Getting Started with montaj GridKnit

Today's explorers use regional and local geophysical compilations of magnetic, radiometric and apparent resistivity data to perform comprehensive geologic interpretations and maximize the usefulness of newly acquired and historic gridded data.

The common problem in creating compilations is determining how to stitch grids characterized by different survey parameters, noise levels, data quantity, sample densities or uncertainties at grid edges. Traditional methods (such as manually adjusting values on neighbouring grids, levelling to an existing low-resolution grid or using various weighting schemes to merge grids) can produce smooth final products but the process is time-consuming and costly.

This How-To Guide provides an overview of the montaj GridKnit™ grid stitching methods and pre-processing tools. For detailed information on working with GridKnit, read the *Working with montaj GridKnit* How-To Guide.

GridKnit Menu

The *Grid Knitting* menu can be found on the *Grid and Image* menu on the main Oasis montaj menu bar.
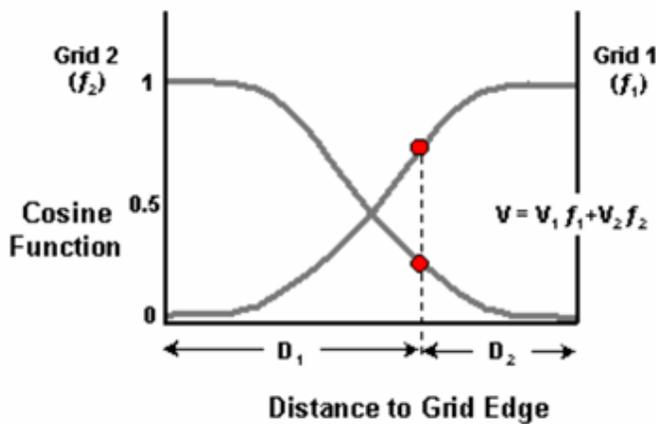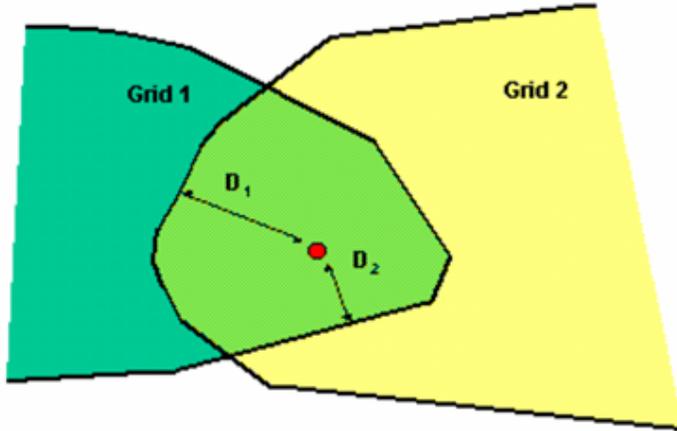
GridKnit menu
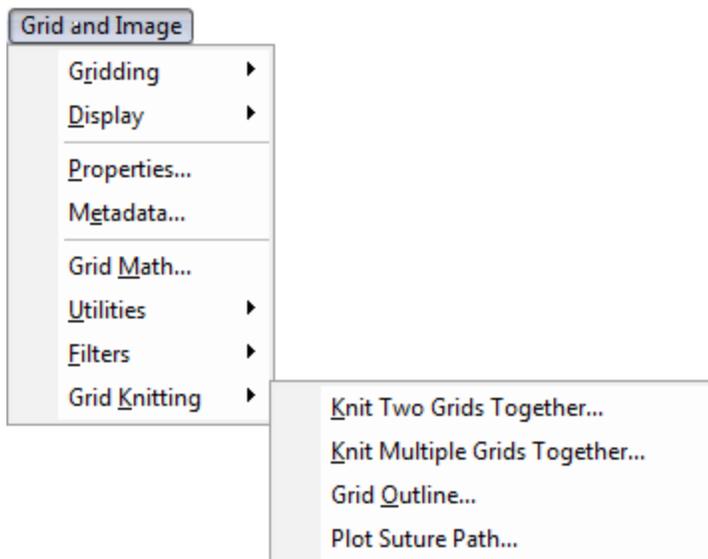
Stitch Method

