A Hierarchical Bayes Model for Combining Precipitation Measurements from Different Sources

Tamre P. Cardoso and Peter Guttorp

NRCSE

Technical Report Series

NRCSE-TRS No. 082

March 1, 2007

A Hierarchical Bayes Model for Combining Precipitation Measurements from Different Sources

Tamre P. Cardoso and Peter Guttorp

ABSTRACT

Surface rain rate is an important climatic variable and many entities are interested in obtaining accurate rain rate estimates. Rain rate, however, cannot be measured directly by currently available instrumentation. A hierarchical Bayes model is used as the framework for estimating rain rate parameters through time, conditional on observations from multiple instruments such as rain gauges, ground radars, and distrometers. The hierarchical model incorporates relationships between physical rainfall processes and collected data. A key feature of this model is the evolution of drop-size distributions (DSD) as a hidden process. An unobserved DSD is modeled as two independent component processes: 1) an AR(1) time-varying mean with GARCH errors for the total number of drops evolving through time, and 2) a time-varying lognormal distribution for the size of drops. From the modeled DSDs, precipitation parameters of interest, including rain rate, are calculated along with associated uncertainty. This model formulation deviates from the common notion of rain gauges as "ground truth"; rather, information from the various precipitation measurements is incorporated into the parameter estimates and the estimate of the hidden process. The model is implemented using Markov chain Monte Carlo methods.

INTRODUCTION

Surface rainfall is an important environmental variable that is incorporated across many areas of study, including meteorology, climatology, agriculture, land use, and hydrology, for example. The various fields of study require estimation of precipitation at a range of temporal and spatial scales. Hydrologists and land use planners may be interested in short-term rainfall in relatively small regional areas for studies involving flood and flash flood forecasting. Researchers in climatology or agriculture may be interested in climatic studies that focus on weekly, monthly, or annual totals over large spatial extents.

Although of considerable interest, surface precipitation rates and amounts are difficult to estimate. Unlike other atmospheric variables, rainfall can display extreme heterogeneity in space and time. Both the occurrence of precipitation and the rate at which it falls may be highly variable, even within a single rain event. As such, due to the importance of and difficultly with estimation, there has been much research in the area of precipitation measurement and estimation.

Precipitation is measured using many different instruments. Some instrumentation measures rainfall directly, while others make indirect measurements of quantities that can be related to rainfall. All of the instrumentation, however, make indirect measurements of the usual

quantity of interest, namely, rain rate. The most common instruments include rain gauges and ground-based scanning radar. Distrometer *in situ* measurements and satellite deployed instrumentation are also used, although less commonly. Each of the instruments has inherent strengths and weaknesses, which affect the observed measurements and subsequent estimation of surface rainfall.

There exist many algorithms, each with its own set of assumptions, for calculating rainfall related parameters from the instrumental observations. Many times the observations obtained from one instrument are used in the calculation of estimates based on the observations from another instrument. Rain gauge data are often considered to be "ground truth" and are used to adjust the estimates based on other instruments. This is done despite the known fact that observations from rain gauges are inherently prone to errors and that they are no more "true" than the observations from other instruments.

Surface precipitation can be defined by a population of falling drops. The distribution of the size diameters of the drops characterizes the population behavior; thus, the drop-size distribution (DSD) forms a basic descriptor in the modeling of rain microphysics. The DSD is used to compute a variety of "derived" properties of rainfall through mathematical relationships; common computed properties include water content, rain rate, and reflectivity. DSD data, however, are not collected on a routine basis due to the expense of distrometer instruments and the limited spatial coverage of an individual instrument. Rain gauge measurements and ground radar images are much more readily available.

The literature is full of modeling approaches for the estimation of surface rainfall, including empirical statistical models (Stern and Coe 1984; Skaugen et al. 1996, for example) and a variety of stochastic models (Bell 1987; Rodriguez-Iturbe and Eagleson 1987; Cox et al. 1988; Rodriquez-Iturbe et al. 1988; Smith 1993, for example). More recently, hierarchical Bayesian models have been used in a number of environmental applications (Royle et al. 1998, Wikle 1998, Berliner 2000, Berliner et al. 2000a, Wikle et al. 2001, Hrafnkelsson 2003). In this article we present a generalized model for surface rainfall that accommodates data from multiple instruments and produces estimates of precipitation parameters and associated uncertainties. Using a hierarchical Bayesian approach, the evolution of DSDs are modeled as an unobserved process, while also incorporating information from data gathered through commonly deployed instruments. This formulation does not depend on having expensive DSD data and departs from

the notion of gauge data as being "ground truth". Rather, an estimate of the drop-size distribution at each point in time, with attendant uncertainty measures, can be used to calibrate all the instruments simultaneously.

HIERARCHICAL MODEL

Surface rainfall parameters through time are estimated using a flexible five-stage hierarchical Bayesian formulation that is based on a model for generating rainfall drop-size distributions. Estimates of the unobserved process, denoted by $N(D)$, are the primary quantity of interest. The first stage of the hierarchy specifies a measurement error model for the observational data, denoted by Z and G, both of which are observations of functions of $N(D)$, with error. The second stage of the hierarchy allows for a time series formulation of the unobserved DSD process. Time series parameters and temporal dynamic terms of the DSD evolution process are modeled in the third stage. The hierarchical Bayes formulation is completed in stages four and five by specifying priors on model parameters.

Empirical Analysis of DSD Data

Exploratory data analysis was conducted on a set of one-minute drop size spectra collected using a Joss-Waldvogel distrometer (JWD) from a site in Eureka, California from January through March 1999. Rain rates were calculated for each of the one-minute spectra to identify rainfall events. Seven events ranging from 12 to 24 hours in duration were chosen for evaluation. Barplots of the binned drops and time series of the total number of drops/minute were evaluated for each event. Two general observations were made:

1. The distributions of drop-sizes were generally unimodal and right-skewed. Based on visual inspection and five number summaries, the distributions of the shapes of the one-minute DSD data were similar both within and across rainfall events. Figure 1 shows barplots for some one-minute DSDs.
2. The total number of drops from one-minute to the next was highly variable both within and across events. Log-transformed time series data suggested an underlying auto-regressive process. Figure 2 shows some time-series plots for total number of drops.

Unobserved Process

A DSD, denoting the number of drops per mm diameter bin interval and per m^3 of air is well defined as

$$N(D_i)_t = \frac{n(D_i)_t}{AtV(D_i)\Delta D_i} \tag{1}$$

where $n(D_i)_t$ is the observed number of drops (raw counts) in each diameter bin interval D_i for sampling interval t; A is the sample surface area in m^2; t is the length of sample collection in seconds; $V(D_i)$ is the terminal fall speed in still air in m/s; and, ΔD is the drop diameter interval in mm.

Now we define the number of drops over diameters D as

$$n(D)_t = n_t \times P(D)_t \tag{2}$$

where n_t is the total number of raw counts over all drop diameters for the time interval t and is assumed independent of $P(D)_t$, a probability density function defining the shape of the DSD over the interval t; $P(D)_t$ gives the probability of finding a drop within a diameter D to $D + dD$.

Substituting equation 2 into equation 1 yields

$$N(D_i)_t = \frac{n_t \times P(D_i)_t}{AtV(D_i)\Delta D_i} \tag{3}$$

Next, define N_t as the total number of drops per m^3 of air. Substitution into equation 3 gives the DSD formulated in terms of total counts as

$$N(D_i)_t = \frac{N_t \times P(D_i)_t}{\Delta D_i} \tag{4}$$

For a rainfall event of duration T, the unobserved states $N(D)_t$ are estimated through models for N_t and $P(D)_t$; $t = 1,\ldots,T$. To be consistent with commonly collected distrometer data, each t represents a one-minute interval.

Model for Total Number of Drops, N_t

Let N_t denote the total number of drops at time t minutes, $t = 1,2,\ldots,T$. To account for the serial correlation and burst variability observed in many of the time series of total number of drops in contiguous DSD spectra, we introduce a model composed of a time-varying mean process integrated with a *GARCH*(1,1) conditional variance process. Let $TN_t = \ln N_t = \mu_{N_t} + \varepsilon_t$.

