

Understanding Compression of Geospatial Raster Imagery

Document Overview

This document was created for the North Carolina Geographic Information and Coordinating Council (GICC), <http://ncgicc.com>, by the GIS Technical Advisory Committee (TAC). Its purpose is to serve as a best practice or guidance document for GIS professionals that are compressing raster images.

This document only addresses compressing geospatial raster data and specifically aerial or orthorectified imagery. It does not address compressing LiDAR data.

Compression Overview

Compression is the process of making data more compact so it occupies less disk storage space. The primary benefit of compressing raster data is reduction in file size. An added benefit is greatly improved performance over a network, because the user is transferring less data from a server to an application; however, compressed data must be decompressed to display in GIS software. The result may be slower raster display in GIS software than data that is not compressed. Compressed data can also increase CPU requirements on the server or desktop.

Glossary of Common Terms

- **Raster** is a spatial data model made of rows and columns of cells. Each cell contains an attribute value identifying its color and location coordinate. Geospatial raster data like satellite images and aerial photographs are typically larger on average than vector data (predominately points, lines, or polygons).
- **Compression** is the process of making a (raster) file smaller while preserving all or most of the data it contains. Imagery compression enables storage of more data (image files) on a disk than if they were uncompressed.
- **Compression ratio** is the amount or degree of reduction of an image's file size. Expressed as the ratio of its original size to the target size: for example if a file was 20 MB and after compression it's 2 MB that equates to a 10:1 compression ratio.
- **Lossless**. True numerical lossless is a class of data compression that allows the original data to be identically reconstructed from the compressed data. Some image file formats, like PNG or GIF, use only lossless compression. Lossless compression is used in cases where it is important that the original and the decompressed data be identical, or where deviations from the original data could be deleterious.
- **Lossy** compression is the class of data encoding methods that uses inexact approximations (or partial data discarding) for representing the content that has been encoded. Such compression techniques are used to reduce the amount of data that would otherwise be needed to store, handle, and/or transmit the represented content. Using well-designed lossy compression technology, a substantial amount of data reduction is often possible before the result is sufficiently degraded and noticed by the user.
- **Decoding** is the process of decompressing the data so it can be viewed. Generally a user of compressed data needs software that can view the raster image.
- **GDAL (Geospatial Data Abstraction Library)** is a library for reading and writing raster geospatial data formats, and is released under the permissive X/MIT style free software license by the Open Source Geospatial Foundation.

Understanding Compression of Geospatial Raster Imagery

Raster Data Compression

The nominal size of a raster image (in bytes) is its height times its width, times the number of samples, times the number of bytes per sample.

For example, the nominal size of an 8-bit RGB image measuring 10,000 x 10,000 is calculated as follows:

10,000 (height) x 10,000 (width) x 3 (samples) x 1 (byte) = 300,000,000 bytes = 300 MB

A compression ratio is the amount or degree of reduction in an image's file size. Compression is expressed as the ratio of its nominal size to the target size (the output size of the file after compression). Compressing an image file at a ratio of 20:1, means that the target size is 5 percent of the nominal size.

Compressed at a ratio of 20:1, the size of the above image is:

$300,000,000 / 20 = 15,000,000 = 15 \text{ MB}$

Compression of raster files is usually done by using an image file compression algorithm. A compression algorithm involves encoding information using fewer bits than the original uncompressed image. There are two types of image file compression algorithms: lossless compression and lossy compression. Both are explained below.

Lossless Compression

Lossless compression is compression without any loss of data quality. The lossless compression algorithm eliminates only redundant information, while preserving a perfect copy of the original uncompressed image.

Lossless compression is used when it is critical that all bits of the original and the decompressed data be identical. This is the case for archival storage, as well as for uncommon workflows where no possible loss of precision is ever acceptable. This is done by rewriting the data in a more space efficient way, removing all kinds of repetitions.

