Artificial Intelligence in Geosciences 4 (2023) 199–208

Contents lists available at ScienceDirect

Artificial Intelligence in Geosciences

journal homepage: www.keaipublishing.com/en/journals/artificial-intelligence-in-geosciences

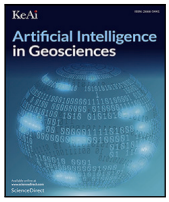

Original research articles

Estimating relative diffusion from 3D micro-CT images using CNNs

Stephan Gärttner [a,*], Florian Frank [a,b], Fabian Woller [a], Andreas Meier [a], Nadja Ray [a,c]

[a] Friedrich-Alexander-Universität Erlangen-Nürnberg, Department Mathematik, Cauerstraße 11, Erlangen, 91058, Germany
[b] Math2Market GmbH, Richard-Wagner-Straße 1, Kaiserslautern, 67655, Germany
[c] Mathematical Institute for Machine Learning and Data Science, Goldknopfgasse 7, Ingolstadt, 49085, Germany

ARTICLE INFO

Dataset link: https://zenodo.org/record/5166669

MSC:
68T07
76M10
76R50

Keywords:
Digital rock
Neural networks
Deep learning
Relative diffusion
Partial saturation

ABSTRACT

In recent years, convolutional neural networks (CNNs) have demonstrated their effectiveness in predicting bulk parameters, such as effective diffusion, directly from pore-space geometries. CNNs offer significant computational advantages over traditional methods, making them particularly appealing. However, the current literature primarily focuses on fully saturated porous media, while the partially saturated case is also of high interest for various applications. Partially saturated conditions present more complex geometries for diffusive transport, making the prediction task more challenging. Traditional CNNs tend to lose robustness and accuracy with lower saturation rates. In this paper, we overcome this limitation by introducing a CNN, which conveniently fuses diffusion prediction and a well-established morphological model that describes phase distributions in partially saturated porous media. We demonstrate the ability of our CNN to perform accurate predictions of relative diffusion directly from full pore-space geometries. Finally, we compare our predictions with well-established relations such as the one by Millington–Quirk.

1. Introduction

Data-driven methods have proven to be powerful instruments for the solution of central tasks in digital rock physics, ranging from segmentation of pore-scale images (Wang et al., 2021) to the prediction of effective quantities such as porosity, permeability, and diffusion (Graczyk and Matyka, 2020; Wu et al., 2019). Especially for the latter case, convolutional neural networks (CNNs) play a dominant role. In Wu et al. (2019), a 2D CNN was trained on artificially generated datasets for diffusion prediction achieving superior accuracy in comparison to heuristic approaches such as the Bruggeman equation. Similar results were obtained in Wang et al. (2020) and Prifling et al. (2021) for the 3D case on artificially generated geometries, underlining the enormous speed-up of several orders of magnitude in comparison to standard lattice Boltzmann diffusion solvers.

However, existing literature mainly focuses on fully saturated samples, in which the complete pore space is equally available for diffusive processes, while the (altering) unsaturated situation is also highly relevant for various applications e.g. involving imbibition or drainage processes. As such, diffusion is restricted to a certain subdomain of the full pore space. In numerical studies (Schulz et al., 2006) as well as in experimental studies (Hunt et al., 2014), this relative diffusion is reported to vary significantly with saturation. In general, structure and connectivity of the pore space also highly influence relative diffusion. In practise, this is often neglected or accounted for in heuristic laws via lumped parameters. In the context of CNNs, we demonstrate that standard CNNs trained on fully saturated pore spaces poorly generalize to partially saturated samples. More precisely, the prediction quality deteriorates monotonously with the saturation level. This behavior is expected due to the increasingly significant structural differences of the geometry available for diffusive transport.

We address the limitations of previous works and show the capabilities of suitably constructed CNNs to perform accurate diffusion predictions also on partially saturated samples in 3D. To address realistic scenarios, we aim at predicting the relative diffusion of real rock micro-CT scans instead of artificially generated samples.

We assume their pore space to be filled with two different static fluid phases, where one is completely wetting (e. g. water) and the other is fully non-wetting (gas), i. e. developing zero-degree contact angles. Moreover, we assume only the wetting phase to allow for diffusive transport. As such, the medium's relative diffusion coincides with its absolute diffusion with respect to the wetting-phase domain. The wetting phase's geometry is expected to admit lower degrees of pore-connectivity and higher complexity with lower saturation, also limiting the applicability of heuristic laws.

In the literature, multiple methods are available to predict the distribution of two fluid phases in the pore space of a porous medium,

* Corresponding author.
E-mail addresses: gaerttner@math.fau.de (S. Gärttner), frank@math.fau.de (F. Frank), nadja.ray@ku.de (N. Ray).

https://doi.org/10.1016/j.aiig.2023.11.001
Received 18 July 2023; Received in revised form 20 October 2023; Accepted 15 November 2023
Available online 1 December 2023

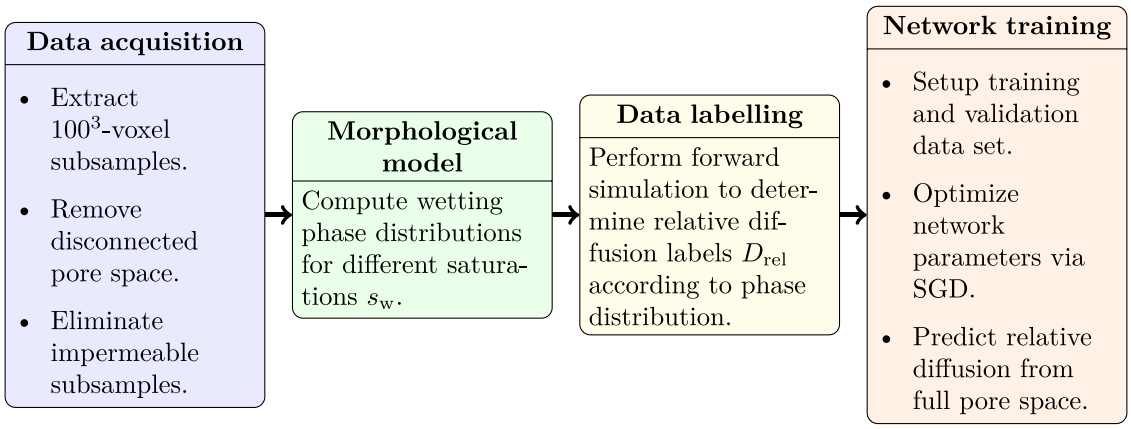

Fig. 1. Overall workflow in flow chart representation.

a comparison of which is found in Pot et al. (2015). According to this study, lattice Boltzmann methods provide a very good approximation to the experimentally measured phase distribution at the expense of large computational overheads, cf. Genty and Pot (2012). A different approach is posed by pore-morphology models (maximum spheres method) as elaborated in Mu et al. (2016), Silin et al. (2011), Hilpert and Miller (2001) and Pot et al. (2015). In this purely geometric approach, balls are placed within the pore space, successively displacing the wetting fluid phase until a desired saturation level is reached. As indicated in Mu et al. (2016), the resulting algorithm can be implemented efficiently while still reflecting the main features of the related experiments, cf. Pot et al. (2015). Consequently, we make use of the second approach to distribute the non-wetting phase in the pore space of our CT samples and integrate it into our data-driven approach.

Our main aim then, is to demonstrate the capability of CNNs to predict relative diffusion, while at the same time circumventing the additional explicit precomputation of the phase distribution.

We note that for the similar problem of permeability prediction, deep neural networks were already successfully used to predict relative permeability from numerous precomputed and carefully-chosen sample properties, cf. Arigbe et al. (2018). Contrary to this approach, we novelly train a CNN to perform relative diffusion predictions directly from the full pore-space geometry. To do so, the network is trained on relative diffusion data computed from the explicit phase distribution as dictated by the morphological model. Our CNN therefore learns to implicitly reflect the influence of saturation on the target quantity. As such, for the first time, it combines the application of a morphological model with the subsequent diffusion computation on the resulting geometry into a monolithic data-driven approach. We evaluate the prediction quality of this approach with a specific focus on the improvement due to the incorporation of the morphological model. Moreover, we compare the performance of our data-driven approach with heuristic laws for the prediction of the relative diffusion such as proposed Bruggeman or Millington–Quirk.

