

AO01 Observations of the Karthala Eruption

Candidate Number: 46162

Supervisors: Dr. E. Carboni and Dr. R.G. Grainger

Word Count: 5287

Abstract

Near-simultaneous observations made by two geostationary satellites in November 2005 are used to stereographically retrieve the cloud-top height of the volcanic eruption plume of Karthala volcano, Grande Comore. Observations were made every 30 minutes, and so the time-evolution of the plume is reconstructed. Comparisons between the speed and accuracy of four different area-matching algorithms run at three different scales¹ are made, and one is found to be unreliable. Results suggest the plume reached up to 20 km, significantly higher than existing analyses have concluded.

1 Introduction

The Karthala volcano makes up a substantial part of the island of Grande Comore, situated roughly halfway between the coasts of Madagascar and Mozambique at 11.75°S 43.38°E. On 24–25 November 2005, Karthala erupted [*Smithsonian Institution*, 2005], producing a large ash plume which travelled South-East and eventually reached the Northern tip of Madagascar before dispersing.

The eruption occurred within a region of observational overlap between the METEOSAT-5² (MET-5) and METEOSAT Second Generation (MSG) satellites, and as such, stereographic methods can be used

¹By scale it is meant the size of the subimage around each pixel in which that pixel's matching occurs.

²see Appendix for additional information about the satellites used in this investigation.

to determine the cloud-top height (CTH) [*Hasler et al.*, 1991], [*Diner et al.*, 2005]. Such information can have a variety of uses, from aviation safety [*Prata*, 2008] to volcanology [*Oppenheimer*, 1998] and climate modelling [*Robock*, 2000].

The benefit of the stereographic methods presented here is that no prior information of the cloud or surrounding atmosphere is required. Radiometric methods typically involve detailed prior knowledge of cloud reflectance and scattering phase functions to be reliable – see for example [*Minnis et al.*, 1993].

[*Prata and Kerkmann*, 2007] discuss in some detail the makeup of the eruption plume, but only general mention is made of the CTH. Their results are based primarily on the direction and speed of wind-induced motion compared to a radiosonde ascent made 1000 kilometres from Karthala near Antananarivo, Madagascar. This investigation will attempt to improve on these results, and also present greater spatial and temporal detail.

2 Method

2.1 Principle of the Method

This investigation is based on the well-known principle of stereoscopic parallax: the apparent shift in position of an object relative to a fixed background due to a change in observation angle. In this case, the object being observed is the Karthala eruption plume, and the fixed background is the Mozambique Channel and surrounding landmasses. The change in obser-

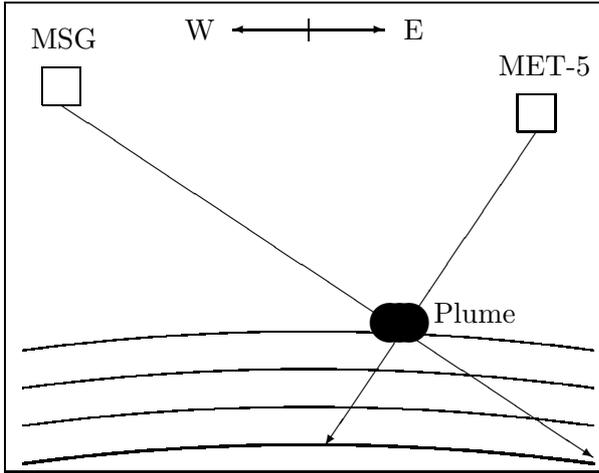


Figure 1: A *simplified and not-to-scale diagram of the relative positions of the two satellites being translated to a shift in the observed position of the plume. The thick curve represents the Earth's surface, whilst the thinner curves represent successive altitudes of reprojection.*

vation angle is due to the differing positions of the two satellites MET-5 and MSG (see Figure 1). However, due to the spherical geometry of the Earth, the different viewing angles result in distortion as well as shifting, especially because one satellite is closer to nadir than the other. As such, application of matching algorithms on the initial satellite images would be futile, and the images are first reprojected to a common longitude/latitude (lon/lat) grid. Once this has been done, matching can be performed, with any deviations from coincidence being attributed to a parallax shift, and the magnitude of this shift being directly related to CTH.

It should be noted that there is a 140 second gap between an MSG image and the corresponding MET-5 image. This will result in a correction needing to be applied to sections of cloud moving above a certain velocity.

It should also be noted that the input images are rescaled by their maxima and minima to the range 0–255. The greater slant in the ray path of the MSG images means it is looking through more of the plume, and hence there is more absorption and so a darker image. This

will be shown to have a significant effect on the matching of one of the algorithms.

2.2 Reprojection

The main difference between the method used here and that described in [Prata and Turner, 1997] is that, instead of determining CTH from a horizontal pixel shift applied to each subimage individually, the images are reprojected at different altitudes before matching the subimages directly. This is basically achieved by changing the effective radius of the Earth before the images are reprojected to a common lon/lat grid (see Figure 1). One benefit of this is that it avoids difficulties due to the spherical geometry of the problem, which would otherwise cause the same pixel shift to relate to a different CTH depending on where in the image was being studied. It also allows some small improvement in speed, as the reprojection can be applied *before* application of the matching algorithms, whereas Prata's method requires separate shifting for each algorithm.

The result is that for each matcher, every pixel has associated with it a column vector of match coefficients at different altitudes. This vector is then searched for the altitude with the optimal match, and this altitude is then assigned to the pixel.

Care must be taken to ensure that homogeneous regions, such as the sea surface around the cloud, are correctly identified as being at ground level. This is because, if the region is sufficiently homogeneous, then *all* altitudes offer an equally good match. For this reason, the search for an optimal match altitude is forced to zero in the case of multiple ideal matches. The downside of this is that highly homogeneous regions within the plume itself will necessarily be flagged as zero altitude.

