

The Elkhorn Slough Historic Habitats Project

Eric Van Dyke

Elkhorn Slough National Estuarine Research Reserve

Abstract

Geospatial analysis is essential to conserving and restoring sensitive habitats. Yet trained technical staff is scarce within organizations that undertake this critical work. Although GIS and remote sensing are central to research, stewardship, and interpretation programs at ESNERR, coordination of geospatial work remains the responsibility of a single staff member. Elkhorn Slough has an enthusiastic cadre of volunteer docents. In recent years, a GIS component was added to our docent curriculum. As a result, a dedicated volunteer "historic habitats team" emerged with the ambitious goal of completing a major GIS project: a detailed baseline habitat classification for our watershed. Using widely available tools, including ArcView 3.x with the Image Analysis extension, the team developed a modular strategy to georeference, interpret, and digitize dozens of historic color infrared aerial photographs and to assemble a composite habitat map. In the future, the team expects to perform sophisticated habitat analysis and change detection.

Introduction

For readers who belong to large, well-funded geospatial analysis teams with access to sophisticated (and expensive) GIS and image processing software, some aspects of the project I'm about to describe are going to sound rather strange. Attempting to accurately georectify aerial photographs using only four ground control points? Digitizing a complex land cover map using nothing but simple overlapping polygons distributed among a number of separate shapefiles? What is this, some kind of contest—like doing GIS with one hand tied behind one's back?

I suspect that many of you will understand, though. Those of us who work in nonprofits or other organizations where staff and resources are scarce learn to attain ambitious goals by inventing innovative methodologies and by extending basic tools. One of the key lessons from this project is that if you must make do with limited resources, you'd do well to depend on a small, dedicated group of conservation-minded volunteers and on ArcView 3.x with its versatile Avenue programming language.

The 45,000 acre Elkhorn Slough watershed, located near Monterey Bay in central California (Figure 1), includes various environmentally sensitive habitat types and is home to several State and Federally listed rare and endangered species. Elkhorn Slough National Estuarine Research Reserve hosts a partnership between Federal, State, and community-based non-profit organizations that share the common goal of conserving and restoring the tidal wetlands and adjoining uplands of the Elkhorn Slough watershed.

Geospatial research was initiated at Elkhorn Slough with the acquisition of hardware, software, and training provided by the National Oceanic and Atmospheric Administration's Protected




Figure 1: Location of the Elkhorn Slough watershed in central California.

Areas GIS program in 1998, and bolstered by the award of a Conservation Technology Support Program grant to the non-profit Elkhorn Slough Foundation in 1999. Today, it would be difficult to find a research or stewardship activity at the Reserve that doesn't include a GIS, GPS, or remote sensing component. A growing segment of the ESNERR research program is our historical ecology project that uses geospatial tools to explore and quantify historical patterns of habitat change and to assist in setting ecologically and historically appropriate habitat restoration goals.

Despite all of this GIS activity at ESNERR, there's only one full-time geospatial professional on the Reserve staff—although as a result of a series of basic GIS and GPS workshops, the majority of the Reserve's employees, interns, and student researchers are now conversant with GIS concepts. More recently, a number of the Reserve's volunteer docents have received GIS training and are actively involved in an ambitious historical ecology research project. This paper reviews the initial work undertaken by this volunteer team and, in particular, some of the innovative GIS techniques that it developed out of necessity.

The Historic Habitats Team

As with many similar organizations, a large proportion of the day-to-day activities at Elkhorn Slough National Estuarine Research Reserve are conducted by trained volunteer docents—everything from leading visitor tours, planting and weeding native landscapes, and facilities maintenance, to active participation in staff and graduate student research projects. All ESNERR volunteers complete a ten-week comprehensive training program.

With the growth of our historical geospatial research, a new element was introduced into the docent training curriculum in 2002: an introduction to historical ecology and historical habitat mapping. As a result, two technically sophisticated volunteers expressed their interest in participating in historical ecology research. By the end of the 2003 training, over a dozen volunteers had expressed a similar interest. Clearly the time had come to incorporate volunteers in the research.

