ProRaster User Guide

Created: June 2022
Last updated: 12 March 2024
ProRaster Version: 2.2.06
Author: Sam Roberts

ProRaster Essential

ProRaster Premium

ProRaster Scientific

ProRaster Essential, ProRaster Premium, and ProRaster Scientific are developed by Roberts Geospatial Engineering - Software for GIS and Big Data, Delivering Innovation, Performance, and Correctness.

Roberts Geospatial Engineering develops innovative geospatial software for big data processing, visualization, and analysis. We strive for performance, accuracy, and correctness. Be confident of the right results every time.

https://robertsgeospatial.com.au/

About Sam Roberts, Developer and Founder, Roberts Geospatial Engineering.

- Geospatial Engineer with 30 years' experience in C++ engineering for GIS, Geophysical and Scientific applications.
- Well-versed in modern C++ programming techniques and have a proven track record of delivering complicated projects on time and on budget, whilst balancing the competing demands of design excellence, performance, efficiency, and future extensibility.
- Extensive data processing experience and focus on raster technology, satellite multispectral data processing and visualization, LiDAR processing, and potential field modelling and processing applications.
- Lead developer on the MapInfo Pro Advanced product, which brought a new generation of raster processing, visualization, and analysis to the desktop. MapInfo Pro Advanced won the Australian Business Award in 2017 for Software Innovation.
- Designed and engineered the MRR raster format, which is a key enabling technology in the product, which was awarded a United States Patent in 2018 (Number US 9877031 B2).

Contents

Welcome to ProRaster

Thank you for purchasing, acquiring, or researching the ProRaster family of products!

The ProRaster product family delivers three core outcomes:

1. ProRaster Essential and ProRaster Premium are designed to render raster data. The primary goal of the rendering engine is to achieve the highest level of rendering quality and accuracy at any scale. The secondary goal is to provide consistently high rendering performance at any scale that provides users with an immersive experience. The hope is that this will lead users to an improved understanding of their data and an enhanced capacity to make discoveries in their raster datasets.

2. ProRaster Scientific ships with multispectral satellite imagery processing and analysis, powered by virtual raster technology.

   ProRaster Scientific is a powerful and efficient platform that enables experts and amateurs to maximise their productivity when processing, rendering, and analysing multispectral satellite imagery. It uses virtual raster technology to chain together processing operations that shepherd your imagery from the original rasters through to advanced products. Do less and achieve more – faster and more efficiently!

   The easy import process applies the initial corrections to your Landsat or Sentinel2 imagery. Create new products by combining imagery spatially (mosaics), temporally (sequences), and by pixel compositing (masking out bad quality pixels and filling the gaps with high quality pixels from temporally adjacent imagery). Then apply processing operations including spectral index computations, polygon clipping and pixel masking, calculator operations and data transformations, reprojection and more. Complete your analysis by computing multidimensional statistics, rendering imagery and spectral indices, and batch outputting imagery frames for video production.

3. ProRaster Scientific ships with a suite of Raster processing operations powered by virtual raster technology.

   Traditional raster processing imports a source raster, applies a processing operation to it, and exports an output raster. This process consumes time and resources, and you can't check the result of your processing operation until the output has been generated. ProRaster Scientific uses virtual raster technology to apply processing operations in real-time, at any scale, to rasters of any size. It takes no time, uses no resources, and you can immediately see if your processing operation has been appropriate and successful.

   Your virtual raster output can be used as the input to another operation, leading to chains of processing operations. Of course, you can "crystallise" your virtual raster at any time by exporting it to a standard raster format like MRR or GeoTIFF.

ProRaster Online Resources

Please make use of the product website to further your understanding of the product. Visit the ProRaster product page at:
https://robertsgeospatial.com.au/products/proraster/

YouTube You will find a rich variety of videos on YouTube that will support your learning. The Roberts Geospatial Engineering YouTube channel has playlists focussed on ProRaster as well as exploring raster data and other related topics. Explore the channel at:

https://www.youtube.com/channel/UCZOyA3Mbodd-x0-PVXW89WA

Get an introduction to ProRaster at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnZxXuFicgsEH4tHPui_yokK

Discover ProRaster basic training at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnapplar54wccXL_fVR1sP8w

Take a deeper dive into ProRaster at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnahZUwTU7zCmyGJn3vBL1W2

See the latest ProRaster features at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnbnBgfFdA2LCqHCTASH506y

ProRaster Product Family

There are three tiers in the ProRaster product family – ProRaster Essential, ProRaster Premium, and ProRaster Scientific.

ProRaster Essential is designed exclusively for business users of the latest version of MapInfo Pro. It has a limited feature set and only supports publishing of algorithms to MapInfo Pro. It is not compatible with other GIS applications or mapping applications in general. Learn more at:
https://robertsgeospatial.com.au/products/proraster/essential/

ProRaster Premium is designed for general users of any GIS or mapping application. It provides an unrestricted set of functionalities and can harness the full power of the rendering algorithm concept. Learn more at:
https://robertsgeospatial.com.au/products/proraster/premium/

ProRaster Scientific is the flagship product that builds on Premium by providing advanced productivity tools for processing, rendering, and analysing satellite multi-spectral data. It also provides a suite of raster processing operations. Learn more at:
https://robertsgeospatial.com.au/products/proraster/scientific/

Get an introduction to ProRaster Scientific at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnYrwpKCdvt3G7zdqIk6Nm4V

See worked examples of ProRaster Scientific at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnb9Rac32a2NwBJPwFQc0Sk1

To see the features in the latest release and the history of ProRaster releases, please visit the product updates page at:
https://robertsgeospatial.com.au/products/proraster/updates/

To see what features might arrive in the future, visit the ProRaster Roadmap page at:
https://robertsgeospatial.com.au/products/proraster/proraster-roadmap/

If you are particularly interested in a new feature, or would like to suggest a new feature, please provide feedback to Roberts Geospatial Engineering.

Hardware Recommendations

ProRaster will run on modest hardware that can run Windows 10 (version 1809 or later) or Windows 11. The primary hardware requirement is that you have a display resolution of at least 1440x900 pixels. If your display resolution does not meet this minimum requirement, then the ProRaster dialogs may exceed the size of the screen.

ProRaster Premium and Scientific allow you to open external windows. Large display surfaces provide more immersive opportunities for data discovery. I recommend 27-inch monitors at 2560x1440 resolution. I am wary of using very high-resolution monitors as they increase the amount of rendering work ProRaster must do without necessarily improving the visualisation of the data (as pixel-scale details are physically very small on-screen).

ProRaster uses the GPU to render imagery, but you do not need a discrete graphics card. A CPU with an integrated GPU is perfectly adequate.

ProRaster uses multithreading to improve performance when rendering and when processing. My recommendation is that you have a quad-core CPU or better. ProRaster Scientific will benefit from more CPU cores and higher performance CPU’s.

You ought to have at least 4GB of RAM. My recommendation is that you have 8GB or more to run ProRaster Essential and Premium, and 16GB or more to run ProRaster Scientific.

If your primary “C” drive is a hard disk, then it is time to upgrade! ProRaster works better if your data is stored on a high-performance SSD, and you ought to buy the largest one you can afford.

Getting Started

Let’s get started straight away! If you have a raster you want to display, open ProRaster and then open Windows File Explorer. Locate the raster you wish to render and drag-and-drop the file onto ProRaster. ProRaster will create a default algorithm for the raster and the algorithm will be rendered in the preview map window. Now save your algorithm by hitting the “Save As…” button and specifying the file name and location. Your algorithm is saved to a small XML file with a MRD extension. See a YouTube demo [here](#).

Raster Formats

Some rasters are stored in a single file, others in two or more files, and some use an entire directory on disk. If you are unsure which file you should drag-and-drop onto ProRaster, then you can drop down the “Open…” menu button and select the “New algorithm for raster…” option. This opens a file browse dialog and helps you choose the right file by suggesting appropriate file extensions for different raster formats. In some cases, you can select any of multiple files associated with a raster and ProRaster will open the raster correctly.

You can find a list of the supported raster formats at:
https://robertsgeospatial.com.au/products/proraster/help/supported-raster-formats/

Overview Cache

When you open a raster in ProRaster for the first time, you might experience a delay before the raster is displayed in the map. This is because ProRaster is preparing an overview pyramid cache for the raster. These cache files are stored alongside the original raster and may have a PPRC, PERC or OVR extension. If you have a very large raster, preparing this cache may take some time. For more information, refer to the chapters “Understanding raster technology” and “Understanding the raster engine”. To exercise control over the building of the overview cache, refer to the chapter “Using the Raster Source Editor dialog”.

You can find a video playlist explaining raster technology at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnaTNLDt-Kq4TkukcQfKT78J

Opening Algorithms

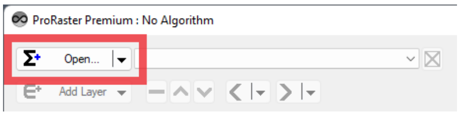

Open Menu Button

If you have created (or been supplied with) an MRD algorithm file, you can open it by hitting the “Open…” menu button, selecting the “Open Algorithm…” menu option, or you can simply drag-and-drop the MRD file onto ProRaster from File Explorer. To easily load algorithms that you have accessed in the past, look for and select an algorithm from the recently accessed list on the Open menu button. Algorithms that have been opened directly from the multispectral product editor appear in a sub-menu.

If you are a MapInfo Pro user and you have displayed a raster in Pro, you may have a GHX file stored alongside that raster. Inside this GHX file is a description of how you rendered the raster in Pro. You can match that rendering style in ProRaster by choosing the “New Algorithm for Raster GHX…” option on the Open menu button. Browse to the raster and ProRaster will build an algorithm matching the rendering style you used in Pro.

If you have already defined a raster source, you can create a default algorithm for it from the Open menu button by selecting the “New Algorithm for Raster Source…” option. Select the raster source you want to display and ProRaster will create a default algorithm for the raster source and the algorithm will be rendered in the preview map window.

You can create a new algorithm for a raster via the “New Algorithm for Raster…” menu option. Browse to the raster file you want to display, and the algorithm will be constructed to display the raster appropriately.

You can create a new algorithm for a WebMap file (WMS or similar) via the “New Algorithm for WebMap” menu option. This will display a sub-menu that mirrors the content of the C:\ProgramData\RGE\WebMaps folder. Select any one of the WebMap files, and the algorithm will be constructed to display the raster appropriately. A selection of WebMap files ship with ProRaster, and you can add any files you develop yourself into the C:\ProgramData\RGE\WebMaps directory in any folder structure that you please to make them visible for selection.

Creating Algorithms

To start building a rendering algorithm manually, drop down the Open menu button and select one of the “New LUT/RGB/Image algorithm” options. The algorithm, which will have a single layer of the type you requested, will be displayed in the tree control. Select any item in the tree to display the property pages associated with that item. You can enable or disable components and layers by toggling the checkbox button in the tree, or by toggling the “Enable” checkbox in the primary property page for each item.

For each of the layer components that you have enabled, you will need to specify a raster source and select a field and band. On the “Component” property page, drop down the Raster Source list box and select a raster source. If you do not have a raster source defined, you can hit the browse button to the right of the list box and browse to a raster. Select an appropriate field and band.

For “LUT Color” and “RGB Color“ layers you also need to specify a data transform that converts data values taken from the raster into color values. For these components, select the “Color” property page and drop down the “Data Transform Type” list box to choose a transform category and then, for that category, drop down the “Data Transform” list box to select a specific transform. For “LUT Color” layers you also need to select a color table. On the “Color” property page, click on the “Color Table” preview graphic to open an in-place selection drop-list.

Previewing the Algorithm

If your algorithm is valid, then it will be rendered in the preview map window. The rendering engine is a multi-threaded tile-based renderer that uses hardware GPU acceleration. It is designed to provide a high quality, high performance, immersive rendering experience. See a YouTube demo here.

In the map window, use the mouse to zoom and pan. Roll the mouse wheel to zoom in and out about the location of the cursor. There is a natural limit to how far you can zoom out and once you reach this limit the map will fit to the algorithm extents. You can force the map to fit to the algorithm extents and redraw by hitting the “Zoom to algorithm extents” button. Pan the map by holding down the left mouse button and dragging.

Zoom Button

Drop down the “Zoom to selected extent” menu button to acquire more zoom options. You can zoom the algorithm extents, to the bounds of any coordinate system used in the algorithm, or to the extent of any raster source used in the algorithm. You can also select a manual zoom option which allows you to define a zoom rectangle in WGS84 geodetic longitude & latitude coordinates.

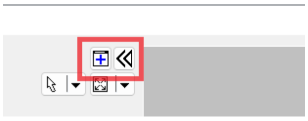

Open Map and Preview Slider

You can collapse the preview map window by clicking the “Close preview window” button. Once it is collapsed (and depending on your product tier) you can hit the “Open preview window” button to restore it, or you can hit the “Open a floating map” button to open an undocked and resizable map window. ProRaster Scientific and ProRaster Premium users can open multiple undocked map windows for each loaded algorithm. To close all undocked map windows for the currently selected algorithm, simply hit the “Open preview window” button. To control whether the ProRaster dialog stays on top of the map windows, see the Options dialog.

When you have multiple algorithms loaded, you can opt to keep all algorithms and their map windows active or only keep the current algorithm active. Only keeping the currently selected algorithm active can reduce memory consumption. To control this, see the Options dialog.

Tooltip Menu Button

As you move the mouse cursor about in a map window, ProRaster can report the raster value at that location as a tooltip. Drop down the tooltip menu button to select the reporting mode – No tooltip, Key Data Value, Key Data Information, or All Data Information. You can toggle the “Show geodetic coordinates” mode to report coordinates in WGS84 longitude & latitude rather than the coordinate system of the algorithm. To acquire a complete report on all the raster data at the cursor location, double click. A report will be displayed in the property sheet.

Editing the Algorithm

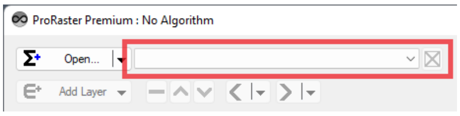

Algorithm List & Close Button

You can have multiple algorithms open in ProRaster simultaneously. To edit an algorithm, firstly make sure it is selected by choosing it in the algorithm drop-list at the top of the main dialog. When you select an algorithm, all the map windows associated with that algorithm are brought to the front of the window order. The algorithm tree is updated, and the appropriate property pages will be displayed. If you want to close an algorithm, hit the “Close the current algorithm” button.

Whenever you make a change to any property, the changes to the algorithm are immediately reflected in the preview map. Some property changes are not reflected immediately – for example when you type text or numbers into any text box. To apply changes that you have made to a number or a text string, hit the return key. In many dialog-based applications, hitting the return key causes the dialog to close. In ProRaster this does not happen, and it is safe to hit the return key to trigger a property update.

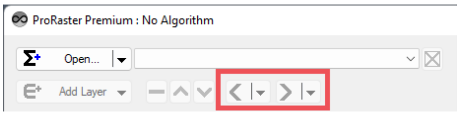

Undo and Redo Buttons

Whenever you make an edit, the change is recorded in the undo-redo list. You can drop down the undo and redo menu buttons to see what changes have been made or hit the button to undo or redo a change.

In the main dialog title bar, the status of the algorithm will be reflected. It will tell you whether the algorithm has been named and whether it has been modified and needs to be saved. To save an algorithm, hit the “Save” or “Save As…” buttons.

Publishing Algorithms

ProRaster Essential and ProRaster Premium users can publish algorithm to the latest version of MapInfo Pro. It is important that you are using the latest release of MapInfo Pro and that you have applied the latest patches and upgrades. See a YouTube demo [here].

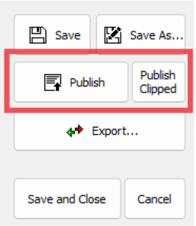

Publish Buttons

When you publish your algorithm, MapInfo Pro is directed by ProRaster to open your MRD algorithm file in a new map. Consequently, it is essential that you have saved your algorithm prior to publishing. If you have not done so, ProRaster may save it to a default location.

Make sure that MapInfo Pro is running. To publish to MapInfo Pro, hit the Publish button. If required, your algorithm will be saved and then you ought to see a new map opened in MapInfo Pro. Your algorithm will be rendered in this new map.

You can do this manually by simply opening a new map in Pro and then adding a new raster table and browsing to your MRD algorithm file. You cannot publish or otherwise consume your algorithms in any software apart from ProRaster and MapInfo Pro. This is because MapInfo Pro has a driver that will read and render MRD algorithms. It is exceedingly unlikely that such a driver will ever be made available for any other software. When you publish, no copy of the algorithm data is made. MapInfo Pro renders the rasters referred to in the algorithm. Consequently, you will be able to zoom in to any level of detail.

The "Publish Clipped" button publishes your MRD algorithm to MapInfo Pro, clipped to a polygon. MapInfo Pro will not render any pixels outside of the clipping polygon. When you hit the "Publish Clipped" button, a file browse dialog opens, and you must select the polygon file. This must be in MapInfo Pro table format. It can contain a complex polygon set but should not contain any other data besides the clipping polygons.

ProRaster Scientific supports publishing algorithms to MapInfo Pro from the Operations menu button. However, MapInfo Pro does not support the more advanced features of algorithms in ProRaster Scientific and it does not support the virtual rasters that are likely to be used in those algorithms.

Exporting Algorithms

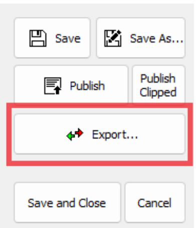

Export Button

You can export your algorithm to a located image raster. You can take this located image raster and display it in any GIS or mapping application. It is essential that you have saved your algorithm prior to exporting. If you have not done so, ProRaster may save it to a default location. See a YouTube demo here.

The Export operation can also generate multiple output images in batch export mode. In this case, an image is created for every unique time point in the algorithm. For more detail, see the comments later in this section.

Choose an appropriate format for the located image raster. MRR is the most efficient format, but to load an MRR into software other than MapInfo requires an MRR GDAL driver. GeoTIFF is a widely supported format. Band Interleaved (BIL) is also offered as a widely supported transportable format.

Depending on the output raster format, you may have storage compression options. I recommend using a lossless compression codec. When outputting to GeoTIFF you can use LZW or DEFLATE lossless compression codecs. LZW is more likely to be compatible with older software. DEFLATE requires you to set a compression level from 1 (least compressed) to 9 (most compressed).

Export Image Dialog

You can output 24-bit (RGB) or 32-bit (RGBA) color format. I recommend using 32-bit RGBA as the alpha band can be used as a transparency mask in applications like QGIS and ArcGIS.

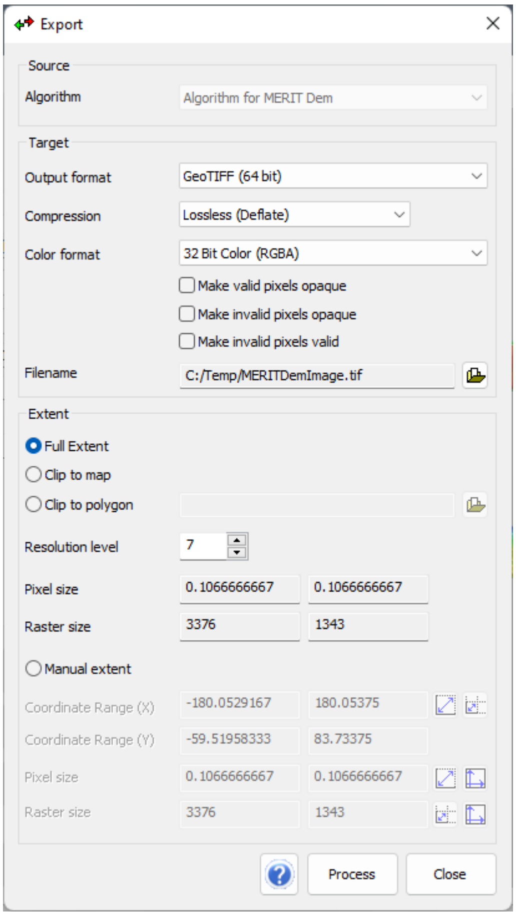

You can output the full extent of the algorithm, or the clipped extents reflected in the preview map window, or you can clip the algorithm to a polygon file. The polygon vectors must be in a MapInfo Pro table format. Complex polygons sets are supported. You can specify a rectangular extent manually. In this case you must enter the coordinate range, in the coordinate system and horizontal units of the algorithm.

Having established the extents of the output raster, you must then establish the cell size. This will determine the raster size (in columns and rows of pixels). This, in turn, will determine how large the output raster file is and how long the export processing operation will take to execute.

When using automated extents, you adjust the cell size by varying the target resolution level. This integer level varies the cell size by a factor of 2 for every integer adjustment. These cells are aligned with the raster cells (if possible) and by maintaining cell alignment you minimise the required interpolation and produce the highest quality output. Dial the resolution level up and down and you will see the pixel size and raster size adjusted accordingly. Once you are happy with your choice, you can proceed to execute the operation.

When using manual extents, you can edit the pixel size or the raster size independently (choose one!). After you edit one of these parameters you should adjust the other parameters to suit. Buttons are provided to –

- Fit the raster size to the range
- Fit the pixel size to the range
- Fit the raster size to the pixel size
- Fit the range to the pixel size
- Fit the pixel size to the raster size
- Fit the range to the raster size

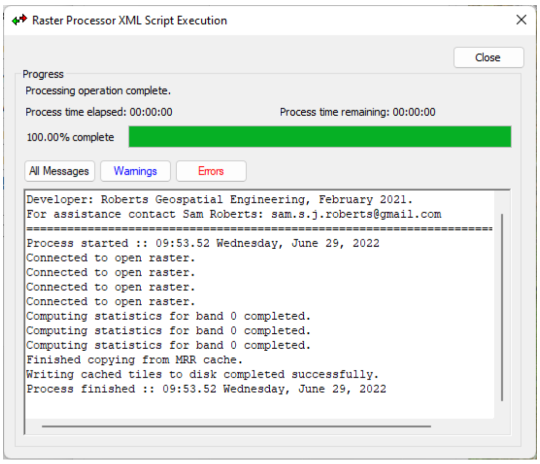

Processing Dialog

When you hit the Process button to execute the export operation, a raster processing dialog will open to report the progress of the operation. You can cancel the operation in this dialog if you wish. If the output raster is large, this processing operation may consume a significant amount of memory and CPU resources. Much like Eddie, the shipboard computer of the *Heart of Gold* thinking about how to make tea, you may find your computer sullen and unresponsive whilst the process is executing.

ProRaster Scientific users, from version 2.2.01 onwards, are able to access a temporal batch export option. An image is output for every unique time point in the algorithm. To activate this option, check the "Batch export by temporal event” switch which is directly under the output filename input controls. The option may be greyed out and unselectable. This occurs when all the rasters in the algorithm have only a single event. At least one raster in the algorithm must have multiple temporal events to trigger this option.

From the rasters that have more than one temporal event, ProRaster will identify all the unique time points in the algorithm. It will then step through those time points in sequence and set each component to a time range that is equal to this time point, or the closest event with a time before it. It operates just like the Time Control Panel in ProRaster Scientific.

Each image will be written using the filename you designated with a numeric suffix. This suffix is designed to be easily identified by video editing software which can automatically import the images in sequence.

A small text file with a “_Time.txt” suffix and extension is also written containing the time of each image in the sequence.

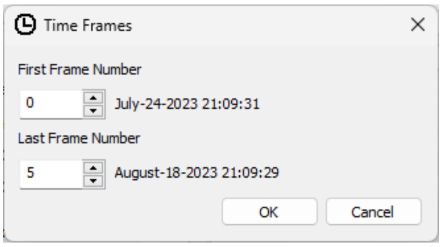

To exercise control over which frames in the temporal sequence will be output, hit the “Batch Export Time Options” button to display the “Time Frames” dialog. You can select the first- and last-time frame to export.

Understanding Raster Technology

A raster, in essence, is a database of values arranged in a regular geometric pattern. Rather than looking up data values by some key or index, we look up data values spatially by computing the coordinate of the cell which we want to interrogate.

A modern raster is likely to be a multi-resolution pyramid of sparsely arranged tiles of equisized cells. To find a cell value, we compute the tile coordinate and acquire the tile, then compute the cell index within the tile and acquire the cell. This can be done for data at any level of resolution.

You can find a series of videos explaining raster technology on YouTube here.

Resolution Levels

The base resolution level (0) is the "primary" level. It is the level from which all other levels are constructed. When we populate a raster, we firstly populate the base level, and then derive the other resolution levels from the base resolution data.

The overview resolution levels (1 upwards) contain data at increasingly lower resolution. In ProRaster, the size of the cells doubles for each resolution level. It follows that the number of cells in each resolution level is 4 times less than the resolution below it. In general, it is not possible to generate high quality overview level data on-the-fly. Consequently, overviews are generated for rasters and cached. They may be stored in the raster file (examples include MRR and COG) or they may be stored in a separate cache alongside the raster (examples include PPRC, PERC, and OVR files).

Overview levels are used to provide scale appropriate data as you zoom out to higher scales. When you zoom in, you will eventually require resolution levels that are higher than the base level. These are called underview resolution levels (-1 downwards). In ProRaster, underview resolution levels are generated on demand from the base level by interpolation. Interpolation methods such as nearest neighbour, bi-linear, and bi-cubic are employed. Underview levels are not permanently cached but may be stored in the raster in-memory tile cache temporarily.

Structure

To structure the data in a raster, we use the concept of fields, bands, components, and events. At the top level is the field concept, which is used to differentiate between the major classes of raster data – Continuous (gridded), Classified (thematic), Imagery, and Imagery using a color palette. The field type restricts the type of bands that may be present, simplifying the treatment of the data for users and software.

The bands of a raster are like the columns in a database table. Each field may contain multiple bands with unique data types. It follows that the bands in a field are likely to be intimately related to each other.

Components are acquired from band data on-the-fly. Some data types contain multiple numeric components that can be separated. For example, a RGB color data type can be separated into red, green, and blue color components.

Events can provide a temporal dimension in a raster. At each event, which occurs at a declared time, data can be stored in the raster. This may replace or augment data from other events.

Invalid Cell Values

In most rasters, some cells will have no data value. In the past these cells were referred to as NULL cells. Often, they contained a special NODATA value. In ProRaster these cells are referred to as invalid, and then we classify their status as either Empty or Null. An empty cell has no data but could have if it was defined. A Null cell has no data and is not allowed to contain data. Cells may contain valid data in some bands, but no data in other bands and cell validity is recorded per-cell and per-band.

Raster data must be able to be retrieved efficiently from storage. Some legacy raster formats are not able to be read efficiently. In these cases, ProRaster may create a PERC cache which duplicates the entire raster, as well as the overview levels. Compression is also a necessary feature of modern rasters. With a tile-based storage structure, rasters can combine efficient data access with data compression.

You can find a playlist of videos discussing modern raster technology at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnaTNLDt-Kq4TkukcQfKT78J

This article discusses modern raster structure and concepts of size:
https://robertsgeospatial.com.au/raster-size-concepts-explained-how-big-is-my-raster-2/

Understanding the Raster Engine

At the heart of ProRaster is a raster engine. It is primarily responsible for –

- mounting rasters in a wide variety of raster formats
- efficiently acquiring raster data
- creating or accessing overview pyramids to ensure that high quality data can be served at larger scales
- generating underview tiles by interpolation to ensure that high quality data can be served at smaller scales
- serving high quality, scale appropriate data to the rendering engine and virtual raster engine in any coordinate system
- caching raster data in memory for efficiency and performance
- acquiring or computing raster statistics

Drivers

To achieve this, the raster engine has a collection of drivers, each of which can load one or more different raster formats. The "native" drivers do most of the work and load most common raster formats. In addition to these drivers, a driver based on the open-source GDAL raster IO library is provided. This driver provides access to a wide variety of formats – both common and uncommon. Two special drivers are provided for virtual rasters (MVR) and rendering algorithms (MRD). The MRD driver allows the raster engine to mount an MRD file as a raster and invoke the rendering engine to populate it on-the-fly.

You can find a list of the formats supported by ProRaster at:
https://robertsgeospatial.com.au/products/proraster/help/supported-raster-formats/

The goal of the engine was to provide modern raster quality, flexibility, and performance, regardless of the format of the original raster. This is always achieved when native drivers are used, and generally achieved when the GDAL driver is used.

You may find that a driver mounts a raster incorrectly, or it interprets the data inappropriately. When you mount a raster, you can provide driver preferences to attempt to correct these issues. These preferences can be declared for a raster source resource and are applied to all the rasters in the raster source. For more information, look at the documentation for the Raster Source Editor dialog.

Resolution Levels

Overview pyramids are created by down-sampling. In ProRaster overview levels always vary by a factor of 2 and always cell-align with lower levels. Consequently, we always down-sample by averaging, if that is appropriate for the field and band data type.

Underview pyramids are created by up-sampling. In ProRaster you can choose nearest neighbour, bi-linear, local bi-cubic, or global bi-cubic up-sampling methods, depending on the field and data type.

When the raster engine provides data on a specific grid to the rendering engine or the virtual raster engine, it will choose an appropriate resolution level, and this may require using up-sampling interpolation. To then shift the data from the cell centres to the requested grid cell centres may require a second interpolation phase to move the data up to half a grid cell. This second interpolation phase always uses nearest neighbour or bi-linear.

Statistics

Rendering rasters often requires statistics. Statistics are used to define data – color transforms and are used to scale hill-shading effects. The raster engine can efficiently compute statistics for rasters for the base level resolution, or any overview resolution level. Best practice is to pre-compute statistics for the base level and to cache them. In ProRaster you may find this data stored in a GHX file alongside the raster. If statistics are not available, they will be computed on the fly by targeting an overview resolution level.

As of version 1.0.0.5, statistics are now stored alongside the raster is a binary STTX file, in addition to the GHX file. However, if both the GHX and STTX files are available when you open a raster, only the STTX file will be used. STTX files provide substantial performance benefits over GHX files.

Caching

The raster tile cache stores raster data in memory after it has been loaded from disk. When the engine requires that data again, it can acquire it directly from the cache rather than loading it from external storage. This improves performance. However, when you are rendering large rasters, it is easy to accumulate a very large amount of cached data very quickly. ProRaster limits the size of the cache, and you can control this from the Options dialog. The cache is also used when executing processing operations – like computing statistics, generating overview levels, or exporting to a located image. For these operations, a large cache may be beneficial.

Whilst a large tile cache generally leads to highest performance, it is important to try not to exceed the physical memory limit of your computer. As soon as the system starts allocating virtual memory, performance will degrade substantially. You may find it is better to use a smaller tile cache to guarantee the physical memory limit is never exceeded rather than use a large tile cache.

Events and the Temporal Dimension

The raster engine supports the temporal dimension via a concept called "Events". Every field in a raster will contain at least one event, and it is possible for a field to contain multiple events.

Each event occurs at a specified time. It is a requirement that the time of each event is later than the time of the previous event. At each event, the field will contain raster data for every band in the field. This data will either be declared as "totally replacing" data from previous events, or it may "partially replace" data from previous events. This is an important distinction that controls how data is accumulated over multiple events.

When you request data from the raster engine (either because you want to render or query it, or you want to use it for some processing operation) you can specify that you want data over a range of events - corresponding to a from-to time range. In this case, the engine will gather data by creating what are called "roll-up tiles". It uses a "painter's" algorithm to accumulate data into the tile, starting at the first event and proceeding forward in time to the last requested event. Data from later events will overprint data from earlier events. There are two important behaviour modifiers associated with this process.

Firstly, the first event used in a roll-up tile may be later than the event the user requested. This depends on whether events are "total replacement" or "partial replacement". The algorithm will not look back in time past a total replacement event. It considers such an event as an impenetrable barrier in time. By definition, there is no relationship between data before a total replacement event and the event data.

Secondly, invalid cells can have different classifications, and this controls how they are used in the painter's algorithm. If a cell is invalid and "empty" then it will not overprint a valid cell (from a previous event). If a cell is invalid and "null" then it will overprint a valid cell and invalidate it. In other words, an empty cell is transparent, and a null cell is opaque. Null cells can be used like a rubber to erase data from previous events. This cell classification is specified by setting the cell value to either zero (empty) or one (null) and setting the cell validity to false.

Consider the following examples:

Weather radar data: Weather radar data could be stored in a raster as a series of total replacement events. The radar information for each event is unique and unrelated to other events. You can step through the events and display the weather radar data like a movie.
Satellite Imagery: You can have a multi-spectral satellite scene that has been collected at multiple time points. Each scene is masked to remove cloud and other compromised pixels. These partial replacement scenes can be rolled up over time so that invalid pixels in earlier scenes are overprinted by valid pixels in later scenes, and invalid (empty) pixels in later scenes do not overprint valid pixels in earlier scenes.

Lidar Elevation: You might regularly conduct LiDAR surveys of an iron ore mine. Once a year you conduct a full site survey which you record as a total replacement event. Every day you conduct surveys of the open pit operation areas and ore stockpile, and these are recorded as partial replacement events. The elevation data for the mine can be recovered for any time by generating a roll-up tile from the previous total replacement survey to the requested event time.

The raster engine can store statistics (for all bands) for each event. These statistics are always computed as a roll-up of all events to the target event. If the statistics are computed at the base resolution level, then they will be stored (either in the raster file, in a GHX file, and/or in a STTX file).

You can find a playlist of videos discussing modern raster technology at:
https://www.youtube.com/playlist?list=PLE8RCT90BXnaTNLDt-Kq4TkukcQfKT78J

You can find a variety of articles on raster technology on the RGE website:
https://robertsgeospatial.com.au/the-structure-of-a-modern-raster/
https://robertsgeospatial.com.au/659-2/
https://robertsgeospatial.com.au/raster-statistics/
https://robertsgeospatial.com.au/rasters-in-mapinfo-pro/
https://robertsgeospatial.com.au/large-raster-rendering-performance-in-mapinfo-pro/

Understanding Algorithms

An algorithm contains global resources and settings and then one or more layers of a particular layer-type. Each layer has several components, each of a particular component-type, and the number and type of the components depends on the layer type. In the user interface the algorithm, layers, and the components of the layers, are all represented by different items in the algorithm tree control. When you select an item in the algorithm tree, the property sheet changes to present the property pages for that item and shows all the properties of that item. Get a background explainer on YouTube here.

This information is written into the MRD algorithm file. This simple text file contains XML script describing the algorithm. You can view this file in a text editor like Notepad++. As you become familiar with the XML schema you may even choose to edit the file manually on occasion.

An MRD algorithm file is largely self-contained. It will not, of course, contain any raster data. The raster files the algorithm uses are referred to by their file path name. It may also refer to color tables by their file path name, but in some cases the color table is written into the MRD file itself. If you use color table maps or color legends in data transforms, these will also be referred to by their file path name.

Although algorithms, layers, and components clearly have many properties, you may notice that not all these properties are written into the MRD file. Where a property is set to the default value, it will not be recorded in the file. This keeps MRD files short and succinct.

Render Order

Algorithms are rendered, layer by layer, using a “painters” algorithm. The layers in the algorithm are rendered in order from first to last. Pixels that have been written by a layer may be overwritten or modified by a subsequent layer. The layers are listed in the MRD file in the order in which they will be rendered. In the algorithm tree the layers are presented so that rendering proceeds from the bottom of the tree to the top. In other words, if you move a layer up the tree, it will overprint layers below it.

Components in a layer cannot be reordered, although individual components can be turned on and off. Most combinations of enabled and disabled components are possible, except that you cannot enable an Opacity component with no other components enabled.

Components

There are three types of layers, and the components found in each of these layers is shown below.

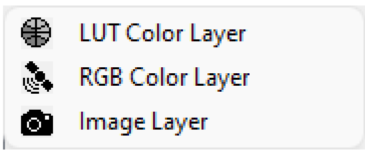

- LUT Color Layer
  - Color
  - Intensity
  - Opacity
  - Mask
- RGB Color Layer
  - Red
  - Green

- Blue
- Pan
- Intensity
- Opacity
- Mask

- Image Layer
  - Image
  - Intensity
  - Opacity
  - Mask

The Color / Red/ Green / Blue/ Image components provide the color (RGB components) of the pixel. The Intensity component applies hill-shading which modifies the intensity of the pixel color. The Opacity component computes the opacity (Alpha component) of the pixel color. Note that opacity is the reverse of transparency. In an 8-bit per component color system, an opacity value of 0 is completely transparent and a value of 255 is completely opaque.

The Pan component provides a mechanism to apply panchromatic sharpening at render time to satellite multispectral data.

The Mask component provides a mechanism to mask a layer using a raster. Pixels are only rendered where the mask raster source is valid.

Note that the Opacity, Pan, and Mask components are not available in ProRaster Essential. The Pan and Mask components are not available in ProRaster Premium.

Layers

A LUT Color layer computes pixel colors by transforming raster data values into color via a data transform and a color look-up table (LUT). This is a three-stage transformation pipeline. Firstly, the data value is transformed to a scaled index, a value between 0.0 and 1.0. Secondly this scaled index is transformed into an integer index matching the number of colors in the color table (from 0 to N-1). Finally, the required RGB pixel color is obtained from the color table.

A RGB Color layer operates like a LUT Color layer, except that the red, green, and blue color components for each pixel are all acquired individually. In this case, the color table for each component is defined by the system and cannot be modified. These color tables contain 256 varying shades of red, green, and blue from no saturation (0) to full saturation (255).

An Image Layer acquires pixel colors directly from a raster without modification. In other words, the raster must have a data band that has a color data type such as RGB or RGBA, and this color data is acquired from the raster directly and without modification.

Raster Sources

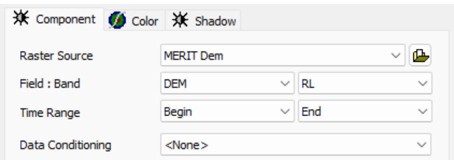

Component Raster Source Selection

All enabled components of a layer must be connected to a raster source. The field, band, and time range for this raster source must be specified. There is no restriction placed on what raster sources can be used in a layer or an algorithm. You can use rasters with different coordinate systems, different cell sizes, and different field and band structures. When a layer is rendered and component raster data is acquired, any coordinate system transformations or raster cell interpolations will be performed on-the-fly to ensure that the rendering engine is provided with the highest quality data that is most appropriate for the scale of the map.

A raster source may contain a single raster, or it may contain multiple rasters. In this case, the rasters must all share the same coordinate system and raster structure. It is likely, but not a requirement, that they will all have the same cell size too. This creates a possibility that there may be a different number of rasters in each of the components in a layer. This is only supported in three different scenarios.

1. Many : One
   Many rasters in the primary component and one raster in each of the secondary components. This is a common scenario – for example you may have multiple rasters in the color component and a single large-scale raster in the intensity component.
2. One : Many
   One raster in the primary component and many rasters in the secondary components. This is an uncommon scenario, and we expect that the components that contain multiple rasters have the same number of rasters. If not, only the minimum number are processed.
3. Many : Many
   Multiple rasters in the primary and secondary components. In this scenario we expect that you will supply the same number of rasters in each of the components. If not, only the minimum number of rasters are processed.

Working with the ProRaster Main Dialog

The ProRaster main dialog allows you to open, create, edit, and view multiple MRD raster rendering algorithm files simultaneously. From the system menu you can access the About dialog which contains links to online resources, and the Options dialog which enables you to customise some of the behaviour of ProRaster.

The left half of the main dialog is devoted to algorithm editing and the right half contains the preview map window. Hit the “Close preview window” button to collapse this half of the dialog and stop rendering. Hit the same button again, now labelled “Open preview window” to reopen the preview map window. If you have external map views open for the algorithm, hitting this button will close them all.

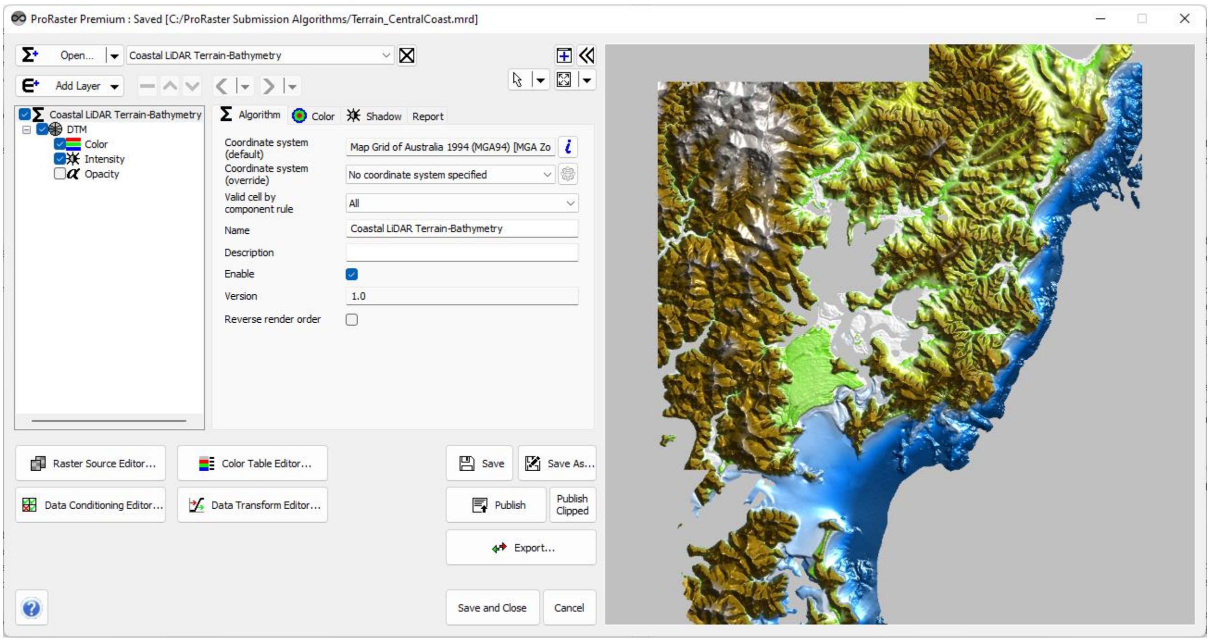

ProRaster Premium Main Dialog with Preview Map

Opening Algorithms

The “Open…” menu button allows you to open and create algorithms. Load an MRD by:

- Hitting the “Open…” menu button and browsing for an MRD file.
- Opening the “Open…” menu button, selecting “Open Algorithm…” and browsing for an MRD file.
- Drag and drop an MRD file from File Explorer onto the ProRaster main dialog.

Creating Algorithms

Create an MRD by:

- Opening the “Open…” menu button, selecting “New Algorithm for Raster…” and browsing for a raster file. This will create a default algorithm for the first field in that raster.
- Opening the “Open…” menu button, selecting “New Algorithm for Raster Source…” and selecting an existing raster source. This will create a default algorithm for the first raster in the raster source.
- Opening the “Open…” menu button, selecting “New Algorithm for Raster GHX…” and browsing for a raster file. The file ought to have GHX file alongside it that has been populated with rendering information by MapInfo Pro Advanced. This will create a matching algorithm for the raster.

- Opening the “Open…” menu button, selecting “New Algorithm for WebMap” and selecting one of the WebMap XML files listed in the sub-menus. This will create a matching algorithm for the raster.
- Opening the “Open…” menu button, selecting “New LUT Algorithm”. This will create an algorithm with a LUT Color layer. Two other items create algorithms with an RGB Color layer or and Image layer.
- Opening the “Open…” menu button and selecting a previously accessed MRD file from the recents list.
- Drag and drop a raster file onto the main dialog. This will create a default algorithm for the first field in that raster.

At the centre top of the main dialog is a drop-list showing the algorithms loaded for editing. The name of each algorithm is reflected in this list (so it is a good idea to uniquely name your algorithms). Change this list selection to edit different algorithms. Close the selected algorithm by hitting the “Close the current algorithm” button. If necessary, you will be prompted to save the changes.

Saving Algorithms

The title bar shows the status of the current algorithm. It will reflect whether the algorithm is unsaved, saved, or modified and will show the file path and name of the MRD algorithm file.

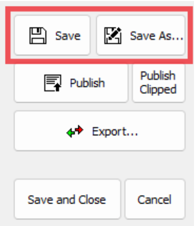

Hit the “Save As..” button to save a new algorithm or to save an existing algorithm in a different location. Once an algorithm has been saved, hit the “Save” button to save it to file. If it has not been saved previously, you will be prompted to browse to a file path and name.

Hit the “Save and Close” button to close the main dialog and exit ProRaster. By default, all modified algorithms will be saved silently before the application quits. If you have created new algorithms, you will be prompted to save them. You can change this behaviour in the Options dialog. If you disable the “Save modified algorithms on OK” option, then you will be prompted to save modified and new algorithms when you hit “Save and Close”, but no algorithms will be silently saved.

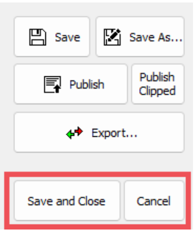

Hit the “Cancel” button (or hit the window close button) to close the main dialog and exit ProRaster. You will be prompted to save modified algorithms. By default, new algorithms that have not been saved will be discarded. You can change this behaviour in the Options dialog. If you enable the “Save new algorithms on Cancel” option, then you will be prompted to save new algorithms when you hit “Cancel”.

Editing Algorithms

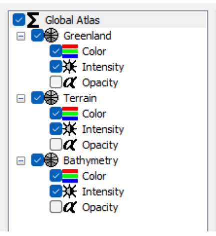

Algorithm Tree

The algorithm tree control reflects the Algorithm – Layer – Component structure of the algorithm. At the top of the tree is the algorithm item. The layers are subordinate to this item and listed from last to first. This reflects the rendering order where layers that are last render over the top of layers that are first. Each layer has two or more components that are subordinate to the layer item. The number and order of the components cannot be changed.

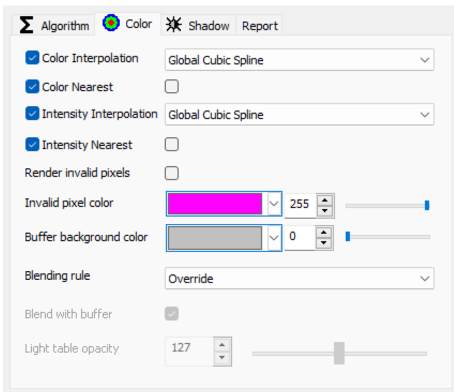

Property sheet

Select any item in the algorithm tree to display its properties. These are displayed in a property sheet next to the algorithm tree. There may be multiple property pages, depending on what item in the tree is selected.

