

Eksploatacja i Niezawodnosc - Maintenance and Reliability

Volume 29 (2027), Issue 3

https://ein.org.pl

The Polish Maintenance Society

Reliable 3D EIT reconstruction for masonry condition monitoring using conditional diffusion and cascaded FEM-space refinement

Grzegorz Kłosowski[1], Anna Hola[2], Monika Kulisz[1,*], Tomasz Rymarczyk[3,4], Jan Sikora[4,3], Konrad Niderla[4,3], Łukasz Skowron[1]

[1] Lublin University of Technology, Poland
[2] Wrocław University of Science and Technology, Poland
[3] Netrix S.A., Poland
[4] WSEI University, Poland
* Corresponding Author: m.kulisz@pollub.pl

Received: 27 July 2026
Revised: 25 September 2026
Accepted: 2 October 2026
Online: 13 November 2026

This is an open access article under the CC BY 4.0 license

Abstract

Three-dimensional electrical impedance tomography (EIT) reconstruction of moisture-related conductivity in masonry requires preservation of both global structure and local finite-element variations. This study proposes a conditional denoising diffusion probabilistic model with cascaded FEM-space residual refiners (cDDPM-SFR). A 448-component voltage vector conditions a 3D U-Net on a 16 × 16 × 16 voxel grid, followed by three multilayer perceptron (MLP) stages producing 3050 outputs. Training used 49,000 synthetic cases and evaluation 1,000. Compared with MLP, mean squared error (MSE) decreased from 0.119 to 0.048, peak signal-to-noise ratio (PSNR) rose from 21.692 to 26.961 dB, structural similarity index (SSIM) from 0.978 to 0.992, and Pearson correlation coefficient from 0.981 to 0.992 (Holm-adjusted $p < 0.001$). Four field acquisitions showed correlations of 0.196–0.935 and mean median absolute deviation of 0.689. Independent dielectric profiles and the EIT surface proxy identified the same moisture-related zone near 360 cm, supporting physical plausibility.

Keywords: *electrical impedance tomography, masonry condition monitoring, conditional diffusion model, residual refinement, inverse problem*

Article citation:
Kłosowski G, Hola A, Kulisz M, Rymarczyk T, Sikora J, Niderla K, Skowron Ł, Reliable 3D EIT reconstruction for masonry condition monitoring using conditional diffusion and cascaded FEM-space refinement, Eksploatacja i Niezawodnosc – Maintenance and Reliability 2027: 29(3) http://doi.org/10.17531/ein/242319

Highlights

- cDDPM-SFR pipeline combines conditional diffusion with cascaded FEM-space refinement.
- A conditional diffusion model reconstructs global 3D conductivity patterns.
- Consensus reconstruction limits acquisition-specific variation in real EIT data.

1. Introduction

1.1. Electrical impedance tomography for condition assessment of masonry structures

Moisture-related deterioration is a recurrent maintenance problem in masonry structures because water ingress activates several interacting degradation mechanisms. Reliable identification of moisture-affected regions is therefore required not only to assess the current condition of a wall, but also to determine the scope and priority of maintenance interventions [1]. Among the methods available for non-destructive condition assessment, electrical impedance tomography can be particularly useful when information on the spatial extent of moisture-related changes is required.

Electrical impedance tomography provides spatial information on the internal electrical properties of a structure from measurements performed at its boundary [2]. During one acquisition sequence, low-amplitude alternating current is injected through selected pairs of surface electrodes, while the resulting voltage differences are recorded between selected remaining electrode pairs. The measured voltages depend on the conductivity distribution within the investigated domain and are related to it through a forward model, usually formulated using the finite element method (FEM) [3]. Reconstruction of the internal conductivity field requires the solution of the corresponding inverse problem, which is nonlinear, ill-posed, and sensitive to measurement noise, electrode contact conditions, and inaccuracies in the assumed geometry [4].

The diagnostic contrast obtained in masonry measurements

is associated with changes in the electrical properties of the pore system. Water containing dissolved ions creates conductive paths through bricks, mortar, and other porous building materials; therefore, an increase in water saturation is generally accompanied by an increase in effective electrical conductivity [5,6]. This relationship is not determined by moisture content alone. Evidence from dielectric moisture measurements also indicates that salt contamination can lead to an overestimation of moisture when material-specific calibration is not applied [7]. The measured response is also affected by pore structure, material composition, temperature, and local heterogeneity between bricks and mortar. For this reason, an EIT reconstruction should primarily be interpreted as a spatial distribution of electrical conductivity associated with moisture-related changes rather than as a direct map of absolute moisture content.

The diagnostic value of tomography also depends on the reliability of the complete measurement and data-processing chain. More generally, studies in non-destructive condition monitoring have shown that both measurement-system parameters and the integration of sensing with signal processing affect signal quality and the reliability of diagnostic decisions [8,9]. In process tomography, electrical measurements combined with appropriate data-processing methods have supported real-time fault detection in reactive crystallisation and machine-learning-based classification of two-phase gas–liquid flow regimes [10,11]. Although these applications do not concern masonry inspection, they illustrate how electrical measurement data can be transformed into condition indicators supporting operational decisions.

Complementary studies in flow metrology and conductive sensing systems have demonstrated that measurement accuracy depends on error-correction procedures [12], sensing-element geometry [13], mechanical integrity and data-acquisition architecture [14,15]. These findings are relevant to field EIT because variations in electrode contact, sensing configuration and signal acquisition may propagate into the reconstructed conductivity field.

1.2. Data-driven reconstruction and representation challenges in three-dimensional EIT

The practical value of electrical impedance tomography in masonry condition assessment depends largely on the reliability of the reconstruction procedure. Boundary voltage measurements do not provide a direct image of the internal conductivity field. Instead, the conductivity distribution must be inferred by solving a nonlinear inverse problem in which different internal configurations may produce similar electrical responses. The problem is additionally underdetermined because the number of reconstructed parameters is usually much larger than the number of available measurements [3,16]. In the system considered in this study, one acquisition cycle provides a 448-component voltage vector, whereas the reference conductivity field is defined over 3050 tetrahedral finite elements. Consequently, each measurement vector must be transformed into an output space almost seven times larger than the input space.

Classical reconstruction algorithms address the instability of the inverse problem by introducing regularisation. Tikhonov regularisation, regularised Gauss–Newton methods, and total-variation approaches restrict the admissible solution by imposing assumptions concerning its magnitude, smoothness, or spatial sparsity [3,17]. These methods provide stable reconstructions for many EIT configurations, but their performance depends on the regularisation parameter and on whether the assumed prior adequately represents the actual conductivity distribution. Tikhonov-type constraints may produce excessively smooth reconstructions and reduce the contrast of material interfaces, whereas conventional total-variation regularisation may introduce staircase artefacts and piecewise-constant regions [17,18]. Such effects are undesirable in masonry imaging, where moisture-related conductivity changes may form gradual transitions and occur against a heterogeneous background determined by bricks, mortar joints, cavities, repairs, and local salt accumulation. Consequently, a single predefined regularisation model may not adequately represent all conductivity patterns encountered during structural assessment.

Machine-learning methods provide an alternative to classical EIT reconstruction by learning the relationship between boundary voltages and conductivity distributions from numerical or experimental data. Fully learned models, including multilayer perceptrons, convolutional networks, and recurrent architectures, can reconstruct conductivity fields without

iterative optimisation for each measurement, which is advantageous in repeated inspections and monitoring applications [4,19]. Most direct regression models, however, generate a single deterministic output and do not explicitly represent alternative conductivity distributions consistent with the measured voltages [20].

For three-dimensional masonry imaging, a multi-sequential long short-term memory (LSTM) model was developed to process several disturbed measurement vectors, thereby improving reconstruction accuracy and robustness to noise [21]. The influence of input representation was further investigated using variable-frequency measurement sequences, showing that frequency-dependent information can improve reconstruction quality and robustness [22]. A multi-head architecture was subsequently proposed to divide the 3050-element FEM output vector into smaller groups reconstructed by separate branches connected to a shared representation of the 448-component measurement vector [23]. This reduced the complexity of each reconstruction task and improved the recovery of local conductivity variations.

Reconstruction accuracy is spatially non-uniform. It has been shown that the effectiveness of three-dimensional electrical tomography depends on both anomaly location and sensing modality [24]. This finding suggests that global quality indicators may conceal local reconstruction errors, particularly in planar wall measurements. Therefore, evaluation should also consider interface preservation and the occurrence of false local extrema.

Diffusion models formulate EIT reconstruction as an iterative conditional denoising process rather than a one-step regression task. DiffusionEIT and subsequent voltage-conditioned approaches demonstrated that measurement-guided denoising can improve image continuity and reduce reconstruction artefacts [20,25]. Zhang et al. proposed DiffEIT, in which a conditional DDPM generates homogeneous, inhomogeneous, and differential sensitivity priors that are subsequently used by a learned reconstruction algorithm [26]. Their approach incorporated measurement-dependent conditioning and reduced the reconstruction procedure to 25 diffusion steps.

