

Multi-domain signal processing and hybrid deep learning for robust ultrasonic tomography in industrial reactors

Monika Kulisz^{1*}, Grzegorz Kłosowski¹, Tomasz Rymarczyk^{2,3},
Konrad Niderla^{2,3}, Jan Sikora^{2,3}

¹ Lublin University of Technology, Nadbystrzycka 38D, 20-618 Lublin, Poland

² Research & Development Centre Netrix S.A., Związkowa 26, 20-704 Lublin, Poland

³ WSEI University, Projektowa 4, 20-209 Lublin, Poland

* Corresponding author's e-mail: m.kulisz@pollub.pl

ABSTRACT

Monitoring industrial reactors using ultrasonic tomography requires solving a strongly nonlinear and ill-posed inverse problem based on limited, noisy boundary measurements. This paper proposes a novel hybrid deep learning architecture that utilizes multi-domain signal processing via the Hilbert transform, Fourier transform, and S-transform (HFS) approximation. Instead of processing raw time-of-flight sequences directly, they were converted into three-channel descriptors, and the training dataset was explicitly augmented with 5% Gaussian noise to make the model resilient to real-world measurement disturbances. Evaluation proved the significant superiority of the proposed approach over the baseline model analyzing raw signals. The HFS architecture reduced the mean squared error from 0.241 to 0.128 and substantially increased the Pearson correlation coefficient from 0.624 to 0.863. This confirms that the proposed approach yields higher spatial localization fidelity, better preserves the geometry of submerged objects, and effectively suppresses background artifacts. The proposed framework provides a reliable and stable data-driven tomographic imaging tool for multiphase processes.

Keywords: industrial reactors, ultrasound tomography, deep learning, Hilbert transform, Fourier transform, S-transform.

INTRODUCTION

Industrial reactors are among the most critical units in chemical and process plants because the phenomena that ultimately determine conversion, selectivity, product consistency, energy efficiency, and operational safety occur inside vessels that are usually opaque and only partially observable during normal operation. Recent process-imaging literature has therefore increasingly emphasized that spatially resolved monitoring is essential wherever point measurements fail to capture internal heterogeneity and transient process dynamics [1]. This requirement is particularly evident in crystallization and multiphase systems, where tomographic methods have been developed specifically to improve the sensitivity of process-state observation and to

better capture internal interfaces, phase boundaries, and local disturbances [2].

The same tendency can be seen in data-driven studies showing that raw tomographic measurements can already support reliable identification of gas–liquid flow regimes relevant to industrial vessels and pipeline systems [3]. From a metrological perspective, recent work has also shown that even a limited liquid-phase presence can significantly distort measurement fidelity in gas-dominated mixtures, which confirms that hidden internal heterogeneity directly affects the reliability of industrial diagnostics [4]. A broader move toward internal-state imaging rather than purely external inspection is likewise visible in recent non-destructive tomographic studies of engineered materials, where volumetric methods

were used to assess internal self-repair efficiency and structural evolution [5].

In industrial operation, however, the importance of internal monitoring is not limited to diagnosis alone. Recent reliability-oriented studies have shown that hidden factors strongly influence maintenance-system readiness and the practical availability of technical systems [6]. This logic extends naturally to pressure equipment and process vessels, where real-time mechanical integrity assessment and predictive maintenance have been linked directly to safer operation, lower risk exposure, and better asset management [7,8]. The consequences of insufficient knowledge about internal degradation are especially clear in recent work on reactor pressure vessel welds, where long-term structural integrity is shown to depend on mechanisms that cannot be inferred reliably from surface-level observations alone [9]. At the same time, internal hydrodynamics remain a key determinant of reactor quality and efficiency. Recent studies have demonstrated ultrasonic monitoring of bubble-column operation, predictive modeling of gas holdup, and tomography-based observation of precipitation processes, showing that internal flow structure and phase evolution directly affect process performance and control quality [10,11]. For this reason, monitoring the interior of industrial reactors should be viewed not as an auxiliary measurement task, but as a central component of safety assurance, product-quality stabilization, and reliability-oriented process supervision.

Among non-invasive sensing approaches, ultrasonic tomography (UST) is particularly attractive for industrial reactors because it can operate in opaque media without interrupting the process and can provide spatially resolved information on sound-speed variation, phase distribution, and evolving internal structures. Previous studies have shown that time-of-flight-based UST is capable of monitoring internal process evolution in industrially relevant reactor conditions [12].

Related studies have also shown that UST measurements can support direct estimation of gas holdup in bubble-column reactors when coupled with data-driven processing [13], while ultrasonic measurements combined with machine-learning methods can be used to identify flow regimes in bubble-column operation [14]. The growing role of ultrasound in reactor engineering is additionally reflected in recent review papers discussing industrial crystallization applications and the

design logic of ultrasonic reactor set-ups [15,16]. Nevertheless, the usefulness of UST remains constrained by the complexity of wave propagation and by the mathematical difficulty of the inverse problem. Recent finite-element studies have shown that ultrasonic bulk-wave behavior in heterogeneous media depends strongly on internal structure and multimodal wave interactions [17].

Other recent work has emphasized the sensitivity of ultrasonic energy transmission to medium properties and propagation conditions, which in reactor environments leads to attenuation, scattering, refraction, and multipath effects [18]. These challenges become even more pronounced in bubbly two-phase systems, where increasing medium complexity substantially complicates ultrasound-based reconstruction and prediction [19]. As a result, reactor-oriented UST must reconstruct a high-dimensional internal field from a limited number of boundary measurements affected by noise and model mismatch. This limitation has been stated explicitly in recent reactor-focused tomography studies, where industrial UST is described as a sparse, noisy, and inherently ill-posed inverse problem [20]. Consequently, the industrial value of UST depends not only on its ability to reveal otherwise inaccessible process states, but also on reconstruction methods that remain stable and informative under limited data and realistic measurement disturbances.

Recent literature shows that artificial intelligence and deep learning are increasingly used not only for post-reconstruction image enhancement, but across the entire UST pipeline, including signal representation, feature extraction, inversion acceleration, denoising, and monitoring [21,22]. In UST, feature extraction directly from measurement signals has been shown to improve reconstruction quality and support more effective learning-based inversion [23], while machine-learning-assisted reflective UST has confirmed that properly designed signal descriptors can still enable accurate internal imaging [24,25].

