An insect-inspired detection algorithm for aerial drone detection

John V. James, Benjamin S. Cazzolato and **Steven Grainger**
School of Mechanical Engineering
University of Adelaide, Australia
{john.james, benjamin.cazzolato, steven.grainger}@adelaide.edu.au

David C. O'Carroll
Department of Biology
Lund University, Sweden
david.ocarroll@biol.lu.se

Steven D. Wiederman
Adelaide Medical School
University of Adelaide, Australia
steven.wiederman@adelaide.edu.au

Abstract

Flying insects display remarkable visual abilities which can serve as inspiration for target detection algorithms. For example, biologically inspired models of neurons in the dragonfly brain have previously been evaluated for their ability to detect moving targets from a ground-based robotic platform. Here we evaluate the performance of a model tuned to a specific biological task (detection of prey and conspecifics), applied to a very different endeavour - the detection of airborne drones as viewed by another drone. We identify the circumstances for success and failure with the prey-detection model, whilst varying the computational complexity of the task with a reduction in spatial sampling. This evaluation provides substantial insight into how the bio-inspired algorithm may be generalised for other automatic target detection tasks, such as in the detection of unmanned aerial vehicles.

1 Introduction

Flying insects have remarkable visual abilities. Using low resolution, blurry compound eyes, they can identify behaviourally relevant features in real time. A prime example of insect visual prowess is the dragonfly. Dragonflies have small brains of about one million neurons and measuring about 2mm across. Despite these relatively limited resources, they can engage in highly dynamic pursuits of prey with high success in complex environments [Olberg *et al.*, 2000]. The algorithm used by the dragonflies to identify the position of prey is effective, efficient and compatible with low resolution and blurry input optics. Such processing seems ideal for mobile robotic applications, where computational resources are at a premium, especially in airborne applications.

Neurons which respond selectively to small moving targets, named small target motion detectors (STMDs), have been discovered through electrophysiological investigation of the visual systems of various insect species including the dragonfly [O'Carroll, 1993; Nordström and O'Carroll, 2006; Keleş and Frye, 2017]. Dragonfly STMDs have broad velocity tuning and respond robustly to moving targets even against moving, cluttered backgrounds [Wiederman and O'Carroll, 2011; Nordström *et al.*, 2006]. The response characteristics of these STMDs have been modelled using an elementary STMD (ESTMD) model [Wiederman *et al.*, 2008]. In prior studies, the computational efficiency and effectiveness of this ESTMD model in detecting small targets during closed loop pursuits was evaluated using both computer simulations and when deployed on a ground-based robotic platform [Bagheri *et al.*, 2017a; 2017b]. However, this prior work was limited to essentially planar pursuits in a ground-based environment.

In the present study, we evaluated the ESTMD model on a dataset comprised of video footage of an airborne, moving drone with video footage captured by another drone. An important characteristic of the ESTMD model is that like the dragonfly it emulates, the processing operates on a low resolution, blurred version of the input image. We varied the spatial resolution, acuity and size tuning of the model to quantify the trade-off between computational cost and detection effectiveness. We also explored several means of improving detection performance. Our results provide insight into the current feasibility and limitations of this approach and identifies avenues for developing a bio-inspired target detection algorithm for the detection of aerial drones.

2 Methods

2.1 Video Capture

Five videos were captured from an airborne Phantom 4 Pro drone of another Phantom 4 Pro drone during a

mostly clear afternoon with patchy cloud. The surrounding area included a mixture of hillside, distant cityscape, parklands and low-lying suburbs. Four videos were captured at 1080p at 120 frames per second and one at 720p at 30 frames per second. Two video sequences were recorded from a stationary drone at heights of either 6m or 12m. The other three videos contained substantial amounts of ego-motion as the observing drone engaged in dog-fight pursuits. For analysis, the field of view extended 84° diagonally, leading to 73.7° horizontal and 45.7° vertical fields of view (16:9 aspect ratio). In total, over 30 minutes of video footage served as inputs to our modelling efforts, consisting of 170,551 frames in total with 57,251 frames corresponding to the stationary and pursuit experiments where the drone was identified in the field of view. We purposely designed experiments to be challenging, by moving the target drone (1) at varying distances, (2) translated at varying velocities, (3) against cluttered backgrounds of cloud and landscape scenery and (4) in the presence of moving distracters (e.g. cars, pedestrians, etc.).

2.2 Computational model

The ESTMD model was implemented as a discrete time model with a 1 ms step size. Figure 1 shows an overview of the model. Additional details of the model implementation can be found in [Bagheri *et al.*, 2017a]. Insect optics were represented by subsampling the image after blurring with a gaussian kernel. Both the distance between subsampled pixels and the blur size were varied. The ESTMD model is monochromatic, using the green channel (recreated from the Bayer mosaic) to emulate the spectral sensitivity of insect vision.

