Re-examining the global distribution of seamounts using neural network techniques

Andrew Valentine[1], Lara Kalnins[2] & Jeannot Trampert[1].
andrew@geo.uu.nl • www.geo.uu.nl/∼andrew

Summary

Many problems in geomorphology require researchers to identify and catalogue occurrences of some particular class of feature. This identification may be straightforward 'by eye', but is typically difficult to automate, since simple mathematical models for the feature fail to adequately account for natural variations between examples.

As an example, we consider the task of identifying seamounts (isolated bathymetric highs of volcanic origin) from global bathymetric data. We propose a new approach to this problem, using neural networks to assimilate the topographic characteristics of hand-picked seamounts. This may then be used to assess whether a particular locations, and thus the network can be used to extrapolate the selection across a regional or global dataset.

Neural networks

Neural networks allow complex mathematical relationships to be discovered and modelled.

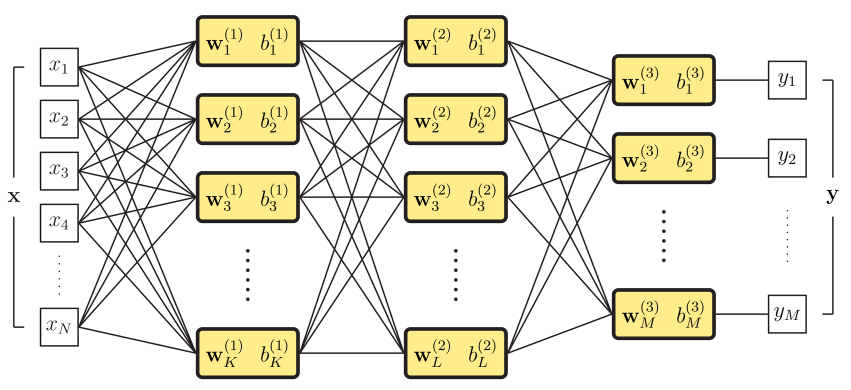

- Individual neurons implement simple mappings: $\mathbf{u} \rightarrow f(\mathbf{w}\cdot\mathbf{u}+b)$ where $\mathbf{w}$ represents some set of tunable *weights*, and b a *bias*.
- By connecting many neurons together, arbitrarily complex mappings $\mathbf{x} \rightarrow \mathbf{y}$ may be represented.
- The mapping is governed by the (independent) weights and biases associated with each neuron.
- Given a set of example inputs $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}$ and corresponding desired outputs, $\{\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_N\}$, it is possible to set up an optimisation problem to find values for the tunable parameters. This is often called *training* the network. No other knowledge of the relationship between the $\{\mathbf{x}_i\}$ and $\{\mathbf{y}_i\}$ is required.
- Once the network has been trained, it may be used to *predict* $\mathbf{y}$ for previously unseen input vectors, $\mathbf{x}$.

The autoencoder

An autoencoder is a particular class of neural network (see schematic, ***right***), designed to find low-dimensional representations of complex datasets (Hinton & Salukhutdinov, 2006).

- Can be considered as a connected encoder-decoder pair;
- Encoder layers have an output vector of lower dimension than their inputs;
- Decoder layers have an output vector of higher dimension than their inputs;
- The network is trained so that the top-level decoder layer produces a *reconstruction* ($\mathbf{s}_i'$) as close as possible to the input data ($\mathbf{s}_i$), for some *training set* $\{s_1, s_2, \ldots, s_N\}$: we minimise

$$E = \sum_i \left|\mathbf{s}_i' - \mathbf{s}_i\right|^2$$

- The encoding process forces the network to identify the characteristic features of the training set, and discard finer details;
- If a new example, $\mathbf{s}_n$ is 'similar' to those in the training set, the network will be able to represent it, and $E_n = \left|\mathbf{s}_n' - \mathbf{s}_n\right|^2$ will be small;
- If the new example is not similar to those in the training set, E_n will be large;
- The autoencoder can therefore be used to assess whether new examples are comparable with those in the training set (as discussed in Valentine & Trampert, 2012).