Therefore, our method is well-suited to simplify workflows in large-scale transport simulations involving local dynamic changes in saturation and/or porous matrix structure. In these scenarios, bulk parameters are to be steadily recomputed from reference volumes. This applies to quasi-static drainage processes as well as dissolution/precipitation scenarios, which both require repeated evaluations of current effective parameters like relative diffusion over time.

This paper is outlined as follows: In Section 2, we describe the methodology used for generating a suitable set of training data. This includes the morphological model for phase distribution in a partially saturated pore space as well as the forward simulation to label the resulting samples with the corresponding diffusion value. Section 2.4 presents the architecture chosen for our CNN. In Section 3, we investigate the performance of our CNN to predict relative diffusion from full pore-space geometries. Therefore, we train on data which take the explicit phase distribution given by the morphological model into account. We moreover compare the performance of our approach with heuristic laws. Finally, a conclusion of our results is presented in Section 4.

Table 1
Sample characteristics. Characteristics of the original Bentheimer sandstone sample: imaging resolution of the CT-scan, res., experimentally determined porosity, ϕ_{exp} ($\pm 0.5\%$), porosity computed from the CT-scan, ϕ_{cmp}, experimentally determined permeability, k_{exp} ($\pm 10\%$), mean capillary diameter, MCD, and interior surface area, A_{cmp}, computed via the Matlab function `isosurface`. All other quantities are determined in Neumann et al. (2021).

res. [µm]	ϕ_{exp}	ϕ_{cmp}	k_{exp} [mD]	MCD [µm]	A_{cmp} [mm]
2.25	22.64%	26.72%	386	30.0	355

2. Methodology and data preparation

In this section, we present the overall workflow, by which our training data set is generated. First, we provide an overview of the sampling procedure, extracting suitable 100^3-voxel training samples from a larger micro-CT scan. Subsequently, the morphological model creating a realistic phase distribution for a given partial saturation value of the pore space is described. The resulting subdomain of the wetting fluid is then used to perform computations determining the sample's relative diffusion. Details on the forward simulation of our workflow as well as the CNN design are provided below. A schematic overview of the complete workflow is given in Fig. 1.

2.1. Sampling and preprocessing

In order to properly assess the prediction quality of our CNN on realistic geometries of partially saturated porous media, we prepare our training data set from segmented micro-CT scans of a sandstone specimen, cf. Neumann et al. (2020). For classification purposes, we summarize the most relevant characteristics of this sample in Table 1. As a basis for this study, we make use of the data set generated in Gärttner et al. (2023) for the purpose of permeability predictions using CNNs. The related software as well as the data set are available in the package RTSPHEM (Gärttner and Frank, 2021). In the following, we outline the major processing steps.

All training samples are derived from a 1000^3 binary voxel image of a Bentheimer sandstone as provided in Neumann et al. (2020, 2021), featuring a resolution of 2.25 µm per voxel. Due to the particularly broad distribution in pore diameter in this type of sandstone (Gong et al., 2020), it is well suited to provide non-trivial wetting-phase geometries. We use a subsample size of 100^3 voxels, which is comparable

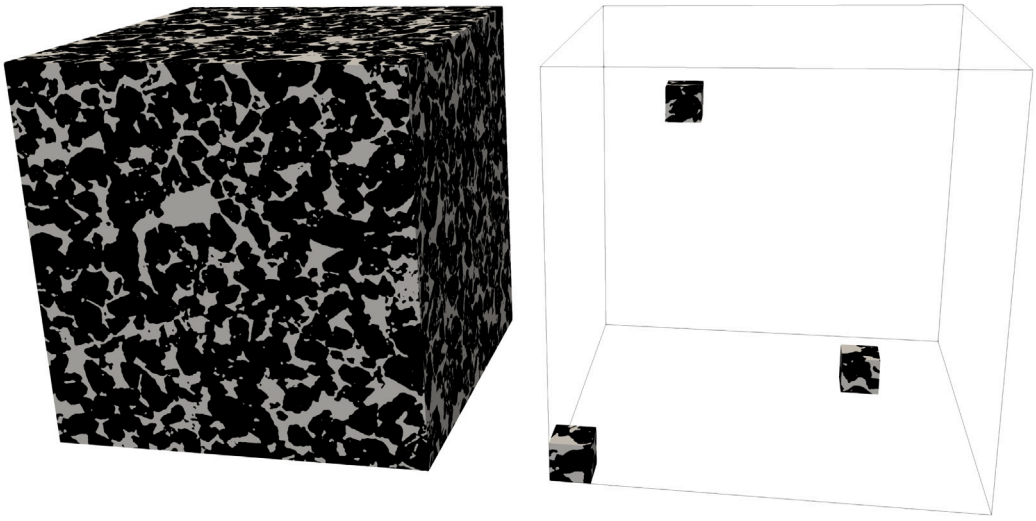

Fig. 2. Original Bentheimer sample (left) and three extracted subsamples of size 100×100×100. The position of the subsample with respect to the cube outline is identical to the relative position within the original sample. The solid matrix is colored black.

to similar studies for diffusion prediction (Wang et al., 2020; Prifling et al., 2021) or permeability prediction (Tang et al., 2022; Hong and Liu, 2020) using neural networks. In Fig. 2, we provide an illustration of the original Bentheimer sample as well as some exemplary extracted subsamples.

The sampling procedure follows the sliding frame approach as presented in Sudakov et al. (2019). As such, more samples are obtained compared to standard dissection while the subsamples can still be considered independent, cf. Hong and Liu (2020). Moreover, the data set is enriched by rotations around the y and z axis.

Finally, impermeable (disconnected) samples are excluded from the data set as their diffusion is zero independent of the saturation level. Likewise, disconnected pore space is removed from the samples since it does not contribute to diffusive transport. In doing so, we guarantee in particular that the morphological algorithm described in Section 2.2 only positions the non-wetting phase in accessible parts of the pore space. Applying the above workflow, a data set encompassing 5000 samples is established.

2.2. Morphological model

In the following, we present the morphological model used to compute the phase distribution of the wetting and non-wetting phase. As already illustrated in Hilpert and Miller (2001), well-established morphology models are capable of providing reasonable approximations of the wetting behavior in porous media. The morphological algorithm used in this paper consists of placing balls of radius r within the pore space, cf. Mu et al. (2016). The radius is derived from the Young–Laplace equation, relating capillary pressure p_c and surface tension γ to the curvature of the separating interface in terms of its spherical radius r, i. e., for zero contact angles, it holds:

$$p_c = 2\frac{\gamma}{r}.$$

In order to reproduce a virtual drainage process with gradually increasing capillary pressure, such balls are placed within the pore space with decreasing radius. This consideration leads to the well-known *maximum-inscribed-sphere* methodology (Silin et al., 2011) that we apply to our samples. This methodology is implemented efficiently following the approach of Mu et al. (2016). Initially, the sample is fully saturated with respect to the wetting phase. In the first step, an Euclidean distance map $d_1 : [0,1]^3 \to \mathbb{R}_{\geq 0}$ is computed measuring the distance of any point in the sample to the solid matrix. Matlab's built-in function `bwdist` is able to compute this function with linear complexity. The maximizer of d_1 corresponds to the center of the first and largest ball $\mathcal{B}_1$ to inscribe into the pore space while the related value of d_1 reflects its radius. Subsequent balls are now additionally required to exhibit a midpoint within the remaining wetting phase. Formally, this is achieved by setting

$$d_i = \begin{cases} 0 & \text{in } \mathcal{B}_{i-1}, \\ d_{i-1} & \text{else.} \end{cases}, \quad i \geq 2,$$

i. e. deleting the last ball from d_{i-1}, and iterating the procedure with d_i. The algorithm terminates when the requested saturation value s_w with respect to the wetting phase is undercut.