The spatiotemporally filtered signal is then separated into two independent, rectified channels for light increments (ON) and light decrements (OFF). The ON and OFF channels are separately filtered with a non-linear temporal filter implemented as:

$$Y_k = u_k - \gamma_{k-1} \tag{1}$$

where u_k is the current input, Y_k is the current output, γ_k is a state variable of the filter given by:

$$\gamma_k = \frac{t_s}{t_s + \tau_k} u_k + \frac{\tau_k}{t_s + \tau_k} \gamma_{k-1} \tag{2}$$

where $t_s = 1ms$ is the model sample time and τ_k is given by:

$$\tau_k = \begin{cases} 5ms, & u_k - \gamma_{k-1} \geq 0 \\ 350ms, & otherwise \end{cases} \tag{3}$$

This temporal filtering represents neuronal adaptation, observed in putative elements of the target-detection pathway. Size selectivity is established by suppressing responses to large features, such as long edges or bars. To implement this tuning, the separate channels are spatially filtered by convolution with a square kernel with strong negative gain. The square kernel of the centre-surround spatial filter, K, was defined by:

$$\begin{aligned} K_{1\ldots n,1} &= -a \\ K_{1\ldots n,n} &= -a \\ K_{1,1\ldots n} &= -a \\ K_{n,1\ldots n} &= -a \\ K_{\frac{n-1}{2},\frac{n-1}{2}} &= 2 \\ K &= 0 \quad elsewhere \end{aligned} \tag{4}$$

where n is the length of the kernel edges and $a = 4/(n-1)$. The size of this kernel was varied between trials to tune the model to targets of different sizes.

One of the channels is then delayed relative to the other using a low pass filter:

$$H(z) = \frac{\frac{1}{51}z + \frac{1}{51}}{z + \frac{49}{51}} \tag{5}$$

To create a matched filter for the spatiotemporal signature of a small, moving target, the ON and OFF channels are multiplied together to produce the final output.

To match the framerate of the input video with the higher update rate of the model, intermediate frames were created by linearly interpolating between frames of the source video. Model outputs between frames were then summed to produce an output at the same framerate as the input video. To ensure that the model outputs were not affected by start-up transients, the first 150 frames were discarded. Three variants of the model were tested: one which correlated the undelayed ON channel with the delayed OFF channel and which responded preferentially to dark targets; one which correlated the delayed ON channel with the undelayed OFF channel to respond preferentially to light targets; and, another which combined the outputs of the light and dark variants leading to sensitivity to both polarities. The temporal low-pass filter in Figure 1 was in the position shown for the variant responsive to dark targets, whereas for light targets this filter was moved to the opposite channel.

2.3 Ground truth

Ground truth values for the position of the drone in the video were prepared by manually selecting the location of the drone at 200ms intervals based on the full resolution imagery. Linear interpolation was then used to fill in the rest of the frames. When testing a range of subsampling resolutions, the target may disappear near the edges of the visual field. Therefore to ensure a fair comparison across varying degrees of subsampling, the ground truth was limited to only those frames where the drone was

