Deep convolutional neural networks as a geological image classification tool

Rafael Pires de Lima[1,2], Alicia Bonar[1], David Duarte Coronado[1], Kurt Marfurt[1], Charles Nicholson[3]

[1]School of Geology and Geophysics, The University of Oklahoma, 100 East Boyd Street, RM 710, Norman, Oklahoma, 73019, USA

[2]The Geological Survey of Brazil – CPRM, 55 Rua Costa, São Paulo, São Paulo, Brazil

[3]School of Industrial and Systems Engineering, The University of Oklahoma, 202 West Boyd Street, RM 124, Norman, Oklahoma, 73019, USA

Keywords

convolutional neural networks, transfer learning, automatization, microfossil identification, petrography

ABSTRACT

A convolutional neural network (CNN) is a deep learning (DL) method that has been widely and successfully applied to computer vision tasks including object localization, detection, and image classification. DL for supervised learning tasks is a method that uses the raw data to determine the classification features, in contrast to other machine learning (ML) techniques that require pre-selection of the input features (or attributes). In the geosciences, we hypothesize that deep learning will facilitate the analysis of uninterpreted images that have been neglected due to a limited number of experts, such as fossil images, slabbed cores, or petrographic thin sections. We use transfer learning, which employs previously trained models to shorten the development time for subsequent models, to address a suite of geologic interpretation tasks that may benefit from ML. Using two different base models, MobileNet V2 and Inception V3, we illustrate the successful classification of microfossils, core images, petrographic photomicrographs, and rock and mineral hand sample images. ML does not replace the expert geoscientist. The expert defines the labels (interpretations) needed to train the algorithm and also monitors the results to address incorrect or ambiguous classifications. ML techniques provide a means to apply the expertise of skilled geoscientists to much larger volumes of data.

INTRODUCTION

Machine learning (ML) techniques have been successfully applied, with considerable success, in the geosciences for almost two decades. Applications of ML by the geoscientific community include many examples such as seismic-facies classification (Meldahl et al., 2001; West et al., 2002; de Matos et al., 2011; Roy et al., 2014; Qi et al., 2016; Hu et al., 2017; Zhao et al., 2017), electrofacies classification (Allen and Pranter, 2016), and analysis of seismicity (Kortström et al., 2016; DeVries et al., 2018; Perol et al., 2018; Sinha et al., 2018), and classification of volcanic ash (Shoji et al., 2018), among others. Conventionally, ML applications rely on a set of attributes (or features) selected or designed by an expert. Features are specific characteristics of an object that can be used to study patterns or predict outcomes. In classification modeling, these features are chosen with the goal of distinguishing one object from another.

Typically, feature selection is problem dependent. For example, a clastic sedimentary rock is most broadly classified by its grain size; therefore a general classification for a rock sample (data) is sandstone if its grain sizes (features) lie from 0.06 mm to 2.0 mm following the Wentworth size class. In this example, a single feature is used to classify the sample, but more complex and/or detailed classification often requires analysis of multiple features exhibited by the sample. An inefficiency of traditional ML approaches is that many features may be constructed while only a subset of them are actually needed for the classification.

The use of explicitly designed features to classify data was the traditional approach in ML applications within the geosciences as in many other research areas. This classification approach works well when human interpreters know and can quantify the features that distinguish one object from another. However, sometimes an interpreter will subconsciously classify features and have difficulty describing what the distinguishing features might be, relying on "I'll know what the object is when I see it". In contrast to feature-driven ML classification algorithms, deep learning (DL) models extract information directly from the raw unstructured data rather than the data being manually transformed.