Numerical example:

An example of seven gray pixels:

128, 127, 126, 121, 124, 123, 120

Can be re-written in shorter numbers requiring less bits like:

128, -1, -1, -5, +3, -1, -3

Examples of LOSSLESS METHODS are:

- Lempel-Ziv-Welsh (LZW) method (.tif)
- MrSID 2:1 Compression method (.sid)

The TAC recommends that lossless compression be used to generate a "master" image from which other derivative images will be made.

When using LizardTech software to create MrSID (.sid) formatted raster files the level of compression generally has a maximum compression ratio of 2:1. This level of compression typically yields numerically lossless compression. The size of the compressed file will vary based on the number of redundant information.

Understanding Compression of Geospatial Raster Imagery

Lossy Compression

Lossy compression is compression with some loss of data quality. The lossy compression algorithm eliminates redundant information as well as irrelevant information and permits only an approximate reconstruction of the original uncompressed image.

Lossy compression is also done by rewriting the data in a more space efficient manner, less important details of the image are manipulated or even removed so that higher compression rates are achieved, therefore reducing file size. Lossy compression is dangerously attractive because it can provide compression ratios of 100:1 to 200:1, depending on the type of information being compressed, but the cost is loss of data.

Numerical example:

An example of seven gray pixels:

128, 127, 126, 121, 124, 123, 120

Can be re-written like:

128 – 6

Result after decompression:

128, 127, 126, 125, 124, 123, 122

Examples of LOSSY METHODS are:

- JPEG method (.JPG)
- JPEG 2000 method (.jp2)
- MrSID Compression method (.sid) (anything above a 2:1 compression ratio)

The advantage of lossy methods over lossless methods is that in some cases a lossy method can produce a much smaller compressed file than any known lossless method while still meeting the requirements of the application.

NC 911 Board's Orthoimagery Example

Since 2010 the North Carolina 911 Board has funded orthophotography projects. In 2010, the Board funded a statewide aerial imagery project and each county received aerial imagery. This ambitious project was a significant learning experience and the State received a lot of comments and questions from local governments on compression ratios.

Starting in 2012, the 911 Board started funding regional projects with approximately a quarter of the State being flown each year. Each county receives imagery in three formats: 1. uncompressed tiff images by tile (largest dataset) 2. The same tiles in .sid format with a 20:1 compression, and 3. A MrSID mosaic compressed to 50:1.

The compression ratio chosen for the image tiles provides an acceptable balance between file size and performance when loaded into a desktop GIS. The higher compression ratio of 50:1, which is used for mosaics, results in a smaller file size which aids in data transfer, e.g. copying to a computer to be used in the field with no internet connection, or copying mosaics between servers, computers, or hard drives.

Understanding Compression of Geospatial Raster Imagery

The mosaics are processed using flat formatting rather than a composite formatting. The flat formatting creates a larger file size (approximately 50% larger) but draws or displays much faster in GIS desktop software. The general consensus is the performance benefits outweigh the increased size.

Archival of Raster Data

While compression is a necessary and useful process for facilitating the interchange of raster images, it is not appropriate for the long term preservation of the raster data. In order to ensure that the original record is accessible for future use, the State Archives of North Carolina requires transfers of geospatial raster data for permanent storage to be in an uncompressed TIFF format with either embedded GeoTIFF metadata or ISO format metadata. As open source conversion applications or programming libraries for JPEG 2000 become more robust, this requirement may be modified to allow lossless JPG 2000 transfers.

Summary

Utilizing compression techniques is an accepted practice to reduce file size of, and therefore effectively manage, large amounts of geospatial imagery. It is important to understand the results that can be expected when compressing imagery. This document was created to help convey that and guide these decisions. Practitioners should take into consideration their workflows and planned use of the imagery before compressing. Trial and error may have to ensue in order to achieve the most desirable product for the intended project or application.