You can enable or disable any layer or component by either hitting the checkbox in the algorithm tree alongside the item or toggling the “Enable” checkbox in the primary property page for each item. If you disable a component, then that component will not contribute when rendering the layer. You can turn components on and off in any combination. If you disable a layer, then that layer will not be rendered. Turning components and layers on and off may also change inherited and propagated properties like the coordinate system.

Add Layer Menu Button

Add a new layer to the top of the algorithm by dropping down the “Add Layer” menu button and selecting the type of layer that you want to add. From this menu button you can also add a new layer for a raster, a raster source, a WebMap, or for a WebMap Overlay. In each case, the layer that is added is most suitable for displaying the selected raster.

You can then layers by selecting the layer you wish to move and hitting either the “Move the layer up” or down buttons. Hit the “Delete the selected layer” button to delete a layer.

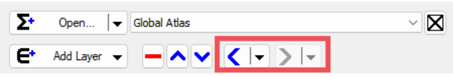

Undo Redo Buttons

All edits you make to an algorithm are recorded in an undo-redo list. Hit the “Undo edit” menu button to undo the latest change. Hit the “Redo edit” to restore a change that you made by hitting undo. Drop the undo and redo menu buttons down to see a description of the change that will be removed or reinstated. To undo or redo to a specific change, select it from the list. All changes up to and including that change will be undone or restored.

When you edit a property in any property page, the changes you make are immediately reflected in the preview map window and any external map windows. Some edits will require entering a text string or a number into an edit box or some other kind of control. In this case, the edit is not committed until you finish typing the text. You can indicate that you have finished editing by hitting the return key. In ProRaster, hitting the return key does not close the application or other dialogs. Alternatively, you can change the focus to another control (in other words, click on something else).

Algorithms, layers, and components can all have names. It is advisable to name the algorithm because it is reflected in the algorithm list box and in the algorithm tree. Layer names are reflected in the algorithm tree and the system may automatically name an algorithm when you assign a raster source to the primary component of the layer. Component names are reflected in the algorithm tree as a [suffix].

If you are using a free or restricted version of ProRaster, then the main dialog may have a different appearance. Where functionality is restricted, buttons may not be visible. If you request an action, you may be presented with an Upgrade dialog instead. You may also see an upgrade property page for selected items in the algorithm tree, indicating that you are not authorised to modify the properties of that item.

User Interface Skins

If you are running ProRaster Scientific you can elect to emulate either ProRaster Essential or ProRaster Premium. Change the user interface skin by selecting the appropriate option from the system menu. To change back, simply close the application and restart it. ProRaster Premium is also able to emulate ProRaster Essential.

Working with the Preview Map

If your algorithm is valid, then it will be rendered in the preview map window. An algorithm can be invalid for a variety of reasons:

- The algorithm is disabled, or all layers are disabled, or there are no layers defined, or all components in layers are disabled.
- You have a component enabled but a raster source has not been selected for that component.
- You have a raster source selected for a component, but the raster could not be loaded. In this case the Field and Band drop-lists will be empty.
- You have not defined a data-color transform for an enabled Color, Red, Green, Blue, or Opacity component.
- Raster data cannot be reprojected into the algorithm coordinate system.
- Raster data, when reprojected into the algorithm coordinate system, is outside the allowable data range of that coordinate system.

The most difficult problems to diagnose are with coordinate systems. For example, imagine you have a LUT Color layer with color drawn from a projected coordinate system (like UTM) and you have a global raster in a geodetic coordinate system in the intensity component. The system determines that the algorithm coordinate system is projected, and it attempts to compute the algorithm data range by converting the range of the global geodetic raster to projected coordinates. These coordinates are likely to be invalid and this causes a rendering failure. In a situation like this you can use the "Zoom to selected extent" button to zoom to the bounds of a coordinate system, the extents of a raster source, or to a manual geodetic coordinate range.

Zooming and Panning

In the map window, use the mouse to zoom and pan. Roll the mouse wheel to zoom in and out about the location of the cursor. There is a natural limit to how far you can zoom out and once you reach this limit the map will fit to the algorithm extents. You can force the map to fit to the algorithm extents and redraw by hitting the "Zoom to algorithm extents" button. Pan the map by holding down the left mouse button and dragging.

Zoom Button

Sometimes the algorithm properties (reflected in the dialog and property pages) can get out of sync with the rendered preview. You can resolve the issue by hitting the "Zoom to algorithm" button, which also causes the algorithm to reset and reload.

Also, it is possible that the status of a layer or component (enabled or disabled) reflected in the algorithm tree control can get out of sync with the same property in the property pages. To resolve this, set and reset the Enable property in the property pages and the tree, properties and preview should all come back into sync.

Map Buttons

You can collapse the preview map window by clicking the "Close preview window" button. Once it is collapsed (and depending on your product tier) you can hit the "Open preview window" button to restore it, or you can hit the "Open a floating map" button to open an undocked and resizable map window. You may be able to open multiple undocked map windows. To close all undocked map windows, simply hit the "Open preview window" button. To control whether the ProRaster dialog stays on top of the map windows, see the Options dialog.

Drop down the "Zoom to selected extent" menu button for more zoom options. You will be able to zoom to the bounds of any coordinate systems used in the algorithm (if those bounds are known). You will be able to zoom to the extents of any raster source used in the algorithm. You can also select the "Zoom manually" option to open a zoom dialog in which you can specify a zoom rectangle in WGS84 longitude and latitude coordinates.

These zoom options are provided both for convenience, and for recovering control of the map window when the true extents of the algorithm cannot be computed. This can occur when the extents of a raster source are outside of the bounds of a coordinate system that those extents must be converted into. This conversion will often fail and return invalid coordinates, and this will cause the map to fail to render. Zooming to a valid extent for the algorithm coordinate system can recover the map window.

The manual zoom dialog now allows you to zoom to the bounding box of a polygon loaded from a MapInfo TAB file, in addition to manually defining a zoom bounding box. Once the polygon has been loaded and the bounds established, you can edit them if you wish. These bounds will be in the coordinate system of the polygon table.

The manual zoom dialog now allows you save the bounding box. Check the box to "Remember" the zoom rectangle and provide a name for it. Thereafter, the zoom rectangle will appear in the zoom drop-menu.

The zoom drop-menu now provides control over algorithm clipping boundaries. A clipping boundary is a rectangle, defined in a coordinate system, to which the algorithm will be clipped. When the algorithm is rendered, output pixels are clipped to these bounds. The clipping bound is saved with the algorithm to an MRD file. You can define a clipping boundary in the manual zoom dialog. Simple check the "Set the algorithm clipping bounds" button when you define your zoom boundary. To turn this clipping boundary off, select the "Zoom to algorithm (unbounded)" option in the zoom drop-menu.

Reporting


Tooltip Reporting Mode

As you move the mouse cursor about in a map window, ProRaster can report the raster value at that location as a tooltip. Drop down the "Tooltip reporting mode" menu button to control the amount of information displayed – No tooltip, Coordinates, Key Data Value, Key Data Information, or All Data Information.

- No tooltip

  No tooltip is displayed as you move the mouse cursor over the preview map or external map window.

- Coordinates

  The location of the cursor is reported. The format of the coordinates can be modified via the "Format Coordinates" sub menu. Choose between Decimal (X,Y), Geodetic Decimal (Latitude, Longitude), and Geodetic (Latitude, Longitude). Note that geodetic coordinates are always presented in a (Y, X) order corresponding to latitude and longitude, unlike all other coordinate systems.

- Key Data Value

  The key value is the value of the raster at the cursor location in the primary component of the layer. If there are multiple layers then the value of the first layer, from the top down, to provide a valid value at the cursor location is reported. The primary component is generally the color component, but if color is disabled it will be the intensity component.

- Key Data Information

  In this mode we report the key value, but also report the raster, field, and band from which that data was acquired. The location of the cursor is also reported, and the same formatting options apply - see “Coordinates”.

- All Data Information

  This is a verbose reporting mode that you would only use sparingly. We report the all the raster values, as well as the raster, field, and band from which that data was acquired. The location of the cursor is also reported, and the same formatting options apply - see “Coordinates”.

When you have chosen an option to format the coordinates in geodetic coordinates, ProRaster will convert the mouse cursor coordinates from the coordinate system of the algorithm to WGS84 latitude & longitude geodetic coordinates.

Report Property Page

To acquire a complete report on all the raster data at the cursor location, double click. A report will be displayed in the algorithm property sheet, on a property page labelled “Report”. This property page is automatically selected for your convenience.

The report is in XML, and you can copy-and-paste it into an editor of your choice. In the report you will find the exact location of the cursor and the coordinate system of the reported coordinates. It will then provide hit-testing information on every raster intercepted at that coordinate. These are reported from the top down (i.e., last rendered to first rendered).

There are two possibilities when reporting raster data at a location:

1. You can report the cell value, at base level resolution, with or without interpolation.
2. You can report the actual cell value used in the rendering engine, derived from an appropriate resolution level, probably acquired by interpolation (it depends on the interpolation properties in the algorithm).

ProRaster reports data type (2) – the actual data value used in the rendering engine. This means the point must have a scale. For each raster intercepted we report the coordinate range of the pixel that has been targeted, as well as the cell value of that pixel.

If the raster is a virtual raster (MVR) or a rendering algorithm (MRD) then the report will drill down into the source rasters that those virtual rasters depend upon. In this case, you will see the value of the virtual raster reported, and then the values of the source rasters reported. As virtual rasters can contain other virtual rasters, these chains can become quite deep.

Working with Raster Data

ProRaster supports a wide variety of raster file formats, depending on the product tier you are using. You can find a complete list at:
https://robertsgeospatial.com.au/products/proraster/help/supported-raster-formats/

There will often be other files alongside the raster:

- If you load the raster into MapInfo Pro, then there will be a table (TAB) file. ProRaster generates these files for your algorithm (MRD) file.
- If you render the raster in MapInfo Pro there will be a property (GHX) file. ProRaster also uses this file to record statistics for rasters.
- All rasters require an overview cache if they do not have one built-in. ProRaster generates PPRC and PERC overview cache files. GDAL generates OVR overview cache files. Formats like MRR and COG (Cloud Optimised GeoTIFF) have the overviews built-in to the primary raster file.

If your raster is stored on a read-only medium, then ProRaster will not be able to write the files alongside the raster. In the case of statistics, it means the GHX file cannot be written, and the statistics will not be stored. If an overview cache file is not present, then it will be generated – every time you open the raster. A temporary file will be created, and it gets deleted when you close the raster. If you want to store rasters on read-only medium, make sure you store the statistics and overview cache files alongside them. Otherwise, you may experience performance and data quality issues.

Raster Sources

Algorithms do not use rasters – they use raster sources. You can see a demo of the raster source editor on YouTube here. Raster sources are declared as global resources, and all the layers and components in the algorithm can connect to them. What is the difference between a raster and a raster source?

- A raster source may contain more than one raster.
- A raster source may contain rasters listed by filename or it may use wildcards (i.e., *.tif) to gather all files from a directory at run time. In other words, you can change the rasters in an algorithm just by moving files to and from a directory rather than editing the algorithm.
- A raster source can target files in zip archives by internal directory and file name.
- A raster source contains a driver preference data package, which is supplied to the raster engine whenever it opens any raster in the raster source. These driver preferences can modify the behaviour of the raster engine at run time.
- A raster source can use relative file paths.

The key difference is that a raster source can contain multiple rasters. So, when you connect a layer component to a raster source, you may be providing it with multiple rasters. There are some rules and expectations around this:

- We require that all the rasters have the same field, band, and event structure.
- We require that all the rasters have the same coordinate system.
- It is likely the rasters will not significantly overlap each other (although they can).
- It is likely the rasters will have a similar cell size (although they do not have to).
- It is likely the rasters will all have the same file format (although they do not have to).

A common scenario is where you have multiple raster images that are tiles of some larger raster. It is easier for the user to create a single layer to render these tiles and to define a raster source that includes all the tiles than to create a layer for each raster.

Where raster sources overlap, the order in which they are processed is important. Rasters in a raster source are processed in the order in which they appear in the list. You can see this list in the Raster Source Editor dialog. That means rasters at the top of the list will be processed first.

In an algorithm, layers are rendered in order using a "Painters" algorithm. This means that layers rendered first may be overprinted by layers that are rendered later. The same is true within a layer. Raster sources are rendered using a "Painters" algorithm. This means that whatever raster is rendered first may be overprinted subsequently. The consequence of this is that rasters within a raster source ought to be ordered from lowest quality and resolution to highest quality and highest resolution, to ensure that the highest quality data is presented.

Commonly, a raster source will use relative paths. In this case, the path must be relative to the file location of the MRD algorithm file. If you use relative paths, then you can store your data and algorithms as a package. When you move that package about, the algorithms remain valid even though the absolute paths of the source rasters may have changed. This is an option you can control from the Options dialog.

It is a good idea to make sure that statistics and overviews are properly prepared for your raster sources. You can do this via the Raster Source Editor dialog which provides three processing operations – Validate, Clean and Prepare. The Validate operation just looks at the raster and reports if any information is missing. The Clean operation deletes cache and other files. The Prepare operation computes statistics and generates overview caches. From here, you can exercise some control over the quality of the data in the overview cache, optimising either for speed or for quality.

For some rasters - TIFF files in particular - there may be a need to provide some additional information to the raster engine when it opens the file. These directives are supplied in driver preferences which you set in the Raster Source Editor dialog via the Advanced button. In some cases, this is critical to generation of appropriate overviews. For example, if you have a Landsat GeoTIFF file the value zero (0) may need to be treated as a null value, indicating an invalid cell. If you do not define this, the zero values will pollute the overview pyramids and blend with actual data to create rubbish on the edge of the scene. I noticed several years ago that all Landsat data uploaded to AWS suffers from this error.

Web Mapping Services

Support for Web Mapping Services (WMS and WMTS) is limited to ProRaster Premium and ProRaster Scientific. ProRaster Essential does not include support for WMS.

Web mapping services are supported via the GDAL "WMS" and "WMTS" drivers. For more information on this driver refer to the GDAL documentation at:

https://gdal.org/drivers/raster/wms.html
https://gdal.org/drivers/raster/wmts.html.

To open a WMS or WMTS, you firstly need to build a local service description XML file. You can treat this XML file as though it is a raster. You can drag-and-drop it onto ProRaster to generate a default algorithm and immediately display the web mapping data . You can add it to a raster source for use in any algorithm.

Note that GDAL provides an option to cache the map tiles on your local hard disk. You should not use this option. Make sure that the <Cache> block in the local service description XML file is completely deleted.

ProRaster will cache tiles in-memory as they are downloaded, providing performance optimisation without hard disk pollution.

A web mapping service relies on a vendor providing data, and that vendor may have terms of use that you ought to carefully consider before consuming the data. As a rule, you are likely to breach legal usage restrictions if you attempt to cache or export the vendors online data.

A collection of WMS local description service XML files is shipped with ProRaster and installed into the C:/ProgramData/RGE/WebMaps folder. You can access these files directly from the user interface via the "Open" menu button or the "Add Layer" menu button. Alternatively, drag and drop these files onto ProRaster to visualise them and explore the data.

You can add your own WebMap local service description files to the WebMaps folder and add sub- folder as necessary. These will be reflected in the UI when you restart ProRaster. The standard set of WebMap files are copied into the WebMaps folder each time you start ProRaster, so if you want to return to the default set of shipped files it is safe to delete the WebMaps folder and simply run ProRaster again.

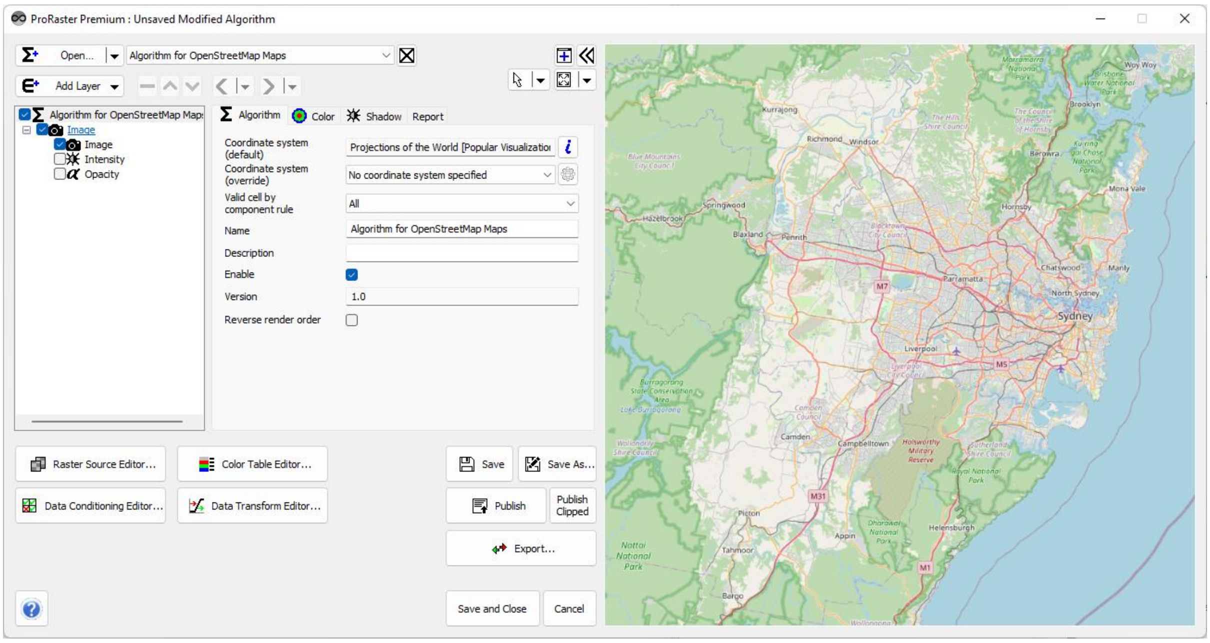

OpenStreetMap WMS layer displayed in ProRaster Premium

WebMap Overlays are a mechanism that is provided to display a WMS mapping service in an algorithm using greyscale and opacity modulation to enable the underlying raster layers to be visualised with a minimum of color degradation. A collection of WebMap Overlays are shipped with ProRaster and installed into the C:/ProgramData/RGE/WebMapOverlays folder. You can access these files from the user interface via the "Add Layer" menu button. If you develop your own overlay rasters you can add them to this folder to make them visible in ProRaster.

An overlay is developed by using a virtual raster (MVR) to convert a WebMap (WMS) raster to greyscale. This is implemented by a Calculator operation. A typical formula is "RED*0.299+GREEN*0.587+BLUE*0.114". The MVR is then displayed in a LUT Color layer and placed on top of all other layers. Blending in the algorithm is enabled. A data transform is used to darken the dark pixels and lighten the light pixels. Opacity modulation is enabled to allow dark pixels to be render opaquely and light pixels to be rendered transparently. You can adjust the scale of the WebMap using the "Adjust Resolution" control on the "Component" page of the "Color" component on the LUT Color overlay layer.

Working with Coordinate Systems

For an introduction to the science of coordinate systems, you can do no worse than visiting the online documentation for many of the major GIS application suppliers.

All raster datasets ought to have a coordinate system so that an application like ProRaster - or a GIS application - can locate the raster appropriately and display it in the context of other data which may use a different coordinate system. In the bad old days, most raster formats had no mechanism to store the coordinate system, but that is not the case today with modern raster formats.

The raster engine will attempt to identify the coordinate system of a raster when it mounts it. Unfortunately, this is a dark art as just about every raster format defines coordinate systems in their own way. The raster engine must interpret this data and convert it into a “standard definition” of a known coordinate system. This is further complicated by the fact that no industry accepted “standard definition” exists. EPSG codes are some kind of standard (see https://epsg.io/) and WKT (Well Known Text) is another “standard” way to express coordinate system data. In MapInfo Pro there is a “Category – Member” system that maps to coordinate system definitions. ProRaster uses the MapInfo system to specify coordinate systems in MRD algorithm files.

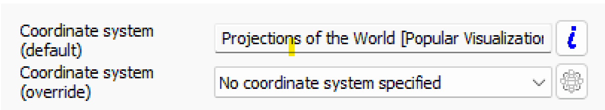

Coordinate System Selection Controls

In the ProRaster algorithm properties you will see two editable fields for coordinate system data at three different levels – for the algorithm, for each layer, and for each component in each layer. The first field is labelled “Coordinate system (default)”. The field shows you a readable text representation of the coordinate system, and to the right of the label is a button you can press to open an information dialog for that coordinate system. The second field is labelled “Coordinate system (override)” and contains a drop-list from which you can select one of the coordinate systems that you have recently used (or select “No coordinate system specified”). To the right of this list is a button that opens the Projection Selector dialog. This dialog helps you to select an appropriate coordinate system from the thousands available, after which it will be reflected in the list.

Component Coordinate System

Every enabled component needs to be connected to a raster source, which contains one or more rasters. Recall that the system assumes that all rasters in a raster source share the same coordinate system (if they don’t, rendering errors may occur). The first (default) field will reflect the coordinate system detected from the rasters in the raster source – if it is defined and valid. If no coordinate system could be identified from the raster source, then the field will state “No coordinate system detected or defined”.

You can use the second (override) field to specify the coordinate system of the raster source. You would only do this is there was no coordinate system defined, or the coordinate system could not be interpreted from the rasters, or the coordinate system was interpreted incorrectly.

Once you override the coordinate system, it will be reflected in the default coordinate system field.

Layer Coordinate System

The (default) coordinate system of a layer is acquired from the coordinate systems of the components in the layer. It acquires the first valid coordinate system, in the order in which they are listed in the algorithm tree. If there are multiple valid coordinate systems to choose from, it will prefer a geodetic (or geographic) coordinate system over any other kind. If you have overridden a coordinate system for a component, then it may get propagated up to the layer.

You can specify the coordinate system in the layer via the override field. The only effect this may have is to modify the default coordinate system of the algorithm. From a rendering perspective, the only coordinate systems that matter are the component coordinate system (because we assume this is the coordinate system of the raster data), and the algorithm coordinate system (because this is the coordinate system that you want to use when rendering).

I do not recommend overriding the coordinate system in the layer. It is better to make sure the coordinate systems are properly defined for each component.

Algorithm Coordinate System

The (default) coordinate system for the algorithm is acquired from the coordinate systems of the layers. It takes the first defined and valid coordinate system in the order in which the layers are defined (from bottom of the render order to top – or bottom of the tree to top). If there are multiple valid coordinate systems to choose from, it will prefer a geodetic (or geographic) coordinate system over any other kind.

In most cases the algorithm coordinate system will be the same as the rasters that are used in that algorithm, and that is generally the desired outcome.

You can use the override field to change the coordinate system of the algorithm. This has a dramatic effect. It will force all the rendering to be performed in the specified algorithm coordinate system, and all the raster data in the algorithm will be reprojected on-the-fly into this coordinate system.

Coordinate System Overflow Issues

When you are reprojecting from one coordinate system to another, data coordinate overflow issues can occur. Consider this (not so) theoretical example:

Imagine I have a satellite scene I want to drape on terrain. So, I have an RGB Color layer that is connected to a Sentinel-2 scene in a UTM projection. In the Intensity component I have a global terrain raster for hill-shading that is in a geodetic (longitude / latitude) coordinate system. The layer picks up the UTM coordinate system and so does the algorithm, and that is what I want. The preview map fails to display. Why?

It fails because the algorithm has failed to work out the data range. It reprojected the extremity coordinates of the geodetic data into the UTM projection, and this generated invalid (possibly infinite) coordinates. Everything went bad, and the map failed to display.

In this case, I can override the coordinate system at the layer or algorithm to the geodetic projection, and it will then display - but I am not using the desired projection. To work around this problem, use the “Zoom to selected extent” button to zoom to a coordinate system bounds, or to a raster source extent, or to a manual extent.

Working with Transparency

In ProRaster, the alpha color channel and all user interface controls refer to opacity. Opacity is the inverse of transparency. In an 8-bit color value, an opacity of 0 is completely transparent and an opacity of 255 is completely opaque. When we combine colors with different alpha values, this is referred to as blending. You can see a demo of transparency on YouTube here.

Blending opacity values occurs in three ways in ProRaster. Firstly, you may have an alpha channel in a located image raster that you want to render. You can honour and pass through these existing alpha values. Secondly, you can modify the opacity of layers in an algorithm so that you can see through layers that overlap. Thirdly, you can modulate opacity using a data – opacity transform on a per-pixel basis. This allows you to make features of interest in the raster opaque and make other data transparent.

Blending

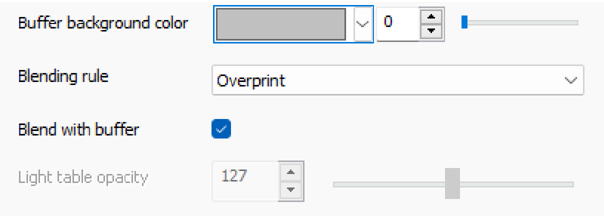

Blending Controls for an Algorithm

Blending is controlled from the algorithm. Select the algorithm item in the tree control, then click on the Color property page. The primary blending control is the “Blending rule”. You must set this property appropriately to perform any blending. It has the following values:

- Override

  Override mode completely disables blending. In this mode, the Opacity components are not used. Even if you enable the opacity component and populate it, it will have no impact on the rendered algorithm. Layer opacity settings are not used in this mode and all layers will be opaque. Alpha channels in located imagery will be ignored and all pixels will be opaque.

- Overprint

  Overprint is different to override in only one respect. When you have an Image layer and the Image component is connected to a located image that has an alpha band, those alpha values will be honoured and passed through the render engine. You can also achieve this using the Blend mode, so Overprint is somewhat redundant.

- Light Table

  Light Table enables a restricted form of layer blending that simulates a light table. (A light table is a glass work desk illuminated from underneath onto which you can place map sheets and overlay them.) In this mode, the opacity of all layers is controlled by the “Light table opacity” setting. Where layers overlap, they will be blended. A particular feature of this mode is that where layers are not overlapped, they will be completely opaque and shown with full color saturation. This reflects the nature of a light table – where there is nothing on top of the map, the map is not obscured.

- Blend

  This enables blending by layer and blending by pixel and blending by image alpha component. In a LUT Color or RGB Color layer, the opacity is determined by the layer opacity multiplied by the pixel opacity computed from the Opacity component (if enabled and defined). In an Image layer the opacity is determined by the layer opacity multiplied by the pixel opacity acquired from the image color (if present).

In summary:

- To turn off all blending, choose Override.
- To keep the alpha values of a located image, choose Overprint or Blend.
- When working with multiple overlapping layers, choose Light Table.
- When you want per-layer opacity control or per-pixel opacity modulation, choose Blend.

From the Color properties of an Image layer, you can turn on “Ignore Color Alpha” to prevent the alpha component of the color data source being used. You would use this option for Image layers in a scenario where you want to enable blending, but you do not want to use the alpha component in the RGBA data source the Image component is connected to. It has no effect for other layer types and no effect unless the blending mode for the algorithm is set to Overprint or Blend.

Blend with buffer

If your algorithm is semi-opaque, it follows that you can see through the algorithm. To what? You can see through to the background map canvas, which is filled with a light grey buffer background color by default. To see the background, you need to enable “Blend with buffer”, so that when your algorithm becomes transparent it fades to the background buffer color.

Note that in Light Table mode you can never see the background canvas as unobscured layers are treated as opaque.

If you disable “Blend with buffer” then the algorithm may appear that it is no longer transparent. Layers that overlay each other can still be faded in and out, but the bottom layer appears to be opaque regardless. In fact, the algorithm is still transparent, but that transparency is only being reflected in the alpha channel of the outputted color pixels. If you were to export this algorithm to a RGBA located image, the color is output as pure, and the alpha channel contains the opacity. When you display the located image in a GIS application, you can then blend the image with the background map canvas value in the GIS application.

In summary:

- Enable “Blend with buffer” when viewing your algorithm in the ProRaster map to make it appear as though it is transparent and blends to background
- Disable “Blend with buffer” when exporting located imagery to ensure that the output colors are pure.

If you have an Image layer connected to a raster that contains an alpha component, then you can use that alpha component by choosing Overprint or Blend mode. You can ignore that alpha component by choosing Override or Light Table mode. If you want to use blending but want to ignore the alpha component in an image, then render the image using a RGB Color layer instead of an Image layer.

Modulating Opacity

With the blending rule set to Blend, you can modulate opacity by enabling and defining an Opacity component in any layer. Note that opacity modulation overrides image opacity color components.

To modulate opacity, you must enable the Opacity component, connect a raster source and select a field and band, and define a data transform. This is just like any other color modulation, but in this case the color table is a table of 256 opacity values from transparent (0) to opaque (255). By default, high data values will map to high opacity values. As an example, if you connect a terrain raster and use a standard transform like Linear, then the high ground will tend to be opaque and the low ground will tend to be transparent. If you want an inverse correlation, then it can be easier to reverse the color table than to reverse the data transform.

One useful opacity modulation is to modify a greyscale maps image so that colors that are whiter (like empty whitespace and parks etc) are transparent and colors that are darker (like roads) are opaque. In this case you want to map low values (black is 0) to high opacity and high values (white is 255) to low opacity. To do so, reverse the opacity color table.

In the image below, the opacity of a terrain raster is modulated by itself. Terrain that is lower will have a lower opacity (higher transparency) and fade to the background color, and higher ground will have a higher opacity (lower transparency).

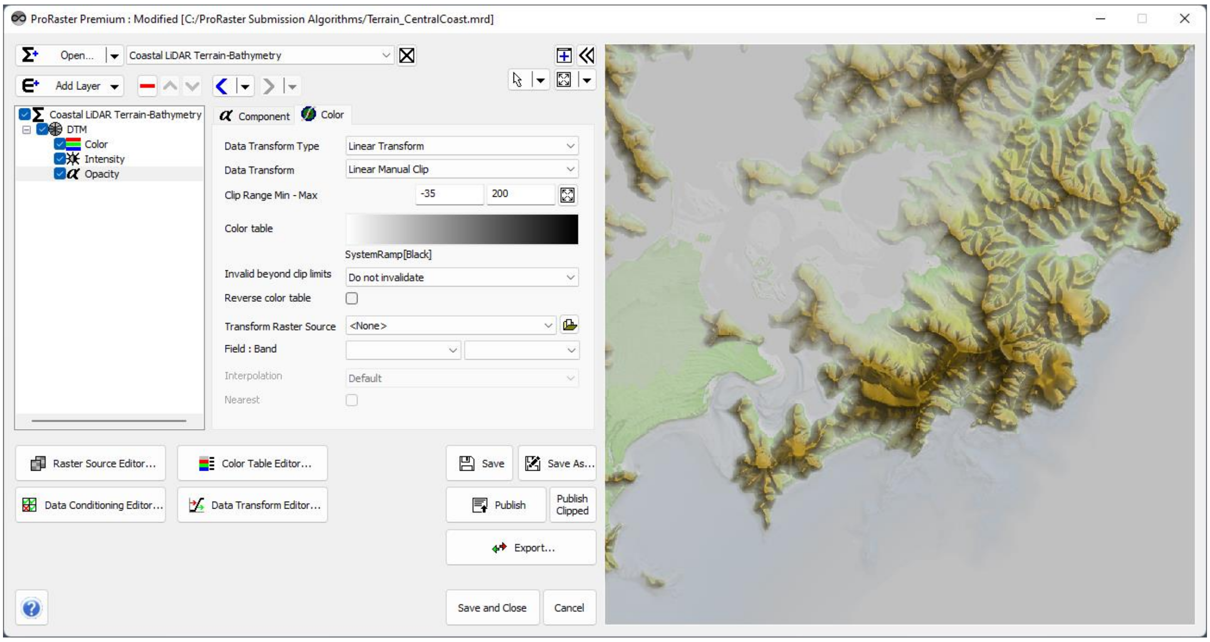

Terrain data, opacity modulated by itself

Working with Global Resources

An algorithm will contain a variety of named global resources that the algorithm will refer to (by that name) in the layer and component properties. These resources are either Raster Sources, Color Tables, Data Transforms, or Data Conditioning Filters.

These resources can be created and stored in ProRaster and then used in any algorithm as desired. For example, you can create a Raster Source and ProRaster will store those details. Thereafter, you can use that raster source in any algorithm that you create. This is convenient if you tend to reuse particular rasters frequently. It also helps you to go back to raster sources you have created in the past easily. When an algorithm uses a resource defined and stored in ProRaster, it gets fully copied into the algorithm MRD file. Consequently, MRD files are self-contained and do not depend on information in ProRaster, although they may depend on files that are referred to in that resource (like a color table file for example).

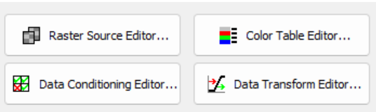

Resource Editor Dialogs

Resources are created and edited by four editor dialogs – the Raster Source Editor, the Color Table Editor, the Data Transform Editor, and the Data Conditioning Editor. Once you have created a resource, you will then refer to it by name in the algorithm property page. For example, for each component you must select a raster source. For a LUT Color layer Color component, you must choose a color table. For any component, you can choose a data conditioning filter to apply to the raster data prior to rendering. And color and opacity components, you need to select a data transform to transform from data to color.

When you create a resource, the first thing to do is to name it. You must specify a unique name for each resource you create (within that resource category). The editor dialogs help you to do this by issuing a warning if you have not chosen a unique name for a resource.

Sometimes, you can take a shortcut and create a resource without using the editor dialogs. For example, you can create an algorithm by browsing to a raster and you can populate a component by browsing to a raster. In both cases, that raster will be added to a new raster source behind the scenes.

Some resources are created on demand according to established rules. For example, there are a wide variety of standard data transforms that you can use, and many of these can be customised with specific clipping ranges and other parameters. These resources are not named or stored; they are simply created on demand.

Resource object definitions as well as temporary application data are stored in binary files on your C:\ drive. If you delete one or more of these files, then you will lose temporary settings (like your Options selections) as well as resource objects that you have created. However, it will not prevent ProRaster from functioning and these filles will be recreated when needed.

Data that is accessed by all users and all apps in the product family is stored here:
C:\ProgramData\RGE\ProRasterUserResources.dat

Data that is local to a user (called sam) and used by all apps in the product family is stored here:
C:\Users\sam\AppData\Local\RGE\ProRasterUserResources.dat

Data that is local to a user (called sam) and used by a specific app (called ProRaster Premium) is stored here:
C:\Users\sam\AppData\Local\RGE\ProRaster Premium\UserResources.dat

When you open an existing algorithm MRD file, it will contain a variety of global resources. In many cases, these can be matched to existing resources that you have defined via the editor dialogs. If so, ProRaster will use this resource when saving the MRD. If the resource is unknown, it will be stored temporarily in ProRaster and removed when the algorithm is closed.

Algorithm Properties

Access the algorithm properties by clicking on the algorithm item at the top of the tree control. The algorithm property sheet will be displayed containing the Algorithm, Color, Shadow, and Report property pages.

What property pages and the properties presented in those pages depends upon the product tier and the product licensing.

Properties can be one of four different types, and this is true of properties for algorithms, layers, and components:

- Local
  The property applies at algorithm level to operations that occur at this level of scope.
- Default
  The property, when enabled at the algorithm level, provides a default value that will be used by subordinate levels (layers and components), unless they specifically override the property themselves.
- Override
  The property, when enabled at the algorithm level, overrides similar properties that may be defined at layer scope or component scope.
- Propagated
  A default value for the property is acquired by propagating up the tree from lower-level components and layers but can be overridden at the algorithm level.

Algorithm page

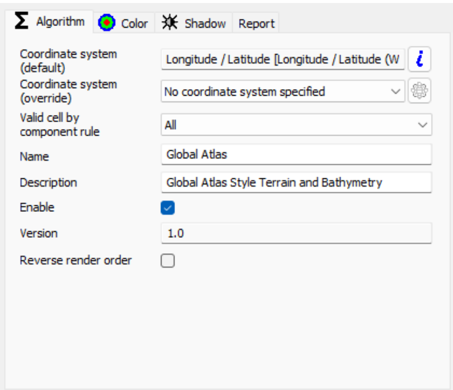

Algorithm Property Page

Name: You should always name your algorithm as this name is reflected in the algorithm list and in the tree. When you export an algorithm, it is a requirement that the algorithm be named, and if it isn't you will be warned.

Description: A short, optional, description of the algorithm.

Coordinate system (default): This is a *propagated* property and is acquired from layers, components, or rasters used in the components. It reflects the coordinate system that the algorithm will be rendered in. All raster data used in the algorithm will be reprojected on-the-fly, if required, from its known coordinate system into this coordinate system for rendering. If you override the coordinate system (see next item), this setting will change to reflect that.

Coordinate system (override): To specify the coordinate system that the algorithm will be rendered in and to override the default coordinate system, select a recently used coordinate system from the drop-list or hit the button to select a coordinate system using the Projection Selection dialog.

Valid cell by component rule: This is a *default* property which can be overridden in individual layers.

When a layer is rendered the rendering engine acquires data at the pixel locations from every enabled and valid component. The data for each component may come from different rasters or different fields and bands, and so the cell validity can vary for each component. We may have valid cell data in one or more components and invalid (empty or null) data in other components.

This rule allows you to control whether a pixel color will be generated, based on whether the component data is valid or invalid. The two most important settings are “All” and “Any”. The “All” setting tells ProRaster that the data from every component must be valid to generate a pixel color. This is like rendering the intersection of all valid data in the layer. The “Any” option means that a color will be generated if any component contributes valid data. This is like rendering the union of all valid data in the layer. The VCBC rule can have the following values:

- All
  The data from every component must be valid to generate a pixel color.
- Primary
  The data from the primary component must be valid to generate a pixel color. The primary component is the first enabled component in the tree. It will generally be a color or image component.
- Secondary
  The data from the secondary component must be valid to generate a pixel color. The secondary component is the second enabled component in the tree. It will generally be Intensity. Note that in a RGB Color layer the Red, Green, and Blue components are all considered primary.
- Any
  A pixel color will be generated if any component contributes valid data.

In addition to these simple rules, there are multi-stage rules. These rules cause the rendering engine to render each layer two or three times. This will reduce rendering performance in most scenarios. In a multi-stage rendering scenario, a pixel that is valid from an earlier stage will not be overprinted by a valid pixel in a later stage.

- All to Any
  This triggers a three-stage rendering pipeline for every layer. We first render using the All rule, then the Primary rule, then the Any rule.
- All to Primary
  This triggers a two-stage rendering pipeline for every layer. We first render using the All rule, then the Primary rule.
- All to Secondary
  This triggers a two-stage rendering pipeline for every layer. We first render using the All rule, then the Secondary rule.
- Primary to Any
  This triggers a two-stage rendering pipeline for every layer. We first render using the Primary rule, then the Any rule.

The VCBC rule you select can result is vastly different rendered outcomes, and it can also have a very serious impact on rendering performance.

- “Any” may be slower than “All”. If you select “All” then the rendering engine finds the intersection of the component data. This may result in a smaller (or empty) area to render and improve performance.
- “All to Any” renders three times per layer and the other multi-stage options render twice per layer. This will naturally result in lower performance (but not always!).
- If you have raster sources that contain multiple files the situation is more complicated!

Consider the following scenario: I have an Image layer with a raster source in the Image component that contains hundreds of geotiff located topographic images in a tiled configuration. I have a global scale terrain raster (with no bathymetry) in the Intensity component onto which I am draping the imagery.

If I use “All” I get high performance because most of the tiff files will not intersect the scene. But my scene is clipped to the shoreline as I have no terrain data offshore. It is also clipped to the edges of the topographic imagery.

If I use “Primary” I get high performance and my topographic imagery now renders offshore. The algorithm accepts pixels with the image but not the intensity components.

If I use “Secondary” I get low performance and all I see is the terrain - my imagery disappears. This is unexpected and explained by the multi-file raster source. For each source raster it renders, the engine finds it has terrain data but, usually, no intersecting topographic imagery. It populates the pixel with terrain only and it will not overwrite that pixel again. So, we get a bad result and a slow result because it is doing plenty of work for every raster in the raster source.

If I use “Any” I get low performance and no onshore topographic imagery. This is unexpected and explained by the multi-file raster source. For each source raster it renders, the engine finds it has terrain data or maybe offshore topographic imagery but, usually, no intersecting onshore topographic imagery. It populates the pixel with terrain or offshore imagery, and it will not overwrite that pixel again. So, we get a bad result and a slow result because it is doing plenty of work for every raster in the raster source.

If I use “All to Any” I see everything, and it is fast. In the first pass I render “All” and any pixels that get populated will be kept. Then I render “Primary” and this fills in the offshore topographic imagery. Then I render “Any” and this fills in the global terrain. The “Any” phase is fast because all the data has already been cached by the previous phases and most pixels are already valid and are not rendered.

If I use “All to Primary” I see the topographic imagery onshore and offshore and terrain clipped to this and it is fast. The “All” phase renders the onshore topographic imagery, then the “Primary” phase renders the offshore topographic imagery.

If I use “All to Secondary” I see the onshore topography and the terrain. The “All” phase renders the onshore topographic imagery, then the “Secondary” phase renders the terrain. Offshore topographic imagery is not rendered.

If I use “Primary to Any” I see the onshore and offshore topographic imagery and the terrain, and it is fast. The “Primary” phase renders the onshore and offshore topographic imagery and the “Any” phase renders the terrain.

In summary:

“All” works : we see onshore imagery intersected with terrain.
“Primary” and “All to Primary” works: we see onshore and offshore imagery intersected with terrain. In this case choose “Primary” because you don’t need multi-stage rendering.
“All to Any” and “Primary to Any” works: we see onshore and offshore imagery and terrain. In this case, choose “Primary to Any” because it only involves two stages, not three.
“All to Secondary” works: we see onshore imagery and terrain.

You can see a video on YouTube explaining this scenario at:
https://www.youtube.com/watch?v=4D_TrSU7ju8

Reverse render order: Reverse the order in which the layers are rendered. By default, they are rendered in the order in which they are listed in the MRD file which corresponds to the bottom of the tree to the top.

Enable: Enable or disable the algorithm from the tree or from this property.

Version: A version number like "1.0", optional and currently not used.

Color page

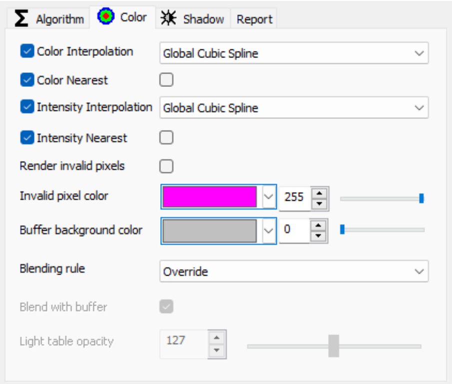

Algorithm Color Property Page

Color Interpolation: This is an *override* property. It will override the same property in layers and components. This sets the interpolation method used to generate underview tiles for color and opacity components. Toggle the checkbox to override the property and select from the drop-list. The following methods are available:

Nearest neighbour: Underview cells are populated with the value of the base level cell that overlaps them. In ProRaster, underview cells never partially overlap a base level cell. Cell values are matched exactly.
Linear: Bi-directional linear interpolation. Will not overshoot or undershoot.
Local cubic spline: Bi-directional cubic spline operator. Will overshoot and undershoot.
Global cubic spline: Bi-directional cubic spline computed for the entire tile (smoothest). Will overshoot and undershoot.
Default: Chooses the highest quality method based on field type and band data type.

This interpolation method will be used for the Color component in a LUT Color layer, the Red, Green, and Blue components in an RGB Color layer, and the Image component in an Image layer. It is also used for Opacity components.

Color Nearest: This is an *override* property. It will override the same property in layers and components. Tile data at any resolution level may have to be interpolated to provide data that is aligned with the pixel grid and centred on the pixel centres. ProRaster works to try to avoid this in most scenarios to maximise data and image quality. When ProRaster needs to interpolate, it will use either nearest neighbour or bi-linear interpolation. The choice will depend on the field type and band data type. If possible, it will default to using bi-linear. You can override this by toggling the override checkbox and checking the nearest neighbour box.

This interpolation method will be used for the Color component in a LUT Color layer, the Red, Green, and Blue components in an RGB Color layer, and the Image component in an Image layer. It is also used for Opacity components.

Intensity Interpolation: This is an *override* property. It will override the same property in layers and components. This sets the interpolation method used to generate underview tiles intensity components. Toggle the checkbox to override the property and select from the drop-list.

Intensity Nearest: This is an *override* property. It will override the same property in layers and components. If possible, it will default to using bi-linear. You can override this by toggling the override checkbox and checking the nearest neighbour box.

Buffer background color: When you render the algorithm to a map or when you export a located image, it will be rendered in a rectangular pixel buffer. This buffer will be pre-filled with pixels of the selected color. Set the RGB components of the color via the color picker control and set the alpha component from 0 to 255.

Render invalid pixels: Raster cells that are invalid are generally not rendered and so will not overwrite the buffer background color. In rare situations, you may want to render invalid cells in some special color. You can enable this here.

Invalid pixel color: The RGBA color used to render invalid cells.

Blending rule: This rule controls what transparency and opacity modulation is allowed in an algorithm. It has the following values:

Override: Completely disables blending.
Overprint: Modulates pixel opacity in an Image layer by the alpha component of the image component raster source.
Light Table: Blends overlapping layers together by simulating a light table.
Blend: Enables blending by layer and blending by pixel and blending by image alpha component.

Blend with buffer: Enable this option to blend your imagery with the buffer background color so that it appears the algorithm fades to background in the map. When you export your imagery, you may want to disable this option to ensure that the background color does not bleed into the imagery color.

Light table opacity: Sets the opacity of every layer in the algorithm from 0 (transparent) to 255 (opaque) when in Light Table blending mode.

Shadow page

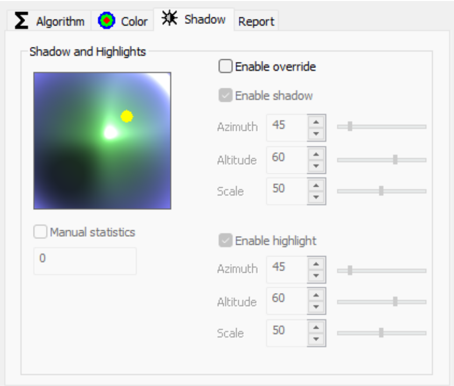

Algorithm Shadow Property Page

Enable override: This is an *override* property that will override the shadow and specular highlight settings for every Intensity component in every layer in the algorithm. Check the box to control these properties from the algorithm.

The appearance of the shadow and specular highlight is determined by the azimuth and altitude of an artificial sun. You can set these parameters for the shadow and for the highlight independently. You can enable one or both features at the same time.

