Aircraft based Stereographic Reconstruction of 3D Cloud Geometry

Tobias Kölling[1], Tobias Zinner[1], and Bernhard Mayer[1]

[1]Ludwig Maximilians Universität, Meteorologisches Institut, München, Germany

Correspondence: T. Kölling (tobias.koelling@physik.uni-muenchen.de)

Abstract. This work describes a method to retrieve location and geometry of clouds using RGB images from a video camera on an aircraft and data from the aircraft's navigation system. Opposed to ordinary stereo methods where two cameras with fixed relative position at a certain distance are used to match images taken at the exact same moment, this method uses only a single camera and the aircrafts movement to provide the needed parallax. Advantages of this approach include a relatively simple installation on an aircraft and the possibility to use different image offsets, even larger than the size of the aircraft. Detrimental effects are the evolution of observed clouds during the time offset between two images as well as the background wind. However we will show that to some extent usable wind information can also be recovered. Our method allows the derivation of cloud top geometry which can be used, e.g., to provide location and distance information for other passive cloud remote sensing products. In addition it can also improve retrieval methods by providing cloud geometry information useful for the correction of 3D illumination effects. We show that this method works as intended by comparison to data from a simultaneously operated lidar system. The stereo method provides lower heights than the lidar method, on average by $126\,\mathrm{m}$. This behaviour is expected as the lidar method has a lower detection limit (leading to greater cloud top heights for the downward view) while the stereo method also retrieves data points on cloud sides and lower cloud layers (leading to lower cloud heights). Systematic errors across the measurement swath contribute less than $50\,\mathrm{m}$.

1 Introduction

In cloud remote sensing applications, where observed clouds and observer are at different locations, cloud detection, cloud location and cloud geometry are parameters of high importance for all consecutive retrieval products. These parameters themselves govern characteristics like cloud mass or temperature and subsequently thermal radiation budget and thermodynamic phase. Typically passive remote sensing using spectral information is used to retrieve cloud properties including cloud optical thickness, effective droplet radius, thermodynamic phase or liquid water content. However, these methods can not directly measure the cloud's location. Finally, unknown cloud location and geometry are the central reason for uncertainties in microphysical retrievals because of the complex impact of 3D structures on radiative transport (e.g. Várnai and Marshak, 2003; Zinner and Mayer, 2006). The classic method of handling complex, inhomogeneous parts of the atmosphere (e.g. typical MODIS retrievals) is to exclude these parts from further processing. This of course can severely limit the applicability of such a method. As shown by Ewald (2016); Ewald et al. (2018) the local cloud surface orientation affects retrieval results. In particular, it is shown that changes in surface orientation and changes in droplet effective radius produce a very similar spectral

response. Thus an independent measurement of cloud surface orientation would very likely improve retrieval results on droplet effective radius.

As location and geometry information is of such a great importance, a couple of different approaches to get this information can be found. Among these are active methods using lidar or radar. Fielding et al. (2014) and Ewald et al. (2015) show how
5 3D distributions of droplet sizes and liquid water content of clouds can be obtained by the use of a scanning radar. Ewald et al. (2015) even visually demonstrate the quality of their results by providing simulated images using the retrieved 3D distributions as input and comparing them to actual photographs. Major downside of this approach is the limited scanning speed. Consequently these methods are especially difficult to employ on fast moving platforms. For this reason, the typical implementations of cloud radar and lidar on aircraft only provide data directly below the aircraft.

10 Passive methods are often less accurate but can cover much larger observation areas in shorter measurement times. They typically either use spectral features of the signal or use observations from multiple directions. MODIS (Moderate Resolution Imaging Spectroradiometer) cloud top height for example uses thermal infrared images to derive cloud top brightness temperatures (Strabala et al., 1994). Using assumed cloud emissivity and atmospheric temperature profiles, cloud top heights can be calculated. Várnai and Marshak (2002) used gradients in the MODIS brightness temperature to further classify observed
15 clouds into "illuminated" and "shadowy" clouds. Another spectral approach has been demonstrated amongst others by Fischer et al. (1991) and Zinner et al. (2018) using oxygen absorption features to estimate the travelled distance of the observed light through the atmosphere. Assuming most of the light gets reflected at or around the cloud surface, this information can be used to calculate the location of the clouds surface.

Other experiments (e.g. Beekmans et al., 2016; Crispel and Roberts, 2018; Romps and Öktem, 2018) use multiple ground
20 based all-sky cameras and apply stereophotogrammetry techniques to georeference cloud fields. Due to the use of multiple cameras, it is possible to capture all images at the same time - a big advantage when observing moving and changing clouds.

Spaceborne stereographic methods have been employed e.g. for the Multi-angle Imaging SpectroRadiometer (MISR) (Moroney et al., 2002) and the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) (Seiz et al., 2006). MISR features 9 different viewing angles which are captured during 7 minutes of flight time. During the long time period of
25 about one minute between two subsequent images the scene can change substantially. Clouds in particular are transported and deformed by wind, which adds extra complexity on stereographic retrievals. The method by Moroney et al. (2002) adresses this problem by tracking clouds along all the perspectives and derivation of a coarse wind field at a resolution of about 70 km. ASTER comes with only two viewing angles but still takes about 64 s to complete one image pair. Consequently, the method by Seiz et al. (2006) uses other sources of wind data (e.g. MISR or geostationary satellite data) to correct for cloud motion
30 during the capturing period.

In this work, we present a stereographic method which uses 2D images taken from a moving aircraft at different times to find the georeferenced location and structure of the cloud surface facing the observer. This method neither depends on estimates of the atmospheric state nor does it depend on assumptions on the cloud shape. Contrasting to the spaceborne methods, our method only takes 1 s for one image pair. Due to the relatively low operating altitude of an aircraft compared to a satellite, the
35 observation angle changes rapidly enough to use two successive images without the application of a wind correction method.

As we employ an 2D imager with a wide field of view, each cloud is captured from many different perspectives (up to about 100 different angles, depending on the distance between aircraft and cloud). Due to the high number of viewing angles, it is possible to derive geometry information of partly occluded clouds. Furthermore, this allows to simultaneously derive an estimate of the 3D wind field and use it to improve the retrieval result.

5   We demonstrate the application of our method to data obtained in the NARVAL-II and NAWDEX field campaigns (Stevens et al., submitted; Schäfler et al., 2018). In these field campaigns, the hyperspectral imaging system specMACS has been flown on the HALO aircraft (Ewald et al., 2016; Krautstrunk and Giez, 2012). The deployment of specMACS, together with other active and passive instrumentation, aimed at a better understanding of cloud physics including water content, droplet growth, cloud distribution and geometry. The main component of the specMACS system are two hyperspectral line cameras.

