

Derisking geologic carbon storage from high-resolution time-lapse seismic to explainable leakage detection

Ziyi Yin¹, Huseyin Tuna Erdinc¹, Abhinav Prakash Gahlot¹, Mathias Louboutin¹, and Felix J. Herrmann¹

<https://doi.org/10.1190/tle42010069.1>

Abstract

Geologic carbon storage represents one of the few truly scalable technologies capable of reducing the CO₂ concentration in the atmosphere. While this technology has the potential to scale, its success hinges on our ability to mitigate its risks. An important aspect of risk mitigation concerns assurances that the injected CO₂ remains within the storage complex. Among the different monitoring modalities, seismic imaging stands out due to its ability to attain high-resolution and high-fidelity images. However, these superior features come at prohibitive costs and time-intensive efforts that potentially render extensive seismic monitoring undesirable. To overcome this shortcoming, we present a methodology in which time-lapse images are created by inverting nonreplicated time-lapse monitoring data jointly. By no longer insisting on replication of the surveys to obtain high-fidelity time-lapse images and differences, extreme costs and time-consuming labor are averted. To demonstrate our approach, hundreds of realistic synthetic noisy time-lapse seismic data sets are simulated that contain imprints of regular CO₂ plumes and irregular plumes that leak. These time-lapse data sets are subsequently inverted to produce time-lapse difference images that are used to train a deep neural classifier. The testing results show that the classifier is capable of detecting CO₂ leakage automatically on unseen data with reasonable accuracy. We consider the use of this classifier as a first step in the development of an automatic workflow designed to handle the large number of continuously monitored CO₂ injection sites needed to help combat climate change.

Introduction

Seismic monitoring of geologic carbon storage (GCS) comes with its own unique challenges. Among these challenges, the need for low-cost highly repeatable, high-resolution, and high-fidelity images ranks chiefly. Densely sampled and replicated time-lapse surveys, which rely on permanent reservoir monitoring systems or on replicated streamer or node surveys, may be able to provide images conducive to interpretation and reservoir management, but these approaches are often too costly and require too much hand-holding to be of practical use for GCS at many injection sites.

To overcome these challenges, we replace the current paradigm of costly replicated acquisition, cumbersome time-lapse processing, and interpretation with a joint-inversion framework that maps time-lapse data to high-fidelity and high-resolution images from sparse nonreplicated time-lapse surveys. We demonstrate that we arrive at an imaging framework that is suitable for automatic detection of pressure-induced CO₂ leakage, which represents one

of the possible leakage scenarios. Rather than relying on meticulous 4D workflows in which baseline and monitoring surveys are processed separately to yield accurate and artifact-free time-lapse differences, our approach exposes information that is shared among the different vintages by formulating the imaging problem in terms of an unknown fictitious common component and innovations of the baseline and monitor surveys with respect to this common component. Because the common component is informed by all time-lapse surveys, its image quality improves when the surveys bring complementary information, which is the case when the surveys are not replicated. In turn, the enhanced common component results in improved images for the baseline survey, the monitor survey(s), and their time-lapse difference(s). Our joint wave-equation-based imaging formulation is versatile and capable of accounting for real data time-lapse issues such as changes in the background velocity model, calibration errors in shot and receiver locations (Oghenekohwo and Herrmann, 2017), and noise (Tian et al., 2018; Wei et al., 2018). The same applies to corrections for the source signature using on-the-fly source estimations (Yang et al., 2020; Z. Yin et al., 2020). We acknowledge, however, that the robustness of our method to such real data issues needs to be validated and is a topic for future study.

To showcase the achievable imaging gains and how these can be used in a GCS setting where CO₂ leakage is of major consideration, we create hundreds of time-lapse imaging experiments involving CO₂ plumes whose behavior is determined by the two-phase flow equations. To mimic irregular flow due to pressure-induced opening of fractures, we increase the permeability in the seal at random locations and pressure thresholds. The resulting flow simulations are used to generate time-lapse data sets that serve as input to our joint imaging scheme. The produced time-lapse difference images are subsequently used to train and test a neural network that as an explainable classifier determines whether the CO₂ plume behaves regularly or shows signs of leakage.

Our contributions are organized as follows. First, we discuss the time-lapse seismic imaging problem and its practical difficulties. Next, we introduce the joint recovery model (JRM) that takes explicit advantage of information shared by multiple surveys. By means of a carefully designed synthetic case study involving saline aquifers made of Bunter Sandstone in the southern North Sea, we demonstrate the uplift of the JRM and how its images can be used to train a deep neural network classifier to detect erroneous growth of the CO₂ plume automatically. Aside from determining whether the CO₂ plume behaves regularly or not, our network also provides class activation mappings (CAMs) that visualize areas in the image on which the network is basing its classification.

¹Georgia Institute of Technology, Atlanta, Georgia, USA. E-mail: ziyi.yin@gatech.edu; herdinc3@gatech.edu; agahlot8@gatech.edu; mlouboutin3@gatech.edu; felix.herrmann@gatech.edu.

