

De-risking Carbon Capture and Sequestration with Explainable CO₂ Leakage Detection in Time-lapse Seismic Monitoring Images

Huseyin Tuna Erdinc,^{*1} Abhinav Prakash Gahlot,^{*2} Ziyi Yin,³
Mathias Louboutin,² Felix J. Herrmann,^{1,2,3}

¹School of Electrical and Computer Engineering, Georgia Institute of Technology

²School of Earth and Atmospheric Sciences, Georgia Institute of Technology

³School of Computational Science and Engineering, Georgia Institute of Technology
{herdinc3, agahlot8, ziyi.yin, mlouboutin3, felix.herrmann}@gatech.edu,

Abstract

With the growing global deployment of carbon capture and sequestration technology to combat climate change, monitoring and detection of potential CO₂ leakage through existing or storage induced faults are critical to the safe and long-term viability of the technology. Recent work on time-lapse seismic monitoring of CO₂ storage has shown promising results in its ability to monitor the growth of the CO₂ plume from surface recorded seismic data. However, due to the low sensitivity of seismic imaging to CO₂ concentration, additional developments are required to efficiently interpret the seismic images for leakage. In this work, we introduce a binary classification of time-lapse seismic images to delineate CO₂ plumes (leakage) using state-of-the-art deep learning models. Additionally, we localize the leakage region of CO₂ plumes by leveraging Class Activation Mapping methods.

Introduction

According to the International Energy Agency and the International Panel on Climate Change report (IPCC 2018), there is a need for a 50 percent reduction of greenhouse gas emissions by 2050 to avoid an increase of 1.5 degrees Celsius of Earth’s average temperature. This can only be achieved by reduced dependence on fossil fuels, use of renewable sources of energy and large-scale global deployment of carbon reduction technologies such as carbon capture and sequestration (CCS). This technology consists of collection, transportation, and injection of CO₂ into an appropriate geologic storage reservoir for extended time periods (tens of years). Especially, unlike other solutions, CCS is considered a relatively low-cost, long-term and imminent solution. However, potential CO₂ leakage from the underground reservoirs due to pre-existing or pressure-induced faults poses risks (Ringrose 2020). Thus, it is necessary to de-risk CCS projects by monitoring CO₂ plumes in order to accurately detect and predict potential leakages as early as possible.

Time-lapse seismic monitoring has been introduced as a reliable technology to monitor the CO₂ dynamics in the Earth’s subsurface during carbon sequestration (Lumley

2001) and is already in use at existing storage sites (Arts et al. 2008; Chadwick et al. 2010; Ringrose et al. 2013; Furre et al. 2017). In essence, sequential (i.e once every 6 months/year/...) seismic datasets, called vintages, are collected in the field over an area covering the storage reservoir. Then, each seismic dataset is inverted to obtain high fidelity images of the subsurface over time (Arts et al. 2008; Ayeni and Biondi 2010; Yin, Louboutin, and Herrmann 2021). The evolution of the CO₂ reservoir can finally be visualized by subtracting the seismic images between different points in time. However, due to the inherently weak and noisy amplitudes of the CO₂ reservoir’s response in those seismic images, detecting the presence of potential irregularities, such as in CO₂ plumes, corresponding to a leakage is a challenging problem. To tackle this difficulty, we propose a machine learning based detection method based on standard binary classification.

Recently, numerous methods leveraging machine learning have been introduced for the detection of CO₂ leakage based on a simple artificial neural network (ANN) (Li et al. 2018a), and a combination of convolutional neural networks (CNN) and Long Short-Term Memory (LSTM) networks (Zhou et al. 2019). While leading to accurate predictions, these methods usually rely solely on the field recorded data rather than the subsurface seismic images. Besides, practical considerations such as repeatability (the ability to record the data in the exact same way every year) hinders their applicability to real world cases. On the other hand, as we rely on visualizing the CO₂ plumes in the seismic image, we can take advantage of advanced seismic imaging techniques designed for non-repeated seismic acquisition such as the joint recovery model (JRM) (Oghenekohwo and Herrmann 2017a; Wason, Oghenekohwo, and Herrmann 2017; Yin, Louboutin, and Herrmann 2021). Additionally, this imaging technique has demonstrated higher fidelity imaging than sequential seismic imaging allowing for easier detection of CO₂ leakage.

