DISSERTATION

AN INTEGRATED RETRIEVAL FRAMEWORK FOR MULTIPLE POLARIZATION, MULTIPLE

FREQUENCY RADAR NETWORKS

Submitted by

Joseph C. Hardin

Department of Electrical and Computer Engineering

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Summer 2015

Doctoral Committee:

Advisor: V. Chandrasekar

Anura P. Jayasumana
Paul Mielke
Margaret Cheney

ABSTRACT

An Integrated Retrieval Framework for Multiple Polarization, Multiple Frequency Radar Networks

Radar networks form the backbone of severe weather and remote sensing in throughout most of the world. These networks provide diverse measurements of weather phenomenon, but ultimately are measuring indirect parameters rather than detecting the physics of the situation. One of the long standing goals of weather remote sensing is to relate the measurements from the various instruments to the physics that give rise to the measurements. Weather radar networks give both a better spatial coverage than single radars, as well as providing multiple looks at the environment. Newly developed radar networks have started to incorporate multiple frequencies and multiple polarizations to take advantage of attributes of different radar frequencies. Raindrops occupy different scattering regimes based on the frequency of the radar being used. Based on this, multiple radars at different wavelengths provide unique information about the microphysical characteristics of the atmosphere. Nonetheless, very little work has been conducted on fusing multiple radar measurements at heterogeneous frequencies to improve microphysical retrievals. This work presents a forward variational algorithm for multiple radar fusion that retrieves microphysical parameters from the atmosphere. The single radar case and the multiple radar case will both be addressed. Ground instrumentation will be used for verification, and the spatial and temporal variability of precipitation microphysics will be discussed.

Table of Contents

List of Tables

xiv

INTRODUCTION

"What you learn from others you can use to follow.

What you learn for yourself you can use to lead."
— Richard Hamming, The Art of Doing Science and Engineering:

Learning to Learn

1.1. MOTIVATION AND BACKGROUND

Weather radars are one of the best and most ubiquitous tools we currently have for predicting weather across the globe. They have been deployed in many different forms to address both operational and research interests for many different purposes. Most common to the atmospheric science community are ground based radars that are used for operational forecasting of the weather in nations around the world. These radars provide a real time view of incoming weather over a large geographic area. The largest deployment of these radars is the NEXRAD network of WSR-88D S-Band radars in the United States, although smaller deployments exist for both operational and research purposes in most industrialized nations. These radars form the backbone of the storm early warning and forecasting system in the US. They are the primary way to detect tornadoes and other severe weather events. In addition, they feed into most numerical weather prediction models allowing them the ability to forecast at long temporal ranges. There has recently been a large expansion of the capabilities of the NEXRAD system by converting them to dual-polarization radars allowing for greatly increased capabilities in every aspect of operation.

In addition to the NEXRAD systems, many universities, government organizations, and research labs deploy research radars used to conduct basic research on the atmosphere. These radars are most commonly ground based, although many airborne systems exist and

are actively fielded. These systems, based on the diversity of uses, tend to be more flexible. This flexibility causes them to be the initial test ground for more advanced features before they are deployed into larger operational networks. One of the larger research networks of these instruments is owned by the Atmospheric Radiation Measurement Program (ARM) funded by the Department of Energy (DOE). They operate multiple fixed sites around the world, as well as multiple mobile instrumentation facilities, with very dense instrumentation allowing for both short and long term field campaigns that attempt to answer basic science questions. Many other deployments of research radars exist both domestically and abroad however. This includes the numerous field campaigns conducted every year by the NASA Global Precipitation Measurement (GPM) Ground Validation (GV) Program, such as the Midlatitude Continental Convective Clouds Experiment (MC3E) and the Iowa Flood Studies (IFLoodS) campaign.

There are also a large number of spaceborne radars that provide us with a global picture of weather. The most popular system was the Tropical Rainfall Measurement Mission (TRMM), which has since been replaced with the Global Precipitation Measurement Mission, a joint mission between NASA and the Japanese space agency (JAXA). This constellation of satellites aim to track the global water budget. It contains two radars and multiple radiometers on the core satellite and additional instrumentation on constellation satellites and is capable of tracking the weather over the majority of the globe.

The longstanding goal of any remote sensing field is to link a physical phenomenon with instrument measurements taken of that phenomenon. In the field of weather radar, the ultimate goal is to use radar measured parameters to retrieve the microphysical properties of the atmosphere. The microphysical properties of the atmosphere include many different parameters. Among these are the temperature of the atmosphere and particles it contains.

The microphysical parameterization of highest importance to us is the drop size distribution of particles (DSD). This quantity tells us how many particles of each size there are in a volume of the atmosphere. A related quantity is the drop shape relationship (DSR) of particles which tells us the approximate shape of the particles as they change sizes. Another parameter of interest is the species composition of these particles which differentiates between snow, ice, graupel, rain, and various other sub-species. Estimation of each of these parameters has spawned it's own research subfield[3–5].

Microphysical properties of the atmosphere can be incorporated into forecasting and numerical weather prediction (NWP) models. Any improvements in the microphysical parameterizations will yield large improvements in NWP which improves the forecasts that are issued operationally. Estimation of these parameters is very difficult as there does not currently exist a reliable way to do *in-situ* measurements that is fielded by any major agency. The relationships between microphysical parameters and remote sensing instrumentation outputs are related by complex nonlinear functions, and based upon parameterizations, are not usually invertible.

Estimating microphysical properties from radar data is not a new idea, and there are a large number of papers on different algorithms[6–9]. The current most popular algorithms for estimation of microphysical properties are primarily based upon empirical relationships between scattering simulations and radar measured data, predominantly utilizing a power law fit. Unfortunately, empirical models are usually tuned for a particular situation and as the situation changes, they often give poor performance. The relationships are usually calculated by using data captured at the ground by a disdrometer, a device capable of measuring the drop size distribution at a small location in space, most often by using an optical array, and calculating equivalent radar measured parameters based on T-Matrix scattering

simulations[10]. These models are heavily dependent upon the assumptions that go into the scattering simulation. To alleviate this, ensemble averages are used, which generalizes the model allowing them to be used in more situations, but reducing their specificity in any given situation. In addition, they only account for a spatially localized patch of data which is not always capable of uniquely representing a microphysical situation. The topic of the variability of these measurements is a longstanding question of interest to the field.

Even if the empirical models were flawless, the measurement process for radar encounters many very significant sources of error that vastly degrades the performance of any empirical algorithm, and indeed often make the measurements completely useless. As the electro-magnetic waves from the radar propagate through the atmosphere, they encounter several sources of distortion. These distortions can cause large errors and discrepancies in the measured parameters which have a very real and direct effect on the microphysical retrieval. The predominant error source is attenuation. Attenuation is the reduction in the power of the wave as some of the energy is absorbed by gases and liquids in the atmosphere. This causes the returned wave to vary as a function of the atmospheric composition of the medium before it. Heuristic based estimators of the DSD require that the beam has been corrected for attenuation and any errors in this attenuation correction will cause large errors in the retrieval of the microphysical properties. There are many other sources of error that we will address in this work as well, all of which serve to obscure the true underlying microphysical properties of the atmosphere when measured from a radar.

A relatively new development in the situation is the deployment of multiple networks of radars operating at varying frequencies (heterogeneous networks) such as the ARM Southern Great Plains (SGP) facility in northern Oklahoma. These networks contain radars at multiple frequencies, often with different scan strategies, strengths, and weaknesses. The

different wavelengths of these radars cause drops to fall into different scattering regimes. We can leverage the increased information this provides to improve our microphysical estimates beyond the normal improvement we get from having multiple independent measurements. Working with multiple frequencies however brings it's own sets of challenges as each radar frequency band, while providing different information, also suffers from differing sets of weaknesses. For instance, as the wavelength decreases we get an increased phase response, but also get a larger amount of phase noise and attenuation.

Nonetheless, correlations between radar measurements and microphysical parameters in the atmosphere do exist, and based on radar measurements we can infer information about the atmosphere. In this work, we will discuss some of the information that we can extract from radar measurements.

1.2. Problem Statement

Given the microphysical properties of a parcel of the atmosphere, we can calculate the value of the measurements that a given radar would see. This is accomplished with the use of scattering simulations and systems theory. Unfortunately, the inverse direction is not so easy. Given a set of radar measurements, the retrieval of the microphysical properties of the atmosphere is an ill-posed, under-constrained problem. We have essentially three mostly independent parameters that are measured, and in return need to derive a non-parametric probability distribution function of drop sizes, a drop shape relation, and a hydrometeor type identification. To simplify the problem, we usually parameterize the drop size distribution with a normalized gamma distribution. This reduces the large problem space of drop size distribution estimation to the estimation of three parameters per range bin.

Further complicating the retrieval of the microphysical parameterization are several sources of error that are unavoidable. Along with the measurement error that is present with any real world measurement, we also suffer from several other more significant errors including attenuation, ground clutter, and backscatter differential phase, among others. The most influential of these errors is attenuation. As radar waves propagate through an atmospheric medium, some of the energy is absorbed by both the atmosphere, and precipitation. The amount of attenuation is dependent both upon the radar wavelength and the precipitation. Any algorithm designed to retrieve the microphysical parameterization needs to be able to correct attenuation using measurements that are themselves corrupted by attenuation. If our assumption of drop shape relationship is incorrect, then this problem becomes even more difficult.

There appears to be some hope when one considers a network of radars operating at different frequencies, as these provide independent overlapping views of the same precipitation. Each radar provides information along a different path which suffers from different attenuation sources yielding unique path integrated attenuations. If we assume our error in estimating attenuation and correcting for it between each radar is independent of other radars, then by integrating multiple views, we reduce the overall uncertainty about the measurement as a function of the number of radars used. If the radars differ in frequency from one another, we get additional information about the size distribution due to the transition between Rayleigh and Mie scattering being frequency dependent, as well as the direct dependency of phase shift on drop size relative to wavelength. This causes the ratio of returns from different frequencies to provide us with estimates of some properties of the environment such as liquid water content. In addition, different frequencies of radar benefit from different unique characteristics. Lower frequency radars suffer much less attenuation than their

shorter wavelength cousins. This makes them a much better estimator of returned power. Higher frequency radars have a much more pronounced phase response in the presence of precipitation making them a much better estimator of rainfall rates especially in light rain scenarios. By combining multiple measurements at multiple frequencies we can attempt to take advantage of the best characteristics of each radar. The samples from different radars lay on incongruous spherical grids however, and so must be interpolated onto a common grid before comparison. This provides yet another source of error that must be factored into the retrieval. How to incorporate these multiple measurements has not been adequately worked out in the literature and so the ideas in this proposal represent a major step forward in the solution to that problem.

1.3. Research Objectives

The main objective of this research is to improve the retrieval of drop size distributions by using multiple frequency heterogeneous dual-polarization radar networks. The development of an algorithm capable of this has to solve several sub-problems. Therefore we can break this goal up into several specific research objectives. These are:

(1) Improve current dual polarization single radar retrievals by using a variational scattering based approach capable of correcting attenuation and retrieving the drop size distribution, while ensuring physical self consistency.

(2) Develop a network microphysical retrieval algorithm that incorporates multiple frequencies and polarizations, capable of working with non time-aligned, non-spacially co-located radars.

(3) Incorporate non-radar instruments such as disdrometers into the retrieval to improve estimates of the microphysical parameters as well as providing a verification source.

7

(4) Characterize the spatial variability of rain microphysics for verification.

1.4. Overview of Dissertation

This document will begin with a discussion of some of the necessary background and prior work in Chapter 2. We will then move on to the development of a single radar retrieval algorithm that operates at multiple frequencies and is capable of retrieving the drop size distribution in Chapter 3. Next, in Chapter 4 we will talk about our strategy for network retrievals and detail a framework for multi-frequency radar network retrievals at multiple frequencies. Then we will use our algorithm and ground disdrometers to characterize the spatial and temporal variability of microphysics in Chapter 5. This will help with verification and fundamental estimation of these parameters. Finally we will wrap up the discussion of the retrievals and examine future directions this work should take. Following the main body of the text, we will address some of the implementation and more practical issues such as quality control and data selection in Appendix A, as well as giving technical descriptions of the radars we will work with in this study.

Prior Work and Background

> What we usually consider as impossible are simply engineering
>
> problems... there's no law of physics preventing them.
>
> — Michio Kaku

The algorithms covered in this dissertation bring together a large number of related but independent theories from different fields. Before examining the multiple radar retrieval algorithm, we need to understand what constitutes the microphysics of precipitation, how the concepts work, and ultimately how microphysical retrievals operate, as well as why multiple radars are beneficial to the solution. To develop this algorithm, we first address how we go from a microphysical parameterization to the radar measured moments and why the reverse process is ill-posed. Next is coverage of Bayesian multi-sensor fusion theory to set the stage for the network paradigm. Some of the more practical applications of microphysical estimation such as rain rate estimation will be covered, including algorithms for microphysical estimation from ground instrumentation. We will briefly step through some of the theory of sensor fusion. Then we will move on to addressing issues that arise when it comes to using and comparing ground direct measurement instrumentation such as disdrometers with remote sensing instruments such as radar. Finally, we conclude by looking at some of the relevant work in the field, followed by a discussion as to the limitations of the current research to set the stage for the development of this body of work.

2.1. Microphysical Explanation for Hydrometeor Formation

While there is a lot that the academic community still does not understand about some of the details of formation of precipitation in the atmosphere, there is a fairly good understanding of the fundamentals. There are different mechanisms for formation of hydrometeors, but

the most important mechanism is the lifting of air. Start with a volume of air, and for the sake of this discussion, assume this volume of air does not exchange energy with it's surroundings, that is, the behavior is adiabatic. Now take this parcel of air, and make it warmer than the air that surrounds it. The ideal gas law is given as

$$PV = nRT \qquad \text{(Ideal Gas Law)}$$

(1)

where P is the pressure of the air mass, V is it's volume, n is the number of moles of the constituent gases, R is the universal gas constant, and T is the temperature of the gas. As we the temperature of the gas increases, the volume will increase proportionately, causing a decrease in density, and thus an increase in buoyancy. This will cause the gas to rise. As the gas rises, the temperature of the surrounding air decreases(up until the tropopause), as does the pressure. This causes the parcel to expand even more to equalize the pressure. As the parcel continues to rise, it's temperature decreases. Up until now we have assumed all of the moisture was in the form of water vapor, however by using the Clausius-Clapeyron equation [11, 12]

$$\frac{de_s}{dT} = \frac{L_v(T)e_s}{R_v T^2} \qquad \text{(Clausius-Clapeyron)}$$

(2)

where e_s is the saturation vapor pressure, T is the temperature, L_v is the latent heat of evaporation, and R_v is the water vapor gas constant, the transition from vapor to liquid water can start to be examined. This equation gives the saturation vapor pressure as a function of T. The relationship between saturation vapor pressure and temperature is exponential.

Therefore as T drops, the saturation vapor pressure decreases significantly. So as the cloud cools off, the relative humidity of the parcel increases exponentially. At some point, the relative humidity will be close to 100%. When this happens, if there are cloud condensation nuclei in the atmosphere, droplets will start to form. If there are not CCN, then this process will wait until it is closer to 108% relative humidity. The difference of these two numbers is governed by the Kelvin equation[13]

(3)
$$ln\frac{p}{p_0} = \frac{2\gamma V_m}{rRT} \qquad \text{(Kelvin's equation)}$$

where p is the actual vapor pressure, p_0 is the saturated vapor pressure, γ is the surface tension, V_m is the molar volume of the liquid, R is the universal gas constant, r is the radius of the droplet, and T is the temperature.

Now that there is liquid cloud droplets starting to form, what happens next? This is where things start to diverge based upon conditions.

The drops eventually reach a size where upward air motion cannot hold them aloft. At this point these drops start to fall. On the way down, they hit other drops. These collisions can cause the two drops to merge(whereby a bigger drop accumulates smaller drops on it's fall) or on occasion can cause drops to break apart. The drops are held together into a spherical shape at small sizes by the surface tension of water. As the drop grows, it starts to change shape, becoming more oblate due to the aerodynamic drag upon the surface of the drop. If the drop grows past a certain point, usually about 8 mm, the drop starts to have a chance of being torn apart by the aerodynamic forces on the drop. When the drop breaks apart, it yields several smaller drops of varying sizes. Additionally, if the air is dry, drops

11

can start to evaporate reducing in size. If the drops all fully evaporate before hitting the ground, the result will be Virga clouds.

The key takeaway point from this very brief overview of thermodynamics and cloud microphysics is that these processes give us a wide variety of drop sizes. Instead of drops asymptotically approaching a single size, there is a behavior where drops grow, and then subsequently shatter into smaller drops. The upward motion of air as well tends to cause sorting of drops which can cause the size distribution of raindrops to be more varied.

2.2. Microphysical Parameterizations

The number of free parameters needed to fully represent the physics of precipitation makes any truly comprehensive approach intractable. The research community has turned to bulk microphysical models to rectify this. These models reduce the behavior of a parcel of atmosphere into a lower dimensional representation of the mean tendencies of the parcel. When dealing with the distribution of the sizes of raindrops in particular, these models reduce the expression for the number of drops of each size in a unit area into a three parameter distribution. We also need to know how these bulk models translate to the measurements present in radar data. We can use T-Matrix scattering theory [14] to simulate what a radar would see when observing an arbitrary parcel of the atmosphere. However, to calculate dual polarization radar measured parameters we also need to know how each drop deforms as it falls due to static tension and aerodynamic drag, as well as it's temperature. In actual practice however, assumptions about shape and temperature are enforced. The goal of our microphysical retrieval will be to retrieve the drop size distribution parameters.

2.2.1. Drop Size Distribution. The Drop Size Distribution (DSD) is a quasi-distribution, In that it is unnormalized with respect to total energy, that represents the number of liquid

drops as a function of drop size. A comparable concept exists that encompasses non-liquid drops called a particle size distribution(PSD). In this work we will focus on liquid drops and so we will only discuss the drop size distribution. Drop size distributions are most commonly represented with a parametric distribution for both analytical, and empirical reasons. One of the first popular parameterizations was as an exponential distribution [15]. The exponential distribution is given by

$$(4) \qquad\qquad N(D) = N_0 \exp\left(\Lambda D\right) \ \mathrm{mm}^{-1}\mathrm{m}^{-3}$$

where Λ is the shape parameter, while N_0 is the intercept parameter. This distribution is most useful when the measurement period is sufficiently long (30 minutes or longer), or the measurement area is sufficiently large. When the measurement period is smaller, the exponential distribution tends towards a poor fit.

The most common parametrization of the drop size distribution for instantaneous measurements is the gamma distribution developed by Ulbrich [16]. This three parameter model is flexible and can accommodate a wide variety of environments, leading to it's popularity. The following equation gives the number of drops having a equivalent volume diameter D per m^{-3} of the atmosphere.

$$(5) \qquad N(D; N_0, \Lambda, \mu) = N_0 D^\mu \exp(-\Lambda D) \qquad \text{for} \qquad 0 < D \leq D_{\max}$$

The parameters of this distribution have units that are not particularly physically intuitive and so an alternative normalized version of this model was created to more closely

relate the parameters to physical measurements [17]. This normalized gamma distribution is given as

$$(6) \qquad N(D; N_W, D_0, \mu) = N_w f(\mu) \left(\frac{D}{D_0}\right)^{\mu} \exp\left(-(3.67 + \mu)\left(\frac{D}{D_0}\right)\right)$$

with

$$(7) \qquad f(\mu) = \frac{6}{3.67^4} \frac{(3.67 + \mu)^{(\mu+4)}}{\Gamma(\mu + 4)}$$

where N_W is the intercept parameter of an equivalent exponential distribution with the same water content and median drop diameter and is related to N_0 by

$$(8) \qquad N_0 = N_w f(\mu) D_0^{-\mu}; \qquad \mathrm{mm}^{-1-\mu}\mathrm{m}^{-3}$$

While D_0 is the median drop diameter that satisfies

$$(9) \qquad \frac{\pi}{6} \int_0^{D_0} D^3 N(D) dD = \frac{\pi}{6} \int_{D_0}^{D_{max}} D^3 N(D) dD$$

D_0 gives the diameter at which half of the water is contained in drops less than diameter D_0 and half of the water is contained in drops larger than D_0. μ is a shape parameter which changes the form of the distribution. Examples of gamma distributions are shown in Figure 2.1 for various values of D_0 and μ.

Shown in Figure 2.2 are example exponential distributions with the same liquid water content as in Figure 2.1. The lines correspond to the exponential version of the lines in the

14
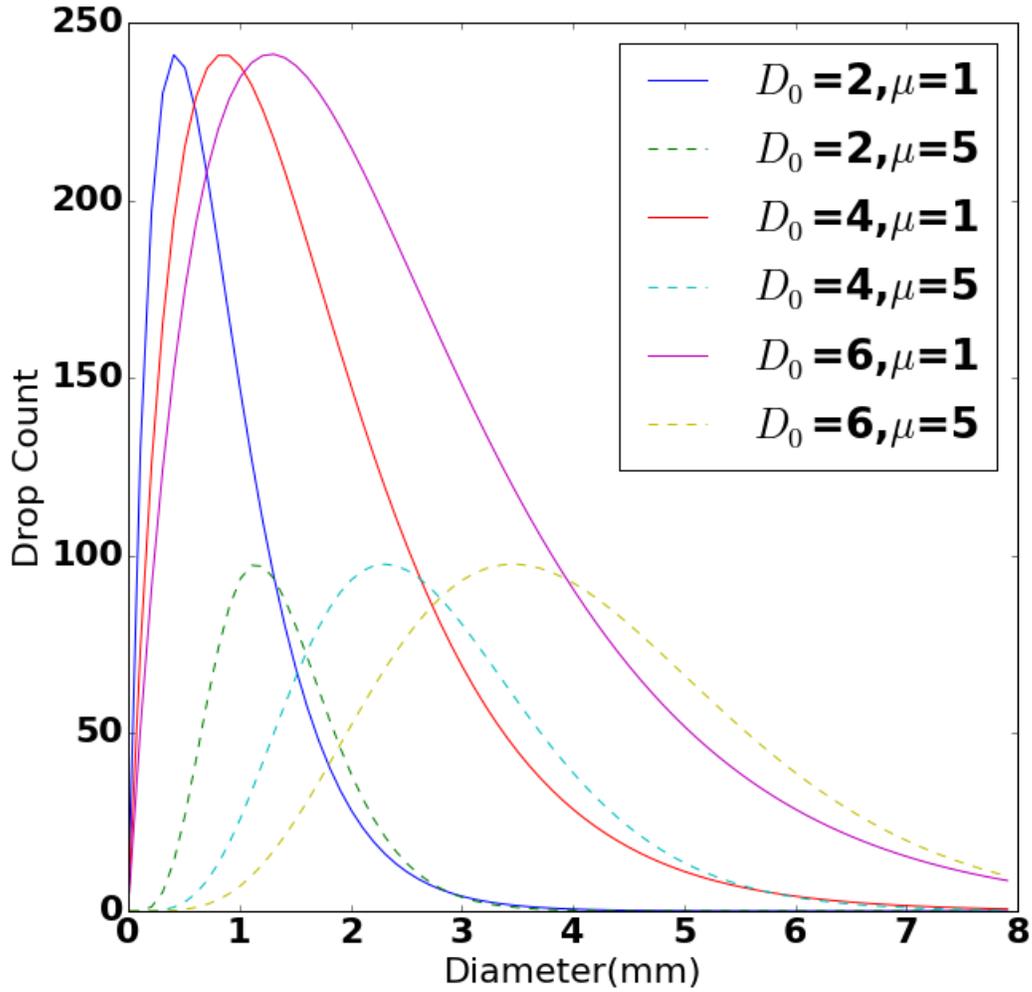
FIGURE 2.1. Shown is example drop size distributions(DSD's) with differing median drop diameter(D_0), and shape(μ) parameters. For all curves, the normalized intercept parameter (N_w) is fixed to 1000. This shows that while both D_0 and μ change the shape of the distribution, μ has a much larger effect on the amount of the distribution contained in the tails.

previous figure. These plots show two possible choices for representing the parameterization of the distribution of raindrops. Given the same total volume of water in the air, there are many physically plausible ways of distributing it into drop sizes.
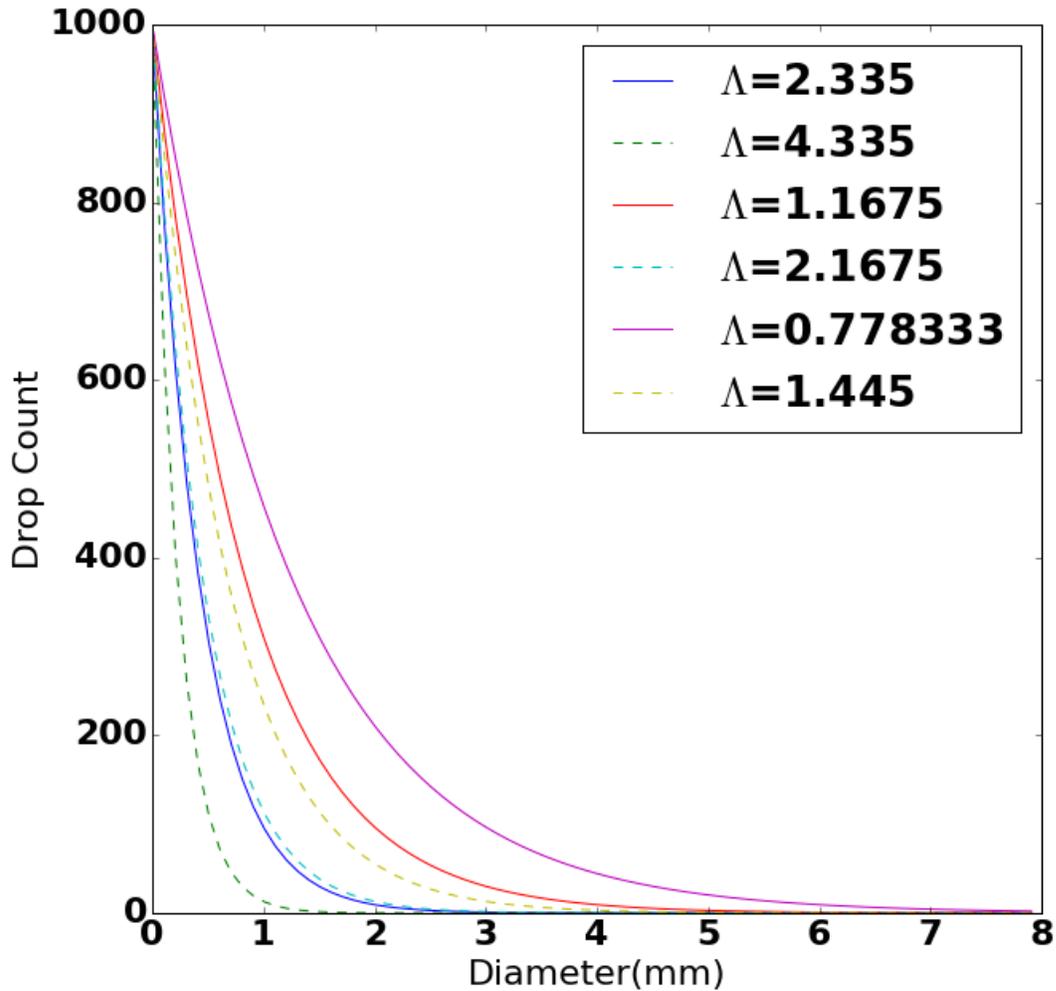
FIGURE 2.2. Example exponential distribution DSDs. Each line in this image corresponds the exponential distribution with the same liquid water content as in the previous figure. The shape parameter Λ is calculated so as to preserve the liquid water content.

The instrument used to measure the drop size distribution of rainfall is called a disdrometer. It proves estimates of DSD's, and occasionally DSR's, as a function of time. These devices use a variety of methods, most commonly either laser or camera based, to count the number of drops of each size, and measure their vertical velocity. Shown in Figure 2.3 is the drop size distribution measured at a disdrometer over the course of a day taken in five minute
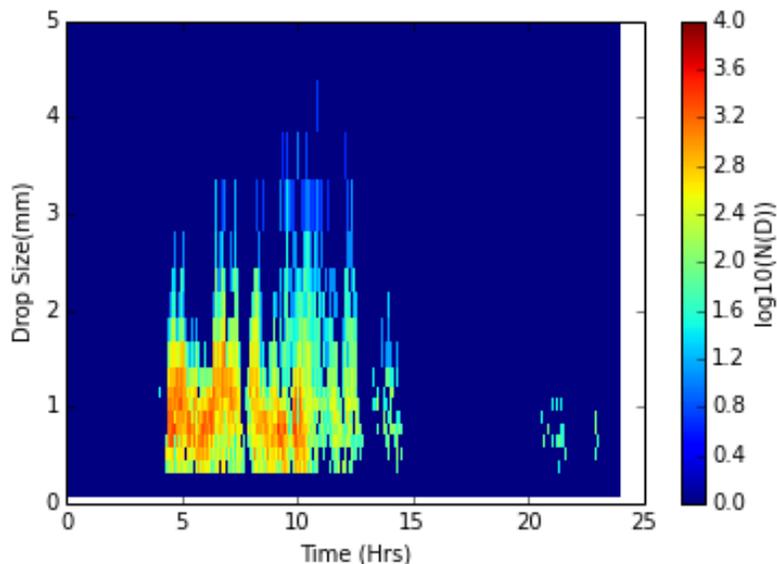
FIGURE 2.3. Drop Size Distribution measurements from a disdrometer using a Parsivel disdrometer from the IFLoodS Field campaign. The plot displays 24 hours of data, captured in five minute increments. The periods between five and thirteen hours are the most active with the largest drops recorded existing around the 4 mm mark.

intervals. From hours five to fifteen, large variations in the average drop size can be seen. These measurements are taken at the ground and unfortunately not directly comparable to radar. They do however give a good idea about the microphysical conditions immediately above them. The data from disdrometers can be used to calculate various microphysical parameterizations, and can contribute to restricting the free parameters for radar retrievals.

We can examine a shorter time instance to see the form of the DSD as shown in Figure 2.4. As mentioned above, the distribution more closely resembles the form of a gamma distribution than that of an exponential distribution. This is expected due to the short five minute time interval over which it was taken. If this sample was taken over a longer time period (several hours for instance), then the DSD would more closely resemble an exponential distribution.
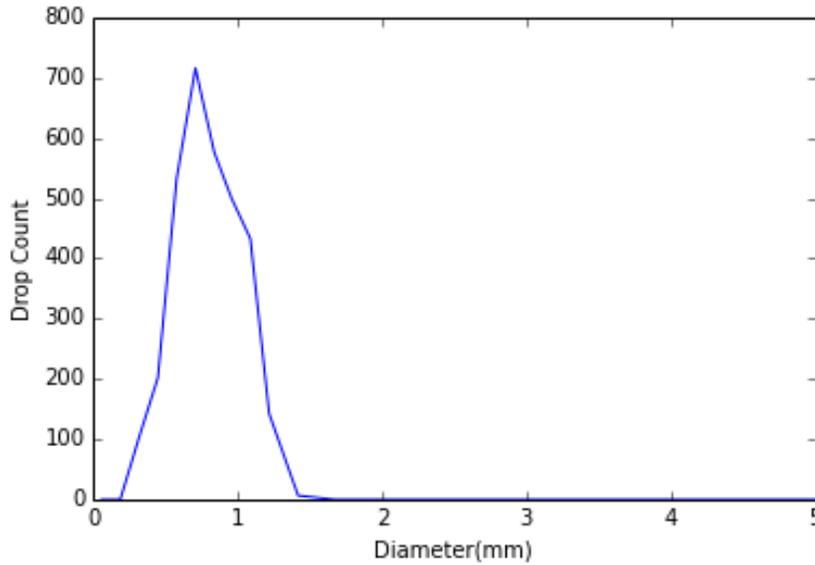
FIGURE 2.4. Single time estimate of a DSD from a disdrometer based upon five minutes of sampling time. The distribution does appear 'gamma-like', with most drops existing below the 1.5 mm mark. Based upon the absence of large drops, it can be inferred the rain-type is stratiform and would result in low reflectivities. The absence of large drops will cause a small $Z + dr$ value.

2.2.2. DROP SHAPE RELATIONSHIP. The shape of raindrops is governed by the contrasting forces of hydrostatic tension and the upward force exerted by the displacement of air caused by the falling motion of the drop. When drops are small they encounter very little air resistance and the hydrostatic tension dominates causing drops to be spherical. As the drops grow in size, they pick up more speed and encounter a larger aerodynamic force. This causes the drops to deform and become oblate spheroids. When the drops grow past a certain point, usually reaching a maximum somewhere below 9 mm, the aerodynamic forces dominate and cause the drop to break up into smaller drops. This deformation of the drops is the primary motivation behind the development of dual polarization weather radar. The increased axis ratio causes increased power returns in the horizontal direction. Measuring the polarimetric variables allows algorithmsto estimate the average axis ratio of the drops which gives a sense of the size of the drops. Combining this with an estimate of the number

18

Raindrops < 0.3 mm ⟶ $Z_{dr} = 0$ dB

$Z_{dr} \approx 1.5$ dB	$Z_{dr} \approx 2.0$ dB	$Z_{dr} \approx 3.6$ dB
2.70 mm	3.45 mm	5.30 mm
$Z_{dr} \approx 4.0$ dB	$Z_{dr} \approx 5.5$ dB	$Z_{dr} \approx 6.3$ dB
5.80 mm	7.35 mm	8.00 mm
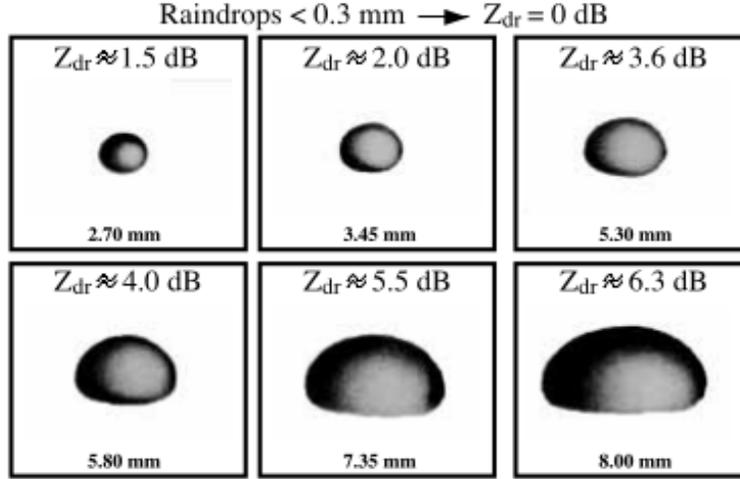
FIGURE 2.5. Drop shapes as diameter increases causing a "hamburger bun" shape. The deformation is caused by the increase in aerodynamic drag on the surface of the drop from passing through the air. As can be seen, as the drop increases in size, the effect gets more pronounced. From [1]

of drops gives important parameters such as rain rate and liquid water content. Figure 2.5 shows an example of raindrop shapes as size increases.

The drop shape relationship (DSR) gives the ratio of the semi-major and semi-minor axes of the raindrop as a function of drop size. Several different models exist and are usually based upon empirical fits to data gathered in wind tunnels or controlled environments. The choice of DSR can have a large influence on the results of any microphysical retrieval algorithm. By changing the variation in the ratio of the major to the minor axis, the DSR changes how the polarimetric variables respond to increases in drop size.

This work will primarily utilize three different models. The first is due to Pruppacher and Beard [4]. It is given as

$$(10) \qquad\qquad\qquad r = 1.03 - 0.062 D_{eq}$$

19

This model will be referred to with the (pb) denotation. Following the work by Prup-pacher and Beard, wind tunnel tests suggested another model attributable to Beard and Chuang [18]

(11) $r = 1.0048 + 5.7 \times 10^{-4} D_{eq} - 2.628 * \times 10^{-2} D_{eq}^2 + 3.682 \times 10^{-3} D_{eq}^3 - 1.677 \times 10^{-4} D_{eq}^4$

which will be referred to as model (bc). More recently, Thurai and Bringi [19] used a hose located on a bridge 80 m above a video disdrometer to determine the drop axis ratio empirically. This model, denoted as (tb) is given as

$$
(12) \quad \frac{b}{a} = \begin{cases} 1 & : D_{eq} < 0.7 \\ -8.5 \times 10^{-3} D_{eq}^4 - 0.1317 D_{eq}^3 + 0.4698 D_{eq}^2 - 0.5165 D_{eq} + 1.173 & : 0.7 < D_{eq} < 1.5 \\ -4.095 \times 10^{-5} D_{eq}^4 + 7.66 \times 10^{-4} D_{eq}^3 - 3.99 \times 10^{-3} D_{eq}^2 - 0.0625 D_{eq} + 1.065 & : D_{eq} > 1.5 \end{cases}
$$

Figure 2.6 shows the axis ratios for each of these models as a function of drop size. At small drop sizes the Thurai and Bringi(tb) and Beard and Chuang(bc) models result in more spherical drops, while the (bc) model predicts slightly more oblate drops for moderate to large drop sizes. These differences manifest themselves in the radar measured parameters, having the largest effect on the polarimetric measurements.

2.3. MOMENT CALCULATION

Given a drop size distribution and a drop shape relationship one can derive the radar measured parameters for a given radar. The fields of interest will be radar reflectivity at horizontal polarization(Z_h), the differential reflectivity(Z_{dr}) and differential phase(Φ_{dp}). Reflectivity is defined to be
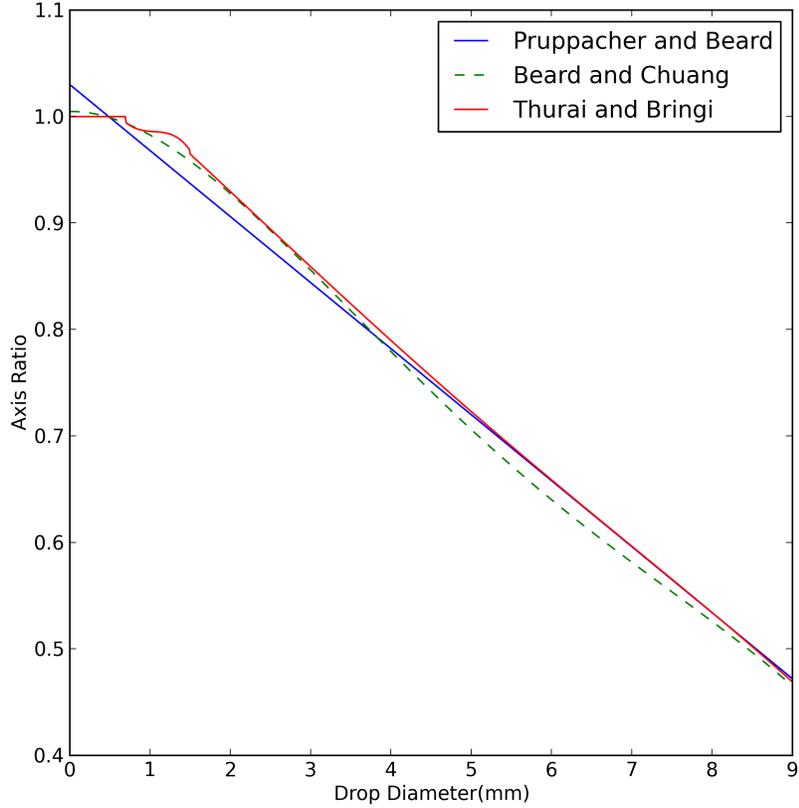
FIGURE 2.6. Drop shape relationship axis ratios. At very large drop sizes, all three relationships converge. The primary difference will be in how the ratio is handled in the 0-4 mm range. It should be noted that the linear equation in the Pruppacher and Beard model causes incorrect results below 0.5 mm.

$$(13) \qquad \zeta_h = \frac{\lambda^4}{\pi^5 |K_p|^2} \int_D |S_{hh}(r, D)|^2 N(D) dD$$

Where $S_{hh}(r, D)$ is the backscattering cross section at horizontal polarization for a drop of diameter D, located at range r, K_p is the dielectric constant of water, and λ is the radar wavelength. Similarly, differential reflectivity is given by

(14)
$$\zeta_{dr} = \frac{\int_D |S_{hh}(r,D)|^2 N(D) dD}{\int_D |S_{vv}(r,D)|^2 N(D) dD} \qquad (7.9a)$$

where $S_{vv}(r,D)$ is the backscattering cross section at vertical polarization. Differential reflectivity is then the ratio of the reflectivities at horizontal and vertical polarizations. It depends directly upon the drop axis ratio and therefore upon the drop size, as these two parameters are related through the drop shape relationship. One can also notice that Z_{dr} is independent of the number of drops.

The final radar measured parameter is the differential phase Φ_{dp}. This gives the cumulative change in waveform phase between horizontal and vertical as the pulse propagates through the atmospheric medium. Differential phase has the rather desirable attribute of being immune to attenuation, as well as a good measure of attenuation. What the radar measures however is not directly the differential phase but rather Ψ_{dp} which takes into account the forward propagation phase change, as well the backscatter differential phase δ_{co}. They are related by

(15)
$$\Psi_{dp}(r) = \Phi_{dp}(r) + \delta_{co}(r)$$

where

(16)
$$\Phi_{dp}(r) = \int_0^r K_{dp}(r) dr + \delta_{sys}$$

K_{dp} denotes the specific differential phase and is given as

(17)
$$K_{dp} = \frac{2\pi}{k_0} \operatorname{Re} \int_D N(D) \left[\hat{h} \cdot \vec{f}(r, D) - \hat{v} \cdot \vec{f}(r, D) \right] dD$$

where δ_{sys} is a constant term representing the built in system phase, and $\hat{f}$ is the forward scattering radar cross section. This is usually an adjustable parameter(or can be compensated for) and so from here on out will be assumed to be zero. Specific differential phase is a spatially localized quantity. It can be hard to estimate however due to the co-mingling of error sources and path integration. One of the most important differences with K_{dp} is that it is a forward scattering quantity. The immunity to attenuation of the parameter allows it's use in correction of path integrated attenuation.

Shown in Figure 2.7 is the equivalent radar measured parameters calculated for the drop size distribution shown previously based upon these equations. In the first panel, the entire day is shown. In the remaining three panels, only the time periods of the day with significant precipitation are shown. Additionally, in the final panel the rain rate for this dataset is displayed. The data has a large amount of variability that is commonly seen in data sampled from real sources.

In the above relationships, the radar parameters were written as a function of their back and forward scattering cross sections S and f. We can evaluate these parameters numerically based on the T-Matrix solution, also called the extended boundary condition method.

There are three fundamental scattering regimes. These are the Rayleigh, Mie, and Optical regimes, each of which exhibits different behavior with respect to electromagnetic radiation. We can normalize the size of a particle with respect to the wavelength of the radiation with the equation
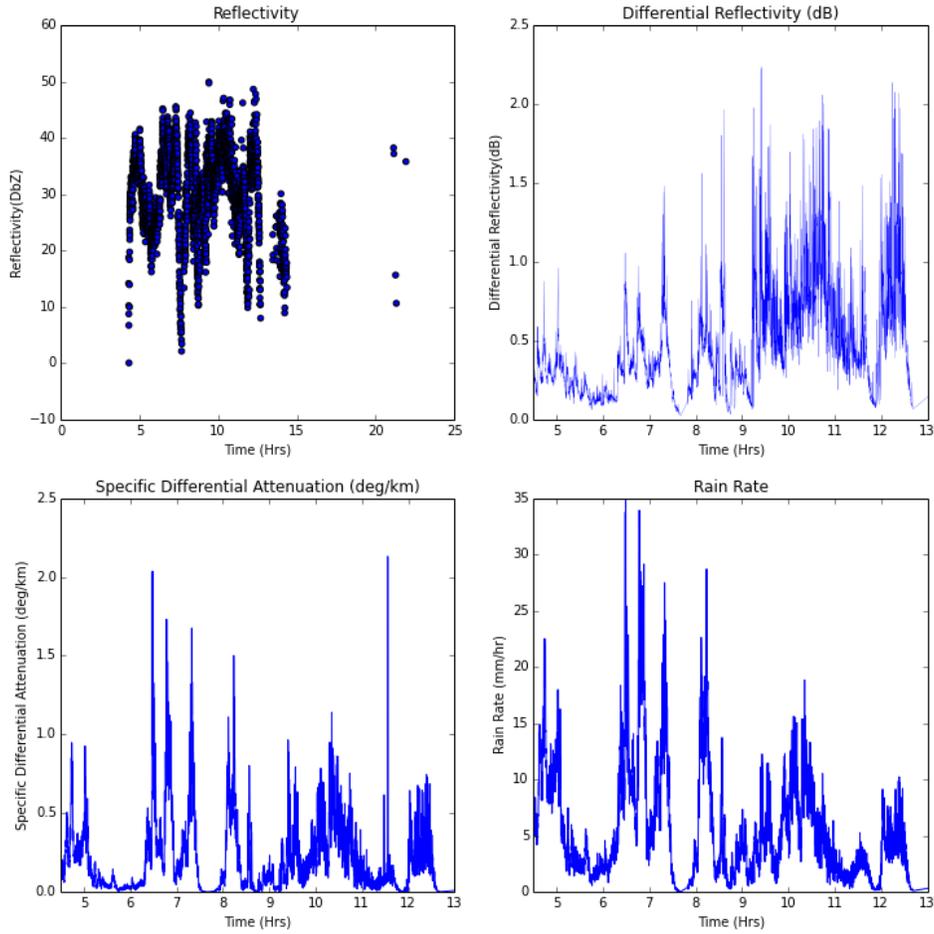
FIGURE 2.7. Results of T-Matrix Scattering for the DSD shown in Figure 2.3. Shown in the top two panels are the two reflectivity measurements. Shown in the bottom two panels are the specific differential attenuation, and the estimated rain rate.

$$\alpha = \frac{\pi D}{\lambda}$$

Where α is our dimensionless size parameter. We can split the different scattering regimes

into three different cases.

(18) $$\alpha \ll 1 \qquad \text{(Rayleigh Scattering)}$$

(19) $$\alpha \approx 1 \qquad \text{Mie Scattering}$$

(20) $$\alpha \gg 1 \qquad \text{Optical Scattering}$$

In each of these regions, the behavior of the particle to electromagnetic waves is different. For S-Band radars, the majority of particles will fall into the Rayleigh scattering regime. As we get to C-band and higher, a larger fraction of the particles will be in the Mie region.

In the Rayleigh regime the backscattered cross section increases roughly as D^6, that is

(21) $$|S_{hh}|^2 \approx D_{eq}^6$$

This is the region most lower frequency radars operate in. One must be careful when using a Rayleigh assumption. The assumption that reflectivity scales as D^6 is used implicitly in much of the radar literature (Oftentimes reflectivity is even defined to be in terms of D^6). While this is correct at lower frequencies, at higher frequencies it can cause a significant portion of mathematical assumptions utilized to be incorrect.

In the Mie region, the electromagnetic wave wraps around the particle and has a reflection in the drop causing a resonance region. These reflections cause constructive/destructive interference. Instead of having a scaling of D_{eq}^6, we get an oscillating behavior. These oscillations dampen as the particle size approaches the optical region. The Mie region represents a significant difficulty in retrievals as it causes the relationship between radar measured parameters and drop size to no longer be invertible as the mapping is no longer one-to-one. Also

the Mie region can cause 'weird' behavior in polarimetric parameters that differ from long held assumptions from using radars at S-Band. For instance, it is possible to get negative specific differential phase values at higher frequencies. We will address these issues more in Chapter 3 and Chapter 5.

