photonics

MDPI

Article

A Digital Twin-Inspired Correction Method for Infrared Detectors

Jiangyu Tian [1], Libing Jin [1,2] and Jun Chang [1,*]

[1] School of Optics and Photonics, Beijing Institute of Technology, Beijing 100081, China; 3120225305@bit.edu.cn (J.T.); 3220235080@bit.edu.cn (L.J.)
[2] Beijing Institute of Space Mechanics & Electricity, Beijing 100086, China
* Correspondence: optics_chang@126.com

Abstract

Infrared focal plane arrays (IRFPAs) often suffer from spatiotemporal nonuniformity that persists after conventional two-point nonuniformity correction (NUC), especially under temperature drift and time-varying readout conditions. These residuals are typically structured, including column-group striping caused by shared column-end circuits and row-wise baseline/common-mode drift induced by row-scanning paths. We propose a structured, digital-twin-inspired detector-side refinement of two-point NUC that augments the bias term with interpretable low-dimensional components: a static column bias vector capturing group-correlated residuals and a row-related structured term consisting of a static row baseline and a frame-synchronous common-mode component with row-dependent sensitivity, while keeping the two-point gain/offset backbone unchanged. Rather than representing a full system-level digital twin of the infrared payload, the proposed framework serves as a detector-side virtual representation of dominant readout-induced structured residual states that can be estimated and updated from calibration data. Experiments on blackbody calibration data across multiple temperature points demonstrate that the column-related structured component significantly reduces group-wise column residuals, the row-related structured component suppresses time-varying row striping, and the combined method improves both column- and row-direction metrics consistently across temperatures.

Keywords: infrared focal plane array (IRFPA); non-uniformity correction (NUC); two-point calibration

1. Introduction

Nonuniformity in IRFPAs mainly arises from pixel-to-pixel variations in gain and offset, together with fixed-pattern noise (FPN) introduced by the readout circuitry. In uniform-field imaging, these effects appear as stationary spatial textures whose characteristics drift with ambient temperature and detector operating conditions. This problem is particularly important in platforms such as unmanned aerial vehicles and satellite payloads, where long-duration operation and environmental variation are common. Conventional two-point NUC, implemented using blackbody or flat-field references, estimates gain and offset independently for each pixel. However, it does not explicitly model system-level structured residuals induced by shared row-driving and column-readout circuits. As a result, residual column striping, row-wise baseline offsets, and low-frequency structured drift often remain after correction.

Received: 10 February 2026
Revised: 6 April 2026
Accepted: 20 April 2026
Published: 21 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

A broad range of NUC methods has been developed for IRFPA imaging. Among them, calibration-based approaches, especially conventional two-point NUC, remain widely used because they are compatible with real-time implementation and hardware deployment [1,2]. These methods estimate pixel-wise gain and offset parameters from two flat-field reference images acquired at known radiometric levels. They are effective for compensating intrinsic pixel-to-pixel response variation, but they are built on an implicit pixel-wise independence assumption. Consequently, they do not account for structured residuals induced by shared circuitry, such as column–channel grouping and row–driver coupling. Under temperature drift or time-varying readout conditions, column–group striping and row-wise baseline drift therefore often persist after correction [3–5]. To improve flexibility, some studies have introduced multi-point calibration or higher-order polynomial fitting [6,7]. However, these approaches still rely on static parameter models and do not explicitly describe how structured residuals evolve under changing temperature conditions.

Another major research direction is scene-based or adaptive NUC. Representative methods update correction parameters online from temporal image statistics [8–10], or impose spatiotemporal consistency through temporal filtering, total-variation regularization, and related constraints [11–13]. These approaches provide a degree of adaptivity without dedicated calibration and can suppress some low-frequency nonuniform artifacts in dynamic scenes. However, their effectiveness often depends on assumptions about scene motion, temporal variation, or mean-drift behavior. In many cases, the correction terms are estimated in the image domain rather than in the original NUC parameter space, which makes them difficult to integrate consistently with calibration-based correction and reduces interpretability for deployment. Related image-domain post-processing methods, including wavelet filtering, principal component analysis, and frequency-domain stripe suppression, have also been used to reduce structured column noise [14–16]. Although such methods may improve image appearance, they do not directly model the physical source of the residuals and are therefore less suitable for embedded IRFPA applications that require rapid and accurate target energy measurement, parameter consistency, and stable real-time operation.

Deep learning and data-driven methods that have emerged in recent years attempt to formulate infrared NUC as an end-to-end mapping problem, leveraging convolutional networks, Transformer architectures, or generative adversarial networks to model and restore images [17,18]. These methods can improve image quality in some cases, but they usually require large training datasets and do not explicitly represent physical factors such as temperature drift or readout-chain structure. In addition, their limited interpretability makes parameter deployment and cross-device generalization more challenging. Against this background, a gap remains between physically interpretable calibration-based correction and adaptive suppression of structured residuals. In particular, there is still a need for a method that introduces column-related periodic residuals and row-wise baseline drift into the NUC framework in a parametric and deployment-consistent manner.

A method is therefore needed that introduces dominant structured residuals into the NUC parameter space in a physically interpretable and deployment-consistent manner. Here, “digital twin” does not denote a full-system-level digital twin of the entire infrared payload. Instead, it denotes a detector-side virtual representation of dominant readout-induced structured residual states associated with shared readout structures. The proposed method is therefore more appropriately described as a digital-twin-inspired detector-side structured correction framework, in which a small number of virtual state variables are linked to physically interpretable residual modes and are estimated and updated from calibration data.

The proposed framework is most directly applicable to IRFPAs whose residuals are dominated by grouped column-end sharing effects and row-scanning-related correlated drift. For detectors with different readout architectures, the same modeling principle can still be used, but the structural parameterization, basis representation, and common-mode formulation should be adapted to the specific hardware topology. Compared with representative scene-based temporal correction methods, the present method operates directly in the NUC parameter space and explicitly targets readout-structure-induced residuals with clear physical interpretation.

2. Materials and Methods

Conventional IRFPA nonuniformity correction is commonly described by a linear two-point model. For pixel (x, y) in frame t acquired at operating condition temperature T, the raw detector response is written as

$$Y(x, y, t; T) = G(x, y; T)\, S(x, y, t; T) + O(x, y; T) + N(x, y, t; T). \tag{1}$$

Here, $Y(x, y, t; T)$ is the detector output, $S(x, y, t; T)$ is the equivalent ideal signal, $G(x, y; T)$ and $O(x, y; T)$ are the pixel gain and offset, respectively, and $N(x, y, t; T)$ collects random noise and unmodeled effects.

