Updated review of paper submitted to Ocean Science: “On the resolutions of ocean altimetry maps” by Ballarotta, Ubelmann, Pujol, Taburet, Fournier, Legais, Faugere, Chelton, Dibarboure, and Picot

J. Thomas Farrar
Department of Physical Oceanography
Woods Hole Oceanographic Institution

April 3, 2019

Note: This was the review of the first version of the manuscript. It was submitted directly to the authors during the revision process and addresses some issues that were not made clear in the first version of the review posted on the OSD website. This is the version the authors responded to in their revision.

This paper presents a new method for assessing the spatial and temporal resolution of a gridded data product and applies the method to the DUACS altimetry product. The paper is well written and the undertaking is worthwhile. The method for assessing resolution is based on estimation of the coherence between independent data and the gridded product. I have a great deal of respect for these authors, and so it is with a lot of discomfort and hesitation that I must say that the method used in the paper is fundamentally flawed, and I strongly recommend against publication. In what follows, I will focus on what is wrong with the proposed method of assessing resolution, because that is essentially the only problem I have with the paper. I would like to express my sincere regret to the authors that I must argue so strongly against publication of their paper– I hope they find this review helpful rather than offensive. I also apologize to all involved for the lengthy review! I felt it was necessary to include some examples to illustrate what is wrong with the methodology. If my assessment is incorrect, I would be happy to be corrected. (In order to begin to convince me, the authors would need to first produce an example in 1-D of how the coherence-based measure of resolution can quantify the filtering properties of a filtering operation performed on the data– as shown in Equation 6 of the paper and below in Equation 14, this looks like an impossible task.)

The method is based on a conceptual model of a linear single-input/single-output system (in the terminology of the book by Bendat and Piersol that the authors cite) that is used to interpret the coherence of the data product with independent data. The conceptual linear system has no noise on the measurement of the system output but does have noise on the measurement of the system input. As described by the authors and by Bendat and Piersol (2010, p. 185), there is an input signal u(t) that goes through a linear system (the mapping) to produce an output signal y(t). The output y(t) is known perfectly (without measurement noise), but the input u(t) is not known perfectly– only a noisy measurement of it is available, x(t)=u(t)+m(t). For brevity, I will refer to this as a ”Case 2” model for interpretation of the coherence. The proposed measure of resolution is the wavenumber (or frequency) at which the coherence of the mapped field with an independent 1-D data record is equal to 0.5; in the Case 2 model, this is the wavenumber at which the signal-to-noise ratio is equal to 1. In my earlier version of this review, I may have been confused about which signal the authors consider as the ‘input’ and which they consider as the 'output'– this distinction does not really matter for the arguments I present below.

The resolution capability of a mapped field should depend on (i) the noise and signal levels in the input data, (ii) the sampling of the data, and (iii) the manipulations performed on the data during the mapping (e.g., filtering). The coherence method addresses (i) but not (iii). It isn't clear whether the method addresses (ii), but the fact that the method can tell us nothing about (iii) means it cannot be a useful measure of resolution by itself.

There are two problems:

1. The most basic problem is that, in 1-D, the coherence should be completely unaffected by the mapping and its filtering and instead depends only on the signal-to-noise ratio of the raw input data. This can be clearly seen from Eqn 6 of the paper, but I have included a derivation below to make this more clear. The derivation shows that the method cannot yield any useful information about the filtering in a 1-D example. The method can give us information about the SNR in the input data, but not about the filtering.

2. A secondary problem arises when the method is applied by computing the coherence between a 1-D sample of SSH against the mapped field to estimate the resolution in that direction. Take the example of computing coherence between independent along-track data and the mapped SSH. The along-track coherence is affected by across-track filtering and temporal filtering, but not by along-track filtering. It thus seems like a grave misinterpretation to say that the 1-D coherence reveals something meaningful about the along-track resolution.

I am going to focus most of this review on the first problem. Below, I give a general derivation that shows that, in a 1-D problem with noise and filtering, the coherence cannot provide a useful measure of resolution. I follow that with 1-D examples, with and without noise. Finally, I include some 2-D examples to illustrate the second problem listed above.

1 The coherence method in 1-D

To see that the coherence-based measure of resolution does not tell us anything directly about the filtering properties of DUACS mapping system, consider a 1-D mapping on a uniform grid. Let the measured SSH be $\tilde{x}{=}x + n$, where x is the true SSH and n is the measurement noise. We estimate the SSH as some linear combination of the measurements:

$$h_n = \sum_{m=-\infty}^{\infty} w_m \tilde{x}_{n-m}. \tag{1}$$

The weighting function w specifies the particular linear combination, and could represent any linear mapping; the weighting function is assumed to be nonzero over only a finite extent, so the summation indices do not actually extend to infinity. Equation 1 is a convolution operation, which we write as,

$$h = w * \tilde{x} = w * (x + n), \tag{2}$$

where $*$ is the discrete convolution operator. Fourier transforming both sides and using the convolution theorem,

$$\hat{h}(k) = \hat{w}(k)\big(\hat{x}(k) + \hat{n}(k)\big), \tag{3}$$

where the hats indicate the Fourier coefficients. $\hat{w}$ is important because it specifies the filtering properties of the mapping– it is the filter transfer function that relates the Fourier transform of the mapped data to the Fourier transform of the input data. This follows directly from Equation 3.

In the paper, the mapped field is compared to independent data (data withheld from the mapping), which we will represent as $y = x + n_0$, where y is the measured SSH and n_0 is the noise on that independent record. The Fourier transform of the independent SSH record is,

$$\hat{y}(k) = \hat{x}(k) + \hat{n_0}(k). \tag{4}$$

No filtering is applied to this independent record.

