



Accepted Articles have been peer-reviewed and accepted for publication in *InterPore Journal*. To make research available as early as possible, accepted manuscripts may be published online prior to completion of the final production process. These versions have not yet undergone final formatting, proofreading, or author corrections and therefore do not represent the **Version of Record**. Minor changes may occur during the final production process. The article will be updated with the final Version of Record when it is published in its assigned journal issue. **Accepted Articles should be cited using their DOI, which will remain the same after publication of the final Version of Record.**

Please cite the following article as:

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

Received on: 15 Dec. 2025

Accepted on: 1 Apr. 2026

Published online (Accepted version): 20 Apr. 2026

1 **A Deep-Learning Iterative Stacked Approach for**
2 **Prediction of Reactive Dissolution in Porous Media**

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

5 ¹School of Energy, Geoscience, Infrastructure and Society (EGIS), Heriot-Watt University
6 ¹EH14 4AS, Edinburgh, United Kingdom

7 **Key Points:**

- 8 • Data-driven deep-learning approach for predicting multiple future states of reactive
9 dissolution in porous media using an iterative strategy;
10 • Proposal of a multi-level network stacking pipeline, where each level is trained to
11 minimize the errors produced by its predecessor;
12 • Comparative analysis among several deep learning algorithms, which are trained on
13 an ensemble of numerical simulation models, encompassing various pore structures
14 and fluid trajectories.

Corresponding author: Marcos Cirne, M.Cirne@nc1.ac.uk

Corresponding author: Ahmed Elsheikh, A.Elsheikh@hw.ac.uk

Abstract

Simulating reactive dissolution of solid minerals in porous media has many subsurface applications, including carbon capture and storage (CCS), geothermal systems and oil & 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.

1 Introduction

Numerical solvers have been extensively used to simulate and understand the effects of reactive dissolution of solid minerals in subsurface porous media, in diverse applications such as CO_2 sequestration (Wang et al., 2023), hydrogen storage (Heinemann et al., 2021), enhanced oil recovery (Esfe & Esfandeh, 2020), radioactive waste disposal (Liang et al., 2021) and geothermal systems (Salimzadeh & Nick, 2019). 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 (Khebzegga et al., 2020). 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 (Zhu et al., 2022; Garnier et al., 2021; Da Wang et al., 2021; Kochkov et al., 2021).

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. CNN-based methods (Alqahtani et al., 2018; Graczyk & Matyka, 2020; A. Li et al., 2020; Santos et al., 2020; Tang et al., 2021), 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) (Yan et al., 2022; Du et al., 2023; He et al., 2020) 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) (Mohajerin & Waslander, 2019; MS & Menon, 2021), convolutional long short-term networks (ConvLSTMs) (Cheng et al., 2023; Feng et al., 2024) and gated recurrent units (GRU) (Ding et al., 2022; Al-Shabandar et al., 2021). A more recent approach named recurrent neural operator (RNO) (Karimi & Bhattacharya, 2024) 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. Other have adopted a hybrid approach, such as Reichstein et al. (2019), who leverages physical process models along with data-driven machine-learning 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 (Koesdwiady et al., 2018). 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 the prediction errors is by stacking multiple networks and perform an iterative process to minimize the overall residual error of the predictions (Kani & Elsheikh, 2017). 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 way, accounting for spatial and temporal features, as a surrogate model to forecast the 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 of Maes et al. (2022). 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 predict an output sequence $\mathcal{Y} = \hat{X}_{t+1:t+n}$ based on the input \mathcal{X} . In other words, we want to train a deep learning model \mathcal{F} with learnable parameters Θ which learns a mapping $\mathcal{F}_\Theta : \mathcal{X} \mapsto \mathcal{Y}$, where \mathcal{Y} represent the ground truth states, trying to minimize 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 as **multi-step forecasting** the cases where $n > 1$, and as **single-step forecasting** when $n = 1$.

2.2 Forecasting Strategies

An in-depth analysis of possible forecasting strategies for multivariate time-series data can be found in Lim and Zohren (2021). Among the existing strategies, there is **direct forecasting**, in which a DL model is trained to forecast each of the n future steps (Makridakis et al., 2018), 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 (Taieb et al., 2012).

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 in time-series forecasting, this strategy has been widely used for spatiotemporal forecasting tasks, such as video prediction (Gao et al., 2022; Oprea et al., 2020) and earth science forecasting (Xu et al., 2021; Nguyen et al., 2023). 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) (Shi et al., 2015) are a class of neural networks designed to capture spatiotemporal dependencies in a time-evolving sequence, combining the strengths of CNNs and 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

151 both encoder and decoder blocks contain the same number of ConvLSTM cells. This type
 152 of architecture is inspired on sequence-to-sequence (Seq2Seq) models for natural language
 153 processing and time-series tasks (Bahdanau, 2014; Sutskever et al., 2014). An example is
 154 depicted in Figure 1, with two ConvLSTM cells in each block. After an input sequence
 155 \mathcal{X} is processed by the encoder block, the final hidden state from the last ConvLSTM cell
 156 forms a compact representation of \mathcal{X} (latent representation), which is fed to the decoder
 157 block to produce the hidden states for each output time step. Finally, a 3D convolutional
 158 layer receives the output from the decoder block and produces the predictions of all phys-
 159 ical properties for each desired time step. More details about the encoding and decoding
 160 processes can be found in (Kakka, 2022).




Figure 1: Encoder-Decoder structure of ConvLSTM cells for multi-step prediction.

161 2.3.2 U-Shaped Fourier Neural Operator

162 Proposed by Wen et al. (2022), the U-Shaped Fourier Neural Operator (U-FNO) is an
 163 extension of the Fourier neural operator (FNO) (Z. Li et al., 2020), designed to solve PDEs
 164 across diverse problems involving computational fluid dynamics. FNOs are known to be
 165 resolution-invariant, meaning that they can be trained on a lower resolution and evaluated
 166 on higher resolution, and yield superior performance against conventional CNNs by operat-
 167 ing directly on the Fourier space (frequency domain), replacing convolution operations by
 168 pointwise multiplications, which are much faster and efficient.