2.4. Rain Rate Estimation

An important real parameter derived from weather radar observations with immediate practical importance is the instantaneous rain rate over a given area. Rain rate is typically measured in mm/hr and has huge implications for flooding, hydrology, droughts and severe weather forecasting. If the drop size distribution in a volume of the atmosphere is known, the instantaneous rain rate due to that volume can be calculated as

$$(22) \qquad R = 0.6\pi 10^{-3} \int_D v(D) D^3 N(D) dD; \ \text{mm h}^{-1}$$

where $v(d)$ is the downward velocity of a drop with diameter D. We usually model velocity with a power law, and in the absence of vertical updrafts we assume terminal velocity can be approximated by the formula

$$(23) \qquad v(D) = \alpha D^{0.67}$$

where α varies from 3.6 to 4.2 [1]. If we insert this back into our equation for rain rate we get the relationship

$$(24) \qquad \qquad R = \alpha_r \int_D D^(3.67)N(D)dD; mmh^{-1}$$

Which means that for a homogeneous grouping of drops of diameter D the rain rate will scale as roughly $D^{3.67}$. If we wish to estimate rain rate from the radar measured parameters there are a few options. We can look at an approximation of the relationships between drop size and radar measured parameters to get a feeling for which is most likely to make the best estimate.

$$(25) \qquad \qquad Z_h \approx nD^6$$

$$(26) \qquad \qquad K_{dp} \approx nD^{4.2}$$

If the precipitation in the atmosphere was composed entirely of homogeneous drop sizes at a single time, any of these algorithms would be sufficient. However, due to variability in the drop sizes, reflectivity will exhibit the most variability in the rainfall rate based upon drop size distribution uncertainty as it has the largest exponent scaling. This suggests that the dual polarimetric variables make much better estimators of rain rate, and indeed the literature supports this view [20–22] .

Most commonly, rainfall is estimated using a power fit to empirical data based upon the dual polarimetric variables. The two most popular dual polarimetric estimators are of the form

(27)
$$R(K_{dp}) = \alpha K_{dp}^{\beta}; \qquad \alpha \approx 129 f^{-\beta}, \beta \approx 0.85$$

(28)
$$R(Z_h, Z_{dr}) = c_1 \zeta_h^{\beta} \zeta_{dr}^{\kappa}$$

where ζ_h and ζ_{dr} are the reflectivity and differential reflectivity in linear units. These relationships vary based upon region, time of year, etc. However disdrometers can be used to measure the drop size distribution and simulate the radar data to empirically fit these relationships with some level of accuracy. This temporal and spatial/climatalogical localization improves the performance of these rain rate estimators for a given time and region.

2.5. Bayesian Multi Sensor Fusion

Given multiple measurements of some underlying distribution, Bayesian estimation theory provides a very natural way to fuse the different measurements. In this section we will briefly cover the Bayesian methodology for data fusion.

Given a set of measurements $Y = Y_0, \cdots, Y_1$ of some underlying parameter X then we wish to form an estimate of the probability distribution of X based upon the measured Y's. We can write out Bayes rule

(29)
$$P(X|Y_1, Y_2, \cdots, Y_K) = \frac{P(Y_1, \cdots, Y_K|X)P(X)}{P(Y_1, \cdots, Y_K)}$$

where $P(X)$ represents our prior, the belief before the measurements about the distribution of the variable X and $P(X|Y)$ is the posterior, or the distribution of X after having accounted for the measurement from Y. In this case $P(Y)$ only serves to normalize the

distribution and so we will discard it. $P(Y|X)$ represents the probability of us measuring Y given the underlying variable X. Based on this we can write

$$(30) \qquad P(X|Y_1, \cdots, Y_K) \propto P(X) \prod P(Y_k|X) P(X)$$

This gives us an equation to update our belief of our posterior value based upon new measured values. We will use this to combine the measurements from multiple radars into one estimate with associated density function. This theory is very simple, but very powerful. One large assumption we are forced to make however is the prior distribution. This codifies pre-experimental beliefs about what the correct answer should be. While Bayesian frameworks can update and work beyond this, a common criticism of Bayesian methods is improper choice of priors. Shown in Figure 2.8 is an example layout for a radar network.

2.6. Previous work on Microphysical Retrievals

There has been significant interest in microphysical retrievals over the decades, and accordingly there has been a large number of proposed algorithms. There are several different taxonomies available for these algorithms and we will focus primarily on two dividing criterion: whether the algorithm takes advantage of polarization, and whether the retrieval uses a single or multiple radars. We start by dividing retrievals into single or multiple radar retrievals, and then in each section focus on whether they take advantage of polarization or not. This is just one choice of taxonomies, and there are many other ways of categorizing retrievals.

2.6.1. Single Radar Retrievals. The vast majority of microphysical retrievals have focused on single radar retrievals. Much of the early work in this area focused more on
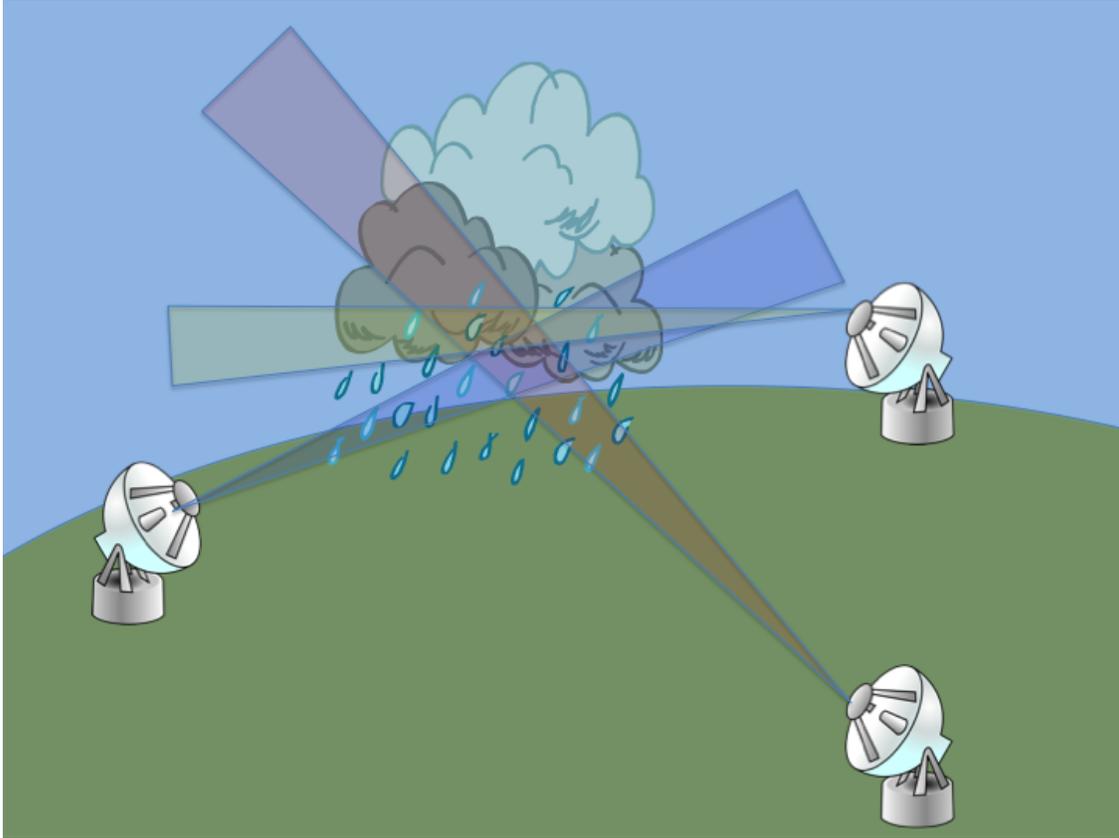
FIGURE 2.8. Conceptual layout of Multi-Frequency Radar Network. Each radar has a different field of view of the storm due to their different locations, while the different frequencies from each radar provide additional information about the size distribution of the particles.

attenuation correction than the microphysics estimation. These two topics tend to be intertwined, as estimation of attenuation is required for the majority of microphysics estimation schemes, and an accurate microphysics estimation helps to refine attenuation estimates. Unfortunately, the reverse is also true. As attenuation contaminates measurements, it makes microphysical estimates more difficult, and to estimate the attenuation to fix this, a good microphysical estimate is needed.

Some of the first work of note in the area was done in 1954 by Hitschfeld and Borden [23]. In their seminal paper they present a multiple pass single polarization algorithm to estimate attenuation. The algorithm briefly works as follows. The rain rate along a ray is estimated

based upon measured reflectivity. They then relate attenuation to the rain rate and use their rain rate estimate to correct the attenuation. Next they use this new reflectivity to estimate rain rate, and thus attenuation, and re-correct reflectivity. This continues until the algorithm stabilizes. Unfortunately, the algorithm only works for small amounts of attenuation, fails at higher frequencies, and does not utilize polarimetric information to get a better rain rate estimation. In addition, while attenuation and liquid water content are indeed closely related, attenuation and rain rate do not share such a close linkage at higher frequencies.

More recently work by Gorgucci et al. [24] developed a heuristic estimation methodology based upon self consistency from scattering simulations between the DSD parameters. This algorithm estimates two of the DSD parameters, D_0 and N_w, based upon the values of Z_h, Z_{dr}, and K_{dp}. It is a three step algorithm. In the first step, for each range bin, an effective slope of the drop shape relationship β is estimated. This is then used in a power law formulation for D_0 and N_w. Based upon these parameters, attenuation is then estimated and corrected in all forward bins. Alternatively, an attenuation correction algorithm can be applied to correct the data, and then the algorithm is used to estimate the parameters after correction. The performance starts to suffer if the assumed μ value is wrong, or a different drop shape relationship is present than the one assumed in the paper. A way around the shape relationship constraint is to scatter for a larger number of shape relationships and μ values. This has the upside of incorporating the variability caused by multiple different shape relationships. The downside is it increases the average error by skewing the distribution away from what would be a correct shape relationship if it were known.

Another algorithm by Le and Chandrasekar uses an iterative dual frequency dual polarization approach for estimating the drop size distribution [6]. It utilizes a two pass approach, where in the first pass it starts at the end of a ray and estimates the drop size distribution

based upon the dual frequency ratio between the two frequencies and the dual polarization parameters in the beams. This iterates backward from the end of the ray to the first range bin calculating the DSD and attenuation at each bin. Then the forward pass calculates the measured radar parameters based upon the drop size distribution. Cost functions representing the different between the measured and estimated retrievals are then minimized to determine the stopping point of the algorithm. While this algorithm has better performance than many others in literature, it is unsuitable for our purposes based upon a few points. First, it heavily relies upon having colocated dual frequency, beam matched radars. This allows it to work in spite of small calibration biases. We however need an algorithm designed to work with single frequency radars, and radars that are not co-located or beam aligned. Although this algorithm does have a single frequency fall-back option, we feel an algorithm designed with the constraints of a single frequency should have the capability to perform better.

Another important algorithm is due to Lim and Chandrasekar. It is an attenuation correction algorithm based upon the Z-PHI method [25]. It uses dual polarization parameters, to estimate attenuation, and constraints the attenuation based upon the total change in Φ_{dp}. It shows very good performance at attenuation correction, but does not attempt to do any sort of microphysical retrieval. This leaves the possibility of separating the two stages, but estimating the error structure of two uncoupled algorithms would be significantly more difficult.

Finally, a newer algorithm due to Yoshikawa et. al. [26] uses a forward variational process to constrain the choice of possible drop size distributions. This algorithm was devised for X-Band radars, although it is not necessarily limited to any particular frequency. It has multiple different phases to address source of different errors. Unfortunately, the algorithm

was never tested on real radar data and so remains a theoretical algorithm. This is the base algorithm we will start with here to extend, and so we will delay extended discussion of this algorithm until Chapter 3.

Many additional algorithms exist for the calculation and retrieval of rain-rate, but as our interest is in the more fundamental DSD parameters we will not address those here.

Most retrieval algorithms that attempt to estimate microphysical parameters here suffer from inability to estimate μ, preferring to focus more on N_w and D_0. This approach has merit as μ has less physical meaning than the other parameters, and mostly focuses to shape the distribution. Additionally, there is a good argument to be made for μ being purely a mathematical artifact and not having any physical meaning.

2.6.2. NETWORK RADAR RETRIEVALS. While a significant amount of work has gone into microphysical retrievals for a single radar, very little work has been done to leverage the mutual information from a network of radars operating at multiple frequencies. Most of the research on networks of radars primarily focuses on "mosaicing" the radar data. This means using same frequency radars and attempting to combine the radar measurements to get a combination of the radar measured fields that better represent the atmosphere by canceling things out like beam blockage. The most commonly used mosaic product is the NEXRAD national mosaic which is a combination of all of the WSR-88D radars used in the CONUS [27–29]. One large flaw with the mosaic method is it's inability to be used with multi-frequency data. As different frequencies return different measurements for the same drop size distribution, any attempt to combine multi-frequencies is inherently flawed. The single frequency constraint utilizes the assumption that the radars are sampling an underlying field with varying sources of error, but that fundamentally they are all sampling the *same* field. At multiple frequencies this assumption breaks down. For instance, in the Mie region, an

X-Band radar will given a different Z_{dr} value than an S-Band radar. Another source of error is the phase measurements. Although K_{dp} is itself independent of direction, most estimation procedures for K_{dp} use a path length approximation that causes K_{dp} to be correlated along a ray artificially, but not laterally along azimuth. While the small errors in reflectivity may be acceptable for more qualitative uses, the fundamental differences in the dual polarimetric parameters at changing frequencies renders multi-frequency mosaic operations on the radar measured parameters useless.

Most mosaic procedures primarily use the mosaic with a max operator to take the highest measurement. A more sophisticated algorithm for single frequency networks is due to Lim and Chandrasekar [30]. This algorithm uses a network of single frequency radars and relies on the (very good) assumption that the specific attenuation at a given location will be the same no matter the direction of propagation. By doing this, it can use multiple looks from different radars to constrain the attenuation at every bin, yielding a accurate attenuation correction algorithm. By using multiple radars to estimate the attenuation, estimating the intrinsic radar parameters will most likely be possible. This algorithm does not however attempt to estimate the microphysical parameters in any way. It also fails to address the multiple frequency case. Nonetheless, this algorithm does lend support to multiple independent views increasing the performance of a network retrieval over a single radar retrieval.

While not necessarily multi-radar, one promising development for multi-frequency retrievals is the use of the dual frequency ratio (DFR) for retrievals. The dual frequency ratio is the ratio of retrieved horizontal power at two different frequencies. The ratio of the two frequencies gives an idea of how where in the scattering spectrum the bulk of the drop sizes are, allowing you to restrict D_0 in particular[31–33]. An example of DFR between K_u and K_a frequencies for a large amount of data is shown in Figure 2.9. A very pronounced (although

not easily characterized) trend can be seen. Dual frequency ratio gives another piece of information in constraining the microphysical parameters. Indeed, the new Global Precipitation Measurement satellite constellation uses DFR on it's Dual Frequency Precipitation Radar (DPR) to improve the retrievals.
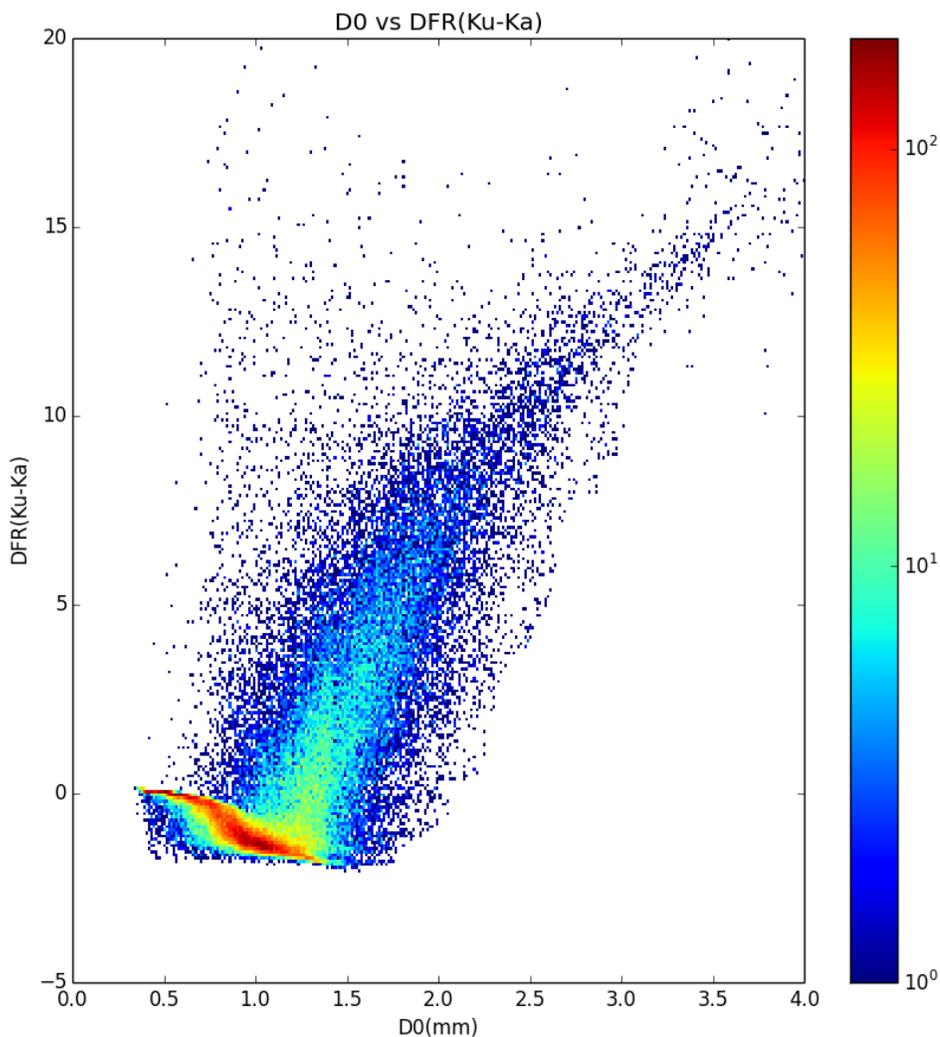
FIGURE 2.9. (Ku-Ka) Dual Frequency Ratio as a Function of D_0. This shows the relationship between reflectivities at the two frequencies is not a linear relationship. The parameters are generated from the IFloodS disdrometer dataset consisting of 17,000 minutes of binned histograms.

2.7. Error Structure in Microphysical Retrievals

There are several sources of error in microphysical retrievals in a network environment. Some of the more dominant errors are listed in Table 2.1. These error sources can be split into three different categories. The first category is instrument error. This is the lowest level of errors, caused by flaws in the instrumentation, or in it's operation. The most significant error in this category is the system bias error type. This is the calibration bias caused by incorrect estimates of the different power and loss parameters. This can result from incorrectly measuring these parameters, or more commonly from changing parameters due to environmental effects such as aging, and changing temperatures. These errors manifest themselves as a constant offset applied to the measured powers. The primary two biases that will affect the retrieval is the bias in returned power at horizontal polarization, and the ratio of this with the bias in returned power at the vertical polarization. Two related errors are the instrument precision, and the instrument noise. These take the form of Gaussian zero mean random noise signals added to the measurements. These error types are common to all observational instruments. The final instrument error of note is the limited sensitivity of the instrument. At further ranges, or lower returned powers, the measurement of returned power will fall below the noise floor of the receiver. When this happens, we can not reliably measure the signal. This has two primary effects on the radar measurements as it affects the retrieval. The first is the imposition of a maximum range. Beyond a certain range, signals no longer return enough power to be measured by the radar. The second effect primarily effects higher frequency radars. At higher frequencies, attenuation is more severe. This reduces the received power by the radar, causing more of the signal to fall below the noise floor. In this way, the attenuation combined with the limited sensitivity cause portions of the return to go "extinct".
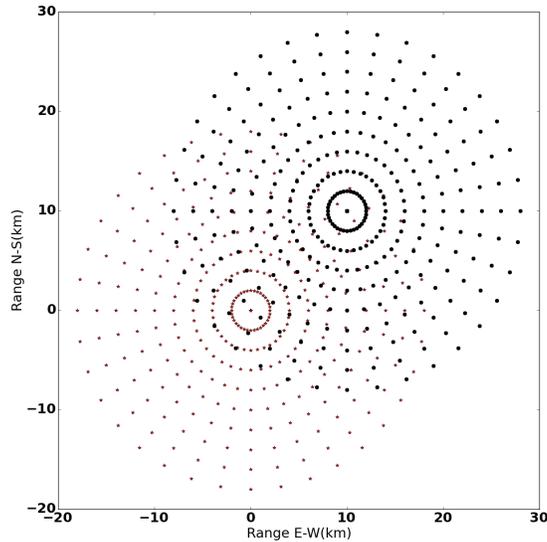
FIGURE 2.10. Non-Coincident Sampling on a Polar Grid. No unique choice of origins will cause the points to line up, making the comparison of non-coincident grids challenging.

The second category is observational sources of error. This is a large category, but the most important ones to this work are attenuation from propagation, beam and time mismatch between radars, and non uniform beam filling. Attenuation is the loss of power caused by the absorption of energy from the hydrometeors. This is primarily governed by the extinction cross section of the drop. Larger drops, and more numerous amounts of drops, increase the attenuation experienced by the electromagnetic ray. As the frequency of the electromagnetic waves increase, the attenuation of the waves increase as well. This makes attenuation a much more significant problem for higher frequency radars, with X-Band and above suffering particularly hard. Algorithms exist to help correct attenuation errors in many different cases, but it is currently still very far from being solved for the general case. For the network retrieval, beam and time mismatch are significant sources of error. The radar measurements are taken on a spherical grid, and if two radars are at two different locations, the grid points will not line up. An example of this is shown in Figure 2.10. In three

TABLE 2.1. Many of the different sources of error in the retrieval process. These are broken down into three different categories based on where they are introduced in the retrieval process.

Sources of Retrival Error		
Instrument	Observational	Retrieval
Instrument Precision	Attenuation	Improper Prior
System Biases	Beam Mismatch	Gridding Errors
Instrument Noise	Time Mismatch	Network Assumptions(IID)
Sensitivity Limits	Beam Block	Improper DSR
	Non-Uniform Beam Filling	Improper DSD Parameterization
	Reflectivity Gradients	Erroneous Covariance Estimates

dimensions, this effect is even more pronounced. This means that each radar is measuring a slightly different thing. Due to the time it takes to scan a volume, these measurements are likely not even at the same time. With proper scanning strategy design, these errors can be minimized somewhat, but never entirely eliminated. As the beams are not along the exact same orientation, the effects they suffer from propagation are different as well. Two errors in this set are closely related. Non-Uniform Beam Filling(NUBF) and Reflectivity Gradients are when the contents of a voxel measured by the radar is not spatially homogeneous within the voxel. While this is strictly always the case, the normal assumption in radar is that the variability is low enough that it does not induce significant errors. This is usually not that outlandish of a suggestion, but at the edges of storms this can cause measurement errors. The final error source is beam blockage. This is where the edge of a beam intersects with some structure such as a tree, tower, or house. This causes the measurements to be erroneous. Beam blockage has a larger effect on power measurements(as the power of the beam can be partially, or wholly blocked), whereas phase measurements suffer less significantly.The retrieval in this work will correct attenuation, and attempt to account for the time and beam mismatch effects. The other errors are unfortunately outside the scope of this work.

The third category is errors from the retrieval algorithm itself. These errors can be split into improper assumptions, and procedural errors. The improper assumptions includes choosing the wrong shape parameter for the retrieval, wrong temperatures, and incorrectly estimating the covariances of the parameters. This includes extra assumptions of independence that may not fully be warranted, as well as assuming a normalized gamma distribution when the distribution may in fact be better represented by another form. Unfortunately, these assumptions must be made to keep the solution to the problems tractable. Otherwise, there are far too many free parameters for the constrained measurement set that the radar records. The procedural errors are errors introduced by manipulations within the retrieval process. One of the largest of these is gridding errors from the interpolation step in the network retrieval stage. This is interpolation error when comparing non-coincident points. Another source of error is imprecision in the scattering operations, and only being able to account for liquid precipitation.

These errors all serve to reduce the accuracy of the microphysical estimation. In the following chapters we will demonstrate how corrections for some of these errors have been incorporated into the algorithms themselves, and hopefully show that the rest of the errors that we are unable to fix do not serve to degrade the microphysical estimation significantly.

2.8. Spatial and Temporal Variability of Rain Microphysics

Finally, a factor that will play a large role in the interpretation and quality of results for a microphysical retrieval is the spatial and temporal variability of the precipitation microphysics.

What is needed is a good way to link the variability of microphysics spatially and temporally. One of the earliest works on this topic was due to Taylor [34] and resulted in the

"Taylor Hypothesis". The Taylor Hypothesis states that the atmosphere is made up of horizontal sheets, stacked vertically, and that as they fall, they undergo an affine transformation. This means no major morphology changes. While this is easily disproven over a large enough vertical section of the atmosphere, if analysis is restricted to areas under the melting layer, it can provide a useful simplification. The primary implication of this is as follows. If the velocity of a storm is known, then a comparison at different times can be made by using the spatial variability to move backwards. More specifically, the autocorrelation function with respect to time has a link to the autocorrelation function with respect to space. This gives us some hope for comparing aerial retrievals with ground truth measured at ground level with a disdrometer or rain gage.

Most research on microphysical spatial variability has revolved around rainfall rates[35, 36]. While this work is interesting, it does not give us enough information about the variability of the underlying microphysical parameterizations. More recently there has been some work on the variability of the microphysical parameterization itself [37], although this particular study used a single polarization radar and some large assumptions about microphysics, limiting the accuracy and generalizablity. Less research has been done using dual polarization radars [38], although some research has been done on the spatial variability of the radar measurements themselves[39].

Fundamentally, truly characterizing the variability of the microphysics requires multiple radars, and a network of ground instrumentation. This is very expensive, and so has not received as much attention as it warrants. This situation was addressed in 2012 by the IFloodS campaign, which provided a very robust dataset with respect to microphysical variability. The spatial variability results from this campaign will be addressed in Chapter 5.

2.9. Summary

This chapter laid out the framework needed to understand the rest of the discussion in following chapters. In particular it discussed what a microphysical parameterization is, and how to use the radar to measure indirect effects of the microphysical situation on the radar wave. The physics behind precipitation formation was covered, to explain why a distribution is required, and why the distribution is so much more complex than the radar measurements see to indicate. Next single radar retrievals and multi-sensor Bayesian fusion based network retrievals were discussed. Next, a discussion of prior work in microphysical retrievals, both at the single and multi radar level was detailed. Following this, a breakdown of the different sources of microphysical errors in the retrieval process was outlined before finally moving on to a discussion of the spatial and temporal variability of precipitation microphysics and the constraints and effects these will have on the radar retrieval. The following chapters will lay out a methodology and algorithm for multi-frequency radar retrievals, and then address results and spatial and temporal variability analyses.

CHAPTER 3

Single Radar Retrieval

"Begin at the beginning," the King said gravely, "and go on till you

come to the end: then stop."

— Lewis Carroll, Alice in Wonderland

This chapter addresses the issue of retrievals when only a single radar is present. The retrieval for a single radar will form the first building block of the network retrieval covered in the next chapter.

A radar network is made up of a constellation of single radars operating at possibly different frequencies. The first step in our microphysical retrieval will be to work data form each radar individually to correct for sources of error and find a microphysical estimation to feed into our network retrieval. Attenuation and most sources of radar errors are localized to a single radar, making this the best place to correct these errors before attempting to combine the output of each radar into the network retrieval. In this chapter we explain and demonstrate a single radar microphysical retrieval algorithm capable of retrieving the drop size distribution from dual-polarized radar data at arbitrary frequencies based upon work by Yoshikawa et al[26]. This algorithm will form the core of our network retrieval algorithm. We will discuss some of the background, and then lay out the mathematical formulation for the algorithm. Finally, we will look at some preliminary results, and point out weaknesses in the current algorithm before discussing proposed additions to the algorithm.

3.1. Background

Given a small volume of atmosphere located at range r from the radar, we can describe the distribution of raindrops in the volume using the normalized gamma distribution discussed in chapter 2. In particular, assume that the distribution of the raindrops is given by

(31)
$$N(D; N_w, D_0, \mu) = N_w f(\mu) \left(\frac{D}{D_0}\right)^{\mu} \exp\left(-(3.67 + \mu)(\frac{D}{D_0})\right)$$

with

(32)
$$f(\mu) = \frac{6}{3.67^4} \frac{(3.67 + \mu)^{(\mu + 4)}}{\Gamma(\mu + 4)}$$

where N_w is the intercept parameter of an equivalent exponential distribution with the same water content and median drop diameter and is related to N_0 by

(33)
$$N_0 = N_w f(\mu) D_0^{-\mu}; \qquad \mathrm{mm}^{-1-\mu} \mathrm{m}^{-3}$$

D_0 is the median drop diameter, and μ is the shape parameter. This formulation is covered in more detail in Chapter 2.

Given a set of three measurements, the distribution can be parameterized with the triplet:

(34)
$$X(i, k) = (D_0(i, k), N_w(i, k), \mu(i, k)).$$

This gives the conditions in the volume given by the shell segment from r_i to r_{i+1}, and angular extent ϕ_k to ϕ_{k+1} and θ_k to θ_{k+1}

Along a radial direction from the radar, break the ray up into cells given by equal angular width, and radial range from some central point, then form a vector of triplets as

(35)
$$\mathbf{X}(k) = (X(0, k), X(1, l), ... X(N, k))$$

This vector represents a simple microphysical parameterization of the precipitation present in this volume of the atmosphere. Now place a radar at the center of this radial, whereit would

typically measure a triplet of radar moments comprised of the horizontal reflectivity(Z_h), the differential reflectivity(Z_{dr}), and the differential phase(Φ_{dp}). Again form a vector of these triplets. If each element of the vector is given by

$$(36) \qquad \mathbf{Y}(i,k) = (Z_h(i,k), Z_{i,dr}(k), \Phi_{dp}(i,k))$$

then the associated ray of measurements is given by the vector

$$(37) \qquad \mathbf{Y}(k) = (Y(0,k), Y(1,k), \ldots, Y(N,k))$$

$\mathbf{Y}(k)$ arises from the drop size distribution $\mathbf{X}(k)$ and can be represented by

$$(38) \qquad \mathbf{Y(k)} = \mathcal{F}\mathbf{X}(k)$$

where $\mathcal{F}$ is a operator defined by the equations (13, 14, 16).

Unless otherwise stated, we will primarily treat rays as independent, and so we will drop the k subscript for now unless explicitly required.

Then we can define a microphysical retrieval as an estimation of the inverse operator $\mathcal{F}^{-1}$ that satisfies

$$(39) \qquad \mathbf{X} \approx \mathcal{F}^{-1}\mathbf{Y}$$

Unfortunately $\mathcal{F}$ is highly nonlinear, and non-invertible.

It should be noted that while this is the definition used for radar based retrievals, different fields define the microphysical parameterization differently. In particular, atmospheric modelers tend to incorporate temperature, species identification, and several other factors. We, however, will just attempt to retrieve the parameters of the drop size distribution.

Several confounding factors make the retrieval more significantly more difficult. To start with, we must account for radar measurement errors, and so our radar measurement triplet for a single gate becomes

(40) $$Y(k) = (Z_h(i) + \sigma_{Z_h}(i), Z_{dr}(i) + \sigma_{Z_{dr}}(i), \Phi_{dp}(i) + \sigma_{\Phi_{dp}}(i))$$

where each of the σ_i's are zero mean random variables with noise power given by $\gamma_{Z_h}, \gamma_{Z_{dr}}, \gamma_{\Phi_{dp}}$. This measurement error can be reliably assumed to be constant in each bin over the course of a sweep and known in advance. In addition to the measurement error, we must factor in attenuation causing additional errors in measured powers. The radar also measures not Φ_{dp}, but at each bin a certain amount of backscattered differential phase given by δ_{hv}. The sum of these quantities is called Ψ_{dp} and is the actual quantity measured by the radar. This means the triplet we actually measure is given by

(41) $$\hat{Z}_h(i) = Z_h(i) + \sigma_{Z_h}(i) + \eta_h - \sum_{t=0}^{i} A_h(t)$$

(42) $$\hat{Z}_{dr}(i) = Z_{dr}(i) + \sigma_{Z_{dr}}(i) + \eta_{dr} - \sum_{i=0}^{k} A_{dr}(i)$$

(43) $$\hat{\Phi}_{dp}(i) = \sum_{t=0}^{i} \Phi_{dp}(i) + \sigma_{\Phi_{dp}}(k) + \delta_{hv}(k)$$

where

$$\eta_h = \text{Horizontal Calibration Bias}$$

$$\eta_{dr} = \text{Differential Calibration Bias}$$

$$\delta_{hv}(k) = \text{Backscatter Differential Phase}$$

These equations illustrate several important points. First, errors in estimating attenuation and differential phase propagates through to all future bins. Second, the differential phase has a backscatter component that adds a noise-like response. If we want to calculate the specific differential phase component, which is given as the range derivative of differential phase, then the operation will amplify the effects of the backscatter phase. This means we must deal with backscatter phase to get a good estimate of specific differential phase.

We can take great care in the calibration of the instruments and minimize the calibration bias to 1dBZ and 0.2 dB for horizontal and differential reflectivity respectively. This will still however cause some issues in the retrieval.

3.2. Algorithm Formulation

Having covered the basics of the retrieved signal and in particular, addressing the sources of noise that can contaminate the measurements, we lay out an algorithm for the retrieval of microphysical parameters using data from a single radar, at an arbitrary frequency. This builds upon work by Yoshikawa et. al. [26].

3.2.1. Mathematical Formulation. The algorithm formulation we are going to use is a greedy variational method. The algorithm attempts to find the parameterization of the drop size distribution that best minimizes the cost function discussed below for each individual bin.

We will use the following notation. We will represent the intrinsic radar parameters(those unaffected by attenuation, calibration biases, or other sources of nosie) as $Y(i)$ for a range bin i. The radar measured variables(those directly from the radar, with all sources of noise) will be represented by $\hat{Y}(i)$, and our reconstructed estimate of the radar measured parameters

will be given by $\tilde{Y}(i)$. Similarly we will represent the true drop size distribution by $X(i)$ for the i^{th} range bin, and our estimate of the drop size distribution in the i^{th} bin as $\tilde{X}(i)$.

Our goal in this stage then is given a set of radar measurements $\hat{Y}(i)$, to recover an estimate of the drop size distribution $\tilde{X}(i)$ for each of the N range bins in a ray. Put another way, for every i between 1 and N, we wish to find

$$(44) \qquad \underset{\tilde{X}(i)}{\arg\min} \left(\hat{Y}(i) - \mathcal{F}\tilde{X}(i) \right)$$

To find this we are going to set up a maximum likelihood based estimator. We start by inserting the maximum likelihood estimate as

$$(45) \qquad p\left(\hat{Y}(i)|X(i) \right) = \mathcal{N}\left(\hat{Y}(i)|\mathcal{F}\hat{X}(i), \mathbf{\Sigma} \right)$$

Which can be rewritten as

$$(46) \qquad p\left(\hat{Y}(i)|X(i) \right) = \mathcal{N}\left(\hat{Y}(i)|\tilde{Y}(i), \mathbf{\Sigma} \right)$$

Where $\mathcal{N}(x, \mu, \Sigma)$ is the normal distribution with mean μ and covariance matrix Σ. Σ is given by

$$
(47) \qquad \Sigma = \begin{pmatrix} \sigma_{Z_h} & 0 & 0 \\ 0 & \sigma_{Z_{dr}} & 0 \\ 0 & 0 & \sigma_{K_{dp}} \end{pmatrix}
$$

This arises by assuming that the covariances of the estimates are independent. We assume that the covariance parameters are known based on empirical measurements. From here we form the likelihood function.

$$
(48) \qquad -\ln p\left(\hat{Y}(i)|X(i)\right) \approx \left(\hat{Y}(i) - \tilde{Y}(i)\right)^{T} \Sigma^{-1} \left(\hat{Y}(i) - \tilde{Y}(i)\right)
$$

Which reduces down to

$$
\begin{aligned}
(49) \qquad J(X(i)) &= -\ln p\left(\hat{Y}(i)|X(i)\right) \\
&\approx \frac{\left(\hat{Z}_h(i) - \tilde{Z}_h(i)\right)^2}{\sigma_{Z_h}^2} + \frac{\left(\hat{Z}_{dr}(i) - \tilde{Z}_{dr}(i)\right)^2}{\sigma_{Z_{dr}}^2} + \frac{\left(\hat{\Phi}_{dp}(i) - \tilde{\Phi}_{dp}(i)\right)^2}{\sigma_{\Phi_{dp}}^2}
\end{aligned}
$$

We then minimize this using any preferred optimization algorithm. For the purposes of this research we have primarily stuck with the "Sequential Linear Least Squares", a constrained optimization technique.

3.2.2. PHYSICAL REALIZABILITY. One of the issues with microphysical retrievals is that most algorithms assume some kind of independence between the two parameters D_0 and N_w. In actuality the two tend to be highly linked. This makes sense intuitively. Consider a fixed water content in the atmosphere. If average drop sizes are large, there will be less of these drops as they compete for the same water content. If there is a large number of drops, then

we would expect this to be a limiting function on the upper size of the drops. We can verify this empirically as well. In Figure 3.1 the two parameters D_0 and N_w are plotted against each other for the IFloodS field campaign.
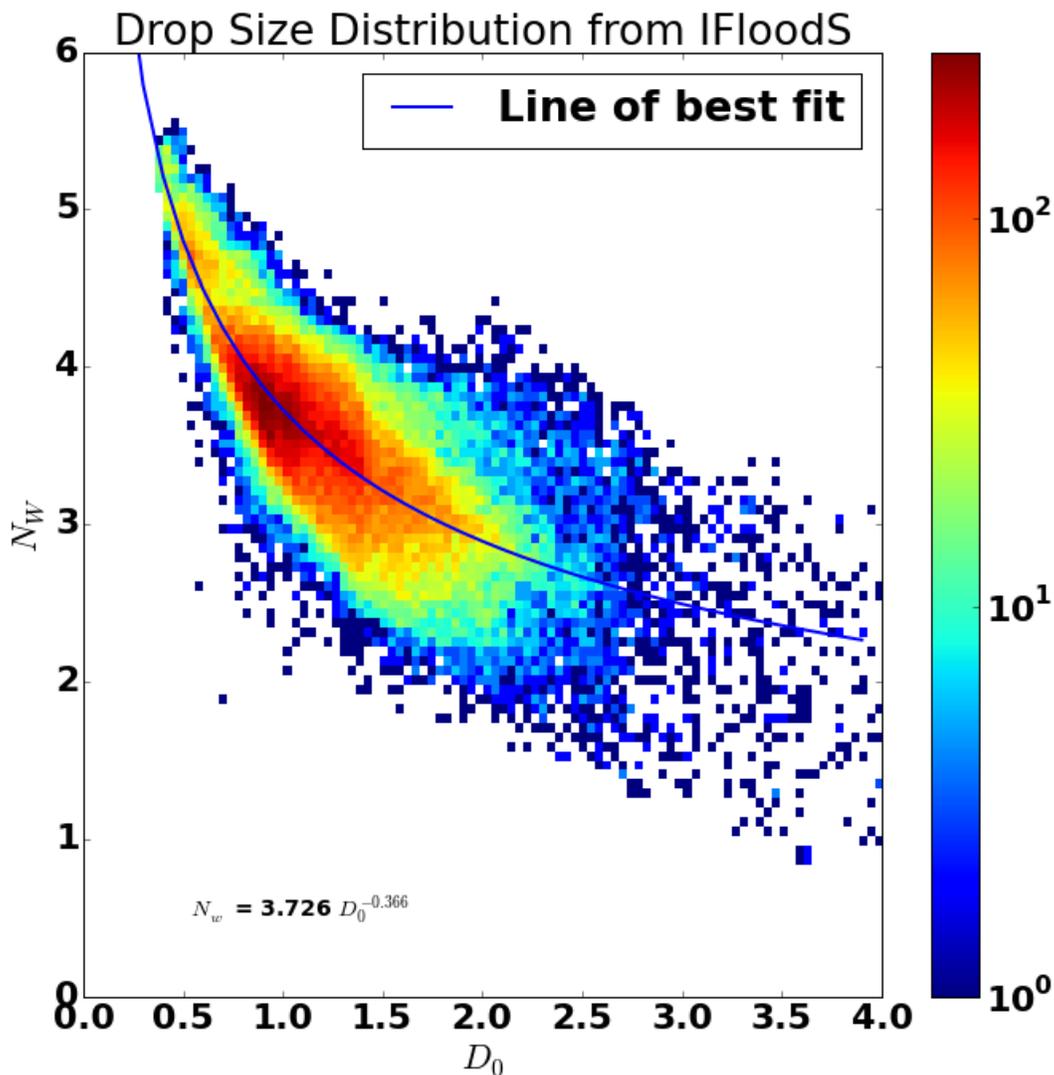
FIGURE 3.1. DSD parameters from IFloodS. This figure shows the relation of the main two parameters D_0 and N_w. The region to which the parameters are constrained can be seen very clearly. N_w is shown on a log scale, as will be the case in most of this work.

We see that the behavior does indeed follow our intuition. This behavior seems to hold true with very few outliers. We would like for our algorithm to preferentially select drop size

distributions from within this region. With this in mind we can introduce a new cost term as

$$(50) \qquad J_{DSD}(X(i)) = \frac{(N_w - \hat{N}_w(D_0))^2}{\alpha_G^2}$$

where $\hat{N}_w$ is an estimator for N_w based upon D_0. In the case of our IFloodS data, our estimator is given by

$$(51) \qquad \hat{N}_w(D_0) = 3.726 D_0^{-0.366}$$

The parameter α_G is tuned empirically to prevent this cost term from becoming too dominant. If we add this term into our cost function, we end up with the new cost function

$$(52)$$
$$J(X(i)) = \frac{\left(\hat{Z}_h(i) - \tilde{Z}_h(i)\right)^2}{\sigma_{Z_h}^2} + \frac{\left(\hat{Z}_{dr}(i) - \tilde{Z}_{dr}(i)\right)^2}{\sigma_{Z_{dr}}^2} + \frac{\left(\hat{\Phi}_{dp}(i) - \tilde{\Phi}_{dp}(i)\right)^2}{\sigma_{\Phi_{dp}}^2} + \frac{\left(N_w - \hat{N}_w(D_0)\right)^2}{\alpha_G^2}$$

3.2.3. BACKSCATTER COMPENSATION. One of the major confounding factors, already discussed is the backscattered component that we measure when attempting to measure Φ_{dp}, given as δ_{hv}. This is the phase change the wave encounters as it is reflected back towards the antenna by a hydrometeor. As δ_{hv} is such a big issue, it behooves us to attempt to estimate it. We can model the backscatter differential phase of a single drop, and indeed of full drop size distributions. We know that δ_{hv} depends on not only the wavelength of the radar, but also the temperature, and the shape and size of the drop. We can plot δ_{hv} as a function of drop diameter for different frequencies and temperatures. This is show in Figure 3.2.
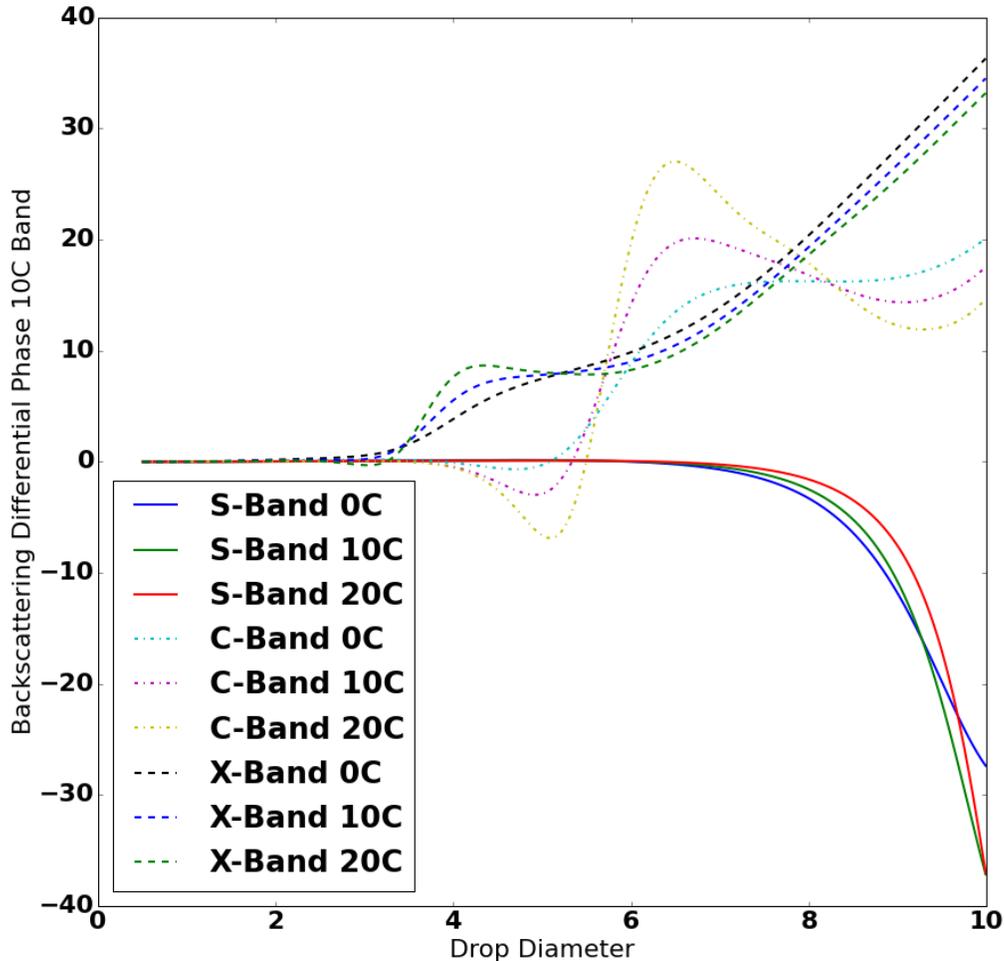
FIGURE 3.2. δ_{hv} for multiple frequencies and temperatures. The results dependent very highly on both temperature and operating frequency. Note that the backscatter differential phase can be negative as well as positive. Three temperatures between 0 C and 10 C are shown.

For small to medium drop sizes, *generally* higher frequencies means a higher δ_{hv}. In the resonance region around 5 mm we see that that δ_{hv} is actually higher at C-Band than at X-Band. C-Band has a higher backscatter differential phase in the region 6-8 mm. This makes up a very small percentage of drops in the real world. It does show that C-Band will be particularly susceptible to large drops in the atmosphere.

We can account for this in the algorithm by modifying our Φ_{dp} calculation to estimate the backscatter differential phase term. This is left as an option, as some radars report a filtered version of Φ_{dp} instead of the true Ψ_{dp} that they measure.

3.2.4. ALGORITHM IMPLEMENTATION DETAILS. This subsection will discuss some of the lower level implementation details of our implementation of the mathematical framework from above. Fundamentally the algorithm is comprised to two inter-dependent parts. The first is the forward scattering operation displayed in Figure 3.3 that implements equations 13, 14, and 16. The second part of the algorithm is the minimization of the forward variational error as shown in Figure 3.4.

Both of these algorithms require certain *a-priori* assumptions. In particular, the choice of drop shape relationship, drop temperature, and shape parameter. Unless otherwise specifies, the algorithm implementation defaults to using the Beard and Chuang [18] drop shape relationship at a drop temperature of 10° C. The shape parameter μ of the DSD will be taken as a default of three.

The scattering operation itself uses an extension of a Python wrapper called PyTMatrix [40] around a Fortran T-Matrix Scattering Library [10]. The parameters for creation of the binned normalized gamma distribution DSD model are listed in Table 3.1.