More recent methods have introduced initial reconstructions and physical consistency constraints into the diffusion process [27]. These studies confirm the usefulness of measurement- and physics-informed conditioning, but they primarily address prior generation, denoising, and voltage consistency. In the reviewed diffusion-based EIT studies, comparatively limited attention has been given to the loss of element-level information caused by mapping an irregular three-dimensional FEM mesh onto a regular voxel grid. Existing studies have focused primarily on measurement conditioning, denoising performance, and physical consistency, while the influence of the spatial representation itself has been discussed less extensively [20,25].

A conditional diffusion model can act as a learned prior in the reconstruction process. Spatial patterns represented in the training data guide the reverse denoising process towards conductivity fields that are statistically plausible with respect to the learned distribution [20,25]. In this sense, explicit analytical regularisation can be supplemented or partly replaced by an implicit, data-driven prior. Reconstruction quality nevertheless depends on the representativeness of the training set, and distribution mismatch may produce artefacts for previously unseen conductivity patterns or measurement disturbances [28].

In three-dimensional EIT, convolutional diffusion models introduce an additional representation challenge. Conventional 3D U-Net architectures operate on regular volumetric grids, whereas EIT forward and inverse models are commonly formulated on irregular finite-element meshes [29,30]. The application of conventional dense 3D convolutional architectures generally requires the FEM conductivity field to be converted into a regular voxel representation. When several finite elements are assigned to the same voxel, their values may be aggregated, for example through averaging, which may remove local element-level differences before the data are processed by the diffusion network.

Similar limitations of voxel-based representations have been reported in three-dimensional learning, where discretisation may cause quantisation loss and discard geometric information retained by the original mesh [31,32]. This representation loss is distinct from the ill-posedness of the EIT inverse problem. The latter originates from incomplete and noisy boundary measurements, whereas voxelisation loss is introduced by the numerical representation used by the reconstruction model. Increasing the capacity of the U-Net or the number of diffusion steps cannot restore information that has already been discarded

during preprocessing.

A finer voxel grid can reduce the number of finite elements merged into a single voxel, but it also increases memory requirements and computational cost, since the number of volumetric cells grows cubically with grid resolution [32]. Direct processing of the tetrahedral mesh would avoid voxelisation, but it requires specialised graph- or mesh-based operations rather than conventional three-dimensional convolutions [29,30]. These considerations suggest that voxel-domain diffusion models may preserve global spatial organisation while remaining limited in their ability to recover local element-level variations lost during mesh-to-voxel conversion.

1.3. Motivation, novelty and paper structure

To address voxelisation-induced information loss, this study proposes cDDPM-SFR, which combines measurement-conditioned diffusion reconstruction on a regular voxel grid with cascaded residual refinement in the native FEM space. The diffusion stage provides the global spatial reconstruction, whereas three successive FEM-space refiners use the original EIT measurements to correct local element-level errors after voxel-to-FEM mapping. The main novelty therefore lies in combining probabilistic volumetric reconstruction with dedicated post-reconstruction correction in the original computational domain.

The approach is evaluated on synthetic data in terms of reconstruction accuracy and structural fidelity, and on repeated field measurements in terms of repeatability and physical plausibility using independent dielectric measurements as an external spatial reference.

The objective of the study was to evaluate whether the proposed cDDPM-SFR architecture improves three-dimensional EIT reconstruction while retaining physically plausible spatial patterns in field measurements. Accordingly, two research questions were formulated. RQ1 examined whether cDDPM-SFR improves synthetic reconstruction accuracy relative to the direct MLP baseline and whether successive FEM-space refinement reduces the remaining element-level reconstruction error. RQ2 examined whether repeated field reconstructions preserve a common dominant spatial organisation and whether the resulting consensus reconstruction is spatially consistent with independent dielectric measurements. H1 hypothesised that cDDPM-SFR would achieve significantly better reconstruction accuracy than the direct MLP baseline and that successive FEM-space refinement would provide incremental improvement. H2 hypothesised that repeated field reconstructions would preserve a common dominant spatial pattern and that the consensus reconstruction would identify a moisture-related zone consistent with the independent dielectric profiles.

The remainder of the paper is organised as follows. Section 2 describes the measurement system, data representation, cDDPM-SFR architecture, experimental protocol, and evaluation procedure. Section 3 presents the synthetic and real-measurement reconstruction results. The Discussion evaluates the agreement between the consensus EIT reconstruction and the dielectric measurements, and the final section summarises the main conclusions and limitations.

2. Material and methods

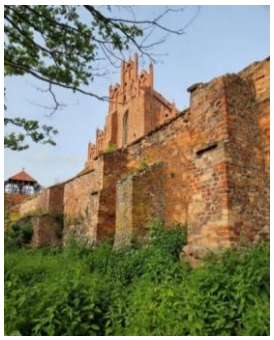
2.1. Research object

The object of the research was a preserved section of the defensive wall in Starogard Gdański, erected in the first half of the 14th century (Figure 1a). The wall consists of a stone foundation and an above-ground section constructed mainly of solid ceramic brick with lime mortar; its thickness ranges from approximately 150 to 200 cm. The construction of the wall is based on a stone foundation with an offset, built of granite erratic boulders of varying sizes, bonded with masonry mortar. The higher, above-ground parts of the fortifications were built of solid ceramic brick on a lime mortar [33]. Currently, the investigated preserved section of the wall is about 130 m long and runs along the Wierzyca River. It performs a crucial structural function – it acts as a retaining wall for the escarpment on which St. Matthew's Church and other buildings are founded (the difference in ground levels on both sides of the wall is over 4 m).

A section located near the belfry of St. Matthew's Church, directly above the stone foundation of the wall, was selected as the measurement area (Figure 1b). This choice was dictated by the good accessibility of the area and the presence of the largest amount of preserved, original brick material in this location.

Electrical impedance tomography measurements were

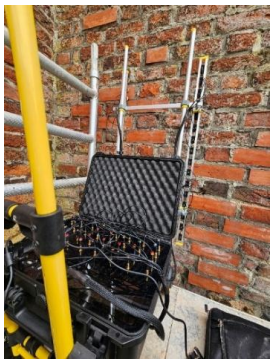
acquired using a dedicated portable EIT system designed and manufactured in the Netrix S.A. laboratory. The measurement setup, shown in Figure 1c, comprised a tomographic acquisition unit, a computer for system control and data recording, and two linear electrode bars attached to the investigated masonry surface. The measurement system consisted of two linear electrode strips, each containing 16 electrodes, giving a total of 32 electrodes attached to the tested masonry surface. The device integrated current excitation, voltage acquisition, electrode switching, and measurement-sequence control, enabling automated in situ operation.



(a)



(b)



(c)

Figure 1. A preserved section of the defensive wall of Starogard Gdański: a) view from the Wierzyca River; b) view of the measurement point on the wall, c) EIT measurement station.

Data acquisition was performed using an adjacent-electrode protocol. Current was sequentially injected through neighbouring electrode pairs, while voltage differences were measured between successive electrodes. Owing to the symmetry of the linear arrangement, 16 current-injection configurations were used, with 28 independent voltage readings acquired for each configuration. Consequently, one complete measurement sequence consisted of 448 voltage values. The sequences were repeated to reduce the influence of random disturbances and improve measurement repeatability. The resulting 448-component vectors were subsequently used as input data for the reconstruction model.

2.2. Data representation and preprocessing

In the proposed reconstruction procedure, 448-component voltage sequence was used as the conditioning information for the generative inverse model. The reconstructed quantity was the internal electrical conductivity distribution associated with moisture-dependent changes inside the masonry. The EIT values were interpreted as relative conductivity indicators rather than direct measurements of absolute moisture content. Therefore, the resulting images were used to identify spatial gradients and zones of increased electrical response, which may correspond to regions of elevated moisture.

The reconstruction model was trained using synthetic EIT datasets generated from a finite-element forward model of the investigated masonry geometry. The available numerical dataset consisted of 49,000 training cases and 1,000 validation cases. For each case, the input consisted of a 448-value voltage vector, whereas the reference output was a conductivity distribution assigned to 3050 tetrahedral finite elements of the computational mesh.

Because the diffusion model uses 3D convolutions, FEM targets were mapped to a $16 \times 16 \times 16$ voxel grid using element-centroid coordinates. The 3050 FEM elements occupied 1191 voxels; when multiple elements mapped to the same voxel, their conductivity values were averaged. A binary field-of-view mask excluded inactive voxels from the loss. Conductivity targets were normalised to the model range, with -1 assigned outside the FEM domain, and voltage vectors were standardised.

2.3. Conditional DDPM with cascaded FEM refiners

The inverse EIT problem was solved using the conditional

DDPM with cascaded FEM-space residual refiners (cDDPM-SFR) — a two-stage pipeline in which the cDDPM stage provides an initial volumetric reconstruction, followed by a cascaded FEM-space refiner that corrects systematic errors in the native finite-element domain (Figure 2).

