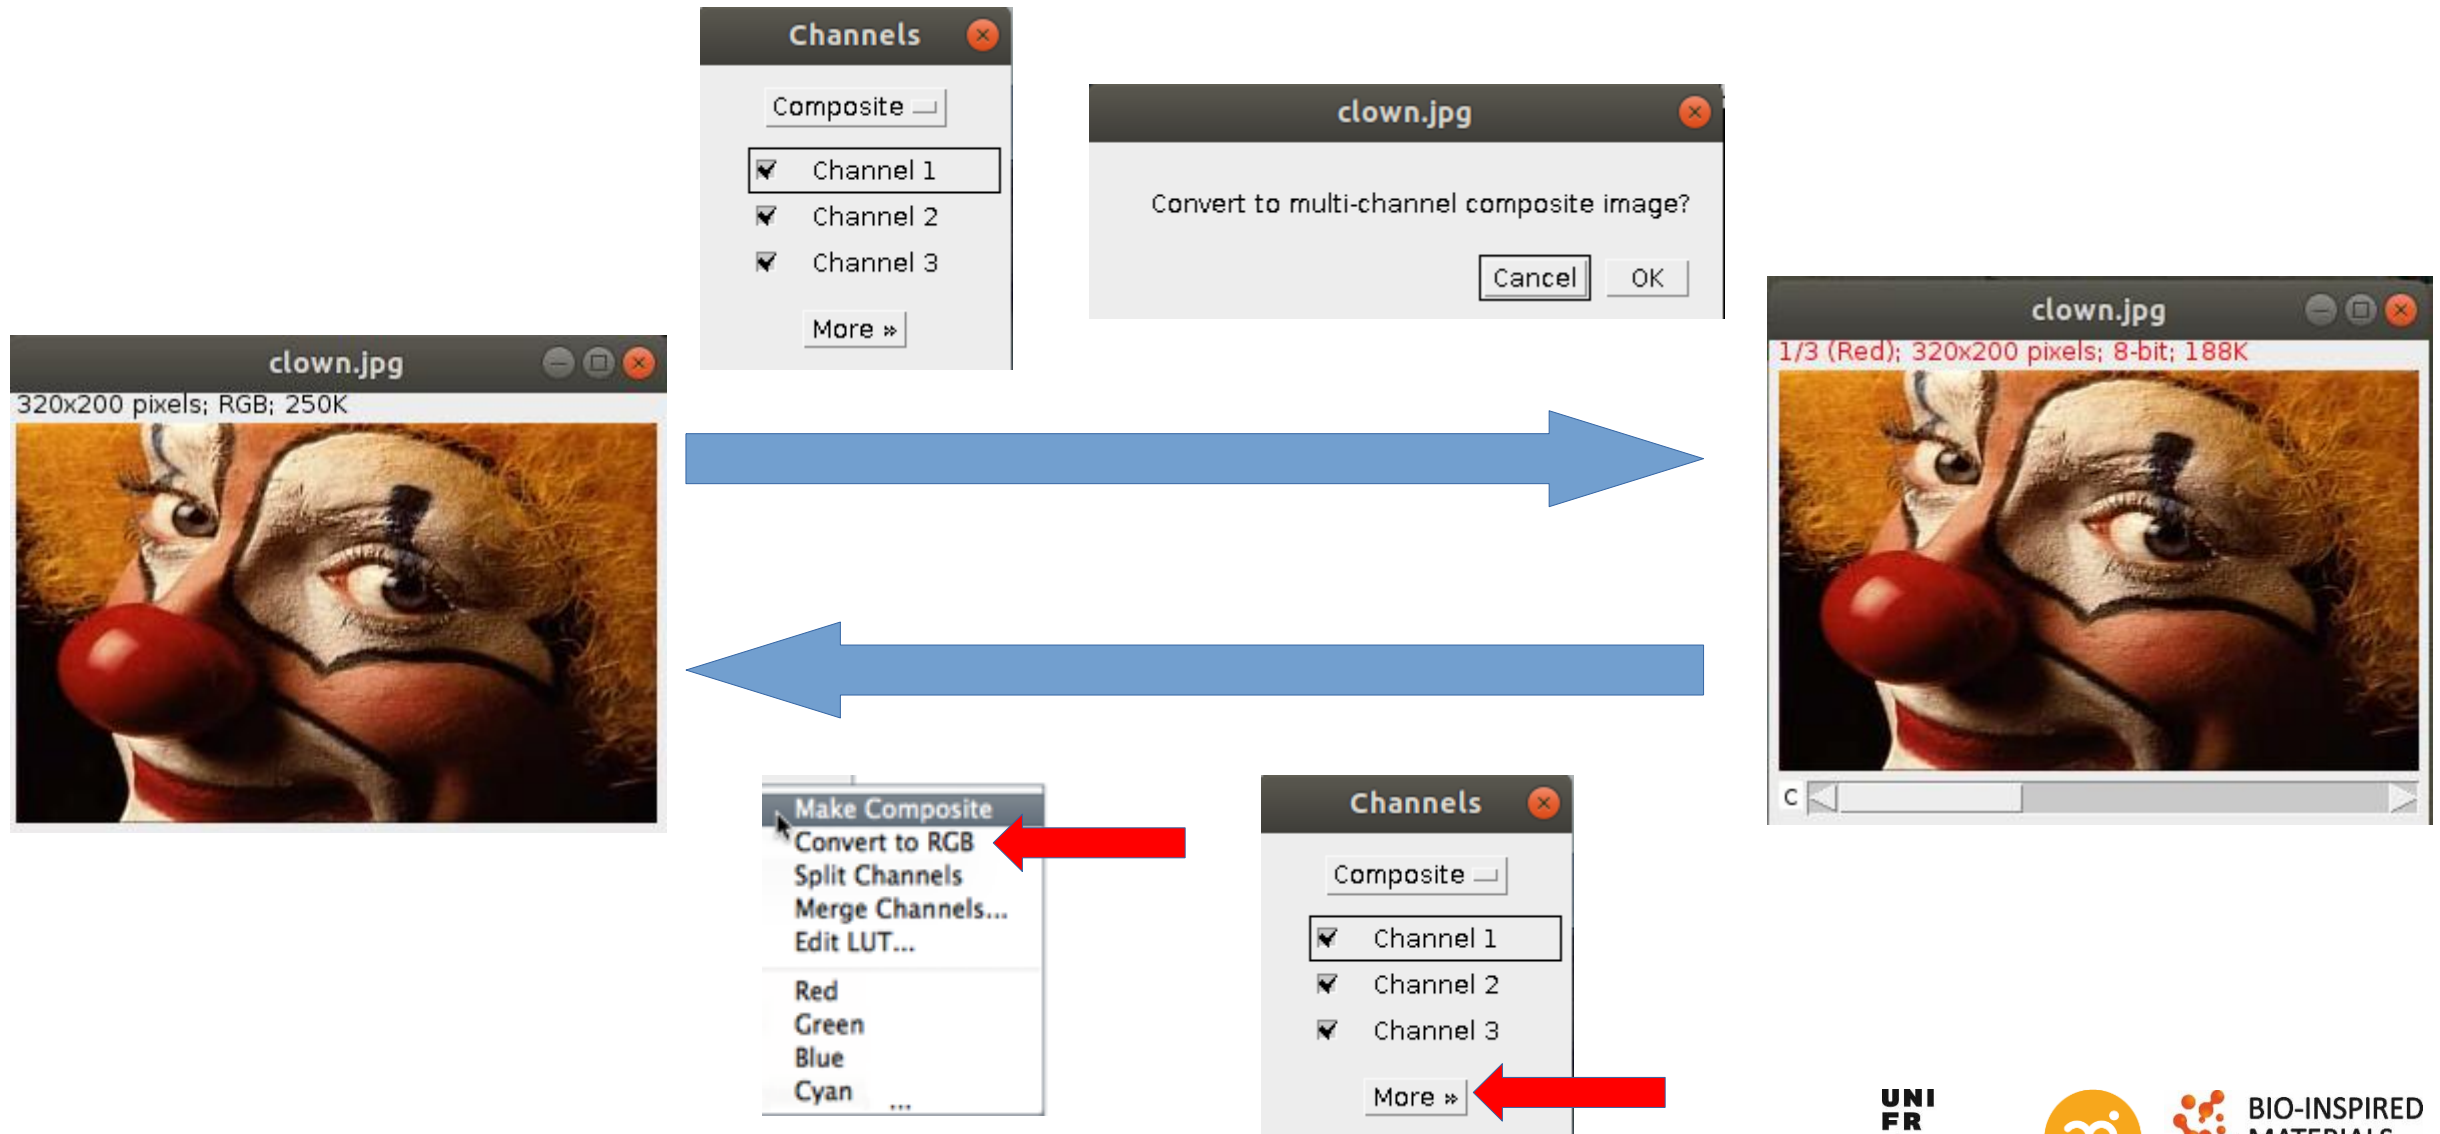
Merge channels

Merge the channels again to an RGB image (Image > color > Merge channels OR Channels tool: more > merge channels)