However, depending on the pore-geometry, this algorithm may lead to a significant undershooting in s_w, e. g. in cases including macropores. In order to guarantee an accurate value of s_w, the radius r_{final} of the last sphere $\mathcal{B}_{\text{final}}$ is tuned to adjust for the desired saturation as precisely as possible, cf. Mu et al. (2016). The resulting subdivision of the pore space into the two fluid domains is illustrated by one of the training samples in Fig. 3.

We subject each of the 5000 samples generated as described in Section 2.1 to the above morphological model for six different saturation levels with respect to the wetting phase, namely $s_w \in \{1, 0.9, 0.8, 0.7, 0.6, 0.5\}$. We note that using these values, the domain of the wetting phase remained connected for each considered sample. Throughout this paper, we denote the arising geometries of the wetting phase by a pair (i, s_w) of the original sample index $i \in \{0, \dots, 4999\}$ and the saturation s_w. As such, the generated data set contains a total of 30 000 samples.

2.3. Forward simulation

To apply a supervised learning approach as outlined in Section 2.4, each sample within the data set derived by the methods of Section 2.1, 2.2 needs to be labeled with a computed relative diffusion D_{cmp} that we use as the reference value.

To this end, in Section 2.3.1, we describe the method used for our diffusive transport simulations on the wetting-phase domain. More precisely, for each of the 30 000 subsamples, a stationary concentration field and the associated flux is computed by solving (1) using a mixed finite element scheme.

In Section 2.3.2, the relative diffusion is calculated by averaging the mass flux through the sample in the direction of transport.

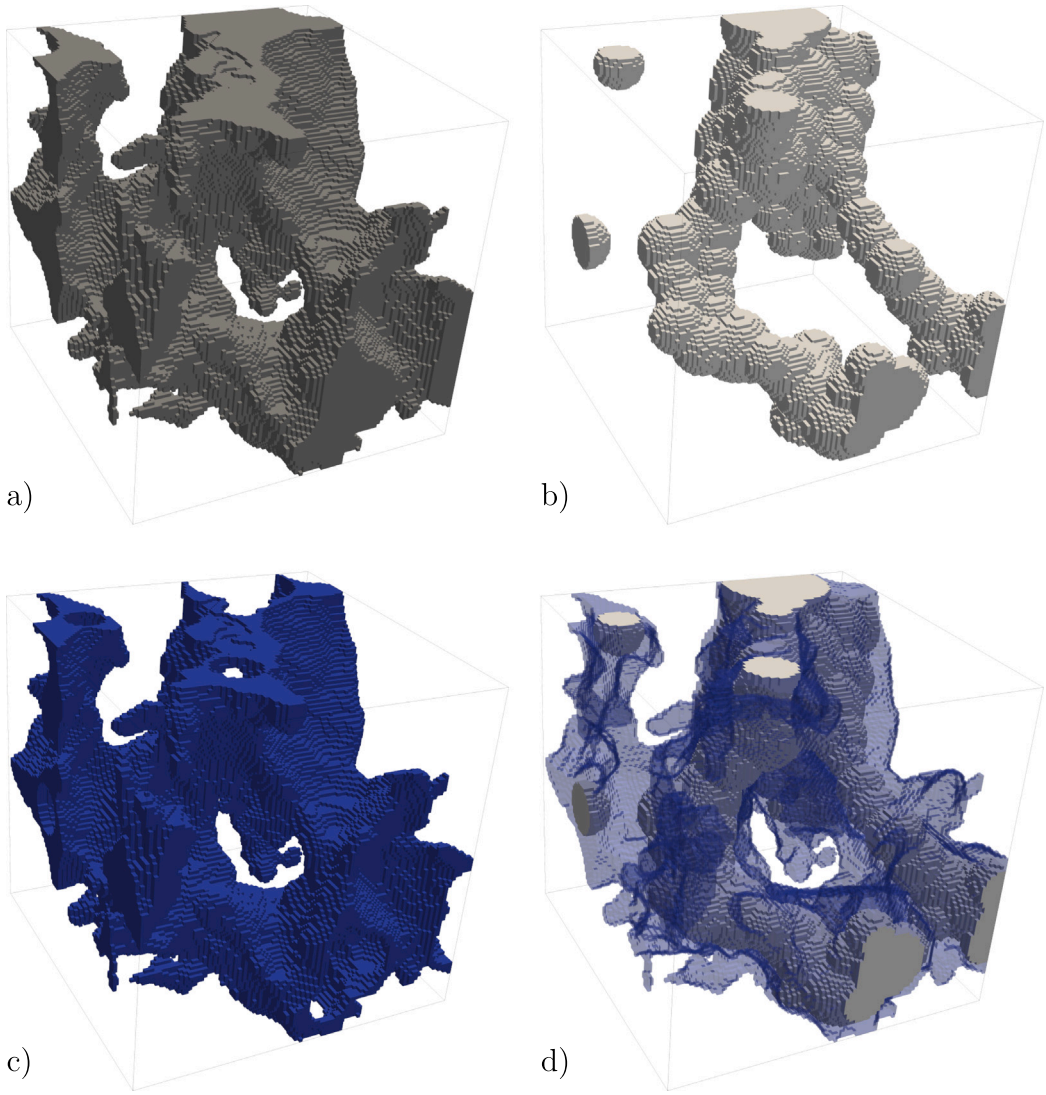

Fig. 3. Illustration of morphological model for target saturation $s_w = 0.5$. Full pore space (a) is accordingly decomposed into a non-wetting phase (b) and a wetting fluid phase (c) of equal volume. Figure (d) presents an overlay of both phases with a translucent wetting phase. In total, the non-wetting phase domain consists of 80 overlapping spheres, ranging in radius between 20.38 and 5.35 voxel lengths.

2.3.1. Computation of the concentration field

We consider the stationary diffusive transport equation in mixed form

$$-\nabla \cdot \boldsymbol{j} = 0 \qquad \text{in } \Omega, \tag{1a}$$

$$\boldsymbol{j} = D_0 \nabla c \qquad \text{in } \Omega, \tag{1b}$$

where $\Omega \subset (0,1)^3$ is the domain that consists of the union of voxels belonging to the wetting phase of one of the considered 100^3 voxel subsamples. In (1), $c = c(x,y,z)$ denotes the concentration field, $\boldsymbol{j} = \boldsymbol{j}(x,y,z)$ the mass flux, and D_0 the molecular diffusion constant. Considering diffusion with respect to the x direction, we impose the following boundary conditions with the inflow boundary $\Gamma^{\text{in}} = \{0\}\times(0,1)\times(0,1) \cap \partial\Omega \neq \emptyset$ and outflow boundary $\Gamma^{\text{out}} = \{1\}\times(0,1)\times(0,1) \cap \partial\Omega \neq \emptyset$:

$$c = 1, \qquad \Gamma^{\text{in}}, \tag{2a}$$

$$c = 0, \qquad \Gamma^{\text{out}}, \tag{2b}$$

$$\nabla c \cdot \boldsymbol{\nu} = 0, \qquad \partial\Omega \setminus (\Gamma^{\text{in}} \cup \Gamma^{\text{out}}), \tag{2c}$$

incorporating in particular homogeneous Neumann boundary conditions at the interior boundary which corresponds to the pore walls, cf. Cooper et al. (2016) and Wang et al. (2020). Since the data preparation methodology of Section 2.1 already eliminated impermeable samples, the problem is well-posed with respect to the imposed boundary conditions.