Other studies have improved reconstruction under practical constraints by using super-resolution convolutional neural networks and by reducing the number of measurement channels with machine-learning support while preserving localization capability. Reactor-oriented tomography has also begun to adopt recurrent and hybrid models, including LSTM-based monitoring frameworks, and recent work has shown that even the choice of loss function can affect reconstruction

quality and energy-efficiency-oriented supervision [26,27]. More broadly, deep learning is rapidly expanding across the ultrasonic workflow [21,28], with advances in CNN-accelerated inversion [29,30], improved low-frequency imaging [31], CNN-based denoising [32], physics-embedded deep models [33,34], multimodal fusion [35], and neural-architecture modifications that improve robustness under sparse and noisy measurements [20,36]. Despite this progress, comparatively little attention has been paid to enriching sparse reactor-oriented time-of-flight data before learning through complementary time, frequency, and time–frequency representations, which motivates the present study.

Standard deep learning approaches to the UST inverse problem typically rely on mapping raw, single-channel 1D measurement sequences directly to high-dimensional tomographic images. However, these basic models struggle with the highly ill-posed nature of the problem, an unfavourable input-output ratio, and a severe domain shift between idealized synthetic simulations and real-world measurements, which are inherently contaminated by environmental noise and highly variable ultrasonic transducer coupling at the reactor surface. To address these limitations and bridge the existing research gap, this study proposes a novel, data-driven UST reconstruction pipeline designed specifically for monitoring industrial processes and reactor dynamics.

The novelty of the proposed methodology is based on three fundamental pillars. First, a

multi-domain preprocessing strategy utilizing the Hilbert, Fourier, and S-transform (HFS) is introduced to convert raw time-of-flight sequences into information-rich descriptors. Second, a hybrid deep learning architecture was employed, combining Convolutional Neural Networks for local feature extraction with Long Short-Term Memory layers for global sequence modelling. Finally, to guarantee high structural fidelity and robustness in practical field conditions, the network was trained on a synthetic dataset explicitly augmented with 5% Gaussian noise. This early noise injection acts as a powerful architectural regularizer, forcing the model to learn generalized and resilient representations that successfully compensate for industrial measurement disturbances.

MATERIALS AND METHODS

The experimental data utilized in this study were acquired using a prototype ultrasonic tomography system (Figure 1) developed and constructed at the Research and Development Center of Netrix S.A., specifically designed for monitoring industrial processes. The tomographic measurements were conducted on a laboratory model of an industrial reactor, which consisted of a liquid-filled (water) tank designed to house various submerged objects acting as internal structural anomalies.

The core of the measurement apparatus comprises 16 ultrasonic transducers evenly distributed along the external perimeter of the reactor.




Figure 1. UST tomograph connected to the reactor model

This strategic placement ensures full volumetric coverage of the cross-section and enables the precise localization of internal phase variations. The system operates by analyzing the time-of-flight (TOF) of acoustic waves propagating through the medium. Each transducer is capable of functioning in both transmission and reception modes. By sequentially exciting individual transducers and recording the arrival times at the remaining sensors, a complete set of TOF measurements is gathered. The total number of independent measurement paths (M) is determined by the number of transducers (n) according to the following equation:

$$M = \frac{n^2 - n}{2} \quad (1)$$

For $n = 16$ transducers, this configuration yields exactly 120 independent TOF measurements per full scanning cycle. This directly corresponds to the 120-dimensional raw input vector provided to the neural network reconstruction model.

From a hardware perspective, the data acquisition system is built upon sophisticated signal processing electronics, including a variable sampling rate analog-to-digital converter (ADC), multi-stage harmonic filtering, and programmable gain amplification. Optimal transducer excitation is achieved using a dedicated high-voltage generator capable of producing square waveforms up to $\pm 100V$.

To isolate the acoustic perturbations caused exclusively by internal anomalies, a differential measurement strategy was employed. Initial reference measurements were conducted in a homogeneous environment (a reactor devoid of objects) to establish a baseline TOF distribution. Subsequently, objects were submerged into the reactor, altering the acoustic wave propagation characteristics (e.g., via scattering and diffraction) and resulting in specific TOF deviations. These TOF variations relative to the baseline background provide the crucial contrast required for the reliable image reconstruction of the internal spatial distributions.

The study addresses the inverse reconstruction problem in UST, where the objective is to estimate the spatial acoustic velocity distribution inside the monitored object from boundary ultrasonic measurements. In the analysed dataset, each observation consists of an input measurement vector and its corresponding target acoustic velocity representation.

Formally, each sample is described as a pair:

$$(\mathbf{x}_i, \mathbf{y}_i), \mathbf{x}_i \in \mathbb{R}^{120}, \mathbf{y}_i \in \mathbb{R}^{3200} \quad (2)$$

where: \mathbf{x}_i is the raw measurement vector composed of 120 time-of-flight (ToF) measurements, and \mathbf{y}_i is the target output vector containing 3200 values describing the reconstructed acoustic velocity pattern in flattened form. The complete dataset contains 35,000 samples, divided into 31,500 training cases and 3,500 test cases, corresponding to a 90/10 split.

This problem is strongly ill-posed and nonlinear, because small perturbations in the measured acoustic signals / transit times may correspond to multiple admissible velocity distributions. In addition, the relatively short input sequence and the high dimensionality of the reconstruction target create an unfavourable input–output ratio.

To ensure the robustness of the reconstruction model and to bridge the potential domain shift between synthetic finite-element simulations and real-world in-situ measurements, the entire set of ToF measurement sequences was explicitly augmented with 5% Gaussian noise prior to the multi-domain HFS transformation. In practical UST field applications, measurements are inherently susceptible to environmental disturbances, electronic noise, and, most critically, highly variable ultrasonic transducer coupling at the heterogeneous at the reactor surface. Training the network exclusively on idealized, clean synthetic data carries a high risk of overfitting, which often leads to unstable reconstructions and spurious artifacts when processing actual measurement data. By integrating severe 5% measurement noise directly into the primary training dataset, the noise acts as a powerful architectural regularizer. This early noise injection forces the model to learn more generalized and resilient features across the time, frequency, and time-frequency domains, significantly stabilizing the highly ill-posed inverse mapping and guaranteeing that the network maintains high structural fidelity when deployed under non-ideal field conditions. For this reason, direct mapping from the raw 120-point signal to the 3200-dimensional acoustic velocity vector may be insufficient to capture subtle acoustic variations. To improve the information content of the input representation, a multi-domain pre-processing strategy was employed.

