134

Enhancing early CO_2 leakage detection using pre-trained deep learning models

G. Pantaleo[1], M. Pipan[1]

[1] University Of Trieste

Summary

This study presents a deep learning framework for early CO_2 leakage detection from time-lapse seismic data, leveraging a pre-trained ResNet152 model originally trained for natural images classification. The framework integrates forward seismic modeling, deep feature extraction, and leakage mass prediction. ResNet152 is adapted to extract embeddings from time-lapse seismic shot gathers, which are compared using a Siamese network setup. Time-lapse embeddings are projected into a 2D latent space using UMAP, enabling anomaly detection and temporal pattern analysis. A multilayer perceptron (MLP) regressor is then trained to estimate leakage mass from the extracted embeddings. Tests on synthetic datasets from the Kimberlina CO_2 storage site simulations show that the model effectively captures the temporal evolution of seismic signal related to CO_2 leakage, particularly during early stages. UMAP projections reveal clear separability between baseline, no-leakage, and leakage scenarios. Despite being trained on non-seismic data, the model produces informative embeddings that support leakage detection. The MLP regressor achieves satisfactory performance under realistic noise levels. This framework demonstrates the potential of using pre-trained CNNs for seismic monitoring tasks and is extendable to other architectures. It offers a computationally efficient, generalizable approach that can serve as an early-warning to flag anomalies and guide targeted surveys.

Enhancing early CO_2 leakage detection using pre-trained deep learning models

Introduction

Underground carbon-dioxide storage plays a crucial role to play in meeting net-zero emissions targets by 2050. The diffusion of Carbon Capture and Storage (CCS) technologies is often hampered by cost constraints, particularly during the characterization and monitoring phase of deep reservoirs, such as aquifers or depleted gas fields. For this reason, there is an urgent need to develop cost-effective monitoring strategies that primarily target key structures and ensure CO_2 containment within the storage complex. Time-lapse seismic monitoring is a widely adopted approach, but its application for early leakage and migration detection beyond the designated storage complex remains challenging due to acquisition challenges, noise interference, and subtle signal changes. This work aims to develop a Deep Learning (DL) framework that could be both computationally cost-effective and efficient for early detection of CO_2 leakage during time-lapse seismic surveys. We explore the potential of using pre-trained Convolutional Neural Networks (CNNs), specifically ResNet152 (He et al. 2016), originally trained on the ImageNet dataset (Deng et al. 2009). Using a pre-trained model, one can avoid the training phase from scratch, which is a computationally expensive and data-intensive task. Instead, we leverage the deep representations and weights learned by the model from a different dataset and task, adapting them to identify changes in time-lapse seismic data, and to provide an initial quantitative estimation of CO_2 leakage.

We adapted ResNet152 to extract meaningful features from seismic images and to monitor the evolution of seismic data over time. To test the proposed framework, we consider hypothetical well leakage scenarios for a candidate CCS site near Kimberlina, CA, USA. We used a dataset based on 3D numerical multi-phase flow simulations (Buscheck et al. 2017; Yang et al. 2019). These simulations model CO_2 and brine leakage from a legacy well into overlying shallow aquifers. The model domain includes three overlying aquifer layers, while excluding the CO_2 injection well, storage reservoir, and primary seal (see Yang et al. (2019) for further details of the Kimberlina model). The dataset comprises 1000 simulations, which each contain a unique aquifer heterogeneity, aquifer and caprock permeability. The output time steps are every 10 years between 0-200 years. The simulation outputs considered in this work comprise CO_2 saturation maps and CO_2 leakage mass for each time step. The saturation maps were converted into P-wave velocity models using rock physics models, providing the input for forward seismic modeling. The CO_2 leakage mass values were used as targets for training a regressor, enabling the framework to provide a final quantitative estimate of the detected leakage.

Deep Learning Framework

We designed a framework that integrates physics-based seismic modeling to simulate time-lapse seismic surveys, with feature extraction and anomaly detection techniques. Seismic data were generated through seismic acoustic forward modeling based on the P-wave velocity models constructed by Wang et al. (2021), for each simulation time step. A total of 21000 forward modeling was performed using a cross-well survey configuration with 3 sources and 50 receivers for each timestep. Gaussian noise at different signal-to-noise ratios (SNR = 30, 20, 5) was added to simulate realistic acquisition conditions, and evaluate the impact of seismic noise on regressor prediction performance.