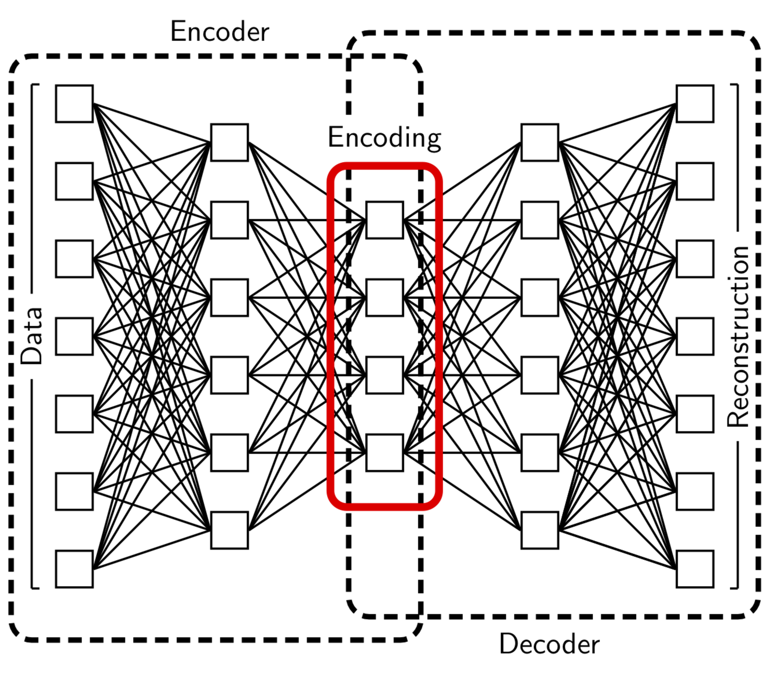

From Valentine & Trampert (2012)

Data: 'patches' of seafloor

To search for seamounts, we systematically test 'patches' (spatial windows) of seafloor bathymetry to assess whether they are likely to be centred on a seamount.

- Data is extracted from the global bathymetry dataset of Smith & Sandwell (1997, v.14.1);
- Each patch is 150 km×150 km, sampled on a 64×64 grid;
- Raw data is projected using a Lambert Azimuthal projection, filtered using a tenth-order Butterworth low-pass filter with frequency cutoff 2.5× Nyquist's frequency, and then resampled onto the 64×64 grid;
- The edges of the patch are downweighted relative to the centre, with scale factor $\alpha(r) = 1 - ar^2$;
- Data extraction is the most computationally expensive part of the process, requiring ∼1 s per patch—although well-suited to a distributed implementation.

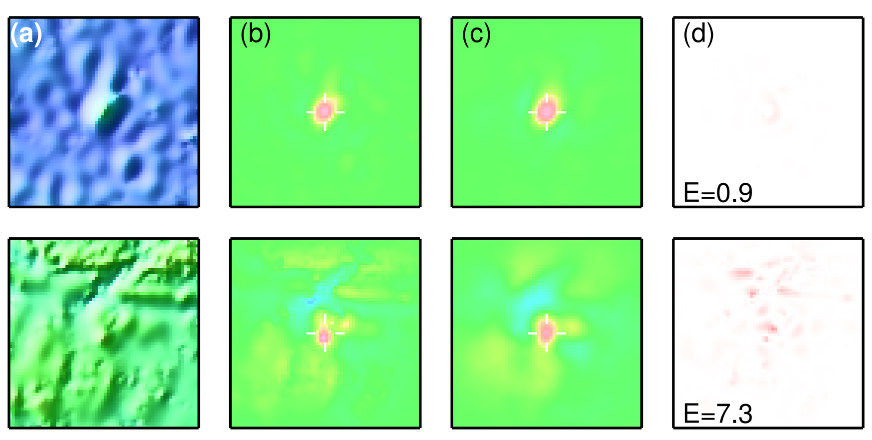

Left: Examples of patches centred on seamounts. From left to right: (a) raw data—150 km×150 km region extracted from global bathymetry; (b) autoencoder inputs—rescaled, weighted bathymetry; (c) inputs reconstructed by trained autoencoder; (d) difference between inputs and outputs.

To construct a training set, we use the seamount catalogue developed by Kim & Wessel (2011) in conjunction with visual inspection to assemble a collection of 1000 seamount-centred patches. This is then used to train an autoencoder that represents 4096-element input data vectors as 64-element encodings.