2.3 Matching Algorithms

In this investigation, four different matching algorithms have been used, although the reader will later see that in practice

only three should be thought of as reliable. The four matching algorithms are: a correlation function like that described in [Prata and Turner, 1997], the Multipoint Matchers Using Means (M2) and Medians (M3) described in [Muller et al., 2002], and a simple root-mean-square (RMS) difference. Of these four, the RMS can theoretically (though not usefully) be used on individual pixels; the rest are area-based matchers and so this investigation will focus on their use at three different scales: 5×5 , 11×11 and 15×15 pixels. It will be shown that the 11×11 scale offers a good compromise between speed and clarity, but for the majority of this report the higher clarity 15×15 images will be presented.

2.3.1 The Correlation Function

In this investigation, the reprojection method on the input images results in differences in the form of the correlation function from [Prata and Turner, 1997], upon which it is based. There are not any pixel shift parameters here, because all shifting is performed prior to the application of matchers. The calculation of the correlation coefficient, however, remains functionally the same:

$$corr(x, y) = \frac{\psi_{MET5,MSG}}{\sqrt{\psi_{MET5,MET5}}\sqrt{\psi_{MSG,MSG}}}$$

where

$$\psi_{view1,view2} = \sum_{j=-a}^a \sum_{k=-a}^a r_{view1}(x+j, y+k)r_{view2}(x+j, y+k)$$

Here, r represents the rescaled brightness temperature (RBT) of a given pixel; a represents the limits of the subimage being matched around the pixel – i.e. for the 5×5 scale, a is 2. A small downside to this is that pixels within a of the image edge cannot be matched, and so outputs will have a border region in which no matches can be found.

2.3.2 The M2 and M3 Matchers

These matching algorithms are exhaustively described in [Muller et al., 2002]. The basic

idea is that, within each subimage, if the match is perfect a given pixel will differ from the mean (or median) of the subimage by the same proportion as its counterpart pixel in the other subimage. This allows the matcher to ignore any differences in scaling between the input images (see section 2.1).

The M2 match coefficient is found as follows:

$$m2(x, y) = \frac{\sum_{j=-a}^a \sum_{k=-a}^a |R_2^{j,k} - C_2^{j,k}|}{\sum_{j=-a}^a \sum_{k=-a}^a |R_2^{j,k}|}$$

where

$$R_2^{j,k} = \frac{r_{MET5}(x+j, y+k) - \langle R \rangle}{R_{\max} - R_{\min}}$$

$$C_2^{j,k} = \frac{r_{MSG}(x+j, y+k) - \langle C \rangle}{C_{\max} - C_{\min}}$$

$\langle R \rangle$ Mean of MET-5 subimage around (x, y)

R_{\max} Maximum value of MET-5 subimage

R_{\min} Minimum value of MET-5 subimage

Similarly, the M3 match coefficient is found as follows:

$$m3(x, y) = \frac{\text{median}^{j,k} |R_3^{j,k} - C_3^{j,k}|}{\text{median}^{j,k} |R_3^{j,k} - 1|}$$

where

$$R_3^{j,k} = \frac{r_{MET5}(x+j, y+k)}{\text{median}(R)}$$

$$C_3^{j,k} = \frac{r_{MSG}(x+j, y+k)}{\text{median}(C)}$$

$\text{median}(R)$ Median of MET-5 subimage around (x, y)

C corresponds equivalently to the MSG subimage throughout, and r and a have their previous meanings.

It is important to note that in a subimage with the same maximum and minimum (i.e. a perfectly homogeneous subimage) the M2 matcher goes to infinity, and the M3 matcher goes to infinity when $R_3^{j,k}$ is 1. Therefore an additional small term is needed in the denominators of $R_2^{j,k}$, $C_2^{j,k}$ and $m3(x, y)$ to ensure the program can run – here, 10^{-10} is used. The actual value is arbitrarily chosen, but as long as it is smaller than 10^{-6} it will not affect the float-precision calculations used here.

Algorithm	5 × 5	11 × 11	15 × 15	Total
Correlation	6	21	37	64
M2	8	23	39	70
M3	5	17	29	51
RMS	4	12	22	38

Table 1: *Time in minutes for the completion of each algorithm at each scale on one pair of 500×250 pixel images. Absolute values will vary with processor speed.*

2.3.3 The RMS Difference

By far the simplest of the four matchers, the idea comes from [Prata and Turner, 1997] where it is presented as a computationally inexpensive alternative to the correlation function. It has the following form here:

$$\text{RMS}(x, y) = \sqrt{\sum_{j=-a}^a \sum_{k=-a}^a (R_{RMS}^{j,k} - C_{RMS}^{j,k})^2}$$

where

$$\begin{aligned} R_{RMS}^{j,k} &= r_{\text{MET5}}(x + j, y + k) \\ C_{RMS}^{j,k} &= r_{\text{MSG}}(x + j, y + k) \end{aligned}$$

Table 1 shows that for the 500 × 250 pixel images used throughout this investigation, the RMS difference method is substantially faster than all the other algorithms. However, it will later be shown that the outputs of this method are noticeably different from those of the other three.