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

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

179 2.3.3 Temporal Attention Unit

180 Developed for video prediction tasks, Temporal Attention Units (TAUs) (Tan et al.,
 181 2023) leverage the ability to capture time evolution in image sequences by introducing
 182 parallelizable attention mechanisms that eliminate the need of recurrent-based units (such
 183 as RNNs and ConvLSTMs), speeding up the training process. In turn, the spatial modules
 184 are represented by simple 2D convolutions. Moreover, not only TAUs account for intra-
 185 frame differences through the mean squared error loss, but also for inter-frame variations

186 by embedding a differential divergence regularization term. The resulting loss function is
 187 expressed in Equation 3:

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

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

190 3 Methodology

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

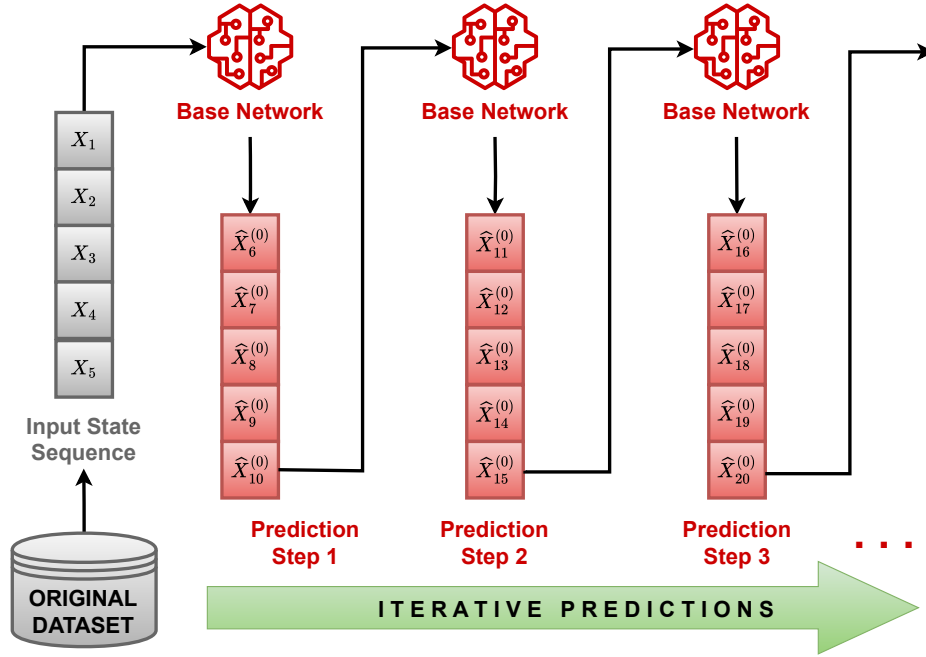


Figure 2: Flowchart of our iterative multi-step approach with network stacking for full evaluation of a reactive dissolution process on simulation data.

199 To improve the initial solution produced by the base network, we propose a multi-
 200 level stacking of trained networks. In this case, the same neural network architecture is
 201 used to train each level, with the same sizes for both input and output tensors as the base
 202 network, as well as the loss function. Considering the base network as Level 0, trained with
 203 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
 204 new dataset comprised of all possible output sequences $\hat{\mathcal{Y}}$ computed by the base network for
 205 each $\mathcal{S}_i \in \mathcal{D}$, representing initial approximations to the true states \mathcal{Y} . The Level 1 network

206 is then trained to receive, as input, one of those responses produced by the base network,
207 and outputs a new approximation (correction) to the ground truth within the same time-
208 step interval, hypothesizing that it will yield less errors than the initial solution. Following
209 this rationale, we can continue the stacking process by training a Level 2 network with the
210 outputs produced by Level 1, and so on until no further improvement is observed. During
211 this process, only one level can be trained at a time.

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

223 In general, let L be the number of levels of correction to be applied on the base solution.
224 In the present study, we evaluate corrections up to $L = 3$ as a practical upper bound to
225 examine whether improvements persist beyond the first correction while keeping the overall
226 training cost manageable; however, this value should be regarded as a tunable design choice
227 rather than an optimal or universal setting. Increasing the number of correction levels
228 may also increase the risk of overfitting, particularly when the residual errors become less
229 structured at higher levels. After the base and all of the L correction networks are trained,
230 the iterative multi-step prediction process can be conducted as illustrated in Figure 3 with
231 $L = 3$. Here, the first 5 input states of a simulation sample are fed into the base network,
232 producing an initial approximation for the subsequent 5 states. Then, this output is fed to
233 the Level 1 network which produces the first level of correction for the initial approximation.
234 Later, this output serves as input to Level 2 network, whose output is forwarded as an input
235 to Level 3 network, which yields the final prediction for time steps 6-10. For the next
236 prediction step, the output from the last network of the stack is fed to the base network
237 to produce an initial approximation for steps 11-15, which is corrected by the subsequent
238 levels in the stack. The whole process is repeated until the simulation end time.

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

244 4 Reactive Dissolution Dataset and Preprocessing

245 Our reactive dissolution dataset (Cirne et al., 2025) consists of 32 numerical simulations
246 generated using `GeoChemFoam` (Maes et al., 2022), an open-source micro-continuum solver
247 for pore-scale reactive transport. Each simulation follows the *Benchmark 2* configuration
248 presented by Maes et al. (2022), which models calcite dissolution in a quasi two-dimensional
249 micromodel. The computational domain consists of a polydisperse disk pack with an initial
250 porosity of approximately $\phi_0 = 0.45$ and an initial permeability of $K_0 = 5.6 \times 10^{-10} \text{ m}^2$.
251 Simulations are performed on a uniform Cartesian mesh with a spatial resolution of $\Delta x, y =$
252 $25 \mu\text{m}$ and a $\Delta z = 150 \mu\text{m}$.

