




OPEN

## Physics-informed machine learning with differentiable programming for heterogeneous underground reservoir pressure management

Aleksandra Pachalieva<sup>1,2</sup>, Daniel O'Malley<sup>2</sup>, Dylan Robert Harp<sup>2</sup> & Hari Viswanathan<sup>2</sup>

Avoiding over-pressurization in subsurface reservoirs is critical for applications like CO<sub>2</sub> sequestration and wastewater injection. Managing the pressures by controlling injection/extraction are challenging because of complex heterogeneity in the subsurface. The heterogeneity typically requires high-fidelity physics-based models to make predictions on CO<sub>2</sub> fate. Furthermore, characterizing the heterogeneity accurately is fraught with parametric uncertainty. Accounting for both, heterogeneity and uncertainty, makes this a computationally-intensive problem challenging for current reservoir simulators. To tackle this, we use differentiable programming with a full-physics model and machine learning to determine the fluid extraction rates that prevent over-pressurization at critical reservoir locations. We use DPFEHM framework, which has trustworthy physics based on the standard two-point flux finite volume discretization and is also automatically differentiable like machine learning models. Our physics-informed machine learning framework uses convolutional neural networks to learn an appropriate extraction rate based on the permeability field. We also perform a hyperparameter search to improve the model's accuracy. Training and testing scenarios are executed to evaluate the feasibility of using physics-informed machine learning to manage reservoir pressures. We constructed and tested a sufficiently accurate simulator that is 400 000 times faster than the underlying physics-based simulator, allowing for near real-time analysis and robust uncertainty quantification.

Reservoir pressure management is essential for injection/extraction operations in the subsurface for resource extraction, carbon sequestration, climate change mitigation, and renewable energy. However, when mishandled, the pressure management strategy can cause induced seismicity, examples of which are the seismic events due to large-scale wastewater re-injection in central Oklahoma<sup>1-3</sup>. Another example of a failed pressure management strategy is the seismic events that followed the injection of large quantities of water at high pressure at the geothermal reservoir in Basel, Switzerland<sup>4-6</sup>. The incident led to public distrust and, finally, the cancellation of the enhanced geothermal systems (EGS) project altogether<sup>7</sup>, which was just one out of multiple EGS projects canceled due to induced seismicity concerns<sup>8</sup>.

Another application that can benefit from a better understanding of pressure management is Geologic CO<sub>2</sub> sequestration (GCS). GCS at a large scale is necessary to reduce anthropogenic CO<sub>2</sub> emissions enough to combat global warming and climate change. To do this efficiently, it is crucial to choose a well-suited reservoir in which 99% of the geologically sequestered carbon will remain sequestered for over 1000 years<sup>9</sup>. This requires a pressure management strategy that minimizes the risk of leakage and potential induced seismicity through wells, faults, and fractures<sup>10-12</sup>. This requires solving complex physics models with sufficient fidelity and enough realizations to understand and forecast the system behavior. This is only feasible if we use methods to speed up the process while keeping the results accurate.

An overwhelming amount of research is focused on using machine learning (ML) to improve the efficiency and accuracy of subsurface energy-related fluid-flow applications<sup>13</sup>. Machine learning has been used to create reduced order models for geologic CO<sub>2</sub> sequestration<sup>14-24</sup>, CO<sub>2</sub> enhanced oil recovery<sup>25-28</sup>, geothermal energy<sup>25,29-32</sup>, geothermal energy<sup>33-37</sup>, and oil and gas extraction<sup>38-42</sup> as summarized in Table 1. For each of these applications, an effective pressure management strategy is required to mitigate the risks associated with injection/

<sup>1</sup>Center for Non-Linear Studies, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. <sup>2</sup>Earth and Environmental Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. ✉email: apachalieva@lanl.gov

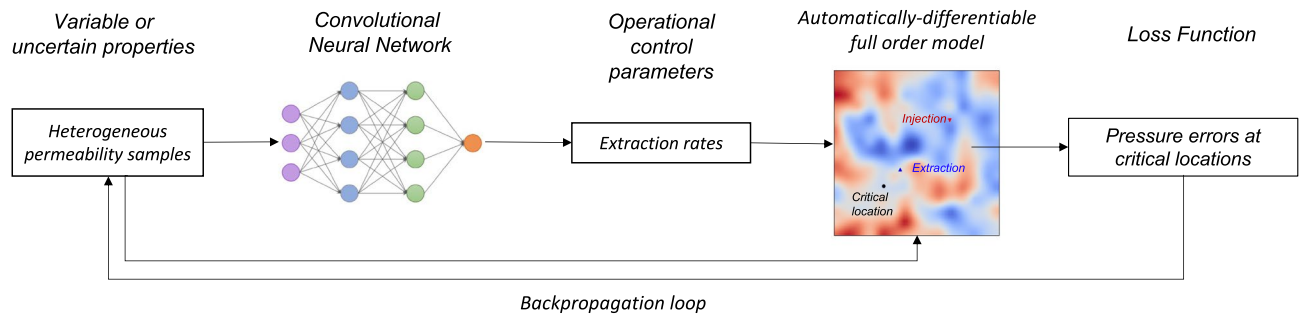
Study	Year	Application				
		CO <sub>2</sub> Sequestration	CO <sub>2</sub> enhanced oil recovery	Geothermal energy	Oil and gas extraction	Injection/extraction
Cihan et al. <sup>11</sup>	(2015)					X
Harp et al. <sup>12</sup>	(2017)					X
Krasnov et al. <sup>25</sup>	(2017)		X	X		
Hegde et al. <sup>38</sup>	(2017)				X	
Chen et al. <sup>14</sup>	(2018)	X				
Menad et al. <sup>15</sup>	(2019)	X				
You et al. <sup>26</sup>	(2019)		X			
Chen et al. <sup>29</sup>	(2019)			X		
Hanga et al. <sup>39</sup>	(2019)				X	
You et al. <sup>30,31</sup>	(2020)			X		
Sinha et al. <sup>16</sup>	(2020)	X				
Harp et al. <sup>45</sup>	(2021)					X
Pirizadeh et al. <sup>28</sup>	(2021)		X			
Sircar et al. <sup>41</sup>	(2021)				X	
Mohamadian et al. <sup>42</sup>	(2021)				X	
Ahmed et al. <sup>20</sup>	(2021)	X				
Ahmed et al. <sup>37</sup>	(2022)			X		
Yan et al. <sup>21–23</sup>	(2022)	X				

**Table 1.** Summary of the main prior studies and their application field.

extraction operations<sup>1,10–12,43–45</sup>. However, the use of machine learning to manage the pressures in a subsurface injection scenario has not been investigated in depth.

