CS230

CNN-based seismic facies classification from 3D seismic data

Wei Chu
Department of Geophysics
Stanford University
weichu@stanford.edu

Iris Yang
Department of Geophysics
Stanford University
iyyang@stanford.edu

Abstract

We present a new three-dimensional (3D) convolutional neural network (CNN) framework to classify seismic facies from 3D seismic data. The new CNN architecture implements a variant of the LeNet-5 design, and improves training accuracies significantly in the early training stage. A sparse sampling scheme to preprocess input data is introduced to improve the computational speed and maintain a desirable receptive field. The new CNN design in combined with the sparse sampling scheme reduces the training time by more than 50%, and achieves a test accuracy of 0.9977. The current 3D CNN-based approach proves to be very promising in making geologically reasonable and consistent predictions for seismic facies.

1 Introduction

The increasing availability of three-dimensional (3D) seismic data demands a workflow for automatic seismic facies classification. Traditionally, the seismic facies interpretation and classification are made by geoscientists. In addition to the tremendous amount of time involved in classifying 3D seismic volumes, the quality of the results is often less controlled due to the variation in geoscientists' experiences. A convolutional neural network (CNN) based procedure that automatically learns and extracts useful features from 3D seismic volumes, and consequently makes seismic facies classification can significantly reduce the human labor and uncertainties, thereby accelerating decision-making time.

2 Related work

A few recent works have attempted to automate the seismic facies classification using CNN. Waldeland and Solberg (2017) presented a binary classification algorithm for salt dome detection by classifying each pixel using its surrounding 3D seismic information. Ildstad and Bormann (2017) extended Waldeland and Solberg (2017)'s strategy by making multi-class seismic facies predictions using similar input data formats and CNN architectures. Our project bases on the "MalenoV" repository and labeled data created by Waldeland and Solberg (2017), and aims to further improve the prediction accuracy and computational speed.

3 Dataset

The 3D seismic dataset for the current work is obtained from a publicly available seismic survey of the Netherlands offshore F3 block (dGB Earth Sciences B.V., 1987). The 3D seismic cube (data) is a

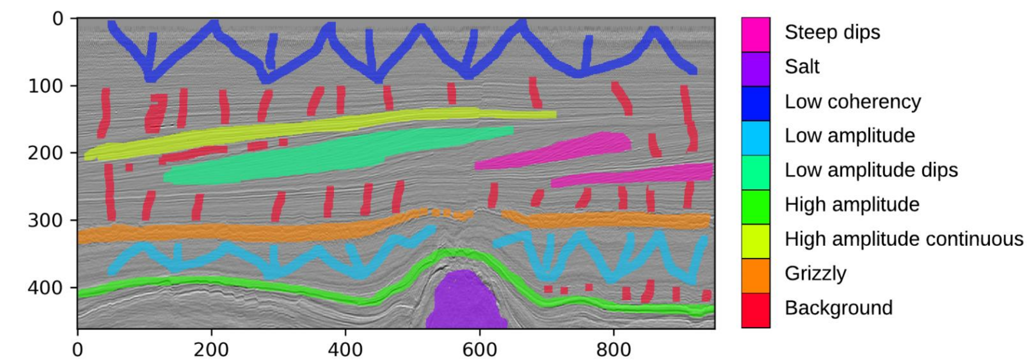

Figure 1: Labeled pixels on the inline 339 section.

3D volume of size $651 \times 951 \times 452$ in the inline, xline, and time dimensions, respectively. As shown in Figure 1, Ildstad and Bormann (2017) labeled 9 classes (8 seismic facies + 1 background class) on pixels from the inline 339 section, which is a two-dimensional (2D) slice. The inline 339 section is 951×452, with a total of 140,111 labeled data points. The training, validation, and test sets are prepared in a way such that each sample is a 3D sub-cube of size $65 \times 65 \times 65$ with the central voxel (pixel) belonging to one of the labeled pixels in the inline 339 section.

4 Methods

4.1 New CNN design: the modified LeNet-5

We made major improvement on the CNN architecture based on Ildstad and Bormann (2017)'s CNN design. The original CNN architecture was fairly crude: four of the five convolutional layers used the same filter design, which is a 50-channel $3 \times 3 \times 3$ filter; no max-pooling layers were included. The new CNN architecture we redesigned, as shown in Figure 2, is a variant of the LeNet-5 named "the modified LeNet-5". The new design incorporated the overarching concept that the number of channels of filters should increase while the filter size should decrease as the layers go deeper. Two max-pooling layers were added in order to enhance local features; the dimension of the filter tensors in the original LeNet-5 was increased by 1 so that 3D convolution can be performed on the input sub-cubes. The new CNN design demonstrates significant improvement in the loss reduction and training accuracy gains in the early training stage; the results are analyzed in details in the "Results and Discussion" section.