An obvious project was just waiting to be undertaken. In 2001, Reserve and Foundation staff jointly completed a detailed habitat classification of the entire 45,000 acre Elkhorn Slough watershed based on digital ortho imagery and extensive ground truthing. Historic 1:12,000 scale color infrared imagery of the watershed was also available; more than 100 separate 10" × 10" color prints flown back in April 1980. A similar classification of these 24-year-old images would provide a valuable look into the past and a means to quantify habitat change trends during the intervening two decades.

The questions were equally obvious. To georectify and interpret those dozens of aerial images, and to digitize the thousands of polygons that would represent the watershed's habitat types, was a dauntingly large task (which is why it had not been undertaken by Reserve staff, despite the obvious value of the resulting mapping and change analysis). Could it make sense to attempt such a project entirely with enthusiastic but untrained volunteers?

After much email and earnest discussion, a group of nine ESNERR volunteers agreed to give it a try. We christened our group the ESNERR "Historic Habitats Team" and reserved a series of upcoming weekends for project planning and focused GIS and remote sensing training sessions.

From the outset, it was clear that a key to our team's success would be our ability to partition the project into as many small, independent pieces as possible—without having it completely disintegrate! Decentralization was essential for a variety of reasons. The work would need to take place at a variety of locations such as GIS labs at local colleges and universities and at volunteers' homes and workplaces, because there simply aren't sufficient computer hardware and software licenses at the Reserve. In addition, volunteers work at different speeds and on their own individual schedules. The "glue" that would allow us to pull all these pieces back together would be the powerful scripting ability provided by ArcView 3.x's Avenue language.

Rectification and Mosaicing

Our first task was to assemble a georeferenced image mosaic of the core watershed from 78 of the overlapping 1980 color infrared aerial photographs. This mosaic would provide the basis for our habitat classification, and will also undoubtedly be used by many other researchers for future projects. Therefore we wanted to achieve the highest quality feasible, recognizing the limitations of the available image processing software (ArcView 8.x and ArcView 3.x with the Image Analysis extension) and the team's experience.

We adopted an innovative methodology that achieves relatively high spatial accuracy and low distortion while requiring the selection of only four ground control points per image. More importantly, our strategy supported decentralization; each image could be independently



Figure 2: A control point (red) placed near the perimeter of an image's effective area (green). Sketching (blue) helps to precisely position control points.




Figure 3: Intermediate image mosaic consisting of effective areas extracted from 78 color infrared aeriels. "Honeycomb" shapefile defines the effective areas.

rectified, repeatedly if necessary, at a remote location. When satisfied with each image's alignment to ground control, only the four control point pairs needed to be saved; these were easily emailed back to the Reserve as a shapefile. Resampling and saving the actual image files was unnecessary at this point, avoiding the need to transfer these extremely large raster files.

We used current, high quality digital orthophotos for ground control, selecting features (buildings, road intersections, etc.) that were readily identifiable on both the 1980 photos and the contemporary ortho imagery. Because we selected only four ground control points per image, thus there would be no least-squares adjustment, it was important to locate their positions as carefully as possible. A very useful technique was to sketch line segments from the reference imagery onto the graphic display—crosses to locate building centers, intersecting lines to precisely locate road intersections, and so forth—in order to increase the precision of control point locations (Figure 2).

Because distortion increases with increasing distance from the center of an aerial photograph, we constructed our final mosaic using only the central portion of each overlapping image. We maximized the spatial accuracy of points within this central effective area (at the expense of points near the outer edges of the image, which would not be included in the final mosaic) by selecting our four control points at well spaced locations near the perimeter of this central region (Figure 2). Because camera tilt was likely the greatest source of distortion within this central effective area, we resampled the image with a four-point plane projective model rather than the traditional 3-point affine model.




Figure 4. Plane projective rectification. Eight control points, four manually placed (red) and four automatically inserted (blue).




Figure 5. Final mosaic composed of 27 orthogonally aligned image tiles. "Checkerboard" defines the boundaries of the tiles.