Let S_a and S_b denote the two calibration radiometric levels, and let $Y_a(x, y)$ and $Y_b(x, y)$ denote the corresponding mean calibration frames. The conventional two-point estimates are then

$$\hat{G}_{2pt}(x, y) = \frac{Y_b(x, y) - Y_a(x, y)}{S_b - S_a}, \tag{2}$$

$$\hat{O}_{2pt}(x, y) = Y_a(x, y) - \hat{G}_{2pt}(x, y) S_a. \tag{3}$$

The corresponding corrected output is:

$$\hat{S}_{2pt}(x, y, t) = \frac{Y(x, y, t) - \hat{O}_{2pt}(x, y)}{\hat{G}_{2pt}(x, y)}. \tag{4}$$

The above formulation implicitly assumes pixel-wise independence, meaning that the gain and offset of each detector element are estimated independently. However, in large-format IRFPAs, the readout chain often contains shared structures, such as grouped column–channel amplification and reused row-driving, reset, or sampling paths. As a result, system-level residuals that are correlated with the hardware architecture may remain even after conventional correction, including grouped column steps, column striping, and row-wise baseline offsets. These residuals are not purely random. Instead, they persist as structured fixed patterns and drift gradually with temperature. Therefore, pixel-wise two-point parameters alone are insufficient to describe such shared-structure-induced residuals, especially in cross-temperature operating conditions.

Consider M uniform-field frames acquired at temperature point T_i. Because the scene is spatially uniform, the desired corrected signal can be approximated by a spatially constant scalar $S_0(T_i)$. After conventional two-point correction, the corrected output can therefore be written as

$$\hat{S}_{2pt}(x, y, t; T_i) = S_0(T_i) + R(x, y, t; T_i), \tag{5}$$

where $R(x, y, t; T_i)$ denotes the residual nonuniformity together with random perturbations.

Because the structured residual is relatively stable over frames while the random component is approximately zero-mean, a steady residual map is estimated by temporal averaging after removing the per-frame global level:

$$\overline{R}(x, y; T_i) = \frac{1}{M}\sum_{t=1}^{M}(\hat{S}_{2pt}(x, y, t; T_i) - \mu(T_i)), \quad (6)$$

Here, $\mu(T_i)$ denotes the frame mean or the region-of-interest mean of the T_i corrected frame at temperature point T_i. This subtraction removes the global constant component and retains the spatially structured residual.

As shown in Figure 1, significant structured residuals persist after conventional two-point NUC. This indicates that the remaining error cannot be fully described by independent pixel-wise gain and offset mismatches alone, but is instead related to correlated structures in the readout chain. Since conventional two-point NUC estimates gain and offset independently for each pixel, it is effective for correcting static pixel-level nonuniformity but lacks explicit terms to capture group-correlated column residuals and row-correlated baseline drift. These observations motivate the introduction of additional structured bias components into the parameter model.

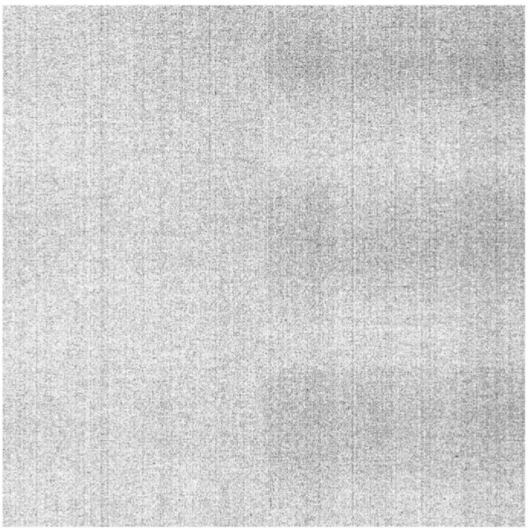


Figure 1. Noise of the detector to a uniform background after conventional two-point NUC.

Figure 2 shows the column-mean profile of the corrected residual and reveals clear group-wise inconsistency along the column direction, with step-like changes or repeating textures appearing approximately every 256 columns. This pattern is consistent with grouped sharing in the column-end readout architecture. As illustrated in Figure 3, large-format readout circuits often divide the columns into fixed groups that share column amplifiers, sample-and-hold stages, correlated double-sampling stages, multiplexing resources, and part of the reference or analog-to-digital conversion chain in order to satisfy bandwidth, power, and area constraints. Because different column groups exhibit different device mismatch, reference-distribution drop, bias-settling behavior, and temperature-drift coefficients, each group introduces a different common bias component. After conventional two-point NUC, this effect appears as systematic inter-group column residuals, which motivates the introduction of a static column bias term.

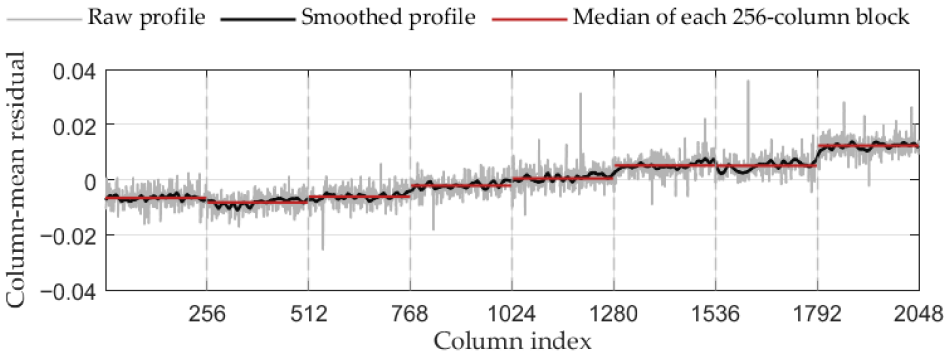


Figure 2. Column-mean profile of the residual after conventional two-point NUC.

