

# PROCEEDINGS OF SPIE

[SPIDigitalLibrary.org/conference-proceedings-of-spie](https://SPIDigitalLibrary.org/conference-proceedings-of-spie)

## Three-dimensional signal decomposition of infrared image sequences and large-scale non-uniformity analysis

Thomassen, Jan, van Rheenen, Arthur

Jan B. Thomassen, Arthur D. van Rheenen, "Three-dimensional signal decomposition of infrared image sequences and large-scale non-uniformity analysis," Proc. SPIE 11536, Target and Background Signatures VI, 115360F (20 September 2020); doi: 10.1117/12.2574573

**SPIE.**

Event: SPIE Security + Defence, 2020, Online Only

# Three-dimensional signal decomposition of infrared image sequences and large-scale non-uniformity analysis

Jan B. Thomassen\* and Arthur D. van Rheenen

Norwegian Defence Research Establishment (FFI), PO Box 25, N-2027 Kjeller, Norway

## ABSTRACT

The three-dimensional noise model is a methodology to analyse the noise of a thermal imaging sensor, such as an infrared (IR) camera. This allows us to decompose a noisy signal into components and quantify properties such as noise equivalent temperature difference (NETD), temporal noise, rain, streaks, or various types of fixed pattern noise. As part of this analysis, it is necessary to identify trends in order to split the data into signal and noise. In this paper we discuss methods to perform this split. We then show that not only the noise, but also the trends contain interesting information and can be used to quantify large-scale non-uniformities in calibrated IR images. We apply this analysis to investigate three different effects that may appear in recorded data: How does the uniformity of the background change when we vary the temperature, the distance, or the lens focus? We have performed a series of laboratory measurements on blackbodies in order to investigate these effects. We find that large-scale non-uniformity may be present even in calibrated images, with an order of magnitude up to  $\Delta T \sim 0.6$  K.

**Keywords:** Thermal imaging sensors, camera calibration, signal analysis, three-dimensional noise model, noise characterisation, infrared cameras

## 1. INTRODUCTION

The three-dimensional noise model was introduced in Ref. 1 as a method to analyse the noise of a thermal imaging sensor, such as an infrared (IR) camera. This allows us to decompose a noisy signal into components and quantify properties such as NETD, temporal noise, rain, streaks, or various types of fixed pattern noise, see e.g. Sec. 12.8. in Ref. 2.

In its basic form, the analysis in Ref. 1 assumes a uniform signal in time and space (focal plane). The various noise components are isolated by taking averages in each direction. This decomposes the signal into averages and noise. As pointed out in Ref. 3, if a trend is present in the signal, for instance due to the camera optics, this trend must be identified and removed in order to isolate the noise from the average background.

In this paper we explore this signal decomposition technique further. The large-scale trends that are present in a three-dimensional signal contains interesting information, even for a recording of a uniformly radiating background, such as a calibration blackbody. Typically an IR camera is calibrated with a two-point calibration using two temperature controlled blackbodies at close distance, with lenses focused at large distance, in order to perform a non-uniformity correction (NUC) and to obtain images (signals) in terms of apparent (or equivalent) temperatures. However, as we will see, even when calibrated in this way, a recording of a uniform background may still contain large-scale non-uniformities if there is a change in conditions such as temperature, distance or lens focus.

The noise cube methodology introduced in Ref. 1 is the subject of Sec. 2. We discuss how to identify and isolate trends in the data, and we show how the large-scale trend can be used to quantify non-uniformity in an image. As examples, we apply this technique to “raw” (i.e. non-calibrated) and calibrated image sequences from two of our IR cameras.

In Sec. 3, we use this technique in three applications with calibrated images: First, when a camera is calibrated at two reference blackbody temperatures, how do uniform scenes at different temperatures appear? Second, it was pointed out in Ref. 4 that NETD is affected by whether irradiation from a scene is flooded or collimated. When we calibrate a camera at short distance, how uniform is a scene at longer distance? And third, how is the uniformity of a recording at different settings for the lens focus?

---

\*Corresponding author. E-mail: Jan.Thomassen@ffi.no, Telephone: +47 63 80 79 21

## 2. THREE-DIMENSIONAL NOISE MODEL METHODOLOGY

Let us start with the theoretical considerations of the noise cube methodology before we apply this on an example with a real camera system.

First, a comment on terminology: In our context, the data is considered to be a sum of “signal” and “noise”, so that data can be decomposed into a signal part and a noise part. It is not always clear what is “signal” and what is “noise”, and this may also depend on the context. This means that we will sometimes be pragmatic in our terminology and say things like “the signal contains a noisy part”.

### 2.1 Basic methodology

The 3D noise model was introduced in Ref. 1 as a framework for describing noise in thermal imaging systems. We will first recall the methodology, and we will follow the terminology and notation from that paper.

Consider a signal in the form of a data cube  $U$  of dimensions  $T \times V \times H$ , where  $T$  is the number of frames (“time”),  $V$  is the vertical extension and  $H$  is the horizontal extension,

$$U = U(t, v, h) \quad \text{where} \quad t = 0, \dots, T-1, \quad v = 0, \dots, V-1, \quad h = 0, \dots, H-1. \quad (1)$$

Such data cubes are elements in a space of 3D arrays. Each element in the cube may be a digital number (for raw data) or a radiance value or temperature (for calibrated data). These data are sampled at discrete points, but sometimes it is convenient to consider an underlying signal on a continuum, so that  $(t, v, h)$  are points in a region of  $\mathbb{R}^3$ :

$$(t, v, h) \in [0, T] \times [0, V] \times [0, H]. \quad (2)$$

Hopefully, this will be clear from the context.

This signal can be decomposed into eight components by the use of an “averaging operator”  $D_i$ , with  $i = t, v$  or  $h$ , which we define in the following way: We consider the  $t$ -direction for simplicity, but the idea is the same for all three directions. First define the mean value  $\bar{U}_t(v, h)$  of  $U$  in the  $t$ -direction as

$$\bar{U}_t(v, h) = \frac{1}{T} \sum_{t=0}^{T-1} U(t, v, h). \quad (3)$$

Then  $D_t$  is defined by

$$D_t U(t, v, h) = [\bar{U}_t(v, h), \dots, \bar{U}_t(v, h)] \quad (\text{repeated } T \text{ times}). \quad (4)$$

$D_t$  is a projection operator on the space of 3D arrays since the average of repeated values is the same value:  $D_t^2 U = D_t U$ . We can use  $D_t$  to split the signal into “noise” and mean value components in the time direction:

$$U = (1 - D_t)U + D_t U. \quad (5)$$

We can always perform this decomposition, but of course the  $(1 - D_t)$  part is proper noise only when no trend is present and the measured scene is uniform. An illustration of the noise decomposition operation  $D_t$  on a noisy signal is shown in Fig. 1. The Figure shows a synthetic signal  $U(t)$ ,  $0 \leq t \leq 500$ , consisting of a sum of a constant (arbitrarily equal to 1) and Gaussian random noise ( $\sigma = 0.25$ ).




Figure 1. Illustration of the noise decomposition operation  $D_t$  on a noisy signal.

