

Vertical profiling of precipitation using passive microwave observations: the main impediment and a proposed solution

Ziad S. Haddad and Kyung-Won Park
Jet Propulsion Laboratory, California Institute of Technology

Abstract

Several methods have been proposed to train microwave radiometers to retrieve precipitation rates estimated by a radar which observed the same location at the same time. These radar-trained passive-microwave algorithms differ in the quantities that are estimated: some estimate the vertically-integrated liquid water, while others estimate the near-surface precipitation. Since it is no more or less credible to estimate the rain rate at the surface than it is to estimate the rain rate at any discrete altitude, it is particularly interesting to quantify to what extent it is indeed feasible to estimate vertical profiles of precipitation from a passive microwave radiometer, what the obstacles are, and what vertical resolution would be achievable. To that end, we selected five study regions, and started by quantifying the vertical variability of rainfall as derived from the Tropical Rainfall Measuring Mission (TRMM) radar. Two cases emerged: a monsoon-like case where the first principal component of the vertical precipitation accounts for about 90% of the variability, and a Mediterranean-like case where the first principal component accounts for about 80% of the variability. A Bayesian approach was applied to the TRMM Microwave Imager measurements co-located with the radar profiles. For the monsoon-like regions, it produced estimates of rain rates at 250 meter vertical increments which compared well with the TRMM radar estimates. For the Mediterranean-like regions, the retrieval errors were very large. We therefore proceeded to identify the main reason for the failure of the straightforward training method. It turns out to be the unknown signature of the sea surface in the portion of the beam that does not contain precipitation. In the problematic Mediterranean case, our original straightforward approach can still be applied to measurements which do not suffer from this identifiable partial-beam-filling. For measurements that do, we derive a filtering approach to neutralize the variability of the partial surface signature and thus overcome the problem.

0. Introduction

The first satellite dedicated to the measurement of precipitation, the Tropical Rainfall Measuring Mission (TRMM), carries a cross-track-scanning ranging Precipitation Radar (PR) with a 220-km swath, as well as a nine-channel passive TRMM Microwave Imager (TMI) with a 700-km swath. One of the main reasons for the inclusion of the radar and the radiometer on the same platform was to examine and refine the assumptions behind the radiometer retrievals of precipitation – for example, by using the direct three-dimensional measurements of the PR to train the radiometer within their common (and relatively narrow) swath, and thus improve the skill of the radiometer at estimating rain outside the common swath. Indeed, several groups have proposed methods to perform this training (see e.g. Masunaga and Kummerow, 2005; Grecu and Olson, 2006; Jiang and Zipser, 2006; Viltard et al, 2006). Historically, the passive-microwave estimation algorithms differ not only in the approach and specifics of each algorithm, but also in the quantities that are estimated: some try to estimate vertically-averaged quantities, most notably the vertically-integrated liquid water, while others try to estimate the near-surface rain rate. Since it is no more or less credible to estimate the rain rate at the surface than it is to estimate the rain rate at any discrete altitude, it seems particularly interesting to try to quantify to what extent it is indeed credible and feasible to estimate vertical profiles of precipitation from a multi-channel passive microwave radiometer, what the obstacles might be, and what vertical resolution would be achievable, with some quantitative assessment of the uncertainty in one's estimates. This paper reports on our attempt to start answering these questions.

To place this project in perspective, it is worth recalling the first attempt to estimate precipitation from space with microwave radiometry, namely the Electronically Scanning Microwave Radiometer (ESMR), a single-frequency (19.35 GHz, i.e. 15.5 mm), single-polarization instrument that was launched into space on the Nimbus-5 sun-synchronous satellite in December 1972. ESMR was the first space-borne instrument to use microwave absorption to estimate precipitation by quantifying the increase in optical depth, which corresponds to higher 19-GHz brightness temperatures. In this case, one channel yields one measurement for each field-of-view, which can then be related to one vertically-integrated precipitation quantity. The second incarnation of ESMR, flown on Nimbus-6, also had a single frequency, 37 GHz. However, the third incarnation, the Scanning Multi-channel Microwave Radiometer (SMMR), flown on Seasat and on Nimbus 7, consisted of a five-channel (0.81, 1.36, 1.66, 2.80, 4.54 cm), dual-polarization instrument. And, indeed, in clear air, the measurements from SMMR's multiple channels were used to estimate the surface wind speed, surface temperature and vertically-integrated precipitable water underlying the measured brightness temperatures, a feat which would not have been possible with a single-channel measurement. SMMR could also measure cloud liquid water content in non-raining conditions. In precipitation, one could contemplate using the multiple channels to obtain vertical information about the distribution of the water, but it is not at all clear to what extent the amount of information in the measurements would allow one to do so, as the radiances now depend on the rain as well as on the surface wind, surface temperature, vertical distribution of water vapor etc.

While SMMR did not benefit from any co-located or simultaneous direct measurements, such as those of a rain radar, to allow one to evaluate the information content of the measured brightness temperatures and how correlated it is with the vertical distribution of the condensed water, TRMM does. Moreover, the fact that TRMM's PR has a reasonably high resolution (of about 4.5 km, after the satellite boost in 2001 to its current 405-km altitude) suggests that one should be able to analyze and perhaps quantify the imperfections in one's passive vertical estimates. To that end, we selected five study regions: the Eastern Mediterranean, the Eastern Atlantic off the continental shelf of North Africa, the Western Atlantic off the Southeastern seaboard of the United States, the Eastern Pacific off Southern California and the Yellow Sea / Sea of Japan region. We started by quantifying the vertical variability of rainfall as derived from the TRMM radar in each of these five regions. This is described in section 1. A straightforward Bayesian estimation approach was applied to the TMI measurements co-located with the radar profiles. The results are described in detail in section 2. We then proceeded to identify the main reason for the apparent failure of the straightforward training method for the Mediterranean-like regions, and derived a filtering approach to get around the problem. This is described in section 3, which also quantifies the residual uncertainty in our estimates. It is important to note that the specific Bayesian retrieval procedures which we apply are not necessarily optimal (except in a narrow sense), because they do depend on the choice we made for the form of the measurement vector. But we do identify the main obstacle to the vertical profiling of precipitation using passive microwave radiometry, and the result does suggest one procedure which produces verifiably improved estimates.

