# Simultaneous temperature estimation and nonuniformity correction from multiple frames

Navot Oz, Omri Berman, Nir Sochen, David Mendelovich, and Iftach Klapp

arXiv:2307.12297v1 [cs.CV] 23 Jul 2023

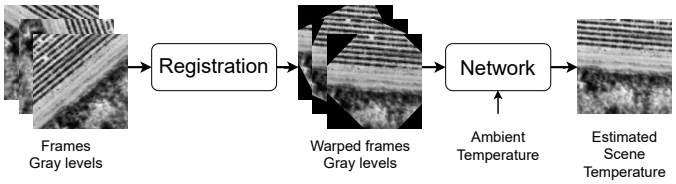


Fig. 1. Estimating the scene temperature from a burst of gray level frames.

***Abstract*—Infrared (IR) cameras are widely used for temperature measurements in various applications, including agriculture, medicine, and security. Low-cost IR camera have an immense potential to replace expansive radiometric cameras in these applications, however low-cost microbolometer-based IR cameras are prone to spatially-variant nonuniformity and to drift in temperature measurements, which limits their usability in practical scenarios.**

**To address these limitations, we propose a novel approach for simultaneous temperature estimation and nonuniformity correction from multiple frames captured by low-cost microbolometer-based IR cameras. We leverage the physical image acquisition model of the camera and incorporate it into a deep learning architecture called kernel estimation networks (KPN), which enables us to combine multiple frames despite imperfect registration between them. We also propose a novel offset block that incorporates the ambient temperature into the model and enables us to estimate the offset of the camera, which is a key factor in temperature estimation.**

**Our findings demonstrate that the number of frames has a significant impact on the accuracy of temperature estimation and nonuniformity correction. Moreover, our approach achieves a significant improvement in performance compared to vanilla KPN, thanks to the offset block. The method was tested on real data collected by a low-cost IR camera mounted on a UAV, showing only a small average error of $0.27^{\circ}C - 0.54^{\circ}C$ relative to costly scientific-grade radiometric cameras.**

**Our method provides an accurate and efficient solution for simultaneous temperature estimation and nonuniformity correction, which has important implications for a wide range of practical applications.**

***Index Terms*—Deep learning, Fixed-Pattern Noise (FPN), Multiframes, Temperature estimation, Nonuniformity correction, Infrared cameras, Microbolometer, space variant nonuniformity**

Navot Oz, Omri Berman and Iftach Klapp are with the Department of Sensing, Information and Mechanization engineering, Agricultural Research Organization, Volcani Institute, P.O. Box 15159, Rishon LeZion 7505101, Israel.

Navot Oz, Omri Berman and David Mendelovich are with the School of Electrical Engineering, Tel Aviv University, Tel Aviv 69978, Israel.

Nir Sochen is with the Department of Mathematics, Tel Aviv University, Tel Aviv 69978, Israel.

This work was supported by the Israeli Ministry of Agriculture's Kandel Program (grant number 20-12-0018).

## I. Introduction

TEMPERATURES is an important indicator for the state of an object. For example, the temperature of a plant is important in deducing information on its well-being [1], [2].

Long-wave infrared (LWIR) imaging is a technique that measures the thermal radiation emitted from an object, commonly known as infra-red (IR) imaging. To avoid noise and improve accuracy, radiometric IR cameras are usually cooled to $200_K$ and below. The cooling apparatus, as well as complex shuttering and control systems, increase the cost of the camera considerably. Although IR imaging is a well-established technique, the high cost of IR cameras prohibits widespread usage.

An alternative approach to radiometric thermal imaging involves the use of low-cost uncooled microbolometer arrays, which can facilitate the creation of inexpensive IR cameras with low energy requirements. Unlike photon-counting detector arrays, microbolometer arrays gauge alterations in electrical resistance resulted from the radiation emission of an object [3]. Each microbolometer in the array is heated by the thermal radiation to a temperature that is reliant on the scene, resulting in each microbolometer having a marginally different temperature based on the observed scene and the incident angle of the radiation. The incident radiation causes a miniscule change in the resistance of the microbolometer. The temperature of the scene is reflected by the variation in resistance of each microbolometer. The infinitesimal changes in resistance detected by each microbolometer are used to create an image that corresponds to the temperature of the observed scene.

Although microbolometer arrays are a useful tool for thermal imaging, they have significant limitations. Space-variant nonuniformity and noise from various sources affect the accuracy of these arrays. The nonuniformity drifts due to the change in ambient temperature, which causes unpredictable errors in the sensor readings.

The lack of a cold shield in the uncooled camera is a prominent cause of nonuniformity [4]. This self-radiation effect is attributable to the housing and lens of the camera, which emit thermal radiation onto the sensor. The self-radiation varies according to the ambient temperature of the camera.

Fixed-pattern noise (FPN) is an additional factor that contributes to nonuniformity in microbolometer arrays. The readout circuitry of these arrays is typically line-based, like charge coupled devices. Even minor differences between line-readers on the same array can result in significant variation between lines in the resulting image [5].

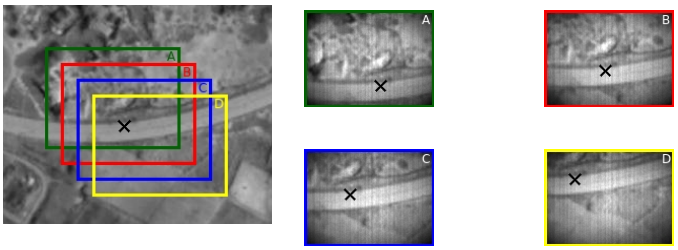


Fig. 2. Simulation of consecutive frames taken during a drone flight from an IR camera. The frames sampled on the left image are marked by colored rectangles. The effects of the spatially variant non uniformity is seen in the frames on the right. The cross on the road appears in different location throughout the frames and is affected differently by the spatially variant nonuniformity.

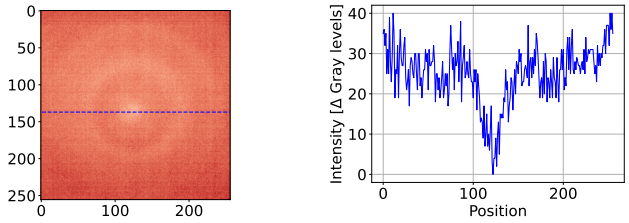


Fig. 3. Example of the nonuniformity in low-cost IR cameras. On the left is an image of a $30^\circ$ blackbody with ambient temperature of $44.7^\circ$, and on the right are the intensities along the blue dashed line.

Noises in the camera has the effect of increasing noise equivalent differential temperature (NEDT), which refers to the minimum detectable change in scene temperature [5]. The NEDT is a measure of the sensitivity of the camera. The higher the NEDT, the less sensitive the camera is to changes in temperature.

An image of a uniform heat source (blackbody) is shown in Fig. 3. The spatially-variant nonuniformity is demonstrated by the radial patterns in the gray levels of the left subfigure. The subfigure on the right plots the gray levels along the blue dashed line, showing the impact of the nonuniformity and the noise on the gray levels.

A widely used application of IR imaging is remote sensing. Remote sensing is the process of acquiring information about an object without making physical contact with it. The information is acquired by measuring the reflected or emitted radiation from the object. The information is then used to deduce the physical properties of the object. Remote sensing is used in a variety of fields (e.g, agriculture, geology, and meteorology).

One common use-case for IR camera is to be mounted on drones. This setup results in high overlap between frames (Sec. V-A). The redundant information can be used to simultaneously improve the accuracy of the temperature estimation and correct nonuniformity in the frames. Fig. 2 illustrates how redundant information between frames is beneficial. The object is affected differently by the nonuniformity at each frame, which means that the true underlying temperature of the object can be extracted.

The aim of this study is twofold: exploiting the redundancy in data and the physical model of the camera to develop a method of estimating scene temperatures using a low-cost IR camera based on microbolometers, and to correct for nonuniformity in the frames.

## II. Related work

Estimating the temperature can be broadly divided into two parts - transforming the output of the camera to temperatures, and correcting nonuniformity in the sensor. Finding a transformation between camera output to temperature is called *thermal calibration*. Correcting the nonuniformity in the sensor is called *nonuniformity correction* (NUC).

### A. Thermal calibration

The raw output of the IR camera is dependent on the object temperature, and the output values themselves are given in gray levels. For example, the gray levels dynamic range in the `Tau2` is $14_{bit}$. The classic approach is to calibrate the camera for different ambient temperatures [6].

A large dataset that has pairs of object and ambient temperatures must be collected for calibration. The gain and offset are calculated from the data per-pixel to attune for the spatially-variant nonuniformity. Thus, the calibration process usually requires considerable time and resources.

Schulz et al. [6] used a single-point correction. Meaning that a single ambient temperature is used, a constant gain is assumed and only the offset is found. Riou et al. [5] suggested a two-point correction that requires two ambient temperatures, but solved for both gain and offset, and it is widely used across industrial IR cameras today. Both methods used a linear regression to extract the gain and offset coefficients. Nugent et al. [7] modeled the gain and offset as polynomial in the temperature of the object and used least-squares to extract the coefficients. Contemporary works adds prior knowledge into the calibration process. Liang et al. [8] found the gain and offset for a given temperature and interpolated the results for other ambient temperatures, Chang and Li [9] incorporated the integration time of each frame as prior knowledge to the calibration.

The calibration data must be collected for each camera separately, because each camera is slightly different due to the manufacturing process. This requires scientific-grade equipment, making the calibration process infeasible for most users.

### B. Nonuniformity correction

As said in Sec. I, the frames of the IR camera suffer from spatially variant nonuniformity. The nonuniformity can be corrected for a single frame, or by combining information from multiple frames (known as scene-based).

*1) Single frame:* A given image contains information that can be exploited for different tasks. Information such as low frequencies [10], recurring patches in the image [11] or the statistical distribution of patches in the image [12]. Some works used a single image to correct the nonuniformity.

Scribner et al. [13] used a NN to find the offset and gain by alternating optimization and gradient descent. Tendero and

Gilles [14] used histogram equalization across the columns in a frame, and then applied a discrete cosine transform to denoise the frame. Cao et al. [15] relied on spatial dependence between adjunct pixels to estimate both the ambient temperature and the correction. Zhao et al. [16] solved an optimization problem, with a constraint on the directional gradients of each frame.

Recent works apply deep learning (DL) methods for single-image NUC. Jian et al. [17] learned the nonuniformity pattern from the filtered high-frequencies of the frames. He et al. [18] trained a CNN that outputs a corrected image end to end. Chang et al. [19] constructed a multi-scale network to reconstruct a corrected frame. Saragadam et al. [20] solved an optimization problem with a NN as a prior, and a physical model as a constraint. Oz et al. [21] modeled the nonuniformity and trained a network based on the physics of the acquisition model.

Single image methods require only a single frame so they are easier to apply, but their performance are degraded as opposed to scene-based methods.

*2) Scene based:* Scene-based works rely on the assumption that the change in ambient temperature is slower than the frame rate, thus the gain and offset are constant between consecutive frames.

Harris and Chiang [22] calculated a shift and normalization terms per-pixel and updated these terms recursively when new frames arrive. Hardie et al. [23] registered the frames with another method and then averaged the results per-pixel. Vera and Torres [24] improved the NN suggested by Scribner et al. [13] with an adaptive learning rate and a different loss function that accounts for multiframe information. Averbuch et al. [25] reformulated the NUC problem into a Kalman filter. Zuo et al. [26] estimated the *irradiance* per-pixel between two frames. Papini et al. [27] approximated the gain and offset from multiple pairs of blurred and sharp images.

The common characteristic between these previous works is that an update step must be performed when new frames arrived, before the correction step. The combined update and correction steps are computationally intensive and pose a constraint on the run-time of the system.

A method based on neural networks to simultaneously estimate the scene temperature and correct the nonuniformity using multiframes information has not been achieved yet.

The study builds on the image acquisition model, which describes the relationship between the observed scene and the output of the camera (Sec. III-A). By leveraging redundant information across multiple frames and ambient temperature data, the study develops a kernel estimation network (KPN) that uses DL techniques to estimate the temperature of each pixel (Sec. IV-A).

The efficacy of the method is demonstrated through tests on real measurements obtained using an uncooled IR camera with those from a scientific radiometric camera. These tests illustrate the ability of the method to correct for nonuniformity and estimate temperatures accurately across different cameras (Sec. V-A).

Our main contributions are: (i) we exploit the redundant information between frames to simultaneously estimate the scene temperature and correct the nonuniformity using a deep neural network; (ii) we impose the physical model of the camera as a constraint on the network to enhance the temperature estimation accuracy; (iii) we incorporate the ambient temperature data as an additional input to the network to further improve the accuracy of the temperature estimation; and (iv) we demonstrate the advantages of using multiple frames over single-frame methods through extensive experiments on synthetic and real data.

## III. Background

We develop the physical image acquisition model of the IR camera in Sec. III-A, and then expand it to multiple frames in Sec. III-B.

### A. Image acquisition

A *blackbody* is an ideal Lambertian surface that emits the maximal radiation at any given wavelength. The spectral density of radiation emitted from a blackbody is described by Planck's law [4]:

$$M_\lambda(T) = \frac{2\pi hc^2}{\lambda^5}\frac{1}{\exp(\frac{hc}{\lambda kT})-1} \quad [W\cdot sr^{-1}\cdot m^{-3}] \tag{1}$$

where $T$ is the temperature of the blackbody in Kelvin, $\lambda$ is the radiation wavelength, $h$ is Planck's constant, $k$ the Boltzmann constant and $c$ the speed of light.

The power emitted over the entire bandwidth is found using Stefan-Boltzmann law [4]:

$$M(T) = \int_0^\infty M_\lambda(T)d\lambda = \sigma\cdot T^4 \quad [W\cdot sr^{-1}\cdot m^{-2}] \tag{2}$$

where $\sigma$ is the Stefan-Boltzmann constant.

The equations above holds for an ideal blackbody. Real objects can never emit the maximal radiation for a given wavelength due to physical constraints (e.g., material, viewing angle). The ratio between the ideal emission and the practical emission of an object is called *emissivity*. Thus, the Stefan-Boltzmann law for radiance power of practical objects is:

$$M(T) = \sigma\cdot\epsilon\cdot T^4 \quad [W\cdot sr^{-1}\cdot m^{-2}] \tag{3}$$

where $\epsilon$ is the emissivity.

Estimating the incident power by an object on a microbolometer is done by integrating over the physical dimensions of the system on Eq. (3). The incident power on the microbolometer can be written as [4]:

$$\phi(T) = \gamma\cdot\sigma\cdot\epsilon\cdot T^4 \quad [W] \tag{4}$$

where $\gamma$ is a coefficient that accounts for the dimensions of the object and field of view of the camera.

The intensities of the pixels in radiometric IR cameras (i.e. gray levels) are linearly proportional to the incident power on the microbolometer. To model the intensities, we consider a small environment near a reference temperature $T_0$ and expand the Stefan-Boltzmann law in Eq. (4) by Taylor series. In Kelvin, the temperature of the object can be considered a small perturbation around a reference temperature, because the reference temperature is usually hundreds of Kelvin, while

$\Delta T$ is usually tens of Kelvin. The Taylor expansion of Eq. (4) is:

$$\begin{aligned} I(t_{obj}) &= \gamma\epsilon\sigma T^4 = \gamma\epsilon\sigma(\Delta T + T_0)^4 \\ &\approx 4\gamma\epsilon\sigma T_0^3 \Delta T + \gamma\epsilon\sigma T_0^4 \\ &\approx g \cdot t_{obj} + d \end{aligned} \tag{5}$$

where $I(t_{obj})$ is the gray level output of the IR camera, $g = 4\gamma\epsilon\sigma T_0^3, d = \gamma\epsilon\sigma T_0^4$ are the gain and offset coefficients. Using the relation between Kelvin and Celsius, we denote $t_{obj} \equiv T - 273.15$ in $^\circ C$.

Eq. (5) shows that the radiation is linear in scene temperature in the small environment near $T_0$, with the term $g$ dependent on the object temperature, and the term $d$ independent of the object temperature.

The incident power $\phi(t_{obj})$ in Eq. (4) changes the temperature of the microbolometer by a small fraction. The change in temperature also changes the electrical resistance of the microbolometer [28]. By applying a constant electrical current on the microbolometer and using Ohm-like law, a mapping between the incident power and the voltage of the microbolometer can be derived [4]. In a low-cost uncooled IR camera, the resistance of the microbolometer changes with the ambient temperature.

To account for the effects of the ambient temperature on the resistance of the microbolometer, the gain and offset of the IR camera are modelled as a function of the ambient temperature [7]:

$$I(t_{obj}, t_{amb}) = g(t_{amb}) \cdot t_{obj} + d(t_{amb}) \tag{6}$$

For a given ambient temperature, the gray levels of pixel $[u, v]$ can be written as:

$$I(t_{obj})[u, v] = g[u, v] \cdot t_{obj}[u, v] + d[u, v] \tag{7}$$

The gain and offset are two dimensional (2D), and together they model the space-variant nonuniformity.

The signal-to-noise ratio of uncooled IR cameras is often low due to noises, with the most dominant noises being $\frac{1}{f}$ and electronic (Johnson) noise [28, Chapter 5]. The $\frac{1}{f}$ noise is more dominant because the camera operates at a low frequency. $\frac{1}{f}$ noise can be modelled as Gaussian [29] with zero mean.

### *B. Multiframes*

Consecutive frames over a brief period of time have an overlap between them (an example of a real UAV pattern with overlapping frames can be seen in Sec. V-A). These consecutive frames are called a burst. The overlap between frames implies that the same object appears in multiple frames. As seen in Eq. (5), the gain and offset are dependent on the pixel location on the sensor, thus different views of the same object can be exploited as redundant information. The redundant information between frames is demonstrated in Fig. 2.

To exploit the redundant information, first an object must have the same coordinates across all frames. To achieve coordinate alignment registration is performed.

Image registration is the process of aligning two or more images of the same scene taken from different viewpoints or at different times [30]. Transforming a source frame toward the coordinate system of another destination frame is called a projective transformation, or an homography [31, Chapter 0]. An homography transformation preserves co-linearity between the frames. Moreover, an homography is invertible and linear by definition [31, Def. 2.9]. In layman terms, the homography preserves the shapes and relations between objects.

After applying the homography on the source frame, an object should have the same coordinates in both the source and destination frames. The transformed source frame is called a warped frame. Expanding to $N$ frames, there exists a set of projective transformations $m_1, \ldots, m_N$ towards a common plane such that the overlap between the frames is maximal [31, Ch. 4]. Objects that appear in the overlapping area will have the same coordinates in every warped frame. For our practical use, we choose a pivot frame for each burst of frames and annotate the pivot frame as $\mathcal{I}$.

An underlying assumption throughout this work is that the gain and offset in Eq. (7) are constant for a series of frames taken over a short duration of time (a second). This assumption holds because the ambient temperature of the camera changes at a much slower rate (several minutes).

Let $X$ be the fourth power of an accurate 2D temperature map, and $I_1, \ldots, I_N$ be a set of $N$ frames of $X$ captured by the IR camera. $I_i$ is in gray levels. $I_i^{x_i, y_i}$ is the value of the pixel in the $[x_i, y_i]$ location of $I_i$. $[u, v]$ are the coordinates of $\mathcal{I}$ the pivot frame. The frames in the burst can be formulated as:

$$\begin{aligned} I_1^{x_1,y_1} &= g^{x_1,y_1}(t_{amb}) \cdot m_1^{-1}\left(X^{u,v}\right) + d^{x_1,y_1}(t_{amb}) \\ &\vdots \\ I_N^{x_N,y_N} &= g^{x_N,y_N}(t_{amb}) \cdot m_N^{-1}\left(X^{u,v}\right) + d^{x_N,y_N}(t_{amb}) \end{aligned} \tag{8}$$

where $m_1, \ldots, m_N$ are the set of homographies that transforms each frame into $\mathcal{I}$ the pivot frame. The zero-mean noise $\mathcal{N}$ was omitted for brevity.