We will show in the following sections that we can efficiently and accurately detect CO₂ from realistic seismic images recovered by JRM on synthetic but representative models of the Earth subsurface. We demonstrate our method using different state-of-the-art deep learning models in a transfer learning setting to classify CO₂ plume seismic images

^{*}These authors contributed equally.

with regular (no-leakage) CO₂ plume or with CO₂ leakage. As CO₂ leakage detection needs trustworthiness, we further unravel the decisions made by our models and utilize Class Activation Mapping (CAM) methods (Zhou et al. 2015) to identify and visualize seismic image areas crucial for model classification results. We show that the CAM result accurately focuses on the leakage portion of the CO₂ plume and reservoir, validating that our network detects leakage based on state of the CO₂ reservoir over time.

Our main contributions are organized as follows. First, we introduce the classification models used for leakage detection and the CAM methods for visualizing the area of interest in the seismic image. Second, we demonstrate the accuracy of our models and qualitatively examine the results of our CAM methods on a realistic synthetic set of CO₂ plume images.

Methodology

In order to speed up the training process and to compensate for the overfitting that may occur with modest sized datasets, we rely on transfer learning (Yosinski et al. 2014) using pre-trained state-of-the-art models as a starting point. In particular, we consider four modern architectures known to achieve high accuracy on standard dataset such as ImageNet-1k (Russakovsky et al. 2015). The models used are VGG (Simonyan and Zisserman 2014), ResNet (He et al. 2016), Vision Transformer (ViT) (Dosovitskiy et al. 2021), and Swin Transformer (Swin) (Liu et al. 2021), all pre-trained on the standardized ImageNet-1k dataset.

VGG: is a convolutional neural network (CNN) model that achieved significant success in The ImageNet Large Scale Visual Recognition Challenge (ILSVRC) competition in 2014 (Simonyan and Zisserman 2014). VGG consists of sequences of convolution and maxpool layers. In our numerical experiments, the VGG16 variant with 16 trainable layers is used.

ResNet: is a CNN architecture with residual connections proposed to solve the vanishing gradient problem in very deep networks (He et al. 2016). ResNet consists of residual blocks and each residual block has convolution layers and shortcut connections performing identity mapping. In our numerical experiments, the ResNet34 variant with 34 trainable layers is used.

ViT: is an architecture based on transformer which is used in the field of Natural Language Processing (NLP) (Vaswani et al. 2017). Internally, the transformer learns a relationship between input token pairs, and uses 16x16 patches of images as input tokens (Dosovitskiy et al. 2021). In our numerical experiments, the tiny ViT variant is used allowing lower memory and computational imprint.

Swin: is a special type of ViT that represents image patches hierarchically by starting from small-sized patches and gradually increasing the size through merging to achieve scale-invariance property (Liu et al. 2021). Compared to ViT, Swin transformer has superior (linear) computational efficiency by computing self-attention within certain patches of a window. In our numerical experiments, tiny Swin variant is used allowing lower memory and computational imprint.

Hyperparameters	VGG16	ResNet34	ViT	Swin
Batch Size	8	8	8	8
Learning Rate	5×10^{-5}	6×10^{-3}	4×10^{-3}	10^{-3}
Exp Decay Rate(γ)	0.95	0.92	0.98	0.98

Table 1: Training hyperparameters for the four models. All models were trained with the same number of epochs and optimizer.