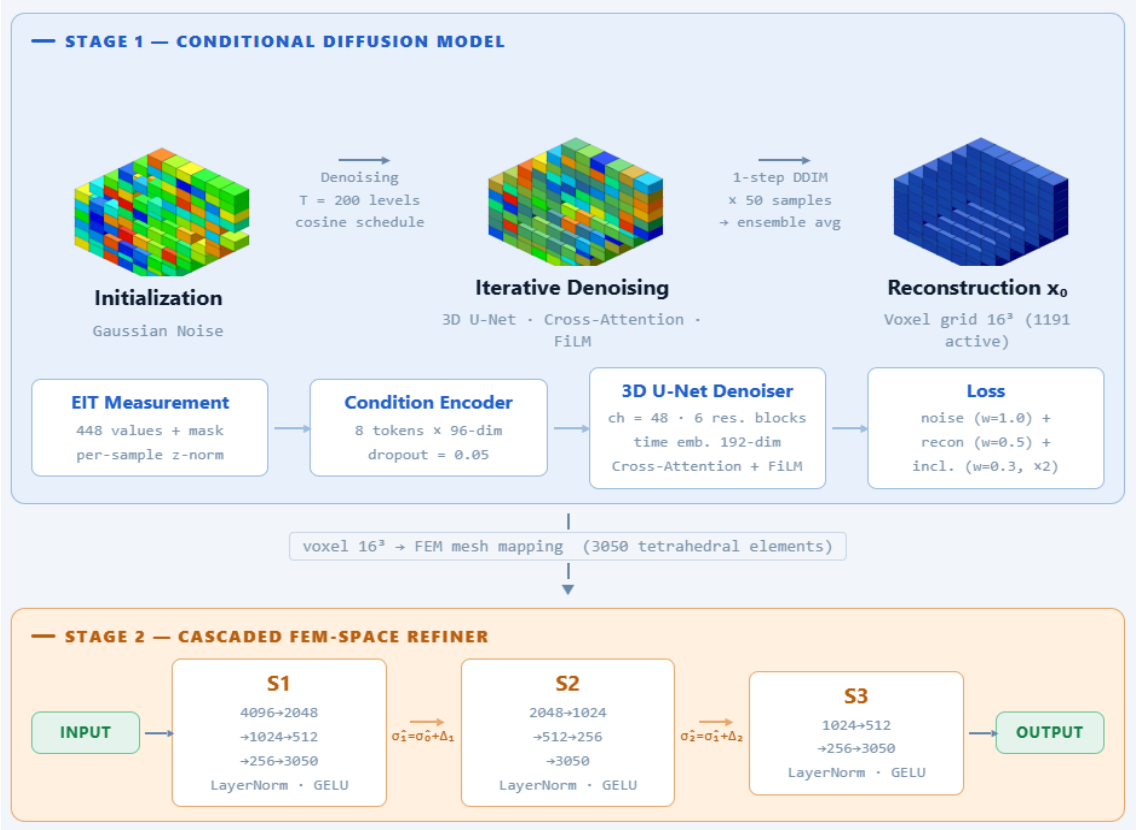


Figure 2. Schematic overview of the proposed conditional diffusion reconstruction pipeline.

In contrast to direct regression methods, which predict the full conductivity image in a single step, the diffusion-based approach reconstructs the image through an iterative denoising process. During training, Gaussian noise is gradually added to the reference conductivity volumes, and the neural network learns to estimate the noise component present at a given diffusion step. During inference, the process is reversed: the reconstruction starts from random noise and is progressively denoised under the control of the EIT voltage measurements.

The diffusion process used 200 noise levels with a cosine noise schedule. This schedule provides a smooth transition between weakly and strongly corrupted states and improves the stability of learning across different noise intensities. At each training iteration, a random diffusion step was selected, the corresponding noisy conductivity volume was generated, and the model was trained to predict the added noise. In this way, the network learned a conditional reverse process that links the measured voltage sequence with a plausible internal conductivity distribution.

The denoising network was implemented as a conditional 3D U-Net receiving the noisy conductivity volume, diffusion-step encoding, and EIT conditioning vector. It used 48 base channels and six residual blocks. Diffusion time was represented by a 192-dimensional embedding, while the 896-dimensional EIT input, comprising 448 voltage values and a 448-element validity mask, was encoded into twelve 96-dimensional conditioning tokens. Measurement conditioning was introduced through cross-attention and Feature-wise Linear Modulation (FiLM), with a dropout rate of 0.05. The network used 3D convolutions, group normalisation, GELU activations, residual blocks, and encoder–decoder skip connections.

Training used a composite objective function combining noise-prediction, clean-volume reconstruction, and inclusion-focused losses (1):

$$L_{\text{total}} = \lambda_{\varepsilon} L_{\varepsilon} + \lambda_{\text{rec}} L_{\text{rec}} + \lambda_{\text{incl}} L_{\text{incl}}, \quad (1)$$

where L_{ε}measures the error in predicting the diffusion noise, L_{rec} measures the reconstruction error of the predicted clean conductivity volume, and L_{incl} gives additional emphasis to

voxels belonging to conductivity inclusions. For the model reported in this study, the corresponding weights were $\lambda_\varepsilon = 1.0$, $\lambda_{\mathrm{rec}} = 0.5$, and $\lambda_{\mathrm{incl}} = 0.3$. All loss components were evaluated within the field-of-view mask so that inactive voxels outside the finite-element domain did not contribute to optimisation.

The individual components were defined as:

$$L_\varepsilon = \frac{\sum_v M_v |\varepsilon_v - \hat{\varepsilon}_v|}{\sum_v M_v} \quad (2)$$

$$L_{\mathrm{rec}} = \frac{\sum_v M_v |x_{0,v} - \hat{x}_{0,v}|}{\sum_v M_v} \quad (3)$$

$$L_{\mathrm{incl}} = \frac{\sum_v M_v I_v w_{\mathrm{incl}} |x_{0,v} - \hat{x}_{0,v}|}{\sum_v M_v I_v},\ w_{\mathrm{incl}} = 2 \quad (4)$$

where x_0 and $\hat{x}_0$ denote the reference and predicted clean conductivity volumes, ε and $\hat{\varepsilon}$ denote the true and predicted diffusion noise, M_v is the field-of-view mask, and I_v identifies voxels belonging to conductivity inclusions.

The complete training objective was therefore:

$$L_{\mathrm{total}} = 1.0 L_\varepsilon + 0.5 L_{\mathrm{rec}} + 0.3 L_{\mathrm{incl}} \quad (5)$$

The network was optimised using the AdamW optimiser. The maximum number of training epochs was 300 with a mini-batch size of 4. Gradient accumulation over two mini-batches was applied, giving an effective batch size of 8. Small Gaussian perturbations with a standard deviation of 0.01 were added to the voltage-conditioning vector during training to improve robustness to measurement variability. The validation loss was monitored on an independent validation set and early stopping was applied with a patience of 20 epochs.

Inference used one-step denoising diffusion implicit model (DDIM) sampling with η = 0 and no classifier-free guidance. It is deterministic for a fixed initial noise field; 50 reconstructions used different initial noise realisations and were averaged before FEM-space refinement.

Reconstruction quality as a function of the number of DDIM sampling steps was evaluated experimentally and is reported together with the other quantitative results in Section 3.1.

Since the 3D U-Net operates on a regular 16×16×16 voxel grid, the mapping from the finite-element mesh (3050 tetrahedral elements) to voxel space introduces an irreversible information loss that limits the achievable reconstruction fidelity. To overcome this bottleneck, a cascaded FEM Refiner was applied as a post-processing stage operating directly in the native 3050-element FEM space.

Three successive MLP networks were trained sequentially (S1→S2→S3), each predicting a residual correction to the output of the previous stage. The input to each stage consisted of the EIT measurement vector (448 elements) concatenated with the current FEM-space prediction (3050 elements), giving a 3498-dimensional input. The output was a 3050-dimensional residual that was added to the current prediction. Stage S1 used a network with hidden layers of sizes 4096, 2048, 1024, 512, and 256 units with LayerNorm and GELU activations. Stages S2 and S3 used progressively smaller architectures with hidden layers of sizes 2048, 1024, 512, 256 and 1024, 512, 256 respectively. Output layer weights were initialised to zero so that each refiner started as an identity mapping. Training of each stage used an inclusion-weighted MSE loss, AdamW optimiser with linear learning-rate warmup, and early stopping. Input augmentation in the form of additive Gaussian noise and random dropout applied to the diffusion ensemble predictions was used during training of Stage S1 to improve generalisation.

The complete cDDPM-SFR inference mapping can be expressed as:

$$\overline{\mathbf{x}_V} = \frac{1}{J}\sum_{j=1}^{J} G_\theta(\mathbf{z}_T^{(j)}, C_\psi(\tilde{\mathbf{v}}, \mathbf{m})), \quad \hat{\mathbf{y}}^{(0)} = P(\overline{\mathbf{x}_V}) \quad (6)$$

$$\hat{\mathbf{y}}^{(k)} = \hat{\mathbf{y}}^{(k-1)} + R_{\varphi_k}([\tilde{\mathbf{v}}, \hat{\mathbf{y}}^{(k-1)}]), \quad k \in \{1,2,3\} \quad (7)$$

$$\hat{\mathbf{y}}_{cDDPM-SFR} = \hat{\mathbf{y}}^{(3)} \quad (8)$$

where G_θ denotes the conditional diffusion generator, C_ψ the measurement-conditioning encoder, $\tilde{\mathbf{v}} \in \mathrm{R}^{448}$ the standardised EIT voltage vector, $\mathbf{m}$ the corresponding measurement-validity mask, and $\mathbf{z}_T^{(j)} \sim \mathrm{N}(0, \mathrm{I})$ the j-th initial noise realisation. The operator P maps the averaged voxel-domain reconstruction onto the 3050-element FEM mesh, while R_{φ_k} denotes the residual predicted by the k-th FEM-space refiner. In the reported configuration, J = 50.