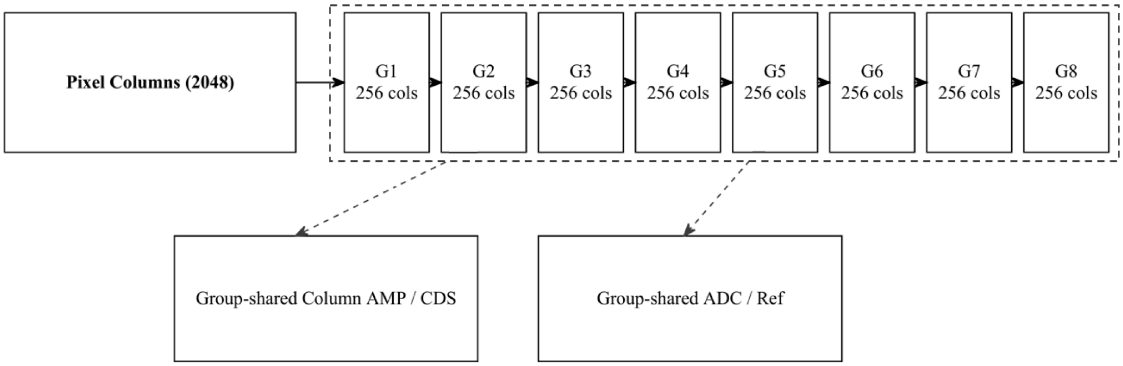


Figure 3. Schematic of grouped column-end readout sharing.

Figure 4 shows that the corrected residual exhibits pronounced baseline inconsistency along the row direction, meaning that different rows have different mean residual levels. Such row-correlated structure is consistent with the row-scanning and row-level readout process illustrated in Figure 5. During sequential row activation, the row driver and decoder, row-selection switches, reset and clamping operations, row-buffer stages, and reference-settling procedures are affected by row-line parasitic, charge injection, residual reset noise, settling-time differences, and row-buffer mismatch. As a result, the floor level or baseline bias varies from row to row. This observation motivates the introduction of a static row baseline term in the structured bias model.

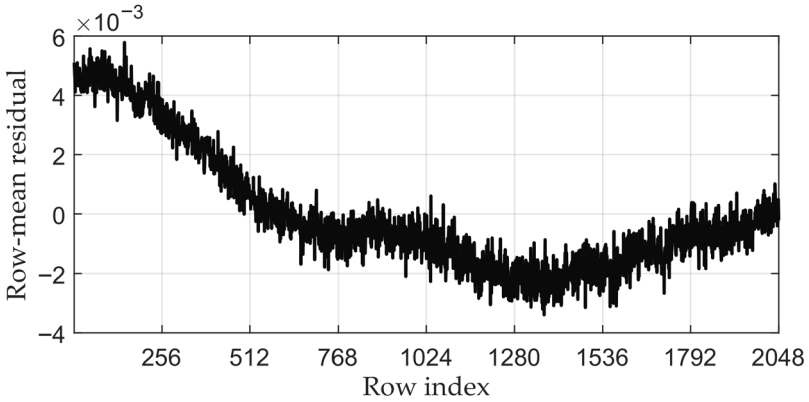


Figure 4. Row-mean profile of the residual after conventional two-point NUC.

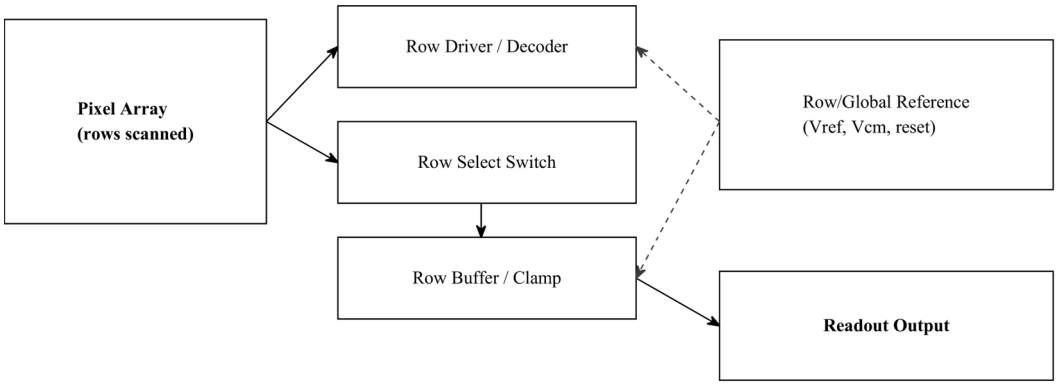


Figure 5. Schematic of row-scanning readout.

Figure 6 shows that the row-mean values vary coherently across the frame sequence rather than remaining constant in time. This behavior indicates that the row residual contains not only a static baseline term but also a time-varying common-mode component that evolves synchronously from frame to frame. This observation motivates the introduction of a row-related structured term in the model.

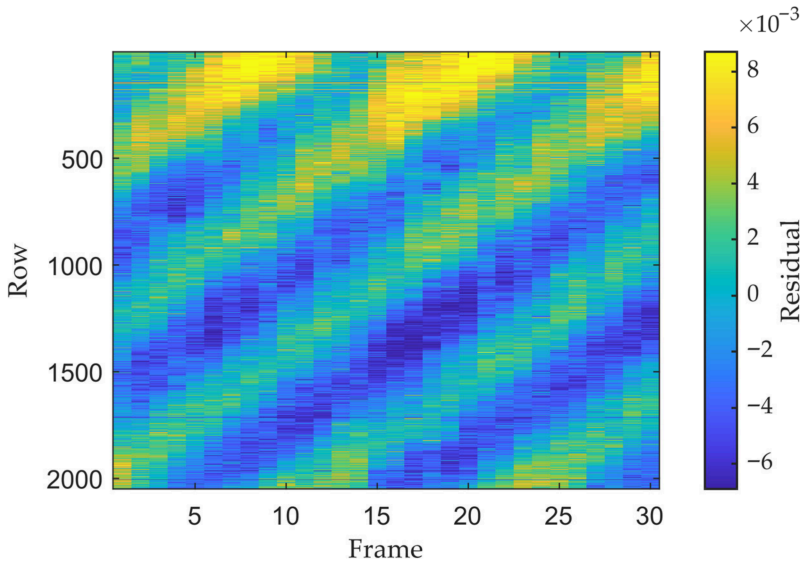


Figure 6. Evolution of the row-mean residual over 30 frames after conventional two-point correction.

Figures 2–6 indicate that the dominant residual after conventional two-point NUC is not purely pixel-wise random noise but is mainly composed of two structured modes: (i) a quasi-static column-related component associated with shared column-end circuitry, and (ii) a row-related structured component consisting of a static row baseline plus a frame-synchronous common-mode component. Accordingly, we keep the conventional two-point estimates $\hat{G}_{2pt}$ and $\hat{O}_{2pt}$ unchanged and augment only the offset branch with three additional low-dimensional terms: a static column bias $b_c(x)$, a static row baseline $b_r(y)$, and a dynamic row term $k_r(y)\ c(t)$. This decomposition preserves the original two-point backbone while explicitly modeling the dominant structured residuals observed in the calibration data.