Change the order of the grayscale channels in the composite dataset

Arrange: Change the order of the layers in the stack (Image > color > Arrange Channels...)

Channels tool: example RGB image



Channels tool: example RGB image

The image illustrates the workflow of the Channels tool in GIMP for splitting and merging an RGB image. It consists of several panels and dialog boxes:

- clown.jpg**: The original image, showing a clown's face. The status bar indicates "1/3 (Red); 320x200 pixels; 8-bit; 188K".
- Channels**: A dialog box with a "Composite" dropdown and three checked channels: "Channel 1", "Channel 2", and "Channel 3". A "More »" button is highlighted with a red arrow.
- Context Menu**: A menu with options: "Make Composite", "Convert to RGB", "Split Channels" (highlighted with a red arrow), "Merge Channels...", and "Edit LUT...". Below are color options: "Red", "Green", "Blue", "Cyan", and "...".
- C1-clown.jpg (RGB)**: The image with only the Red channel visible. The status bar indicates "320x200 pixels; 8-bit; 62K".
- C2-clown.jpg (RGB)**: The image with only the Green channel visible. The status bar indicates "320x200 pixels; 8-bit; 62K".
- C3-clown.jpg (RGB)**: The image with only the Blue channel visible. The status bar indicates "320x200 pixels; 8-bit; 62K".
- Merge Channels**: A dialog box with input fields for "C1 (red)", "C2 (green)", and "C3 (blue)", each set to its respective channel name. Other channels (C4-C7) are set to "*None*". There are checkboxes for "Create composite" (checked), "Keep source images", and "Ignore source LUTs". "Cancel" and "OK" buttons are at the bottom.
- Channels**: A second dialog box, identical to the first, with "Channel 1", "Channel 2", and "Channel 3" checked. The "More »" button is highlighted with a red arrow.
- Context Menu**: A second menu, identical to the first, with "Merge Channels..." highlighted by a red arrow.

Blue arrows indicate the flow of the process: from the original image to the split channels, and from the split channels back to the merged composite image.

3. Z-stacks


1. One file including the entire Z-stack

Native: TIFF

Non-Native

- lsm (Zeiss): Use LSM toolbox
- lif (Leica): Use Bio-Formats plugin
- ...

2. Sequence = a number of 2D images (same XY size, same bitdepth) in a single folder



File > Import > Image Sequence

Enter or browse the folder path

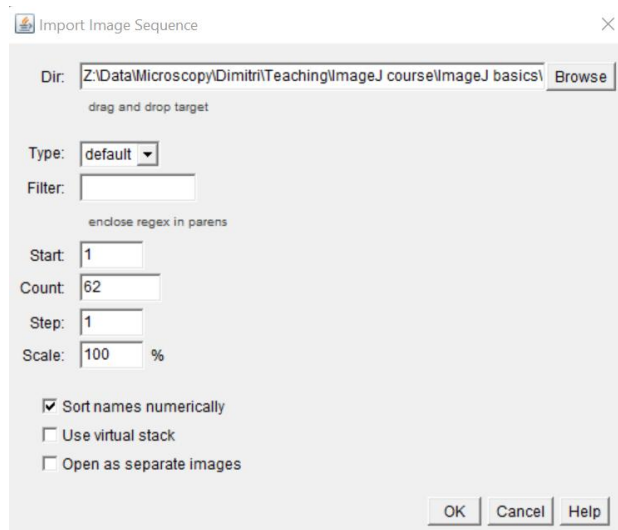
Possibility to reduce the stack

Import options


Opening sequences

EXERCISE

Open Example 2 (the folder) and import the sequence



- **File > Import > Image Sequence**
- Locate the folder
- (you do not see the actual files in the folder)
- Regex Filter: allows filename filtering (e.g. tif will only include files that have tif in the filename)



All images must have the same size! (X, Y and bitdepth!)

Watch out for **OS generated thumbnail files**


Possibility to open as virtual stack

Opening sequences

EXERCISE

Open Example 2 (the folder) and import the sequence

1/124 (Data000); 128x107 pixels; 8-bit; 1.6MB



Stack of 124 Slices, now looking at slice 1
X = 128
Y = 107
Z = 124

Works exactly the same if you would have opened a multi-image file (eg. Tiff)

What is the difference between TIF and TIFF?

Move through the stack

Operations on Z-stacks

Image > stacks

Add Slice, Delete Slice, Next Slice, Previous Slice, Set Slice...

Image > Stacks > Make montage

Produces a single grid-image containing the individual images that compose stacks and hyperstacks

Image > stacks > Stacks to Images

Releases the n images in the one stack window into n windows (watch out with large stacks!)

Image > stacks > Images to Stacks

Takes all open Images and puts them into 1 stack. Regex filter possible.

Image > Stacks > Z project

Projects the entire stack onto a 2D image

Image > Stacks > 3D project


Maximum intensity projection of the stack

Image > Stack > Tools

Combine, Reduce, Make substack, ...

Image > Stacks > Tools > Grouped Z project...Output: a new stack, but with each slice the average/max/sum or each group

Z-Stacks: Extended depth of field

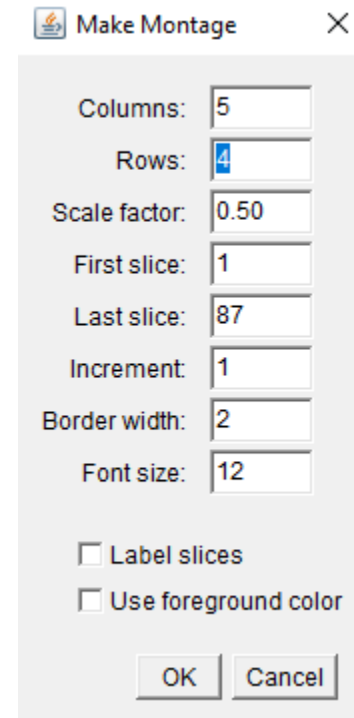



Montage tool

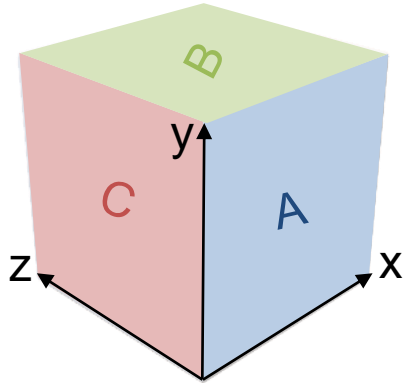
EXERCISE

Try out the tools in the Images > Stack menu and with Example 2


Open a stack, then:
Image > stack > Make montage...