Modeling reservoir pressure management is challenging considering the complex heterogeneity of the reservoirs and the uncertainties of the systems' input parameters. This complex heterogeneity typically requires high-fidelity physics-based models to make CO<sub>2</sub> predictions. Furthermore, characterizing the heterogeneity accurately is fraught with parametric uncertainty. Accounting for both heterogeneity (which contributes to the complexity and high computational cost of the physics model) and uncertainty demands many realizations. Performing many realizations makes this a computationally-intensive problem that is challenging for current reservoir simulation workflows. For some applications, such as the oil and gas industry, tens of thousands of wells have been drilled, resulting in large amounts of data and the development and usage of data-driven machine learning models<sup>40,46–49</sup>. However, this is not the case for applications such as GCS, where few wells are in use, data is scarce, expensive, and time-consuming to obtain<sup>14,50,51</sup>. For such data-limited applications, physics constraints are often introduced into the machine learning algorithms to regularize the neural networks training and thus augment the lack of data<sup>52,53</sup>. This approach is called physics-informed neural networks (PINN)<sup>54</sup>. A limitation of the PINN approach is that if the physics are not trustworthy – there are no guarantees that the computation will quickly (or at all) converge to the correct solution. An incorrect solution could misguide the pressure management machine learning model. A major limitation of most traditional numerical models is the calculation of parameter gradients from high-fidelity physics reservoir simulations. Most fluid and transport simulators, which can simulate subsurface fluid injection/extraction rely on finite-difference gradients to evaluate many physics-based parameters<sup>55–59</sup>. This is why such simulators are built without the use of differentiable programming (DP) and automatic differentiation (AD) techniques which are standard in machine learning approaches such as PINNs. Computing finite-difference gradients for highly-dimensionalized models (e.g., those with heterogeneous permeability fields) is computationally inefficient and often prevents the traditional physical models from being included in machine learning workflows. A solution to this problem is using DP and AD that takes advantage of the chain rule to evaluate complex derivatives more efficiently<sup>60</sup>, including for implementing trustworthy numerical models based on traditional methods such as finite difference/element/volume<sup>45,61</sup>.

We developed a physics-informed machine learning (PIML) framework that determines the fluid extraction rates for dedicated wells to maintain the pressure at a critical location during water injection. Our framework uses physics-informed machine learning with differentiable programming, and solves a full-physics model that is used for the training and testing of the machine learning process. We identify the following four significant differences between our approach and the prior studies: (1) we consider a full-physics model using the DPFEHM framework<sup>62</sup> that accounts for the permeability field's heterogeneity, which was overlooked or deemed too expensive by previous studies as discussed in<sup>45</sup>; (2) our physics model is differentiable, which in alternative approaches, such as PINNs, do not rigorously guarantee that the physical constraints will be satisfied; (3) we use a steady-state equation that looks at the long-term impact of the injection/extraction. Harp et al.<sup>45</sup> considered only a fixed time in the future, which limited the ability to evaluate the sustainability of the injection process; and (4) DPFEHM allows us to combine a rigorous, trustworthy physics model with Convolutional Neural Networks (CNN) that have built-in AD. In addition, DPFEHM minimizes the time for code development by automatically creating the execution code from a short description of the equations of interest.



**Figure 1.** (Color online) Workflow diagram of physics-informed machine learning framework for managing reservoir pressures at a critical location during subsurface fluid injection. The key innovation of the surrogate model is the automatically-differentiable full-order model that allows for heterogeneity.

One way to develop a full-physics model for pressure management in the context of CO<sub>2</sub> sequestration would require the completion of the following steps: (1) development of a simple pressure management PIML framework solving the analytical Theis model<sup>63</sup> as shown by Harp et al.<sup>45</sup>; (2) development of a pressure management PIML framework using a full-physics model with heterogeneous permeability field, which we demonstrate in the current study; and (3) adding a multi-phase model. Our single-phase pressure management model is more relevant to wastewater injection. However, including a full-physics model in the ML workflow and accounting for heterogeneity are essential steps towards pressure management for multi-phase injection scenarios such as CO<sub>2</sub> sequestration. Furthermore, without this PIML framework, we would be unable to do uncertainty quantification (UQ) in real-time since thousands (or more) partial differential equation-constrained optimization problems would be required.

## Methods

Operators at underground reservoir sites require pressure management systems to make informed decisions for the injection/extraction rates to minimize the risk of leakage and induced seismicity and maximize the reservoir's performance (e.g., maximize the net fluid injected). The traditional approach to constructing pressure management systems uses full-order physics models that do not allow for UQ in real-time. Existing alternatives, such as the work of Harp et al.<sup>45</sup>, used a highly simplified homogeneous model which neglects heterogeneity. In reality, the subsurface is highly heterogeneous, so for the model to be trustworthy, heterogeneity must be accounted for. Our PIML framework uses a full-physics model and ML combined with DP, making this approach computationally and practically feasible.

A schematic representation of the proposed PIML framework is shown in Fig. 1. We use this framework to determine the fluid extraction rates at an extraction well to maintain the pressure at a critical location in a reservoir with a heterogeneous permeability field. We achieve this by using a neural network model (NNM) trained on a set of permeability fields that, along with the NNM-predicted extraction rate, act as an input to the full-physics DPFEHM model. The NNM is trained to determine the extraction rates for a heterogeneous permeability field. The physics constraints needed for the training process are incorporated through the DPFEHM framework, which provides the physics information in our PIML framework. Our framework has similarities to the works of Harp et al.<sup>45</sup> and Srinivasan et al.<sup>64</sup>, that use physics-informed neural networks with physical constraints used during the NNM training. A significant difference in our approach is that instead of using a differentiable, simple analytical solution, we use a differentiable numerical physics model. Alternative approaches such as PINNs do not have rigorous guarantees that the physical constraints will be satisfied.

We executed a suite of training scenarios using a single-phase model with heterogeneous permeability fields randomly generated using a Gaussian distribution function. Since the model has many parameters, computing finite-difference gradients become infeasible, and the only efficient approach to solve the problem is by using reverse-mode automatic differentiation. We perform a hyperparameter search by varying the batch size (i.e., the number of training samples included in each gradient calculation) and the learning rate (i.e., the step size controlling the rate at which each iteration moves towards the minimum of the loss function).

**Physics model.** We consider the pressure change during injection/extraction of a single-phase fluid flow in a heterogeneous permeability field. Such subsurface reservoirs are modeled using the following equation

$$\nabla \cdot (k(x) \cdot \nabla h) = f, \quad (1)$$

where  $k(x)$  are heterogeneous permeability fields depending on the position  $x$ ,  $h$  is the pressure head and  $f$  is the injection/extraction. This is a steady-state equation which allows us to evaluate the long-term impact of the injection and the extraction on the pressure head.

The equation is solved using DPFEHM<sup>62</sup>. DPFEHM uses the standard two-point flux finite volume approximation, which results in a trustworthy solution to the physics model. In addition, DPFEHM has built-in AD that makes the transition between the physics model and the machine learning model seamless.

**PIML framework.** As shown in Fig. 1, our PIML framework trains an NNM to estimate the extraction rates at dedicated extraction wells to maintain predefined pressures at critical locations during fluid injection. This is necessary near faults with high induced seismicity risk or at abandoned wellbores with leakage potential. To make the model more realistic, we randomly initialized heterogeneous permeability fields using a Gaussian distribution function. The PIML framework is part of the DPFEHM package, which is available at <https://github.com/OrchardLANL/DPFEHM.jl>.

The PIML workflow consists of the following steps:

1. Generate a training dataset that consists of  $N_b$  batches and  $N_s$  samples per batch. We use heterogeneous permeability samples, randomly initialized using a Gaussian distribution function.
2. Construct a CNN with an input layer that accepts a permeability field and an output layer that estimates the fluid extraction rates at the extraction well.
3. Calculate a loss function that quantifies the error between the model's overpressure and the target overpressure at a critical location.
4. Train the CNN to determine the extraction rates that minimize the error between the model's overpressure and the target overpressure at a critical location by adjusting the CNN model parameters based on the loss-function parameter gradients.