A systematic search for seamounts

The trained autoencoder can represent previously unseen patches containing seamount-like structures; however, non-seamount topography cannot be represented.

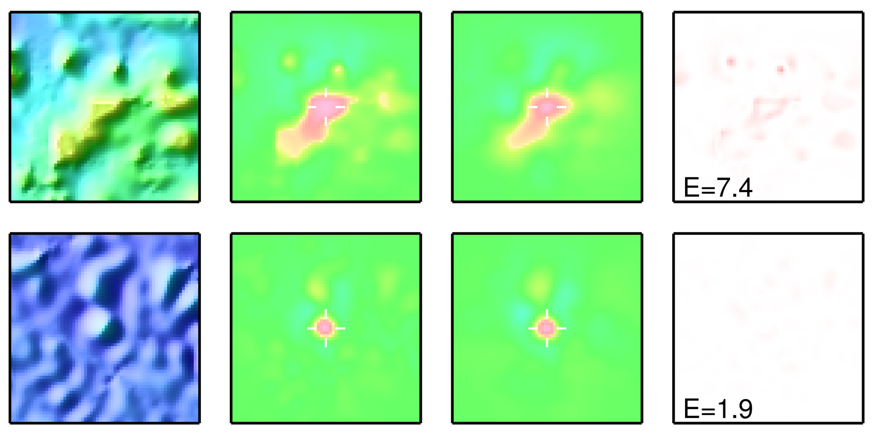

Left: Seamount-like patches can be encoded and reconstructed by the trained network with relatively little loss of information: the difference between inputs and outputs is small.

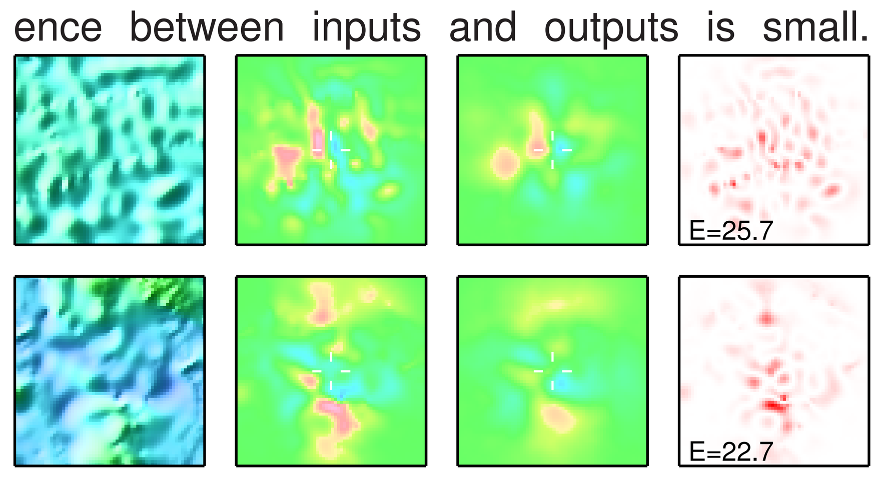

Right: Attempting to encode and reconstruct patches that are not obviously centred on seamounts leads to large reconstruction errors.

We can use this to quantify how well a given site matches the collection of seamounts used during network training. At each point on a 2D grid, we extract a patch, encode and reconstruct it, and compute the reconstruction error. We therefore obtain a grid of reconstruction error, and low values are assumed to correspond to seamounts.

Example — Pacific Ocean

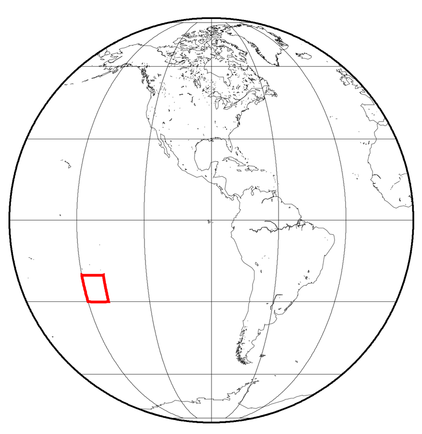

To demonstrate the method, we apply it to a 10°×10° region in the Pacific Ocean (as outlined on map, ***left***). The Sandell & Smith 1-minute bathymetry for the region is shown, ***right***.

