

On Measuring the Degree of Irregularity in an Observing Network

CHARLES A. DOSWELL III

National Severe Storms Laboratory, Norman, Oklahoma

SONIA LASHER-TRAPP

School of Meteorology, University of Oklahoma, Norman, Oklahoma

(Manuscript received 19 September 1995, in final form 5 June 1996)

ABSTRACT

Meteorological observing networks are nearly always irregularly distributed in space. This irregularity generally has an adverse impact on objective analysis and must be accounted for when designing an analysis scheme. Unfortunately, there has been no completely satisfactory measure of the degree of irregularity, which is of particular significance when designing artificial sampling networks for empirical studies of the impact of this spatial distribution irregularity. The authors propose a measure of the irregularity of sampling point distributions based on the gradient of the sums of the weights used in an objective analysis. Two alternatives that have been proposed, the fractal dimension and a "nonuniformity ratio," are examined as candidate measures, but the new method presented here is considered superior to these because it can be used to create a spatial "map" that illustrates the spatial structure of the irregularities in a sampling network, as well as to assign a single number to the network as a whole. Testing the new measure with uniform and artificial networks shows that this parameter seems to exhibit the desired properties. When tested with the United States surface and upper-air networks, the parameter provides quantitative information showing that the surface network is much more irregular than the rawinsonde network. It is shown that artificial networks can be created that duplicate the characteristics of the surface and rawinsonde networks; in the case of the surface network, however, a *declustered* version of the observation site distribution is required.

1. Introduction

As noted in Koch et al. (1983), Smith et al. (1986), and Barnes (1994), the *degree* of irregularity in an observational array's distribution can have a large impact on the way an objective analysis (OA) is done and how successful it is likely to be. In fact, Buzzi et al. (1991) have developed a method to minimize the negative impact of irregularity in the spatial sampling. Empirical tests of OA schemes often are conducted to support the choices of the OA method and any of its associated parameters. These empirical tests usually make use of an analytic input function with which to compare the analyzed values. On one hand, it is logical to sample the analytic function with the *actual* station distribution (e.g., Smith et al. 1986). When doing this, however, there is some degree of uncertainty regarding the generality of the results; the given results might depend to some unknown extent on the specific station distribution under consideration and its position in relation to the

analytic function. If, instead, an *artificial* network is used in empirical tests (e.g., Barnes 1994, hereinafter B94), it becomes possible to remove the effects of a particular realization by performing the tests on a number of different, but statistically similar, station distributions. The problem with this latter approach is that there has not been a simple way to compare the artificial and real networks. In other words, there has been no common measure of irregularity between the two, such that it can be said with confidence that the artificial network is "similar" in some sense to the real network. The objective of the present study is to find a measure of irregularity that allows the comparison of such artificial distributions with real data arrays.

Two measures of the degree of irregularity proposed by other authors are investigated and found inadequate, for reasons described in section 2. A new measure of irregularity is proposed in section 3, based on an idea presented in the appendix of Doswell and Caracena (1988, hereinafter DC88), and various tests of the proposed measure are presented. Section 4 contains two practical examples of using the measure, using the U.S. surface and rawinsonde networks pictured in Fig. 1, and section 5 concludes with a summary of the results of this work and additional topics for future research.

Corresponding author address: Dr. Charles A. Doswell III, NOAA/ERL/National Severe Storms Laboratory, 1313 Halley Circle, Norman, OK 73069.
E-mail: doswell@nssl.uoknor.edu

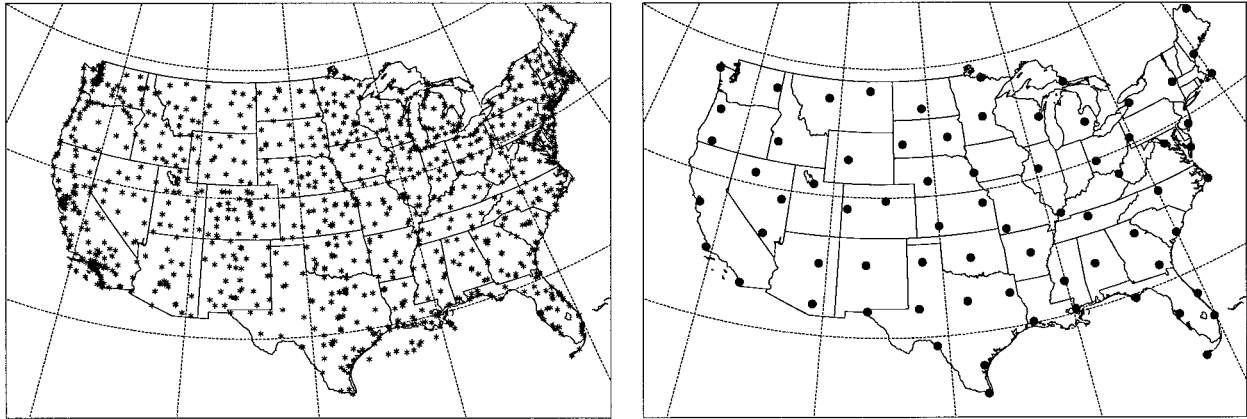


FIG. 1. (a) The U.S. surface observation network sites and (b) the U.S. upper-air network sites, as of fall 1993.

2. Previously proposed measures of irregularity

a. The fractal dimension

Lovejoy et al. (1986) have proposed using the fractal dimension to characterize the distribution of a geophysical data array. When considering a station distribution in a two-dimensional embedding space (as on the surface of the earth), the fractal dimension (denoted D_m) should be two for a uniform distribution of data points. Real, irregular data distributions (Fig. 1) should have a fractal dimension between zero and two, with the degree of inhomogeneity being measured by $2 - D_m$. The correlation dimension D_c is often used as an approximation for the fractal dimension because it is easier to calculate (Grassberger and Procaccia 1983; Lovejoy et al. 1986; Korvin et al. 1990). Determining the correlation dimension consists of counting the number of stations n within a series of circles of increasing radii r around each point in the observation lattice, so that $n = n(r)$.¹ We followed the recommendation of Korvin et al. (1990), who noted that r should not exceed one-third of the *largest* interstation distance to produce a reliable estimate of D_c . By finding the average of $n(r)$ over all the stations (avoiding counting the same distance twice), denoted $\langle n(r) \rangle$, a plot of $\ln \langle n(r) \rangle$ versus $\ln r$ can be created. The correlation dimension D_c is the slope of a line fitted to the data on such a plot. Lovejoy et al. (1986) used the correlation dimension method to find a fractal dimension of approximately 1.75 for the World Meteorological Organization (WMO) network and presented this measure as a guide in determining detectability limits. Because the fractal dimension addresses

the inhomogeneity of the network, we investigated it as a candidate measure of irregularity to compare simulated and real networks.

We applied the correlation dimension technique to fictitious and real station distributions, including the surface and upper-air stations of the contiguous United States. When considering stations near the boundaries of the finite area data lattice, we obtain values of $n(r)$ that differ significantly from those within the interior of the data domain. This property of finite data domains is well known: Barnes (1964), Achtemeier (1986), DC88, and Pauley (1990) all recognized that data lattice boundaries create difficulties. The standard approach (although by no means the only one) is to erect what Cressie (1991, p. 607) called “guard areas” inside the perimeter of the data lattice. In other words, one only considers information from stations within the *interior* of the data lattice; one chooses a guard *barrier* (a term we have used in preference to Cressie’s term, *area*) such that the results near the edge of the guard barrier are indistinguishable from those deeper within the data lattice.