The ε_t's are an independent sequence of correlated errors from an unknown distribution. The ε_t values, conditional on all previous values are assumed to be normally distributed: $\varepsilon_t \left| \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots \sim N\left(0, \sigma_{\varepsilon_t}^2\right)\right.$. Under a $GARCH(1,1)$ model $\sigma_{\varepsilon_t}^2 = a_0 + a_1 \varepsilon_{t-1}^2 + b_1 \sigma_{\varepsilon_{t-1}}^2$ where the parameters $a_0, a_1,$ and b_1 are subject to the constraints $a_0 > 0, a_1, b_1 \geq 0,\ a_1 < 1,$ and $a_1 + b_1 \leq 1$ (Alexander 1998). By substitution,

$$\begin{aligned}
&\varepsilon_1 \left| \varepsilon_0 \sim N\left(0, a_0 + a_1 \varepsilon_0^2 + b_1 \sigma_{\varepsilon_0}^2\right)\right. \\
&\varepsilon_2 \left| \varepsilon_1, \varepsilon_0 \sim N\left(0, a_0 + a_1 \varepsilon_1^2 + b_1 \sigma_{\varepsilon_1}^2\right)\right. \\
&\vdots \\
&\varepsilon_t \left| \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots, \varepsilon_1, \varepsilon_0 \sim N\left(0, a_0 + a_1 \varepsilon_{t-1}^2 + b_1 \sigma_{\varepsilon_{t-1}}^2\right)\right.
\end{aligned}$$

Given the conditional distribution of $\varepsilon_t \left| \varepsilon_{t-1}, \varepsilon_{t-2}, \ldots, \varepsilon_1, \varepsilon_0\right.$, the joint probability for ε is:

$$\begin{aligned}
\left[\varepsilon_t, \varepsilon_{t-1}, \ldots, \varepsilon_0\right] &= \left[\varepsilon_t \left| \varepsilon_{t-1}, \varepsilon_{t-2} \ldots, \varepsilon_0\right.\right]\left[\varepsilon_{t-1} \left| \varepsilon_{t-2}, \varepsilon_{t-3} \ldots, \varepsilon_0\right.\right] \ldots \left[\varepsilon_0\right] \\
&= \left[\prod_{i=1}^{t} \left[\varepsilon_{t+1-i} \left| \varepsilon_{t-i}, \varepsilon_{t-i-1}, \varepsilon_{t-i-2}, \ldots, \varepsilon_0\right.\right]\right]\left[\varepsilon_0\right] \\
&= \left[\prod_{i=1}^{t} \frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_{t+1-i}}} \exp\left\{-\frac{1}{2\sigma_{\varepsilon_{t+1-i}}^2}\left(\varepsilon_{t+1-i}^2\right)\right\}\right]\frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\varepsilon_0^2\right)\right\}
\end{aligned}$$

Assume that conditional on μ_{N_t} and the parameters $\theta_N = \left\{a_0, a_1, b_1, \sigma_{\varepsilon_0}^2\right\}$, the $\ln N_t$'s are independent and are distributed as $N\left(\mu_{N_t}, \sigma_{\varepsilon_t}^2\right)$. Then the joint distribution of $\ln N$ is:

$$\begin{aligned}
\left[TN\right] &= \left[TN_1, \ldots, TN_T \left| \mu_{N_1}, \ldots, \mu_{N_T}; \sigma_{\varepsilon_1}^2, \ldots, \sigma_{\varepsilon_T}^2\right.\right] \\
&= \left[\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_t}} \exp\left\{-\frac{1}{2\sigma_{\varepsilon_t}^2}\left(TN_t - \mu_{N_t}\right)^2\right\}\right] \times \frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_0}} \exp\left\{-\frac{1}{2\sigma_{\varepsilon_0}^2}\left(TN_0 - \mu_{N_0}\right)^2\right\}
\end{aligned}$$

Set $TN_0 = \mu_{N_0}$. This reduces $P\left(\varepsilon_0\right)$ to the constant term $\frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_0}}$ and yields:

$$\left[\prod_{t=1}^{T} \frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_t}} \exp\left\{-\frac{1}{2\sigma_{\varepsilon_t}^2}\left(TN_t - \mu_{N_t}\right)^2\right\}\right] \times \frac{1}{\sqrt{2\pi}\sigma_{\varepsilon_0}}$$

Priors on GARCH parameters

Since the *GARCH* parameters are constrained to be greater than zero, each coefficient is assigned a lognormal prior:

$$a_0 : \text{Lognormal}\left(\alpha_0, \tau_{\alpha_0}^2\right)$$
$$a_1 : \text{Lognormal}\left(\alpha_1, \tau_{\alpha_1}^2\right)$$
$$b_1 : \text{Lognormal}\left(\beta_1, \tau_{\beta_1}^2\right)$$

The hyperparameters $\alpha_0, \alpha_1, \beta_1, \tau_{\alpha_0}^2, \tau_{\alpha_1}^2$, and $\tau_{\beta_1}^2$ are taken as fixed, with values that may vary by application.

Mean process for number of drops

The mean for TN_t, μ_{N_t}, is modeled as a first-order autoregressive process. Let $\mu_{N_t} = a_N \mu_{N_{t-1}} + \eta_{N_t}$ where η_{N_t} is an independent random normal process with mean 0 and variance σ_η^2. The parameters are $\theta_\mu = \left\{a_N, \sigma_\eta^2\right\}$. Conditional on a_N and $\mu_{N_{t-1}}$,

$$\left[\mu_{N_t} \middle| a_N \mu_{N_{t-1}}\right] \sim N\left(a_N \mu_{N_{t-1}}, \sigma_\eta^2\right); \;\; \sigma_\eta^2 \sim IG\left(\mu_\eta, \tau_\eta^2\right)$$

The parameter a_N is estimated assuming $a_N \sim N\left(\alpha_N, \tau_{\alpha_N}^2\right)$. The hyperparameters $\alpha_N, \tau_{\alpha_N}^2, \mu_\eta$, and τ_η^2 are taken as fixed, with values that may vary by application.

DSD shape, $P(D)_t$

The DSD shape is modeled as a lognormal distribution with time varying parameters ξ_t and ϕ_t^2. Conditional on ξ_t and ϕ_t^2, the D_t's are assumed independent yielding the joint distribution for *D* as

$$\left[D \middle| \xi, \phi^2\right] = \prod_{t=1}^{T} \frac{1}{\sqrt{2\pi} d_t \phi_t} \exp\left\{-\frac{\left(\ln d_t - \xi_t\right)^2}{2\phi_t^2}\right\}$$

where $\xi_t \sim N\left(\mu_\xi, \sigma_\xi^2\right); \;\; \phi_t^2 \sim IG\left(\alpha_\phi, \beta_\phi\right)$. The hyperparameters $\mu_\xi, \sigma_\xi^2, \alpha_\phi$, and β_ϕ are taken as fixed, with values that may vary by application.

Data Components

The hierarchical model incorporates two common data sources, namely, rain gauge and ground radar observations. Since given a DSD the instantaneous (true) rain rate and a derived (true) radar reflectivity can be calculated, we develop a measurement error model for both the rain gauge and radar observations.

Gauge Observations, G

Rain gauges measure the amount of rainfall accumulating over time at a fixed location in space. If the amount of rainfall measured at a point is measured without error, the gauge measurement, *G*, can be equated to the appropriate time integral of rain rate. The predominant error in gauge measurements is a systematic bias induced by winds, which often results in an underestimate of the surface rainfall. Thus, define a measurement error model for gauges as a function of the true rain rate and wind speed (taking other systematic biases as negligible). Define a gauge observation as

$$G_t = \left\{ c\left(w_t\right) \int_{t-\Delta t}^{t} R_s ds \right\} + \varepsilon_{G_t}$$

where R_t is the derived instantaneous rain rate at time *t*; $c\left(w_t\right)$ are gauge type-specific coefficients ($c \le 1$) primarily based on wind speed, *w*; and ε_{G_t} are independent measurement errors modeled as $N\left(0, \sigma_G^2\right)$; $\sigma_G^2 : IG\left(\mu_G, \tau_G^2\right)$.

