The color of sea level:

The importance of spatial variations in spectral shape for assessing the significance of trends. *

Chris W. Hughes and Simon D. P. Williams
National Oceanography Centre, Joseph Proudman Building,
6 Brownlow St.,Liverpool L3 5DA, UK
cwh@noc.ac.uk; sdwil@noc.ac.uk

Abstract

We investigate spatial variations in the shape of the spectrum of sea level variability, based on a homogeneously-sampled 12-year gridded altimeter dataset. We present a method of plotting spectral information as color, focusing on periods between 2 and 24 weeks, which shows that significant spatial variations in the spectral shape exist, and contain useful dynamical information. Using the Bayesian Information Criterion, we determine that, typically, a 5th order autoregressive model is needed to capture the structure in the spectrum. Using this model, we show that statistical errors in fitted local trends range between less than 1 and more than 5 times what would be calculated assuming "white" noise, and the time needed to detect a 1 mm/yr trend ranges between about 5 years and many decades. For global-mean sea level, the statistical error reduces to 0.1 mm/yr over 12 years, with only 2 years needed to detect a 1 mm/yr trend. We find significant regional differences in trend from the global mean. The patterns of these regional differences are indicative of a sea level trend dominated by dynamical ocean processes, over this period.

*An edited version of this paper was published by AGU. Copyright 2010 American Geophysical Union: Hughes, Chris W. and Simon D. P. Williams, 2010: The importance of spatial variations in spectral shape for assessing the significance of trends, J. Geophys. Res. (115), doi:10.1029/2010JC006102.

1 Introduction

Sea level varies on time scales from seconds to millions of years as a result of a wide variety of physical processes [*Harrison*, 2002]. The recent interest in climate change and global sea level has led to a focus on trends in sea level over time scales of a few years to centuries. The aim of this work has mainly been to understand the causes and effects of global sea level rise, although there is now increasing interest in the regional patterns of sea level change. For example, loss of mass due to melting of ice from a particular region results in a characteristic pattern of vertical land movement and change in the gravitational field, which produces a sea level "fingerprint", and there have been several attempts to use these fingerprints to infer the size of meltwater sources responsible for recent sea level rise [*Mitrovica et al.*, 2001; *Plag*, 2006; *Marcos and Tsimplis*, 2007]. Attribution of such patterns to climate change assumes that they are not being swamped by dynamical variability in the ocean. Even more fundamentally, it assumes that the fit of a linear trend to the measured sea level curve is statistically significant.

At the same time, sea level and related variables (bottom pressure and density) are the crucial quantities being measured in efforts to monitor the large-scale ocean circulation, for example the flow through Drake Passage [*Woodworth et al.*, 2006] and the prototype monitoring systems for the North Atlantic meridional overturning circulation (MOC) [*Cunningham et al.*, 2007; *Bingham and Hughes*, 2008; *Bingham and Hughes*, 2009] . While the primary aim of such systems is to look for secular or step-like changes in the large-scale ocean circulation, they are operating within a system which has variability over a wide range of length and time scales, many of which may be unrelated to the slow, basin-scale variability which is the intended target of the monitoring system. Detection of a meaningful change in the MOC, for example, requires that the signal be distinguishable from the background natural variability.

These are just two examples, but the point is quite generic: when fitting a curve to any time series, assessment of the statistical error in the fit depends on the spectrum of the time series. For this reason, it is important to understand the shape of spectra of sea level. Such an understanding can also be valuable in understanding the physics of the processes which lead to those spectra.

So, we would like to have a model for the natural background spectrum of sea level variability—ideally a wavenumber and frequency spectrum. This is the subject discussed in detail by *Wunsch* [2009], who focuses on the measurement of ocean transport variability in the North Pacific. Here, we focus purely on the frequency spectrum of global sea level variability, and in par-

ticular we address the question of whether there is a canonical spectral shape which can be used in the assessment of whether fitted trends (or other functions of time) are statistically significant given the background variability. In doing this, we are in effect treating each time series in isolation, and therefore not taking advantage of any possible correlations between time series, or of the fact that, for example, some parts of the spectrum are likely to be due to signals which are correlated (with or without lags) over large distances whereas other parts will represent purely local variability. Our purpose is first to find out over which scales the shape of the spectrum varies, before worrying about the combination of information in time and space domains (although we will consider as a special case the globally-averaged sea level time series). From the nature of the satellite altimeter signal being considered, we expect that the variability will, in many places, be dominated by the influence of oceanic mesoscale dynamics, although other processes are also clearly present. The simplest, and most useful, result would be to find a particular function of frequency which could be used to represent the spectrum at any point.

Past work has suggested that there might be such a function. *Stammer* [1997], using sea level measurements from the TOPEX/POSEIDON altimeter, emphasized the similarity between spectra from different regions of the ocean, while pointing out some differences in western boundary currents and in the near-equatorial region, and *Zang and Wunsch* [2001; ZW hereafter] built upon this work to suggest a canonical frequency-wavenumber spectrum which, when scaled by a spatially-varying amplitude factor, was offered as a first order description of the oceanic spectrum anywhere. On the other hand, *Le Traon* [1991] showed from a regional analysis of Geosat data in the North Atlantic that the energetic Gulf Stream region contained relatively more energy at high frequencies (periods less than about 80 days) than regions of the eastern North Atlantic at similar latitudes. We now have much longer time series of altimetry, and can look into these spectra in more detail.

In seeking to investigate the spatial variability of frequency spectra, we run into the difficulty that it is impractical to display a spectrum for each point in the ocean, and area-averaging requires the introduction of preconceived ideas of areas over which the spectrum is assumed to be homogeneous. We address this by introducing a novel method of displaying spectral information in color, exploiting the way the eye is adapted to extracting useful information from the spectrum of light.

We will show that there are significant spatial variations in the shape of the spectrum, which cannot be adequately described by one or two parameters. In much of the tropics, Rossby wave theory gives a useful aid to understanding the observed spectrum, as suggested by Lin et al. (2008), but

at higher latitudes this interpretation breaks down and more complex spatial variations are seen. In fact, in order to describe the spectral variation globally, we find that a fifth order autoregressive model is required (i.e. 6 parameters are needed to describe the spectrum, 5 shape parameters and 1 amplitude). Having fitted such a model to each chosen time series of sea level anomalies, we can then estimate the time required for a sea level trend (we assume 1 mm/year as an example of a typical difference from the global mean) to rise above the statistical background variability.

2 The sea level rainbow

We use the delayed-time, reference, merged sea-level anomaly, gridded product from AVISO, which can be downloaded via http://www.aviso.oceanobs.com/en/data/product-information/duacs/ssaltoduacs-products/index.html. Time series at each grid point consist of 630 weekly values over the period 5 April 1995 to 25 April 2007, in order to avoid a gap in coverage by the ERS1 satellite before this period, allowing for homogeneous sampling over the period considered. This dataset is, at all times, derived from a combination of one altimeter in the TOPEX/POSEIDON orbit and one in the ERS 35-day repeat orbit. The data are provided on a 1/3 degree Mercator grid covering latitudes between 82°S and 82°N, but the optimal interpolation used means that signals at wavelengths shorter than about 100 km are suppressed (the effective cutoff wavelength is somewhat latitude dependent, being longer at lower latitudes). As part of the AVISO processing, the data have been corrected for tides and a barotropic model has been used to estimate and subtract the barotropic response to atmospheric pressure and wind stress forcing. This will reduce the signal mainly at the high frequencies which dominate such barotropic motions. In high latitude coastal regions, where such variations are the dominant signal, the model explains about 40–50% of the variance remaining after application of the traditional Inverted Barometer response [*Carrère and Lyard*, 2003], and similar accuracy can be inferred for this component of the variability in deeper regions, although a smaller percentage of variance is explained in deeper regions as a result of other processes becoming more important. For each time series, we have simultaneously fitted and removed annual and semiannual sinusoids and a linear trend, following which we have used linear interpolation to fill a maximum of 5 gaps of up to 5 weeks each (gaps are mainly due to seasonal sea ice, but also result from occasional instrumental problems). Time series with more or longer gaps are discarded. Periodograms have then been calculated from Fourier transforms of the time series at each grid point. We perform no explicit spectral band

averaging, as this is implicit in the conversion of spectrum to color.

In order to display the resulting spectra, we exploit the fact that the canonical spectrum proposed by ZW is, with a rescaling of frequency, close to the spectrum of white light. The ZW spectrum as a function of frequency σ is proportional to $\sigma^{-1/2}$ at periods longer than 100 days, and proportional to σ^{-2} at shorter periods, so at high frequencies it is what is usually referred to as "red" noise. It is a common misconception that the spectrum of visible white light is equivalent to a "white" noise spectrum, in which the Power Spectral Density (PSD) is constant as a function of frequency. In fact, the visible spectrum is often plotted as a Spectral Power Distribution (SPD) function, as a function of wavelength (which is inversely proportional to frequency). There is a factor of σ^{-2} difference between these two representations so, while the SPD of white light is almost flat across the visible part of the spectrum, the corresponding PSD is proportional to σ^{-2}. White light is "red" noise. A spectrum of light which was truly "white" noise would actually appear to the eye as quite an intense blue. This unfortunate misunderstanding leads to a problem of nomenclature. To be clear, colors mentioned here will all be visually-perceived colors unless enclosed in inverted commas, in which case the color referred to is that of the conventional definition for colored noise.

Despite this problem, the true relationship between spectrum and color turns out to be rather useful here. We are used to interpreting small differences from white light ("red" noise) as colors, so if the sea level spectrum was exactly "red" noise, a translation into a light spectrum would render it as a grey scale: low amplitude variability would be black, and high amplitude white, with greys at intermediate values. With this translation, then, colors indicate that the shape of the spectrum differs from the ZW shape.

This would not work if we were to map the spectrum of low frequency (say interannual) sea level variability on to the visible spectrum, as (according to ZW), this spectrum is much less steep than σ^{-2}, so the resulting colors would all be towards the blue end of the color range (though still on the "red" side of "white" noise). Annual and semiannual signals are removed in our time series analysis, which will produce a dip at the corresponding part of the spectrum, so to avoid any effect of this removal we chose to map the range of periods 2–24 weeks onto the visible wavelength range of 380–760 nm. This mapping means that low frequency variability (periods longer than semiannual) is not represented in the plotted spectrum, as the mapping results in this variability shifting to the invisible infrared part of the electromagnetic spectrum.

The picture which results from this remapping of the sea level spectrum is shown in Fig. 1a. There is a lot of information in this figure, so it is worthwhile to spend some time describing it in detail.

