Detection: how many pixels are required?

van Rheenen, Arthur, Thomassen, Jan

SPIE.

Detection – how many pixels are required?

Arthur D. van Rheenen[1a], Jan B. Thomassen[a]

[a]Norwegian Defence Research Establishment, P. O. Box 25, N-2027 Kjeller, Norway

ABSTRACT

In this work, we studied how detection range of a small target changes as a function of the number of defined target pixels. The intuitively simplest method to look for a target in IR images is to look for the "hottest" (highest intensity) pixel. Image noise or sun glint may cause false detection. To build in some robustness against such detractions, one considers looking for the hottest group (blob) of contiguous pixels. A blob could be any shape; here we consider square blobs, 1x1, 2x2, 3x3, etc. pixels. One expects the average blob intensity to decrease with blob size. On the other hand, the noisiness of the background blobs also decreases with blob size. The net result is that initially, for small blob sizes, the detection range increases before falling again for larger blob sizes. We demonstrate this by analyzing IR recordings of a small vessel sailing outbound until it "disappears". We develop a simple model that supports these observations. The model is based on a synthetically generated sequence of images of a receding target, uses a basic sensor characteristic, and Johnson's detection criterion.

Keywords: detection range, contrast, background noise, small targets

1. INTRODUCTION

Detection, recognition, and identification are often lumped together when speaking of target task performance of sensors: DRI range is a familiar term. Detection differs from the other two tasks in that it does not require identifying any level of detail on the target – just locating it. In other words, the requirements on the spatial resolution of the sensor are mild for detection, only a sense of bearing (target direction) is required. On the other hand, sensor sensitivity is a crucial performance indicator. Larger pixels tend to have better signal to noise performance for extended objects because the signal strength increases linearly with the number of received photons, whereas the noise increases as the square root of the photon number.

In this work, we investigate whether combining multiple small pixels, which we shall call blobs, into larger pixels indeed results in larger detection ranges, and whether there is an optimal blob size, which maximizes the detection range. We limit ourselves to square blobs. Possibly, blobs with aspect ratios more similar to the target aspect ratio result in longer detection ranges.

We organize the paper as follows. First, we present detection ranges obtained from analyzing long IR recordings made during the 2016 Minotauros trial on Crete. During the experiment, a small boat sailed either out on a straight course from the sensors or in from a large distance. Detection ranges were estimated for different blob sizes and showed a maximum when the blob size was several pixels on a side. Next, we develop a bare-bones model that predicts detection ranges as a function of blob size. The model is too simple to predict accurate ranges, but predicts that this range function has a maximum. In the last section, we discuss the results and mention how the model could be improved.

2. EXPERIMENTS

Under auspices of NATO STO SET-211 an IR trial (Minotauros) was performed in September 2016 on the isle of Crete. One of the experiments consisted of a small boat sailing on a straight course, out from the sensor platforms or in towards the sensor platforms. We made IR recordings; both LW and MW, at low frame rates (1 fps) and range measurements were made at short intervals, allowing for co-registration of frame number and range. The sensors we used are commercially available IrCam Equus cameras – their details are presented in Table 1.

---

[1] arthur-d.vanrheenen@ffi.no

Table 1. Specifics of the IrCam Equus327 L/SM IR-cameras used in the experiments.

Characteristic	LWIR	MWIR
Optics [mm]	50/200	50/200
F-number	2	2
Detector pitch [μm]	16	15
Focal plane (H x V) [pixels]	640 x 512	640 x 512
IFOV [μrad]	320/80	300/75
Horizontal FOV [degrees]	11.7/2.9	11.0/2.7
Spectral range [μm]	7.5 – 9.9	3.0 – 5.0
Integration time [ms]	0.1	1

Before analyzing the raw recordings, we applied two pre-processing steps: non-uniformity correction and bad-pixel removal. The first step equalizes the response of all pixels and the second identifies pixels with a weak or no response by replacing their reading with the median response of pixels surrounding the offending pixel.

