Doctoral Dissertation

博士論文

Deep Learning-Based Prediction and Upscaling of Rock Physical Properties
「深層学習を用いた岩石物性予測モデルとアップスケーリング手法の構築」

Yaotian Guo
The Graduate School of Sciences and Technology for Innovation
Yamaguchi University
山口大学大学院創成科学研究科

2026 年 9 月

Abstract

Subsurface reservoir evaluation requires estimating how much fluid can be stored underground, how fluids migrate during injection or fluid withdrawal, and how rock properties affect electrical monitoring and mechanical stability. To make these estimates, models are needed that upscale microstructural information from pore geometry, pore throats, grain contacts, and the connected solid phase to the rock properties used at the core and reservoir scales. Such models are important for hydrocarbon recovery, geothermal energy use, carbon capture and storage (CCS), and underground hydrogen storage (UHS). The main target properties are porosity, which represents storage capacity; absolute permeability, which controls fluid flow; electrical conductivity, which is related to electrical monitoring; and bulk and shear moduli, which control mechanical response. These properties are scale dependent because pore connectivity, pore-throat geometry, grain contacts, and the connectivity of the solid phase vary with the observation scale. High-resolution X-ray computed tomography (CT) images can resolve these microstructural features and can therefore be used to compute flow, electrical conduction, and elastic deformation by direct numerical simulation (DNS). However, high-resolution scans cover only a limited volume and therefore cannot fully capture centimeter-scale heterogeneity within the core. Low-resolution CT images can scan wider rock regions, but resolution loss and partial-volume effects obscure the pore geometry required for DNS. This dissertation addresses this resolution–field-of-view trade-off by combining digital rock physics (DRP), deep learning (DL), and stochastic rock physics modeling (S-RPM) to upscale rock properties from the pore scale to the core scale and further to the field scale.

Chapter 1 presents the background and objectives of this study. It reviews the scale dependence of rock properties, the resolution–field-of-view trade-off in CT imaging, and the limitations of deterministic porosity–property relationships. A conventional porosity–property relationship assigns one property value to one porosity value, but real rocks can have different permeability, electrical conductivity, and elastic moduli even at the same porosity because their pore geometry and solid structure differ. This motivates the focus of this dissertation on microstructure-based property prediction and on upscaling methods that account for property uncertainty caused by rock microstructure.

Chapter 2 develops a workflow for upscaling absolute permeability from the pore scale to the core scale. Permeability values are first computed on segmented high-resolution CT subvolumes using the lattice Boltzmann method (LBM). The corresponding low-resolution CT subvolumes at the same locations are then used as input data, and the DNS-derived permeabilities are used as training labels for a convolutional neural network (CNN). The trained CNN is applied to the low-resolution CT image of the whole core to predict a three-dimensional permeability map. This permeability map is then passed to a Darcy flow solver, which calculates flow through the whole core and gives the upscaled core-scale permeability. In the Boise sandstone verification, the workflow produced upscaled permeabilities of 5.85 D and 1.53 D for two Boise sandstone samples. The meter-scale coreflooding value was 2.8 D, which lies between the two predicted values. This relationship is consistent with the fact that the two samples were taken from higher- and lower-porosity regions. A larger 8.7 cm Boise sandstone core was then used to test whether the workflow could be applied to a wider core region and to examine the effects of training dataset size and subvolume size. Subvolumes of 100^3, 200^3, and 300^3 voxels were compared to evaluate how subvolume size affects permeability-map prediction and the resulting upscaled permeability. Within the three tested cases, larger subvolumes captured connected pore networks more consistently, reduced the influence of local outliers, and stabilized CNN training. Among the three cases, the 300^3-voxel subvolume size produced the most stable prediction and gave an upscaled permeability of 12.2 D, compared with a coreflooding value of 8.9 D. These results confirm that the proposed workflow can be applied to a wider rock core and can provide a stable estimate of its upscaled permeability. They also show that upscaling accuracy depends not only on CNN regression accuracy but also on whether the subvolumes used to construct the permeability map contain sufficiently connected pore networks. In other words, the effective REV size of the subvolumes is also a controlling factor.

Chapter 3 extends the permeability-prediction workflow developed in Chapter 2 to multiphysics property prediction in a heterogeneous reservoir sandstone. In this workflow, low-resolution CT images are used as inputs and DNS-derived properties from high-resolution CT images are used as training labels; therefore, paired high-resolution and low-resolution CT images from the same locations are required. To obtain such paired data experimentally, a millimeter-scale sub-core must be cut from a centimeter-scale rock sample for high-resolution scanning, and the resulting image must then be registered to the low-resolution image of the original core. This process damages the original core, making it difficult to perform subsequent macro-

scopic measurements on the same sample. For fragile reservoir rocks, the core may also break during sub-core extraction. To address this problem, a Downsampling Generative Adversarial Network (DSgan) is introduced to generate low-resolution counterparts from high-resolution CT images. This makes it possible to compute physical properties from high-resolution CT images by DNS, generate the corresponding low-resolution images by DSgan-based downsampling, and use the paired data for training. SE-Atten-Res3D, a three-dimensional residual network designed to emphasize image features relevant to property prediction, is then used to predict absolute permeability, porosity, electrical conductivity, bulk modulus, and shear modulus simultaneously. For all five properties, the test-data R^2 values between predicted and reference values were above 0.80. The model also reproduced both the mean porosity–property trends and the scatter of property values observed at the same porosity. Core-scale upscaling using the predicted local property maps further showed that effective core-scale properties depend not only on the magnitude of local properties but also on their spatial connectivity. The porosity–property scatter obtained from the SE-Atten-Res3D predictions is then used to fit deterministic mean trends and porosity-dependent property variability, which define the Stochastic Rock Physics Model (S-RPM) used in Chapter 4.

Chapter 4 uses the permeability S-RPM to include microstructure-induced permeability uncertainty in field-scale CO_2 storage simulation. Conventional porosity-based rock physics models, such as the Kozeny–Carman relationship for permeability, assign one permeability value to one porosity value and therefore cannot represent the permeability uncertainty caused by differences in pore structure at the same porosity. This permeability uncertainty is caused by internal rock heterogeneity and can affect flow pathways, CO_2 migration, and trapping behavior. Therefore, the S-RPM adds a spatially correlated stochastic term to the Kozeny–Carman relationship. This stochastic term is obtained from the residuals around the mean porosity–permeability trend fitted to the Chapter 3 prediction results and is modeled as a spatially correlated Gaussian random field. At the core scale, the S-RPM reconstructs property distributions and spatial maps using porosity alone as input. The resulting upscaled properties are consistent with the direct SE-Atten-Res3D upscaling results. The permeability S-RPM is then applied to CO_2 storage simulation in the Johansen aquifer using the vertical-equilibrium model in the MATLAB Reservoir Simulation Toolbox (MRST). In the basic comparison, two permeability fields are examined: a deterministic Kozeny–Carman permeability field as the conventional baseline and an S-RPM permeability field calibrated from the Chap-

ter 3 prediction results. Additional sensitivity cases are conducted using amplified S-RPM permeability fields. The calibrated S-RPM case produces CO_2 volume estimates close to the deterministic case because the fitted log-permeability variability is small and is further averaged vertically in the vertical-equilibrium model and laterally over the spreading CO_2 plume. In 20 sensitivity cases using amplified S-RPM permeability fields, the stochastic amplitude and correlation length are varied. These simulations show that storage and trapping estimates change when connected high- and low-permeability regions are large enough to influence CO_2 plume migration pathways. In particular, the case with an amplitude scale factor of 5 and a correlation length of 10 cells shows CO_2 migration behavior similar to previously reported Johansen aquifer simulations. At an intermediate correlation length, channeling along high-permeability regions reduces trapped CO_2 volume by up to 32%. By contrast, short-correlation-length fluctuations are averaged at the plume scale, whereas long-correlation-length fluctuations tend to act as broad permeability increases or decreases over the model area.

Chapter 5 summarizes the main findings of this dissertation. The main contribution is an upscaling workflow that connects pore-scale DNS, core-scale DL prediction, S-RPM fitting, and field-scale reservoir simulation. High-resolution CT-based simulations provide training data, while low-resolution CT images provide core-scale heterogeneity. Deep learning learns the relationship between low-resolution CT images and DNS-derived property values, making it possible to predict rock properties from whole-core low-resolution CT images. The S-RPM then carries the property variability learned from DL predictions into field-scale reservoir simulation, where CT images cannot be used directly. This workflow shows that rock properties can be upscaled from the pore scale to the core scale and further to the field scale while accounting for mean porosity–property trends, the spatial connectivity of local properties, and microstructure-induced property variability.

要旨

地下貯留層の評価では，地下にどれだけ流体を貯留できるか，注入または生産時に流体がどのような経路で移動するか，さらに電気的モニタリングや力学的安定性に関わる岩石物性を評価する必要がある。そのためには，地下岩石の孔隙，固体の微細構造に含まれる物理情報を，コアスケールおよび貯留層スケールの物性値にアップスケールするモデルが必要である。このようなモデルは，炭化水素回収，地熱エネルギー利用，二酸化炭素回収・貯留（CCS），および地下水素貯蔵（UHS）において重要である。対象となる主な岩石物性は，貯留容量を表す空隙率，流体流動を支配する絶対浸透率，電気的モニタリングに関わる電気伝導率，力学応答を支配する体積弾性率およびせん断弾性率である。これらの物性値はいずれもスケール依存性を持ち，孔隙の連結性，孔隙スロートの形，固体粒子間の接触，および固体相の連結構造がスケールによって異なるため，得られる物性値も変化する。高解像度 X 線 CT 画像では，これらの微細構造を捉えることができるため，流動，電気伝導，弾性変形を直接数値シミュレーション（DNS）により求めることができる。一方で，撮影範囲が小さいため，センチメートルスケールのコア内部に存在する不均質性を十分にカバーできない。低解像度 CT 画像では，より広い岩石領域を撮影できるが，解像度の限界と部分体積効果により，DNS に必要な孔隙形状を十分に保持できない。本論文では，この解像度と視野範囲のトレードオフに対応するため，デジタル岩石物理（DRP），深層学習（DL），および確率論的岩石物理モデリング (S-RPM) を組み合わせ，岩石物性を孔隙スケールからコアスケール，さらにフィールドスケールへアップスケールするワークフローを構築した。

第 1 章では，これらの背景として, 岩石物性のスケール依存性，CT 画像における解像度と視野範囲のトレードオフ，および決定論的な空隙率–物性関係式の限界について述べた。一般的な空隙率–物性関係式は，ある空隙率に対して一つの物性値を与える。しかし実際の岩石では，同じ空隙率であっても，孔隙形状や固体相構造が異なれば，浸透率，電気伝導率，弾性率は異なり、不確定性が存在する。この問題に対処するため、本論文では、微細構造に基づく物性予測を行い、物性の不確定性を考慮した物性値のアップスケーリングを行う。

第 2 章では，孔隙スケールからコアスケールへの絶対浸透率アップスケーリング手法を構築した。まず，セグメンテーションした高解像度 CT サブボリュームに対して格子ボルツマン法（LBM）を用いて浸透率を計算した。次に，同じ位置に対応する低解像度 CT サブボリュームを入力データとし，DNS で得た浸透率を教師データとして畳み込みニューラルネットワーク（CNN）を学習させた。学習済み CNN をコア全体の低解像度 CT 画像に適用し，三次元浸透率マップを予測した。さらに，得られた浸透率マップを Darcy 流ソルバーに入力し，コア全体を通る流れを計算す

ることで，コアスケールのアップスケール浸透率を求めた。Boise 砂岩を用いたワークフロー検証では，2 つの Boise 砂岩サンプルのアップスケール浸透率として 5.85 D および 1.53 D が得られた。メートルスケールのコアフラッディング実験値は 2.8 D であり，2 つの予測値の中間に位置した。この関係は，2 つのサンプルがそれぞれ高孔隙率領域と低孔隙率領域から採取されたことに対応している。さらに，8.7 cm の Boise 砂岩コアを用いて，このワークフローがより広い範囲を含む岩石コアにも適用できるかを検証し，そして学習データ数とサブボリュームサイズの影響を調べた。100^3，200^3，300^3 ボクセルのサブボリュームを比較し，サブボリュームサイズが浸透率マップの予測とアップスケール浸透率に与える影響を調べた。テストした 3 ケースの中では，サブボリュームが大きいほど，連結した孔隙ネットワークを含みやすく，外れ値的な挙動が減少し，CNN の学習も安定した。300^3 ボクセルのサブボリュームで最も安定した予測が得られ，コアフラッディング値 8.9 D に対して 12.2 D のアップスケール浸透率を得た。これらの結果から，提案したアップスケーリングワークフローが広い範囲を含む岩石コアにも適用できることを確認した。また，アップスケーリング精度は CNN の回帰精度だけでなく，浸透率マップを構成するサブボリュームが連結した孔隙ネットワークを十分に含むか，すなわち実質的な REV サイズにも依存することが示された。

第 3 章では，第 2 章で構築した浸透率予測ワークフローを，不均質な貯留層砂岩のマルチフィジックス物性予測へ拡張した。この手法では，低解像度 CT 画像を入力，高解像度 CT 画像から得た DNS 物性値を教師データとして用いるため，同じ位置に対応する高解像度 CT 画像と低解像度 CT 画像のペアが必要となる。通常，このペアを実測で作成するには，センチメートルスケールの岩石サンプルからミリメートルスケールのサブコアを切削して取り出し，さらに両画像を位置合わせする必要がある。しかし，この方法では元のコアを損傷するため，その後に同じコアを用いて巨視的物性を測定することが難しくなる。さらに，脆弱な貯留層岩石では，サブコアの切削時に岩石コアが砕ける可能性がある。この問題に対処するため，高解像度画像から低解像度 CT 画像に相当する画像を生成する Downsampling Generative Adversarial Network（DSgan）を導入した。後の文これにより，DNS で物性値を計算した高解像度 CT 画像から，DSgan によるダウンサンプリングで低解像度画像を生成し，両者を対応付けた学習データを作成できる。次に，マルチ物性値予測するための回帰モデルとして，SE-Atten-Res3D を用いて，絶対浸透率，空隙率，電気伝導率，体積弾性率，せん断弾性率を同時に予測した。テストデータにおける予測値と参照値との決定係数は，すべての物性で 0.80 以上で高い予測精度を達成できていることが確認できた。また，平均的な空隙率–物性トレンドだけでなく，同一空隙率における物性値のばらつきも再現された。さらに，予測した局所物性マップを用いてコアスケールアップスケーリングを行い，コア全体の有効物性値が局所物性値の

大きさだけでなく，それらの空間的なつながりにも影響されることを示した。最後に，SE-Atten-Res3D の予測結果から得られた空隙率–物性散布を用いて，決定論的な平均トレンドと空隙率に依存する物性のばらつきをフィッティングし，第 4 章で用いる Stochastic Rock Physics Model（S-RPM）を定式化した。

第 4 章では，第 3 章で得た S-RPM、において、浸透率の S-RPM をフィールドスケールの CO_2 貯留シミュレーションに適用した。従来の空隙率ベースの岩石物理モデル (Kozeny-Carmen 経験式) では，ある空隙率に対して一つの浸透率を与えるため，同じ空隙率でも孔隙構造の違いによって生じる浸透率のばらつきを表現できない。この浸透率のばらつきは，岩石内部の不均質性に由来するため，流体の流れ道や CO_2 の移動・トラップ挙動に影響する可能性がある。そこで S-RPM では，Kozeny-Carmen 経験式に，空間相関を持つ確率的ばらつきの項を加えた。このばらつき項は，第 3 章で得られた空隙率–物性散布から，平均トレンドに対する残差として求め，空間的に相関するガウスランダム場としてモデル化した。コアスケールでは，空隙率のみを入力として S-RPM により物性分布と空間マップを再構築した。そのアップスケール値が SE-Atten-Res3D による直接アップスケーリング結果と整合することを確認した。さらに，MATLAB Reservoir Simulation Toolbox（MRST）の垂直平衡モデルを用いて，Johansen 帯水層における CO_2 貯留シミュレーションを行った。この計算では，基本ケースとして，従来モデルである決定論的 Kozeny – Carman 浸透率場と，第 3 章の予測結果から較正した S-RPM 浸透率場を比較した。さらに，小スケール不均質性の影響を調べるため，ばらつきを増幅した S-RPM 浸透率場を用いて感度解析を行った。較正 S-RPM では対数浸透率のばらつきが小さく，さらに垂直平衡モデルにおける鉛直方向の平均化と，CO_2 プルームの広がりに伴う水平方向の平均化を受けるため，CO_2 存在量は決定論的ケースと近い結果となった。一方，増幅した S-RPM 浸透率場を用いた 20 通りの感度解析では，ばらつきの大きさと相関長を変化させた。これらの解析により，高浸透率・低浸透率の連続領域が CO_2 プルームの移動経路に影響する大きさを持つ場合，貯留量やトラップ量の予測が変化することを示した。特に，ばらつきの増幅係数を 5，相関長を 10 セルとしたケースでは，既往の Johansen 帯水層シミュレーションで報告されている CO_2 移動傾向と同様の挙動が得られた。また，相関長が中程度の場合，高浸透率領域に沿ったチャネリングが生じ，トラップされた CO_2 体積は最大 32%減少した。これに対し，短い相関長のばらつきはプルームスケールで平均化され，長い相関長のばらつきは領域全体の浸透率シフトとして働く傾向を示した。

第 5 章では，本論文の成果をまとめた。本研究の主な貢献は，孔隙スケール DNS，コアスケール深層学習予測，S-RPM フィッティング，フィールドスケール貯留層シミュレーションを接続するスケールアップ手法を示した点にある。高解像度 CT 画像に基づく物理シミュレーションは教師データを与え，低解像度 CT 画像はコアス

ケールの不均質性を与える。深層学習は，低解像度 CT 画像と DNS 由来の物性値の対応関係を学習し，コア全体の低解像度 CT 画像から岩石物性値を予測可能にする。さらに，S-RPM は，深層学習で得られた物性値のばらつきを，CT 画像が直接利用できないフィールドスケールの貯留層シミュレーションへ反映させる。このワークフローにより，平均的な空隙率–物性トレンド，局所物性値の空間的なつながり，および微細構造に由来する物性のばらつきを考慮しながら，岩石物性を孔隙スケールからコアスケール，さらにフィールドスケールへアップスケールできることを示した。

Contents

1 Introduction **4**
1.1 Background and Motivation 4
1.2 Multiscale Upscaling of Rock Properties 6
1.3 Digital Rock Physics and the Resolution–Field-of-View Trade-off . . . 9
1.4 Deep Learning for CT-Based Property Prediction 11
1.5 From Core-Scale Digital Rock Physics to Field-Scale Stochastic Rock Physics 14
1.6 Dissertation Outline 16

2 Deep Learning-Based Permeability Upscaling from Pore Scale to Core Scale **18**
2.1 Overview 18
2.2 Rock Samples 21
2.2.1 Boise sandstone cores for workflow development 21
2.2.2 Larger centimeter-scale Boise sandstone core for workflow verification 22
2.3 Method 23
2.3.1 Overview of the upscaling workflow 23
2.3.2 Training dataset acquisition and subvolume characterization . 26
2.3.3 Registration of images at different resolutions 31
2.3.4 Direct flow simulation for permeability labels 32
2.3.5 CNN models for permeability prediction 33
2.3.6 Loss functions and accuracy measurements 34
2.3.7 Darcy flow solver for whole-core upscaling 35
2.4 Workflow Development and Model Selection 36
2.4.1 Permeability labels from direct flow simulation 36
2.4.2 Evaluation of loss functions 37
2.4.3 Performance of different CNN architectures 38

2.4.4 Porosity prediction 42
2.4.5 Permeability maps predicted by the trained CNN 43
2.4.6 Whole-core upscaled permeability 47
2.5 Workflow Verification on a Larger Core 50
2.5.1 LBM permeability distributions for different subvolume sizes . 52
2.5.2 Effect of training dataset size on prediction accuracy 55
2.5.3 Prediction accuracy for permeability and porosity 57
2.5.4 Permeability–porosity relationship and permeability maps . . 60
2.5.5 Darcy-scale pressure and velocity fields 63
2.5.6 Upscaled permeability of the larger core 66
2.6 Summary 69

3 Registration-Free Multiphysics Upscaling and Stochastic Rock Physics Modeling 72
3.1 Overview 72
3.2 Sample J7 and Multiscale CT Imaging 73
3.3 Method 74
3.3.1 Overview of the Proposed Upscaling Framework 74
3.3.2 Training Dataset Preparation 77
3.3.3 Numerical Simulations for Physical Properties 79
3.3.4 DSgan-Based Generation of LR-Like CT Subvolumes 82
3.3.5 Multi-Property Prediction Using SE-Atten-Res3D 85
3.3.6 Whole-Core Property Mapping and Upscaling 88
3.3.7 Stochastic Rock Physics Model 90
3.4 Results and Discussion 93
3.4.1 DNS-Derived Physical Properties of HR Subvolumes 93
3.4.2 Validation of DSgan-Based Downsampling 97
3.4.3 Evaluation of Multi-Property Prediction 97
3.4.4 Core-Scale Property Mapping and Porosity–Property Relations 100
3.4.5 S-RPM Reconstruction of Core-Scale Property Maps 104
3.4.6 Core-Scale Upscaling 104
3.5 Summary 112

4 Field-Scale CO_2 Storage Simulation Using a Stochastic Rock Physics Model 114
4.1 Overview 114

4.2 Permeability Formulation for the RPM and S-RPM 116
4.2.1 Deterministic and Stochastic KC Models 116
4.2.2 Spatial Correlation and VE Permeability Assembly 118
4.3 Vertical-Equilibrium Reservoir Simulation Model 119
4.3.1 Governing Equations and Numerical Solution 119
4.3.2 CO_2 Volume Accounting 121
4.4 Johansen Aquifer Simulation Setup 121
4.5 Deterministic, Calibrated, and Amplified Stochastic Cases 123
4.5.1 Representative KC-Based Cases 123
4.5.2 Porosity–Permeability Dispersion Relative to DNS Calibration 126
4.6 Sensitivity to Stochastic Amplitude and Correlation Length 127
4.7 Discussion 135
4.8 Summary 137

5 Conclusions **139**
5.1 Main Findings and Contributions 139
5.2 Limitations and Future Work 141

Acknowledgement **143**

Appendix A Neural network architectures **144**

References **149**

Chapter 1

Introduction

1.1 Background and Motivation

Subsurface geological formations are used for hydrocarbon recovery Lake et al. (2014); Koroteev and Tekic (2021), geothermal energy development Tester et al. (2006); Rybach (2010); Wu et al. (2024), carbon capture and storage (CCS) Bachu (2008); Raza et al. (2019), and underground hydrogen storage (UHS) Bade et al. (2024); Taiwo et al. (2024). In these applications, the pore space stores and transmits fluids, whereas the solid rock frame supports mechanical loads and deforms in response to changes in pore pressure. Reservoir evaluation therefore requires rock properties that describe storage capacity, fluid transport, electrical response, and mechanical deformation.

Porosity quantifies the fraction of the rock volume occupied by pores and is therefore related to storage capacity, whereas absolute permeability describes how readily fluid flows through the connected pore network. Permeability depends on pore connectivity, pore-throat size and shape, and flow-path tortuosity. Consequently, two rock volumes with similar porosity can have different permeability because their pore networks can differ in connectivity and pore-throat size. Under the same pressure gradient, rock volumes with different permeability transmit different amounts of fluid. These differences in flow rate affect injection and production rates, pressure propagation, and plume migration in CCS and UHS.

Beyond porosity and permeability, subsurface characterization also requires the evaluation of electrical and mechanical properties. Electrical conductivity is needed to interpret resistivity and electromagnetic surveys because these measurements depend on connected pore water, fluid saturation, pore-water salinity, temperature, and alteration minerals Archie et al. (2003); Bergmann et al. (2012); Munoz (2014);

Arnason (2000). Bulk and shear moduli describe the deformation of the rock frame and are used to evaluate stress redistribution, reservoir compaction, caprock integrity, and induced seismicity Andrä et al. (2013); Dvorkin et al. (2011); Blunt et al. (2013); Mavko et al. (2020). These properties are evaluated for the same rock, but they are controlled by different structural features: permeability by connected pore paths and narrow throats, electrical conductivity by connected water-filled pores, and elastic moduli by the continuity and contacts of the solid skeleton. A reservoir-characterization framework should therefore estimate flow, electrical, and mechanical properties from a common three-dimensional representation of the rock microstructure.

The main difficulty is that rock structure varies from the scale of individual mineral grains and pores to the scale of cores and geological formations. Rock microstructure may include nanometer-scale clay coatings, micrometer-scale pore throats, millimeter-scale laminations, centimeter-scale vugs or fractures, and meter- to kilometer-scale depositional heterogeneity. Figure 1.1 illustrates this hierarchy of scales and the corresponding modeling approaches typically used at each level, from molecular dynamics at the nanoscale to field-scale grid simulation. High-resolution imaging and pore-scale simulation can represent individual pores and grain contacts, but only within a small rock volume. Core- and field-scale models cover larger rock volumes by using lower-resolution images or coarser numerical cells, so individual pores, pore throats, and grain contacts are not represented. Because these omitted structures control local fluid flow, electrical conduction, and load transmission, their effects must be included in the effective properties assigned to the larger-scale model. As the observation volume increases, it contains different combinations of pores, throats, layers, fractures, and connected pathways; consequently, the measured or computed effective properties can change with scale Bear (2013); Worthington (2004); Blunt (2001); Clauser (1992); Goel et al. (2016). A permeability value calculated from a small pore-scale image may not be representative of a centimeter-scale core, and a core-scale value may not directly represent the effective property required in a field-scale reservoir simulation.

This dissertation therefore develops an upscaling framework that uses physical-property information obtained from small, detailed rock images to estimate local property distributions and effective properties for larger rock volumes. In this dissertation, upscaling comprises two stages. Local property values obtained from small high-resolution regions are first estimated at many positions across the whole core to construct a three-dimensional property distribution. The governing equa-

tions for fluid flow, electrical conduction, or elastic deformation are then solved on this distribution to calculate one effective property for the whole core. The first stage focuses on absolute permeability because fluid transport is a primary process in the subsurface applications considered here, and permeability is the rock property used in Darcy's law to quantify this transport. This stage establishes the sequence of calculating local permeability from HR images, predicting its three-dimensional distribution over the whole core from LR images, and calculating the effective whole-core permeability from that distribution. After this permeability-based workflow is established, Chapter 3 extends the upscaling framework to permeability, porosity, electrical conductivity, bulk modulus, and shear modulus. The property variation quantified from these core-scale results is then incorporated into the stochastic rock physics modeling developed for field-scale applications, where continuous three-dimensional CT images are unavailable.

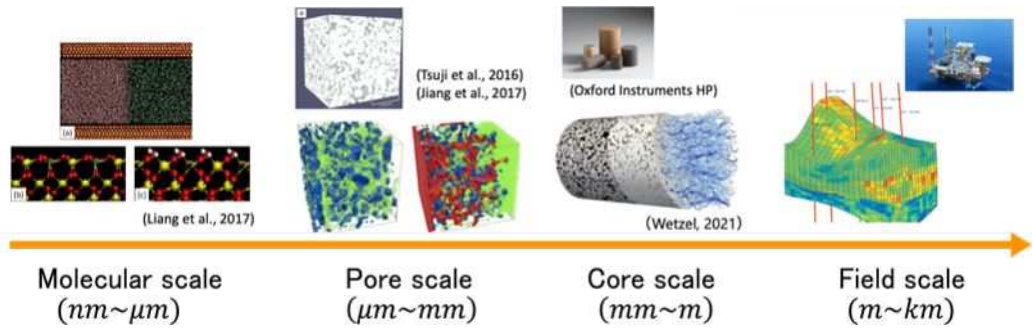


Figure 1.1: Conceptual multiscale workflow for subsurface reservoir evaluation. From left to right: molecular dynamics simulations at the nanoscale, digital rock physics and fluid flow simulations at the pore scale, physical rock samples and 3D imaging at the core scale, and large-scale grid modeling at the field scale. The arrows indicate the transfer of physical-property information between the successive scales.

1.2 Multiscale Upscaling of Rock Properties

Conventional grid-based upscaling starts from a numerical model in which permeability or another rock property has already been assigned to every fine cell. Groups of neighboring fine cells are then replaced by fewer coarse cells so that the subsequent flow simulation can be performed with a smaller computational grid. The conventional approaches reviewed here include theoretical homogenization, averaging-based methods, local numerical upscaling, and renormalization-based methods Hornung (1997); Durlofsky (1991); Wen and Gómez-Hernández (1996); Christie and Blunt

(2001); Eichel et al. (2005); Dehghan Khalili et al. (2013). In its classical form, theoretical homogenization assumes periodic small-scale heterogeneity and a clear separation between the small and large length scales; under these conditions, it derives the effective coefficients used in the larger-scale governing equation Hornung (1997); Durlofsky (1991). In averaging-based methods, all fine-cell values inside one coarse block are replaced by a single arithmetic, harmonic, geometric, or related average Wen and Gómez-Hernández (1996). Local numerical upscaling temporarily retains the fine cells, imposes a pressure difference across the block, and solves for the resulting fluid flow Durlofsky (1991); Wen and Gómez-Hernández (1996). The fine block is then replaced by one coarse-cell permeability that produces the same total flow under the same pressure difference. Renormalization-based methods perform this replacement in stages: a small group of neighboring cells is first replaced by one effective block, and the resulting blocks are combined again until the target grid size is reached Dehghan Khalili et al. (2013). The appropriate coarse-cell value depends on how the fine-cell permeabilities are arranged relative to the direction of flow. When high- and low-permeability layers are aligned with the flow, they provide parallel flow paths and the effective permeability is represented by the arithmetic average; when the flow must cross the layers successively, the layers act in series and the harmonic average applies Wen and Gómez-Hernández (1996). In local numerical upscaling, the boundary conditions determine how fluid enters, leaves, or crosses the sides of the fine block; changing these conditions can therefore change the calculated coarse-cell permeability Durlofsky (1991); Wen and Gómez-Hernández (1996). During repeated coarsening, a narrow connected high-permeability path or a low-permeability barrier can be merged with the surrounding cells. The coarse model may then predict a different flow path or total flow because the original connection or barrier is no longer represented Wen and Gómez-Hernández (1996); Christie and Blunt (2001). For the averaging, local numerical, and renormalization methods described above, the calculation starts from property values assigned to a fine grid and produces property values on a grid containing fewer, larger cells.

Reduced-order and machine-learning methods have also been developed to convert an existing detailed permeability field into a coarser representation. Onimisi et al. (2023) used the SPE10 fine-grid permeability model as input, represented it with a reduced set of POD or NPOD modes, and remapped the reconstructed values to a coarser grid by nearest-neighbor interpolation. Elmorsy et al. (2023) grouped eight neighboring digital-rock blocks, used their directional permeabilities as inputs to a PINN containing a Darcy–Ohm-based analytical relation, and predicted one

effective permeability for the combined block. Repeating this eight-to-one operation produced permeability estimates for progressively larger digital-rock volumes. Yousefzadeh and Ahmadi (2023) generated detailed permeability fields containing channel-like structures and trained a ConvLSTM model to convert them into coarser permeability fields. The model parameters were adjusted so that the pressure drops at the wells estimated by the fast marching method were similar for the detailed and coarse permeability fields. These studies coarsen permeability information that is already available throughout the detailed model.

CT-based upscaling starts from a different data condition: detailed properties can be calculated only in limited high-resolution image regions, whereas a low-resolution CT image covers a larger part or all of the core. Dehghan Khalili et al. (2013) registered HR and LR CT images so that porosity measured at the two resolutions could be compared at the same rock locations. The resulting HR–LR porosity relations were used to calibrate the whole-core LR image as a porosity field, and porosity–permeability relations derived from the HR images were then applied to this field to assign permeability throughout the core. The resulting permeability field was finally upscaled using renormalization and Darcy-flow calculations. Botha and Sheppard (2016) calculated permeability in high-resolution regions and fitted statistical relations between those values and the open-pore fraction, open-pore size, and formation factor measured in the corresponding low-resolution regions. The fitted relations were then evaluated repeatedly along the whole-core low-resolution image to assign a local permeability to each image region, producing a three-dimensional permeability map for subsequent whole-core upscaling. Menke et al. (2020) trained an Extra-Trees model to predict the permeability of small image blocks from their phase volume fractions and pore connectivity in three directions. Each larger test volume was divided into these small blocks, one predicted permeability was assigned to every block, and Darcy flow was solved through the resulting blockwise permeability distribution to obtain the effective permeability of the larger volume. These CT-based studies transfer HR-derived permeability information to larger images through porosity or selected image-derived descriptors that can be evaluated over the larger domain. Chapter 2 adopts this general strategy by using permeability calculated in a limited high-resolution region to estimate local permeability throughout a larger low-resolution core. Chapter 2 differs by using each raw three-dimensional LR grayscale subvolume itself as the input to a machine-learning model, rather than first reducing the image region to porosity or a small number of selected descriptors. The model therefore receives all grayscale voxel values and their three-dimensional

arrangement, including image information that would be lost if the subvolume were summarized only by selected descriptors. This approach also avoids segmentation of the LR CT image before property prediction. After training, the model predicts one local permeability for each LR subvolume of the whole core. These values are assembled into a three-dimensional whole-core permeability distribution, and Darcy flow through this distribution is used to calculate the effective permeability of the whole core.

The studies reviewed above focus mainly on permeability. Electrical conductivity and elastic moduli are subject to the same scale problem because electrical conductivity depends on connected water-filled pathways, whereas bulk and shear moduli depend on the continuity and contacts of the solid skeleton. Multiphysics reservoir evaluation therefore requires electrical and elastic property distributions in addition to permeability. Having established the permeability-upscaling workflow in Chapter 2, Chapter 3 extends the framework to whole-core prediction and upscaling of permeability, porosity, electrical conductivity, bulk modulus, and shear modulus. Both chapters rely on combining detailed physical information from small high-resolution images with the larger spatial coverage of low-resolution images; the next section explains why these two image types must be used together.

1.3 Digital Rock Physics and the Resolution–Field-of-View Trade-off

X-ray computed tomography (CT) and digital rock physics (DRP) provide a direct means of evaluating rock properties from three-dimensional images. In DRP, a three-dimensional CT image is processed to identify the pore and solid regions, and the resulting digital rock geometry is used as the computational domain for numerical simulation. Direct numerical simulation (DNS) then solves the governing equations of fluid flow, electrical conduction, or elastic deformation directly within the pore or solid geometry represented by the image. For permeability, the Lattice Boltzmann Method (LBM) and finite-volume methods are commonly used to solve pore-scale flow problems Raeini et al. (2012); Jiang and Tsuji (2014, 2015, 2017). For electrical conductivity and elastic moduli, the finite element method (FEM) can be used to solve the corresponding potential or mechanical equilibrium equations on voxelized rock images Arns et al. (2002); Garboczi (1998); Andrä et al. (2013). In this dissertation, a subvolume denotes a small cubic region extracted from a

three-dimensional CT image.

The reliability of DNS results depends on both image quality and spatial resolution. Image noise, reconstruction artifacts, and segmentation errors can alter phase boundaries, while insufficient spatial resolution can blur or remove pore throats and grain contacts. CT imaging presents a trade-off between spatial resolution and field of view (FOV). Increasing spatial resolution generally reduces the physical volume that can be imaged in a single scan. In this dissertation, high-resolution (HR) CT denotes images with sufficient detail for pore- or grain-scale simulation within a limited volume, whereas low-resolution (LR) CT denotes images that cover a larger core but leave the smallest pore and grain structures unresolved. HR CT images can represent pore throats, connected pores, and grain contacts with sufficient detail for DNS, but their FOV is limited to a small part of the rock sample. LR CT images can cover an entire centimeter-scale core and show variations in density or texture across the sample, while pore throats, grain contacts, and other small structures remain unresolved. In these LR images, the partial-volume effect further blurs phase boundaries because a single voxel can contain both pore and solid material Withers et al. (2021); Wildenschild and Sheppard (2013); Cnudde and Boone (2013). LR CT images therefore retain centimeter-scale variations in grayscale intensity and texture, but their unresolved pore and grain geometry is generally insufficient for direct pore-scale flow, electrical, or elastic simulation. Permeability, in particular, is controlled by narrow pore throats and the connectivity of the flow network. Loss of a narrow pore throat can reduce or interrupt a calculated flow path, whereas artificial connection of separate pores can create a flow path that is absent from the rock; either error can bias the computed permeability. Previous studies have shown that absolute and relative permeability depend strongly on pore-space geometry Fredrich et al. (1993); Pape et al. (1999). Electrical conductivity is affected for the same structural reason: loss or artificial connection of water-filled pore pathways changes the paths available for electrical current and can bias the computed conductivity Arns et al. (2001). Elastic-modulus calculations likewise depend on the geometry of the solid skeleton: if low resolution blurs grain boundaries or grain contacts, the simulated load-bearing structure differs from that of the rock and the calculated moduli can be biased Arns et al. (2002); Garboczi (1998); Andrä et al. (2013). Porosity can sometimes be estimated from coarser grayscale images because it depends mainly on pore volume, whereas DNS of transport and elasticity requires connected pore paths and grain contacts to be represented with sufficient accuracy.

The image-based upscaling problem is therefore to combine the pore-scale phys-

ical information available in small HR regions with the whole-core spatial coverage available in the LR image. HR and LR CT images provide different parts of the data required for property prediction. The LR scan covers the whole core, whereas the HR scan covers only a smaller region inside that core. Within the region imaged at both resolutions, the same small rock volume is available as a detailed HR image and as a coarser LR image. A physical-property value is calculated from the HR image of that volume, and the corresponding LR image is retained as the input paired with that value. Only the overlapping region provides such image–property pairs; the remaining parts of the whole-core LR image have no corresponding HR-derived property labels and are therefore the regions where the trained model is later used for prediction. To construct each pair correctly, the HR and LR scans must be aligned so that the LR image and the HR-derived property value correspond to the same rock volume. This alignment is referred to as image registration. In Chapter 2, the HR and LR scans contain an overlapping region of the same core, allowing registered LR subvolumes and HR-derived permeability values to be combined into paired training data.

Because LR images do not resolve pore throats and grain contacts directly, property prediction requires a method that can extract information from their three-dimensional grayscale patterns. The next section introduces deep learning as the method used to learn from the paired data constructed in the overlapping HR–LR region and to predict properties throughout the whole-core LR image.

1.4 Deep Learning for CT-Based Property Prediction

Deep learning (DL) has become an effective tool for extracting structural information from rock CT images. Convolutional neural networks (CNNs) learn spatial features directly from image data, and 3D CNNs account for the three-dimensional arrangement of voxels within CT images. CNN-based methods in DRP have been used both for processing rock CT images and for predicting physical properties from them. Image-processing applications include super-resolution enhancement of micro-CT images Roslin et al. (2022) and reconstruction of high-fidelity digital-rock images Bizhani et al. (2022), while a broader review of deep-learning applications in pore-scale imaging is given by Wang et al. Wang et al. (2021a). Physical-property prediction studies have used CNNs to estimate permeability and multiple rock properties

and to predict velocity fields, relative-permeability behavior, and capillary pressure from digital rock images Santos et al. (2020); Alqahtani et al. (2021); Telvari et al. (2023); Ko et al. (2023); Xie et al. (2023); Chen et al. (2023); Wang and Zai (2023); Khan and Khanal (2024); Kang et al. (2024). Convolutional–recurrent models have been used for property-curve prediction, whereas attention- and transformer-based models have been introduced to weight learned image features and combine information over larger spatial ranges Xie et al. (2023); Temizel et al. (2025); Emeliana et al. (2025); Tang et al. (2026).

DL models can represent relationships in which a physical property depends on several spatial features of a CT image rather than on a single descriptor such as porosity. This capability is relevant because permeability and electrical conductivity depend on pore connectivity, tortuosity, and local constrictions, whereas elastic moduli depend on the geometry and contacts of the solid skeleton. DL models can also receive raw grayscale CT subvolumes directly, without first converting every voxel into a binary pore or solid label. In LR CT images, many voxels contain both pore and solid material, so assigning each voxel entirely to one phase introduces uncertainty. Using the original LR grayscale image avoids selecting a pore–solid threshold before prediction and preserves the grayscale variations produced by partial-volume voxels.

Most image-based property-prediction studies use images acquired at the same voxel size for both model training and application and predict properties for image volumes of comparable physical size. In this dissertation, physical-property values calculated from small HR subvolumes are used as training labels, while the trained model is applied to LR subvolumes that collectively cover a much larger whole-core image. Training therefore requires paired HR-derived property values and LR images from the same rock regions. Chapter 2 constructs each training pair by combining the permeability calculated by LBM from a segmented HR subvolume with the raw LR subvolume representing the same registered rock region. A 3D CNN is trained on these pairs to predict permeability from the three-dimensional LR voxel pattern. After training, the whole-core LR image is divided into subvolumes with the same input size, and the CNN predicts one local permeability for each subvolume. The predicted values form a three-dimensional whole-core permeability distribution, and Darcy flow through this distribution gives the effective permeability of the whole core. Chapter 2 also evaluates the effects of subvolume size and training-sample number on prediction and upscaling accuracy. Smaller subvolumes provide more training samples and display stronger local permeability variation, but they are

more likely to contain atypical regions dominated by pore space or solid material and to omit flow paths that extend beyond the subvolume. Larger subvolumes contain more representative pore networks, but they provide fewer samples and require more expensive DNS calculations. Chapter 2 therefore evaluates both local prediction accuracy and whole-core upscaling accuracy for different subvolume sizes and training-sample numbers.

Chapter 3 addresses the two limitations remaining after Chapter 2: the need for physically registered HR–LR image pairs and the restriction of the prediction target to permeability. In Chapter 3, the centimeter-scale core is first scanned at 20 μm/voxel, after which a millimeter-scale sub-core is physically extracted and scanned at 1 μm/voxel. Many pore and grain features visible in the 1 μm/voxel image are blurred within the 20 μm/voxel image and therefore cannot be matched individually across the two scans. In addition, the HR volume is obtained from a physically extracted sub-core, and fragile sandstone can be damaged or fractured during cutting. The loss of individually identifiable pore-scale features in the LR scan, together with possible cutting damage to the extracted sub-core, makes reliable voxel-by-voxel registration difficult. GAN-based image-translation methods have been used to transform rock CT images between different resolutions and acquisition characteristics You et al. (2021); Fugui et al. (2025); Ugolkov et al. (2025); Hamidian et al. (2025); Yan et al. (2025); Zhang et al. (2025). Chapter 3 uses this image-transformation capability in the downsampling direction to construct training inputs without registering the two physical scans. A Downsampling Generative Adversarial Network (DSgan) is trained with unpaired sets of raw HR and LR subvolumes and converts each raw HR subvolume into an LR-like subvolume with the size and grayscale characteristics of an actual LR input. The LR-like input is generated from the raw HR subvolume whose segmented version is used to calculate the five physical-property labels. The generated input and the five labels therefore represent the same rock region without physical registration between the actual HR and LR scans. The second extension in Chapter 3 is the prediction of multiple physical properties from the same LR-like input. SE-Atten-Res3D then uses these pairs to predict permeability, porosity, electrical conductivity, bulk modulus, and shear modulus from each LR-like input. SE-Atten-Res3D combines three-dimensional residual blocks with Squeeze-and-Excitation blocks that assign weights to learned feature channels before the five properties are predicted Hu et al. (2018). The trained model is then applied to all LR subvolumes of the whole-core image to generate the five three-dimensional property distributions. The predicted permeability and

electrical-conductivity distributions are used in Darcy-flow and electrical FEM calculations, and the predicted elastic-modulus distributions are used in mechanical FEM calculations to obtain the corresponding whole-core effective properties.

Chapter 3 also predicts porosity on the same whole-core grid as the other physical properties, allowing property variation among subvolumes with similar porosity to be quantified. Subvolumes with similar porosity can have different permeability, electrical conductivity, and elastic moduli because their pore connectivity, tortuosity, grain contacts, and solid-frame continuity differ. This core-scale property variation becomes important in the next stage of the dissertation because both the Chapter 2 and Chapter 3 workflows still require three-dimensional CT images. At the field scale, continuous CT images are unavailable whereas porosity is commonly available; the next section explains how the S-RPM uses the core-scale property variation to represent the resulting property uncertainty.

1.5 From Core-Scale Digital Rock Physics to Field-Scale Stochastic Rock Physics

Field-scale reservoir models are constructed from well logs, seismic interpretation, geological models, and measurements from a limited number of cores rather than from continuous three-dimensional CT images. Porosity is commonly estimated throughout the field-scale model from well logs, seismic interpretation, or geological modeling. Other rock properties are therefore often assigned from porosity using deterministic rock physics relations. Classical examples include the Kozeny–Carman relation for permeability Carman (1937), Archie's law for electrical conductivity Archie et al. (2003), and the critical porosity model or Gassmann's equations for elastic properties Nur et al. (1995); Gassmann (2007); Mavko et al. (2020). These relations describe the expected change in each physical property as porosity changes. Once each relation has been calibrated, the porosity assigned to a model cell can be used directly to calculate the corresponding permeability, electrical conductivity, or elastic modulus. Nevertheless, they usually define a single deterministic value for a given porosity. The resulting curve therefore represents the mean porosity–property relation but does not include the property variation associated with different pore and grain structures at the same porosity.

Porosity alone does not determine a unique value of permeability, electrical conductivity, or elastic modulus. At the same porosity, permeability can vary because

of different pore-throat sizes, connectivity, and tortuosity. Electrical conductivity can vary because of differences in the connectivity and tortuosity of water-filled pore pathways. Elastic moduli can vary because of differences in grain contacts, cementation, and the continuity of the solid skeleton Farquharson et al. (2015); Avseth et al. (2000); Drobek et al. (2011). Chapter 3 quantifies this property variation by fitting a mean porosity–property relation and comparing the predicted properties of subvolumes with similar porosity with the corresponding mean value. The differences from the mean quantify the part of the property variation that is not explained by porosity alone. Well logs, seismic interpretation, and geological models do not identify the local pore and grain structures responsible for this property variation. A deterministic relation preserves the property change caused by porosity, but it assigns the same property value to all cells that have the same porosity. It therefore omits the property variation caused by different local pore and grain structures at similar porosity. Because the locations and structures responsible for this omitted variation are unknown at the field scale, the local property values remain uncertain. Porosity alone also does not determine where property values lie above or below the mean porosity–property relation. The spatial arrangement of these local variations can create or interrupt preferential flow paths and can change the distribution of stiffness in the model. Consequently, fluid migration, pressure buildup, trapping, stress, and deformation can depend on both the magnitude and spatial arrangement of the property variation.

Stochastic rock physics and uncertainty-quantification methods represent reservoir properties as distributions or multiple possible spatial fields rather than as one fixed field. Previous studies differ in the field data and statistical methods used to construct these probabilistic descriptions. Mukerji et al. (2001) and Eidsvik et al. (2004) combined seismic and well information with Bayesian and Markov models to estimate spatial probabilities of lithofacies and pore fluids. Spikes et al. (2007) incorporated rock-physics relations into Bayesian seismic inversion, whereas Grana and Mukerji (2012) propagated well-log measurement uncertainty to petrophysical-property and facies estimates by Monte Carlo simulation. de Figueiredo et al. (2018) combined rock-physics models, Gibbs sampling, and geostatistical simulation to estimate spatially correlated elastic and petrophysical properties. Across these studies, seismic attributes, well-log measurements, rock-physics relations, and statistical or geostatistical models are used to estimate uncertain facies, pore fluids, and rock properties and to describe their spatial variation. In contrast, this dissertation uses the property variation quantified from the Chapter 3 whole-core CT-based maps,

which are predicted by a model trained with DNS-derived labels. The magnitude of the stochastic variation is therefore calibrated from microstructure-related property variation represented in these CT-based data.

To incorporate this CT-derived property variation into field-scale models that contain porosity but no continuous CT data, this dissertation formulates a Stochastic Rock Physics Model (S-RPM). For each model cell, the S-RPM first calculates the mean property value from the local porosity using the corresponding deterministic porosity–property relation. A stochastic variation is then added around this mean value, and its standard deviation at each porosity is calibrated from the property variation quantified in Chapter 3. A spatially correlated random field determines the three-dimensional arrangement of these positive and negative variations.

Applying this procedure to a three-dimensional porosity field generates a corresponding property field in which the mean follows the deterministic porosity–property relation and the magnitude of the local variation is calibrated from the Chapter 3 CT-based data. Chapter 3 formulates the S-RPM separately for permeability, electrical conductivity, bulk modulus, and shear modulus, generates whole-core property maps from the predicted porosity field, and compares their effective properties with those obtained from the SE-Atten-Res3D maps. Chapter 4 then applies only the permeability formulation of the S-RPM, using the Kozeny–Carman relation for the mean permeability, to CO_2 storage simulations of the Johansen aquifer. The simulations compare CO_2 migration and trapping obtained with deterministic and stochastic permeability fields.

1.6 Dissertation Outline

This dissertation is organized into five chapters.

Chapter 1 presents the background and motivation of the study. It explains the scale dependence of rock properties, reviews conventional grid-based, CT-based, and deep-learning-assisted upscaling, and introduces the transition from core-scale image analysis to field-scale stochastic rock physics modeling.

Chapter 2 establishes and validates a CNN–Darcy workflow for upscaling absolute permeability from pore-scale calculations to the whole-core scale. Each permeability value calculated by LBM from a segmented HR CT subvolume is paired with the raw LR CT subvolume representing the same rock region, and these pairs are used to train a CNN. The trained CNN predicts the three-dimensional permeability distribution of the centimeter-scale core, and a Darcy-flow solver calculates

its effective whole-core permeability. This chapter also investigates how the number of training samples and the size of subvolumes affect prediction accuracy and the reliability of the upscaled permeability.

Chapter 3 removes the requirement for physically registered HR–LR training pairs and extends the workflow to whole-core prediction and upscaling of permeability, porosity, electrical conductivity, bulk modulus, and shear modulus in a heterogeneous reservoir sandstone. DSgan converts each HR subvolume into an LR-like input so that training pairs can be constructed without physical HR–LR image registration. Each LR-like input remains paired with the five labels calculated from the same HR subvolume, and SE-Atten-Res3D uses these pairs to predict the five whole-core property distributions. Darcy-flow and FEM calculations are then performed on the predicted distributions to obtain the effective whole-core permeability, electrical conductivity, bulk modulus, and shear modulus. The property variation observed among subvolumes with similar porosity is then used to formulate separate S-RPMs for permeability, electrical conductivity, bulk modulus, and shear modulus.

Chapter 4 applies the permeability formulation of the S-RPM to field-scale CO_2 storage simulations of the Johansen aquifer. The same reservoir porosity field is used to generate a deterministic Kozeny–Carman permeability field and stochastic permeability fields whose variation is calibrated from the Chapter 3 property data. These fields are supplied to simulations of CO_2 injection and post-injection migration. The simulations compare how deterministic and stochastic permeability fields affect CO_2 migration and the partitioning between mobile and trapped CO_2.

Chapter 5 summarizes the main findings and conclusions of the dissertation. It summarizes the contributions and limitations of the multiscale workflow and discusses its extension to other rock types, multiphase properties, and field-data-constrained modeling.

Chapter 2

Deep Learning-Based Permeability Upscaling from Pore Scale to Core Scale

2.1 Overview

Fluid flow in porous rocks is important in subsurface engineering applications, including oil and gas recovery Lake et al. (2014), geothermal energy Tester et al. (2006), and geologic carbon storage Juanes et al. (2006). Absolute permeability characterizes fluid transport in porous media and depends strongly on pore-space geometry. Pore-structure heterogeneity occurs over a wide range of length scales, resulting in scale-dependent permeability Worthington (2004). Scale dependence and heterogeneity therefore complicate permeability prediction.

Three-dimensional imaging allows rock heterogeneity to be examined at multiple scales: nano-CT imaging at the nanometer scale Goral et al. (2019), micro-CT imaging at the pore scale Blunt et al. (2013), and medical CT imaging at the core scale Zhai et al. (2020). Previous studies have examined relationships among these scales and the scale dependence of permeability Krause et al. (2011); Dehghan Khalili et al. (2013); Menke et al. (2021). Anto-Darkwah et al. Anto-Darkwah et al. (2021); Anto-Darkwah and Rabinovich (2022) estimated small-scale sub-core properties by a top-down method using coreflooding data. Dehghan Khalili et al. Dehghan Khalili et al. (2013) investigated carbonate reservoir rocks and used porosity–permeability relationships derived from micro-CT images to estimate permeability at a larger scale through renormalization. Menke et al. Menke et al. (2021) applied multivariate

regression to upscale and predict permeability from the pore scale to the Darcy scale by deriving an upscaled porosity–permeability relationship. A remaining task is to estimate core-scale permeability directly from low-resolution CT images while accounting for small-scale heterogeneity.

Core-scale permeability prediction requires permeability data calculated from high-resolution CT images and the larger rock volume covered by low-resolution CT images. Low-resolution CT images provide greater sample coverage and capture core-scale heterogeneity, whereas narrow pore throats and connected pore paths become blurred or unresolved at lower resolution Botha and Sheppard (2016). Direct numerical simulation (DNS) solves the governing equations of fluid dynamics on detailed three-dimensional rock images to obtain pore-scale flow fields. Previous studies have shown that absolute and relative permeability depend strongly on pore-space geometry Fredrich et al. (1993); Pape et al. (1999). Accurate pore-scale flow simulation therefore requires high-resolution CT images that resolve pore geometry and the main flow paths. Some volume-averaged properties, including porosity and elastic properties, are less sensitive to local pore geometry and can be estimated from lower-resolution CT images Arns et al. (2007); Chawshin et al. (2022).

DNS methods such as the lattice Boltzmann method (LBM) and finite volume method can solve fluid flow in pore spaces extracted from high-resolution CT images Raeini et al. (2012); Jiang and Tsuji (2014); Jiang et al. (2014); Jiang and Tsuji (2015, 2017); Jiang et al. (2021b). These simulations require substantial computational resources and time. High-resolution imaging also restricts the field of view to a small rock volume, whereas low-resolution CT imaging provides greater spatial coverage. Permeability upscaling therefore combines permeability data calculated from high-resolution CT images with the larger spatial coverage of low-resolution CT images. The resulting coarse-scale model must reproduce the averaged flow response of the fine-scale model. Classical homogenization provides effective properties for periodic structures Xu (2012), whereas natural rocks commonly exhibit non-periodic heterogeneity. Accordingly, the upscaling model must account for rock heterogeneity across the available image scales Menke et al. (2021).

CNNs can extract three-dimensional image features and learn nonlinear relationships between CT images and physical properties. Neural networks have been used to analyze rock CT images Da Wang et al. (2019); Wang et al. (2021b), reconstruct synthetic rock images using generative adversarial networks You et al. (2021), and predict petrophysical properties such as permeability, relative permeability, and fluid-phase distributions Suzuki et al. (2022); Wang et al. (2021b). Santos et al.