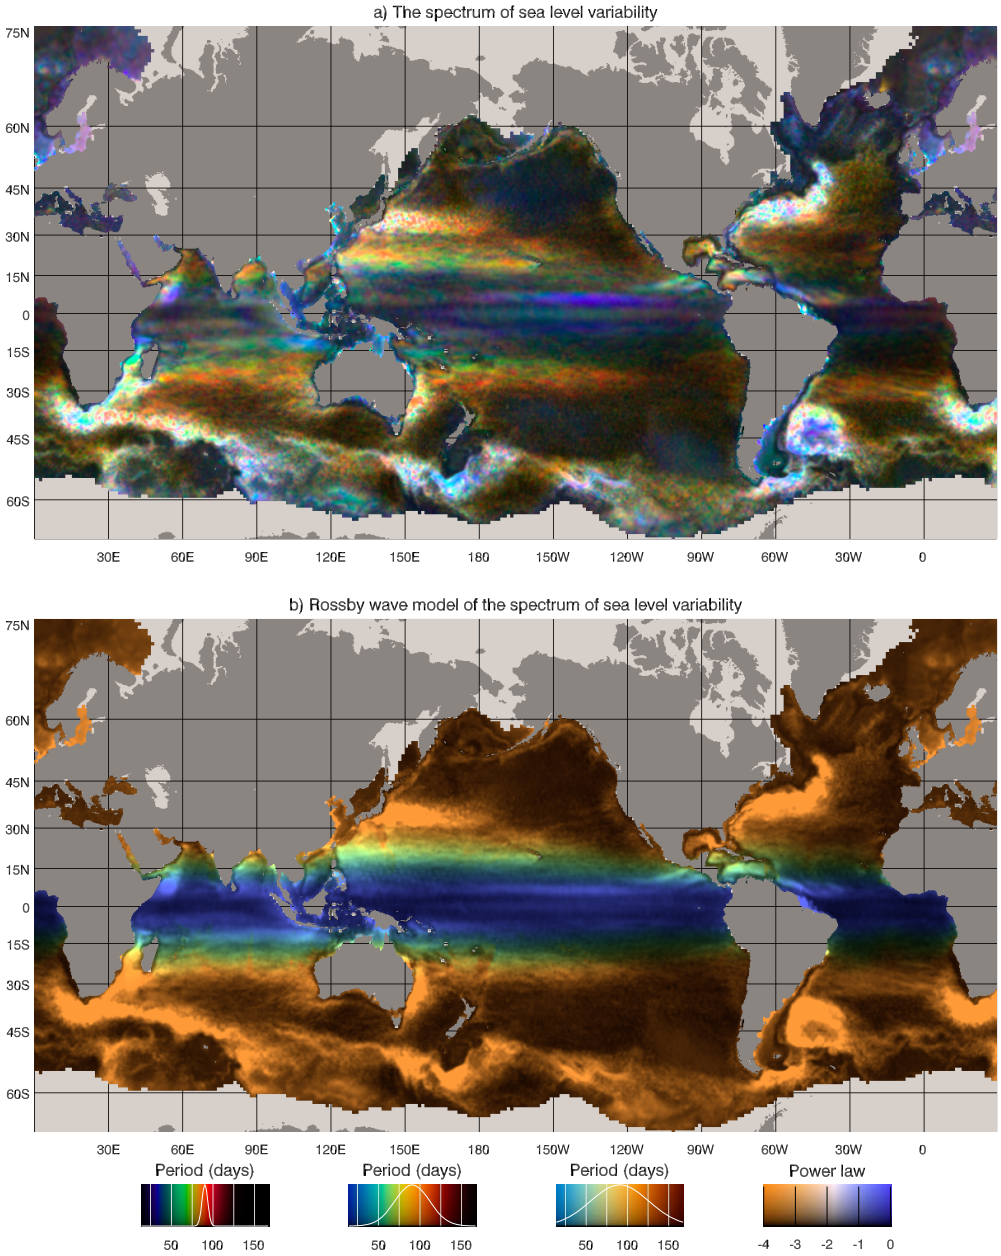

Figure 1: The spectrum of sea level variability presented as the color and brightness of the equivalent spectrum of light. The range of periods 2 weeks to 24 weeks is mapped to light wavelengths 380 to 760 nm (see Appendix A for a detailed description of this mapping). a) The spectrum based on 12 years of altimeter data. b) Brightness as above, but with color derived from a simple spectral model in in which power is proportional to frequency to the power 0 at low frequencies and to the power -4 at high frequencies. The dividing frequency is 1.4 times the maximum allowed linear first baroclinic Rossby wave frequency, derived from the Rossby radius taken from the atlas of *Chelton et al.* [1998]. The scale bars show colors (at all brightnesses) corresponding to spectra which are gaussian, as a function of period, with three different widths illustrated by the white curves, and the color of spectra in which power spectral density (as a function of frequency) is proportional to frequency raised to different powers.

First, in order to interpret the plot, note that there is both brightness and color information here. Brightness is related to the amplitude of sea level variability over the range of frequencies mapped, while color relates to the shape of the spectrum. We chose the scaling with the idea that the ZW spectrum should be close to white or grey, and examples of the actual color of the ZW spectrum are given by the two greys which represent land and missing data (mainly due to sea ice). The fact that the ocean region has no overall blue or red tint confirms that the ZW spectrum is a good first-order appoximation. Other colors do not have a simple one-to-one relationship with frequency, just as they do not in light. For example, orange can be sodium orange light at a particular frequency, or it can be a mixture of red and yellow light. There is thus no simple general interpretation of a given color, and it is impossible to provide a scale for this figure which gives quantitative information about amplitudes, or which covers the possible meanings of all possible colors. Rather, the figure should be interpreted as a qualitative guide to variations in amplitude (brightness) and shape (color) of the spectrum, with some information about dominant periods of variability to be inferred from the color.

Some more detailed guidance can be obtained by looking at the scale bars plotted at the bottom. These show the colors (at all different brightnesses) which would be produced by spectra of various different shapes. The first three scales show the colors resulting from a Gaussian SPD (i.e. Gaussian as a function of period, not the PSD as a function of frequency), with various widths as shown by the curves drawn over the scales, and with peaks at various periods. The fourth scale shows the colors produced by different power-law spectra. Thus, as anticipated, a PSD in which power is proportional to σ^{-2} ("red" noise) appears as approximately white or grey, with steeper power laws (e.g. σ^{-3} or σ^{-4}) appearing as a kind of brick-red color, and more gentle power laws (σ^{-1} or σ^{0}, which is "white" noise) appearing as quite a strong blue. This shows that simple power-law noise cannot be responsible for many of the quite intense colors which appear in the figure, which must therefore mean that these colors indicate the presence of a peak in the SPD (though not necessarily in the PSD). More complicated spectra can produce colors not in any of these scales. For example, a doubly-peaked SPD, with peaks near both short and long period ends of the range, would produce a mixture of red and blue colors which may be responsible for some of the pink and purple regions seen in the figure.

The feature which immediately stands out in Figure 1a is the double rainbow pattern, in which the color starts as blue at latitudes of a few degrees north and south, and shifts through the colors of the rainbow to red at latitudes of about 30° north and south. This is visible in all ocean basins,

although less clearly in the Atlantic where variability is weaker, and is even quite clear in miniature in the South China Sea. This is the feature which was noted by *Lin et al.* [2008], who explained it in terms of first mode baroclinic Rossby waves. Such waves have a maximum possible frequency $\sigma = \beta R_o/2$ where β is the northward gradient of the Coriolis parameter f, and $R_o = c_1/f$ is the first baroclinic Rossby radius (c_1 is the speed of first mode baroclinic gravity waves). *Lin et al.* [2008] find a close agreement between what they call the peak in the sea level spectrum and this critical frequency as calculated from observation-based estimates of the Rossby radius. However, because they integrate the spectrum over bins which are equally-spaced in period, the peaks they refer to are peaks in the SPD (plotted as a function of period), not the PSD (a function of frequency). The PSD may simply display a change of slope at the reference frequency, rather than a peak.

Lin et al. [2008] explain this rainbow as the accumulation of Rossby wave energy at the Rossby wave turning latitudes. However, while the linear theory is clearly an important factor and Rossby wave dynamics do help explain the observations, it is also clear that much of what is seen here is not linear waves. For example, as *Stammer* [1997] pointed out, the bright blue stripes at about $\pm 5^\circ$ in the Pacific, corresponding to a period of about 30 days, are the signature of Tropical Instability Waves [*Legekis*, 1977; *Qiao and Weisberg*, 1995]. Further south, the various brightly-colored stripes north of New Zealand and east of Australia are the signature of eddies produced by instability (mainly barotropic instability) of zonal jets in the region, producing periods in the region of 70 days [*Qiu and Chen*, 2004; *Qiu et al.*, 2009].

To investigate this further, we have plotted the colors which would result from this theory in Figure 1b. We calculated the shortest-permitted period for first-mode baroclinic Rossby waves as $T = 4\pi/(\beta R_o)$, where we take R_o from the atlas of *Chelton et al.* [1998]. We then assumed a PSD more extreme than the ZW spectrum, with power proportional to σ^0 at periods longer than T, and proportional to σ^{-4} at shorter periods (plotted as an SPD, this produces a symmetrically-peaked spectrum). The amplitude of the spectrum was chosen so that the total power at periods shorter than 24 weeks was the same as that in the observed sea level signal. As a result, Fig. 1a shows colors given by the Rossby wave theory, but brightness from the observations. In fact, we found that the predicted colors tended to be too red using the pure theory, so the figure actually uses a reference period of $T/1.4$ rather than T. This is, to some degree, subjective, and different factors appear to work better in different regions, but the best value to choose is clearly somewhat smaller than T, but larger than $T/1.5$.

Equatorward of 30° north and south, there is broadly a good match between the figures, showing that Rossby wave theory is a useful first order

description of what is seen. However, the reduction of the chosen shortest-permitted period by a factor of 1.4, which is required to match the observations, means that the observed peak in the SPD is at periods shorter than any allowed by the linear theory. This is consistent with the observation that many of the regions of large variability are associated with instabilities. Both barotropic and baroclinic instability processes, on a beta plane, are intimately associated with Rossby wave dynamics, so it makes sense that Rossby waves should play a role, but that nonlinear processes might permit the appearance of energy at periods prohibited by the linear dispersion relation. The Rossby wave model also leads to a rather smooth gradation of color in comparison with the observations. This is also consistent with the idea that localized instability processes are important over certain regions. However, some level of smoothness is inevitable, since the Rossby radius information is only provided at a resolution of 1 degree. Overall, we have the impression (but can be no more definite than that) that the required adjustment factor is larger in the most energetic regions, and smaller in the quieter regions, consistent with the quiet regions being closer to linear wave dynamics. Our choise of 1.4 for the adjustment factor is weighted towards being correct for the most energetic regions, as these stand out most clearly (as they are brighter) in the plot.

Equally clearly, the Rossby wave interpretation breaks down at higher latitudes. This is understandable, as linear Rossby wave speeds drop below 1 cm s^{-1} poleward of about 40°[*Chelton and Schlax*, 1996]; , making advective processes much more dominant at these latitudes and resulting in much stronger interactions with bottom topography, especially in the Southern Ocean (e.g. *Hughes et al.*, [1998]). A similar argument was made by *Le Traon* [1991], who suggested that the increased high frequency energy in the Gulf Stream region, which shows as the blue and white colors in our plot, is the result of energetic mesoscale turbulence, whereas the low-frequency dominated eastern North Atlantic (brown-orange-red in our plot) is closer to a linear wave regime. More recently, *Tulloch et al.* [2009] also came to the conclusion that linear Rossby wave theory ceases to be a useful description of the observations at latitudes polewards of about 30°, based on a fit of observations to Rossby wave dispersion relations and a scaling comparing propagation speed with eddy velocities.

