

Maximum Likelihood Spectral Fitting: the Batchelor Spectrum

Barry Ruddick
Ayal Anis
Keith Thompson

Department of Oceanography
Dalhousie University
Halifax, N.S., B3H 4J1
CANADA

April 28, 1999

DRAFT-- Submitted, Journal of Atmospheric and Oceanic Technology

Abstract

We describe a simple technique for fitting spectra that is applicable to any problem of adjusting a theoretical spectral form to fit observations. All one needs is a functional form for the theoretical spectrum, and an estimate for the instrumental noise spectrum. The method, based on direct application of the Maximum Likelihood approach, has several advantages over other fitting techniques:

1. It is unbiased in comparison with other least-squares or cost function-based approaches.
2. It is insensitive to dips and wiggles in the spectrum. This is because the range of wavenumbers used in the fit does not vary, and the built-in noise model forces the routine to ignore the spectrum as it gets down towards the noise level.
3. Error bars. There is a theoretical estimate for the variance of the fitted parameters, based on how broad or narrow the likelihood function is in the vicinity of its peak.
4. We calculate statistical quantities that indicate how well the observed spectrum fits the theoretical form. This is extremely useful in automating analysis software, to get the computer to automatically flag "bad" fits.

The method is demonstrated using data from SCAMP (Self-Contained Autonomous Microstructure Profiler), a free-falling temperature microstructure profiler. Maximum Likelihood fits to the Batchelor spectrum are compared to the SCAMP-generated fits and other least-squares techniques, and also tested against pseudo-data generated by Monte-Carlo techniques.

Pseudo-code outlines for the spectral fit routines are given.

1. Introduction: Turbulent Microstructure - what are χ and ϵ ?

When a thermally-stratified fluid is stirred, fluid parcels are moved from warm to cold regions (and vice-versa), where they become thermal anomalies. As a result, the mean-square thermal perturbations $\overline{(T\Theta)^2}$ grow. The budget equation for the variance of these thermal anomalies is, neglecting a divergence term (Osborn and Cox, 1972):

$$\frac{\partial}{\partial t} \overline{(T\Theta)^2} = -2\overline{(w\Theta T')} \frac{\partial \bar{T}}{\partial z} - \chi_\theta, \quad (1.)$$

where primes denote turbulent quantities, κ_T the thermal diffusivity, and $\chi_\theta \equiv +2\kappa_T \overline{(\nabla T')^2}$.

The first term on the right of (1), the product of turbulent heat flux and average thermal gradient, describes the production of thermal anomalies by stirring.

As soon as thermal anomalies are formed, they begin to blend into the background by molecular heat diffusion. The rate of decrease of thermal variance due to molecular diffusion, is represented by the second term on the right, $-\chi_\theta$. If the production and dissipation terms balance on average, the turbulent heat flux can be estimated by measuring χ_θ and the mean temperature gradient (Osborn and Cox, 1972). Microstructure profiles of a single component of temperature gradient (i.e., $\partial T'/\partial z$) on cm or smaller scales are often used to infer χ_θ and hence heat flux.

Similarly, stirring imparts kinetic energy ($TKE = \frac{\rho}{2} \overline{(u'^2 + v'^2 + w'^2)}$) to the turbulent motions (first term on the right, below), and this energy is partially expended in increasing the potential energy of the system (second term), with the rest being dissipated by molecular viscosity (ϵ , third term). The TKE budget equation expressing this balance is, approximately (c.f., Ruddick et al 1997)

$$\frac{\partial}{\partial t} \overline{TKE} = -\bar{\rho} \overline{(u'w')} \frac{\partial \bar{u}}{\partial z} - g \overline{(\rho'w')} - \bar{\rho} \epsilon$$

where

$$\epsilon = -\nu \overline{\frac{\partial u'_i}{\partial x_j} \left(\frac{\partial u'_i}{\partial x_j} + \frac{\partial u'_j}{\partial x_i} \right)}$$
(2.)

is the rate of viscous dissipation, and is proportional to the mean-square straining of the turbulent eddies. It is usually measured by direct observation of microscale shears. The rhs expresses, respectively, the increase of TKE by stirring, the conversion to mean potential energy by turbulent buoyancy flux, and energy loss by viscous dissipation.

Observations of ε , usually obtained directly via observations of microscale shear, can be used to infer an upper bound for the turbulent buoyancy flux (Osborn, 1980). When both χ and ε are observed, then the combination can be used to infer mixing efficiency (Oakey, 1982, Moum, 1996), and have been used to test for unequal mixing of heat and salt (Ruddick et al, 1997, St. Laurent et al, 1999)

Dillon and Caldwell (1980) and Oakey (1982) showed that it is possible to estimate χ_θ directly plus ε indirectly by fitting observed temperature gradient spectra to the theoretical "Batchelor" form, whose cutoff wavenumber depends on Epsilon and thermal diffusivity (section 2). An instrument used for this purpose is described in section 3.

The main content of this paper, section 4, is a description of a relatively simple and useful technique for fitting theoretical spectra to observations. It was developed in the context of Batchelor spectral fitting, but should be useful in the context of fitting any theoretical spectrum to observations. It is based on the direct application of the Maximum Likelihood Estimation (MLE) principle, resulting in an efficient and unbiased estimator. The technique has three additional advantages over more conventional least-squares techniques:

- Least squares techniques assume Gaussian errors, while MLE does not.
- Error bars - a theoretical estimate for the standard error of the fit parameters (in this case, the Batchelor wavenumber and ε), and
- Rejection criteria - a statistical measure to indicate when the observed and theoretical spectra are too poor a match to be accepted.

Examples of the application of the technique to Batchelor spectral fitting are given in section 5. Monte-Carlo tests of the method, in comparison with least-squares and other cost-function techniques, are shown in section 6. Discussion is in section 7, and the MLE meta-code is given in Appendix A.

2. Batchelor spectrum

Batchelor (1959) discussed the small-scale behaviour of passively convected, diffusive scalars like temperature and salinity in the presence of turbulence. He noted that for scales smaller than the Komolgorov viscous cutoff scale, $(\nu^3/\epsilon)^{1/4}$, the flow appears as a randomly varying straining motion with root mean square strain rate (and corresponding inverse timescale) $\tau^{-1} = 0.5 (\epsilon/\nu)^{1/2}$. The strain concentrates the large-scale temperature gradient to form a "micro-front", and molecular diffusion then acts to keep the front from becoming infinitely sharp (see figure 1). Strain and diffusion balance when the frontal width is the diffusion scale, $\ell = \sqrt{\kappa_T \tau}$, and eliminating τ in favour of ϵ gives

$$\ell^{-1} \equiv k_B = \left(\frac{\epsilon}{\nu \kappa_T^2} \right)^{1/4}, \quad (3.)$$

the "Batchelor cutoff wavenumber". Note that the units of k_B are in radians per meter; the conversion factor 2π radians per cycle must be applied to convert to cpm.




Figure 1 - Straining field in a mean temperature gradient, and "micro-front" produced by the convergence.

Batchelor (1959) was then able to solve the advection-diffusion equation for temperature, driven by turbulent strain and a large-scale temperature gradient, to derive an analytic expression for the high-wavenumber part of the temperature gradient spectrum:

$$S_B(k; k_B, \chi_\theta) = (q/2)^{1/2} \chi_\theta k_B^{-1} \kappa_T^{-1} f(\alpha)$$

where

$$S_B(\alpha) = \alpha \left[e^{-\frac{\alpha^2}{2}} - \alpha \int_\alpha^\infty e^{-\frac{x^2}{2}} dx \right] \quad (4.)$$

$$\alpha = k k_B^{-1} \sqrt{2q}$$

The universal constant q is in the range 3.4-4.1 (Dillon and Caldwell, 1980, Oakey, 1982), and the value 3.4 is used here in order to match that used in the SCAMP software. Dillon and Caldwell (1980) note that a percentage error in q leads to twice the percentage error in ϵ .

This spectrum is a self-similar form that scales with the Batchelor wavenumber, k_B , rising slowly to the Batchelor cutoff, and falling exponentially beyond it. Changing χ_θ shifts the spectrum vertically, and changing ϵ with χ_θ fixed shifts the spectrum without change of shape along a line of slope -1 (Figure 2).




