

Analysis of Tools Used to Quantify Droplet Clustering in Clouds

BRAD BAKER AND R. PAUL LAWSON

SPEC, Inc., Boulder, Colorado

(Manuscript received 10 December 2009, in final form 10 April 2010)

ABSTRACT

The spacing of cloud droplets observed along an approximately horizontal line through a cloud may be analyzed using a variety of techniques to reveal structure on small scales, sometimes called clustering, if such structure exists. A number of techniques have been applied and others have been suggested but not yet rigorously defined and applied. In this paper techniques are studied and evaluated using synthetic droplet spacing data. For the type of small-scale structure (clustering) modeled in this study, the most promising analysis approach is to use a combination of the power spectrum and the fishing statistic. Standard deviations and confidence intervals are determined for the power spectrum, the pair correlation function, and a modified fishing statistic. The clustering index and the volume-averaged pair correlation are shown to be less usefully normalized forms of the fishing statistic.

1. Introduction

The possibility that spatial organization of cloud droplets caused by their inertia in turbulent flow could lead to significant effects on clouds as a whole via more variable droplet growth rates through competition for the vapor, enhanced collision and coalescence, and/or radiative effects has been discussed and debated in the literature (Shaw et al. 1998; Pinsky and Khain 2001; Grabowski and Vaillancourt 1999; Knyazikhin et al. 2005; Marshak et al. 2005). A number of researchers have investigated droplet spacing in clouds using data obtained from droplet-counting probes mounted on aircraft (Baker 1992; Pinsky and Khain 2001; Chaumat and Brenguier 2001; Marshak et al. 2005; Kostinski and Shaw 2001), on a balloon (Lehmann et al. 2007), and in a wind tunnel (Saw et al. 2008). The questions that researchers have attempted to address are 1) whether there is some organization (clustering) in excess of what would be expected from turbulent entrainment and mixing and 2) whether such structure has a significant effect on cloud microphysics. Shaw et al. (2002) discuss mathematical tools for analyzing droplet spacing data for the purpose of investigating clustering. This paper describes a continuation of those efforts by evaluating

various analysis techniques, both theoretically and via synthetic data. In addition to evaluating the usefulness of each technique, confidence intervals are determined, making the tests more rigorous.

The techniques to be evaluated are the fishing statistic F (Baker 1992), the technique of Marshak et al. (2005), and three techniques discussed in Shaw et al. (2002): the clustering index CI (Chaumat and Brenguier 2001), the pair correlation PC , and the volume-averaged pair correlation $VAPC$. We also evaluate the power spectrum PS . An estimate of the power spectrum was used by Pinsky and Khain (2001) among other techniques. All the tests, with the exception of Marshak's, are mathematically related and as such contain similar information. The difference is in how usefully that information is displayed and how sensitive the method is to structure in the data.

2. Fishing statistic, clustering index, and volume-averaged pair correlation function

The fishing statistic was developed to detect structures in droplet concentration, derived from time series of droplet counts, on spatial scales too small for the structure to be detected by eye. It is a hypothesis-testing statistic that tests whether a given set of data points is Poisson distributed and is referred to as either a "statistic" or "test" throughout this manuscript. Where cloud concentration is uniform, a series of counts per sampled volume will be Poisson distributed. For aircraft-mounted

Corresponding author address: Brad Baker, SPEC, Inc., 3022 Sterling Circle, Suite 200, Boulder, CO 80301.
E-mail: brad@specinc.com




FIG. 1. Synthetic droplet spacing data converted to concentration (counts per distance) time series. (a) Actual concentration modeled, which consists of alternating constant-concentration blocks of 1-cm length, with concentrations of 20 and 30 cm^{-1} . (b) Synthesized data sequence of simulated measured concentrations assuming random droplet placements given the true concentrations in (a).

probes, measurements in time are directly related to sampled volume if the true airspeed and probe sample area are constant. We calculate and display measured time series of counts per sampled volume using the distance traveled by the aircraft d as the independent variable. The size of the distance bin used L may be varied and is the independent variable for the hypothesis testing statistics. Equation (1) displays the formulas for F , CI, and VAPC as functions of L :

$$F(L) = \left(\frac{\frac{V_X}{M_X} - 1}{\sqrt{\frac{2}{N_X - 1}}} \right), \quad \text{CI}(L) = \left(\frac{V_X}{M_X} - 1 \right),$$

$$\text{VAPC}(L) = \frac{\left(\frac{V_X}{M_X} - 1 \right)}{M_X}, \quad (1)$$

where X are the N_X data points of a counts-per-distance bin (distance bin length = L) series. That is, a long sample of particle counts is divided in distance bins of length L , where X is the number of counts in each bin and N_X is the total number of bins in the sample (N_X decreases as L increases); V_X is the variance of X and M_X is the mean of X . All three statistics, F , CI, and VAPC, are based on the dispersion index (V_X/M_X) minus its expected value (i.e., one) under the null hypothesis that X are Poisson distributed. Thus, all three statistics have expected values of zero. The differences among these tests lie only in their normalizations. Note that F is normalized by the standard deviation σ of the dispersion index. This has the advantage that the statistics of F (e.g., the

probability of any certain value under the null hypothesis) are approximately independent of L and N_X , whereas for CI and VAPC the probability of any certain value under the null hypothesis varies greatly with L and N_X .

This advantage is demonstrated with the help of synthetic data and is shown in Figs. 1 and 2. Droplet spacing data are modeled using a random number generator. In this case the minimum L is 10 μm , the average distance between droplets is 400 μm , the average concentration C_{ave} is equivalently 25 cm^{-1} , and 2 m of data are synthesized ($L_{\text{max}} = 2$ m). Structure is built in by making the average distance between droplets vary in alternating 1-cm blocks. In this case, the average distance between droplets was 500 and 333 μm in the alternating blocks. Figure 1a shows the ideal time series of concentration that is being modeled while Fig. 1b shows one model realization of that time series, calculated with 1-cm distance bins. Because of the randomness of small volume observations, it is not easy to detect the structure by eye even though the time series was calculated at the resolution of the imposed structure and the distance bins were aligned with the actual structure. This time series and four additional random realizations were analyzed via the three statistics F , CI, and VAPC. The results are shown in Fig. 2. On each figure a dashed line represents the value of 3σ away from the expected mean under the null hypothesis. Because of the imposed structure, all three tests show values exceeding the 3σ curve, peaking at 0.5 cm, half the size of the alternating blocks. All three tests are equivalent in information content; however, because of the normalization it is easier to quickly detect that information with the fishing test (i.e., the peak is more distinct).