The optimization process itself is very computationally intensive. To reduce this, a downscaling factor is applied to reduce the radial resolution of the data by a factor of d_f. This factor is variable, and depends upon the spatial variability of the radar data, as well as the underlying radial sampling resolution of the radar. Test cases in this paper most often use a downscaling factor of 2 or 4 depending on the original gate spacing. The downscaling operation is applied using a simple mean filter. Empirical tests showed that the downsampling had very little effect on the overall retrieval.
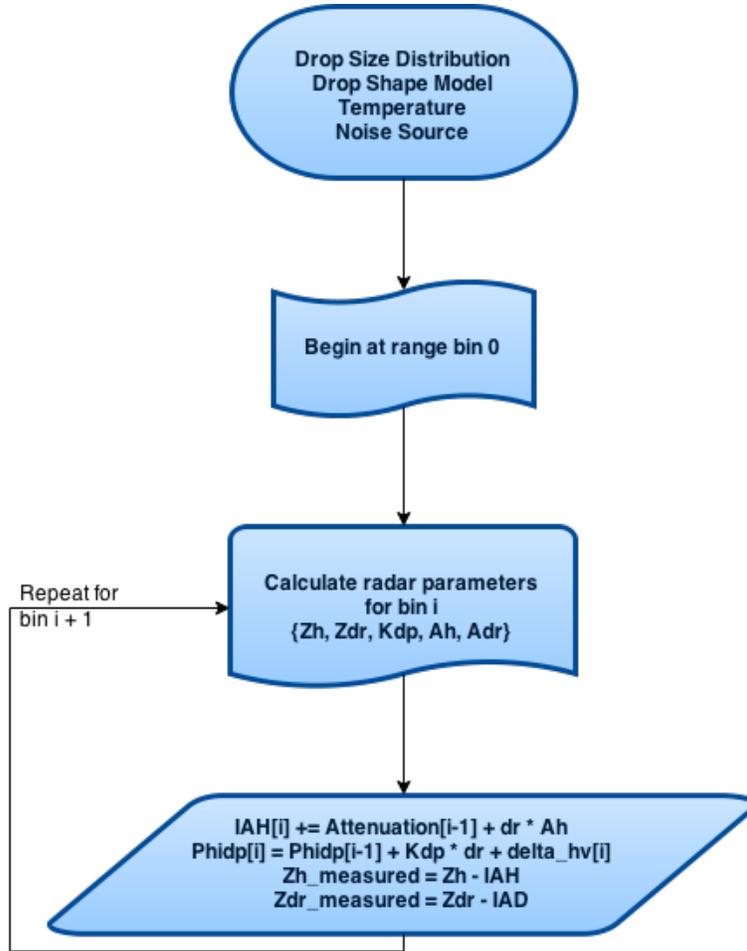
FIGURE 3.3. Flowchart of the forward scattering process. The process starts at the closest range bin to the radar and then proceeds radially outward. The result of the process is the conversion of a set of disdrometer parameterizations into a set of radar measured parameters.

A suitable minimization algorithm should not make a significant difference in the quality of the retrieval, however it can have a large effect on the computational complexity. For the single radar retrieval, a Truncated Newton Conjugate Gradient[41, 42] method is used. This is a bounded retrieval algorithm which showed good computational performance on this retrieval. The DSD parameters are bounded to the region $(0, 6)$ in mm and dB respectively. The retrieval has a tolerance of $5.0e^{-08}$ or a maximum of 100 iterations as it's stopping condition.

TABLE 3.1. PyTMatrix Scattering Kernel Parameters. These parameters are used to create a binned model of a normalized gamma distribution for processing by the T-matrix code.

Scattering Parameter	Value
DSD Resolution	0.26 mm
D_{max}	8 mm
Orientation	20° Gaussian Canting
Orientation Averaging	Canting Fixed Average
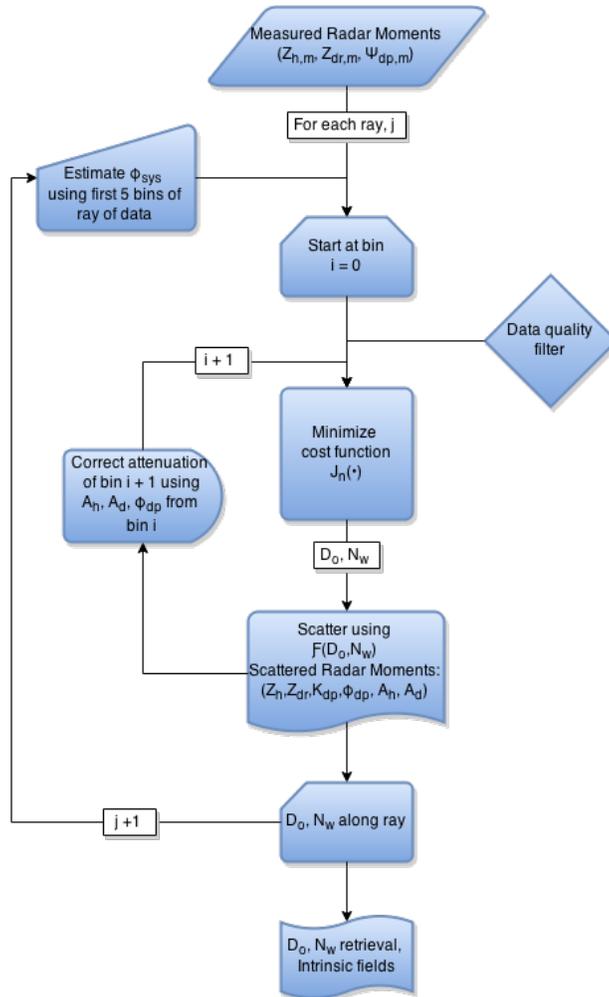
FIGURE 3.4. Single Radar Retrieval Implementation Flowchart. This shows the algorithm flow for the single radar retrieval as it is implemented in this work.

3.3. Results

Fundamentally, one of the big difficulties in deriving retrieval algorithms is also the most challenging limitations for evaluating the accuracy of the results of any retrieval; A lack of ground truth data. What this means is that without actual DSD measurements in a cloud, we must either rely on using synthetic data, or use more indirect methods.

Synthetic data can be generated by assuming a drop size distribution(DSD), drop shape relationship, temperature, and using the T-Matirx method [10] to solve the forward scattering problem as addressed in Chapter 2. Of particular importance is how to determine the distribution of DSD parameters. It is important to choose parameters that are physically realizable. Also we would like for the spatial variation to be representative of what is to be expected from real world precipitation. Using data from disdrometers constrains the choice of parameters to a physically realizable space. To account for the spatial variation expected from real storms we use the method due to Chandrasekar et. al. [2]. This method uses a retrieval algorithm to calculate an expected drop size distribution field. We can then constrain the choice of DSD's to within our physically realizable space based on disdrometer data. Once we have a drop size distribution field, we can calculate the radar equivalent parameters as well as applying different sources of error and noise.

Using synthetic data to validate our algorithm runs the risk as with all synthetic data, of encoding the beliefs and assumptions of our model, into our data, causing the tests to not represent a truly independent verification. The largest of these assumptions is that the DSD is accurately represented by a gamma model. Although the gamma model is one of the best models we currently have, recent research by Ignaccolo and Michele [43] has shown that the gamma model fails to accurately represent the microphysics in over 55% of the authors

captured data. Additionally, we often model the noise sources as independent, though this is most likely not true for all but the most basic of sources.

While this does not invalidate using synthetic data to validate a model, what it does do is inform us that we need more than just synthetic data to truly build confidence in any retrieval framework. Using real data however brings it's own set of challenges. In particular, we do not have a ground truth field! If we don't have ground truth, what do we compare our retrieval with to estimate it's accuracy? A common choice in the literature is to use a derived rainfall field from the retrieval, along with ground based rain gages. Another method is to use the retrieved results and compare them to the microphysical parameters from ground disdrometers. An additional method is to compare the scattered results from the retrieved parameters with the actual radar measurements.

In the following sections we will show results using several different evaluation methodologies to demonstrate the accuracy of our approach.

3.3.1. SIMULATED DATA RESULTS. Currently, the only way to get an accurate measure of retrieval accuracy with respect to the actual drop size distribution is to use synthetic data. In this section, we will look at the results of the retrieval on synthetic data sets . This allows us to characterize retrieval results as the different sources of distortion change and attempt to estimate trends in behavior. The dataset we will use was generated by using a retrieval algorithm based on Gorgucci [9]. The original fields were measured with the CSU-CHILL radar. The resulting DSD parameterizations were constrained using several months of disdrometer data obtained during GPM field campaigns to a physically realizable space. This gives us the microphysical field shown in Figure 3.5.

Scattering calculations require several parameters be fixed *a-priori*. The default values for these parameters, unless otherwise specified is shown in Table 3.2.

56
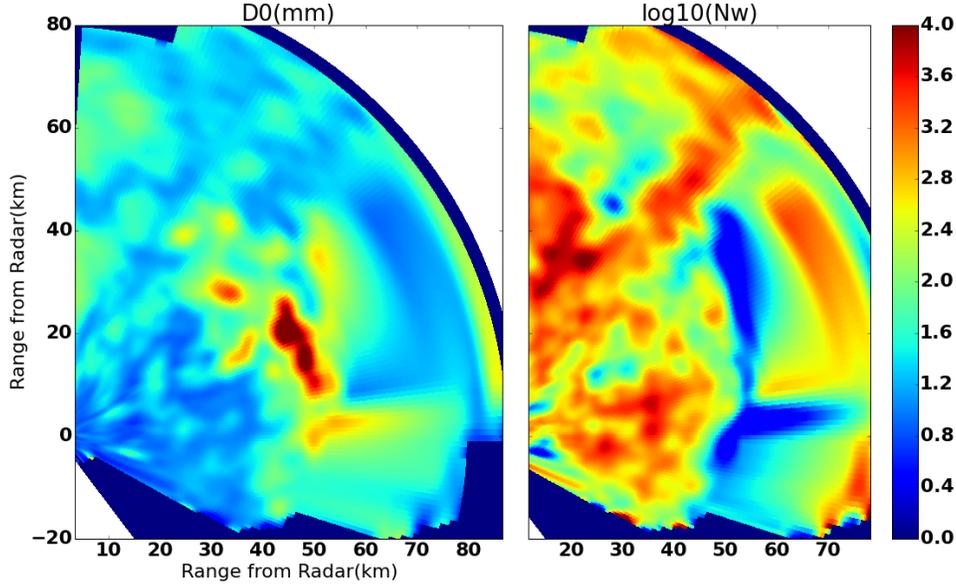
FIGURE 3.5. Microphysical parameter field derived from CSU-CHILL radar measurements. This field provides the base for the simulated data results.

TABLE 3.2. Default scattering parameters for T-Matrix scattering for the forward solution. Values are chosen based on commonly accepted values for reasonable calibration of a radar.

Scattering Parameter	Value
Z_h noise variance	1dBZ
Z_{dr} noise variance	0.2 dB
Φ_{dp} noise variance	1.5°
Temperature	10°C
Shape Relationship	Beard and Chuang(BC) [18]

The values of these parameters were chosen based upon expected levels of error in radar data, and so as to be comparable to other literature on microphysical retrievals [26]. These parameters can have a large effect. Much of our synthetic results section will attempt to calculate how much an error in fixing these values affects our retrieval accuracy. In this subsection we will focus primarily on mean absolute error [44]. As we are primarily interested

57

in the general trends each error source has on our retrieval, this should be enough. At S-Band frequencies (2.7 GHz) with no noise added beyond attenuation, we get the equivalent radar measured parameters that are shown in Figure 3.6.
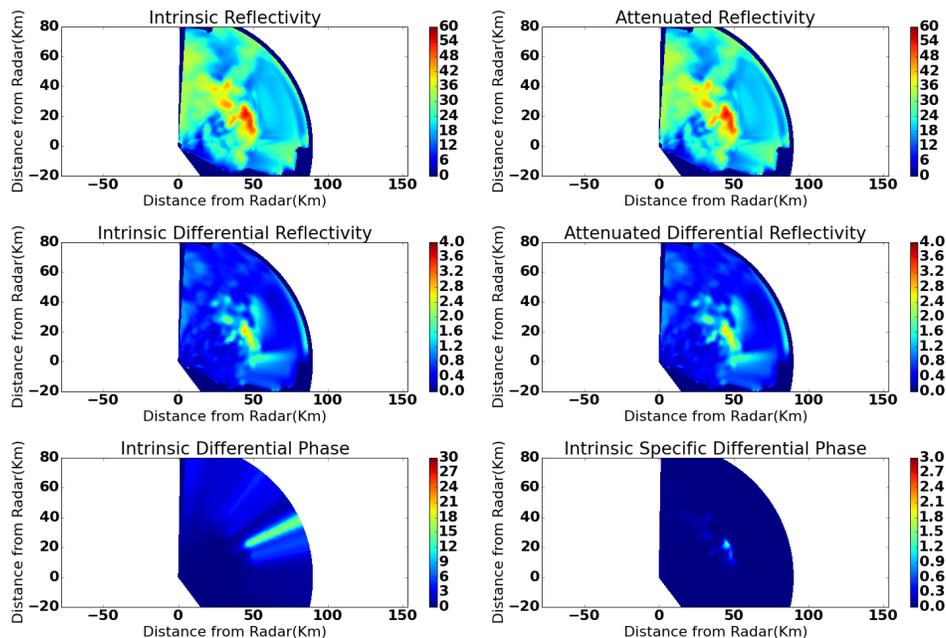
FIGURE 3.6. Scattering simulation fields at S-Band for synthetic data. Shown are both intrinsic and attenuated fields. At S-Band the degree of attenuation is not significant however.

The first step of our retrieval is a down sampling step for computational reasons. In our retrieval it is important to account for the effect the down sampling itself has on the retrieval accuracy. We can characterize the effect this has by calculating the retrieval error of a downsampled field with no sources of noise added. Without downsampling, the algorithm is able to exactly retrieve the fields so any error is due purely to the effect of downsampling. This error will be related to the spatial variability of the microphysics along the ray. For instance, if we are down sampling by a factor of two, and two adjacent range bins differ greatly, then any possible choice of retrieval will show a large error relative to one or the other bin. The results of this are listed in row one, of Table 3.3. We can see that the mean

absolute error we get from this is on the order of a hundredth of a millimeter for D_0, and roughly double that for N_w for the lower frequency bands, and higher for higher frequency bands. We can compare this to the error of predicting one microphysical parameter in one bin, by repeating the previous bin. The MAE error for our synthetic dataset is 0.022 for D_0 and 0.03 for N_w using this method. Half of this error gives us the cost of using the mean of the bins. We can see that these are similar to the numbers given in Table 3.3. We will assume that this is a lower error bound for any down sampling retrieval strategy using this dataset.

A much more realistic test is to add in standard levels of measurement and sampling noise, and then see how well the retrieval performs. If we model our noise as zero mean Gaussian noise, with standard deviances given in Table 3.2, and perform our retrieval, we get the results given in row two of Table 3.3 for each frequency band.

We can see that we do get a significant increase in the error rates of the retrieved microphysical parameters as expected. Unless explicitly stated, all of the other error tests for the synthetic data will have this same level of noise added, and as such, each increase in error rate should have the Additive noise case as it's lower bound.

While there exists a large number of ways to calibrate the power received by weather radars, all of the popular methods are costly, time consuming, or inaccurate. More frequently, the calibration procedures are done infrequently. What this means for us is that it is common for radar data to have a constant offset applied to Z_h and Z_{dr} that biases the data away from it's true value. We can usually correct for some of this error in post-processing, but inevitably some of the error will remain in the data. In rows three and four we calculate the retrieval accuracy for each frequency band with a fixed amount of calibration error added. In the third row we show results when reflectivity has $1dB$ of positive bias added. In row four we

TABLE 3.3. Scattered Simulation Retrieval Mean Absolute Error (MAE).

	$S - Band$		$C - Band$		$X - Band$		$Ku - Band$	
	D_0	N_w	D_0	N_w	D_0	N_w	D_0	N_w
Base	1.33e-02	1.97e-02	1.52e-02	2.3e-02	1.6e-02	2.63e-02	7.3e-02	0.11
Additive Noise	0.16	0.38	0.16	0.38	0.14	0.33	0.27	0.45
1 db Z_h bias	0.16	0.38	0.16	0.37	0.16	0.33	0.46	0.59
0.2 db Z_{dr} bias	0.28	0.57	0.25	0.52	0.24	0.50	0.41	0.71
Wrong Shape	0.17	0.44	0.15	0.41	0.14	0.38	0.18	0.40
Wrong $\mu = -1$	0.22	0.60	0.43	0.60	0.46	0.62	0.60	0.74
Wrong $\mu = 1$	0.22	0.41	0.22	0.40	0.22	0.39	0.31	0.48
Wrong $\mu = 5$	0.19	0.48	0.19	0.47	0.17	0.43	0.22	0.46
Wrong $\mu = 7$	0.24	0.57	0.23	0.56	0.22	0.54	0.22	0.50

show the case where differential reflectivity has $0.2dB$ added in separately. These correspond to realistic levels of calibration error for well calibrated radars. While occasionally radars can have higher errors, we assume we can fix it down to these levels in post processing.

In row five labeled "Wrong Shape" we retrieve assuming the Beard and Chuang [18] DSR, but calculate the scattered values using the Thurai and Bringi [19] DSR. The improper choice of DSR's will most heavily effect Z_{dr} in the forward scattering case. Finally in the last several rows we calculate the retrieval error by scattering with a different μ value than our retrieval algorithm assumes. We show several values in increments of two around the retrieval assumed value of three.

We see that in all cases, our retrieval is fairly resilient to these error sources. In particular, we seem to be very resilient to reflectivity calibration. Differential reflectivity calibration has a much more significant effect. For most sources of error, the lower a frequency, the less it appears to affect the data, with K_u suffering the most significant degradation. In particular we notice that when the actual μ is much lower than we expect, we experience the most significant errors.

We can see that a reflectivity calibration error of 1 dBZ did not have a large effect. A natural question then is what level of calibration error at horizontal will cause the accuracy of the algorithm to degrade? We can vary the level of calibration introduced to examine how the error structure changes as the calibration error increases. We can see in Figure 3.7 the retrieval behavior as the calibration error increases.

We can see that the error in estimation for the bottom three frequencies remains fairly low, although it does increase as a function of the bias for both D_0 and N_w. The error for D_0 barely increases, while the error for N_w see less than a 50% increase. The error for X-Band increases the quickest of the lowest three frequencies.

The K_u band retrieval however appears to be considerably more sensitive to calibration biases. The error for both of the microphysical parameters increases very quickly compared to the other three frequencies. The errors increase by a factor of three for N_w and more than a factor of six for D_0. The K_u band radar operates with more of it's DSD spectrum in the Mie band, and as such, we expect the resonances created in that region to create more of a difficulty with the retrieval.

Similarly we can look at the behavior as the differential reflectivity calibration error increases. The results for this are shown in Figure 3.8. The results for increasing differential calibration bias are a little more predictable. We see all the frequencies have a similar behavior, showing a nearly linear increase in error as a function of the bias.

It is possible to shine a little more light on the situation by looking at two more error metrics, the normalized standard error (NSE) and the normalized bias (NB) of the retrieval. NSE is given in Table 3.4, while NB is given in Table 3.5.

The normalized standard error measure tells a slightly different story. While negative values of μ remain the largest error case, there is an increased level of uniformity between
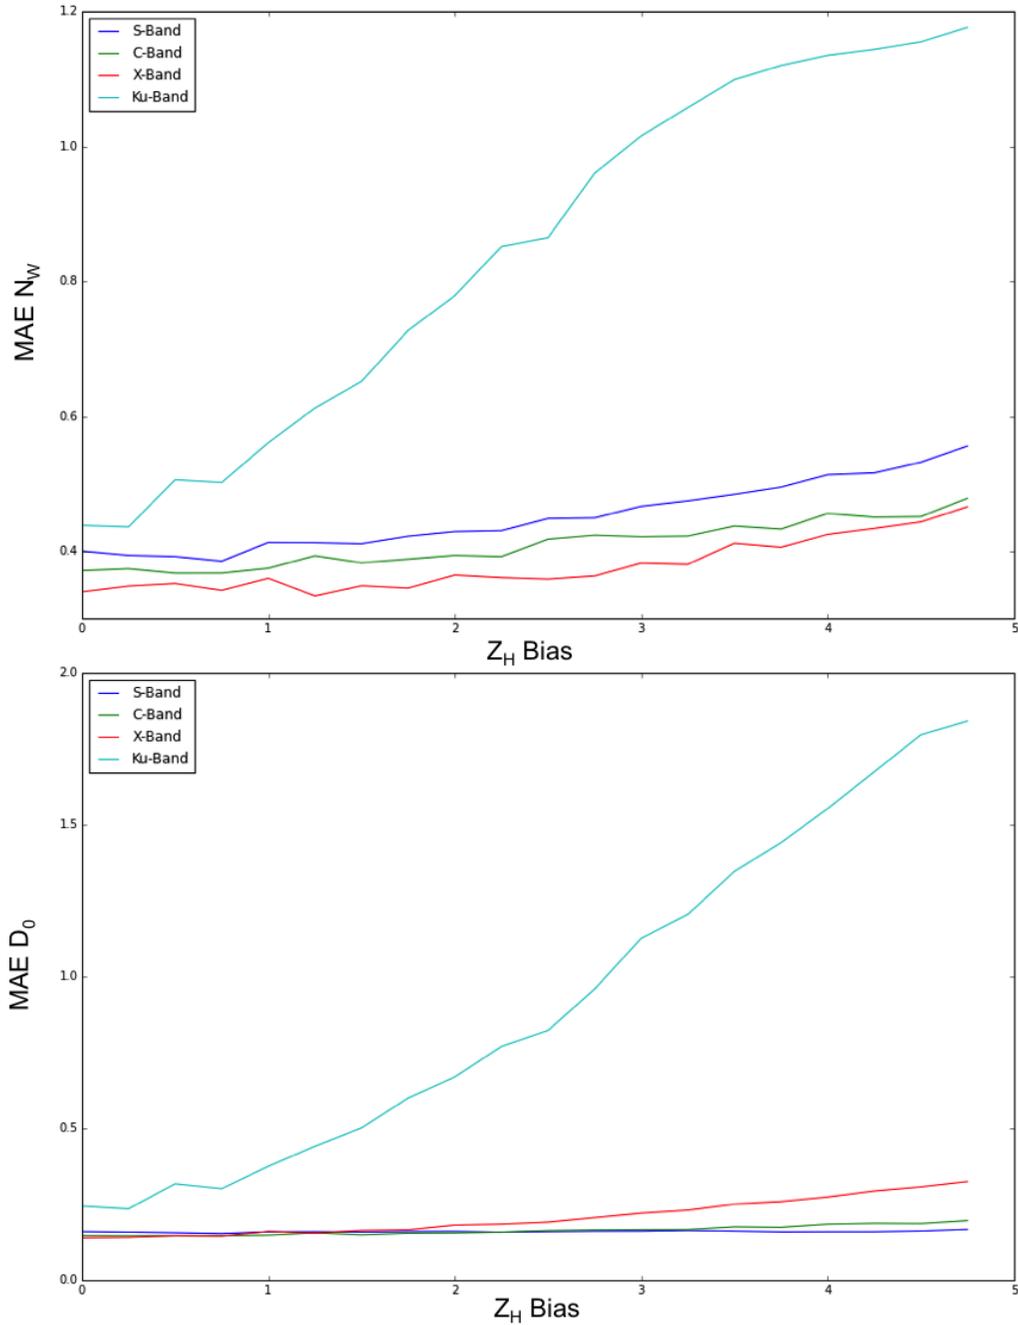
FIGURE 3.7. Mean Absolute Error as a function of reflectivity calibration error. K_u shows more serious degredatoin than the three lower frequencies. The other three frequencies show limited degredation as a function of the calibration.

the frequencies for most of the other cases. This suggests that for the MAE case, a few

outliers are causing a disproportionate amount of the error.
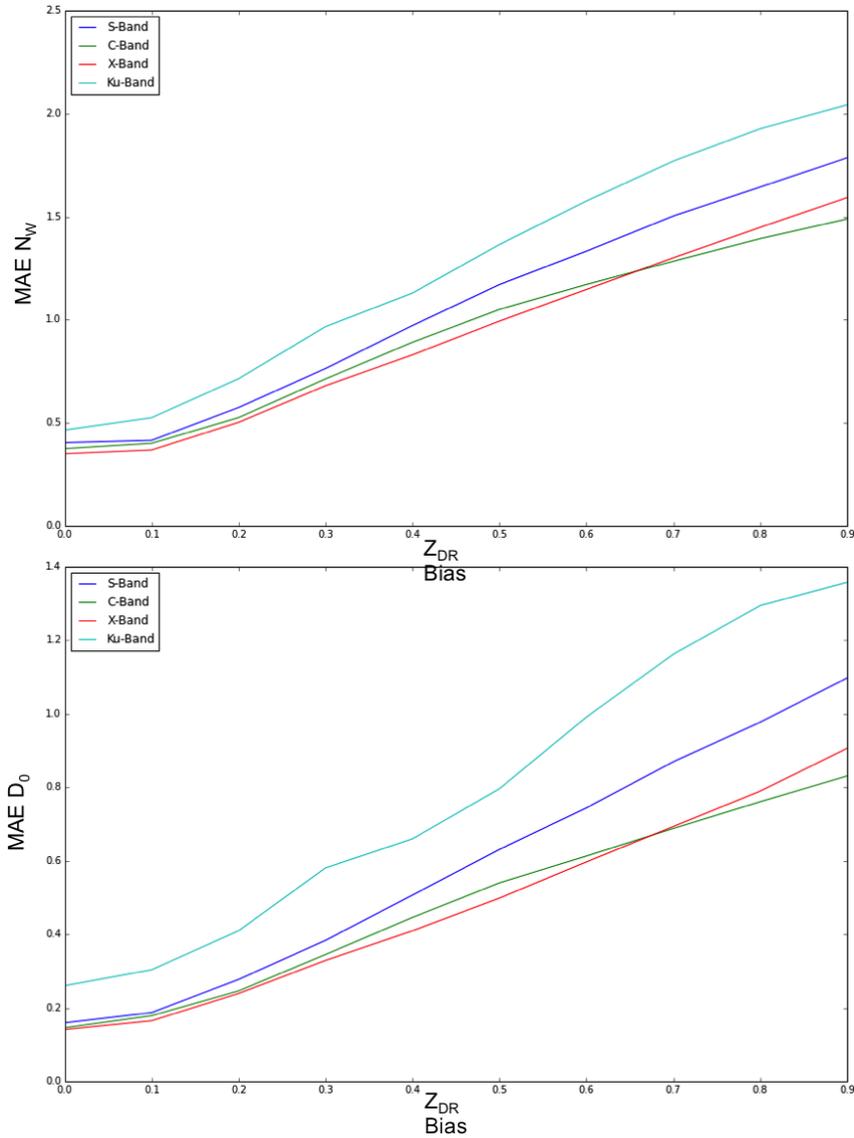
FIGURE 3.8. MAE as a function of differential reflectivity calibration error. The behavior of the four frequencies shows a much more similar behavior than the results from the reflectivity claibration case.

Table 3.5 which shows the normalized bias in the retrieval paints a much prettier picture. While again, the negative value of μ contains significant errors, having a bias of about 20%, the rest of the error sources do not appears to cause too much of a degredation. The K_u band suffers more severely than the other frequencies in the bias cases, but has similar performance

63

TABLE 3.4. Scattered Simulation Retrieval Normalized Standard Error (NSE)

	$S - Band$		$C - Band$		$X - Band$		$Ku - Band$	
	D_0	N_w	D_0	N_w	D_0	N_w	D_0	N_w
Base	9.2e-03	7.6e-03	1.00e-02	8.4e-03	1.0e-02	9.4e-03	4.0e-02	4.0e-02
Additive Noise	0.11	0.14	0.11	0.14	0.096	0.13	0.19	0.18
1 db Z_h bias	0.11	0.14	0.11	0.14	0.10	0.12	0.31	0.22
0.2 db Z_{dr} bias	0.19	0.20	0.18	0.20	0.15	0.17	0.26	0.24
Wrong Shape	0.11	0.17	0.11	0.16	0.09	0.15	0.12	0.15
Wrong $\mu = -1$	0.29	0.24	0.30	0.23	0.31	0.24	0.41	0.30
Wrong $\mu = 1$	0.14	0.16	0.15	0.16	0.15	0.15	0.24	0.20
Wrong $\mu = 5$	0.13	0.19	0.13	0.18	0.12	0.17	0.15	0.19
Wrong $\mu = 7$	0.16	0.22	0.16	0.22	0.15	0.21	0.15	0.20

TABLE 3.5. Scattered Simulation Retrieval Normalized Bias (NB) results.

	$S - Band$		$C - Band$		$X - Band$		$Ku - Band$	
	D_0	N_w	D_0	N_w	D_0	N_w	D_0	N_w
Base	4.8e-04	7.2e-05	9.60e-05	4.7e-04	3.8e-04	-1.9e-04	4.1e-02	2.7e-03
Additive Noise	3.5e-03	-0.02	-1.9e-03	-0.02	-5.5e-03	-0.01	-0.06	0.05
1 db Z_h bias	-1.7e-04	-0.06	-0.02	-0.05	-0.04	-0.03	0.22	0.12
0.2 db Z_{dr} bias	-0.18	0.19	-0.17	0.18	-0.14	0.15	-0.24	0.2
Wrong Shape	0.05	-0.08	0.03	-0.07	0.02	-0.06	-0.01	-0.04
Wrong $\mu = -1$	-0.28	0.21	-0.29	0.21	-0.30	0.22	-0.40	0.28
Wrong $\mu = 1$	-0.11	0.09	-0.11	0.08	-0.13	0.13	-0.22	0.15
Wrong $\mu = 5$	0.08	-0.10	0.08	-0.10	0.07	-0.10	-0.01	-0.04
Wrong $\mu = 7$	0.13	-0.16	0.13	-0.17	0.12	-0.16	0.06	-0.11

for the other cases. Again the increased error with respect to the calibration bias is likely due to pushing it into nonlinearities in the scattering regime.

While we have looked primarily at the absolute error to get a sense of the magnitude of the distortions created in the retrieval, it is instructive to look at the actual values of the retrieval error. In particular, if we compare the errors of the two microphysical parameters as shown in Figure 3.9, we see an interesting trend. There is a pseudo-linear negative trend between

the two parameters. This means as we make an increasing level of error in estimating one parameter, we expect to make an error in the opposite direction in the other microphysical parameter. While this is not strictly positive for the results of the retrieval itself, what it does do is reduce errors in some derived applications of the retrieval. For instance, if estimating rain rate from the retrieval, we can expect these offsetting errors to reduce the overall level of error we would have expected by looking purely at the MAE results.

Based on this error analysis using synthetic data, this retrieval should be able to return the drop size distribution parameters within an error bound of 0.2 for the median drop diameter D_0, and within 0.4 for the normalized intercept parameter N_w given in logarithmic units.

3.3.2. SINGLE RADAR MATCHING. One of the measures of success for a retrieval algorithm is how well the retrieved drop size distribution parameters can reconstruct the signal that the radar measures. With real radar measurements, we have no way of knowing with certainty the microphysical situation in the atmosphere, and as such we cannot directly compare the retrieved DSD. Instead, we must use more indirect comparisons. As such we can perform the microphysical retrieval, then calculate the radar equivalent measurements for that microphysical field. By comparing the radar measured field, and our re-scattered field, we can get a sense of how well the microphysical field represents the radar measured data.

One issue with fitting measured data is the issue of noise. To perfectly reconstruct data, the algorithm would have to generate a DSD that explains the noise, as well as the true signal. As the noise is a random signal with little to no correlation with the underlying signal, this is not likely what we wish to recreate. Unfortunately, separating out the noise from the true signal is not a trivial(or necessarily even possible) operation. As such, the
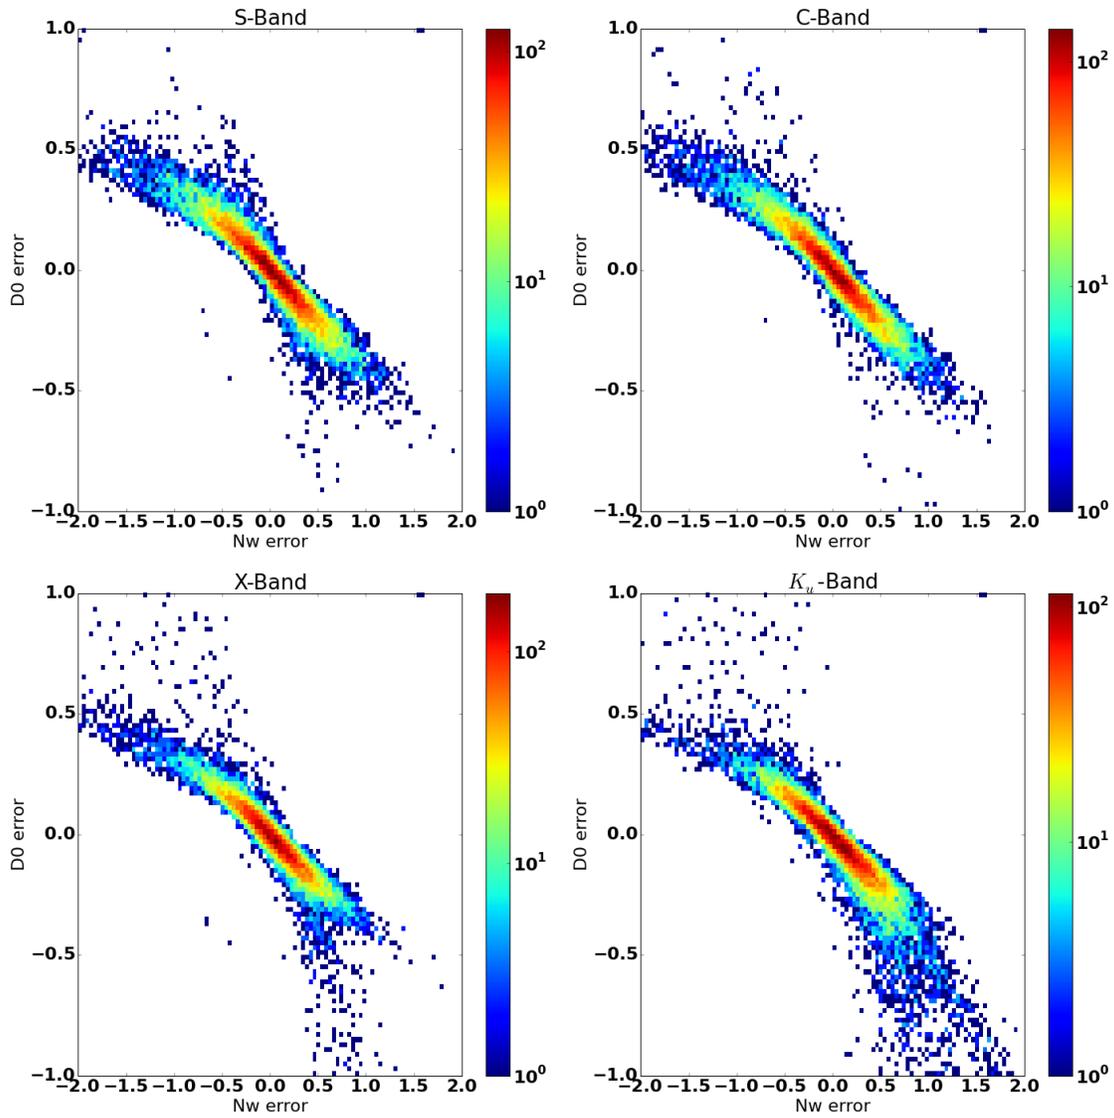
FIGURE 3.9. Simulated Data Microphysical Errors from the two DSD parameters. There is a very well defined negative correlation between the errors in the two parameters. This indicates that the errors will offset each other for derived products such as rain-rate.

comparison with measured data will have some baseline level of error that is unavoidable. As shown in the previous section, the algorithm is fairly resilient to noise sources however.

The verification of this algorithm will use three radars operating at different frequencies as examples. It will cover S, X, and Ku-Band radars.

3.3.2.1. *NPOL(S-Band) IFloodS.* The first radar examined is NASA's Polarimetric Radar (NPOL). This is a deployable S-Band radar that is deployed to various field campaigns to provide ground validation for the various GPM-GV field campaign. More details about NPOL are in Appendix A. The data file demonstrated here is a Range-Height Indicator scan (RHI) from the IFloodS field campaign.

The radar measured moments can be seen in figure 3.10. This is a storm with both stratiform and convective regions. The melting layer is visible around 3.5 km. Additionally, the radar experiences beam blockage and ground clutter for the first few degrees. At a range of 35 to 45 kilometers there is a strong precipitation core that reaches down to ground level.

As the radar cannot directly measure the underlying DSD field, a microphysical retrieval is performed on the data, and then the results are put through the T-Matrix scattering process to estimate the radar measured parameters based upon that microphysical parameterization. This gives the results in figure 3.11. In the left column is the scattered moments of the retrieved DSD, while in the right column is the radar measured parameters.

It is evident in these images that the ground clutter at low elevations is causing the retrieval to return erroneous results. Once out of ground clutter however the microphysical retrieval accurately reproduces the measured radar fields. In the twin core structure, it is even able to return a representative DSD parameterization above the melting layer (ML). This is likely due to strong updrafts causing the return in the core to be liquid, despite being above the ML. Reflectivity and Differential Reflectivity show a high level of agreement between the two. Although Φ_{dp} is not shown here, it too agrees fairly well. Instead of Φ_{dp}, K_{dp} is shown. The algorithm matches the reconstruction between Φ_{dp}, not K_{dp} and so we allow K_{dp} to differ between the retrieval and the measured field. As K_{dp} is an estimated quantity for the radar a difference between the two is not an error. It is evident in the
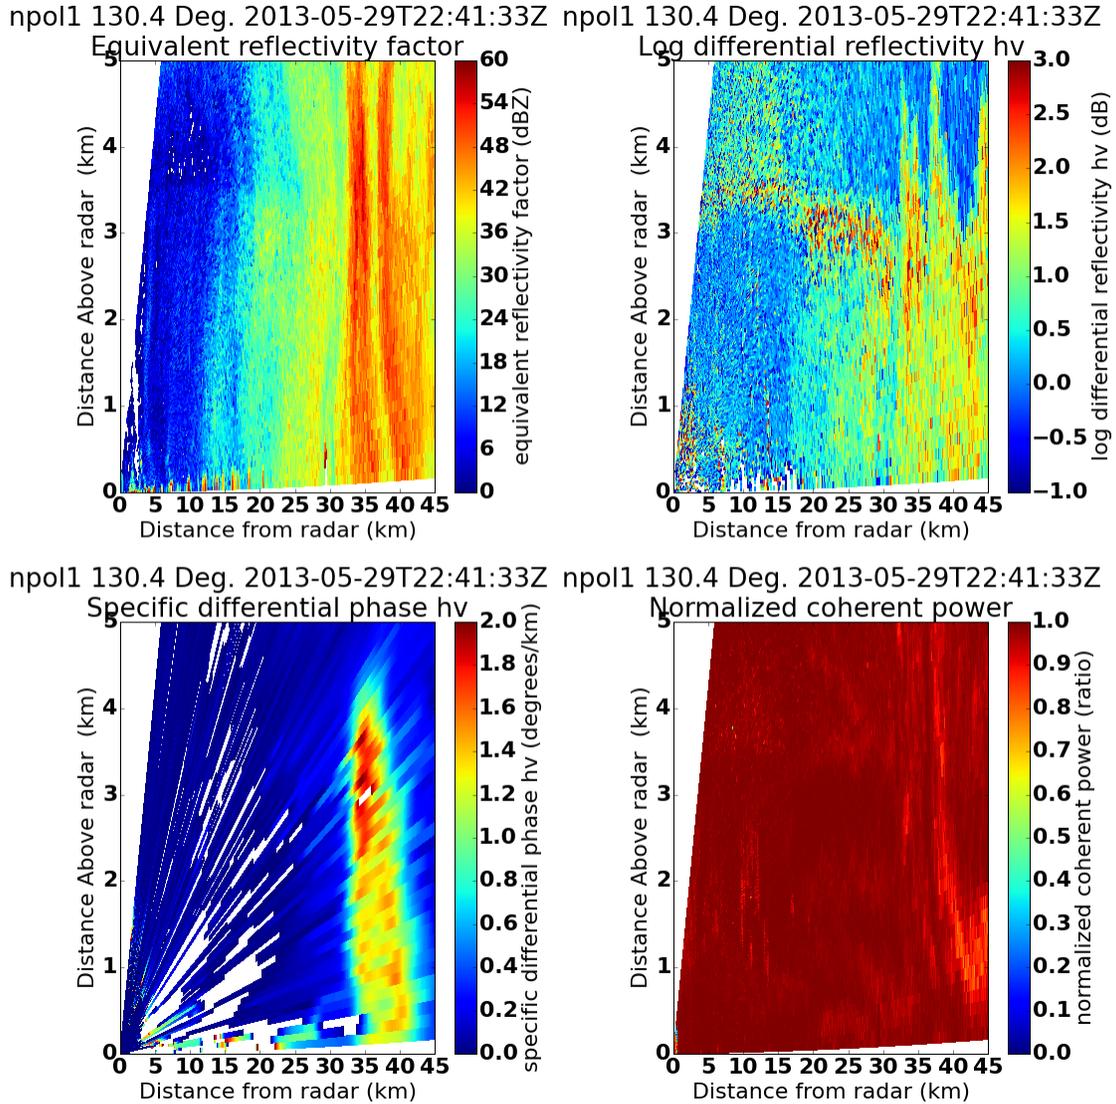
FIGURE 3.10. NPOL IFloodS RHI Situational Overview. The melting layer is visible around 3.5 km. In the upper panels are the two power measurements. In the lower left is the radar derived K_{dp} measurement. The bottom right panel shows normalized coherent power, also commonly called signal quality index (SQI).

storm cores that the retrieval algorithm provides a much finer spatial localization of K_{dp} than that recorded by the radar. The estimation method the radar uses for K_{dp} is based on an extended version of the $Z - \Phi$ method [1] due to Wang et. al. [45]. This algorithm, while better than most, does introduce a level of spatial smoothing common to the vast majority
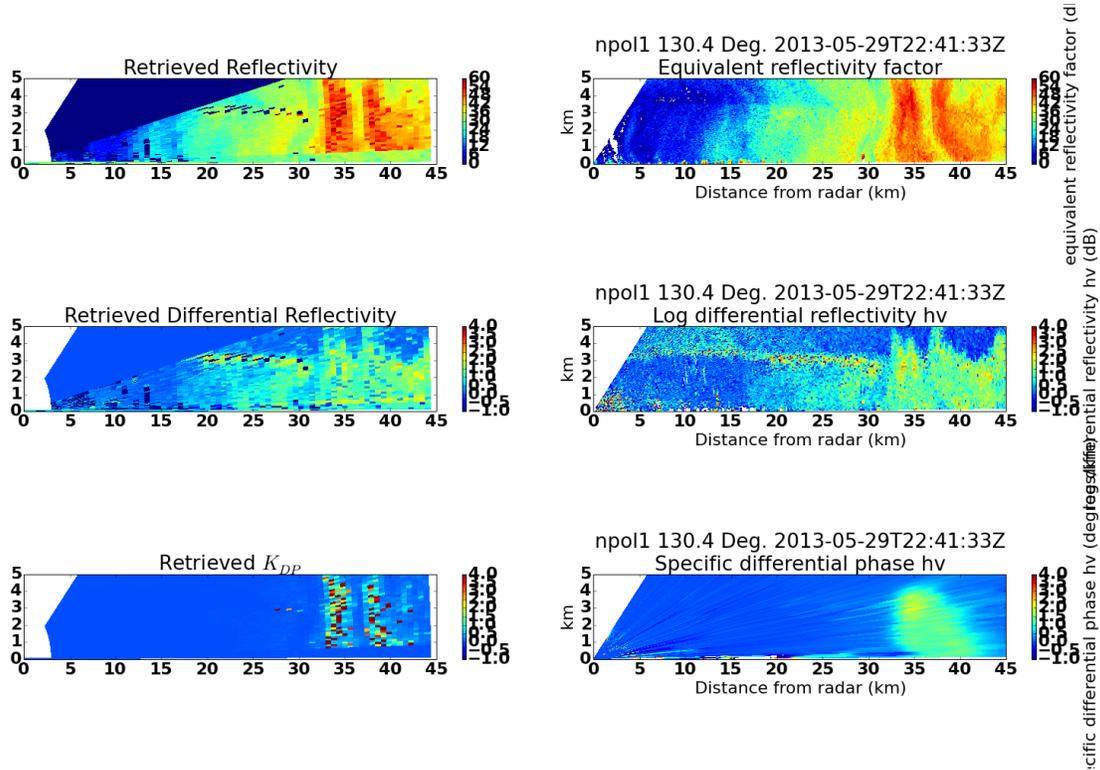
FIGURE 3.11. Results of microphysical retrieval process for NPOL IFloodS test scan. The left panes show the T-Matrix scattered parameters based on the retrieved microphysical parameters. The right panel shows radar measured parameters.

of K_{dp} estimation algorithms. Reducing the spatial smoothing action of most estimation algorithms by reducing the path length they estimate over, results in a higher susceptibility to backscatter differential phase as discussed in Chapter 2.

More specifically, Figure 3.12 shows the retrieval results for a single ray at an elevation of 2.3°. There are several features of note in this plot. The first is the behavior over the first 12 km. This part of the scan is heavily degraded by ground clutter and shows reduced ρ_{hv} and NCP values. As such, the algorithm filters many of these range bins out and so does not suffer from attempting to estimate the erroneous Φ_{dp} shown. This can be seen in the form of "holes" in the graph. The retrieval is to capture all major trends in the data. The
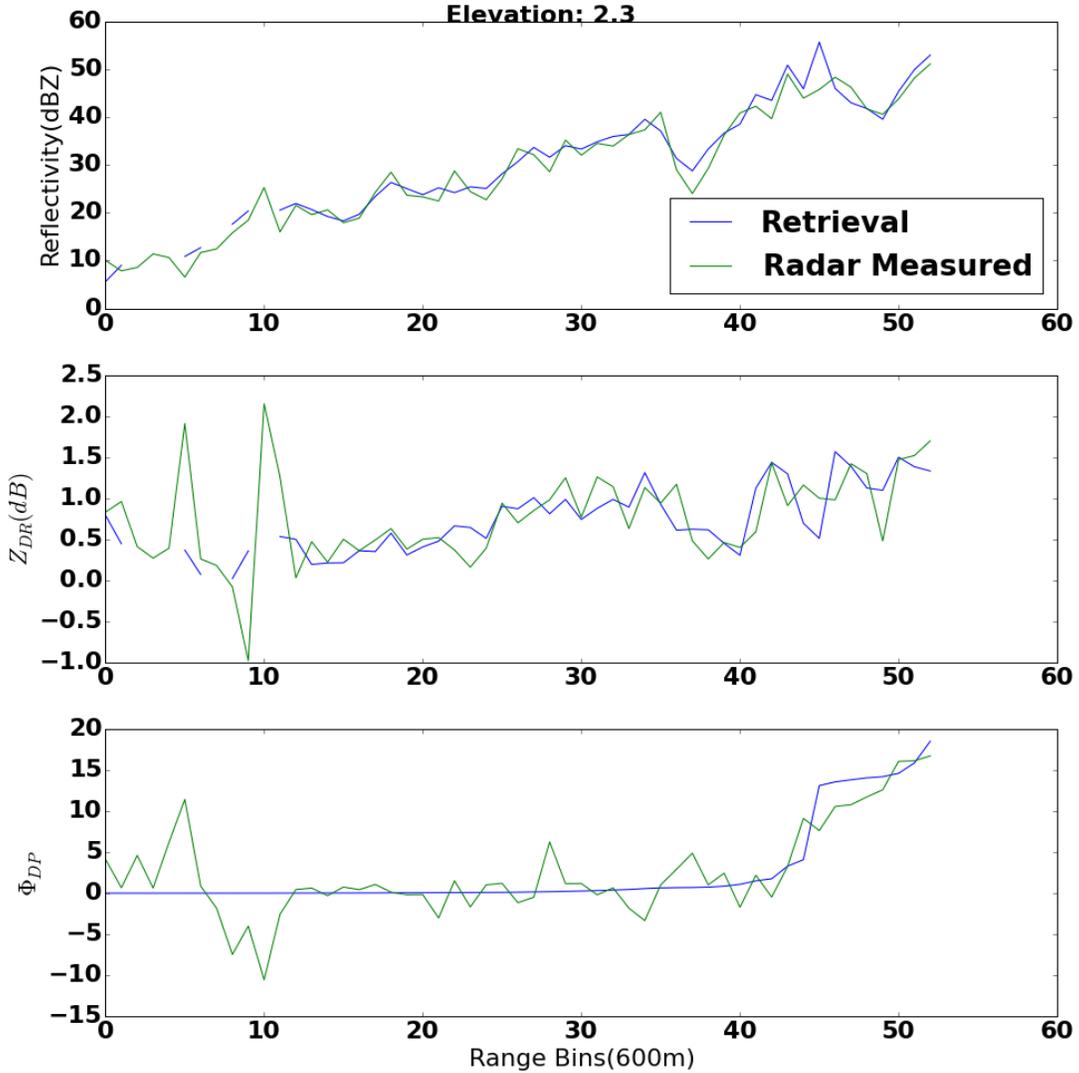
FIGURE 3.12. NPOL Single Ray Retrieval. Shown is a single ray of the retrieval corresponding to 2.4°. Overlaid is the radar measured parameters. The beginning of the ray shows some ground clutter, which the microphysical retrieval rejects.

