

Deep Learning-Based Inversion of Near-Field Overpressure Signatures for Low-Boom Supersonic Transport

Yu-Huan Peng[1,2], Ke-Shi Zhang[1,2,*] , Qing Chen[1,2] & Zhong-Hua Han[1,2]

[1]School of Aeronautics, Northwestern Polytechnical University, Xi'an, Shaanxi, 710072, PR China
[2]National Key Laboratory of Aircraft Configuration Design, Xi'an, Shaanxi, 710072, PR China

Abstract

Low-boom inverse design aims to minimize ground sonic-boom loudness and requires an efficient way to infer feasible near-field overpressure signatures from prescribed far-field ground waveforms. Existing methods remain strongly dependent on design experience and are limited by high computational cost, unstable shock inversion, and insufficient representation of local waveform features. To address these issues, this work proposes a low-dimensional feature-mapping framework that predicts near-field latent variables and reconstructs target near-field overpressure signatures through the trained decoder. The proposed method consists of three steps: a one-dimensional convolutional autoencoder (1D-CAE) encodes near-field overpressure signatures into latent variables, segment-wise principal component analysis (PCA) extracts low-dimensional far-field features, and a multilayer perceptron (MLP) establishes a nonlinear mapping from low-dimensional far-field features to near-field latent variables. Using the JAXA-Wing-Body (JWB) low-boom configuration from the Second AIAA Sonic Boom Prediction Workshop (SBPW-2) as the test case, the far-field waveform propagated from the inverted near-field overpressure signature reaches a perceived loudness level (PLdB) of 76.72, corresponding to a 2.44 PLdB reduction from the baseline. The results demonstrate that the proposed method establishes an effective correspondence between low-dimensional far-field features and near-field overpressure signatures, enabling efficient target near-field signature generation for conceptual-stage low-boom supersonic transport (SST) design.

Keywords: Supersonic Transport; Low-Boom Inverse Design; Near-Field Overpressure Signature Inversion; Deep Learning; Segment-Wise PCA

1. Introduction

Civil supersonic transport can significantly reduce intercontinental flight time. However, overland operation remains constrained by sonic boom [1]. During supersonic flight, shock and expansion waves generated by the aircraft propagate through the atmosphere and evolve into ground-level sonic-boom signatures under nonlinear propagation effects [2]. Excessive sonic-boom loudness affects public acceptance and has long limited commercialization of civil supersonic aircraft. Although the NASA X-59 demonstrator is designed to show the potential of aerodynamic shaping to reduce the ground perceived loudness level to approximately 75 PLdB [3], more stringent quiet-boom targets below 70 PLdB still require improved suppression and inverse design methods [4].

Low-boom inverse design has been explored through equivalent-area shaping, direct numerical optimization, surrogate-assisted design, and target-waveform-based methods [1,5]. Among them, target-waveform-based inverse design connects ground loudness requirements with aerodynamic shape modification by prescribing a far-field ground waveform and seeking a corresponding near-field overpressure signature. Related target near-field generation methods include waveform parameterization, forward-propagation search, reverse propagation, and reduced-order inversion. Waveform-parameterization methods modify waveform parameters to generate near-field targets [6], while revised Whitham theory supports rapid near-field prediction and target evaluation [7]. However, prescribed near-field targets remain experience-dependent and may not correspond to minimum ground loudness.

* **Corresponding author, E-mail: zhangkeshi@nwpu.edu.cn**

Forward-propagation-based search methods evaluate candidate near-field signatures through repeated solutions of the generalized Burgers equation, leading to high computational cost [6]. Reverse propagation can infer near-field overpressure signatures from far-field or mid-field waveforms, but the inverse process is sensitive to regularization errors and numerical dissipation, which may degrade weak shocks and local high-frequency details [8,9]. POD-based inversion combined with the generalized Burgers equation has demonstrated the feasibility of low-dimensional sonic-boom inversion, but linear modal representations remain limited in capturing local nonlinear features, such as sharp peak–trough transitions and tail recovery [10]. Therefore, an efficient and robust method is needed to recover feasible near-field overpressure signatures from target far-field ground waveforms while preserving local acoustic features.

Deep neural networks offer a promising solution because of their strong capability in nonlinear mapping and inverse modeling [11,12]. They have been widely used in flow prediction [13], flow reconstruction [14], and aerodynamic inverse design [15,16]. In low-boom research, neural networks have been applied to far-field-to-near-field inversion [17] and multi-fidelity prediction of near-field overpressure distributions [18]. Nevertheless, most existing methods still rely on high-dimensional end-to-end mappings, which leads to higher model complexity and training-data demand as well as the risk of local nonphysical oscillations.

To address these issues, this paper proposes a low-dimensional feature-mapping framework for near-field overpressure inversion. Rather than directly mapping a far-field pressure sequence to a near-field pressure sequence, the proposed method decomposes the inversion problem into low-dimensional compression of far-field ground waveforms, mapping from far-field features to near-field latent variables, and decoder-based reconstruction of near-field overpressure signatures. A one-dimensional convolutional autoencoder (1D-CAE) is used to construct a compact latent representation of near-field overpressure signatures. Segment-wise principal component analysis (PCA) is applied to extract far-field features while emphasizing boom-sensitive regions. A multilayer perceptron (MLP) is then used to map far-field features to near-field latent variables. The target near-field overpressure signature is reconstructed by the trained decoder and verified through forward propagation. By reformulating the high-dimensional inversion problem in a low-dimensional feature space, the proposed framework improves inversion stability and provides an efficient framework for generating target near-field overpressure signatures in conceptual-stage low-boom inverse design.

2. High-Fidelity Sonic-Boom Simulation

The near-field and far-field waveform data used in this study are generated through a high-fidelity sonic-boom simulation workflow. The workflow consists of two stages: CFD simulation for near-field overpressure signatures prediction and far-field sonic-boom propagation based on the augmented Burgers equation.

2.1 Near-Field CFD Simulation

To obtain high-fidelity near-field overpressure signatures, a zonal mesh composed of a near-surface region and a wave-capturing region is generated. The grid lines in the wave-capturing region are aligned with the Mach angle to reduce numerical dissipation during sonic-boom propagation. The Mach-cone-type mesh is shown in Figure 1.

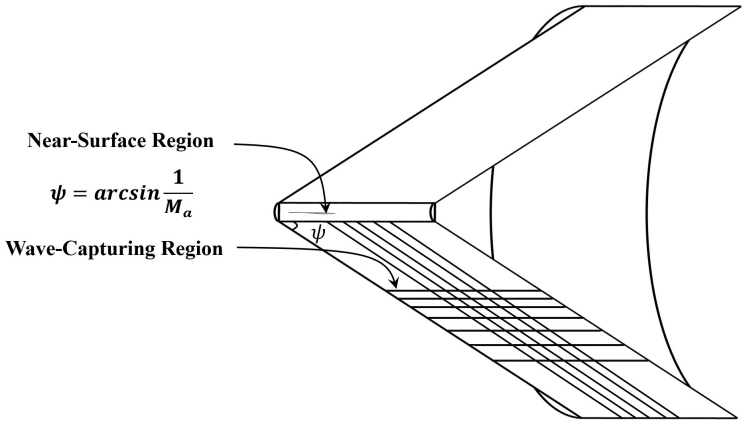


Figure 1 – Mach-cone-type mesh

The JAXA-Wing-Body (JWB) low-boom configuration from the Second AIAA Sonic Boom Prediction Workshop (SBPW-2) is selected as the test case [19], as shown in Figure 2. The reference length of the JWB configuration is 38.7 m, and the half-model reference area is 32.8 m². The total number of mesh cells is 22,428,911. The freestream condition follows the SBPW-2 specification: Mach number 1.60, cruise altitude 15,760 m, and angle of attack 2.3067°.