At this point it is worth discussing a subtle point that leads to an apparent inconsistency here and later in this paper. Shaw et al. (2002) and Kostinski and Shaw (2001) show that if one assumes that PC is a purely local measure of structure, then F , CI, and VAPC represent integrated effects of structure on all smaller scales since they are related to PC by integration. They also suggest that VAPC should be more local than F or CI, which we find is not the case (Fig. 1) since they are all essentially the same except for their normalizations. The subtlety is that no measure is universally or fundamentally “the” local measure of structure. Any measure may be more or less local depending on the form of the structure being detected. Below, we show that for periodic structure the PC is less local than F or the power spectrum. Other forms of structure could presumably be constructed for which the opposite is true.

3. The Marshak technique

The technique of Marshak et al. (2005) divides the measurement period into subsections of equal size L and counts the number N of subsections containing one or more droplets of a given size¹ R ; N is therefore a function of R and L . The data are fit to a power-law function as

$$N(R, L) \propto L^{-D(R)}, \quad (2)$$

where the exponent D is a function of droplet size R .

Marshak et al. (2005) interpret the observed values of D as follows: $D(R) = 1$ implies that the droplets of size R are Poisson distributed (i.e., not clustered), $0 < D(R) < 1$ implies droplets of size R are clustered, and $D(R) = 0$ is a trivial case of sparse droplets of size R (no distance bin of size L contains more than one drop of size R); $D(R) > 1$ does not occur. We will show that these interpretations are incorrect and that the correct interpretation is associated with analysis over regions containing both dense and sparse cloud; D is related to the relative fractions of dense versus sparse cloud, regardless of whether the dense and sparse cloud regions are large and adjacent or interspersed on small scales.

Using the model described in section 2 but with different values of L_{\max} and C_{ave} , two 400-m-long segments of data were synthesized ($L_{\max} = 400$ m), one clustered and one completely homogeneous. In both cases the mean droplet spacing² is the same, 10 cm ($C_{\text{ave}} = 0.1 \text{ cm}^{-1}$). In the clustered case the mean droplet spacing varies

¹ A size R refers to a range of sizes around the value R .

² This spacing is as measured by an instrument traversing the cloud and is distinct from the nearest neighbor spacing in three-dimensional space.




FIG. 2. Results of applying (a) the fishing test, (b) the clustering index, and (c) the volume-averaged pair correlation to five different realizations of modeled data as shown in Fig. 1. The dashed line in each plot represents 3σ from the expected value of zero under the null hypothesis of a pure Poisson process with no concentration variations, which also represents about the 99th percentile.

between 6.7 and 20 cm ($C_{\text{ave}} = 0.15 \text{ cm}^{-1}$ and 0.05 cm^{-1} , respectively) in alternating 100-cm blocks of cloud.

Figure 3 shows the fishing statistic F and $N(R, L)$ for both the homogeneous data and the clustered data. Also

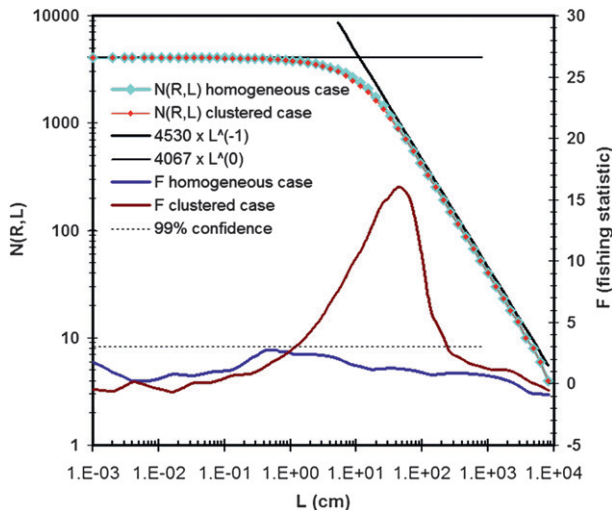


FIG. 3. Results of calculating the fishing statistic and $N(R, L)$ for both homogeneous and clustered model data, as described in the text. The dashed horizontal line at $F = 3$ represents 3σ for the fishing statistic under the null hypothesis, which is also about the 99% confidence level that the null hypothesis (i.e., that the data are homogeneous) is violated. Straight lines representing power-law relationships with exponents of 0 and -1 are also shown.

shown are power-law curves with $D = 1$ and $D = 0$ and the 3σ line for the fishing statistic. For the homogeneous case the fishing statistic remains below 3, as expected at least 99% of the time. For the clustered case the fishing statistic greatly exceeds 3, indicating the structure. The maximum is at 50 cm, which equals one-half the block size, as expected. However, $N(R, L)$ is nearly the same for both the homogeneous and the clustered cases, following a power law with exponent of -1 ($D = 1$) at L substantially larger than the average interparticle spacing (droplets are dense), and following a horizontal line ($D = 0$) at L substantially smaller than the average interparticle spacing (sparse droplets). This demonstrates that the technique is not effective at detecting clustering, leaving only the question of how does a slope D between 0 and 1 come about?

A straight line on the log-log plot, indicating a power law with slope $|D|$ between 0 and 1, may be obtained over a finite domain of L values by combining two otherwise homogeneous regions side by side, one where the droplets are dense and one where they are sparse. An example is shown in Fig. 4, which shows $N(R, L)$ for a modeled period of cloud that consists of adjacent regions, 100 m with a mean droplet spacing of 2.5 cm and 300 m with a mean droplet spacing of 500 cm. Except for the large-scale structure, there is no clustering in these synthesized data. The result is similar to that shown in Fig. 1 of Marshak et al. (2005), which was interpreted as implying clustering. Using the synthetic data, the slope




FIG. 4. Result of calculating $N(R, L)$ on model data that consist of adjacent homogenous regions of different concentration, as described in the text. Straight lines representing power-law relationships with exponents of -1 and -0.63 are also shown.

can be adjusted to various values between 0 and 1 by adjusting the relative sizes of the adjacent regions. Increasing the relative size of the sparse region compared to the size of the dense region decreases the slope $|D|$.