A Siamese neural network architecture (Chicco 2020) was implemented using a pre-trained ResNet152 as the feature extractor (Fig. 1). The original ResNet152, trained for image classification on natural images, was modified by removing its final fully connected classification layer. The output of the last convolutional block, a 2048-dimensional feature vector, was used as the embedding representation of the input seismic data.

Each input consists of three shot gather images combined into a single 3-channel image. The model processes pairs of time-lapse seismic data, a baseline survey and a monitor survey, to produce embeddings that capture key features of each input. The Siamese setup consists of two identical pre-trained ResNet152 models, with the same configuration, parameters, and weights, operating in parallel. This setup ensures consistent and reliable comparison of the resulting embeddings. From the resulting embeddings, we compute the Euclidean distance to quantify the degree of change between the baseline and monitor images. These high-dimensional embeddings are then reduced to a two-dimensional latent space using Uniform Manifold Approximation and Projection (UMAP; McInnes et al. 2018), enabling cluster visualization and anomaly detection. Finally, a multilayer perceptron (MLP) regressor is trained to predict the leakage mass, a physically interpretable parameters derived from flow simulations. This final step is crucial, as it can also provide quantitative information for decision-making in seismic monitoring.

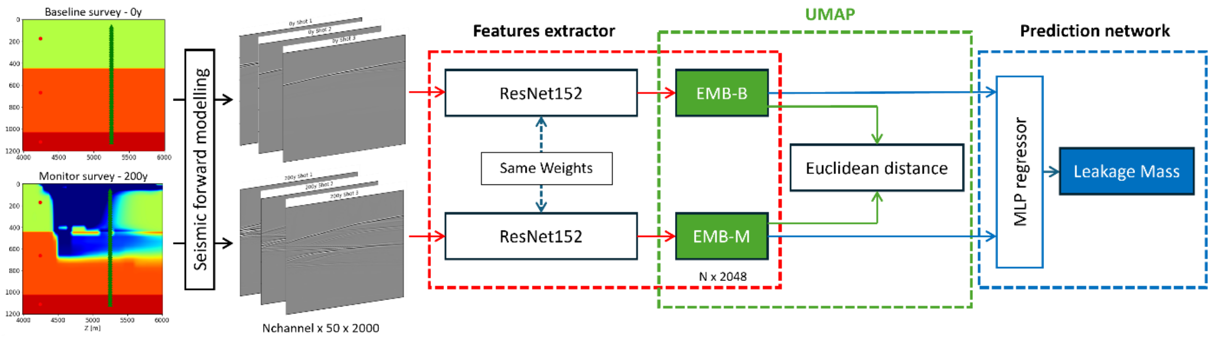

Figure 1 *Overview of the proposed DL framework for time-lapse seismic monitoring. From left: examples of P-wave velocity models from baseline (0 years) and monitor (200 years) surveys used as input to forward modeling to generate synthetic seismic shot gathers. These are processed by a Siamese ResNet152 to extract embeddings (EMB-B and EMB-M, with dimensions batch size × 2048); then they are compared via Euclidean distance and projected into a 2D latent space using UMAP; the Multilayer Perceptron (MLP) regressor predicts the CO_2 leakage mass from the embeddings.*

Early leakage detection and leakage mass prediction

The UMAP space was constructed using embeddings from both baseline and monitor surveys. These projections provide insight into how the Siamese ResNet152 is able to extract deep and meaningful features of the input data in relation to CO_2 leakage progression. In the training set (Fig. 2 left), embeddings from baseline surveys form a well-defined cluster, while monitor embeddings, colored by survey time, are progressively distributed along a certain trajectory. The size of each marker reflects the Euclidean distance between the monitor embedding and its corresponding baseline embedding, larger markers indicate greater deviation. This distribution aligns with simulation time, suggesting that the model captures the temporal evolution of leakage scenarios, particularly at early timesteps such as 10 and 20 years after injection (t10 and t20 in Fig. 2). In contrast, later timesteps become increasingly dispersed in the latent space, indicating greater variability and change as leakage progresses. In the test set (Fig. 2 right), we projected the unseen embeddings into the same UMAP latent space. Embeddings from monitors without leakage remain close to the baseline cluster, while those from leakage scenarios project into a separate region, forming a structure similar to that observed for the training set. This demonstrates the ability of the model to generalize and the stability of the embedding space, although ResNet152 was originally trained for classification task on non-seismic images. These results indicate that the extracted embeddings are meaningful and discriminative for the task of leakage detection.