Z-Stacks: Reslice (orthogonal rotation)



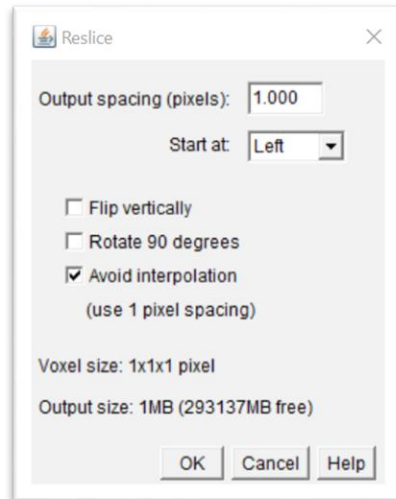
(Plane A) XY view



X=128
Y=107
Z=124

Bold: on-screen depicted dimensions

Image > stack > Reslice...



Output spacing (pixels): 1,000
Start at: Top

Reslice from top/bottom
(Plane B) XZ view

X=128
Y=124
Z=107



Reslice from left/right
(Plane C) YZ view

X=107
Y=124
Z=128



Output spacing (pixels): 1,000
Start at: Left

Other tools:

Radial reslice

orthogonal reconstructions of a stack by rotating a line ROI around one end of its center. Useful for data with rotational symmetry

Dynamic reslice

Creates an arbitrary cross section along a user-defined line


4. Hyperstacks

Hyperstacks are multidimensional images, extending image stacks to four (4D) or five (5D) dimensions:


- x (width),
- y (height),
- z (slices),
- c (channels or wavelengths)
- t (time frames)

Hyperstacks are displayed in a window with 2 or 3 labelled scrollbars. Similarly to the scrollbar in stacks, including a play/pause icon.


3D stack (z=5)



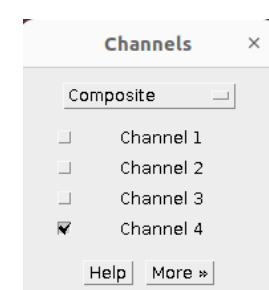
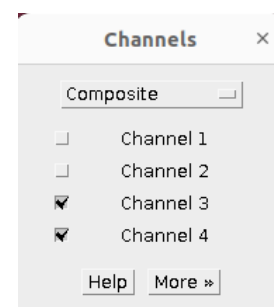
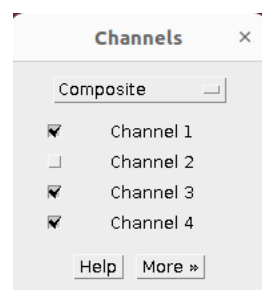
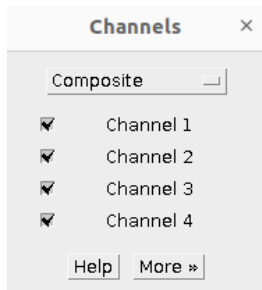




4D stack (z=5, C=2)



5D stack (z=5, C=2, t=51)



Hyperstacks




Videos/timelapse

Out of the box, ImageJ has limited support (no codecs, no audio). However, it can open/close uncompressed AVI formats.

Videos/timelapse

Can be understood as a 3D stack where the third dimension is not spatial but temporal



28/36; 156x124 pixels; 8-bit; 680K

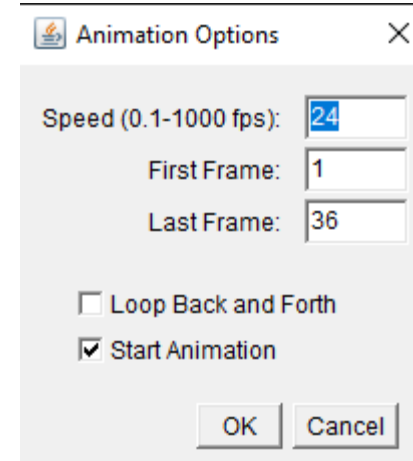
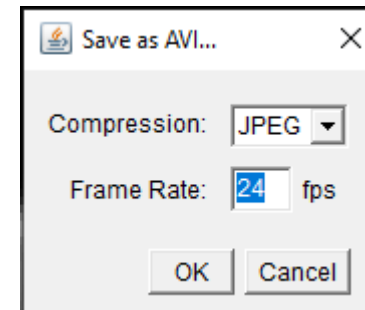


Image > Stacks >
Animation >
Animation options

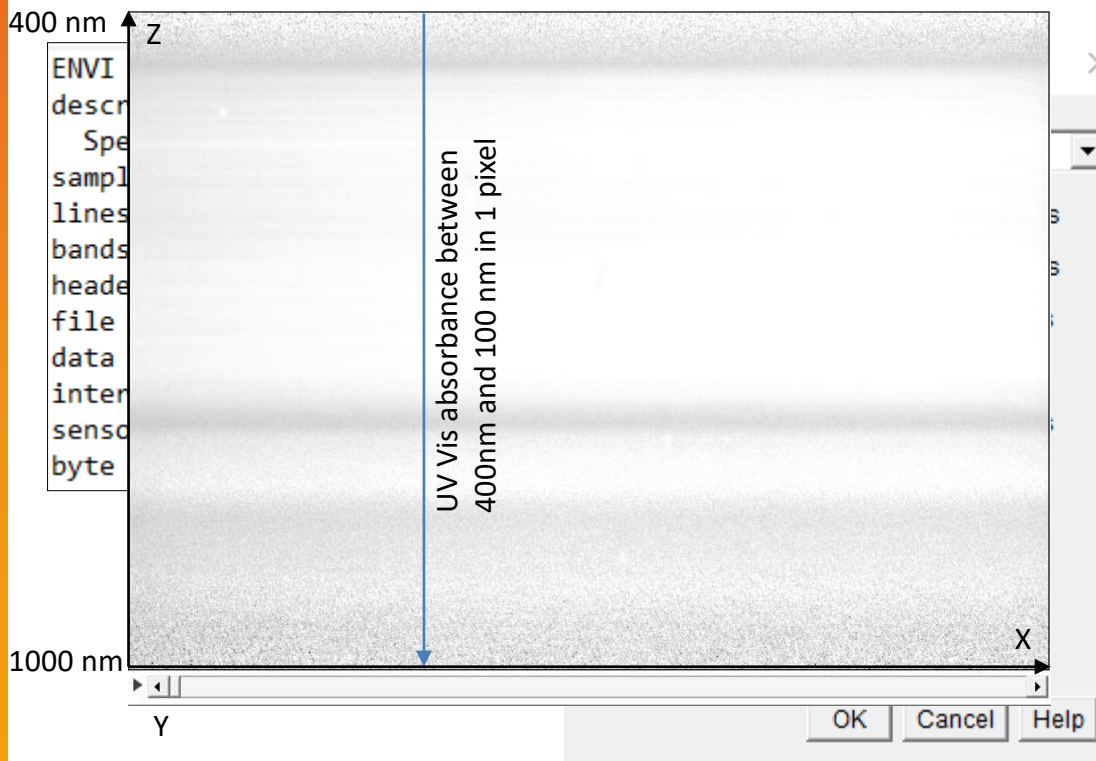


File > Save as... > Avi...

5. Custom multi-dimensional datasets





FIJI Raw import does not interpret your data, it just reads it (remember Lecture 1).

Example: Cytoviva dataset



XY: image dimension
Z: spectral dimension

Image > stack > Plot Z-axis profile...


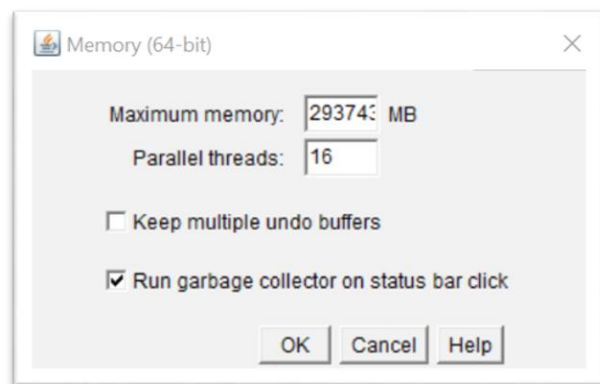


6. Virtual stacks

- Virtual stacks are disk resident (as opposed to RAM resident) datasets
- The only way to load image sequences that do not fit in your RAM.