Having erected a guard barrier near the edges of the data lattice, another problem arises, however. Fitting a straight line to the points in the plot using least squares is a straightforward procedure, but problems arise when deciding *which* points to use in the fitting process if the entire profile is not linear. Results of this process are shown in Fig. 2 for both the surface and upper-air networks. For the upper-air network, two different parts of the curve appear linear, yielding very different fractal dimensions of 1.97 and 4.46. The concept of a network having different fractal dimensions over different scales is not new (e.g., Tessier et al. 1994) but complicates the use of the fractal dimension for comparing the irregularity of station distributions. Furthermore, the fractal dimension in either section of the plot can be changed by making small changes in which points to consider in the line-fitting. Using the last 12 data points at the top of the plot instead of the last 30 yields a fractal

¹ The radius r has been nondimensionalized for our calculations. The value of 190.5 km used for the nondimensionalization is an artifact of the algorithm we have used to convert latitude–longitude locations to x – y coordinates on a map projection. That algorithm is based the Limited-Area, Fine-Mesh Model (LFM) grid spacing of 190.5 km at the standard latitude used in the polar stereographic projection, expressing coordinates in grid units.

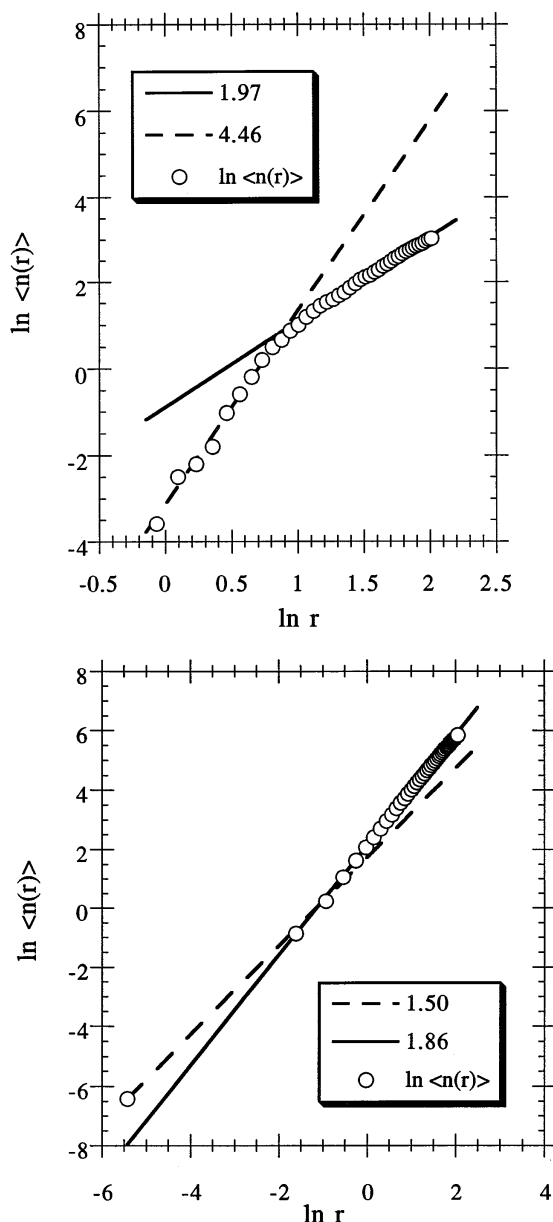


FIG. 2. Results of the fractal dimension method for (a) the U.S. upper-air and (b) the U.S. surface observation networks pictured in Fig. 1. The slopes of the fitted lines are indicated in the legend boxes.

dimension of 1.72 instead of 1.97. The fractal dimension of the surface network (Fig. 2b) also shows some indication of multiple fractal dimensions (1.50 or 1.86, again depending on which points are chosen for the fitted line). The uncertainty in the fractal dimension associated with the choice of points for the line-fitting is as large as or larger than the difference between the upper-air and surface networks. Considering this point and the possible uncertainty about which fractal dimension to use, we conclude that the subjectivity associated with this measure is unacceptable for comparing irregularity among different spatial distributions.

b. The nonuniformity ratio

In B94, considerable use was made of a nonuniformity ratio ρ , proposed by Smith et al. (1986), which they defined as

$$\rho = \frac{E - M}{M},$$

where E is what Smith et al. call the “equivalent uniform station spacing” (defined as the spacing derived by distributing the original number of stations uniformly over the data domain) and M is the mean distance to each station’s nearest neighbor in the real array.² A uniform sampling array would have $\rho = 0$, and the greater the irregularity, the larger ρ would become. It certainly can be argued that our proposed measure is not substantially different from ρ . However, ρ is a single number intended to represent the nonuniformity of the data array as a whole. By using the proposed measure described in the following section, the irregularity can also be displayed over the domain, to provide a picture of how the data density varies in space. In our opinion, this conveys more information about the nonuniformity of data than does any single number.

An equally important point is the fact that ρ is independent of the OA scheme that will be used with the data, while the measure proposed in the next section can be “matched” to the OA scheme by making use of the same parameters. Thus, this measure truly assesses the impact that the irregularity has on the OA scheme’s results and can provide feedback on the appropriate choice of parameters to minimize the effect of the irregularity on the OA.

3. A new measure of irregularity

Two ideas have contributed to our proposed measure. First, Barnes (1964) showed a figure (his Fig. 4) displaying the number of stations influencing the analysis as a function of space. Barnes used a “radius of influence” in his OA scheme: Stations outside this radius were not considered in the analysis. Thus, the spatial constancy of the number of stations within this radius reflects, in a crude way, the uniformity of stations. Where that number is relatively constant (as in the center of the United States), the stations are relatively uniform. Barnes’s figure shows that the major contribution to nonuniformity of rawinsonde observations is that associated with data domain boundaries. Outside of the land area of the United States, the data density drops precipitously. There are clusters and voids within the

² As used in B94, the uniformity ratio is not quite the same as that defined by Smith et al. B94 uses the average of the six nearest neighbors, rather than the single nearest neighbor, when calculating M . This should provide a somewhat smoother result than that of Smith et al.

