Universiteit Utrecht

What Does The Mantle Remember Of Its Convection History?

Mantle convection simulations are a powerful tool for investigating many outstanding problems in geophysics. Models can be used to test many theories about the Earth, including how the planet formed and cooled, how the geodynamo works, and how the mantle and crust affect one another. Models are constrained by observations, but for early Earth history we can only construct theories based on today's observations and the geological record. Here we investigate what information the present state of mantle convection may contain about early Earth history.

Suzanne Atkins, Paul Tackley*, Andrew Valentine, Jeannot Trampert
Utrecht University, Netherlands; *ETH Zürich, Switzerland
s.e.atkins@uu.nl

The Problem

To estimate the initial conditions for Earth, we need a method to identify which conditions influence the final convection state to a detectable extent and a method to invert final convection state for those initial conditions.

Why?

Convection simulations are problems where the same initial conditions with the same boundary conditions always produce the same final result. But what are the initial conditions and boundary conditions needed to produce an Earth-like result?
If we can constrain some parameters reliably, we can begin to investigate large-scale mantle convection and investigate more complex models and smaller-scale variations.

Uncertain & Unknown Initial Conditions + Unknown Boundary Conditions = Earth-like model

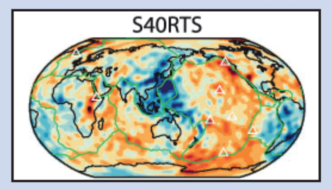

S40RTS: Ritsema et al. 2011

Convection is also a non-linear problem. Small changes to initial or boundary conditions may make a large difference to the final convection state (the "Butterfly Effect"). Large changes in initial conditions may make almost no difference to final state. Parameters may interact so that their effects cancel each other out.

