# Camera Based Visibility Estimation

Thomas Sutter[1], Fabian Nater[1] and Christian Sigg*[2]

[1]Logitech Upicto GmbH
[2]Swiss Federal Office of Meteorology and Climatology MeteoSwiss

**Abstract**

MeteoSwiss operates a network of panorama cameras positioned along Switzerland's main flight paths. This network supports the meteorologist in creating general aviation forecasts, which are a summary of expected prevailing visibility and cloud base height along the path. We present preliminary results for estimating visibility automatically from panorama camera images. The estimation algorithm is based on Koschmieder's law, which relates apparent contrast of an object to its distance from the observer. Local contrast information is computed from image patches using a standard measure for human contrast perception. The observation distance is computed from the known camera position and a digital elevation model. We evaluate the algorithm by comparison to expert judgements provided by three trained observers, on a data set of diverse weather conditions and topographies. We show that the estimate of the algorithm is typically in agreement with the expert judgments. We also discuss the limitations of the algorithm and present cases where the estimation accuracy is unsatisfactory. Finally, we present an automated image quality control mechanism that can reliably identify some of those cases.

## 1 Introduction

MeteoSwiss, the Federal Office of Meteorology and Climatology, operates a network of panorama cameras positioned along Switzerland's main flight paths. Each camera is mounted on a rotating head and acquires a panoramic sequence of images once every ten minutes (see Figure 1). One purpose of the camera network is to support the meteorologist in creating general aviation forecasts (GAFORs). The GAFOR predicts the relevant weather conditions for visual flight, and is a summary of expected prevailing visibility and cloud base height along the path. The current visibility and cloud base height form the basis of the forecast.

*christian.sigg@meteoswiss.ch (corresponding author)

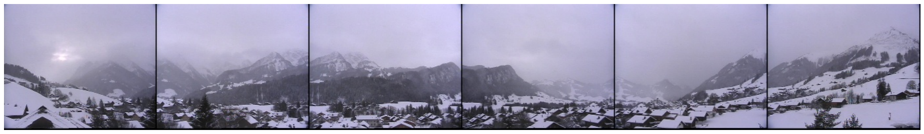

Figure 1: A panoramic sequence of six images obtained by the camera stationed at Chateau d'Oex on January 30th 2015 at 8:00 a.m. The local image contrast falls quickly with increasing distance from the camera.

*Visibility* is defined as the greatest distance at which a black object of suitable dimensions (located on the ground) can be seen and recognized when observed against the horizon sky [1]. The object is no longer visible when enough *airlight* is scattered into the path from the object to the observer. Since the panorama cameras record natural scenes without dedicated black reference objects, the closely related definition of *visual range* is more useful for our application. Visual range is defined as the distance at which the contrast of a given object with respect to its background is just equal to the contrast threshold of an observer [1]. The apparent object contrast falls below the perceptible threshold when light emanating from the object towards the observer is attenuated sufficiently, compared to the scattered-in airlight.

In the MeteoSwiss network, visibility is currently estimated from visual observations and measured using transmissometers. The latter are laser-based measurement systems for the transmittance, based on the attenuation of emitted light due to scattering. The transmissometer provides a local measurement of the atmosphere at the measurement site.

MeteoSwiss is engaged in an effort to increase the spatial and temporal availability of the visibility parameter, which necessitates migrating from manual observations to automated measurements. We present preliminary results for estimating visibility automatically from panorama camera images. The potential benefit is two-fold. First, we re-use the already existing camera network, instead of operating additional instruments. Second, panorama cameras observe more of the atmosphere compared to transmissometers, leading to more representative measurements for weather situations where the atmosphere is significantly non-homogeneous.

## 2 Estimation Method

The estimation algorithm is based on Koschmieder's law, which relates the apparent contrast $C_x$ of an object to its distance $x$ from the observer

$$C_x = C_0 e^{-\sigma x} \tag{1}$$

$C_0$ is the inherent object contrast and $\sigma$ is the *extinction coefficient* that characterises atmospheric absorption and scattering, cf. Equation (9.9) in [1].