Analysis begins with identifying the smallest area around the target, common to all frames that the target does not leave during the recording. Ideally, the boat sails in a straight line, but in reality it swerved a bit. Additionally, two areas, one on either side of the target area, are identified (see Fig. 1). The two areas on the side are used to calculate a median intensity-pitch profile, a background that is subtracted from the pixel intensities in the target area. Ideally, the target pixels in the target area will now have a non-zero intensity (contrast), and the background pixels in the target area will have intensity zero. Obviously, the background is not homogeneous: there is noise and the sea has wave structure. The next step in the image analysis is to slide one pixel at a time, a window of B x B pixels over the target area and calculate the average intensity. The highest intensity and the position of the window with the highest intensity is recorded. This is repeated for all frames in the recording and for several blob sizes B: 1, 2, 3, …, 10, 13, 16, 20, 25, and 30 pixels.
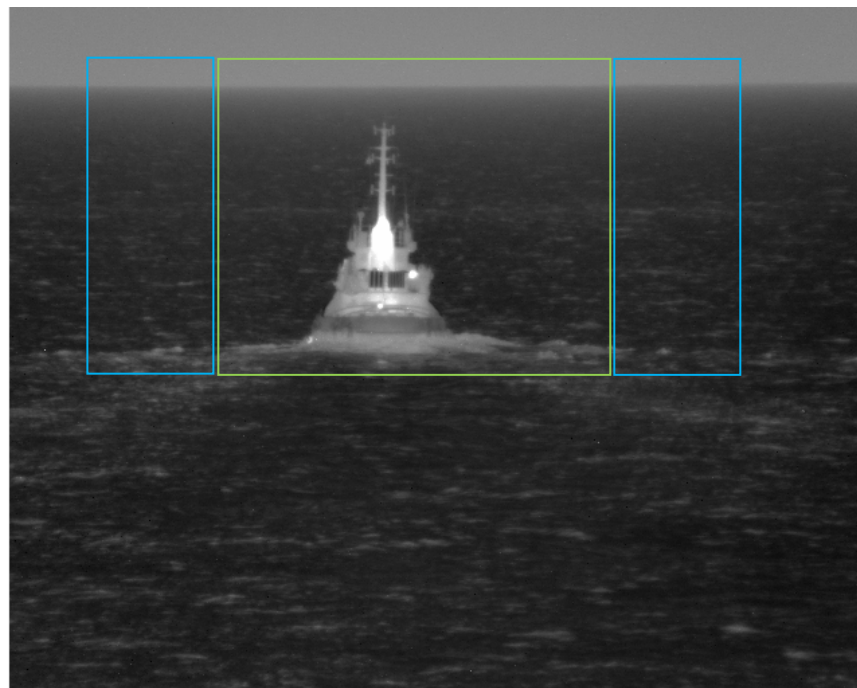

Figure 1. Example of a frame with the target box (green) and the two boxes (blue) that are used to calculate the median intensity-pitch profile of the background. This background profile is subtracted from the pixel intensities in the target box.

We use two approaches to estimate the detection ranges. The first is based on the intensity-range curves that we obtained for each blob size. As the range increases, the intensity (contrast) decreases and reaches a noisy plateau at large range. We define the detection range as the range where the decaying part of the curve meets the plateau.
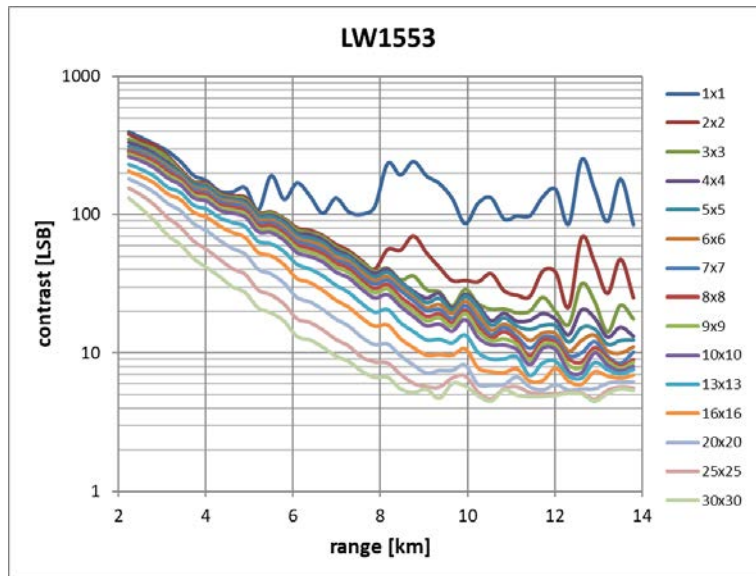

Figure 2. Mean blob intensity (contrast) versus range, here for an LW recording with a 200-mm lens. The legend on the side indicates the blob sizes in pixels. The detection range for the largest blob size (30 x 30 pixels) is about 8.5 km.

