

CamStatsPRO Viewshed Methodology

Methodology Version: 1.0 · 23 September 2026

What Is a Viewshed?

A viewshed is a map of the terrain that a camera can see. It is a geometric statement about the land, built from public elevation data, your camera's specifications, and obstructions that block the camera's view, including vegetation and any buildings or other obstructions identified during camera calibration.

What Does the Orange Shaded Area Represent?

The orange layer on the map represents terrain within the camera's modelled field of view and analysis range that has an unobstructed line of sight to the lens. It answers one question: could a straight line reach this spot without the land, the vegetation, or a mapped obstruction getting in the way?

The viewshed does not, however, guarantee that an observer can always 'see' the content at a location, because content may not be lit, in focus, large enough to make out, or unobscured by the weather.

What Do We Build It From?

Ground Elevation:

Ground elevation is derived from public federal survey data, including USGS 3DEP (United States Geological Survey 3D Elevation Program) and NASA SRTM (National Aeronautics and Space Administration Shuttle Radar Topography Mission) data. The 3DEP data is 1/3 arc-second (~10 meters) and the SRTM is 1 arc-second (~27 to 30 meters) at latitudes in the continental United States. Within 2 kilometers of the camera we request 3DEP on a 3-meter grid. The USGS service resamples that from the best source it holds for the location (1-meter lidar where it exists, 10-meter otherwise), so a 3-meter grid is the resolution we analyse at rather than a guarantee of 3-meter survey coverage. That fine data does two things, and it is worth being precise about which: it is averaged into the main 10-meter grid, where it makes those cells more accurate but no smaller, and it is kept at full 3-meter resolution for a separate close-range calculation that the Detail map is drawn from. The wider map sheets are computed on the 10-meter grid. SRTM is used where 3DEP does not reach, or when a 3DEP request fails. Elevation figures are orthometric (i.e., height above the geoid): 3DEP is referenced to NAVD88 (North American Vertical Datum of 1988) in the continental United States, and SRTM to the EGM96 (Earth Gravitational Model 1996) geoid.

Vegetation:

Elevation data describes bare ground. A viewshed computed on bare ground tells you what the terrain permits, which is not the same as what the camera will see. So we lay LANDFIRE (Landscape Fire and Resource Management Planning Tools, a joint USGS and US Forest Service program) Existing Vegetation Height over the elevation data and compute against that

combined surface. This applies to every viewshed we compute in the continental United States, which is the extent LANDFIRE covers; outside it, or on the occasion the vegetation service cannot be reached, the analysis falls back to bare terrain. Where the underlying elevation is SRTM we add only the portion SRTM is missing, about half the canopy height (0.51). SRTM is radar, and its C-band phase center sits partway into the canopy, so an SRTM 'ground' elevation in forest already includes some of the trees; adding the full height would count them twice. That 0.51 is a CamStats calibration parameter, not a published constant: we measured it over 100 points of Adirondack forest, comparing SRTM against 3DEP with LANDFIRE height as the reference. It is one site and one forest type. The 3DEP needs no correction at all.

Camera Specifications:

We collect the position (latitude and longitude), the height above the ground it stands on, the direction it faces, the horizontal width of its view, and its tilt. Before processing, the camera parameters can be manually calibrated by aligning the camera image with a 3D rendering of the same view. Once aligned, the user is given the option of manually digitizing features on the landscape (Vegetation, buildings, etc.) that obstruct the terrain view.

What Goes Into the Analysis

Four things determine the shape of the orange area, and nothing else does:

Terrain:

USGS 3DEP, or NASA SRTM where 3DEP does not reach.

Vegetation:

LANDFIRE Existing Vegetation Height.

Camera geometry:

location, height, azimuth, field of view and tilt.

Local obstructions:

the trees, buildings and other features traced during calibration.

Everything else on a CamStats map is geographic context, drawn on the sheet but taking no part in the visibility calculation: roads, water, place names, protected areas, building footprints and the rest. A building shows on the map because it helps you read the landscape. It blocks the view only if it was traced during calibration.

How Is the Viewshed Calculated?

