

Rainfall Estimation from a Combination of TRMM Precipitation Radar and GOES Multispectral Satellite Imagery through the Use of an Artificial Neural Network

TIM BELLERBY

Department of Geography, University of Hull, Hull, United Kingdom

MARTIN TODD

School of Geography, University of Oxford, Oxford, United Kingdom

DOM KNIVETON

Department of Geography, University of Leicester, Leicester, United Kingdom

CHRIS KIDD

School of Geography, University of Birmingham, Birmingham, United Kingdom

(Manuscript received 27 September 1999, in final form 14 February 2000)

ABSTRACT

This paper describes the development of a satellite precipitation algorithm designed to generate rainfall estimates at high spatial and temporal resolutions using a combination of Tropical Rainfall Measuring Mission (TRMM) precipitation radar (PR) data and multispectral Geostationary Operational Environmental Satellite (GOES) imagery. Coincident PR measurements were matched with four-band GOES image data to form the training dataset for a neural network. Statistical information derived from multiple GOES pixels was matched with each precipitation measurement to incorporate information on cloud texture and rates of change into the estimation process. The neural network was trained for a region of Brazil and used to produce half-hourly precipitation estimates for the periods 8–31 January and 10–25 February 1999 at a spatial resolution of 0.12 degrees. These products were validated using PR and gauge data. Instantaneous precipitation estimates demonstrated correlations of ~ 0.47 with independent validation data, exceeding those of an optimized GOES Precipitation Index method locally calibrated using PR data. A combination of PR and GOES data thus may be used to generate precipitation estimates at high spatial and temporal resolutions with extensive spatial and temporal coverage, independent of any surface instrumentation.

1. Introduction

There are numerous applications in meteorology and hydrology for which accurate estimation of rainfall at small spatial and temporal scales (daily or subdaily estimates at resolutions of 25 km or better) would be invaluable. Such information is available for limited areas using combinations of ground-based radar and dense networks of rain gauges. For large areas of the globe, however, the in situ infrastructure necessary for this form of precipitation-monitoring network is not in place.

Satellite-based precipitation-monitoring techniques

are well established. However, their ability to generate products at high spatial and temporal resolutions is limited [see Barrett and Beaumont (1994) for a general view]. Continuous, high-temporal resolution satellite data are, of course, only available from instruments mounted on geostationary platforms. This restriction currently limits the available wavelengths to those in the infrared (IR) and visible (VIS) parts of the electromagnetic spectrum. The lack of VIS data at night has generally restricted geostationary precipitation-monitoring techniques to the use of IR data alone. Although satellite IR algorithms benefit from high temporal sampling, IR radiances from cloud tops have only an indirect relationship with surface rainfall, resulting in weak statistical relationships between cloudiness and precipitation. The most commonly employed techniques count cloudy pixels that are colder than a given threshold tem-

Corresponding author address: Tim Bellerby, Department of Geography, University of Hull, Hull HU6 7RX, United Kingdom.
E-mail: T.J.Bellerby@geo.hull.ac.uk

perature. It is assumed that such pixels are associated with probably precipitating cumulonimbus clouds possessing cold, high tops. High-level cirrus and other non-precipitating cloud forms interfere with this simple relationship. Further difficulties are introduced by the lack of any direct information on rain rates beneath precipitating clouds. Accordingly, many simple IR algorithms rely on the effects of substantial scale averaging to improve accuracy, for example the Geostationary Operational Environmental Satellite (GOES) Precipitation Index (GPI; Arkin and Meisner 1987).

The relationships between IR-derived cloud indices and rainfall are variable in both space and time (Bellerby and Barrett 1993). An acknowledgement of this variability has led to the development of continuously calibrated IR precipitation algorithms. Todd et al. (1995, 1999) describe an IR rainfall estimation technique, designed for operational hydrological applications, in which IR rain–no-rain thresholds and rain rates are optimized in space and time through a comparison with real-time rain gauge observations. It was observed that optimum IR thresholds over East Africa are variable in time and space as a result of variability in local-scale rainfall and cloud characteristics. The resulting optimized IR threshold fields identified rainfall events more accurately than did fixed threshold fields. Herman et al. (1997) describe another technique whereby 10-day IR estimates of rainfall are corrected using gauge data. Unfortunately, all of these methods are sensitive to the density and spatial distribution of the gauge network and are applicable only to regions where high-quality real-time gauge data are available.

Satellite passive microwave (PMW) sensors such as the Special Sensor Microwave Imager (SSM/I) and, more recently, spaceborne precipitation radar systems (PR) are able to provide significantly more accurate estimates of instantaneous rain rates but suffer from the poor temporal sampling associated with platforms in low Earth orbit. This low sampling rate makes these instruments most suitable for estimation of accumulated rainfall over periods of one month or greater. To account for limitations in the accuracy of rainfall estimates from both polar-orbiting PMW and geostationary IR satellite data, a number of methods have been developed that combine these two data types. Adler et al. (1993) developed an adjusted GPI in which a correction factor is derived from the comparison of PMW and GPI estimates for coincident time slots over some extended period (e.g., one month). This correction is then retrospectively applied to all the hourly GPI estimates during that period. Xu et al. (1999) have recently extended this work to develop the Universally Adjusted GPI (UAGPI) method in which both the monthly IR threshold and IR conditional rain rate are optimized using coincident PMW and IR data, again over an extended period. Todd and Kniveton (1999) developed this principle further by calibrating the IR–rain rate relationship at each 1° grid cell using coincident IR and PMW data. This method

has been successfully tested over the global Tropics and subtropics. However, all these combined IR–PMW methods rely to some extent on scale averaging in the calibration of IR parameters. More important, when using PMW SSM/I data for calibration, estimates produced at the temporal frequency of the geostationary IR imagery are likely to retain the bias associated with inadequate sampling of the diurnal rainfall cycle by SSM/I.

It is well known that rainfall estimates from sun-synchronous polar-orbiting satellites are subject to bias in regions for which the diurnal cycle of rainfall is pronounced. Morrissey and Janowiak (1996) show that for pentad and monthly estimates of tropical rainfall, the temporal subsampling associated with satellite estimation introduces a conditional bias in which the algorithm overestimates high rainfall and underestimates low rainfall. The magnitude of this conditional bias is dependent on a number of factors, including sampling rate, integration area, and study region. For a 3-hourly sampling scheme, the conditional bias is very close to zero, but it increases markedly for a 12-hourly scheme. Li et al. (1996) assessed the errors in estimation of monthly rainfall associated with a variety of sampling schemes using data from the Darwin, Australia, rain radar. It is shown that, as a result of the diurnal rainfall cycle, estimates are sensitive not only to the rate of sampling but to the starting time of the scheme. Again, a 3-hourly scheme shows minimal bias in comparison with a 12-hourly scheme, for which bias ranged from +20% to –20%, depending on sampling starting time. Soman et al. (1995) produced very similar results using the Darwin radar data from two 20-day periods. It is likely, therefore, that sampling errors will be larger over land, for which the diurnal cycle is more pronounced.