Applying all combinations of  $D_i$  operators, we get the 3D noise decomposition

$$U = S + N_t + N_v + N_h + N_{vh} + N_{th} + N_{tv} + N_{tvh}, \quad (6)$$

where the terms are defined by:

$$\begin{aligned} S &= D_t D_v D_h U, \\ N_t &= (1 - D_t) D_v D_h U, \\ N_v &= D_t (1 - D_v) D_h U, \\ N_h &= D_t D_v (1 - D_h) U, \\ N_{vh} &= D_t (1 - D_v) (1 - D_h) U, \\ N_{th} &= (1 - D_t) D_v (1 - D_h) U, \\ N_{tv} &= (1 - D_t) (1 - D_v) D_h U, \\ N_{tvh} &= (1 - D_t) (1 - D_v) (1 - D_h) U. \end{aligned} \quad (7)$$

These quantities are full three-dimensional arrays, but they may be constant in one or more directions. The subscripts indicate which direction(s) have true variations. Thus,  $S(t, v, h) = S$ , a global constant,  $N_t(t, v, h)$  has variation only in the time direction but is constant in the pixel plane ( $v, h$ ) for a given frame  $t$ ,  $N_{vh}(t, v, h)$  has variation across the pixel plane but is constant in the  $t$ -direction. The last component  $N_{tvh}(t, v, h)$  has variations in all directions. It is related to the NETD of the detector.

## 2.2 Identifying and removing trends

The methodology described in Sec. 2.1 can always be applied to split a recorded signal into noise components. However, even for uniform scenes, such as a blackbody, the signal may contain large-scale trends, for example from the optics or the detector electronics. As described in Ref. 3, we need to identify these trends and handle them appropriately, typically by performing a low-order polynomial fit and removing it from the signal. In this Section, we elaborate on this technique and describe how we have used it on our data.

An illustration of the situation is shown in Fig. 2. This is a slight modification of our previous synthetic example in Fig. 1. Here we have added a “trend” in the form of a low frequency sine wave to the constant and noise. Applying the same decomposition, we see that a trend is still present in the “noise” term  $(1 - D_t)U$ .




Figure 2. Illustration of the noise decomposition operation  $D_t$  on a noisy signal with a large-scale trend.

For a real uniform scene, the large-scale trends show up in the vertical and horizontal noise components  $N_v$  and  $N_h$ , and in the pixel plane component  $N_{vh}$ . The trend in  $N_{vh}$  tends to go in a diagonal direction, since trends that are axis-aligned are contained in  $N_v$  and  $N_h$ . Thus, a trend in the  $(v, h)$ -plane is given by a sum,

$$f(v, h) = f_v(v) + f_h(h) + f_{vh}(v, h), \quad (8)$$

where  $f_v$  and  $f_h$  are separate trends in  $N_v$  and  $N_h$ , respectively, and  $f_{vh}$  is the trend in  $N_{vh}$ .

In order to identify the trends, low-degree polynomial fits can be made to  $n_v(v) = N_v(*, v, *)$ ,  $n_h(h) = N_h(*, *, h)$ , and  $n_{vh}(v, h) = N_{vh}(*, v, h)$ , where ‘\*’ means “any index”. One can implement these polynomial fits with weighted least squares. Let us consider the  $v$ -direction to be specific. We seek polynomials of the form  $f(v) = \sum f_i B_i^{(d)}(v)$ , where  $f_i$  are the coefficients and  $B_i^{(d)}(v)$  are the basis functions for polynomials of a given degree  $d$ . Thus, we find the best fit  $f(v)$  to the function  $n_v(v)$  by minimizing the integral

$$\int_0^V |f(v) - n_v(v)|^2 w(v) dv \quad (9)$$

with respect to the coefficients  $f_i$ . Here,  $w(v)$  is a weight function. As we will see, in some cases, depending on the context, the signal  $n(v)$  may have unusual behaviour at the boundary. E.g. the raw signal from our IRCAM LWIR camera has consistently higher values at the top and bottom rows of the sensor matrix. We may then choose a weight function  $w$  that is zero on the boundaries, that is,  $w(v) = 0$  for  $v = 0$  and  $v = V$ . We tend to use the weight function\*

$$w(v) = \frac{2}{V} \sqrt{(V-v)v}, \quad (10)$$

which is smooth, but just zeroing out the end points,

$$w(v) = \begin{cases} 0, & v = 0 \\ 1, & 0 < v < V \\ 0, & v = V \end{cases} \quad (11)$$

will also work. For  $n_{vh}$ , a two-step process of first fitting in the  $v$ -direction, then in the  $h$ -direction may be used. There is no universal method of choosing the appropriate polynomial degrees for the different components. This must be determined by trial and error on a case-by-case basis. For the data set collected for this paper – to be described in more detail in Sec. 2.4 – we have found that the polynomial degrees that work well are  $d_v = 4$ ,  $d_h = 6$ , and  $\mathbf{d}_{vh} = (3, 3)$ , respectively. These degrees were high enough to describe the low-frequency trends.

When the trend functions  $f_v(v)$ ,  $f_h(h)$  and  $f_{vh}(v, h)$  are found in this way, we can subtract them from  $N_v$ ,  $N_h$  and  $N_{vh}$  and add them to  $S$ :

$$\begin{aligned} N_v(t, v, h) &\rightarrow N_v(t, v, h) - f_v(v), \\ N_h(t, v, h) &\rightarrow N_h(t, v, h) - f_h(h), \\ N_{vh}(t, v, h) &\rightarrow N_{vh}(t, v, h) - f_{vh}(v, h), \\ S(t, v, h) &\rightarrow S(t, v, h) + f_v(v) + f_h(h) + f_{vh}(v, h). \end{aligned} \quad (12)$$

$S$  is no longer a constant, but contains trends from the signal  $U$ . The  $N$ -terms contains noise, but may also contain components of the signal that are not described by a low order polynomial, like the mentioned jump in the top and bottom rows for raw signals.

Figure. 3 now illustrates how the new decomposition works.




Figure 3. Illustration of the new decomposition into signal and noise.

### 2.3 Calibration

Calibrated IR images have noise levels that are significantly reduced compared to raw images, but may still have non-trivial large-scale trends – even for recordings of uniform targets. We applied a two-point calibration procedure in our analysis. We are considering laboratory measurements in this paper, not real target scenes, so any complexity coming from atmospheric effects will be ignored. In particular, we consider the atmospheric transmission to be unity:  $\tau_{atm} = 1$ . As a result, the absolute values for equivalent temperatures that were estimated are not always correct. However, errors in temperature *differences* should be small.

\*This choice is inspired by the weight function  $w(x) = \sqrt{1-x^2}$  on the interval  $[-1, 1]$  associated with the inner product of Chebyshev polynomials of the second kind.<sup>5</sup>

A simplified model of the signal received at a detector element (pixel) from the target is then given by<sup>6</sup>

$$S_{tgt} = S_0 + K \int_0^\infty N_{tgt}(\lambda) r_s(\lambda) d\lambda, \quad (13)$$

where

- $S_0$  is the camera offset,
- $K$  is the gain,
- $N_{tgt}(\lambda)$  is the spectral radiance from the target,
- $r_s(\lambda)$  is the relative spectral response of the detector, and
- $\lambda$  is wavelength.

For calibration, the targets used were a single blackbody set at two different temperatures. The emissivity of the blackbody was high,  $\epsilon \approx 1$ , see Sec. 2.4 for more details of the blackbody. Thus  $N_{tgt}(\lambda) = N_{bb}(\lambda, T)$ , where

$$N_{bb}(\lambda, T) = \frac{2c^2h}{\lambda^5 (e^{hc/\lambda kT} - 1)}. \quad (14)$$

Here,  $c$  is the speed of light in vacuum,  $h$  is Planck's constant and  $k$  is Boltzmann's constant.