Figure 2: Batchelor Spectrum for various k_B , including estimated SCAMP instrumental noise level ($3 \cdot 10^{-7}$ ($^{\circ}\text{C}/\text{m})^2/\text{cpm}$), with χ_{θ} constrained according to eq. (9). The k_B value corresponding to each curve is indicated with a plus, and the approximate corresponding dissipation level, ε , is shown by the second logarithmic scale below the k -axis. The effects of changing χ_{θ} and ε on the spectrum are indicated by arrows. An observed spectrum is shown.

Dillon and Caldwell (1980) made well-resolved observations of temperature gradient spectra, and tested these against the hypothesized Batchelor scaling, forming ensemble-averaged dimensionless spectra. They found the observed gradient spectra matched the Batchelor form exceptionally well for intense turbulent events (Cox numbers $\chi_{\theta}/6\kappa_T(\bar{T}_z)^2$ greater than 2500), and deviated at low wavenumbers for Cox numbers less than 500. In all cases the peak and cutoff agreed well with theory.

Oakey (1982) made direct velocity-based observations of ε simultaneously with temperature gradient measurements, and compared the ε found from Batchelor fits to selected, well-resolved records with the velocity-based ε . He found the agreement to be excellent, proving that ε can be estimated from well-resolved temperature gradient spectra.

3. SCAMP

SCAMP (Self-Contained Autonomous Microstructure Profiler) is a small, hand-held profiling instrument manufactured by Precision Measurement Engineering of Encinitas, CA. It free-falls slowly through the water column at a nominal rate of 0.1 m/s and measures T at 100 Hz in order to resolve the Batchelor spectral peak. The instrument also can be fitted with microscale conductivity, accurate T , conductivity, and dissolved oxygen sensors, and has a pressure sensor for instrument depth.

Processing software (HOST) provided with the system estimates χ_θ and ε by fitting the Batchelor form to the measured temperature gradient spectrum. The HOST software uses a weighted least squares algorithm, adjusting *both* χ_θ and ε to minimize the cost function

$$C1 = \sum_{k_i} ((S_{\text{obs}} - S_B) / S_B)^2 \quad (5)$$

which is the square of the difference between Batchelor spectrum (S_B) and observed (S_{obs}), weighted by the squared Batchelor spectrum. The wavenumber range used for the fit is selected to avoid regions dominated by instrumental noise.




Figure 3 - SCAMP fit example, using C1, equation 5. MLE fit (section 4) is shown for comparison. Note the cutoff for the SCAMP fit (C1) to avoid the noise-dominated high-wavenumber region. Same observational spectrum is used in figures 2 and 3.

Poor convergence of the HOST fit algorithm (version 1.08 - since improved) led us to develop improved fit techniques. An additional goal was to develop fitting algorithms with rejection criteria that could be completely automated, making it unnecessary to manually edit the results for poor fits. These should be useful in a variety of spectral fitting contexts, and we describe them in section 4.

We began by constructing an instrumental noise model spectrum, and including that in the theoretical fitted spectrum. We also constrained χ_θ to match the integrated temperature gradient spectrum less modelled noise, leaving a single parameter to be varied, k . Both these improvements are described in more detail in section 4.1. The remainder of the fitting procedure involved generating a family of Batchelor spectra plus noise curves for all values

of k_B and selecting the one curve that provides the "best" match to the observed spectrum (figure 2).

Cost Function approach

Prior to implementing the MLE algorithm (section 4.3), we experimented with a variety of least-squares based cost functions, the first being weighted least squares, C1, described in the previous section. (A cost function is simply a measure of misfit.) Another cost function tried was:

$$C3 = \sum_k (\ln(S_{\text{obs}}) - \ln(S_{\text{th}}))^2 \quad (6.)$$

with $S_{\text{th}} = S_B + S_n$ denoting the theoretical spectrum, and S_n denoting instrumental noise.

This adjusts parameters to minimize the mean squared distance in log space between the observed and theoretical spectra. Leuketina and Imberger (1999) minimize a similar quantity, involving the mean squared distance in log space along a slope of +1. This paper also has an excellent discussion of the issues relating to instrument drop speed, fine structure contamination, instrumental noise, and the range of wavenumbers to be used in the fit.

We found that C3 works reasonably well, except that our Monte-Carlo tests found it to be biased low.

We experimented with several other cost functions, some of which amounted to least-squares with different kinds of weighting. They are described in section 7, where they are compared with each other and the MLE method using Monte-Carlo generated data sets. We eventually developed the MLE algorithm because we realized that experimenting with cost functions was too ad-hoc, with little theoretical guidance.

4. Maximum Likelihood Spectral Fitting

Fitting a theoretical spectral form to observations is a common problem, and is usually approached by a least-squares technique. Most time series textbooks discuss the estimation of Auto-Regressive-Moving Average (ARMA) models from time series of observations via least-squares estimation; however, those (linear) techniques are difficult to adapt to the fitting of a parametrically-defined spectrum such as the Batchelor form. Leuketina and Imberger (1999) discuss fitting the Batchelor spectrum to an observed spectrum, and describe their experiences with a variety of least-squares-based cost functions. They also clearly discuss issues relating to instrument noise, finestructure contamination, and the range of wavenumbers to be used in the fit. The spectral fitting techniques we describe in this section (allowing for instrument noise in the fitted spectrum, reduction to a single-parameter fit, and direct application of the Maximum Likelihood method) yield a stable, robust, and accurate algorithm that seems to be an improvement over least squares. Although the techniques are not new, they have not been described in the context of the common and important task of fitting a spectrum, and so are described below. We describe general improvements in section 4.1, then we give a "textbook-style" introduction to estimating the mean by the method of maximum likelihood in section 4.2, and discuss spectral fitting by direct application of MLE in section 4.3. A significant advantage of MLE is that, when the model is valid, the shape of the likelihood function yields an estimate for error bars on the fitted parameters. This is discussed in section 4.4.

Another aspect of fitting a spectrum to observations is that of deciding when a fit is too poor to be accepted -- sometimes the theoretical spectrum cannot be made to fit the observations acceptably well for any value of the parameters. Priestley (1981, section 6.2.6) discusses several "goodness-of fit" tests designed to answer the question: "Does the spectrum computed from observations conform to the theoretical form hypothesized?" However, most of the tests described require that the spectral form be fully specified *a priori*, without being fit to the observations. Since two parameters, the variance and the Batchelor wavenumber, are constrained by the observations, the tests cannot be used. An exception is Bartlett's homogeneity of chi-squared test (Priestley, pp 486-488, Bartlett, 1954), which is, according to Priestley, "...asymptotically valid even when [the spectrum] is not fully specified but involves parameters which must be estimated from the data". This test involves testing the ratio of the observed and theoretical spectrum to see if it is consistent with a χ_d^2 probability distribution. The approach is similar to the tests described

in section 4.5, where we discuss various measures of spectral misfit and our experiences with them.

4.1 General Improvements

1. Allow for instrumental noise spectrum $S_n(k)$ as a component of the theoretical spectrum.

It is important to recognize that instrumental noise will act as a lower limit to any measured spectrum. Leuketina and Imberger (1999), limit the range of their fit to wavenumbers not dominated by noise. If the fitted wavenumber range is varied with the trial k_B , then the cost function has local minima due to the statistical dips and wiggles in the observed spectrum, making it more difficult to locate the global minimum, or to decide which minimum to select.

If the instrumental noise spectrum $S_n(k)$ can be estimated from bench tests, from quiet sections of field records, or any other way, then it is simple to include it in the theoretical spectrum that is fit to observations:

$$S_{th}(k) = S_B(k) + S_n(k) \quad (7.)$$

The simple modification (7.) has the benefit that the fitted spectrum only changes in the regions where the theoretical spectrum exceeds the noise level, and remains constant where it lies below (figure 2). Thus, any cost function does not change at locations where the signal is less than the noise, and the fit "ignores" those regions. For example, the derivative of the cost function $C3$ with respect to k_B can be easily calculated as:

$$\frac{\partial C3}{\partial k_B} = -2 \left(\frac{1}{1 + S_n / S_B} \right) \frac{\partial}{\partial k_B} \ln(S_B) \quad (8.)$$

The second factor on the right goes from 1 to 0 where the signal becomes less than the noise, giving a smoothly varying weighting factor that depends on the expected signal/noise ratio.