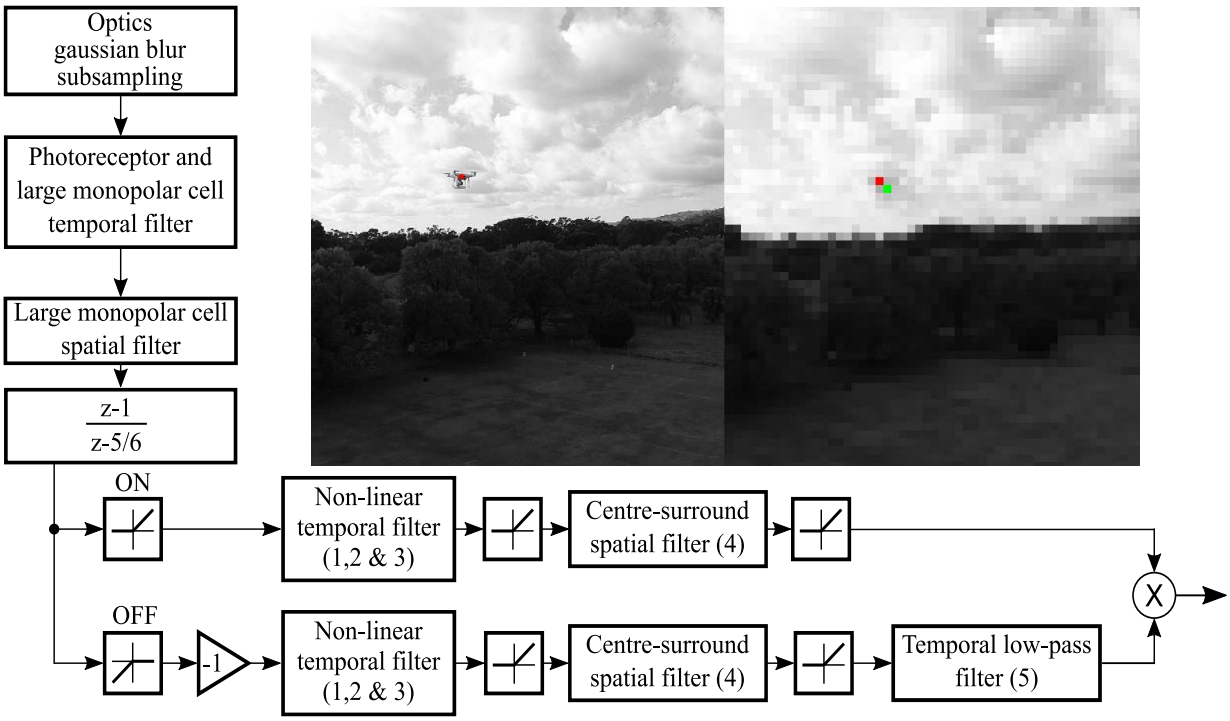

Figure 1: Model overview and example images. The left image shows an example of a high resolution input video frame with the ground truth location marked in red. The right image shows an example subsampled and blurred input image (0.69° per model unit, 0.55° gaussian half-width blur) with red and green pixels indicating the ground truth and the location of the maximum model output on this frame respectively. Model processing begins with linear temporal and spatial filtering followed by rectification into separate channels of opposite polarity, non-linear temporal filtering of each channel and lastly correlation between the two opposite polarity channels. The model variant shown here is sensitive to dark targets; the low-pass filter on the OFF channel is used to delay that channel relative to the ON channel prior to correlation. The model thus responds most strongly to light increments which are preceded by light decrements, as may result from a moving dark target.

visible at all subsampling resolutions tested. The speed of the drone in the image plane, v, on each frame was estimated from the ground truth using:

$$v = \frac{\sqrt{(x_t - x_{t-1})^2 + (y_t - y_{t-1})^2}}{\Delta t} \quad (6)$$

where x_t, x_{t-1}, y_t and y_{t-1} are respectively the current and previous horizontal and vertical coordinates of the drone and Δt is the time between frames. Ground truth values for the size of the drone were created by manually drawing a rectangular bounding box around the drone at 600ms intervals. The drone size was taken to be equal to the length of the diagonal of the bounding box. Linear interpolation was used for intermediate frames.

2.4 Performance evaluation

We evaluate the performance metric entitled 'Recall', by dividing the number of frames on which the drone was detected, at varying levels of precision, by the total number of frames on which the target drone was visible. When the spacing between model units was small ($\leq 0.613°$), the target drone was taken to be detected if the frame's maximum model output was within a 1° radius of the ground truth location. With wider model unit spacing, we considered the target drone detected if it was within a 3x3 pixel region in the subsampled image centred on the ground truth location. While this resulted in regions of different sizes, it ensured that responses were not missed due to delays inherent in the model processing. These two analysis regimes are referred to as High resolution and Low resolution respectively. Frames where the drone was not visible were excluded from the analysis.

3 Results

The model performance was evaluated for two different groups of parameters, (1) High resolution and (2) Low resolution. Low resolution settings are intended to reflect more physiologically realistic values. The values of model sampling resolution, degree of optical blur and the size of the centre-surround kernel that is included in

each of these groups is defined in Table 1.

3.1 Comparison of dark, light and combined model variants

Combining outputs from the light and dark model variants provided a small, overall improvement in Recall at both High and Low resolution in most conditions. Figure 2 shows the Recall for each variant when using the best combination of parameters from Table 1 for each of the five videos. We used the best combination of parameters for each video, rather than a single set of parameters for all videos, to strike a balance between the models having inappropriate size tuning for a particular video on the one hand and changing parameters with infeasible frequency (e.g. per frame) on the other. Improvements in performance between the variants was more pronounced in the Low resolution group of parameter variations. Interestingly, with parameters in the High resolution group, the dark selective variant performed better in videos with significant ego-motion, suggesting that on average the light selective variant contained more responses to background than to the target in these cases. This is a surprising outcome as the drone was white in colour. Because of the similarity between model variants, the remainder of the results presented in this paper are limited to the combined variant.

3.2 Image characteristics affecting performance for the combined model variant

As expected for a model tuned to the size and velocity of a moving target, Recall was influenced by the apparent speed and size of the drone in the image plane. Figure 3 shows the performance across all frames in the four 1080p videos in three apparent size ranges. The Recall revealed the velocity tuning of the system, where in the Low models (A-C) performance decreased at either low or high target speeds (across all target sizes tested). As expected, the models struggled to detect the moving target when it was in the lower ranges of apparent target size (A,D). Recall in the Low resolution models was dramatically improved across the entire speed range when the target size was within the larger range 2.9°-4.3°(C). In the High resolution models, velocity tuning was less pronounced, especially in the mid and large target size ranges (E,F). Recall was greater for larger apparent sizes across a wider range of apparent speeds. At High resolution, the model produced responses to artefactual flicker resulting from rotation of the drone's rotors, which locally give the appearance of rapid motion due to successive increments and decrements in luminance, which may contribute to the relatively high detection performance for low speeds.

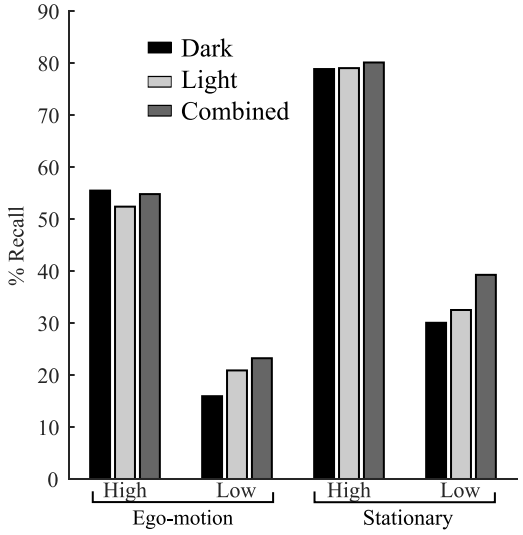

Figure 2: Comparison of model variants for High resolution and Low resolution. The three videos featuring substantial ego-motion (i.e. movement of the pursuing drone) showed lower performance compared to the two videos where the observer drone was stationary. Recall shown is the average achieved in all videos in the two classes if using parameters from Table 1 which give the best recall for each separate video, i.e. one set of parameters from the High resolution set and one from the Low resolution set for each of the five videos.