The simulations of this paper use lowest-order mixed finite elements (Raviart–Thomas elements, RT_0/P_0). To solve the arising linear system, a preconditioned MINRES solver is used with a relative tolerance of 10^{-8}. The software is implemented in the distributed-parallel framework of mfem (Anderson et al., 2021). As we will illustrate in Section 2.3.2, simulations performed on an unrefined mesh already provide sufficient accuracy for our purposes. In Fig. 4, we illustrate the concentration fields and diffusive fluxes for one sample at saturation levels $s_w = 1$ and $s_w = 0.5$. The effective diffusion D_{cmp} according to (3) requires the computation of the mass flux $\boldsymbol{j}$, which is an explicit unknown in the mixed finite element method.

2.3.2. Diffusion estimation and validation

In order to obtain the relative diffusion of our data samples, we average the normal outward flux $\boldsymbol{j} \cdot \boldsymbol{\nu}$ over the outflow boundary

$$D_{\text{cmp}} = \frac{1}{D_0\,|\Gamma^{\text{out}}|} \int_{\Gamma^{\text{out}}} \boldsymbol{j} \cdot \boldsymbol{\nu}\, \mathrm{d}\sigma \tag{3}$$

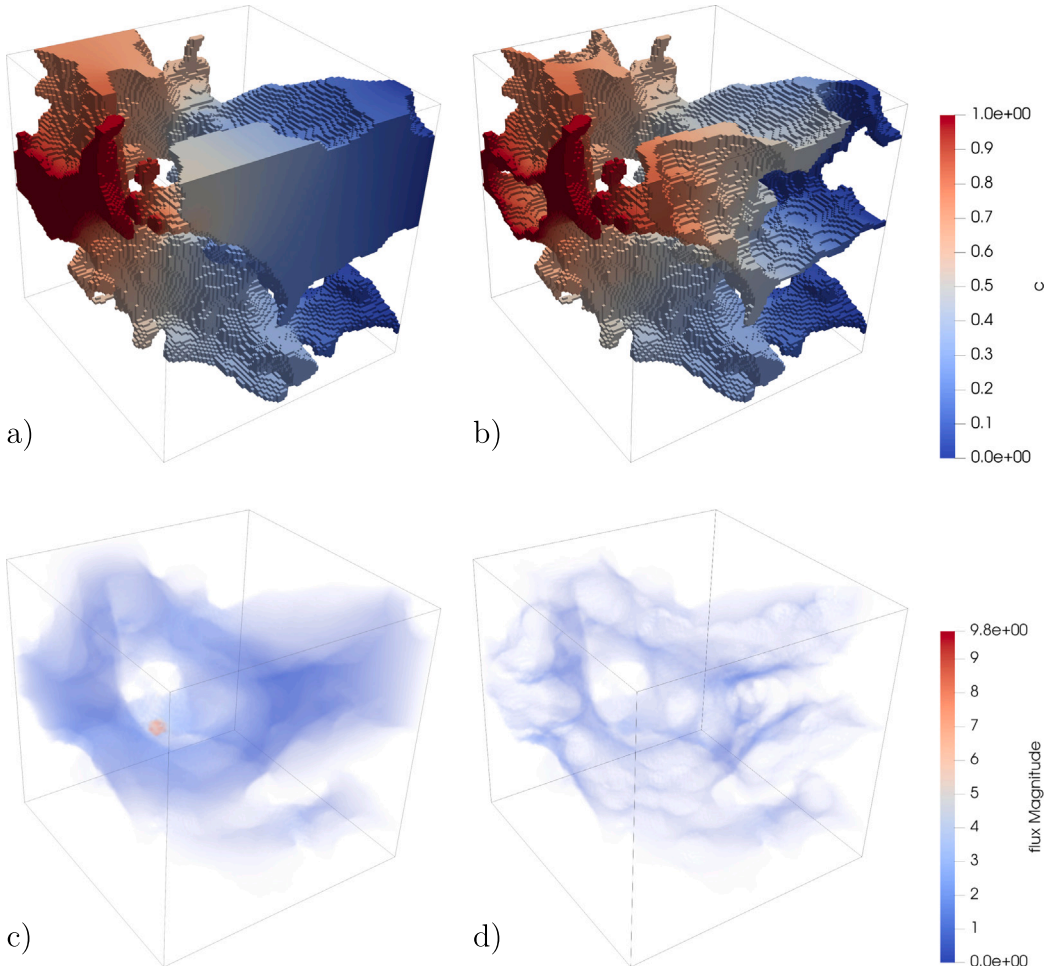

Fig. 4. Simulated concentration field c (Figures (a) and (b)) and mass flux magnitude $|j|$ (Figures (c) and (d)) on a 100^3 voxel cube. Figures (a), (c) illustrate Sample $(0, 1)$, and Figures (b), (d) illustrate Sample $(0, 0.5)$ corresponding to the wetting-phase distribution of Sample $(0, 1)$ at $s_w = 0.5$.

and normalizing it to the molecular diffusion constant D_0, cf. (1), Wang et al. (2020). We note that D_{cmp} is an intrinsic property of the geometry and therefore scale-invariant. Reformulating the diffusion problem above for a sample of arbitrary linear size $\alpha > 0$, i.e. replacing (2a) by $c = \alpha$, a simple rescaling of the unit-solution proves this property.

For validation, we compare the results of our implementation to the software package TauFactor (Cooper et al., 2016). Building upon a finite difference discretization using ghost nodes for efficiently prescribing Dirichlet boundary conditions, TauFactor is capable of computing diffusion from voxelized images and was also previously applied in similar studies (Prifling et al., 2021). We consider samples of different saturation levels and investigate the convergence of solutions with respect to global mesh refinement, i. e. halving the discretization length with each refinement level. In Table 2, the individual diffusion values are presented, the coarsest (rl 0) referring to the original grid resolution 100^3. Comparing the results of the finest resolution (rl 2), both approaches deliver relative diffusion values that differ from each other by less than 1.5%, cf. Table 2, over the whole range of considered saturations. However, values computed by TauFactor seem to converge significantly slower towards the respective reference value under mesh refinement. This behavior is expected despite the agreeing theoretical order of convergence due to the explicitly discretized flow variable in our RT_0/P_0 finite element ansatz, resulting in a higher approximation quality at the expense of a larger number of unknowns.

Our diffusion simulations required 59.3 MINRES iterations on average to converge to a relative residual of 10^{-8} (min: 44, max: 69). Performing two simulations in parallel each using five cores on an Intel Xeon E5-2630 v4 CPU, all 30 000 forward simulations required approximately 15 h to complete. In comparison, TauFactor performed the analogous calculations within two hours using the same resources. However, the results displayed in Table 2 indicate that mfem at refinement-level zero delivers more accurate results than the once refined TauFactor calculation, more than compensating for the increase in computation time. Moreover, the parallelism of mfem allows for strong scalability in comparison to TauFactor. Consequently, we use the mixed discretization results due to their superior computation time/accuracy ratio.

2.4. Machine learning model

In the following, we introduce our CNN architecture. For sake of clarity, we refrain from a general introduction to artificial neural networks here and refer to the abundant literature, e. g. Aggarwal (2018) and Alzubaidi et al. (2021). Our network is based on the experiments of Wang et al. (2020) and relies on standard and well-established structural components such as convolutional (`conv`), max-pooling (`MP`), dense (`dense`) and batch normalization layers (`BN`). In Fig. 5, we provide a graphical representation of our CNN, Table 3 presents the complete set of hyperparameters used.

In comparison to the architecture suggested by Wang et al. (2020), we dispense with dropout layers and use instead substantial L^2-regularization of the weights to reduce overfitting which slightly improved performance in our setting. As such, large amplification factors are penalized by the overall cost functional improving robustness.

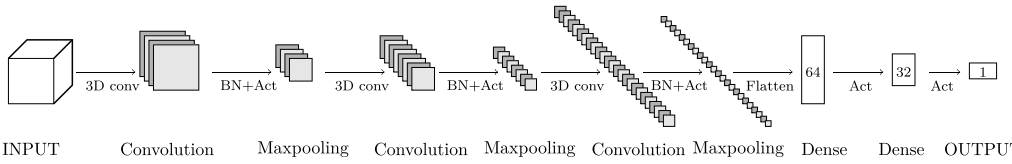

Fig. 5. Graphical representation of CNN structure. The abbreviation 'Act' labels the activation function. For further details, we refer to Table 3. For convenience, the representation style is similar to Wang et al. (2020).