The squared coherence between the mapped SSH h and the independent SSH record y is,

$$\gamma^2 = \frac{\langle \hat{y}^* \hat{h} \rangle^2}{\langle \hat{y}^* \hat{y} \rangle \langle \hat{h}^* \hat{h} \rangle}, \tag{5}$$

where the superscript asterisk denotes the complex conjugate. To estimate the coherence and gain, we must estimate $\langle \hat{y}^* \hat{y} \rangle$, $\langle \hat{h}^* \hat{h} \rangle$, and $\langle \hat{y}^* \hat{h} \rangle$. They are given by,

$$\langle \hat{y}^* \hat{h} \rangle = \hat{w} \big(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{x} \rangle + \langle \hat{x}^* \hat{n} \rangle + \langle \hat{n}_0^* \hat{n} \rangle \big), \tag{6}$$

$$\langle \hat{y}^* \hat{y} \rangle = \langle \hat{x}^* \hat{x} \rangle + \langle \hat{x}^* \hat{n}_0 \rangle + \langle \hat{n}_0^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{n}_0 \rangle, \tag{7}$$

$$\langle \hat{h}^* \hat{h} \rangle = \hat{w}^2 \big(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{x}^* \hat{n} \rangle + \langle \hat{n}^* \hat{x} \rangle + \langle \hat{n}^* \hat{n} \rangle \big). \tag{8}$$

(Note that we are assuming the mapping operation does not introduce a phase shift, and thus that $\hat{w}$ is real.) Now, if we assume the noise is uncorrelated with the true SSH and that the noise between different passes/instruments is uncorrelated, then,

$$\langle \hat{x}^* \hat{n} \rangle = \langle \hat{x}^* \hat{n}_0 \rangle = \langle \hat{n}_0^* \hat{n} \rangle = 0. \tag{9}$$

Then,

$$\langle \hat{y}^* \hat{h} \rangle = \hat{w} \langle \hat{x}^* \hat{x} \rangle, \tag{10}$$

$$\langle \hat{y}^* \hat{y} \rangle = \langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{n}_0 \rangle, \tag{11}$$

$$\langle \hat{h}^* \hat{h} \rangle = \hat{w}^2 \left(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}^* \hat{n} \rangle \right). \tag{12}$$

Now, we can substitute into the coherence expression to obtain,

$$\gamma^2 = \frac{\langle \hat{y}^* \hat{h} \rangle^2}{\langle \hat{y}^* \hat{y} \rangle \langle \hat{h}^* \hat{h} \rangle} = \frac{\langle \hat{x}^* \hat{x} \rangle^2}{\big(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{n}_0 \rangle \big) \big(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}^* \hat{n} \rangle \big)}. \tag{13}$$

If we assume that the noise in the input data has the same spectrum as the noise in the independent data ($\langle \hat{n}_0^* \hat{n}_0 \rangle = \langle \hat{n}^* \hat{n} \rangle$), this simplifies somewhat to,

$$\boxed{\gamma^2 = \frac{\langle \hat{x}^* \hat{x} \rangle^2}{\big(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}^* \hat{n} \rangle \big)^2}.} \tag{14}$$

We can immediately see two important things: (1) when the noise variance equals the true SSH variance, the squared coherence has a value of 0.25 (not at 0.5 as assumed in the paper), and (2) the coherence is independent of the Fourier transform of the weighting function ($\hat{w}$), which is the quantity that contains all of the information about how the mapping filters the data. This derivation clearly shows that, with or without noise, the coherence tells us nothing about the filtering in the 1-D example. Based on this, it should be impossible to infer anything about the smoothing in a 1-D example from the coherence. Note that, if either the noise on the independent data or the noise in the input data (prior to mapping) is assumed to be zero, we recover the situation assumed in the paper, in which the squared coherence is 0.5 when the SNR=1.

The transfer function between the mapped SSH h and the independent SSH record y is,

$$H = \frac{\langle \hat{y}^* \hat{h} \rangle}{\langle \hat{y}^* \hat{y} \rangle} = \frac{\hat{w} \langle \hat{x}^* \hat{x} \rangle}{\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{n}_0 \rangle}. \tag{15}$$

If the noise is small relative to the true SSH signal, the transfer function approaches the filter transfer function of the mapping. (If we used the actual input data instead of the independent, withheld data, the transfer function would be identical to the filter transfer function of the mapping.)

For completeness, we can also write down the ratio of the two spectra:

$$\frac{\langle \hat{h}^* \hat{h} \rangle}{\langle \hat{y}^* \hat{y} \rangle} = \frac{\hat{w}^2 \left(\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}^* \hat{n} \rangle \right)}{\langle \hat{x}^* \hat{x} \rangle + \langle \hat{n}_0^* \hat{n}_0 \rangle}. \tag{16}$$

If the noise in the input data has the same spectrum as the noise in the independent data ($\langle \hat{n}_0^* \hat{n}_0 \rangle = \langle \hat{n}^* \hat{n} \rangle$), this simplifies to a nice result,

$$\frac{\langle \hat{h}^* \hat{h} \rangle}{\langle \hat{y}^* \hat{y} \rangle} = \hat{w}^2. \tag{17}$$

The above derivation was developed for a 1-D situation, but it would come out exactly the same if we have a multidimensional mapping, a multidimensional Fourier transform, and compute the coherence against a multidimensional version of the withheld data. In the paper, there is only 1-dimensional data to compare to the multidimensional mapping, and this adds an additional difficulty related to problem (2) mentioned in the introduction.

2 Examples to illustrate behavior of proposed measure of resolution

I have prepared four examples to illustrate the behavior of the proposed measure of resolution and to compare it to more commonly used measures (the gain and the spectral ratio method). For the 1-D case, there are examples with and without noise. The examples also have no sampling errors. The examples thus focus on what the proposed measure of resolution can tell us about the filtering properties of the mapping algorithm. The four examples are: (1) a 1-D case without noise, (2) and 1-D case with noise, (3) a 2-D case with a white spectrum (no noise), and (4) a 2-D case with a red spectrum (no noise). In all four examples, I mapped the data using three different smoothers, a Gauss-Markov smoother (also known as optimal interpolation), a Gaussian weighted average smoother, and a quadratic loess smoother. At least a few of the authors are familiar with these smoothers, and the exact details of the smoothers are not important. The loess and Gaussian smoother parameters were chosen such that they filter the data with a 25-km half-power filter cutoff in one spatial direction (nominally the “along-track” direction). The Gauss-Markov smoother uses a Gaussian autocovariance function and assumes the measurement contains a small amount of white noise, and it has similar but not identical filtering properties (with an autocovariance function spatial decay timescale of 25 km). In the 1-D example with noise, I also varied the filter cutoff wavelength to examine the sensitivity of the coherence to the filtering.