In our experience, it's better to slightly overestimate the noise than underestimate it. Otherwise, most fit techniques will spend effort inappropriately trying to fit to noise. The

MLE fit in figure 3 was done using the simplest of noise models -white noise, and in this case the fit appears to work well and is not sensitive to the chosen noise level. A more detailed and accurate noise model, as is used in all later examples, allows successful fits in situations with a much lower signal/noise ratio. This noise model was made from a parametric fit based on the SCAMP thermistor circuit and theoretical/electronics considerations (M. Head, pers. comm. 1999).

2. Reduce the number of free parameters adjusted

Dillon and Caldwell (1980), and Oakey (1982) showed how χ_θ is constrained by the integrated temperature gradient spectrum, after allowance for instrument noise:

$$\chi_\theta = 6\kappa_T \int_0^\infty (S_{obs} - S_n) dk \quad (9.)$$

In the case of fitting the Batchelor spectrum, only a single parameter, k_B , remains to be adjusted. This amounts graphically to sliding the Batchelor curve along a slope of -1, and selecting the k_B and corresponding ε values that give the "best" fit (figure 2).

The advantages of reducing the dimension of the parameter space searched are many. First, valuable observations are not "used" in adjusting parameters that can be determined in other ways. Second, the adjusted value of a parameter might be inconsistent with other constraints, making it difficult to decide what to believe. Third, searching in multiple dimensions to minimize a computed quantity can be difficult and time-consuming, especially if that quantity is not a quadratic. The SCAMP software fit algorithm searched in two dimensions using the Marquardt-Levy algorithm, and did not converge rapidly enough in the k_B direction, leading to significant inaccuracies in that parameter.

4.2 A Simple Example: MLE for Gaussian noise, and relationship with least squares

The material in this section can be found in several statistics texts, and is included to remind the reader how MLE leads to the familiar least squares minimization approach in the case of Gaussian errors. The MLE concepts are then applied to the specific case of spectral fitting in section 4.3.

Suppose we measure the length of an object many times, each time with an error. If the error is distributed as a Gaussian, then the probability of obtaining a measured length y is given by:

$$f(y) = C \exp\left[-\frac{(y - \mathcal{L})^2}{2\sigma^2}\right], \quad (10.)$$

where \mathcal{L} is the true length, σ is the standard deviation of the error (here assumed constant), and C is a normalizing factor to make the integral of f equal 1. If we make a histogram of a large number of measurements, it will look something like figure 4a.

What should we do if we don't know the true length, \mathcal{L} , and wish to find the best estimate of it from our observations? The expected probability density function (pdf) of the measured length is shown for $a \sim 1.7$ in figure 4b. The Maximum Likelihood Estimation (MLE) procedure is analogous to considering the family of pdfs for all values of a , and picking the one that agrees "best" with the observed histogram -- i.e., by sliding the Gaussian curve (b) along the y-axis until it agrees best with the histogram (a). For MLE, "best" agreement is the value of a for which the particular set of observations *is most probable*.

Consider a single measurement – the probability of finding it in the i^{th} bin of fig. 4a, centered on y_i is the pdf times the bin width:

$$p_i = f(y_i)\Delta y \quad (11.)$$

If we make N measurements, $y_i = y_1, y_2, \dots, y_N$, the combined probability P of obtaining that particular set of measurements is the product of the probability for each measurement:

$$\begin{aligned} P &= \prod_{i=1}^N p_i \\ &= \prod_{i=1}^N f(y_i)\Delta y \end{aligned} \quad (12.)$$

Obviously, if the data set contains lots of measurements that are far from a , then it is a very unlikely data set: $f(y)$ is small for most points, and so P is small.

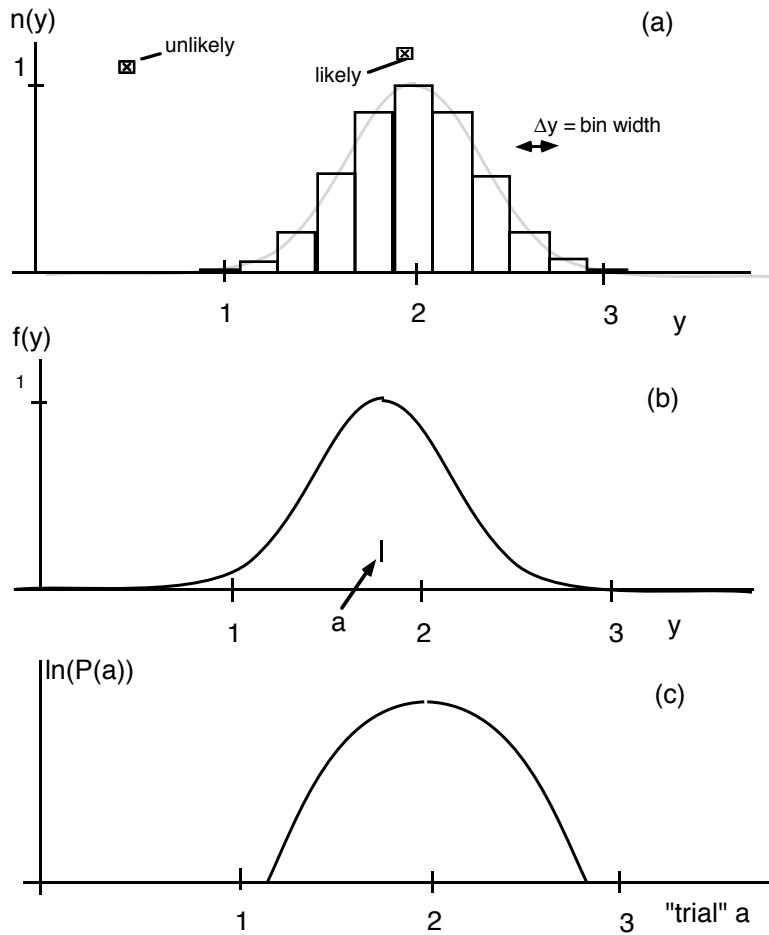


Figure 4. (a) Histogram of observations of a quantity with true value $\mathcal{f} = 2$, and Gaussian error. (b) Hypothesized Gaussian pdf with $a \sim 1.7$. (c) The $\ln(\text{probability})$ of our observations as a function of the hypothesized value of a , showing how $\ln(P)$ is maximized as a is changed.

Now we turn the question around - we know the y , but not \mathcal{f} :

Given N observations of y , what is the value of a that makes the observed y most likely?

To answer this, we guess a trial value of a , compute the likelihood P according to eq. 12, and then repeat for all possible values of a . We choose the a that gives the largest P – the one with the Maximum Likelihood. Maximising P is the same as maximizing $\ln(P)$, so it's more convenient to maximize the sum of the log of individual probabilities:

$$\begin{aligned}
\ln(P) &= \ln\left(\prod_{i=1}^N p_i\right) \\
&= \sum_{i=1}^N \ln(p_i) \\
&= \sum_{i=1}^N \ln(f(y_i)) + N \ln(\Delta y)
\end{aligned} \tag{13}$$

and substituting the Gaussian pdf (10):

$$\begin{aligned}
\ln(P) &= \sum_{i=1}^N \ln\left\{C\Delta y \exp\left[-\frac{(y_i - a)^2}{2\sigma^2}\right]\right\} \\
&= \sum_{i=1}^N \ln\left\{\exp\left[-\frac{(y_i - a)^2}{2\sigma^2}\right]\right\} + N \ln C + N \ln \Delta y \\
&= \sum_{i=1}^N -\frac{(y_i - a)^2}{2\sigma^2} + N \ln C + N \ln \Delta y
\end{aligned} \tag{14}$$

Note that σ , N , C , and Δy don't change when a is changed (C is important when the pdf is non-Gaussian, or when estimating σ). So for *Gaussian* error distribution, the a that makes the observed data set "most likely" is the one that minimises the sum of squares $S \equiv \sum_{i=1}^N \frac{(y_i - a)^2}{2\sigma^2}$. If we compute S as a function of the trial a , we will in this case find S is

a parabola (figure 4c). Minimizing S :

$$\begin{aligned}
0 &= \frac{\partial}{\partial a} \sum_{i=1}^N (y_i - a)^2 \\
&= \sum_{i=1}^N (-2)(y_i - a) \\
\Rightarrow a &= \frac{1}{N} \sum_{i=1}^N y_i
\end{aligned} \tag{15}$$

Thus, the Maximum Likelihood estimate for a is the usual formula for the mean.