4.2 Input data preprocessing: sparse sampling

When generating sub-cubes for the input data, we implemented a feature that sparsely samples from the raw seismic cube at a specified step size to preprocess the input data. Initially, the $65 \times 65 \times 65$ sub-cubes were generated by continuously sampling around the labeled central voxels. The sparse sampling scheme improves the computational speed by reducing the sub-cube size. Figure 3 shows the comparison between the continuous sampling scheme, and the sparse sampling scheme at a step size of 2 in each dimension, resulting in a sub-cube of size $33 \times 33 \times 33$. The new sub-cube essentially keeps the same receptive field as the old $65 \times 65 \times 65$ sub-cube, but the data size is reduced by a factor of 8. When the input data size is not an issue, the sparse sampling scheme also enables storing more spatial information, potentially improving the prediction accuracy for large-scale features.

5 Results and Discussion

We tested four models using different CNN models and parameters. In all the cases, 40,000 training, 10,000 validation and 10,000 test samples were randomly drawn without replacement from 141,111 total data points. Because only a few epochs were trained, a fixed learning rate of 0.001 was used

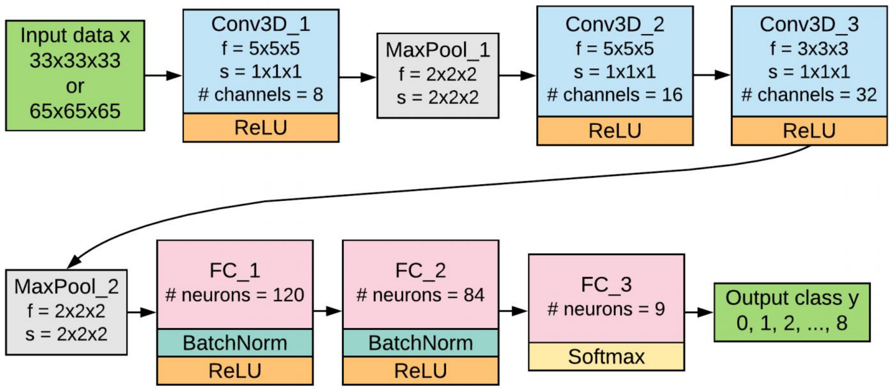

Figure 2: The modified LeNet-5 3D CNN architecture.

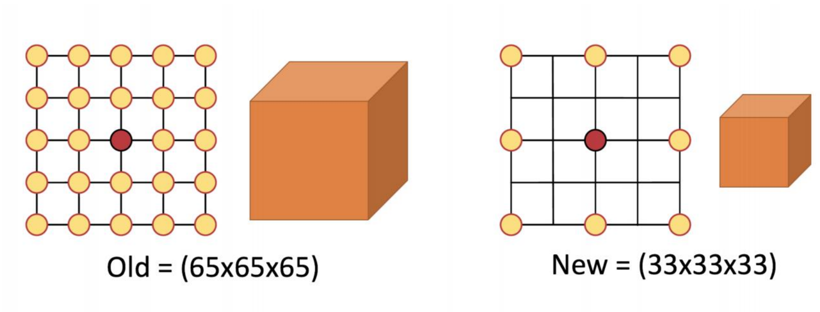

Figure 3: Illustrations of sampling schemes. Left: the original continuous sampling scheme and the resulting sub-cube of size $65 \times 65 \times 65$; right: the sparse sampling scheme around the labeled central voxel in red at a step size of 2 along each dimension and the resulting sub-cube of size $33 \times 33 \times 33$.

except for the base case, where it was initially set to an adaptive learning rate. A cross-entropy loss was defined, and an Adam optimizer was used to perform the stochastic gradient descent. The mini-batch size was set to 32 for size $65 \times 65 \times 65$ sub-cubes, and 128 for size $33 \times 33 \times 33$ sub-cubes. Two epochs were trained in each case. All the training and predicting processes were run on the Stanford CEES GPU cluster with 8 Nvidia Tesla K80 GPUs.

