PCP-GAN: Property-Constrained Pore-Scale Image Synthesis via Conditional Generative Adversarial Networks

Arash Rabbani, Lecturer in AI

Ali Sadeghkhani, PhD student

Brandon Bennett, Lecturer in Knowledge Representation

School of Computer Science, University of Leeds

Masoud Babaei, Senior Lecturer in Subsurface Energy Engineering, University of Manchester

InterPore 2026, Nantes, France

dataflowlab.org

Data Flow Lab

Laris Karklis/The Washington Post

The Pore-Scale Imaging Bottleneck

Characterising porous media at the pore scale is fundamental to predicting transport and multiphase flow, yet imaging is expensive and sample availability is inherently limited.

- Core plugs are recovered only at discrete depths.
- Diagenesis causes strong depth-dependent textural variability.
- Micro-CT and thin-section petrography yield sparse datasets.

The Key Question

Can a generative model synthesise statistically representative pore-scale images at prescribed petrophysical conditions?

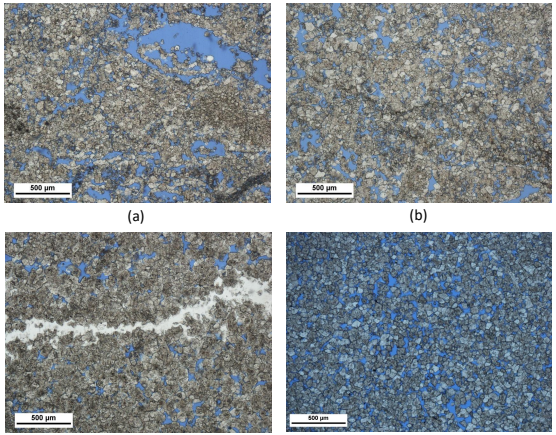

Figure: Thin-section petrography at four formation depths

A standard GAN generates samples from a learned distribution with no external control. A **conditional GAN (cGAN)** augments the generator G and discriminator D with a condition vector c encoding petrophysical targets.

Relevant condition variables:

- **Porosity (ϕ)**: continuous scalar.
- **Formation type**: discrete lithofacies label.
- **FVF / MRL**: fibre density and orientation (Chapter 2).

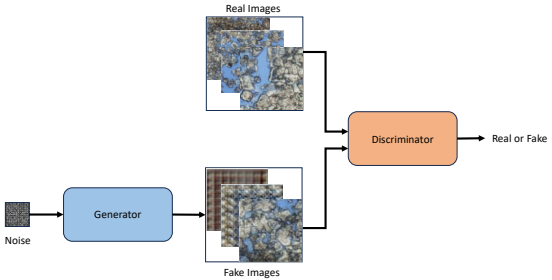

Figure: Standard GAN vs. conditional GAN architecture

Three-stage pipeline:

1. **Label** the dataset using domain-specific tools (U-Net segmentation, Rose diagrams for orientation).
2. **Train** a multi-conditional GAN conditioned on the labelled petrophysical properties.
3. **Deploy** for representative image generation at any specified petrophysical state.

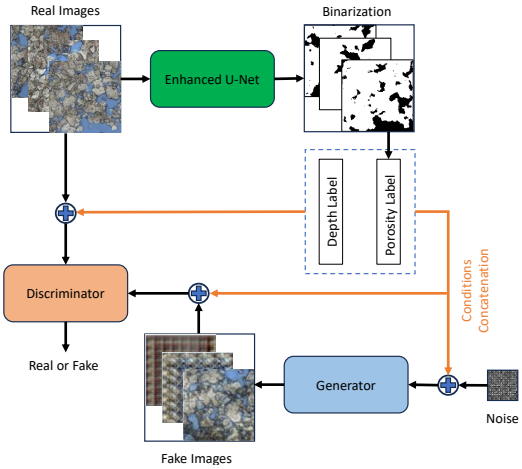

Figure: PCP-GAN workflow

PCP-GAN Architecture: Dual Conditioning on Porosity and Facies

Condition vector:

1. **Porosity** $\phi \in [0, 1]$ (continuous).
2. **Formation depth** (discrete, 4 lithofacies classes).

Network configuration:

- Generator G: 5 transposed convolutional layers, $\sim$50 M parameters.
- Discriminator D: 5 convolutional layers with dropout for regularisation.