Santos et al. (2020) used a CNN to predict pore-scale flow velocity fields. Rabbani and Babaei Rabbani and Babaei (2019) estimated the permeability of pore-network models from geometric properties using an artificial neural network. Alqahtani et al. Alqahtani et al. (2021) predicted permeability directly from CT images using 3D CNN regression.

This chapter develops and verifies a CNN-based workflow for core-scale permeability mapping and upscaling. The CNN is trained using low-resolution raw CT subvolumes and permeability labels calculated by LBM from the corresponding high-resolution CT subvolumes. The trained CNN predicts the whole-core permeability map, which is used in a Darcy flow solver to calculate the upscaled permeability. This procedure combines permeability calculated from resolved pore structures with the larger rock volume represented by low-resolution CT images.

The chapter consists of two stages. The first stage establishes the workflow using two small Boise sandstone core samples scanned at two resolutions. High-resolution CT subvolumes are segmented and used for LBM permeability calculations, while the corresponding low-resolution raw CT subvolumes are used as CNN inputs. Five loss functions—mean squared error (MSE), mean absolute error (MAE), Huber loss, mean squared logarithmic error (MSLE), and log-cosh loss—and six CNN architectures—Conv8, VGG16, VGG19, GoogLeNet-V1, GoogLeNet-V3, and ResNet34—are evaluated. The loss-function and network comparisons identify the model used for whole-core permeability prediction and Darcy-scale upscaling.

The second stage verifies the workflow using a larger 8.7 cm Boise sandstone core and examines the effects of training dataset size and subvolume size. The high-resolution CT image acquired from the central region of the core is partitioned into 100^3, 200^3, and 300^3 voxel subvolumes. The corresponding low-resolution raw CT subvolumes are used for CNN training, and the trained CNN is applied to the low-resolution CT image of the whole core. The analysis examines the number of training data points, the average pore-body count per subvolume, CNN prediction accuracy, and the resulting upscaled permeability. Small subvolumes provide more training data but may contain pore structures that are not representative of connected core-scale flow. Larger subvolumes contain more representative pore-network information but provide fewer training samples and require more computation. These effects are evaluated to determine the subvolume size and training dataset required for permeability prediction and whole-core upscaling.

2.2 Rock Samples

2.2.1 Boise sandstone cores for workflow development

Boise sandstone was selected for workflow development because its pore structure and petrophysical characteristics are suitable for permeability upscaling studies. Boise sandstone has a range of pore and grain sizes, a well-characterized mineral composition, and a relatively simple geological structure, which facilitates the initial development of the workflow. Boise sandstone consists of approximately 40% quartz and 50% feldspars, with a brine porosity of about 24%. The meter-scale permeability of the parent Boise sandstone sample was 2.8 D, as measured by a coreflooding experiment. Reported permeability values for Boise sandstone range from 0.8 to 7.7 D Weinbrandt et al. (1975); Tripathi et al. (2018); Dong et al. (2010); Velcin et al. (2020). The pore-size and grain-size distributions are centered at approximately 50 and 100 μm, respectively, based on thin-section optical images.

Two small core samples, denoted Sample I and Sample II, were extracted from the meter-scale Boise sandstone sample. Each sample had a diameter of 6.8 mm and a length of approximately 1.0 cm (Figure 2.1). The local permeability distributions of the two samples were measured using a Core Laboratories Pressure Decay Profile Permeameter (PDPK-400). The measured permeability ranged from 4 to 10 D, which is higher than the meter-scale coreflooding result. This difference is consistent with possible overestimation by the PDPK method Otosigbo (2017) and local sampling of relatively high-permeability regions. Micro-CT imaging of the dry samples was conducted using an Xradia Versa XRM-500 scanner at two resolutions: 2 μm/voxel for high-resolution imaging and 8 μm/voxel for low-resolution imaging.

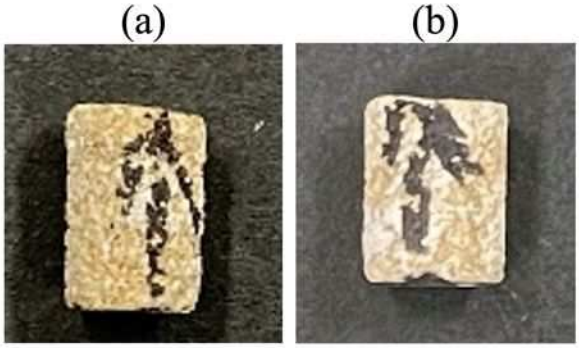


Figure 2.1: Boise sandstone core samples used for workflow development: (a) Sample I and (b) Sample II. The black arrow indicates the horizontal direction.

2.2.2 Larger centimeter-scale Boise sandstone core for workflow verification

A larger Boise sandstone core was used to verify the workflow after the basic upscaling procedure had been established. The same rock type was used so that the verification focused on the effects of core size, training dataset size, and subvolume size rather than lithological differences. The grain-size distribution, pore structure, and petrophysical characteristics of Boise sandstone have been characterized in previous studies Louis et al. (2024); Cheung et al. (2012); Li and Xie (2021).

The sample was imaged using an Xradia Versa XRM-500 scanner at two resolutions. The complete core was scanned at 20 μm/voxel, covering a diameter of 2.0 cm and a length of 8.7 cm (Figure 2.2(a)). The central region of the sample was scanned at 3 μm/voxel to obtain the pore geometry required for pore-scale flow simulation (Figure 2.2(b)). The high-resolution scan was used to calculate permeability labels using LBM. The corresponding low-resolution raw CT subvolumes were used for CNN training, and the full low-resolution CT image was used for whole-core permeability prediction.

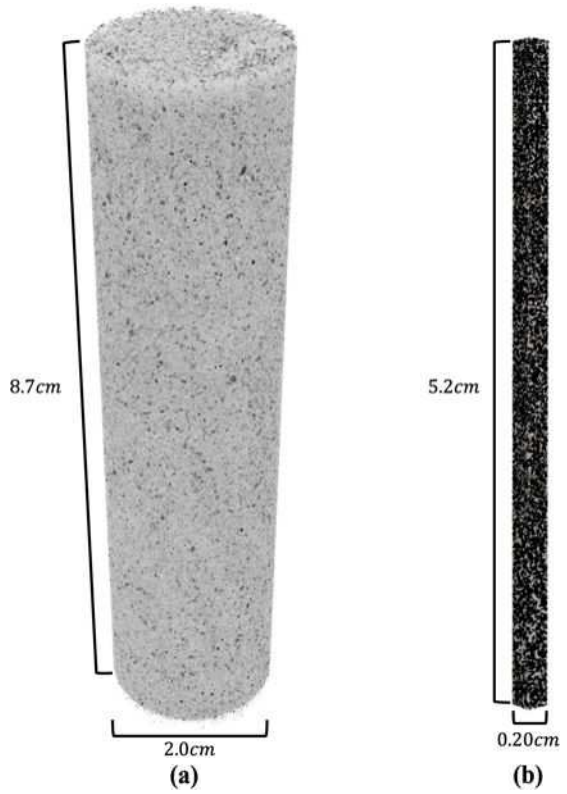


Figure 2.2: Large Boise sandstone core scanned at two resolutions for workflow verification: (a) low-resolution CT image of the entire core and (b) high-resolution CT image of the central region.

2.3 Method

2.3.1 Overview of the upscaling workflow

The permeability upscaling method consists of three main parts. First, a CNN is trained using paired data consisting of low-resolution raw CT subvolumes and permeability values calculated by LBM from the corresponding segmented high-resolution CT subvolumes. Second, the trained CNN predicts the permeability distribution of the whole core directly from low-resolution raw CT images. Third, the predicted permeability map is used in a Darcy flow solver to calculate the upscaled permeability.

High-resolution CT images resolve the pore throats and connected flow paths required for LBM calculation, whereas low-resolution CT images cover the core-scale domain. The workflow therefore associates each permeability value calculated from

a high-resolution CT subvolume with the low-resolution raw CT subvolume at the same location. The low-resolution images are used without binarization, retaining grayscale information associated with unresolved pores and partial-volume effects.

The basic permeability upscaling workflow used for the two small core samples is shown in Figure 2.3. The procedure is as follows:

1. Scan the whole 1.0 cm-long core sample at low resolution.
2. Select a central region of interest and scan it at high resolution to resolve the pore geometry.
3. Apply image filtering and segmentation to the high-resolution CT images to obtain binary pore geometry.
4. Extract subvolumes from the segmented high-resolution CT images.
5. Calculate the permeability of each high-resolution subvolume using LBM.
6. Register the high-resolution and low-resolution CT images and identify the corresponding low-resolution subvolumes.
7. Construct a training dataset consisting of low-resolution raw CT subvolumes and the corresponding LBM permeability values.
8. Train the CNN models using the paired dataset.
9. Predict the permeability map of the whole core from the low-resolution raw CT images.
10. Calculate the flow flux in the whole core using a Darcy flow solver.
11. Calculate the upscaled permeability of the whole core.

The same workflow was subsequently applied to the larger 8.7 cm core. The larger-core case differs in the CT resolutions, core dimensions, and subvolume-size design, as described in Section 2.3.2.

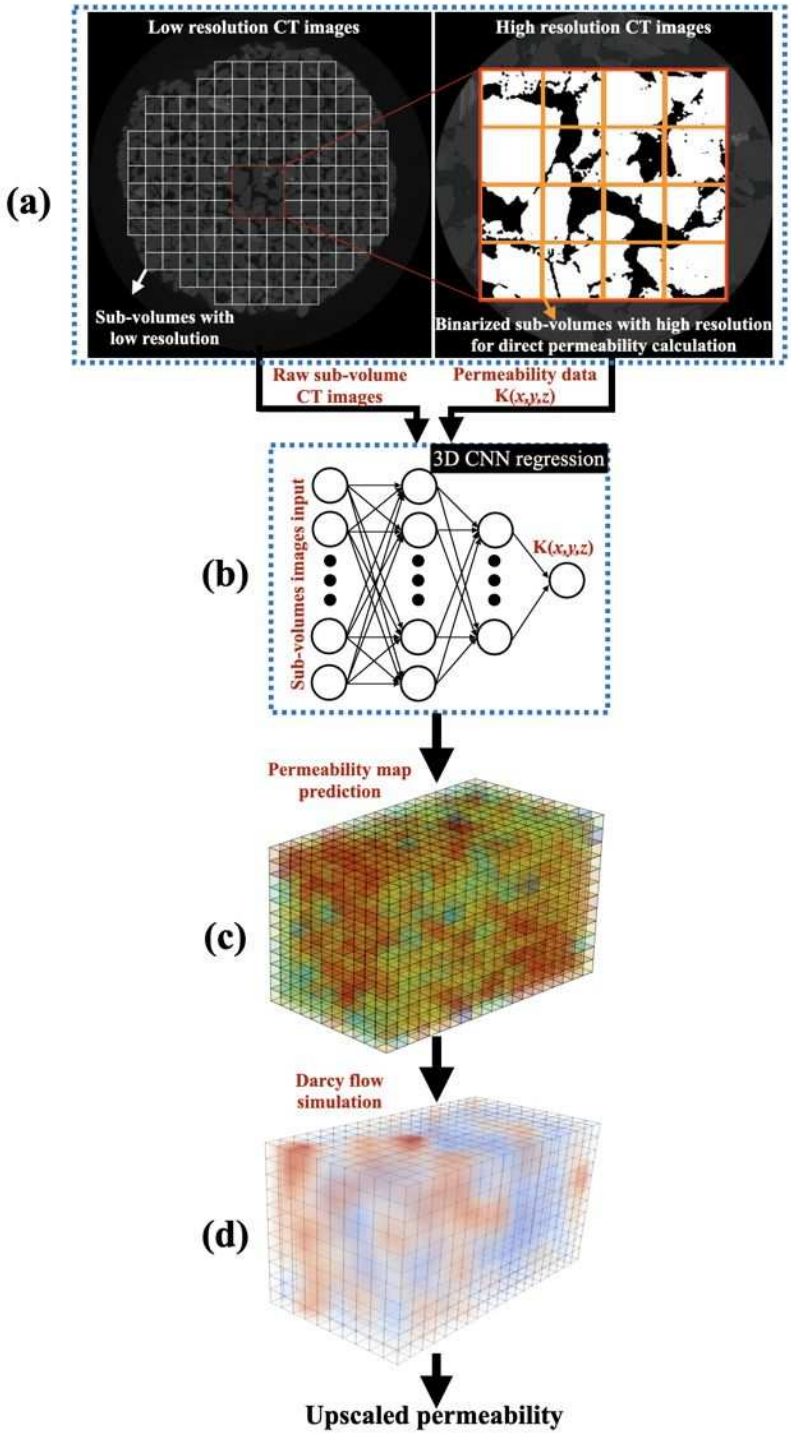


Figure 2.3: Basic permeability upscaling workflow: (a) micro-CT scanning and preparation of the training dataset; (b) CNN training using CT images and permeability labels; (c) whole-core permeability map predicted by the trained CNN; and (d) flow-flux distribution calculated using the Darcy flow solver.

2.3.2 Training dataset acquisition and subvolume characterization

For each workflow-development sample, the whole core was first imaged at 8 μm/voxel to cover the complete sample (Figure 2.4). The central region was then selected as the region of interest and imaged at 2 μm/voxel to obtain the pore geometry required for LBM permeability calculation. A non-local means filter was applied to the high-resolution CT images to reduce image noise while preserving phase boundaries. The filtered images were segmented using a threshold selected from a local minimum in the grayscale-intensity histogram. The high-resolution CT images were segmented for pore-scale flow simulation, whereas the low-resolution raw CT images were used directly as CNN inputs.

A binary pore-geometry model with dimensions of $690 \times 690 \times 5000$ voxels was extracted from the reconstructed image of each cylindrical sample. The binary model was divided into 200^3 high-resolution subvolumes. This subvolume size was selected to be smaller than the representative elementary volume so that the dataset included a range of porosity and permeability values Hong and Liu (2020). Adjacent subvolumes overlapped by 100 voxels during extraction to increase the number of training samples.

The permeability of each high-resolution subvolume was calculated using the LBM solver described in Section 2.3.4. The corresponding low-resolution raw CT subvolume was extracted at the same location. Because the high-resolution and low-resolution scans differed by a factor of four in voxel size, each 200^3 high-resolution subvolume corresponded to a 50^3 low-resolution subvolume. The resulting dataset consisted of pairs of low-resolution raw CT subvolumes and the corresponding permeability values calculated from the high-resolution subvolumes using LBM. A total of 8000 subvolumes were obtained for each sample. Ninety percent of the data were used for training, and the remaining 10% were used for testing.

Because adjacent subvolumes overlapped during extraction, the random 90/10 split may place spatially adjacent or partially overlapping subvolumes in both the training and test datasets. The test accuracy in the workflow-development stage is therefore interpreted primarily as interpolation accuracy within the sampled high-resolution region, rather than as generalization accuracy for spatially independent regions. The subsequent application to the larger core evaluates the workflow over a wider core-scale domain.

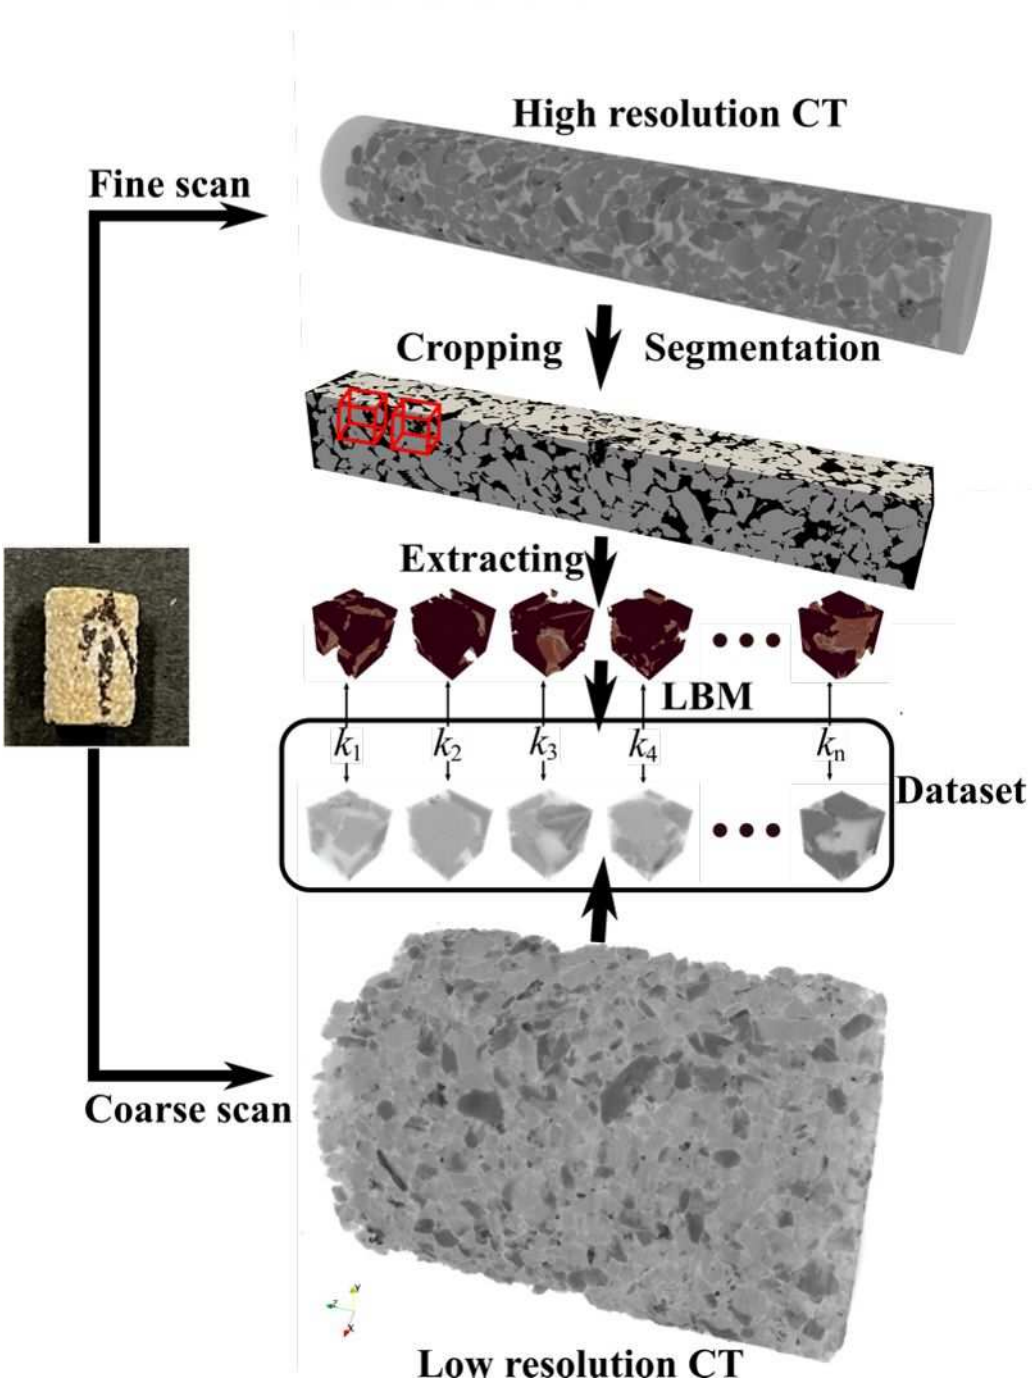


Figure 2.4: Training dataset acquisition for the workflow-development samples. Each data pair consists of a low-resolution raw CT subvolume and the permeability calculated by LBM from the corresponding registered high-resolution subvolume.

For the larger verification core, the high-resolution CT image was filtered and segmented using the same procedure. A binary pore-geometry model with dimensions of $660 \times 650 \times 17230$ voxels was reconstructed from the high-resolution cylindrical

region. The model was partitioned into three high-resolution subvolume sizes: 100^3 voxels for Case I, 200^3 voxels for Case II, and 300^3 voxels for Case III (Figure 2.5). Adjacent subvolumes overlapped by 50 voxels during extraction.

Absolute permeability was calculated using LBM for single-phase flow in the z-direction. The resulting label was therefore the scalar effective permeability along the core flow direction. Porosity was calculated directly from the binary high-resolution CT image as the fraction of pore voxels.

The corresponding low-resolution raw CT subvolumes were extracted at the same physical locations. Because the voxel-size ratio between the high-resolution and low-resolution scans was approximately 6.7, the corresponding low-resolution subvolume sizes were 15^3, 31^3, and 46^3 voxels for Cases I, II, and III, respectively. The total numbers of extracted subvolumes were 48,960, 34,000, and 21,120 for Cases I, II, and III, respectively (Table 2.1). Ninety percent of each dataset was used for training, and the remaining 10% was used for testing.

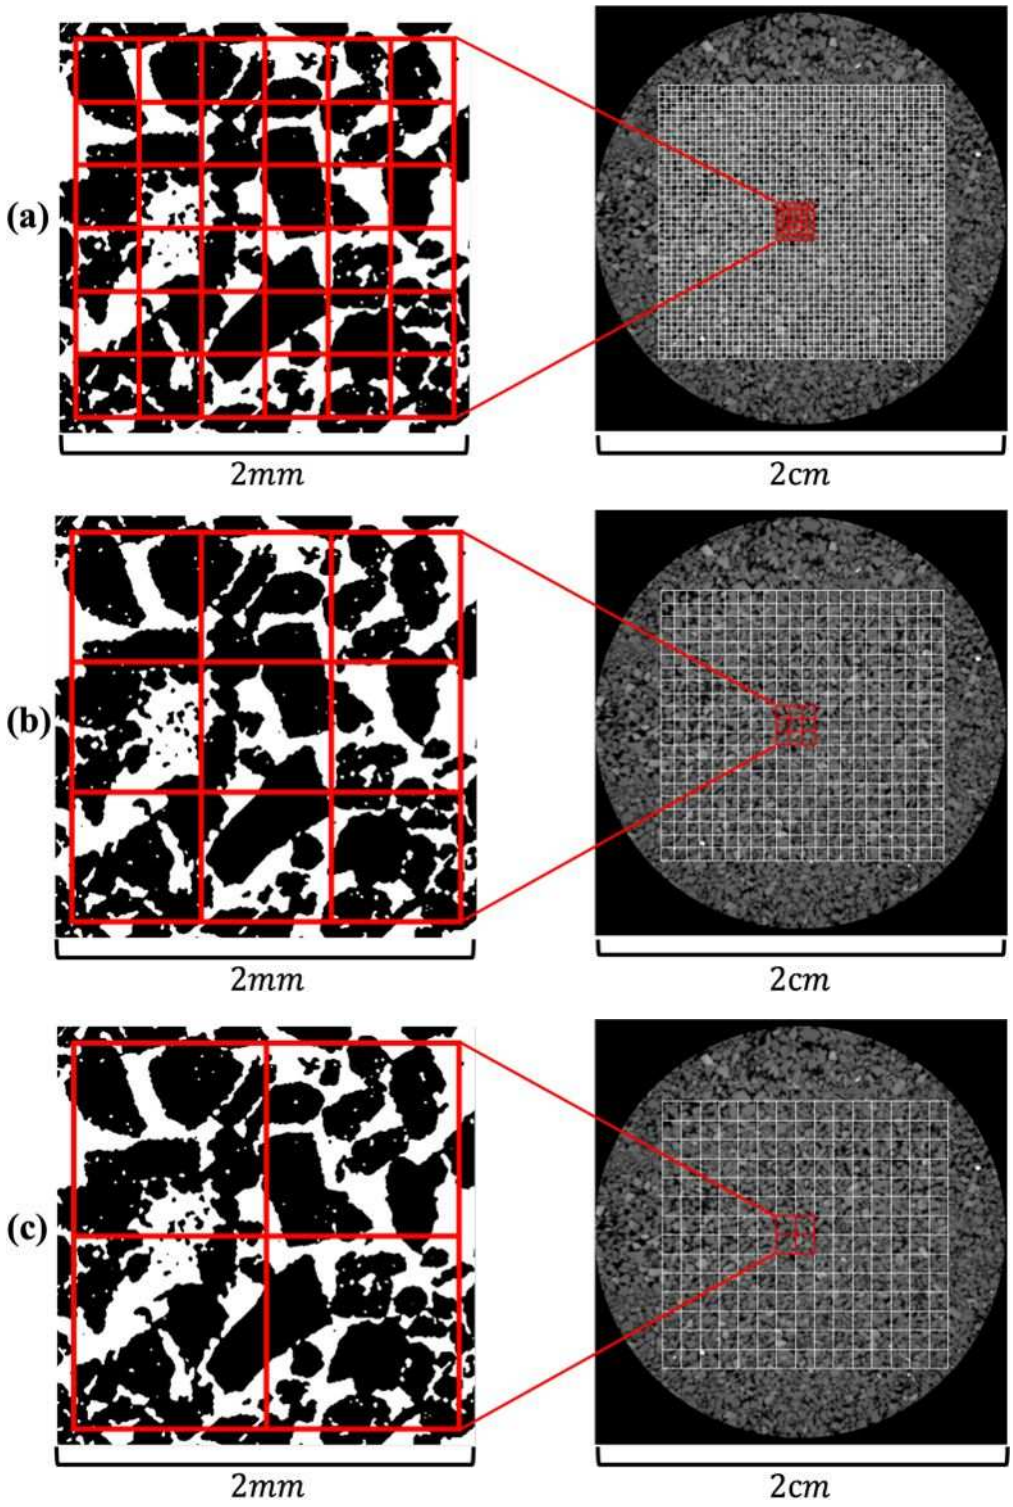


Figure 2.5: Subvolume extraction for the larger verification core. High-resolution binary CT subvolumes are shown on the left, and the corresponding low-resolution raw CT subvolumes are shown on the right for (a) 100^3 voxels (Case I), (b) 200^3 voxels (Case II), and (c) 300^3 voxels (Case III).

Case	I	II	III
High-resolution subvolume size (3 μm/voxel)	100^3	200^3	300^3
Low-resolution subvolume size (20 μm/voxel)	15^3	31^3	46^3
Number of subvolumes	48,960	34,000	21,120
Average pore-body count from PNM	74	421	1240

Table 2.1: Subvolume dimensions, number of subvolumes, and average number of pore bodies for the larger verification core.

The average pore-body count in each subvolume was examined using pore network modeling (PNM). The 100^3, 200^3, and 300^3 subvolumes contained approximately 74, 421, and 1240 pore bodies on average, respectively. The pore-body count was used as a structural indicator of subvolume representativeness. A subvolume defined only by voxel dimensions depends on the CT resolution and the characteristic grain size of the rock. For example, a 100^3 subvolume at 3 μm/voxel contains approximately 74 pore bodies in the Boise sandstone studied here, whereas a rock with larger grains would contain fewer pore bodies within the same voxel dimensions. The average pore-body count therefore provides information that cannot be obtained from voxel dimensions alone.

Coordination-number distributions were derived from the pore network models to characterize pore connectivity (Figure 2.6). The coordination number represents the number of throats connected to each pore body and is related to local flow connectivity. Case III had the highest average coordination number, 3.98, and a relatively narrow distribution. Case I had a lower average coordination number, 3.09, and a broader distribution.

Small subvolumes preferentially sample local pore-rich or solid-rich regions and may contain truncated, dead-end, or disconnected pore branches. These structures reduce the apparent coordination number. Larger subvolumes include a greater portion of the connected pore network, reducing the influence of truncation at the subvolume boundaries. The resulting coordination-number distribution is therefore more representative of the pore network within the larger rock sample. This difference affects both the stability of the calculated permeability labels and the image features available for CNN training.

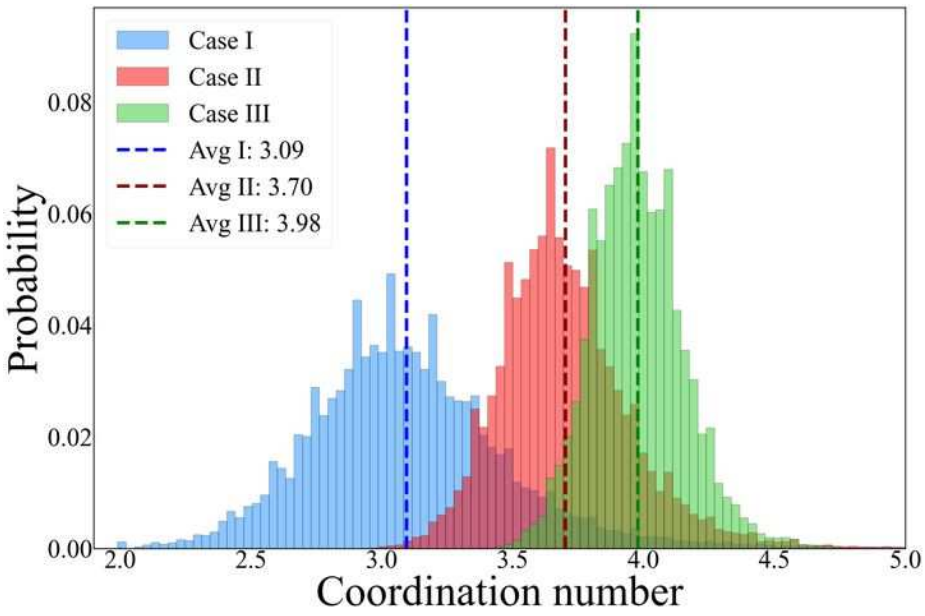


Figure 2.6: Coordination-number distributions for pore networks extracted from each subvolume size. A higher coordination number indicates a larger number of throats connected to each pore body. Case III has the highest average value and the narrowest distribution.

Porosity was also calculated for 400^3, 500^3, and 600^3 high-resolution cubic subvolumes using the same binary CT image. Together with Cases I–III, these data were used to examine how the porosity distribution changes for subvolume sizes from 100^3 to 600^3 voxels. The 5th–95th percentile interval denotes the range between the 5th and 95th percentiles of the porosity distribution and is used as a measure of its width.

2.3.3 Registration of images at different resolutions

The high-resolution and low-resolution CT scans of each sample have different voxel sizes and fields of view. The same physical region must therefore be identified in both images. During scanning, the sample position was adjusted to maintain a consistent scanning angle between the two scans. The high-resolution image was then registered to the low-resolution image using CT slices located at the same depth from the top of the core.

Figure 2.7 shows an example of registration between the two resolutions. The high-resolution region is located within the low-resolution whole-core image, and the corresponding low-resolution raw CT subvolumes are extracted from the registered locations. The registration associates each LBM permeability value calculated from

a high-resolution subvolume with the corresponding low-resolution CNN input.

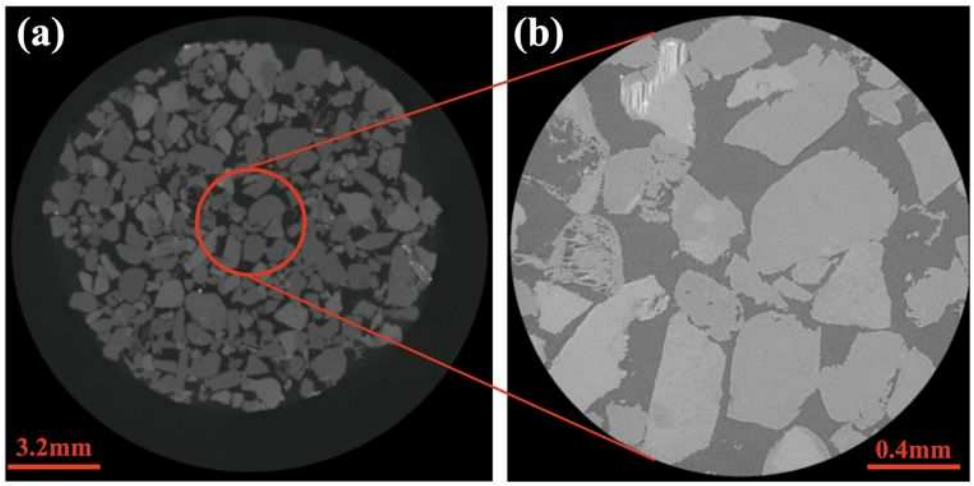


Figure 2.7: Registration of CT images at two resolutions: (a) low-resolution CT image of the whole core and (b) high-resolution CT image of the central region.

2.3.4 Direct flow simulation for permeability labels

Direct pore-scale simulation has been used to obtain flow properties from digital rock CT images Boek and Venturoli (2010); Yang and Boek (2013); Jiang and Tsuji (2017); Yang et al. (2019). In this chapter, the lattice Boltzmann method (LBM) was used to calculate single-phase flow in the segmented pore space. The LBM can simulate single-phase and multiphase flow directly on three-dimensional digital rock images Liu et al. (2021); Li et al. (2021); Jiang et al. (2021a); Zhang et al. (2022); Jiang et al. (2022). It approximately recovers the Navier–Stokes equations from kinetic theory Shan et al. (2006).

In the LBM, the particle distribution function $f_i(\mathbf{x}, t)$ represents the probability of finding a fluid pseudo-particle with lattice velocity $\mathbf{e}_i$ at position $\mathbf{x}$ and time t:

$$f_i(\mathbf{x} + \mathbf{e}_i \delta t, t + \delta t) = f_i(\mathbf{x}, t) + \Omega_i(\mathbf{x}, t), \tag{2.1}$$

where $\mathbf{e}_i$ is the lattice velocity in the ith direction, δt is the time step, and Ω_i is the collision operator. A three-dimensional lattice with 19 velocity vectors (D3Q19) was used Guo and Shu (2013). The multiple-relaxation-time (MRT) collision operator was adopted to improve numerical stability and reduce viscosity dependence Pan et al. (2006); d'Humières (1992).

Single-phase flow was simulated until steady state was reached. Convergence was determined by monitoring the change in the average pore velocity. The simulation

was considered converged when the relative change in average velocity was below 10^{-5}. The scalar permeability in the imposed flow direction, k, was calculated as

$$k = \frac{\mu \, \bar{v} \, L}{\Delta P}, \tag{2.2}$$

where $\bar{v}$ is the mean velocity in the flow direction, μ is the dynamic viscosity, L is the domain length in the flow direction, and ΔP is the equivalent pressure drop over the domain length.

A constant body force was applied to represent the pressure gradient, and no-slip boundary conditions were imposed on the grain surfaces. Therefore, ΔP in Eq. (2.2) represents the pressure drop equivalent to the imposed body force over the length L. The pressure gradient was set sufficiently low to maintain Stokes-flow conditions.

Periodic boundary conditions were applied at the outer boundaries of the computational subvolume. In the imposed flow direction, the sample was mirrored to maintain pore-space connectivity across the periodic boundary. The periodic boundary condition can affect transverse flow components and the off-diagonal components of the permeability tensor. In this chapter, only the diagonal permeability component aligned with the imposed flow direction is used for CNN training and Darcy-scale upscaling Guibert et al. (2016).

2.3.5 CNN models for permeability prediction

Convolutional neural networks (CNNs) were used for regression from low-resolution raw CT subvolumes to scalar permeability values. CNNs apply a sequence of nonlinear transformations through convolutional layers LeCun et al. (1998); Alqahtani et al. (2021). Six three-dimensional CNN architectures were evaluated during workflow development: Conv8, VGG16, VGG19 Simonyan and Zisserman (2014), GoogLeNet-V1, GoogLeNet-V3 Szegedy et al. (2015), and ResNet34 He et al. (2016). The detailed network architectures are provided in Appendix A.

Conv8 was the shallowest architecture evaluated in this chapter and was used as the baseline model. It consists of four convolution blocks, each containing a 3D convolutional layer, batch normalization Ioffe and Szegedy (2015), and a rectified linear unit (ReLU) activation function Nair and Hinton (2010). The convolution blocks are followed by global average pooling (GAP) and dense regression layers. The GAP layer converts the final feature map into a vector and passes it to the fully connected regression layers. The 3D convolutional layers extract image features in the three spatial directions.

VGG16 and VGG19 are deeper convolutional architectures based on repeated small convolutional filters. In the three-dimensional versions used here, $3 \times 3 \times 3$ filters are used throughout the convolution blocks. Stacking these filters increases the receptive field and the number of nonlinear transformations. VGG16 contains 16 weight layers and five max-pooling layers, whereas VGG19 contains three additional convolutional layers.

GoogLeNet uses inception modules in which convolutional filters of different sizes are applied in parallel and their outputs are concatenated. GoogLeNet-V3 factorizes convolutional operations to reduce the number of calculations and increase network depth. For the regression problem in this chapter, the final classification layers were replaced by fully connected layers and a scalar regression output. The auxiliary classification branches used in the original GoogLeNet architectures were removed.

ResNet uses shortcut connections to address vanishing gradients and degradation in deep networks He et al. (2016). A basic ResNet block contains two convolutional layers, each followed by batch normalization and ReLU activation. The shortcut connection adds the input of the block to the output of the second convolutional layer before the final ReLU activation. This structure allows the network to learn residual mappings.

All networks were implemented using TensorFlow Abadi et al. (2015). Based on the workflow-development comparison, ResNet34 was used for the larger-core verification.

2.3.6 Loss functions and accuracy measurements

Five loss functions were tested during CNN training: mean squared error (MSE), mean absolute error (MAE), Huber loss, mean squared logarithmic error (MSLE), and log-cosh loss. Their definitions are listed in Table 2.2. Here, y_i is the LBM permeability label for the ith sample, y_i^p is the predicted value, and n is the number of samples in the dataset.

MAE and MSE are commonly used for regression problems Botchkarev (2018). Huber loss is quadratic for residuals below the threshold δ and linear for larger residuals. The threshold was set to $\delta = 1.0$. MSE applies a quadratic penalty to large residuals. MSLE calculates residuals after logarithmic transformation, reducing the contribution of large absolute differences. Log-cosh loss approaches MSE for small residuals and MAE for large residuals.

Although different loss functions were used during training, MAE was used as the

common metric for comparing prediction accuracy. The loss-function test described in Section 2.4.2 used the Conv8 network. The MAE calculated from the test dataset was used to select the loss function for the subsequent comparison of the six CNN architectures.

Loss function	Formula
MSE	$\frac{1}{n}\sum_{i=1}^{n}(y_i - y_i^p)^2$
MAE	$\frac{1}{n}\sum_{i=1}^{n}\lvert y_i - y_i^p\rvert$
Huber	$\frac{1}{n}\sum_{i=1}^{n}\begin{cases}\frac{1}{2}(y_i - y_i^p)^2, & \lvert y_i - y_i^p\rvert \le \delta,\\ \delta\left(\lvert y_i - y_i^p\rvert - \frac{1}{2}\delta\right), & \lvert y_i - y_i^p\rvert > \delta,\end{cases}$
MSLE	$\frac{1}{n}\sum_{i=1}^{n}\left[\log(y_i+1) - \log(y_i^p+1)\right]^2$
Log-cosh	$\frac{1}{n}\sum_{i=1}^{n}\log\left[\cosh(y_i^p - y_i)\right]$

Table 2.2: Definitions of the loss functions used for CNN training.

2.3.7 Darcy flow solver for whole-core upscaling

The permeability values of subvolumes in the high-resolution central region were calculated directly using LBM. The remaining low-resolution core domain was divided into non-overlapping subvolumes of the same physical size, and their permeability values were predicted using the trained CNN. The resulting whole-core permeability map was used in a Darcy flow solver to calculate the effective permeability of the core.

The MATLAB Reservoir Simulation Toolbox (MRST) Lie (2019b,a) was used to solve the single-phase flow problem. The non-overlapping subvolumes defined the Cartesian grid, and the governing equations were discretized using a two-point finite volume scheme. For an incompressible single phase with constant viscosity,

$$\nabla \cdot \mathbf{v} = q, \qquad \mathbf{v} = -\frac{\mathbf{K}}{\mu}\nabla P, \tag{2.3}$$

where $\mathbf{v}$ is the Darcy velocity, q is the source or sink term, $\mathbf{K}$ is the permeability tensor, μ is the dynamic viscosity, and P is pressure.

The CNN-predicted subvolume permeability was assigned to the diagonal permeability component aligned with the imposed flow direction, and the off-diagonal components were neglected. Fixed pressure boundary conditions were imposed on

the inlet and outlet faces to generate a pressure gradient. No-flow boundary conditions were imposed on the remaining faces. After convergence, the pressure in each grid cell and the flux across each cell face were obtained. The upscaled permeability was calculated from the total flux through a cross-section of the core using Darcy's law.

2.4 Workflow Development and Model Selection

2.4.1 Permeability labels from direct flow simulation

The permeability values calculated from the high-resolution subvolumes using LBM were used as labels for CNN training. Figure 2.8 shows the histograms of calculated permeability on a logarithmic scale for Sample I and Sample II. Because the two samples were extracted from different locations in the same meter-scale Boise sandstone sample, their permeability distributions differ. The mean permeabilities were 10.8 D for Sample I and 5.03 D for Sample II. The standard deviation was smaller for Sample I than for Sample II. Thus, the permeability values of Sample I are more concentrated, whereas those of Sample II are distributed over a wider range.

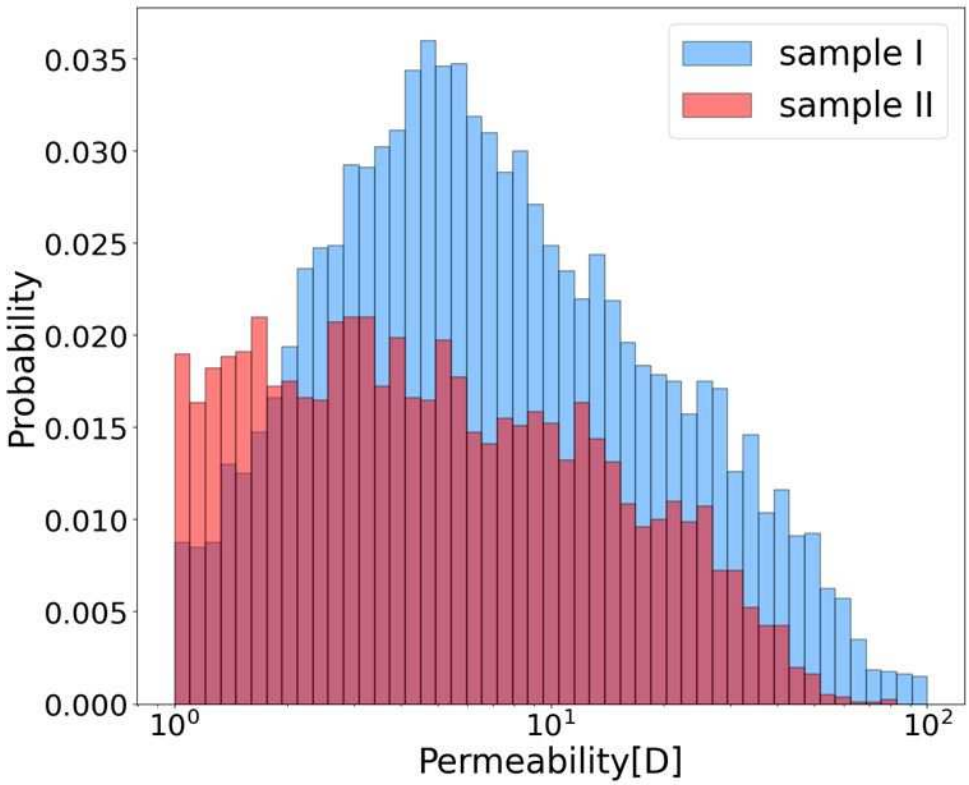


Figure 2.8: Histograms of the calculated permeability on a logarithmic scale for Sample I and Sample II.

2.4.2 Evaluation of loss functions

The five loss functions listed in Table 2.2 were evaluated using the same training and testing datasets. Table 2.3 summarizes the MAE obtained when the Conv8 network was trained on Sample I using each loss function. Huber loss gave the lowest MAE. In the training dataset, some high-porosity subvolumes have large permeability values, which produce large residuals during training. Huber loss reduces the contribution of these residuals compared with MSE Alqahtani et al. (2020); Huber (1992). Huber loss was therefore used for the subsequent network comparison, while MAE was retained as the common accuracy metric.

Loss function	MSE	MAE	Huber	MSLE	Log-cosh
Accuracy (MAE)	1.8957	1.9384	1.7639	2.8808	1.9488

Table 2.3: Accuracy of the five tested loss functions for Sample I using Conv8. The reported metric is MAE.

2.4.3 Performance of different CNN architectures

After selecting Huber loss, the six CNN architectures were evaluated for permeability prediction. Figures 2.9 and 2.10 compare the LBM-calculated permeability values with the CNN-predicted values for Sample I. The red diagonal line represents $y = x$. The predictions from GoogLeNet and ResNet are closer to this line than those from Conv8 and VGG, particularly in the high-permeability range. ResNet34 gave the highest accuracy for the test dataset. Therefore, only the ResNet34 prediction result is shown for Sample II in Figure 2.11.

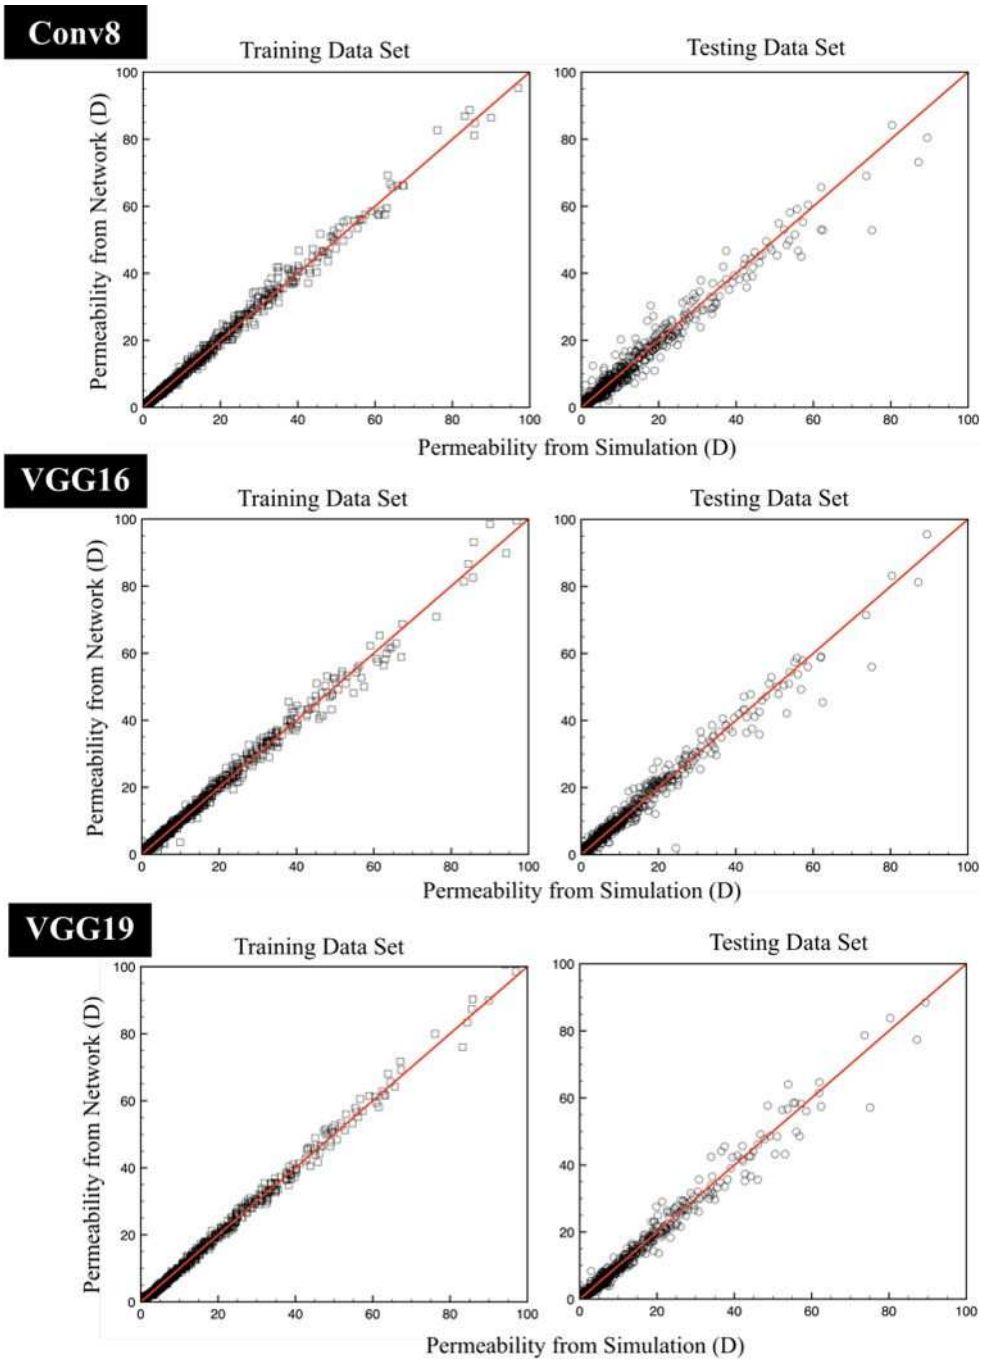


Figure 2.9: Permeability values calculated by LBM plotted against CNN predictions for Sample I using Conv8, VGG16, and VGG19. The red diagonal line represents $y = x$.

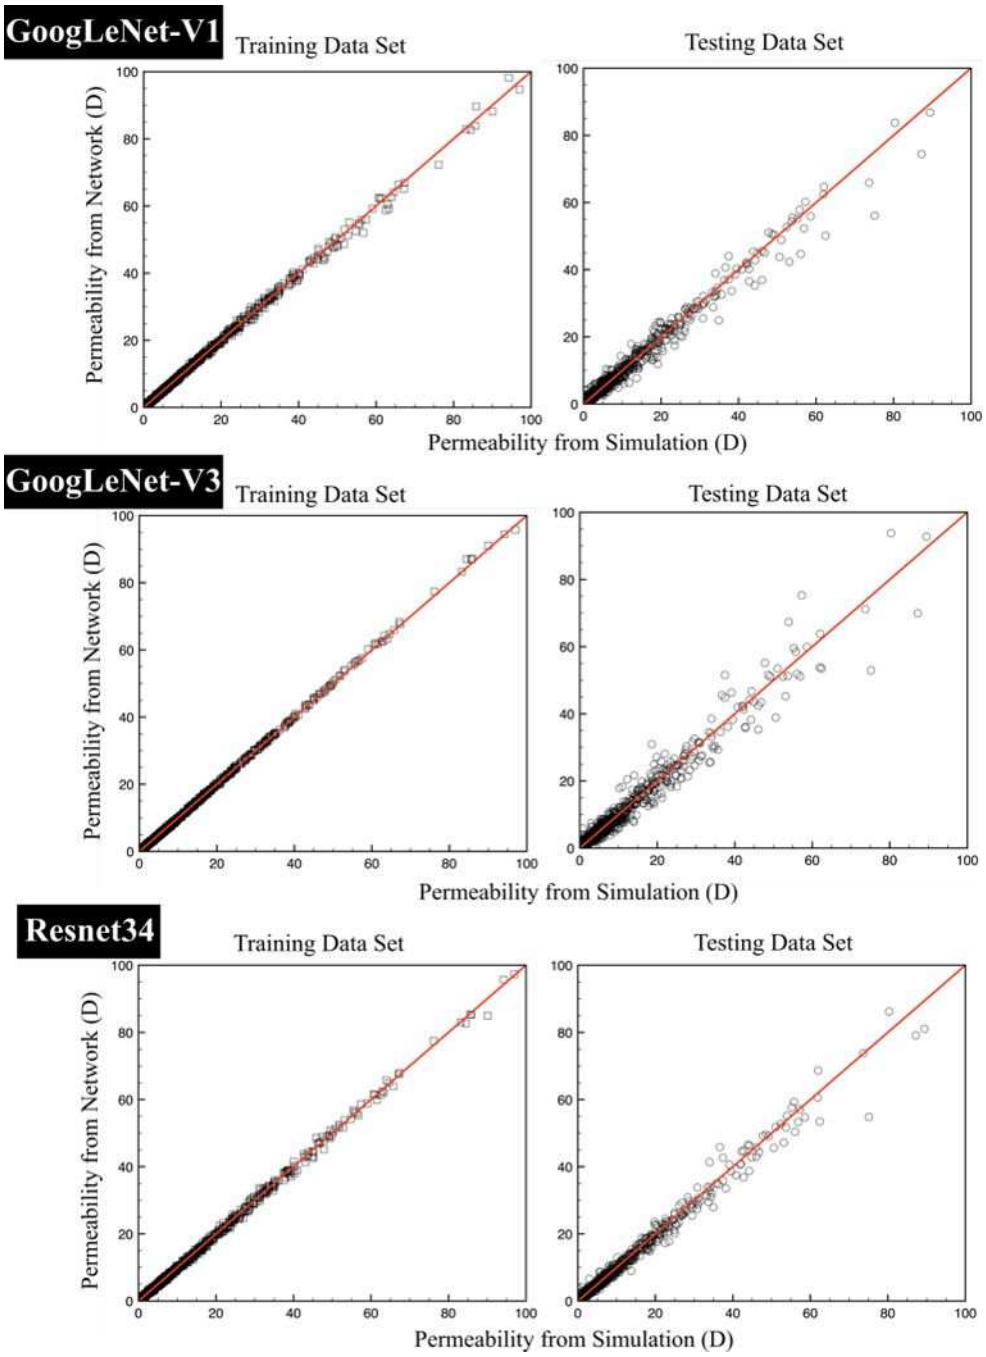


Figure 2.10: Permeability values calculated by LBM plotted against CNN predictions for Sample I using GoogLeNet-V1, GoogLeNet-V3, and ResNet34. The red diagonal line represents $y = x$.

Table 2.4 summarizes the regression accuracy and training time of the six networks. ResNet34 achieved the lowest MAE for both samples, with values of 0.4546 for Sample I and 0.2249 for Sample II. Prediction accuracy generally increased with network depth. However, the comparison between VGG16 and VGG19 shows that

adding convolutional layers to a conventional sequential network does not consistently improve accuracy. In deeper conventional networks, backpropagation becomes more difficult, and additional layers can increase the risk of convergence to local optima or overfitting. The shortcut connections in ResNet34 reduce these problems and allow the deeper network to be trained more effectively.

High-permeability subvolumes make a large contribution to MAE. Such subvolumes are relatively rare, and the number of corresponding training samples is limited. Their prediction errors therefore increase the MAE, particularly for Sample I. In Sample II, the permeability values are lower and more concentrated, resulting in a lower MAE.

The training was performed on a workstation with 64 GB of RAM and one Quadro RTX 5000 GPU. ResNet34 required the longest training time among the tested networks, whereas GoogLeNet required less time than the other deep networks. Based on the prediction results, ResNet34 was selected for whole-core permeability prediction and for the subsequent verification using the larger core.