The losses and accuracies are plotted as a function of the mini-batch number in Figure 4. In the legend, the "base" model refers to the old CNN architecture and parameters used in Ildstad and Bormann (2017); the "new" models refer to the modified LeNet-5 architecture shown in Figure 2; the "65" and "33" refer to the lengths of each dimension in the sub-cubes; the "1" and "2" refer to the step sizes used in sampling from the raw seismic volume. As shown in Figure 4 (a) and (b), the new CNN design demonstrates a significant improvement in the performance in the early training stage: the loss decreases more quickly, and the training accuracy increases more rapidly within the first few mini-batches. The use of a new CNN architecture is even more advantageous when the computational power is limited: the training accuracy can achieve more than 0.9 within 100 mini-batches, whereas it needs over 1,000 mini-batches in the base case.

Table 1 compares the performance metrics of the above four models. The "Base" model uses the old CNN design and an adaptive learning rate, and gives a test accuracy of 0.9855 at the end of the second epoch. Although the accuracy is fairly high after two epochs, the early training stage performance, as shown in Figure 4, is not satisfactory due to its CNN design and the initial learning rate setup. The "New 1" model, which uses the new CNN design, shows significant improvement in the performance in the early training stage in Figure 4, but suffers from a much slower training time due to a bigger CNN architecture. The "New 2" and "New 3" models kept the new CNN design, but used smaller input sub-cubes. The reduction of the input data size effectively reduces the the training

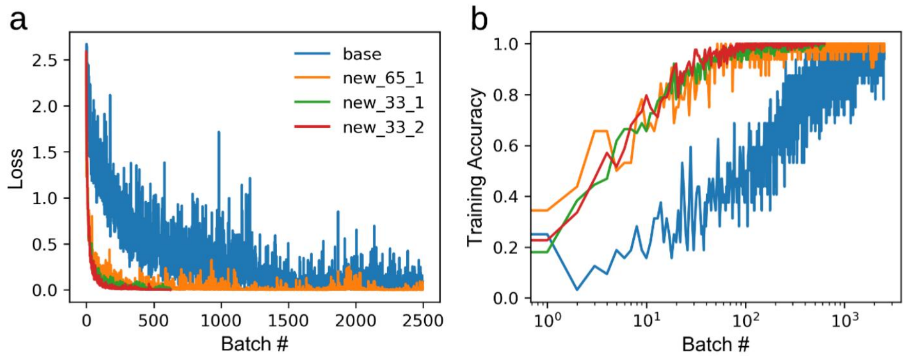

Figure 4: Losses (a) and training accuracies (b) plotted against the mini-batch number for four CNN models. The mini-batch number is plotted in the normal scale in (a), and in the logarithmic scale in (b).

time by more than 50%, and achieves higher computational efficiency. Finally, comparing the original continuous sampling scheme in the "New 2" model with the sparse sampling scheme at a step size of 2 in the "New 3" model, we observe that the sparse sampling scheme results in higher validation and test accuracies because it allows for a deeper receptive field. In summary, the new CNN design in combined with the sparse sampling scheme is the best model for seismic facies classification in the current study.

As a verification, we made seismic facies predictions on the entire inline 339 section using the trained "New-3" model. Figure 5 shows the comparison between the labeled data and predicted results. We observe very good match between the labels and predictions, verifying a test accuracy of more than 0.99. We further made predictions on two xline 2D sections where labels are not available. Figure 6 shows the classified results on xline 450 and 610 sections. The top plots show the raw seismic sections, and the bottom plots show the classified results. We observe good agreement in the seismic facies sequence between the xline 450 and 610 sections; the "salt" facies, which are usually present inside antiforms in nature, were labeled within expected regions in both sections. We conclude that the CNN model is able to generate geologically reasonable and consistent predictions.

6 Conclusion

In conclusion, we found that the 3D CNN framework is very promising in making geologically reasonable and consistent predictions for seismic facies from 3D seismic data. The new CNN design significantly improves training accuracies in the early training stage. The "New-3" model, which combines the new CNN design and the sparse sampling scheme, reduces the training time by more than 50%, and achieves a test accuracy of 0.9977, making itself the best CNN model in the current study.

