Neuromantic User Guide V1.4

Darren Myatt*,
Cybernetics,
School of Systems Engineering
University of Reading

July 12, 2007

1 Introduction

This is the user guide for the Neuromantic application, a freeware program designed for reconstructing 3D models of neurons from serial image stacks. It is designed to streamline the process of semi-manual and semi-automatic reconstruction via an intuitive interface to navigate the stack and other useful features such as auto-focussing.

1.1 System Requirements

Neuromantic is designed to run well on most desktop Windows machines, but if not enough physical memory is available to hold the entire stack in bitmap format then performance will be significantly degraded. For example, a stack with image resolution 2731×3750 and 107 images will require just over 1 Gigabyte of RAM to be memory resident, which means that the machine should probably have at least 1.5 Gigabytes of RAM in order to run smoothly.

The application has been developed on a Windows XP platform, although it should work as well on most other versions of Windows. Additionally, the application is known to work in WINE for Linux. In order for the 3D graphics to function a recent version of OpenGL should be installed. If OpenGL is not present on a machine then the 3D Window will be unavailable, but this does not significantly reduce the functionality of the program.

Neuromantic was primarily designed to be used with a three button mouse/trackball. It can be used without one, such as on a laptop, but the interface will not feel as intuitive.

*d.r.myatt@reading.ac.uk

2 Working with Neuromantic

2.1 Loading an image stack

Neuromantic is designed to work with single images or stacks formed from sequentially numbered images of the type 'ImageNamexxx.yyy'. Select 'File→Load Stack' from the main menu or press F2 and select the first image in the sequence. Loading will then begin: depending on the stack size and format, this may take a while the first time.

The image formats currently supported by Neuromantic are:

- Windows Bitmap (.bmp)
- JPEG (.jpeg, .jpg)
- TIFF (.tif, .tiff)
- Multi-image Tiff files (.tif, .tiff)

Note that currently all images are converted to 8 bit grayscale upon loading in order to save memory. Also, 16 bit TIFF files are not currently supported, but this problem will be updated soon.

If you have an image that you need Neuromantic to load and it refuses to load or loads incorrectly, please email the image to d.r.myatt@reading.ac.uk stating the problem and I'll fix it asap.

2.2 Morphology Files

At the current time, Neuromantic only reads and writes SWC morphology files, although other file formats will be added in the near future, notably the MorphML XML format.

It has been confirmed that the neuronal simulation package NEURON[1] will import the SWC files produced by Neuromantic successfully.

3 Neuromantic Basics

Once a stack has successfully loaded it will appear in the main window. Use the mouse wheel to adjust the magnification level and hold down the right mouse button to drag the stack around. The magnification level can also be changed by using the drop down box on the tool bar.

The bar labelled "Scroll Through Image Stack" can be used, astoundingly enough, to scroll through the individual images in the stack. This can also be achieved (in all modes) by holding down the middle mouse button and moving the mouse vertically.

Neuromantic currently has two main modes (down from three in versions prior to V1.4.0) - semi-manual and semi-automatic. The current mode is indicated by which of the toolbar buttons at the top left is currently depressed, and

[1] http://www.neuron.yale.edu/neuron/

defaults to semi-manual. The mode may be changed via these toolbar buttons or by pressing SPACE to switch between them.

4 Image Processing

Neuromantic provides on-the-fly image processing that can be applied to the image stack to significantly improve visibility and can make visually resolving ambiguities much simpler. These options reside on the right hand panel.

Note that the image processing is only performed as an image is being drawn: no alteration is made to the underlying stack data, and therefore repeatedly adjusting the contrast will not lead to any degradation of the image data.

When dealing with Golgi stained (or similar) stacks from a TLB microscope that exhibit dark dendrites on a light background, it is generally advantageous to select the 'invert' option: this should expose a lot more detail.

The two sliding bars define the stretching of the luminosity histogram - by experimenting with these very low contrast features can be made much more visible. However, be careful of repeatedly altering the contrast throughout a reconstruction, as it may bias the estimation of dendritic radius when reconstructing by hand.

Figure 1 demonstrates the highly significant difference that inversion/contrast adjustment can make on a Golgi-stained stack.