Model	Number of layers	MAE (Sample I)	MAE (Sample II)	Training time (Sample
Conv8	8	1.4226	0.8016	4 h 59 min
VGG16	16	1.2031	0.6939	13 h 22 min
VGG19	19	1.1348	0.9163	19 h 24 min
GoogLeNet-V1	22	1.1729	0.6141	5 h 2 min
GoogLeNet-V3	25	0.7589	0.4332	4 h 16 min
ResNet34	34	0.4546	0.2249	45 h 51 min

Table 2.4: Regression accuracy and training time for the tested CNN architectures.

Figure 2.11 compares the ResNet34 predictions with the LBM-calculated permeability values for Sample II. The predicted values are distributed close to the $y = x$ line, consistent with the MAE of 0.2249 reported in Table 2.4.

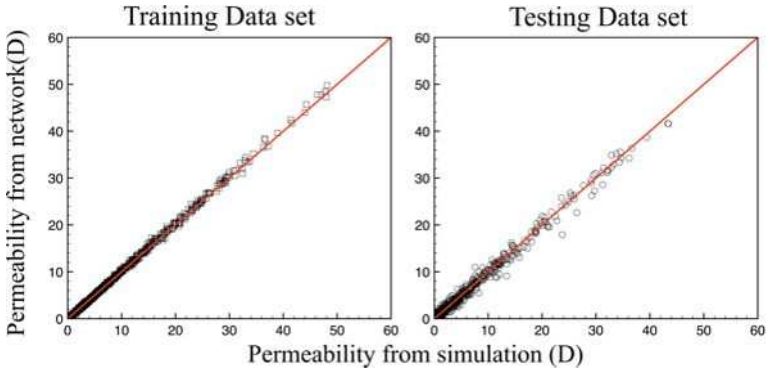


Figure 2.11: Permeability values calculated by LBM plotted against ResNet34 predictions for Sample II. The red diagonal line represents $y = x$.

2.4.4 Porosity prediction

Porosity was also predicted directly from the low-resolution raw CT subvolumes using ResNet34. The reference porosity values were calculated from the same segmented high-resolution subvolumes used for the LBM simulations. Figures 2.12 and 2.13 compare the reference and predicted porosity values for Samples I and II.

The MAE values were 0.00193 for Sample I and 0.00152 for Sample II. These errors are smaller than those obtained for permeability. Porosity is primarily determined by the pore-volume fraction and the grayscale distribution of the CT image, whereas permeability also depends on pore connectivity and flow-path tortuosity. The results show that porosity can be predicted directly from raw low-resolution CT images without segmenting the low-resolution images.

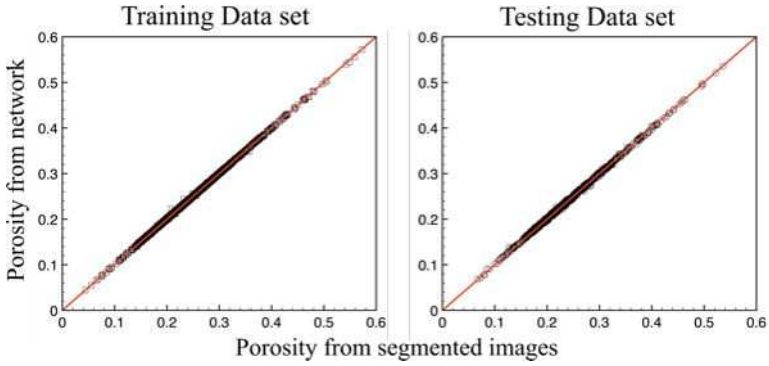


Figure 2.12: Porosity calculated from segmented high-resolution CT images plotted against ResNet34 predictions for Sample I. The red diagonal line represents $y = x$.

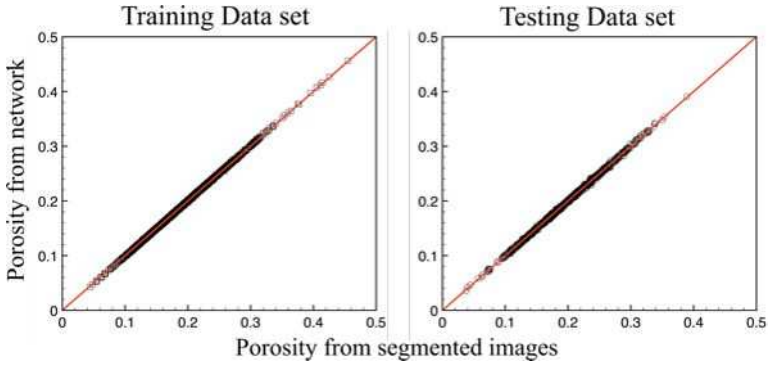


Figure 2.13: Porosity calculated from segmented high-resolution CT images plotted against ResNet34 predictions for Sample II. The red diagonal line represents $y = x$.

2.4.5 Permeability maps predicted by the trained CNN

After ResNet34 was trained, the low-resolution whole-core domain was divided into non-overlapping subvolumes, and the permeability of each subvolume was predicted directly from its raw CT values. This prediction included subvolumes outside the high-resolution training region. Figure 2.14 shows the predicted permeability distributions for Samples I and II. The predicted distributions retain the difference observed in the LBM labels: Sample I has a higher permeability distribution than Sample II.

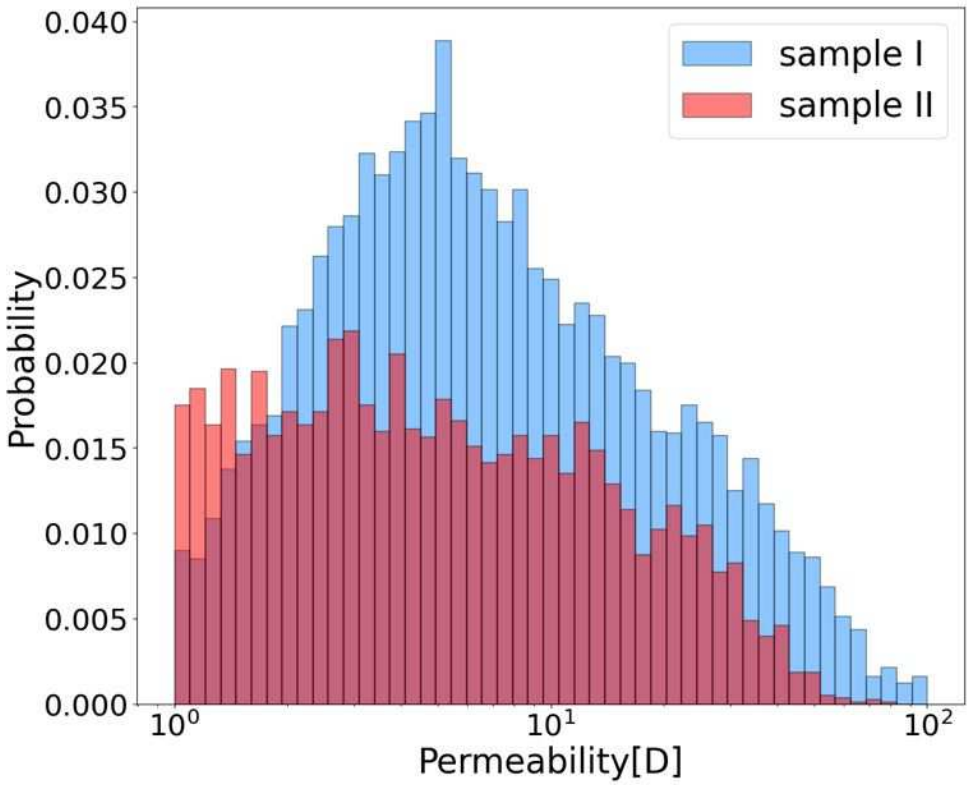


Figure 2.14: Histograms of the predicted permeability on a logarithmic scale for Sample I and Sample II.

The permeability labels used for training were calculated along the core-axis direction. To estimate the diagonal permeability components in the two perpendicular directions, the trained network was also applied to CT subvolumes rotated along the corresponding axes. The off-diagonal components of the permeability tensor were omitted.

Figure 2.15 shows the predicted whole-core permeability maps and the rectangular regions used for the Darcy flow calculation. Sample I has higher predicted permeability than Sample II, consistent with the LBM-derived label distributions. The rectangular computational regions were extracted from the cylindrical samples to exclude the boundary region, where the CT image quality was lower. The prediction was performed directly from raw low-resolution images and did not require segmentation of these images.

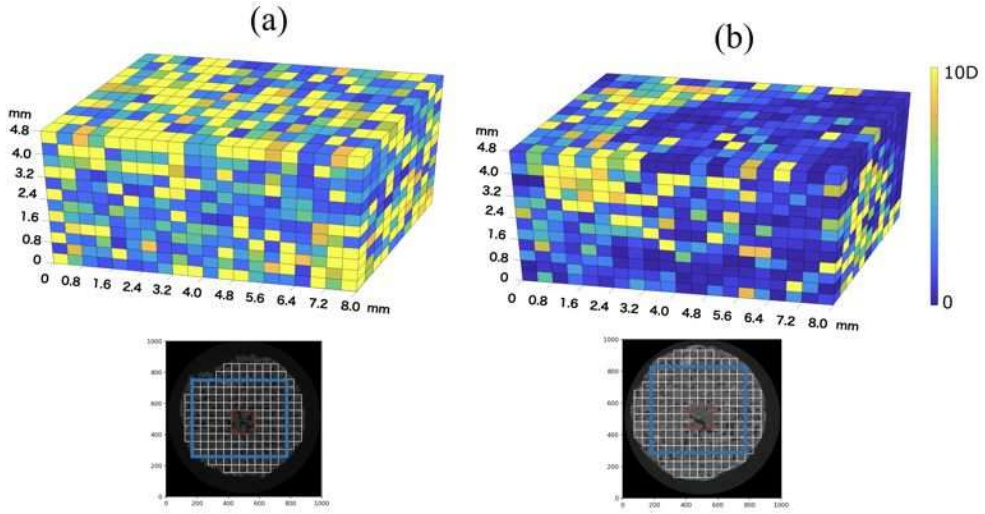


Figure 2.15: Permeability maps for the whole cores and the rectangular regions extracted from the original cylindrical samples for Darcy flow calculation: (a) Sample I and (b) Sample II.

The permeability–porosity relationships were examined for the LBM-calculated data and the CNN-predicted whole-core data. Figures 2.16 and 2.17 show the results for Samples I and II, respectively. The solid lines show fits obtained using the Kozeny–Carman (KC) model,

$$k = C\frac{\phi^3}{(1-\phi)^2}, \tag{2.4}$$

where C is a parameter related to pore geometry and ϕ is porosity. The model was fitted to the computed data using the least-squares method. The fitted values of C were 154.5 for Sample I and 255.9 for Sample II.

Both the calculated and predicted data show an increase in permeability with porosity. They also show a range of permeability values at a given porosity that is not represented by a single KC-model curve. The CNN-predicted data reproduce the distribution observed in the LBM-calculated data.

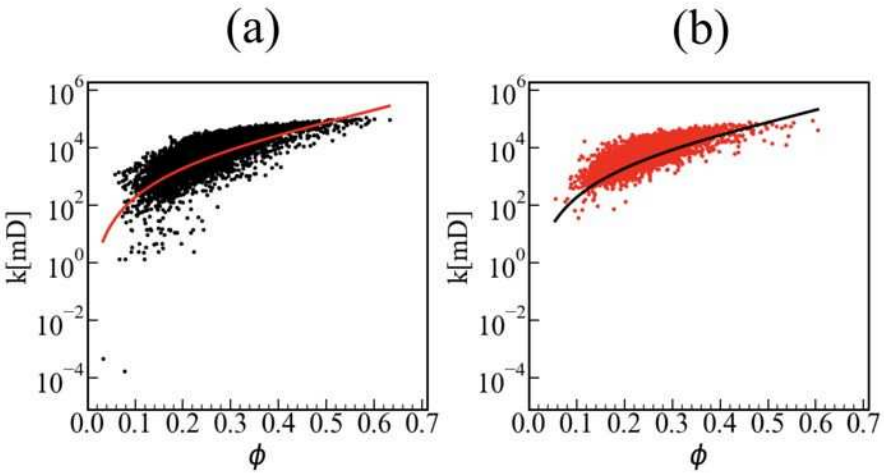


Figure 2.16: Permeability–porosity relationship for Sample I: (a) calculated LBM data and (b) predicted data. The solid line shows the KC-model fit.

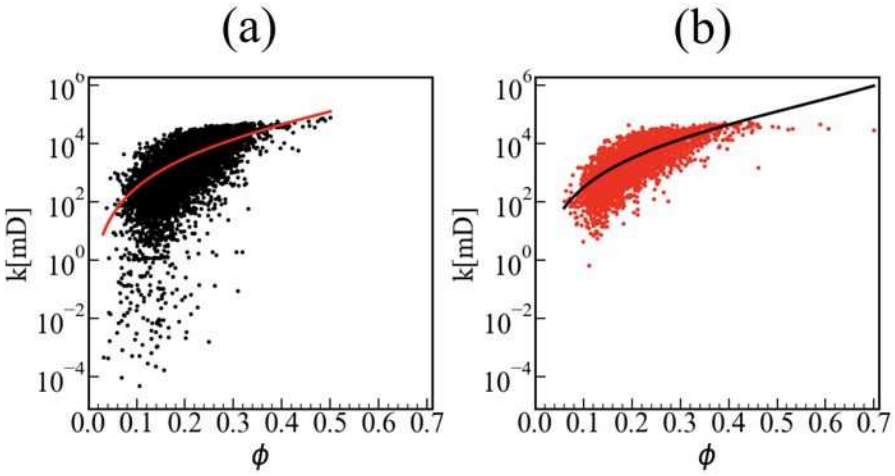


Figure 2.17: Permeability–porosity relationship for Sample II: (a) calculated LBM data and (b) predicted data. The solid line shows the KC-model fit.

For comparison, the permeability data were also examined using a porosity-based scaling relationship derived from permeability bounds Martys et al. (1994):

$$k = \frac{2\phi_2}{s^2}\,(\phi_1)^f\,, \tag{2.5}$$

where s is the specific surface area, defined as the pore-surface area divided by the total sample volume, ϕ_2 is the solid-volume fraction, $\phi_1 = \phi - \phi_c$ is the pore-volume fraction above the critical porosity, ϕ_c is the critical porosity, and f is a

fitting exponent. For the sandstone samples, ϕ_c was set to zero Martys et al. (1994). The fitted values of f were 1.66 for Sample I and 1.84 for Sample II.

Figure 2.18 shows that a range of permeability values remains after scaling. Thus, porosity-based scaling does not reproduce the full permeability distribution. The CNN predictions retain the permeability variation contained in the LBM-calculated data.

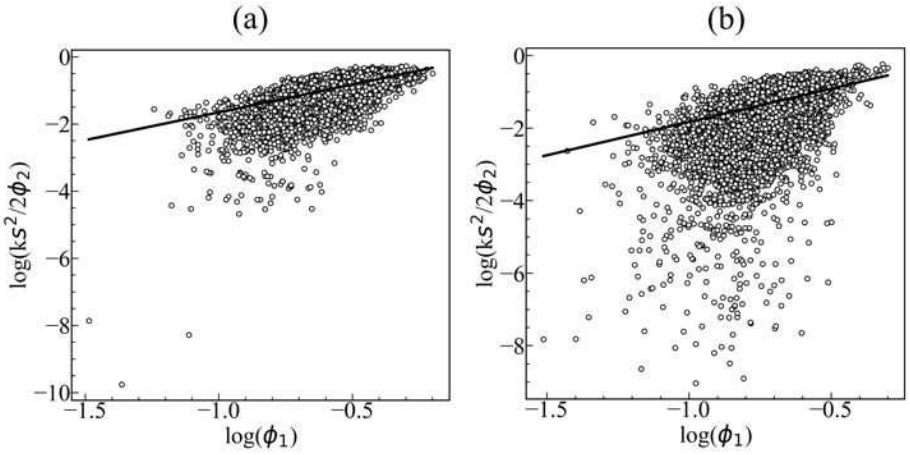


Figure 2.18: Relationship between scaled permeability and porosity for (a) Sample I and (b) Sample II. The solid line shows the fit to the scaled data using $k = \frac{2\phi_2}{s^2}(\phi_1)^f$. The fitted values of f are 1.66 and 1.84 for Samples I and II, respectively.

2.4.6 Whole-core upscaled permeability

The predicted permeability maps were used as input to the Darcy flow solver. Figure 2.19 shows the pressure distributions for Samples I and II. Sample I has a relatively uniform pressure gradient along the flow direction, consistent with its more homogeneous permeability map. Sample II has a less uniform pressure drop because its permeability distribution is more heterogeneous.

Figure 2.20 shows the LBM velocity streamlines in the two high-resolution central regions. Sample I contains more connected high-velocity flow paths through larger pores. Sample II shows lower velocities and more tortuous flow paths, consistent with its lower permeability.

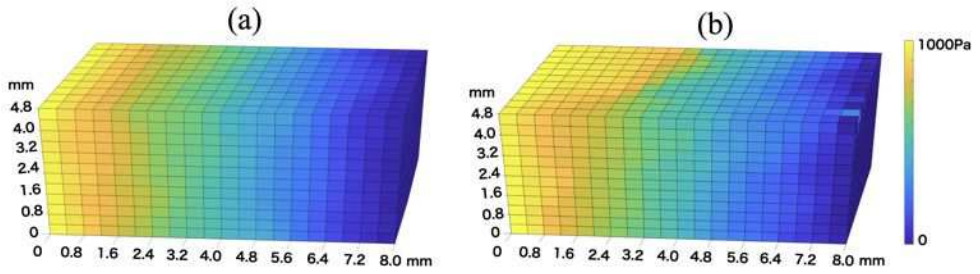


Figure 2.19: Pressure distributions of the whole cores obtained from the Darcy flow solver: (a) Sample I and (b) Sample II.

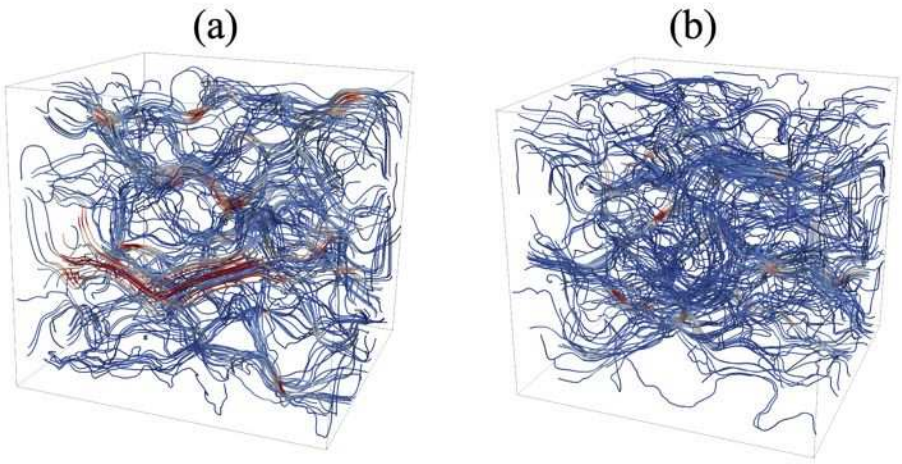


Figure 2.20: Velocity streamlines in the two central regions obtained from direct LBM flow simulation: (a) Sample I and (b) Sample II. Red and blue indicate high and low velocity, respectively.

For comparison, PNM was used to calculate permeability from the high-resolution central regions and the low-resolution whole-core images. PNM represents the pore space using idealized pore bodies and throats and calculates flow through the resulting network Blunt (2001). Table 2.5 summarizes the permeability values obtained by PNM, LBM, CNN-based upscaling, and coreflooding.

The CNN–Darcy workflow gave whole-core permeabilities of 5.85 D for Sample I and 1.53 D for Sample II. Relative to the LBM permeabilities of the high-resolution central regions, 4.83 D and 1.89 D, the corresponding differences were 21% and 19%.

PNM gave 5.15 D and 0.75 D for the high-resolution central regions. For Sample II, the small pore bodies are not represented adequately by the simplified pore-network elements, resulting in a lower PNM permeability than the LBM value. For the low-resolution whole-core images, PNM gave 12.8 D and 3.53 D. At this resolu-

tion, pore throats are represented by intermediate grayscale voxels, and their classification during segmentation affects the reconstructed network. The CNN-based workflow uses the raw low-resolution images and does not require this segmentation step.

Method	Sample I	Sample II
k of center region (PNM; high resolution)	5.15 D	0.75 D
k of whole core (PNM; low resolution)	12.8 D	3.53 D
k of center region (LBM; high resolution)	4.83 D	1.89 D
k of whole core (CNN; low resolution)	5.85 D	1.53 D
Meter-scale coreflooding data	2.8 D	

Table 2.5: Large-scale permeability values obtained by different methods.

The meter-scale coreflooding permeability of the parent Boise sandstone was 2.8 D, which lies between the CNN-upscaled values for Samples I and II. Samples I and II were extracted from relatively high- and low-porosity regions, respectively, which is consistent with their higher and lower predicted permeabilities.

The transferability of the trained network within the same rock type was also examined. The network trained using Sample I was applied to the low-resolution CT image of Sample II. The predicted permeability map in Figure 2.21 is similar to the map obtained using the network trained on Sample II. The resulting upscaled permeability was 2.47 D. This value is higher than the 1.53 D obtained using the Sample II network because the Sample I training dataset contains more high-porosity and high-permeability subvolumes. The value of 2.47 D is also closer to the coreflooding value than the low-resolution PNM result of 3.53 D.

This result shows that a network trained on one sample can be applied to another sample of the same rock type when their pore structures have similar characteristics. Different rock types require dedicated training because their porosity ranges, pore connectivity, and CT-image characteristics differ. A database of trained networks for different rock types could support permeability estimation directly from raw CT images.

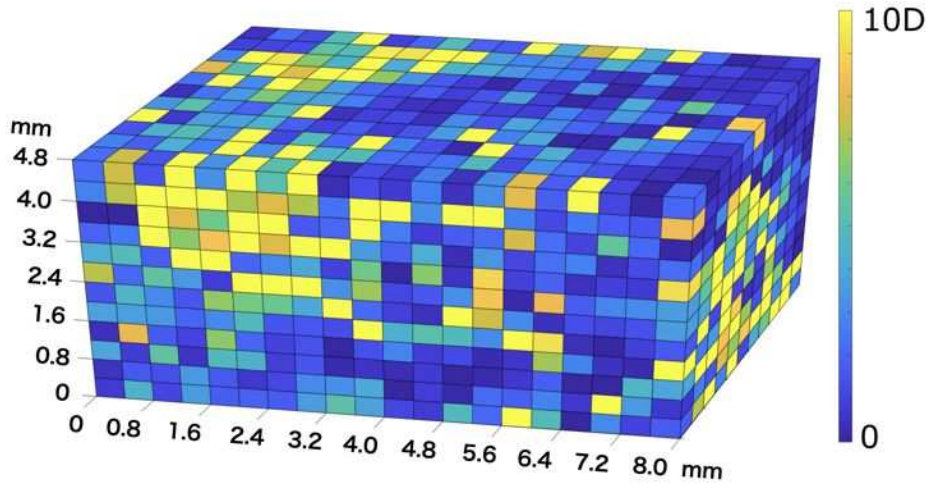


Figure 2.21: Permeability map of Sample II predicted by the neural network trained using data extracted from Sample I.

2.5 Workflow Verification on a Larger Core

This section verifies the permeability upscaling workflow using the larger centimeter-scale Boise sandstone core. The verification focuses on how the subvolume size and the number of training data points affect CNN prediction accuracy, permeability-map reliability, and the final Darcy-scale upscaled permeability. The verification workflow is shown in Figure 2.22.

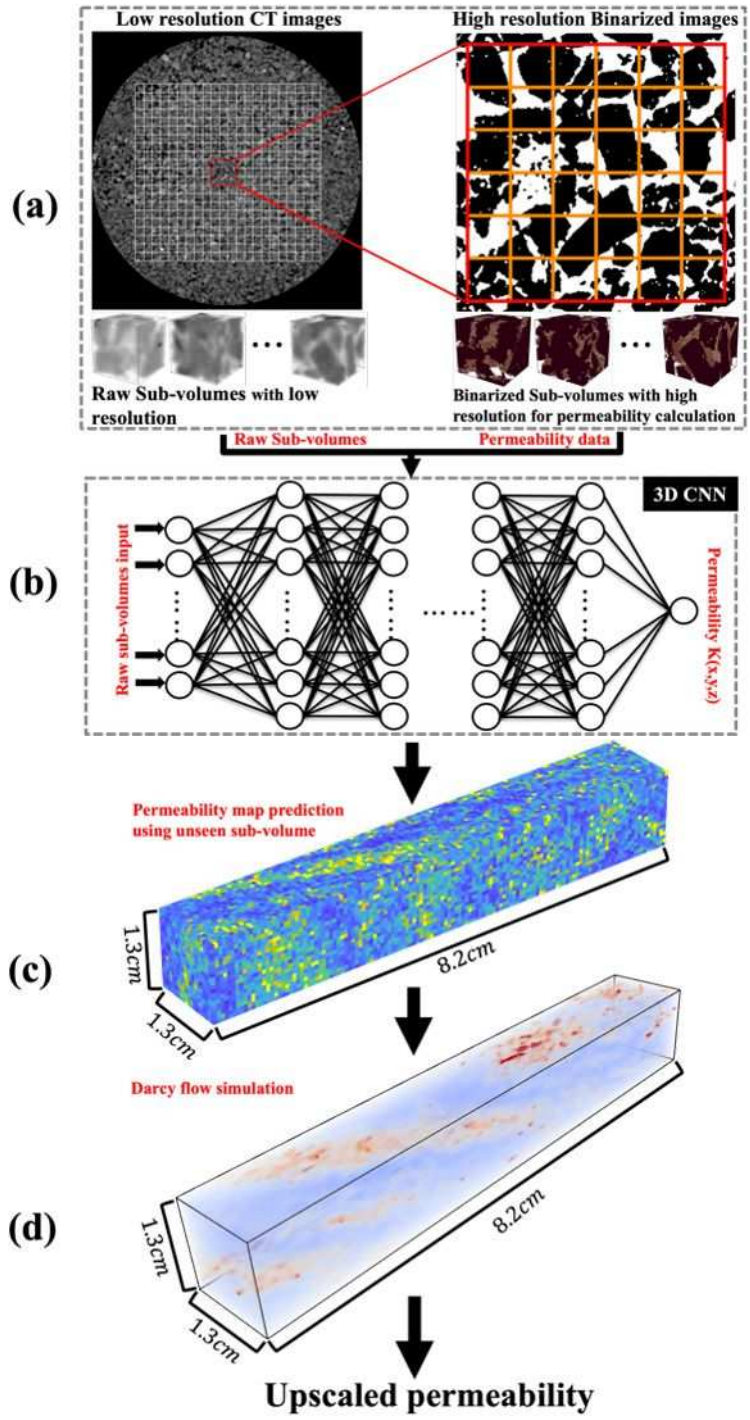


Figure 2.22: Permeability upscaling workflow for the larger core used in workflow verification: (a) CT scanning at two resolutions and creation of training datasets; (b) CNN training using raw CT data and permeability labels; (c) prediction of whole-core permeability maps by the trained CNN; and (d) calculation of flow-flux distribution using the Darcy flow solver. The red box in (a) indicates the selected region for high-resolution analysis, and the orange grid lines indicate the subvolume partition used for permeability computation.

2.5.1 LBM permeability distributions for different subvolume sizes

The larger core was partitioned into three high-resolution subvolume sizes to examine how the subvolume size affects the distribution of LBM-derived permeability labels. Figure 2.23 shows the LBM permeability distributions for Cases I, II, and III. The standard deviations of permeability were 20.73, 12.88, and 8.05 D for Cases I, II, and III, respectively. The distribution therefore becomes narrower as subvolume size increases. This behavior is expected because larger subvolumes average out local heterogeneity and reduce the occurrence of extreme values.

From the perspective of CNN training, data diversity and structural representativeness must be balanced. Case I provides the largest number of training data points and the broadest permeability distribution. At first glance, this broad diversity may appear beneficial for learning a wide range of local CT patterns. However, small subvolumes can contain disconnected flow paths, isolated pores, or nearly solid regions that are not representative of the connected flow network at core scale. These outlier subvolumes can cause overestimation or underestimation of permeability and reduce prediction reliability. Cases II and III reduce this problem because larger subvolumes average out local heterogeneities and reduce the probability of encountering outlier subvolumes. The loss of local diversity is compensated by improved physical representativeness. Therefore, an appropriate subvolume size is required to balance data diversity, training-data quality, and permeability-map prediction accuracy.

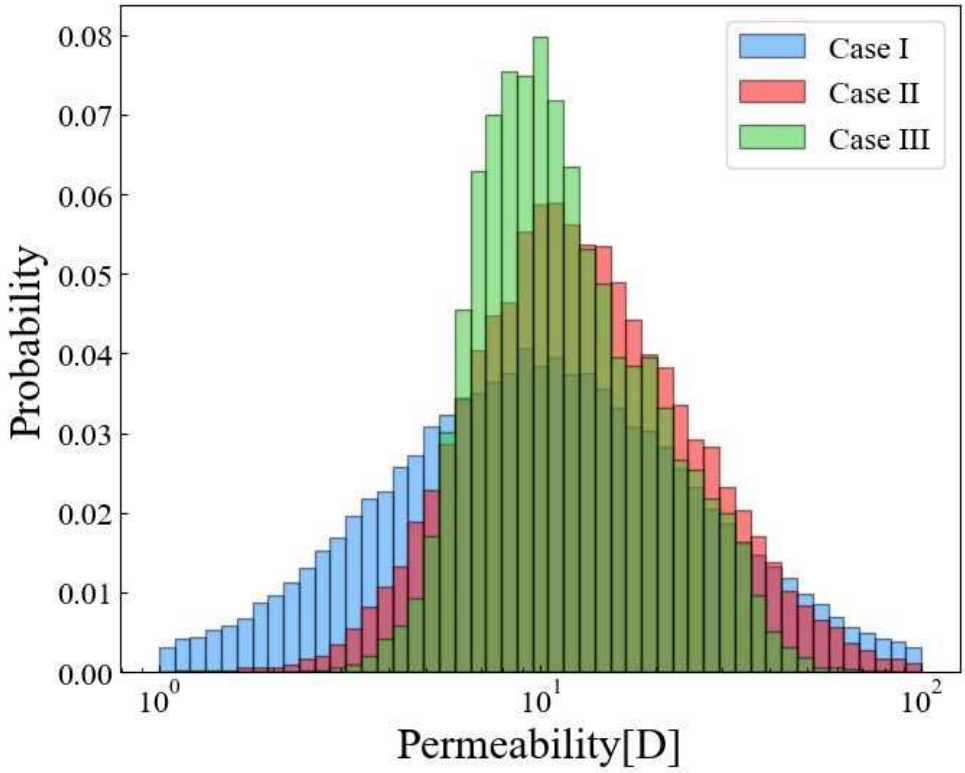


Figure 2.23: Histograms of the calculated permeability on a logarithmic scale for Case I, Case II, and Case III.

Figure 2.24 shows the mean porosity, the 5th–95th percentile interval, and the number of extracted subvolumes for subvolume sizes from 100^3 to 600^3 voxels. The mean porosity remained close to 29% over the tested size range. The 5th–95th percentile width decreased from approximately 31 percentage points for the 100^3 subvolumes to 8.66 percentage points for the 300^3 subvolumes and 5.71 percentage points for the 600^3 subvolumes. This indicates that local pore-volume variation is progressively averaged within larger support volumes.

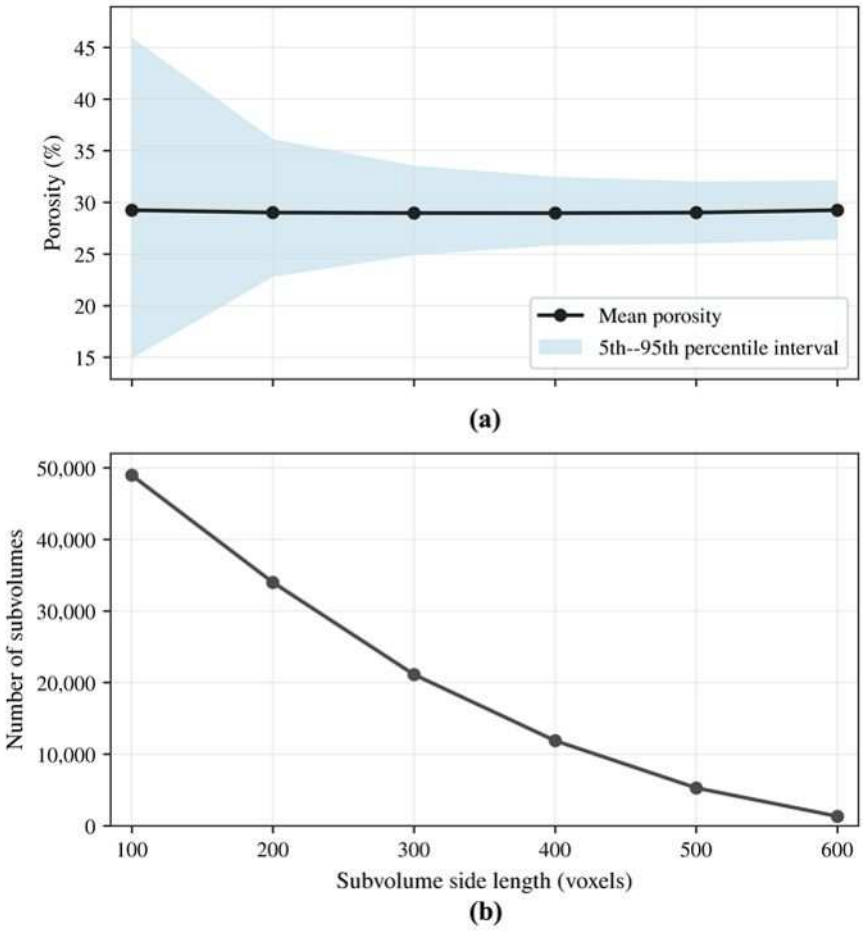


Figure 2.24: Effect of high-resolution subvolume size on porosity statistics and sample availability. The upper panel shows the mean porosity and the 5th–95th percentile interval. The lower panel shows the number of extracted subvolumes.

The reduction in the 5th–95th percentile interval from 100^3 to 300^3 indicates that extremely high- and low-porosity local patches are no longer dominant at the 300^3 support size. This is relevant to permeability prediction because such local patches can generate permeability labels controlled by isolated pores, disconnected flow paths, or nearly solid regions rather than by the connected pore network of the core. At 300^3, these extreme local effects are reduced, while 21,120 subvolumes remain available for training and whole-core mapping.

Increasing the size beyond 300^3 continues to narrow the porosity interval, but the additional narrowing mainly comes from averaging local porosity contrasts within each larger subvolume. In the present workflow, each subvolume is assigned one local permeability value. Therefore, when the subvolume size is increased to 400^3,

500^3, or 600^3, local high- and low-porosity regions that would be treated as separate local blocks at 300^3 are combined into a single larger block. This reduces the local property variation that the permeability map can represent. The number of available subvolumes also decreases from 21,120 at 300^3 to 11,880, 5,280, and 1,320 at 400^3, 500^3, and 600^3, respectively.

2.5.2 Effect of training dataset size on prediction accuracy

ResNet34 was used for the larger-core verification because it showed the best performance in the first-stage model comparison. The network was trained for 300 epochs, and the model parameters giving the best validation performance were stored for prediction. Prediction accuracy was evaluated using the coefficient of determination, R^2. An R^2 value near 1 indicates a close fit between the predicted and reference values, whereas a value near 0 indicates poor regression performance Draper and Smith (1998).

Figure 2.25 shows the relationship between the number of training data points and prediction accuracy for permeability and porosity. For all three subvolume sizes, R^2 increases as the number of training data points increases. Case III reaches an R^2 value close to 1 with fewer training data points and stabilizes rapidly. Case II also reaches high accuracy, but requires more training data points than Case III. Case I has lower initial accuracy and continues to improve as more data are added, but even with 48,960 data points it does not reach the accuracy obtained with the larger subvolumes. In Case II, convergence is achieved with approximately 20,000 data points. In Case III, satisfactory accuracy is achieved with approximately 10,000 data points. Porosity prediction follows the same trend: Case III achieves very high accuracy with relatively few data points, Case II follows closely, and Case I requires more data to reach comparable accuracy.

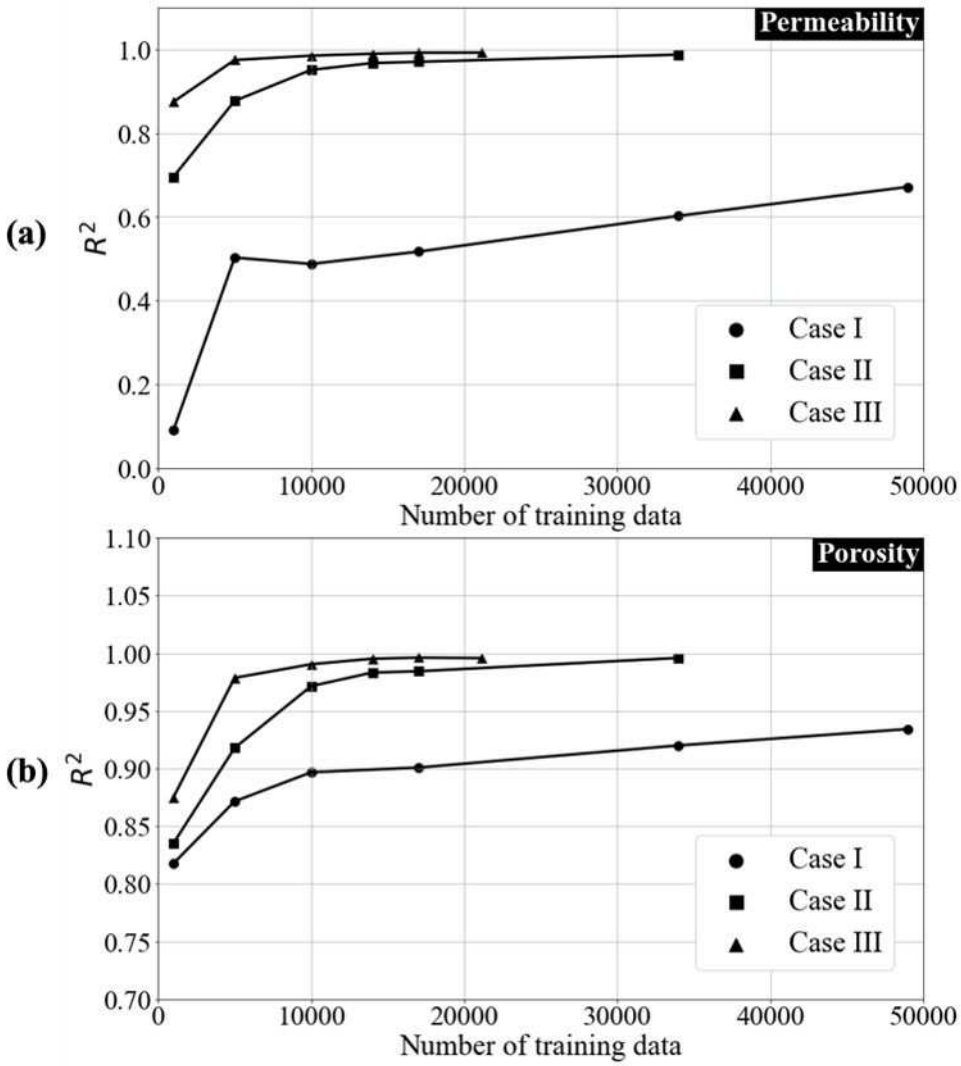


Figure 2.25: Relationship between the number of training data points and prediction accuracy R^2 for (a) permeability and (b) porosity in Cases I, II, and III.

The improvement in accuracy with larger subvolumes is attributed to the inclusion of more representative pore-network features. A larger subvolume contains more pores and throats, which reduces the effect of isolated local features and provides a more stable relationship between low-resolution CT texture and effective permeability. However, larger subvolumes also increase computational cost during CNN training. Table 2.6 shows that the training time increased from 8.3 hr for Case I to 39.3 hr for Case II and 89.8 hr for Case III on a single NVIDIA A6000 GPU. Memory usage also increased with subvolume size because of the larger input dimensionality and feature complexity.

Case	I	II	III
High-resolution subvolume size (voxels)	100^3	200^3	300^3
Training time (hr)	8.3	39.3	89.8
Memory usage (GB)	23.3	25.8	33.2

Table 2.6: CNN training time and memory usage for each subvolume-size case.

The 300^3 case already required 89.8 hr of training and 33.2 GB of GPU memory. Increasing the high-resolution subvolume size from 300^3 to 400^3, 500^3, and 600^3 would increase the number of voxels in one LBM domain and one corresponding CNN input by factors of 2.37, 4.63, and 8.00, respectively. This would increase the cost of generating each permeability label and training the CNN. At the same time, the larger support volumes would combine more local heterogeneity into each single label, reducing the local property variation represented in the whole-core permeability map. Since the 300^3 case already gave the highest prediction accuracy among Cases I–III, the complete workflow was evaluated up to this size.

2.5.3 Prediction accuracy for permeability and porosity

After the training accuracy converged, the prediction performance was evaluated for the three subvolume-size cases by comparing CNN-predicted permeability values with LBM-calculated permeability values for the test dataset (Figure 2.26). Case I shows the largest scatter, and many predicted values fall below the $y = x$ line, indicating a systematic underestimation of permeability for unseen subvolumes. This is because smaller subvolumes contain fewer representative rock features, and the characteristics of individual subdomains can vary greatly. In Case II, the predictions are more tightly clustered around the $y = x$ line, indicating improved accuracy and better generalization compared with Case I. Case III shows the best agreement with the $y = x$ line, making it the most reliable case among the three subvolume sizes for permeability prediction. This result indicates that larger subvolumes capture more representative rock features and reduce the variability between subdomains, thereby improving the learning of effective permeability.

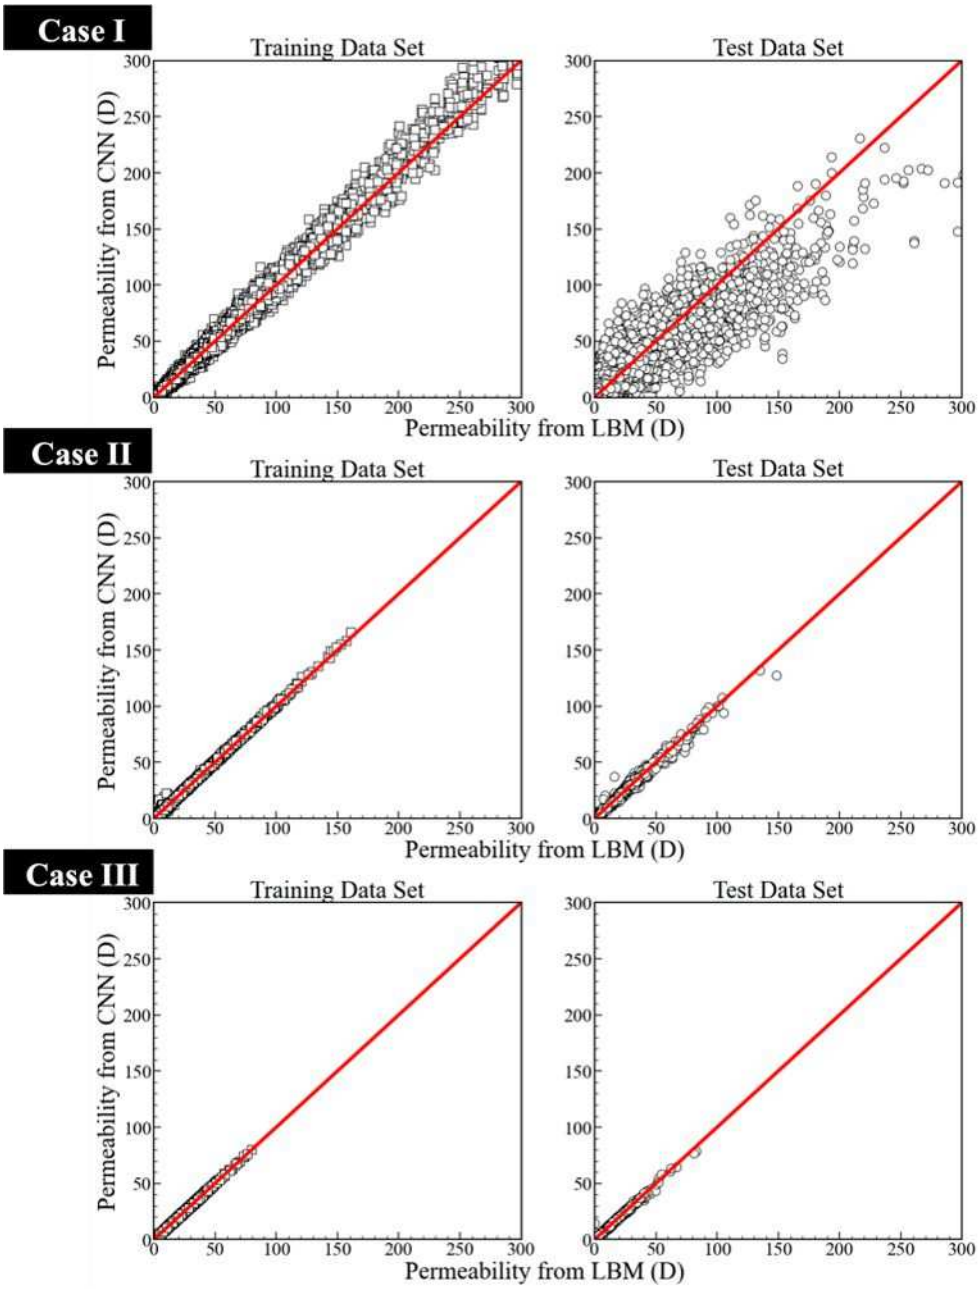


Figure 2.26: Permeability calculated by LBM plotted against CNN-predicted permeability for Cases I, II, and III. The red diagonal line represents $y = x$.

The porosity prediction results show the same tendency (Figure 2.27). Porosity is easier to predict than permeability because it is primarily determined by the grayscale histogram, whereas permeability is also influenced by pore-space connectivity and flow-path tortuosity. Nevertheless, Case I still shows a relatively large prediction error, whereas Cases II and III improve substantially, with Case III showing predictions very close to the reference porosity calculated from the high-resolution

binary CT image. These results indicate that subvolume size and structural representativeness affect the prediction accuracy of both geometric and transport properties.

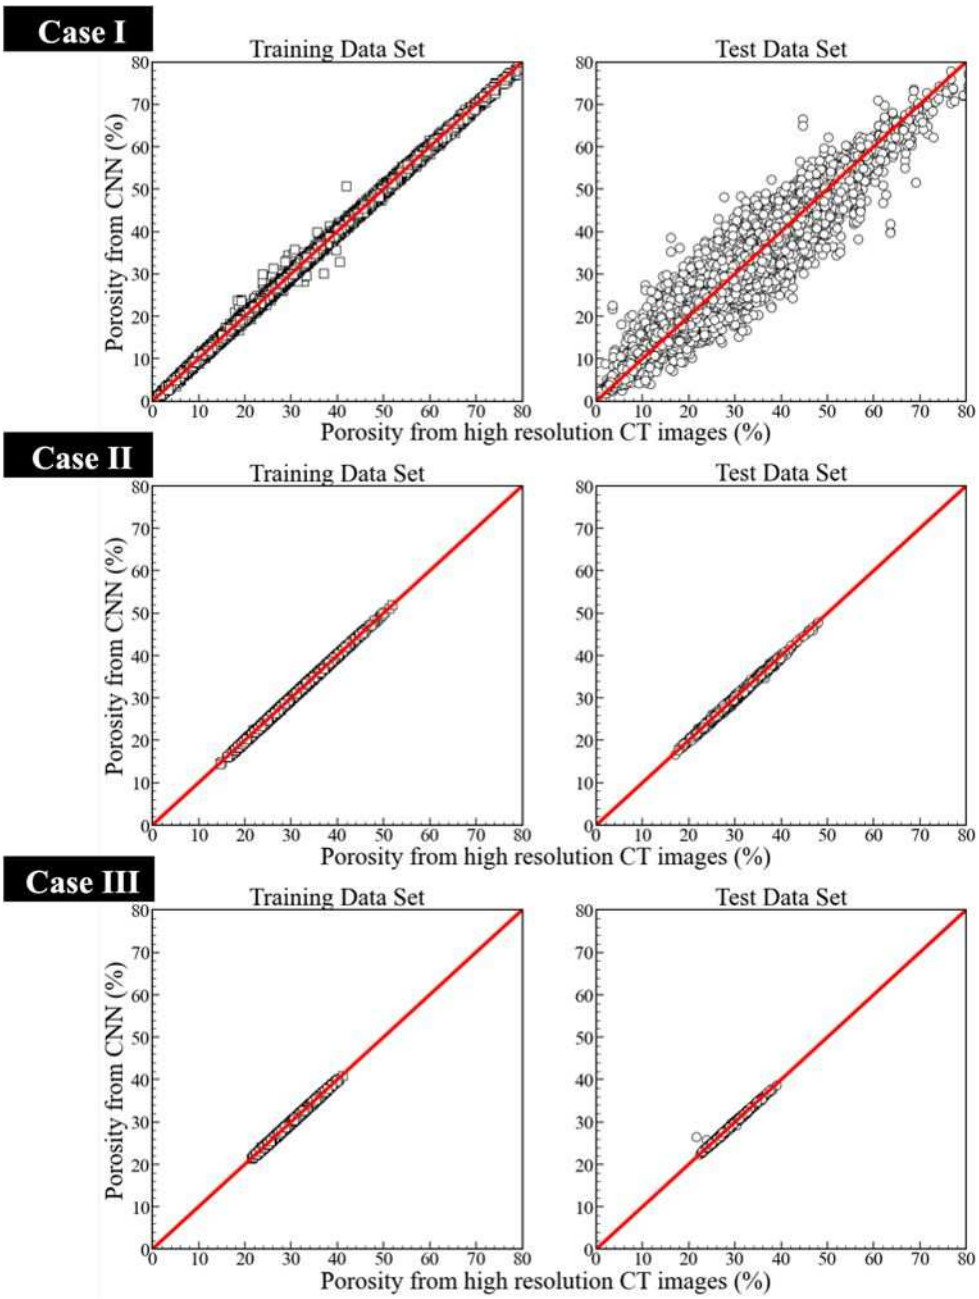


Figure 2.27: Porosity calculated from high-resolution binary CT images plotted against CNN-predicted porosity for Cases I, II, and III. The red diagonal line represents $y = x$.

2.5.4 Permeability–porosity relationship and permeability maps

The permeability–porosity relationships for Cases I, II, and III are shown in Figure 2.28. The solid line shows the Kozeny–Carman (KC) model fit. The KC model captures only a mean trend and does not reproduce the spread of permeability at a given porosity. In Case I, the KC fitting curve exhibits a downward shift because numerous outliers are present within the 100^3 subvolumes, leading to a bias in the fitted relationship. As the subvolume size increases to 200^3 and 300^3 voxels, both the LBM-calculated and CNN-predicted data become more tightly clustered around the KC-model line, and the predicted data closely mirror the calculated data. The CNN-based workflow captures both the overall trend and the distributional spread because it extracts three-dimensional image information related to pore microstructure and flow-path connectivity, rather than relying on porosity alone. This distinction is important for upscaling because the Darcy-scale flux depends on the spatial organization of high- and low-permeability regions, not only on their average porosity.

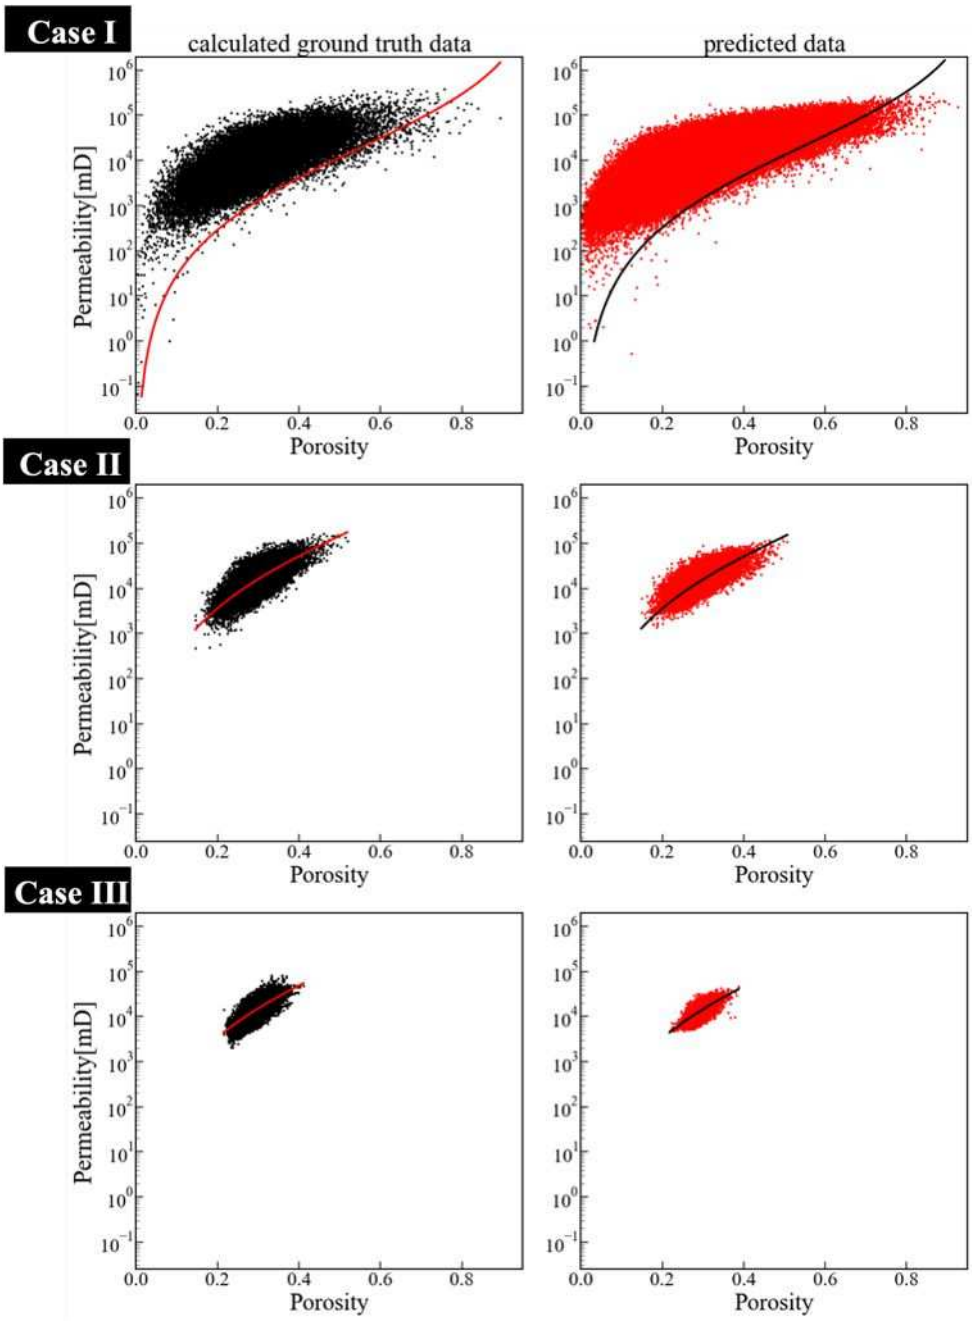


Figure 2.28: Relationship between permeability and porosity for Cases I, II, and III. The solid line shows the KC-model fit.