Table 2
Accuracy of numerical schemes. Relative diffusion of different samples and saturation levels computed by mfem and TauFactor for three different refinement levels (rl). By index increments of three, we avoid considering rotated versions of the same sample, cf. Section 2.1, i. e. different transport directions of the same geometry.

Sample	s_w	mfem rl 0	mfem rl 1	mfem rl 2	TauFactor rl 0	TauFactor rl 1	TauFactor rl 2
(0,1)	1	0.133	0.133	0.133	0.129	0.132	0.133
(3,0.9)	0.9	0.202	0.202	0.202	0.197	0.200	0.201
(6,0.8)	0.8	0.0902	0.0906	0.0908	0.0876	0.0895	0.0905
(9,0.7)	0.7	0.0598	0.0602	0.0604	0.0579	0.0595	0.0601
(12,0.6)	0.6	0.0731	0.0740	0.0745	0.0696	0.0726	0.0738
(15,0.5)	0.5	0.0267	0.0272	0.0275	0.0245	0.0264	0.0272

Table 3
Layer structure of our PhyCNN.

Block	Layers	Learnables
input1	image input 100 × 100 × 100	–
conv1	`conv`(32,5) – `BN` – `LeakyReLU`(0.1) – `MP`(5,5)	4 096
conv2	`conv`(64,5) – `BN` – `LeakyReLU`(0.1) – `MP`(4,4)	256 192
conv3	`conv`(100,3) – `BN` – `LeakyReLU`(0.1)	173 100
dense1	`dense`(64) – `LeakyReLU`(0.1)	800 064
dense2	`dense`(32) – `LeakyReLU`(0.1)	2 113
output	`regression`(1)	–

Nomenclature:
`conv`(N,K): convolutional layer with N channels and $K \times K \times K$ kernel size;
`BN`: batch normalization layer;
`MP`(N,P): maxPooling Layer, size N stride P;
`dense`(N): dense layer with N neurons;
`regression`(N): regression layer with N neurons;
`LeakyReLU`(α): leaky rectified linear unit, slope α on negative inputs.

In Section 3.3, CNNs as described here are trained to predict relative diffusion values from the full pore space. As the morphological model, cf. Section 2.2, explicitly entered the computation of those values, it is implicitly incorporated by the network.

3. Evaluation

In this section, we perform and analyze our machine learning studies. More precisely, we evaluate the prediction quality of our CNN, cf. Section 2.4, in two different scenarios. At first, we perform the training on the fully saturated samples ($s_w = 1$) and corresponding labels and evaluate how the CNN generalizes to the partially saturated situation. We refer to this setup as CNN 1. In the second setup (CNN 2), we train a structurally identical CNN to perform predictions of *relative* diffusion directly from the full pore space, i. e. train with the labels of the partially saturated pore-space. Since our data labels are computed relying on the morphological model of Section 2.2, it is imprinted on the training data set. As such, we evaluate the CNNs' capability to implicitly incorporate the action of the underlying morphological model.

In order to allow for a transparent grading of our results, we first introduce important characteristic quantities with respect to our data set.

3.1. Evaluation metrics

In the following, we display the statistic parameters σ (standard deviation), R^2 (coefficient of determination), and MSE (mean-squared error). Measuring the spread of a data distribution, the standard deviation σ is defined as

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}\left(t_i - \bar{t}\right)^2}$$

for N real-valued quantities t_i. Thereby, $\bar{t} = \frac{1}{N}\sum_{i=1}^{N} t_i$ denotes the arithmetic mean value.

Moreover, the coefficient of determination

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(t_i - y_i)^2}{\sum_{i=1}^{N}(t_i - \bar{t})^2}$$

for targets t_i and corresponding predictions y_i is used to quantify the degree of correlation between two data sets. Values of R^2 close to one indicate a good agreement.

Similarly, the mean-squared error (MSE) is an absolute measure of accuracy given by

$$\mathrm{MSE} = \sum_{i=1}^{N}\frac{(t_i - y_i)^2}{N}.$$

3.2. Data statistics on relative diffusion

In Fig. 6, we present the saturation dependent relative diffusion D_{rel} of 18 randomly chosen samples. It is evident that all displayed curves show a monotonous behavior, since an increase of accessible pore space facilitates transport processes. However, each sample exhibits a unique saturation–diffusion relation varying significantly over the full data set: On average over the full 5000 samples, 54.7% of the original diffusion are still preserved down to a saturation of $s_w = 0.5$. Thereby, the values range between 2.8% and 77.4% with a standard deviation of $\sigma = 0.0870$. Furthermore, we note that the D_{rel}/s_w-curves depicted in Fig. 6 exhibit convex as well as concave behavior. These observations underpin the importance of sophisticated methods for the prediction of relative diffusion. Moreover, the data set is considered representative of a broad class of diffusive transport scenarios in partially saturated media providing a suitable training foundation for our CNNs outlined in Section 3.3.

As a well-known heuristic approach to relative diffusion, we compare our experimental findings to the Bruggeman equation (Tjaden et al., 2016):

$$\frac{D_{rel}}{D_{abs}} = s_w^{\alpha} \tag{4}$$

relating the relative diffusion D_{rel} to the diffusion of the fully saturated sample D_{abs}. The exponent α is a free parameter of the model. However, an explicit value can be derived for special obstacle geometries such as

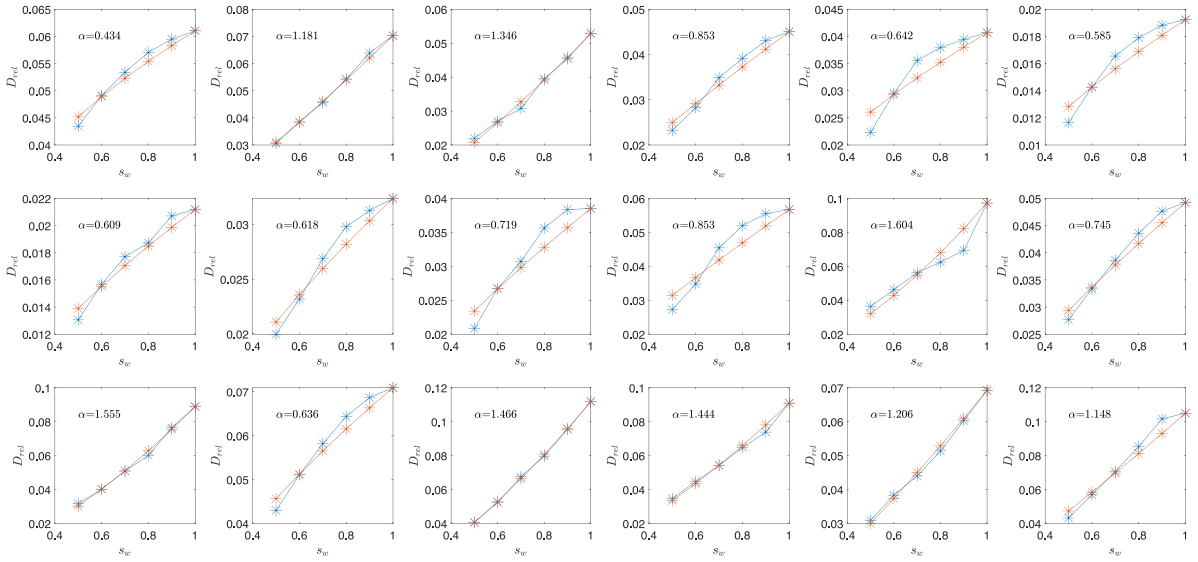

Fig. 6. Relative diffusion D_{rel} over partial saturation s_{w} for 18 different samples. Blue data points refer to the computed value, red curves are related to a fit using the Bruggeman equation (4) for the displayed parameter α. On average over all 5000 samples, 54.7% of the diffusion is maintained when restricting the available volume for diffusive transport from $s_{\mathrm{w}} = 1$ to $s_{\mathrm{w}} = 0.5$ with values ranging from 77.4% to 2.8% ($\sigma = 0.0870$).