GridKnit enables you to stitch together any two grids, even if the grids have different cell sizes and/or transforms. The montaj GridKnit extension delivers two advanced methods that produce fast, high-quality results for virtually any type of geophysical grid. The blending method provides a quick way of merging two grids via standard smoothing functions. The suturing method enables you to automatically or manually define a suture path for estimating errors and then applies a proprietary multi-frequency correction along the path. This method is used for joining grids that have a narrow overlap area and relatively small anomalies. Suturing provides more accurate results and greater control than blending.

Blending Method

Blending Grids: Cosine Function

Grid 1

Grid 2

D_1

D_2

Grid 2
(f_2)

1

Grid 1
(f_1)

**Cosine
Function**

0.5

$V = V_1 f_1 + V_2 f_2$

0

D_1

D_2

Distance to Grid Edge

The Blending method uses a blending function over the area of overlap so that transition from one grid to the other is smooth. Except for the optional removal of a static offset or trend, the grids beyond the overlap regions remain unchanged.

The blending function determines the weighting of one grid against the other inside the overlap region. The function works by calculating the distance to the edge of each grid for each data point. The diagram on the right shows the distance to the edge of Grid 1 (from a single data point) as D_1 and the distance to Grid 2 as D_2. The blending function uses a "cosine" function which varies smoothly from 0 to 1 — the function takes on a value of 0.5 at positions midway between two grids, and whose derivative approaches 0 at both ends. If a position is equidistant between both edges, its value is the average of the grid values found at that point. The diagram (right) shows the value (on the cosine curve) for Grid 1 is around 0.7 and the value for Grid 2 is 0.3. This means that when the blending function calculates the value for the data point, it uses 0.7 of the value of the point on grid 1 and 0.3 of the value of the point on Grid 2.

⚠️ *Where the edges of Grid #1 and Grid #2 cross at any single point, the blending scheme breaks down, since by definition both grids are full owners of that point. In this case the average of the two grid values at that point will be used.*
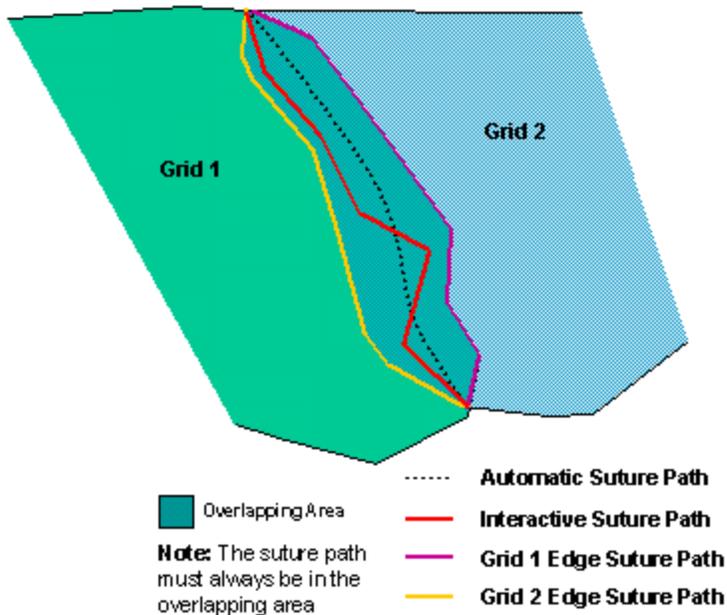
Suture Method

The suture method defines a suture line to determine where the grids will be joined. The line must lie completely within the overlapping area of the two grids. The area outside of the overlap remains unchanged.

Along the suture line, a mismatch in the grid values is corrected by adjusting the grids on either side of the path. For instance, for a point on the suture path where the value for Grid #1 is larger than the value for Grid #2, the average value is used to eliminate discrepancies. Points adjacent to the path point might then be adjusted to produce a smooth transition between the two grids.