Given a DSD, the instantaneous rain rate can be derived based on meteorological principles as a function of the third power of drop diameter (Battan 1973). The rain rate, *R*, for a given *t*, is estimated as

$$R_t = c_R \frac{\pi}{6} \int_0^{\infty} D^3 V\left(D\right) N\left(D\right)_t \, dD$$

where *D* is the drop diameter in mm; *V*(*D*) is a deterministic function for the terminal velocity for drops with diameter *D*; $N\left(D\right)_t$ is the DSD giving the number of drops per mm diameter *D* per m^3 at time *t*; and c_R is an additional constant to account for units.

Given the model for $N(D_i)_t$ (equation 4), summing over all ΔD and substituting $N_t \times P(D)_t$ for $N(D)_t$ yields:

$$R_t = c_R \frac{\pi}{6} \int_0^\infty D^3 V(D) \left[N_t f(D)_t \right] dD \tag{5}$$

where N_t is the total drop concentration at time *t* and $f(D)_t$ is the pdf for D_t defined through $P(D)_t$.

Assuming an independent Gaussian measurement error model, and conditional on rain rate as derived from a DSD and parameter $\theta_G = \{\sigma_G^2\}$, the elements of *G* are independent and in each case $G_t : N(m_{G_t}, \sigma_G^2)$ where m_{G_t} is the integral over the time interval between gauge measurements, of the instantaneous rain rate defined above.

Ground Radar Observations, Z

The equivalent radar reflectivity observations from ground radar, *Z*, are a function of the average returned power and range of the radar scan (Battan 1973). If there are no measurement errors, *Z* will conform to the reflectivity derived from a DSD. To accommodate various factors that can lead to errors when estimating Z, we construct a simple measurement error model for equivalent radar reflectivity as a function of the derived reflectivity and random errors; $Z_t = Z_{D_t} + \varepsilon_{Z_t}$ where Z_{D_t} is the meteorologically derived reflectivity at time *t* and ε_{Z_t} are independent measurement errors modeled as $N(0, \sigma_Z^2)$; $\sigma_Z^2 : IG(\mu_Z, \tau_Z^2)$.

The treatment of derived reflectivity, Z_D, is similar to that for rain rate. Given a DSD, Z_D can be derived as the sixth moment of the DSD (Battan 1973). Derived reflectivity at time *t* is estimated as

$$Z_{D_t} = c_Z \int_0^\infty D^6 N(D)_t dD$$

where *D* is the drop diameter in mm; $N(D)_t$ is the current drop spectra giving the number of drops per mm diameter *D* per m^3 at time *t*; and c_Z is a constant related to sample volume and

equals one for a one-m^3 sample. As for gauge observations, summing over all ΔD and substituting $N_t \times P(D)_t$ for $N(D)_t$ yields:

$$Z_{D_t} = c_Z \int_0^\infty D^6 \left[N_t f(D)_t \right] dD$$

where variables are as previously defined. In practice, Z_{D_t} is estimated as a sum over binned drop sizes. Assuming this independent Gaussian measurement error model, and conditional on Z_D derived from a DSD and parameter $\theta_Z = \{\sigma_Z^2\}$, the elements of Z are independent and in each case $Z_t \sim N\left(Z_{D_t}, \sigma_Z^2\right)$.

Model Summary

We use the components presented above to construct the hierarchical model formulation. At the highest level we assume, based on a series of conditional independence assumptions, that given a DSD ($N(D)$), the gauge (G) and radar (Z) observations are conditionally independent with the following factorization:

$$\begin{aligned}\left[N(D), G, Z\right] &= \left[N(D)\right]\left[G, Z \middle| N(D)\right] \\ &= \left[N(D)\right]\left[G \middle| N(D)\right]\left[Z \middle| N(D)\right]\end{aligned}$$

The components $\left[G \middle| N(D)\right]$ and $\left[Z \middle| N(D)\right]$ represent likelihoods based on the gauge and radar data, respectively. The first component, $\left[N(D)\right]$ represents the prior probability for the DSD process. The conditional probabilities for the full model, including the parameters and hyperparameters, are summarized in Table 1.

Table 1. Hierarchical Model Summary.

Variables	Conditional Probabilities for Model Components[a]
Observational Data	$\left[G_t \middle\| N(D)_t\, \theta_G\right]$[b] $\left[Z_t \middle\| N(D)_t\, \theta_Z\right]$
Hidden Process	$\left[TN_t \middle\| \mu_{N_t}, \theta_N\right]$[c] $\left[D_t \middle\| \theta_D\right]$
Temporal Dynamics for mean of *TN*	$\left[\mu_{N_t} \middle\| \theta_\mu\right]$
Model Parameters	$\left[\theta_G, \theta_Z, \theta_N, \theta_D, \theta_\mu \middle\| \theta_H\right] = \left[\theta_G \middle\| \theta_{H_G}\right]\left[\theta_Z \middle\| \theta_{H_Z}\right]\left[\theta_N \middle\| \theta_{H_N}\right]\left[\theta_\mu \middle\| \theta_{H_\mu}\right]\left[\theta_D \middle\| \theta_{H_D}\right]$ where θ_H is a collection of hyperparameters and $\theta_G = \left\{\sigma_G^2\right\} = \left[\sigma_G^2 \middle\| \mu_G; \tau_G^2\right]$ $\theta_Z = \left\{\sigma_Z^2\right\} = \left[\sigma_Z^2 \middle\| \mu_Z; \tau_Z^2\right]$ $\theta_N = \left\{a_0, a_1, b_1, \sigma_{\varepsilon_0}^2\right\} = \left[a_0 \middle\| \alpha_0; \tau_{\alpha_0}^2\right]\left[a_1 \middle\| \alpha_1; \tau_{\alpha_1}^2\right]\left[\beta_1 \middle\| \beta_1; \tau_{\beta_1}^2\right]\left[\sigma_{\varepsilon_0}^2\right]$ $\theta_\mu = \left\{a_N, \sigma_{\eta_t}^2\right\} = \left[a_N \middle\| \alpha_N; \tau_{\alpha_N}^2\right]\left[\sigma_{\eta_t}^2 \middle\| \mu_\eta; \tau_\eta^2\right]$ $\theta_D = \left\{\xi_t, \phi_t^2\right\} = \left[\xi_t \middle\| \mu_\xi; \sigma_\xi^2\right]\left[\phi_t^2 \middle\| \alpha_\phi; \beta_\phi\right]$
Hyperparameters	$\theta_H = \left\{\mu_G, \tau_G^2, \mu_Z, \tau_Z^2, \alpha_0, \alpha_1, \beta_1, \tau_{\alpha_0}^2, \tau_{\alpha_1}^2, \tau_{\beta_1}^2, \alpha_N, \tau_{\alpha_N}^2, \mu_\eta, \tau_\eta^2, \mu_\xi, \sigma_\xi^2, \alpha_\phi, \beta_\phi\right\}$

[a] The bracket notation $\left[\bullet \middle\| \bullet\right]$ is used as a short hand for denoting conditional distributions.

[b] $N(D)_t$ is general notation for $N_t \times P(D)_t$.

[c] $TN_t = \ln N_t$.

Parameter Estimation

We used a Markov Chain Monte Carlo (MCMC) implementation to sample from the posterior distribution of the hierarchical model. The appropriate Markov chain is constructed using a single-component Metropolis-Hastings algorithm based on the hierarchical structure of the model, and on the conditional probabilities associated with the various model components.

Starting values, hyperparameter values, the length of burn-in, and the number of iterations must be determined for each application. Starting values are random, based on the prior distributions of the parameters; the distributions can be altered by the choice of the fixed hyperparameters. The required number of iterations for a given application depends on the length of the burn-in, the autocorrelation of the parameters, and the number of samples needed for parameter characterization. Convergence is monitored using visual analysis of plotted parameters in two ways: 1) by performing multiple runs using different random starts; and 2) running multiple chains using the same starting values.

APPLICATION TO EUREKA RAINFALL

Data

The Eureka site consists of three sources of data: distrometer, rain gauge, and ground radar data. The distrometer data were collected using a JWD instrument. The distrometer was located near the Eureka radar site at $40°48'\text{N}/124°09'\text{W}$. Each data file contained one hours worth of raw bin counts collected at one-minute intervals. The raw bin counts were converted to one-minute DSD spectra counts $N_{D_i}\,(\text{m}^{-3}\text{mm}^{-1})$ by adjusting raw counts for the instrument collection area, time, terminal fall speed, and drop diameter.