3.3 Effect of subsampling and blur on performance of the combined model variant

For detection in mobile robotic applications, it is desirable to minimise the complexity of the detection algorithm to minimise power and weight. We therefore tested how reducing complexity affected Recall by increasing spacing between model units and the strength of Gaussian blur of the image. Reducing spatial resolution reduced detection performance, but this was offset by a decrease in computational demands (Figure 4A). Note the quadratic reduction in complexity compared to the linearly decreasing performance. Therefore, if a lower Recall is feasible, this can have huge impacts in computational complexity. At fine resolutions, Recall was greatest with no blur. As the subsampling of the image becomes coarser, some blurring of the input image results in improved Recall (Figure 4A,B).

3.4 Performance of the combined model variant with reduced precision

One possible implementation of this algorithm is to identify segments of the image for further processing. Here, it is not only the maximum value but rather the N top

Table 1: Model parameters. High resolution settings comprised each combination of the parameters in the first row. Similarly for Low resolution settings.

Category	Spacing between model units (degrees)	Blur kernel size (degrees)	Centre-surround kernel size (model units)
High resolution	0.038° -0.613°	0 - 3.070°	3 - 57
Low resolution	0.691° - 3.454°	0.614° - 3.070°	3 - 13

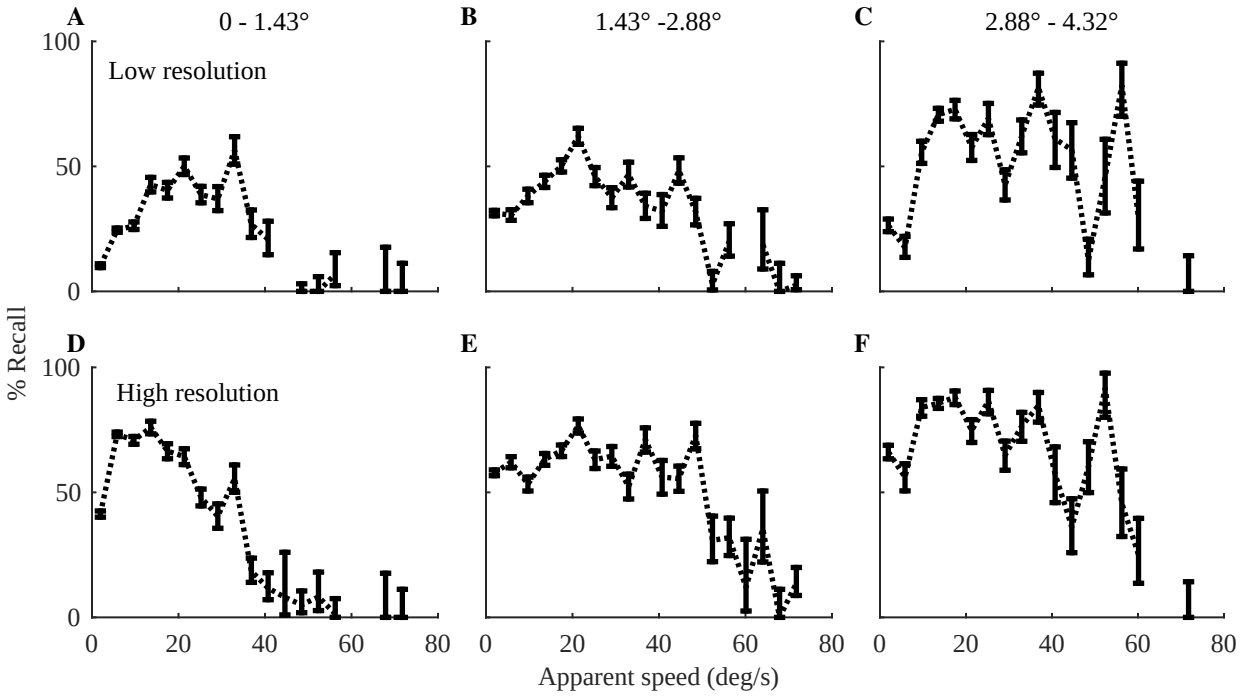

Figure 3: Recall for the combined model for different apparent sizes and velocities of the target for Low, A-C, or High, D-F, resolution. Separate panels correspond to apparent drone size of A) 0 to 1.43°; B) 1.43° to 2.88°; and, C) 2.88° to 4.32°. The stated upper limit of each range is excluded from the range. D-F have matching drone size ranges. Bins where the 95% confidence interval for the detection rate was wider than 30% are omitted. Error bars show the 95% confidence interval for the Recall assuming a Bernoulli distribution.

values that could subserve further processing. We therefore examined whether the target was contained within the top 10 values (Figure 5A). Cumulative recall can be increased significantly by including additional ranks. However, the prominence of the contribution of rank 1 compared to the others indicates that generally if the model output associated with the drone is not the greatest output in the scene, it is overwhelmed by many other strong responses to the background. Continuing to include additional ranks further improves Recall in both High resolution and Low resolution settings, albeit at different rates (Figure 5B). The benefit of including additional ranks quickly diminishes and it is not feasible to achieve 100% recall in this way; even including the top 1,000 ranks achieves only approximately 90% recall in this case. For 0.614° spacing, 1,000 ranks covers more than 12% of the visual field.