Let us define the integrated irradiances  $N_1$  and  $N_2$  of the two blackbodies at temperatures  $T_1$  and  $T_2$  as

$$N_i = \int_0^\infty N_{bb}(\lambda, T_i) r_s(\lambda) d\lambda, \quad i = 1, 2, \quad (15)$$

and the *time average* of the measured signals as  $S_1$  and  $S_2$ . We then get two equations of the form

$$S_i = S_0 + KN_i, \quad (16)$$

which can be solved for  $K$  and  $S_0$ :

$$K = \frac{S_2 - S_1}{N_2 - N_1}, \quad (17)$$

$$S_0 = S_1 - KN_1. \quad (18)$$

$K$  and  $S_0$  are both arrays in the  $(v, h)$ -plane:  $K = K(v, h)$  and  $S_0 = S_0(v, h)$ . They can be used to convert a full 3D data cube of a measured target  $S_{tgt}(t, v, h)$  into radiance values by inverting Eq. (16) for each time slice:

$$N_{tgt}(t, v, h) = \frac{S_{tgt}(t, v, h) - S_0(v, h)}{K(v, h)}. \quad (19)$$

This radiance cube can then be converted to a 3D data cube of apparent temperatures by solving Eq. (15) numerically for each  $(t, v, h)$ .

## 2.4 Measurements and implementation

The sensors used in this analysis were two IRCAM cameras – one MWIR and one LWIR – mounted with 50 mm lenses. Sensor sizes were  $512 \times 640$  pixels. Recordings were done with a total of 100 frames per sequence. Thus, a full data cube will have the pixel dimensions  $T \times V \times H = 100 \times 512 \times 640$ . From these cameras we have access to “raw” data, which means 14-bit numbers for each pixel. So the LSB value of a pixel is an integer in the range  $[0, 16383]$ . For calibration, bad pixels must be identified and handled, which will be discussed later. A table with the characteristics are shown in Tab. 1.

The target that was recorded for all the applications in this paper was an SR-800N blackbody from CI Systems, size  $40 \times 40$  cm, with a high emissivity surface – for all practical purposes  $\epsilon = 1$ . The surface of this blackbody has high uniformity. Temperatures can be set to a wide range of values, and in our case temperatures from about  $15^\circ\text{C}$  to  $50^\circ\text{C}$  were used. This blackbody was used both as a target and as a calibrating blackbody.

Table 1. Characteristics of the IRCAM cameras.

	LWIR	MWIR
Model	Equus 327k L pro	Equus 327k SM pro
Spectral range	7.9 – 9.6 $\mu\text{m}$	3.0 – 5.0 $\mu\text{m}$ (with filter)
Detector type	MCT	InSb
Array size	640 $\times$ 512	
Pixel pitch	16 $\mu\text{m}$	15 $\mu\text{m}$
A/D bit resolution	14 bits	
Focal length	50 mm	
Horizontal Field of View (FOV)	11.7 $^\circ$	11.0 $^\circ$
Instantaneous FOV (IFOV)	320 $\mu\text{rad}$	300 $\mu\text{rad}$
Integration time	0.1 ms	1.0 ms

In Ref. 3, a part of the text was devoted to describing efficient ways to compute the various noise components in order to “speed up computation times”. Reference 3 is from 2005. However, efficiency is not a problem today, and we used a brute force approach for our computations. Our analysis code was implemented in Python with packages from SciPy and NumPy.<sup>7</sup> Operations on NumPy arrays are very efficient, and noise decomposition of data cubes of dimensions  $(T, V, H) = (100, 512, 640)$  were quite fast – about 2.3 s – when run on a Windows 10 laptop with Intel Core i5 CPU and 16 GB RAM memory.

## 2.5 Example: IRCAM MWIR and LWIR cameras

As an example of 3D noise cube analysis, let us consider both raw and calibrated signals from our IRCAM cameras. For these measurements, the distance from the camera to the target blackbody was  $r = 1.9$  m and the temperature was  $T = 30$   $^\circ\text{C}$ .

Visualising the results of a decomposition is not trivial.  $N_t$ ,  $N_v$  and  $N_h$  are inherently one-dimensional and may be shown as a function plot.  $N_{vh}$ ,  $N_{th}$  and  $N_{tv}$  are two-dimensional and may be shown as surface plots or images. Even better if the surface plot is interactive, so that the observer may rotate the axes to see the data from all angles. Nevertheless, we find that a “normal” surface plot, with a careful attention to position and scale of the axes provides a useful visualisation. For instance,  $N_{vh}$  has a number of high and low spikes that stand out from the values clustered around the mean, which is necessary to clip for a useful visualisation.  $N_{tvh}$  is inherently three-dimensional, but can be visualised by a surface plot of a time slice.  $S$  will also be visualised by a time slice, since this will show the trend in later examples.

### 2.5.1 Raw signals, LWIR

The result of the noise decomposition of a raw LWIR signal, without any handling of the trends, is shown in Fig. 4. The mean value  $\bar{S}$  and standard deviations  $\sigma_i$  for the noise components  $N_i$  are shown in Tab. 2. All the  $N_i$ 's have zero mean. Trends are quite notable in  $N_v$  and  $N_h$ , which clearly reveals a fall-off of signal intensity towards the edges of the focal plane array. A trend is present also in  $N_{vh}$ , but it is much harder to see because of the noise pattern with a lot of spikes. Also visible in  $N_v$  is a significantly stronger signal from the top and bottom rows, which is probably due to the readout electronics. All noise components related to time variation –  $N_t$ ,  $N_{tv}$ , and  $N_{th}$  – are small, as measured relative to  $N_{vh}$ .

We can apply the trend removing procedure previously described to these components. The result is shown in Fig. 5. The mean value and standard deviations are now as shown in Tab. 3. Most notably,  $\sigma_v$  and  $\sigma_h$  has gone down.  $\sigma_{vh}$  has not changed much, which indicates that most of the variation in this component is caused by the tails of the distribution. It turns out that the mean values of  $N_{vh}$ ,  $N_v$  and  $N_h$  are no longer exactly zero – they are 0.61,  $-0.23$  and 21.83, respectively. This happens because the polynomial trends that are subtracted from them are calculated from least squares minimisation, with no restrictions on the mean values.