A graphical control shows the position of the sun as a filled yellow circle, and the position of the specular highlight sun as a star. You can drag the sun around the control by left clicking, holding, and dragging. When you release the button, the properties will be updated. You can drag the highlight sun around the control by right clicking, holding, and dragging. In both cases the preview graphic is updated to show what the effect looks like.

In addition to moving the sun angle, you can change the scaling, or depth, of the shadow and highlight. To ensure that shadow is approximately uniform, regardless of the scale or units of the raster values, the scaling factor depends on a statistic computed from the raster. If this is unsuitable, you can manually override this value. Specifying a suitable value is largely guess work and trial and error.

Enable shadow: Check this box to enable hill-shading with artificial shadow.
Azimuth (shadow): The sun azimuth can be varied between 0 and 360 degrees.
Altitude (shadow): The sun altitude can be varied between 0 and 90 degrees.
Scale (shadow): The shadow depth can be varied between 0 and 100.

Enable highlight: Check this box to enable specular highlights.
Azimuth (highlight): The sun azimuth can be varied between 0 and 360 degrees.
Altitude (highlight): The sun altitude can be varied between 0 and 90 degrees.
Scale (highlight): The highlight scaling can be varied between 0 and 100.

Manual statistics: Enable this option and define the average cell-to-cell difference to set scaling manually.

Report page

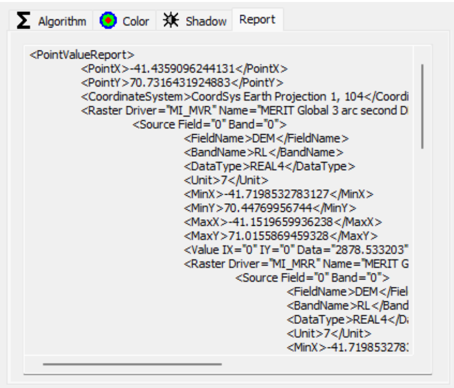

Algorithm Report Property Page

The report page contains no editable properties. If you double click in the map, the report page will be opened and populated with an XML report. You can copy and paste this out into any XML or text editor for inspection. The report contains information on the data values of all the rasters that contribute to the algorithm at the location you double clicked.

Layer Properties

Access the layer properties by clicking on the layer item in the tree control. The layer property sheet will be displayed containing the Layer, and Color property pages. Properties can be one of four different types – local, default, override and propagated.

What property pages and the properties presented in those pages depends upon the product tier, the product licensing, and the layer type.

Layer page

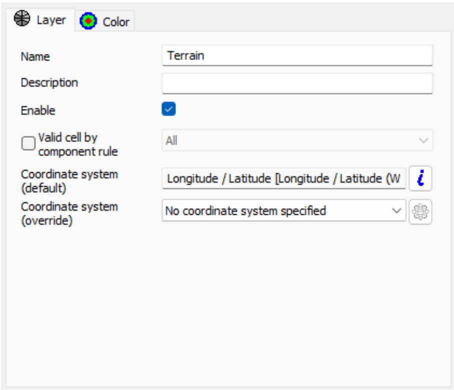

Layer Property Page

Name: You are encouraged to name your layers as this is reflected in the tree. Layers are named by default when you first connect a raster source to the primary component of a layer.

Description: A short, optional, description of the layer.

Enable: Enable or disable the layer from the tree or from this property. This will prevent the layer being rendered.

Valid cell by component rule: A default value for the VCBC rule is acquired from the algorithm properties, but you can override it for individual layers. Override the property by checking the box then selecting the VCBC rule. For more information, see the section on algorithm properties.

Coordinate system (default): This is a *propagated* property and is acquired from components, or rasters used in the components. In turn, it may be propagated up to the algorithm. If you override the coordinate system (see next item), this setting will change to reflect that. Note that this coordinate system, apart from it's role in propagating a coordinate system to the algorithm, is never used when actually rendering.

Coordinate system (override): To specify the coordinate system and to override the default coordinate system, select a recently used coordinate system from the drop-list or hit the button to select a coordinate system using the Projection Selection dialog.

Report Cell Value: Turn this option off to prevent the layer contributing to cell value reporting. This affects the cell value tooltip and the cell value report acquired by double clicking on the map. Reporting is enabled by default, except when you add an Image layer to an existing algorithm. The new image layer will have reporting disabled by default.

A common scenario in which you would use this option is when you are displaying raster data with a partially transparent WebMap Image layer overlain. Turning off reporting on the image layer ensures that the underlying raster value is reported instead of the RGB color of the image layer pixels (which may be of no interest).

Color page

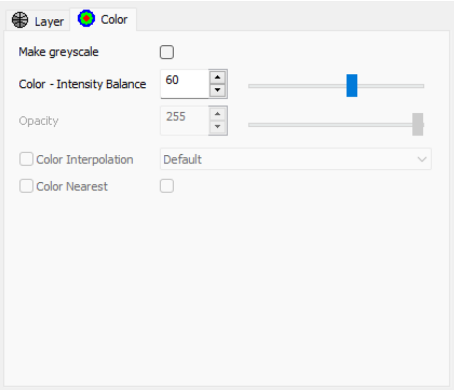

Layer Color Property Page

Make greyscale: Check the box to modify the color of the rendered layer to greyscale.

Color – Intensity balance: This control is provided to achieve a visually pleasing balance between the saturation of the color in the layer and the intensity of the hill-shading. The value ranges from 0 to 100. At zero, color saturation is minimum and hill-shading intensity is maximum. From 0 to 50 the color saturation increases until it reaches a maximum at 50. From 50 to 100 the hill-shading intensity decreases until it reaches a minimum at 100. For example, drag the control up to about 60 – 70 so fade out the hill-shading a bit and focus more on the color imagery.

Opacity: The opacity of the layer from 0 (transparent) to 255 (opaque). You will need to change the algorithm blending mode to "Blend" for this control to be accessible.

Color Interpolation: This is an *override* property. It will override the same property in components. This sets the interpolation method used to generate underview tiles for color and opacity components. Toggle the checkbox to override the property and select from the drop-list. The following methods are available - Nearest neighbour, Linear, Local cubic spline, Global cubic spline, and Default.

This interpolation method will be used for the Color component in a LUT Color layer, the Red, Green, and Blue components in an RGB Color layer, and the Image component in an Image layer. It is also used for Opacity components.

Color Nearest: This is an *override* property. It will override the same property in components. Tile data at any resolution level may have to be interpolated to provide data that is aligned with the pixel grid and centred on the pixel centres. ProRaster works to try to avoid this in most scenarios to maximise data and image quality. When ProRaster needs to interpolate, it will use either nearest neighbour or bi-linear interpolation. The choice will depend on the field type and band data type. If possible, it will default to using bi-linear. You can override this by toggling the override checkbox and checking the nearest neighbour box.

This interpolation method will be used for the Color component in a LUT Color layer, the Red, Green, and Blue components in an RGB Color layer, and the Image component in an Image layer. It is also used for Opacity components.

Ignore Color Alpha: This control is only visible for Image layers and is only enabled when the algorithm blending mode is enabled. It has no effect in LightTable mode. In Overprint or Blend mode, it can be used to ignore the alpha values that may be acquired from the color data source. It allows this alpha data to be ignored even when the algorithm uses blending. The opacity of the layer may still be controlled via the Opacity slider and edit controls.

Component Properties

Access the component properties by clicking on the component item in the tree control. The component property sheet will be displayed containing the Component, Color, and Shadow property pages. Properties can be one of four different types – local, default, override and propagated.

What property pages and the properties presented in those pages depends upon the product tier, the product licensing and the layer and component type.

Component Page

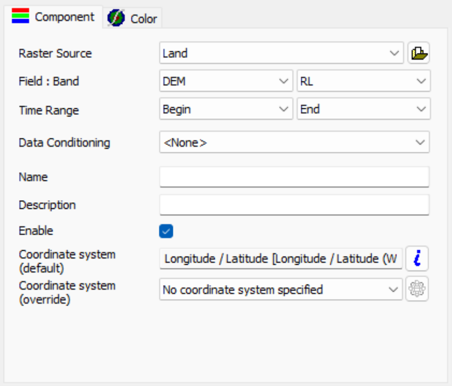

Component Property Page

Name: The component name is reflected in the tree. Components are named by default when you first connect a raster source.

Description: A short, optional, description of the component.

Enable: Enable or disable the component from the tree or from this property. If you disable a component, then it will not contribute to the rendering of a layer. For a layer to render you must enable at least one component, and that cannot be Opacity.

Coordinate system (default): This is a *propagated* property and is acquired from rasters used in the component. It may be propagated up to the layer and algorithm. If you override the coordinate system (see next item), this setting will change to reflect that.

Coordinate system (override): To specify the coordinate system and to override the default coordinate system, select a recently used coordinate system from the drop-list or hit the button to select a coordinate system using the Projection Selection dialog. By doing so, you are telling ProRaster what the coordinate system of the rasters is in the raster source used by the component.

Raster Source: Drop down the raster source list to select a raster source for this component. If you have not yet created a raster source, you can hit the Browse button to browse to a raster file.

Field : Band: Once you have defined the raster source, you will be presented with a list of fields and a list of bands for that field. Select a field and band.

Time Range: Once you have defined the raster source and the field, you will be presented with a list of time points (events) in the raster. Select the time range that you want to display from beginning to end inclusive. When you change the time range, by default it will be changed in all components in the layer. This is controlled by a checkbox to the immediate left of the time range controls.

Data Conditioning: Select a data conditioning filter, if desired, from the list of filters you have prepared with the Data Conditioning Filter Editor. This filter will be applied to the raster data during rendering.

Adjust Resolution: When data is acquired from the raster it will be taken from an appropriate resolution level. This option allows you to take the data from a resolution level that is up to three levels above the appropriate level.

To use this feature, set the adjustment factor to a value greater than zero. Resolution adjustment can be useful when working with WebMap (WMS) rasters. Often an adjustment of “1” is appropriate for these layers as it makes the text readable and reduces the IO requirements.

Color Page

Data Transform Type: To transform from data to color space, you can select from a wide variety of system defined transforms, data transforms you have hand crafted in the Data Transform Editor, or data transforms in color tables, maps, and legends. Firstly, select the data transform category shown in this list, and then select the required transform (below).

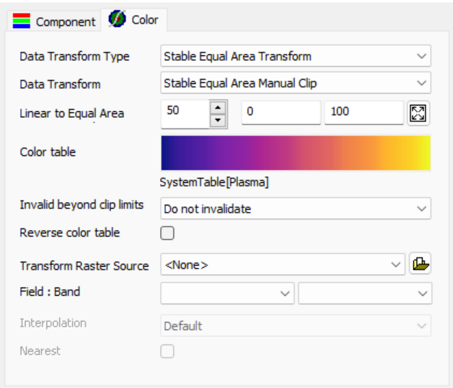

Component Color Property Page

Data Transform: For any data transform category there may be multiple data transforms defined. For system transforms, these allow you to vary clipping ranges and other parameters. The categories and transforms available are listed below. In a RGB Layer, the Red Component Color page will have a check box next to the Data Transform Type that allows you to mirror any changes to the data transform in the Red component in the Green, Blue, and Pan components.

Local Transforms: If you have loaded an MRD and the data transforms in that MRD are not standard or matching a known transform, then they will be listed here.

No Transform *(Pass Scaled, Pass Index, Pass Value, Rotate Index)*: A variety of restricted transforms that you can use when the data is already partially transformed. Pass Scaled can be used if your data is already scaled between zero and one. Pass Index can be used if your data is already a zero-based integer index matching the number of colors in the color table. Pass value can be used if your data is already a color. Rotate index can be used if your data is a set of zero-based integers. From a set of N colors (matching the number of colors in the color table), it uses the modulus function to assign a color index from each integer index.

Color transform *(Color bits)*: A linear transform with minimum at zero and maximum at 2^N-1. Enter the number of bits of color resolution.

Linear Transform *(Linear, 1% clip, 5% clip, 10% clip, 0.5%Q clip, 2%Q clip, 5%Q clip, Manual Clip)*: A linear transform where the minimum and maximum are established from data statistics, or input manually. The percentage clip options clip the bottom N% and top N% of data range. The percentile clip options clip the bottom N%Q and top N%Q of the data distribution.

Whenever a manual clip option is selected, you must enter the minimum and maximum values of the transform range. You can hit the “Fit” button to automatically set the minimum and maximum values to the statistical range of the data displayed in the map. This option allows you to stretch the color over the visible data range rather than the entire data range and create a more visually appealing local rendering.

Logarithmic Transform *(Various options + Manual Clip)*: A logarithmic transform where the minimum and maximum are established from data statistics, or input manually. Log base 10 is used in all cases. The transform will be linear between a specified (+/-) range about zero. These ranges are selectable – 1, 10, 100, 1000, 0.1, 0.01, 0.001, 0.0001.

Equal Area Transform *(Equal Area, 1% clip, 5% clip, 10% clip, 0.5%Q clip, 2%Q clip, 5%Q clip, Manual Clip)*: An equal area transform, stretching color evenly over the data distribution, where the minimum and maximum are established from data statistics, or input manually. You can control the transform using the “Linear to Equal Area” option which ranges from 0 to 100. At zero the transform is a linear transform, and it transitions to a full equal area transform at 100.

Stable Equal Area Transform *(Stable Equal Area, 1% clip, 5% clip, 10% clip, 0.5%Q clip, 2%Q clip, 5%Q clip, Manual Clip)*: An equal area transform, stretching color evenly over the data distribution, where the minimum and maximum are established from data statistics, or input manually. The transform uses a equal frequency histogram instead of a equal width histogram. These histograms are resistant to outliers and highly skewed data distributions.

Sigmoid Lightening Transform *(Lighten, 1% clip, 5% clip, 10% clip, 0.5%Q clip, 2%Q clip, 5%Q clip, Manual Clip)*: A sigmoid function where the minimum and maximum are established from data statistics, or input manually.

Sigmoid Darkening Transform *(Darken, 1% clip, 5% clip, 10% clip, 0.5%Q clip, 2%Q clip, 5%Q clip, Manual Clip)*: A sigmoid function where the minimum and maximum are established from data statistics, or input manually.

Ranges *(Ranges, Manual Clip, Quantile, About Mean, About Median, About Mode, Deviations, Jenks)*: A specified number of ranges where the minimum and maximum are established from data statistics, or input manually. Ranges are based on equal width data ranges, between the minimum and maximum data values or the clipping limits entered manually. Quantile ranges are based on data distribution. Ranges about the mean, median, and mode are centred on those statistical values and half the ranges are above and half are below. The range width is based on statistics. Deviations about the mean are centred on the mean and range widths are one standard deviation. Jenks natural breaks uses a statistical algorithm to determine the best distribution of ranges for the dataset.

Stored Transforms: In this category are listed all the transforms that you have created manually using the Data Transform Editor dialog.

System color maps: Lists all the color maps in the C:\ProgramData\RGE\ColorTables\SystemMaps directory. You can specify a clipping range for these maps, even though this does not affect the color mapping. When you select a color map, the color table is automatically changed to match.

User defined color tables: Lists all the color maps in the C:\ProgramData\RGE\ColorTables\UserTables directory. You can specify a clipping range for these maps, even though this does not affect the color mapping.

External folder: Every external color table folder that you have connected to ProRaster will be listed as a separate category. All the color maps in that folder will be listed. You can specify a clipping range for these maps, even though this does not affect the color mapping.

Color Table: This control shows a preview of the currently selected color table. Click on the control to drop down a list of all color tables. Recently selected tables are shown at the top, and then all the color tables connected to ProRaster. Double click to select a color table.

Invalid beyond clip limits: Data values that outside of the clipping limits of the transform are normally assigned the bottom and top colors in the color table. This option enables you to prevent rendering pixels that are outside the data transform clipping range. You can clip beyond the limits or within the limits in a variety of combinations.

Reverse color table: Reverse the color table. This is particularly useful for Opacity components and can be a convenient alternative to reversing the data transform.

Transform Raster Source: Most transforms rely on some statistics derived from a raster. You can choose to acquire those statistics from a different raster to the one you are rendering. Select a raster source or browse to a raster or select <None> to disable this feature.

Field : Band: If you choose a transform raster source, you must also choose a field and band from which to acquire statistics.

Interpolation: If this property has not been overridden at the layer or algorithms levels, you can set the interpolation method used to generate underview tiles for color and opacity components. Select from the drop-list. The following methods are available - Nearest neighbour, Linear, Local cubic spline, Global cubic spline, and Default.

Nearest: If this property has not been overridden at the layer or algorithms levels, you can direct the system to use nearest neighbour interpolation when interpolating from tile to pixel grid centres. Otherwise, the default is bi-linear unless this is not supported by the field type and band data type.

Shadow Page

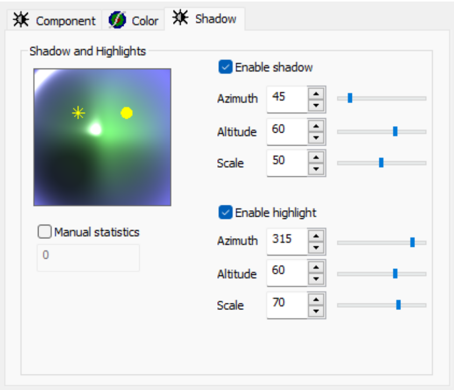

Component Shadow Property Page

This property can be controlled from the algorithm. If so, the controls will be disabled. Otherwise, you can control the shadow and specular highlighting for the layer from here.

The appearance of the shadow and specular highlight is determined by the azimuth and altitude of an artificial sun. You can set these parameters for the shadow and for the highlight independently. You can enable one or both features at the same time.

A graphical control shows the position of the sun as a filled yellow circle, and the position of the specular highlight sun as a star. You can drag the sun around the control by left clicking, holding, and dragging. When you release the button, the properties will be updated. You can drag the highlight sun around the control by right clicking, holding, and dragging. In both cases the preview graphic is updated to show what the effect looks like.

In addition to moving the sun angle, you can change the scaling, or depth, of the shadow and highlight. To ensure that shadow is approximately uniform, regardless of the scale or units of the raster values, the scaling factor depends on a statistic computed from the raster. If this is unsuitable, you can manually override this value. Specifying a suitable value is largely guess work and trial and error.

Enable shadow: Check this box to enable hill-shading with artificial shadow.
Azimuth (shadow): The sun azimuth can be varied between 0 and 360 degrees.
Altitude (shadow): The sun altitude can be varied between 0 and 90 degrees.
Scale (shadow): The shadow depth can be varied between 0 and 100.

Enable highlight: Check this box to enable specular highlights.
Azimuth (highlight): The sun azimuth can be varied between 0 and 360 degrees.
Altitude (highlight): The sun altitude can be varied between 0 and 90 degrees.
Scale (highlight): The highlight scaling can be varied between 0 and 100.

Manual statistics: Enable this option and define the average cell-to-cell difference to set scaling manually.

Using the Raster Source Editor dialog

In an algorithm, all rasters are stored in raster source resources, and those resources are referred to by name in the algorithm. In ProRaster you can take shortcuts to access rasters, but behind the scenes your rasters will still be in a raster source. Creating raster sources is easy and they will save you time in the future as you will always be able to come back to a specific raster in ProRaster without trying to remember where you have squirrelled it away on disk. See a YouTube demo here.

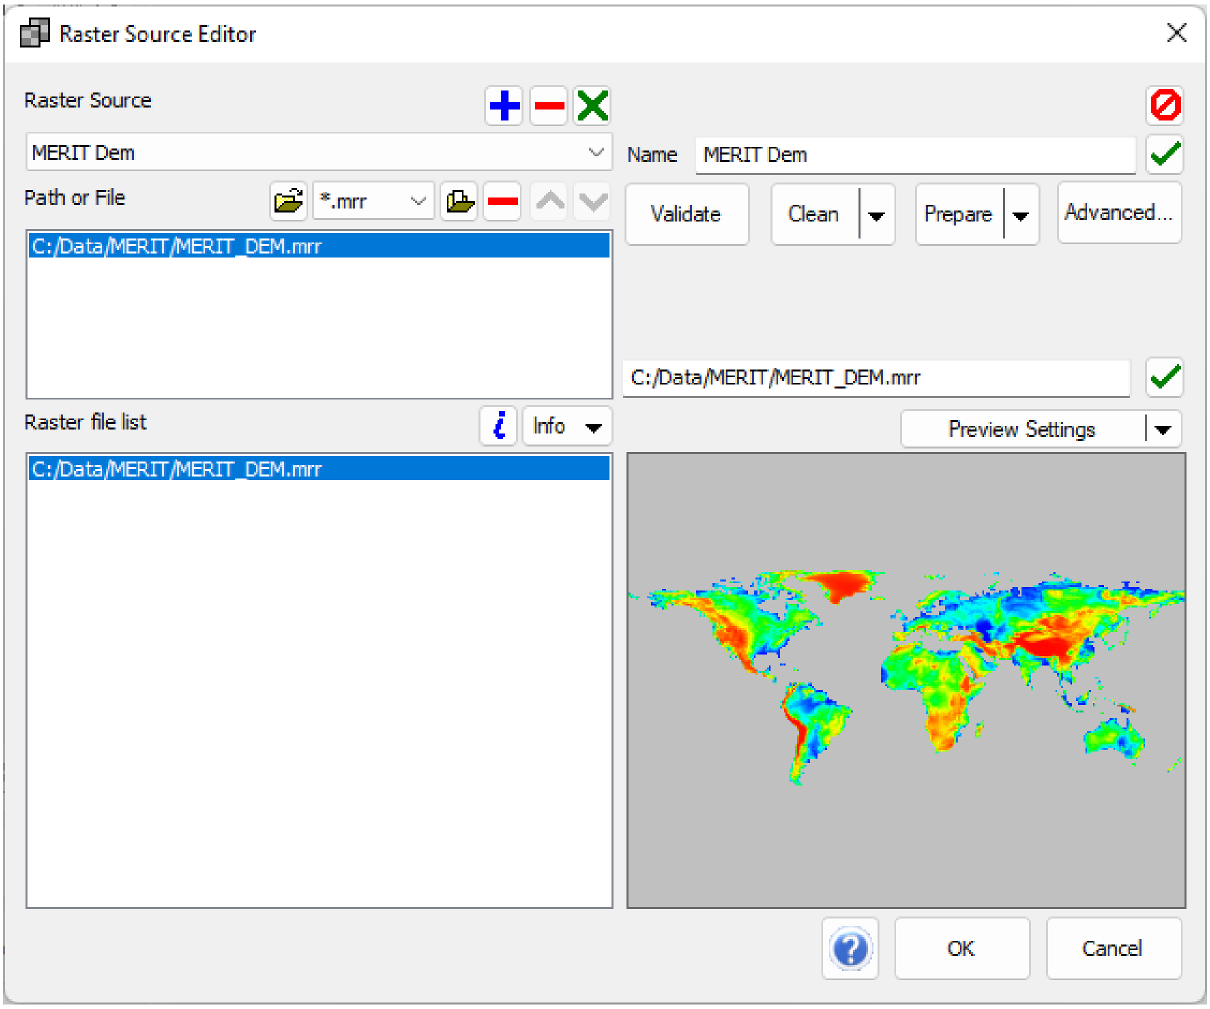

The Raster Source Editor Dialog

Saved raster source definitions are stored on your C: drive in a binary file. You can only edit this file using the Raster Source Editor dialog. In addition to saved raster sources that you have created, you may have temporary raster sources. These are created in two ways – firstly, when you browse for a raster in a component property page or when you create an algorithm for a raster, and secondly when you open an MRD algorithm file. When you open an MRD the raster source definitions it contains will be compared against your saved definitions. If no match can be found, then a temporary raster source is created. When you close the algorithm, it will be destroyed.

Defining a raster source

The Raster Source Editor dialog shows you all the raster source definitions, by unique name, in the system (for the current algorithm). Temporary raster sources are shown above the dotted line and permanent (saved) rasters sources are below it. Select any raster source in the drop-list to edit or view it.

Buttons are provided to add a new raster source, delete the selected raster source, copy the selected raster source and to delete all raster sources. No changes are made until you hit the “OK” button, so if you want to reverse your edits, just hit “Cancel”.

When you add a new raster source, you must immediately give it a unique name. The dialog will indicate whether the name you have assigned is unique. You can then specify what raster, or rasters, will be used in the raster source. Note that rasters are remembered by file path and name – no copy of the rasters is made.

A raster source will commonly contain a single raster. Less frequently, you may define raster sources that contain multiple raster files. When you connect a multi-file raster source to a layer, all the rasters will be rendered using the layer properties. In that sense, it is an alternative to creating a layer for every raster and is much more convenient.

All rasters in a raster source must all share the same coordinate system and raster structure (meaning that if you pick Field 0, Band 3 – there is such a field and band in every raster, and it contains the same kind of data). It is likely that they will all have the same cell size. As you can connect a different raster source to every component in a layer, there may be a different number of rasters in each of the components. This is supported in three different scenarios – Many : One, One : Many, and Many : Many (see Understanding Algorithms).

To add rasters to a raster source, you need to add items to the “Path or File” list. You can drag-and-drop a raster file from Windows File Explorer onto the dialog. Alternatively, hit the “Add a new file” button to browse for a file (or multiple files). As you add new items to this list, you will see new rasters appear in the “Raster file list”.

Instead of adding raster by full path and name, you can add all the rasters of a particular type in a folder by using a path with a wildcard. To do this, drop down the “Selected file extension” list and choose the raster extension type you want to select. Then hit “Add a new folder” and browse to the folder you want to target. All the rasters of that type in the folder will be added to the raster file list.

It is important to understand that the raster source definition only consists of the paths and files in the “Path or File” list. When you use an extension wildcard only one path string is written into the algorithm, even though it may result in rendering 500 geotiff files (for example). Also, the number of files that will be rendered depends on how many are there at the time – so the algorithm can change by adding and removing rasters from the target directory.

You can delete items from the :Path or File” list or reorder items. When you do so, the raster file list will be updated to reflect your changes. The order that rasters appear in this list is important if the rasters overlap. Rasters are rendered in the order in which they appear in the raster file list using a “painters” algorithm, meaning that rasters that are rendered first may subsequently be overprinted that rasters that are rendered later. It follows that you should order the rasters from lowest to highest resolution, or from lowest quality to highest quality.

When you add a new raster source in the Raster Source Editor and hit OK to save the modifications and close the editor, the raster source will automatically be opened into a default algorithm and rendered.

Rasters in a ZIP archive

Using a raster source you can target rasters, in a restricted list of supported formats, that are in a ZIP archive file. Firstly, add the ZIP archive file to the “Path or File” list. There are three possibilities:

1. The ZIP file contains a single raster in the root folder.

The raster engine may be able to open the ZIP file directly and read the raster. If so, you ought to be able to see a preview of the raster in the ZIP archive.

2. The ZIP file contains multiple rasters in the root folder.

   By default, the raster engine will only open the first file it discovers in the archive. You can direct it to open all rasters by manually editing the path. For example, if it contains Arc ASCII rasters then you can add “/*.asc” to the path. For example: C:/MyData/MyProject/ZIPFile.zip/*.asc. Hit the “Apply and save changes” button, which should turn to a green tick.

   The rasters in the ZIP archive should now be listed in the raster file list, and you ought to be able to preview them.

3. The ZIP file contains one or more rasters in sub-folders in the archive.

   The same trick will find rasters in sub-folders in the archive. For example, if it contains an Arc ASCII raster then you can add “/*.asc” to the path. For example: C:/MyData/MyProject/ZIPFile.zip/*.asc. Hit the “Apply and save changes” button, which should turn to a green tick. You should be able to see rasters, even if located in sub-folders in the archive.

   You can also target a specific raster in an archive by naming it in full. For example:
   C:/ MyData/MyProject/ZIPFile.zip/srtm_35_01/srtm_35_01.asc

   This will load the srtm_35_01 Arc ASCII raster from the srtm_35_01sub-folder of the ZIPFile.zip archive file.

The editor dialog includes a preview map window. To see a preview, click on a raster in the raster file list. A default rendering of that raster will then be displayed. You can zoom and pan about this preview in the normal way. To customise the field, band, and time range you can drop down the “Preview Settings” menu button for options.

Information Report

Info Controls

You can obtain information about a raster by selecting it in the raster file list and then hitting the information report button. You can customise the detail in this report by dropping down the Info menu button and selecting the required detail level. When you hit the information button a Raster Processor dialog is opened, and the report is dumped into the messages section of the dialog. You can resize this dialog and copy and paste text out of it.

The available detail levels are Brief, Standard, Detailed, Detailed + Statistics, and Complete (to text file). If you select “Detailed + Statistics” then you get summary statistical information for every band, but this does not include the histogram data. If you select “Complete” then you will also get the histogram data for each band. As this can create a lengthy report, you will be prompted to browse to an output text file into which the report is written.

Validate, Clean, and Prepare Operations

Before you can use a raster in ProRaster, or in any modern application, you need to do some preparation work. It is highly recommended that you ensure your raster has an overview pyramid cache and has high quality statistics stored and ready to be used. If it needs to, ProRaster will generate these caches on-the-fly when the raster is opened. However, that does not mean that it generates them appropriately. You can exercise control over the generation of these assets using the Validate, Clean, and Prepare operations.

The Validate operation will examine each raster in the currently selected raster source and report whether it has the required overview and statistics cache files. When you hit the “Validate” button a Raster Processor dialog opens, and the operation is executed. If any information is missing, you will see warnings and error messages in the operation report.

The Clean operation will delete any existing cache and ancillary files. Drop down the Cache menu button to control what files will be removed. By default, it removes data cache files which contain the overview pyramid and GHX files which contains statistics. When you hit the “Clean” button a Raster Processor dialog opens, and the operation is executed. Before any files are permanently deleted, it will prompt you for confirmation.

The Prepare operation will generate overview pyramid cache files and compute statistics, as required. Drop down the Prepare menu button to control cache file generation. You can choose Speed (which is the default) or Quality. This setting will modify what kind of compression is used in the cache. For data, if you choose Quality, ProRaster will use lossless compression that is slower but results in a smaller file size. For imagery, choosing Quality will result in lossless compression, whereas choosing Speed may result in lossy compression. When you hit the “Prepare” button a Raster Processor dialog opens, and the operation is executed.

For all these operations, you need to ensure that the rasters that are going to be processed are not open elsewhere. This includes in this instance of ProRaster, or any other instance of ProRaster you have running! If you have the raster displayed in an algorithm, then the processing operations may not be able to modify the files. Also ensure they are not open in other applications like MapInfo Pro or QGIS.

Advanced Driver Preferences

Hit the Advanced button to open the Driver Preferences dialog. This presents a wide range of options that you can override.

Driver Preferences are supplied to the raster engine whenever it opens a raster. They will be supplied when the engine opens any of the rasters in the raster source definition. The driver preferences may modify the way the engine opens the raster, or the way a driver in the engine opens a raster. This can result in profound changes. In some cases, those changes impact the content of the overview pyramid cache and statistics. Therefore, if you make changes to driver preferences, you will probably then want to run the Clean and Prepare operations to update those cache files.

For example, Landsat data files used to be supplied as geotiff files with cells outside the data scene given a value of zero. However, it was not a declared invalid cell value. With driver preferences you can specify that zero is the invalid cell value. If you do not, all those zero values infect the overview pyramid and get blended with actual data values, producing polluted data up the pyramid. Landsat scenes on AWS initially had this problem.

The general structure of driver preferences is that each property has a default value, and if you want to change the value you indicate that you want to override the property and can then edit the value. Some properties target the engine, some apply to all drivers, and some target specific drivers. Some of the useful properties are listed below:

Engine, Driver ID: In general, you do not specify which driver will open a raster. The engine will figure it out. However, some rasters can be opened by a Native driver and by the GDAL driver. In some cases, this may result in a different treatment of the data. You can specify a driver ID here to force the engine to use that driver. For a list of all driver ID's, see this web page. For example, to make GDAL load a raster, use the MI_GDAL_GENERIC code.

Coordinate offset: You can translate a raster in X and Y by specifying a coordinate offset. For example, if you have a global raster in a geodetic coordinate system you can offset it by 180 degrees in X to duplicate it to the east and provide a world-wrap style rendering.

All Drivers, Set coordinate system: This provides another mechanism to set the coordinate system of a raster if the engine cannot determine it or incorrectly determines it. Enter the coordinate system as a MIF string, which you can find in the Projection Selector dialog.

All Drivers, Set cell validity from alpha: If the raster has an alpha band, use this band to set cell validity.

All Drivers, Set field type: Some drivers may not be able to properly determine the field type of a raster. Use this to set it manually.

Geotiff driver, Specify the null value: Set a specified cell value to null.

Image driver (JPEG, PNG), Treat Image Palette as Image: Load ImagePalette raster fields as Image field type.

Image driver (JPEG, PNG), Use overview cache files: Allow overview pyramid cache files to be generated and permanently stored for images. Normally a temporary cache file is used.

Using the Color Table Editor dialog

Color tables are used exclusively in the LUT Color layer as the source of color for the Color component. See a demo on YouTube here. Opacity components use a fixed table, as do the Red, Green, and Blue components in an RGB Color layer. To select a color table for use in a LUT Color layer, you must first make that table available to ProRaster using the Color Table Editor dialog.

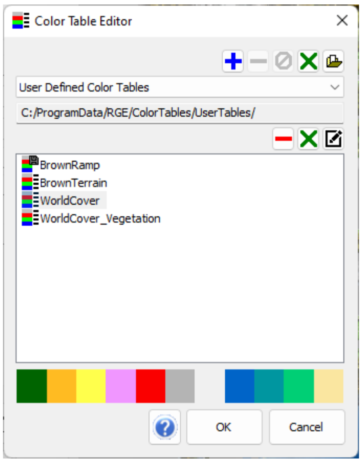

Color Table Editor Dialog

The Color Table Editor manages the color table resources that ProRaster can access. It provides a preview of all color tables under management, allows you to connect new tables that exist in external directories, create new tables, copy existing tables, and provides tools to edit color tables, maps, and legends.

There are four different kinds of color tables – ramps, tables, maps, and legends. A color table ramp is a programmatically defined color table. The first and last color are defined, and all the intervening colors are created by interpolation in either RGB or HSL space. A color table is a file on disk and, even though it can be in a variety of different formats, ProRaster will only write color tables in one format – the XLUT format. The colors in the table may all be defined individually, or they may be created by interpolation when loaded. A color table map is a color table on file that includes a data map, generally listing a data value associated with every color value. Multiple formats are supported for reading, but maps are only ever written in XLUT format. A color table legend is a list of data values, ranges, or strings along with rendering information including color. ProRaster supports the LEG format for reading and writing legends.

Category list

The editor lists all the color tables managed by the system by category. Select a category, and a list of all the color tables therein will be displayed. The icon for each table reflects whether it is a ramp, on disk, contains a data map, or a legend. Select a color table in the list and a preview will be rendered at the base of the dialog. The category list contains the following categories:

Local Algorithm Color Tables: These are color tables that have been loaded from the current MRD algorithm that have not been matched with any known color table under management.

Recently Selected Color Tables: This is a list of color tables that you have recently used and duplicates color tables from other categories.

System Color Ramps: A color ramp is a simple table that interpolates between two defined colors. These are programmatically created and are not stored in any file.

System Color Tables: The color tables that were shipped with ProRaster and are stored in the C:\ProgramData\RGE\ColorTables\SystemTables directory.

System Color Maps: The color table maps that were shipped with ProRaster and are stored in the C:\ProgramData\RGE\ColorTables\SystemMaps directory.

User Defined Color Tables: Color tables, color table maps, and color table legends that have been created by you and saved into the default directory at C:\ProgramData\RGE\ColorTables\UserTables.

External Folder […]: An external directory that contains color table files that you have linked to ProRaster. You can link multiple external folders.

In the editor, a button is provided to "Add a new external directory". If you have color tables files in a directory that you have created, or you have color tables in files in a directory installed by another software package, then you can link this to ProRaster and, thereafter, any color tables that ProRaster can read will be available for use in your algorithms. A second button is provided to "Remove the selected external directory" so that ProRaster forgets about its contents.

The "Clear recently accessed list" button simply erases the list of recently selected color tables. This does not actually remove any color tables.

The "Delete the selected color table" button allows you to delete files from external folders or from the user defined color tables directory. Care should be taken. The files will not be deleted until you hit "OK".

Creating color tables

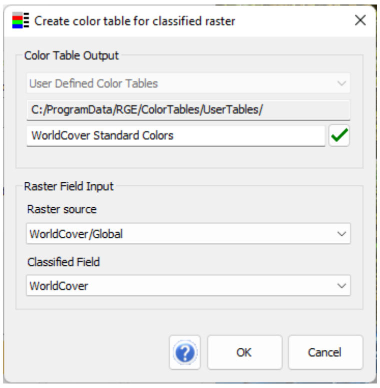

New Classified Color Table

If you have the "user defined color tables" category selected, or any external folder category selected, then you can hit the button to "Create a new color table for a classified raster". This option is designed to generate color table legends for classified and image palette fields for a selected raster source. When rendering classified data, you can use a LUT Color layer and color transform and modulate by the legend produced by this option. Furthermore, you can edit the legend and remove items to prevent certain classes being rendered.

Hit the "Create a new color table for a classified raster" button. A dialog is presented showing the output file path and name. As always, you must specify a unique name for the color table. Then choose a raster source from the list and choose a classified or image palette field from that raster source. Hit OK and the new color table legend, using discrete value matching, will be created.

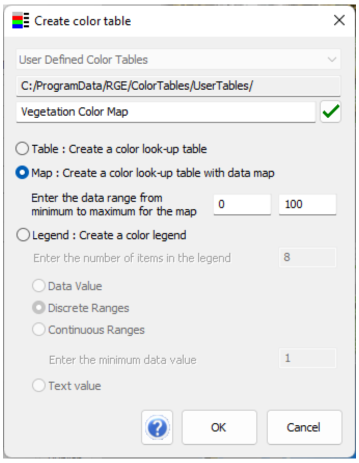

Create Color Table

Two more creation options are provided – to "Create a new color table in the selected folder" and to "Create a new color table from the selected table". The first option will create a new table in the user defined or external directory. The second option will create a new table in the user defined directory by copying at least some elements of another table. Both options will then open the "Create Color Table" dialog.

The dialog will show you where the color table will be written. Your first task is to choose a unique name for the table which will double as a filename. The dialog will indicate to you whether the name you have selected is unique.

Secondly, you need to specify what kind of table it will be – a Table, Map or Legend. If you choose Legend, then you need to specify the number of items in the legend and whether the legend will use data values, data ranges, continuous data ranges, or data text strings.

If you are copying a color table, then the colors and data in that table will be used to populate the new table as applicable, regardless of what type the new color table will be. Hit OK to create the table and you will be returned to the main dialog. Your new color table will be visible in the list of color tables. You can double click on the table in the list to edit it or select the table and hit the “Edit” button.

Editing color tables

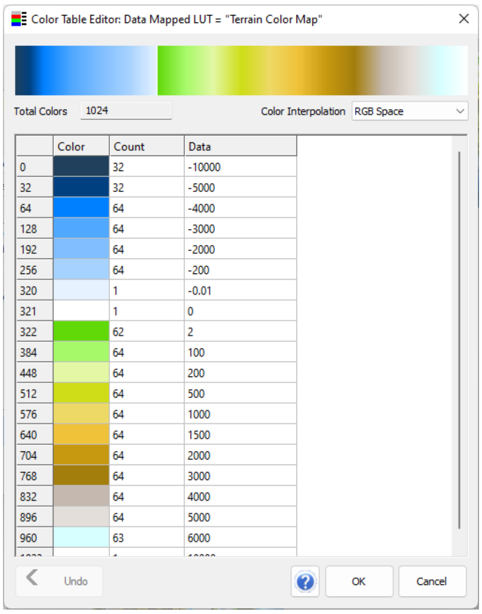

Color Table Editor

The editor dialog contains a preview of the color table at the top and a spreadsheet of the data. An “Undo” button is provided to undo changes. The total number of colors in the table is reported under the color preview. For a standard color table, you can also change the interpolation method between RGB and HSL color space. In the spreadsheet, you can duplicate the current row by navigating to the rightmost cell of that row and hitting the return key (the row will be duplicated below). You can delete the previous row by navigating to the leftmost cell of a row and hitting the delete key (the row above will be deleted). You can also delete a row by clicking on the row header cell and hitting the Delete key.

You can build a color table or map by defining every color as a row in the sheet, or you can define a set of key colors and allow the colors in between to be generated by interpolation.

To specify a color, click on the color cell in the sheet. Select one of the pre-defined colors or hit the “Custom” button to open a more advanced color selection dialog. Note that you only need to specify the RGB components of the color. The alpha component is not used. Whenever you make any change to the table, the preview graphic is updated to reflect the change.

The row index tells you zero based color index of the color in that row. The “Count” column is the number of colors defined by the row. The colors for the row will be found by interpolation from the row color and the next row color. The count value should be one or more, and the count value for the last row ought to be one.

A color table map has an additional “Data” column which allows you to associate data values with colors in the table. The data values must be monotonically increasing, and you should always define a data value for the top and bottom colors. It is not necessary to define a data value for each key color in the table, but it is likely that you will. To use a table like this in an algorithm, select it as both the data transform and the color table.

A color table legend is a more traditional arrangement where each row in the spreadsheet contains a color and a data value or range of data values, and the table consists of only so many colors as there are rows in the table. No color interpolation is used.

Legends that contain a data value or text string per row are generally used to colorise integer data, such as is commonly found in classified rasters. Data values in the raster that do not match any value in the legend may not be rendered.

Legends that contain a data value range can be used to colorise continuous data. Ranges can be discrete or continuous. In the first case, any data that is >= the minimum and <= the maximum is considered in the range. In the second case, any data that is >= the minimum and < the maximum is considered in the range. The “Continuous range” legend is designed for ranges that touch each other but do not overlap. In either case, any data that is not inside a range will not be rendered.

Using the Data Conditioning Filter dialog

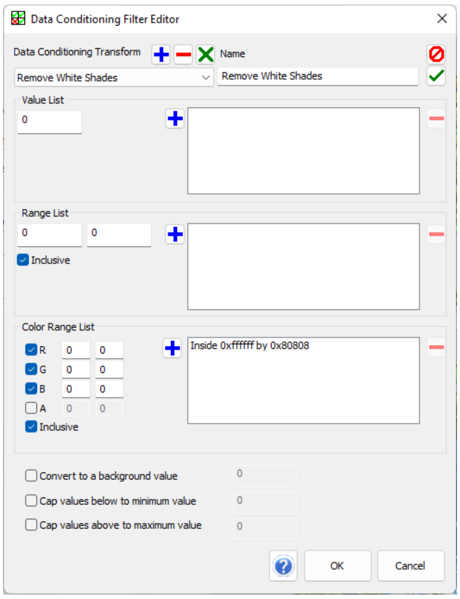

Data Conditioning Filter Editor Dialog

In every component, you can select a Data Conditioning Filter to apply to the raster data (specified by your selection of the raster source, field, band, and time range). This filter is applied to the data prior to the data – color transform. See a demo on YouTube here.

The primary purpose of a Data Conditioning Filter is to remove invalid or unwanted data values from the data stream prior to rendering.

The editor dialog presents you with a list of known filters. First in the list, above the dotted line, will be any filters that have been loaded from the current MRD algorithm. Below the dotted line will be any filters that you have defined and saved via the editor dialog. Select a filter to edit it.

Buttons are provided to add a new filter, delete the current filter, or clone the current filter. You would use the clone operation to save a copy of a filter that you have loaded from an MRD file. You can also delete all filters. If you do this accidentally, make sure you hit Cancel to exit the editor dialog. No changes are saved until you exit the dialog by hitting the OK button.

When you add a new filter, you must immediately give it a unique name, because filters are referred to by name in the algorithm XML. As you type the name, the dialog will indicate whether it is unique and valid or non-unique.

These are the operations that are allowed, and the order in which they are applied:

1. Reject values
   You can define one or more values that will be rejected. If a data value matches one of the listed values to reject, it will be registered as invalid.
2. Reject ranges
   You can define one or more ranges of values that will be rejected. If a data value is >= the minimum value in a range and <= the maximum value of that range, it will be registered as invalid. This can also be reversed so that values are rejected if outside of the range rather than inside the range.
3. Reject color ranges
   For 1, 2, 3, and 4-byte color data values, like RGB or RGBA, you can define color ranges that will be rejected. In this case, you define the color you want to reject and an envelope about this color that will also result in rejection. For example, if the color component value is 64 and the envelope is 8, then values from 56 to 72 will be rejected. The envelop will be clipped at zero and 255. This can also be reversed so that values are rejected if outside of the range rather than inside the range.
4. Transform
   Apply a transform to valid data of the form Y = A.X + B. For example, if you have a terrain raster you can use this transform on the Z band to convert between feet and metres or apply a sea level change offset.
5. Convert to background
   If a data value is invalid at this point in the processing, it can now be converted to a “background value” and made valid.
6. Cap values below minimum
   If a data value is valid and below a minimum cap, it is set to the minimum cap value.
7. Cap values above maximum

If a data value is valid and above a maximum cap, it is set to the maximum cap value.

You can then define the six filter parameters as described above – values, value ranges, color ranges, background value and caps. For example, enter a new value and then hit the “Add” button to add it to the list of rejected values.

Once you have defined your data conditioning filters, you will be able to select them by name for every component in your rendering algorithm.