In step 2, the PIML framework trains a CNN based on LeNet-5<sup>65</sup> to determine the extraction rates at an extraction well to maintain pressure at critical locations, such as faults with induced seismicity risk or abandoned wellbores with leakage potential, during fluid injection at dedicated injection wells with heterogeneous permeability fields. LeNet-5's architecture consists of a convolutional encoder with two convolutional layers and a dense block with three fully-connected layers. For this work, we use a slightly modified LeNet-5 CNN with the following architecture:

```
model = Chain(Conv((5, 5), 1=>6, relu),
              MaxPool((2, 2)),
              Conv((5, 5), 6=>16, relu),
              MaxPool((2, 2)),
              flatten,
              Dense(1296, 120, relu),
              Dense(120, 84, relu),
              Dense(84, 1)) |> f64
```

We use two convolutional layers, two subsampling layers, a flattening operation, and a fully connected dense block. Each convolutional layer uses a 5x5 kernel and a max pooling subsampling layer which calculates the maximum of the values present in each kernel. The first convolutional layer has 6 output channels, and the second one has 16. The activation function with a stride of 2 reduces the dimensionality through downsampling by a factor of 4. We flatten the output of the convolutional block by taking the four-dimensional input and transforming it into two-dimensional input to be compatible with the fully-connected dense block. The dense block has three fully-connected layers, with 120, 84, and 1 outputs, respectively. Relu activation functions are used for all hidden layers as  $\sigma(x) = \max(0, x)$ , with  $x$  being the input and  $\sigma(x)$  being the output of the neuron. The last dense layer outputs a single value associated with the extraction rate,  $Q_{\text{ext}}$ .

During preliminary investigations, we explored the use of more complex deep neural networks (such as the VGG16<sup>66</sup>) with a larger number of neurons and/or more hidden layers; however, they did not improve the convergence speed and thus the performance of the PIML framework. On the contrary, in some cases, they led to slower convergence and worse overall results. The reason for this is most likely that larger networks are more prone to complex response surfaces, which increases the chance for the training to become trapped in local minima. By contrast, our empirical results suggest that LeNet-5 provides a good balance between strong performance and ease of training. As we transition to more complex physics models, the ease of training becomes more critical when the data sets become smaller. Nonetheless, the optimal architecture for these problems remains an exciting area for research.

In step 3, the loss function is defined as the sum of the squared errors between the simulated and the target overpressures at a critical location as

$$\mathcal{L}(\theta) = \sum_i^{N_b} \sum_j^{N_s} \left[ \Delta h(Q_{\text{NN}}(\theta, k_j(x)), k_j(x)) - \Delta h^{\text{target}} \right]^2, \quad (2)$$

with  $N_b$  being the number of batches,  $N_s$  the number of samples per batch, and  $\Delta h^{\text{target}}$  the target overpressure. The predicted overpressure  $\Delta h$  is a function of the injection rate  $Q_{\text{NN}}$  and the permeability  $k_j(x)$  at a specific position.  $Q_{\text{NN}}$  is calculated using the LeNet-5 CNN model and takes two parameter,  $\theta$  being the model parameters and  $k_j(x)$  permeability. From the loss function we calculate the root-mean-square error (RMSE) as

$$\text{RMSE}(\theta) = \sqrt{\frac{\mathcal{L}(\theta)}{N_b N_s}}. \quad (3)$$

In step 4, the NNM is trained to minimize the loss function  $\mathcal{L}(\theta)$  as

$$\min_{\theta} = \mathcal{L}(\theta). \quad (4)$$

using an Adam optimizer from the Julia Flux package<sup>67</sup>. We use the DPFEHM package to solve the full-physics model as shown in Eq. 1. Within the DPFEHM framework, we automatically differentiate the physics model using Julia's Zygote package<sup>68</sup>.

**Differentiable programming.** Traditional physics models are rarely automatically differentiable and require the usage of finite difference methods to compute parameter gradients<sup>55–59</sup>. This is computationally inefficient and makes incorporating these models in a machine learning workflow infeasible when the number of parameters is large. When looking at a large number of physics-model parameters in a PIML framework, choosing the most efficient way to calculate gradients is essential for the algorithm's success. This is because computing the gradient of the loss function is central to training machine learning models.

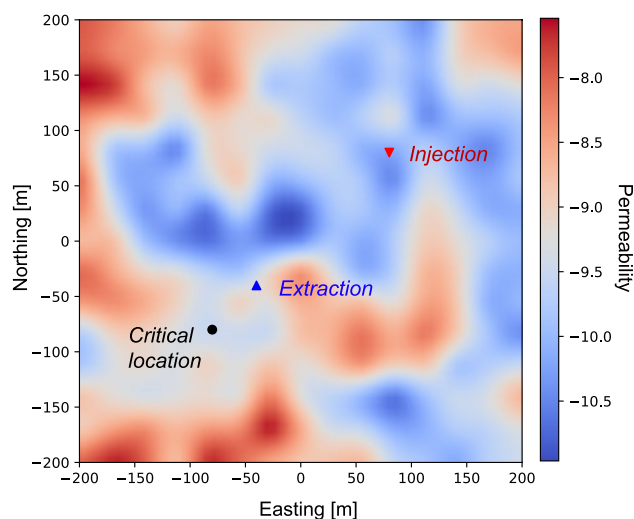
In computational fluid dynamics, finite-differences or numerical differentiation are often used to compute gradients, but that could be computationally expensive, especially when the number of input parameters grows. The number of model runs required to compute the gradient is proportional to the number of input parameters. When using finite difference (FD), the quality of the solution is also greatly influenced by the truncation and round-off errors associated with different finite difference formulas. An alternative is using differentiable programming (DP), and in particular, reverse-mode AD, which utilizes the chain rule to calculate complex derivatives. It breaks a computer program down into elementary operations, which allows the derivatives to be evaluated accurately to working precision. Reverse-mode AD becomes more efficient when the number of model inputs is large (e.g., the parameters of a neural network) and the number of model outputs is small (e.g., the loss), which can be demonstrated with the following example: The LeNet-5 model uses approximately 1 million model parameters, which gradients can be obtained using AD in a single forward and backward model pass regardless of the number of model parameters; in comparison, using a central finite-difference model, we would have to calculate two forward simulations per model parameter which results in 2 million forward calls. In AD, the advantage of using reverse-mode AD (akin to adjoint methods) comes from its computational cost being independent of the number of design variables. By contrast, the cost of the forward sensitivity analysis increases linearly with the number of design variables<sup>69</sup>. Depending on the simulator and the application, the AD adjoint simulations can be more expensive than the forward simulations; however, the discrepancies in execution time have been decreasing with the development of DP<sup>60,70</sup>. For a physics simulator, which typically involves a solution of a system of equations, the backward pass requires only the solution of a linear system of equations (even if the underlying equations are nonlinear). So in this setting, the backward pass is rarely more expensive than the forward pass.

## Results

The physics model includes an injection well, an extraction well, and a critical location. The injection well injects water at a rate of 1.0 MMT/y (million metric tons per year), which is equivalent to  $0.031688 \text{ m}^3/\text{s}$ . The extraction well protects the critical location where a target overpressure (change in pressure from pre-injection conditions) is set to 0.0 MPa. The permeability field is randomly initialized using a multivariate Gaussian distribution function with a correlation length of 50m and log standard deviation equal to 1. A schematic representation of the simulation configuration is shown in Fig. 2, where we indicate the positions of the injection, extraction, and critical location. The permeability field is illustrated using a color map, where blue and red denote low and high permeability, respectively. The size of the simulated domain is 400 meters on each side. Our PIML framework trains the neural network to achieve the target overpressure at the critical location for a given permeability field. We solve a steady-state equation that captures the long-term impact of the injection/extraction. The background reservoir pressure in our simulation is equal to 19.0 MPa, which is in line with MPC 26-5 well located in Kemper County, Mississippi. MPC 26-5 well is part of the ECO<sub>2</sub>S project and its depth is 1791 meters<sup>71,72</sup>.