A particularly striking example of topographic influence occurs south of the Antarctic Circumpolar Current at about 150°W, where two large, dark blocks of reddish and bluish color occur on the northern and southern flanks respectively of the Pacific Antarctic Rise, the divide between red and blue occuring exactly at the ridge crest. This can be seen more clearly in Figure 2, which shows a section of Figure 1a replotted at a brighter level, with

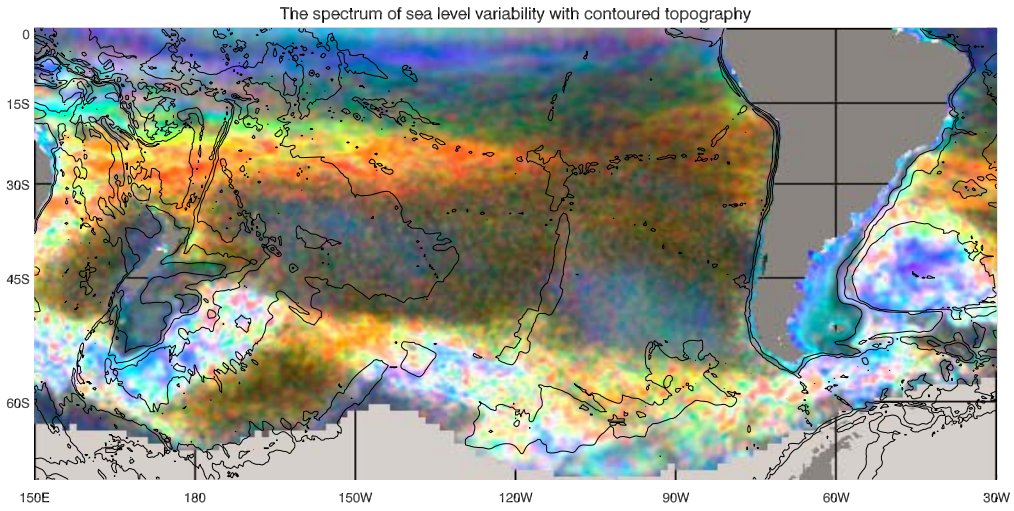

Figure 2: Enlarged section of Figure 1a, shown six times brighter and with topographic contours superimposed. Contours are at depths of 1, 3 and 5 km. The 3 km contour near to 150°W, 60°S represents the crest of the Pacific Antarctic Rise.

topographic contours overlaid for reference. Another high-latitude feature which stands out is the pair of milky bluish clouds which occur in the eastern North and South Pacific. These are large-scale barotropic modes driven by local wind stress [*Fu*, 2001], the bluish color being representative of the greater high-frequency content in wind stress than is found in internally-driven ocean processes at these latitudes; a third such region southwest of Australia is masked by the high eddy variability found there. In principle, these modes should be removed from the altimetry by the barotropic model correction applied by AVISO, but in practice this will only reduce the variability, not completely remove it.

A similar signal is seen on the continental shelf regions, especially around the UK and east of Argentina, but also on the New Zealand and Australian shelves, and the Canadian Atlantic shelf region. Variability in these regions has a characteristic blue color, reflecting the dominance of atmospheric forcing. A similar spectral plot of atmospheric pressure fluctuations, taken as a proxy for the vector wind stress (not shown) is uniformly blue with a latitude-dependent brightness. It is notable that the blue of the shelf seas and the (usually) redder color of the deep ocean are in many places separated by a dark line of low variability. This dark line is particularly clear around most of the Atlantic, but also around Australia, the South China Sea, and high latitude Pacific regions. It invariably occurs near the steepest part of the

continental slope, and represents a decoupling of the deep ocean and shelf sea variability at these frequencies.

Concern has been expressed that this "decoupling" may be an artifact of the altimeter data gridding process, especially as it often occurs close to land, and the altimeters are known to be less reliable (or even produce no data) close to land. In order to test this in more detail, we have taken 10 years of TOPEX/POSEIDON along-track data (364 cycles of 9.916 days each, starting in October 1992) from the AVISO along-track corrected sea level anomaly dataset, downloaded from the same website as the gridded data. Ten short sections of track have been identified which cross the continental slope in the Grand Banks and Argentine Shelf regions (Figure 3). Each section is approximately 480 km long (75 repeat-sampled points) and is centered on the minimum variability in corrected sea level at periods shorter than 24 weeks (17 cycles). The right-hand panel of Figure 3 shows how the standard deviation of this high frequency sea level varies along each track section, together with the ocean depth (dashed lines). The minimum variability (at the centre of each section) always occurs close to the continental shelf edge.

To see how the shape of the spectrum varies to either side of this minimum, we divided each track into five sections of equal length: sections one and two on the shelf, section three centred on the variability minimum, and sections four and five in deeper water. Spectra were then calculated as averages over all five tracks in each region, for each section (one to five) separately. The resulting spectra are shown in the left hand panels of Figure 3, with dotted lines for the two shelf section, a solid line for the central (minimum variability) spectrum, and dashed lines for the two deep ocean sections. In both regions it is clear that, in comparison with the central section, variability increases preferentially at high frequencies on the shelf, and at low frequencies in deeper water, so the along-track analysis confirms the shift in spectral slope from steep (corresponding to red or white colors in Figure 1a) in deep water, through a minimum of variability near the shelf edge, to shallow (blue in Figure 1a) on the shelf.

There are many other features which could be investigated in more detail, but we will point out just two more. Off the Pacific coast of Central America there are green-blue streaks. As noted by *Chelton et al.* [2004], these are regions of strong episodic wind stress and wind stress curl resulting from winds blowing from the Gulf of Mexico and Caribbean Sea, being steered through the mountains, and presumably driving unusual ocean variability, although still close to the appropriate Rossby wave cut-off frequency. Finally, along the Indian Ocean coast of southern Africa, and extending out to the south of Africa, there is a green-blue region, suggesting a dominant period of about 50–70 days. This pathway is consistent with the suggestion of

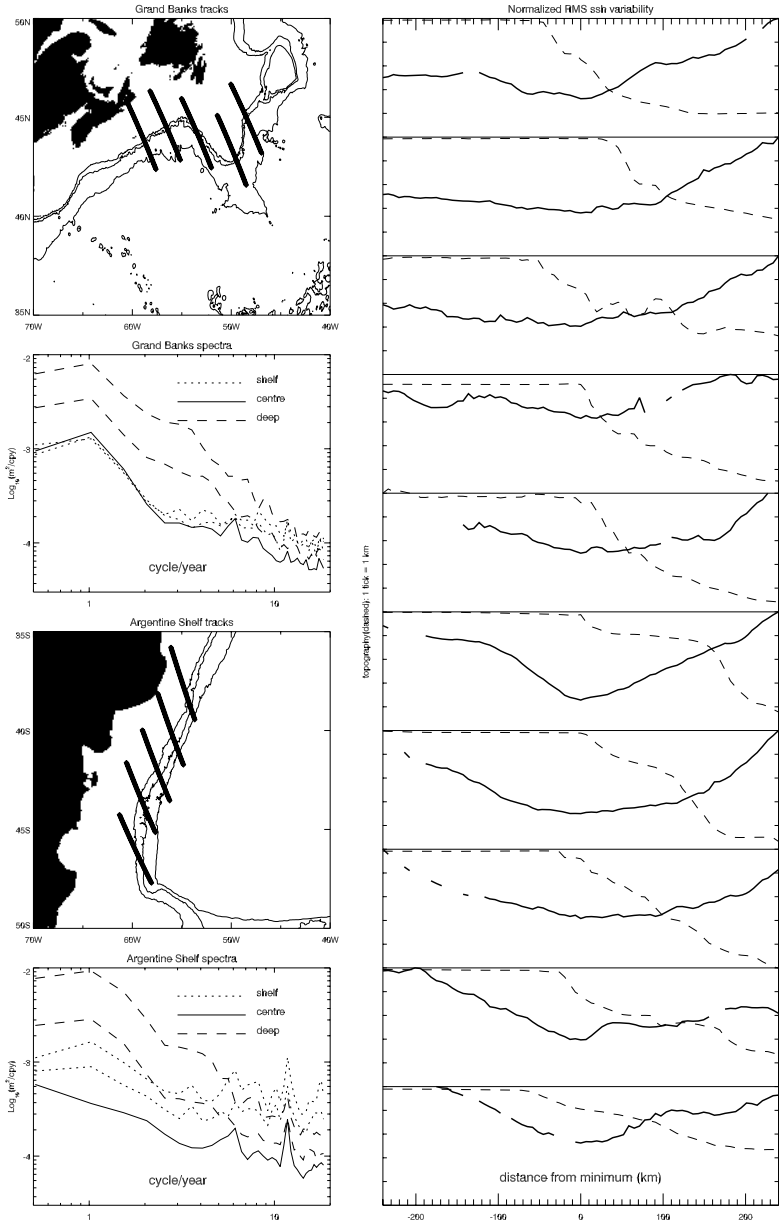

Figure 3: Along-track analysis of TOPEX/POSEIDON altimeter data in the Grand Banks and Argentine Shelf regions. Panels on the left show the five sections of track from each region, and sea level frequency spectra as an average across all five sections, but divided into five equal-length regions along the length of each track. The solid line is averaged along the middle sections of the five tracks (where sea level variability is minimum), the two dashed lines are spectra from the deep-ocean side of this minimum, and the two dotted lines from the shelf sea side. The right-hand panel shows topography (dashed) along each track section (tracks shown, top to bottom, in order from northeast to southwest), and root-mean-square sea level variability (solid) at periods shorter than 24 weeks. The latter has been normalized so that the largest value is the same for each section.

Schouten et al. [2002] that the shedding of eddies at the Agulhas retroflection is catalysed by the propagation of eddies down the African coast, producing 4–5 eddies per year (which would be a slightly longer period of 73–91 days).

It can be seen that this plot synthesizes a lot of information about ocean dynamics, and initial tests (not shown) suggest that it would make a good model diagnostic, making it possible to quickly recognize processes and features which are being modeled well or badly. For our present purposes, however, the main point is that the color, and hence the shape of the spectrum, varies significantly and in some cases over rather short distances. Latitude (or Rossby wave cut-off frequency) is a good predictor of color in many of the low-latitude deep-ocean regions, but outside these regions, and where particularly strong variability is encountered, there is no obvious way to choose regions over which the frequency spectrum will be homogeneous. If we are to build a model of the sea level spectrum, it must have quite a large degree of flexibility to be able to cope with these variations.