1. Vertical variability of precipitation

We selected five study regions between latitudes 30°N and 36.23°N: the Eastern Mediterranean, extending from 10°E to 35°E, from 1 January to 30 April; the Yellow Sea / Sea of Japan area, extending from 120°E to 135°E, from 1 May to 30 September; the Eastern Pacific, extending from 135°W to 110°W, from 1 December to 30 April; the Western Atlantic, extending from 80°W to 65°W, from 1 June to 30 September; and the Eastern Atlantic, extending from 25°W to 10°W, from 1 November to 31 March. In each region, we compiled a database of co-located radar and radiometer measurements as illustrated in figure 1. The figure shows four contiguous and overlapping 10.7-GHz fields-of-view (thin solid lines) that include two consecutive beams from two consecutive scans, along with four contiguous and overlapping 19.4-GHz fields-of-view (dashed lines), four contiguous 37-GHz fields-of-view (thick solid lines) and the eight 85.5-GHz fields-of-view that fit within the 37-GHz footprints. We average the measured brightness temperature at each frequency and polarization measured in all beams, and co-register the result with the precipitation profile estimated by the radar within the 3 cross-track \times 3 along-track = 9 radar beams that are closest to the center of the radiometer pattern (the radar-estimated precipitation is averaged at each altitude across the nine radar beams). The TRMM radar product contains a land/ocean flag, and only retain measurements where all the radar beams within the four 3-dB 10.7-GHz fields of view were devoid of land.

Thus each regional database would allow one to estimate, from a set of measured TMI brightness temperatures as in figure 1, a vertical precipitation profile at a horizontal resolution corresponding to 3x3 contiguous radar beams, i.e. about 13.5 x 13.5 km². Indeed, each database associates to each set of co-located radiometer measurements (as in figure 1) a vector of vertical rain-rate profiles obtained by horizontally averaging the rain-rate estimates which the TRMM standard combined radar/radiometer algorithm calculates within the closest 9-tuple of contiguous radar beams. The combined algorithm is described in Haddad et al, 1997, and the most recent validation analysis of its estimates can be found in Wolff and Fisher, 2008. Each vertical rain-rate profile has up to 80 bins, representing the rain rate within each of eighty 250m-thick layers. To quantify the vertical variability of the rain rate (and consequently keep the size of each database to the minimum necessary to reconstruct the precipitation profiles adequately), we performed a principal component analysis on the rain-rate profiles in each region. The results are summarized in table 1.

In all regions, the first principal component is essentially the vertical average of the rain rates, as was observed in several previous analyses of the global TRMM data (see e.g. Coppens et al, 2000). For the Western Atlantic and the Yellow Sea / Sea of Japan regions, the magnitude of the coefficients of the second principal component does not change much in height, but the coefficients are positive from the surface to 2.25 km, and negative from 2.5 km up: that principal component is thus essentially the difference between the rain rate below about 2.25 km and the rain rate above that height (as was observed in the global analysis in Coppens et al, 2000). For the other three regions, the sign of the coefficients of the second principal component also changes only once, but it does so at 2 km rather than 2.25 km. In all regions, the first three principal components account for over 96% of the vertical variability of the precipitation. That is why we decided, as a first step, to use each database to match measured brightness temperatures to the first three principal components of the rain only, and produce an estimate of the vertical rain-rate profile from the estimates of the first three principal components along with the mean values over the entire database for the higher-order components. Thus, given measured brightness temperatures \vec{T} , and calling \vec{S} the vector of principal components of the rain \vec{R} , we start by estimating the first three rain principal components S_1 , S_2 , and S_3 using the Bayesian estimators

$$\hat{S}_n = \frac{1}{K} \sum_{i \in \text{database}} S_n^{(i)} e^{-0.5 (\vec{T} - \vec{T}_i)' C^{-1} (\vec{T} - \vec{T}_i)}, \quad n = 1, 2 \text{ or } 3 \quad (1)$$

in which K is the normalizing constant $K = \sum_i \exp[-0.5(T - \vec{T}_i)' C^{-1}(T - \vec{T}_i)]$, and C is the sum of the observation error variance in T (a diagonal matrix) and the conditional covariance matrix of \vec{T} given \vec{R}_i . In practice, to make sure the problem stays tractable and to control the size of the database, for C we diagonalized the covariance matrix of \vec{T} for each region, and subsequently retained only the first three principal components of the brightness temperatures, which we then used to compute the quadratic distances in the exponent of (1). Once \hat{S}_1 , \hat{S}_2 and \hat{S}_3 have been determined, we use the means of the remaining components of \vec{S} and reconstruct \vec{R} by multiplying \vec{S} by the appropriate change-of-basis matrix.

We used data from 2007 to build the databases along with the statistical analyses, and applied the resulting Bayesian procedure to data from 2006.

2. Vertical profiling using the databases – first approach

The results of our first (straightforward Bayesian matching) approach are illustrated in figure 2a, which shows estimated vs actual rain rates for 16,374 data points in the Yellow Sea region, at altitudes of 0, 2000 and 4000 meters above sea level. The biases are below 11%, and the r.m.s. uncertainty in the estimated rain rate (i.e. the rms deviation, from unity, of the ratio of estimated to true rain rate) is about 70%, a quite encouragingly low value. Figure 2b similarly compares our estimates vs actual rain rates for 17,918 data points in the Western Atlantic region.

In contrast, the results over the Mediterranean seem quite discouraging, as figure 3 (representing 13718 samples) illustrates. The vertically-integrated and near-surface comparisons are very disappointing, with a scatter so large that the mismatch cannot be characterized merely by bias or variance. However the comparison at 4000m suggests that some measurement instances may be more amenable than others to estimate the underlying vertical distribution of rain, as evidenced by a small but noticeable portion of samples clustering around the diagonal. To investigate the possible causes of the vastly poorer performance of our approach in this case compared to the Yellow-Sea case, we first verified that the principal-components reconstruction was not the culprit. Figure 4 compares the actual near-surface rain rates with those that we reconstructed using the exact values of the first three principal components and the means of all the higher-order principal components, for a subset of data collected during 2006 over each of the two regions. Manifestly, the reconstructed rain rates over the Mediterranean show a broader scatter, consistent with the smaller portion of the total variability captured by the first three principal components (see table 1), but the increased scatter (unaccompanied by any bias) cannot explain the glaring mismatch in figure 3.

We next examined the other side of the database, namely the brightness temperatures in the two regions. Table 2 shows the coefficients of the first two principal components in the two cases. Comparing the two coefficients with the largest magnitudes (highlighted in bold) to the remaining ones, for each region, confirms that the first principal component can be considered an emission index, while the second can be considered a scattering index. More interesting is the main contrast between the coefficients for the two regions, namely, in the case of the Mediterranean, the large sensitivity of the first principal component to the 85H channel with no sensitivity to the 85V channel: because the precipitation signature is largely un-polarized, contrary to that of the sea surface, this suggests a strong sensitivity of the brightness temperatures to the sea-surface signal – a signal that seems absent in the case of the Yellow Sea, and could very well explain the success of our first passive-microwave profiling attempt in the latter region and its failure in the former.