The PIML algorithm uses 512, 1024, 2048, or 4096 training and disjoint testing samples per epoch depending on the batch size, where the batch size ranges from 32 to 256. The testing data is sampled only once from a multivariate Gaussian distribution, while the training data is sampled in every epoch. Our PIML model solves a physics problem using a unique permeability field for each random sample in each epoch and uses that information to compare the predicted and the target overpressure in the loss function as given in Eq. 2. This characteristic of our framework ensures that the ML model has an (almost) infinite supply of unique datasets, formally called synthetic datasets. That is, the model *never* sees the same data point twice, so there is little risk of overfitting, a unique characteristic of the DP that allows for data to be generated on the fly compared to other ML models that use just a fixed set of data obtained from simulations or/and measurements. This is possible because our DP approach has the physics model in the training loop. Therefore, generating the training data (synthetic data) and running the model is unnecessary before the training begins. The total number of executed epochs is 10,000, but similar results could be obtained using a much smaller number of epochs on the order of 4,000. The training samples are randomly initialized before training each epoch, making the total number of unique training samples for the batch size of 256 equal to 40,960,000 throughout the ML training. The test samples are initialized only once before starting the first epoch. We ran the training on a central processing unit (CPU) with an AMD EPYC 7702P 64-Core processor without parallelization paradigms. The smallest training ran for about

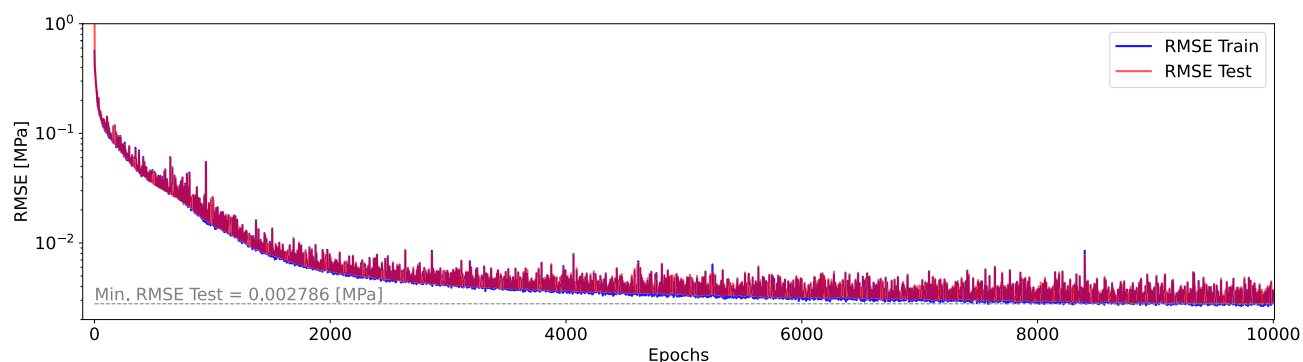




**Figure 2.** (Color online) Diagram of a single-phase pressure management with heterogeneous permeability field including one injection and one extraction wells, and a critical location. This figure was created using Julia Version 1.6.2<sup>73,74</sup>, and the DPFEHM framework<sup>62</sup>.

	bsize = 32	bsize = 64	bsize = 128	bsize = 256
$lr = 1e^{-3}$	0.018658	0.013856	0.010190	0.009823
$lr = 3e^{-4}$	0.005779	0.015159	0.005228	0.006147
$lr = 1e^{-4}$	0.005332	0.003719	0.002852	0.002786

**Table 2.** Hyperparameter search to determine the optimal ML parameters. Minimum testing RMSE for a variety of batch sizes,  $bsize \in [32, 64, 128, 256]$ ; and learning rates,  $lr \in [1e^{-3}, 3e^{-4}, 1e^{-4}]$ . The results are given in pressure units [MPa].

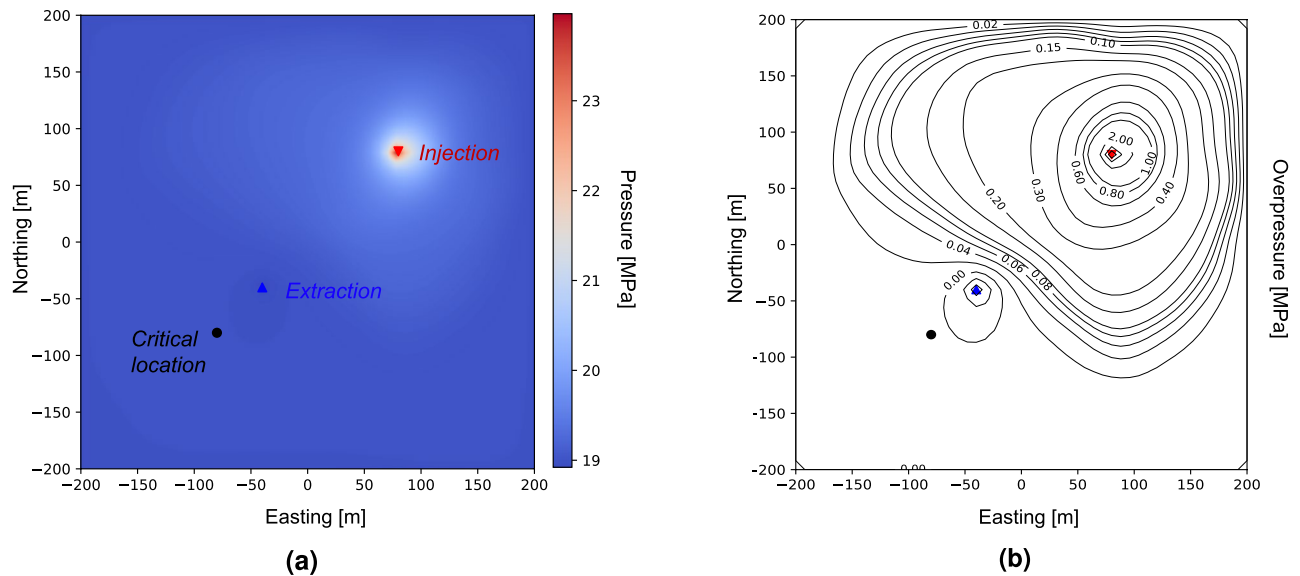


**Figure 3.** (Color online) RMSE for the best training set from the hyperparameter search ( $bsize = 256$ , and  $lr = 1e^{-4}$ ). The minimum RMSE is equal to 0.002786 MPa overpressure at the critical location. This figure was created using Julia Version 1.6.2<sup>73,74</sup>, and the DPFEHM framework<sup>62</sup>.

6 hours on a CPU, while the training with 256 batch size took about 48 hours for 10,000 epochs. We could have stopped the training at epoch 4,000 since we do not see significant improvement of the results. This would have reduced the time for the ML training by more than half.

We perform a hyperparameter search in which we test the RMSE for a variety of batch sizes ( $bsize \in [32, 64, 128, 256]$ ), and learning rates ( $lr \in [1e^{-3}, 3e^{-4}, 1e^{-4}]$ ) as shown in Table 2. The hyperparameter search shows that by decreasing the learning rate from  $lr = 1e^{-3}$  to  $lr = 1e^{-4}$ , the RMSE decreases three to four times for each batch size. The best results are achieved using a batch size and learning rate equal to 256 and  $1e^{-4}$ , respectively. The results can be potentially improved by further decreasing the learning rate. However, we did not expand the search since the model was producing results that were within acceptable limits.