Instead of using the raw measurement sequence alone, each input vector was transformed into a richer multi-representational descriptor that combines information from the time, frequency, and time–frequency domains. This preprocessing stage was designed to increase the sensitivity of the reconstruction model to both global spectral characteristics and local non-stationary disturbances. For each raw measurement vector

$$\mathbf{x} = [x_1, x_2, \dots, x_{120}]^T \quad (3)$$

three complementary signal representations were generated.

The first representation was based on the analytic signal obtained with the Hilbert transform. The Hilbert envelope emphasizes instantaneous amplitude changes and highlights abrupt variations in the acoustic response:

$$x_a[n] = x[n] + j\mathcal{H}\{x[n]\} \quad (4)$$

$$h[n] = |x_a[n]| \quad (5)$$

where: $\mathcal{H}\{\cdot\}$ denotes the Hilbert transform and $h[n]$ is the envelope signal.

This representation is particularly useful for capturing local irregularities that may be associated with acoustic velocity disturbances.

The second representation was obtained from the discrete Fourier transform of the measurement vector:

$$X[k] = \sum_{n=0}^{N-1} x[n] e^{-\frac{j2\pi kn}{N}} \quad (6)$$

$$f[k] = |X[k]| \quad (7)$$

where: $N = 120$.

The Fourier magnitude spectrum provides a compact description of the global frequency content of the signal and allows the model to exploit spectral signatures related to different acoustic velocity patterns.

The third representation was a simplified time–frequency descriptor approximating the S-transform. In the implemented pipeline, this was achieved using a spectrogram-based representation. The signal was analysed with a window length of 20 samples, zero overlap, and $N_{FFT} = 128$. The magnitude of the resulting time–frequency matrix was then averaged over the frequency axis to obtain a one-dimensional descriptor:

$$s[m] = \frac{1}{K} \sum_{k=1}^K |S[m, k]| \quad (8)$$

where: $S[m, k]$ is the spectrogram magnitude at time index m and frequency bin k , and K is the number of frequency bins.

This component captures localised non-stationary behaviour that is not directly visible in either the raw signal or the global Fourier spectrum.

Unlike the standard S-transform, which utilizes computationally expensive frequency-dependent Gaussian windows, the implemented approximation employs a spectrogram with a fixed 20 sample window, zero overlap, and $N_{FFT} = 128$, followed by frequency averaging. This simplification drastically reduces the computational burden and produces a flattened 1D sequence that matches the required input dimensions, while successfully preserving the core capability of simultaneous time-frequency localization for non-stationary phenomena.

Because the three derived descriptors have different effective resolutions, they were resampled to a common length using one-dimensional interpolation. In the implemented model, the target sequence length was set to $C = 100$. Thus, each of the three channels was interpolated to 100 samples:

$$\tilde{\mathbf{h}} \in \mathbb{R}^{100}, \tilde{\mathbf{f}} \in \mathbb{R}^{100}, \tilde{\mathbf{s}} \in \mathbb{R}^{100} \quad (9)$$

and then concatenated into a single feature matrix:

$$\mathbf{C}_{HFS} = [\tilde{\mathbf{h}}, \tilde{\mathbf{f}}, \tilde{\mathbf{s}}] \in \mathbb{R}^{100 \times 3} \quad (10)$$

This representation can be interpreted as a compact three-channel sequence, where each channel describes a different physical aspect of the same measurement. The common-length interpolation ensures that all samples have identical input dimensions and can be processed by the same neural architecture.

The choice of $C = 100$ is consistent with the actual implementation and dataset size. It offers a compromise between preserving signal variability and avoiding unnecessary dimensional inflation. Compared with using the original 120-point raw vector directly, the HFS representation provides a more structured and information-rich input space for learning the inverse mapping.

For comparison purposes, a baseline single-channel representation was also considered, in which the raw 120-point measurement vector

was fed directly into a recurrent model without HFS transformation.

To approximate the nonlinear mapping between the measurement representation and the acoustic velocity distribution, a hybrid deep learning architecture was employed. As illustrated in Figure 2, the proposed model consists of three main processing stages: a convolutional feature extraction block, a recurrent processing block for sequence modelling, and a dense output stage. Such a design was selected because the inverse UST problem requires both effective extraction of local signal patterns and the ability to capture long-range dependencies across the transformed measurement sequence.

For the HFS-based configuration, the network input is the three-channel sequence $C_{HFS} \in \mathbb{R}^{100 \times 3}$, where each channel corresponds to one complementary physical-domain representation of the same measurement. At the initial stage, the input sequence is processed by a one-dimensional convolutional layer with a kernel size of 3 and same padding. This convolutional operation enhances the extraction of local structures and short-range dependencies in the HFS representation, which is particularly important for identifying subtle changes related to local acoustic velocity perturbations. The convolutional output is subsequently normalized and passed through a nonlinear activation layer to improve training stability and increase representational capacity.

The deeper part of the model consists of recurrent layers based on LSTM. These layers are responsible for learning the sequential relationships embedded in the transformed measurement

signal. Their inclusion is justified by the fact that the HFS representation preserves an ordered structure, even though it originates from multiple transformed domains. In practice, the recurrent component enables the model to integrate local descriptors into a coherent global representation, which is necessary for reconstructing the full acoustic velocity field from a limited number of boundary measurements.

To reduce the risk of overfitting, dropout regularization was introduced between selected recurrent layers. This is especially important in the present problem because the model predicts a high-dimensional output vector of length 3200 from a relatively compact input sequence. The final part of the network is a fully connected regression layer containing 3200 output neurons. This layer performs the projection from the learned latent representation into the target acoustic velocity space and produces the reconstructed acoustic velocity vector in flattened form.