The Tropical Rainfall Measuring Mission (TRMM) provides a significant new platform for precipitation monitoring studies (Kummerow et al. 1998). The TRMM satellite is in a non sun-synchronous orbit and thus, over some extended period, provides sampling of the complete diurnal cycle of rainfall. The TRMM PR is the first spaceborne precipitation radar and is believed to provide the most accurate estimates of instantaneous rain rate available to date. A combination of high spatial resolution, high accuracy, and the absence of sun-synchronicity make the TRMM PR an ideal candidate for use in combination with geostationary data in a multisourced satellite precipitation algorithm. Given the quality of the PR data source, it seems reasonable to attempt to extract the maximum information possible from the geostationary imagery. The *GOES-8* and *-9* sensors have four channels in the mid- and thermal IR wavelengths in addition to a VIS channel (Table 1), each of which contains information on cloud characteristics and/or atmospheric moisture. The European Organisation for the Exploitation of Meteorological Satellites (EUMETSAT) Meteosat Second Generation system to be launched in 2001 will provide an extended range of

TABLE 1. Characteristics of GOES-8 and EUMETSAT Meteosat Second Generation image data.

		GOES									
Channel		1		2		3		4		5	
Wavelength (μm)		0.65		3.9		6.7		11		12	
		Meteosat Second Generation									
Wavelength (μm)	0.6	0.8	1.6	3.8	6.2	7.3	8.7	9.7	10.8	12.0	13.4

channels (Table 1). Image textures and rates of change have additionally been shown to be of value in precipitation estimation.

Hsu et al. (1999) demonstrated the utility of using artificial neural networks (ANN) to generate functions linking IR and VIS image characteristics (including image texture characteristics) to precipitation. Training their network against an extensive radar-rain gauge network, they generated daily, monthly (both at 0.25° resolution), and hourly (1.25° resolution) precipitation products that demonstrated high correlations (~ 0.7 – 0.9) against validation data. These results suggest that a relatively strong, if transient, relationship does exist between VIS/IR satellite image characteristics and precipitation. The TRMM PR provides a means to determine and to update dynamically this kind of empirical relationship to produce precipitation estimates at high temporal and spatial resolutions independent of a surface rainfall-measuring infrastructure. This paper describes the development of an empirical algorithm that combines TRMM PR data with both VIS and multichannel IR GOES data using an ANN approach.

2. Methodology

a. Neural networks

ANN provides a convenient and powerful means of performing nonlinear classification and regression (Freeman and Skapura 1991; Skapura 1996). An ANN operates by partitioning its input space into subdomains, which may then be associated with a desired set of output values. The most common network structures in current use are multilayer feed-forward networks (MFFN). These networks divide their input space using a set of hyperplanes:

$$\mathbf{w}_i \cdot \mathbf{x} = c_i. \tag{1}$$

Here \mathbf{x} is a vector of input values, \mathbf{w}_i is a vector normal to hyperplane i , and c_i is a scalar constant defining the location of the plane with respect to the origin. The actual classification is performed using nonlinear functions based on (1), the most common form being the sigmoidal function:

$$n_i(\mathbf{x}) = f(\mathbf{w}_i \cdot \mathbf{x}), \quad \text{where } f(s) = \frac{1}{1 + e^{-s+c}}. \tag{2}$$

Each expression $n_i(\mathbf{x})$ responds to a given input pattern by changing its output state. By analogy to biological systems, these expressions are known as artificial

neurons. A sufficiently large collection (or layer) of artificial neurons may be used to define an arbitrarily complex partition of the input space. To complete the mapping, the vector $\mathbf{n}(\mathbf{x}) = [n_i(\mathbf{x})]$ must be associated with a desired output value. If the possible outputs form a discrete set, this mapping may be achieved using a second layer of sigmoidal functions. To generate a continuous output, however, it is preferable to use the linear combination $\sum W_i^o n_i(\mathbf{x})$ for an appropriate set of coefficients W_i^o (known as a linear artificial neuron). The artificial neurons directly providing output values are said to form the output layer of the network. Each layer of artificial neurons before the output layer is known as a hidden layer. Mathematically, a feed-forward network with two hidden layers may be expressed as

$$\begin{aligned} \mathbf{n}_1 &= \mathbf{f}^{h_1}(\mathbf{W}^{h_1} \mathbf{x}) & \mathbf{n}_2 &= \mathbf{f}^{h_2}(\mathbf{W}^{h_2} \mathbf{n}_1) \\ \mathbf{o} &= \mathbf{f}^o(\mathbf{W}^o \mathbf{n}_2), \quad \text{where } \mathbf{f}^l: (s_i) \rightarrow [f_i^l(s_i)]. \end{aligned} \tag{3}$$

Here \mathbf{o} is the vector of output values, \mathbf{n}_l is the vector of outputs from hidden layer l , \mathbf{x} is the vector of input values, \mathbf{W}^{h_1} , \mathbf{W}^{h_2} , and \mathbf{W}^o are matrices whose rows contain the normal vectors \mathbf{w} for each neuron in the hidden and output layers, respectively, and $f_i^l(s)$ is the function describing the behavior of the i th neuron in layer l . This formulation may be trivially extended to any number of hidden layers.

b. An artificial neural network for precipitation estimation

The configuration of the MFFN constructed for this study is shown in Fig. 1. Forty-five input values (defined below) define the essential characteristics of the GOES image data. These input values are scaled to a range approximating $[0, 1]$ and connected to each of 200 sigmoidal artificial neurons in the first hidden layer. The first hidden layer is connected to a further 100-neuron hidden layer, and precipitation estimates are formed from a linear combination of the outputs from this second layer. Both zero and negative outputs from the final linear combination are taken to indicate zero precipitation; a decision that provides considerable advantages from the point of view of the network training process.

