Appendix: Spectral Stabilisation

When we wish to examine the relation between different geophysical fields it is desirable to have each of them in a common representation. We can achieve this by using a down-sampling approach implemented in the spectral domain, which has the additional merit of suppressing noise and thereby increasing the reliability of the longest wavelength components of the field. Such long wavelength components play an important role in many comparisons between fields

For studies at the continental scale, most available datasets are oversampled relative to the information content required by the particular methodology. For Australian datasets available, the applicable resolutions are:

Gravity each pixel: 800 x 800 m (Bacchin et al., 2008)

Magnetic intensity each pixel: 80 x 80 m (Milligan et al., 2010)

Onshore topography each pixel: 30 x 30 m (internal Geoscience Australia data)

Combined on and offshore topography each pixel: 250 x 250 m (internal Geoscience Australia data)

Our approach has been to work with a 4096 x 4096 pixel representation of each field using a resolution with pixels 1.28 km square. This is sufficient to retain much detail, but also is a very manageable size for further spectral domain processing, such as extraction of derivatives or upward continuation.

Rather than simply reduce the data density in each direction by the appropriate factor, we employ an ensemble of multiple realisations for a dataset produced by sub-sampling. By selecting existing pixels from the dataset for larger 'master pixels', we can achieve multiple realisations for the random processes that yielded the original data (Bendat and Piersol, 2010).This approach provides stable long wavelength components at the expense of some spatial resolution and the highest frequency components.

Consider a two-dimensional dataset with nX x nY original pixels, with a target number of pixels of X x Y. We thus wish to subsample by a factor of n in each dimension. The first stage is to pad the dataset with additional zeroes so that number of pixels in each dimension is an integer multiple of n. As a simple example, consider a 7 x 6 pixel image that we wish to subsample by a factor of 4 in each dimension. We begin by padding this to 8 x 8 pixels.

In the general case, we now define master pixels where the number of master pixels in each dimension is the number of zero-padded pixels divided by n. For the simple example, we have 2 master pixels in each dimension. We now select, from our data, each possible value from each sub-pixel that relates to each master pixel. This provides n^2 realisations of the original data. For our trivial example, we have 4 x 4 = 16 realisations in the spatial domain. From these n^2 realisations in the spatial domain, we take the 2-dimensional Fourier transform (FFT) of each of these realisations, to compile 16 Fourier representations of the sub-sampled data. Each of these realisations is an equally valid representation of the original data in the frequency domain. From these multiple realisations of the subsampling, we can average to produce a stabilised representation of the data in the spectral domain. This averaging is performed for each Fourier component in the final set to produce a single value with reduced noise compared with the individual subsamples. In addition to the average, we can collect other statistics on the various wavelength components.

As an example, consider the 5th edition of the magnetic map of Australia, with variable reduction to the pole to the pole using the algorithm of Cooper and Cowan (2005). The preliminary step is to project the data to an area-preserving projection (Albers Equal-Area projection: Snyder, 1987, European Petroleum Survey Group coordinate reference 3577). The dataset is then padded with zeroes and centred in a space of 64000 x 64000 pixels. The pixels in the original dataset pixels are 80 m x 80 m; by applying a 16 x 16 sub-sampling procedure there are 256 separate realisations of the magnetic data that can be averaged to provide a spatial resolution of 1280 m x 1280 m pixels. After sub-sampling and averaging these 256 frequency-domain realisations, we can obtain statistics about the distribution of Fourier components. The level of detail recovered after inverse Fourier transforming the data to the spatial domain is higher (but also smoother) than the resolution achieved by spatial down-sampling (Figure A.1).

This example shows a portion of the magnetic anomaly dataset in Western Australia, with a rich collection of features at many spatial scales and orientation. Figure A.1(A) displays the original dataset employing 80 m x 80 m pixels. The other two images show the effect of down-sampling to 1280 m x 1280 m, i.e. a reduction in sampling density of 16 in both latitude and longitude. Figure A.1(B) illustrates the result with spectral stabilisation (using 256 spectral sub-samples) and Figure A.1(C) uses simple 16-fold data resampling.

