Extracting information from seismic waveforms using a neural network approach

Ralph de Wit[1], *Paul Käufl*[1] *& Jeannot Trampert*[1]

[1] Dep. of Earth Sciences, Utrecht University, The Netherlands

Contact: rdewit@geo.uu.nl, kaeufl@geo.uu.nl

Universiteit Utrecht

NWO Netherlands Organisation for Scientific Research

Introduction

We aim to invert seismic waveforms by using artificial neural networks. Neural networks can be viewed as non-linear filters and are very common in speech, handwriting and pattern recognition. We use neural networks to extract information on the seismic source and Earth structure which is contained in a seismogram.

We show the general concepts of two applications here. The first application can be viewed as a pattern recognition problem. The goal is to perform full seismic waveform inversion and invert for Earth structure and seismic source. A second application treats the seismogram as a discrete time series by using recurrent neural networks, which are the non-linear equivalent of recursive filters. Here the aim is to use the first arriving seismic waves to predict later arriving seismic phases.

To be able to train a neural network, a so-called training set is needed. We construct such a training set by drawing many random Earth models (and seismic sources) from a prior model probability density function (pdf) and solving the forward problem for each of these models, thus generating synthetic seismograms. For global 2D and 3D Earth models, we aim to use spectral-element methods, such as AXISEM (Nissen-Meyer et al., 2007) and in a later stage SPECFEM3D (Tromp et al., 2008).

Seismic waveform inversion through pattern recognition

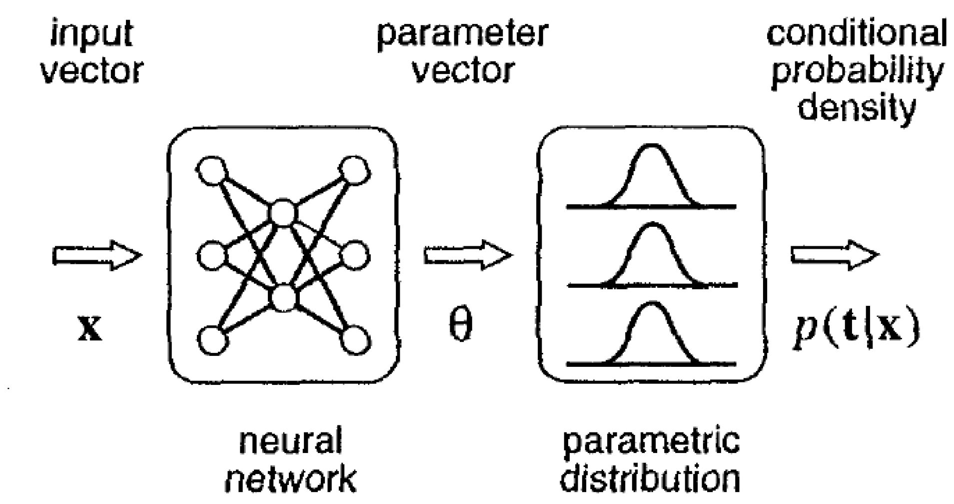

Figure 1: A Mixture Density Network, as depicted in Bishop (1995).

We use a Mixture Density Network (MDN, Figure 1) to obtain marginal posterior pdfs of our model parameters, thereby acquiring fully probabilistic information on the model. An MDN can approximate an arbitrary conditional pdf as a linear combination of Gaussian kernels (Bishop, 1995):

$$p(\mathbf{t}|\mathbf{x},\mathbf{w}) = \sum_{j=1}^{M} \alpha_j(\mathbf{x};\mathbf{w})\phi_j(\mathbf{t}|\mathbf{x};\mathbf{w}) \tag{1}$$

where $\mathbf{w}$ are the adjustable parameters in the neural network used, M is the number of Gaussian kernels, α_j are the mixing coefficients, which can be interpreted as the relative importance of the j^{th} kernel, and ϕ_j are the spherical Gaussian kernels of the form

$$\phi_j(\mathbf{t}|\mathbf{x};\mathbf{w}) = \frac{1}{(2\pi\sigma_j^2(\mathbf{x};\mathbf{w}))^{c/2}} \exp\left\{-\frac{(\mathbf{t}-\boldsymbol{\mu}_j(\mathbf{x};\mathbf{w}))^2}{2\sigma_j^2(\mathbf{x};\mathbf{w})}\right\} \tag{2}$$

where c is the dimensionality of the output vector $\mathbf{t}$. The parametric distribution, which is described by the means $\boldsymbol{\mu}_j$, the variances σ_j^2 and α_j, is the output of a conventional neural network. We use a feed-forward Multi-Layer Perceptron (MLP) with hyperbolic tangents as non-linear activation functions. Network training corresponds to the minimisation of the negative logarithm of Eq. (1) for a training data set. The minimisation is done using either on-line or batch learning methods, such as gradient descent, quasi-Newton or conjugate gradient methods.