CAM Methods

Deep learning models for classification are notoriously treated as “black boxes” as they do not expose their internal knowledge or operations to its users and do not provide interpretable results. To solve this problem, CAM based saliency maps (heatmaps) were introduced to highlight the most class-discriminative regions of to-be-classified input images (Zhou et al. 2015). Since CO₂ leakage requires high fidelity, transparent and interpretable models, we use CAM to further make our model results explainable and highlight the regions of the seismic image that are most relevant to the classification results. In our study, we considered two CAM methods. First, Grad-CAM (Selvaraju et al. 2019), a gradient-based CAM method considered as the state-of-the-art in terms of explainability of neural networks for classification. This CAM method extracts gradients from a specific layer of a model and computes the weighted average of that specific layer’s activations. Second, we consider Score-CAM (Wang et al. 2020), a perturbation based CAM method. Score-CAM also computes the weighted average of activations of a user-specified layer but, unlike Grad-CAM, Score-CAM relies on propagating (forward pass through the network) a masked input image where the mask is obtained via upsampling the activations of the user-defined layer. This CAM method provides high accuracy and interpretable heatmaps and alleviates potential noise and spread present in the gradient used for the Grad-CAM heatmaps.

Numerical Case Study

To generate the training dataset of CO₂ plume evolution, we used five 2D vertical slices extracted from the 3D Compass velocity model (E. Jones et al. 2012) shown in Fig. 1(a). This model is a synthetic but realistic model representative of the complex geology of the southeast of the North Sea. The dimension of each model (slice) used in our work is 2131 X 4062 m². We used FwiFlow (Li et al. 2020), to simulate the CO₂ flow dynamics and JUDI (Witte et al. 2019) to model the seismic data and compute the seismic images of the subsurface.

Time-lapse reservoir and seismic simulation

We consider a realistic two well setting with a fixed injection well injecting CO₂ and a production well extracting brine from subsurface storage reservoir. Injection of supercritical CO₂ into saline aquifers is an example of multi-phase flow in porous media. While we understand more complicated geothermal, geochemical and geomechanical process may eventually be considered to model the CO₂ dynamics, we follow the two-phase immiscible incompressible flow




Figure 1: Workflow for CO₂ Leakage Monitoring

physics, which in its leading order describes the process of supercritical CO₂ displacing brine in the pore space of the rock. The system is governed by conservation of mass and Darcy’s law. We refer to the existing literature (Li et al. 2020; Wen, Tang, and Benson 2021) (Li et al. 2020) for more details about this physical system.

Using empirical relation and the Kozeny-Carman equation (Costa 2006), the acoustic properties (velocity and density) from the Compass model were converted into permeability and porosity (Fig. 1(b)) to simulate the multi-phase flow (CO₂ and brine in porous media) in the reservoir. We used FwiFlow.jl (Li et al. 2020) to solve multi-phase flow equations based on the finite volume method. We simulated the CO₂ flow for a duration varying between 7 to 12 years (Fig. 1(c)). The reservoir was initially filled with saline water and we injected compressed CO₂ at the rate of 1 MT/day into the reservoir for all simulations. In order to mimic CO₂ leakage, we then created a fracture at a random location along the top seal of the reservoir when the pressure induced by the CO₂ injection reaches a threshold of 15 MPa. We then converted back these simulated CO₂ saturation snapshots over time into wave properties with the patchy saturation model (Avseth, Mukerji, and Mavko 2010) to obtain time-lapse subsurface models (Fig. 1(d)). We used this model because at higher pressure condition, local fluid flow slows down resulting in an acoustic velocity trend which follows patchy saturation (Li et al. 2018b).

Based on these models, we then simulated the baseline seismic survey corresponding to the initial stage (before the injection of CO₂) and the monitor seismic survey corresponding to the final stage at the end of the reservoir simulation (Fig. 1(e)). As mentioned in the introduction, it is very difficult to exactly replicate the baseline and monitor surveys. In order to mimic the realistic scenario in the field, the baseline and monitor datasets were simulated using different acquisition geometries (position of the measurements). Finally, we recovered the time-lapse seismic images using JRM (Oghenekohwo and Herrmann 2017b; Watson, Oghenekohwo, and Herrmann 2017; Yin, Louboutin, and Herrmann 2021) to alleviate potential noise and inaccuracies in the seismic images in the case of non-replicated

time-lapse surveys. These recovered images along with the label (leakage/no-leakage) serve as the input to the classification network. We generated a total of 1000 leakage and 870 no-leakage scenarios, and computed the baseline, monitor and difference images with the JRM method in each case.

