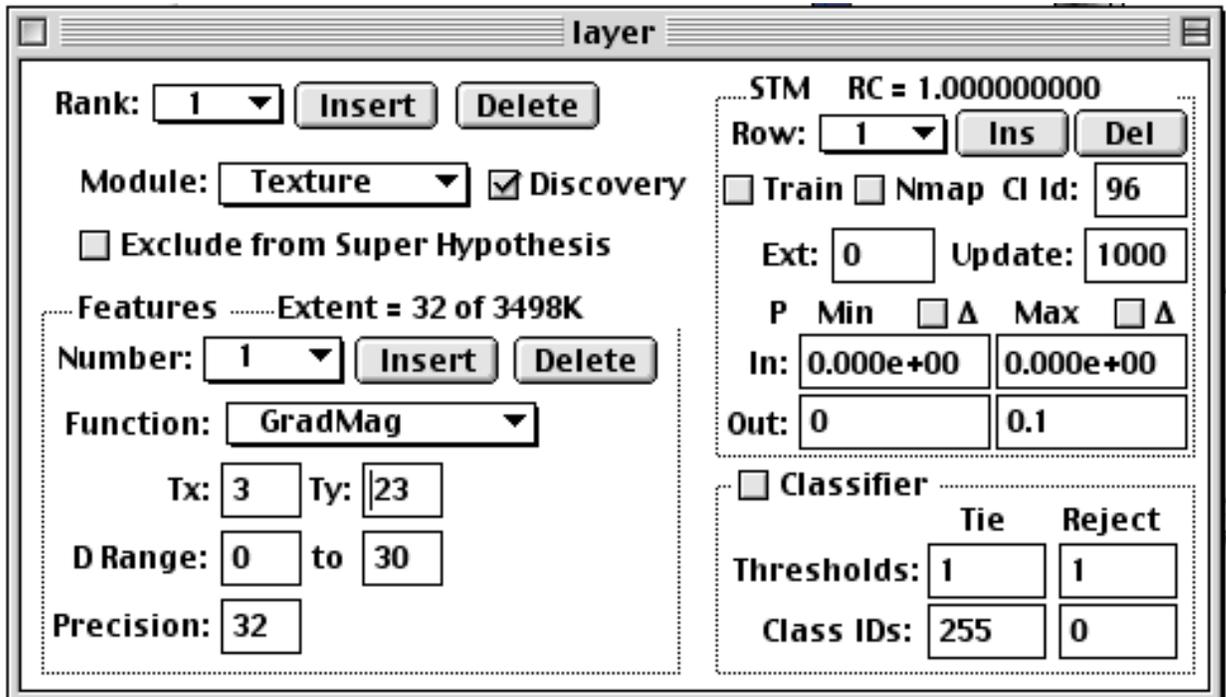


Apéndice A – Mac ALISA Manual Versión 12.2

Se incluye en este apéndice la versión del manual usada. Aunque el documento es limitado, resultó factible operar con los diversos parámetros para obtener los resultados presentados.



The screenshot shows a window titled "layer" with the following configuration options:

- Rank:** 1 (dropdown), **Insert** (button), **Delete** (button)
- Module:** Texture (dropdown), **Discovery**
- Exclude from Super Hypothesis**
- Features** Extent = 32 of 3498K
- Number:** 1 (dropdown), **Insert** (button), **Delete** (button)
- Function:** GradMag (dropdown)
- Tx:** 3, **Ty:** 23
- D Range:** 0 to 30
- Precision:** 32
- STM** RC = 1.000000000
- Row:** 1 (dropdown), **Ins** (button), **Del** (button)
- Train** **Nmap** **CI Id:** 96
- Ext:** 0, **Update:** 1000
- P** **Min** **Δ** **Max** **Δ**
- In:** 0.000e+00, 0.000e+00
- Out:** 0, 0.1
- Classifier**
- Thresholds:**

	Tie	Reject
	1	1
- Class IDs:**

255	0
-----	---

The **ALIAS** Corporation welcomes you to MacALISA!