Eq. (8) formulates the acquisition process of a frame as projecting the temperature map $X$ using an inverse of the homography $m_i$, and than sampling the projected $X$ by applying the gain $g$, offset $d$ and noise $\mathcal{N}$. Notice that an object will be sampled at different coordinates for each frame, and since the gain and offset are spatially-variant the object will have different gain and offset for each frame.

The result in Eq. (8) means that an object appearing in pixel $[u, v]$ of the temperature map $X$ will have multiple representation with different values of $g, d$ and $\mathcal{N}$, enabling the use of redundant information between the frames.

Redundant information between frames was used for many image restoration tasks such as super-resolution [32], [33], denoising [34] and debluring [35]. Many recent works in the area use DL for either the alignment [36], the fusion between frames [37], or both [38], [39], [40].

## IV. Temperature estimation

The proposed method simultaneously estimates the scene temperature and corrects nonuniformity from a burst of consecutive frames.

An overview of the method is presented in Fig. 1. The IR camera captures overlapping gray-level frames. A burst of

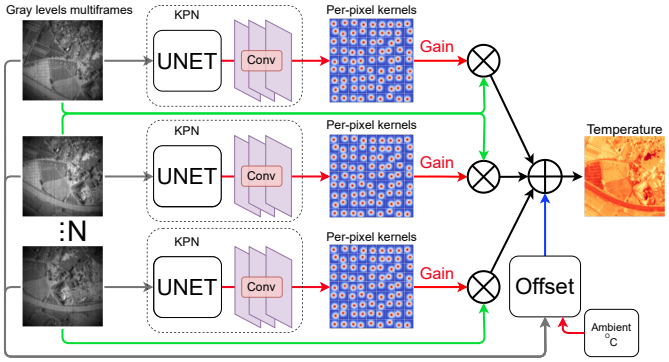


Fig. 4. Schematics of the model. The gray-level multiframes are fed into the kernel prediction network (KPN), and the KPN outputs the per-pixel kernels for each frame. Each frame is divided to overlapping patches with the same support as the kernels. The patches and the kernels are multiplied element-wise and each product is summed, resulting in a 2D gain map for each frame. All the 2D gain maps are summed depth-wise, resulting in a single 2D map. The offset, a single scalar value, is added to the single 2D map to get the estimated temperature map. A detailed description of the network architecture and an enlarged figure of the network Fig. 32 is in the supplementary material.

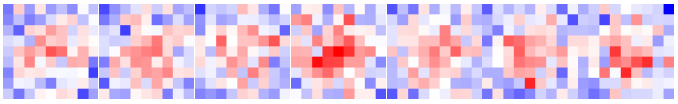

Fig. 5. Example of kernels estimated by the network, each kernel for a different frame in the burst. The dimensions of the kernels in the figure are $9 \times 9$ and they are predicted for the center pixel of 7 consecutive frames. The middle kernel is from the reference frame. Red has a higher magnitude, and blue has a lower magnitude.

consecutive gray-level frames are registered towards $\mathcal{I}$, the pivot frame that has the maximal overlap to all other frames. The registered gray-level frames are the input to the network, along with the ambient temperature. The output of the network is a 2D map of the estimated scene temperatures.

### A. Network

In Eq. (8) we show that different views of the same object have usable redundant information. To exploit the redundancy, these different perspectives require accurately mapping the frames to $\mathcal{I}$ the pivot frame.

Naïvely, a temperature map $\hat{X}$ can be estimated from Eq. (8) by:

$$\begin{aligned} \hat{X}^{u,v}_{naive} &= \frac{1}{N}\sum_{i=1}^{N}\left[\frac{1}{g^{u-x_i,v-y_i}}\tilde{I}^{u,v}_i - \frac{d^{u-x_i,v-y_i}}{g^{u-x_i,v-y_i}}\right] \longrightarrow \\ \hat{X}^{u,v}_{naive} &= \frac{1}{N}\sum_{i=1}^{N}\left[G^{u-x_i,v-y_i}\tilde{I}^{u,v}_i + D^{u-x_i,v-y_i}\right] \end{aligned} \tag{9}$$

for 2D coefficient maps $G$ and $D$.

The naïve approach requires exact registration between frames. The information must be located on the exact same coordinates across all frames. Inaccurate registration leads to artifacts or ghosting, as well as inexact temperature estimation. Even with a robust registration framework there is always some degree of misalignment between frames, so the naïve approach is unsuitable for practical use.

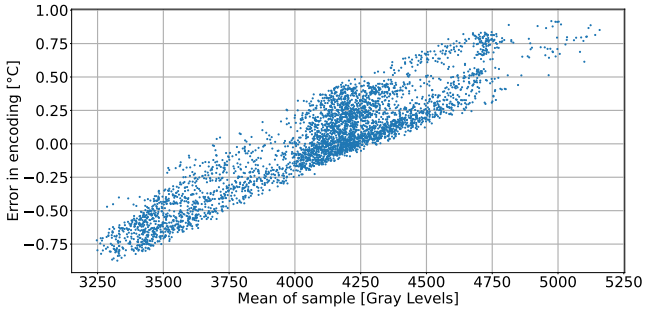


Fig. 6. The difference between the true temperatures and the offset block's estimate from Eq. (10), for different average input gray levels.

The method for temperature estimation proposed in this work is robust to misalignment between frames. The frames are registered towards $\mathcal{I}$ using any off-the-shelf registration method, then fed into a neural network that predicts a kernel for each pixel in every frame of the burst. The kernels are then applied on overlapping patches around each pixel by an inner product between the patch and kernel. Our method is based on kernel prediction networks (KPN) proposed by De Brabandere et al. [41]. Fig. 5 shows kernels predicted by the network. The kernels compensate for misalignment between frames by spatially shifting their center to compensate for shifts.

The architecture of the temperature estimation network is based on UNET [42], with the kernel prediction block attached to the rear end of the decoder. The kernel prediction block is composed of three $1 \times 1$ convolution layers with activations, and is described in Table II at the supplementary material. The entire network architecture is detailed in the supplementary material.

Although the KPN corrects nonuniformity, its temperature estimation is inaccurate. To improve the temperature estimation to match radiometric cameras, we used the ambient temperature as prior information to calibrate the output of the network. The offset between the gray level frames and the temperatures was modeled as a polynomial of the mean of the gray level frames and the ambient temperature:

$$\tilde{d}\left(\tilde{I}, t_{amb}\right) = \frac{1}{N}\sum_{n=1}^{N}\left[\sum_{i,j=0}^{\nu}\delta_{i,j}\cdot \text{Mean}\left(\tilde{I}_n\right)^i \cdot t^j_{amb}\right] \tag{10}$$

for $\tilde{d}_n$ the offset for frame $n$, $\text{Mean}\left(\tilde{I}_n\right)$ the spatial mean of the grey-level $n$'th frame, $t_{amb}$ the ambient temperature, $\delta_{i,j}$ the coefficients of the polynomial, and $\nu$ the degree of the polynomial.

The offset block was implemented by a fully connected layer that was jointly trained with the network. We found that a polynomial of degree $\nu = 4$ offers sufficient improvement in the accuracy of the temperature estimation, and that training the offset block separately from the network does not offer significant improvement. Fig. 6 shows the results of the offset block. The error between the temperature estimation of the offset block and the GT temperature is shown. The error is sub-degree Celsius, and the offset is accurate enough to calibrate the output of the network.

The following equation describes the temperature estimation by applying KPN to the image acquisition model. To combine the information from multiple frames, the gain term in Eq. (8) is generalized as KPN, and the information from all frames are used. The kernels applied to each pixel handles the nonuniformity and noise, while the offset term in Eq. (10) handles the thermal calibration:

$$\hat{X}^p = \sum_{n=1}^{N} \left\langle \mathcal{K}_n^p, S^p\left(\tilde{I}_n^p\right)\right\rangle + \tilde{d}\left(\tilde{I}, t_{amb}\right) \tag{11}$$

for $N$ the number of frames in a burst, $\mathcal{K}$ the kernel of size $K \times K$, and $S(\cdot)$ a function that samples a $K \times K$ patch around a pixel $p$ in the support of the frames.

A scheme of the model is shown in Fig. 4. The registered burst of frames is fed into the network, which outputs a kernel for each pixel in each frame. These kernels serve as the *gain* in Eq. (11). The registered frames are also fed to the offset block along with the ambient temperature, which outputs the *offset* term in Eq. (11). The gain is applied to the frames and the results are depth-wise summed. The scene temperature estimation is obtained by adding the offset term to the result of the depth-wise summation.

### B. Loss functions

The loss is comprised of a fidelity term, a gradient smoothness term and a structural term:

The structural term $\mathcal{L}_{\mathcal{SSIM}}$ maximizes the commonly-used structural similarity metric (SSIM). SSIM loss improves results in image-restoration tasks [43].

The fidelity and gradient terms are similar to Mildenhall et al. [44], with the exception that the $L_1$ loss is used instead of the $L_2$, because $L_1$ was shown to be more robust to outliers [45]. The loss function is formulated as:

$$\begin{aligned}\mathcal{L} = &\left|\left|M(\hat{X}) - M(X)\right|\right|_1 + \\ &\lambda_1 \left|\left|M(\nabla\hat{X}) - M(\nabla X)\right|\right|_1 + \\ &\lambda_2 \cdot \mathcal{L}_{\mathcal{SSIM}}(M(\hat{X}), M(X))\end{aligned} \tag{12}$$

for $\hat{X}$ the temperature estimated by the network, $X$ the GT temperature, $M$ a mask of valid pixels in the registration process, $\lambda_1, \lambda_2$ hyperparameters to balance to losses and $\nabla$ the magnitude of the Sobel operators. The mask is produced by the registration algorithm.

The final values of the hyperparameters were set to $\lambda_1 = 0.1$ and $\lambda_2 = 0.01$.

### C. Synthetic data

The network was trained with synthetic data in a supervised manner. The inputs to the network were created from accurate 2D temperature maps collected using a scientific-grade IR camera (`A655sc`). A degradation model of a low-cost IR camera (`Tau2`) was applied to the temperature maps, transforming them to grey-level frames. As a result, the network trained on transforming gray-level frames to accurate temperature maps.

The goal of the degradation model was to faithfully transform temperature maps into gray level maps, allowing the supervised training process of the network. The modeling process had three stages. First, collecting data with the IR camera in a controlled environment. The second stage was to find per-pixel coefficients using the image acquisition model in Sec. III-A. The last stage was to use adjunct pixel dependencies as a constraint on the degradation model.

The degradation model required frames of objects with known temperature by a `Tau2` in different ambient temperatures. To collect this data, a `Tau2` was placed inside an environmental chamber in front of a scientific-grade blackbody (`SR-800N`). The blackbody and environmental chamber were cycled to different pairs of $(t_{amb}, t_{obj})$, and frames were acquired at the different permutations. Fig. 3 is an example from the collected data.

The `Tau2` was modeled by the image acquisition model in Eq. (6). The calibration was done according to Nugent et al. [7], by using a third-degree polynomial to approximate the coefficients $g, d$. For each pixel in the sensor, Eq. (6) can be formulated as:

$$I_p(t_{obj}, t_{amb}) = \sum_{i=0}^{3}\left(g_{i,p} \cdot t_{amb}^i \cdot {t_{obj}}_p^4 + d_{i,p} \cdot t_{amb}^i\right) \tag{13}$$

for $g_{i,p}, d_{i,p}$ the i'th gain and offset coefficients at pixel $p$, respectively. Eq. (13) can be rewritten as matrix multiplication:

$$\begin{aligned} T_p^n = &\quad [{t^4_{obj}}_n \quad \dots \quad {t^4_{obj}}_n t^3_{amb\,n} \quad 1 \quad \dots \quad t^3_{amb\,n}] \\ C_p = &\quad [g_{0,p} \quad \dots \quad g_{3,p} \quad d_{0,p} \quad \dots \quad d_{3,p}]^T \\ I_{N,p} \equiv &\; T_{N,p} \cdot C_p \end{aligned} \tag{14}$$

for $T_p^n$ that contains the appropriate temperatures of the n'th sample of a permutation, and $T_{N,p}$ a matrix with all the temperatures corresponding to all the samples of the permutation as rows, and $I_{N,p}$ is a matrix with all the acquired samples as rows. Eq. (14) is solved using least-squares to find the coefficients:

$$C_p = T^{\dagger}_{N,p} \cdot I_{N,p} \tag{15}$$

for $T^{\dagger}_{N,p}$ the Moore-Penrose pseudo-inverse.

Stacking all the 2D coefficient maps $C_p$ into a 3D tensor $\underline{\underline{C}}$, with $\underline{\underline{C}}[0]$ being the 2D map of coefficient $g_0$ etc.

The degradation model described in Eq. (15) is per-pixel, thus unique to each camera. Meaning that nonuniformity will also be modeled by the coefficients (e.g., dead pixels, fixed-pattern noise). This realization limits the usability of the degradation model only for the specific camera that collected the data.

To enable the degradation model to generalize for different cameras, the final stage in the degradation model exploits the circular symmetry of the nonuniformity and uses the dependency between neighboring pixels to enable the degradation model to generalize for other cameras.

Nonuniformity has a circular symmetry around the middle of the frame [4]. This is due to the ambient temperature of the camera, generating radiation from the chassis and lens, which is also reflected onto the sensor. Rays of thermal radiation from the body of the camera travels to the sensor and affect each pixel differently. The superposition of these rays on each pixel creates the circular symmetry of the nonuniformity. An example of the circular symmetry can be seen in Fig. 3.

The spatial dependency was modeled as a radial map around the middle of the frame. The radial map was constructed from two mesh-grids $\underline{\underline{H}}, \underline{\underline{W}}$ with dimensions the same as the frames. Each row of $\underline{\underline{H}}$ and each line in $\underline{\underline{W}}$ runs from $-0.5$ to $0.5$, such that $\underline{\underline{H}} = \underline{\underline{W}}^T$. The radial map is defined as:

$$\begin{aligned} \underline{\underline{R}} &= \sqrt{\underline{\underline{H}}^2 + \underline{\underline{W}}^2}, \qquad \underline{\underline{H}}, \underline{\underline{W}}, \underline{\underline{R}} \in \mathcal{R}^{h,w} \\ \underline{\underline{R}} &= \sqrt{\begin{bmatrix} -0.5 & \dots & -0.5 \\ \vdots & \ddots & \vdots \\ 0.5 & \dots & 0.5 \end{bmatrix}^2 + \begin{bmatrix} -0.5 & \dots & 0.5 \\ \vdots & \ddots & \vdots \\ -0.5 & \dots & 0.5 \end{bmatrix}^2} \end{aligned} \tag{16}$$

for $h, w$ the dimensions of the frames. The power of the matrix is performed element-wise.

The coefficient maps are modeled as:

$$\underline{\underline{\hat{C}}}[i] = \sum_{j=0}^{M} m_i \cdot \underline{\underline{R}}^j, \qquad m_i \in \mathcal{R},\ \underline{\underline{C}}, \underline{\underline{R}} \in \mathcal{R}^{h,w} \tag{17}$$

for $m_i$ the spatial coefficient, and $M$ the number of spatial coefficients. Least-squares is solved to find the spatial coefficients.

Estimating a frame from a given temperature map is performed by:

$$\hat{I}(t_{obj}, t_{amb}) = T(t_{obj}, t_{amb}) \cdot \hat{C} \tag{18}$$

for $T(t_{obj}, t_{amb})$ the temperatures vector in Eq. (14).

The final degradation model was noiseless and only contained low frequencies. Random fixed pattern noise and Gaussian noise were added to the model during training. This enabled the network to converge to a general solution, applicable on different cameras with various degradation profiles.

### *D. Training procedure*

The network was trained on a single Nvidia Titan A100. The network was written in Python 3.10 [46] using Pytorch 1.13 [47]. The seed was set to 42, and the CUDNN backend was set to deterministic mode. The network was trained using the ADAM optimizer [48] with a learning rate of $10^{-4}$. The learning rate was halved on a validation loss plateau of more than 3 epochs. The network was run for 60 epochs with batches of 16, meaning each epoch was roughly 800 iterations. Early stopping was applied for a validation loss plateau of 8 epochs. The weights were initialized using the orthogonal scheme [49] with a scaling of $10^{-1}$. The hyperparameter search was ran twice with different seeds (42 and 24), and the best results were chosen.

Multiframes were simulated by randomly sampling homographies for each frame in the dataset, creating different views of the same frame. The inverse homographies were used to register all the views towards the original frame, which was set as $\mathcal{I}$ the pivot frame. The sampled homographies either created a random walk from one side of the temperature map to the other side, or a hover above a random point in the frame. The overlap between the different views were randomly set between $60\%$ to $80\%$, similar to a UAV flight scenario, as seen in the scenario described in Fig. 11. The homography and frame warping was implemented with the package Kornia v0.67.

Imperfect registration was simulated by randomly adding perturbations to the inverse homographies - random translation of up to $\pm 2$ pixels and noise from the distribution $\mathcal{N}(0, 5 \cdot 10^{-}5)$ to the perspective elements of the homography (commonly known as $h_{31}, h_{32}$). Random horizontal and vertical flips, and $90^\circ$ rotations were applied to the frames before the homography sampling.

The gray level frames were cropped to $128 \times 128$ patches before entering the network. For validation, a constant cropping was applied around the middle of the frame, and no other augmentations were applied.

Random Gaussian noise with $\sigma^2 = 5$ gray levels and FPN were generated for each frame (Sec. IV-E). FPN was generated as:

$$\begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}_{h \times 1} \cdot \begin{bmatrix} U[u_{\min}, u_{\max}] \\ \vdots \\ U[u_{\min}, u_{\max}] \end{bmatrix}^T_{1 \times w} \tag{19}$$

where $U$ is uniform distribution. $u_{\min}, u_{\max}$ were chosen as $u_{\min} = 0.9, u_{\max} = 1.01$. The Gaussian noise and FPN were only generated once for each frame and used throughout the entire validation process. This was done for reproducible of results between experiments.

Normalization to range [0,1] is applied to both the temperature map and the gray level frame. Throughout the training and validation sets, the maximal and minimal values of the temperature maps and the maximal and minimal values of the gray level frames were obtained. To apply the normalization on the temperature maps:

$$\bar{X} = \frac{X - X_{\min}}{X_{\max} - X_{\min}} \tag{20}$$

where $\bar{X}$ is the normalized input and $X_{\min}, X_{\max}$ are the minimal and maximal temperatures over all datasets.
To apply the normalization for the gray level:

$$\bar{I}(t_{amb}) = \frac{I(t_{amb}) - I_{\min}}{I_{\max} - I_{\min}} \tag{21}$$

where $\bar{I}$ is the normalized gray level frame and $I_{\min}, I_{\max}$ are the minimal and maximal gray levels over all datasets.

The following pipeline summarizes the creation of samples for the network. First, an accurate temperature map is sampled from the dataset. $N$ homographies are randomly sampled and applied to the temperature map to create an overlapping burst of frames. The model described in Sec. IV-C is applied to each frame in the burst to turn in to a gray level frame (Eq. (18)). The same FPN is applied to all frames in the burst (Eq. (19)), and random noise is applied to each frame in the burst separately. Finally, normalization is applied to the ambient temperature (Eq. (20)) and overlapping gray level frames (Eq. (21)), and both are passed to the network.