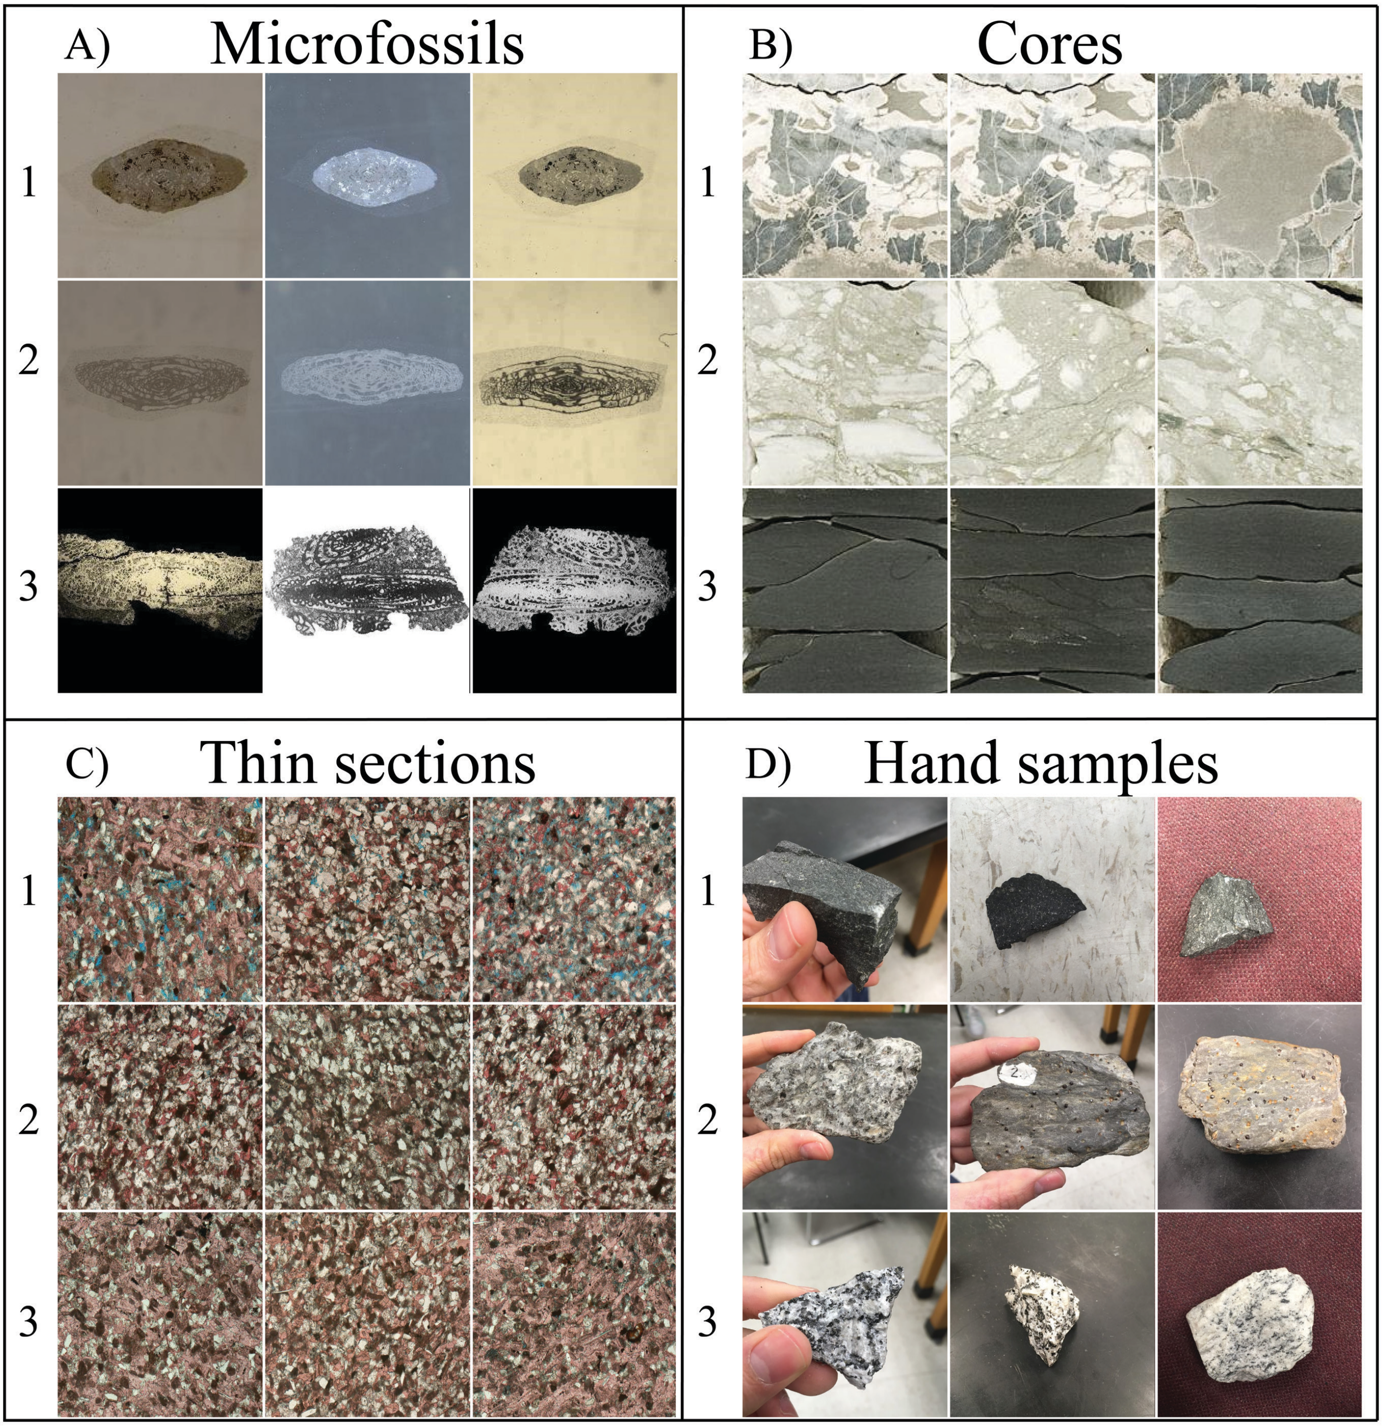

Figure 1: Examples of the data used in this study. A) Three of the seven Fusulinids groups (Beedeina (1), Fusulinella (2), and Parafusulina (3)). B) Three of the five lithofacies (bioturbated mudstone-wackestone (1), chert breccia (2), and shale (3)). C) Reservoir quality classes (high (1), intermediate (2), and low (3)) D) Three of the six rock sample groups (basalt (1), garnet schist (2), and granite (3)). Samples were interpreted by professionals working with each separate dataset.

Because of their greater complexity (and resulting flexibility and power) convolutional neural networks (CNN) usually requires more training data than traditional ML processes. However, when expert-labeled data are provided, non-experts can use the CNN models to generate highly accurate results (e.g. TGS Salt Identification Challenge | Kaggle, 2019).

DL applications in the geosciences require experts to first define the labels used to construct the necessary data sets as well as identify and address any ambiguous results and anomalies. In order to bring awareness and provide basic information regarding CNN models, DL techniques, and the necessity of expert-level knowledge needed to utilize these advancements, we applied these methods to four different geologic tasks. Figure 1 shows samples of different types of data that can be interpreted and labeled by experienced geologists. We use such interpretations to train our models. In this manuscript, we show how CNN can aid geoscientists with microfossil identification, core descriptions, petrographic analyses, and as a potential tool for education and outreach by creating a simple hand specimen identification application.

CONVOLUTIONAL NEURAL NETWORKS AND TRANSFER LEARNING

Recent CNN research has yielded significant improvements and unprecedented accuracy (the ratio between correct classifications and the total number of samples classified) in image classification and are recognized as leading methods for large-scale visual recognition problems, such as the annual ImageNet Large Scale Visual Recognition Challenge (ILSVRC, Russakovsky et al. (2015)). Specific CNN architectures have been the leading approach for several years now (e.g., Szegedy et al., 2014; Chollet, 2016; He et al., 2016; Huang et al., 2016; Sandler et al., 2018). Researchers noted that the parameters learned by the layers in many CNN models trained on images exhibit a common behavior – layers closer to the input data tend to learn general features, such as edge detecting/enhancing filters or color blobs, then there is a transition to more specific dataset features, such as faces, feathers, or object parts (Yosinski et al., 2014; Yin et al., 2017). These general-specific CNN layer properties are important points to be considered for the implementation of transfer learning (Caruana, 1995; Bengio, 2012; Yosinski et al., 2014). In transfer learning, first a CNN model is trained on a base dataset for a specific task. The learned features (model parameters) are repurposed, or transferred, to a second target CNN to be trained on a different dataset and task (Yosinski et al., 2014).