253 The flow field is incompressible and governed by the Darcy-Brinkman-Stokes equa-
254 tions under isothermal conditions, while the transport of a single reactive species obeys an
255 advection-diffusion equation with a surface reaction rate $R = k_c c$ applied at the fluid-

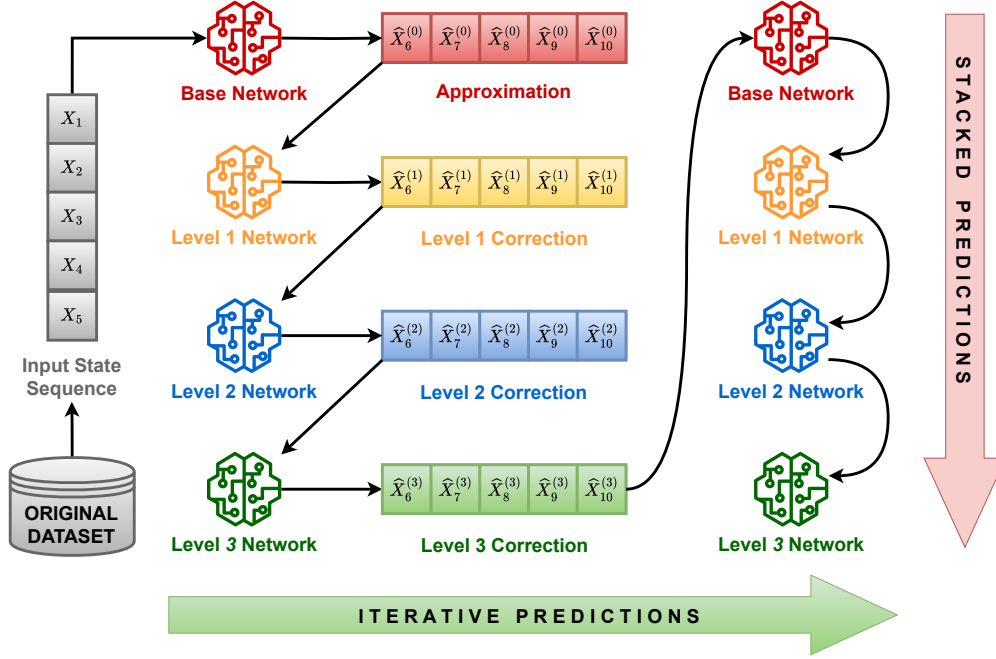


Figure 3: Flowchart of our iterative multi-step approach with network stacking for predicting reactive dissolution.

256 solid interface. A representative parameter set corresponding to the *conical wormhole*
 257 *regime* is as follows: kinematic viscosity $\nu = 1 \times 10^{-6} \text{ m}^2/\text{s}$, diffusion coefficient $D =$
 258 $1 \times 10^{-9} \text{ m}^2/\text{s}$, inlet concentration $c_i = 0.01 \text{ kmol}/\text{m}^3$, stoichiometric coefficient $\zeta = 1$, cal-
 259 crite density $\rho_s = 2710 \text{ kg}/\text{m}^3$, molar mass $M_{ws} = 100 \text{ kg}/\text{kmol}$, Kozeny–Carman constant
 260 $k_f = 2 \times 10^{11} \text{ m}^2$, a volumetric flow rate $Q = 5.2174 \times 10^{-12} \text{ m}^3/\text{s}$, and reaction rate con-
 261 stant $k_c = 8.1632 \times 10^{-6} \text{ m}/\text{s}$. These conditions correspond to Péclet and kinetic numbers of
 262 $\text{Pe} = 1$ and $\text{Ki} = 1$, respectively, representing the transition between compact and wormhole
 263 dissolution regimes. Boundary conditions follow Maes et al. (2022): a constant inflow rate
 264 at the inlet, fixed outlet pressure, and no-slip/no-flow boundaries on the top and bottom
 265 walls.

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

279 We note that, although each simulation provides a large number of spatial samples over
 280 time, the present dataset contains only 32 distinct pore geometries generated within a single
 281 benchmark family. Therefore, this dataset should be interpreted as a controlled test bed for

282 methodological evaluation rather than as an exhaustive representation of the broader class of
 283 pore structures encountered in reactive porous media. In our train/validation split, complete
 284 simulations (rather than individual time steps) were separated, so the validation set consists
 285 of unseen pore geometries that were not used during training. As a result, the reported
 286 performance reflects generalization to new realizations within the same benchmark configuration,
 287 but does not yet establish robustness across substantially different pore-structure classes,
 288 reaction regimes, or boundary conditions. Extending the dataset to cover broader
 289 geometric descriptors and multiple flow-reactive regimes remains an important direction for
 290 future work.




Figure 4: Input features from one dataset sample representing the dissolution states for time steps 0, 30, 60 and 90 (from left to right).

291 In essence, GeoChemFoam solves the reactive transport in porous media in a quasi-
 292 static regime, as described by Maes et al. (2022). Given a porosity map **eps** at a time
 293 step t , it first calculates the steady-state velocity (i.e., it produces \mathbf{U}_x and \mathbf{U}_y at time
 294 step t). Then, it calculates a steady-state concentration \mathbf{C} at time step t , which includes
 295 the contributions of injection, diffusion, and reactive dissolution at the fluid-solid interface.
 296 From these calculations, it estimates a local reaction rate in each computational fluid and
 297 the new **eps** field for the time step $t + 1$, which in turn is used to calculate the flow once
 298 again to produce a new **eps** map at time step $t + 2$, and so on.

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

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

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

305 5 Performance Evaluation of Deep Learning Models

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

309 5.1 Training and Evaluation Settings

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

318 Concerning our proposed multi-level stacking of neural networks method, for each DL
 319 algorithm described in Section 2.3, we perform corrections up to Level $L = 3$. The models
 320 were trained on a NVIDIA Titan RTX with 24 GB of memory for a total of 100 epochs
 321 (along with a patience rate of 20 epochs for early stopping), using a batch size of 4, a learning
 322 rate of 0.0005 and the Adam optimizer (Kingma, 2014) with moment values $\beta_1 = 0.9$ and
 323 $\beta_2 = 0.999$. The mean squared error was adopted as the loss function for ConvLSTM,
 324 U-FNO and the intra-frame difference term of TAU. Regarding the latter, the α constant
 325 for the regularization term (inter-frame difference) was fixed at 0.1.