3.5 Presence of strong distracters during camera movement

Targets may be missed as a result of low response to the target itself, leading to false negatives, or due to overwhelming responses to background features, leading to false positives. We sought to determine whether false positives contributed to the lower performance achieved in videos with ego-motion versus those with a stationary observing drone. We examined the responses to background features during movements which led to the target response not being maximal, such as gradual panning

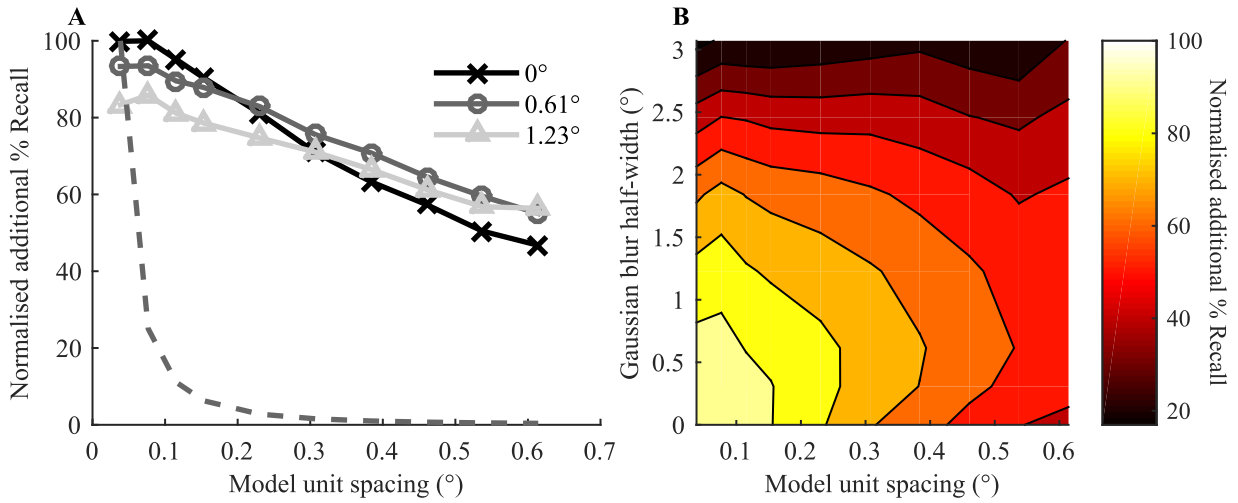

Figure 4: Effect of spacing between model units and blur. The 'additional Recall' shown here is the Recall achieved across all frames in the four 1080p videos less the Recall which would be achieved through random guessing. Data shown is normalised by dividing the additional Recall for a given blur and sub-sampling by the maximum additional Recall achieved for any blur and sub-sampling combination tested. A) Normalised additional Recall as a function of spacing between model units. Solid lines show normalised additional Recall for different Gaussian blur half-widths. The dashed line corresponds to the normalised number of model units in the scene, corresponding to computational demands. B) Normalised additional recall for different spacing and blur half-widths.