In addition to the HFS-based model, a baseline architecture was considered in order to assess the contribution of the proposed multi-domain preprocessing strategy. In this reference configuration, the network received the raw 120-point measurement vector as a single-channel sequence and used a more compact recurrent architecture without the full HFS transformation. This comparison allowed the influence of the enriched signal representation on reconstruction performance to be evaluated directly. From a methodological point of view, the proposed architecture therefore serves not only as a regression tool, but also as a framework for examining whether multi-domain

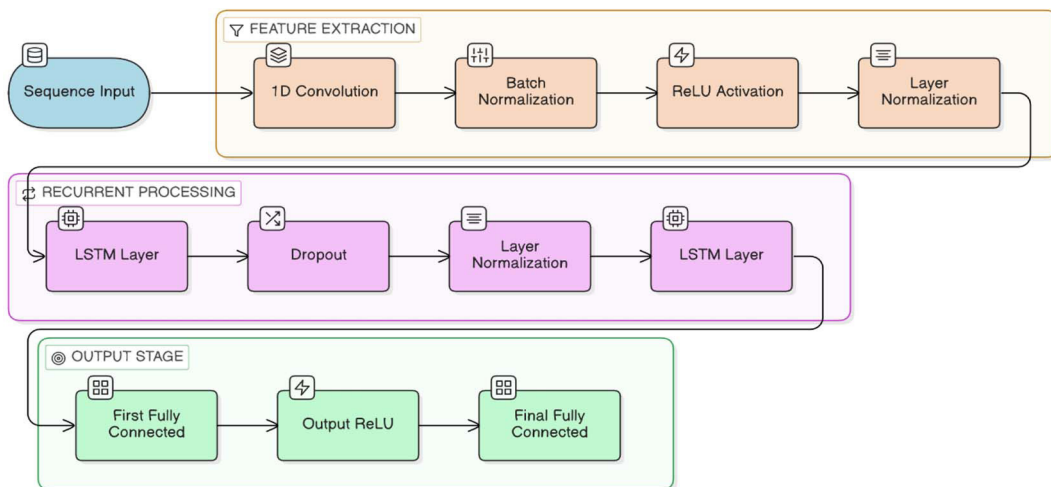


Figure 2. Schematic overview of the proposed hybrid 1D CNN-LSTM neural network architecture for UST image reconstruction

signal embedding improves the learnability of the UST inverse problem.

Specifically, the recurrent component in the final model consists of two LSTM layers containing 1024 and 2048 units, respectively. To reduce the risk of overfitting, dropout regularization with a rate of 0.2 was introduced between the layers. The network was trained using the Adam optimizer with a default learning rate of 0.001, which was selected due to its good convergence properties in deep nonlinear models and its robustness in handling high-dimensional parameter spaces. The maximum number of training epochs was set to 3000, allowing the model sufficient time to converge even in the presence of the strong non-linearity characteristic of inverse UST mapping. A mini-batch size of 64 was used in the HFS-based model, which provided a suitable balance between computational efficiency and gradient stability. During training, the input samples were shuffled at every epoch to reduce ordering effects and improve generalization.

To monitor model quality during optimization, validation performance was evaluated periodically at fixed intervals. Early stopping was employed as a regularization mechanism, so that training was terminated automatically when the validation error ceased to improve for a pre-defined number of checks. This strategy prevented unnecessary continuation of optimization after convergence and limited overfitting to the training data. The best-performing model parameters were retained and saved together with the training history for subsequent inference and analysis.

To comprehensively evaluate the quality of the UST image reconstructions, four evaluation metrics were employed: Mean squared error (MSE), peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and the Pearson correlation coefficient (PCC). These metrics were carefully selected to reflect the physical and diagnostic aspects of the acoustic reconstructions. It is crucial to emphasize that UST serves as a tomographic imaging tool designed to visualize relative differences in internal acoustic velocity, rather than to quantify absolute physical properties of multiphase flows or submerged objects. Consequently, the PCC must be interpreted primarily as a measure of the spatial distribution consistency, which is a key observable in industrial anomaly diagnostics. Similarly, SSIM serves as an indicator of the spatial localization fidelity, evaluating whether the structure and geometry of submerged objects and

phase boundaries are properly preserved. Finally, MSE and PSNR act as complementary global measures quantifying the overall reconstruction fidelity and robustness to noise.

RESULTS

Figure 3 presents five representative reconstruction cases. The first column shows the reference acoustic velocity distributions, the second column presents the reconstructions obtained with the baseline simple model operating on the raw measurement sequence, and the third column shows the reconstructions generated by the proposed HFS-based model. The comparison was intended to illustrate the visual differences between the reference patterns and the outputs of both neural architectures under representative test conditions.

As illustrated in Figure 3, the HFS-based network provides a visibly closer approximation of the reference acoustic velocity distributions compared to the baseline architecture, which processes only the raw measurement sequence. A detailed qualitative inspection of the representative cases reveals three primary visual advantages of the proposed multi-domain preprocessing strategy. First, in terms of localization accuracy, the HFS-based network more precisely recovers the spatial position of regions with altered acoustic velocity. In contrast, the baseline model tends to over-smooth these regions, resulting in blurred boundaries and a noticeable loss of spatial precision. Second, regarding geometric fidelity, the advantage of the HFS-based approach becomes particularly evident in more complex, irregular patterns. In such instances, the proposed model better preserves edge sharpness and maintains the separation between distinct structural components, whereas the baseline reconstructions show stronger merging and distortion of neighboring anomalies. Third, the HFS-based architecture demonstrates significantly improved artifact suppression. The baseline model often exhibits undesired fluctuations in regions that should remain relatively neutral (the background), producing a noisier and less reliable reconstruction. By leveraging the combined time, frequency, and time-frequency descriptors, the HFS-based network effectively mitigates these background fluctuations, yielding considerably cleaner tomograms with fewer spurious variations.”

