

A DEEP-LEARNING ITERATIVE STACKED APPROACH FOR PREDICTION OF REACTIVE DISSOLUTION IN POROUS MEDIA

Marcos Cirne¹ , Hannah P. Menke¹ , Alhasan Abdellatif¹ , Julien Maes¹ , Florian Doster¹ , Ahmed H. Elsheikh¹ 

¹Institute of GeoEnergy Engineering (IGE), School of Energy, Geoscience, Infrastructure and Society (EGIS), Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom

Correspondence to:

Ahmed H. Elsheikh,
A.Elsheikh@hw.ac.uk

How to Cite:

Cirne, M., Menke, H.P., Abdellatif, A., Maes, J., Doster, F., & Elsheikh, A.H. (2026). A Deep-Learning Iterative Stacked Approach for Prediction of Reactive Dissolution in Porous Media. *InterPore Journal*, 3(3), IPJ020926-7. <https://doi.org/10.69631/j1re7z03>

RECEIVED: 15 Dec. 2025

ACCEPTED: 1 Apr. 2026

AVAILABLE ONLINE: 20

Apr. 2026

VERSION OF RECORD: 2

Sept. 2026

ABSTRACT

Simulating reactive dissolution of solid minerals in porous media has many subsurface applications, including carbon capture and storage (CCS), geothermal systems and oil and gas recovery. As traditional direct numerical simulators are computationally expensive, it is of paramount importance to develop faster and more efficient alternatives. Deep-learning-based solutions, most of them built upon convolutional neural networks (CNNs), have been recently designed to tackle this problem. However, these solutions were limited to approximating one field over the domain (e.g. velocity field), not accounting for the coupled evolution of multiple interacting fields, including concentration, porosity and flow rates. In this manuscript, we present a novel deep learning approach that incorporates both temporal and spatial information to predict the future states of the dissolution process at a fixed time-step horizon, given a sequence of input states. The overall performance, in terms of speed and prediction accuracy, is demonstrated on a numerical simulation dataset, comparing its prediction results against state-of-the-art approaches, also achieving a speedup around 10^4 over traditional numerical simulators.

KEYWORDS

Reactive Dissolution, Deep Learning, Iterative Stacking, Time-Series Data

WHY THIS PAPER MATTERS

Predicting how underground rocks dissolve over time is essential for technologies such as carbon capture and storage (CCS) and geothermal energy, but traditional computer simulations based on physical models are often slow and computationally expensive. In this work, we developed an artificial intelligence approach that learns from existing data how the dissolution process evolves over time, enabling predictions of future conditions up to 10,000 times faster than conventional methods. Our work could help scientists and engineers exploring climate and energy solutions evaluate more scenarios, accelerating the development of sustainable subsurface technologies.



© 2026 The Authors

This is an open access article published by InterPore under the terms of the Creative Commons Attribution 4.0 International license (CC BY 4.0 Deed)
<https://creativecommons.org/licenses/by/4.0/>

1. INTRODUCTION

Numerical solvers have been extensively used to simulate and understand the effects of reactive dissolution of solid minerals in subsurface porous media across a wide range of applications, including CO_2 sequestration (41), hydrogen storage (15), enhanced oil recovery (9), radioactive waste disposal (25) and geothermal systems (34). Due to the intrinsic complexity of these processes, which are governed by a set of highly non-linear partial differential equations (PDEs), it is computationally expensive to simulate (19). Recently, deep learning (DL) algorithms have become a prominent tool for speeding up the modelling process, while at the same time generating highly accurate simulations of subsurface fluid dynamics (6, 12, 21, 45).

Typically, DL algorithms for subsurface applications rely on data-driven approaches that require a significant number of examples so that these algorithms can properly learn the underlying physics of the phenomena to be studied. Convolutional neural network (CNN)-based methods (2, 13, 23, 35, 40), which employ a series of spatial convolutions to extract meaningful features from image-like data, have been widely adopted to predict specific properties of porous media (such as porosity, permeability and fluid flow). Alternatively, physics-informed neural networks (PINNs) (8, 14, 44) have been adopted to embed the physical laws that govern a given dataset as a prior information into deep neural networks (DNNs).

When it comes to the task of forecasting the future evolution of nonlinear dynamic systems, both categories of neural networks naturally struggle to yield reasonable predictions. Several approaches attempt to combine those types of networks with time-dependent units, including recurrent neural networks (RNNs) (27, 30), convolutional long short-term memory networks (ConvLSTMs) (4, 10) and gated recurrent units (GRU) (1, 7). A more recent approach named recurrent neural operator (RNO) (18) was used to map functions rather than discrete data points (thus reducing computational costs) for prediction of reactive flow, but this approach is only applied within large-scale domains. Others have adopted a hybrid approach, such as Reichstein et al. (33), who leverages physical process models along with data-driven machine-learning (ML) algorithms. The forecasted lead time is bound to the studied phenomenon, and may vary from milliseconds to hundreds of years. However, these methods can only be trained to predict a fixed (limited) amount of future steps, regardless of the time unit.

In order to yield predictions from an initial state for long-term horizons, those dynamic systems employ an iterative strategy in which the output derived from a prediction is used to comprise the input for the subsequent prediction and so on. However, this approach suffers from error propagation with each future step, as the input distributions are more likely to shift away from the distribution under which those systems were trained (22). Although there are alternatives to avoid the pitfalls of iterative strategies, they are still highly vulnerable to performance degradation due to the intrinsic uncertainties in forecasting further time steps. A possible heuristic to minimize prediction errors is by stacking multiple networks and perform an iterative process to minimize the overall residual error of the predictions (17). In other words, at each level of the stack, the corresponding network tries to improve the results achieved in the previous level.

This paper presents an approach to predict the dynamical evolution of reactive dissolution in porous media. Given an ensemble of numerical pore-scale simulations containing different trajectories for the dissolution process, we train a DL algorithm in a data-driven and supervised manner that accounts for spatial and temporal features and serves as a surrogate model to forecast future dissolution states. From a sequence of input states, the algorithm is first trained to predict a fixed

amount of output states. To assess the quality of the dissolution forecasting, an iterative stacked strategy is adopted, and the outputs are evaluated against a ground truth by means of similarity and error metrics. Moreover, a multi-level stacking approach will be investigated as an attempt to reduce the error accumulation.

The novel contributions of this paper are twofold:

- We develop an iterative stacked framework that performs simultaneous prediction of multiple fields and can be applied to both single-step and multi-step scenarios, while yielding accurate results for reactive dissolution simulations that are orders of magnitude faster than traditional numerical solvers;
- We conduct a comparative analysis among different DL architectures in terms of speed and prediction accuracy.

The remainder of this manuscript is organized as follows: Section 2 defines the problem addressed by our work, including the principles of reactive dissolution, forecasting strategies and deep learning algorithms; Section 3 details our proposed methodology for multi-step prediction of reactive dissolution, as well as the DL methods used in our experiments; Section 4 describes our dataset and the preprocessing steps prior to training the DL algorithms; Section 5 discusses the produced results; finally, we conclude the paper in Section 6.