4. Confidence intervals for the pair correlation, the power spectrum, and the fishing statistic

Before we can evaluate the efficiency of the remaining tests, the pair correlation and the power spectrum, confidence intervals must be established under the null hypothesis of random droplet placement with equal probability everywhere. As a first step toward determining confidence intervals for the power spectrum and pair correlation, we revisit the statistics of the fishing test. This allows confirmation and extension of the results of Baker (1992) and makes the fishing statistic more sensitive.

a. The fishing statistic

Baker (1992) showed that under the null hypothesis, the basic fishing statistic has a mean value of 0 and standard deviation σ of 1. It was also found that the 99th percentile was about at or below the value of 3. The basic fishing statistic is not used on real data in Baker (1992) or in the remainder of this paper. Instead we use a modified fishing statistic, which for each L averages four F values determined with four different starting points, equally spaced within the first interval (i.e., at 0, $L/4$, $L/2$, $3L/4$ into the series). It was mentioned in Baker (1992) that this further reduces σ and the 99th percentile values, but the

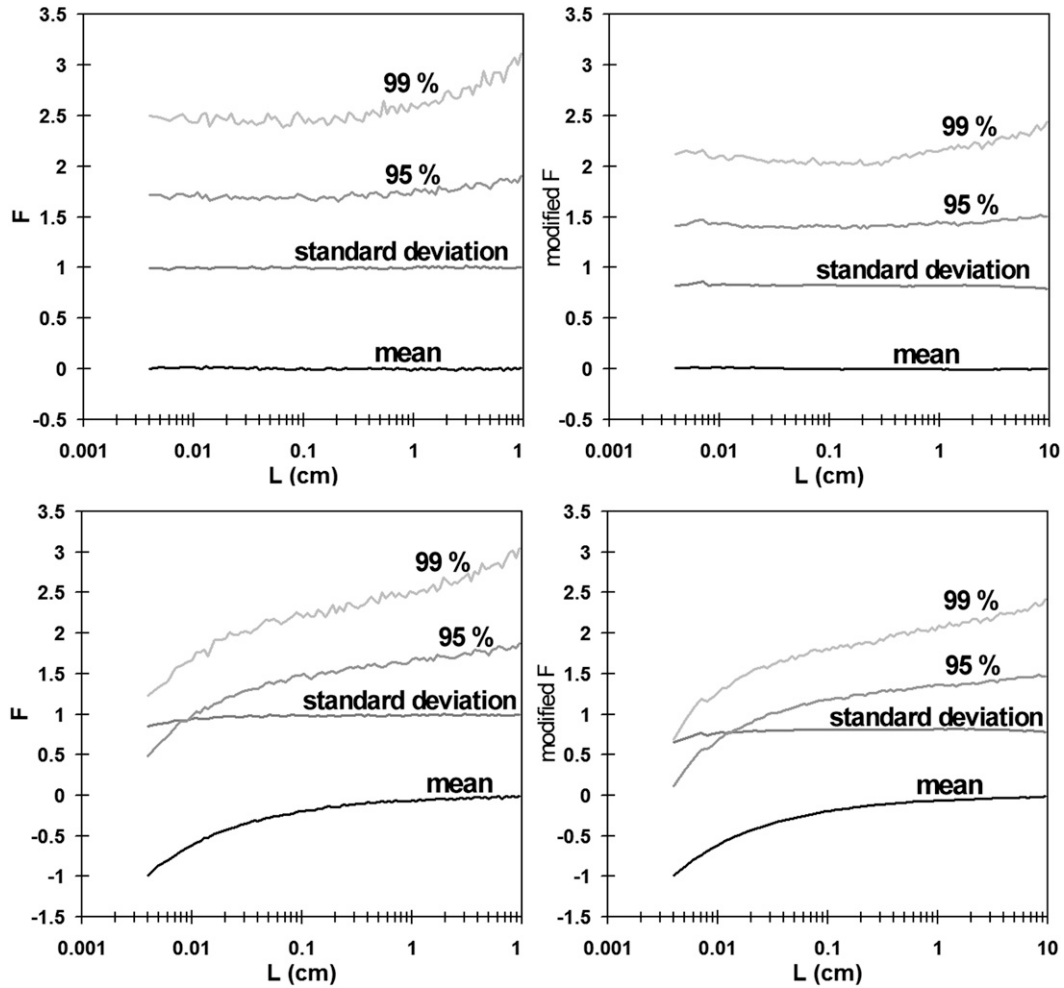


FIG. 5. The mean, σ , and 95th and 99th percentiles of both the (left) basic and (right) modified fishing statistics applied to 10 000 model (top) ideal and (bottom) noncoincident homogeneous time series of length 0.8 m and 10 cm^{-1} concentration.

amount was not quantified. With the vastly improved computing power currently available, sufficient data can be modeled and processed to more precisely determine these statistics. Further reduction could be obtained by averaging more F values, up to a maximum of L since there can be a maximum of L different starting points. However, the computing time increases proportionally with the number of F values calculated. Thus, averaging four values is a good compromise.

Figure 5 shows the result of applying both the basic fishing test and the modified test to 10 000 synthesized, homogenous time series of 1 m length ($L_{\text{max}} = 100\ 000$ distance bins or clock ticks at $10\text{-}\mu\text{m}$ resolution) and mean concentration C_{ave} of 10 cm^{-1} (0.01 per tick). The mean, σ , and 95th and 99th percentiles are shown. Even with a mean concentration of 0.01 droplets per $10\text{-}\mu\text{m}$ minimum-resolution distance bin (i.e., 10 cm^{-1}), there are occurrences of more than one droplet in $10 \mu\text{m}$.

However, droplet-counting probes cannot distinguish such coincident events. Therefore, each time series is modeled without coincidence (i.e., more than one droplet in $10 \mu\text{m}$ is counted as one) as well as modeling the ideal case, which, for comparison, allows more than one droplet in $10 \mu\text{m}$.

For the ideal time series and basic fishing statistic, the mean and standard deviation are 0 and 1, respectively, and are scale independent. The 95th and 99th percentiles are somewhat scale dependent. The 95th percentile remains below 2 for all scales. The 99th percentile remains below 3 for most scales except the largest where it exceeds 3 slightly. These results are consistent with Baker (1992). Similarly, the results for the basic statistic applied to noncoincident data are consistent with the results of Baker (1992). At large and medium scales the results are the same as for ideal data, while at small scales all the statistics are reduced.

TABLE 1. The values of the relevant model parameters (C_{ave} , L_{max}) used and the pair correlation statistics (mean, σ , and 95th and 99th percentiles) calculated from 1000 random realizations for each case.