### *E. Data*

The dataset used for training consisted of $12,897$ frames, while the validation set was composed of $4,723$ frames, all of which were captured by an unmanned aerial vehicle (UAV)

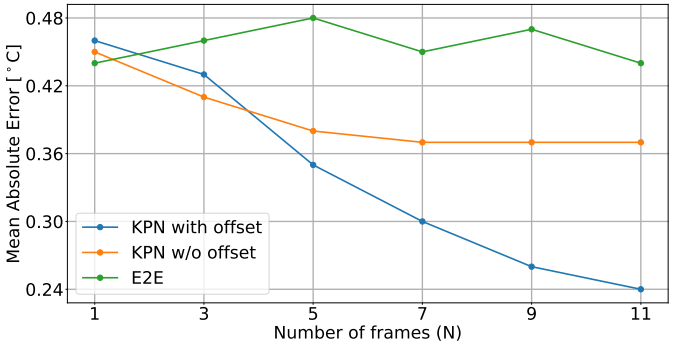


Fig. 7. MAE in $^{\circ}C$ as a function of the number of frames $N$ for different network configurations.

flying at a height of $70-100$ meters above various agricultural fields in Israel. Only clear and in-focus frames were selected for the dataset manually by a human user.

The noise variance was established by analyzing the measurements taken in the environmental chamber. All frames were stacked depth-wise, and the variance of each pixel was calculated, resulting in a 2D variance map. The mean of the variance map was $5$ gray levels, which was used as the noise variance in the network training. The influence of $t_{amb}$ and $t_{obj}$ on $\sigma^2$ was determined to be insignificant.

To prevent data leakage between the training and validation sets and evaluate the network's ability to generalize to new data, the validation sets were captured at the same locations as the training sets but on different days. This validation approach was maintained across all training schemes to ensure a fair comparison between different experiments. The same split between the training and validation datasets was maintained throughout the study.

## V. Results

To show the efficacy of our method, we compared the mean absolute error (MAE) of the temperature estimation with different blocks of the network, namely with and without the offset estimation block, and using an end-to-end (E2E) network instead of the KPN architecture. The results are displayed as a function of the number of frames $N$ in Fig. 7. While our method can handle misalignments between frames, other methods require a perfect alignment, which is impossible in real-world scenarios. As a results, comparing temperature estimation with other methods is impossible, and we can only compare NUC between our method and other methods on the $\mathcal{I}$ pivot frame. Moreover, the other methods are not radiometric, and thus cannot be used for temperature estimation.

Fig. 7 demonstrates the superiority of our method, evident by the low MAE for almost every number of frames. The E2E was a UNET [42] architecture similar to our KPN network. The main difference was that the last layer estimated the per-pixel result instead of outputting the kernel. An hyper-parameter search was also performed for the E2E solution for fair comparison (number of channels and normalization). The MAE of the E2E network in Fig. 7 is unaffected by the number of frames, while the KPN results improve with the number of frames, indicating that the E2E network only uses the reference frame. Moreover, the MAE results for E2E were worse than for the KPN network.

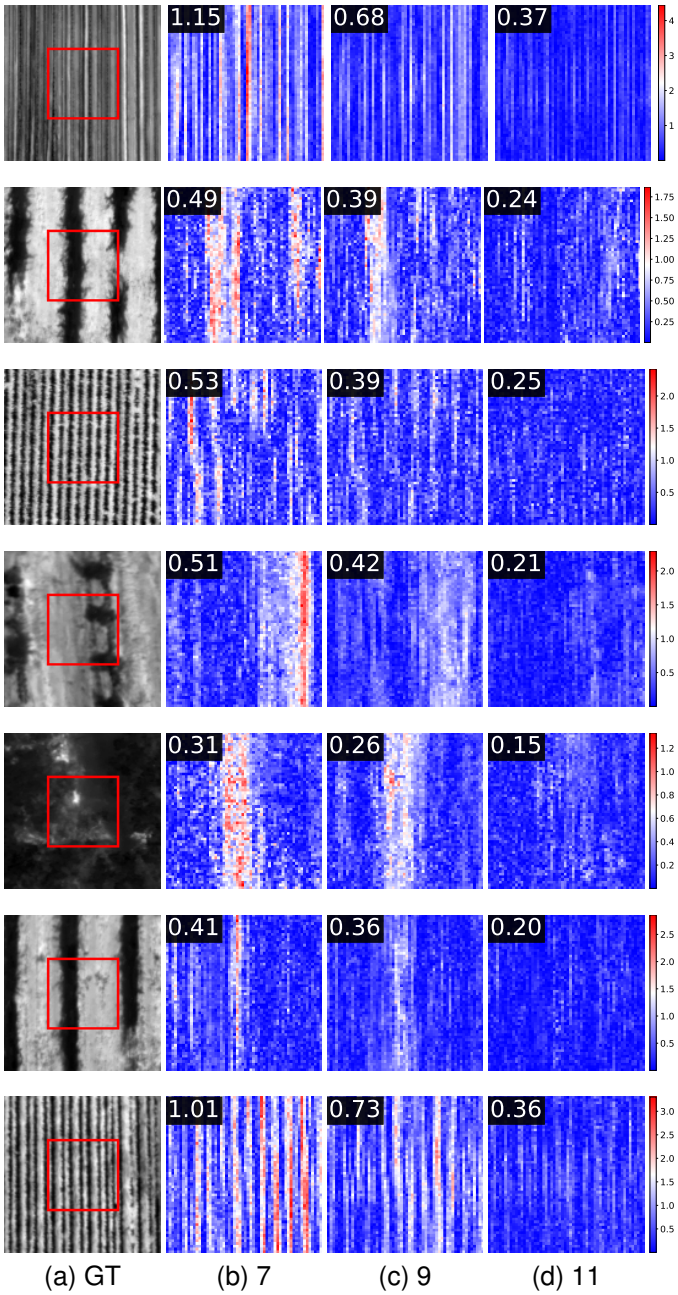


Fig. 8. Difference between our temperature estimation the GT. The left-most figure is the GT. The next figures are the zoom-in of the area inside the red rectangle. The nubmer of frames used for the estimation is bellow each map, from left to right 7, 9 and 11 frames. The MAE between the GT and our estimation is written on the top-left corner of each map.

The offset block greatly improves the results with the number of frames as seen in Fig. 7, with more than a $0.1^{\circ}C$ improvement for KPN and more than $0.2^{\circ}C$ improvement relative to E2E for $N = 11$, indicating that the offset block is beneficial. The results suggest that increasing the number of frames without the offset block reaches a plateau around

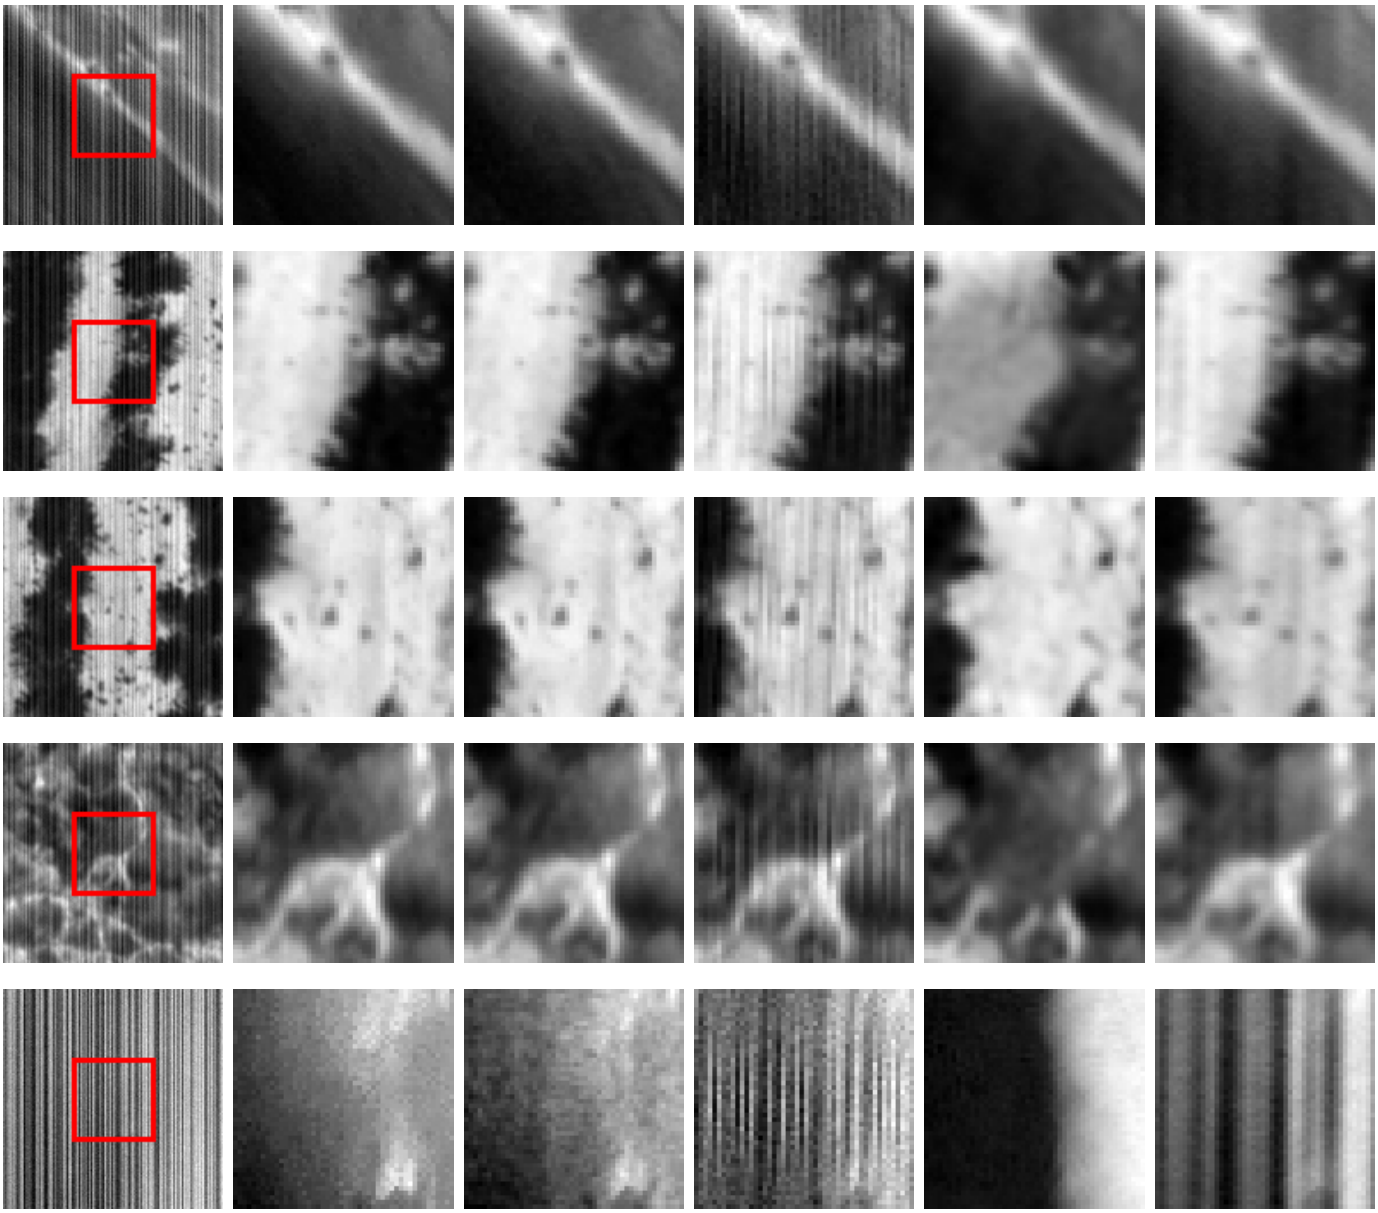


Fig. 9. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results.

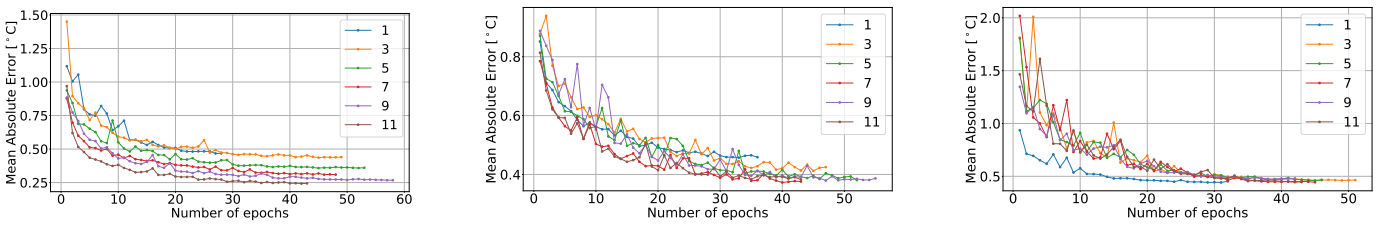


Fig. 10. Convergence of the validation MAE loss for different architectures - KPN with offset (a), without offset (b), and E2E (c). Each color represents a different number of frames $N$.

$N = 5$ and does not improve the results further, in contrast to the offset block which continues to improve the results with more frames. Since the offset block is lightweight, it offers significant improvement with little computational cost.

The effect of the number of frames $N$ is shown in Fig. 8. The figure shows per-pixel error in temperature estimation for different number of frames. The left-most figure (a) shows the GT temperature map. The absolute difference per-pixel between the GT and the estimation of our method for the area inside the red rectangle is shown in the next figures (b), (c), (d). Each figure shows the estimated temperature for a different number of frames - (b) $N = 7$ frames, (c) $N = 9$ frames and (d) $N = 11$ frames. The color bar on the right of each row shows the error range for the row in $^\circ C$. The MAE in $^\circ C$ between the GT and the estimation is written in the top-left corner of each difference map. Each row is a different frame. The improvement caused by the number of frames is clear by the homogeneity in the difference map and the MAE decreasing as a function of $N$. More examples are available in the supplementary material on Figures 20 to 24.

Because other methods are not radiometric and essentially only improve the appearance of a frame, we could only compare NUC results with other methods. Fig. 9 displays NUC results of different methods. The first column (a) shows the reference sample frame. The second column (b) shows the GT temperature map. The third column (c) shows the results of our method. The fourth column (d) shows the results of ADMIRE [14] performed on each frame separately and than registered and averaged. The fifth column (e) shows the estimation of DeepIR [20] and the sixth column (f) shows the estimation of He et al. [18]. All results were obtained with $N = 11$.

Our method NUC is better than other methods, as evident in Fig. 9. ADMIRE [14] fails to rectify the FPN, DeepIR [20] hallucinates details (e.g., the deformation in the junction on the fourth row, or the abrupt black to white edge on the fifth row). He et al. [18] fails to handle the FPN. Both DeepIR [20] and He et al. [18] oversmoothes the results. These methods have low fidelity, and thus are unable to serve for the purpose of temperature estimation. More results are available in the supplementary material on Figures 14 to 19.

Fig. 10 depicts the convergence of the validation MAE loss of the E2E and the KPN with the offset block and the KPN without the offset block, as a function of the number of frames. Notice that the loss for the E2E networks convergence to roughly the same value for all number of frames, while the KPN-based networks achieve different results as a function of the number of frames $N$. When comparing the convergence with and without the offset block it seems that the offset block has a smoothing effect on the validation loss. This effect might happen because the KPN can concentrate on correcting the NUC, while the offset block handles the temperature estimation.

### A. Real data

We validated the effectiveness of the proposed method on real data. Two cameras, `A655sc` and `Tau2`, were attached to a DJI Matrice 600 UAV and both captured the same scenes in nadir view at a height of $50_m$ above ground at a vertical speed of $10_{m/s}$. The `A655sc` is a scientific-grade radiometric camera which outputs a temperature map of the scene, while the `Tau2` outputs a gray level map corresponding to the radiation flux. An image of the setup can be found in Fig. 31 of the supplementary material. Notice that the `Tau2` used for the experiment is not the one used to collect the calibration data in Sec. IV, which further strengthens the generality and robustness of the proposed method.

The frame rate for the `Tau2` was set to $30_{Hz}$, the resolution of the `Tau2` is $336 \times 256$, the focal length was $9.8_{mm}$ and the sensor size is $4.4_{mm}$ per 256 pixels (in the direction of the flight). The ground sampling distance was $\frac{50 \cdot 4.4}{9.8 \cdot 256} = 0.087_{m/pix}$. The drone passed $\frac{10}{30} = 0.33_m$ every frame. This means that an object moved $\frac{0.33}{0.087} = 3.80_{pix}$ between consecutive frames. Thus, an object could appear in $\frac{256}{3.80} \cong 67$ frames. The `A655sc` field of view was much larger than the `Tau2` field of view, so a frame of the `A655sc` contained multiple frames of the `Tau2`.

The `A655sc` requires accurate ambient parameters to produce a valid temperature map. The ambient temperature and humidity were gathered from a nearby weather station ($28.4^\circ C$ and $32\%$ respectively). The emissivity was tuned using an accurate temperature sensor placed in the scene.

Both cameras were focused to infinity. The height of flight at $50_m$ above ground ensured that the difference in depth of field due to objects was negligible. `A655sc` captured $1,192$ frames at $5_{Hz}$ and the `Tau2` captured $7,152$ frames at $30_{Hz}$.

The frames of the `Tau2` were divided to overlapping groups of 7 frames each. We used 7 frames due to hardware limitations. The frames of each group were registered towards the middle frame of the group. The registration was performed by SIFT feature-matching using the Python package Kornia V0.6.10. The registered frame groups were the input to the network.

The output of the network was the estimated temperature map of the scene. These estimated temperature maps were registered to the `A655sc` temperature maps by hand-picking correspondence points. The final registration was performed using the Python package OpenCV V4.5.1. The GT and estimated temperature maps are presented in Figures 25 to 30 in the supplementary material.

Six results are presented in Fig. 11 and four more are presented in Fig. 13 in the supplementary material.

We present the difference maps between the estimated and the GT temperature maps, produced by the proposed method, in each subfigure. The GT maps are in gray, and the color scale from blue to red is the magnitude of the difference, with blue denoting low and red denoting high errors. The upper-left corner of the upper image displays the MAE of the difference map as a white number.