2.4 Observation Channel

All outputs presented in the results section of this report were obtained through stereo matching of infrared images at 10.8 μm. As shown in Figure 2, the volcanic plume is presented in high contrast in this channel. Compare this to the 0.6 μm visible channel, where the presence of bright cloud saturates the RBT range. This results in the plume being barely discernable to the eye, and the matching algorithms have a similar problem. As a result, Figure 3 shows how the lack of contrast in the

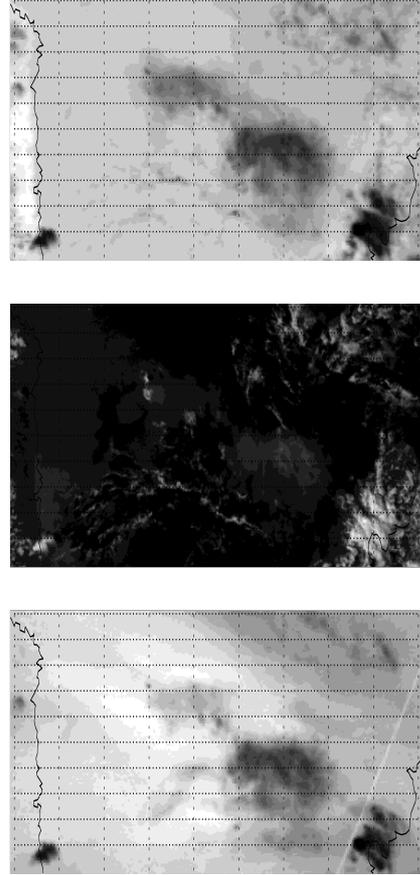


Figure 2: *Rescaled MET-5 input images at 12:00:00, reprojected at sea-level to a common lon/lat grid. Top: Infrared (10.8 μm). Middle: Visible (0.6 μm). Bottom: Water Vapour (7.3 μm).*

visible channel results in a substantial proportion of the plume being missed. MSG is capable of high resolution observation in the visible, but as the images are reprojected to the same resolution on the lon/lat grid this does not alleviate the problem.

There is a third option in the water vapour channel at 7.3 μm, and it has similar levels of contrast to the infrared, in the sense that the plume is discernable (see bottom of Figure 2). However, the background contains significant mobile spatial structure due to the varying levels of moisture in the atmosphere, and the matchers process this as having significant *vertical* structure due to non-negligible motion

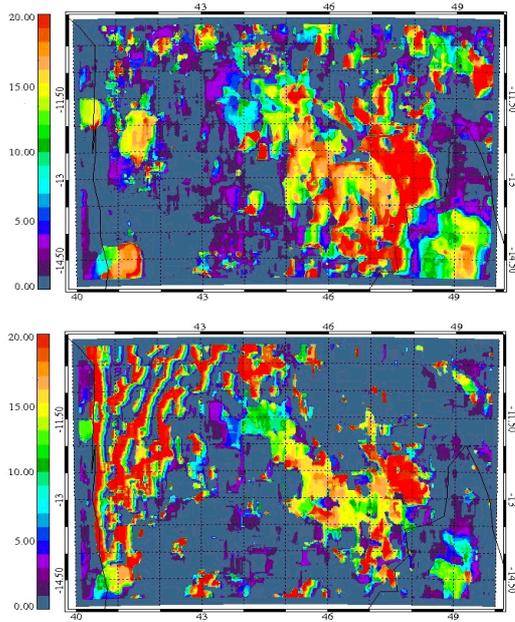


Figure 3: CTH obtained from 15×15 scale correlation function at 12:00:00. Height is in kilometres. Top: Infrared ($10.8 \mu\text{m}$). Bottom: Visible ($0.6 \mu\text{m}$).

over the time gap between MSG and MET-5 measurements. This means that, although the plume may be well matched, the output images are very difficult to interpret.

3 Results and Discussion

3.1 Intercomparison of Matchers

Figure 4 presents the four different matchers run at 3 different scales for the timeslot 12:00:00³. The first thing to note is that the 5×5 scale images are too sensitive to small-scale contrast changes, and result in an extremely noisy output.

Comparison between the 11×11 and 15×15 images will show some small differences, but for

³Outputs will be presented referring to the timecode of the MET-5 image, rather than the preceding MSG image. To work out the equivalent MSG timecode, subtract 140 seconds. All timecodes are in UTC.

the most part the outputs are consistent. The main differences are that the 11×11 images are slightly noisier and slightly more susceptible to labelling plume features as ground (see section 2.2). The 15×15 outputs are smoother, which makes more sense from a physical perspective, as although there is some stratification between the two sections of the bifurcating plume, the interior of the plume would be expected to vary smoothly rather than discontinuously. However, the 11×11 images should be sufficient in situations where speed is a factor.

Comparing between matching algorithms, the correlation function, M2 and M3 matchers all appear relatively consistent, with the leading (SE) edge being labelled as the highest (reddest) part of the plume, and the separating lower section (NW) being labelled primarily in the light-blue to yellow.

The RMS results are significantly different in several ways. Firstly, the majority of the SW part of the plume is completely missing, being labelled as ground-level. Secondly, much of the rest of the plume is labelled at 20km, the cut-off of the algorithms' range. Thirdly, those regions which are labelled at a given altitude seem to be extremely homogeneous, with the image being primarily made up of large blocks of colour, rather than the smoother shading of the other three matchers. The reason for all this is the RMS algorithm's greater susceptibility to the different brightness temperature scales of the two satellites' images.