We notice in Fig. 2 that the plateaus are rather noisy, and that the extracted detection range comes with some uncertainty. To provide some consistency, we calculate the mean plateau value for the highest range values as well as the standard deviation (noisiness). Further, we fit the low-range part of the curves to an exponential, and calculate for what range the exponential meets the mean plateau value plus a factor times the standard deviation. We plot (Fig. 3) the estimated detection ranges (blue, red, green) for three values of the factor multiplying the standard deviation: 2, 3, and 5, respectively.
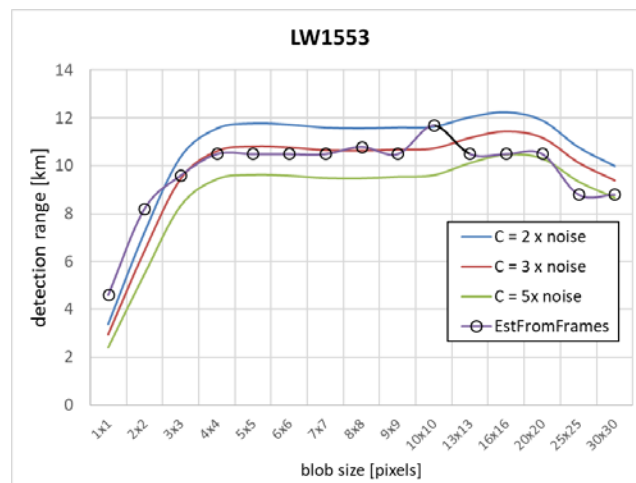

Figure 3. Detection range versus blob size for three different uncertainty factors (see text): blue-2, red-3, and green-5. The black circles are detection ranges estimated by inspection of frames. When the blob does not track the target anymore, the target has reached the detection range. The colored curves are obtained using the first method, the black circles are the result of the second method.

The other method we used to estimate the detection range is to look for the frame (range) when the blob loses track of the target. Since we kept a record of the location of the hottest blob in the image, limited to the "green" box in Fig. 1, we can overlay that blob on the frames and just visually inspect the frames to determine at what range the blob loses/starts track in case of an outbound/inbound run. We are not interested in the exact value of the detection range; we want to demonstrate that there is a maximum as a function of blob size. Therefore, we limit the search of the frames to 20 frames evenly spaced over the recording. This limits the amount of "manual" work to be performed. The detection range is determined as the interpolated range between the last frame with track and the first frame without track. Fig. 4 shows a sequence of frames, in this case of an inbound boat. Overlaid are the location (green circle) of the 30x30-pixel blob with the largest mean intensity, as well as the horizon (yellow line). We see that initially, starting at the top left, the blob misses the target, which is at the horizon, roughly in the middle. When the boat has gotten closer than 9 km, the blob tracks the boat all the way in. The estimated detection range is $0.5 \cdot (9 + 8.5) \sim 8.8$ km. The detection ranges thus determined are also plotted in Fig 3, indicated by the black circles. It is not clear why the blob does not "find" the target at 9 km or some of the larger ranges. The target is observable. Apparently, very bright single pixels distract also the larger blobs.
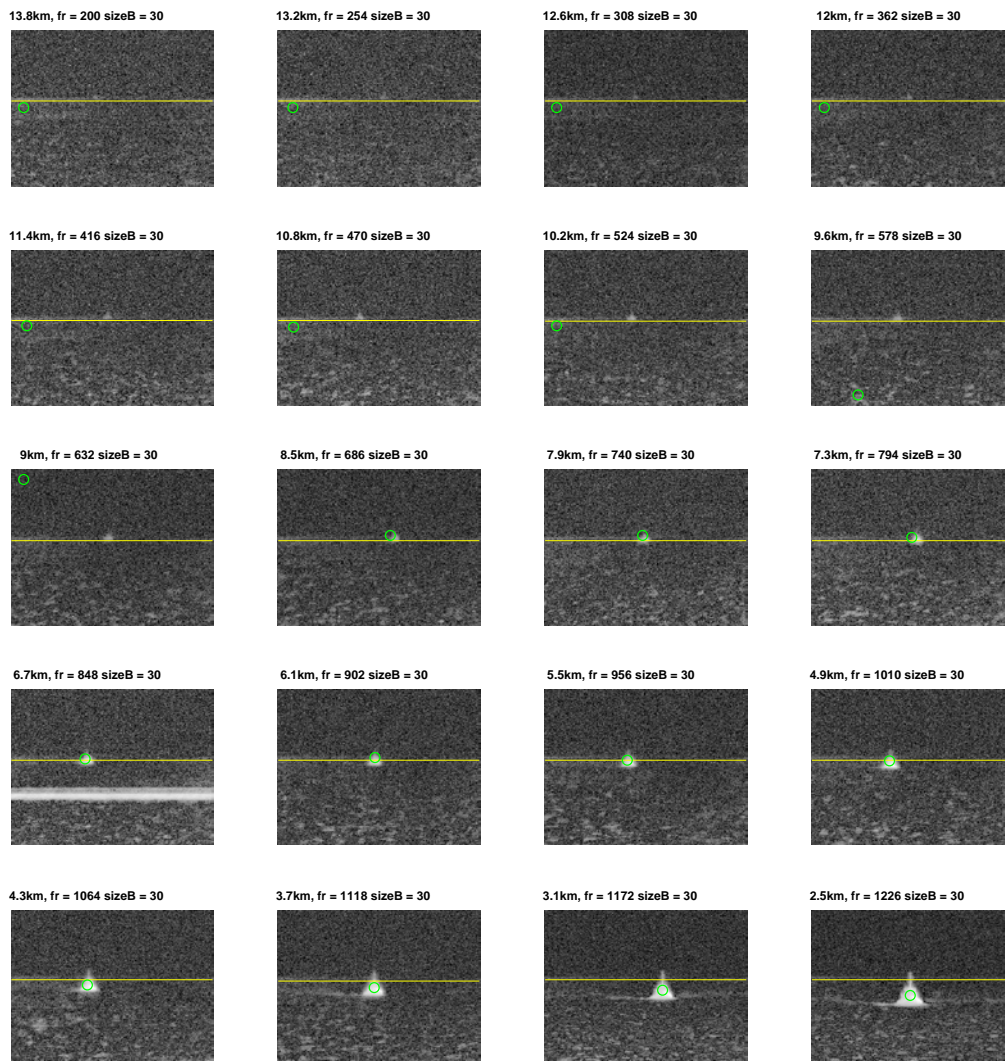