Seismic monitoring with time-lapse imaging

To keep track of CO₂ plume development during GCS projects, multiple time-lapse surveys are collected. Baseline surveys are acquired before the supercritical CO₂ is injected into the reservoir. These baseline surveys, denoted by the index $j = 1$, are followed by one or more monitor surveys, collected at later times and indexed by $j = 2, \dots, n_v$ with n_v the total number of surveys.

Seismic monitoring of GCS brings its own unique set of challenges that stem from the fact that its main concern is (early) detection of possible CO₂ leakage from the storage complex. To be successful in this task, monitoring GCS calls for a time-lapse imaging modality that is capable of

- detecting weak time-lapse signals associated with small rock-physics changes induced by CO₂ leakage;
- attaining high lateral resolution from active-source surface seismic data to detect vertically moving leakage;
- handling an increasing number of not perfectly calibrated seismic surveys collected over long periods of time (approximately 50 to 100 years);
- reducing costs drastically by no longer insisting on replication of time-lapse surveys to attain high degrees of repeatability; and
- lowering the cumulative environmental imprint of active-source acquisition.

Monitoring with the JRM. To meet these challenges, we choose a linear imaging framework where observed linearized data for each vintage are related to perturbations in the acoustic impedance via

$$\mathbf{b}_j = \mathbf{A}_j \mathbf{x}_j + \mathbf{e}_j \quad \text{for } j = 1, 2, \dots, n_v. \quad (1)$$

In this expression, the matrix $\mathbf{A}_j$ stands for the linearized Born scattering operator for the j^{th} vintage. Observed linearized data, collected for all shots in the vector $\mathbf{b}_j$, are generated by applying the $\mathbf{A}_j$'s to the (unknown) impedance perturbations denoted by $\mathbf{x}_j$ for $j = 1, 2, \dots, n_v$ with the noise term $\mathbf{e}_j$. The task of time-lapse imaging is to create high-resolution, high-fidelity, true-amplitude estimates for the time-lapse images, $\{\mathbf{x}_j\}_{j=1}^{n_v}$, from nonreplicated sparsely sampled noisy time-lapse data.

We argue that our choice for linearized imaging is justified for four reasons. First, CO₂ injection sites undergo baseline studies involving vintage data and possible follow-up surveys, which means that accurate information on the background velocity model is generally available. Second, changes in the acoustic parameters induced by CO₂ injection are typically small, so it suffices to work with one and the same background model for the baseline and monitor surveys. Third, when the background model is sufficiently close to the true model, linearized inversion, which corresponds to a single Gauss-Newton iteration of full-waveform inversion, converges quadratically. Fourth, because the forward model is linear, it is conducive to the use of the JRM where inversions are carried out with respect to the common component, which is shared between all vintages, and innovations with respect to the common component. Because the common component represents an average, we expect this joint imaging method to be relatively

robust with respect to kinematic changes induced by time-lapse effects or by lack of calibration of the acquisition (Oghenekohwo and Herrmann, 2017).

By parameterizing time-lapse images, $\{\mathbf{x}_j\}_{j=1}^{n_v}$, in terms of the common component, $\mathbf{z}_0$, and innovations with respect to the common component, $\{\mathbf{z}_j\}_{j=1}^{n_v}$, we arrive at the JRM where representations for the images are given by

$$\mathbf{x}_j = \frac{1}{\gamma} \mathbf{z}_0 + \mathbf{z}_j \quad \text{for } j = 1, 2, \dots, n_v. \quad (2)$$

Here, the parameter γ controls the balance between the common component, $\mathbf{z}_0$, and innovation components, $\{\mathbf{z}_j\}_{j=1}^{n_v}$ (X. Li, 2015). Compared to traditional time-lapse approaches, where data are imaged separately or where time-lapse surveys are subtracted, inversions for time-lapse images based on the above parameterization are carried out jointly and involve inverting the following matrix:

$$\mathbf{A} = \begin{bmatrix} \frac{1}{\gamma} \mathbf{A}_1 & \mathbf{A}_1 & & \\ \vdots & & \ddots & \\ \frac{1}{\gamma} \mathbf{A}_{n_v} & & & \mathbf{A}_{n_v} \end{bmatrix}. \quad (3)$$