The suture method uses a multi-frequency approach to spread corrections over the two grids in proportion to the wavelength of the mismatch found along the suture path. This ensures that the transition from one grid to the other remains smooth, regardless of the amplitude and wavelength of the features that the suture path crosses.

Locating a Suture Path

The suture line must be located inside the overlap area between the two grids. GridKnit provides four methods of determining the position of the suture line in the overlap area between the two grids.

- Automatic
- Interactive
- Grid 1 edge
- Grid 2 edge

The diagram on the right illustrates where the suture path is located using each method. The automatic suture path is located equidistant between the two grid edges. The interactive suture line is defined by the user and can be located anywhere in the overlap area. The grid 1 suture option uses the edge of Grid 1 as the suture path. Likewise, the Grid 2 option uses the edge of grid 2 as the suture path.

Automatic Suture Path

If the automatic option is selected, the suture line will bisect the overlap region; every point along the line will be at approximately equal distance from the borders (grid edges) of the overlap region. Use this method when you want to place the suture line in the middle of the two grids and use both grids equally to determine the values of the output grid.

Interactive Suture Path Definition

If the interactive option is selected, the user must define a path using the current map as a reference. (If there is no current map, one will be requested).

All suture paths must, by definition, lie within regions of overlap. The start and end points of the path within these regions is pre-determined by the geometry of the overlap, and cannot be altered. Those parts of the user-path not intersecting the overlapping grid sections are ignored. If the "start" and/or "end" points are not present in the user-path, the user-path is automatically altered to include them. It does this by joining the first and last points of the user-path within the overlap to the fixed start and end points. The resulting path may not cross any dummy positions, or any non-overlapping region of the data or an error occurs.

A single user-path may be used to define the path through two or more separate overlap sections by including at least one "outside" point as a space indicator between each path section through an overlap area. (All "outside" points are ignored when calculating the suture path from a user-path.)

Postage Stamp Stitching (Grid Edge Path)

"Postage stamp" stitching is used to incorporate a smaller grid into a larger grid; this is accomplished using the Grid Edge path selections. Selecting the Grid Edge path corresponding to that grid may preserve the contents of the entire smaller grid – the suture path will entirely enclose the smaller grid.

⚠️ *The path creation algorithm may run into difficulty when the Grid Edge path is selected if the boundaries of the grid are not sufficiently smooth. It may be necessary to smooth jagged edges, remove "peninsulas" or fill in "bays" using interpolation (see Pre-processing Algorithms) in order for the path to be determined unambiguously. The automatic option produces smoother paths due to the averaging effects of using edges from both grids simultaneously, and finding the midpoints between them.*

Suture Weighting

A value in the range 0 to 1 determines how corrections are apportioned between the grids. A cosine function is used to weight the grids. The figure below shows how the weight function works. The numbers on the side of the chart below represent the proportion of weighting for each grid. For example, to calculate the suture point referring to the grid value (dot) closest to the bottom, the weighting function would use 0.25 of grid 1 and 0.75 of grid 2. Similarly, if a value of the "Proportion of Grid 1" is 0, then all corrections are applied to Grid 1. A value of 0 indicates that all corrections are applied to Grid 2, and value of 0.5 (dot in the centre) means corrections are shared equally between both grids.

By default, the weighting is even between the grids. However, by forcing all the corrections toward one grid or the other, the user can keep one grid unaltered up to and including the suture path.