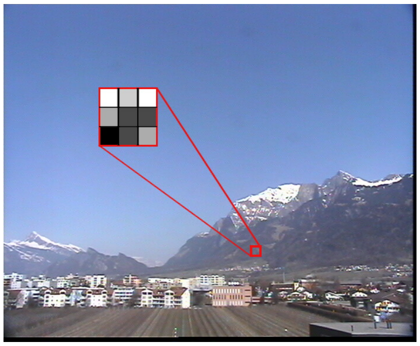

(a) The apparent contrast is computed on image patches obtained by sliding a window across the camera image.

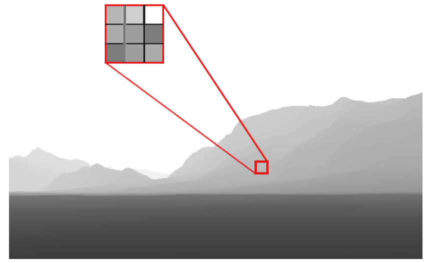

(b) The distance is computed from the corresponding image patch of the *depth map*, where pixel intensities encode the distance from the camera.

Figure 2: Koschmieder's law (1) relates apparent contrast of an object to its distance from the observer.

The local contrast information $C_x$ is computed from camera image patches (see Figure 2a) using a standard measure for human contrast perception [4]. This measure is based on the root-mean-squared contrast of the luminance of the patch pixels, center-weighted by a raised cosine window. See Equations (1)–(3) in [4] for details.

The observation distance $x$ is computed from the known camera position and the EU-DEM digital elevation model [2]. The free camera parameters of rotation and field of view are determined once during the algorithm setup phase with the help of a GUI (see Figure 3). The DEM is then rendered using the camera parameters to produce a depth map that corresponds to the camera image (see Figure 2b).

We also investigated using *pose estimation* [5] to determine all camera parameters (including position). In this approach, a set of landmarks (such as a mountain top) with known coordinates are identified in the camera image, and the camera parameters are then determined automatically by numerical optimization. However, we found that the interactive matching takes about the same manual effort and achieves a greater precision, because easily identifiable but distant landmarks often led to bad numerical conditioning for the pose optimization step.

Using a sliding window approach, the local contrast is computed for each image patch, and the corresponding depth is taken as the median of all depth values in the patch. Patches that contain the sky or a depth discontinuity are discarded. As can be seen from Figure 4, the maximum observed contrast decays exponentially with distance, in accordance with Koschmieder's law (1). The rate of decay is estimated by fitting the parametrized curve

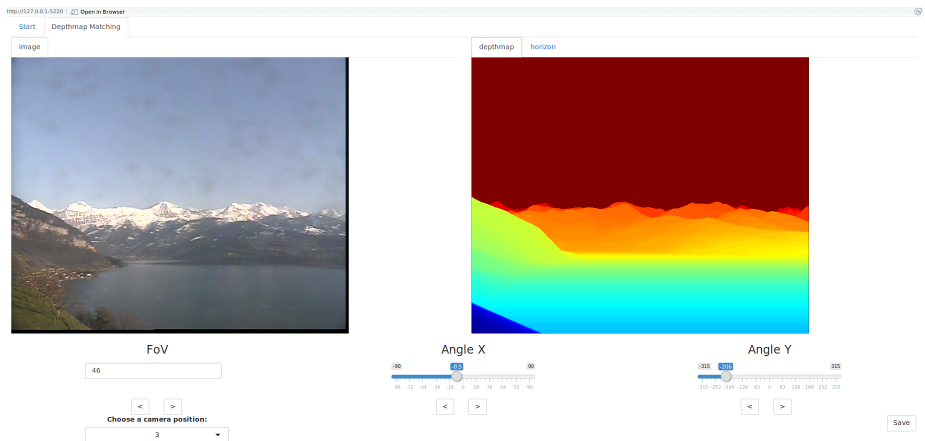


Figure 3: The camera parameters of rotation and field of view are determined by interactive matching of the rendered topography to a reference camera image. The GUI provides a side-by-side view of the camera image and the rendered topography, or an overlay of the rendered horizon onto the camera image (not shown). The camera parameters are optimized interactively until a sufficient matching accuracy is achieved.