A set of processed 10-minute distrometer data was also available. The 10-minute data are drop spectra that have been averaged over 10 contiguous one-minute periods. The observed 10-minute drop spectra were used to calculate rain rates and reflectivity values to use in place of gauge and radar observations for model verification.

Hourly gauge accumulation data (TD-3240) were obtained from the National Climatic Data Center (NCDC) on-line service. The data included measurements from 14 gauging stations in the vicinity of Eureka, California over the period 1 January 1999 through 31 March 1999, of which nine stations were in the vicinity of the radar and distrometer. Data were collected using Fischer-Porter precipitation gauges with automated readouts. Precipitation data from the NCDC database are quality checked and edited, as necessary, by an automated and a manual edit

(Hammer and Steurer 1998). To coincide with the time domain only implementation of the model, gauge data were limited to measurements from the gauging station located closest to the distrometer at 40°49'N/124°10'W .

Level II WSR-88D radar scans were ordered from NCDC for the Eureka radar site, which is located at 40°29'54"N/124°17'31"W , for periods coinciding with rainfall events identified by gauge and distrometer data. There was a radar scan approximately every six minutes for several elevation increments. The spherical coordinate Level II data were interpolated to a 2km x 2km Cartesian coordinate system of equivalent reflectivity (*dBZ*) values using standard programs for NEXRAD radar. The interpolated data were only available for one of the identified rainfall events; the availability of radar data defined, and limited, the time period over which the model was implemented.

Appropriate reflectivity values needed to be extracted from the 2km x 2km gridded values to correspond with the fixed spatial location of the distrometer and closest rain gauge. Reflectivity in the time domain was extracted from the spatial data by first identifying the closest pixel located vertically above the distrometer and nearby rain gauge. Given the uncertainty in the vertical flow of rainfall above the surface of the earth to a point on the ground, an equivalent reflectivity was calculated as the aerial average of first and second order neighbors using the identified pixel and the eight neighboring pixels. Issues of beam blockage along the path of the radar to the location of the distrometer were considered to be negligible due to the gradually increasing nature of the coastal terrain of Northern California. Therefore, the reflectivity values were extracted from the lowest radar beam (tilt elevation of 0.5^{o}).

We implemented the hierarchical Bayes model using data from a nine hour (550 minutes) rainfall event on 6 February 1999. For the course of the model run there were nine hourly gauge readings and 92 ground radar observations.

Specification of Hyperparameters, Initial Values, and Constraints on Parameters

We used exploratory analysis of distrometer data from five February rainfall events to come up with informative means and variances for the fixed hyperparameters. The constant hyperparameter values are summarized in Table 2 with model results.

In addition to the hyperparameter values, we specified several $t = 0$ starting values as follows: $\varepsilon_0 = 0,\ \sigma^2_{\varepsilon_t} = 0.001,$ and $\mu_{N_0} = 3.0$. We based initial conditions for the remaining parameters on random draws from the appropriate prior distributions.

We introduced cutoffs for two types of constraints; namely, constraints due to rainfall dynamics/instrumentation and constraints due to statistical considerations. Constraints based on rainfall dynamics included upper bounds for three parameters:

1. We limited the total number of drops in a one-minute interval to 6.0, about 2.5% above the maximum value observed from the distrometer data. This corresponds to a total of $\exp(6) \approx 403$ raw count drops when summed over the 20 drop diameter bins. We applied this cutoff to the parameters TN_t and μ_{N_t}.
2. We limited the upper bound for drop diameter size such that $P(D \geq 5) \approx 0.000001$. This limitation was imposed since the rain event of interest represented light rainfall. The distrometer data had no raw counts beyond a mean drop diameter of 3.15 mm.

Proposals for parameter updating are drawn from uniform distributions centered at the current parameter values. This implementation made it possible to draw negative-valued proposals, depending on the current value of the parameter. Constraints based on statistical considerations thus included lower bounds on variance parameters (>0) and $\xi_t > 0$. Required model constraints were also placed on the three *GARCH* coefficients.

Estimating m_G and Z_D

At each time period for which there were gauge or radar observations, the data components of the model required computation of the two quantities m_G and Z_{D_t}, the true equivalent gauge value and the true derived reflectivity, respectively. Both of these values are deterministic quantities based on the current state of the parameters that define the hidden DSD process.

We calculate the true equivalent gauge value as the sum over the previous 60 minutes of each of the instantaneous rain rates in mm/min derived from the state of the DSD before each point in time. The instantaneous rain rates defined by equation 5 are approximated by summing over the 20 drop diameter bins that would typically be observed using a JWD instrument as follows:

$$R_t = \frac{3.6\pi}{6000} \sum_{i=1}^{20} \left[\frac{\exp\{TN_t\} \times P(D_i)_t}{\Delta D_i} \right] D_i^3 V(D_i) \Delta D_i$$

where $P(D_i)_t$ is based on the cdf of a lognormal distribution with parameters ξ_t and ϕ_t^2. The terminal fall speed, $V(D_i)$, was calculated as $3.78D_i^{0.67}$ (Smith 1993), where D_i is the mean drop diameter for bin *i*. Then

$$m_{G_t} = \frac{\sum_{i=t-60+1}^{t} R_i}{60}$$

True derived reflectivity is calculated as an average, over the previous 10 minutes, of the sixth moment of the DSD. As with the equivalent gauge calculation, we approximate derived reflectivity by summing over the 20 drop diameter bins using

$$Z_t = \sum_{i=1}^{20} \left[\frac{\exp\{TN_t\} \times P(D_i)_t}{\Delta D_i} \right] D_i^6 \Delta D_i$$

Then

$$Z_{D_t} = \frac{\sum_{i=t-10+1}^{t} Z_i}{10}$$

Z_{D_t}, in units of mm^6m^{-3}, are converted to *dBZ* as $10\log_{10}(Z_{D_t})$.

Tuning, Thinning, and Convergence

Tuning the chain was optimized by making many small runs (5000 to 15000 iterations in length) and monitoring the acceptance rates for each of the parameters. Adjustments to the acceptance rates were made by either increasing or decreasing the upper and lower bounds on the uniformly distributed parameter proposals. After much experimentation, bounds on the uniform proposals were adjusted to yield acceptance rates ranging from about 0.42 on the low end to about 0.85 on the high end. Most parameters had acceptance rates in the range of 0.60 – 0.65. While these values exceed the commonly recommended values of 0.1 to 0.5 for updating multidimensional components (Roberts 1996), we found that mixing and convergence performed better at these higher levels.

We used thinning to obtain roughly independent samples from the joint posterior distribution. We used *gibbsit* (Raftery and Lewis 1995) to get estimates of the ideal spacing between iterations and the number of burn-in iterations, as well as the number of iterations for a desired precision. In most cases a thinning value of 150 was more than sufficient to obtain

approximately independent samples, although on occasion, a parameter thinning value would hit 300 or so.

We monitored convergence by visual inspection of plotted output. The number of burn-in iterations recommended by *gibbsit* was generally low and always < 5000. Convergence was also assessed by looking at the path of several runs using both the same and different starting values.

Results

Model Verification

There is no way to specifically validate the model in terms of comparing outputs to "true" values. Prior to running the fully implemented model, however, we conducted model verification by assessing the ability of the model to realistically reproduce targeted signals and integrated parameters associated with surface rainfall. We replace NEXRAD radar observations with derived reflectivity calculated from 10-minute distrometer data. Rain gauge data are omitted for this assessment. To maximize the inputs we use a 10-minute time step, thus insuring a data value at each time step.

The verification run includes 49 derived reflectivity observations computed from the one set of distrometer data. The MCMC run was ended after 280,000 iterations following 100,000 burn-in iterations. Based on graphical output, all parameters converged within the 100,000 burn-in period. Results for a typical model verification run showing time independent parameter estimates and time varying parameter estimates (using the 800 MCMC samples generated after applying a thinning value of 350) are shown in Tables 2 – 3, respectively. Further, using the posterior parameter estimates for the hidden process we can compare calculated integrated rainfall parameters with values obtained from the 10-minute distrometer data. Figures 3 and 4 show the comparisons for rain rate and reflectivity. Note that although we compare the model output with quantities derived from distrometer data, we do not consider the distrometer data as ground truth. Distrometer data are subject to measurement errors.

