


# Georeferencing an Aerial Image of Point Mugu State Park in Ventura County, CA

by Liz Baumann  
Fall 2015, Final Project  
CVEN 5382, GIS Spatial Database Development

## Table of Contents

Introduction.....	2
Objectives.....	2
Processing.....	2
1. Establishing control features.....	2
2. Georeferencing image to control.....	3
3. Adding more features in other coordinate systems.....	5
Considerations.....	6
Conclusions.....	6
Acknowledgements.....	7



## ***Introduction***

For this study I chose to examine the area represented by an aerial image of a portion of Ventura County centered on Point Mugu State Park, located about 30 miles west of Los Angeles, California. This region includes a variety of terrain, with elevation ranging from sea level to over 3000 feet. The state park occupies most of the area and is bounded on the south by the Pacific Ocean. Visible in the image just west of the park lies the Oxnard Plain, a flat farming region near sea level; and on the northeast rising above the plain lies the Conejo Valley city of Newbury Park, CA. The Santa Monica Mountains stretch eastward through Point Mugu State Park towards Los Angeles.

In this paper I will describe the procedures I used to transform the raw aerial image I found online of this area to a georeferenced raster that can be used in ArcMap alongside other spatial data, including data referenced with different coordinate systems.

## ***Objectives***

My goals for this project were to apply ArcGIS georeferencing tools to an image and use appropriate coordinate system alignment procedures to bring other map layers with different coordinate systems in to a common map document. The intent was for this exercise to be representative of practical situations I might encounter outside of class. I would find a raw aerial image online, bring it into ArcGIS, and process this image so it aligned with and could be used analytically or visually alongside other layers in an ArcMap document. I would employ techniques learned in the GIS Spatial Database Development course, focusing primarily on geoprocessing and working with data from different coordinate systems.

## ***Processing***

I divided my processing into 3 steps:

1. Establishing control features
2. Georeferencing image to control
3. Adding more features in other coordinate systems

### **1. Establishing control features**

For control features and to serve as primary orientation for my map, I found a set of roads and streams for the study area, a portion of Ventura County, and added them as layers to an ArcMap document. These control features were in a NAD27 geographic coordinate system, with a State Plane projected coordinate system appropriate for the area. To limit my control layer to a region roughly the same size as my image, I created a polygon feature class representing my image boundary and clipped roads and streams to this polygon. To create the image polygon, I located bounding latitude and longitude values in the image metadata, supplied them to ArcMap's "Go to XY" tool, placed markers representing each corner of my image, and connected the markers by editing features. (Note, I could also have supplied a file of the coordinates, which could be more useful if I would need to do this procedure more systematically or with multiple images.) After clipping my roads and streams to this area and zooming to the resulting extent of the clipped area, I was ready to add my image to the map document.

## 2. Georeferencing image to control

After adding the image to my ArcMap document, I opened the georeferencing toolbar and selected 'fit to display'. Having the control layers roughly the same size as my image was important in this step to help make the image roughly align with the control. I could see that my image was translated to the north of the roads layer, while rotation and scale were comparable to the roads. From here, I would add control points to hone the alignment or "register" the image, and once satisfied with the alignment, save the image by choosing "rectify" from the toolbar.

Reviewing my rough initial fit, I could see that my image didn't line up anywhere to the control layers, but it was close enough that I could see where to perform the control point linking. The roads in my image were always somewhere to the north of the corresponding roads in my control layer, though it varied from northeast to northwest. I suspected this was due to there being distortions beyond translation, rotation and scaling (such as bending or curving) in my image, but also could be due in part to the range of elevations in my study area leading to relief displacement or inconsistent scale at different elevations in my image. Ordinarily one would orthorectify an image with different elevations like this, but I would not be attempting orthorectification here; I will address this point further in the 'Considerations' section.


As to the georeferencing procedure, I was ready to begin the task of adding control points to register the image to my other map layers. I did not find my streams layer particularly useful here, as the corresponding streams in my image were usually either too wide, too covered by vegetation, or too ambiguous as to whether I was looking at a stream or a road. Therefore my control points primarily involved the linking of road centerlines. I attempted to place points in each corner of my image and distribute them widely across my area, though in a few places I found this to be impossible, such as the ocean to the southwest, a few areas in the mountains which lack well defined features, some areas to the east which appeared somewhat blurred on the image, and portions of the farmland to the northwest with few roads. When I had amassed a respectable number of initial points, I opened up the Links table on the georeferencing toolbar to review and refine my points as well as choose my transformation method.