These additional components are incorporated directly into the NUC parameter model so that the correction remains consistent with the readout architecture. Let the region of interest have size $H \times W$. To represent column-correlated fixed-pattern residuals caused by shared column-end resources, such as column amplifiers or analog-to-digital conversion

channels, we define a static column bias vector $b_c(x)$ and expand it to the two-dimensional detector grid $B_c(x, y)$,

$$B_c(x, y) = b_c(x). \tag{7}$$

To model the row-related residual, we decompose it into a static row baseline and a frame-synchronous temporal component. Specifically, let $b_r(y)$ denote the time-invariant row baseline, let $c(t)$ denote the frame-synchronous common-mode scalar, and let $k_r(y)$ denote the row-dependent sensitivity to that common mode. The resulting row-related structured term is written as

$$R_{row}(x, y, t) = b_r(y) + k_r(y)\, c(t). \tag{8}$$

In this formulation, $b_r(y)$ captures the static row-to-row baseline difference, $c(t)$ represents the frame-dependent common disturbance shared across rows, and $k_r(y)$ describes how strongly each row couples to that disturbance. The column term $b_c(x)$ is independent of time and represents the static column-related residual associated with the shared column-end circuitry. To make the decomposition identifiable, $b_r(y)$ is defined as the time-invariant part of the row residual, whereas $c(t)$ describes only the temporal fluctuation around that baseline.

With the original two-point calibration estimates $\hat{G}_{2pt}(x, y)$ and $\hat{O}_{2pt}(x, y)$ kept fixed, the proposed method modifies only the offset branch. The effective time-dependent offset is defined as

$$O_{eff}(x, y, t) = \hat{O}_{2pt}(x, y) + b_c(x) + b_r(y) + k_r(y)\, c(t). \tag{9}$$

The corresponding structured two-point correction is then

$$\hat{S}_{prop}(x, y, t) = \frac{Y(x, y, t) - O_{eff}(x, y, t)}{\hat{G}_{2pt}(x, y)}. \tag{10}$$

This formulation keeps the original two-point gain term unchanged while explicitly compensating the dominant column-related and row-related residual modes in the offset branch.

The key idea of this model is that column-related and row-related structured residuals are not treated as image-domain post-processing targets. Instead, they are introduced explicitly into the NUC parameter space and participate directly in the correction formula, thereby preserving physical interpretability and deployment consistency. Unlike multi-point calibration methods with explicit temperature-function fitting, the present formulation does not impose a direct temperature-dependent bias function $b(\cdot; T)$. Rather, it captures the dominant structured residuals across different temperatures and frame sequences through a combination of a static column term and a row-related structured term, while preserving the main parameter backbone and implementation simplicity of conventional two-point NUC.

The parameter estimation is performed in a sequential manner. First, the static column term $b_c(x)$ is estimated from the steady residual map. Second, after subtracting the column term, the static row baseline $b_r(y)$ is estimated from the temporally averaged row residual. Third, the row-sensitivity vector $k_r(y)$ is treated as a fixed calibration quantity. Finally, for each frame, the frame-synchronous common-mode scalar $\hat{c}(t)$ is estimated from the column- and baseline-corrected row residual by a sensitivity-weighted least-squares projection:

$$\hat{c}(t) = \frac{\sum_{y=1}^{H} k_r(y) \frac{1}{W} \sum_{x=1}^{W} R(x, y, t)}{\sum_{y=1}^{H} k_r(y)^2 + \varepsilon}, \tag{11}$$

where ε is a small positive constant for numerical stability. This projection isolates the frame-synchronous row-mode amplitude while reducing leakage from the static column and row components.

On this basis, we define two energy-ratio metrics to quantify the relative contributions of the column-related structured component and the row-related structured component at temperature point T_i:

$$\eta_c(T_i) = \frac{\| R_c(\cdot,\cdot) \|_2^2}{\| \overline{R}(\cdot,\cdot;T_i) \|_2^2}, \eta_{rt}(T_i) = \frac{\| \overline{R}_{rt}^{*}(\cdot,\cdot;T_i) \|_2^2}{\| \overline{R}(\cdot,\cdot;T_i) \|_2^2}, \tag{12}$$

where $\overline{R}(\cdot,\cdot;T_i)$ denotes the steady residual map at temperature point T_i, $R_c(\cdot,\cdot)$ denotes the column-related structured component, and

$$\overline{R}_{rt}^{*}(x,y;T_i) = b_r(y) + k_r(y)\,\overline{c}(T_i), \tag{13}$$

with

$$\overline{c}(T_i) = \frac{1}{M_i}\sum_{t=1}^{M_i} \hat{c}(t). \tag{14}$$

where M_i is the number of frames acquired at temperature point T_i. Here, $\overline{R}_{rt}^{*}(x,y;T_i)$ represents the row-related structured component constructed using the average common-mode level at that temperature point.

3. Results

The experimental data were obtained from a blackbody-based uniform-field calibration sequence. At each temperature point, 30 raw frames were acquired. The temperature points were selected according to equal-step spacing in radiant exitance derived from the blackbody radiation law. To prevent defective pixels from biasing the statistics, all frames were preprocessed by bad-pixel correction before quantitative evaluation. Defective pixels and outliers were detected using a robust thresholding rule and replaced by the corresponding frame mean.

For two-point calibration, the third and tenth temperature points were selected in order to avoid possible nonlinearity or saturation near the ends of the temperature range. The detailed partition of temperature points is provided in Supplementary Table S1. The gain and offset parameters of the two-point NUC model were estimated from the mean images of the 30 frames acquired at these two calibration points. In the figures, the horizontal axis is displayed using a mapped temperature obtained from equal-step radiant-exitance reparameterization. This mapping is used only for visualization and alignment and does not alter the original temperature grouping or the acquisition sequence.

Figure 7 shows that the column-related structured component accounts for a substantial fraction of the post-correction residual across multiple temperature points. This result indicates that the dominant column-wise residual after conventional two-point NUC is largely associated with bias introduced by the shared column-end readout chain. The row-related structured component also contributes consistently across temperatures, indicating that time-varying row drift is a systematic residual mode rather than an incidental pixel-level fluctuation.