New DL applications often require large volumes of data, however the combination of CNNs and transfer learning allows the reuse of existing DL models to novel classification problems with limited data, as has been demonstrated in diverse fields, such as botany (Carranza-Rojas et al., 2017), cancer classification (Esteva et al., 2017), and aircraft detection (Chen et al., 2018). Analyzing medical image data, Tajbakhsh et al. (2016) and Qayyum et al. (2017) found that transfer learning achieved comparable or better results than training a CNN model with randomly initialized parameters. As an example, training the entire InceptionV3 (Szegedy et al., 2015) with 1000 images (five classes, 50 original images for each class, four copies of each original image) with randomly initialized parameters can be 10 times slower than the transfer learning process (11 minutes vs 1 minute on average for five executions) using a Nvidia Quadro M2000 (768 CUDA Cores). On a CPU (3.60 GHz clock speed), training the entire model can take up to 2 hours whereas transfer learning can be completed within a few minutes. We also noticed that transfer learning is easier to train. During the speed comparison test, transfer learning achieved high accuracies (close to 1.0) within 5 epochs (note the dataset is very simple with most of the samples being copies of each other). Successful applications of computer vision technologies in different fields suggest that ML models could be extremely beneficial for geologic applications, especially those in the category of image classification problems.

For the examples we present in this paper (Figure 1), we rely on the use of transfer learning (Yosinski et al., 2014) using the MobileNetV2 (Sandler et al., 2018) and InceptionV3 as our base CNN models. Both MobileNetV2 and InceptionV3 were trained on ILSVRC. Therefore, the CNN models we used were constructed based on inputs of 3-channels (RGB) of 2D photographic images. We randomly select part of the data to be used as a test set maintaining the same proportion of samples per class as in the training set. The data in the test set is not used during the

Table 1: Summary of test accuracy for the examples in this study.

Dataset	Number of training samples	Number of test samples	Number of output classes	MobileNetV2 Accuracy	InceptionV3 Accuracy
Microfossils (Fusulinids)	1480	184	7	1.00	1.00
Core	227	28	5	1.00	0.97
Petrographic thin-sections	194	31	3	0.81	0.81
Rock samples	1218	151	6	0.98	0.97