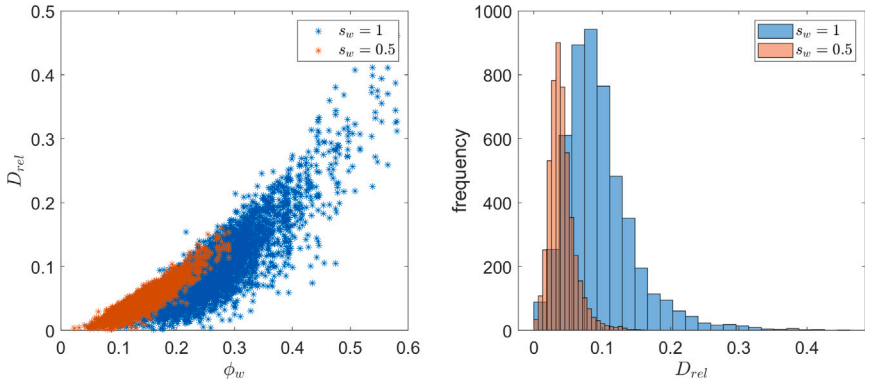

Fig. 7. Correlation between relative diffusion D_{rel} and partial porosity with respect to the wetting phase ϕ_{w} for samples exhibiting different saturation levels s_{w} (left). Distribution of diffusion values in a histogram, again itemized by saturation level (right).

balls ($\alpha = 1.5$) and cylinders ($\alpha = 2$), cf. Tjaden et al. (2016). Due to convexity ($\alpha > 1$)/concavity ($\alpha < 1$), cf. Fig. 6, a single parameter α cannot approximate all curves equally well. This is reflected in the large range of optimal parameter choices $\alpha \in (0.434, 1.604)$ for the individual $D_{\mathrm{rel}}/s_{\mathrm{w}}$-relations, cf. Fig. 6, again underlining the diversity displayed in the data set.

To illustrate the richness of our data set with respect to the covered porosity and diffusion range, we present further statistics in Fig. 7. The left image presents the dependence of the relative diffusion D_{rel} on the partial porosity $\phi_w := s_w \cdot \phi$ of the sample with respect to the volume occupied by the wetting phase. On the right, a histogram illustrates the distribution of D_{rel} in the generated data set. Apparently, the data is well-distributed over a large parameter range and is therefore considered representative, again promising high generality and robustness of our CNNs. In order to regularize the data for training, data points within the sparsely sampled regions $D_{\mathrm{rel}} < 0.01$ or $D_{\mathrm{rel}} > 0.3$ are disregarded for CNN training and validation performed in Section 3.3, see also Fig. 7. However, this procedure only discards 332 samples in total, i. e. reducing the size of our data set by approximately 1%.

3.3. CNN performance

In the following, we evaluate the prediction quality of our CNNs in the setups CNN 1 and CNN 2. In both studies, the networks described in Section 2.4 are trained over 13 epochs using a stochastic gradient descent optimizer, cf. MATLAB (2021). The step size is chosen as $\eta = 0.001$ initially and reduces by 60% every four epochs. As a cost functional, the mean squared error (MSE, cf. Section 3.1) is used supplemented by an L^2 regularization weighted by a factor of 0.02. The training procedure performed on an Nvidia Geforce Titan X graphics chip completed within 2.5 h. Throughout the study, samples were distributed among the training and validation data set using a key of 90%/10%. The implementation is performed in Matlab R2021a (MATLAB, 2021) using the *Deep Learning Toolbox*.

In CNN 1, the network was trained for predicting the *absolute diffusion*. As illustrated in Fig. 8, the network was able to perform accurate predictions on the validation data set achieving $R^2 = 89.96\%$. Comparing to $R^2 = 93.87\%$ on the training data set, the CNN is moderately prone to overfitting. We note that the measured accuracy is lower than in the experiments (Prifling et al., 2021; Wang et al., 2020) which is probably due to the use of CT-data instead of artificially generated geometries.

To investigate the generalization capabilities of CNN 1 to relative diffusion, we subject the network to data samples related to partial saturations $s_{\mathrm{w}} \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$. More precisely, predictions were performed on 500 samples per saturation level that are derived from the validation data set for $s_{\mathrm{w}} = 1$ and are therefore independent of the training data. The CNN is given the geometry of the wetting phase distribution as the input. As illustrated in Table 4, the CNN proves

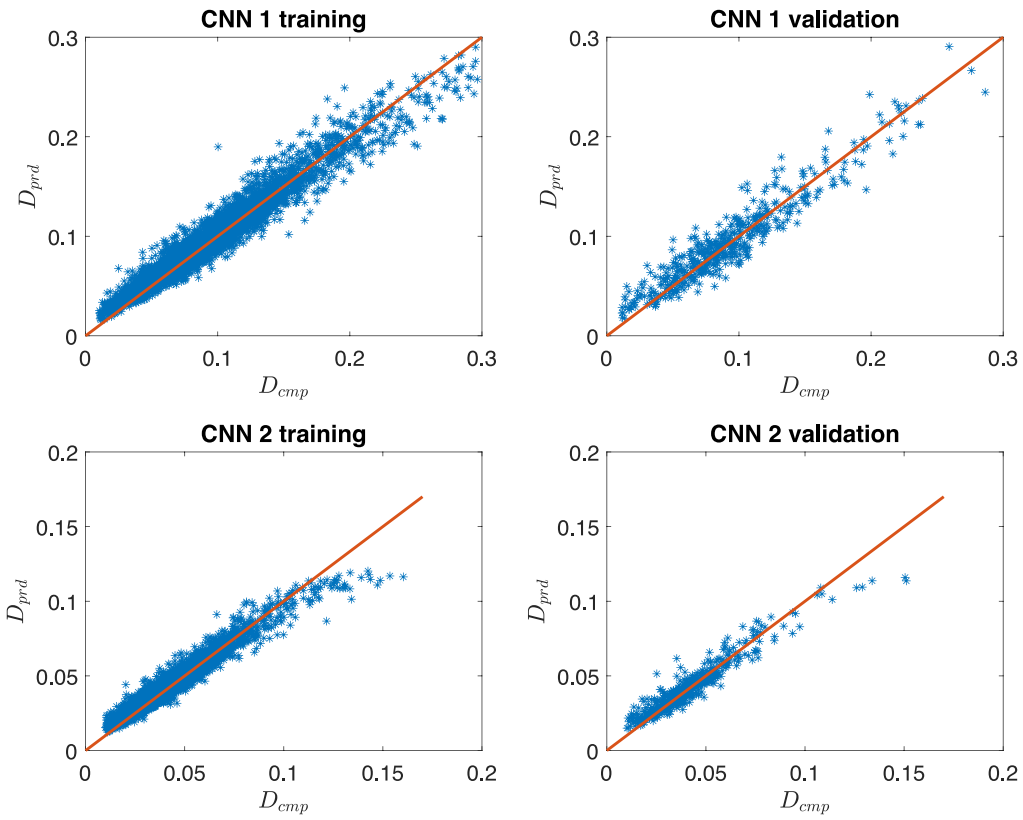

Fig. 8. Accuracy of CNN predictions: Correlation plots between computed D_{cmp} and predicted relative diffusion D_{prd} for CNN 1 and CNN 2 for validation as well as training data. The red line always reflects perfect correlation.

Table 4
Deteriorating prediction quality on partially saturated geometries. Prediction accuracy of CNN 1 on a validation data set for five different saturation levels s_{w}. For each level, the set encompasses 500 samples independent of the training data.