Figure 4. Frames, cropped to the "green" box in Fig 1 at constant range intervals of about 600 m. Overlaid are the location of the hottest 30x30 pixel blob in that frame (green circle) and the location of the horizon (yellow line). In this recording, the boat is inbound. Initially, the blob misses the target, but when it gets closer, closer than 9 km the blob tracks the target consistently. Here we estimate the detection range as 8.8 km (mean of 8.5 and 9 km). The range, frame number and blob size are indicated above each frame.

We observe in Fig. 3, that the two methods give about the same results when we choose an appropriate factor. As mentioned above, we are not so much interested in the numerical values, rather in the trend: the detection range increases with the blob size initially, and decreases again for large blob sizes. In this particular example, the maximum is rather broad, spanning a good range of blob sizes, from 4 to 20.

We have a limited number of other examples. There are several reasons for this. First, it is time-consuming to send out a boat to long range, traveling at relatively low speed (10 knots). Second, even though it was rather humid during the trial, limiting visibility in LW, MW is much less affected, and therefore even longer detection ranges are expected, requiring even longer boat runs. Third, a black body was mounted on the boat, visible from the sensor platforms on outbound trips, but not on inbound trips. This increases the visibility of the target, and makes the target more of a point target, on outbound runs, again requiring long boat runs before the boat "disappears". There are competing lens requirements: shorter focal length for detection run or long focal length to answer the question "how far can I see with the best lens?"

In Fig. 5 we collected results from analyses of other runs. Most recordings show the same behavior as LW1553, Fig. 4, a broad maximum for a range of blob sizes larger than 1. LW1935 shows the best detection ranges are obtained for very small blobs, 1x1 and 2x2. There is a local maximum for larger blob sizes. As before, the two methods yield comparable results.

For runs 2046 and 1575 we could not extract detection ranges in the MW band. In the case of 2046, the boat started its inbound run at 9 km and in run 1575 it turned at 15 km. In both cases, the boat was still "visible" in the MW at these maximum ranges. Run 2025 was stopped at 15 km range, where the boat was still "visible" for all blob sizes less than 25x25 pixels.