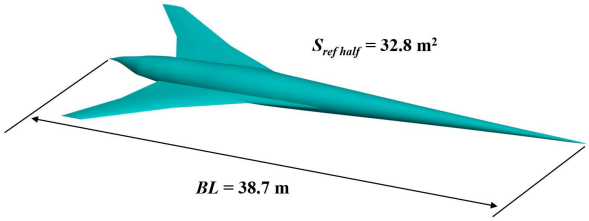


Figure 2 – Geometry of the JWB low-boom configuration

The CFD simulation is performed by solving the Reynolds-Averaged Navier-Stokes (RANS) equations coupled with the *k*-*ε* turbulence model. The simulated near-field overpressure signature is compared with the workshop reference data provided by the workshop in Figure 3. The CFD results agree well with the reference waveforms, reproducing the main shock locations, pressure peaks, troughs, and waveform recovery are accurately captured. Therefore, the CFD simulation provides reliable near-field source signatures for subsequent far-field sonic-boom propagation.

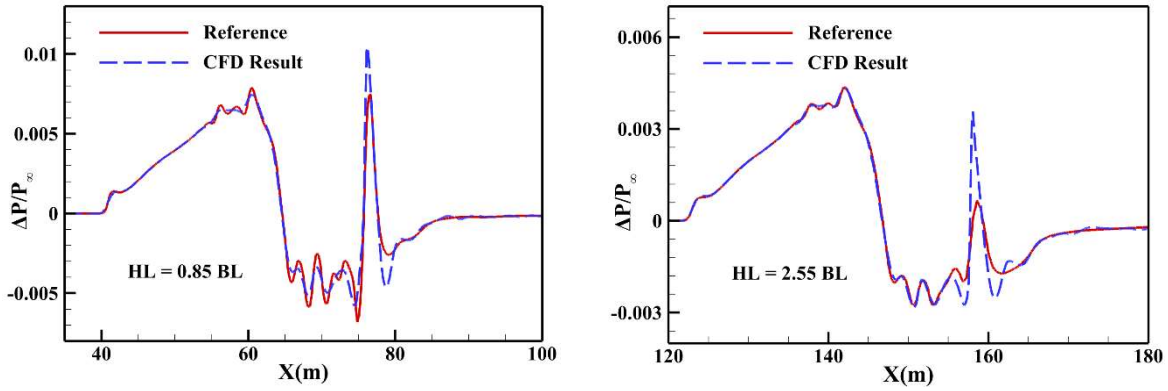


Figure 3 – Comparison of JWB near-field overpressure signature with reference data

2.2 Far-Field Sonic-Boom Propagation

After the near-field overpressure signature is obtained using the CFD simulation, it is propagated to the ground using bBoom, an in-house sonic-boom propagation tool based on the augmented Burgers equation [20]. The propagation model accounts for the main nonlinear and atmospheric effects, including geometric spreading, atmospheric stratification, thermoviscous absorption, and molecular relaxation. Through equation transformation, the fluid governing equations are converted into acoustic formulations. By neglecting third-order infinitesimal terms and considering geometric acoustic spreading and atmospheric effects, the one-dimensional propagation equation can be obtained as follows:

$$\frac{\partial p'}{\partial s} = -\frac{1}{2G}\frac{\partial G}{\partial s}p' + \frac{\beta p'}{\rho_0 c_0^3}\frac{\partial p'}{\partial t'} + \frac{\delta}{2c_0^3}\frac{\partial^2 p'}{\partial t'^2} + \sum_j \frac{(\Delta c)_j \tau_j}{c_0^2(1+\tau_i\frac{\partial}{\partial t'})}\frac{\partial^2 p'}{\partial t'^2} \tag{1}$$

where p' denotes sound pressure, s is the acoustic ray coordinate, and $t' = t - \int ds/c_0$ represents the retarded time of the sonic-boom waveform, with c_0 denoting the ambient speed of sound. The terms on the right-hand side describe the atmospheric effects considered during sonic-boom propagation, including geometric acoustics, atmospheric stratification, thermoviscous absorption, and molecular relaxation. In the acoustic formulation, sound propagation is assumed to follow acoustic rays.

Using this propagation method, bBoom propagates the near-field overpressure signature to the ground. As shown in Figure 4, the far-field ground waveform agrees well with the SBPW-2 reference data, confirming the validity of the sonic-boom simulation workflow.

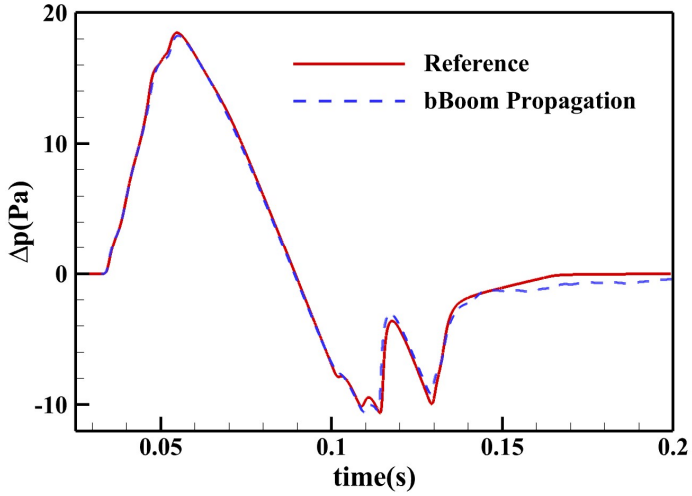


Figure 4 – Validation of bBoom far-field propagation against the SBPW-2 reference result

3. Low-Dimensional Mapping for Near-Field Overpressure Inversion

Figure 5 illustrates the proposed target-oriented framework, which consists of "target far-field ground waveform–low-dimensional feature mapping–near-field latent-variable prediction–decoder-based near-field reconstruction." In this study, near-field inversion refers to the overall process of inferring a feasible near-field overpressure signature from a prescribed far-field ground waveform, whereas reconstruction specifically denotes the decoder-based generation of the near-field waveform from predicted latent variables.

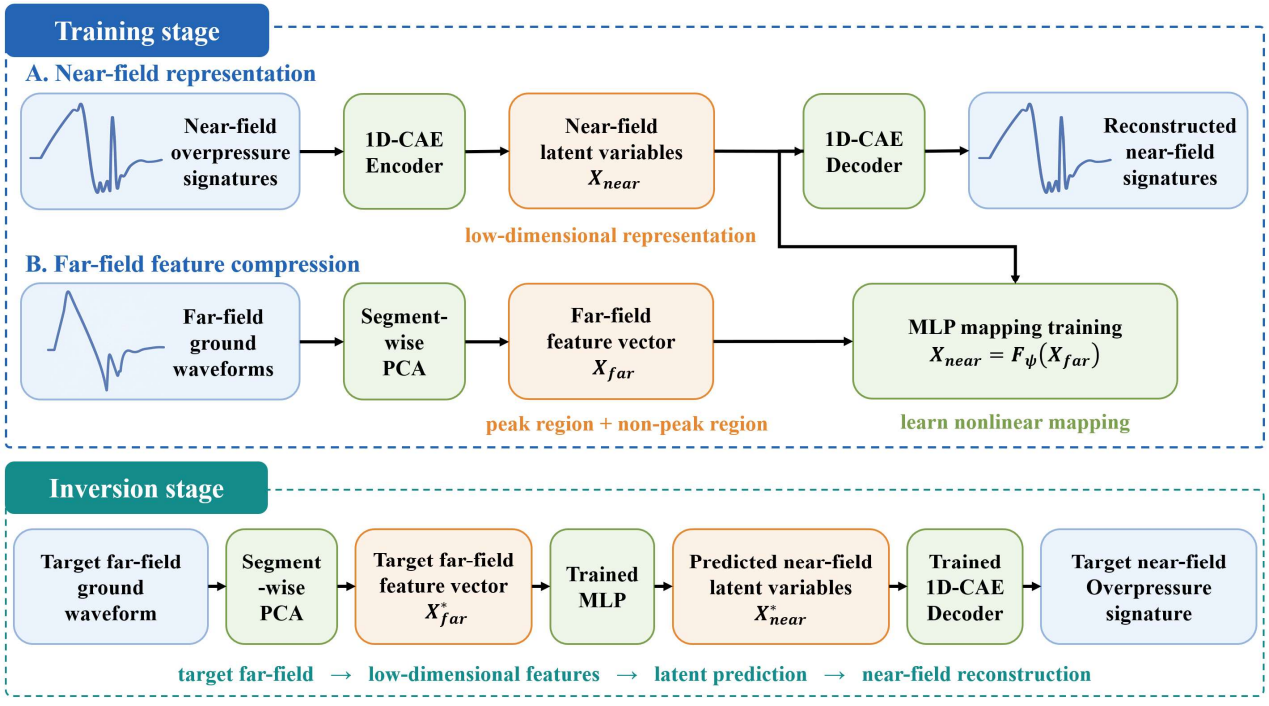


Figure 5 – Proposed near-field inversion framework (training & inversion stages)

