



OPEN

## Combined person classification with airborne optical sectioning

Indrajit Kurmi, David C. Schedl & Oliver Bimber✉

Fully autonomous drones have been demonstrated to find lost or injured persons under strongly occluding forest canopy. Airborne optical sectioning (AOS), a novel synthetic aperture imaging technique, together with deep-learning-based classification enables high detection rates under realistic search-and-rescue conditions. We demonstrate that false detections can be significantly suppressed and true detections boosted by combining classifications from multiple AOS—rather than single—integral images. This improves classification rates especially in the presence of occlusion. To make this possible, we modified the AOS imaging process to support large overlaps between subsequent integrals, enabling real-time and on-board scanning and processing of groundspeeds up to 10 m/s.

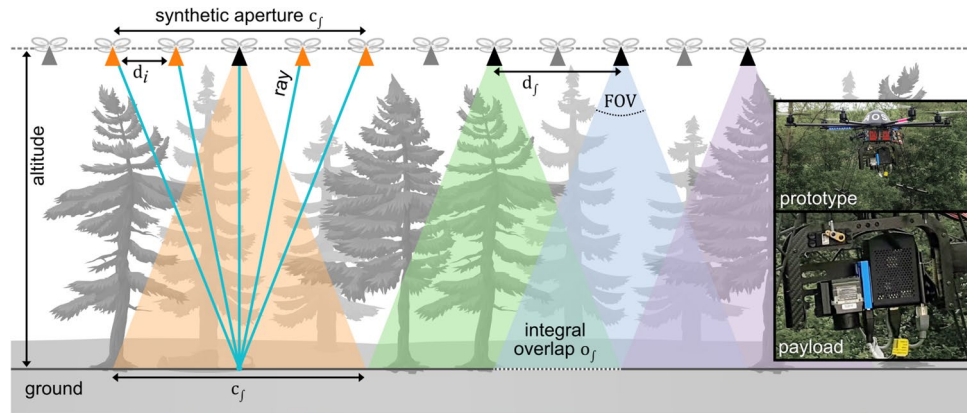
Synthetic aperture (SA) sensing is a widely used technique for emulating wide aperture sensors where they would be physically infeasible. It acquires individual signals of multiple or a single moving small-aperture sensor to computationally combine them to improve, for instance, resolution, depth of field, frame rate, contrast, or signal-to-noise ratio. SA sensing has been utilized with diverse sensors in a wide range of applications, such as radar<sup>1–3</sup> (obtaining weather-independent images and reconstructing geospatial depth), radio telescopes<sup>4,5</sup> (observing large celestial phenomena in outer space), microscopes<sup>6</sup> (reconstructing a defocus-free 3D volume using interferometry), sonar<sup>7–10</sup> (generating high-resolution mappings of underwater objects and seafloors), ultrasound<sup>11,12</sup> (non-intrusive intravascular 3D imaging), laser<sup>13,14</sup> (earth observation utilizing shorter wavelengths using LiDAR), and optical imaging<sup>15–22</sup> (acquiring structured light fields with large camera arrays for various post-processing steps, such as refocusing, computation of virtual views with maximal synthetic apertures, and varying depth of field).

Airborne optical sectioning (AOS)<sup>23–31</sup> is an effective wide-synthetic-aperture aerial imaging technique that can be deployed using camera drones. It allows virtual mimicking of a wide aperture optic of the shape and size of the scan area (possibly hundreds to thousands of square meters) that generates images of an extremely shallow depth of field above an occluding structure, such as a forest. These images are computed by integrating regular single pictures that are captured by the drone and allow optical slicing through dense occlusion (caused by leaves, branches, and bushes). In each slice, AOS can reveal targets, such as artifacts, objects, wildlife, or persons, which would remain occluded for regular cameras. Compared to alternative airborne scanning technologies, such as LiDAR<sup>32–34</sup> and synthetic aperture radar<sup>1–3</sup>, AOS is cheaper, wavelength-independent, and offers real-time computational performance for occlusion removal. We have applied AOS within the visible<sup>23</sup> and the thermal<sup>26</sup> spectra, and demonstrated its usefulness in archeology<sup>23</sup>, wildlife observation<sup>27</sup>, and search and rescue (SAR)<sup>30,31</sup>. By employing the randomly distributed statistical model in Ref.<sup>25</sup>, we explained AOS' efficiency with respect to occlusion density, occluder sizes, number of integrated samples, and size of the synthetic aperture.