2. PRELIMINARIES

2.1. Problem Statement

Let $\mathcal{S} = \{X_1, X_2, \dots, X_N\}$, where $X_i \in \mathbb{R}^{C \times H \times W}$, be a numerical simulation of a reactive dissolution process composed of N consecutive time steps of state maps of size $H \times W$ for each of the C physical properties. The governing equations for pore-scale reactive dissolution follow the formulation proposed by Maes et al. (28). Also, let m be the size of an input sequence of consecutive states and t a starting point from which we want to predict the next n subsequent states. The problem of forecasting consists of taking an input sequence $\mathcal{X} = X_{t-m+1:t} = \{X_{t-m+1}, X_{t-m+2}, \dots, X_{t-1}, X_t\}$ and predicting an output sequence $\mathcal{Y} = \hat{X}_{t+1:t+n}$ based on the input $\mathcal{X}$. In other words, we aim to train a DL model $\mathcal{F}$ with learnable parameters Θ to learn a mapping $\mathcal{F}_\Theta : \mathcal{X} \mapsto \mathcal{Y}$, where $\mathcal{Y}$ represents the ground truth states, by minimizing the error between the set of predicted states $\hat{\mathcal{Y}} = \mathcal{F}_\Theta(\mathcal{X})$ and the ground truth states $\mathcal{Y}$ according to a loss function $\mathcal{L}$. Therefore, the optimal set of parameters Θ^* for $\mathcal{F}$ is stated as in **Equation 1**:

$$\Theta^* = \arg \min_{\Theta} \mathcal{L}(\hat{\mathcal{Y}}, \mathcal{Y}) \quad (1)$$

We refer to cases where $n > 1$ as **multi-step forecasting** and to those where $n = 1$ as **single-step forecasting**.

2.2. Forecasting Strategies

An in-depth analysis of possible forecasting strategies for multivariate time-series data can be found in Lim and Zohren (26). Among the existing strategies, there is **direct forecasting**, in which a DL model is trained to forecast each of the n future steps (29), but it does not consider the relationships among the predicted states $\hat{X}_{t+i}$, also being limited to a maximum forecast horizon n . Another strategy is known as **iterative forecasting** (also called **recursive forecasting**), in which a trained model outputs one step ahead, then uses this output to comprise the inputs for the prediction of the next step, which yields the output for the subsequent step, and so on. However, this strategy is highly prone to error accumulation, especially for longer horizons in which all inputs are forecasted values rather than actual observations (38).

To circumvent those issues, there are alternate strategies that predict multiple steps at the same time. The most widely used is known as **Multiple-Input Multiple-Output (MIMO)**, which preserves the dependency between the forecasted values, also reducing the error accumulation problem from the iterative strategy (up to time step $t + n$).

Besides being applied to time-series forecasting, this strategy has also been widely used for spatiotemporal forecasting tasks, such as video prediction (11, 32) and earth science forecasting (31, 43). Even though the forecasting horizon is limited by the model, one can perform recursive steps to yield predictions for longer time steps.

2.3. Deep Learning Methods

2.3.1. Encoder-Decoder ConvLSTMs


Convolutional Long Short-Term Memory Networks (ConvLSTMs) (36) are a class of neural networks designed to capture spatiotemporal dependencies in a time-evolving sequence, combining the strengths of CNNs and Long Short-Term Memory networks (LSTMs) into a single method. Given an input sequence $\mathcal{X}_t$, as well as the hidden state $\mathcal{H}_{t-1}$ and the cell state $\mathcal{C}_{t-1}$ from the previous time step, the gate mechanisms contained in a ConvLSTM cell – represented by i_t (input gate), f_t (forget gate), and o_t (output gate) – control the amount of information that is going to be propagated (or forgotten) for the next states. The process of computing the next hidden and cell states from an input sequence $\mathcal{X}$ can be mathematically described as shown in **Equation 2**, where σ is the sigmoid activation function, $*$ represents the convolution operations, $\circ$ stands for the Hadamard product (element-wise product), W and b are weights and biases for each of the gates and the cell state.

$$\begin{aligned} i_t &= \sigma(W_{xi} * \mathcal{X}_t + W_{hi} * \mathcal{H}_{t-1} + W_{ci} \circ \mathcal{C}_{t-1} + b_i) \\ f_t &= \sigma(W_{xf} * \mathcal{X}_t + W_{hf} * \mathcal{H}_{t-1} + W_{cf} \circ \mathcal{C}_{t-1} + b_f) \\ \mathcal{C}_t &= f_t \circ \mathcal{C}_{t-1} + i_t \circ \tanh(W_{xc} * \mathcal{X}_t + W_{hc} * \mathcal{H}_{t-1} + b_c) \\ o_t &= \sigma(W_{xo} * \mathcal{X}_t + W_{ho} * \mathcal{H}_{t-1} + W_{co} \circ \mathcal{C}_t + b_o) \\ \mathcal{H}_t &= o_t \circ \tanh(\mathcal{C}_t) \end{aligned} \quad (2)$$

To perform multi-step prediction with ConvLSTMs, a sequence of ConvLSTM cells is structured as an encoder-decoder architecture (also known as **Seq2Seq** model), where both encoder and decoder blocks contain the same number of ConvLSTM cells. This type of architecture is inspired on sequence-to-sequence (Seq2Seq) models for natural language processing and time-series tasks (3, 37). An example is depicted in **Figure 1**, with two ConvLSTM cells in each block. After an input sequence $\mathcal{X}$ is processed by the encoder block, the final hidden state from the last ConvLSTM cell forms a compact representation of $\mathcal{X}$ (latent representation), which is fed to the decoder block to produce the hidden states for each output time step. Finally, a 3D convolutional layer receives the output from the decoder block and produces the predictions of all physical properties for each desired time step. More details about the encoding and decoding processes can be found in P. R. Kakka (16).

2.3.2. U-Shaped Fourier Neural Operator

As proposed by Wen et al. (42), the U-Shaped Fourier Neural Operator (U-FNO) is an extension of the Fourier neural operator (FNO) (24), designed to solve PDEs across diverse problems involving computational fluid dynamics. Fourier neural operators are known to be resolution-invariant, meaning that they can be trained on a lower resolution and evaluated on a higher resolution, and yield superior performance against conventional CNNs by operating directly on the Fourier space (frequency domain), replacing convolution operations by pointwise multiplications, which are much faster and efficient.



Originally, the core idea of FNOs is to learn mappings of functions (also known as operators) defined in a continuous domain, rather than operating on discretized functions on 2D grids, like images. Still, FNOs are able to operate on image-like data (as used in our work), despite not leveraging the full potential of the process of operator-learning.

To increase the overall performance in multiphase flow problems, the aforementioned authors introduced the U-Fourier layer as an upgrade of the original FNO architecture, combining the advantages of both CNN- and FNO-based models, increasing both training and test accuracies. On the other hand, this improvement in accuracy comes at the expense of the flexibility of training and testing at different resolutions. Moreover, U-FNO's usually take a longer time to be trained than traditional FNO's.