1. Setup

Platform. MacALISA is available for any 68K Macintosh or PowerPC running under System 7.1 or higher. Please make sure you order the version you need. Although it comes from the factory with the preferred RAM allocation set to about 7800K, it runs quite satisfactorily with 3000K. Below that, you will be unable to keep more than a few images on your screens at a time. Use the Get Info box to allocate the appropriate amount of RAM for your application.

MacALISA now supports realtime color video input. MacALISA will recognize a camera or a VCR connected to a video acquisition board or to the video input ports on Macintoshes with built-in video-acquisition hardware, and even a Connectix QuickCam connected to a serial port. Note that QuickTime and QuickTime Power Plug extensions must be enabled on your system, even if you are not actually using the color features or color video input.

Monitors. You will need at least one color monitor supporting 256 colors. However, if you are going to process realtime input from a color video camera or VCR, you will need to set a monitor capable of displaying at least thousands of colors, and preferably millions of colors, to use the color features correctly.

The more screen area you have available, the more comfortable using MacALISA will be, because typically you will want to have A LOT of images on the screen at one time. If you have a small B&W or gray-scale monitor in addition to a color monitor, that is very useful. You can use the small monitor to display the Module Parameters window, and the color monitor to display images.

Getting More Information about MacALISA. For a description of the capabilities and applications of MacALISA, you are invited to browse through the home pages of Professor Peter Bock in the Department of Electrical Engineering and Computer Science at the George Washington University in Washington DC (URL = <http://www.seas.gwu.edu/faculty/pbock>) Inquires about MacALISA and other ALISA products may be directed to Mr. James Hubshman at The ALIAS Corporation (email = JRHals@mail.idt.net).

2. Conventions and Restrictions

Accessing Images. Currently MacALISA can read 8-bit images in four different file formats: Raw, Photoshop 2.0, BOB (a special image format used by DeBabelizer), and RGB. It can also display and process color images captured directly from a video camera or VCR. By default, MacALISA saves images in Photoshop format, but you can change the saved image format by simply selecting another format via the Options menu. Note that MacALISA does not process RGB image files directly; when you open an RGB image file, MacALISA displays it as a **luminance** image with 256 levels of gray. Although video images captured directly from a camera or VCR *are* displayed in color, they are processed just like any other RGB image. (See Direct Video Image Processing below.)

If you save images in either BOB or Photoshop format, then MacALISA will never have to ask you for the size of an image when trying to read it. Also, if the file is a (non-text) histogram or a SQUARE raw image, MacALISA will detect this and read the image in without asking for its dimensions.

WARNING: MacALISA assumes a raw image is square if its filesize is a perfect square, which is NOT a foolproof test. So, if possible, you should avoid using rectangular images whose areas are perfect squares. If this cannot be avoided, check "Always Ask Size" in the Options menu.

If you ask MacALISA to open a rectangular image whose area is not a perfect square, you will be prompted to specify its size in rows and columns. However, if you are opening a set of rectangular images automatically one after another (File menu item "Open > Images...") once you have specified the size of the first image, MacALISA will assume that all subsequent images of the same area have the same aspect ratio.

WARNING: If MacALISA is opening a set of non-square images and eventually encounters an image that has the same area as the first image (which determined the "default" area), but has a different aspect ratio, MacALISA will not recognize this and open the image incorrectly, switching rows and columns.

MacALISA can process a single image that is already open in the active window, or it can automatically process an entire set of images that are not yet open. The commands that process entire sets of images are those which end with the word "All...", such as "Transform All..." and "Train All..." under the Process menu. The entire set of images to be processed must be in a single folder, to which you will be asked to navigate before the process is begun. Once you have selected the correct folder, ALL the images in that folder will be processed. Thus, you may wish to create special folders for processes, containing only those images you wish to have processed during a particular step in your experiment.

Viewing and Saving. If you specify the "All..." option for a process, and you do not check "Pause After Each" in the Options menu, you will have only a moment to view the processed image before its window closes and the window for the next image opens. If you have elected to save the processed images, this may not matter, because you are probably only interested in generating the new set of processed images without viewing them during the generating process. On other other hand, if you have not elected to save the processed images, unless you are training you are not accomplishing very much, except to observe the sequence of windows as they flash very briefly on the screen one after the other. If you need more time to see each image, check the "Pause After Each" option, and each image will remain on the screen until you hit any key.