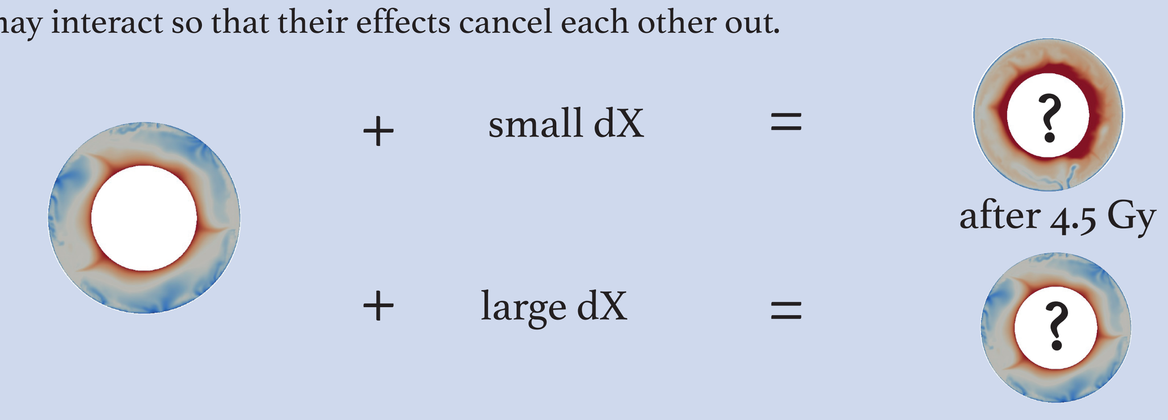

Model Characteristics & Initial Conditions

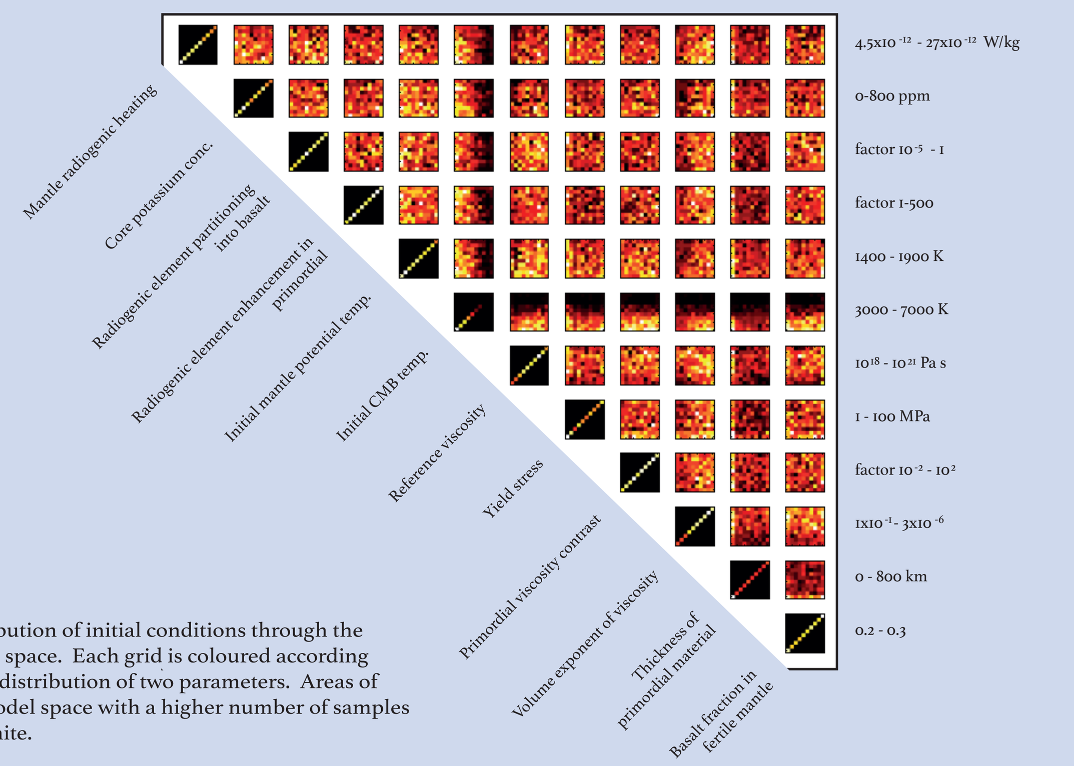

Distribution of initial conditions through the model space. Each grid is coloured according to the distribution of two parameters. Areas of the model space with a higher number of samples are white.

We investigate how 12 different mantle convection conditions (thermal, chemical and compositional), plus mineral physics parameters, affect mantle convection. We vary all of the parameters at once, allowing us to investigate which parameters have a dominant effect and how different parameters interact.

Each condition is drawn independently from a uniform distribution with ranges wider than those expected for the Earth. Major element oxide components for each of basalt, harzburgite and primordial rock types are also varied. Some combinations of conditions, particularly high core temperature + low viscosity, cause the models to crash because convection becomes too vigourous.

Our Solution

By running thousands of convection simulations with widely varying initial conditions, we can train a neural network to identify subsequent patterns that depend on particular initial conditions and to approximate the inverse function between final convection state and initialisation conditions.

The model

We run the mantle convection simulation StagYY 3000 times. Each run varies 12 initialisation parameters independently. The models have a 2D spherical annulus geometry with resolution 64x512 cells and free-slip boundaries, are compressible, and form basaltic crust through melting of fertile mantle material. Composition is tracked with 500,000 tracers. Major element ratios of each rock type (basalt, harzburgite and primordial) also vary in each model, with mineral physics properties calculated usng the Perple_X package (Connolly, 2009).

The temperature and density fields calculated by StagYY can then be used to train our network.

The neural network

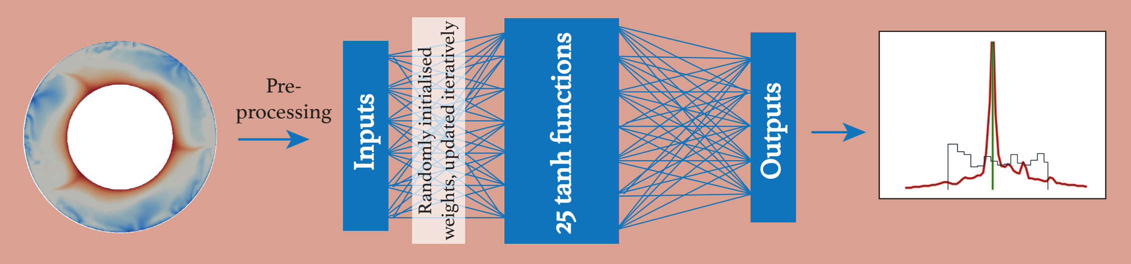

Why neural networks?

Neural networks can approximate complex non-linear functions. They are also capable of interpolating through high-dimension parameter space, meaning high dimension problems, such as this, can be attempted with far fewer data points than would be require for other probabilistic approaches, such as Monte-Carlo methods.