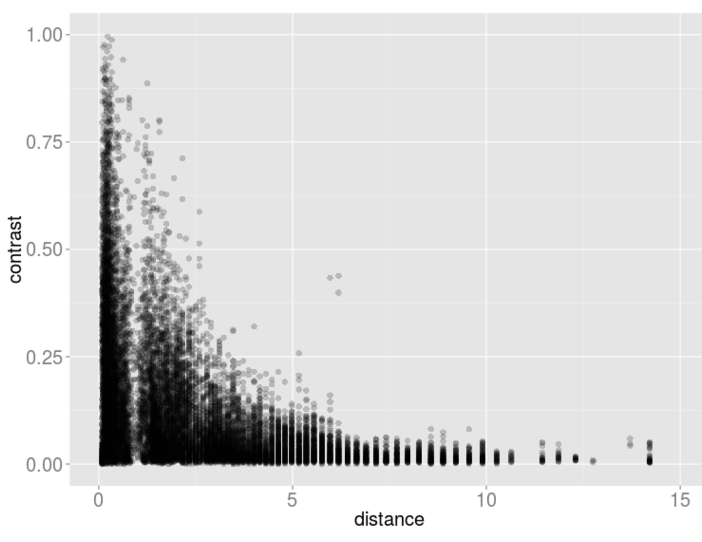


Figure 4: Distance versus contrast pairs for patches extracted from the image sequence depicted in Figure 1. The outliers at approx. 6 km are due to objects in the view of the camera that are not present in the DEM, and thus have an incorrect distance assigned to the contrast value. Manual correction or masking of the generated depth map can remove these outliers (not shown).

$$(\alpha - \beta)\, e^{-\sigma x} + \beta \tag{2}$$

into the 90th percentile of the contrast values, where $\sigma$ is again the extinction coefficient and $\beta$ is the contrast value of opaque fog. $\alpha$ is estimated from the contrast distribution of objects closest to the camera. Fitting the curve to a high percentile instead of the maximum provides robustness against outliers in the scatter plot.

Finally, the visibility range $P$ is estimated as

$$P = -\frac{\ln \epsilon}{\sigma} \tag{3}$$

where $\epsilon$ is the minimal contrast threshold that the human eye can detect, cf. Equation (9.6) in [1]. $\epsilon$ is usually assumed to be in the range of 0.02 to 0.2 [1].

# 3 Automated Quality Control

For a fully autonomous application of the algorithm, it is necessary to detect camera images that are unsuitable for visibility estimation. All images therefore pass through an automated quality control (QC) step, and the visibility is not estimated for image sequences of insufficient quality. These are images where the illumination of the scene is insufficient, that are not properly aligned, or images where the lens is partially or fully occluded.

A misalignment occurs if the camera rotation head becomes blocked, e.g during a snow storm. The image no longer aligns with the depth map, which would lead to invalid depth contrast pairs in the visibility estimation step. The QC corrects small misalignments und detects large ones by means of *phase correlation* [3] of the current image to a reference image. The reference is the same image that has already been used in the depth map alignment step.

Partial or full lens occlusion occurs if water droplets, snow or ice collects on the lens. Lens occlusions are detected by *interest point matching* for images in the panoramic sequence that have no scene overlap. Interest points are detected using the Good Features to Track algorithm [7] and matched using BRISK descriptors [6]. Matching interest points with similar pixel coordinates reliably indicate lens occlusions.

# 4 Evaluation

We evaluate the algorithm by comparison to judgements provided by three expert observers, on a data set of diverse weather conditions and topographies. Using the labeling tool depicted in Figure 5, each expert provided estimates of the *prevailing visibility*, the visibility which holds for at least half of the field of view. The data set contains 130 image sequences from five cameras located in different topographies (Swiss Plateau, valley and mountain pass), ranging from