To explain this, first consider a region near the Western edge of the plume, with the darkness increasing towards the interior. Due to the relative positions of the satellites and Karthala, the darker MSG plume is projected further to the East than the MET-5 plume at low altitudes of reprojection (see Figure 1). Therefore, the darker, more central regions of the MET-5 plume are initially overlaid on the lighter, more exterior regions of the MSG plume. The difference in brightness temperature rescaling means that this presents a *better* match than at a greater altitude, where the interiors of both plumes are overlaid. This explains the loss of much of the SW part of the plume as being at

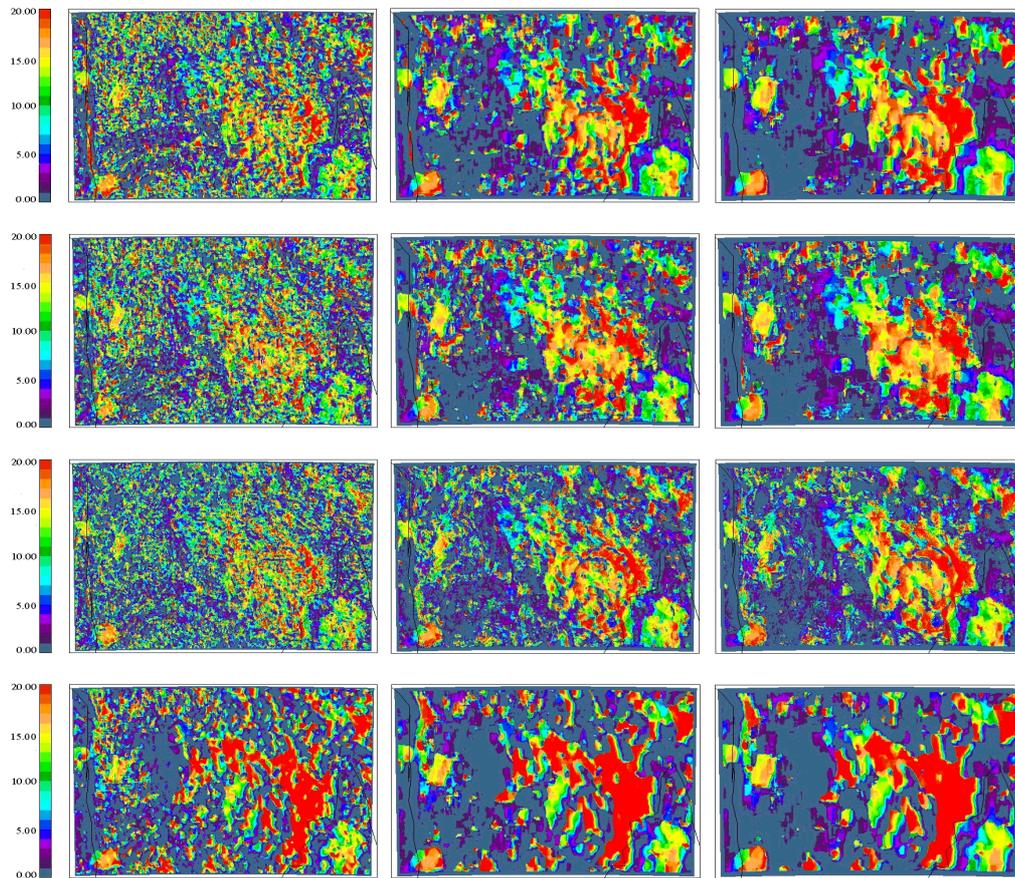


Figure 4: *CTH* obtained from different matchers and scales at 12:00:00. Height is in kilometres. Left Column: 5×5 scale. Middle Column: 11×11 scale. Right Column: 15×15 scale. Top: Correlation Function. Second Row: M2 Matcher. Third Row: M3 Matcher. Bottom: RMS Difference. Lon/lat grid has been removed for ease of viewing.

ground-level.

Now, consider instead the Eastern edge of the plume, with the darkness gradient in the opposite direction. Here, the matches get successively *better* with increasing altitude, as successively darker regions of the interior of the MET-5 image are overlaid on the brighter exterior of the MSG image. This is why the Eastern edge of the plume is presented as being at 20km, the highest altitude the algorithm can produce. The ‘blockiness’ of the image can then be explained due to regions of the images being in one of these two regimes – i.e. either nearer the Eastern edge, or nearer the Western edge, with a sharp discontinuity between the two.

In comparing the correlation function, M2 and M3 matchers, there is no independent source of verification of the plume altitude, and so none can readily be identified as the best. [Seiz *et al.*, 2006] describe the multipoint matchers as the M2 having faster run-time but lower coverage than M3. This investigation found the opposite to be true, with the M2 matcher being the slowest of all the algorithms used, but with the greatest coverage (fewest internal ground-level flags) and also the smoothest output. However, a rigorous comparison will not be presented here, and all three algorithms will be used together as an internal consistency test for later analyses.

3.2 Selective Masking

The outputs of the algorithms could be presented unmodified, but their clarity could be improved by judicious use of selective masks to filter out the noise whilst retaining the pertinent information about the plume. Within the context of the time-evolution of the images the plume is discernable, but especially with single images it might be profitable to present images with much of the non-plume distractions removed. With multi-channel data, a selective mask like that described in [Prata and Grant, 2001] could be applied, but such data are only available from MSG, not MET-5. Hence, this investigation has made use

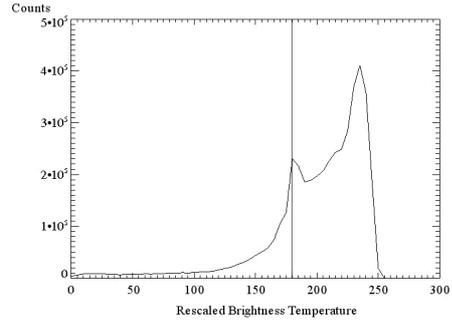


Figure 5: A histogram of RBT over all MET-5 input images. Binsize is 5.

of two different selective masks.

The first, and simpler, of the two masks is the application of an RBT filter to the input images. On the assumption that the plume is one of the darkest objects in the images, removal of pixels above a certain RBT threshold should select the plume over the background noise and some of the non-plume clouds. Analysis over all timeslots suggests that an upper threshold of 180 (non-inclusive) should be suitable to remove the main peaks of sea-surface background and cloud without risking removing plume features (see Figure 5).