2.3.3. Temporal Attention Unit

Developed for video prediction tasks, Temporal Attention Units (TAUs) (39) leverage the ability to capture time evolution in image sequences by introducing parallelizable attention mechanisms that eliminate the need of recurrent-based units (such as RNNs and ConvLSTMs), thus speeding up the training process. In turn, the spatial modules are represented by simple 2D convolutions. Moreover, not only TAUs account for intra-frame differences through the mean squared error loss, but also for inter-frame variations by embedding a differential divergence regularization term. The resulting loss function is expressed in **Equation 3**:


$$\mathcal{L}(\hat{y}, y) = \sum \|\hat{y} - y\|^2 + \alpha \mathcal{L}_{reg}(\hat{y}, y) \quad (3)$$

where $\mathcal{L}_{reg}$ represents the Kullback-Leibler divergence between the probability distributions of the inter-frame differences from $\hat{y}$ and y , and α is a weight term defined empirically.

3. METHODOLOGY

Without loss of generality and assuming $m = 5$ input steps and $n = 5$ output steps, we first train a base network by making it receive a sequence of 5 "perfect" input states and predict the subsequent 5 states (i.e., without successive iterations), as done in a traditional MIMO approach. After the training process, to evaluate the full evolution of the dissolution process on a given simulation $\mathcal{S}$ containing N total steps, we conduct successive iterative predictions, as illustrated in **Figure 2**, by taking the first 5 input states from $\mathcal{S}$, yielding outputs at time steps 6-10. In turn, this output is fed as an input to the same base network to produce the states at time steps 11-15, and the process is repeated iteratively.

To improve the initial solution produced by the base network, we propose a multi-level stacking of trained networks. In this case, the same neural network architecture is used to train each level, with the same sizes for both input and output tensors as the base network, as well as the loss




function. Considering the base network as Level 0, trained with a dataset $\mathcal{D} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{t_s}\}$, where t_s is the number of training samples, we generate a new dataset comprised of all possible output sequences $\hat{\mathcal{Y}}$ computed by the base network for each $\mathcal{S}_i \in \mathcal{D}$, representing initial approximations to the true states $\mathcal{Y}$. The Level 1 network is then trained to receive, as input, one of those responses produced by the base network, and outputs a new approximation (correction) to the ground truth within the same time-step interval, hypothesizing that it will yield less errors than the initial solution. Following this rationale, we can continue the stacking process by training a Level 2 network with the outputs produced by Level 1, and so on until no further improvement is observed. During this process, only one level can be trained at a time.

Conceptually, this multi-level stacking strategy should be interpreted as a sequential residual-correction procedure rather than as a single deeper end-to-end network. The base network first produces an approximate prediction over the target horizon, and each subsequent level is trained on the prediction distribution generated by the previous level in order to learn a correction over the same interval. In this sense, the later levels are exposed to structured residual errors that arise from the preceding predictor, which differs from simply increasing the depth or width of a single jointly optimized architecture. The usefulness of this strategy is expected to depend on the existence of a learnable residual structure in the errors of the previous level. Conversely, if the base model already captures the dynamics well, or if the residuals become dominated by noise and distribution shift, additional correction levels may provide limited gains or even degrade performance.

In general, let L be the number of levels of correction to be applied on the base solution. In the present study, we evaluate corrections up to $L = 3$ as a practical upper bound to examine whether improvements persist beyond the first correction while keeping the overall training cost manageable; however, this value should be regarded as a tunable design choice rather than an optimal or

universal setting. Increasing the number of correction levels may also increase the risk of overfitting, particularly when the residual errors become less structured at higher levels. After the base and all the L correction networks are trained, the iterative multi-step prediction process can be conducted as illustrated in **Figure 3** with $L = 3$. Here, the first 5 input states of a simulation sample are fed into the base network, producing an initial approximation for the subsequent 5 states. Then, this output is fed to the Level 1 network which produces the first level of correction for the initial approximation. Later, this output serves as input to Level 2 network, whose output is forwarded as an input to Level 3 network, which yields the final prediction for time steps 6-10. For the next prediction step, the output from the last network of the stack is fed to the base network to produce an initial approximation for steps 11-15, which is corrected by the subsequent levels in the stack. The whole process is repeated until the simulation end time.



Although the methodology was illustrated for $m = n = 5$, it can be easily adapted to receive any m input states and predict n output states. For greater values of m and n , the overall performance of the framework tends to drop due to larger error accumulations. Conversely, lower values of those variables yield more accurate predictions, despite leading to a longer training process and total inference time to predict all time steps.

4. REACTIVE DISSOLUTION DATASET AND PREPROCESSING

Our reactive dissolution dataset (5) consists of 32 numerical simulations generated using GeoChemFoam (28), an open-source micro-continuum solver for pore-scale reactive transport. Each simulation follows the *Benchmark 2* configuration presented by Maes et al. (28), which models calcite dissolution in a quasi 2D micromodel. The computational domain consists of a poly-disperse disk pack with an initial porosity of approximately $\phi_0 = 0.45$ and an initial permeability of $K_0 = 5.6 \times 10^{-10} \text{ m}^2$. Simulations are performed on a uniform Cartesian mesh with a spatial resolution of $\Delta x, y = 25 \mu\text{m}$ and a $\Delta z = 150 \mu\text{m}$.


The flow field is incompressible and governed by the Darcy-Brinkman-Stokes equations under isothermal conditions, while the transport of a single reactive species obeys an advection–diffusion equation with a surface reaction rate $R = k_c c$ applied at the fluid–solid interface. A representative parameter set corresponding to the *conical wormhole regime* is as follows: kinematic viscosity $\nu = 1 \times 10^{-6} \text{ m}^2/\text{s}$, diffusion coefficient $D = 1 \times 10^{-9} \text{ m}^2/\text{s}$, inlet concentration $c_i = 0.01 \text{ kmol/m}^3$, stoichiometric coefficient $\zeta = 1$, calcite density $\rho_s = 2710 \text{ kg/m}^3$, molar mass $M_{ws} = 100 \text{ kg/kmol}$, Kozeny–Carman constant $k_f = 2 \times 10^{11} \text{ m}^2$, a volumetric flow rate $Q = 5.2174 \times 10^{-12} \text{ m}^3/\text{s}$, and reaction rate constant $k_c = 8.1632 \times 10^{-6} \text{ m/s}$. These conditions correspond to Péclet and kinetic numbers of $Pe = 1$ and $Ki = 1$, respectively, representing the transition between compact and wormhole dissolution regimes. Boundary conditions follow that of Maes et al. (28): a constant inflow rate at the inlet, fixed outlet pressure, and no-slip/no-flow boundaries on the top and bottom walls.

The 2D simulation results are stored at 100 equally spaced time steps, resulting in time-evolving dissolution state maps of size 260×260 . For the **ML** algorithms tested in this work, all maps were cropped to 256×256 by removing the first and last two rows and columns from each state map. Each of these simulations conveys a particular rock sample with its own pore structure, as well as a distinct fluid trajectory over the course of the dissolution process. An example is illustrated in **Figure 4**. These maps encompass four different input properties: **C**, the concentration of the acidic solution used in the dissolution; **eps**, a function indicating the volume fraction of the pore space occupied by pore in each voxel, which is proportional to the amount of fluid contained in each voxel; **U_x**, the magnitude and direction of flow of the acidic solution in the horizontal axis of a 2D Cartesian plane; and similarly **U_y**, for the vertical axis. The average time to produce each simulation was approximately 3 hours using 24 CPUs of 3 GHz each. All simulations employ the same flow and reactive transport conditions, where both Péclet and kinetic numbers are equal to 1.

