

PARAMETER INVERSION ALGORITHM FOR AERO-ENGINE γ-PHOTON FLOW FIELD IMAGES BASED ON UNET-GAN ARCHITECTURE

Hou LiYing[1,2], Liu JianTang[1], Song YuLu[1]& Xiao Hui[1,2] ✉
([1]College of Automation Engineering, Nanjing University of Aeronautics and Astronautics, 29 Jiangjun Road, Nanjing, China.
[2]Shenzhen Research Institute, Nanjing University of Aeronautics and Astronautics, Shenzhen, China.
xiaohui@nuaa.edu.cn)

Abstract

γ-photon emission tomography offers a promising non-intrusive solution for monitoring the internal flow fields of aero-engine nozzles. However, the stochastic nature of photon detection results in sparse grayscale density maps that lack direct thermodynamic insights. To address this, this paper proposes a high-precision parameter inversion algorithm based on a Conditional Generative Adversarial Network (cGAN) architecture. The framework employs a U-Net generator with a ResNet101 backbone to map probabilistic photon distributions to continuous multi-channel flow parameters. A novel Physical Aware Enhancement (PAE) module, integrating multi-scale dilated convolutions and dual-attention mechanisms, is embedded to suppress statistical noise and enforce gradient continuity, effectively eliminating non-physical artifacts. Experimental validation on the SW60B turbojet engine demonstrates that the proposed method achieves a Peak Signal-to-Noise Ratio (PSNR) of 33.86 dB and a Structural Similarity Index (SSIM) of 0.86. Comparative analysis confirms its superiority over conventional generative models in maintaining the structural integrity and physical fidelity of the reconstructed flow fields.

Keywords: Aero-engine monitoring, Parameter inversion, GANs, Flow field reconstruction

1. Introduction

Real-time knowledge of the thermodynamic state inside an operating aero-engine nozzle is central to predictive maintenance, performance verification, and the development of next-generation propulsion architectures. Conventional intrusive measurement methods, including thermocouple rakes, pressure probes, and hot-wire anemometers, are fundamentally constrained by the extreme thermal and mechanical environment of the nozzle section: sensor lifetime is severely limited at gas temperatures exceeding 1,500 K, probe-induced flow blockage alters the quantity being measured, and spatial resolution is inherently confined to discrete probe locations [1, 2]. Non-contact optical diagnostics such as schlieren imaging, laser Doppler velocimetry, and planar laser-induced fluorescence circumvent probe-induced disturbance but require transparent optical access windows that are structurally incompatible with the metallic casings of production aero-engines, introducing considerable integration complexity that precludes routine deployment [3].

γ-photon flow-field imaging exploits the high penetrating power and electrical neutrality of 511 keV annihilation photons to perform volumetric flow measurements through opaque engine casings without any physical contact. When a β^+-emitting radionuclide tracer introduced into the engine airflow undergoes positron decay, the emitted positron thermalises and annihilates with a local electron, producing two collinear 511 keV gamma-ray photons emitted in opposite directions. An external ring of coincidence detectors captures simultaneous photon pairs, each defining a line of response (LOR) through the annihilation site. By statistically accumulating millions of LORs and applying the Ordered Subsets Expectation Maximisation (OSEM) reconstruction algorithm, a two-dimensional map of positron-annihilation density is recovered, as illustrated in Figure 1 [4, 5]. While this grayscale positron-flux image is statistically proportional to the local tracer concentration and hence to the local flow density, it contains no direct spectral encoding of temperature, pressure, or

velocity. Recovering these thermodynamic quantities from the grayscale observation defines what the present paper terms the parameter inversion problem.

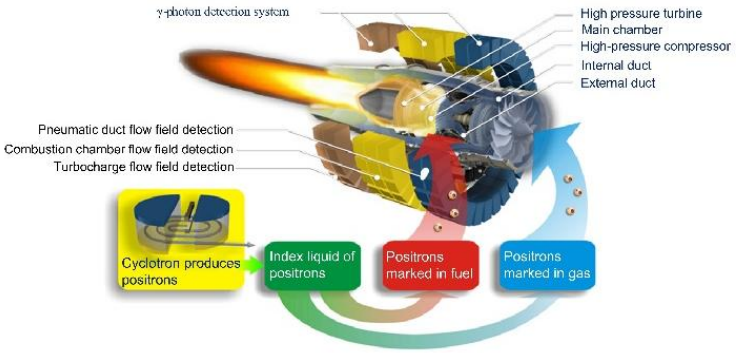


Figure 1 – Schematic diagram of the γ-photon online monitoring system for aero-engine flow fields.

Parameter inversion from γ-photon flow-field images is inherently ill-posed for two principal reasons. First, the probabilistic nature of positron emission and coincidence detection introduces statistical noise and spatial blur into the reconstructed image that are absent from the smooth physical fields sought: the sharp edges present in the true thermodynamic field are blurred into continuous probability density gradients by the Gaussian-like point spread function of the imaging system. Second, the multi-channel parameter output encoding spatially resolved distributions of temperature, pressure, and velocity simultaneously demands global coherence across the entire image while requiring local precision at physically active boundaries such as jet edges, shear layers, and thermal gradient zones [6]. Conditional Generative Adversarial Networks (cGANs) offer a principled approach to this problem: unlike regression-based networks that predict the conditional mean and produce blurred outputs, cGANs learn the full conditional distribution of the target field given the input observation, with the adversarial objective promoting sharp, high-frequency outputs consistent with the true distribution [7, 8, 9, 10].