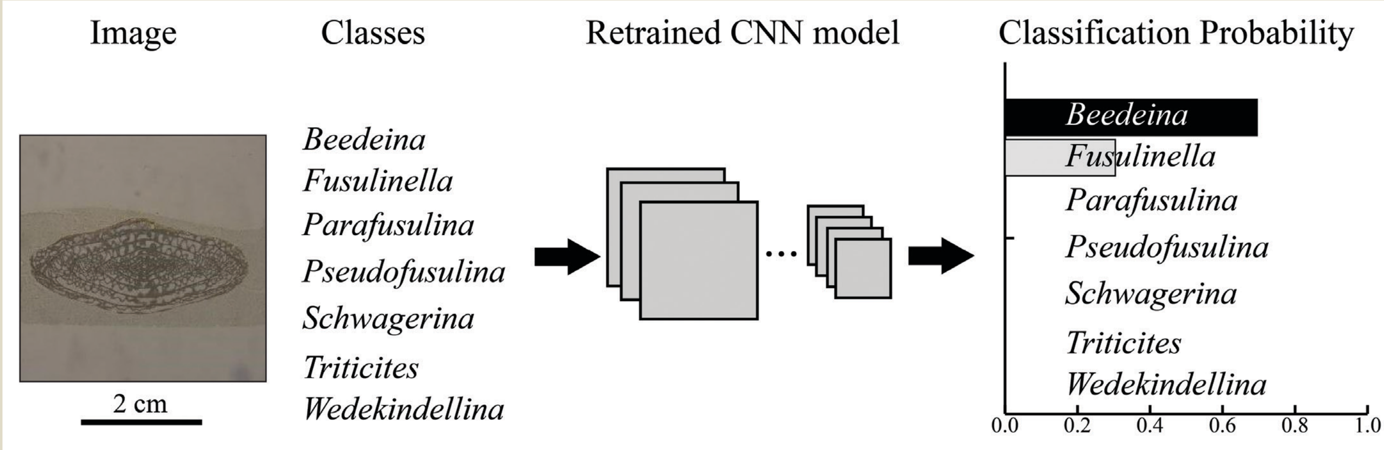

Figure 2: An example of the classification process. In this example, a thin-section image that should fit one of the seven fusulinid genera is analyzed by the model. The model outputs the probability assigned to each of the possible classes (all probabilities summing to 1.0). The term "classes" here is used in the ML sense rather than the biological one. In the example provided, our model provided a high probability for the same class as the human expert. Note that in the implementation we use the model will classify any image as one of the seven learned classes – even if the image is clearly not a fossil. This highlights the importance of a domain expert intervention.

computational process for model training; rather, it is used to evaluate the quality and robustness of the final model. Due to limited space, we refrained showing the CNN mistakes and many of the steps necessary for data preparation.

CNN-ASSISTED FOSSIL ANALYSIS

Biostratigraphy has become a less common focus of study in the discipline of paleontology (Farley and Armentrout, 2000, 2002), but the applications of biostratigraphy are necessary for understanding age-constraints for rocks that cannot be radiometrically dated. Access to a specific taxonomic expert to accurately analyze fossils at the species-level can be as challenging as data acquisition and preparation. Using labeled data from the University of Oklahoma Sam Noble Museum and iDigBio portal, we found that fusulinids (index fossils for the Late Paleozoic) can be accurately classified with the use of transfer learning. Accurate identification of a fusulinid depends on characteristics that must be observed and exposed along the long axis of the (prolate spheroid-shaped) fusulinid. We used a dataset of 1850 qualified images including seven different fusulinid genera. After retraining the CNN model, we obtained an accuracy for the test set (10% of the data) of 1.0 for both retrained MobileNetV2 and InceptionV3 (Table 1). Figure 2 shows a schematic view of the classification process.

CNN-ASSISTED CORE DESCRIPTION

Miles of drilled cores are stored in boxes in enormous warehouses, many of which have either been neglected for years or never digitally described. Core-based rock-type descriptions are important for understanding the lithology and structure of subsurface geology. Using several hundred feet of labeled core from a Mississippian limestone in Oklahoma (data from Suriamin and Pranter, 2018 and Pires de Lima et al., 2019), we selected a small sample of 285 images from five distinct lithofacies to be classified by the retrained CNN models. Pires de Lima et al. (2019) describes how a sliding window is used to generate CNN input data, cropping small sections from a standard core image. We used 10% of the data as the test set and achieved an accuracy of 1.0 using the retrained MobileNetV2 and an accuracy of 0.97 using the retrained InceptionV3 (Table 1).

CNN-ASSISTED RESERVOIR QUALITY CLASSIFICATION USING PETROGRAPHIC THIN SECTIONS

Petrography focuses on the microscopic description and classification of rocks and is one of the most important techniques in sedimentary and diagenetic studies. Potential information gained from thin section analysis compared to hand specimen descriptions include mineral distribution and percentage, pore space analysis, and cement composition. Petrographic analyses can be laborious even for experienced geologists. Using a total of 161 photomicrographs of parallel Nicol polarization of thin sections from the Sycamore Formation shale resource play in Oklahoma, we classified these images as representatives of high, intermediate, and low reservoir quality depending on the percent of calcite cement and pore space. We used 20% of the images in the test set and obtained a test set accuracy of 0.81 for both the retrained

MobileNetV2 and the retrained InceptionV3 (Table 1).