326 For the quantitative evaluation of the iterative predictions of each output property,
 327 we calculate the Pearson correlation coefficient (PCC) to assess the similarity between the
 328 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
 329 stated in Equation 5:

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

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

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

339 5.2 Model Statistics

340 Table 1 shows some statistics about the average training epoch times, numbers of
 341 trainable parameters, final model sizes and average forward times (i.e., the elapsed time

342 for processing input data through a network to produce the next 5 time steps) for each
 343 DL model. Although TAU has the highest numbers of trainable parameters and a rela-
 344 tively larger model size, it still achieved the lowest training and forward times, due to its
 345 recurrence-free structure.

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

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

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

351 Analyzing the total runtime of each algorithm to produce a 100-step simulation, TAU
 352 takes an average time of 0.66 seconds, followed by ConvLSTM (1.2 seconds) and U-FNO
 353 (1.58 seconds). These represent a speedup with an order of magnitude between 10^3 and 10^4
 354 when compared to the average time taken by GeoChemFoam (3 hours).

355 5.3 Iterative Prediction and Model Stacking Analysis

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

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

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

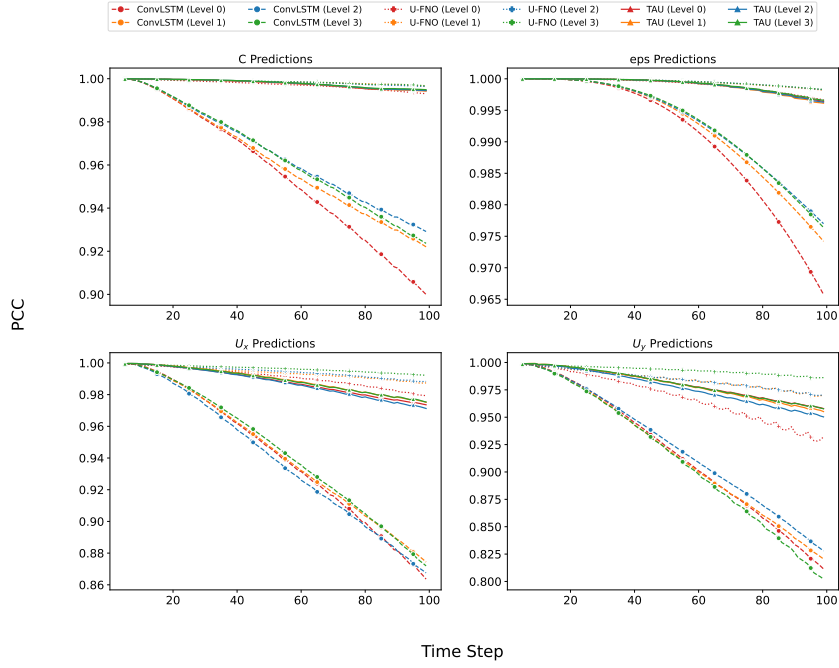


Figure 5: Average correlation scores of all samples from the training set for the iterative predictions produced by each algorithm.

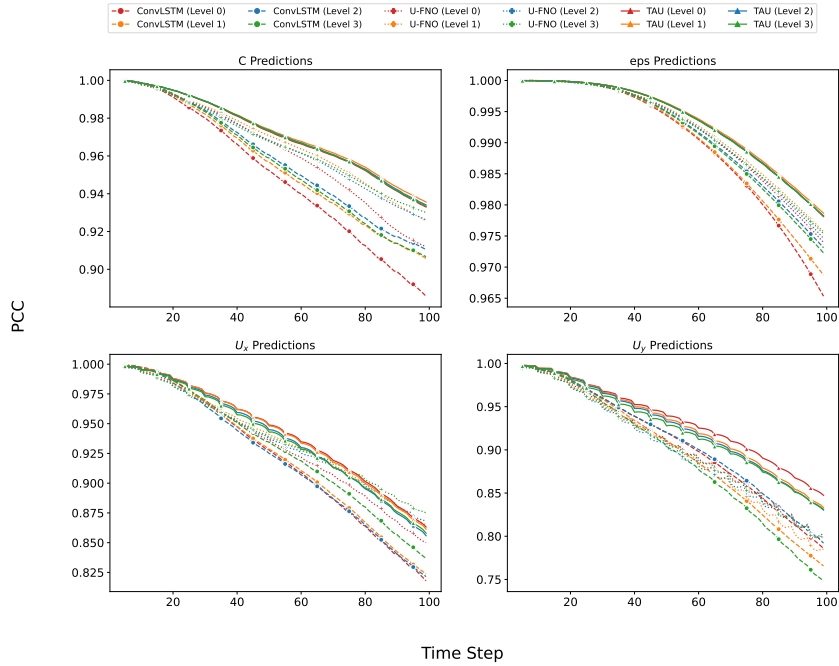


Figure 6: Average correlation scores of all samples from the validation set for the iterative predictions produced by each algorithm.

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

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

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

405 5.4 Relevance of Engineered Features

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

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

419 5.5 Qualitative Analysis

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

425 The predictions of C for all algorithms and all network levels are displayed in Figure 9,
426 along with their respective difference maps to the ground truth ($\hat{Y} - Y$). We can notice a
427 clear evolution on the ConvLSTM results, as the results at Level 0 contain too much noise,
428 which is mostly mitigated over the subsequent levels. However, there is a slight decrease in
429 the PCC score from Level 2 to Level 3. Regarding U-FNO, the best results were achieved




Figure 7: Average correlations for all models in the validation set and each output property, comparing the performances of all Level 0 networks with and without the engineered features, respectively represented by solid and dashed lines.