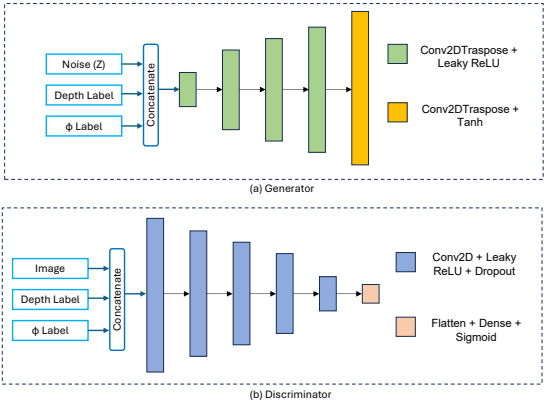

Figure: Generator and discriminator architectures

Pore Morphology Diversity across the Carbonate Sequence

Connected porosity is impregnated with blue epoxy for thin-section imaging, enabling colour-based segmentation by U-Net.

The dataset covers a wide spectrum of pore types: interparticle, intercrystalline, intraparticle, and vuggy, each with distinct pore-size distributions and shape factors.

Motivation for Conditioning on Depth

A single unconditional GAN cannot represent this lithofacies diversity; formation depth must be an explicit condition variable.

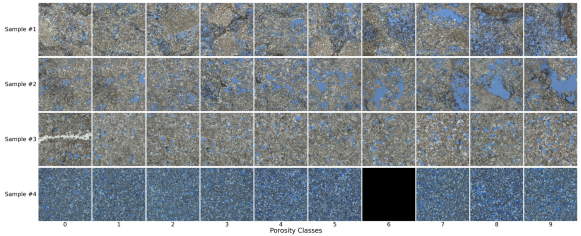

Figure: Diverse pore morphologies across the carbonate sequence

Adversarial training proceeds jointly across all four formations.

- Balanced dataset: 160 images per porosity class per formation.
- The adversarial loss and the ϕ-tracking error are monitored concurrently.

Physical Validation Metric

Convergence of the ϕ-tracking curve is the primary indicator of physically meaningful training progress; the adversarial loss alone is insufficient.

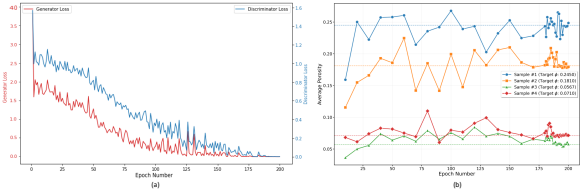

Figure: Loss curves and porosity-tracking during training

Quantitative validation: measured porosity of generated images versus the prescribed target ϕ.

- $R^2 = 0.95$ across all four formations.
- Mean absolute error MAE ≈ 0.01.

Significance

PCP-GAN generates rock textures at *any* specified porosity, enabling systematic population of the ϕ-space for upscaling studies.

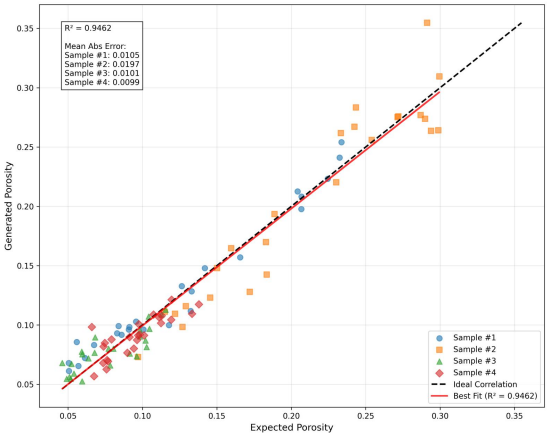

Figure: Target vs. measured porosity: fidelity analysis

3D Reconstruction: Pore Space at Prescribed Porosity

Each column shows a single ϕ target: the training 2D thin-section (top), the reconstructed 3D volume (middle), and the extracted pore space rendered as a 3D surface (bottom).

Porosity increases monotonically from left to right ($\phi = 0.057$ to 0.155), demonstrating continuous and well-controlled generation across the conditioning range.

Key Observation

Pore-space topology and grain-packing texture both evolve consistently with ϕ, confirming that the generator has learnt physically meaningful structure rather than superficial colour statistics.