However, directly applying standard cGAN architectures to γ-photon flow-field parameter inversion leaves three domain-specific challenges unaddressed. First, the probabilistic imaging process produces noise patterns that a generator lacking explicit noise suppression will propagate or amplify in the recovered parameter field. Second, thermodynamic consistency requires long-range spatial dependencies, as the temperature distribution at the nozzle periphery is physically coupled to the centreline distribution, yet standard convolutional encoders with limited receptive fields cannot efficiently model these cross-regional relationships. Third, the PatchGAN-style discriminator evaluates local authenticity at fixed patch scales but remains insensitive to smooth-gradient violations at thermodynamic boundaries, which is precisely the artefact category that erroneous parameter inversion produces [11]. The present paper addresses all three challenges through the Physical Aware Enhancement (PAE) module, a multi-component attention block inserted symmetrically into both the generator and the discriminator. The PAE module integrates four-branch multi-scale dilated convolution for noise-robust feature construction, a NonLocalBlock for global spatial dependency modelling, and CBAM dual-attention for physically focused feature refinement. The full network is optimised under WGAN-GP with an adaptive-weight VGG perceptual loss evaluated at the relu3_3 feature layer, and trained on a high-fidelity paired dataset constructed by coupling ANSYS CFD simulations with GATE Monte Carlo γ-photon transport modelling for the SW60B turbojet engine.

2. Methodology

2.1 Network Architecture

The overall cGAN framework is illustrated in Figure 2. The generator G adopts the encoder–decoder topology of U-Net [12], modified by substituting the standard convolutional encoder with a ResNet101 backbone [13] pre-trained on ImageNet. This transfer learning strategy provides hierarchical feature representations spanning a wide range of spatial scales: shallow layers encode low-level luminance edges and fine texture, intermediate layers capture mid-scale geometric

structure, and deep layers encode high-level semantic content. The decoder employs PixelShuffle sub-pixel convolution [14] with ICNR weight initialisation for upsampling, which rearranges feature channels into spatial pixels and is free of the checkerboard artefacts that transposed convolutions introduce through uneven gradient overlap. At each encoder–decoder interface, a two-layer channel expansion block (ratio ni → 2ni → ni) consolidates cross-channel interactions before the features are passed through skip connections that concatenate high-resolution encoder feature maps with upsampled decoder activations. Three PAE modules are embedded at the low-, mid-, and high-resolution levels of the decoder, serving complementary roles: the low-resolution PAE enforces global tonal consistency across the recovered field; the mid-resolution PAE smooths colour gradients within coherent flow regions; and the high-resolution PAE preserves fine boundary definition at jet edges and shear layers.

The discriminator D is a fully convolutional PatchGAN-style network whose successive strided convolution blocks produce a spatial grid of authenticity scores, enabling evaluation of physical plausibility at multiple receptive-field scales simultaneously. Spectral normalisation is applied to every convolutional layer to enforce the 1-Lipschitz condition required by the Wasserstein formulation and to prevent discriminator weight explosion during training. A single PAE module is inserted immediately after the discriminator's first convolutional block, before any spatial downsampling. This placement is deliberate: operating at full input resolution, the physically-aware discriminator becomes sensitive to gradient-continuity violations at thermodynamic boundaries that purely data-driven PatchGAN discriminators would otherwise overlook, providing gradient feedback that explicitly penalises non-physical reconstructions throughout training.

The network is optimised under a Wasserstein adversarial loss with gradient penalty (WGAN-GP), with the discriminator updated five times per generator step to ensure stable and informative gradient feedback. The generator loss combines the Wasserstein adversarial term with a VGG perceptual loss evaluated at the relu3_3 feature layer, which provides a balance between low-level spatial detail preservation and semantic structural consistency. An adaptive weight decay strategy progressively reduces the perceptual loss contribution as training matures, shifting emphasis from structural fidelity in the early training phase toward adversarial sharpening in the later stages.