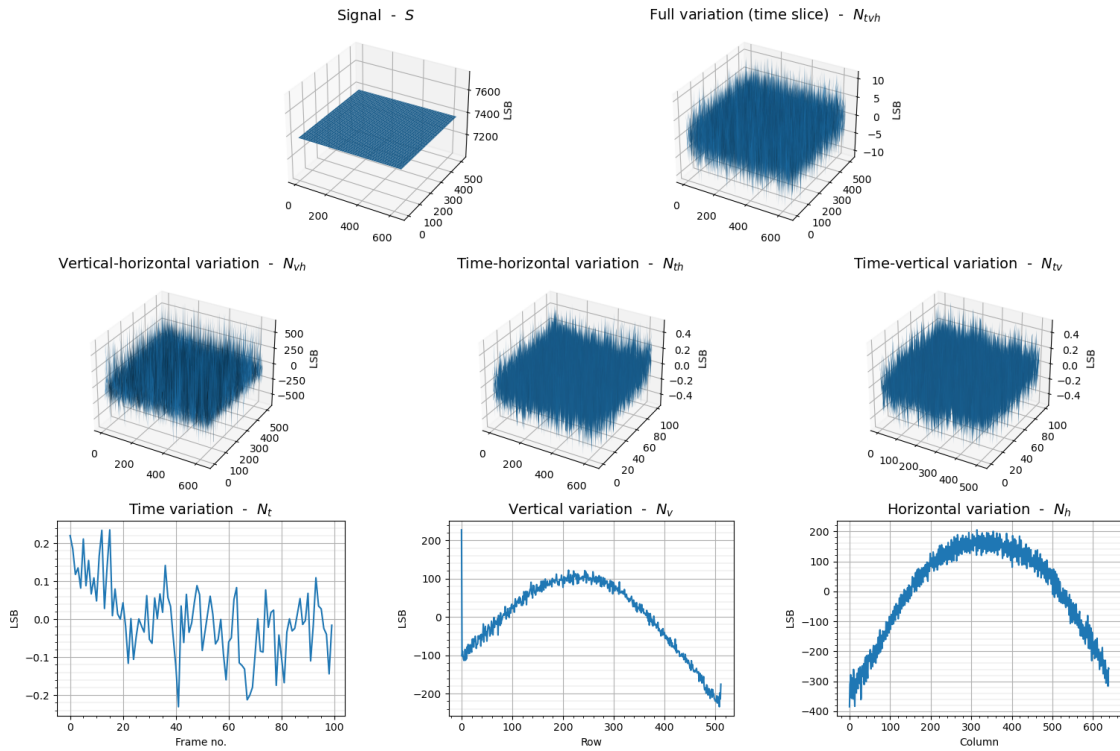


Figure 4. Noise decomposition of a raw LWIR signal.

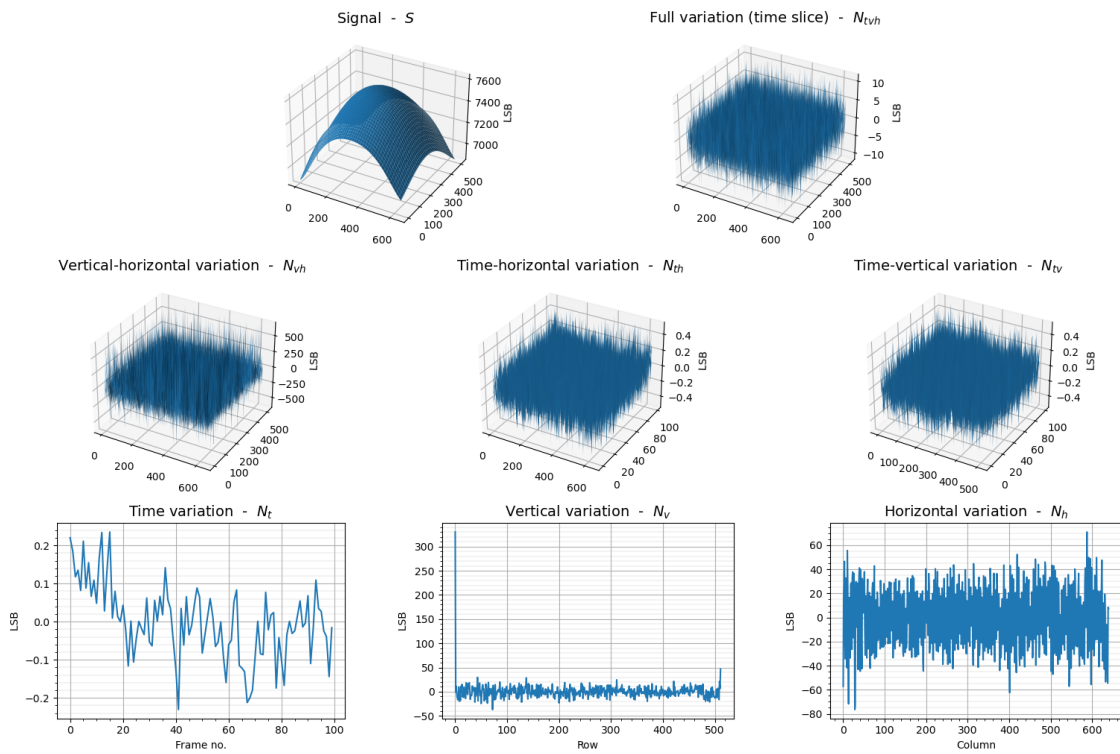


Figure 5. Trends removed from  $N_v$ ,  $N_h$  and  $N_{vh}$  and added to  $S$ . LWIR.



Table 2. Mean value and standard deviations for the raw LWIR signal. The units are LSB.

$\bar{S} = 7376.43$		$\sigma_{tvh} = 3.82$
$\sigma_{vh} = 220.30$	$\sigma_{th} = 0.18$	$\sigma_{tv} = 0.18$
$\sigma_t = 0.10$	$\sigma_v = 92.72$	$\sigma_h = 152.22$

Table 3. Mean value and standard deviations for the detrended raw LWIR signal. Units are LSB.

$\bar{S} = 7354.22$		$\sigma_{tvh} = 3.82$
$\sigma_{vh} = 220.53$	$\sigma_{th} = 0.18$	$\sigma_{tv} = 0.18$
$\sigma_t = 0.10$	$\sigma_v = 17.44$	$\sigma_h = 25.78$

Still visible are the higher signals from the top and bottom rows in  $N_v$ . These stand out from what appears to be a background level of random noise, so they are easily isolated and can be subtracted from  $N_v$ . The remaining signal now looks “random-like”, see Fig. 6. For this remaining component, the standard deviation drops to  $\sigma_v = 9.35$ .

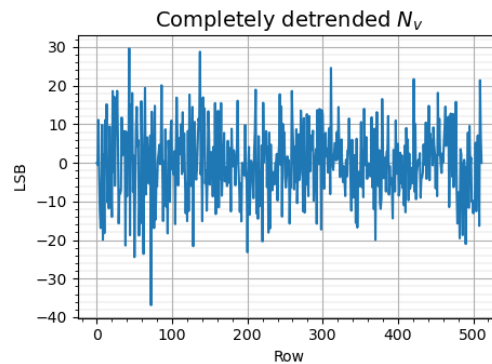


Figure 6. Noisy part of  $N_v$ . The noise appears to be random. LWIR.

The new  $N_h$  component does not have any remaining large-scale trends after the polynomial part is removed, but a close examination shows that it has a “small-scale” trend – it appears to have a high frequency non-random component. This is easily verified by a Fourier transformation. Removing this frequency component now reveals an apparently random signal. Figure 7 shows this for the first 100 columns of  $N_h$ . The new standard deviation is  $\sigma_h = 13.09$ .

Finally, let us consider  $N_{vh}$ . A large number of (positive) spikes in this component makes it difficult to see




Figure 7. High frequency component of  $N_h$ . First 100 columns are shown. LWIR.

the trend. However, a look at the least squares polynomial fit (degrees (3, 3)), reveals a non-trivial contribution, as shown in Fig. 8.

Least squares polynomial fit to  $N_{vh}$




Figure 8. Least squares polynomial fit to  $N_{vh}$ .

The values of the spatial noise components  $\sigma_{vh}$ ,  $\sigma_v$  and  $\sigma_h$  are large compared to residual component  $\sigma_{tvh}$ , which is related to NETD. This reflects the fact that uncalibrated signals are very noisy and have poor image quality. Later, when we consider calibrated signals, these ratios improve considerably.<sup>†</sup>

### 2.5.2 Raw signals, MWIR

The raw signal from the MWIR camera is similar to the LWIR. These MWIR measurements were also done at distance  $r = 1.9$  m with  $T = 30$  °C. For this signal we have the mean value and standard deviations as shown in Tab. 4.

Table 4. Mean value and standard deviations for the detrended raw MWIR signal. Units are LSB.