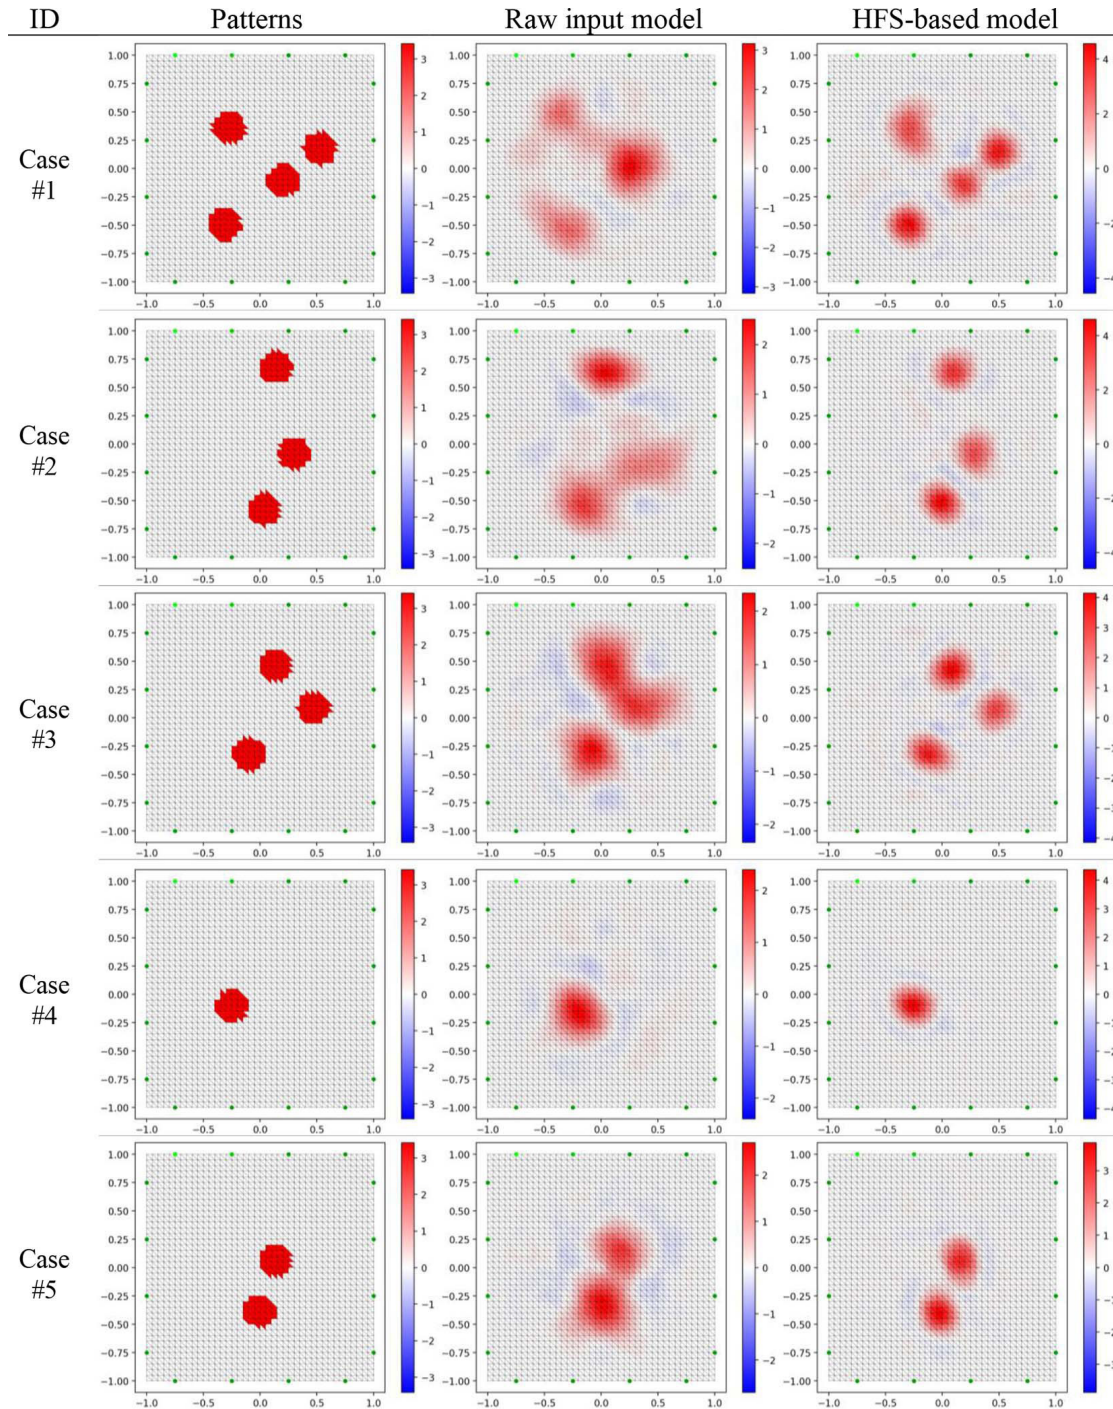


Figure 3. Results of the reconstructions using baseline and HFS-based model

Table 1 summarizes the quantitative reconstruction results for the same five representative cases shown in Figure 3. The table compares the values of MSE, PSNR, SSIM, and Pearson correlation coefficient (PCC) obtained for the baseline simple model and the proposed HFS-based model. For each metric, an additional column with the average value across the five analysed cases can also be included in order to provide a compact summary of the overall reconstruction tendency.

The quantitative results summarized in Table 1 strongly support these visual observations. The proposed HFS-based architecture consistently outperforms the baseline model across all evaluated metrics, both for the five representative cases and in the average results calculated over the 1000 validation cases. Notably, the global error was significantly reduced, with the average MSE dropping from 0.241 in the baseline model to 0.128 in the HFS-based model. This substantial improvement

Table 1. Comparison of reconstruction quality indicators for a baseline and HFS-based model

Metrics	Input type	Case #1	Case #2	Case #3	Case #4	Case #5	Average of 1000 validation cases
MSE	Baseline model	0.536	0.418	0.332	0.097	0.202	0.241
	HFS-based model	0.217	0.188	0.174	0.038	0.092	0.128
PSNR	Baseline model	2.707	3.783	4.782	10.102	6.926	7.037
	HFS-based model	6.621	7.247	7.602	14.129	10.329	10.048
SSIM	Baseline model	0.392	0.387	0.397	0.452	0.417	0.407
	HFS-based model	0.393	0.391	0.415	0.468	0.425	0.422
PCC	Baseline model	0.523	0.485	0.633	0.709	0.677	0.624
	HFS-based model	0.839	0.808	0.829	0.896	0.868	0.863

in overall reconstruction fidelity and robustness to noise is further reflected in the PSNR, which increased from an average of 7.037 to 10.048.

From a diagnostic perspective, the most critical improvement was observed in the PCC, which surged from an average of 0.624 in the baseline to 0.863 in the proposed model. Since the PCC serves as a primary measure of the spatial distribution consistency, this significant increase demonstrates the superior capability of the multi-domain representation in correctly capturing acoustic velocity variations, which is a key observable in industrial anomaly diagnostics. Furthermore, the SSIM showed a consistent improvement (from an average of 0.407 to 0.422), indicating that the HFS transformation enhances the spatial localization fidelity, thereby better preserving the structure and geometry of submerged objects and phase boundaries. Overall, the integration of time, frequency, and time-frequency descriptors effectively stabilized the highly ill-posed UST inverse mapping, demonstrating a clear advantage over relying solely on raw single-channel measurement sequences. This observation confirms that PCC and SSIM should be interpreted jointly, since PCC reflects global correlation of the reconstructed field, whereas SSIM provides stricter information about local structural preservation.