Images saved automatically as part of a <process> "All..." are placed in the same folder as the source images with an appropriate extension appended to the associated source image file name. After the process is completed, you can view the folder by Date, and select and move the new images to a convenient folder. Note that an RGB image is always saved as a luminance image, not the RGB image.

Storage Requirements. A word of caution about saving your files. When you're running a classifier experiment with many classes and trying lots of different configurations, it's very easy to use up a lot of disk storage. If you also want to save the class maps on which your experiment performance metrics will be based, you're likely to eat up a fair amount of disk space rather quickly.

However, we do not recommend that you solve this problem by simply throwing away your MacALISA files. Murphy knows, that strategy will just come right back and kick you in the face. What's our advice? Get a bigger hard disk, or load a large-capacity removable cartridge before you begin a big project. Also note that compression utilities like AutoDoubler can reduce STMs and class maps to tiny fractions of their expanded sizes.

Aborting a Process. Pressing Command-Period (or actually Command-<any key>) will abort an on-going process after the current image has finished processing.

Unfinished Business. MacALISA is an evolving product. There are many functions listed in the menus and galleries that have not been implemented yet. Those menu items that are dimmed (but should be available in the current context) simply have not yet been implemented, such as all the printing commands in the File menu. None of the items in the Calculate menu works yet (except "Custom" under the STM item, which is intended for power users only). Each version of MacALISA adds more capability, and The ALIAS Corporation will keep you informed about updates.

3. Doing Things Other Than Training or Testing

3.1 Viewing Images

Opening and Closing. Images (including saved histograms) may be opened and displayed one at a time or in sets using the "Open Image" or "Open Images..." items in the File menu, respectively. (See Conventions above.) Images may be closed one at a time using the close box on the image window or the "Close Image" item in the File menu. All open images may be closed at once using the "Close Images" item. An image may be saved with its current name or a new name with "Save Image" or "Save As Image..." in the File menu, respectively.

Video Input. If you have a video camera or VCR connected to your Mac, you may activate the capture of a sequence of images via the File Menu under New > Image > Video (Shortcut is Command D). Activating video capture will open a small window that displays the sequence of images coming from the camera or VCR as fast as your video-acquisition equipment can capture them. You can freeze a captured image by clicking on the window, and restart capture again by clicking on the window again. At the current time, to display *and* process color correctly, the video input window must be positioned somewhere on the monitor that displays the menubar. Moreover, to apply transforms to these images at the current time, you must configure the monitor to display thousands or millions of colors, because Apple uses indexed color with the system palette to display only 256 colors. (Also see Direct Video Image Processing at the end of this manual.)

Image Info Window. The Image Info window (which is opened by default when MacALISA is launched) displays an enlarged view of a 9x9 pixel field around the location of the cursor in the active window of any image other than a histogram. (See Making Histograms below.) Below the enlarged view is an information line that provides the following information: <pixel value>: [<row number> , <column number>] x <magnification factor>

Magnification. The image magnification factor of any active window can be increased or decreased using the "In", "Out", or "Normal" submenu items for the "Magnification" item in the Window menu. The "Set..." item has not been implemented yet. With direct video input, the magnification factor controls the size of the sampled image, from 120 by 160, to 640 by 480. The larger the image, the more time it will take to process a selected transformation.

Palettes. The display palette for the active image window may be set via the submenu items for the "Palette" item in the Window menu. The "Anomaly" palette is the standard white-blue-yellow-red-black ALISA anomaly detection palette. The "Custom" palette is also currently identical to the "Grayscale" palette, but will eventually offer the user the ability to customize the palette. The Color IDs (pixel values) of the "Classifier" palette include the standard set of anomaly detection colors from ID 000 to 031, a set of reds of varying intensity from ID 032 to 063, a set of grays of varying intensity from ID 224 to 255, and 16 bands of constant colors:

<u>Color IDs</u>	<u>Color</u>	<u>Color IDs</u>	<u>Color</u>
096-103	red	160-167	burgundy
104-111	dark blue	168-175	light gray
112-119	cyan	176-183	light green
120-127	dark green	184-191	light blue
128-135	orange	192-199	purple
136-143	magenta	200-207	teal
144-151	brown	208-215	khaki
152-159	yellow	216-223	dark gray

In future versions each band will include 8 decreasing saturations of the corresponding color.

