ArcMap

About georeferencing

Related topics

expand all

Raster data is commonly obtained by scanning maps or collecting aerial photographs and satellite images. Scanned map datasets don't normally contain spatial reference information (either embedded in the file or as a separate file). With aerial photography and satellite imagery, sometimes the locational information delivered with them is often inadequate, and the data does not align properly with other data you may have. Thus, in order to use some raster datasets in conjunction with your other spatial data, you often need to align it, or **georeference**—it to a map **coordinate system**. A map coordinate system is defined using a map projection (a method by which the curved surface of the earth is portrayed on a flat surface).

When you georeference your raster dataset, you define its location using map coordinates and assign a coordinate system. Georeferencing raster data allows it to be viewed, queried, and analyzed with other geographic data.

Learn about geographic coordinate systems

Learn about projected coordinate systems

Learn about supported map projections

▶ Aligning the raster

Generally, you will georeference your raster dataset using existing spatial data (target data), such as a **vector** feature class, that resides in the desired map **coordinate system**. This assumes that there are features in your spatial data that are also visible in the raster—for example, street intersections or building corners.

The process involves identifying a series of ground control points—known x,y coordinates—that link locations on the raster dataset with locations in the spatially referenced data (target data). The control points are used to build a polynomial transformation that will convert the raster dataset from its existing location to the spatially correct location. The connection between one control point on the raster dataset (the “from point”) and the corresponding control point on the aligned target data (the “to point”) is called a link.

The number of links you need to create depends on the complexity of the polynomial transformation you plan to use to transform the raster dataset to map coordinates. However, adding more links will not necessarily yield a better registration. If possible, you should spread out the links over the entire raster dataset rather than concentrating them in one area. Typically, having at least one link near each corner of the raster dataset and a few throughout the interior produces the best results.

Generally, the greater the overlap between the raster dataset and target data, the better the alignment results, because you’ll have more widely spaced points with which to georeference the raster dataset. For example, if your target data only occupies one quarter of the area of your raster dataset, the points you could use to align the raster dataset would be confined to that area of overlap. Thus, the areas outside the overlap area are likely not properly aligned.

Keep in mind that your georeferenced data is only as accurate as the data to which it was aligned. To minimize errors, you should georeference to data that is at the highest resolution and largest scale for your needs.

▶ Transforming the raster

When you've created enough links, you can transform—or warp—the raster dataset to permanently match the map coordinates of the target data. Warping uses a polynomial transformation to determine the correct map coordinate location for each cell in the raster.

Use a first-order—or affine—transformation to shift, scale, and rotate your raster dataset. This generally results in straight lines on the raster dataset mapped as straight lines in the warped raster dataset. Thus squares and rectangles on the raster dataset are commonly changed into parallelograms of arbitrary scaling and angle orientation.

A first-order transformation will probably handle most of your georeferencing requirements. With a minimum of three links, the mathematical equation used with a first-order transformation can exactly map each raster point to the target location. Any more than three links introduces errors, or residuals, that are distributed throughout all the links. However, you should add more than three links because if one link is positionally wrong, it has a much greater impact on the transformation. Thus, even though the mathematical transformation error may increase as you create more links, the overall accuracy of the transformation will increase as well.

The higher the transformation order, the more complex the distortion that can be corrected. However, transformations higher than third order are rarely needed. Higher-order transformations require more links and thus will involve progressively more processing time. In general, if your raster dataset needs to be stretched, scaled, and rotated, use a first-order transformation. If, however, the raster dataset must be bent or curved, use a second- or third-order transformation.

Interpreting the root mean square error

The degree to which the transformation can accurately map all control points can be measured mathematically by comparing the actual location of the map coordinate to the transformed position in the raster. The distance between these two points is known as the residual error. The total error is computed by taking the root mean square (RMS) sum of all the residuals to compute the RMS error. This value describes how consistent the transformation is between the different control points (links). Links can be removed if the error is particularly large, and more points can be added.

While the RMS error is a good assessment of the accuracy of the transformation, don’t confuse a low RMS error with an accurate registration. For example, the transformation may still contain significant errors due to a poorly entered control point.

Resampling the raster dataset

While you might think each cell in a raster dataset is transformed to its new map coordinate location, the process actually works in reverse. During georeferencing, a matrix of "empty" cells is computed using the map coordinates. Then, each empty cell is given a value based on a process called resampling.

The three most common resampling techniques are nearest neighbor assignment, bilinear interpolation, and cubic convolution. These techniques assign a value to each empty cell by examining the cells in the ungeoreferenced raster dataset. Nearest neighbor assignment takes the value from the cell closest to the transformed cell as the new value. It’s the fastest resampling technique and is appropriate for categorical or thematic data. Bilinear interpolation and cubic convolution techniques combine a greater number of nearby cells (4 and 16, respectively) to compute the value for the transformed cell. These two techniques use a weighted averaging method to compute the output transformed cell value and thus are only appropriate for continuous data such as elevation, slope, aerial photography, and other continuous surfaces.

Resampling is part of the georeferencing process; however, it can also be used to adjust the cell size of a raster dataset or can be part of a geoprocessing operation to alter the values of cells.

Should you rectify your raster?

You can permanently transform your raster dataset after georeferencing it by using the Rectify command on the Georeferencing toolbar. Rectify creates a new raster dataset that is georeferenced to map coordinates. You can save this in GRID, TIFF, or IMAGINE format.

ArcMap doesn’t require you to rectify your raster dataset to display it with other spatial data. You might choose to rectify your raster dataset if you plan to perform analysis with it or want to use it with another software package that doesn’t recognize the external georeferencing information created by ArcMap.