Satellite precipitation estimation poses a considerable challenge to a simple neural network approach because of presence of many contradictory exemplars in the training data. These exemplars are cases for which similar inputs are associated with very different outputs and occur particularly around the edges of rain areas where

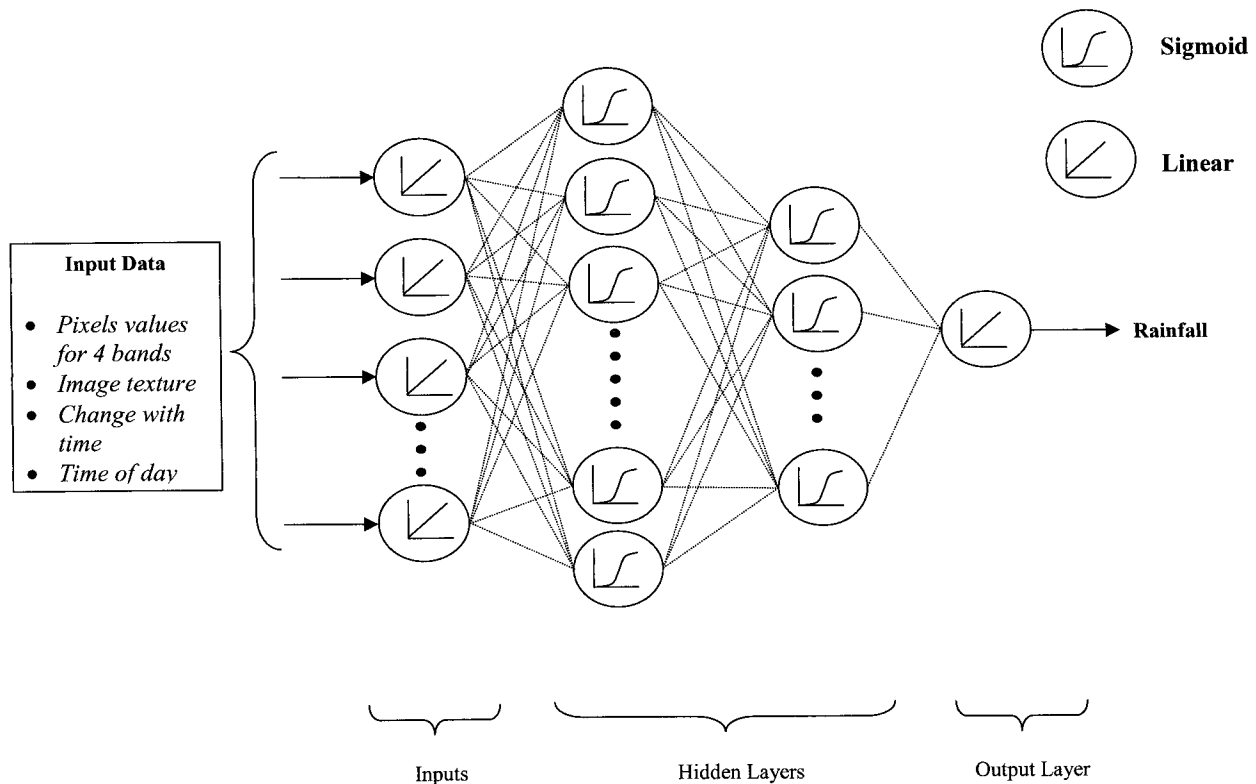


FIG. 1. The structure of the neural network used in this study.

cloud features may vary only marginally across the rain-no-rain boundary. In many cases these problems were found to prevent a single-layer feed-forward network from effectively training against the example data. The addition of a second hidden layer effectively overcame this problem by allowing sufficient additional complexity in network functionality to resolve the contradictions inherent in the training data. A two-layer network is an effective means of providing considerable additional complexity for a relatively modest additional computational cost and is generally most effective when the input data cluster within the hyperspace they define.

New-generation GOES data in the visible, shortwave IR, and thermal IR wavelengths can provide useful information on cloud characteristics related to rainfall processes. The relationship of the visible channel to cloud thickness and rainfall has long been established (Lovejoy and Austin 1979), and the 3.9- μm channel contains information on the state of cloud water (Scorer 1989) that, in conjunction with the 10.7- μm channel, can help to distinguish nonraining clouds (Vicente 1996). Radiances recorded in the 6.7- μm channel are strongly related to upper-level moisture. The longwave IR channel (10.7 μm) provides information on cloud-top temperature. The 12.0- μm channel was not available for this study, which is unfortunate because the split thermal IR window channels have been used to aid cirrus cloud identification (Inoue 1987) and for estimation of low-

level moisture (Chesters et al. 1987). Future work will include both thermal IR channels and can only be expected to improve rainfall estimation. Visible data are affected by the solar zenith angle and are unavailable at night; 3.9- μm data are partially affected by solar radiation and will change their meaning according to time of day and between day and night. To account for day-night differences, missing nighttime visible data were represented by a zero input. This input was sufficiently different from other valid visible pixel values to enable the neural network to separate day and night. Including local time as an additional input augmented this separation. In addition to assisting with day-night discrimination, the local-time input allowed the network to make a first-order allowance for solar zenith angle to the extent that this angle remained a constant function of local time throughout a training period.

Previous work has suggested that additional information on rainfall intensities can be obtained from the texture and rate of change of cloud-top temperatures (Griffith et al. 1978; Wu et al. 1985; Adler and Negri 1988), and this approach is a key component of operational rainfall estimation schemes (Scofield 1987). Convective cloud systems tend to produce greater rainfall during their developing stage rather than during the mature or decaying stages, such that rapidly decreasing cloud-top temperatures can be indicative of rapid vertical uplift and intense rainfall. Similarly, a "coarse"

TABLE 2. Inputs to the neural network.

Inputs	Description
1	Time of day
2–4	Brightness temperature for the pixel centered on the precipitation measurement for each of the three available IR bands
5	Visible albedo for the pixel centered on the precipitation measurement or zero if no visible image is available
6–8	Mean pixel value over a 3 by 3 pixel square for each available IR band
9	Mean pixel value over a 3 × 3 pixel square for the visible band, or zero if no visible image is available
10–12	Std dev of pixel values over a 3 by 3 pixel square for each available IR band
13	Std dev of pixel values over a 3 by 3 pixel for the visible band, or zero if no visible image is available
14–16	Mean pixel value over a 5 by 5 pixel square for each available IR band
17	Mean pixel value over a 5 by 5 pixel square for the visible band, or zero if no visible image is available
18–20	Std dev of pixel values over a 5 by 5 pixel square for each available IR band
21	Std dev of pixel values over a 5 by 5 pixel square for the visible band, or zero if no visible image is available
22	No. of IR band-4 pixels in a 3 by 3 pixel square with a brightness temperature less than 220 K
23	No. of IR band-4 pixels in a 5 by 5 pixel square with a brightness temperature less than 220 K
24–45	Inputs 2–23 calculated for the previous GOES image (one-half hour previous to the current image)

texture can indicate the presence of convective cloud “turrets” often associated with high rain rates. In this study, we included inputs relating to the local texture and rate of change in each of the four GOES channels. After, Hsu et al. (1999), information on cloud texture was incorporated by computing statistics (mean and variance) for 3 × 3 and 5 × 5 pixel rectangles surrounding the measurement location. Given that cloud-fraction information has been proven useful by the GPI, the number of 10.7- μm pixels colder than 220 K were also computed for these rectangles. A threshold of 220 K is more severe than the 235-K threshold used in the GPI. However, because this input is intended to provide supplemental information, not to derive the rain–no-rain boundary, it is reasonable to screen for somewhat colder cloud tops, more likely to be associated with convective storms.