While traditional time-lapse imaging approaches strive toward maximal replication between surveys to suppress acquisition-related artifacts, imaging with the JRM, which entails inverting the underdetermined system in equation 3 using structure-promotion techniques (e.g., via ℓ_1 -norm minimization), improves the image quality of the vintages themselves in situations where the surveys are not replicated. This occurs in cases where $\mathbf{A}_i \neq \mathbf{A}_j$ for $\forall i \neq j$, or in situations where there is significant noise. This remarkable result was shown to hold for sparsity-promoting denoising of time-lapse field data (Tian et al., 2018; Wei et al., 2018), for various wavefield reconstructions of randomized simultaneous-source dynamic (towed-array) and static (ocean-bottom cable/node) marine acquisitions (Oghenekohwo and Herrmann, 2017; Oghenekohwo et al., 2017; Wason et al., 2017), and for wave-based inversion, including least-squares reverse time migration (RTM) and full-waveform inversion (Herrmann and Oghenekohwo, 2017; Oghenekohwo, 2017). The observed quality gains in these applications can be explained by improvements in the common component resulting from complementary information residing in nonreplicated time-lapse surveys. This enhanced recovery of the common component in turn improves the recovery of the innovations and therefore the vintages themselves. The time-lapse differences themselves also improve, or at the very least, remain relatively unaffected when the surveys are not replicated. Relaxing replication of surveys obviously leads to reduction in cost and environmental impact. Below, we show how GCS monitoring also benefits from this approach.

Monitoring with curvelet-domain structure promotion. To obtain high-resolution and high-fidelity time-lapse images, we

invert the system in equation 3 (Witte et al., 2019b; M. Yang et al., 2020; Z. Yin et al., 2021) with

$$\begin{aligned} \underset{\mathbf{z}}{\text{minimize}} \quad & \lambda \|\mathbf{Cz}\|_1 + \frac{1}{2} \|\mathbf{Cz}\|_2^2 \\ \text{subject to} \quad & \|\mathbf{b} - \mathbf{Az}\|_2 \leq \sigma, \end{aligned} \quad (4)$$

where $\mathbf{C}$ is the forward curvelet transform, λ is the threshold parameter, and σ is the magnitude of the noise. At iteration k and for $\sigma = 0$, solving equation 4 corresponds to computing the following iterations:

$$\begin{aligned} \mathbf{u}_{k+1} &= \mathbf{u}_k - t_k \mathbf{A}_k^\top (\mathbf{A}_k \mathbf{z}_k - \mathbf{b}_k) \\ \mathbf{z}_{k+1} &= \mathbf{C}^\top \mathcal{S}_\lambda (\mathbf{C} \mathbf{u}_{k+1}), \end{aligned} \quad (5)$$

where $\mathbf{A}_k$, with a slight abuse of notation, represents the matrix in equation 5 for a subset of shots randomly selected from sources in each survey. The vector $\mathbf{b}_k$ contains the extracted shot records from $\mathbf{b}$ and the symbol $\top$ refers to the adjoint. The dynamic steplength t_k is given by $t_k = \|\mathbf{A}_k \mathbf{z}_k - \mathbf{b}_k\|_2^2 / \|\mathbf{A}_k^\top (\mathbf{A}_k \mathbf{z}_k - \mathbf{b}_k)\|_2^2$ (Lorenz et al., 2014). Sparsity is promoted via curvelet-domain soft thresholding, $\mathcal{S}_\lambda(\cdot) = \max(|\cdot| - \lambda, 0) \text{sign}(\cdot)$, where λ is the threshold. The vector $\mathbf{z}_k$ contains the baseline and innovation components.

Numerical case study: Bunter Sandstone in the southern North Sea

Before discussing the impact of high-resolution and high-fidelity time-lapse imaging with the JRM on the downstream

task of automatic leakage detection with a neural network classifier, we first detail the setup of our numerical experiments using techniques from simulation-based acquisition design as described by Z. Yin et al. (2021). To generate realistic time-lapse data and training sets for the automatic leakage classifier, we follow the workflow summarized in Figure 1. In this approach, use is made of proxy models for seismic properties derived from real 3D-imaged seismic and well data (Jones et al., 2012). With rock physics, these seismic models are converted to fluid-flow models that serve as input to two-phase flow simulations. The resulting data sets, which include pressure-induced leakage, will be used to create time-lapse data used to train our classifier. For more detail, refer to the caption of Figure 1.

Proxy seismic and fluid-flow models. Among the various CO₂ injection projects, GCS in offshore saline aquifers has been most successful in reaching scale and in meeting injection targets (Michael et al., 2010). For that reason, we consider a proxy model derived from real 3D-imaged seismic and well data (Jones et al., 2012) and representative for CO₂ injection in the south of the North Sea involving a saline aquifer made of the highly permeable Bunter Sandstone. This area, which is actively being considered for GCS (Kolster et al., 2018), consists of the following three geologic sections (see Figure 2 for the permeability and porosity distribution):

- 1) the highly porous (average 33%) and permeable (more than 170 mD) Bunter Sandstone reservoir of about 300–500 m thick. This section, denoted by red colors in Figure 2, corresponds to the saline aquifer and serves as the reservoir for CO₂ injection;




Figure 1. Simulation-based monitoring design framework. (a) Starting with a proxy model for the wavespeed and density, the workflow proceeds by converting these seismic properties into (b) permeability and porosity. (c) These fluid-flow properties are used to simulate CO₂ plumes that behave regularly or exhibit leakage outside the storage complex. Induced changes by the CO₂ plume for the wavespeed and density are depicted in (d) and serve as input to simulations of time-lapse seismic data (S/N 8.0 dB) and shot-domain time-lapse differences (S/N –31.4 dB). Imaging results for regular and irregular plume developments are plotted in (f) and serve as input to (g) the deep neural classifier, which determines whether the flow behaves regularly or leaks. Activation mappings in (h) show regions on which the network is basing its classification. As expected, the activation mapping is diffusive in case of regular CO₂ plume development and focused on the leakage location when CO₂ plume behaves irregularly.