Figure 2.29 shows the predicted permeability maps of the larger core for the three subvolume sizes. The maps are consistent with the raw CT structure: less dense regions are assigned higher permeability and denser regions are assigned lower permeability. The maps also show how the choice of subvolume size affects the spatial resolution and apparent heterogeneity of the predicted permeability distribution.

Case I exhibits a highly granular permeability distribution and preserves strong small-scale variability, but this variability includes unstable local heterogeneity and prediction errors. As the subvolume size increases, noise-like local variability is suppressed and larger-scale flow structures become more visible. Case III produces the smoothest permeability distribution, with a more uniform and coherent macroscopic structure. Figure 2.30 compares the permeability map with the corresponding raw CT rock model, demonstrating the spatial alignment between the predicted permeability field and the observed rock structure.

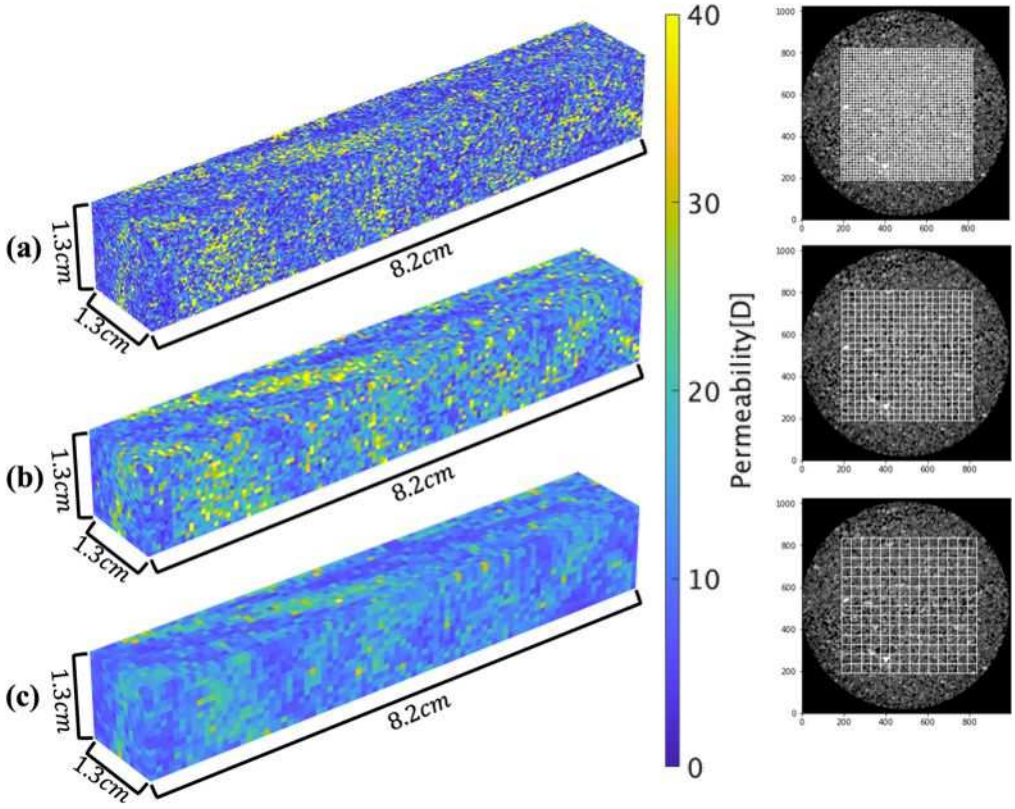


Figure 2.29: Permeability maps for the larger core and the rectangular regions extracted from the cylindrical sample for Darcy flow calculation: (a) Case I, (b) Case II, and (c) Case III.

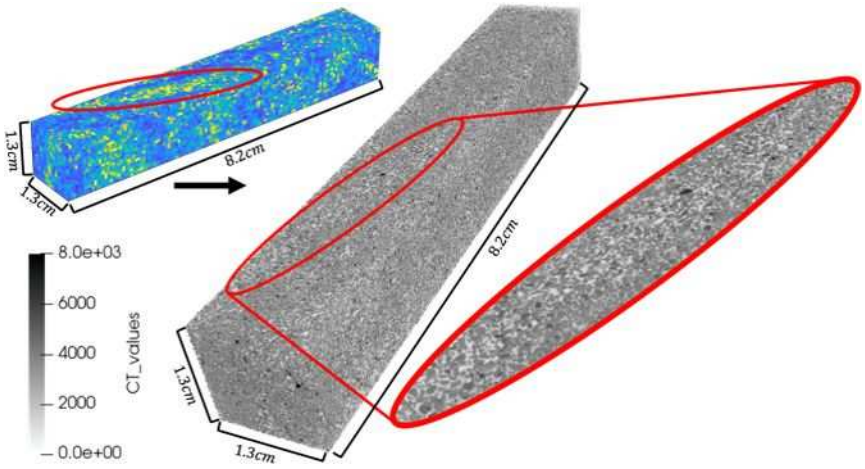


Figure 2.30: Comparison between the predicted permeability map and the corresponding raw CT rock model.

2.5.5 Darcy-scale pressure and velocity fields

The three permeability maps were used as input to the Darcy flow solver to compute pressure and flux fields. Figure 2.31 shows the resulting pressure distributions and the pressure differences between Cases I and III and between Cases II and III. The pressure distributions in all three cases show the general trend of decreasing pressure from the inlet to the outlet. The pressure-difference maps show that the pressures in Cases I and II are generally higher than those in Case III from the inlet to the midsection of the core, whereas the pressure near the outlet tends to be lower than that in Case III. Case I exhibits larger local pressure fluctuations because the small subvolume size preserves strong local permeability variations. Case II is more uniform, but still retains local variability relative to Case III. Case III produces the smoothest pressure distribution because small-scale heterogeneity is averaged over larger subvolumes. These results indicate that increasing the subvolume size improves the overall uniformity of the pressure distribution.

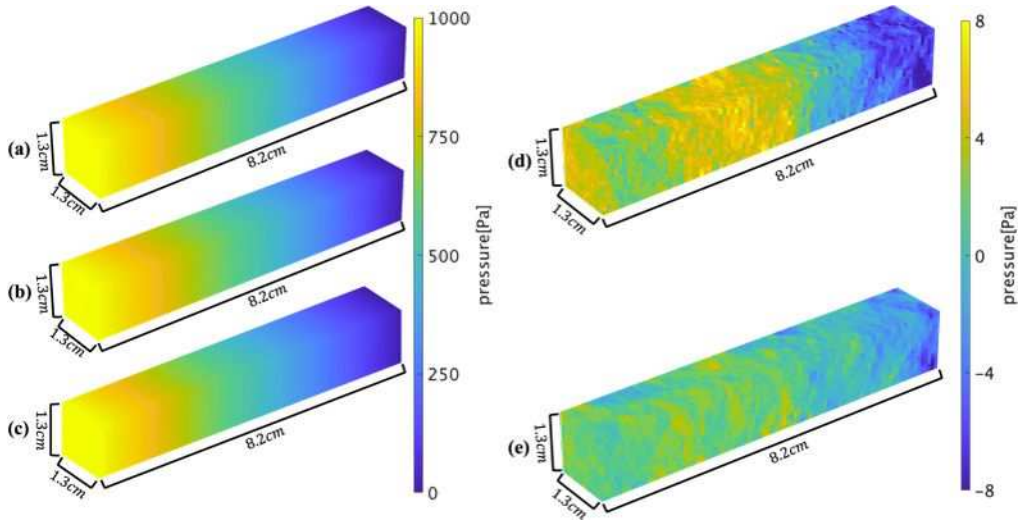


Figure 2.31: Pressure distributions of the larger core obtained from the Darcy flow solver: (a) Case I, (b) Case II, (c) Case III, (d) pressure difference between Case I and Case III, and (e) pressure difference between Case II and Case III.

For a perfectly homogeneous porous medium, the pressure decreases linearly from inlet to outlet. Figure 2.32 compares the simulated average pressure with the ideal linear pressure distribution and shows the permeability variation along the core flow direction. The pressure difference from the ideal linear profile correlates with the permeability distribution. The pressure difference exhibits a distinct pattern: it decreases in the inlet region, increases between approximately 0.5 cm and 4.4 cm, and then decreases again toward the outlet. Below 0.5 cm, the permeability is lower than the average, which increases flow resistance and causes the average pressure to drop more steeply than the ideal linear profile. Between 0.5 cm and 4.4 cm, the permeability exceeds the average value, reducing the required pressure gradient and producing a more gradual pressure decrease. At approximately 4.4 cm, permeability decreases sharply and remains below the average toward the outlet, again requiring a larger pressure gradient. This behavior indicates that the predicted permeability map controls the Darcy-scale pressure field in a physically consistent manner.

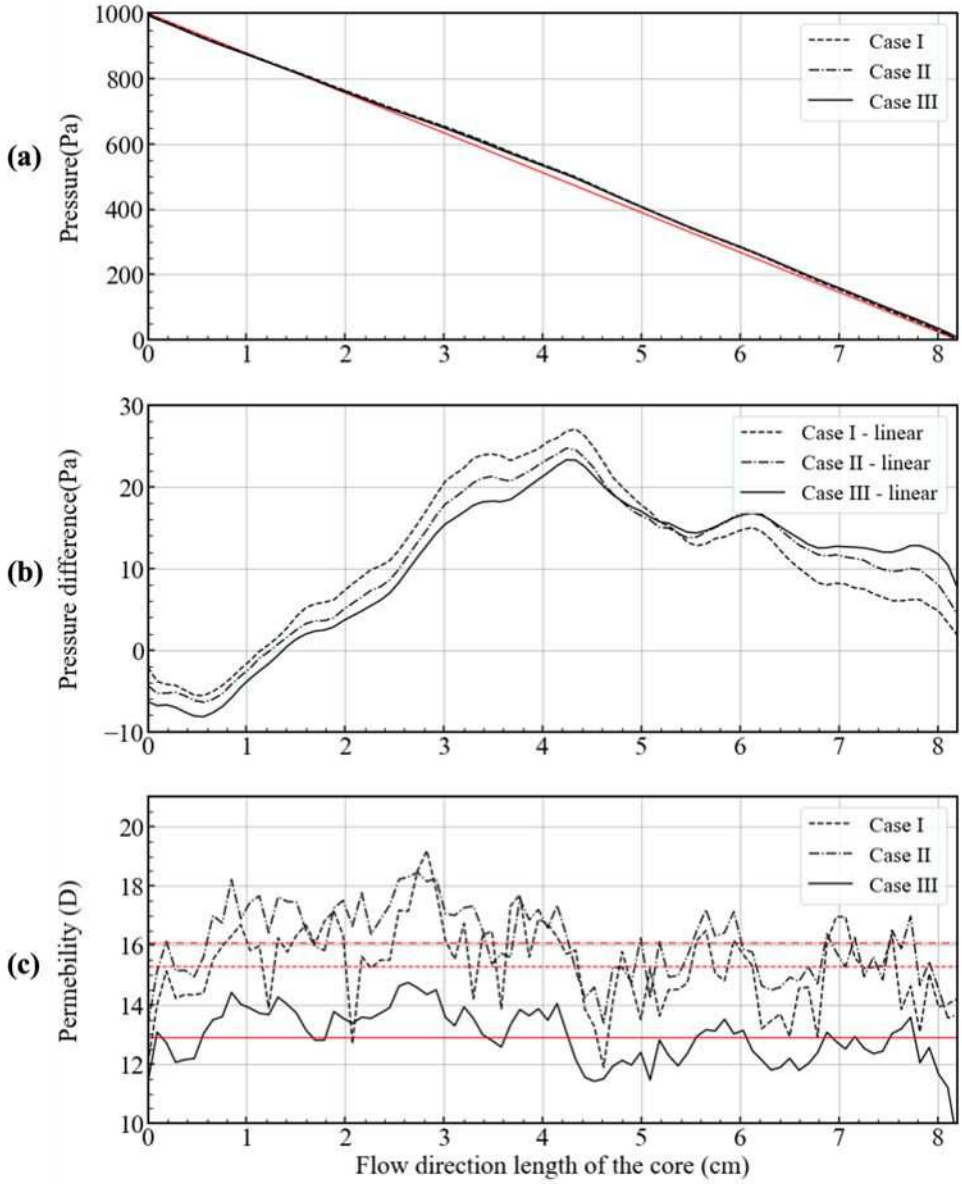


Figure 2.32: Variation of properties along the core flow direction: (a) pressure, where the red line represents the ideal linear pressure distribution; (b) pressure difference between average pressure and ideal linear pressure; and (c) permeability, where the red line represents the average permeability.

Figure 2.33 shows dimensionless velocity distributions normalized by the mean flow velocity. Case I contains strong localized velocity variations, including regions of negative velocity and local backflow, because many subvolumes are dominated by either flow paths or solid regions. Case II shows a moderately smoothed velocity distribution, in which localized variations are still apparent but less pronounced because of intermediate averaging effects. Case III gives the smoothest velocity

field, with more evenly distributed variations and reduced fine-scale heterogeneity because of the averaging effect of larger subvolumes. The velocity-field comparison demonstrates that subvolume size affects not only the CNN prediction accuracy, but also the physical response calculated by the Darcy flow solver.

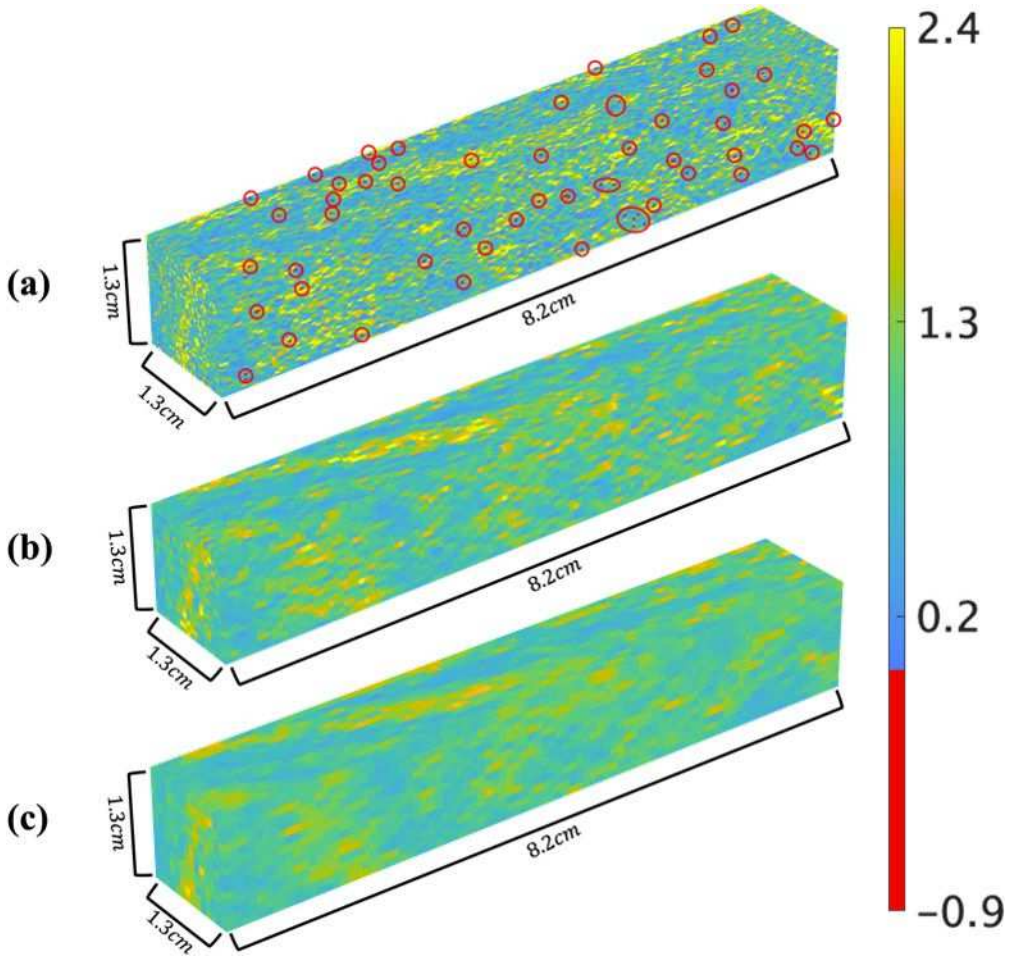


Figure 2.33: Dimensionless velocity distributions normalized by the mean flow velocity of the larger core obtained from the Darcy flow solver: (a) Case I, (b) Case II, and (c) Case III.

2.5.6 Upscaled permeability of the larger core

The upscaled permeability of the larger core was calculated from the Darcy-scale flux. For comparison, the arithmetic average of subvolume permeability was also calculated, and PNM was used to estimate the permeability of the high-resolution central region. PNM represents the pore space as a network of pores and throats and solves transport equations on simplified pore geometries Blunt (2001). In this study,

the pore network model was reconstructed using the SNOW algorithm Gostick et al. (2019), and fluid simulations were performed with OpenPNM Gostick et al. (2016). Figure 2.34 shows the reconstructed pore network.

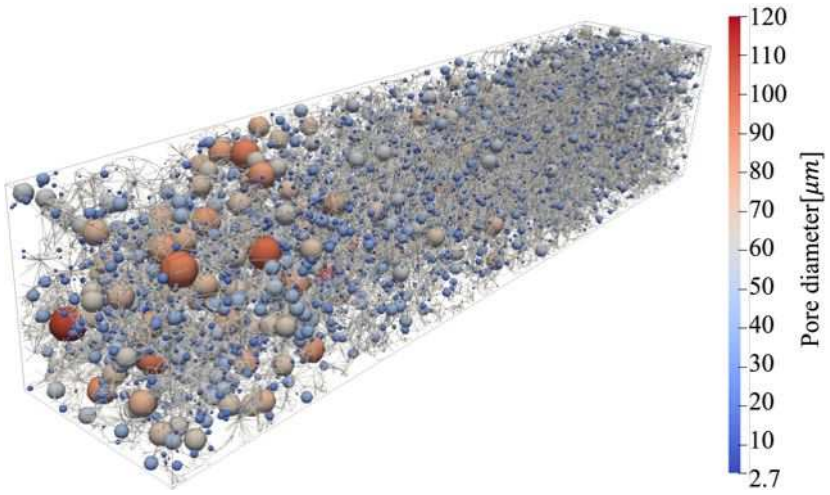


Figure 2.34: Three-dimensional visualization of the pore network model reconstructed from the high-resolution CT image. Spheres represent pore bodies and are colored by diameter. Lines represent throats connecting pore bodies.

Table 2.7 summarizes the large-scale permeability values obtained by different methods for the larger core. The coreflooding experiment measured a whole-core permeability of 8.9 D. The CNN–Darcy workflow gave upscaled permeability values of 10.9 D, 14.2 D, and 12.2 D for Cases I, II, and III, respectively. The arithmetic average permeability values were 15.3 D, 16.1 D, and 12.9 D. The PNM result for the high-resolution central region was 19.1 D, which is substantially higher than the coreflooding result. This overestimation occurs because PNM simplifies the pore space into idealized pore bodies and throats. Errors in the definition of pores and throats can strongly affect transport-property prediction Liu et al. (2022), especially in rocks with relatively small pore sizes such as Boise sandstone.

Case	I	II	III
k of whole core (CNN–Darcy upscaled)	10.9 D	14.2 D	12.2 D
Arithmetic average permeability	15.3 D	16.1 D	12.9 D
k of central region (PNM)		19.1 D	
Coreflooding experimental data		8.9 D	

Table 2.7: Large-scale permeability values obtained by different methods for the larger core. Here, 1 D $\approx 9.869 \times 10^{-12}$ m^2.

The difference between the Darcy-upscaled permeability and the arithmetic average decreases as subvolume size increases. This occurs because local heterogeneity is progressively averaged out in larger subvolumes. As the permeability field becomes more uniform, the flux distribution becomes closer to that of a homogeneous medium, and the effective permeability approaches the arithmetic average. However, the most homogeneous field is not necessarily the most physically informative; excessive averaging can remove flow-relevant heterogeneity. Therefore, the appropriate subvolume size must balance structural representativeness, prediction accuracy, computational cost, and preservation of heterogeneity.

Case I gives an upscaled permeability of 10.9 D, which appears closest to the coreflooding result. However, this apparent agreement should not be interpreted as evidence that Case I is the most reliable case. The CNN prediction accuracy is lowest in Case I, and the prediction plot in Figure 2.26 shows that many predicted values fall below the $y = x$ line, indicating systematic underestimation for unseen subvolumes. In addition, the flux map in Figure 2.33 shows unstable local velocity variations. The small subvolumes overemphasize local heterogeneity and can introduce noise-like artifacts into the Darcy simulation. Therefore, the selection of subvolume size should be based not only on closeness to the experimental value, but also on intrinsic indicators such as prediction accuracy, residual behavior, and flux-field stability. Case II improves prediction accuracy, but its permeability histogram contains many high-permeability subvolumes that can reconnect to form preferential flow paths, resulting in an upscaled permeability of 14.2 D. Case III yields 12.2 D, which is closer to the experimental value than Case II and is supported by the highest CNN accuracy. In the tested range, the 300^3 subvolume size, corresponding to an average of approximately 1240 pore bodies, provided the most stable and reliable result for this Boise sandstone core.

This result also indicates that pore-body count is a more transferable indicator of subvolume adequacy than voxel size alone, because voxel size depends on CT resolution and the characteristic grain size of the rock. Specifying the subvolume in terms of pore-body count anchors the criterion to the topology of the flow network. Therefore, the subvolume should be large enough to contain a statistically representative connected pore network, while remaining small enough to resolve core-scale heterogeneity.

In terms of computational efficiency, PNM still requires several hours for network reconstruction and flow calculation. The CNN-based workflow requires approximately 40 min to obtain the upscaled permeability for Case III after the model

is trained: 38 min for generating the permeability map and 2 min for the Darcy flow simulation. This difference indicates that the proposed workflow has advantages in computational time, spatial coverage, and prediction reliability once an appropriate trained network is available.

2.6 Summary

This chapter developed and verified a CNN-based workflow for upscaling absolute permeability from high-resolution pore-scale information to core-scale effective permeability inferred from low-resolution CT images. The workflow uses LBM to calculate permeability labels from segmented high-resolution CT subvolumes, trains CNNs using the corresponding low-resolution raw CT subvolumes, predicts whole-core permeability maps directly from low-resolution CT images, and calculates the upscaled permeability using a Darcy flow solver. The central purpose of this chapter was to establish whether DNS-derived pore-scale permeability information can be transferred to larger low-resolution CT domains while preserving the influence of local pore-structure heterogeneity.

In the workflow-development stage, two small Boise sandstone core samples were used to establish the basic method. High-resolution CT images were segmented and used to obtain LBM permeability labels, whereas low-resolution raw CT images were used directly as CNN inputs without segmentation. Five loss functions and six three-dimensional CNN architectures were tested. Huber loss gave the best performance in the loss-function test, and ResNet34 gave the best prediction accuracy among the tested architectures. The trained ResNet34 model predicted both porosity and permeability from low-resolution raw CT images. Porosity was predicted with small MAE values because it mainly depends on volumetric grayscale information, whereas permeability was more difficult to predict because it also depends on pore-space connectivity and flow-path tortuosity.

The predicted permeability maps from the trained CNN were used in a Darcy flow solver to calculate whole-core upscaled permeability. The CNN–Darcy workflow gave upscaled permeability values of 5.85 D and 1.53 D for Samples I and II, respectively. The meter-scale coreflooding permeability of the parent Boise sandstone was 2.8 D, which lies between the two local sample predictions. This agreement is reasonable because Samples I and II were extracted from regions with different local porosity and permeability. A transfer test was also performed by applying the network trained on Sample I to Sample II. The resulting upscaled permeability was

2.47 D, indicating that a trained network can be transferred to an unseen sample of the same rock type to some extent, although separate training is still required for different lithologies.

The workflow-development results demonstrate that low-resolution raw CT images contain useful grayscale information related to unresolved pore structure and partial-volume effects. The comparison with the KC model and PNM further shows that permeability cannot be represented reliably by porosity or simplified pore-network geometry alone. The KC model captures only an average permeability–porosity trend and cannot reproduce the scatter caused by pore connectivity and pore-throat geometry. PNM can provide useful estimates when high-resolution pore structures are available, but its accuracy depends strongly on how pore bodies and throats are defined. By contrast, the CNN-based workflow uses three-dimensional CT-image information related to pore microstructure and flow-path connectivity, allowing it to capture both the trend and the scatter of the permeability–porosity relationship.

In the workflow-verification stage, a larger 8.7 cm Boise sandstone core was used to examine the effects of training dataset size and subvolume size. Three high-resolution subvolume sizes were tested: 100^3, 200^3, and 300^3 voxels. These corresponded to average pore-body counts of approximately 74, 421, and 1240, respectively. As subvolume size increased, the LBM-derived permeability distributions became narrower, pore-network connectivity became more stable, and CNN prediction accuracy improved. Case I provided the largest number of training data points and the broadest permeability distribution, but the small subvolumes were not sufficiently representative of connected flow pathways and produced the lowest prediction accuracy. Cases II and III reduced the influence of outlier subvolumes by including more representative pore-network structures. Case III reached high prediction accuracy with relatively few training data points, although it required the longest training time because of the larger input size.

The Darcy-scale verification showed that subvolume size affects not only local CNN prediction accuracy, but also the final pressure field, velocity field, and upscaled permeability. Case I preserved strong local permeability variations and produced unstable local velocity fluctuations, including local backflow. Case III produced smoother pressure and velocity fields because small-scale heterogeneity was averaged over larger subvolumes. The CNN–Darcy upscaled permeability values were 10.9 D, 14.2 D, and 12.2 D for Cases I, II, and III, respectively, compared with the coreflooding value of 8.9 D. Although Case I appears closest to the experimental

value, it is not the most reliable case because it has the lowest CNN prediction accuracy and unstable Darcy-scale velocity fields. Case II overestimates permeability because many high-permeability subvolumes can reconnect to form preferential flow paths. Case III gives the most stable and reliable result among the tested cases because it combines the highest CNN prediction accuracy with a physically smoother Darcy-scale response.

For the Boise sandstone studied here, the 300^3 high-resolution subvolume size, corresponding to an average of approximately 1240 pore bodies, provided the best balance between structural representativeness, prediction accuracy, preservation of meaningful heterogeneity, and upscaling reliability. This result also shows that pore-body count is a more transferable indicator than voxel size alone, because it is tied to the topology of the connected pore network rather than to a specific CT resolution. Therefore, the optimal subvolume size should not be defined only by voxel dimensions, but also by the number of pore bodies and the representativeness of connected flow pathways contained in each subvolume.

The limitations of the workflow are also clear. The optimal subvolume size identified in this chapter was obtained for Boise sandstone under single-phase flow conditions, and its generality to other rock types has not yet been established. Rocks with very small, isolated, or poorly connected pores, such as low-porosity crystalline rocks, may be difficult to analyze using the same CT-resolution settings. In addition, the trained CNN is specific to the rock type, pore structure, porosity range, and image-acquisition conditions. For different lithologies, the training procedure must be repeated, or a broader database of trained networks must be constructed. Despite these limitations, the CNN–Darcy workflow provides a physically consistent route for integrating high-resolution pore-scale permeability information with centimeter-scale low-resolution CT heterogeneity. This workflow forms the basis for the multiphysics extension and stochastic field-scale modeling developed in the following chapters.

Chapter 3

Registration-Free Multiphysics Upscaling and Stochastic Rock Physics Modeling

3.1 Overview

Chapter 2 developed a CNN-based workflow that linked permeability labels derived from HR CT-based LBM simulations to LR CT images and used the predicted permeability map for whole-core upscaling Jiang et al. (2023); Guo et al. (2025). That workflow required HR and LR images of the same rock volume to be spatially registered. At the 20-fold voxel-size difference used in this chapter, reliable registration is difficult because pore-scale landmarks visible in the HR image are not resolved in the LR image. This chapter removes the registration requirement by using a Downsampling Generative Adversarial Network (DSgan) to generate LR-like inputs directly from HR CT subvolumes. The generated LR-like images are constrained toward the grayscale characteristics of the actual 20 μm/voxel LR CT domain while remaining derived from the same HR subvolumes as the DNS reference labels. Each generated LR-like input is therefore paired with reference physical properties calculated from the corresponding HR subvolume without physical HR–LR image registration.

SE-Atten-Res3D is trained on these paired data to predict permeability, porosity, electrical conductivity, bulk modulus, and shear modulus from LR-like CT subvolumes. The five target properties depend on different components of the pore–solid structure, including pore volume, connected transport paths, and the load-bearing solid framework. SE-Atten-Res3D combines 3D residual blocks, SE

channel-attention modules, and multi-stage feature aggregation. The SE modules compute channel weights from globally pooled features and rescale the intermediate feature tensors before regression. The trained ensemble is applied to the whole-core LR CT volume to generate five property maps, which are then used in Darcy-flow and FEM calculations of the whole-core effective properties.

The final part of the chapter formulates an S-RPM from the porosity–property scatter in the predicted core-scale maps. Deterministic porosity–property relationships describe the mean trend but do not represent the residual property variation observed among cells with similar porosity. This residual variation reflects microstructural differences that are not described by porosity alone, including pore connectivity, pore-throat constrictions, and the geometry of the load-bearing solid framework. The S-RPM represents this residual variation with a porosity-dependent amplitude and a spatially correlated stochastic field calibrated from the 640,000 predicted cell-wise porosity–property pairs. The stochastic term is calibrated from the residuals between the SE-Atten-Res3D predictions and the fitted deterministic porosity–property relationships. For each property, the fitted deterministic relationship is combined with a porosity-dependent residual amplitude and a standardized spatial random field. Given a porosity field and a prescribed spatial-correlation model, the S-RPM reconstructs three-dimensional property fields that retain the calibrated porosity–property spread. The reconstructed S-RPM maps are compared with the SE-Atten-Res3D maps in terms of their porosity–property distributions, spatial variations, and whole-core effective properties. The S-RPM therefore provides a porosity-based stochastic representation for larger domains in which continuous CT images are unavailable.

The remainder of this chapter describes Sample J7 and the multiscale CT data, the DNS reference calculations, DSgan and SE-Atten-Res3D, whole-core property mapping and upscaling, and the S-RPM formulation. The S-RPM reconstruction is evaluated by comparing its porosity–property distributions and whole-core effective properties with those obtained from the SE-Atten-Res3D maps.

3.2 Sample J7 and Multiscale CT Imaging

A sandstone core sample was obtained from the Frio Formation at a depth of approximately 1,887 m in Texas, USA, where CO_2-EOR operations are conducted. This sample is referred to as "Sample J7" in this dissertation. Compared with the Boise sandstone samples used in Chapter 2 Jiang et al. (2023); Guo et al. (2025),

Sample J7 shows greater spatial variation in sedimentary texture and pore structure. Sample J7 is a quartz-rich, clay-bearing sandstone composed of 70.6% quartz, 16.1% feldspar, 11.7% clay, and 1.6% other minor minerals. The observed heterogeneity includes variations in pore connectivity, local mineral distribution, and load-bearing grain contacts, which can produce spatial variations in permeability, electrical conductivity, and elastic moduli. Sample J7 was selected to evaluate the proposed workflow on a reservoir-core sandstone with clear microstructural heterogeneity.

CT imaging was performed on the whole core and an extracted sub-core at two voxel sizes. The whole core was first scanned at 20 μm/voxel to cover the centimeter-scale core domain. Before the LR scan, numbered circular reference marks were placed on the core surface to identify the approximate region selected for HR imaging. The resulting LR image provides the input domain for whole-core property mapping, but it does not resolve the pore-scale geometry required for DNS. After the LR scan, a millimeter-scale sub-core was extracted from the marked region and scanned at 1 μm/voxel (actual voxel size: 0.999 μm/voxel) to resolve pore-scale geometry within a smaller field of view. The HR image resolves narrow pore throats, connected pores, grain contacts, and solid grains that define the computational domains for the DNS-based flow, electrical, and elastic-property calculations. At 20 μm/voxel, partial-volume effects blur pore–solid boundaries and narrow pore throats, so the LR image cannot be used directly as the computational domain for pore-scale DNS.

The 20-fold difference in voxel size and the difference in field of view make precise registration of the HR and LR images difficult. Sub-core extraction can also damage the original core, and fragile reservoir sandstones may fracture during cutting. The registration-free procedure described in the following section therefore generates LR-like inputs computationally from the HR subvolumes rather than relying on physically registered HR–LR image pairs.

3.3 Method

3.3.1 Overview of the Proposed Upscaling Framework

The objective of the proposed workflow is to estimate multiple physical properties from LR CT images and to calculate whole-core effective properties. The workflow consists of three main parts. First, a Downsampling Generative Adversarial Network (DSgan) is used to generate LR-like images from HR CT images, allowing paired

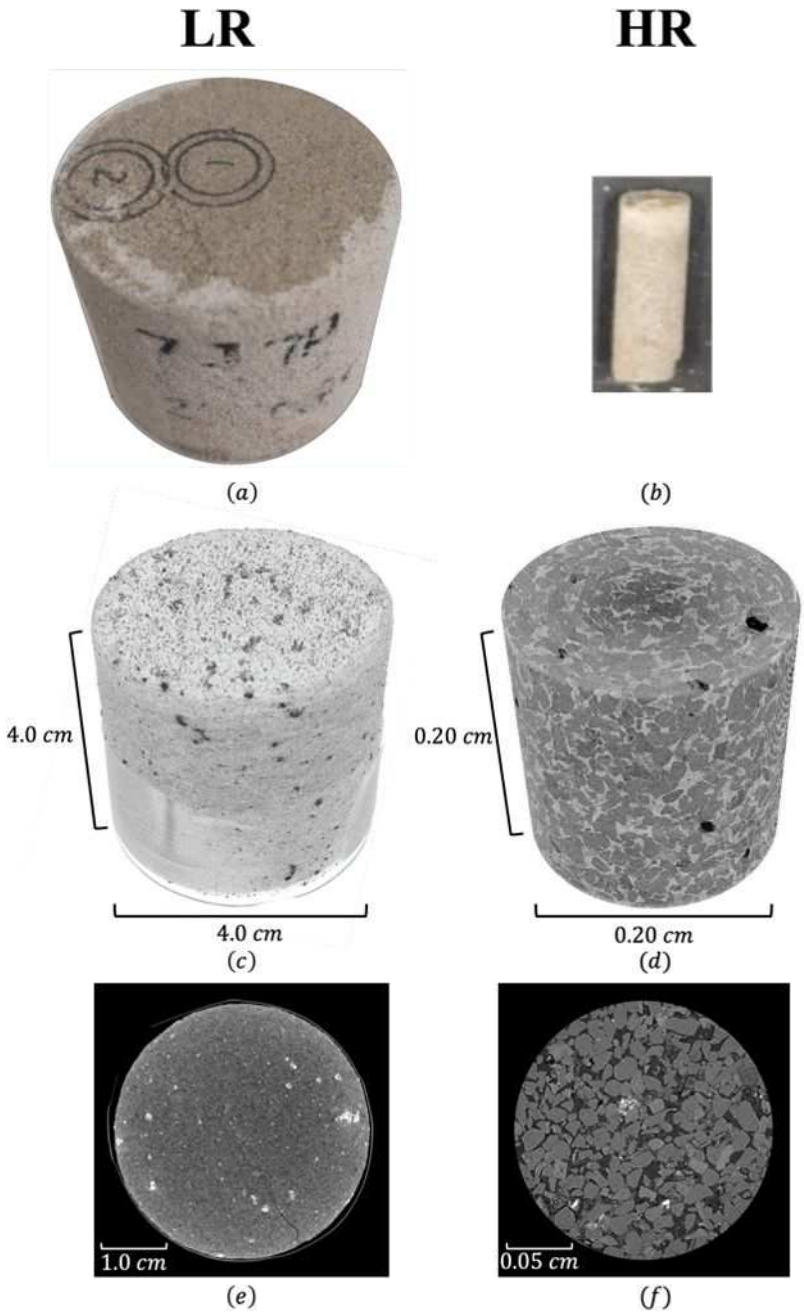


Figure 3.1: Multiscale CT images of Sample J7. (a, c) Three-dimensional volume rendering and a representative two-dimensional cross-sectional slice of the whole-core LR image acquired at 20 μm/voxel. (b, d) Three-dimensional volume rendering and a representative two-dimensional cross-sectional slice of the extracted sub-core HR image acquired at 1 μm/voxel. The numbered marks on the core surface indicate the approximate region selected for sub-core extraction. Scale bars are 1.0 cm in (c) and 0.05 cm in (d).

datasets to be constructed without physical HR–LR image registration. Second, SE-Atten-Res3D is trained as a multi-output regression model to predict five physical properties: permeability k, porosity ϕ, electrical conductivity σ^e, bulk modulus K, and shear modulus G. Third, the trained model is applied to the whole-core LR CT image to generate three-dimensional property maps, which are then used to calculate whole-core effective properties. The upscaled permeability is calculated using a Darcy-flow solver in the MATLAB Reservoir Simulation Toolbox (MRST), whereas the upscaled electrical conductivity and elastic moduli are calculated using the finite-element method (FEM).

The detailed workflow is shown in Figure 3.2 and consists of the following steps:

1. The whole sandstone core, approximately 4.0 cm in length and 2.0 cm in radius, is scanned at 20 μm/voxel to cover the whole-core volume. Raw grayscale LR subvolumes are extracted from the LR CT image.

2. A sub-core sample, approximately 0.20 cm in length and 0.10 cm in radius, is extracted from the whole core and scanned at 1 μm/voxel. Raw grayscale HR subvolumes are extracted from the HR CT image.

3. The HR CT image is processed by noise reduction and segmentation to construct a binary pore–solid model. Binary HR subvolumes are extracted from the segmented image.

4. Physical-property reference labels are obtained from the binary HR subvolumes. Permeability k is calculated by the lattice Boltzmann method (LBM), electrical conductivity σ^e, bulk modulus K, and shear modulus G are calculated by FEM, and porosity ϕ is obtained directly from the segmented pore volume.

5. Subsets of raw grayscale LR and HR subvolumes are used as unpaired datasets to train DSgan.

6. The trained DSgan generator converts raw grayscale HR subvolumes into LR-like grayscale subvolumes.

7. The paired dataset for SE-Atten-Res3D is constructed by linking each generated LR-like subvolume with the physical-property labels calculated from the corresponding binary HR subvolume.

8. SE-Atten-Res3D is trained using the paired dataset. The model takes an LR-like CT subvolume as input and predicts k, ϕ, σ^e, K, and G simultaneously.

9. The trained SE-Atten-Res3D ensemble is applied to the raw grayscale LR subvolumes extracted from the whole-core LR CT image to predict three-dimensional maps of the five physical properties.

10. The predicted property maps are used for whole-core upscaling. The permeability map is used in a Darcy-flow calculation in MRST to estimate k_{upscaled}, whereas the electrical-conductivity and elastic-property maps are used in FEM calculations to estimate $\sigma^e_{\text{upscaled}}$, K_{upscaled}, and G_{upscaled}.

3.3.2 Training Dataset Preparation

The training data were prepared in two stages. First, unpaired grayscale datasets were prepared to train DSgan. Second, paired LR-like inputs and HR-derived physical-property labels were constructed to train SE-Atten-Res3D.

For DSgan training, raw grayscale subvolumes were extracted from both the whole-core LR CT image and the HR CT image. The whole-core LR CT image was divided into an $80 \times 80 \times 100$ array of non-overlapping subvolumes, yielding 640,000 raw LR subvolumes of 15^3 voxels each (Figure 3.3(a)). Because the LR voxel size was 20 μm, each LR subvolume represented a cubic physical domain of (300 μm)^3.

The raw HR CT image was sampled using 300^3-voxel subvolumes (Figure 3.3(b)). Because the HR voxel size was 1 μm, each HR subvolume also represented a physical domain of (300 μm)^3. Adjacent 300^3-voxel HR windows overlapped by 50 voxels. This extraction yielded 12,000 raw HR subvolumes. From the raw LR and raw HR subvolume sets, 1,000 subvolumes were randomly selected from each domain and used as the unpaired training datasets for DSgan.

The binary HR subvolumes used for physical-property calculation were prepared from the segmented HR CT image. The raw HR image was first processed using a non-local means filter to reduce image noise while preserving phase boundaries. The filtered HR image was then segmented using a grayscale histogram-based threshold to separate the pore and solid phases. A binary pore–solid model with dimensions of $1400 \times 1350 \times 1900$ voxels was cropped from the cylindrical region of the segmented HR image. This binary model was divided into 300^3-voxel subvolumes using the same spatial positions and the same 50-voxel overlap as the corresponding raw

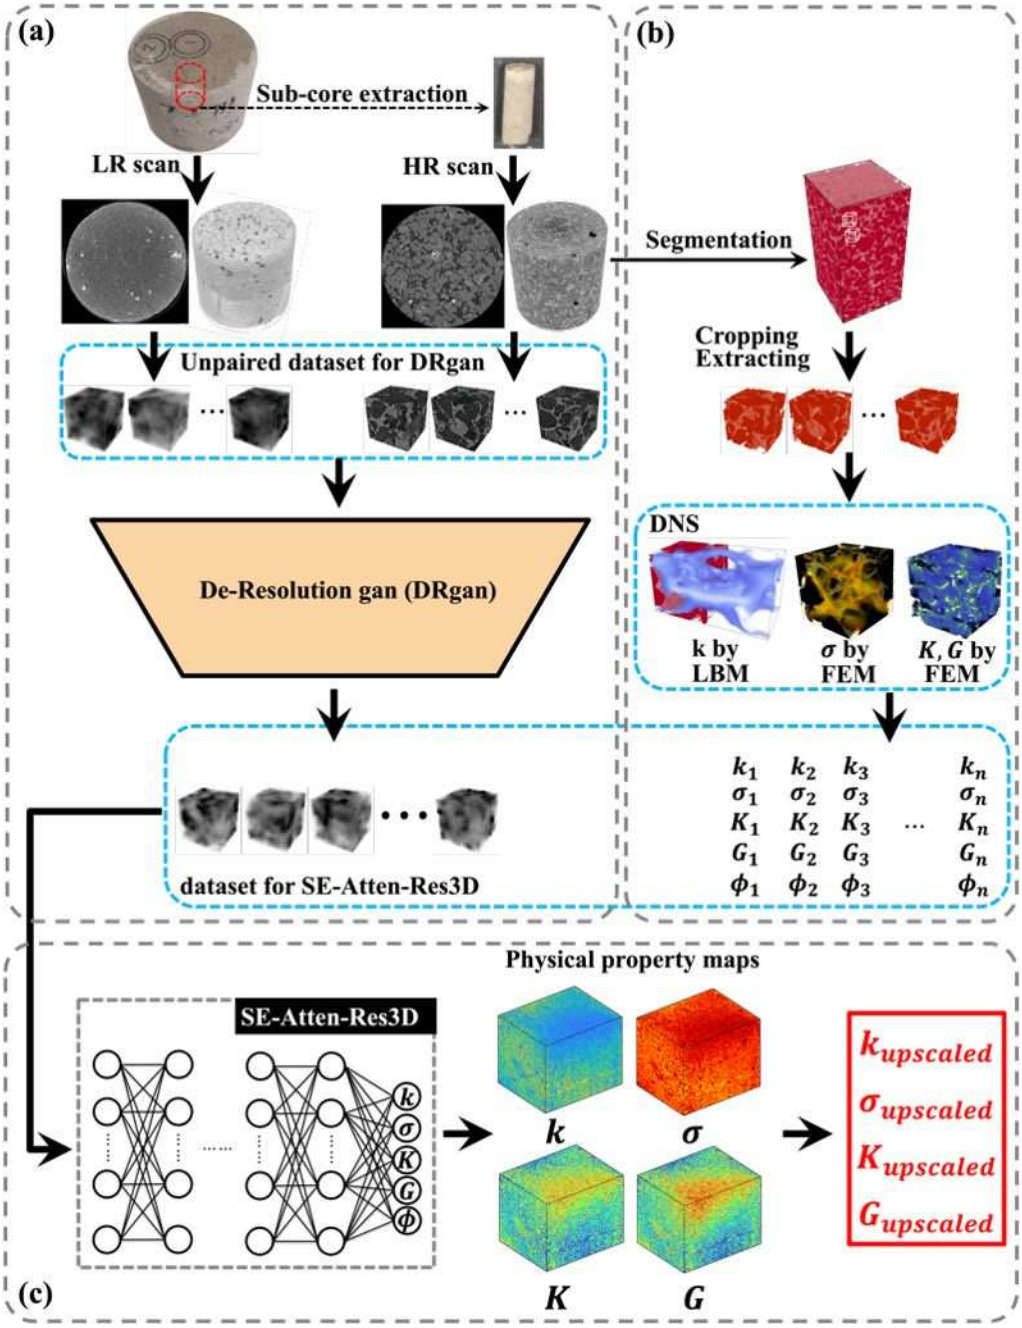


Figure 3.2: Workflow of the proposed multiscale upscaling method. (a) Whole-core LR CT scanning, sub-core extraction and HR CT scanning, and preparation of unpaired RLR and RHR subvolumes for DSgan training. (b) Segmentation of the HR CT image, extraction of binary HR subvolumes, and calculation of the reference labels: k by LBM; σ^e, K, and G by FEM; and ϕ from the segmented pore volume. (c) Conversion of RHR subvolumes into CLR subvolumes by DSgan, construction of paired CLR–label data, training and whole-core application of SE-Atten-Res3D, and calculation of k_{upscaled}, $\sigma^e_{\text{upscaled}}$, K_{upscaled}, and G_{upscaled}. Subscripts $1, \ldots, n$ index the paired HR subvolumes and their corresponding labels.

HR subvolumes. The binary HR subvolumes were used for the physical-property calculations described in Section 3.3.3.

After DSgan training, the trained generator was applied to all 12,000 raw grayscale HR subvolumes. Each 300^3-voxel raw HR subvolume was converted into a 15^3-voxel LR-like grayscale subvolume representing the same $(300\ \mu\text{m})^3$ physical domain. Because the generated LR-like input and the corresponding binary HR subvolume originate from the same spatial position in the HR image, the physical-property labels calculated from the binary HR subvolume can be paired directly with the generated LR-like input without manual HR–LR image registration.

The final SE-Atten-Res3D dataset consisted of 12,000 pairs. Each pair contained one 15^3-voxel LR-like input and five HR-derived physical-property labels: k, ϕ, σ^e, K, and G. For model training and evaluation, 90% of the pairs were used for training and the remaining 10% were used for testing.

3.3.3 Numerical Simulations for Physical Properties

For each 300^3-voxel HR subvolume, k, σ^e, K, and G were calculated by DNS, whereas ϕ was obtained directly from the segmented pore volume. The binarized CT images were used directly as voxel-based Cartesian computational grids. Each cubic voxel was treated as one computational element and assigned to either the pore phase (0) or the solid phase (1). The calculation methods for permeability, electrical conductivity, and elastic moduli are described below.

1. **Permeability (k):** The LBM was used to calculate single-phase fluid flow within the pore space of each HR subvolume. A constant pressure gradient was applied across the domain, no-slip boundary conditions were imposed on the solid surfaces, and the simulation was continued until steady state was reached. The detailed LBM settings and numerical parameters are given in Chapter 2 and our previous studies Jiang et al. (2023); Guo et al. (2025). The absolute permeability k was calculated using Darcy's law:

$$k = \frac{\mu \langle u \rangle}{\nabla P}, \tag{3.1}$$

   where μ is the dynamic viscosity of the fluid, $\langle u \rangle$ is the volume-averaged velocity, and ∇P is the applied pressure gradient.

2. **Electrical conductivity (σ^e):** The effective electrical conductivity was determined using a voxel-based FEM Garboczi (1998) based on the variational

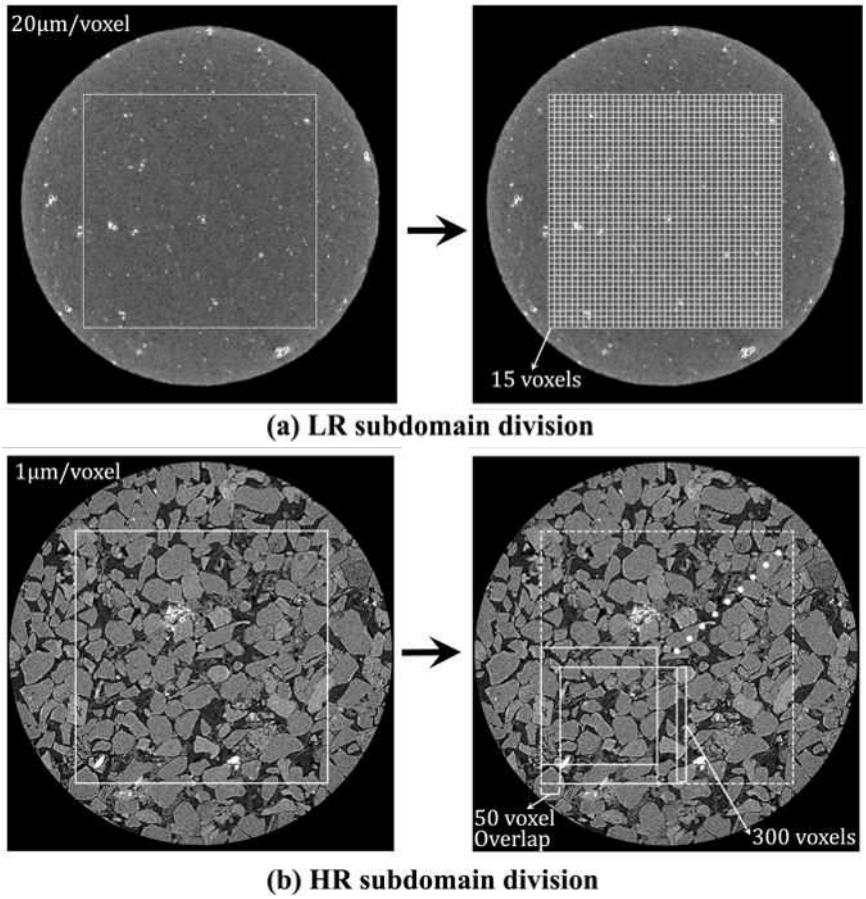


Figure 3.3: Subvolume extraction from the LR and HR CT images. (a) A cuboidal region in the whole-core LR CT image acquired at 20 μm/voxel was divided into non-overlapping 15^3-voxel subvolumes. (b) A region in the HR CT image acquired at 1 μm/voxel was sampled using 300^3-voxel windows with a 50-voxel overlap between adjacent windows. Both subvolume sizes correspond to a physical domain of (300 μm)^3. The grids and boxes are illustrated on representative cross-sectional slices.

principle of minimum energy dissipation. For a single cubic voxel, the local electrical energy dissipation W_{elec} is written as

$$W_{\text{elec}} = \int_0^1 \int_0^1 \int_0^1 \left[\frac{1}{2} e_p \sigma_{pq} e_q\right] dx\, dy\, dz, \tag{3.2}$$

where x, y, and z are the local coordinates within the voxel, e_p is the p-th component of the local electric field, and σ_{pq} is the local conductivity tensor assigned to the voxel.

The water-saturated pore phase was assigned an isotropic conductivity of $\sigma_{\text{water}} = 0.3$ S/m Hilhorst (2000), whereas the solid phase was assigned $\sigma_{\text{solid}} = 1.0\times10^{-10}$ S/m Telford et al. (1990). In both phases, the electrical-conductivity tensor was assumed to be isotropic, with identical diagonal components and zero off-diagonal components.

A uniform macroscopic electric field E was applied to the subvolume under periodic boundary conditions (PBC). The nodal voltages were solved using the conjugate-gradient (CG) algorithm to minimize the total electrical energy until convergence. After convergence, the effective electrical conductivity σ^e was calculated from macroscopic Ohm's law:

$$\sigma^e = \frac{\langle J \rangle}{E}, \tag{3.3}$$

where $\langle J \rangle$ is the volume-averaged current density.

3. **Elastic properties (bulk modulus K and shear modulus G):** The effective elastic moduli were calculated using a voxel-based three-dimensional linear-elastic FEM solver Garboczi (1998) based on minimization of the elastic strain energy. Mechanical equilibrium was obtained by minimizing the total elastic strain energy W_{elas}. Using Voigt notation ($\alpha, \beta = 1, \ldots, 6$) Nye (1985), the elastic strain energy in a single voxel is written as

$$W_{\text{elas}} = \int_0^1 \int_0^1 \int_0^1 \left[\frac{1}{2} \epsilon_\alpha C_{\alpha\beta} \epsilon_\beta\right] dx\, dy\, dz, \tag{3.4}$$

where ϵ_α is the α-th component of the local strain and $C_{\alpha\beta}$ is the local stiffness tensor. For an isotropic phase, $C_{\alpha\beta}$ was specified from the bulk modulus K and shear modulus G Norris (2006).

The pore-water phase was assigned a bulk modulus of $K_{\text{w}} = 2.18$ GPa and

a shear modulus of $G_{\mathrm{w}} = 0$ GPa Mavko et al. (2020). The elastic moduli of the solid phase were estimated from the mineralogical composition of Sample J7 using the Voigt–Reuss–Hill averaging scheme Hill (1963). The resulting solid-phase moduli were $K_{\mathrm{s}} = 40.1$ GPa and $G_{\mathrm{s}} = 33.2$ GPa.

A prescribed macroscopic strain tensor was applied to each subvolume under PBC. The normal strain components were set to $\epsilon_{xx} = \epsilon_{yy} = \epsilon_{zz} = 0.01$, and the off-diagonal tensor strain components were set to $\epsilon_{xy} = \epsilon_{xz} = \epsilon_{yz} = 0.005$. For this prescribed strain state, the finite-element problem was solved by minimizing the elastic energy using the CG method.