To verify that that is indeed the explanation for the success of our straightforward approach in the first case and its failure in the second, we decided to sort the data according to discrepancy in the 85-GHz polarizations. Spencer, Goodman and Hood (1989) were first to identify the need for a polarization correction for water surfaces, and devised a linear combination of the 85-H and 85-V channels that produces a polarization-corrected temperature, in this case given by $PCT_{85} = 1.818 T_{85V} - 0.818 T_{85H}$. Rather than correcting for polarization, we would like to identify when the difference in the two polarizations is greatest, and we therefore use our heuristic weighted-polarization-difference-in-precipitation discriminant WPDiP, empirically derived as $WPDiP = T_{85V} - 0.83 T_{85H}$ and empirically verified to fall generically below 50.5K for fields-of-view that are filled with rain, and well above 50.5K for clear sea surface. The right panels in figure 5 reproduce the previous comparisons for samples over the Mediterranean satisfying $WPDiP > 50.5$, while the left panels plot only those points with $WPDiP < 50.5$. The graphs confirm that the problem is the all-too-frequently inhomogeneous composition of the field of view, when the radiometer captures some precipitation in part of the beam as well as clear surface in the remainder of the field of view.

3. Neutralizing the surface contribution in a given beam

Having the quantitative means to detect the problem is good, but it only filters those measurements where our straightforward approach produces good results (illustrated in the left panels of figure 5) and it does not help one make estimates for the surface-contaminated cases (the right panels of figure 5). One could try to characterize the radiometric signature of the surface, for example by studying no-rain areas near the rainy beam of interest, and then subtract this contribution or otherwise account for it in the rain retrieval. The main problem with such an approach is that one cannot establish a priori the extent of the inhomogeneity within a given beam. In other words, if λ is the fraction of the field of view that contains rain, with \vec{R} the rain rates in this rainy portion, and if \vec{T} represents the measured brightness temperatures, the difficulty arises from the fact that we measure

$$\vec{T} = \lambda T(\text{rain}) + (1 - \lambda) T(\text{clear}) \quad (2)$$

while the corresponding estimates \hat{R} should be

$$\hat{R} = \lambda \vec{R} + (1 - \lambda) 0 \quad (3)$$

The best way to solve this problem would be to transform the measurement \vec{T} so as to make the contribution of the clear-surface term (the second summand on the right-hand side of (2)) vanish.

To that end, we performed a principal-components analysis of the clear-air brightness-temperature measurements in the Mediterranean case. Since the variability of the radiances is dominated by the surface wind and temperature, one would expect that, of the eight principal components, six should have relatively small variances. Indeed, the eigenvalues in decreasing order turn out to be

$$539.7 > 145.2 > 15.2 > 9.4 > 5.6 > 0.75 > 0.33 > 0.13$$

Thus, if we call U the orthogonal matrix which diagonalizes the covariance C_T of \vec{T} , U_i the i^{th} column of U (see table 3), and μ_i the mean value of $U_i^t \cdot \vec{T}$ (i.e. of the i^{th} principal

component) under clear conditions, then $U_8^t \cdot \vec{T}$ can be considered constant and equal to μ_8 (up to an r.m.s. error of $0.13^{1/2} = 0.36$ K, i.e. far less than the measurement error), and similarly for $U_7^t \cdot \vec{T}$, ... up to $U_3^t \cdot \vec{T}$. In rain, it is therefore natural to use, instead of the eight components of \vec{T} , the six transform measurements $U_3^t \cdot \vec{T}$, ... , $U_8^t \cdot \vec{T}$, so that we would then be trying to estimate (3) from

$$U_i^t \cdot \vec{T} \approx \lambda U_i^t \cdot T(\text{rain}) + (1 - \lambda) \mu_i, \quad i = 3, \dots, 8 \quad (2')$$

Instead of the original eight measurements, we end up with only six quantities with which to estimate the rain vertical principal components, but the advantage of the latter six is that they are almost entirely insensitive to the surface contribution to the original eight: this should prevent (or at least strongly constrain) the database comparison process from latching onto the portion of the variability that is due to surface effects – and consequently help ensure that the rain estimation is reflecting the sensitivity of the brightness temperature to the precipitation itself, thus producing estimates that are better matched to the rain. Figure 6 illustrates the results that we obtained with the same samples as in figure 5, using the six principal components with the lowest variance as in (2') in the Bayesian procedure described in the last paragraph of section 1. While the performance is vastly improved, it is manifestly still imperfect – the right panels of figure 6 show a rather wide scatter. It is almost certain that this residual uncertainty can be reduced significantly, because the estimation procedure which we implemented is still far from optimal, since we are again (as in the original application) using the first three principal components of the transformed observations only: this is necessary to keep the size of the database manageable, but it is still a rough approximation to the exact procedure summarized in equation (1). Implementing the latter exactly would require one to bin any observation according to an appropriately fine discretization of our 6-dimensional space, an operation which would require a prohibitively large number of samples if we have any hope of ending up with more than one or two samples in each 6-dimensional bin. We are actively studying ways to improve that aspect of the approach. This is an important issue, because the fact that the top three principal components capture most of the variability in our measurement vector does not necessarily imply that they reflect adequately that aspect of the variability which is most sensitive to the precipitation. Nevertheless, at any given altitude the bias in our new retrievals is less than 5% and the r.m.s. uncertainty is about 80% (while the uncertainty in the vertically-integrated rain rate is about 51%), very encouragingly low values.

While other approaches have been considered to account for the partial-beam-filling problem (see e.g. Chiu and Petty, 2006, and the references therein), our approach is closer in its physical and mathematical justification to the method developed by Grody (1991) to detect precipitation – with the notable difference that our goal is to detect and neutralize the contribution from the clear surface within the partially rainy field of view. The approach does succeed in producing estimates of the vertical profiles of precipitation, and the uncertainty in these altitude-specific estimates is demonstrably only slightly greater than the uncertainty in our estimates from measurements that do not include a significant clear-surface contribution.