A combination of pixel data and textural statistics from four channels and cloud fractions for two different rectangles gives rise to 22 input values. To incorporate temporal change into the scheme, these 22 values were computed both for the image corresponding to the precipitation event and for the previous image, transmitted 30 min earlier. Each precipitation measurement thus was derived from 45 different inputs, including local time (Table 2).

c. Training data

Multispectral satellite IR imagery was obtained every 30 min from the *GOES-8* satellite for a window covering central Amazonia for the two periods 8–31 January 1999 and 10–25 February 1999, inclusive. The original data had a spacing of 4 km and were remapped to a linear latitude–longitude projection with a grid spacing of 0.04°. These data were then combined with coincident TRMM precipitation radar data to form a training set for the neural network.

The TRMM PR operates at 13.8 GHz with a swath width of 215 km and horizontal and vertical resolutions of 4.3 km (at nadir) and 0.25 km (at nadir), respectively. The observable range within the atmosphere is from the surface to 15-km altitude (Kummerow et al. 1998). The PR provides an excellent source of training data. However, it is not without disadvantages and cannot be taken fully to represent ground truth, especially for the exact delineation of the edges of rain areas. Although an extensive validation program has been established, we currently are awaiting the full results of this program. Preliminary studies suggest that the PR is performing to expected standards (Oki et al. 1998). In a comparison with ground-based radar (rain gauges) the PR tended to underestimate the rainfall range (total), although distinguishing satellite and ground radar (gauge) errors proved to be problematic. There are a number of potential problems with PR estimates of rain rate that can explain some of the observed errors. First, the accuracy of the TRMM PR is very much dependent upon the ability to identify the correct cloud droplet distribution. The current processing measures rely upon a limited set of parameters that may not be optimal for all cloud microphysical conditions. Second, the design specification for the rain–no-rain boundary was 0.7 mm h⁻¹, but the operational use of the PR suggests that the signal-to-noise ratio is better than predicted, such that a lower rain–no-rain threshold (0.5 mm h⁻¹) can be achieved (Kummerow et al. 1998). Third, the effects of surface backscatter (exacerbated toward the edge of the PR swath) mean that the PR estimates of rainfall used in this study are for an atmospheric layer extending from 2000 to 4000 m. Under certain conditions, processes of evaporation and/or orographic enhancement may occur below this layer, which may introduce errors when compared with direct surface observations. Last, the radar beam can undergo attenuation as it passes through the rain system, and this attenuation may be severe where significant ice exists aloft, leading to an underestimation of lower-level rainfall.

The radar data are available at a number of stages of processing. For this study, level 2A25 PR profile data were chosen. At this stage, power and noise values have been converted via an apparent reflectivity to give an estimate of vertical rainfall rate profile for each beam at each resolution cell of the radar (TRMM Science Data and Information System 1999). For the training dataset,

the average rain rate between the heights of 2 and 4 km in the atmosphere was calculated for each 0.04°-resolution cell from the level-PR2A25 products. The 2–4-km average data were chosen to remove the influence of residual surface contamination. These data range in value from 0.0 to 3000 mm h⁻¹ and are associated with a rain flag that indicates rain or no-rain status and the rain type assumed in the rain-rate retrieval. For the purposes of this study, all potentially rainy pixels (both possible and certain) were included as nonzero rainfall estimates.

In theory, it should be possible to train a neural network to produce precipitation estimates at the spacing of the GOES image data. In practice, however, the relationship between image data and precipitation field at this resolution proved too noisy to allow an effective training of the neural network structures investigated in this study. For this reason, the remapped 0.04° PR precipitation data were reduced in effective spacing to 0.12° through the application of a three-by-three averaging filter. The filtered product was computed on the original 0.04° grid to avoid reducing the size of the training set. The network was thus trained to estimate precipitation over a 0.12° rectangular area centered on a given pixel location.

Because the relationship between rainfall and cloud patterns is known to vary over space and time with changing meteorological conditions, it is important to allow for this variation in the training process. One approach would be to incorporate location and date as additional inputs, augmenting the time-of-day parameter already present. This approach proves to be infeasible, however. With date, time, and location all present in the input dataset, there is a risk that the network will formulate a relationship between just these temporal and locational parameters and precipitation, to the exclusion of the satellite inputs. Such a relationship might be very efficient at reproducing the precipitation patterns present within the PR swaths but would not necessarily be at all effective at estimating precipitation beyond the swath coverage.

Given the undesirability of including date and location as inputs, the alternative solution is to restrict that spatial extent over which the training set is generated and periodically to retrain the network. The choice of training area and retraining period is a trade-off between the need to provide the widest possible variety of exemplars effectively to train the network and the need to capture the changing nature of the precipitation–cloud-cover relationship. For example, it was found that the network would very effectively train to datasets generated using single PR swaths, encapsulating the precipitation–cloudiness relationship for single storm systems to a very high degree of accuracy. However, the resulting functions proved to be not sufficiently general to generate effective precipitation estimates between overpasses.

Successful satellite precipitation methods such as

UAGPI recalibrate on at least a monthly basis. It thus was decided to train separately the network for each of the two periods of data availability: 8–31 January and 10–24 February. The training area was then selected to be as small as possible while still including a sufficient number of PR swaths to cover all parts of the diurnal cycle and thus to represent effectively the diurnal variation in the visible and 3.9-μm channels. The 3.1° square (80 by 80 0.04° pixels) training area selected to meet these criteria is shown as area A in Fig. 2.