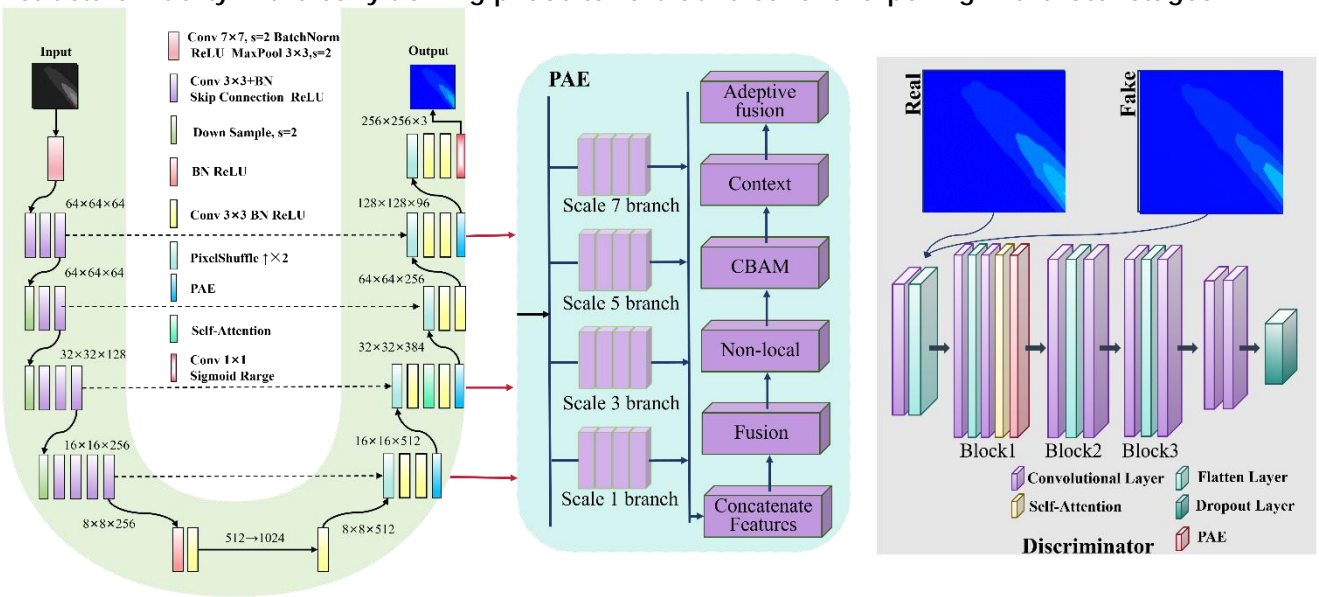


Figure 2 – Architecture of the proposed UNet-GAN framework for flow field parameter inversion.

2.2 PAE Module

The PAE module receives an input feature tensor and transforms it through four sequential stages: multi-scale feature construction, global dependency modelling, dual-attention refinement, and adaptive feature fusion, as illustrated in Figure 3.

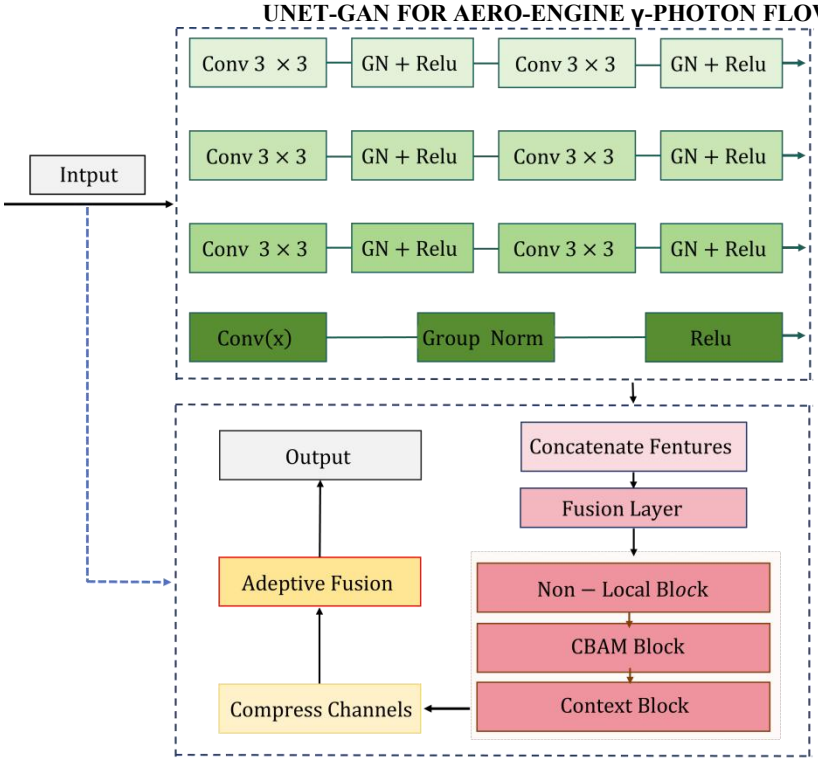


Figure 3 – Structure of the PAE module.

Upon entering the PAE module, feature maps are first processed by a multi-scale dilated convolutional fusion network comprising four parallel branches. Branch1 applies a 1×1 pointwise convolution for fast dimensionality reduction and inter-channel mixing. Branch2 stacks two consecutive 3×3 convolutions, achieving a 5×5 effective receptive field suited to medium-scale nozzle geometry. Branch3 employs a 3×3 convolution with dilation rate r=2, achieving a 5×5 effective footprint with sparser sampling and reduced parameter count. Branch 4 uses a 3×3 convolution with dilation rate r=3, extending the effective receptive field to 7×7 to capture large-scale jet cross-sectional structure. The feature maps from the four branches are concatenated along the channel dimension to form a multi-scale feature representation, as shown in formula (1).

$$F_{\text{multi}} = \left[F_1; F_3; F_5; F_7\right] \tag{1}$$

The concatenated feature map then undergoes feature fusion and smoothing through a 3×3 convolution followed by GroupNorm normalisation and ReLU activation, resulting in the fused feature F_{fused}, as shown in formula (2).

