Skyline-Based Camera Orientation Estimation and Directed Viewshed Computation from Geotagged Images

Evan Twarog - etwarog@stanford.edu
Stanford University, Electrical Engineering Department

EE 367

Location Knowledge from Imagery

We estimate camera orientation from a single geotagged images by aligning the observed skyline with a synthetic skyline generated from a digital elevation model (DEM).

Motivation

Left: AlertCA is a network of cameras that monitors for wildfires. Geolocating ignitions requires knowledge of nearby landscapes. Image courtesy of AlertCA.

Applications include environmental monitoring tasks such as wildfire detection, geolocation of observed events, and GPS-denied navigation.

Related Work and Methods

Monoplotting: Approach to geolocate features in a single photograph using a DEM.

Related Work:

1. **Bozzini et al.** [1] developed the WSL Monoplotting Tool for georeferencing oblique images by calibrating the camera and intersecting image rays with a DEM.
2. **Chen et al.** [2] extend skyline matching by incorporating internal ridge features and estimating camera roll to improve robustness when skylines are occluded.
3. **Fedorov et al.** [3] identify mountain peaks by aligning photo edges with DEM-rendered panoramas from the approximate camera location.

Modern Approaches:

- Computer vision algorithms to specific components, such as skyline extraction.
- Extend to other features in the landscape other than the skyline (Better for flatter regions).

References

[1] C. Bozzini, M. Conedera, and P. Krebs, *A New Monoplotting Tool to Extract Georeferenced Vector Data and Orthorectified Raster Data from Oblique Non-Metric Photographs*, International Journal of Heritage in the Digital Era, vol. 1, no. 3, 2012.
[2] Y. Chen, G. Qian, K. Gunda, H. Gupta, and K. Shafique, *Camera Geolocation From Mountain Images*, Proc. 18th International Conference on Information Fusion, Washington, DC, 2015.
[3] R. Fedorov, P. Fraternali, and M. Tagliasacchi, *Mountain Peak Identification in Visual Content Based on Coarse Digital Elevation Models*, Proc. ACM MAED Workshop, Orlando, FL, 2014.

Georeferencing Approach

① Extract Image Metadata

- Key Fields: Focal Length, GPS Coordinates
- Calculate Field of View

② Extract Image "Skyline"

- Apply gaussian blur and Laplacian filtering on grayscale image to separate sky from high-gradient terrain
- Apply thresholding and morphological erosion to create binary sky mask
- Enforce a single horizon per image column
- Perform multiple pass to reject unreliable columns

③ Generate Viewshed and Synthetic Skyline

- Use ray tracing to determine which DEM cells are visible from observer location.
- Synthetic skyline corresponds to the DEM cell with the maximum elevation angle in each direction.

④ Estimate Orientation via Cross-Correlation Between Extracted and Synthetic Skylines

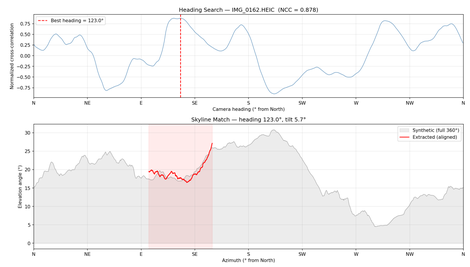

An example output from the cross correlation. The maximum cross correlation corresponds to the estimated viewing direction.

⑤ Create "Imaged" Viewshed from Estimated Orientation

Case Study: A Walk in Alaska

Skyline Extraction:

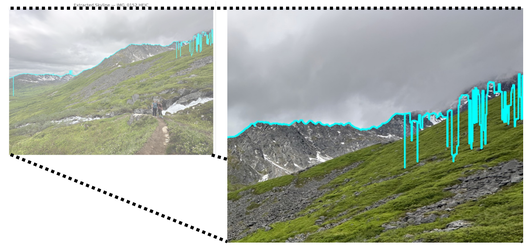

Skyline extraction can be challenging in foggy or snow-covered terrain. Despite partial skyline detection failures, we maintain robust direction estimates.

Cross-Correlation:

Mean	Min	Max
0.922	0.878	0.963

Values of the peak cross-correlation values across the 6-image set

- The method consistently identifies a dominant correlation peak corresponding to the correct viewing direction.

Output Viewshed Results:

- Quantitative: Limited because we do not have the original image capture direction and orientation.
- Qualitative: All images match the orientation that the image was captured

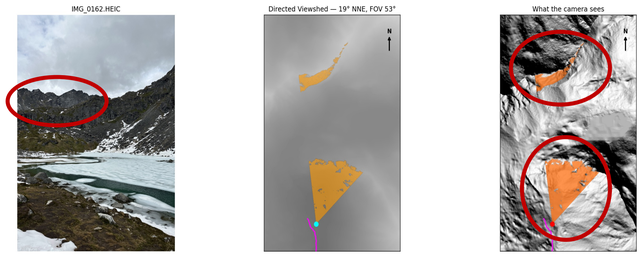

Two mountain ranges are captured in this image. The derived viewshed identifies these two ranges clearly.

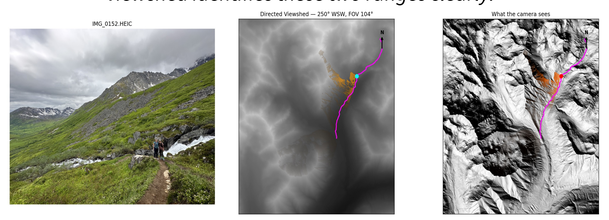

Even with horizon extraction failures [see above], we accurately estimate the image orientation.