7 Contribution

We modified around 200-300 out of 1554 lines of code of the "MalenoV" repository, mostly in changing the CNN design, implementing the sparse sampling scheme, adding callbacks for monitoring training progress, and modifying the visualization functions. We also wrote around 100 lines of code in a separate file to process and visualize data. Wei was primarily in charge of designing and implementing the new CNN architecture and the sparse sampling scheme; he also helped with testing models and visualization. Iris was primarily in charge of testing models, setting up and running GPUs, processing and visualizing data, and write-ups.

Table 1: Performance metrics of four CNN models. New 3 was considered the best model because of its high accuracies and computational efficiency.

Model	Sub-cube size	Step size	Mini-batch size	# of para.	Running time	Train'g acc.	Val. acc.	Test acc.
Base	65	1	32	297,697	7min16s	0.9736	0.9857	0.9855
New 1	65	1	32	8,478,817	29min28s	0.9974	0.9854	0.9853
New 2	33	1	128	522,337	167s	0.9959	0.9488	0.9455
New 3	33	2	128	522,337	187s	0.9990	0.9961	0.9977

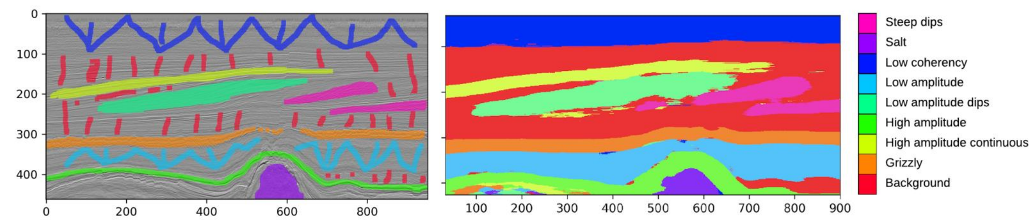

Figure 5: Labeled data (left) and predicted results (right) on the same inline 339 2D section. The "New-3" model was used for predictions.

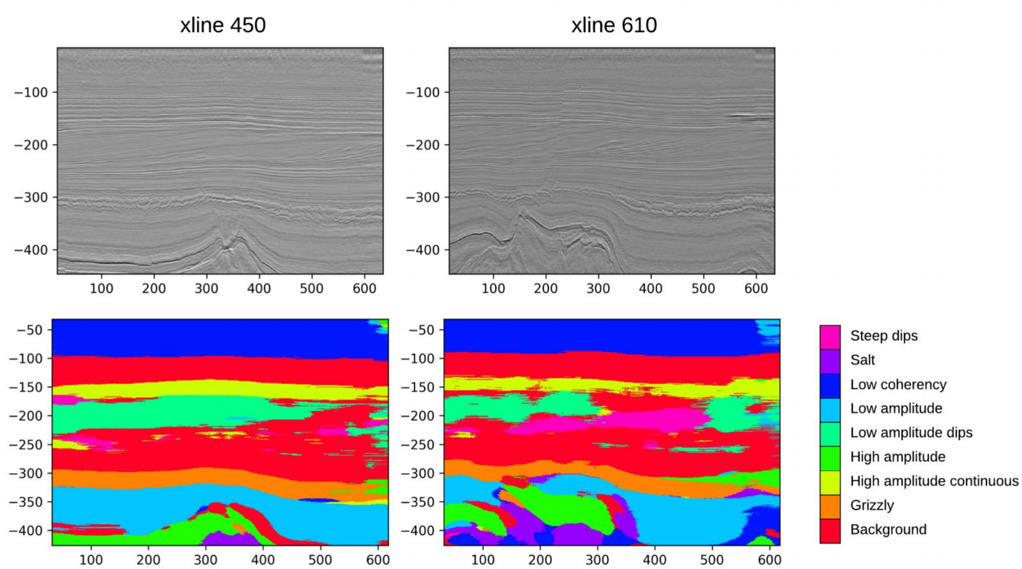

Figure 6: Raw seismic data (top) and prediction results (bottom) using the New 3 model on xline 450 (left) and xline 610 (right) seismic sections.

References

dGB Earth Sciences B.V., 1987, Netherlands offshore f3 block. (https://www.opendtect.org/osr/Main/NetherlandsOffshoreF3BlockComplete4GB).

Ildstad, C. R., and P. Bormann, 2017, Using 3d segy seismic and 3d convolutional neural networks to map and classify seismic facies: https://github.com/bolgebrygg/MalenoV.

Waldeland, A., and A. Solberg, 2017, Salt classification using deep learning: Presented at the 79th EAGE Conference and Exhibition 2017.