Understanding Compression of Geospatial Raster Imagery

References

- ArcGIS Resource Center
<http://resources.arcgis.com/en/help/main/10.2/index.html#/009t00000021000000>
- Geotiff specification
<http://trac.osgeo.org/geotiff>
- GDAL image format information including maximum file sizes
http://www.gdal.org/formats_list.html
- Jpeg2000 specification
<http://www.jpeg.org/jpeg2000>
- LizardTech's MrSID Technology
<http://www.lizardtech.com>
- Speed comparisons
<http://gis.stackexchange.com/questions/14262/speed-of-various-raster-data-formats>
- Wikipedia
<https://www.wikipedia.org>
- Geospatial Data Abstraction Library
<http://www.gdal.org>

Understanding Compression of Geospatial Raster Imagery

Appendix 1 - Proprietary Compression Formats (MrSID, ECW)

Multi-resolution seamless image database (MrSID) is a proprietary compression format from LizardTech that enables rendering and manipulation of imagery without sacrificing image quality. MrSID enables compression ratios, true multiple resolutions, selective decompression, seamless mosaicking, and improved image browsing and manipulation.

The nominal size of an 8-bit RGB image measuring 1000 x 2000 is calculated as follows:

$1000 \text{ (width)} \times 2000 \text{ (height)} \times 3 \text{ (samples)} \times 1 \text{ (byte)} = 6,000,000 \text{ bytes} = 6 \text{ MB}$

A compression ratio is the amount or degree of reduction in an image's file size, expressed as the ratio of its nominal size to the target size (the intended size after compression). Compressing an image at a ratio of 20:1 means that the target size is 5 percent of the nominal size.

Compressed at a ratio of 20:1, the size of the above image is:

$6,000,000 / 20 = 300,000 = 300 \text{ KB}$

According to LizardTech MrSID compression ratios of 2:1 are lossless. This means only half the storage space is needed and yet the numerically identical original data is still retained.

For further storage savings, MrSID technology's lossy compression can yield typical ratios of up to 20:1 while still offering a level of image quality that makes data indistinguishable to the eye for most datasets. LizardTech considers compression ratios up to 20:1 to be "visually lossless." At 20:1 there is numerical data loss, but with respect to what the image is being used for, the amount of loss is considered imperceptible.

Higher compression ratios are possible but are lossy. Depending on how much image quality needed to retain, and depending on the type of original imagery, ratios of 40:1 and beyond can be used.

The MrSID format has been extended to MrSID Generation 4 (MG4) to support multispectral imagery, enabling users to compress 4-band NAIP data, 8-band Landsat data or even 224-band AVIRIS data.

When using LizardTech software to create MrSID (.sid) formatted raster image files, the level of compression generally has a maximum compression ratio of 2:1. This level of compression typically yields numerically lossless compression. The size of the compressed file will vary based on the number of redundant information.

Flat vs. Composite Mosaic when using MrSID

The MrSID mosaic format is a seamless collection of multiple raster inputs. A mosaic is an efficient and practical means for loading a large physical volume and quantity of images into a single manageable file or layer. The three most important factors to consider when encoding a mosaic are scale, resolution, and file size.

When creating a mosaic of multiple images in MrSID file format, there are two output formats, flat or composite, that may be used to create the resulting image. A composite mosaic indexes the separate, original images to create the combined image. Images retain the original tile structure, but are contained in a single file. A flat mosaic combines all input tiles into a single output image without any reference to the original tiles. Regardless of format, it is up to the user to determine the output compression ratio based upon their requirements.

Users should note that the two mosaic types are different in terms of compression statistics. Flat format statistical output for compression will report what the user supplied as the output target compression. Composite statistical reporting could be misleading because it is based on average compression per input file. For example if a target

Understanding Compression of Geospatial Raster Imagery

compression at 50:1 is set but the rectangular bounding extent of the output contains a significant amount of “No Data”, statistics will yield a compression ratio greater than 50:1. The reason being that “No Data” is highly compressed and therefore contributes to the overall statistical average.