Recently, we have proven that classification of partially occluded persons in forests using aerial thermal images is significantly more effective when AOS is used to integrate single images before classification than when classification results of single images are combined<sup>30</sup>. We have also demonstrated the real-time application of AOS in fully autonomous and classification-driven adaptive SAR operations<sup>31</sup>. Currently, we achieve average precision scores of 86.0–92.2% under realistic conditions<sup>30,31</sup>.

In this article, we show how the chance of detecting persons can be increased by combining multiple classification results from strongly overlapping AOS integral images. Combined classification has been widely studied for improving overall classifier performance (improving accuracy, dealing with diverse and noisy datasets, etc.)<sup>35–41</sup>. Combination mechanisms can be used at various levels of the classification<sup>42–52</sup> (e.g., at the data, feature, or decision levels). Decision-level combination is particularly popular due to its simplicity (no extensive knowledge of classifiers required) leading to lower complexity of the combination method. Most decision-level combination techniques are broadly categorized based on the output of the classifier (e.g., abstract, rank, and measurement). The most informative are measurement-based (also known as score-based) combination methods<sup>36</sup>. Approaches for measurement-based combination can be further categorized into adaptive<sup>47–50</sup> and non-adaptive

Institute of Computer Graphics, Johannes Kepler University, 4040 Linz, Austria. ✉email: oliver.bimber@jku.at



**Figure 1.** Synthetic aperture sampling with airborne optical sectioning: sampling parameters and current drone prototype with payload (right).

methods<sup>39,51,52</sup>. Adaptive techniques are based mainly on evolution, artificial intelligence algorithms, or fuzzy set theories, while non-adaptive combination techniques apply simple rules for combination, such as sum, product, maximum, and median<sup>39,51,52</sup>. No theoretical or empirical evidence of the general superiority of any particular combination scheme exists, and even simple combination schemes have been shown to improve accuracy in various systems<sup>40,41</sup>.

To guarantee real-time rates on low-performance on-device mobile processors, we consider non-adaptive, measurement-based combination techniques. We demonstrate in “[Combined classification](#)” and “[Results](#)” sections that (compared to independent classifications in single integral images) the product of median and maximum confidence scores of combined classifications significantly suppresses false while boosting true detections. However, to enable a combined classification initially, the AOS imaging process must support large overlaps between subsequent integral images, which was not hitherto feasible. We discuss optimal AOS sampling parameters and how we achieve these overlaps in the “[Continuous 1D synthetic aperture sampling](#)” and “[Combined classification](#)” sections. We start with a brief review of the AOS principle in the “[Airborne optical sectioning](#)” section.

### Airborne optical sectioning


As shown in Fig. 1, AOS captures thermal radiation with a drone that samples forest within the range of a synthetic aperture (SA) at flying altitude. This results in multiple geotagged aerial thermal images and consequently in unstructured thermal light-field rays formed by image pixels and camera poses on the SA<sup>53,54</sup>. With known camera intrinsics, drone poses, and a representation of the terrain (e.g., a digital elevation model or a focal plane approximation<sup>28</sup>), each ray’s origin on the ground can be reconstructed. Averaging all rays with the same origin results in a focused and widely occlusion-free integral image of thermal sources (e.g., persons) on the ground.

There is a statistical chance that a point on the forest ground is unoccluded by vegetation from multiple perspectives, as explained by the probability model in Ref.<sup>25</sup>. Thus, depending on density, more or fewer rays of a surface point contain information of random occluders, while others carry the constant signal of the target. Integrating multiple rays deemphasizes the occlusion signal and emphasizes the target signal. Since the remaining occlusion only lowers the contrast of the target<sup>25</sup>, reliable classification of the target is possible—even under strong occlusion conditions<sup>30</sup>.

AOS relies on the camera’s pose information while capturing images within a certain SA. In Ref.<sup>30</sup>, the SA was a 2D sampling area, and precise pose estimation was achieved with computer-vision-based reconstruction techniques. This resulted in high-quality integral images but was—for two reasons—not usable in time-critical applications, such as search and rescue: Sampling a 2D area led to long flying times, and precise computer-vision-based pose estimation was not possible in real time. All computations were carried out offline and after recording at predefined SA waypoints. This was significantly improved in Ref.<sup>31</sup> by sampling along short 1D SAs (linear flight paths) and by using instant but imprecise sensor readings from the drone (barometer altitude, non-differential GPS location, and compass orientation) for real-time computations directly on the drone. Despite the imprecision caused by lower-dimensional sampling and imprecise pose estimation, person classification performance was similar to that of 2D SA sampling with precise pose estimation<sup>31</sup>.