4. Conclusions

The first interesting conclusion from this study is that the five regions considered, all lying above 30°N, fall in one of two cases: a monsoon-like case (which includes the Yellow Sea / Sea of Japan and the Western Atlantic) where the first principal component of the vertical precipitation accounts for about 90% of the variability and the second accounts for another 6% or so, and a Mediterranean-like case (including the Eastern Pacific, Eastern Atlantic and the Mediterranean) where the first principal component accounts for 80% of the variability and the second for 11%. When a straightforward Bayesian estimation approach was applied to the TMI measurements co-located with the radar profiles, for the monsoon-like regions it produced estimates of rain rates at 250 meter vertical increments which, when compared to the TRMM radar estimates, showed a negligible bias and a standard deviation of between 60% and 70%. For the Mediterranean region, the errors with this direct approach were very large and prompted us to look for the main cause of the degraded performance and a possible solution. Using a weighted-polarization-difference-in-precipitation discriminant, we established that the problematic beams are the ones observing the clear surface in a significant portion of the field of view. The solution we propose relies on selecting the small-eigenvalue principal components obtained from an analysis of clear-surface returns, and applying the database-Bayesian retrieval using that smaller set of (linearly) transformed measurements. The results show a markedly improved performance in the Mediterranean region.

Thus the second main conclusion is that we can effectively classify each vector of measured brightness temperatures according to whether it is affected by partial clear-surface radiances, and when the unaffected measurements are considered separately, the vertical profiles of the underlying precipitation can be adequately retrieved using a straightforward Bayesian approach. On the other hand, vertical precipitation profiles can be retrieved from measurements that are only partially filled with precipitation, as long as the partial clear-surface signature is neutralized in the measurements – a process which we accomplish by linearly transforming the original measured brightness temperatures and retaining a subset of the transformed measurements, discarding those that are most sensitive to the variable surface emissivity.

It is important to add that the estimation procedure that we implemented is not optimal for several reasons: foremost, the approach treats the combined-radar-radiometer vertical profiles of precipitation as the reference against which the passive measurements are to be trained – this is a shortcoming because the radar-derived profiles do not constitute absolute truth (see, e.g., Wolff and Fisher, 2008, for the most recent and comprehensive description of the comparison of the retrievals with ground validation data); another reason that our approach is not optimal is that we only used the first three principal components of the observations, instead of the full vector of measured (or transformed, in the modified surface-neutralizing method) brightness temperatures, to keep the size of the database – and the mechanics of the estimation procedure – manageable; while we were careful to include only measurements over water, we did not classify the observations according to near-surface wind or sea-surface temperature or rain regime, and our radar-based land/ocean classifier does make simplifications which cannot prevent some

misclassification; perhaps most important, we performed the principal component analyses on the rain rates and the brightness temperatures themselves, tacitly implying that Gaussian statistics govern the joint behavior of these variables, when it is not necessarily true that the marginal distributions are normal, let alone the joint density functions. An approach which avoids these simplifications should improve the performance of the retrievals. With the validity of the concept established and the main problems identified, the next task is therefore to develop an effective retrieval algorithm that is as nearly optimal as possible.

5. Acknowledgements

We wish to thank C. Kummerow for pointing out similarities with other groups' results. This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

6. References

Chiu, J.C., and G.W. Petty, 2006: Bayesian retrieval of complete posterior PDFs of oceanic rain rate from microwave observations. *J. Appl. Meteor. Clim.* **45**, 1073-1095.

Coppens, D., Z.S. Haddad and E. Im, 2000: Estimating the uncertainty in passive-microwave rain retrievals. *J. Atmos. Ocean. Tech.* **17**, 1618-1629.

Greco, M., and W.S. Olson, 2006: Bayesian Estimation of Precipitation from Satellite Passive Microwave Observations Using Combined Radar-Radiometer Retrievals. *J. Appl. Met. Clim.* **45**, 416-433.

Grody, N.C., 1991: Classification of snow cover and precipitation using the Special Sensor Microwave Imager. *J. Geophys. Res.* **96**, 7423-7435.

Haddad, Z.S., E.A. Smith, C.D. Kummerow et al, 1997: The TRMM 'Day-1' radar/radiometer combined rain-profiling algorithm, *J. Met. Soc. Japan* **75**, 799-809.

Jiang, H., and E. J. Zipser, 2006: Retrieval of hydrometeor profiles in tropical cyclones and convection from combined radar and radiometer observations. *J. Appl. Meteor.* **45**, 1096-1115.

Masunaga, H., and C.D. Kummerow, 2005: Combined radar and radiometer analysis of precipitation profiles for a parametric retrieval algorithm. *J. Atmos. Ocean. Tech.* **22**, 909-929.

Spencer, R.W., H.M. Goodman and R.E. Hood, 1989: Precipitation retrieval over land and ocean with SSM/I: Identification and characteristics of the scattering signal. *J. Atmos. Ocean. Tech.* **6**, 254-273.

Viltard, N., C. Burlaud and C.D. Kummerow, 2006: Rain retrieval from TMI brightness temperature measurements using a TRMM PR-based database. *J. Appl. Meteor. Clim.* **45**, 455-466.

Wolff, D.B., and B.L. Fisher, 2008: Comparisons of instantaneous TRMM ground validation and satellite rain-rate estimates at different spatial scales. *J. Appl. Meteor. Clim.* **47**, 2215-2237.

Figure Captions

Figure 1. Typical radiometer fields-of-view pattern for one sample in our database: four 10.4-GHz beams (thin solid lines representing the 3-dB beam width), four 19.4-GHz beams (dashed lines), four 37-GHz beams (thick solid lines), and eight 85.5-GHz beams (shaded). This example illustrates a pattern along the satellite track, with the arrow indicating the direction of motion. The grid spacing approximates the resolution of the radar beams, with the 3x3 beams that would be associated to the depicted radiometer pattern shaded in light blue.

Figure 2a: Estimated vs actual rain rates in the Yellow Sea / Sea of Japan region, at 0m ASL (top left), 2000m ASL (top right), and 4000m ASL (bottom left).

Figure 2b: Estimated vs actual rain rates in the Western Atlantic, at 0m ASL (top right), 2000m ASL (top right), and 4000m ASL (bottom left).

Figure 3: Estimated vs actual rain rates in the Mediterranean region, at 0m ASL (top right), 4000m ASL (bottom left), and vertically integrated (top left).

Figure 4: Comparison of actual surface rain rates with those obtained from the rain profiles reconstructed using the first three principal components only.

Figure 5: Estimated vs actual rain rates in the Mediterranean region, at 2000m ASL and vertically-integrated, for mostly rainy samples (left) and significantly surface-contaminated ones (right).

Figure 6: Estimated vs actual rain rates for the Mediterranean using the original direct Bayesian method (left panels) and the modified method using the six combinations of the radiances which vary least in clear conditions (right panels).