3 Modeling the spectra

The one-parameter model (amplitude) of ZW certainly cannot fit the variation of color of the spectrum, because such a model can only vary in brightness. What we need is a multiparameter model, but it is not clear a priori how many parameters are optimal. A way to address this problem is to use the Bayesian Information Criterion (BIC) of *Schwarz* [1978], which (under certain, rather weak assumptions required to find an analytical solution) gives a formula which balances the quality of fit against the number of fitted parameters in a manner which gives a Bayesian-optimal solution. The BIC differs from the related Akaike Information Criterion (AIC) [*Akaike*, 1974] in that the BIC penalizes additional parameters more strongly as the amount of available data increases, although the two criteria approach the problem of model selection from quite different viewpoints. A useful discussion of these and related criteria can be found in *Liddle* [2007], who finds the BIC to be more reliable. We actually tried both AIC and BIC and found that the BIC appeared to produce more stable estimates, with less spatial variability than the AIC.

To vary the number of model parameters we use the family of autoregressive (AR) spectral models which are capable of describing ever more complex spectra as the number of parameters increases. For example, the order 1 function AR(1) is a two-parameter model with one shape parameter and an amplitude, AR(2) has two shape parameters and an amplitude, etc. One property of these models is that all tend towards a "white" spectrum

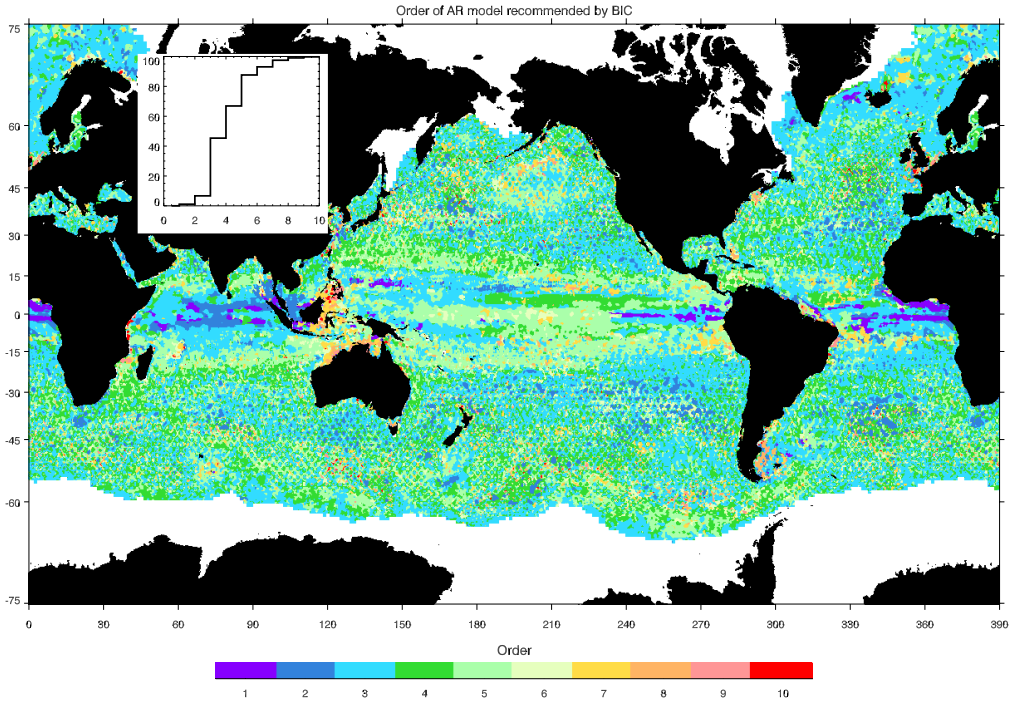

Figure 4: The order of the autoregressive model which is recommended by the Bayesian Information Criterion at each point. The inset is a cumulative distribution function (as percentage of grid points) for the recommended order.

at the low frequency extreme, which is a sensible null hypothesis for parts of the spectrum which are unconstrained by the data.

For each time series, we apply the BIC to estimate the optimal order of AR model to fit. AR models are fitted to the time series using the least-squares covariance method, which minimises the forward prediction error [*Kay*, 1988; *Stoica and Moses*, 1997]. Figure 4 shows the order predicted by this procedure as a map and as a cumulative distribution function. It is clear from this that AR(1) and AR(2) are rarely sufficient (less than 7% of cases). About 81% of points recommend AR(3), AR(4) or AR(5), so AR(5) or less is sufficient for 88% of points. The spatial distribution of the optimal order shows interesting structure, particularly in the tropics. We might speculate that some of the tropical regions requiring higher order fits represent regions in which multiple baroclinic modes play a significant role, but that remains a question for future investigation.

Based on this distribution, we choose to fit the time series everywhere with an AR(5) model. We do this rather than fit the optimal number of parameters at each point, in order to avoid sharp jumps in color at boundaries between

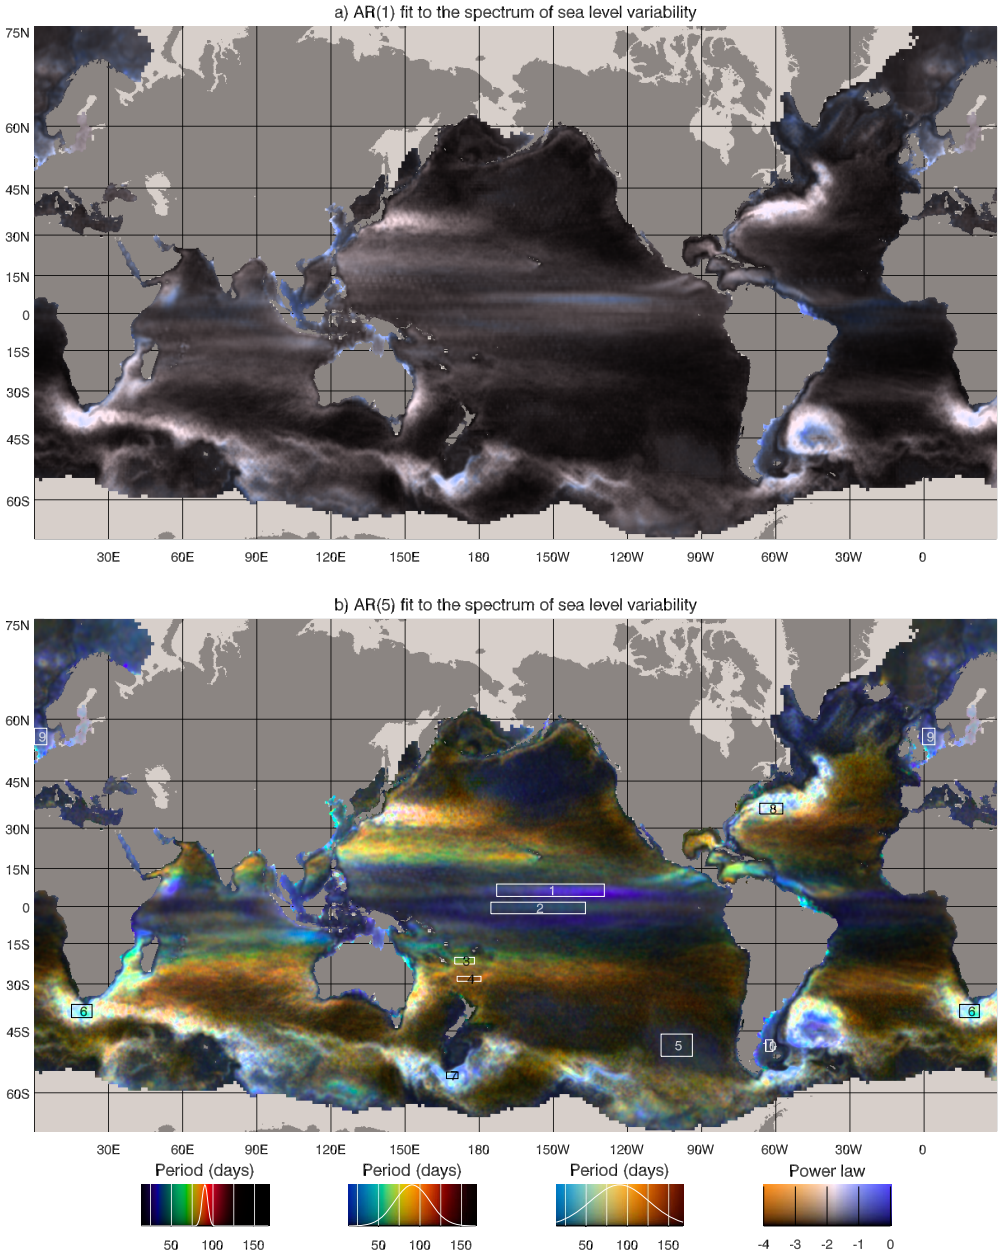

Figure 5: As for Figure 1a, but showing the colors resulting from a) AR(1) and b) AR(5) fits to the time series.

parameter numbers. This means that some spectra will be over-fitted, but those will tend to be the smoother spectra for which over-fitting means that the AR parameters themselves will be poorly constrained, while doing little harm to the shape of the spectral model. We illustrate the colors which result from both AR(1) and AR(5) fits in Figure 5.

It is clear that most of the color variation has been well captured by the AR(5) fit. Some of the reds near to $\pm 30^\circ$ are less intense than the raw spectral plot, but most features are retained. There are suggestions of a decrease in statistical noise in the colors, for example in the greater organization of colors in the region to the west and northwest of Hawaii, and in the more uniform blue-white color of the most energetic regions, although it is also possible that some of these "noisy" features from the raw colors are real, and are underfitted by the AR(5) model.

Figure 5a shows the spectra which result from an AR(1) fit to the time series. This produces practically no color over most of the ocean, balancing out the steeper and shallower parts of the spectrum to produce something close to σ^{-2} behaviour, and showing that this 2-parameter model actually produces colors similar to the 1-parameter ZW mode. An exception occurs in the regions where very high frequencies are most important, where some blue regions appear. Other AR models of order 2 to 4 produce more colorful plots (not shown), but it is not until AR(5) is used that the full rainbow effect is convincingly reproduced.

In order to illustrate the amount of structure in the actual and fitted spectra, we have selected 10 regions (labelled boxes in Figure 5b), over which the color suggests that it is meaningful to take a spatial average of the spectrum. Figure 6 shows the spectra averaged over these regions, and the corresponding average AR(5) fitted spectrum. For reference, the boundary between shaded and unshaded regions in the figure represents the ZW spectrum at arbitrary (constant) amplitude. The quality of fit is generally good, especially in the highest energy regions 6, 7 and 8. The AR(5) fit correctly captures the shoulder in the spectra at periods shorter than 1 year, which is responsible for much of the color seen in Figure 1, although in some cases the raw spectra show a sharp local peak at that shoulder which the AR(5) model cannot capture. The model also captures fairly well the second, long period plateau and rise in boxes 1 and 2, which is a measure of the strong interannual variability (including El Niño) in the tropical Pacific. The case where the spectral structure is least well-fitted is box 10 on the Argentine Shelf, which is a region where the BIC suggests almost half the points require AR orders of 7–9.