The nonlinear mapping between the far-field features and near-field latent variables can be expressed as

$$X_{near} = F_{\psi}(X_{far}) \tag{2}$$

where X_{far} denotes the low-dimensional far-field feature vector obtained by segment-wise PCA compression of the far-field ground waveform, X_{near} denotes the near-field latent-variable vector encoded from the near-field overpressure signature by the 1D-CAE encoder, and F_{ψ} denotes the nonlinear mapping function learned by the MLP. For a prescribed target far-field ground waveform, the low-dimensional far-field features are first extracted, the corresponding near-field latent variables are then predicted, and the target near-field overpressure signature is reconstructed by the trained

1D-CAE decoder. Compared with an end-to-end method that directly maps a far-field pressure sequence to a near-field pressure sequence, this framework significantly reduces both the input and output dimensions, thereby improving inversion stability and reducing nonphysical waveform oscillations.

3.1 Sample Construction for Paired Near-/Far-Field Signatures

Paired samples consisting of near-field overpressure signatures and their corresponding far-field ground waveforms are first constructed. The baseline near-field overpressure signature is obtained by high-fidelity CFD simulation and then described using key points, shock positions, waveform amplitudes, and other parameters. A series of near-field overpressure signatures is generated within reasonable perturbation ranges of these baseline parameters, and each near-field sample is input into bBoom for far-field propagation to obtain the corresponding far-field ground waveform, as shown in Figure 6. The generated paired dataset is used for model training.

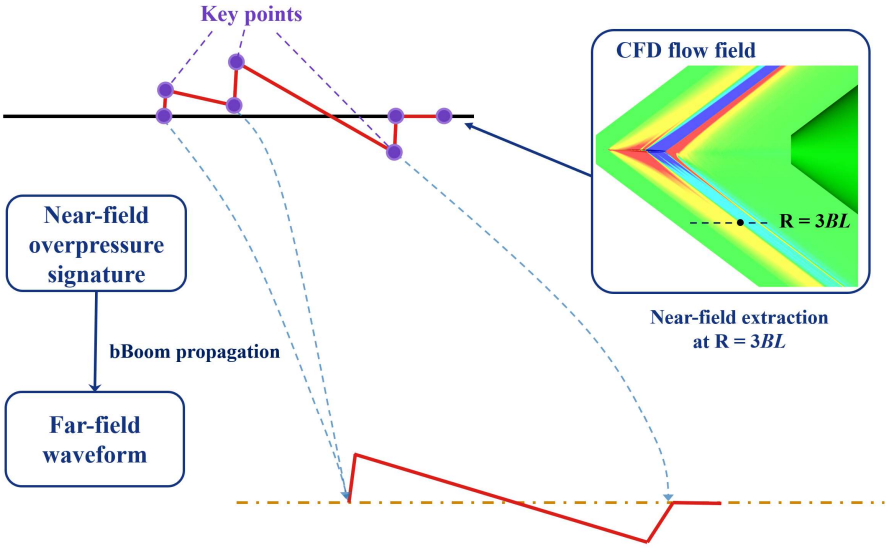


Figure 6 – Paired near-field/far-field sample generation

3.2 Near-Field Latent Extraction via 1D-CAE

The near-field overpressure signature contains multiple features, including the leading shock peak, negative-pressure trough, local oscillations, and tail decay. If the discrete near-field pressure sequence is directly used as the MLP output, the output dimension increases substantially, making the mapping more difficult to train. In addition, the network may ignore local waveform details, such as shocks, troughs, and tail oscillations, when optimized under a global-error constraint. Therefore, this paper uses a one-dimensional convolutional autoencoder (1D-CAE) to learn a low-dimensional representation of the near-field overpressure signature, compressing the original high-dimensional near-field sequence into latent variables that contain the main physical features and reconstructing the near-field overpressure signature through the decoder.

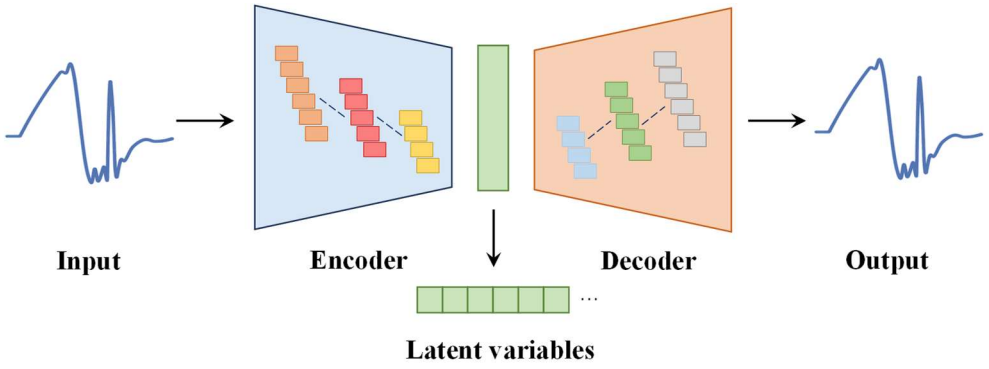


Figure 7 – Architecture of the 1D-CAE for near-field representation

The 1D-CAE consists of a convolutional encoder, a latent fully connected layer, and a transposed-convolution decoder, as shown in Figure 7. The encoder progressively extracts local waveform features of the near-field overpressure signature through multiple Conv1d, BatchNorm1d, and ReLU layers, while stride-based downsampling compresses the sequence length. The convolutional features are then flattened and mapped to a low-dimensional latent space through a fully connected layer, yielding the near-field latent variables. In the decoder, the latent variables are first back-projected into the convolutional feature space through a fully connected layer, and ConvTranspose1d, BatchNorm1d, and ReLU layers are subsequently used for stepwise upsampling to recover the near-field overpressure waveform. A Sigmoid activation function is adopted in the output layer to obtain the normalized reconstruction. Because repeated downsampling and upsampling may cause deviations in sequence length, an automatic length-alignment mechanism is incorporated in the forward pass to ensure that the reconstructed output and the input signal have the same number of discrete points.

3.3 Far-Field Feature Compression via Segment-Wise PCA

Far-field ground waveforms usually contain many discrete sampling points. Using them directly as the MLP input would significantly increase the network parameter scale and training difficulty. Meanwhile, different regions of a far-field ground waveform contribute differently to the perceived loudness of the ground sonic boom. The peak, trough, shock rising edge, and tail-recovery region are closely related to sonic-boom intensity, whereas relatively smooth regions mainly reflect the overall waveform trend. If a unified global PCA is used for dimensionality reduction, it can achieve a small overall reconstruction error, but the principal components are mainly determined by regions with larger variance over the full waveform. Local features near peaks, weak-shock regions, and tail-recovery processes may therefore be weakened because of their short duration and relatively small sampling proportion.

Accordingly, segment-wise PCA is adopted in this paper to compress the far-field ground waveform into low-dimensional far-field features. As shown in Figure 8, sonic-boom-sensitive regions are first identified according to the amplitude statistics of the far-field ground waveforms in the training set. The peak and its neighboring sampling points are defined as the peak region, while the remaining sampling points are defined as the non-peak region. PCA models are then built separately for the two regions to obtain the principal-component coefficients of the peak and non-peak regions. Finally, the two sets of coefficients are concatenated to form the low-dimensional far-field features. While controlling the input dimension, this method assigns greater representational capacity to the peak, shock, and local waveform-variation regions, thereby improving the stability and effectiveness of the nonlinear mapping from far-field features to near-field latent variables.