Using the Data Transform Editor dialog

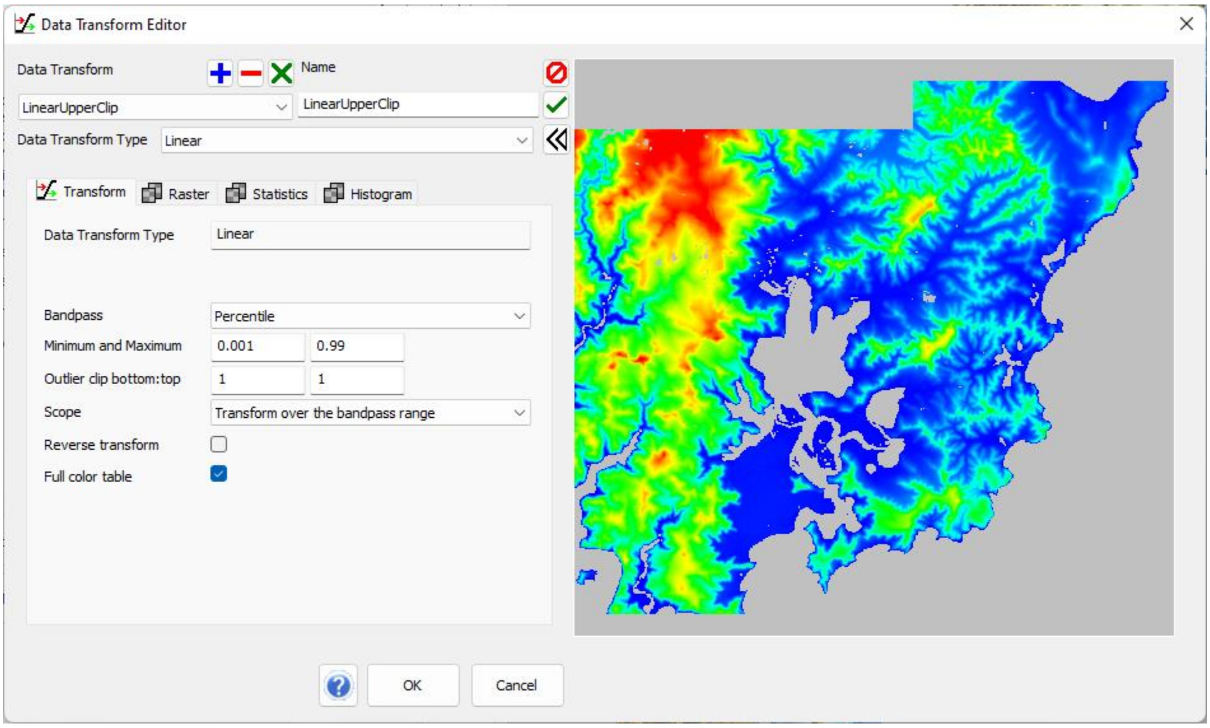

Data Transform Editor Dialog

A data transform is used to convert from raster data to a value between zero and one, which is then scaled to an integer index, then used to acquire a color from a look-up table. See a demo on YouTube here.

The data transform is a mathematical function. It may be expressed mathematically, like Linear, Logarithmic, or Sigmoid, or it may be defined by a table of coordinates (raster data, scaled index 0 - 1) which the software can use to estimate the scaled index for any input data value. Equal Area is an example of the latter, and the editor provides you with a set of tools to create and edit your own function tables.

The component property pages provide a rich set of data transforms that are easy to define and likely to be sufficient for the needs of most users. The Data Transform Editor dialog provides a more advanced capability with additional flexibility. In particular:

- For the standard transforms, there are more flexible methods of defining the bandpass clipping extents.
- The bandpass clipping extents can be decoupled from the transform.
- For logarithmic transforms you can edit both the log base and linear transition value.
- Some transforms can be reversed.
- You can use the full color table for range transforms.
- A rich set of table-based transforms are provided.

In addition, the editor dialog provides a preview of the transform on a raster of your choice. For this raster you can see the summary statistics presented in a property page as well as the distribution histogram presented either as a spreadsheet or graphical view. If statistics have not been computed for the raster, you can compute them by hitting the “Compute Statistics…” button on the “Raster” property page.

The data transform editor dialog presents you with a list of known transforms. First in the list, above the dotted line, will be any transforms that have been loaded from the current MRD algorithm. Below the dotted line will be any transforms that you have defined and saved via the editor dialog. Select a transform to edit it.

Buttons are provided to add a new transform, delete the current transform, or clone the current transform. You would use the clone operation to save a copy of a transform that you have loaded from an MRD file. You can also delete all transforms. If you do this accidentally, make sure you hit Cancel to exit the editor dialog. No changes are saved until you exit the dialog by hitting the OK button.

When you add a new transform, you must immediately give it a unique name, because transforms are referred to by name in the algorithm XML. As you type the name, the dialog will indicate whether it is unique and valid or non-unique.

Transform types

You can then set the data transform type, by selecting from the drop-list provided. Once you select a transform type a property sheet will be displayed allowing you to edit the properties of the transform.

Pass Scaled: Your data values are assumed to lie between 0 and 1. The data transform will simply convert this value to a color index.

Pass Index: Your data values are assumed to be an unsigned zero-based integer matching the number of colors in the color table. The data transform uses the data value to look up the color.

Pass Value: Your data value is a color. No modification to the data is made.

Rotate Index: Supply the number of colors through which to rotate. Your data values are consecutive integers. The data transform is the modulus of the data value and the number of colors.

Color bits linear: Supply the number of bits (N) of color, creating a linear transform from 0 to 2^N-1. For example, imagery is typically in 8-bit color.

Linear: A linear transform between two values. These values are determined by the “Bandpass” settings, which can be “None”, “Data Values”, “Percentage”, or “Percentile”. If you select “None” then the bandpass values are acquired from statistics (minimum and maximum cell values). Select “Data Values” to enter the minimum and maximum bandpass values manually. Otherwise, you can express the values as a percentage (0 – 1) or percentile (0 – 1). Note that you can reverse the transform by reversing the minimum and maximum bandpass values.

If you choose None, Percentage, or Percentile, then you can also specify the number of histogram bins to clip at the bottom and top of the range. This can be a useful way to clip out the outliers at the edges of the dataset without having to work out what those outliers are. Just clip off one or two bins at the bottom and top and improve the data – color transform.

You can reverse a linear transform easily enough by swapping the minimum and maximum bandpass values. For this transform and some others, you can also toggle this with the “Reverse transform” button. In many cases it is more convenient to reverse the color table than reverse the transform.

Logarithmic: A logarithmic transform between two values, see Linear. Here, you can specify the base of the logarithm (the default is base 10). The logarithmic transform is undefined about zero, so you can define a range about zero where the transform is linear. Enter this as the “Minimum”.

Sigmoid: A sigmoid transform between two values, see Linear. Specify the sigmoid K factor between -1 and +1.

Equal Area: An equal area (or histogram equalised) transform between two values, see Linear. You can scale the severity of the transform up and down – from linear to equal area – by adjusting the Percentage Scale between 0 and 1.

Stable Equal Area: An equal area (or histogram equalised) transform between two values, see Linear. You can scale the severity of the transform up and down – from linear to equal area – by adjusting the Percentage Scale between 0 and 1.

Equal Area uses an "Equal Width" histogram which results in faster rendering. Stable Equal Area uses a "Equal Frequency" histogram which can produce higher quality transforms, particularly in the presence of outliers or highly skewed distributions.

Ranges: Spread about the mean: Enter the number of ranges (N). The transform will generate (N/2) ranges spread evenly from the mean to the maximum bandpass value, and (N/2) ranges spread evenly from the minimum bandpass value to the mean.

If you specify N ranges, then you are usually indicating that you want to display N colors, and you will probably supply a color table with N colors in it. You can supply a color table with more than N colors. In that scenario, you can check the button to use the "Full color table" and the transform will interpolate between your ranges to use all the colors.

Ranges: Spread about the median: As previously, spread about the median.

Ranges: Spread about the mode: As previously, spread about the mode.

Ranges: Deviations about the mean: Enter the number of ranges (N). The transform will generate (N/2) ranges above and below the mean. The width of each range is equal to the one standard deviation.

Ranges: Jenks Natural Breaks: Enter the number of ranges. The Jenks natural breaks algorithm is used to compute the best position of the breaks and the width of each range.

Table: The table options create a table of coordinates on the data transform where the x-coordinate is the data value, and it maps to the y-coordinate which is a scaled index between 0 – 1. However, ProRaster provides you with multiple ways to define the data values, see below.

Table: Index over range: In this table the data values are not entered manually. Instead, you define the minimum and maximum value by using the bandpass parameter tools (see Linear, above), and the data values are assigned linearly across this range. All you edit is the scaled index for each entry in the table.

This option can be difficult to work with because the data range is split into N equal width ranges, where N is equal to the number of rows in the table. The index for each row is applied at the minimum value of each range.

When you choose one of these table options, you will see a "Table" property page appear in the property sheet. On this page is a spreadsheet which you can edit. At the base of the sheet are controls to set the number of rows and to reset the table to default values for those rows. An undo mechanism is provided so you can undo changes.

You can add and delete rows in the spreadsheet by editing. Add a row by moving to the rightmost column and hitting the return key. The row will be duplicated below the selected row. Delete a row by moving to the leftmost column and hitting the delete key. This will delete the row above the selected row.

Note that the scaled index values do not have to be monotonically increasing.

Table: Index for value: In this table you edit both the data value and the scaled index, so you have complete manual control over the table.

The data values need to be monotonically increasing, but the scaled index values do not.

Table: Index for percentage: In this table a percentage, between 0 and 1, is used to map to data values. This percentage is stretched between the bandpass limits.

The percentage values need to be monotonically increasing, but the scaled index values do not.

Table: Index for percentile: : In this table a percentile, between 0 and 1, is used to map to data values. This percentile is stretched between the bandpass limits. Recall that a percentile is a percentage of the number of samples, equivalent to a statistical quantile for a data distribution.

The percentile values need to be monotonically increasing, but the scaled index values do not.

Table: Match value: This table, and all the following tables, are referred to as legend transforms. In these tables the scaled index does not need to be declared and is not editable. You will almost certainly use these tables with a color table that has a matching number of colors (the same number as the rows in the table) and it follows that the color for each row in the table will be the same as the zero-based row index. In fact, you may find it more convenient to create a *legend color table map* in the Color Table Editor dialog which creates a transform in exactly the same way, but also allows you to specify the color for each row.

This table declares data values which map to individual colors. Data values that do not match the values listed in the table will not be assigned a color appropriately. You ought to use a table like this for integer data, like classification codes for example, rather than continuously variable data like terrain.

The data values need to be monotonically increasing.

Table: Match value range: In this table you declare a set of data ranges, from minimum to maximum for each range. Data values that fall into a range are assigned the color for that range.

Within a range, the minimum value ought to be <= the maximum. However, ranges do not have to be monotonically increasing and they can overlap. Where a data value can be assigned to multiple ranges, it will be assigned to the first in the table.

Table: Match continuous value range: In this table we make sure that the ranges are monotonically increasing and do not overlap. A data value must be >= the minimum value and < the maximum value t be assigned to a range.

Table: Match value string: In this table we supply text strings as the data values. You would use these to transform against a string field in a raster. This might occur in a classified raster, for example.

Once you have defined your data transform, you will be able to select it by name in the component “Color” property page under the “Stored Transforms” category.

Transform preview

You can apply your transform to a raster in real-time as you design it. The right-hand side of the dialog is a preview map window in which you can pan and zoom in the usual way. A button is provided to open and close the preview window (it cannot be undocked). Select the raster to view in the preview by opening the “Raster” property page and selecting a raster source. You can select the field, band, and time range using the "Preview Settings" menu button.

The preview raster will be displayed using a LUT Color layer algorithm. The current data transform is applied to convert the raster data to scaled value, and then a standard "pseudocolor" color table is used to provide color. If you open the data transform editor when you are working on an algorithm, the raster you were using in the algorithm will be selected as the default.

Preview raster statistics

You can access statistics for the raster source you have selected for the preview. Select the "Statistics" property page to see summary statistics of the raster. This includes the following information:

Sample Count: Total number of samples, Number of valid samples, Number of invalid samples, Number of invalid samples that were Not A Number, Number of invalid samples that are empty, Number of invalid samples that are null.

Summary: Minimum value, Maximum value, Mean, Variance, Standard deviation, and Signal to Noise ratio.

Distribution: Lower quartile (25%), Upper quartile (75%), Interquartile Range, Median, and Mode.

Spatial: Minimum, Maximum, Mean, Variance and Standard deviation. These are not standard statistical measures – they measure the difference between a cell value and the adjacent cell value.

Select the "Histogram" property page to see distribution statistics. These can be presented as a spreadsheet or as a graph. You can apply a logarithmic stretch to the vertical axis of the graph.

The spreadsheet contains the distribution histogram data. Each row of the sheet corresponds to one histogram bin. In ProRaster, histograms are either "Equal width" where all the bins have the same width and there are a large number of bins, or "Equal frequency" where bins have a variable width, tend to have the same number of samples in each bin, and there are fewer bins.

Each bin covers a defined data range. This is reflected in the Bottom and Top columns in the sheet. Even so, the data samples within that range may only cover a subset of the range. This is reflected in the Minimum and Maximum columns in the sheet. The number of samples recorded in each bin is shown in the Count column. The Cumulative Count column records the percentage of samples up to and including that bin (from 0 – 1).

You can compute statistics for the selected preview raster field, band, and time range. On the "Raster" property page, hit the "Compute Statistics…" button to execute the operation. A progress dialog will be displayed, and you can cancel the operation by hitting the "Cancel" button on the progress dialog. If the "Compute Statistics…" button is not enabled, then the statistics for the preview raster are already computed and ought to be visible in the Statistics and Histogram property pages.

Statistics for the preview raster are computed from the base level resolution data. This data will be permanently cached in the GHX file and in the STTX file associated with the raster. The next time you access the raster, the statistics will be automatically available. Statistics can be computed for each event in the raster, but even though you can select any range of events for the preview raster, the statistics are always computed for the roll-up of events from the first event in the field to the selected second event.

Using the Projection Selector dialog

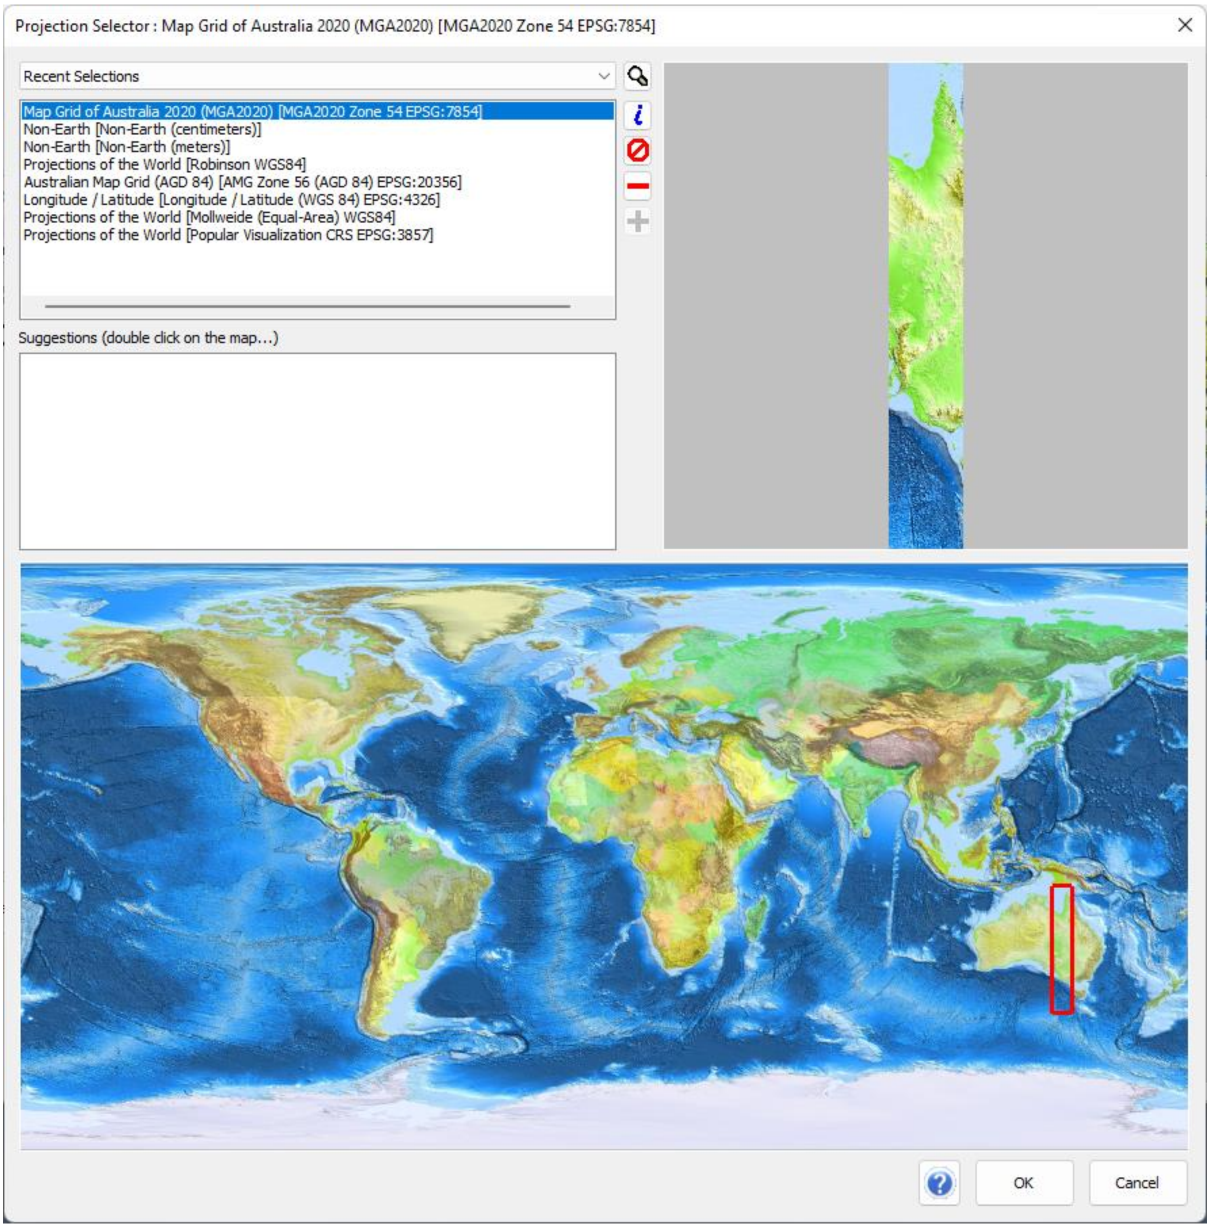

Projection Selection Dialog

The Projection Selector dialog, elsewhere known as the Projection Explorer, allows you to choose a coordinate system from a large range of predefined and user defined coordinate systems.

With this dialog you can:

- Select a coordinate system or select “no coordinate system”.
- Define a new coordinate system.

You make a coordinate system selection by selecting the appropriate category in the category drop-list and then the appropriate member in the member list located immediately below it. The selected coordinate system will be displayed in the title bar of the dialog. If you wish to choose “no coordinate system”, a button is provided to clear the category and member lists.

As an alternative to navigating these lists, the dialog can help you find and select a coordinate system in a few ways. Firstly, it records all the coordinate systems you have selected in the past and displays them in the

“Recent Selections” category. Secondly, you can search using a text string which will list all coordinate systems that may approximately match that string. Finally, you can search by double-clicking on the world map to see a list of coordinate systems which may be suitable choices for that location.

In the MapInfo Pro ecosystem, coordinate systems are defined in a text file called MAPINFOW.PRJ which you will find in the Professional directory of your MapInfo Pro installation. In this file, coordinate systems are separated into more than 150 categories (or groupings), and in each category there will be one or more members, each of which represents a coordinate system.

Each member will have a descriptive string and a set of numeric (and sometimes text) parameters which define all the coordinate system elements such as the spheroid, datum, and projection. To understand the parameters and learn how to define a coordinate system using parameters, you will need to refer to the MapInfo Pro User Guide.

This long-standing mechanism for using coordinate systems and projections in MapInfo Pro is at odds with more recent systems which are more likely to use EPSG codes or a well-known-text representation of the coordinate system data. Support for EPSG codes is provided by embedding the code in the member description text preceded by a “\p” character sequence. To learn more about EPSG codes and find coordinate system definitions by code, you can explore these websites - https://epsg.org/home.html, https://epsg.io/, https://spatialreference.org/.

The main dialog

The dialog is split into three zones – the “control” zone in the top left, the “world map” zone along the bottom and the “projected map” zone in the top right.

The control zone contains the category and member lists which define the currently selected coordinate system. This will be reflected in the dialog title. You can get detailed information about the coordinate system by clicking the information button located to the immediate right of the member list. You can search for coordinate systems by clicking the search button located to the right of the category list. The clear button allows you to select “no coordinate system”. The category list contains two special categories – the “Recent Selections” category, located at the very top of the list, and the “User defined coordinate systems” category, located at the very bottom of the list.

Every time you use the dialog and hit “Ok” to exit, your selected coordinate system is written into the recent selections list. The next time you use the dialog, your most recent selection will be front and centre. This data is permanently cached to disk and is shared across all software developed by RGE. If you wish, you can delete all members from this list by hitting the delete button located to the right of the member list.

If you select the user defined coordinate system category, then the “add” button to the right of the member list will be enabled, and you will be able to add a new coordinate system definition. These are permanently cached to disk and shared across all software developed by RGE. If you wish, you can delete the selected member from this list by hitting the delete button located to the right of the member list. You can also edit an existing member by double clicking on it in the list.

Below the member list is the suggestions list, which will be populated with candidate coordinate systems when you double click on a location in the world map. You can click on any suggestion to select that coordinate system.

The map windows

The world map is an interactive map control showing the topography, bathymetry and countries of the world using a geographic coordinate system. The projected map is an interactive map control showing world topography and bathymetry data, reprojected into the selected coordinate system.

In either map, use the mouse wheel to zoom in and out about the mouse cursor location. To pan, click and hold the left mouse button and drag, then release the button to end panning. Both maps will automatically fit to the data extents when you zoom out beyond the maximum extent.

The world map also displays countries as transparent colored polygons, superimposed over the topography. As you zoom in, the opacity of the countries layer is increased so that they become the dominant feature of the map. As you move the mouse over the map, the geographic coordinates (longitude and latitude) are displayed in a tooltip. The tooltip also reflects the country and the province of the country that the mouse cursor is hovering over. If you left double click the mouse in the world map, the suggestions list box will be populated with coordinate system candidates for that location. If you click on any coordinate system in the suggestions list, it will become the selected coordinate system in the category and member lists.

The projected map displays the world topography and bathymetry reprojected into the selected coordinate system. It does not display any countries and the tooltip only reflects the cursor location. To display data in the selected coordinate system usually requires clipping the world raster to the geographic bounds of the coordinate system. This feature is not available for all coordinate systems. When reprojection is not supported the projected map uses geographic coordinates and the topographic data is displayed in greyscale rather than color. When the geographic bounds are known they are rendered in the world map as a red box.

Searching

The search dialog provides a mechanism to search for a coordinate system by approximate string matching. It searches through the category and member string collection and, as you type a string, returns a list of candidate coordinate system definitions that match the search string. For example, you can search for a coordinate system by EPSG code by simply typing in the code as the search string. Alternatively, you might enter the name of a country to find any country specific coordinate systems. If you select a candidate coordinate system and hit “Ok” the main dialog will update to reflect that selection.

The information dialog provides detailed information on the selected coordinate system. This includes the category and member strings as well as the EPSG code (if known), the projection, datum, spheroid, and horizontal units. The coordinate system expressed as a MapInfo Interchange Format (MIF) string is also shown as well as the geographic (longitude and latitude) bounds. In some cases, the bounds will also be known in projected coordinates.

To add a user defined coordinate system, select the “User defined coordinate systems” category (which is the very last item on the categories list), then click the add button. A dialog is presented in which you can define the member string, the EPSG code and the parameters string.

The member string is purely descriptive and will be reflected in the member list once you add the new coordinate system. If you know the EPSG code, then enter this in the appropriate edit box rather than adding it to the member string, as the system will automatically add it to the member string for you. The most important information you must supply is the parameters string. This is difficult to construct manually, and to do so you will need to refer to the detailed information provided in the MapInfo Pro Users Guide. If you hit “Ok” the new coordinate system will be added to the user defined coordinate systems member list and main dialog will update to reflect that selection.

If you double click on a user defined coordinate system in the members list, the same dialog will open and allow you to edit the existing definition.

You can use this dialog to explore coordinate systems and projections, and to preview maps using the selected coordinate systems. You are limited to those spheroids, datums and projections that the MapInfo Pro system supports and for more information on this, consult the MapInfo Pro User Guide.

The world map is displayed using a geographic coordinate system based on the WGS84 spheroid. A presentation like this is useful for displaying the entire surface of the Earth, but results in increasing distortion as you move north and south of the equator. There are many other projections that address this issue in different ways, and the projected map will show you a preview of how the world will be rendered using the selected projection. A good place to start is the “Projections of the World” category where you can see the effect of projections like Mollweide, Sinusoidal or Popular Visualisation (used by Google Maps).

Many projections are specifically designed for a particular region of the Earth. Outside of this region, attempting to project coordinates may result in mathematical failure. For projections like this, the dialog needs to know the geographic bounds of the region. If the bounds are not known, it will not attempt to preview the projection. One of the most popular local projections is Universal Transverse Mercator (UTM). UTM can be applied to any region and to facilitate this there are many standard “zones” which define bounded regions. The category “Universal Transverse Mercator (WGS84)” is a good place to find UTM projections suitable for our region of interest.

Configuring Options

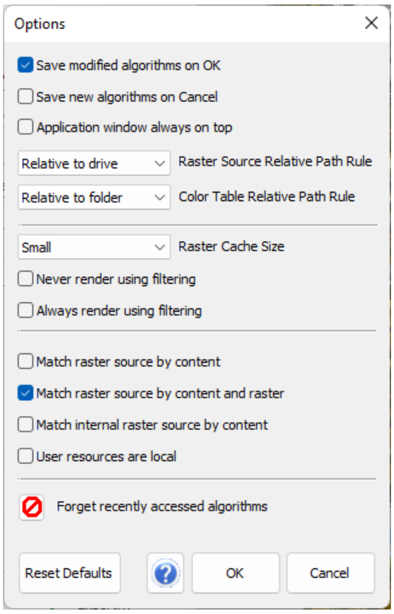

Options Dialog

Access the Options dialog from the system menu. From here you can modify some of the default behaviour of ProRaster. To return all options back to their standard defaults, hit the “Reset Defaults” button and hit “OK”.

Note that there are several options associated with matching resources loaded in an MRD to existing resources. The use of all these options is under review and subject to change.

Save modified algorithms on OK: When you hit “Save and Close” to exit ProRaster the default behaviour is to silently save any modified algorithms that have a known filename. You will also be prompted to save any new algorithms. Check this option to prevent algorithms being silently saved. If checked, you will be prompted to save modified algorithms.

Save new algorithms on Cancel: When you hit “Cancel” to exit ProRaster you will be prompted to save any modified algorithms. By default, new and unsaved algorithms will be silently discarded. Check this option to prevent discarding new algorithms. You will be prompted to save them.

Application window always on top: When you have undocked external map windows, check this option to keep ProRaster on top of these map windows. By default, the main dialog can be obscured by external map windows.

Only load and show the current algorithm: By default, every algorithm that is open in ProRaster will remain active and in memory. External map windows displaying algorithms will remain visible and active, regardless of which algorithm is currently selected. You can enable this option and force ProRaster to keep only the currently selected algorithm loaded. All other algorithms, and their map windows, will be closed. The primary reason to do this is to allow more memory to be allocated to the current algorithm. When you select a different algorithm the map windows for this algorithm will be restored.

Raster source relative path rule: An MRD file will contain references to files on disk. You can choose to write file names as absolute paths or relative paths. An absolute path contains the drive, folder and file name and extension. A relative file name contains a path that is relative to the location of the MRD file. If you choose an absolute path, then your algorithm will become invalid if the files it depends upon are moved. On the other hand, you can move your algorithm and it will remain valid. If you choose a relative path then, if you keep your algorithms and data packaged together, you can move them about, and the algorithms will remain valid.

This option refers to the rasters listed in raster source definitions. Choose “Absolute” to write absolute paths. Choose “Relative to drive” to write a relative path if the data is on the same drive as the MRD, or absolute otherwise. This is the default. Choose “Relative to folder” to write a relative path if the data is in the same folder as the MRD, or in a folder subordinate to it, or absolute otherwise.

Color table relative path rule: This option refers to color table files. These are generally located in system directories. You may also create unique color tables for projects and store them with the project data. . Choose “Absolute” to write absolute paths. Choose “Relative to drive” to write a relative path if the data is on the same drive as the MRD, or absolute otherwise. Choose “Relative to folder” to write a relative path if the data is in the same folder as the MRD, or in a folder subordinate to it, or absolute otherwise. This is the default.

Raster cache size: All raster data is stored in memory, at least temporarily, in a tile cache. It is faster to retrieve raster data from the cache than to reload it from disk. ProRaster keeps raster data in the cache as long as possible to maximise performance. Set the cache size to Small, Medium, Large, or "I want it all!".

- "Small" sets the cache size to the maximum of 512MB or 25% of free physical memory.

  The amount of free memory will change each time you run ProRaster. It depends on the amount of physical RAM installed and the amount of RAM being consumed by the operating system and other software.

- "Medium" sets the cache size to 1024MB or 50% of free physical memory.
- "Large" sets the cache size to 2048MB or 87.5% of free physical memory.
- "I want it all!" sets the cache size to 2048MB or 87.5% of physical memory.

  The amount of physical memory is the amount of physical RAM you have installed. As it is unlikely that 87.5% of this physical memory will be free (Windows will be hogging it), choosing this option will probably force the operating system to use virtual memory to support large raster caches.

Never render using filtering: Force the GPU rendering engine to use nearest neighbour sampling for bitmaps.

Always render using filtering: Force the GPU rendering engine to use bilinear sampling for bitmaps.

Match raster source by content: Deprecated.

Match raster source by content and raster: Under review.

Match internal raster source by content: Under review.

User resources are local: Deprecated.

Forget recently accessed algorithms: Clears the list of recently accessed MRD files, which normally appear on the Open menu button.

Reset defaults: Reset all options back to their default values. Hit OK to save the changes.

ProRaster Scientific

ProRaster Scientific is built upon the ProRaster Premium platform and extends it with a wide range of additional features and capabilities. See a summary on YouTube here.

ProRaster Scientific contains a powerful new processing interface for satellite multispectral data. In this interface, you can import and organise your satellite multispectral data scenes. It provides a plethora of ways to render, combine, process, and analyse that multispectral data.

ProRaster Scientific provides deep support for virtual raster technology. It ships with a range of processing operations, many of which generate virtual rasters that can then be rendered, used as a source for additional processing, or analysed.

ProRaster Scientific extends ProRaster Premium with new mapping features and new rendering algorithm capabilities. See a demo of these new features on YouTube here.

Mapping Extensions

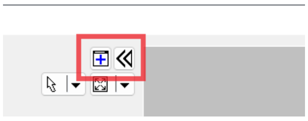

You can open multiple floating maps for each algorithm that is loaded. To open an additional floating map for the currently displayed algorithm, just press the "Open a floating map" button. When you select an algorithm, all the floating map windows open for that algorithm are brought to the front. When you hit the "Open Preview button", all the floating maps will be closed, and the preview window will be opened.

You can use the Time Control Panel dialog to control which time point (event) is being displayed in an algorithm. If you use the satellite multispectral analysis and rendering capabilities in ProRaster Scientific, you will routinely create and display rasters that include a time dimension. The Time Control Panel dialog makes it easy to manipulate this temporal dimension and change the rendered time point.

The Export operation has been extended to allow batch export of imagery using time as the iteration trigger. This operates just like the Time Control Panel and can be used to create numbered sequences of images of your multispectral data that can be automatically turned into MP4 videos by video editing software.

Algorithm Extensions – Mask Component

A new component called "Mask" is available in all layers. You can use the mask component to provide a raster rendering mask for the layer. When enabled, pixels in the layer will only be rendered where the mask raster source is valid and not zero. For example, the mask bands generated for satellite multispectral scenes can be used to prevent rendering of all pixels that are compromised by cloud or are identified as water or snow.

To use the mask component, just specify the raster source, field, band, and event range. You can apply a data conditioning filter to the mask raster source. Pixels are masked if the mask raster pixel is invalid, or it is valid and the data value is zero, or the data value is zero or invalid after data conditioning.

You can generate a mask raster via the "Create Raster Mask" option on the processing menu. This uses a raster calculator operation to generate a virtual raster that has valid values where you want to display pixels and invalid values where you want to mask pixels.

Algorithm Extensions – Pan Component

A new component called “Pan” is available in the RGB Color layer. You can use this component to apply panchromatic sharpening to any RGB Layer. Panchromatic sharpening uses a panchromatic multispectral color band with a higher spatial resolution to “sharpen” multispectral color bands with a lower spatial resolution. It has long been a feature of Landsat rendering where the 15m resolution Pan band is used to sharpen the 30m resolution spectral bands.

Pan-sharpening is performed by taking the RGB color components, converting to HSL color space, replacing the Luminosity with the scaled Pan band data, and then converting back to RGB color space.

To apply pan-sharpening, specify the raster source, field, band, and event range. You can apply a data conditioning filter to the mask raster source. You also need to provide a data transform to scale the pan data. You can use this transform to control the pan-sharpening process and adjust the output color levels.

Virtual Raster Technology

ProRaster Scientific makes extensive use of virtual raster technology. It uses virtual rasters to implement the processing stages in the multispectral satellite data product pipeline. It also generates virtual rasters to implement many of the processing operations available on the Processing menu.

A virtual raster is just like a normal raster file, except that it does not contain any actual raster data. However, it contains instructions that tell the virtual raster engine how to construct that raster data, on the fly and on demand. In other words, a virtual raster is a recipe for making a raster in real time. In most cases, a virtual raster will draw source data from other raster files, usually modifying it in some way, before emitting it.

In ProRaster, you can use a virtual raster wherever you might otherwise use a normal raster file. You can display virtual rasters by including them in rendering algorithms. You can use virtual rasters as input to processing operations. You can analyse virtual rasters and acquire statistics over regions. And a virtual raster can be a raster input source to another virtual raster, allowing you to build chains of virtual rasters to execute complex processing operations.

ProRaster can load three different kinds of virtual rasters – MVR, MRD, and VRT.

The MVR format (MapInfo Virtual Raster) is the primary virtual raster format used in ProRaster. The format is being actively developed and extended. An MVR file is a small text file containing information in XML format.

The MRD format is a ProRaster rendering algorithm. It may be somewhat counterintuitive, but these are also virtual rasters. When the virtual raster engine fields a request for data from an MRD, it invokes the rendering engine to render the algorithm and returns the data as RGBA imagery. You can use an MRD as input to a processing operation or as a raster input to another rendering algorithm. For example, you can execute a “clip to polygon” operation on an MRD to restrict rendering to the polygon set. Another useful technique is to build a multi-layer MRD using blending and then drape that MRD on a terrain surface.

The VRT format is a virtual raster format supported by the GDAL open-source raster library. ProRaster provides support for rendering and processing VRT files but does not provide any support for creating or editing VRT files. For further information on VRT, search the online GDAL documentation.

An MVR virtual raster is structured just like a normal raster, which means it needs to have fields, bands, and events, cell size and tile origin defined, and all the other properties of a raster clearly defined. However, most of the raster properties can be derived from the raster sources that are referred to by the virtual raster, so there is often no need to explicitly define them in the MVR file. Usually, these files are quite simple and simply link raster sources to field and band definitions via some processing operation.

Virtual raster data is cached by the raster engine just like any other raster data. If the requested data is not in the cache, then the virtual raster engine will be invoked to provide the data. This will trigger a process that may be deep and complex, depending on how that virtual raster data is defined. For example, virtual rasters are often arranged in chains where each MVR link in the chain draws on data provided by the link before it. Each link in the chain executes some processing operation on the data before passing it through.

Ultimately, data will usually be acquired from one or more source rasters (in rare cases virtual rasters define their data algorithmically). The raster data may have to be refactored in multiple ways – by reprojection, resampling and relocation. The virtual raster engine will acquire source data at different resolutions from throughout the raster data pyramid. For this reason, all source rasters ought to have a complete overview pyramid.

Satellite Multispectral Raster Data Processing, Rendering, and Analysis

ProRaster Scientific is a powerful platform for working with multispectral satellite raster data, such as Landsat and Sentinel 2. It helps you to prepare your multispectral scenes, manage your scene collection, render scenes, create products from scenes that implement waterfall processing pipelines, and analyse multispectral raster data products to acquire quantitative and qualitative measures of a wide variety of earth parameters.

A key concept to understand is that the software uses virtual rasters to implement processing pipelines for multispectral data scenes. This means there is no raster data duplication – all you require are the original scene rasters, and the virtual raster engine works from that original data, every time. It also means that processing operations are not executed in the normal way – they simply define a new link in the virtual raster chain. No time spent waiting!

The key design aims for the satellite multispectral data tool were -

1. **Ease of use**. Working with multispectral scenes has always been complicated and opaque. ProRaster automates the low-level processing so that you can immediately begin rendering and analysing raster data, confident in the correctness and quality of that data.
2. **Low technical knowledge requirement**. You no longer need to be a data scientist to work with multispectral data. ProRaster does all the hard work for you and, if you follow some key guidelines, you can perform data analysis like a professional.
3. **Low system resource consumption**. Multispectral data consumes storage space, and it is expensive to acquire and manage. By using virtual rasters, ProRaster ensures that there is no data duplication and storage requirements are minimised.
4. **High level of efficiency and productivity**. Working with multispectral data in ProRaster is fast. The virtual raster engine only evaluates what it needs to, on demand. That means it is only ever doing the minimum amount of work, and you are not sitting about waiting for operations to execute.
5. **High level of correctness**. By implementing indexes, transforms, and corrections to scenes in virtual rasters, based on the metadata supplied with those scenes, ProRaster takes all the drudgery out of data preparation. And that means no operator mistakes!
6. **Low entry point**. ProRaster Scientific is designed to democratise access to multispectral data by lowing the cost of entry by two orders of magnitude.

Processing Operations

ProRaster Scientific ships with a growing list of processing operations that allow you to modify raster data. In most cases, those operations generate a virtual raster output, and so the result can only be used in ProRaster itself. However, to provide interoperability with other software, you can “crystallise” a virtual raster and convert it to a raster in one of a variety of widely supported formats.

The processing operations support the satellite multispectral data analysis, and many of the same processing operations will be found in the multispectral product editor. Many of the operations support batch processing (where multiple input files are processed using the same parameters to generate multiple output files). The following processing operations are available -

1. Export to MRR
   Convert any raster format to an MRR format, filtering by field, band, and event.
2. Export to Transportable Raster Format
   Convert any raster format to a collection of simple rasters in a transportable format like GeoTIFF, filtering by field, band, and event.
3. Copy From
   (A subset of the Export capability) Create a lower resolution copy of a raster by extracting data directly from an overview level.
4. Statistics Analysis
   Compute statistics for any raster, filtering by field and band, and at any resolution level.
5. Join
   Join similar rasters, often tiles generated from some other process, together.
6. Clip to polygon
   Clip a raster to a complex polygon set (contain holes, islands, multiple parts etc).
7. Clip to raster
   Clip a raster to another raster.
8. Reproject
   Change the coordinate system of a raster.
9. Resample
   Change the cell size of a raster.
10. Realign
    Change the cell edge position of a raster.
11. Calculator
    Combine the values of multiple rasters using a calculator expression.
12. Create Raster Mask
    Create a raster suitable for use as a mask by defining a calculator expression.
13. Conditioning
    Apply a data conditioning operator to a raster.
14. Transform
    Apply a non-linear data transform to a raster band.
15. Image Time Series
    Combine multiple MRD algorithms into a time series virtual raster.

ProRaster Scientific Satellite Multispectral Processing

Welcome to the world of satellite multispectral data processing, rendering, and analysis! ProRaster Scientific gives you the tools you need to process, visualise, and analyse multispectral imagery.

To get started with multispectral imagery in ProRaster Scientific, you will follow a processing pipeline like this –

1. Download scenes from your data provider.

   The Earth Observing Satellites acquire data in swaths called a "scene". Each scene covers a static footprint on the Earth and the satellite will typically revisit that scene regularly and reacquire it many times over the life of the bird. This data is processed and then distributed to users as a package containing rasters for each spectral band and other information. This is a scene, and it is your starting point. A tile is a rectangular subset of a scene.

   You can download Landsat scenes for free from USGS Earth Explorer - https://earthexplorer.usgs.gov/

   You can download Sentinel 2 tiles for free from the Copernicus Browser accessed via the Copernicus Data Space Ecosystem –
   https://dataspace.copernicus.eu/

2. Unpack and store your scenes on your local drive or network.

   The scene will be delivered as a compressed archive. You will need to decompress (unpack) this archive and store the scene data package on your local drive or network. How you organise this is up to you.

3. Create a multispectral scene database and add the scenes to the database.

   In ProRaster, add each scene you download to your scene database. The database is a small file that links to the original raster data – it does not contain a copy of the raster data. ProRaster will generate a "scene assembly virtual raster" for each new scene/tile and store it alongside the database. You will access all data from the scene via this assembly MVR from now on.

4. Render the scenes.

   Take a look at your new scene! Open the scene database browser and select the scene. Right click and from the pop-up context menu, select "Render Scene". Choose a band combination to highlight specific attributes of the Earth. Your scene will then be displayed in ProRaster.

5. Create a multispectral product database.

   In ProRaster we distinguish between a basic scene and a multispectral data product. A product starts life as one or more scenes, to which you can apply processing operations in a waterfall cascade pattern. Before you can create a product, you need to create a database to store it in.

6. Browse for and select scenes from which to create a multispectral product.

   Open a scene database browser and select a scene. Right click and from the pop-up context menu, select "Create scene product". A new product will be added to your product database and the new product will be opened in the product editor dialog.

7. Edit the product and apply processing operations.

   You can apply a growing list of processing operations to the scene in your product including – Mask, Clip, Reproject, Calculator, and Index. You can generate a rendering algorithm, compute statistics, and export the data to a raster. Define these operations in any order you wish to build a waterfall cascade processing pattern. Thereafter, when you change any operation, all the downstream operations will get updated.

As your level of expertise improves, you will go on to combine products into more advanced products and combine rendering algorithms into more advanced visualisations. You will employ the "Create Product Sequence" wizard to create and edit your products.

To learn more about multispectral satellite imagery I recommend the textbooks available for download on the Earth Observation Australia website.

Compatibility

ProRaster Scientific supports different satellite platforms either directly, by providing automatic scene import, or indirectly by supporting manual scene importation. If you are interested in seeing direct import support added for any additional platforms, register your interest by contacting Roberts Geospatial Engineering.

ProRaster Scientific supports the Landsat platform, including the following families of birds.

- Landsat 1, 2, 3
- Landsat 4, 5
- Landsat 7
- Landsat 8, 9

The USGS distributes Landsat data and has done so for more than 50 years. Over the years the birds and instruments have changed. The processing has changed with it, as has the delivery package. ProRaster Scientific supports the following processing packages.

- Precollection
- Collection 1
- Collection 2 (current)

Landsat is delivered as a science product with a defined level of processing. ProRaster Scientific supports the following science products.

- Level 1 (top of atmosphere)
- Level 2 (bottom of atmosphere / surface reflectance / surface temperature)

ProRaster Scientific does not currently directly support the U.S. analysis ready data (ARD), Level 2 spectral indices, and Level 3 science products. This does not mean you cannot use these products in ProRaster, but you will need to work with them manually.

ProRaster Scientific supports the ESA Sentinel 2 platform, including the following families of birds.

- Sentinel 2A, 2B

ProRaster Scientific supports the PDGS IDP processing package. It supports versions 2.xx, 3.xx, 4.xx, and 5.xx processing baselines. The delivery of processing baseline 4.X was a major upgrade and is supported in ProRaster Scientific.

ProRaster Scientific supports the following Sentinel2 products.

- Level 1C (top of atmosphere)
- Level 2A (bottom of atmosphere / surface reflectance)

If you have multispectral data from any other platform, then you may be able to use it in ProRaster Scientific by manually importing the scenes. Even if this is not possible, you still have the option to use the Processing menu to manually process the data.

You may be able to download Landsat and Sentinel 2 scenes from a variety of other sources. If you do so, please be aware that ProRaster expects to find certain metadata with the scene that provides correction parameters and data. It also expects the data to be in certain arrangements.

As existing platform data delivery evolves, and new platforms are launched, ProRaster Scientific will adapt to provide support. As of mid-2023, we know that Landsat Next will launch in 2030 as a constellation of three birds. Landsat Next will provide a 6-day temporal repeat time and will have 29 spectral bands that include the traditional Landsat 8/9 and Sentinel2 bands, whilst adding new spectral bands. Spatial resolution will improve to 10/20 metre spectral bands and 60 metre thermal bands.

Limitations

ProRaster has no internal limitations that will limit the number of scenes you can work with. The primary limitation is the amount of installed physical RAM. Opening scenes consumes memory, and keeping enough scene data in the tile cache to optimise performance will consume all your memory! I recommend having at least 16GB of RAM. If you are in the market for a new PC, buy all the RAM you can afford. ProRaster has been tested working with products that source data from more than 500 scenes. The upper limit would appear to be many times higher than this.

An important practical limitation is storage space. Each scene will consume about 1.0 to 1.5GB of storage, and it is a good idea to have your scenes on an SSD. If you are running ProRaster on a laptop with only a 256GB main SSD, you will quickly run out of high-performance storage space.

CPU performance considerations may limit the number of scenes that you can work with effectively, although existing performance bottlenecks are being continually mitigated through regular software updates.

Concepts, Terms, and Outcomes

Radiance

When a satellite with a multispectral instrument observes the Earth, it measures the radiance of the Earth in discrete spectral bands that cover a certain range of electromagnetic wavelengths. For example, it measures radiance in the red, green, and blue bands, and from these bands we can build a “Natural Color” representation of the imagery. Typically, the spectral bands will extend beyond the visible light wavelengths, particularly to longer infrared wavelengths. Landsat includes a thermal camera that measures even longer wavelength thermal bands.

Radiance vs Reflectance

Radiance, the light the instrument sees, is ultimately converted into reflectance, which is purely a property of the material being observed. From a scientific standpoint, reflectance values are the numbers we want to work with. They typically range between 0 and 1.

Spatial Resolution

These measured radiance values are supplied to us as rasters, and rasters have a cell size that we refer to as the spatial resolution of the data. This cell size may reflect the natural spatial resolution of the instrument that acquired the data. The resolution tends to vary by instrument and by spectral band.

ProRaster links the "base resolution" of a scene to the "Blue" spectral band. For example, a Landsat 9 scene has a base cell size of 30x30 metres, and a Sentinel 2 scene has a base cell size of 10x10 metres. The panchromatic band, if supplied, will be higher resolution (15 metres in Landsat 9). Longer wavelength spectral bands will be lower resolution, and thermal bands tend to be lower still.
ProRaster will happily work with spectral bands that are different resolutions, but the same is not true for other software systems. For this reason, data providers often resample their spectral bands to a common resolution.

Spectral Resolution

In addition to a spatial resolution, the data will also have a spectral resolution. This reflects how many levels of grey the instrument can differentiate between, generally expressed as a number of bits. 8-bit imagery, for example, has 256 levels of grey.