The MAE values span $0.27^\circ C - 0.54^\circ C$, indicating a high accuracy of temperature estimation comparable to the `A655sc` precision of ($\sim 0.5^\circ C$). This was obtained without applying any thermographic corrections or NUC to the `Tau2` data, relying solely on the radiation flux as gray levels. The

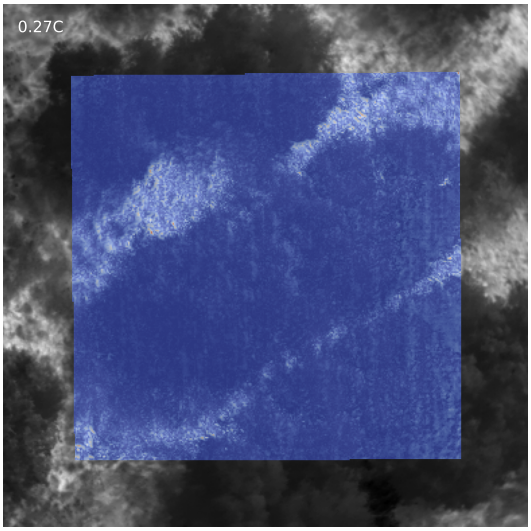


(a)

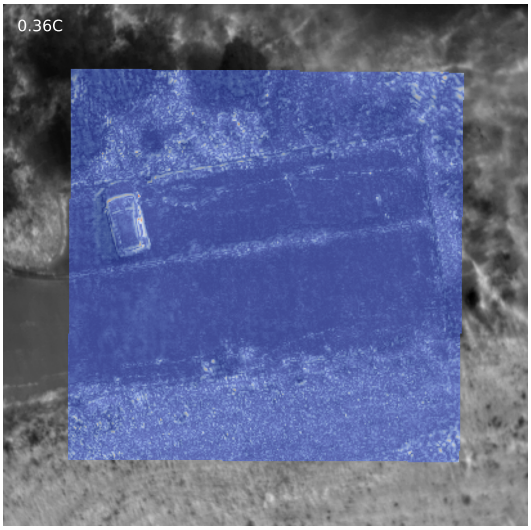


(b)

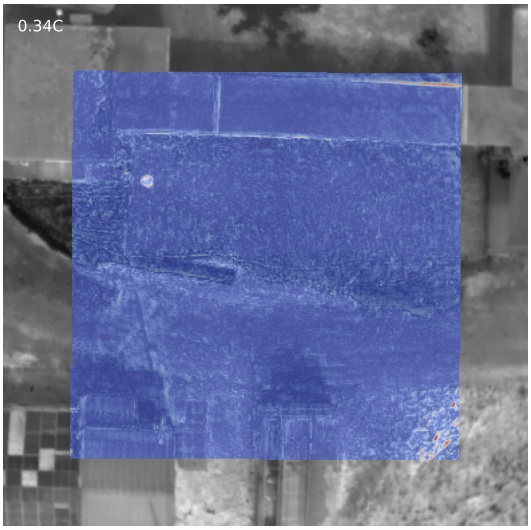


(c)

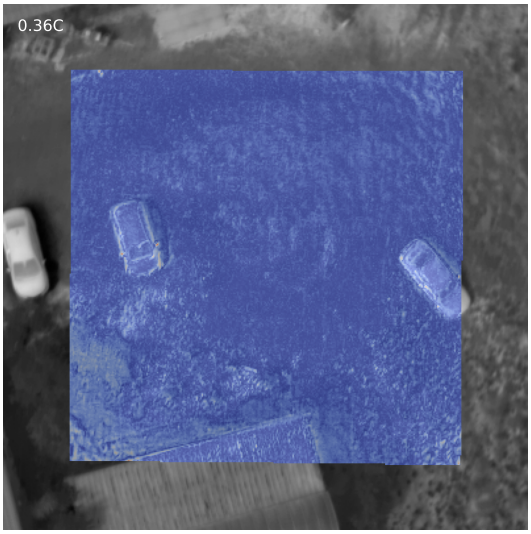


(d)

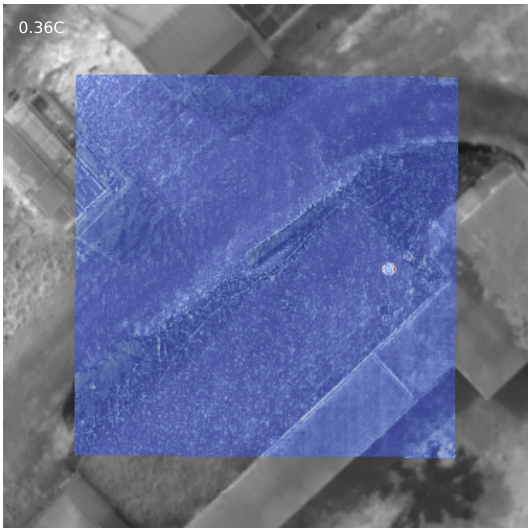


(e)

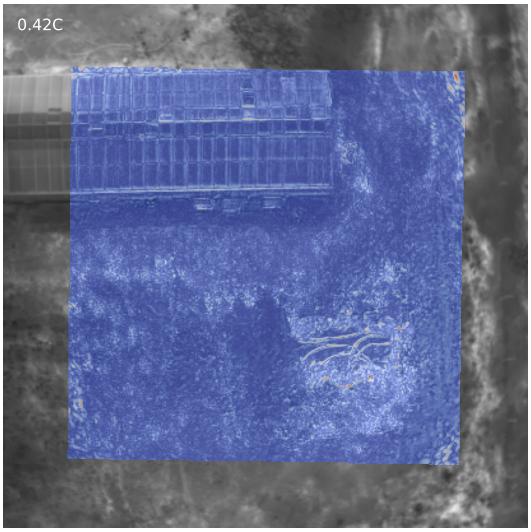


(f)

Fig. 11. Various results on real data. The gray background is the GT temperature map and the colored map is the difference between the GT and estimated temperature maps. The number on the top-left corner is the MAE between the estimated and GT temperature maps.

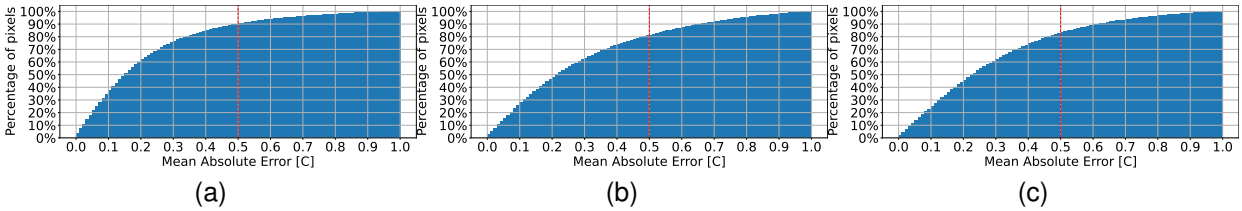


Fig. 12. MAE of pixels as a cumulative function for real data. The y-axis is the percentage of pixels with MAE less than the value on the x-axis. The caption of each subfigure is the corresponding figure in Fig. 11.

supplementary material provides the detailed configuration of the `Tau2`.

The histogram of the MAE between the GT temperature map and the estimated temperature map is shown in Fig. 12. The dashed red lines indicate the $0.5°C$ threshold. All three examples shows that more than $80\%$ of the pixels have MAE less than $0.5°C$. This further solidifies the effectiveness of the proposed method.

## VI. Conclusion

We have presented a novel method for simultaneous temperature estimation and non uniformity correction in IR imaging, based on a DL that incorporates the physical model of the sensor. The method uses redundant information between multiple overlapping frames to infer the scene temperature and correct non-uniformity, without requiring any calibration or external reference. The method also exploits prior knowledge of the camera ambient temperature, which is measured by a built-in sensor, to improve the accuracy and robustness of the estimation.

We have evaluated the performance of the method on synthetic and real data and compared it with existing methods. The results show that the method can achieve high accuracy and low error, and can handle various scenarios, such as changing ambient temperature, moving objects, and complex backgrounds.

We showed that performance improve with the number of frames, showing the benefits of exploiting the redundant information between frames. The training process introduced misalignments between frames, which were handled by the method and did not affect the performance. The method can also generalize well to different camera models and settings and can be easily adapted to different applications. This was demonstrated by real data collected with a different camera mounted on a UAV. The MAE error with the real UAV data were $0.27°C - 0.54°C$, which is comparable to the accuracy of scientific-grade cameras.

The method offers a simple and effective solution for improving the quality and reliability of low-cost uncooled infrared imaging and can potentially enable new applications that require accurate and consistent temperature measurements.

## Acknowledgments

The authors are deeply grateful for the help of Ohaliav Keisar for the UAV data collection.

The authors thank Dr. Yaffit Cohen and Dr. Eitan Goldstein for the UAV data used in this work; and Moti Barak, Lavi Rosenfeld and Liad Reshef for the design and construction of the environmental chamber.

## Funding

The research was funded by the Israeli Ministry of Agriculture's Kandel Program under grant no. 20-12-0030.

## Disclosures

The authors declare no conflicts of interest.

## References

[1] H. D. Adams, M. Guardiola-Claramonte, G. A. Barron-Gafford, J. C. Villegas, D. D. Breshears, C. B. Zou, P. A. Troch, and T. E. Huxman, "Temperature sensitivity of drought-induced tree mortality portends increased regional die-off under global-change-type drought," *Proceedings of the National Academy of Sciences*, vol. 106, no. 17, pp. 7063–7066, 2009. [Online]. Available: https://www.pnas.org/doi/abs/10.1073/pnas.0901438106

[2] H. G. Jones, "Monitoring plant and soil water status: established and novel methods revisited and their relevance to studies of drought tolerance," *Journal of Experimental Botany*, vol. 58, no. 2, pp. 119–130, 09 2006. [Online]. Available: https://doi.org/10.1093/jxb/erl118

[3] R. Bhan, R. Saxena, C. Jalwania, and S. Lomash, "Uncooled infrared microbolometer arrays and their characterisation techniques," *Defence Science Journal*, vol. 59, p. 580, 11 2009.

[4] M. Vollmer and K.-P. Mllmann, *Infrared Thermal Imaging: Fundamentals, Research and Applications*, 2nd ed. Wiley-VCH, 2018.

[5] O. Riou, S. Berrebi, and P. Bremond, "Non uniformity correction and thermal drift compensation of thermal infrared camera," *Thermosense XXVI*, vol. 5405, p. 294, 2004.

[6] M. Schulz and L. Caldwell, "Nonuniformity correction and correctability of infrared focal plane arrays," *Infrared Phys. Technol*, vol. 36, pp. 763–777, 1995.

[7] P. W. Nugent, J. A. Shaw, and N. J. Pust, "Correcting for focal-plane-array temperature dependence in microbolometer infrared cameras lacking thermal stabilization," *Optical Engineering*, vol. 52, p. 061304, 2013.

[8] K. Liang, C. Yang, L. Peng, and B. Zhou, "Nonuniformity correction based on focal plane array temperature in uncooled long-wave infrared cameras without a shutter," *Applied Optics*, vol. 56, p. 884, 2 2017.

[9] S. Chang and Z. Li, "Calibration algorithm for cooled mid-infrared systems considering the influences of ambient temperature and integration time," *Applied Optics*, vol. 58, p. 8118, 10 2019.

[10] N. Oz, N. Sochen, O. Markovich, Z. Halamish, L. Shpialter-Karol, and I. Klapp, "Rapid super resolution for infrared imagery," *Optics Express*, vol. 28, p. 27196, 2020.

[11] A. Shocher, N. Cohen, and M. Irani, "Zero-shot super-resolution using deep internal learning," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 12 2017, pp. 3118–3126. [Online]. Available: https://ieeexplore.ieee.org/document/8578427

[12] T. R. Shaham, T. Dekel, and T. Michaeli, "Singan: Learning a generative model from a single natural image," *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2019-Octob, pp. 4569–4579, 2019.

[13] D. A. Scribner, K. A. Sarkady, M. R. Kruer, J. T. Caulfield, J. D. Hunt, and C. Herman, "Adaptive nonuniformity correction for IR focal-plane arrays using neural networks," in *Infrared Sensors: Detectors, Electronics, and Signal Processing*, T. S. J. Jayadev, Ed., vol. 1541, International Society for Optics and Photonics. SPIE, 1991, pp. 100 – 109. [Online]. Available: https://doi.org/10.1117/12.49324

[14] Y. Tendero and J. Gilles, "Admire: a locally adaptive single-image, non-uniformity correction and denoising algorithm: application to uncooled ir camera," *Infrared Technology and Applications*, vol. 8353, pp. 580–595, 5 2012. [Online]. Available: https://doi.org/10.1117/12.912966

[15] Y. Cao and C.-L. Tisse, "Single-image-based solution for optics temperature-dependent nonuniformity correction in an uncooled long-wave infrared camera," *Optics Letters*, vol. 39, no. 3, p. 646, Jan. 2014. [Online]. Available: https://doi.org/10.1364/ol.39.000646

[16] J. Zhao, Q. Zhou, Y. Chen, T. Liu, H. Feng, Z. Xu, and Q. Li, "Single image stripe nonuniformity correction with gradient-constrained optimization model for infrared focal plane arrays," *Optics Communications*, vol. 296, pp. 47–52, 2013.

[17] X. Jian, C. Lv, and R. Wang, "Nonuniformity correction of single infrared images based on deep filter neural network," *Symmetry*, 2018.

[18] Z. He, Y. Y. Cao, Y. Dong, J. Yang, and C.-L. Tisse, "Single-image-based nonuniformity correction of uncooled long-wave infrared detectors: a deep-learning approach," *Applied Optics*, vol. 57, p. D155, 2018.

[19] Y. Chang, L. Yan, L. Liu, H. Fang, and S. Zhong, "Infrared aerothermal nonuniform correction via deep multiscale residual network," *IEEE Geoscience and Remote Sensing Letters*, vol. 16, pp. 1120–1124, 2019, https://owuchangyuo.github.io/files/DMRN.rar¡br/¿Code.

[20] V. Saragadam, A. Dave, A. Veeraraghavan, and R. G. Baraniuk, "Thermal image processing via physics-inspired deep networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [Online]. Available: https://github.com/vishwa91/DeepIR

[21] N. Oz, N. Sochen, D. Mendelovich, and I. Klapp, "Improving temperature estimation in low-cost infrared cameras using deep neural networks," *Arxiv*, 2022.

[22] J. Harris and Y.-M. Chiang, "Nonuniformity correction of infrared image sequences using the constant-statistics constraint," *IEEE Transactions on Image Processing*, vol. 8, no. 8, pp. 1148–1151, 1999.

[23] R. C. Hardie, M. M. Hayat, E. Armstrong, and B. Yasuda, "Scene-based nonuniformity correction with video sequences and registration," *Appl. Opt.*, vol. 39, no. 8, pp. 1241–1250, Mar 2000. [Online]. Available: https://opg.optica.org/ao/abstract.cfm?URI=ao-39-8-1241

[24] E. Vera and S. Torres, "Fast adaptive nonuniformity correction for infrared focal-plane array detectors," *EURASIP Journal on Applied Signal Processing*, 2005.

[25] A. Averbuch, G. Liron, and B. Z. Bobrovsky, "Scene based non-uniformity correction in thermal images using kalman filter," *Image and Vision Computing*, vol. 25, pp. 833–851, 6 2007.

[26] C. Zuo, Q. Chen, G. Gu, and X. Sui, "Scene-based nonuniformity correction algorithm based on interframe registration," *J. Opt. Soc. Am. A*, vol. 28, no. 6, pp. 1164–1176, Jun 2011. [Online]. Available: https://opg.optica.org/josaa/abstract.cfm?URI=josaa-28-6-1164

[27] S. Papini, P. Yafin, I. Klapp, and N. Sochen, "Joint estimation of unknown radiometric data, gain, and offset from thermal images," *Applied Optics*, vol. 57, p. 10390, 2018.

[28] P. W. Kruse, *Uncooled thermal imaging arrays, systems, and applications*. SPIE, 2001.

[29] R. F. Voss, "Linearity of $\frac{1}{f}$ noise mechanisms," *Phys. Rev. Lett.*, vol. 40, pp. 913–916, Apr 1978. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.40.913

[30] M. d. Berg, O. Cheong, M. v. Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed. Santa Clara, CA, USA: Springer-Verlag TELOS, 2008.

[31] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*. Cambridge: Cambridge University Press, 2004.

[32] S. Farsiu, M. Robinson, M. Elad, and P. Milanfar, "Fast and robust multiframe super resolution," *IEEE Transactions on Image Processing*, vol. 13, no. 10, pp. 1327–1344, 2004.

[33] S. Kim and W.-Y. Su, "Recursive high-resolution reconstruction of blurred multiframe images," in *[Proceedings] ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*, 1991, pp. 2977–2980 vol.4.

[34] M. Zhang and B. K. Gunturk, "Multiresolution bilateral filtering for image denoising," *IEEE Transactions on Image Processing*, vol. 17, no. 12, pp. 2324–2333, 2008.

[35] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce, "Non-uniform deblurring for shaken images," *International Journal of Computer Vision*, vol. 98, pp. 168–186, 6 2012.

[36] D. DeTone, T. Malisiewicz, and A. Rabinovich, "Deep image homography estimation," 2016. [Online]. Available: https://arxiv.org/abs/1606.03798

[37] C. Godard, K. Matzen, and M. Uyttendaele, "Deep burst denoising," in *15th European Conference of Computer Vision ECCV*, 2018.

[38] G. Bhat, M. Danelljan, L. V. Gool, and R. Timofte, "Deep burst super-resolution," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, jun 2021, pp. 9205–9214. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR46437.2021.00909

[39] B. Wronski, I. Garcia-Dorado, M. Ernst, D. Kelly, M. Krainin, C.-K. Liang, M. Levoy, and P. Milanfar, "Handheld multi-frame super-resolution," *ACM Trans. Graph.*, vol. 38, no. 4, jul 2019. [Online]. Available: https://doi.org/10.1145/3306346.3323024

[40] M. Deudon, A. Kalaitzis, I. Goytom, M. R. Arefin, Z. Lin, K. Sankaran, V. Michalski, S. E. Kahou, J. Cornebise, and Y. Bengio, "Highres-net: Recursive fusion for multi-frame super-resolution of satellite imagery," *ArXiv*, vol. abs/2002.06460, 2020.

[41] B. De Brabandere, X. Jia, T. Tuytelaars, and L. Van Gool, "Dynamic filter networks," in *30th International Conference on Neural Information Processing Systems (NIPS)*, 2016.

[42] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.

[43] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Transactions on Computational Imaging*, vol. 3, no. 1, pp. 47–57, 2017.

[44] B. Mildenhall, J. T. Barron, J. Chen, D. Sharlet, R. Ng, and R. Carroll, "Burst denoising with kernel prediction networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[45] S. Anwar, S. Khan, and N. Barnes, "A deep journey into super-resolution: A survey," *ACM Comput. Surv.*, vol. 53, no. 3, may 2020. [Online]. Available: https://doi.org/10.1145/3390462

[46] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009.

[47] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," *NIPS*, 2017.

[48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: http://arxiv.org/abs/1412.6980

[49] A. M. Saxe, J. L. McClelland, and S. Ganguli, "Exact solutions to the nonlinear dynamics of learning in deep linear neural networks," in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, Y. Bengio and Y. LeCun, Eds., 2014. [Online]. Available: http://arxiv.org/abs/1312.6120

[50] D. Hendrycks and K. Gimpel, "Gaussian error linear units (gelus)," *arXiv preprint arXiv:1606.08415*, 2016.

[51] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 448–456. [Online]. Available: https://proceedings.mlr.press/v37/ioffe15.html

[52] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1874–1883, 2016.

## APPENDIX

A schematic diagram of the whole network is given in Fig. 4 and enlarged in Fig. 32. Bellow is a detailed description of the network architecture, from the UNET encoder-decoder part to the kernel estimation block and the offset block.

First we describe the UNET encoder-decoder part. We use a tensor with $N$ channels of grey-level frames as the input to the network. The input tensor undergoes a $3 \times 3$ convolution (conv) layer that encodes it from $N$ channels to $\mu$ channels without any activation function. The encoded features then pass through the encoder and decoder parts of the network, where the number of channels is multiplied by a factor of $\xi$ at each level. The encoder and decoder blocks consist of three $3 \times 3$ conv layers each. The first two layers in each block have GeLU [50] and batch normalization (norm) [51], while the last layer has neither activation nor norm. The last layer in each block produces $\mu \times \xi^i$ channels, where $i$ is the level index. The encoder block also applies an average pooling layer with a $\xi \times \xi$ window and stride $\xi$ to reduce the spatial resolution, while the decoder block uses a pixel shuffle layer [52] with an upsample factor of $\xi$ to increase it. We concatenate the encoder block output before pooling with the pixel shuffle output at each level before feeding it to the decoder block. The encoder and decoder block structures are shown in Table I.

After the last decoder block in the UNET, we add a kernel estimation block that generates $N \times K \times K$ channels using three $1 \times 1$ conv layers. The first two layers have GeLU, and the last layer has no activation. The kernel estimation block structure is shown in Table II. We reshape the output of this block as kernels of size $K \times K$ for each frame. We then sample a patch of size $K \times K$ around each pixel in each frame and compute the inner product of the corresponding kernel and patch, as in the first term of Eq. (11). We sum the inner products from all the frames for each pixel.

To map the temperature estimation to the camera range, we use an offset block that takes the means of all the input gray level frames as input and outputs a single scalar. The offset block is a fully-connected layer that acts as a polynomial function of the input. The offset block is explained in more detail in Sec. IV-A.

The final temperature estimation is obtained by adding the offset scalar to the pixel-wise summation of the gain from the kernel estimation block.

The scale factor for decoder and encoder blocks is $\xi \equiv 2$ and the number of channels is $\mu \equiv 64$ throughout the work. The number of levels was empirically set to $4$.

Fig. 13 shows more results of the proposed method on real data.