Runs 1553 and 2046 are inbound runs and the other runs are outbound runs. LW2046 and MW2046 were recorded with 50-mm lenses, the other recordings with 200-mm lenses.
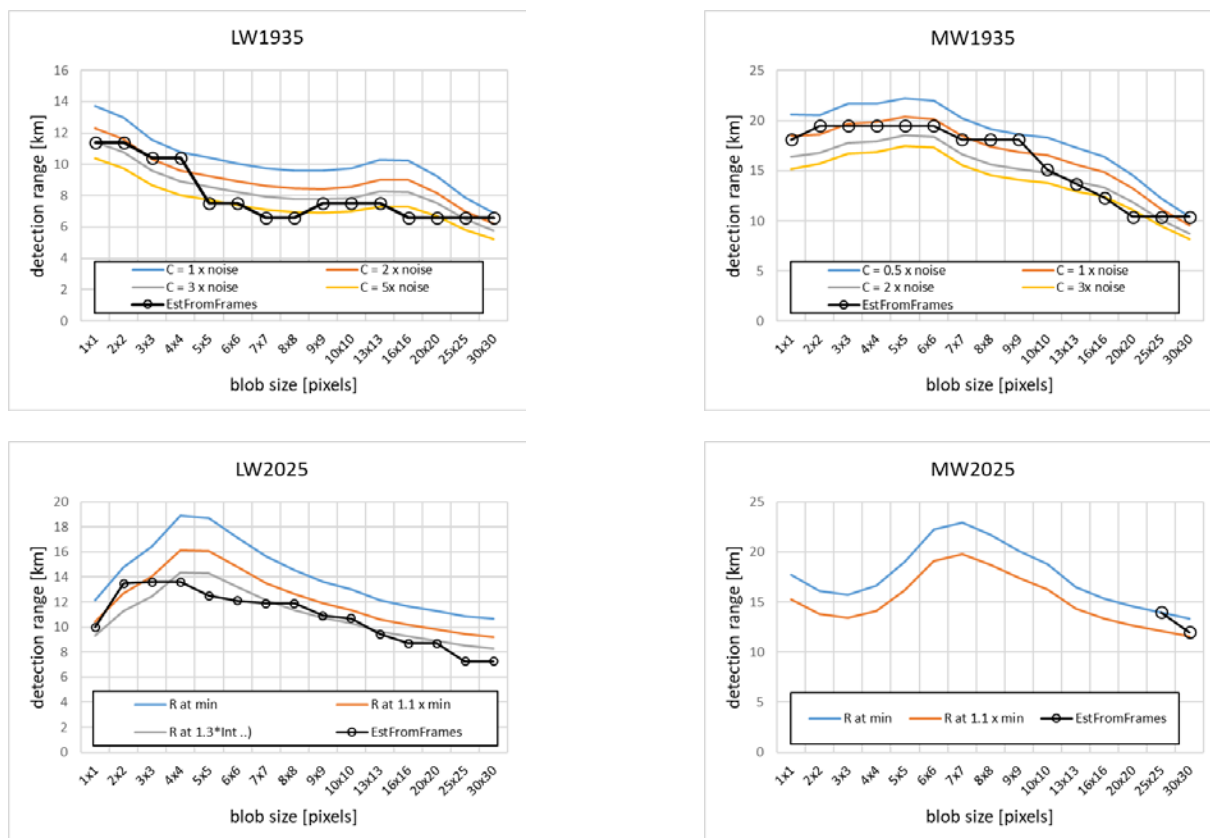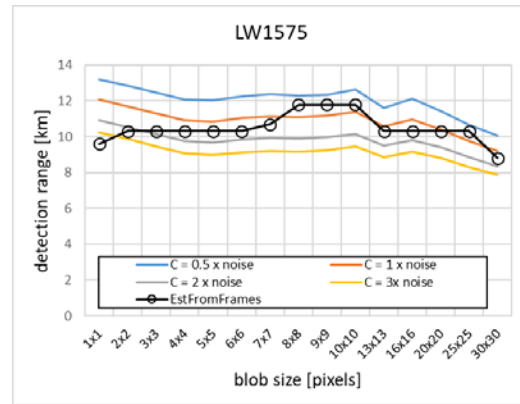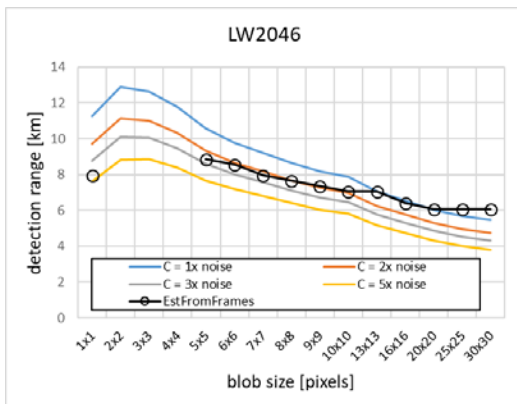
Figure 5. Detection range versus blob size for six different recordings. The colored curves are obtained from the intensity-range plots for each blob size, using different multiplying factors for the plateau noisiness as a threshold. The black circles and lines are detection range estimates obtained from visual inspection of a selection of frames, determining when (at what range) the different blobs stopped tracking the boat