$$F_{\text{fused}} = ReLU(GN(\text{Conv3}\times 3(F_{\text{multi}}))) \tag{2}$$

Subsequently, the fused features are fed into a NonLocalBlock to establish long-range dependencies between any two spatial positions in the feature map. Unlike local convolutions, the NonLocalBlock captures global information in a single step by computing an all-pairs similarity matrix and performing attention-weighted aggregation. Three 1×1 projection convolutions θ, ϕ, and g map the input to query, key, and value embeddings respectively. The output at spatial position i aggregates information from all positions j weighted by their normalised query–key inner product, with a residual connection preserving the original feature information, and can be expressed using formula (3).

$$Z_i = \alpha \cdot W\left(\sum_{\forall j}\frac{\exp(\theta(F_i')^T\phi(F_j'))}{\sum_{\forall k}\exp(\theta(F_i')^T\phi(F_k'))}g(F_j')\right) + F_i' \tag{3}$$

This mechanism ensures that the temperature distribution at one side of the nozzle cross-section remains globally aware of the distribution at the other side, preventing the cross-region colour drift that arises when the generator lacks long-range context.

The feature map is subsequently processed by a Convolutional Block Attention Module (CBAM) [15], which employs a dual attention pathway comprising channel attention followed by spatial

attention to accurately pinpoint and enhance key feature regions. The overall CBAM structure is depicted in Figure 4.

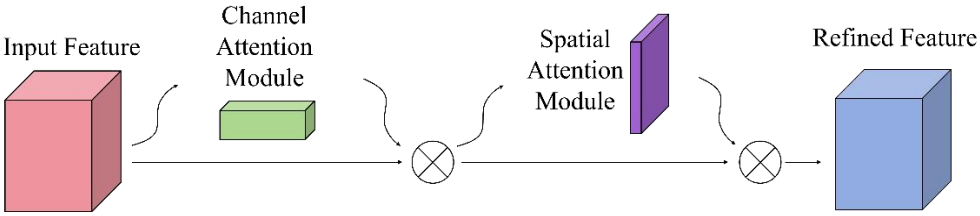


Figure 4 – Structure of the CBAM Attention Mechanism.

The channel attention mechanism first extracts two channel descriptors via global average pooling and global max pooling, as shown in formula (4) and formula (5) respectively.

$$F_{\text{avg}} = \frac{1}{H \times W} \times \sum_{i=1}^{H} \sum_{j=1}^{W} X_c(i,j) \tag{4}$$

$$F_{\max} = \max_{i,j} X_c(i,j) \tag{5}$$

These two descriptors are then processed through a shared two-layer MLP with intermediate channel count C/r where r=16, and combined with sigmoid gating to produce the channel attention weights M_C, as shown in formula (6).

$$M_c(X) = \sigma\left(\text{MLP}(F_{\text{avg}}) + \text{MLP}(F_{\max})\right) \tag{6}$$

The channel attention structure is illustrated in Figure 5.

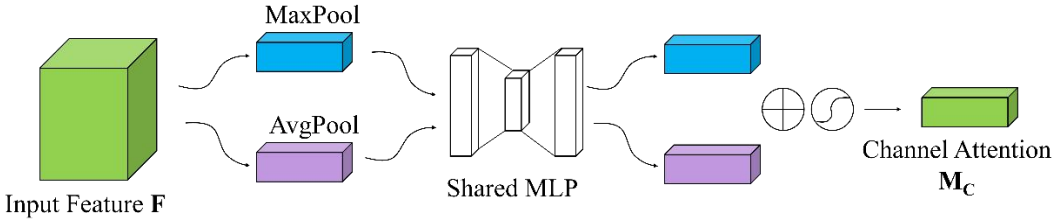


Figure 5 – Structure of the Channel Attention Submodule.

The spatial attention mechanism then operates on the channel-refined feature map.Average pooling and max pooling are applied along the channel dimension to produce two spatial descriptors, as shown in formula (7) and formula (8).

$$F^s_{\text{avg}} = \frac{1}{C} \times \sum_{c=1}^{C} X_c(i,j) \tag{7}$$

$$F^s_{\max} = \max_c X_c(i,j) \tag{8}$$

These descriptors are concatenated and passed through a 7×7 convolution with sigmoid gating to produce the spatial attention map M_s, as shown in formula (9).

$$M_s(X) = \sigma\left(f^{7\times7}\left(\left[F^s_{\text{avg}}; F^s_{\max}\right]\right)\right) \tag{9}$$

The spatial attention structure is illustrated in Figure 6.

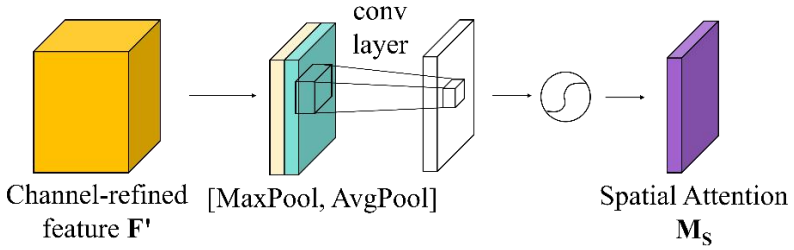


Figure 6 – Structure of the Spatial Attention Submodule.