The posterior means of the three GARCH coefficients are all lower than those specified for the prior distributions. The variability for the a_1 and b_1 coefficients is slightly less than that for the priors. The mean of the AR(1) coefficient is slightly higher than that specified in the

prior. The variance parameter for the AR(1) errors, σ_η^2, is considerably higher than that specified in the prior, as is the measurement error variance for radar (σ_Z^2).

The mean values for *TN* and μ_N are about the same and posterior estimates of *TN* are close to expected counts although the model tends to underestimate the counts. In general, the DSD distribution parameters ξ_t and ϕ_t^2 are more variable than estimates based on the distrometer data. On average, the posterior estimates of ξ_t are similar to the values estimated from the distrometer data. The model tends to slightly underestimate ϕ_t^2.

The posterior estimates for both integrated parameters (rain rate and reflectivity) are more variable than values calculated from the distrometer data realization. This is expected since these estimates include the distrometer data realization with some uncertainty. Aside from the extra variability, the model estimates are a good match with the integrated quantities calculated from distrometer observations. The model does a good job of capturing the overall means and the signals present in the time series.

Full Model Run

The full model run includes 9 hourly rain gauge observations and 92 effective reflectivity observations from ground radar. The MCMC run was ended after 245,000 iterations following a burn-in period of 135,000 iterations. After applying a thinning value of 350, there were 700 MCMC samples for posterior estimates. Convergence for the full model run is much slower than that for the verification run, taking up to about 60,000 iterations for a few parameters. With the exception of the AR(1) variance term, σ_η^2, all parameters appear to converge by 60,000 iterations. The σ_η^2 variance hardly moves until 60,000 iterations, at which point the value starts bouncing around and then settles down at about 0.5, followed by more intermittent bouncing. This pattern could be due to slow mixing or identifiability issues.

Results for a typical full model run showing time independent parameter estimates and time varying parameter estimates are shown in Tables 2 – 3, respectively, along side the verification run results. Time series results for integrated parameters calculated from the estimated parameters for the hidden DSD process are shown in Figures 5 – 8.

Except for a_0, the posterior estimates for the *GARCH* parameters are similar to those from the verification run. The generally larger standard errors reflect the greater uncertainty in

the one-minute model, which does not have data at every time step. The mean error terms, ε_t, are not significantly different from zero based on 95% confidence intervals for the means. The posterior distributions are relatively symmetric and approximately normal, which is in line with the *GARCH* model for ε_t. The within time posterior means for TN_t and μ_{N_t} are essentially the same. The posterior distributions for TN_t are nearly symmetrical and appear to be approximately normally distributed.

The representative ξ_t distributions all have means close to the mean for the prior distribution. The posterior distributions are fairly symmetric; values for ξ_t are variable, ranging from -1.82 to 1.01 and do contain the observed estimates based on the one realization of distrometer data. The representative ϕ_t^2 distributions all have means close to 0.18, the mean of the prior distribution. The posterior distributions are right-skewed; values for ϕ_t^2 are variable, ranging from 0.05 to 1.03 and do contain the observed estimates calculated from the distrometer data.

The in-series variability for TN_t, ξ_t, and ϕ_t^2 is greater than that for the observed distrometer data. Except for ϕ_t^2, the time series values estimated from the distrometer data are nearly centered in the posterior model results. As with the verification run, the full model underestimates the DSD shape parameter, ϕ_t^2.

From the posterior parameter estimates for the hidden DSD process, we can estimate the integrated values for rain rate and derived reflectivity. The time series of calculated rain rates show higher variability than that produced from the one realization of distrometer data. The large variability makes it difficult to compare the signals, though there appears to be some amount of agreement in the gross signal features for rain rate (Figure 5).

The time series plots of derived reflectivity also show higher variability than that produced from one realization of distrometer data, though not to the same extent as for rain rate (Figure 6). The model output consistently overestimates distrometer derived values. A mean 1 normalized plot for reflectivity (Figure 7) removes the bias and clearly shows some similarity in the signals compared with the distrometer based calculations. While the cause of the bias is unknown, it's possible that the observed bias indicates poor calibration of the radar.

An alternate visualization of the calculated rain rate parameter based on model output displays histograms for rain rates from 12 points along the time series (Figure 8). Corresponding values calculated from the distrometer data are also shown. The individual histograms contain the distrometer-based rain rates, although the distrometer-based values are in the lower ends of the distributions for most time points.

Table 2. MCMC Summary Statistics for Time Independent Parameters.

			MCMC Posterior Mean (SE)	
Model Parameter	**Prior Mean**	**Prior SD**	**Verification Run**	**Full Run**
GARCH Coefficients a_0	0.023	0.0062	0.0080 (0.00012)	0.049 (0.0012)
a_1	0.242	0.1007	0.200 (0.00082)	0.164 (0.0020)
b_1	0.742	0.1119	0.677 (0.00099)	0.684 (0.0031)
AR(1) Process a_N	0.95	0.04	0.981 (0.00018)	0.990 (0.0002)
σ_η^2	0.005	0.0007	0.418 (0.0023)	0.478 (0.0033)
Measurement Errors σ_G^2	5.0	1.02	---	5.82 (0.044)
σ_Z^2	2.86	1.429	31.68 (0.132)	104.01 (0.865)

Table 3. MCMC Summary Statistics for Time Varying Parameters for three time points. For the verification run, times are t = 9, 24, 39. For the full run, t = 90, 235, 391.

	Verification Run		**Full Model Run**	
Model Parameter	**Observed Value**	**Posterior Mean (SE)**	**Observed Value**	**Posterior Mean (SE)n**
Total # drops				
TN_{t1}	6.41	6.00 (0.026)	4.37	5.00 (0.056)
TN_{t2}	7.59	6.36 (0.029)	5.90	5.28 (0.055)
TN_{t3}	6.64	6.29 (0.028)	5.08	5.04 (0.051)
Mean # drops				
$\mu_{N_{t1}}$	---	6.00 (0.026)	---	5.01 (0.052)
$\mu_{N_{t2}}$	---	6.34 (0.028)	---	5.25 (0.049)
$\mu_{N_{t3}}$	---	6.30 (0.026)	---	5.05 (0.048)
Errors for TN_t				
ε_{t1}	---	-0.0001 (0.008)	---	-0.017 (0.0214)
ε_{t2}	---	-0.005 (0.008)	---	-0.016 (0.0221)
ε_{t3}	---	-0.012 (0.010)	---	0.023 (0.0217)
Error Variance				
$\sigma^2_{\varepsilon_{t1}}$	---	0.050 (0.0025)	---	0.320 (0.0077)
$\sigma^2_{\varepsilon_{t2}}$	---	0.071 0.0036)	---	0.329 (0.0084)
$\sigma^2_{\varepsilon_{t3}}$	---	0.0800 (0.0040)	---	0.319 (0.0075)
Lognormal Scale Parameter Prior mean: -0.50 SD: 0.400				
ξ_{t1}	-0.54[a]	-0.68 (0.009)	-0.56[b]	-0.49 (0.015)
ξ_{t2}	-0.59	-0.46 (0.009)	-0.61	-0.54 (0.015)
ξ_{t3}	-0.63	-0.49 (0.009)	-0.63	-0.53 (0.015)
Lognormal Shape Parameter Prior mean: 0.182 SD: 0.0857				
ϕ^2_{t1}	0.316[a]	0.155 (0.0018)	0.26[b]	0.18 (0.003)
ϕ^2_{t2}	0.338	0.176 (0.0023)	0.34	0.19 (0.003)
ϕ^2_{t3}	0.371	0.166 (0.0023)	0.37	0.19 (0.004)

[a] Estimated from 10-minute distrometer data.
[b] Estimated from 1-minute distrometer data.

DISCUSSION

The model results for the Eureka rainfall are promising. Using hierarchical Bayes methods we are able to integrate information from multiple data sources for purposes of estimating rain rate and related parameters. While we are unable to validate the posterior estimates of TN_t, ξ_t, and ϕ_t^2, the main parameters for the unobserved DSD process, we can compare their values with the one available distrometer realization. For TN_t, many of the MCMC realizations, though not all, produce time series signals similar to that observed for the distrometer realization with model output that is always more variable. This higher variability is due, in part, to the higher uncertainty that is introduced when observations from two different data sources are considered. The DSD scale and shape parameters capture the gross features of the drop spectra, but the time series for ξ_t and ϕ_t^2 are many times more variable than those produced from distrometer estimates. A comparison of autocorrelations for these parameters based on distrometer data versus model output suggests that there is more structure in the time evolution of the DSD shape than what is accounted for in the current model.