2.1 1-D example, without noise

In the first example, the input signal is a random realization of a process having a red spectrum (k^{-2} power law), sampled on a uniform spatial grid. The input data were "mapped" (smoothed) to the same grid as the sampling positions. There is no measurement noise and no sampling error. In this case, the only thing limiting the resolution of the mapped fields is the filtering inherent in the mapping algorithm. One could imagine this as a case where there is a single along-track pass of data, and they have been mapped to a regular spatial grid (along-track).

Figure 1 (top) shows the spectrum of the input data and of the three mapped fields. The mapped fields have less variance at high wavenumbers because of the filtering inherent in the mapping.

The filtering inherent in the mapping is quantified more directly by computing the cross-spectral gain (or relative amplitude of variability coherent between the input data and the mapped data). The gain is equivalent to the magnitude of the filter transfer function of the filters, and is shown in Figure 1 (middle panel). The Gaussian smoother has a filter transfer function that resembles a Gaussian function (as is expected, because the Fourier transform of the spatially truncated Gaussian weighting function is approximately a Gaussian). The loess smoother has, as expected, a steeper filter roll-off (meaning it decreases from one toward zero more abruptly near the half-power point) with a small but noticeable filter sidelobe at about twice the half-power cutoff wavenumber (Schlax and Chelton, 1992, their Figure 1). There is another, even smaller sidelobe of the loess filter that has a maximum value near wavenumbers of 0.12-0.13. The Gauss-Markov smoothed estimate has a very steep roll-off, and it has a half-power filter cutoff that is at a slightly higher wavenumber than the loess and Gaussian smoothed fields. All of these same features can be clearly seen in the spectra, as well (upper panel of Figure 1).

Based on conventional understandings of the term "resolution", we would say that the resolution, as defined by the half-power point of the filtering, is highest in the Gauss-Markov mapping and is the same in the loess and Gaussian weighted average mappings. (One might reasonably argue the resolution of the Gaussian mapping is worse because it noticeably attenuates wavelengths longer than the half-power cut-off wavelength.) For example, this perspective is similar to the definition of resolution used by the SWOT project when they state the raw (downlinked) resolution of the in-swath SWOT data will be 1-km resolution and 500-m posting– the onboard processor will filter the data, which fundamentally limits the resolution.

Now, if we turn our attention to the squared coherence (bottom panel of Figure 1), and apply the authors' coherence-based definition of resolution, we would reach the opposite conclusion: the coherence with the Gaussian smoother has higher values at high wavenumbers, which the authors would interpret as superior resolution. The squared coherence is almost one, even at wavenumbers three times larger than the half-power point, where the variance is more than 100 times less than the actual (unfiltered) variance. Under the definition of resolution proposed by the authors, the resolution of the mapping would be 6.5km wavelength for the Gaussian mapping, 14.3km wavelength for the loess mapping, and 20km wavelength for the Gauss-Markov mapping. The maximum possible resolution is 6km wavelengths because the sample spacing in this example