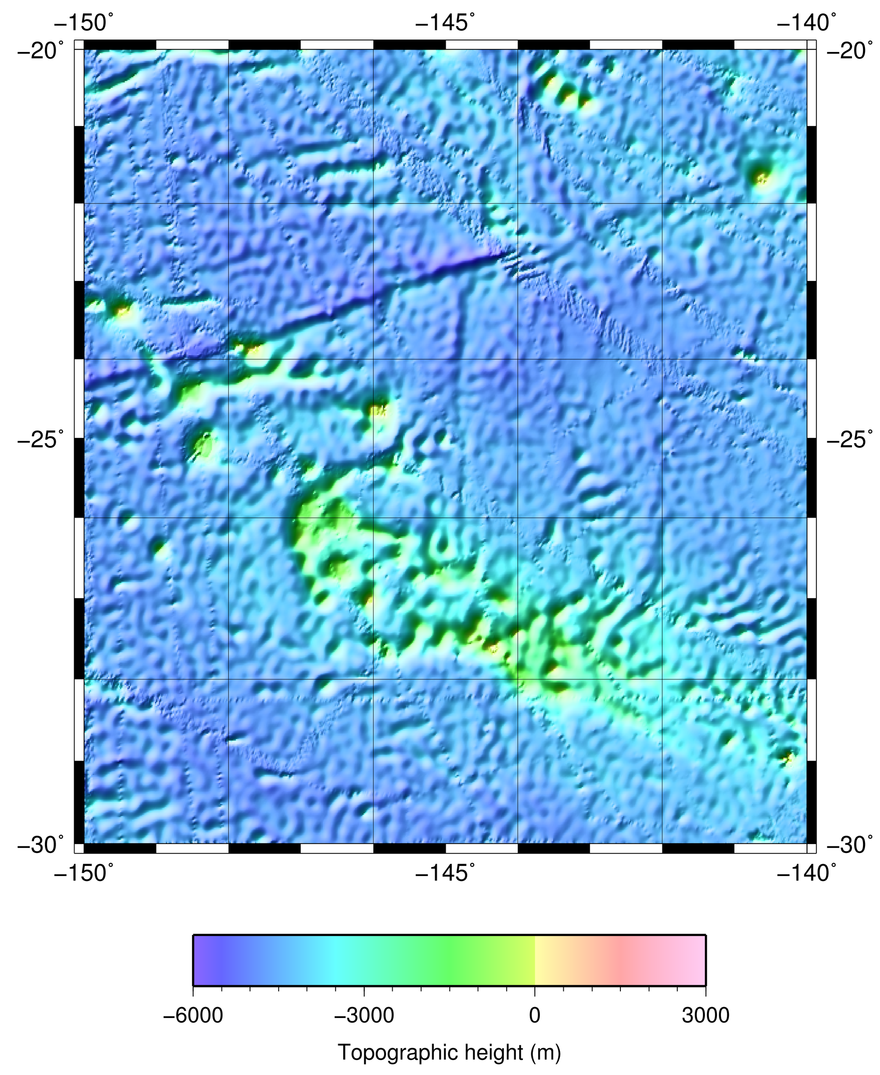

Below: A pre-existing catalogue: seamounts identified by Kim & Wessel (2011). Note that this catalogue is derived from vertical gravity gradient (VGG) data, and this may lead to some discrepancy with seamounts identifiable in bathymetric data.