Suture Weighting
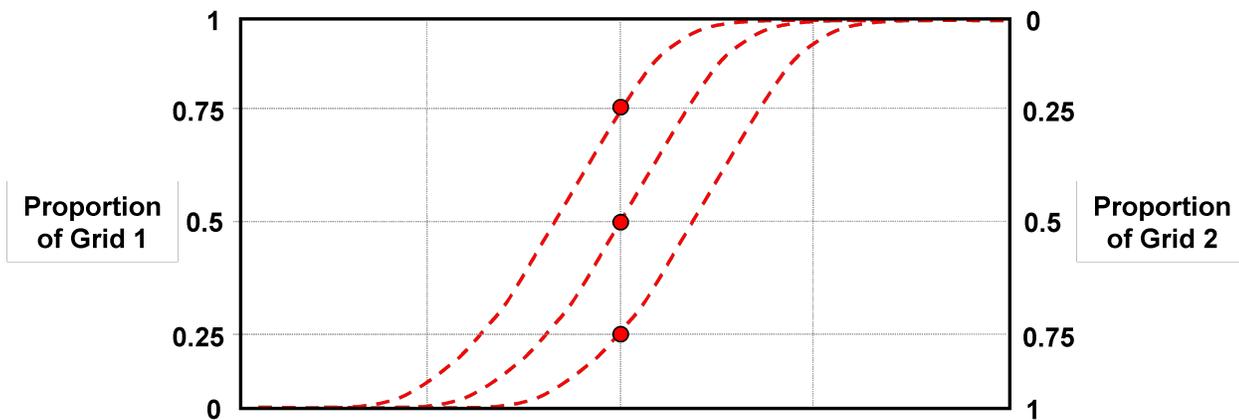
Multi-Frequency Suturing Algorithm (Fast Fourier Transform)

To ensure that the suturing process creates the smoothest possible transition between grids, GridKnit, uses a Multi-Frequency Fourier Transform (FFT) function. A Fourier transform breaks a single, complex curve into a family of 'sine' type curves, each with a unique frequency and amplitude.

The diagram below illustrates how the function works in suturing. GridKnit first selects the grid values along the suture line for each grid. The grid values (frequency) for the points from each grid along the suture path are plotted. In the diagram below, Grid 1 values along the suture line are represented by the red line and Grid 2 values by the blue line. The system calculates the difference (purple line) between the values of Grid 1 and Grid 2.
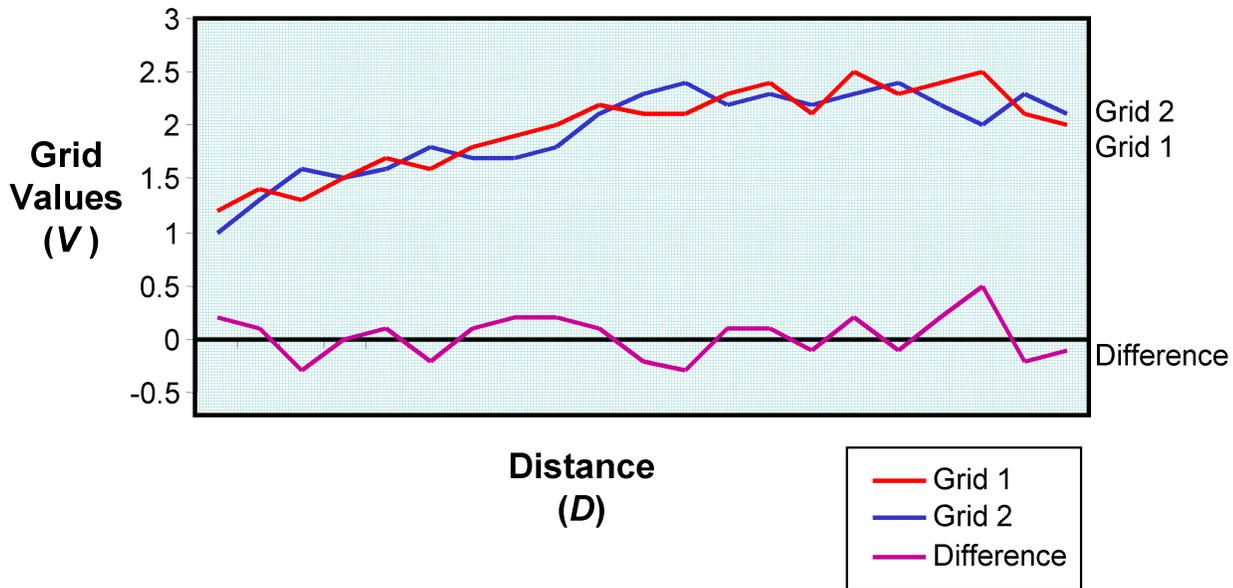
A Fourier transform function is used to split the difference curve into many curves, each representing a single frequency (or wavelength). Next, a cascading "smooth" correction surface is applied to each frequency (or wavelength) curve. The correction surface (grid) applied to each wavelength is proportional to the wavelength. This proportional correction means that narrow features on a grid are given a narrow correction and wider features are given a wider correction. The correction surface for each frequency is added to the existing overlapping grid area so that the grids join perfectly at the edges.