Training

The seismic difference images (difference between baseline and monitor recovery results) were converted to 224x224 gray-scale images with bi-linear interpolation and transformed into three channel images where each channel is a copy of the actual gray-scale image. For the classification, the image dataset was randomly split into an 80% training set and 20% test set. The training set was then further divided into two parts, one for model parameter updating (training) and another for hyperparameter tuning (validation). The training hyperparameters from this second part are summarized in Table 1. For training, we replaced the last fully connected layers (classification layers) of each model with a new fully connected layer. We then trained the network (Fig. 1(g)) in two steps. First, we only trained the last classification layer, by freezing all the other layers, for 100 epochs. Since most of the layers are fixed and do not need gradient updates, this first stage is extremely cheap and computationally efficient. Second, we further trained the full model for an additional 30 epochs to allow fine-tuning of all layers for our specific classification task. Following standard practices in classification settings, we used a binary cross-entropy loss function and the Adam optimizer (Kingma and Ba 2015) for all models. Finally, after the training (Fig. 1(h)), we implemented the CAM based methods (Fig. 1(i)). We used the last convolutional layer activations for the CNN models, and the activations preceding the last attention layer for the transformer-based models.

Analysis

We show on Table 2, different performance metrics on our testing dataset, after training our four networks, with means and confidence intervals after 15 different runs. In detail, we show standard metrics such as accuracy, precision, and recall. Additionally, we also show F1 score (Chinchor 1992),

Model	Accuracy	Precision	Recall	F1	ROC-AUC
VGG16	0.920 ± 0.089	0.941 ± 0.133	0.921 ± 0.081	0.927 ± 0.075	0.920 ± 0.076
ResNet34	0.948 ± 0.020	0.982 ± 0.028	0.928 ± 0.044	0.948 ± 0.040	0.967 ± 0.019
ViT	0.857 ± 0.018	0.910 ± 0.102	0.820 ± 0.098	0.859 ± 0.036	0.923 ± 0.023
Swin	0.836 ± 0.036	0.881 ± 0.108	0.818 ± 0.078	0.841 ± 0.076	0.909 ± 0.007

Table 2: Comparison of performance (for precision and recall, positives represent leakage whereas negatives are no leakage) on the test dataset for our four neural networks. The highest performance for each metric is highlighted in bold.




Figure 2: Grad-CAM and Score-CAM saliency maps overlaid on the corresponding input seismic image containing a CO₂ plume from leakage. The CO₂ plume can be seen on the seismic image as the high amplitude event at 1.3km depth and 1.8km in X.

that combine recall and accuracy, and area under curve of receiver operating characteristic (ROC-AUC) (Bradley 1997) to further evaluate the classification performance of models. We observe in Table 2 that the CNN models outperform the transformer variants in all the metrics by a significant margin and that ResNet34 achieves the best performance in all the measures of evaluation. This result is consistent with the literature, hinting that despite being very accurate on a specific task, transformers do not generalize well with our modest sized dataset (Dosovitskiy et al. 2021). Additionally, we observe that all models lead to better precision compared to recall (more false negatives than false positives). This discrepancy can be attributed to the fact that certain leakage images have very small CO₂ leakage areas (up to a single pixel) in the seismic images and are consequently very difficult to detect.

Second, we show in Fig. 2 the CAM results of each model on a single seismic image from our test dataset. The high amplitude area shows the regions of the seismic images that are most important to the classifier. As expected, those heatmaps provide an explainable representation of the classification as the high amplitudes align with the CO₂ leakage part of the seismic image. We observe that for the CNN, the saliency maps are well centered on the CO₂ leakage portion despite being very coarse. Because of this coarseness, both Grad-CAM and score-CAM provide similar results. On the other hand, transformer-based networks lead to more fo-

cused saliency maps that target the location of the CO₂ leakage extremely well. We observe in that case, the Score-CAM leads to reduction of aliases and noise compared to the Grad-CAM results. This can be linked to the potential presence of noise in the gradients of the transformers as the networks are very deep (Wang et al. 2020).