Φ_{dp} profile appears to properly reject the backscatter differential phase while still capturing the underlying signal.

Figure 3.13 shows quantitative results for how well the retrieved parameterizations reproduced the radar measured fields. Results are shown for four different statistical metrics as discussed in Appendix A. The error is shown for each retrieved ray, with increasing ray

numbers corresponding to increasing elevations of the radar. The results in this graph are restricted to rays above the ground clutter and beam blockage.

Reflectivity can be seen to have a very low level of error by all of the metrics, as well as having no apparent bias. Differential Reflectivity has a very low NSE error, with a low level of bias that slow goes negative at increasing elevations. This is not completely surprising. The scattering operations are calculated based upon a horizontally pointing radar. Correcting this can be accomplished by generating a new set of scattering tables at each elevation angle, or calculating a bias offset term. A similar effect on the normalized bias can be seen in Φ_{dp}. Some of this error can be attributed to including the melting layer in these calculations however. The MAE error shows that reflectivity, and differential reflectivity errors are both less than the assumed level of measurement noise for a radar.

Finally the MASE metric shows that the retrieval is indeed adding information beyond a naive predictor.

While it is instructive to see the metrics for each individual ray, the number that is more useful to compare with other retrievals is the composite number over all of the rays. This composite scores for the entire scan are given in Table 3.6. This table reaffirms the small negative bias in the differential reflectivity reconstruction. This is due to the elevation angle. Errors are small for all parameters, with no appreciable bias for reflectivity or differential phase. The mean absolute error shows low error levels for both reflectivity and differential reflectivity. The differential phase mean absolute error is somewhat higher, but can somewhat be attributed to backscatter differential phase. Finally, as noted before, the MASE error shows the retrieval is indeed adding information over a naive estimator.

Shown in Figure 3.14 is the retrieved drop size distribution parameterization. The first few degrees are contaminated by the ground clutter and beam blockage. The bottom of the
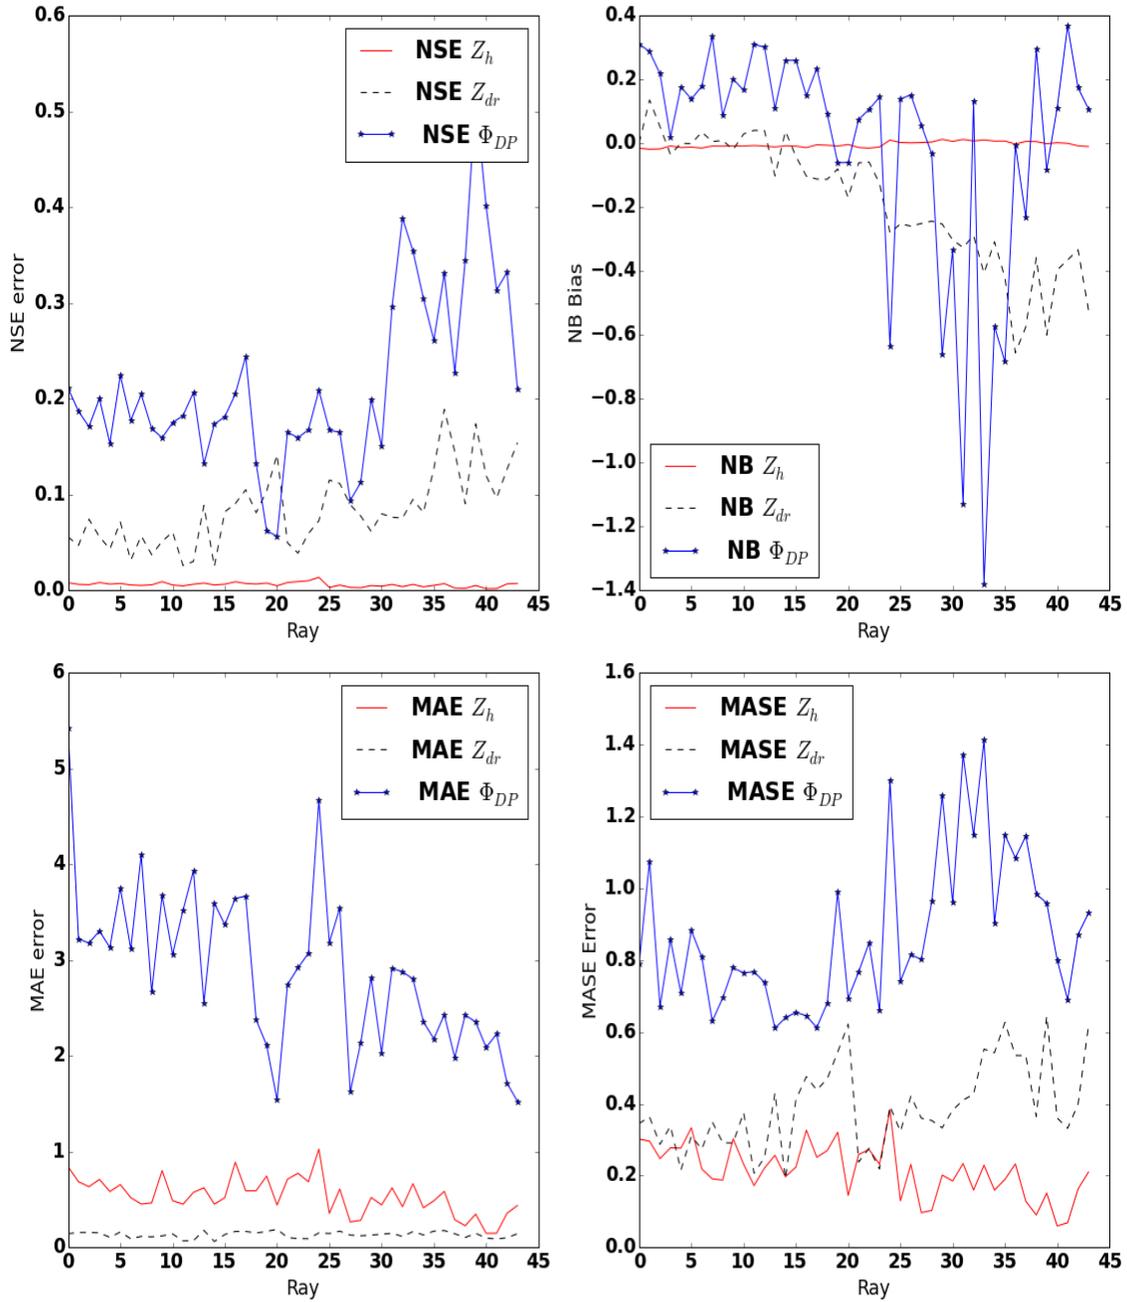
FIGURE 3.13. Error Metrics for NPOL IFloodS Single Radar Retrieval. The increase in ray number corresponds to an increase in elevation. As expected, an increase in elevation is correlated with an increase in error. This is caused by the geometry of the drops changing due to a difference in angle of view.

melting layer is visible around 3km, as a region with significantly larger drops that are fewer in number. This matches what would be expected. As frozen particles melt, they are often

TABLE 3.6. NPOL IFloodS Entire Retrieval Error Metrics. This shows the recreation error with respect to the four metrics(Normalized Standard Error, Normalized Bias, Mean Absolute Error, Mean Absolute Scaled Error). The error metrics are taken using the entire retrieval as input.

Moment	NSE	NB	MAE	MASE
Z_h	0.0058	-0.0048	0.54	0.21
Z_{dr}	0.083	-0.18	0.13	0.39
Φ_{dp}	0.21	-0.0048	2.9	0.87

larger in size than the equivalent completely liquid particles. After they are able to pick up a certain amount of fall velocity, they often break apart into more numerous amounts of small droplets. In the two cores we see a large number of medium sized drops. This matches with expectations for a strong updraft. In the updraft, vertical velocities are high enough to keep larger drops suspended long enough for them to grow to larger sizes. In addition, the updraft lofts air up allowing for an increased rate of drop formation. The areas above and to either side of the cores are likely either frozen particles, or areas where smaller drops were ejected from the core and have much slower fall velocities. The lowered fall velocity would decrease the rate at which these drops are able to grow by collision/coalescence.

Finally Figure 3.15 shows a scatterplot of the microphysical parameterization. The data follows the general trend expected for the data to maintain physical realizability. It also shows instances where the optimization fails to completely converge and forces N_w to 6 to try and match ground clutter. The observed trend is consistent with both theoretical and empirical expectations.

This subsection shows the algorithm returns a physically plausible microphysical parameterization that is capable of recreating the radar measured scene. Data matched not only the radar measured parameters, but the parameterization matched expectations about the physical behavior of the microphysics relative to the storm features. In particular, it was shown
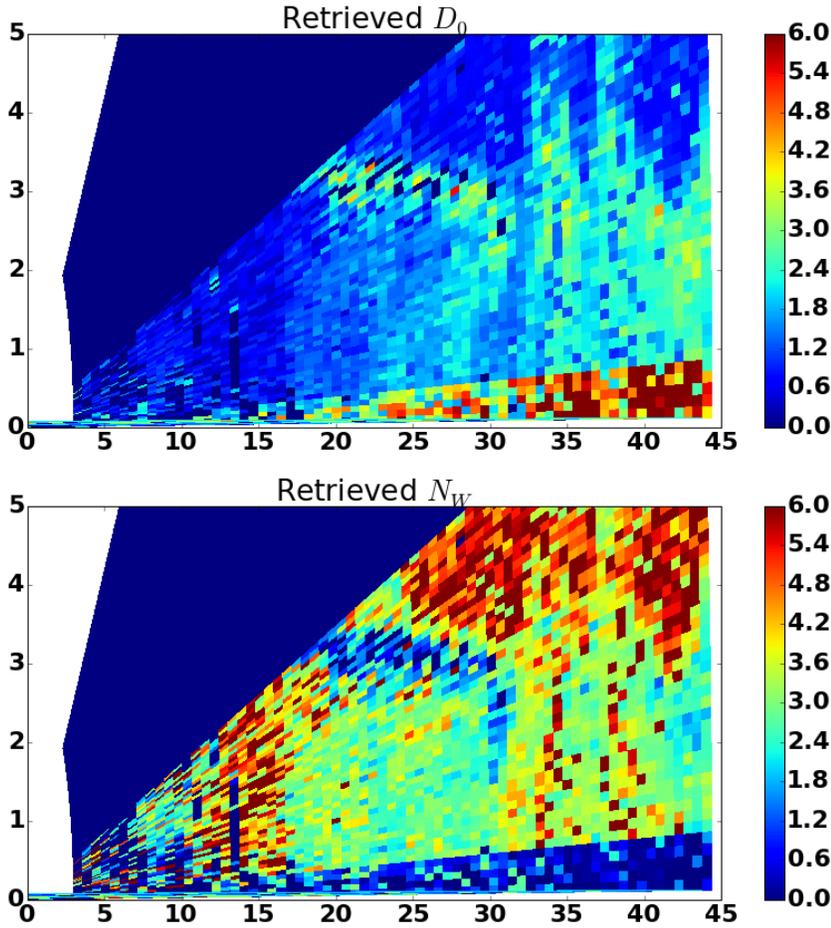
FIGURE 3.14. NPOL retrieved Drop Size Distribution parameters. The top pane shows the median drop diameter, while the bottom parameter shows the normalized intercept parameter. The effect of strong beam blockage is visible at low elevations.

that the behavior right below the melting layer was consistent with physical expectations, as was the behavior in two updraft towers. The auxillary benefit of finer K_{dp} localization was also seen in the displayed parameters.

3.3.2.2. *NASA D3R(Ku-Band).* The NASA Dual-frequency Dual-polarized Doppler Radar(D3R) is a K_u/K_a band solid state pulse compression radar operated by NASA as part of it's GPM
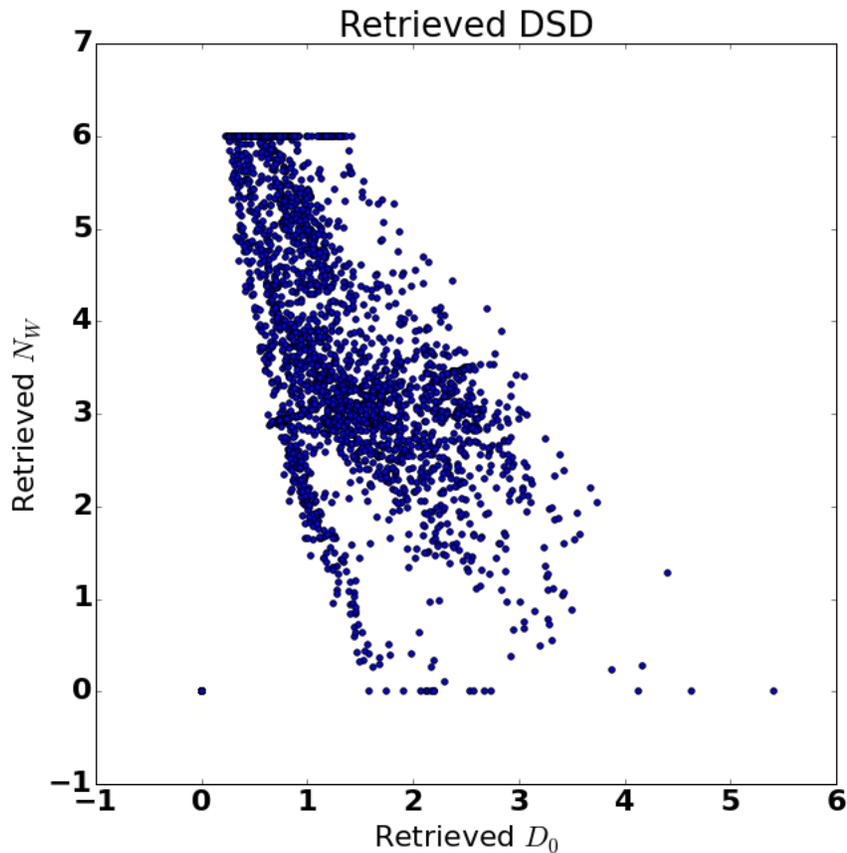
FIGURE 3.15. Self Consistency of DSD Retrieval. The overall trend is consistent with both theoretical and empirical expectations about the behavior of the two parameters with respect to each other. Outliers are visible where the optimization failed to converge, causing N_w to be set to six, which is the maximum value allowed.

Ground Validation Field Campaigns[46]. This subsection will show results from the K_u band radar using data from the IFloodS field campaign discussed at more length in Appendix A. In particular we will use an RHI scan from the 29th of May 2013 at 22:41Z. This is closely temporally and spatially matched to the NASA NPOL S-Band scan displayed above. The primary metric we will be evaluating in this section is the ability of the microphysical parameterization to reconstruct the radar measured moments.
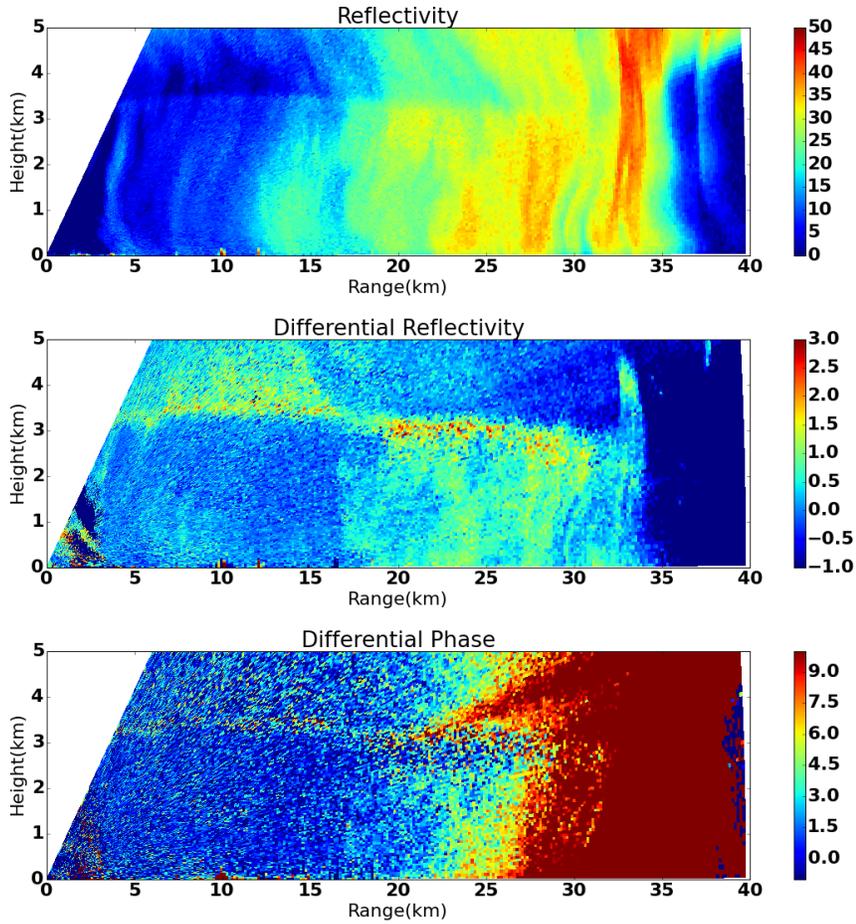
FIGURE 3.16. D3R(ku) Radar Measured Parameters. The first two panes show the power measurements, while the third shows differential phase. The melting layer is visible at about 3 km. There is some beam blockage at low elevations, caused by a feature visible at 10 km. All measurements are directly from the radar with no corrections applied.

Figure 3.16 shows the radar measured moments. This scan shows the two convective cores spaced 10 km apart as seen in the NPOL scan, with the melting layer visible just above 3km. . It does shows a weaker Z_{dr} signature, primarily due to differential attenuation, and a significantly higher Φ_{dp} as expected from a higher frequency. Additionally, these scan show a much higher attenuation than the S-Band case. In particular, after the second core the signal is almost completely extinct. Also visible is high level of clutter and anomalous propagation close to the radar, extending to about 4 km. For the retrieval results presented here, this

clutter has been filtered out as it does not represent liquid precipitation. Additionally, at 10 km a feature can be seen which causes beam blockage at lower elevations.
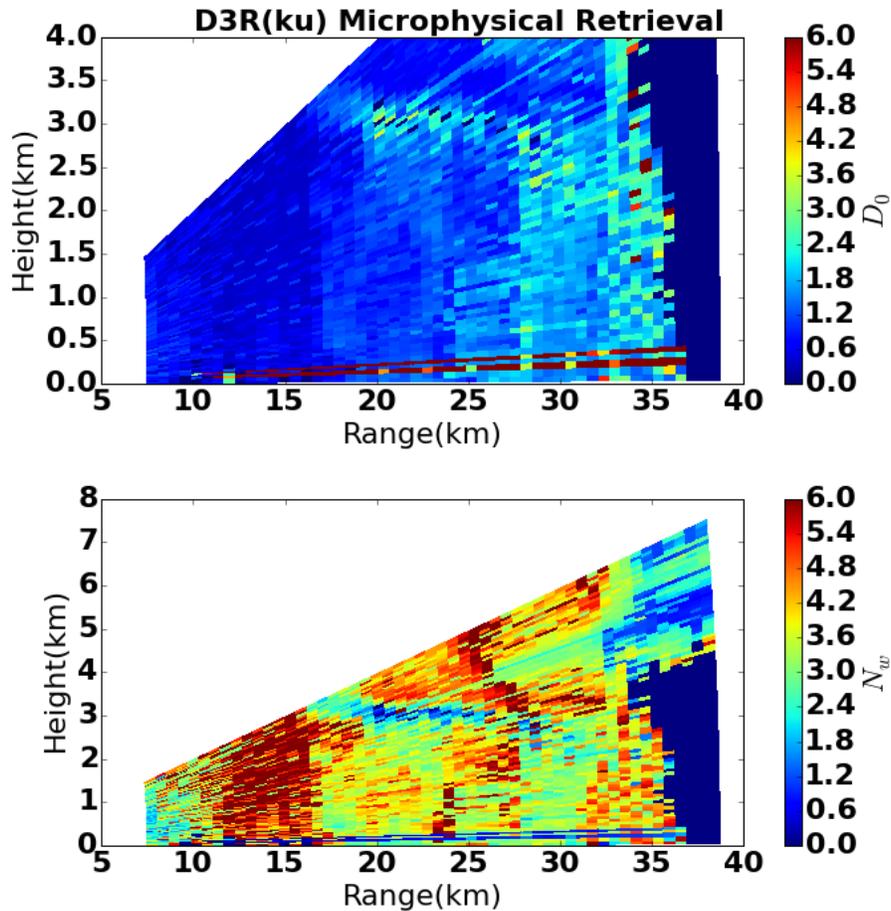
FIGURE 3.17. D3R(ku) Retrieved Microphysical Parameterization Results. The top pane shows median drop diameter, while the bottom pane shows the normalized intercept parameter. Similar features can be seen as in the S-Band case. The extinction that happens at 37 km causes the signal to be unretrievable.

The retrieved microphysical parameterization, as shown in Figure 3.17 shows primarily stratiform type rain in the early ranges with increasing median drop diameters in the two convective towers. As convective storms have higher median drop diameters than their stratiform counterparts, this matches what a visual identification confirms.
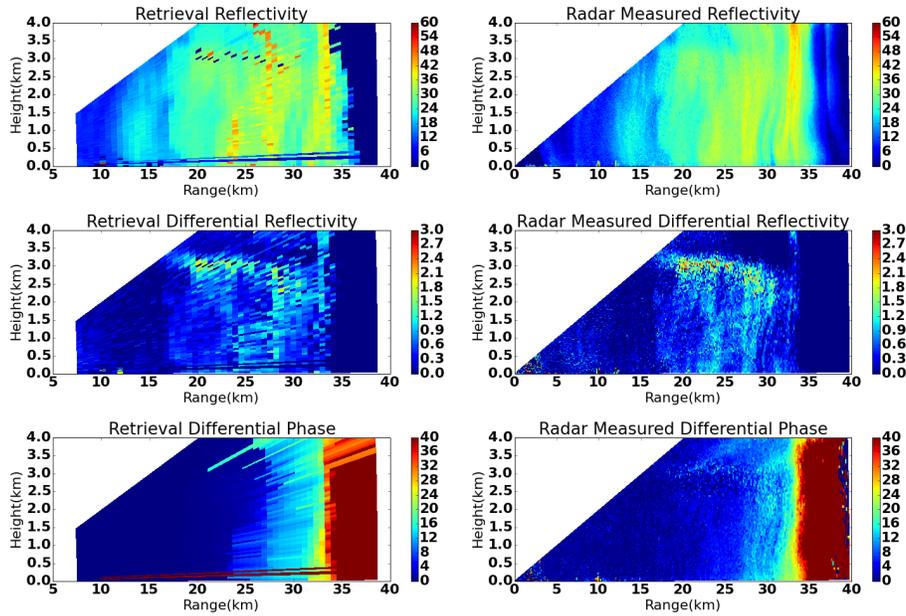
FIGURE 3.18. D3R(ku) Retrieved Scattered Moments. The left panes show the recreated measurements based on T-Matrix scattering of the retrieved microphysical parameterization. The right panes show the equivalent radar measured fields. There is good morphological agreement, although the area in the melting layer shows anomalies as expected.

The retrieval results after scattering are shown in Figure 3.18 in the left panes, while the radar measured parameters are in the right panes. After accounting for downsampling, the radar is able to recreate the fields well. In particular the areas below the melting layer shown a very high level of agreement. The convective towers are well represented, and the area below the melting layer is accurately reproduced. There is some anomalous results in the melting layer, but as the algorithm is designed to work only in liquid precipitation, this is not of concern.

Figure 3.19 shows retrieval comparisons for a single ray taken at 1.4°. The plots display good agreement on all parameters. At about 24 km the plots shown a short lived deviation. This is most likely due to non-uniform beam filling at the edge of the convective tower.

Of more interest for evaluating the algorithm is the overall error metric results. Table 3.7 lists the NSE and NB scores for the comparison of the retrieved fields with the radar
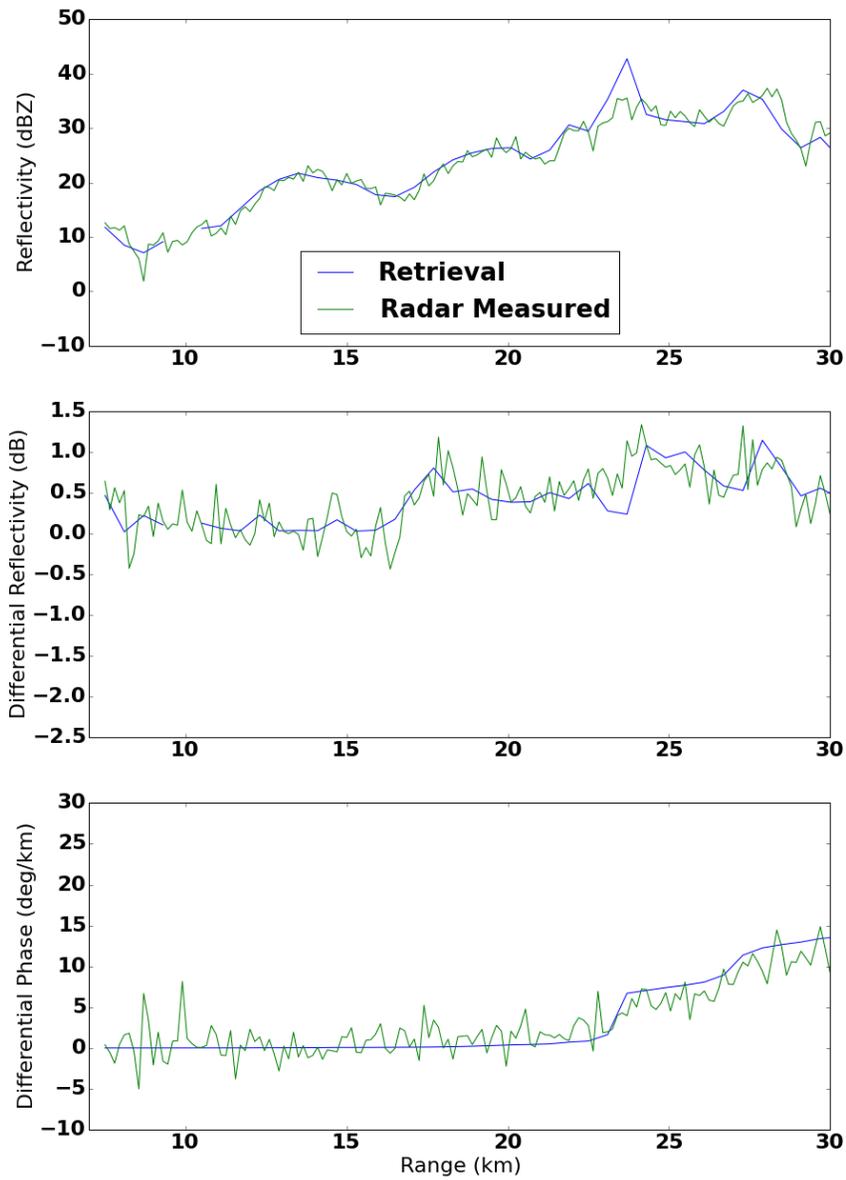
FIGURE 3.19. D3R(ku) Ray Retrieval at 1.4∘. The top two panes show reflectivity and differential reflectivity, while the bottom shows differential phase. Good agreement is seen between all the parameters, with only a small deviation at about 24 km.

measured fields. These error metrics use the data under 3 km and before 32.5 km. These restricts the effect of the melting layer and areas where the beam has gone extinct, or entered a low signal to noise region as this can cause radar measurements to be inaccurate. Both of the power fields, reflectivity and differential reflectivity, show very good with results with no major biases or significant errors(both are under 1% recreation error). The differential phase field shows only a 4% normalized standard error, but it does have a 19% normalized bias. Some of this can be due to the high level of clutter seen at the beginning of the rays, as well as some effect of system phase, which has not been completely accounted for.

TABLE 3.7. D3R(Ku) Overall Error Metrics. The comparison is based on data under 3 km in height, and within 32.5 km of the radar to avoid errors due to the melting layer and frozen precipitation, as well as signal extinction.

	Z_h	Z_{dr}	Φ_{dp}
NSE	0.003	0.047	0.040
NB	0.008	0.0197	−0.188

Overall this scan demostrates the algorithm's ability to work on higher frequency radars. The microphysical field calculated did indeed recreate the radar measured fields with a high level of accuracy, showing very few errors, and only a small level of bias in the differential reflectivity field.

3.3.3. CROSS RADAR COMPARISONS. Another way to gauge the accuracy of the single radar retrieval is to compare retrievals between two radars with overlapping measurement areas. A comparison of the retrievals for the IFloodS case showing retrievals from both the S-Band NPOL, and the Ku-Band D3R radars is shown in Figures 3.20 and 3.21. A good agreement in storm morphology can be observed. In particular the agreement in the median drop diameter fields is very good above the beam blockage contamination. The normalized

intercept parameter graphs also show a good agreement on the storm morphology, though the Ku-Band radar estimates an area of slightly higher values in the storm center. This is possibly due to signal to noise ratio (SNR) issues as the area has a lower reflectivity than much of the surrounding storm.
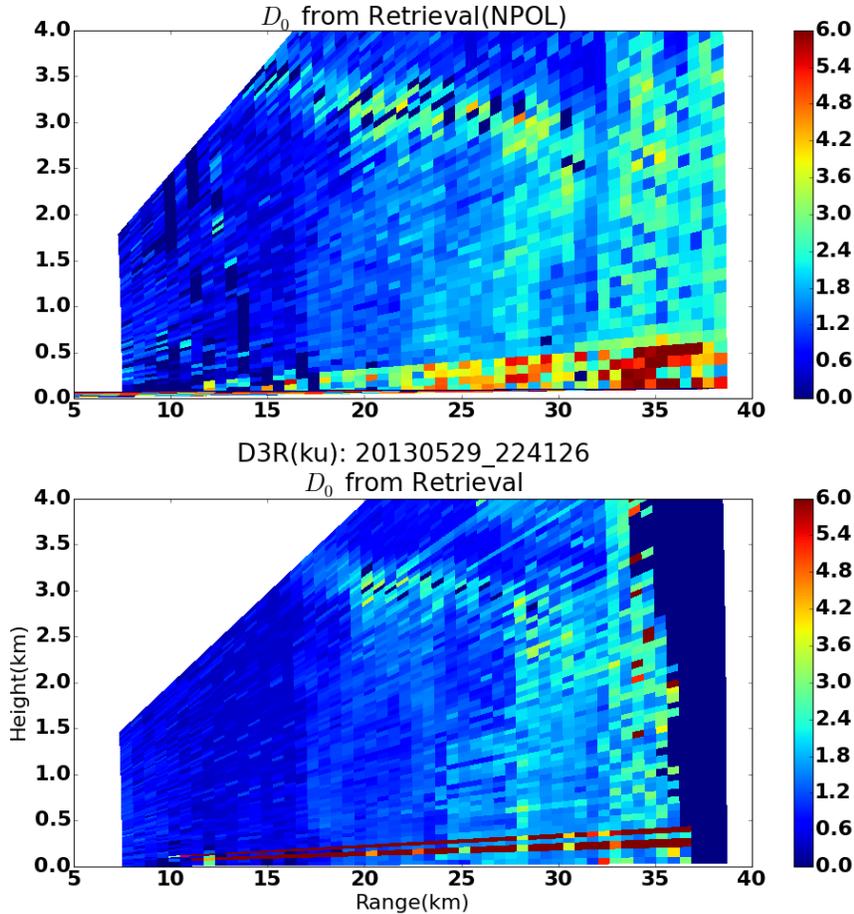
FIGURE 3.20. D3R-NPOL median drop diameter comparison. Data is shown for microphysical retrievals from the IFloodS case described in the text. The morphology matches, as well as overall values.

Figure 3.22 shows the statistical distribution of microphysical parameters between the two retrievals. There is a good agreement between the two distributions in general. However, D3R does appear to classify far fewer regions as very low N_w. Again, this is believed to be
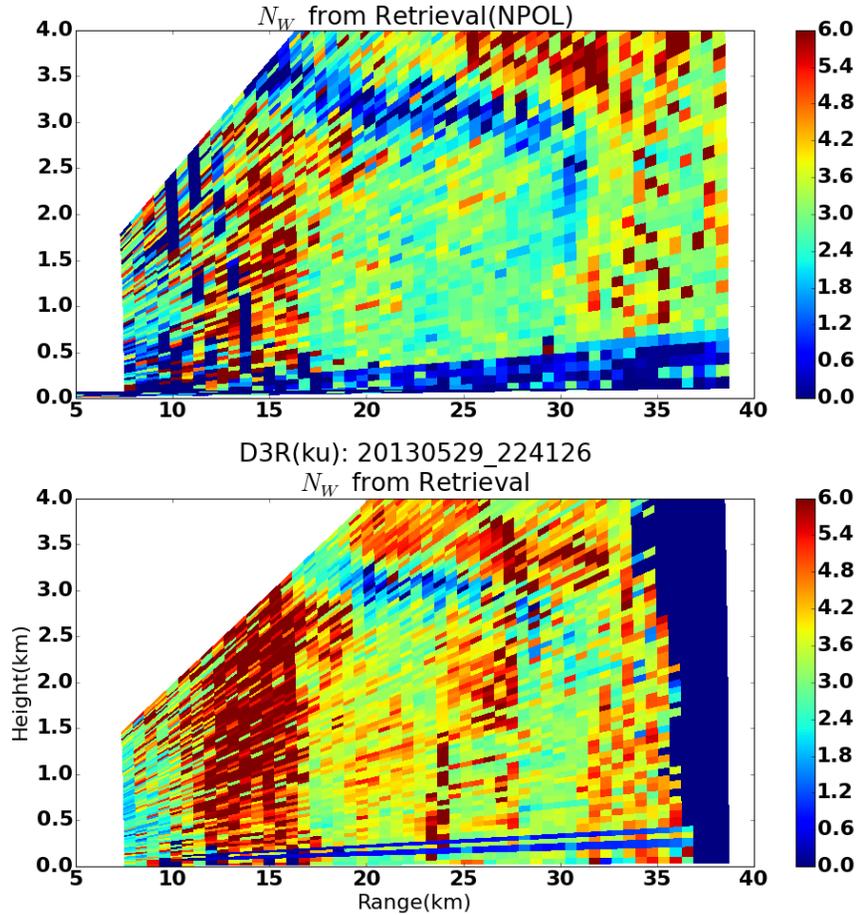
FIGURE 3.21. D3R-NPOL normlized intercept parameter comparison. Overall morphology agrees, though absolute values differ in the lower signal to noise region seen between twenty and thirty kilometers.

due to the lower signal to noise ratio in the fainter portions of the storm between twenty and thirty kilometers.

3.3.4. COMPARISON WITH GROUND INSTRUMENTATION. Most of the comparisons used so far to show validity of the retrieval have been forced to use indirect comparisons to show that the retrieved DSD adequately captures the radar measured parameters. This evaluation framework has it's limitations. In particular, it assumes uniqueness of the mapping between DSD fields, and radar fields. Due to noise in the measurements, evaluating the recreation
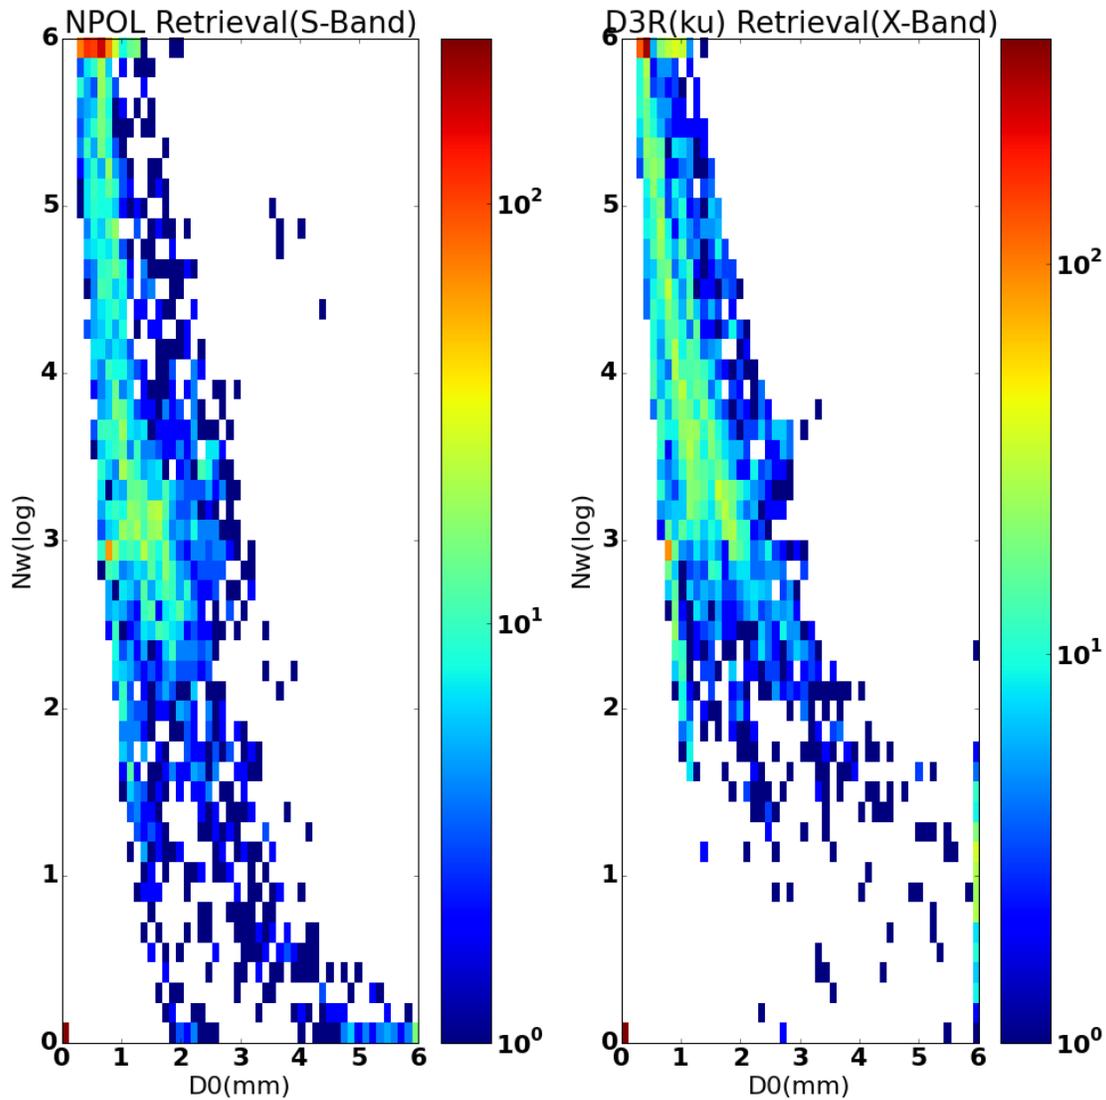
FIGURE 3.22. D3R-NPOL Microphysics Retrieval Comparison. While overall morphology agrees, the lower portion of the curve shows some disagreement between the two frequencies of retrieval. In particular, the higher frequency displays a smaller number of low intercept parameter values.

of the radar measure fields involves a noisy comparison as well. This sub-section will use another indirect comparison, although indirect in a different way.

Disdrometers as discussed in Chapter 2 are a ground based instrument that uses a variety of techniques to count rain droplets as they pass through a sensor area, giving the size of each drop, and often the speed. These are usually binned measurements and so the histogram

of the counts over a time period gives us an approximation of the drop size distribution for that time period passing through the sensor. Using the velocity, the data can be normalized to get a "representative sample" of the drop size distribution immediately above the sensor.

Unfortunately, what disdrometers do not measure is the drop size distribution in the air. This makes it somewhat hard to directly compare radar data with disdrometer data. This subsection will use disdrometers to compare with radar returns to provide another verification point keeping this in mind. It is important to remember that the measurement from the disdrometer is not measuring the same area, time, and distribution, that the radar is seeing. As the precipitation falls to the ground, it undergoes changes. Most comparisons of disdrometers with radars assume the Taylor Hypothesis[34], which loosely states that a storm is made up of multiple horizontal sheets that are acted upon by an affine transformation as they fall. Then given sufficient sampling, one could recreate the above earth conditions by working backwards, and thus time and space will act as a proxy for each other. While we know this is not true, modeling the storm dynamics as drops fall is not feasible. Then comparing radar and disdrometer data will have a certain amount of intrinsic error in the comparison. The closeness of the comparison is somewhat correlated with the spatial and temporal variability of the microphysical structure, and will be addressed more in Chapter 5.

The data this comparison will use is from a Parsivel Laser Disdrometers as seen in Figure 3.23 to provide the ground data. A case from IFloodS that took place on May 26th, 2013 is shown in Figure 3.24.

Applying the methodology given in Bringi and Chandrasekar[1] to the disdrometer measured DSD's from two disdrometers gives the D_0 and N_w time series shown in Figure 3.25 in panels (a) and (b). Displayed are the time series for two different co-located disdrometers. In

FIGURE 3.23. Parsivel Disdrometer at IFloodS *Courtesy of Iowa Flood Center*. This type of disdrometer uses an optical array to measure the drop size distribution as well as the velocities.
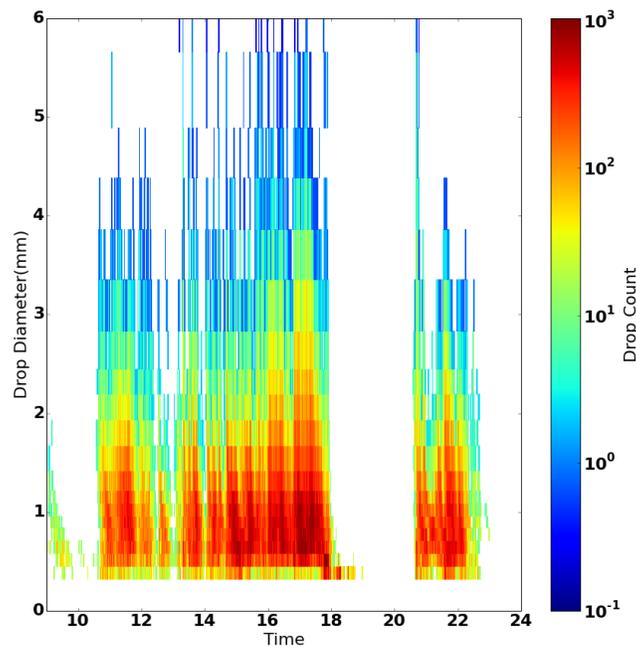
FIGURE 3.24. Drop Size Distribution Example from Parsivel from May 26th, 2013.

Panel (c), the overall parameterization is shown. The two disdrometers, denoted APU3 and APU4, are located about 100 meters apart. The two disdrometers show good agreement for

most time periods, although there are some time periods where the agreement between the disdrometers suffers, such as the time between 16Z and 18Z. Panel (d) shows the difference in microphysical estimates between the two disdrometers. Although most of the points match between the two instruments, there are a significant number of points with large differences. Some of this is explained by spatial and temporal variability of the microphysics, and the rest should be due to sampling errors and wrong-distribution errors.
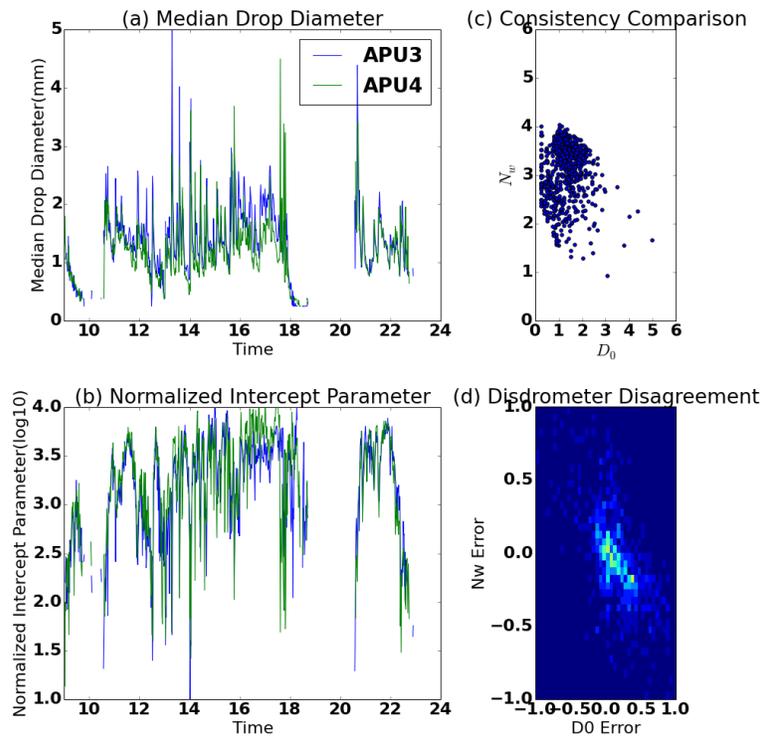
FIGURE 3.25. Microphysical Parameterization Time Series from Two Disdrometers. Shown in the left panels is the time series of the estimated microphysical parameterization. Panel (c) shows the self consistency comparison of the two disdrometers. Panel (d) shows the relative errors between the two disdrometers.

To characterize the between device sampling error, Table 3.8 shows the error rates using disdrometer APU3 to predict APU4. The normalized errors are less than four percent, with a between instrument bias of up to twenty percent. Mean absolute error is shown as well

for comparison with other algorithms. These numbers form a target for a radar retrieval to reach when compared with a disdrometer.

TABLE 3.8. Between Disdrometer Error for Co-located Disdrometers. This compares the error between co-located disdrometers 3 and 4. Shown are results for both the native time sampling of 1-minute increments, as well as a 7 minute windowed average.

	NSE 1-minute	NSE 7-minute	NB 1-minute	NB 7-minute
D_0	0.0156	0.024	0.155	-0.215
N_w	0.0399	0.005	-0.195	0.026

The disdrometer dataset from IFloodS contains data that is sampled in 1 minute intervals. This provides for a very noisy dataset with a large variance, much too noisy to characterize the large spatial area that the radar measurements cover. Also, due to not knowing the exact fall trajectory of the drops, there is a variable time delay between disdrometer and radar measurements. For instance, in the one minute period that the disdrometer samples data over, the average distance a 1 mm drop will fall is 123 meters, while a much larger 4 mm drop will fall about 500 meters. To slightly minify the effect of these two error sources, we will downsample the disdrometer data using a sliding average with a width of 7 minutes centered around the time of each radar sample. Each data point will be the average of two minutes into the future to account for the time taken to sample the volume by the radar, as well as 4 minutes into the past to capture some of the variability due to the fall speed. To illustrate the effect this has on the data, original, and temporally averaged data for the median drop diameter of APU4 is shown in Figure 3.26.