10 Depending on the particular measurement purpose, additional imagers are added. The hyperspectral imagers are operating in the wavelength range of $400 - 1000\,\mathrm{nm}$ and $1000 - 2500\,\mathrm{nm}$ at a spectral resolution of a few nanometers. Further details are described by Ewald et al. (2016). During the measurement campaigns discussed in this work, the two sensors where looking in nadir perspective and have been accompanied by a 2D RGB imager with about twice the spatial resolution and field of view. In this work, we focus on data from the 2D imager, because it allows observing the same cloud from different angles.

15   In section 2 we briefly explain the measurement setup. Section 3 introduces the 3D reconstruction method and section 4 presents a verification of our method. For geometric calibration of the camera we use a common approach. Nonetheless, as the geometry reconstruction method is very sensitive to calibration errors, we provide a short summary of our calibration process in appendix A. We used the OpenCV library (Bradski, 2000) for important parts of this work. Details are listed in appendix B.

2   Measurement Setup

20 During the NARVAL-II and NAWDEX measurement campaigns specMACS was deployed on board the HALO aircraft. As opposed to Ewald et al. (2016), the cameras have been installed in a nadir looking perspective. The additional 2D imager (Basler acA2040-180kc camera + Kowa LM8HC objective) has been set up to provide a full field-of-view of approximately $70°$ with 2000 by $2000\,\mathrm{pixels}$ and data acquisition frequency at $1\,\mathrm{Hz}$. To cope with the varying brightness during and between flights, the camera's internal exposure control system has been used.

25   Additionally, the WALES lidar system (Wirth et al., 2009), the HALO Microwave Package HAMP (Mech et al., 2014), the Spectral Modular Airborne Radiation measurement sysTem SMART (Wendisch et al., 2001) and an AVAPS dropsonde system (Hock and Franklin, 1999) was part of the campaign specific aircraft instrumentation. The WALES instrument is able to provide an accurate cloud top height and allows to directly validate our stereo method as described in section 4.

3   3D reconstruction

30 The goal of our 3D reconstruction method is to find georeferenced points which are part of a cloud surface at a specific time in an automated manner. Input data are geometrically calibrated images from a 2D camera fixed to the aircraft. As the aircraft
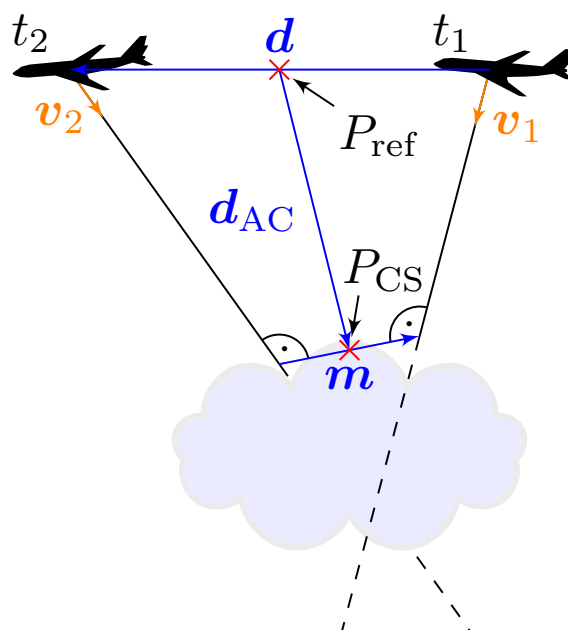
Figure 1. Schematic drawing of the stereographic geometry. Images of clouds are taken at two different times from a fast moving aircraft. Using aircraft location and viewing geometry, a point P_{CS} on the clouds surface can be calculated. Note that the drawing is not to scale: $\boldsymbol{d}$ is typically around 200 m, $\boldsymbol{d}_{AC}$ in the order of 5 km and $\boldsymbol{m}$ is a description of mis-pointing and in the order of only a few meters.

flies, pictures taken at successive points in time show the same clouds from different perspectives. A schematic of this geometry is shown in fig. 1. The geometric calibration of the camera and the rigid mounting on the aircraft allows to associate each sensor pixel with a viewing direction in the aircraft's frame of reference. Using the aircraft's navigation system, all relevant distances and directions can be transformed into a geocentric reference frame in which most of the following calculations are performed.

5   The reconstruction method contains several constants which are tuned to optimize its performance. Their values are listed in table 1.

In order to perform a stereo positioning, a location on a cloud must be identified in multiple successive images. We start by identifying points which are likely to be detectable in another image. The surrounding of these points must show a contrast in both directions of the image plane. Ideally, these points should be evenly distributed over the image. For the identification of

10   these points, we use the method by Shi and Tomasi (1994). Candidate points are sorted by quality, points which have better candidates at a distance of less then r_{min} are removed from the list and the remaining best N_{points} are taken. For these initial points, matches in the following image are sought using the optical flow algorithm described by Lucas and Kanade (1981). In particular, we use a pyramidal implementation of this algorithm as introduced by Bouguet (2000). If no match can be found, the point is rejected.

The locations of the two matching pixels define the viewing directions v_1 and v_2 in fig. 1. The distance travelled by the aircraft between two images is indicated by d. Under the assumption that the aircraft travels much faster than the observed clouds, an equation system for the position of the point on the cloud's surface P_{CS} can be found. In principle, P_{CS} is located at the intersection of the two viewing rays along v_1 and v_2, but as opposed to 2D space in 3D space there is not necessarily

5   an intersection, especially in presence of inevitable alignment errors. We relax this condition by searching for the shortest distance between the viewing rays. The shortest distance between two lines can be found by introducing a line segment which is perpendicular to both lines. This is the mis-pointing vector m. The point on the cloud's surface P_{CS} is now defined at the center of this vector. If for further processing a single point for the observer location is needed, the point P_{ref} at the center of both aircraft locations is used.

10   This way, many points potentially located on a cloud's surface are found. Still, these points contain a number of false correspondences between two images. During turbulent parts of the flight, errors in synchronization between aircraft navigation system and camera will lead to errors in calculated viewing directions. To reject these errors, a set of filtering criteria is applied (the threshold values can be found in tab. 1). Based on features of a single P_{CS}, the following points are removed:

- P_{CS} position is behind the camera or below ground

15   - absolute mis-pointing $|m| > m_{\text{abs}}$

- relative mis-pointing $|m|/|d_{\text{AC}}| > m_{\text{rel}}$

Figure 2 shows long tracks starting at a location on the cloud surface. These tracks follow the relative cloud position through up to 30 captured images. The tracks are generated from image pairs by repeated tracking steps originating at the t_2 pixel position of the previous image pair. Using these tracks, additional physics based filtering criteria can be defined.

