

Building low frequency model with Deep Neural Network for Constrained Sparse Spike Inversion without structural model

Introduction

The low frequency model (LFM) provides the absolute elastic background that Constrained Sparse Spike Inversion (CSSI) cannot recover from band-limited seismic. Conventional LFM building interpolates well curves onto an interpreted structural framework, an approach whose reliability falls away where the structure itself is uncertain or hard to construct. This paper presents a workflow that builds the LFM directly with a Deep Feed-forward Neural Network (DFNN), trained on well curves with attributes derived from seismic partial stacks and the trend from seismic interval velocity. The trained network predicts elastic properties, i.e., P-Impedance, V_p/V_s and Density volumes, which after median edge-preserving and high-cut filtering serve as the LFM for input to CSSI. Because the prediction is guided by seismically derived attributes rather than by interpolation across an interpreted framework, no structural model is required. The workflow shortens model turnaround from days or weeks to hours and is well suited to complex geology and to the exploration phase, where well control is sparse and interpretation is not yet refined.

Several methods reduce this dependence. Hak and Mesdag (2015) introduced lateral variability by incorporating seismic velocity into the LFM through rock physics relationships. Pendrel and Schouten (2018) estimated impedance from facies using Bayesian analysis, bypassing explicit well curve interpolation. Iterative LFM updating with early inversion passes can further suppress side lobes around high-contrast features (Jarvis, 2006; Mesdag et al., 2010). Building on this trend toward more data-driven, less interpretation-dependent models, DFNN is used to generate the LFM. Deep networks have already proven effective at predicting reservoir properties such as porosity (Malik, 2019), water saturation and net pay (Colwell and Kjøsnes, 2018). Here, following the approach of Hassan et al. (2021), the network predicts the elastic background itself, removing the structural-model requirement. The remainder of this paper sets out that workflow step by step.

Method and Workflow

The DFNN is a supervised, multi-layer feed-forward network (Figure 1). Information flows in one direction, from an input layer through two or more hidden layers to an output layer; with two or more hidden layers the network is considered “deep”, and the additional non-linear layers allow it to represent higher-order relationships between inputs and outputs that a shallow network cannot. Each node forms a weighted sum of the outputs of the previous layer, adds a bias term, and passes the result through a non-linear activation function. Stacking non-linear layers in this way lets the network approximate the complex, non-unique mapping between seismic response and elastic property without it having to be specified analytically.

Training is the process of finding the weights and biases that make the network reproduce the known elastic properties at the wells. It begins from an initial set of random weights and iteratively minimises an objective function measuring the misfit between the actual and predicted training values, solved as a non-linear inverse problem (Bekti, 2020). Optimisation proceeds by gradient descent and terminates on a minimum-error criterion, using either the conjugate gradient or steepest descent algorithm. In practice a small number of layers with a moderate node count is sufficient for elastic-property prediction.




Figure 1 Architecture of the Deep Feed-forward Neural Network used to build the LFM. Seismic-derived attributes and well measurements enter the input layer and pass through non-linear hidden layers to predict the elastic properties (P-Impedance, Vp/Vs and Density) at the output layer; bias nodes feed each downstream layer.

The workflow that wraps this network around the LFM-building problem (Hassan et al., 2021) proceeds as follows:

1. Data conditioning and alignment. As in any inversion project, the pre-stack gathers or partial angle stacks are conditioned and aligned, and accurate well ties are established. Trim statics on the gathers, careful volume alignment between the partial stacks, and reliable well ties are prerequisites for a dependable DFNN result; any misalignment left here is learned by the network and propagates into the predicted volumes.
2. Input and target definition. The network inputs are the elastic-property well curves together with volume attributes computed from the seismic partial angle stacks and the seismic interval velocity volume. The wells and interval velocity contribute the low-frequency trend, while the partial stacks supply the laterally varying amplitude information that allows the trend to depart from a simple interpolation. The targets are the elastic properties to be predicted, typically P-Impedance, Vp/Vs and Density. The analysis window is set to the range of the well curves so that the network is trained and applied over a consistent interval.
3. Network training. A supervised learning pass optimises the weights against the objective function. Training combines the seismic-derived attributes so that the network captures the relationship between seismic response and elastic property at the well locations, then propagates that relationship laterally across the survey using the attributes as a guide.
4. Validation and hyper-parameter selection. To guard against over-training, blind-well validation is combined with percentage-based validation, in which the training points are randomly split into training and validation subsets (commonly 70% for training and 30% for validation) and the error of each is tracked independently. This analysis guides the choice of hyper-parameters such as the number of iterations, the number of attributes, and L1/L2 regularisation. The optimal iteration count is read off where the validation error stabilises; continuing past that point lowers the training error but raises the validation error, the signature of over-fitting.
5. Noise suppression. Once the elastic-property volumes are produced, a median edge-preserving lateral filter is applied. This suppresses trace-to-trace random noise while preserving values across steeply dipping zones and fault edges, producing a smoother volume and avoiding the jittering that a subsequent high-cut filter would otherwise amplify.
6. Band-limiting to the LFM. A high-cut filter then reduces the volumes to the low-frequency band required by the model. The cut-off is selected from a seismic amplitude-spectrum QC of the partial stacks, or by letting the inversion test the optimal merge frequency, so that the LFM and the band-limited seismic combine without a gap or an overlap in the spectrum.

