A Fully-Automated High Performance Image Registration Workflow to Support Precision Geolocation for Imagery Collected by Airborne and Spaceborne Sensors

D. A. White[1], C. R. Davis[1]

[1]Geographic Information Science & Technology Group
Oak Ridge National Laboratory
1 Bethel Valley Road Knoxville, TN 37831
Telephone: +1 865-574-7647
Email: {whiteda1,daviscr}@ornl.gov

1. Introduction

Deriving precise coordinates from airborne and spaceborne imagery, with uncertainty estimates, is a very challenging and time consuming task. That task is significantly more difficult to accomplish when the imagery comes from one or more sensors that have questionable and/or incomplete photogrammetric metadata. Before precision geolocation activities can take place, photogrammetric metadata must be complete and consistent such that images are correctly registered to one another and the Earth's surface.

For projects with a limited number of images from a small number of sensors, a manual registration solution executed through commercially-available desktop software is appropriate and recommended. However, what happens when you are dealing with a very large and rapid stream of images from a wide variety of sensors with differing modalities and wildly varying photogrammetric metadata? An automated, high performance image registration workflow is being built at Oak Ridge National Laboratory to meet this need. This paper describes the core concepts underlying its creation and discusses all of the steps in the automated workflow and the software being used to accomplish them.

2. Core Development Concepts

The automated workflow is being built around several core concepts: (1) Develop required applications using only open source, in-house, or government-furnished software, (2) leverage well-established photogrammetric and computer vision techniques to reduce risk, (3) expose all components as services that can communicate with one another, (4) use high performance computing architectures and paradigms wherever possible, (5) strictly adhere to and take full advantage of metadata standards for the National Imagery Transmission Format (NITF) to enable interoperability, and (6) keep the system as flexible as possible through extensive use of plugin frameworks. These concepts are at the heart of the workflow described below and drive every aspect of the system's creation.

3. Registration Workflow

The automated workflow consists of five steps: trusted source selection, global localization, image registration, sensor model resection and uncertainty propagation, and enhanced metadata generation, as can be seen in fig 1. Each step builds off of the preceding one and will be discussed in turn.

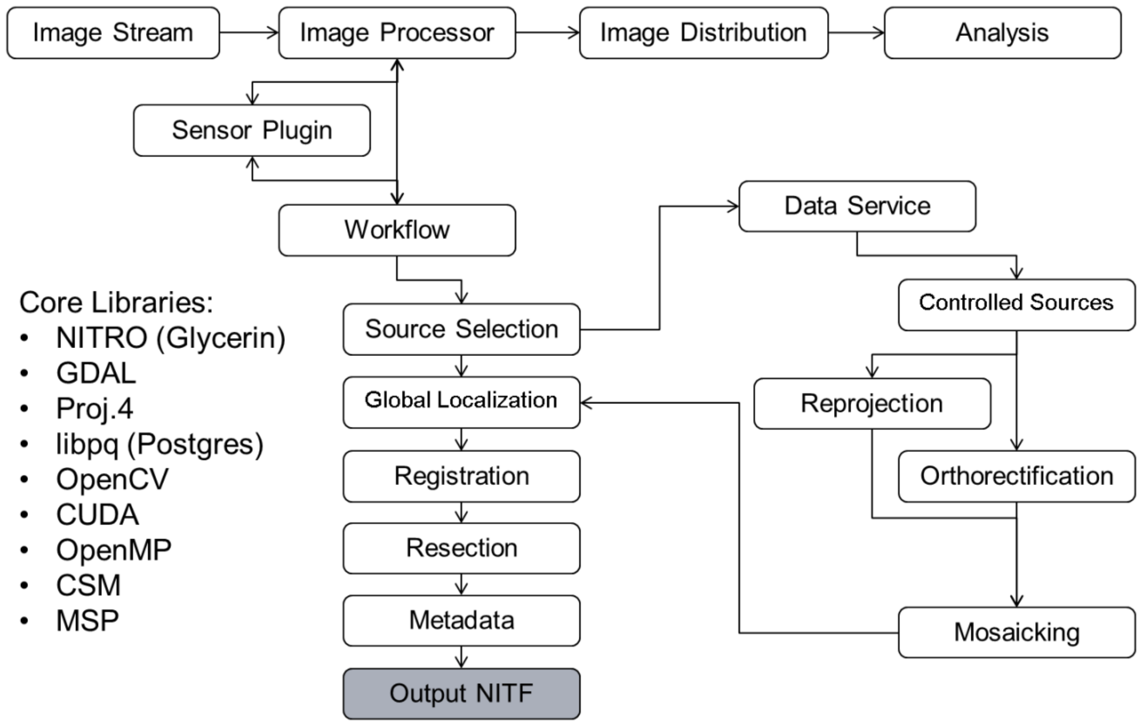

Figure 1. The automated image registration workflow.

3.1 Trusted Source Selection

Trusted source selection is the process by which the initial estimate of an image's location, based on available photogrammetric metadata, is used to determine which parts of which sources of controlled data will be used to aid in the registration process. Controlled sources consist of well-characterized orthoimagery, elevation, and Earth surface deformation (geoid). We constructed a PostgreSQL+PostGIS database, which consists of a series of tables that contain pointers to repositories of tiled versions the sources as well as polygon footprints. Using a spatial bounding box query to the database, buffered by a tuneable percentage to account for incorrect initial image placement, a list of pointers to specific tiles are returned for each source, which are then used to build mosaics on the fly. Those mosaics are used as inputs in the next two steps. MongoDB and Apache Spark are currently being evaluated as database alternatives.