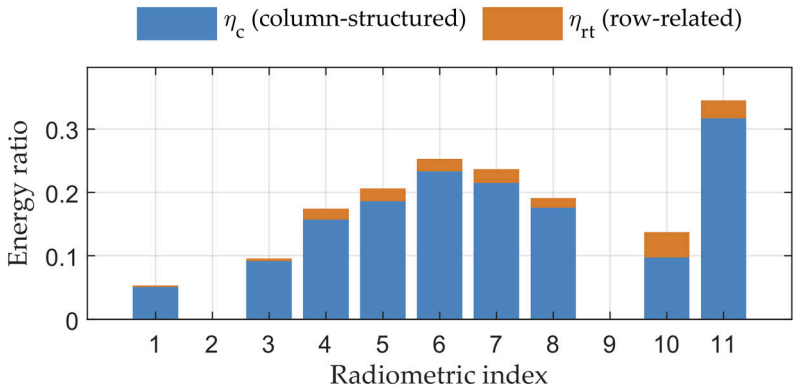


Figure 7. Ratio of structured residual components vs. radiometric level.

To further verify the existence of time-varying row-direction residual drift and evaluate its effect on correction, two-dimensional heatmaps are presented at the validation temperature points to show the evolution of the row-mean residual with frame index. The per-frame row-mean residual is defined as

$$\bar{r}(y,t) = \frac{1}{W}\sum_{x=1}^{W}\left(\hat{S}_{prop}(x,y,t) - \mu(T_i)\right). \tag{15}$$

Under Baseline, that is, conventional two-point correction only, the row-mean residual exhibits clear stripe-like temporal structures. After the row-related structured correction is introduced, these structures are substantially weakened, and the overall temporal fluctuation amplitude of the row-mean residual is significantly reduced, as shown in Figures 8 and 9. In addition, the estimated frame-synchronous common-mode scalar exhibits stable periodic oscillations under Baseline, whereas its amplitude is effectively suppressed after correction. This result indicates that the proposed row model captures and removes the dominant temporal drift using only a very small number of additional parameters.

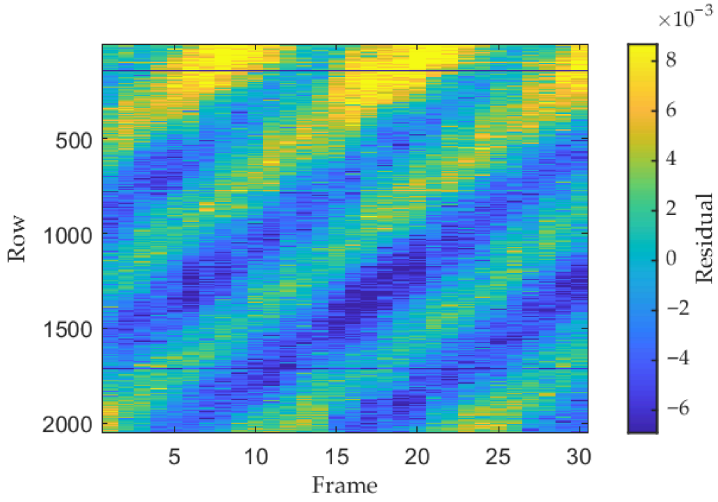


Figure 8. Evolution of the row-mean residual over 30 frames under Baseline.

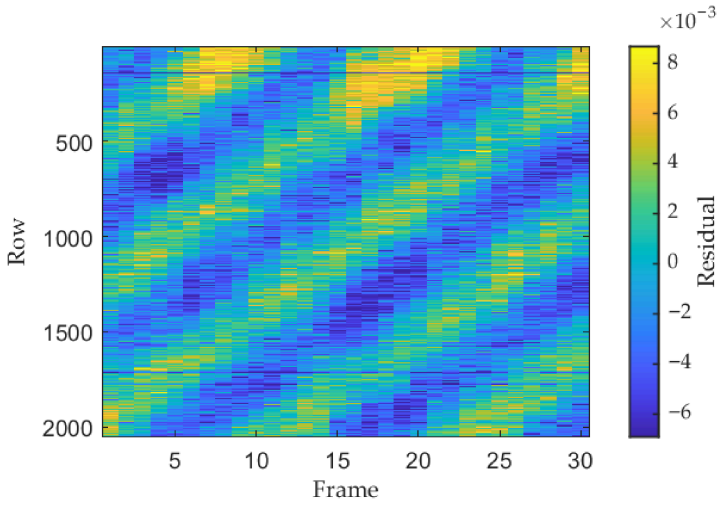


Figure 9. Evolution of the row-mean residual over 30 frames after the row-related structured correction.

Figures 10 and 11 compare the spectra of the frame-synchronous common-mode scalar $\hat{c}(t)$ before and after the row-related structured correction, showing that the high-frequency components and the overall fluctuation amplitude of $\hat{c}(t)$ are effectively suppressed after correction.

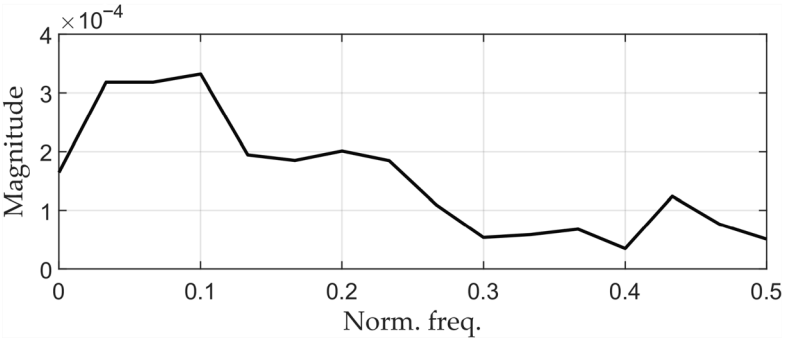


Figure 10. The magnitude of $\hat{c}(t)$ under Baseline.

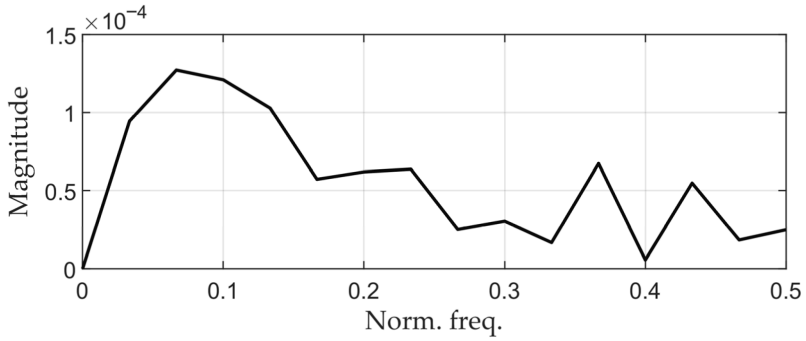


Figure 11. The magnitude of $\hat{c}(t)$ after the row-related structured correction.

