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

Evan Twarog

Abstract—We present an approach for estimating camera orientation from geotagged photographs in mountainous terrain and computing the corresponding directed viewshed, the precise region of terrain visible in the image. Given a photograph with GPS coordinates and EXIF-derived field of view, we (1) extract a one-dimensional skyline profile using gradient-based image segmentation, (2) synthesize a 360-degree horizon profile from a high-resolution digital elevation model (DEM) via ray casting, and (3) determine the camera heading and tilt by maximizing normalized cross-correlation (NCC) between the two profiles over all candidate orientations. A directed viewshed identifies the terrain region likely visible within the photograph's field of view. We demonstrate the method on photographs taken during a backcountry hike in the Chugach Mountains of Alaska using 2 m a Polar Geospatial Center ArcticDEM dataset. Because the original image heading was not available in the metadata, performance was evaluated by comparing angular offsets between a synthetic skyline and the extracted, oriented skyline. The mean azimuth discrepancy was $2.95°$, while the mean elevation discrepancy was $0.98°$.

Index Terms—Skyline matching, viewshed analysis, digital elevation model, camera orientation, geotagged photography, normalized cross-correlation

1 Introduction

Geotagged photographs captured in mountainous terrain encode rich geometric information: the GPS location of the camera, the focal length from image metadata, and the visual content of the scene itself. The image user may be interested in later understanding more about what part of the landscape was captured, potentially for event registration. Applications include event location estimation from video or image feeds, particularly when only a single sensor is available such as in remote regions. An example includes wildfire ignition location when it is only captured with a single camera.

This paper addresses the problem of recovering camera orientation from a single geotagged photograph and a digital elevation model (DEM). The horizon line, the sky-terrain boundary visible in a mountain photograph serves as a unique angular fingerprint of the observer's location and viewing direction. By matching this extracted skyline to a synthetic 360-degree horizon profile rendered from the DEM, we can determine the camera heading.

Our contributions are as follows:

1) A pipeline from raw geotagged photographs to georeferenced viewshed maps, requiring only standard image processing and no machine learning.
2) A signal-processing approach to skyline extraction using Laplacian gradients and morphological filtering.
3) A directed viewshed visualization that maps each photograph's content to the specific terrain it depicts.
4) An error estimation approach based on manually matching corresponding features between the oriented synthetic skyline and the extracted skyline.

We demonstrate the pipeline on photographs taken during a hike in the Chugach Mountains of Alaska (Fig. 3), using a 2 m Polar Geospatial Center Arctic DEM [1].

• *E. Twarog is with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305.*
E-mail: etwarog@stanford.edu

2 Related Work

2.1 Photo-to-Terrain Alignment

The problem of aligning photographs to 3D terrain models has received significant attention. Baboud et al. [2] proposed an automatic photo-to-terrain alignment system for annotating mountain pictures, using edge detection on rendered DEM panoramas and matching against image edges via a robust alignment procedure. Their method estimates the camera position and orientation, but requires an approximate initialization.

Baatz et al. [3] and Saurer et al. [4] tackled large-scale visual geo-localization in mountainous terrain, extracting skyline contours from query images and matching them against a database of rendered DEM silhouettes. Their approach demonstrated that skyline shape alone provides sufficient information for geo-localization over large geographic regions. Our work differs in that we assume the GPS position is known (from EXIF metadata) and focus on recovering the viewing direction, which simplifies the problem from a 3-DOF search to a 1-DOF heading sweep.

Fedorov et al. [5] presented a method for mountain peak identification based on edge matching between photographs and DEM-rendered views. Their work demonstrated the viability of coarse DEMs for skyline matching, though they

focused on peak labeling rather than viewshed computation.

2.2 Terrain Panorama Rendering and Monoplotting

Bozzini et al. [6] developed a monoplotting tool for extracting georeferenced data from oblique photographs using DEMs. Their work established the georeferencing framework that connects image coordinates to terrain coordinates through camera geometry and DEM data. Steger, Steger and Schar [7] introduced HORAYZON, an efficient ray-tracing algorithm for computing horizon profiles and sky view factors from DEMs, demonstrating that high-resolution terrain visibility computation is tractable for large datasets.