A paradox inherent in our strategy of selecting control points near the perimeter of the central effective areas, and of including only this central portion of each image in the final mosaic, is that the determination of effective areas is only possible after all the images have been georeferenced! Therefore we actually performed the rectification task twice; a quick initial pass with four preliminary control points to locate the imagery for determination of effective areas, and then a final, refined placement with the control points near the perimeter of the effective area. After the first pass, an Avenue script applied the Spatial Analyst *proximity* function to calculate the region closest to the center of each (first pass, roughly placed) image. The resulting "honeycomb" polygon shapefile was used to guide the selection of control point locations and also served as the mask for assembling the mosaic.

When (second pass) control point selection was complete for all 78 images, an Avenue script assembled the intermediate mosaic by rectifying each image according to the saved control points and extracting each effective area according to the honeycomb mask (Figure 3). The script inserted four additional control points at midpoints between the four original control points. The result of resampling to these eight control points using a second-order polynomial transformation was a plane projective rectification model (Figure 4). Images were resampled to a consistent 0.5 m cell size using the cubic convolution method, and image boundaries were rounded outward to the nearest meter, yielding pixels aligned to even half meter coordinates.

A final Avenue script divided the assembled mosaic into 27 3km square north/south aligned tiles, enabling individual sections of the completed mosaic to be displayed in ArcView without triangular "no data" borders (Figure 5). Because the tiles were aligned to a uniform grid and the

cell size was unchanged, we performed this second resampling using the nearest neighbor method.

By randomly selecting a total of 100 independent points on the mosaic and comparing their locations to corresponding features on the reference orthophotos, we assessed the average spatial accuracy of our rectified imagery to be better than 5.4 meters.

Interpretation and Digitizing

Visual interpretation and classification of habitat types from historic aerial imagery is challenging even for experienced analysts. The process of defining and capturing the resulting geographic features as polygons in a GIS can also be complex and error-prone. Yet our team has demonstrated that by adopting simple, straightforward methodologies, dedicated beginners can produce accurate, professional results. Much as we did with the preparation of image mosaics, we avoided many complexities by partitioning and decentralizing individual tasks, and then relying on automated Avenue scripts for integration and assembly of the final habitat map.

Our first, prerequisite task was to define an appropriate classification system. We strictly adhered to the principle of "what you see is what you get"; every class in our habitat taxonomy corresponds to a set of regions that appear spectrally and texturally similar on the imagery. Definition of mapping units is never based on characteristics that are not visually apparent. Saltwater is visually indistinguishable from freshwater, for example, and thus share a single classification unit, as do grasslands and abandoned agricultural fields. Because aerial imagery captures land cover rather than land use, we classify "built-up" areas according to their dominant vegetation type (e.g. "woodland", "scrub", "grassland") unless these man-made features form a majority of the ground coverage; in that case, we classify them generally as "development" rather than that attempting to guess the particular land use (e.g. "residential", "commercial", or "industrial"). Another consequence of adopting the WYSIWYG principle is that our taxonomy deviates from standard habitat classification systems (e.g. *A Land Use and Land Cover Classification System for Use with Remote Sensor Data*, Anderson et al. 1996) due to regional differences. Maritime chaparral and coastal sage scrub, for example, are two habitat types that are unique to our central California coastal region but that are quite distinct on our color infrared imagery. Traditional classification systems would combine these two within a common "scrub/shrub" class. Conversely, although softwood and hardwood forests are distinguished in traditional classification systems, hardwoods (Blue Gum Eucalyptus) and softwoods (mostly Monterey Pine) in our region are spectrally and texturally indistinguishable, so we combined them in a single "eucalyptus/conifer forest" class. To clarify the visual characteristics of each of our habitat classes, we distributed "cheat sheets" to team members that illustrate representative examples (Figure 6). We applied a "greater than 50% cover" rule to characterize mixed habitat regions; for example a grassland with scattered oaks is classified as "oak woodland" if trees cover more than half the ground, but "grassland" if they cover less. Interpretation was performed at a consistent 1:2500 scale and using a minimum mapping unit of approximately 400m².