	E Mediterranean		Yellow Sea		E Pacific		W Atlantic		E Atlantic	
1 st eigenv.	36.1	78.7%	51.4	90.9%	14.3	80.7%	52.4	90.4%	19.3	82.2%
2 nd eigenv.	6.31	13.7%	3.34	5.9%	2.1	11.6%	3.64	6.3%	2.58	10.5%
3 rd eigenv.	1.77	3.9%	0.89	1.6%	0.68	3.9%	0.96	1.7%	0.83	3.5%
4 th eigenv.	0.57	1.2%	0.26	0.5%	0.23	1.3%	0.27	0.5%	0.24	1%
5 th eigenv.	0.25	0.5%	0.13	0.2%	0.1	0.6%	0.14	0.2%	0.11	0.5%
6 th eigenv.	0.16	0.3%	0.09	0.2%	0.06	0.4%	0.09	0.2%	0.08	0.3%
7 th eigenv.	0.12	0.3%	0.07	0.1%	0.05	0.3%	0.07	0.1%	0.06	0.2%
8 th eigenv.	0.1	0.2%	0.06	0.1%	0.04	0.2%	0.06	0.1%	0.05	0.2%

Table 1: Eigenvalues and the proportion of the variance corresponding to the first eight principal components of the vertical rain-rate vector for each region.

	Yellow Sea / Sea of Japan		Mediterranean	
	T_1'	T_2'	T_1'	T_2'
10V	0.241	-0.141	0.125	-0.161
10H	0.440	-0.249	0.224	-0.288
19V	0.304	0.043	0.262	-0.126
19H	0.572	0.055	0.497	-0.265
37V	0.213	0.264	0.304	-0.057
37H	0.491	0.410	0.658	-0.009
85V	-0.167	0.503	0.046	0.554
85H	-0.114	0.650	0.302	0.704

Table 2: Coefficients of the first two principal components of the brightness temperatures for two regions in our study.

	U_1	U_2	U_3	U_4	U_5	U_6	U_7	U_8
10V	0.1388	0.4847	-0.2307	0.0139	-0.2358	0.5459	-0.5410	0.2146
10H	0.2453	0.7383	-0.2495	0.1922	0.1730	-0.3583	0.3454	-0.1344
19V	0.2200	0.0930	0.1659	-0.4695	-0.2492	0.3319	0.2283	-0.6860
19H	0.4267	0.1259	0.4378	-0.5508	0.3275	-0.1736	-0.1434	0.3860
37V	0.2192	-0.0274	0.1042	0.0532	-0.6443	0.0597	0.5251	0.4930
37H	0.4948	-0.0899	0.4792	0.6041	-0.1147	-0.0840	-0.2727	-0.2408
85V	0.1941	-0.1528	-0.3853	-0.2624	-0.4618	-0.5950	-0.3751	-0.1171
85H	0.5997	-0.4042	-0.5257	0.0493	0.3217	0.2620	0.1566	0.0386

Table 3: Coefficients of the principal components of the brightness temperatures in clear air in the Mediterranean. Because U_1 and U_2 capture over 95.6% of the variability in the clear-air measurements, the deviations of (U_3, \dots, U_8) in rain from their mean values in clear air should be due to the precipitation.