Formally, for stage $k \in \{1,2,3\}$ each refiner was trained by minimising an inclusion-weighted mean squared error over the $N = 3050$ finite elements:

$$L_{S_k} = \frac{1}{N}\sum_{e=1}^{N} w_e \left(y_e - \hat{y}_e^{(k)}\right)^2, \quad w_e = 1 + (\beta - 1)\cdot \mathbb{1}[y_e > \tau] \quad (9)$$

where y_e is the reference conductivity of element e, $\hat{y}_e^{(k)}$ the prediction of stage k, τ = 1.80 the inclusion threshold (20% of the training-target range above the background value BG = 1.0), and β = 5.0 the inclusion-weighting factor. The same threshold and weighting factor were used for all three refinement stages (S1, S2, S3).

As a reference method (baseline model), a fully connected multilayer perceptron (MLP) was trained to directly map the EIT voltage measurement vector to the finite-element conductivity distribution. The network architecture consisted of five hidden layers with 1024, 1024, 512, 512, and 256 units respectively, followed by a linear output layer of 3050 units. Each hidden layer applied batch normalisation, a ReLU nonlinearity, and dropout regularisation with rates of 0.30 and 0.20 for the first and second hidden layers respectively.

A Tikhonov-regularised linear reconstruction was additionally used as a classical reference method. The regularisation parameter was selected by five-fold cross-validation on the training data, and the final reconstruction operator was evaluated on the same 1,000 validation cases as the learned models.

To quantitatively evaluate the reconstruction performance of the generative model, complementary metrics capturing both pointwise errors and structural fidelity were utilized. The evaluation was based on the Mean Squared Error (MSE), Mean Absolute Error (MAE), Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and the Pearson Correlation Coefficient (PCC). These metrics are defined as follows. Let $\hat{\boldsymbol{\sigma}} = [\hat{\sigma}1, \dots, \hat{\sigma}N]^{\top}$ denote the reconstructed conductivity vector and $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_N]^{\top}$ the reference vector, both of length $N = 3050$.

$$\mathrm{MSE} = \frac{1}{N}\sum_{i=1}^{N} (\hat{\sigma}_i - \sigma_i)^2 \tag{10}$$

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N} |\hat{\sigma}_i - \sigma_i| \tag{11}$$

$$\mathrm{PSNR} = 10\log_{10}\left(\frac{DR^2}{\mathrm{MSE}}\right),\ DR = \max(\boldsymbol{\sigma}) - \min(\boldsymbol{\sigma}) \tag{12}$$

$$\mathrm{PCC} = \frac{\sum_{i=1}^{N} (\hat{\sigma}_i - \bar{\hat{\sigma}})(\sigma_i - \bar{\sigma})}{\sqrt{\sum_{i=1}^{N} (\hat{\sigma}_i - \bar{\hat{\sigma}})^2 \cdot \sum_{i=1}^{N} (\sigma_i - \bar{\sigma})^2}} \tag{13}$$

$$\mathrm{SSIM} = \frac{(2\mu_{\hat{\sigma}}\mu_{\sigma} + C_1)(2\mathrm{cov}(\hat{\sigma}, \sigma) + C_2)}{(\mu_{\hat{\sigma}}^2 + \mu_{\sigma}^2 + C_1)(s_{\hat{\sigma}}^2 + s_{\sigma}^2 + C_2)} \tag{14}$$

where $\mu\hat{\sigma}, \mu\sigma$ are the means, $s\hat{\sigma}^2, s\sigma^2$ the variances, and $\mathrm{cov}(\hat{\boldsymbol{\sigma}}, \boldsymbol{\sigma})$ the covariance of the reconstructed and reference vectors; the stability constants are $C_1 = (0.01 \cdot DR)^2$ and $C_2 = (0.03 \cdot DR)^2$, following Wang et al. [34].

In line with the GREIT/EIDORS perspective on EIT image assessment, PCC is interpreted as a measure of vertical moisture-gradient consistency, while SSIM reflects the localization fidelity and the structural preservation of the moisture-affected zones. Simultaneously, MSE and PSNR quantify the absolute pointwise reconstruction errors and the signal-to-noise ratio.

The primary statistical comparison between the MLP baseline and the final S3 reconstruction was performed for MSE, PSNR, SSIM, and PCC using the complete set of 1,000 paired validation cases. For each of these four metrics, the uncertainty of the mean was quantified using a percentile bootstrap with 10,000 paired case-level resamples (random seed 4444). The same case indices were retained for both compared models in each bootstrap sample, thereby preserving the paired structure of the validation data. Differences between the MLP baseline and the final S3 reconstruction were additionally assessed using two-sided Wilcoxon signed-rank tests. To account for multiple comparisons across the four primary metrics, p-values were adjusted using the Holm procedure. FEM elements within individual reconstructions were not treated as independent observations.

2.4. Experimental validation protocol for real masonry measurements

The applicability of the trained cDDPM-SFR model to experimental data was evaluated using four complete EIT acquisition sequences obtained from the investigated historic masonry wall. Three sequences were acquired with an excitation current of 300 μA, whereas one sequence was acquired at 400 μA. Each acquisition retained the complete 448-position measurement structure required by the reconstruction model. The repeated 300 μA sequences were used to assess short-term measurement and reconstruction repeatability, while the 400 μA sequence provided an additional observation at a higher excitation level.

Each voltage vector was processed independently by the cDDPM and the three successive FEM-space refinement stages, S1, S2, and S3. This produced four final conductivity fields representing repeated observations of the same wall region. Pairwise comparisons between the reconstructed fields were used to assess reconstruction repeatability. In addition, an elementwise median of the four S3 fields was calculated to obtain a consensus reconstruction less sensitive to disturbances occurring in individual acquisition sequences.

Formally, let $x_e^{(m)}$, m = 1,…,4, denote the S3-stage

conductivity value at FEM element e (e = 1,…,3050) obtained from the m-th acquisition (three sequences at 300 μA and one at 400 μA). The consensus reconstruction was defined as the elementwise median across acquisitions,

$$\hat{x}_e = \text{median}\left(x_e^{(1)}, x_e^{(2)}, x_e^{(3)}, x_e^{(4)}\right), \quad e = 1, \ldots, 3050 \quad (15)$$

and the associated reconstruction uncertainty was quantified by the elementwise median absolute deviation (MAD),

$$\text{MAD}_e = \text{median}\left(\left|x_e^{(1)} - \hat{x}_e\right|, \left|x_e^{(2)} - \hat{x}_e\right|, \left|x_e^{(3)} - \hat{x}_e\right|, \left|x_e^{(4)} - \hat{x}_e\right|\right) \quad (16)$$

An independent dielectric method was used to provide a spatial reference for the real-data evaluation. Dielectric measurements were performed along two vertical profiles corresponding to the left and right EIT electrode arrays. Each profile contained 16 measurement points located between 310 and 403 cm above ground level. The dielectric method provided mass moisture content, U_m, expressed as a percentage.



Figure 3. Non-destructive measurement of wall moisture using the GANN UNI 2 meter equipped with a B50 dielectric ball probe.

Measurements were performed using a GANN UNI 2 meter with a B50 ball probe (Figure 3). The method is based on changes in the dielectric properties of masonry with moisture content. The instrument provides a dimensionless reading X, which was converted to mass moisture content using an object-specific relationship previously established against gravimetric measurements: $f(x) = 0.0699e^{0.034X}$. The calibration yielded $R^2 = 0.7448$, supporting its use for semi-quantitative moisture assessment. Because dielectric readings may also be affected by material heterogeneity, porosity, and salt content, they were treated as an independent spatial reference rather than an absolute ground truth.

3. Results

3.1. Reconstructions of simulation measurements

Table 1 presents four representative validation cases selected to illustrate the reconstruction performance of the baseline model and the cDDPM-SFR under different spatial configurations of the conductivity field. The four cases were selected to represent distinct reference-field geometries differing in boundary orientation, spatial extent and geometric complexity; selection was based on the reference configurations rather than on the magnitude of improvement achieved by cDDPM-SFR. Although they do not represent the full diversity of the validation dataset, they enable a direct qualitative comparison of boundary localization, structural continuity, and the occurrence of reconstruction artefacts.

In Case #1, both models reproduced the separation between the upper and lower conductivity regions. Nevertheless, the baseline reconstruction contained additional local variations in the lower part of the domain, visible as light-blue areas that were not present in the reference pattern. The cDDPM-SFR substantially reduced these fluctuations and produced a more homogeneous reconstruction of the low-conductivity region. The greatest qualitative improvement can be observed in Cases #2 and #3. The baseline model generated several isolated high-intensity regions and irregular surface distortions, particularly close to inclined and non-uniform interfaces. These artefacts resulted in a fragmented representation of the reconstructed structures. In contrast, the cDDPM-SFR preserved the continuity of the conductivity regions and more accurately reproduced their spatial extent and boundary orientation.