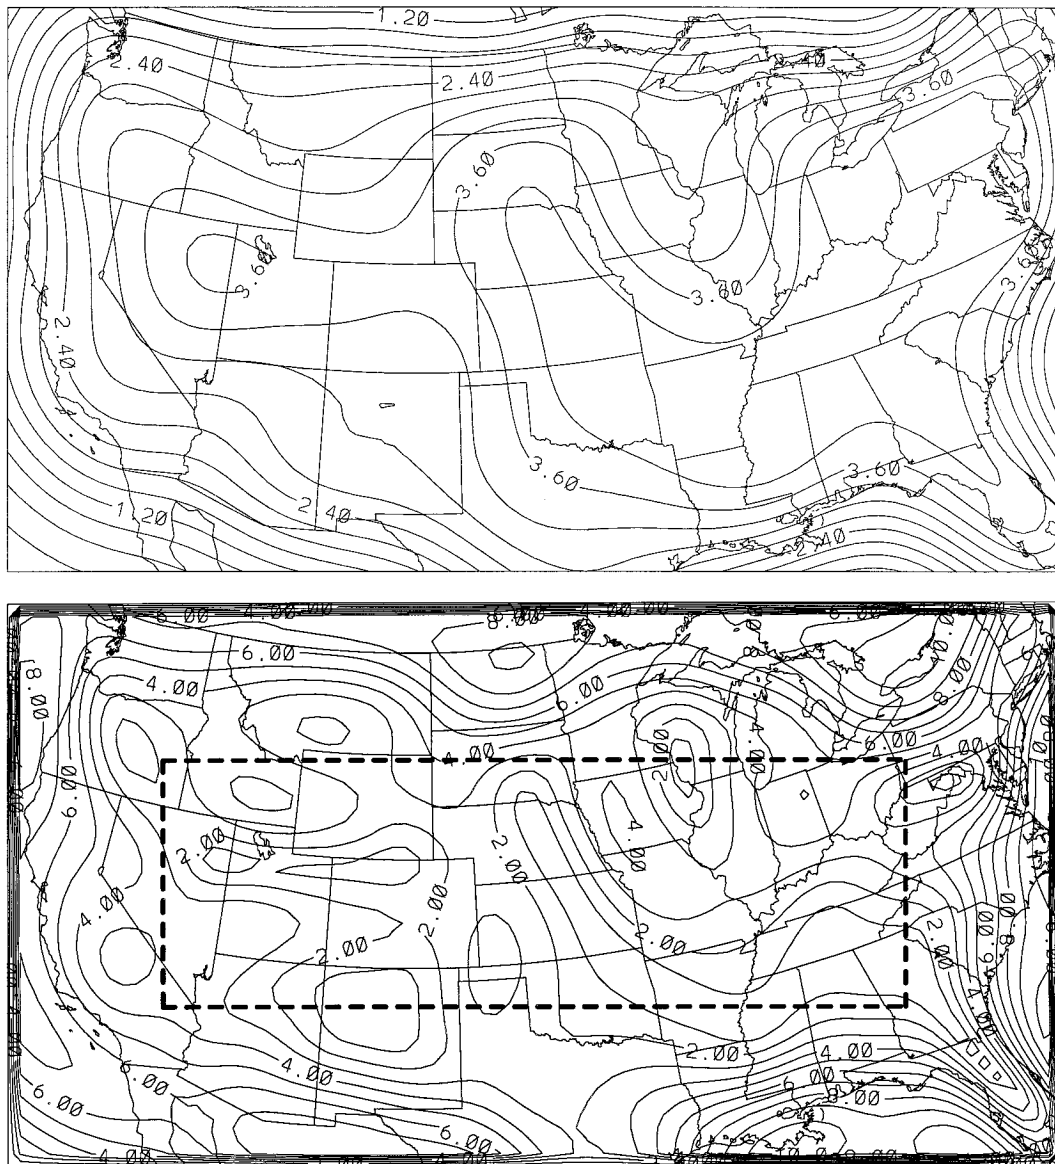


FIG. 3. Distributions of (a) the sum of the weights (nondimensional) and (b) the proposed measure μ (values shown are 10 times the value of a centered difference over two dimensionless grid intervals) based on Eq. (2) in the text, for the fall 1993 upper-air network. The dashed rectangle in (b) depicts the area within which the values shown in Table 2 were computed.

interior of the country also, and a finer contour interval would make this more obvious.

The other aspect of the idea was explored tentatively in DC88 in the appendix. Specifically, they showed that the gradient of the averaging function used in distance-dependent weighted averaging contains a term involving the gradient of the *normalizing factor*, which is simply the sum of the weights affecting any given grid point. Figure 3a shows that in regions of quasi-uniform data, the gradient of the sum of the weights affecting the analysis should be quite small; in regions of substantial irregularity, the gradient would be large and could affect the calculation of data gradients.

Therefore, we consider the magnitude of the gradient of the sum of the weights; that is,

$$\mu \equiv \left| \nabla \left[\sum_{k=1}^n w_k \right] \right|, \tag{1}$$

(where n is the number of stations considered and w_k is the weight assigned to the k th station at the analysis point in question) to be a candidate parameter for estimating the degree of irregularity in a station distribution. Although the selection of a weighting function

is a potentially troublesome issue, it should be clear that unless the selection is done poorly, many different functions all should give roughly comparable results. We have chosen to use the Gaussian weighting function proposed by Barnes (1964)

$$w_k = e^{-(R_k/\lambda)^2}, \quad (2)$$

(where R_k is the Euclidean distance from the analysis point to the k th data point and λ is the shaping parameter of the scheme) largely because of its convenience and familiarity. Determination of the *shaping parameter* λ is considered below.

The examples shown in this paper are all for a single-pass OA scheme, but the proposed method can be adapted for multipass schemes. For purposes of this paper (testing the proposed measure), we consider it sufficient to employ any particular OA scheme; the single-pass, Gaussian weighting has been chosen for convenience. We believe that if some other scheme is being used for OA, that scheme is the one to use for measuring the irregularity of the sampling network. Multipass OA techniques require calculating an inverse Fourier transformation on the known final response function of the multipass scheme, to find the single-pass weighting function equivalent to the multiple-pass scheme. Once that equivalent single-pass weighting function is known, (1) can be used to calculate the μ values as described in the following sections.

a. Some preliminary issues

Certain parameters must be set before calculating the values of μ described by (1), and unwise choices of these parameters may render the measure useless. Thus, we now describe the experimentation that has led to the choices we advocate.

1) GUARD BARRIER

The notion of a guard barrier has been introduced already, in the context of noting the effects of the data lattice boundary upon the results of the fractal dimension method. In the context of our method, we observe that the shaping parameter λ in (2) determines a length scale of importance: the e -folding distance for the weights. The parameter λ determines the “reach” of the weighting scheme; for example, the weighting scheme gives a weight less than 0.0183 for all points beyond a Euclidean distance of 2λ . This means that a sum of the weights will not “feel” the boundaries very much until it is within about 2λ – 3λ . If the guard barrier is chosen to be somewhere in this range, the average value of the sum of the weights will not be affected adversely by the data lattice boundaries. After some experimentation (Fig. 4) with a uniform square grid, which should yield $\mu = 0$, we have chosen a guard barrier of $4\Delta d$, equiv-




FIG. 4. The average μ versus λ for a uniform square grid (with $\Delta d = 1.0$) for three different values of the guard barrier distance: $2\Delta d$, $4\Delta d$, and $6\Delta d$.

alent to about 3λ (where Δd is the median *data* spacing). The choice of $4\Delta d$ is a compromise between guard barriers of $2\Delta d$ and $6\Delta d$: $4\Delta d$ (with $\lambda = 1.3$; see the next section for a discussion of how λ is chosen) gives a more accurate depiction of μ than the $2\Delta d$ case without sacrificing so much of the data domain as in the $6\Delta d$ case.

2) SHAPING PARAMETER

Given the foregoing experiments, it appears that by making λ small enough in the interiors of our theoretically “uniform” data grids, it is indeed possible to drive μ to quite low values. When λ is too large, the interior of the data domain still “feels” the data boundaries; however, it is not obvious that we would necessarily want to make λ extremely small, since that implies excessive weighting on values quite close to the analysis point. Ordinary OA considerations suggest that making λ too small gives an excessively “noisy” analysis. Our results (Fig. 4) show that when λ is too small, the μ values increase owing to spurious waves that appear in the field of the sums of the weights because of a Moiré-like effect. These results show that the smallest value of the average μ for the uniform grid occurs at $\lambda = 1.3$, which is 1.3 times the median data spacing (Δd). This value was endorsed by Pauley and Wu (1990) and is within the range of values advocated by Caracena et al. (1984). It is important to note that one should use the same value of λ in the irregularity measure as that used in one’s OA scheme. For examining the theoretically uniform square data lattices with unit spacing, we use $\lambda = 1.3$. As we have shown, the guard barrier that is suited best to a value of $\lambda = 1.3$ is $4\Delta d$; these choices make the interior values of μ sufficiently small for any practical purposes.