Figure 2. (a) Permeability and (b) porosity derived from a 2D slice of the Compass model.

- 2) the primary seal (permeability 10^{-4} – 10^{-2} mD) made of the Rot Halite Member, which is 50 m thick and continuous (black layer in Figure 2); and
- 3) the secondary seal made of the Haisborough group, which is more than 300 m thick and consists of low-permeability (15–18 mD) mudstone (purple section in Figure 2).

To arrive at the fluid-flow models, we consider 2D subsets of the 3D Compass model (Jones et al., 2012) and convert these seismic models to fluid-flow properties (see Figure 1b) by assuming a linear relationship between compressional wavespeed and permeability in each stratigraphic section. For further details on the conversion of compressional wavespeed and density to permeability and porosity, we refer to empirical relationships reported in Klimentos (1991). During conversion, an increase of 1 km/s in compressional wavespeed is assumed to correspond to an increase of 1.63 mD in permeability. From this, porosity is calculated with the Kozeny-Carman equation (Costa, 2006)

$$K = \phi^3 \left(\frac{1.527}{0.0314 * (1 - \phi)} \right)^2, \text{ where } K \text{ and } \phi \text{ denote permeability}$$

(mD) and porosity (%) with constants taken from the Strategic UK CCS Storage Appraisal Project report.

Fluid-flow simulations. To model CO_2 plumes that behave regularly and irregularly, the latter due to leakage, we solve the two-phase flow equations numerically² for both pressure and concentration (D. Li and Xu, 2021; D. Li et al., 2020). To mimic possible pressure-induced CO_2 leakage, we increase the permeability at random distances away from the injection well within

the seal from 10^{-4} to 500 mD when the pressure exceeds approximately 15 MPa. At that depth, the pressure is below the fracture gradient (Ringrose, 2020). Because pressure-induced fractures come in different sizes, we also randomly vary the width of the pressure-induced fracture openings from 12.5 to 62.5 m. Examples of fluid-flow simulations without and with leakage are shown in Figure 1c.

Rock-physics conversion. To monitor temporal variations in the plume's CO_2 concentration seismically, we use the patchy saturation model (Avseth et al., 2010) to convert the CO_2 concentration to decrease in compressional wavespeed and density. These changes are shown in Figure 1d. The fact that these induced changes in the time-lapse differences in seismic properties are relatively small in spatial extent (approximately 800 m for the plume and less than 62.5 m for the leakage) and amplitude (1.68% time-lapse change in the acoustic impedance) calls for a time-lapse imaging modality with small normalized root-mean-square (Nrms) values (Kragh and Christie, 2002).

Time-lapse seismic simulations. Training and validating automatic detection of CO_2 leakage from the storage complex requires the creation of realistic synthetic time-lapse data sets that contain the seismic imprint of regular as well as irregular (leakage) plume development. To this end, baseline surveys are simulated prior to CO_2 injection for different subsets of the Compass model. Monitor surveys are simulated 200 days after leakage occurs to verify that potential leakage can be detected automatically early on. For regular plume development, we shoot monitor surveys for each subset at random times after CO_2 injection. To strike a balance between acquisition productivity and time-lapse image quality, use is made of dense semipermanent acoustic monitoring at the seafloor with 25 m receiver spacing. Contrary to expensive permanent reservoir monitoring systems with multicomponent geophones, our system works with hydrophones connected to underwater buoys located 2 m above the ocean bottom. Aside from being relatively low cost, this system also avoids complications arising from elastic wave interactions at the seabed. Time-lapse acquisition costs are further reduced by nonreplicated coarse shooting with the source towed at 10 m below the ocean surface. Subsampling artifacts are reduced by using a randomized technique from compressive sensing where 32 sources are located at nonreplicated jittered (Herrmann and Hennenfent, 2008) source positions, yielding an average source sampling of 125 m. Given this acquisition geometry, linear data are generated³ with equation 1 for a 25 Hz Ricker wavelet and with the band-limited noise term set so that the data's signal-to-noise ratio (S/N) is 8.0 dB. This noise level leads to an extremely poor S/N of –31.4 dB for time-lapse differences in the shot domain. See Figure 1e.