The complete dual-attention computation is articulated in formula (10) and formula (11).

$$X' = M_c(X) \otimes X \tag{10}$$

$$X'' = M_s(X') \otimes X' \tag{11}$$

By first recalibrating which feature channels are most informative and then identifying where in the spatial map they are most relevant, CBAM concentrates representational capacity on the physically active regions such as jet boundaries, temperature cores, and turbulent wake structures where accurate parameter inversion is most critical.

After attention weighting, the feature map undergoes further refinement within a ContextBlock, which fuses global contextual statistics captured through global average pooling with local spatial detail preserved through local convolutions, generating adaptive blending weights for each spatial position. The resulting high-dimensional features are then channelled into a two-stage 1×1 channel-compression module to eliminate redundant dimensions introduced by the multi-branch concatenation. Finally, an AdaptiveFusion gate combines the enhanced features with the original PAE input through learned channel-wise and spatial-wise weights, with a fixed residual coefficient on the original input branch.

3. Experiments

3.1 Simulation Data Generation Pipeline

Acquiring paired (γ-photon image, ground-truth flow-parameter field) data from a live operating engine is currently infeasible at the volumes and precision required for deep learning. A two-stage high-fidelity simulation pipeline is therefore constructed. In Stage 1, ANSYS Fluent computational fluid dynamics simulations of the SW60B single-shaft turbojet engine nozzle are performed across six throttle settings spanning 30% to 100% rated thrust, covering the full operational envelope of the engine. At each setting, post-processing extracts cross-sectional distributions of temperature T (K), static pressure P (Pa), and velocity magnitude V (m/s) at multiple axial stations of the nozzle exit plane. After normalisation and quality screening for numerical convergence, 42 base flow-field maps y are retained as ground-truth targets.

In Stage 2, each thermodynamic map is converted into a three-dimensional volumetric phantom representing the spatial distribution of the β^+-emitting radionuclide tracer, and this phantom is input to the GATE (Geant4 Application for Tomographic Emission) Monte Carlo simulation platform [16], as illustrated in Figure 7. GATE models the complete physical chain: radionuclide decay statistics, positron range, 511 keV photon production, photon transport through the engine metal structure, and coincidence detection by a ring of 52 detector modules with 312 crystals per ring, a detection ring diameter of 180 mm, energy window 350–625 keV, and time resolution 1.5 ns. After simulating 10^7 coincidence events per phantom, OSEM reconstruction (20 subsets, 4 iterations) recovers 128×128 pixel grayscale positron-flux images. A super-resolution noise-suppression module [17] then sharpens spatial resolution from 128×128 to 256×256 and reduces statistical granularity without introducing artificial smoothing, yielding the final network input x.

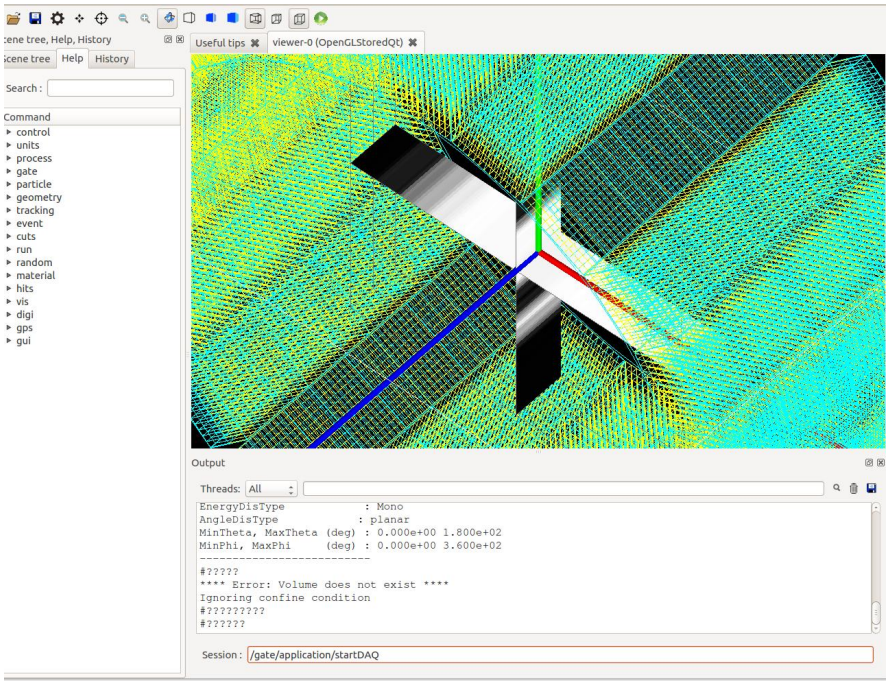


Figure 7 – Configuration of the GATE simulation platform.