RECONSTRUCTIONS OF REAL MEASUREMENTS

In this study, a total of 10 real-world measurement cases were acquired and evaluated to test the model’s performance under practical

industrial conditions. Figure 4 displays reconstructions for five representative cases out of this set. Crucially, all real measurements were entirely excluded from the training dataset, which consisted solely of synthetic data augmented with 5% Gaussian noise. This ensured a rigorous evaluation of the model’s robustness against domain shift and actual transducer coupling variability. Figure 4 displays reconstructions derived from actual measurements acquired through the HFS-based model with subsequent processing. To ensure high-fidelity visualization and to mitigate artifacts associated with spatial discretization, the raw reconstructed data were subjected to a dedicated three-stage post-processing pipeline. This procedure was developed to reduce mesh-dependent jaggedness, suppress background noise, and recover the physically plausible organic morphology of the internal inclusions.

The first step was to apply primary spatial smoothing to the reconstructed distribution. This distribution had high-frequency artifacts and uneven edges because of the discrete unstructured triangular finite-element mesh. A distance-weighted Gaussian filter was applied directly to the mesh to lessen these effects. For each element, a new value was calculated by adding up the intensities of its neighbors, with the weights based on how far apart the triangle incenters were from each other. This step worked like a low-pass spatial operator, putting broken signals together into more coherent areas and lowering local variance. Adaptive intensity thresholding was used in the second stage to get the most important areas. The first smoothing connected

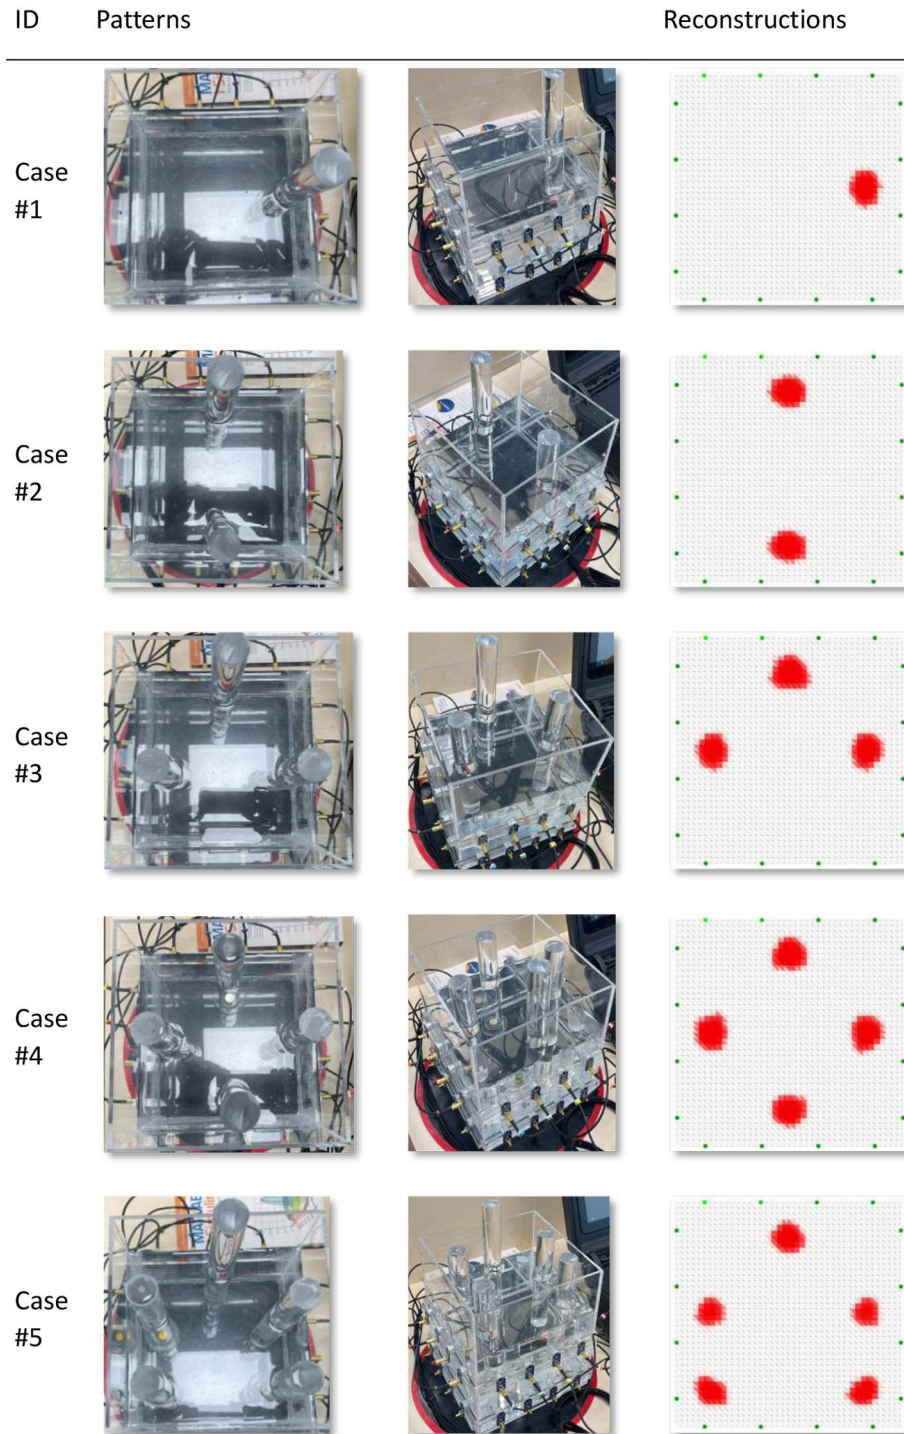


Figure 4. Reconstructions based on real measurements obtained using the HFS-based model with post-processing

broken structures, but it also made the response wider and added diffuse halos with background noise. So, a strict relative criterion, like 60% of the global maximum intensity, was used to set a proportional threshold. We set elements below this threshold to zero, which got rid of halos and background noise and left only the dense high-intensity cores that corresponded to the

physical inclusions. In the third stage, secondary anti-aliasing improved the shape of the inclusions that were taken out. Thresholding separated the important structures, but it also made sharp and unnatural boundaries that showed the angular mesh geometry. A second Gaussian filter with less variance was used on the isolated cores to bring back morphological realism. This

operation gently smoothed the threshold boundaries without making the inclusions noticeably bigger. This made the cross-sectional profiles smoother and more physically plausible.