For Case #4, the baseline model already provided a relatively accurate representation of the two principal regions. Consequently, the visual improvement obtained with the diffusion model was less pronounced than in Cases #2 and #3. Nevertheless, the conditional diffusion reconstruction exhibited a more coherent transition between the upper and lower parts of the domain and contained fewer isolated deviations from the expected conductivity values. The qualitative comparison indicates that the proposed model does not merely smooth the

reconstructed field. Instead, it suppresses spurious local extrema while retaining the position, continuity, and geometry of the moisture-related conductivity regions.

Table 1. Results of the reconstructions using baseline model and cDDPM-SFR.

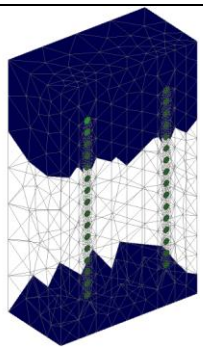
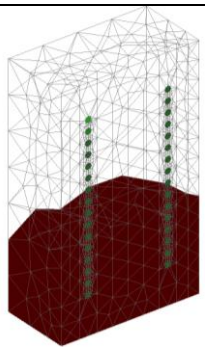
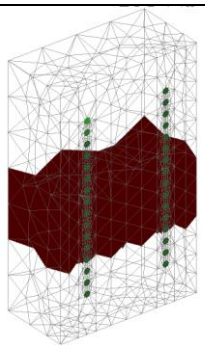
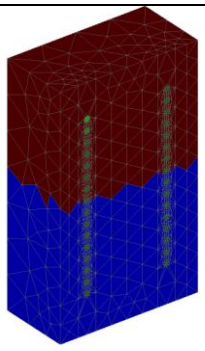
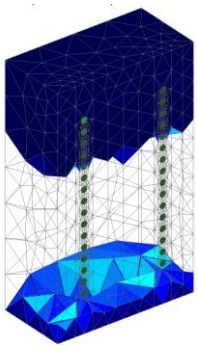
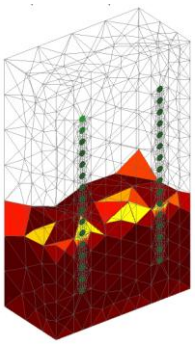
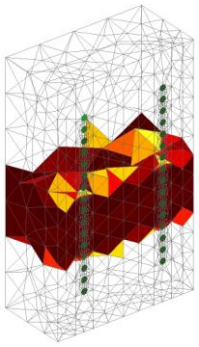
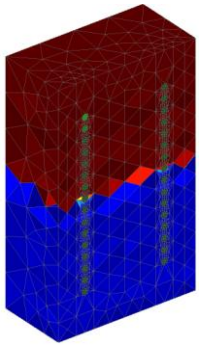
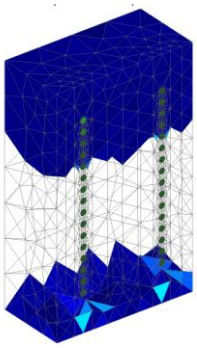
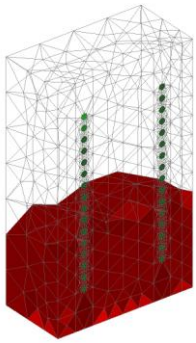
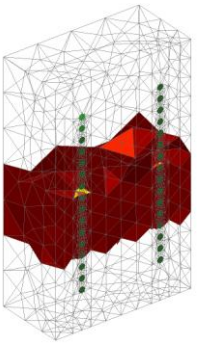
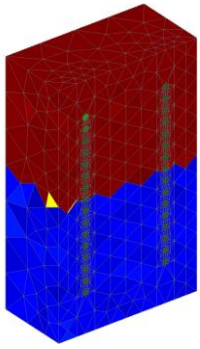
#	Case #1	Case #2	Case #3	Case #4
Pattern				
Baseline model				
cDDPM-SFR				

Table 2. Comparison of reconstruction quality indicators.

	#	Case #1	Case #2	Case #3	Case #4	average of 1000 validation cases
MSE	Baseline model	0.125	0.093	0.135	0.079	0.119
	cDDPM-SFR	0.081	0.018	0.048	0.025	0.048
PSNR	Baseline model	21.048	22.334	20.734	23.038	21.692
	cDDPM-SFR	22.921	29.409	23.701	27.902	26.961
SSIM	Baseline model	0.981	0.985	0.977	0.989	0.978
	cDDPM-SFR	0.988	0.997	0.989	0.996	0.992
PCC	Baseline model	0.984	0.988	0.981	0.993	0.981
	cDDPM-SFR	0.989	0.997	0.989	0.997	0.992

The quantitative results presented in Table 2 confirm the qualitative observations. Across the complete set of 1,000 paired validation cases, mean MSE decreased from 0.11887 for the direct MLP baseline to 0.04793 for the final S3 reconstruction, corresponding to a 59.7% reduction. The mean paired difference was −0.07094 (95% CI: −0.08274 to −0.05949). PSNR increased by 5.268 dB (95% CI: 4.565–5.951 dB), while the mean paired differences for SSIM and PCC were 0.01468 (95% CI: 0.01175–0.01804) and 0.01118 (95% CI: 0.00860–0.01403), respectively. All four Wilcoxon signed-rank tests remained significant after Holm correction (adjusted $p < 0.001$).

The representative cases in Table 2 illustrate these improvements for different spatial configurations, with the largest gains observed for Cases #2 and #3, characterised by inclined and irregular conductivity boundaries. Stage-wise analysis further showed a progressive improvement across the FEM-space refiners: mean MSE decreased from 0.09072 after S1 to 0.05086 after S2 and 0.04793 after S3, while PSNR increased from 23.406 to 26.612 and 26.961 dB. The largest gain occurred between S1 and S2, whereas S3 provided a smaller additional improvement. This stage-wise analysis quantifies the incremental contribution of successive FEM-space refiners but does not isolate the contribution of the diffusion stage itself.

To isolate the separate contributions of the diffusion stage and the FEM-space refiners, the accuracy of the diffusion output alone — mapped directly from the 16×16×16 voxel grid back to the native 3050-element FEM mesh, without any refinement — was additionally evaluated on the validation set and compared with the accuracy obtained after each successive refinement stage (Table 3).

Table 3. Ablation isolating the contribution of the diffusion stage from that of the cascaded FEM-space refiners.

Reconstruction stage	PSNR [dB]	MAE	PCC
Diffusion output only (voxel→FEM, before refinement)	18.03	0.106	0.946
Direct MLP baseline	21.69	0.158	0.981
S1 (first FEM-space refiner)	23.41	0.107	0.985
S1+S2	26.61	0.078	0.992
S1+S2+S3 (final cDDPM-SFR)	26.96	0.073	0.992

The diffusion output alone, mapped back onto the native FEM mesh without refinement, reached PSNR = 18.03 dB, MAE = 0.106, PCC = 0.946 — below the direct MLP baseline (PSNR = 21.69 dB). These metrics were evaluated over all 3050 physical FEM elements in the original conductivity units and are therefore not directly comparable with the normalised voxel-domain metrics reported in Table 4. The first FEM-space refiner (S1) already closed this gap and exceeded the direct MLP baseline, while S2 and S3 provided further improvements, reaching PSNR = 26.96 dB and PCC = 0.992 for the complete cDDPM-SFR pipeline. These results support complementary roles of the diffusion stage and successive FEM-space refinement: the diffusion stage alone does not reach baseline-level accuracy at the element level, and it is the FEM-space refiners that recover and exceed it.

To justify the one-step DDIM configuration, reconstruction quality was additionally evaluated in the voxel domain before voxel-to-FEM mapping as a function of the number of DDIM sampling steps, using a single initial-noise realisation (ensemble size = 1) (Table 4).

Table 4. Sensitivity of voxel-domain reconstruction quality to the number of DDIM sampling steps before voxel-to-FEM mapping.

DDIM sampling steps	PSNR [dB]	SSIM	MAE	PCC
1	26.82	0.977	0.0092	0.977
5	24.68	0.968	0.0204	0.969
10	23.46	0.961	0.0267	0.962
20	22.55	0.955	0.0333	0.957
50	21.05	0.944	0.0459	0.946
100	20.11	0.935	0.0658	0.938
200	18.50	0.914	0.0858	0.919

Increasing the number of DDIM steps consistently reduced reconstruction quality: PSNR decreased from 26.82 dB at one step to 18.50 dB at 200 steps, while SSIM and PCC decreased and MAE increased. Thus, single-step sampling was optimal for the reported model.

The influence of ensemble size (the number of averaged independent noise realisations) on reconstruction quality and computational time was evaluated separately at a fixed single sampling step, for N = 1, 5, 10, 20 and 50, reporting both voxel-domain and FEM-space accuracy together with the measured single-CPU reconstruction time for the 50-case validation set (Table 5).