In actually applying the filter, care must be taken. Since the filter is dependent upon the state of the *input* images, the filter must be applied over all altitudes equally. Therefore, the final mask will include contributions from both images in their Eastern-most and Western-most projections. In this way, none of the plume is missed, but on the other hand the mask will necessarily be too broad.

The second mask is the application of a contrast filter selectively applied at the retrieved altitude. The contrast of a pixel is here defined as the difference between the maximum and minimum RBT within the pixel’s subimage. The plume, as a dark object on a bright background, has substantial contrast across its edges. Therefore, application of a filter to remove sections of the image containing low contrast should select for the edge of the plume. Another analysis over all timeslots produces

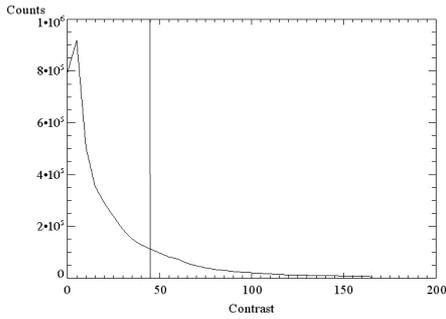


Figure 6: A histogram of contrast over all MET-5 input images. Binsize is 5.

a possible threshold of 45 (non-inclusive), but there is much greater scope for different interpretations with this mask (see Figure 6), and so the choice made here was because it visibly seems to work for the largest number of images.

This filter has the advantage that it is applied at the individual pixel level, and so the retrieved altitude of the pixel can be used to determine which altitude projection of the input images to apply the filter to. This removes the excessive coverage of the RBT mask, but the disadvantage is that the more homogeneous regions in the interior of the plume are masked out, as there is little contrast in these regions. This means that interpretation of the outputs requires a certain amount of ‘filling in the blanks’, as the interior state of the plume must be interpolated from the edges. On the plus side, there is much greater confidence that the true extent of the plume is being presented. A comparison of the two masks is presented in Figure 7.

3.3 CTH Retrieval and Plume Evolution

Data from the 25th November 2005 eruption of Karthala volcano were analysed in 30 minute intervals from 00:30:00 to 18:00:00 UTC, at which point the plume was too dispersed for continued matching. Limitations of space preclude presentation of all three reliable matchers at all three scales under both masks and over

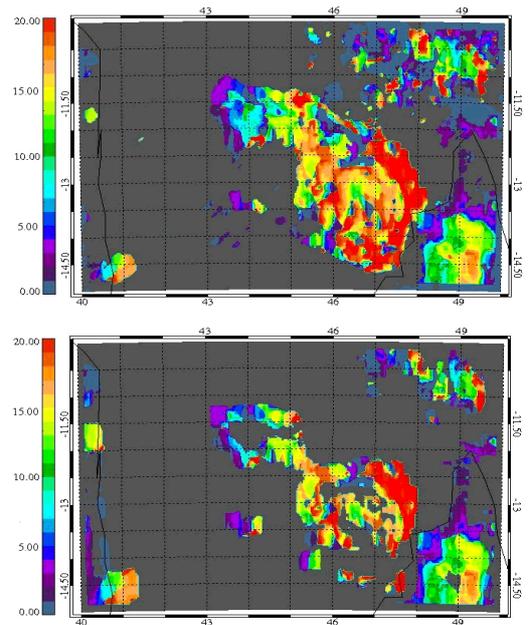


Figure 7: CTH obtained from 15×15 scale correlation function. Height is in kilometres. Top: Brightness masked. Bottom: Contrast masked. Dark grey represents masked-out pixels. Clearly shown are the excessive coverage of the RBT mask and the blank interior regions of the contrast mask.

all timeslots. Therefore, only the most complete matcher – M2 – is presented here, and only at the scale of greatest clarity – 15×15 pixels. It is also only presented under the RBT mask, as this allows for easier interpretation, even though as previously discussed it is less selective. Finally, the outputs are only presented at hourly intervals, as this should allow the reader to ascertain the time-evolution of the plume without necessitating a large number of very small images.

The time evolution of the eruption plume is shown in Figure 8. Certain features are of particular interest. Firstly, the plume bifurcation as described in [Prata and Kerkmann, 2007] can be observed at around 13:30:00. In subsequent times, the lower plume quickly disperses, with the upper plume mixing with or overlay-