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

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

445 5.6 Estimation of Bulk Properties

446 To convey a better understanding of the overall dissolution dynamics in our case study,
 447 we analyze the evolution of two of the bulk properties used for modelling reactive dissolution
 448 at the field-scale: porosity and permeability. Such an analysis is intended to check the
 449 physics-consistency of our trained models, rather than an attempt to directly learn from
 450 those bulk properties. For each sample from our validation set, we selected 10 *eps* maps,
 451 starting at time step 5, with an offset of 10 time steps between each pair of consecutive
 452 maps, ending at time step 95. We repeated this process for the *eps* maps predicted by
 453 each of the DL algorithms at the aforementioned time steps. The porosity and permeability




Figure 8: Ground truth maps at time steps 0 and 100 for C and eps .

454 values for all maps are then calculated using GeoChemFoam, and the results achieved by
 455 each algorithm are compared against the ones obtained from the original data.

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

460 **5.6.1 Porosity Estimation**

461 Figure 11 shows the average error in porosity considering all models from the training
 462 set. To enhance the overall visualization of the average error curves, we will only display
 463 the values from time step 50 onwards. Looking at the average evolution over the sampled
 464 time steps (Figure 11a), we observe that the curves from U-FNO and TAU are the closest
 465 from the original porosity values, especially at late time steps, where the error tends to
 466 be higher due to the error accumulation problem of the iterative approach. Conversely,
 467 the curves from ConvLSTM are farther away from the ground truth, even though it shows
 468 some improvement over the network levels. In the RMSE curves (Figure 11b), we observe
 469 that TAU Level 3 achieved the lowest errors on time steps 25 and 35, but ends up with a
 470 higher error than its Level 0 counterpart and all U-FNO variants (except U-FNO Level 0).




Figure 9: Predictions of C at time step 100 for a case from the validation set, along with their respective PCC scores and difference maps to ground truth (Figure 8c). Best results for each network level are in bold.

471 Both ConvLSTM and U-FNO showed a more consistent evolution over the network levels,
472 whereas the U-FNO Level 3 having lowest errors at the end of the dissolution.

473 For the validation set, similar results as the training set are achieved for the average
474 error plots (Figure 12a). Here, TAU Level 3 was the closest curve to the ground truth
475 at the end of the dissolution, although the curves from its remaining levels and the ones
476 from U-FNO were slightly farther away, but still constantly lying inside the error interval.
477 Concerning the RMSE plots (Figure 12b), we can also notice an error reduction on higher
478 network levels for all algorithms, especially after time step 45. This reduction is more
479 noticeable for ConvLSTM, whose Level 2 network achieved significantly lower results than
480 its Level 0 counterpart, at the same time it yielded the lowest errors between time steps
481 25 and 55. TAU produced very close results among all levels, where Level 2 achieved the
482 lowest error rates at late time steps, but still being slightly worse than ConvLSTM Level 2.
483 Last, U-FNO showed the highest errors at early time steps (5 to 25), achieving the lowest
484 error rates for late time steps (65 to 95).

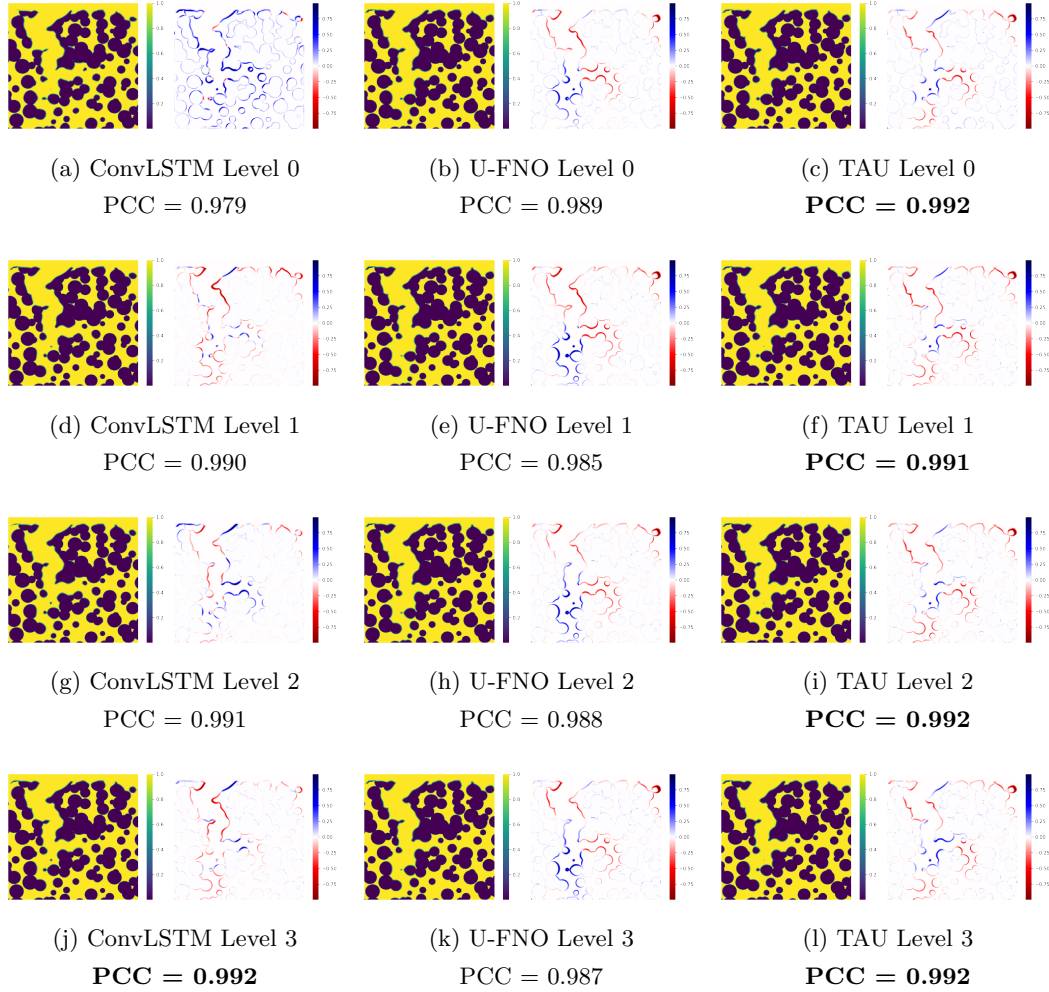


Figure 10: Predictions of ϵ_{ps} at time step 100 for a case from the validation set, along with their respective PCC scores and difference maps to ground truth (Figure 8d). Best results for each network level are in bold.