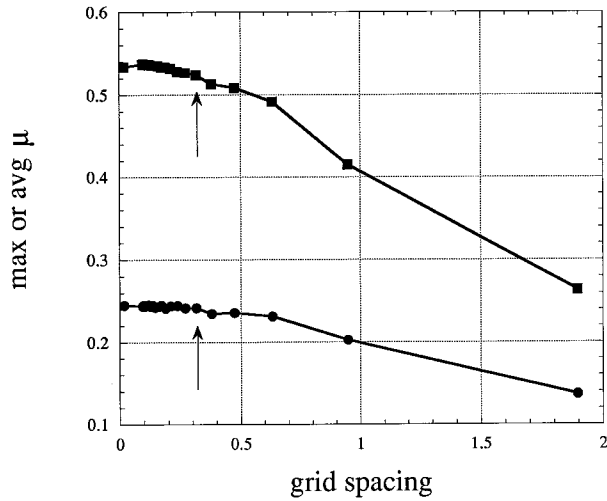


FIG. 5. Maximum (squares) and average (circles) values of μ for the upper-air network as a function of computational grid spacing. Plotted points correspond to particular values of $\Delta d/N$, where N ranges from 1 to 20 in unit steps, plus a last value (far left) at $N = 100$. Arrows indicate values at a grid spacing of $\Delta d/6$.

We have chosen not to use a “radius of influence” or “cutoff radius” in the analysis. Therefore, all data points are included in the sum of the weights at any point in the computational domain. Clearly, points far away contribute virtually nothing to the sum, which therefore will be dominated by the data distribution near any specific point in question. If a cutoff radius is used in one’s OA scheme, however, the same cutoff radius should be used when testing the irregularity, to keep it “matched” to the OA scheme.

3) COMPUTATIONAL GRID SPACING

In our calculation of μ using (1), the gradients are computed with second-order finite differences on a square computational grid, and the computational grid spacing has an effect on the size of the μ values. The maximum and average values of μ increase as the grid spacing decreases (Fig. 5); a smaller grid spacing is able to detect more of the real value of the magnitude of the gradients. Similar results were found for the surface network (not shown). Apparently, the true value of μ can be found only in the limit as the computational grid spacing approaches zero. To maximize the accuracy of μ values while keeping computational costs associated with a large number of grid points within bounds, we have chosen a computational grid spacing of $\Delta d/6$; most of the value of the gradient (99% in this example) is captured at this point (see Fig. 5). The $\Delta d/6$ criterion can be applied to most data distributions encountered in meteorology, unless it is obvious a priori that the distribution is pathologically irregular (large voids combined with intense clustering of sample points). This choice obviously is related to issues of resolution discussed in DC88.

TABLE 1. Displacement experiments and the corresponding average μ .

Experiment	x displacement	y displacement	Avg μ
Control	0.0	0.0	0.000005
1	0.5	0.5	0.000010
2	-0.3	0.2	0.000006
3	-0.05	-0.4	0.000007
4	0.2	0.0	0.000010
5	0.2	-0.5	0.000005

b. Tests with uniform data distributions

To conduct a “control” experiment with our method, we evaluate the maximum, minimum, and average μ values for a 27×15 uniform square grid (as an example of a fictitious, uniform data distribution). The computed average value of μ is slightly greater than zero (~ 0.00000517); this corresponds to the minimum plotted in Fig. 4 for $\lambda = 1.3$. For a uniform data distribution, μ should be zero, so this control experiment confirms this supposition, at least within the finite computational limits of real experiments. It will be shown in later experiments (see, e.g., Table 2) that values of μ for different networks are possible on the order of 1, so this μ value of 0.00000517 is indeed a very small number in comparison to the range of values possible and thus can be effectively considered to be zero. A test performed with twice the number of grid points in the uniform grid with the same unit grid spacing produces an average μ of approximately 0.00000487, which demonstrates only a modest dependence of μ on n .

We also have tested the effects of the computational grid having some specific spatial relationship to the data sampling array. These “displacement” experiments consist of shifting the sampling sites in relation to the computational grid and evaluating the effect on μ . Five different displacements of the sampling sites are shown in Table 1, along with the corresponding average μ values. Because the sampling sites are uniformly distributed, the average values of μ are expected to be zero as in the control case, and indeed they are very small, albeit with slight variation. The results of these experiments reveal that the average μ and, hence, our choice of λ and the guard barrier are not affected significantly by a displacement of the data points relative to the computational grid.

c. Tests with artificial irregular distributions

We can create increasingly irregular distributions in a manner comparable to B94 to test the applicability of our measure. The distributions start with a uniform square array of sampling sites with unit spacing, which are displaced according to

$$\begin{aligned} x &= x_0 \pm n_r D \\ y &= y_0 \pm n_r^* D, \end{aligned} \tag{3}$$

where x and y are the new locations of the data point originally located at (x_0, y_0) , n_r and n_r^* are pseudorandom numbers uniformly distributed between 0 and 1, and D is the *scatter distance*,³ the maximum amount the point can be moved in either of the x or y directions. For each grid point, four random numbers between zero and one are generated: The first is the *amount* the grid point is moved in the x direction and the second is the *sign* of that movement (a random number less than 0.5 means movement in the negative x direction); the third and fourth are the same except they apply to the y direction. The algorithm to generate random numbers is an adaptation of the method described by Press et al. (1986), which they assert to be free of sequential correlation. As D increases, so should the irregularity of the distribution, at least up to a “saturation” point (see below). Some examples of the artificial distributions are shown in Fig. 6. We have created 20 realizations for each size increment of D by starting the pseudorandom number generator with a different seed for each realization and have averaged the results over the set of 20 realizations to find typical results for D values of that magnitude. The scatter distance D is allowed to vary from 0.1 to 100 by increments of 1.0, except between 0.1 and 2.0, where the increment is 0.1.

In the process of our experimentation, it became clear that we need to decide how to deal with points that are scattered outside of the original data boundaries. Therefore, all of our experiments are done with three different “boundary conditions”: 1) “dispersive,” in which the data points are allowed to be scattered outside of the original data boundaries; 2) “reflective,” in which points that would have been scattered outside a boundary are reflected that same distance back inside the boundary; and 3) “periodic,” in which the data points are allowed to exit a boundary but reenter the domain at the opposite boundary, such that the point is as far inside the one boundary as it would have been outside the other boundary.

The results for the reflective and periodic boundaries tend to be very similar in most cases, but the dispersive case behaves differently, owing to a decrease in the number of points within the original boundaries. That the dispersive case would behave differently could have been anticipated just by looking at the three kinds of distributions at $D = 10$ (Fig. 7). In the dispersive case, the overall density of data within the original data domain boundaries decreases as D increases.

The maximum, minimum, and average μ values using each of the different boundary conditions (Fig. 8) show that by a $D \sim 1.5$, the irregularity has attained a maximum. This can be considered a sort of “saturation” of the irregularity; increasing D further simply moves

points around without materially affecting the irregularity of the distribution. Figure 10 in B94 shows basically the same result.⁴ Making $D > 1.5$ reveals no discernible trend in the reflective and periodic boundary cases; however, for the dispersive boundary case, the average μ starts to decrease again. This effect obviously is due to the decreasing number of points within the computational grid.