Note 1: for the 2025-run, the boat sailed out to 15 km. In MW recording, all the blobs tracked the boat out to 15 km, except for the two largest blob sizes, 25x25 and 30x30 pixels. For those blob sizes we could estimate the detection range, but not for the other blob sizes. Therefore, only two range estimates from the frames are available.
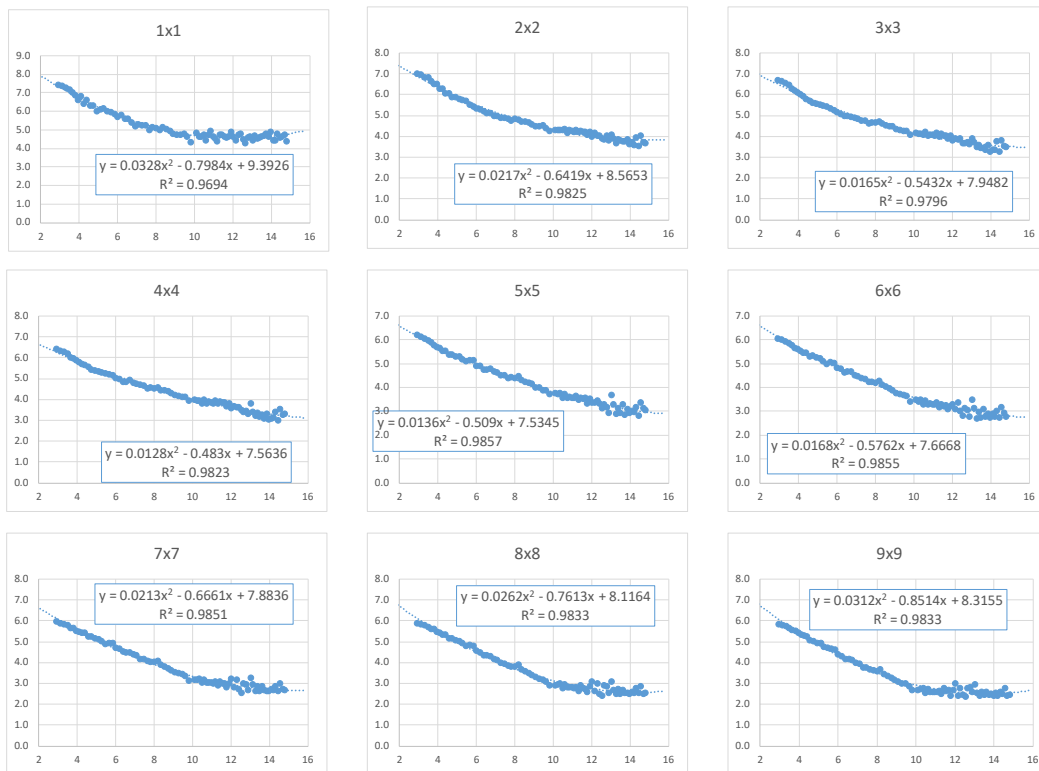

Figure 6. Illustration of how we determined detection ranges for LW2025. This approach was also used to analyze MW2025. The vertical axis is the natural logarithm of the blob intensity and the horizontal axis the range in km. The plot titles indicate the blob size. The R^2-number indicates the goodness of fit (ideally equal to 1). The coefficients are used to find the range at the minimum.

Note 2: The intensity-range curves, the equivalent of Fig. 2, for the 2025-runs did not show a clear plateau value for the smaller blob sizes. Without any further analysis, it was not possible to estimate for what range the decaying part of the

curve hits the plateau value. For these two recordings, LW2025 and MW2025, we chose to model the logarithm of the decaying intensity curve as a quadratic in the range, and assumed that the range for which the parabola reaches its minimum value is the detection range. There is no justification for using a quadratic other than that it is the simplest function that fits the data very well. We illustrate this method in Fig. 6. For larger blob sizes, larger than 7x7, we see a clear plateau value and there is no need to apply this approach. This does not take any uncertainty in the determination of the plateau level into account. Possibly, it would be more realistic to find the range at which the parabola hits 1.1 or 1.3 times the minimum value. Looking at the plot for LW2025 in Fig. 5, this method also seems to agree reasonably well with range estimates from visual inspections of frames. The shape of the curve is similar, but numerical agreement depends on the choice of the threshold.

3. MODEL

In this section, we build a very simple model that will support the observations we made in the previous section, Experiments. We model the boat as a square collection of pixels, N pixels on a side (say, N x N = 100 x 100 pixels), with a uniform intensity C (measured in LSB), placed in the center of a much larger square of pixels (say of size $2N$: 200 x 200 pixels) with zero intensity. Let us assume that the range at which we observe this configuration is 1 km. To simulate the target at larger ranges we step, not slide, a window of size M x M pixels over the original target scene and calculate the mean intensity in each of the new blobs. The new target is now of size (N/M) x (N/M) pixels. This new target is again placed in the middle of a new box of size $2N$ x $2N$ pixels with zero intensity. This is our new scene, now at range $R = M$ km. Note that at the edges the background intensity is mixed in with the target pixel intensity because of the larger footprint of the new pixels. By letting the value of M range from 2 to 25 in steps of 1 we simulate the target "sailing out" from 1 to 25 km.