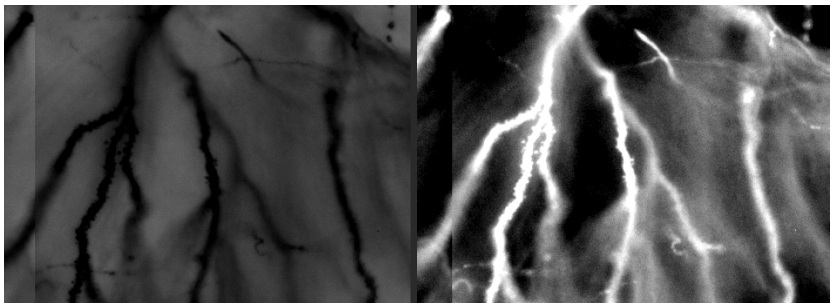

Figure 1: The left panel shows the actual luminosity of the original image stack: much of the detail is obscured. Through inversion, dendritic spines become significantly more visible, as in the right panel.

There are also multiple image interpolation options for the main window that can be switched between as follows:

- F9 - No interpolation
- F10 - Bilinear interpolation
- F11 - Bicubic interpolation

Bicubic interpolation is superior to bilinear interpolation, which is only really included for completeness. Using bicubic interpolation allows more effective and intuitive segmentation on thin dendrites when working with high levels of zoom.

5 Basic Stack Navigation

Navigation of the image stack including translation, zooming, changing the current image and auto-focussing may all be achieved through mouse control alone to maximise productivity. In both reconstruction modes, the basic controls for navigating around the stack are identical:

- Hold down the right mouse button to drag the stack around.
- Scrolling the mouse wheel (or pressing - and + on the keyboard) will zoom in and out of the stack. The scroll direction that causes a zoom in can be flipped by Options→Controls→Reverse Mouse Wheel Zoom.
- Press and hold the middle mouse wheel and then dragging the mouse vertically to change between different stack images (i.e., moving through the z axis) - this mirrors the functionality of the stack bar underneath the main image.
- Clicking the middle mouse wheel will perform an auto-focus at the given point over ± 5 stack images. Holding SHIFT while clicking will focus through the entire stack.

6 Basic Reconstruction Editing

In both reconstruction modes individual compartments or entire branches may be selected and edited as desired. Both the selection operations and the varying editing options will be discussed in detail in the next two subsections.

6.1 Selection operations

Left click somewhere along a compartment to select it: the mouse cursor changes to indicate when a compartment will be selected by a left click. The compartment node (white square) will be highlighted, along with the line linking the compartment node to its parent, and the current image will be changed to the z value associated with that compartment (if an image stack is loaded).

If a compartment node does not have a parent then it will just be displayed as a square alone.

Pressing SHIFT whilst left clicking will add compartments to the current selection, rather than replacing it. Performing SHIFT+Left Click on a compartment that is already selected will deselect it.

Pressing CTRL whilst left clicking will select the entire branch the given compartment lies upon: all descendant compartments will be selected, and ancestor compartments will be selected either to 1) the last branch point or 2) the

tree root (back until a parentless node is reached). The mode is defined by the current setting in Options→Branch Selection.

The branch selection performed when pressing CTRL+SHIFT will also stop if another selected compartment is reached is any direction. Therefore, to select a given subtree, select the beginning compartment of the subtree, then hold CTRL+SHIFT and click on one of its child compartments: the entire subtree will then be selected. Similarly, by selecting multiple compartments with SHIFT before performing a branch select it is simple to select any specific part of a tree.

If desired, you could then inverse the selection with CTRL+I and delete all other compartments, leaving you only with the given subtree.

Pressing ALT whilst left clicking will select all compartments of the same *type* as the clicked compartment, thus allowing easy selection of the entire basal or apical tree etc.

Pressing Ctrl+D deselects all nodes (and is also available from Edit→Deselect All).

There are several specific selection options in the Edit→Autoselect menu that allow the selection of trifurcations, nodes with no parent and lone nodes (those with neither parent nor child).

6.2 Editing compartments

The x, y position of an existing compartment may be modified by left dragging it. To modify the Z position of a compartment(s), select the compartment(s) and move to the desired new stack image (in either order), then select Edit→Set Z To Current Slice, or use the shortcut CTRL+C.

You may alter the radius of a compartment by selecting it (left click), then holding down CTRL and dragging with the middle mouse button pressed.

Altering connectivity is achieved in the following two ways:

Add Connection Select the desired parent compartment (i.e. the compartment node that will be further up the tree after connection), then select Edit→Connect Node. Now left click on the child compartment node that wish to connect the first node to, and you should see the appropriate connection appear. If the child node already had a parent, then it will be replaced by this operation.

Remove connection Select the two nodes which have an undesired connection between them, and then select Edit→Disconnect Selected. This function will also work with greater than two nodes by removing any connection shared between two selected nodes.

7 Manual Mode

Manual mode is accessed via the black cross toolbar button, and Neuromantic will automatically start up in this mode. In manual mode, compartments are

added to the reconstruction one at a time by the user, in a manner similar to the normal reconstruction methods in Neurolucida and similar applications.