Imaging with JRM versus RTM. Given the simulated time-lapse data sets with and without leakage, time-lapse difference images are created according to two different imaging scenarios, namely via independent RTM, conducted on the

²We used the open-source software FwiFlow.jl (D. Li et al., 2020; D. Li and Xu, 2021) to solve the two-phase flow equations for both the pressure and concentration.

³We used the open-source software JUDI.jl (Witte et al., 2019a; Louboutin et al., 2022) to model the wave propagation. This Julia package implements highly optimized propagators using Devito (Louboutin et al., 2019; Luporini et al., 2020, 2022).

baseline and monitor surveys separately, and via inversion of the JRM (cf. equations 3 and 4). To limit the computational cost of the Bregman iterations (equation 5), four shot records are selected per iteration at random from each survey for imaging (W. Yin et al., 2008; Witte et al., 2019b; Yang et al., 2020; Z. Yin et al., 2021). This limits the cost of the joint inversion to the equivalent of three RTMs. The recovered baseline images are shown in Figure 3a for RTM and Figure 3b for JRM. For the leakage scenario, the time-lapse differences are plotted in Figures 3c and 3d for RTM and JRM, respectively. For the regular plume, the time-lapse differences are plotted in Figures 3e and 3f for RTM and JRM, respectively. From these images, it is clear that joint inversion leads to relatively

artifact-free recovery of the vintages and time-lapse differences. This observation is reflected in the Nrms values, which improve considerably as shown by the histograms in Figure 4 for 1000 imaging experiments. Not only do the Nrms values shift toward the left, their values are also more concentrated when inverting time-lapse data with the JRM. Both features are beneficial to automatic leakage detection.

Deep neural network classifier for CO₂ leakage detection

The injection of supercritical CO₂ into the storage complex perturbs the physical, chemical, and thermal environment of the reservoir (Newell and Ilgen, 2019). Because CO₂ injection increases the pressure, this process may trigger CO₂ leakage across the seal




Figure 3. RTM versus inversion JRM. (a) RTM image of the baseline; (b) JRM image of the baseline; (c) time-lapse difference and CO₂ plume for independent RTM images with leakage; (d) time-lapse obtained by inverting the time-lapse data jointly with leakage; (e) time-lapse difference and CO₂ plume for independent RTM images without leakage; (f) time-lapse obtained by inverting the time-lapse data jointly without leakage. Notice improvement in the time-lapse image quality. This improvement is reflected in the Nrms values that decrease from 8.48% for RTM to 3.20% for JRM.




Figure 4. Nrms values for 1000 time-lapse experiments.




Figure 5. Confusion matrix for classifier trained on recovery images from JRM.

when the pressure increase induces opening of pre-existing faults or fractures zones (Pruess, 2006; Ringrose, 2020). To ensure safe operation of CO₂ storage, we develop a quantitative leakage detection tool based on a deep neural classifier. This classifier is trained on time-lapse images that contain the imprint of CO₂ plumes that behave regularly and irregularly. In case of irregular flow, CO₂ escapes the storage complex through a pressure-induced opening in the seal, which causes a localized increase in permeability (shown in Figure 3d).

Because time-lapse differences are small in amplitude and strongly localized laterally when leakage occurs, highly sensitive learned classifiers are needed. For this purpose, we follow Erdinc et al. (2022) and adopt the Vision Transformer (ViT) (Dosovitskiy et al., 2021). This state-of-the-art classifier originated from the field of natural language processing (Vaswani et al., 2017). Thanks to their attention mechanism, ViTs have been shown to achieve superior performance on image classification tasks where image patches are considered as word tokens by the transformer network. As a result, ViTs have much less image-specific inductive bias compared to convolutional neural networks (Dosovitskiy et al., 2021).

To arrive at a practical and performant ViT classifier, we start from a ViT that is pretrained on image tasks with 16×16 patches and apply transfer learning (Yosinski et al., 2014) to fine-tune this network on 1576 labeled time-lapse images. Catastrophic forgetting is avoided by freezing the initial layers, which are responsible for feature extraction, during the initial training. After the initial training of the last dense layers, all network weights are updated for several epochs while keeping the learning rate small. The labeled (regular versus irregular flow) training set itself consists of 1576 time-lapse data sets divided equally between regular and irregular flow.

After the training is completed, baseline and monitor surveys are simulated for 394 unseen earth models with regular and irregular plumes. These simulated time-lapse data sets are imaged with JRM by inverting the matrix in equation 3 via Bregman iterations in equation 5. The resulting time-lapse difference images (see Figures 3d and 3f for two examples) serve as input to the ViT classifier. Refer to Figure 5 for performance, which corresponds to a 2×2 confusion matrix. The first row denotes the classification results for samples with regular plume (negative samples), where 193 (true negative) out of 206 samples are classified correctly. The second row denotes the classification results for samples with CO₂ leakage over the seal (positive samples), where 147 (true positive) out of 188 samples are classified correctly. Because JRM recovers relatively artifact-free time-lapse differences, the classifier does not pick up too many artifacts related to finite acquisition as CO₂ leakage. This leads to much fewer false alarms for CO₂ leakage.