The shapes of the spectra are interesting. In most cases there is a clear shoulder, with the spectrum steeper than σ^{-2} at higher frequencies (in boxes

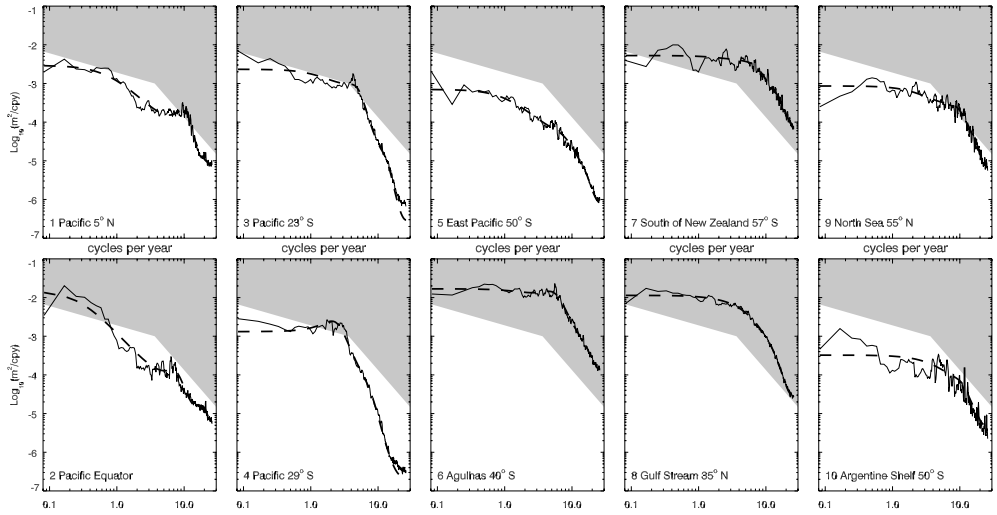

Figure 6: Power spectral densities as a function of frequency, for spectra averaged over the 10 boxes shown in Figure 5b. Dotted lines show the corresponding AR(5) fit to the data. The shading covers the region above the line represented by the *Zang and Wunsch* [2001] spectrum, at an arbitrary chosen amplitude.

3 and 4 it comes close to σ^{-5} while the Agulhas box gives σ^{-3}), and flatter than $\sigma^{-1/2}$ at lower frequencies. This behaviour was behind our choice of power laws used in the Rossby wave model illustrated in Figure 1b. In the Gulf Stream and the East Pacific (boxes 8 and 5) there is more of a smooth curve than a sharp shoulder. The steep slope at high frequencies may be partly a result of the removal of some high-frequency barotropic variability, and partly a result of filtering by the mapping procedure. This is particularly likely in boxes 5, 9 and 10, each of which are blue regions in Figure 1a. The shoulder in these regions occurs at frequencies higher than 10 cycles per year, as it does in box 1. Elsewhere, however, the steep spectral slope extends over too wide a frequency range to be entirely due to these effects.

Having settled on an AR(5) fit to the data, we now have model spectra at each point, which can be used to calculate meaningful statistical errors on sea level trends.

4 Implications of spectral shape for trend estimates

Least-squares fitting of a linear trend to a time series with no true trend, but with a spectrum of variability, will not usually produce zero trend. Given a model of the spectrum, it becomes possible to determine the size of the

false trend estimate which can be expected to result from this estimation procedure. This is an estimate of the statistical error in the fitted trend.

It is worth noting at this point a few limitations on what this estimate will tell us. First, it is purely statistical error, with no account taken of systematic error due to instrument drift, which can only be constrained by comparison with tide gauges to an estimated level of 0.4 mm/yr globally [*Leuliette et al.*, 2004]. Second, even if a trend is shown to be significantly above the statistical error, that is no proof that it is a true secular trend. It simply shows that it is meaningful to interpret the trend as a significant component of the time series over the period of observations. There is nothing to say the trend will not reverse given longer time series. Third, conclusions concerning significance of trends determined from longer time series are contingent on the accuracy of the spectral model at longer periods than those constrained by the observations. The AR sequence of models extrapolates the spectrum at long periods using a "white" noise model. Often, this will tend to underestimate statistical errors, as most geophysical time series tend to be somewhat redder than "white" noise. However, if the longest measured periods correspond to a peak in the spectrum (as might be expected where El Niño is the dominant signal, since it typically shows peak variability at around 3–7 years), the spectral model will overestimate errors in longer time series.

The statistical error in trends will depend on the amplitude of variability, and also on the shape of the spectrum. If a single spectral shape were an adequate description of the ocean spectrum everywhere, then only amplitude would matter, and statistical errors would be a constant multiple of the errors calculated assuming the spectrum was "white" (i.e. that each point in the time series is independent). In Figure 7a, we take the statistical error for a trend fitted to the 12-year time series, as estimated from the AR(5) model at each point, and divide it by the error estimated using the equivalent "white" noise model. This shows the importance of spatial variations in spectrum shape for the accurate estimation of trend errors. The ratio varies from less than 1 in a few regions, to more than 5 in the tropical Pacific. Thus, using even the best compromise for the spectral shape will result in error estimates which are wrong by more than a factor of 2. Accounting for spectral shape variation is important.

It is interesting that the tracks of the TOPEX/POSEIDON and Jason altimeters are clearly visible in Figure 7a as regions which have a lower error ratio than the surrounding ocean, particularly clearly in the South Pacific where they appear as green lines surrounding yellow diamonds. In our experience, this is unusual for diagnostics based on this mapped product, which rarely show any evidence of "trackiness". It demonstrates that, even after the gridding procedure, there is more information, and therefore a tighter

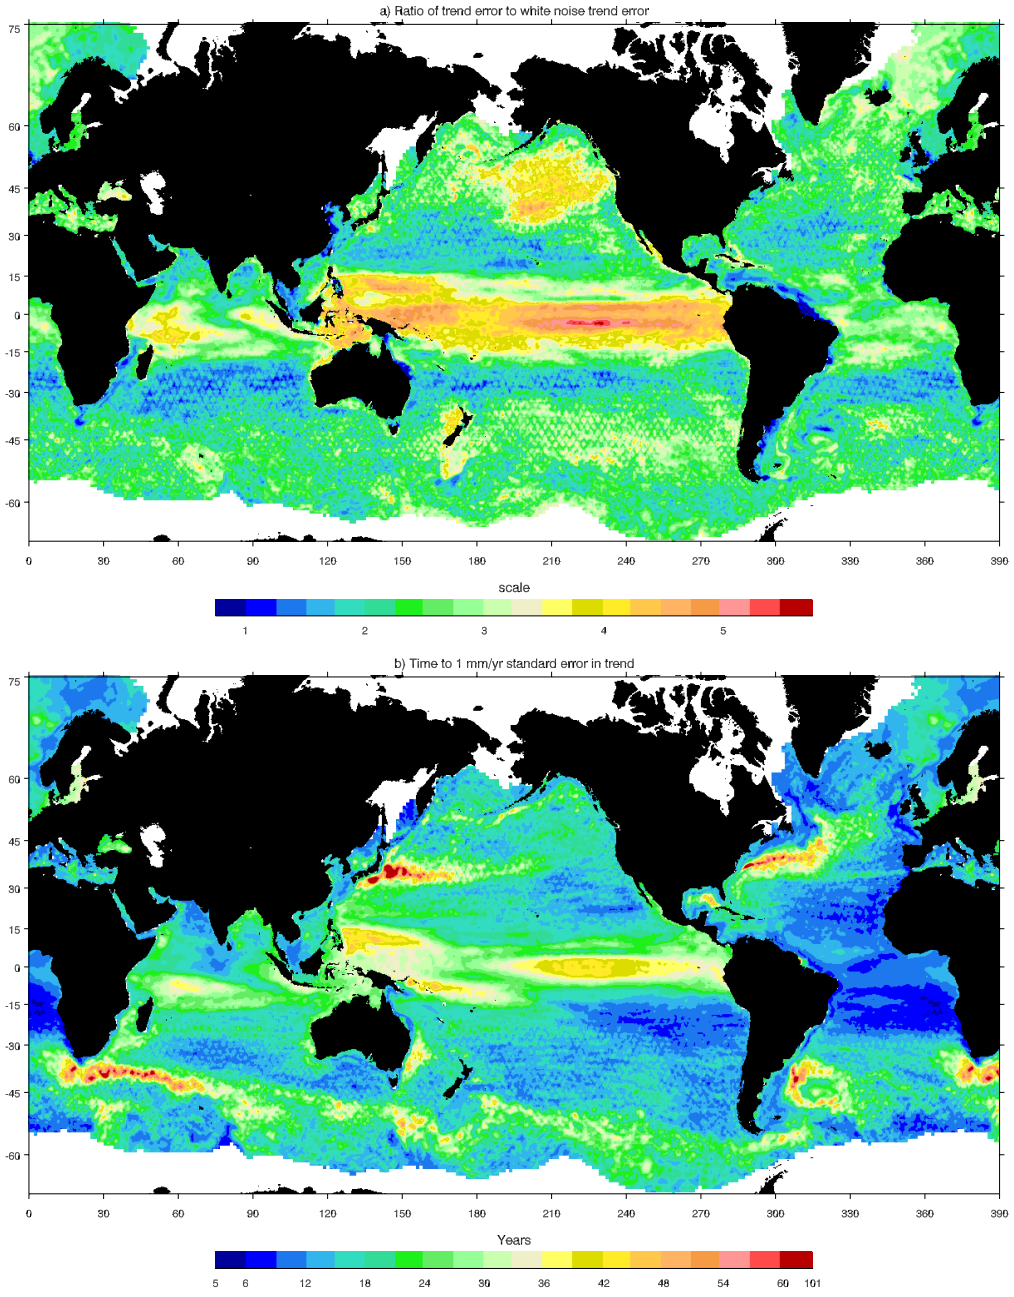

Figure 7: a) The ratio of statistical error in sea level trend over 12 years, estimated from the AR(5) fit to the data, to the equivalent error assuming a white noise model. b) The length of sea level time series which would be necessary for the statistical error in trend to reduce to 1 mm/yr, based on the AR(5) fit to the time series at each point.

constraint on trend estimates, along the more frequently sampled TOPEX tracks (the satellite return period is just under 10 days along these tracks, compared with 35 days along the ERS/Envisat tracks).

In Figure 7b we show the length of time series required for the statistical error on trend to drop below 1 mm/yr, a typical accuracy requirement for distinguishing causes of regional patterns in sea level trends. This drops below 12 years only in the quietest regions of the ocean, including much of the Atlantic continental slope. Elsewhere, the estimate relies on extrapolation of the spectrum and is therefore more speculative. In many regions the time required is several decades (reaching a maximum of 101 years at a point in the Kuroshio extension). Around the southern North Sea and English Channel, estimates vary between about 15 and 25 years, consistent with the 20–25 years implied by Figure 3 of *Shennan and Woodworth* [1992], who performed an empirical analysis of the scatter in observed trends from tide gauges in this region. As we note above, the time may well be an overestimate in the El Niño region because the long period end of the spectrum may well be a peak (the counterintuitive result that this region appears blue in Figure 1a is because interannual periods are not reflected in this plot: they are not mapped to visible colors but to the infrared). The results are also broadly consistent with *Douglas* [1991], who recommended the use of tide gauge time series at least 50 years long for the estimation of meaningful global trends, and estimated a standard error of 0.1 mm/yr for the mean rate calculated from an average of 21 tide gauge records with minimum length 60 years. We clearly cannot assess this in detail from our short time series, but the number appears to be of about the right size.