The effective bulk modulus K and shear modulus G were calculated from the volume-averaged stress and the prescribed macroscopic strain using Hooke's law:

$$K = \frac{\langle \tau_{xx} \rangle + \langle \tau_{yy} \rangle + \langle \tau_{zz} \rangle}{3\left(\epsilon_{xx} + \epsilon_{yy} + \epsilon_{zz}\right)}, \tag{3.5}$$

$$G = \frac{1}{3}\left(\frac{\langle \tau_{xz} \rangle}{2\epsilon_{xz}} + \frac{\langle \tau_{yz} \rangle}{2\epsilon_{yz}} + \frac{\langle \tau_{xy} \rangle}{2\epsilon_{xy}}\right), \tag{3.6}$$

where $\langle \tau_{ij} \rangle$ is the volume-averaged stress component.

These calculations provided the reference labels k, ϕ, σ^e, K, and G for all HR subvolumes. The labels were paired with LR-like grayscale inputs generated by DSgan as described in the following subsection.

3.3.4 DSgan-Based Generation of LR-Like CT Subvolumes

The objective of DSgan is to convert raw HR CT subvolumes into LR-like CT subvolumes. The model consists of a generator G, a fixed trilinear interpolator I, and a discriminator D (Figure 3.4). A raw HR CT subvolume is denoted by $RHR \in \mathbb{R}^{1\times300\times300\times300}$, and a raw LR CT subvolume is denoted by $RLR \in \mathbb{R}^{1\times15\times15\times15}$. The generator converts RHR into a converted LR-like subvolume, $CLR \in \mathbb{R}^{1\times15\times15\times15}$:

$$CLR = G(RHR). \tag{3.7}$$

The adversarial objective trains the generator so that the score map $D(CLR)$ approaches the LR-domain target map of ones:

$$\mathcal{L}_{G,\mathrm{adv}} = \mathbb{E}\left[\left(D(G(RHR)) - \mathbf{1}\right)^2\right]. \tag{3.8}$$

For an input subvolume X, $D(X)$ denotes a one-channel score map of size $8 \times 8 \times 8$. The symbol $\mathbf{1}$ denotes a target map with the same dimensions as $D(X)$. The expectation $\mathbb{E}$ denotes averaging over the training samples and all locations in the score map.

The discriminator distinguishes raw LR subvolumes RLR from converted LR-like subvolumes CLR. During discriminator training, the target score is 1 for RLR and 0 for CLR. Using the least-squares GAN objective Mao et al. (2017), the discriminator loss is

$$\mathcal{L}_D = \frac{1}{2}\mathbb{E}\left[(D(RLR) - \mathbf{1})^2 + D(G(RHR))^2\right]. \tag{3.9}$$

Because the HR and LR training datasets are unpaired, the adversarial loss constrains the distribution of CLR toward the raw LR domain but does not require each CLR subvolume to preserve the structure of its input RHR subvolume. A cycle-consistency loss was therefore included to constrain the correspondence between RHR and CLR Zhu et al. (2017). The fixed interpolator I resamples CLR onto the 300^3 HR grid and produces a cycle-reconstructed HR subvolume CHR:

$$CHR = I(CLR) = I(G(RHR)). \tag{3.10}$$

The cycle-consistency loss is defined as the L_1 difference between CHR and the original RHR input:

$$\mathcal{L}_{\text{cyc}} = \mathbb{E}\left[\|CHR - RHR\|_1\right] = \mathbb{E}\left[\|I(G(RHR)) - RHR\|_1\right]. \tag{3.11}$$

The discriminator and generator were updated using separate objectives. The discriminator was updated by minimizing Eq. 3.9, whereas the generator was updated by minimizing

$$\mathcal{L}_{G,\text{total}} = \mathcal{L}_{G,\text{adv}} + \lambda\mathcal{L}_{\text{cyc}}. \tag{3.12}$$

Figure 3.5(a) shows the generator architecture. The generator uses reflection padding Zhu et al. (2017), three-dimensional convolutional layers Tran et al. (2015), instance normalization Ulyanov et al. (2016), ReLU activations Nair and Hinton (2010), and residual blocks He et al. (2016). Reflection padding was applied before the first and final convolutional layers. The first convolutional layer converts the single-channel RHR input into a multi-channel feature tensor. Three strided three-dimensional convolutional layers then reduce the spatial dimensions from $300 \times$

300×300 to $25 \times 25 \times 25$. Twelve residual blocks are applied at this stage. After the residual blocks, three convolutional layers transform the feature tensor from $25 \times 25 \times 25$ to $15 \times 15 \times 15$. Reflection padding and the final convolutional layer then produce the single-channel CLR output.

Figure 3.5(c) shows the fixed trilinear interpolator. The interpolator resamples the 15^3 CLR grid onto the 300^3 HR grid using coordinate mapping, neighborhood identification, and trilinear weighting. In coordinate mapping, each HR-grid voxel location is mapped to a continuous position in the CLR grid Jaderberg et al. (2015); Paszke et al. (2019). The eight CLR voxels surrounding the mapped position are then identified. The corresponding CHR voxel value is calculated as the distance-weighted average of these eight CLR voxel values Lehmann et al. (2002). This procedure produces the 300^3 CHR volume used in the voxel-wise cycle-consistency calculation.

Figure 3.5(b) shows the discriminator architecture. The discriminator processes each 15^3 input using three-dimensional convolutional layers Tran et al. (2015), instance normalization Ulyanov et al. (2016), and LeakyReLU activations Maas et al. (2013). The first convolutional layer uses a stride of 2 and reduces the spatial dimensions from 15^3 to 8^3. Instance normalization is not applied after this first convolutional layer. The subsequent convolutional layers use a stride of 1 and retain the 8^3 spatial dimensions while changing the number of feature channels. The final convolution produces a one-channel 8^3 score map without a sigmoid activation.

To reduce training oscillation, the discriminator was updated using a buffer of previously generated CLR volumes rather than only the most recent generator outputs Shrivastava et al. (2017). The buffer stored 50 previously generated CLR volumes. The cycle-consistency weight was set to $\lambda = 6$. The generator and discriminator were optimized using Adam Kinga et al. (2015) with an initial learning rate of 0.0002 Zhu et al. (2017). The learning rate was kept constant for the first 200 epochs and was linearly reduced to zero over the following 100 epochs. The unpaired DSgan training dataset consisted of 1,000 raw LR subvolumes and 1,000 raw HR subvolumes. After training, the generator converted each RHR subvolume into a CLR subvolume with grayscale characteristics of the raw LR CT domain. Each CLR subvolume was then paired with the five physical-property labels calculated from the binary HR subvolume at the corresponding location.

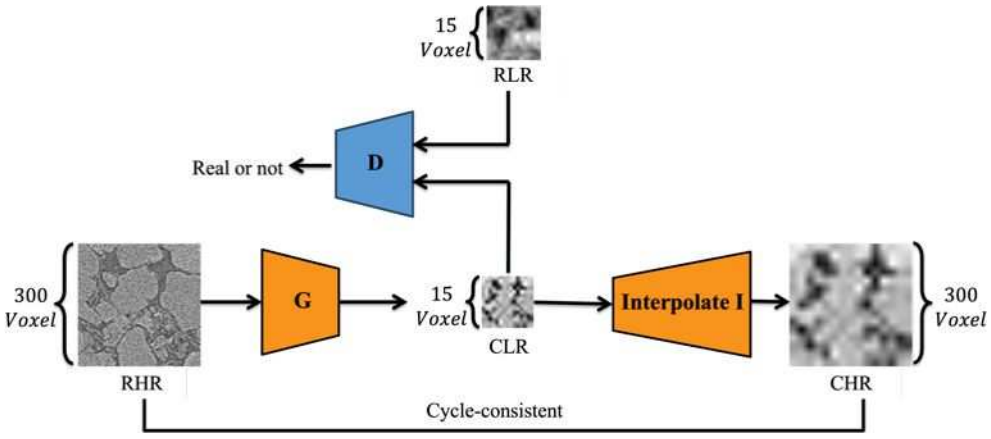


Figure 3.4: Schematic of DSgan training. The generator G converts an RHR subvolume of 300^3 voxels at 1 μm/voxel into a CLR subvolume of 15^3 voxels. The discriminator D evaluates CLR subvolumes together with unpaired RLR subvolumes acquired at 20 μm/voxel. The fixed trilinear interpolator I resamples CLR onto the 300^3 HR grid to obtain CHR for the cycle-consistency loss. The displayed images are representative two-dimensional slices of the three-dimensional subvolumes.

3.3.5 Multi-Property Prediction Using SE-Atten-Res3D

Predicting physical properties from DSgan-generated LR-like subvolumes requires a model that can extract image features related to pore and solid geometry from LR CT images. A shared three-dimensional CNN does not explicitly recalibrate feature channels before predicting properties governed by different components of the pore–solid structure. In multi-task learning, interference among task-related features is discussed as feature interference or negative transfer Hou and Cao (2022); Wang et al. (2019); Tang et al. (2026). SE-Atten-Res3D therefore uses channel recalibration before the final regression layers. The model combines three-dimensional residual blocks He et al. (2016), Squeeze-and-Excitation (SE) channel-attention modules Hu et al. (2018), and multi-stage feature aggregation Hariharan et al. (2015) (Figure 3.6).

Before being input into the network, each CT subvolume was normalized independently. Let $\mathbf{V}_{\text{in}}$ denote the input tensor of raw CT values. The normalized input tensor $\mathbf{V}_{\text{norm}}$ was calculated as

$$\mathbf{V}_{\text{norm}} = \frac{\mathbf{V}_{\text{in}} - \mu_{\text{in}}}{\sigma_{\text{in}}}, \tag{3.13}$$

where μ_{in} and σ_{in} are the mean and standard deviation of the voxel values within the 15^3 input subvolume.

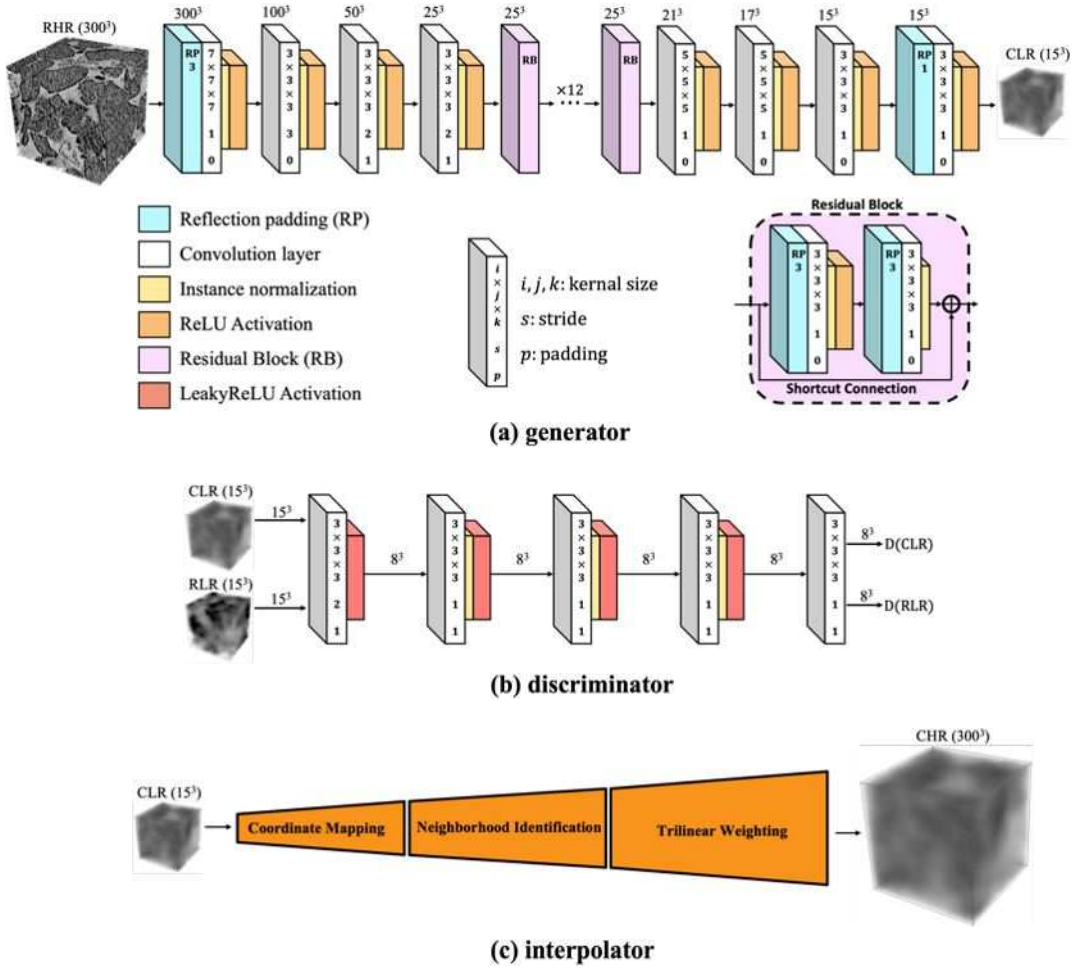


Figure 3.5: Detailed components of DSgan. (a) Generator G: reflection padding, three-dimensional convolution, instance normalization, ReLU activation, and 12 residual blocks transform RHR from 300^3 voxels to CLR at 15^3 voxels. (b) Discriminator D: three-dimensional convolutional layers process RLR or CLR subvolumes of 15^3 voxels and produce a one-channel 8^3 score map. (c) Fixed trilinear interpolator I: coordinate mapping, neighborhood identification, and trilinear weighting resample CLR from the 15^3 grid onto the 300^3 HR grid to obtain CHR. Numbers above the layer blocks indicate the spatial dimensions; i, j, and k denote kernel dimensions, s denotes stride, and p denotes padding.

An SE channel-attention block is embedded in each residual unit. For an intermediate feature tensor $\mathbf{F} \in \mathbb{R}^{C\times D\times H\times W}$, where C is the number of channels and D, H, and W are the spatial dimensions, the SE block first computes a channel descriptor by global average pooling:

$$\mathbf{s} = \mathrm{GAP}(\mathbf{F}) \in \mathbb{R}^{C}. \tag{3.14}$$

The channel weights are then calculated from $\mathbf{s}$ using two fully connected layers:

$$\mathbf{w} = \mathrm{Sigmoid}\left[\mathbf{W}_2\,\mathrm{ReLU}\left(\mathbf{W}_1\mathbf{s}\right)\right] \in \mathbb{R}^{C}, \tag{3.15}$$

where $\mathbf{W}_1 \in \mathbb{R}^{(C/r)\times C}$ and $\mathbf{W}_2 \in \mathbb{R}^{C\times(C/r)}$ are the weights of the two fully connected layers and $r = 4$ is the reduction ratio. The ReLU activation Nair and Hinton (2010) is applied between the two fully connected layers, and the sigmoid function maps the weights to the range $(0, 1)$. The feature tensor is then rescaled channel-wise:

$$\tilde{\mathbf{F}} = \mathbf{w} \odot \mathbf{F}, \tag{3.16}$$

where $\odot$ denotes channel-wise multiplication.

Image features related to the target properties occur at different spatial scales. SE-Atten-Res3D therefore extracts features through four residual stages with 32, 64, 128, and 256 channels. The output of each stage is reduced by global average pooling, and the four pooled vectors are concatenated:

$$\mathbf{F}_{\mathrm{agg}} = \left[\mathrm{GAP}(\mathbf{S}_1) \parallel \mathrm{GAP}(\mathbf{S}_2) \parallel \mathrm{GAP}(\mathbf{S}_3) \parallel \mathrm{GAP}(\mathbf{S}_4)\right], \tag{3.17}$$

where $\mathbf{S}_m$ is the output tensor of the m-th residual stage and $\parallel$ denotes vector concatenation. Because the four stage outputs contain 32, 64, 128, and 256 channels, $\mathbf{F}_{\mathrm{agg}}$ contains 480 components.

The aggregated vector was passed through a shared fully connected block. The block consists of a fully connected layer with 256 nodes, ReLU activation, and dropout, followed by a fully connected layer with 128 nodes, ReLU activation, and dropout. Dropout was applied after each ReLU activation with a probability of $p = 0.4$ Srivastava et al. (2014). A final linear layer with five output nodes produced the standardized target values of k, ϕ, σ^e, K, and G.

Ten SE-Atten-Res3D models were trained independently using the same architecture, training dataset, loss function, optimizer, and hyperparameters but different

random seeds Lakshminarayanan et al. (2017). Each seed controlled weight initialization, mini-batch ordering, and dropout masks. The ten training runs therefore produced reproducible but different trained models. Each input subvolume was passed through the ten trained models, producing ten predictions for each physical property. For each property, these ten predictions were averaged and the average was used as the final prediction.

3.3.6 Whole-Core Property Mapping and Upscaling

After SE-Atten-Res3D training, the ensemble was applied to the 640,000 non-overlapping LR subvolumes extracted from the whole-core LR CT image. These subvolumes formed an $80 \times 80 \times 100$ Cartesian grid. Each grid cell was assigned the predicted local values of k, ϕ, σ^e, K, and G for the corresponding subvolume. The whole-core effective properties calculated from these maps were k_{upscaled}, $\sigma^e_{\text{upscaled}}$, K_{upscaled}, and G_{upscaled}.

For permeability and electrical-conductivity upscaling, directional property values were required along the x, y, and z axes of the Cartesian grid. Each LR subvolume was therefore evaluated in three orientations, including the original orientation and two rotated orientations, using the trained SE-Atten-Res3D ensemble. The resulting directional values were assigned to the diagonal components of the local permeability and electrical-conductivity tensors:

$$\mathbf{k}(\mathbf{x}) = \operatorname{diag}\left(k_x(\mathbf{x}), k_y(\mathbf{x}), k_z(\mathbf{x})\right), \qquad \boldsymbol{\sigma}^e(\mathbf{x}) = \operatorname{diag}\left(\sigma^e_x(\mathbf{x}), \sigma^e_y(\mathbf{x}), \sigma^e_z(\mathbf{x})\right). \tag{3.18}$$

Only diagonal components were used because the rotation-based prediction provides directional effective values along the Cartesian grid axes, whereas the off-diagonal components were not estimated.

For permeability upscaling, the predicted permeability tensor field $\mathbf{k}(\mathbf{x})$ was used in a Darcy-flow calculation. The single-phase Darcy-flow equation was solved using MRST with a two-point finite-volume formulation Lie (2019a):

$$\nabla \cdot \mathbf{u} = 0, \qquad \mathbf{u} = -\frac{1}{\mu}\mathbf{k}(\mathbf{x})\nabla p, \tag{3.19}$$

where $\mathbf{u}$ is the Darcy velocity, p is pressure, $\mathbf{k}(\mathbf{x})$ is the local permeability tensor, and μ is fluid viscosity. Fixed pressures were applied to the inlet and outlet faces normal to the flow direction, and no-flow boundary conditions were applied to the remaining faces. After the pressure field was solved, k_{upscaled} in the flow direction was

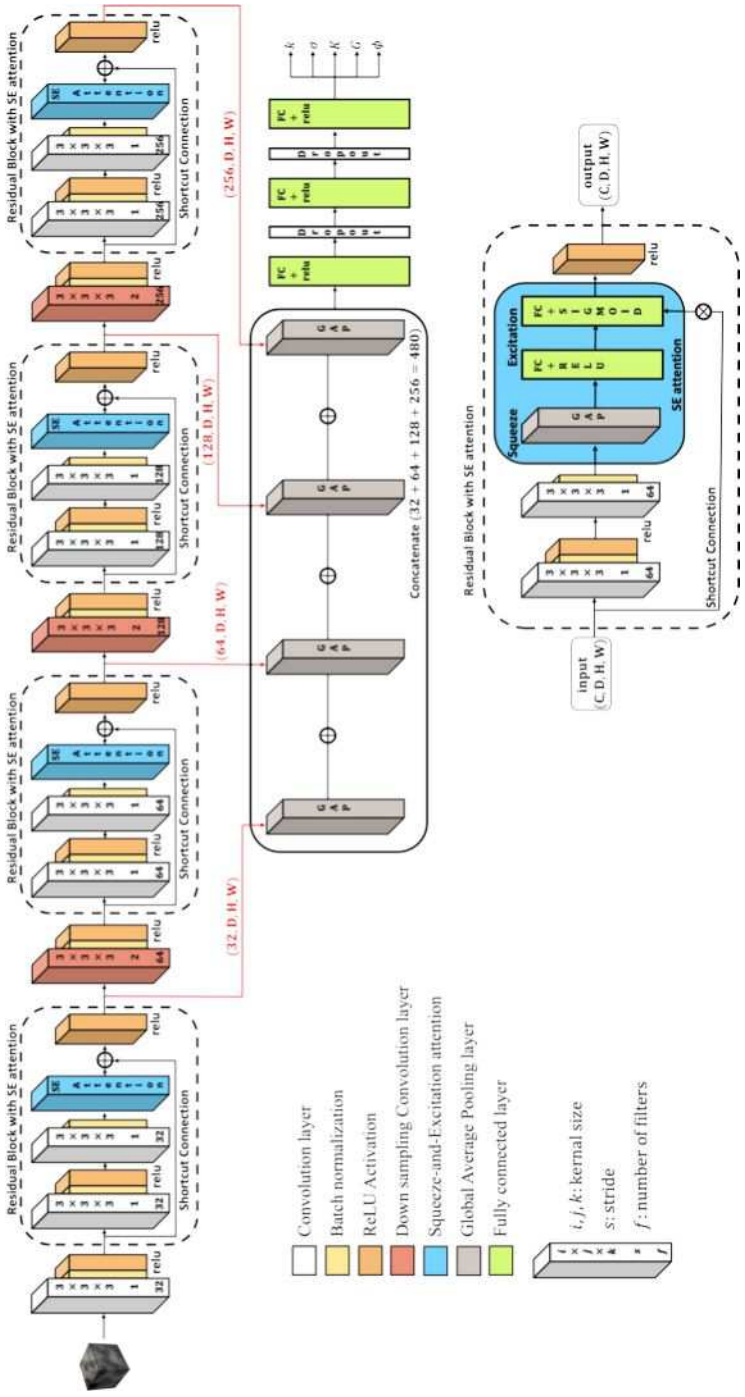


Figure 3.6: Architecture of SE-Atten-Res3D. The network contains four residual stages with SE channel-attention modules and channel numbers of 32, 64, 128, and 256. The output of each stage is reduced by global average pooling, and the four pooled vectors are concatenated into a 480-component feature vector. The aggregated vector is passed through shared fully connected and dropout layers to predict k, ϕ, σ^e, K, and G. The inset shows the residual block with the SE module. C denotes the number of channels, and D, H, and W denote the spatial dimensions.

calculated from the total volumetric flux through the outlet face and the imposed pressure difference using Darcy's law.

For electrical-conductivity upscaling, the FEM formulation described in Section 3.3.3 was applied to the whole-core electrical-conductivity tensor field. Each grid cell was assigned the local tensor $\boldsymbol{\sigma}^e(\mathbf{x})$ defined in Eq. 3.18. A macroscopic electric field was applied to the whole-core domain under PBC, and the electric-potential field was obtained by minimizing the total electrical energy using the CG method. After convergence, $\sigma^e_{\text{upscaled}}$ in the applied-field direction was calculated from the volume-averaged current density using Eq. 3.3.

For elastic upscaling, the FEM formulation described in Section 3.3.3 was applied to the predicted whole-core elastic-property maps. Each grid cell was assigned the local bulk modulus $K(\mathbf{x})$ and shear modulus $G(\mathbf{x})$, and the local stiffness tensor was constructed from these values. The same prescribed macroscopic strain tensor and PBC used for the HR subvolume calculations were applied to the whole-core grid. The finite-element problem was solved by minimizing the elastic strain energy using the CG method. After convergence, K_{upscaled} and G_{upscaled} were calculated from the volume-averaged stress and the prescribed macroscopic strain using Eqs. 3.5 and 3.6, respectively.

3.3.7 Stochastic Rock Physics Model

The S-RPM was formulated from the porosity–property relationships obtained from the SE-Atten-Res3D whole-core prediction maps. The calibration dataset consisted of 640,000 cell-wise porosity–property pairs, (ϕ, k), (ϕ, σ^e), (ϕ, K), and (ϕ, G). The observed porosity range in this dataset was $0.128388 \leq \phi \leq 0.329788$.

For each physical property q, the S-RPM is written as

$$Y_q(\mathbf{x}) = \mu_q\left(\phi(\mathbf{x})\right) + \mathcal{U}_q(\mathbf{x}), \qquad q \in \{k, \sigma^e, K, G\}, \tag{3.20}$$

where Y_q is the modeled property variable, $\mu_q(\phi)$ is the fitted deterministic porosity–property relationship, and $\mathcal{U}_q(\mathbf{x})$ is the stochastic residual around that relationship. Permeability and electrical conductivity were modeled in base-10 logarithmic space, whereas the elastic moduli were modeled in linear space:

$$Y_k = \log_{10} k_{\text{mD}}, \qquad Y_{\sigma^e} = \log_{10} \sigma^e, \qquad Y_K = K, \qquad Y_G = G. \tag{3.21}$$

The deterministic permeability relationship was based on the Kozeny–Carman

form Carman (1937); Costa (2006). In logarithmic form, the S-RPM permeability model is

$$\log_{10} k_{\mathrm{mD}}(\mathbf{x}) = a_k + b_k \log_{10}\left[\frac{\phi(\mathbf{x})^3}{(1-\phi(\mathbf{x}))^2}\right] + \mathcal{U}_k(\mathbf{x}). \tag{3.22}$$

The electrical-conductivity relationship was based on Archie's law Archie et al. (2003); Winsauer et al. (1952):

$$\log_{10} \sigma^e(\mathbf{x}) = \log_{10}\left(a_{\sigma^e}\sigma_w\right) + b_{\sigma^e}\log_{10}\phi(\mathbf{x}) + \mathcal{U}_{\sigma^e}(\mathbf{x}), \tag{3.23}$$

where $\sigma_w = 0.3$ S/m is the pore-fluid conductivity.

For the elastic properties, the reference porosity dependence was defined using the critical-porosity model Nur et al. (1995) combined with Gassmann's equation Gassmann (2007); Mavko et al. (2020). The dry-frame moduli are

$$K_{\mathrm{dry}}(\phi) = K_s\left(1-\frac{\phi}{\phi_c}\right), \qquad G_{\mathrm{dry}}(\phi) = G_s\left(1-\frac{\phi}{\phi_c}\right), \tag{3.24}$$

where $\phi_c = 0.4$, $K_s = 40.1$ GPa, and $G_s = 33.2$ GPa. The saturated moduli are

$$K_{\mathrm{sat}}(\phi) = K_{\mathrm{dry}}(\phi) + \frac{\left[1-K_{\mathrm{dry}}(\phi)/K_s\right]^2}{\dfrac{\phi}{K_w} + \dfrac{1-\phi}{K_s} - \dfrac{K_{\mathrm{dry}}(\phi)}{K_s^2}}, \qquad G_{\mathrm{sat}}(\phi) = G_{\mathrm{dry}}(\phi), \tag{3.25}$$

where $K_w = 2.18$ GPa.

The S-RPM elastic-property models are then written as

$$K(\mathbf{x}) = a_K K_{\mathrm{sat}}\left(\phi(\mathbf{x})\right) + b_K + \mathcal{U}_K(\mathbf{x}), \tag{3.26}$$

and

$$G(\mathbf{x}) = a_G G_{\mathrm{sat}}\left(\phi(\mathbf{x})\right) + b_G + \mathcal{U}_G(\mathbf{x}). \tag{3.27}$$

For all four properties, the stochastic residual is represented as

$$\mathcal{U}_q(\mathbf{x}) = S_q\left(\phi(\mathbf{x})\right)P(\mathbf{x}), \qquad S_q(\phi) = c_q + d_q\phi, \qquad q \in \{k, \sigma^e, K, G\}, \tag{3.28}$$

where $S_q(\phi)$ gives the porosity-dependent standard deviation of the residual for property q, and $P(\mathbf{x})$ is a dimensionless spatial random field with zero mean and unit variance. The same $P(\mathbf{x})$ is used for all four properties, whereas $S_q(\phi)$ is calibrated separately for each property.

The coefficients a_q and b_q were first obtained by fitting the corresponding deter-

ministic porosity–property relationship to the 640,000 cell-wise pairs. After fitting the deterministic relationship, the residual for cell i was calculated as

$$\mathcal{U}_{q,i} = Y_{q,i} - \mu_q(\phi_i), \tag{3.29}$$

where i denotes the cell index.

The porosity range was divided into 19 equal-width bins. Only bins containing more than 10 cells were retained for the residual standard-deviation fitting. For each retained bin b, the standard deviation of the corresponding residuals was calculated and denoted by $s_{q,b}$. Each bin therefore provided one calibration point $(\phi_{\mathrm{mid},b}, s_{q,b})$, where $\phi_{\mathrm{mid},b}$ is the midpoint of the porosity interval. A least-squares line was fitted to these calibration points:

$$s_{q,b} \approx c_q + d_q \phi_{\mathrm{mid},b}. \tag{3.30}$$

The fitted intercept c_q and slope d_q define the porosity-dependent residual amplitude $S_q(\phi)$.

The fitted coefficients used in the S-RPM are listed in Table 3.1.

Table 3.1: Fitted coefficients used in the S-RPM. The coefficients a_q and b_q define the deterministic porosity–property relationship, whereas c_q and d_q define the porosity dependence of the residual standard deviation.

Property	Property variable	a_q	b_q	c_q	d_q
k	$\log_{10} k_{\mathrm{mD}}$	4.4764	0.9208	0.1602	-0.3914
σ^e	$\log_{10} \sigma^e$	1.6121	2.2035	0.0669	-0.1682
K	K	0.7664	3.8844	0.1903	-0.0421
G	G	0.8250	2.2072	0.1588	1.3642

The spatial distribution of the stochastic residual was assigned using the random field $P(\mathbf{x})$ generated by the Spectral Randomization Method. The field is approximated as

$$P(\mathbf{x}) \approx \sqrt{\frac{2\sigma_P^2}{N_{\mathrm{modes}}}} \sum_{m=1}^{N_{\mathrm{modes}}} \cos\left(\mathbf{k}_m \cdot \mathbf{x} + \varphi_m\right), \tag{3.31}$$

where $N_{\mathrm{modes}} = 1000$ is the number of cosine modes. The phase φ_m was sampled independently from a uniform distribution on $[0, 2\pi]$. For independent uniformly distributed phases, each unscaled cosine mode has zero mean and variance $1/2$. The prefactor in Eq. 3.31 therefore scales the variance of the random field to σ_P^2. In this study, $\sigma_P^2 = 1$, so $S_q(\phi)$ directly represents the standard deviation of the stochastic residual at porosity ϕ.

The wave vector $\mathbf{k}_m$ controls the direction and spatial wavelength of each cosine

mode. Its components along the three cell-index coordinate directions were sampled as

$$k_{m,j} = \frac{\sqrt{2}}{\ell_j}\epsilon_{m,j}, \qquad \epsilon_{m,j} \sim \mathcal{N}(0,1), \qquad j \in \{1,2,3\}, \tag{3.32}$$

where ℓ_j is the correlation-length parameter along direction j, expressed in grid-cell units. The corresponding wavelength of a mode is

$$\lambda_{m,j} = \frac{2\pi}{|k_{m,j}|}. \tag{3.33}$$

Increasing ℓ_j reduces the typical magnitude of $k_{m,j}$ and therefore produces slower spatial variation in $P(\mathbf{x})$ along that direction. Decreasing ℓ_j produces shorter-wavelength variations. Because the stochastic residual is given by $\mathcal{U}_q(\mathbf{x}) = S_q(\phi)P(\mathbf{x})$, the correlation-length parameters control the spatial size of contiguous positive and negative residual regions.

For the core-scale reconstruction in this chapter, the correlation-length vector was set to

$$\boldsymbol{\ell} = [2,1,1] \quad \text{cells}, \tag{3.34}$$

which produced spatial variations comparable in scale to those observed in the SE-Atten-Res3D prediction maps. The property fields reconstructed by the S-RPM were subsequently upscaled using the same Darcy-flow and FEM procedures described in Section 3.3.6.

3.4 Results and Discussion

3.4.1 DNS-Derived Physical Properties of HR Subvolumes

The procedures described in Section 3.3.3 were applied to the segmented HR subvolumes to obtain the reference physical-property labels. For each HR subvolume, five reference labels were obtained: ϕ, k, σ^e, K, and G.

Figure 3.7 shows a representative DNS result for one HR subvolume. The selected subvolume has $\phi = 24.702\%$, $k = 1015.952$ mD, $\sigma^e = 0.022087$ S/m, $K = 18.118$ GPa, and $G = 12.438$ GPa. Figure 3.7 shows the segmented geometry and the velocity, current-density, and stress fields for this subvolume. The velocity field is confined to the connected pore space and increases locally where the flow path narrows. The current-density field is concentrated in the connected water-filled pore phase. The stress field is carried by the connected solid phase. Higher stress

appears in thin solid bridges separating adjacent pores, where the applied load is transmitted through a smaller solid cross-sectional area. These field patterns provide a qualitative check on the physical consistency of the DNS calculations.

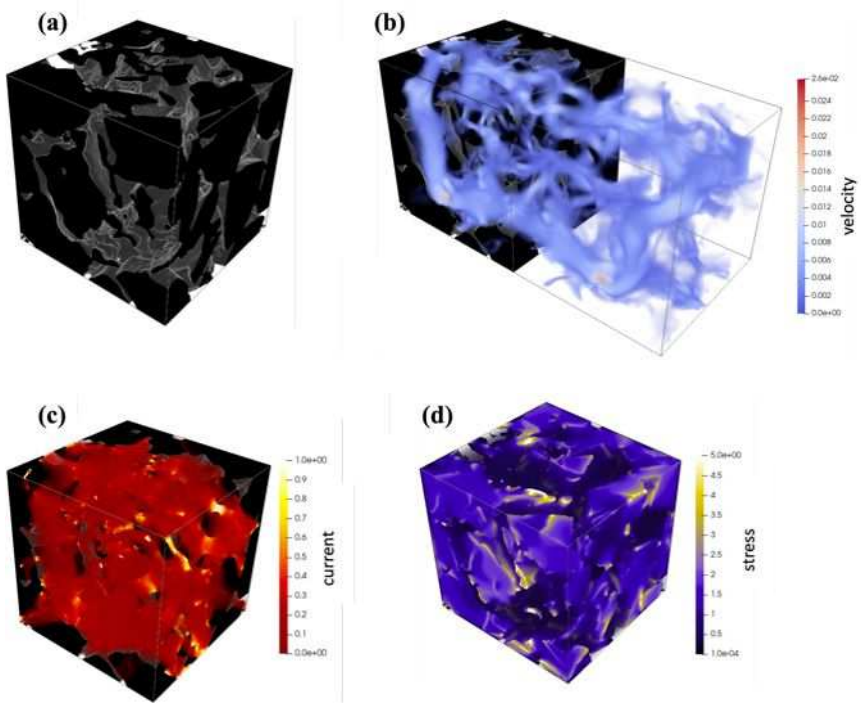


Figure 3.7: Representative DNS result for an HR subvolume whose five reference labels are close to the corresponding median values of the 12,000-subvolume dataset. Panel (a) shows the segmented three-dimensional pore–solid geometry. Panel (b) shows the velocity-magnitude field in the connected pore space. Panel (c) shows the current-density-magnitude field in the water-filled pore phase. Panel (d) shows the stress-magnitude field in the solid phase.

Figure 3.8 shows the relative-frequency histograms of the DNS-derived properties and porosity, and Table 3.2 summarizes their statistics. Here, P_5 and P_{95} denote the 5th- and 95th-percentile values, respectively, and the ratio P_{95}/P_5 expresses the spread between these two values as a dimensionless factor. Skewness was used to describe the asymmetry of each distribution and was calculated as

$$\text{Skewness} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{x_i - \bar{x}}{s}\right)^3, \tag{3.35}$$

where x_i is the property value of the i-th subvolume, $\bar{x}$ is the mean, s is the standard deviation, and N is the number of subvolumes. A skewness value close to zero indicates an approximately symmetric distribution, whereas positive and negative

values indicate tails toward larger and smaller values, respectively.

Table 3.2: Statistics of HR-derived physical-property labels for 12,000 HR subvolumes. P_5 and P_{95} denote the 5th and 95th percentiles, respectively.

Property	Unit	Mean	Median	P_{95}/P_5	Skewness
Porosity ϕ	%	24.638	24.611	1.646	-0.078
Permeability k	mD	1110.887	1028.660	6.876	0.888
Conductivity σ^e	S/m	0.023	0.022	3.217	0.416
Bulk modulus K	GPa	18.091	17.927	1.584	0.437
Shear modulus G	GPa	12.615	12.494	1.976	0.255

The porosity values range from 12.010% to 36.221%, showing that the HR subvolumes include a broad range of local pore-volume fractions. The small difference between the mean and median, together with the skewness close to zero, indicates that the porosity distribution is nearly symmetric.

The transport properties show larger P_{95}/P_5 ratios than porosity. For permeability, $P_{95}/P_5 = 6.876$, meaning that the 95th-percentile permeability is approximately seven times the 5th-percentile value. Together with the positive skewness, this indicates a tail toward high-permeability subvolumes. Electrical conductivity also shows a positive skewness, although its P_{95}/P_5 ratio is smaller than that of permeability.

The larger spread of k than σ^e can be interpreted from their different dependence on pore-throat radius. For laminar flow through an ideal cylindrical throat, the Hagen–Poiseuille relation gives a flow rate proportional to the fourth power of throat radius. In contrast, electrical current through a water-filled cylindrical path is proportional to cross-sectional area and therefore to the square of radius. Local differences in throat radius or constriction can therefore affect permeability more strongly than electrical conductivity. Both properties require connected pore pathways, but their sensitivity to throat size differs because of these radius dependences.

The elastic moduli were obtained from the volume-averaged stress–strain relation under prescribed strain conditions and are primarily related to the connected solid framework in each HR subvolume. The P_{95}/P_5 ratio is smaller for K than for G, indicating that K varies less than G among the HR subvolumes. This difference is consistent with their deformation modes. K describes resistance to volumetric compression, whereas G describes resistance to shape change. Because shear deformation depends on the geometry and orientation of solid contacts, G shows a larger relative spread than K.

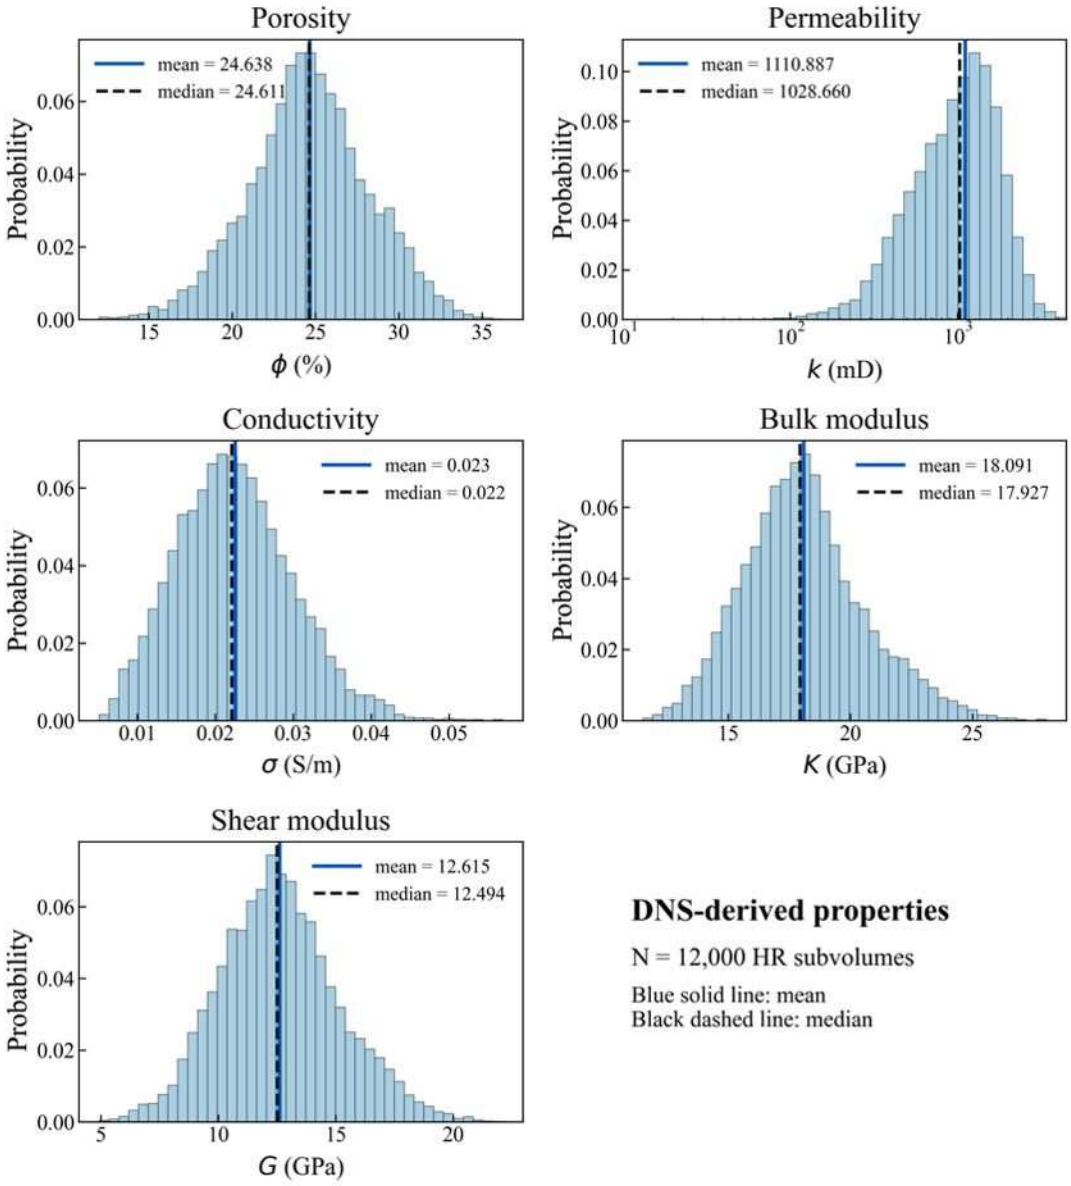


Figure 3.8: Relative-frequency histograms of the five reference labels obtained for 12,000 HR subvolumes: porosity, permeability, electrical conductivity, bulk modulus, and shear modulus. The permeability axis is logarithmic. The blue solid and black dashed vertical lines indicate the mean and median values, respectively.

3.4.2 Validation of DSgan-Based Downsampling

The DSgan-based downsampling procedure described in Section 3.3.4 was evaluated using qualitative image comparison and quantitative CT-value histogram comparison between CLR and RLR. For comparison, a Gaussian-LR volume (GLR) was created from RHR by applying a three-dimensional Gaussian filter followed by downsampling to the LR grid.

Figure 3.9 compares representative RHR, GLR, CLR, and RLR images. The displayed RHR, GLR, and CLR panels are derived from the same source subvolume, whereas RLR is an unpaired example from the actual LR CT domain. CLR retains visually identifiable grain and pore patterns from RHR while showing reduced sharp contrast and coarser grayscale transitions similar to those observed in RLR. GLR retains the smoothed source pattern but does not reproduce the local grayscale variations observed in RLR. In this comparison, CLR is closer to RLR than GLR in grayscale texture.

For the quantitative comparison, normalized CT-value histograms were constructed for CLR and RLR (Figure 3.10). The CLR histogram was calculated from all 12,000 CLR subvolumes, whereas the RLR histogram was calculated from the actual 20 μm/voxel LR CT data. The agreement was evaluated using the histogram overlap coefficient, which measures the shared area between two normalized histograms. The overlap coefficient was 0.9158, indicating that 91.6% of the normalized histogram area was shared by CLR and RLR.

The qualitative comparison and histogram overlap coefficient show that CLR has a CT-value distribution close to RLR while retaining visually identifiable grain and pore patterns from RHR.

3.4.3 Evaluation of Multi-Property Prediction

The SE-Atten-Res3D ensemble was evaluated using the paired LR-like inputs and HR-derived reference labels prepared in the previous sections. The dataset contained 12,000 paired subvolumes, of which 10,800 were used for training and 1,200 were used for testing.

Figure 3.11 and its continuation compare the HR-derived reference values with the SE-Atten-Res3D ensemble predictions. In both the training and test panels, the predictions follow the diagonal trend over the sampled range of each property. The test points are more widely distributed around the diagonal than the training points, but the diagonal relation is maintained for all five properties. The training R^2

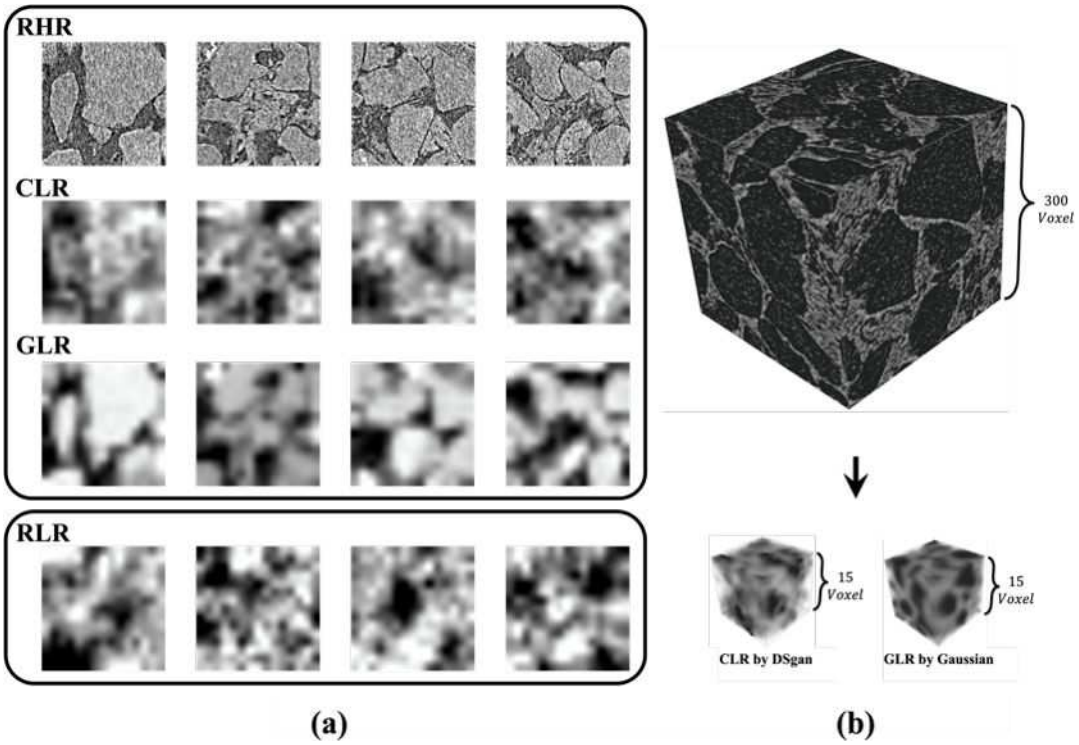


Figure 3.9: Qualitative comparison of DSgan-based downsampling. Panel (a) shows representative two-dimensional slices of RHR, GLR, CLR, and RLR. Panel (b) shows volume renderings of the source RHR subvolume and the corresponding CLR and GLR subvolumes.

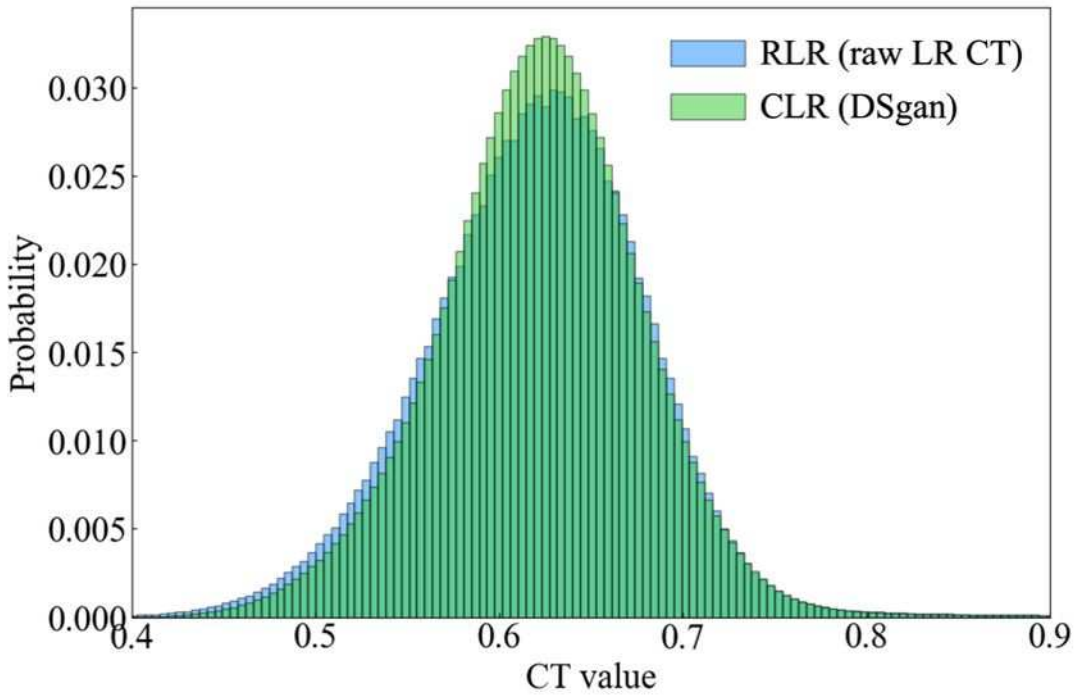


Figure 3.10: Normalized CT-value histograms used for quantitative validation of DSgan-based downsampling. The blue histogram represents RLR values from the raw 20 μm/voxel LR CT data, and the green histogram represents CLR values from all 12,000 DSgan-converted LR-like subvolumes.

values range from 0.98 to 0.99. The test values range from 0.80 to 0.86: $R_k^2 = 0.80$, $R_{\sigma^e}^2 = 0.81$, $R_K^2 = 0.86$, $R_G^2 = 0.81$, and $R_\phi^2 = 0.85$. These results show that the LR-like inputs contain image information related to the HR-derived reference properties.

The differences in test R^2 can be interpreted from the image features required for each property. K and ϕ show the highest test values, with $R_K^2 = 0.86$ and $R_\phi^2 = 0.85$. Porosity is directly related to void fraction, and K is related to the amount of load-bearing solid phase; these volume-fraction-related features remain partly represented in the coarsened input. Permeability and electrical conductivity show $R_k^2 = 0.80$ and $R_{\sigma^e}^2 = 0.81$. These transport properties depend on connected pore pathways, tortuosity, and local pore-throat constrictions. In the 15^3-voxel LR-like input, narrow pore throats and local constrictions are represented only as coarse grayscale variations. The shear modulus shows $R_G^2 = 0.81$. Because shear deformation depends on the geometry and orientation of solid contacts, G is more difficult to predict than ϕ and K from the coarsened input.

The SE-Atten-Res3D ensemble achieved $R^2 \geq 0.80$ for all five properties on the test set. This performance supports application of the trained ensemble to the whole-core LR CT image in the following subsection.

3.4.4 Core-Scale Property Mapping and Porosity–Property Relations

The trained SE-Atten-Res3D ensemble was applied to all LR subvolumes in the whole-core LR CT volume to generate core-scale maps of the five physical properties. The spatial distributions of the predicted properties are shown in Figure 3.12. The five maps are displayed on the same subvolume grid.

Regions with higher predicted porosity generally show higher k and σ^e and lower K and G, whereas lower-porosity regions show the opposite pattern. The red-circled regions correspond to local low-porosity zones and show lower k and σ^e together with higher K and G. A larger low-porosity region is visible near the upper-right part of the displayed volume. In this region, both k and σ^e decrease, whereas K and G increase. The spatial contrast is more pronounced in G than in K, consistent with the greater sensitivity of shear deformation to the geometry and continuity of solid contacts. The dashed rectangle marks a crack-like zone with higher ϕ, k, and σ^e and lower K and G.

To examine the first-order porosity dependence of the predicted properties, the

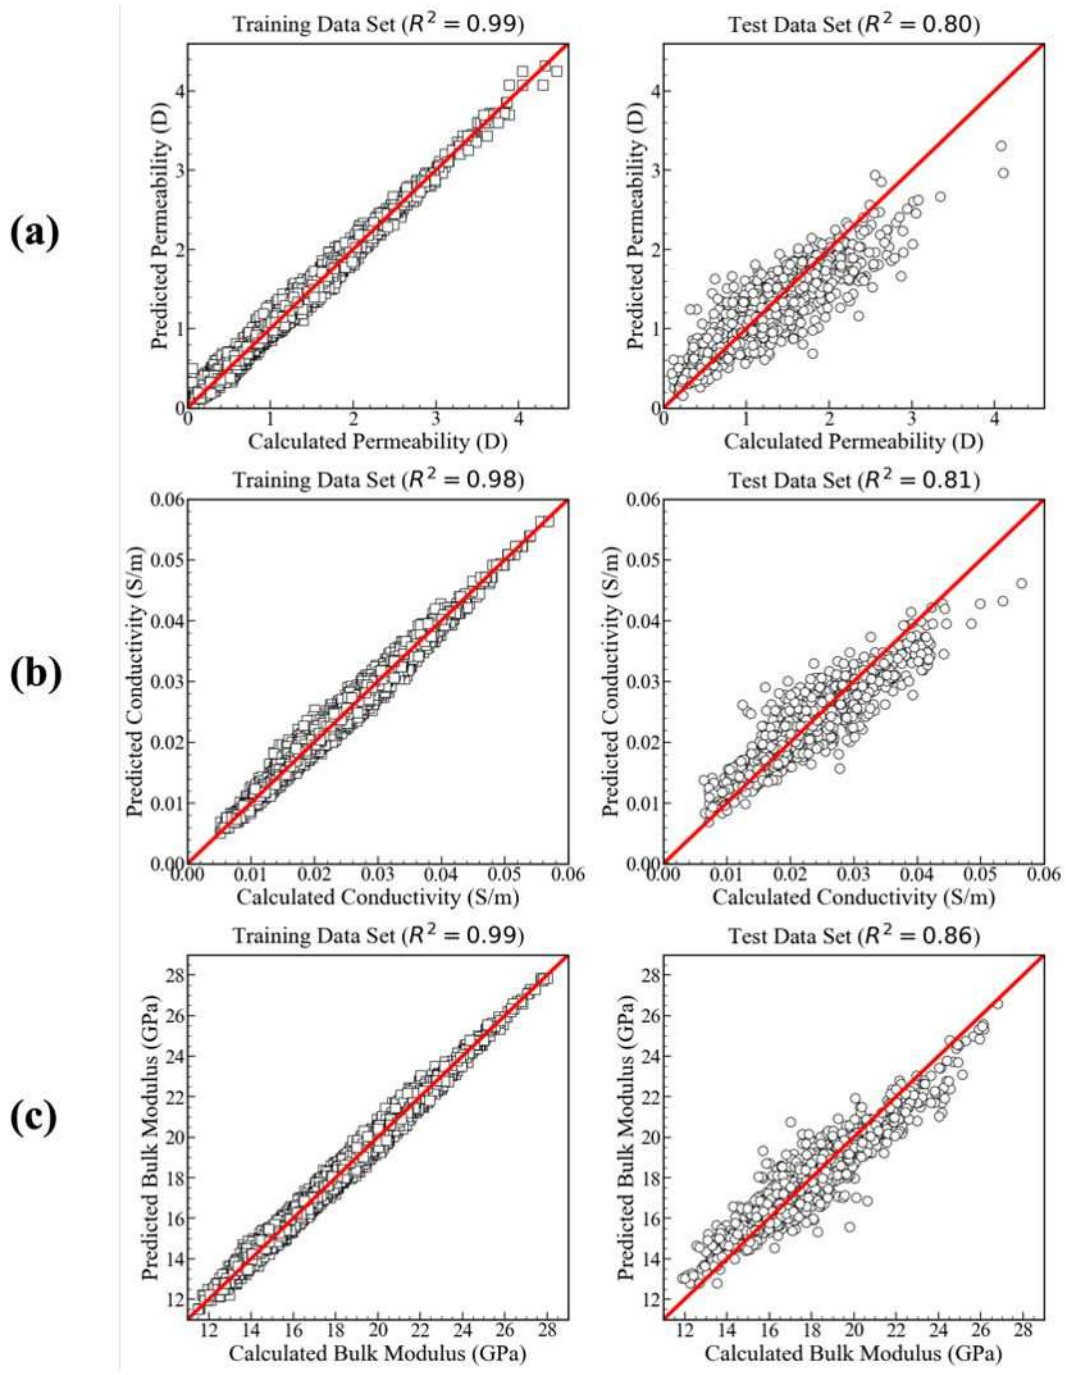


Figure 3.11: Comparison between HR-derived reference values and SE-Atten-Res3D ensemble predictions from the corresponding 15^3-voxel LR-like inputs for k, σ^e, and K. The left panels show the training samples, and the right panels show the test samples. The red line denotes the diagonal relation ($y = x$).