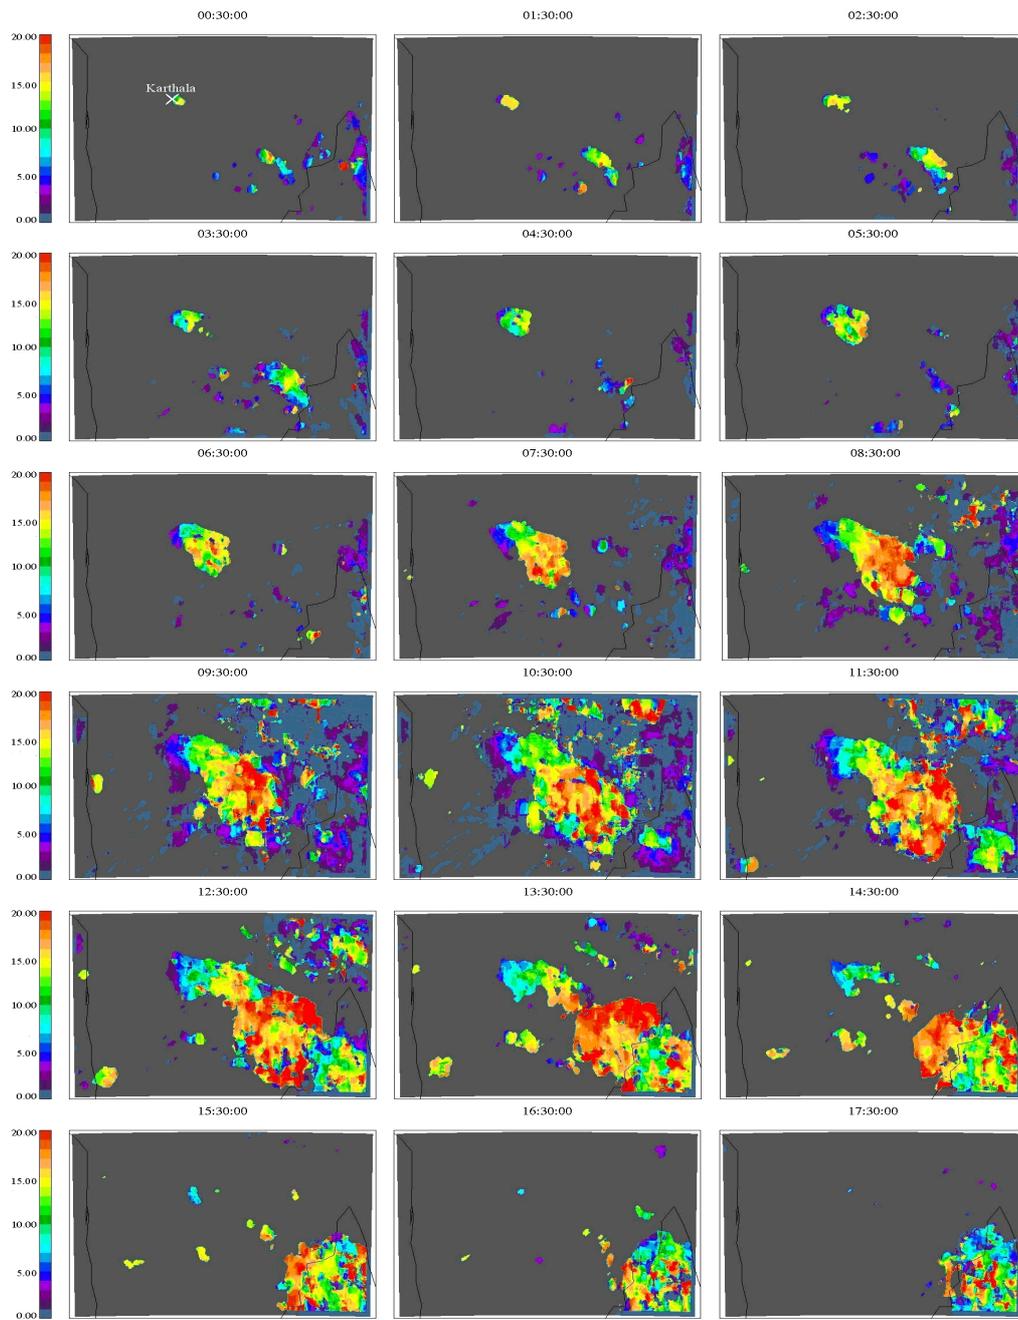


Figure 8: *CTH* obtained from 15×15 scale, *RBT*-masked *M2* matcher from 00:30:00 to 17:30:00 UTC in hourly intervals. Height is in kilometres. Lon/lat grid has been removed for ease of viewing. Dark grey represents masked-out pixels.

ing a cloud bank over Northern Madagascar until they become indistinguishable. The stratification between the plumes agrees quite well with Prata. Prata describes the lower plume as being at 8km, and the upper at 12km. This investigation found that the lower plume is mostly between 4 and 12 km, and the upper plume mostly between 14 and 20 km.

Also of interest is the early motion of the plume. The images show that for the first few hours the plume rose but did not significantly move. Then, after about 07:30:00 the higher regions of the plume reached an altitude of strong South-Eastward winds and the upper plume rapidly spread in this direction. This suggests significant stratification of the wind profile, which is contrary to the findings of Prata and so suggest the Antananarivo radiosonde data are not sufficiently near to the plume to accurately describe its motion.

3.4 Error Analysis

Detailed quantitative analyses of the results of this investigation would require independent data for verification, but as has already been discussed these data are not readily available. However, analysis of potential sources of uncertainty will give some measure of the confidence with which the results can be treated.

A source of uncertainty already mentioned here is the motion of the plume in the 140 second window between an MSG scan and the corresponding MET-5 scan. Using simple geometric considerations, it was found that increments of approximately 12 ms^{-1} would result in a 1 km height error. Tropospheric winds would not normally be expected to exceed 11.8 ms^{-1} in the tropics [Velden *et al.*, 1997], so although some aspects of the cloud may require shifting downwards by 1 km, it is not likely to be more than this.

Another potential source of uncertainty would be the breadth of the correlation function (or equivalent) around the altitude of best match. An approximate analysis (see Appendix), using inverse theory considerations from [Rodgers, 2000], came to the conclusion

that the majority of the plume is actually limited by the discretisation of the reprojections (i.e. error range of 1 km or less), although some of the non-plume surroundings have error ranges of 10 km or more.

It might be suggested that the precision limit ought to be down to the pixel size of the input images rather than the discretisation of the reprojection algorithm. Using simple geometric considerations, this would give a single pixel precision varying from 1.4 to 3.5 km from West to East (due to the slant of the ray paths). However, the area matchers described herein are actually capable of sub-pixel accuracy, because a 1 km reprojection needn't shift every pixel to alter the match coefficient, so long as at least one pixel in the subimage is shifted.