CNN-ASSISTED ROCK SAMPLE ANALYSIS

By creating a simple website, the general population could have immediate access to a rock identification tool using transfer learning technology. For this work in progress, we used smartphones to acquire 1521 pictures of six different rock types, using five different hand samples for each one of the rock types. We took pictures with different backgrounds, as visually depicted in Figure 1, however all pictures were taken in the same classroom. After retraining the CNN models, we obtained an accuracy for the test set (10% of original data) of 0.98 using the retrained MobileNetV2 and 0.97 using the retrained InceptionV3 (Table 1). We note that our model does not perform well with no-background images (i.e., pictures in which the rock sample is edited and seems to be within a white or black canvas) as such images were not used in training.

CONCLUSIONS AND FUTURE WORK

Although gaining popularity and becoming established as robust technologies in other scientific fields, transfer learning and CNN models are still novel with respect to application within the geoscience community. In this paper, we used CNN and transfer learning to address four potential applications that could improve data management, organization, and interpretation in different segments of our community. We predict that the versatile transfer learning and deep learning technologies will play a role in public education and community outreach, allowing the public to identify rock samples much as they currently can use smart phone apps to identify visitors to their bird feeder. Such public engagement will increase geological awareness and provide learning opportunities for elementary schools, outdoor organizations, and families.

For all of our examples, we were able to achieve high levels of accuracy (greater than 0.81) by repurposing two different CNN models originally assembled for generic computer vision tasks. We note that the examples and applications demonstrated here are curated, and therefore we expected highly accurate results. We presented demonstrations with limited classes and relatively well-controlled input images, so near perfect accuracies cannot necessarily be expected in an open, free-range deployment scenario. Regardless, the ability to create distinctive models for specific sets of images allows for a versatile application.

The techniques we have shown could greatly improve the speed of monotonous tasks such as describing miles of core data with very similar characteristics or looking at hundreds of thin sections from the same geologic formation. While the tasks are performed by the computer, the geoscience expert is still the most important element in every analysis in order to create the necessary datasets and provide quality control of the generated results. In the end, the expert validates the correctness of the results and looks for anomalies that are poorly represented by the target classes. We believe ML can help maintain consistency in interpretations and even provide a resource for less common observations and data variations, such as previously overlooked fossil subspecies and unique mineralogical assemblages in small communities and private collections, thereby building and reconciling a more complete international database. By combing expert knowledge and time efficient technology, ML methods can accelerate many data analysis processes for geologic research.

ACKNOWLEDGEMENTS

We thank the iDigBio initiative for providing access to the community for biodiversity collections data. Rafael acknowledges CNPq (grant 203589/2014-9) for the financial support and CPRM for granting the leave of absence allowing the pursuit of his Ph.D. studies. We thank Roger J. Burkhalter from the University of Oklahoma Sam Noble Museum of Natural History for providing the Fusulinids images used in this manuscript.

REFERENCES

ALLEN, D. B., AND M. J. PRANTER, 2016, Geologically constrained electrofacies classification of fluvial deposits: An example from the Cretaceous Mesaverde Group, Uinta and Piceance Basins: AAPG Bulletin, v. 100, no. 12, p. 1775–1801, doi:10.1306/05131614229.

BENGIO, Y., 2012, Deep Learning of Representations for Unsupervised and Transfer Learning, *in* I. Guyon, G. Dror, V. Lemaire, G. Taylor, and D. Silver, eds., Proceedings of ICML Workshop on Unsupervised and Transfer Learning: PMLR, Proceedings of Machine Learning Research, p. 17–36.

CARRANZA-ROJAS, J., H. GOEAU, P. BONNET, E. MATA-MONTERO, AND A. JOLY, 2017, Going deeper in the automated identification of Herbarium specimens: BMC Evolutionary Biology, v. 17, no. 1, p. 181, doi:10.1186/s12862-017-1014-z.