Pre-processing Tools

The montaj GridKnit extension offers three trend removal options, "none", "to each other" and "to a regional grid". With the "to each other" method, trends are calculated based on the two grids being stitched only, not to some third regional grid. This is the default behaviour. With the "to a regional grid" method, trends are calculated with respect to a regional grid. The regional grid must overlap the entire areas of the grids to be stitched together. The name of the grid, the detrending method used, and whether to remove the regional grid from the final product, can be specified by clicking the *Options* button.

Trend Removal "To Each Other" Method

The Trend Removal dialog enables you to remove a static shift or trend from one or both grids before stitching. Two new trend removal options, 2nd and 3rd order trends, have been added to the system. This is in addition to the previously available static (0 order) and slope (1st order). The zero-order or static correction removes a single average value from the entire grid. The first-order correction removes a best-fitting plane or slope. Higher orders fit 2nd or 3rd order polynomials in X and Y.

The trend options enable you to set a different trend order for each grid. This option is useful if you are connecting a small grid to a master grid and want to adjust the trend of the smaller grid to the master grid. In this case, you would select none for the master grid and another trending method (static or slope) for the smaller grid.

Trend Removal "To a Regional Grid" Method

The Regional Trend Removal dialog enables you to remove a regional trend based on a regional grid. The regional grid must overlap the entire areas of the grids to be stitched together.

Two new trend removal options, 2nd and 3rd order trends, have been added to the system. This is in addition to the previously available static (0 order) and slope (1st order). The zero-order or static correction removes a single average value from the entire grid. The first-order correction removes a best-fitting plane or slope. Higher orders fit 2nd or 3rd order polynomials in X and Y.

If the RMS Error Tolerance is specified, then each trend order, up to the specified order, will be tested, and only the lowest order required reducing the RMS difference between the regional and the detrended grid below the tolerance is used.

For example: If T1 is the trend calculated for the input grid G1, and TR is the trend calculated for the regional grid GR, then

RMS Tolerance > RMS Error ((GR – TR) – (G1 – T1))

or, stated another way, the input grid G1 is detrended using T1, then the trend from the regional is added. The error is now the difference between the regional and this retrended grid. If no RMS Error Tolerance is specified, then the specified trend is applied immediately.

Point Usage

The calculation for the Trend Order correction may be done in one of four ways; using all the points in the grid, using the edge points, using only the points which overlap with the other grid, or using only the edge points which overlap with the other grid.

It is not recommended to use the overlap point selection for the first-order (slope) correction in cases where the overlap region is small, or narrow in one particular direction. In these situations, the trend in a small region of the grid may be quite different from that found in the remainder and removing it may seriously alter the appearance of the full grid.

⚠️ *If the "edge points" option is selected, and there is more than one data area in the grid data, then only the first section that the algorithm encounters is used in the calculation.*

Interpolation Options

The GridKnit *Interpolation Options* dialog enables the user to set the spline method to use during interpolation and the maximum gap to interpolate across. Interpolation is used to fill gaps between points in individual grids prior to grid stitching. Trend calculation and removal are performed subsequent to any interpolation. (Interpolation is optional).