Saturation s_{w}	0.9	0.8	0.7	0.6	0.5
R^2	91.17%	91.31%	86.67%	69.23%	34.88%

capable of preserving its prediction accuracy across a decent range of saturations $s_{\text{w}} \in \{0.7, 0.8, 0.9, 1\}$. Below these values, the geometry of the wetting phase domain seems to differ significantly from the full pore spaces in the training data, resulting in decreased precision. At $s_{\text{w}} = 0.5$, the coefficient of determination declined to $R^2 = 34.88\%$, rendering the CNN hardly usable in practice. Considering, for example, quasi-static drainage processes which naturally feature a wide range of intermediate saturation levels, this approach requires the repeated explicit computation of the current phase distribution and subsequent prediction therefrom by a specialized CNN. As we demonstrate in the following, both steps can be fused into a single CNN rendering the explicit application of the morphological model superfluous.

On CNN 2, we train a CNN of the same structure as outlined in Section 2.4 to predict the relative diffusion related to a sample with saturation $s_{\text{w}} = 0.5$ directly from the full pore-space geometry. By this second study, we assess the ability of the CNN to take the underlying morphological model (cf. Section 2.2) into account. We note that this experiment can be conducted analogously for any other saturation level. However, for $s_{\text{w}} = 0.5$, the CNN of our first study trained on full pore-space geometries only achieved unsatisfactory results, cf. Table 4. As illustrated in the correlation plots of Fig. 8, CNN 2 achieves a similar performance as CNN 1 reflected in $R^2 = 93.22\%$ on the training data and $R^2 = 89.99\%$ on the validation data despite a few outliers for particularly large relative diffusion values. Moreover, the accuracy changes only marginally when performing the predictions from the raw-samples which did not undergo the removal of disconnected pore space, cf. Section 2.1. Therefore, samples are not required to be preprocessed in application. We note that the obtained approximation quality is significantly higher in comparison to well-established heuristic laws. First, we compare our CNN-generated results to the Bruggeman equation (4) computing relative diffusion from absolute diffusion and saturation. Given the precomputed values for D_{abs} and the target saturation $s_{\text{w}} = 0.5$, the optimal value $\alpha = 1.285$ was obtained via curve fitting. Even in this specific setup (fitting for a single saturation level), the approximation quality obtained by the Bruggeman relation results in $R^2 = 84.21\%$. Similar accuracy is obtained by the Millington–Quirk type relation (Millington and Quirk, 1961)

$$D_{\text{rel}} = s_{\text{w}}^{\alpha} \phi^{\beta} \tag{5}$$

relating saturation s_{w} and porosity ϕ to the target quantity via two adjustable parameters (α, β). Using the optimal values $(1.6094, 1.5340)$ obtained by curve fitting of the samples with $s_{\text{w}} = 0.5$, the approximation quality is similar to the one obtained by the heuristic law according to Bruggeman and yields $R^2 = 83.39\%$. However, the Millington–Quirk relation is much simpler to apply as the computation of ϕ is significantly cheaper than the determination of D_{abs}. In summary, we conclude that CNNs are able to perform relative diffusion predictions equally accurate as standard absolute diffusion predictions, cf. Table 5, outpacing simple heuristic approaches. Moreover, we note that using an Nvidia Geforce Titan X graphics chip for inferencing, CNN 2 is capable of predicting relative diffusion for a single sample in roughly 30 ms. Therefore, it renders large numbers of such evaluations computationally feasible as for example required in the simulation of large-scale transport phenomena in porous media with evolving phase distribution.

4. Conclusions

In this work, we investigated the feasibility of convolutional neural networks for relative diffusion prediction on partially saturated CT

Table 5
Final results. Comparing the prediction quality of CNN 2, Millington and Quirk (1961) and Bruggeman (Tjaden et al., 2016) on 50%- relative diffusivity. Both heuristic laws are applied with the optimal respective parameters.

Law	CNN 2 (test data)	Millington–Quirk	Bruggeman
R^2	89.99%	84.21%	83.39%

samples. As our study indicated, classical approaches trained on fully saturated samples corresponding to $s_w = 1$ only generalize well to large saturation levels (approximately $s_w \geq 0.7$) and drastically lose their approximation quality for smaller values. We therefore conclude that diffusion prediction on partially saturated samples is not well-covered by existing data driven methods. In order to predict the relative diffusion, such models additionally require the explicit application of a suitable morphological model to obtain the phase distribution at a certain saturation level as an extra preprocessing step. This workflow naturally comes with additional computational overhead and costs.

With our approach we have proven the ability of convolutional neural networks to combine the application of a morphological model with the diffusion prediction into one single monolithic data-driven approach. Our network was successfully trained to predict the relative diffusion related to a specific saturation level directly from the complete pore space without performing any additional preprocessing tasks. The qualitative and quantitative evaluation of our approach demonstrates that our CNN are capable of performing the prediction task equally accurate and robust as standard diffusion prediction on fully saturated pore spaces and are superior to classical heuristic relations such as suggested by Bruggeman or Millington-Quirk. In summary, we

- directly predict the relative diffusion from full pore-space geometries via CNNs.
- consider a monolithic approach including a morphological model at no extra computational overhead.
- outperform the prediction accuracy of Bruggeman or Millington–Quirk laws.

The resulting speed-up of our approach is well-suited to simplify workflows in large-scale transport simulations and specifically benefits setups that involve changes in phase distribution and geometry evolution, accordingly requiring numerous evaluations of relative diffusion over time. Moreover, in a straight-forward approach solely relying on the results obtained in this paper, interpolation between multiple CNNs each specialized to a certain saturation level is practicable. As such, our concepts are expected to generalize to a continuum of saturation levels rendering them also feasible for diffusion predictions in media with dynamically (quasi-statically) changing saturation levels as occurring in drainage processes.

The approach presented in this paper allows for further improvements with respect to two major aspects. First, using a more intricate network structure, an additional input parameter might be used to adjust the network's prediction for different contact angles between both fluid phases. As such, it generalizes the ideally non-wetting behavior of the secondary fluid investigated in this study. Second, physics-informed/physics-based neural networks have demonstrated superior accuracy in the prediction of complex physical phenomena in comparison to plain data-driven approaches. For the related task of permeability assessment (in fully saturated samples), the additional supply of different pore-space characteristics such as porosity, surface area or constrictivity indices to the network have shown to improve its performance, cf. Tembely et al. (2021), Wu et al. (2018) and Gärttner et al. (2023). Therefore, suitably adapted pore-space quantities are expected to increase prediction accuracy in our setup as well.

CRediT authorship contribution statement

Stephan Gärttner: Conceptualization, Data curation, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Florian Frank:** Conceptualization, Data curation, Software, Supervision, Visualization, Writing – original draft, Writing – review & editing. **Fabian Woller:** Software. **Andreas Meier:** Software. **Nadja Ray:** Conceptualization, Supervision, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Links to the data used/produced in this paper will be made available within the framework of RTSPHEM (Gärttner and Frank, 2021), https://zenodo.org/record/5166669, (GPL-3.0 license, v1.1.0) upon publication.

Acknowledgments

S. Gärttner and A. Meier were supported by the DFG Research Training Group 2339 Interfaces, Complex Structures, and Singular Limits.

N. Ray was supported by the DFG Research Training Group 2339 Interfaces, Complex Structures, and Singular Limits and the DFG Research Unit 2179 MadSoil.

F. Frank and F. Woller received support from the Competence Network for Scientific High Performance Computing in Bavaria (KONWIHR).

Computer code availability

The code used in this paper will be made available within the Matlab-compatible framework of RTSPHEM (Gärttner and Frank, 2021), https://zenodo.org/record/5166669, (GPL-3.0 license, v1.1.0) upon publication.

References

Aggarwal, C.C., 2018. Neural Networks and Deep Learning. Springer International Publishing, http://dx.doi.org/10.1007/978-3-319-94463-0.

Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M.A., Al-Amidie, M., Farhan, L., 2021. Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. J. Big Data 8 (1), http://dx.doi.org/10.1186/s40537-021-00444-8.