20   Each of these tracks contains many P_{CS} points which should all describe the same part of the cloud. As clouds move with the wind, the points P_{CS} do not necessarily have to refer to the same geocentric location. In fact, figure 3 shows similar tracks to figure 2, but now these tracks consist of the corresponding cloud surface points P_{CS} in a 3D geocentric coordinate system. These tracks do not jump randomly around some center as they would do if they are only influenced by measurement errors. Instead, they roughly follow some preferred direction.

25   For the observation period (up to 30 s) it is assumed that the wind moves parts of a cloud on almost straight lines at a relatively constant velocity (which may be different for different parts of the cloud). Then, sets of P_{CS} can be filtered for unphysical movements. The filtering criteria are

- *velocity jumps*: the fraction of maximum to median velocity of a track must be less than v_{jump}

- *count*: the number of calculated P_{CS} in a track must be above a given minimum N_{min}

30   - *distance uncertainty*: the distance d_{AC} between aircraft and cloud may not vary more than d_{abs} or the relative distance variation with respect to the average distance of a track must be less than d_{rel}
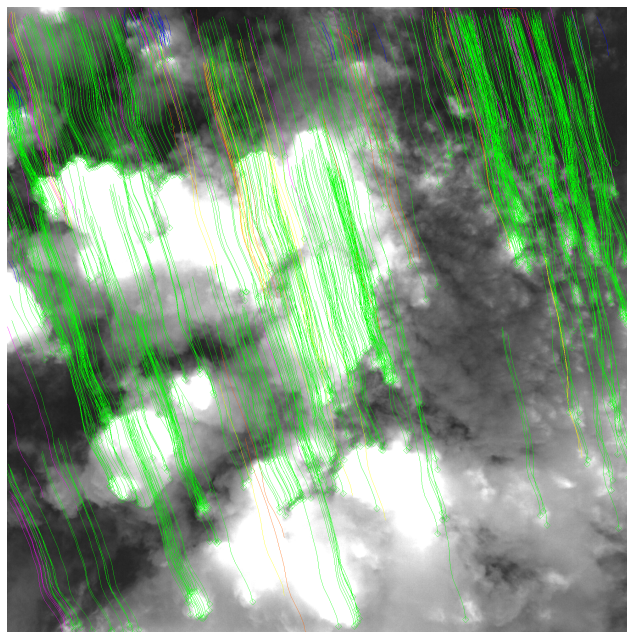
Figure 2. Image point tracking, every line in this image represents a cloud feature which has been tracked along up to 30 images. The images used have been taken on the NAWDEX flight RF07 (2016-10-06 09:32:15 UTC) in an interval of 1 s. The color indicates if the given thresholds have been met as follows: green $\hat{=}$ OK; yellow $\hat{=}$ mis-pointing exceeded; pink $\hat{=}$ velocity jumps too large; blue $\hat{=}$ not enough points; orange $\hat{=}$ distance uncertainty too large.
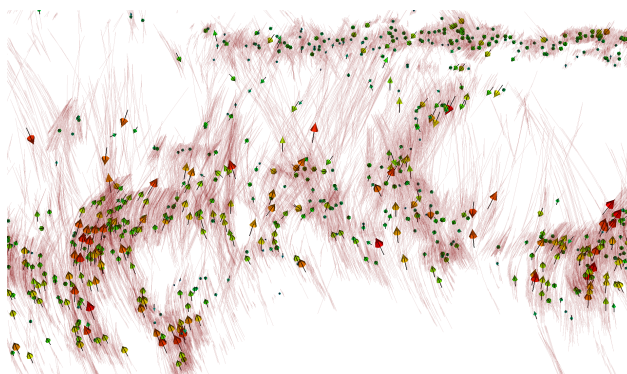
Figure 3. This figure shows a closeup of the retrieved data from the scene in figure 2. The view is not from the camera perspective, but sideways, with a slight downwards tilt. In dark red, P_{CS} of a track are connected to a line. The arrows show a local average of the resulting point movement vectors.

During measurements close to the equator, typical during the NARVAL-II campaign, the sun is frequently located close to the zenith. In this case, specular reflection of the sunlight at the sea surface produces bright spots, known as sunglint and illustrated in figure 4. Due to waves on the ocean surface, these regions of the image also produce strong contrasts. It turns

Table 1. Filter thresholds

name	value
N_{points}	1000
r_{min}	5 px
m_{abs}	20 m
m_{rel}	1.5×10^{-3}
v_{jump}	3
N_{min}	5
d_{abs}	250 m
d_{rel}	7×10^{-2}
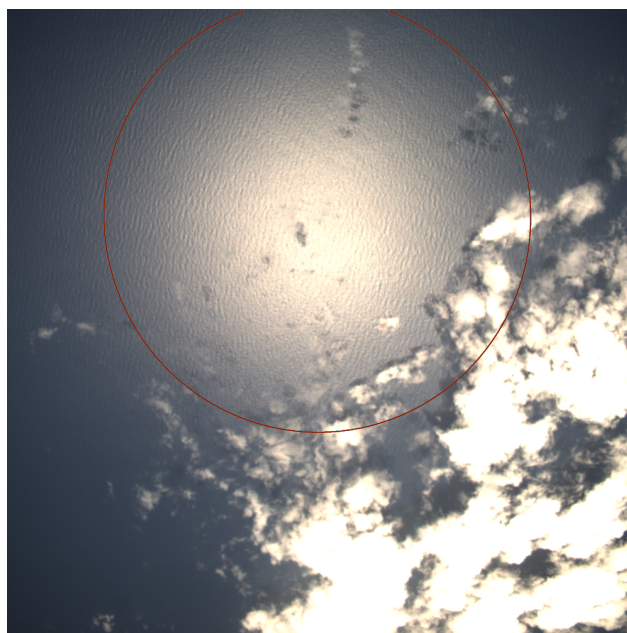
Figure 4. At low latitudes, close to the local noon as on the NARVAL-II flight RF07 (2016-08-19 15:06:13 UTC), the specular reflection of the sun on the ocean surface (sunglint) produces bright spots and high contrasts on the waves tails. While the bright spots can visually hide clouds, the contrasts create useless initial tracking points. The latter are mitigated by calculating the region of a potential sunglint (shown as red contour) and masking that region before the images are processed.