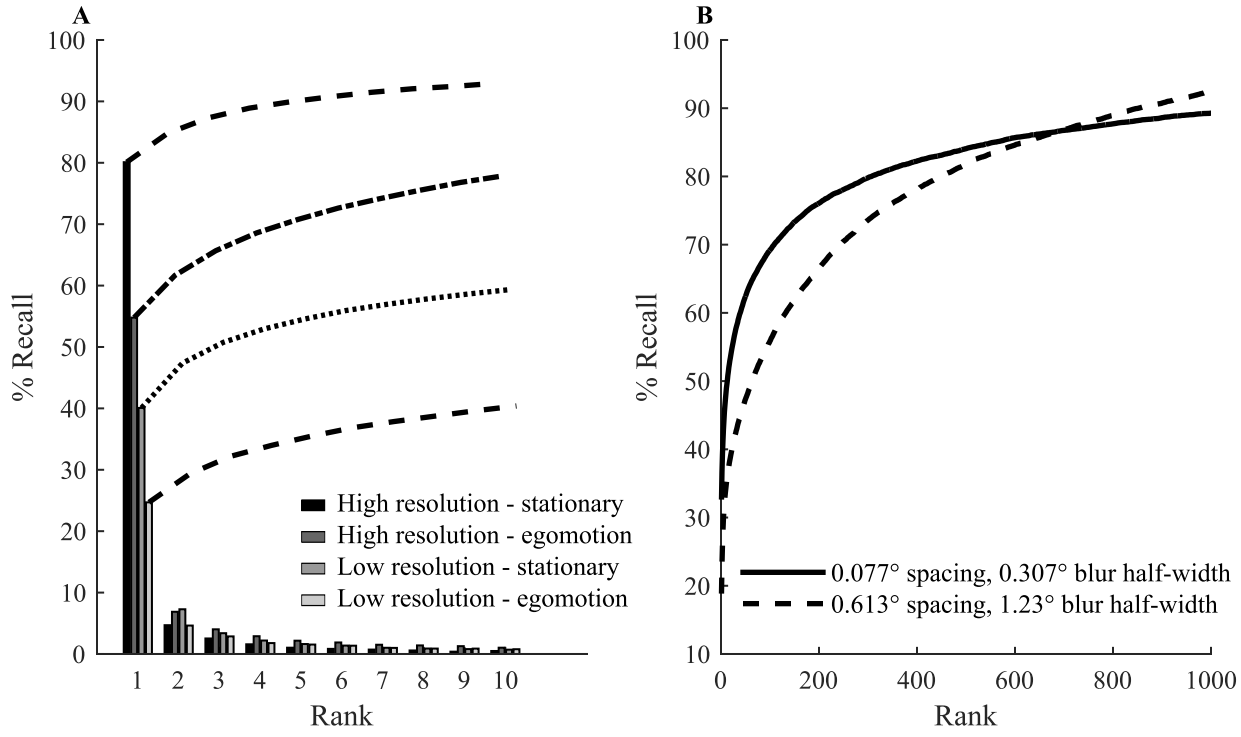

Figure 5: Inclusion of additional ranks of output per frame. A) Bars indicate the incremental improvement in recall for including additional ranks. Lines show the cumulative recall as additional ranks are included. B) Recall achievable by including many additional ranks of output on a video with significant ego-motion.

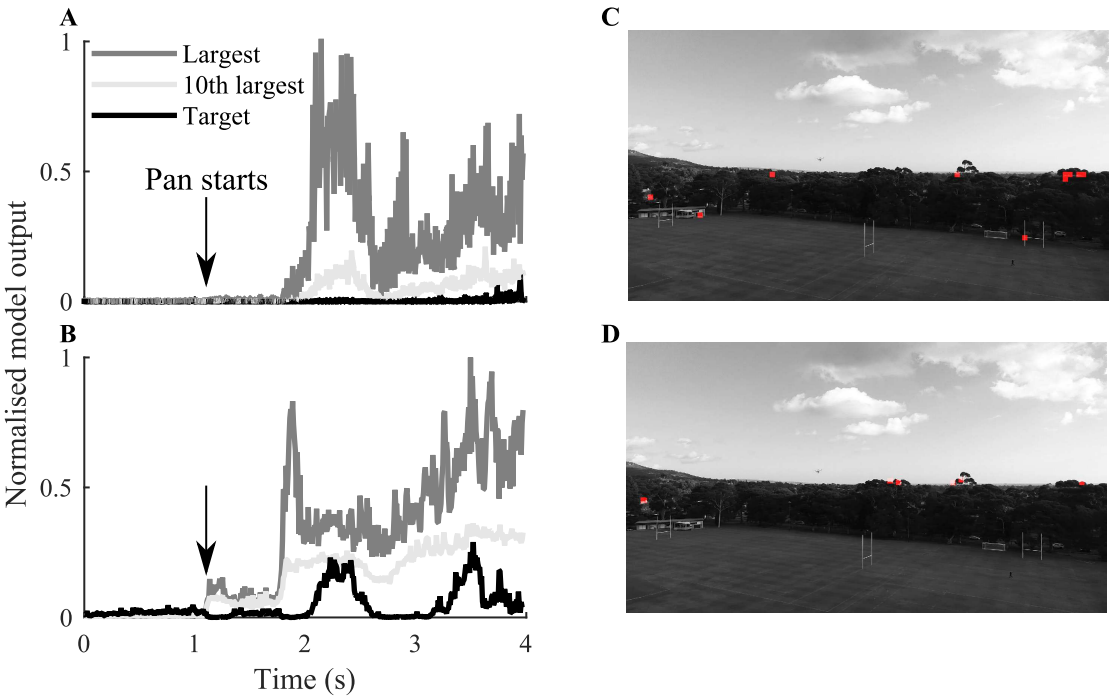

Figure 6: Example ESTMD model output in response to video segments that include strong distracting features. A) Trace of model outputs for Low resolution settings for part of a video where the observer and target drone are initially stationary until shortly after 1 second. After that time, the target drone begins moving and the observer pans to track it. Panning becoming swifter around 2 seconds. Lines show the largest response on each frame, 10^{th} largest and the responses generated by the target. B) Location of top 10 model outputs on a frame during panning matching the trace in (A). (C) and (D) are equivalents of (A) and (B) but for High resolution settings. Note the difference in false positive locations between (B) and (D).

movements. During panning movements, background features produce strong distracting responses for both High and Low resolution models (Figure 6A,C).

Examples of such features are regions with high local contrast, such as tops of trees viewed against sky (Figure 6B,D). Although a substantial response to the target may be generated, this is overwhelmed by many distracting background features as shown by the target response in Figure 6C falling below the 10^{th} largest non-target response for most of the time window following the commencement of a panning movement.