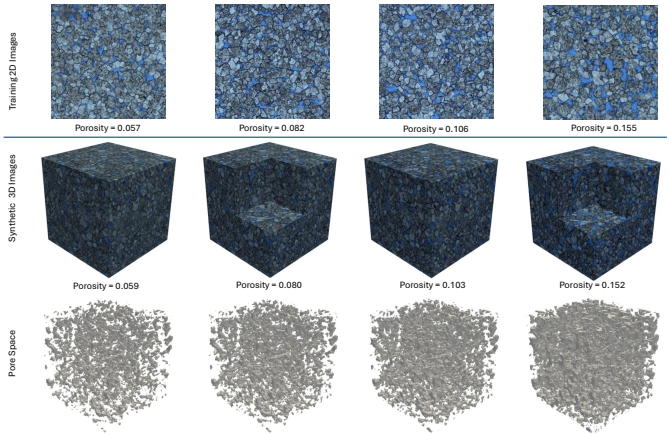

Figure: Training 2D images, synthetic 3D volumes, and extracted pore space at four prescribed porosities

Mineralogy Preservation: RGB Colour Fidelity

Pore-scale images carry mineralogical information beyond binary pore/grain segmentation:

- **White**: anhydrite cement patches.
- **Grey**: dolomite matrix.
- **Blue**: epoxy-impregnated porosity.

PCP-GAN is trained on RGB images, reproducing mineralogical assemblages rather than a binary pore mask alone.

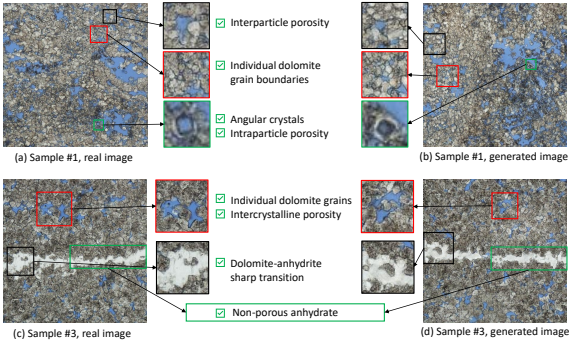

Figure: Real vs. synthetic: detailed mineral-phase comparison

Morphological Validation: Pore Geometry Metrics

Visual realism is necessary but not sufficient. The generated pore space must reproduce measurable morphological descriptors:

- Pore-size distribution
- Specific surface area
- Tortuosity

Agreement across all three confirms that the generative model captures pore geometry at multiple length scales, not merely global porosity.

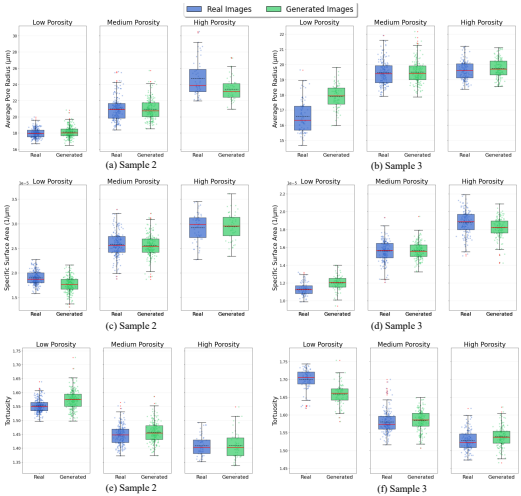

Figure: Morphological metrics: real vs. generated

Representativeness error: deviation of a sub-image's ϕ and K from the bulk core plug values.

- Random physical sub-images: error 37–700%.
- PCP-GAN generated images: error 2–12%.

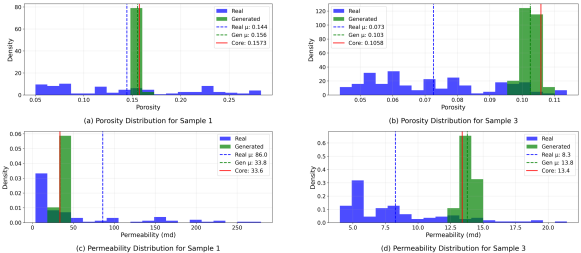

Figure: Porosity–permeability distributions: real sub-samples vs. generated

Upscaling Implication

Synthetically generated images are better representatives of the bulk core than randomly cropped physical sub-samples, reducing the REV requirement for pore-network modelling.