3.2 Making Transformed Images

MacALISA currently offers a large set of transforms, including most of the "standard" ALISA transforms. New transforms are regularly added to the repertoire. To recommend a transform to be included in MacALISA, send an algorithm via email to "JRHals@mail.idt.net". Eventually, MacALISA will allow you to "plug in" your own custom transforms.

Source Images. The image in the active window (including a *frozen* direct-video-input image) may be transformed using the "Transform" item under the Process menu. After selecting a transform and specifying its parameters, the transform will be applied, and the transformed image will appear in a new window with an appropriate extension appended to the source-image file name. An RGB image *file* is displayed in grayscale as its luminance image, from which the Hue, Saturation, Lightness, Brightness, Red, Green, and Blue images may be extracted with the corresponding transforms. A *direct* video image is displayed in color in its own window (See Video Input in Section 3.1).

An entire folder of images (or a continuous sequence of direct video images) may be transformed using the "Transform All..." item under the Process menu in a similar fashion. However, before you specify the parameters of the transform, you will be asked to navigate to the folder containing the images to be transformed.

Parameters. Different transforms require different parameters. However, most transforms will ask for at least three parameters: token size, dynamic range, and precision. The token size specifies the size of the receptive field for the transform, *i.e.*, the size of the window over which the transform is applied. The dynamic range parameter defines the dynamic range for the transformed image as a subset of the dynamic range of the source image, and the precision parameter defines the number of equal-sized bins over the specified dynamic range into which the transformed image will be quantized. In addition, you may check the "Equalize" box if you want the feature image to be displayed with the specified precision and dynamic range expanded to the full range of 256 gray levels.

3.3 Making Histograms

In a fashion similar to making feature images or transformed images, MacALISA will generate a histogram for the image in the active window ("Histogram" command), the histograms for all the images in any folder ("Histogram All..." command), or a histogram of the sum of the counts and sum of the squares of the counts for all images in any folder ("Accumulate Histogram..." command). The histograms report the counts for each of the 256 levels in an image (gray scales, normalities, confidences, or class IDs). The vertical axis is automatically scaled to accommodate the maximum count. Often this corresponds to a level that, because of its relatively large count, makes the bars for the counts of other levels rather small and hard to see. Future versions will offer some scaling options; for now you can at least determine the actual count for any bin (and other information) by placing the cursor over the bin and reading the information line at the bottom of the Image Info window:

bin number> : <percentage of total counts between 0 and bin>% <count>/<total counts>

4. Training and Testing

The primary purpose of MacALISA is to allow you to train on sets of images that belong to one or more classes and then test other images to see which class they belong to. Training with direct video is similar to training with image files; the differences are explained in the section entitled Video Training and Testing at the end of this manual.

4.1 The Layer Configuration

Role and Definitions. The configuration of a single layer of a multi-rank ALISA system, including any STM weights that have been acquired as a result of training (or, for that matter, as a result of any method) is called a **layer configuration**. The ALISA Module Parameters window is the focal point for defining the **layer configuration parameters**. In this window you specify the number of ranks for the layer, select the type of module for each rank (texture or geometry), select frequency or discovery learning, specify the modules features and their parameters, define the operating and display parameters for each row of the STM and whether it is to be trained, and specify the classifier parameters. A layer configuration can be saved at any time (Save As > Layer...). Likewise, at any time you may open a previously saved layer configuration (Open > Layer...), but of course this will replace the layer configuration that is currently in memory. You will always be warned about actions that will overwrite the parameters and STMs and be given a chance to Cancel

t h e a c t i o n .