Figure 6: “Cheat sheet” to assist classification of “no-canopy” cultivated areas.

A second key feature of our habitat classification system is its binary, hierarchically nested structure: generalized “superclasses”, each comprising a pair of increasingly specific subclasses at the next lower level. This binary hierarchy closely resembles a traditional dichotomous key. Hierarchy allowed us to simplify the classification process by allowing the interpreter to focus on just one habitat class at a time. Probably the most obvious characteristic visible on our 1980 color infrared imagery is the presence or absence of a tree or shrub canopy. Regions dominated by trees and shrubs appear darker and show obvious texture, while regions with few or no trees or shrubs (e.g. grasslands and unvegetated areas) are lighter and show minimal texture. Therefore our initial, highest-level task consisted of interpreting the entire watershed solely on the basis of inclusion in this “tree/shrub canopy” superclass and digitizing simple polygons to surround included regions. To complete this top level of the class hierarchy, a custom Avenue script was then run to merge overlapping polygons, remove any polygons smaller than the minimum mapping unit, and then create complementary, adjoining “non-canopy” polygons in a separate shapefile complete the coverage (Figure 7).

As digitizing at any higher level in the classification hierarchy is completed, the remaining effort is reduced to a pair of smaller, independent tasks corresponding to the two subclasses at the next lower level. Thus after the set of polygons containing those portions of the watershed with tree or shrub canopy has been delineated, further classification to “woodland” or “scrub” is only performed within the boundaries of these “canopy” polygons. Similarly, “woodland” polygons will be more specifically classified as “orchard”, “riparian woodland”, “eucalyptus/conifer woodland”, or “oak woodland” at successively lower levels of the hierarchy. Conversely, classification of “open water”, “cultivated areas”, “grasslands”, and “bare ground” need only be performed within the boundaries of the separate (automatically generated) set of “non-canopy” polygons. The entire classification hierarchy is shown in Table 1.




Figure 7: Final digitized “canopy” layer (green) and automatically generated “no-canopy” layer (gray) in separate shapefiles.




Figure 8: Complex "canopy" area surrounding islands of "non-canopy" (asterisks) created by digitizing two overlapping polygons.

A principal advantage of our binary, hierarchical interpretation methodology is a dramatic reduction in the complexity of “heads-up” digitizing, an otherwise difficult and error-prone task. Traditionally, polygons representing multiple habitat types are digitized and edited simultaneously within a single coverage. Adjacency is maintained through snapping and other interactive topology maintenance rules. The creation of accidental gaps, overlaps, and “slivers” is a practically inevitable result of this complexity. At the same time, the class membership of each polygon must be identified through manual text or menu driven entry, another frequent source of errors. Our methodology avoids these pitfalls by isolating each habitat type in its own separate shapefile. Digitizing is restricted to single part polygons, all of the same class. Creation of complex polygons is accomplished by surrounding “holes” with multiple overlapping polygons (Figure 8). Manual merging of overlapping polygons within the shapefile is not necessary. In addition to entirely avoiding adjacency and class naming problems, this strategy facilitates project decentralization; each habitat type can be digitized independently. Habitat classes (and therefore shapefiles) are related according to a rule of strict containment down the classification hierarchy; the boundaries of polygons contained in a subclass are automatically clipped to the boundaries of polygons in the parent class. For example the extent of an “oak woodland” polygon will never extend beyond the boundaries of higher-level “woodland” and “canopy” polygons. Therefore exterior boundaries shared by a subclass and its superclass never need to be re-digitized. After digitizing any new interior boundaries, the lower-level polygon need only be extended to overlap the extent of the higher-level polygon where it will be clipped (Figure 9). In ArcView 3.x, this clipping is performed in real time by a custom interactive stream digitizing tool; in ArcView 8.x it is completed by Avenue during the cleanup process.




Figure 9: Digitized eucalyptus layer (blue). Overlap (blue crosshatch) beyond boundary of "canopy" layer (green) is automatically removed.