3.2 Dataset Construction

The original dataset comprised 42 base paired samples derived from the CFD simulation pipeline described above. Given the limited scale of flow-field simulation data, training deep generative models directly on such a small corpus would risk severe overfitting. A multi-level data augmentation strategy was therefore designed to bridge the gap between idealised CFD simulations and the stochastic characteristics of real γ-photon flow-field imaging conditions.

Basic spatial transforms including random cropping, horizontal and vertical flipping, and small-angle rotation within ±5° were applied to simulate variability in detector geometry and cross-sectional sampling plane positioning across different monitoring setups. Photometric and noise augmentations including random brightness perturbation with factor 0.8–1.2, contrast adjustment with factor 0.7–1.3, Gaussian blur with $\sigma \in [0.5, 1.5]$ px, and additive Gaussian noise with $\sigma \in [0, 0.05]$ were incorporated to replicate the statistical and resolution variability introduced by the γ-photon imaging process under different tracer activity levels and acquisition conditions, including Poisson noise contributions inherent in low-count coincidence event acquisition. Large-scale datasets including COCO and ImageNet were additionally incorporated to increase sample diversity and improve generalisation capability.

After augmentation, the dataset expanded to 3,368 paired images, which were split into 2,422 training images and 946 validation images, approximately a 7:3 ratio. Representative examples from the training set are shown in Figure 8, and examples from the validation set are shown in Figure 9. All images were resized to 256×256 pixels, stored in PNG format to preserve quality, and normalised using ImageNet channel statistics to standardise the inputs. The dataset used single-channel grayscale γ-photon images as inputs and three-channel RGB thermodynamic parameter maps as outputs.

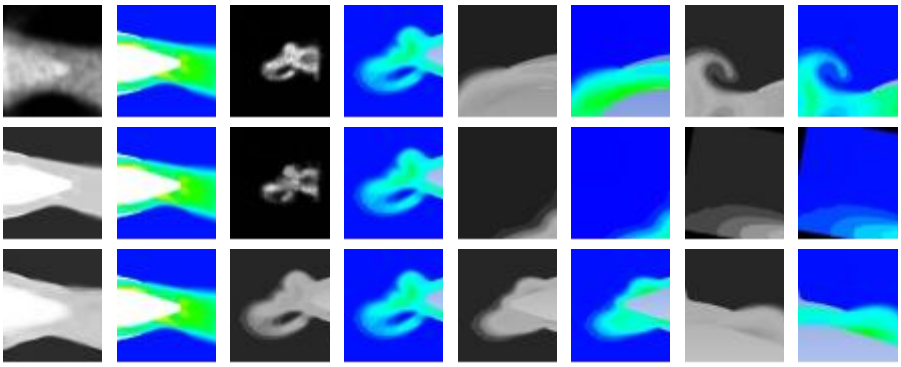


Figure 8 – Selected training set images.

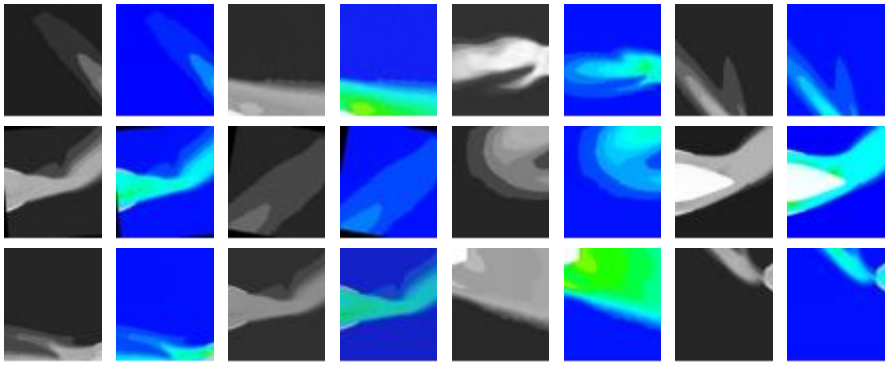


Figure 9 – Selected validation set images.

3.3 Training Setup

Training proceeds in two consecutive phases. In the pre-training phase spanning 25 epochs, the generator is optimised alone under the perceptual loss L_perc with no discriminator involved. This phase establishes a well-conditioned initial mapping from the photon-flux input domain to the thermodynamic parameter-field output domain, substantially reducing the risk of mode collapse

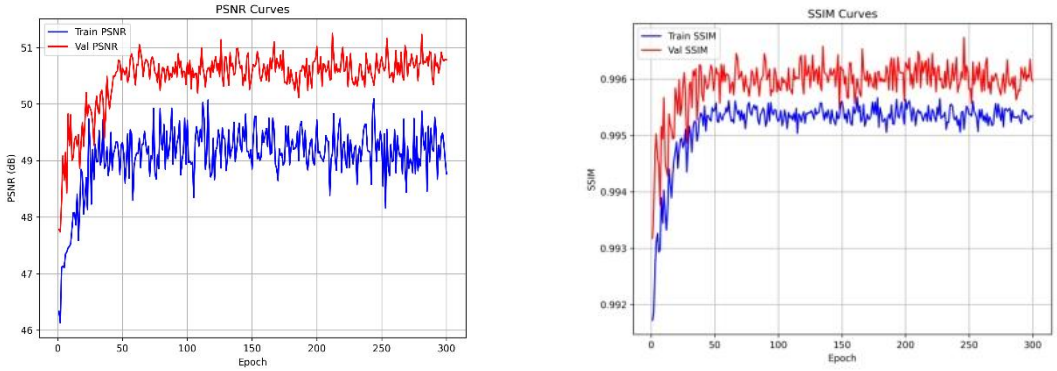
and gradient instability that can occur when an untrained generator faces a well-trained discriminator in the early adversarial phase.