	A	B	C	D	E	F	G	H
C_{ave} (# per tick)	0.005	0.01	0.01	0.05	0.1	0.5	1.0	5.0
L_{max} (ticks)	80 000	40 000	60 000	40 000	60 000	80 000	80 000	80 000
$1/(C_{\text{ave}}\sqrt{L_{\text{max}}})$	0.71	0.50	0.41	0.10	0.041	0.0071	0.0035	7.1×10^{-4}
Mean	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
σ	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Ideal	0.71	0.50	0.41	0.10	0.041	0.0071	0.0035	7.1×10^{-4}
Noncoincident	0.71	0.50	0.41	0.098	0.039	0.0054	0.0021	2.4×10^{-5}
95th%	1.3	0.89	0.72	0.17	0.067	0.012	0.0058	0.0012
Noncoincident	1.3	0.88	0.71	0.16	0.064	0.0089	0.0034	4.3×10^{-5}
99th%	2.0	1.3	1.1	0.24	0.096	0.016	0.0082	0.0016
Noncoincident	2.0	1.3	1.1	0.23	0.091	0.013	0.0048	6.4×10^{-5}

The modified fishing test behaves very much the same as the basic test except the σ and percentiles are reduced. For example, σ is reduced from 1 to near 0.8. This represents an increase in sensitivity. So instead of a value of 3, which was used in Baker (1992), the value of 2.4 (about 3σ) is used as the threshold for rejection of the null hypothesis of homogeneous data. The 99th percentile is generally below 2.4 for the modified test and is well below 2.4 for the small (cm) scales of interest. Thus, this lower rejection criterion still represents better than 99% significance. For the sake of brevity in the remainder of this manuscript, the modified fishing test will be referred to as the fishing test.

The results presented in the previous paragraph were found to be independent of the synthetic data parameters (i.e., C_{ave} , and L_{max}) by testing at various values of those parameters, including specifically cases A–F of Table 1. Only the percentiles were found to vary slightly in case C where they increased slightly, presumably because of the low number of total droplets in that case.

b. The power spectrum

The power spectrum is the square of the Fourier transform, implemented in this study via the fast Fourier transform (Cooley and Tukey 1965) applied to $X(d)$, the series of counts per interval at the highest resolution possible (10 μm in our model). Dividing by the variance of the data series normalizes the power spectrum. Under the null hypothesis of droplets randomly placed with equal probability everywhere, this results in an expected value of 1 at all scales L . Here L refers to the wavelengths of the Fourier components rather than the size of the distance bins as in the sections above. However, we use the same notation since in each case L refers to the length scale at which the particular technique is being calculated. The power spectrum is very noisy and therefore is smoothed (averaged) using a variable width

window. The averaging window size Nw_{ps} varies depending on the value of L as

$$Nw_{\text{ps}} = 2 \left[\text{floor} \left(\sqrt{\frac{L_{\text{max}}}{L}} \right) - 1 \right] + 1, \quad (3)$$

where floor represents rounding down to the nearest integer and L_{max} is the length of the data series. This choice, while arbitrary, provides maximum resolution where the data points are far apart and greater smoothing where the data points are denser and thus noisier.

Figures 6 and 7 show the average result of applying the variance normalized power spectrum to 10 000 time series of synthetic homogeneous data, each 60 000 time




FIG. 6. The mean (gray line at 1), σ (dashed black line also at 1), and 95th and 99th percentiles (black lines at 3 and 4.6, respectively) averaged over 10 000 normalized but not smoothed power spectra of random (Poisson distributed white noise) time series of length 0.6 m and concentration 10 cm^{-1} .

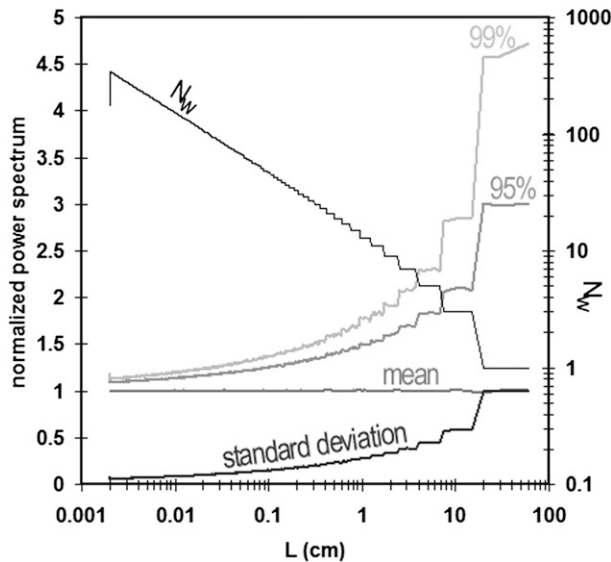


FIG. 7. The mean, σ , and 95th and 99th percentiles averaged over 10 000 normalized and smoothed power spectra of random (Poisson distributed white noise) time series of length 0.6 m and concentration 10 cm^{-1} . Also shown using the right vertical axis, as the thin black line angling from upper left to lower right, is the width (number of data points) of the averaging window used to smooth the power spectra (N_{wps}).

steps long (0.6 m at $10\text{-}\mu\text{m}$ resolution; i.e., $L_{\max} = 0.6$ m or 60 000 ticks) with a mean concentration of 0.01 events per time step ($C_{\text{ave}} = 10 \text{ cm}^{-1}$). The mean, σ , and 95th and 99th percentiles are shown in Figs. 6 and 7. Figure 6 shows the raw (normalized but not smoothed) power spectrum results; Fig. 7 shows the smoothed power spectrum results. The data displayed in Fig. 6 were also averaged (smoothed) using the same variable width window as was used for the power spectrum smoothing (in Figs. 7 and 8), as well as being an average of 10 000 realizations, since without smoothing these data are quite noisy, like the power spectrum itself.

For the raw power spectrum, the standard deviation is the same as the mean (i.e., 1). The 95th percentile is 3, and the 99th percentile is 4.6 for all L . This scale independence is a desirable characteristic in a statistic. However, as exemplified in Fig. 8, the raw power spectrum is too noisy, with many occurrences of values greatly exceeding 3σ . This is expected, as there are 25 000 independent data points. The smoothed power spectrum, while having scale dependent statistics, is well behaved relative to the 3σ curve (Fig. 8). The standard deviation of the smoothed power spectrum equals $\sqrt{N_{wps}^{-1}}$. For consistency with the fishing test, though with slightly less significance, 3σ is used as the threshold to indicate when the power spectrum deviates from the null hypothesis of droplets randomly placed with equal probability everywhere.




FIG. 8. Normalized power spectra, both raw (gray) and smoothed (black), for a homogeneous (white noise) time series of length 2.4 m (240 000 ticks at $10\text{-}\mu\text{m}$ resolution) with mean concentration of 10 cm^{-1} . The 3σ curves for the raw (dashed) and smoothed (dotted) cases are also displayed.