1. Virtual stacks are read-only, so changes made to the pixel data are not saved when you switch to a different slice
2. Commands like Crop [X] may create a RAM issue since any stack generated from commands that do not generate virtual stacks will be RAM resident.

Edit > options > memory & threads will allow you to change the RAM allocated




RAM for all other software

Total RAM of PC

Memory	3.7 GiB
Processor	Intel® Core™ i7 CPU M 620 @ 2.67GHz × 4
Graphics	NVS 5100M/PCIe/SSE2
OS type	64-bit

RAM for ImageJ





Measuring in 3D

Note on non-isometric data (LSM, FIB, ...)

Prerequisites:

- binary images (see lecture III of this series)
- Proper axial and spatial resolution set (Image > Properties)

When the axial resolution (in Z) is not the same as the spatial resolution (in XY):
Image > Properties



Original data

A549_PCL200.tif

Channels (c): 1

Slices (z): 49

Frames (t): 1

Note: $c \cdot z \cdot t$ must equal 49

Pixel width: 0.3603982 micron

Pixel height: 0.3603982 -

Voxel depth: 0.6059463 -

Frame interval: 0 sec

Origin (pixels): 0,0,0

Invert Y coordinates

Global

OK Cancel

After segmentation (from iLastik)

A549_PCL200-t0-channel0_Simple Se...

Channels (c): 1

Slices (z): 49

Frames (t): 1

Note: $c \cdot z \cdot t$ must equal 49

Pixel width: 1.0000 pixel

Pixel height: 1.0000 -

Voxel depth: 1.0000 -

Frame interval: 0 sec

Origin (pixels): 0,0,0

Invert Y coordinates

Global

OK Cancel

3D Objects counter

Analyze > 3D OC options

Allows to set the Measurements that will be performed

Turn off



3D-OC Set Measurements

Parameters to calculate:

<input checked="" type="checkbox"/> Volume	<input checked="" type="checkbox"/> Surface
<input checked="" type="checkbox"/> Nb of Obj. voxels	<input checked="" type="checkbox"/> Nb of Surf. voxels
<input checked="" type="checkbox"/> Integrated Density	<input checked="" type="checkbox"/> Mean Gray Value
<input checked="" type="checkbox"/> Std Dev Gray Value	<input checked="" type="checkbox"/> Median Gray Value
<input checked="" type="checkbox"/> Minimum Gray Value	<input checked="" type="checkbox"/> Maximum Gray Value
<input checked="" type="checkbox"/> Centroid	<input checked="" type="checkbox"/> Mean distance to surface
<input checked="" type="checkbox"/> Std Dev distance to surface	<input checked="" type="checkbox"/> Median distance to surface
<input checked="" type="checkbox"/> Centre of mass	<input checked="" type="checkbox"/> Bounding box

Image parameters:

Close original images while processing (saves memory)

Show masked image (redirection required)

Maps' parameters:

Dots size

Font size

Show numbers

White numbers

ResultsTable parameters:

Store results within a table named after the image (macro friendly)

Redirect to:

OK Cancel

3D Objects counter

Analyze > 3D Objects Counter


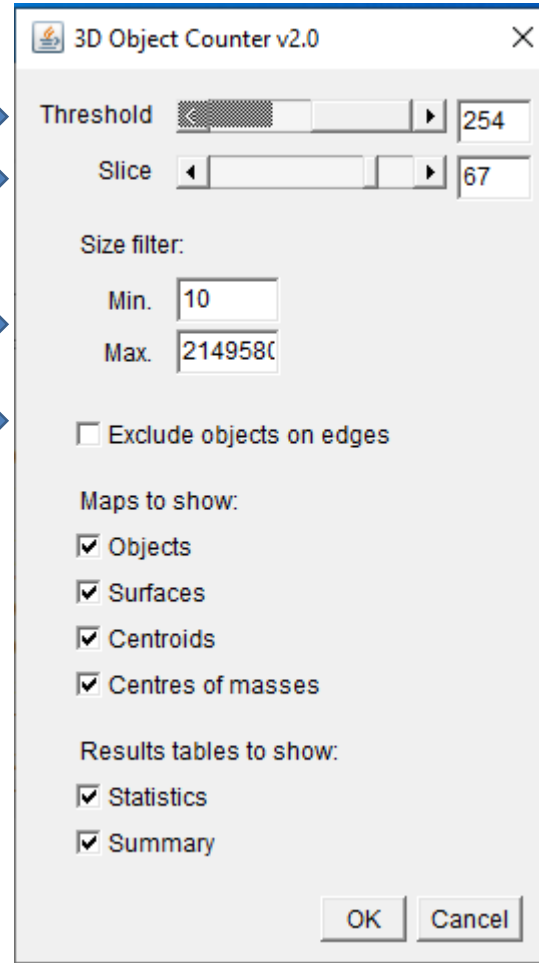
Similar to 'Measure particles', but: Threshold is asked

(Poor) thresholding attempt →
Move through stack →

Eliminate noise or large clumps →



Eliminate partial objects →

Output



3D Objects counter: Output

Surfaces



N: 21495808	Min: 0
Mean: 0.0186	Max: 17
StdDev: 0.447	Mode: 0 (21445801)
Value: ---	Count: ---

Check the Look up table

3D Objects counter: Output

Statistics for A549_PCL200-t0-channel0_Simple Segmentation Stage 2-1.tiff

File Edit Font

Volume (micron ³)	Surface (micron ²)	Nb of obj. voxels	Nb of surf. voxels	IntDen	Mean	StdDev	Median	Min	Max	X	Y	Z	Mean dist. to surf. (micron)	SD
1408.099	1201.742	17891	3974	4562205	255	0	255	255	255	382.541	60.942	35.786	7.534	2.4
2118.878	1497.572	26922	6070	6865110	255	0	255	255	255	197.877	89.038	33.429	9.103	3.3
647.894	585.138	8232	2013	2099160	255	0	255	255	255	155.839	63.630	34.102	5.357	1.3
643.565	558.641	8177	2092	2085135	255	0	255	255	255	172.481	142.474	33.814	5.120	10.9
697.950	649.447	8868	2294	2261340	255	0	255	255	255	151.261	37.186	36.641	5.656	1.5
1707.412	1195.269	21694	4806	5531970	255	0	255	255	255	250.222	80.860	36.413	7.609	2.2
747.534	633.551	9498	2437	2421990	255	0	255	255	255	195.318	34.086	36.335	5.452	1.2
1255.649	993.930	15954	3658	4068270	255	0	255	255	255	96.251	102.124	37.521	6.826	2.0
682.367	581.667	8670	2050	2210850	255	0	255	255	255	148.327	164.184	36.652	5.336	1.0
1001.277	758.804	12722	3025	3244110	255	0	255	255	255	198.978	170.228	37.802	6.027	1.4
598.783	537.068	7608	1945	1940040	255	0	255	255	255	162.643	195.234	39.588	5.080	1.1
644.431	582.112	8188	2171	2087940	255	0	255	255	255	234.675	177.968	41.452	5.394	1.5

3D Objects counter

EXERCISE

Open Example 7 and calculate the volume of the objects using the 3D object counter.