Ultimately the estimated DSD parameters are used to produce estimates of the integrated quantities rain rate and derived reflectivity. The verification runs do a good job of capturing the features of these quantities; the full model runs capture the gross features, but again produce higher variable time series output and overestimates of derived reflectivity. The bias seen in derived reflectivity could be due to a number of factors; most notably, differences may be related to sampling errors associated with distrometer data. Derived reflectivity is more adversely affected by small differences in drop-size distributions since Z_D is a function of D^6 while *R* is a function of D^3. Either overestimation of larger drops by the model, or under recording of larger drops by the distrometer could lead to these observed differences. Verification runs produced better consistency between model derived reflectivity and distrometer derived values; use of 10-minute distrometer data helps compensate for distrometer sampling errors that arise due to the small sample volume collected in a one-minute interval.

There is no measurable truth when it comes to determining rain rates at the earth's surface. For each instrument deployed in the field, there are claims and beliefs as to how well they perform their respective tasks, in light of uncertainty. Instead of using rain gauge data as "ground truth", the posterior model distributions capture multiple sources of uncertainty that

reflect uncertainty associated with one or more sources such as: the inability to directly measure quantities of interest; measurement and sampling errors associated with the multiple instruments; sparse gauge data; and, incorrect or inadequate assumptions and specifications related to equations used to represent the underlying physical processes.

As a proof of concept, the hierarchical Bayes model is able to generate reasonable estimates for most of the modeled parameters when applied to the Eureka data set. Using existing model runs as baselines we can further characterize the observed uncertainties by making stepwise modifications to model stages or by incorporating additional data. In the short term, a temporal dynamic stage should be developed for the evolution of DSD through time. Also, we should run the model for other rainfall events. The particular Eureka event consisted of relatively light rainfall. Determining how well the model performs for heavier rainfall events will provide some idea as to the portability/usability of this model under a variety of rainfall scenarios. In the longer term, some ideas for expanding and using this modeling approach include:

1. If applicable, incorporate wind corrections for gauge observations in future model runs. Gauge catchments are inherently affected by winds; a wind effect would be one component of the overall uncertainty observed in the model output.
2. Incorporate a vertical correction model component for radar reflectivity observations, based on an underlying process model, to compensate for the fact that ground-based radar provides information about precipitation above the earth's surface. Include data from vertical pointing radar to augment the vertical correction model component.
3. Expand the time only model to a space-time model. Such a model could provide a tool that allows for the creation of spatial rainfall maps with uncertainty estimates, while also providing the opportunity for more expanded analysis of the component uncertainties.
4. Use model output to estimate better *Z-R* relationships dynamically.
5. Use model output to assess the value added by using other types of precipitation measurements (or by removing some instruments).

ACKNOWLEDGEMENTS

We would like to thank Dr. Sandra Yuter, who provided the distrometer and NEXRAD radar data sets, for her help with data pre-processing and her invaluable insights related to instrumentation, meteorological properties of precipitation, and relationships between the two.

REFERENCES

Alexander C. Volatility and correlation: measurement, models, and applications. In Alexander C, editor, *Risk Management and Analysis. Vol. 1: Measuring and Modelling Financial Risk.* John Wiley Sons Ltd, 1998.

Battan LJ. *Radar observation of the atmosphere*. University Chicago Press, Chicago, revised edition, 1973.

Bell TL. A space-time stochastic model of rainfall for satellite remote-sensing studies. J*ournal of Geophysical Research*, 92(D8):9631-9643, 1987.

Berliner LM. Hierarchical Bayesian modeling in the environmental sciences. *Allgemeines Statistishe Archiv,* 84:141-153, 2000.

Berliner LM, Levine RA, and Shea DJ. Bayesian climate change assessment. *Journal of Climate*, 13(21):3805-3820, 2000a.

Cox DR, S. RF, and IshamV. A simple spatial-temporal model of rainfall. *Proc. R. Soc. Lond.*, 415:317-328, 1988.

Hammer G and Steurer P. *Data documentation for hourly precipitation data TD-3240*. Technical report, National Climatic Data Center, February 11, 1998.

Hrafnkelsson B. Hierarchical modeling of count data with application to nuclear fall-out. *Journal of Environment and Ecological Statistics*, 10:179-200, 2003.

Raftery A and Lewis S. Implementing MCMC. In Gilks W, Richardson S, and Spiegelhalter D, editors, *Markov Chain Monte Carlo in Practice*, pages 115-130. Chapman Hall, London, 1996.

Roberts GO. Markov chain concepts related to sampling algorithms. In Gilks W, Richardson S, and Spiegelhalter D, editors, *Markov Chain Monte Carlo in Practice*, pages 45-57. Chapman Hall, London, 1996.

Rodriguez-Iturbe I, Cox DR, F.R.S., and Isham V. A point process model for rainfall: further developments. *Proc. R. Soc. Lond.*, 417(283-298), 1988.

Rodriguez-Iturbe I and Eagleson PS. Mathematical models of rainstorm events in space and time. *Water Resources Research*, 23(1):181-190, 1987.

Royle JA, Berliner LM, Wikle CK, and Milliff R. A hierarchical spatial model for constructing wind fields from scatterometer data in the Labrador Sea. In *Case Studies in Bayesian Statistics IV*, pages 367-382. Springer-Verlag, 1998.

Skaugen T, Creutin JD, and Gottschald L. Reconstruction and frequency estimates of extreme daily areal precipitation. *Journal of Geophysical Research*, 101(D21):26287-26295, 1996.

Smith JA. Marked point process models of raindrop-size distributions. *Journal of Applied Meteorology*, 32:284-296, 1993.

Stern R and Coe R. A model fitting analysis of daily rainfall data. *Journal of the Royal Statistical Society A*, 147(Part 1):1-34, 1984.

Wikle CK, Milliff RF, Nychka D, and Berliner M. Spatiotemporal hierarchical Bayesian modeling: tropical ocean surface winds. *Journal of the American Statistical Association*, 96(454):382-397, 2001.

Wikle CK. Hierarchical Bayesian space-time models. *Environmental and Ecological Statistics*, 5:117-154, 1998.

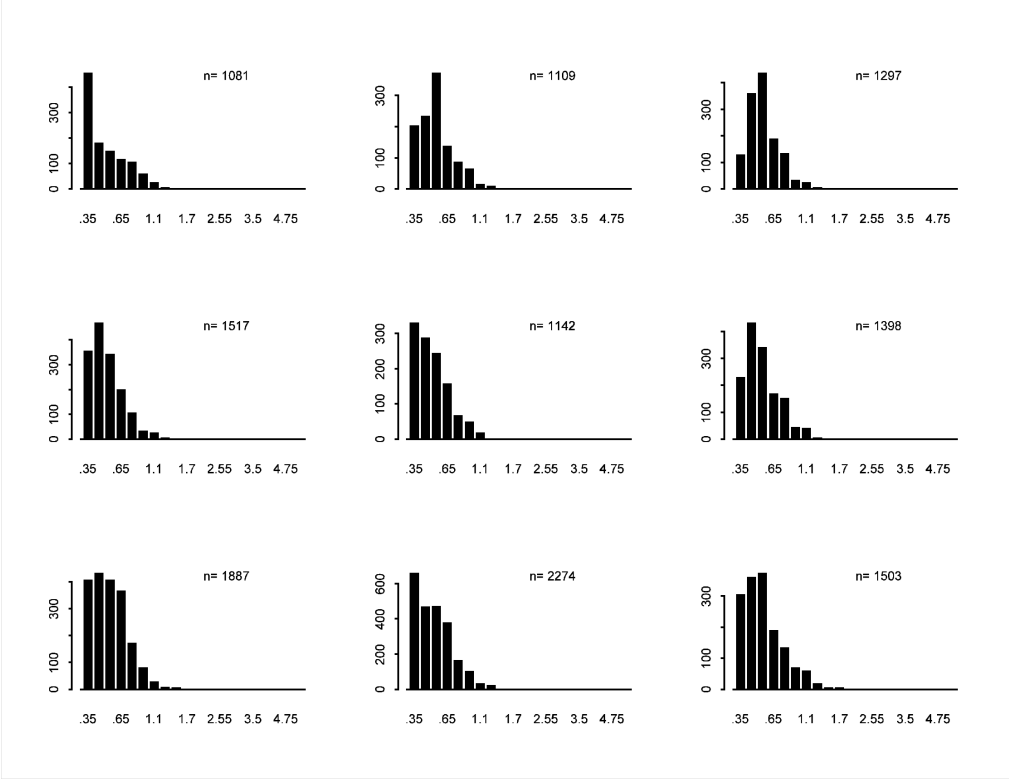

Figure 1. Barplots for nine consecutive minutes of distrometer DSD data from a Eureka rainfall event; *n* is the total number of drops in each one-minute interval.