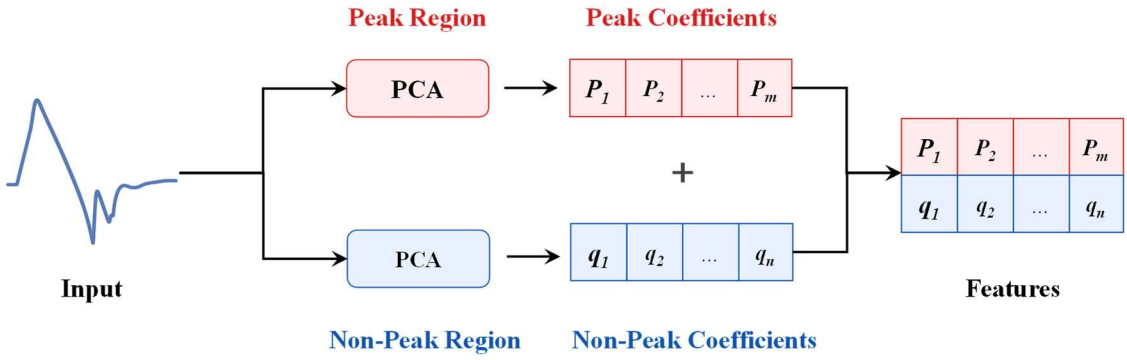


Figure 8 – Segment-wise PCA for low-dimensional far-field features

3.4 Nonlinear Far-to-Near Mapping with MLP

After obtaining the low-dimensional far-field features and near-field latent variables, a multilayer perceptron (MLP) is constructed to learn the nonlinear mapping between them. The MLP input is the low-dimensional far-field features obtained by segment-wise PCA, and the output is the near-field latent variables extracted by the 1D-CAE. The network consists of an input layer, several hidden layers, and an output layer. Nonlinear activation functions are used between layers to fit complex nonlinear relationships, as shown in Figure 9.

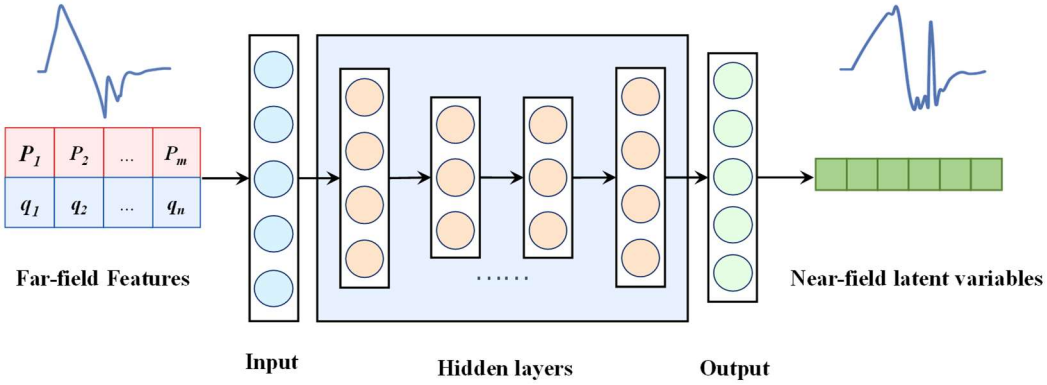


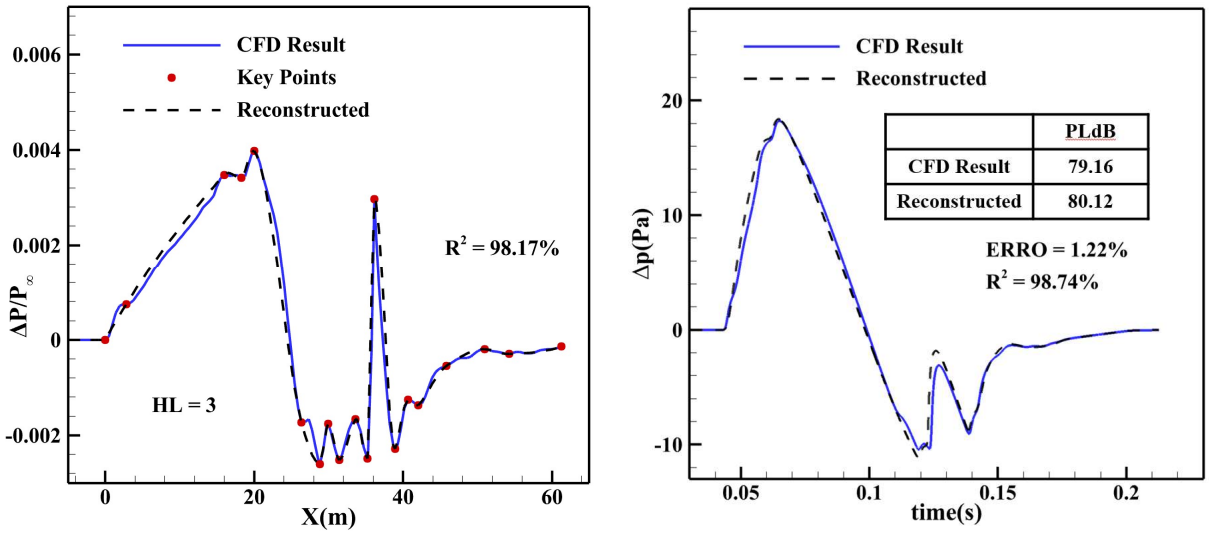
Figure 9 – MLP mapping from low-dimensional far-field features to near-field latent variables

Compared with directly predicting the complete near-field pressure sequence, this strategy significantly reduces the output dimension and assigns the complex waveform reconstruction task to the trained decoder, thereby improving inversion stability. Given a target far-field ground waveform, the inversion process is decomposed into three sequential steps: extraction of low-dimensional far-field features, MLP-based prediction of near-field latent variables, and decoder-based reconstruction of the near-field overpressure signature. Because the MLP predicts near-field latent variables rather than the complete near-field pressure sequence, the decoder maps the predicted latent variables back to the near-field overpressure-signature space. This helps suppress local nonphysical oscillations that are prone to occur in direct high-dimensional prediction and improves the smoothness and propagation verifiability of the generated target near-field overpressure signature.

4. JWB-Based Evaluation of the Inversion Framework

4.1 JWB Near-Field Parameterization and Sample Space

The JWB near-field overpressure signature extracted at R = 3BL is parameterized with 19 key points and reconstructed by Akima interpolation. As verified in Figure 10, the parameterized near-field overpressure signature preserves the main features of the CFD simulation, achieving R^2 = 98.17% for near-field reconstruction and R^2 = 98.74% for the propagated far-field waveform. The relative PLdB error is only 1.22%, indicating that the parameterization is sufficiently accurate for sample generation.



(a) Near-field reconstruction (b) Far-field propagation

Figure 10 – Near-field parameterization and propagation verification

To construct the sample space shown in Figure 11, the boundary key points are fixed, while the ordinates of the remaining 17 key points are perturbed by ±15% and sampled using LHS. After envelope correction, 3000 paired near-field/far-field samples are generated through Akima reconstruction and bBoom propagation, forming the training basis for the inversion model.

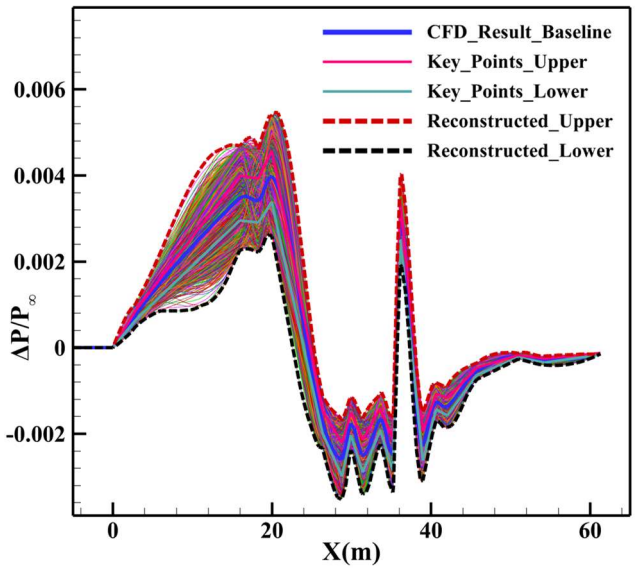


Figure 11 – Near-field overpressure sample space

4.2 Evaluation of 1D-CAE Latent Representation

The 1D-CAE is evaluated to determine a compact but accurate near-field latent-variable space. The 150-point near-field overpressure signatures are preprocessed into 136-point inputs, and the dataset is split into training, validation, and test sets at a ratio of 7:2:1, with normalization fitted only on the training set. As shown in Table 1, the 32-64-128 encoder configuration achieves the best balance between reconstruction accuracy and model complexity.