2.3 Viewshed Analysis

Viewshed analysis—determining the set of terrain points visible from a given observer location—is a fundamental GIS operation [8], [9]. Traditional viewshed algorithms compute omnidirectional (360-degree) visibility. Our directed viewshed restricts computation to a wedge defined by the matched camera heading and field of view, producing a map of exactly the terrain depicted in a given photograph.

3 METHOD

Our pipeline consists of four stages: metadata extraction, skyline extraction, synthetic skyline generation with viewshed computation, and skyline matching with directed viewshed computation.

3.1 Metadata Extraction

From each HEIC or JPEG photograph, we extract EXIF metadata including GPS coordinates (latitude, longitude) and the 35 mm-equivalent focal length f_{35}. The horizontal field of view is computed as

$$\mathrm{HFOV} = 2\arctan\left(\frac{36\,\mathrm{mm}}{2\,f_{35}}\right), \tag{1}$$

where 36 mm is the width of the standard 35 mm full-frame sensor (36×24 mm) and f_{35} is read from the EXIF `FocalLengthIn35mmFilm` field [10].

3.2 Skyline Extraction

We extract the sky-terrain boundary from each image using a gradient-based approach. While neural network methods for skyline segmentation exist, we adopt a classical signal-processing pipeline for simplicity and to avoid training-data requirements. Fig. 5 shows the extracted skyline (cyan) overlaid on three photographs spanning varied conditions: a valley with cloud-obscured peaks (IMG_0141), a steep slope captured with a wide-angle lens (IMG_0152), and a portrait-orientation view across a frozen lake (IMG_0162).

3.2.1 Gradient-Based Sky Detection

The input RGB image is converted to grayscale and smoothed with a 9×9 box filter followed by a 15×15 median filter to suppress texture while preserving the horizon edge. We compute the Laplacian $\nabla^2 I$ and threshold it: pixels with $|\nabla^2 I| < \tau$ (we use $\tau = 6$) are classified as candidate sky regions, exploiting the observation that sky regions have near-zero second derivatives while terrain exhibits high-frequency gradients.

3.2.2 Morphological Refinement

The binary sky mask is refined by morphological erosion with a 9×3 rectangular kernel, which removes thin spurious sky regions. A per-column scan then enforces a single contiguous sky region at the top of each column: for each column, we identify the first sky-to-terrain transition and set all pixels below it to terrain.

3.2.3 Profile Extraction and Smoothing

The skyline profile $p[x]$ is defined as the row index of the first terrain pixel in each column x. This raw profile is smoothed by a multi-pass outlier rejection algorithm: columns where $|p[x] - p[x-1]| > \Delta_{\max}$ (we use $\Delta_{\max} = 20$ pixels) are flagged, and contiguous runs of flagged columns are replaced by linear interpolation. A final median filter ($k = 15$) removes residual jitter. Three passes ensure that neighbors of removed spikes are also cleaned.

3.3 Synthetic Skyline and Viewshed

Once we extract the skyline, we use the GPS location metadata to then construct a synthetic skyline we compare against.

3.3.1 Ray Casting

From the observer's pixel position on the DEM, we cast $N = 3{,}600$ rays uniformly spaced in azimuth (0.1° steps) using compass bearing convention (0° = North, 90° = East). For each ray at azimuth θ, the column and row increments are

$$\Delta c = \sin\theta, \qquad \Delta r = -\cos\theta, \tag{2}$$

where the negative sign accounts for image coordinates having row indices increasing downward (southward). Each ray is traced outward step by step up to a maximum radius $R_{\max} = 15\,\mathrm{km}$. At each step s, the slope to the terrain point is

$$\alpha_s = \frac{z(r_s, c_s) - z_{\mathrm{obs}}}{s \cdot \Delta_{\mathrm{cell}}}, \tag{3}$$

where $z_{\mathrm{obs}} = z(r_0, c_0) + h_{\mathrm{obs}}$ includes the observer height ($h_{\mathrm{obs}} = 1.7\,\mathrm{m}$) and Δ_{cell} is the DEM cell size (2 m).