Within the range of values tested, which includes cases A through G of Table 1, the statistics of the power spectrum, described above, are independent of C_{ave} and L_{\max} and also independent of whether the data are ideal or noncoincident.

c. The pair correlation

We implement the pair correlation function [Eq. (4)] following Eq. (12) in Shaw et al. (2002). This expression of the pair correlation reveals its close relationship with the autocorrelation as discussed in Shaw et al. (2002), which in turn demonstrates its relationship to the power spectrum, since the autocorrelation is the Fourier transform of the power spectrum and visa versa:

$$\eta(L) = \frac{\sum_{d=1}^{L_{\max}} X(d)X(d+L)}{L_{\max}\bar{X}^2} - 1, \quad (4)$$

where periodic boundary conditions are used for $X(d)$, the counts per distance bin, at its highest resolution ($10 \mu\text{m}$ in our models); \bar{X} refers to its mean. Note that L again refers to the distance scale of the technique, but here it is a lag, or shift, as opposed to a Fourier component's wavelength or distance-bin size, as in the power spectrum and fishing test, respectively.

We determine the statistics of the pair correlation function in the same way as described above for the power spectrum and the fishing statistic, by generating homogeneous synthetic data and compiling the statistics. Like the power spectrum, the pair correlation is noisy where the data points become dense. Therefore, the pair correlation function is smoothed (averaged) using essentially




FIG. 9. Statistics (mean, σ , and 95th and 99th percentiles) compiled over 1000 pair correlation functions (raw not smoothed) applied to model homogeneous data as described in the text.

the same filter as for the power spectrum, with the averaging window size varying as

$$Nw_{pc} = 2 \left[\text{floor} \left(\sqrt{\frac{L}{L_{min}}} \right) - 1 \right] + 1, \quad (5)$$

where L_{min} is the sampling resolution, which is modeled as $10 \mu\text{m}$ in these examples.

Analogous to Figs. 6 and 7, which show statistics of the raw and smoothed power spectrum, Figs. 9 and 10 show the mean, σ , and 95th and 99th percentiles of 1000 pair correlations applied to synthetic homogeneous data, each 60 000 time steps long (0.6 m at $10\text{-}\mu\text{m}$ resolution; i.e., $L_{max} = 0.6$ m or 60 000 ticks), with a mean concentration of 0.01 events per time step ($C_{ave} = 10 \text{ cm}^{-1}$), for both the raw (Fig. 9) and smoothed (Fig. 10) pair correlation functions. The size of the averaging (smoothing) window Nw_{pc} is also shown in Fig. 10. The mean is zero for both the raw and smoothed pair correlation functions, as expected for random data. The σ and the 95th and 99th percentiles for the raw pair correlation, in this case, are about 0.41, 0.72, and 1.1, respectively. At this concentration, there is little difference between ideal and noncoincident data. These statistics for the pair correlation are independent of L but vary with C_{ave} and L_{max} , as shown in Table 1, for a wide range of the $C_{ave}\text{-}L_{max}$ parameter space. For each case shown in Table 1, the synthetic data are homogeneous (both ideal and noncoincident) and the results for the raw (not smoothed or averaged) pair correlation are shown. The main results of these variations, summarized in Table 1 are (i) the mean of the pair correlation for this random homogeneous data is always zero, (ii) σ of the ideally




FIG. 10. Statistics (mean, σ , and 95th and 99th percentiles) compiled over 1000 smoothed pair correlation functions applied to model homogeneous data as described in the text. The thin black line running in steps from the lower left to the upper right is the width (Nw_{pc} , right axis) of the averaging window used to smooth the pair correlation function.

modeled data is equal to $(C_{ave} \sqrt{L_{max}})^{-1}$, (iii) the 95th percentile is always less than two standard deviations and the 99th percentile is always less than three standard deviations, and (iv) the noncoincident data vary only slightly from the ideal data when C_{ave} is small but deviate as C_{ave} increases. Therefore, as long as the chance of coincidences is low, we can use 3σ with greater than 99% significance as the threshold for rejecting the null hypothesis that droplets are randomly placed with equal probability everywhere. For the smoothed (averaged) pair correlation, σ decreases as $\sqrt{Nw_{pc}}$, so the threshold is given as $3/(C_{ave} \sqrt{Nw_{pc} L_{max}})$.

Figure 11 shows the pair correlation function, raw and smoothed, applied to a single model run of homogeneous data. This example shows the large variations of the raw pair correlation that occasionally exceed 3σ . The smoothed pair correlation remains below 3σ .

5. Evaluation of the various droplet clustering tests

The statistics of the fishing statistic, pair correlation function, and power spectrum tests have been determined, allowing with high confidence the rejection of the null hypothesis when a signal clearly exceeds 3σ for any of the tests. The determination of these statistics opens the way to evaluate the relative merits of each test by comparing the response of each test to various synthetic inhomogeneous data.

a. Periodic structure

Figure 12 shows the results of applying all three tests to the same synthetic data series and is typical of what is




FIG. 11. Pair correlations, both raw (gray) and smoothed (black), for a homogeneous (white noise) time series of length 0.4 m (40 000 ticks at $10\text{-}\mu\text{m}$ resolution) with mean concentration of 10 cm^{-1} . The 3σ curves for the raw (dashed) and smoothed (dotted) cases are also displayed.

found by repeated random realizations. The synthetic data simulates alternating blocks, each 1 cm long and homogenous, with lower, then higher concentrations, 0.7 and 1.3 cm^{-1} respectively. The total length of the time series is 2.4 m (i.e., there are 120 pairs of low and high concentration blocks).

The fishing test results are as expected; a clearly significant peak at about half the block size (0.5 cm) that drops off steeply toward larger L and slowly toward smaller L . The power spectrum shows a very significant peak at 2 cm and another smaller peak at approximately 4 cm. The main peak at 2 cm is the expected signal since that is the repetition length of the square wave. The pair correlation shows an oscillation with many peaks exceeding 3σ starting at 2 cm, as expected, since the power spectrum has a sharp localized peak and the pair correlation is essentially its Fourier transform. Thus, these simulations show that for periodic structure, the pair correlation test is less localized than the fishing and power spectrum tests.