Table 1 – 1D-CAE reconstruction performance for different encoder configurations

Encoder convolution-block channel configuration	R^2	RMSE
16-32-64	99.89%	5.96×10^{-5}
32-64-128	**99.94%**	$\mathbf{4.41 \times 10^{-5}}$
16-32-64-128	99.94%	4.54×10^{-5}
32-64-128-256	99.89%	5.94×10^{-5}
16-32-64-128-256	99.92%	5.28×10^{-5}

Latent-dimension sensitivity is further examined over 4–48 dimensions. The reconstruction accuracy improves with increasing latent dimension but becomes nearly saturated after 32 dimensions. Therefore, the final 1D-CAE adopts 32 latent variables and obtains a root mean square error (RMSE) of 1.54×10^{-5} and an R^2 of 99.99% on the test set, confirming high-fidelity near-field representation.

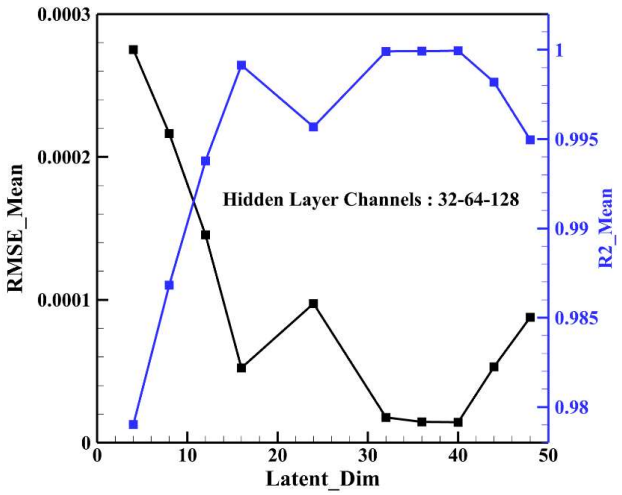


Figure 12 – Latent-dimension sensitivity

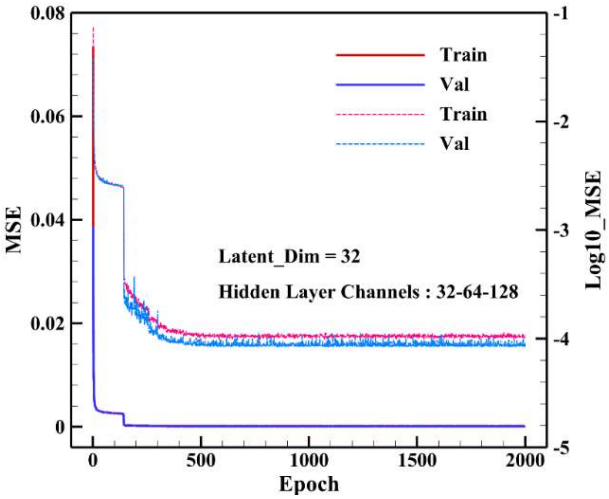


Figure 13 – 1D-CAE training convergence

Figure 14 further shows that the reconstructed samples reproduce the major peaks, troughs, and tail oscillations. Thus, the trained decoder provides a reliable near-field reconstruction space for the subsequent mapping from low-dimensional far-field features to near-field latent variables.

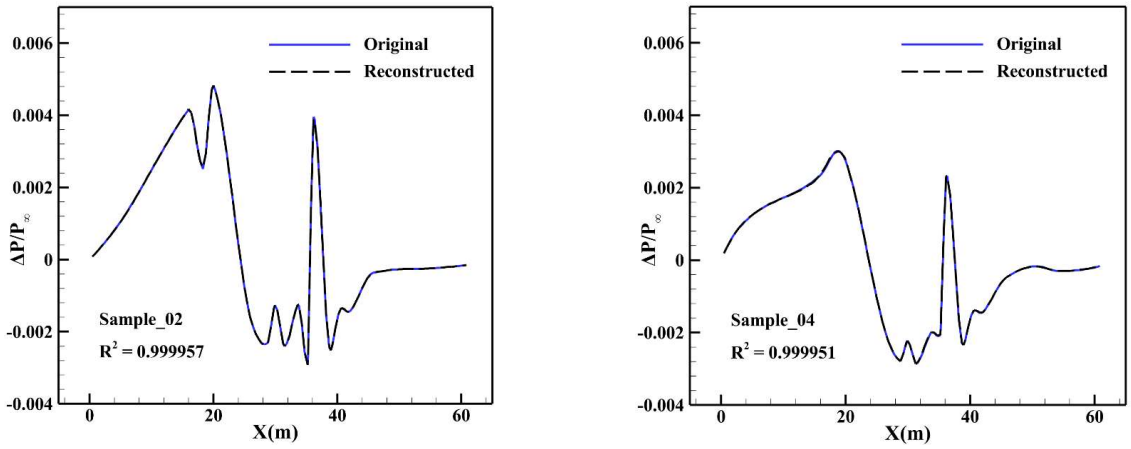


Figure 14 – Reconstruction of two representative test samples

4.3 Evaluation of Segment-Wise PCA Features

Segment-wise PCA is used to compress the approximately 1,945-point far-field waveform into low-dimensional far-field features. A sensitivity analysis of 40–100 principal components shows that the reconstruction improvement becomes limited after 80 components; therefore, an 80-dimensional representation is adopted (Figure 15).

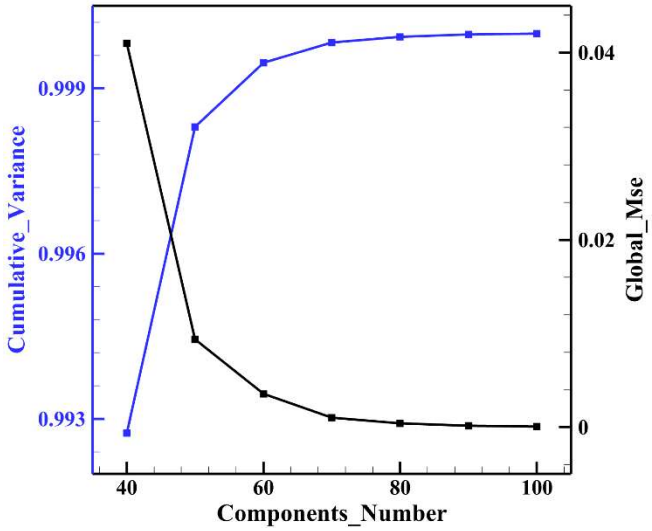


Figure 15 – Far-field PCA dimension sensitivity

Based on the training-set amplitude distribution, the peak-sensitive region is identified using the 98.5% quantile and expanded by 10 sampling points on each side. The final representation assigns 20 components to the peak region and 60 to the non-peak region, preserving peak, trough, and tail-recovery features while controlling the MLP input dimension (Figure 16).