4.3 MLE model for Spectral fits

The exponential in the Gaussian pdf led in (14) to the sum of squares as the function to minimize. We don't need to assume a gaussian pdf for the observations, but can work with any other form of pdf. Going back to equation 13, we maximize $\ln(P) = \sum_{i=1}^N \ln(f(y_i))$,

where f is the expected pdf of y . (Note that the $N\Delta y$ factor is a constant, and so can be ignored.)

Our statistical model assumes the observed spectrum, S_{obs} , is equal to its theoretical value (the Batchelor spectrum plus instrument noise), but is subject to statistical variability (Jenkins and Watts, 1968), and is distributed as a χ^2_d pdf, with d degrees of freedom, such that d depends on the spectral technique, window, and averaging methods used (Jenkins and Watts, 1968):

$$Y_i \equiv \frac{S_{obs}}{S_{th}} \sim \frac{\chi_d^2}{d} \quad (16.)$$

So the pdf for S_{obs} is:

$$f(S_{obs}) = \frac{d}{S_{th}} \chi_d^2 \left(\frac{dS_{obs}}{S_{th}} \right) \quad (17.)$$

with S_{th} given by (7). The normalizing factor in front makes the integral of f equal 1. From equation 13, the MLE for k_B is found by maximizing

$$C11 \equiv \ln(P) = \sum_{i=1}^N \ln \left(\frac{d}{S_B(k_i; k_B, \chi_\theta) + S_n(k_i)} \chi_d^2 \left(\frac{dS_{obs}(k_i)}{S_B(k_i; k_B, \chi_\theta) + S_n(k_i)} \right) \right) \quad (18.)$$

as a function of k_B , with χ_θ constrained by (9). In our experience with fitting a Batchelor spectrum, (18) has a single maximum, so that any simple search technique will suffice. This may not be true for spectra badly contaminated with finestructure, or for other theoretical spectral shapes, or for much noisier or more poorly resolved signals. If a spectral form with multiple adjustable parameters is fit, then more sophisticated search techniques would need to be used.

There are two special cases of interest. First, when the raw periodogram is used, so that the degrees of freedom for the spectral estimator is $d=2$, then

$$\chi_2^2(y) = \frac{1}{2} \exp(-y/2) \quad (19.)$$

(cf Jenkins and Watts, 1968), and C11 reduces to

$$C12 = -\sum_{i=1}^N \frac{S_{obs}}{S_{th}} - \sum_{i=1}^N \ln(S_{th}) + N \ln(2) \quad (20.)$$

in agreement with Brillinger (1985). We verified that our computed C11 formula agreed with C12 for $d=2$.

Second, when d becomes asymptotically large (in practice, d must exceed about 60), then the $\frac{1}{d} \chi_d^2$ pdf becomes asymptotically normal with mean 1 and variance $2/d$. The MLE estimate C11 becomes a weighted least squares estimate:

$$C11|_{d \rightarrow \infty} \equiv -\sum_{i=1}^N \left(\frac{S_{obs} - S_{th}}{S_{th}} \right)^2 - \sum_{i=1}^N \ln(S_{th}) + N \left(\frac{1}{4} + \sqrt{\frac{d}{4\pi}} \right) \quad (21.)$$

which, aside from the second term, is equivalent to C1.

Figure 5, lower left panel, shows the likelihood function $\exp(C11)$ as a function of the trial k_B , computed for a SCAMP data segment with a large signal/noise ratio. The likelihood function is remarkably close to a Gaussian shape, supporting the estimated error bars (shown) that are described in section 4.4. The lower right panel shows the SCAMP temperature gradient spectrum, computed with 6 degrees of freedom using a Hanning window, and the MLE fitted spectrum corresponding to the peak in the likelihood function. We also show the parametric noise model adopted for the fits, and the alternative power law fit (section 4.5).

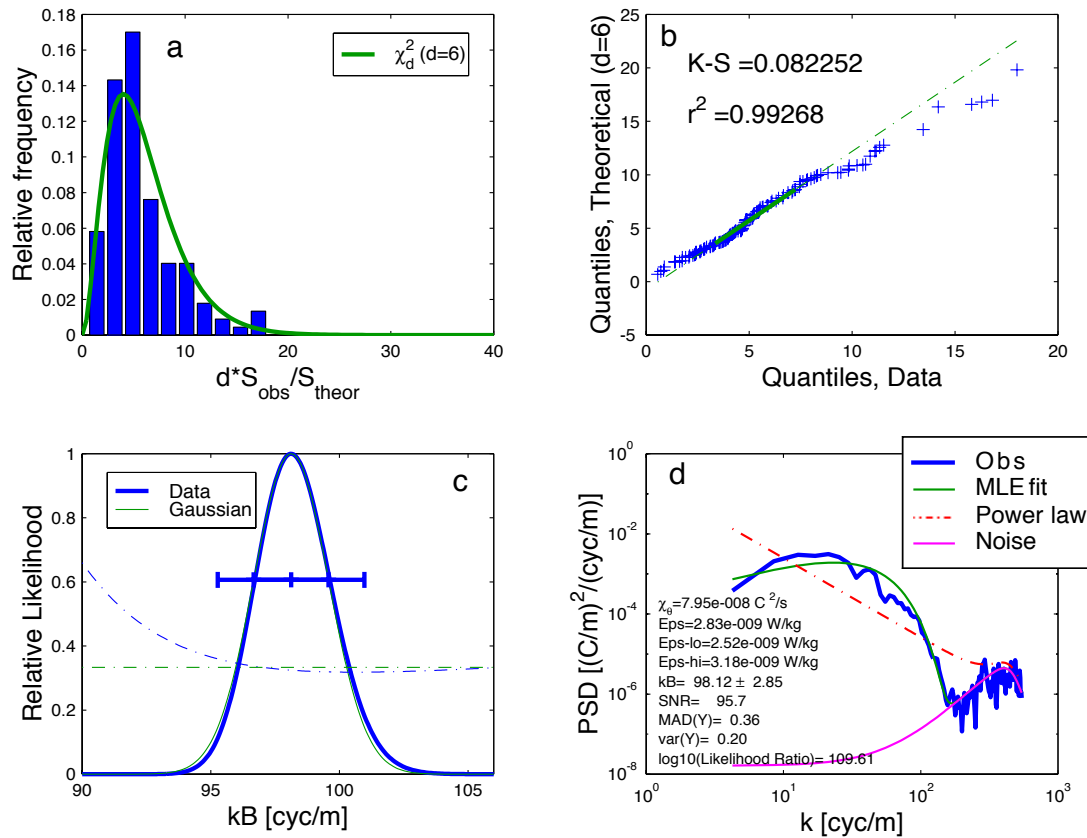


Figure 5. MLE fit to observed spectrum. **(c)** Relative likelihood function $\exp(C11)$ (eq. 18) normalized to a maximum of 1 versus trial k_B , and Gaussian function with same second derivative at maximum. Horizontal error bars indicate 1 standard deviation according to eq. (22) by inner vertical bars, plus approx. 95% confidence limits by outer vertical bars. Horizontal dash-dot line indicates the theoretical level of the variance (eq. 23), and the concave upwards dash-dot curve is the total variance (23) as a function of trial k_B . **(d)** Observed spectrum (thick curve) and MLE fit (upper thin curve). Noise model is shown as the lower thin curve. Dash-dot line is least-squares fit to a power law plus noise used for likelihood ratio computation (section 4.5). See text for definitions of SNR, MAD, $\text{var}(Y)$, and LikelihoodRatio. MAD(Y) and Var(Y) in the legend are computed for the range of wavenumbers for which signal exceeds noise. **(a)** Histogram of the ratio of observed and theoretical spectra, and theoretical pdf for comparison. **(b)** Quantile-quantile plot of observed spectral ratio and theoretical pdf.

4.4 Error Bars on fitted parameters

For the case described in section 4.2, with measurements y_i distributed about the true value \mathcal{f} with a Gaussian pdf, the sum of squares $S \equiv \sum_{i=1}^N \frac{(y_i - a)^2}{2\sigma^2}$ and the log-likelihood function (14) are parabolic in a with curvature $\frac{N}{2\sigma^2}$. So the likelihood function is a Gaussian shape with variance σ^2/N .