With a model for the spectrum at each grid point, we are now in a position to calculate whether the observed trend at each point is statistically significant. In Figure 8, we plot the observed trend (based on a joint least-squares fit of annual and semiannual cycles plus linear trend) over the 12-year analysis period, and the trend minus its (quasi-) global average (area weighted), which is 3.0 mm/yr. Superimposed on these plots are contours representing the lines where the observed trend is equal to plus or minus two standard deviations of the estimated statistical error in trend. In other words, values redder than the black contour are in the top 2.3% of the trends which would be expected by chance from the observed spectrum with no true trend, and values bluer than the red contour are in the bottom 2.3%. We will use this as our definition of a significant trend.

We see from Figure 8a, that significant positive trends are common, but significant negative trends are rare, the main exceptions being in the far northern Pacific, and south of South Africa. Around much of the North Atlantic coast the trend over this period is not significant. Put another way,

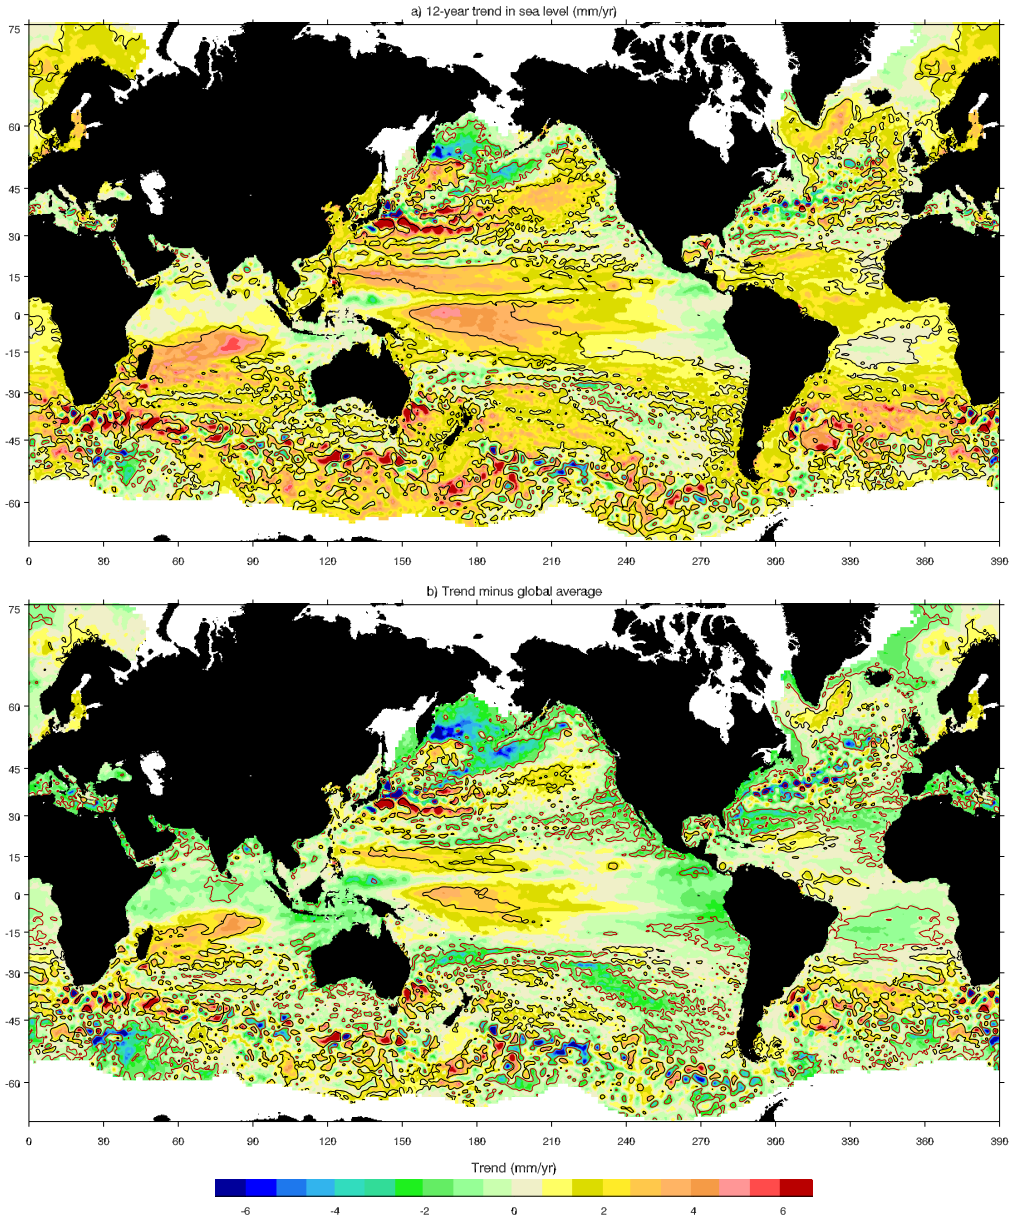

Figure 8: a) The trend in sea level over the chosen 12-year period starting in April 1995, calculated by joint least-squares fitting of a linear trend plus annual and semiannual cycles. b) The sea level trend minus its global average. Each panel has contours superimposed where the trend is twice (black) or minus twice (red) the standard error in trend estimated from the AR(5) fit to the time series at each point.

Figure 8b shows that the difference of the trend from the global average is significantly negative around the north Atlantic coast, as well as a number of other regions.

It is interesting to compare Figure 8a with the tropical Pacific and Indian Ocean trends derived from altimetry and tide gauges by *Church et al.* [2006]. The latter were calculated for a different period: 1993–2001, and give a completely different result, the main feature of which is a large sea level rise (faster than 15 mm/yr) in the vicinity of Indonesia, and a fall along the equator east of the date line. The fact that the pattern is completely different in Figure 8a is consistent with the fact that none of the negative trends in this region are identified as being statistically significant, and that the regional difference pattern in Figure 8b is statistically significant over a rather small part of the tropics. These regions are subject to large interannual variability associated with El Niño and the Indian Ocean dipole, and the trends over less than a decade are unlikely to be representative of the long term.

The patterns in Figure 8b, which show the regional variations in trends, do not look like those which would be expected from sea level fingerprint analysis to be associated with melting ice. For example, melting of Greenland ice would produce a pattern of sea level rise which is negative over a region near to Greenland, stretching as far south as Nova Scotia and Ireland [*Mitrovica et al.*, 2001], and positive further south. The observed trend, while negative close to Greenland, remains negative close to the coast out about 20°N, and becomes positive far too quickly offshore. It seems clear that this signal is better interpreted as an ocean dynamical process, for which the difference between shallow and deep water is a much more natural boundary. Similarly, the large, negative trend in the far northern Pacific is not a good spatial match to the fingerprint of Alaskan glacier melt, and is far too large to be attributed to this source. A case could be made for decreasing sea level around Antarctica, with an increase elsewhere, but this is far from convincing.

Equally the North Atlantic pattern is not very suggestive of the pattern associated with changes to the Meridional Overturning Circulation. The pattern identified by *Häkkinen* [2002] in deep water on the western side of the basin is not prominent, and while the (relative) decrease in sea level on the continental shelf is consistent with what *Bingham and Hughes* [2009] would predict for a weak slowing of the overturning, it seems as large on the eastern shelf as on the western shelf, counter to their expectations for such a mode. Although the patterns appear to be dynamical, they remain consistent with the conclusion of *Wunsch and Heimbach* [2006] that there are small fluctuations rather than any indication of a wholesale change in the overturning.

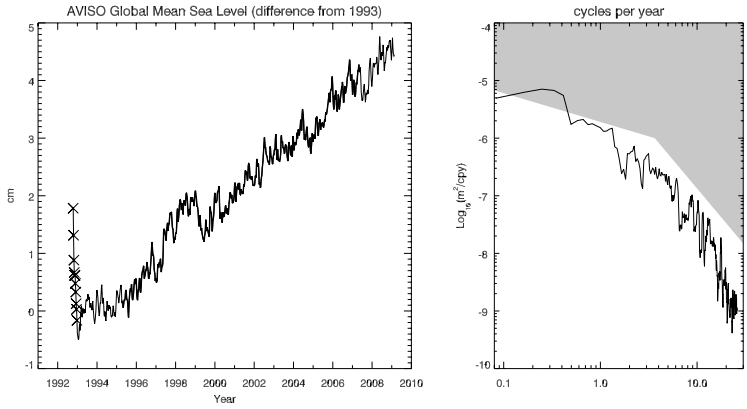

Figure 9: The time series and spectrum of the global-average sea level. The first 11 points in the time series, marked with crosses, are from a period during which there is a known error in the altimeter boresight pointing. The part of the time series used in this analysis, and in the calculation of the spectrum, is marked by a thicker line. The spectrum is shown as in Figure 3, except that there is a factor of 1000 difference in the ordinate amplitude and shading.

5 Spatial average sea level

We have, in this analysis, been considering the time series independently of each other. There is no simple way of taking the information thus generated to assess whether regional averages of sea level trends are statistically significant, as this depends also on the spatial covariances of the time series. Instead, a time series of the chosen regional average must be considered (with the region chosen on objective grounds, without reference to the observed patterns), and that time series can then be subjected to the same analysis as for any other single time series. An obvious example is the time series of global-average sea level (actually averaged over the region for which the time series is sufficiently complete as defined in Section 2, and with area-weighting of each grid point). We show this time series (with annual and semiannual cycles removed), and the spectrum of the detrended time series, in Figure 9. The thick line in the left hand panel is the part of the time series used in this study. The first eleven points, marked with crosses, are from an initial period of the TOPEX/POSEIDON altimeter during which the altimeter boresight pointing error was anomalously large [*Fu et al.*, 1994], and are clearly anomalous. The spectrum has been plotted on a scale three orders of magnitude smaller than that used in Figure 6, reflecting the fact that variability about the trend is much smaller than at a typical single point in the ocean.

For the analysis period, the standard deviation of the global sea level time series minus fitted trend is 2.50 mm. This small variability is the result of two processes. First, there is simple statistical noise reduction. This would happen most effectively if each time series was independent of every other one, as the average of a large number of independent time series tends toward zero. The second process is associated with spatial covariance. Change in global-average sea level requires a change in the volume of the ocean, whereas the variation at each grid point can be much greater as a result of redistribution of water with no net volume change. If the time series were all independent (except for a spatially-constant trend), then noise reduction would best be achieved by averaging the time series weighted inversely by the standard "error" represented by the noise. If we do such an average, then we obtain a time series with standard deviation (after detrending) of 3.83 mm assuming "white" noise, or 3.74 mm using our spectral estimate of the trend error at each point. From this it is clear the global average is reducing variability by more than just statistical averaging of noise, it is genuinely reflecting the fact that much of the variability is due to redistribution of volume.