Panchromatic band

Some platforms include a panchromatic spectral band, which covers a wide range of wavelengths. This pan band generally has a higher spatial resolution than the other spectral bands. For example, Landsat 7, 8, 9 return a 15x15 metre pan band in addition to the 30x30 metre resolution spectral bands.

Digital Numbers

All data providers supply the spectral data as Digital Numbers. Typically, this is a representation of the data scaled into a 16-bit integer. These numbers have no units and, as a rule, should not be used just as they are. The scene typically provides metadata that enables the user to convert the DN values to radiance via a relatively simple mathematical expression. This expression will require constant values, which may vary by scene, that are specified in the metadata.

Top of Atmosphere (TOA)

From radiance observed in space, the first major correction is to convert the digital numbers to top of atmosphere radiance, reflectance, and brightness temperature (in Kelvin for thermal bands). The reflectance correction includes a correction for the sun angle. This may be provided per pixel as a raster. ProRaster will perform this correction if required.

Bottom of Atmosphere (BOA) / Surface Reflectance (SR)

This secondary correction takes TOA data and attempts to remove atmospheric effects such as aerosol scattering and thin clouds. There are multiple inputs to this processing, and the inputs can vary by methodology. Inputs can include topographic DEM, MODIS derived water vapor and ozone models, air temperature, thermal bands, and sun angle.

BOA is considered the best data for scientific analysis, but it may not be available for all scenes. ProRaster does not compute BOA. You will need to source this data directly from your provider.

TOA Corrections

Landsat Level 1 scenes will be corrected to TOA in the scene assembly MVR. This is performed for spectral bands and thermal bands. The Landsat online documentation details how these corrections should be made for each of the Landsat birds. Each correction is based on constants that are supplied as metadata for each scene.

To compute reflectance involves a correction for the sun angle. Historically, this was supplied as a single value at the centre of the scene and was applied to all pixels in the scene. Collection 2 data includes a raster for the sun angle across the scene. When supplied, this is used by ProRaster to compute a per-pixel sun angle correction.

QA

All scenes are provided with quality assurance information, both for the scene as a whole and for each pixel in the scene. For example, the metadata may record the percentage of the scene that is covered by cloud. In addition, raster QA data may be provided that identifies cloud in each pixel. Other parameters like radiometric saturation, water, snow, and defective pixels may also be recorded. This information is generally quite difficult to tweeze out of the scene rasters. In ProRaster, it is used to mask pixels to prevent them from being rendered or used in data analysis and statistics.

Choosing DN vs TOA vs BOA

What spectral target should you use? ProRaster will give you spectral data as digital numbers, top of atmosphere corrected radiance and reflectance and brightness temperature, and bottom of atmosphere surface reflectance and surface temperature (if available). What data should you choose for rendering and analysis?

In other software packages, the desire to do as little work as possible influences the decision. So, you can technically use DN values to compute an index like NDVI, and it is easier to do so as no pre-processing is required. However, in ProRaster it is just as easy to use BOA as TOA or DN. My advice is to use the most processed data you have available. So, use reflectance and temperature rather than radiance. Use BOA rather than TOA, if you have it. Use DN values only if you have no other option.

It follows that I would recommend you download Level 2 (BOA) scene data from your provider rather than Level 1 (TOA) scene data. Unfortunately, for both Sentinel 2 and Landsat, your choice can force you to make a compromise.

In the case of Sentinel 2, the BOA data is superior as it contains a QA band and so you can mask out cloud and other unwanted pixels. This is so critical that I think it rules out using TOA. However, band 10 (Cirrus) is not included in the BOA product. Even so, my recommendation is to download BOA (Level 2A) Sentinel 2 tiles.

In the case of Landsat 8 and 9, the BOA product does not contain any additional critical information. Unfortunately, it does not contain the bands 8 (Pan), 9 (Cirrus), and 11 (TIRS2)! The most grievous loss, in my view, is the Pan band which prevents you from employing pan-sharpening. I find it difficult to make a recommendation for Landsat. It depends on how much you need, or would benefit from, BOA data over TOA data.

Tile vs Scene

Some satellite data is distributed as scenes, and some is distributed as tiles. A scene contains the full spatial extent of the raster data that the satellite collected in each imaging operation. Typically, a scene will have a rhombus shape, reflecting the north-south orbit of the bird and the rotation of the Earth. In contrast, a tile is a rectangular sub-section of a scene. It may be in the same coordinate system as the scene, but frequently it is in a different coordinate system. Only data from a single scene is distributed in a tile, even though the tile may overlap the footprint of multiple scenes. Consequently, the coverage in a tile varies depending on which scene is contributing data to the tile.

Scene vs Mosaic

A satellite will acquire data in a swath that is called a scene. However, some providers do not ship data as complete scenes. Instead, they split the data up into tiles which are in fixed locations. Consequently, the coverage in a tile may vary depending on where the satellite was in its orbit. Regardless of whether it is a tile or a scene, when you import it into ProRaster we consider it a scene from that point forward.

A mosaic is a collection of adjacent scenes. Scenes will generally overlap adjacent scenes a certain amount, so there are no gaps in a mosaic. Depending on how many scenes you gather, your mosaic may cover a region, a continent, or the whole planet. It is important to remember that adjacent scenes may not be contemporaneous. They may have been acquired on a different orbit of the bird. As a rule, the larger your mosaic, the broader the range of times at which the scenes were acquired. Remember that the birds travel in the NE direction, so scenes may be approximately contemporaneous from south to north, but from east to west this will not be possible.

Also, a mosaic may combine scenes that are supplied with different coordinate systems. In this scenario, it may be necessary to reproject each scene to a suitable common coordinate system. ProRaster will perform this reprojection on the fly.

Composited Scenes

When you mask a scene to remove low quality pixels – whether they be saturated, invalid, or cloud affected, or whether they represent surfaces that you do not want to include, like water and snow – you will almost always end up with a scene that has pixel gaps in it. You can fill these gaps with data acquired nearly contemporaneously. This may be an earlier or later revisit of the bird to that scene, or it may be a visit to that scene from a sister bird from the same family. So, we can composite multiple instances of a scene together, hopefully as contemporaneous as possible, to fill the gaps in the desired scene with quality data from other visitations of that scene. This is a composited scene. At the end of this process, we have an improved dataset that may be more suitable for rendering an analysis.

The virtual raster that represents the composited scene will take multiple input sample points per pixel, select one based on your rules, and will output a single data point per pixel. Usually, you will choose to mask out cloud and other unwanted pixels, and then fill the gaps by assigning new pixel values based on temporal proximity. Alternatively, you can use statistics to assign pixel values – choosing either the minimum or maximum, or computing the median or the mean. Or, you can choose a pixel value based on a Spectral Index calculation – selecting the pixel as it has the minimum or maximum index value, or is closest to some expected value.

Sequences

A sequence is a collection of revisits to the same scene recorded as separate temporal events. You can think of each event as a frame in a movie, except that the frame stores a dataset rather than an image. As a raster can contain multiple events, a sequence is still considered a single dataset and, as such, can be used in any processing operation.

By combining these concepts – scene, mosaic, composited, and sequence - we can build a set of possible datasets that can be created as a multispectral product.

Scene: A virtual raster representing a satellite scene or tile containing multiple fields, with multiple bands per field, containing corrected and uncorrected spectral data, QA data, ancillary data and imagery, and mask bands.

Composited Scene: A virtual raster for a scene that merges two or more masked revisits of the scene together, as contemporaneously as possible, preserving high quality pixels and rejecting low quality pixels to populate as many pixels with high quality data as possible.

Scene Sequence: A virtual raster for a scene that contains two or more independent temporal events, where each event represents the scene surveyed at a particular revisit time.

Composited Scene Sequence: A virtual raster for a scene that contains two or more independent temporal events, where each event is a composited scene that merges two or more nearly contemporaneous revisits of the scene together and represents the scene surveyed at a designated revisit time.

Mosaic: Two or more adjacent scenes, as contemporaneous as possible, merged into a single dataset.

Composited Mosaic: A mosaic where one or more contributing scenes is a composited scene that merges two or more nearly contemporaneous revisits of the scene together.

Mosaic Sequence: A virtual raster for a mosaic that contains two or more independent temporal events, where each event represents the mosaic surveyed at a particular designated revisit time.

Composited Mosaic Sequence: A virtual raster for a mosaic that contains two or more independent temporal events, where each event is a composited mosaic, and represents the mosaic surveyed at a designated revisit time.

Spectral Dataset

A spectral dataset is a virtual raster that is derived from one of the eight dataset descriptions above, that preserves the spectral bands and other bands from the original dataset. Furthermore, the original scenes that were used to create the dataset are all spectrally compatible.

We can argue about what the definition of spectral compatibility is. Are the spectral bands collected by Landsat 7 compatible with the spectral bands collected by Landsat 9? What about the spectral bands collected by Sentinel 2 vs Landsat? In ProRaster, I take a very strict definition of spectral compatibility. This is out of necessity, because if two scenes are combined in a dataset then it is a requirement that the virtual rasters have the exact same field and band structure. The consequence of this is that I consider scenes to be spectrally compatible if they come from the same satellite family, and they use the same processing, and they are the same product. So, for example, you can mix and match Sentinel 2A and 2B scenes if they are both TOA or BOA. You can mix and match Landsat 8 and 9 scenes if they use the same Collection, Processing Level and Product.

If scenes are not spectrally compatible, you cannot combine them into a single dataset. However, you can still treat them individually and then combine downstream products from those scenes, like spectral index computations or rendering algorithms. In that instance, you need to be aware of the possible incompatibilities.

Spectral Band Combinations

If we have red, green, and blue spectral bands then we can use a RGB layer to display the data and connect the appropriate spectral bands to the respective components. This is referred to as a “Natural Color” band combination. By using different spectral band combinations, we can generate imagery that highlights certain features of the earth, soil, vegetation, and water. ProRaster allows you to choose from about 20 commonly used band combinations, or to define your own. Not all these combinations will be available for all scenes – it depends on what spectral bands are present in the scene. For example, you will have a much richer choice of combinations in a Landsat 9 scene compared to a Landsat 5 scene.

You can add new spectral band combinations to the system for your own use via the RGB Combinations Editor dialog. You can also assign custom data transforms to the different components to customise the way you render spectral data, right out of the box, in ProRaster Scientific.

Index Computations

Index computations, drawing data from one or more spectral bands, use a mathematical expression to compute a spectral index. Scientists have published hundreds of these formulas, but there are just a couple of dozen that are in common use. For example, the Normalised Difference Vegetation Index (NDVI) is used to compute and compare the amount of green vegetation. The formula to compute the index is:

NDVI = (NIR-Red)/(NIR+Red)

ProRaster contains scores of these index computations as generic formulas. If the scene or dataset you are processing contains the required spectral bands, then ProRaster will allow you to select and compute the index. Note that some index formulas assume that the inputs are scaled to between 0 and 1 (i.e., reflectance).

Difference Computations

If you have data acquired at two or more times, then you can compute the difference between the data acquired at those times. Generally, the data will be an Index or some other kind of computed value. With a sequence data product, you can compute the difference between consecutive time points (events) and create a new sequence data product that quantifies how the data changes over time. You can also compute the difference between a key time (event) and each other time.

Sometimes a difference is computed as a ratio rather than as a difference. This is supported in the Difference operation as the calculator expression used to compute the difference between two events in a product sequence is editable by the user.

Rendering

ProRaster is designed to render raster data, and rendering satellite multispectral data is just what it likes to do.

You will see options to “Display Scene” or “Render Scene” (or whatever data product you are working with). In the first case, the virtual raster will be opened in the main ProRaster interface using a default Look-up Color layer. This will allow you to explore the fields, bands, and events of the MVR and render each band individually.

In the second case, the virtual raster will be opened in the main ProRaster interface using an RGB color layer. A dialog will be presented asking you to choose from a list of valid band combinations, or you can define your own combination. You will also be able to enable pan-sharpening, if available. You will be able to enable masking, if available.

Pan-sharpening in ProRaster is only available when rendering. It is not a processing operation and cannot be used as any kind of input to a processing operation.

Masking in ProRaster is available in any layer and can be used to prevent rendering of pixels that overlap invalid pixels in the mask raster. Although it can be used with any kinds of rasters, its main purpose is to mask out compromised pixels in multispectral data.

Clipping and Masking

When analysing data, bad pixels in means bad statistics out, and a compromised result. Therefore, you will want to mask out bad pixels aggressively and you will want to use polygons to clip the dataset to only that area that you wish to study. In ProRaster, this is achieved by "Clipping to raster" (masking) or "Clipping to polygon".

To help mask out bad pixels, ProRaster builds a Mask field that contains some standard bands. It populates these bands by interrogating the QA band data associated with the scene. For example, you will often see these bands - Pixels (clear), Land (clear), Water (clear), Snow (clear). If you only want pixels that are designated Land, and you want to reject all compromised pixels, choose "Land (clear)". If you just want to reject band pixels, choose "Pixels (clear)".

This data is supplied by the data provider and is of mixed quality. As a rule, older processing is not as good as later processing. You might find that an old Landsat scene in Precollection format may have much poorer QA data than the same scene reprocessed to Collection 2 format. Even so, often the Land – Water – Snow differentiation is poor. In my view, the pixel labelling is rarely aggressive enough and many pixels that are designated as clear of cloud may be partially cloud affected.

Statistics

For any raster dataset you can compute statistics. This will compute summary statistics - like the minimum, maximum, mean – as well as distribution statistics which include a distribution histogram and other measures like the standard deviation, median, mode and quartiles. These statistics can be used to form the basis of your data analysis. You can track how the statistics change over time and quantify the changes.

Export

ProRaster Scientific creates MVR virtual rasters that cannot be loaded, displayed, or used in other products. To get your raster data out of ProRaster and into other software, you can use the export capability. This will write the raster data to a new raster file in a transportable format.

The MRR format is best suited for exporting MVR files. An MRR can contain multiple fields, bands, and events, and can store all the data in an MVR in a single raster file. However, MRR is difficult to use outside of ProRaster and MapInfo Pro.

ProRaster supports exporting to transportable raster formats like GeoTIFF and Band Interleaved (BIL). These formats are not able to support the complicated field, band, and event structure on an MVR raster. Consequently, when you export to a transportable format every data band (by field and event) is written to a separate raster file. Even so, concepts like Classified fields are not properly supported by other raster formats or software.

MVR Chains

An MVR virtual raster is connected to one or more sources of input raster data. It will usually modulate that data in some way – by applying a calculator expression for example – and then emit it again as output. An MVR can be the source raster to another MVR. This allows ProRaster to build chains of MVR rasters that implement multiphase processing operations.

Creating chains of MVR rasters does not invoke any processing. However, when you consume the MVR at the end of the chain – by rendering it, computing statistics, exporting it, or in any other way that requires cell values to be evaluated – the entire processing chain will be invoked to evaluate raster cell data. This is performed in real time, on demand, as data is only ever requested over a particular spatial range and at a particular spatial resolution. A spatially restricted data request such as this can usually be responded to promptly.

Scene Assembly MVR

When you download a scene from a provider, you will be supplied with a collection of rasters – one for each spectral band (possibly at multiple spatial resolutions), as well as other rasters for QA data, ancillary data like sun angle, and pre-rendered imagery.

After importing that scene into a scene database, you will have one additional raster – this is the scene assembly MVR. From then on, ProRaster will interact with the scene rasters through this virtual raster. It provides a host of services including –

- Directing the raster engine to load the source rasters through a specific driver, if needed.
- Supplying the raster engine with driver preferences for each source raster.
- Directing the raster engine to build overview level pyramid caches, if needed.
- Gathers the spectral bands into multi-banded fields.
- Implements transformations to convert DN to TOA radiance and reflectance.
- Implements transformations to convert thermal DN to TOA temperature.
- Implements transformations to convert DN to BOA reflectance and temperature.
- Gathers the ancillary raster data into multi-banded fields.
- Picks out the individual QA values for each pixel and exposes them in multi-banded fields.
- Builds a mask field that uses the QA data to identify pixels of different classification and quality.

The scene assembly MVR will correct historical and current processing mistakes. For example –

- Landsat data used to ship without a defined null value, injecting zero values into the raster.
- Landsat COG files were shipped with the zero values incorrectly included in the overview levels, making that data technically incorrect.
- Sentinel 2, which uses JPEG2000 as a raster format, currently delivers overview levels for QA data that are technically incorrect. The overviews blend data values that cannot be combined. ProRaster computes a new overview level cache for the QA data in Sentinel 2 scenes.

When the scene assembly MVR is created, you can opt to duplicate the scene data to rasters in MRR format. This may provide improvements to correctness and performance.

Importing Scenes

ProRaster Scientific provides automatic import for scenes from directly supported platforms. This makes importing easy and the scene assembly MVR is automatically generated for the scene. For unsupported platforms, there is a manual import process. You will be required to supply some key information about the platform and the scene, as well as identify all the source rasters and link them to spectral channels. Most multispectral platforms can be supported via the manual import process.

Revisiting Scenes

Landsat 8/9 and Sentinel2 A/B revisit a scene every 16 days. The birds in a family are staggered by 8 days, so within a family you get a revisit every 8 days. Rows to the north and south will be imaged at almost the same time.

Every second path is surveyed by a single bird every second day, but from one path to the next adjacent path or previous adjacent path there may be a gap of 7 or 9 days, either before or after, for a single bird.

With two birds and a set of up to 8 adjacent paths, you can acquire a path every day (from West to East). Each path will be visited by a different bird (first one bird, then the second, repeating). Across up to 8 paths, you ought to be able to acquire data across a span of up to 8 days (and as many rows North and South as you want).

Across more than 8 paths you can acquire all scenes over no more than an 8-day period. So, even for a large mosaic, you can achieve quite a compact temporal range of no more than 8 days. In other words, the satellite family completely surveys the Earth every 8 days.

If you were to combine both Landsat and Sentinel2, you may be able to improve the revisit time. If they are staggered perfectly then revisit time could be reduced to 4 days. However, I do not have any guidance for this yet.

Industry Terms vs ProRaster Scientific Terms

The Earth Observation industry sometimes uses terms that are different to ProRaster Scientific. Here are some terms used in ProRaster and the equivalent industry term.

RGB Color layer
An RGB Color layer in a rendering algorithm is also called a “Color Composite layer”.

RGB/Band Combination
An RGB/Band Combination, which combines three standard spectral bands in an RGB Color layer is also called a “False color composite”.

Composited product
I used to refer to compositing as “collating” and you will see me use this term in older videos. In Google Earth Engine, compositing is referred to as “reducing”.

Mosaic / Merge
An image Mosaic, or the process of merging imagery, is also called “image sewing” or “image stitching”.

Stitch / Histogram matched
A stitched image that uses spectral scaling or histogram matching is also called a “Mosaic”.

Pan-sharpening
Pan-sharpening is also called “spectral blending” or “fusion”.

Mosaic
A mosaic is also called “spatial blending” or “fusion”.

Sequence
A sequence product is also called “temporal blending” or “fusion”. A sequence product is often referred to as multidimensional (the additional dimension being time).

Scene sequence
A scene sequence (containing two frames) is also called a “bi-temporal image”.

Mosaic Sequence
A mosaic sequence is also called a “multi-temporal image”.

Combining imagery spatially

Sometimes this is referred to as spatial blending or fusion. There are many different and confusing terms that describe the combining of images. These are the terms that are used in ProRaster Scientific.

A *Join* operation combines two or more images and combines them into a single image. Where images overlap the value of one of the images is used (usually by establishing an image order). The raster structure, the cell size, and the coordinate system of the input images is expected to be the same. Sometimes this kind of operation is called image *sewing*.

A *Mosaic* operation combines two or more (usually multispectral) images and combines them into a single image. Where images overlap the value of one of the images is used (usually by establishing an image order). The raster structure of the input images is expected to be the same. The cell size and the coordinate system of the input images can vary.

A *Merge* operation combines two or more images into a single image. Where images overlap, a rule is applied to determine the final pixel values. This may assign the value from one of the images or it may combine the values from multiple images by *blending* (for example, by averaging). The raster structure of the input images is expected to be the same. The cell size and the coordinate system of the input images can vary.

A *Stitch* operation combines two or more images into a single image. Each input image may be *levelled* prior to the stitch, which may apply a multidimensional offset and scale to each pixel, to ensure values are compatible in level and range. Where images overlap, a rule is applied to determine the final pixel values. This may assign the value from one of the images or it may combine the values from multiple images by *blending* (for example, by averaging). Where two images overlap, the edge may be feathered to reduce the visual impact of the edge. The operation will likely target a single specified data band from each input raster. The cell size and the coordinate system of the input images can vary. Sometimes this kind of operation is called a *Mosaic*.

A *Histogram Matched* operation uses histogram matching to transform the data values of all the input raster bands into the data space of the reference raster. Histogram matching transforms a data value by converting it to a percentile (from the input raster histogram) and then back to a data value (from the reference raster histogram). It may be used as a strategy in levelling input images prior to a stitch operation.

Combining imagery temporally

In ProRaster Scientific this is called “compositing”, but I have previously used the term “collating”, and it is also called “temporal blending”, “fusion”, and “reducing”. A composited product is a pixel-based composite. When we composite, we (usually) mask out pixels in the primary image that are impacted by cloud or some other kind of noise or unwanted property and replace these pixels by selecting the “best” pixel from a set of images acquired over a specified period. The best pixel is likely to be chosen as it is cloud free and the acquisition time of that pixel is closest in time to the acquisition time of the primary image. This process may introduce artefacts.

The Multispectral Scene Database

After you download a satellite multispectral scene from a data provider, your first step will be to add it to a multispectral scene database. You can do this automatically or manually, depending on whether your scene is compatible. This performs three key tasks –

1. In an automatic import, the scene is parsed, and metadata is acquired. ProRaster works out what the scene is – what platform, what bird, what processing and product. It works out what spectral bands are available, what QA data is supplied, and what ancillary data is available. If ProRaster recognises and supports the scene, then it can be added to a scene database. This will write a small amount of metadata into the database, including a record of the location of all the raster source files in the scene.

   In a manual import, it is your job to supply all the information about the platform, bird, and scene. You need to identify all the input rasters and link them to supported spectral bands.

2. An MVR virtual raster file (the scene assembly MVR) is prepared for the scene and written into a folder subordinate to the database. This MVR gathers all the data associated with the scene and combines it into a single virtual raster. It applies spectral transformations to convert from Digital Numbers to radiance, reflectance, temperature. It may make top of atmosphere corrections. It picks out information from the QA data and presents it as individual raster bands. It uses the QA data to prepare a mask field that can be used to identify pixel classifications and pixel quality.

3. The source raster files are finalised. It is a requirement of ProRaster that each source raster in the scene has a full overview data pyramid. If it is not supplied, or it is supplied but is invalid, then it will be generated. This will produce a .PPRC cache file for each scene raster file.

Creating a scene database

You interact with scene databases via the “Browse Scenes” menu button. You can see this explained on YouTube here.

Drop down the menu button to reveal the options for creating, editing, managing, and deleting scene databases.

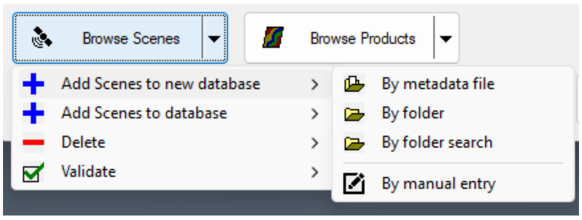

Create a new scene database by selecting “Add Scenes to new database” and selecting one of four methods for adding one or more scenes to the database. The first three methods – by metadata file, folder, and folder search - perform an automatic import. The last method – by manual entry – initiates a manual import.

Before you browse to the scene(s), you need to create and name your new scene database. Your database must have a unique name. The scene database file will use this name. If the name is unique then you will see a green tick. Hit OK to create the scene database.

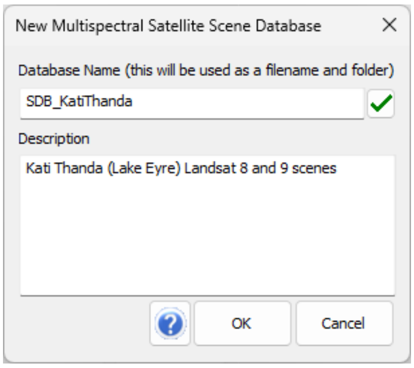

When doing an automatic import, you now must browse to the scene (or scenes) that you want to add to the database. There are three ways to do this – by metadata file, by folder, or by folder search.

1. Browse for and select the key metadata file associated with the scene. A single scene will be imported.

   For a Landsat scene, the key metadata file is the *_MTL.txt file. These files have been distributed with each Landsat scene from the beginning of the program. These days the data is mirrored in JSON and XML files.

   For a Sentinel 2 scene, the key metadata file is an XML file that is named according to the processing level. Look for files named "MTD_MSIL1C.xml " and "MTD_MSIL2A.xml".

2. Browse for the root folder that the scene is stored in. ProRaster will search for the metadata and import a single scene.

   This is an easier option and pushes the work of finding the metadata file back to ProRaster. Simply browse to the root folder for the scene and ProRaster will search for the information it needs. It relies on you storing each scene in a unique folder. This is a practice that you ought to follow religiously. Do not dump the data from multiple scenes into a single folder!

3. Browse for a folder. ProRaster will search in this folder, and in all its sub-folders, for any scenes recursively. Every scene that it finds will be imported.

   This is the easiest option. Using this method, you can add as many scenes as you want to your database. Simply browse to a folder underneath which you are storing all the scenes you want to import. ProRaster will detect and import all the scenes it finds under this folder. To take advantage of this, it makes sense to keep your scene folders organised into a project hierarchy so that you can easily add the scenes you want for a particular project to a scene database.

Once you have specified where to find the scene(s), a processing dialog will open, and you will be able to monitor the progress of ProRaster as it imports the scene(s).

You can create as many scene databases as you want. How you organise your scenes is up to you. You may want to add all scenes to a single large database, or you may want to create a database for each project that you undertake.

- A scene database file has a "<your name>.msssdb" file name. You must specify the name when you create the database. You cannot change the name thereafter.
- Every scene database you create will be located in the "C:\Users\<your user name>\RGE\MSS" folder on your main drive. You cannot move this file.

- When you create a scene database it creates a new folder in the "C:\Users\<your user name>\RGE\MSS" folder of the same name. For example, if I create a new scene database called "SDB_KatiThanda.msssdb" then a new folder "C:\Users\<your user name>\RGE\MSS\SDB_KatiThanda" will be created.
- The new folder will contain an MVR virtual raster file for every scene in the database. This is the key, and only, raster resource that you will interact with for the scene. You will never refer to the original raster files distributed with the scene. These are all accessed via this key MVR file.
- No raster data from the scene is ever written into the scene database. Consequently, the scene database will be a small file, even if it contains thousands of scenes.
- When you add a scene, the overview pyramid data cache will be prepared (if required). This takes a few moments per scene. The cache files are written into the same directory as the original scene raster files. This is the only modification to the original scene folders that is made. Once these rasters have been generated, they will not be regenerated.
- You can add a scene to two or more scene databases without penalty.
- When you delete a database both the .msssdb file and the folder and all its contents will be deleted.
- When you delete a database the original scene folder and all its contents remain untouched.

From these points, you can conclude that the scene database is a relatively low value resource. Databases can be created, and recreated, easily. However, once you consume a scene in a multispectral product, that product will refer to the scene database by name. If that database is deleted, then the product will become orphaned.

Scene data duplication

By default, when you import a scene the scene raster data will not be duplicated. The scene assembly MVR will refer directly to the original source rasters in the scene. However, there are options to duplicate the scene raster data which may be of interest in advanced scenarios. These options currently apply to Landsat and Sentinel2 scenes, but not to scenes that are manually imported.

When ProRaster Scientific accesses raster data in a scene, it may take multiple steps to improve, correct, augment, or replace that raster data.

- The raster may be missing critical information. For example, a TIFF file may not record the raster null value. In scenarios like this, ProRaster will use driver preferences to provide critical missing information at run time. This information is stored in the scene assembly MVR.
- The raster may not have an overview cache. ProRaster will create an overview cache and store it alongside the original raster. The cache file will have a PPRC extension.
- The raster overview data may not be valid or entirely correct. For example, if null values are not identified, they may be used to construct the overview and combined with valid data. The result is incorrect data values. Another common mistake is to generate overviews from QA data bit masks, and then to simply average those values. This generates invalid bit mask values in the overviews. ProRaster may choose to correct these errors by generating a new overview cache which will be used in preference to the overview data in the original raster.
- If ProRaster computes statistics for a raster from the base level data, it will store those statistics for future reference. The statistics will be stored alongside the original raster in a file with an STTX extension. An ASCII file with a GHX extension may also be written.
- For efficiency and correctness, ProRaster can replace a scene raster with a new raster in MRR format. This is called a *duplicate*. Thereafter the scene assembly raster will refer to the MRR duplicate raster and not the original scene raster.

ProRaster Scientific provides control over the creation of duplicate rasters. This is controlled from the MSS Options dialog, which is accessed from the system menu. Duplicate rasters can be stored alongside the original raster in the scene folder, or they can be stored in a duplicate cache folder of your choosing. There are four options.

1. Duplication: Minimal intervention

   The goal of this mode is to prevent, where possible, writing of any files to the scene folder. ProRaster will use the original scene rasters and the rasters ought to provide overview levels.

   If any duplicates already exist for the scene, they will be used. They may be in the scene folder or in the duplicate cache folder. If overviews are not available, PPRC cache files will be written into the scene folder. No new duplicates will be generated.

2. Duplication: Recommended

   The goal of this mode is to provide a balance between minimising storage requirements and maximising correctness. It is the default mode.

   If any duplicates already exist for the scene, they will be used. They may be in the scene folder or in the duplicate cache folder. If overviews are not available, PPRC cache files will be written into the scene folder. If overviews need to be corrected, replacement overview cache files will be generated. No new duplicates will be generated.

3. Duplication: Full on demand

   The goal of this mode is to provide duplication of the source raster data.

   If any duplicates already exist for the scene, they will be used. They may be in the scene folder or in the duplicate cache folder. If any duplicates are missing, they will be generated.

4. Duplication: Full regeneration

   The goal of this mode is to correct and replace duplicates. It should only be used temporarily, and then the mode can be changed back to “Full on demand”.

   All duplicates will be regenerated.

The final option is to specify the folder that the duplicate rasters will be written into. By default, this folder is not specified, and in this case all duplicates are stored in the scene folder alongside the original scene raster. If you specify a duplicate cache folder, the duplicates from each scene will be written into a unique folder in the folder you specify. The system will create and manage these sub-folders.

Manually adding a scene

To manually add a scene to a scene database, choose the “By Manual Entry” option on the Browse Scenes menu button to either add a scene to a new database, or add a scene to an existing database. The MSS Scene Entry dialog will be displayed.

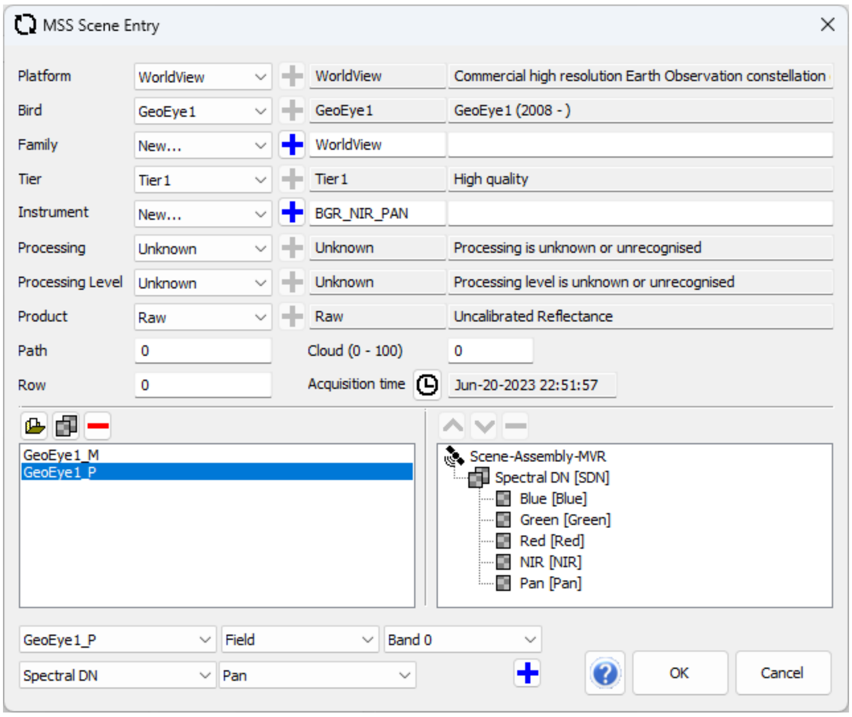

You are encouraged to enter appropriate information that describes the platform, bird, processing, and scene and you will need to identify the physical raster files on disk and specify which spectral bands they represent.

The "Platform" identifies the satellite constellation and may encompass multiple birds and families of birds. For example, the "Landsat" platform encompasses all Landsat bird's past, present, and future. Select a platform from the drop list. If you need to define a new platform, select "New…" and then you can type in the name of the new platform and a description (if you wish). Once you have entered the new platform name and description, make sure you hit the blue "+" button to store the new data. The blue button will turn grey, and your new platform will be displayed in the drop list.

Repeat this exercise for the Bird, Family, Tier, Instrument, Processing, Processing Level, and Product. The Bird identifies the satellite that acquired the scene. Within a platform, birds may be grouped by Family. All birds in a family are expected to use the same instruments. The Tier refers to the quality of the data and it is usually appropriate to select Tier 1 (high quality). The Instrument is a code name for the spectral instrument, or combination of instruments, on the bird. Processing refers to the processing epoch used to process the scene data. Processing Level refers to the level of corrections or processing operations applied to the scene data. For example, the processing level may be raw digital numbers, top of atmosphere reflectance, or surface reflectance. The Product refers to the output data product that was generated by the processing.

This data is used by ProRaster Scientific when you try to combine multiple scenes together into a product like a mosaic or sequence. Scenes need to have matching parameters. It is up to you to enter these parameters appropriately. I recommend you enter Platform, Bird, and Product data as a minimum.

The Path and Row parameters are mashed together to identify the tile. You may not have specific orbital path and row information, but you may have some identifying information for the tile. You can enter this into the Path parameters and leave Row blank if you wish.

You can enter a Cloud percentage coverage (from 0 to 100) for the scene if you wish (and if you know what it is). This will help you order scenes by quality when creating products.

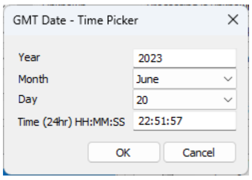

You need to enter the acquisition time for the scene. This is the UTC time (Coordinated Universal Time) at which the scene was acquired by the bird. It is not necessary to be completely accurate, but if you want to create sequence products then the scene must have a valid acquisition date and time. When you hit the time button, the GMT Date – Time Picker dialog will be displayed and you can enter a date and time.

Having defined the scene information, you now need to identify the scene source rasters and build the scene assembly virtual raster by linking source raster bands to spectral bands. Firstly, add all the source rasters to the raster list. You can browse to select one or more rasters, choose an existing raster source, or drag and drop rasters from File Explorer into the list. The order of the rasters in this list is not critical.

The Raster Source Editor gives you the option to link driver preferences with raster files. For example, a commonly used option for TIFF files is to specify an invalid cell value (null value). If you require any driver preferences, you will need to build your raster sources first in the Raster Source Editor dialog, and then select them in the MSS Scene Entry dialog.

Now you can build the scene assembly MVR, one band at a time. As you build the MVR the fields and bands will be displayed in the tree. You can reorder the bands by moving them up or down, or you can delete bands, by selecting them in the tree and hitting the appropriate button above.

Using the three drop lists at the top, select the source raster, source raster field, and source raster band. Now link this band to a spectral field and band using the two drop lists at the bottom. You can choose one of the following fields listed below. The list shows the code identifier, the field name, and the field description.

SDN:	Spectral DN	"Spectral digital numbers"
STREFDN:	Spectral TOA REF DN	"Spectral top of atmosphere reflectance digital numbers"
STRAD:	Spectral TOA RAD	"Spectral top of atmosphere radiance"
STREF:	Spectral TOA REF	"Spectral top of atmosphere reflectance"
SBREFDN:	Spectral BOA REF DN	"Spectral bottom of atmosphere reflectance digital numbers"
SBREF:	Spectral BOA REF	"Spectral bottom of atmosphere reflectance"
TDN:	Thermal DN	"Thermal digital numbers"
TTTDN:	Thermal TOA TEMP DN	"Thermal top of atmosphere temperature digital numbers"
TTRAD:	Thermal TOA RAD	"Thermal top of atmosphere radiance"
TTT:	Thermal TOA TEMP	"Thermal top of atmosphere temperature"
TBTDN:	Thermal BOA TEMP DN	"Thermal bottom of atmosphere temperature digital numbers"
TBT:	Thermal BOA TEMP	"Thermal bottom of atmosphere temperature"
MASK:	Mask	"Pixel validity mask"

The spectral fields allow you to choose one of the following spectral bands. The list shows the code identifier, the band name, and the band description.

Coastal:	Coastal Aerosol	"Coastal Aerosol Blue"
Blue:	Blue	"Visible Blue"
Green:	Green	"Visible Green"
Red:	Red	"Visible Red"
Pan:	Pan	"Panchromatic"
RDEG1:	Red Edge High	"Vegetation Red Edge High Frequency"
RDEG2:	Red Edge Mid	"Vegetation Red Edge Mid Frequency"
RDEG3:	Red Edge Low	"Vegetation Red Edge Low Frequency"
NIR:	NIR	"Near Infra-red"
NWIR:	Narrow NIR	"Narrow Near Infra-red"
Water:	Water Vapor	"Water Vapor"

Cirrus:	Cirrus	"Cirrus Cloud"
SWIR1:	SWIR High	"Short wave Infra-red High Frequency"
SWIR2:	SWIR Low	"Short wave Infra-red Low Frequency"

The thermal fields allow you to choose one of the following thermal bands. The list shows the code identifier, the band name, and the band description.

TIRS1:	Thermal High	"Thermal Infra-red High Frequency"
TIRSHG:	Thermal High Gain	"Thermal Infra-red High Frequency High Gain"
TIRS2:	Thermal Low	"Thermal Infra-red Low Frequency"

The mask field allows you to choose one of the following mask bands. The list shows the code identifier, the band name, and the band description.

MPC:	Pixels (Clear)	"Unobstructed pixels"
MPLC:	Pixels (+ largely clear)	"Largely unobstructed pixels"
MPPC:	Pixels (+ partly clear)	"Partly obstructed pixels"
MLC:	Land (Clear)	"Unobstructed land pixels"
MWC:	Water (Clear)	"Unobstructed water pixels"
MSC:	Snow (Clear)	"Unobstructed snow pixels"
MEX:	Extant	"Observed pixels"

Once you have selected the source raster, field, and band, and the destination field and band, the "+" button will turn blue. Hit the button to add the band to your scene assembly MVR, and it will be displayed in the tree. As you add new bands, they will automatically be sorted into fields.

Once you have defined all the parameters, check it over thoroughly. Make sure everything is correct. Then hit the OK button to add the scene to the database. The scene assembly MVR will be written, and the scene database browser dialog will open. If you need to import multiple scenes, the new parameters that you defined will be saved so that you will not need to enter them again (whilst ProRaster Scientific is running).

Browsing a scene database

You interact with scene databases via the "Browse Scenes" menu button. See a demo on YouTube here.

If you click on the button, it will open a browser for a scene database. If you have more than one scene database, then a dialog will open allowing you to choose which scene database you wish to browse. Double click on the database in the list or select the database and hit OK.

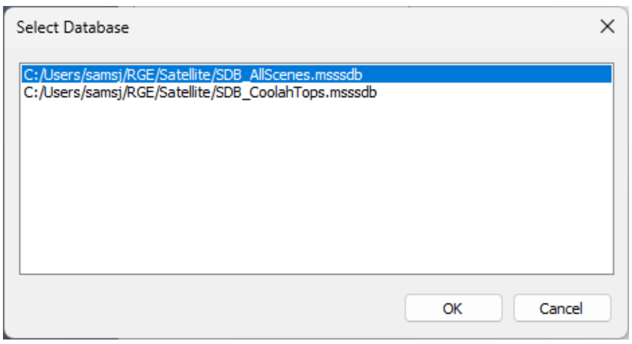

Selecting a Scene Database

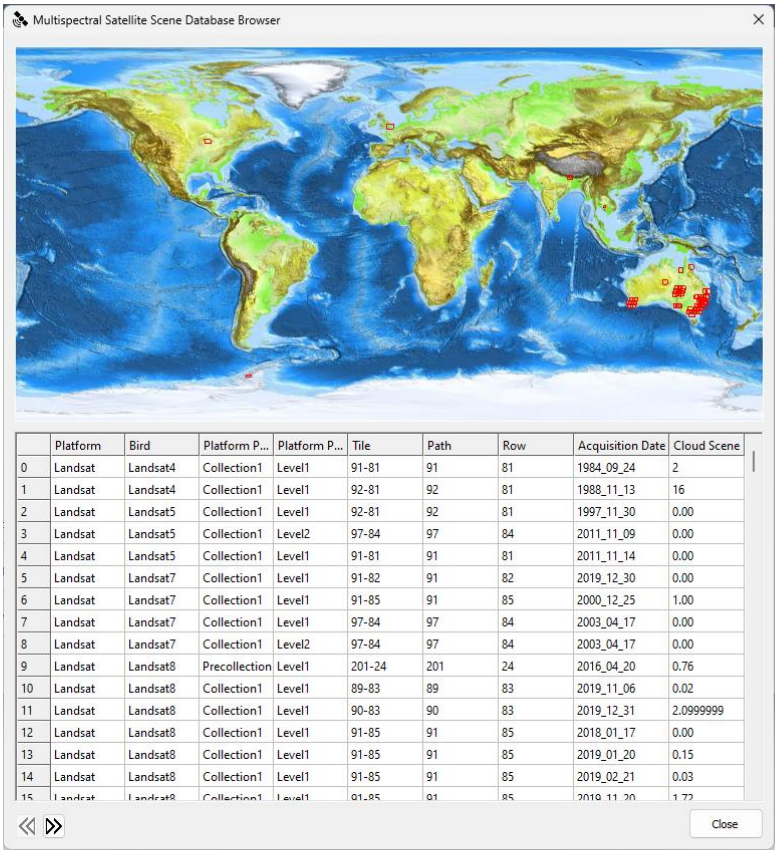

Scene Database Browser

The scene database browser dialog contains a world map at the top and a spreadsheet at the bottom. Each scene in the database is represented by a row in the spreadsheet.

The map shows the terrain and countries of the Earth. The spatial extent of each scene in the database is shown on the map as a red rectangle. Zoom the map with the mouse roller and pan by left clicking and holding, then dragging. As you move the cursor about the map a tool tip sows the coordinate of the cursor. If you are hovering over a scene, the ID of the scene will be shown also.

The purpose of the browser is to allow you to find and select a scene, or a group of scenes, from which you want to create a product. To help you achieve this, you can make queries. You can make a query from the map or from the spreadsheet. A query will probably reduce the number of scenes in the spreadsheet, which is reflected in the map. At the bottom left are two buttons that allow you to execute or reverse a query.

If you left mouse button double-click in the map you can make a spatial query that selects all the scenes that intersect the location that you clicked. You can also zoom and pan in the map, then right mouse button click to bring up the context menu and hit “Select visible scenes”. This will execute a spatial query that selects all scenes that intersect the map bounds.

In the spreadsheet you can execute a query by selecting one or multiple rows (scenes) in the sheet, and then hit the execute query button (forward arrow). This will collapse the spreadsheet to only those selected scenes.

Often you will want to find scenes by filtering by platform, bird, processing, and other scene parameters. Click on the column header to sort the spreadsheet by the column. Scroll down to find the parameters you are interested in. Left mouse button double click on the cell and the browser will select all the rows that have matching data in that column. You can then hit the execute query button to collapse the sheet to those selected scenes.

The map may not always display all scenes in the database. This will depend on the map projection employed. Right click in the map and, from the context menu, hit "Projection". You can choose from Standard, WebMap, Globe, North Pole, and South Pole.

Right click on the spreadsheet column header to control the column widths and to choose which columns are displayed in the sheet. These data items are acquired from the metadata for the scene.

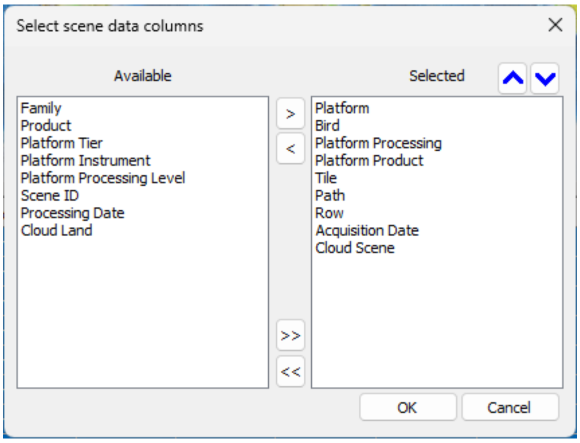

If you select a scene and right click on it, a context menu will provide you with options to display or render that scene or to create a product from that scene, or that group of selected scenes.

Editing a scene database

Drop down the "Browse Scenes" menu button to see options for editing a scene database.

You can add a new scene, or many new scenes, to a database. As before when creating a new database, this can be by browsing for the scene metadata file, browsing for the scene folder, or browsing for a root folder and adding all scenes that ProRaster can discover under that folder.

You can easily delete duplicate scenes from a database. If you use the folder search option, it is easy to add a scene to a database more than once. There is no harm in this, and the problem is easily rectified by choosing the "Delete > Duplicate Scenes" option.

You can delete a database by selecting the "Delete > Database" option. The database selector dialog will open and you can choose the database to target. The database, and the folder containing the key MVR files will be sent to the recycle bin.

The "Validate" option gives you some options to validate the scenes in your database. "Validate" will make sure that all the files it expects in every scene are present and reachable. "Regenerate" will regenerate the key MVR file for each scene in the database.

Rendering a scene

In the scene browser, select a scene and right click to bring up a context menu. You will be presented with options to "Display scene", "Render scene" and "Export scene". For a demo see YouTube [here](#).

Choose “Display scene” to open the key scene MVR file in a default algorithm in the main ProRaster interface. It chooses a Look-up Table color layer, so you can explore all the fields, bands, and events of the raster. If you wish to keep the algorithm, use the normal mechanism to save it to an MRD file.