Class activation mapping-based saliency map

While our ViT classifier is capable of achieving good performance (see Figure 5), making intervention decisions during GCS projects calls for interpretability and trustworthiness of our classifier (Hooker et al., 2019; Mackowiak et al., 2021; Zhang et al., 2021). To enhance these features, we take advantage of CAMs (Zhou et al., 2016). These saliency maps help us identify the discriminative spatial regions in each image that support a particular class decision. In our application, these regions correspond to areas where the classifier deems the CO₂ plume to behave irregularly (if the classification result is leakage). By overlaying time-lapse difference images with these maps, interpretation is facilitated, assisting practitioners to make decisions on how to proceed with GCS projects and take associated actions. Figure 6 illustrates how the Score CAM approach (Wang et al., 2020) serves this purpose⁴. Figure 6a shows the CAM result for a time-lapse difference image classified as a CO₂ leakage (Figure 3d). Despite few artifacts around the image, the CAM clearly focuses on the CO₂ leakage over the seal, which could potentially alert the practitioners of GCS. When the plume is detected as growing regularly, the CAM result is diffusive (shown in Figure 6b). This shows that the classification decision is based on the entire image and not only at the plume area. The scripts to reproduce the experiments are available on the SLIM GitHub page: <https://github.com/slimgroup/GCS-CAM>.

⁴We used the open-source software PyTorch library for CAM methods (Gildenblat and contributors, 2021) to calculate the CAM images.

Discussion and conclusion

As a first step in the development of scalable automatic workflows for seismic monitoring of GCS, we propose a methodology for low-cost time-lapse imaging that exploits commonality between baseline and monitor surveys through the JRM. By means of carefully designed realistic synthetic time-lapse seismic experiments, we have shown that highly repeatable, high-resolution, and high-fidelity images are achievable without insisting on replication of the baseline and monitor surveys. Because our method relies on a joint-inversion methodology, it also averts labor-intensive 4D processing to compensate for less-than-ideal acquisitions. Aside from establishing our claim of relaxing the need for replication empirically, through hundreds of synthetic time-lapse experiments yielding significant improvements in time-lapse image quality and Nrms values, we also showed that a deep neural classifier can be trained to detect CO₂ leakage automatically. While the classification results are encouraging, false positives and negatives remain. We argue that these may be acceptable because decisions to intervene, e.g., to stop injection of CO₂, typically involve other complementary sources of information such as pressure drops at the wellhead. In future work, we plan to extend our methodology to different leakage scenarios and quantification of uncertainty. We also intend to further investigate robustness of the proposed joint imaging methodology with respect to calibration errors and variations in the source signature within and across different surveys. Finally, interpretability of the neural classifier's output and different leakage scenarios and their impact on the shape of the CO₂ plume also will be investigated further. ■■

Acknowledgments

We would like to thank Ramesh Neelamani and Bram Willemsen for the thorough reviews and valuable suggestions to our original manuscript. We would like to thank Charles Jones and Philipp A. Witte for constructive discussions. The carbon capture and storage project information is taken from the Strategic UK CCS Storage Appraisal Project, funded by the UK Department of Energy and Climate Change, commissioned by the Energy Technologies Institute (ETI), and delivered by Pale Blue Dot Energy, Axis Well Technology, and Costain. The information contains copyright information licensed under ETI Open License. This research was carried out with the support of Georgia Research Alliance and partners of the ML4Seismic Center.

Data and materials availability

Data associated with this research are available and can be accessed at <https://doi.org/10.5281/zenodo.7222318>.

Corresponding author: ziyi.yin@gatech.edu

References

Avseth, P., T. Mukerji, and G. Mavko, 2010, Quantitative seismic interpretation: Applying rock physics tools to reduce interpretation risk: Cambridge University Press, <https://doi.org/10.1017/CBO9780511600074>.




Figure 6. CAM for time-lapse difference images with a leaking plume and with a regular plume.