5.6.2 Permeability Estimation

485

486

487

488

489

490

491

492

493

494

495


Figure 13 showcases the results for the permeability estimation on the training set. From the curves in both plots, we can observe that the multi-level stacking was not enough to produce a consistent evolution over the levels for all algorithms. This is corroborated by the fact that the closest curves to the ground truth (Figure 13a) were obtained from TAU Level 0 and U-FNO Level 0. Moreover, when analyzing the RMSE curves for each algorithm (Figure 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, where the Level 0 network was 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.

496


497

498

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




(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 grey); (b) RMSE evolutions.




(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 grey); (b) RMSE evolutions.




(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.



(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.

499 curve was the only one to remain inside the error margin bounds during all the dissolution.
500 About ConvLSTM, the best results were achieved by its Level 0 network.

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

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

516 6 Conclusions

517 This paper presented a data-driven method that leverages deep learning to predict the
518 temporal evolution of reactive dissolution in porous media within a controlled benchmark
519 setting. An iterative stacked MIMO approach was adopted to produce the future states of
520 a dissolution process, starting with an initial set of "perfect" inputs, generating an output
521 which is then used as input to predict the subsequent states, and so on. To mitigate
522 the overall errors of the predictions, a multi-level stacking approach was proposed, where
523 each level is trained to correct the errors produced by the previous level network. Three
524 different algorithms (ConvLSTM, U-FNO and TAU) were tested in a dataset comprised of
525 32 numerical simulation models.

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

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

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

554 Appendix A Mean Squared Error (MSE) Scores for Iterative Prediction

555 To quantify the error magnitudes of the iterative predictions, we also conducted an
556 analysis of the evolution of MSE scores. Figures A1 and A2 show the MSE scores of the
557 iterative predictions on the training and validation sets, respectively. Compared to the
558 results discussed in Section 5.3, the ranking of all methods on both scenarios remained the
559 same as the ones produced by the PCC metric. These results indicate that TAU not only
560 has a higher linear relationship to the ground truth, but also yields the smallest errors on
561 its predictions among the tested algorithms.

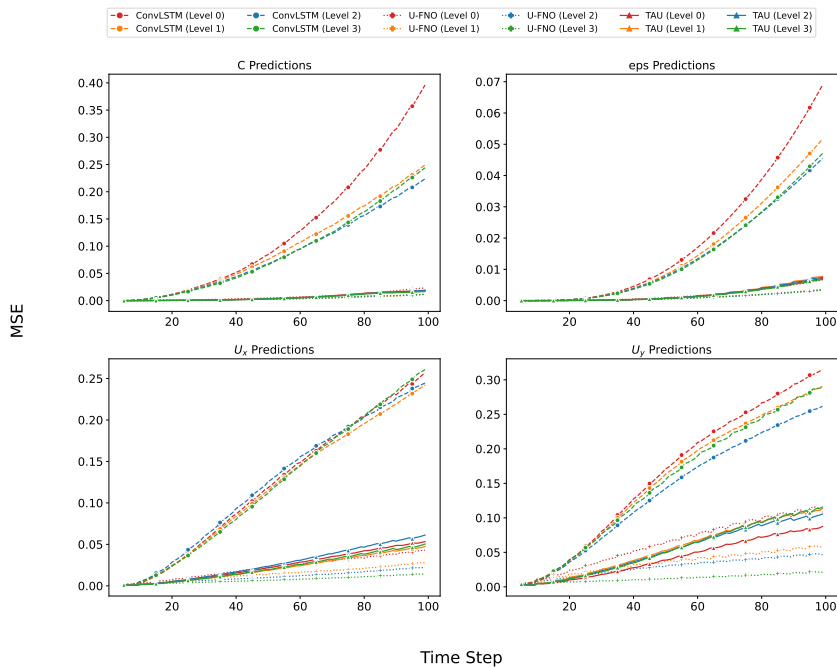


Figure A1: Average MSE scores of all samples from the training set for the iterative predictions produced by each algorithm.

562 Open Research Section

563 *Data Availability Statement:* The source code used to reproduce all results for our
564 iterative stacked method, including pre-trained models for all ML algorithms described
565 in this work, can be found at <https://github.com/ai4netzero/ReactiveDissolution>.
566 The supporting dataset for reactive dissolution is publicly available at <https://zenodo.org/records/14974428>
567 under the Creative Commons Attribution International 4.0 li-
568 cense (Cirne et al., 2025).

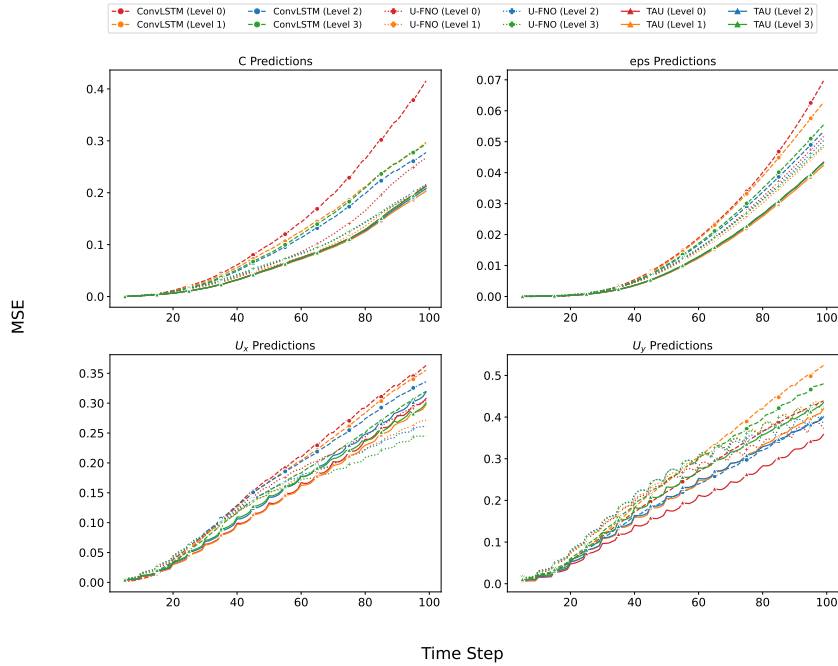


Figure A2: Average MSE scores of all samples from the validation set for the iterative predictions produced by each algorithm.