Is 50 M parameters necessary for carbonate textures?

- **Model B (25 M)**: fails to recover large-scale pore clusters and grain boundaries.
- **Model A (38 M)**: textural repetition artifacts, indicating insufficient capacity to represent the full pore-morphology diversity.

Remark

The heterogeneity of carbonate pore systems demands sufficient network depth and filter count; simpler architectures adequate for sandstone may not transfer.

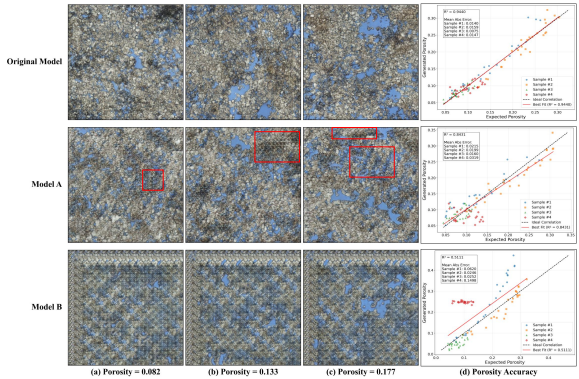

Figure: Ablation: effect of network capacity

The knee meniscus is a fibrous porous medium whose transport and mechanical properties are governed by collagen fibre architecture:

- Fibre volume fraction (FVF) controls solid-phase density.
- Fibre orientation (anisotropy) determines permeability directionality and mechanical anisotropy.
- μCT imaging is slow and donor tissue is scarce.

3D Challenge: Full 3D convolution on $512 \times 512 \times 512$ volumes is computationally intractable at training scale.

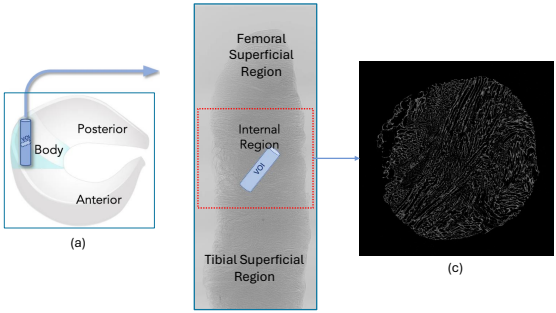

Figure: Meniscus anatomy and μCT cross-sections

Architectural innovation: represent a $512 \times 512 \times 32$ sub-volume as a single 32-channel 2D image, enabling standard 2D convolutions.

Inter-slice continuity loss penalises structural discontinuities between adjacent planes:

$$\mathcal{L}_{\text{cont}} = \sum_{z} \left| x_z - x_{z+1} \right|$$

This enforces pore-throat and fibre connectivity across the through-plane direction, which a purely 2D loss cannot capture.

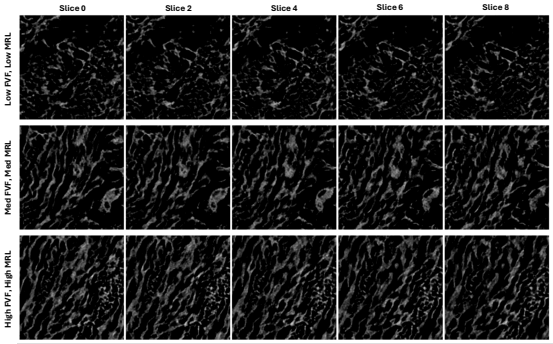

Figure: Inter-slice continuity profiles: real vs. generated

Condition Variables: Fibre Volume Fraction and Alignment

Two independent microstructural descriptors condition the generator, analogous to ϕ and formation type in Chapter 1:

- **FVF**: fibre volume fraction, governing bulk permeability magnitude.
- **MRL** (Mean Resultant Length): circular-statistics measure of fibre directional concentration, governing permeability anisotropy.

Independent control of fibre density and fibre alignment is essential for structure–property mapping.