Figures 14 to 19 compares the results of the proposed method to ADMIRE [14], DeepIR [20] and He [18] methods. Figures 18 and 19 specifically show the hallucination effect of DeepIR [20] method.

Figures 20 to 24 displays the absolute error per-pixel as a function of number of frames, both in quantitatively and qualitatively.

Figures 25 to 30 are the original images used for the real data results in Fig. 11. On the left of each figure is the GT temperature map acquired by the `A655sc`, and on the right is the estimated temperature map by the proposed method. The raw data cannot be displayed because it consists of 7 frames.

Fig. 31 shows the UAV used for the real data experiments.

Fig. 32 is an enlarged version of Fig. 4.

Table III specifies the parameters of `Tau2` that were used throughout all the experiments.

TABLE I
ARCHITECTURE OF THE ENCODER AND DECODER BLOCKS.

| **Type** | **BN** [51] | **GeLU** [50] | **Kernel** | **Output** |
|---|---|---|---|---|
| Conv2D | $\surd$ | $\surd$ | 3 | $\mu \times h \times w$ |
| Conv2D | $\surd$ | $\surd$ | 3 | $\mu \times h \times w$ |
| Conv2D | $\times$ | $\times$ | 3 | $(\xi \times \mu) \times h \times w$ |

TABLE II
ARCHITECTURE OF THE KERNEL PREDICTOR BLOCK.

| **Type** | **GeLU** [50] | **Kernel** | **Output** |
|---|---|---|---|
| Conv2D | $\surd$ | 1 | $\mu \times h \times w$ |
| Conv2D | $\surd$ | 1 | $\mu \times h \times w$ |
| Conv2D | $\times$ | 1 | $(N \times K^2) \times h \times w$ |

TABLE III
THE FLIR TAU2 SETTINGS AS DESCRIBED IN TAU2 QUARK SOFTWARE IDD

| Function | State | Function | State |
|---|---|---|---|
| FFC Mode | Auto | FPS | 4 ($60_{Hz}$) |
| FFC Period | 0 | CMOS Depth | 0 ($14_{bit}$ w/o AGC) |
| Isotherm | 0 | LVDS | 0 |
| DDE | 0 | LVDS Depth | 0 ($14_{bit}$) |
| T-Linear | 0 | XP | 2 ($14_{bit}$) |
| AGC | Manual | Brightness Bias | 0 |
| Contrast | 0 | Brightness | 0 |
| ACE | 0 | SSO | 0 |
| Gain | High | | |

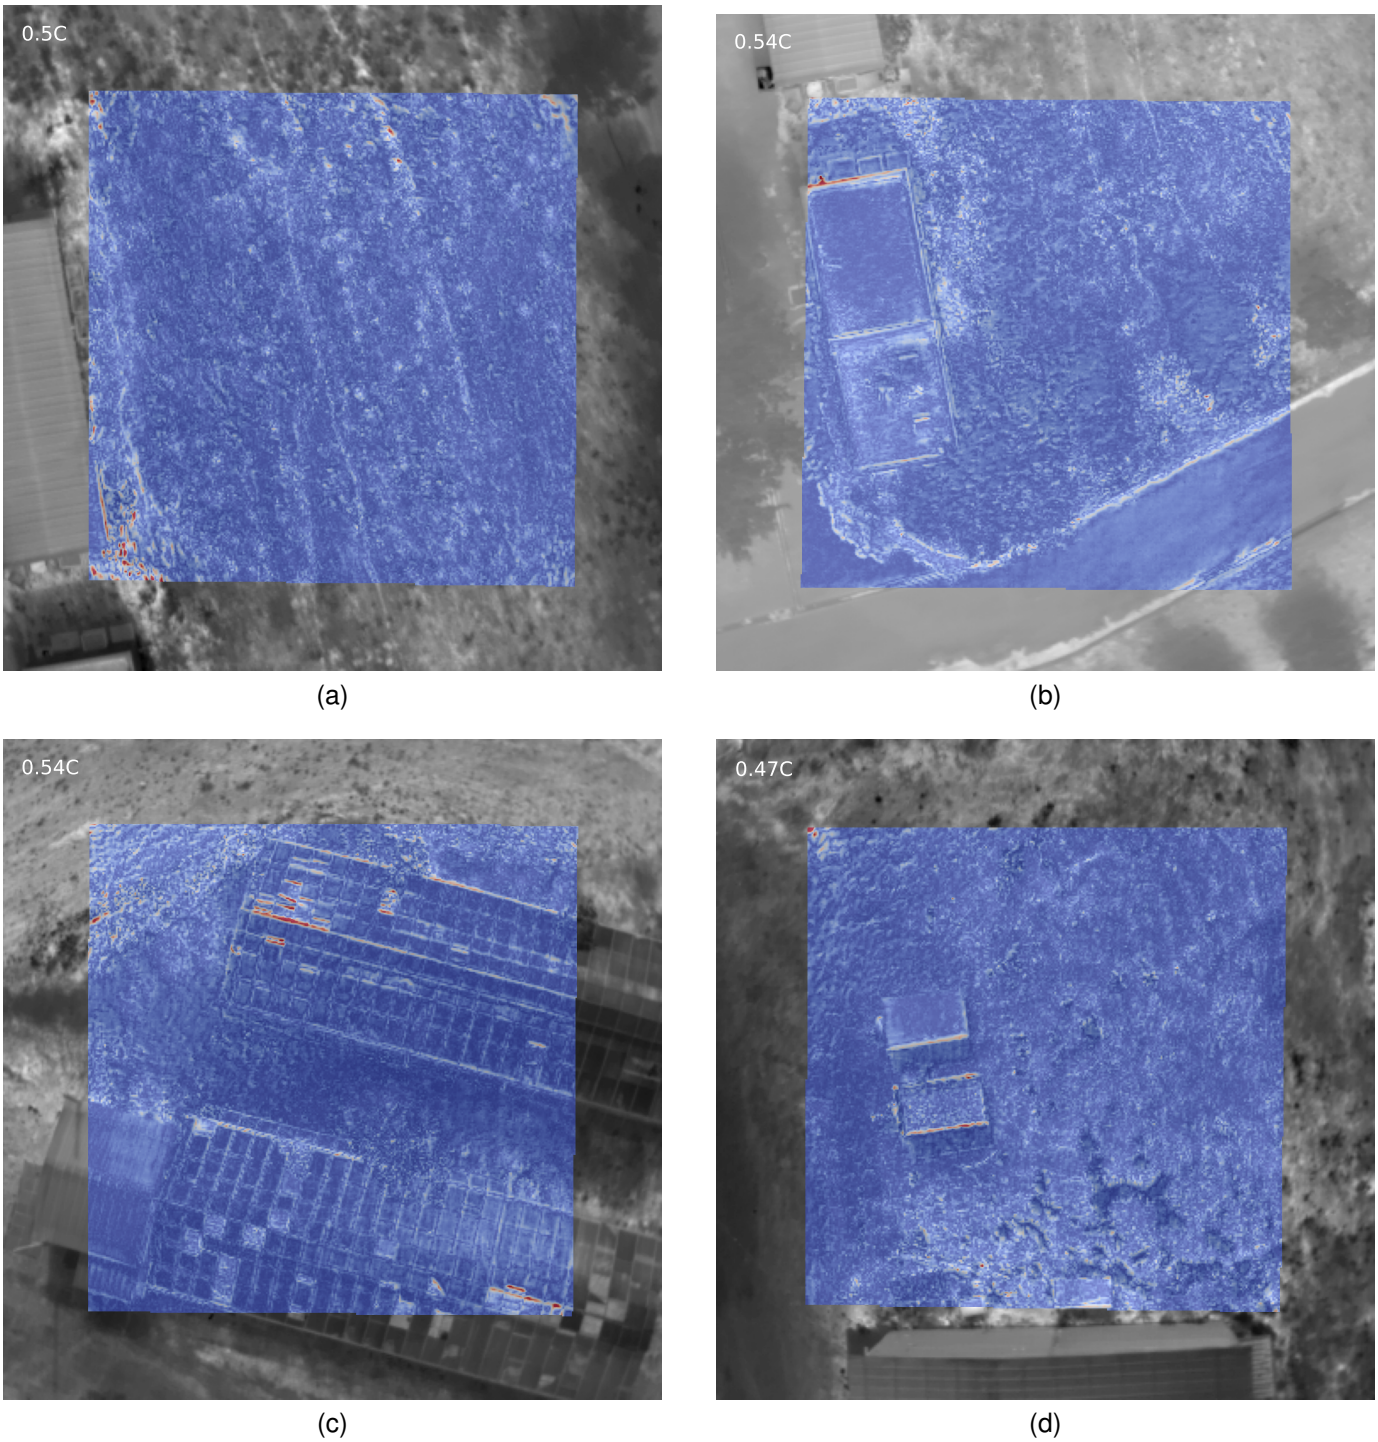


Fig. 13. Various results on real data. The gray background is the GT temperature map and the colored map is the difference between the GT and estimated temperature maps. The number on the top-left corner is the MAE between the estimated and GT temperature maps.

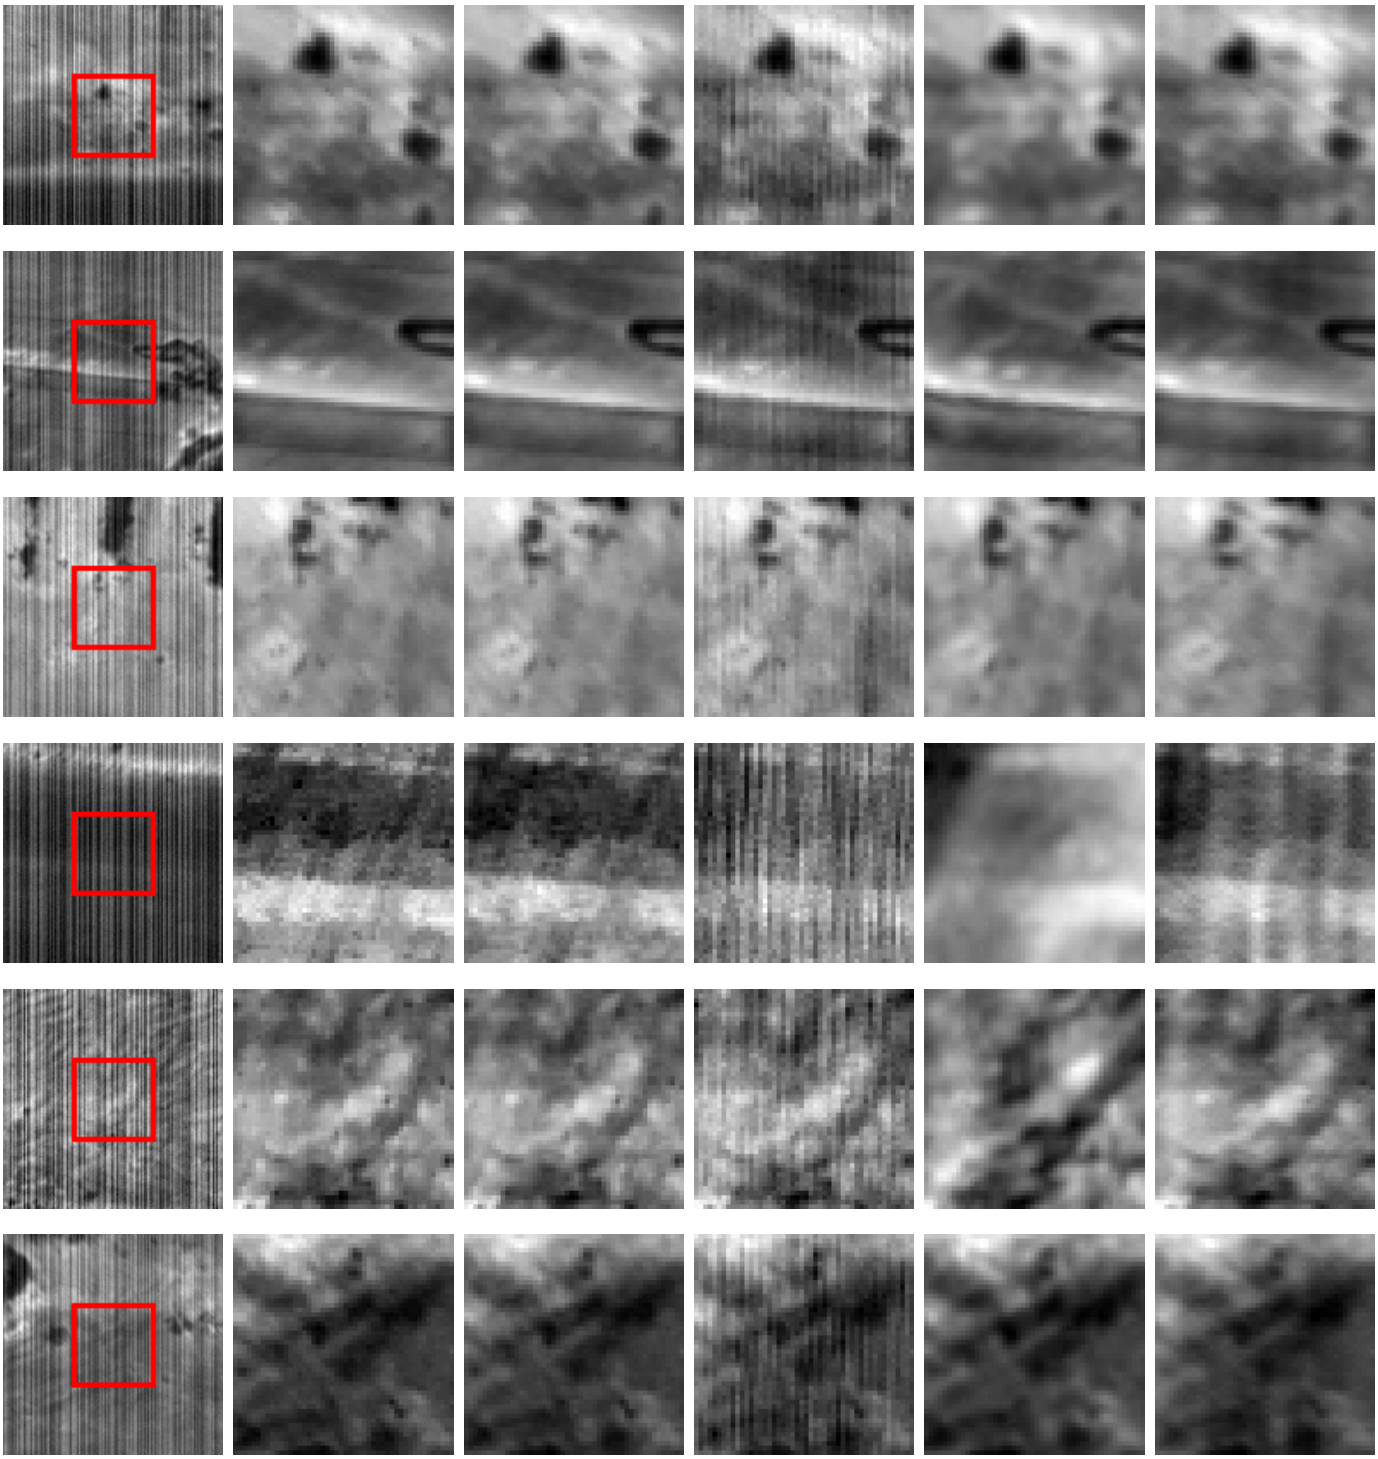


Fig. 14. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results.

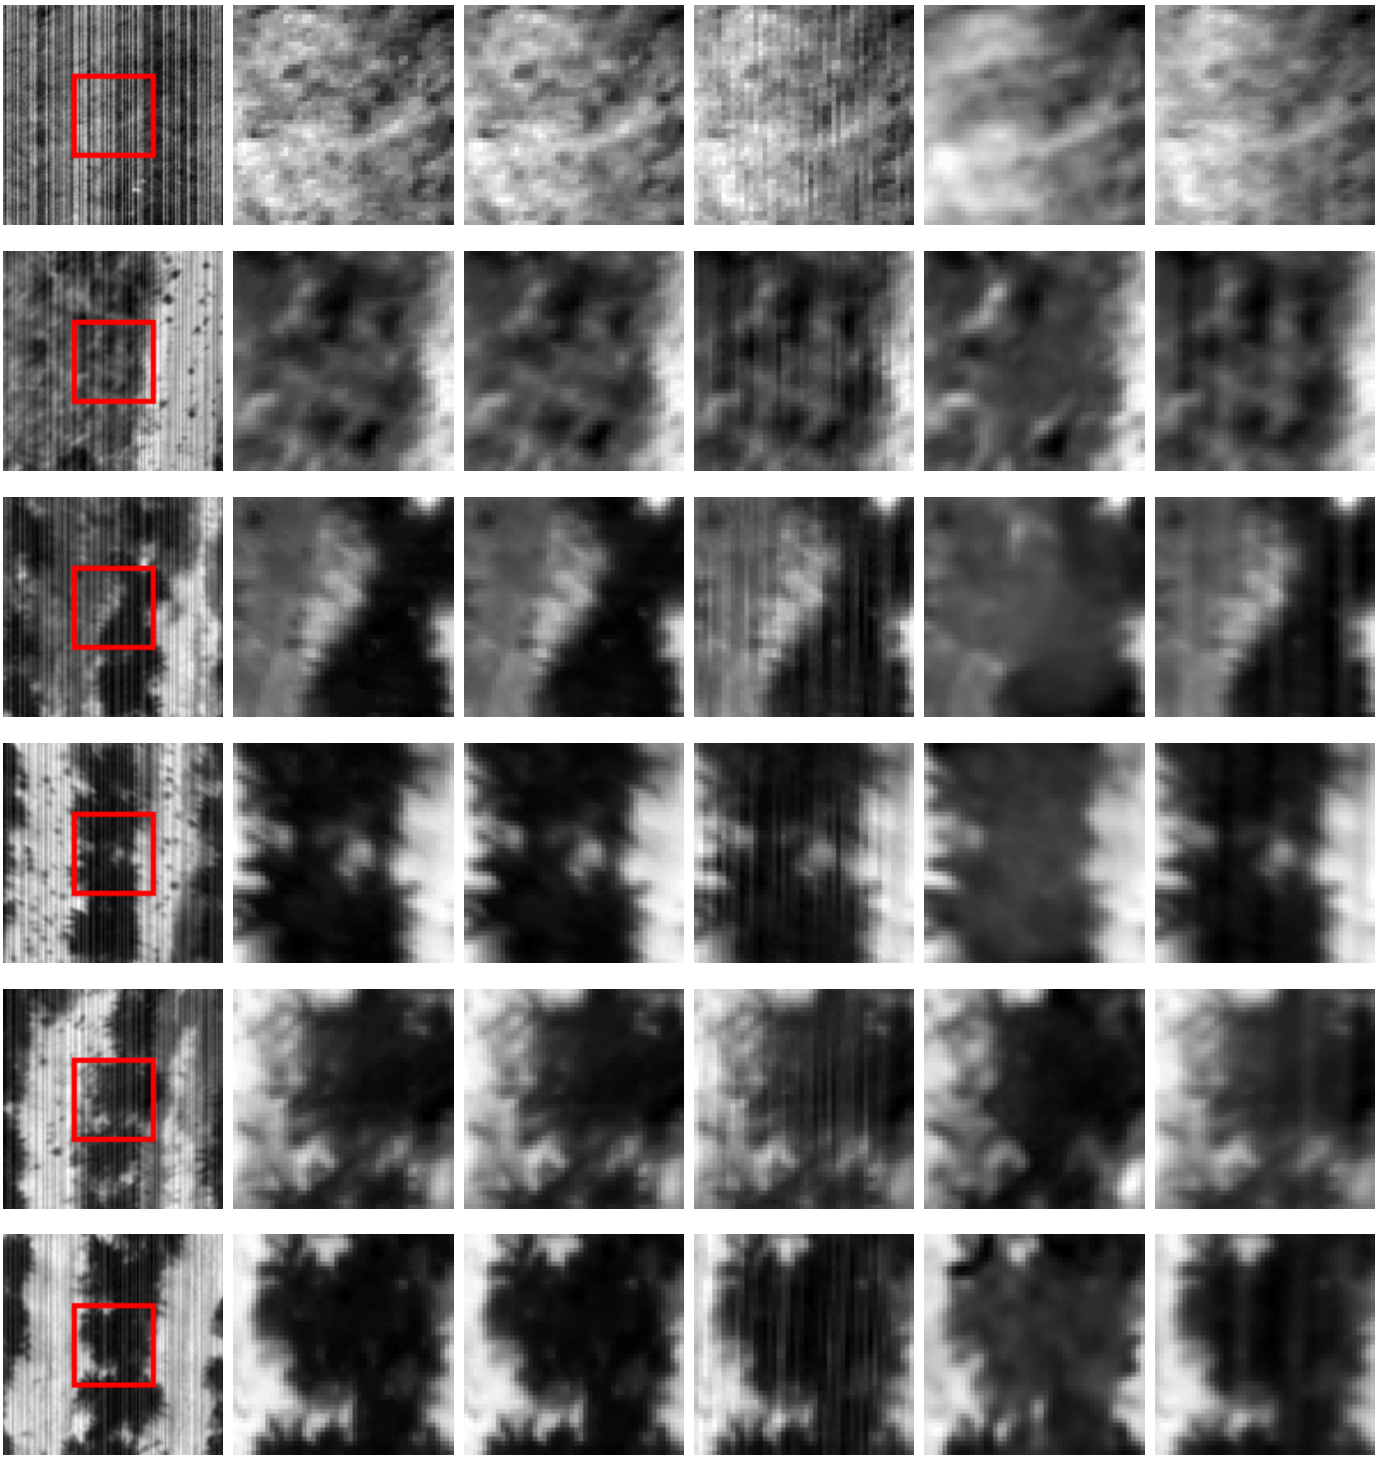


Fig. 15. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results.