569 *Software Availability Statement:* The experiments on porosity and permeability estimation
 570 were run on version 5.1 of GeoChemFoam, available at [https://github.com/](https://github.com/GeoChemFoam)
 571 [GeoChemFoam](https://github.com/GeoChemFoam) under the GNU General Public License (GPL-3.0) (Maes et al., 2022).

572 Acknowledgments

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

577 References

- 578 Alqahtani, N., Armstrong, R. T., & Mostaghimi, P. (2018). Deep learning convolutional
 579 neural networks to predict porous media properties. In *SPE Asia Pacific Oil and Gas*
 580 *Conference and Exhibition* (p. D012S032R010).
- 581 Al-Shabandar, R., Jaddoa, A., Liatsis, P., & Hussain, A. J. (2021). A deep gated re-
 582 current neural network for petroleum production forecasting. *Machine Learning with*
 583 *Applications*, 3, 100013.
- 584 Bahdanau, D. (2014). Neural machine translation by jointly learning to align and translate.
 585 *arXiv preprint arXiv:1409.0473*.
- 586 Cheng, M., Fang, F., Navon, I. M., & Pain, C. (2023). Ensemble Kalman filter for GAN-
 587 ConvLSTM based long lead-time forecasting. *Journal of Computational Science*, 69,
 588 102024.
- 589 Cirne, M., Menke, H., Abdellatif, A., Maes, J., Doster, F., & Elsheikh, A. H. (2025).
 590 *Reactive dissolution in porous media [Dataset]*. Zenodo. Retrieved from [https://](https://doi.org/10.5281/zenodo.14974428)
 591 doi.org/10.5281/zenodo.14974428 doi: 10.5281/zenodo.14974428

- 592 Da Wang, Y., Blunt, M. J., Armstrong, R. T., & Mostaghimi, P. (2021). Deep learning in
593 pore scale imaging and modeling. *Earth-Science Reviews*, *215*, 103555.
- 594 Ding, C., Zhou, Y., Pu, G., & Zhang, H. (2022). Low carbon economic dispatch of power
595 system at multiple time scales considering GRU wind power forecasting and integrated
596 carbon capture. *Frontiers in Energy Research*, *10*, 953883.
- 597 Du, H., Zhao, Z., Cheng, H., Yan, J., & He, Q. (2023). Modeling density-driven flow in
598 porous media by physics-informed neural networks for CO₂ sequestration. *Computers
599 and Geotechnics*, *159*, 105433.
- 600 Esfe, M. H., & Esfandeh, S. (2020). 3D numerical simulation of the enhanced oil recovery
601 process using nanoscale colloidal solution flooding. *Journal of Molecular Liquids*, *301*,
602 112094.
- 603 Feng, Z., Tariq, Z., Shen, X., Yan, B., Tang, X., & Zhang, F. (2024). An encoder-decoder
604 ConvLSTM surrogate model for simulating geological CO₂ sequestration with dynamic
605 well controls. *Gas Science and Engineering*, *125*, 205314.
- 606 Gao, Z., Tan, C., Wu, L., & Li, S. Z. (2022). SimVP: Simpler yet better video predic-
607 tion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
608 Recognition* (pp. 3170–3180).
- 609 Garnier, P., Viquerat, J., Rabault, J., Larcher, A., Kuhnle, A., & Hachem, E. (2021). A
610 review on deep reinforcement learning for fluid mechanics. *Computers & Fluids*, *225*,
611 104973.
- 612 Graczyk, K. M., & Matyka, M. (2020). Predicting porosity, permeability, and tortuosity of
613 porous media from images by deep learning. *Scientific Reports*, *10*(1), 21488.
- 614 He, Q., Barajas-Solano, D., Tartakovsky, G., & Tartakovsky, A. M. (2020). Physics-informed
615 neural networks for multiphysics data assimilation with application to subsurface
616 transport. *Advances in Water Resources*, *141*, 103610.
- 617 Heinemann, N., Alcalde, J., Miocic, J. M., Hangx, S. J. T., Kallmeyer, J., Ostertag-Henning,
618 C., . . . Rudloff, A. (2021). Enabling large-scale hydrogen storage in porous media—the
619 scientific challenges. *Energy & Environmental Science*, *14*(2), 853–864.
- 620 Kakka, P. R. (2022). Sequence to sequence AE-ConvLSTM network for modelling the
621 dynamics of PDE systems. *arXiv preprint arXiv:2208.07315*.
- 622 Kani, J. N., & Elsheikh, A. H. (2017). DR-RNN: A deep residual recurrent neural network
623 for model reduction. *arXiv preprint arXiv:1709.00939*.
- 624 Karimi, M., & Bhattacharya, K. (2024). A learning-based multiscale model for reactive flow
625 in porous media. *Water Resources Research*, *60*(9).
- 626 Khebzegga, O., Iranshahr, A., & Tchelepi, H. (2020). Continuous relative permeability
627 model for compositional simulation. *Transport in Porous Media*, *134*, 139–172.
- 628 Kingma, D. P. (2014). Adam: A method for stochastic optimization. *arXiv preprint
629 arXiv:1412.6980*.
- 630 Kochkov, D., Smith, J. A., Alieva, A., Wang, Q., Brenner, M. P., & Hoyer, S. (2021). Ma-
631 chine learning-accelerated computational fluid dynamics. *Proceedings of the National
632 Academy of Sciences*, *118*(21), e2101784118.
- 633 Koesdwiady, A., El Khatib, A., & Karray, F. (2018). Methods to improve multi-step time
634 series prediction. In *2018 International Joint Conference on Neural Networks (IJCNN)*
635 (pp. 1–8).
- 636 Li, A., Chen, R., Farimani, A. B., & Zhang, Y. J. (2020). Reaction diffusion system
637 prediction based on convolutional neural network. *Scientific Reports*, *10*(1), 3894.
- 638 Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., & Anand-
639 kumar, A. (2020). Fourier neural operator for parametric partial differential equations.
640 *arXiv preprint arXiv:2010.08895*.
- 641 Liang, S.-Y., Lin, W.-S., Chen, C.-P., Liu, C.-W., & Fan, C. (2021). A review of geochemical
642 modeling for the performance assessment of radioactive waste disposal in a subsurface
643 system. *Applied Sciences*, *11*(13), 5879.
- 644 Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Philo-
645 sophical Transactions of the Royal Society A*, *379*(2194), 20200209.
- 646 Maes, J., Soullaine, C., & Menke, H. P. (2022). Improved volume-of-solid formulations