Thus, we have verified that μ increases when the irregularity increases. Given this initially satisfactory result, it is useful to evaluate how our artificial station distributions compare with those characterized by complete spatial randomness (i.e., exhibiting nearest-neighbor distributions described by a Poisson distribution, as detailed in the appendix). Using the nearest-neighbor distributions for $D = 0.1, 0.5, 1.0,$ and 5.0 (Fig. 9), it is clear that as D increases, the distribution approaches that of a Poisson random variable. Using the Pearson test (also described in the appendix), the distributions for $D = 0.1$ and 0.5 are rejected as being Poisson at the 0.01 significance level, but the distributions for $D = 1.0$ and 5.0 are accepted as good fits to the Poisson distribution at the 0.01 level. Thus, our method of creating “random” distributions proves to be quite comparable to true spatially random sampling for $D \geq 1.0$.

A final test addresses the dependence of μ on n , the number of data sampling sites in the distribution but this time using an irregular sampling distribution. Using $D = 0.5$, 20 different realizations of irregular distributions are created using twice the number of data points as used in the previous experiments. The average $\mu = 0.939$, which is only slightly different than the value found above with half the sampling sites (0.954). We conclude that μ does not depend strongly on n .

4. Two practical examples

This section describes how one can use the proposed measure to create artificial distributions with the same amount of irregularity as real data networks. Once the artificial distributions have been verified to be as irregular as the data network in question, those artificial distributions can be used to test how well an OA scheme responds to irregularities in the data sampling.

a. U.S. upper-air network

The appropriate λ for the upper-air network is computed as 1.3 times the median from the nearest-neighbor

³ Note that D is scaled to the grid spacing (Δ) of the initially uniform array of sampling sites to be perturbed so that, for example, a D value of 5.0 corresponds to $D = 5\Delta$.

⁴ What Barnes calls the “offset” is the same as our “scatter distance,” or D . It does appear in Barnes’s Fig. 10 that “saturation” occurred at an offset of 2, whereas our results suggest that saturation occurs around $D = 1$. When we run a test comparing the methods with *identical* station distributions, they both “saturate” around $D = 1$. Thus, whatever the source of the discrepancy, they appear to be giving similar results for a given station distribution.

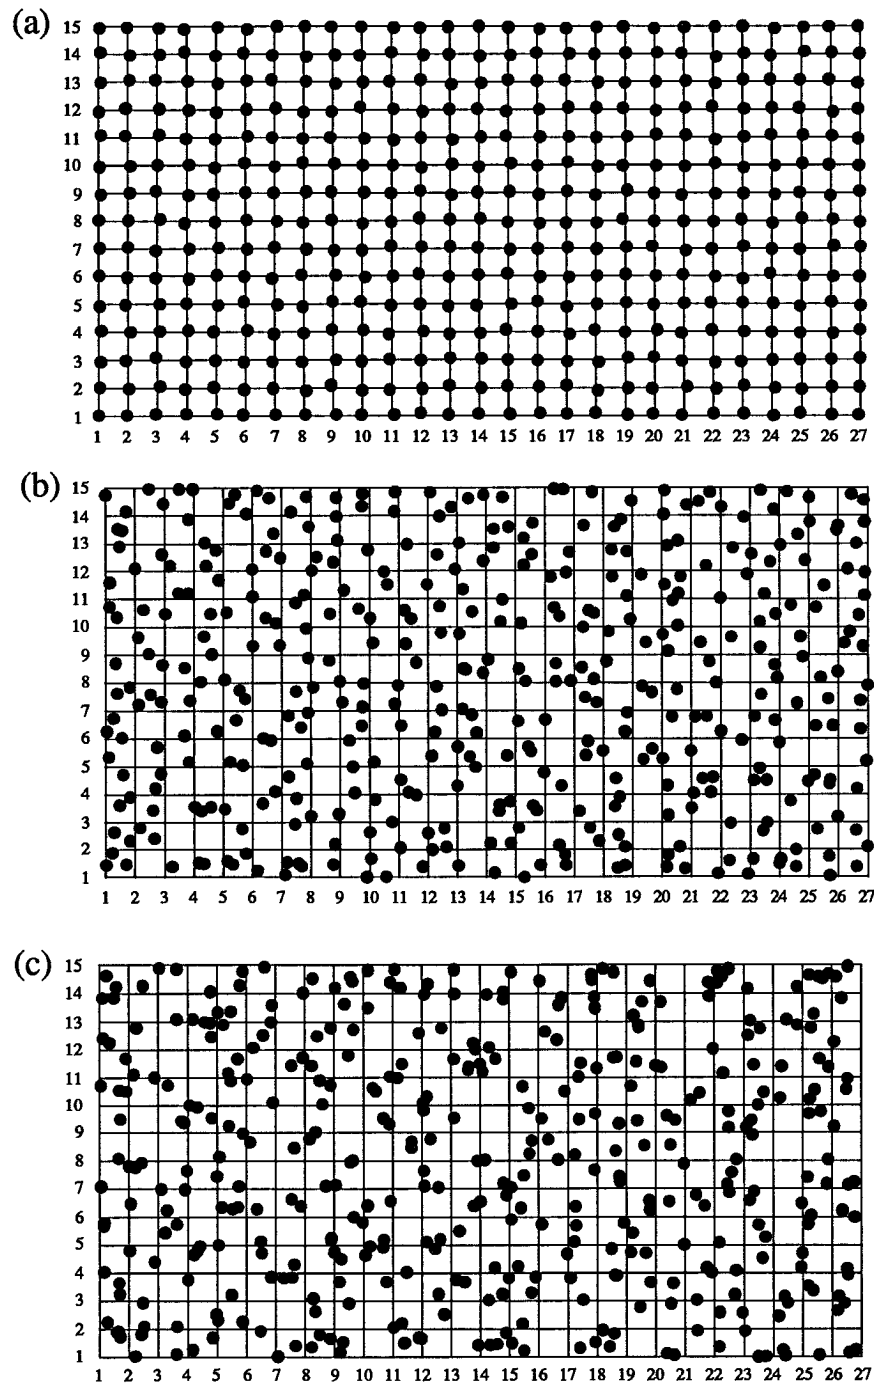


FIG. 6. Examples of increasingly irregular distributions created by perturbing uniform data points from their original locations (represented by the grid). Distributions are shown for (a) $D = 0.1$, (b) $D = 0.5$, and (c) $D = 1.0$.

distribution and is about 470 km. For the upper-air network, it is necessary to change the guard barrier to $2\Delta d$ (~ 723 km), because Δd (the median of the station spacing) is so large for this distribution that the $4\Delta d$ value suggested earlier does not leave much of an interior part of the dataset to evaluate. The computational grid spac-

ing is $\Delta d/6$, or about 60 km. The average μ is 0.24 for the upper-air network; by comparing this value to the μ values in Fig. 8 for the artificial distributions, we see that artificial distributions having the same amount of irregularity can be created using $D \approx 0.15$. Thus, we are able to create artificial distributions that are com-