Figure 10: Composite habitat map.

Upon completion of digitizing at each level of the classification hierarchy, an Avenue script performs cleanup on the shapefile, merging overlapping polygons, exploding multi-part polygons, and flagging polygons that don't meet the minimum mapping unit threshold. The script also performs non-interactive enforcement of the containment rule by clipping to the boundaries of higher-level polygons if necessary. Finally, the script creates a separate shapefile containing complementary shapes (the "non-" class) to complete the coverage at that level.

When digitizing all levels of the classification hierarchy has been completed, the final composite habitat map comprises the stack of shapefiles corresponding to the full set of habitat classes (i.e. the non-overlapping "leaves" of the classification tree, Figure 10). Flattening the hierarchy into a single combined coverage is not necessary.

Conclusions

The Elkhorn Slough historic habitats project has provided an unambiguous answer to the question of whether a team of enthusiastic volunteers can complete an ambitious GIS and remote sensing project with minimal resources—at least from a technical perspective. A number of innovations emerged that simplify and streamline the tasks of constructing a spatially accurate historic aerial photo mosaic and of interpreting and digitizing the mosaic to produce an accurate habitat classification map. Our strategy of partitioning and decentralizing tasks, and then integrating the results through automated Avenue scripts, has proven particularly effective.

Yet difficulties of an entirely different nature have accompanied the project from its inception. We dramatically underestimated the investment that would be necessary to focus and sustain the energies of a diverse and independent team of volunteers within our highly decentralized project structure. Portions of the project would be completed efficiently, yet the overall schedule frequently dragged as key tasks remained incomplete. Sometimes it felt like the quickest course might be to forget about teamwork and just do it all oneself. Effective project coordination has turned out to be at least as critical to our success as the development of clever technical methodologies. Too bad Avenue doesn't have a set of requests for managing people as it does for managing polygons and pixels!

The team's most frequently asked question, during those tedious sessions of painstaking manual interpretation, has been: "can't this be done automatically"? From historical aerial photography, with its limited spectral resolution and inconsistent color balance, the level of classification accuracy that we're seeking probably can't be readily obtained. With the completion of our historical baseline map, we hope to begin to explore this question by acquiring contemporary multispectral imagery and expanding our skills into the areas of automated classification and change analysis techniques. With luck, the historic habitats team will have some interesting results to report again next year.

Acknowledgments

Thanks to Aneita, Ann, Gary, Jackie, Kelly, Kermit, Meghan, Natasha, Sarah, Sharon, and everyone else who helped.

Table 1. Elkhorn Slough Historic Habitats Classification Taxonomy

- [1a] canopy (tree/shrub layer > 50% cover)
 - [2a] shrubs (shrub layer > 50% cover)
 - [3a] **Maritime Chaparral** (uniform rusty red)
 - [3b] **Coastal Scrub**, dunes or inland (mottled gray/violet)
 - [2b] trees (tree layer > 50% cover)
 - [4a] **Orchard** (patterned texture, red to pink)
 - [4b] non-cultivated
 - [5a] **Riparian Woodland** (pinkish red)
 - [5b] non-wetland (red)
 - [6a] **Eucalyptus/Conifer Woodland** (deep shadow, greenish red)
 - [6b] **Oak Woodland** (rounded crowns, red)
- [1b] non-canopy (< 50% tree/shrub layer cover)
 - [7a] **Open Water**, slough, lake, pond; tidal or fresh
 - [7b] non-water
 - [8a] **Cultivated** row crops, bare plowed land (patterned, red to pink to white)
 - [8b] non-cultivated
 - [9a] vegetated
 - [10a] **Marsh**, tidal or fresh (gray-green)
 - [10b] **Grassland/Vegetated Old Field** (pale pink, tan)
 - [9b] unvegetated
 - [11a] **Developed** residential, industrial, greenhouse
 - [11b] undeveloped
 - [12a] **Unvegetated Wetland** (bluish gray)
 - [12b] **Bare Ground/Dry Pond** (white)