Consider, again for Gaussian observations y_i , the MLE estimate of \mathcal{f} . By formula (15), this is the average of the y_i . For Gaussian y_i , the average is also distributed as a Gaussian, with mean \mathcal{f} and variance σ^2/N , the same shape as the likelihood function above. Thus, for Gaussian observations, the likelihood function and the estimated mean have the same shape pdf. This analogy between the likelihood function and the pdf of the estimated parameter allows us to use the curvature of the log-likelihood function to estimate the variance in any MLE-estimated parameter. This analogy is supported by the observation that the computed likelihood function $\exp(C11)$ is usually found to be a good match to a Gaussian function with the same curvature (see figure 5, lower left panel).

Priestley (1981, Section 5.2.2) derives from the Cramer-Rao inequality an expression for the variance of the Maximum Likelihood estimator (18) for k_B , in terms of the curvature of the log-likelihood function at the maximum:

$$\text{var}(k_B) \geq -\frac{1}{E(\partial^2 \ln(P)/\partial k_B^2)} \approx -\frac{1}{(\partial^2(C11)/\partial k_B^2)}. \quad (22.)$$

Thus the curvature of C11 at the peak gives a lower bound estimate for the standard error of the fitted parameter k_B , and corresponding error bounds on ϵ . This formula works for any parameter estimated via MLE fitting, and can be extended to estimate a covariance matrix for multiple parameter fits (Priestley, 1981).

4.5 Measures of misfit

There are many situations in which the observed spectrum simply has the wrong shape - one that fails to match any of the family of theoretical spectra to the expected accuracy. This may be because a data segment contains two or more mixing events with different ϵ (leading to multiple or indistinct cutoff), because the event is evolving and not in a stationary balance, or for some other reason. In such cases, the spectral fit to k_B should be

rejected, although integral variance measures such as χ_θ are probably valid. We present here two possible rejection criteria -- objective statistical measures that identify poorly fitting spectra. These should allow automated editing of processed data. We also discuss other goodness of fit measures that did not work very well.

PDF tests on Spectral ratio

When the theoretical and observed spectra match perfectly with the exception of statistical sampling error, the quantity $Y_i \equiv S_{obs}(k_i)/S_{th}(k_i)$ is distributed as d^{-1} times a χ_d^2 pdf, which has an expected value of 1 and an expected variance of $2/d$, where d is the equivalent number of degrees of freedom for the spectral estimator used to compute S_{obs} . A test of the statistical model (16) is, therefore, to test for this pdf. Figure 5(a) shows the observed histogram of Y_i , along with the χ_d^2 pdf for $d=6$ (the number of degrees of freedom appropriate to the spectral window used). The comparison is favorable, and certainly makes the point that the spectral errors are far from Gaussian. A more incisive test to see if dY_i has a χ_d^2 pdf is the quantile-quantile (Q-Q) plot in the upper right panel (c) of figure 5. A perfect match between theoretical pdf on the y-axis and the observed histogram on the x-axis would result in a 1:1 straight line. The Q-Q plot is quantified by the r^2 value of the straight-line fit, and/or the Komolgorov-Smirnov (K-S) test (Priestley, 1981, pp 480-481). The spectral fit is good but not perfect (see the low observed values in the range of 50-90 cpm in figure 5d), yet the Q-Q plot is quite close to a straight line, with an r^2 value of 0.97 and a K-S value of 0.08, indicating an acceptable fit. Looking at the histogram or the Q-Q plot requires manual editing, and the r^2 value from the Q-Q plot tended to be high (well above 0.9) even for rather poor fits, and so did not serve as a useful "rejection statistic". We believe that this is because the Y_i statistic tends to be dominated by the large number of spectral estimates from the noise-dominated wavenumbers. The sensitivity of these measures may be improved by restricting their calculation to wavenumbers for which the signal exceeds the noise, but we did not experiment with that restricted calculation.

Because $Y_i \cdot d$ should have a χ_d^2 pdf,

$$E\{\text{var}(Y_i)\} = 2 / d. \quad (23.)$$

Hence, when the spectrum fits perfectly we expect the variance of Y_i to equal $2/d$ and to increase from that value as the spectrum deviates from the theoretical shape. The "cost function" (23), computed for all wavenumbers as a function of k_B , is shown in figure 5,

lower left panel, along with the expected value. While the minimum occurs fairly close to the MLE-fit k_B , and is close to (in this case slightly below) the theoretical value, the minimum is very broad compared to the MLE maximum, suggesting that the error may be larger than the MLE error. This is found to be true using Monte Carlo tests in section 6.

We also computed the variance of Y_i over the range of wavenumbers for which the fitted spectrum (S_B) exceeds the noise (S_n) -- the value is shown in the legend of figure 5d. This seems to be a potentially useful indicator of spectral misfit (analogous to residual sum of squares) - a rejection criterion would be a number significantly larger than that expected. However, experience in using this measure on real data found it to not be as robust as a similar measure, the Mean Absolute Deviation (MAD), also computed over the wavenumber range where signal/noise exceeds 1.

$$MAD2 \equiv \frac{1}{n} \sum_{k_i=k_1}^{k_n} \left| \frac{S_{obs}}{S_{th}} - \left\langle \frac{S_{obs}}{S_{th}} \right\rangle \right|, \quad (24.)$$

In (24), k_n indicates the wavenumber where signal matches noise. Experience with both (23) and (24) on real SCAMP data leads us to prefer (24) as being more robust, with MAD2 values greater than 1.2 (for $d=6$) indicating unacceptable fits. (Note: We call this quantity MAD2 to distinguish it from a similar quantity MAD1, described in section 6.)

One advantage held by $\text{Var}(Y)$ is the theoretical result that for a perfect fit $\text{Var}(Y)$ should be $2/d$. We determined the equivalent "perfect fit" value for $MAD2(Y)$ by computing the expected value of $\left| \frac{\chi_d^2}{d} - \left\langle \frac{\chi_d^2}{d} \right\rangle \right|$ in a Monte-Carlo fashion. We found that an upper bound for $MAD2(Y)$ in the case of perfect fit is $\sqrt{2/d}$, the square root of $E(\text{Var}(Y))$. We take twice this value as our rejection level.

Spectral smoothing – the selection of d

Although formula (18) is valid for general dof, there appears to be little advantage to doing a lot of spectral smoothing. The situation is analogous to the task of fitting a straight line to, say, 20 points. One could group them in clumps of 5 and fit a straight line to the four averaged points. The statistical fitting error is not reduced, but some bias is introduced, due to lumping of different data points into single averaged values. For the Batchelor spectrum, with a sharp cutoff, spectral smoothing will smooth the cutoff, and therefore

may introduce bias into the fit. The advantage to smoothing is that more spectral smoothing will reduce the scatter of the observed spectrum and hence reduce the expected value of $\text{var}(Y_i)$ and $\text{MAD}(Y)$, making misfit spectra easier to detect.

Alternative Spectral Model and the Likelihood Ratio

The thin dash-dot line in figure 5d is a MLE fit of S_{obs} to a power law $S_{\text{th}} = A \cdot k^{-b} + S_n$, with A constrained to preserve the spectral variance. The idea is to compare the Batchelor spectrum fit (which has a sharp cutoff), to a simpler one with no cutoff - a simple power law. If the fit to the Batchelor spectrum does not provide a significantly better fit than the power-law fit, then that segment should be rejected as not having a clear cutoff. The two spectral fits are compared by computing the log-likelihood (18) for the power law fit, subtracting it from the log-likelihood for the Batchelor fit, and translating the difference to base 10 for easier interpretation. The likelihood ratio, shown in the legend of fig 5, is $\log_{10}(P_{\text{Batch}}/P_{\text{power}})$. Segments with the Batchelor fit being significantly more likely are accepted. Although the likelihood values from two different spectral models (each with one free parameter) cannot be compared with formal statistical validity, both likelihoods are computed with the same statistical model (eq. 18), using the same noise model, and a theoretical spectrum with one free parameter. The likelihood ratio appears to be a useful indicator of poor spectral fits. It appears from experience that ratios greater than 10^2 (i.e., \log_{10} values larger than 2) can be confidently accepted.