Conclusion

We have introduced an interpretable deep-learning method for CO₂ leakage detection with very high accuracy on a synthetic but realistic model of a CO₂ sequestration reservoir. First, we showed through four state-of-the-art models that we can detect potential CO₂ leakage from the recovered time-lapse seismic images. Second, we demonstrated that CAM provides an interpretable and accurate visualization of the CO₂ plume in case of leakage. Additionally, we showed that transformer-based models (ViT, Swin) led to more focused CAM and that Score-CAM provided cleaner and therefore more explainable heatmaps. On the other hand, we found that standard CNNs led to better classification results and therefore better leakage detection. In particular, ResNet model performed best and achieved a very high score above 90% in every evaluation metric. Future work will focus on improving the classification network to achieve higher accuracy in leakage detection and on refining the heatmaps for better explainability.

Acknowledgments

This research was carried out with the support of Georgia Research Alliance and partners of the ML4Seismic Center. The authors thank Philipp A. Witte at Microsoft for the constructive discussion.

References

- Arts, R. J.; Chadwick, A.; Eiken, O.; Thibeu, S.; and Nooner, S. L. 2008. Ten years' experience of monitoring CO₂ injection in the Utsira Sand at Sleipner, offshore Norway. *First Break*, 26.
- Avseth, P.; Mukerji, T.; and Mavko, G. 2010. *Quantitative seismic interpretation: Applying rock physics tools to reduce interpretation risk*. Cambridge university press.
- Ayeni, G.; and Biondi, B. 2010. Target-oriented joint least-squares migration/inversion of time-lapse seismic data sets. *Geophysics*, 75.
- Bradley, A. P. 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7): 1145–1159.
- Chadwick, A.; Williams, G.; Delepine, N.; Clochard, V.; Labat, K.; Sturton, S.; Buddensiek, M.-L.; Dillen, M.; Nickel, M.; Lima, A. L.; Arts, R.; Neele, F.; and Rossi, G. 2010. Quantitative analysis of time-lapse seismic monitoring data at the Sleipner CO₂ storage operation. *The Leading Edge*, 29(2): 170–177.
- Chinchor, N. 1992. MUC-4 Evaluation Metrics. In *Proceedings of the 4th Conference on Message Understanding*, MUC4 '92, 22–29. USA: Association for Computational Linguistics. ISBN 1558602739.
- Costa, A. 2006. Permeability-porosity relationship: A reexamination of the Kozeny-Carman equation based on a fractal pore-space geometry assumption. *Geophysical Research Letters*, 33(2).
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Hounsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*.
- E. Jones, C.; A. Edgar, J.; I. Selva, J.; and Crook, H. 2012. Building Complex Synthetic Models to Evaluate Acquisition Geometries and Velocity Inversion Technologies. *European Association of Geoscientists & Engineers*, cp-293-00580.
- Furre, A.-K.; Eiken, O.; Alnes, H.; Vevatne, J. N.; and Kær, A. F. 2017. 20 Years of Monitoring CO₂-injection at Sleipner. *Energy Procedia*, 114: 3916–3926. 13th International Conference on Greenhouse Gas Control Technologies, GHGT-13, 14-18 November 2016, Lausanne, Switzerland.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- IPCC. 2018. Global warming of 1.5°C. An IPCC Special Report on the impacts of global warming of 1.5°C above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. *In Press*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. arXiv preprint arXiv:1412.6980.
- Li, B.; Zhou, F.; Li, H.; Duguid, A.; Que, L.; Xue, Y.; and Tan, Y. 2018a. Prediction of CO₂ leakage risk for wells in carbon sequestration fields with an optimal artificial neural network. *International Journal of Greenhouse Gas Control*, 68: 276–286.
- Li, D.; Wei, J.; Di, B.; Ding, P.; Huang, S.; and Shuai, D. 2018b. Experimental study and theoretical interpretation of saturation effect on ultrasonic velocity in tight sandstones under different pressure conditions. *Geophysical Journal International*, 212: 2226–2237.
- Li, D.; Xu, K.; Harris, J. M.; and Darve, E. 2020. Coupled Time-Lapse Full-Waveform Inversion for Subsurface Flow Problems Using Intrusive Automatic Differentiation. *Water Resources Research*, 56(8): e2019WR027032. E2019WR027032 10.1029/2019WR027032.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. *In ICCV*.
- Lumley, D. E. 2001. Time-lapse seismic reservoir monitoring. *GEOPHYSICS*, 66(1): 50–53.
- Oghenekohwo, F.; and Herrmann, F. J. 2017a. Highly repeatable time-lapse seismic with distributed Compressive Sensing—mitigating effects of calibration errors. *The Leading Edge*, 36(8): 688–694. (The Leading Edge).
- Oghenekohwo, F.; and Herrmann, F. J. 2017b. Improved time-lapse data repeatability with randomized sampling and distributed compressive sensing. In *EAGE Annual Conference Proceedings*. (EAGE, Paris).
- Ringrose, P. 2020. *How to store CO₂ underground: Insights from early-mover CCS Projects*, volume 129. Springer.
- Ringrose, P.; Mathieson, A.; Wright, I.; Selama, F.; Hansen, O.; Bissell, R.; Saoula, N.; and Midgley, J. 2013. The In Salah CO₂ Storage Project: Lessons Learned and Knowledge Transfer. *Energy Procedia*, 37: 6226–6236. GHGT-11 Proceedings of the 11th International Conference on Greenhouse Gas Control Technologies, 18-22 November 2012, Kyoto, Japan.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252.
- Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2019. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *International Journal of Computer Vision*, 128(2): 336–359.
- Simonyan, K.; and Zisserman, A. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv preprint arXiv:1409.1556*.

- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention Is All You Need. *In Advances in Neural Information Processing Systems*, 6000–6010.
- Wang, H.; Wang, Z.; Du, M.; Yang, F.; Zhang, Z.; Ding, S.; Mardziel, P.; and Hu, X. 2020. Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks. *In CVPR*.
- Wason, H.; Oghenekohwo, F.; and Herrmann, F. J. 2017. Low-cost time-lapse seismic with distributed compressive sensing—Part 2: impact on repeatability. *Geophysics*, 82(3): P15–P30. (Geophysics).
- Wen, G.; Tang, M.; and Benson, S. M. 2021. Towards a predictor for CO₂ plume migration using deep neural networks. *International Journal of Greenhouse Gas Control*, 105: 103223.
- Witte, P. A.; Louboutin, M.; Kukreja, N.; Luporini, F.; Lange, M.; Gorman, G. J.; and Herrmann, F. J. 2019. A large-scale framework for symbolic implementations of seismic inversion algorithms in Julia. *Geophysics*, 84(3): F57–F71. (Geophysics).
- Yin, Z.; Louboutin, M.; and Herrmann, F. J. 2021. Compressive time-lapse seismic monitoring of carbon storage and sequestration with the joint recovery model. *In SEG Technical Program Expanded Abstracts*, 3434–3438. (IMAGE, Denver).
- Yosinski, J.; Clune, J.; Bengio, Y.; and Lipson, H. 2014. How transferable are features in deep neural networks?. *In Advances in neural information processing systems*, 3320–3328.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2015. Learning Deep Features for Discriminative Localization. *In IEEE CVPR*, 2921–2929.
- Zhou, Z.; Lin, Y.; Zhang, Z.; Wu, Y.; Wang, Z.; Dilmore, R.; and Guthrie, G. 2019. A data-driven CO₂ leakage detection using seismic data and spatial-temporal densely connected convolutional neural networks. *International Journal of Greenhouse Gas Control*, 90: 102790.