$\bar{S} = 7122.28$		$\sigma_{tvh} = 6.23$	
$\sigma_{vh} = 197.83$	$\sigma_{th} = 0.30$	$\sigma_{tv} = 0.31$	
$\sigma_t = 0.26$	$\sigma_v = 10.86$	$\sigma_h = 30.08$	

Two notable differences compared to the LWIR is related to the vertical component  $N_v$ . First of all, there is a much smaller increase in the signal for the top and bottom rows. On the other hand, for an extended part at the bottom of the sensor there appears to be an increase in a band of about 64 pixels high, see Fig. 9 which shows  $N_v$  in this case. This second trend could be removed by applying a piecewise polynomial fit, however it essentially disappears in the calibrated signal.

### 2.5.3 Calibrated signals

Let us now consider calibrated signals. We apply two blackbody recordings to convert a data cube to radiance values, as described in Sec. 2.3. With calibration, the problem of bad pixels turn up. Bad pixels can be treated with varying degrees of sophistication. Taking a minimal approach, we need to fix pixels in the gain matrix  $K$  that are zero because we divide with  $K$  in the conversion to radiance. We do this by replacing small (and negative!) values by median filtered values. The same signal from the LWIR camera as discussed above will now decompose as shown in Fig. 10. Calibration temperatures were  $T_1 = 25$  °C and  $T_2 = 30$  °C.

Many features from the raw signal disappears, because fixed patterns of pixels that are consistently high or low are “calibrated out”. There are still trends that can be identified by the polynomial fitting methods. After this trend removal we end up with a decomposition as shown in Fig. 11. There is still a non-random periodic component in  $N_h$ , but the amplitude is quite small, less than 0.01 K.

<sup>†</sup>We have not adjusted the camera integration time, or other settings, to minimise noise.

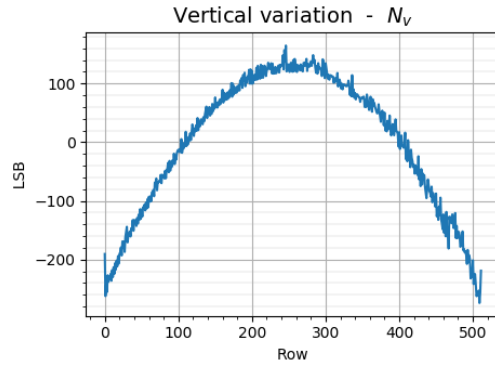


Figure 9.  $N_v$  for the MWIR camera.

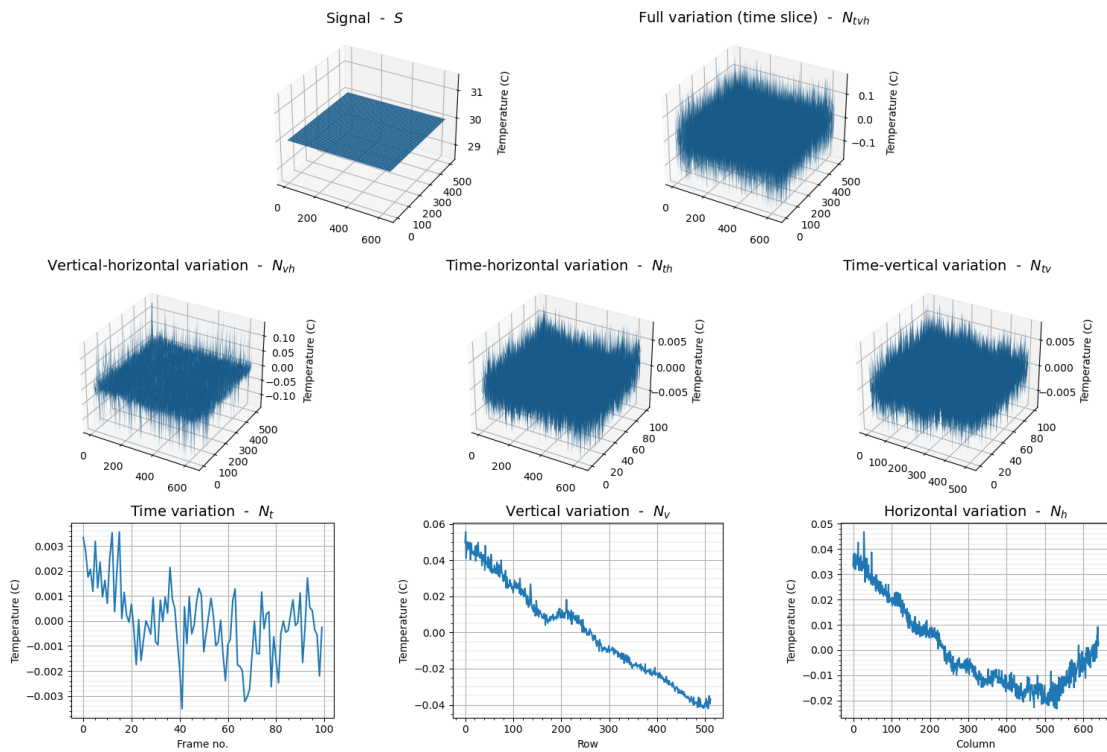


Figure 10. Decomposition of a calibrated signal from the LWIR camera. Units are degrees Celsius.

The mean value and standard deviations of the calibrated LWIR signal is shown in Tab. 5. From the Table we can find from  $\sigma_{tvh}$  that the NETD is 58 mK at  $T = 30^\circ\text{C}$  and integration time 0.1 ms.

The  $S$ -part of the signal, which in our formulation contains the trends, is the interesting part to us. We can

Table 5. Mean value and standard deviations for the calibrated LWIR signal. Units are degrees Celsius.

$\bar{S} = 30.0$		$\sigma_{tvh} = 0.058$
$\sigma_{vh} = 0.046$	$\sigma_{th} = 0.003$	$\sigma_{tv} = 0.003$
$\sigma_t = 0.001$	$\sigma_v = 0.003$	$\sigma_h = 0.003$

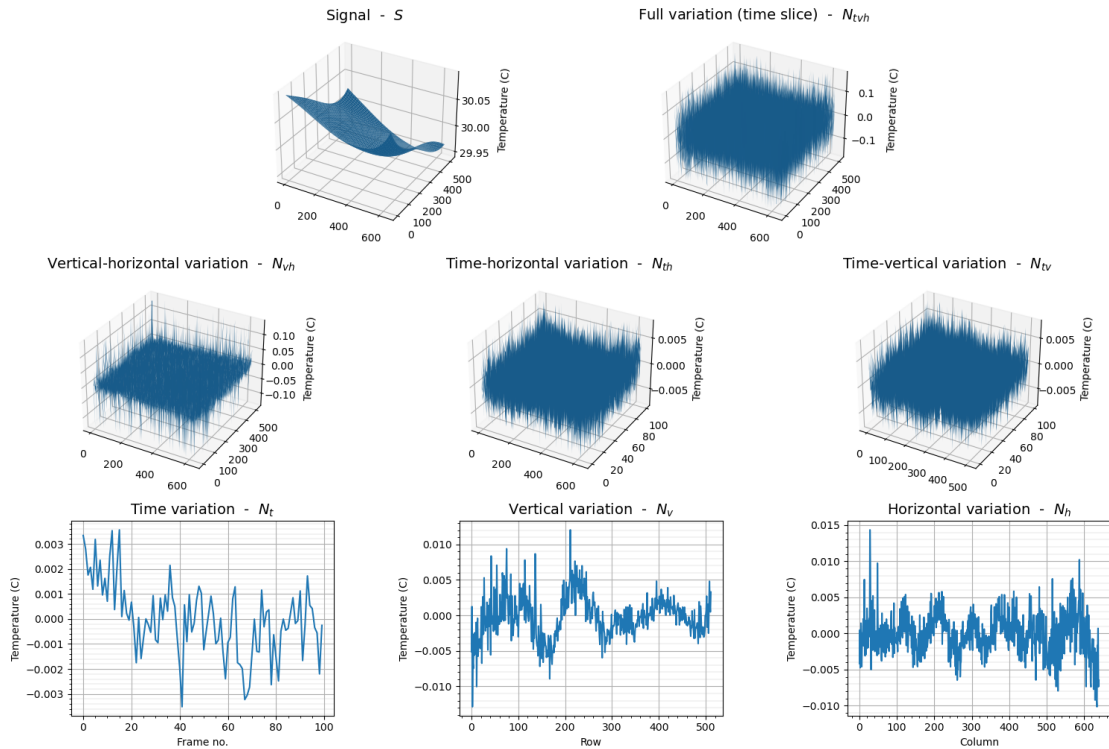


Figure 11. Calibrated signal with trends removed. LWIR.