7. Integration with CSSI. The resulting P-Impedance, Vp/Vs and Density LFM volumes are passed directly into the Constrained Sparse Spike Inversion. Inversion parameterisation gives control over the seismic/synthetic correlation, and the inverted elastic volumes can be compared against the direct DFNN prediction as a further, independent QC. Because the DFNN prediction follows the seismically imaged structure, the resulting LFM honours fault-controlled and stratigraphic discontinuities without any explicit structural model.

The defining feature of the workflow is that steps 2 to 7 contain no structural-modelling stage. Property variation away from the wells is determined by the seismic attributes the network has learned to map, not by interpolation across interpreted horizons and faults. The same sequence applies whether the geology is simple or complex, which removes the single most time-consuming and uncertainty-laden stage of conventional LFM building.

Discussion

The workflow carries several practical consequences. First, turnaround drops sharply. The conventional sequence of interpretation, structural modelling and LFM building can take days to weeks (Figure 2), because the structural model must be built and quality-controlled before any LFM can be interpolated onto it. With the DFNN, each elastic-property model is produced in a few hours once the data are conditioned, so several models can be generated and compared within the time a single conventional model would take.

Second, the LFM becomes consistent with the seismic image. Property boundaries fall on the faults, and stratigraphic edges actually present in the data rather than on a separately interpreted framework, which is particularly valuable where faulting is dense or geometries are complex and where an interpreter would otherwise have to make many subjective choices.

Third, the approach fits the exploration phase, where well control is limited and seismic interpretation is still maturing. A usable LFM can be produced before a detailed structural interpretation exists, allowing inversion and quantitative interpretation to begin earlier in the life of a prospect.




Figure 2 Deterministic reservoir characterization workflow. Conventional LFM building covering seismic interpretation, structural modelling, and the LFM itself may take days to weeks; building the LFM with machine learning shortens this to hours.

These benefits depend on disciplined QC. Blind-well validation, in which one or more wells are withheld entirely from training and used only to test the prediction, is the strongest safeguard against a network that merely memorises its training data; percentage-based validation and the error curves support the choice of hyper-parameters but do not by themselves prove that the prediction generalises. The principal trade-off is that the network's lateral predictions are only as good as the conditioning, alignment and well ties that feed it, so the front end of the workflow (step 1) remains decisive: errors introduced before training cannot be recovered afterwards. The method also depends on attributes that genuinely carry low-frequency information; where the partial stacks and velocities are weak or noisy, the prediction reverts toward the well trend and the advantage over conventional interpolation narrows.

Conclusion

This paper has described a workflow that builds the LFM for Constrained Sparse Spike Inversion directly with a Deep Feed-forward Neural Network, removing the need for a structural model. The network trains on well curves and incorporates attributes from seismic partial stacks together with the trend from seismic velocity; median edge-preserving and high-cut filtering then reduce the predicted elastic volumes to the LFM band for input to CSSI. Because property variation is driven by seismic attributes rather than by interpolation across an interpreted framework, the LFM is consistent with the seismic image and honours fault-controlled and stratigraphic discontinuities without explicit structural modelling. The workflow has a shorter turnaround than conventional LFM building, is robust in complex geological settings, and is well suited to the exploration phase, provided that data conditioning, well ties and blind-well validation are carried out with care.

References

- Bekti, R. [2020] Deep machine learning application for supervised facies classification. EAGE/AAPG Digital Subsurface for Asia Pacific Conference, Sep 2020, Volume 2020, 1-5.
- Colwell, T. and Kjøsnes, Ø. [2018] Comparative study of deep feed forward neural network applications for seismic reservoir characterization. First EAGE/PESGB Workshop Machine Learning, Volume 2018, 1-4.
- Hak, B. and Mesdag, P. [2015] Incorporating seismic velocity data in AVO/AVA low frequency models by honoring local geology. 77th EAGE Conference and Exhibition, Volume 2015, cp-451-00380.
- Hassan, T.M.S.T., Lee, C.S., Bekti, R.P.A. and Ting, J. [2021] Building low frequency model with Deep Learning for seismic inversion in complex geology without structural model. 82nd EAGE Annual Conference & Exhibition, Volume 2021, 1-5.
- Jarvis, K. [2006] Integrating well and seismic data for reservoir characterization: Risks and rewards. Australian Earth Sciences Convention 2006, Extended Abstracts, MV6.
- Malik, J. [2019] Application of machine learning for porosity prediction using 3D seismic data. Conference Proceeding APGCE 2019, Volume 2019, 1-5.
- Mesdag, P., Marquez, D., De Groot, L. and Aubin, V. [2010] Updating low frequency model. 72nd EAGE Conference & Exhibition, Extended Abstracts, F027.
- Pendrel, J. and Schouten, H. [2018] A facies-first workflow for the estimation of elastic-reservoir properties. SEG International Exposition and 89th Annual Meeting 2019, SEG Technical Program Expanded Abstracts: 386-390.