out that such contrasts are preferred by the Shi and Tomasi algorithm for feature selection, but are useless in order to estimate the cloud surface geometry. To prevent the algorithm from tracking these points, the image area in which bright sunglint is to be expected is estimated using the bidirectional reflectance distribution function (BRDF) by Cox and Munk (1954) included
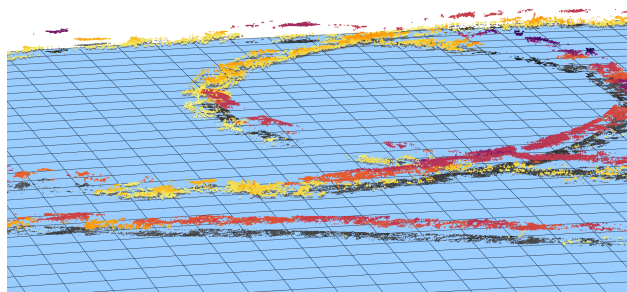
Figure 5. The collection of all $\bar{P}_{\mathrm{CS}}$ form a point cloud. Here, a scene from the second half of the NARVAL-II flight RF07 is shown. The colors indicate the points height above the WGS84 reference ellipsoid (indicated as blue surface), which also highlights the ability to detect multiple cloud layers at once. The gray dots are a projection of the points onto the surface to improve visual perception.

in the libRadtran package (Mayer and Kylling, 2005; Emde et al., 2016). The resulting area (indicated by a red line in fig. 4) is masked out of all images before any tracking is performed. Masking out such a large area from the camera image seems to be a wasteful approach. In fact it is acceptable, because due to the large viewing angle of the camera, all masked-out clouds are almost certainly visible at a different time in another part of the image. Therefore, these clouds can still be tracked in the

5 presence of sunglint.

After all filtering, a final mean cloud surface point $\bar{P}_{\mathrm{CS}}$ is derived from each track as the centroid of all contributing cloud surface points. The collection of all $\bar{P}_{\mathrm{CS}}$ form a point cloud in a Cartesian 3D reference coordinate frame which is defined relative to a point on the earth's surface. Such a point cloud is shown in figure 5. This point cloud can then be used as a starting point for further processing, including usage as reference points to verify other distance measurement techniques as well as a

10 starting point for surface reconstruction.

4 Verification

4.1 Lidar comparison

Cloud top height information derived from the WALES lidar (Wirth et al., 2009) is used to verify the accuracy of the described method. While the stereo method provides P_{CS} at arbitrary positions in space, the lidar data is defined on a fixed grid ("cur-

15 tain") beneath the aircraft. To match lidar measurements to related stereo data points, we collect all stereo points which are horizontally close to a lidar measurement. This can be accomplished by defining a vertical cylinder around the lidar beam with 150 m radius. Every stereo derived point which falls into this cylinder with a time difference of less than 10 s is considered as stereo point related to the lidar measurement. As the (almost) nadir pointing lidar is observing cloud top heights only, we use the highest stereo point inside the collection cylinder. Figure 6 compares the measured distance between aircraft and clouds

20 from the WALES lidar and the stereo method. Generally, there is a good agreement between both methods. Because of the
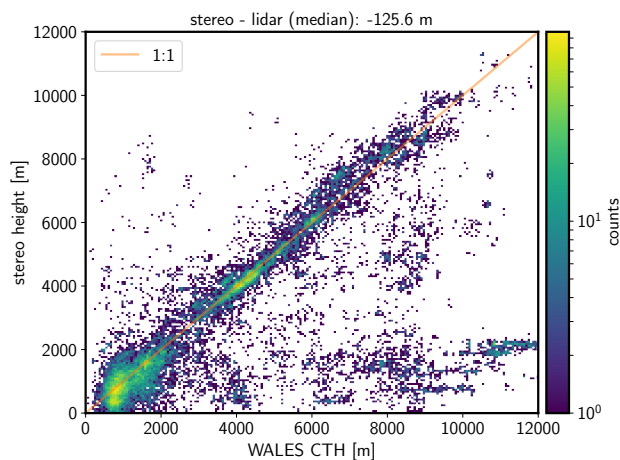
Figure 6. Comparison of cloud top height, measured with the WALES lidar and the stereo method. The most prominent outliers, present in the region of high lidar CTH and low stereo height can be attributed to thin, mostly transparent cirrus layers and cumulus clouds below. While the lidar detects the ice clouds, the stereo method retrieves the height of the cumulus layer below.

Figure 7. Thin cirrus layer above cumulus cloud field during NAWDEX RF10 (2016-10-13 10:32:10 UTC). Scenes like this are the typical cause for large outliers in the lidar comparison. The lidar detects the signal reflected by the cirrus layer while the stereo method uses the contrast of the lower cumulus clouds.

way data is collected for this comparison, the automated comparison over many whole flights still suffers from the possibility that lidar and stereo method derived heights are related to different cloud layers in a multi-layer situation. Due to its high sensitivity, the lidar method strongly favors the top height of the highest cloud while the stereo method is able to look at lower cloud layers along a slanted path. Thus, we expect that false correspondences typically relate larger lidar heights to lower stereo

5   heights. The most prominent instance of this effect is the region of WALES CTH > 6000 m and low stereo height in fig. 6. This region can be attributed almost exclusively to parts of the flight in which high thin cirrus clouds are found above a low layer of cumulus clouds, like the scene shown in figure 7. Additionally, as the stereo method relies on contrasts in the images, it works best at cloud edges which are often lower than cloud tops. The median bias is approximately 126 m for all compared flights with lower stereo heights as expected.

10   ## 4.2   Across track stability

Errors in the sensor calibration could lead to systematic errors in the retrieved cloud height with respect to lateral horizontal distance relative to the aircraft (perpendicular to flight track). In order to assess these errors, data from a stratiform cloud deck observed between 09:01:25 and 09:09:00 UTC during NAWDEX flight RF11 on 2016 Oct 14 has been sorted by the average across track pixel position. While the cloud deck features a lot of small scale variation, it is expected to be almost horizontal on

15   average. Figure 8 shows a 2D histogram of all collected data points. From visual inspection of the histogram, apart from about 50 px at the sensors borders, no significant trend can be observed. To further investigate the errors, a 2^{nd}-order polynomial has been fitted to the retrieved heights. This polynomial is chosen to cover the most likely effect of sensor mis-alignment which should contribute to a linear term and distortions in the optical path which should contribute to a quadratic term. The difference between left and right side of the sensor of 21 m corresponds to less than $0.1°$ of absolute camera misalignment and

20   the curvature of the fit is also small compared to the overall dimensions of the observed clouds.