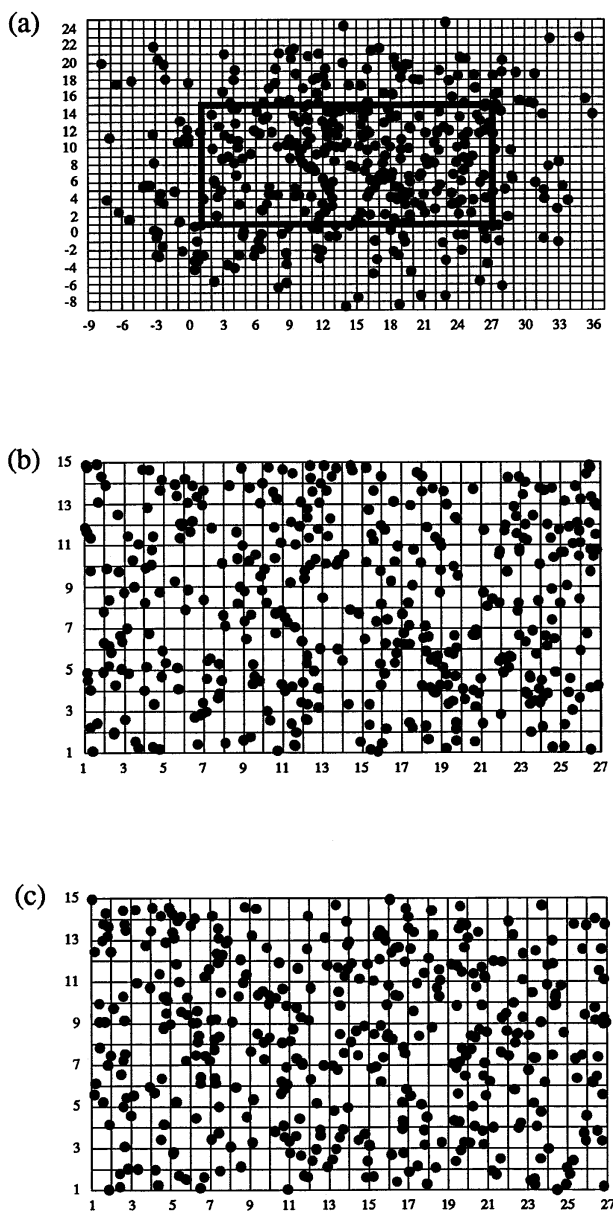


FIG. 7. Examples of artificial distributions for $D = 10$ using (a) dispersive boundaries, (b) periodic boundaries, and (c) reflective boundaries. The heavy solid line in (a) represents the original 27×15 data domain.

parable in terms of irregularity to the upper-air network with which to test an OA scheme.

It is of some interest to note, within this context, that the upper-air network has been undergoing some perturbations as a result of the modernization efforts within the National Weather Service. Using our method for characterizing the degree of irregularity of the distribution, Table 2 reveals that the changes in station siting have not significantly changed the regularity of the network yet. For those in the meteorological community who feel that the greater the degree of sampling irreg-

ularity, the lower confidence one can have in data analysis, any shuffling of the station sites is a major concern. With our proposed measure of irregularity, the irregularity of the resulting network can be monitored.

b. U.S. surface network

After using our proposed measure of irregularity on the surface network,⁵ we find the average value of μ to be 2.69, using a $\lambda \sim 56$ km, a guard barrier of $4\Delta d$, or about 173 km, and a computational grid spacing of $\Delta d/6$, or about 7 km. Comparing this μ value to those in Fig. 8, we find a curious result: We are unable to duplicate the amount of irregularity in the surface network by the random scattering process we have used to create the artificial distributions.

To understand the reason for this result, we again fit Poisson curves to the surface network's nearest-neighbor distribution and judge the goodness-of-fit with the Pearson test. The surface network is rejected as being Poisson at the 0.01 level of significance (Fig. 10). Considering Fig. 10, the surface network appears to be too clustered (many stations have very close nearest neighbors) to be considered spatially random. It is this clustering that makes the surface network so very irregular, so much that we are unable to duplicate it with the artificial distributions.

Based on the preceding results, we decided to *decluster* the surface network to decrease the irregularity. Our simple declustering algorithm is as follows: A cluster is defined by counting the number of stations within a certain distance, the *declustering radius*, of any given station. More than one station within the declustering radius constitutes a cluster. When a cluster is detected, a station in the cluster is removed as determined by the original ordering in the station listing. After the first of the stations in a cluster is removed, the cluster is tested again and stations are removed repeatedly until only one station in the original cluster remains.

Declustering the surface network does decrease the irregularity (see Fig. 11). Using our simple algorithm with different declustering radii, it was found that when it is declustered to remove stations less than 60 km apart, the irregularity is low enough to be duplicated by the artificial distributions. This is revealed in Fig. 12, with μ values for the declustered surface network overlaid onto the artificial network results originally shown in Fig. 8c. Thus, artificial distributions can be created with the same amount of irregularity as that of the declustered surface network using $D \approx 0.85$. The values of μ for the declustered surface network and the total surface

⁵ The surface network that we use includes all of the possible reporting stations. Generally, not all of the sites report at any given time, so the actual station density of real data sites will be slightly less than this density used here. When we test the actual reporting sites on several different days, the change in μ is less than 10%.

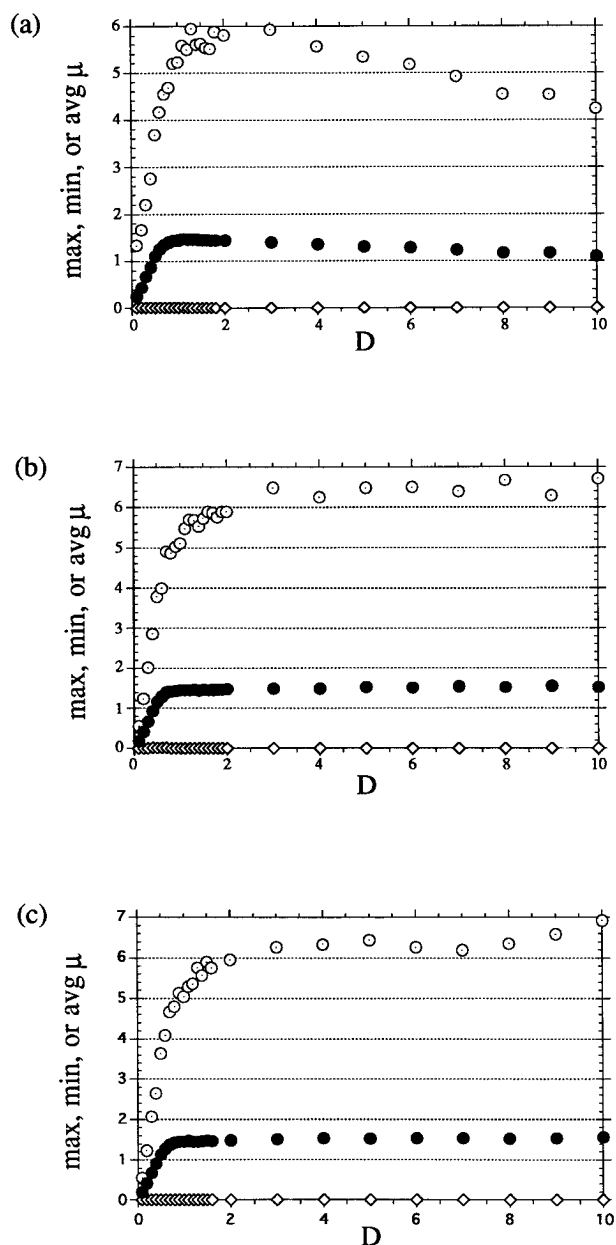


FIG. 8. Averaged maximum (open circles with dots), minimum (diamonds), and average (filled circles) μ values for the 20 simulations of artificial distributions versus D using (a) dispersive boundaries, (b) periodic boundaries, and (c) reflective boundaries.

network are also presented in Table 2 with the upper-air network results. It should be noted that even the declustered surface network fails the objective Pearson goodness-of-fit test for a Poisson distribution, although the visual appearance of the fit (not shown) is considerably better than that of Fig. 10.