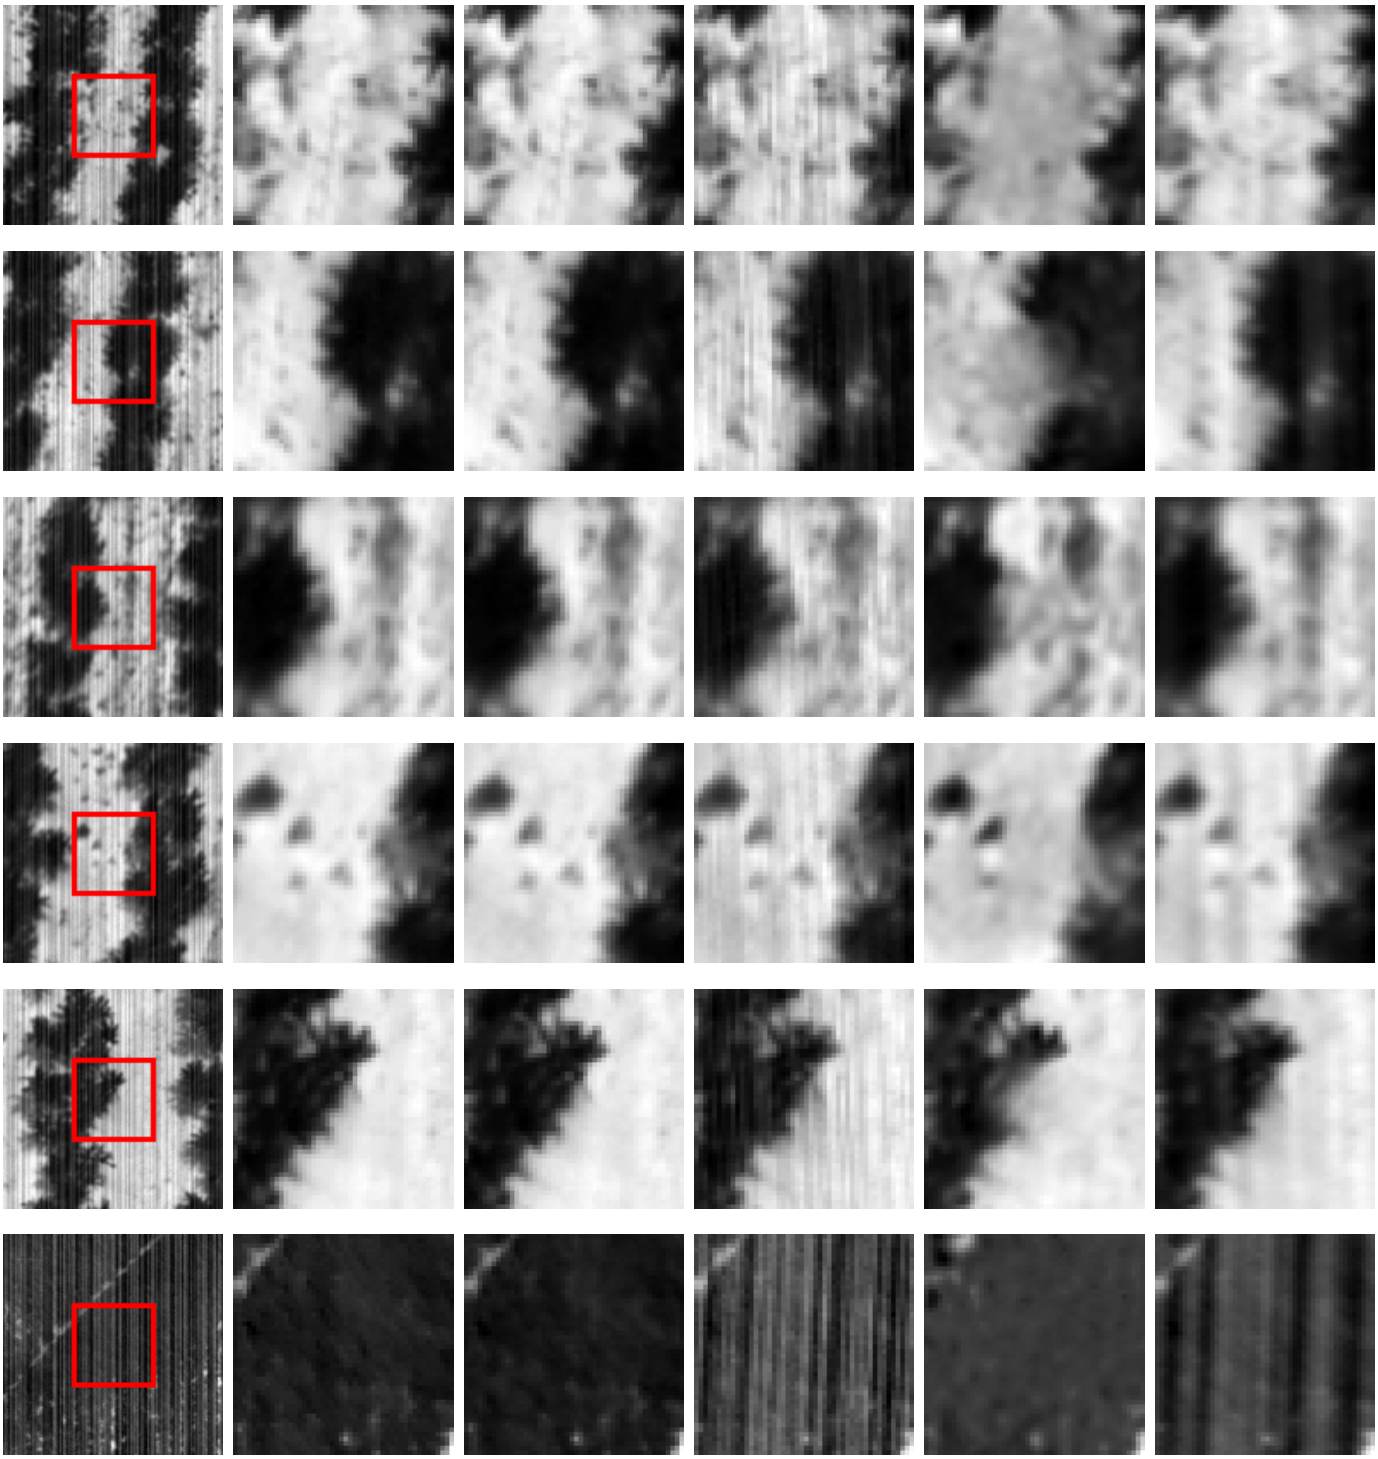


Fig. 16. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results.

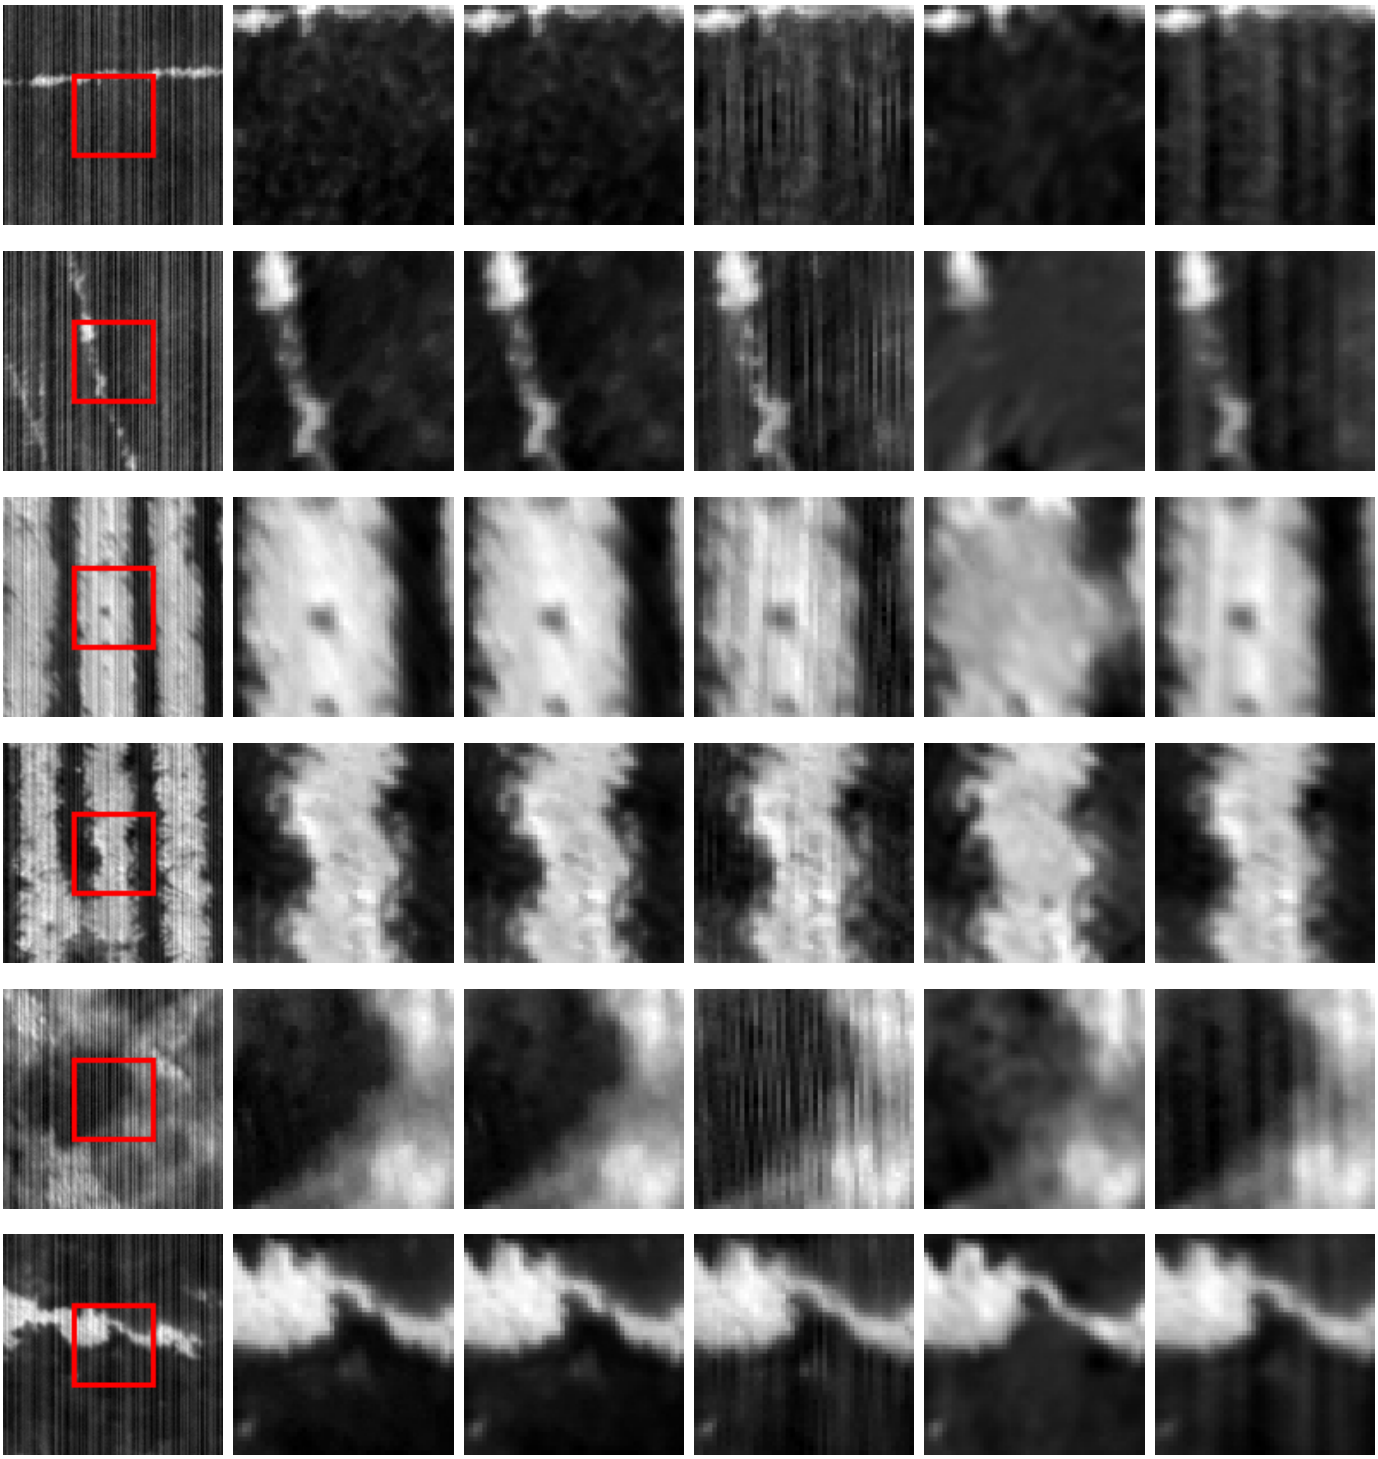


Fig. 17. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results.

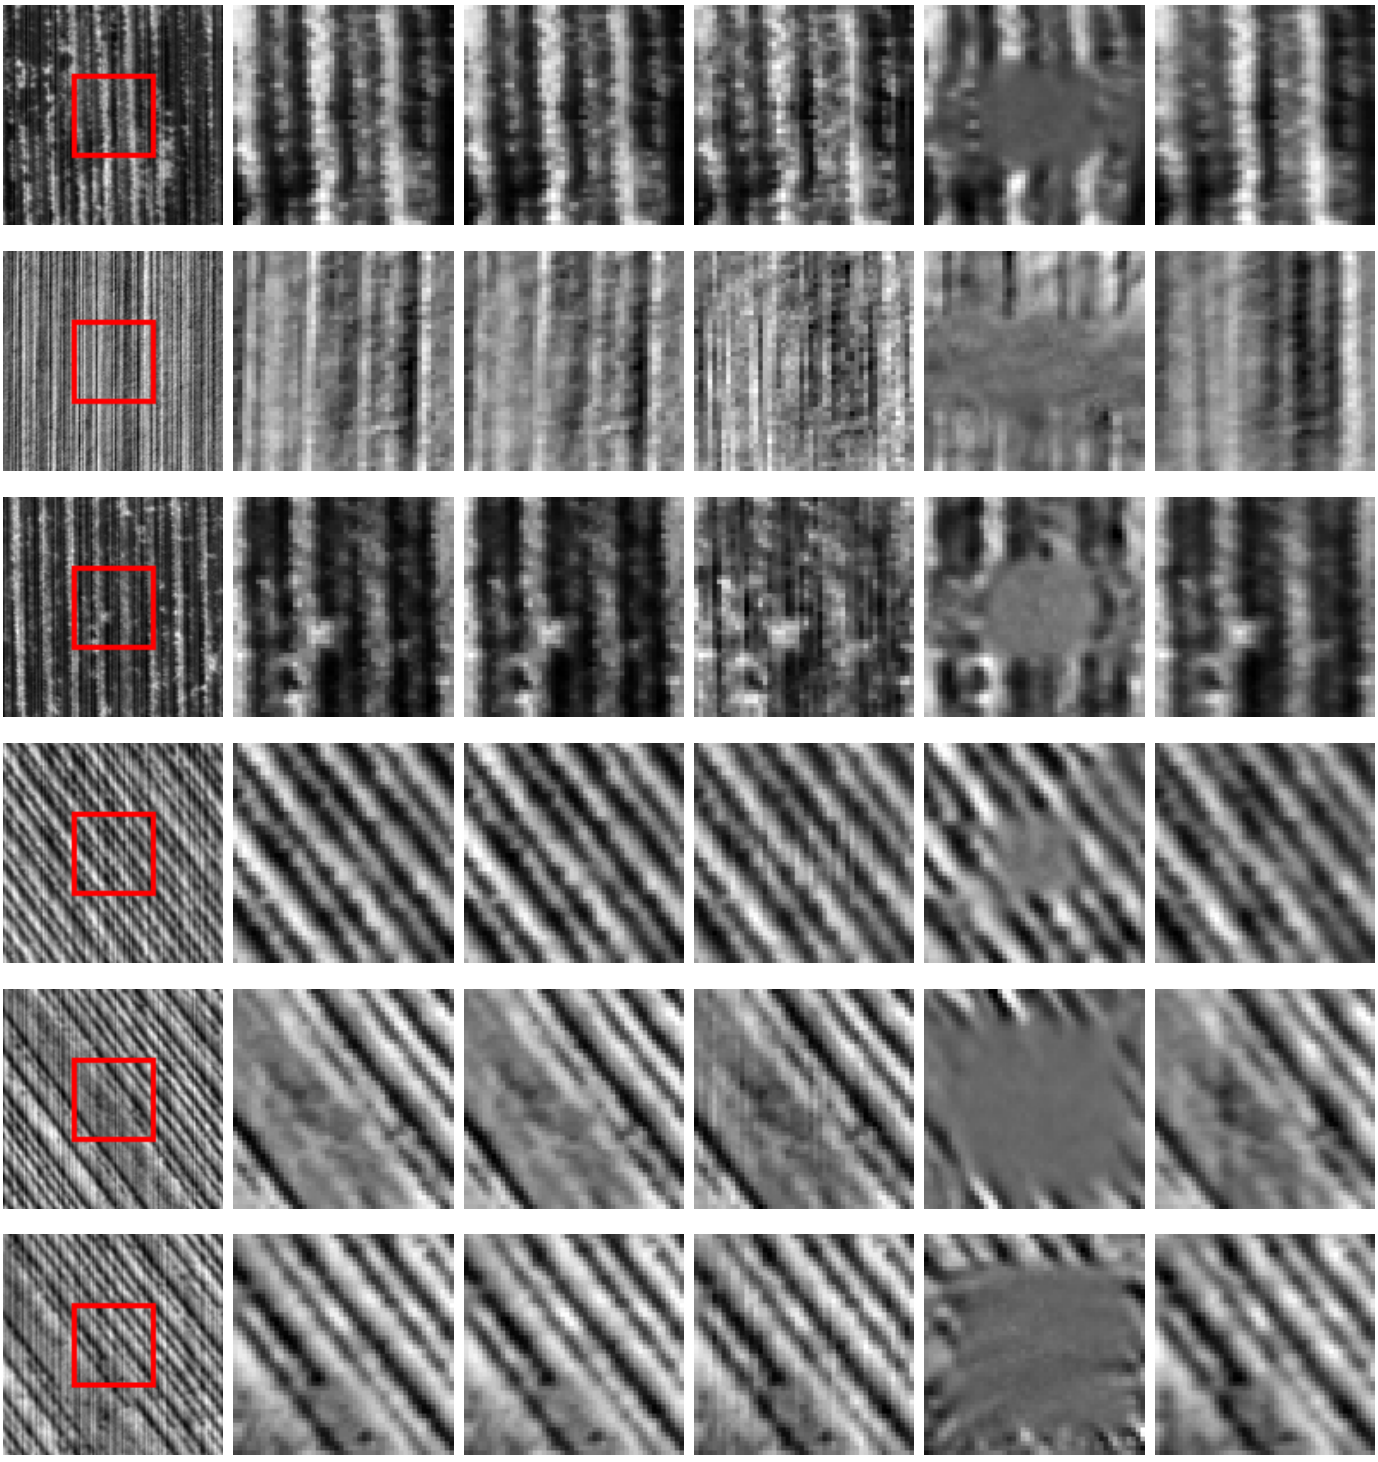


Fig. 18. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results. These results displays the hallucination effect of DeepIR [20] method.