For each of the synthetic scene images the pixel with the largest intensity is found and its intensity recorded. In addition, a blob of B x B pixels is slid, not stepped, over each scene image and the intensity of the hottest blob is recorded. We have chosen the following values for B: 1, 2, 3, …, 10, 13, 16, 20, 25, and 30, as in the experiment above.

Initially, we assume zero extinction, so that a reduction in blob intensity with range is purely geometric. A simple extension of this experiment is to introduce an extinction coefficient that is constant over the waveband one is interested in, assuming for instance good, medium, and poor atmospheric transmission.

Note that the maximum blob size (30 x 30 pixels) is chosen to be smaller than the size (100 x 100 pixels) of the original target at 1 km. With homogeneous target intensity, the maximum blob intensity for all blob sizes in the 1-km scene is still C. The intensity of the hottest blob remains constant with increasing range until the range has increased so much that the apparent target size (N/R) is less than the blob size B. This occurs when the range R exceeds N/B. At these larger ranges, there are now roughly $(N/R)^2$ pixels with intensity C, and $B^2 - (N/R)^2$ pixels with intensity zero. Hence, the mean blob intensity, neglecting smearing at the target edges, equals $C \cdot (N/R)^2/B^2$ in the hottest blob.

$$I_B\left(R\right) = \frac{N^2 C}{B^2 R^2} \tag{1}$$

So far, we have discussed the scene. Next, we will introduce the detection metric. This involves a basic sensor model, the Minimum Resolvable Temperature Difference function, and the Johnson criterion for detection. Since we assume there is no atmospheric extinction (this is not a limitation, only a simplification), the scene intensities, which are really radiances, are not attenuated as they propagate towards the sensor. We chose the background to have zero intensity; hence, the signals from the target are really contrasts, radiance contrasts. We assume here, that the sensor is calibrated, so that an arriving radiance contrast may be converted to an apparent temperature difference by a simple division by the sensor gain, G:

$$\Delta T_A = \frac{1}{G} I_B(R) = \frac{1}{G} \frac{N^2 C}{B^2 R^2} \tag{2}$$

The Minimum Resolvable Temperature Difference curve describes the trade-off between spatial resolution (level of detail) and the contrast in an image. The curve is a camera characteristic and depends among other things on the sensitivity and the detector pitch-optics combination. As an example, we use a model that has been fitted to data measured on both our LW and MW IrCam cameras[1]:

$$\Delta T_A = q \cdot \exp(p \cdot f_X) \tag{3}$$

Here, f_X is the spatial frequency measured in cycles/mrad and p and q are model parameters. Upon substitution of the expression (Eq. 3) for the apparent temperature difference into Eq. 2 and solving for the spatial frequency one obtains:

$$f_X = \frac{1}{p} \ln\left(\frac{1}{q} \frac{1}{G} \frac{N^2 C}{B^2 R^2} \right) \tag{4}$$

Note that this estimate of the spatial frequency, given in units of cycles/mrad, is found by finding the spatial frequency that corresponds with the apparent temperature difference on the MRTD curve. This is the straightforward Johnson approach.

The number of cycles on target, N_{cy}, is the product of the maximum spatial frequency (cycles/mrad) and the target size in mrad, $B \cdot IFOV$, where $IFOV$ is the instantaneous field of view, is defined in Eq. 5:

$$N_{cy} = f_X \cdot B \cdot IFOV = \frac{B \cdot IFOV}{p} \ln\left(\frac{1}{q} \frac{1}{G} \frac{N^2 C}{B^2 R^2} \right) \tag{5}$$

According to Johnson[2], detection requires 1 cycle on target, but let us denote this criterion as N_D, the number of cycles required for detection. Solving Eq. 5 for the detection range R_D we obtain:

$$R_D = \frac{1}{B} \sqrt{\frac{N^2 C}{q \cdot G}} \exp\left(-\frac{p \cdot N_D}{2 \cdot B \cdot IFOV} \right) \tag{6}$$

Inspection of Eq. 6 reveals that the first factor on the right hand size suggests a decrease of the range with blob size, whereas the exponent actually increases with blob size B. The two competing factors suggest that there must be an extreme value for R as a function of B. It can be shown that R reaches a maximum when the blob size is chosen to be equal to: $p \cdot N_D / (2 \cdot IFOV)$. This observation would suggest there is an optimum blob size for a given target task (N_D) with a given sensor (p, $IFOV$).