3.3.2 Viewshed and Skyline Profile

A point is *visible* if its elevation angle exceeds the maximum elevation angle of all preceding points along the same ray. The omnidirectional viewshed is the binary map of all visible cells across all rays. The synthetic skyline $\hat{e}(\theta)$ records the maximum elevation angle $\max_s \alpha_s$ along each azimuth θ, converted to degrees:

$$\hat{e}(\theta) = \arctan\Big(\max_s \alpha_s\Big) \cdot \frac{180^\circ}{\pi}. \tag{4}$$

3.4 Skyline Matching

3.4.1 Angular Coordinate Conversion

The extracted pixel skyline $p[x]$ is converted to angular coordinates using the camera's FOV. For image column x in an image of width W:

$$\phi(x) = \frac{2x - (W-1)}{W-1} \cdot \frac{\mathrm{HFOV}}{2}, \tag{5}$$

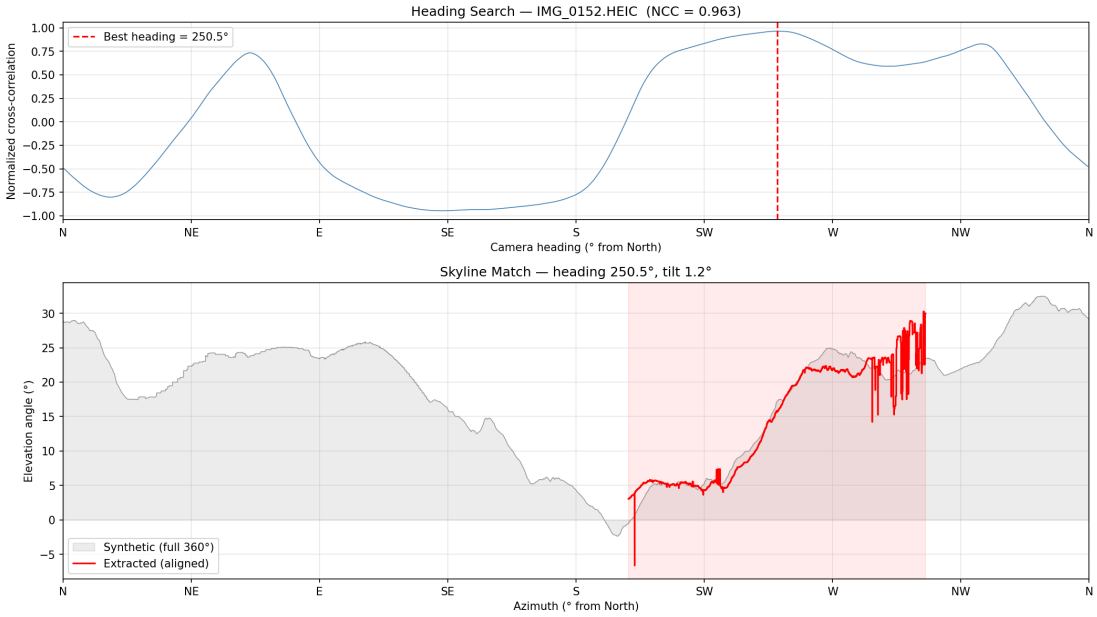

Fig. 1: Example matching result for IMG_0152. **Top**: Normalized cross-correlation versus candidate heading, showing a sharp peak at 251°. **Bottom**: Overlay of the extracted photo skyline (red) and the DEM-derived synthetic skyline (blue) at the best-fit heading, with NCC = 0.963.

(a) IMG_0141: Looking up a valley.

(b) IMG_0152: Looking across a slope.