Coincident PR and GOES image data were collected for PR swaths intersecting the training area during each training period. The training set for 8–31 January contained 9766 points from 11 individual swaths; the 10–25 February training data comprised 9177 points from 12 swaths.

d. Training the network

An MFFN may be trained using a simple gradient descent approach. Given an input vector \mathbf{x} and a known output vector \mathbf{y} , it is possible to define errors for the outputs \mathbf{o} from each layer of neurons in the network, with the absolute errors defined for output neurons *back-propagated* to hidden layer neurons according to the interconnection weightings defined by the coefficients \mathbf{W} :

$$\begin{aligned}\delta^o &= \mathbf{y} - \mathbf{o}, & \delta^{h_2} &= (\mathbf{W}^o)^T \delta^o, & \text{and} \\ \delta^{h_1} &= (\mathbf{W}^{h_2})^T \delta^{h_2}.\end{aligned}\quad (4)$$

Here δ^o denotes errors in the output of the output layer, δ^{h_2} denotes errors in the output of the second hidden layer, and δ^{h_1} denotes errors in the output of the first hidden layer. The \mathbf{W}^T denotes the transpose of \mathbf{W} , \mathbf{W}^o denotes the weights for the output layer, and \mathbf{W}^{h_2} denotes the weights for the second hidden layer. These absolute errors are used as the basis of a first-order estimate of the gradient of the instantaneous sum-square error $|\mathbf{y} - \mathbf{o}|$ with respect to the coefficient matrices \mathbf{W}^h and \mathbf{W}^o . These matrices are then updated by moving a short distance η along the estimated gradient vector to effect a small reduction in sum-square error:

$$\begin{aligned}\mathbf{W}^{o'} &= \mathbf{W}^o + \eta \delta^o \left(\frac{d\mathbf{f}^o(s)}{ds} \Big|_{s=\mathbf{W}^o \mathbf{n}_2} \right) \mathbf{n}_2^T \\ \mathbf{W}^{h_2'} &= \mathbf{W}^{h_2} + \eta \delta^{h_2} \left(\frac{d\mathbf{f}^{h_2}(s)}{ds} \Big|_{s=\mathbf{W}^{h_2} \mathbf{n}_1} \right) \mathbf{n}_1^T \\ \mathbf{W}^{h_1'} &= \mathbf{W}^{h_1} + \eta \delta^{h_1} \left(\frac{d\mathbf{f}^{h_1}(s)}{ds} \Big|_{s=\mathbf{W}^{h_1} \mathbf{x}} \right) \mathbf{x}^T.\end{aligned}\quad (5)$$

This procedure is repeated for every example mapping (\mathbf{x}, \mathbf{y}) in the training dataset, and the whole procedure is iterated until the sum-square error (or the incremental improvement in the same) reduces to an acceptable level.




FIG. 2. The geographical areas used for network training (A) and for validation (B). The smaller box within area A delineates the extent of Fig. 4.

As noted above, the function that links satellite imagery to precipitation is relatively complex and must be derived from noisy data. Training time may be greatly improved by setting the error δ^o associated with a particular example output to zero if that output is either zero or negative and the observed precipitation value is zero. This procedure removes a considerable restriction on the particular function that the network must represent. Only the positive parts of the function are specified. The parts of the function corresponding to zero rainfall may take on any negative value.

3. Results

a. Comparison with simple multispectral precipitation estimation methods

The satellite precipitation neural network was trained separately for the January and February data, respec-

tively. When training a complex neural network against noisy input data, it is essential to ensure that the non-linear function it represents has general applicability. This generality may be achieved by testing network outputs against a second set of exemplars not employed in the training process. For this purpose additional coincident TRMM and GOES input data were extracted for a smaller (40-by-80 pixel) training area adjacent to the primary training area (area B in Fig. 2). In addition to covering a different geographical area, these validation datasets contained a different set of PR swaths and thus represented a different temporal as well as spatial distribution of example data. Training was continued until the correlation of network products to PR data for the independent validation dataset (area B) ceased to improve by more than 1% in 1000 iterations. The network converged in 8000 iterations for the January training period and 13 000 iterations for February. In each case, the value of the training rate parameter η was initially set at 0.001 and was decreased to 0.0005 after 5000 iterations and to 0.0001 after a further 2000 iterations.

The performance of the two-layer neural network model was compared with two simpler models: linear regression of precipitation against all 45 input parameters and a simplified neural network containing 200 neurons in a single hidden layer. Table 3 lists both dependent correlations (for the “area-A” training datasets) and independent correlations (for the “area-B” validation datasets) for each of these models. The two-layer neural network outperforms linear regression for both training periods, indicating the successful identification of nonlinear functions in the training data. The single-layer network fares considerably less well and is sig-

TABLE 3. Correlations between model outputs and PR measurements computed for each of three models: linear regression, a simple neural network with a single hidden layer of 200 neurons, and the full two-layer neural network.

	8–31 Jan	10–25 Feb
Correlations with training data (area A)		
Linear regression	0.55	0.37
Single-layer ANN	0.51	0.54
Two-layer ANN	0.58	0.59
Correlations with validation data (area B)		
Linear regression	0.51	0.38
Single-layer ANN	0.51	0.40
Two-layer ANN	0.55	0.46




FIG. 3. Example hourly precipitation estimates for 1215 local time (LT) on 25 Jan 1999 from (left) PR (1219 LT overpass), (center) ANN, and (right) optimized GPI. The grayscale indicates rainfall rates in tenths of a millimeter. The grayscale has been truncated at 4 mm h^{-1} to enable detail to be seen in all three images. Some points in the PR image actually exceed 30 mm h^{-1} .

nificantly outperformed by the linear regression technique for the January training period.

b. Comparison with more-conventional satellite precipitation techniques

The neural network was used to produce gridded half-hourly precipitation estimates for the two periods 8–31 January 1999 and 10–25 February 1999 using the weights derived from training the network against the PR data for each respective period. To provide some measure of independent validation of ANN estimates against PR data, a second set of products was produced for 10–25 February using the weights derived from training over the longer period of 8–31 January.

To compare the performance of the ANN against a more conventional satellite precipitation technique, equivalent estimates were also produced for an optimized GPI modeled after the UAGPI (Xu et al. 1999). At present, cloud fraction is the most common indicator of rainfall rate used in IR satellite precipitation algorithms. By optimizing the temperature threshold and rainfall rate for the GPI over the study area for local conditions using the available PR data, it would be possible to assess the additional precipitation information present in multispectral data over and above cloud fraction. Separate temperature thresholds and conditional rain rates were derived for each training period using coincident 0.04° PR and GOES IR data. The brightness temperature threshold was chosen to produce the same number of raining pixels (as defined by a 0.5-mm h^{-1} cutoff) from the coincident dataset as did the PR. These pixels might not be the same ones—only the total number was optimized. Dividing the total PR-observed rainfall in the coincident dataset by the number of pixels colder than the threshold yielded a conditional rainfall rate (Xu et al. 1999). This process gave rise to temperature thresholds and conditional rain rates of 225 K and 3.1 mm h^{-1} for January and 228 K and 3.6 mm h^{-1} for February. The UAGPI typically derives parameters

over 1° squares. However, a PR swath is much narrower than an SSM/I swath and would leave 1° calibration areas with insufficient overpasses (none at all in many cases). Although mathematically the optimized GPI and UAGPI are identical, we refer to the optimized GPI as such because the calibration is conducted at a different scale than that specifically suggested for the UAGPI by Xu et al. (1999).