For the application we have in mind here, seismograms would serve as the input vector $\mathbf{x}$ and Earth structure or seismic source as the target vector $\mathbf{t}$. The network is trained by using a large synthetic data set that we constructed using spectral-element methods. Once the network has been trained, it can be presented with new unseen input data, in this case real seismograms. As output we then obtain the posterior pdf $p(\mathbf{t}|\mathbf{x})$ which represents our final state of knowledge on the model parameters. An example of such an application is shown for a toy problem in Figure 2.

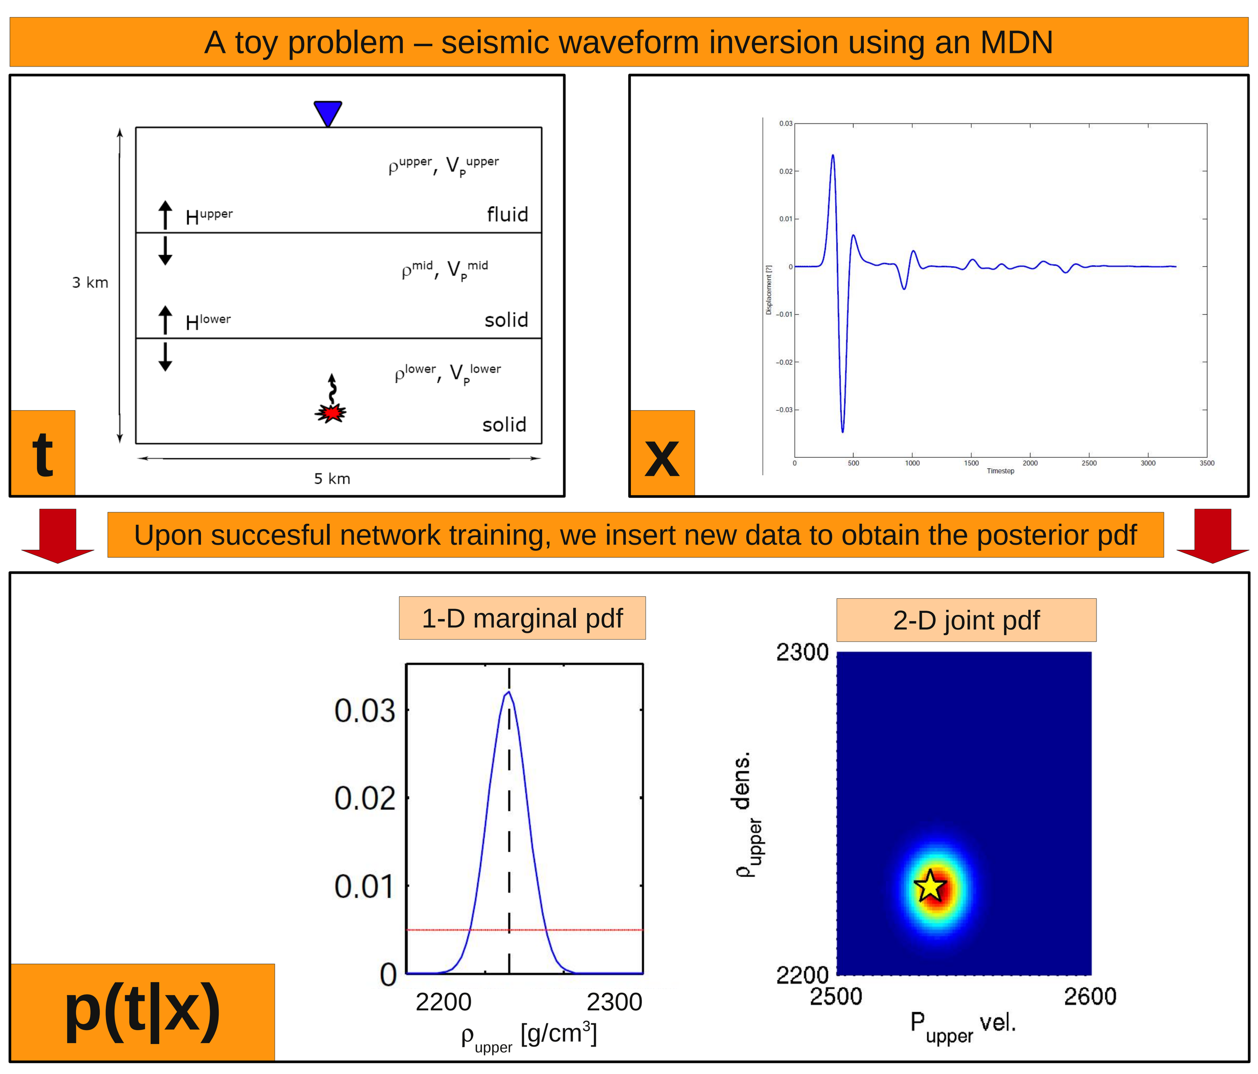

Figure 2: A simple example of seismic waveform inversion by using a Mixture Density Network. (Top left) eight model parameters (ρ,V_P in three layers and two interface depths) are varied in the training set. (Top right) An example of a seismic waveform in the training data set. (Bottom) One MDN is trained on the 1-D marginal pdf for ρ_{upper}, another on the 2-D marginal pdf for ρ_{upper} and $V_{P,upper}$.