Figure 3.26 shows the decrease in variance in the data due to downsampling, while the mean is maintained. All radar comparisons will be made to this new averaged time series. For each radar file with data near the disdrometer, the retrieval algorithm was run and the
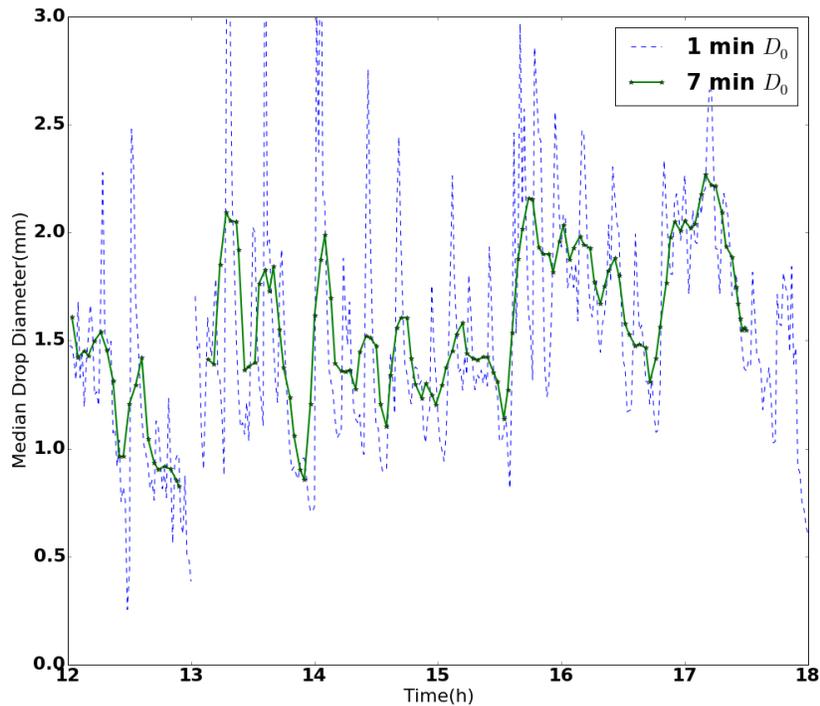
FIGURE 3.26. Median Drop Diameter: The Effect of Temporally Downsampling on Time Series. While the mean stays the same, as expected, the averaging operation decreases the variance. This helps to reduce the effect of errors caused by limited sampling resolution.

microphysical parameterization performed. To compare the radar data with the disdrometer the radar data must be similarly averaged, or else it is likely the bin being compared, did not actually have any drops falling on the disdrometer. After the retrieval, an annular region is averaged centered over each disdrometer consisting of 3km in range, and 8 degrees in azimuth, which corresponds to between 2km and 3.3km for the disdrometer results presented here. There were 134 radar files with data over the disdrometer in the time period of interest.

Shown in Figure 3.27 is the results of comparing the time series from the disdrometer with the time series from the radar retrieval. Good agreement can be seen between the retrieval and the disdrometer. The general behavior matches well, although the variance of the radar
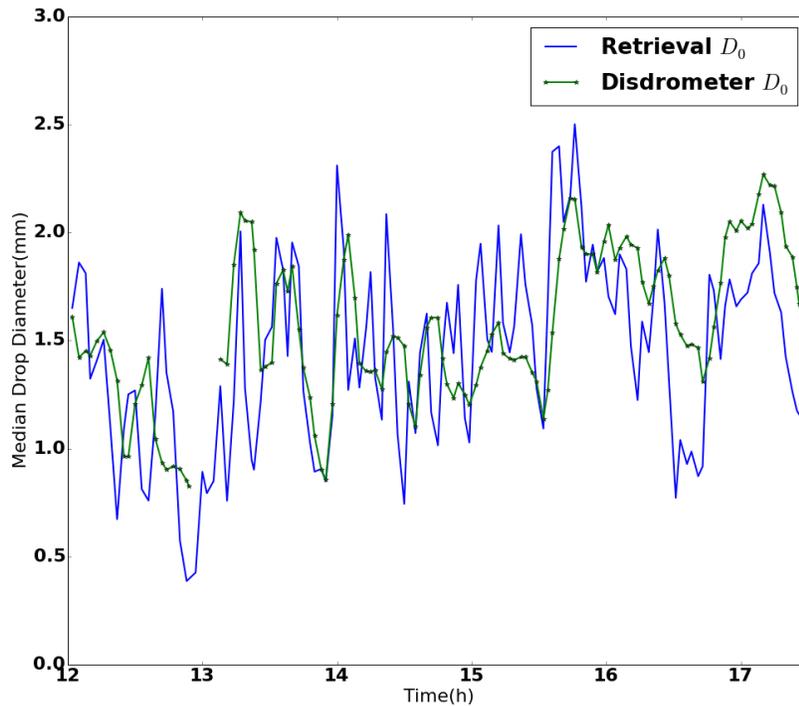
FIGURE 3.27. A Comparison of D_0 Between the Radar Retrieval and Disdrometer APU4 Measured DSD based on 134 radar sweeps.

retrieval estimates is significantly higher. Some of this is due to temporal averaging of the radar data. In particular, a persistent half-millimeter under-estimate takes place around 16:30Z in the figure and appears to last for several time steps. It is not known whether this is a mistake from the radar retrieval, or due to the spatial variability of the microphysics.

Unfortunately, the comparison of the retrieval with N_w is not quite as good. As shown in Figure 3.28, the comparison of the retrieved intercept parameter with that from the disdrometer has a larger absolute error. Again the mean of the data appears correct, but the retrieval shows larger fluctuations in it's estimate of N_w. Part of this could be due to spatial variability. The phase response at S-band is also lower and so estimation of N_w is expected to be harder at this lower frequency than some of the higher frequencies.
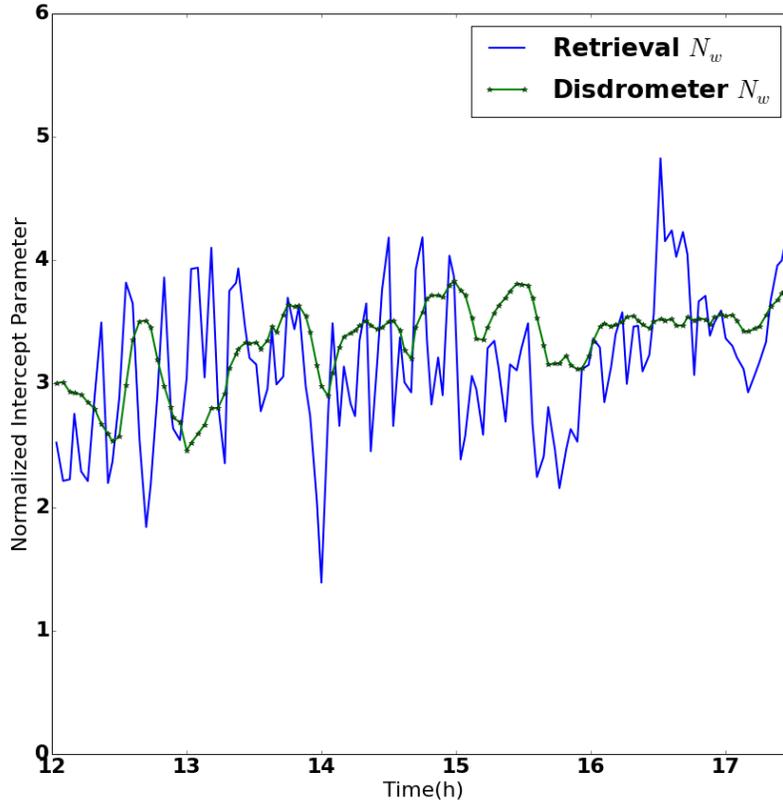
FIGURE 3.28. A Comparison of N_w Between the Radar Retrieval and Disdrometer APU4 Measured DSD.

TABLE 3.9. Relative Disdrometer Locations with Respect to NPOL. Disdrometers three and four are co-located. While other disdrometers exist, these are teh only ones within range of both NASA radars during the field campaign.

	Range(km)	Azimuth(∘)
APU2	15.2	128.7
APU3	24.6	129.9
APU4	24.9	130.5
APU5	48.1	131.8

For a more comprehensive look at the performance, this analysis was repeated for five

of the disdrometers, including a pair of co-located disdrometers. The relative locations of

these disdrometers with respect to the radar is shown in Table 3.9. Disdrometers two and five are located at their own sites, while disdrometers three and four are co-located (Within a kilometer of each other). This means while the general trend for all disdrometers should match loosely, the two co-located disdrometers will have a much higher agreement in their measurements. Also, the difference between the two co-located disdrometers gives us a rough measure of sampling error and spatial variability effects on estimation.

TABLE 3.10. Radar Retrieval Error with respect to Disdrometer. The two error metrics shown are the ones most commonly used to estimate the error between radar retrievals and ground instruments.

	N_w **NSE**	N_w **NB**	D_0 **NSE**	D_0 **NB**
APU2	0.030	-0.0052	0.0238	-0.1922
APU3	0.029	-0.082	0.019	-0.058
APU4	0.026	0.1104	0.0187	-0.086
APU5	0.038	-0.044	0.039	-0.229

Finally, the composite results for all disdrometers is shown in Table 3.10. The table shows the results are on par with the previously shown between disdrometer errors for 1 minute sampling. Also displayed in Table 3.8 is cross-disdrometer results for the 7-minute sequences used in this analysis. The error between the radar retrieval and the disdrometer results are higher for the 7 minute case, which is expected. The temporal averaging reduces the effect of spatial variability on the disdrometers.

Although this analysis was only for a single frequency of radar, it shows the radar retrieval has very good agreement with ground instrumentation with errors of between 2% and 4% with biases of between 0.4% and 22%. In particular, the retrieval does a much better job of retrieving a D_0 profile in agreement with disdrometers than it does with N_w. It is expected

as frequency of the radar increases, the estimation of N_w should improve due to the increased phase response upon which the N_w estimation is very sensitive.

3.4. Summary

This chapter laid out a retrieval framework for estimating microphysical parameterizations of liquid precipitation from radar observations based on solving the forward variational problem. The framework applies to arbitrary frequency radars. An error characterization was given based on different possible sources of noise and contamination using a simulated dataset. The algorithm was shown to be very resillient to realistic levels of measurement noise in the three radar parameters. It was also shown to work even in the presence of reasonable calibration errors in the radar measurements. The algorithm was then applied to several real datasets for different frequency radars. The ability of the retrieved drop size distribution to recreate the radar measured fields was used as a performance metric. The retrieved fields in all cases showed good agreement with radar measured fields. Finally a comparison between a series of radar retrievals and ground instrumentation showed very good agreement between the retrieval derived parameters, and the parameters measured at the ground.

CHAPTER 4

Network Retrieval

"Diversity and independence are import because the best collective decisions are the produdct of disagreement and contest, not consensus or compromise"

— James Surowiecki, The Wisdom of Crowds

4.1. Introduction and Background

The core subject of this dissertation is the improvement of microphysical retrievals by leveraging radar networks instead of single radar measurements. First, a justification of why a network approach offers improvements is provided, followed by discussion of one of the common problems when working with networks of radar instrumentation, non-coincident grids. An algorithm for generating network datasets from single radar measurements is developed, creating a way to characterize the performance of a network retrieval. Following this, a framework for multi-frequency non-colocated radar networks is presented. Results are shown for the test dataset, along with error characterization.

4.1.1. Benefits of a Network Approach. Ultimately, the goal of a network retrieval is to retrieve more accurate values than is possible from a single radar. Fundamentally, the network retrieval brings us three different advantages. The first benefit is to reduce errors by having multiple measurements from different angles of each range bin. Multiple measurements helps to cancel out the effect of measurement noise from any given radar. Additionally, based on independence of path integrated attenuation, multiple measurements of the same range bin helps to offset errors due to attenuation mis-estimation. For single frequency networks the intrinsic field can be assumed to be the same. This means each

radar provides an independent estimate of the same quantity. Indeed, the different between each radar also provides an additional quantity of information. By estimating attenuation, accurate rain rates can be calculated [47].

The second benefit is to help offset beam blockage, blind regions, and other areas where signal is not available to a single radar. This is a common issue with most radars in operational and research use. Full beam blockage eliminates the possibility of any measurement happening past the blocked location, while partial beam blockage corrupts the power measurements from the radar, causing them to be severely underestimated. Networks extend the range beyond what any one radar can manage, while being significantly more cost effective. This is one of the core tenents of the Collaborative Adaptive Sensing of the Atmosphere (CASA) program[48]. Multiple radars provide an increased field of view, as well as minimizing blind regions due to blockage and curvature of the earth effects(where the beam curves upward relative to the earth, causing low lying atmospheric phenomenon to be invisible to the radar.)

The third is to leverage the information returned from measurements at different frequencies. Each frequency has it's own benefit and by leveraging a range of frequencies we can get a better retrieval than the single frequency retrieval. For instance, low frequency radars suffer from very low levels of attenuation relative to higher frequency radars. Where a high frequency radar would record inaccurate measurements due to attenuation, the lower frequency radar will provide much more accurate measures. For this reason, lower frequencies often have a much higher coverage area. On the other hand, due to the increased phase response at higher frequencies, estimation of rain rates and phase effects can be significantly more accurate at high frequencies. Higher frequency radars can also be operated at lower powers and generally at higher range resolutions. In Chapter 5, an explanation for the

benefits of multiple frequency measurements of the same field will be explained based upon different scattering regimes.

Averaging, while straight forward, is not necessarily the most elegant, or intelligent, way of combining this data as shown in the next section. An ideal fusion strategy needs a more nuanced approach to combining the parameters that uses the physics of the measurements.

4.1.2. OBJECTIVE ANALYSIS: COMING TO A COMMON GRID. One of the crucial steps in any network rterieval is to bring the measurements from the different radars onto a common grid. As each radar is located at a different location, it is geometrically impossible to use a spherical grid. The most common choice of grids is the common cartesian grid. To convert measurements taken on a spherical grid to a cartesian grid adds errors to the underlying measurement due to the interpolation process. This process is called Objective Analysis [49].

There exists a large number of different methodologies for gridding radar data, but the vast majority follow the "Region of Influence" approach. Simply put, to determine the value at a grid point on the cartesian grid, you find all radar measurements for that field within a distance r_i denoted the region of influence (ROI). Then the value of the measurement at that grid point is the weighted average of all of the points within the ROI. The weighting function varies, but is usually similar to the form

$$(53) \qquad\qquad w_{i,j} = \frac{f(r_{i,j}^2)}{\sum_i w_{i,j}}$$

where the distance $r_{i,j}$ is the distance from the ith point within the region of influence to the grid point j. For instance, the very common Cressman interpolation kernal is given by

95

$$w_{i,j} = \exp \frac{-r_{i,j}^2}{\kappa}$$

with κ representing the fall-off parameter that controls how tight the spatial weighting is.

a common extension to the traditional objective analysis algorithms is the addition of an adaptive ROI that changes based on proximity to the radar. Close to a radar, points are very dense and a large ROI causes an excessive amount of smoothing. Far from the radar however, points become much more sparse and often a nearby point cannot even be located within the ROI leading to "holes" in the interpolation. The same ROI used in different locations then gives different results. Newer algorithms allow the ROI to increase as grid points move further away from the nearest radar. This approach performs significantly better as can be seen in the results by Trapp and Doswell [49]. These same algorithms can be used to merge multiple radars at low frequencies(Primarily S-Band) by treating the measurements of each radar as a cloud of points. Some algorithms, such as the Barnes[50] and Cressman[51] interpolation algorithms that are most commonly used, are based around a ROI based interpolation scheme, with iterative updates to minimize the interpolation error in the reconstructed field. This however only works for the scalar quantities such as the power measurements and K_{dp}. When attempting to grid Φ_{dp}, the path dependence of the parameter makes ROI based approaches useless. Similarly for velocity, which is a vector quantity, the gridding operators do not have the capability to handle the gridding of vector data. There are other approaches for vector quantities not covered here as this work does not utilize velocity in any meaningful way.

The multi-frequency case further complicates the problem. The objective analysis algorithms all assume a true underlying field, that the radar measurements are approximations of. While this is true of the microphysical fields, it is not true of the radar measured parameters at multiple frequencies. Given the same microphysics, different frequencies will take different measurements. Indeed, the difference in measurements for differing frequencies is something that both this work, and the greater literature will seek to exploit. This means that any algorithm that attemps to merge multiple frequency radar data at different locations must work at the microphysical level, and not at the radar measurement level. As the relation between the different frequency measurements is highly nonlinear, it is not feasible to renormalize the data for any networks except possibly lower frequency S-band and X-band combined networks. The resonance properties at C-band, as well as the very high nonlinearities at higher bands makes this approach nontractable.

In this work, a different approach will be taken. Multiple radar observations will be used to derive co-estimates of microphysical parameters that work for each radar.

4.2. Generation of a Network Dataset from a Single Radar

The same troubles with characterizing a retrieval algorithm that presented themselves in Chapter 3 with respect to the lack of ground truth, also complicate the analysis of the network retrieval algorithm. This section lays out a methodology for generating network datasets with arbitrary characteristics from an underlying radar scan, essentially turning the data from one low frequency radar into a network of multi-frequency radars suitable for analysis. This type of dataset is vital for accurate error characterization of the network retrieval algorithm, and forms an important tool for cross comparing radars in a network setup.
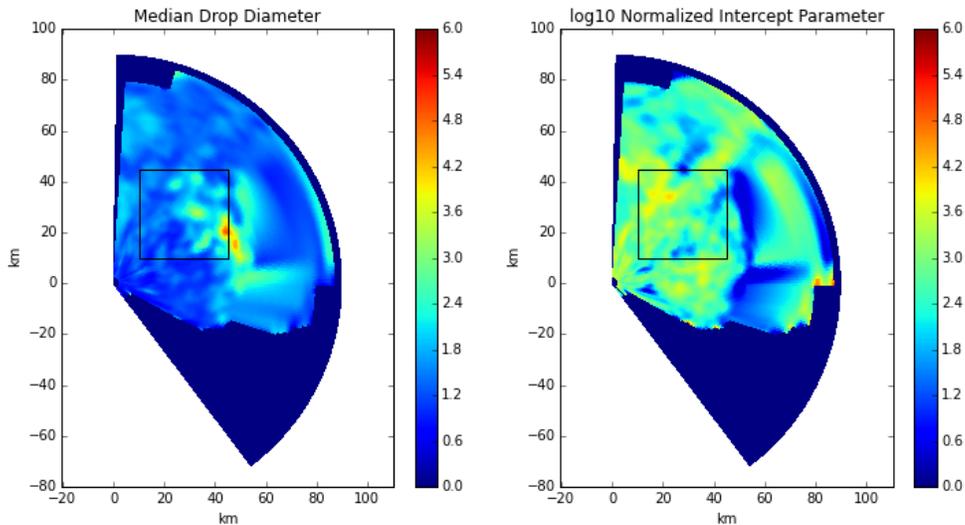
FIGURE 4.1. Base Microphysical Parameterization Derived from S-Band Radar Observations by a microphysical retrieval based on data originally sampled by CSU-CHILL. This dataset is courtesy of Le and Chandrasekar and based upon work by Lim and Chandrasekar [2].

Lim and Chandrasekar showed a method for simulating X-Band radar measurements based on measurements at S-Band [52]. Their work showed three different methods for generating X-Band measurements. The first relied on modeling the nonlinearities between S-Band and X-Band Z_{dr} and then frequency scaling the phase measurements, while assuming reflectivity was constant between the two. This gives the intrinsic fields, which can then be attenuated based on the relationship between K_{dp} and attenuation. Their second method used a look-up table to find microphysical tuples that matched the measured parameters. The third approach used a full microphysical retrieval, and then T-Matrix scattering. In what follows we extend this method significantly to handle the network approach.

The first step is to perform a microphysical retrieval at S-Band to get the field of parameters. While an accurate retrieval is desired, small errors are acceptable for this use case. The goal of the microphysical retrieval is to get a microphysical field that represents the true spatial variability of the radar parameters instead of being a purely synthetic spatial

allocation of rain. In this case, a different algorithm than the single radar retrieval algorithm developed in Chapter 3 is used. This is to prevent the encoding of *a-priori* assumptions. The results of this step are shown in Figure 4.1. This is based upon a dataset provied by Minda Le and Chandrasekar.
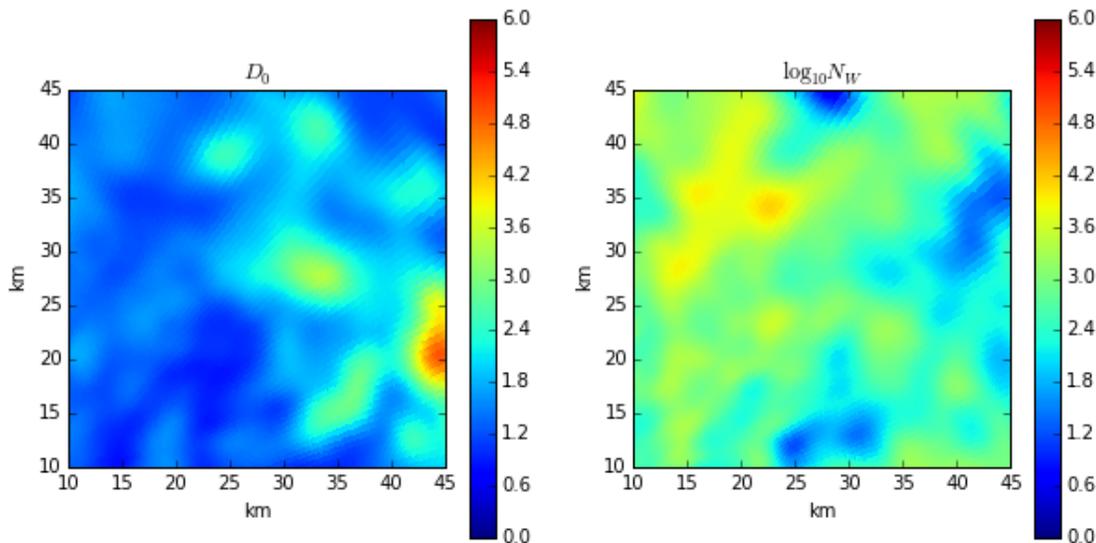
FIGURE 4.2. Microphysical region that will be cross-gridded to non-colocated spherical coordinate systems representing each radar. Data is located on a spherical grid.

Next, the area in the square region is isolated as the area of interest. This is chosen as a 60 km square (3600 km^2 area) to mimic several operational multi-frequency networks. It is also chosen as a region with both active and quiet areas. In particular, it is chosen such that some of the radars will experience significant path integrated attenuation at areas where other radars will experience very little. There is also a mix of weaker areas and stronger areas making this a good test case. The isolated region is shown in figure 4.2.

The next step is taking this square region (where the samples are still on a spherical coordinate system) and choosing the placement of the radar network. In figure 4.3, four

radar positions are displayed with 40 km range rings displayed around each radar point. Each of these radars is 20 km from grid center located at rotations of 90° from north. Next, new spherical grids are created centered at each radar position. For the test setup here, a 120° degree section of this grid at one degree angular increments, and 100 m range increments is used. This limits the data processing required as the area of interest is the center region. However, this limitation is not required and could be arbitrarily removed depending on user needs. This 120° field of view from each radar gives us a large amount of overlap, while also providing areas where less than the full four radars have data. This is desired for the test case, as operational radars often have blind spots, and indeed this is touted as the benefit of a radar network.

The benefit of this spherical-to-spherical conversion is a lower overall error than the intermediate step of converting to a Cartesian grid. Every interpolation step adds an amount of error. Most interpolation algorithms are low-pass filters, and the averaging process removes higher frequency spatially localized information. By limiting the number of intermediate steps, the end to end error will be reduced in most cases. In addition, the microphysical parameters D_0 and N_w are linked, and gridding each one individually introduces errors that move the distribution away from a physically realizable one. For dense griddings, this is not a strong source of error, but the more interpolation steps involved, the more change that these two parameters lose their underlying relationship to each other.

The results of the spherical-to-spherical cross gridding are shown in Figure 4.4. Shown is the microphysical field for each radar overlaid on top of one another. It is clear that there is very good agreement, with edges of each radar's field of view not showing any perceptible difference from the next radar.
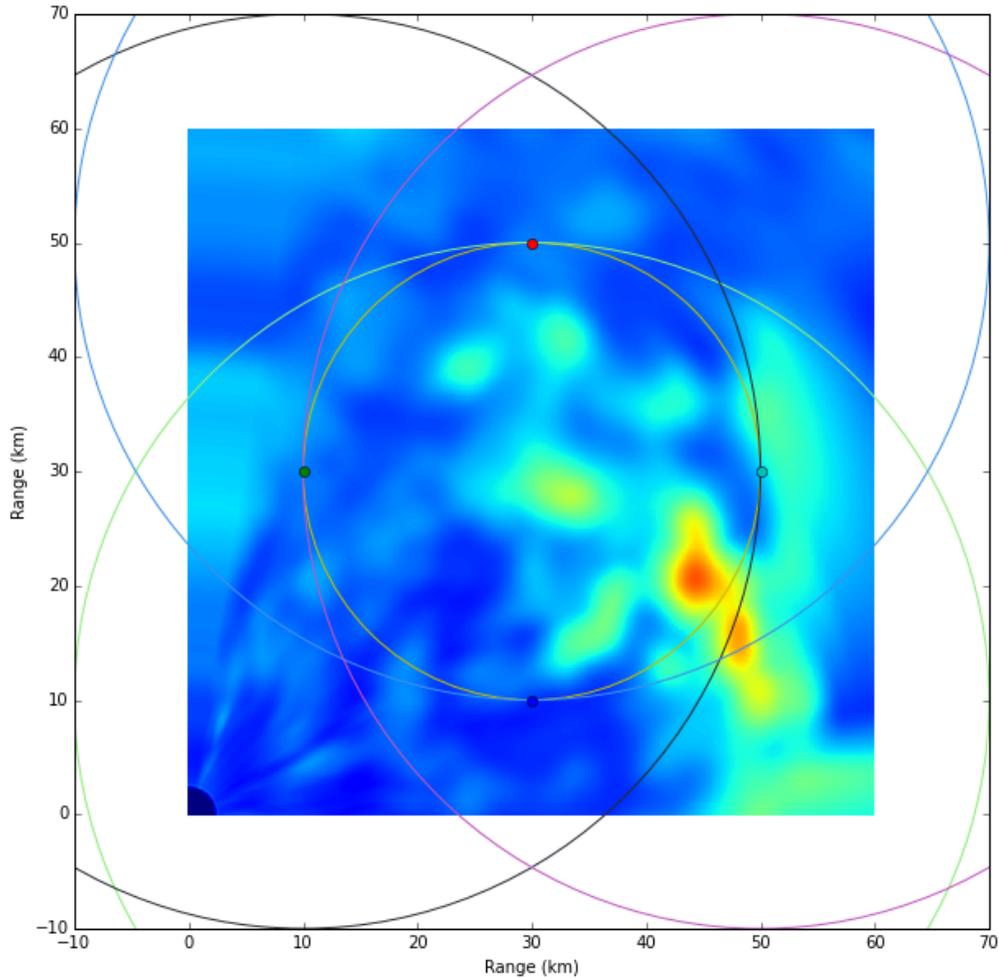
FIGURE 4.3. Locations and fields of view of the four simulated radars. The image shows the median drop diameter field. The four colored dots show the location of each radar. The simulated network will have a 120 degree field of view for each radar centered on grid center. Shown on the map is a range ring of 40 km centered on each radar. The underlying data is on a spherical grid. Finally a 10 km range ring is shown in the center to represent a hypothetical intensive operations period area(IOP).

The next step is to use T-matrix scattering to calculate the radar measured parameters at each of these bins. For this, frequencies for each radar must be chosen. In this case, radar 1, the southern-most radar, will be S-Band. Rotating around the circle clockwise
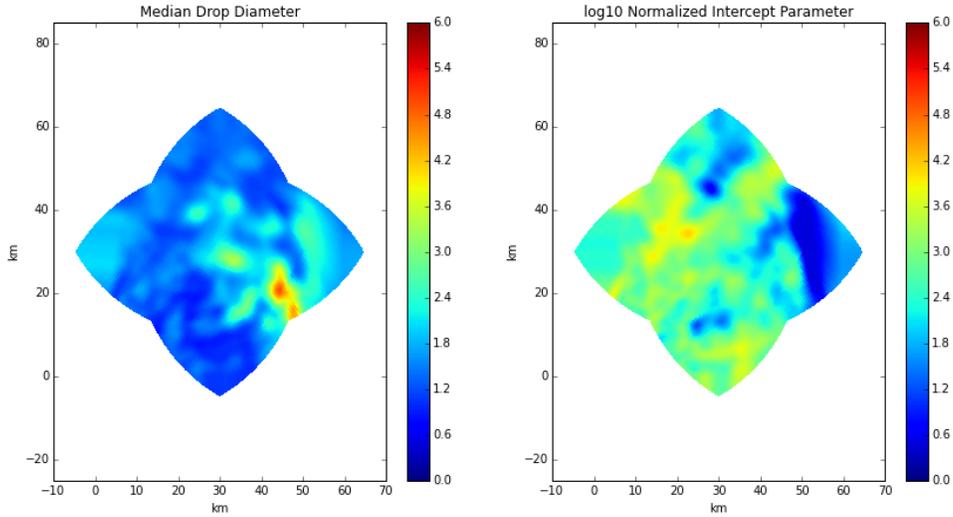
FIGURE 4.4. Overlapping microphysics from each radar grid showing good agreement between new coordinate systems. Left panel shows the median drop diameter, while the right panel shows the normalized intercept parameter on a log scale. Radar Z-order is determined by frequency, with higher frequency being displayed at higher relative layering.

will be radar 2 at C-Band, radar 3 at X-Band, and radar 4 at Ku-band located in the easternmost position. Each radar ray is independent, and path integrated variables such as path integrated attenuation(PIA), path integrated differential attenuation(PIDA), and Φ_{dp} are calculated, as well as the point measurements of Z_h, Z_{dr}, and K_{dp}. Both intrinsic fields(Those without attenuation), as well as the radar measured fields(fields with attenuation applied) are calculated. It is important to note that these fields will not match the underlying field this dataset was derived from due to the differing frequencies, locations, and path integrations.

Shown in Figure 4.5 is the result of the entire network generation process. By column, each of the radar measured parameters is shown, while by row each of the different radars is displayed. A few things become immediately evident. First, the morphology of reflectivity
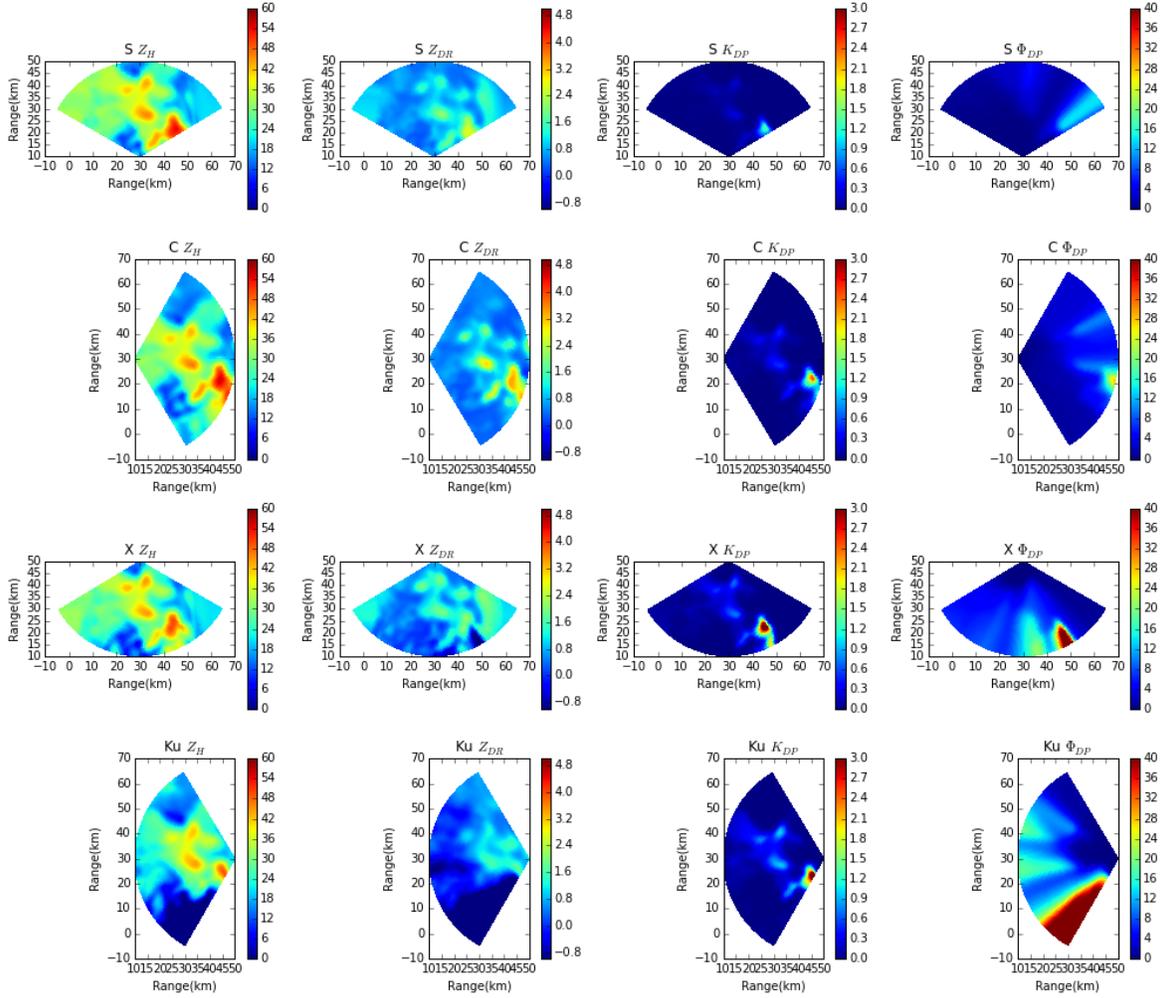
FIGURE 4.5. Radar measured fields for each radar in the network generated by the algorithm. There are four radars of differing frequencies going down the page, in the order of (S,C,X,Ku). The radar measured fields Z_h, Z_{dr} K_{dp} and Φ_{dp} are displayed column wise for each radar.

matches on all of the radars, with differences due primarily to attenuation. Examining Ku-Band Reflectivity, the region located at around 200° experiences heavy attenuation as it goes through a storm core, ultimately resulting in an extinct signal. Although this is synthetic data, and so extinction could be accounted for and fixed(As there is no simulated noise floor), any value of reflectivity below zero is truncated out of the dataset in the retrieval process to simulate more realistic conditions. A similar phenomenon exists for Z_{dr} linked to differential

attenuation, as can be seen on the Ku-Band plot. As frequency increases, the attenuation is expected to increase. We see visually noticeable attenuation effects in the X-band as well.

The next major difference is the phase components. As expected, higher frequencies have significantly higher phase responses. For lower frequencies (X and below), a usual and safe assumption is that K_{dp} scales linearly with frequency. As K_{dp} is not affected by attenuation, it is commonly used to correct attenuation at higher frequencies, as well as provide a measurement immune to calibration biases of the radar. This dataset provides a multi-frequency dataset for the network retrieval, as well as ground-truth to characterize the performance and errors in the retrieval.

Shown in figure 4.6 is another overlaid plot for reflectivity and differential reflectivity for all of the radars. In the top row, the intrinsic variables are plotted. These are the values of the field, that have not had attenuation applied. Large discontinuities are noted in both reflectivity and differential reflectivity. This means, even given a perfect attenuation correction algorithm and a perfect K_{dp} estimation algorithm, the underlying fields will not match. This has one very large consequence for this work. Any attempt to merge multiple measurements from each radar must be done on a common field. The intrinsic fields for each radar, are not the same. This means, to combine multiple radars at different frequencies, we must work in the microphysical domain, as this is the common field from which these measurements are derived.

An idea to get around this requirement is to try to work with derivative scaled fields. For instance, one could scale K_{dp} by frequency, and try to estimate the nonlinear relationship between differential reflectivity at the different frequencies and work with a scaled version of the radar measured parameters. Unfortunately, even if restricted to a normalized gamma distribution, the relationship between parameters at each frequency varies highly based on
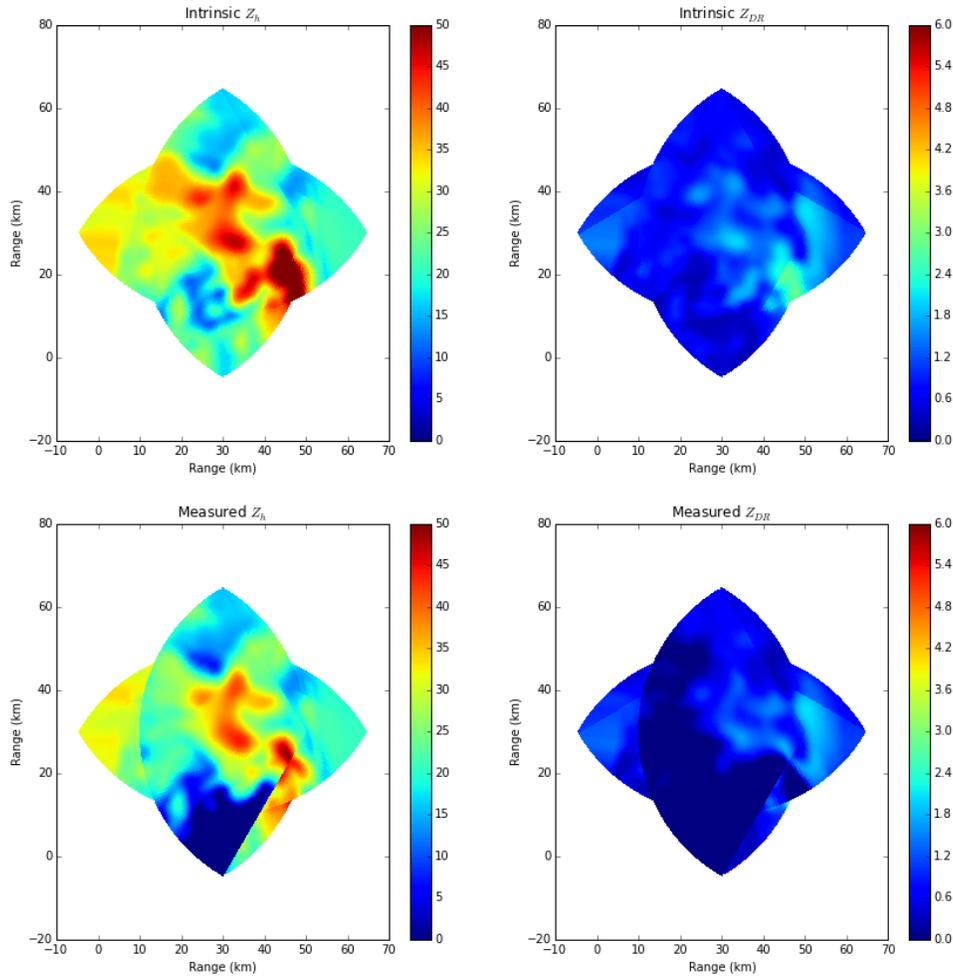
FIGURE 4.6. The two power measurements overlaid for all four radars, both intrinsic in the top row, and measured in the bottom row. This demonstrates why merging must be done at the microphysical level, instead of the radar measurement level.

temperature, and shape relationship. This means multi-frequency "point-cloud" approaches are inadequate for multiple frequency cases. While this is true at lower frequencies(K_u and below), it is even more true when Mie scattering at higher frequencies gets involved. While Mie scattering can cause resonance regions in C-Band, and nonlinearities at X-Band and Ku-Band, at higher frequencies the results get even more complicated. As we will show in

Chapter 5, at Ka-Band it is possible to get negative K_{dp}'s, and phase components no longer scale linearly with frequency. Even power measurements at higher frequencies suffer due to the Mie regime resonances, causing the relationships to be very challenging to predict.

4.3. Network Retrieval Framework

The problems with combining multiple frequency radar measurements at the moment level was described above. It is much more natural to combine the radar measurements at the microphysical level, where the underlying fields match. Unfortunately, radars don't measure the underlying microphysical field. However, the correct microphysical retrieval should minimize the observed error between all radars involved. Additionally, there should be one unique microphysical field with this property as the radar measurements are all derived from the same microphysical observations, corrupted by noise and measurement error, and passed through nonlinear processes. The network retrieval framework in this work will leverage this to combine multiple radar measurements.

The network retrieval framework can be split up into the following steps

(1) Single Radar Retrieval to estimate intrinsic fields for each radar.

(2) Definition of analysis grid.

(3) For every grid point, find covering patch points of each radar.

(4) Minimize fusion cost function subject to patch point weighting.

We'll start with a mathematical definition of the framework before moving on to implementation details.

4.3.1. Mathematical Development of the Framework. Assume a network of n radars, indexed by i, each operating at a frequency f_i, located at cartesian points $l_i = (x, y, z)$

with at least partially overlapping fields of view. Each radar yields a series of measurements over a spherical grid centered at point p_i with indices $r_i(j,k), \theta_i(j,k)$ given by the tuples

$$(55) \qquad \tilde{Y}_i(j,k) = (\tilde{Z}_{h,i}(j,k), \tilde{Z}_{dr,i}(j,k), \tilde{K}_{dp,i}(j,k))$$

Where each of these measurements is contaminated by attenuation and measurement error as given in equations 13, 14, 16. The first step is to pass these through the single radar retrieval developed in Chapter 3. This is used to calculate three things. First is an estimate of the intrinsic radar fields, free of attenuation and measurement noise. The second is estimation of a spatially localized K_{dp}. Third is the recreation error described later on in section 4.3.2, which gives a measure of the confidence of the estimation of each of the intrinsic fields. Additionally, a microphysical parameterization is provided, but not utilized yet in this retrieval to keep it as general as possible. It is possible to substitute in another algorithm capable of estimating the intrinsic fields accurately, as well as a good spatially localized estimate of K_{dp} at this point. Denote estimates of the intrisic microphysical field for each radar i at grid point j,k by

$$(56) \qquad \hat{Y}_i(j,k) = (\hat{Z}_{h,i}(j,k), \hat{Z}_{dr,i}(j,k), \hat{K}_{dp,i}(j,k))$$

Next, define a cartesian analysis grid with resolution δ_x, δ_y indexed by (u,v). For simplicity, it will be assumed that $\delta_x = \delta_y = \delta_g$ for the rest of this explanation. Then for every point (u,v) in the analysis grid define the set

$$\zeta_i(u, v) = \{q_i | (u, v) \in \mathcal{H}(q_i)\} \tag{57}$$

where

$$q_i = (p_1, p_2, p_3, p_4) \tag{58}$$

is a tuple of four points on the spherical grid defined for radar i, and

$$\mathcal{H}(q_i) = \{\sum_{i=1}^{|q_i|} \alpha_i x_i \mid (\forall i\colon \alpha_i \geq 0) \wedge \sum_{i=1}^{|q_i|} \alpha_i = 1\} \tag{59}$$

is the convex hull of the set of points. This gives the set of points of each radar that form the covering patches. Then form the set

$$\chi(u, v) = \{\min(\mathcal{H}(\zeta_i(u, v)))\forall i \in \mathcal{N}_4\} \tag{60}$$

where min denotes the minimum area covering. This is the set of bounding patches from each set of radar measurements with the smallest area. Each point in q_i corresponds to the nearest corner point in the bounding patch. These represent the points of four independent measurements from each of the n radars. The goal is to find the microphysical parameters (D_0, N_w) that best explains all of these measurements. An alternative approach would be to use the nearest neighboring point to the analysis grid point from each radar. Unfortunately,

this fails in areas of high spatial variability, and is also highly sensitive to noise and errors in each grid point. By using the minimally covering patch, the variability of the measurement is taken into account, as well as reducing the errors from the single radar retrieval, as well as anomalous errors due to underlying instrument measurement errors.

Next we setup the cost function based on minimizing the MSE error between the measurements, scaled by a weighting function discussed in the next subsection. This cost function takes the form

$$
\begin{aligned}
J_n(\chi(u,v)) = & \sum_{i=1,t=1}^{i=4,t=4} \left(\frac{\left(\bar{Z}_h(f_i, D0, Nw) - \tilde{Z}_{h,i}(p_{i,t}) \right)^2}{w_{h,i}(p_{i,t})} \right) \\
(61) \quad & + \sum_{i=1,t=1}^{i=4,t=4} \left(\frac{\left(\bar{Z}_{dr}(f_i, D0, Nw) - \tilde{Z}_{dr,i}(p_{i,t}) \right)^2}{w_{dr,i}(p_{i,t})} \right) \\
& + \sum_{i=1,t=1}^{i=4,t=4} \left(\frac{\left(\bar{K}_{dp}(f_i, D0, Nw) - \tilde{K}_{dp,i}(p_{i,t}) \right)^2}{w_{dp,i}(p_{i,t})} \right)
\end{aligned}
$$

Where the $\bar{Z}_h(f_i, D0, Nw)$ denotes the results of scattering at frequency f_i the microphysical parameters D_0 and N_w for the relevant radar parameter (One of reflectivity, differential reflectivity, and specific differential phase). The tilde notation $\tilde{Z}$ denotes the intrinsic field estimate from the single radar retrieval. The weighting functions in the denominators will be addressed in the next subsections.

Given this cost function, the solution for the network retrieval $\mathcal{G}$ at each grid point (u, v) is given by

$$(62) \qquad \mathcal{G}(u, v) = \underset{D_0, N_w}{\arg \min} \, J_n(\chi(u, v))$$

The network retrieval process is shown graphcially in the flowchart in Figure 4.7 for a four radar, four frequency case such as that in the test case developed above.

To modify the network retrieval to use an alternative single radar retrieval, or algorithm capable of estimating intrinsic fields, requires only modifying the inclusion of the recreation error described below. Additionally, as the radar combines the measurements at the microphysical level using the output after scattering, it is not limited to Rayleigh scattering regimes like many other algorithms. This means it is possible to incorporate elements such as negative K_{dp}'s at higher frequencies such as K_a. It also does not rely on *a-priori* gridding of the data. In addition, there should be some level of resistance to errors caused by advection and differences in temporal sampling between the radars by using the covering patch approach instead of the nearest neighbor estimation.