The viewshed is calculated using the R3 Algorithm (The exhaustive class of viewshed algorithm: a line of sight is tested from the observer to every cell, rather than to the perimeter with the interior interpolated). This approach places a virtual observer at your camera's position and height. Next, a virtual line is cast from the camera to every cell within a specified range. If the ray connects to the observation cell without the land, vegetation, or mapped obstructions getting in the way, the cell is marked as being visible. Cells that are obstructed are marked as being hidden from the virtual observation point. Corrections for the Earth's curvature and the

atmosphere are applied at every cell. First, the fact that the Earth curves away beneath a level sight line (~eight meters over ten kilometers) is factored into the model. Second, the fact that the atmosphere bends light slightly downward and returns a little of it, is modelled by using a refraction coefficient ($k = 0.13$), which is consistent with industry standards. Atmospheric refraction bends the line of sight slightly downward, partially offsetting the effect of the Earth's curvature.

Finally, we add two safeguards of our own, which are not part of R3 itself. They are intended to stop isolated artifacts in the elevation model from casting long false shadows: the elevation data is lightly smoothed before analysis, and an obstruction must persist across three consecutive samples along a ray before the cell is declared hidden. Three samples is three cells of whichever grid the sheet is computed on (i.e., about 9 meters on the Detail map, about 30 meters on the wider sheets, and about 90 meters where SRTM is used). These safeguards can also suppress very narrow real terrain features, so they trade a small amount of local detail for greater stability in the resulting viewshed.

What We Do Not Model

- The viewshed analysis assumes a clear weather day and does not account for atmospheric conditions (e.g., precipitation, haze, etc.) that may obstruct the view.
- The analysis does not inherently account for seasonal variations in vegetation (i.e., leaf on/off). Users may re-trace vegetation from winter imagery and reprocess, but that replaces the existing viewshed rather than producing a second seasonal version: there is one current viewshed per camera.
- Human-made structures (e.g., buildings, towers, etc.) may be incorporated as digitized obstructions but are not automatically included in the viewshed analysis.
- The elevation grid sets the scale of terrain detail the model can resolve: about 3 meters within 2 kilometers on the Detail map, about 10 meters on the wider sheets, and about 30 meters where SRTM is used. Features smaller than the source data or the analysis grid may not be represented reliably. Grid resolution is not the same thing as horizontal or vertical accuracy.
- The model assumes an undistorted lens view.
- The model accounts for the height of the camera, but it does not add height at the far end. A cell is tested at the modelled surface (i.e., bare ground in the open, the top of the canopy in forest), so a location shown as visible does not mean the ground beneath the trees is visible, nor that a built feature (e.g., a tower) at that location would be.
- The model does not account for the lighting conditions at a distant location, and the degree to which a location is lit or in shadow.
- Changes after processing (e.g., vehicles, cranes, vegetation growth, construction, or movement of the camera).

Disclaimer

The viewshed provides a reasonable approximation of the locations that are visible in a camera view, as well as the context of what lies beyond the visible field of view. It is not intended for safety-critical decisions, or for legal, boundary or privacy determinations. It may not be suitable, without independent verification, for surface feature detection or recognition claims, or for regulatory submissions that require surveyed accuracy.

Assumptions and Parameters

Physical Constants

Earth radius	6,371,000 m
Refraction coefficient (k)	0.13
Effective Earth radius	$R / (1 - k) \approx 7,323,000$ m

Elevation and Vegetation

Primary elevation	USGS 3DEP, 1/3 arc-second (~10 m)
Near field	3DEP at 3 m, within 2 km of the camera
Fallback elevation	NASA SRTM, 1 arc-second (~27-30 m)
Vertical datum	orthometric — NAVD88 for 3DEP in the continental US, EGM96 for SRTM
Vegetation	LANDFIRE 2024 Existing Vegetation Height, 30 m (continental US)
Canopy already in SRTM	0.51 — a CamStats calibration parameter; only the remainder is added

Camera

Mounting height	as entered; 4 m assumed if not given
Range	as entered; 20 km assumed if not given
Maximum range	100 km
Horizontal field	from the azimuth pair you set
Vertical field	derived from the horizontal field and image shape
Image shape, if unknown	16:9 assumed

Numerical

Elevation pre-smoothing	Gaussian, $\sigma = 0.7$ cells
Samples needed to declare a cell hidden	3 consecutive

Sources and Credit

USGS 3D Elevation Program (3DEP)	Ground elevation across the United States
NASA SRTM, served through AWS Terrain Tiles	Fallback elevation surface outside 3DEP coverage. Required attribution: "SRTM data courtesy of the U.S. Geological Survey."
LANDFIRE 2024 Existing Vegetation Height	The vegetation layer, across the continental United States
WSL Monoplotting Tool	Established the single-photograph georeferencing technique our obstruction tracing follows

Built in Python, on NumPy, SciPy, Numba and rasterio/GDAL.