Under the assumption that atmospheric extinction may be accounted for by a simple Beer-like expression, i.e. a factor $exp(-\alpha \cdot R)$ for the IR band, the observed blob intensity will be a slight modification of Eq. 1:

$$I_B(R) = \frac{N^2 C}{B^2 R^2} \exp(-\alpha R) \tag{7}$$

Carrying through the procedure described above we find a simple modification of Eq. 5:

$$N_{cy} = \frac{B \cdot IFOV}{p} \ln\left(\frac{1}{q} \frac{1}{G} \frac{N^2 C}{B^2 R^2} \exp(-\alpha R) \right) \tag{8}$$

It is not possible to solve this equation analytically for R. However, we observe that $exp(-\alpha R)$ is a monotonically decreasing function of R, just as the factor $1/R^2$ and therefore does not fundamentally change the behavior. The detection ranges will still have a maximum, as a function of the blob size. Hence, there is a best blob size for detection.

Using typical values for the parameters p, N_D, and $IFOV$ (0.5 mrad/cycle, 1 cycles on target, 0.075 mrad/pixel) we get an optimal blob size for detection between 3 and 4 pixels. This number is not unlike the optimal blob size observed in Figs. 3 and 5.

4. SUMMARY

We analyzed LW and MW recordings made during the September 2016 Minotauros trial. In particular, we estimated detection ranges from recordings of an inbound or outbound small boat, sailing a straight course. In addition to looking for the hottest pixel in each frame and recording its intensity and location in the frame, we looked for the hottest blobs of pixels in each frame, with a blob defined here as contiguous square group of pixels, measuring BxB pixels, with B = 2, 3, …, 10, 13, 16, 20, 25, 30. From the intensity versus range curves for each blob, we estimated a detection range, defined as the range where the decaying intensity reaches a background plateau value. As a separate check, we visually inspected a selection of frames on which the location of the hottest blob was overlaid, checking whether the blob tracked the boat. The two methods both showed that there is an optimum blob size, for which the detection range is maximal. This blob size is not necessarily equal to 1x1 pixel, rather, depending on the recording, somewhere between 3 and 10 pixels.

When the boat in run 1553, see Fig. 4, is at range 10.2 km, for instance, its image roughly consists of 20 pixels (3 wide, 7 tall), i.e. is an extended target. This implies that even at the longest range, 13.8 km, the target is extended, even though its pixel count is halved. This range is beyond the detection ranges we estimated. This is typical for the recordings: when the target disappears, it is still an extended target. Run 1553 is an inbound run, i.e. without a visible black-body source.

For outbound runs, when the black body is visible, one would expect hot-spot detection to be favored over blob detection. The limited size black body (20 cm in diameter) quickly reduces to a single pixel at range 2.5 km for the longest focal length used here. Analysis of run 1935 seems to support this conjecture, as the best detection ranges are obtained for the smallest blob sizes. Analysis of the other outbound runs, 1575 and 2025 do not show this clear tendency.

We developed a simple model in which we implemented a scene with a square target, simulating a receding target by resampling the close-range scene. We modeled the sensor behavior with the minimum resolvable temperature curve, a curve that relates spatial resolution to sensitivity. Finally, we used Johnson's detection criterion to calculate the detection range for different blob sizes. We developed an analytical expression for this range, under some simplifying assumptions, which shows that the detection range has an absolute maximum value for a particular blob size, $p \cdot N_D/(2 \cdot IFOV)$. The model supports the observations that we made.

One of the simplifying assumptions was a zero-extinction atmosphere. We showed that inclusion of a Beer-like extinction is straightforward and does not alter the fundamental behavior, although it may have an effect on numerical values.

Preliminary experimenting with scene generation, where we included non-uniform target intensity and noise in the scene, did not show fundamentally different behavior, suggesting that although the model is rather simple, it seems to include the most important elements to describe the detection problem discussed here.

We showed that for reasonable values of the parameters, the optimal blob size is numerically not unlike the optimum blob sizes we found in our analyses of the recordings.

In this paper, we considered square blobs of pixels. Curves of detection range versus blob size vary in shape, even though weather conditions during the recordings were rather similar: 25 ºC and 60-70% relative humidity. Possibly, blobs with different aspect ratios, ones that better match the target under consideration, could show different detection-range versus blob-size curves, with wider or narrower maxima. This could be investigated.

REFERENCES

[1] van Rheenen, A. D., P. Taule, J. B. Thomassen, and E. Blix Madsen, "MRTD- man versus machine, Proc. SPIE 10625, 106250N (2018)

[2] Sjaardema, T. A., C. S. Smith, and G. C. Birch, "History and evolution of the Johnson criteria", Sandia Report SAND2015-6368