In summary, users of compressed raster files have trade-offs in terms of file size, processing time, and performance as follows:


Flat mosaic:

- Larger file size
- Longer process time to create
- Quicker rendering or display time in end-user software


Composite mosaic:

- Smaller file size
- Quicker process time to create
- Slower rendering or display time in end-user software

Below are example images using MrSID compression at different scales:



Understanding Compression of Geospatial Raster Imagery



Understanding Compression of Geospatial Raster Imagery



Understanding Compression of Geospatial Raster Imagery



Understanding Compression of Geospatial Raster Imagery



Understanding Compression of Geospatial Raster Imagery



ECW Format

In order to accommodate all of the aerial imagery each year for quality control, it is necessary to convert the images to a smaller size for storage. Since NC Geodetic Survey does not have a copy of MrSid for conversion and there are free options available, the ECW format was chosen.

Using a 20:1 compression ratio results in the following:

OC6i0_37_000_20069302_20130220_0328R0.tif	Tiff format	293,048 kb
OC6i0_37_000_20069302_20130220_0328R0.ecw	ECW format	11,515 kb

Renaming the tiff world file (.tfw) extension to .ecw allows georeferencing of the ECW file.

IrfanView (<http://www.irfanview.com>) can be used to perform batch conversions on the tiff files and there are many file explorer variants that will allow batch extension renaming.

Samples are shown below.

Understanding Compression of Geospatial Raster Imagery

20:1 ECW at 100% zoom



Understanding Compression of Geospatial Raster Imagery

1:1 Tiff at 100% zoom



Understanding Compression of Geospatial Raster Imagery

20:1 ECW at 214% zoom



Understanding Compression of Geospatial Raster Imagery

1:1 Tiff at 214% zoom



Understanding Compression of Geospatial Raster Imagery

Appendix 2 - Open Compression Methods (LZW, Deflate (LZ77), JPEG, JPEG 2000)

The GDAL (<http://www.gdal.org>) raster library has been incorporated into several software projects, both proprietary and open source, see <https://trac.osgeo.org/gdal/wiki/SoftwareUsingGdal>. The GDAL library and utility programs have several options for image compression, both lossless, and lossy. The `gdal_translate` utility, http://www.gdal.org/gdal_translate.html, can be used to compress geotiff images with lossless compression internally (LZW, Deflate, Packbits) or it can be used to convert to lossless JPEG2000 format. To get the best lossless compression with geotiff requires some research into the command line options, see <http://linfiniti.com/2011/05/gdal-efficiency-of-various-compression-algorithms/>. Some compression methods are more suited to aerial imagery than others.

GDAL lossless compression methods efficiency:

Compression Method	File Size
Uncompressed Tiff	293 MB
Packbits	294 MB
LZW	179 MB
Deflate	162 MB
JPEG2000 (Lossless)	167 MB
JPEG quality 100 Tiff	158 MB

There are also lossy compression options using `gdal_translate`. The most commonly used compressions are JPEG-compressed Tiff and Jpeg2000 lossy. The jpeg compressed tiff quality level can be controlled with the default level of 75 and additional compression can be achieved by using the `PHOTOMETRIC` option.

GDAL Lossy compression methods efficiency:

Compression Method	File Size
Uncompressed Tiff	293 MB
Jpeg Compressed Tiff (default + photometric)	21 MB
JPEG2000 (default options)	73 MB

Geotiff files can be tiled internally for greater reading speed, and overviews can be added to the geotiff images using the “`gdaladdo`” command, <http://www.gdal.org/gdaladdo.html>, to speed reading at different scales of display.

Understanding Compression of Geospatial Raster Imagery

The GDAL library reads and writes over 100 image formats as commonly compiled, see http://www.gdal.org/formats_list.html, and can be compiled with external licensed libraries to read/write proprietary formats MrSID, ECW , and more efficient versions of JPEG2000.