Figure 3 compares near-simultaneous instantaneous rainfall estimates from the PR, the ANN, and the optimized GPI for 25 January 1999. The ANN and optimized GPI estimates were produced from the 1215 local time (LT) GOES image (and its predecessor in the former case); the PR overpass occurred at 1219 LT. The higher-rainfall areas identified by the ANN bear some resemblance to those identified in the PR swath, especially on the southeast side of the rain system. However, the ANN estimates for these areas are much smaller in magnitude (37-mm h^{-1} peak intensity as opposed to 4 mm h^{-1}).

The ANN and optimized GPI products were compared with the available PR data at 0.12° spacing (Table 4). Five error statistics were computed for each training period: root-mean-square error (rmse) mean bias, correlation, correlation for PR-identified raining points, and skill. For precipitation algorithms, the skill score is defined as the combined fraction of correctly identified rain and correctly identified “no rain” expressed as a percentage of the total number of estimates made. Because of uncertainty in PR rainfall estimates at very low rain rates, the rain–no-rain threshold was set at 0.5 mm h^{-1} . Table 4 additionally lists error statistics for February products generated using models trained to the January dataset. The ANN demonstrates higher correlations with PR data than does the optimized GPI, although for the January training period it shows a higher bias and consequently a higher rmse. Possible reasons for this are discussed later in this section with the results from the validation against ground data. In the case of the validation using February data of January-trained

TABLE 4. Error statistics for instantaneous 0.12° ANN and optimized (Opt.) GPI precipitation estimates validated against dependent and independent PR data.

	Correlation	Correlation (rainfall > 0.5 mm)	Rmse (mm)	Bias (mm)	Skill
Jan (4187 points)					
ANN	0.53	0.27	1.57	-0.19	0.77
Opt. GPI	0.44	0.19	1.48	-0.002	0.89
Feb (3097 points)					
ANN	0.56	0.52	1.09	-0.004	0.91
Opt. GPI	0.33	0.25	1.30	-0.11	0.91
Feb trained against Jan data (3097 points)					
ANN	0.41	0.41	1.16	-0.03	0.86
Opt. GPI	0.34	0.29	1.24	-0.06	0.92

models, although the overall correlations are somewhat reduced, the advantages of the ANN method over the optimized GPI remain evident.

Table 5 lists validation statistics computed at three different spacings: 0.12° (as in Table 4), 0.48°, and 1.0°, estimates for the lower spatial resolutions being created by averaging the high-resolution 0.12° product. The ANN method produces consistently better correlations than does the optimized GPI at all three spatial scales.

A second validation exercise was conducted using independent rain gauge data, obtained from a subregion in the central zone of the study window. The gauge network was established as part of the TRMM-Large-Scale Biosphere-Atmosphere (LBA) Experiment in Amazonia and provides a high density of gauges over a number of small regions, each approximating the size

of a single satellite IR or PR footprint (Rutledge 1999). In total, 40 gauges are distributed over the subregion (Fig. 4) covering 13 individual cells in the 0.04° grid. The gauges are of a tipping-bucket design and provide estimates of rain rate with a maximum resolution of 1 min. These data were used for validation of instantaneous rain rate and have been accumulated over 24-h periods for validation of satellite daily rainfall estimates. Observations were averaged over all 13 0.04° cells, providing a quantity effectively averaged over a single grid cell with a spacing of 0.14°. When comparing validation results at this scale with previously published statistics, it is instructive to note that this area approximates that of a single SSM/I 85-GHz footprint. Instantaneous estimates of rainfall rate were derived from gauge data by averaging over a 15-min period either side of the satellite imaging time and then averaging over all gauges. Because the gauge data did not form a contiguous 0.12° square, it was not possible to compare directly an average of gauge values with a single ANN output. Separate ANN outputs were computed for the 0.12° squares surrounding each of the 13 0.04° grid squares containing at least one rain gauge. These values were then averaged to supply an estimate for the mean areal precipitation for the area covered by the gauges. This process would be expected to introduce some loss in the variability of the ANN precipitation estimates. Optimized GPI outputs were available at 0.04° spacing and were simply averaged over the 13 0.04° cells containing rain gauges.

The results of the gauge validation are presented in Table 6. For instantaneous estimates produced from 30-min imagery, ANN outperforms the optimized GPI substantially in terms of both correlation and correlation while raining. The optimized GPI shows a slightly higher skill at identifying the rain-no-rain boundary. This result is unsurprising, because the optimized GPI brightness temperature threshold was selected to optimize this discrimination. Rain-no-rain delineation is a problematic process, because a very small change in estimated precipitation may move a considerable number of measurements from one category to the other.

The rmse and biases for both techniques are domi-

TABLE 5. Error statistics for the ANN and optimized (Opt.) GPI techniques compared at a range of spatial scales.

	Correlation	Rmse (mm)	Bias (mm)
0.12°			
Jan (3126 data points)			
ANN	0.53	1.57	-0.19
Opt. GPI	0.44	1.48	-0.002
Feb (2710 data points)			
ANN	0.56	1.09	-0.004
Opt. GPI	0.33	1.30	-0.11
0.48°			
Jan (157 data points)			
ANN	0.75	1.00	-0.30
Opt. GPI	0.69	0.93	-0.10
Feb (156 data points)			
ANN	0.80	0.58	-0.06
Opt. GPI	0.53	0.80	-0.07
1°			
Jan (39 data points)			
ANN	0.89	0.60	-0.24
Opt. GPI	0.83	0.60	-0.03
Feb (40 points)			
ANN	0.93	0.17	-0.06
Opt. GPI	0.81	0.35	-0.07




FIG. 4. TRMM-LBA rain gauge network. Figures shown are the number of gauges in each 0.04° cell.

TABLE 6. Error statistics for precipitation estimates produced using the ANN and optimized (Opt.) GPI method by comparison with rain gauge data.