5. Summary and future work

We have shown that it is possible to create artificial networks that closely match the characteristics of the

U.S. upper-air and *declustered* U.S. surface datasets by starting with a uniform grid of points and performing the appropriate perturbations. Therefore, we believe that our method of characterizing the degree of irregularity in a sampling array enables meteorologists to do empirical experiments with artificial networks with some assurance that their artificial networks have similar sampling characteristics to the real networks. Our approach to measuring the degree of irregularity of station distributions is simple both in principle and in practice so that it should be possible to execute an analysis of the irregularity in a dataset routinely before doing an objective analysis, and we recommend that those doing OA make it a practice to do so.

Future efforts in this area might well include a systematic exploration of analyses done with triangular computational grids. In B94, it was noted that in a triangular array of sites each site has six equidistant nearest neighbors, whereas in a square array, each site has only four equidistant nearest neighbors. Hence, in this restricted sense, a triangular array is “more uniform” than a square grid.

It also would be useful to explore alternative methods for declustering data networks for the purpose of achieving a roughly uniform distribution of points for objective analysis purposes. For example, the “superob” method (DiMego 1988) of replacing station clusters with a single station having the average location coordinates of stations within a cluster might well give somewhat better results than the simple scheme we have used. Also, it remains to be seen how one might create an artificial network with the distribution characteristics of the actual surface network before declustering. We believe that a method for artificial clustering the results of a “perturbations” experiment can be developed.

Finally, we have indicated that station distributions might have important impacts on objective analysis, owing to the gradient of the sum of the weights term as described in the appendix of DC88. It would be useful to know precisely at what degree of irregularity the OA is affected significantly from this term. As noted in DC88, when this term becomes important, the ordering of objective analysis and differentiation becomes important in gradient computations. Most schemes computing derivatives diagnostically do the objective analysis first, which DC88 contended is the improper order for irregular station distributions. Thus, some empirical testing with quantitative knowledge of the degree of irregularity would be valuable in deciding the validity of doing the objective analysis first.

Acknowledgments. We appreciate the helpful critiques of an earlier version of the manuscript by Dr. S. Barnes (NOAA–Forecast Systems Laboratory) and Prof. J. J. Stephens (The Florida State University). We benefited from discussions with Prof. M. Richman (University of Oklahoma) and from the critical comments contributed by Dr. H. Brooks and Mr. P. Spencer (NSSL). Finally,

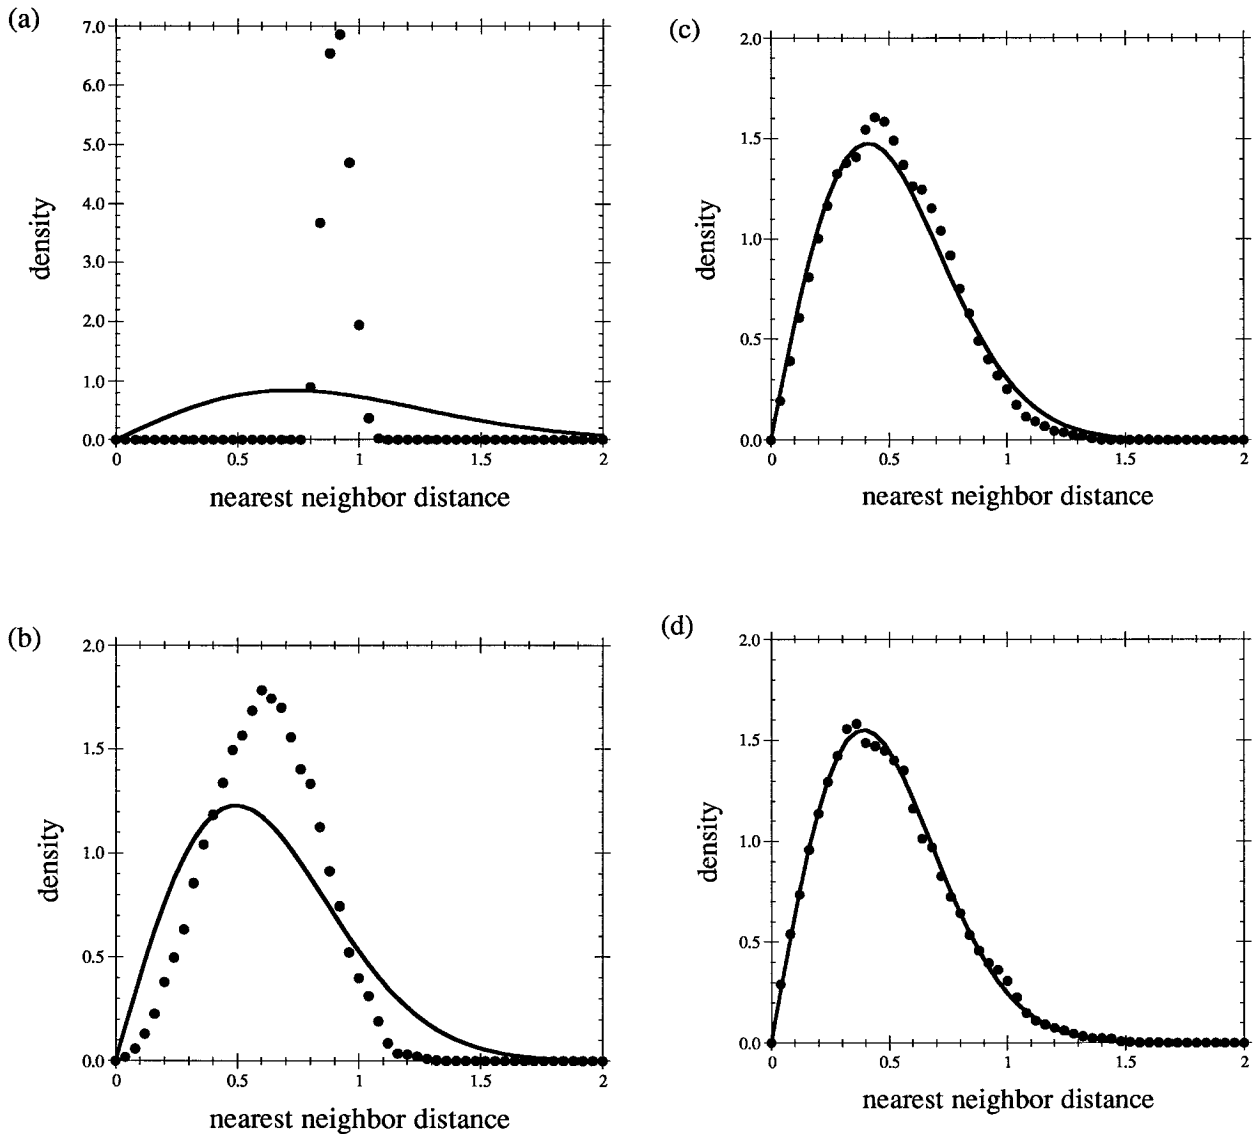


FIG. 9. Poisson curves fit to the averaged artificial distributions for (a) $D = 0.1$, (b) $D = 0.5$, (c) $D = 1.0$, and (d) $D = 5.0$. Solid lines denote the theoretical Poisson values while filled circles represent the observed values.

we wish to thank the anonymous reviewers for their suggestions clarifying the presentation. This work is based in part on the junior author's Master's thesis research, which was partially supported by a Patricia Roberts Harris Fellowship through the Department of Ed-

ucation. We also obtained partial support from the Center for Analysis and Prediction of Storms, CAPS (University of Oklahoma).