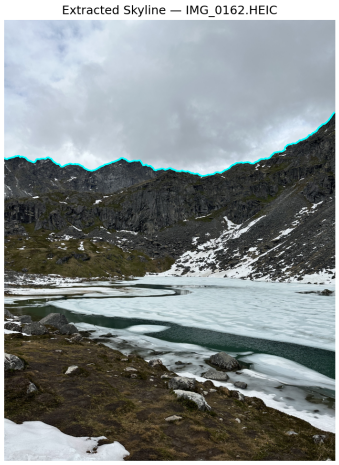

(c) IMG_0162: Looking up at a series of ranges.

Fig. 2: Extracted skyline profiles (cyan) overlaid on source photographs. While in most images we successfully extract a full skyline, in limited instances, the skyline extraction fails due to cloud occlusion and snow coverage.

giving the relative azimuth from the image center (negative = left, positive = right). The elevation angle is

$$e(x) = \frac{(H-1)/2 - p[x]}{(H-1)/2} \cdot \frac{\text{VFOV}}{2}, \tag{6}$$

where H is the image height.

3.4.2 Heading Search via NCC

We search over candidate headings $h \in [0°, 360°)$ at 0.5° steps. For each h, the photo's relative azimuths are mapped to absolute compass bearings $\theta_x = (h + \phi(x)) \bmod 360°$, and the synthetic skyline is interpolated at these bearings. The optimal tilt (camera pitch) is estimated as the mean residual:

$$\hat{t}(h) = \frac{1}{W} \sum_{x=0}^{W-1} \left[\hat{e}(\theta_x) - e(x)\right]. \tag{7}$$

The normalized cross-correlation between the tilt-corrected photo skyline and the synthetic skyline is

$$\text{NCC}(h) = \frac{\sum_x \left(\hat{e}(\theta_x) - \bar{\hat{e}}\right)\left(e'(x) - \bar{e'}\right)}{\sqrt{\sum_x \left(\hat{e}(\theta_x) - \bar{\hat{e}}\right)^2 \sum_x \left(e'(x) - \bar{e'}\right)^2}}, \tag{8}$$

where $e'(x) = e(x) + \hat{t}(h)$ is the tilt-shifted photo skyline [11]. The best heading is $h^* = \arg\max_h \text{NCC}(h)$.

3.5 Directed Viewshed

Given the matched heading h^* and the horizontal FOV, we compute a *directed viewshed*: a viewshed restricted to the angular wedge $[h^* - \mathrm{HFOV}/2,\ h^* + \mathrm{HFOV}/2]$. This uses the same ray-casting algorithm as the omnidirectional viewshed but limits rays to the wedge. The heading h^* (a true compass bearing) is converted back to a grid bearing for ray casting:

$$h_{\mathrm{grid}} = (h^* + \gamma) \bmod 360^\circ. \tag{9}$$

The directed viewshed represents a visible terrain wedges constrained by heading and FOV.

4 Analysis, Evaluation and Comparison to Other Methods

4.1 Analysis

No ground-truth compass headings are available (the phone's `GPSImgDirection` EXIF field was not populated). We evaluate matching quality using three complementary criteria: (1) the NCC score and peak sharpness of the heading search, (2) visual verification of the directed viewshed against the photograph content, and (3) manual point correspondence measurements of angular offsets between the extracted and synthetic skylines (Section 5.6).

4.2 Evaluation

Benchmarking skyline-based alignment methods remains challenging. Several prior works evaluate performance using controlled datasets where camera pose is known or approximately reconstructed. Baboud et al. [2] evaluated alignment accuracy using manually verified correspondences between image skylines and rendered terrain silhouettes, reporting angular errors in camera orientation. Similarly, Saurer et al. [4] evaluated skyline matching performance by measuring localization accuracy across large geographic databases of rendered terrain silhouettes. For our approach, we assess our oriented, extracted skyline against a synthetic skyline generated by the high-resolution DEM. In future work, we may consider collecting similar imagery but record compass headings and lines of bearing to key terrain features.