### **Attempt #1: 3rd Order Polynomial Transformation**

I was starting with 20 control points. The default transformation method, 1st order polynomial (Affine), produced a root mean squared error (RMSE) of 166. Recalling that an RMSE less than raster cell size is ideal, I found I had a long way to go to reach that guideline, as my cell size was about 19. I also recalled that there are different methods, and especially given my older aerial camera photo, range of elevations, and observation that the image placed inconsistently in reference to the roads layer, I gave myself the freedom to explore the other options. I reviewed that the 1st order Affine polynomial will correct for translation of origin, rotation and scale differences, but if an image has additional levels of distortion like bending or curving, a higher order polynomial can accommodate these characteristics. My observations led me to think higher order polynomials were justified. The 3rd order polynomial, which requires 10 control points, brought me to an RMSE of 46. I noted there were some other options for methods I could also try, namely Splines and "Adjust", but I decided to work with the 3rd order polynomial for now and see if I could reach an RMSE under 19. I looked over my 20 points, both in terms of each residual and reviewing where I'd placed them. I found 2 or 3 instances where I had placed points poorly or there might have been changes to road layouts - given the 20 years that have passed since the image was taken, I was mindful of the possibility of true road changes. I removed these problem points, and also added a few more points to compensate, arriving at 24 control points. From here, I noticed that adding additional points were not really helping improve my result, both in that the RMSE remained high, even getting higher, and that in my map I wasn't seeing improvements to misalignments. I turned instead to removing control points with high residuals, and found after a few iterations of removal I was at an RMSE of 16, using 15 control points.


Link								
Total RMS Error: Forward:16.3459								
	Link	X Source	Y Source	X Map	Y Map	Residual_x	Residual_y	<Residual>
<input checked="" type="checkbox"/>	1	1.906331	3.037567	1683429.65...	230758.4...	-11.8133	25.443	28.0518
<input checked="" type="checkbox"/>	2	4.447958	4.516323	1710538.43...	246166.5...	-19.1277	14.9469	24.2751
<input checked="" type="checkbox"/>	3	4.464648	4.062908	1710637.60...	241269.7...	20.7284	-12.3013	24.1037
<input checked="" type="checkbox"/>	4	2.436508	2.376266	1689050.16...	223726.6...	9.84912	-17.5603	20.1338
<input checked="" type="checkbox"/>	5	4.109596	1.546040	1706480.35...	214902.4...	-8.86868	17.7752	19.8648
<input checked="" type="checkbox"/>	6	1.070721	3.504218	1674449.78...	235869.0...	0.696462	-19.7039	19.7162
<input checked="" type="checkbox"/>	7	4.577322	1.657312	1711411.77...	215981.0...	6.86367	-12.5698	14.3216
<input checked="" type="checkbox"/>	8	4.796616	4.626228	1714342.80...	247370.5...	10.5348	-9.29164	14.0469
<input checked="" type="checkbox"/>	9	4.793822	4.140821	1714140.37...	242131.1...	-9.08818	8.90613	12.7246
<input checked="" type="checkbox"/>	10	3.314189	4.429256	1698489.55...	245217.6...	5.85609	-8.54011	10.3551
<input checked="" type="checkbox"/>	11	4.440396	3.588931	1710256.86...	236280.4...	-7.40231	2.80183	7.91482
<input checked="" type="checkbox"/>	12	1.079503	3.974079	1674563.66...	240956.8...	5.94974	2.76558	6.56109
<input checked="" type="checkbox"/>	13	0.829995	2.599912	1671779.28...	226228.9...	-1.01483	6.11931	6.20289
<input checked="" type="checkbox"/>	14	1.446647	4.605547	1678572.44...	247675.5...	-3.88323	4.46869	5.92019
<input checked="" type="checkbox"/>	15	2.603803	1.315721	1690665.44...	212405.6...	0.720075	-3.25952	3.33811

The numbers were favorable, but how did this fit look visually? Turning to a detailed review of the image versus the roads layer, I found that results were good in most areas, but there were some clearly misaligned roads; not too surprisingly, the poorest fits were in the mountains. I recalled that a low RMSE is not necessarily a better fit, and thought of adding more control points again, but then also remembered that simply adding more points to my chosen method was not improving my results before, and maybe I should try something else instead.



### Attempt #2: Splines Transformation

I remembered there were some other transformation options I hadn't tried: Splines and Adjust. Splines were described in class lectures and go by some other names like rubbersheeting and piecewise polynomial transformation. The procedure basically fits each point exactly and tries to establish a smooth, continuous fit for in-between-points areas; i.e. it fits 'locally', and RMSE is 0 by definition for Splines. The ordered polynomial transformations approach fitting in a more global and formulaic way, which means that local point fits sometimes do not line up well and are the source for the RMSE. Splines also needed 10 control points, and to start I went back to my 24 point fit, and added more points until I was satisfied with the fit upon review. I noticed as I linked additional road points in the mountains that my near-sea level Pacific Coast highway which runs between the mountains and ocean became unfit, requiring additional links... again I wondered if this was an artifact of relief displacement and using georeferencing only when I ideally should also be incorporating orthorectification in image transformation. In the end with splines I had amassed 40 points! but my fit was quite good, at least by the measures I had to compare it to.