4.3   Wind data comparison

An important criteria that we use to identify reliable tracking points is based on the assumption that the observed movement of the points can be explained by a smooth transport due to a background wind field. The thresholds for this test are very tolerant, so the requirements for the accuracy of the retrieved wind field are rather low. However, a clear positive correlation between

25   the observed point motion and the actual background wind would underpin this assumption substantially. In order to do so, we compare the stereo wind against an ECMWF reanalysis in a layer in which many stereo points have been found.

In the following, the displacement vectors of every track have been binned in 1 min bins along the flight track and 200 m bin in the vertical. To reduce the amount of outliers, bins with less than 100 entries have been discarded. Inside the bins, the upper and lower 20% of the wind vectors (counted by magnitude) have been dropped. All remaining data has been averaged to

30   produce one mean wind vector per bin. In figure 9 the horizontal component of these vectors is compared to ECMWF reanalysis data at about 2000 m above ground with horizontal sampling of $0.5°$. The comparison shows overall good agreement according to our goal to consider a quantity in the stereo matching process which roughly behaves like the wind. The general features of wind direction and magnitude are captured. Deviations may originate from multiple sources including the time difference
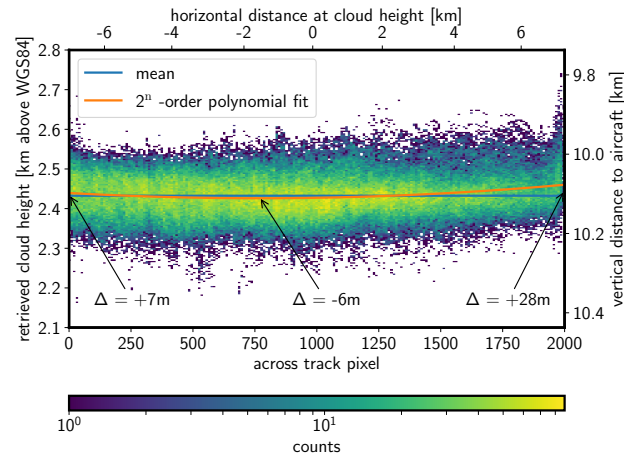
Figure 8. In the time from 09:01:25 to 09:09:00 UTC during NAWDEX flight RF11 on 2016 Oct 14, a stratiform cloud deck has been observed. The parabolic fit shows that a small systematic variation can be found beneath the noise (which is due to small scale cloud height variations). Compared to the overall dimensions of the observed cloud (≈ 14 km) and the uncertainty of the method, these variations are small. It may still be noted, that data from the edges of the sensor (≈ 50 px on each side) should be taken with care.
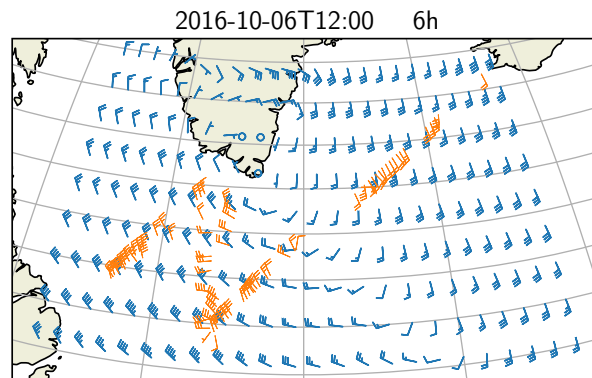
Figure 9. Horizontal wind at about 2000 m above ground. Comparison between ECMWF reanalysis (blue) and stereo derived wind (orange). Comparing grid points with co-located stereo data, the mean horizontal wind magnitude is 15.1 m/s in ECMWF and 13.4 m/s in the stereo dataset. This amounts to a difference of 1.7 ± 4.5 m/s in magnitude and $6.0 \pm 33°$ in direction. The shown deviations are standard deviations over all grid points with co-located data.

between reanalysis and measurement, representativity errors and uncertainties of the measurement principle. These results corroborate the assumption that the observed point motion is related to the background wind and filtering criteria based on this assumption can be applied.

5 Conclusions

The 3D cloud geometry reconstruction method described in this work is able to produce an accurate set of reference points on the observed surface of clouds. This has been verified by comparison to active remote sensing in the nadir perspective. Using data from the observation of a stratiform cloud field, we could verify that no significant systematic errors are introduced by
5   looking in off-nadir directions. Even for sunglint conditions cloud top heights can be derived: Because of the wide field-of-view of the sensor there are always viewing directions not affected by the sunglint.

As a visible contrast suited for point matching is a central requirment of the method, it is able to provide position information at many but not every point of a cloud. Especially flat cloud tops can show very little contrast and are hard to analyse using our method. In future, we will integrate other position datasets like the distance measurement technique using O_2A-absorption as
10   described by Zinner et al. (2018) which is expected to work best in these situations. In combining multiple datasets, the low bias and angular variability of the stereo method can even help to improve uncertainties of other methods.

While the wind information derived as part of the stereo method constituted a byproduct of this work the results look promising. After some further investigations about its quality and possibly additional filtering, this product might further add valuable information to the campaign dataset.

15   During the development of this method, it became clear that a precise camera calibration (relative viewing angles in the order of $0.01°$) is crucial to this method. A permanent time-synchronization between the aircraft position sensors and the cameras, accurate at the order of tens of milliseconds is indispensable as well. It should be noted that this does involve time stamping each individual image to cope with inter-frame jitter as well as disabling any image stabilization inside the camera. For upcoming measurement campaigns, improvements may be achieved by optimizing the automatic exposure of the camera
20   for bright cloud surfaces in stead of relying on the built-in exposure system. Furthermore, it would be useful to re-investigate the proposed method with a camera system operating in the near-infrared which would most likely profit from higher image contrasts due to lower Rayleigh scattering in this spectral region.

Data availability. specMACS data is available at https://macsserver.physik.uni-muenchen.de after requesting a personal account.

Appendix A: Geometric Camera Calibration

25   As the distance between aircraft and observed clouds is typically much larger than the flight distance between two images, the 3D reconstruction method relies on precise measurements of camera viewing angles. To allow analysis as presented in this paper, frame rates of about 1 Hz are required. At that frame rate, a change in distance between cloud and aircraft of 100 m at a distance of 10 km results in approximately $0.01°$ difference of the relative viewing angle or about $1/3$ px. Consequently, achieving accuracies in the order of 100 m or below requires both, to average over many measurements in order to get sub-pixel
30   accuracy and to remove any systematic error in the geometric calibration to less than $1/3$ px. This is only possible if distortions in the cameras optical path can be understood and corrected.