1. Check calibration
2. Do the analysis
3. Change the settings and repeat

Image > Properties... (for 3D spatial and axial settings)
Analysis > 3D object counter
Analysis > 3D OC settings

3D suite (plugin)



Help > Update ... >
Plugins > 3D suite

3D ImageJ Suite

<https://sites.imagej.net/Tboudier/>

Analysis ▶

Binary ▶

Filters ▶

Relationship ▶

Segmentation ▶

3D Manager V4 (testing)

3D Manager V4 Macros

3D Manager

3D Manager Options

Spatial ▶

Tools ▶

See next slide

(Morphological) filters in 3D

Local Linear filters in 3D

Measuring distances (e.g. border to border)

Binary segmentation tools (e.g. 3D watershed)

(experimental)

(scripting)

ROI 3D manager (see next)

3D suite (plugin)

▶	3D Intensity Measure
▶	3D Centroid
▶	3D Volume
▶	3D Surface
▶	3D Distance Contour
	3D Feret
	3D Compactness
	3D Ellipsoid measure
	3D RDAR
▶	3D Mesh Measure (slow)
▶	3D Ellipsoid Fitting
□	3D Convex Hull (slow)

Input

Raw data and binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Binary mask

Output

Intensity stats of each object

Position of centroid of each object (X,Y,Z)

Volume of each object

Surface of each object

Distance stats between center and shell

Caliper distances in 3D and ortho-planes

Sphericity and 3D compactness

Goodness of fit measurements


Ellipsoid: how much is sticking out

Fitting measures to ellipsoid

3D convex hull

3D suite (plugin)

ROI3D manager




1. 3D segment (use binary Data!!)

you get a new window with your objects in different shades

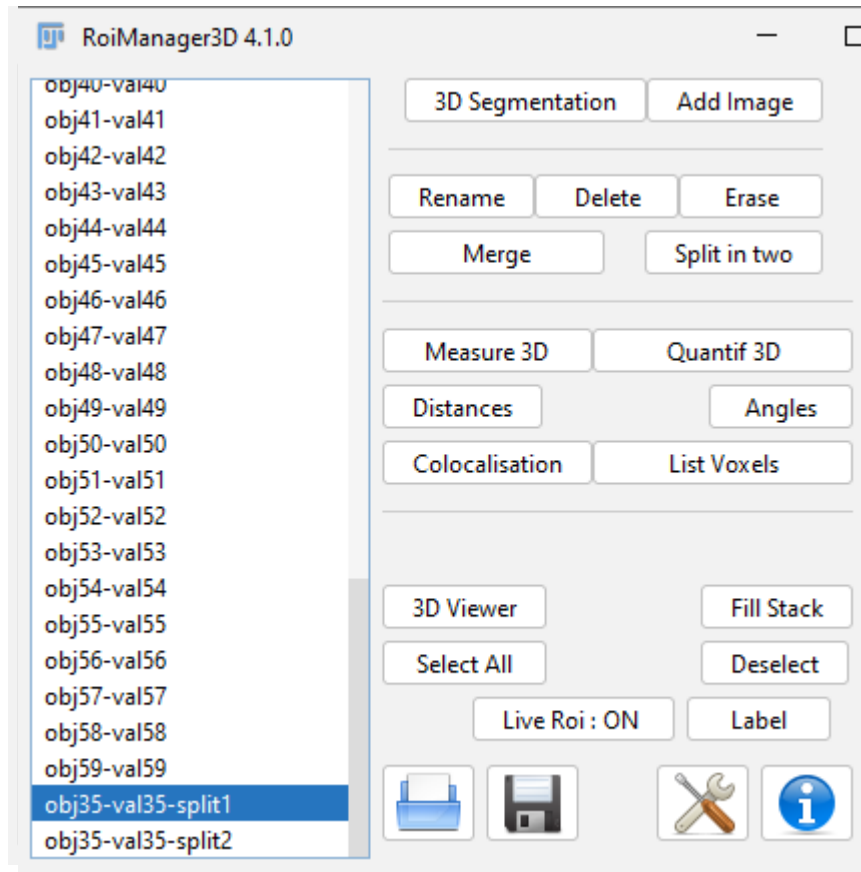
2. Add an image

this adds the objects

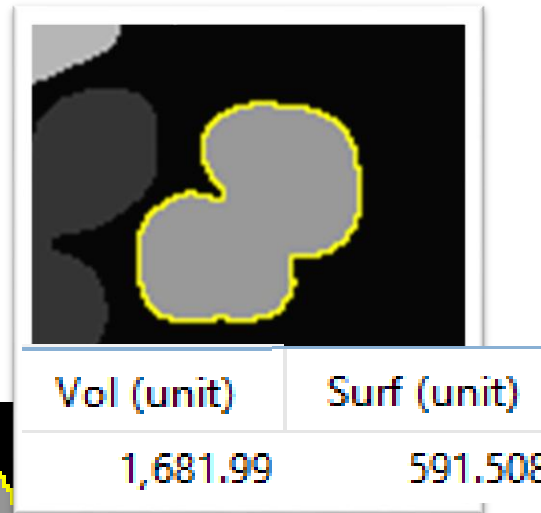
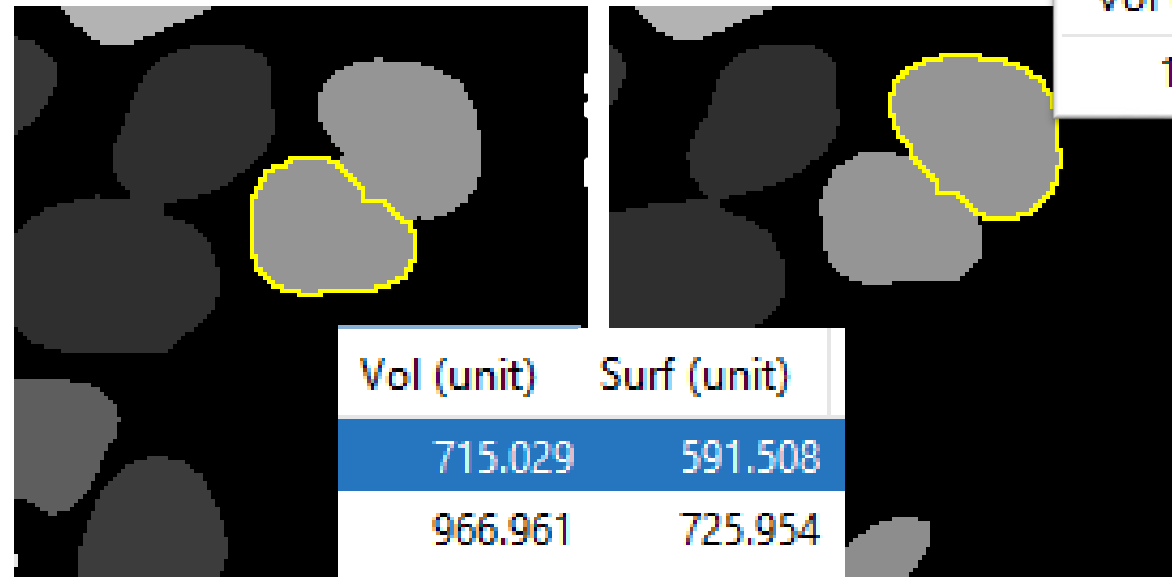


3D suite (plugin)

ROI3D manager



1. 3D segment (use binary Data!!)
2. Add an image
3. Click "Live ROI: OFF" (makes it "ON")
4. From the list, select obj35-val35
5. Then click "split in two"



3D suite (plugin)

EXERCISE

Open Example 7 and calculate the volume of the objects using the 3D manager of 3D suite.
Try to split some objects in the 3D suite

Image > Properties... (for 3D spatial and axial settings)

Analysis > 3D object counter (and 3D OC settings)

Plugins > 3D suite > 3D manager



- Segmentation
- Add image

Visualization of 3D data




Visualizing 3D data

1. 2D depictions
2. Renderings
 1. Surface rendering
 2. Volume rendering




Visualizing 3D data: Depth encoded

AFM image




(= depth coded, Fire LUT)

AFM as 3D stack



3D stack



Depth coded 3D stack




Image > Hyperstacks >
temporal color-coded

Visualizing 3D data: Orthogonal views and depth coding

EXERCISE

Open Example 3.






Try both:

Depth-encoded Color



Image > Hyperstacks > Temporal color-code
/ choose a LUT (e.g. Grays)

Orthogonal views

Image > stack > orthogonal views




Visualizing 3D data: Orthogonal views



Orthogonal views

The intersection point of the three views follows the location of the mouse click and can be controlled by clicking and dragging in either the XY, XZ or YZ view.