4.2 Defining the Layer Configuration Parameters

Defaults at Launch. When you first launch MacALISA, the ALISA Module Parameters window appears, presenting a default layer configuration with one feature (gradient magnitude) and one STM row for a Texture Module using frequency learning. All these parameter settings are arbitrary, and all the weights in the single row of the STM are zero (indicated by a Saturation of 0.0%). Chances are none of these parameters is what you want, so you will most probably want to specify new values for each of them.

Module Type. Probably the first decision you will want to make is the module type. In the Module box at the top of the Layer Parameters window, select Texture or Geometry. Specifying the parameters for a Geometry Module is explained separately in a subsequent section of this manual. If you select the Texture module type, it is now up to you to decide which to specify next: the Feature parameters in the Feature Frame (the left side of the Layer Parameters window) or the STM in the STM Frame (the upper right quadrant of the Layer Parameters window). The Classifier parameters in the Classifier Frame (the lower right quadrant of the Layer Parameters window) can be set and reset at any time, either during training to observe the progress of learning, or during testing to generate your class maps (*i.e.*, the results of classification). This is explained in more detail below.

Feature Specification. To specify the features, select the Function for the feature Number 1 and specify all its parameters. Note that the unit of precision is not bits, but states. In MacALISA the number of states for a feature may be any positive integer, limited only by the available RAM. The upper and lower bounds of the Dynamic Range are expressed as percentages. Any real number in the interval $\{0,100\}$ is okay, although the upper bound must be greater than the lower bound.

WARNING: If the lower bound is not zero, MacALISA will use one of the states for the feature to accumulate the counts for all feature values between zero and the lower bound. Likewise, if the upper bound is not 100, MacALISA will use one of the states for the feature to accumulate the counts for all feature values between the upper bound and 100. This is not a bug. It is a standard ALISA feature that allows some non-linearity in the feature quantization at either end or both ends of the dynamic range. However, note that if you specify 2 states for the precision (a bit unusual) and a dynamic range whose lower bound is not zero and whose upper bound is not 100, insufficient resources have been requested to provide all four states: an underflow state, 2 central states (what you asked for), and an overflow state.

Repeat this process for each feature you want in your layer configuration. If you decide you don't want one of the features after all, delete it with the Delete button. The extent of the STM for these features is reported in the number after the word "Extent =". Make sure you have enough RAM to accommodate this extent for EACH row in the STM you have specified (*i.e.*, four times the product of the extent and the number of rows).

Single-class Parameter Specification (anomaly detection). Now go to the STM Frame and define an STM for single-class anomaly detection. Only one row is needed, which should be already there by default. (You might want to check that there is only one row specified for the STM by clicking on the Row index box.). The Class ID is unimportant, because it is only used for defining the class color in class maps for multi-class classification, not anomaly detection. If you want MacALISA to do anomaly detection during training to watch the progress of learning, make sure the N Map (Normality Map) checkbox is checked. Remember, without checking the Pause After Each option, you'll have to be content with a brief glimpse of each normality map. To set Pmin and Pmax, see Section 4.4. Note that the values for Pmin and Pmax have no effect on training or testing, just on the dynamic range of the *display* of normality maps.

Because this example assumes you are doing anomaly detection, the Classifier will not be used, so make sure the checkboxes for Classifier is unchecked, which it should already be by default.

Multi-class Parameter Specification. For multi-class training, it should be pretty clear now what to do. When you are setting up the layer configuration parameters for the module, just keep inserting rows in the STM, specifying the parameters for each one, until you have included all the rows (classes) you want.

Learning Policy. You have one more decision to make before you train your module: do you want to use frequency or discovery learning? If you want to use discovery learning, make sure the Discovery checkbox is checked; if you want to use frequency learning, make sure the Discovery checkbox is NOT checked.

You can Save the new layer configuration using the Save As > Layer.... menu commands.

4.3 Training

The Train Checkbox. First make sure you click the Train checkbox for the STM row you want to train, so that MacALISA knows to update this row of the STM. You are allowed to check the Train checkboxes for several rows at the same time, but training multiple rows simultaneously with the same data is very unusual!