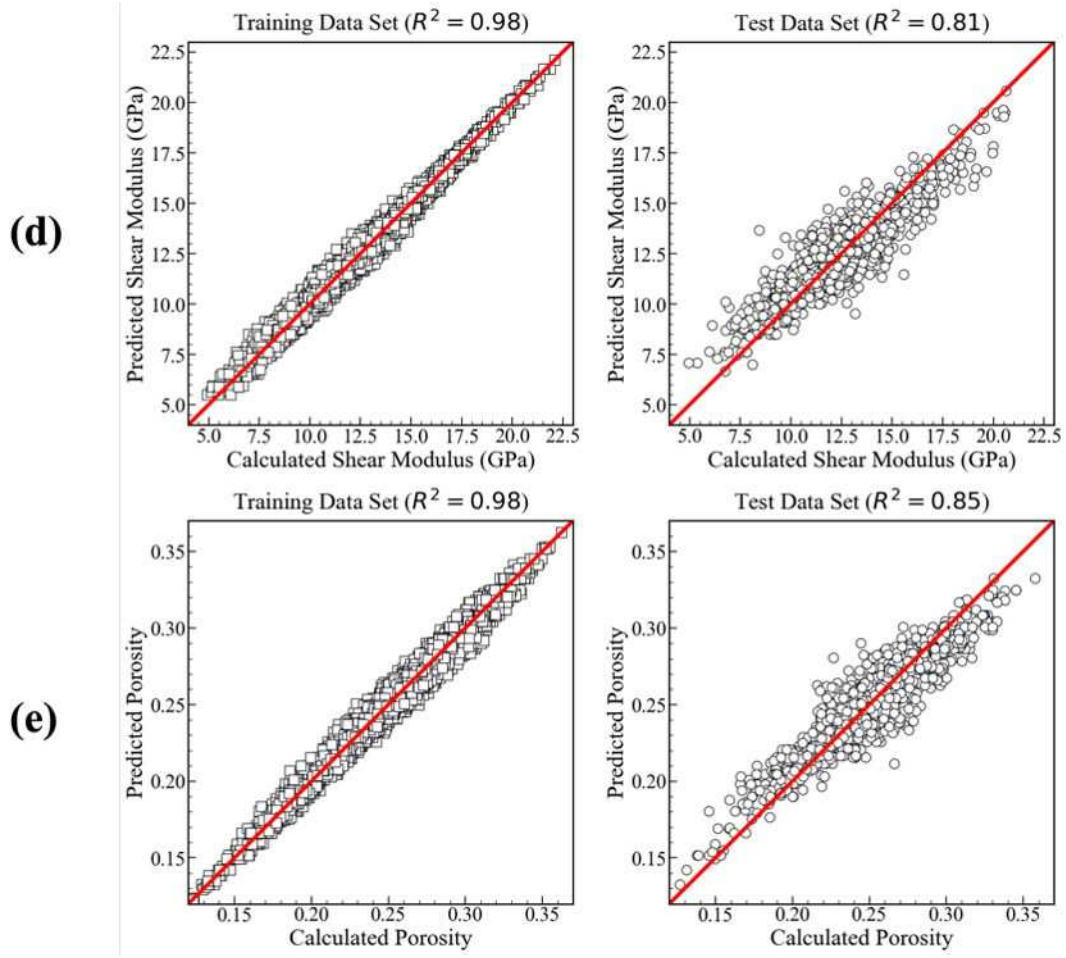


Figure 3.11: Continued comparison between HR-derived reference values and SE-Atten-Res3D ensemble predictions for G and ϕ. The red line denotes the diagonal relation ($y = x$).

porosity–property relationships were compared with conventional rock-physics relationships. The permeability–porosity relationship was compared with the Kozeny–Carman model Carman (1937); Costa (2006),

$$k = C_{\mathrm{KC}} \frac{\phi^3}{(1-\phi)^2}, \tag{3.36}$$

where C_{KC} is an empirical permeability coefficient. The fitted C_{KC} values were 36324.94 mD for the DNS-derived labels and 39064.04 mD for the whole-core predictions. Both datasets show increasing permeability with porosity, consistent with an increase in pore volume and connected flow paths.

The electrical-conductivity–porosity relationship was compared with Archie's law Archie et al. (2003); Winsauer et al. (1952),

$$\sigma^e = \frac{\sigma_w}{a} \phi^m, \tag{3.37}$$

where σ_w is the pore-fluid conductivity, a is a dimensionless empirical coefficient, and m is the cementation exponent. For the DNS-derived labels, the fitted values were $1/a = 1.3475$ and $m = 2.0848$; for the whole-core predictions, they were $1/a = 1.6121$ and $m = 2.2035$. The predicted conductivity increases with porosity, consistent with ionic conduction through connected fluid-filled pore space.

For the elastic properties, the porosity dependence was compared with the critical-porosity model Nur et al. (1995) combined with Gassmann's equation Gassmann (2007); Mavko et al. (2020). The dry-frame moduli were expressed as

$$K_{\mathrm{dry}} = K_s \left(1 - \frac{\phi}{\phi_c}\right), \qquad G_{\mathrm{dry}} = G_s \left(1 - \frac{\phi}{\phi_c}\right), \tag{3.38}$$

where ϕ_c is the critical porosity and K_s and G_s are the bulk and shear moduli of the solid matrix. The saturated bulk modulus was then calculated using Gassmann's equation, and the saturated shear modulus was set equal to the dry shear modulus:

$$K_{\mathrm{sat}} = K_{\mathrm{dry}} + \frac{\left(1 - K_{\mathrm{dry}}/K_s\right)^2}{\dfrac{\phi}{K_w} + \dfrac{1-\phi}{K_s} - \dfrac{K_{\mathrm{dry}}}{K_s^2}}, \qquad G_{\mathrm{sat}} = G_{\mathrm{dry}}, \tag{3.39}$$

where K_w is the bulk modulus of the pore water. The same parameters as those used in the DNS calculations were adopted: $\phi_c = 0.4$ Nur et al. (1995), $K_s = 40.1$ GPa, $G_s = 33.2$ GPa, and $K_w = 2.18$ GPa Mavko et al. (2020). The predicted bulk

and shear moduli decrease with increasing porosity, following the first-order trends described by the critical-porosity/Gassmann relationships. This decrease reflects the reduction in the load-bearing solid fraction as pore volume increases.

The predicted whole-core properties follow the porosity-dependent trends described by the reference rock-physics relationships. Higher porosity is associated with higher permeability and electrical conductivity and with lower bulk and shear moduli. At the same time, cells with similar porosity retain a finite spread in k, σ^e, K, and G. This residual variation forms the basis of the S-RPM reconstruction examined in the following subsection.

3.4.5 S-RPM Reconstruction of Core-Scale Property Maps

The SE-Atten-Res3D predictions show that subvolumes with similar porosity can have different values of k, σ^e, K, and G. The S-RPM described in Section 3.3.7 was used to reconstruct this residual variation from the predicted porosity field.

Figure 3.14 shows the porosity–property distributions reconstructed by the S-RPM together with the fitted deterministic relationships. The reconstructed values follow the same porosity dependence and show a similar overall spread to the SE-Atten-Res3D predictions for all four properties.

Figure 3.15 shows the three-dimensional property maps reconstructed by the S-RPM. The reconstructed maps retain the broad porosity-controlled spatial pattern observed in the SE-Atten-Res3D maps. Regions with higher porosity generally show higher k and σ^e and lower K and G, whereas lower-porosity regions show the opposite pattern.

These comparisons show that the S-RPM reconstructs the porosity–property trends, residual spread, and broad spatial variations represented in the SE-Atten-Res3D predictions. The whole-core effective responses calculated from the SE-Atten-Res3D and S-RPM maps are compared in the following subsection.

3.4.6 Core-Scale Upscaling

The SE-Atten-Res3D and S-RPM property maps were used separately to calculate the whole-core effective values of k, σ^e, K, and G using the Darcy-flow and FEM solvers described in Section 3.3.6.

Figures 3.16 and 3.17 show the velocity, pressure, current-density, and stress fields calculated from the SE-Atten-Res3D and S-RPM maps, respectively. Corresponding panels show similar spatial patterns of high and low field values.

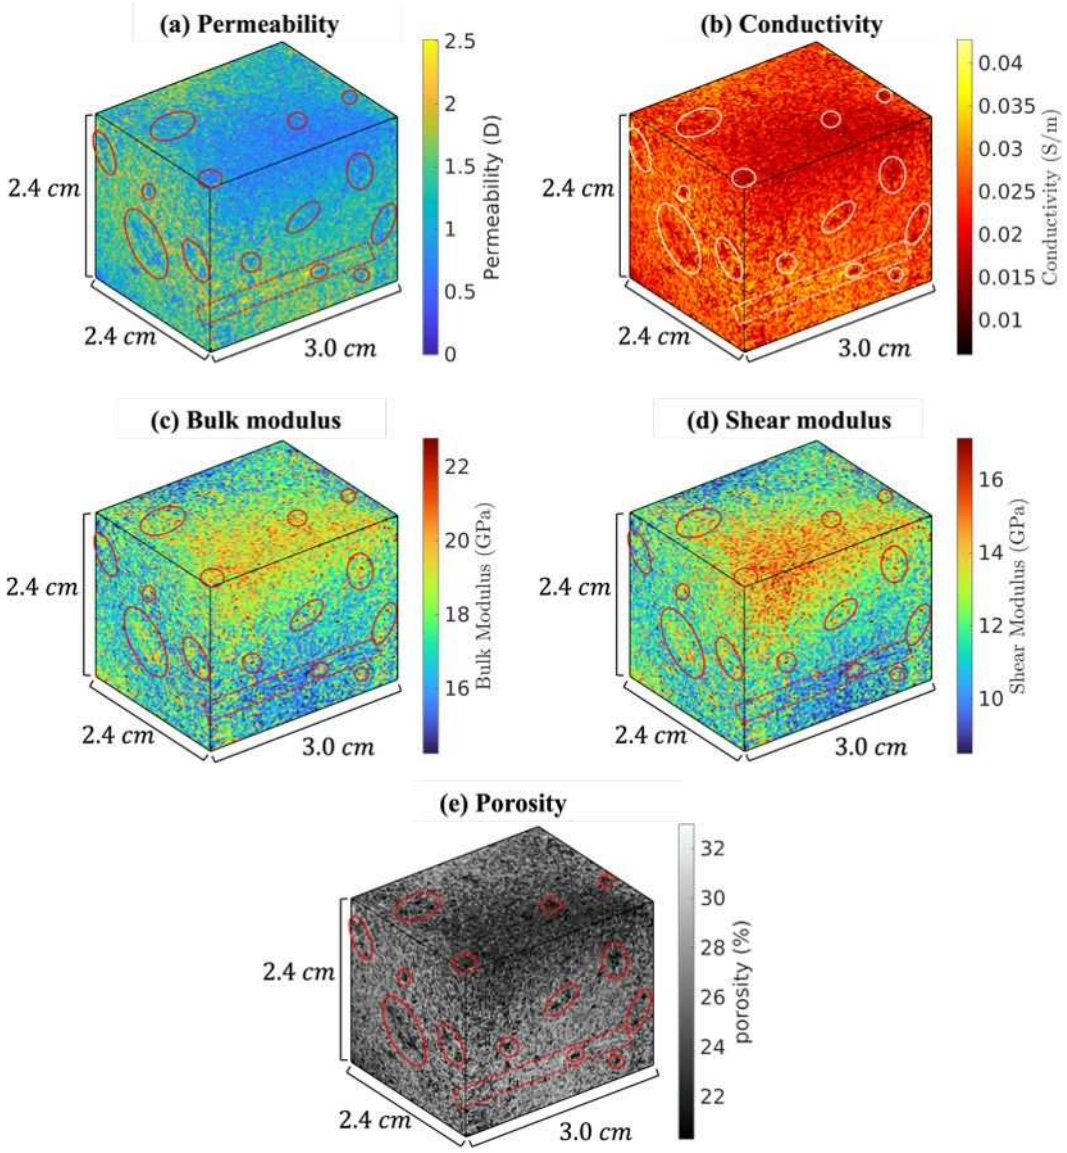
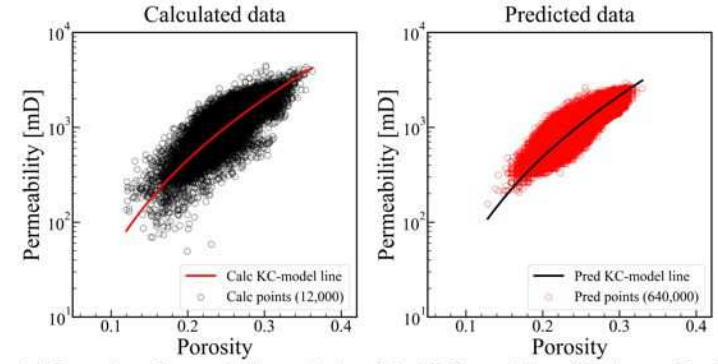
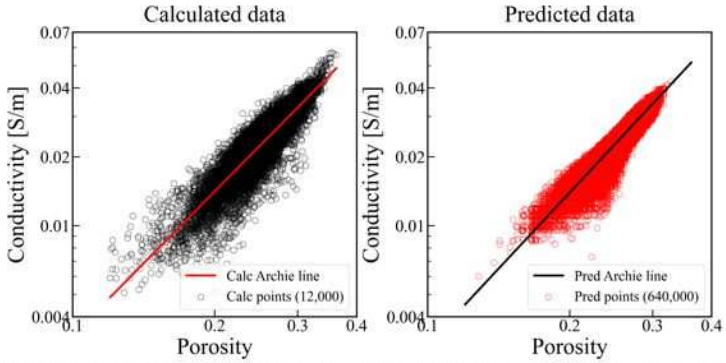


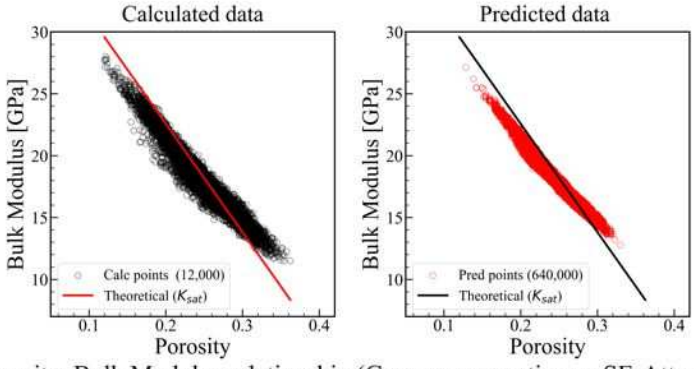
Figure 3.12: Spatial distributions of the predicted physical properties for Sample J7: (a) permeability, (b) electrical conductivity, (c) bulk modulus, (d) shear modulus, and (e) porosity. Red circles mark local low-porosity regions, and the dashed rectangle marks the crack-like zone discussed in the text.



(a) Porosity–Permeability relationship (KC model vs SE-Atten-Res3D)

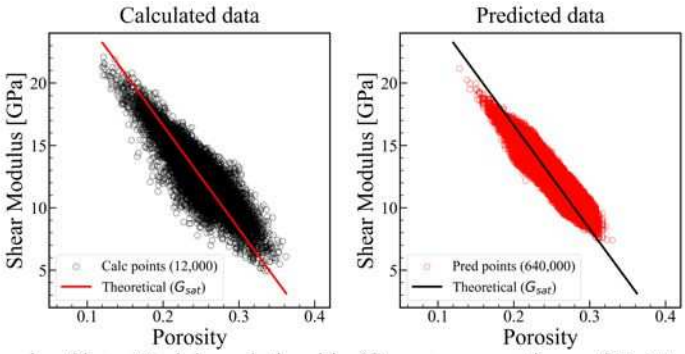


(b) Porosity–Conductivity relationship (Archie model vs SE-Atten-Res3D)



(c) Porosity–Bulk Modulus relationship (Gassmann equation vs SE-Atten-Res3D)

Figure 3.13: Porosity–property relationships for the DNS-derived labels and the SE-Atten-Res3D whole-core predictions. The left panels show DNS-derived values, and the right panels show predicted values for the 640,000 subvolumes in the whole-core LR CT volume. The relationships shown are (a) porosity–permeability, (b) porosity–electrical conductivity, and (c) porosity–bulk modulus. Solid curves indicate reference trends based on the Kozeny–Carman relationship, Archie's law, and the critical-porosity/Gassmann framework.



(d) Porosity–Shear Modulus relationship (Gassmann equation vs SE-Atten-Res3D)

Figure 3.13: Porosity–shear modulus relationship for the DNS-derived labels and the SE-Atten-Res3D whole-core predictions. The solid curves indicate the reference trend based on the critical-porosity model.

Table 3.3 compares the two map-based upscaling results, the arithmetic mean of the local SE-Atten-Res3D predictions, and two deterministic RPM estimates obtained from representative porosity values. Using the SE-Atten-Res3D upscaled values as the reference, the S-RPM values differ by 0.10% for k, 0.09% for σ^e, 0.004% for K, and 0.28% for G. The S-RPM maps therefore reproduce the four whole-core effective properties obtained from the SE-Atten-Res3D maps within 0.3%.

Relative to the SE-Atten-Res3D upscaled values, the arithmetic means are 3.0% higher for k and 1.3% higher for σ^e, whereas they are 1.8% lower for K and 2.2% lower for G. The difference arises because arithmetic means weight all cells equally and do not represent transport-path connectivity or heterogeneous stress–strain fields. For the transport properties, locally high values contribute to the whole-core response only when they form connected transport paths. For the elastic properties, the effective response depends on how local stiffness variations are connected through the load-bearing solid framework.

The deterministic RPM estimates were calculated using two representative porosity values: $\phi = 0.25$, the mean porosity of the center-core region, and $\phi = 0.33$, measured for the whole-core sample. The center-core RPM estimates differ from the SE-Atten-Res3D upscaled values by 5.4–19.0% among the four properties. The whole-core-porosity RPM estimates are 156.4% higher for k and 74.4% higher for σ^e, whereas they are 37.2% lower for K and 53.2% lower for G. These differences show that a single representative porosity does not contain the local property variation and spatial arrangement required to reproduce the map-based whole-core effective properties.

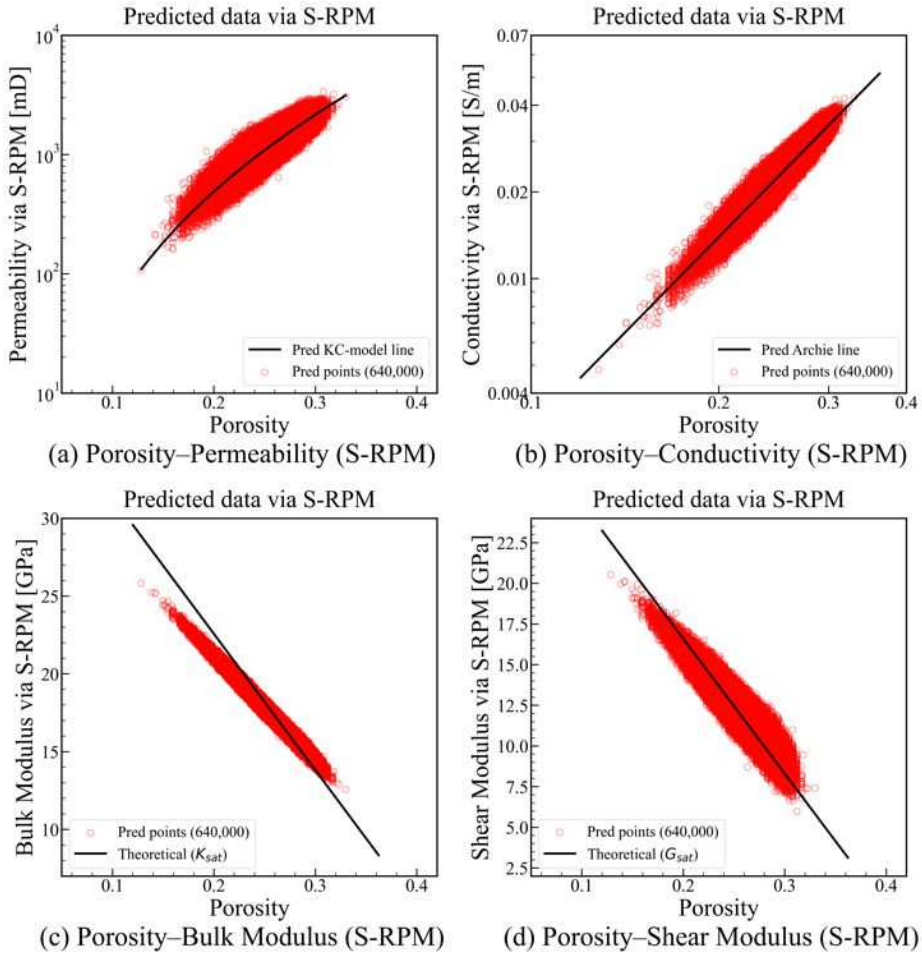


Figure 3.14: Porosity–property distributions reconstructed by the Stochastic Rock Physics Model (S-RPM). Red open circles represent the 640,000 S-RPM reconstructed values for (a) permeability, (b) electrical conductivity, (c) bulk modulus, and (d) shear modulus. Solid black lines denote the fitted deterministic relationships: (a) Kozeny–Carman, (b) Archie's law, and (c, d) critical-porosity/Gassmann-based elastic models.

Table 3.3: Comparison of whole-core effective properties obtained from the SE-Atten-Res3D and S-RPM maps, arithmetic means of local SE-Atten-Res3D predictions, and deterministic single-porosity RPM estimates.

Method	k (mD)	σ^e (S/m)	K (GPa)	G (GPa)
SE-Atten-Res3D upscaled	1234.4844	0.024264	17.6524	12.2459
S-RPM upscaled	1233.2709	0.024286	17.6517	12.2803
Arithmetic mean (SE-Atten-Res3D)	1271.5671	0.024573	17.3368	11.9818
RPM ($\phi = 0.25$, center core)	1000.3335	0.021707	18.6690	12.9065
RPM ($\phi = 0.33$, whole core)	3165.2565	0.042313	11.0871	5.7270

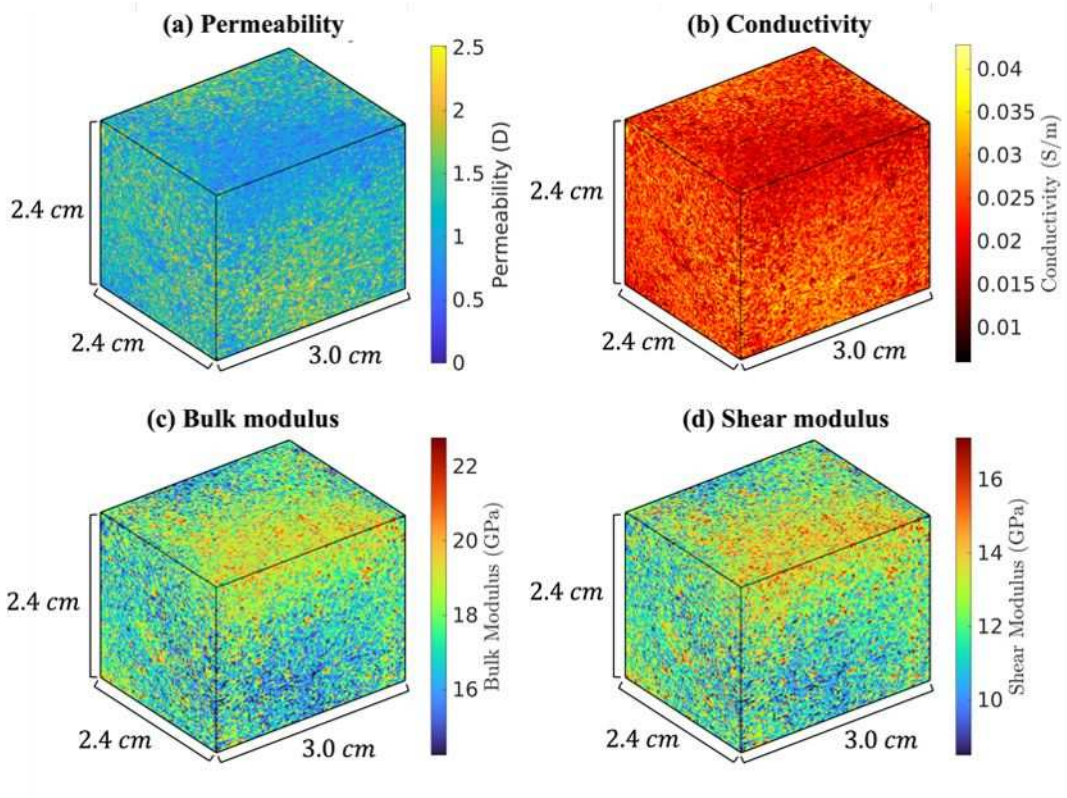


Figure 3.15: Three-dimensional property maps reconstructed by the Stochastic Rock Physics Model (S-RPM): (a) permeability, (b) electrical conductivity, (c) bulk modulus, and (d) shear modulus.

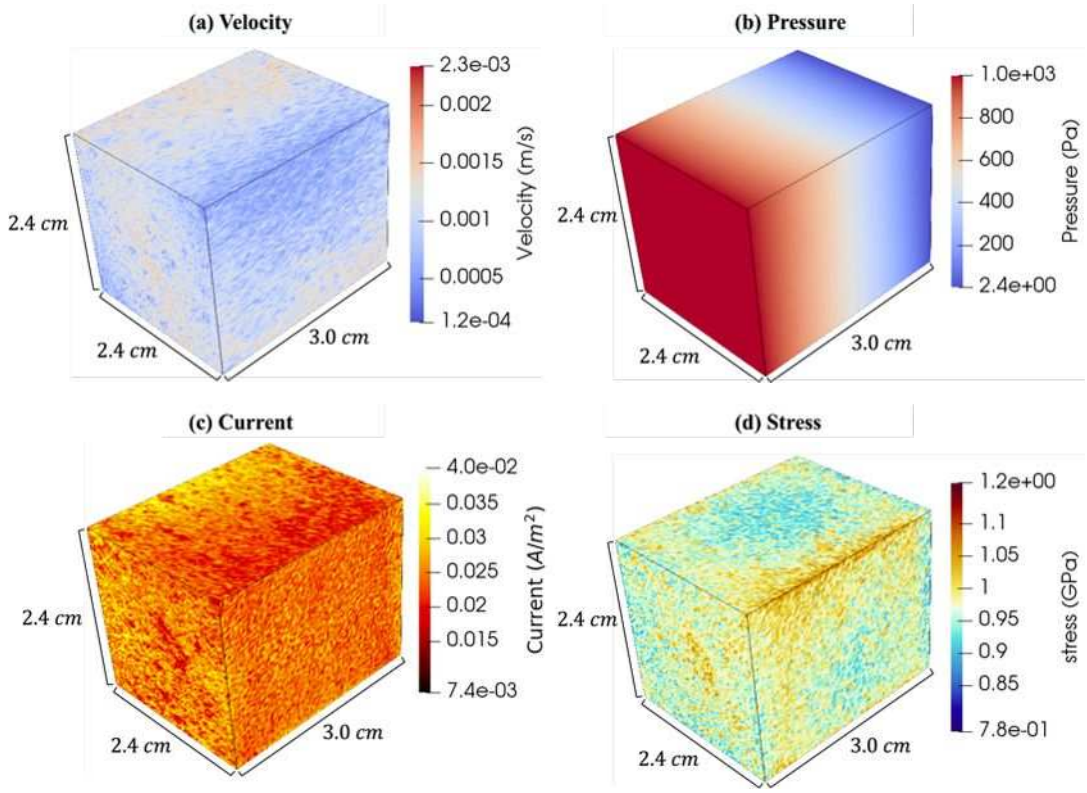


Figure 3.16: Spatial distributions of the physical fields calculated from the SE-Atten-Res3D property maps. (a) Velocity and (b) pressure fields obtained from the Darcy solver for permeability upscaling. (c) Current-density field obtained from the FEM solver for electrical-conductivity upscaling. (d) Stress field obtained from the FEM solver for elastic-modulus upscaling.

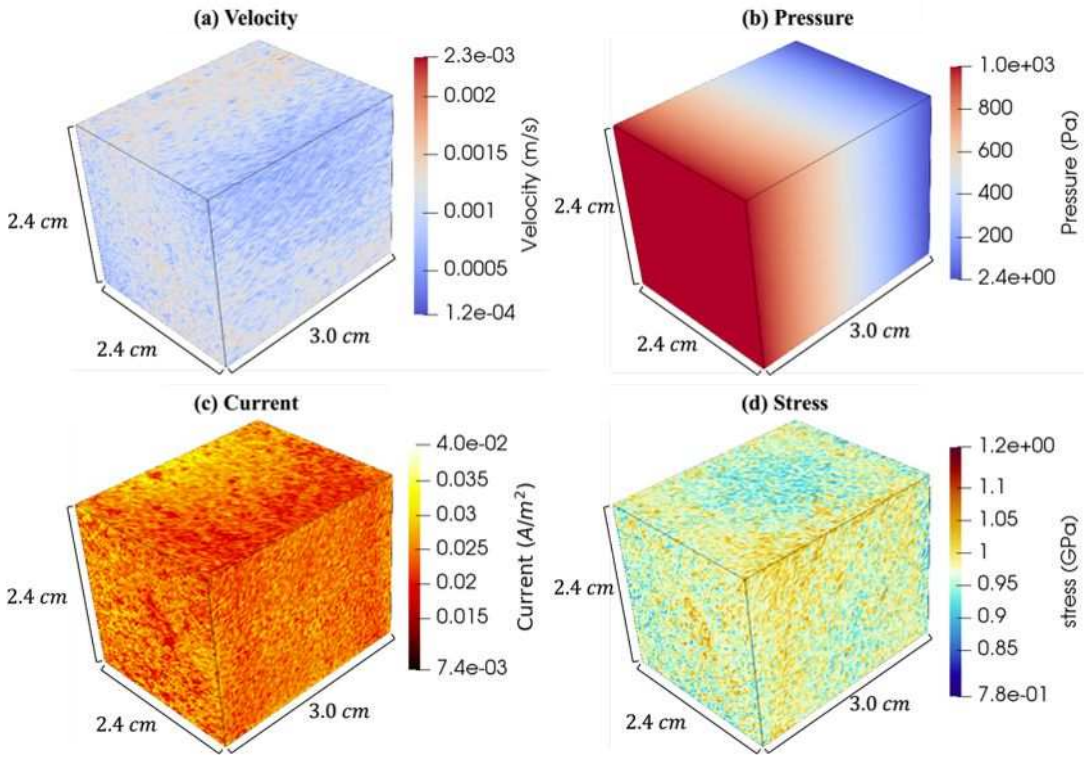


Figure 3.17: Spatial distributions of the physical fields calculated from the S-RPM property maps. (a) Velocity and (b) pressure fields obtained from the Darcy solver for permeability upscaling. (c) Current-density field obtained from the FEM solver for electrical-conductivity upscaling. (d) Stress field obtained from the FEM solver for elastic-modulus upscaling. Corresponding panels use the same color scales as Figure 3.16.

3.5 Summary

This chapter developed a multiscale workflow combining DSgan, SE-Atten-Res3D, and the S-RPM to transfer DNS-derived pore-scale physical information to core-scale property prediction and whole-core effective-property calculation for Sample J7. DSgan constructs registration-free paired datasets, SE-Atten-Res3D predicts five core-scale property maps from LR-like CT inputs, and the S-RPM represents the predicted porosity–property scatter using a porosity-conditioned stochastic residual. The main findings are summarized as follows:

1. **Registration-free paired dataset construction using DSgan:** DSgan converted 300^3-voxel HR subvolumes into 15^3-voxel LR-like subvolumes, allowing each input to be paired directly with physical-property labels calculated from the corresponding HR subvolume without physical HR–LR image registration. The generated CLR images showed a CT-value distribution close to the actual 20 μm/voxel LR CT data, with a histogram overlap coefficient of 0.9158, and showed LR-like grayscale texture more clearly than the Gaussian-smoothed downsampling baseline.

2. **Simultaneous multi-property prediction and core-scale mapping:** The SE-Atten-Res3D ensemble predicted permeability k, porosity ϕ, electrical conductivity σ^e, bulk modulus K, and shear modulus G from 15^3-voxel LR-like CT inputs, with test-set R^2 values of 0.80–0.86. The predicted whole-core maps were used in Darcy-flow and FEM calculations. The resulting effective properties differed from the arithmetic means of the local predictions, showing that the whole-core response depends on the spatial arrangement of local properties.

3. **Porosity-conditioned stochastic modeling using the S-RPM:** The S-RPM represents property variation at a given porosity using a residual amplitude calibrated from the SE-Atten-Res3D porosity–property pairs and a spatially correlated random field. The reconstructed property maps retained the porosity–property trends and residual spread of the SE-Atten-Res3D predictions. Their whole-core effective properties differed by less than 0.3% from those obtained directly from the SE-Atten-Res3D maps.

The workflow has two limitations in its present form. First, the SE-Atten-Res3D model and S-RPM coefficients were calibrated using Sample J7. Their applicability

to other rock types, including carbonates and shales, requires separate validation. Second, the physical-property calculations assume single-phase flow and fully saturated conditions. Multiphase flow, partial saturation, and capillary pressure can alter transport and elastic responses. Future work should test the workflow under multiphase and partially saturated conditions and evaluate its transferability to other geological formations.

Chapter 4

Field-Scale CO_2 Storage Simulation Using a Stochastic Rock Physics Model

4.1 Overview

Chapter 2 developed a deep-learning workflow for upscaling permeability from pore-scale flow simulations to centimeter-scale core CT images. Chapter 3 extended the workflow to the simultaneous prediction of multiple rock properties and formulated a stochastic relationship for each property from its conditional distribution at a given porosity. This chapter applies the S-RPM formulated for permeability to a field-scale CO_2 storage simulation of the Johansen aquifer. The deterministic KC model, the J7-based S-RPM, and amplified S-RPM cases are compared under identical reservoir-simulation conditions.

The porosity field is taken from the NPD5 Johansen dataset Eigestad et al. (2009). For the KC-based cases, absolute permeability is generated from the NPD5 porosity field using deterministic and stochastic KC relationships. The deterministic KC model assigns a single permeability value to each porosity and defines the RPM baseline. For the S-RPM formulated for permeability, a spatially correlated Gaussian term is added to $\log_{10} k$ around the deterministic KC mean. The porosity-dependent standard deviation $\sigma_k(\phi)$ is fitted to the permeability variation observed at a given porosity in the J7 DNS dataset. In the sensitivity analysis, σ_{scale} multiplies $\sigma_k(\phi)$ and controls the width of the permeability distribution around the deterministic KC mean. Increasing σ_{scale} increases the contrast between high- and

low-permeability cells. The correlation length ℓ controls the distance over which high or low permeability values remain spatially correlated, and increasing ℓ enlarges the corresponding regions. These two parameters are varied to evaluate how the magnitude and spatial scale of permeability variation affect plume migration and the partitioning of V_{trap}, V_{res}, and V_{mob}.

The deterministic and stochastic KC relationships are applied to the NPD5 porosity field, and the resulting three-dimensional permeability fields are averaged along each vertical column to construct the VE top-surface model. The spatially correlated Gaussian fields are generated using the spectral randomization method. Each permeability field is then used in the same Johansen injection and post-injection simulation. The grid, porosity, fluid properties, well conditions, boundary conditions, and time schedule are fixed across the cases, so differences in the calculated plume behavior arise from the permeability fields. Two S-RPM settings are examined. The J7-based S-RPM case uses the fitted function $\sigma_k(\phi)$ with $\sigma_{\mathrm{scale}} = 1$ and $\ell = [2, 1, 1]$ cells. The amplified cases vary σ_{scale} and ℓ to examine how the magnitude and spatial scale of permeability variation affect the CO_2 volume partition.

The simulations use the NPD5 Johansen aquifer model distributed with the MRST `co2lab` module Eigestad et al. (2009); Nilsen et al. (2015). CO_2 is injected through a single well for 100 years, followed by 400 years of post-injection migration. The incompressible VE model solves the pressure equation on the top-surface grid and represents the CO_2 plume by its current height h and historical maximum height $h_{\max}$ in each vertical column. The CO_2 inventory is divided into the plume-retreat immobilized volume V_{trap}, the residual volume within the current plume region V_{res}, and the mobile volume within the current plume region V_{mob}. The plume-retreat immobilized volume is computed from the difference between $h_{\max}$ and h. Accordingly, V_{trap} is the residual CO_2 left behind in the region where the plume has retreated. The same VE fluid and trapping formulation is used in all cases, allowing the effects of the permeability fields on these three volumes to be compared directly. All stochastic fields are generated using the fixed random seed 42. Each $(\sigma_{\mathrm{scale}}, \ell)$ pair is therefore represented by one permeability realization.

The J7-based S-RPM produces nearly the same plume and volume histories as the deterministic KC model because the fitted permeability variation is small and is further reduced by VE column averaging. Changing σ_{scale} and ℓ changes the permeability contrast and spatial organization, which modifies the plume path and the partitioning between plume-retreat immobilized and mobile CO_2. The response to ℓ is non-monotonic. At $\ell = 1$ and 3 cells, the plume reaches both the left-side

plume branch and the right-side fault-adjacent branch after crossing many short-range permeability regions. At $\ell = 10$ cells, increasing σ_{scale} keeps the current plume concentrated near the left-side branch close to the injector and reduces spreading toward the right-side fault-adjacent branch. The overlap between h and $h_{\max}$ therefore increases, and the plume-retreat thickness $h_{\max} - h$ decreases. At $\ell = 30$ cells, broad high- and low-permeability regions intersect both branches over a larger area, and the plume again reaches both sides of the model. These results show how the magnitude and spatial scale of permeability variation control the CO_2 volume partition in the present VE model. The permeability-dependent plume response is compared with the Johansen simulations of Eigestad et al., in which different permeability representations modified the CO_2 front and swept region. This comparison is used as an external numerical check that the present simulations reproduce the permeability-dependent plume response reported for the Johansen FormationEigestad et al. (2009).

4.2 Permeability Formulation for the RPM and S-RPM

4.2.1 Deterministic and Stochastic KC Models

For a granular porous medium, the classical KC relation is Carman (1937)

$$k = \frac{d_p^2}{180\tau^2} \frac{\phi^3}{(1-\phi)^2}, \tag{4.1}$$

where d_p is a characteristic grain diameter, τ is tortuosity, ϕ is porosity, and k is absolute permeability. To represent the J7 data, both the intercept and the exponent of the KC relationship are fitted. The fitted relationship is written in logarithmic form as

$$\log_{10} k_{\text{mD}} = a + b \log_{10} \left[\frac{\phi^3}{(1-\phi)^2} \right], \tag{4.2}$$

where k_{mD} is permeability expressed in millidarcy. Least-squares fitting to the J7 dataset gives

$$a = 4.4764, \qquad b = 0.9208. \tag{4.3}$$

Applying Eq. (4.2) to the local porosity of each active cell defines the deterministic RPM, in which one permeability value is assigned to each porosity.

The deterministic logarithmic mean is defined as

$$\mu_k(\phi) = a + b \log_{10}\left[\frac{\phi^3}{(1-\phi)^2}\right]. \tag{4.4}$$

The deterministic permeability field is written as

$$\log_{10} k_{\mathrm{mD}}(\mathbf{x}) = \mu_k\big(\phi(\mathbf{x})\big), \tag{4.5}$$

where $\mathbf{x}$ denotes the spatial position of a reservoir cell. This field defines the deterministic KC case used in the reservoir simulations.

The S-RPM formulated for permeability represents the variation in k at a given porosity by adding a stochastic term to the deterministic KC mean. The resulting log-permeability relationship is

$$\log_{10} k_{\mathrm{mD}}(\mathbf{x}) = \mu_k\big(\phi(\mathbf{x})\big) + \sigma_k\big(\phi(\mathbf{x})\big) G(\mathbf{x}), \tag{4.6}$$

where $G(\mathbf{x})$ is a spatially correlated Gaussian random field with zero mean and unit variance. The porosity-dependent standard deviation is

$$\sigma_k(\phi) = \max\left(s_0 + s_1\phi, 0.01\right), \qquad s_0 = 0.1602, \quad s_1 = -0.3914. \tag{4.7}$$

A lower bound of 0.01 is applied to maintain a positive standard deviation over the fitted porosity range. The deterministic KC relation defines the conditional mean, whereas $\sigma_k(\phi)G(\mathbf{x})$ represents the permeability variation around that mean.

The coefficients in Eqs. (4.3) and (4.7) were fitted to the J7 DNS dataset. The dataset contains 640,000 porosity–permeability pairs on a $100 \times 80 \times 80$ Cartesian grid. Across the fitted dataset, the standard deviation of $\log_{10} k - \mu_k(\phi)$ is 0.0619. This value corresponds to a multiplicative variation of approximately 15% in permeability at one standard deviation.

For the sensitivity analysis, the stochastic term is scaled by σ_{scale}:

$$\log_{10} k_{\mathrm{mD}}(\mathbf{x}) = \mu_k\big(\phi(\mathbf{x})\big) + \sigma_{\mathrm{scale}}\sigma_k\big(\phi(\mathbf{x})\big) G_\ell(\mathbf{x}). \tag{4.8}$$

Here, $G_\ell(\mathbf{x})$ denotes a spatially correlated Gaussian random field with correlation length ℓ. The J7-based S-RPM case uses $\sigma_{\mathrm{scale}} = 1$. Setting $\sigma_{\mathrm{scale}} = 0$ recovers the deterministic KC model. Values of $\sigma_{\mathrm{scale}} > 1$ broaden the permeability distribution around the KC mean and are used to quantify the resulting change in the simulated

CO_2 inventory.

4.2.2 Spatial Correlation and VE Permeability Assembly

The spatial covariance of $G_\ell(\mathbf{x})$ is

$$\mathrm{Cov}\left[G_\ell(\mathbf{x}), G_\ell(\mathbf{x}')\right] = \exp\left\{-\pi\left[\frac{(x_1 - x_1')^2}{\ell_1^2} + \frac{(x_2 - x_2')^2}{\ell_2^2} + \frac{(x_3 - x_3')^2}{\ell_3^2}\right]\right\}, \quad (4.9)$$

where $\mathbf{x} = (i, j, k)$ denotes the grid-cell index and ℓ_1, ℓ_2, and ℓ_3 are correlation lengths expressed in numbers of cells. The Gaussian field is generated using the spectral randomization method. For the Gaussian covariance in Eq. (4.9), a zero-mean, unit-variance field is approximated by a finite sum of random cosine modes:

$$G_\ell(\mathbf{x}) \approx \sqrt{\frac{2}{N_{\mathrm{modes}}}} \sum_{m=1}^{N_{\mathrm{modes}}} \cos\left(\mathbf{k}_m \cdot \mathbf{x} + \varphi_m\right), \quad (4.10)$$

where $\varphi_m \sim \mathcal{U}(0, 2\pi)$ and each component of the wave vector is sampled independently as

$$k_{m,j} = \frac{\sqrt{2}}{\ell_j}\varepsilon_{m,j}, \qquad \varepsilon_{m,j} \sim \mathcal{N}(0,1), \qquad j = 1, 2, 3. \quad (4.11)$$

The simulations use $N_{\mathrm{modes}} = 1000$, for which the generated field variance remains within approximately 5% of unity in the reported cases. The J7-based S-RPM case retains the numerical setting $\ell = [2, 1, 1]$ used for the original random-field generation. The sensitivity cases use isotropic correlation lengths $\ell_1 = \ell_2 = \ell_3 = \ell$, with $\ell \in \{1, 3, 10, 30\}$ cells, and vary σ_{scale} independently.

The permeability tensor assigned to each cell is

$$\mathbf{K}(\mathbf{x}) = \mathrm{diag}\left(k(\mathbf{x}),\, k(\mathbf{x}),\, 0.1k(\mathbf{x})\right). \quad (4.12)$$

The ratio $k_v/k_h = 0.1$ is used in all reservoir simulations. The three-dimensional rock properties are averaged arithmetically along each vertical column using MRST's `averageRock` routine to construct the VE top-surface model.

4.3 Vertical-Equilibrium Reservoir Simulation Model

4.3.1 Governing Equations and Numerical Solution

The reservoir-scale problem is modeled as immiscible two-phase flow of brine and CO_2 in a porous medium. For phase $\alpha \in \{w, c\}$, where w denotes brine and c denotes CO_2, mass conservation is

$$\frac{\partial(\phi\rho_\alpha S_\alpha)}{\partial t} + \nabla \cdot (\rho_\alpha \mathbf{u}_\alpha) = q_\alpha, \tag{4.13}$$

where the Darcy velocity is

$$\mathbf{u}_\alpha = -\frac{k_{r\alpha}(S_\alpha)}{\mu_\alpha}\mathbf{K}\left(\nabla p_\alpha - \rho_\alpha \mathbf{g}\right). \tag{4.14}$$

Here, S_α is saturation, ρ_α is density, μ_α is viscosity, $k_{r\alpha}$ is relative permeability, p_α is pressure, q_α is the source term, $\mathbf{K}$ is the absolute permeability tensor, and $\mathbf{g}$ is the gravitational acceleration vector. The phase saturations satisfy $S_w + S_c = 1$.

Pore-scale capillary pressure is neglected in the present VE model. The Johansen aquifer is laterally extensive relative to its thickness, and CO_2 is less dense than brine. The VE formulation assumes instantaneous vertical segregation of CO_2 and brine within each column. The current plume thickness is represented by $h(\mathbf{x}_\parallel, t)$, measured downward from the caprock, where $\mathbf{x}_\parallel$ denotes position on the top-surface grid. The historical maximum,

$$h_{\max}(\mathbf{x}_\parallel, t) = \max_{0 \le \tau \le t} h(\mathbf{x}_\parallel, \tau), \tag{4.15}$$

records the maximum plume thickness reached at each top-surface cell. Vertical integration of the CO_2 mass balance gives

$$\frac{\partial}{\partial t}\left[\bar{\phi}(\mathbf{x}_\parallel)\bar{S}_c(h, h_{\max})\right] + \nabla_\parallel \cdot \mathbf{u}_{\parallel,c}(h) = q_c(\mathbf{x}_\parallel, t), \tag{4.16}$$

where $\bar{\phi}$ is the vertically averaged porosity, $\bar{S}_c(h, h_{\max})$ is the column-averaged CO_2 saturation, and $\mathbf{u}_{\parallel,c}$ is the lateral VE CO_2 flux. The column-averaged porosity and horizontal permeability are

$$\bar{\phi}(\mathbf{x}_\parallel) = \frac{1}{H(\mathbf{x}_\parallel)}\int_0^{H(\mathbf{x}_\parallel)} \phi(\mathbf{x}_\parallel, z)\, dz, \qquad \bar{\mathbf{K}}_\parallel(\mathbf{x}_\parallel) = \frac{1}{H(\mathbf{x}_\parallel)}\int_0^{H(\mathbf{x}_\parallel)} \mathbf{K}_\parallel(\mathbf{x}_\parallel, z)\, dz, \tag{4.17}$$

where $H(\mathbf{x}_\parallel)$ is the local aquifer thickness. This arithmetic column averaging reduces short-range vertical permeability variation before the VE flow calculation.

The lateral VE flux is driven by the top-surface pressure gradient and buoyancy. The lateral CO_2 flux is

$$\mathbf{u}_{\parallel,c} = -\lambda_c(h)\bar{\mathbf{K}}_\parallel \left[\nabla_\parallel p_t + (\rho_c - \rho_w) g \nabla_\parallel h\right], \tag{4.18}$$

where p_t is the top-surface pressure, $g = |\mathbf{g}|$, and $\lambda_c(h)$ is the vertically integrated CO_2 mobility. The primary VE variables are the top-surface pressure p_t, the current plume thickness h, and the historical maximum thickness $h_{\max}$.

Both phases are treated as incompressible. The CO_2 and brine densities and viscosities are constant. Pressure acts as a Lagrange multiplier that enforces the divergence constraint on the total velocity in the incompressible two-phase system. At each global time step, the VE pressure equation is solved on the top-surface grid. The top-surface pressure equation is written as

$$-\nabla_\parallel \cdot \left[\bar{\mathbf{K}}_\parallel(\mathbf{x}_\parallel)\lambda_{\text{tot}}(h)\left(\nabla_\parallel p_t - \rho_{\text{eff}}(h) g \nabla_\parallel z_t - \mathbf{g}_\perp(h)\right)\right] = q(\mathbf{x}_\parallel, t), \tag{4.19}$$

where z_t is the caprock depth, $g = |\mathbf{g}|$, $\lambda_{\text{tot}}(h) = \lambda_c(h) + \lambda_w(h)$ is the total mobility, and $\rho_{\text{eff}}(h)$ is the mobility-weighted density. The term involving $\nabla_\parallel z_t$ represents the top-surface gravity contribution associated with caprock slope. The term $\mathbf{g}_\perp(h)$ is the hydrostatic interface contribution associated with the CO_2–brine density difference. The vertically integrated phase mobilities used in the sharp-interface VE model are

$$\lambda_c(h) = \frac{k_{rc,m} h}{\mu_c}, \qquad \lambda_w(h) = \frac{k_{rw,m}\left[H - h - s_{rc}(h_{\max} - h)\right]}{\mu_w}, \tag{4.20}$$

where s_{rc} is the residual CO_2 saturation and $k_{rw,m}$ and $k_{rc,m}$ are endpoint relative permeabilities.

Summing the phase mass balances under incompressibility and $S_w + S_c = 1$ yields an elliptic pressure equation. The transient part of the problem remains in the hyperbolic height equation, Eq. (4.16). Pressure and transport are solved sequentially. The pressure equation is solved by the mixed-hybrid mimetic finite-difference method implemented in MRST. The height equation is then advanced by an explicit upwind VE transport solver with internal CFL sub-stepping. The reported simulations use the accelerated VE transport kernel when available.

4.3.2 CO_2 Volume Accounting

The reported CO_2 volumes are computed from the VE state variables h and $h_{\max}$. The pore volume per unit column height is defined as

$$\mathrm{pv}(\mathbf{x}_\parallel) = \bar{\phi}(\mathbf{x}_\parallel) A(\mathbf{x}_\parallel), \tag{4.21}$$

where pv is the pore volume per unit column height and A is the top-surface cell area. The summation is taken over all active VE columns. The three reported volumes represent the plume-retreat immobilized CO_2, the residual CO_2 within the current plume region, and the mobile CO_2 within the current plume region:

$$V_{\mathrm{trap}} = s_{rc} \sum_{\mathbf{x}_\parallel} \left[h_{\max}(\mathbf{x}_\parallel, t) - h(\mathbf{x}_\parallel, t) \right] \mathrm{pv}(\mathbf{x}_\parallel), \tag{4.22}$$

$$V_{\mathrm{res}} = s_{rc} \sum_{\mathbf{x}_\parallel} h(\mathbf{x}_\parallel, t) \mathrm{pv}(\mathbf{x}_\parallel), \tag{4.23}$$

$$V_{\mathrm{mob}} = (1 - s_w - s_{rc}) \sum_{\mathbf{x}_\parallel} h(\mathbf{x}_\parallel, t) \mathrm{pv}(\mathbf{x}_\parallel). \tag{4.24}$$

Here, s_{rc} is the residual CO_2 saturation and s_w is the residual water saturation. The total in-place CO_2 volume under reservoir conditions is

$$V_{\mathrm{CO_2}}^{\mathrm{tot}} = V_{\mathrm{trap}} + V_{\mathrm{res}} + V_{\mathrm{mob}} = \sum_{\mathbf{x}_\parallel} \mathrm{pv}(\mathbf{x}_\parallel) \left[s_{rc} h_{\max}(\mathbf{x}_\parallel, t) + (1 - s_w - s_{rc}) h(\mathbf{x}_\parallel, t) \right]. \tag{4.25}$$

During injection, $h = h_{\max}$ in the invaded columns, giving $V_{\mathrm{trap}} = 0$. After injection stops, parts of the plume retreat, $h < h_{\max}$, and the difference $h_{\max} - h$ contributes to V_{trap}. Accordingly, V_{trap} denotes the plume-retreat immobilized volume. All reported volumes are expressed under reservoir conditions.

4.4 Johansen Aquifer Simulation Setup

The simulations use the NPD5 Johansen corner-point model distributed with the MRST `co2lab` module Eigestad et al. (2009); Nilsen et al. (2015). The full three-dimensional grid has dimensions $100 \times 100 \times 11$, corresponding to 110,000 grid cells, of which 29,128 are active. The VE top-surface grid contains one top-surface cell for each active vertical column. Porosity is read from the NPD5 porosity dataset and mapped to the active cells. The deterministic KC, J7-based S-RPM, and amplified

S-RPM permeability fields are generated from the NPD5 porosity field using the relationships in Section 4.2. Thus, the three representative KC-based cases differ in their absolute-permeability fields while sharing the remaining model inputs. Table 4.1 summarizes the horizontal-permeability statistics of the three representative KC-based fields.