We note that, although each simulation provides a large number of spatial samples over time, the present dataset contains only 32 distinct pore geometries generated within a single benchmark family. Therefore, this dataset should be interpreted as a controlled test bed for methodological evaluation rather than as an exhaustive representation of the broader class of pore structures encountered in reactive porous media. For the training/validation split, complete simulations (rather than individual time steps) were separated, so the validation set consists of unseen pore geometries that were not used during training. As a result, the reported performance reflects generalization to new realizations within the same benchmark configuration, but does not yet establish robustness across substantially different pore-structure classes, reaction regimes, or boundary conditions. Extending the dataset to cover broader geometric descriptors and multiple flow-reactive regimes remains an important direction for future work.

In essence, GeoChemFoam solves the reactive transport in porous media in a quasi-static regime, as described by Maes et al. (28). Given a porosity map **eps** at a time step t , it first calculates the steady-state velocity (i.e., it produces **U_x** and **U_y** at time step t). Then, it calculates a steady-state concentration **C** at time step t , which includes the contributions of injection, diffusion, and reactive dissolution at the fluid–solid interface. From these calculations, it estimates a local reaction rate in each computational fluid and the new **eps** field for the time step $t + 1$, which in turn is used to calculate the flow once again to produce a new **eps** map at time step $t + 2$. This iterative process is then repeated for subsequent time steps.

Aside from the original features, three extra input features were engineered as an attempt to improve the performance of all DL models: 1) **Magnitude of velocity**: defined as $U = \sqrt{U_x^2 + U_y^2}$; 2) **C Scaled**, a log-transformation over the concentration values; and 3) **Combined Filter**, a binary mask based on C and eps constraints which shows the portions of the grains in a porous matrix



that are being dissolved at a particular time step t . Given a position (i, j) in a map, the combined filter is calculated according to **Equation 4**:

$$\text{Filter}(i, j) = \begin{cases} 1 & \text{if } C(i, j) \geq 10^{-4} \text{ and } 0.01 \leq \text{eps}(i, j) \leq 0.99, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

5. PERFORMANCE EVALUATION OF DEEP LEARNING MODELS

In this section, we will discuss the quantitative and qualitative results of our proposed method on all algorithms described in Section 2.3. The source code is available at <https://github.com/ai4netzero/ReactiveDissolution>.

5.1. Training and Evaluation Settings

From the 32 simulation samples in our dataset, 24 were randomly selected as the training set and the remaining 8 as the validation/test set. Before the DL models are trained, the entire data is normalized with respect to the mean and standard deviation values of the training set. All input features, except for C Scaled and the Combined Filter, have their original values subtracted by their respective means and the results are each divided by their respective standard deviations. On the other hand, a min-max normalization is applied to the output values so that they all belong within the $[0, 1]$ range, making it suitable to adopt the sigmoid activation function for all DL models.

Concerning our proposed multi-level stacking of neural networks method, for each DL algorithm described in Section 2.3, we perform corrections up to Level $L = 3$. The models were trained on a NVIDIA Titan RTX with 24 GB of memory for a total of 100 epochs (along with a patience rate of 20 epochs for early stopping), using a batch size of 4, a learning rate of 0.0005 and the Adam optimizer (20) with moment values $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The mean squared error (MSE) was adopted as the loss function for ConvLSTM, U-FNO and the intra-frame difference term of TAU. Regarding the latter, the α constant for the regularization term (inter-frame difference) was fixed at 0.1.

For the quantitative evaluation of the iterative predictions of each output property, we calculate the Pearson correlation coefficient (PCC) to assess the similarity between the predicted state $\hat{Y} \in \mathbb{R}^{H \times W}$ and the ground truth state $Y \in \mathbb{R}^{H \times W}$ at each time step, as stated in **Equation 5**:

$$\text{PCC}(\hat{Y}, Y) = \frac{\text{cov}(\hat{Y}, Y)}{\sigma_{\hat{Y}} \cdot \sigma_Y} \quad (5)$$

where $\text{cov}(\hat{Y}, Y)$ is the covariance between $\hat{Y}$ and Y and σ represents the standard deviation of a state map.

It is important to emphasize that the present inference framework is purely data-driven and does not explicitly enforce physical constraints or conservation laws during prediction. Accordingly, the long-horizon results reported here should be interpreted as surrogate predictions within the range of conditions represented by the training data, rather than as guaranteed physically constrained forecasts under arbitrary reactive transport settings. Incorporating physics-aware constraints into the training objective or architecture is a promising avenue for improving robustness in future work.

5.2. Model Statistics

Table 1 shows some statistics about the average training epoch times, number of trainable parameters, final model sizes and average forward times (i.e., the elapsed time for processing input data through a network to produce the next 5 time steps) for each DL model. Although TAU has the highest numbers of trainable parameters and a relatively larger model size, it still achieved the lowest training and forward times, due to its recurrence-free structure.

Table 1: Performance statistics for all tested models. Best results for each category are highlighted in bold.

DL Model	Epoch Train Time (min)	Total Train Time (h)	Trainable Parameters	Model Size (MB)	Forward Time (ms)
ConvLSTM	6.44	10.73	4 M	13	60
U-FNO	10.98	18.30	6 M	25	79
TAU	4.22	7.03	11 M	135	33

DL: deep learning; ConvLSTM: convolutional long short-term memory network; U-FNO: U-shaped Fourier neural operator; TAU: temporal attention unit.

Despite being the most lightweight model, ConvLSTM has a 50% slower training time per epoch than TAU, also taking nearly twice as long for the forward operation. Finally, U-FNO achieved the slowest training and forward times, but the resulting model after the training process is about five times smaller than TAU and almost two times bigger compared to ConvLSTM.

Analyzing the total runtime of each algorithm to produce a 100-step simulation, TAU takes an average time of 0.66 seconds, followed by ConvLSTM (1.2 seconds) and U-FNO (1.58 seconds). These represent a speedup with an order of magnitude between 10^3 and 10^4 when compared to the average time taken by GeoChemFoam (3 hours).

5.3. Iterative Prediction and Model Stacking Analysis

Figure 5 shows the average correlations of the iterative predictions on all training samples for all algorithms and correction levels. For all properties, it can be observed that both TAU and U-FNO are more robust to error accumulation, with the latter achieving slightly higher correlations at late time steps. Considering the effects of model stacking, ConvLSTM achieved its best results at Level 2 correction (except for U_x prediction). Still, it was not capable of performing better than the base networks (Level 0) of the other two algorithms. Overall, TAU showed no significant improvement when applying multi-level stacking for all cases, as evidenced by the fact that the correlations do not systematically increase over the correction levels. On the other hand, U-FNO showed a more consistent evolution from the base network to Level 3 in all scenarios, as the average correlations at the late time steps typically increase after each correction.