4 Conclusions

Three of the stereo matching algorithms presented here have proven to be produce consistent and physically reasonable results, with good spatial and temporal detail. Without independent verification the reliability of these results are not known, but consistency across three quite different methods is a positive sign. Future work could include usage of different and more sophisticated matchers (eg: hierarchical matching like [Rampapriyan *et al.*, 1986] or the M4 matcher from [Muller *et al.*, 2007]) to further improve this confidence, but the benefit of the matchers included here is that they are easily adapted to the geometry of the problem and are not computationally expensive.

Additional work on the effect of wind-induced motion, perhaps using methods from [Leese *et al.*, 1970] or [Horváth and Davies, 2001], or the NestedMax algorithm from [Moroney *et al.*, 2002], could be pursued to better quantify the wind correction.

Finally, independent verification could be sought from the application of the matching to geological features of known altitude. However, these features would have to be strongly contrasted in the infrared to be matchable.

Acknowledgements

The author would like to thank Dr. Carboni and Dr. Grainger for their assistance and supervision throughout this project. Thanks also go to the other members of the Earth Observation Data Group, Atmospheric, Oceanic and Planetary Physics, The University of Oxford for their ideas and advice. The author specifically would like to thank M. Clerici for the use of his altitude reprojection algorithm, which was a cornerstone of this investigation.

References

- [Diner et al., 2005] Diner, D.J., Braswell, B.H., Davies, R., Gobron, N., Hu, J., Jin, Y., Kahn, R.A., Knyazikhin, Y., Loeb, N., Muller, J.-P., Nolin, A.W., Pinty, B., Schaaf, C.B., Seiz, G., and Stroeve, J., (2005), The value of multiangle measurements for retrieving structurally and radiatively consistent properties of clouds, aerosols and surfaces, *Remote Sens. Environ.*, *97*, 495–518
- [EUMETSAT] European Organisation for the Exploitation of Meteorological Satellites, www.eumetsat.int
- [Hasler et al., 1991] Hasler, A.F., Strong, J., Woodward, R.H., and Pierce, H. (1991), Automatic Analysis of Stereoscopic Satellite Image Pairs for Determination of Cloud-Top Height and Structure, *J. Appl. Meteorol.* *30*:257–281
- [Horváth and Davies, 2001] Horváth, Á., and Davies, R., (2001), Feasibility and Error Analysis of Cloud Motion Wind Extraction from Near-Simultaneous Multiangle MISR Measurements, *J. Atmos. Oceanic Technol.* *18*, 591–608
- [Leese et al., 1970] Leese, J.A., Novak, C.S., and Clarke, B.B., (1970), An Automated Technique for Obtaining Cloud Motion from Geosynchronous Satellite Data using Cross Correlation, *Journal of Applied Meteorology*, *10*, 118–132
- [Minnis et al., 1993] Minnis, P., Liou, K.-N., and Takano, Y., (1993), Inference of Cirrus Cloud Properties Using Satellite-observed Visible and Infrared Radiances Pts 1 & 2, *J. Atmos. Sci.* *50*, 1279–1322
- [Moroney et al., 2002] Moroney, C., Davies, R., and Muller, J.-P., (2002), Operational Retrieval of Cloud-Top Heights Using MISR Data, *IEEE Transactions on Geoscience and Remote Sensing*, *40*, No.7, 1532–1540, doi:10.1109/TGRS.2002.801150
- [Muller et al., 2002] Muller, J.-P., Mandanayake, A., Moroney, C., Davies, R., Diner, D.J., and Paradise, S. (2002), MISR Stereoscopic Image Matchers: Technique and Results, *IEEE Transactions on Geoscience and Remote Sensing*, *40*, pp. 1547–1559
- [Muller et al., 2007] Muller, J.-P., Denis, M.-A., Dundas, R.D., Mitchell, K.L., Naud, C., and Mannstein, H. (2007) Stereo cloud-top heights and cloud fraction retrieval from ATSR-2, *International Journal of Remote Sensing*, *28*, No.9, 1921–1938
- [Oppenheimer, 1998] Oppenheimer, C., (1998), Volcanological applications of meteorological satellites, *Int. J. Remote Sensing*, *19*, No.15, 2829–2864
- [Prata, 2008] Prata, A.J., (2008), Satellite detection of hazardous volcanic clouds and the risk to global air traffic, *Nat. Hazards*, Springer Science and Business Media, doi:10.1007/s11069-008-9273-z
- [Prata and Grant, 2001] Prata, A.J., and Grant, I.F., (2001), Retrieval of microphysical and morphological properties of volcanic ash plumes from satellite data: Application to Mt. Ruapehu, New Zealand., *Q.J.R. Meteorol. Soc.*, *127*(576B), 2153–2179
- [Prata and Kerkmann, 2007] Prata, A.J., and Kerkmann J. (2007), Simultaneous retrieval of volcanic ash and SO₂ using MSG-SEVIRI measurements, *Geophys. Res. Lett.*, *34*, L05813, doi:10.1029/2006GL028691
- [Prata and Turner, 1997] Prata, A.J., and Turner J.P. (1997), Cloud-Top Height Determination Using ATSR Data, *Rem. Sens. Env.*, *59*(1), 1–13.
- [Rampapriyan et al., 1986] Rampapriyan, H.K., Strong, J.P., Hung, Y., and Murray, C.W.Jr., (1986), Automatic matching of SIR-B Images for elevation mapping, *IEEE Trans. Geosci. Remote Sens.* GE-24(4):462–472
- [Robock, 2000] Robock, A., (2000), Volcanic Eruptions and Climate, *Rev. Geophys.* *38*, 191–218
- [Rodgers, 2000] Rodgers, C.D., (2000), Inverse Methods for Atmospheric Sounding: Theory and Practice, World Scientific
- [Seiz et al., 2006] Seiz, G., Davies, R., and Grün, A., (2006), Stereo cloud-top height retrieval with ASTER and MISR, *Int. J. Remote Sens.* *27*(9):1839-1853
- [Smithsonian Institution, 2005] Smithsonian Institution, (2005), Karthala, *Bull. Glob. Volcanism Netw.* *30*(11), 2–4
- [Velden et al., 1997] Velden, C.S., Hayden, C.M., Nieman, S.J., Menzel, W.P., Wanzong, S., and Goerss, J.S., (1997), Upper-tropospheric winds derived from geostationary satellite water vapor observations, *Bull. Amer. Meteor. Soc.*, *78*, 173–195