Additional runs with decreasing concentration differences confirm what seems apparent in Fig. 12, namely that the power spectrum is considerably more sensitive than the fishing test and the pair correlation to this type of structure.

b. Random vortices

The periodic model data simulations presented in the previous section are useful but the results could be misleading. The power spectrum is specifically good at detecting periodicity; however, small-scale structure in real clouds is not expected to manifest as periodic. Therefore, another model was constructed that more closely simulates what the inhomogeneous structure is

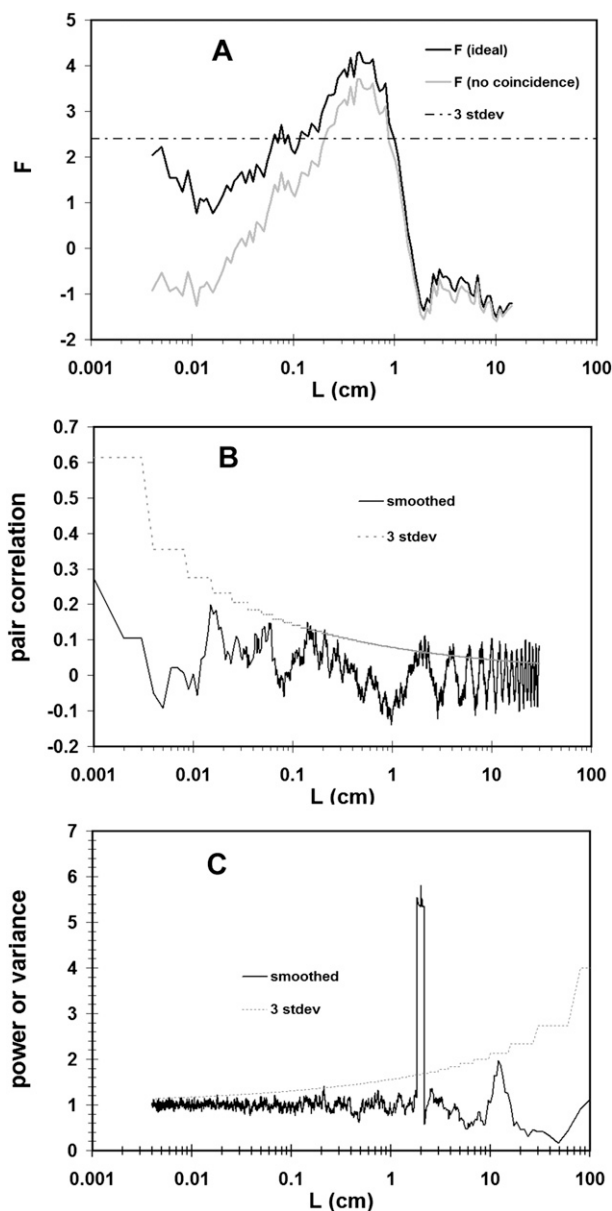


FIG. 12. The results of applying the (a) fishing statistic, (b) pair correlation, and (c) power spectrum to a series of modeled droplet data 2.4 m long ($L_{\text{max}} = 240\,000$ ticks) with alternating blocks, each 1 cm (1000 ticks) long, of higher ($1.3C_{\text{ave}}$) and lower ($0.7C_{\text{ave}}$) concentration, where $C_{\text{ave}} = 10\text{ cm}^{-1}$ (0.01 per tick). Both ideal and noncoincident data are shown for the fishing statistic, whereas for the power spectrum and pair correlation the difference between ideal and noncoincidence is negligible at this concentration.

expected to look like if it is caused by small-scale vortices, as suggested in the literature (Shaw et al. 1998).

The vortex model simulates homogenous concentration everywhere, except that at random locations a vortex-like structure is inserted. The vortex-like structures have an inner block of one-half the structure's total size, with

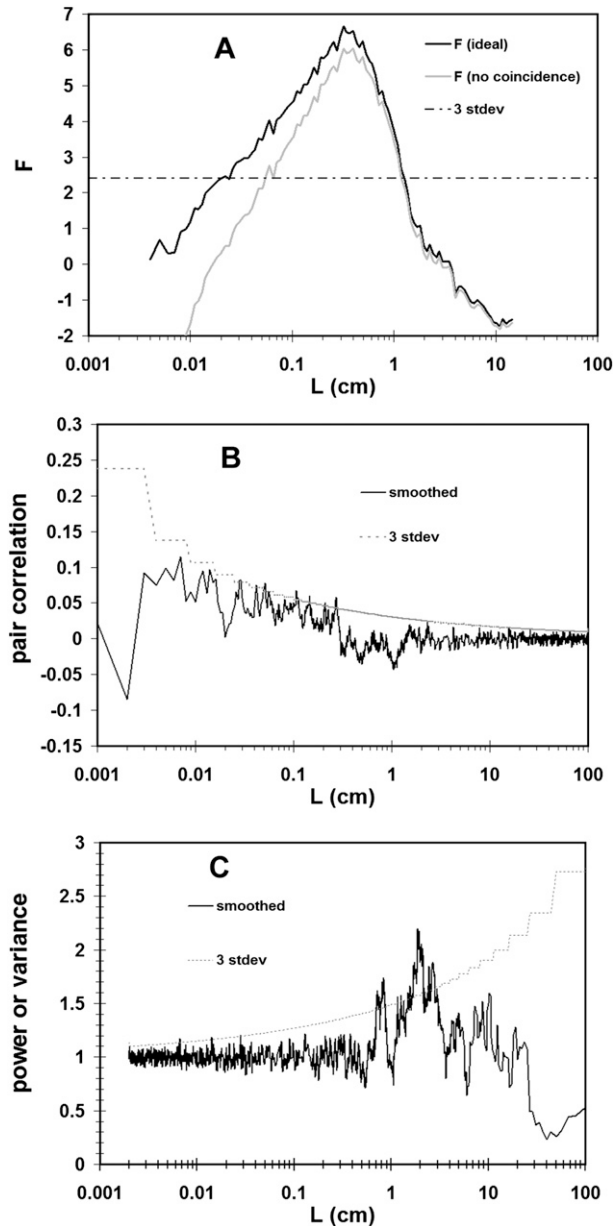


FIG. 13. The results of applying the (a) fishing statistic, (b) pair correlation, and (c) power spectrum to a series of modeled droplet data 4 m long ($L_{\max} = 400\,000$ ticks), $C_{\text{ave}} = 2\text{ cm}^{-1}$, with randomly placed vortex structures as described in the text and shown in Fig. 14. Both ideal and noncoincident data are shown for the fishing statistic, whereas for the power spectrum and pair correlation the difference between ideal and noncoincidence is negligible.

reduced concentration (RC_{ave}) and outer blocks on each side of the inner block of size one-quarter the structure's total size. The outer blocks each have an enhanced concentration $[(2 - R)C_{\text{ave}}]$, where R can vary between 0 and 1].