- 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**, no. 2, L02318, <https://doi.org/10.1029/2005GL025134>.
- Dosovitskiy, A., L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, et al., 2021, An image is worth 16 × 16 words: Transformers for image recognition at scale: Presented at International Conference on Learning Representations, <https://openreview.net/forum?id=YicbFdNTTy>, accessed 29 November 2022.
- Erdinc, H. T., A. P. Gahlot, Z. Yin, M. Louboutin, and F. J. Herrmann, 2022, De-risking carbon capture and sequestration with explainable CO₂ leakage detection in time-lapse seismic monitoring images: Association for the Advancement of Artificial Intelligence Fall Symposium, <https://slim.gatech.edu/Publications/Public/Conferences/AAAI/2022/erdinc2022AAAIddc/erdinc2022AAAIddc.pdf>, accessed 29 November 2022.
- Gildenblat, J., and contributors, 2021, PyTorch library for CAM methods: GitHub, <https://github.com/jacobgil/pytorch-grad-cam>.
- Herrmann, F. J., and G. Hennenfent, 2008, Non-parametric seismic data recovery with curvelet frames: *Geophysical Journal International*, **173**, no. 1, 233–248, <https://doi.org/10.1111/j.1365-246X.2007.03698.x>.
- Herrmann, F. J., and F. Oghenekohwo, 2017, Improved time-lapse data repeatability with randomized sampling and distributed compressive sensing: 79th Conference and Exhibition, EAGE, Extended Abstracts, <https://doi.org/10.3997/2214-4609.201701389>.
- Hooker, S., D. Erhan, P.-J. Kindermans, and B. Kim, 2019, A benchmark for interpretability methods in deep neural networks: Proceedings of the 33rd International Conference on Neural Information Processing Systems, 9737–9748.
- Jones, C. E., J. A. Edgar, J. I. Selva, and H. Crook, 2012, Building complex synthetic models to evaluate acquisition geometries and velocity inversion technologies: 74th Conference and Exhibition, EAGE, Extended Abstracts, <https://doi.org/10.3997/2214-4609.20148575>.