In general, the sequential combination of broad spatial smoothing, strict core extraction, and localized boundary refinement closes the gap between discrete mathematical reconstruction and continuous physical reality. The proposed post-processing pipeline, along with the strong HFS-based reconstruction model, reduces artifacts caused by spatial discretization, suppresses measurement disturbances and background noise, and creates clean, sharp, and geometrically consistent tomograms that accurately show the internal structures of the reactor, even in less-than-ideal industrial settings.

CONCLUSIONS

This study addressed the highly ill-posed and nonlinear inverse problem in UST, where the objective is to reconstruct the spatial acoustic velocity distribution from boundary ToF measurements. To overcome the limitations of learning directly from raw, single-channel sequences, a novel multi-domain preprocessing strategy based on the Hilbert transform, Fourier transform, and S-transform approximation (HFS-based model) was proposed. The enriched descriptors were subsequently processed by a hybrid deep learning architecture, effectively combining local feature extraction with recurrent long-range sequence modelling.

A critical aspect of the proposed methodology was its orientation towards practical, non-ideal measurement conditions. By explicitly augmenting the training dataset with 5% Gaussian noise prior to the HFS transformation, the network was forced to learn generalized and resilient representations. This early noise injection acted as a strong architectural regularizer, effectively bridging the domain shift between synthetic finite-element simulations and real-world field conditions. It successfully compensated for inherent environmental disturbances and highly variable ultrasonic transducer coupling, significantly stabilizing the inverse mapping.

Both qualitative and quantitative evaluations demonstrated the clear superiority of the multi-domain HFS-based architecture over the baseline simple model. The integration of time, frequency, and time-frequency descriptors allowed

the network to achieve a substantial reduction in the global error, nearly halving the average MSE (from 0.241 to 0.128) and increasing the PSNR (from 7.037 to 10.048). Most notably, the PCC surged from an average of 0.624 in the baseline model to 0.863. From a diagnostic perspective, this proves the model's exceptional capability to accurately capture the spatial distribution consistency of acoustic velocity, which is a crucial parameter in industrial anomaly diagnostics. Furthermore, visual inspections and improved SSIM (increasing from 0.407 to 0.422) confirmed that the HFS-based approach yields higher spatial localization fidelity, better preserves the geometry of submerged objects and phase boundaries, and effectively suppresses background artifacts.

In conclusion, the results confirm that transforming raw ultrasonic measurement sequences into a multi-domain HFS descriptor fundamentally improves the learnability and robustness of the UST inverse problem. The proposed hybrid framework provides a reliable, data-driven tomographic imaging tool capable of visualizing relative differences in internal acoustic properties. Future work will focus on the deployment of this architecture in extensive industrial process monitoring campaigns and the potential integration of calibration models for absolute quantitative assessment of internal multiphase flows.

REFERENCES

1. Hampel U, About L, Banasiak R, Schleicher E, Soleimani M, Wondrak T, et al. A review on fast tomographic imaging techniques and their potential application in industrial process control. *Sensors* 2022;22:2309. <https://doi.org/10.3390/S22062309>
2. Wajman R. The concept of 3D ECT system with increased border area sensitivity for crystallization processes diagnosis. *Sensor Review* 2021;41:35–45. <https://doi.org/10.1108/SR-10-2019-0254>
3. 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. <https://doi.org/10.24425/BPASTS.2024.148842>
4. Tomaszewska-Wach B, Rzasca M. A correction method for wet gas flow metering using a standard orifice and slotted orifices. *Sensors* 2021;21. <https://doi.org/10.3390/S21072291>
5. Korzeniewska E, Sekulska-Nalewajko J, Goławski J, Irzmańska E, Adamus-Włodarczyk A, Kozikowski