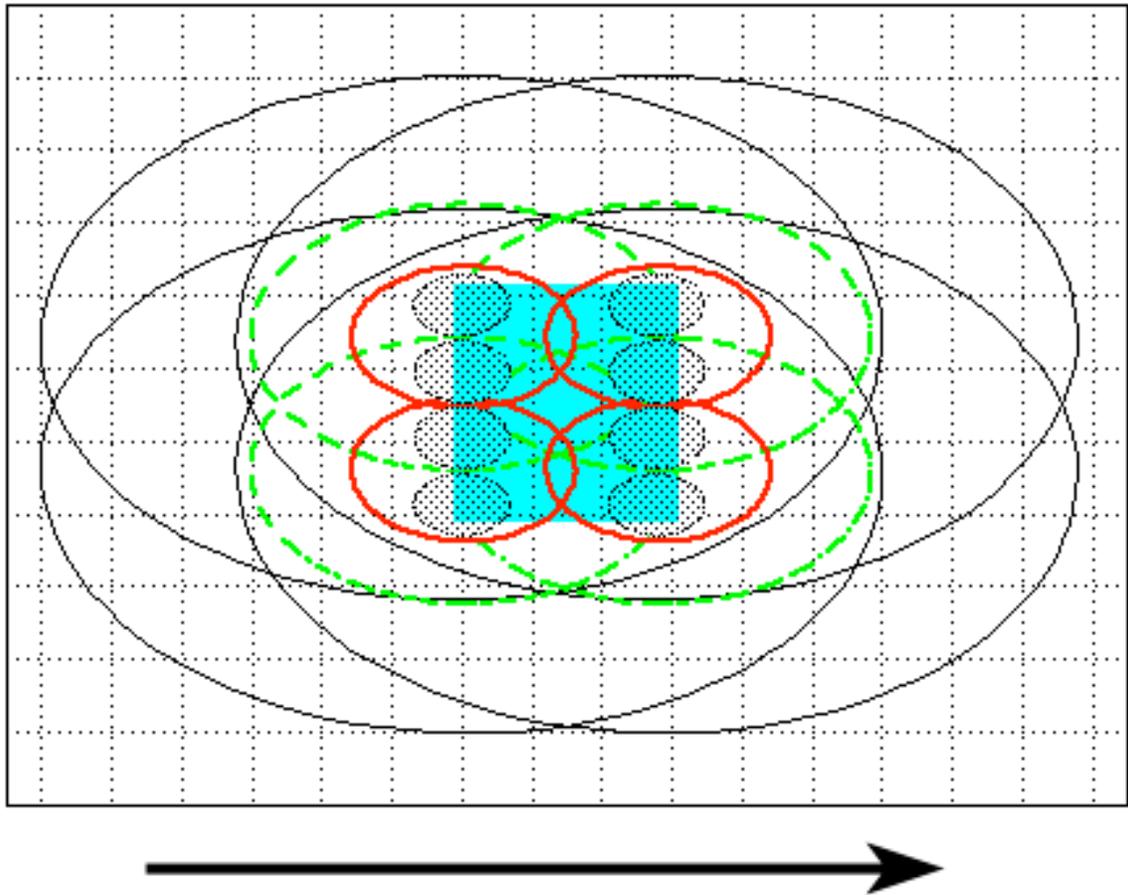


Figure 1

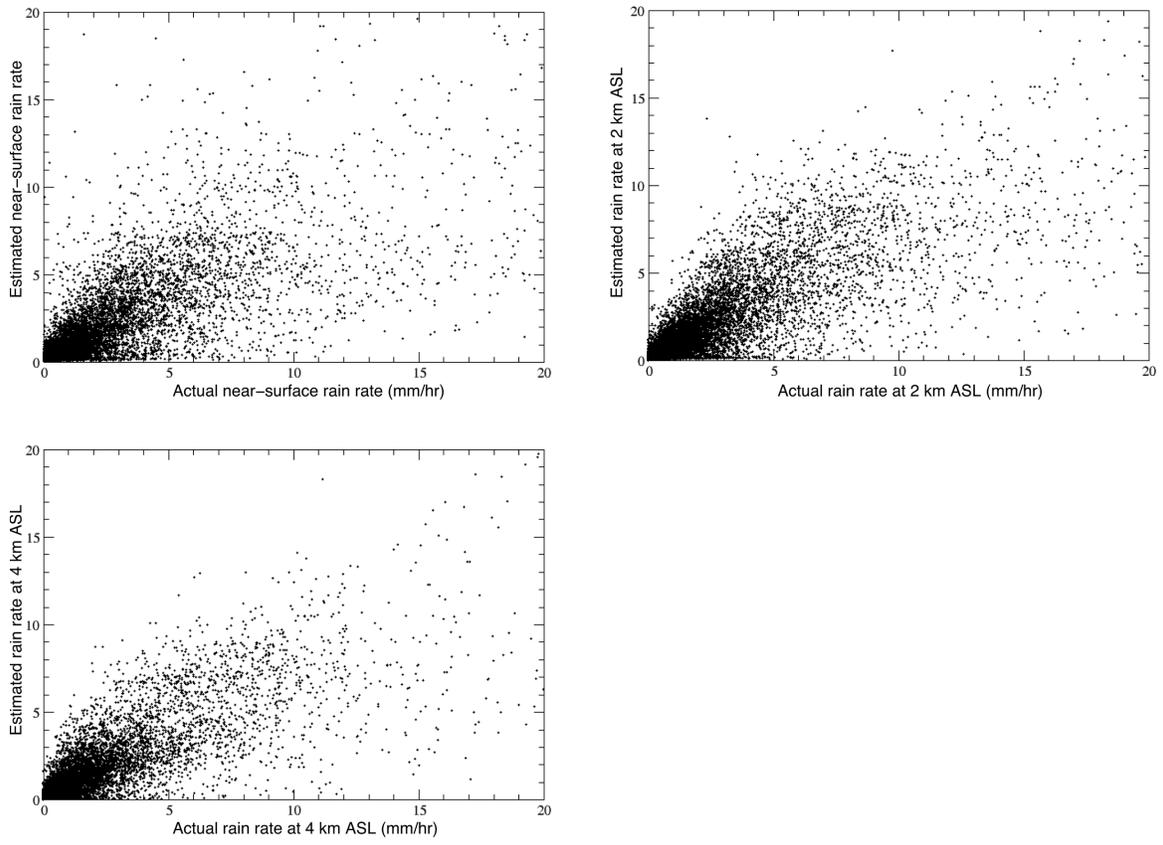


Figure 2a

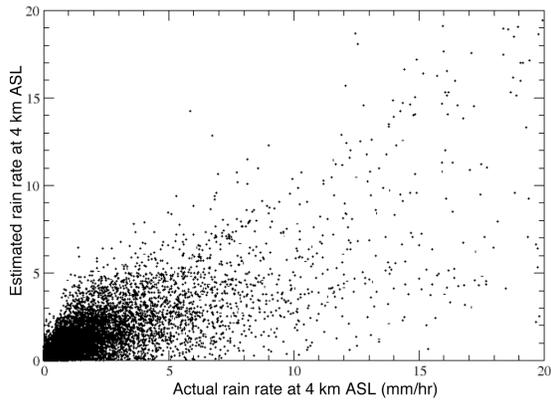
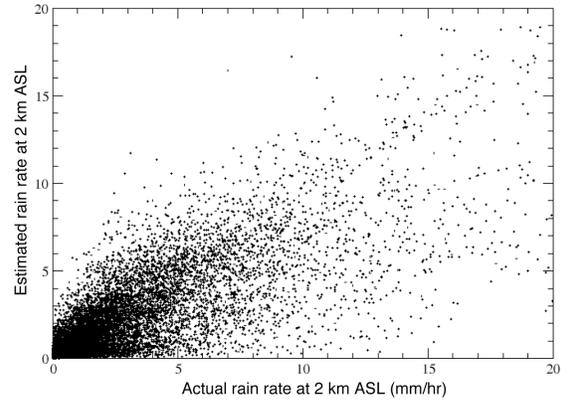
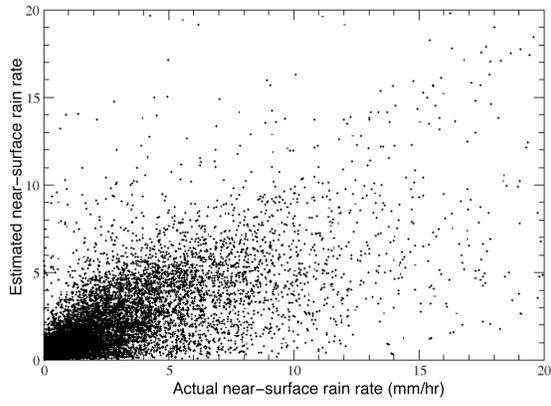