Comparison of rejection criteria

Of the rejection criteria that we tried, only $\text{Var}(Y)$, $\text{MAD}(Y)$ (both computed only over the wavenumber range where signal exceeds noise), and the likelihood ratio seemed to reliably indicate unacceptable spectral fits. We investigated and compared the behaviour of each on real SCAMP data sets, looking to see if one or the other measure seemed more reliable or robust than the others.

Figure 6 shows in the top row histograms and cumulative distribution functions for each of the measures, computed from a SCAMP data set. It is comforting to see that the rejection criteria we have settled on are not severe – only a relatively small fraction of the fits would be rejected (about 22% in this data set). The correspondence between the three measures is shown in the second row. It appears that $\text{Var}(Y)$ and $\text{MAD}(Y)$ are well-correlated (panel

f), although $\text{Var}(Y)$ is proportional to the cube of $\text{MAD}(Y)$, rather than the expected second power. This is likely due to $\text{Var}(Y)$ being more sensitive to outliers. The other measure, $\log_{10}(\text{likelihood ratio})$, is better correlated with $\text{MAD}(Y)$ than $\text{Var}(Y)$, although the scatterplots (d) and (e) are similar. We conclude that the most successful pair of rejection measures is $\text{MAD}(Y)$ and $\log_{10}(\text{likelihood ratio})$.




Figure 6. (a.) Histogram and cumulative distribution function of $\text{Var}(Y)$ from a data set with 5498 segments and $d=6$. Dotted vertical line shows theoretical "perfect fit" variance ($2/d = 1/3$), and vertical thin line shows our suggested rejection criterion (5.7). (b.) As for (a), but for $\text{MAD}(Y)$. Here the vertical dotted line is $\sqrt{2/d}$, an approximate upper bound for perfect fits verified by Monte-Carlo methods. The thin vertical line indicates our suggested rejection limit of $2\sqrt{2/d}=1.2$. (c.) As for (a), but for $\log_{10}(\text{likelihood ratio})$. The thin vertical line indicates our suggested rejection limit of 2. Note that values *larger* than 2 should be accepted. (d.) Scatterplot of $\log_{10}(\text{likelihood ratio})$ vs $\text{Var}(Y)$. Solid lines indicate suggested acceptance values for each measure in (d-f). (e.) Scatterplot of $\log_{10}(\text{likelihood ratio})$ vs $\text{MAD}(Y)$. (f.) Scatterplot of $\text{Var}(Y)$ vs $\text{MAD}(Y)$.

Summary of rejection criteria

For SCAMP Batchelor fits we suggest that a given data segment be rejected if any of the following criteria are true:

1. Integrated signal/noise ratio is less than 1.3.
2. $MAD2(Y)$ computed over the wavenumber range where signal exceeds noise, is larger than $2\sqrt{2/d}$.
3. $\text{Log}_{10}(\text{likelihood ratio})$ is less than 2.

In the data set of figure 6, these criteria result in the rejection of 22.2% of the data segments.

5. Further Examples of MLE fits.

Figure 7 shows a MLE fit to a data segment with a low signal level ($S/N=1.3$), but with a clear Batchelor spectrum. The fitted spectrum plus noise is a very good match to the observed spectrum, the likelihood function is a reasonable match to a Gaussian, and the estimated error in Batchelor wavenumber (about 9%) is larger than in figure 5. The spectral ratio (Y) is a good match to the theoretical pdf, and the Q-Q plot and it's associated statistics so indicate. The "rejection indicators" both indicate that the fit is acceptable: the Batchelor fit is 10^{24} times more probable than a power law fit, and the MAD is 0.38, which is small. The variance of Y is also small, and close to the theoretical value.




Figure 7. Example of real data fit with very low signal level. All indications point to a good fit. See figure 5 caption for detailed information on each panel.

Figure 8 shows a fit to a SCAMP data segment that is unacceptable, and is flagged as such by the Likelihood ratio. The deviation between the observed and theoretical spectra is somewhat larger than expected from time series theory. The variance of the spectral ratio (0.6) is significantly larger than the theoretical value of 0.33, and even more so when

computed over the range of wavenumbers where signal exceeds noise (1.88). The MAD is also moderately large, 0.78, and the Likelihood ratio computation states that a Batchelor spectral fit is $10^{-0.95}$ times more likely (i.e., about 10 times less likely) than a simple power law. This matches what the eye says – the spectrum is the "wrong shape". Yet the histogram and the Q-Q plot in the upper quadrants of the plot seem to say that the fit is acceptable. This is because these plots are dominated by the large number of data points in the wavenumber range for which noise exceeds signal.




Figure 8. Example of a "bad" fit – one for which the spectral shape is not an acceptable match to the Batchelor form for any value of k_B . . See figure 5 caption for detailed information on each panel.

6. Monte-Carlo Tests

For each of an array of k_B and χ_θ values covering a range suitable for SCAMP fitting, we generated 300 random samples of Batchelor spectra plus noise conforming to the statistical model (17.), with $d=6$ degrees of freedom. These represent data sets for which "good" fits should be possible, and for which we knew the answers. These were then analyzed by our fitting routines, using the MLE algorithm and a variety of least-squares and cost-function algorithms, and the bias (estimated minus true value) and standard deviation, both normalized by the true value of k_B , were computed and expressed in percent. The purpose was to:

1. test the MLE method for bias, and to see if the predicted standard deviation of the estimated k_B agreed with the estimated fit error from eq. (22), and
2. to compare the MLE fit bias and standard deviation with those from the other fit algorithms.

In figure 9a, the gray-scale contours show the computed standard deviation of the k_B estimates. These should be compared with the standard error of k_B estimated from formula (22) in figure 9b. The agreement is excellent except in the lower right-hand corner, where the fit becomes noisy due to the coincidence of the noise and Batchelor peaks.

Figure 9a also shows the estimated bias in percent of the MLE estimate of k_B (heavy black contours). It can be seen that the bias is, in practical terms, zero, with the exception of the lower right corner, where all fitting methods had a difficult time.




Figure 9. (a) Standard deviation in percent of the MLE estimated k_B as a function of the "true" k_B and χ_θ , shown in grey-scale with white contours. The MLE bias in percent is shown as heavy black contours. (b) Estimated standard deviation of the MLE estimated k_B as a function of the "true" k_B and χ_θ , computed from the average of 300 realizations of eq. (22). (i.e., from the average curvature of the 300 log-likelihood functions.) Both panels show k_B scaled by the Sampling Nyquist wavenumber and χ_θ (expressed as signal/noise ratio) on the top and right-hand axes respectively.

The following least-squares and cost-function estimates for k_B were made using the Monte-Carlo data sets. In each case, the improvements described in section 4.1 were

implemented: the theoretical spectrum included instrument noise, and χ_θ was constrained by the observed variance less noise, leaving a single-parameter fit.

1. C11 (MLE fit)
2. C3 (eq. 6.) Least squares minimization in log space.
3. LS. Weighted least squares minimization on Batchelor plus noise spectra. (As for C1, eq. 5, with S_B replaced by S_{th})
4. MAD1: Mean value of $|Y-1|$ minimized This is a robust version of LS.
5. Var(Y) minimized. (eq. 23.)
6. MAD2: MAD($Y - \langle Y \rangle$) minimized. Like MAD1, but use the computed mean of Y rather than its expected value of 1(Eq. 24.). This is a robust version of Var(Y).