APPENDIX

Testing Distributions against the Poisson Distribution

According to Cressie (1991, 602 ff. and 633 ff.), "the distribution theory for nearest-neighbor distances . . . under complete spatial randomness is well-known." In a two-dimensional Cartesian space, the distribution function of the station to station distance has a density given by

TABLE 2. Maximum, minimum, and average values of μ for theoretically uniform and real sampling networks.

Sampling network	Max μ	Min μ	Avg μ
Uniform square grid	0.00	0.00	0.00
U.S. upper-air network, fall 1993	0.51	0.01	0.24
U.S. upper-air network, Nov 1994	0.55	0.01	0.23
U.S. upper-air network, Feb 1995	0.51	0.01	0.23
Total U.S. surface network	34.04	0.00	2.69
Declassified U.S. surface network	5.67	0.01	1.38




FIG. 10. Poisson curve fit to the nearest-neighbor distribution for the U.S. surface network pictured in Fig. 1.

$$g(x) = 2\pi\xi xe^{-\pi\xi x^2}, \quad x > 0,$$

where x is the distance from a station to its nearest neighbor and ξ is the intensity parameter, which can be approximated by the average data density over the domain. This distribution is derived by assuming that the station distribution is described by a homogeneous Poisson process, whereby the probability of having a station in a given small area dx^2 is given by ξ and that probability is essentially constant over the domain.

Poisson curves are fit to the nearest-neighbor distributions computed from the distribution to be tested using the method of least squares, which involves solving

$$\xi = \frac{\sum_i x_i g_i e^{-\pi\xi x_i^2}}{\pi \sum_i x_i^3 g_i e^{-\pi\xi x_i^2} + 2\pi \sum_i x_i^2 e^{-2\pi\xi x_i^2} - 2\pi^2 \xi \sum_i x_i^4 e_i^{-\pi\xi x_i^2}}$$

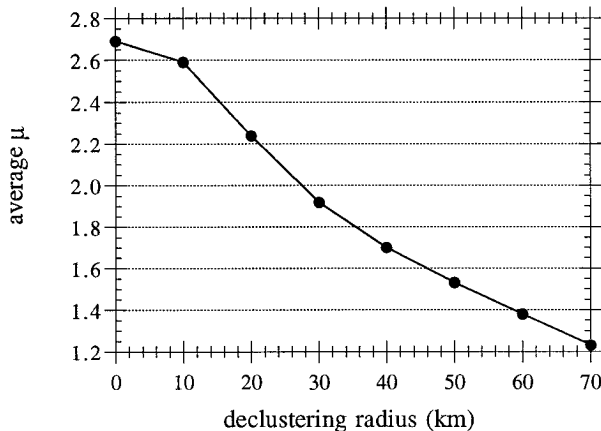


FIG. 11. Average μ as a function of the declustering radius used to decluster the U.S. surface network.




FIG. 12. As in Fig. 8c except the declustered surface network values of μ (denoted by squares) are overlaid. The declustering radius is 60 km.

iteratively for ξ (the parameter of the Poisson distribution); g is the distribution function described above. Despite its formidable appearance, the iterative solution converges rapidly. The fit of the Poisson curve to the nearest-neighbor distribution is judged by the Pearson test statistic C_1 , defined by

$$C_1 = \sum_{i=1}^k \frac{(X_i - n\hat{p}_{i_0})^2}{n\hat{p}_{i_0}},$$

where k is the number of classes in the nearest-neighbor distribution, X_i is the observed number in each nearest-neighbor category, n is the total number of points, and \hat{p}_{i_0} is the theoretical Poisson probability for each nearest-neighbor category. As long as $\hat{p}_{i_0} > 5$ for each nearest-neighbor category, C_1 can be treated as a chi-squared random variable (Larsen and Marx 1986), and the associated hypothesis test is to reject the distribution as being Poisson if

$$C_1 \geq \chi_{1-\alpha, k-2}^2$$

at the α level of significance, with $k - 2$ degrees of freedom.

REFERENCES

Achtmeier, G. L., 1986: The impact of data boundaries upon a successive corrections objective analysis of limited-area datasets. *Mon. Wea. Rev.*, **114**, 40–49.

Barnes, S. L., 1964: A technique for maximizing details in numerical weather map analysis. *J. Appl. Meteor.*, **3**, 396–409.

—, 1994: Applications of the Barnes objective analysis scheme. Part I: Effects of undersampling, wave position, and station randomness. *J. Atmos. Oceanic Technol.*, **11**, 1433–1448.

Buzzi, A., D. Gomis, M. A. Pedder, and S. Alonso, 1991: A method to reduce the adverse impact that inhomogeneous station distributions have on spatial interpolation. *Mon. Wea. Rev.*, **119**, 2465–2491.

Caracena, F., S. L. Barnes, and C. A. Doswell III, 1984: Weighting function parameters for objective interpolation of meteorological data. Preprints, *10th Conf. Weather Forecasting and Analysis*, Clearwater Beach, FL, Amer. Meteor. Soc., 109–116.

- Cressie, N., 1991: *Statistics for Spatial Data*. John Wiley and Sons, 900 pp.
- DiMego, G. J., 1988: The National Meteorological Center regional analysis system. *Mon. Wea. Rev.*, **116**, 977–1000.
- Doswell, C. A., III, and F. Caracena, 1988: Derivative estimation from marginally sampled vector point functions. *J. Atmos. Sci.*, **45**, 242–253.
- Grassberger, P., and I. Procaccia, 1983: Measuring the strangeness of strange attractors. *Physica*, **9D**, 189–208.
- Koch, S. E., M. des Jardin, and P. J. Kocin, 1983: An interactive Barnes objective map analysis scheme for use with satellite and conventional data. *J. Climate Appl. Meteor.*, **22**, 1487–1503.
- Korvin, G. D., M. Boyd, and R. O'Dowd, 1990: Fractal characterization of the South Australian gravity station network. *Geophys. J. Int.*, **100**, 535–539.
- Larsen, R. J., and M. L. Marx, 1986: *An Introduction to Mathematical Statistics and Its Applications*. Prentice-Hall, 630 pp.
- Lovejoy, S., D. Schertzer, and P. Ladoy, 1986: Fractal characterization of inhomogeneous geophysical measuring networks. *Nature*, **319**, 43–44.
- Pauley, P. M., 1990: On the evaluation of boundary errors in the Barnes objective analysis scheme. *Mon. Wea. Rev.*, **118**, 1203–1210.
- , and X. Wu, 1990: The theoretical, discrete, and actual response of the Barnes objective analysis scheme for one- and two-dimensional fields. *Mon. Wea. Rev.*, **118**, 1145–1210.
- Press, W. H., S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, 1986: *Numerical Recipes*. Cambridge University Press, 818 pp.
- Smith, D. R., M. E. Pumphry, and J. T. Snow, 1986: A comparison of errors in objectively analyzed fields for uniform and non-uniform station distributions. *J. Atmos. Oceanic Technol.*, **3**, 84–97.
- Tessier, Y., S. Lovejoy, and D. Schertzer, 1994: Multifractal analysis and simulation of the global meteorological network. *J. Appl. Meteor.*, **33**, 1572–1586.