Table 5. Sensitivity of reconstruction quality and computational time to the ensemble size (number of averaged independent noise realisations).

Ensemble size N	Voxel PSNR [dB]	Voxel PCC	FEM PSNR [dB]	FEM PCC	Time [s] (50 cases, CPU)
1	26.78	0.974	16.99	0.932	6
5	28.19	0.977	17.85	0.944	26
10	28.53	0.977	17.93	0.944	60
20	28.62	0.977	18.01	0.945	116
50	28.72	0.977	18.03	0.946	280

Quality improved with ensemble size but with strongly diminishing returns, while computational time scaled approximately linearly with N. The largest gain occurred between N=1 and N=5 (voxel PSNR +1.41 dB, FEM PSNR

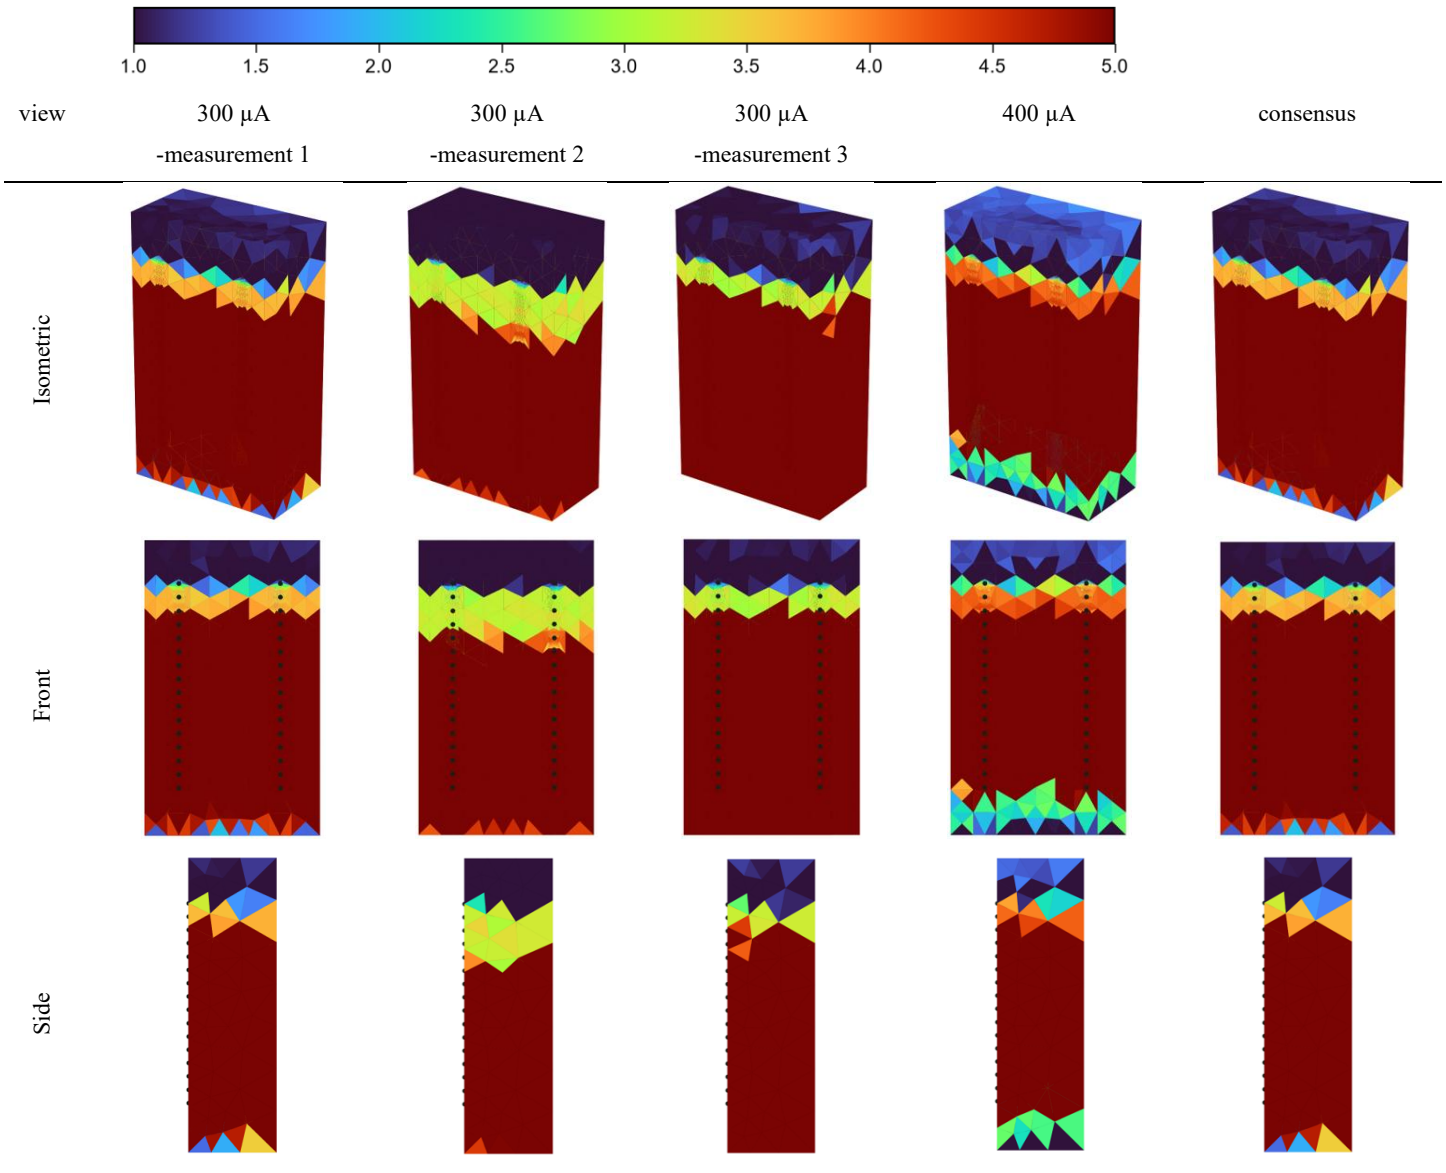
+0.86 dB) for a five-fold increase in computation (6 s to 26 s); increasing N five-fold again from 10 to 50 (60 s to 280 s) raised FEM PSNR by only 0.10 dB and FEM PCC by 0.002. The ensemble size of 50 used for the reported model and for the real-measurement reconstructions therefore lies well past the point of diminishing returns; N=50 was nonetheless retained throughout for consistency with the originally reported results. Timings were measured on a single CPU core for one complete pass over the 50-case validation set and, consistent with the field-measurement timing reported above, should be read as descriptive rather than as a repeated hardware-controlled benchmark.

The Tikhonov-regularised reconstruction achieved a mean PSNR of 22.83 dB, MAE of 0.163, and PCC of 0.977, compared with 21.69 dB, 0.158, and 0.981 for the direct MLP and 26.96 dB, 0.073, and 0.992 for cDDPM-SFR, respectively (Table 6). Thus, cDDPM-SFR outperformed both the classical regularised reconstruction and the direct MLP baseline, supporting the benefit of the proposed diffusion-and-refinement pipeline over both linear and direct nonlinear reconstruction approaches.

Table 6. Comparison with a classical Tikhonov-regularised linear reconstruction baseline.

Method	PSNR [dB]	MAE	PCC
Tikhonov-regularised linear reconstruction	22.83	0.163	0.977
Direct MLP baseline	21.69	0.158	0.981
cDDPM-SFR (S1+S2+S3)	26.96	0.073	0.992

Table 7. Comparison of cDDPM-SFR reconstructions obtained from real EIT measurements.



3.2. Reconstructions of Real EIT measurements

Table 7 presents the cDDPM-SFR reconstructions obtained independently for the four experimental EIT measurements conducted in-situ at a single chosen measurement location on a real-world object. Although all reconstructions indicate a general vertical differentiation of the relative conductivity field, differences can be observed in the position, continuity, and extent of the reconstructed zones.

The first 300 µA measurement produced a conductivity distribution that differed visibly from the remaining acquisitions. In particular, the transition between the higher- and lower-index regions occurred at a different elevation and exhibited a less consistent spatial course. The second and third 300 µA measurements generated substantially more similar conductivity fields. Both reconstructions preserved a comparable vertical arrangement of the principal regions and showed similar positions of the transition zone. Their strong similarity was also confirmed quantitatively by a Pearson correlation coefficient of 0.935. The reconstruction obtained at 400 µA retained the same general vertical organisation of the conductivity field but differed locally from the repeated 300 µA results. Pairwise correlations with the three 300 µA reconstructions ranged from 0.598 to 0.758, indicating moderate-to-strong correspondence.

Across the four field acquisitions, elementwise MAD of the relative EIT-derived conductivity index had mean 0.689, median 0.505, 95th percentile 2.778 and maximum 3.790. Pairwise PCC ranged from 0.196 to 0.935, showing heterogeneous repeatability. MAD describes dispersion across these four acquisitions, including different excitation currents; it is not a calibrated uncertainty or confidence interval.