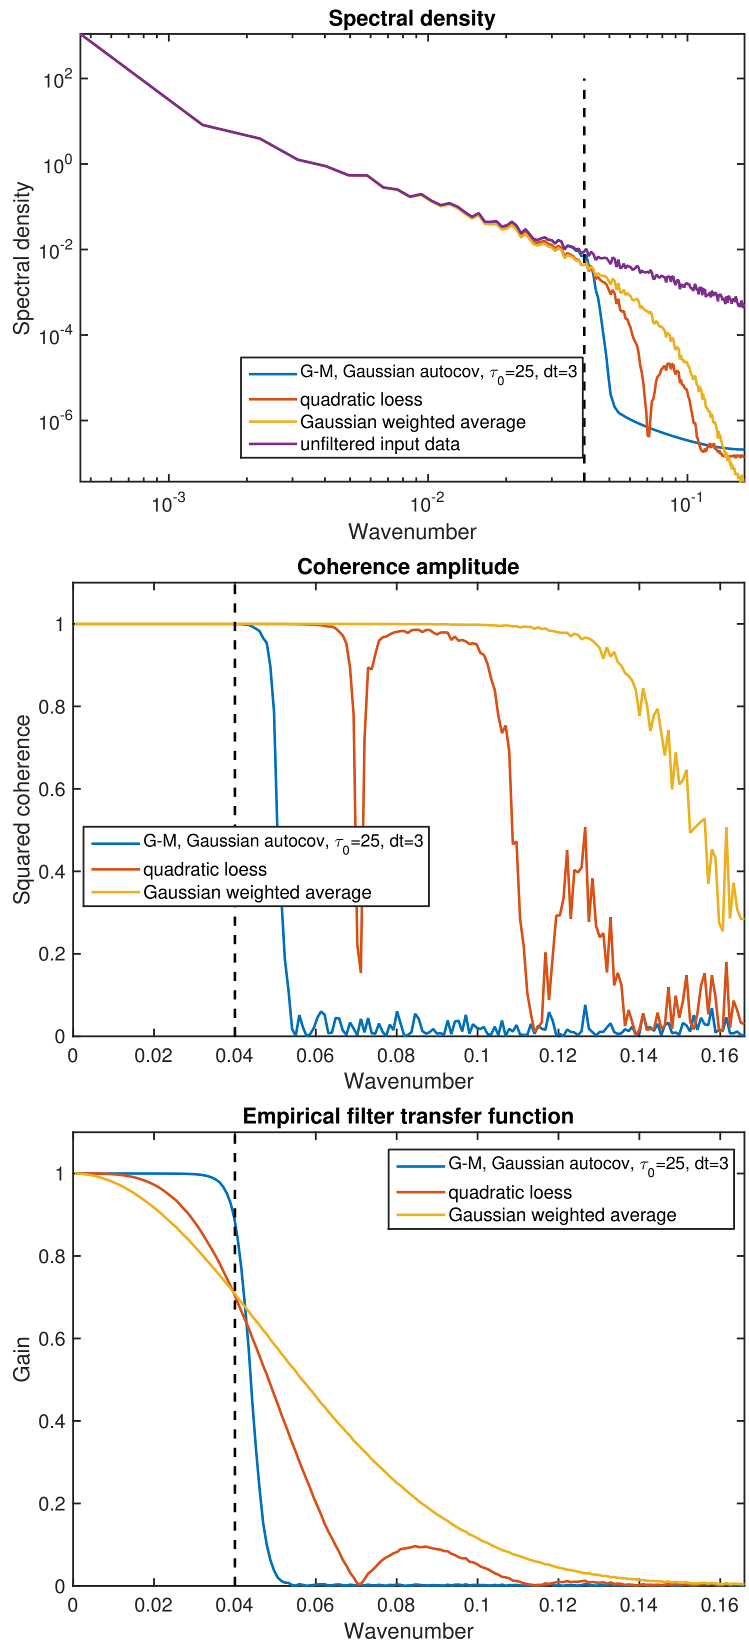

Figure 1: 1D example without noise. Upper panel: Spectra of input signal and mapped fields. Middle panel: squared coherence of the mapped field with the input signal. Lower panel: spectral gain computed between the input signal and the mapped field. The black dashed line marks the theoretical half-power wavenumber of the loess and Gaussian smoothers.

is 3km.

The coherence in this noise-free example should have been one at all wavenumbers for all three mappings (Equation 14). I believe the reason the coherence is not one has to do with numerical errors (such as roundoff errors) that occur at the higher wavenumbers where the amplitude of the filtered signal approaches zero. Since the filter transfer function of the Gaussian has the slowest rolloff, numerical issues only gradually become a problem. They abruptly become a problem for the Gauss-Markov filter because its filter transfer function drops abruptly to zero through the filter cutoff wavenumber. The numerical issues for the loess smoother become problematic at the wavenumbers that correspond to nodes between the sidelobes, and less of a problem near the extrema of the sidelobes. The existence of two sidelobes can be seen in the figure. (Consistent with the idea that numerical errors are influencing the high-wavenumber coherence, I found that the coherence at the high wavenumbers changed when I did the following: (1) force Matlab to use a different FFT algorithm, which should affect the error of the FFT, and (2) reduced the number of points in the time series, which should decrease the error of the FFT. I suspect that most of the numerical errors are in the mapping.)

In summary, this first example shows that the coherence-based measure of resolution fails to be useful in a simple, but relevant, example. The example is relevant because the filtering of the mapping procedure is an important aspect of the resolution of the mapped field, and the coherence criterion tells us almost nothing about this filtering.

2.2 1-D example, with noise

This example is identical to the previous one, but independent realizations of random noise were added to both the independent data and to the input data. We did this for two different noise values– in one case the noise spectrum intersects the spectrum of the true SSH at lower wavenumbers than the filter half-power point, and in the other case the noise spectrum intersects the spectrum of the true SSH at higher wavenumbers than the filter half-power point. These cases are referred to as the high-noise and low-noise cases respectively. To make it easy to see how different filtering properties appear in the coherence, I also included two sets of filter cutoffs (25-km and 15-km wavelengths) for each noise case.

In the low noise case, the results are unsurprisingly very similar to the noise-free case (Figure 2). The spectral gain provides a qualitatively useful but quantitatively inaccurate description of the filtering, almost correctly identifying the half-power points and grouping the mappings into two clusters according to their filtering properties. In contrast, the coherence measure of resolution identifies resolutions of 20.6km, 14.8km, 12.3km, and has three mappings clustered around 9km (including both the 15 and 25km Gaussian mappings).

The high-noise case is qualitatively different than the low-noise case. In the high-noise case, all of the coherence curves tend to collapse onto one another (Figure 3), as expected from Equation 14. Also, the gain does not provide a very useful measure of the filtering. Because we have constructed this example so that the noise in the input data has the same spectrum as the noise in the independent data ($\langle \hat{n}_0^* \hat{n}_0 \rangle = \langle \hat{n}^* \hat{n} \rangle$), the spectral ratio method does well in both of these examples with noise, as expected from Equations 16-17.