4.3 Comparison to Other Methods

Unlike global image-based geo-localization approaches such as Baatz et al. and Saurer et al., our method assumes the camera position is known and solves only for orientation, reducing the search space and enabling efficient skyline alignment using simple cross-correlation. Compared to monoplotting approaches that require calibrated camera models, our pipeline requires only approximate field-of-view metadata.

5 Results

5.1 Dataset and Setup

We tested the approach on photographs captured during a backcountry hike in the Chugach Mountains near Anchorage, Alaska. Images were taken with an iPhone (4032×3024 or 3024×4032 pixels) in both landscape and portrait orientation. DEM data were obtained from the Polar Geospatial Center ArcticDEM dataset.

5.2 Matching Quality

NCC scores range from 0.877 to 0.963 (mean 0.931), indicating strong agreement between extracted and synthetic skylines. The estimated tilt values are small ($|t| < 6^\circ$), consistent with hand-held photographs.

5.3 Peak Sharpness

The reliability of the heading estimate depends not only on the magnitude of the NCC peak, but also on whether similar terrain features are misidentified as the correct heading. In each case examined, the NCC correctly identified the approximate orientation when compared against similar features elsewhere in the skyline. The peak NCC values were typically above 0.90, whereas "secondary" peaks associated with similar terrain rarely exceeded 0.80. Figure 1 shows the NCC-versus-heading curve for IMG_0152.

While this approach performed well in featured terrain with high-resolution DEMs, it raises questions about the robustness of the algorithm when only lower-resolution *a priori* knowledge is available or when the terrain contains fewer distinctive features. Future work will examine the impact of both factors.

5.4 Directed Viewshed Visualization

Figs. 3 and 4 show directed viewshed outputs for two photographs taken at different locations along the hike. Each three-panel figure shows the original photograph (left), the directed viewshed wedge overlaid on the grayscale DEM (center), and the viewshed on a hillshade terrain rendering (right). The viewshed wedges align visually with the terrain features visible in the photographs. The hike track (magenta) and observer position (cyan dot) provide geographic context.

5.5 Robustness Considerations

The skyline extraction proved robust across varying conditions, including partial snow cover, overcast skies, and the presence of a frozen lake in the foreground (IMG 0162).The multi-pass outlier rejection effectively removed artifacts from clouds, snow patches, and foreground objects that might otherwise corrupt the horizon profile.

The lowest NCC score (0.877 for IMG_0162) corresponds to a narrow-FOV portrait-orientation image looking roughly northward, where the available skyline arc is only 53°—a smaller matching footprint that still produced a correct match based on visual verification.

5.6 Manual Angular Offset Analysis

To quantify the local agreement between the extracted and synthetic skylines beyond the global NCC metric, we performed a manual point-correspondence analysis. For each image, we overlaid both the extracted skyline (cyan) and the NCC-aligned synthetic skyline (red, dashed) on the original photograph. An operator then selected three pairs of corresponding landmark points—each pair consisting of a point on the extracted skyline and a point on the synthetic skyline at the same visually identifiable terrain feature (e.g., ridge apex, saddle point, or peak tip).