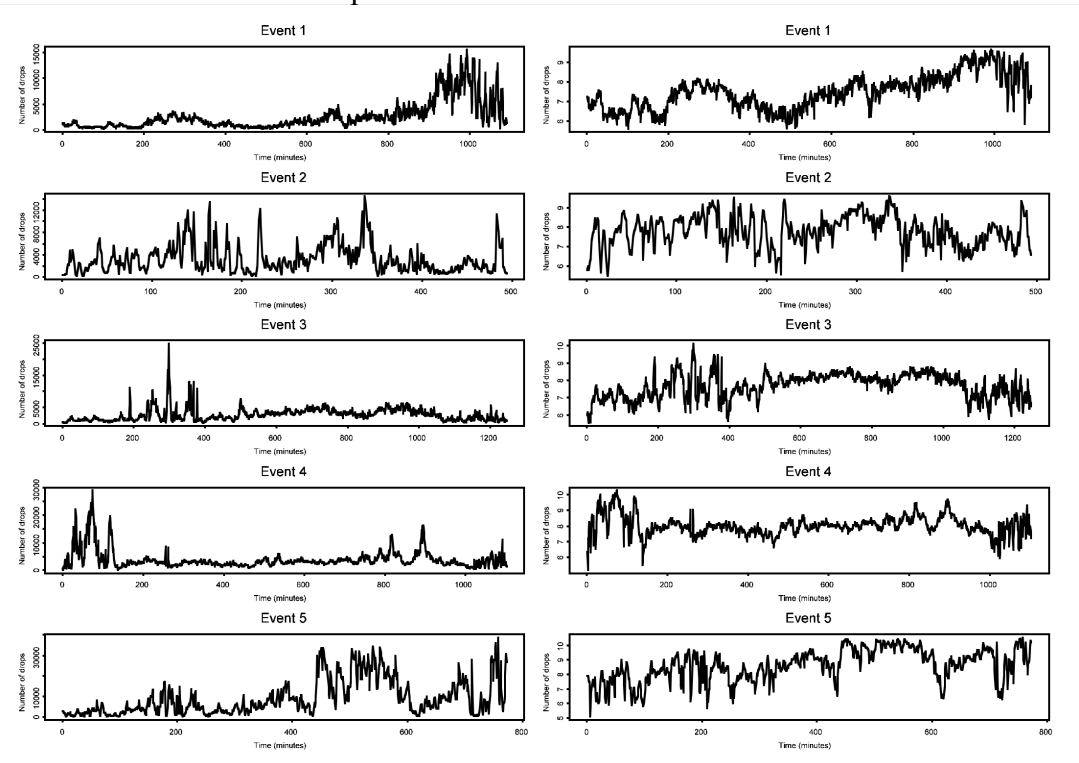

Figure 2. Time series of the one-minute total number of drops (left) and log transformed total number of drops (right) for each of five Eureka rainfall events.

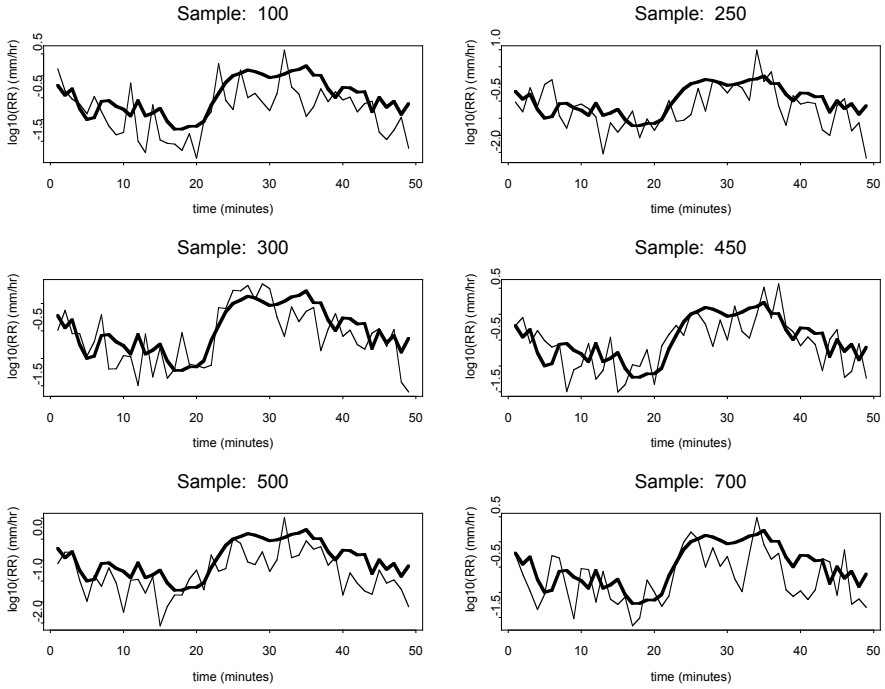

Figure 3. Times series plots of posterior $\log_{10}$ rain rate estimates calculated for some representative samples from the validation run (light lines) with rain rate estimates from the distrometer data (heavy lines).

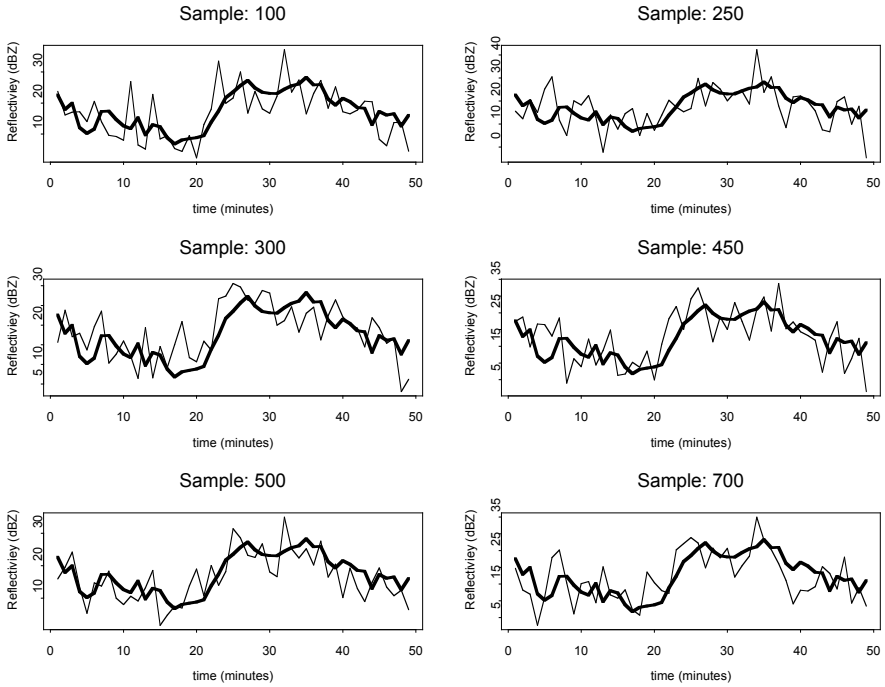

Figure 4. Times series plots of posterior reflectivity estimates calculated for some representative samples from the validation run (light lines) with reflectivity estimates from the distrometer data (heavy lines).