XY and XZ coordinates are displayed in the title of the projection panels. The mouse wheel changes the screen plane in all three views.



How to get rid of the marker lines?

Image > Overlay > hide overlay (or remove overlay)

3D rendering

Note: renderings require **interpretation** by the user. Hence, they are the convolution of the raw scientific data and the feature the user would like to see.

1. Surface rendering

= binary threshold-based

2. Volume rendering

= percentage threshold-based


Never publish only renderings.
Always provide the raw data (i.e.
orthogonal views)

Surface rendering: Isosurfaces

Isosurface


A three-dimensional analogue of an isoline. It is a surface that represents points of a constant value within a volume.

- Step 1: Creating an isoline by thresholding
- Step 2: voxels to mesh by marching cubes
- Step 3: Mesh to rendering through shaders




Isosurfaces: Step 1: Thresholding the voxels

Original



Binary



0	30	31	30	31	32	31	30	30	30	40	42	42	43	81	82	81	82	101	111
5	38	37	36	37	37	38	37	39	39	40	45	54	66	74	92	99	99	120	124
5	35	35	34	36	37	38	35	38	41	43	47	63	74	72	97	108	109	123	116
7	34	36	35	35	39	38	36	38	43	48	57	67	77	80	86	118	113	134	118
7	35	33	34	38	37	38	38	39	39	50	59	71	89	102	105	105	112	127	131
8	34	35	35	37	40	37	40	40	43	54	74	77	89	102	109	110	105	115	144
6	34	36	37	35	39	40	41	42	49	57	72	82	87	109	124	123	114	114	121
4	35	35	35	37	40	44	43	47	57	64	75	82	88	102	114	128	124	129	137
6	36	38	37	38	43	42	46	47	60	77	88	96	95	108	124	128	115	138	131
5	34	36	38	40	43	42	47	52	75	90	92	102	105	113	123	118	114	127	146
7	37	39	42	44	45	50	61	77	88	91	97	104	112	112	122	123	120	122	131
7	37	39	40	48	56	58	68	72	89	102	107	103	121	113	133	128	130	135	111
7	38	41	45	53	68	78	74	84	97	104	118	114	125	120	124	133	138	151	137
0	41	46	53	57	72	84	80	83	120	113	117	125	122	129	134	127	137	131	138
4	47	52	61	63	89	103	92	100	110	120	119	125	128	132	128	127	143	136	134
1	54	64	68	87	98	102	109	111	105	121	122	131	132	127	124	129	127	141	141
8	56	70	75	89	97	104	114	111	112	129	129	131	119	126	158	149	131	148	128
4	86	97	90	101	109	110	117	109	109	127	128	129	127	122	153	155	132	139	125
7	85	108	100	108	115	112	135	127	124	129	131	140	141	134	141	154	144	126	126
6	80	100	112	111	112	116	126	147	142	147	144	148	147	147	141	150	140	127	124





Threshold = 83

Edge only = isoline



A threshold is calculated

- Pixel value > threshold, the voxel is considered to contain the signal (=object).
- Pixel value < threshold, the voxel is considered not to contain the signal (=background).
- This classification system is binary; it defines each voxel as containing either 100% or 0% of the signal
- Once classified, a surface is defined as the boundary between the pixels (=isoline)

Isosurfaces: Step 2: Isoline/Voxel to mesh conversion





Mesh





Marching squares

Isosurfaces: Step 2: Voxel to mesh conversion

Marching squares




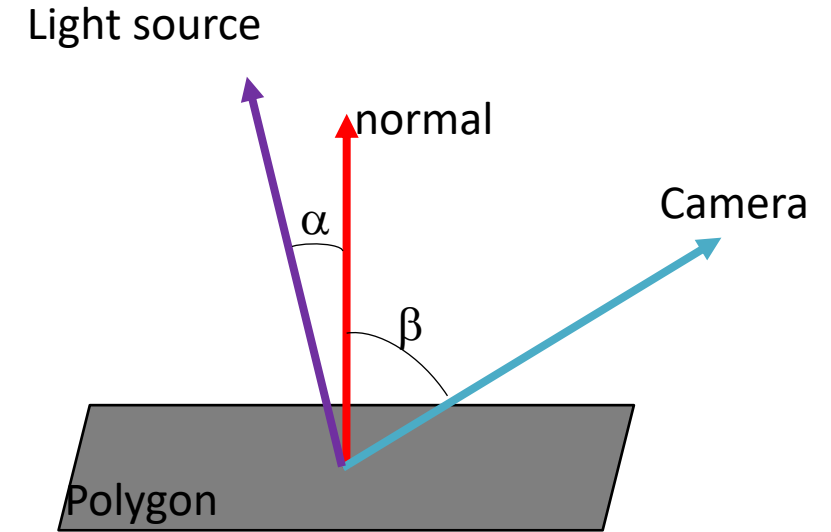


Marching cubes



This is called a polygon

Intensities -> Binary -> 64 predefined values / marching cubes

Isosurfaces: Step 3: Reflection and intensity



Normal: vector perpendicular to the polygon

The normal defines the color (or shade) of the polygon


Polygons rendered without shader (flat)

α = angle between light and normal
 β = angle between camera and normal

The smaller the difference between the angles
The brighter the polygon



Isosurfaces: Step 3: Illumination

No shader





Gouraud shading

Bilinear interpolation of the intensities (color) between two normals




Phong shading

Barycentric interpolation of the normals themselves



Modern hardware: use Phong (better than Gouraud, but a bit more intensive computing)


Isosurface: towards ray tracing



① Sphere equation: $(\vec{p} - \vec{c}) \cdot (\vec{p} - \vec{c}) = r^2$ Intersection:

Ray equation: $\vec{r}(t) = \vec{o} + t\vec{d}$ $(\vec{o} + t\vec{d} - \vec{c}) \cdot (\vec{o} + t\vec{d} - \vec{c}) = r^2$

$$t^2 (\vec{d} \cdot \vec{d}) + 2(\vec{o} - \vec{c}) \cdot t\vec{d} + (\vec{o} - \vec{c}) \cdot (\vec{o} - \vec{c}) - r^2 = 0$$



② Illumination Equation (Blinn-Phong) with recursive Transmitted and Reflected Intensity:


$$I = k_a I_a + I_i \left(k_d (\vec{L} \cdot \vec{N}) + k_s (\vec{V} \cdot \vec{R})^n \right) + \underbrace{k_t I_t + k_r I_r}_{\text{recursion}}$$

③ Snell's law: $\frac{\sin \theta_1}{\sin \theta_2} = \frac{v_1}{v_2} = \frac{n_2}{n_1}$ $n_{air} \sin \theta_i = n_{glass} \sin \theta_t$ refraction coefficients:
 $n_{air} = 1, n_{glass} = 1.5$

④ Area Light Simulation: $I_{light} = \frac{\#(\text{visible shadow rays})}{\#(\text{all shadow rays})}$

The more bounces, the more realistic the image becomes

Isosurface: Example




ImageJ 3D viewer

Isosurface and (very basic) volume renderer
Good quality, but limited
Buggy (in my view)

But:
export as STL, wavefront ==> 3D printer

And volume calculation



Commercial renderer

Avizo/Amira/Imaris
Very flexible, commercial software
Good quality, extensive renderer

Available through SciIT
(BioNano workstation)



Open source ray-tracer

Blender 2.82 cycles renderer
Realistic rendering possible
Slow

Free to download

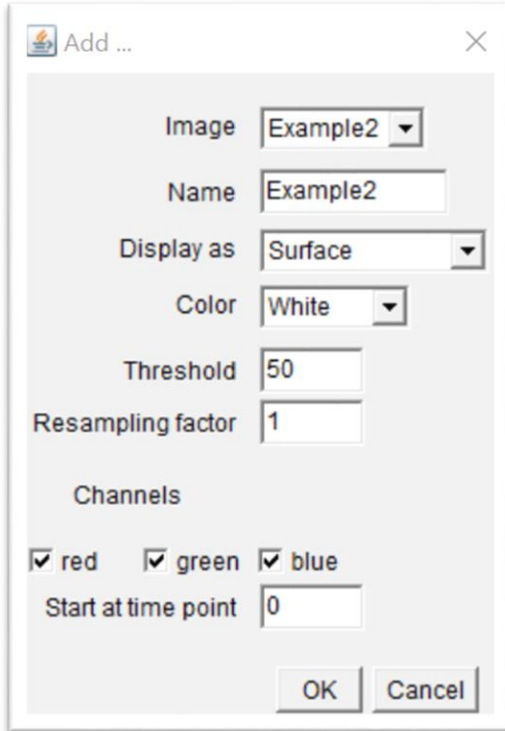
Isosurface: Ray-tracing and GANs



Isosurface: surface rendering

EXERCISE

Open Example 2 and try out the 3D viewer.