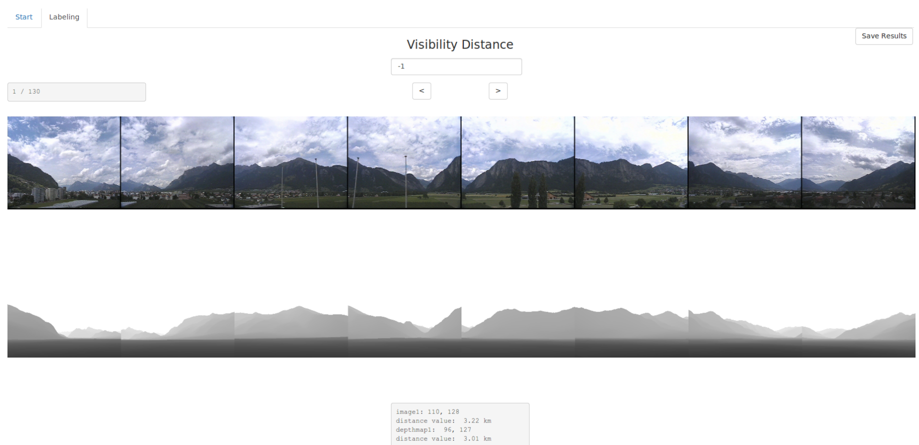


Figure 5: The experts provided estimates of the prevailing visibility with the help of a labeling tool. Panoramic image sequences were presented in randomized order. The expert can read depth values in the image or in the depth map by using mouse clicks, and enter the prevailing visibility estimate in the text box. Labels for the whole data set are exported in CSV format.

winter to summer and fog to clear sky. The images were presented in randomized order, and five image sequences were presented twice to roughly assess the labeling consistency of each expert.

Our limited evaluation of intra-expert consistency suggests that all experts are consistent in their judgments (see Figure 6), but also suggests that there is a significant inter-expert bias. Figure 7 compares the algorithm estimates with the expert labels on 26 image sequences from the Chateau d'Oex camera. Overall there is a good agreement between the algorithm and the experts, but there are also significant differences. An analysis of the image sequences 11 to 15 reveals a common cause. Since those images were recorded in spring, there was still snow on the mountain tops, while the valley was already free of snow. The partial snow cover creates a high local contrast for the distant mountain tops, which had a marked influence on the algorithm but not the experts.

Table 1 summarizes the overall performance of the algorithm, compared to the three experts. We used three different error measures in the evaluation.

1. **GAFOR state error.** The difference in GAFOR states between two estimates. The error is zero if the estimates are in the same GAFOR state, and counts the number of states if they are different. E.g. estimates of Oscar and Mike result in a state error of two.

2. **Absolute distance error.** The difference in kilometers between estimates.

3. **Relative distance error.** The difference in kilometers between esti-

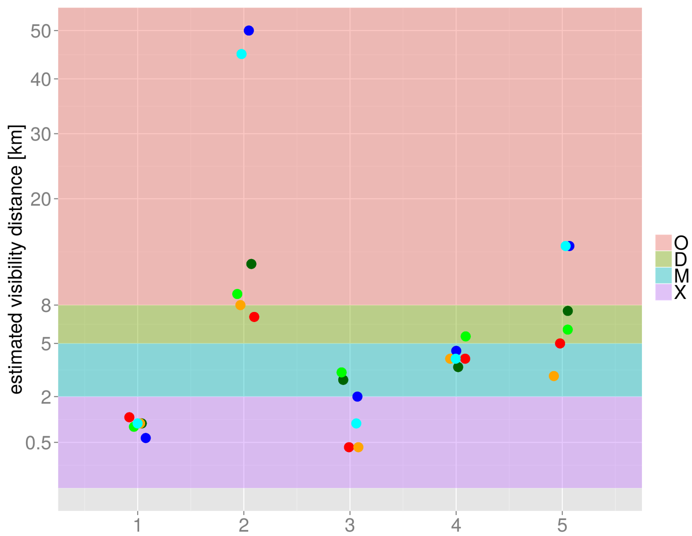


Figure 6: Intra-expert labeling consistency for five images of the data set. Labels from the same expert are in light – dark blue, light – dark green and orange – red, respectively. The background colors correspond to the GAFOR classes Oscar (> 8 km), Delta (> 5 km), Mike (> 2km ) and X-ray (< 2km).