Appendix

Additional Satellite Information

METEOSAT-5 – A geostationary satellite (now decommissioned) with nadir above 63°E and nominal orbital height 35,800 km. Carried a radiometer capable of observing in the infrared at 10.8 μm , in the visible at 0.6 μm and in the water-vapour channel at 7.3 μm . Made measurements at 00:30:00 and every 30 minutes subsequently.

METEOSAT Second Generation, a.k.a. METEOSAT-8 – A geostationary satellite with nadir above 0°W and nominal orbital height 42,000 km. Carries the SEVIRI instrument. Makes measurements at 00:12:40 and every 15 minutes subsequently.

SEVIRI – Spin-Enhanced Visible and Infrared Imager - The radiometer onboard MSG, capable of observations in 12 channels including 10.8 μm , 0.6 μm and 7.3 μm . Also capable of High Resolution observation in the visible channel, with twice the resolution of the other channels and MET-5 (see section 2.4 for reasoning why this was not used). Above data from [EUMETSAT].

Inverse Theory Error Analysis

Let \mathbf{y} be the measured variable, and \mathbf{x} the variable to be retrieved. Then:

$$\mathbf{y} = \mathbf{F}(\mathbf{x}) + \varepsilon$$

Where \mathbf{F} is the forward model which links the two variables, and ε is the noise. In this investigation, \mathbf{y} is the correlation coefficient (or equivalent) and \mathbf{x} is the CTH. Then if $\mathbf{S}_\mathbf{x}$ and $\mathbf{S}_\mathbf{y}$ are the covariances on these variables, inverse theory [Rodgers, 2000] in the absence of *a priori* information concludes:

$$\mathbf{S}_\mathbf{x} = \left(\mathbf{K}^T \mathbf{S}_\mathbf{y}^{-1} \mathbf{K} \right)^{-1} \quad \left[\mathbf{K} = \frac{\partial \mathbf{F}(\mathbf{x})}{\partial \mathbf{x}} \right]$$

Then if σ_x and σ_y are the standard errors on these variables, in this simple 1-D case it collapses down to;

$$\sigma_x = \left(\frac{\partial F(x)}{\partial x} \cdot \frac{1}{\sigma_y} \right)^{-1}$$

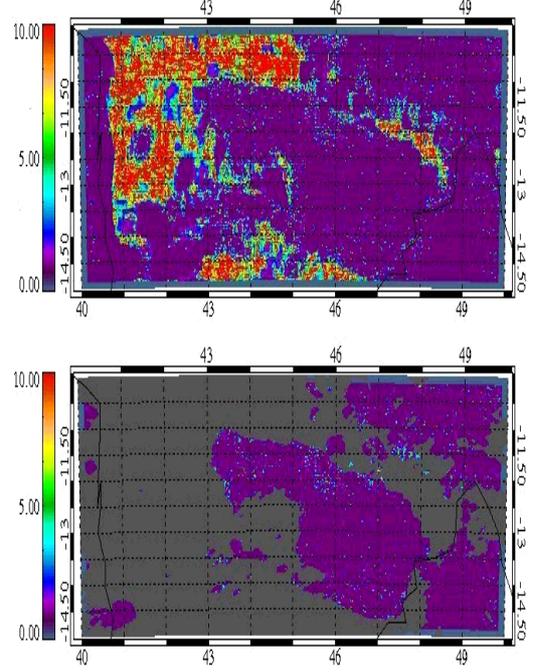


Figure 9: *Standard altitude error in kilometres from inverse theory considerations on the 15×15 scale correlation function at 12:00:00. Values below 1 km are set to 1 km (purple), and values above 10 km are set to 10 km (red). Top: Unmasked. Bottom: Brightness masked.*

To calculate an approximate error for y , an input image was seeded with random noise distributed between ± 1 RBT units and then matched to the same, unseeded image. The mean of the correlation coefficients was then subtracted from 1 (perfect correlation) to give the standard error σ_y .

To calculate the gradient of the forward model, the difference between the correlation coefficients of nearest neighbours around the altitude of best match were used in the standard linear interpolation:

$$\frac{\partial F(x)}{\partial x} \approx \frac{\Delta F(x)}{\Delta x}$$

Some care is necessary in this method when the altitude of best match is zero or 20 km, as only

one nearest neighbour can be used (there is no -1^{th} or 21^{st} coefficient). Care must also be taken when the gradient is zero, as this sends σ_x to infinity – here, such values are forced to 10 km, but in examining the results this does not occur within the plume, only on the sea-surface.

Figure 9 shows the standard altitude error calculated from this method. Note that errors below 1 km (the reprojection discretisation) are set to 1 km, and errors above 10 km are set to 10 km. The figure shows that a small number of pixels within the interior of the plume have error ranges higher than 1 km, as might be expected since the interior is harder to match. However, a general error of 1 km is suitable for most of the plume.