4.3.2. RECREATION ERROR. An important aspect of the single radar retrievals is how well does the estimated microphysics represent the measured radar parameters. Parameterizations which more closely resemble the radar measured moments have a higher likelihood of being correct than microphysical parameterizations that do not closely match the radar parameters. This can be used as a thresholding parameter to denote success of the retrieval, or as a measure of confidence in a more probabilistic sense. For the single radar retrieval, a recreation error term for the i^{th} radar, and the point (j, k) is defined as

$$(63) \qquad r_e(f_i, j, k) = \max \left(r_e^h(f_i, j, k), r_e^d(f_i, j, k), r_e^p(f_i, j, k) \right)$$
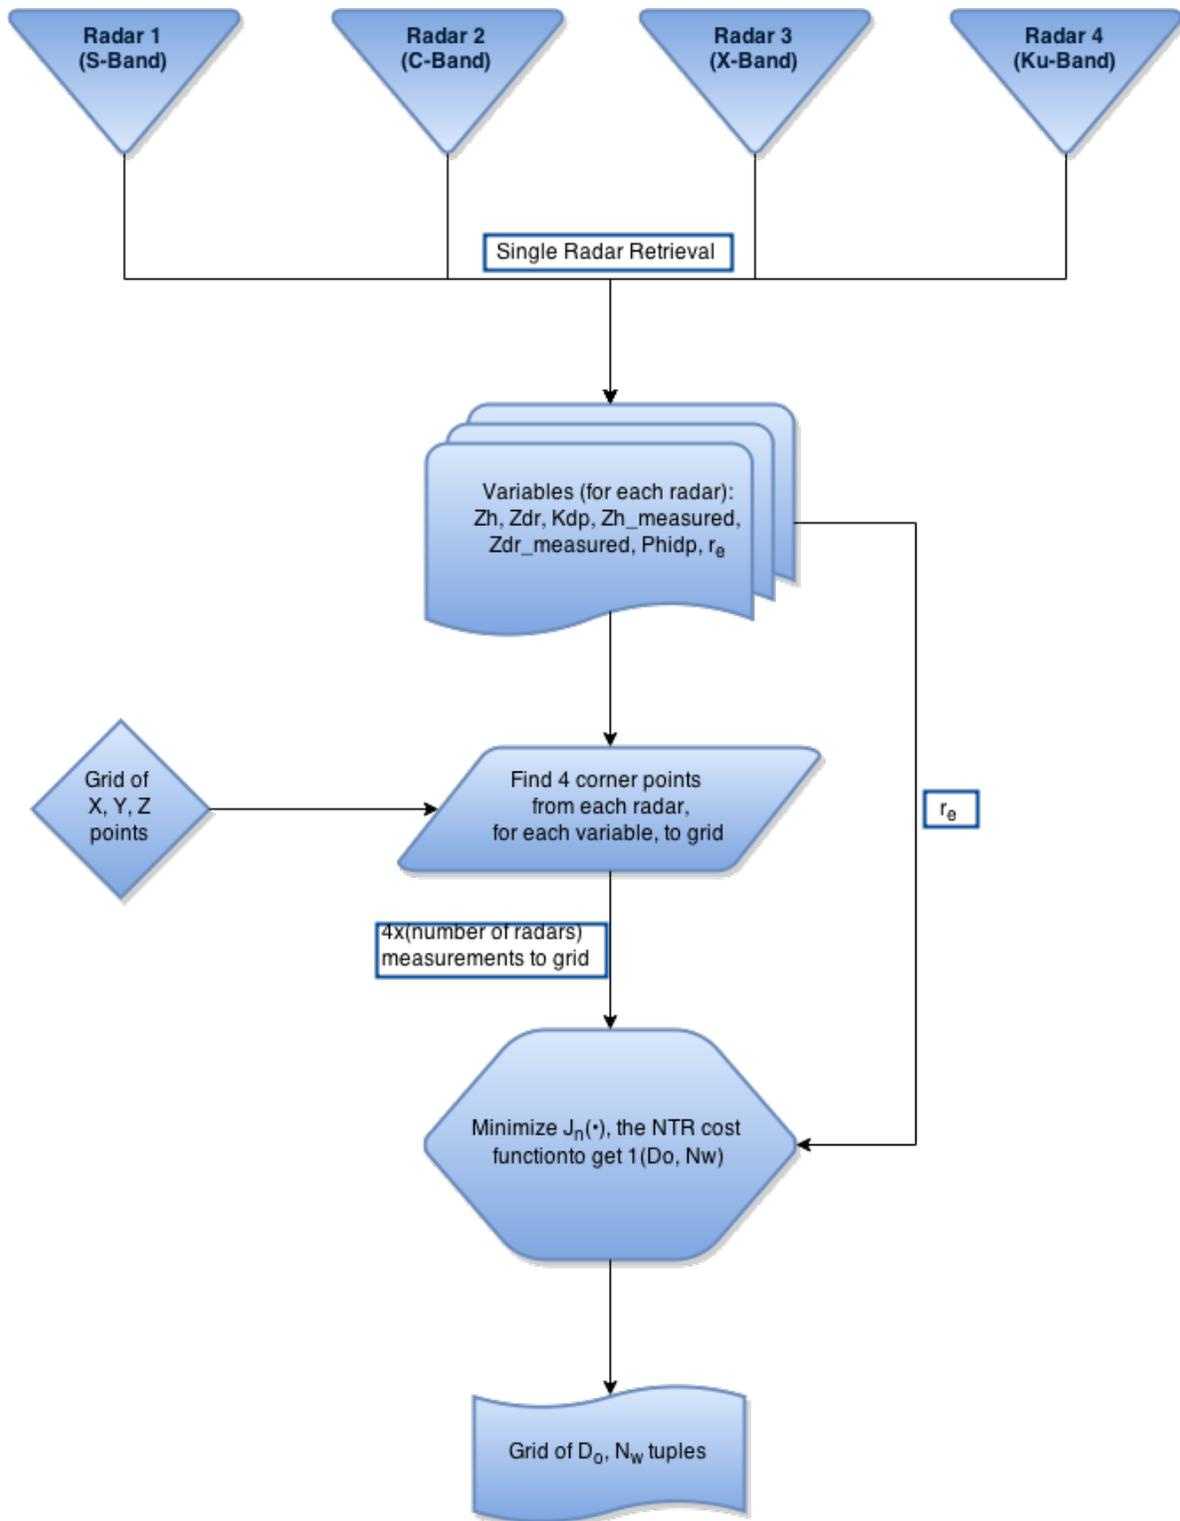
FIGURE 4.7. Flowchart of overall network retrieval process for a four frequency four radar network. This process scales out to n radars.

where the three arguments are the component recreation error of reflectivity, differential reflectivity, and differential phase respectively. The component recreation errors are of the form

$$(64) \qquad r_e^h(f_i, j, k) = \frac{Z_{hm}(f_i, j, k) - \hat{Z}_{hm}(f_i, j, k)}{Z_{hm}(f_i, j, k)}$$

This represents the normalized error of the recreation of each of the radar measured parameters. If used as a threshold for the network retrieval algorithm to filter out data used in the fusion process, a cutoff value must be established. Empirically, for the radars used in the test dataset, and the D3R/NPOL case, a value of 0.3 was found empirically to work. This will be dependent upon the noise in the radar measurements however. A particularly noise phase measurement, as seen on some cheaper radars, can cause the phase term to dominate. In this case, a rescaling of the phase term by the expected noise power is suggested. If used as a confidence parameter in a Bayesian update step, the recreation error needs to be scaled by the expected variance of the measurement.

The recreation errors for each radar in the test dataset are shown in Figure 4.8. In general, the single radar retrievals accurately represent the radar measured parameters. In low signal regions however, there is a noticeable increase in recreation error, especially for the S-Band radar. This is likely due to the phase terms being so small that any deviation is a significant error percent wise. This could be detected and accounted for by an automatic adjustment. Alternatively, the approach taken here was to allow those values to be filtered out, instead placing more trust in higher frequency measurements that had a stronger phase response in those regions. The second significant area of error is around the extinction point to the southwest of Radar 4, displayed in the bottom right panel. Right before the signal
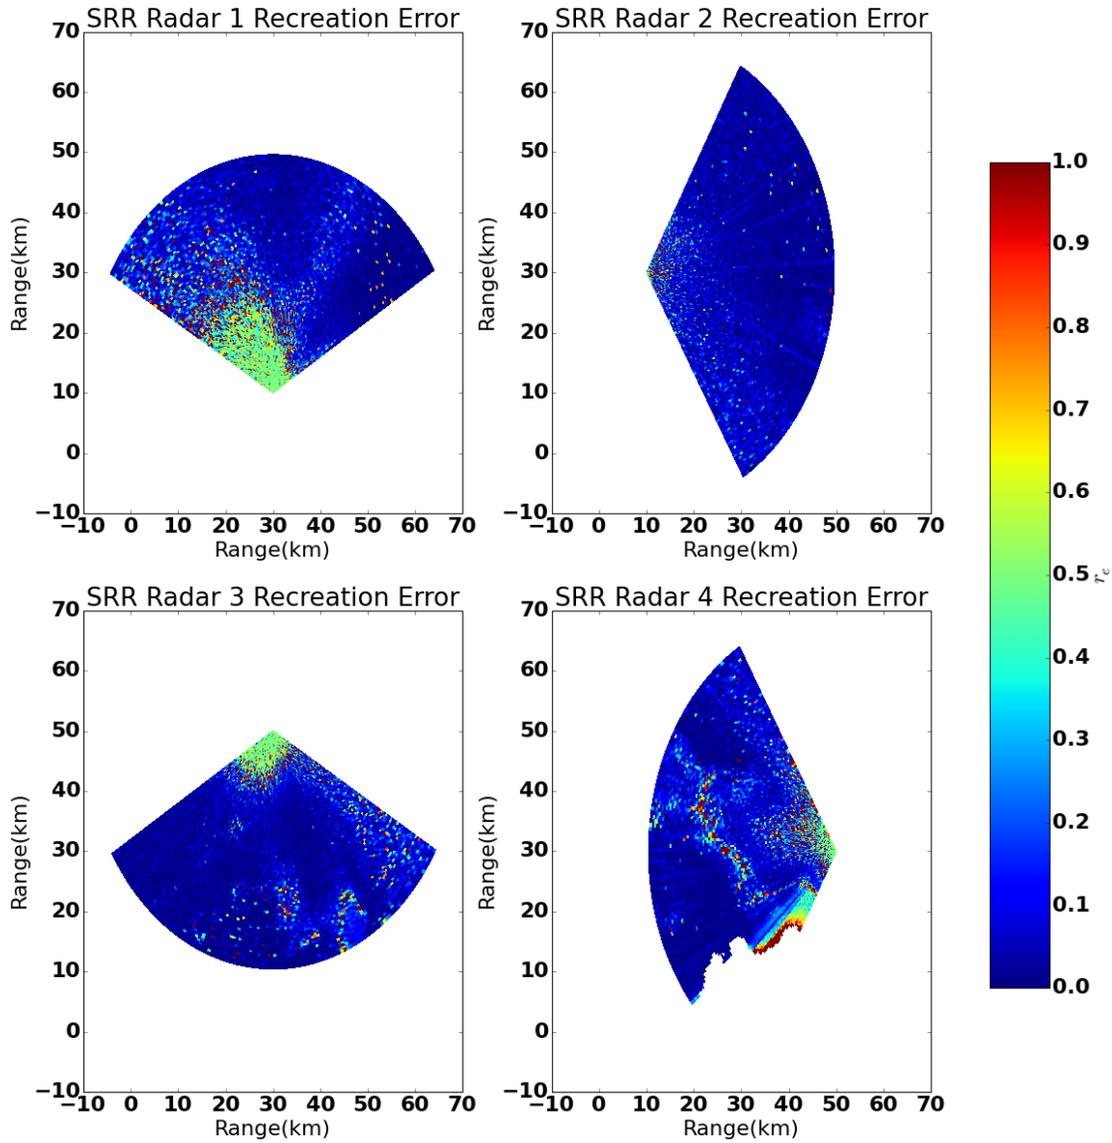
FIGURE 4.8. Recreation error r_e for the four radars in the test dataset. In order the frequencies of the radars are (S,C,X,Ku)-Band. Recreation errors are generally low, although there are significant increases in very low signal regions, as well as in regions of extinction at higher frequencies.

goes extinct, errors start to significantly increase. This is likely due to declining signal level.

The further radialy from the radar along these rays have no signal and so a recreation error

cannot be calculated(or alternatively, can be assumed infinite).

4.3.3. WEIGHTING FUNCTION. The final key piece of the network retrieval is the weighting functions w_h, w_d, and w_p in the cost function. Assuming all measurements were truly independent, these could be calculated by calculating the local linearizatino of the jacobian of the scattering function and estimating an inverse operator at that point. Unfortunately, this is both computationally expensive, does not adequately factor in non random errors or recreation errors, and is numerically highly unstable. This also assumes the estimated intrinsic field contains only minor errors, or else the estimation of the local linearization of the jacobian is incorrect.

A more tractable and stable approach is to use a combination of estimated variance of the instrument measurement error, recreation error, and distance from the grid point. Given the set of parameters $re_i(p_t)$ denoting the recreation error described above, $r((u, v), (p_t))$ the distance from the analysis grid point to the radar measurement at p_t, and the estimated instrument variance, which is assumed to be known (and on the order of (0.8, 0.2, 1.5) for (Z_h, Z_{dr}, Φ_{dp}) for most radars) the weight parameters can be written as

$$
w_{h,i}(p_t) = \begin{cases} \alpha_{h,i} r((u, v), (p_t))^\beta \sigma_h^\gamma, & \text{if } re_i(p_t) \leq 0.3 \\ \\ \nu_{h,i} \exp^{\rho re_i(p_t)} r((u, v), p_t)^\beta \sigma_h^\gamma & , \text{else} \end{cases}
$$

The weight parameters are of the same form for each measurement. The parameters (α, β, γ) can be tuned based on an optimization specific to each radar network, or empirically set based on assumed prior knowledge.

4.4. RESULTS

The test case used to evaluate this retrieval will be the generated test case shown previously in this chapter. This is a four frequency, non co-located case with radars at S,C,X,

and Ku-Band. Figure 4.9 shows the median drop diameter retrieval results from applying the network retrieval to this dataset. The resolution grid is chosen to be the study area listed above, where the resolution grid points have a uniform spatial resolution of 500 m. The first panel shows network retrieval, while the second panel shows the ground truth. The following four panels show the four component single radar retrievals. Examining the component retrievals show a large number of visually discernable regions of error. Most of the retrievals over-estimate the region of large median drop diameter in the southeast corner of the image. Also, the component retrievals appear more granular, having a higher level of spatial variation. The network retrieval successfully reduces this over estimation of D_0 in the southeast corner, while also having a spatial variability that more closely matches the ground truth. The K_u-Band radar shows an area of significantly higher errors in the southeast corner where the signal goes extinct. This area of high error does not propagate into the network retrieval as the recreation error filter removes it. The region in the lower left that is displayed as missing data is due to no radar field of view containing data in that region.
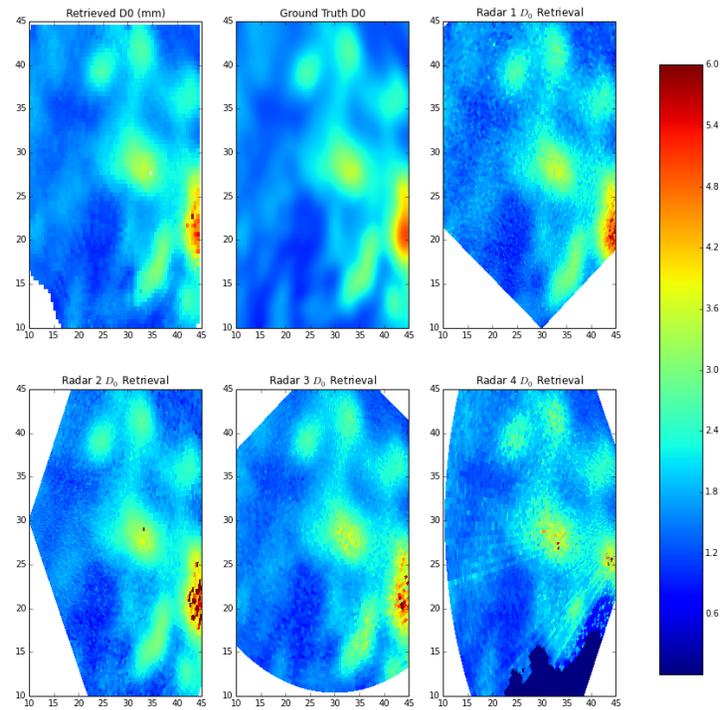
FIGURE 4.9. Median Drop Diameter retrieval results. The first pane is the retrieved field. The second pane is the ground truth field. The following four panels are the single radar retrievals of median drop diameter for each of the four individual radars.
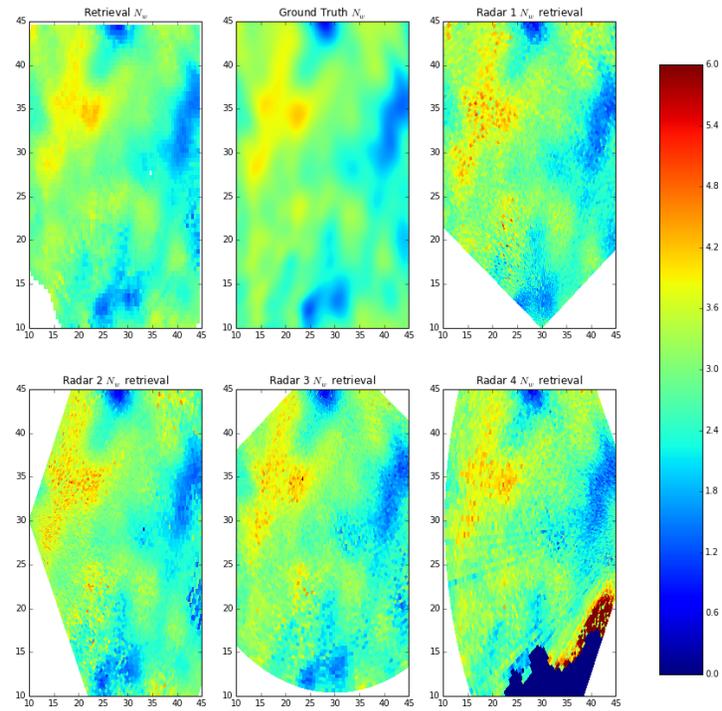
FIGURE 4.10. Normalized Intercept Parameter retrieval results. The first pane is the retrieved field. The second pane is the ground truth field. The following four panels are the single radar retrievals of the normalized intercept parameter for each of the four individual radars.

Figure 4.10 shows similar results for the normalized intercept parameter N_w. Again, the spatial variability of the network retrieval appears more in line with the ground truth, while rejecting an area of high error caused by signal extinction in the K_u-Band retrieval. The spatial variability of the component retrievals appears much higher, especially in the upper left region of the images. Each of the component retrieval exhibits a very "spotty" return, with high bin to bin variability. The network retrieval successfully removes the spurious variability, while correctly retrieving the natural variability of the underlying field.

Figure 4.11 shows a histogram of the errors in the median drop diameter retrieval for each of the component radars, normalized to provide a density estimation. The errors for the component radars S, C, and X are centered around zero showing little to no systematic bias in the retrieval. The K_u band radar has a bimodal distribution due to the large errors encountered around the region of signal extinction, causing a Gaussian distribution to be a poor fit. Overlaid on each histogram is a gaussian fit to the distribution. The parameters for these fits are shown in Table 4.1. We can see that the lowest three frequency radars all have a standard deviation of D_0 error of between 1.44 and 1.63.

Similarly, Figure 4.12 shows histograms of the errors in estimating the normalized intercept parameter N_w for each of the component radars. Again, the three lower frequency radars all display no appreciable bias in the results, with the X-Band radar showing the tightest coupling of errors around zero. The coefficients for the Gaussian fits overlaid on the histograms are shown in row two of Table 4.1. A similar behavior is seen in K_u for this parameter as that seen for D_0.

The distribution of errors for the network retrieval is shown in Figure 4.13. While each of the component retrievals can only fairly be evaluated over their individual fields of view, the network retrieval errors may be taken over the entire domain where data is present within
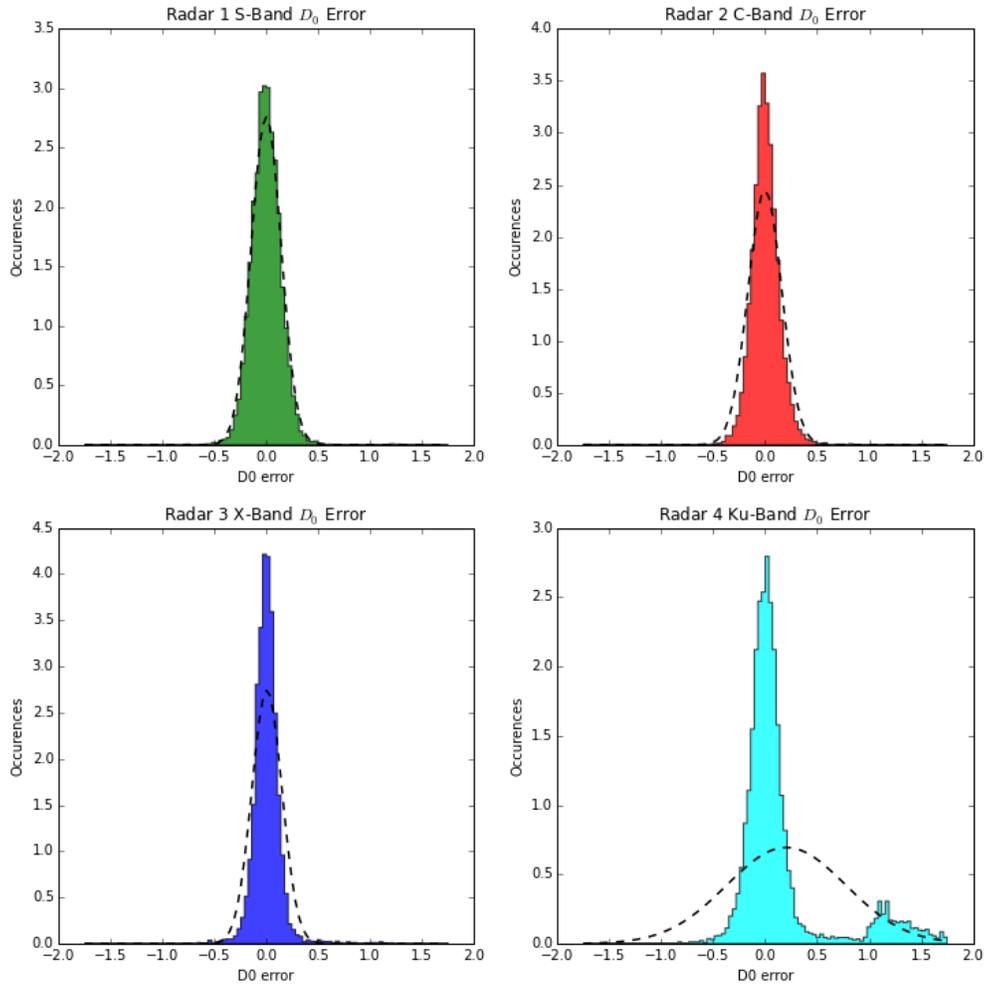
FIGURE 4.11. Normalized histograms of single radar retrieval errors for the median drop diameter parameter. Overlaid are gaussian approximations to the distribution.

the area of study. A Gaussian fit to each of the distributions is overlaid on the plots, with coefficients given in Table 4.2. While there is a very tiny increase in the mean value of the error, the standard deviation of the error decreases significantly compared to any of the component retrievals.
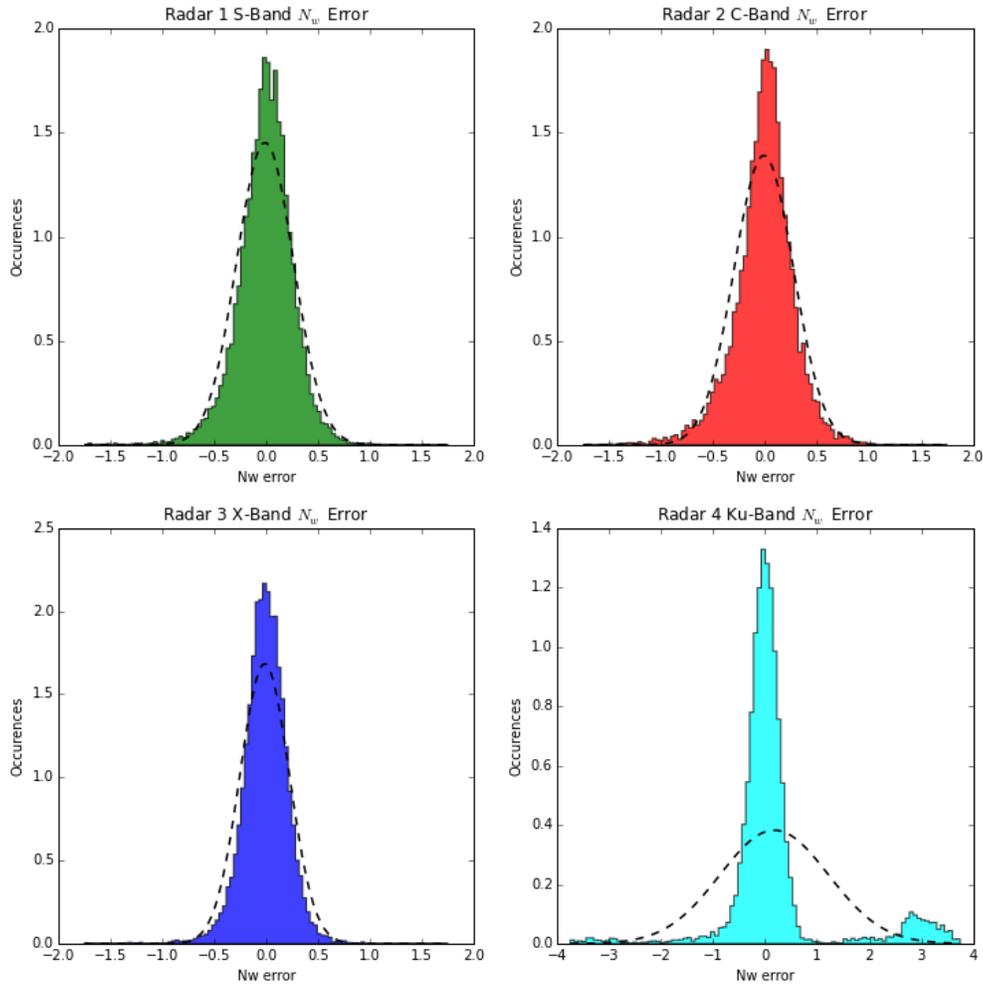
FIGURE 4.12. Normalized histograms of single radar retrieval errors for the normalized intercept parameter. Overlaid are gaussian approximations to the distribution.

Additionally, by incorporating multiple different radars, the field of view of the network retrieval is significantly increased over the individual component radars.

While the distribution of the errors in the retrieval gives a good measure of the performance of the retrieval within the dataset, the measure is not scale independent. This precludes future comparison with other retrievals. Table 4.3 contains the error metric scores

TABLE 4.1. Coefficients of a Gaussian fit to the distribution of errors in single radar retrievals. μ is the mean of the distribution, while σ is the standard deviation. It is clear there is a good fit for three of the component radars, whereas extinction in the K_u-Band radar causes a bimodal distribution

	$S-Band$		$C-Band$		$X-Band$		$Ku-Band$	
	μ	σ	μ	σ	μ	σ	μ	σ
D_0	-0.0015	0.144	0.0029	0.163	0.0057	0.154	0.203	0.575
N_w	-0.012	0.274	-0.013	0.287	-0.0157	0.237	1.044	1.044
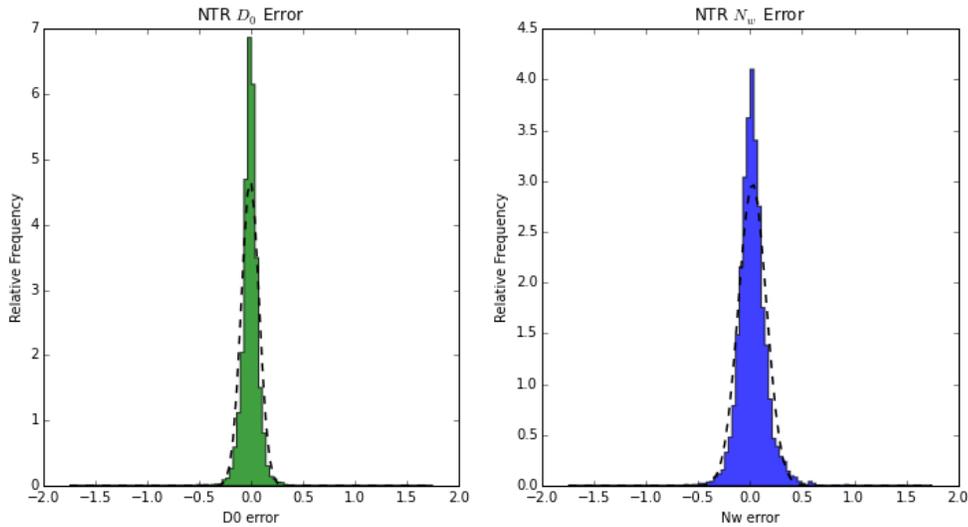
FIGURE 4.13. Normalized histograms of network retrieval errors. Overlaid are gaussian approximations to the distribution.

TABLE 4.2. Coefficients of a Gaussian fit to the distribution of errors in the network retrieval. The NTR exhibits a slight increase in the mean error, with a significant decrease in the standard deviation of the errors.

Parameter	μ	σ
D_0	-0.011	0.085
N_w	0.022	0.134

for four of the commonly used metrics for retrieval success. As expected from the distribution of errors, the three low frequency radars have negligable bias. Even the K_u band radar has a low bias despite the area of signal extinction causing extraneous errors. The following discussion will refer primarily to the three lowest frequency radars unless explicitly mentioned. The normalized standard error in the single radar retrievals is also low, in the range of 5-8 %. The mean absolute error indicates most of the expected errors in D_0 estimation should be less than 0.1mm, and less than 0.2 dB for the normalized intercept parameter. Finally the MAPE metric is provided for reference.

TABLE 4.3. Single Radar Retrieval Error Metrics for Network Retrieval Dataset

	S − Band		C − Band		X − Band		Ku − Band	
	D_0	N_w	D_0	N_w	D_0	N_w	D_0	N_w
NSE	0.059	0.077	0.063	0.071	0.048	0.065	0.167	0.213
NB	-8.3e-04	-4.67e-03	-1.72e-03	-4.5e-03	3.06e-03	-6.2e-03	0.11	0.067
MAE	0.109	0.198	0.105	0.208	0.09	0.163	0.308	0.550
MAPE	0.063	0.087	0.067	0.073	0.469	0.081	0.174	0.211

Table 4.4 lists the same error metric scores for the combined network retrieval. The error in D_0 is halved over any of the other radars, with an even better improvement for the normalized intercept parameter. Biases did slightly increase, but are both under 1% which is acceptable. Finally, the mean absolute error indicates that errors in D_0 estimation should be primarily less than 0.055 mm, while errors in N_w being predominantly below 0.1 dB. Overall, each of metrics improved by about a factor of two when using the network retrieval as compared with the single radar retrievals. Finally, the MAPE score is shown for reference.

TABLE 4.4. Network Retrieval Overall Error Measures.

Parameter	NSE	NB	MAE	MAPE
D_0	0.031	-0.0057	0.055	0.032
N_w	0.034	-0.007	0.096	0.036

The grid resolution for this test case was set at 500m in both directions to match a common grid size used for network radar mosaics. The retrieval however is not limited to any particular resolution. Figure 4.14 shows the retrieval at 500 m, 250 m, and 100 m resolutions. There is an increase in quality seen moving to 250 meters resolution, and a visual improvement in quality moving to 100 m. In this way the retrieval can match an analysis grid of interest. However, accuracy will not particularly change with higher spatial resolutions beyond a point. The original radar retrievals are conducted at a range resolution of 400 m and an angular resolution of 1°. This means as the resolution increases, the choice of corner points starts to become the same between different resolution grid points. While this provides a more visually pleasing image, it does not increase the accuracy of the retrieval, as at this point, data is merely being replicated.
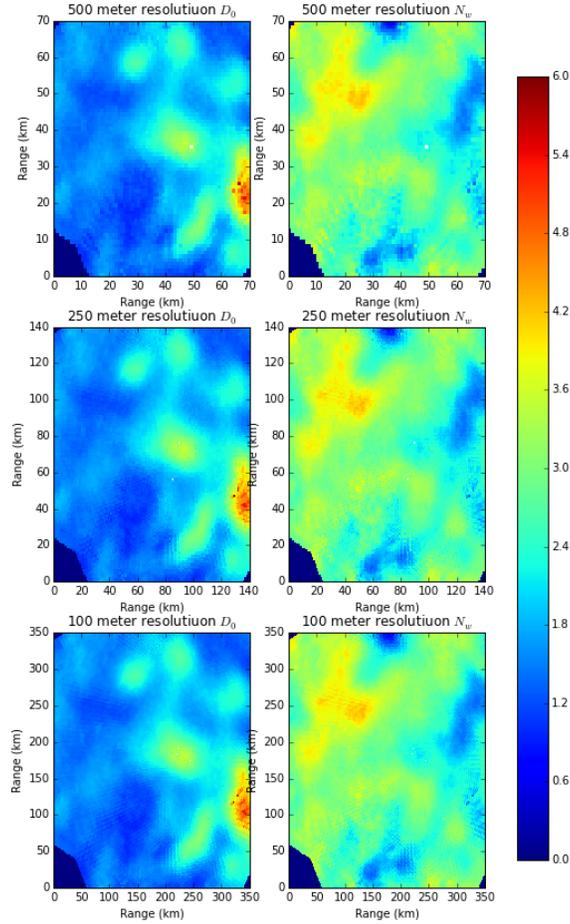
FIGURE 4.14. The retrieval grid can have arbitrary resolution and grid spacing. In this example, three different resolution grids are chosen. The top grid is 500m resolution, while the middle is 250m resolution, and the bottom retrieval is at a resolution of 100m.

4.5. SUMMARY

This chapter presented a framework for microphysical retrievals using a multiple frequency non-colocated radar networks. The formulation used radar measured parameters to estimate a global microphysical parameterization by combining radar measurements at a microphysical level. The accuracy of the method was demonstrated on a simulated network

dataset using four radars, each at a different frequency. It was shown that the algorithm is capable of handling many of the major network error sources such as beam blockage, extinction of the signal at higher frequencies, disagreements between individual radars, and differences in the intrinsic field due to scattering regime and path integrated effects. In this case, the network approach decreased relevant error metrics by a factor of two on average, with a larger coverage area than any individual radar in the network. It was also shown that the retrieval is able to avoid propagation of errors from mis-estimated intrinsic fields.

This approach extends to larger networks with an arbitrary number of radars at different frequencies. Additionally, by utilizing a covering patch approach to fusion instead of a traditional cloud of points approach, a small level of tolerance to advection and temporal errors is introduced. While this work used the single radar retrieval in Chapter 3 as a base for the intrinsic field estimation, the network formulation here is not bound to this constraint. In place of the recreation error provided by the single radar retrieval for use in the weighting function of the cost equation, a constraint based on internal self consistency of the measurements could likely be used.

In addition, a new methodology for creating synthetic test sets for radar networks from a single radar was developed. This method grids data sampled onto a spherical grid, onto another coincident spherical grid. In the process, data can be adjusted to simulate arbitrary frequency radars.

CHAPTER 5

Contributions of Non-Radar Ground

Instrumentation

"No one is actually dead until the ripples they cause in the world die

away.."

— Terry Pratchet, Reaper Man

Although the focus of this work is radar retrievals, other instruments that are deployed

to measure the atmosphere can contribute significantly to the estimation of microphysics,

and aid the radar retrieval in multiple ways. One of the more commonly used companion

instruments is networks of dual tipping bucket rain gauges. These gauges can be used both

to validate radar derived estimates of precipitation from the national NEXRAD network, as

well as be incorporated into the estimation of rain rates directly[53]. While in general the

estimated rain rate at a rain gauge will be more accurate than a radar estimate, it also only

measures a small area and can be subject to it's own sources of error. A radar may have a

higher average level of error at any given point, but will have a significantly larger coverage

area. Furthermore, networks of radars can be deployed to estimate precipitation over very

large coverage areas with high temporal resolution [54].

These rain gauges, although useful, typically only measure instantaneous and cumulative

rain rates, not microphysical parameterizations. Another instrument, called a disdrometer,

can be deployed to measure the drop size distribution of rainfall. The three most popular

types of these are the Joss-Waldvogel (JWD), Autonomous Parsivel Unit (APU), and 2-D

Video Disdrometer (2DVD) shown in Figure 5.2.

Each of these instruments provides a binned DSD over some sampling period. A common

sampling period for field projects is one or five minute intervals. Estimates of DSDs were

FIGURE
5.1. (a) Joss
Waldvogel
Disdrometer
(Courtesy
nasa.gov)

FIGURE
5.2. (b) Au-
tonomous
Parsivel Unit
(Courtesy
nasa.gov)

FIGURE
5.3. (c) 2D
Video Dis-
drometer
(Courtesy
nasa.gov)

FIGURE 5.4. The three most common disdrometer types used in operational and research settings. The Joss-Waldvogel is an acoustic-impact device. The Parsivel unit is based upon the use of a laser coupled to an optical measurement array. Finally, the Video Disdrometer uses orthogonal line scan cameras.

used as a verification source in Chapter 3. The rest of this chapter will show how ground instrumentation can contribute to the network retrieval by constraining parameters, and providing increased information about the microphysical fields.

5.1. DROP SIZE DISTRIBUTION CONSTRAINTS

One of the best uses for disdrometers in the network retrieval is as a source of additional constraints on the DSD parameterization. These constraints limit the search space of possible DSD parameterizations, as well as increasing the accuracy by providing increased information about the scattering parameters. This section will address three key constraints that disdrometers can provide.

The 2DVD disdrometer type operates by using orthogonal line scan cameras. This allows it to take images of drops as they fall through a collection area. These images can be processed to calculate drop axis ratios for drops of different sizes, as well as estimating the range of canting angles. These results can be used to choose between the different drop shape

relationships (DSR) described in the literature. These drop shape relationships can vary by location and time, and consensus in the community has not been reached on the best model. The Thurai and Bringi [19] DSR was derived through the use of a 2DVD. These instruments are commonly deployed during NASA GPM-GV field campaigns, as well as other research field campaigns.
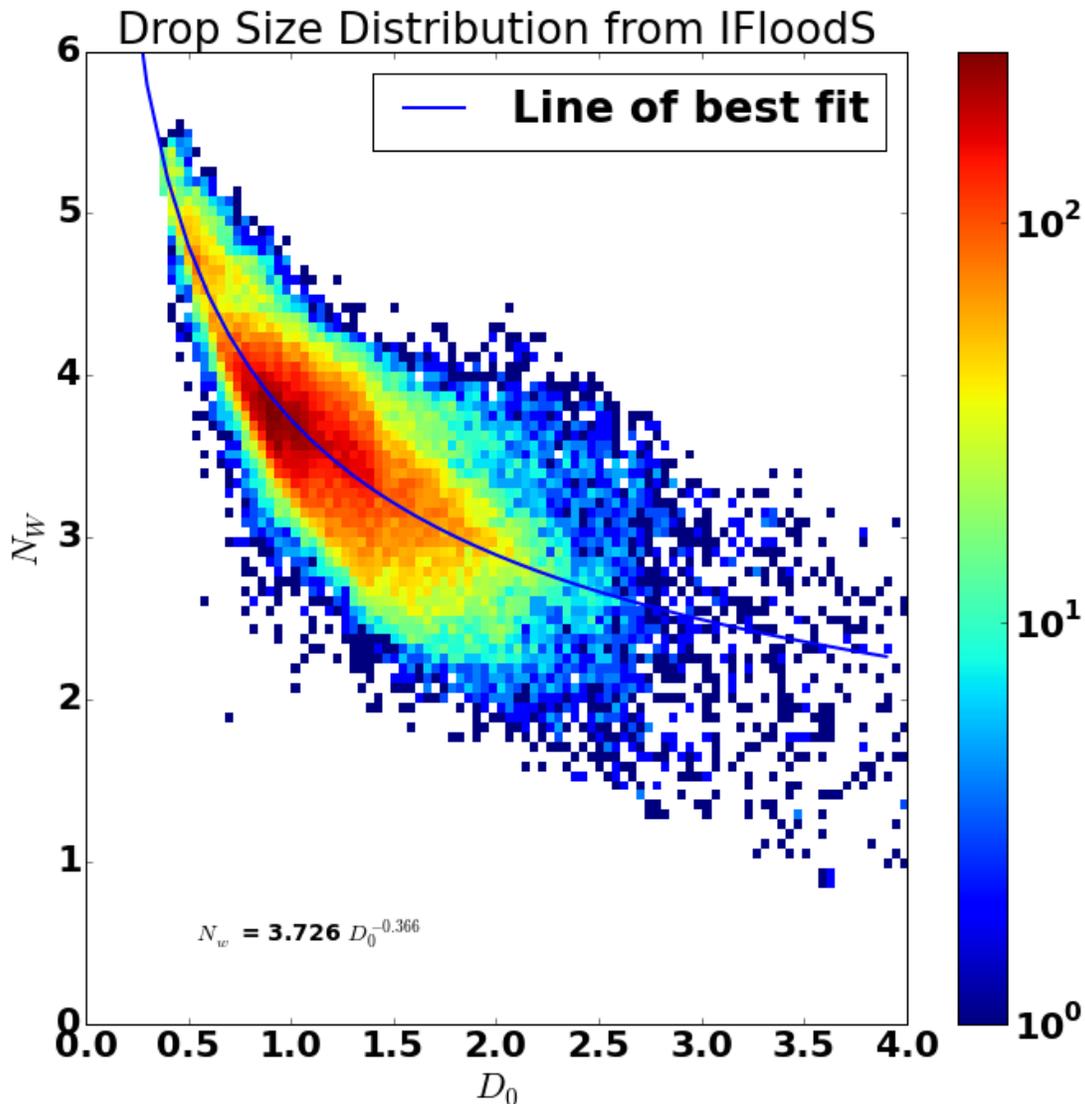
FIGURE 5.5. DSD distribution during IFloodS Field Campaign as derived from Autonomous Parsivel Units. A power law fit to the data is overlaid with the displayed parameterization.

A second constraint provided by disdrometers is a restriction of the physically realizable regions in the parameter space. It is known that the entire parameter space for (D_0, N_w) is not reachable by physically realistic drops. However, the constrained area can vary by geographic region, time of year, and storm type [55]. By deriving a more accurate bound on this physically realizable region, the cost function in both the single and the network retrieval can be augmented. Figure 5.5 shows the distribution of DSD parameters during the IFloodS field campaign measured by the network of APU's. A power law fit to the data is overlaid on the data based on the Levenberg-Marquardt algorithm described in Appendix 2. The fit to the data was incorporated into the single radar retrieval in the NPOL IFloodS example in Chapter 3. In general, the effect on bins that were retrieved accurately without the additional constraint is negligible. The addition of the constraint however did help to remove several anomalous results where the retrieved DSD was not physically realizable.

The third constraint of interest to the retrieval process is the selection of the shape parameter μ. In the implementation of the retrievals in this work, the value of μ defaults to a value of 3. Chapter 3 characterized the performance of the single radar retrieval for differing values of μ, showing a modest decrease in retrieval accuracy as the true value of μ diverged from the assumed value. The measured DSD at the disdrometer can be used to select a more accurate value of μ for each portion of the storm within range of the disdrometer.

The procedure for retrieving μ from the disdrometer measurements is covered in Bringi and Chandrasekar [1]. First, liquid water content is estimated. Next, the value of D_0 is given as the drop size diameter value for which half of the water content is above, and half of the water content is below. Next, the normalized intercept parameter N_w is derived based on it's relation with the liquid water content and median drop diameter. Finally, μ is chosen

so as to minimize the error between the modeled normalized gamma distribution with those parameters, and the actual measured histogram of drop sizes.

5.2. Spatial and Temporal Variability of Microphysics

One of the topics of interest in the atmospheric remote sensing community is how the distributions of microphysical parameters are structured, and how they vary. For instance, it is known that different climatic regions have different drop size distributions, and that the parameterization of these DSD's can vary on a seasonal scale. In addition, there is a high amount of variability between the DSDs encountered in convective storms compared to those sampled in stratiform type storms. Even the variability of the parameters differs between the two storm types. Indeed, classification schemes for these storms based on differences in DSD parameters have been proposed Penide et al. [56].

The spatial and temporal variability of the microphysics has a large effect on the single radar retrieval in Chapter 3, and the network retrieval developed in Chapter 4. The downsampling operation in the single radar retrieval relies on neighboring range samples to have a high level of correlation. If this does not hold, the decimation filter, which is a low pass filter, will cause excessive loss of spatially localized features. This degrades the overall estimation, and can cause large iterative errors to develop. The degree to which the downsampling operation affects the end result is directly related to the spatial variability of the microphysics.

The network retrieval is affected by the variability in two different ways. First, the fusion of measurements from the corners of the covering patch assumes the variability of the region covered by the patch is such that the measurements are all highly correlated with the retrieved microphysical parameters. This is implicit in the assumption that there is a

(D_0, N_w) field that minimizes the disagreement between the different radar measurements, and thus the cost function. Second, the time that each radar samples a given area is different than the time that the other radars sample that area. While this difference can be minimized with creative scanning strategies, it can never be fully eliminated. If the storm has a high enough level of temporal variability, advection will dominate the measurements and the radars will essentially be measuring completely different fields.

The following subsections lay out a framework for estimating this spatial variability using radars and disdrometers. We will then compare the variabilities, and discuss a relevant theory behind comparing ground and aerial observations.

5.2.1. TAYLOR HYPOTHESIS. Taylor was one of the first individuals to study the issue of comparing ground based measurements with radar observations of the atmosphere [34]. A key insight of the work was that if you assume that the atmosphere consists of vertically stacked stratified layers, and assume that operations on these layers are affine transformations, then differences in spatial location are equivalent to differences in time, and are related by the autocorrelation function of the storm structure. More specifically, given a storm with mean velocity η, the temporal correlation of the storm at time τ is equal to the spatial correlation of the storm at spatial lag $r = \eta\tau$. If one then factors in the affine transformation condition, changes in time can be directly compared to changes in space, and thus to changes in height.

It is known that storms change morphology as they precipitate out of the atmosphere due to a large number of physical processes, such as sedimentation, evaporation, and collision-coalescence. While the affine transformation condition is unrealistic in the most general case [57], it can be valid on small enough spatial and temporal scales.

5.2.2. VARIABILITY OF SCATTERED MOMENTS. Based on the Taylor hypothesis, there is a link between the temporal and spatial variability of the microphysics. There are two different methods to compare the variability of the measurements. The correlations can be compared in the radar measurement domain by simulating scattering on the disdrometer measurements, or alternatively on the microphysical fields after a radar retrieval has been performed. First we will address the results in the radar measurement domain.

Data from 30 days of the IFloodS campaign for both disdrometer and S-Band radar (NPOL) measurements were chosen as the dataset. For each disdrometer, radar measurements were simulated using the PyDisdrometer package at S-Band frequency. Then the average of the daily autocorrelations over the scattered parameters was calculated to give a temporal decorrelation time, defined as the lag where the autocorrelation drops below $1/e$. Next, the cross correlation between each pair of disdrometers was calculated for the three radar scattered parameters. A trend line was fit to this data, and the decorrelation distance was determined, again with a $1/e$ threshold.

Next the radar data above each disdrometer in the lowest elevation scan was processed in a similar way to calculate temporal and spatial decorrelation times/distances for the corresponding radar measurements. The results of this are shown in Table 5.1.

In general the results from the radar have a much larger correlation distance, and longer correlation time, with the exception of K_{dp}. This is likely due to a larger measurement area for the radar compared to the cross section of the storm sampled by the disdrometer. Unfortunately, the placement of the disdrometers was not optimal for this type of study, as large distances existed between many of the disdrometers. The maximum distance between disdrometers was over 100 km, although many pairs and triplets of co-located (within 500 meters) disdrometers are included in the dataset. A likely explanation for the higher temporal

TABLE 5.1. Scattered Moment Autocorrelation Times and Distances. Autocorrelations are calculated over 30 days of data from the NASA IFloodS campaign. The decorrelation value is taken to be $1/e$.

	Spatial Average (km)	Temporal Average (min)
APU Z_h	3.75	10.74
Radar Z_h	12.27	19.73
Overall Z_h	8.68	15.60
APU Z_{dr}	2.33	8.32
Radar Z_{dr}	5.60	10.37
Overall Z_{dr}	4.05	9.35
APU K_{dp}	3.30	13.32
Radar K_{dp}	8.00	4.34
Overall K_{dp}	5.53	10.06

correlation in K_{dp} for the disdrometer is the path specific nature of K_{dp} estimation. As the storm moves, the estimation algorithm can have significantly varying K_{dp}.

Figure 5.6 shows the distribution of spatial cross correlations between each pair of disdrometers, and between each pair of radar points. As expected, there is a sharp drop off with distance, with radar having a higher correlation with distance. Similarly, Figure 5.7 shows the results for differential reflectivity. A sharper gradient is observed with the correlations, but the the general trend as well as the increased correlations for the radar measurements remain. Finally, Figure 5.8 shows comparable results for specific differential phase, which exhibits similar trends to the previous two plots.