The consensus field is characterised by relatively high index values throughout most of the lower and central parts of the investigated volume and lower values in the upper zone. The boundary between these regions is spatially non-uniform and changes elevation across the wall surface. Local intermediate-index areas are visible near both electrode arrays, indicating that the transition is gradual rather than represented by a single horizontal interface.

4. Discussion

To assess whether the real-data EIT reconstruction was physically plausible, the consensus field was compared qualitatively with dielectric measurements. The dielectric data were not used during model training or reconstruction. For this comparison, an EIT-derived surface proxy was calculated at each electrode position as the inverse-distance-weighted mean of the consensus values in the 12 nearest FEM elements. This uncalibrated quantity is expressed in arbitrary units and represents local variation in the reconstructed field rather than electrical conductivity or moisture content.

Table 8. Results of dielectric moisture measurements along the left and right electrode arrays.

Height above ground [cm]	Left electrode array			Right electrode array		
	Electrode no.	GANN UNI 2, X [–]	Mass moisture content, U_m [%]	Electrode no.	GANN UNI 2, X [–]	Mass moisture content, U_m [%]
403	1	140.90	8.41	32	137.40	7.47
397	2	140.30	8.24	31	137.40	7.47
391	3	140.50	8.30	30	135.20	6.93
384	4	133.10	6.45	29	139.60	8.05
378	5	140.60	8.33	28	131.60	6.13
372	6	141.00	8.44	27	146.50	10.18
366	7	140.40	8.27	26	133.00	6.43
360	8	155.90	14.01	25	151.00	11.86
354	9	129.00	5.61	24	139.20	7.94
347	10	124.10	4.75	23	136.20	7.17
341	11	145.20	9.74	22	129.20	5.65
335	12	129.90	5.79	21	127.40	5.32
329	13	124.80	4.87	20	140.80	8.39
323	14	141.50	8.59	19	140.80	8.39
316	15	133.60	6.56	18	132.00	6.22
310	16	133.40	6.52	17	134.50	6.77

The dielectric measurements showed a heterogeneous moisture distribution (Table 8). The mass moisture content ranged from 4.75% to 14.01% along the left array and from 5.32% to 11.86% along the right array. Both profiles reached their

maximum at 360 cm above ground level, with values of 14.01% and 11.86%, respectively. The occurrence of the maximum at the same elevation on both sides suggests a broader moisture-related zone rather than an isolated point response.

The consensus EIT reconstruction likewise showed a vertically differentiated and spatially non-uniform field. Relatively high index values occupied most of the lower and central parts of the investigated volume, whereas lower values occurred in the upper zone. The transition between these regions changed elevation across the wall, which is consistent with moisture transport through heterogeneous masonry rather than along a single horizontal boundary.

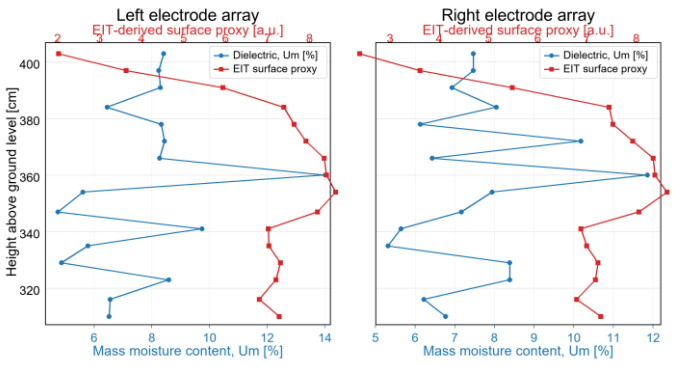


Figure 4. Qualitative comparison of the dielectric moisture profiles and the surface proxy extracted from the uncalibrated cDDPM-SFR consensus reconstruction.

An increase in the EIT-derived surface proxy was visible near 360 cm, corresponding to the elevation at which both dielectric profiles reached their maximum (Figure 4). From a general spatial perspective, the two methods therefore provided a consistent picture: both indicated a heterogeneous distribution and identified the central part of the investigated region as the location of the principal moisture-related feature. The agreement observed along both electrode arrays supports the physical plausibility of the anomaly reconstructed by EIT.

This agreement should be interpreted at the level of moisture-related zones and general spatial tendencies rather than as point-by-point equivalence. The dielectric probe provides a local, predominantly near-surface measurement, whereas EIT integrates information from a larger three-dimensional volume. In addition, electrical conductivity is influenced by water content and dissolved ions, while dielectric measurements may be affected by porosity, material composition and salts. These differences explain local deviations without contradicting the common location of the broader moisture-related zone.

The repeated EIT measurements further support the stability of this interpretation. The second and third measurements at 300 µA produced highly similar fields (Pearson coefficient 0.935), while the 400 µA reconstruction retained the same general vertical organisation. The elementwise median reduced acquisition-specific variation and preserved the features shared by the individual measurements. The dielectric results independently corroborated the principal spatial feature identified by EIT, particularly the moisture-related zone near 360 cm. The EIT reconstruction should nevertheless be interpreted as a map of relative conductivity variation rather than as a quantitative map of absolute moisture content.

A further limitation concerns the domain gap between the synthetic training data and the field measurements. The numerical dataset was generated using a fixed FEM geometry and controlled conductivity distributions, whereas historic masonry exhibits variability in brick and mortar properties, background conductivity, salt concentration, electrode contact conditions and local geometry. Gaussian perturbations introduced during training approximate measurement variability but cannot reproduce all sources of field noise and material heterogeneity. Consequently, the experimental results should be interpreted as evidence of feasibility and spatial plausibility rather than as confirmation that the synthetic training distribution fully represents the electrical behaviour of historic masonry. The agreement with the independent dielectric profiles provides external support for the identified spatial pattern but does not eliminate possible simulation-to-real distribution shift.

The archived CPU run took 11.55 s for the four field acquisitions with 50-member averaging, including refiner loading. This single recorded duration is descriptive, not a repeated hardware-controlled benchmark. A controlled ensemble-size and timing study remains necessary.

5. Conclusion

This study developed and evaluated a two-stage cDDPM-SFR pipeline for three-dimensional EIT reconstruction of moisture-related conductivity fields in masonry. The measurement-conditioned diffusion model recovered the global spatial

organisation of the field on a voxel grid, while three residual refinement stages used the original voltage vector to correct the result after its transfer to the 3050-element FEM representation.

For the complete set of 1,000 synthetic validation cases, cDDPM-SFR consistently outperformed the direct MLP baseline. Mean MSE decreased from 0.11887 to 0.04793, corresponding to a 59.7% reduction, while PSNR increased by an average of 5.268 dB. The paired improvements were statistically significant for MSE, PSNR, SSIM, and PCC after Holm correction (all adjusted $p < 0.001$). Stage-wise analysis further showed a progressive reduction in MSE from 0.09072 after S1 to 0.05086 after S2 and 0.04793 after S3, confirming an incremental benefit of successive FEM-space refinement. The reconstructed fields also showed improved continuity and fewer isolated local artefacts, particularly for inclined and irregular interfaces.

Application to four experimental measurement sequences showed that the dominant vertical organisation of the conductivity field was retained despite local inter-acquisition variability. The second and third measurements at 300 μA were highly similar (Pearson coefficient 0.935), while the 400 μA reconstruction preserved the same general spatial pattern. Because interpretation based on a single acquisition would have been sensitive to measurement-specific deviations, an elementwise median was used to obtain a consensus field that limited the influence of locally atypical values and retained the repeatable spatial features.

The independent dielectric measurements supported the physical plausibility of the real-data reconstruction. Both dielectric profiles reached their maximum moisture content at 360 cm above ground level, where an increase in the EIT-derived surface proxy was also observed. These results demonstrate the potential of cDDPM-SFR for three-dimensional identification of moisture-related conductivity variations in masonry. The reconstructed spatial information may support the selection of wall regions for targeted inspection, maintenance intervention, or repeated condition monitoring.

Future research should extend the experimental validation to multiple masonry structures representing different materials, geometries, moisture distributions, salt concentrations, and environmental conditions. Particular attention should be given to establishing a quantitative relationship between the reconstructed conductivity indices and gravimetrically determined moisture content. A larger number of repeated acquisitions should also be used to quantify reconstruction uncertainty and sensitivity to excitation current, electrode contact variability, temperature, and measurement noise. Further work will compare the proposed voxel-to-FEM pipeline with graph- and mesh-based diffusion architectures that operate directly on irregular finite-element domains. Finally, longitudinal field measurements should be conducted to determine whether cDDPM-SFR can reliably identify temporal changes in masonry condition and support maintenance planning, intervention prioritisation, and post-repair assessment.