Choose “Render scene” to open the key scene MVR file in a default algorithm in the main ProRaster interface. In this case it uses an RGB layer and it needs to assign spectral bands to the red, green, and blue color components. The following dialog is displayed to enable you to exercise control over the band assignments.

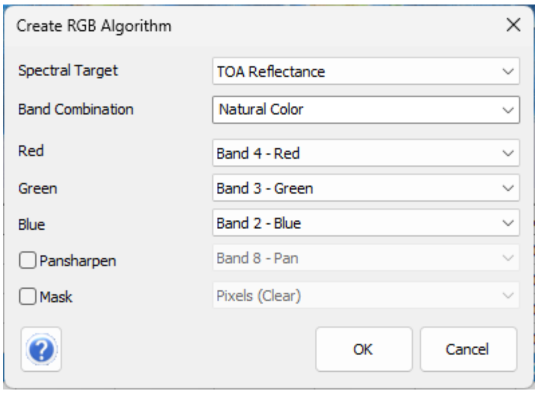

Firstly, select the spectral target. ProRaster will automatically select the most advanced spectral target, and this is usually the best choice. You may be presented with the following options, depending on what spectral data is available in the scene –

- Digital Numbers
- TOA Digital Numbers (top of atmosphere)
- BOA Digital Numbers (bottom of atmosphere / surface reflectance)
- TOA Radiance (top of atmosphere)
- TOA Reflectance (top of atmosphere)
- BOA Reflectance (bottom of atmosphere / surface reflectance)

The choice of spectral target may modify the number of spectral bands available.

Secondly, choose a band combination. There are about 20 standard combinations, but how many you see in the list will be determined by what spectral bands you have available in your scene. Regardless of what band combination you choose, you can always manually change the band assignments (both in this dialog, and later in the algorithm editor interface). ProRaster currently offers the following band combinations.

Natural Color	Red, Green, Blue	
Natural Color with Atmospheric Removal	SWIR2, NIR, Green	
Natural Color with Atmospheric Removal High	SWIR1, NIR, Green	
False Color	Blue, Green, Red	
False Color (Urban)	SWIR2, SWIR1, Red	
False Color Infrared (Vegetation)	NIR, Red, Green	
False Color (Vegetation and Water)	NIR, SWIR2, Coastal	
False Color (Vegatative Analysis)	NIR, RDEG2, Red (Landsat 1-5 MSS)	
Shortwave Infrared	SWIR2, NIR, Red	
Short-Wave Infrared Narrow	SWIR2, NWIR, Red	
Color Infrared	RDEG2, Red, Green	(Landsat 1-5 MSS)
Agriculture	SWIR1, NIR, Blue	
Healthy Vegetation	NIR, SWIR1, Blue	
Vegetation Analysis	SWIR1, NIR, Red	

Geology	SWIR2, SWIR1, Blue
Bare Earth	SWIR1, Green, Blue
Fire Scar	SWIR2, NIR, Blue
Land/Water	NIR, SWIR1, Red
Bathymetric	Red, Green, Coastal
Atmospheric Penetration	SWIR2, SWIR1, NIR
Atmospheric Penetration B	SWIR2, SWIR1, NWIR
RGB(NIR,SWIR1,SWIR2)	NIR, SWIR1, SWIR2

Select the desired band combination, and the appropriate spectral bands will be chosen for the Red, Green, and Blue components. You can modify these as you see fit. See the next section, "Using the RGB Combinations Editor", to see how to add new band combinations of your own design. You can also use this editor to set the default data transforms for the Red, Green, Blue, and Pan components.

If the scene contains a Panchromatic band, then you will be able to turn on pan-sharpening. The appropriate pan spectral band is chosen.

If the scene contains a Mask field, then you can turn on raster masking and choose a mask band. ProRaster may have prepared one or more of the following mask bands.

Pixels (Clear)	Unobstructed pixels retained
Pixels (+ largely clear)	Largely unobstructed pixels retained
Pixels (+ partly clear)	Partly obstructed pixels retained
Land (Clear)	Unobstructed pixels classified as Land retained
Water (Clear)	Unobstructed pixels classified as Water retained
Snow (Clear)	Unobstructed pixels classified as Snow (or Ice) retained
Extant	All pixels containing data

Hit OK and the algorithm will be constructed and opened in the main ProRaster interface. If you wish to keep the algorithm, use the normal mechanism to save it to an MRD file.

An alternative option may also be presented. If you have a rendering algorithm open in the main ProRaster interface, then you will be presented with an option to "Render scene to layer". This will take you through the same procedure as "Render Scene", except that the scene will be rendered in a new layer that is added to the existing open rendering algorithm. Using this technique, you can build a multi-layered algorithm, one scene per layer, in which you can control the color data transform for each scene individually.

When you render any spectral data via this method, metadata is written into the algorithm listing all the valid band combinations for the chosen spectral field. This information is reflected on the Layer property page and from here you can change the RGB band combination without having to go back to the scene browser or product editor.

If an algorithm is open (and it may be empty), you will see an option to "Use in Algorithm". This will add the scene assembly virtual raster on the selected row(s) to the raster source collection for that algorithm. You can then easily refer to it when adding layers and configuring components.

You will also see an option "Tile ID to clipboard". This simply copies the Tile ID for the selected row(s) to the clipboard. You can use these Tile ID's in the Create Product Sequence wizard. They can be pasted directly into the appropriate form in the wizard, or into a temporary text file.

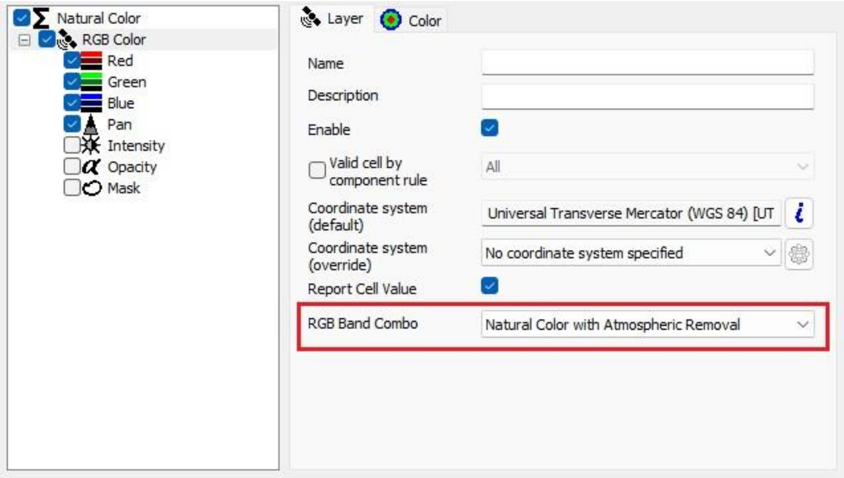

The “Export scene” option will take you to the Export processing operation dialog. For a description of this dialog, see the Processing menu help section. You can use this operation to convert the scene key virtual raster into a crystallised raster in either MRR format, or a transferrable raster format. In the latter case, each band (of each field and event) is output to a separate raster and all the rasters are placed into a structured folder. You can modify the output by selecting the resolution level to target, choosing to export all events or just the latest event, and selecting the fields and bands to export.

The “Delete Scene” option will delete all the selected scenes from the scene database. This removes the entry from the database and deletes the scene assembly virtual raster associated with each scene. It does not delete or in any way modify the original scene data.

Using the RGB Combinations Editor

From the “Browse Scenes” menu button in the main interface, select the “RGB Combinations Editor” menu item to open the RGB Combinations Editor dialog. In this dialog you can –

- Add new Red-Green-Blue band combinations.
- Delete band combinations you have defined previously.
- Change the default data transform for the Red, Green, Blue, and Pan components, when the system creates a rendering algorithm using an RGB Layer for a spectral raster.

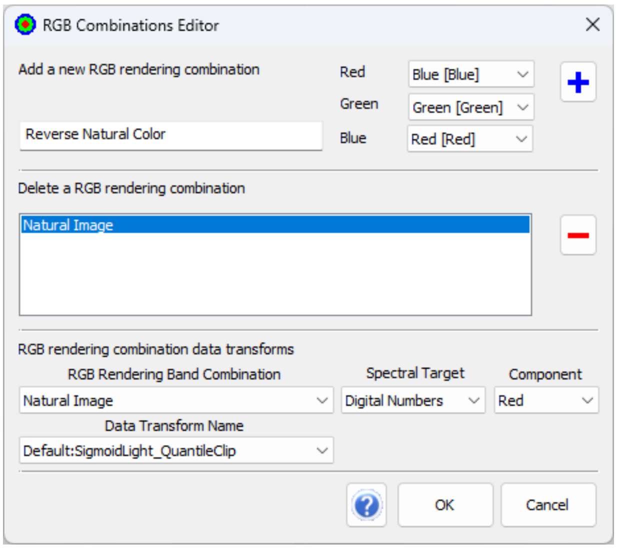

In the top section of the dialog, you can add new band combinations. Give your new band combination a unique name and then associated the desired spectral bands with each of the Red, Green, and Blue components. Once you have defined all four required parameters, hit the "+" button to create the new band combination.

In the central section of the dialog, you can delete any band combinations that you have defined previously. The system does not allow you to delete system defined band combinations, but you can delete any band combinations you have manually defined.

Your new band combination will be visible whenever you render a spectral raster from the Multispectral Scene Database browser dialog or the Multispectral Product Database editor dialog. To be visible, the spectral dataset must contain the required spectral bands used in the combination.

In the bottom section of the dialog, you can change the default data transform that is used in the rendering algorithms you generate for spectral rasters. In a rendering algorithm for spectral data, you must define a data transform for spectral data values to the system primary color tables, for the Red, Green, Blue and Pan components (if defined). For every band combination an appropriate default transform will be assigned, but you can reassign these to a different transform of your own design if you wish.

Firstly, choose the RGB band combination that you want to edit. You can edit both the new combinations you have added via this dialog, or the system combinations. Then, select the spectral target. Provision is made to define different transforms for spectral digital numbers, radiance, and reflectance, because the data range for each of these kinds of data are likely to be different. Then select the component you want to edit.

When you select a component, the name of the data transform is displayed. The list will contain the default system defined transforms as well as any data transforms that you have added and defined manually using the Data Transform Editor. It follows that you ought to have created an appropriate data transform (or multiple) in the Data Transform Editor first. In most cases you will assign the same data transform to all four components, but you can assign a different transform to each component if you wish.

Hit "OK" to save the changes you have made in this dialog. These settings are stored in your "user hidden data application" resources file. If you hit "Cancel" to exit the dialog, all the changes you have made are discarded.

Scene Quality (unusable scenes)

When a scene is imported it is automatically assigned a quality index based on the amount of cloud in the scene. In the scene database browser spreadsheet, scenes with poor quality will be displayed with a light red background. A special quality index is called "unusable". In the scene browser you can tag selected scenes as "unusable" or untag them. This tag is used when creating products from scenes – either in the scene database browser or via the Create Product Sequence Wizard. Scenes tagged "unusable" will not be included in any products made via these methods.

Creating a product

The main purpose of the scene database browser dialog is to enable you to select one or more scenes from which to create a multispectral product. Building a product takes you on a road towards data analysis of the scene(s) . The product is written into a multispectral product database and can be edited using the multispectral product editor dialog.

From the scene database browser, you can make five types of products directly.

Scene

The simplest product, consisting of a single scene. The raster input to the product is just the key MVR virtual raster associated with the scene.

Select any scene (row) in the database and right click to create a scene product.

Composited Scene

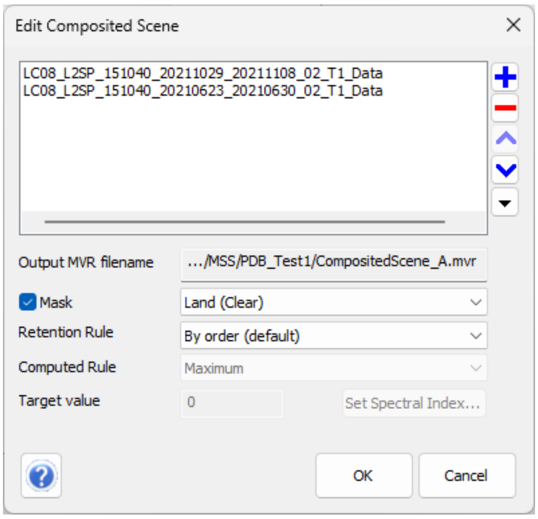

A single scene, constructed by merging two or more nearly contemporaneous visits of the same tile. You usually specify a raster mask to differentiate between high-quality pixels (which are retained) and low-quality pixels (which are discarded). Low-quality pixels discarded from one scene may be replaced by high-quality pixels from another scene in the stack. The raster input to this product is a new MVR that combines the input scenes using a raster mask operation.

Sort the database by tile and then by acquisition time. Select two or more scenes that are spectrally matching (same satellite family, same processing and processing level, same product, same tile) and as contemporaneous as possible. Right click to create a composited scene product.

We provide three general methods for compositing. The first method is masking and filling. Each scene in the composite is masked to remove low-quality pixels (usually cloud affected). This stack of masked scenes is then merged. Working from the bottom to the top, valid pixels are retained and replace any existing valid pixels. Scenes at the top of the list will overprint scenes lower in the list. Use the sorting options to order the list appropriately. Choose “By order” as the retention rule.

The second method is statistical. We assign or compute a value for each pixel by statistical analysis of the pixels from all the contributing scenes. You can assign the minimum or maximum or compute the mean or median. The order of the scenes is not important. You do not have to apply a mask although it is still advisable. Choose “Minimum”, “Maximum”, “Mean”, or “Median” as the retention rule.

The third method is assignment by computed value. For each scene we supply a computed raster derived from the data in that scene. We assign a value from the scene in which the computed value is the minimum, the maximum, or the closest to an expected value. Currently, we allow you to supply a Spectral Index as the computed raster. Choose “By computed value” as the retention rule. Select the spectral index to compute. Set the computed rule appropriately and the expected value if needed.

Scene Sequence

A single tile, with multiple events. Combine two or more visits of the same tile by declaring each visit to be an event in the raster. This provides a new dimension to the raster dataset – a time dimension. It can be imagined as frames in a movie. The raster input to this product is a new MVR that combines the input scenes and declares the events.

Sort the database by tile and then by acquisition time. Select two or more scenes that are spectrally matching (same satellite family, same processing and processing level, same product, same tile) and at the acquisition times that you want to target. Right click to create a scene sequence product.

Mosaic

Two or more adjacent scenes, merged to produce a single raster dataset. Where tiles overlap, the order of the scenes in the mosaic determines which data will be retained in areas of overlap. The raster input to this product is a new MVR that merges the input scenes.

Sort the database spatially. Select two or more scenes that are spectrally matching (same satellite family, same processing and processing level, same product, same tile) and as contemporaneous as possible. Right click to create a mosaic product.

Composited Mosaic

A mosaic, where one or more of the scenes is composited from two or more scenes. The raster input to this product is a new MVR that combines and merges the input scenes using a raster mask operation.

Sort the database spatially. Select two or more scenes that are spectrally matching (same satellite family, same processing and processing level, same product, same tile) and as contemporaneous as possible. Right click to create a composited mosaic product.

Sometimes, when you select a group of scenes, you are given an option to create a [Partial] product. This indicates that it can make a product of the indicated type from the selected scenes, if one or more of the scenes is not included. If you choose this option, the scenes that are not suitable for inclusion in the product will be excluded.

The options for compositing scenes in the mosaic are the same as for a composited scene product.

The Multispectral Product Database

You interact with product databases via the “Browse Products” menu button. For a demo see YouTube here.

Drop down the menu button to reveal the options for creating and deleting product databases.

Creating a product database

Create a new product database by selecting “Create new product database”. You need to create and name your new product database. Your database must have a unique name. The product database file will use this name. If the name is unique then you will see a green tick. Hit OK to create the product database.

You can create as many product databases as you want. How you organise your products is up to you. You may want to add all products to a single large database, or you may want to create a database for each project that you undertake.

- A product database file has a “<your name>.msspdb” file name. You must specify the name when you create the database. You cannot change the name thereafter.
- Every product database you create will be located in the “C:\Users\<your user name>\RGE\MSS” folder on your main drive. You cannot move this file.
- When you create a product database it creates a new folder in the “C:\Users\<your user name>\RGE\MSS” folder of the same name. For example, if I create a new product database called “PDB_KatiThanda.msspdb” then a new folder “C:\Users\<your user name>\RGE\MSS\PDB_KatiThanda” will be created. If you have matching scene and product databases for projects, it is a good idea to differentiate them with some naming prefix like SDB_ and PDB_.
- The new folder will contain all the MVR virtual raster files and all other files that are created for every product in the database.
- When you delete a database both the .msspdb file and the folder and all its contents will be deleted.

You can delete individual products from a product database (in the product browser) or delete the entire database. In the first case, all of the output files generated by the product will be moved to the recycle bin. In the second case, both the .msspdb file and the entire output folder will be moved to the recycle bin.

Browsing a product database

Click on the Browse Products button to open a browser for a product database. If you have more than one product database, then a dialog will open allowing you to choose which product database you wish to browse. Double click on the database in the list or select the database and hit OK.

The product database browser is very much like the scene database browser. It contains a world map showing the bounding box of each product in the database and a spreadsheet containing the products.

The primary purpose of the product browser is to select a product for editing. Double click on a product to open the multispectral product editor for that product, or right click and select “Edit Product”.

Right click on a product and select “Delete Product” to delete that product from the database. The product will be removed from the database. Optionally, all the output files generated by the product can be moved to the recycle bin.

Creating a product

In the same way that you can select multiple scenes in the scene browser and create a new product from those scenes, so too can you select multiple products in the product browser and create a new product from those products. The new product is written into a multispectral product database and can be edited using the multispectral product editor dialog.

From the product database browser, you can make four types of advanced products.

Scene Sequence

This is an alternative way to produce a scene sequence product. You can also do this from the scene browser where you select the scenes directly. The selected products must be spectrally compatible. Select two or more scene products in the database and right click to create a scene sequence product.

Composited Scene Sequence

A composited scene sequence brings multiple composited scene products together into a sequence. The selected products must be spectrally compatible. It is also likely that each composited scene product will use the same mask band. Select two or more composited scene products in the database and right click to create a composited scene sequence.

Spectral compatibility means each scene has the same Platform, Family, Processing, Processing Level, and Product. It not only ensures that data is comparable across scenes, but it also ensures that the field and band structure of the rasters is the same. This is a critical issue for virtual raster construction. Here is an example for Landsat.

Platform	Landsat
Family	Landsat 8 & 9
Processing	Collection 2
Processing Level	L2SP (Level 2, Surface reflectance and temperature)
Product	Level 2

Spectral compatibility for Sentinel 2 is easier to achieve as you only need to make sure that all your scenes are either Level 1C or Level 2A.

Consider an example using Landsat 8 & 9. Each bird revisits each tile every sixteen days, and the orbits are staggered so that each tile on the Earth is revisited by a bird every eight days, or four times per month. Each individual scene is likely to be affected by cloud cover, so to build a more complete dataset with the most high-quality pixel coverage, you choose to composite the four scenes for each month into a composited scene. You sort the scene based on proximity to the middle of the month, so each composited scene is the best representation of the Earth in that tile on the 15^{th} of the month. Having started with four data values for each pixel (one per scene), you now have only one data value per pixel.

Over a period of a year for this tile, you build twelve composited scenes, one for each month of the year. You can then bring those together into a sequence, allowing you to visualise the temporal changes easily and to perform the same data analysis on each composited scene in the sequence. Ultimately, you will compute statistics for this dataset and build up a quantitative picture of how various measured parameters have changed over the course of that year.

Mosaic Sequence

A mosaic sequence brings multiple mosaic products together into a sequence. The selected products must be spectrally compatible. Select two or more mosaic products in the database and right click to create a mosaic sequence.

When you are combining tiles into a mosaic, it is a more difficult challenge to acquire scenes that are approximately contemporaneous. Tiles in a north-south direction lie along the bird's orbit and are collected in sequence, but tiles to the east and west (which lie on a different orbital path) may be surveyed days earlier or later. In the case of the Landsat 8 & 9 family and the Sentinel 2 family, tiles on different orbital paths may be separated by up to eight days.

It is likely that your mosaic will be affected by cloud cover and that this may vary substantially from one tile to the next. If you plan to use the mosaic for data analysis, you may wish to mask the cloud out in the product editor, which will reduce the coverage of your dataset.

Composited Mosaic Sequence

A composited mosaic sequence brings multiple composited mosaic products together into a sequence. The selected products must be spectrally compatible. It is also likely that each composited scene product will use the same mask band. Select two or more composited mosaic products in the database and right click to create a composited mosaic sequence.

To improve the quality of the data in a mosaic, you can combine repeat observations of one, or more, or all tiles in the mosaic. For example, if you were to gather all the scenes collected over a period of a month for a mosaic of tiles, then you can composite them and collapse the dataset to a single value per pixel, with a higher percentage of quality data. This is a composited mosaic. If you were to do this for every month in a year, you can then bring those twelve composited mosaic datasets together into a sequence.

If an algorithm is open (and it may be empty), you will see an option to "Use in Algorithm". This will add the product virtual raster on the selected row(s) to the raster source collection for that algorithm. You can then easily refer to it when adding layers and configuring components.

You will also see an option "Tile ID to clipboard". This simply copies the Tile ID for all the scenes in the product in the selected row(s) to the clipboard. You can use these Tile ID's in the Create Product Sequence wizard. They can be pasted directly into the appropriate form in the wizard, or into a temporary text file.

The Multispectral Product Editor

Access the product editor dialog by creating a new product in the scene browser, or by editing a product in the product browser. For a demo see YouTube here.

The editor allows you to build new data products by applying processing operations to the original spectral raster product. As you add operations, you build a waterfall style cascade where the output from an operation is dependent on all the operations up the waterfall. If you update or modify an operation, then all the operations downstream of that modified operation will be updated.

On the left of the dialog is a tree control which represents the operation waterfall. On the right is a property sheet that changes content depending on what operation you select in the tree control.

Every item in the tree control is an operation. An operation has inputs (usually a spectral raster) and outputs (usually a raster). It applies an operation to the inputs to generate the outputs. In almost all cases, both the inputs and output will be virtual rasters. All these virtual rasters will be stored in the product database folder, and they will have an automated file naming convention. In general, you are not able to modify the naming of inputs or outputs.

All operations have a common property page that show you the name and description of the operation. You can also see an XML script which is used to generate the operation and lists all the properties of the operation including the inputs and the outputs. Hit the “Modify and Run Operation” button, or simply double click on the operation item in the tree control, to open a dialog to edit the operation properties. If you modify these properties and hit OK, then all downstream operations in the product will be updated.

Most operations have a preview property page which displays a map showing a default rendering of the product up to and including the application of that operation. You can use the mouse wheel to zoom this map. Click, hold, and drag the mouse left button to pan the map.

If you right click on an operation in the tree control, you will see a context menu which might display the following options. The actual number of options that are offered can vary according to the operation selected.

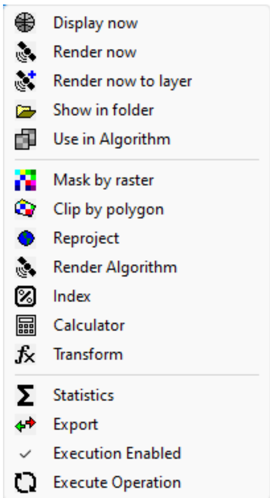

“Display now” sends the output of the selected operation to the ProRaster main interface and is displayed in a default algorithm. A Look-up Table color layer will be used so you can explore each of the fields, bands, and events in the output raster. If you want to keep the algorithm you must save it in the normal way.

“Render now” sends the output of the selected operation to the ProRaster main interface and is displayed in a default algorithm. An RGB color layer will be used, and you can choose a band combination to assign spectral bands to rendering components. If you want to keep the algorithm you must save it in the normal way.

“Render now to layer” will be offered if you have a rendering algorithm open in the main ProRaster interface. In this case, the product is rendered in a new layer that is added to the existing algorithm. In all other respects, this option works the same as the “Render now” option.

“Show in folder” opens a Windows Explorer application and navigates it to the folder that contains the output files from the operation.

If an algorithm is open (and it may be empty), you will see an option to “Use in Algorithm”. This will add the product virtual raster to the raster source collection for that algorithm. You can then easily refer to it when adding layers and configuring components.

“Mask by raster” applies a mask to the input. You will probably use one of the mask bands in the original scene raster data.

“Clip by polygon” clips the input to a complex polygon set.

“Reproject” changes the coordinate system of the raster. At the same time you should change the cell size to a suitable cell size in the new coordinate system. You can also realign the cell boundaries.

“Render Algorithm” generates a spectral rendering algorithm, just like the “Render now” option, but saves it to the product database folder. It can be used as input to other products.

“Index” computes one or more spectral indices. The input must be spectral, and the output will be non-spectral continuous data.

"Difference" computes the difference between events in a raster, either between each event (a moving difference) or between a named key event and each other event (a fixed difference).

"Calculator" applies a calculator expression to the input, incorporating other raster data as required. The output will be non-spectral continuous data.

"Transform" applies a non-linear data transformation to a single selected band.

"Statistics" computes statistics for the input raster dataset. You can compute statistics for any selection of fields and bands and events. Once computed, you will be able to browse the statistics in the property sheet. There is no raster output from this operation.

"Export" exports the input virtual raster to a crystallised raster in either MRR format or a transportable raster format.

"Execution Enabled / Execution Paused" refers to the execution of statistics and export operations. When you edit an operation, the system will execute that operation and then recursively execute all child operations in the cascade. Most operations simply write virtual rasters and take little time, but statistics and export operations may take considerable time to complete. You can toggle this option to pause the execution of these kinds of operations. Later, you will need to trigger the execution of these operations manually.

"Execute Operation" will execute the selected processing operation and then recursively execute all child operations in the cascade. If you have paused execution of statistics and export operations, or the product is out of data for any other reason, then you can use this option to manually trigger execution.

Editing a product

At the top of the tree is the original product (Scene, Composited Scene, etc…). The product was created by choosing one or more scenes or one or more products and, possibly, defining other properties like a mask band. You can edit this item and change the scene or product, or collection of scenes or products, that were used to generate the product.

If you edit a Scene product, you will be presented with the scene browser dialog and can select a replacement scene. In all other cases, you will be presented with a dialog that shows the list of scenes in the product and allows you to add new scenes, delete existing scenes, and reorder the scenes. An example is shown below for a Composited Scene.

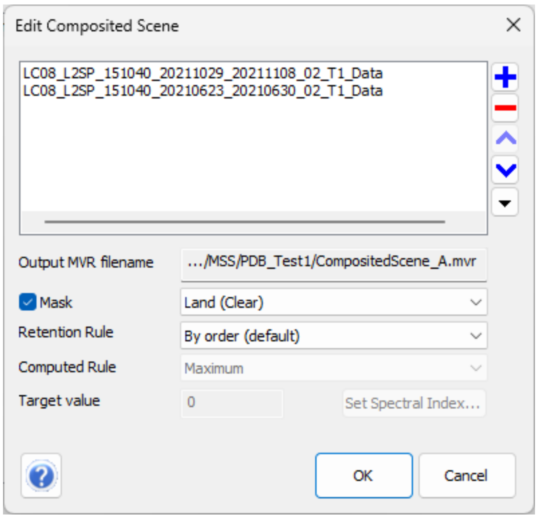

The scenes are listed from "top" to "bottom". In most products, scenes are composited using a painter's algorithm where data from scenes underneath are overprinted by data from scenes on top. The scene list reflects this order. When you add a scene to the list it will open the scene database browser and you will be able to select a scene.

You can reorder the scenes manually by selecting a scene and clicking the move up or move down buttons. A more effective way to order the scenes is to drop down the order menu button. This provides the following options.

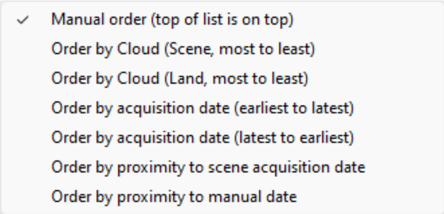

In almost all cases it is advantageous to order by cloud coverage. You can use cloud coverage for the entire scene, or cloud coverage for the land classified pixels in the scene, if this information is available. The scenes with the least cloud coverage will be placed at the top so that this higher quality data overprints other scenes that are compromised by higher levels of cloud coverage.

Alternatively, you can order by acquisition date. Usually, you would order the scenes from earliest (oldest) to latest (most recent). In a scene sequence product this order is enforced. In a composited scene or composited mosaic, you are combining multiple masked nearly contemporaneous scenes. You may want to sort the scenes by proximity to a central (ideal) scene or a specified date. In other words, you are trying to use as much data that is as close to the ideal acquisition date as possible. To sort the scenes so that they are closest to a specific scene, select the scene in the list then choose the appropriate sorting option. The manual option will present a dialog for you to enter the target date and time.

There are restrictions on what scenes you can choose to add to any existing product. The system may force you to select a scene that is the same tile as the original scene. It may force you select scenes that are spectrally compatible with other scenes in the product. If you choose a scene that is spectrally different, then the downstream processing operations may need to be edited to correct problems that arise from changing the field and band structure of the raster.

If you edit a product that draws from multiple source products (like a composited scene sequence for example), you will be presented with a dialog that shows the list of products in the product and allows you to add new products and delete existing products. An example is shown below for a Mosaic Sequence.

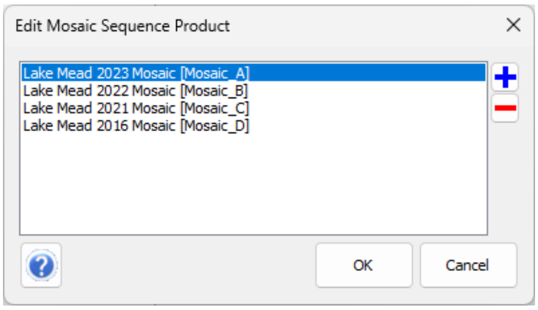

Hit the “Add” button to add a new product to the sequence. You will be prompted to select a product database, and then to browse for the product. Select the product in the spreadsheet and hit OK.

Hit the “Delete” button to remove a product from the sequence. This does not delete the original product.

All the products you edit with this dialog will be sequences and so the order of the products is fixed and determined by the product event time. The products in the list will automatically be sorted from latest time to earliest time (top to bottom).

Mask

When you add or edit a Mask operation you will see the following dialog.

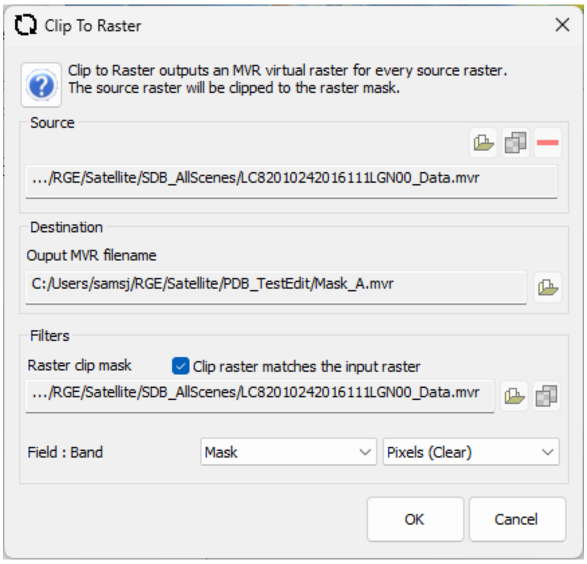

Both the input raster filename and the output raster filename are fixed and not editable. The system will define these parameters. All you need to do is to specify the masking raster, field, and band. Usually, this mask raster will be the input raster. Pixels are masked where the mask raster cells are invalid or are valid and zero.

Clip to Polygon

When you add or edit a Clip to Polygon operation you will see the following dialog.

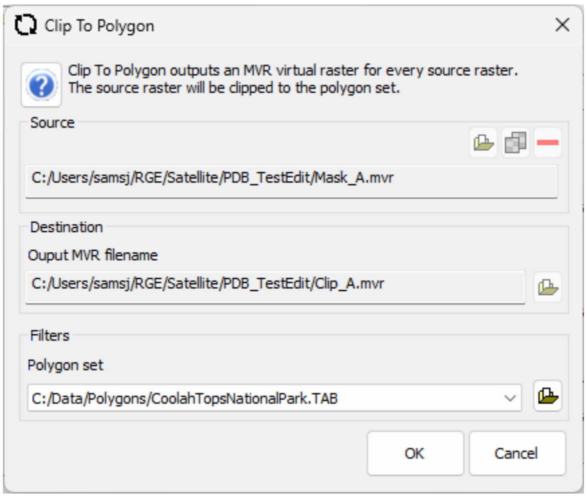

Both the input raster filename and the output raster filename are fixed and not editable. The system will define these parameters. All you need to do is to specify the polygon to which the raster will be clipped. You can select from a list of recently accessed polygon files, or browse to a new file. The polygon file must be in a MapInfo Pro TABLE format.

Reproject

When you add or edit a Reproject operation you will see the following dialog.

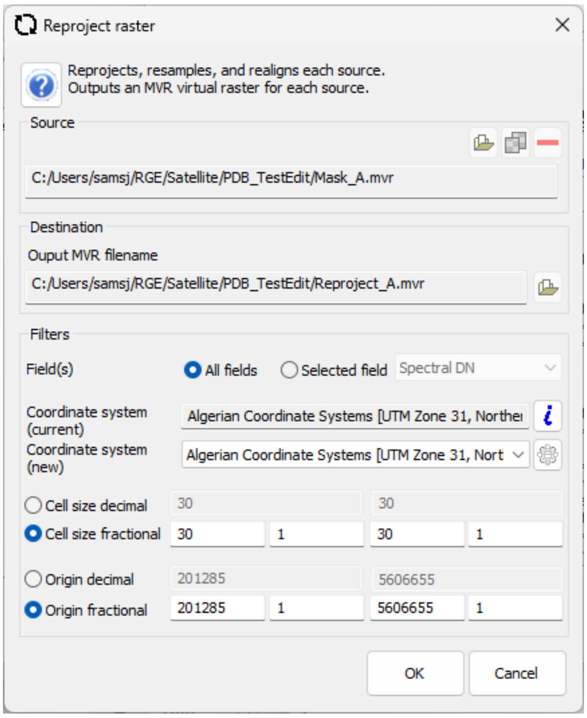

Both the input raster filename and the output raster filename are fixed and not editable. The system will define these parameters. The current coordinate system is shown, and you can select a new coordinate system from a list of recently used coordinate systems, or by opening the projection explorer dialog.

When you change coordinate system you should also consider changing the cell size to values that are equivalent in the horizontal units of the new coordinate system. For example, if the raster was 30 meter resolution and the new coordinate system is geographic, then change the cell size to one arc second (1/3600 of a degree).

You can also change the cell alignment. This is less critical. Regardless of what number you choose, you will only ever be able to shift the edges of the cells by up to half a cell width.

Render Algorithm

When you add or edit a Render Algorithm operation you see the familiar spectral band combination selection dialog.

Index

When you add or edit an Index operation you will see the following dialog.

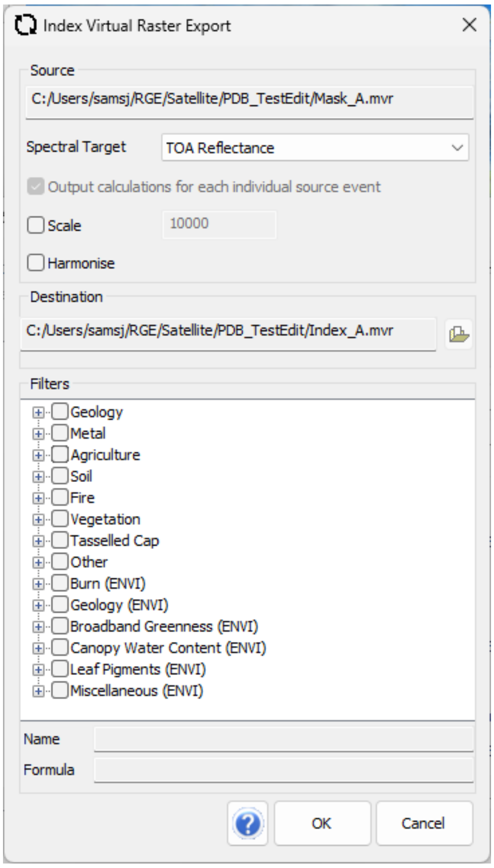

Both the input raster filename and the output raster filename are fixed and not editable. The system will define these parameters.

You need to choose the spectral target. By default, ProRaster chooses the highest value spectral target.

You can apply some conditioning to the spectral band data prior to consuming it in the index formula. The "Scale" option allows you to divide by a scaling factor, usually to convert the values to numbers between 0 and 1. The "Harmonise" option will clamp all input values to within 0 and 1 (meaning values <0 will be set to 0 and values >1 will be set to 1).

Finally, select the index, or the group of indices, that you wish to compute. There are a broad range of index formulas, grouped by their application. Most of these formulas have been acquired from the IDB Index Database (https://www.indexdatabase.de/). Others, where indicated, have been acquired from the ENVI online documentation. Many of these are duplicates. Click on the index in the tree to see its name, description, and formula.

A selection of commonly used indices is provided under the "Common" grouping. You can choose one of these indices by checking the box and then selecting an index from the drop list. Otherwise, to choose multiple indices or any of a much broader range of indices, uncheck the box and select the indices you want to compute in the tree list.

There is no immediate penalty for computing multiple indices, rather than just one. For example, if you are looking for an Index that accurately maps a fire scar, you can select all the Fire and Burn related indices and then display them in ProRaster and make a visual comparison. By doing so, you may discover that one index works better for your study area than the others. Thereafter you can focus on that index. The penalty for choosing many indices only comes at the end of the processing chain when you compute statistics or export the raster. At this point, the system must do all the work to evaluate all the indices.

If you choose a single index and choose to display that index in the algorithm rendering engine, ProRaster Scientific will match the index name (for example NDVI or NBRT1) against the names of available color table maps. If a match is found, it will be used. Appropriate color table maps are provided for all the common spectral indices. These color table maps provide both a color table and a data transform and enable you to render a spectral index consistently.

You can provide your own color table maps for spectral indices or override the existing maps. The system will search for appropriately named color table maps in the UserTables folder, any externally linked folders, and finally the SystemMaps folder, in that order. It will use the first matching table it finds. The standard tables are shipped in the SystemMaps folder, so you can override one of these tables by placing a table of the same name in the UserTables folder or any externally linked folder.

Difference

When you add or edit a Difference operation you will see the following dialog.

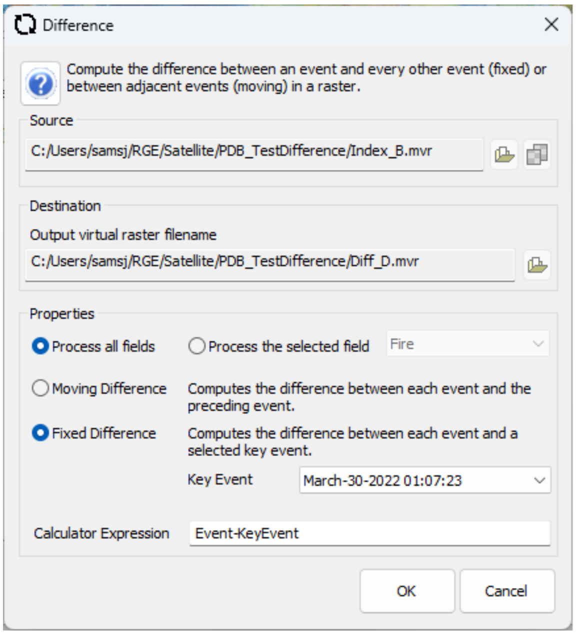

Both the input raster filename and the output raster filename are fixed and not editable. The system will define these parameters. The input raster must contain two or more temporal events (a sequence) and the output raster will contain either the same number of events, or one less.

You can compute the difference for all fields and bands in the input raster, or for a selected field and all bands in that field. The operation will only process fields that are Continuous. Image, Image Palette, and Classified fields will not be processed.

You can compute a moving difference, which is the difference between each event and the preceding event. The first event generated in the output raster will correspond to the second event in the input raster.

You can compute a fixed difference, which is the difference between each event and a key event that you select. Usually, the key event will be the first event in the sequence.

The difference operation is executed by calculator operations defined in the output virtual raster. You can edit the calculator expression if you wish. The standard expression is –

“Event – KeyEvent”

In this expression, “Event” and “KeyEvent” are variables that refer to the current event and the key event (or preceding event). These variable names cannot be changed, but you can change the expression. For example, you can reverse the sign of the difference by using –

“KeyEvent – Event”

You can compute the absolute value using –

“abs(Event - KeyEvent)”

You can introduce a conditional expression to output only differential values of interest - for example this expression only outputs difference values that are greater than 0.1 –

“cond(Event–KeyEvent>0.1,Event–KeyEvent,null)”

Calculator

When you add or edit a Calculator operation you will see the following dialog.

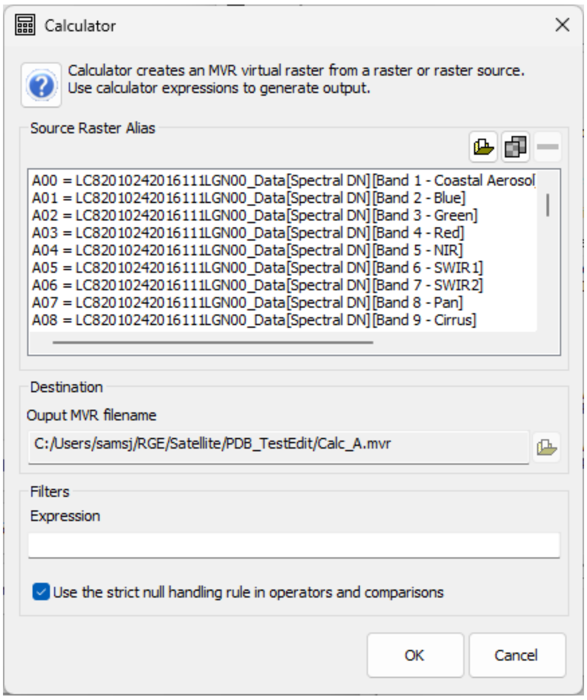

The output raster filename is fixed and not editable. The system will define this parameter. Alias codes are automatically created for all the fields and bands of the input raster. You can add other rasters and raster sources to this and define additional alias codes if you want to incorporate data from other rasters in the calculation.

Your main task is to create or edit the calculator expression, which acts on raster variables (alias codes) to produce a single banded output.

Transform

When you add or edit a Transform operation you will see the following dialog.

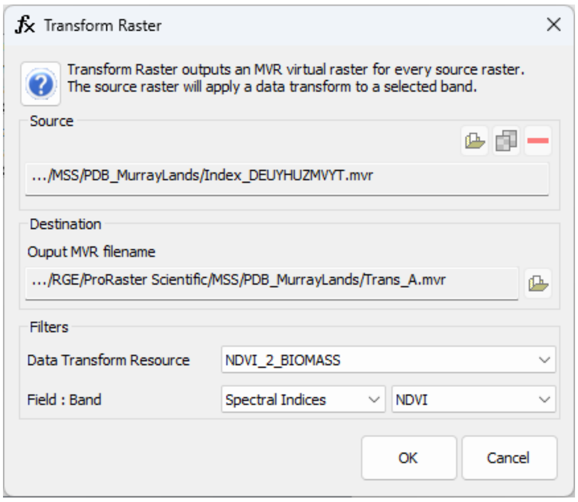

Both the input raster filename and the output virtual raster filename are fixed and not editable. The system will define these parameters.

You need to define a data transform using the Data Transform Editor. You will then be able to select this transform from the list. Select the field and band that the operation will target.

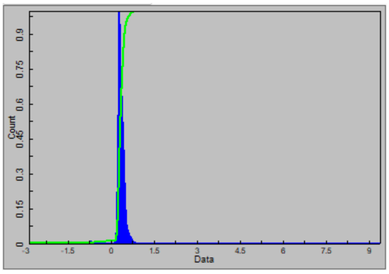

As an example, consider transforming Normalised Difference Vegetation Index (NDVI). The histogram shown to the left is a typical NDVI distribution with the bulk of values between 0 and 1, but with a tail of outliers that have negative values and large positive values. I will use a transform to convert the NDVI values to a proxy measurement of "Feed On Offer (FOO)".

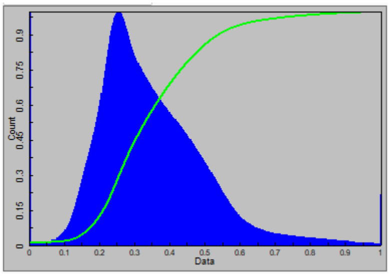

	Value	Index
0	-1000000	0
1	0.1	0
2	0.8	1
3	1000000	1

My first example transforms all the values up to 0.1 to zero and all the values from 0.8 upwards to 1. The data transform is defined as a "Table: Index for value" which applies a linear transform between the data value pairs (input NDVI value, output biomass value). The NDVI values have been clipped and redistributed over the range 0 – 1, a proxy measurement for FOO.

The relationship between NDVI and FOO is unlikely to be linear. It will depend on the ground conditions and will vary between localities. It may be logarithmic, which means that as NDVI increases FOO increases less quickly (low values of NDVI have less feed value). It may be exponential, which means that as NDVI increases FOO increases more quickly (low values of NDVI have more feed value).

In the following example I have used a logarithmic transform that places less value on lower values of NDVI. The histogram is skewed to the left.

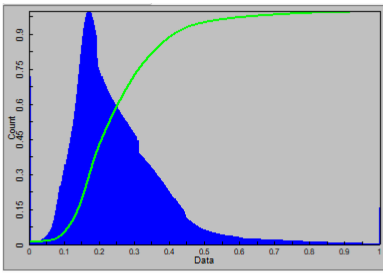

	Value	Index
0	-1000000	0
1	0.1	0
2	0.2	0.091
3	0.3	0.194
4	0.4	0.31
5	0.5	0.444
6	0.6	0.6
7	0.7	0.783
8	0.8	1
9	1000000	1

In the following example I have used an exponential transofrm that places more value on lower values on NDVI. The histogram is skewed to the right. In both cases, the distribution has discontinuities at the transform control points.