Each compartment is defined by a 3D point, a radius value and a pointer to its parent compartment. Each compartment has only a single parent, and therefore a strict tree structure is enforced on the reconstruction.

Once you have located your desired dendrite via navigating through the stack, click the middle mouse button to auto-focus on the dendrite. By default, this auto-focus is constrained to five slices above and below the current slice to help stay focussed on the correct dendrite. To focus over the entire stack, hold down SHIFT while clicking the middle mouse button. If required, such as when the auto-focus focusses on the wrong dendrite, press S to return to the last focus depth.

To define a compartment point, click and hold the mouse button on one edge of the desired dendrite and then drag the mouse to the opposite edge, thereby providing an estimate for the diameter of the dendrite at this point. A red line will be shown as this is happening.

Once the left mouse button is released the red line will disappear and a compartment point (a white square) will appear showing the estimated midline of the dendrite at this point. The z (depth) position of a compartment point is set to the stack slice that was selected at the time that the compartment was defined.

If the process is then repeated for a section slightly further down the dendrite a new compartment point will appear and a line will join the two compartment points so far defined, showing a progression down the dendrite. Carrying on this procedure will result in the entire dendrite being segmented.

If you make a mistake at any point, pressing DELETE or A will delete the last defined compartment point.

The parent of any newly defined compartment is always set to the currently selected compartment, so in order to add a branching point simply click on the required compartment to select it, and then continue adding compartments from there.

The type of the compartment that is currently being added is defined by the radio buttons on the Edit Mode tab of the right hand panel. Also note that the compartment type of previously defined points may be edited using Tree Edit mode.

8 Semi-automatic Mode

To improve productivity, a semi-automatic tracing mode has now been added to Neuromantic. This mode allows the application to automatically trace lengths of dendrite indicated by the user and measure their radius.

Select semi-automatic tracing by going clicking on the AUTO toolbar button at the top left of the main window. As you move the mouse cursor over the main window you will now notice that there is a red dot following it - this dot demonstrates the size of the dendrite that you want to follow. To increase/decrease

this size, alter the top slider on the right hand panel.

On the right panel, there are two radio buttons that specify for either following dark or light dendrites - make sure the correct one of these is selected for your data or tracing will go very wrong. Also, remember that pressing "Invert" on the image processing panel *does not invert the underlying stack data*, so if you're looking at an inverted stack of a Golgi-stained neuron the "Dark dendrites" option should still be selected.

The basic process of semi-automatic tracing is as follows:

- Navigate to the correct slice by auto-focussing on the desired point (middle mouse button click achieves this, as in Tree mode).
- Click and hold the left mouse button: this will begin the tracing. It should remain held down until the tracing is complete.
- Move the mouse roughly down the dendrite - you will see blue boxes pop up (currently of a default size of 128x128 pixels). These patches are the area that is currently being image processed and routed. Each box represents (currently) a stack of 11 patches from the current slice - 5 to the current slice + 5. There must be an unbroken line of such patches down the dendrite, and the speed at which they pop up is dependent on the computer speed, so move more slowly if you have to.
- As you move the mouse down the dendrite auto-focus on the dendrite by clicking the middle mouse button (I admit this is a little fiddly at the moment, as you need to be pressing the left and middle mouse buttons at once). Also, you can still translate the stack by also holding the right mouse button.
- If a current route has currently been calculated to the position the mouse cursor is at, a line showing the basic trace will be shown. Once a tracing has begun, the application is constantly calculating routing information, so if no route currently exists then waiting for a few moments will help.
- Release the mouse button to complete the tracing.

If the algorithm won't correctly trace a neurite in one go, segment it in multiple traces - the next trace will automatically connect up to the last as long as the start point is nearby. Also, remember that if the trace ends up horribly wrong, you can simply press CTRL+Z to undo the change and start again.

Figure 2, which shows an image captured partway through making a trace should make the process clearer.