Training Sets

Preprocessing

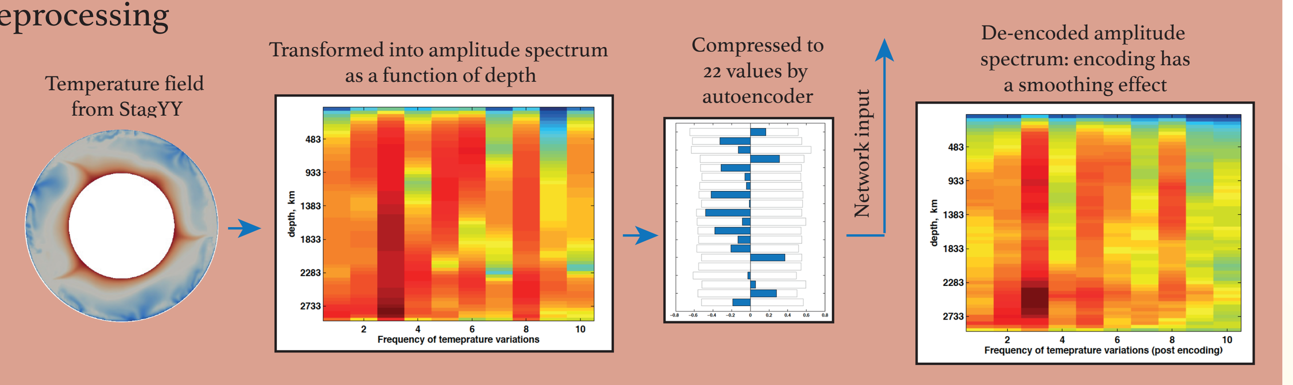

We train our neural networks using three different data sets: temperature, composition (fraction of basaltic material in each cell), and thermal heterogeneity spectrum. Each is calculated by StagYY starting from known initial conditions. For temperature and composition, the field is transformed into the frequency domain. We then train an auto-encoding neural network (Valentine & Trampert, 2012) to compress the amplitude spectrum.

How do neural networks work?

A neural network consists of randomly initialised weights linking several layers of functions. Our network is a mixture density network, with a layer of input functions, a layer of 25 tanh functions and a layer of output Gaussian functions. Each input element contributes to each tanh function and each tanh function contributes to the mean and standard deviation of each Gaussian. The Gaussians are then summed to form a probability density function which is the network estimation of whatever we are trying to find for that input.

To train the network, we have a training set of inputs each with a known target value. The error between the known target and the network estimate is calculated at each iteration. The contribution of each of these weights to the error in the output is calculated, and updated to reduce the error. We use the iRPROP+ algorithm of Igel & Huskens (2000) to update the network. We use a committee of 40 networks, as this removes some of the effects of over-training. Over-training occurs when the network tries to fit a function to the training data points too accurately, at the expense of the generalised performance.

Once it has been trained, the network can be shown new, unknown data for which it can calculate the likelihood of that new data being the result of any initial condition.

Results

We can train a network to find various rheological and composition parameters from the resulting convection patterns.