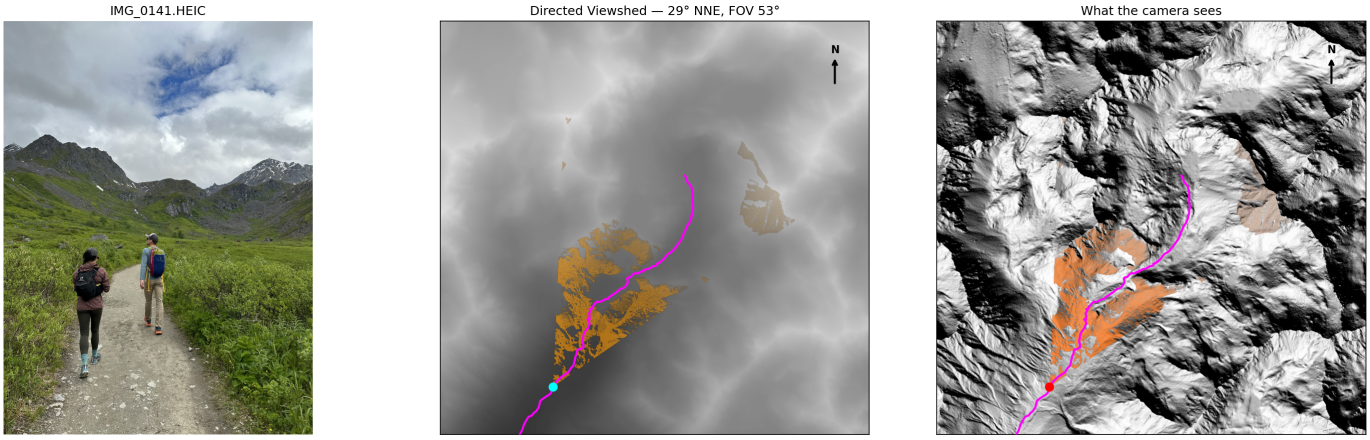

Fig. 3: Directed viewshed for IMG_0141 (heading 29° NNE, FOV 53°). A view looking up a valley including distant peaks. The viewshed wedge (orange) captures the valley and ridgeline visible in the photograph. The hillshade rendering (right) clearly shows the valley morphology.

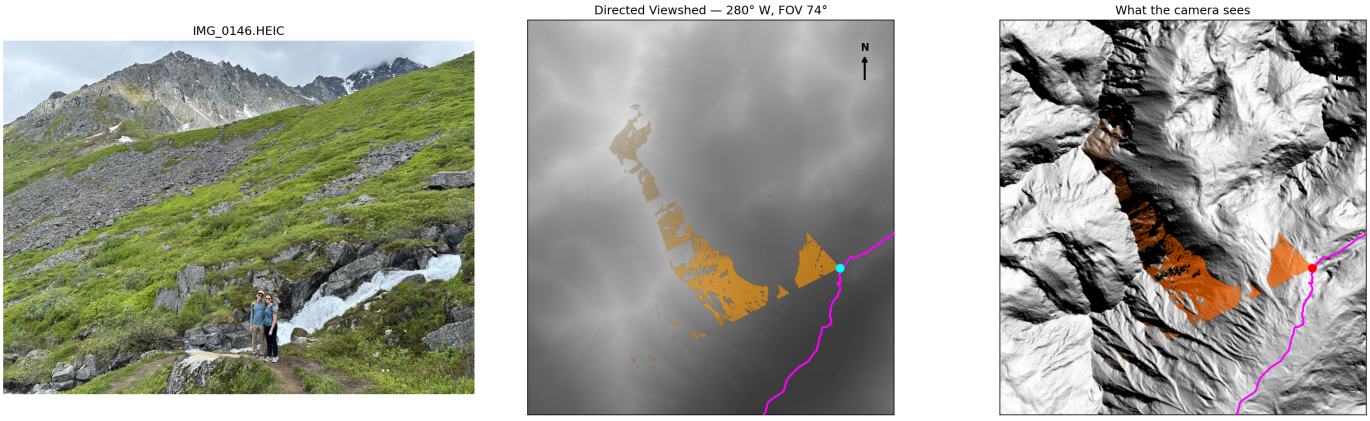

Fig. 4: Directed viewshed for IMG_0146 (heading 280° W, FOV 74°). This image captures two distinct horizons with a closer gentle ridge giving way to a more jagged mountain ridge. The viewshed extends along the valley axis, matching the visible terrain. Disconnected visible patches beyond the foreground ridge correspond to distant peaks visible over the ridge.

(a) IMG_0141: Looking up a valley.

(b) IMG_0146: Looking across a slope.

(c) IMG_0162: Looking up at a series of ranges.

Fig. 5: Extracted skyline profiles (cyan) and synthetic skyline profiles (red) overlaid on original images.