3.6 Use of thresholds and the maximum operation to improve precision

For the detection events produced by the algorithm to be useful, these must be sufficiently reliable indicators of the presence of a target. In other words, there must not be excessive numbers of false positives. This is not reflected in the Recall metric. We investigated whether a threshold can be used to reduce the number of false positives. Here we define target responses to be the maximum model output within the region described in 'Performance Evaluation' under Methods on each frame. Outputs outside this region are defined as background responses. For this analysis, we included frames of the videos where no drone was present and so only background responses were present. Applying a small threshold to the model outputs excludes a high proportion of background responses without removing a significant number of target responses (Figure 7A). In the example shown in Figure 7, almost all target responses were above a threshold of about 10^{-7} whereas almost all background responses were not, which is reflected in the conditional probabilities of responses being over the threshold given that they are a target response or a background response. If considering only the maximum response on each frame, then applying a threshold excludes background and target responses at a much more similar rate (Figure 7B). This is reflected in similar conditional probabilities of these responses being over a threshold given that they are the maximum response on a frame and are either a target response or background response. Without applying a maximum operation, the probability of a given output on a subsampled pixel corresponding to a target

when considering all responses on each frame is very low for each threshold tested (Figure 7C). Even at thresholds which exclude the majority of target responses, model outputs are not reliable indicators of the presence of a target: compare 7A and 7C at high thresholds. If only considering responses which are the maximum on each frame, detection outputs are much more reliable indicators of the presence of a target: compare 7C and 7D. Combining the maximum operation with a threshold can further increase reliability, as shown by the solid line in 7D moving above the dotted line for a range of thresholds.

4 Discussion

The ESTMD model achieved much more reliable detection when the observer was stationary than when there was significant ego-motion. Unsurprisingly, the complex outdoor setting provided many spatial features which, when viewed by a camera which is smoothly changing its direction, can match the spatiotemporal profile that ESTMD is responsive to. The observing drone in this study tended to engage in long, relatively gradual panning motions, which created substantial periods of large-scale background motion. Conversely, insects use various strategies to limit apparent background motion. Examples of these include using short, rapid changes of gaze rather than slow panning movements [van Hateren and Schilstra, 1999; Geurten *et al.*, 2010; Lin and Leonardo, 2017]; using sideways translational movements in place of rotations [Olberg *et al.*, 2007; Kassner and Ribak, 2018]; and, predictive tracking [Mischiati *et al.*, 2015]. These strategies may reduce the prevalence of false positives. Object detection in insects is also influenced by behavioural state [Maimon *et al.*, 2009] and differences in optic flow [Zabala *et al.*, 2012]. The algorithm described in this paper did not make use of these potential additional sources of information.

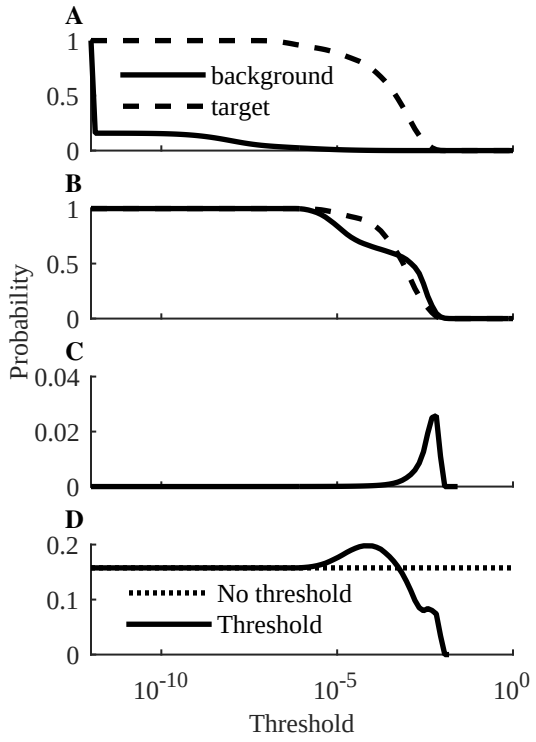

Figure 7: An example of the effect of using thresholds and a maximum operation on precision for a video with ego-motion and High resolution. A) The probability of a response being over threshold given that it is either a target or background response. B) The probability of the maximum value for the frame being over threshold given that it either a target or background response. C) The probability of a response that is over threshold corresponding to a target. D) Probability of the maximum response for the frame corresponding to a target, with or without also applying a threshold.

In its current form, the ESTMD model would not selectively respond to target drones in the visual environment due to the low precision. However, it is noteworthy that under Low resolution settings, the model discarded more than 99.8% of the incoming image information before the first temporal filter. This form of modelling therefore, is likely well suited as a pre-filter, removing redundant information before further feature identification methods can be applied. Detection methods which are computationally efficient may be useful to generate proposals for more reliable yet expensive methods [Ren *et al.*, 2017]. It must be noted that we made no attempt to tune the bio-inspired algorithm for the sizes and velocities expected with spatiotemporal signatures of the target drone. Moreover, no attempt was made here to adjust the velocity tuning of the model to suit the movement profile of the drone. Our results indicated that performance may be improved by combining parallel pathways tuned to either light or dark targets. Future research will investigate how the application of different tuning techniques, as well as the implementation of pursuit strategies that minimise ego-motion, improve the ability to detect moving drones in complex visual environments.