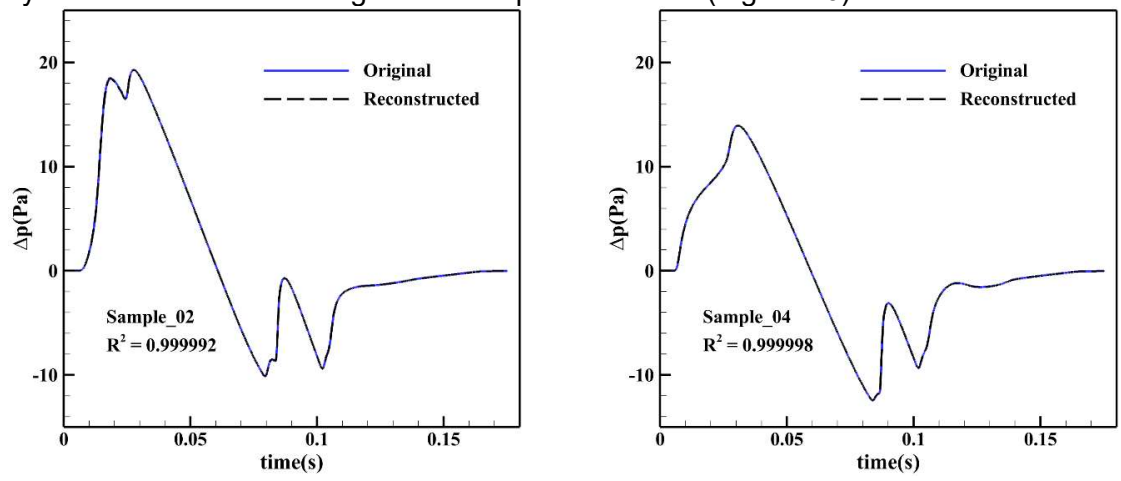


Figure 16 – Segment-wise PCA reconstruction of far-field samples

4.4 Training and Testing of the MLP Inversion Mapping

The MLP is used only to learn the nonlinear mapping from the 80-dimensional far-field features to the 32-dimensional near-field latent variables, because the 1D-CAE decoder has already been shown to reconstruct near-field overpressure signatures accurately.

The network contains four hidden layers with 256, 192, 128, and 64 neurons and is trained for 2,000 epochs using AdamW, with a learning rate of 8×10^{-4}, a weight decay of 5×10^{-5}, a batch size of 128, and MSE loss. The convergence in Figure 17 indicates stable training.

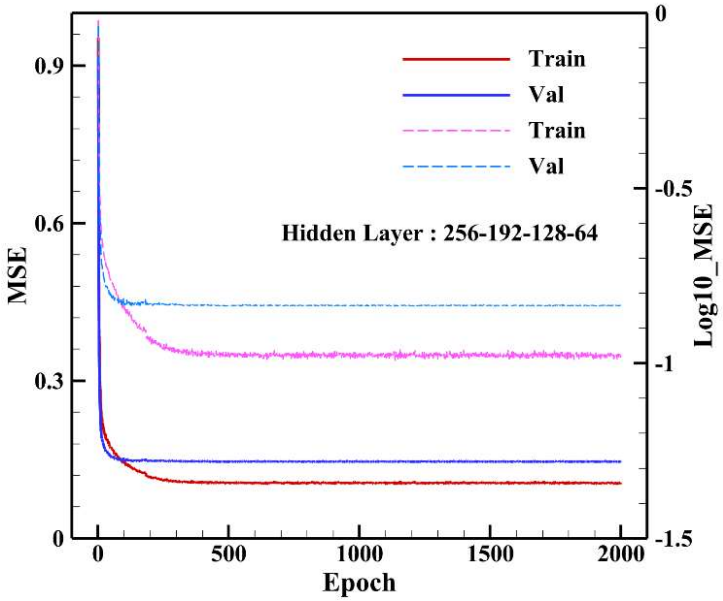
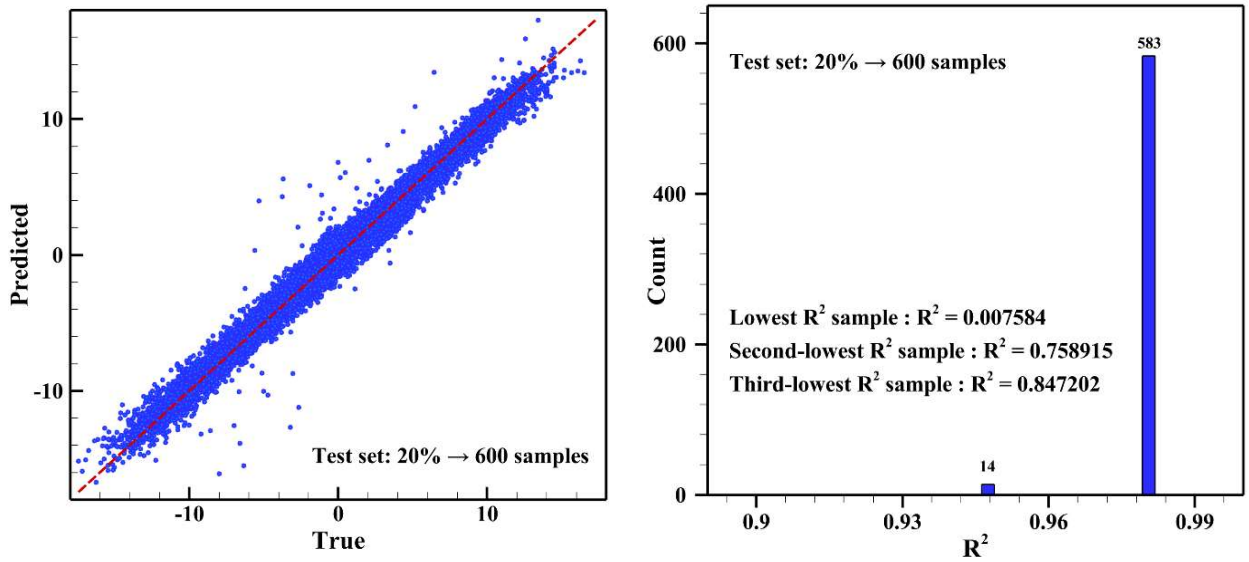


Figure 17 – Training convergence of the MLP mapping model

During testing, predictions are inverse-transformed and clipped within the training-set latent-variable bounds to avoid unstable extrapolation. The 19,200 predicted latent-variable values are distributed close to the 45° line, and only three test samples exhibit relatively large prediction errors, demonstrating that the mapping accuracy is sufficient for target near-field generation (Figure 18).



(a) Latent-variable prediction consistency (b) Test-sample R^2 distribution

Figure 18 – MLP prediction performance for near-field latent variables

5. Low-Boom Target Waveform Construction and Near-Field Inversion Verification

5.1 Sonic-Boom Characteristics of Typical Waveforms

To establish the acoustic rationale for the target-waveform construction, six representative far-field sonic-boom waveforms are examined, including the N-wave sonic boom (NWSB), flat-top sonic boom (FTSB), ramped-front sonic boom (RFSB), multi-shock sonic boom (MSSB), smoothed low-peak

multi-shock sonic boom (SLP-MSSB), and reduced-shock wide-tail multi-shock sonic boom (RSWT-MSSB), as shown in Figure 19. These waveforms represent different low-boom control mechanisms, from shock-strength redistribution to pressure-recovery reshaping, and therefore provide a reference for identifying the far-field features that should be modified in the JWB target waveform.

