Georeferencing of Historical Imagery

Collin Giguere; cdgiguere@umass.edu
Sowmya Vasuki Jallepalli; sjallepalli@umass.edu
Data Science for the Common Good Fellows, 2021

Motivation

Aerial images of the earth taken from drones, aircraft, or planetary satellites, can provide rich insight into how lands and waterways change over time due to human development or naturally occurring processes and events. In order for these images to be useful, each pixel needs to be associated with a precise latitude and longitude point. This can be done via a process known as georeferencing. Unfortunately, adding georeferencing information to an aerial image is a time consuming task that today can only be performed by trained human analysts. This manual process involves identifying points which have remained consistent between two images, the aerial image and an already georeferenced image. These points are called Ground Control Points or GCPs. GCPs are used to warp these images and propagate the geospatial information to the rest of the pixels on the image.

Analysing land use has a wide range of applications in regional planning, creating maps of regions, and studying historical changes from one point in time to another. In this project, performed in partnership with UMass Libraries, Department of Environmental Conservation and advised by Pixel Forensics Inc., our challenge was to develop a fully automated approach to georeferencing aerial images by comparison to fully georeferenced commercial satellite imagery.

Related Works

W. Tabib, "A SIFT Matching Approach for Initial Registration," December, 2013. [Online] http://www.cs.cmu.edu/afs/cs.cmu.edu/academic/class/15463-f13/www/final_proj/www/wtabib/iccppaper_final.pdf

Dataset

We were given a corpus of nearly five thousand aerial images covering the entirety of the Commonwealth of Massachusetts taken 70 years ago; this was known as the MacConnell Aerial Images Collection. Just over 400 of them had been manually georeferenced by students already. Important information is missing about how these photos were taken, such as the height at which the plane was flying and the type of camera used. The original images are not even available! Only scans of prints, most of which have been manually marked up with land use information. We compared these images to 2019 USGS aerial photos from MassGIS.

We were also given one latitude/longitude point that was somewhere near the center of each MacConnell image, as well as the flight path of the imaging platform. One unexpectedly significant attribute of the images is that adjacent images overlap by up to 50%.

Approach

To find our potential GCPs, we first adjust for color differences between the pseudo-grayscale historical images (scanning introduces color anomalies) and the modern colored images. We apply histogram matching to adjust the color intensity of the modern images to the historical images, then both are converted to true grayscale. The historical images are split into 16 tiles and matched to a grid of 9 aerial images in a 4-1 mapping to spread matches evenly across the historic image. This mapping is precomputed via the points known to be in each historical image.

We detect keypoints in the four historical tiles and the single aerial image using Harris Corner Detection. Next, SIFT is applied to compute descriptors for each keypoint. This yields 128 length vectors to describe each keypoint. A brute force matching algorithm is used to determine the best matches of descriptors, or the closest vectors by euclidean distance, for keypoints between the two images. Finally, the RANSAC algorithm is applied to separate inliers from outliers in the matches and the inliers are recommended as GCPs.

Results

We delivered two image processing pipelines to our partners. One completes the original task; georeferencing a historic image via comparison to fully georeferenced commercial satellite images. This approach gives an average error of about 100m displacement from the manual approaches mapping, which is generally too large to allow the system's recommendations to be used for georeferencing without human review. Instead, a human analyst will quickly inspect these automatically generated pairs of points, and decide whether or not to adjust it to line up correctly, or scrap the pair altogether.

We also discovered that we could run our pipeline on pairs of overlapping historic images, where one image was already georeferenced. Our pipeline then could recommend GCPs in the overlapping area, and do so with an average accuracy of under 30 meters! This was a great success and well within the acceptable range of error for automatic georeferencing without human review. The pipeline for this approach varies slightly from the original only in that it does not split either image at all; matches are computed for descriptor-keypoint pairs across the entirety of both images. This pipeline can be applied to any pair of overlapping aerial images, not just the ones in the MacConnell set.