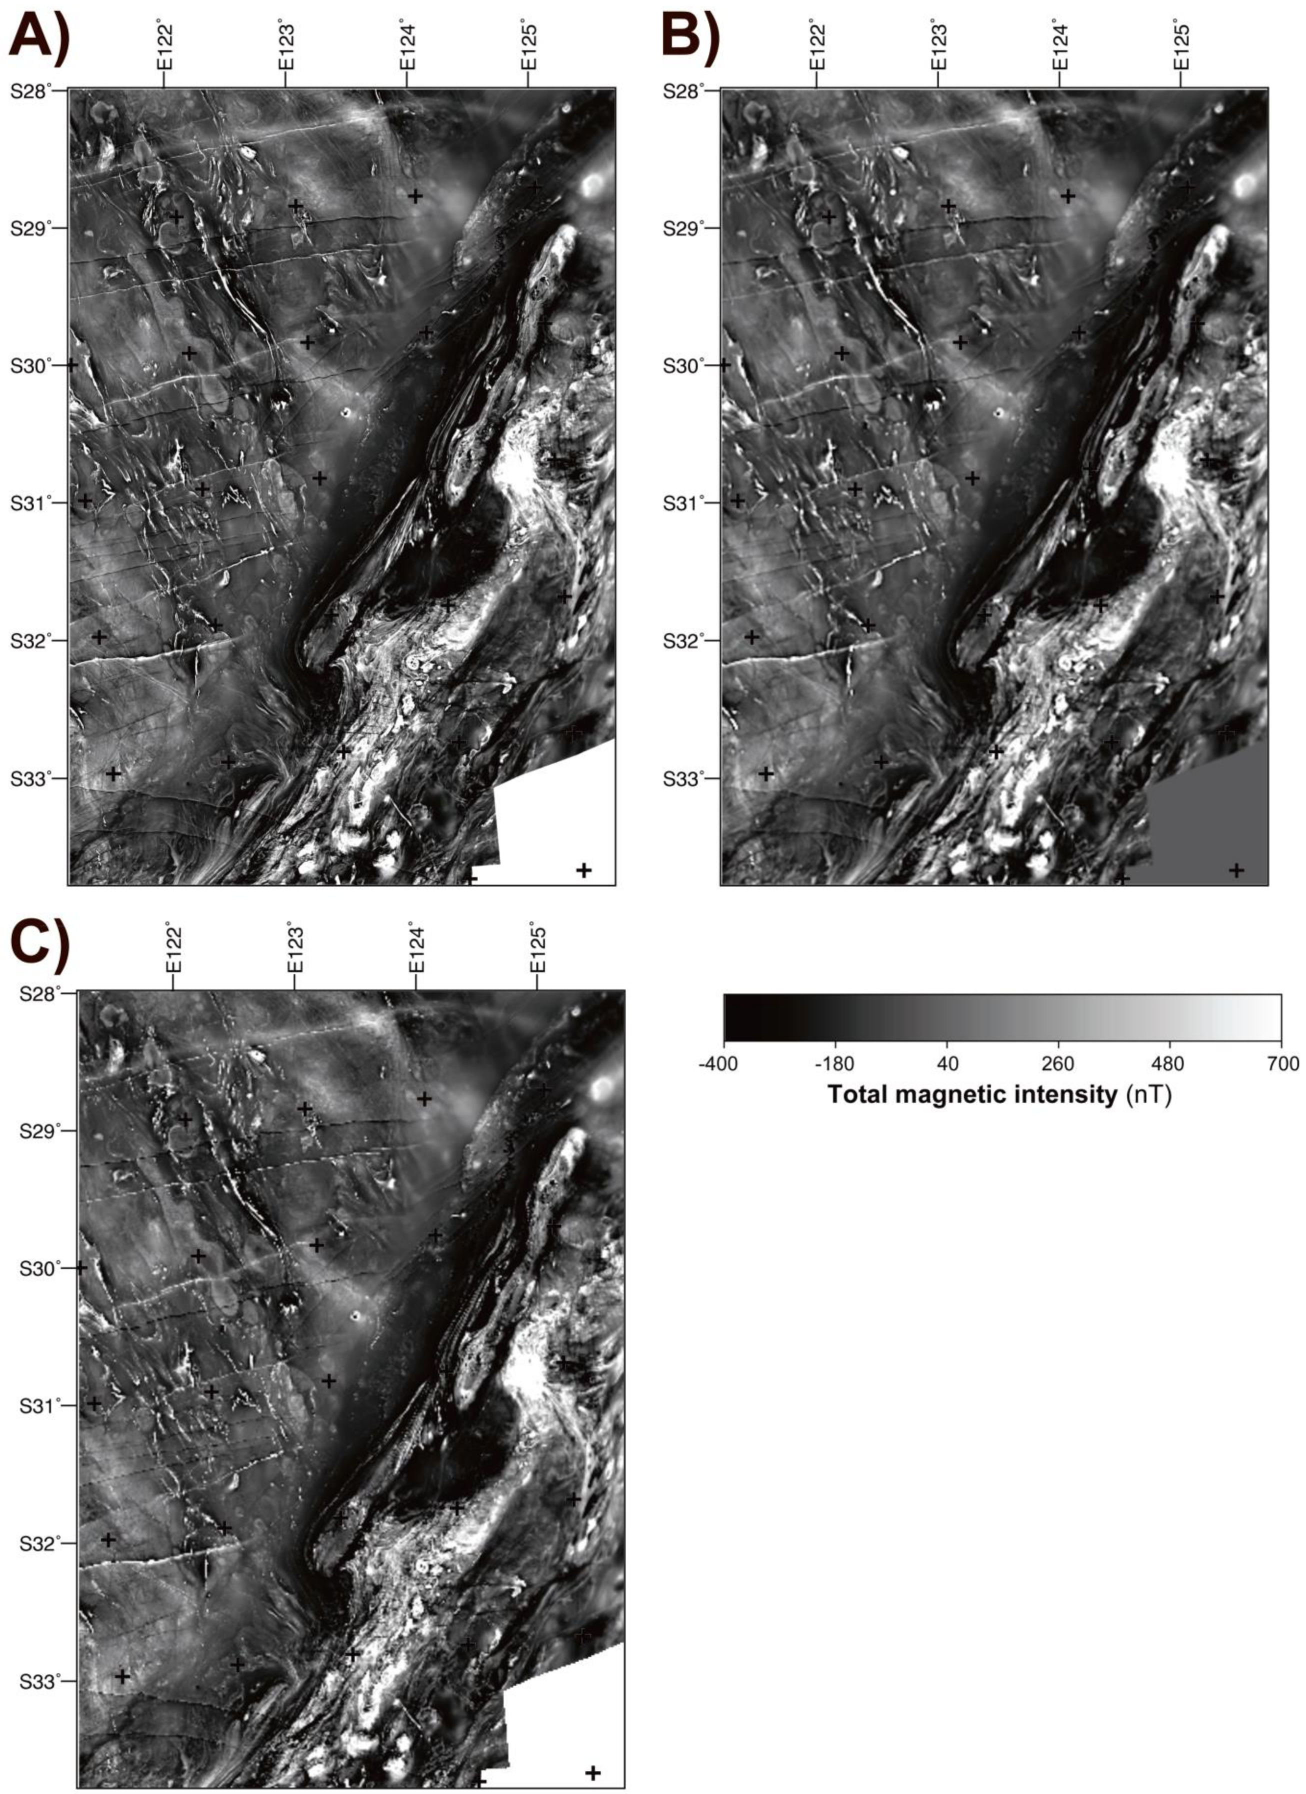

Figure A.1: A) Magnetic data in Western Australia extracted from the magnetic anomaly map of Australia, using the original 80 x 80 m pixels. B) Same data extent as A) but 1280 x 1280 m pixels, derived from spectral domain averaging of 256 realisations. C) Same data extent as A) but 1280 x1280 m pixels, spatially resampled only.

From Figure A.1 we see that with the spectral stabilisation approach it is possible to preserve the strong contrasts in the original image more effectively and retain detail at a level approaching the pixelation employed. The procedure requires significantly more computational effort than simple spatially resampling, but has the added benefit of providing statistics for the spectral components.

The representation of data in this fashion has an additional benefit beyond stabilising the longest wavelength components: by down-sampling both gravity and magnetic data in the same fashion but to varying degrees, we can maintain comparable frequency content between the two datasets. This assists in comparison of these two datasets, since most geophysical filtering techniques work in the spectral domain.

The process of spectral stabilisation can also be employed in a hierarchical fashion to provide a representation of the variation in spectral content across the continent. Thus we can successively halve windows across the continental dataset, and use 64-fold, then 16-fold and finally 4-fold averaging for the spatial spectra at each step. In this way a stable calculation of the longest wavelength features that can be represented at each window scale is achieved, along with a measure of the uncertainty in the 1-D power spectra. This process has been employed by Chopping and Kennett (2015) in a study of the variation of Curie depth across the continent (see Chapter 4), and also provides the foundation for comparisons of the gravity and topography fields (Chapter 12).

This text is taken from *The Australian Continent: A Geophysical Synthesis*,
by B.L.N. Kennett, R. Chopping and R. Blewett, published 2018 by ANU Press
and Commonwealth of Australia (Geoscience Australia), Canberra, Australia.