The performance of these estimators is compared to that of the MLE in figure 10. It can be seen that:

- C11 (MLE) has the smallest bias (less than 1/4%), and the smallest standard deviation overall. (Figure 10a)
- C3 has, in most of the $k_B - \chi_\theta$ space, similar standard deviation to the MLE estimate. However, in the high k_B low χ_θ region it has much larger standard deviations, up to about 100%. In the low k_B high χ_θ corner the standard deviation is also high. In addition, C3 has a fairly uniform bias of about -2%. (Figure 10b)
- LS suffers from a 5% bias and rather large standard deviation at the low and high ends of the k_B scale. (Figure 10c)
- MAD1 minimization is only slightly inferior in performance to MLE in standard deviation, but is biased about 2% high. (Figure 10d)
- Var(Y) minimization has unacceptably high standard deviation and bias. (Figure 10e)
- MAD2 minimization is very noisy and biased, similar to Var(Y) minimization. (Figure 10f)

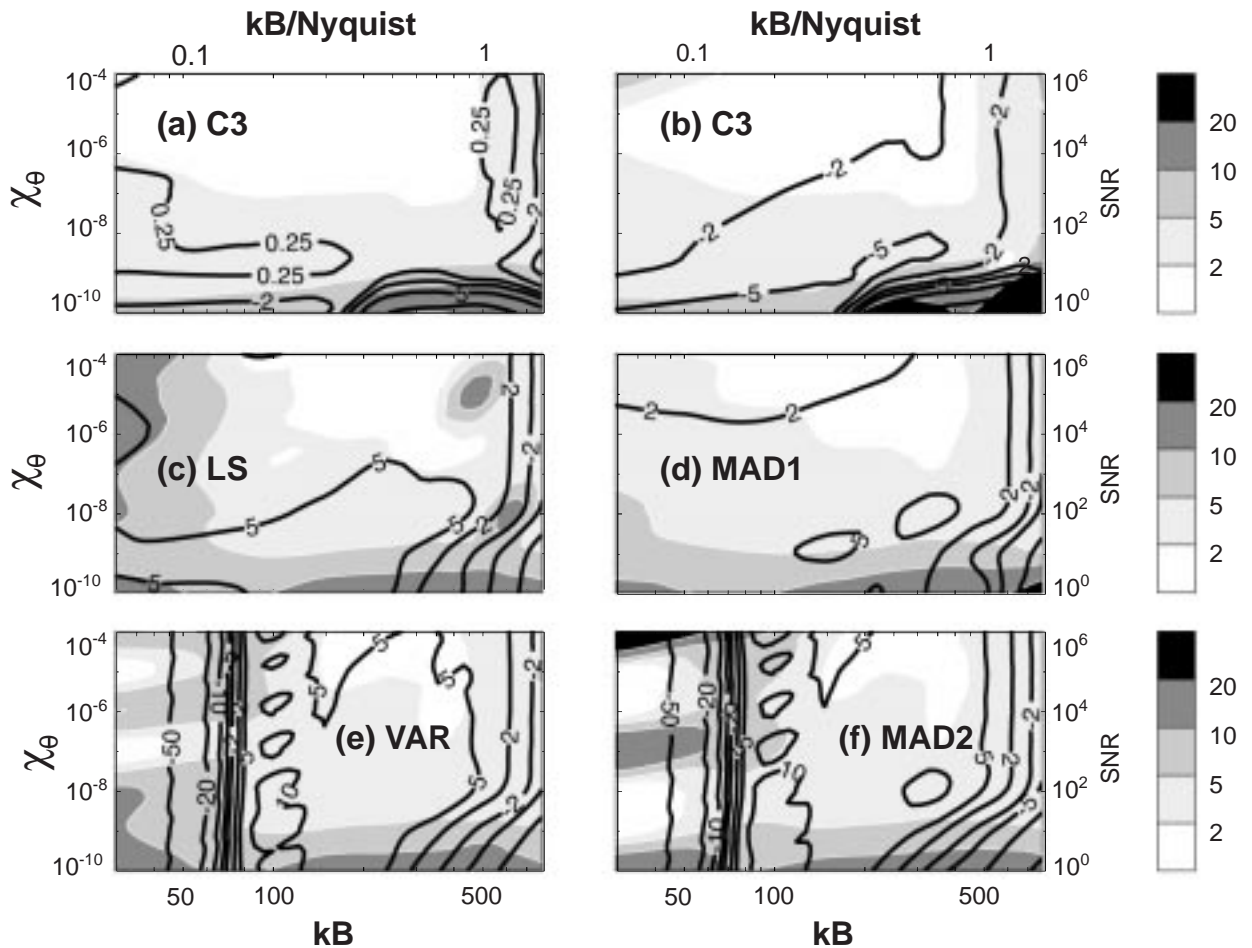


Figure 10 (a) Standard deviation, in percent, of the MLE estimated k_B as a function of the "true" k_B and χ_θ , shown in grey-scale with white contours. The MLE bias in percent is shown as heavy black contours. (b) As for (a), but for log-space least-squares fitting, C3. (c) As for (a), but for weighted least squares. (d) As for (a), but for MAD1 minimization. (e) As for (a), but for Variance(Y) minimization. (f) As for (a), but for MAD2 minimization. All panels show k_B scaled by the sampling Nyquist wavenumber and \square (expressed as signal/noise ratio) on the top and right-hand axes respectively.

The conclusion is that, while some estimators work nearly as well as MLE (MAD1 in particular), none perform quite as well or are as unbiased as MLE, and the MLE error

estimate (22) is supported by the Monte-Carlo tests. MAD1 deserves further study, as it's robustness properties could make it less sensitive to "poorly shaped" spectra.

7. Discussion

A method to parametrically fit theoretical spectra to observations was developed, and demonstrated with fits to the Batchelor spectrum. The useful techniques include:

1. Explicit incorporation of an instrumental noise spectrum in the fitted model. This helps the fit routines to automatically ignore regions dominated by instrument noise.
2. Use of available constraints to reduce the dimension of the fitted parameter space.
3. Direct application of the Maximum Likelihood technique to yield efficient, unbiased spectral fits. This is important because the statistical error in spectral estimation is not Gaussian, and so least squares is not the most appropriate technique.
4. The curvature of the log-likelihood function near its peak gives an estimate of the error in the fitted parameters.
5. We develop and discuss several measures of spectral misfit, useful for rejecting data segments for which the best fit spectrum is unacceptable. These allow automated processing of large numbers of spectra, for which manual editing is tedious.

The MLE technique was compared with other least-squared and cost-function techniques using Monte-Carlo generated data sets that satisfy the expected statistical distribution. The standard deviation and the bias of fitted Batchelor wavenumber are estimated by repeated fits for each of several values of true Batchelor wavenumber and signal level. The MLE technique is found to be the best one, with bias typically less than 1/4%, the smallest standard deviation overall, and estimated fit error in agreement with the standard deviation.

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), and Canadian GLOBEC. It is a pleasure to acknowledge collaborative discussions with and generous advice from Neil Oakey and Dan Kelley. Mike Head kindly provided the SCAMP noise routine, and the algorithms behind much of the SCAMP processing system. Adam Webber, supported by NSERC and the Nova Scotia links program, performed many of the early computations during the development of the fitting algorithms. A. Anis was partly supported by a grant from the German-Israel Foundation.

Appendix A MLE Spectral fit Pseudo-Code Outline

MATLAB versions of the m-files used in this paper for MLE spectral fitting, and also for computing Batchelor spectra, may be requested directly from the authors.

1. Compute noise spectrum vs frequency (call instrument noise subroutine); change to wavenumbers using drop speed of instrument.
2. Compute χ_θ using observed spectral variance less noise variance (eq. 9).
3. Loop for trial k_B :
 4. Compute trial Batchelor spectrum, eq. 4 (call Batch_Spec subroutine).
 5. Add computed noise spectrum from line 1.
 6. Calculate log likelihood function C11 at each wavenumber (eq. 18). (Call Chi2pdf transformation function). Sum C11 over all wavenumbers.
 7. Compute $\text{Var}(Y = S_{\text{obs}}/S_{\text{th}}$ summed over all wavenumbers (eq. 23).
8. End of Loop.
9. Compute location of maximum in C11. This is the best-fit k_B .
10. Compute curvature in C11 function at maximum.
11. Use eq. 22 to compute standard error in k_B .
12. Plot $\exp(C11)$ as likelihood function. Plot Gaussian curve (eq. 10) with same standard deviation and peak. Plot error bars corresponding to line 11. Plot line 7 ($\text{Var}(Y)$ versus trial k_B and theoretical minimum value $2/d$).
13. Compute ϵ and upper/lower bounds on ϵ from k_B results in lines 9, 11 (eq. 3).
14. Compute integrated S/N ratio from lines 1 and 2.
15. Compute $Y = S_{\text{obs}}/S_{\text{th}}$ at each k , for k_B of best fit.
16. Compute $\text{Var}(Y)$ (eq. 23) over wavenumbers for which S_B exceeds S_n .
17. Compute $\text{MAD}(Y-1)$ (eq. 24) over wavenumbers for which S_B exceeds S_n .
18. Loop for power law fit:
 19. Repeat steps 5-7 using power law instead of Batchelor spectrum. Constrain the variance of power law plus noise spectrum to match the observed variance.
20. End of power law loop.
21. Compute location of C11 maximum to find best power law.
22. Take difference of $\log(\text{likelihood})$ from lines 9 and 21. Convert to \log_{10} to get likelihood ratio.
23. Plot best fit spectrum, noise spectrum, and observed spectrum. Superpose best-fit power law spectrum. Annotate this plot with results from lines 2,13,9,11,14, 16, 17.
24. Plot histogram of $d*Y$, and χ_d^2 pdf for comparison.
25. Plot Q-Q plot of $d*Y$ compared with χ_d^2 cdf. Compute r^2 value and K-S test value.