We use methods provided by the OpenCV library (Bradski, 2000) to perform the geometric camera calibration, our notation is chosen accordingly. Geometric camera calibration is done by defining a parameterized model which describes how points in world coordinates are projected onto the image plane including all distortions along the optical path. Generally, such a model includes extrinsic parameters which describe the location and rotation of the camera in world space and intrinsic parameters

5   which describe processes inside of the camera's optical path. Extrinsic parameters can differ between each captured image while intrinsic parameters are constant as long the optical path of the camera is not modified. After evaluation of various options for the camera model, we decided to use the following:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = R \begin{pmatrix} X \\ Y \\ Z \end{pmatrix} + \boldsymbol{t} \tag{A1}$$

$$x' = x/z \tag{A2}$$

10   $$y' = y/z \tag{A3}$$

Where X, Y, Z are the world coordinates of the observed object, R and $\boldsymbol{t}$ the rotation and translation from world coordinates in camera centric coordinates and x, y, z are the object location in camera coordinates. x' and y' are the projection of the object points onto a plane at unit distance in front of the camera. The distortion induced by the lenses and the window in front of the camera is accounted for by adjusting x' and y' to x'' and y'':

15   $$r^2 = x'^2 + y'^2 \tag{A4}$$

$$x'' = x'(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + s_1 r^2 + s_2 r^4 \tag{A5}$$

$$y'' = y'(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + s_3 r^2 + s_4 r^4 \tag{A6}$$

Here k_1 to k_3 describe radial lens distortion and s_1 to s_4 add a small directed component according to the thin prism model. During evaluation of other options provided by OpenCV, no significant improvement of the calibration result was found using

20   more parameters. Finally, the pixel coordinates can be calculated by a linear transformation (which is often called "camera matrix"):

$$u = f_x x'' + c_x \tag{A7}$$

$$v = f_y y'' + c_y \tag{A8}$$

Here, f_x and f_y describe the focal lengths and c_x and c_y describe the principal point of the optical system. In this model, there

25   are 6 extrinsic parameters (rotation matrix R and displacement vector $\boldsymbol{t}$) and 11 intrinsic parameters ($k_1...k_3, s_1...s_4, f_x, f_y, c_x, c_y$).

We use the well-known chessboard calibration method to calibrate this model which is based on Zhang (2000). The basic idea is to relate a known arrangement of points in 3D world space to their corresponding locations on the 2D image plane by a model as described above and solve for the parameters by fitting it to a set of sample images. The internal corners of a rectangular chessboard provide a good set of such points as they are defined at intersections of easily and automatically

30   recognizable straight lines. Furthermore, the intersection of two lines can be determined to sub-pixel accuracy which improves
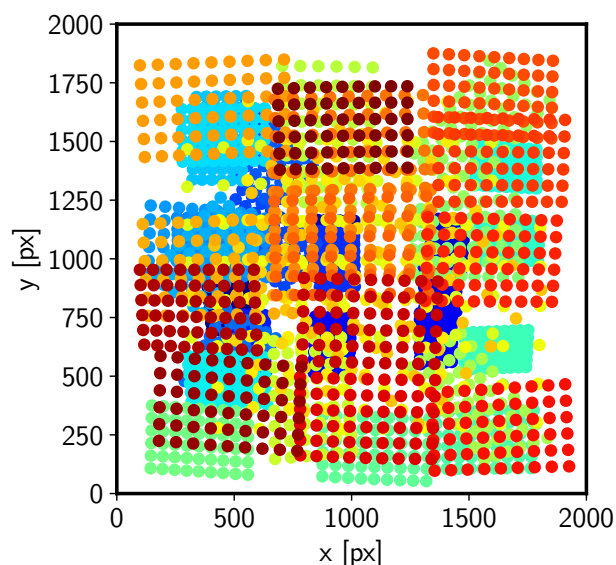
Figure A1. Locations of each chessboard corner during calibration on the image sensor plane.

the calibration performance substantially. While the extrinsic parameters have to be fitted independently for every image, the intrinsic parameters must be the same for each image and can be determined reliably if enough sample images covering the whole sensor are considered.

To evaluate the success of the calibration, the reprojection error can be used as a first quality measure. The reprojection error is defined as the difference of the calculated pixel position using the calibrated projection model and the measured pixel position on the image sensor. To calculate the pixel position, the extrinsic parameters have to be known, thus the images which have been used for calibration are used to calculate the reprojection error as well. This makes this test susceptible to falsely return good results due to an over fitted model. We use many more images (of which each provide multiple constraints to the fit) than parameters to counter this issue and have validated the stereo method (which includes the calibration) against other sensors to ensure that the calibration is indeed of good quality. Nonetheless, a high reprojection error would indicate a problem in the chosen camera model.

We have taken 62 images of a 9 by 6 squares chessboard pattern with 65 mm by 65 mm square size on an aluminium composite panel using the system assembled in aircraft configuration. The images have been taken such that the chessboard corners are spread over the whole sensor area, Figure A1 shows the pixel locations of all captured chessboard corners. After previous experiments with calibration targets made from paper and cardboard, it became clear that small ripples which inevitably appeared on the cardboard targets render the reprojection errors unusably high. Using the rigid aluminium composite material for the calibration target let the average reprojection error drop by an order of magnitude to a very low value of approximately 0.15 pixels, which should be enough to reduce systematic errors across the camera to less than 50 m. The per pixel reprojection error is shown in Figure A2 for every chessboard corner captured.
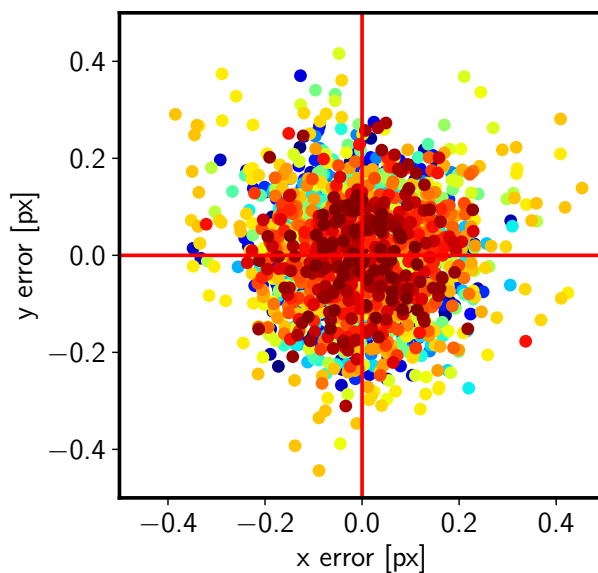
Figure A2. Reprojection error: Difference between the actual position of the chessboard corners and the calculated positions of the chessboard corners after applying the camera calibration. The average reprojection error is about $0.15\,\mathrm{pixels}$. Note that the actual chessboard corner locations can be found far into sub-pixel accuracy by following the lines along the edges of the squares.