CARUANA, R., 1995, Learning Many Related Tasks at the Same Time with Backpropagation, *in* G. Tesauro, D. S. Touretzky, and T. K. Leen, eds., Advances in Neural Information Processing Systems 7: MIT Press, p. 657–664.

CHEN, Z., T. ZHANG, C. OUYANG, Z. CHEN, T. ZHANG, AND C. OUYANG, 2018, End-to-End Airplane Detection Using Transfer Learning in Remote Sensing Images: Remote Sensing, v. 10, no. 1, p. 139, doi:10.3390/rs10010139.

CHOLLET, F., 2016, Xception: Deep Learning with Depthwise Separable Convolutions: CoRR, v. abs/1610.0.

DEVRIES, P. M. R., F. VIÉGAS, M. WATTENBERG, AND B. J. MEADE, 2018,

Deep learning of aftershock patterns following large earthquakes: Nature, v. 560, no. 7720, p. 632–634, doi:10.1038/s41586-018-0438-y.
ESTEVA, A., B. KUPREL, R. A. NOVOA, J. KO, S. M. SWETTER, H. M. BLAU, AND S. THRUN, 2017, Dermatologist-level classification of skin cancer with deep neural networks: Nature, v. 542, no. 7639, p. 115–118, doi:10.1038/nature21056.
FARLEY, M. B., AND J. M. ARMENTROUT, 2000, Fossils in the Oil Patch: Geotimes, p. 14–17.
FARLEY, M. B., AND J. M. ARMENTROUT, 2002, Tools, Biostratigraphy becoming lost art in rush to find new exploration: Offshore, p. 94–95.
HE, K., X. ZHANG, S. REN, AND J. SUN, 2016, Deep Residual Learning for Image Recognition, *in* 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR): IEEE, p. 770–778, doi:10.1109/CVPR.2016.90.
HU, S., W. ZHAO, Z. XU, H. ZENG, Q. FU, L. JIANG, S. SHI, Z. WANG, AND W. LIU, 2017, Applying principal component analysis to seismic attributes for interpretation of evaporite facies: Lower Triassic Jialingjiang Formation, Sichuan Basin, China: Interpretation, v. 5, no. 4, p. T461–T475, doi:10.1190/INT-2017-0004.1.
HUANG, G., Z. LIU, AND K. Q. WEINBERGER, 2016, Densely Connected Convolutional Networks: CoRR, v. abs/1608.0.
KORTSTRÖM, J., M. USKI, AND T. TIIRA, 2016, Automatic classification of seismic events within a regional seismograph network: Computers & Geosciences, v. 87, p. 22–30, doi:10.1016/J.CAGEO.2015.11.006.
DE MATOS, M. C., M. (MOE) YENUGU, S. M. ANGELO, AND K. J. MARFURT, 2011, Integrated seismic texture segmentation and cluster analysis applied to channel delineation and chert reservoir characterization: Geophysics, v. 76, no. 5, p. P11–P21, doi:10.1190/geo2010-0150.1.
MELDAHL, P., R. HEGGLAND, B. BRIL, AND P. DE GROOT, 2001, Identifying faults and gas chimneys using multiattributes and neural networks: The Leading Edge, v. 20, no. 5, p. 474–482, doi:10.1190/1.1438976.
PEROL, T., M. GHARBI, AND M. DENOLLE, 2018, Convolutional neural network for earthquake detection and location: Science Advances, v. 4, no. 2, p. e1700578, doi:10.1126/sciadv.1700578.
PIRES DE LIMA, R., F. SURIANAM, K. J. MARFURT, AND M. J. PRANTER, 2019, Convolutional neural networks as aid in core lithofacies classification: Interpretation, v. 7, no. 3, doi:10.1190/int-2018-0245.1.
QAYYUM, A., S. M. ANWAR, M. AWAIS, AND M. MAJID, 2017, Medical image retrieval using deep convolutional neural network: Neurocomputing, v. 266, p. 8–20, doi:10.1016/J.NEUCOM.2017.05.025.