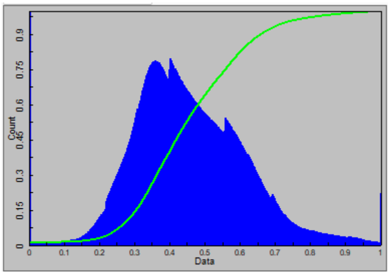

	Value	Index
0	-1000000	0
1	0.1	0
2	0.2	0.217
3	0.3	0.4
4	0.4	0.556
5	0.5	0.69
6	0.6	0.801
7	0.7	0.909
8	0.8	1
9	1000000	1

Compute statistics for the transformed raster, then use the WeightedSum measure to sum these values to acquire a proxy for FOO across the entire raster. Convert to average per cell by dividing by the number of valid samples. Assuming the cell area is in square metres, multiply by (10000/CellArea) to convert to average proxy FOO per hectare.

Statistics

When you add or edit a Statistics operation you will see the following dialog.

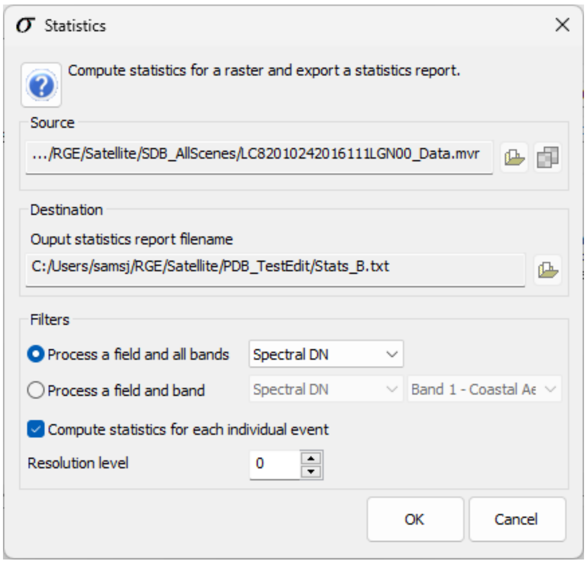

Both the input raster filename and the output statistics filename are fixed and not editable. The system will define these parameters.

All you need to do is to select the field to process, or the specific field and band. Usually, you will compute statistics for each event in the raster if multiple events are present. Each field-event-band combination you target will generate a statistics dataset. You can specify the resolution level to target (0 is the base resolution). It may be advantageous to target a higher resolution initially to test the operation, but generally you will compute statistics at the base resolution.

When you hit OK on the statistics dialog, a progress dialog will open showing the progress of the operation as it computes statistics. Remember, that this operation forces the entire virtual raster operation chain to be evaluated and draws data from all the way back to the original scene raster files. This is where your computer needs to do some real work!

The operation will generate three output files with different extensions. The *.STTX file is a binary file containing the statistics data. This is only used by ProRaster. The *.TXT file is a text file contain a complete report of each statistics dataset that was computed. This is in a readable format, so you can view it in a text editor and browse through the data. The *.CSV file is a text file in a comma delimited format. It contains a header line and then a single line for every computed dataset containing summary statistics. The columns are -

- Field, Band, StartTime, EndTime, FirstEvent, LastEvent
  The zero based field index and zero based band index.
  The time of the first event and last event, in seconds after 00:00, Jan 1 1970 UTC.
  The zero based index of the first and last event.
- FieldName, BandName, StartTimeDate, EndTimeDate
  The name of the field and band.
  The time of the first event and last event, in a readable format.
- Resolution, CellSizeX, CellSizeY, CellArea
  The data pyramid resolution level.
  The width and height of the cells in the raster at the resolution level.
  The area of a cell at the resolution level, in whatever the units of the cell size is, squared.
- Mode
  The statistics calculation mode which can be - No, Count, Summary, Distribution, Spatial.
- SampleCount, ValidSampleCount, InvalidSampleCount
  The total number of samples in the dataset, the number of samples that contain valid data values and the number of samples that contain invalid (null) sample values.
- InvalidNumberCount, EmptySampleCount, NullSampleCount
  The number of invalid samples that contain invalid numbers (like infinity), the number of invalid samples that are designated “empty”, and the number of invalid samples that are designated “null”.
- SampleArea, ValidSampleArea, InvalidSampleArea
  The total area of all samples, the total area of all valid samples, and the total area of all invalid samples. Area is expressed in units of the cell size, squared.
- InvalidNumberArea, EmptySampleArea, NullSampleArea
  The total area of invalid samples that contain invalid numbers, the total area of invalid samples that are designated “empty”, and the total area of invalid samples that are designated “null”.
- Minimum, Maximum, Mean, Sum
  The minimum valid sample value in the dataset, maximum valid sample value, and mean (average) valid sample value. The sum of all values (equivalent to the Mean multiplied by the ValidSampleCount).
- Variance, StandardDeviation, Signal2Noise
  The variance and standard deviation of the distribution of the valid samples in the dataset, measuring the variability of sample values from the mean. Signal2Noise is the mean divided by the standard deviation and measures the ratio of signal to noise in the dataset.
- LowerQuartile, UpperQuartile, InterquartileRange
  The data values at the 25th percentile, 75th percentile and the data range between these values. In other words, 25% of all data is below the LowerQuartile, and 25% of all data is above the UpperQuartile, and 50% of all data is between these values.
- Median, Mode, Skew, Integration

The Median is the value at the centre of the distribution, equivalent to the 50th percentile. The Mode is value that appears most frequently in a dataset, and an approximation can be computed from the distribution histogram. The Skewness of the distribution is approximated by computing 3*(Mean-Median) /Standard Deviation.
The Integration is the area under the curve. It is computed from an equal width histogram as the sum of all the histogram bin data widths multiplied by the sample count in each bin.

- Cell2CellMinimum, Cell2CellMaximum, Cell2CellMean
  These values are an approximate measure of the change in data values between horizontally adjacent samples.
- Cell2CellVariance, Cell2CellStandardDeviation
  These values are an approximate measure of the change in data values between horizontally adjacent samples.

In addition to these three output files, when you compute statistics at the base resolution level (0), the statistics may be permanently stored alongside the source raster. The statistics will be stored in an STTX file as well as a GHX file.

In the Product Editor, when you select a Statistics operation in the operation waterfall, you will see three new property pages – Summary, Spatial, and Histogram. These property pages present the calculated statistics and provide you with a basic mechanism to browse and inspect them.

A statistics operation may compute statistics for individual Field, Events, and Bands. On the top of each page is a drop list showing each statistics dataset available. You can select any dataset from this list or hit the arrow buttons to the right to skip forward and backwards through the list. When you select a dataset, each property page updates to show that dataset.

On the Summary page there are two buttons that allow you to view the *.TXT report file and the *.CSV data file. The report file will be opened in your default text editor and the CSV file will be opened in your default spreadsheet application.

The Summary page shows the summary statistics including the sample counts, summary statistics, and distribution statistics.

The Spatial page shows the cell size and area, the sample count areas, and the cell-to-cell spatial statistics. On this page you can also compute any value for a percentile, or any percentile from a value. Specify what kind of number you are entering (value or percentage) and whether you want a report for the data above or below this value. Enter percentages as values between 0 and 1. If you type in a value, you will see a report like this –

"88.5% percent of data is equal to or higher than the value 1.5. This corresponds to 1540201 cells or an area of 4385939828."

"11.5% percent of data is less than the value 1.5. This corresponds to 163020 cells or an area of 385939828."

If you type in a percentage (for example, 0.5 = 50%) then the report will compute the value corresponding to that percentage and give a similar report.

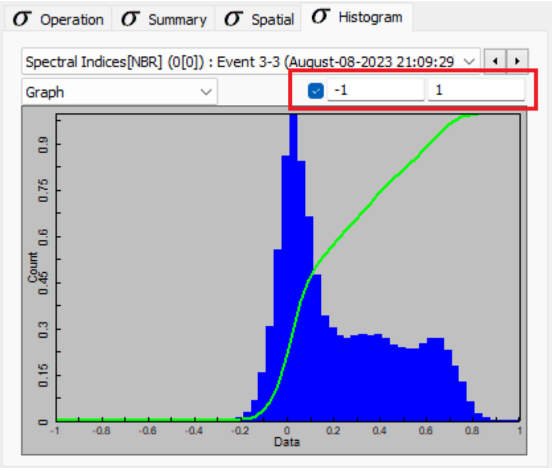

The Histogram page shows the distribution histogram. You can display it as a graph or as a spreadsheet. When graphing, you can apply varying degrees of log scaling to the vertical axis to enhance the detail in the lower frequency bins.

You can also define the minimum and maximum extents of the X axis. This option is useful when you are looking at temporal statistics because, as you flip through the time points, the histogram is displayed over a constant X axis range making it easier to visually compare. To set the axis extent, click the check box at the top of the graph and enter the minimum and maximum axis values (as shown in the image). To return to the full extent, just turn the check box off.

You can apply a data transform to the dataset instantly by selecting a data transform from the “Transformed Sum” combo box. Then, hit the button to copy the data to the clipboard. You can paste the comma delimited data into a spreadsheet for analysis. In addition to some basic information, the operation computes the sum of the transformed values. Normally, the sum of all values is computed as the mean value multiplied by the number of valid samples. In this case, the data distribution histogram is used to compute an approximate sum. It is usually accurate to within about 1-2% of the true value. This is a quick alternative to using a Transform operation (and computing statistics). It can be used to quickly gauge the effect of a data transformation.

Export

The Export operation writes a copy of a raster to a new raster file (or files) in a designated raster format. It is primarily designed to convert (or crystallise) an MVR virtual raster into a raster file.

You may want to export data so that it can be viewed or consumed in another software package, or by another user. Another application is to improve rendering performance of an MVR by converting it to an MRR.

For additional information, see the [export operation](#).

Create Product Sequence Wizard

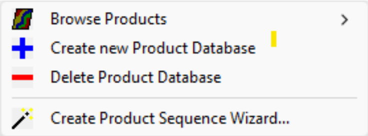

The Create Product Sequence Wizard is designed to turbocharge your productivity and make it easy to create, edit, and recreate temporal sequence products. It automates what would otherwise be a manual product creation and editing process. For a demo visit YouTube [here](#).

Access the Create Product Sequence Wizard from the Product Database drop-menu button.

The goal of the wizard is to make is easy to set up, extend, and modify long baseline studies of spectral indices computed for polygonal regions. You can use the wizard to create the following kinds of products –

- Scene Sequence
  A product with spatial coverage that is limited to a single scene or tile. Multiple acquisitions of this tile are combined in chronological order into a sequence.

- Mosaic Sequence
  A product with spatial coverage that usually covers multiple scenes or tiles. Multiple acquisitions of these tiles are combined in chronological order into a sequence.

- Composited Scene Sequence
  A product with spatial coverage that is limited to a single scene or tile. Multiple acquisitions of this tile are combined in chronological order into a sequence. At each temporal event in the sequence, multiple acquisitions of the tile are combined by masking poor quality pixels and retaining high quality pixels.

- Composited Mosaic Sequence
  A product with spatial coverage that usually covers multiple scenes or tiles. Multiple acquisitions of these tiles are combined in chronological order into a sequence. At each temporal event in the sequence, multiple acquisitions of the tile are combined by masking poor quality pixels and retaining high quality pixels.

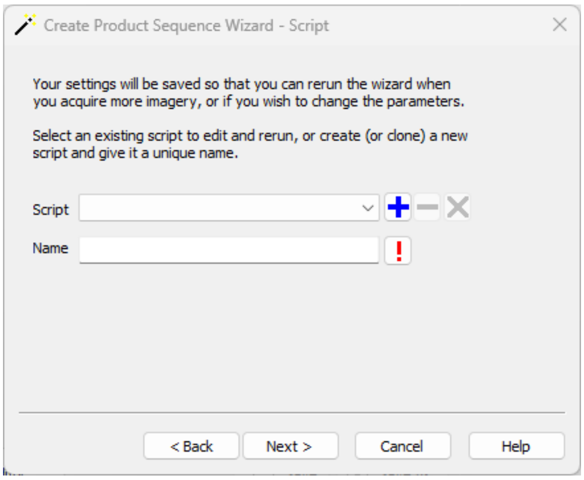

When you run the wizard, your input parameters are saved to a script that you can reuse and edit at any time in the future. You must give each script a unique name. You can create new scripts or edit, clone, and delete scripts from within the wizard. You can also delete all scripts and wizard report files from the MSS Options dialog.

On the "Script" page you need to select the script that you wish to edit. This may be an existing script, a script that you have made a copy of (cloned), or a new script.

A list of existing scripts is displayed in the script drop list. Use the buttons to the right of the drop list to add a new script, delete the selected script, or clone the selected script.

The name of the script is shown below the script drop list. Each script must have a unique name. If the name you have chosen for your script is unique, you will see a green tick. If it is not unique, you will see a red exclamation mark. Edit the name of the script to ensure it is unique and either hit 'return' or click the button to assign the name to the script.

When you run the wizard for the first time, there will be no script or name displayed. Hit the blue button to add a new script, then modify the name as you please.

Depending on your choices on this page, the wizard may execute in one of three modes.

1. Create mode.

   A new script is created with default properties. You will be required to correctly define all properties. When executed, new products are generated.

2. Clone mode.

   A new script is created with properties copied from an existing script. You can change any properties. When executed, new products are generated.

3. Edit mode.

   An existing script is edited. You can change most properties. When executed, intermediate products from the prior execution are deleted and new intermediate products are created. The final sequence product is modified but not recreated. Any processing operations associated with the product are retained.

The inputs to a scene sequence product are a list of scene assembly virtual rasters. There are no intermediate products. The inputs to all other sequence products are a list of intermediate products. For example, if you create a mosaic sequence then your inputs will be mosaic products. These are referred to as intermediate products. The final sequence product is referred to as the final product. The wizard will create both the intermediate and the final products. Note that it is possible to create a final product that is not a sequence.

When you edit an existing script, you are electing to re-run the wizard using modified parameters. The intermediate products that were generated by the earlier run are deleted from the product database and the final product is modified. Any additional processing operations you have previously defined for your final product will be preserved.

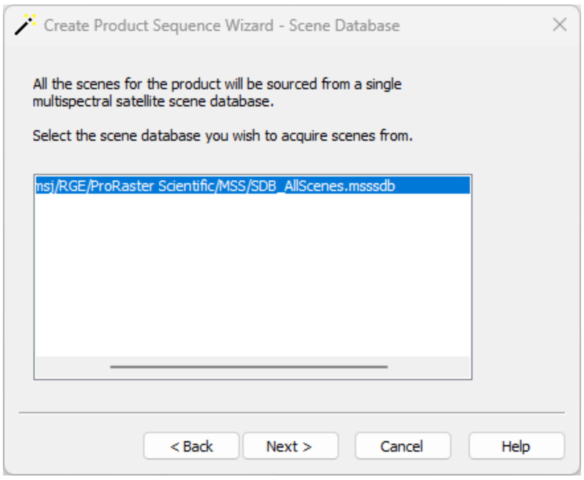

On the "Scene Database" page, select the scene database that will be the source for all scenes in the final product.

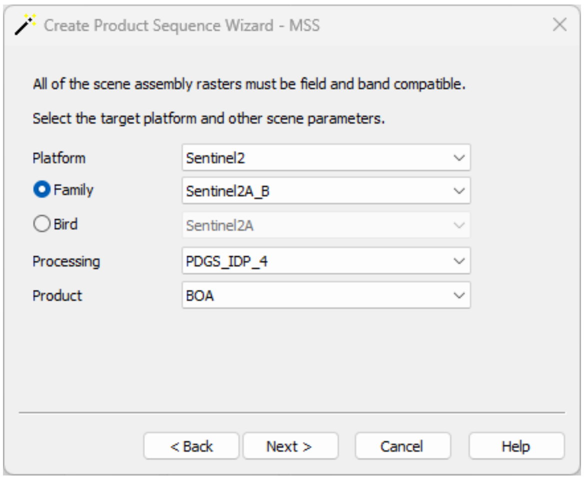

On the "MSS" page you will begin to define a query that will select scenes from the scene database that you have chosen.

You will only be given options that are represented by scenes in the database. For example, if there are no Landsat scenes in the database then no Landsat platform options will be presented. As you make selections, the number of available subordinate selections will become restricted by availability.

Firstly, select the satellite platform. It follows that only scenes from one platform can be used in a product.

Then, specify whether you want to target a specific bird, or all the birds within a family. Birds within a family are known to be compatible and so can be combined into products.

Select the processing level and the spectral product (for example Bottom of Atmosphere (BOA), Top of atmosphere (TOA) or Uncorrected (Raw).

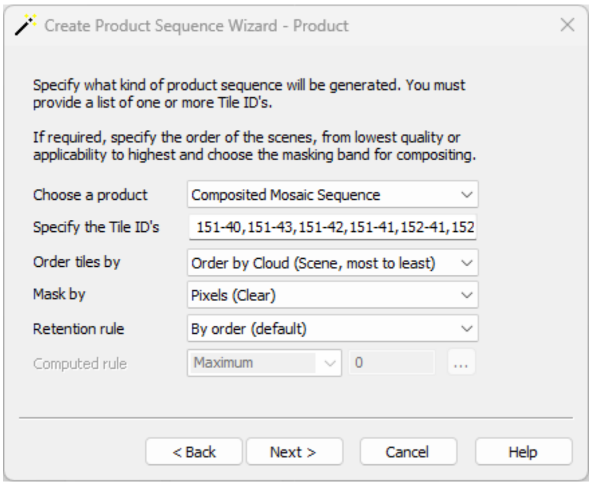

On the "Product" page define what kind of sequence product you want to produce. You can select a Scene Sequence, Mosaic Sequence, Composited Scene Sequence, or Composited Mosaic Sequence.

Regardless of what kind of product you are creating, you must now refine the scene query by declaring what scenes you wish to use. In the wizard, these are referred to as tiles which is a more general term encompassing scenes (a full image capture) and tiles (a clipped partial image capture).

Specify the tiles by entering a string of comma delimited Tile IDs. For example, for Landsat you might have a string like "99-80,99-79,98-80,98-79". In this case the Tile ID is a mashup of the path and row numbers. For Sentinel2 you might have a string like "04QHJ,04QHH,04QGJ,04QGH". In this case the Tile ID matches the Sentinel2 tile delivery system.

You can find the Tile IDs in the multispectral scene database browser spreadsheet. If you select one or more scenes in the browser and then right click, you will be given an option to copy the Tile IDs to the clipboard. It is a good idea to use a text editor to construct your Tile ID string in this way, and then copy the final string into the wizard. You can do the same thing in the multispectral product database browser. This is a more useful option as all the Tile IDs used in mosaics will be copied to the clipboard.

For mosaic products you must also specify a rule that defines the order in which tiles are assigned to each sequence event. For example, in a mosaic you will have multiple tiles in each event that may partially overlap. It is common to order these by cloud coverage, so that the least cloudy tiles are on top and overlap the cloudier tiles. When you are compositing you will likely have multiple revisits of each tile. In this case it is common to sort by temporal proximity to the sequence event date. This ensures that high quality data that is closest to the event data is used in preference to data that is more temporally distant from the event date.

When compositing, define the method by choosing a retention rule. Choose "By Order" for masking and filling, and make sure you define the mask band and scene order. For statistics retention choose minimum, maximum, mean, or median. You do not have to mask and the order of the scenes is not important. Choose "By computed value" to assign by a spectral index.

When compositing, you can select a band to use for masking. This is used to mask out low-quality pixels – for example pixels that are affected by cloud – so that high-quality pixels acquired at a different time can be used instead. If you do not want to mask, choose the "Extant" band.

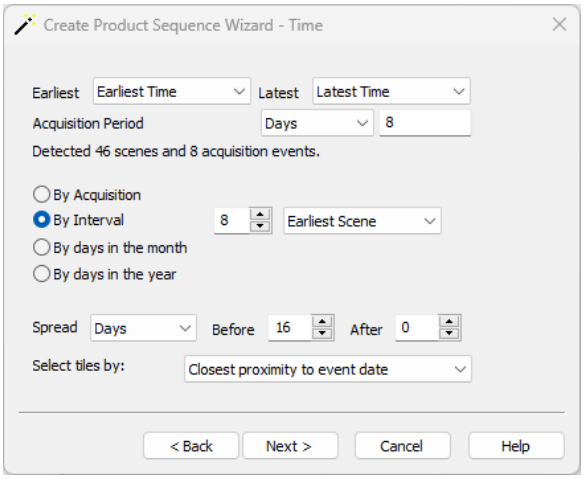

On the "Time" page you will specify the rules that determine how many sequence events the wizard will generate, and the time of those events, and how tiles will be allocated to each of those events.

The wizard will interrogate the scene database and compile a list of all tiles in the database that match the query you defined in the previous pages. These are presented in two drop lists, sorted by acquisition time.

Select the first and last tile you want considered for inclusion in the product. For convenience, you can use the "Earliest Time" and "Latest Time" options to automatically select the full temporal range of tiles. This can be useful if you acquire more scenes and rerun the wizard later to update the product. If you are creating a Scene Sequence product, there will be no other properties to define. In this case, each tile becomes an event in the sequence.

From the tile acquisition times, the wizard will attempt to define a set of acquisition events. To shape this process, you can define the maximum time span of an acquisition event (the acquisition period). For some mosaics, acquisition events may be clearly detectable and useful for defining the sequence event times. For others, they may be of little or no use.

Next, define the rule that establishes the sequence event times. The wizard will create an intermediate product for each sequence event – these can be thought of as frames in a movie. For some mosaics you can use the acquisition events. The system assigns the event time as the acquisition time of the temporally central tile. If this is not a suitable approach, you must define a regular event schedule.

The "By Interval" schedule uses the time of the first tile by default, then generates a new event every Nth day (determined by the interval you specify) until the event time exceeds the time of the last scene. If you want to start the sequence on a particular day, you can define the start date-time manually.

The "By days in the month" option will generate events on specific days in each month. Specify the days as comma delimited numbers from 1 to 31. For example, "1,16" would generate an event at the beginning of each month and in the middle of each month. By default, the first event will be on or after the first scene acquisition time. You can override this default if you want.

The final option is "By days in the year" which is similar, except that the days are counted from the beginning of the year. For example, "1,92,183,274" would generate events quarterly throughout the year.

The wizard must determine which tiles are assigned to each event. A tile may be assigned to multiple different events (this is common when compositing). You can define a time "Spread" ranging from before the event time to after the event time. This spread can be a specified as a number of acquisitions or as days. For example, if you are scheduling events in the middle of every month, then you might define the spread as being 16 days before the event and 16 days after it, capturing all the tile data acquired over the course of that month.

When creating a Mosaic Sequence, you have an additional “Select Tiles by” option to specify whether duplicated tiles are selected, and how they are selected. If you are creating a mosaic from Landsat scenes, then there is no reason to include multiple revisits of the same scene in the mosaic. When multiple revisits of a scene are present, you can specify a selection rule. Usually, you would choose the scene with the least cloud, or the scene that was acquired closest to the event acquisition time. If you are creating a mosaic from Sentinel2 tiles, you may want to include multiple revisits of the same tile, as it is likely that these were acquired from different scenes and the coverage in each tile differs. In this case, choose the “All Tiles” option.

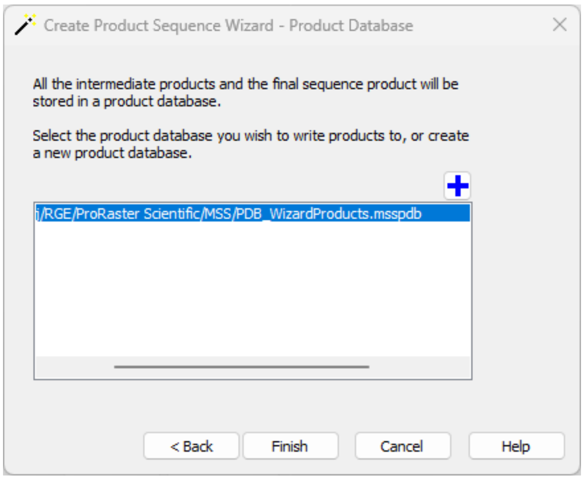

On the “Product Database” page you can select the multispectral product database to target. All the intermediate products and the final product will be written into this dataset.

You can elect to create a new product database if you wish – just hit the blue add button and define the new database name.

If you are editing an existing script, you must target the same product database as was targeted in the previous run. This is automatically selected and not editable.

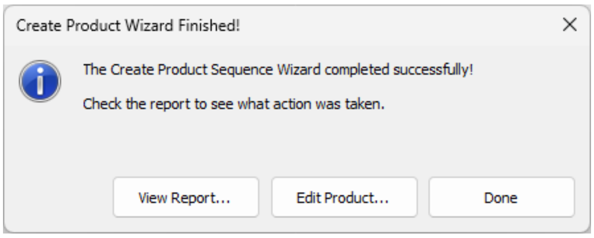

When you click “Finish” the wizard will execute. This will usually only take a few moments. A small dialog will open reporting whether the wizard succeeded or failed. Regardless of the outcome, you should then hit the “View Report…” button to review the report generated by the wizard. This is a simple ASCII text file and will open in your default text file editor. The report file will be written into the C:\Users\<YourUserName>\RGE\ProRaster Scientific folder. You can delete them manually as you see fit. Alternatively, if you delete all wizard scripts via the MSS Options dialog, then all reports are deleted at the same time.

If the wizard completes successfully, you should hit the “Edit Product…” button to open a product editor dialog for the final sequence product. You can use the editor to review the product and render it and begin your analysis.

If you have run the wizard in “Edit Mode” then any existing operations in the product operation cascade will now require updating. When you hit “Edit Product…” to open the product editor, the system will automatically trigger an update. If there are existing operations that compute statistics or export rasters, this update may take time to execute.

If you hit the “Done” button to exit, you will need to trigger a manual update of the product in the product editor later. This is critical. If you do not do this then all products in the processing cascade will be invalid. You can trigger this update manually by right clicking on the top item in the tree control and selecting “Execute Operation”. The whole product cascade will then be updated.

To help you get started, consider these recommendations.

- To create a Scene product, use the Scene Database Browser.

- To create a Scene sequence product, use the Scene Database Browser.
- To create a small Mosaic, use the Scene Database Browser.
- To create a large Mosaic, use the wizard.
  Order scenes in the mosaic by cloud.
  Select scenes in the mosaic by temporal proximity to the event time.
- To create a Mosaic Sequence, use the wizard.
  Order scenes in the mosaic by cloud.
  Select scenes in the mosaic by temporal proximity to the event time.
  Gather tiles from before the event time.
- To create a Composited Scene, use the Scene Database Browser.
- To create a Composited Scene Sequence or Composited Mosaic or Composited Mosaic Sequence, use the wizard.
  Order scenes in the mosaic by temporal proximity to the event time.
  Gather tiles from before the event time or straddling the event time.

Time Control Panel

The Time Control Panel is a floating (modeless) dialog that can be used to manipulate the time dimension in a rendering algorithm. It is available in ProRaster Scientific and ProRaster Premium.

Access the Time Control Panel by pressing the button to the right of the Undo and Redo buttons on the main dialog. Pressing the button again will close the dialog, or you can close the dialog in the normal way.

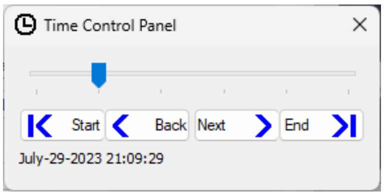

The Time Control Panel automatically connects to the algorithm that you are currently editing. Every time point (event) that is detected in the raster sources used in the rendering algorithm will be represented by the slider bar. You can use the slider bar to select a time point or press the buttons to step forward and back and to the start or end time.

ProRaster supports rasters that contain a temporal dimension. This means that the raster contains data recorded at one or more times. These are called events. Each event is classified as either a “partial replacement" or “total replacement” event, and each event has a date and time associated with it.

For every component in a rendering algorithm, you must specify the first and last event (inclusive) that you wish to render. In the user interface, the events are shown as a drop list containing the time of the event. When the system is asked to acquire data over a range of events, it performs a collation process. This is like a “painters’ algorithm” where data from the first event may be overprinted by data in a later event. There are some special rules that control this process.

1. When collating a range of events, the system only looks back to the latest “total replacement” event, if present. A total replacement event is like erasing the etch-a-sketch. All data acquired before it is ignored.
2. Valid data (pixels) recorded at a later time will overprint valid data recorded at an earlier time.
3. Invalid data can be classified as ‘empty’ (which is the default for all invalid data cells) or ‘null’. Data that is empty does not overprint valid data. Data that is null will overprint (erase) valid data.

Only two raster formats support the time dimension and events. These are the MRR format and the MVR virtual raster format. For all other raster formats, you can assign an event time to the raster using advanced raster preferences in the raster source editor dialog.

The multispectral satellite imagery processing tools in ProRaster Scientific routinely generate virtual rasters with multiple events. For example, you might use the Create Product Sequence wizard to create a mosaic sequence product that contains multiple events. In a product like this, all the events are designated total replacement events.

The Time Control Panel interrogates the algorithm and looks at all the rasters that are in use. It creates a list of all the unique event times, ignoring any rasters that only contain a single event. These times are sorted, and each unique time point becomes selectable from the control panel. When you select a time point it will choose the event in each raster source that has a time less than or equal to the time point selected. It assigns this event as both the first and last event in the range, implicitly assuming that each event is a total replacement event.

ProRaster Scientific Processing Operations

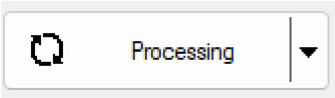

The Processing menu button provides access to a list of processing operations. Some of these operations produce raster outputs or some other kind of data for analysis, but most will produce an MVR virtual raster as output, which is then immediately opened for viewing in the main algorithm rendering interface. In other words, in most cases, the processing operation will take little or no time, regardless of the size or complexity of the inputs.

Some of the operations support batch operations. In this scenario, there are multiple data sources, and the same processing is applied to each source. An MVR is created for each input source, with a filename that reflects the source filename.

Most operations allow you to choose input rasters by either selecting an existing raster source or browsing to an existing raster file or to multiple raster files. It is vital to understand the difference between a raster source and a raster file – a raster source may point to multiple raster files, and this can change as raster files are added to, or removed from a folder. A raster source may also use driver preferences which modify the way the rasters are mounted by the raster engine. Some operations treat a raster source containing multiple files as a single entity, and some operate on each raster file in the raster source separately.

The virtual rasters produced by these operations can only be displayed or used in ProRaster. To take an MVR to another application or GIS system you will have to convert it (crystallise it) to a raster file in a transportable raster format. The “Export” operation provides this functionality.

The Processing menu button contains options to publish the current, saved, MRD rendering algorithm to MapInfo Pro or to publish the algorithm, clipped to a polygon, to MapInfo Pro. This functionality replaces the publishing buttons that are found in ProRaster Essential and ProRaster Premium. However, MapInfo Pro does not support the more advanced features in MRD algorithms that are used in ProRaster Scientific. Also, MapInfo Pro does not support many of the advanced features in the MVR virtual rasters that are likely to be consumed in the algorithms. Consequently, do not be surprised if the algorithm fails to load into MapInfo Pro or fails to render properly, or at all.

Export Operation

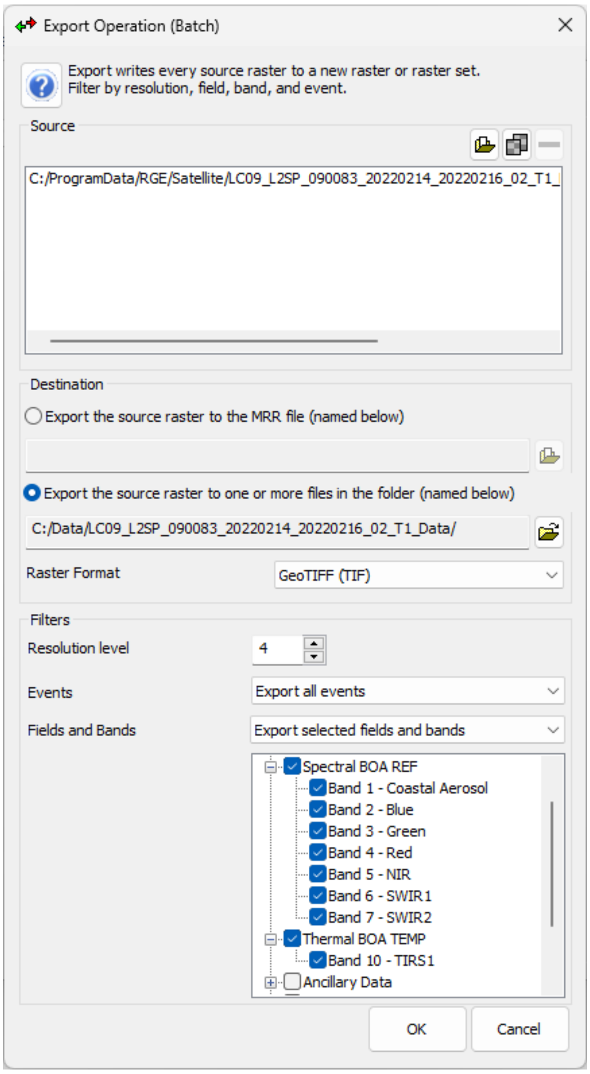

The Export operation writes a copy of a raster to a new raster file (or files) in a designated raster format. It is primarily designed to convert (or crystallise) an MVR virtual raster into a raster file. In ProRaster Scientific, chains of virtual rasters are used to build datasets for rendering and analysis. These virtual rasters cannot be used outside of ProRaster. However, you can export virtual rasters to other raster formats, thus “crystallising” the data, and use those rasters in other applications.

MVR virtual rasters support multiple fields, bands, and events. The only other raster format that supports these concepts is the MRR format. Consequently, the export operation will either copy the input raster to a single MRR format raster file, or it will copy each band of the input raster to a separate raster file in a transportable raster format like GeoTIFF. In this second scenario, the operation creates a folder for each input raster containing folders for all the fields and events, containing each band in a separate file.

In this operation, each raster you select as input is processed individually. If you select a raster source, then each raster file in that raster source is processed individually. The Export operation supports batch processing where multiple raster inputs can be processed serially.

You can either export each input raster to an MRR format raster, or to a transportable raster format. If you choose MRR format for the output raster, then the operation will generate a single MRR output file. Use the browse dialog to specify the output filename. This file will contain all the fields, events, and bands that you have chosen to export. You will be able to use this MRR file in ProRaster and in MapInfo Pro. You may struggle to use it in any other applications as it is likely to contain multiple fields and events and these features are unsupported by the GDAL driver used to access MRR format rasters in other applications.

When you export to MRR format, you can specify what compression codec is employed to compress the data in the MRR. Choose from “None”, “Quality”, “Balanced”, or “Speed”. If you choose “None” then the data remains uncompressed. This may result in large output files. “Quality” uses lossless compression for imagery and data fields. Data fields are compressed to the highest level obtainable to minimise the file size. Imagery files may be large. “Balanced” uses high quality lossy compression for imagery (JPEG) and a balanced lossless compression for data fields. “Speed” uses low quality lossy compression for imagery and a fast but poor lossless compression for data fields.

The alternative is to export to a transportable raster format. You can select either GeoTIFF (TIF) or Band Interleaved (BIL). In both cases, the export process will generate a single file for every band in the virtual raster. This may result in dozens of output rasters. For this reason, you must specify an output folder for your rasters. This will become the root folder for the export and each field and event will be a sub-folder of that root folder. The bands will be exported to rasters in those sub-folders.

When you select a folder, it is advisable to delete all the existing content in that folder. You can do this from the browse dialog. The name of the source raster is copied to the clipboard so that if you create a destination folder, you can paste the source name into it for convenience.

You can apply output filters including the resolution level, the event selection, and the field and band selection.

The default base resolution level is zero. Each increase in the resolution level doubles the cell size and reduces the output raster size by a factor of four. Conversely, each decrease in the resolution level halves the cell size and increases the output raster size four times. In an MVR, bands in a field may have different “natural” cell sizes. However, all bands are exported at the same cell size. For example, in satellite multispectral data the bands may have different cell sizes. The color bands are used to establish the MVR cell size, but other bands like Thermal or Panchromatic may have lower or higher resolution than this.

You can export all events or only the last “rolled-up” event. In the first case, each event is exported individually and there is no roll-up of events. In the second case, the data is rolled-up from the first event to the last event.

You can export all fields and all bands, or you can select individual fields and bands (in any combination) for export. Use the checkboxes in the tree control to turn on and off fields (and all bands in those fields) and to turn individual bands on and off.

Hit OK to execute the export. The dialog will close, and a raster processor dialog will open and display the progress of the processing operation. You can cancel the operation from this dialog if you wish.

Statistics Analysis Operation

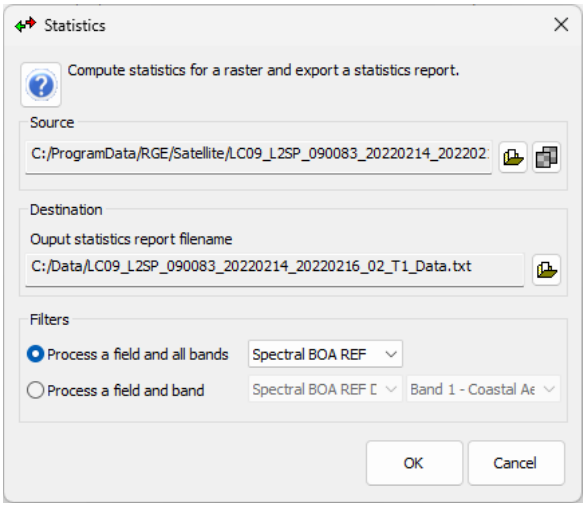

The Statistics Analysis operation will compute statistics for a selected field or bands in a raster input. It will write a text file report containing those computed statistics. The operation only supports the processing of a single raster file. If you select a raster source that contains multiple raster files, only the first file will be processed. Statistics analysis does not support batch processing.

Select a source raster for processing by either browsing to the file or by selecting a raster source. If you select a raster source that references multiple input rasters, only the first raster will be processed.

The operation will target a single selected field. You can compute the statistics for all bands in that field or for one selected band.

The computed statistics may be cached in an STTX file that is saved alongside the source raster. If so, you can view the statistics data using the Raster Source Editor dialog. The statistics will also be exported to a text report file. You need to specify the filename for this report.

The report will contain summary statistics for every band. This includes the number of valid cells, the minimum and maximum values, the mean, and many more. It will also contain histogram data. This is written into the report in a comma-delimited format.

The histogram contains many bins, each of which covers a numeric range. Put together, the bins cover the whole data range from the minimum to the maximum values in the raster. The width of each bin can vary. For this reason, we report the value for the bottom of the bin and the top of the bin. Within each bin, we also record the range of the actual data samples that lie in the bin. These are reported as the value bottom and value top. The number of samples that lie in the bin is reported. Also, the (cumulative) percentage of samples that lie in that bin, and all previous bins.

Join Operation

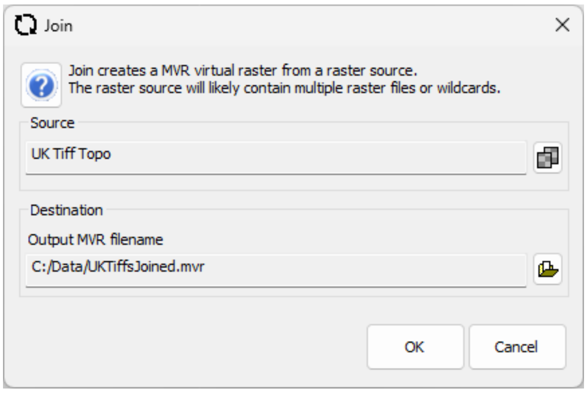

The Join operation outputs an MVR virtual raster that is a copy of an input source. Select a raster source containing multiple raster files as the input source. Join does not allow you to browse to select input rasters, so you must construct a suitable raster source in the Raster Source Editor first. Join does not support batch processing.

Put simply, the Join operation converts a multi-file raster source into a single MVR. Many of the processing operations do this, but Join does not modify the data.

The primary application of this is to take a collection of rasters that are in a tiled arrangement (or postage stamp arrangement if you prefer) and bring them together into a single raster. In this scenario, it is likely that all the source rasters will have the same coordinate system, and the same cell size and it is likely that the cells will align and not overlap. In fact, it is a pre-requisite that the rasters all share the same coordinate system. However, the cell size and alignment do not have to match. Where cells in rasters overlap, the order of the rasters in the list determines what value will be retained in a cell where there are multiple possible inputs.

Clip to Polygon Operation

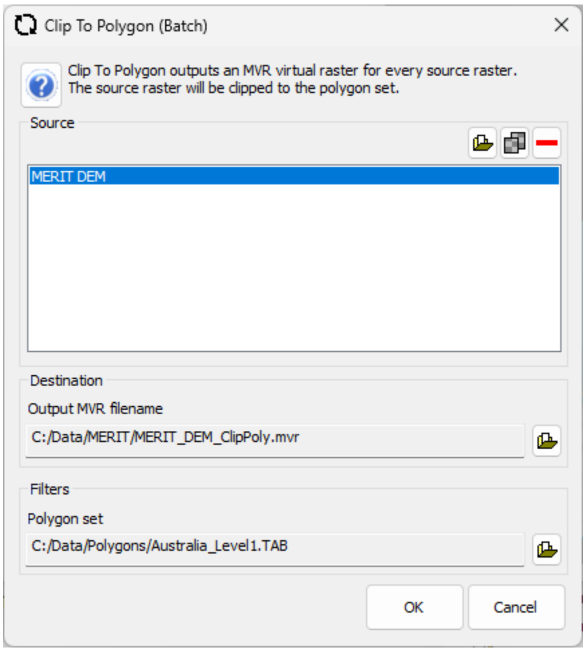

The Clip to Polygon operation outputs an MVR virtual raster that is a copy of an input source, clipped to a complex polygon set. An input source can be a raster source or a raster file. If the raster source contains multiple raster files, then they are merged. Clip to Polygon supports batch processing where multiple raster inputs can be processed serially.

This operation is used to clip a raster, or a raster source containing one or more rasters, to a complex polygon set loaded from a MapInfo Pro table file.

The polygon set can contain multiple polygons and the polygons can have multiple parts and holes. The coordinate system of the polygons and the rasters do not have to match.

If you clip a raster source that references multiple raster files, then this operation is like combining a Join operation with a clipping function. Consequently, there is no need to Join rasters and then clip them – this operation performs those two functions in one step.

Clip to Raster Operation

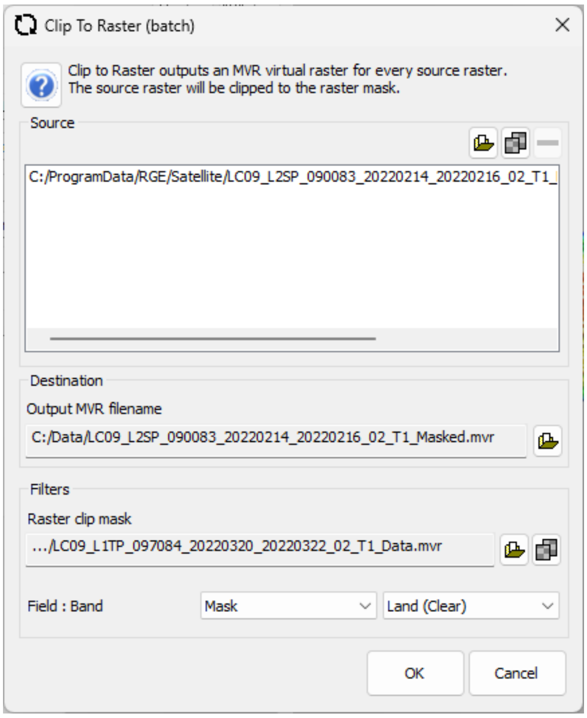

The Clip to Raster operation outputs an MVR virtual raster that is a copy of an input source, clipped to a raster mask. An input source can be a raster source or a raster file. If the raster source contains multiple raster files, then they are merged. Clip to Raster supports batch processing where multiple raster inputs can be processed serially.

This operation applies the cell validity of the mask raster to the source raster. Valid cell data will only be retained in the output MVR where cell data is valid in both the source raster and the mask raster. Where either the source raster or the mask raster cells are invalid, the output MVR cells will be invalid.

This is just the same as using the Mask component in the rendering algorithm to prevent the rendering of pixels where they are not defined in the mask raster. However, in this case, that masking is actually applied to the raster directly, so that for operations performed on the MVR thereafter, the masked pixel data will not exist.

If you clip a raster source that references multiple raster files, then this operation is like combining a Join operation with a clipping function. Consequently, there is no need to Join rasters and then clip them – this operation performs those two functions in one step.

Reproject – Resample – Realign Operation

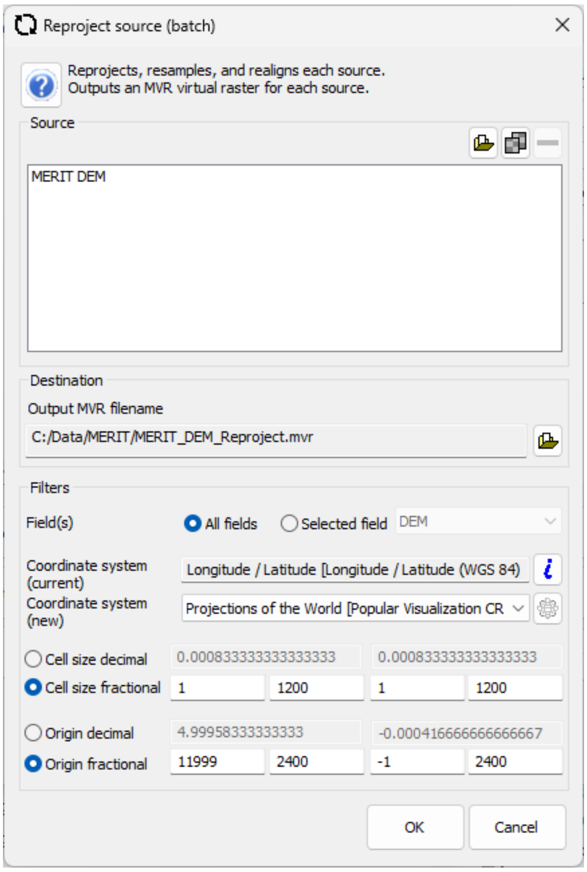

The Reproject operation outputs an MVR virtual raster that is a copy, or a partial copy, of an input source with a different coordinate system to the input source. An input source can be a raster source or a raster file. If the raster source contains multiple raster files, then they are merged. Reproject supports batch processing where multiple raster inputs can be processed serially.

The Resample operation outputs an MVR virtual raster that is a copy, or a partial copy, of an input source with a different base resolution cell size to the input source. An input source can be a raster source or a raster file. If the raster source contains multiple raster files, then they are merged. Reproject supports batch processing where multiple raster inputs can be processed serially.