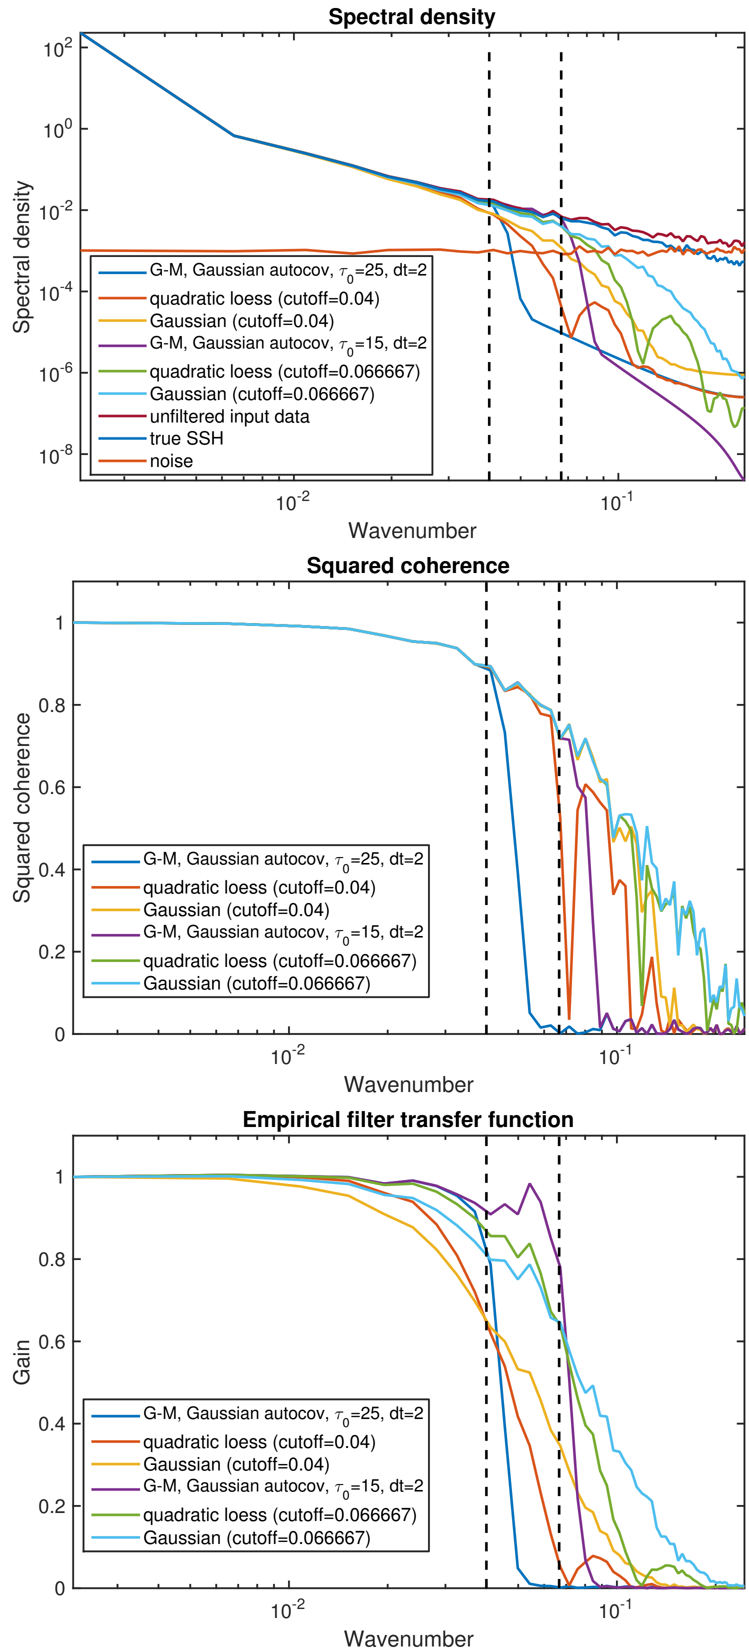

Figure 2: 1D example, with "low noise". Upper panel: Spectra of input signal and mapped fields, including a partitioning of the signal and noise in the input signal. Middle panel: squared coherence of the mapped field with the input signal. Lower panel: spectral gain computed between the input signal and the mapped field. The black dashed lines marks the theoretical half-power wavenumber of the loess and Gaussian smoothers. Note that this figure uses a log scale for frequency, unlike some other figures.

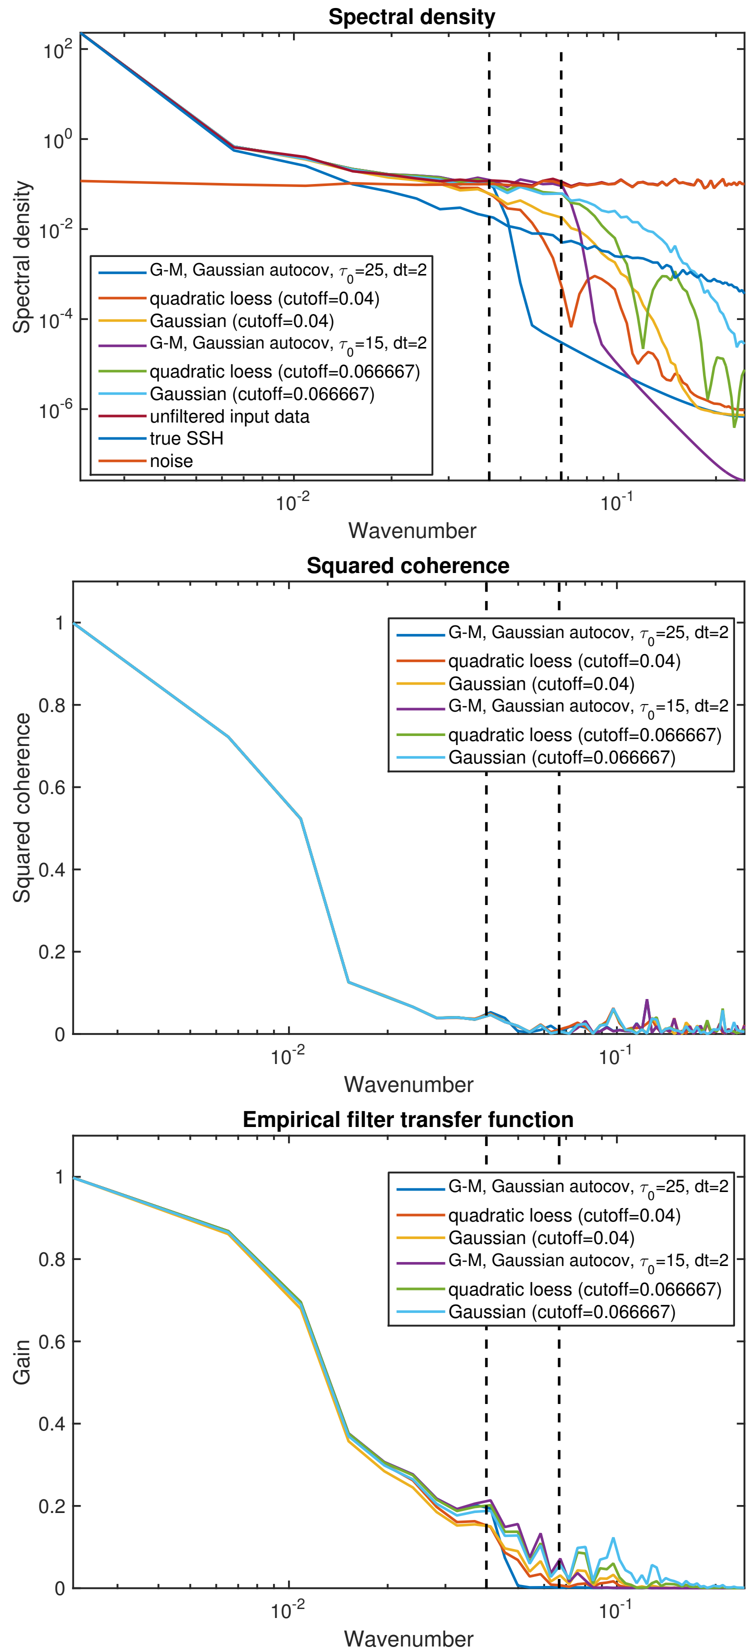

Figure 3: 1D example, with "high noise". Upper panel: Spectra of input signal and mapped fields, including a partitioning of the signal and noise in the input signal. Middle panel: squared coherence of the mapped field with the input signal. Lower panel: spectral gain computed between the input signal and the mapped field. The black dashed lines marks the theoretical half-power wavenumber of the loess and Gaussian smoothers. Note that this figure uses a log scale for frequency, unlike some other figures.