Figure 6 illustrates the results for the validation set. The plots show that TAU achieved the best iterative predictions for all cases (except for U_x prediction). Apart from ConvLSTM, the correlations had a reasonable drop when compared to those from the training data. However, this drop is much larger for the U-FNO, suggesting that this network is much more prone to over-fitting than the others.

Regarding the performance evolution along the correction levels, ConvLSTM still yielded its best results on Level 2 network, except for U_x prediction, in which Level 3 was more robust to error accumulation. However, it achieved the lowest correlations in almost all scenarios, being only superior to U-FNO in the prediction of U_y . The validation curves for U-FNO, unlike the results for the training set, did not show a consistent evolution over the levels. For instance, in the prediction of C , the Level 1 network achieved the best results until time step 80, where it was surpassed by Level 3. Moreover, Level 1 prevailed as the best correction network during all time steps of eps prediction. Finally, the Level 1 network from TAU achieved the highest correlations for predictions of C



and *eps*, as well as for U_x predictions until time step 50. Despite its better performance against ConvLSTM and U-FNO, the multi-level stacking strategy did not cause a significant improvement over the TAU Level 0 network. As the algorithms were trained using the MSE loss function, a more carefully designed loss function, tailored to capture reactive transport properties, could potentially lead to more systematic improvements, especially in achieving progressively better results at higher correction levels.


More broadly, the results indicate that the effectiveness of the multi-level stacking strategy is architecture-dependent in the present study. For ConvLSTM, the base predictor appears to leave residual errors that remain structured enough to be partially corrected by the subsequent levels, which explains the more consistent gains observed across the correction stages. For TAU, the Level 0 predictor already performs strongly, leaving less room for systematic improvement; in that case, the remaining residuals may be too small or too irregular for later correction networks to provide noticeable benefits. In contrast, the validation behavior of U-FNO is less consistent, suggesting that its correction levels may be more sensitive to distribution shift and overfitting. Therefore, the proposed stacking framework should not be interpreted as uniformly beneficial across all model classes, but rather as a correction mechanism whose utility depends on the residual-error structure left by the base architecture and on the chosen training objective.

To ratify the robustness of our method, we also evaluated our regression results with a complementary metric (MSE), as described in the **Supplementary Material** (available online). Typically, MSE and PCC capture different aspects of performance in regression tasks. The former is mostly associated with numerical accuracy, which is important in forecasting tasks, while the latter analyzes directions and trends of predictions, regardless of scale. This is particularly useful in our case, as we are predicting four different property fields.

5.4. Relevance of Engineered Features

Figure 7 illustrates the effects of including the three engineered features described in Section 4 along with the four original input properties. To provide a clear visualization of these effects, we will only discuss the results with respect to Level 0 networks. From the plots, we notice that for both ConvLSTM and U-FNO, there is a performance drop when transitioning from 7 to 4 input features, except for U_x predictions for ConvLSTM. Moreover, U-FNO with 4 input features achieved better results than ConvLSTM with 7 input features.


Conversely, TAU yielded slightly better performances without the engineered features, also achieving superior results than the other two methods. However, a further analysis with respect to its internal parameters (including the α term of the loss function) must be conducted to confirm whether the engineered features are actually relevant. Nevertheless, the plots demonstrate TAU's capability to understand the evolution of the reactive dissolution process with a smaller set of input features.



5.5. Qualitative Analysis

In this section, we will conduct a qualitative analysis on a sample from the validation set, comprising a particular pore geometry. Herein, we will consider the predictions at the last time step (100) after several iterative steps are performed for each algorithm. To simplify our analysis, we only show the predictions for C and eps fields. **Figure 8** shows the ground truth maps for those two properties at time steps 0 and 100.

The predictions of C for all algorithms and all network levels are displayed in **Figure 9**, along with their respective difference maps to the ground truth ($\hat{Y} - Y$). A clear evolution can be observed in the ConvLSTM results, as the Level 0 results contain too much noise, which is largely mitigated over the subsequent levels. However, there is a slight decrease in the PCC score from Level 2 to Level




3. Regarding U-FNO, the best results were achieved at Level 0, which yielded a PCC of 0.004 lower than ConvLSTM Level 2. Moreover, all levels produced similar shapes of dissolution channel, only missing the right branching at the bottom-center part of the map. Finally, although not having benefitted from the multi-level stacking approach, TAU achieved the highest scores for all levels, where even its Level 0 network performed better than any other level from both ConvLSTM and U-FNO. Looking at the difference maps, we can also notice a smaller range of errors compared to the other two algorithms. With respect to the dissolution shape, it was able to capture some of the right branching, but not as much as ConvLSTM did.

Figure 10 shows the predictions for ϵ_{ps} at time step 100. In general, all networks managed to yield nearly perfect predictions when compared to the ground truth. The PCC scores for ConvLSTM strictly increase over all levels. At Level 0, we can notice a prevalence of overpredictions all over the map, which is mitigated at the subsequent levels. On the other hand, neither U-FNO nor TAU showed significant differences from one level to another, with the latter achieving the best PCC scores at each level, having only tied with ConvLSTM at Level 3.

5.6. Estimation of Bulk Properties

To provide a better understanding of the overall dissolution dynamics in our case study, we analyzed the evolution of two of the bulk properties used for modelling reactive dissolution at the field-scale: porosity and permeability. Such an analysis is intended to check the physics-consistency of our trained models, rather than to directly learn from those bulk properties. For




each sample from the validation set, we selected 10 ϵ_{ps} maps, starting at time step 5, with an interval of 10 time steps between each pair of consecutive maps, ending at time step 95. We repeated this process for the ϵ_{ps} maps predicted by each of the DL algorithms at the aforementioned time steps. The porosity and permeability values for all maps were then calculated using GeoChemFoam, and the results achieved by each algorithm were compared against those obtained from the original data.

For all cases discussed in this section, we will show how the errors in porosity and permeability for each algorithm evolve due to the dissolution process in terms of: 1) their respective average values; and 2) the root mean squared error (RMSE) versus the values from the ground truth (referred to as "Original Data" in the subsequent plots).

5.6.1. Porosity Estimation

Figure 11 shows the average error in porosity considering all models from the training set. To enhance the overall visualization of the average error curves, only the values from time step 50 onwards are shown. Looking at the average evolution over the sampled time steps (Fig. 11a), we observe that the curves from U-FNO and TAU are the closest from the original porosity values, especially at late time steps, where the error tends to be higher due to the error accumulation problem of the iterative approach. Conversely, the curves from ConvLSTM are farther away from the ground truth, even though they show some improvement over the network levels. In the RMSE




curves (Fig. 11b), we observe that TAU Level 3 achieved the lowest errors at time steps 25 and 35, but ended up with a higher error than its Level 0 counterpart and all U-FNO variants (except U-FNO Level 0). Both ConvLSTM and U-FNO showed a more consistent evolution over the network levels, with U-FNO Level 3 achieving the lowest errors at the end of the dissolution.