To further evaluate the effect of the static column bias term, Figures 12 and 13 compare the column-mean residual bands before and after column correction at a representative validation temperature point. Under the Baseline, the residual band exhibits clear group-wise oscillation and within-group fluctuation, indicating that the dominant column-direction residual remains strongly correlated with the grouped column-end readout structure. After introducing the static column bias term, the periodic banding is substantially weakened

and the inter-group step variation is reduced. This result shows that the proposed column model effectively absorbs the dominant structured residual associated with shared column-end resources.

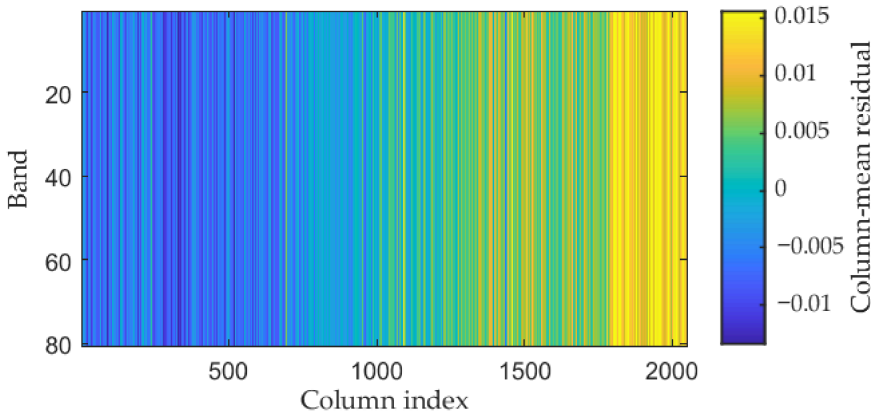


Figure 12. Column-mean residual band under Baseline.

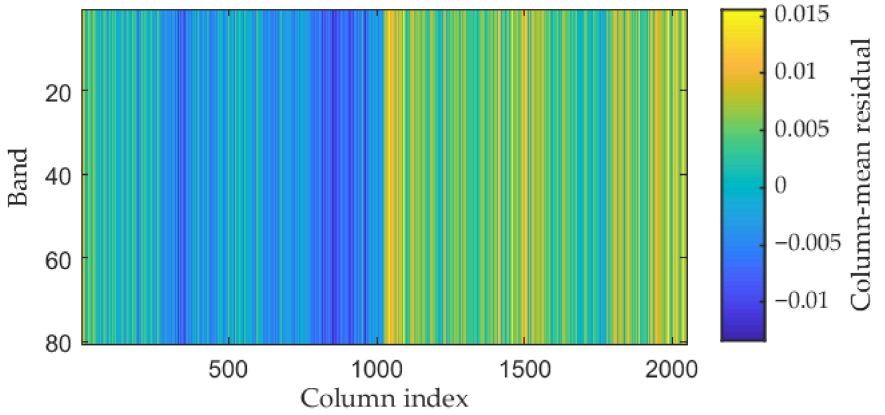


Figure 13. Column-mean residual band after applying static column bias correction.

To isolate the role of each structural term, four correction settings were compared: Baseline, corresponding to conventional two-point NUC; column bias correction only; row-related structured correction only; and the combined correction. At the validation temperature points, the Baseline column-mean curve exhibits pronounced within-group fluctuations. After the static column bias term is introduced, the overall oscillation of the column-mean curve is significantly reduced, showing that this term effectively absorbs the dominant column-wise structured residual induced by the shared column-end readout chain. By contrast, the row-related structured term has only a limited effect on the column-mean curve, which is consistent with the expectation that the row model has weak projection onto column statistics. For the row-mean curve, the Baseline result shows evident row-wise baseline offsets and low-frequency disturbances induced by temporal drift. After introducing the row-related structured term, the standard-deviation metric of the row-mean curve decreases significantly, and the temporal stripe patterns in the heatmap are suppressed. Finally, the combined correction suppresses both column-wise structured residuals and row-wise temporal drift, yielding more consistent performance across validation temperatures.

The static column bias term mainly improves the column-direction metrics, whereas the row-related structured term mainly improves the row-direction metrics. The combined

correction yields consistent improvements in both directions. The curve-based evaluation metrics used in this study are defined as follows:

$$\overline{R}_c^*(x) = \frac{1}{H}\sum_{y=1}^{H}\overline{R}(x,y), \overline{R}_r^*(y) = \frac{1}{W}\sum_{x=1}^{W}\overline{R}(x,y), \tag{16}$$

$$\sigma_c = \mathrm{std}\left(\overline{R}_c^*\right), \Delta_c = \max\left(\overline{R}_c^*\right) - \min\left(\overline{R}_c^*\right), \tag{17}$$

$$\sigma_r = \mathrm{std}\left(\overline{R}_r^*\right), \Delta_r = \max\left(\overline{R}_r^*\right) - \min\left(\overline{R}_r^*\right). \tag{18}$$

The standard-deviation-based metrics σ_c and σ_r are more robust to global fluctuations, whereas the spike-sensitive metrics Δ_c and Δ_r respond more strongly to a small number of abnormal column or row outliers. Slight inconsistencies may therefore appear at individual validation points. Nevertheless, both the standard-deviation metrics and the time-varying heatmaps consistently show that the row-related structured term suppresses systematic drift in a clear and reproducible manner.

Quantitative results are summarized in Table 1. At both validation temperatures, the combined correction method consistently reduces the residual nonuniformity in both the column and row directions. At 39.4 °C, the column-direction metric decreases from 0.0054 to 0.0047 and the row-direction metric decreases from 0.0021 to 0.0013, corresponding to improvements of 13.2% and 37.3%, respectively. At 50.7 °C, the column-direction metric decreases from 0.0073 to 0.0060 and the row-direction metric decreases from 0.0021 to 0.0012, corresponding to improvements of 17.2% and 40.8%, respectively. Overall, the results show that the proposed method remains effective across different temperature points and is particularly advantageous in suppressing row-direction residual fluctuations.

Table 1. Comparison of column/row residual standard deviations and improvement rates for different correction strategies at the validation temperature points.