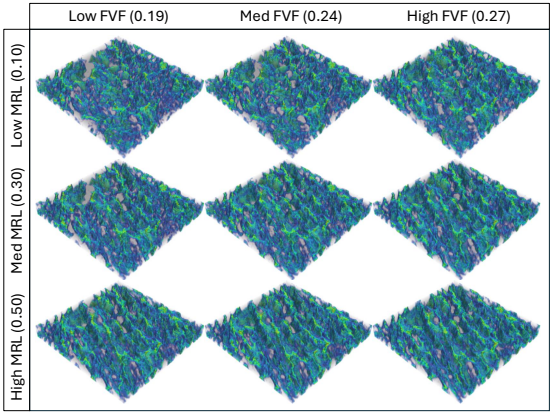

Figure: Conditional synthesis grid: FVF vs. MRL

Fidelity of prescribed vs. measured properties in generated volumes:

- FVF: $R^2 = 0.82$, MAE $= 0.014$.
- MRL: $R^2 = 0.88$, MAE $= 0.040$.

High accuracy is maintained across the full range of tissue heterogeneity, from isotropic low-density regions to highly aligned high-density zones near the tibial surface.

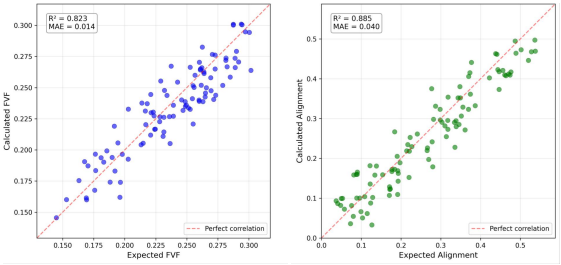

Figure: Expected vs. measured FVF and MRL

Phase-Space Coverage: No Mode Collapse

A generative model for structure–property studies must sample the full FVF–MRL phase space without mode collapse.

- Rare high-density (high FVF) configurations reproduced.
- Rare high-alignment (high MRL) configurations reproduced.
- Coverage matches the empirical distribution of the μCT dataset.

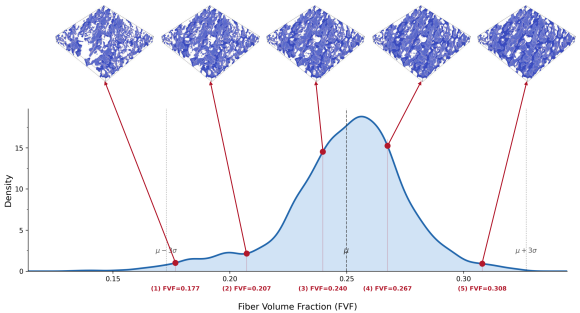

Figure: FVF distribution: real dataset vs. generated ensemble

Power-law scaling fitted across the generated ensemble:

$$K = 1.68\,\phi^{n}$$

with exponent n varying with fibre alignment:

- $n \approx 4.8$ (isotropic, low MRL).
- $n \approx 3.7$ (highly aligned, high MRL).

Physical Interpretation

Fibre alignment creates preferential flow pathways that reduce the sensitivity of K to ϕ at low porosities, flattening the scaling exponent.

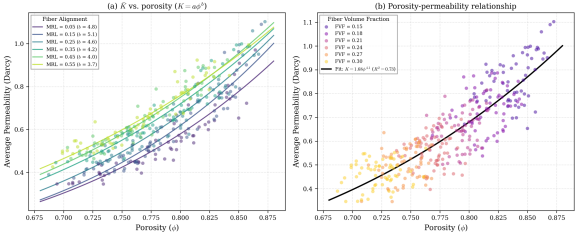

Figure: Porosity–permeability scaling laws across MRL range

Chapter 1: Carbonate Reservoir Characterisation

PCP-GAN synthesises petrographically realistic pore-scale images across four carbonate lithofacies with precision porosity control ($R^2 = 0.95$, MAE ≈ 0.01). Generated images are more representative of the bulk plug than random physical sub-samples, directly benefiting pore-network modelling and upscaling workflows.

Chapter 2: Fibrous Porous Media (Meniscus)

The slices-as-channels architecture enables tractable 3D synthesis with independent FVF/MRL control. Pore-network modelling of 360 generated volumes reveals a MRL-dependent shift in the porosity–permeability scaling exponent, a structure–property relationship inaccessible from the limited physical dataset alone.