- 647 for micro-continuum simulation of mineral dissolution at the pore-scale. *Frontiers in*
648 *Earth Science*, 10, 917931.
- 649 Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). Statistical and machine learning
650 forecasting methods: Concerns and ways forward. *PloS One*, 13(3), e0194889.
- 651 Mohajerin, N., & Waslander, S. L. (2019). Multistep prediction of dynamic systems with
652 recurrent neural networks. *IEEE Transactions on Neural Networks and Learning Sys-*
653 *tems*, 30(11), 3370–3383.
- 654 MS, M., Vishnu Mohan, & Menon, V. (2021). Measuring viscosity of fluids: a deep learning
655 approach using a CNN-RNN architecture. In *Proceedings of the first international*
656 *conference on ai-ml systems* (pp. 1–5).
- 657 Nguyen, T., Shah, R., Bansal, H., Arcomano, T., Madireddy, S., Maulik, R., ... Grover,
658 A. (2023). Scaling transformer neural networks for skillful and reliable medium-range
659 weather forecasting. *arXiv preprint arXiv:2312.03876*.
- 660 Oprea, S., Martinez-Gonzalez, P., Garcia-Garcia, A., Castro-Vargas, J. A., Orts-Escolano,
661 S., Garcia-Rodriguez, J., & Argyros, A. (2020). A review on deep learning tech-
662 niques for video prediction. *IEEE Transactions on Pattern Analysis and Machine*
663 *Intelligence*, 44(6), 2806–2826.
- 664 Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., &
665 Prabhat. (2019). Deep learning and process understanding for data-driven Earth
666 system science. *Nature*, 566(7743), 195–204.
- 667 Salimzadeh, S., & Nick, H. (2019). A coupled model for reactive flow through deformable
668 fractures in enhanced geothermal systems. *Geothermics*, 81, 88–100.
- 669 Santos, J. E., Xu, D., Jo, H., Landry, C. J., Prodanović, M., & Pyrcz, M. J. (2020).
670 PoreFlow-Net: A 3D convolutional neural network to predict fluid flow through porous
671 media. *Advances in Water Resources*, 138, 103539.
- 672 Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., & Woo, W.-c. (2015). Convo-
673 lutional LSTM network: A machine learning approach for precipitation nowcasting.
674 *Advances in Neural Information Processing Systems*, 28.
- 675 Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to Sequence Learning with Neural
676 Networks. *arXiv preprint arXiv:1409.3215*.
- 677 Taieb, S. B., Bontempi, G., Atiya, A. F., & Sorjamaa, A. (2012). A review and comparison
678 of strategies for multi-step ahead time series forecasting based on the NN5 forecasting
679 competition. *Expert Systems with Applications*, 39(8), 7067–7083.
- 680 Tan, C., Gao, Z., Wu, L., Xu, Y., Xia, J., Li, S., & Li, S. Z. (2023). Temporal attention unit:
681 Towards efficient spatiotemporal predictive learning. In *Proceedings of the IEEE/CVF*
682 *Conference on Computer Vision and Pattern Recognition* (pp. 18770–18782).
- 683 Tang, M., Liu, Y., & Durlofsky, L. J. (2021). Deep-learning-based surrogate flow modeling
684 and geological parameterization for data assimilation in 3D subsurface flow. *Computer*
685 *Methods in Applied Mechanics and Engineering*, 376, 113636.
- 686 Wang, W., Xie, Q., An, S., Bakhshian, S., Kang, Q., Wang, H., ... Yuan, B. (2023).
687 Pore-scale simulation of multiphase flow and reactive transport processes involved in
688 geologic carbon sequestration. *Earth-Science Reviews*, 104602.
- 689 Wen, G., Li, Z., Azizzadenesheli, K., Anandkumar, A., & Benson, S. M. (2022). U-fno—an
690 enhanced fourier neural operator-based deep-learning model for multiphase flow. *Ad-*
691 *vances in Water Resources*, 163, 104180.
- 692 Xu, L., Chen, N., Chen, Z., Zhang, C., & Yu, H. (2021). Spatiotemporal forecasting
693 in earth system science: Methods, uncertainties, predictability and future directions.
694 *Earth-Science Reviews*, 222, 103828.
- 695 Yan, B., Harp, D. R., Chen, B., & Pawar, R. (2022). A physics-constrained deep learning
696 model for simulating multiphase flow in 3D heterogeneous porous media. *Fuel*, 313,
697 122693.
- 698 Zhu, L.-T., Chen, X.-Z., Ouyang, B., Yan, W.-C., Lei, H., Chen, Z., & Luo, Z.-H. (2022).
699 Review of machine learning for hydrodynamics, transport, and reactions in multiphase
700 flows and reactors. *Industrial & Engineering Chemistry Research*, 61(28), 9901–9949.