Anderson, R., Andrej, J., Barker, A., Bramwell, J., Camier, J.-S., Dobrev, J.C.V., Dudouit, Y., Fisher, A., Kolev, T., Pazner, W., Stowell, M., Tomov, V., Akkerman, I., Dahm, J., Medina, D., Zampini, S., 2021. MFEM: A modular finite element library. Comput. Math. Appl. 81, 42–74. http://dx.doi.org/10.1016/j.camwa.2020.06.009.

Arigbe, O., Oyeneyin, B., Arana, I., Droubi, M., 2018. Real-time relative permeability prediction using deep learning. J. Pet. Explor. Prod. Technol. 9, http://dx.doi.org/10.1007/s13202-018-0578-5.

Cooper, S.J., Bertei, A., Shearing, P.R., Kilner, J.A., Brandon, N.P., 2016. TauFactor: An open-source application for calculating tortuosity factors from tomographic data. SoftwareX 5, 203–210. http://dx.doi.org/10.1016/j.softx.2016.09.002.

Gärttner, S., Alpak, F.O., Meier, A., Ray, N., Frank, F., 2023. Estimating permeability of 3D micro-CT images by physics-informed CNNs based on DNS. Comput. Geosci. 27, 245—262. http://dx.doi.org/10.1007/s10596-022-10184-0.

Gärttner, S., Frank, F., 2021. RTSPHEM – Reactive Transport Solver in Porous Homogenized Evolving Media. Department Mathematik, Friedrich-Alexander-Universität Erlangen-Nürnberg, http://dx.doi.org/10.5281/zenodo.5166669, URL https://github.com/cupperfreeze/RTSPHEM/, Accessed: October 20, 2023.

Genty, A., Pot, V., 2012. Numerical simulation of 3D liquid–gas distribution in porous media by a two-phase TRT lattice Boltzmann method. Transp. Porous Media 96 (2), 271–294. http://dx.doi.org/10.1007/s11242-012-0087-9.

Gong, L., Nie, L., Xu, Y., 2020. Geometrical and topological analysis of pore space in sandstones based on X-ray computed tomography. Energies 13 (15), 3774. http://dx.doi.org/10.3390/en13153774.

Graczyk, K.M., Matyka, M., 2020. Predicting porosity, permeability, and tortuosity of porous media from images by deep learning. Sci. Rep. 10 (1), http://dx.doi.org/10.1038/s41598-020-78415-x.

Hilpert, M., Miller, C.T., 2001. Pore-morphology-based simulation of drainage in totally wetting porous media. Adv. Water Resour. 24 (3), 243–255. http://dx.doi.org/10.1016/S0309-1708(00)00056-7, Pore Scale Modeling.

Hong, J., Liu, J., 2020. Rapid estimation of permeability from digital rock using 3D convolutional neural network. Comput. Geosci. 24 (4), 1523–1539. http://dx.doi.org/10.1007/s10596-020-09941-w.

Hunt, A.G., Ghanbarian, B., Ewing, R.P., 2014. Saturation dependence of solute diffusion in porous media: Universal scaling compared with experiments. Vadose Zone J. 13 (4), http://dx.doi.org/10.2136/vzj2013.12.0204, vzj2013.12.0204.

MATLAB, 2021. Version 9.10.0.1602886 (R2021a). The MathWorks Inc., Natick, Massachusetts.

Millington, R.J., Quirk, J.P., 1961. Permeability of porous solids. Trans. Faraday Soc. 57, 1200–1207. http://dx.doi.org/10.1039/tf9615701200.

Mu, Y., Sungkorn, R., Toelke, J., 2016. Identifying the representative flow unit for capillary dominated two-phase flow in porous media using morphology-based pore-scale modeling. Adv. Water Resour. 95, 16–28. http://dx.doi.org/10.1016/j.advwatres.2016.02.004, Pore scale modeling and experiments.

Neumann, R., Andreeta, M., Lucas-Oliveira, E., 2020. 11 Sandstones: raw, filtered and segmented data. http://dx.doi.org/10.17612/f4h1-w124, http://www.digitalrocksportal.org/projects/317.

Neumann, R.F., Barsi-Andreeta, M., Lucas-Oliveira, E., Barbalho, H., Trevizan, W.A., Bonagamba, T.J., Steiner, M., 2021. High accuracy capillary network representation in digital rock reveals permeability scaling functions. Sci. Rep. 11 (1), http://dx.doi.org/10.1038/s41598-021-90090-0.

Pot, V., Peth, S., Monga, O., Vogel, L.E., Genty, A., Garnier, P., Vieublé-Gonod, L., Ogurreck, M., Beckmann, F., Baveye, P.C., 2015. Three-dimensional distribution of water and air in soil pores: Comparison of two-phase two-relaxation-times lattice-Boltzmann and morphological model outputs with synchrotron X-ray computed tomography data. Adv. Water Resour. 84, 87–102. http://dx.doi.org/10.1016/j.advwatres.2015.08.006.

Prifling, B., Röding, M., Townsend, P., Neumann, M., Schmidt, V., 2021. Large-scale statistical learning for mass transport prediction in porous materials using 90,000 artificially generated microstructures. Front. Mater. 8, http://dx.doi.org/10.3389/fmats.2021.786502.

Schulz, V.P., Mukherjee, P., Becker, J., Wiegmann, A., Wang, C.-Y., 2006. Numerical evaluation of effective gas diffusivity – saturation dependence of uncompressed and compressed gas diffusion media in PEFCs. ECS Trans. 3 (1), 1069–1075. http://dx.doi.org/10.1149/1.2356226.

Silin, D., Tomutsa, L., Benson, S., Patzek, T., 2011. Microtomography and pore-scale modeling of two-phase fluid distribution. Transp. Porous Media 86, 495–515. http://dx.doi.org/10.1007/s11242-010-9636-2.

Sudakov, O., Burnaev, E., Koroteev, D., 2019. Driving digital rock towards machine learning: Predicting permeability with gradient boosting and deep neural networks. Comput. Geosci. 127, 91–98. http://dx.doi.org/10.1016/j.cageo.2019.02.002.

Tang, P., Zhang, D., Li, H., 2022. Predicting permeability from 3D rock images based on CNN with physical information. J. Hydrol. 606, 127473. http://dx.doi.org/10.1016/j.jhydrol.2022.127473.

Tembely, M., AlSumaiti, A.M., Alameri, W.S., 2021. Machine and deep learning for estimating the permeability of complex carbonate rock from X-ray micro-computed tomography. Energy Rep. 7, 1460–1472. http://dx.doi.org/10.1016/j.egyr.2021.02.065.

Tjaden, B., Cooper, S.J., Brett, D.J., Kramer, D., Shearing, P.R., 2016. On the origin and application of the Bruggeman correlation for analysing transport phenomena in electrochemical systems. Curr. Opin. Chem. Eng. 12, 44–51. http://dx.doi.org/10.1016/j.coche.2016.02.006, Nanotechnology / Separation Engineering.

Wang, Y.D., Blunt, M.J., Armstrong, R.T., Mostaghimi, P., 2021. Deep learning in pore scale imaging and modeling. Earth-Sci. Rev. 215, 103555. http://dx.doi.org/10.1016/j.earscirev.2021.103555.

Wang, H., Yin, Y., Hui, X.Y., Bai, J.Q., Qu, Z.G., 2020. Prediction of effective diffusivity of porous media using deep learning method based on sample structure information self-amplification. Energy AI 2, 100035. http://dx.doi.org/10.1016/j.egyai.2020.100035.

Wu, H., Fang, W.-Z., Kang, Q., Tao, W.-Q., Qiao, R., 2019. Predicting effective diffusivity of porous media from images by deep learning. Sci. Rep. 9, 20387. http://dx.doi.org/10.1038/s41598-019-56309-x.

Wu, J., Yin, X., Xiao, H., 2018. Seeing permeability from images: Fast prediction with convolutional neural networks. Sci. Bull. 63 (18), 1215–1222. http://dx.doi.org/10.1016/j.scib.2018.08.006.