Open Questions and Future Work

- Physics-informed loss terms (Stokes, Darcy) to enforce flow consistency during training.
- Extension to wettability and multiphase saturation as additional condition variables.
- Transfer learning across lithofacies to reduce labelled data requirements.

Thank you

Questions and discussion

University of Leeds
Data Flow Lab Research Group
www.DataFlowLab.org
a.rabbani@leeds.ac.uk

InterPore 2026, Nantes, France

To condition G on ϕ, ground-truth porosity labels must be computed reliably from RGB thin-section images.

- An **enhanced U-Net** segments blue-epoxy pores from the mineral matrix (Dice coefficient $= 0.95$).
- Patch size is determined via **Representative Elementary Volume (REV)** analysis: the minimum patch area at which ϕ stabilises.

REV as Training Unit

The REV defines the natural patch size for training: sub-REV patches yield biased porosity labels and degrade conditioning accuracy.

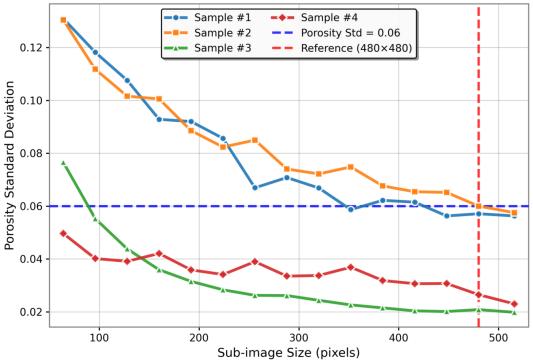

Figure: REV analysis for patch size selection

Visual Fidelity across Lithofacies

Real vs. synthetic comparison at two representative depths:

- **Grainstone (1879 m)**: inter-particle porosity and grain boundary morphology reproduced.
- **Microcrystalline dolomite (1918 m)**: anhydrite inclusions and intercrystalline pore network reproduced.

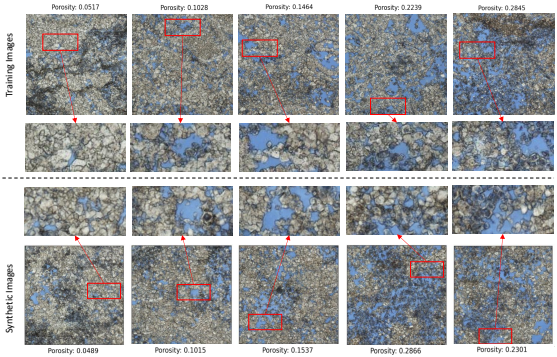

Figure: Sample reconstruction: Formation 1879 m

Two-point correlation $S_2(r)$ is a well-established microstructure descriptor sensitive to fibre diameter, inter-fibre spacing, and clustering.

Validation Result

Generated $S_2(r)$ curves overlap with the real curves across all FVF–MRL conditioning combinations, confirming correct pore-throat geometry and connectivity.

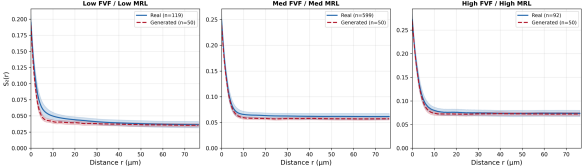

Figure: Two-point correlation: real vs. generated

How are 3D patches assigned ground-truth condition labels?

- **FVF**: fibre voxels are segmented in each z-slice; the mean fraction defines the patch label.
- **MRL**: fibre orientation angles are extracted via gradient decomposition; the resultant length of the angular distribution is computed per patch.

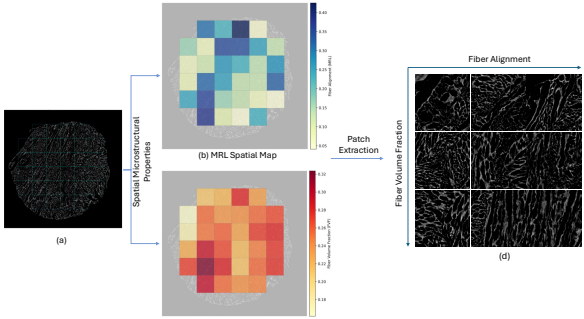

Figure: Spatial maps of FVF and MRL across the tissue