use this component to describe large-scale non-uniformity of a signal as the conditions have changed relative to the conditions when calibrated. This will be explored in the next Section. A contour map of (a slice of) the calibrated signal  $S$  looks like Fig. 12. The temperature map shows regions in the image plane with temperatures in the same interval. Each interval has a range of 0.1 K.

A contour map like this is a tool to quantify and visualise large-scale non-uniformities in the signal.

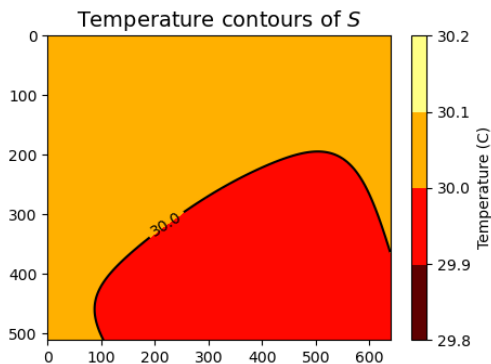


Figure 12. Contour map of a time slice of the calibrated signal  $S$ , after conversion to temperatures.

### 3. APPLICATIONS

We can use the contour map calculation from the last Section to investigate any deviation from uniformity that may arise when the conditions of the measurements depart from those at the time of calibration. We consider three different effects: 1) the temperature of the scene is outside the range defined by the cold and hot calibration blackbodies, 2) the calibration is done with blackbodies at close range but the scene is at long distance, and 3) the calibration is done with one setting of the lens focus, but the scene is recorded with a different focus setting.

To find out, we have made a series of controlled measurements in the laboratory. The scenes that were recorded were in all cases the blackbodies from CI systems, as mentioned in Sec. 2.4. The surfaces of these blackbodies are highly uniform. The distances to these targets were from  $r = 0.1$  m to about  $r = 2$  m. The scenes are therefore somewhat artificial, but should nevertheless indicate whether the various effects exist, and what their magnitudes are.

#### 3.1 Variations with temperature

Two-point calibration fixes the signal to two blackbodies, uniformly radiating at two different temperatures. The calibrated signal is by definition uniform at these two points. What about signals that are measured at other points, either inside or outside the range defined by the two temperatures? To find out, we made a series of measurements of our calibration blackbody at different temperatures, from  $25^\circ\text{C}$  to  $50^\circ\text{C}$  in increments of  $5^\circ\text{C}$ . We could then pick two calibration measurements and one target scene from this series and look for effects. More than one measurement was taken at each temperature, so that a target measurement could always be found that was different from any of the calibration measurements. The distance to the blackbodies were  $r = 1.9$  m, which allowed us to fill the image plane with the uniform blackbody.

First let us choose  $25^\circ\text{C}$  and  $35^\circ\text{C}$  for calibration points. Temperature maps for MWIR and LWIR at the various temperatures are shown in Figs. 13 and 14, respectively. In these Figures, the first and third plots are the same temperatures as the calibration points, marked with blue and red borders. These are fairly uniform to within  $\Delta T \approx 0.1$  K. Why these two plots are not exactly uniform is not clear – perhaps a small deviation from  $\epsilon = 1$  for the blackbodies or other environmental effects in our laboratory. In any case, this level of temperature difference represents a kind of “baseline” non-uniformity for our measurements. Increasing the temperature of the target to  $T = 50^\circ\text{C}$  increases the non-uniformity of the target to about  $\Delta T \approx 0.5$  K. The target that lies within the calibration interval – the second plot at  $T = 30^\circ\text{C}$  – also has a small non-uniformity of  $\Delta T \approx 0.2$  K for MWIR, slightly less for LWIR.

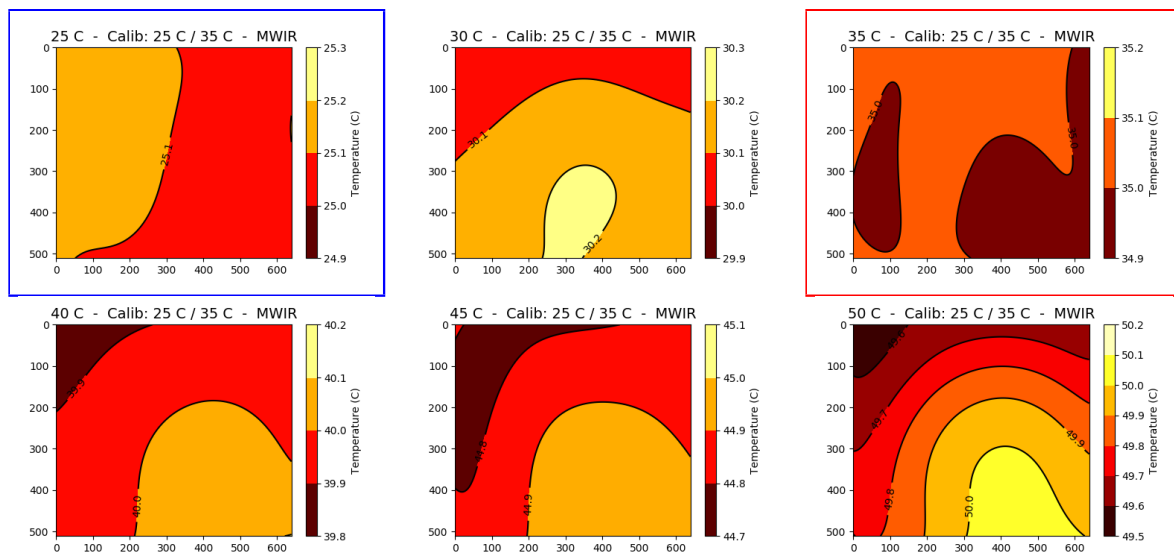


Figure 13. Temperature maps for MWIR calibrated at  $25^\circ\text{C}$  and  $35^\circ\text{C}$  (blue and red borders).