References

1. Proietti N, Calicchia P, Colao F, De Simone S, Di Tullio V, Luvidi L, Prestileo F, Romani M, Tatì A. Moisture Damage in Ancient Masonry: A Multidisciplinary Approach for In Situ Diagnostics. *Minerals* 2021; 11(4): 406. https://doi.org/10.3390/MIN11040406.
2. Tallman TN, Smyl DJ. Structural health and condition monitoring via electrical impedance tomography in self-sensing materials: a review. *Smart Materials and Structures* 2020; 29(12): 123001. https://doi.org/10.1088/1361-665X/ABB352.
3. Brazey B, Haddab Y, Zemiti N. Robust imaging using electrical impedance tomography: review of current tools. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2022; 478(2258): 20210713. https://doi.org/10.1098/rspa.2021.0713.
4. Dimas C, Alimisis V, Uzunoglu N, Sotiriadis PP. Advances in Electrical Impedance Tomography Inverse Problem Solution Methods: From Traditional Regularization to Deep Learning. *IEEE Access* 2024; 12: 47797-47829. https://doi.org/10.1109/ACCESS.2024.3382939.
5. Smyl D, Hallaji M, Seppänen A, Pour-Ghaz M. Quantitative electrical imaging of three-dimensional moisture flow in cement-based materials. *International Journal of Heat and Mass Transfer* 2016; 103: 1348-1358. https://doi.org/10.1016/J.IJHEATMASSTRANSFER.2016.08.039.
6. Hallaji M, Seppänen A, Pour-Ghaz M. Electrical resistance tomography to monitor unsaturated moisture flow in cementitious materials. *Cement and Concrete Research* 2015; 69: 10-18. https://doi.org/10.1016/J.CEMCONRES.2014.11.007.
7. Kendzierawska W, Trochonowicz M, Szostak B. Moisture investigations in heritage structures – the influence of masonry salt content on

the reliability of dielectric measurements. *Budownictwo i Architektura / Civil and Architectural Engineering* 2025; 24(4): 25014. https://doi.org/10.35784/BUD-ARCH.8768.

8. Jurdziak L, Błażej R, Kirjanów-Błażej A, Rzeszowska A, Kostrzewa P. Improving the effectiveness of the DiagBelt+ diagnostic system - analysis of the impact of measurement parameters on the quality of signals. *Eksploatacja i Niezawodność – Maintenance and Reliability* 2024; 26(3): 187275. https://doi.org/10.17531/EIN/187275.
9. Kozłowski E, Antosz K, Mazurkiewicz D, Sęp J, Żabiński T. Integrating advanced measurement and signal processing for reliability decision-making. *Eksploatacja i Niezawodność – Maintenance and Reliability* 2021; 23(4): 777-787. https://doi.org/10.17531/EIN.2021.4.20.
10. Aghajanian S, Rao G, Ruuskanen V, Wajman R, Jackowska-Strumillo L, Koiranen T. Real-time fault detection and diagnosis of CaCO3 reactive crystallization process by electrical resistance tomography measurements. *Sensors* 2021; 21(21): 6958. https://doi.org/10.3390/S21216958.
11. Wajman R, Nowakowski J, Łukiański M, Banasiak R. Machine learning for two-phase gas-liquid flow regime evaluation based on raw 3D ECT measurement data. *Bulletin of the Polish Academy of Sciences: Technical Sciences* 2024; 72(1): e148842. https://doi.org/10.24425/BPASTS.2024.148842.
12. Tomaszewska-Wach B, Rzasa M. A correction method for wet gas flow metering using a standard orifice and slotted orifices. *Sensors* 2021; 21(7): 2291. https://doi.org/10.3390/S21072291.
13. Rzasa MR, Czapla-Nielacna B. Analysis of the Influence of the Vortex Shedder Shape on the Metrological Properties of the Vortex Flow Meter. *Sensors* 2021; 21(14): 4697. https://doi.org/10.3390/S21144697.
14. Korzeniewska E, Zawiślak R, Przybył S, Sarna P, Bilska A, Mączka M. Prototype of Data Collector from Textronic Sensors. *Sensors* 2023; 23(24): 9813. https://doi.org/10.3390/S23249813.
15. Pawłowski S, Plewako J, Korzeniewska E. Field Modeling of the Influence of Defects Caused by Bending of Conductive Textronic Layers on Their Electrical Conductivity. *Sensors* 2023; 23(3): 1487. https://doi.org/10.3390/S23031487.
16. Zhang M, Ma Y, Huang N, Ge H. Survey of EIT Image Reconstruction Algorithms. *Journal of Shanghai Jiaotong University (Science)* 2022; 27: 211-218. https://doi.org/10.1007/s12204-021-2333-1.
17. Shi Y, Zhang Y, Wang M, Lou Y, Zheng S. Reconstruction of conductivity distribution variation with Tikhonov and wavelet frame combined method for electrical impedance tomography. *Transactions of the Institute of Measurement and Control* 2023; 45(14): 2658-2668. https://doi.org/10.1177/01423312231159216.
18. Chen Y, Dong F, Tan C. Space-constrained optimized Tikhonov regularization method for 3D hemorrhage reconstruction by open magnetic induction tomography. *Physics in Medicine & Biology* 2022; 67(22): 225012. https://doi.org/10.1088/1361-6560/AC9E42.
19. Li Y, Zhu Q, Liu Z. Deep Learning for Image Reconstruction in Electrical Tomography: A Review. *IEEE Sensors Journal* 2025; 25(9): 14522-14538. https://doi.org/10.1109/JSEN.2025.3554696.
20. Shi S, Kang R, Liatsis P. A Conditional Diffusion Model for Electrical Impedance Tomography Image Reconstruction. *IEEE Transactions on Instrumentation and Measurement* 2025; 74: 4504816. https://doi.org/10.1109/TIM.2025.3550245.
21. Kulisz M, Kłosowski G, Rymarczyk T, Hoła A, Niderla K, Sikora J. The use of the multi-sequential LSTM in electrical tomography for masonry wall moisture detection. *Measurement* 2024; 234: 114860. https://doi.org/10.1016/J.MEASUREMENT.2024.114860.
22. Kłosowski G, Kulisz M, Rymarczyk T, Skowron Ł, Olszewski P, Niderla K. Application of machine learning in electrical process tomography with variable frequency measurement sequences. *Measurement* 2025; 247: 116770. https://doi.org/10.1016/J.MEASUREMENT.2025.116770.
23. Kłosowski G, Hoła A, Kulisz M, Rymarczyk T, Niderla K, Sikora J. A multi-head neural network architecture for EIT inverse problem resolution in wall moisture imaging. *Measurement* 2026; 283: 122107. https://doi.org/10.1016/J.MEASUREMENT.2026.122107.
24. Przysucha B, Wójcik D, Rymarczyk T, Król K, Kozłowski E, Gąsior M. Analysis of Reconstruction Energy Efficiency in EIT and ECT 3D Tomography Based on Elastic Net. *Energies* 2023; 16(3): 1490. https://doi.org/10.3390/EN16031490.
25. Liu J, Shi F, Xiong H, Zhou Y. DiffusionEIT: Diffusion Model for Electrical Impedance Tomography. *IEEE Transactions on Instrumentation and Measurement* 2024; 73: 4509209. https://doi.org/10.1109/TIM.2024.3472911.
26. Zhang T, Wang Z, Wang Q. A Sensitivity Prior Generation With Conditional Diffusion Model for EIT Reconstruction. *IEEE Signal*

Processing Letters 2025; 32: 3042-3046. https://doi.org/10.1109/LSP.2025.3590324.

27. Zhang X, Rong Z. A Data-Constrained and Physics-Guided Conditional Diffusion Model for Electrical Impedance Tomography Image Reconstruction. *Sensors* 2026; 26(5): 1728. https://doi.org/10.3390/S26051728.
28. Hu J, Song B, Fessler JA, Shen L. Patch-Based Diffusion Models Beat Whole-Image Models for Mismatched Distribution Inverse Problems. *arXiv preprint* 2024; https://arxiv.org/abs/2410.11730v1.
29. Herzberg W, Hauptmann A, Hamilton SJ. Domain independent post-processing with graph U-nets: applications to electrical impedance tomographic imaging. *Physiological Measurement* 2023; 44(12): 125008. https://doi.org/10.1088/1361-6579/AD0B3D.
30. Liu Z, Wang J, Shan Q, Liu D. GraphEIT: Unsupervised Graph Neural Networks for Electrical Impedance Tomography. *IEEE Transactions on Computational Imaging* 2024; 10: 1559-1570. https://doi.org/10.1109/TCI.2024.3485517.
31. Hu Z, Bai X, Shang J, Zhang R, Dong J, Wang X, Sun G, Fu H, Tai CL. VMNet: Voxel-Mesh Network for Geodesic-Aware 3D Semantic Segmentation. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*; Montreal, QC, Canada. 2021; p.15468-74. https://doi.org/10.1109/ICCV48922.2021.01520.
32. Xu J, Zhang R, Dou J, Zhu Y, Sun J, Pu S. RPVNet: A Deep and Efficient Range-Point-Voxel Fusion Network for LiDAR Point Cloud Segmentation. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*; Montreal, QC, Canada. 2021; p.16004-16013. https://doi.org/10.1109/ICCV48922.2021.01572.
33. Sypek A, Sypek R. *Zamki i obiekty warowne Pomorza Gdańskiego*. Warsaw: Oficyna Wydawnicza Alma-Press; 2004.
34. Wang Z, Bovik AC, Sheikh HR, Simoncelli EP. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing* 2004; 13(4): 600–612. https://doi.org/10.1109/TIP.2003.819861.