Appendix B: OpenCV usage

The image processing of this work has been done with help of the OpenCV library (Bradski, 2000). The most important functions used are:

- `calibrateCamera` to calculate the camera calibration coefficients from a set of chessboard images

5   - `goodFeaturesToTrack` to find pixels which are most likely good candidates to be identified in the following image

- `calcOpticalFlowPyrLK` to find the corresponding pixel in the following image

Author contributions. T. Kölling was in charge of the presented measurements, developed the stereo reconstruction method and prepared the manuscript. T. Zinner and B. Mayer prepared the field campaigns, provided valuable input during the development of the method and contributed to the final version of the manuscript.

10   *Competing interests.* The authors declare that they have no conflict of interest.

Acknowledgements. This work was supported by the DFG (Deutsche Forschungsgemeinschaft, German Research Foundation) through the Project 264269520 "Neue Sichtweisen auf die Aerosol-Wolken-Strahlungs-Wechselwirkung mittels polarimetrischer und hyper-spektraler Messungen". This work was supported by the Max Planck Society, and the DFG (Deutsche Forschungsgemeinschaft, German Research Foundation) through the HALO Priority Program SPP 1294 "Atmospheric and Earth System Research with the Research Aircraft HALO

5 (High Altitude and Long Range Research Aircraft)". The authors thank Hans Grob for his help during the measurements.

References

Beekmans, C., Schneider, J., Läbe, T., Lennefer, M., Stachniss, C., and Simmer, C.: Cloud photogrammetry with dense stereo for fisheye cameras, Atmospheric Chemistry and Physics, 16, 14 231–14 248, https://doi.org/10.5194/acp-16-14231-2016, https://www.atmos-chem-phys.net/16/14231/2016/, 2016.

Bouguet, J.-y.: Pyramidal implementation of the Lucas Kanade feature tracker, Intel Corporation, Microprocessor Research Labs, 2000.

Bradski, G.: The OpenCV Library, Dr. Dobb's Journal of Software Tools, 2000.

Cox, C. and Munk, W.: Measurement of the Roughness of the Sea Surface from Photographs of the Sun's Glitter, J. Opt. Soc. Am., 44, 838–850, https://doi.org/10.1364/JOSA.44.000838, http://www.osapublishing.org/abstract.cfm?URI=josa-44-11-838, 1954.

Crispel, P. and Roberts, G.: All-sky photogrammetry techniques to georeference a cloud field, Atmospheric Measurement Techniques, 11, 593–609, https://doi.org/10.5194/amt-11-593-2018, https://www.atmos-meas-tech.net/11/593/2018/, 2018.

Emde, C., Buras-Schnell, R., Kylling, A., Mayer, B., Gasteiger, J., Hamann, U., Kylling, J., Richter, B., Pause, C., Dowling, T., and Bugliaro, L.: The libRadtran software package for radiative transfer calculations (version 2.0.1), Geoscientific Model Development, 9, 1647–1672, https://doi.org/10.5194/gmd-9-1647-2016, https://www.geosci-model-dev.net/9/1647/2016/, 2016.

Ewald, F.: Retrieval of Vertical Profiles of Cloud Droplet Effective Radius using solar reflectance from cloud sides, Ph.D. thesis, LMU München, 2016.

Ewald, F., Winkler, C., and Zinner, T.: Reconstruction of cloud geometry using a scanning cloud radar, Atmospheric Measurement Techniques, 8, 2491–2508, https://doi.org/10.5194/amt-8-2491-2015, https://www.atmos-meas-tech.net/8/2491/2015/, 2015.

Ewald, F., Kölling, T., Baumgartner, A., Zinner, T., and Mayer, B.: Design and characterization of specMACS, a multipurpose hyperspectral cloud and sky imager, Atmospheric Measurement Techniques, 9, 2015–2042, https://doi.org/10.5194/amt-9-2015-2016, http://www.atmos-meas-tech.net/9/2015/2016/, 2016.

Ewald, F., Zinner, T., Kölling, T., and Mayer, B.: Remote Sensing of Cloud Droplet Radius Profiles using solar reflectance from cloud sides. Part I: Retrieval development and characterization, Atmospheric Measurement Techniques Discussions, 2018, 1–35, https://doi.org/10.5194/amt-2018-234, https://www.atmos-meas-tech-discuss.net/amt-2018-234/, 2018.

Fielding, M. D., Chiu, J. C., Hogan, R. J., and Feingold, G.: A novel ensemble method for retrieving properties of warm cloud in 3-D using ground-based scanning radar and zenith radiances, Journal of Geophysical Research: Atmospheres, 119, 10–912, https://doi.org/10.1002/2014JD021742, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JD021742, 2014.

Fischer, J., Cordes, W., Schmitz-Peiffer, A., Renger, W., and Mörl, P.: Detection of Cloud-Top Height from Backscattered Radiances within the Oxygen A Band. Part 2: Measurements, Journal of Applied Meteorology, 30, 1260–1267, https://doi.org/10.1175/1520-0450(1991)030<1260:DOCTHF>2.0.CO;2, https://doi.org/10.1175/1520-0450(1991)030<1260:DOCTHF>2.0.CO;2, 1991.

Hock, T. F. and Franklin, J. L.: The NCAR GPS Dropwindsonde, Bulletin of the American Meteorological Society, 80, 407–420, https://doi.org/10.1175/1520-0477(1999)080<0407:TNGD>2.0.CO;2, https://doi.org/10.1175/1520-0477(1999)080<0407:TNGD>2.0.CO;2, 1999.

Krautstrunk, M. and Giez, A.: The Transition From FALCON to HALO Era Airborne Atmospheric Research: Background – Methods – Trends, pp. 609–624, Springer Berlin Heidelberg, Berlin, Heidelberg, https://doi.org/10.1007/978-3-642-30183-4_37, https://doi.org/10.1007/978-3-642-30183-4_37, 2012.