The BIC, when applied to this spectrum, recommends the use of an AR(5) model of the spectrum. With such a model, we find that the time to 1 mm/yr accuracy (one standard deviation) is only 107 weeks. Simultaneous fitting of a trend plus AR(5) model to the time series results in a trend of 3.2 mm/yr with a standard error of 0.1 mm/yr. This statistical error estimate is 3.5 times the error estimate which would result assuming "white" noise, but remains significantly smaller than the true error, estimated from systematic effects [*Leuliette et al.*, 2004] to be about 0.4 mm/yr. Both trend and statistical error are consistent with that reported by *Cazenave et al.* [2008] for the 15-year period from the beginning of 1993 (3.1 ± 0.1 mm/yr). It is also worth noting that this is the trend of the raw time series, without any correction for the effect of glacial isostatic adjustment on the volume of the ocean (the changing shape of the ocean basins is such that, if the sea surface did not move, the volume of water would be increasing at a rate equivalent to about 0.3 mm/yr of sea level rise [*Peltier*, 2001]).

6 Summary and conclusions

We have investigated the shape of the spectrum of sea level variability based on a homogeneous subset of 12 years of the AVISO gridded altimeter dataset. We find that, while the spectrum proposed by ZW is a reasonable first order description of the average spectral shape, there are significant spatial variations in shape. These variations mean that errors in trend estimates are

not simply related to the variance of detrended sea level time series, so the variation in spectral shape must be explicitly accounted for when estimating trends (or fits to any other chosen function of time).

Spectra typically show very steep power laws at high frequencies (a few weeks to, usually, a few months), perhaps as steep as power proportional to frequency to the power -5 in places. There is often a shoulder in the spectrum, sometimes a peak and sometimes a more gradual curvature, before it settles to a gentler power law at lower frequencies. In some places, the spectrum then steepens again to an interannual peak. There are spatial variations in the amount of structure in the spectrum, as well as in the positions of the break points and steepness of power laws, and we find that a fifth-order autoregressive model is necessary to capture the spectral structure over much of the ocean.

A good visual representation of spatial variations in spectral shape and amplitude can be presented (Figure 1) by exploiting the way the eye represents spectral information as color, although only at relatively high frequencies which provide sufficient degrees of freedom. We limit this representation to periods shorter than 24 weeks for this reason. This shows a rainbow effect over latitudes equatorward of about 30°. While this is consistent with the significance of the baroclinic Rossby wave cut-off frequency, as suggested by *Lin et al.* [2008], we argue that many of the energetic features giving rise to this rainbow have already been identified as instability processes, and are therefore nonlinear effects, albeit with Rossby wave dynamics as an important component.

The visualization also emphasizes the complexity of spectra at higher latitudes, and the difference between deep ocean and shelf seas. Shelf seas typically have a bluer spectrum, indicative of the dominance of atmospheric forcing, and there is often a minimum of variance separating the different colors of deep ocean and shelf seas, showing that deep ocean sea level variations are decoupled from coastal variations at these relatively high frequencies.

Trends in sea level over the altimetry period show significant spatial variations, which appear from their patterns to be of dynamical origin. Statistical significance of these patterns means that it is a meaningful exercise to use ocean models and observations to determine the physical cause of these variations.

Appendix: Converting spectra to colors

This Appendix summarizes how to go from a spectrum to a color displayed on a monitor. The information here was mostly gleaned from Charles Poynton's

Color FAQ and Gamma FAQ documents hosted at http://www.poynton.com.

The first stage is to determine tristimulus values X, Y, and Z. These can be considered to be the three numbers which completely specify a color and its brightness. Each number represents a weighted integral of the light spectrum, with Y in particular corresponding to the perception of brightness known as luminance. The procedure for calculating these values will be described at the end. For now, assume you have X, Y, and Z, and want to work out how to display the corresponding color on a monitor, by sending it three bytes corresponding to r, g and b values.

This is an inverse problem. You know the color you wish to achieve, and you know (in principle) the colors of the three (red, green, blue) phosphors used by the display. The question is to find the brightnesses of each of these (more precisely, the number to send to the software to produce the appropriate brightness) you need in order to produce the correct color. To work this out, it is useful to work with an intermediate set of numbers R, G, B which are linear functions of X, Y, Z. These are conventionally specified such that $R = G = B = 1$ corresponds to white light at full brightness, according to the monitor's chosen definition of white (this may actually be quite a bluish color). We now seek the matrix $\mathbf{A}$ such that

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \mathbf{A} \cdot \begin{pmatrix} R \\ G \\ B \end{pmatrix}. \tag{1}$$

Since $\mathbf{A}$ consists of 9 numbers, we need 9 independent pieces of information to determine it. However, the information (the colors of the three phosphors) needed to work out X, Y, Z from R, G and B is conventionally given in a slightly strange form. The color of each phosphor is specified in terms of its chromaticity x, y, $z = 1-x-y$ where $x = X/(X+Y+Z)$, $y = Y/(X+Y+Z)$, $z = Z/(X+Y+Z)$, which tell you about the color without luminance information (only x and y may be given as z can be derived from them), and relative weights for the phosphors are specified by the chromaticity chosen to represent white. So we have chromaticities x_r, y_r, z_r for red, x_g, y_g, z_g for green, x_b, y_b, z_b for blue, and x_w, y_w, z_w for white, giving eight independent numbers (because $z = 1 - x - y$). The final piece of information is the fact that $R = G = B = 1$ gives white light with $Y = 1$. This is sufficient, after some manipulation, to specify $\mathbf{A}$, which can then be inverted to give the conversion from XYZ to RGB:

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \mathbf{A}^{-1} \cdot \begin{pmatrix} X \\ Y \\ Z \end{pmatrix}. \tag{2}$$

As an example (as used in the calculations in this paper), Poynton recommends the use of a standard he calls Rec. 709, for which

$$\begin{array}{llll} x_r = 0.640 & x_g = 0.300 & x_b = 0.150 & x_w = 0.3127 \\ y_r = 0.330 & y_g = 0.600 & y_b = 0.060 & y_w = 0.3290 \\ z_r = 0.030 & z_g = 0.100 & z_b = 0.790 & z_w = 0.3582, \end{array} \tag{3}$$

which leads to

$$\mathbf{A} = \begin{pmatrix} 3.2409699 & -1.5373832 & -0.49861079 \\ -0.96924375 & 1.8759676 & 0.041555082 \\ 0.055630032 & -0.20397685 & 1.0569714 \end{pmatrix}. \tag{4}$$

and

$$\mathbf{A}^{-1} = \begin{pmatrix} 0.41239081 & 0.35758433 & 0.18048081 \\ 0.21263903 & 0.71516866 & 0.072192319 \\ 0.019330821 & 0.11919473 & 0.95053222 \end{pmatrix}. \tag{5}$$

Note: this is our calculation, which differs from Poynton's in the 5th significant figure, perhaps because the white chromaticities given in (A3) are truncated at 4 significant figures.

Inserting XYZ into (A2) gives us the RGB values, but there is a potential problem: what if one or more values is either negative or greater than 1? We will consider how to fix negative values first, and then values greater than 1.

It may seem surprising that negative values can arise. The reason for this is that no three colors are sufficient to produce all possible colors. While the XYZ values derived from any power spectrum must be positive, and the transform from RGB to XYZ is bound to result in positive values, that is not the case for the inverse transform. Physically, what is meant by a negative value for R (say) is that the target color cannot be matched using the chosen set of three colors, but if some red was added to the target color then a match could be found. The range of possible colors for a given set of red, green and blue chromaticities is called the gamut, so negative values are referred to as an "out of gamut" error.

So, we have a color which cannot be displayed. What can we do about it? Our chosen option is to do as little damage as possible to the color and brightness, but accept a reduction in saturation of the color. This is achieved by 1) adding just enough white (equal amounts of R, G and B) to remove the negative value, and then 2) rescaling the luminance back to its original value by applying the same factor to R, G and B. This is achieved by calculating $\delta = -\min(R, G, B, 0)$. If $\delta > 0$ then calculate fixed values R_1, G_1 and B_1 from

$$(R_1, G_1, B_1) = (R + \delta, G + \delta, B + \delta)\frac{Y}{Y + \delta}. \tag{6}$$

Now we have to deal with large values. These could have been avoided by scaling everything at the beginning, but that is often not what is wanted as it can result in pictures which are dark except for very small regions. Given that we are going to allow large values to occur, how should we rescale them for display?

There are various options, depending on the application. In our pictures, where preserving color is more important than brightness, our preferred option is to rescale R, G and B by dividing by the largest of the three, if the largest is greater than 1. This means that brightness can be compromised, but leaves meaningful colors at all points (meaning that white cannot occur as a result of over-saturation of the brightness).

So far we have been working in linear RGB. That means that images can be added by adding the RGB components (before applying out of gamut corrections and limiting the maximum), and the correct colors will result. However, for reasons of efficiency, it is usual for the monitor to have a nonlinear response to input numbers, with brightness being proportional to input number raised to some power γ (but with a linear tail added at low values). This means that the inverse function must be applied to RGB to give the nonlinear $R'G'B'$ which must be presented to the monitor. While the actual response usually has a γ of about 2.5, it is usual to extend the dynamic range slightly by using a somewhat smaller γ for the inverse transform, the standard value being $20/9 = 2.2\dot{2}$, or $1/\gamma = 0.45$.

The inverse transform takes the form

$$\begin{aligned} R' = &\quad hR &\quad R \geq R_c \\ R' = &\quad fR^{1/\gamma} + g &\quad R \leq R_c, \end{aligned} \tag{7}$$

where the coefficients are chosen so that $R' = 1$ where $R = 1$ and at the matching value $R = R_c = 0.018$, the curve of R' against R is continuous and has a continuous first derivative. This results in

$$\begin{aligned} h = &\quad 4.506813168 \\ g = &\quad -0.09914989 \\ f = &\quad 1.09914989. \end{aligned} \tag{8}$$

Once this transform has been applied to give $R'G'B'$ each in the range 0 to 1, these can be linearly rescaled and converted into byte variables in the range 0 to 255 to give the usual rgb values used in 24-bit color representation. This is the final product, to be sent to the display.

Finally, we return to the issue of how to determine the tristimulus values X, Y and Z. These are determined from the Spectral Power Distribution

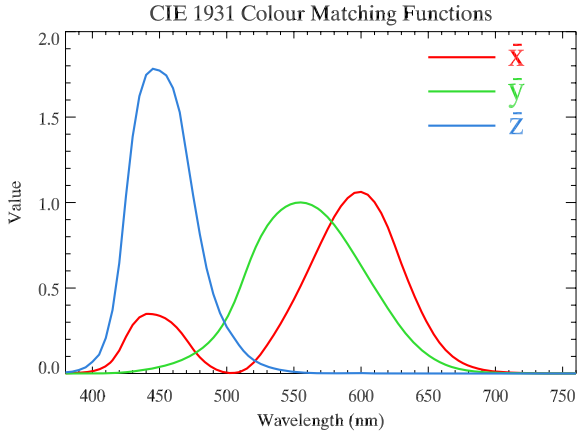

Figure 10: The CIE 1931 color matching functions $\overline{x}$, $\overline{y}$ and $\overline{z}$.