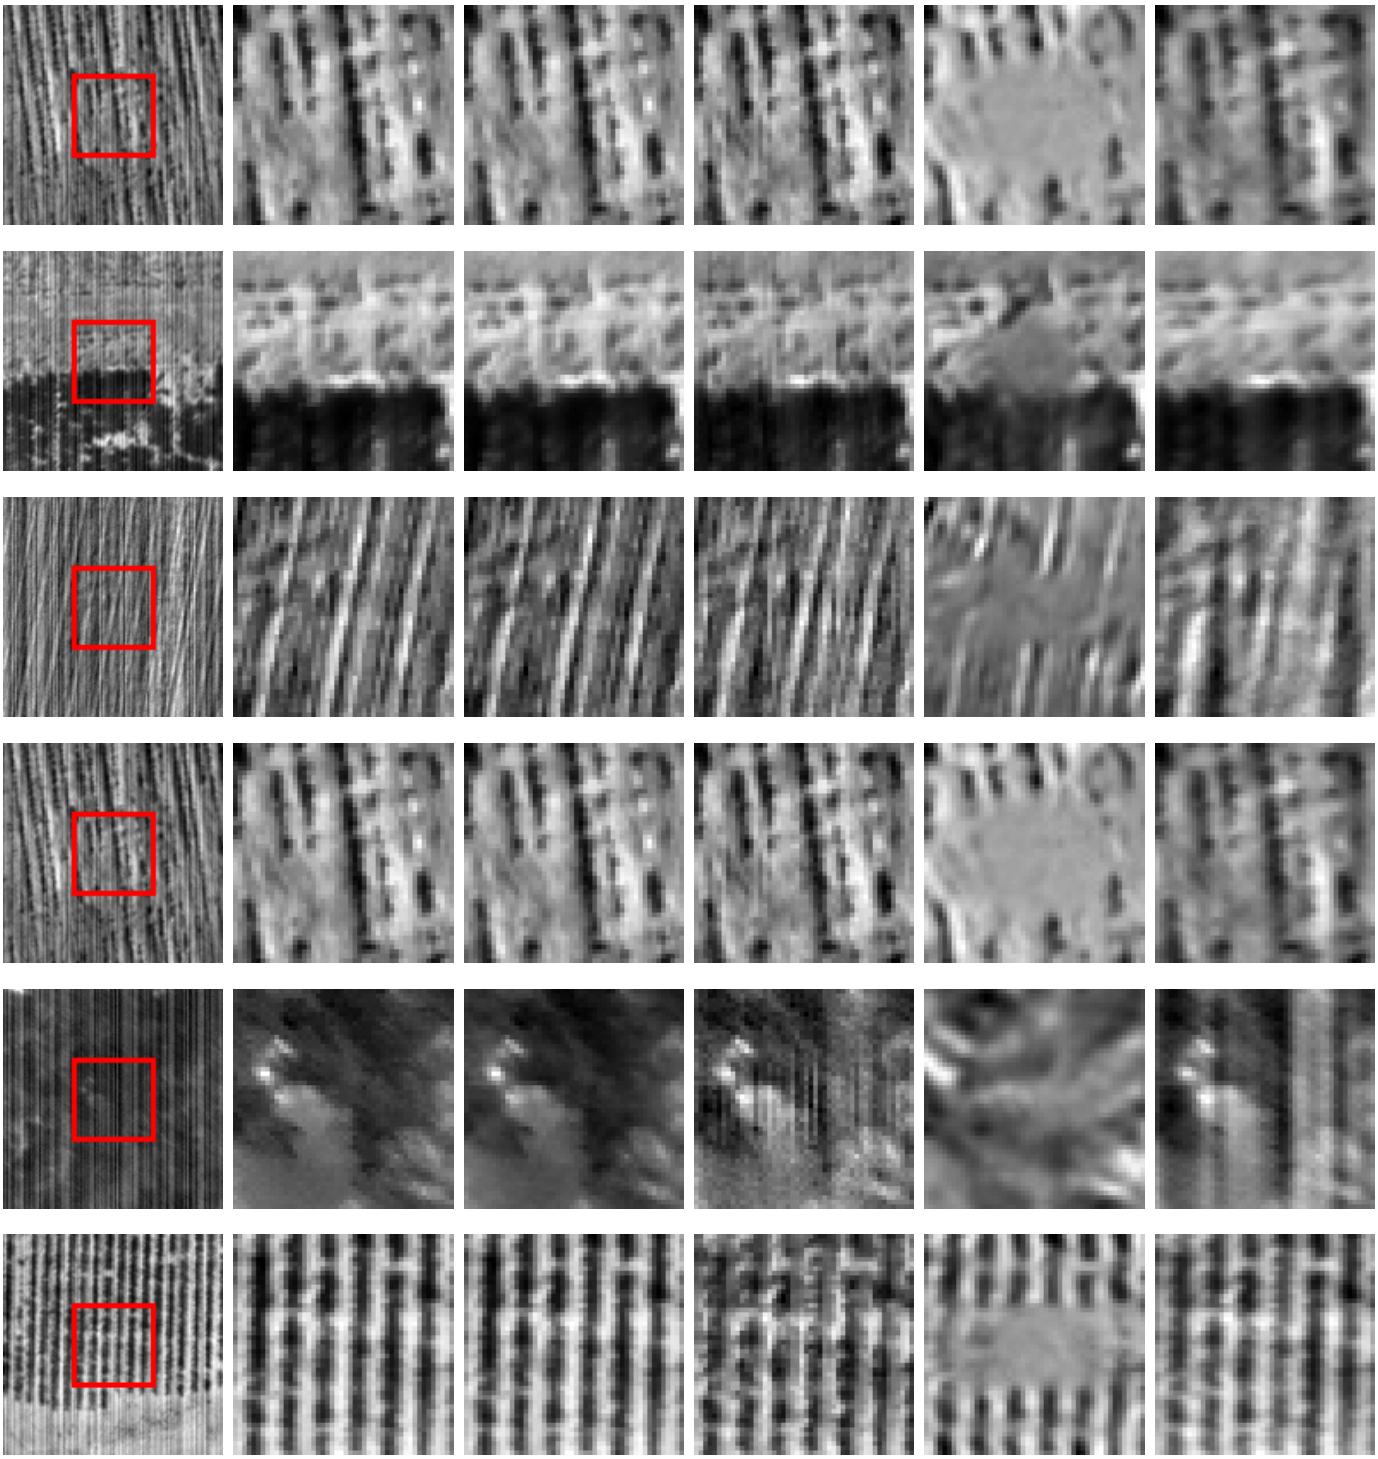


Fig. 19. Zoomed-in results of different methods. The left-most figure is the reference frame with a red rectangle. The following figures are the results of the area inside the red rectangle. $N = 11$ for all results. These results displays the hallucination effect of DeepIR [20] method.

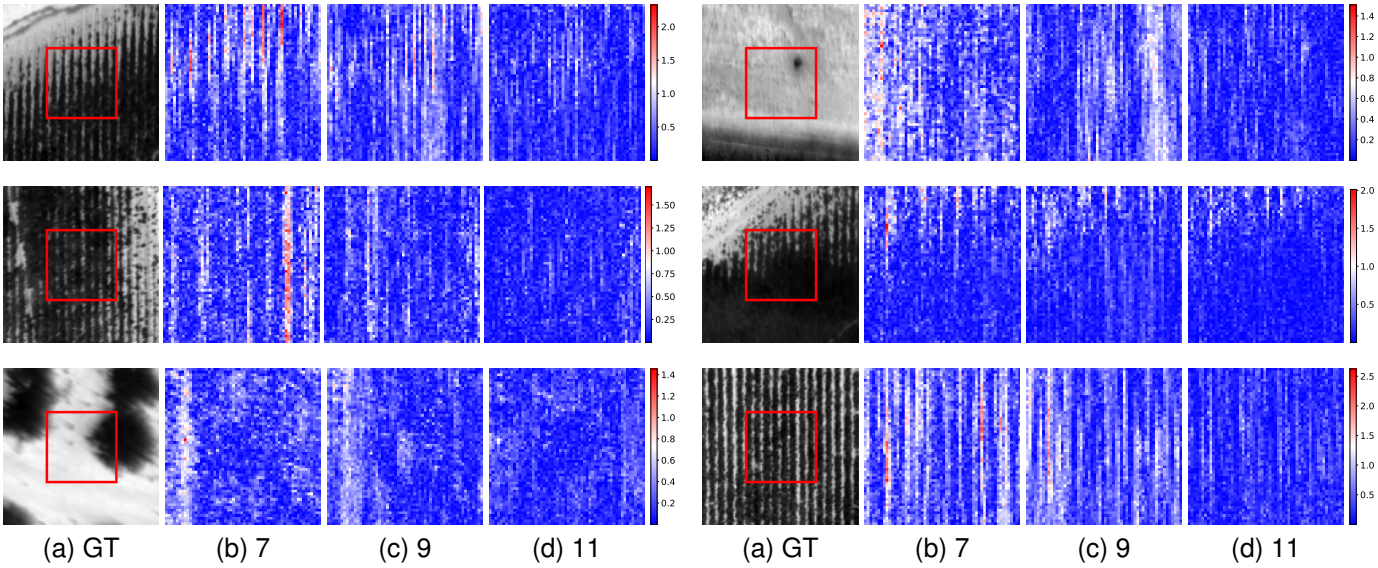


Fig. 20. Difference between the temperature estimation with our method and the ground truth. The left-most figure is the ground truth. The following figures are the zoom-in of the area inside the red rectangle. The number bellow the difference map is the number of frames used for the temperature estimation, from left to right 7, 9 and 11 frames.

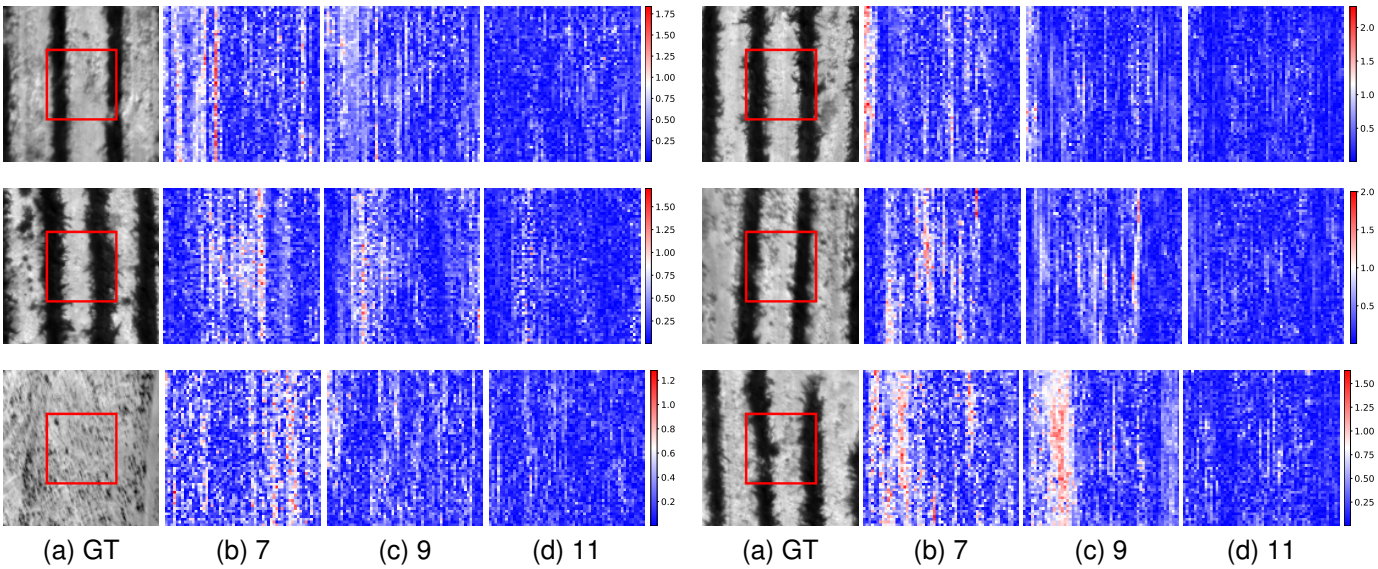


Fig. 21. Difference between the temperature estimation with our method and the ground truth. The left-most figure is the ground truth. The following figures are the zoom-in of the area inside the red rectangle. The number bellow the difference map is the number of frames used for the temperature estimation, from left to right 7, 9 and 11 frames.

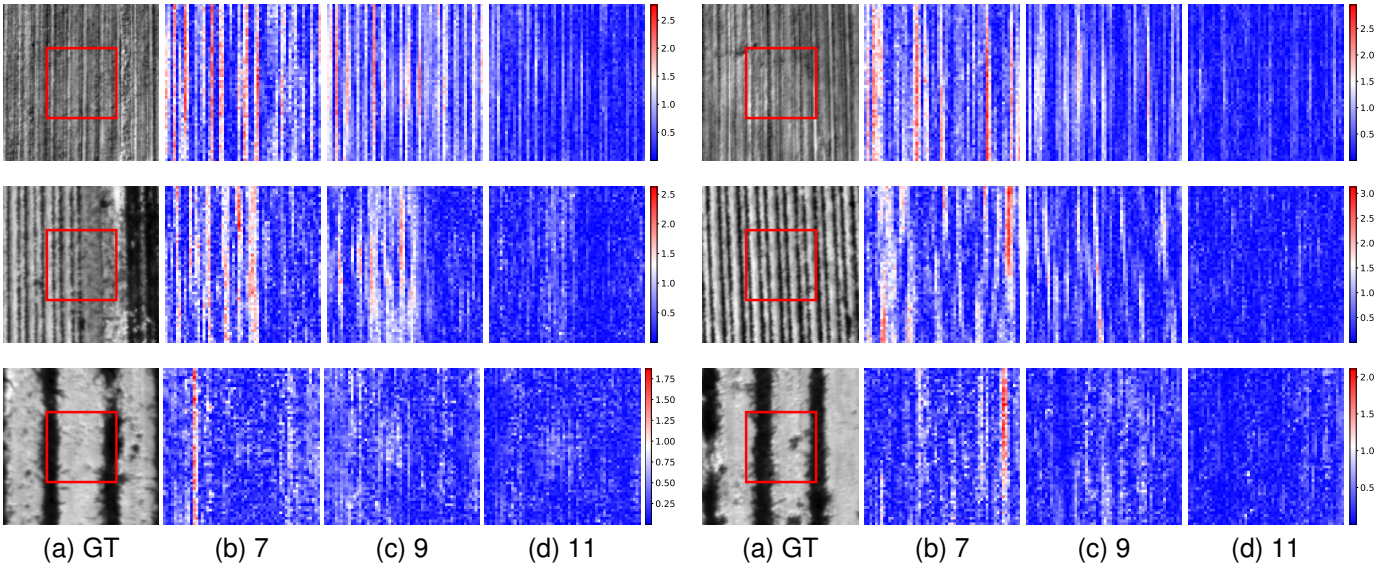


Fig. 22. Difference between the temperature estimation with our method and the ground truth. The left-most figure is the ground truth. The following figures are the zoom-in of the area inside the red rectangle. The number bellow the difference map is the number of frames used for the temperature estimation, from left to right 7, 9 and 11 frames.

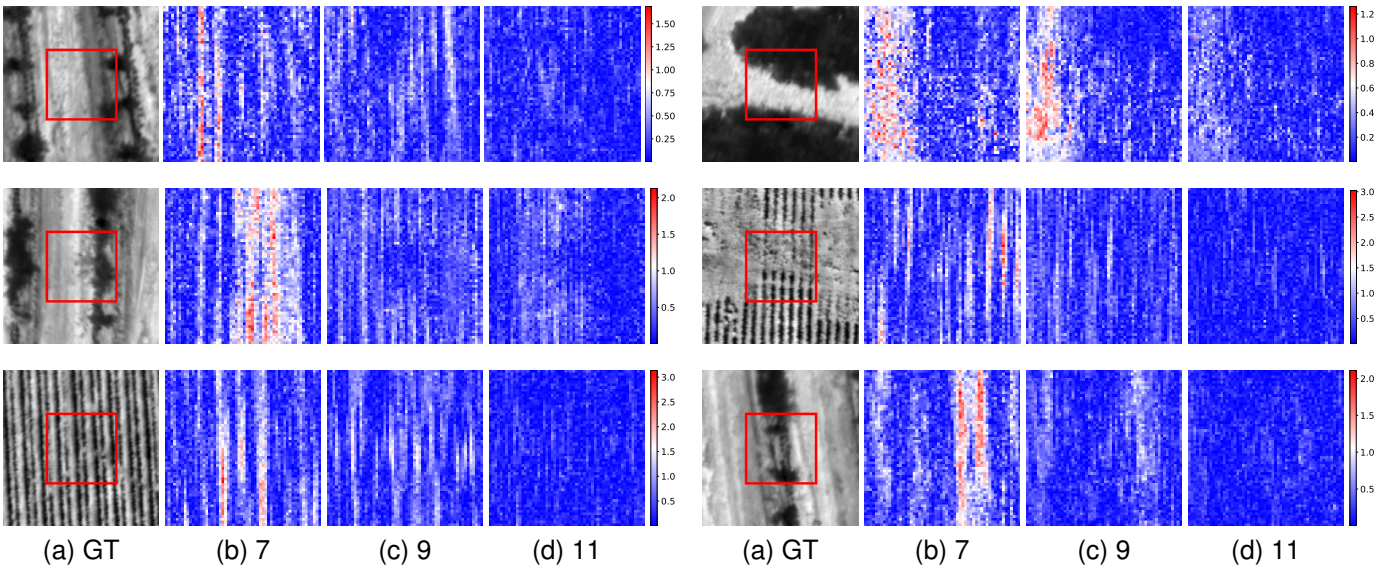


Fig. 23. Difference between the temperature estimation with our method and the ground truth. The left-most figure is the ground truth. The following figures are the zoom-in of the area inside the red rectangle. The number bellow the difference map is the number of frames used for the temperature estimation, from left to right 7, 9 and 11 frames.