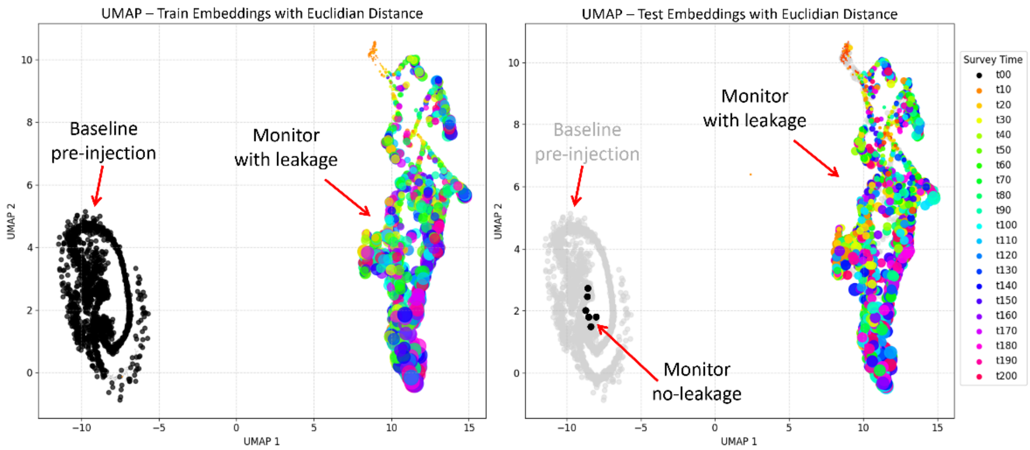

***Figure 2** UMAP projections of train (left) and test (right) embeddings. Baseline and no-leakage monitors cluster together, while leakage scenarios form a distinct cluster, highlighting the ability of the model to separate leakage and no-leakage surveys.*

As the noise level increases (i.e., lower SNR), the global structure of the UMAP embeddings becomes progressively less defined and harder to interpret (Fig. 3a–c). At high SNR (30), the latent space maintains a coherent structure, with visible separation between different surveys. However, at SNR = 20, the clustering begins to degrade, and by SNR = 5, the separation between leakage and no-leakage scenarios becomes indiscernible. This degradation in the latent space could affect downstream tasks, such as regression or classification. In the final stage of our framework, the embeddings are used as input to a regression model that predicts leakage mass for each time-lapse survey analyzed. The predicted versus true values are shown in Fig. 3d–f. At SNR = 30, the regressor achieves good accuracy, with a coefficient of determination $R^2 = 0.77$, with predicted leakage mass closely matching ground truth values. At SNR = 20, the model maintains good predictive performance ($R^2 = 0.61$), particularly for lower leakage mass values, although it tends to underestimate larger leakage events. At SNR = 5, the correlation between predicted and true leakage mass is strongly reduced ($R^2 = 0.08$). This highlights the reduced ability of the model to generalize under high-noise conditions and underscores the importance of robust denoising and preprocessing techniques when applying the framework to real field data.