Acknowledgments

This work was supported with supercomputing resources provided by the Phoenix HPC service at the University of Adelaide. This research received funding from Australian Research Council DECRA (DE150100548), Australian Research Council Future Fellowship (FF180100466) and and the Swedish Research Council (VR 2014-4904). The authors wish to

thank the Adelaide Botanic Gardens and the University of Adelaide Unmanned Research Aircraft Facility (URAF) for supporting this research.

References

[Bagheri *et al.*, 2017a] Zahra M. Bagheri, Benjamin S. Cazzolato, Steven Grainger, David C. O'Carroll, and Steven D. Wiederman. An autonomous robot inspired by insect neurophysiology pursues moving features in natural environments. *Journal of Neural Engineering*, 14:046030, August 2017.

[Bagheri *et al.*, 2017b] Zahra M. Bagheri, Steven D. Wiederman, Benjamin S. Cazzolato, and David C. O'Carroll. Performance of an insect-inspired target tracker in natural conditions. *Bioinspiration and Biomimetics*, 12:025006, February 2017.

[Geurten *et al.*, 2010] Bart R.H. Geurten, Elke Braun, and Martin Egelhaaf. A syntax of hoverfly flight prototypes. *Journal of Experimental Biology*, 213:2461-2475, July 2010.

[Kassner and Ribak, 2018] Ziv Kassner and Gal Ribak. Role of side-slip flight in target pursuit: blue-tailed damselflies (*Ischnura elegans*) avoid body rotation while approaching a moving perch. *Journal of Comparative Physiology A*, 204:561-577, June 2018.

[Keleş and Frye, 2017] Mehmet F. Keleş and Mark A. Frye. Object-Detecting Neurons in Drosophila. *Current Biology*, 27:680-687, March 2017.

[Lin and Leonardo, 2017] Huai-Ti Lin and Anthony Leonardo. Heuristic Rules Underlying Dragonfly Prey Selection and Interception. *Current Biology*, 27:1124-1137, April 2017.

[Maimon *et al.*, 2009] Gaby Maimon, Andrew D. Straw, and Michael H. Dickinson. Active flight increases the gain of visual motion processing in Drosophila. *Nature Neuroscience*, 13(3):393-399, February 2010.

[Mischiati *et al.*, 2015] Matteo Mischiati, Huai-Ti Lin, Paul Herold, Elliot Imler, Robert Olberg, and Anthony Leonardo. Internal models direct dragonfly interception steering. *Nature*, 517:393-U410, January 2015.

[Nordström *et al.*, 2006] Karin Nordström, Paul D. Barnett, and David C. O'Carroll. Insect detection of small targets moving in visual clutter. *PLoS Biology*, 4(3):378-386, February 2006.

[Nordström and O'Carroll, 2006] Karin Nordström and David C. O'Carroll. Small object detection neurons in female hoverflies. *Proceedings of the Royal Society B*, 273:1211-1216, May 2006.

[O'Carroll, 1993] David C. O'Carroll. Feature-detecting neurons in dragonflies. *Nature*, 362:541-543, April 1993.

[O'Carroll *et al.*, 2011] David C. O'Carroll, Paul D. Barnett, and Karin Nordström. Local and global responses of insect motion detectors to the spatial structure of natural scenes. *Journal of Vision*, 11(14):1-17, December 2011.

[Olberg *et al.*, 2007] Robert M. Olberg, R.C. Seaman, M.I. Coats, and A.F. Henry. Eye movements and target fixation during dragonfly prey-interception flights. *Journal of Comparative Physiology A*, 193:685-693, July 2007.

[Olberg *et al.*, 2000] Robert M. Olberg, A.H. Worthington and, K.R. Venator. Prey pursuit and interception in dragonflies. *Journal of Comparative Physiology A*, 186:155-162, February 2000.

[Ren *et al.*, 2017] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions On Pattern Analysis And Machine Intelligence*, 39(6):1137-1149, June 2017.

[van Hateren and Schilstra, 1999] Johannes H. van Hateren and C. Schilstra. Blowfly Flight and Optic Flow II. Head Movements During Flight. *The Journal of Experimental Biology*, 202:1491-1500, June 1999.

[Wiederman *et al.*, 2008] Steven D. Wiederman, Patrick A. Shoemaker, and David C. O'Carroll. A model for the detection of moving targets in visual clutter inspired by insect physiology. *PLoS One*, 3(7):1-11, July 2008.

[Wiederman and O'Carroll, 2011] Steven D. Wiederman and David C. O'Carroll. Discrimination of Features in Natural Scenes by a Dragonfly Neuron. *The Journal of Neuroscience*, 31(19):7141-7144, May 2011.

[Zabala *et al.*, 2012] Francisco Zabala, Peter Polidoro, Alice Robie, Kristin Branson, Pietro Perona, and Michael H. Dickinson. A Simple Strategy for Detecting Moving Objects during Locomotion Revealed by Animal-Robot Interactions. *Current Biology*, 22:1344-1350, July 2012.