Recurrent architectures for seismic waveform prediction

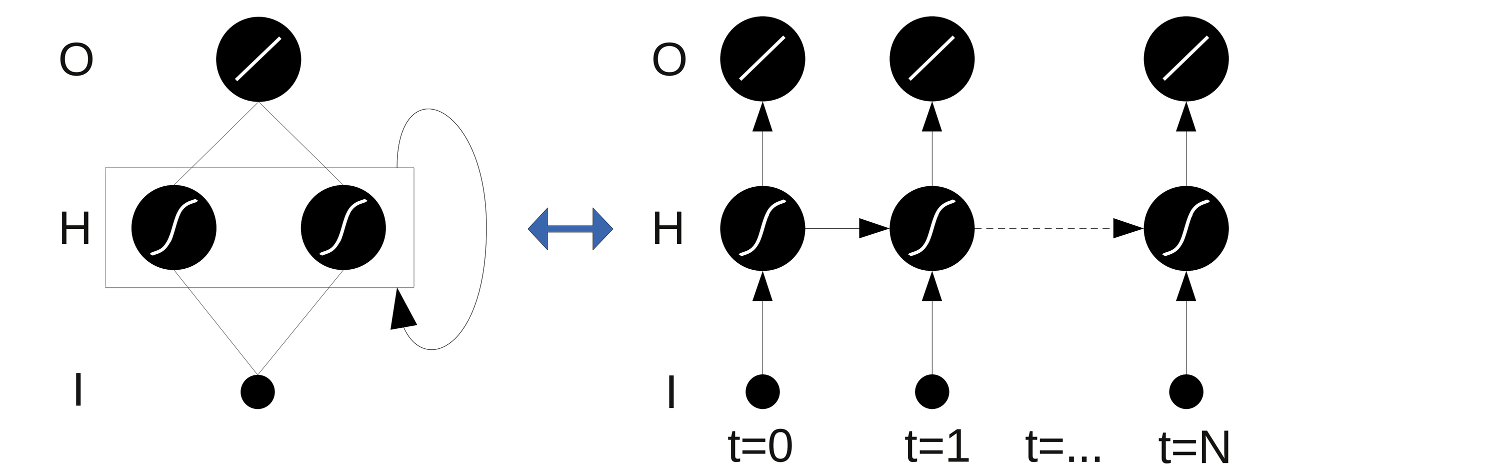

Figure 3: Left panel: Example of a two-layer recurrent artificial neural network with one input node (I), two sigmoid units in the hidden layer (H) and a linear output node (O). Right panel: A recurrent neural network unfolded in time.

It has previously been shown (Hammer, 2000) that recurrent neural networks (RNNs) have very general function approximation properties. We thus use a recurrent neural network architecture similar to the one shown in Figure 3 to extract information from seismic waveforms. We therefore treat the recorded signal as a discrete time-series which is fed into the network sequentially, that is a new input is given to the network at each time-step. A recurrent network is able to 'memorize' previous inputs by feeding hidden layer output signals back into the neurons at the next time-step. Thus at each time-step the output of a recurrent network $\mathbf{y}(t)$ is a function of the complete history of the input signal $\mathbf{x}(t)$ so far:

$$y_o^t = f\left(\mathbf{x}^t, \mathbf{x}^{t-1}, \ldots, \mathbf{x}0\right) = \sum_{h=1}^{H} w_{ho}\, g\left(a_h^t\right) \tag{3}$$

$$a_h^t = \sum_{i=1}^{I} w_{ih} x_i^t + \sum_{h'=1}^{H} w_{h'h}\, g\left(a_{h'}^{t-1}\right), \tag{4}$$

where w_{ij} denotes the connection weight from the ith to the jth unit and $g(\cdot)$ is a non-linear, sigmoid activation function. An optimal set of network parameters is found during the training stage using an on-line back-propagation algorithm.

Although in theory the whole signal history is available to the network at each time-step, given the number of hidden units is large enough, in practice the amount of past values that can be used for training is limited, due to a problem known as vanishing error gradient. This can be overcome by using an extension to the standard RNN architecture called Long-Short-Term-Memory (Hochreiter and Schmidhuber, 1997).