2.2.1 Summary of 1-D examples

The coherence should be completely independent of the filtering and should only contain information about the SNR (Equation 14). In the 1-D example without noise, the coherence seems to indirectly contain information about the filtering, which appears to be related to the SNR associated with noise from numerical errors. In this case, application of the coherence method yields estimates of the resolution that range from 6.5-20 km for three different smoothers with nearly equivalent half-power wavelengths of about 25 km. In the more realistic cases that contain noise, the coherence method applied to the high-noise case (when the wavenumber at which SNR=1 is at a lower wavenumber than the filter cutoff) does give important information about the SNR in the input data and the independent validation data, but it tells us nothing at all about the differences in the filtering in the the six different smoothing scenarios. The coherence method yields no reliable information about the smoothing in the low-noise case (when the wavenumber at which SNR=1 is at a higher wavenumber than the filter cutoff). In this latter case, for example, the Gaussian smoothers with 25-km and 15-km half-power wavelengths are both identified as having about 9-km resolution using the coherence method.

The derivation and examples given above illustrate my basic objection to the coherence-based measure of resolution: it tells us nothing useful about the filtering in the DUACS mapping. It potentially can provide useful information about the SNR of the input data, but the SNR of the input data will be the same regardless of what mapping parameters are applied to the data. If the method is not robust in 1-D, we cannot safely apply it in 3-D.

2.3 2-D example with a white spectrum

We could really stop this discussion here. However, I imagine the authors might wonder, as I did when thinking about the above example, why the coherence-based measure of resolution seems to provide reasonable results when applied to the DUACS system. The next two examples show another problem with the coherence-based measure of resolution that is related to the fact that the filtering and spectrum in one dimension (e.g., time) has important effects on the coherence in the other dimension (e.g., space).

In this example, the input signal is a random realization of a process having a white spectrum, sampled on a uniform 2-D grid meant to represent space in one dimension and time in the other. The input data were "mapped" (smoothed) to the same space-time grid as they were sampled on. There is no measurement noise and no sampling error. In this case, the only thing limiting the resolution of the mapped fields is the filtering inherent in the mapping algorithm. One could imagine this case as a case where there are repeated along-track passes of data, and they have been mapped to a regular space-time grid (along-track). Alternatively, instead of thinking of one dimension as time, we could imagine the two dimensions as along- and across-track directions. (Please note that these examples have noisier and more coarsely resolved spectral estimates because I could not afford the computer power/time to use a very large number of samples.)

In this second example, the smoothing in the "along-track" direction is the same as in the 1-D example. In the other dimension (call it time), there is some smoothing (10-unit half-power wavelength for the loess and Gaussian smoothers, and similar for the Gauss-Markov mapping), but the exact value doesn't matter.

In analogy to the approach used in the paper, I took a single "along-track" sample of the input data and of the 2D mapped data (from the central "time" of the mapped domain, to avoid edge