Select surface
Select a color

Set threshold

Do not downsample (value =1)

1. Plugins > 3D viewer
2. Select Display as surface, color (your choice) and resampling factor of 1)
3. Change the threshold (Edit > Adjust threshold). Set it to 50


FIJI



Avizo



Imaris




File > Export surfaces (Wavefront, STL, ...)

◆ □ 3D printer

Isosurfaces: example



Advantages:

- Computationally fast
- Good 3D interpretation



Disadvantages:

- Noise effects: only one signal (e.g. LSM channel, segmented/thresholded)
- Hence: not suitable for noisy data (e.g. electron tomography)
- Preferably: thresholded/segmented (binary) data



Main disadvantage: A decision for every voxel must be made.

This can produce:

- false positives (spurious surfaces)
- false negatives (erroneous holes in surfaces)

3D rendering

Never publish only renderings.
Always provide the raw data.

Note: renderings require **interpretation** by the user. Hence, they are the convolution of the raw scientific data and the feature the user would like to see.

1. Surface rendering
= binary threshold-based
2. **Volume rendering**
= percentage threshold-based

Direct volume rendering methods generate images of a 3D volumetric data set without explicitly extracting geometric surfaces from the data (Levoy 1988).

Volume rendering offers the possibility for displaying weak or fuzzy surfaces. This frees one from the requirement to make a decision whether a surface is present or not.

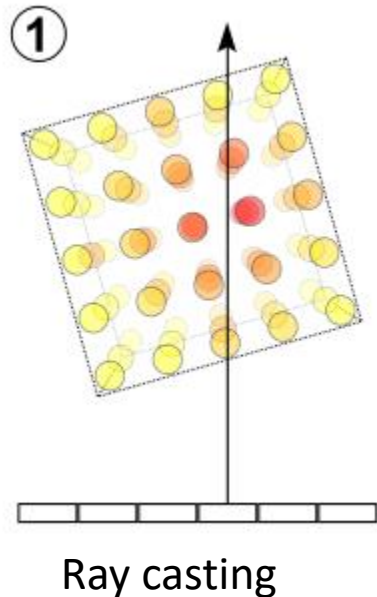
Every voxel should contribute to the image

How does it work?

1. **VOLUME RAY-CASTING (or ray marching)**: Cast imaginary rays through the entire 3D stack
2. DEFINE TRANSFER FUNCTION: setup rules for color and alpha (opacity)
3. DEFINE EDGES AND LIGHT SOURCE: shading
4. ACCUMULATE THE DATA

Volume rendering: 1. Ray casting & interpolation

For each pixel of the final image, a ray of sight is shot ("cast") through the volume.
At non-orthogonal angles, **interpolation** is needed



Volume rendering: Example - Maximum intensity projection

projects in the **visualization plane** the voxels with maximum intensity that fall in the way of **parallel rays** traced from the **viewpoint** to the plane of projection





110	48	38	37	38	37	39	37	38	38	38	40	43	49	75	87	92	97	92	107	110	←
124	45	38	37	36	37	37	38	37	39	39	40	45	54	66	74	92	99	99	120	124	←
123	35	35	35	34	36	37	38	35	38	41	43	47	63	74	72	97	108	109	123	116	←
134	37	34	36	35	35	39	38	36	38	43	48	57	67	77	80	86	118	113	134	118	←
135	37	35	33	34	38	37	38	38	39	39	50	59	71	89	102	105	105	112	127	135	←
144	38	34	35	35	37	40	37	40	40	43	54	74	77	89	102	109	110	105	115	144	←
127	36	34	36	37	35	39	40	41	42	49	57	72	82	87	109	124	123	114	114	127	←
137	34	35	35	35	37	40	44	43	47	57	64	75	82	88	102	114	128	124	129	137	←
138	36	36	38	37	38	43	42	46	47	60	77	88	96	95	108	124	128	115	138	135	←
140	35	34	36	38	40	43	42	47	52	75	90	92	102	105	113	123	118	114	127	140	←
132	37	37	39	42	44	45	50	61	77	88	91	97	104	112	112	122	123	120	122	132	←
135	37	37	39	40	48	56	58	68	72	89	102	107	103	121	113	133	128	130	135	117	←
151	37	38	41	45	53	68	78	74	84	97	104	118	114	125	120	124	133	138	151	137	←
138	40	41	46	53	57	72	84	80	83	120	113	117	125	122	129	134	127	137	131	138	←
143	44	47	52	61	63	89	103	92	100	110	120	119	125	128	132	128	127	143	136	134	←
142	51	54	64	68	87	98	102	109	111	105	121	122	131	132	127	124	129	127	141	142	←
158	58	56	70	75	89	97	104	114	111	112	129	129	131	119	126	158	149	131	148	128	←
155	64	86	97	90	101	109	110	117	109	109	127	128	129	127	122	153	155	132	139	129	←
154	67	85	108	100	108	115	112	135	127	124	129	131	140	141	134	141	154	144	126	120	←
150	96	89	108	113	111	112	116	126	147	143	147	144	140	147	147	141	150	140	137	136	←

For each sampling point: RGBA is computed (Red, Green, Blue and Alpha)

Volume rendering: step 2: Sampling and interpolation

For each pixel of the final image, a ray of sight is shot ("cast") through the volume.
At non-orthogonal angles, **interpolation** is needed




Volume rendering: step 2: Sampling and interpolation

Nearest Neighbour


= unweighted

- Take the value of the closest voxel

1D NN: closest of two points



2D NN: closest pixel of four corners of a square




Linear

= Center of mass


- Take the linear average of the two pixels the ray is intersecting

1D Linear: Center of mass of two points



Bilinear: Center of mass of square corners

Trilinear: Center of mass of cube lattice points




Cubic

→ Center of mass


= Lagrange polynomials, cubic splines or cubic convolution

1D Cubic: Center of mass of 3rd degree polynomial



Bicubic: Center of mass of 16 pixels

Tricubic: Center of mass of 64 pixels




Volume rendering: Example - Maximum intensity projection

projects in the **visualization plane** the voxels with maximum intensity that fall in the way of **parallel rays** traced from the **viewpoint** to the plane of projection

Image > Stack > 3D Project...