The recurrent approach may have advantages over a classical feed-forward approach: (1) The reduced input dimensionality leads to a smaller number of network parameters, leading to faster convergence during the training stage. (2) It is not necessary to determine the sequence length in advance, which makes it feasible to use the network for on-line tasks, such as real-time earthquake localization and source-inversion as well as the prediction of later arriving phases.

In order to assess the capabilities of a RNN to perform the latter task, we set up a very simple toy problem. We simulate wave propagation in a 2D homogeneous medium and vary the location of a moment-tensor point source. Synthetic waveforms are recorded at a single station and fed into a LSTM recurrent network, which is trained on the arrival time of the S phase. The trained network is then tested by processing previously unseen P wave sequences which give rise to S wave arrival time predictions. Exemplary results from a trained network are given in Figure 4.

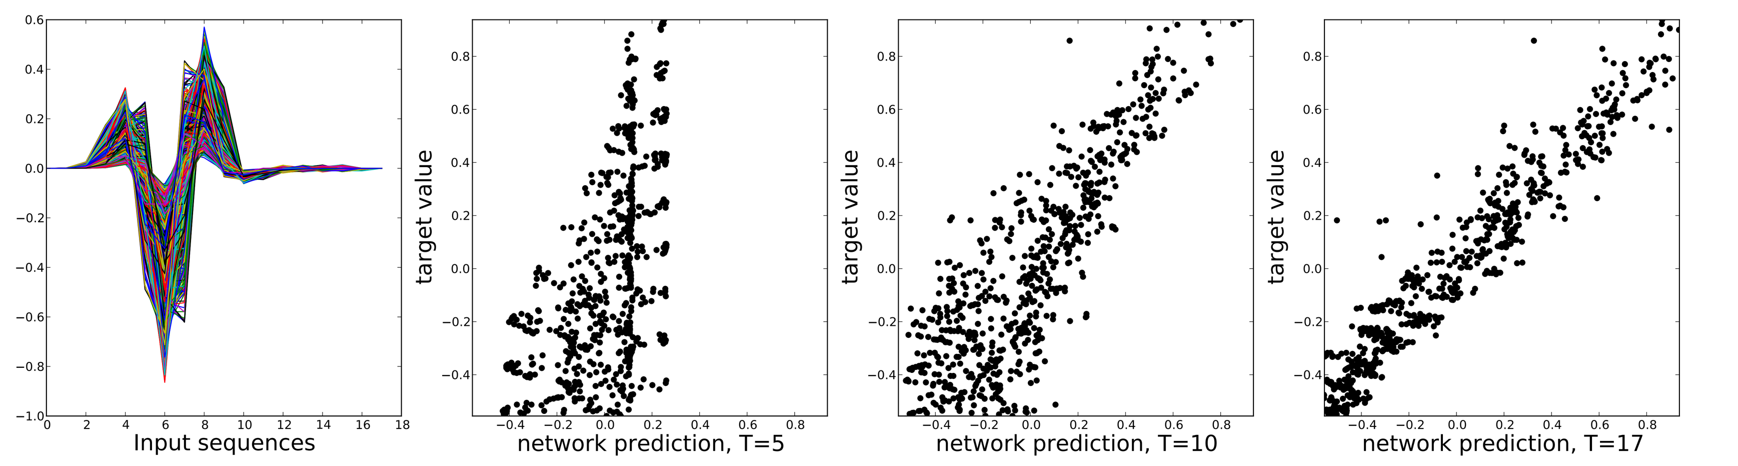

Figure 4: Leftmost panel: Synthetic P waves are given as an input to a LSTM recurrent network sequentially while training it on S wave arrival times as targets. Shown on the right are network predictions from a trained network on an independent test set after 5, 10 and 17 timesteps have been presented to the network. Each dot corresponds to one of the traces shown on the left. As can be seen, the more of each sequence is available to the network, the more accurate the prediction gets.

References

Bishop, C. M. (1995). *Neural networks for pattern recognition.* New York: Oxford University Press.

Hammer, B. (2000). On the Approximation Capability of Recurrent Neural Networks. In *Neurocomputing*, Volume 31, pp. 107–123.

Hochreiter, S. and J. Schmidhuber (1997). Long Short-Term Memory. *Neural Computation 9*, 1735–1780.

Nissen-Meyer, T., A. Fournier, and F. A. Dahlen (2007). A two-dimensional spectral-element method for spherical-earth seismograms-I. Moment-tensor source. *Geophys. J. Int. 168*, 1067–1092.

Tromp, J., D. Komatitsch, and Q. Liu (2008). Spectral-Element and Adjoint Methods in Seismology. *Communications in Computational Physics 3*, 1–32.