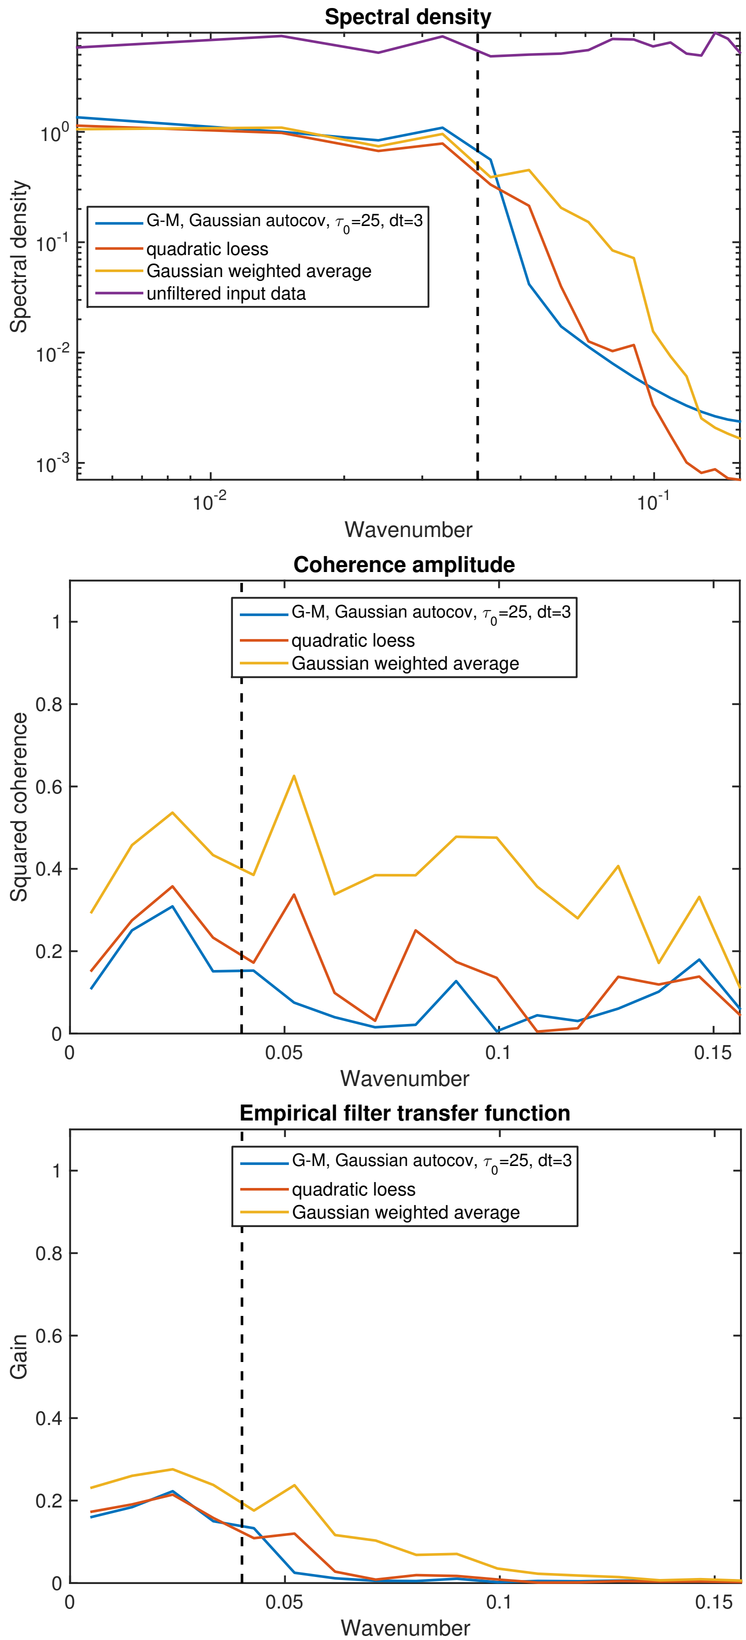

Figure 4: 2D example for a variable that has a white spectrum. Upper panel: Spectra of along-track samples of input signal and mapped fields. Middle panel: squared coherence of the mapped field with the input signal. Lower panel: spectral gain computed between the input signal and the mapped field. The black dashed line marks the theoretical half-power wavenumber of the loess and Gaussian smoothers in the along-track direction.

effects in the mappings). The spectra of the along-track input data and mapped data are shown in Figure 4 (upper panel). The variance of all three mapped fields is reduced at all wavenumbers relative to the input data, but the variance reduction is greatest at the highest wavenumbers. This is easy to understand– it just reflects the fact that the along-track filtering affects the larger along-track wavenumbers but the temporal filtering affects all along-track wavenumbers equally and accordingly causes a uniform reduction in variance. (This is analogous to how the across-track smoothing of SWOT data should reduce the contributions of noise to the along-track spectrum.)

The estimates of gain between the along-track input data and mapped data tell a similar story (Figure 4, middle panel). The mapped fields are attenuated relative to the input data at all along-track wavenumbers because of the temporal filtering, and they are even more attenuated at the high along-track wavenumbers because of the along-track filtering.

The squared coherence is not especially interesting. Squared coherence in a particular wavenumber band can generally be interpreted as the fraction of the variance at that wavenumber in the input record (the raw along-track data) that can be accounted by multiplying the Fourier transform of the output record (the along-track mapped data) by some complex-valued constant (the value of the transfer function at that wavenumber). So, we can see another aspect of what we saw in the spectra and gain plots– the coherence is reduced at all wavenumbers because the temporal filtering reduced the variance in the mapped fields at all wavenumbers relative to the raw along-track record, and thus there is only limited ability of the along-track mapped data to account for the variance in the raw along-track data. (I struggled to understand this, but I found this simple example helpful: imagine if the temporal averaging were very extreme such that the mapped data is very close to the time-mean SSH anomaly; in that case, we would expect very low along-track coherence with an along-track pass of SSH anomaly that would be dominated by eddies and variability.)

If we did try to apply the coherence-based measure of resolution in this example, we would conclude that the wavelength resolution of the Gauss-Markov and loess mappings is larger than the domain size (1800 km), and that the Gaussian mapping might resolve about 50 km wavelengths.

2.4 2-D example with a red spectrum

I imagine the authors still might wonder, as I did when thinking about the above example, why the coherence-based measure of resolution seems to provide reasonable results when applied to the DUACS system. The coherence in Figure 1 of the paper, for example, does not look like the coherence in Figure 1 or in Figure 4. Instead, Figure 1 of the paper has high coherence at low wavenumbers and low coherence at high wavenumbers. I think the explanation is that the variabilty in SSH has a red spectrum in space and time, so that the low wavenumbers tend to be associated with energetic low frequencies, and the low frequencies tend to be associated with energetic low wavenumbers. This example, with a signal that has a red spectrum in both wavenumber and frequency, is meant to illustrate how the coherence-based measure of resolution applied to a red signal spectrum combined with the multiple-dimension filtering of the DUACS system can lead to results that seem reasonable, even though the apparently reasonable looking results are basically accidental.

In this example, the input signal is a random realization of a process having a red spectrum in both wavenumber and frequency (spectrum proportional to $k^{-2}\omega^{-2}$), sampled on a uniform 2-D grid meant to represent space in one dimension and time in the other. All other aspects are identical to the second example.

I again took a single "along-track" sample of the input data and of the 2D mapped data. The spectra of the along-track input data and mapped data are shown in Figure 5 (upper panel). Unlike the second example with white spectra, the variance of the three mapped fields is not noticeably reduced at low wavenumbers. This is not difficult to understand– there is relatively little variance at the high frequencies in the temporal dimension, so the temporal low-pass filtering associated with the mapping has relatively little effect on the variance in the along-track direction.

The gain plots in this example (Figure 5, middle panel) look more similar to the 1-D example than to the 2-D example with a white noise signal. The half-power points one would infer from the gain are at similar wavenumbers to the theoretical 25km half-power point of the loess and Gaussian smoothers (29km wavelength for the loess, 27km wavelength for the Gaussian, and 22.7km wavelength for the Gauss-Markov mapping). This good agreement is totally accidental. If the frequency spectrum of the SSH were different, or if the temporal smoothing were different, the gain between the along-track input data and mapped data would change, and the point where it is equal to $\sqrt{0.5}$ (i.e., the half-power point) would be different.

The squared coherence in this example looks qualitatively similar to Figure 1 of the paper, with high coherence at low along-track wavenumbers and low coherence at high wavenumbers. The coherence-based definition of resolution would yield along-track wavenumbers of 8.5km wavelength for the Gaussian mapping, 14.7km wavelength for the loess mapping, and 19.9km wavelength for the Gauss-Markov mapping.