For the validation set, similar results as the training set are achieved for the average error plots (Fig. 12a). Here, TAU Level 3 was the closest curve to the ground truth at the end of the dissolution, although the curves from its remaining levels and the ones from U-FNO were slightly farther away, but still constantly lying inside the error interval. Concerning the RMSE plots (Fig. 12b), an error reduction can also be observed at higher network levels for all algorithms, especially after time step 45. This reduction is more noticeable for ConvLSTM, whose Level 2 network achieved significantly lower results than its Level 0 counterpart, at the same time it yielded the lowest errors between time steps 25 and 55. For TAU, very close results were observed among all levels, with Level 2 achieving the lowest error rates at late time steps, but still performing slightly worse than ConvLSTM Level 2. Finally, U-FNO showed the highest errors at early time steps (5 to 25), while achieving the lowest error rates for late time steps (65 to 95).

5.6.2. Permeability Estimation

Figure 13 shows the results for the permeability estimation on the training set. From the curves in both plots, it can be observed that the multi-level stacking was insufficient to produce a consistent




(a) Absolute Porosity Evolution (Training Set)




(b) Porosity Error Evolution vs. Original Data (Training Set)

Figure 11: Porosity error analysis, considering the averages for each time step across all samples from the training set. (a) Absolute porosity evolutions vs. ground truth (with error margin bounds in gray); (b) RMSE evolutions. *ConvLSTM*: convolutional long short-term memory networks; *U-FNO*: U-shaped Fourier neural operator; *TAU*: temporal attention units.



(a) Absolute Porosity Evolution (Validation Set)



(b) Porosity Error Evolution vs. Original Data (Validation Set)

Figure 12: Porosity error analysis, considering the averages for each time step across all samples from the training validation set. (a) Absolute porosity evolutions vs. ground truth (with error margin bounds in gray); (b) RMSE evolutions. *ConvLSTM*: convolutional long short-term memory networks; *U-FNO*: U-shaped Fourier neural operator; *TAU*: temporal attention units.

evolution over the levels for all algorithms. This is corroborated by the fact that the closest curves to the ground truth (**Fig. 13a**) were obtained from TAU Level 0 and U-FNO Level 0. Moreover, when analyzing the RMSE curves for each algorithm (**Fig. 13b**), ConvLSTM Level 3 yielded the worst results during all the dissolution steps. Regarding the TAU curves, the Level 3 network achieved the lowest error rates until the last time step, at which point the Level 0 network performed slightly better. For the U-FNO, the Level 3 network was the best among all networks between time steps 25 and 65, being later surpassed by its Level 1 counterpart.

With respect to the estimation on the validation set, the evolution curves from **Figure 14a** show that the U-FNO curves were the closest ones to the ground truth. Unlike the training set, the TAU curves were farther away from the ground truth, where the Level 0 curve was the only one to remain inside the error margin bounds during all the dissolution. For ConvLSTM, the best results were achieved by its Level 0 network.


The discrepancy between the high voxel-wise agreement observed for the *eps* maps and the less systematic behavior of the derived permeability estimates is physically plausible. Although the predicted *eps* fields often achieve very high similarity and correlation scores, permeability is a highly non-linear and connectivity-sensitive bulk property that can be strongly affected by localized errors in narrow flow pathways, channel connectivity, and hydraulically important throat regions. Consequently, small spatial discrepancies in critical regions may lead to non-monotonic changes in permeability, even when the overall voxel-wise agreement remains high. Since the networks in this work were not trained with an explicit objective on bulk properties, improvements in local field prediction do not necessarily translate into monotonic improvements in macroscopic flow estimates across stacking levels.

Considering the RMSE curves from **Figure 14b**, the lowest errors were achieved by ConvLSTM Level 0 (time steps 15 to 45) and TAU Level 0 (time steps 55 to 95). Furthermore, ConvLSTM Level 3 produced the highest errors between time steps 45 to 95. For the same interval, U-FNO Level 3 achieved the lowest results among all its levels, although still being slightly worse than TAU Level 0.


6. CONCLUSIONS

This paper presented a data-driven method that leverages deep learning to predict the temporal evolution of reactive dissolution in porous media within a controlled benchmark setting. An iterative stacked MIMO approach was adopted to produce the future states of a dissolution process, starting with an initial set of “perfect” inputs, generating an output which is then used as input to predict the subsequent states, with this process repeated iteratively. To mitigate the overall errors of the predictions, a multi-level stacking approach was proposed, where each level is trained to correct the errors produced by the previous level network. Three different algorithms (ConvLSTM, U-FNO and TAU) were tested in a dataset comprised of 32 numerical simulation models.

Although error accumulation in recursive strategies for prediction of time-series data is still an open problem, all algorithms showed high correlation scores, especially regarding the predictions of *C* and *eps*, even at late time steps. Moreover, the multi-stacking approach was successful at improving the results from the base model (Level 0) of each algorithm for the majority of the cases. On the other hand, TAU did not benefit so much from this pipeline, which emphasizes the need for further investigation on the weight of the regularization term of its loss function. Even so, its Level 0 network was capable of achieving higher correlations than the other algorithms (and their corrections) for all predicted properties (except U_x), at the same time it achieved faster training and forward times.




(a) Absolute Permeability Evolution (Training Set)




(b) Permeability Error Evolution vs. Original Data (Training Set)

Figure 13: Permeability error analysis, considering the averages for each time step across all samples from the training set. (a) absolute permeability evolutions vs. ground truth (with error margin bounds in grey); (b) RMSE evolutions. *ConvLSTM*: convolutional long short-term memory networks; *U-FNO*: U-shaped Fourier neural operator; *TAU*: temporal attention units.



(a) Average Permeability Evolution (Validation Set)



(b) Permeability Error Evolution vs. Original Data (Validation Set)

Figure 14: Permeability error analysis, considering the averages for each time step across all samples from the validation set. (a) absolute permeability evolutions vs. ground truth (with error margin bounds in grey); (b) RMSE evolutions. *ConvLSTM*: convolutional long short-term memory networks; *U-FNO*: U-shaped Fourier neural operator; *TAU*: temporal attention units.

Despite the high correlation scores for *eps* prediction, there is still room for improvement in estimating bulk properties from the predicted *eps* maps, which would bridge the gap between pore-scale interactions and macroscopic flow and transport behavior and provide a broader understanding of such phenomena in porous media. This is particularly true when assessing the error rates for porosity and permeability estimation, which were expected to decrease over each network level, following similar patterns to the *eps* correlation plots. However, as none of the networks was calibrated to be aware of the overall bulk properties at a given time step, the evolution of those error rates ended up showing a pattern that was uncorrelated with the *eps* predictions. Nevertheless, all algorithms managed to produce low-magnitude error rates for porosity and permeability, and similar evolutions to their respective ground truths, considering an "average pore geometry" across all samples from our dataset. Hence, within the benchmark setting considered here, our method shows strong potential as a fast surrogate to complement traditional numerical solvers, especially when taking into account the reported speed-up and lowered computational expense.

Future directions for this work include: 1) analysis of gradient accumulation to improve iterative predictions at late time steps; 2) in-depth study of bulk property estimation; 3) application of the proposed method in larger-scale domains and in 3-dimensional simulations, which only requires a change of the base architecture, preserving the original idea of multi-level stacking.

STATEMENTS AND DECLARATIONS

Supplementary Material

All supplementary or supporting information provided by the authors will be published online alongside the main article. When such material is available, it will also be indicated here along with a link to the pdf for downloading.