QI, J., T. LIN, T. ZHAO, F. LI, AND K. MARFURT, 2016, Semisupervised multiattribute seismic facies analysis: Interpretation, v. 4, no. 1, p. SB91–SB106, doi:10.1190/INT-2015-0098.1.
ROY, A., A. S. ROMERO-PELÁEZ, T. J. KWIATKOWSKI, AND K. J. MARFURT, 2014, Generative topographic mapping for seismic facies estimation of a carbonate wash, Veracruz Basin, southern Mexico: Interpretation, v. 2, no. 1, p. SA31–SA47, doi:10.1190/INT-2013-0077.1.
RUSSAKOVSKY, O. ET AL., 2015, ImageNet Large Scale Visual Recognition Challenge: International Journal of Computer Vision, v. 115, no. 3, p. 211–252, doi:10.1007/s11263-015-0816-y.
SANDLER, M., A. HOWARD, M. ZHU, A. ZHMOGINOV, AND L.-C. CHEN, 2018, MobileNetV2: Inverted Residuals and Linear Bottlenecks: ArXiv e-prints.
SHOJI, D., R. NOGUCHI, S. OTSUKI, AND H. HINO, 2018, Classification of volcanic ash particles using a convolutional neural network and probability: Scientific Reports, v. 8, no. 1, p. 8111, doi:10.1038/s41598-018-26200-2.
SINHA, S., Y. WEN, R. A. PIRES DE LIMA, AND K. MARFURT, 2018, Statistical controls on induced seismicity: Unconventional Resources Technology Conference, doi:10.15530/urtec-2018-2897507-MS.
SURIAMIN, F., AND M. J. PRANTER, 2018, Stratigraphic and lithofacies control on pore characteristics of Mississippian limestone and chert reservoirs of north-central Oklahoma: Interpretation, p. 1–66, doi:10.1190/int-2017-0204.1.
SZEGEDY, C., W. LIU, Y. JIA, P. SERMANET, S. E. REED, D. ANGUELOV, D. ERHAN, V. VANHOUCKE, AND A. RABINOVICH, 2014, Going Deeper with Convolutions: CoRR, v. abs/1409.4.
SZEGEDY, C., V. VANHOUCKE, S. IOFFE, J. SHLENS, AND Z. WOJNA, 2015, Rethinking the Inception Architecture for Computer Vision: CoRR, v. abs/1512.0.
TAJBAKHSH, N., J. Y. SHIN, S. R. GURUDU, R. T. HURST, C. B. KENDALL, M. B. GOTWAY, AND J. LIANG, 2016, Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning? IEEE Transactions on Medical Imaging, v. 35, no. 5, p. 1299–1312, doi:10.1109/TMI.2016.2535302.
TGS Salt Identification Challenge | Kaggle, n.d.: <https://www.kaggle.com/c/tgs-salt-identification-challenge> (accessed January 10, 2019).
WEST, B. P., S. R. MAY, J. E. EASTWOOD, AND C. ROSSEN, 2002, Interactive seismic facies classification using textural attributes and neural networks: The Leading Edge, v. 21, no. 10, p. 1042–1049, doi:10.1190/1.1518444.
YIN, X., W. CHEN, X. WU, AND H. YUE, 2017, Fine-tuning and visualization of convolutional neural networks, *in* 2017 12th IEEE Conference on Industrial Electronics and Applications (ICIEA): IEEE, p. 1310–1315, doi:10.1109/ICIEA.2017.8283041.
YOSINSKI, J., J. CLUNE, Y. BENGIO, AND H. LIPSON, 2014, How transferable are features in deep neural networks? Advances in Neural Information Processing Systems, v. 27, p. 3320–3328.
ZHAO, T., F. LI, AND K. J. MARFURT, 2017, Constraining self-organizing map facies analysis with stratigraphy: An approach to increase the credibility in automatic seismic facies classification: Interpretation, v. 5, no. 2, p. T163–T171, doi:10.1190/INT-2016-0132.1.

Accepted June 2019