TABLE 1: Mean Angular Offset Between Extracted and Synthetic Skylines (3 manually selected point pairs per image)

Image	Mean $\lvert\Delta\theta_{\text{az}}\rvert$ (°)	Mean $\lvert\Delta\theta_{\text{el}}\rvert$ (°)	Mean $\Delta\theta$ (°)
IMG_0143	1.27	1.09	2.45
IMG_0146	0.74	0.59	1.04
IMG_0152	4.11	1.98	4.68
IMG_0141	1.34	0.73	1.53
IMG_0155	1.87	1.29	2.42
IMG_0162	8.37	0.16	8.37
Overall	**2.95**	**0.98**	**3.31**

The pixel displacement $(\Delta p_x, \Delta p_y)$ between each pair was converted to angular offsets using the camera's field of view:

$$\Delta\theta_{\text{az}} = \frac{\Delta p_x}{W} \cdot \text{HFOV}, \qquad \Delta\theta_{\text{el}} = -\frac{\Delta p_y}{H} \cdot \text{VFOV}, \quad (10)$$

where W and H are the image dimensions in pixels. The total angular offset is $\Delta\theta = \sqrt{\Delta\theta_{\text{az}}^2 + \Delta\theta_{\text{el}}^2}$.

Table 1 reports the per-image mean offsets across the three measured point pairs. Five of six images show mean total offsets below $2.5°$, with azimuthal and elevation components typically under $2°$. The outlier is IMG_0162, a portrait-orientation image of a frozen alpine lake with a narrow FOV (53°), where horizontal offsets of 5–10° suggest a systematic mismatch. This image also had the lowest NCC (0.878), and visual inspection indicates the extracted skyline deviates where snow and rock produce ambiguous gradients along the horizon.

Across all 18 measurements, the mean azimuthal offset is $2.95°$ and the mean elevation offset is $0.98°$, yielding a mean total angular error of $3.31°$. Excluding IMG_0162, the mean total drops to $2.42°$. Sources of error include: (1) skyline extraction artifacts from clouds, snow, or foreground objects; (2) DEM resolution limitations (2 m cells smooth fine ridge features); (3) lens distortion not modeled in the linear FOV-to-pixel mapping; and (4) subjectivity in manual point selection, particularly where the skyline curves gradually.

6 Discussion and Limitations

The results demonstrate that skyline matching provides a reliable signal for estimating camera orientation when the observer location is known. In most images, the normalized cross-correlation (NCC) between the extracted skyline and the DEM-derived skyline produced a sharp, well-defined peak, indicating that the skyline profile acts as a distinctive geometric signature of the terrain. The directed viewshed visualizations further confirm that the estimated headings align with the terrain features visible in the photographs.

Despite these encouraging results, several limitations remain. The approach for skyline extraction is sensitive to visual ambiguity at the sky-terrain boundary. Clouds, snow patches, and foreground objects can introduce gradients that resemble horizon edges, which may distort the extracted skyline profile. Recent work has demonstrated improved skyline extraction using neural-network-based segmentation methods.

An additional limitation is that the geometric accuracy of the synthetic skyline depends on the resolution and fidelity of the digital elevation model. While the 2 m ArcticDEM provides high spatial resolution, fine terrain features such as sharp ridge crests or narrow peaks may still be smoothed by the DEM grid, introducing small discrepancies between the rendered skyline and the true visual horizon.

The current method assumes an idealized pinhole camera model with a linear mapping between image pixels and viewing angles derived from the EXIF field-of-view metadata. In practice, consumer cameras may exhibit lens distortion or digital correction that deviates from this model. These effects can introduce systematic angular offsets, particularly for wide-angle lenses.

Finally, the evaluation presented here relies on manual feature correspondences between the synthetic and extracted skylines because ground-truth camera headings were not available in the EXIF metadata. While this approach provides a reasonable estimate of alignment accuracy, a more rigorous evaluation would involve datasets with known camera poses or multiple images of the same scene to enable consistency checks.

Future work could address these limitations by incorporating lens distortion calibration, using multi-image constraints along a hiking track, or integrating learning-based skyline segmentation methods to improve robustness under challenging visual conditions.