Figure 3 illustrates the rate at which the RMSE decreases with the number of epochs. The results are for the best hyperparameters we tested during PIML training, which uses a batch size of 256 and a learning rate equal to



**Figure 4.** (Color online) (a) Resulting pressure distribution within the reservoir. The positions of the injection/extraction and critical locations are indicated; (b) Overpressure values within the prescribed domain extracted from the PIML framework. The blue and red triangles denote respectively the position of the extraction and the injection wells, respectively, while the circle depicts the critical location at which a target pressure is set. These figures were created using Julia Version 1.6.2<sup>73,74</sup>, and the DPFEHM framework<sup>62</sup>.

$1e^{-4}$ . The figure shows the training RMSE error in blue and the testing RMSE in red. We observe a rapid decrease in the RMSE at the beginning of the training. In later epochs (after approximately 4000 epochs), the RMSE starts to decrease much slower, reaching a plateau around 8000 epochs with an overall error reduction of 99%. The minimum RMSE reached using the testing data samples is 0.002786 MPa, which is small in comparison to the overpressure caused by the injection and makes the model sufficiently accurate to be used in practice. That is, other factors beyond errors in the injection rate, such as uncertainty about the heterogeneity, will impact the decision-making process for the extraction rate, and this model could be used to help find an appropriate extraction rate given uncertainty by predicting an ensemble of extraction rates for a variety of permeability fields.

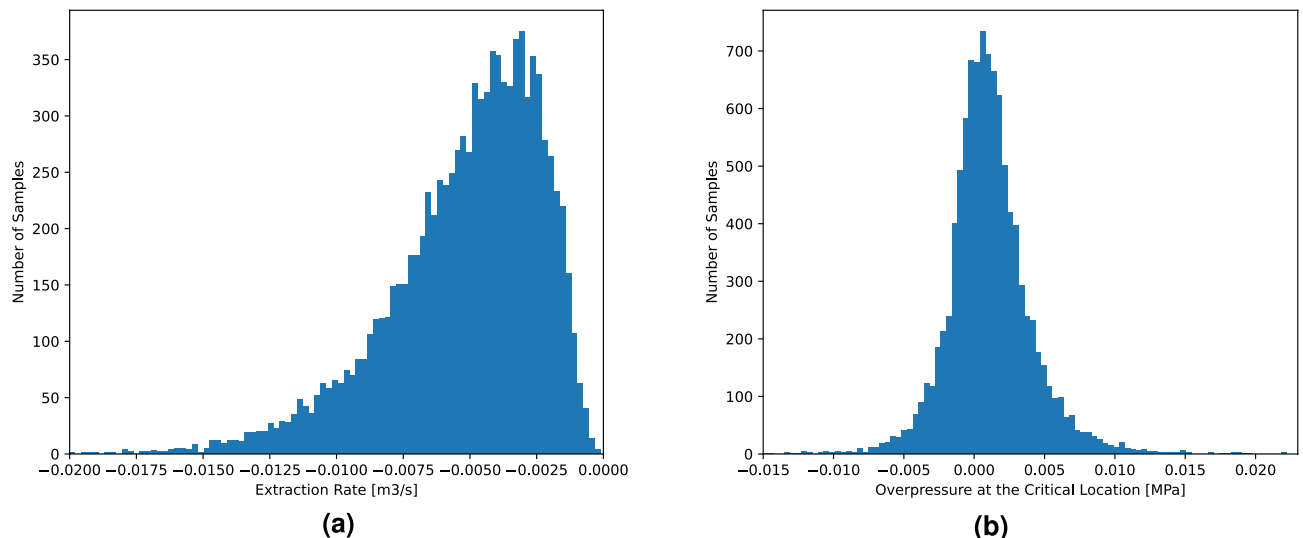
After training the PIML frameworks, we can use them to generate extraction rates and overpressures at the critical location given a heterogeneous permeability field. In Fig. 4, we illustrate the resulting pressures and overpressures for a random permeability field within the domain. We used the same permeability field as shown in Fig. 2, and the PIML framework for which we achieved the lowest RMSE (bsize = 256 and lr =  $1e^{-4}$ ). The background pressure of the MPC 26-5 well is 19.0 MPa, which is reflected in the total pressure of the well depicted in Fig. 4a. Fig. 4b shows that the prescribed pressure at the critical location, 0.0 MPa, is obtained by the PIML framework with good accuracy and, despite the small deviation, still falls into an acceptable range. As expected, the overpressure rates at the injection well are considerably larger, up to 6 MPa. At the same time, at the extraction well, we achieve negative overpressure, which protects the critical location from exceeding the prescribed 0.0 MPa.

Next, we look at the distribution of extraction rates and overpressures at the critical location after running 10,000 random heterogeneous field samples through our PIML framework as shown in Fig. 5. The distribution function of the extraction rate, depicted in Fig. 5a, is skewed and shows that most of the samples require an extraction rate in the range of  $-0.009 \text{ m}^3/\text{s}$  to  $-0.001 \text{ m}^3/\text{s}$  in 90% of the time. Compared to our injection rate of 1.0 MMT/year, the extraction rate is equal to  $-0.283824$  MMT/year, and  $-0.031536$  MMT/year, respectively. Fig. 5b illustrates that the overpressure at the critical location deviates very slightly from the prescribed overpressure, specifically in the range of  $-0.003$  MPa to 0.006 MPa in 90% of the time. These results confirm that the PIML framework can provide useful information for the operators of pressure management systems.

Overall, solving the physics model on a CPU (AMD EPYC 7702P 64-Core) takes  $5.9e^{-3}$  seconds, and obtaining the extraction rates using the trained PIML model on a GPU (NVIDIA RTX A6000) takes only  $1.4e^{-8}$  seconds (equivalent to 14 nanoseconds) per sample. Our framework solves the constrained optimization model more than 400 000 times faster than solving the physics model allowing for near real-time analysis and robust uncertainty quantification. We obtained these results using a total number of 50 000 to fully utilize the GPU, and we excluded the time needed to train the PIML model, since this is done a priori and does not affect the reservoir operators.

## Discussion

We have demonstrated that the PIML framework can successfully manage subsurface pressures by controlling fluid extraction in heterogeneous permeability fields. Such operations are relevant for resource extraction (oil and gas), climate mitigation ( $\text{CO}_2$  sequestration), and production of renewable energy (geothermal energy). A similar PIML approach has already been evaluated by<sup>45</sup> using homogeneous permeability. We have added a significantly more complex physics model using the DPFEHM framework to ensure a seamless transition



**Figure 5.** (Color online) **(a)** Extraction rates in cubic meters per second [ $\text{m}^3/\text{s}$ ]; **(b)** Overpressure at the critical location in pressure units [MPa]. These figures were created using Julia Version 1.6.2<sup>73,74</sup>, and the DPFEHM framework<sup>62</sup>.

between the execution of the physics model and the machine learning model. We show that a PIML framework with built-in AD can train an NNM to determine the fluid extraction needed to achieve pressure management goals while injecting large amounts of fluid into the subsurface. This has the potential to help operators at sites effectively manage reservoir pressures and can be coupled with decision-making strategies to attain operational efficiencies<sup>75–79</sup>.