Train and Train All.... You can train a row (class) in the STM one image at a time, or you can train a row in the STM with an entire folder of images or a sequence of video input images automatically, one after another.

If you want to train one image, open this image using the "Open Image" command in the File Menu or activate the direct video input and select this window. Then choose the "Train" command from the Process menu. Depending on the speed of your Mac, the number of features you have specified, and the complexity of the feature algorithms, MacALISA will train this image in somewhere from 0.5 to 20 seconds.

If you want to train an entire set of images automatically one after the other, assemble these images into a folder somewhere on some volume. (Make sure there are no documents but the training images in this folder!) Then choose the "Train All..." command from the Process menu, navigate to the folder where the training images are, click Open, and specify the image size, if asked. After training starts, you can watch the images go by as they are learned. Also you can observe the hypotheses that MacALISA generates if you have the "N Map" or "Classifier" boxes checked. These hypotheses can also be automatically saved for you.

Untrain and Untrain All.... It is sometimes useful to **un**learn a lesson. The "Untrain" and "Untrain All..." commands in the Process menu let you do this. Simply follow the same steps as for training, and when the untraining process is finished, the feature vectors counts for the untraining set will have been subtracted from the STM. Untraining is often used to measure the effects of subsets of a large training class on detection or classification performance by eliminating them one at a time. Untraining can also be used to recover from training errors.

Saving. When training (or untraining) is complete, the STM will be saturated (to some extent) with useful statistics about whatever it is you trained with. You should save the layer configuration with the new STM data before you do anything else.

You can reload this (or any other) saved layer configuration at any time. Having done so, you can test with it, retrain any existing rows in it, train any existing rows with additional training images, untrain subsets of the trainings for one or more rows, and add and delete additional rows. You **cannot** change the Features or Feature parameters without zeroing all the STMs. Changing the parameters makes the allocation of the feature space invalid and is not allowed.

Zeroing the STM Counts. At any time you can **zero** the counts in one or more existing rows in an STM in order to retrain these rows with new or modified training images or simply remove these rows (classes) from consideration during classification. To do this, simply select the appropriate row in the STM, zero it using the Zero Counts command in the Edit menu (you will be warned about the consequences of this action), click the Train checkbox for this row, choose "Train" or "Train All..." in the Process menu, and then proceed as described above for training. To train any existing row with additional training images, simply check its Train checkbox and choose "Train" or "Train All..." from the Process menu.

Multi-class Training. A multi-class STM is trained one row (class) at a time, so make sure that you have checked the Train checkbox for the row you are about to train, and that the Train checkboxes for all the other rows are unchecked. Otherwise MacALISA will update more than one row during this training pass.

When the training for one row has finished, uncheck the Train checkbox for that row, cycle to another row, check the Train checkbox for that row, and initiate training again. Repeat this for every row in the STM that you want to train or retrain.

For multi-class training, you probably will want to UNcheck Anomaly Detect for all the rows during training, or you will get quite a stream of normality maps marching across your screen (one for each row), as well as chew up a lot of time with the classifier. You need not set the Classifier parameters until you are ready to test. On the other hand, occasionally you may, in fact, want to turn on the classifier during training that you can see the class maps that are generated to observe the progress of learning (use the Pause After Each option). If you want to observe the class maps during training, set the Class ID for each class (row) for the color you want, based on the color table listed at the end of Section 3.1. You can read about setting the Classifier parameters in the Section 4.4.

Specifying and Training a Geometry Module. Specifying the layer configuration parameters for a Geometry Module is almost the same as for the Texture Module, except, of course, the Geometry Module uses only the class-difference feature, which has only one parameter to specify: the undersampling interval. Specifying the STM parameters and Classifier parameters is the same as with the Texture Module.

The multi-class training procedure for a Geometry Module is the same as described above for the Texture Module. It is certainly also possible to do anomaly detection with a Geometry Module, although this is not often done. Usually the purpose of an experiment with a Geometry Module is to classify several different kinds of structures in a class map, either canonically or secularly.