I also tried the "Adjust" transformation method, which is like a compromise between splines and polynomial. My result with Adjust was similar to the Splines, and I decided to move ahead with my Splines fit onto the next step, which was to save my results. I could either modify my existing image by applying the georeferencing, or save the transformation permanently in a new raster by using the Rectify option; I chose the latter, and selected cubic convolution as the interpolation type since it is recommended for aerial or satellite images where you want the (continuous) smoothing advantage of using multiple cells to calculate cell values like bilinear interpolation, but with a bit sharper of a resulting image from using 16 rather than 4 surrounding cells.

### 3. Adding more features in other coordinate systems


My final objective was to review and explore how to ensure that layers in different coordinate systems align and work with each other in ArcMap. My map thus far was all in the geographic NAD27 coordinate system with a State Plane projection. I found two other spatial datasets for the area that were in WGS84 and NAD83 geographic coordinate systems. Reviewing the procedures for working with coordinate systems, the NAD83 is a nice go-between for the other two systems, and with a set of NAD27 layers in a dataframe, one could set the dataframe to NAD83 before adding the WGS84 layer in order to have the layers align in the map document.

To test this, I first added the WGS84 layer, a map of trails, to my existing NAD27 dataframe. A warning about the unmatched geographic coordinate systems popped up, and ArcMap tried to align it by choosing a transformation. My trails were off by a small but noticeable distance versus the corresponding roads. I reviewed that the discrepancy is largely due to origin differences in local Kansas-centered NAD27 and global earth-centered WGS84. There is a well established grid-to-grid transformation called NADCON that is applied between NAD27 and NAD83, and NAD83 to WGS84 transformations work well as they are both earth-centered, therefore NAD83 can serve as a nice go-between when working with NAD27 and WGS84 coordinate systems. So I set my dataframe to NAD83, re-added my WGS84 layer, and now found that the trails aligned well.

NAD27 Dataframe:



NAD83 Dataframe:



Next I added a fire perimeter layer, where the spatial reference was in NAD83. Now that my dataframe was in NAD83, there were no alignment problems to be seen in adding this layer.

I had previously found it a bit confusing to understand the implications of what Esri describes as 'on-the-fly' versus inherent spatial reference transformations. A take-away that I realized from reviewing this practical exercise is that my final data frame could be in a geographic coordinate system that none of my layers are in, and I don't necessarily have to create multiple versions of features or rasters each transformed to the GCS I want to use.

## ***Considerations***

Below are a few considerations that came up during the previously described procedures.

### **1. Relief Displacement and Georeferencing**

Throughout the georeferencing process I kept in mind that I was not orthorectifying my image, nor using a stereo pair of images nor aerial triangulation of multiple images to try to take into account relief displacement. I was only georeferencing a single image, and I would probably encounter some signs that relief displacement was present since I had a considerable range of elevations represented in the image. I reviewed the concepts in that the range of elevations on a single aerial photo would mean the image has different scales across it, with larger scales at higher elevations than at lower elevations. Wanting to learn more, I looked up what was needed to perform orthorectification in ArcMap: I would need a high resolution elevation model and RPCs (rational polynomial coefficients) present in the metadata of my image. I knew I could find DEMs, though possibly not high enough in resolution, but my photo was an older aerial photo which did not have RPCs. I decided to go forward and see what I could accomplish with georeferencing only, again interested in how the inconsistent scaling would manifest in the georeferencing process. I also realized that exploring this non-ideal situation could be a valuable learning experience and probably akin to situations which GIS analysts and GIS developers encounter frequently in the real world.

### **2. Image Purpose**

Consideration for how the image is intended to be used after georeferencing also came to mind. If I was going to use it for a rough land use classification of the area, precision of registration is probably not that crucial. Using the image for background is a bit of a mixed bag; on the one hand, alignment seems less important since it is mainly for display, but on the other, it would quickly leave a bad impression if roads don't line up.


### **3. Georeferencing Combined With Coordinate Transformation**

When I reached the point of adding my additional trails and fire layers, I briefly wondered if there would be circumstances where applying coordinate transformation to my control layers first and then performing georeferencing would be advisable. In my case, I was not looking to permanently transform my layers to NAD83 or WGS84, but I felt that if such a permanent transformation were to be planned, that one might want to georeference the image to the final GCS rather than NAD27 as in my example.

## ***Conclusions***

In this exercise, I explored a practical situation by example, whereby I found an ungeoreferenced aerial image online, transformed it using different georeferencing techniques, and combined it with features in other coordinate systems in an ArcMap document. I further developed my skills and knowledge about georeferencing and gained new understandings in working with different coordinate systems in ArcMap. I also researched how one would use ArcMap tools to orthorectify

an image using a DEM. The rasters that resulted from my georeferencing efforts were useful and meaningful, especially when employing the Splines transformation method and with patience to link a copious number of control points.



## ***Acknowledgements***

- [Aerial Image](#) from USGS – EROS via Earth Explorer: single-frame 8-bit RGB image, 304.68 mm focal length, scale 66000, 10/24/1995
- [Ventura County GIS data](#), from 2015
- National Park Service: Santa Monica Mountains Public Trails; from ArcGIS Online
- California Department of Forestry Fire and Resource Assessment Program ([CDF-FRAP](#))
- Esri ArcGIS