The computational cost is modest enough that it can be run on a single CPU with AMD EPYC7702P 64-Core processor without parallelization. However, the full-physics model and the heterogeneity add complexity to the model that is only feasible because of the DPFEHM framework and its built-in AD. The number of training epochs can be reduced depending on the accuracy demands, which can be a helpful strategy when considering a more complex physics model. Our results showed that, though we trained for 10,000 epochs, a much smaller number could be used to obtain accuracy consistent with realistic operational goals. This could be important when using more complex physics models (such as multi-phase flow) to keep the computational cost manageable. Parallelization could also be exploited in these cases.

We tested different convolutional neural networks, one of them was the VGG Net, a very deep CNN often used for image recognition<sup>66</sup>. To execute the VGG Net more efficiently, we used a hybrid approach, in which the ML model was executed on the GPU while the physics model was executed on the CPU. Even though the VGG Net is significantly deeper, it did not improve the RMSE of the training and sometimes performed worse. However, such a neural network could be more efficient when training on even more complex physics models such as multi-phase models needed for investigating climate mitigation and, in particular  $\text{CO}_2$  sequestration. These models might require a more complex network to deal with the more complex physics present in these cases, but this is a topic for future research.

In addition, we batch-parallelized the model using Julia's parallel map operation implemented as a `pmap`-function. Since our physics model is only modestly expensive, the communication between processes was longer than the time needed for the computation. Therefore, the `pmap` did not improve the performance. However, this technique could be used more effectively when investigating more complex physics models, such as multi-phase flow. Again, this is a topic for future research.

We anticipate that other machine learning models such as support vector regression, random forest, or supervised machine learning, will not outperform the proposed PIML model with DP. Taking into account the definition of our problem – learning an extraction rate from given permeability data (resembling a  $400 \times 400$  pixel image), injection rates and a fixed pressure at a critical location, we think that using CNNs is a natural choice. They are able to handle a large number of images (in our case permeability fields) and have the flexibility to train using a large number of model parameters. SVR, random forest and supervised ML would not exploit the image structure of the problem. In addition, we would not get better performance with either of these techniques. For the level of full-physics and ML model complexity, the performance we get is really good. We would like to mention that previous works were unable to solve in reasonable time much simpler problems and having a heterogeneous PIML model for pressure management was often infeasible.

Here, we would like to emphasize some of the advantages and the limitations of our PIML framework:

- + the framework is able to predict extraction rates to manage subsurface pressures in heterogeneous permeability fields;
- + the framework can use a large number of ML model parameters due to the built-in AD in DPFEHM;



- + the framework uses a well-know LeNet-5 CNN, that has a simple architecture, it is fast and predicts the extraction rates with an acceptable accuracy;
- + the framework uses a dataset which is produced by solving the full-physics model. Solving the physics model is fully integrated in the training loop, which allows for data to be generated on the fly compared to other ML models that use just a fixed set of data;
- + the framework gives a direct solution on how to control subsurface pressures;
- + the physics model can be batch-parallelized using Julia's parallel map operation implemented as a `pmap`-function, while the ML model can be executed using GPUs. Both of these have been already implemented and tested. For the current problem we did not see significant speed-up, since the physics and the ML models are only modestly expensive;
- the framework does not address issues such as maximizing storage capacity by tuning the injection/extraction rates;
- the full-physics model solves a single-phase steady-state fluid flow with heterogeneity, and currently cannot be used to solve multi-phase flows, which are essential for CO<sub>2</sub> sequestration and other applications.”

In order to extend the framework to multi-physics fluid flows, we would need to perform the following steps: (1) implement a multi-phase full-physics model as part of the DPFEHM framework, (2) use more sophisticated ML model to be able to capture the complexity coming from the physics model and increase the number of parameters and the depth of the network, (3) use CPU and GPU parallelization paradigms to improve the performance and ensure the feasibility of the problem. We plan to pursue this problem in the future.

## Conclusions

We applied a PIML framework to a subsurface pressure management problem that considers the pressure change during injection/extraction. We considered a single-phase steady-state fluid flow with heterogeneity that looks at the long-term impact of the injection/extraction on the reservoir. In our PIML framework, a convolutional NNM is trained to determine fluid extraction rates at a dedicated critical reservoir location during the injection. We performed a hyperparameter search, which showed that decreasing the learning rate while increasing the batch size improves the results. In conclusion, we would emphasize the following observations:

- A PIML framework can train a convolutional NNM to manage reservoir pressures with heterogeneous permeability fields resulting in small deviations from the target overpressure. We accomplished this by combining a differentiable full-physics simulator with a convolutional neural network.
- This problem is only feasible because of the DPFEHM framework, which has a built-in AD. To solve the physics model, we use the DPFEHM framework as part of the PIML framework.
- DPFEHM allows us to bridge the gap between numerical models and machine learning techniques because both are compatible with the same AD frameworks
- The hyperparameter search shows that decreasing the learning rate and increasing the batch size is beneficial for bringing down the RMSE of the ML training. Our results can be potentially improved by further decreasing the learning rate.
- We tested a hybrid implementation of the PIML framework using a deeper convolutional neural network (VGG Net). In that case, the physics model is solved on a CPU, while the machine learning model is solved on a GPU. The results showed no improvement in the ML training, concluding that the relatively simple LeNet network suffices for this problem.
- We also batch-parallelized the combined NNM and physics model using Julia's parallel map operation (`pmap`). However, for the single-phase steady-state flow problem with heterogeneity, the added communication costs outweigh the benefits of parallelization. This is because the physical model, in this case, is relatively inexpensive. We anticipate that for more complex physics models, such as introducing a multi-phase flow, will benefit from such parallelization.
- A natural next step in this line of research is to study this problem in the context of multi-phase flows, which are essential for CO<sub>2</sub> sequestration, among other applications.

## Data availability

All data in this paper is synthetically generated. The computer code for generating and analyzing the data are available at <https://github.com/OrchardLANL/DPFEHM.jl>.

Received: 21 June 2022; Accepted: 19 October 2022

Published online: 04 November 2022

## References

1. Zoback, M. D. Managing the seismic risk posed by wastewater disposal. *Earth* **57**, 38–43 (2012).
2. Keranen, K. M., Weingarten, M., Abers, G. A., Bekins, B. A. & Ge, S. Sharp increase in central Oklahoma seismicity since 2008 induced by massive wastewater injection. *Science* **345**, 448–451 (2014).
3. McNamara, D. E. *et al.* Earthquake hypocenters and focal mechanisms in central Oklahoma reveal a complex system of reactivated subsurface strike-slip faulting. *Geophys. Res. Lett.* **42**, 2742–2749 (2015).
4. Baer, M. *et al.* Earthquakes in Switzerland and surrounding regions during 2006. *Swiss J. Geosci.* **100**, 517–528 (2007).
5. Deichmann, N. *et al.* Earthquakes in Switzerland and surrounding regions during 2007. *Swiss J. Geosci.* **101**, 659–667 (2008).
6. Dyer, B., Schanz, U., Ladner, F., Haring, M. & Spillman, T. Microseismic imaging of a geothermal reservoir stimulation. *Lead. Edge* **27**, 856–869 (2008).