You will note that the first time you select the Geometry Module type, MacALISA goes to sleep for a few seconds, and the watch cursor is active. During this hiatus MacALISA is using a storage mapping function to define the stimulus domain of the Geometry Module. This only happens when you select a Geometry Module for the first time after launching MacALISA.

Masked Training. As a manual selective attention mechanism, the "Masked Train" and "Masked Train All..." commands in the Process menu allow you to train only those pixels in the training image that correspond spatially to pixels on an associated mask that have a specified value. This value is defined as the Class ID for the reject class. All pixels in the training image that do not correspond to a mask pixel with this value will be ignored, *i.e.*, will not cause update of the STM.

A mask for each and every training image must be stored in a separate folder, must be the same size as the corresponding training image, and must have the same filename as the corresponding training image. There must be exactly one mask for each training image. Before training begins, you be be asked to navigate to the folder containing the mask(s).

4.4 Testing

Accessing a Layer Configuration. Testing images means either classifying and/or anomaly detecting. Before you can test any images, you must either load a layer configuration file, which contains all the layer parameters and the STM weights from a prior training session, or you must have just finished training an STM, so that the layer parameters and STM weights are already in memory. To load a previously saved layer configuration file, use the Open > Layer..." File menu command.

Before you test an image, you will want to set the anomaly detection and/or classification parameters in the Classifier Frame to appropriate values. Here is some guidance for this.

Pmin and Pmax. Pmin and Pmax are thresholds that define the range of normality values (expressed as probabilities) that will be displayed within the bounds of the color or grayscale range of the normality map. Any feature vector with a normality less than Pmin will be set to black in the normality map; any feature vector with a normality greater than Pmax will be set to white in the normality map. Values between Pmin and Pmax will be displayed as colors within the range of the palette selected under the Window menu. Pmin and Pmax affect only the display of the normality map, not the STM. They are used to visualize the range of the results of anomaly detection and are often useful to investigate possible decision criteria for specific applications.

Setting Pmin and Pmax is a bit of an art. If you don't know how, as a rule of thumb, start by setting Pmin equal to zero and Pmax equal to one-half of the reciprocal of the STM extent. Then experiment with the values until you get normality maps that represent what you want.

Classifier Parameters. In the Classifier frame, you can specify the thresholds for ties and rejects, as well as the Class IDs for the colors or gray levels that will be used to display ties and rejects. Both thresholds lie in the real confidence interval $\{0,1\}$. If you set both thresholds to unity, no feature vectors will ever be classified as rejected, and only those feature vectors whose class conditional probabilities are numerically identical will be classified as tied. As you lower the thresholds, more and more feature vectors will be classified as rejected or tied. If you set both thresholds to zero, all feature vectors will be tied (not rejected, generating an all-white class map).

Testing one or many images. To test a single image, open it and/or make sure it is the active window by clicking on it, and then choose "Test" from the Process menu (or press Command-T).

To test many images, choose "Test All..." from the process menu and navigate to the folder where you have placed your test images. MacALISA will then proceed to test all these images, one after another.

Cross Validation. (currently disabled) The "Cross Validate" and "Cross Validate All..." commands facilitate "leave-one-out" classification, commonly called cross-validation. During a Cross Validate process, each image is first untrained, then tested, and finally retrained. This allows a set of training images to be treated as "unique" test images, one at a time, and eliminates the need to have a separate set of unique images for control purposes.

Before initiating cross-validation, make sure you check ONLY the Train box(es) of the row(s) for which the images were trained.

Direct Video Image Processing. You can transform, train, or test continuous sequences of images from a video camera simply by starting video capture and using the Train All..., Transform All..., or Test All... menu items. This process is stopped by clicking on the video window (or actually anywhere), but you *MUST HOLD THE MOUSE DOWN* until the current video image has been completely processed. Just quickly clicking on the window will not work, because the Apple Event Queue is not checked until the processing of the current image is finished. (It would be too expensive to check the queue more often.) You will know when to release the mouse, because the training or testing process will simply halt; then you can let go.

Good Luck! And thanks for Asking ALISA.