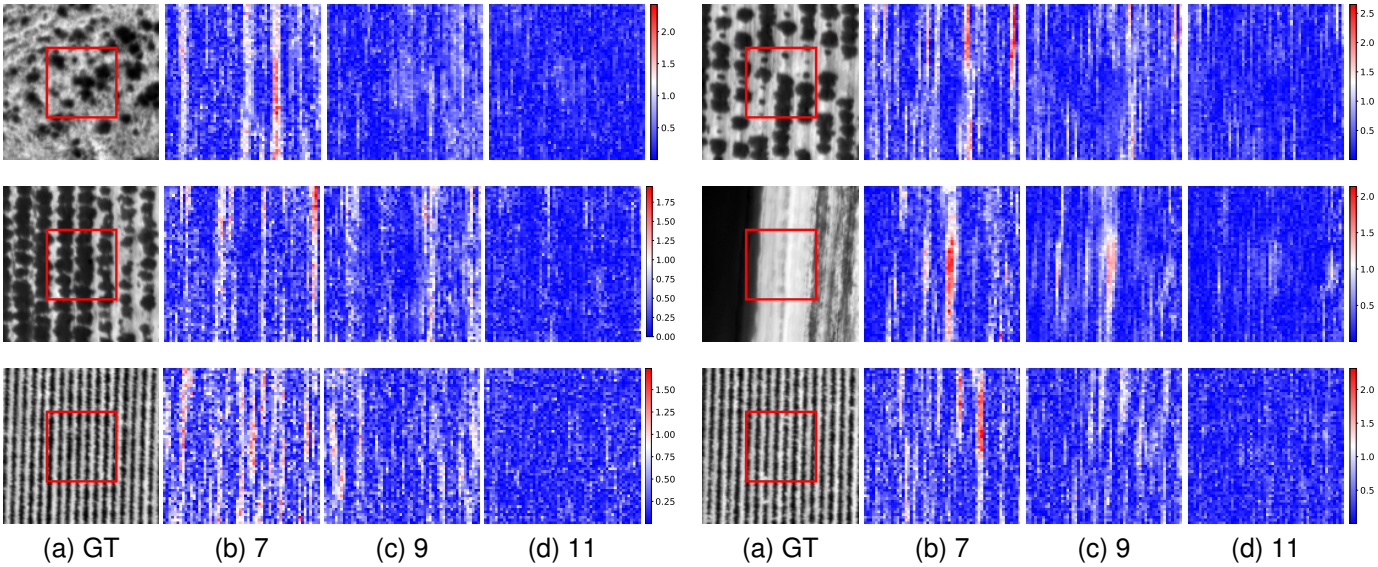


Fig. 24. Difference between the temperature estimation with our method and the ground truth. The left-most figure is the ground truth. The following figures are the zoom-in of the area inside the red rectangle. The number bellow the difference map is the number of frames used for the temperature estimation, from left to right 7, 9 and 11 frames.

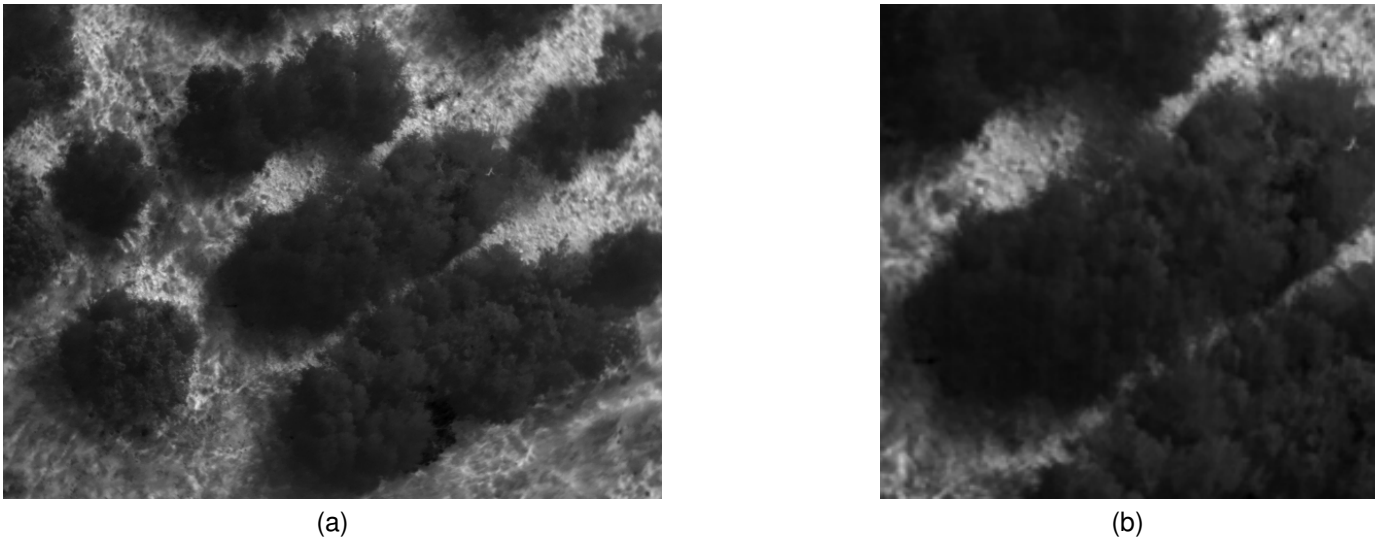


Fig. 25. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (a).

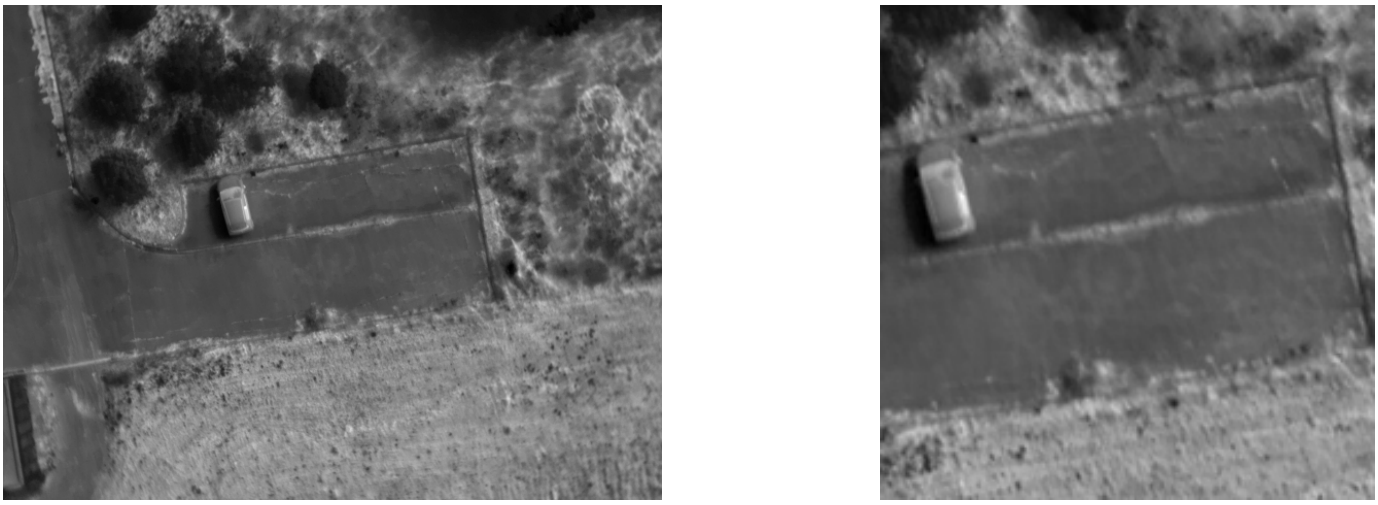

(a) (b)

Fig. 26. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (b).

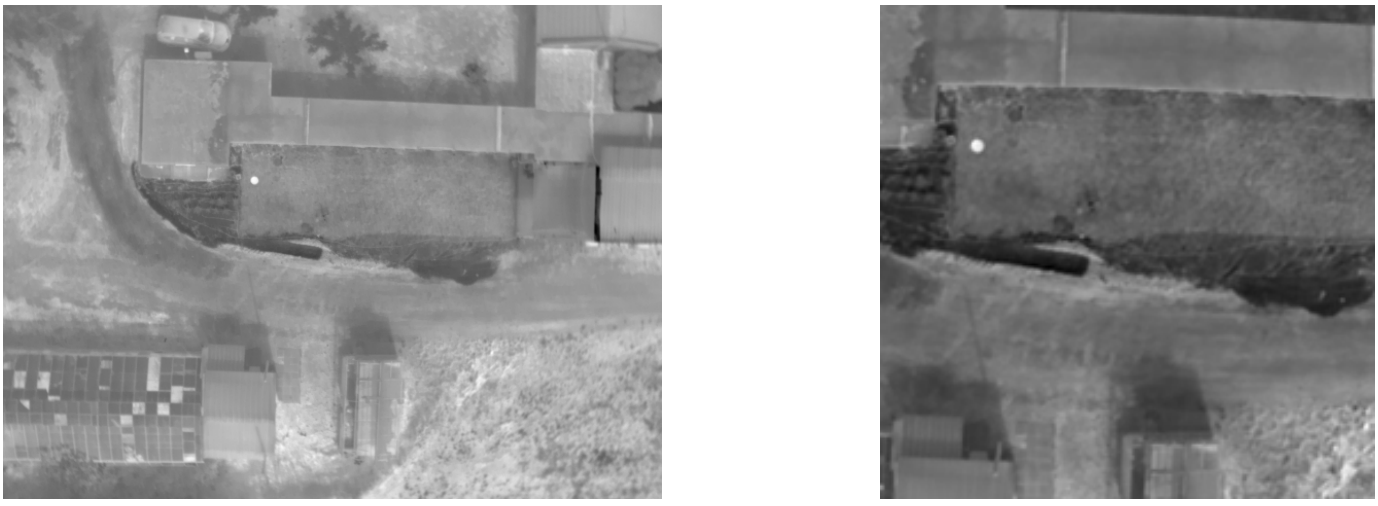

(a) (b)

Fig. 27. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (c).

(a)

(b)

Fig. 28. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (d).

(a)

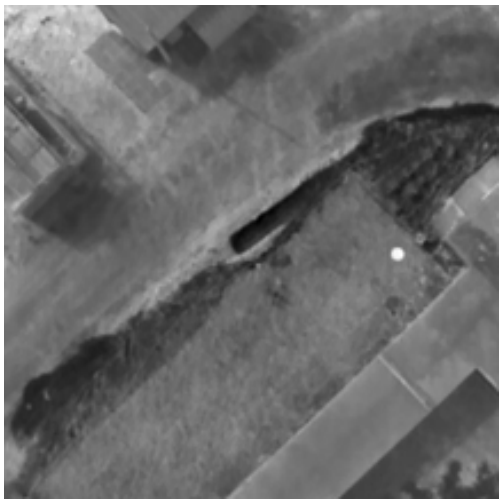

(b)

Fig. 29. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (e).

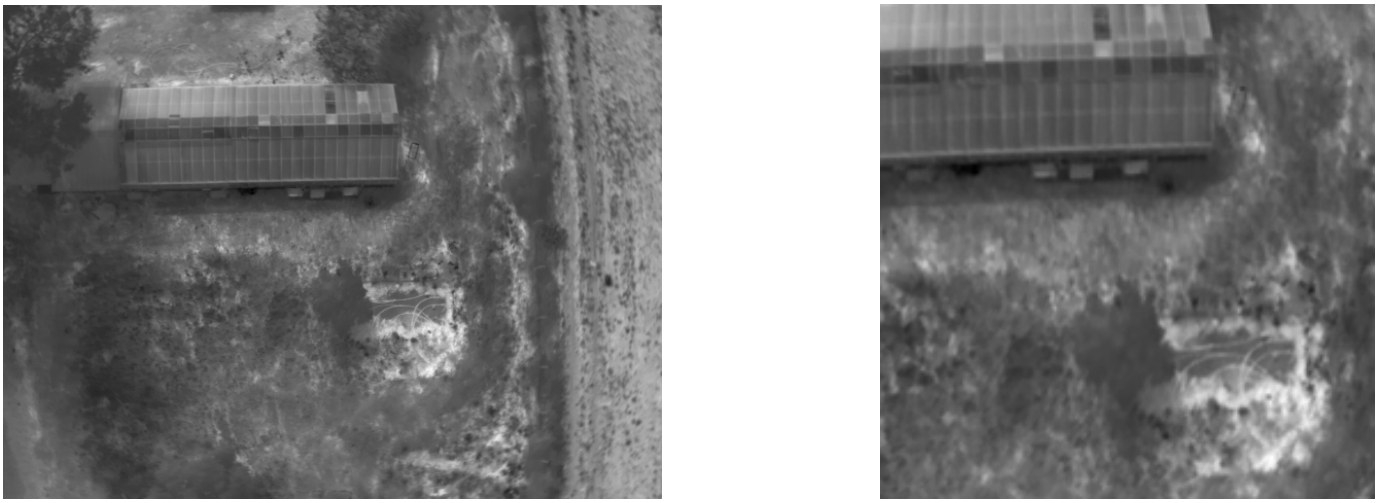

(a) (b)

Fig. 30. Ground truth (left) and estimated (right) temperature maps for the result in Fig. 11 (f).

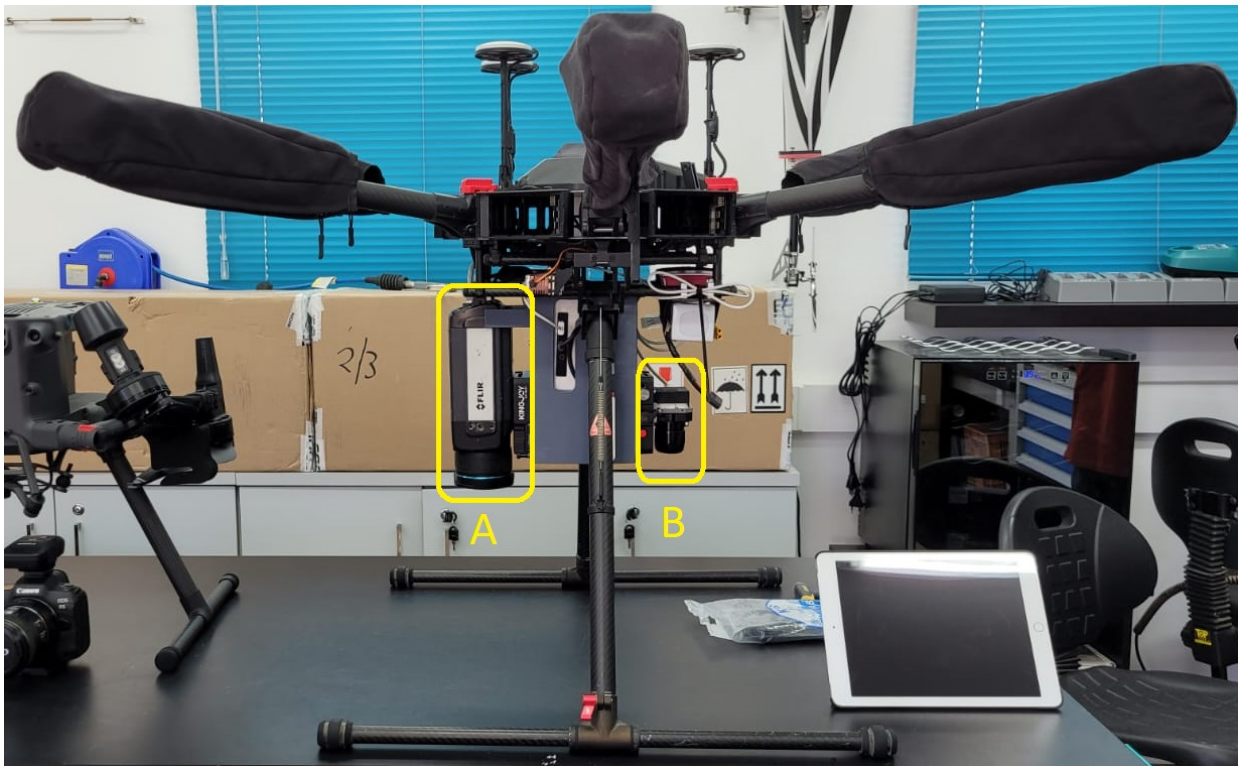


Fig. 31. UAV used for the collection of real data. `A655sc` is mounted on the left (marked A) and `Tau2` is mounted on the right (marked B).

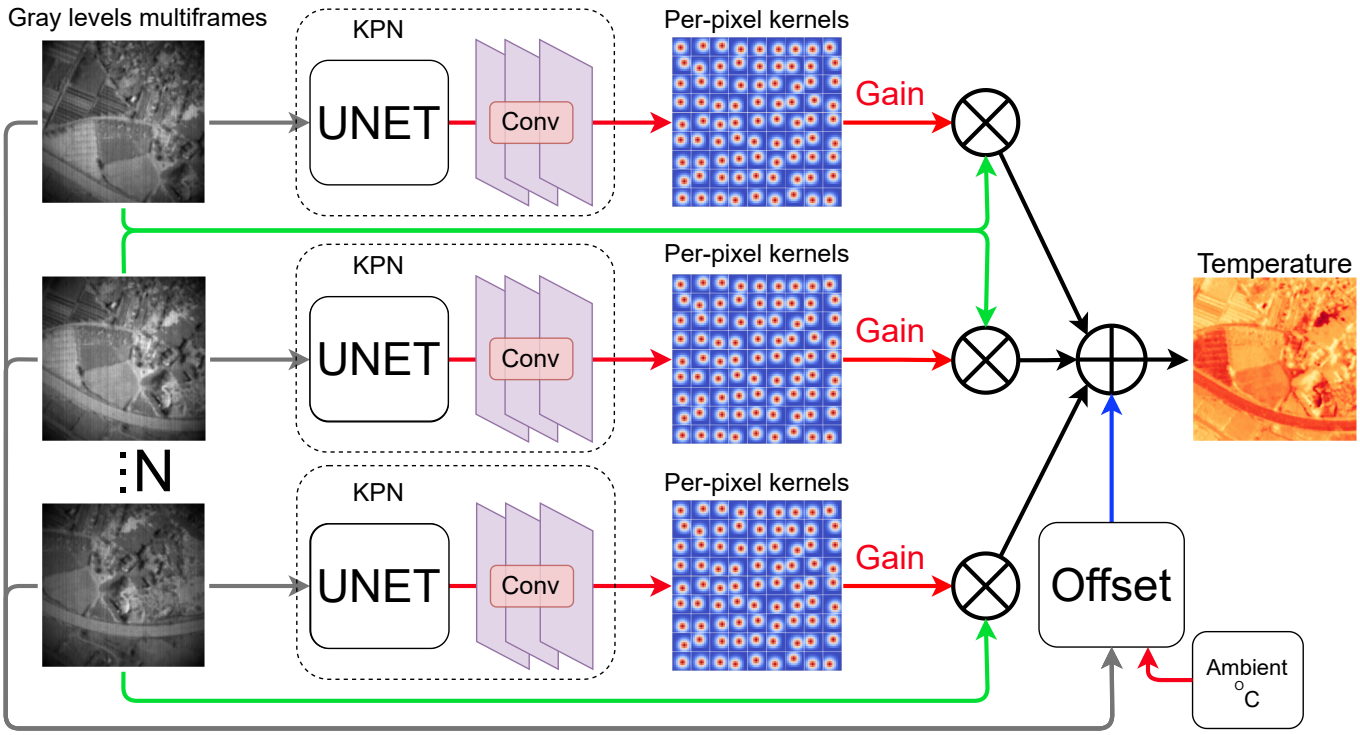


Fig. 32. Schematics of the model. The gray-level multiframes are fed into the kernel prediction network (KPN), and the KPN outputs the per-pixel kernels for each frame. Each frame is divided to overlapping patches with the same support as the kernels. The patches and the kernels are multiplied element-wise and each product is summed, resulting in a 2D gain map for each frame. All the 2D gain maps are summed depth-wise, resulting in a single 2D map. The offset, a single scalar value, is added to the single 2D map to get the estimated temperature map.