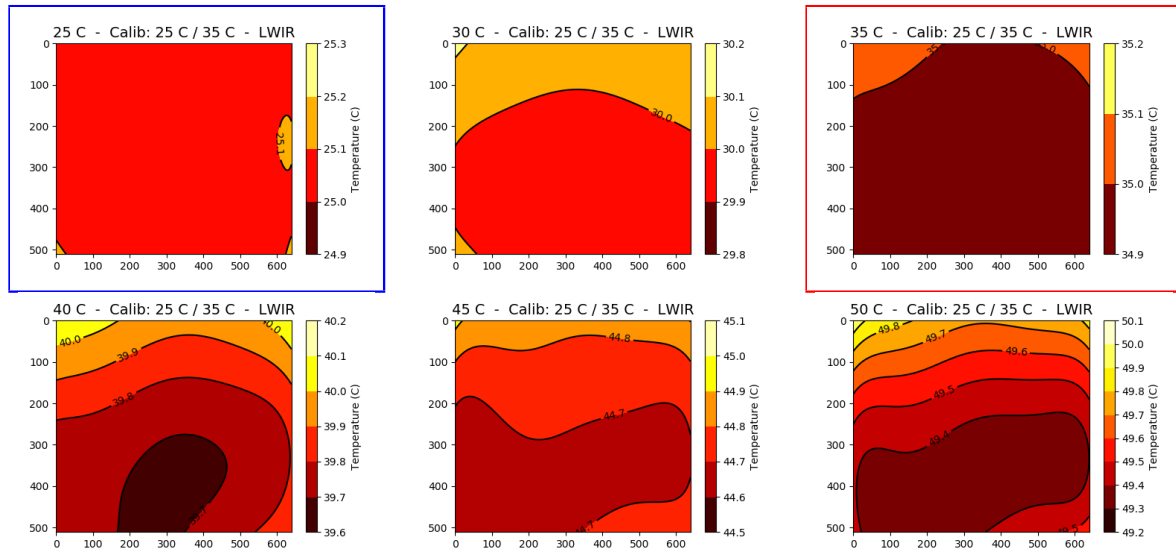


Figure 14. Temperature maps for LWIR calibrated at 25 °C and 35 °C (blue and red borders).

If we choose calibration points at the high end of the range, 40 °C and 50 °C, we find the maps shown in Figs. 15 and 16. Similar effects are present in this case, with more non-uniformity showing up at the low end, farthest from the calibration points. However, for LWIR, the most prominent non-uniformity appears at  $T = 35\text{ }^{\circ}\text{C}$ , just outside the range.

Why do we see these effects? One possible explanation is that the response of the detector elements is not exactly linear, so that extrapolation outside the calibration interval will gradually get worse. Another possibility is uncertainties in the calibration procedure, i.e. measurement uncertainties in the input to the calibration calculations. An uncertainty of  $\sim 0.1\text{ K}$  at each calibration point – 10 K apart – might lead to extrapolation errors of  $\sim 0.5\text{ K}$  at the far end of the range. Also, if the blackbody has an emissivity that is slightly less than 1 ( $\epsilon < 1$ ), then small reflections from the environment may be visible in the image of the blackbody recordings used for calibration, which is amplified when extrapolated outside the range. Furthermore, absorption in, or

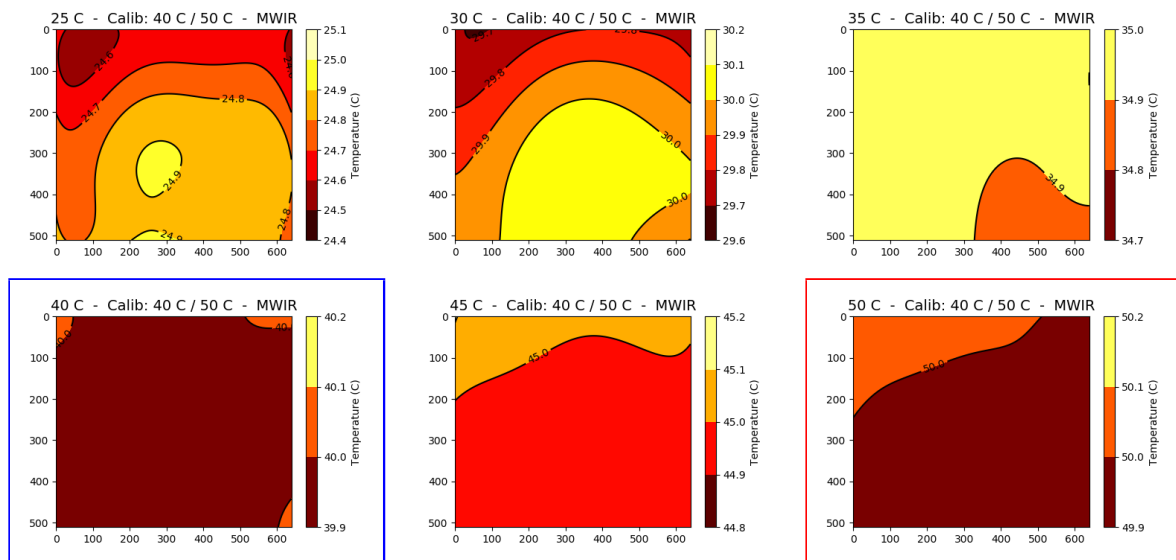


Figure 15. Temperature maps for MWIR calibrated at 40 °C and 50 °C (blue and red borders).

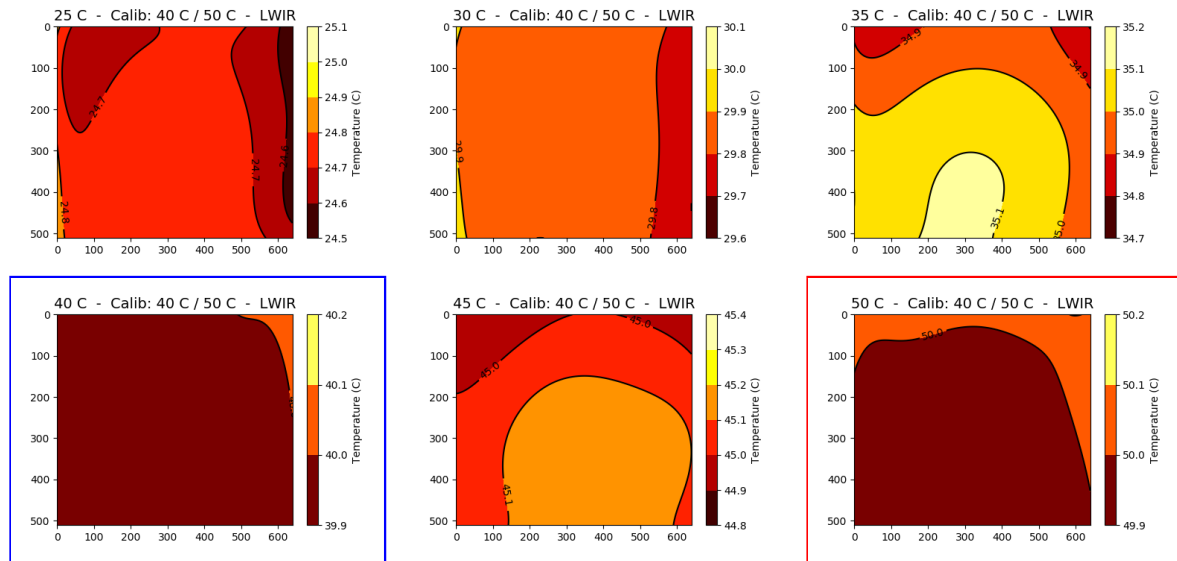


Figure 16. Temperature maps for LWIR calibrated at 40 °C and 50 °C (blue and red borders).

emission from, the optics may also contribute to these effects. Whatever the explanation, we need to be careful with measurements that are extrapolated too far outside the calibration interval.

### 3.2 Flooded vs. collimated irradiation

It was pointed out in Ref. 4 that the NETD of a sensor that is irradiated with collimated radiation from a (uniform) scene is lower than the same sensor with flooded irradiation. Their analysis was done to compare noise levels of sensor systems with and without optics. In our case, we are considering a camera with a lens, i.e. a full optical system, but there could conceivably be a similar phenomenon. Calibration to blackbodies are usually done at close range – a few centimetres, i.e. ”flooded” – while target scenes are at a distance – at the order of kilometres, i.e. ”collimated”.