FIG. 14. Synthetic droplet spacing data converted to concentration (counts per distance) time series. (a) The actual concentration modeled. (b) A synthesized data sequence of simulated measured concentrations assuming random droplet placements given the true concentrations in (a).

Figure 13 shows the results of applying all three tests to a time series produced using the vortex model. The time series is 4 m long ($L_{\max} = 400\,000$) with $C_{\text{ave}} = 2\text{ cm}^{-1}$. The average distance between the randomly placed structures is 25 cm and their size is 2 cm (2000 ticks) with $R = 0.25$. Figure 14 shows the first 1.5 m of the time series, showing both the ideal concentration being modeled and the random realization, at 0.5-cm resolution. The fishing test and power spectrum both indicate structure at scales consistent with their characteristics determined via the periodic model. That is, the power spectrum peaks near the size of the vortex structure (2 cm) and the fishing statistic at near one-half the block size (i.e., between 0.25 and 0.5 cm). In this case, there is poor indication of the structure via the pair correlation function. This may result from the information being spread over the domain, as it must be since the information is still fairly localized for the power spectrum and the two are related by the Fourier transform. Further runs indicate that even without periodicity, the power spectrum is about as sensitive as the fishing statistic to these model data and the pair correlation is much less sensitive.

c. Effect of large-scale structure

To investigate the effect of large-scale structure on the various tests, the random vortex model was further modified to include a large-scale structure. The large-scale structure is imposed by making the average concentration of the first half of the data series lower than the average concentration of the second half. That is, the concentration of the first half equals SC_{ave} , while the concentration of the second half equals $(2 - S)C_{\text{ave}}$.

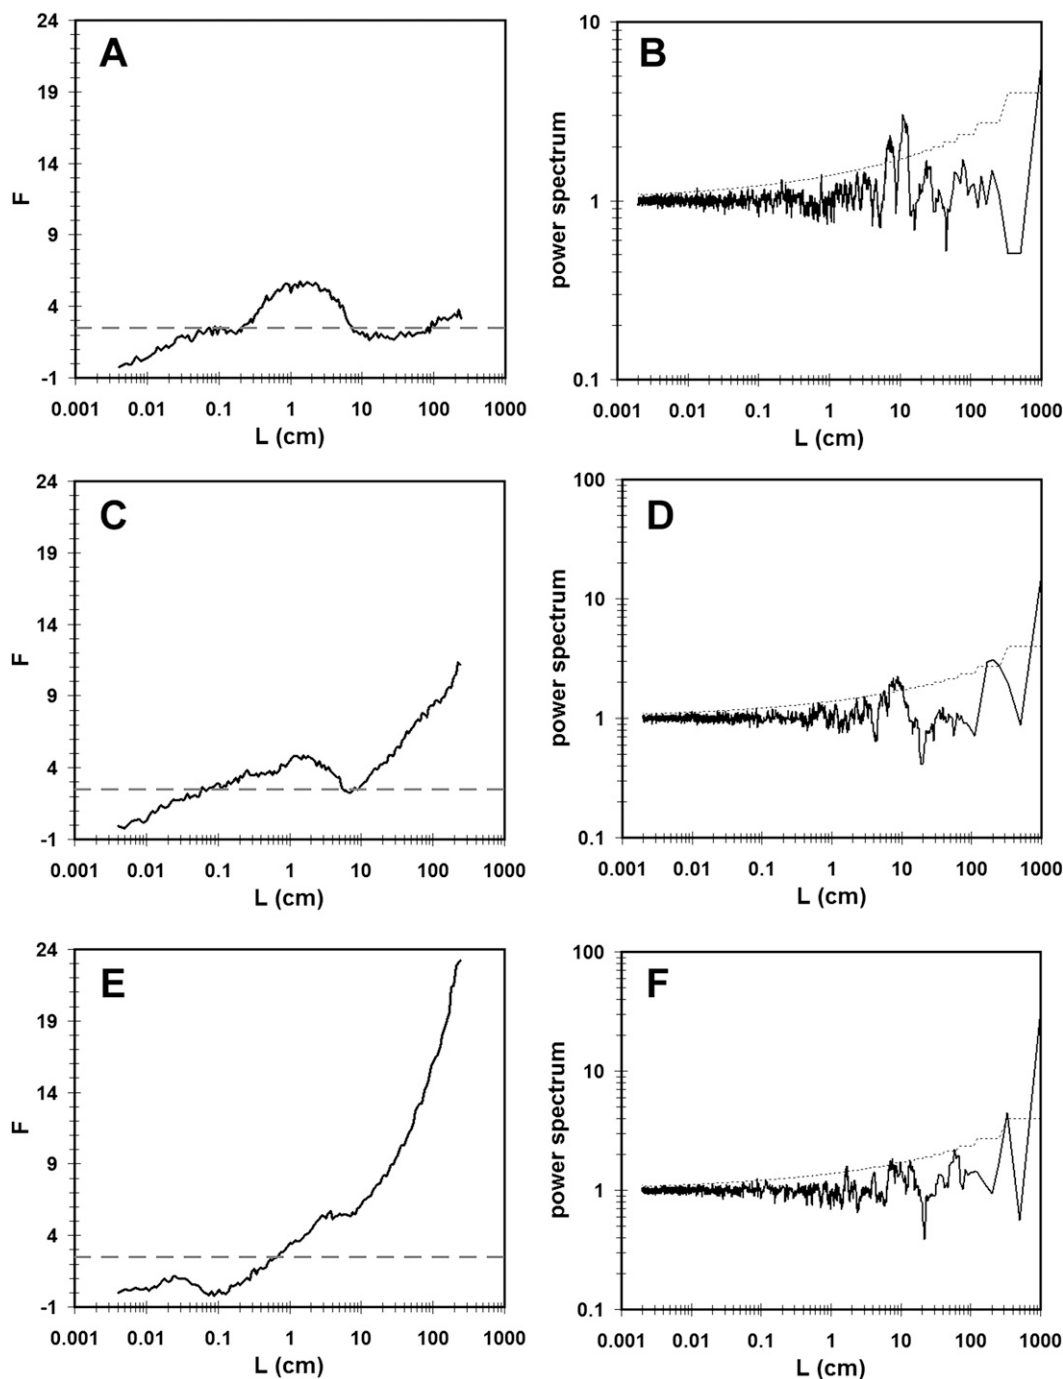


FIG. 15. Comparisons of the fishing statistic and the power spectrum applied to model data that contain both random vortices and large-scale structure as described in the text. The strength of the large-scale structure is (a),(b) 10% ($S = 0.1$), (c),(d) 15% ($S = 0.15$), and (e),(f) 20% ($S = 0.2$). Results from noncoincident data are shown but ideal data results are not significantly different at the concentrations modeled.