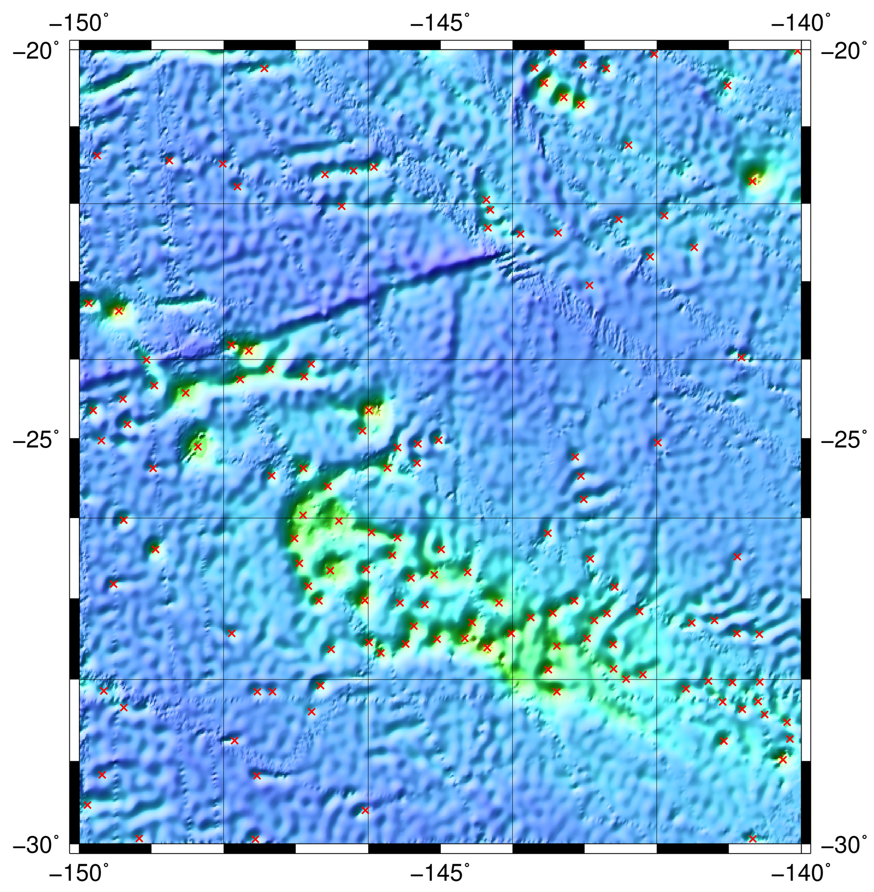

Our method involves testing each point on the 1-minute bathymetric grid to assess the likelihood that it lies at the centre of a seamount.

- We extract a patch centred on each point in the grid, encode and reconstruct this using the trained network, and compute the reconstruction error;
- The resulting grid of reconstruction error can be seen ***below left***; where this is low (yellow/red colours), the network is able to accurately represent local bathymetry;
- Since the network is trained only on seamount patches, low reconstruction error is indicative of seamount-like topography;
- ***Below right*** the same error grid is shown as contours overlain on bathymetry.