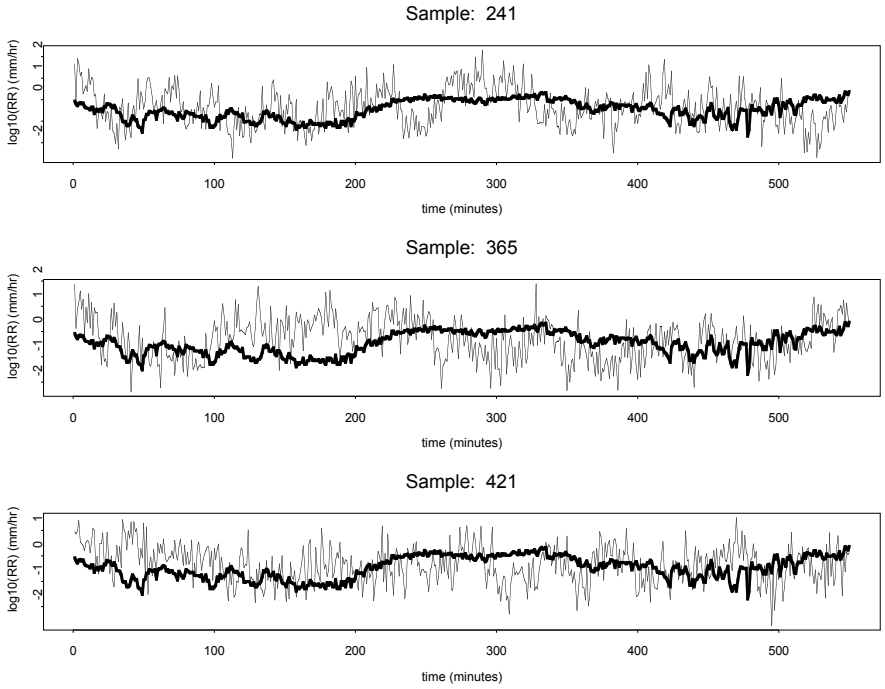

Figure 5. Time series plots of full model run posterior $\log_{10}$ rain rate estimates calculated for selected samples (light lines) with logged rain rate values based on distrometer data (heavy lines).

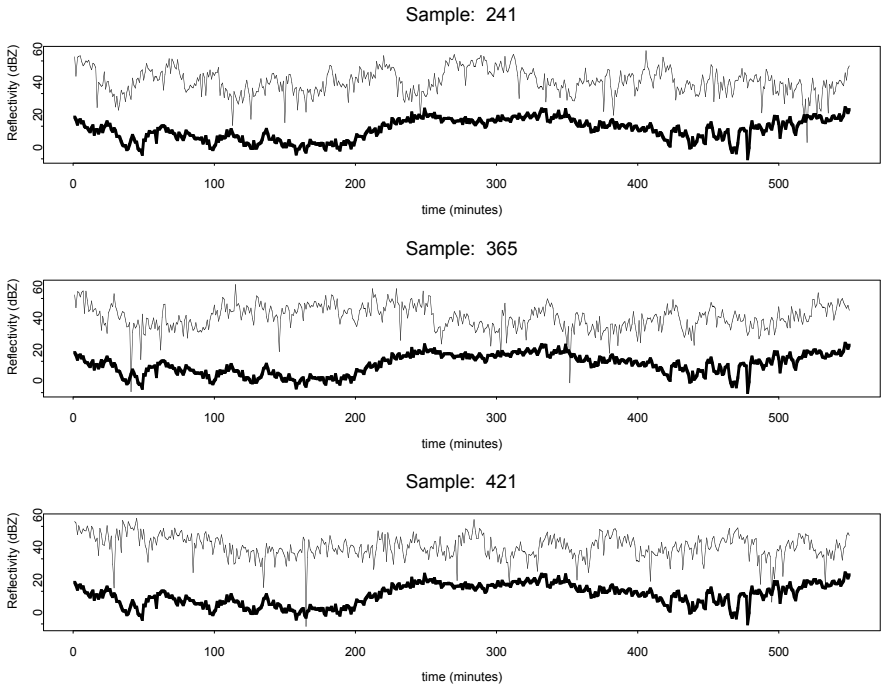

Figure 6. Time series plots of full model run posterior reflectivity estimates calculated for selected samples (light lines) with reflectivity values based on distrometer data (heavy lines).

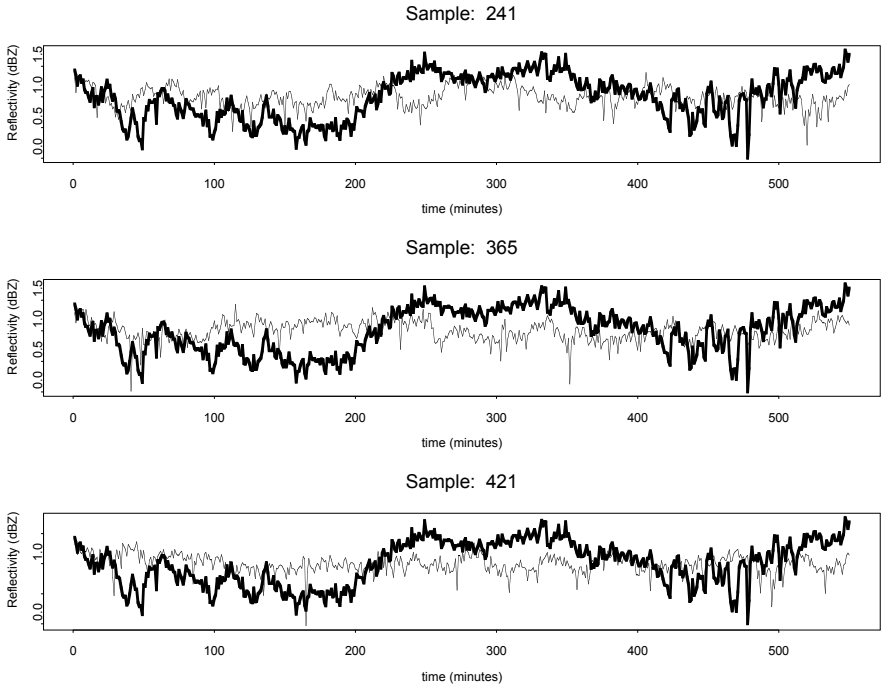

Figure 7. Time series plots of full model run posterior reflectivity (normalized to mean 1) estimates calculated for selected samples (light lines) with reflectivity values based on distrometer data (heavy lines).

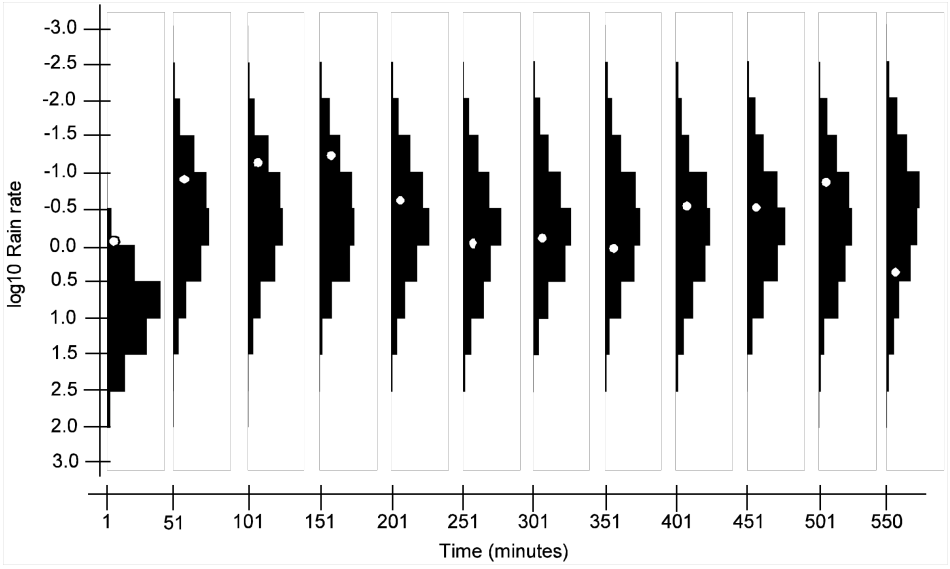

Figure 8. Frequency histograms of the full model run posterior distribution of $\log_{10}$ rain rate estimates for selected time points within the 550 minute modeled event. All histograms are scaled with equivalent *y*-axes. The white circles correspond to the values of rain rate calculated from one realization of distrometer data.