5.2.3. VARIABILITY OF MICROPHYSICAL PARAMETERS. Another way to evaluate the variability of the microphysics is to use disdrometer estimates of the microphysical parameters, and results from a radar microphysical retrieval. For this analysis, we will use the
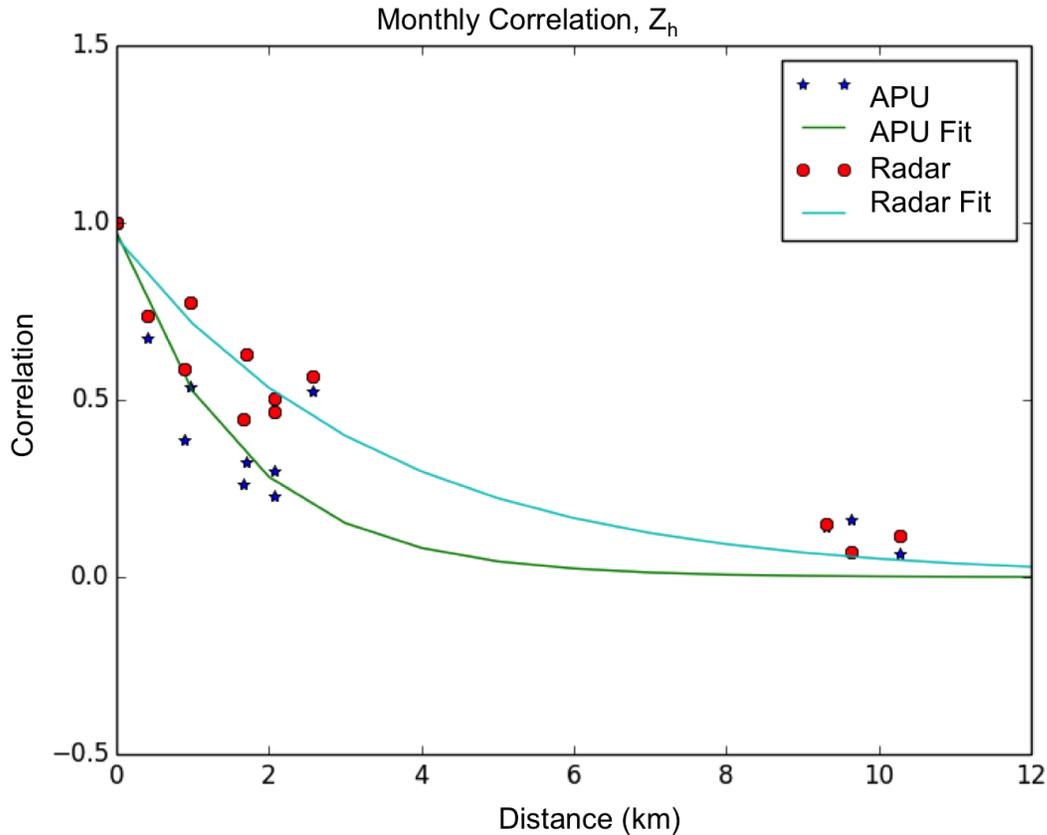
FIGURE 5.6. Spatial Correlation Distribution for Z_h based on a month of radar and disdrometer data. Radar measured moments are used for the radar label, while T-Matrix scattered radar moments are used from the disdrometer to compare disdrometer based correlations.

IFloodS dataset, although this time we will restrict the analysis to a 90 minute window during which a storm passed over one of the disdrometers on June 29th, 2013. The disdrometer measurements will be used to calculate a temporal correlation length, while the radar retrieval measurements will be used to calculate a spatial correlation distance.

The Drop Size Distribution from the APU01 disdrometer for the 90 minute period is shown in Figure 5.9. This shows the majority of the DSD is under 2mm drop size, although portions of the time have drops approaching 6 mm.

Next the microphysical parameters are estimated using the approach described in Appendix B. This gives a microphysical parameterization at 1 minute increments over the 90
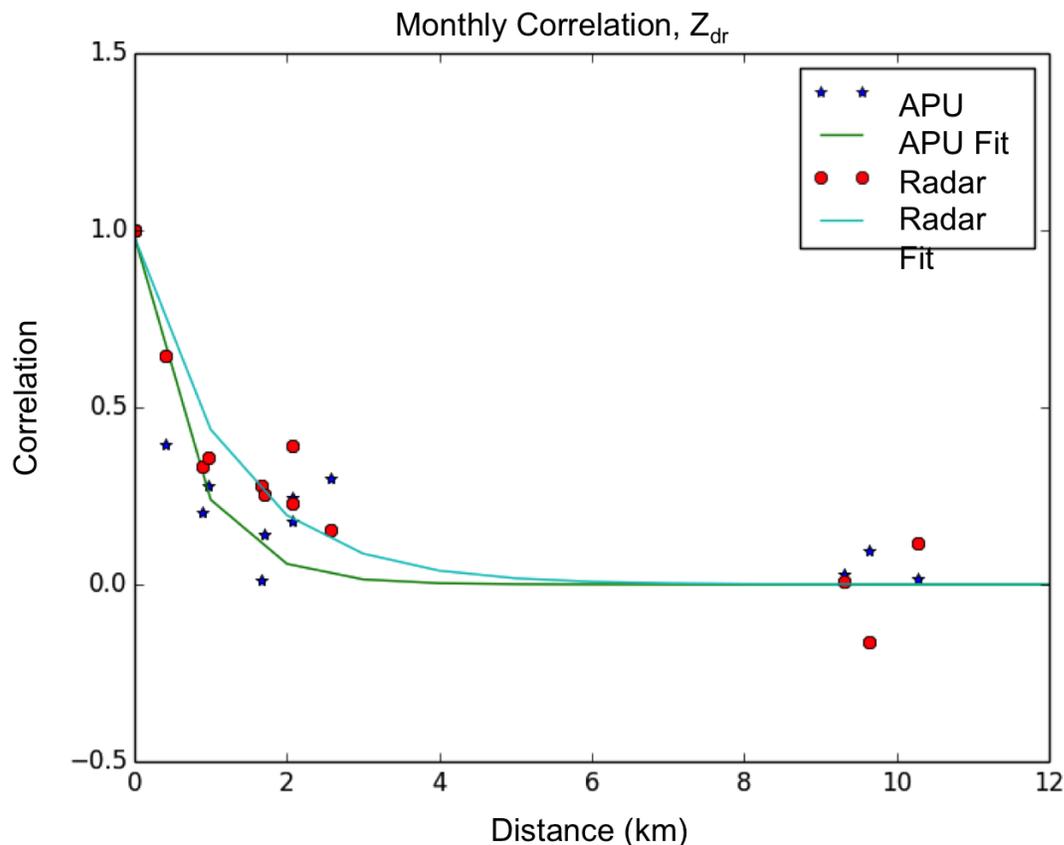
FIGURE 5.7. Spatial Correlation Distribution for Z_{dr} based on a month of radar and disdrometer data. Radar measured moments are used for the radar label, while T-Matrix scattered radar moments are used from the disdrometer to compare disdrometer based correlations.

minute period. The results of this operation are shown in Figure 5.10. The time period starts out with larger drops detected, before the median drop size decreases as the trailing stratiform portion of the storm passes over the disdrometer.

The autocorrelations for each of these figures are shown in Figures 5.11 and 5.12, with the decorrelation threshold displayed as the strong bar. The median drop diameter exhibits a correlation time of 15 minutes for this case, while the normalized intercept parameter decorrelates significantly more quickly with a time of only 3 minutes.

The spatial autocorrelation function of the radar retrieved microphysical parameterization using a combined S-Band/Ku-Band network retrieval is shown in Figures 5.13, and 5.14,
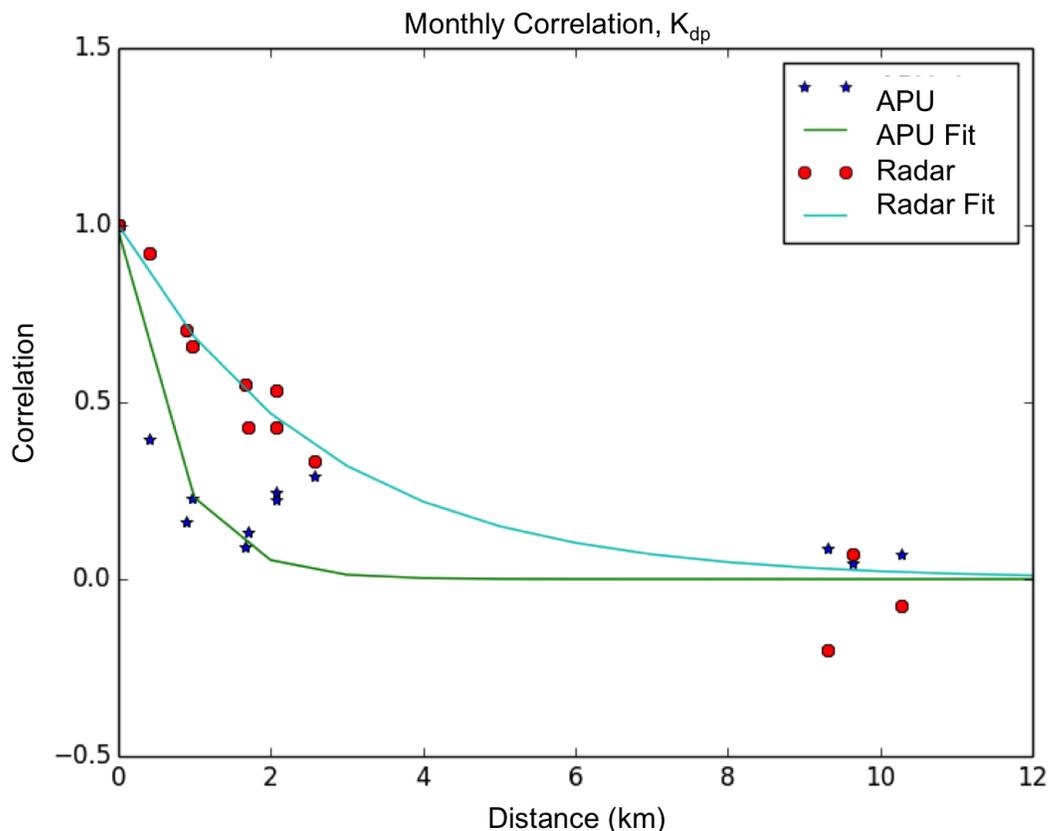
FIGURE 5.8. Spatial Correlation Distribution for K_{dp} based on a month of radar and disdrometer data. Radar measured moments are used for the radar label, while T-Matrix scattered radar moments are used from the disdrometer to compare disdrometer based correlations.

along with decorrelation threshold $1/e$. The median drop diameter becomes decorrelated at a lag of 10, which corresponds to 5 km. Similarly, the normalized intercept parameter becomes decorrelated at a lag of 3, which corresponds to a distance of 1.5 km. Similarly to the temporal result, the variability in N_w appears to be significantly larger than that of D_0.

Next, one can examine the differences between the two correlation scales. Starting with the relationship assumed by the Taylor hypothesis

$$r(D_0) = \eta\tau(D_0)$$ (65)
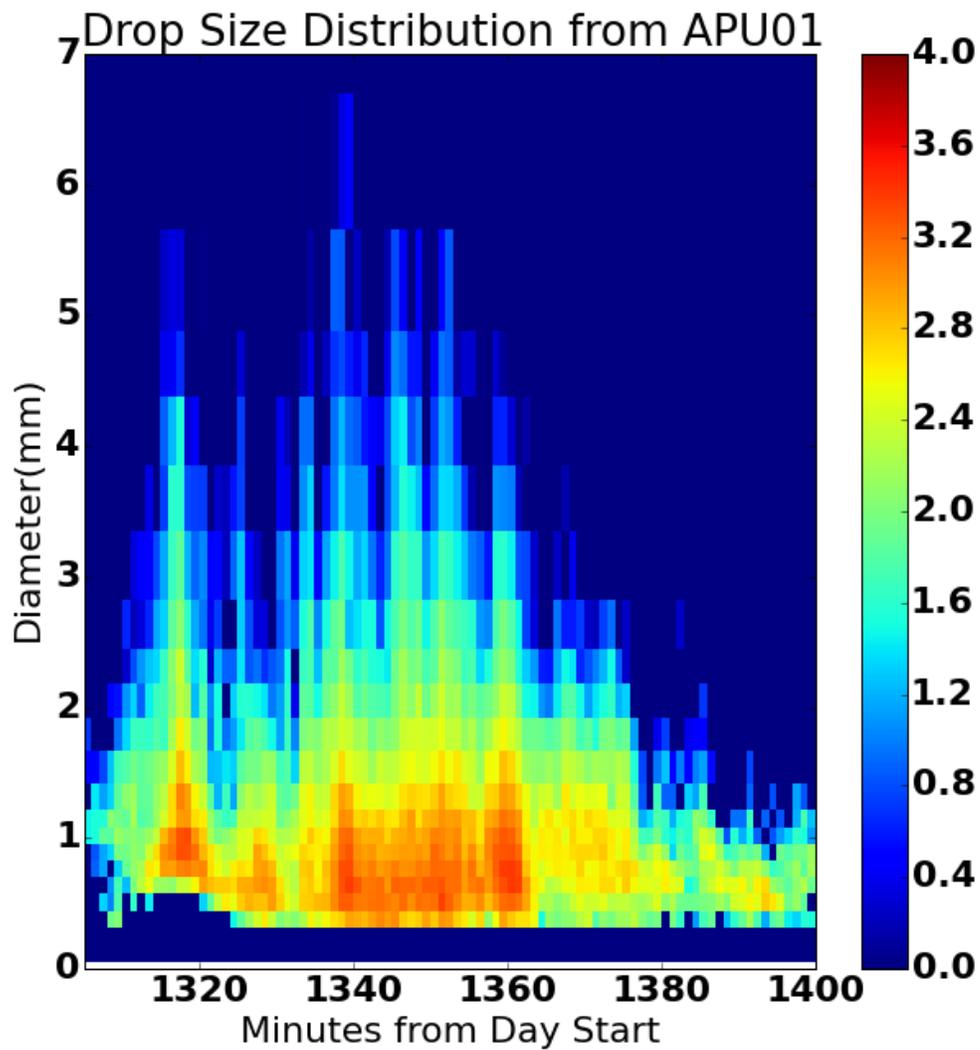
FIGURE 5.9. Drop Size Distribution from APU01 on June 29th around 21Z. Intensity of the diagram is drop count per bin on a $\log_{10}$ scale.

this can be rearranged to give

$$\frac{r(D_0)}{\tau(D_0)} = \eta$$

(66)
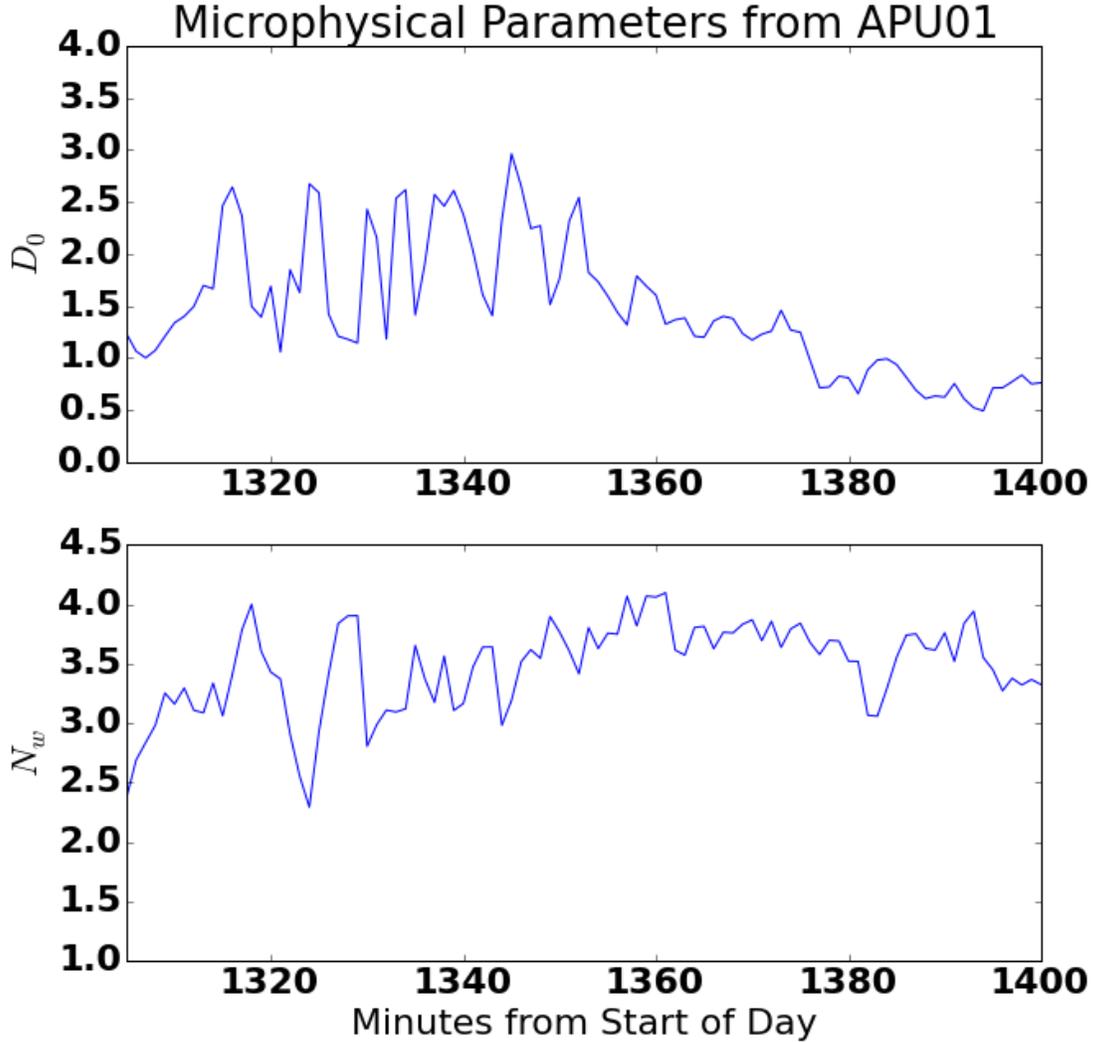
FIGURE 5.10. Microphysical Parameterization from APU01 on June 29th around 21Z. The top panel shows the median drop diameter, while the bottom parameter shows the normalized intercept parameter.

with a similar rearrangement for the normalized intercept parameter. In other words, the temporal and the spatial correlations are related to each other by the storm velocity η. While this relationship was only strictly formulated for the 2D Space-Time autocorrelation function, we will apply it to the 1D case. Solving for η using the above calculated decorrelation times and distances gives values of $\eta = (5.5, 8.3)$ m/s for the two parameter pair (D_0, N_w)
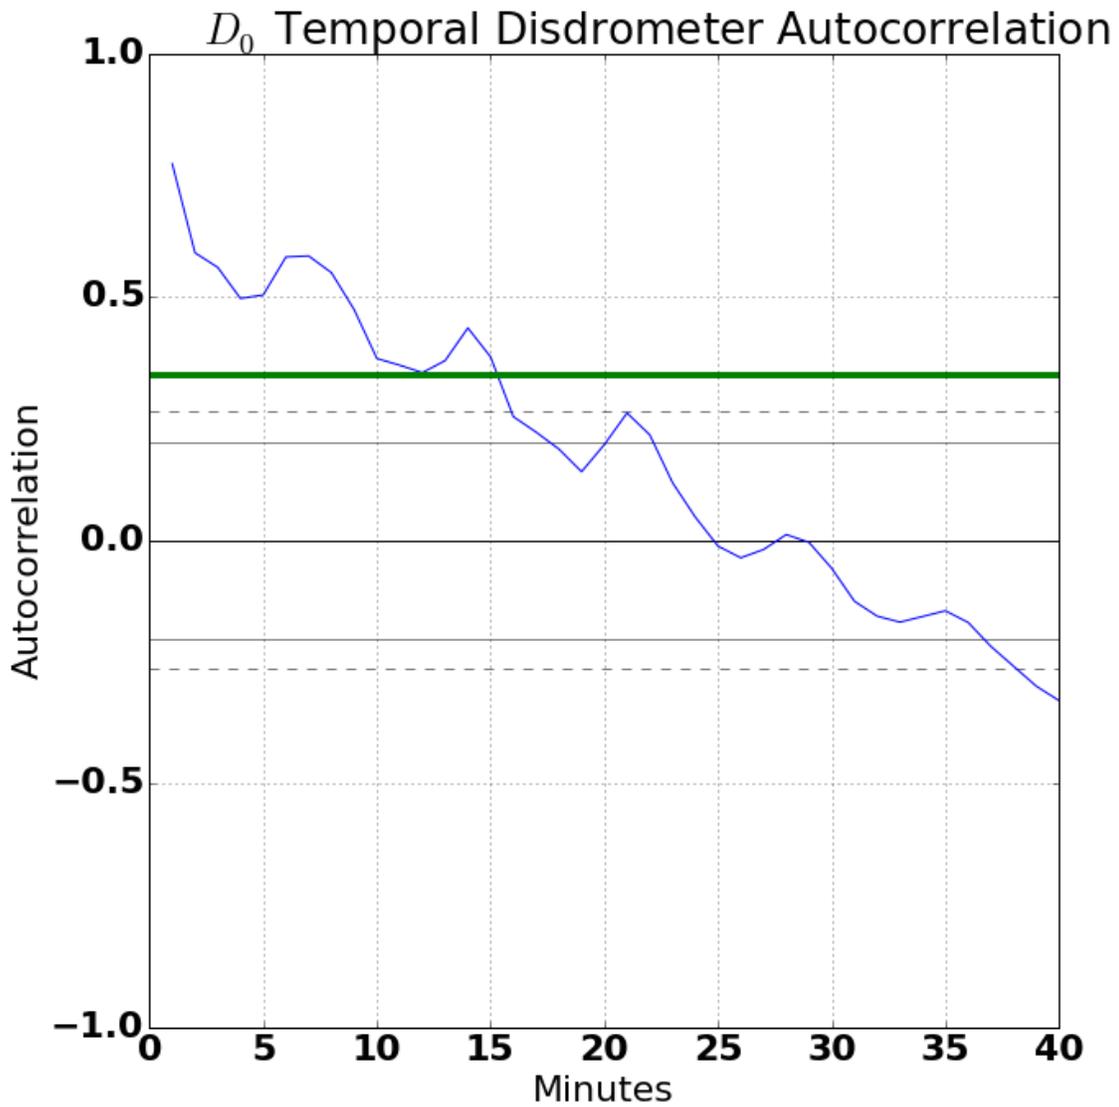
FIGURE 5.11. Temporal autocorrelation for the median drop diameter parameter. The strong bar corresponds to the decorrelation threshold 1/e. The decorrelation threshold is crossed at 15 minutes.

respectively. The average storm motion from radar measurements at the height of the melting layer, which is assumed to be the dominant generating physical driver, was about 6 m/s. This is surprisingly good agreement.
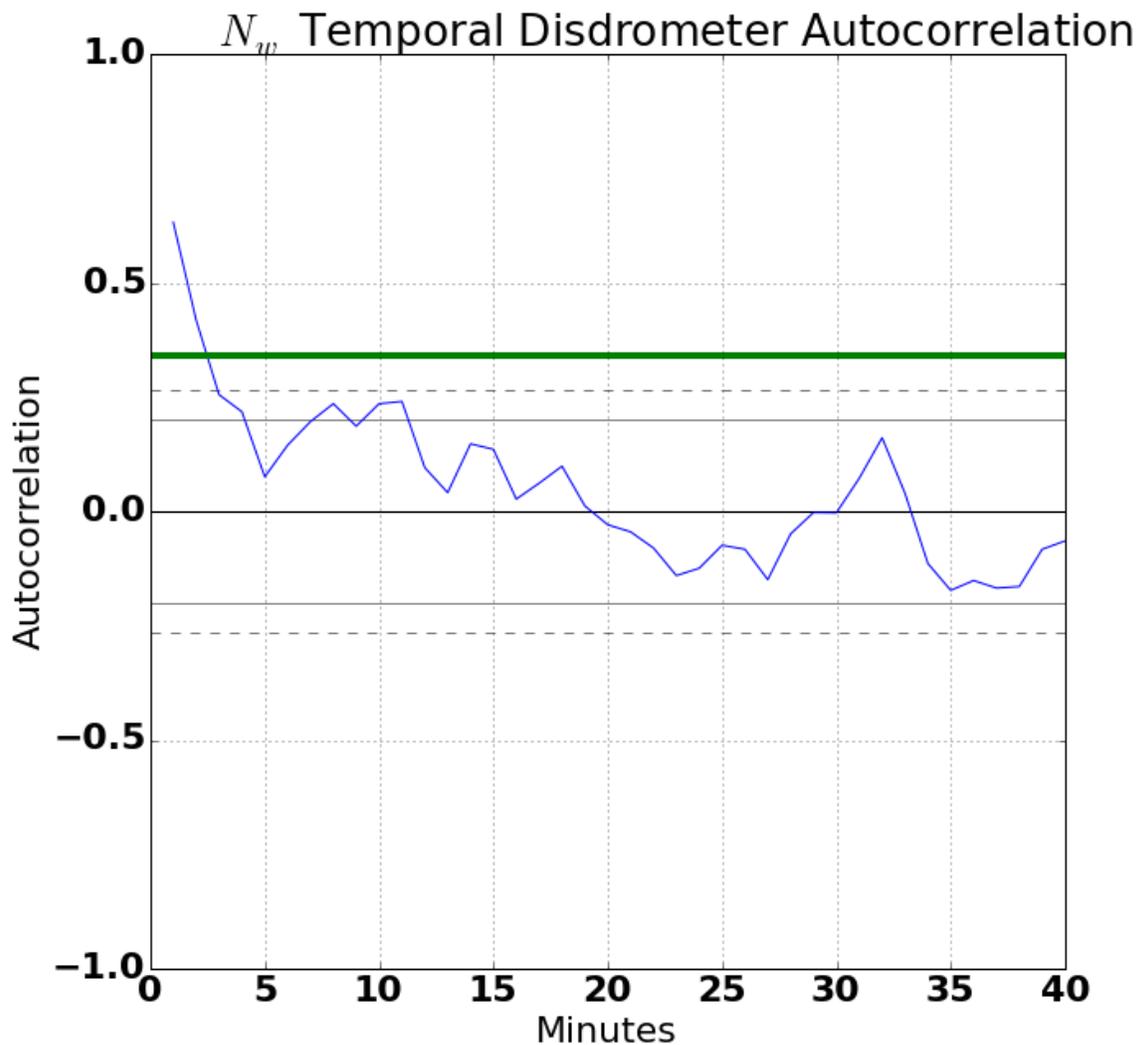
FIGURE 5.12. Temporal autocorrelation for the normalized intercept parameter. The strong bar corresponds to the decorrelation threshold $1/e$. The decorrelation threshold is crossed at 3 minutes.

5.2.4. MULTI-FREQUENCY SCATTERING RESULTS. The final contribution of ground instrumentation considered here is based upon the use of scattering results to find new relationships between multiple frequency observations of the same underlying microphysical fields. A lot of assumptions go into the empirical relations used by researchers in the field. A large number of these assumptions are derived from years of working with S-Band radar data.
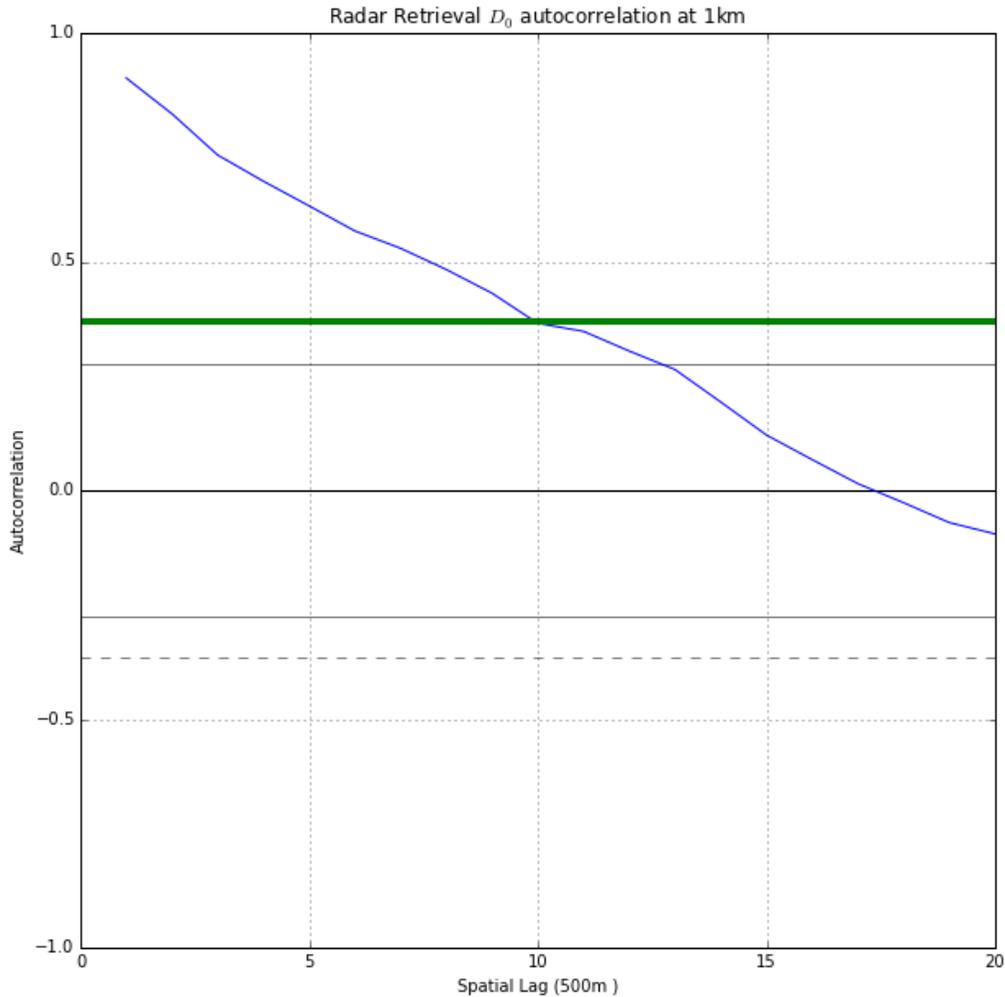
FIGURE 5.13. Spatial Autocorrelation for radar retrieved median drop diameter at a height of 1km. Each lag corresponds to 500m distance. The measurements become decorrelated at 5km.

Many of these relations have become canon, and unfortunately are used in situations where they are not necessarily true. This subsection will examine some of these assumptions as they relate to relations between radar moments as frequencies change.

To show the effect of changing frequency, the entire disdrometer dataset from IFloodS was passed through PyDisdrometer to generate equivalent radar measurements at K_u and
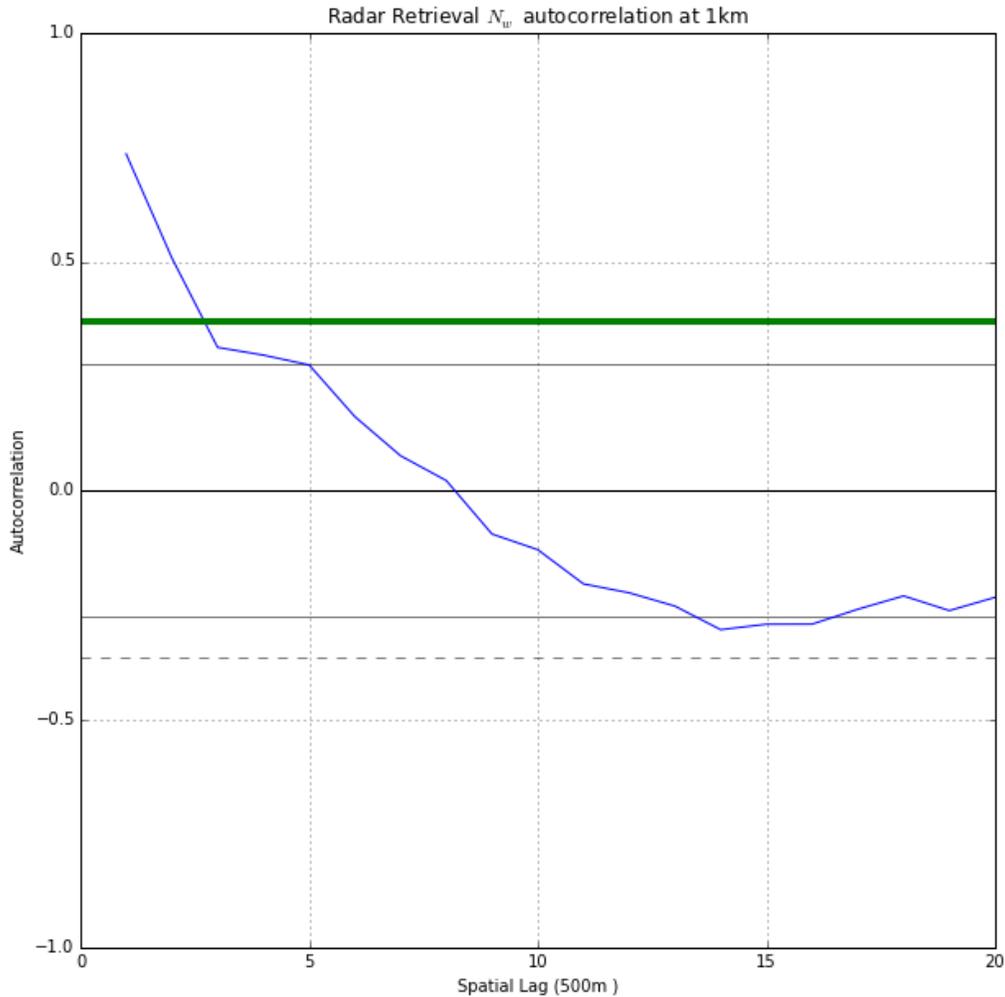
FIGURE 5.14. Spatial Autocorrelation for radar retrieved normalized inter-
cept parameter at a height of 1km. Each lag corresponds to 500m. The
measurements become decorrelated at 1.5km.

K_a frequencies. Scattering was done on the binned model, and then D_0 was subsequently

estimated for the purposes of labeling. Figure 5.15 and Figure 5.16 show the reflectivity of

this dataset as a function of median drop diameter for K_u and K_a. Overall, the trend is as

expected, however the K_a does exhibit a much higher variability in reflectivity at large drop

sizes. This is due to resonances in the Mie Region.
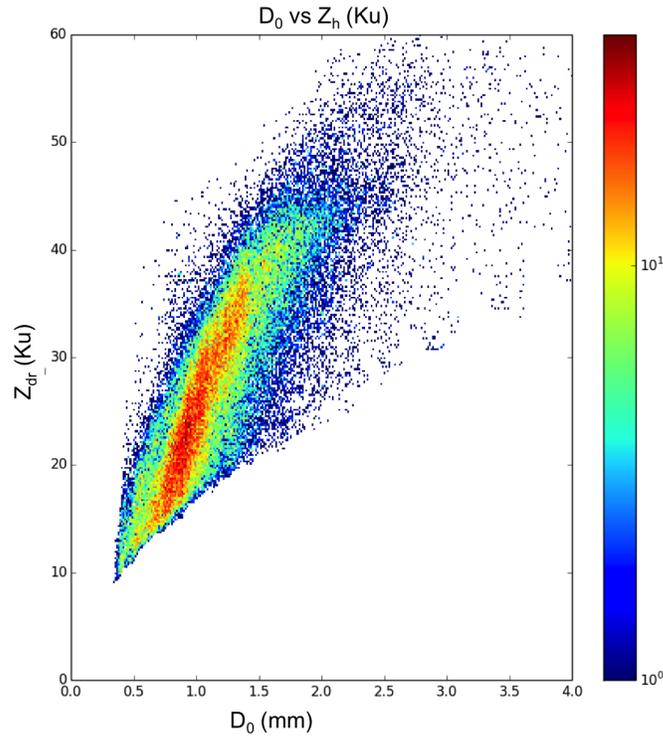
FIGURE 5.15. Reflectivity at K_u as a function of median drop diameter.

Figure 5.17 shows the dual frequency ratio of the two reflectivities. This is just the ratio of the two frequencies(or the difference in log domain). This shows that Ka has a slightly higher reflectivity for smaller drops, but as the drop size increases beyond 1.5 mm, the trend reverses and K_u has a much higher measured reflectivity. Many different retrieval algorithms, including those used on the GPM satellite, leverage the dual frequency ratio of co-located, beam-aligned radar systems.

There is more to be gleaned from this dataset however than somewhat well known results for DFR. Figure 5.18 shows the differential reflectivity at K_a band. It shows increasing Z_{dr} at small drop sizes, before plateauing at 0.5 dB. This is caused by resonances in the Mie region. As verification for this Figure 5.19 shows reflectivity and differential reflectivity from the D3R radar's K_a band during IFloodS. This shows that the differential reflectivities do
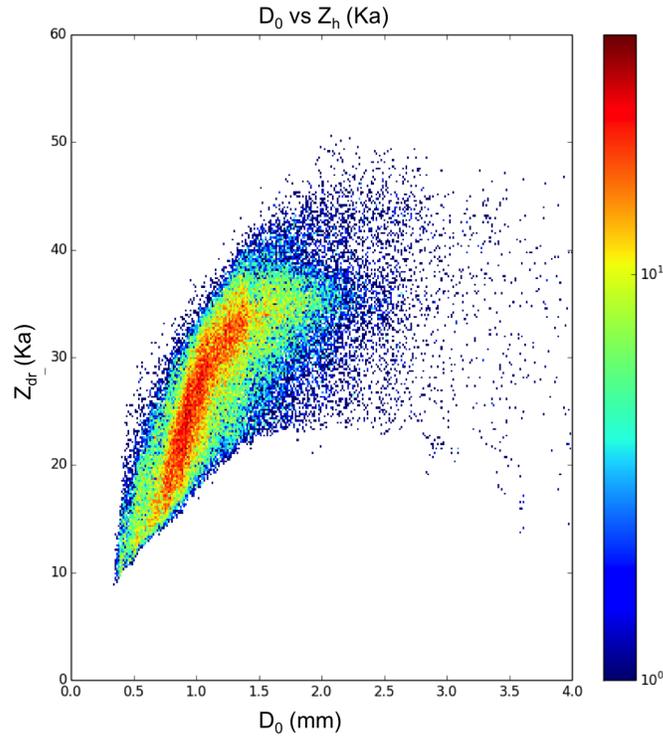
FIGURE 5.16. Reflectivity at K_a as a function of median drop diameter.

indeed plateau, although from the radar measurements at about 0.75 dB. The difference is likely due to a difference of drop shape relationship. This implies that power law type relationships for rain-rate estimation at K_a will run into estimation issues when based purely on Z_{dr} as some have proposed.

Next, K_{dp} shows some very interesting behavior that flies in the face of commonly held dogma. Figure 5.20 shows the specific differential phase results for the dataset. The interesting part of this graph is the negative K_{dp} values. These are generally not present at any of the lower frequencies for liquid precipitation, and indeed many algorithms do not even allow values of K_{dp} to be negative.

Finally, Figure 5.21 shows the ratio of specific differential phases between the two frequencies. It is generally assumed that K_{dp} scales linearly with frequency. This mostly holds
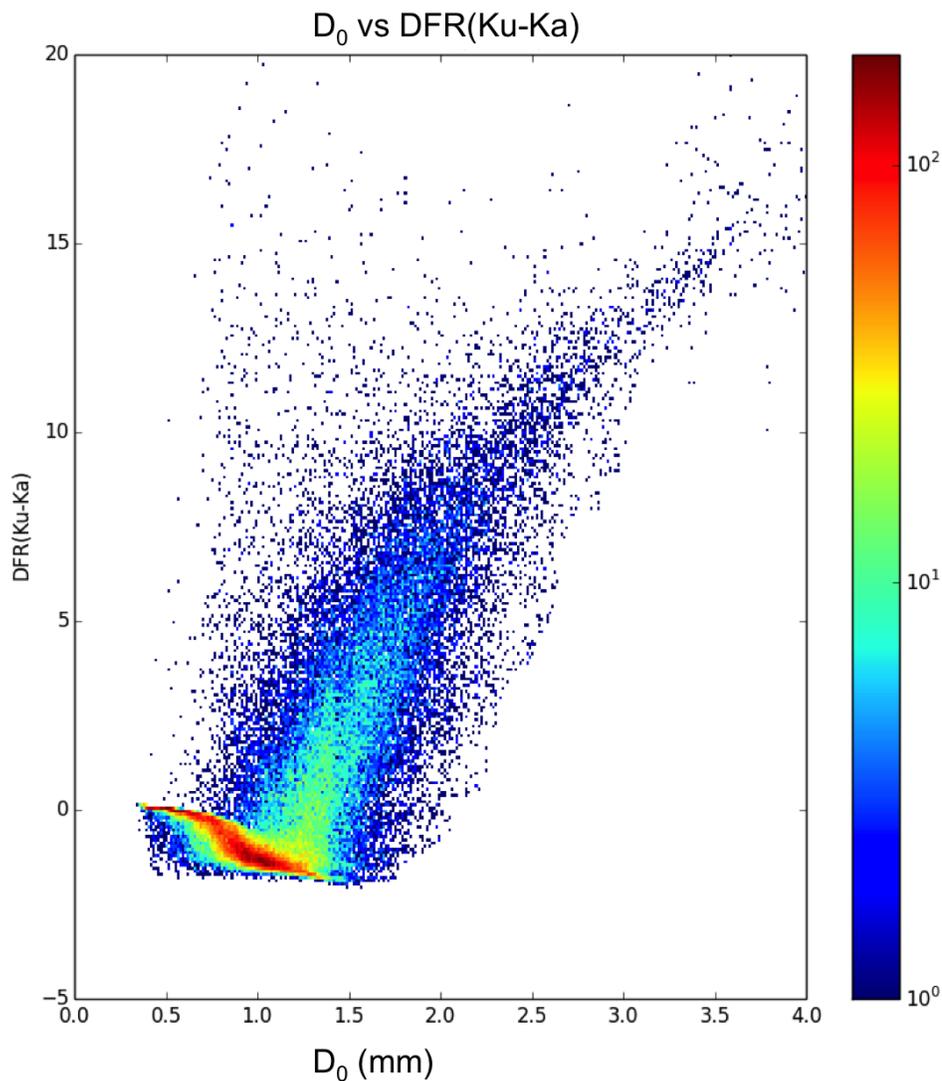
FIGURE 5.17. Dual Frequency Ratio(DFR)(Ku-Ka) for the IFloodS disdrometer dataset. This is repeated here for reader convenience.

true for lower frequencies. At K_a however this assumption fails. The ratio of the two frequencies used for scattering is roughly 2.7, and so the ratio of K_{dp} would be expected to be 2.7. Instead, the figure shows that the ratio is not constant, and changes with drop size. This result is not completely unknown to the literature, but is not as widely disseminated as it should be.
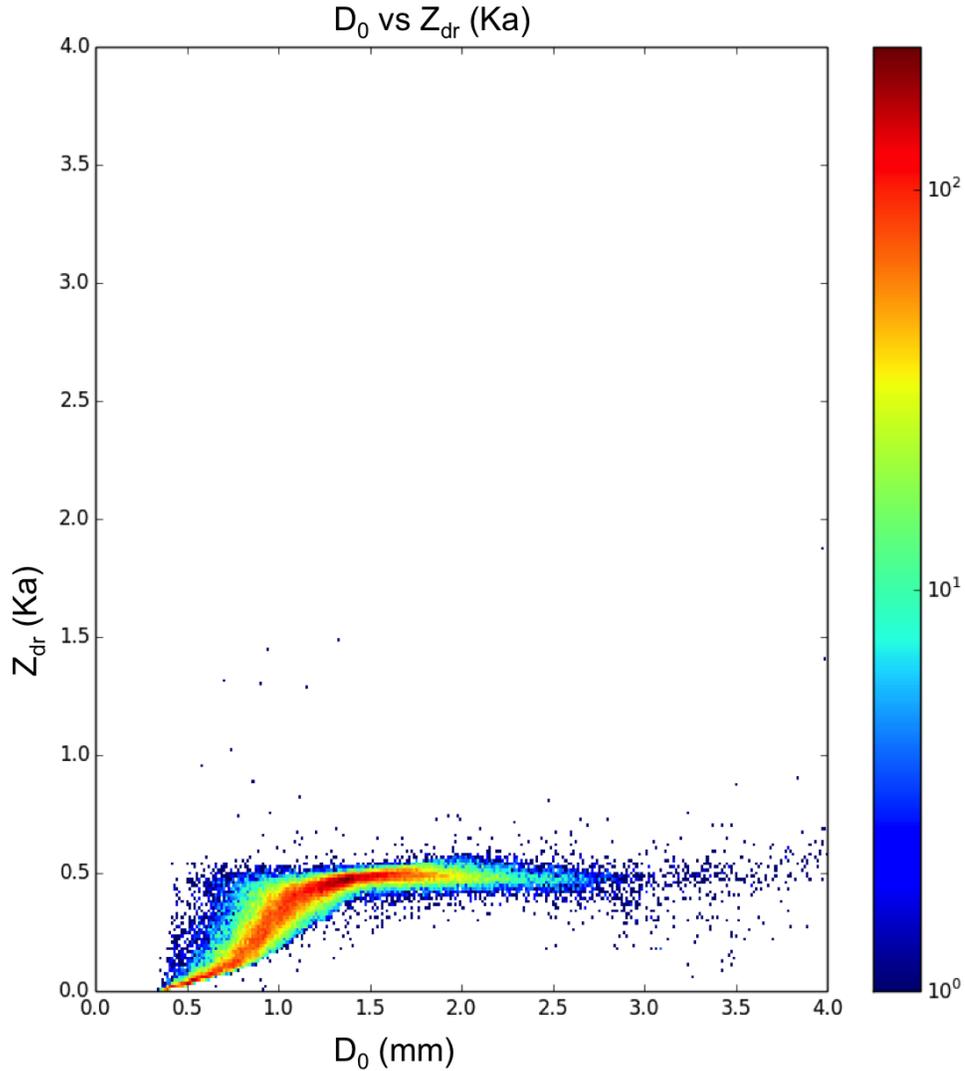
FIGURE 5.18. Differential reflectivity at K_a for the IFloodS dataset based on T-Matrix Scattering of the binned histogram. The differential reflectivity has the interesting feature of maxing out at a value of 0.5 dB.

5.3. SUMMARY

Although not as commonly deployed as radar networks for research and operational purposes, this discussion has shown that ground measurements of drop size distributions using disdrometers has a place not only as a verification tool for radar retrievals, but also as a source of information for the instrument fusion in both the single and networked radar cases.
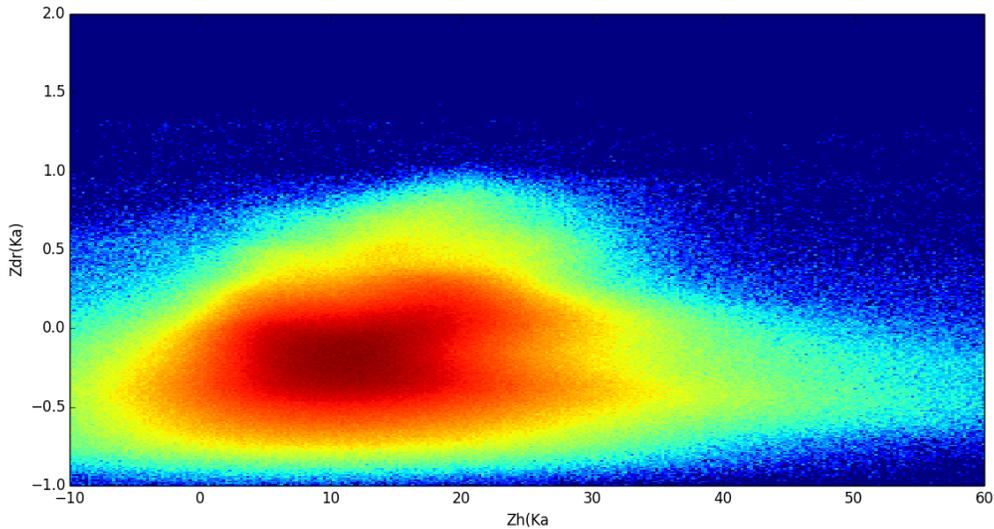
FIGURE 5.19. Reflectivity vs differential reflectivity from D3R's K_a band radar during IFloodS. Results are based on T-Matrix Scattering.