3.2 Global Localization

Global localization is the process by which a reduced resolution version of a problematic image is matched against a reduced resolution control orthoimage to quickly ascertain a reasonable guess of actual location and to do a coarse update of available photogrammetric metadata. The rough guess is then used to reduce the footprint of the control data and thus restrict the search space during full-resolution image registration. The problematic image is first orthorectified using an OpenMP-enabled, parallelized, cross-platform C++ application that leverages GDAL, Proj.4, and the elevation and geoid control data returned during trusted source selection. It is then matched against the control orthoimage using one of three techniques called through a plugin framework built on top of OpenCV: patch-based normalized cross-correlation, point-based ORB, or patch-based normalized mutual information (ported from Insight Toolkit)—depending on modality. Hardware acceleration via CUDA is employed heavily to reduce processing time. The coarse update is accomplished through a traditional resection process (see below).

Both the resection and orthorectification processes rely upon the presence of a sensor model. To keep this element generic, the Community Sensor Model plugin framework is used. That framework provides a common set of widely-used photogrammetric functions exposed through a single API. An application using that framework does not have to know or care what sensor it is dealing with.

3.3 Image Registration

Image registration is a significantly more complex and compute intensive version of global localization. If necessary, the spatial resolution of one or both images is adjusted so that both begin on an even footing. Then image pyramids are built for both so that matching can take place at multiple scales to boost the success rate.

The registration process deviates from global localization in several respects. First, due to potentially large differences in viewing geometry between the problematic image and the control orthoimage, the latter is projected up into the coordinate space of the former using the sensor model, elevation data, and geoid data. This turns a 3D problem into a 2D problem through perspective matching. Second, both images are broken up into tiles to support faster matching through distributed computing and hardware acceleration as well as to make the matching problem local instead of global. Third, the end goal of the matching process is to produce a large set of match points that can be used during resection, so additional work must be done when employing patch-based matching algorithms instead of point-based ones. For the former, interest point operators like Harris must be used within a pair of tiles to find common points. The registration process does leverage the same hardware-accelerated image matching capabilities employed during global localization.

Once a large match point list has been generated across all scales, the sample/line coordinates of those points in the problematic image and the longitude/latitude/height above ellipsoid coordinates of those points in the control image (calculated by projecting them back onto the ground via the sensor model) are supplied to the resection process to more accurately refine the sensor model's behavior. Using rigorous sensor models throughout the registration process, within a hardware-accelerated environment, is to our knowledge a novel approach.

3.4 Sensor Model Resection and Uncertainty Propagation

Resection involves the application of unified least squares (ULS) to make changes to a small set of sensor model parameters so that the model has the best agreement possible with a set of user-supplied control points. These points are usually generated through a manual selection process to ensure high quality, but a spatially bounded and localized matching process that is automated produces enough viable points that resection can take place. This is the same approach used to create structure from motion. The resection process generates parameter-specific uncertainty estimates as a by-product and those estimates can be translated from sensor space to ground space in order to compute uncertainties for real-world coordinates.

3.5 Enhanced Metadata Generation

The remaining challenge is capturing the behavior of the adjusted sensor model, with its improved geolocation accuracy and associated spatial uncertainty estimates, and storing it in a standardized way that can be leveraged by a large number of desktop applications—all while leaving the original metadata intact to preserve its pedigree. To accomplish this, imagery is read from and written to NITF, a tightly-regulated hierarchical format that permits the inclusion of extensive metadata and multiple images in the same file.

There are standardized metadata containers to support many types of information, but the ones of chief interest to this project are IGEOLO, BLOCKA, RPC00B, RSM, and SENSRB. The first two describe a four-vertex polygon that covers the footprint of the image in a real-world coordinate system, with the latter being able to do so with higher fidelity. RPC00B and RSM are two forms of replacement sensor models, which for the purposes of this paper are generic sensor models built by fitting observations from the actual model to a set polynomials that have either a fixed (RPC) or variable (RSM) number of terms. Both can communicate spatial uncertainty, but RSM does so more robustly. With that said, more software packages support RPC, so both are generated.

Underlying all of the NITF components of the project is an open source library called NITRO, which can handle all aspects of reading and writing these very complicated files. We accelerated performance for multi-threaded and distributed architectures using OpenMP, including the I/O steps, which can take a significant amount of time for large images. We call this new high performance version Glycerin.

4. Conclusion

When precision geolocation has to be done on a very large scale and at a very rapid pace, leveraging imagery that may have only a limited amount of photogrammetric metadata, traditional manual solutions break down. An automated, high performance registration workflow is being built to address this need. It is adhering to an open source philosophy, using well-known techniques, and following well-established standards to ensure the maximum possible level of utility.

5. Acknowledgements

Prepared by Oak Ridge National Laboratory, P.O. Box 2008, Oak Ridge, Tennessee 37831-6285, managed by UT-Battelle, LLC for the United States Department of Energy under contract number DEAC05-00OR22725.