7. Deichmann, N. & Giardini, D. Earthquakes induced by the stimulation of an enhanced geothermal system below Basel (Switzerland). *Seismol. Res. Lett.* **80**, 784–798 (2009).
8. Majer, E. L. *et al.* Induced seismicity associated with enhanced geothermal systems. *Geothermics* **36**, 185–222 (2007).
9. Metz, B., Davidson, O., De Coninck, H., Loos, M. & Meyer, L. *IPCC Special Report on Carbon Dioxide Capture and Storage* (Cambridge University Press, Cambridge, 2005).
10. Buscheck, T. A. *et al.* Combining brine extraction, desalination, and residual-brine reinjection with CO<sub>2</sub> storage in saline formations: Implications for pressure management, capacity, and risk mitigation. *Energy Proc.* **4**, 4283–4290 (2011).
11. Cihan, A., Birkholzer, J. T. & Bianchi, M. Optimal well placement and brine extraction for pressure management during CO<sub>2</sub> sequestration. *Int. J. Greenh. Gas Control* **42**, 175–187 (2015).
12. Harp, D. R. *et al.* Development of robust pressure management strategies for geologic CO<sub>2</sub> sequestration. *Int. J. Greenh. Gas Control* **64**, 43–59 (2017).
13. Viswanathan, H. S. *et al.* From fluid flow to coupled processes in fractured rock: Recent advances and new frontiers. *Rev. Geophys.* **60**, e2021RG000744 (2022).
14. Chen, B., Harp, D. R., Lin, Y., Keating, E. H. & Pawar, R. J. Geologic CO<sub>2</sub> sequestration monitoring design: A machine learning and uncertainty quantification based approach. *Appl. Energy* **225**, 332–345 (2018).
15. Menad, N. A., Hemmati-Sarapardeh, A., Varamesh, A. & Shamshirband, S. Predicting solubility of CO<sub>2</sub> in brine by advanced machine learning systems: Application to carbon capture and sequestration. *J. CO<sub>2</sub> Util.* **33**, 83–95 (2019).
16. Sinha, S. *et al.* Normal or abnormal? Machine learning for the leakage detection in carbon sequestration projects using pressure field data. *Int. J. Greenh. Gas Control* **103**, 103189 (2020).
17. Wang, Z., Dillmore, R. M. & Harbert, W. Inferring CO<sub>2</sub> saturation from synthetic surface seismic and downhole monitoring data using machine learning for leakage detection at CO<sub>2</sub> sequestration sites. *Int. J. Greenh. Gas Control* **100**, 103115 (2020).
18. Amar, M. N. & Ghahfarokhi, A. J. Prediction of CO<sub>2</sub> diffusivity in brine using white-box machine learning. *J. Petrol. Sci. Eng.* **190**, 107037 (2020).
19. Liu, R. & Misra, S. Machine learning assisted exploration and production of subsurface energy and carbon geo-sequestration: A review. *Earth Space Sci. Open Arch.* **56** (2020).
20. Ahmmed, B., Karra, S., Vesselinov, V. V. & Mudunuru, M. K. Machine learning to discover mineral trapping signatures due to CO<sub>2</sub> injection. *Int. J. Greenh. Gas Control* **109**, 103382 (2021).
21. Yan, B., Chen, B., Harp, D. R., Jia, W. & Pawar, R. J. A robust deep learning workflow to predict multiphase flow behavior during geological CO<sub>2</sub> sequestration injection and post-injection periods. *J. Hydrol.* **607**, 127542 (2022).
22. Yan, B., Harp, D. R., Chen, B., Hoteit, H. & Pawar, R. J. A gradient-based deep neural network model for simulating multiphase flow in porous media. *J. Comput. Phys.* **463**, 111277 (2022).
23. Yan, B., Harp, D. R., Chen, B. & Pawar, R. A physics-constrained deep learning model for simulating multiphase flow in 3D heterogeneous porous media. *Fuel* **313**, 122693 (2022).
24. Thanh, H. V., Yasin, Q., Al-Mudhafar, W. J. & Lee, K.-K. Knowledge-based machine learning techniques for accurate prediction of CO<sub>2</sub> storage performance in underground saline aquifers. *Appl. Energy* **314**, 118985 (2022).
25. Krasnov, F., Glavnov, N. & Sitnikov, A. A machine learning approach to enhanced oil recovery prediction. In *International Conference on Analysis of Images, Social Networks and Texts*, 164–171 (Springer, 2017).
26. You, J. *et al.* Assessment of enhanced oil recovery and CO<sub>2</sub> storage capacity using machine learning and optimization framework. In *SPE Europe featured at 81st EAGE Conference and Exhibition* (OnePetro, 2019).
27. Cheraghi, Y., Kord, S. & Mashayekhizadeh, V. Application of machine learning techniques for selecting the most suitable enhanced oil recovery method; challenges and opportunities. *J. Petrol. Sci. Eng.* **205**, 108761 (2021).
28. Pirizadeh, M., Alemohammad, N., Manthouri, M. & Pirizadeh, M. A new machine learning ensemble model for class imbalance problem of screening enhanced oil recovery methods. *J. Petrol. Sci. Eng.* **198**, 108214 (2021).
29. Chen, B. & Pawar, R. J. Characterization of CO<sub>2</sub> storage and enhanced oil recovery in residual oil zones. *Energy* **183**, 291–304 (2019).
30. You, J., Ampomah, W. & Sun, Q. Development and application of a machine learning based multi-objective optimization workflow for CO<sub>2</sub>-EOR projects. *Fuel* **264**, 116758 (2020).
31. You, J. *et al.* Machine learning based co-optimization of carbon dioxide sequestration and oil recovery in CO<sub>2</sub>-EOR project. *J. Clean. Prod.* **260**, 120866 (2020).
32. Siler, D. L., Pepin, J. D., Vesselinov, V. V., Mudunuru, M. K. & Ahmmed, B. Machine learning to identify geologic factors associated with production in geothermal fields: A case-study using 3D geologic data, Brady geothermal field. *Nevada. Geotherm. Energy* **9**, 1–17 (2021).
33. Li, Y., Júlíusson, E., Pálsson, H., Stefánsson, H. & Valfells, A. Machine learning for creation of generalized lumped parameter tank models of low temperature geothermal reservoir systems. *Geothermics* **70**, 62–84 (2017).
34. Rezvanbehbahani, S., Stearns, L. A., Kadivar, A., Walker, J. D. & van der Veen, C. J. Predicting the geothermal heat flux in Greenland: A machine learning approach. *Geophys. Res. Lett.* **44**, 12–271 (2017).
35. Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F. & Repetto, D. Machine learning reveals cyclic changes in seismic source spectra in Geysers geothermal field. *Sci. Adv.* **4**, eaa02929 (2018).
36. Tut Haklıdır, F. S. & Haklıdır, M. Prediction of reservoir temperatures using hydrogeochemical data, Western Anatolia geothermal systems (Turkey): A machine learning approach. *Nat. Resour. Res.* **29**, 2333–2346 (2020).
37. Ahmmed, B. & Vesselinov, V. V. Machine learning and shallow groundwater chemistry to identify geothermal prospects in the Great Basin, USA. *Renew. Energy* **197**, 1034–1048 (2022).
38. Hegde, C. & Gray, K. Use of machine learning and data analytics to increase drilling efficiency for nearby wells. *J. Nat. Gas Sci. Eng.* **40**, 327–335 (2017).
39. Hanga, K. M. & Kovalchuk, Y. Machine learning and multi-agent systems in oil and gas industry applications: A survey. *Comput. Sci. Rev.* **34**, 100191 (2019).
40. Hajizadeh, Y. Machine learning in oil and gas; a SWOT analysis approach. *J. Petrol. Sci. Eng.* **176**, 661–663 (2019).
41. Sircar, A., Yadav, K., Rayavarapu, K., Bist, N. & Oza, H. Application of machine learning and artificial intelligence in oil and gas industry. *Petroleum Res.* **6**, 379–391 (2021).
42. Mohamadian, N. *et al.* A geomechanical approach to casing collapse prediction in oil and gas wells aided by machine learning. *J. Petrol. Sci. Eng.* **196**, 107811 (2021).
43. Axelsson, G., Stefánsson, V., Björnsson, G. & Liu, J. Sustainable management of geothermal resources and utilization for 100–300 years. In *Proceedings World Geothermal Congress*, vol. 8 (2005).
44. Weingarten, M., Ge, S., Godt, J. W., Bekins, B. A. & Rubinstein, J. L. High-rate injection is associated with the increase in us mid-continent seismicity. *Science* **348**, 1336–1340 (2015).
45. Harp, D. R., O'Malley, D., Yan, B. & Pawar, R. On the feasibility of using physics-informed machine learning for underground reservoir pressure management. *Expert Syst. Appl.* **178**, 115006 (2021).
46. Cao, Q. *et al.* Data driven production forecasting using machine learning. In *SPE Argentina Exploration and Production of Unconventional Resources Symposium* (OnePetro, 2016).
47. Mohaghegh, S. D. *Data-Driven Reservoir Modeling* (SPE, 2017).