7 Conclusion

We presented a pipeline for estimating camera orientation from geotagged mountain photographs using skyline matching against DEM-derived horizon profiles. The method requires no machine learning, no compass data, and no manual annotation—only a GPS-tagged photograph and a DEM. Using normalized cross-correlation over a brute-force heading sweep, we achieve NCC scores of 0.88–0.96 on photographs from an Alaska backcountry hike, with sharp correlation peaks indicating unambiguous heading estimates.

The directed viewshed output provides a novel visualization linking photographic content to geographic terrain, with potential applications in outdoor navigation, geological survey, search-and-rescue, ecological monitoring, and augmented reality annotation of mountain scenes. Future work could incorporate lens distortion modeling, and learning-based skyline segmentation to improve robustness under challenging visual conditions.

Poster Session Feedback

The primary feedback received during the poster session was the need to assess performance and accuracy despite the absence of the original image heading. We incorporated this feedback by using a manual feature-matching approach to estimate the angular offset between the synthetic skyline and the extracted skyline.

Acknowledgments

The author thanks the Stanford EE 367 course staff for guidance. ArcticDEM data were provided by the Polar Geospatial Center under NSF-OPP awards 1043681, 1559691, and

1542736. Hike track recorded via GPS during a hike in the Chugach Mountains.

REFERENCES

[1] C. Porter, P. Morin, I. Howat, M.-J. Noh, B. Bates, K. Peterman, S. Keesey, M. Schlenk, J. Gardiner, K. Tomko, M. Willis, C. Kelleher, M. Cloutier, E. Husby, S. Foga, H. Nakamura, M. Platson, M. Wethington, C. Williamson, G. Bauer, J. Enos, G. Arnold, W. Kramer, P. Becker, A. Doshi, C. D'Souza, P. Cummens, F. Laurier, and M. Bojesen, "ArcticDEM, Version 4.1," 2023. [Online]. Available: https://doi.org/10.7910/DVN/3VDC4W

[2] L. Baboud, M. Cadík, E. Eisemann, and H.-P. Seidel, "Automatic photo-to-terrain alignment for the annotation of mountain pictures," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2011, pp. 41–48.

[3] G. Baatz, O. Saurer, K. Köser, and M. Pollefeys, "Large scale visual geo-localization of images in mountainous terrain," in *European Conference on Computer Vision (ECCV)*. Springer, 2012, pp. 517–530.

[4] O. Saurer, G. Baatz, K. Köser, L. Ladický, and M. Pollefeys, "Image based geo-localization in the alps," *International Journal of Computer Vision*, vol. 116, no. 3, pp. 213–225, 2016.

[5] R. Fedorov, P. Fraternali, and M. Tagliasacchi, "Mountain peak identification in visual content based on coarse digital elevation models," in *Proceedings of the 3rd ACM International Workshop on Multimedia Analysis for Ecological Data (MAED)*. ACM, 2014, pp. 7–11.

[6] 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, pp. 499–518, 2012.

[7] C. R. Steger, B. Steger, and C. Schär, "Horayzon v1.2: an efficient and flexible ray-tracing algorithm to compute horizon and sky view factor," *Geoscientific Model Development*, vol. 15, no. 17, pp. 6817–6840, 2022. [Online]. Available: https://gmd.copernicus.org/articles/15/6817/2022/

[8] L. De Floriani and P. Magillo, "Algorithms for visibility computation on terrains: A survey," *Environment and Planning B: Planning and Design*, vol. 30, no. 5, pp. 709–728, 2003.

[9] P. F. Fisher, "Extending the applicability of viewsheds in landscape planning," *Photogrammetric Engineering and Remote Sensing*, vol. 62, no. 11, pp. 1297–1302, 1996.

[10] Camera & Imaging Products Association, "Exchangeable image file format for digital still cameras (Exif) version 3.0," https://www.cipa.jp/std/documents/download_e.html?DC-008-2023-E, 2023, CIPA Standard DC-008-2023.

[11] J. P. Lewis, "Fast normalized cross-correlation," *Vision Interface*, vol. 10, no. 1, pp. 120–123, 1995.