Figure 15 shows the results of applying the fishing and power spectrum tests to three such synthetic data series (with $S = 0.90, 0.85,$ and 0.80). Except for varying S , the model parameters are fixed. The length is 10 m, the

average distance between vortices is 90 cm, and the vortex size is 9 cm.

The pair correlation for these model runs is not shown since, as shown in Fig. 13, the vortex-like structures are

not well revealed even in the absence of large-scale structure. For $S = 0.90$, the large-scale structure is indicated by both the fishing statistic and power spectrum but does not interfere with the detection of the small-scale structure. For $S = 0.85$, the large-scale structure is nearly dominant while the small-scale structure remains just barely detected, for both the fishing and power spectrum tests. For $S = 0.80$, the large-scale structure obscures the detection of the small-scale structure. One can see the effect of the small-scale structure using both tests, but it is unlikely that this structure could be detected without a priori knowledge of its existence. These simulations suggest that the fishing statistic and the power spectrum are roughly equivalent in their ability to detect small-scale structures in the presence of a superimposed large-scale structure.

6. Summary, discussion, and conclusions

Homogeneous synthetic data were used to determine the standard deviations and confidence intervals of the fishing statistic, the power spectrum, and the pair correlation function under the null hypothesis of droplets randomly placed with equal probability everywhere. Synthetic data containing imposed structure were used to evaluate the usefulness of these functions as hypothesis testing statistics. The three functions were examined to determine their respective abilities to reject the null hypothesis and detect the scale(s) of the nonrandom structure. The results are presented within the context of droplets observed in clouds via a rapidly moving aircraft, since this type of simulation is applicable to typical in situ investigations of clouds. However, this situation may be generalized to any sequence of data that are exponentially distributed, or equivalently, for which the number of events per fixed interval is Poisson distributed, under the null hypothesis. For example, the pair correlation was used to search for structure in the spacing of raindrops by Larsen et al. (2005), and similarly by Larsen (2007) to test the spacing of aerosols. However, any process that is expected, or suspected, to be random in time or space could be tested with the tools described herein.

Three types of structure were modeled: a periodic square wave, random vortices, and random vortices with a superimposed large-scale jump. Although mathematically related, and thus containing the same information, the three tests differ dramatically in the presentation of information. For these types of structure, the pair correlation is the least sensitive and yields the least spatial information about the structure. For the periodic structure, the power spectrum is more sensitive than the fishing statistic. For the two models with random vortices, the

fishing statistic and the power spectrum are roughly equivalent in sensitivity, scale resolution, and detection of small-scale structures despite a superimposed large-scale structure, although the presentation of information is quite different for the two tests. A practical approach is to use both the fishing statistic and the power spectrum to analyze the possible effects of droplet clustering that may resemble the type of structures modeled herein. We do not suggest abandoning the pair correlation function entirely. Instead, the pair correlation test should be tried as well, in the event that real structures are sufficiently different than those modeled herein and are of such form that the pair correlation is the more useful tool.

The clustering index and the volume-averaged pair correlation are shown to be equivalent to the fishing statistic but with different normalizations. The normalization of the fishing statistic makes it the most practical choice of the three for studying clustering in droplet spacing data. The analysis presented here suggests that the technique described by Marshak et al. (2005) is not applicable to the investigation of clustering.

Acknowledgments. This work was supported under NSF Grant ATM 0342486. We thank Jean-Louis Brenguier and several anonymous reviewers for their insightful comments that helped improve our presentation of this work.

REFERENCES

- Baker, B. A., 1992: Turbulent entrainment and mixing in clouds: A new observational approach. *J. Atmos. Sci.*, **49**, 387–404.
- Chaumat, L., and J. L. Brenguier, 2001: Droplet spectra broadening in cumulus clouds. Part II: Microscale droplet concentration heterogeneities. *J. Atmos. Sci.*, **58**, 642–654.
- Cooley, J. W., and J. W. Tukey, 1965: An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, **19**, 297–301.
- Grabowski, W. W., and P. Vaillancourt, 1999: Comments on the “Preferential concentration of cloud droplets by turbulence: Effects on the early evolution of cumulus cloud droplet spectra.” *J. Atmos. Sci.*, **56**, 1433–1436.
- Knyazikhin, Y., A. Marshak, M. L. Larsen, W. J. Wiscombe, J. V. Martonchik, and R. B. Myneni, 2005: Small-scale drop size variability: Impact on estimation of cloud optical properties. *J. Atmos. Sci.*, **62**, 2555–2567.
- Kostinski, A. B., and R. A. Shaw, 2001: Scale-dependent droplet clustering in turbulent clouds. *J. Fluid Mech.*, **434**, 389–398.
- Larsen, M. L., 2007: Spatial distributions of aerosol particles: Investigation of the Poisson assumption. *J. Aerosol Sci.*, **38**, 807–822.
- , A. B. Kostinski, and A. Tokay, 2005: Observations and analysis of uncorrelated rain. *J. Atmos. Sci.*, **62**, 4071–4083.
- Lehmann, K., H. Siebert, M. Wendisch, and R. A. Shaw, 2007: Evidence for inertial droplet clustering in weakly turbulent clouds. *Tellus*, **59B**, 57–65.

- Marshak, A., Y. Knyazikhin, M. Larsen, and W. J. Wiscombe, 2005: Small-scale drop-size variability: Empirical models for drop-size-dependent clustering in clouds. *J. Atmos. Sci.*, **62**, 551–558.
- Pinsky, M., and A. P. Khain, 2001: Fine structure of cloud droplet concentration as seen from the Fast-FSSP measurements. Part I: Method of analysis and preliminary results. *J. Appl. Meteor.*, **40**, 1515–1537.
- Saw, E. W., R. A. Shaw, S. Ayyalasomayajula, P. Y. Chuang, and A. Gylfason, 2008: Inertial clustering of particles in high-Reynolds-number turbulence. *Phys. Rev. Lett.*, **100**, 214501, doi:10.1103/PhysRevLett.100.214501.
- Shaw, R. A., W. C. Reade, L. R. Collins, and J. Verlinde, 1998: Preferential concentration of cloud droplets by turbulence: Effects on the early evolution of cumulus cloud droplet spectra. *J. Atmos. Sci.*, **55**, 1965–1976.
- , A. B. Kostinski, and M. L. Larsen, 2002: Towards quantifying droplet clustering in clouds. *Quart. J. Roy. Meteor. Soc.*, **128**, 1043–1057.