Data, Code & Protocol Availability

Data Availability Statement: The source code used to reproduce all results for our iterative stacked method, including pre-trained models for all ML algorithms described in this work, can be found at <https://github.com/ai4netzero/ReactiveDissolution>. The supporting dataset for reactive dissolution is publicly available at <https://zenodo.org/records/14974428> under the Creative Commons Attribution International 4.0 license (5).

Software Availability Statement: The experiments on porosity and permeability estimation were run on version 5.1 of GeoChemFoam, available at <https://github.com/GeoChemFoam> under the GNU General Public License (GPL-3.0) (28).

Author Contributions

Marcos Cirne: Investigation, Methodology, Software, Validation, Visualization, Writing - original draft. **Hannah P. Menke:** Conceptualization, Data Curation, Supervision, Validation, Writing - Review & Editing. **Alhasan Abdellatif:** Validation, Writing - Review & Editing. **Julien Maes:** Software, Validation, Writing - Review & Editing. **Florian Doster:** Conceptualization, Validation, Writing - Review & Editing. **Ahmed Elsheikh:** Conceptualization, Funding Acquisition, Investigation, Methodology, Project Administration, Resources, Supervision, Validation, Writing - Review & Editing.

Conflicts of Interest

There are no conflicts of interest to disclose.

Funding Received

This work is funded by the Engineering and Physical Sciences Research Council's ECO-AI Project grant (reference number EP/Y006143/1), with additional financial support from the PETRONAS Centre of Excellence in Subsurface Engineering and Energy Transition (PACESET).

ORCID IDs

Marcos Cirne	 https://orcid.org/0000-0001-7266-9211
Hannah P. Menke	 https://orcid.org/0000-0002-1445-6354
Alhasan Abdellatif	 https://orcid.org/0000-0002-4547-378X
Julien Maes	 https://orcid.org/0000-0002-3248-9758
Florian Doster	 https://orcid.org/0000-0001-7460-573X
Ahmed Elsheikh	 https://orcid.org/0000-0003-3580-7309

REFERENCES

1. Raghad Al-Shabandar, Ali Jaddoa, Panos Liatsis, and Abir Jaafar Hussain. A deep gated recurrent neural network for petroleum production forecasting. *Machine Learning with Applications*, 3:100013, 2021. doi:10.1016/j.mlwa.2020.100013.
2. Naif Alqahtani, Ryan T Armstrong, and Peyman Mostaghimi. Deep learning convolutional neural networks to predict porous media properties. In *SPE Asia Pacific Oil and Gas Conference and Exhibition*, page D012S032R010. SPE, 2018. doi:10.2118/191906-MS.
3. Dzmitry Bahdanau. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. arXiv:1409.0473.
4. Meiling Cheng, Fangxin Fang, Ionel M Navon, and Christopher Pain. Ensemble Kalman filter for GAN-ConvLSTM based long lead-time forecasting. *Journal of Computational Science*, 69:102024, 2023. doi:10.1016/j.jocs.2023.102024.
5. Marcos Cirne, Hannah Menke, Alhasan Abdellatif, Julien Maes, Florian Doster, and Ahmed H. Elsheikh. Reactive dissolution in porous media [Dataset], 2025. doi:10.5281/zenodo.14974428.
6. Ying Da Wang, Martin J Blunt, Ryan T Armstrong, and Peyman Mostaghimi. Deep learning in pore scale imaging and modeling. *Earth-Science Reviews*, 215:103555, 2021. doi:10.1016/j.earscirev.2021.103555.
7. Can Ding, Yiyuan Zhou, Guang Pu, and Hongrong Zhang. Low carbon economic dispatch of power system at multiple time scales considering GRU wind power forecasting and integrated carbon capture. *Frontiers in Energy Research*, 10:953883, 2022. doi:10.3389/fenrg.2022.953883.
8. Honghui Du, Ze Zhao, Haojia Cheng, Jinhui Yan, and QiZhi He. Modeling density-driven flow in porous media by physics-informed neural networks for CO₂ sequestration. *Computers and Geotechnics*, 159:105433, 2023. doi:10.1016/j.compgeo.2023.105433.
9. Mohammad Hemmat Esfe and Saeed Esfandeh. 3D numerical simulation of the enhanced oil recovery process using nanoscale colloidal solution flooding. *Journal of Molecular Liquids*, 301:112094, 2020. doi:10.1016/j.molliq.2019.112094.
10. Zhao Feng, Zeeshan Tariq, Xianda Shen, Bicheng Yan, Xuhai Tang, and Fengshou Zhang. An encoder-decoder ConvLSTM surrogate model for simulating geological CO₂ sequestration with dynamic well controls. *Gas Science and Engineering*, 125:205314, 2024. doi:10.1016/j.jgsce.2024.205314.
11. Zhangyang Gao, Cheng Tan, Lirong Wu, and Stan Z Li. SimVP: Simpler yet better video prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3170–3180, 2022. doi:10.1109/cvpr52688.2022.00317.
12. Paul Garnier, Jonathan Viquerat, Jean Rabault, Aurélien Larcher, Alexander Kuhnle, and Elie Hachem. A review on deep reinforcement learning for fluid mechanics. *Computers & Fluids*, 225:104973, 2021. doi:10.1016/j.compfluid.2021.104973.
13. Krzysztof M Graczyk and Maciej Matyka. Predicting porosity, permeability, and tortuosity of porous media from images by deep learning. *Scientific Reports*, 10(1):21488, 2020. doi:10.1038/s41598-020-78415-x.
14. QiZhi He, David Barajas-Solano, Guzel Tartakovsky, and Alexandre M Tartakovsky. Physics-informed neural networks for multiphysics data assimilation with application to subsurface transport. *Advances in Water Resources*, 141:103610, 2020. doi:10.1016/j.advwatres.2020.103610.