48. Balaji, K. et al. Status of data-driven methods and their applications in oil and gas industry. In *SPE Europec Featured at 80th EAGE Conference and Exhibition (OnePetro)*, 2018).
49. Xiong, F., Ba, J., Gei, D. & Carcione, J. M. Data-driven design of wave-propagation models for shale-oil reservoirs based on machine learning. *J. Geophys. Res. Solid Earth* **126**, e2021JB022665 (2021).
50. Claproot, M. et al. Workflow using sparse vintage data for building a first geological and reservoir model for CO<sub>2</sub> geological storage in deep saline aquifer A case study in the St. Lawrence Platform, Canada. *Greenh. Gases Sci. Technol.* **2**, 260–278 (2012).
51. Mishra, S. et al. Maximizing the value of pressure monitoring data from CO<sub>2</sub> sequestration projects. *Energy Proc.* **37**, 4155–4165 (2013).
52. Yang, Y. & Perdikaris, P. Adversarial uncertainty quantification in physics-informed neural networks. *J. Comput. Phys.* **394**, 136–152 (2019).
53. Meng, X., Li, Z., Zhang, D. & Karniadakis, G. E. PPINN: Parareal physics-informed neural network for time-dependent PDEs. *Comput. Methods Appl. Mech. Eng.* **370**, 113250 (2020).
54. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707 (2019).
55. Pruess, K. TOUGH2-A general-purpose numerical simulator for multiphase fluid and heat flow. Lawrence Berkeley National Laboratory, LBNL Report Number LBL-29400 (1991).
56. White, M. & Oostrom, M. STOMP subsurface transport over multiple phases: Users guide. Tech. Rep., Pacific Northwest National Lab.(PNNL), Richland, WA (United States) (1997).
57. Harbaugh, A. W. *MODFLOW-2005, the US Geological Survey modular ground-water model: The ground-water flow process* (US Department of the Interior, US Geological Survey Reston, VA, USA, 2005).
58. Zvoloski, G. FEHM: A control volume finite element code for simulating subsurface multi-phase multi-fluid heat and mass transfer May 18, 2007 LAUR-07-3359 (2007).
59. CMG, G. Advanced compositional and unconventional reservoir simulator Version 2018. CMG Ltd., CM Group, Editor (2018).
60. Baydin, A. G., Pearlmutter, B. A., Radul, A. A. & Siskind, J. M. Automatic differentiation in machine learning: A survey. *J. Mach. Learn. Res.* **18**, 1–43 (2018).
61. Wu, H., Greer, S. & O'Malley, D. Physics-embedded inverse analysis with automatic differentiation for the Earth's subsurface. arXiv preprint [arXiv:2208.04426](https://arxiv.org/abs/2208.04426) (2022).
62. O'Malley, D., Harp, D. R. & Vesselinov, V. V. DPFEHM.jl. accessed from <https://github.com/OrchardLANL/DPFEHM.jl> (2022).
63. Theis, C. V. The relation between the lowering of the piezometric surface and the rate and duration of discharge of a well using ground-water storage. *EOS Trans. Am. Geophys. Union* **16**, 519–524 (1935).
64. Srinivasan, S. et al. A machine learning framework for rapid forecasting and history matching in unconventional reservoirs. *Sci. Rep.* **11**, 1–15 (2021).
65. LeCun, Y. et al. Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1**, 541–551 (1989).
66. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556) (2014).
67. Innes, M. Flux: Elegant machine learning with Julia. *J. Open Source Softw.* **3**, 602 (2018).
68. Innes, M. Don't unroll adjoint: Differentiating SSA-form programs. arXiv preprint [arXiv:1810.07951](https://arxiv.org/abs/1810.07951) (2018).
69. Giles, M., Ghate, D. & Duta, M. C. Using automatic differentiation for adjoint CFD code development. Post SAROD Workshop-2005. (2005).
70. Innes, M. et al. A differentiable programming system to bridge machine learning and scientific computing. arXiv preprint [arXiv:1907.07587](https://arxiv.org/abs/1907.07587) (2019).
71. Duguid, A. et al. CO<sub>2</sub> well construction: Lessons learned from United States Department of Energy sponsored projects. In *14th Greenhouse Gas Control Technologies Conference Melbourne*, 21–26 (2018).
72. Riestenberg, D. Survey of existing wellbores in and around Kemper County, Mississippi (Deliverable 4.1). Tech. Rep., Southern States Energy Board, Peachtree Corners, GA (United States) (2018).
73. Bezanson, J., Edelman, A., Karpinski, S. & Shah, V. B. Julia: A fresh approach to numerical computing. *SIAM Rev.* **59**, 65–98 (2017).
74. Julia Core Development Team. Julia Github Repository. accessed from <https://github.com/JuliaLang/julia> (2022).
75. Martin, J. J. *Bayesian Decision Problems and Markov Chains* (Wiley, 1967).
76. Charnes, A., Cooper, W. W. & Rhodes, E. Measuring the efficiency of decision making units. *Eur. J. Oper. Res.* **2**, 429–444 (1978).
77. Geng, L., Chen, Z., Chan, C. W. & Huang, G. H. An intelligent decision support system for management of petroleum-contaminated sites. *Expert Syst. Appl.* **20**, 251–260 (2001).
78. Ben-Haim, Y. *Info-Gap Decision Theory: Decisions Under Severe Uncertainty* (Elsevier, 2006).
79. Zhang, L., Wu, X., Ding, L., Skibniewski, M. J. & Yan, Y. Decision support analysis for safety control in complex project environments based on Bayesian Networks. *Expert Syst. Appl.* **40**, 4273–4282 (2013).

## Acknowledgements

AP acknowledges funding from Center for Nonlinear Studies (CNLS) at Los Alamos National Laboratory. DO acknowledges support from Los Alamos National Laboratory's Laboratory Directed Research and Development Early Career Award (20200575ECR). The authors acknowledge support from the US Department of Energy's Science-informed Machine-learning for Accelerating Real-Time decisions (SMART) initiative.

## Author contributions

A.P.: Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing—original draft, Writing—review and editing. D.O.: Conceptualization, Formal analysis, Software, Methodology, Writing—review and editing. D.R.H.: Software, Writing—review and editing. H.V.: Funding acquisition, Writing—review and editing.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to A.P.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022