The Realign operation outputs an MVR virtual raster that is a copy, or a partial copy, of an input source with a different base resolution cell alignment to the input source. An input source can be a raster source or a raster file. If the raster source contains multiple raster files, then they are merged. Reproject supports batch processing where multiple raster inputs can be processed serially.

The Reproject, Resample, and Realign operations all operate on the same principle and use the same dialog, although not all controls are enabled in all operations.

Reproject allows you to change the coordinate system of a raster. When you do, it is likely that you will also want to change the cell size (resample) and change the cell edge coordinates (realign). The dialog allows you to make all these changes at once.

The current coordinate system of the raster is shown on the dialog, and an information button is provided to provide details about this coordinate system. You can choose any recently used coordinate system by selecting from the new coordinate system drop-list. Alternatively, hit the Projection Explorer button to browse to a new coordinate system.

It is likely that the change of coordinate system will trigger a change in cell size. For example, if the current coordinate system is geographic and the cell size is expressed in degrees and the new coordinate system is UTM, then you will want an equivalent cell size in metres. Because of the nature of MVR rasters, it is not necessary to get this cell size conversion perfectly correct. In fact, all that the cell size really determines is the size of the raster at the base resolution. This is the resolution level at which you will likely compute statistics or export to a different format, so it is important. However, it is not critical to rendering or other processing operations.

You can enter the cell size as a decimal value (width in X and height in Y) or as a fraction. In the latter case, you specify a signed integer numerator and an unsigned integer denominator. For example, a cell size of one arc-second is (1/3600) degrees. If possible, it is a good idea to express the cell size as a fraction. The raster engine uses fractional mathematics to eliminate or minimise round-off errors in floating point numbers.

If you change the cell size, it is likely that you will want to specify the cell alignment. You do this by providing X and Y coordinates for the tile origin. Just like cell sizes, you can specify these coordinates as decimal values or as fractions, and fractions are preferred.

When you move the tile origin for a virtual raster, you do not actually relocate the raster data. All you are doing is specifying where the bottom left corner of cell (0,0) is at base resolution. The tile coordinate system is infinite, so you can put it anywhere. However, you will usually place it at the bottom left corner of the raster.

Moving the tile origin does change the coordinates on which the cell boundaries lie. For example, if you have a cell size that is 100x100m and you place the tile origin at (0,0) then the cell edges will lie on values that follow the series 0, 100, 200,… etc and the cell centre will lie on 50, 150,250,… etc. If you want the cell centre to lie on the 0, 100,200,… values then move the origin to (50,50).

When you change the coordinate system, cell size or cell alignment, it will trigger interpolation in the virtual raster engine. At the end of the day, the engine will try to populate a raster mesh with the best possible data from a source raster. It will do this by targeting a resolution level in the source raster that is the same or better resolution than the destination mesh. It will interpolate data from this level using nearest neighbour, bi-linear, or bi-cubic methods depending on the field type, band data type, and user preferences.

If you choose the Resample operation, the coordinate system information will be shown but not editable. The cell size and tile origin controls will be editable. If you choose the Realign operation, only the tile origin controls will be editable.

Calculator Operation

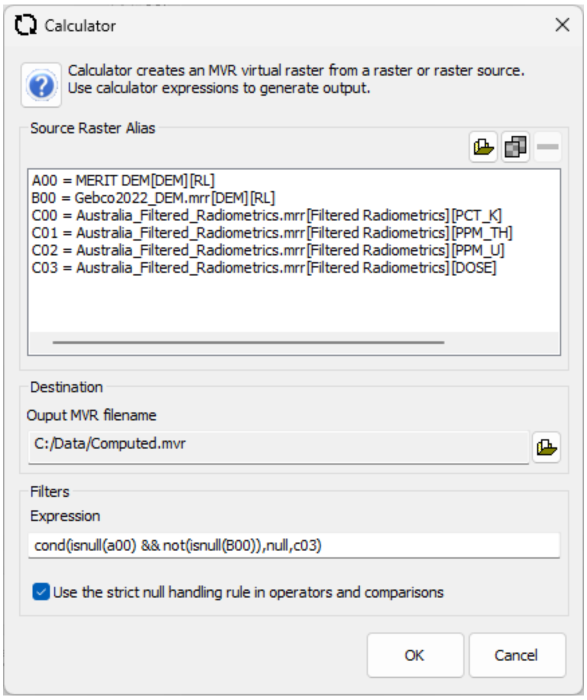

The Calculator operation outputs an MVR virtual raster that applies a calculator expression to the input data. The MVR will contain a single output band. You can select data from any inputs as variables in the expression. Those inputs can either be a raster source (which may contain multiple raster files), or a single raster file you have selected by browsing. Calculator does not support batch processing.

The Calculator evaluates the mathematical and logical expression that you have defined, at every raster location requested by the virtual raster engine. For example, when rendering, the virtual engine will evaluate the expression at every pixel centre. Calculator creates an MVR with a single Continuous field that contains one 32-bit floating point real value band. Be cognisant of the fact that the expression result will always be returned as a 32-bit real value.

Raster cells can either contain a valid value or an invalid value. In the calculator, an invalid value is referred to as a "null". It is likely that your expression will generate either valid values or "null", depending on the inputs and the evaluation of the expression logic. You may want to test to see if an input value is valid within the expression, to decide whether to process the value or not. To do this you can use the function "isnull()" or you can use a logical expression and the constant "null".

You can define input variables that are linked to a raster, field, and band. When the Calculator evaluates the expression at a requested (X,Y) location, it will acquire the value of each input variable raster at that location. The Calculator can do this regardless of the coordinate systems of the input and output rasters, the cell sizes, cell data types etc. It will acquire data in an appropriate projection, from an appropriate resolution level, using an appropriate interpolation technique, in real-time. Note that because a variable has a defined field and band, there is no "band notation" syntax supported in the calculator expression, and the expression only emits a single band output.

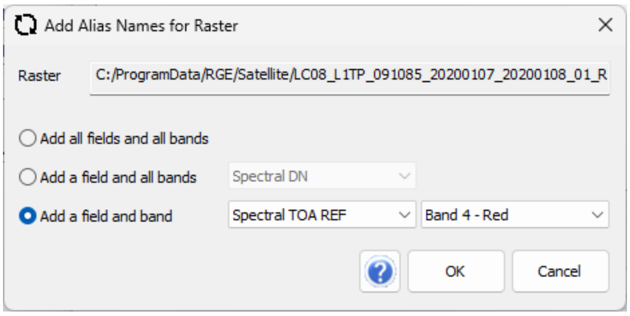

When you browse to a raster or select a raster source, a new dialog will open asking you to define the fields and bands that you want to include in your calculator expression.

The Calculator dialog will automatically generate a unique alias for every raster input. The alias will look like "A00" or "B14". You will use these alias names when you refer to an input raster in the calculator expression. In the dialog, select the alias in the list, and the alias name is copied to the clipboard so you can paste it into your expression.

If an input variable is invalid ("null") and it is used in a mathematical function, then the result will be invalid. For example, the expression "sin(A00)" will return the sin() of the variable A00 when it is valid, and it will return "null" otherwise. If an invalid value is consumed by a mathematical operator, then there are two different outcomes that you can control. For example, the expression "A00 + B00" adds two grids together. If either A00 or B00 is invalid, then we can either return “null” (strict null handling) or we can return that valid value and treat the invalid value as though it does not exist (relaxed null handling). In this example, relaxed null handling would result in valid data over the union of the two rasters and strict null handling would result in valid data over the intersection of the two rasters. This rule affects the operators (+, -, *, /), the comparison operators (<, <=, >=, >), and the ”cond”, “min” and “max” functions as described below.

+, -
If one of the input values is invalid, it is treated as though it is zero.
*
If one of the input values is invalid, it is treated as though it is one.
/
If either the numerator or denominator is invalid, it is treated as though it is one.
<, <=, >, >=
Invalid values are treated as though they are -infinite.
Cond
If the condition is invalid, return the result for a false condition.
Min
If one of the input values is invalid, return the other value.
Max
If one of the input values is invalid, return the other value.

The output data type will be a real number (32-bit floating point) by default. This is an appropriate data type choice for most calculator expressions. In some cases, where the input raster is imagery, you may require an output color data type. In these cases, make sure to check the “Output as a color data type” button to generate an RGB output data type.

The Calculator supports brackets “(“ and “)”. If you use a left bracket, you must close the bracket with a right bracket. Do not use square brackets “[]” or curly braces “{}”.

The Calculator supports two constants – “pi” (3.141592654) and “null” (a magic number). To use other constants in your expression, enter them as though they are real numbers. You can use the unary minus sign (e.g., -2.45) but it is a good idea to enclose values that use unary minus in brackets (e.g., A00+(-2.45)*B00). Otherwise, the expression parser gets confused and you may get an incorrect result. You can also use exponential notation using either “e” or “E” (e.g., 2.45e-2, -1e10, -1e-03, 2.45E-2).

Example expressions

To add two rasters together, use “A00 + B00”. With strict null handling, you will get a result only at the intersection of the two rasters. With relaxed null handling, you will get the union.

To merge two rasters together, use “A00 + B00” and use relaxed null handling.

To find the difference between two rasters, use “A00 – B00” and use strict null handling.

To average two rasters, use “(A00 + B00)/2” and use relaxed null handling.

To transform a raster Y = A*X + B, use “2.5*A00 + 2000” and use strict null handling.

To output a result where a raster is less than a constant, use “cond(A00<100,A00,null)” and use strict null handling.

To output a result where a raster is between two constants, use “cond(A00>=100 and A00<=200,A00,null)” and use strict null handling.

If you want to explicitly handle invalid values, then you can nest your conditions like this.
"cond(A00=null,null,cond(A00<100,A00,null))"
Rather than test against "null", you can use the "isnull()" function.
"cond(isnull(A00),null,cond(A00<100,A00,null))"
This is equivalent to saying "if A00 is null, then output null, otherwise if A00 is less than 100 output A00, otherwise output null".

Here is a complete list of all supported syntax and functions.

Constants

pi   3.141592654
null   -1E32 (a magic number representing an invalid cell value)

Mathematical Operators

+   addition
-    subtraction
*   multiplication
/   division
^   power (A00^2 raises A00 to the power 2)

Brackets

()   Always close your brackets and do not use square brackets or curly braces.

Unary Minus

-    Negates the value of the operand. I recommend you surround the operand in brackets.

Logical Operators

=   Equal to. You can also use "==".
< >   Not equal to. Do not use "!=".
<   Less than
<=   Less than or equal to
>=   Greater than or equal to
>   Greater than
&   And. You can also use "&&" or "and".
| Or. You can also use "||" or "or".

Conditional Functions

cond(exp,a,b)   Returns "a" if the expression evaluates true, or "b" otherwise.
You can also use an "IF exp THEN a ELSE b ENDIF" syntax.
isnull(a)   Returns true if "a" is invalid ("null").
any(x,…)   Returns true if the operand x is equal to any of the following operands, or within the range of any of the following operands. Specify a range of values as "x1:x2". For example, use this expression to return several values and bands of values and invalidate the rest – cond(any(A00,0,10,50,100:200,500:1000,5000:10000),A00,null)

Standard Numeric Functions

cos(x)   Cosine of an angle of x radians.
sin(x)   Sine of an angle of x radians.
tan(x)   Tangent of an angle of x radians.
acos(x)   Inverse of cosine.
asin(x)   Inverse of sine.
atan(x)   Inverse of tangent.

atan2(y,x)	Arc tangent of y/x, in radians.
cosh(x)	Hyperbolic cosine of x.
sinh(x)	Hyperbolic sine of x.
tanh(x)	Hyperbolic tangent of x.
acosh(x)	Inverse of hyperbolic cosine.
asinh(x)	Inverse of hyperbolic sine.
atanh(x)	Inverse of hyperbolic tangent.
pow(x,y)	x raised to the power y.
sqrt(x)	Square root of x.
exp(x)	Eulers number raised to the power x.
log(x)	Common base 10 logarithm of x, see log10().
ln(x)	Natural logarithm of x, see log().
alog(x)	Inverse common logarithm, pow(10,x).
aln(x)	Inverse natural logarithm, exp(x).
abs(x)	Absolute value of x.
mod(x)	Floating point remainder of x/y, see fmod().
ceil(x)	Rounds x upward, returning the smallest integral value that is not less than x.
floor(x)	Rounds x downward, returning the largest integral value that is not greater than x.

These functions call the standard C or C++ functions of the same name (except where indicated). For help on these standard mathematical functions consult any C++ online reference.

Numeric Functions

round(x)	Rounds x up or down to the nearest integer.
min(a,b)	Returns minimum of a and b.
max(a,b)	Returns the maximum of a and b.
degtorad(x)	Converts x from degrees to radians.
radtodeg(x)	Converts x from radians to degrees.
rand()	Returns a random number between 0 and 1.
unit(x)	Clamps x to the range>=0 and <=1. Values below 0 are returned as zero and values above 1 are returned as one.
clamp(x,a,b)	Clamps x to the range>=a and <=b. Values below a are returned as a and values above b are returned as b.
clip(x,a,b)	Clips x to the range>=a and <=b. Values below a are returned as null and values above b are returned as null.

Even if a function takes no parameters, make sure you include the trailing closed brackets.

Bitwise Functions

bitand(a,b)	Bitwise "and" operator on the operands cast to unsigned 64-bit integers.
bitor(a,b)	Bitwise "and" operator on the operands cast to unsigned 64-bit integers.
bitxor(a,b)	Bitwise "xor" operator on the operands cast to unsigned 64-bit integers.
bitnot64(a)	Bitwise "not" operator on the operand cast to unsigned 64-bit integer.
bitnot32(a)	Bitwise "not" operator on the operand cast to unsigned 32-bit integer.
bitnot16(a)	Bitwise "not" operator on the operand cast to unsigned 16-bit integer.
bitnot8(a)	Bitwise "not" operator on the operand cast to unsigned 8-bit integer.
bitshl(a,b)	Bitwise shift left "b" bits on "a" cast to unsigned 64-bit integer.
bitshr(a,b)	Bitwise shift right "b" bits on "a" cast to unsigned 64-bit integer.

Logical Functions

not(a)	Returns true when a is equal to zero.

Color Manipulation

rgba(r,g,b,a)	Converts r, g, b, alpha components (0–255) to a 32 bit 0xABGR value.
red(col)	Return the red component (0-255) from a 0xFFBGR or 0xABGR value.
green(col)	Return the green component (0-255) from a 0xFFBGR or 0xABGR value.
blue(col)	Return the blue component (0-255) from a 0xFFBGR or 0xABGR value.
alpha(col)	Return the alpha component (0-255) from a 0xABGR value.
hue(r,g,b)	Convert color components (0-255) from RGB to HSI space and return hue.
saturation(r,g,b)	Convert color components (0-255) from RGB to HSI space and return saturation.
intensity(r,g,b)	Convert color components (0-255) from RGB to HSI space and return intensity.

Raster Constants

GetXOrigin()	Returns the x coordinate of the bottom left corner of the bottom left cell of the output raster.
GetYOrigin()	Returns the y coordinate of the bottom left corner of the bottom left cell of the output raster.
GetCellSizeX()	Returns the width of the base resolution cell size of the output raster.
GetCellSizeY()	Returns the height of the base resolution cell size of the output raster.
GetCols()	Returns the number of columns in the base resolution level of the output raster.
GetRows()	Returns the number of rows in the base resolution level of the output raster.

Even though these functions take no parameters, make sure you include the trailing closed brackets.

Raster Variables

GetX()	Returns the X coordinate of the centre of the pixel being evaluated.
GetY()	Returns the Y coordinate of the centre of the pixel being evaluated.

Even though these functions take no parameters, make sure you include the trailing closed brackets.

Create Raster Mask Operation

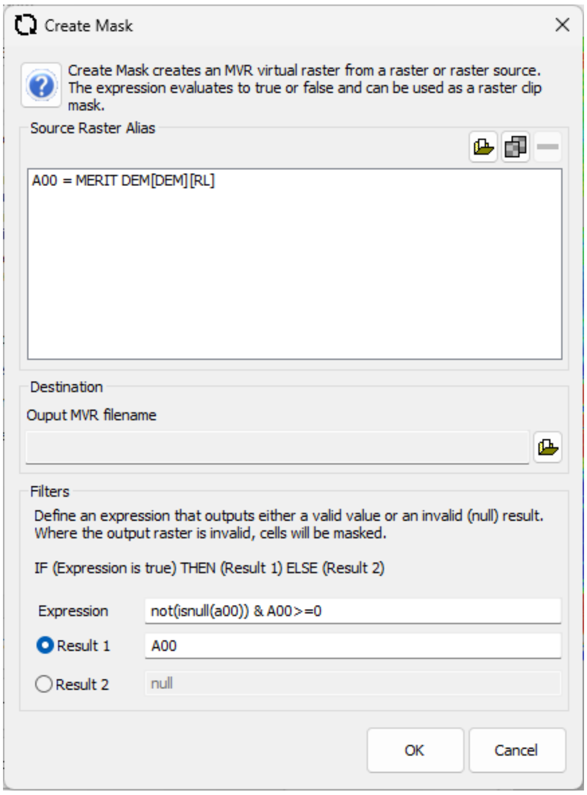

The Raster Mask operation outputs an MVR virtual raster that can be used in the Mask component of a rendering algorithm. The mask raster will contain a single output band. The mask will prevent rendering of pixels where the mask raster is invalid, or where the mask raster is valid and the cell value is zero, or evaluates to zero after data conditioning, or is invalid after data conditioning.

The output MVR applies an IF-THEN-ELSE calculator expression to the input data. One of either the THEN or ELSE results must be “null” or zero – which generates a cell that will trigger masking. You can select data from any inputs as variables in the expression. Those inputs can either be a raster source (which may contain multiple raster files), or a single raster file you have selected by browsing. Raster Mask does not support batch processing.

Raster Mask creates a singled banded continuous field raster that evaluates the calculator expression – “cond(expression,result1,result2)”

The calculator engine uses strict null handling. If this simplified approach is not sufficient for your purposes, you can always generate your own raster mask expression in the Calculator operation. For more help on generating alias names and on calculator expressions, see the Calculator operation documentation.

Conditioning Operation

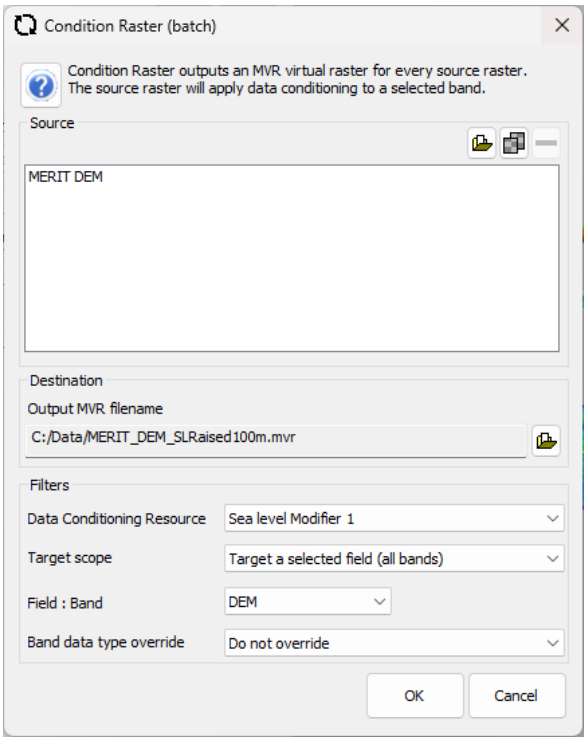

The Conditioning operation outputs an MVR virtual raster that applies a data conditioning filter to a selected field, or to a selected band, of an input raster source.

Each raster source you select, and each raster file you select by browsing, are treated as a separate input. If a raster source contains multiple raster files, then the output MVR will merge these files. Condition supports batch processing where multiple raster inputs can be processed serially.

You can choose a data conditioning resource from any that you have defined using the Data Conditioning Editor.

Choose the field that it will be applied to. You can apply it to all bands in the field or to a particular band. You can also override the output data type of the bands just in case the data type of the source band is not suitable for the output band. This will depend on what kinds of operations the data conditioning filter applies. For example, if you apply a transform operation, this may require you to generate floating point output data bands.

Image Sequence Operation

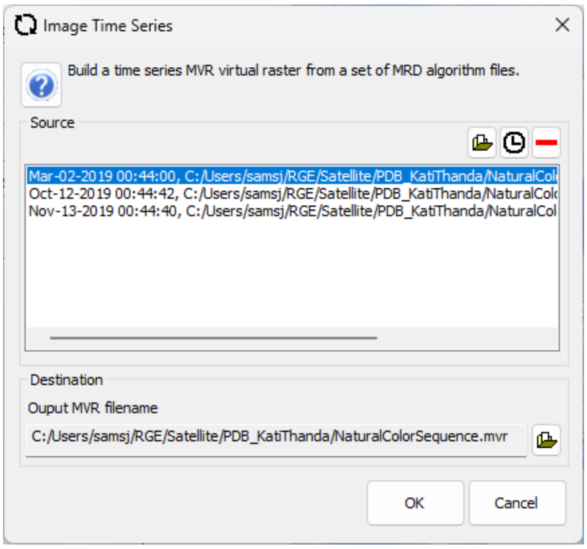

The Image Sequence operation outputs a virtual raster that contains an image field and an event for each input MRD rendering algorithm supplied.

Select multiple input MRD rendering algorithm files that you have previously prepared for a specific region. For example, you may have a satellite scene mosaic. You have acquired scenes over multiple time periods and prepared a mosaic for each period. For each of these mosaics, you have generated an MRD rendering algorithm. Now you want to combine these algorithms.

Browse to each input MRD file. It will be added to the list and automatically placed in order by increasing time. The event time is acquired from the MRD and reflected in the list. You can edit and modify this time, which may change the order of the files in the list. You can also delete items from the list. Once you have added all the files you require, specify the output MVR filename, and hit OK. The new MVR file will immediately be opened and displayed in ProRaster.

The MVR will be displayed using an Image layer and each MRD input will generate one slide (or frame) in the series. Use the time range controls in the Image component to select each slide and display it, and to flick through the events.

Transform Operation

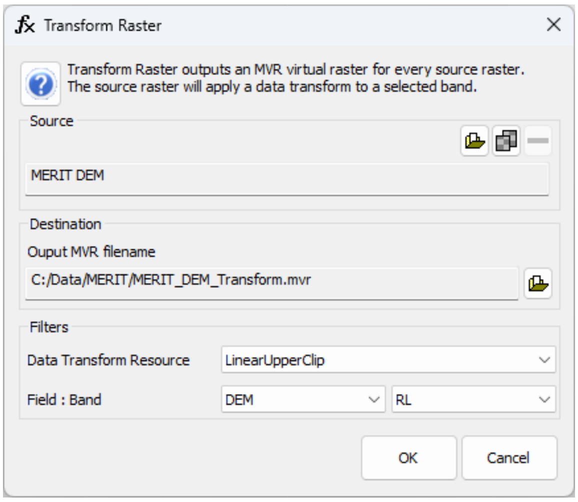

The Transform operation outputs a virtual raster for each selected input raster or raster source that transforms the values of a single selected band.

In single input mode, you are required to browse to the input raster, or select a raster source. You must specify the output virtual raster. In batch mode you can process multiple input rasters and you must specify the folder into which all output virtual rasters will be written. The field and band structure of all the input rasters ought to be the same.

Select the field and band that will be processed. If the raster contains multiple events, all events in the raster will be processed.

Finally, specify a data transform. You must create the data transform using the Data Transform Editor. Once you have done so, it will be available for selection. The data transform will be applied to the input band data values to generate the output data values. In single input mode, when you hit OK the new MVR file will immediately be opened and displayed in ProRaster.

Most data transforms that you can create in the Data Transform Editor emit transformed values between 0 and 1. This value (referred to as an *index*) is then scaled to acquire a value from a color look-up table. In this case, you are likely to want to transform the data values from some range and distribution to another range and distribution. If so, you ought to use one of the following four data transform types that allow you to modify the *index* value, which you can set to the output data values. For more information on the data transforms, see the documentation for the Data Transform Editor.

Table: Index over range: Specify the input data range of the transform by setting bandpass limits. The input data range is equally divided over the rows of the table. The minimum input data value for each row will map to the output data (index) value you specify.

Table: Index for value: Specify the input data value for each row, and the output data value it maps to. This is the most likely transform you will employ.

Table: Index for percentage: Specify the input data range of the transform by setting bandpass limits. Specify the input data value for each row in the table as a percentage (0 – 1) of the input data range. This will map to the output data value (index) in that row.

Table: Index for percentile: As above, except that the input data values are a percentile (0 - 1) of the data distribution. Percentiles are transformed to data values via a distribution histogram.

Configuring Satellite Multispectral Options

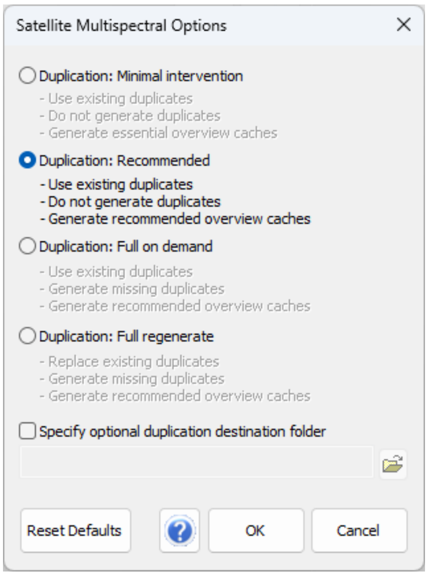

Access the MSS Options dialog from the system menu in ProRaster Scientific. From here you can modify some of the default behaviour of ProRaster Scientific. To return all options back to their standard defaults, hit the "Reset Defaults" button and hit "OK".

Duplication: These controls allow you to duplicate scene rasters. When you duplicate a scene raster, the duplicate raster is used in the scene assembly MVR rather than the original scene raster. All duplicate rasters are written in MRR format.

For information on these four options, see the section on creating scene databases.

Duplicate rasters can be written into the original scene folder, alongside the original raster that they duplicate. Alternatively, you can store your duplicate rasters in a folder of your own choosing. Underneath this folder, the system will create a unique folder for each scene for which you generate duplicate rasters. Check the "Specify optional duplication destination folder" button, then browse to the required folder, to specify a folder. Or simply uncheck the button to indicate that new duplicates should be stored in the original scene folder.

Tips and Advanced Techniques

Use Drag and Drop

Get straight into rendering your data by using drag and drop. Drop your raster file onto ProRaster from Windows File Explorer to get started. Use drag and drop in the Raster Source Editor to populate raster sources faster.

Know your data

When you create or take delivery of a new raster, take some time to get to know your data! Drag and drop it into ProRaster and see how ProRaster displays the raster. What are the field types and band data types? Are these appropriate, or should you use data preferences to modify the defaults? Is the coordinate system correct?

Organise your data

For any project, keep your data, your custom color tables, and your MRD algorithms separate and organised. Use relative paths (see Options dialog) to make sure that MRD algorithms remain valid, even if you shift your project directory to an archive location. Keep global datasets that are useful across multiple projects on hand and accessible and use them in your algorithms. Keep your raster sources organised – I tend to use a "Project/Type/File" naming convention for raster sources so that I can find data faster.

Check out the Information Report

From the Raster Source Editor, get familiar with the Information Report option. To see everything, ask for a "Complete" report. You will see statistics, classification tables, metadata and much more besides. Learn how to interpret the statistical information and assess your data quality. Take a look at the classification table, when present, to check your data quality.

I want more statistics!

Use the Data Transform Editor to find more statistics. For the raster being previewed, you can see all the summary statistics for a band and the distribution histogram as either a table or a graph.

Clean and Prepare

Use the Clean and Prepare processing operations in the Raster Source Editor to make sure your new raster has statistics computed and cached and has an overview cache prepared. If you want smaller files, or higher quality imagery in the overview, use the "Favour Quality" option.

Layers vs Raster Source

If you want to display multiple rasters in an algorithm, you can either create a layer for each raster or you can create a raster source containing multiple rasters and render it with a single layer. If the rasters are not all the same coordinate system and structure, then use layers. If the rasters are imagery or classified, then use a raster source. If the rasters are some continuous data (like terrain) then it depends. To use a raster source, you will need a data transform that is suitable for all the rasters, and you may need to adjust hill-shading manually.

Controlling hill-shading depth across multiple rasters in a raster source
When you use hill-shading with multiple rasters in a raster source, the shadow depth in each of the rasters may differ and you will see some inconsistency between rasters. It occurs because the statistics used are local,

not global. To remedy this, on the Shadow page turn on "Manual statistics". In the edit box, enter a value for the cell-to-cell average change. If you look at the statistics page for the raster in the Data Transform Editor, this will be the "Mean" value in the "Spatial" section. Twiddle the number until you get the desired result.

Create a color table for terrain

Imagine you want 16 color bands spread over known intervals for a terrain raster. In other words, you want values from 0 to 100 to be brown, 100 to 250 to be green, etc. Use the Color Table Editor to create the table either from scratch or by copying an existing table (when you copy an existing table, you inherit its colors). Select a "Legend" with 16 items and "Continuous Ranges". Now edit the new table and adjust the minimum and maximum values for each range. Set the color for each range as desired. Now display your raster using a LUT Color layer, choose the new color table as the data transform and as the color table.

What if you want continuously variable color, but you still want the color transitions to occur at precise levels? This time, create a Map. In this file you need to specify each key color and the data value that color corresponds to. All the colors in between the key colors will be created by interpolation. Make sure you have a data value at the first and last color.

MRD is a raster

An MRD algorithm, saved to disk, can then be loaded back into ProRaster as a raster! When you do so, it will look like it is an image field raster. You ought to use an Image layer to render it. What is the point? One application is transparent layers. If you have multiple transparent layers and you want to drape the whole thing on terrain, then a good approach is to build an algorithm with all the layers and no terrain, then render that algorithm in a new algorithm draped on terrain. Another application is to output some imagery, like maps data for example, to a greyscale algorithm. Then, read this back in as a raster and use the greyscale to opacity modulate the original data. In other words, you are opacity modulating a layer by the greyscale equivalent of itself.

Tips for Displaying Classified Rasters

A classified raster will generally have a color band or individual red, green, and blue color bands. ProRaster will automatically display a raster like this using an Image layer or an RGB Color layer. Usually, a column in the classification table will contain a label. If so, this will be reported as the cell value in the map window.

If there is no color band, then you will need to assign colors manually via a LUT Color layer. Use the color table editor to create a color table legend using discrete value matching, for each class in the raster. If you modulate by the Index band (always band zero), then the classes are integer values from 0 to N-1. Look at the statistics to see what the range of values is.

If you want to display only a selection of the classes, then the best way to proceed is to use the color table editor to create a color table legend for this raster. Then you can edit this color table and remove the classes that you do not wish to display. Now render the raster using a LUT Color layer, modulate by the Index band, and choose your new color table as the data transform and as the color table.

You can also use data conditioning to remove certain colors from an image, and this can be used as a strategy with an Image layer to quickly remove certain classes. Another option, when using a LUT Color layer, is to use the clipping limits to clip out all classes except one or remove certain classes.

Time and Events

ProRaster supports the time dimension. Unfortunately, your raster formats do not. Only two formats support multiple events - MRR and MVR. To build an MRR with multiple events requires writing code using the

MapInfo Pro Advanced raster API. To build an MVR requires documentation that no one has ever published or shared. I plan to provide tools to do this in a future version of ProRaster Scientific.

WebMap (WMS) rasters

Place your local service description XML files in the C:\ProgramData\RGE\WebMaps folder so that they appear in the user interface, and you can add them to algorithms via the "Open" menu button or "Add Layer" menu button. If the text is hard to read, or to reduce the amount of IO that is done, use the "Adjust Resolution" option on the primary Component property page. A setting of 1 is usually appropriate.

WebMap Overlays

Place your greyscale WebMap virtual rasters in the C:\ProgramData\RGE\WebMapOverlays folder so that they appear in the user interface, and you can add them to algorithms via the "Add Layer" menu button. If the text is hard to read, or to reduce the amount of IO that is done, use the "Adjust Resolution" option on the primary Component property page. A setting of 1 is usually appropriate.

Long File Names

Windows has been around for a long time and over the years the definition of a valid filename has changed. Today, there are three kinds of filenames – Short, Long and Very Long. In this context, a filename is the full filename including the driver, path, name, and extension.

For example: "C:\Users\Sam\MSS\SDB_Sentinel2.msssdb".

Short filenames come to us from the mind of Bill Gates and date back to the old 8.3 format used in DOS. You can still convert a Long filename to a Short filename – it will be truncated and modified. Long filenames are what we use today, but many people do not know that they are still limited. A Long filename should not exceed 260 characters in length. This includes the drive specification, path, and filename plus extension. Very Long filenames can be up to 32767 characters long – the final word in long filenames.

ProRaster supports Short and Long filenames. It does not support Very Long filenames. This is because the operating system support for these filenames seems to be patchy at best. Even File Explorer does not support Very Long filenames. Try to create a folder that has a total length of more than 260 characters, and you will find File Explorer will not accept it. Even if you have Very Long folders and files, although you can see them in File Explorer, it will limit your ability to work with them.

This is an issue because Satellite Multispectral scenes tend to use very verbose folder and filename structures. The worst offender is Sentinel2. In a typical workflow you will download your Sentinel2 scene and then extract the data to a location on your hard disk. The data will be extracted to a folder that is named after the unique ID for the scene. The actual data files are in sub-folders under that primary folder. A filename in this scene will typically have a length of 160 characters – not including the drive specification and the path you placed the scene folder into. In other words, you have only 100 characters to use yourself.

It is easy to extract Sentinel2 data and create filenames that are too long. When this happens, you will find that ProRaster will fail to identify the scene and add it to a scene database. Your only solution is to shorten the file path.

One thing you can do with Sentinel2 scenes is shorten the scene root folder name. This is a direct copy of the scene unique ID. However, it is packed full of useless information. Consider the folder name below. It is 66 characters long and all the unnecessary information in it is highlighted.

S2A_MSIL2A_20230429T004701_N0509_R102_T53HQA_20230429T044157.SAFE

_MSIL	Unnecessary
T004701	Acquisition time, not needed
_N0509	Processing baseline, not needed
_R102	Orbit, not needed
T044157	Processing time, not needed
.SAFE	Unnecessary

As each satellite cannot visit a particular tile twice in a day, all the detailed information including the time of acquisition and processing, the orbit and processing baseline is unnecessary. This unique ID could be reduced to “S2A2A_20230429_T53HQA_20230429” and still be unique. And you save 35 characters which you can use in your folder structure.

Adjusting Color Transforms in RGB Layers

When displaying imagery including satellite multispectral imagery in an RGB Color Layer, you need to provide transforms from data space to color space for the Red, Green, Blue, Pan, and Opacity components. The Pan component is often adjusted in the same way as the color components whereas the Opacity component will have an independent transform. Some simple tips can help you to adjust these transforms to improve the visual appearance of your imagery.

In the diagrams below I assume you have a simple transform defined between a lower and upper cutoff.

For all components:
Move the lower cutoff or the upper cutoff downwards to lighten the imagery.
Move the lower cutoff or the upper cutoff upwards to darken the imagery.
Steepen the transform to lighten the imagery.
Flatten the transform to darken the imagery.

For all components:
Move the lower cutoff or the upper cutoff downwards to decrease the contrast.
Move the lower cutoff or the upper cutoff upwards to increase the contrast.
Steepen the transform to decrease the contrast.
Flatten the transform to increase the contrast.

For a color component:
Move the lower cutoff or the upper cutoff downwards to increase the color.
Move the lower cutoff or the upper cutoff upwards to decrease the color.
Steepen the transform to increase the color.
Flatten the transform to decrease the color.

For the Pan component:
Move the lower cutoff or the upper cutoff downwards to lighten the imagery.
Move the lower cutoff or the upper cutoff upwards to darken the imagery.
Steepen the transform to lighten the imagery.
Flatten the transform to darken the imagery.

Aside from moving the upper and low cutoffs, to adjust the transform you need to use either a sigmoid transform or an equal area transform.

Adjust the K factor up to steepen the Sigmoid lightening transform.
Adjust the K factor down to flatten the Sigmoid lightening transform.
Adjust the K factor up to flatten the Sigmoid darkening transform.
Adjust the K factor down to steepen the Sigmoid darkening transform.

Adjust the factor up to steepen the equal area transform.
Adjust the factor down to flatten the equal area transform.

ProRaster Product Tiers Compared

ProRaster Essential is designed for business users of the latest version of MapInfo Pro, currently MapInfo Pro 2021. If you have a MapInfo Pro 2021 subscription, or you have purchased a copy of the latest version of MapInfo Pro, you will be able to load the MRD files created by ProRaster Essential into MapInfo for visualisation. A streamlined user interface with a restricted set of capabilities provides an easy way for business users to render rasters in MapInfo Pro.

ProRaster Premium is designed for all GIS users. It provides a full set of rendering algorithm editing tools and can export rendered images to standard file formats - like geotiff - for display in a wide variety of GIS and scientific applications. If you have a MapInfo Pro 2021 subscription, or you have purchased a copy of the latest version of MapInfo Pro, you will be able to load the MRD files created by ProRaster Premium into MapInfo for visualisation

ProRaster Scientific builds upon the features of ProRaster Premium and adds powerful satellite multi-spectral imaging tools that streamline satellite multi-spectral raster data display and analysis. It also adds a range of raster processing operations.

Compared to ProRaster Premium, ProRaster Essential contains the following restrictions:

- Exporting the rendered algorithm to a located image raster is not supported. No export capability is provided.
- The Data Transform Editor is not supported.
- The Data Conditioning Editor is not supported.
- External, resizable map windows are not supported. Only the docked preview map window is supported.
- Multiple layers in an algorithm are not supported. You can load algorithms with multiple layers, but you will not be able to edit or manipulate them.
- Opacity components are not supported.
- Layer transparency and blending are not supported.
- Event selection is not supported.
- Data conditioning filters are not supported.
- Transform raster sources (basing the data-color transform on the statistics of a different raster source) are not supported.
- Format support in the GDAL driver is very limited. For example, the WMS driver is not supported, preventing use of web mapping services.

Compared to ProRaster Scientific, ProRaster Premium contains the following restrictions:

- You cannot open more than one external map window per algorithm.
- Pan and Mask components are not supported.
- The Processing menu is not available.
- The Multispectral Satellite Scene database is not available.
- The Multispectral Satellite Product database is not available.

ProRaster Scientific allows users to open multiple map windows per algorithm. It has mechanisms to create virtual rasters for satellite multi-spectral scenes that prepare the data for rendering and analysis. It has mechanisms for rendering satellite multi-spectral data using common band ratios and has a large library of band ratios and formulas for rendering analysis. It allows analysis of data over a specified polygon. It also provides a variety of processing operations, mostly implemented using virtual raster technology.

Notes on MapInfo Pro compatibility

ProRaster generates rendering algorithms (MRD files) and virtual rasters (MVR files), both of which can be opened in the latest version of MapInfo Pro. MapInfo Pro supports these files by treating them as though they are rasters and opens them into the MapInfo Pro Advanced raster engine. However, there are numerous compatibility issues that may prevent these files from opening or from being rendered properly.

In MRD files and MVR files you may encounter the following problems:

- Dozens of raster file formats that are supported by the GDAL driver are not supported in MapInfo Pro. The most glaring omissions is WebMap rasters (WMS and WMTS etc), which are not supported by the raster engine in MapInfo Pro.
- Coordinate System support and reprojection of rasters in MapInfo Pro is problematic at times. Some common coordinate systems are not well supported and reprojection between coordinate systems can cause issues, especially near the limits of the bounds.
- Support for standard data types is incomplete in MapInfo Pro and can cause failure in both MRD and MVR files.
- Driver preferences specified in raster sources may not be supported (coordinate offset, event time), or may not be implemented for all raster formats.

The following features in MRD files are not supported in MapInfo Pro:

- Mask Component (in all layer types). This extension was added to ProRaster Scientific to better support satellite multispectral data.
- Pan Component (in RGB layers). This extension was added to ProRaster Scientific to better support satellite multispectral data.
- Algorithm clipping bounds (established by data zooming manually).
- Color transform clipping beyond and between data limits.
- Data conditioning transform (Y=A.X+B) feature.
- Resolution level adjustment in components.
- The “Rotate” color data transform.

The following features in MVR files are not supported in MapInfo Pro:

- Raster masking. This extension was added to ProRaster Scientific to better support satellite multispectral data.

The following features in MRD files are not properly supported in MapInfo Pro:

- Opacity (transparency) handling. Opacity and alpha pixel values are poorly supported in MapInfo Pro. There is a specific bug where the raster handler seems to dilute the raster color with the alpha values, resulting in an invalid representation of the pixel color data.
- Opacity modulation. Whilst this is supported, the raster handler issues described above make it unusable.
- Intensity scaling in multi-layer algorithms. This can result in in appropriate intensity scaling when the layers use rasters of different cell size.
- “No transform” color transforms. The Pass, Pass Scaled and Pass Index transforms may not produce the desired effect.
- Table color transforms. Whilst supported, these transforms may not produce the desired effect.
- Range color transforms. Whilst supported, these transforms may not produce the desired effect.
- “Valid cell by component” rule “All to…” extensions are not supported.
- Outlier bin clipping. Whilst supported, these may not produce the desired effect.

- Support for rasters in ZIP file archives. Support for named rasters in ZIP file archives is not available in MapInfo Pro.
- “Ignore color alpha” in image layers. This feature is not supported.
- Bandpass clipping of color data maps. Whilst supported, these transforms may not produce the desired effect.
- GeoTIFF rasters can be mislocated half a cell to the south and east. This is due to a failure to handle the “PixelIsPoint” flag and make the appropriate adjustment. This can affect rasters produced by ARC GIS and can also affect Landsat data.

Despite this long list of issues, most MRD and MVR files will load and display properly in the latest version of MapInfo Pro. However, as raster code development and maintenance has ceased at MapInfo, the number of incompatibilities can be expected to increase over time.

Copyright

Information in this document is subject to change without notice and does not represent a commitment on the part of the vendor or its representatives. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, without the written permission of Roberts Geospatial Engineering, 7 Birdland Close, Copacabana, NSW 2251, Australia.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

© 2023 Roberts Geospatial Engineering. All rights reserved.
Contact and support information is located at: https://robertsgeospatial.com.au/

Open-Source Attribution and Copyright Notices

Products named herein may be trademarks of their respective manufacturers and are hereby recognized. Trademarked names are used editorially, to the benefit of the trademark owner, with no intent to infringe on the trademark.

zlib

Copyright (C) 1995-2022 Jean-loup Gailly and Mark Adler

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.
Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly　　Mark Adler
jloup@gzip.org　　madler@alumni.caltech.edu

IJG jpeg-9e

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-2022, Thomas G. Lane, Guido Vollbeding.
All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:
(1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
(2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".
(3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of commercial products, provided that all warranty or liability claims are assumed by the product vendor.

7-Zip LZMA

LZMA SDK is placed in the public domain.

Anyone is free to copy, modify, publish, use, compile, sell, or distribute the original LZMA SDK code, either in source code form or as a compiled binary, for any purpose, commercial or non-commercial, and by any means.

libpng

Copyright (c) 1995-2019 The PNG Reference Library Authors.
Copyright (c) 2018-2019 Cosmin Truta.
Copyright (c) 2000-2002, 2004, 2006-2018 Glenn Randers-Pehrson.
Copyright (c) 1996-1997 Andreas Dilger.
Copyright (c) 1995-1996 Guy Eric Schalnat, Group 42, Inc.

The software is supplied "as is", without warranty of any kind, express or implied, including, without limitation, the warranties of merchantability, fitness for a particular purpose, title, and non-infringement. In no event shall the Copyright owners, or anyone distributing the software, be liable for any damages or other liability, whether in contract, tort or otherwise, arising from, out of, or in connection with the software, or the use or other dealings in the software, even if advised of the possibility of such damage.
Permission is hereby granted to use, copy, modify, and distribute this software, or portions hereof, for any purpose, without fee, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated, but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This Copyright notice may not be removed or altered from any source or altered source distribution.

libgeotiff

libgeotiff © 2021 Niles D. Ritter.

EFAL

Licensing terms not supplied.

ECW JPEG2000 SDK

ECW by ERDAS © 1993-2021 Intergraph Corporation, part of Hexagon Geospatial AB and/or its suppliers. All rights reserved.

LZ4

Licenced under the BSD 3-Clause license. The license can be downloaded from http://opensource.org/licenses/BSD-3-Clause.

GDAL

GDAL © 1998-2022 Frank Warmerdam, Even Rouault, and others

GDAL/OGR is licensed under an MIT style license with the following terms:

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

https://gdal.org/license.html#
https://raw.githubusercontent.com/OSGeo/gdal/master/LICENSE.TXT

MrSID

MrSID, MrSID Decompressor and the MrSID logo are trademarks of LizardTech, a Celartem company, used under license. Portions of this computer program are copyright © 1995-2021 LizardTech, A Celartem Company, and/or the University of California or are protected by US patent no. 5,710,835 and are used under license. All rights reserved. MrSID is protected under US and international patent and copyright treaties and foreign patent applications are pending. Unauthorized use or duplication prohibited.

OpenStreetMap

© 2021 OpenStreetMap contributors, CC-BY-SA; see OpenStreetMap http://www.openstreetmap.org (license available at http://www.opendatacommons.org/licenses/odbl/index.html) and CC-BY-SA http://creativecommons.org/licenses/by-sa/2.0

ColorCET

ProRaster ships with color tables developed by the CET project. Visit https://colorcet.com/ to find out more. Refer to: Peter Kovesi. Good Colour Maps: How to Design Them. arXiv:1509.03700 [cs.GR] 2015

Web Mapping Services Terms of Use Restrictions

By creating a local service description file XML file, you can use the WMS driver in GDAL to access web mapping services. Use of these services may be limited by terms of use supplied by the vendor of the data and, as a consumer of the vendors data, you should carefully consider the legality of accessing that data in ProRaster, modifying the appearance of the data, or exporting that data from ProRaster in any form. You should actively seek out and refer to the terms of use and licensing rules of the data vendor before accessing the web map service.

You can find the Google Maps Platform Terms of Service at:
https://cloud.google.com/maps-platform/terms
https://www.google.com/help/terms_maps/

You can find the Microsoft Bing Maps Platform APIs Terms of Use at:
https://www.microsoft.com/en-us/maps/product

You can find the OpenStreetMap Foundation Services Terms of Use at:
https://wiki.osmfoundation.org/wiki/Terms_of_Use