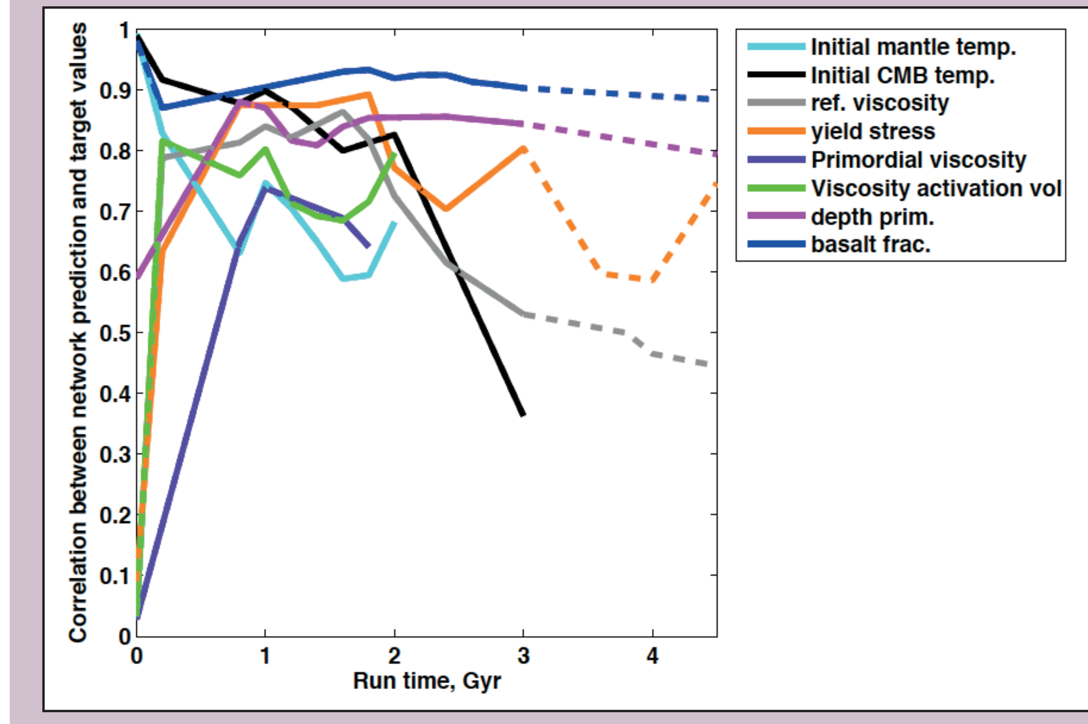

Network ability to map between StagYY ouputs and starting conditions as a function of time

Network Perfomance

We use two different measures of network performance: the correlation between the network predicted values; and the Kullback-Leibler divergence.

Correlation ≈ Accuracy

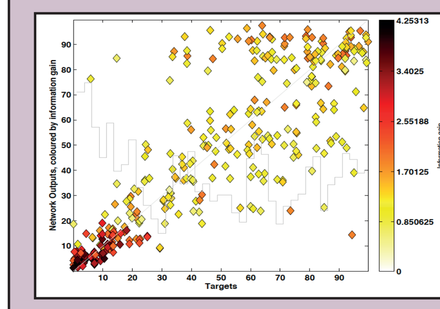

In an ideal case, the maximum likelihood of the network output should perfectly match the target value for every element in the test set. This is not possible because of noise, interference between model parameters, incomplete data space coverage or simply the lack of signal in the data. This scatter plot is for the true yield stress and the network prediction after 4.5 Gy, coloured according to information gain.

Kullback-Leibler Divergence: Information gain ≈ strength of signal

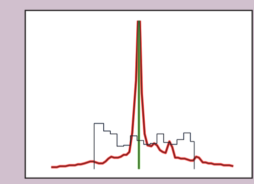

The entropy of a distribution is the degree of predictability of samples from that distribution. A Gaussian with a low standard deviation has low entropy. The relative entropy between two distributions gives a measure of the information gain. The figure to the left has a high relative entropy between the prior (blue) and posterior (red). A K.L divergence of 1.16 bits corresponds to a halving of Gaussian standard deviation.

$$KLdiv(R \,||\, B) = \sum R \log_2 (R / B)$$

Which conditions can we trace in StagYY outputs?

For 4.5 Gy
- Yield stress
- Basalt fraction
- Depth of primordial material

For 2 Gy
- Initial CMB temperature
- Reference viscosity
- Viscosity activation volume

For 1 Gy
- Initial mantle potential temperature

Our networks can find a mapping between various initialisation and intrinsic characteristics and the state of the mantle which results from these initialisation conditions. The fraction of basalt (basalt frac.) in fertile mantle and initial thickness of layers of primordial material (depth prim.) maintain a strong signal in the composition field. The yield stress maintains a strong signal in the temperature field for 4.5 Ga, because high yield stress leads to stagnant-lid type planets, whilst low yield stresses produce more continents allowing the mantle to cool. Signals from initial temperature (initial CMB temperature and mantle potential temperature) are lost after around 1 Gy and 2 Gy respectively as the heat fluxes stabilise. A lack of signal may also be caused by too few training sets or low resolution and can be improved.

Models with low viscosity require more time-steps to reach 4.5 Gy. We currently only have enough completed models runs to train networks up to 3 Gy. However, we can provisionally extend our investigation until we have enough runs. Using the composition field, we can train a network with output profiles from 2 Gy of run time, then show it patterns produced after 4.5 Gy of run time. This works well for composition parameters, such as basalt fraction. We can also remove scaling factors, such as 1D temperature profile, and train networks to recognise the pattern types (eg. 3rd vs. 4th order symmetery in the convection patterns) associated with each initial condition. This is less successful, but does work for yield stress. This method also allows us to apply our networks to tomographic models by assuming that variations observed are primarily due to lateral temperature variations.

What happens with a tomography model?

Assuming that shear speed variations represent temperature variations to the first order, we take the tomographic model S40RTS (Ritsema et al. 2011) and calculate the amplitude spectrum as a function of depth. This is then preprossesed in the same way as our synthetic StagYY temperature fields. All scaling is removed, leaving just the degree 1-10 amplitude spectrum. This means that we can use a network trained with temperature to make an estimate for Earth values.