- P. Evaluation of self-repair efficiency of polymers containing microcapsules using optical coherence tomography. *Compos Struct* 2025;351:118525. <https://doi.org/10.1016/J.COMPSTRUCT.2024.118525>
6. Kozłowski E, Borucka A, Oleszczuk P, Jałowiec T. Evaluation of the maintenance system readiness using the semi-Markov model taking into account hidden factors. *Eksploracja i Niezawodność* 2023;25. <https://doi.org/10.17531/EIN/172857>
 7. Sierra S, Wong J, Dawson K, Podskarbi M, Knezivic D. Real time mechanical integrity for pressure vessels and other critical assets. *American Society of Mechanical Engineers, Pressure Vessels and Piping Division (Publication) PVP* 2022;5. <https://doi.org/10.1115/PVP2022-84348>
 8. Karapanagiotis C, Heimann J, Duffner E, Charmi A, Schukar M, Hashemi S, et al. Towards predictive maintenance of hydrogen pressure vessels based on multi-sensor data. *E-Journal of Nondestructive Testing* 2024;29. <https://doi.org/10.58286/30513>
 9. Emerson JN, Marrero-Jackson EH, Nemets GA, Okuniewski MA, Wharry JP. Nuclear reactor pressure vessel welds: A critical and historical review of microstructures, mechanical properties, irradiation effects, and future opportunities. *Mater Des* 2024;244:113134. <https://doi.org/10.1016/J.MATDES.2024.113134>
 10. Franco EE, Henao Santa S, Cabrera JJ, Laín S. Air flow monitoring in a bubble column using ultrasonic spectrometry. *Fluids* 2024;9:163. <https://doi.org/10.3390/FLUIDS9070163>
 11. Hazare SR, Patil CS, Vala S V., Joshi AJ, Joshi JB, Vitankar VS, et al. Predictive analysis of gas hold-up in bubble column using machine learning methods. *Chemical Engineering Research and Design* 2022;184:724–39. <https://doi.org/10.1016/J.CHERD.2022.06.007>
 12. Koulountzios P, Rymarczyk T, Soleimani M. Ultrasonic time-of-flight computed tomography for investigation of batch crystallisation processes. *Sensors* 2021;21:639. <https://doi.org/10.3390/S21020639>
 13. Hu J, Li N, Wang L, Yang P, Yang Y, Quan Y. Direct estimation of gas holdup in gas-liquid bubble column reactors using ultrasonic transmission tomography and artificial neural processing. *Meas Sci Technol* 2022;33. <https://doi.org/10.1088/1361-6501/AC5D78>
 14. Wongsaraj W, Thong-Un N, Hansot J, Shoji N, Treenuson W, Kikura H. The application of ultrasonic measurement and machine learning technique to identify flow regime in a bubble column reactor. *Sensors International* 2025;6:100294. <https://doi.org/10.1016/J.SINTL.2024.100294>
 15. Xiang L, Fu M, Wang T, Wang D, Xv H, Miao W, et al. Application and development of ultrasound in industrial crystallization. *Ultrason Sonochem* 2024;111:107062. <https://doi.org/10.1016/J.ULTSONCH.2024.107062>
 16. Adamou P, Harkou E, Villa A, Constantinou A, Dimitratos N. Ultrasonic reactor set-ups and applications: A review. *Ultrason Sonochem* 2024;107:106925. <https://doi.org/10.1016/J.ULTSONCH.2024.106925>
 17. Dorval V, Leymarie N, Imperiale A, Demaldent E. Multi-modal characterization of ultrasonic bulk wave properties in heterogeneous textured media through finite element computations. *J Phys Conf Ser* 2024;2904:012005. <https://doi.org/10.1088/1742-6596/2904/1/012005>
 18. Ma Y., Jiang Y., Li CA, Ma Y, Jiang Y, Li C. A universal model for ultrasonic energy transmission in various media. *Sensors* 2024;24:6230. <https://doi.org/10.3390/S24196230>
 19. Wada Y, Hirose Y, Sibamoto Y. A numerical study on machine-learning-based ultrasound tomography of bubbly two-phase flows. *Ultrasonics* 2024;141:107346. <https://doi.org/10.1016/J.ULTRAS.2024.107346>
 20. Rymarczyk T, Kulisz M, Kłosowski G, Wójcik D, Kowalski M, Król K. Improving the reliability of industrial reactors by using differential neural network architecture in ultrasonic tomography. *Eksploracja i Niezawodność – Maintenance and Reliability* 2025;28. <https://doi.org/10.17531/EIN/208883>
 21. Cantero-Chinchilla S, Wilcox PD, Croxford AJ. Deep learning in automated ultrasonic NDE – Developments, axioms and opportunities. *NDT & E International* 2022;131:102703. <https://doi.org/10.1016/J.NDTEINT.2022.102703>
 22. Song K, Feng J, Chen D. A survey on deep learning in medical ultrasound imaging. *Front Phys* 2024;12:1398393. <https://doi.org/10.3389/FPHY.2024.1398393/XML>
 23. Xiao D, Yu ACH. Beamforming-integrated neural networks for ultrasound imaging. *Ultrasonics* 2025;145:107474. <https://doi.org/10.1016/J.ULTRAS.2024.107474>
 24. Majerek D, Rymarczyk T, Wójcik D, Kozłowski E, Rzemieniak M, Gudowski J, et al. Machine learning and deterministic approach to the reflective ultrasound tomography. *Energies (Basel)* 2021;14:7549. <https://doi.org/10.3390/EN14227549>
 25. Liu Z, Zhou X, Yang H, Zhang Q, Zhou L, Wu Y, et al. Reconstruction of reflection ultrasound computed tomography with sparse transmissions using conditional generative adversarial network. *Ultrasonics* 2025;145. <https://doi.org/10.1016/J.ULTRAS.2024.107486>
 26. Kłosowski G, Rymarczyk T, Niderla K, Kulisz M, Skowron Ł, Soleimani M. Using an LSTM network to monitor industrial reactors using electrical capacitance and impedance tomography – a hybrid approach. *Eksploracja i Niezawodność – Maintenance and Reliability* 2023;25:2023. <https://doi.org/10.17531/EIN.2023.1.11>

27. Kulisz M, Kłosowski G, Rymarczyk T, Słonec J, Gauda K, Cwynar W. Optimizing the neural network loss function in electrical tomography to increase energy efficiency in industrial reactors. *Energies* 2024, Vol 17, Page 681 2024;17:681. <https://doi.org/10.3390/EN17030681>
28. Luijten B, Chennakeshava N, Eldar YC, Mischi M, van Sloun RJG. Ultrasound Signal Processing: From Models to Deep Learning. *Ultrasound Med Biol* 2023;49:677–98. <https://doi.org/10.1016/J.ULTRASMEDBIO.2022.11.003>
29. Kleman C, Anwar S, Liu Z, Gong J, Zhu X, Yunker A, et al. Full waveform inversion-based ultrasound computed tomography acceleration using two-dimensional convolutional neural networks. *J Non-destruct Eval Diagn Progn Eng Syst* 2023;6. <https://doi.org/10.1115/1.4062092/1160002>
30. Qu X, Ren C, Yan G, Zheng D, Tang W, Wang S, et al. Deep-learning-based ultrasound sound-speed tomography reconstruction with tikhonov pseudo-inverse priori. *Ultrasound Med Biol* 2022;48:2079–94. <https://doi.org/10.1016/J.ULTRASMEDBIO.2022.05.033>
31. Lei M, Zhang W, Zhang T, Wu Y, Gao D, Tao X, et al. Improvement of low-frequency ultrasonic image quality using an enhanced convolutional neural network. *Sens Actuators A Phys* 2024;365:114878. <https://doi.org/10.1016/J.SNA.2023.114878>
32. Wang T, Yu C. Research on ultrasonic echo signal denoising via integration of adaptive variational mode decomposition algorithm and convolutional neural network. *Insight: Non-Destructive Testing and Condition Monitoring* 2024;66:567–75. <https://doi.org/10.1784/INSI.2024.66.9.567>
33. Ren J, Li J, Liu C, Chen S, Liang L, Liu Y. Deep learning with physics-embedded neural network for full waveform ultrasonic brain imaging. *IEEE Trans Med Imaging* 2024;43:2332–46. <https://doi.org/10.1109/TMI.2024.3363144>
34. Ahmadi M, Biswas D, Lin M, Vrionis FD, Hashemi J, Tang Y. Physics-informed machine learning for advancing computational medical imaging: integrating data-driven approaches with fundamental physical principles. *Artif Intell Rev* 2025;58:297. <https://doi.org/10.1007/S10462-025-11303-W/TABLES/7>
35. Zamiela C, Jiang Z, Stokes R, Tian Z, Netchaev A, Dickerson C, et al. Deep multi-modal u-net fusion methodology of thermal and ultrasonic images for porosity detection in additive manufacturing. *J Manuf Sci Eng* 2023;145. <https://doi.org/10.1115/1.4056873/1156667>
36. Liu Z, Zhang J, Li Z, Ding M, Yuchi M, Qiu W. Reflection ultrasound computed tomography with sparse data by residual diffusion models. 2025;13412:308–15. <https://doi.org/10.1117/12.3045310>