Original stack



Maximum intensity (brightest point)



Advantages

computationally fast

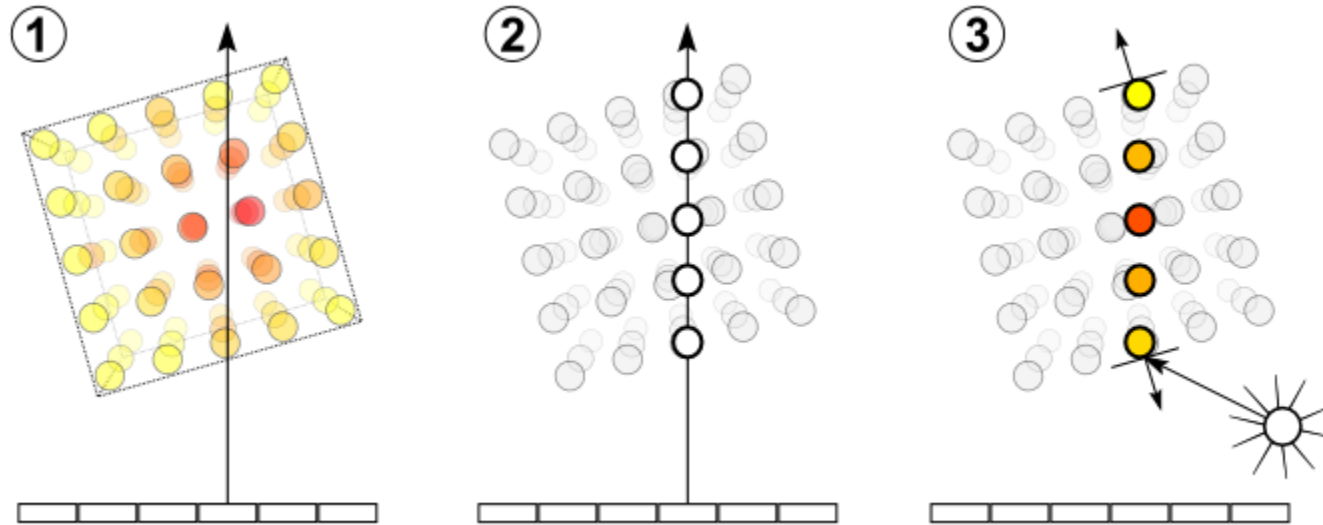
Disadvantages

May not provide a good sense of depth of the original data.

Two MIP renderings from opposite viewpoints are symmetrical images

No difference between left or right, front or back.

Volume rendering: step 3: shading




Shading

For each sampling point, a gradient of illumination values is computed. These represent the orientation of local surfaces within the volume. The samples are then *shaded* (i.e. coloured and lit) according to their surface orientation (normal) and the location of the light source in the scene.

Each sampling point is shaded according to its normal

Note: thresholding needed!


Imaris



Avizo



Volume rendering: step 4: compositing



Compositing

After all sampling points have been shaded, they are composited along the ray of sight, resulting in the final colour value for the pixel that is currently being processed.

$$L_o(\mathbf{x}, \omega_o, \lambda, t) = L_e(\mathbf{x}, \omega_o, \lambda, t) + \int_{\Omega} f_r(\mathbf{x}, \omega_i, \omega_o, \lambda, t) L_i(\mathbf{x}, \omega_i, \lambda, t) (\omega_i \cdot \mathbf{n}) d\omega_i$$

The total spectral radiance

\mathbf{x} = position

ω_o = direction (angle)

λ = wavelength

T = time point

The emitted spectral

radiance

The bidirectional

reflectance distribution
function

The spectral radiance

Volume rendering: Maximum intensity projection

EXERCISE

Open Example 2 and try out the Volume viewer (plugins > volume viewer)

Mode: **Projection (3)** Interpolation: **Trilinear (1)** z-Aspect: **1.0** Sampling: **1.0** **Background** **Snapshot** **Reset**


- Max Projection (2)
- Slice (0)
- Slice & Borders (1)
- Max Projection (2)**
- Projection (3)
- Volume (4)

- Projection (3)**
- Slice (0)
- Slice & Borders (1)
- Max Projection (2)
- Projection (3)**
- Volume (4)

Transfer Function (TF): Color & Alpha

Grayscale (1)


Draw LUT RGB R G B



- Volume (4)**
- Slice (0)
- Slice & Borders (1)
- Max Projection (2)
- Projection (3)
- Volume (4)**

1D 2D Grad 2D MD 3D Fill

Draw the alpha graph of the 1D-TF(lum)



global alpha offset

auto clear

Light

Drag sphere to change direction, double click to change color of light.






object color

ambient

diffuse

specular

shine



Volume rendering: Projection


EXERCISE

Open Example 2 and try out the Volume viewer (plugins > volume viewer)

Mode: Projection (3) Interpolation: Trilinear (1) z-Aspect: 1.0 Sampling: 1.0 Background Snapshot Reset


Projection

Alpha without transfer function adjustments



Projection

Alpha with 1D transfer function adjustments



Transfer Function (TF): Color & Alpha

Grayscale (1)

Draw LUT RGB R G B

1D 2D Grad 2D MD 3D Fill

Draw the alpha graph of the 1D-TF(lum)

global alpha offset a... cl...



Volume rendering: Projection


EXERCISE

Open Example 2 and try out the Volume viewer (plugins > volume viewer)

Mode: **Projection (3)** Interpolation: **Trilinear (1)** z-Aspect: **1.0** Sampling: **1.0** **Background** **Snapshot** **Reset**

Projection

Alpha with 2D transfer function adjustments



Transfer Function (TF): Color & Alpha

Original (0)

Draw LUT RGB R G B

1D 2D Grad 2D MD 3D Fill

Pick alpha (& color) and draw the 2D-TF(lum,grad). Draw with Alt key to erase.

pick color

Volume rendering: Volume

EXERCISE

Open Example 2 and try out the Volume viewer (plugins > volume viewer)

Mode: **Volume (4)** Interpolation: **Trilinear (1)** z-Aspect: **1.0** Sampling: **1.0** **Background** **Snapshot** **Reset**

- Slice (0)
- Slice & Borders (1)
- Max Projection (2)
- Projection (3)
- Volume (4)**

Projection

Threshold and set compositing effects

The image shows a 3D volume rendering of a bear. The bear is rendered in a semi-transparent, greenish-blue color, allowing internal structures to be visible. The bear is positioned within a 3D coordinate system with axes labeled x, y, and z. The bear's body is semi-transparent, revealing internal structures like the spine and ribs. The bear is rendered in a semi-transparent, greenish-blue color, allowing internal structures to be visible. The bear is positioned within a 3D coordinate system with axes labeled x, y, and z.

1D 2D Grad 2D MD 3D Fill

Draw the alpha graph of the 1D-TF(lum)

global alpha offset auto clear

Light

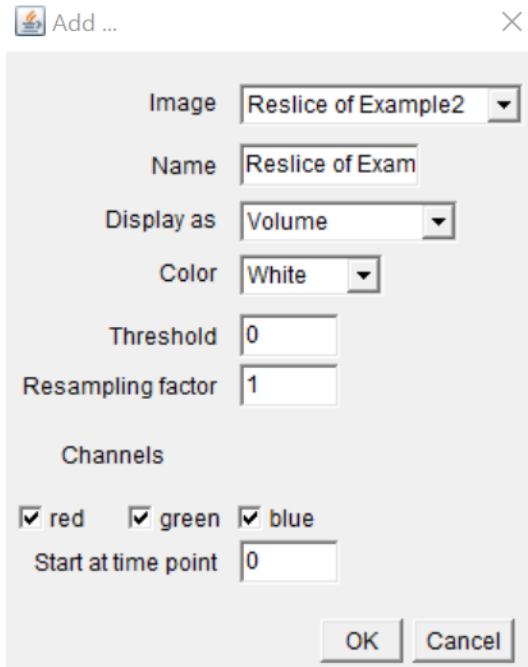
Drag sphere to change direction, double click to change color of light.

object color ambient diffuse specular shine

Volume rendering

EXERCISE

Open Example 2




Select Volume
Select a color


NO threshold
Do not downsample (value =1)

1. Plugins > 3D viewer
2. Select Display as volume, color (your choice) and resampling factor of 1)
3. No need to set a threshold


FIJI



Avizo



Imaris



Volume rendering: Imaris


BioNano has a workstation dedicated to Image rendering (amipc22.unifr.ch)

Soft Matter physics has also a workstation


More number cruncher available at Biology, Medicine, (physics?)

Imaris: dedicated to 3D LSM data


Surface rendering



Volume rendering





Object analysis



Volume rendering: Aviso

BioNano has a workstation dedicated to Image rendering (amipc22.unifr.ch)

Aviso: dedicated to 3D non-fluorescent 3D and 4D data



Z-Stacks

✓ Congratulations,
You finished Part IV, 3D