The slider underneath the dendrite radius slider is the Neurite Threshold, which determines which pixels will and will not be counted as neurites, and therefore considered when routing. On a single image this may be set to zero, but on large stacks tuning it appropriately allows much faster semi-automatic tracing. The higher (more to the right) the slider is, the less pixels will be considered, and the faster the routing will occur. Selecting Test→Routing Progress

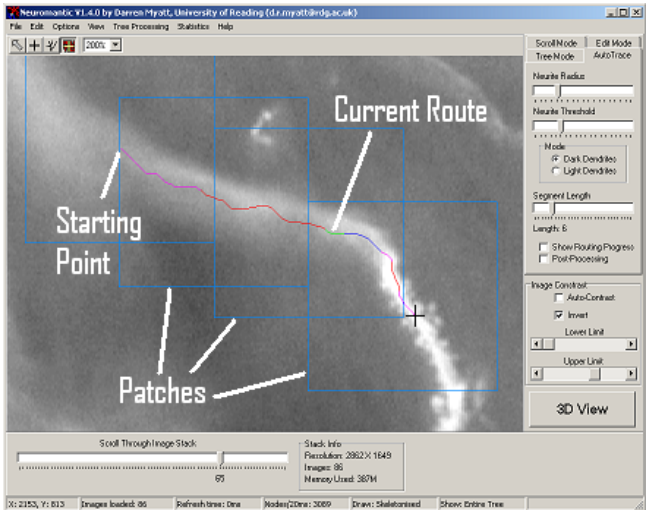

Figure 2: A labelled screenshot taken partway through a semi-automatic tracing.

Display→Routed Nodes will show you which pixels are actually currently being considered during tracing by colouring them in shades of blue (based on the cost function at that point), and thus let you tune this parameter correctly for your data. Hopefully, though, this will be done automatically in future versions. Figure 3 shows what the routing progress looks like when the parameter is set to about the right level.

If the threshold is too high, however, you might see something like Figure 4: it is apparent that the algorithm is not considering pixels that should be counted as neurites, as there is not a solid line between the start and end point.

Conversely, if the threshold is too low you may just end up with all the patches filled completely with blue, in which case the algorithm will work significantly more slowly than it should.

The Segment Length slider bar determines the subsampling rate of the final calculated path. This is by default set to 5 pixels long, but should in general be set high enough that you don't end up with a ridiculously large number of compartments but low enough so that you still approximate the shape of the dendrite well.

9 Bug reporting

Please report any bugs/errors to d.r.myatt@reading.ac.uk and, if possible, describe how they may be replicated to make my life a little easier.

Before you do so please check the README.txt associated with the latest version of Neuromantic to be sure that the error is not already known and/or fixed.

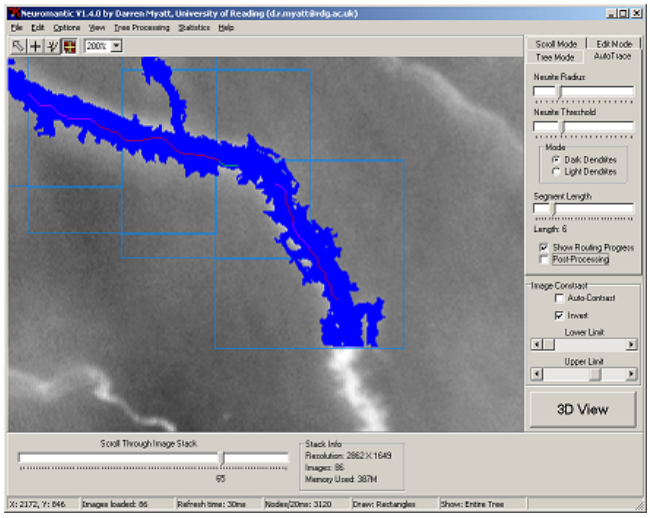

Figure 3: A figure showing the proportion of nodes routed when the neurite threshold is set to approximately the correct value.

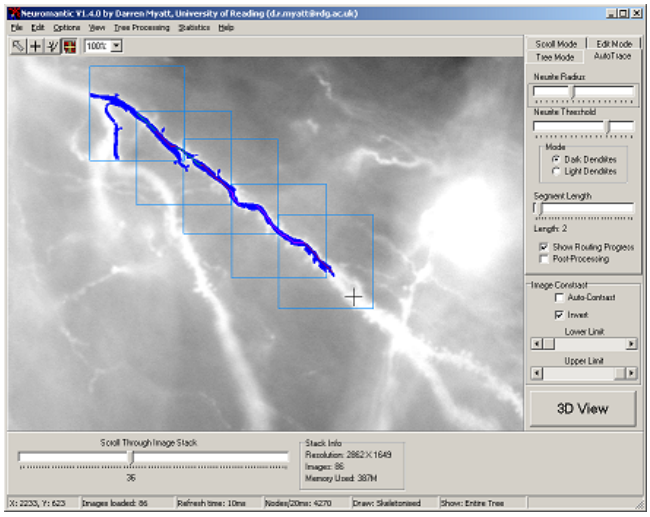

Figure 4: A figure showing the proportion of nodes routed when the neurite threshold is set too high.