Lucas, B. D. and Kanade, T.: An Iterative Image Registration Technique with an Application to Stereo Vision, in: Proceedings of the 7th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'81, pp. 674–679, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1981.

Mayer, B. and Kylling, A.: Technical note: The libRadtran software package for radiative transfer calculations - description and examples of use, Atmospheric Chemistry and Physics, 5, 1855–1877, https://doi.org/10.5194/acp-5-1855-2005, https://www.atmos-chem-phys.net/5/1855/2005/, 2005.

Mech, M., Orlandi, E., Crewell, S., Ament, F., Hirsch, L., Hagen, M., Peters, G., and Stevens, B.: HAMP – the microwave package on the High Altitude and LOng range research aircraft (HALO), Atmospheric Measurement Techniques, 7, 4539–4553, https://doi.org/10.5194/amt-7-4539-2014, https://www.atmos-meas-tech.net/7/4539/2014/, 2014.

Moroney, C., Davies, R., and Muller, J. P.: Operational retrieval of cloud-top heights using MISR data, IEEE Transactions on Geoscience and Remote Sensing, 40, 1532–1540, https://doi.org/10.1109/TGRS.2002.801150, 2002.

Romps, D. M. and Öktem, R.: Observing clouds in 4D with multi-view stereo photogrammetry, Bulletin of the American Meteorological Society, 0, https://doi.org/10.1175/BAMS-D-18-0029.1, https://doi.org/10.1175/BAMS-D-18-0029.1, 2018.

Schäfler, A., Craig, G., Wernli, H., Arbogast, P., Doyle, J. D., McTaggart-Cowan, R., Methven, J., Rivière, G., Ament, F., Boettcher, M., Bramberger, M., Cazenave, Q., Cotton, R., Crewell, S., Delanoë, J., Dörnbrack, A., Ehrlich, A., Ewald, F., Fix, A., Grams, C. M., Gray, S. L., Grob, H., Groß, S., Hagen, M., Harvey, B., Hirsch, L., Jacob, M., Kölling, T., Konow, H., Lemmerz, C., Lux, O., Magnusson, L., Mayer, B., Mech, M., Moore, R., Pelon, J., Quinting, J., Rahm, S., Rapp, M., Rautenhaus, M., Reitebuch, O., Reynolds, C. A., Sodemann, H., Spengler, T., Vaughan, G., Wendisch, M., Wirth, M., Witschas, B., Wolf, K., and Zinner, T.: The North Atlantic Waveguide and Downstream Impact Experiment, Bulletin of the American Meteorological Society, 99, 1607–1637, https://doi.org/10.1175/BAMS-D-17-0003.1, https://doi.org/10.1175/BAMS-D-17-0003.1, 2018.

Seiz, G., Davies, R., and Grün, A.: Stereo cloud-top height retrieval with ASTER and MISR, International Journal of Remote Sensing, 27, 1839–1853, https://doi.org/10.1080/01431160500380703, https://doi.org/10.1080/01431160500380703, 2006.

Shi, J. and Tomasi, C.: Good features to track, in: 1994 Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, pp. 593–600, https://doi.org/10.1109/CVPR.1994.323794, 1994.

Stevens, B., Ament, F., Bony, S., Crewell, S., Gross, S., Hirsch, L., Mayer, B., Wendisch, M., Wirth, M., Bakan, S., Brück, H.-M., Ehrlich, A., Ewald, F., Farrell, D., Forde, M., Gödde, F., Grob, H., Hagen, M., Hansen, A., Jacob, M., Jäkel, E., Jansen, F., Klepp, C., Klingebiel, M., Kölling, T., Konow, H., Mech, M., Peters, G., Rapp, M., Wing, A., and Wolf, K.: A high-altitude long-range aircraft configured as a cloud observatory - the NARVAL expeditions, Bulletin of the American Meteorological Society, submitted.

Strabala, K. I., Ackerman, S. A., and Menzel, W. P.: Cloud Properties inferred from 8 - 12-μm Data, Journal of Applied Meteorology, 33, 212–229, https://doi.org/10.1175/1520-0450(1994)033<0212:CPIFD>2.0.CO;2, https://doi.org/10.1175/1520-0450(1994)033<0212:CPIFD>2.0.CO;2, 1994.

Várnai, T. and Marshak, A.: Observations of Three-Dimensional Radiative Effects that Influence MODIS Cloud Optical Thickness Retrievals, Journal of the Atmospheric Sciences, 59, 1607–1618, https://doi.org/10.1175/1520-0469(2002)059<1607:OOTDRE>2.0.CO;2, https://doi.org/10.1175/1520-0469(2002)059<1607:OOTDRE>2.0.CO;2, 2002.

Várnai, T. and Marshak, A.: A method for analyzing how various parts of clouds influence each other's brightness, Journal of Geophysical Research: Atmospheres, 108, https://doi.org/10.1029/2003JD003561, http://dx.doi.org/10.1029/2003JD003561, 4706, 2003.

Wendisch, M., Müller, D., Schell, D., and Heintzenberg, J.: An Airborne Spectral Albedometer with Active Horizontal Stabilization, Journal of Atmospheric and Oceanic Technology, 18, 1856–1866, https://doi.org/10.1175/1520-0426(2001)018<1856:AASAWA>2.0.CO;2, https://doi.org/10.1175/1520-0426(2001)018<1856:AASAWA>2.0.CO;2, 2001.

Wirth, M., Fix, A., Mahnke, P., Schwarzer, H., Schrandt, F., and Ehret, G.: The airborne multi-wavelength water vapor differential absorption lidar WALES: system design and performance, Applied Physics B, 96, 201, https://doi.org/10.1007/s00340-009-3365-7, https://doi.org/10.1007/s00340-009-3365-7, 2009.

Zhang, Z.: A Flexible New Technique for Camera Calibration, vol. 22, pp. 1330–1334, https://www.microsoft.com/en-us/research/publication/a-flexible-new-technique-for-camera-calibration/, 2000.

Zinner, T. and Mayer, B.: Remote sensing of stratocumulus clouds: Uncertainties and biases due to inhomogeneity, Journal of Geophysical Research: Atmospheres, 111, https://doi.org/10.1029/2005JD006955, https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2005JD006955, 2006.

Zinner, T., Schwarz, U., Kölling, T., Ewald, F., Jäkel, E., Mayer, B., and Wendisch, M.: Cloud geometry from oxygen-A band observations through an aircraft side window, Atmospheric Measurement Techniques Discussions, 2018, 1–20, https://doi.org/10.5194/amt-2018-220, https://www.atmos-meas-tech-discuss.net/amt-2018-220/, 2018.