T (°C)	Base σ_c	Base σ_r	Col σ_c	Row σ_r	Col_Row σ_c	Col_Row σ_r	Improve σ_c	Improve σ_r
39.4	0.0054	0.0021	0.0047	0.0013	0.0047	0.0013	13.2%	37.3%
50.7	0.0073	0.0021	0.0060	0.0012	0.0060	0.0012	17.2%	40.8%

4. Discussion

The proposed framework is most directly applicable to IRFPAs whose residuals are dominated by grouped column-end sharing effects and row-scanning-related correlated drift. For detectors with different readout architectures, the same modeling principle can still be used, but the structural parameterization, basis representation, and common-mode formulation should be adapted to the specific hardware topology and dominant residual modes.

The experimental results show that conventional two-point NUC still leaves significant structured residuals under uniform-field conditions. These residuals are manifested mainly as group-periodic patterns along the column direction and baseline inconsistency along the row direction. When interpreted in the context of the readout-chain architecture, these phenomena are consistent with system-level errors introduced by shared readout paths. At the column end, sharing the analog front-end, reference distribution, and conversion resources across fixed column groups leads to group-correlated bias drift. Along the row direction, row scanning and the reuse of row buffering, clamping, reset, and sampling paths introduce row-by-row floor-level differences. Under frame-synchronous disturbances, such as reference-level fluctuation, supply ripple, or timing coupling, these row-related offsets further exhibit a time-varying common-mode component.

Conventional two-point NUC estimates gain and offset independently for each pixel and therefore compensates pixel-level static nonuniformity only. It does not explicitly absorb the correlated bias terms generated by shared readout structures and is therefore prone to residual striping and low-frequency baseline drift under temperature drift or operating-point variation. This interpretation is consistent with the established view that column- and row-correlated fixed-pattern noise originates from readout-structure sharing. It also implies that purely image-domain destriping methods, although they may improve image appearance, cannot be integrated consistently into the NUC parameter space and are less suitable for embedded real-time deployment.

Compared with representative scene-based temporal correction methods, the proposed method has a different objective and operating regime. Scene-based methods usually estimate correction terms from scene motion or temporal image statistics and are therefore strongly influenced by assumptions about scene content and temporal behavior. By contrast, the present method introduces structural residual terms directly into the NUC parameter space and uses calibration-consistent low-dimensional variables that are explicitly tied to the readout architecture. Here, "digital twin" does not denote a full system-level digital twin of the entire infrared payload. Instead, it denotes a detector-side virtual representation of dominant readout-induced structured residual states. This formulation improves interpretability and deployment consistency because the added parameters retain direct physical correspondence to shared readout structures. Beyond correction performance, the proposed method is also intended to improve the radiometric accuracy of high-precision infrared detectors, providing a physically interpretable and calibration-consistent detector-level output. This makes the corrected signal more reliable as a preliminary input for downstream photonics-related processing, such as precise radiometric analysis, optical system characterization, and other tasks that depend on accurate detector-level calibration.

The proposed method is also deployment-friendly because it does not replace the existing two-point NUC pipeline. Instead, it augments the offset term with only a small number of structural parameters. For an array of size $H \times W$, the additional model parameters scale linearly as $W + 2H$, rather than requiring full-frame iterative optimization or large end-to-end models. During online operation, only a lightweight common-mode statistic and the corresponding low-dimensional compensation need to be updated. The framework is therefore computationally efficient, memory-efficient, and scalable to larger-format infrared detectors. To further evaluate deployability, Table 2 compares the proposed method with representative NUC strategies in terms of parameter count, online computation, and residual nonuniformity from the experimental data. The comparison shows that the proposed method preserves the original two-point NUC backbone, adds only a small number of low-dimensional parameters, and achieves low extra online cost while maintaining residual performance close to that of more complex approaches. All residual nonuniformity values in Table 2 were obtained by evaluating each method on the same blackbody calibration dataset used throughout this study. For the four-point fitting NUC baseline, the four calibration temperature points were selected as those that yielded the lowest residual nonuniformity among all available temperature points, so as to provide a favorable comparison for this method. The proposed method nonetheless achieves lower residual nonuniformity than the four-point fitting baseline while requiring far fewer additional parameters and negligible extra online computation. Under an ideal fully pipelined FPGA implementation with sufficient resources, the three methods can all reach one output pixel per clock. Therefore, the more meaningful comparison is resource-normalized throughput. In this sense, conventional two-point NUC and the proposed method require one per-pixel multiplication, whereas a typical four-point polynomial fitting implementation requires approximately three per-pixel multiplications. Hence, under the same

DSP budget, the four-point fitting method would support about one-third of the parallel pixel lanes and approximately three times the frame time. For a 2048 × 2048 array at 200 MHz, if 12 lanes are available for two-point NUC, the estimated frame time is about 1.75 ms/frame for both two-point NUC and the proposed method, versus about 5.24 ms/frame for four-point fitting.

Table 2. Comparison of parameter count, online computation, and residual nonuniformity for different NUC methods on the 2048 × 2048 IRFPA.

Method	Parameters	Online Computation	Residual Nonuniformity
Two-Point NUC	$2 \times 2048 \times 2048$	Per pixel: 1 add + 1 multiply	0.0362
Four-Point Fitting NUC	$\geq 4 \times 2048 \times 2048$	Per pixel: LUT/polynomial + MACs	0.0022
Proposed Method	$2 \times 2048 \times 2048 + 6144$	Two-point NUC + per frame: 1 add + 1 multiply	0.0011

Although the proposed framework has clear advantages in interpretability and deployability, several limitations and extensions deserve further study. First, higher-order within-group column residual shapes may still appear under some conditions. Such behavior may be associated with higher-level reference-settling effects, interleaved sampling, odd-even column phase differences, or multi-channel interleaved readout. Future work may therefore introduce intra-group basis expansions, such as low-order orthogonal bases or piecewise polynomials, while preserving low-dimensional and deployable parameterization.

Second, the present study performs structured fine-tuning mainly on the bias term in order to remain consistent with two-point NUC deployment. For some detectors, however, temperature dependence in the gain term may also be non-negligible. An important extension would therefore be to incorporate temperature- or operating-point-driven gain fine-tuning under conditions of stability and identifiability, and to study its coupling with the bias-related structural terms.

Finally, the current validation is based primarily on uniform-field calibration sequences. Future studies should incorporate a wider range of operating-condition variations and dynamic-scene data in order to evaluate the generalization ability of the digital-twin-inspired framework across multiple operating domains. Such studies may also support more robust abnormal row and column detection and more reliable online fine-tuning strategies.