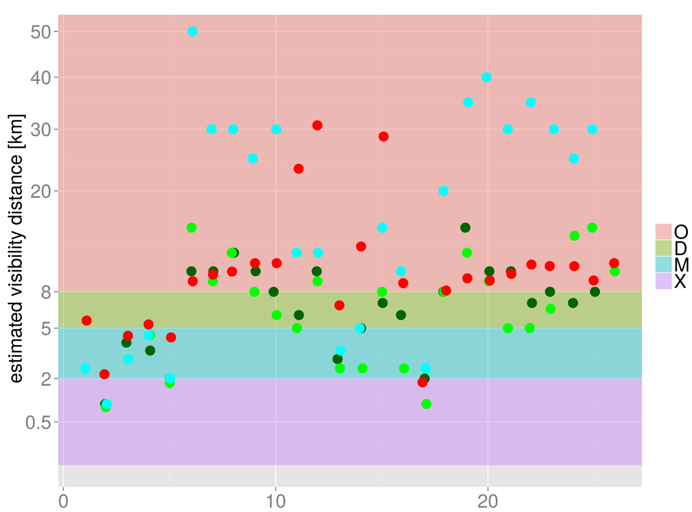


Figure 7: Comparison of the algorithm estimates (red) to the expert labels (cyan, light and dark green) for 26 images from the Chateau d'Oex camera. The estimates of the algorithm are typically in agreement with the experts, but there are significant differences for the image sequences 1 and 11 to 15.

| | Algorithm | Expert 1 | Expert 2 | Expert 3 |
|---|---|---|---|---|
| GAFOR state error | 0.47 | 0.60 | 0.50 | 0.52 |
| Absolute distance error [km] | 3.56 | 6.35 | 6.18 | 9.53 |
| Relative distance error [km] | 0.71 | 0.97 | 0.92 | 1.29 |

Table 1: Overall performance of the algorithm, compared to the three experts. See the text for a description of the error measures and results.

mates, after a square root transform of each distance. This error measure accounts for the fact that differences have a bigger impact on low visibilities.

The results for the experts are computed as follows. For each image sequence, the median estimate of all three experts is calculated. Every expert's estimate is then compared to the median, and the errors are averaged over the whole data set of 130 image sequences. The results for the algorithm are computed similarly, by comparison to the median estimate of all experts and averaging of the error over the whole data set. Table 1 shows that the algorithm achieves a similar overall error compared to the experts. It also shows that expert one and expert two are more similar to each other than to expert three.

# 5 Conclusions

The camera based visibility estimation algorithm is currently in the pre-operational phase. It fulfills our requirement of a small setup effort per camera and doesn't need any tuning of its parameters to local conditions. It produces visibility estimates that are comparable to human observers for a wide range of topographies and weather conditions. And last but not least, the implementation of the algorithm is efficient enough to provide estimates for the whole camera network every ten minutes with modest hardware resources.

But our evaluation has also identified conditions where the estimation error is large. These are conditions that violate basic assumptions of the underlying physical model, such as nighttime, significantly non-diffuse illumination or a systematic bias in the object contrast. The automated quality control is currently able to identify only a subset of the conditions where estimating the visibility is not possible.

It also remains to be seen if the camera network is reliable enough for this new application. Intervention and maintenance procedures are already established to deal with hardware or software malfunctions in the network.

Our project goals take this uncertainty into account. The algorithm estimates are meant to support the meteorologist in creating GAFORs by providing dense spatial and temporal visibility information. The estimates therefore don't have to satisfy guaranteed error limits before they become useful. The preliminary results presented in this report give us confidence that camera based visibility estimation can generate a substantial new benefit from existing infrastructure.

## Acknowledgments

Produced using Copernicus data and information funded by the European Union - EU-DEM layers.

## References

[1] *Manual on the Global Observing System.* WMO-No. 544. World Meteorological Organization, Geneva, 2003.

[2] European Environment Agency. Digital elevation model over europe (eu-dem). `http://www.eea.europa.eu/data-and-maps/data/eu-dem`, 2013.

[3] Hassan Foroosh, Josiane B Zerubia, and Marc Berthod. Extension of phase correlation to subpixel registration. *IEEE transactions on image processing*, 11(3):188–200, 2002.

[4] Robert A Frazor and Wilson S Geisler. Local luminance and contrast in natural images. *Vision research*, 46(10):1585–1598, 2006.

[5] Robert M Haralick, Hyonam Joo, Chung-Nan Lee, Xinhua Zhuang, Vinay G Vaidya, and Man Bae Kim. Pose estimation from corresponding point data. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(6):1426–1446, 1989.

[6] Stefan Leutenegger, Margarita Chli, and Roland Y Siegwart. Brisk: Binary robust invariant scalable keypoints. In *2011 International conference on computer vision*, pages 2548–2555. IEEE, 2011.

[7] Jianbo Shi and Carlo Tomasi. Good features to track. In *Computer Vision and Pattern Recognition, 1994. Proceedings CVPR'94., 1994 IEEE Computer Society Conference on*, pages 593–600. IEEE, 1994.