3 Conclusion

The along-track filtering properties of the mapping schemes should be the same in all three examples. (For example, we can analytically derive the filtering for the Gaussian weighted average.) The half-power point of the filtering would not be a good specification of the resolution, in general, because the resolution will also depend on the sampling and on the SNR of the input data. However, in these examples, which do not have any noise or sampling errors, the effective resolution should be determined only by the filtering. The filter half-power points of the along-track filtering were at a wavelength of about 25km for all three mapping schemes.

I suppose all measures of resolution have their drawbacks (e.g., the spectral ratio approach discussed in the paper is subject to some of the same issues, such as sensitivity of inferred along-track resolution to temporal filtering), but I do not see any theoretical basis for the coherence-based measure of resolution. In the cases we considered, the coherence-based measure yielded along-track wavelength resolutions ranging from 6.5km (close to the Nyquist wavelength) to >1800km (the domain size) for cases in which the actual resolution was about 25km, and I am confident that almost any result could be obtained by varying the signal spectrum and/or the temporal filtering. In a case with no noise and no sampling errors, the along-track wavenumber resolution should not depend on the signal spectrum or the temporal filtering. I cannot think of any relevant case in which it would make sense to define the resolution as the wavenumber where the squared coherence between the mapped product and a one-dimensional sample of the input data is equal to 0.5.

I think that the conclusions of the paper are not quantitatively useful and that it would be counterproductive to publish this. I feel really bad having to say that, because it is clear the authors worked hard to make a high-quality paper. Except for the one methodological flaw, it is

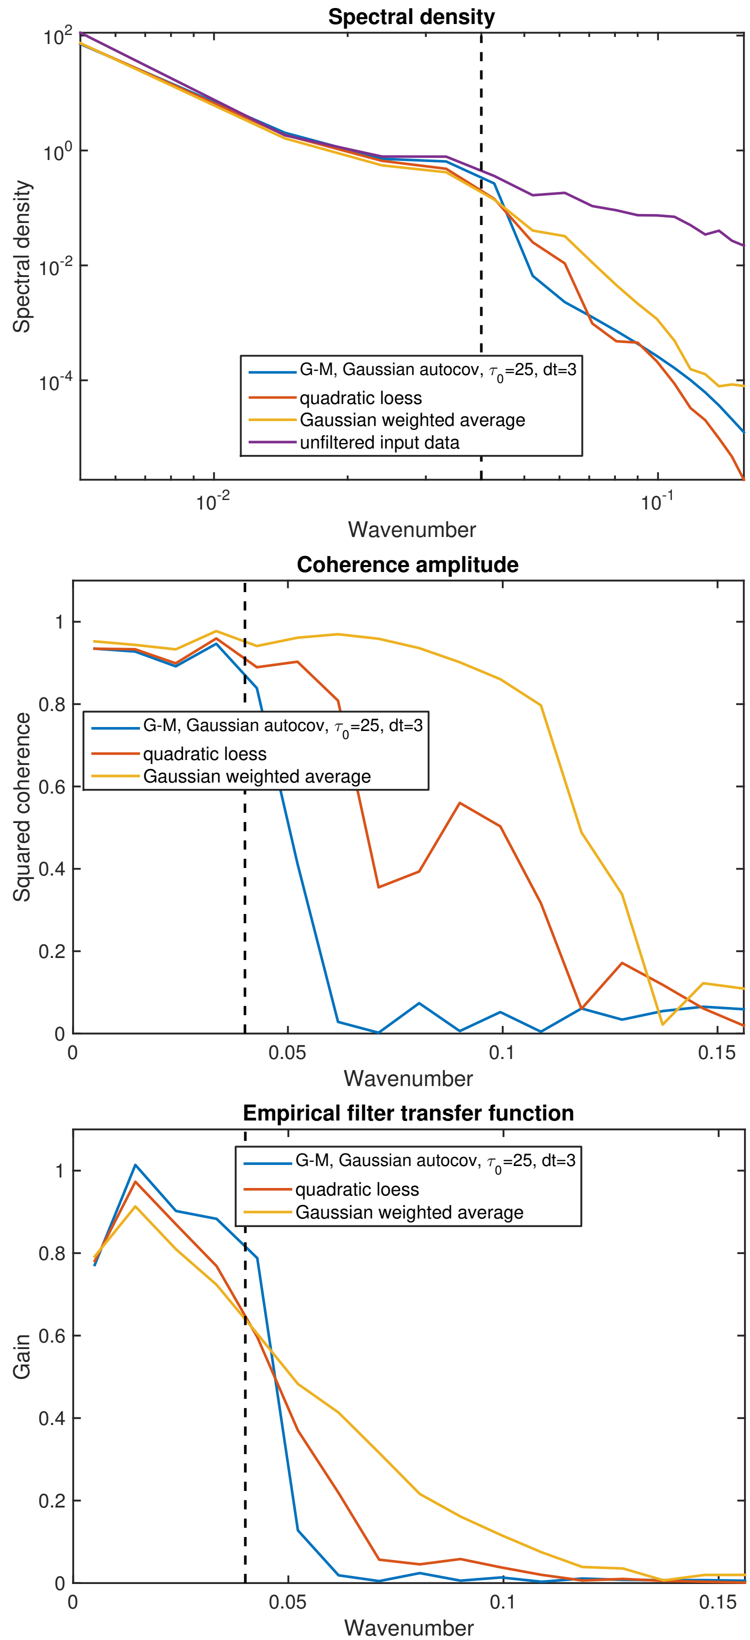

Figure 5: 2D example for a variable that has a spectrum that is red in wavenumber and frequency. Upper panel: Spectra of along-track samples of input signal and mapped fields. Middle panel: squared coherence of the mapped field with the input signal. Lower panel: spectral gain computed between the input signal and the mapped field. The black dashed line marks the theoretical half-power wavenumber of the loess and Gaussian smoothers in the along-track direction.

an excellent paper in all other respects.

I have focused on the fact that the coherence-based measure of resolution does not take into account the filtering by the mapping system. The coherence-based measure of resolution also does not adequately take account of the sampling, and this might also be a serious issue.

Acknowledgements

I thank Dudley Chelton for constructive critical comments on two earlier versions of this review.