The single and network radar retrieval algorithm makes some *a-priori* assumptions about parameters such as the shape relationship and shape parameter μ. Ground instrumentation serves as an anchor allowing estimation of these parameters to be input into the retrieval process. Additionally, the region of physically realizable DSD's that is used as a term in the cost function of the single radar retrieval benefits from a more accurate geographically localized estimation of the distribution. This can be used either to augment the cost function of the retrieval, or as a verification source for the retrieval.

There remains an open question in the field as to how best link the measurements at the ground to measurements taken at higher altitudes by radars. Results were given showing a possible approach to tie together the two measurements by using the autocorrelation functions of each instrument along with the storm velocity to compare variabilities in both space and time. While this does not solve the spatial and temporal variability issues, it does provide a method by which the measurements can be linked and compared.
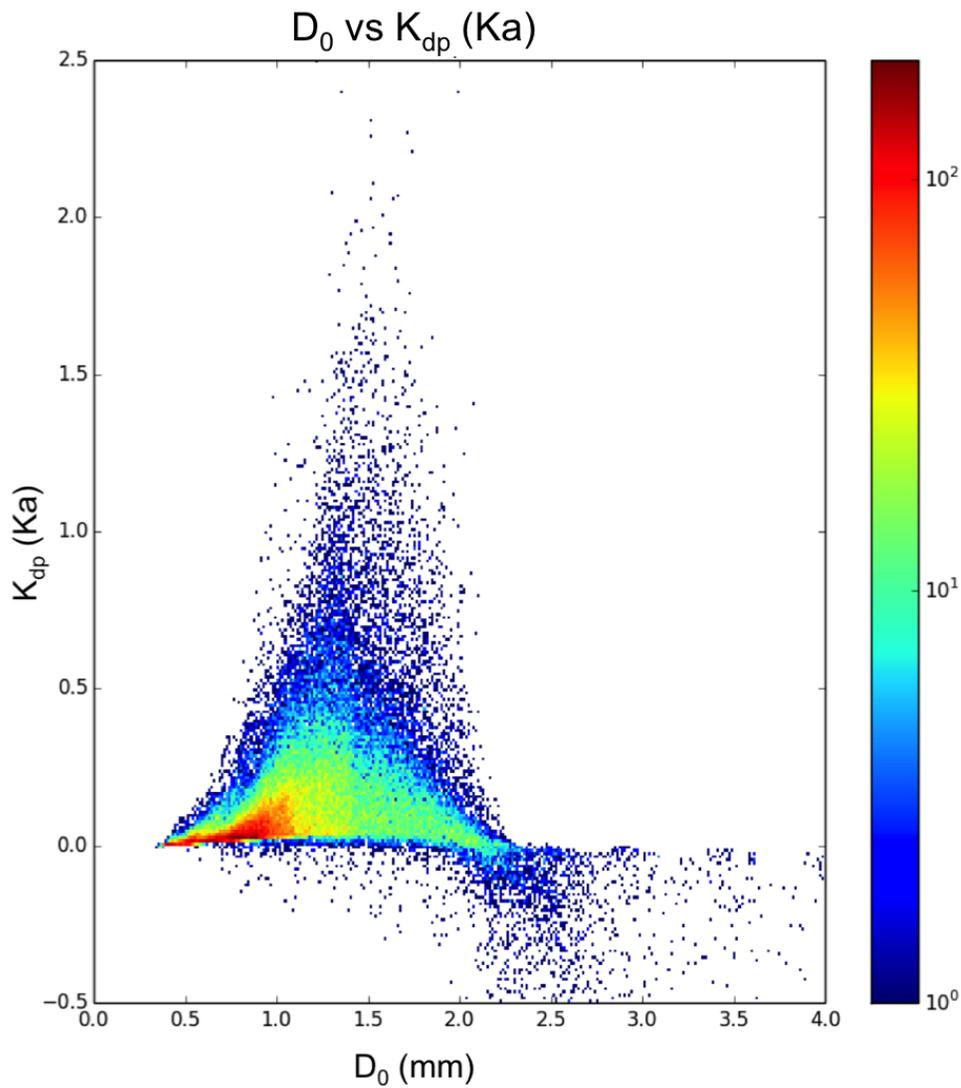
FIGURE 5.20. Specific Differential Phase at K_a for the IFloodS dataset. Of particular importance is the prevalence of negative K_{dp} values, which are not present at lower frequencies for liquid water.

Finally the scattering combined with disdrometer and radar measurements showed how new insights can be gained by cross comparing the different datasets. In particular, this section looked at how the behavior of the radar measurements changes with frequency as the frequency increases.
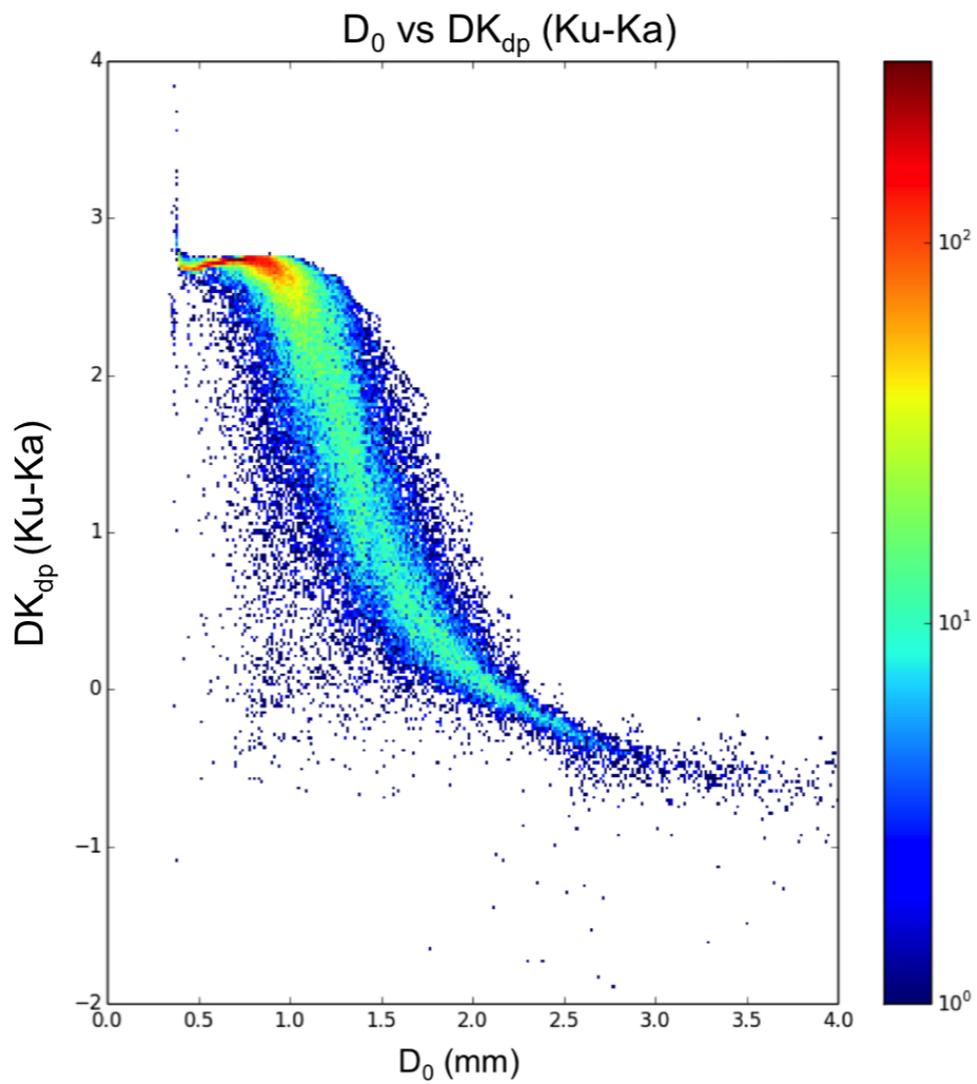
FIGURE 5.21. Ratio of K_{dp} between Ku and Ka for the IFloodS dataset. Data is generated based on T-Matrix Scattering.

SUMMARY AND CONCLUSIONS

"The fact that we live at the bottom of a deep gravity well, on the
surface of a gas covered planet going around a nuclear fireball 90
million miles away and think this to be normal is obviously some
indication of how skewed our perspective tends to be."

— Douglas Adams

Radar networks are making up an increasingly large portion of both the operational remote sensing capabilities of the world, as well the research networks. Networks such as NEXRAD have shown the benefits of combining multiple radars to get an increased field of view as well as well as multiple independent measurements of the atmosphere. Quasi-operational networks such as CASA demonstrate the benefit of incorporating different frequency radars into the equation. Utilizing lower cost, smaller radars for a network allows for increased spatial and temporal sampling. Finally, networks such as the temporary networks set up during the NASA GPM-GV field campaigns, and the more permanent Southern Great Plains facility operated by the Department of Energy's Atmospheric Radiation Measurement (ARM) program have shown that having multiple frequencies allows for discovering more information about precipitation. While a large amount of work has gone into utilizing multi-frequency information to calculate the drop size distribution parameters, the majority of it has focused on the beam-aligned case using dual frequency ratios. For instance, this is the approach used by the newly launched GPM-Core satellite.

This work presents two new algorithms for microphysical retrievals. In Chapter 3 a radar retrieval was formulated that uses information from a single radar to solve the forward

variational scattering problem. The algorithm provides an estimate of the underlying microphysical parameters D_0 and N_w that give rise to the radar measurements. The algorithm was shown to be accurate at a wide range of frequencies with low levels of error. An error characterization was given based on common error sources in radar measurements using both simulated data as well as data from operational research networks. The retrievals gave very good agreement both with recreating the radar measured fields, as well as when compared to estimates of the microphysical parameters measured at the ground by disdrometers.

Next, in Chapter 4 a framework for using non-colocated multiple frequency dual-polarized radar networks was developed. The framework works at the microphysical level, and extends the results from the single radar retrievals, while also admitting other intrinsic field estimation algorithms to be used. The benefits of combining measurements at the microphysical level was examined. To characterize the performance of the algorithm, a method of generating simulations of network radar data from a single radar measurement file was given. This method allows arbitrary locations and frequencies of radars to be generated.

Finally, Chapter 5 shows how ground instrumentation can be used to improve the retrieval performance. In particular, it was shown how the information provided by the different disdrometers can help to constrain the parameter space for the retrieval, while also eliminating some of the necessary *a-priori* assumptions that go into the retrieval. Then, a comparison of spatial and temporal variability between ground instrumentation, radar measurements, and retrieval results was examined.

6.1. Future Work

There are several desirable modifications that could be made to the retrieval process to improve performance in real world scenarios. The following are suggested improvement areas:

- Incorporate a hydrometeor identification algorithm to differentiate between liquid precipitation and frozen precipitation, or sources of clutter. This would reject some of the more spurious cases of error that can be caused by hydrometeors such as melting hail that are not easily detected by the simple quality control filters used in this work.

- Early in development of the retrieval algorithm, it was found that when the calibration of a radar is significantly off, the retrieval has issues with convergence as no DSD could adequately explain the radar measurements. This can be turned around and an optimization algorithm developed to recalibrate the radar based on minimization of global cost functions in the DSD retrieval optimization.

- Preliminary results exist on operational radar networks using the network retrieval framework. These networks however do not contain disdrometers and so it is hard to fully characterize the performance of the algorithm. A full characterization in a multi-frequency radar network with supporting ground instrumentation is required in the future.

- Currently the algorithm is computationally intensive. There exists a number of trade-offs that can reduce the computation time, at the expensive of memory and accuracy. A study detailing these trade-offs could reduce the computation time of the algorithm significantly.

- Currently the disdrometers are primarily used for verification of the algorithm and constraint of some of the *a-priori* assumptions. However a ground disdrometer provides a valid estimate of the microphysical field. It should be possible to directly incorporate this into the retrieval process, either at the minimization of the cost function, or through Bayesian fusion of multiple instruments

Bibliography

[1] V. Bringi and V. Chandrasekar, *Polarimetric Doppler Weather Radar*. Cambridge University Press, 1st ed., 2001.

[2] V. Chandrasekar, S. Lim, and E. Gorgucci, "Simulation of x-band rainfall observations from s-band radar data.," *Journal of Atmospheric & Oceanic Technology*, vol. 23, no. 9, 2006.

[3] H. Liu and V. Chandrasekar, "Classification of hydrometeors based on polarimetric radar measurements : development of fuzzy logic and neuro-fuzzy systems, and in situ verification," *Journal of Atmospheric and Oceanic Technology*, vol. 17, no. 2, 2000.

[4] H. R. Pruppacher and K. V. Beard, "A wind tunnel investigation of the internal circulation and shape of water drops falling at terminal velocity in air," *Quarterly Journal of the Royal Meteorological Society*, vol. 96, no. 408, pp. 247–256, 1970.

[5] D. N. Moisseev and V. Chandrasekar, "Nonparametric estimation of raindrop size distributions from dual-polarization radar spectral observations," *Journal of Atmospheric and Oceanic Technology*, vol. 24, pp. 1008–1018, 2014/01/17 2007.

[6] M. Le and V. Chandrasekar, "Raindrop size distribution retrieval from dual-frequency and dual-polarization radar," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 50, pp. 1748–1758, May 2012.

[7] M. N. Anagnostou, E. N. Anagnostou, J. Vivekanandan, and F. L. Ogden, "Comparison of two raindrop size distribution retrieval algorithms for x-band dual polarization observations," *Journal of Hydrometeorology*, vol. 9, pp. 589–600, 2014/01/17 2008.

[8] S.-G. Park, V. N. Bringi, V. Chandrasekar, M. Maki, and K. Iwanami, "Correction of radar reflectivity and differential reflectivity for rain attenuation at x band. part

i: Theoretical and empirical basis," *Journal of Atmospheric and Oceanic Technology*, vol. 22, pp. 1621–1632, 2014/01/17 2005.

[9] E. Gorgucci, V. Chandrasekar, and L. Baldini, "Microphysical retrievals from dual-polarization radar measurements at x band," *Journal of Atmospheric and Oceanic Technology*, vol. 25, pp. 729–741, 2014/01/17 2008.

[10] M. I. Mishchenko and L. D. Travis, "T-matrix computations of light scattering by large spheroidal particles," *Optics communications*, vol. 109, no. 1, pp. 16–21, 1994.

[11] R. Clausius, "Ueber die bewegende kraft der wärme und die gesetze, welche sich daraus für die wärmelehre selbst ableiten lassen," *Annalen der Physik*, vol. 155, no. 3, pp. 368–397, 1850.

[12] É. Clapeyron, *Mémoire sur la puissance motrice de la chaleur*. J. Gabay, 1834.

[13] H. R. Pruppacher, J. D. Klett, and P. K. Wang, *Microphysics of clouds and precipitation*. Taylor & Francis, 1998.

[14] M. Mischenko and L. Travis, "T-matrix computations of light scattering by nonspherical particles: a review," *J. Quant. Spectrosc. Radiat. Transfer*, vol. 55, pp. 535–575, 1996.

[15] J. S. Marshall and W. M. K. Palmer, "The distribution of raindrops with size," *Journal of meteorology*, vol. 5, no. 4, pp. 165–166, 1948.

[16] C. W. Ulbrich and D. Atlas, "Rainfall microphysics and radar properties: Analysis methods for drop size spectra," *Journal of Applied Meteorology*, vol. 37, pp. 912–923, 2014/01/14 1998.

[17] P. T. Willis, "Functional fits to some observed drop size distributions and parameterization of rain," *Journal of the Atmospheric Sciences*, vol. 41, no. 9, pp. 1648–1661, 1984.

[18] K. V. Beard and C. Chuang, "A new model for the equilibrium shape of raindrops," *Journal of the Atmospheric Sciences*, vol. 44, pp. 1509–1524, 2014/05/11 1987.

[19] M. Thurai and V. N. Bringi, "Drop axis ratios from a 2d video disdrometer," *Journal of Atmospheric and Oceanic Technology*, vol. 22, pp. 966–978, 2014/01/21 2005.

[20] A. Ryzhkov and D. Zrnic, "Comparison of dual-polarization radar estimators of rain," *Journal of Atmospheric and Oceanic Technology*, vol. 12, no. 2, pp. 249–256, 1995.

[21] K. Aydin, V. Bringi, and L. Liu, "Rain-rate estimation in the presence of hail using s-band specific differential phase and other radar parameters," *Journal of Applied Meteorology*, vol. 34, no. 2, pp. 404–410, 1995.

[22] V. Chandrasekar, V. Bringi, N. Balakrishnan, and D. Zrnic, "Error structure of multiparameter radar and surface measurements of rainfall. part iii: Specific differential phase," *Journal of Atmospheric and Oceanic Technology*, vol. 7, no. 5, pp. 621–629, 1990.

[23] W. Hitschfeld and J. Bordan, "Errors inherent in the radar measurement of rainfall at attenuating wavelengths," *Journal of Meteorology*, vol. 11, pp. 58–67, 2014/01/17 1954.

[24] E. Gorgucci, V. Chandrasekar, and L. Baldini, "Can a unique model describe the raindrop shape–size relation? a clue from polarimetric radar measurements," *Journal of Atmospheric and Oceanic Technology*, vol. 26, pp. 1829–1842, 2014/01/17 2009.

[25] S. Lim and V. Chandrasekar, "A dual-polarization rain profiling algorithm," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 44, pp. 1011–1021, April 2006.

[26] E. Yoshikawa, V. Chandrasekar, and T. Ushio, "Raindrop size distribution (dsd) retrieval for x-band dual-polarization radar," *Journal of Atmospheric and Oceanic Technology*, vol. 31, pp. 387–403, 2014/05/11 2013.

[27] C. Langston, J. Zhang, and K. Howard, "Four-dimensional dynamic radar mosaic," *Journal of atmospheric and oceanic technology*, vol. 24, no. 5, pp. 776–790, 2007.

[28] J. Zhang, K. Howard, C. Langston, S. Vasiloff, B. Kaney, A. Arthur, S. Van Cooten, K. Kelleher, D. Kitzmiller, F. Ding, *et al.*, "National mosaic and multi-sensor qpe (nmq) system: Description, results, and future plans," *Bulletin of the American Meteorological Society*, vol. 92, no. 10, pp. 1321–1338, 2011.

[29] J. Zhang, K. Howard, and J. Gourley, "Constructing three-dimensional multiple-radar reflectivity mosaics: Examples of convective storms and stratiform rain echoes," *Journal of Atmospheric and Oceanic Technology*, vol. 22, no. 1, pp. 30–42, 2005.

[30] V. Chandrasekar and S. Lim, "Retrieval of reflectivity in a networked radar environment.," *Journal of Atmospheric and Oceanic Technology*, vol. 25, no. 10, pp. 1755 – 1767, 2008.

[31] V. Bringi, V. Chandrasekar, J. Hubbert, E. Gorgucci, W. Randeu, and M. Schoenhuber, "Raindrop size distribution in different climatic regimes from disdrometer and dual-polarized radar analysis," *Journal of the atmospheric sciences*, vol. 60, no. 2, pp. 354–365, 2003.

[32] S. Matrosov, A. Heymsfield, and Z. Wang, "Dual-frequency radar ratio of nonspherical atmospheric hydrometeors," *Geophysical research letters*, vol. 32, no. 13, 2005.

[33] S. J. Munchak and A. Tokay, "Retrieval of raindrop size distribution from simulated dual-frequency radar measurements," *Journal of Applied Meteorology and Climatology*, vol. 47, no. 1, pp. 223–239, 2008.

[34] G. I. Taylor, "The spectrum of turbulence," *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 164, no. 919, pp. 476–490, 1938.

[35] I. Zawadzki, "Statistical properties of precipitation patterns," *Journal of Applied Meteorology*, vol. 12, no. 3, pp. 459–472, 1973.

[36] I. Zawadzki, "Fractal structure and exponential decorrelation in rain," *Journal of Geophysical Research: Atmospheres (1984–2012)*, vol. 92, no. D8, pp. 9586–9590, 1987.

[37] G. W. Lee, A. W. Seed, and I. Zawadzki, "Modeling the variability of drop size distributions in space and time," *Journal of applied meteorology and climatology*, vol. 46, no. 6, pp. 742–756, 2007.

[38] G. W. Lee, "Sources of errors in rainfall measurements by polarimetric radar: Variability of drop size distributions, observational noise, and variation of relationships between r and polarimetric parameters," *Journal of Atmospheric and Oceanic Technology*, vol. 23, no. 8, pp. 1005–1028, 2006.

[39] U. Germann and I. Zawadzki, "Scale-dependence of the predictability of precipitation from continental radar images. part i: Description of the methodology," *Monthly Weather Review*, vol. 130, no. 12, pp. 2859–2873, 2002.

[40] J. Leinonen, "High-level interface to t-matrix scattering calculations: architecture, capabilities, and limitations," *Opt. Express*, vol. 22, 2014.

[41] S. G. Nash, "Newton-type minimization via the lanczos method," *SIAM Journal on Numerical Analysis*, vol. 21, no. 4, pp. 770–788, 1984.

[42] S. J. Wright and J. Nocedal, *Numerical optimization*, vol. 2. Springer New York, 1999.

[43] M. Ignaccolo and C. D. Michele, "Phase space parameterization of rain: The inadequacy of gamma distribution," *Journal of Applied Meeorology and Climatology*, vol. 53, no. 2, pp. 548–562, 2014.

[44] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International journal of forecasting*, vol. 22, no. 4, pp. 679–688, 2006.

[45] Y. Wang and V. Chandrasekar, "Algorithm for estimation of the specific differential phase," *Journal of Atmospheric and Oceanic Technology*, vol. 26, no. 12, pp. 2565–2578, 2009.

[46] V. Chandrasekar, M. R. Schwaller, M. Vega, J. R. Carswell, K. V. Mishra, R. Meneghini, and C. Nguyen, "Scientific and engineering overview of the nasa dual-frequency dual-polarized doppler radar (d3r) system for gpm ground validation.," in *IGARSS*, pp. 1308–1311, 2010.

[47] A. Ryzhkov and D. Zrnic, "Precipitation and attenuation measurements at a 10-cm wavelength," *Journal of Applied Meteorology*, vol. 34, no. 10, pp. 2121–2134, 1995.

[48] F. Junyent, V. Chandrasekar, D. McLaughlin, E. Insanic, and N. Bharadwaj, "The casa integrated project 1 networked radar system," *Journal of Atmospheric and Oceanic Technology*, vol. 27, no. 1, pp. 61–78, 2010.

[49] R. J. Trapp and C. A. Doswell III, "Radar data objective analysis," *Journal of Atmospheric and Oceanic Technology*, vol. 17, no. 2, pp. 105–120, 2000.

[50] S. L. Barnes, "A technique for maximizing details in numerical weather map analysis," *Journal of Applied Meteorology*, vol. 3, no. 4, pp. 396–409, 1964.

[51] G. P. Cressman, "An operational objective analysis system," *Monthly Weather Review*, vol. 87, no. 10, pp. 367–374, 1959.

[52] V. Chandrasekar, S. Lim, and E. Gorgucci, "Simulation of x-band rainfall observations from s-band radar data.," *Journal of Atmospheric and Oceanic Technology*, vol. 23, no. 9, pp. 1195 – 1205, 2006.

[53] M. Steiner, J. A. Smith, S. J. Burges, C. V. Alonso, and R. W. Darden, "Effect of bias adjustment and rain gauge data quality control on radar rainfall estimation," *Water Resources Research*, vol. 35, no. 8, pp. 2487–2503, 1999.

[54] Y. Wang and V. Chandrasekar, "Quantitative precipitation estimation in the casa x-band dual-polarization radar network," *Journal of Atmospheric and Oceanic Technology*, vol. 27, no. 10, pp. 1665–1676, 2010.

[55] V. N. Bringi, V. Chandrasekar, J. Hubbert, E. Gorgucci, W. L. Randeu, and M. Schoenhuber, "Raindrop size distribution in different climatic regimes from disdrometer and dual-polarized radar analysis," *Journal of the Atmospheric Sciences*, vol. 60, pp. 354–365, 2014/01/17 2003.

[56] G. Penide, A. Protat, V. V. Kumar, and P. T. May, "Comparison of two convective/stratiform precipitation classification techniques: Radar reflectivity texture versus drop size distribution–based approach," *Journal of Atmospheric and Oceanic Technology*, vol. 30, no. 12, pp. 2788–2797, 2013.

[57] B. Li, A. Murthi, K. P. Bowman, G. R. North, M. G. Genton, and M. Sherman, "Statistical tests of taylor's hypothesis: an application to precipitation fields," *Journal of Hydrometeorology*, vol. 10, no. 1, pp. 254–265, 2009.

[58] A. Tokay, A. Kruger, and W. F. Krajewski, "Comparison of drop size distribution measurements by impact and optical disdrometers," *Journal of Applied Meteorology*, vol. 40, no. 11, pp. 2083–2097, 2001.

[59] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *Journal of the Society for Industrial & Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.

APPENDIX A

Description of Field Experiments and Datasets

"Everything went great right up to the explosion. "

— Andy Weir, The Martian

The proper development and evaluation of this algorithm requires a rich dataset. In this appendix we will discuss two different field projects that CSU participated in, and the resulting dataset that we will use to test and validate our algorithm. The first dataset is from the MidLatitude Continental Convective Clouds Experiment, while the second one is the Iowa Flooding Experiment. After discussing both field campaigns we will provide some description of the various instruments we will use. We will conclude with a discussion of some of the algorithms that we used to derive auxillary data.

A.1. Midlatitude Continental Convective Clouds Experiment

The Midlatitude Continental Convective Clouds Experiment(MC3E) was a joint field experiment between NASA and the Department of Energy(DOE). The field project took place during the summer of 2011 in northern Oklahoma centered around the DOE Atmospheric Radiation Measurement (ARM) Climate Research Facility's Southern Great Plains(SGP) site. The purpose of the experiment was primarily to study convective initiation of thunderstorms in the mid-latitudes at multiple scales. The experiment consisted of a dense network of active and passive instruments. The roster of instruments includes three X-SAPR X-band radars, one C-SAPR C-band radar, the NASA NPOL S-Band radar, as well as an array of rain gauges and disdrometers. In addition to the ground radars and disdrometers, multiple radiometers, cloud radars, and millimeter wave instrumentation was available. The NASA ER-2 and University of North Dakota Citation aircraft were also flying coordinated flights.

The radars had an overlapping field of view providing us with a good dataset for multi-frequency analysis. The layout of the subset of the instruments that we are most interested in is shown in figure A.1.

FIGURE A.1. MC3E Instrument Layout. Shown are the ground scanning radars and several auxillary instruments.

While this dataset provides the ideal setup, unfortunately there are some issues with the data quality and operational status of the radars. A large portion of the ground instrumentation was purchased with recent stimulus funding and scaled up very quickly. This causes some issues with calibration. Fortunately, the majority of these issues can be fixed in post

processing with some care. More issues exist with the operational status of the instrumentation. Instruments were frequently broken. Despite all of this, the MC3E dataset provides a huge opportunity for research. The dataset was able to capture many different convective cases, as well as some stratiform cases. There was hail, rain, tornadoes, and every variation of weather one could expect from a summer in the midwest.

A.2. Iowa Flood Experiment

The Iowa Flooding experiment(IFLOODS) was a joint experiment between NASA and the Iowa Flood Center to study variability of rainrate at sub 1km scales. There was a variety of instrumentation, though the core of the experiment was the combination of NASA's NPOL radar, the NASA D3R radar, and a very dense array of disdrometers and rain gauges aligned along a radial. Four X-Band radars were also present from the University of Iowa, but initial analysis of the data quality showed it was of unacceptable quality and so we will not for now include it into this analysis. A map showing the layout of some of the instrumentation by the Iowa Flood Center is shown in Figure A.2.

The primary reason for our interest in the IFLOODS dataset is the dense sampling of disdrometers. This allows us a very good dataset with a large amount of ground truth to test our algorithm with. The disdrometers are arrayed along a radial of the radar to the southeast displayed in yellow in Figure A.2. The NASA NPOL and D3R radars are located in the center of the image. The range rings are spaced 25 km. An array of tipping bucket rain gauges are displayed to the northeast as well, although these are out of range of the D3R radar.

A.3. Radar Descriptions

A listing of the different radars and some of their specifications is shown in Table A.1.
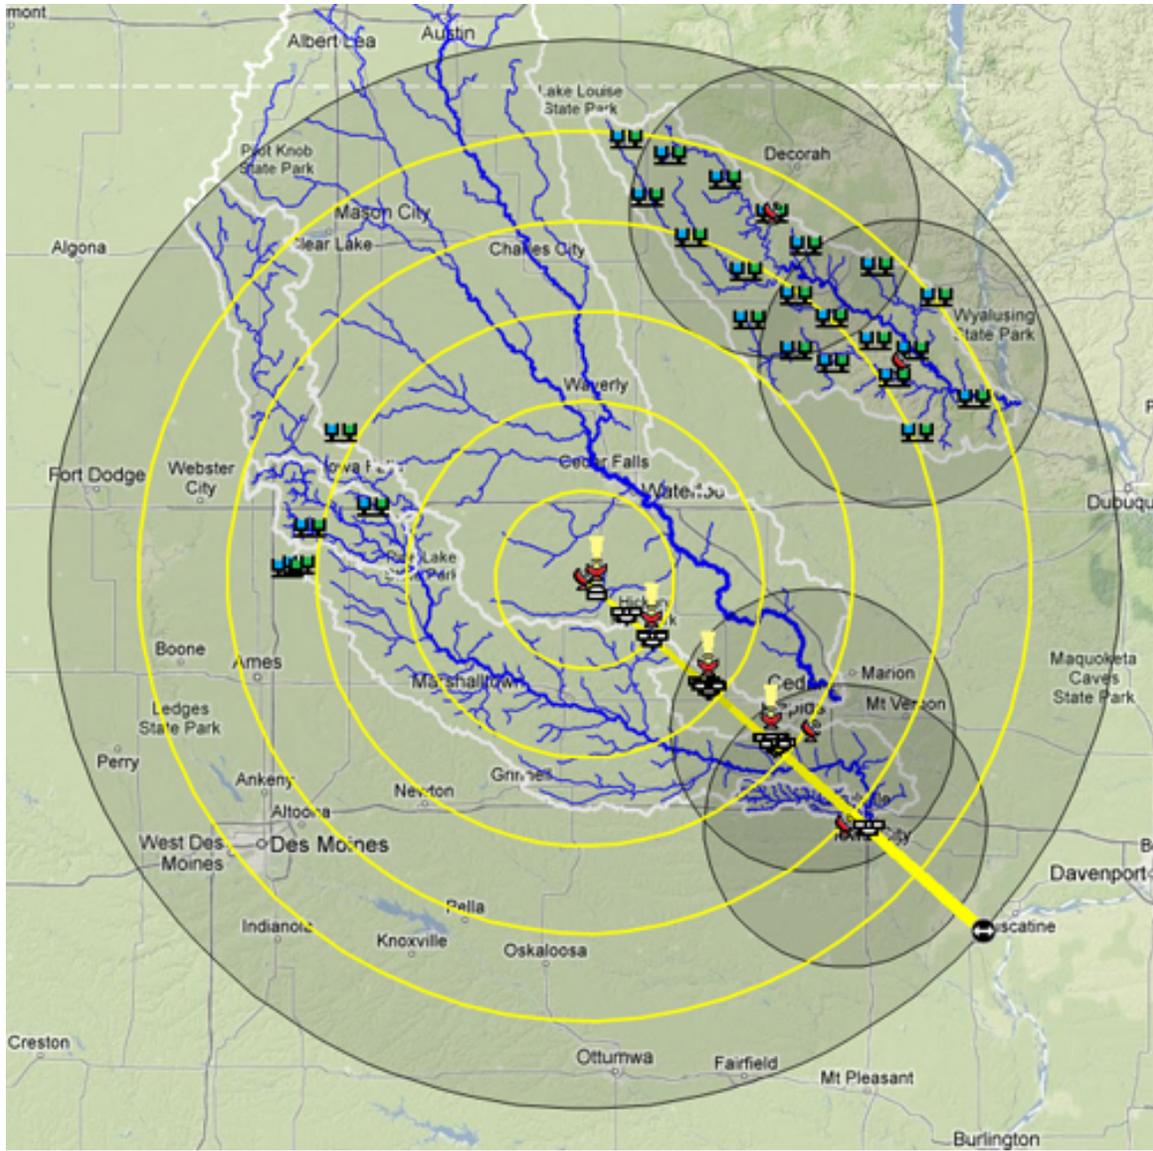
FIGURE A.2. IFLOOD Instrument Layout. Courtesy of Iowa Flood Center

	X-SAPR	C-SAPR	NPOL	CSU-CHILL(S)	CSU-CHILL(X)	D3R(Ku)
Antenna Diameter	2.4m	4.27m	8.5m	8.5m	8.5m	1 m
Antenna Beam Width	0.9°	0.98°	0.9°	1.0°	0.33°	0.86°
Antenna Gain	45 dBi	45 dBi	43 dBi	43 dBi	53 dBi	43 dBi
Polarization type	Sim	Sim	Sim + Alt	Sim + Alt	Sim	Sim/Alt
Operating Frequency	9.5GHz	5.625 GHz	2.725 GHz	2.725 GHz	9.41GHz	13.91 GHz
Transmitter Peak Power	200kW	250kW	850 kW	1MW/channel	25kW	160 W
Minimum Pulse Width	200ns	200ns	200ns	0.2-1.6μs	-	0-100μs
Max PRF	5kHz	5kHz	-	1250 MHz	2GHz	2 kHz
Receiver Dynamic Range	80dB	80dB	-	80 dB	90dB	90 dB
Max Range	40km	120km	300km	300 km	-	40 km

TABLE A.1. Radar Characteristics

A.4. Error Metrics

The literature on forecasting uses a large number of different error metrics, each with their own strengths and weaknesses. As we do not have access to the original datasets and algorithms most papers have utilized, comparing the results between papers is not necessarily straightforward. Additionally, each metric sometimes has slightly different implementations with regards to some of the implementation details such as missing data.

We will use the following error metrics. A more detailed discussion of each metric can be found in the literature[44].

Let

$$(67) \qquad \vec{Y} = (Y_0, Y_1, ...Y_i, ..., Y_n)$$

senote a series of measurements and

$$(68) \qquad \vec{F} = (F_0, F_1, ..., F_i, ..., F_n)$$

denote the output of an estimator for the series $\vec{Y}$. The error term for each retrieval point is then

$$(69) \qquad e_i = Y_i - F_i$$

The error metrics used here are almost all normalized error metrics. The scale the error to the value being forecasted. This is a desirable trait, as otherwise the error metric used

is dependent on the scale of the data being used, making it impossible to cross compare performance over different cases and with other algorithms.

A.4.1. MEAN SQUARE ERROR (MSE). Mean square error is one of two non-normalized metrics used in this work, and is only included for completelyness. In general, it highly depends on the underlying data field, and cannot be compared across multiple studies. The formula for it is given as

$$(70) \qquad\qquad \text{MSE} = \frac{1}{n} \sum_{i=1}^{n} e_i^2$$

In general, this work will not use MSE to draw any conclusions, and just provides it for reference.

A.4.2. MEAN ABSOLUTE ERROR (MAE). The second non-normalized metric is the mean absolute error. This metric gives us the error expected for a given dataset, but again is scale dependent so it cannot be cross compared between studies easily. It is calculated as

$$(71) \qquad\qquad \text{MAE} = \frac{1}{n} \sum_{i=1}^{n} |e_i|$$

What the MAE scores provides is a good way to compare performance on the same dataset between different algorithms and test conditions over the same underlying dataset.

A.4.3. MEAN ABSOLUTE SCALED ERROR (MASE). Mean Absolute Scaled Error (MASE) is one of the more interesting metrics included here. Although not as common as any of the other metrics discussed, it has quite a few desirable conditions. Simply put, the MASE score

of a predictor is how much better it does than a naieve one step predictor. This means it has a level of scale invariance. A score of 1 for MASE denotes that the estimator has the same performance as a one-step estimator that uses the previous value of the measurement. Scores lower denote a better estimator, while scores higher denote a worse estimator. This score metric works best when there is a measure of wide sense stationarity in the measured sequence. Unfortunately, as no other papers report MASE scores, there is no way of comparing it with the literature on microphysical retrievals.

The formula for MASE score is given as:

$$\text{(72)} \qquad \text{MASE} = \frac{1}{n} \sum_{i=1}^{n} \frac{|e_i|}{\frac{1}{n-1} \sum_{i=2}^{n} |Y_i - Y_{i-1}|}$$

A.4.4. NORMALIZED STANDARD ERROR (NSE). Normalized Standard Error (NSE), also called Normalized Absolute Error, is a scaled error metric commonly used in much of the retrieval work in the literature. Although several different forms are all called by the same name, the one used here will be given by

$$\text{(73)} \qquad \text{NSE} = \frac{\frac{1}{n} \sum_{i=1}^{n} |e_i|}{\frac{1}{n} \sum Y_i}$$

A.4.5. NORMALIZED BIAS (NB). The second metric most commmonly seen in the retrieval literature is the Normalized Bias (NB). It provides a good measure of systematic over or under estimates of parameters. The form for NB used in this work is

$$(74) \qquad \text{NB} = \frac{\sum_{i=1}^{n} e_i}{\sum_{i=i}^{n} Y_i}$$

APPENDIX B

TOOL DEVELOPMENT AND OPEN SOURCE CONTRIBUTIONS

OF THIS WORK

> "Empowerment of individuals is a key part of what makes open source
> work, since in the end, innovations tend to come from small groups,
> not from large, structured efforts."
>
> — Tim O'Reilly

This work, as in most dissertation level research efforts, included a significant amount of custom designed software. What is different in this research is the outlook we took on the software development. Most academic software is written in a haphazard manner, designed to be used to develop results for the next journal paper, before being discarded. This leads to very large amounts of time being wasted re-implementing the software by other researchers. In addition, this allows the all too common defense when results can't be recreated that "You implemented the algorithm incorrectly". There are many reasons for the situation to be as it is. These usually stem from lack of funding and time to polish software to be released to the world, as well as the issue of responsibility if the software turns out to be incorrect.

When software is released, it is often passed along through e-mails, or if one is lucky, on an academic website. These efforts are great, but can be very cumbersome to maintain. If a user wishes to ask a question, it needs to go through e-mail. If multiple users have the same question, it entails multiple similar responses by the authors of the software. It also makes it very hard to have a *community* of users that can help support each other. This can be a major impediment to openly releasing code on the internet for other researchers to utilize.

Partially in response to this, the academic community has started to release code on social version control sites such as Github or Bitbucket.These sites provide code hosting, version

control, and community features such as wiki's, documentation, issue and bug tracking, as well as the ability to contribute back to other peoples code bases. This makes it much more simple for users to contribute back bug fixes, and interact with the authors of the software. In addition, there are traffic tracking metrics that show level of interest in a project, views, and how many people have downloaded the code to modify. This allows for metrics that can be presented to funding agencies and universities to help support the writing of this software. These systems have had a major impact on how academic code is written and shared recently.

Another confounding factor is the lack of recognition for software writers. Time spent providing and supporting software is essentially a *sunk cost*, in that it does not count towards publication records or many other metrics used for career progression in academia. This stops many well-meaning academics from being able to contribute to the advancement of the field in a much more concrete way than publishing papers.

Recently, there has been a large push to rectify this by issuing Digital Object Identifier(DOI) numbers. These traditionally have been used by papers to provide a unique string to identify electronic documents. Many different open source companies have started issuing these numbers to allow an easier way to cite, and track citations of, academic open source software. In this way, one can get "citations" for your software, which provides a useful metric for advancement in the academic field.

In this appendix, we will detail two different libraries created as part of this work that have been subsequently open sourced and released to the academic community. These libraries are supported and documented, and available publically on Github. The first library,

PyDisdrometer is a library for working with disdrometer data that has seen uptake by different research groups and organizations, while the second PythonRadarTools is a smaller set of tools released to the community that is brand new.

B.1. PyDisdrometer

PyDisdrometer is an open source python library created at CSU as part of this research to enable researchers to work with and process disdrometer data. Disdrometers, as mentioned earlier in this dissertation, are ground based devices that measure the drop size distribution(DSD) of hydrometeors at ground level. This software fills a major gap as there is no previously existing open source software to work with disdrometer data that we are aware of.

PyDisdrometer is object oriented and revolves around the concept of the DSD as a class. Modules that are responsible for reading in file formats all return an instantiation of a DropSizeDistribution class object. This object is self contained and contains methods that are able to operate on the object itself, allowing for a much more straightforward API. This allows functions access to internal data about the drop size distribution, minimizing the number of parameters that must be passed to any given function.

PyDisdrometer is designed to perform several different tasks. These include:

- File Processing
- Microphysics Parameter Estimation
- Radar Scattering Simulation
- Generation of optimal dual and single polarization rain rate estimators.

In the following section we will discuss the architecture of PyDisdrometer and discuss the implementation of this functionality.

B.1.1. PYDISDROMETER ARCHITECTURE.

B.1.1.1. *File Input/Output.* PyDisdrometer currently has file readers for several major formats. This includes the following:

- Joss-Waldvogel

- Parsivel TeleFormat

- Parsivel NASA Ground Validation(MC3E/IFloodS/IPHEX campaigns, Wallops Experimental)

- 2D Video Disdrometer NASA Ground Validation(MC3E/IFloodS/IPHEX field campaigns)

Each of these data formats provides a slightly different set of metadata. PyDisdrometer attempts to fill in as much derived data as possible. Currently NASA Ground Validation(NASA-GV) format files do not actually provide the rain rate as reported by the disdrometers. For these, methods are provided that calculate rainrates based on the vertical velocity relationships given in [1] as

$$(75) \qquad\qquad v(D) = 9.65 - 10.3 \exp\left(-0.6D\right) \text{ ms}^{-1}$$

which lets us estimate the rain rate as

$$(76) \qquad\qquad R = 0.6\pi * 10e^{-3} \int_0^{D_{max}} v(D)D^3 N(D) dD \text{ mmh}^{-1}$$

If a file includes the vertical velocity, then this is used to calculate the rain rate in preference over the terminal velocity assumption.

A second issue used in several of the file readers is that of drop bounces. This occurs when drops strike the disdrometer and then bounce into the sampling area. This is detectable by

comparing the measured velocity with the terminal velocity. Based on the method by Tokay et. al.[58], all drops whose velocity deviates by more than 50% from terminal velocity are eliminated from the measurements. This processing can only take place on files where the drop velocity is reported individually for each bin.

All file readers return a DropSizeDistribution object.

B.1.1.2. *Scattering.* DropSizeDistribution objects have the ability to calculate radar measured parameters based on the binned drop size distribution. The scattering code is based on the T-Matrix technique[14], and utilizes the PyTMatrix library[40] written by Jussi Leinonen, which wraps an earlier Fortran library written by Michael Mishchenko[10].

Scattering is accomplished using the DropSizeDistribution.calculate_radar_parameters function. It contains different options for simulating scattered radar data from the drop size distribution. The primary option to change is the wavelength of the radar frequency to simulate. As it uses the PyTMatrix library for scattering, any of the built in frequencies listed in table B.1 are accepted.

TABLE B.1. PyDisdrometer Frequency List. Frequencies are drawn from underlying PyTMatrix library.

Option	wavelength
wl_S	11.0 cm
wl_C	5.53 cm
wl_X	3.33 cm
wl_Ku	2.2 cm
wl_Ka	84.3 cm
wl_W	.319 cm

The second major option is the Drop Shape relationship. Three different options are supported and listed in table

These shape relationships are shown in figure B.1.

Shape function	Shape Relationship
DSR.pb	Pruppacher and Beard[4]
DSR.bc	Beard and Chuang [18]
DSR.tb	Thurai and Bringi [19]
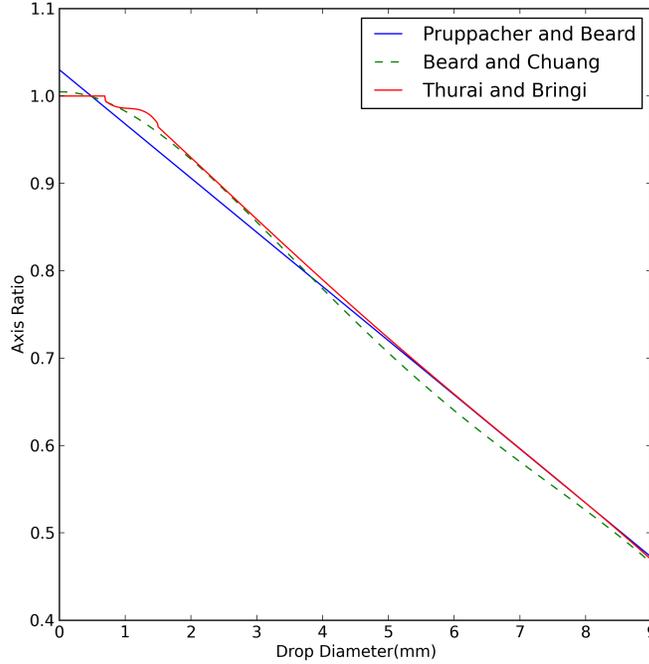
FIGURE B.1. Drop Shape Relationships

Scattering is performed on the binned histogram itself(rather than on the estimated parametric distribution) which allows us to handle distributions that do not fit the normalized gamma assumption. Based on the underlying T-Matrix code many different dual-polarization parameters are estimated and stored in a self describing dictionary inside the DropSizeDistribution object. These are listed in table B.3.

B.1.1.3. *Rain Rate Estimator Fitting.* PyDisdrometer has the ability to calculate power law fits for coefficients for rain rate estimators. This includes all permutations of Z_h, Z_{dr} and K_{dp}. A common mistake when estimating these fits is the linearization of the equations by

TABLE B.3. PyDisdrometer Radar Moments

Radar Moments		
Reflectivity	Differential Reflectivity	Specific Differential Phase
Attenuation	Differential Attenuation	Linear Depolarization Ratio
Backscatter Differential Phase		

taking the logarithm of both sides. While this approach works if there are no error sources of any kind in the measurements, it incorrectly weights the variance of the error. To accurately fit the parameters, PyDisdrometer uses a solution of the nonlinear least squares minimization of the equation

$$(77) \qquad S(\beta) = \sum_{i=1}^{m} (y_i - f(x_i, \beta))^2$$

The algorithm used for this is Levenberg-Marquardt[59]. This provides a more accurate solution to the equation. All rain rate estimators will be of the form

$$(78) \qquad R = aZ_h^\alpha Z_{dr}^\beta K_{dp}^\gamma$$

where each of the parameters α, β, γ are allowed to be zero if explicitly set by the user. This has the effect of removing that parameter from the estimation. This process has proven to be very accurate in practice based on disdrometers from several different field campaigns that were co-located with rain gauges.

B.1.1.4. *Microphysics Estimation.* Disdrometers return binned drop counts and not microphysical parameters, nor parameterized distributions. PyDisdrometer implements several algorithms to estimate different microphysical parameters. These are called from the Drop-SizeDistribution.calculate_dsd_parameterization function.

We can start by calculating the total liquid water content of the distribution as

$$(79) \qquad W = \frac{\pi}{6}\rho_w \int_0^{D_{max}} D^3 N(D)dD$$

where ρ_w is the density of water.

The mass weighted median drop diameter(D_m) is given as:

$$(80) \qquad D_m = \frac{\langle D^4 \rangle}{\langle D^3 \rangle}$$

where $\langle \cdot^N \rangle$ represents the Nth moment of the distribution. We calculate a related quantity, the median drop diameter(D_0) as the value that solves

$$(81) \qquad \frac{\pi}{6} \int_0^{D_0} D^3 N(D)dD = \frac{1}{2}W$$

We calculate the normalized drop total as

$$(82) \qquad N_t = \int_0^{D_{max}} N(D)dD$$

Finally we calculate the normalized intercept parameter as

$$(83) \qquad N_W = \frac{256}{\pi \rho_w} \frac{W}{D_m^4}$$

178