15. Niklas Heinemann, Juan Alcalde, Johannes M. Miocic, Suzanne J. T. Hangx, Jens Kallmeyer, Christian Ostertag-Henning, Aliakbar Hassanpouryouzband, Eike M. Thaysen, Gion J. Strobel, Cornelia Schmidt-Hattenberger, Katriona Edlmann, Mark Wilkinson, Michelle Benthams, R. Stuart Haszeldine, Ramon Carbonell, and Alexander Rudloff. Enabling large-scale hydrogen storage in porous media—the scientific challenges. *Energy & Environmental Science*, 14(2):853–864, 2021. doi:10.1039/D0EE03536J.
16. Priyesh Rajesh Kakka. Sequence to sequence AE-ConvLSTM network for modelling the dynamics of PDE systems. *arXiv preprint arXiv:2208.07315*, 2022. arXiv:2208.07315.
17. J Nagoor Kani and Ahmed H Elsheikh. DR-RNN: A deep residual recurrent neural network for model reduction. *arXiv preprint arXiv:1709.00939*, 2017. arXiv:1709.00939.
18. Mina Karimi and Kaushik Bhattacharya. A learning-based multiscale model for reactive flow in porous media. *Water Resources Research*, 60(9), 2024. doi:10.1029/2023WR036303.
19. Ouassim Khebzegga, Alireza Iranshahr, and Hamdi Tchelepi. Continuous relative permeability model for compositional simulation. *Transport in Porous Media*, 134:139–172, 2020. doi:10.1007/s11242-020-01440-x.
20. Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. arXiv:1412.6980.
21. Dmitrii Kochkov, Jamie A Smith, Ayya Alieva, Qing Wang, Michael P Brenner, and Stephan Hoyer. Machine learning-accelerated computational fluid dynamics. *Proceedings of the National Academy of Sciences*, 118(21):e2101784118, 2021. doi:10.1073/pnas.2101784118.
22. Arief Koesdwiady, Alaa El Khatib, and Fakhri Karray. Methods to improve multi-step time series prediction. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018. doi:10.1109/IJCNN.2018.8489402.
23. Angran Li, Ruijia Chen, Amir Barati Farimani, and Yongjie Jessica Zhang. Reaction diffusion system prediction based on convolutional neural network. *Scientific Reports*, 10(1):3894, 2020. doi:10.1038/s41598-020-60853-2.
24. Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020. arXiv:2010.08895, doi:10.2139/ssrn.4944664.
25. Suu-Yan Liang, Wen-Sheng Lin, Chan-Po Chen, Chen-Wuing Liu, and Chihhao Fan. A review of geochemical modeling for the performance assessment of radioactive waste disposal in a subsurface system. *Applied Sciences*, 11(13):5879, 2021. doi:10.3390/app11135879.
26. Bryan Lim and Stefan Zohren. Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A*, 379(2194):20200209, 2021. doi:10.1098/rsta.2020.0209.
27. Vishnu Mohan M S and Vivek Menon. Measuring viscosity of fluids: A deep learning approach using a CNN-RNN architecture. In *Proceedings of the First International Conference on AI-ML Systems, AIMLSys* '21, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3486001.3486232.
28. Julien Maes, Cyprien Soullaine, and Hannah P Menke. Improved volume-of-solid formulations for micro-continuum simulation of mineral dissolution at the pore-scale. *Frontiers in Earth Science*, 10:917931, 2022. doi:10.3389/feart.2022.917931.
29. Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. Statistical and Machine Learning forecasting methods: Concerns and ways forward. *PloS One*, 13(3):e0194889, 2018. doi:10.1371/journal.pone.0194889.
30. Nima Mohajerin and Steven L Waslander. Multistep prediction of dynamic systems with recurrent neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 30(11):3370–3383, 2019. doi:10.1109/TNNLS.2019.2891257.
31. Tung Nguyen, Rohan Shah, Hritik Bansal, Troy Arcomano, Sandeep Madireddy, Romit Maulik, Veerabhadra Kotamarthi, Ian Foster, and Aditya Grover. Scaling transformer neural networks for skillful and reliable medium-range weather forecasting. *arXiv preprint arXiv:2312.03876*, 2023. arXiv:2312.03876, doi:10.52202/079017-2196.
32. Sergiu Oprea, Pablo Martinez-Gonzalez, Alberto Garcia-Garcia, John Alejandro Castro-Vargas, Sergio Orts-Escolano, Jose Garcia-Rodriguez, and Antonis Argyros. A review on deep learning techniques for video prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):2806–2826, 2020. doi:10.1109/TPAMI.2020.3045007.
33. Markus Reichstein, Gustau Camps-Valls, Bjorn Stevens, Martin Jung, Joachim Denzler, Nuno Carvalhais, and Prabhat. Deep learning and process understanding for data-driven Earth system science. *Nature*, 566(7743):195–204, 2019. doi:10.1038/s41586-019-0912-1.
34. S Salimzadeh and HM Nick. A coupled model for reactive flow through deformable fractures in enhanced geothermal systems. *Geothermics*, 81:88–100, 2019. doi:10.1016/j.geothermics.2019.04.010.

35. Javier E Santos, Duo Xu, Honggeun Jo, Christopher J Landry, Maša Prodanović, and Michael J Pyrcz. PoreFlow-Net: A 3D convolutional neural network to predict fluid flow through porous media. *Advances in Water Resources*, 138:103539, 2020. doi:10.1016/j.advwatres.2020.103539.
36. Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-kin Wong, and Wang-chun Woo. Convolutional lstm network: a machine learning approach for precipitation nowcasting. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'15, page 802–810, Cambridge, MA, USA, 2015. MIT Press. URL: <https://dl.acm.org/doi/10.5555/2969239.2969329>.
37. Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. Sequence to sequence learning with neural networks. *arXiv preprint arXiv:1409.3215*, 2014. arXiv:1409.3215.
38. Souhaib Ben Taieb, Gianluca Bontempi, Amir F Atiya, and Antti Sorjamaa. A review and comparison of strategies for multi-step ahead time series forecasting based on the NN5 forecasting competition. *Expert Systems with Applications*, 39(8):7067–7083, 2012. doi:10.1016/j.eswa.2012.01.039.
39. Cheng Tan, Zhangyang Gao, Lirong Wu, Yongjie Xu, Jun Xia, Siyuan Li, and Stan Z Li. Temporal attention unit: Towards efficient spatiotemporal predictive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18770–18782, 2023. doi:10.1109/cvpr52729.2023.01800.
40. Meng Tang, Yimin Liu, and Louis J Durlowsky. Deep-learning-based surrogate flow modeling and geological parameterization for data assimilation in 3D subsurface flow. *Computer Methods in Applied Mechanics and Engineering*, 376:113636, 2021. doi:10.1016/j.cma.2020.113636.
41. Wendong Wang, Qiuhe Xie, Senyou An, Sahar Bakhshian, Qinqun Kang, Han Wang, Xinyue Xu, Yuliang Su, Jianchao Cai, and Bin Yuan. Pore-scale simulation of multiphase flow and reactive transport processes involved in geologic carbon sequestration. *Earth-Science Reviews*, page 104602, 2023. doi:10.1016/j.earscirev.2023.104602.
42. Gege Wen, Zongyi Li, Kamyar Azizzadenesheli, Anima Anandkumar, and Sally M Benson. U-FNO—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163:104180, 2022. doi:10.1016/j.advwatres.2022.104180.
43. Lei Xu, Nengcheng Chen, Zeqiang Chen, Chong Zhang, and Hongchu Yu. Spatiotemporal forecasting in earth system science: Methods, uncertainties, predictability and future directions. *Earth-Science Reviews*, 222:103828, 2021. doi:10.1016/j.earscirev.2021.103828.
44. Bicheng Yan, Dylan Robert Harp, Bailian Chen, and Rajesh Pawar. A physics-constrained deep learning model for simulating multiphase flow in 3D heterogeneous porous media. *Fuel*, 313:122693, 2022. doi:10.1016/j.fuel.2021.122693.
45. Li-Tao Zhu, Xi-Zhong Chen, Bo Ouyang, Wei-Cheng Yan, He Lei, Zhe Chen, and Zheng-Hong Luo. Review of machine learning for hydrodynamics, transport, and reactions in multiphase flows and reactors. *Industrial & Engineering Chemistry Research*, 61(28):9901–9949, 2022. doi:10.1021/acs.iecr.2c01036.