To see if we could find a distance effect, we made calibration measurements at close distance to the lens,  $r = 0.1$  m, and target scene recordings of our blackbody at  $r = 1.90$  m. Calibration temperatures were  $T_1 = 25$  °C and  $T_2 = 35$  °C. The horizontal FOV for our cameras with 50 mm lenses were about 11 ° for the MWIR and 11.7 ° for the LWIR. Irradiation spanning these angles is not perfectly collinear, but should be sufficient to indicate whether such an effect exists. The distance  $r = 1.9$  m was such that the image plane was fully covered by the surface of the target.

Figure 17 shows temperature maps of scenes of blackbodies at 35 °C – MWIR and LWIR – for both close up and distant scenes. Figure 17 appears to verify that there is an effect. In particular the LWIR plot shows a clear non-uniformity with  $\Delta T \approx 0.6$  K.

It is conceivable that stray light in the form of photons entering the aperture from a large angle could scatter into the field of view (FOV) and contribute to the signal at the sensor’s focal plane, even if their angle of incidence is outside the FOV. Another possible explanation for this is that the surface of the blackbody is not perfectly uniform.

### 3.3 Variations with lens focus

Finally, we investigate the effects of lens focus. Normally, calibration measurements are done with the lenses focused at approximately infinity, which is roughly where targets would be – even though the calibration blackbodies are at close range, i.e. they are strongly out of focus.

One may sometimes observe that calibrating and then changing focus leads to a circular region in the image with different levels. This is most visible when auto gain and offset is turned on in the camera viewer software. How strong is this effect really?




Figure 17. Temperature maps for flooded vs. collimated irradiation.

To find out, we took two “normal” calibration measurements at distance  $r = 0.1$  m, with  $T_1 = 18^\circ\text{C}$ ,  $T_2 = 28^\circ\text{C}$ . The lens was focused at a longer distance – in our case  $\sim 2$  m. This is not infinity, but it is significantly longer than the minimum focus distance. The target was the blackbody at the same distance  $r = 0.1$  m, with  $T_{tgt} = 25^\circ\text{C}$ . A set of target measurements were made with the lens focus set to three different values: minimum, maximum (infinity), and 6 m, which is somewhere in between.

Figure 18 shows temperature maps for these different lens focus settings for MWIR and LWIR. The Figure shows a difference between the MWIR and LWIR. For MWIR we essentially have no effect. However, for LWIR there is a strong effect, in particular for the long focusing distances, with  $\Delta T \approx 0.6$  K. It is not clear why there is such a difference between MWIR and LWIR, but it is possible that it is related to differences in optical design and materials in the MWIR and LWIR lenses.

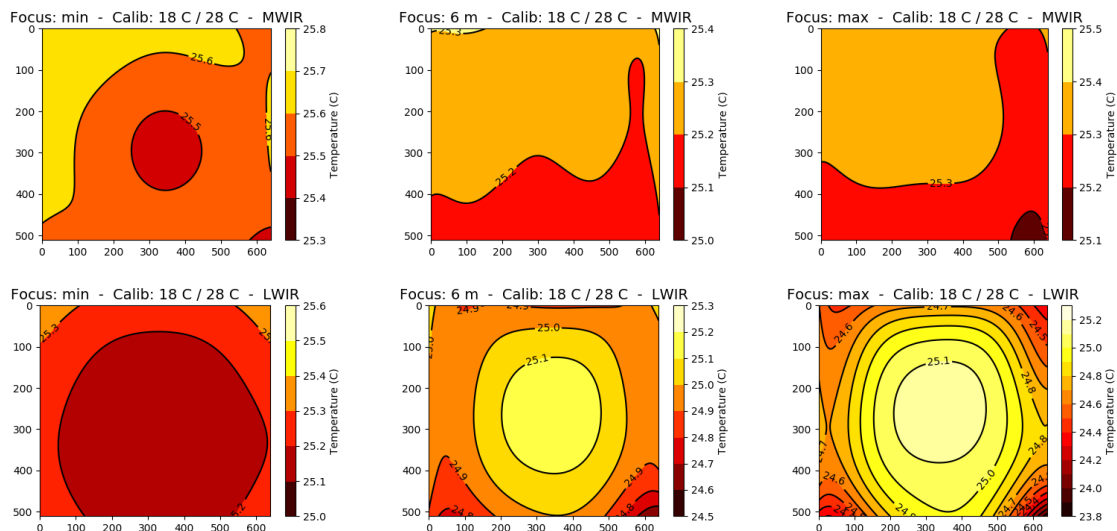


Figure 18. Temperature maps for different lens focus settings.



## 4. DISCUSSION AND CONCLUSION

We have seen that the 3D noise model leads to a decomposition that can be used not only to analyse noise components, but also other properties of the signal. We have used this methodology to extract large-scale trends and thereby non-uniformity. For a uniform scene such as a blackbody, a temperature map can be made in order to visualise and quantify these large-scale non-uniformities in terms of temperature differences.

One possible way these effects could be a problem is when computing contrasts between a target and the background. If the target is located in the centre of the image, its contrast with background pixels off centre could be *over- or underestimated* by perhaps a few tenths of degrees. This could be undesirable depending on the application.

Our measurements were done under controlled laboratory conditions that doesn't necessarily reflect what is seen in the field. Nevertheless, it seems that this is a "proof of concept", and that the effects could in some cases be large enough to influence the precision of calibrated images. A possible way to proceed from here is to see if similar effects are present in "realistic" measurements, such as a background measurement of a uniform cloudy sky.<sup>‡</sup>

## ACKNOWLEDGMENTS

We thank Lars Trygve Heen for reading and commenting on the manuscript.

## REFERENCES

- [1] D'Agostino, J. A. and Webb, C. M., "Three-dimensional analysis framework and measurement methodology for imaging system noise," in [*Infrared Imaging Systems: Design, Analysis, Modeling, and Testing II*], Holst, G. C., ed., **1488**, 110 – 121, International Society for Optics and Photonics, SPIE (1991).
- [2] Holst, G. C., [*Electro-optical imaging system performance*], JCD Publishing, Oviedo, FL, and SPIE Press, Bellingham, WA, sixth ed. (2017).
- [3] O'Shea, P. and Sousk, S., "Practical issues with 3D noise measurements and application to modern infrared sensors," in [*Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XVI*], Holst, G. C., ed., **5784**, 262 – 271, International Society for Optics and Photonics, SPIE (2005).
- [4] Rogers, R. K., Edwards, W. D., Waddle, C. E., Dobbins, C. L., and Wood, S. B., "Reporting NETD: why measurement techniques matter," in [*Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXV*], Holst, G. C., Krapels, K. A., Ballard, G. H., Jr., J. A. B., and Jr., R. L. M., eds., **9071**, 280 – 297, International Society for Optics and Photonics, SPIE (2014).
- [5] Abramowitz, M. and Stegun, I. A., [*Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*], ch. 22, Dover, New York City, ninth dover printing, tenth gpo printing ed. (1964).
- [6] Stark, E., Heen, L. T., and Wikan, K., "SIMVEX 2001 trial – radiant intensity contrast," Tech. Rep. FFI/RAPPORT-2002/02568, FFI (2002).
- [7] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Jarrod Millman, K., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C., Polat, İ., Feng, Y., Moore, E. W., Vand erPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nature Methods* **17**, 261–272 (2020). URL: <https://scipy.org>.

---

<sup>‡</sup>One would have to be careful about the position of the sun, so that there is no gradient in the IR radiation from atmospheric scattering.