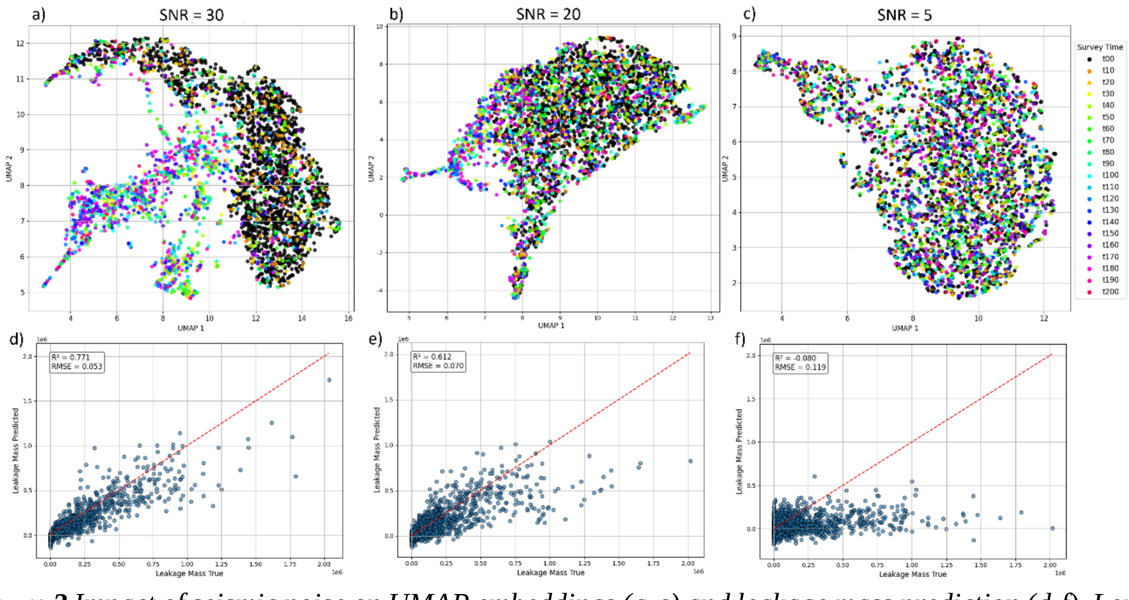

***Figure 3** Impact of seismic noise on UMAP embeddings (a-c) and leakage mass prediction (d-f). Lower SNR leads to less structured embeddings and reduced regression performance.*

Conclusions

This study demonstrates that a pre-trained ResNet152, originally developed for natural image classification, can successfully extract meaningful features from synthetic time-lapse seismic shot gathers. The model is capable of capturing subtle variations in the seismic response associated with CO_2 leakage, particularly in the early stages of migration. Although not trained on geophysical data, the extracted embeddings proved effective in detecting time-lapse seismic changes. The use of UMAP for dimensionality reduction offered a powerful tool for visual interpretation of the latent space, clearly separating baseline, no-leakage, and leakage scenarios. The embedding structure also reflected temporal progression, with early timesteps appearing more compact and later timesteps becoming increasingly dispersed as leakage advanced. In addition to qualitative analysis, the MLP regressor demonstrated that the embeddings extracted by the ResNet152 model are informative to support quantitative estimation of CO_2 leakage mass. The MLP regressor, trained on these embeddings, maintains predictive capability in the presence of moderate levels of noise, indicating that the learned representations capture relevant features of the seismic data despite the absence of domain-specific training.
The proposed approach is generalizable and can be extended to other pre-trained CNNs. Variants of ResNet (e.g., ResNet50, ResNet101) or alternative architectures such as DenseNet, EfficientNet, or Vision Transformers trained on natural image datasets could be tested to further improve feature extraction performance, embedding robustness, and computational efficiency. Overall, the proposed DL framework offers a promising and computationally efficient approach for early leakage detection in time-lapse seismic monitoring. Although it may not replace conventional monitoring systems, it could have strong potential as a complementary early-warning or decision-support tool to flag anomalies and guide targeted surveys.

References

Buscheck, T.A., Mansoor, K., Yang, X., Wainwright, H., Carroll, S. [2019]. Downhole pressure and chemical monitoring for CO2 and brine leak detection in aquifers above a CO2 storage reservoir. International Journal of Greenhouse Gas Control. **91**. 102812. 10.1016/j.ijggc.2019.102812.

Chicco, D. [2020]. Siamese neural networks: an overview. *Artificial Neural Networks*, Methods in Molecular Biology, **2190** (3rd ed.). Humana, New York, NY. Springer Protocols, Humana Press, pp. 73–94, 10.1007/978-1-0716-0826-5_3.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L. [2009]. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248–255).

He, K., Zhang, X., Ren, S., Sun, J. [2016]. Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–7778, https://doi.org/10.1109/cvpr.2016.90.

McInnes, L., Healy, J., Melville, J. [2018]. UMAP: Uniform Manifold Approximation and Projection for dimension reduction. arXiv preprint, arXiv:1802.03426.

Wang, Z., Dilmore, R., Harbert, W., Huang, L., Lin, Y., Feng, S, Wang, S. [2021]. Kimberlina 1.2 Velocity Models. US Department of Energy NETL. 10.18141/1832899.

Yang, X., Buscheck, T.A., Mansoor, K., Wang, Z., Gao, K., Huang, L., Appriou, D., Carroll, S.A. [2019]. Assessment of geophysical monitoring methods for detection of brine and CO2 leakage in drinking water aquifers. Int. J. Greenhouse Gas Control. **90**, 102803. https://doi.org/10.1016/j.ijggc.2019.102803.