Figure 2b

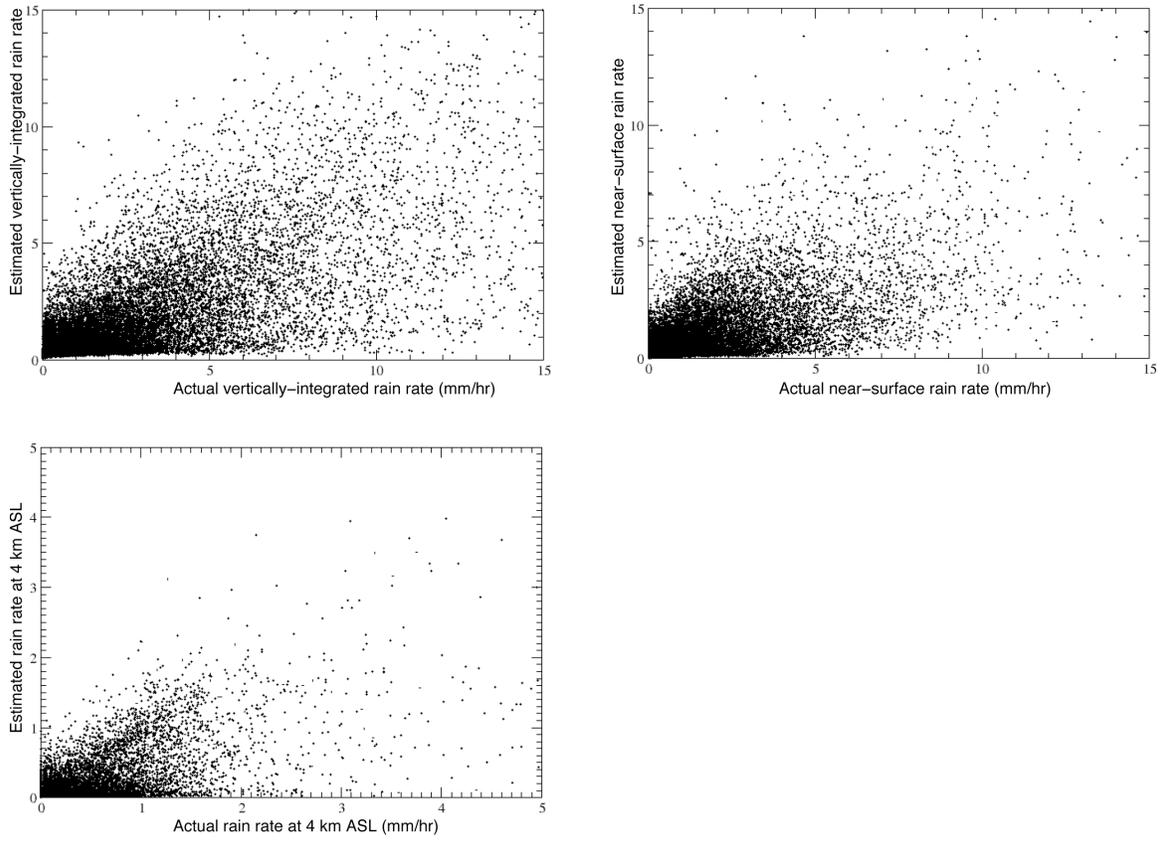


Figure 3

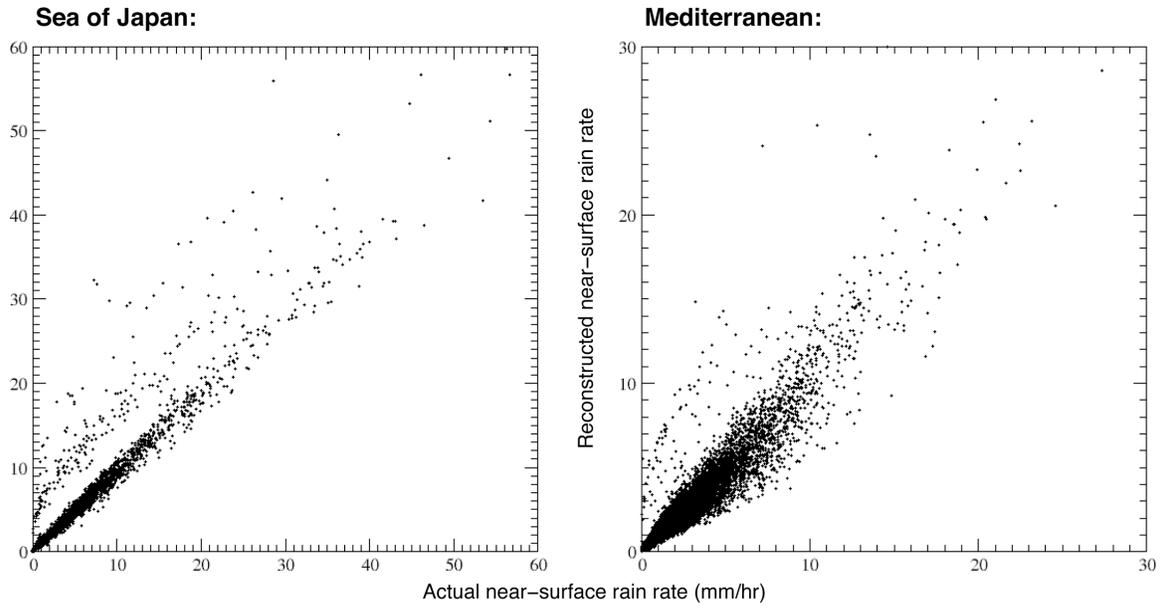