	Correlation	Correlation (rainfall > 0.5 mm)	Rmse (mm)	Bias (mm)	Skill
30 min.					
Jan					
ANN	0.46	0.46	1.56	-0.07	0.81
Opt. GPI	0.39	0.28	1.66	-0.25	0.86
Feb					
ANN	0.48	0.48	1.69	-0.37	0.84
Opt. GPI	0.41	0.33	1.62	-0.19	0.85
1 h					
Jan					
ANN	0.54	0.58	1.30	-0.04	0.81
Opt. GPI	0.46	0.41	1.35	-0.24	0.84
Feb					
ANN	0.60	0.59	1.36	0.33	0.86
Opt. GPI	0.46	0.29	1.38	0.18	0.86
3 h					
Jan					
ANN	0.62	0.55	2.62	-0.08	0.70
Opt. GPI	0.59	0.56	2.79	-0.61	0.78
Feb					
ANN	0.66	0.63	2.73	-0.81	0.77
Opt. GPI	0.61	0.45	2.52	-0.45	0.84
1 day					
Jan					
ANN	0.82	0.72	5.28	-0.12	0.91
Opt. GPI	0.84	0.78	7.04	-3.21	0.83
Feb					
ANN	0.94	0.93	6.85	-3.65	0.96
Opt. GPI	0.85	0.78	6.97	-2.50	0.96




FIG. 5. A comparison of instantaneous rainfall estimates from the satellite neural network algorithm and rain gauge data for the combined period 8–31 Jan and 10–25 Feb 1999. Estimates produced were averaged over all 13 grid cells that contain gauge data. The line of equality is also included for clarity.

nated by high rainfall events in the data. Although neither technique produces high values for rainfall, the presence of such events in the training data will increase the optimized GPI conditional rainfall rate while having relatively little effect on the trained neural network. This effect is particularly noticeable in February, when the negative bias for the ANN technique is considerably higher than the optimized GPI; this high bias pushes the rmse for the former above that for the latter, even though the ANN displays a higher correlation with the gauge data. It is, of course, questionable to what extent higher rainfall rates from even a relatively dense surface network truly represent mean areal precipitation over a large area. It is notable that 33% of the total rainfall observed in the gauge data was supplied by events with rainfall intensities over 10 mm h^{-1} , but the equivalent statistic for the PR training data was only 20%. The bias of the ANN technique with respect to gauge data does not seem to be related directly to the representation of high rainfall events in the training data: 16% of rainfall in the January training dataset was provided by events in excess of 10 mm h^{-1} and 29% was provided by such events in the February training data. The fraction of higher rainfall events in the gauge data was the same (33%) for both periods.

Table 6 also supplies statistics for 1-h, 3-h, and daily time steps. As expected, temporal aggregation improves the results from both techniques and reduces the relative

advantage of ANN over an optimized GPI, until for daily estimates it is unclear which is the better technique. If an ANN technique were to be used to produce data on a daily time step, it might be better specifically to train the neural network using data with that temporal resolution.

Figure 5 plots instantaneous precipitation estimates against gauge data. This plot clearly shows the difference in the relationship between ANN estimates and gauge measurements at low and high rainfall rates. The data points corresponding to lower intensity rainfall events are reasonably uniformly scattered about the 1:1 line; higher rainfall events are scattered much more randomly. Events with intensities between 5 and 20 mm h^{-1} may be particularly difficult for ANN to characterize. However, the three highest rainfall events are characterized as being associated with above-average rainfall.

Figures 6 and 7 plot time series of 30-min precipitation estimates for the two training periods. Although both methods clearly recognize high rainfall events to some degree, they underestimate their magnitude. A few high rainfall events (especially those occurring at the start of 28 January and at the start of 23 February) were missed entirely by both techniques. The ANN technique may be more responsive to low rainfall, tracking a number of low-intensity events that are missed entirely by the optimized GPI.




FIG. 6. Time series of instantaneous rainfall for 8–31 Jan 1999 produced by ANN (upper) and the optimized GPI (lower).

4. Conclusions

A combined TRMM-PR-*GOES-8* precipitation estimation algorithm has been developed and validated against both independent PR data and gauge data. Rain rates were estimated at very high spatial and temporal resolutions that have historically posed the greatest problems for satellite precipitation methods. Instantaneous precipitation estimates at 0.12° resolution demonstrated correlations of ~ 0.4 – 0.5 with respect to independent validation data. The ANN technique performed consistently better than an optimized GPI in terms of correlations with validation data. This result indicates that multispectral geostationary satellite imagery contains information pertaining to low-to-moderate rainfall rates over and above that provided by simple cloud fraction. The validation statistics compare very well with previous studies conducted at similarly small scales and demonstrate the potential of this method for applications where such high-resolution estimates are required.

The ANN method shared with the optimized GPI problems with identifying high rain rates. This problem could be due to the lack of example high rainfall rates in training data but seems at least in part to be due to a lack of characteristic information pertaining to high rainfall events in the image data itself. ANN did not

directly account for missing high rainfall by adjusting rainfall estimates across the board, in the manner of algorithms such as UAGPI. This lack of direct accounting meant that any improvement in bias and rmse statistics for ANN over the optimized GPI was either marginal or favored the optimized GPI. For applications (such as operational hydrology) that require a zero long-term bias, a simple multiplicative correction factor could be derived from the training dataset and uniformly applied to the precipitation estimates. This approach would reduce long-term bias but would not, of course, overcome the nonlinear underestimation of high rainfall events.

Unlike ground-based radar and dense gauge networks, TRMM PR data are routinely available over the global Tropics and subtropics, enabling the potential application of the ANN technique throughout this region. In addition, the ANN method incorporates multispectral VIS and IR satellite information, in contrast to most existing empirical algorithms that combine PMW and IR data. Thus, the method makes full use of the range of spectral channels on board of recent and future geostationary satellite systems. If real-time estimates were required, a rolling calibration window could be employed (e.g., the last eight days including the current day.) A further improvement would be to




FIG. 7. Time series of instantaneous precipitation for 10–25 Feb 1999 produced by ANN (upper) and optimized GPI (lower).

incorporate data from GOES channel IR-5. These data form part of an IR split window that could be very helpful in discriminating cloud thickness.

Note that MFFN is not the only network structure that may be used for this kind of problem. For example, Hsu et al. (1999) used a modified version of the counter-propagation network. Network structures differ most significantly in training time. Given the complexity of the network design employed in this study, the choice of an optimum network structure to represent image–rainfall relationships best warrants further investigation.

The authors are currently undertaking further refinement of the neural network approach and more extensive validation of estimates. We are also testing the method in a range of contrasting climatic zones.