Table 4.1: Horizontal permeability statistics over the 29,128 active three-dimensional cells for the three representative KC-based permeability fields.

Case	Minimum (mD)	Median (mD)	Maximum (mD)
Deterministic KC	74.9	769.9	1,856
J7-based S-RPM	51.7	776.4	2,576
Amplified S-RPM ($\sigma_{\text{scale}} = 5$, $\ell = [10, 10, 10]$)	≈ 10	≈ 780	$\approx 15{,}000$

The simulations use the following constant fluid properties. The CO_2 viscosity is 0.05664 cP, the brine viscosity is 0.30860 cP, the CO_2 density is 686.54 kg m^{-3}, and the brine density is 975.86 kg m^{-3}. The residual CO_2 saturation is $s_{rc} = 0.20$, the residual water saturation is $s_w = 0.10$, and the endpoint relative permeabilities are $k_{rw,m} = 0.2142$ and $k_{rc,m} = 0.85$. A single injector is completed in cell $(i, j, k) = (51, 51, 6)$. The injection rate is

$$Q_{\text{inj}} = 1.4 \times 10^4 \text{ m}^3 \text{ day}^{-1}, \tag{4.26}$$

which corresponds to approximately 5.11×10^6 m^3 yr^{-1}. Injection continues for 100 years, giving a cumulative injected volume of approximately 5.11×10^8 m^3 under the model conditions. A 400-year post-injection period gives a total simulation time of 500 years. The outer time step is 2 years during injection and 5 years during post-injection migration.

Hydrostatic pressure is imposed on the open lateral boundaries during injection. The boundary pressure is given by

$$p \,\big|_{\partial\Omega_{\text{open}}} = \rho_w |\mathbf{g}| z_{\text{face}}, \tag{4.27}$$

where z_{face} is the depth of the boundary face. This condition allows brine to leave the domain as CO_2 is injected. At 100 years, the injector and open pressure boundaries are deactivated. The post-injection calculation is driven by gravity under closed lateral boundaries.

The initial plume height and historical maximum plume height are zero every-

where. The gravitational acceleration is 9.8066 m s^{-2}. The random seed is fixed at 42. Each run produces 131 saved states: 50 injection-phase states and 80 migration-phase states, plus the initial state.

4.5 Deterministic, Calibrated, and Amplified Stochastic Cases

4.5.1 Representative KC-Based Cases

The deterministic KC case is obtained by setting $G(\mathbf{x}) \equiv 0$ in Eq. (4.6). The horizontal permeability distribution has a median of 770 mD, with $P_{10} = 460$ mD and $P_{90} = 1145$ mD. The permeability ranges from 74.9 to 1856 mD.

At 500 years, the deterministic KC case gives

$$V_{\text{trap}} = 1.183 \times 10^8 \text{ m}^3, \qquad V_{\text{res}} = 8.73 \times 10^7 \text{ m}^3, \qquad V_{\text{mob}} = 3.06 \times 10^8 \text{ m}^3. \quad (4.28)$$

The cumulative injected volume is 5.11×10^8 m^3, and the final mass-balance error is less than 0.2%. This result is used as the baseline for the S-RPM cases.

The J7-based S-RPM case uses $\sigma_{\text{scale}} = 1$ and $\ell = [2, 1, 1]$ cells. The generated field has mean$(G) = 0.0003$ and std$(G) = 0.9832$. The standard deviation of the pointwise logarithmic ratio $\log_{10}(k_{\text{S-RPM}}/k_{\text{KC}})$ is 0.073, corresponding to an approximately 18% local permeability variation at one standard deviation.

Figures 4.1 and 4.2 compare the deterministic KC and J7-based S-RPM permeability fields. The J7-based S-RPM adds short-range variation around the deterministic field, while the two $\log_{10} k$ distributions remain close at the plotted scale.

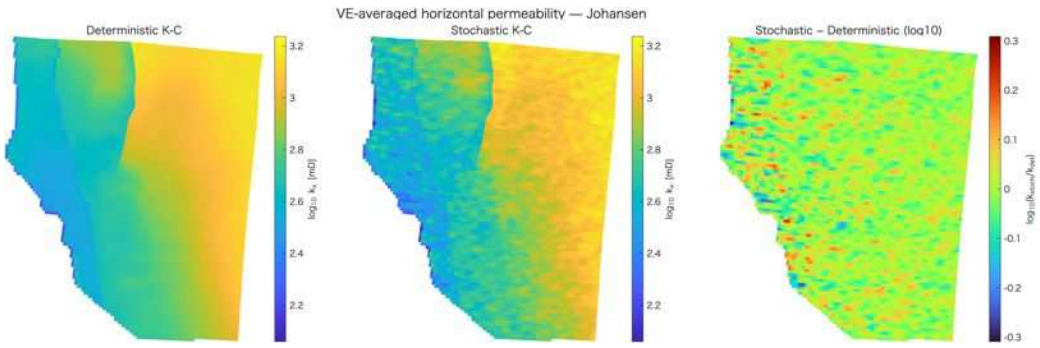


Figure 4.1: VE-averaged horizontal permeability $\log_{10} k_x$ for the deterministic KC model, the J7-based S-RPM, and their pointwise log-ratio. The stochastic variation is small relative to the porosity-controlled KC trend.

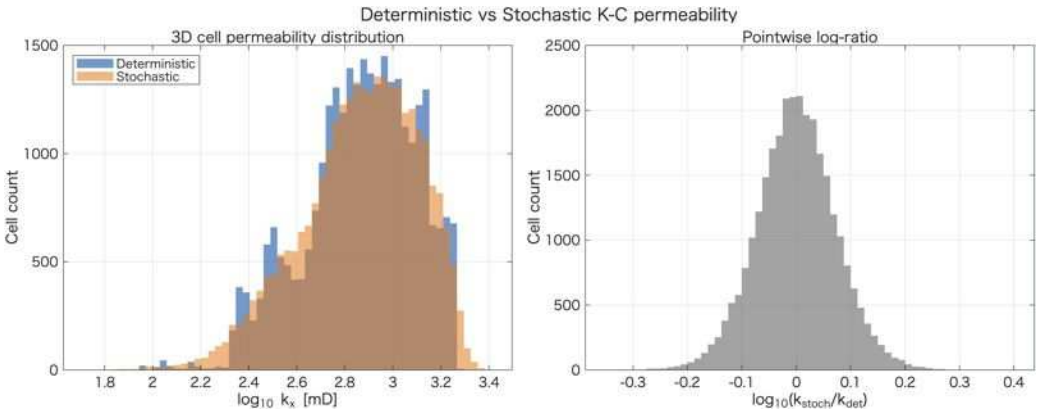


Figure 4.2: Histograms of $\log_{10} k_x$ over the full three-dimensional grid and of the pointwise log-ratio $\log_{10}(k_{\text{S-RPM}}/k_{\text{KC}})$ for the J7-based S-RPM.

The final values of V_{trap}, V_{res}, and V_{mob} agree with the deterministic KC case to three significant figures. The reasons for this agreement are discussed in Section 4.7.

The representative amplified case uses $\sigma_{\text{scale}} = 5$ and $\ell = [10, 10, 10]$ cells. The standard deviation of $\log_{10} k - \mu_k(\phi)$ becomes 0.358, corresponding to a factor of approximately 2.3 in permeability at one standard deviation. This parameter set is used as a representative amplified case before the full sensitivity analysis.

Table 4.2 compares the final volume components. Relative to the deterministic KC case, V_{trap} decreases by 32%, from 118.3×10^6 to 80.5×10^6 m^3, while V_{res} and V_{mob} increase. V_{mob} increases from 305.7×10^6 to 335.1×10^6 m^3, and V_{res} increases from 87.3×10^6 to 95.8×10^6 m^3.

Table 4.2: Final CO_2 volume components at 500 years for the deterministic KC case and the representative amplified S-RPM case.

Inventory component (10^6 m^3)	Deterministic KC	Representative amplified S-RPM
V_{trap}	118.3	80.5
V_{res}	87.3	95.8
V_{mob}	305.7	335.1

The volume histories are shown in Fig. 4.3, and the final plume maps are shown in Fig. 4.4. The representative amplified permeability map contains connected high-permeability regions near the injector-side plume path. The corresponding plume remains concentrated near these regions at 500 years, resulting in lower V_{trap} and higher V_{res} and V_{mob} than the deterministic KC case.tugi

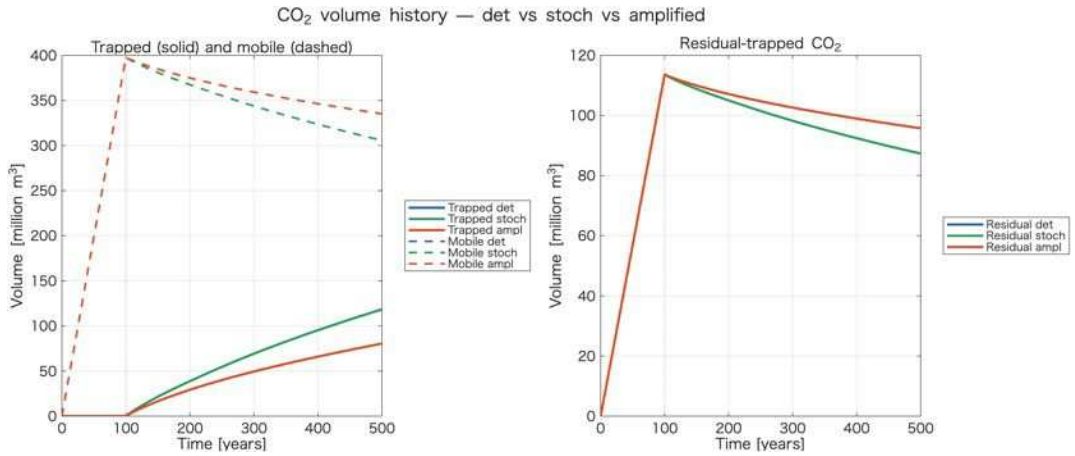


Figure 4.3: CO_2 volume histories for the deterministic KC, J7-based S-RPM, and representative amplified S-RPM cases. The J7-based S-RPM gives the same volume histories as the deterministic KC case at the plotted scale, whereas the representative amplified S-RPM case gives a lower V_{trap} and a larger V_{mob} at 500 years.

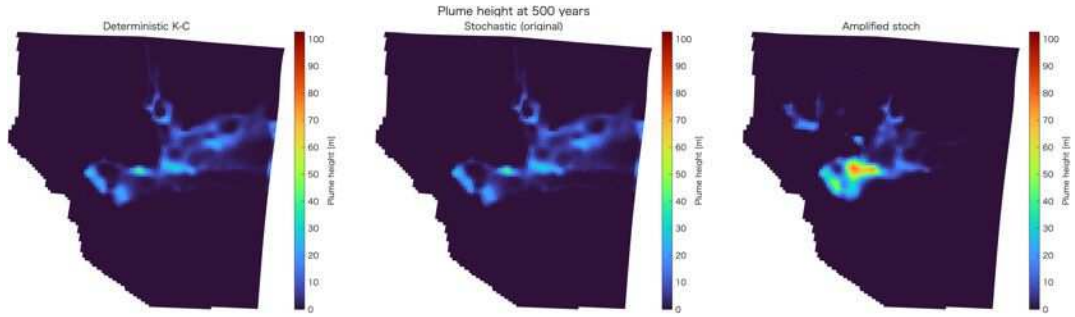


Figure 4.4: Current plume height h at 500 years for the deterministic KC, J7-based S-RPM, and representative amplified S-RPM cases. The representative amplified S-RPM case retains a localized current plume near the injector-side branch.

Eigestad et al. showed that different permeability representations modify the CO_2 front and swept region in the Johansen Formation. Figure 4.5 compares their permeability-representation cases with the present stochastic KC simulations at $\ell = 10$ cells. In the present model, changing σ_{scale} modifies the spatial arrangement of permeability and changes the current plume geometry. This comparison provides an external numerical check of the permeability-dependent plume response in the present simulations.

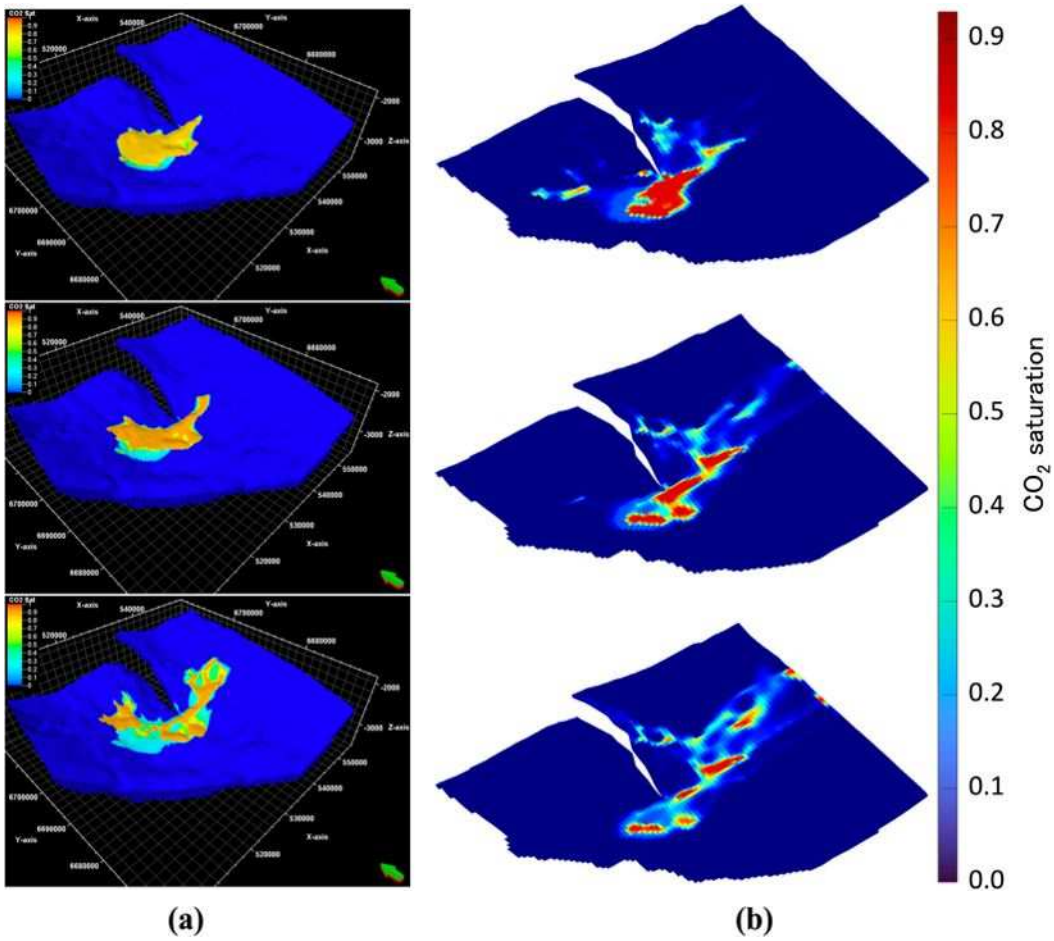


Figure 4.5: Permeability-dependent CO_2 plume patterns in the Johansen Formation. (a) Saturation distributions from Eigestad et al. for heterogeneous permeability, homogeneous 500 mD, and homogeneous 1000 mD cases. (b) Present VE simulations using the stochastic KC model at $\ell = 10$ cells with different values of σ_{scale}. Each row represents a permeability-sensitivity case within its own simulation model.

4.5.2 Porosity–Permeability Dispersion Relative to DNS Calibration

Figure 4.6 compares the porosity–permeability distributions of the deterministic KC, J7-based S-RPM, and representative amplified S-RPM fields. The deterministic KC field lies on the fitted mean relationship. The J7-based S-RPM forms a band around the KC mean that follows the width of the J7 data. The representative amplified case extends beyond the J7-based uncertainty bands. Table 4.3 quantifies the width of each distribution using the standard deviation of $\log_{10} k - \mu_k(\phi)$.

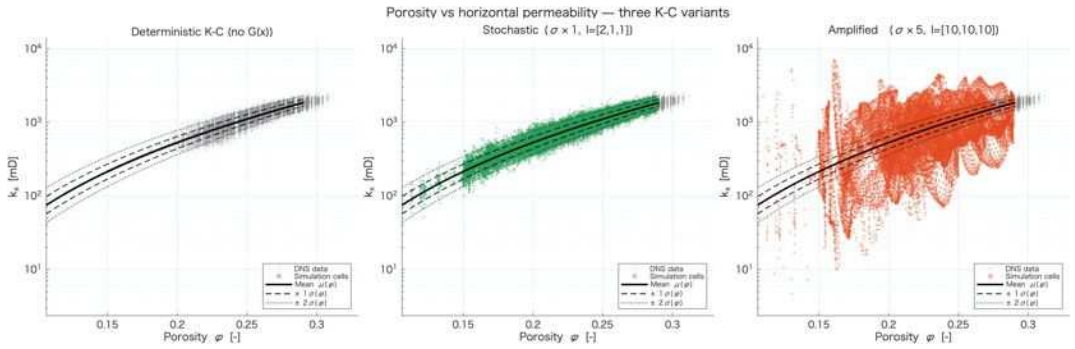


Figure 4.6: Porosity–permeability distributions for the deterministic KC, J7-based S-RPM, and representative amplified S-RPM fields. The J7 DNS reference data and the fitted mean and uncertainty bands are overlaid. The representative amplified case is used as a sensitivity case because its distribution extends beyond the J7-based uncertainty bands.

Table 4.3: Standard deviation of $\log_{10} k - \mu_k(\phi)$ for the J7 DNS data and the generated permeability fields.

Population	Standard deviation of $\log_{10} k - \mu_k(\phi)$
Deterministic RPM	0.0000
J7 DNS dataset	0.0619
J7-based S-RPM	0.0729
Representative amplified S-RPM	0.3583

The standard deviation is 0.0619 for the J7 DNS data and 0.0729 for the generated J7-based S-RPM field. Increasing σ_{scale} broadens the permeability distribution beyond the J7-based uncertainty bands shown in Fig. 4.6. The amplified cases are used to quantify how this broader permeability distribution changes the VE plume and volume inventory.

4.6 Sensitivity to Stochastic Amplitude and Correlation Length

The sensitivity analysis uses

$$\sigma_{\text{scale}} \in \{0, 1, 2, 5, 10\}, \qquad \ell \in \{1, 3, 10, 30\} \text{ cells.} \tag{4.29}$$

These combinations give 20 simulations. For $\sigma_{\text{scale}} = 0$, the stochastic term vanishes and all four values of ℓ reproduce the deterministic KC field. All sensitivity sim-

ulations use the same grid, porosity, fluid properties, injection schedule, boundary conditions, and random seed.

Figure 4.7 shows the VE-averaged horizontal permeability maps for all 20 cases. Rows correspond to $\sigma_{\text{scale}} = 0, 1, 2, 5, 10$. Columns correspond to $\ell = 1, 3, 10, 30$ cells, and the $\sigma_{\text{scale}} = 0$ row is identical across columns. The spatial extent of high- and low-permeability regions increases from $\ell = 1$ to $\ell = 30$ cells.

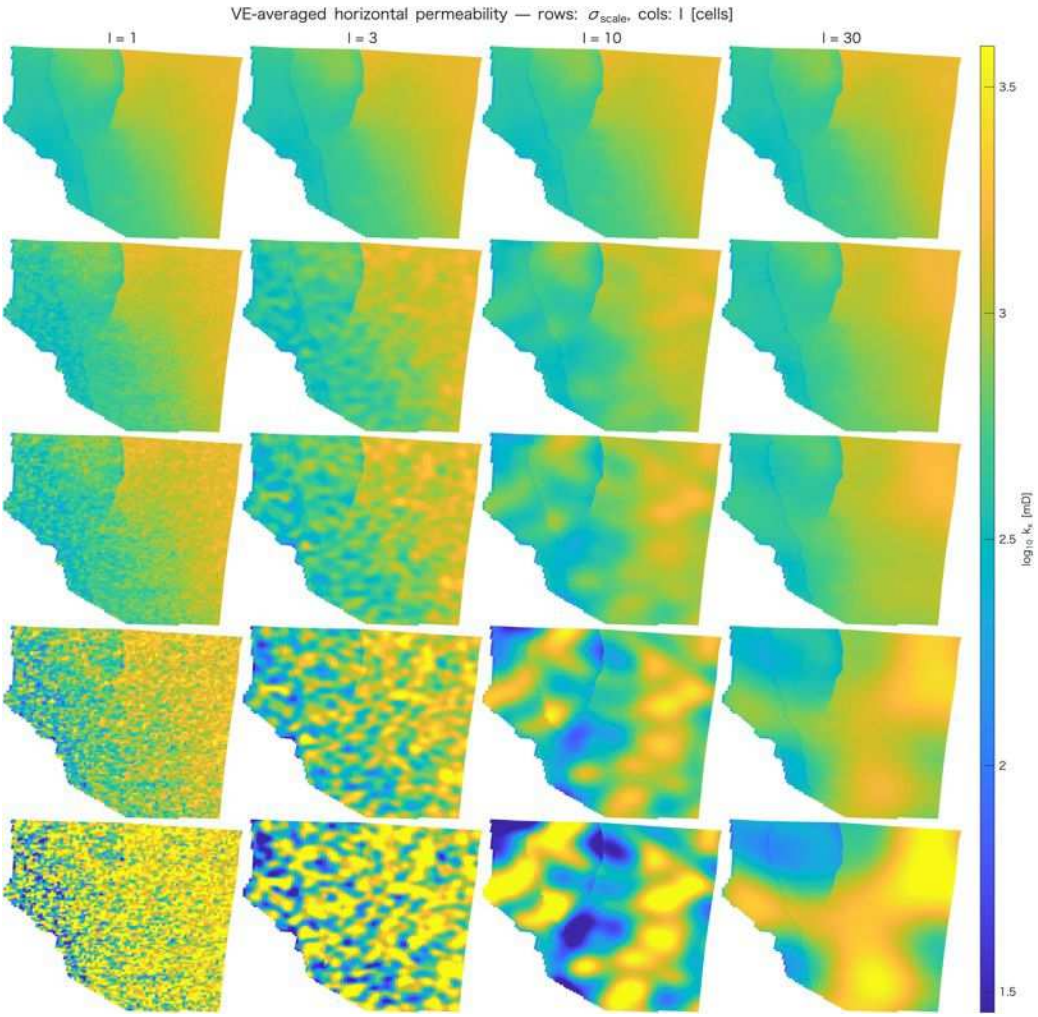


Figure 4.7: VE-averaged horizontal permeability $\log_{10} k_x$ maps for the full sensitivity grid. Rows correspond to $\sigma_{\text{scale}} = 0, 1, 2, 5, 10$, and columns correspond to $\ell = 1, 3, 10, 30$ cells. A common color scale is used for all panels.

Tables 4.4–4.6 list the final CO_2 volumes at 500 years. Among the three inventory components, V_{trap} shows the largest range across the parameter grid. At $\ell = 1$ and $\ell = 3$, increasing σ_{scale} increases V_{trap}. At $\ell = 10$ cells, V_{trap} decreases to 80.45×10^6 m^3 at $\sigma_{\text{scale}} = 5$ and increases to 87.95×10^6 m^3 at $\sigma_{\text{scale}} = 10$. At $\ell = 30$

cells, V_{trap} increases from 118.32×10^6 to 174.87×10^6 m^3 as σ_{scale} increases from 0 to 10.

Table 4.4: Plume-retreat trapped CO_2 volume V_{trap} at 500 years (10^6 m^3).

$\sigma_{\mathrm{scale}} \backslash \ell$	1	3	10	30
0	118.32	118.32	118.32	118.32
1	121.09	121.69	103.59	125.93
2	125.47	125.78	92.66	134.03
5	146.40	140.89	80.45	158.04
10	173.97	160.48	87.95	174.87

Table 4.5: Mobile CO_2 volume at 500 years (10^6 m^3).

$\sigma_{\mathrm{scale}} \backslash \ell$	1	3	10	30
0	305.68	305.68	305.68	305.68
1	303.53	303.06	317.14	299.76
2	300.12	299.88	325.64	293.47
5	283.84	288.13	335.14	274.79
10	262.40	272.89	329.30	261.70

Table 4.6: Residual CO_2 volume at 500 years (10^6 m^3).

$\sigma_{\mathrm{scale}} \backslash \ell$	1	3	10	30
0	87.34	87.34	87.34	87.34
1	86.72	86.59	90.61	85.65
2	85.75	85.68	93.04	83.85
5	81.10	82.32	95.75	78.51
10	74.97	77.97	94.09	74.77

Figures 4.8–4.10 summarize the dependence of the final volume components on σ_{scale} and ℓ. Figure 4.11 shows the corresponding current plume height h at 500 years.

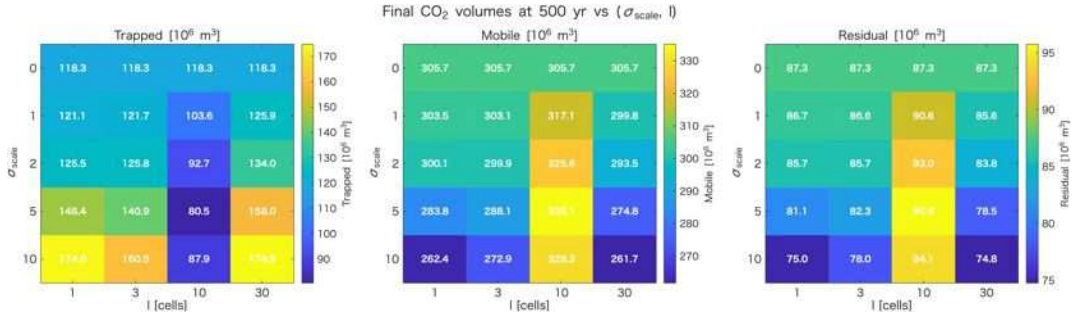


Figure 4.8: Final values of V_{trap}, V_{mob}, and V_{res} at 500 years over the σ_{scale}–ℓ grid. Values inside the panels are in 10^6 m^3.

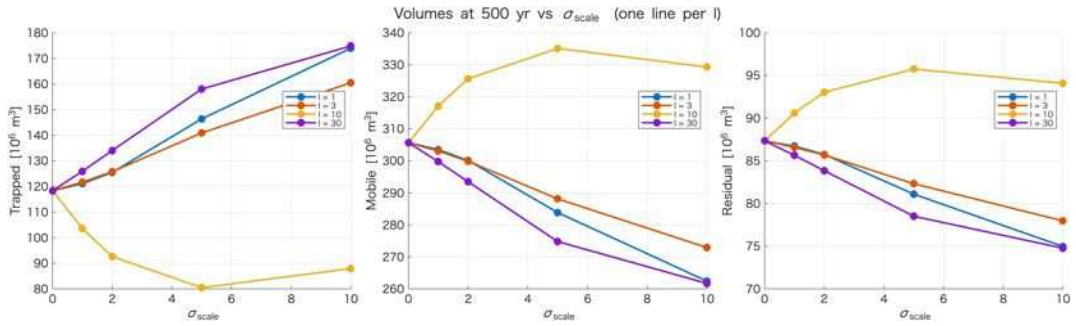


Figure 4.9: Final CO_2 volume components at 500 years as functions of σ_{scale}. Each curve corresponds to one correlation length.

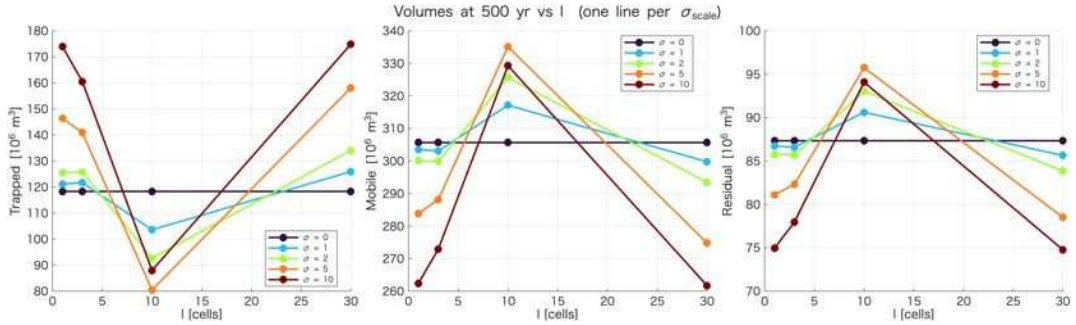


Figure 4.10: Final CO_2 volume components as functions of correlation length, with one curve for each σ_{scale}. The minimum in V_{trap} occurs near $\ell = 10$ cells for the non-zero stochastic amplitudes.

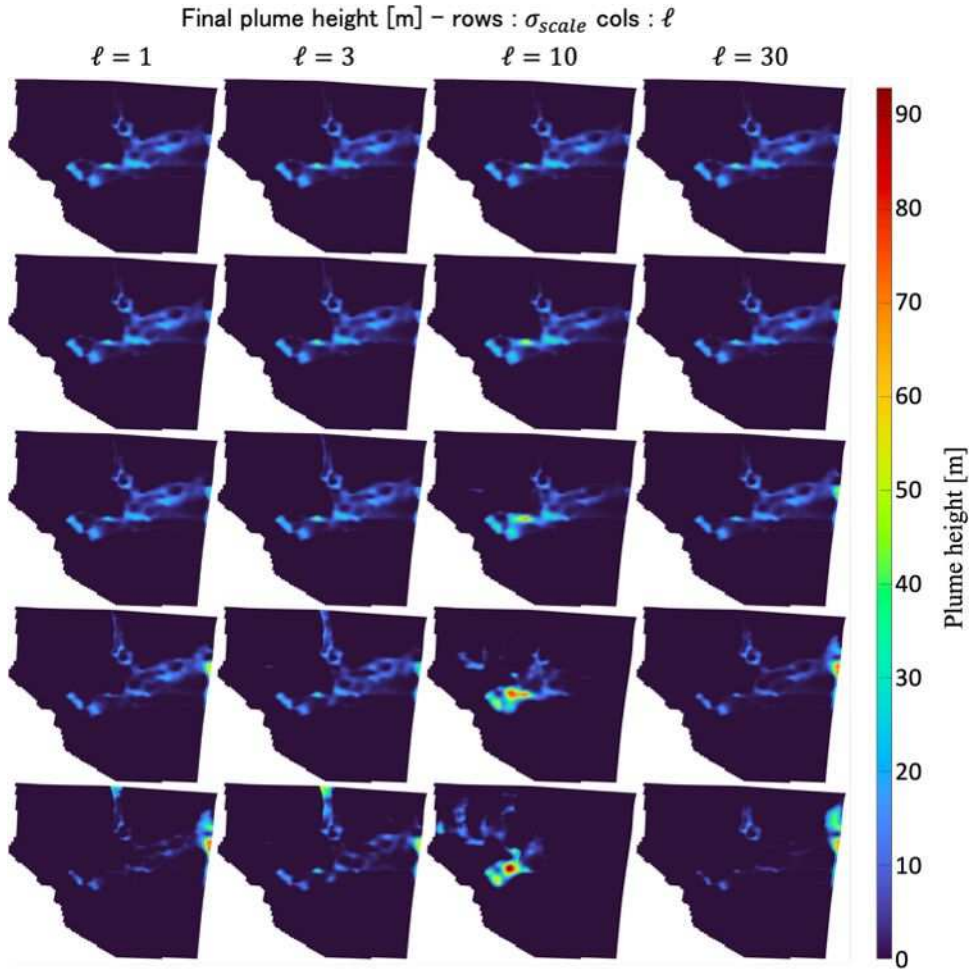


Figure 4.11: Current plume height h at 500 years for the full sensitivity grid. Rows correspond to $\sigma_{\text{scale}} = 0, 1, 2, 5, 10$, and columns correspond to $\ell = 1, 3, 10, 30$ cells.

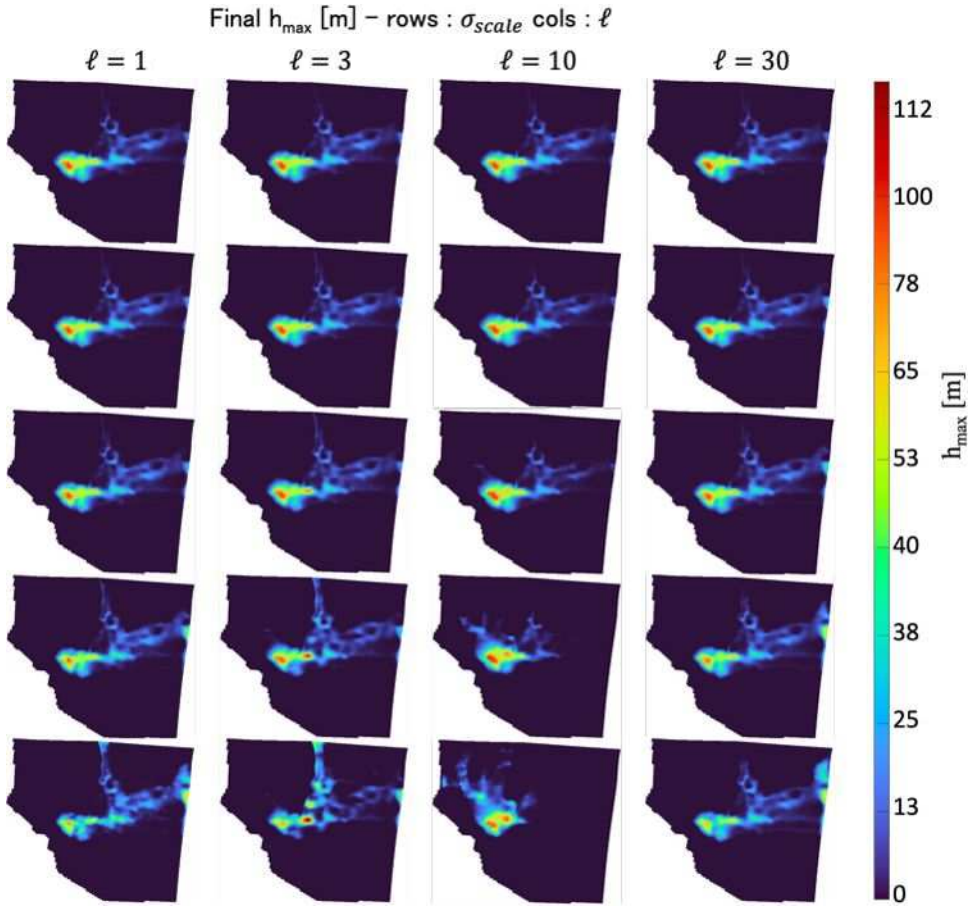


Figure 4.12: Historical maximum plume height $h_{\max}$ at 500 years for the full sensitivity grid. Rows correspond to $\sigma_{\text{scale}} = 0, 1, 2, 5, 10$, and columns correspond to $\ell = 1, 3, 10, 30$ cells.

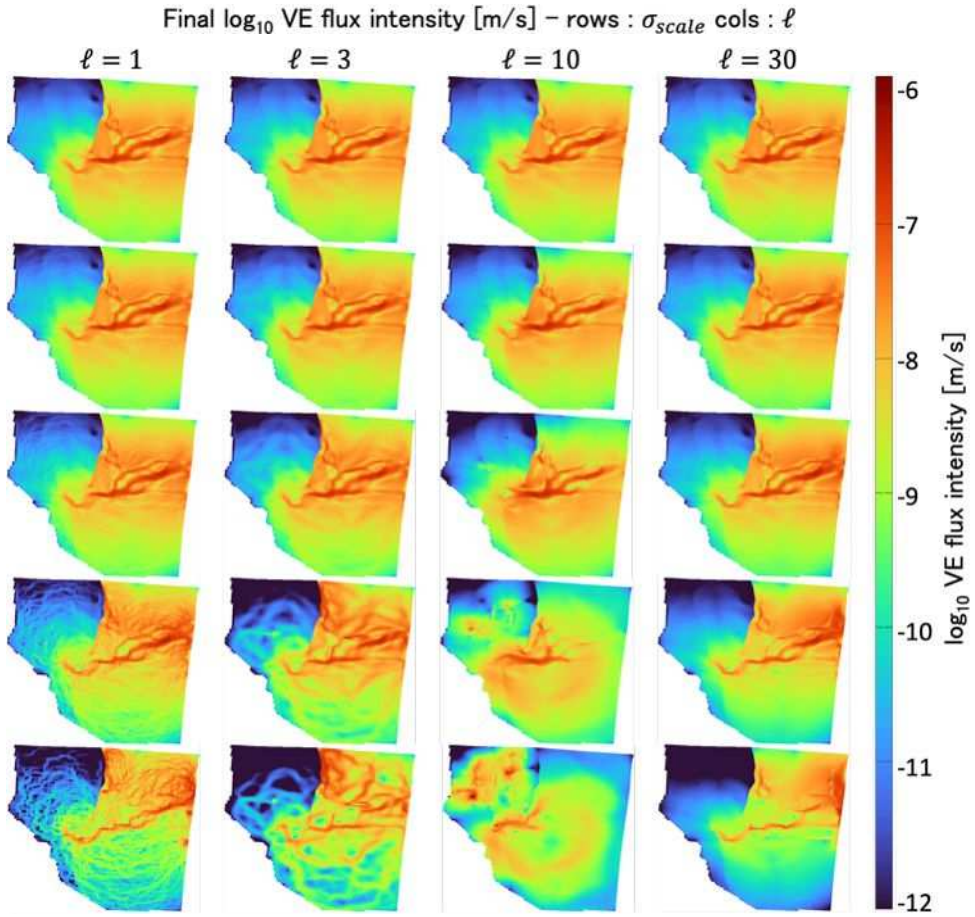


Figure 4.13: Magnitude of the VE total flux at 500 years for the full sensitivity grid. Rows correspond to $\sigma_{\text{scale}} = 0, 1, 2, 5, 10$, and columns correspond to $\ell = 1, 3, 10, 30$ cells.

Figures 4.11 and 4.12 distinguish the current plume footprint from the historical swept footprint. The current plume height h represents the CO_2 plume present at 500 years, whereas the historical maximum height $h_{\max}$ records the region reached before that time. Since V_{trap} is proportional to $h_{\max} - h$ in Eq. (4.22), a smaller $h_{\max} - h$ indicates a larger overlap between the current plume and the historical footprint. For $\ell = 1$ and 3 cells, the current plume and the historical footprint extend through both the left-side plume branch and the right-side fault-adjacent branch. The flow-speed map shows distributed flow over these branches. This produces a broad swept footprint and a larger plume-retreat contribution to V_{trap}. For $\ell = 10$ cells, increasing σ_{scale} concentrates the current plume near the left-side branch close to the injector and reduces spreading toward the right-side fault-adjacent branch. The historical footprint is also spatially limited compared with the $\ell = 1$, 3, and 30

cases. The overlap between h and $h_{\max}$ is therefore large, which reduces $h_{\max} - h$ and gives the minimum V_{trap} in the sensitivity sweep. For $\ell = 30$ cells, broad high- and low-permeability regions intersect both plume branches. The current plume and the historical footprint again extend through both sides of the model, and the plume-retreat contribution to V_{trap} increases. The non-monotonic response with respect to ℓ therefore arises from the spatial relation among permeability organization, the fault-adjacent branch, and the overlap between h and $h_{\max}$. Figure 4.14 shows that the width of the porosity–permeability distribution increases primarily with σ_{scale}. The covariance model prescribes the nominal variance of the Gaussian field, while ℓ changes the spatial scale of high- and low-permeability regions. Cases with similar permeability histograms can produce different plume geometries when the spatial arrangement of permeability differs.

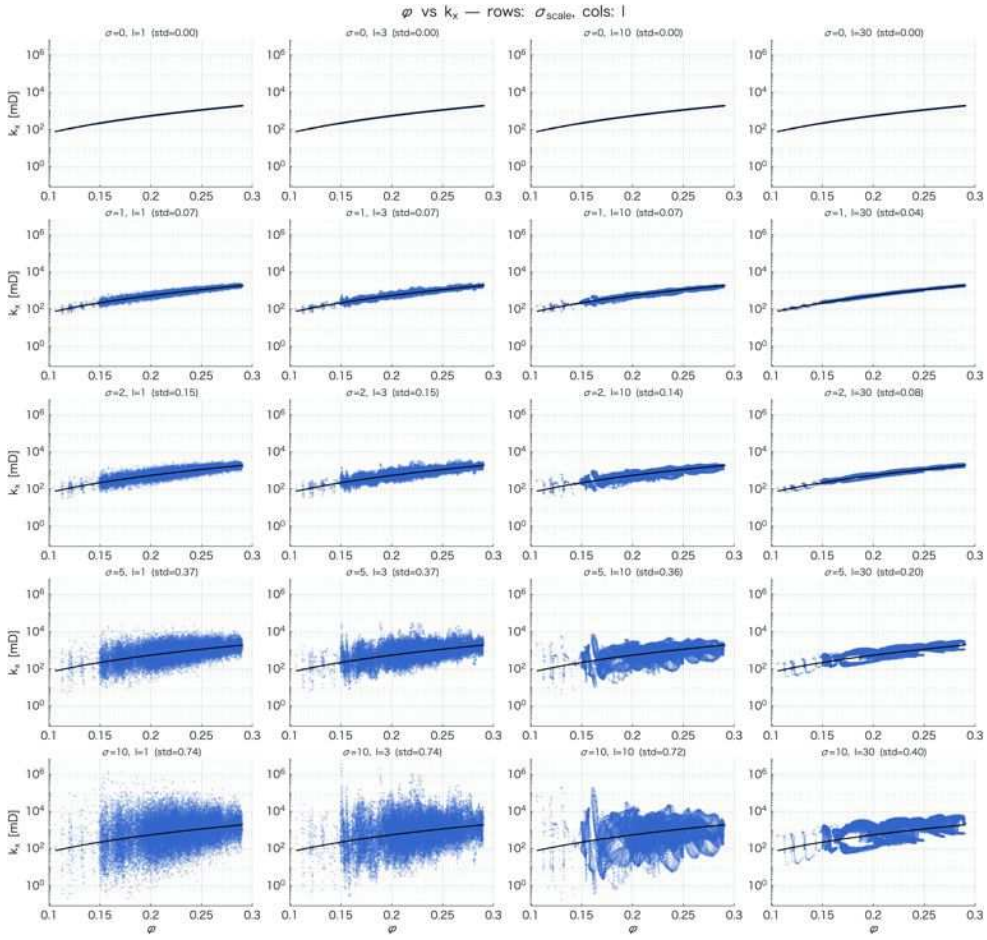


Figure 4.14: Porosity–permeability scatter plots for the full sensitivity grid. The scatter amplitude increases with σ_{scale}, while the marginal distribution is nearly independent of ℓ.

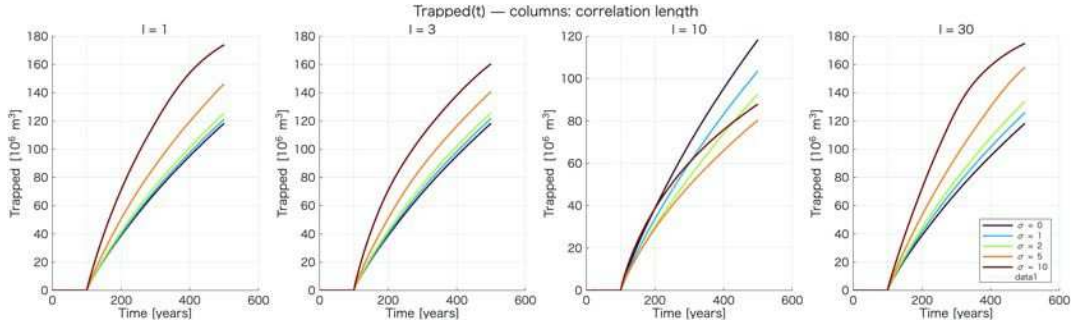


Figure 4.15: Trapped-volume time histories grouped by correlation length and colored by σ_{scale}. Differences in V_{trap} emerge during injection and remain through the 400-year migration period.

The sensitivity results are interpreted by comparing the permeability maps, the current plume height h, the historical maximum height $h_{\max}$, and the flow-speed field. The sensitivity sweep gives the minimum V_{trap} at $\ell = 10$ cells for the non-zero stochastic amplitudes.

For $\ell = 1$ and 3 cells, high- and low-permeability regions occur over short distances. The current plume and the historical footprint spread through both the left- and right-side structural pathways. This produces a broader swept region and a larger plume-retreat wake, consistent with the increase in V_{trap} as σ_{scale} increases.

For $\ell = 10$ cells, increasing σ_{scale} keeps the current plume concentrated near the left-side pathway around the injector and reduces spreading toward the right-side fault pathway. The current plume therefore overlaps strongly with $h_{\max}$, and the wake thickness $h_{\max} - h$ becomes small. This explains why V_{trap} decreases to 80.45×10^6 m^3 at $\sigma_{\text{scale}} = 5$, while both V_{res} and V_{mob} increase to 95.75×10^6 m^3 and 335.14×10^6 m^3, respectively.

For $\ell = 30$ cells, the high- and low-permeability regions are broad and intersect the structural migration pathways over a larger area. The plume again spreads through both sides of the model, producing a broader historical footprint and a larger plume-retreat wake. The response therefore returns toward the behavior of the $\ell = 1$ and 3 cases in terms of V_{trap}. The result remains a single realization generated with seed 42.

4.7 Discussion

The J7-based S-RPM and deterministic KC cases give the same final volume components to three significant figures. This agreement follows from three features of

the calculation. First, the J7 data define a small permeability variation around the KC mean; the generated S-RPM field has a pointwise log-ratio standard deviation of 0.073, corresponding to an approximately 18% local permeability variation. Second, the VE formulation averages permeability along each vertical column, reducing the contribution of short-range vertical variation to the top-surface permeability field. Third, the J7-based field has a short horizontal correlation length relative to the plume footprint, so spatial integration over the plume area reduces its contribution to V_{trap}, V_{res}, and V_{mob}.

The amplified cases show that the effect of permeability variation depends on both σ_{scale} and ℓ. Increasing σ_{scale} broadens the permeability distribution, whereas ℓ controls the spatial arrangement of high- and low-permeability regions. Cases with similar porosity–permeability distributions can therefore produce different plume geometries and different volume partitions. The $\ell = 10$ cases illustrate this effect: the current plume remains concentrated near the left-side branch close to the injector, and the overlap between h and h_{max} reduces the plume-retreat contribution to V_{trap}. In contrast, the $\ell = 1$, 3, and 30 cases spread through both plume branches and produce larger plume-retreat residual volumes.

The comparison with Eigestad et al. provides an external numerical check for the reservoir-scale implementation. Their Johansen simulations showed that changes in permeability representation modify the CO_2 front and swept region. The present amplified S-RPM cases reproduce the same class of permeability-dependent plume response, although the comparison is morphological rather than a one-to-one volume benchmark.

The interpretation is conditioned by the VE formulation and the single-realization stochastic design. The VE model assumes instantaneous vertical segregation and represents the plume by top-surface variables h and h_{max}. CO_2 and brine are treated as incompressible phases, and the volume accounting is based on residual saturation and plume-height history. Each $(\sigma_{\mathrm{scale}}, \ell)$ pair is represented by one permeability realization generated with seed 42. Large-ℓ cases are therefore sensitive to the position of broad high- and low-permeability regions relative to the injector and the fault-adjacent branch. The amplified cases extend beyond the J7 permeability variation and are used to evaluate parameter sensitivity.

4.8 Summary

This chapter applied the S-RPM formulated for permeability to field-scale CO_2 storage simulations of the Johansen aquifer. The deterministic RPM was defined as the fitted KC mean relationship between porosity and permeability. The S-RPM added a porosity-dependent standard deviation fitted to the J7 data and a spatially correlated Gaussian field to the KC mean. The Gaussian fields were generated using the spectral randomization method, and the resulting permeability fields were used in the MRST incompressible VE model.

The J7-based S-RPM reproduced the width of the permeability distribution around the KC mean, with standard deviations of 0.0729 for the generated field and 0.0619 for the J7 data. The final volume components agreed with the deterministic KC case to three significant figures. This agreement results from the small J7-based permeability variation, VE column averaging, and spatial integration over the plume area.

The amplified simulations showed that increasing σ_{scale} and changing ℓ modify the plume path and the partitioning among V_{trap}, V_{res}, and V_{mob}. The sensitivity sweep over $\sigma_{\text{scale}} \in \{0, 1, 2, 5, 10\}$ and $\ell \in \{1, 3, 10, 30\}$ cells produced a non-monotonic dependence of V_{trap} on ℓ. For $\ell = 1$ and 3 cells, V_{trap} increased with σ_{scale}. At $\ell = 10$ cells, increasing σ_{scale} kept the current plume near the injector-side pathway, reduced the wake thickness $h_{\max} - h$, decreased V_{trap} by 32%, and increased both V_{res} and V_{mob} at $\sigma_{\text{scale}} = 5$. At $\ell = 30$ cells, broad high- and low-permeability regions intersected the structural migration pathways over a larger area, and V_{trap} increased for seed 42. The permeability-dependent plume response was also compared with the Johansen simulations of Eigestad et al. Their permeability-representation study showed that changing permeability modifies the CO_2 front and swept region. The present amplified S-RPM simulations reproduce the same class of permeability-dependent plume response, providing an external numerical check of the reservoir-scale implementation.

The chapter quantifies how permeability variation represented by the S-RPM affects field-scale plume migration and CO_2 volume partitioning. For the J7-based permeability variation and the present VE averaging, the deterministic KC and S-RPM volume inventories agree to three significant figures. Within the tested parameter grid, the $\ell = 10$ cases reduce plume-retreat immobilization and increase both residual and mobile CO_2 inventories by keeping the current plume near the injector-side pathway. The workflow transfers the permeability variation obtained

from J7 digital-rock data to reservoir-scale permeability fields and evaluates its effect on the Johansen VE simulation.

Chapter 5

Conclusions

5.1 Main Findings and Contributions

This dissertation developed a multiscale upscaling workflow that uses physical-property values calculated from high-resolution CT images to predict three-dimensional property maps from low-resolution CT images covering whole rock cores. The porosity–property relationships and the variation of each property at the same porosity were then incorporated into the S-RPM, allowing three-dimensional property maps to be generated from porosity data at scales where continuous CT images are unavailable. The permeability component of the S-RPM was further used to examine how permeability variation affects CO_2 migration and storage at the field scale.

Chapter 2 first developed and validated a CNN–Darcy workflow for upscaling permeability from low-resolution CT images. Each permeability value calculated by LBM from a segmented high-resolution CT subvolume was paired with the corresponding raw low-resolution CT subvolume to train the CNN. The trained CNN predicted a permeability map for the whole core, and a Darcy-flow calculation was then used to obtain the upscaled core permeability. The workflow was then evaluated using an 8.7 cm Boise sandstone core to examine its applicability to a larger rock sample. As part of this larger-core study, the effects of subvolume size and the number of training samples were examined using 100^3, 200^3, and 300^3-voxel high-resolution subvolumes. Among the three cases, the 300^3-voxel case gave the highest prediction accuracy and the smoothest pressure and velocity fields. For the tested Boise sandstone, the 300^3-voxel subvolumes contained approximately 1240 pore bodies on average and showed more consistent pore connectivity than the 100^3-voxel subvolumes, which contained approximately 74 pore bodies. Thus, for

this Boise sandstone, subvolumes containing a larger connected pore network gave better predictions even though they provided fewer training samples.

Chapter 3 extended the workflow from permeability prediction to the simultaneous prediction of permeability, porosity, electrical conductivity, bulk modulus, and shear modulus. Because the 1 μm and 20 μm CT images were difficult to match reliably at the same physical location, DSgan converted each 300^3-voxel high-resolution subvolume into a 15^3-voxel low-resolution-like image. For each high-resolution subvolume, the DSgan-generated low-resolution-like subvolume was paired with the permeability, porosity, electrical conductivity, bulk modulus, and shear modulus calculated from the same source subvolume. This produced a paired training dataset without physically registering the high- and low-resolution CT images. Using these paired data, SE-Atten-Res3D achieved test-set R^2 values of 0.80–0.86 for all five properties and predicted their three-dimensional distributions throughout the whole core. Darcy and FEM calculations were then applied to these local property maps to obtain whole-core upscaled values for permeability, electrical conductivity, bulk modulus, and shear modulus.

The S-RPM used the predicted porosity map together with the variation of each property at the same porosity to reconstruct three-dimensional maps of permeability, electrical conductivity, bulk modulus, and shear modulus. These reconstructed maps retained both the average porosity–property relationships and the spread of property values at the same porosity observed in the DL predictions. The whole-core upscaled values calculated from the S-RPM maps differed by less than 0.3% from those calculated from the DL-predicted maps for all four properties.

Chapter 4 used the permeability component of the S-RPM to investigate whether the permeability variation obtained from J7 would affect CO_2 migration and storage in the Johansen vertical-equilibrium model. Two simulations were compared under the same reservoir and injection conditions: one used the deterministic Kozeny–Carman (KC) permeability field, and the other used a permeability field generated by the S-RPM calibrated using the J7 porosity–permeability data. The plume migration patterns and the values of V_{trap}, V_{res}, and V_{mob} were similar in the two cases because the permeability variation represented by the J7 data was small. In addition, the vertical-equilibrium model used one vertically averaged permeability value for each column, further reducing the influence of small local permeability differences.

A sensitivity analysis was then performed by varying two S-RPM parameters: σ_{scale}, which controls the amplitude of permeability variation, and ℓ, which controls

the spatial correlation length of the permeability field. Different combinations of σ_{scale} and ℓ were used to examine how the amplitude and spatial correlation length of permeability variation affect CO_2 migration and the values of V_{trap}, V_{res}, and V_{mob}. In this analysis, V_{trap} represents CO_2 remaining in regions previously occupied by the plume after the plume retreats, V_{res} represents CO_2 held at residual saturation within the current plume region, and V_{mob} represents CO_2 that remains mobile. Relative to the deterministic case, the case with $\sigma_{\text{scale}} = 5$ and $\ell = 10$ cells reduced V_{trap} from 118.3 to 80.5 million m^3, increased V_{res} from 87.3 to 95.8 million m^3, and increased V_{mob} from 305.7 to 335.1 million m^3. These results show that both the magnitude of permeability variation and the spatial continuity of high- and low-permeability regions can change the CO_2 plume path and the three CO_2 volumes, V_{trap}, V_{res}, and V_{mob}.