In the adversarial training phase spanning 100 epochs, generator and discriminator are trained jointly under the composite losses defined in Section 2.1, with the discriminator updated five times per generator step. The Adam optimiser with $\beta_1 = 0$ and $\beta_2 = 0.9$ is used for both networks. The learning rate follows a cosine annealing schedule decaying from an initial value of 5×10^{-5} to a final value of 5×10^{-6} over the 100 adversarial epochs. Batch size is set to 8. The model checkpoint yielding the highest validation PSNR is saved as the final model; this optimal checkpoint is reached at iteration 212 of the adversarial phase. All experiments are conducted on a single NVIDIA RTX 4090 GPU with 24 GB GDDR6X memory.

4. Results And Discussion

4.1 Algorithm Model Training

The evolution of PSNR and SSIM over the 100 adversarial training epochs, evaluated separately on the training and validation partitions, is shown in Figure 10. Both metrics increase steadily through the early adversarial epochs and plateau near the optimal checkpoint at iteration 212, after which only marginal improvement is observed. The validation PSNR and SSIM curves track their training counterparts closely throughout training—the gap between training and validation PSNR remains below 0.8 dB across all epochs—confirming that the model generalises to unseen flow configurations rather than memorising the training set. The plateau behaviour and the close train–validation alignment together indicate stable GAN convergence without mode collapse, consistent with the stabilising effect of the WGAN-GP gradient penalty. The optimal model checkpoint at iteration 212 is used for all subsequent evaluation.



(a) PSNR curves (b) SSIM curve

Figure 10 – Training convergence analysis: (a) PSNR curves; (b) SSIM curve.

4.2 Qualitative and Quantitative Results

Figure 11 presents a side-by-side qualitative comparison for the direct-jet flow configuration, which is characterised by a high-temperature central jet core surrounded by a monotonically decreasing temperature field extending to the cooler outer boundary. DeOldify exhibits visible hue discontinuities at the jet-core boundary: the probabilistic imaging process produces an abrupt positron-flux transition at the jet edge, and the DeOldify decoder, constrained by a limited receptive field that cannot distinguish between a physical boundary and a statistical noise artefact, maps nearby pixels to discontinuous hue values, producing stair-like colour banding at the high-to-medium temperature interface. DDColor generates a more saturated output but introduces erroneous colour assignment in the wake region, assigning warm hues to the cool outer boundary in violation of the physically expected monotonic temperature decay from the nozzle centreline. The proposed method recovers smooth, physically coherent colour gradients throughout the recovered field: the transition from the high-temperature core hue to the low-temperature peripheral hue is spatially gradual and continuous, consistent with the CFD-predicted temperature decay, and the jet boundary is rendered with sharp spatial definition without gradient discontinuities.

These qualitative improvements directly reflect the role of the PAE module, as the NonLocalBlock connects physically coupled cross-sectional regions to prevent hue inconsistencies, while CBAM concentrates representational capacity on the jet boundaries and temperature cores where accurate inversion is most critical.

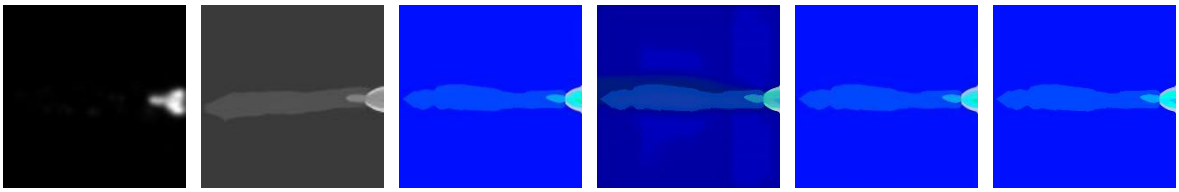


Figure 11 – Colorization results for the direct-jet flow configuration.

Table 1 reports the mean PSNR and SSIM metrics over the complete 946-image validation set at the optimal checkpoint. The proposed method achieves PSNR of 33.68 dB and SSIM of 0.86, outperforming both baselines across both metrics. DDColor performs substantially worse than DeOldify on both PSNR and SSIM despite its more complex dual-decoder architecture, confirming that architectural complexity alone does not guarantee improved performance in the structured thermodynamic inversion domain when domain-specific physical inductive biases are absent.

Table 1 – Quantitative comparison of colorization results.