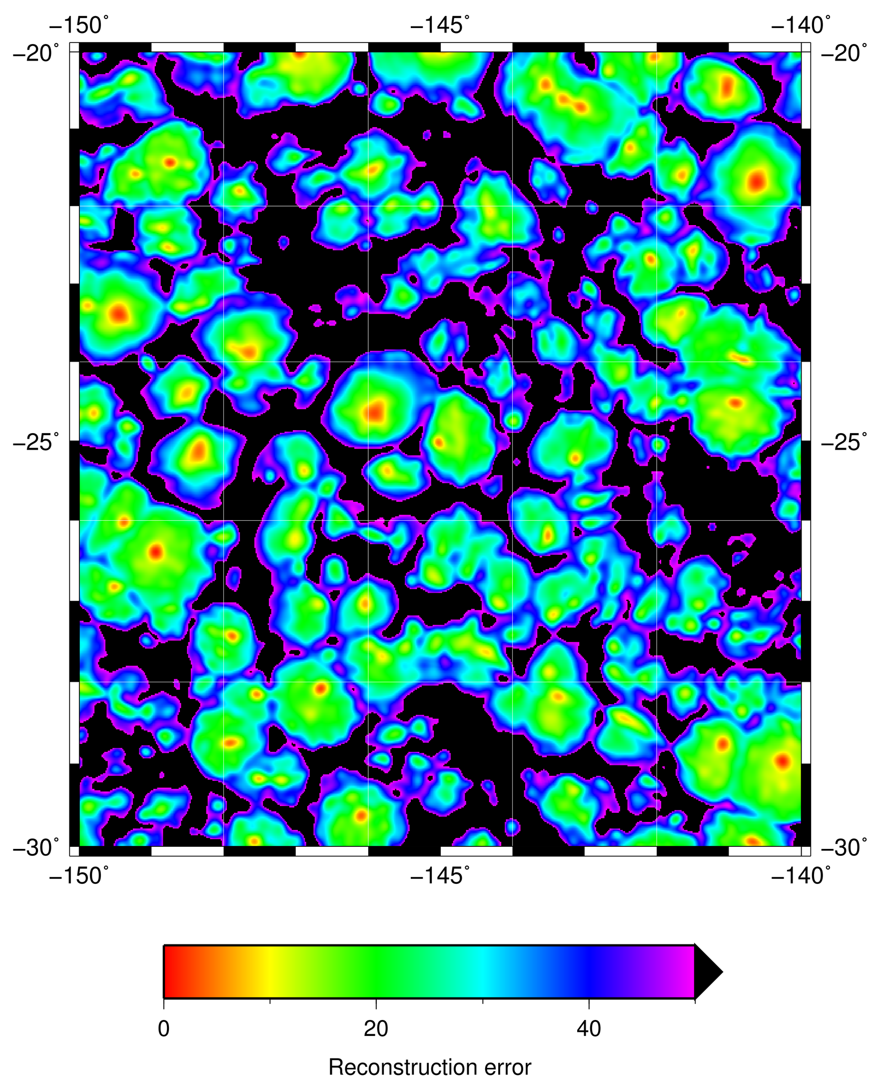

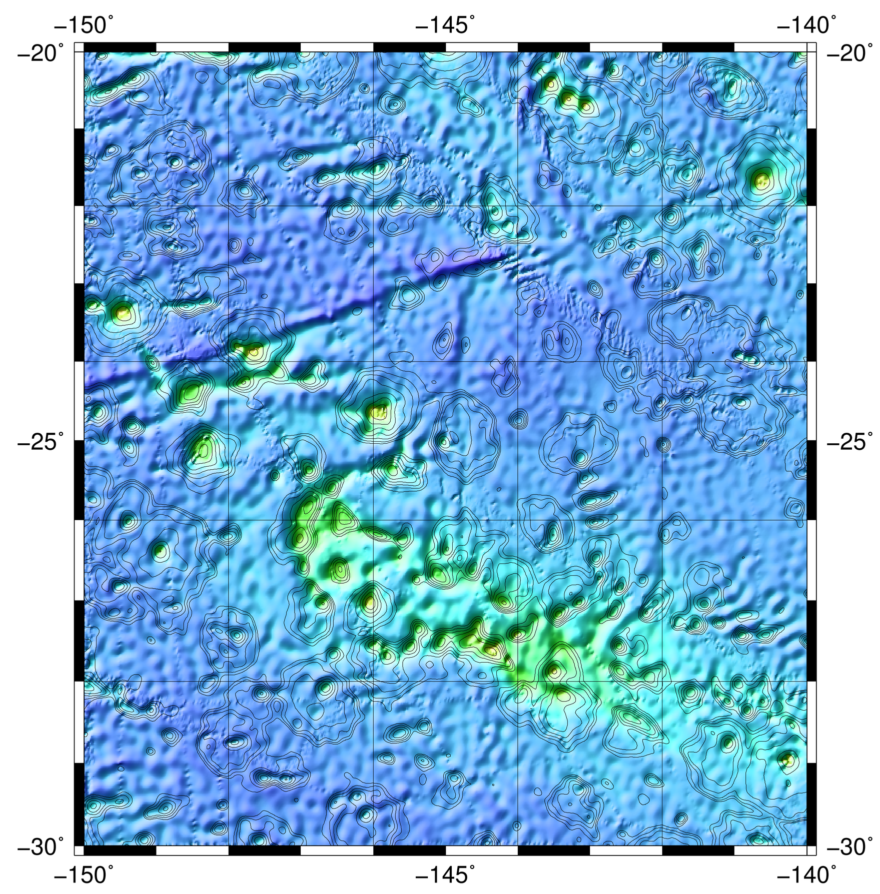

In general, the contours of low reconstruction error correlate well with the occurrence of visible seamounts in the bathymetric data. In principle, we can locate discrete seamounts by identifying local minima in the error surface. However, doing so effectively presents some challenges: see later panel.