Across the three technical chapters, high-resolution CT-based calculations provided physical-property data, and DL combined these data with low-resolution CT images to predict whole-core property maps. The S-RPM then used the porosity–property variation identified at the core scale to generate property fields, and its permeability component was applied in field-scale CO_2 storage simulations to evaluate plume migration and storage behavior.

5.2 Limitations and Future Work

Future work can be organized into two directions: improving the pore-to-core workflow and extending its application from the core scale to the field scale.

For pore-to-core upscaling, the workflow was evaluated using centimeter-scale Boise and J7 sandstone cores. A useful next step would be to examine whether the same workflow remains effective for CT images covering larger rock samples. DSgan was evaluated for a 20-fold difference in voxel size, from 1 to 20 μm/voxel, and its applicability across larger resolution gaps could also be examined. Applying the workflow to carbonates and other rocks with different pore structures would help clarify how broadly the method can be used. These additional rock types could also be used to examine how the suitable subvolume size changes with pore size and pore connectivity. Further optimization of the ResNet34 and SE-Atten-Res3D architectures may improve prediction accuracy. The current framework predicts single-phase absolute permeability and, for water-saturated rock, electrical conductivity and effective bulk and shear moduli. Future work could extend the framework to relative permeability, capillary pressure, wettability effects, and electrical and elastic prop-

erties at different fluid saturations.

For core-to-field upscaling, the Johansen simulations using the S-RPM calibrated with J7 data and the deterministic KC model produced similar results because the permeability variation in J7 was small. A useful next step would be to calibrate the S-RPM using digital-rock data from rocks with larger permeability variation and examine how this variation changes CO_2 plume migration and storage at the field scale. In addition, a field-scale reservoir such as Johansen contains multiple layers and regions separated by faults. Different rock types in these regions have different porosity–permeability relationships. Accordingly, a field-scale S-RPM needs to represent multiple porosity–permeability relationships. Future work could extend the S-RPM by determining separate porosity–permeability relationships for each rock type or layer from core measurements and well-log data. These relationships could then be assigned to the corresponding reservoir regions based on mapped layer and fault information. Future work could also combine full three-dimensional multiphase simulations with geochemical reaction models to represent vertical layering and faults directly and to account for capillary pressure, dissolution trapping, and mineral trapping.

Overall, this dissertation provides an upscaling workflow that connects pore-scale physical-property calculations to core-scale property prediction and, for permeability, to field-scale reservoir evaluation.

Acknowledgement

I would like to express my sincere gratitude to Professor Fei Jiang of the Graduate School of Sciences and Technology for Innovation, Yamaguchi University, who served as my doctoral advisor, for his continuous guidance and support throughout my doctoral research. His extensive knowledge, constructive advice, and encouragement have been invaluable throughout my research, and his guidance has greatly contributed to the completion of this dissertation.

I would also like to express my sincere appreciation to Professor Shinsuke Mochizuki, Professor Kenichiro Tanoue, Professor Masato Mikami, and Professor Ryoya Shiraishi for serving as members of my dissertation committee and for their valuable comments and suggestions.

I would like to thank the Japan Organization for Metals and Energy Security (JOGMEC) and ENEOS Corporation for providing valuable research opportunities and support related to this study.

This work was supported by JSPS KAKENHI Grant Numbers JP20H02069, JP22K03927, JP24H00440, and JP21H05202, and by the JST SPRING Grant Number JPMJSP2111.

I am grateful to all members of the Fluid Engineering Laboratory at Yamaguchi University for their cooperation, discussions, and support throughout my doctoral study.

Finally, I would like to express my deepest gratitude to my mother, Xueyan Huai, for her endless support and encouragement throughout my life. Without her unconditional support, encouragement, and belief in me, I would not have been able to achieve these research outcomes or complete this dissertation.

I would also like to express my appreciation to everyone who has supported me throughout my seven years of study at Yamaguchi University. Although I cannot mention everyone individually, I sincerely appreciate their kindness and support.

Appendix A

Neural network architectures

The neural network architectures used in this study are presented here. In the following figures, ReLU refers to the rectified linear unit activation function Nair and Hinton (2010). A block containing five numbers indicates a convolutional layer: the first three numbers represent the kernel size in the three spatial dimensions, the fourth number represents the stride size, and the fifth number represents the number of filters. MP denotes a max-pooling layer, with the kernel size shown inside the block. BN denotes batch normalization Ioffe and Szegedy (2015). GAP denotes global average pooling. FC denotes filter concatenation in the inception modules. In the dense blocks, DP denotes a dropout layer applied after the dense layer. The size of each dense layer is shown at the top of the corresponding dense block.

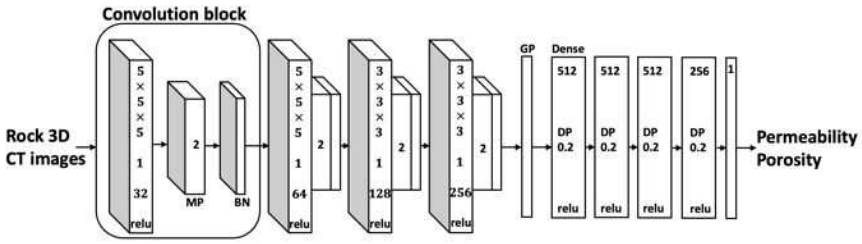


Figure A.1: Conv8 network architecture.

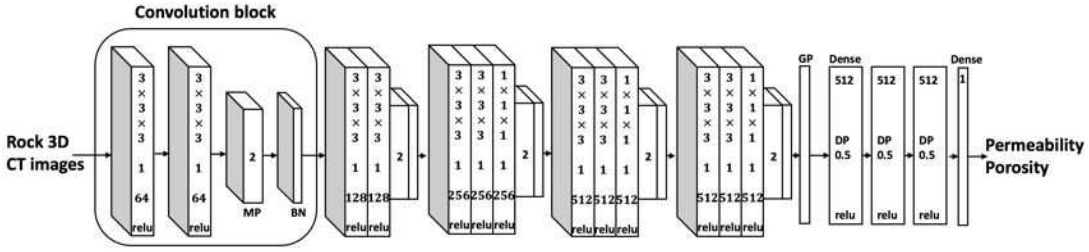


Figure A.2: VGG16 network architecture.

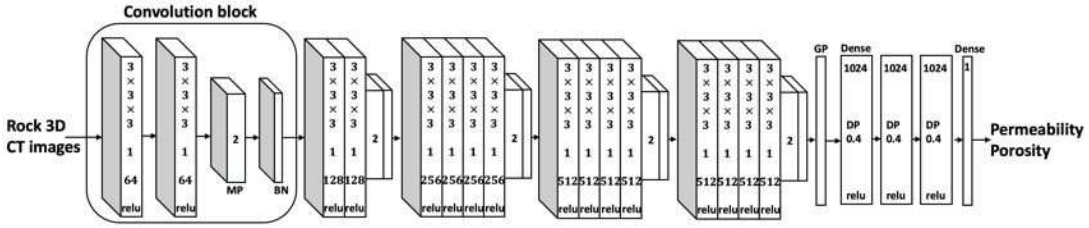


Figure A.3: VGG19 network architecture.

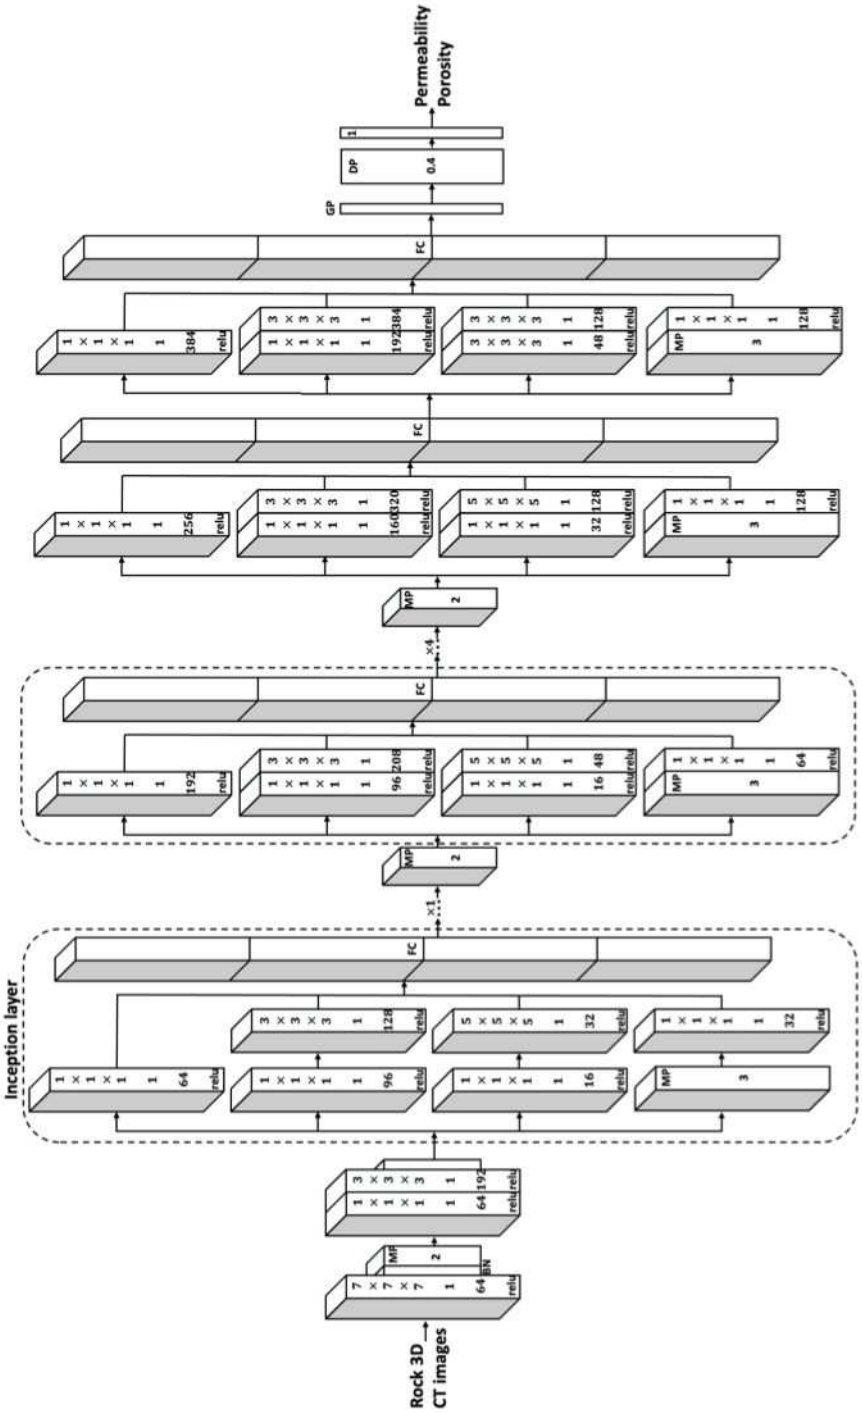


Figure A.4: GoogLeNet-V1 network architecture.

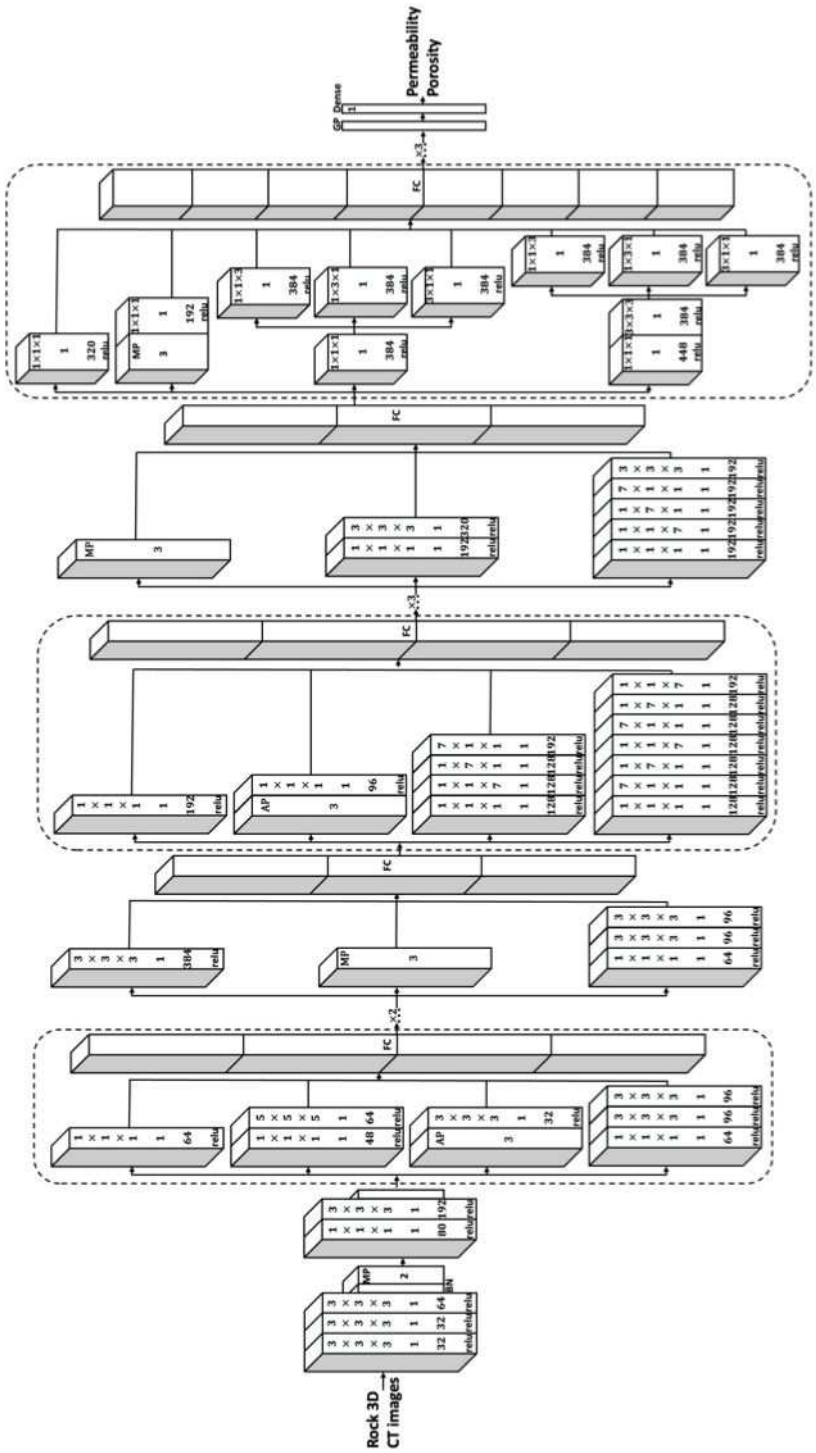


Figure A.5: GoogLeNet-V3 network architecture.

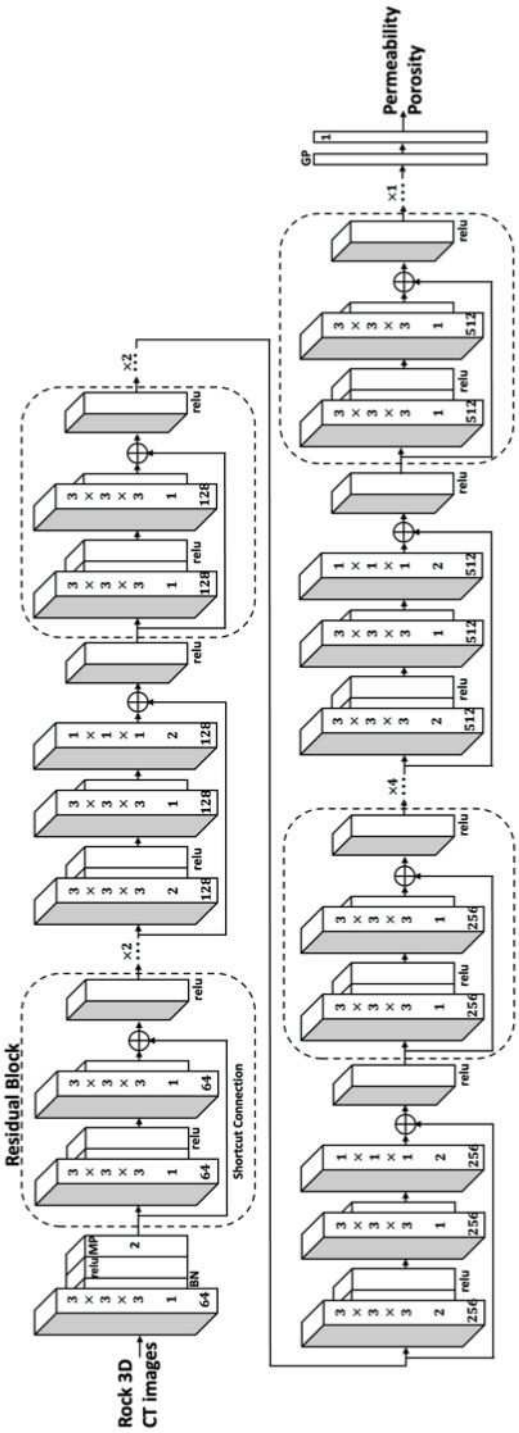


Figure A.6: ResNet34 network architecture.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. DOI: 10.48550/arXiv.1603.04467; Software available from tensorflow.org.

Alqahtani, N., Alzubaidi, F., Armstrong, R. T., Swietojanski, P., and Mostaghimi, P. (2020). Machine learning for predicting properties of porous media from 2d X-ray images. *Journal of Petroleum Science and Engineering*, 184:106514. DOI: 10.1016/j.petrol.2019.106514.

Alqahtani, N. J., Chung, T., Wang, Y. D., Armstrong, R. T., Swietojanski, P., and Mostaghimi, P. (2021). Flow-based characterization of digital rock images using deep learning. *SPE Journal*, 26(04):1800–1811. DOI: 10.2118/205376-PA.

Andrä, H., Combaret, N., Dvorkin, J., Glatt, E., Han, J., Kabel, M., Keehm, Y., Krzikalla, F., Lee, M., Madonna, C., et al. (2013). Digital rock physics benchmarks —part II: Computing effective properties. *Computers & Geosciences*, 50:33–43. DOI: 10.1016/j.cageo.2012.09.008.

Anto-Darkwah, E., Benson, S., and Rabinovich, A. (2021). An improved procedure for sub-core property characterization using data from multiple coreflooding experiments. *International Journal of Greenhouse Gas Control*, 105:103226. DOI: 10.1016/j.ijggc.2020.103226.

Anto-Darkwah, E. and Rabinovich, A. (2022). Modeling imbibition coreflooding in heterogeneous cores with sub-core scale hysteresis. *Advances in Water Resources*, 164:104214. DOI: 10.1016/j.advwatres.2022.104214.

Archie, G. E. et al. (2003). The electrical resistivity log as an aid in determining some reservoir characteristics. *SPE reprint series*, pages 9–18. DOI: 10.2118/942054-G.

Arnason, K. (2000). The resistivity structure of high-temperature geothermal systems in iceland. *Proceedings of the World⋯*. URL: http://www.geothermal-energy.org/pdf/IGAstandard/WGC/2000/R0416.PDF.

Arns, C. H., Knackstedt, M. A., Pinczewski, M. V., and Lindquist, W. (2001). Accurate estimation of transport properties from microtomographic images. *Geophysical research letters*, 28(17):3361–3364. DOI: 10.1029/2001GL012987.

Arns, C. H., Knackstedt, M. A., Pinczewski, W. V., and Garboczi, E. J. (2002). Computation of linear elastic properties from microtomographic images: Methodology and agreement between theory and experiment. *Geophysics*, 67(5):1396–1405. DOI: 10.1190/1.1512785.

Arns, C. H., Madadi, M., Sheppard, A. P., and Knackstedt, M. A. (2007). Linear elastic properties of granular rocks derived from Xray-CT images. In *2007 SEG Annual Meeting*. OnePetro. DOI: 10.1190/1.2792823.

Avseth, P., Dvorkin, J., Mavko, G., and Rykkje, J. (2000). Rock physics diagnostic of North Sea sands: Link between microstructure and seismic properties. *Geophysical research letters*, 27(17):2761–2764. DOI: 10.1029/1999GL008468.

Bachu, S. (2008). CO_2 storage in geological media: Role, means, status and barriers to deployment. *Prog. Energy Combust. Sci.*, 34(2):254–273. DOI: 10.1016/j.pecs.2007.10.001.

Bade, S., Taiwo, K., Ndulue, U., Tomomewo, O., and Oni, B. (2024). A review of underground hydrogen storage systems: Current status, modeling approaches, challenges, and future prospective. *Int. J. Hydrogen Energy*, 80:449–474. DOI: 10.1016/j.ijhydene.2024.07.187.

Bear, J. (2013). *Dynamics of fluids in porous media*. Courier Corporation, Mineola, NY. URL: https://store.doverpublications.com/products/9780486656755.

Bergmann, P., Schmidt-Hattenberger, C., Kiessling, D., Rücker, C., Labitzke, T., Henninges, J., Baumann, G., and Schütt, H. (2012). Surface-downhole electrical resistivity tomography applied to monitoring of co2 storage at ketzin, germany. *Geophysics*, 77(6):B253–B267. DOI: 10.1190/geo2011-0515.1.

Bizhani, M., Ardakani, O. H., and Little, E. (2022). Reconstructing high fidelity digital rock images using deep convolutional neural networks. *Scientific reports*, 12(1):1–14. DOI: 10.1038/s41598-022-08170-8.

Blunt, M. J. (2001). Flow in porous media—pore-network models and multiphase flow. *Current opinion in colloid & interface science*, 6(3):197–207. DOI: 10.1016/S1359-0294(01)00084-X.

Blunt, M. J., Bijeljic, B., Dong, H., Gharbi, O., Iglauer, S., Mostaghimi, P., Paluszny, A., and Pentland, C. (2013). Pore-scale imaging and modelling. *Advances in Water resources*, 51:197–216. DOI: 10.1016/j.advwatres.2012.03.003.

Boek, E. S. and Venturoli, M. (2010). Lattice-Boltzmann studies of fluid flow in porous media with realistic rock geometries. *Computers & Mathematics with Applications*, 59(7):2305–2314. DOI: 10.1016/j.camwa.2009.08.063.

Botchkarev, A. (2018). Performance metrics (error measures) in machine learning regression, forecasting and prognostics: Properties and typology. *arXiv preprint arXiv:1809.03006*. DOI: 10.48550/arXiv.1809.03006.

Botha, P. W. and Sheppard, A. P. (2016). Mapping permeability in low-resolution micro-CT images: A multiscale statistical approach. *Water Resources Research*, 52(6):4377–4398. DOI: 10.1002/2015WR018454.

Carman, P. (1937). Fluid flow through granular beds. *Trans. Inst. Chem. Eng.*, 15:150–166. DOI: 10.1016/S0263-8762(97)80003-2.

Chawshin, K., Berg, C. F., Varagnolo, D., and Lopez, O. (2022). Automated porosity estimation using CT-scans of extracted core data. *Computational Geosciences*, 26(3):595–612. DOI: 10.1007/s10596-022-10143-9.

Chen, X., Yang, J., Ma, L., Rabbani, A., and Babaei, M. (2023). Predicting 3D physical properties from a single 2D slice based on convolutional neural networks: 2D-slice-to-3D-properties for porous rocks. *Water Resources Research*, 59(9):e2023WR035521. DOI: 10.1029/2023WR035521.

Cheung, C., Baud, P., and Wong, T.-f. (2012). Effect of grain size distribution on the development of compaction localization in porous sandstone. *Geophys. Res. Lett.*, 39(21). DOI: 10.1029/2012GL053739.

Christie, M. A. and Blunt, M. J. (2001). Tenth spe comparative solution project: A comparison of upscaling techniques. *SPE Reservoir Evaluation & Engineering*, 4(04):308–317. DOI: 10.2118/72469-PA.

Clauser, C. (1992). Permeability of crystalline rocks. *Eos Trans. AGU*, 73(21):233–238. DOI: 10.1029/91EO00190.

Cnudde, V. and Boone, M. N. (2013). High-resolution X-ray computed tomography in geosciences: A review of the current technology and applications. *Earth-Science Reviews*, 123:1–17. DOI: 10.1016/j.earscirev.2013.04.003.

Costa, A. (2006). Permeability-porosity relationship: A reexamination of the Kozeny-Carman equation based on a fractal pore-space geometry assumption. *Geophysical research letters*, 33(2). DOI: 10.1029/2005GL025134.

Da Wang, Y., Armstrong, R. T., and Mostaghimi, P. (2019). Enhancing resolution of digital rock images with super resolution convolutional neural networks. *Journal of Petroleum Science and Engineering*, 182:106261. DOI: 10.1016/j.petrol.2019.106261.

de Figueiredo, L. P., Grana, D., Bordignon, F. L., Santos, M., Roisenberg, M., and Rodrigues, B. B. (2018). Joint Bayesian inversion based on rock-physics prior modeling for the estimation of spatially correlated reservoir properties. *Geophysics*, 83(5):M49–M61. DOI: 10.1190/geo2017-0463.1.

Dehghan Khalili, A., Arns, J.-Y., Hussain, F., Cinar, Y., Pinczewski, W., Arns, C. H., et al. (2013). Permeability upscaling for carbonates from the pore scale by use of multiscale X-ray-CT images. *SPE Reservoir Evaluation & Engineering*, 16(04):353–368. DOI: 10.2118/152640-PA.

d'Humières, D. (1992). Generalized lattice-Boltzmann equations. *Rarefied gas dynamics*. DOI: 10.2514/5.9781600866319.0450.0458.

Dong, J.-J., Hsu, J.-Y., Wu, W.-J., Shimamoto, T., Hung, J.-H., Yeh, E.-C., Wu, Y.-H., and Sone, H. (2010). Stress-dependence of the permeability and porosity of sandstone and shale from TCDP Hole-A. *International Journal of Rock Mechanics and Mining Sciences*, 47(7):1141–1157. DOI: 10.1016/j.ijrmms.2010.06.019.

Draper, N. and Smith, H. (1998). *Applied regression analysis*. Wiley, New York. DOI: 10.1002/9781118625590.

Drobek, T., Strobel, J., Park, S., Lieu, K., Altermann, W., Lemmel, H., Lindner, P., and Stark, R. (2011). Small angle neutron scattering analysis of porous reservoir rocks. *Proceedings of the Society of Core Analysts*. URL: https://jgmaas.com/SCA/2011/SCA2011-48.pdf.

Durlofsky, L. J. (1991). Numerical calculation of equivalent grid block permeability tensors for heterogeneous porous media. *Water resources research*, 27(5):699–708. DOI: 10.1029/91WR00107.

Dvorkin, J., Derzhi, N., Diaz, E., and Fang, Q. (2011). Relevance of computational rock physics. *Geophysics*, 76(5):E141–E153. DOI: 10.1190/geo2010-0352.1.

Eichel, H., Helmig, R., Neuweiler, I., and Cirpka, O. A. (2005). Upscaling of two-phase flow processes in porous media. In *Upscaling multiphase flow in porous media: from pore to core and beyond*, pages 237–257. Springer. DOI: 10.1007/1-4020-3604-3_12.

Eidsvik, J., Avseth, P., Omre, H., Mukerji, T., and Mavko, G. (2004). Stochastic reservoir characterization using prestack seismic data. *Geophysics*, 69(4):978–993. DOI: 10.1190/1.1778241.

Eigestad, G. T., Dahle, H. K., Hellevang, B., Riis, F., Johansen, W. T., and Øian, E. (2009). Geological modeling and simulation of co2 injection in the johansen formation. *Computational Geosciences*, 13(4):435–450. DOI: 10.1007/s10596-009-9153-y.

Elmorsy, M., El-Dakhakhni, W., and Zhao, B. (2023). Rapid permeability upscaling of digital porous media via physics-informed neural networks. *Water Resour. Res.*, 59(12):e2023WR035064. DOI: 10.1029/2023WR035064.

Emeliana, Y., Loveni Melkisedek, B., Nurma Yulita, I., and Ary Dharmawan, I. (2025). Residual attention U-Net for predicting fluid dynamics and permeability in two-dimensional digital porous rocks. *Machine Learning: Science and Technology*, 6(4):045063. DOI: 10.1088/2632-2153/ae2564.

Farquharson, J., Heap, M. J., Varley, N. R., Baud, P., and Reuschlé, T. (2015). Permeability and porosity relationships of edifice-forming andesites: a combined field and laboratory study. *Journal of Volcanology and Geothermal Research*, 297:52–68. DOI: 10.1016/j.jvolgeores.2015.03.016.

Fredrich, J., Greaves, K., and Martin, J. (1993). Pore geometry and transport properties of Fontainebleau sandstone. *International journal of rock mechanics and mining sciences & geomechanics abstracts*, 30(7):691–697. DOI: 10.1016/0148-9062(93)90007-Z.

Fugui, L., Yongfei, Y., Haiyuan, Y., Liu, T., Yunwei, T., Hai, S., Jun, Y., et al. (2025). Pore structure properties characterization of shale using generative adversarial network: Image augmentation, super-resolution reconstruction, and multi-mineral auto-segmentation. *Petroleum Exploration and Development*, 52(5):1262–1274. DOI: 10.1016/S1876-3804(25)60640-9.

Garboczi, E. J. (1998). Finite element and finite difference programs for com-puting the linear electric and elastic properties of dig-ital images of random materials. *Report NISTIR*. DOI: 10.6028/NIST.IR.6269.

Gassmann, F. (2007). On elasticity of porous media. *Classics of Elastic Wave Theory*, pages 389–408. DOI: 10.1190/1.9781560801931.ch3p.

Goel, G., Abidoye, L. K., Chahar, B. R., and Das, D. B. (2016). Scale dependency of dynamic relative permeability-satuartion curves in relation with fluid viscosity and dynamic capillary pressure effect. *Environmental Fluid Mechanics*, 16(5):945–963. DOI: 10.1007/s10652-016-9469-y.

Goral, J., Walton, I., Andrew, M., and Deo, M. (2019). Pore system characterization of organic-rich shales using nanoscale-resolution 3D imaging. *Fuel*, 258:116049. DOI: 10.1016/j.fuel.2019.116049.

Gostick, J., Aghighi, M., Hinebaugh, J., Tranter, T., Hoeh, M., Day, H., Spellacy, B., Sharqawy, M., Bazylak, A., Burns, A., et al. (2016). OpenPNM: A pore network modeling package. *Comput. Sci. Eng.*, 18(4):60–74. DOI: 10.1109/MCSE.2016.49.

Gostick, J., Khan, Z., Tranter, T., Kok, M., Agnaou, M., Sadeghi, M., and Jervis, R. (2019). PoreSpy: A python toolkit for quantitative analysis of porous media images. *J. Open Source Softw.*, 4(37):1296. DOI: 10.21105/joss.01296.

Grana, D., Pirrone, M., and Mukerji, T. (2012). Quantitative log interpretation and uncertainty propagation of petrophysical properties and facies classification from rock-physics modeling and formation evaluation analysis. *Geophysics*, 77(3):WA45–WA63. DOI: 10.1190/geo2011-0272.1.

Guibert, R., Horgue, P., Debenest, G., and Quintard, M. (2016). A comparison of various methods for the numerical evaluation of porous media permeability tensors from pore-scale geometry. *Mathematical Geosciences*, 48(3):329–347. DOI: 10.1007/s11004-015-9587-9.

Guo, Y., Jiang, F., Tsuji, T., Kato, Y., Shimokawara, M., Esteban, L., Seyyedi, M., Pervukhina, M., Lebedev, M., and Kitamura, R. (2025). Machine learning-based upscaling of rock permeability from pore scale to core scale: effect of training dataset size and sub-core volumes. *Earth Energy Science*. DOI: 10.1016/j.ees.2025.11.008.

Guo, Z. and Shu, C. (2013). *Lattice Boltzmann method and its application in engineering*, volume 3. World Scientific. DOI: 10.1142/8806.

Hamidian, M., Miri, R., and Fazeli, H. (2025). Generative adversarial network-based super-resolution of subsurface rock images: Visual, petrophysical, and flow simulation assessment. *Advances in Water Resources*, page 105184. DOI: 10.1016/j.advwatres.2025.105184.

Hariharan, B., Arbeláez, P., Girshick, R., and Malik, J. (2015). Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 447–456. DOI: 10.1109/CVPR.2015.7298642.

He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778. DOI: 10.1109/CVPR.2016.90.

Hilhorst, M. A. (2000). A pore water conductivity sensor. *Soil Science Society of America Journal*, 64(6):1922–1925. DOI: 10.2136/sssaj2000.6461922x.

Hill, R. (1963). Elastic properties of reinforced solids: some theoretical principles. *Journal of the Mechanics and Physics of Solids*, 11(5):357–372. DOI: 10.1016/0022-5096(63)90036-X.

Hong, J. and Liu, J. (2020). Rapid estimation of permeability from digital rock using 3D convolutional neural network. *Computational Geosciences*, 24(4):1523–1539. DOI: 10.1007/s10596-020-09941-w.

Hornung, U. (1997). *Homogenization and porous media*, volume 6. Springer Science & Business Media. DOI: 10.1007/978-1-4612-1920-0.

Hou, Z. and Cao, D. (2022). Estimating elastic parameters from digital rock images based on multi-task learning with multi-gate mixture-of-experts. *Journal of Petroleum Science and Engineering*, 213:110310. DOI: 10.1016/j.petrol.2022.110310.

Hu, J., Shen, L., and Sun, G. (2018). Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141. DOI: 10.1109/CVPR.2018.00745.

Huber, P. J. (1992). Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer. DOI: 10.1007/978-1-4612-4380-9_35.

Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR. DOI: 10.48550/arXiv.1502.03167.

Jaderberg, M., Simonyan, K., Zisserman, A., et al. (2015). Spatial transformer networks. *Advances in neural information processing systems*, 28. DOI: 10.48550/arXiv.1506.02025.

Jiang, F., Guo, Y., Tsuji, T., Kato, Y., Shimokawara, M., Esteban, L., Seyyedi, M., Pervukhina, M., Lebedev, M., and Kitamura, R. (2023). Upscaling permeability using multiscale X-ray-CT images with digital rock modeling and deep learning techniques. *Water Resour. Res.*, 59(3):e2022WR033267. DOI: 10.1029/2022WR033267.

Jiang, F., Liu, H., Chen, X., and Tsuji, T. (2022). A coupled LBM-DEM method for simulating the multiphase fluid-solid interaction problem. *Journal of Computational Physics*, page 110963. DOI: 10.1016/j.jcp.2022.110963.

Jiang, F., Matsumura, K., Ohgi, J., and Chen, X. (2021a). A GPU-accelerated fluid–structure-interaction solver developed by coupling finite element and lattice Boltzmann methods. *Computer Physics Communications*, 259:107661. DOI: 10.1016/j.cpc.2020.107661.

Jiang, F. and Tsuji, T. (2014). Changes in pore geometry and relative permeability caused by carbonate precipitation in porous media. *Physical Review E*, 90(5):053306. DOI: 10.1103/PhysRevE.90.053306.

Jiang, F. and Tsuji, T. (2015). Impact of interfacial tension on residual CO2 clusters in porous sandstone. *Water Resources Research*, 51(3):1710–1722. DOI: 10.1002/2014WR016070.

Jiang, F. and Tsuji, T. (2017). Estimation of three-phase relative permeability by simulating fluid dynamics directly on rock-microstructure images. *Water Resources Research*, 53(1):11–32. DOI: 10.1002/2016WR019098.

Jiang, F., Tsuji, T., and Hu, C. (2014). Elucidating the role of interfacial tension for hydrological properties of two-phase flow in natural sandstone by an improved lattice Boltzmann method. *Transport in porous media*, 104(1):205–229. DOI: 10.1007/s11242-014-0329-0.

Jiang, F., Yang, J., Boek, E., and Tsuji, T. (2021b). Investigation of viscous coupling effects in three-phase flow by lattice Boltzmann direct simulation and machine learning technique. *Advances in Water Resources*, 147:103797. DOI: 10.1016/j.advwatres.2020.103797.

Juanes, R., Spiteri, E., Orr Jr, F., and Blunt, M. (2006). Impact of relative permeability hysteresis on geological CO2 storage. *Water resources research*, 42(12). DOI: 10.1029/2005WR004806.

Kang, Q., Li, K.-Q., Fu, J.-L., and Liu, Y. (2024). Hybrid LBM and machine learning algorithms for permeability prediction of porous media: A comparative study. *Comput. Geotech.*, 168:106163. DOI: 10.1016/j.compgeo.2024.106163.

Khan, M. and Khanal, A. (2024). Machine learning assisted prediction of porosity and related properties using digital rock images. *ACS Omega*, 9(28):30205–30223. DOI: 10.1021/acsomega.3c10131.

Kinga, D., Adam, J. B., et al. (2015). Adam: A method for stochastic optimization. In *International conference on learning representations (ICLR)*. California;. DOI: 10.48550/arXiv.1412.6980.

Ko, D., Ji, H., and Ju, Y. (2023). Prediction of 3D velocity field of reticulated foams using deep learning for transport analysis. *Transp. Porous Media*, 148(3):577–604. DOI: 10.1007/s11242-023-01961-1.

Koroteev, D. and Tekic, Z. (2021). Artificial intelligence in oil and gas upstream: Trends, challenges, and scenarios for the future. *Energy AI*, 3:100041. DOI: 10.1016/j.egyai.2020.100041.

Krause, M., Perrin, J.-C., Benson, S. M., et al. (2011). Modeling permeability distributions in a sandstone core for history matching coreflood experiments. *Spe Journal*, 16(04):768–777. DOI: 10.2118/126340-PA.

Lake, L. W., Johns, R., Rossen, B., Pope, G. A., et al. (2014). *Fundamentals of enhanced oil recovery*, volume 1. Society of Petroleum Engineers Richardson, TX. DOI: 10.2118/9781613993286.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. (2017). Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30. DOI: 10.48550/arXiv.1612.01474.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324. DOI: 10.1109/5.726791.

Lehmann, T. M., Gonner, C., and Spitzer, K. (2002). Survey: Interpolation methods in medical image processing. *IEEE transactions on medical imaging*, 18(11):1049–1075. URL: http://jmlr.org/papers/v15/srivastava14a.html.

Li, G. and Xie, C. (2021). Calculation of matrix permeability from velocity and attenuation of ultrasonic S-wave. *J. Geophys. Eng.*, 18(6):984–994. DOI: 10.1093/jge/gxab065.

Li, S., Jiang, F., Wei, B., Hou, J., and Liu, H. (2021). Prediction of three-phase relative permeabilities of Berea sandstone using lattice Boltzmann method. *Physics of Fluids*, 33(6):063302. DOI: 10.1063/5.0050727.

Lie, K.-A. (2019a). *An introduction to reservoir simulation using MATLAB/GNU Octave: User guide for the MATLAB reservoir simulation toolbox (MRST)*. Cambridge University Press, Cambridge. DOI: 10.1017/9781108591416.

Lie, K.-A. (2019b). *Single-Phase Flow*, page 111 – 112. Cambridge University Press. DOI: 10.1017/9781108591416.

Liu, H., Lu, Y., Li, S., Yu, Y., and Sahu, K. C. (2021). Deformation and breakup of a compound droplet in three-dimensional oscillatory shear flow. *International Journal of Multiphase Flow*, 134:103472. DOI: 10.1016/j.ijmultiphaseflow.2020.103472.

Liu, Y., Gong, W., Zhao, Y., Jin, X., and Wang, M. (2022). A pore-throat segmentation method based on local hydraulic resistance equivalence for pore-network modeling. *Water Resour. Res.*, 58(12):e2022WR033142. DOI: 10.1029/2022WR033142.

Louis, L., Boyd, P., Hofmann, R., and Saxena, N. (2024). On the role and evolution of local grain size heterogeneity during confined compression of Boise sandstone as seen by X-ray micro-CT imaging. *Rock Mech. Rock Eng.*, pages 1–22. DOI: 10.1007/s00603-024-04152-6.

Maas, A. L., Hannun, A. Y., Ng, A. Y., et al. (2013). Rectifier nonlinearities improve neural network acoustic models. In *Proc. ICML*, page 3. Atlanta, GA. [PDF Available Online].

Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Paul Smolley, S. (2017). Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2794–2802. DOI: 10.1109/ICCV.2017.304.

Martys, N. S., Torquato, S., and Bentz, D. (1994). Universal scaling of fluid permeability for sphere packings. *Physical Review E*, 50(1):403. DOI: 10.1103/PhysRevE.50.403.

Mavko, G., Mukerji, T., and Dvorkin, J. (2020). *The rock physics handbook*. Cambridge university press. DOI: 10.1017/9781108333016.

Menke, H. P., Maes, J., and Geiger, S. (2020). Upscaling the porosity-permeability relationship of a microporous carbonate to the darcy scale with machine learning. *arXiv preprint arXiv:2010.00397*. DOI: 10.48550/arXiv.2010.00397.

Menke, H. P., Maes, J., and Geiger, S. (2021). Upscaling the porosity–permeability relationship of a microporous carbonate for Darcy-scale flow with machine learning. *Scientific Reports*, 11(1):2625. DOI: 10.1038/s41598-021-82029-2.

Mukerji, T., Jørstad, A., Avseth, P., Mavko, G., and Granli, J. (2001). Mapping lithofacies and pore-fluid probabilities in a North Sea reservoir: Seismic inversions and statistical rock physics. *Geophysics*, 66(4):988–1001. DOI: 10.1190/1.1487078.

Munoz, G. (2014). Exploring for geothermal resources with electromagnetic methods. *Surveys in geophysics*, 35(1):101–122. DOI: 10.1007/s10712-013-9236-0.

Nair, V. and Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. In *Icml*. URL: https://icml.cc/Conferences/2010/papers/432.pdf.

Nilsen, H. M., Lie, K.-A., and Andersen, O. (2015). Analysis of co2 trapping capacities and long-term migration for geological formations in the norwegian north sea using mrst-co2lab. *Computers & geosciences*, 79:15–26. DOI: 10.1016/j.cageo.2015.03.001.

Norris, A. (2006). The isotropic material closest to a given anisotropic material. *Journal of Mechanics of Materials and Structures*, 1(2):223–238. DOI: 10.2140/jomms.2006.1.223.

Nur, A. M., Mavko, G., Dvorkin, J., and Gal, D. (1995). Critical porosity: The key to relating physical properties to porosity in rocks. In *SEG technical program expanded abstracts 1995*, pages 878–881. Society of Exploration Geophysicists. DOI: 10.1190/1.1887540.

Nye, J. F. (1985). *Physical properties of crystals: their representation by tensors and matrices*. Oxford university press. URL: https://global.oup.com/academic/product/physical-properties-of-crystals-9780198511656.

Onimisi, T., Lashore, B., Akanji, L., and Gomes, J. (2023). A constrained proper orthogonal decomposition model for upscaling permeability. *Int. J. Numer. Methods Fluids*, 95(6):899–916. DOI: 10.1002/fld.5171.

Otosigbo, G. (2017). A comparative study of profile permeability and air permeability in reservoir characterization: The case of GX 2 well in western Niger delta basin. *FUNAI Journal of Science and Technology*, 3(1):110. URL: https://cir.nii.ac.jp/crid/1370298755605847703.

Pan, C., Luo, L.-S., and Miller, C. T. (2006). An evaluation of lattice Boltzmann schemes for porous medium flow simulation. *Computers & fluids*, 35(8-9):898–909. DOI: 10.1016/j.compfluid.2005.03.008.

Pape, H., Clauser, C., and Iffland, J. (1999). Permeability prediction based on fractal pore-space geometry. *Geophysics*, 64(5):1447–1460. DOI: 10.1190/1.1444649.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32. DOI: 10.48550/arXiv.1912.01703.

Rabbani, A. and Babaei, M. (2019). Hybrid pore-network and lattice-Boltzmann permeability modelling accelerated by machine learning. *Advances in water resources*, 126:116–128. DOI: 10.1016/j.advwatres.2019.02.012.

Raeini, A. Q., Blunt, M. J., and Bijeljic, B. (2012). Modelling two-phase flow in porous media at the pore scale using the volume-of-fluid method. *Journal of Computational Physics*, 231(17):5653–5668. DOI: 10.1016/j.jcp.2012.04.011.

Raza, A., Gholami, R., Rezaee, R., Rasouli, V., and Rabiei, M. (2019). Significant aspects of carbon capture and storage–a review. *Petroleum*, 5(4):335–340. DOI: 10.1016/j.petlm.2018.12.007.

Roslin, A., Marsh, M., Piché, N., Provencher, B., Mitchell, T., Onederra, I., and Leonardi, C. (2022). Processing of micro-CT images of granodiorite rock samples using CNN, part i: Super-resolution enhancement using a 3D CNN. *Miner. Eng.*, 188:107748. DOI: 10.1016/j.mineng.2022.107748.

Rybach, L. (2010). The future of geothermal energy and its challenges. *Proc. World Geothermal Congr.*, 29. URL: https://www.academia.edu/download/82112420/3109.pdf.

Santos, J. E., Xu, D., Jo, H., Landry, C. J., Prodanović, M., and Pyrcz, M. J. (2020). PoreFlow-Net: A 3D convolutional neural network to predict fluid flow through porous media. *Advances in Water Resources*, 138:103539. DOI: 10.1016/j.advwatres.2020.103539.

Shan, X., Yuan, X.-F., and Chen, H. (2006). Kinetic theory representation of hydrodynamics: a way beyond the Navier–Stokes equation. *Journal of Fluid Mechanics*, 550:413–441. DOI: 10.1017/S0022112005008153.

Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., and Webb, R. (2017). Learning from simulated and unsupervised images through adversarial training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2107–2116. DOI: 10.1109/CVPR.2017.241.

Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*. DOI: 10.48550/arXiv.1409.1556.

Spikes, K., Mukerji, T., Dvorkin, J., and Mavko, G. (2007). Probabilistic seismic inversion based on rock-physics models. *Geophysics*, 72(5):R87–R97. DOI: 10.1190/1.2760162.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958. DOI: 10.5555/2627435.2670313.

Suzuki, A., Fukui, K.-i., Onodera, S., Ishizaki, J., and Hashida, T. (2022). Data-driven geothermal reservoir modeling: Estimating permeability distributions by machine learning. *Geosciences*, 12(3):130. DOI: 10.3390/geosciences12030130.

Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9. DOI: 10.1109/CVPR.2015.7298594.

Taiwo, G., Tomomewo, O., and Oni, B. (2024). A comprehensive review of underground hydrogen storage: Insight into geological sites, economics, barriers, and outlook. *J. Energy Storage*, 90:111844. DOI: 10.1016/j.est.2024.111844.

Tang, Z., Qiu, S., Cai, Z., Law, C.-L., Mao, X., and Xu, P. (2026). Multi-gate mixture-of-different-experts model for the prediction of multiple properties in multi-phase rock. *Advances in Geo-Energy Research*, 19(2):182–196. DOI: 10.46690/ager.2026.02.06.

Telford, W. M., Geldart, L. P., and Sheriff, R. E. (1990). *Applied geophysics*. Cambridge university press. DOI: 10.1017/CBO9781139167932.

Telvari, S., Sayyafzadeh, M., Siavashi, J., and Sharifi, M. (2023). Prediction of two-phase flow properties for digital sandstones using 3D convolutional neural networks. *Adv. Water Resour.*, 176:104442. DOI: 10.1016/j.advwatres.2023.104442.

Temizel, C., Odi, U., Li, K., Liu, L., Tutun, S., and Santos, J. (2025). Permeability prediction using vision transformers. *Mathematical and Computational Applications*, 30(4):71. DOI: 10.3390/mca30040071.

Tester, J. W., Anderson, B. J., Batchelor, A., Blackwell, D., DiPippo, R., Drake, E., Garnish, J., Livesay, B., Moore, M., Nichols, K., et al. (2006). The future of geothermal energy. *Massachusetts Institute of Technology*, 358. URL: http://naturalresources.house.gov/uploadedfiles/testertestimony04.19.07.pdf.

Tran, D., Bourdev, L., Fergus, R., Torresani, L., and Paluri, M. (2015). Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497. DOI: 10.1109/ICCV.2015.510.

Tripathi, D. N., Hathon, L. A., and Myers, M. T. (2018). Exporting petrophysical properties of sandstones from thin section image analysis. In *SPWLA 59th Annual Logging Symposium*. OnePetro. URL: https://onepetro.org/SPWLAALS/proceedings-pdf/SPWLA18/SPWLA18/D043S006R001/1185334/spwla-2018-nnn.pdf.

Ugolkov, E., He, X., Kwak, H., and Hoteit, H. (2025). Super-resolution of 3D micro-CT images using generative adversarial networks: enhancing resolution and segmentation accuracy. *Computers & Geosciences*, page 106018. DOI: 10.1016/j.cageo.2025.106018.

Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2016). Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*. DOI: 10.48550/arXiv.1607.08022.

Velcin, H., Dautriat, J., Sarout, J., Esteban, L., and Godel, B. (2020). Experimental reactivation of shear-fractured Berea and Boise sandstones by brine or liquid CO2 injection at depth. *Journal of Geophysical Research: Solid Earth*, 125(2):e2019JB018281. DOI: 10.1029/2019JB018281.

Wang, F. and Zai, Y. (2023). Image segmentation and flow prediction of digital rock with U-net network. *Adv. Water Resour.*, 172:104384. DOI: 10.1016/j.advwatres.2023.104384.

Wang, Y., Blunt, M., Armstrong, R., and Mostaghimi, P. (2021a). Deep learning in pore scale imaging and modeling. *Earth-Sci. Rev.*, 215:103555. DOI: 10.1016/j.earscirev.2021.103555.

Wang, Y. D., Blunt, M. J., Armstrong, R. T., and Mostaghimi, P. (2021b). Deep learning in pore scale imaging and modeling. *Earth-Science Reviews*, 215:103555. DOI: 10.1016/j.earscirev.2021.103555.

Wang, Z., Dai, Z., Póczos, B., and Carbonell, J. (2019). Characterizing and avoiding negative transfer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11293–11302. DOI: 10.48550/arXiv.1811.09751.

Weinbrandt, R., Ramey, H., and Casse, F. (1975). The effect of temperature on relative and absolute permeability of sandstones. *Society of Petroleum Engineers Journal*, 15(05):376–384. DOI: 10.2118/4142-PA.

Wen, X.-H. and Gómez-Hernández, J. J. (1996). Upscaling hydraulic conductivities in heterogeneous media: An overview. *Journal of hydrology*, 183(1-2):ix–xxxii. DOI: 10.1016/S0022-1694(96)80030-8.

Wildenschild, D. and Sheppard, A. P. (2013). X-ray imaging and analysis techniques for quantifying pore-scale structure and processes in subsurface porous medium systems. *Advances in Water resources*, 51:217–246. DOI: 10.1016/j.advwatres.2012.07.018.

Winsauer, W. O., Shearin Jr, H., Masson, P. y., and Williams, M. (1952). Resistivity of brine-saturated sands in relation to pore geometry. *AAPG bulletin*, 36(2):253–277. DOI: 10.1306/3D9343F4-16B1-11D7-8645000102C1865D.

Withers, P. J., Bouman, C., Carmignato, S., Cnudde, V., Grimaldi, D., Hagen, C. K., Maire, E., Manley, M., Du Plessis, A., and Stock, S. R. (2021). X-ray computed tomography. *Nature Reviews Methods Primers*, 1(1):18. DOI: 10.1038/s43586-021-00015-4.

Worthington, P. F. (2004). The effect of scale on the petrophysical estimation of intergranular permeability. *Petrophysics-The SPWLA Journal of Formation Evaluation and Reservoir Description*, 45(01):59–72. URL: https://onepetro.org/petrophysics/article-abstract/45/01/59/170670/The-Effect-of-Scale-on-the-Petrophysical.

Wu, J., Xu, H., Xiong, B., Fang, C., Wang, S., Zong, P., and Liu, D. (2024). Effects of mineral displacement on geothermal reservoir properties at high temperatures identified using micro-CT and digital volume correlation. *Nat. Resour. Res.*, 33:1613–1623. DOI: 10.1007/s11053-024-10361-1.

Xie, C., Zhu, J., Wang, J., Yang, J., and Song, H. (2023). Direct prediction of relative permeability curve from 3D digital rock images based on deep learning approaches. *Int. J. Rock Mech. Min. Sci.*, 170:105544. DOI: 10.1016/j.ijrmms.2023.105544.

Xu, Z.-J. (2012). Homogenization and upscaling for diffusion, heat conduction, and wave propagation in heterogeneous materials. *Communications in Theoretical Physics*, 57(3):348. DOI: 10.1088/0253-6102/57/3/04.

Yan, P., He, X., Ma, Z., Gerke, K. M., Karsanina, M. V., and Teng, Q. (2025). Fusion reconstruction of porous materials using 2D high-and 3D low-resolution images via latent representation regression generative adversarial network. *Construction and Building Materials*, 494:143196. DOI: 10.1016/j.conbuildmat.2025.143196.

Yang, J. and Boek, E. S. (2013). A comparison study of multi-component lattice Boltzmann models for flow in porous media applications. *Computers & Mathematics with Applications*, 65(6):882–890. DOI: 10.1016/j.camwa.2012.11.022.

Yang, L., Yang, J., Boek, E., Sakai, M., and Pain, C. (2019). Image-based simulations of absolute permeability with massively parallel pseudo-compressible stabilised finite element solver. *Computational Geosciences*, 23(5):881–893. DOI: 10.1007/s10596-019-09837-4.

You, N., Li, Y. E., and Cheng, A. (2021). 3D carbonate digital rock reconstruction using progressive growing GAN. *Journal of Geophysical Research: Solid Earth*, 126(5):e2021JB021687. DOI: 10.1029/2021JB021687.

Yousefzadeh, R. and Ahmadi, M. (2023). Fast marching method assisted permeability upscaling using a hybrid deep learning method coupled with particle swarm optimization. *Geoenergy Sci. Eng.*, 230:212211. DOI: 10.1016/j.geoen.2023.212211.

Zhai, H., Xue, Z., Park, H., Aizawa, Y., Baba, Y., and Zhang, Y. (2020). Migration characteristics of supercritical CO2 microbubble flow in the Berea sandstone revealed by voxel-based X-ray computed tomography imaging analysis. *Journal of Natural Gas Science and Engineering*, 77:103233. DOI: 10.1016/j.jngse.2020.103233.

Zhang, Y., Jiang, F., and Tsuji, T. (2022). Influence of pore space heterogeneity on mineral dissolution and permeability evolution investigated using lattice Boltzmann method. *Chemical Engineering Science*, 247:117048. DOI: 10.1016/j.ces.2021.117048.

Zhang, Z., Wang, X., Niu, L., Wang, D., Li, Y., and Du, D. (2025). Unveiling the roles of sub-resolution pore structures on multiphase fluid flow characteristics in porous media via the super-resolution GAN and the CFD simulation. *Chemical Engineering Research and Design*. DOI: 10.1016/j.cherd.2025.08.009.

Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232. DOI: 10.1109/ICCV.2017.244.