Method	PSNR↑	SSIM↑
DeOldify	31.16	0.83
DDColor	14.04	0.62
Our Method	33.86	0.86

5. Conclusion

This paper has addressed the problem of parameter inversion for aero-engine γ-photon flow field images, recovering continuous multi-channel thermodynamic parameter fields from noisy single-channel positron-flux observations. A conditional GAN framework has been proposed in which a ResNet101 U-Net generator maps γ-photon images to three-channel flow-parameter fields under Wasserstein adversarial training with gradient penalty. The central contribution is the Physical Aware Enhancement (PAE) module, deployed symmetrically in both the generator and the discriminator. The PAE module integrates three complementary mechanisms: four-branch multi-scale dilated convolution that captures thermodynamic features across a hierarchy of spatial scales within a single processing block; a NonLocalBlock that establishes long-range spatial dependencies in a single feed-forward pass, ensuring global thermodynamic coherence across the recovered field; and CBAM dual-attention that concentrates representational capacity on the physically active regions such as jet boundaries, temperature cores, and turbulent wake structures where accurate inversion is most critical. An adaptive-weight training schedule utilising the relu3_3 VGG feature layer progressively shifts emphasis from structural fidelity to adversarial sharpening as training matures. Experimental evaluation on the direct-jet flow configuration of the SW60B engine demonstrates that the proposed method achieves PSNR of 33.86 dB and SSIM of 0.86, outperforming DeOldify and DDColor baselines retrained on the same data, validating the effectiveness of the physical consistency enforcement strategy. Future work will extend the two-dimensional cross-sectional inversion framework to full three-dimensional volumetric parameter field recovery, exploiting the complete spatial information captured by multi-ring PET detector arrays.

6. Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (General

Program) (G. No. 62571239), Shenzhen Natural Science Foundation (General Program) (G. No. JCYJ20250604190734042), the Joint Fund of the National Natural Science Foundation of China (G. No. U2441285), and the Experimental Technology Research and Development Project of Nanjing University of Aeronautics and Astronautics (G. No.SYJS202604Z).

7. Contact Author Email Address

The contact author email address is: xiaohui@nuaa.edu.cn

8. 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] Hampel U, et al. High resolution gamma ray tomography scanner for flow measurement and non-destructive testing applications. *Rev. Sci. Instrum.*, Vol. 78, 103704, 2007.
[2] Bruggemann J, Gross A and Pate S. Non-intrusive visualization of optically inaccessible flow fields utilizing positron emission tomography. *Aerospace*, Vol. 7, No. 52, 2020.
[3] Tropea C, Yarin A L and Foss J F (Eds). *Springer Handbook of Experimental Fluid Mechanics*. Springer, Berlin, 2007.
[4] Ongie G, Jalal A, Metzler C A, Baraniuk R G, Dimakis A G and Willett R. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, Vol. 1, No. 1, pp 39–56, 2020.
[5] Adler J and Oktem O. Solving ill-posed inverse problems using iterative deep neural networks. *Inverse Problems*, Vol. 33, No. 12, 124007, 2017.
[6] McCann M T, Jin K H and Unser M. Convolutional neural networks for inverse problems in imaging: a review. *IEEE Signal Processing Magazine*, Vol. 34, No. 6, pp 85–95, 2017.
[7] Goodfellow I, et al. Generative adversarial nets. *Proc Advances in Neural Information Processing Systems*, pp 2672-2680, 2014.
[8] Isola P, Zhu J-Y, Zhou T and Efros A A. Image-to-image translation with conditional adversarial networks. *Proc IEEE/CVF CVPR*, pp 1125–1134, 2017.
[9] Wang T-C, Liu M-Y, Zhu J-Y, Tao A, Kautz J and Catanzaro B. High-resolution image synthesis and semantic manipulation with conditional GANs. *Proc IEEE/CVF CVPR*, pp 8798–8807, 2018.
[10] Vitoria P, Raad L and Ballester C. ChromaGAN: adversarial picture colorization with semantic class distribution. *Proc IEEE/CVF WACV*, pp 2434–2443, 2020.
[11] Lim B, Son S, Kim H, Nah S and Lee K M. Enhanced deep residual networks for single image super-resolution. *Proc IEEE/CVF CVPR Workshops*, pp 136–144, 2017.
[12] Ronneberger O, Fischer P and Brox T. U-net: Convolutional networks for biomedical image segmentation. *Proc Medical Image Computing and Computer-Assisted Intervention*, pp 234-241, 2015.
[13] He K, Zhang X, Ren S and Sun J. Deep residual learning for image recognition. *Proc IEEE/CVF CVPR*, pp 770–778, 2016.
[14] Shi W, Caballero J, Huszar F, Totz J, Aitken A P, Bishop R, Rueckert D and Wang Z. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. *Proc IEEE/CVF CVPR*, pp 1874–1883, 2016.
[15] Woo S, Park J, Lee J Y, et al. Cbam: Convolutional block attention module. *Proc European Conference on Computer Vision*, pp 3-19, 2018.
[16] Jan S, Santin G, Strul D, Staelens S, Assie K, Autret D, Avner S, Barbier R, Bardies M and Bloomfield P M. GATE: a simulation toolkit for PET and SPECT. *Physics in Medicine and Biology*, Vol. 49, No. 19, pp 4543–4561, 2004.
[17] Hui X, Quan L, Xu Y, et al. Research on a noise-suppression super-resolution enhancement module for positron flow field images. *Scientific Reports*, Vol. 15, No. 1, 21443, 2025.