- Klimentos, T., 1991, The effects of porosity-permeability-clay content on the velocity of compressional waves: *Geophysics*, **56**, no. 12, 1930–1939, <https://doi.org/10.1190/1.1443004>.
- Kolster, C., S. Agada, N. Mac Dowell, and S. Krevor, 2018, The impact of time-varying CO₂ injection rate on large scale storage in the UK Bunter Sandstone: *International Journal of Greenhouse Gas Control*, **68**, 77–85, <https://doi.org/10.1016/j.ijggc.2017.10.011>.
- Kragh, E., and P. Christie, 2002, Seismic repeatability, normalized rms, and predictability: *The Leading Edge*, **21**, no. 7, 640–647, <https://doi.org/10.1190/1.1497316>.
- Li, D., and K. Xu, 2021, FwiFlow.jl v0.3.1: Zenodo, <https://doi.org/10.5281/zenodo.5528428>.
- Li, D., K. Xu, J. M. Harris, and E. Darve, 2020, Coupled time-lapse full-waveform inversion for subsurface flow problems using intrusive automatic differentiation: *Water Resources Research*, **56**, no. 8, e2019WR027032, <https://doi.org/10.1029/2019WR027032>.
- Li, X., 2015, A weighted ℓ_1 -minimization for distributed compressive sensing: PhD thesis, University of British Columbia.
- Lorenz, D. A., F. Schopfer, and S. Wenger, 2014, The linearized Bregman method via split feasibility problems: Analysis and generalizations: *SIAM Journal on Imaging Sciences*, **7**, no. 2, 1237–1262, <https://doi.org/10.1137/130936269>.
- Louboutin, M., M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Veslesko, and G. J. Gorman, 2019, Devito (v3.1.0): An embedded domain-specific language for finite differences and geophysical exploration: *Geoscientific Model Development*, **12**, 1165–1187, <https://doi.org/10.5194/gmd-12-1165-2019>.
- Louboutin, M., P. Witte, Z. Yin, H. Modzelewski, C. da Costa, Kerim, and P. Nogueira, 2022, JUDI.jl v3.1.9: Zenodo, <https://doi.org/10.5281/zenodo.7086719>.
- Luporini, F., M. Louboutin, M. Lange, N. Kukreja, rhodrin, G. Bisbas, V. Pandolfo, et al., 2022, Devito v4.7.1: Zenodo, <https://doi.org/10.5281/zenodo.6958070>.
- Luporini, F., M. Louboutin, M. Lange, N. Kukreja, P. Witte, J. Hückelheim, C. Yount, P. H. J. Kelly, F. J. Herrmann, and G. J. Gorman, 2020, Architecture and performance of Devito, a system for automated stencil computation: *ACM Transactions on Mathematical Software*, **46**, no. 1, 1–28, <https://doi.org/10.1145/3374916>.
- Mackowiak, R., L. Ardizzone, U. Kothe, and C. Rother, 2021, Generative classifiers as a basis for trustworthy image classification: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2971–2981, <https://doi.org/10.1109/CVPR46437.2021.00299>.
- Michael, K., A. Golab, V. Shulakova, J. Ennis-King, G. Allinson, S. Sharma, and T. Aiken, 2010, Geological storage of CO₂ in saline aquifers — A review of the experience from existing storage operations: *International Journal of Greenhouse Gas Control*, **4**, no. 4, 659–667, <https://doi.org/10.1016/j.ijggc.2009.12.011>.
- Newell, P., and A. G. Ilgen, 2019, Overview of geological carbon storage (GCS), in P. Newell and A. G. Ilgen, eds., *Science of carbon storage in deep saline formations*: Elsevier, 1–13, <https://doi.org/10.1016/B978-0-12-812752-0.00001-0>.
- Oghenekohwo, F. O., 2017, Economic time-lapse seismic acquisition and imaging — Reaping the benefits of randomized sampling with distributed compressive sensing: PhD thesis, University of British Columbia.
- Oghenekohwo, F., and F. J. Herrmann, 2017, Highly repeatable time-lapse seismic with distributed compressive sensing — Mitigating effects of calibration errors: *The Leading Edge*, **36**, no. 8, 688–694, <https://doi.org/10.1190/tle36080688.1>.
- Oghenekohwo, F., H. Wason, E. Esser, and F. J. Herrmann, 2017, Low-cost time-lapse seismic with distributed compressive sensing — Part 1: Exploiting common information among the vintages: *Geophysics*, **82**, no. 3, P1–P13, <https://doi.org/10.1190/geo2016-0076.1>.
- Pruess, K., 2006, On CO₂ behavior in the subsurface, following leakage from a geologic storage reservoir: Lawrence Berkeley National Laboratory.
- Ringrose, P., 2020, How to store CO₂ underground: Insights from early-mover CCS projects: Springer.
- Tian, Y., L. Wei, C. Li, S. Oppert, and G. Hennenfent, 2018, Joint sparsity recovery for noise attenuation: 88th Annual International Meeting, SEG, Expanded Abstracts, 4186–4190, <https://doi.org/10.1190/segam2018-2996474.1>.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, 2017, Attention is all you need: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 6000–6010.
- Wang, H., Z. Wang, M. Du, F. Yang, Z. Zhang, S. Ding, P. Mardziel, and X. Hu, 2020, Score-CAM: Score-weighted visual explanations for convolutional neural networks: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, <https://doi.org/10.1109/CVPRW50498.2020.00020>.
- Wason, H., F. Oghenekohwo, and F. J. Herrmann, 2017, Low-cost time-lapse seismic with distributed compressive sensing — Part 2: Impact on repeatability: *Geophysics*, **82**, no. 3, P15–P30, <https://doi.org/10.1190/geo2016-0252.1>.
- Wei, L., Y. Tian, C. Li, S. Oppert, and G. Hennenfent, 2018, Improve 4D seismic interpretability with joint sparsity recovery: 88th Annual International Meeting, SEG, Expanded Abstracts, 5338–5342, <https://doi.org/10.1190/segam2018-2975383.1>.
- Witte, P. A., M. Louboutin, N. Kukreja, F. Luporini, M. Lange, G. J. Gorman, and F. J. Herrmann, 2019a, A large-scale framework for symbolic implementations of seismic inversion algorithms in Julia: *Geophysics*, **84**, no. 3, F57–F71, <https://doi.org/10.1190/geo2018-0174.1>.
- Witte, P. A., M. Louboutin, F. Luporini, G. J. Gorman, and F. J. Herrmann, 2019b, Compressive least-squares migration with on-the-fly Fourier transforms: *Geophysics*, **84**, no. 5, R655–R672, <https://doi.org/10.1190/geo2018-0490.1>.
- Yang, M., Z. Fang, P. Witte, and F. J. Herrmann, 2020, Time-domain sparsity promoting least-squares reverse time migration with source estimation: *Geophysical Prospecting*, **68**, no. 9, 2697–2711, <https://doi.org/10.1111/1365-2478.13021>.
- Yin, W., S. Osher, D. Goldfarb, and J. Darbon, J., 2008, Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing: *SIAM Journal on Imaging Sciences*, **1**, no. 1, 143–168.
- Yin, Z., M. Louboutin, and F. J. Herrmann, 2021, Compressive time-lapse seismic monitoring of carbon storage and sequestration with the joint recovery model: First International Meeting for Applied Geoscience & Energy, SEG/AAPG, Expanded Abstracts, 3434–3438, <https://doi.org/10.1190/segam2021-3569087.1>.
- Yin, Z., R. Orozco, P. Witte, M. Louboutin, G. Rizzuti, and F. J. Herrmann, 2020, Extended source imaging — A unifying framework for seismic and medical imaging: 90th Annual International Meeting, SEG, Expanded Abstracts, 3502–3506, <https://doi.org/10.1190/segam2020-3426999.1>.
- Yosinski, J., J. Clune, Y. Bengio, and H. Lipson, 2014, How transferable are features in deep neural networks?: *Proceedings of the 27th International Conference on Neural Information Processing Systems*, 2, 3320–3328.
- Zhang, Y., P. Tiño, A. Leonardis, and K. Tang, 2021, A survey on neural network interpretability: *IEEE Transactions on Emerging Topics in Computational Intelligence*, **5**, no. 5, 726–742, <https://doi.org/10.1109/TETCI.2021.3100641>.
- Zhou, B., A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, 2016, Learning deep features for discriminative localization: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2921–2929, <https://doi.org/10.1109/CVPR.2016.319>.