Figure 4

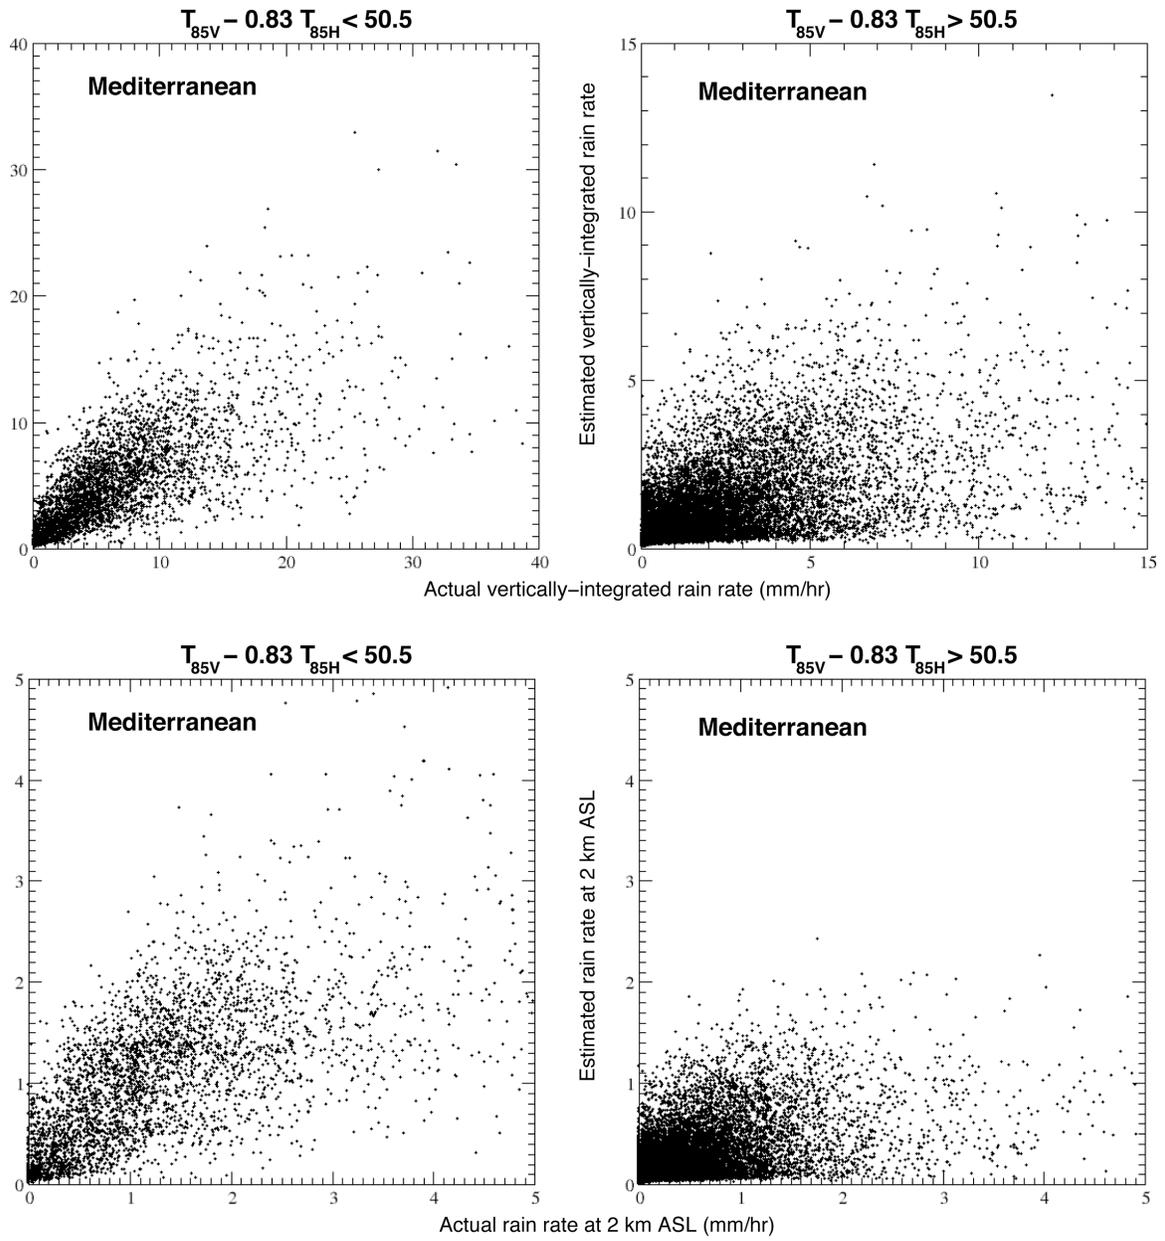


Figure 5

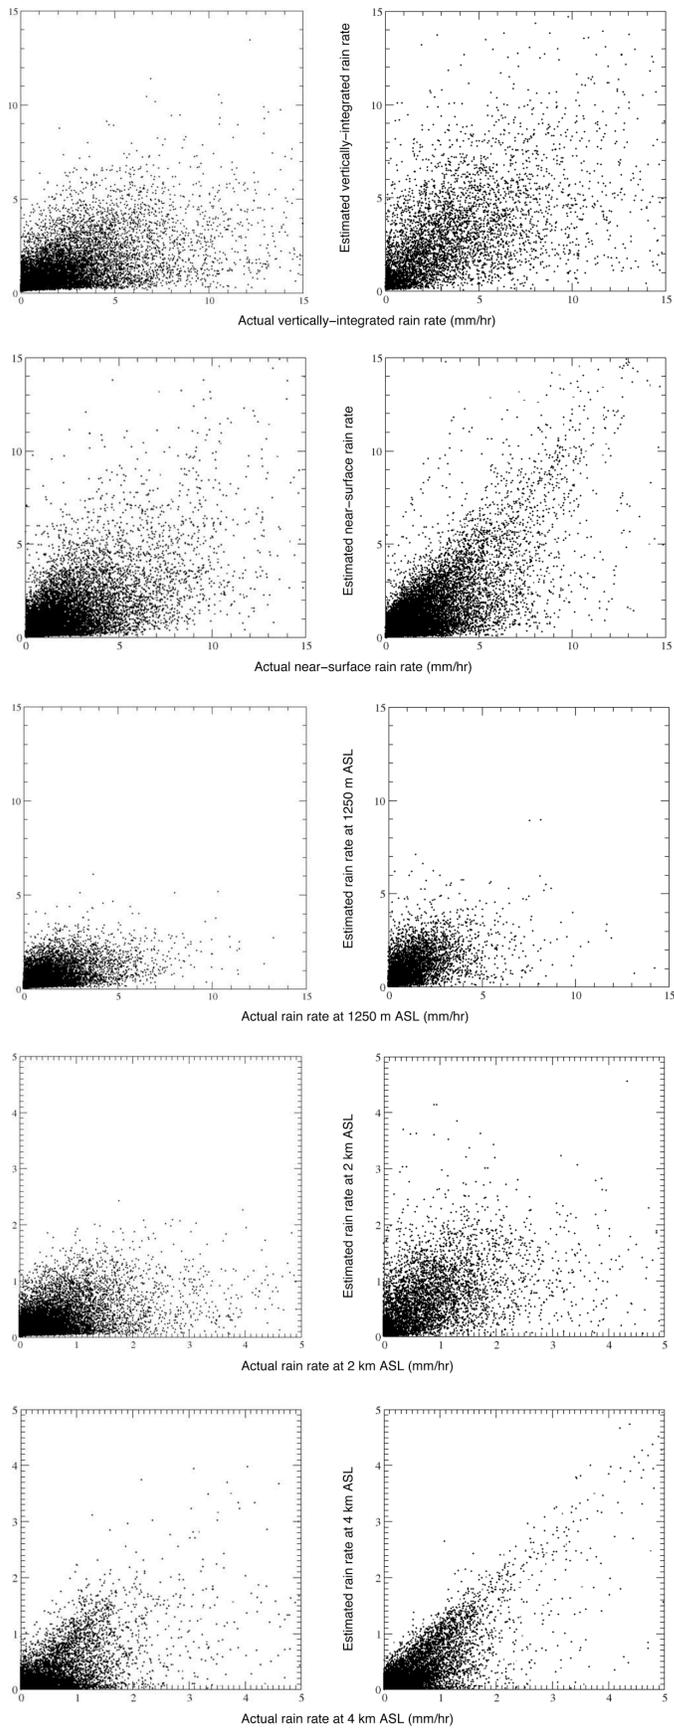


Figure 6