Grids may have gaps or holes inside the overlap region. The presence of gaps will affect the behaviour of the stitching methods. In the blending method, the gaps will be interpreted as grid edges, so that the contribution of the grid with the gap in it will be reduced to 0 at the gap edges. In the suture method, the automatic suture path will tend to avoid the gaps, because it chooses the path midway between the edges of the two grids. These effects may be reduced or eliminated by choosing to interpolate the gaps prior to the stitching process. By default, the interpolation value is left blank, and no interpolation is performed.

Interpolation Method

The spline method determines how values are predicted in areas without data. Three methods are available, Linear, Cubic, and Akima. The linear spline assumes a simple linear variation between the values bounding the gap. The cubic spline is a minimum-curvature method.

In cases where the data contains rapid changes in gradient, the minimum curvature spline may produce undesirable highs or lows in the gap. The Akima spline method does not suffer as much from this problem, although it does tend to produce sharper corners around actual data points and the resulting grid tends to be less smooth.

Gap Size

Grids may have gaps or holes inside the overlap region. The presence of gaps will affect the behaviour of the stitching methods. If the blending method finds a gap in the overlapping area of a grid, it fills in the gap by using the values from the other grid.

In the suture method the suture path will tend to avoid the gaps, because it chooses the path midway between the edges of the two grids. Furthermore, a long-wavelength feature lying along the suture path may be interpreted as two or more short-wavelength features, with less than optimal results. These effects may be reduced or eliminated by choosing to interpolate the gaps prior to the stitching process.

If the gap size is set to a value "x", then all gaps in the data smaller than size "x" will be interpolated prior to stitching. If the gap size is set to 0.0, all gaps will be interpolated. If the gap size is not specified, no interpolation is performed prior to stitching.

If the gap size is set to a value "X", then all gaps in the data smaller than size "X" will be interpolated prior to stitching. If the gap size is set to 0.0, all gaps will be interpolated. If the gap size is not specified, no interpolation is performed prior to stitching. Using a gap of 0.0 can produce undesirable effects if large "bays" appear in the data, including streaking in the output grid in filled-in areas. Interpolation is done across rows, and then across columns.
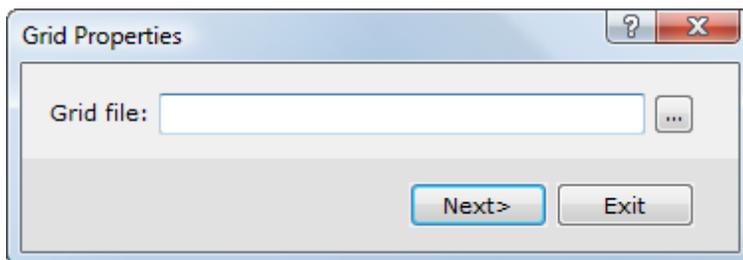
Grid Cell Size

The cell size of the output grid can be specified. If the cell size parameter is left blank, the cell size is taken from the first grid. For example, if you have two grids with different cell sizes, the cell size and transform of the first grid will have primacy and GridKnit will automatically convert the grid cell size of the second grid to match that of the first. In addition, the offsets are adjusted (even to the primary grid) to ensure that grid lines in the output overlie the zero axes. Use the Grid info dialog to determine the cell size of your grids.

To View Grid Cell Size:

1. On the *Grid and Image* menu, select **Properties**.

   The *Grid Properties* dialog appears.

2. Using the [...] button, locate the grid of interest and click the **Next** button. The *Grid Properties* dialog appears.
3. This dialog displays the basic grid information for the selected grid. The grid cell size is represented by the X point separation and Y point separation parameters.

To Change Grid Cell Size:

1. To change the cell size of a grid, on the *Grid and Image* menu, select *Gridding* and then select **Re-grid a grid**.

2. Select the *Input grid name* and specify the *Output grid name*. Modify the cell size. Click the **OK** button.

3. The input grid will be re-gridded with the new cell size.

How-To Guide Publication Date: 18/01/2013

Copyright 2013 *Geosoft Inc. All rights reserved.*