(SPD) function $I(\lambda)$, which describes the power emitted per unit area of the object (illuminant) per nanometre of wavelength, as a function of wavelength (units W m^{-2} nm^{-1}). This is convolved with the CIE 1931 color matching functions [*CIE*, 1971] $\overline{x}$, $\overline{y}$, $\overline{z}$ (also functions of wavelength) to give XYZ values, thus

$$X = \int_0^\infty I(\lambda)\overline{x}(\lambda)\, d\lambda, \tag{9}$$

with analogous expressions for Y and Z. The color matching functions are tabulated at 5 nm intervals, and are shown in Figure 10. A spectrum in which I is independent of wavelength gives a color quite close to white, and results in equal values for X, Y and Z.

The input to this convolution is the SPD $I(\lambda)$ rather than the Power Spectral Density (PSD) $S(\sigma)$. The two can be related by noting that the power P between two wavelengths λ_1 and λ_2 is given by

$$P = \int_{\lambda_1}^{\lambda_2} I(\lambda)\, d\lambda, = \int_{\sigma_2}^{\sigma_1} S(\sigma)\, d\sigma, \tag{10}$$

where, for light, $\sigma = c/\lambda$, so that $\sigma_1 = c/\lambda_1$ and $\sigma_2 = c/\lambda_2$. Noting that $d\lambda = -c\, d\sigma/\sigma^2$, this allows us to identify that

$$I = S\frac{\sigma^2}{c}. \tag{11}$$

When the spectrum is not light, but some other quantity (such as sea-level variability), the amplitude becomes rather arbitrary, and the speed of light in (A11) is clearly no longer relevant. We use a periodogram PSD

estimate (normalized to be in units of m^2/Hz) multiplied by σ^2 (frequency in Hz) to obtain the equivalent of I, and then linearly rescale the chosen range of periods onto the 380–760 nm range of the CIE color matching functions, before proceeding as for a light spectrum. The resulting values of XYZ are then normalized for our plots, dividing by a factor of 1.6×10^6 (the factor is six times smaller for Figure 2). We also brighten the darker regions by dividing each of X, Y, Z by $Y^{0.2}$ (equivalent to raising luminance Y to the power 0.8 while keeping chromaticity constant), as this makes it possible to display a wider range of brightness. We have investigated the impact of smoothing the PSD estimate, but it makes little difference to the final image.

Acknowledgments

Thanks to Mark Tamisiea and Jo Williams for discussions and input during the preparation of this paper. Thanks also to two anonymous referees and the editor, Frank Bryan, who made valuable suggestions. This work was funded by the NERC, as part of Oceans 2025 and the National Centre for Earth Observation.

References

[1] Akaike, H. (1974), A new look at the statistical model identification. *IEEE Trans. Automatic Control, 19* 716–723 doi:10.1109/TAC.1974.1100705.

[2] Bingham, R. J. and C. W. Hughes (2008), Determining North Atlantic meridional transport variability from pressure on the western boundary: A model investigation. *J. Geophys. Res., 114*, C09008, doi:10.1029/2007JC004501.

[3] Bingham, R. J. and C. W. Hughes (2009), The signature of the Atlantic meridional overturning circulation in sea level along the east coast of North America. *Geophys. Res. Lett., 36*, L02603, doi: 10.1029/2008GL036215.

[4] Carrère, L. and F. Lyard (2003), Modeling the barotropic response of the global ocean to atmospheric wind and pressure forcing — comparisons with obervations. *Geophys. Res. Lett., 30*, 1275, doi: 10.1029/2002GL016473.

[5] Cazenave, A., A. Lombard and W. Llovel (2008), Present-day sea level rise: A synthesis. *C. R. Geosci., 340* 761–770.

[6] Chelton, D. B. and M. G. Schlax (1996), Global observations of oceanic Rossby waves. *Science, 272* 234–238.

[7] Chelton, D. B., R. A. deSzoeke, M. G. Schlax, K. El Naggar and N. Siwertz (1998), Geographical variability of the first-baroclinic Rossby radius of deformation. *J. Phys. Oceanogr., 28*, 433–460.

[8] Chelton, D. B., M. G. Schlax, M. H. Freilich and R. F. Milliff (2004), Satellite measurements reveal persistent small-scale features in ocean winds. *Science, 303*, 978–983.

[9] Church, J. A., N. J. White and J. R. Hunter (2006), Sea-level rise at tropical Pacific and Indian Ocean islands. *Global Planet. Change, 53*, 155–168, doi: 10.1016/j.gloplacha.2006.04.001.

[10] Cunningham, S. A., T. Kanzow, D. Rayner, M. O. Baringer, W. E. Johns, J. Marotzke, H. R. Longworth, E. M. Grant, J. J.-M. Hirschi, L. M. Beal, C. S. Meinen and H. L. Bryden (2007), Temporal variability of the Atlantic Meridional Overturning Circulation at 25.5°N. *Science, 317* 935–938, doi:10.1126/science.1141304.

[11] CIE (Commission Internationale de l'Eclairage) (1971), Colorimetry, official recommendations of the International Commission on Illumination. Publication CIE No. 15 (E-1.3.1). Paris, France: Bureau Central de la CIE, 4 Av. du Recteur Poincaré, 75782 Paris Cedex 16.

[12] Douglas, B. C. (1991), Global sea level rise. *J. Geophys. Res., 96*(C4), 6981–6992.

[13] Fu, L. -L. (2003), Wind-forced intraseasonal sea level variability of the extratropical oceans. *J. Phys. Oceanogr., 33*, 436–449.

[14] Fu, L. -L., E. J. Christensen, C. A. Yamarone Jr., M. Lefebvre, Y. Ménard, M. Dorrer and P. Escudier (1994), TOPEX/POSEIDON mission overview. *J. Geophys. Res., 99*, 24,269–24,381.

[15] Häkkinen, S. (2001), Variability in sea surface height: A qualitative measure for the meridional overturning in the North Atlantic. *J. Geophys. Res., 106*(C7), 13,837–13,848.

[16] Harrison, C .G. A. (2002), Power spectrum of sea level changes over fifteen decades of frequency. *Geochem. Geophys. Geosyst., 3*(8), 1047, doi: 10.1029/2002GC000300.

[17] Hughes, C. W., M. S. Jones and S. Carnochan (1998), Use of Transient Features to Identify Eastward Currents in the Southern Ocean. *J. Geophys. Res., 103* 2929–2944.

[18] Kay, S. M. (1988), Modern Spectral Estimation: Theory and Application, Englewood Cliffs, NJ, Prentice-Hall.

[19] Legeckis, R. (1977), Long waves in the Eastern Equatorial Pacific Ocean: A view from a geostationary satellite. *Science, 197*, 1179–1181.

[20] Le Traon, P. Y. (1990), Time scales of mesoscale variability and their relationship with space scales in the North Atlantic. *J. Mar. Res., 49*, 467–492.

[21] Leuliette, E. W., R. S. Nerem and G. T. Mitchum (2004), Calibration of TOPEX/Poseidon and Jason altimeter data to construct a continuous record of mean sea level change. *Mar. Geod., 27*, 79–94.

[22] Liddle, A. R. (2007), Information criteria fro astrophysical model selection. *Mon. Not. R. Astron. Soc. 377*, L74–L78, doi: 10.1111/j.1745.3933.00306.x

[23] Lin, X., J. Yang, D. Wu and P. Zhai (2008), Explaining the global distribution of peak-spectrum variability of sea surface height. *Geophys. Res. Lett., 35*, L14602, doi:10.1029/2008GL034312.

[24] Marcos, M. and M. N. Tsimplis (2007), Forcing of coastal sea level rise patterns in the North Atlantic and the Mediterranean Sea. *Geophys. Res. Lett., 34*, L18604, doi:10.1029/2007GL030641.

[25] Mitrovica, J. X., M. E. Tamisiea, J. Davis and G. A. Milne (2001), Recent mass balance of polar ice sheets inferred from patterns of global sea-level change. *Nature, 409*, 1026–1029.

[26] Peltier, W. R. (2001), Global Glacial Isostatic Adjustment and Modern Instrumental Records of Relative Sea Level History, pp 65–95 in *Sea Level Rise: History and Consequences*, B. C. Douglas, M. S. Kearney and S. P. Leatherman (eds), Academic Press.

[27] Plag, H. P. (2006), Recent relative sea-level trends: an attempt to quantify the forcing factors. *Phil. Trans. Roy. Soc. A, 364*, 821–844.

[28] Qiu, B. and S. Chen. (2004), Seasonal modulations in the eddy field of the South Pacific. *J. Phys. Oceanogr., 34*, 1515–1527.

[29] Qiu, B., S. Chen and W. S. Kessler. (2009), Source of the 70-day mesoscale eddy variability in the Coral Sea and the North Fiji Basin. *J. Phys. Oceanogr., 39*, 404–420.

[30] Qiao, L. and R. H. Weisberg (1995), Tropical instability wave kinematics: Observations from the Tropical Instability Wave Experiment. *J. Geophys. Res., 100*, 8677–8693.

[31] Schouten, M. W., W. P. M. de Ruijter and P. J. van Leeuwen (2002), Upstream control of Agulhas Ring shedding. *J. Geophys. Res., 107*, 3109, doi:10.1029/2001JC000804.

[32] Schwarz, G. (1978), Estimating the dimension of a model. *Ann. Stat., 6*, 461–464.

[33] Shennan, I. and P. L. Woodworth (1992), A comparison of late Holocene and twentieth century sea-level trends from the UK and North Sea region. *Geophys. J. Int., 109*, 96–105.

[34] Stammer, D. (1997), Global characteristics of ocean variability estimated from regional TOPEX/POSEIDON altimeter measurements. *J. Phys. Oceanogr., 27*, 1743–1769.

[35] Stoica, P. and R. Moses (1997), Introduction to Spectral Analysis, Prentice-Hall.

[36] Tulloch, R., J. Marshall and K. Shafer Smith (2009), Interpretation of the propagation of surface altimetric observations in terms of planetary waves and geostrophic turbulence. *J. Geophys. Res., 114*, C02005, doi:10.1029/2008JC005055.

[37] Woodworth, P. L., C. W. Hughes, D. L. Blackman, V. N. Stepanov, S. J. Holgate, P. R. Foden, S. Mack, G. W. Hargreaves, M. P. Meredith and G. Milenevsky (2006), Antarctic Peninsula sea levels: A real-time system for monitoring Drake Passage transport. *Antarctic Sci., 18*, 429–436, doi:10.1017/S0954102006000472.

[38] Wunsch, C. (2009), The oceanic variability spectrum and transport trends. *Atmosphere-Ocean, 47*(4), 281–291, doi:10.3137/OC310.2009.

[39] Wunsch, C. and P. Heimbach (2006), Estimated decadal changes in the North Atlantic Meridional Overturning Circulation and heat flux 1993–2004. *J. Phys. Oceanogr., 36*, 2012–2024.

[40] Zang, X. and C. Wunsch,(2001), Spectral description of low-frequency oceanic variability. *J. Phys. Oceanogr., 31*, 3073–3095.