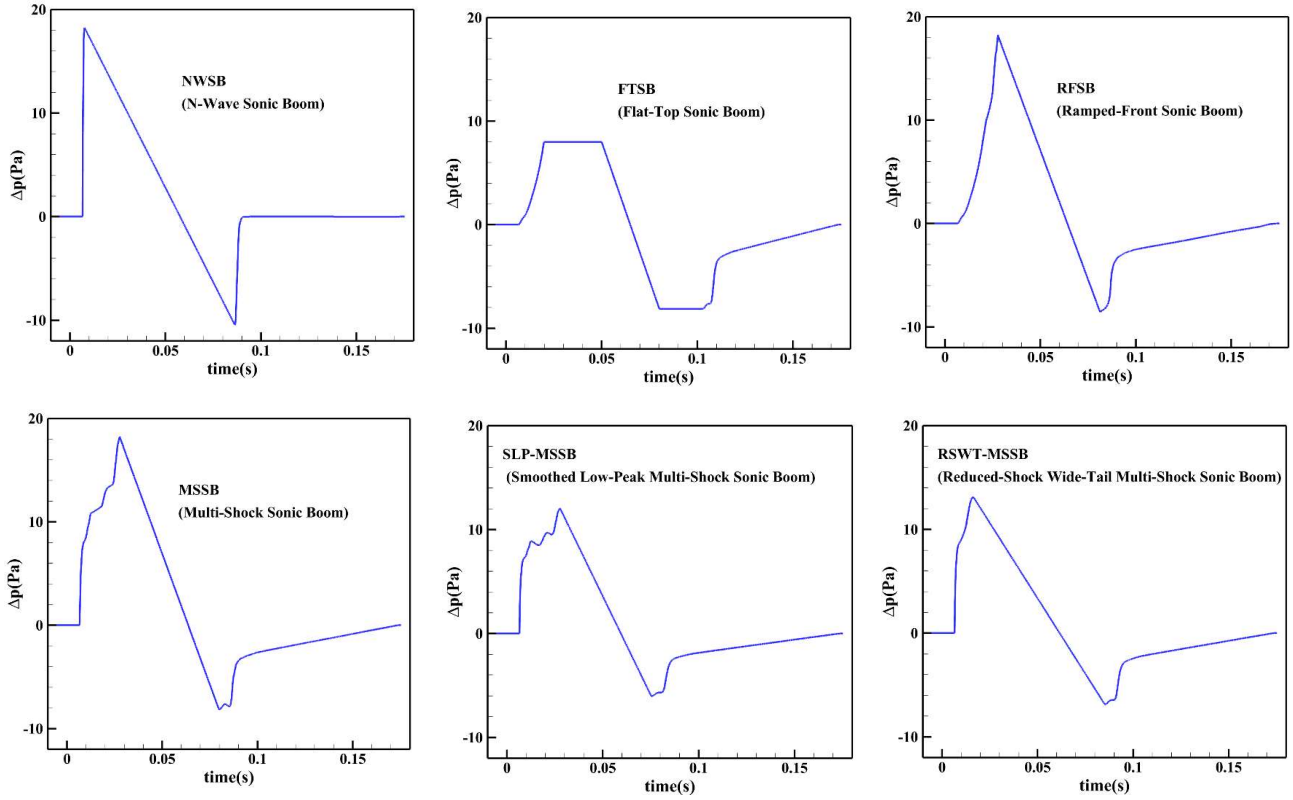


Figure 19 – Six typical far-field ground waveforms

The PLdB comparison in Table 2 indicates that effective loudness reduction cannot be achieved by modifying a single waveform feature alone. The NWSB produces the highest PLdB because the pressure rise and recovery are concentrated into strong leading and trailing shocks. FTSB and RFSB reduce loudness by blunting the peak or weakening the leading shock, but their reductions remain limited when the peak-trough amplitude and tail-recovery behavior are not controlled simultaneously. MSSB redistributes strong shocks into multiple weaker shocks, whereas SLP-MSSB further smooths local peaks and RSWT-MSSB combines reduced shock strength with a widened recovery tail. This trend shows that a practical low-boom target should simultaneously consider leading-shock blunting, peak-trough control, weak-shock redistribution, and tail-recovery smoothing.

Table 2 – PLdB of typical far-field ground waveforms

Waveform type	PLdB	Waveform type	PLdB
NWSB	95.81	FTSB	85.80
RFSB	91.15	MSSB	93.31
SLP-MSSB	88.28	RSWT-MSSB	83.46

5.2 JWB Low-Boom Target Waveform

Following the above acoustic criterion, the JWB low-boom target far-field waveform is constructed from the baseline far-field waveform through localized and physically constrained waveform modification. The objective is not to obtain a lower PLdB by simple amplitude scaling, but to reduce loudness-sensitive local gradients while preserving the main time-domain topology of the original JWB waveform. Therefore, the initial and final pressure levels, the overall pressure-recovery trend, and the principal waveform duration are retained, whereas the leading-shock region, peak-trough transition, and tail-recovery segment are selectively adjusted, as illustrated in Figure 20.

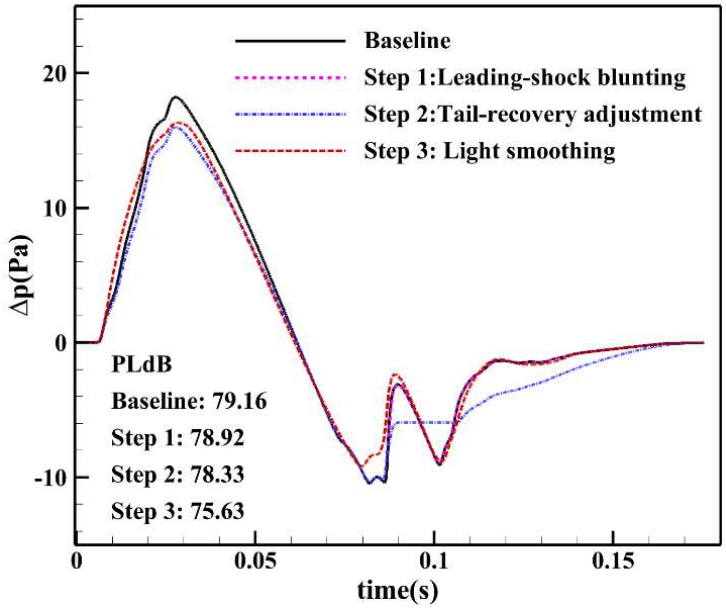
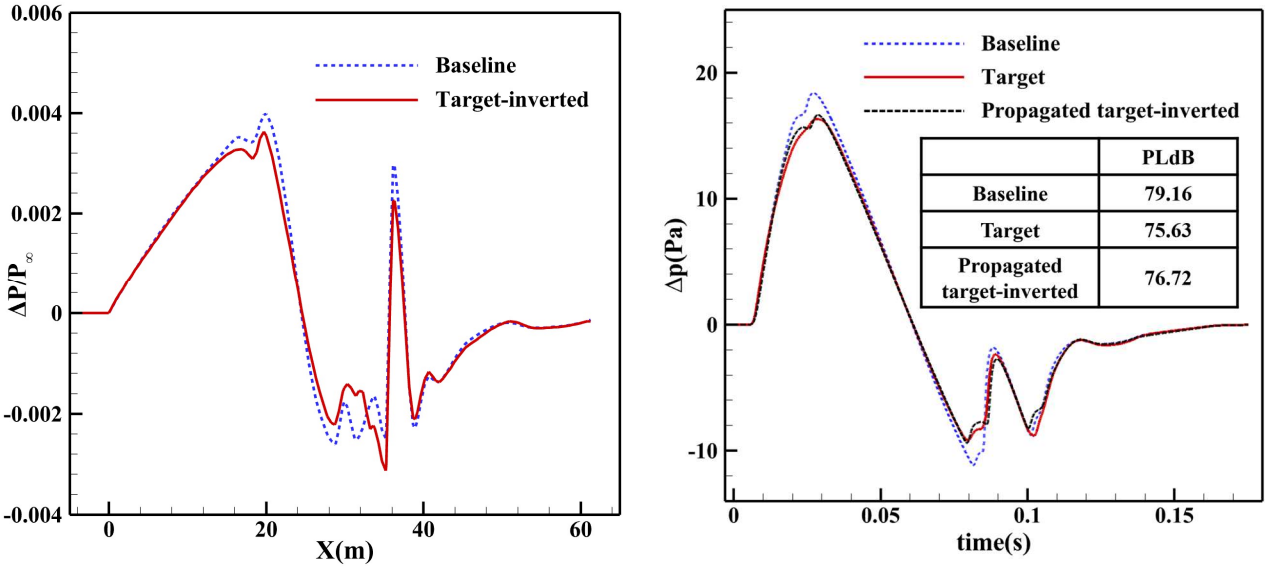


Figure 20 – Construction of the JWB low-boom target waveform

Three sequential modification steps are introduced. First, peak reduction and slope blunting are applied to the leading-shock region to weaken the rapid concentration of waveform energy, reducing PLdB from 79.16 to 78.92. Second, the pressure-recovery process after the negative-pressure trough is smoothed to reduce the contribution of the tail region, further decreasing PLdB to 78.33. Third, light smoothing is applied to suppress residual local discontinuities without changing the overall waveform morphology, yielding the final target PLdB of 75.63. Compared with the baseline waveform, the constructed target achieves a 3.53 PLdB reduction while maintaining a morphology compatible with the JWB sample space, and is therefore used as the prescribed far-field input for near-field inversion.