5. Conclusions

This paper addresses the problem of spatiotemporal structured residual noise that persists after two-point NUC in IRFPAs under temperature drift and operating-point variations, and proposes a digital-twin-inspired detector-side structured correction framework for NUC refinement. The proposed approach preserves the calibration and deployment form of the conventional two-point NUC backbone while introducing low-dimensional bias components consistent with the shared readout architecture. Along the column direction, a static column bias vector captures group-correlated residuals introduced by the shared column-end chain. Along the row direction, a row-related structured term, consisting of a static row baseline and a frame-synchronous common-mode component, captures row-wise baseline differences and time-varying correlated drift. Here, “digital twin” does not denote a full system-level digital twin of the entire infrared payload. Instead, it denotes a detector-side virtual representation of dominant readout-induced structured residual states that can be estimated and updated from calibration data. Experimental results show

that the dominant post-NUC residual is not purely random noise, but is largely governed by group-correlated column bias and row-wise common-mode drift associated with shared readout structures. The proposed framework is particularly suitable for IRFPAs exhibiting grouped column-end sharing and row-scanning-related correlated residuals. For other detector architectures, the same modeling principle can still be used, although the structural parameterization should be adapted to the specific readout topology. Overall, the method is interpretable, deployment-friendly, and scalable, with low online computational cost, making it suitable for real-time correction in larger-format infrared detectors.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/photonics13040396/s1, Table S1: Partition of temperature points for training, validation, and two-point calibration.

Author Contributions: J.T. and L.J. contributed equally to this work. Conceptualization, J.T. and J.C.; methodology, J.T. and L.J.; software, L.J.; validation, L.J. and J.C.; formal analysis, L.J.; investigation, L.J.; resources, J.C.; data curation, L.J.; writing—original draft preparation, J.T.; writing—review and editing, J.T. and J.C.; visualization, L.J.; supervision, J.C.; project administration, J.C.; funding acquisition, J.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by National Key R&D Program of China (2021YFC2202100).

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. (The data are not publicly available due to privacy or ethical restrictions.).

Conflicts of Interest: The author declares no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

IRFPA	Infrared focal plane array
FPN	Fixed pattern noise
NUC	Non-uniformity correction
Two-Point NUC	Two-point non-uniformity correction
CDS	Correlated double sampling
ROI	Region of interest
ADC	Analog-to-digital converter
UAV	Unmanned aerial vehicle

References

1. Orżanowski, T. Nonuniformity correction algorithm with efficient pixel offset estimation for infrared focal plane arrays. *SpringerPlus* **2016**, *5*, 1831. [CrossRef] [PubMed]
2. Wang, L.; Zhou, J.; Liao, S.; Wang, Y.; Li, S.; Xia, Y. A real-time two-point non-uniformity correction method for infrared images based on prior knowledge. *IEEE Photonics J.* **2026**, *18*, 7800209. [CrossRef]
3. Sui, J.; Jin, W.-Q.; Dong, L.-Q.; Wang, X. A new nonuniformity correction algorithm for infrared line scanners. *Proc. SPIE* **2006**, *6207*, 341–348.
4. Kudinov, I.; Kholopov, I. Scene-based non-uniformity fixed pattern noise correction algorithm for infrared video sequences. *CEUR Workshop Proc.* **2020**, *2665*, 135–140.
5. Li, M.; Wang, Y.; Sun, H. Single-frame infrared image non-uniformity correction based on wavelet domain noise separation. *Sensors* **2023**, *23*, 8424. [CrossRef] [PubMed]
6. Wang, R.; Chen, P.; Tsien, P. An improved nonuniformity correction algorithm for infrared focal plane arrays which is easy to implement. *Infrared Phys. Technol.* **1998**, *39*, 15–21. [CrossRef]
7. Holst, S.; Holst, L. Two-point and multi-point nonuniformity correction of infrared focal plane arrays. *Opt. Eng.* **1999**, *38*, 2071–2078.

8. Hardie, R.C.; Hayat, M.M.; Armstrong, E.; Yasuda, B. Scene-based nonuniformity correction with video sequences and registration. *Appl. Opt.* **2000**, *39*, 1241–1250. [CrossRef] [PubMed]
9. Boutemedjet, A.; Deng, C.; Zhao, B. Robust Approach for nonuniformity correction in infrared focal plane array. *Sensors* **2016**, *16*, 1890. [CrossRef] [PubMed]
10. Lv, B.; Tong, S.; Liu, Q.; Sun, H. Statistical scene-based non-uniformity correction method with interframe registration. *Sensors* **2019**, *19*, 5395. [CrossRef] [PubMed]
11. Li, J.; Qin, H.; Yan, X.; Zeng, Q.; Yang, T. Temporal-spatial nonlinear filtering for infrared focal plane array stripe nonuniformity correction. *Symmetry* **2019**, *11*, 673. [CrossRef]
12. Barral, A.; Arias, P.; Davy, A. Fixed pattern noise removal for multi-view single-sensor infrared camera. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Waikoloa, HI, USA, 3–8 January 2024; pp. 1669–1678.
13. Godoy, S.E.; Pezoa, J.E.; Torres, S.N. Noise-cancellation-based nonuniformity correction algorithm for infrared focal-plane arrays. *Appl. Opt.* **2008**, *47*, 5394–5399. [CrossRef] [PubMed]
14. Wang, E.; Jiang, P.; Hou, X.; Zhu, Y.; Peng, L. Infrared stripe correction algorithm based on wavelet analysis and gradient equalization. *Appl. Sci.* **2019**, *9*, 1993. [CrossRef]
15. Lu, D.; Teng, L.; Ren, J.; Tan, J.; Wang, M.; Wang, L.; Gu, G. Scene-based nonuniformity correction method using principal component analysis for infrared focal plane arrays. *Appl. Sci.* **2023**, *13*, 13331. [CrossRef]
16. Zeng, Q.; Qin, H.; Yan, X.; Yang, T. Fourier domain anomaly detection and spectral fusion for stripe noise removal of TIR imagery. *Remote Sens.* **2020**, *12*, 3714. [CrossRef]
17. Jian, X.; Lv, C.; Wang, R. Nonuniformity correction of single infrared images based on deep filter neural network. *Symmetry* **2018**, *10*, 612. [CrossRef]
18. Mou, X.; Zhu, T.; Zhou, X. Visible-image-assisted nonuniformity correction of infrared images using the GAN with SEBlock. *Sensors* **2023**, *23*, 3282. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.