Amplitude spectrum a a function of depth for S40RTS

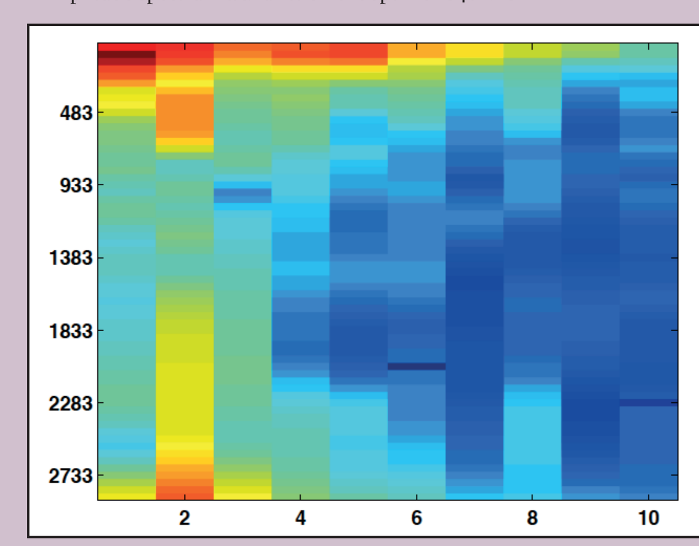

Network estimate for **initial mantle temperature** from S40RTS.

These have low information gain, but five different networks give consistent results between 1400-1500 K.

KL divergence ~0.35

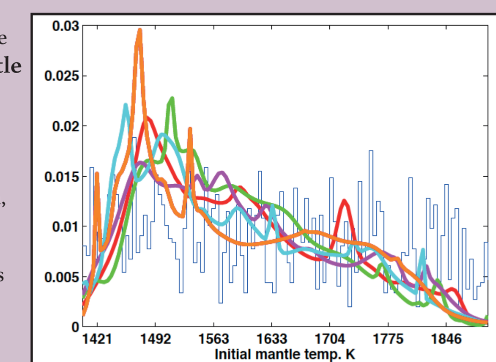

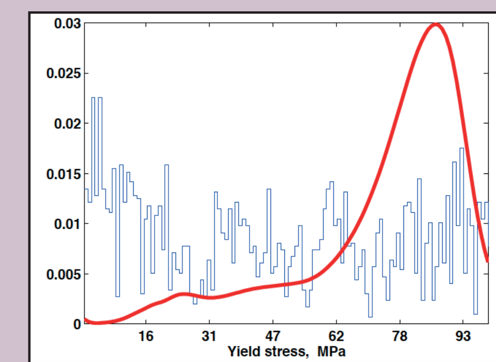

Network estimate for **yield stress** from S40RTS.

This is the best constrained result.

K-L divergence 1.29

Network estimate for **reference viscosity.**

These have low information gain. Four different networks give similar results around 10^{18}.

K-L divergence ~0.3

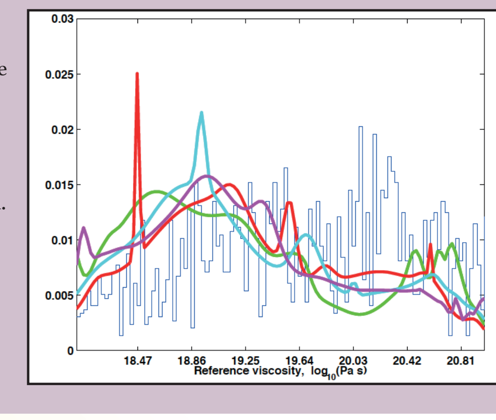

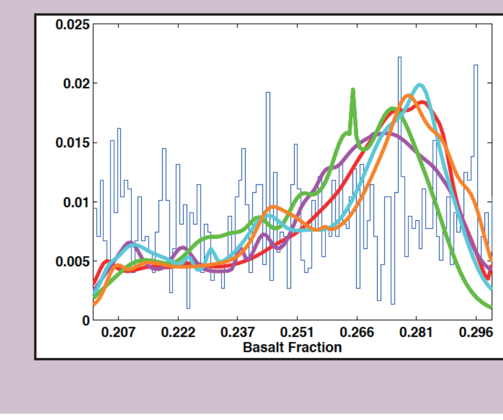

Network estimate for **basalt fraction.**

These have low information gain, but five different networks give consistent results around 0.28.

K-L divergence ~0.35