5.3 Near-Field Inversion and bBoom Verification

Given the constructed low-boom target far-field waveform, the trained inversion framework is applied without retraining. The target pressure sequence is first processed using the same standardization and segment-wise PCA models fitted on the training set to obtain an 80-dimensional far-field feature vector. The trained MLP then predicts the corresponding 32-dimensional near-field latent variables, and the 1D-CAE decoder reconstructs the target near-field overpressure signature. In this way, the prescribed far-field acoustic objective is converted into a near-field source signature that can be verified through forward propagation.



(a) near-field comparison (b) far-field verification via bBoom

Figure 21 – Near-field inversion result and far-field propagation verification

The target-inverted near-field overpressure signature is compared with the baseline near-field overpressure signature in Figure 21(a). The main shock locations, negative-pressure trough, and tail-oscillation pattern are preserved, indicating that the reconstructed signature remains consistent with the aerodynamic source characteristics represented by the JWB sample space. Meanwhile, local amplitude adjustments appear in the leading-shock and recovery regions, reflecting the low-boom modifications imposed on the target far-field waveform rather than a simple scaling of the baseline signature. To verify propagation consistency, the inverted near-field overpressure signature is propagated to the ground using bBoom. As shown in Figure 21(b), the propagated far-field waveform approaches the prescribed target in peak value, trough value, and tail-recovery trend. The resulting loudness is 76.72 PLdB, which is 2.44 PLdB lower than the baseline value of 79.16 PLdB and 1.09 PLdB higher than the target value of 75.63 PLdB. The remaining discrepancy is mainly attributed to MLP latent-variable prediction errors and decoder reconstruction errors. Nevertheless, the result demonstrates that the proposed feature-to-latent inversion framework can generate a propagation-consistent target near-field overpressure signature with reduced sonic-boom loudness from a prescribed low-boom far-field waveform.

6. Conclusions and Outlook

This paper proposes a low-dimensional feature-mapping framework for near-field overpressure signature inversion from prescribed far-field ground waveforms. The high-dimensional far-field-to-near-field inversion problem is transformed into a low-dimensional nonlinear mapping between segment-wise PCA far-field features and 1D-CAE near-field latent variables. An MLP is used to establish the corresponding relationship between far-field features and near-field latent variables, and the trained decoder reconstructs the target near-field overpressure signature, thereby providing an efficient target near-field generation route for conceptual low-boom inverse design.

For the JWB configuration, the 1D-CAE reconstructs near-field overpressure signatures with high accuracy, segment-wise PCA preserves key waveform structures in an 80-dimensional representation, and the MLP gives stable predictions on the test set. After the low-boom target far-field waveform is inverted into a near-field overpressure signature and propagated to the ground, the resulting PLdB is 76.72, which is 2.44 PLdB lower than the baseline value of 79.16 PLdB and 1.09 PLdB higher than the target value of 75.63 PLdB. These results verify the feasibility of the proposed method.

The present sample space is mainly based on a single JWB configuration, and the geometric realizability of the inverted signature still needs to be verified by combining shape optimization with high-fidelity CFD. Future work will extend the dataset to different configurations and flight conditions to improve model generalization, introduce physical constraints and multi-solution screening to enhance inversion robustness, and construct a closed-loop low-boom design process consisting of target far field, target near field, shape optimization, and propagation verification.

7. Acknowledgment

This work was supported by the National Key Research and Development Program of China under Grant No. 2023YFB3002800 and the National Natural Science Foundation of China under Grant No. 52472385.

8. Contact Author Email Address

YuHuan Peng: pyhuan@mail.nwpu.edu.cn

Keshi Zhang: zhangkeshi@nwpu.edu.cn

Qing Chen: cqaero@mail.nwpu.edu.cn

Zhonghua Han: hanzh@nwpu.edu.cn

9. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] Han Z H, Qiao J L, Zhang L W, et al. Recent progress of efficient low-boom design and optimization methods. *Progress in Aerospace Sciences*, 2024, 146: 101007. DOI:10.1016/j.paerosci.2024.101007.

[2] Bonavolonta G, Lawson C, Riaz A. Review of sonic boom prediction and reduction methods for next generation of supersonic aircraft. *Aerospace*, 2023, 10(11): 917. DOI:10.3390/aerospace10110917.

[3] Park M A, Carter M B. Low-Boom Demonstrator near-field summary for the Third AIAA Sonic Boom Prediction Workshop. *Journal of Aircraft*, 2022, 59(3): 563–577. DOI:10.2514/1.C036323.

[4] Morgenstern J, Norstrud N, Stelmack M, et al. Final report for the Advanced Concept Studies for Supersonic Commercial Transports Entering Service in the 2030 to 2035 Period, N+3 Supersonic Program. *NASA John H. Glenn Research Center at Lewis Field*, 2010.

[5] Minelli A, el Din I S, Carrier G. Inverse design approach for low-boom supersonic configurations. *AIAA Journal*, 2014, 52(10): 2198–2212. DOI:10.2514/1.J052834.

[6] Ding Y L, Han Z H, Qiao J L, et al. Inverse design method for low-boom supersonic transport with lift constraint. *AIAA Journal*, 2023, 61(7): 2840–2853. DOI:10.2514/1.J062410.

[7] Chen Q, Han Z H, Zhang K S, et al. Revised Whitham theory for enhanced sonic boom prediction without equivalent area calculation. *AIAA Journal*, 2026, (in press). DOI:10.2514/1.J066640.

[8] Ma C, Huang J T, Liu G, et al. Inverse near-field sonic boom signature prediction for supersonic aircraft. *Journal of Aerodynamics*, 2023, 41(4): 1–10. DOI:10.7638/kqdlxxb-2021.0411.

[9] Gu Y, Huang J T, Chen S, et al. Sonic boom inversion technology based on inverse augmented Burgers equation. *Acta Aeronautica et Astronautica*, 2023, 44(2). DOI:10.7527/S1000-6893.2021.26258.

[10] Zhang Y, Huang J, Gao Z, et al. Inverse design of low boom configurations using proper orthogonal decomposition and augmented Burgers equation. *Chinese Journal of Aeronautics*, 2019, 32(6): 1380–1389. DOI:10.1016/j.cja.2019.03.017.

[11] Liu X, Yang S, Sun H, et al. Review of deep learning-based aerodynamic shape surrogate models and optimization for airfoils and blade profiles. *Physics of Fluids*, 2025, 37(4). DOI:10.1063/5.0268466.

[12] Brunton S L, Kutz J N, Manohar K, et al. Data-driven aerospace engineering: Reframing the industry with machine learning. *AIAA Journal*, 2021, 59(8): 2820–2847. DOI:10.2514/1.J060131.

[13] Guo X, Li W, Iorio F. Convolutional neural networks for steady flow approximation. *Proc 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2016: 481–490.

[14] Fukami K, Fukagata K, Taira K. Super-resolution reconstruction of turbulent flows with machine learning. *Journal of Fluid Mechanics*, 2019, 870: 106–120. DOI:10.1017/jfm.2019.238.

[15] Sekar V, Zhang M, Shu C, et al. Inverse design of airfoil using a deep convolutional neural network. *AIAA Journal*, 2019, 57(3): 993–1003. DOI:10.2514/1.J057894.

[16] Yilmaz E, German B J. Conditional generative adversarial network framework for airfoil inverse design. *Proc AIAA Aviation Forum*, 2020: 2020-3185.

[17] Chen S, Qiu J, Yang H, et al. Deep learning for inverse design of low-boom supersonic configurations. *Advances in Aerodynamics*, 2023, 5(1): 247–266. DOI:10.1186/s42774-023-00145-1.

[18] Wang Y T, Luo X, Liu H, et al. Sonic boom prediction of supersonic passenger aircraft based on multi-fidelity deep neural network. *Acta Aeronautica et Astronautica*, 2025, 46(20): 1–21. DOI:10.7527/S1000-6893.2025.31936.

[19] Park M A, Nemec M. Nearfield summary and statistical analysis of the Second AIAA Sonic Boom Prediction Workshop. *Journal of Aircraft*, 2019, 56(3): 851–875. DOI:10.2514/1.C034866.

[20] Qiao J L, Han Z H, Zhang L W, et al. Far-field sonic boom prediction considering atmospheric turbulence effects: An improved approach. *Chinese Journal of Aeronautics*, 2022, 35(9): 208–225. DOI:10.1016/j.cja.2021.01.013.