Subroutines:

Instrument noise routine – computes S_n as a function of frequency

Batch_Spec subroutine – computes Batchelor spectrum as function of k given χ_θ and k_B .

Chi2pdf subroutine – computes the χ_d^2 pdf. Since this is simply an exponential function and a power, with a normalization constant that involves a Gamma function of d , it may save time to compute the Gamma function once and store the result, leaving only the power and exponential functions to be evaluated.

References

Bartlett, M. S. (1954). Problèmes de l'analyse spectrale des séries temporelles stationnaires. *Publ. Inst. Statist. Univ. Paris*, **III-3**, 119-134.

Batchelor, G.K. (1959) Small-scale variation of convected quantities like temperature in turbulent fluid. Part 1. General discussion and the case of small conductivity. *J. Fluid Mech.* **5** 113-133.

Brillinger, David R. (1985). Fourier Inference: Some methods for the Analysis of Array and Nongaussian Series Data. *Water Resources Bulletin*, **21**, No. 5, 743-756.

Dillon, Thomas M., and D. R. Caldwell, (1980). The Batchelor Spectrum and Dissipation in the Upper Ocean. *Journal of Geophysical Research*, **85**, No. C4, 1910-1916.

Flannery, B. P., Press, W. H., Teukolsky, S. A., Vetterling, W. T. (1995). *Numerical Recipes in C*. Cambridge University Press, printed in the United States.

Jenkins, G..M., and D. G. Watts (1968) Spectral Analysis and its Applications. Holden-Day, San Francisco, 525 pp.

Leuketina, D.A., and J. Imberger (1999). Determining Kinetic Energy Dissipation from Batchelor Curve Fitting. *J. Atmos. Ocean. Tech.*, submitted.

Moum, J. N., 1996: Efficiency of mixing in the main thermocline. *J. Geophys. Res.*, **101**, 12 057–12 069.

Oakey, N. S., 1982: Determination of the rate of dissipation of turbulent energy from simultaneous temperature and velocity shear microstructure measurements. *J. Phys. Oceanogr.*, **12**, 256–271.

Osborn, T. R., 1980: Estimates of the local rate of vertical diffusion from dissipation measurements. *J. Phys. Oceanogr.*, **10**, 83–89.

_____, and C. S. Cox, 1972: Oceanic fine structure. *Geophys. Fluid Dyn.*, **3**, 321–345.

Priestley, M.B.(1981). *Spectral Analysis and Time Series*. Academic Press, London

Ruddick, B.R., D. Walsh, and N. Oakey, 1997. Variations in Apparent Mixing Efficiency in the North Atlantic Central Waters. *J. Phys. Oceanogr.*, **27**, 2589-2605.

St. Laurent, L., and R. W. Schmitt, 1999. The contribution of salt fingers to vertical mixing in the North Atlantic Tracer Release Experiment. *J. Phys. Oceanogr.*, in press.

Figure Captions

Figure 1 - Straining field in a mean temperature gradient, and "micro-front" produced by the convergence.

Figure 2: Batchelor Spectrum for various k_B , including estimated SCAMP instrumental noise level ($3 \cdot 10^{-7}$), with χ_θ constrained according to eq. (9). The k_B value corresponding to each curve is indicated with a plus, and the approximate corresponding dissipation level, ϵ , is shown by the second logarithmic scale below the k-axis. The effects of changing χ_θ and ϵ on the spectrum are indicated by arrows. An observed spectrum is shown.

Figure 3 - SCAMP fit example, using C1, equation 4. MLE fit (section 4) is shown for comparison. Note the cutoff for the SCAMP fit (C1) to avoid the noise-dominated high-wavenumber region. Same observational spectrum is used in figures 2 and 3.

Figure 4. (a) Histogram of observations of a quantity with true value $\mathcal{f} = 2$, and Gaussian error. (b) Hypothesized Gaussian pdf with a ~ 1.7 . (c) The $\ln(\text{probability})$ of our observations as a function of the hypothesized value of a, showing how $\ln(P)$ is maximized as a is changed.

Figure 5. MLE fit to observed spectrum. (c) Relative likelihood function $\exp(C11)$ (eq. 18) normalized to a maximum of 1 versus trial k_B , and Gaussian function with same second derivative at maximum. Horizontal error bars indicate 1 standard deviation according to eq. (22) by inner vertical bars, plus approx. 95% confidence limits by outer vertical bars. Horizontal dash-dot line indicates the theoretical level of the variance (eq. 23), and the concave upwards dash-dot curve is the total variance (23) as a function of trial k_B . (d) Observed spectrum (thick curve) and MLE fit (upper thin curve). Noise model is shown as the lower thin curve. Dash-dot line is least-squares fit to a power law plus noise used for likelihood ratio computation (section 4.5). See text for definitions of SNR, MAD, $\text{var}(Y)$, and LikelihoodRatio. MAD(Y) and Var(Y) in the legend are computed for the range of wavenumbers for which signal exceeds noise. (a) Histogram of the ratio of

observed and theoretical spectra, and theoretical pdf for comparison. **(b)** Quantile-quantile plot of observed spectral ratio and theoretical pdf.

Figure 6. (a.) Histogram and cumulative distribution function of $\text{Var}(Y)$ from a data set with 5498 segments and $d=6$. Dotted vertical line shows theoretical "perfect fit" variance ($2/d = 1/3$), and vertical thin line shows our suggested rejection criterion (5.7). (b.) As for (a), but for $\text{MAD}(Y)$. Here the vertical dotted line is $\sqrt{2/d}$, an approximate upper bound for perfect fits verified by Monte-Carlo methods. The thin vertical line indicates our suggested rejection limit of $2\sqrt{2/d}=1.2$. (c.) As for (a), but for $\log_{10}(\text{likelihood ratio})$. The thin vertical line indicates our suggested rejection limit of 2. Note that values *larger* than 2 should be accepted. (d.) Scatterplot of $\log_{10}(\text{likelihood ratio})$ vs $\text{Var}(Y)$. Solid lines indicate suggested acceptance values for each measure in (d-f). (e.) Scatterplot of $\log_{10}(\text{likelihood ratio})$ vs $\text{MAD}(Y)$. (f.) Scatterplot of $\text{Var}(Y)$ vs $\text{MAD}(Y)$.

Figure 7. Example of real data fit with very low signal level. All indications point to a good fit. See figure 5 caption for detailed information on each panel.

Figure 8. Example of a "bad" fit – one for which the spectral shape is not an acceptable match to the Batchelor form for any value of k_B . . See figure 5 caption for detailed information on each panel.

Figure 9. (a) Standard deviation in percent of the MLE estimated k_B as a function of the "true" k_B and χ_θ , shown in grey-scale with white contours. The MLE bias in percent is shown as heavy black contours. (b) Estimated standard deviation of the MLE estimated k_B as a function of the "true" k_B and χ_θ , computed from the average of 300 realizations of eq. (22). (i.e., from the average curvature of the 300 log-likelihood functions.) Both panels show k_B scaled by the Sampling Nyquist wavenumber and χ_θ (expressed as signal/noise ratio) on the top and right-hand axes respectively.

Figure 9. (a) Standard deviation in percent of the MLE estimated k_B as a function of the "true" k_B and χ_θ , shown in grey-scale with white contours. The MLE bias in percent is shown as heavy black contours. (b) Estimated standard deviation of the MLE estimated k_B as a function of the "true" k_B and χ_θ , computed from the average of 300 realizations of eq. (22). (i.e., from the average curvature of the 300 log-likelihood functions.) Both panels show k_B scaled by the Sampling Nyquist wavenumber and χ_θ (expressed as signal/noise ratio) on the top and right-hand axes respectively.