Acknowledgments. The authors thank Dr. C. Kummerow of the NASA Goddard Space Flight Center for advice, M. Bellerby of P&L Systems International, Ltd., for answering numerous questions on neural network design and implementation, and anonymous referees for helpful suggestions. TRMM data for this project were obtained from the NASA Tropical Rainfall Measuring Mission Science Data and Information System (TSDIS). GOES and gauge data were obtained from the NASA TRMM–LBA Field Campaign.

REFERENCES

- Adler, R. F., and A. J. Negri, 1988: A satellite infrared technique to estimate tropical convective and stratiform rainfall. *J. Appl. Meteor.*, **27**, 30–38.
- , —, P. R. Keehn, and I. M. Hakkarinen, 1993: Estimation of monthly rainfall over Japan and surrounding waters from a combination of low-orbit microwave and geosynchronous IR data. *J. Appl. Meteor.*, **32**, 335–356.
- Arkin, P. A., and B. N. Meisner, 1987: The relationship between large-scale convective rainfall and cold cloud over the Western Hemisphere during 1982–84. *Mon. Wea. Rev.*, **115**, 51–74.
- Barrett, E. C., and M. J. Beaumont, 1994: Satellite rainfall monitoring: An overview. *Remote Sens. Rev.*, **11** (1–4), 49–60.
- Bellerby, T. J., and E. C. Barrett, 1993: Progressive refinement: A strategy for the calibration by collateral data of short-period satellite rainfall estimates. *J. Appl. Meteor.*, **32**, 1365–1378.
- Chesters, D., W. D. Robinson, and L. W. Uccellini, 1987: Optimized retrievals of precipitable water from the VAS split window. *J. Climate Appl. Meteor.*, **26**, 1059–1066.
- Freeman, J. A., and D. M. Skapura, 1991: *Neural Networks: Algorithms, Applications, and Programming Techniques*. Addison-Wesley, 401 pp.
- Griffith, C. G., W. L. Woodley, P. G. Gruber, D. W. Martin, J. Stout, and D. N. Sikdar, 1978: Rain estimation from geosynchronous satellite data—visible and infrared studies. *Mon. Wea. Rev.*, **106**, 1153–1171.
- Herman, A., V. B. Kumar, P. A. Arkin, and J. V. Kousky, 1997: Objectively determined 10-day African rainfall estimates created for famine early warning systems. *Int. J. Remote Sens.*, **18**, 2147–2159.

- Hsu, K., H. Gupta, X. Gao, and S. Sorooshian, 1999: Estimation of physical variables from multichannel remotely sensed imagery using a neural network: Application to rainfall estimation. *Water Resour. Res.*, **35**, 1605–1618.
- Inoue, T., 1987: An instantaneous delineation of convective rainfall areas using split window data of NOAA-7 AVHRR. *J. Meteor. Soc. Japan.*, **65**, 469–481.
- Kummerow, C., W. Barnes, T. Kozu, J. Shiue, and J. Simpson, 1998: The Tropical Rainfall Measuring Mission (TRMM) sensor package. *J. Atmos. Oceanic Technol.*, **15**, 809–817.
- Li, Q., R. L. Bras, and D. Veneziano, 1996: Analysis of Darwin rainfall data: Implications on sampling strategy. *J. Appl. Meteor.*, **35**, 372–385.
- Lovejoy, S., and G. L. Austin, 1979: The delineation of rain areas from visible and infrared satellite data for GATE and mid-latitudes. *Atmos.–Ocean.*, **17**, 1048–1054.
- Morrissey, M. L., and J. E. Janowiak, 1996: Sampling-induced conditional biases in satellite climate-scale rainfall estimates. *J. Appl. Meteor.*, **35**, 541–548.
- Oki, R., K. Furukawa, S. Shimizu, Y. Suzuki, S. Satoh, H. Hanado, K. Okamoto, and K. Nakamura, 1998: Preliminary results of TRMM: Part I, a comparison of PR with ground observations. *Marine Tech. Soc. J.*, **32**, 13–23.
- Rutledge, S. A., cited 1999: A science overview for the NSF component of TRMM-LBA/Brazil. [Available online at <http://cloud1.arc.nasa.gov/trmmlba/overview.html>.]
- Scofield, R. A., 1987: The NESDIS operational convective precipitation estimation technique. *Mon. Wea. Rev.*, **115**, 1773–1792.
- Scorer, R. S., 1989: Cloud reflectance variations in AVHRR channel-3. *Int. J. Remote Sens.*, **10**, 675–686.
- Skapura, D. M., 1996: *Building Neural Networks*. ACM Press, 282 pp.
- Soman, V. V., J. B. Valdes, and G. North, 1995: Satellite sampling and the diurnal cycle statistics of Darwin rainfall data. *J. Appl. Meteor.*, **34**, 2481–2490.
- Todd, M. C., and D. R. Kniveton, 1999: A combined satellite infrared and passive microwave technique for estimation of small-scale rainfall over the global Tropics and subtropics. *Proc. Int. Union of Geophysics and Geodesy 99*, Birmingham, United Kingdom, IUGG, B.285.
- , E. C. Barrett, M. J. Beaumont, and J. Green, 1995: Satellite identification of rain days over the upper Nile River basin using an optimum infrared rain/no-rain threshold temperature model. *J. Appl. Meteor.*, **34**, 2600–2611.
- , —, —, and T. J. Bellerby, 1999: Estimation of daily rainfall over the upper Nile River basin using a continuously calibrated satellite infrared technique. *Meteor. Appl.*, **6**, 201–210.
- TRMM Science Data and Information System, cited 1999: Volume 4: File specifications for TRMM products—level 2 and level 3. NASA Goddard Space Flight Center. [Available online at <http://tsdis.gsfc.nasa.gov/tsdis/tsdis.html>.]
- Vincente, G. A., 1996: Algorithm for rainfall rate estimation using a combination of GOES-8 11 and 3.9 measurements. Preprints, *Eighth Conf. on Satellite Meteorology and Oceanography*, Atlanta, GA, Amer. Meteor. Soc., 274–278.
- Wu, R., J. A. Weinmann, and R. T. Chin, 1985: Determination of rainfall rates from GOES satellite images by a pattern recognition technique. *J. Atmos. Oceanic Technol.*, **2**, 314–330.
- Xu, L., X. Gao, S. Sorooshian, P. A. Arkin, and B. Imam, 1999: A microwave infrared threshold technique to improve the GOES Precipitation Index. *J. Appl. Meteor.*, **38**, 569–579.