To date<sup>30,31</sup>, person classification with AOS has only been achieved in discrete (non-overlapping) integral images, for which 30 single images were required at 1 m intervals before they could be integrated and classified after a 30 m flight distance and 30 s of flight time. Due to slow recording speed and long image transfer times from the camera to the drone’s processor, the ground surface covered did not overlap in multiple integral images. Consequently, the probability of detecting a person on the ground surface relied solely on a single classification chance. If under unfavorable occlusion conditions, a person was not detected in the corresponding integral image, they were never found.

In the Sections below, we present the theoretical and practical foundations of continuous 1D synthetic aperture sampling (i.e., the real-time capturing, processing, and evaluation of integral images with a large ground



**Figure 2.** Sampling distances  $d_i$  of single images with the same image disparity and  $FOV$  at two different flight altitudes  $h_1$  and  $h_2$ . Note that a higher altitude does not result in more coverage on the ground, with a maximum  $D_f$ .

surface overlap); it increases the chance of detecting a person proportionally to the amount of integral image overlap by combining multiple classification results of the same spot on the ground.

### Continuous 1D synthetic aperture sampling

For continuous 1D synthetic aperture sampling, various sampling parameters, such as flying speed and altitude, imaging and processing speed, and camera field of view can be considered (cf. Figs. 1, 2).

The distance between two subsequent integral images depends on the drone's flying speed  $v_f$  [m/s] and the required processing time  $t_p$  [s]:

$$d_f = v_f \cdot t_p \tag{1}$$

where  $d_f$  increases with flying speed and processing time. For instance, for a computational performance (including integral image computation and classification) of  $t_p = 0.5$  s (i.e., two processing passes per second),  $d_f$  is 0.5 m, 2 m, 3 m, and 5 m for  $v_f$  of 1 m/s, 4 m/s, 6 m/s, and 10 m/s, respectively.

The ground coverage of each integral image depends on the flight altitude  $h$  [m] and the camera's field of view  $FOV$  [°]:

$$c_f = 2 \cdot h \cdot \tan(FOV/2). \tag{2}$$

Thus, for  $h = 35$  m above ground level (AGL) and a  $FOV$  of 43°, for example,  $c_f$  is 27.6 m. Note that  $c_f$  represents one dimension/direction of the 2D area covered by each integral image on the ground.

The amount of integral image overlap (i.e., how often the same surface point is covered by subsequent integral images) is:

$$o_f = c_f/d_f. \tag{3}$$

At a  $FOV$  of 43° and an increasing  $d_f$  of 0.5 m, 2 m, 3 m, and 5 m, for instance,  $o_f$  decreases: 55.2, 13.8, 9.2, and 5.5. In principle,  $o_f$  represents the number of times the same person can be detected while being scanned. Note that  $o_f$  should not fall below 1, as this results in imaging gaps (ground portions not being covered at all).

The sampling density of integral images is the number of single images  $N$  being integrated:

$$D_f = N = c_f/d_i, \tag{4}$$

where  $d_i$  [m] is the sampling distance of single images. It correlates with the efficiency of occlusion removal, and has an upper limit, as discussed in Ref.<sup>25</sup>. Note that  $t_p$  also increases proportionally with  $N$ . Furthermore, the rate at which single images are recorded should be equal to or higher than the rate at which integrals are computed ( $d_f \geq d_i$ ). If  $d_f < d_i$ , subsequent integral images will not change.

The integration time (i.e., the time required to capture  $N$  images that are combined into one integral image) is:


$$t_f = c_f/v_f = d_i \cdot D_f/v_f. \tag{5}$$

For  $c_f = 27.6$  m,  $t_f$  is 27.6 s, 6.9 s, 4.63 s, and 2.8 s for flying speeds  $v_f$  of 1 m/s, 4 m/s, 6 m/s, and 10 m/s. Note that a large  $t_f$  is unfavorable for the detection of fast-moving persons, as they introduce motion blur to the integral images that might not be classified correctly.

The sampling distances of single images that cause the same image disparity required for effective occlusion removal (as explained in Ref.<sup>25</sup>) at different flight altitudes (but at the same  $FOV$ ) are linearly related (cf. Fig. 2):

$$d_{i1}/d_{i2} = h_1/h_2. \tag{6}$$

Thus, to match the occlusion removal efficiency of  $d_{i1} = 1$  m sampling distance at  $h_1 = 35$  m, AGL requires a sampling distance of  $d_{i2} = 28.6$  m at  $h_2 = 1000$  m AGL. Both synthetic aperture size and ground region covered scale proportionally to  $h_1/h_2$  to achieve the same  $D_f$  at these two altitudes. Covering the same ground region at



**Figure 3.** Increase in projected occlusion density ( $\tilde{D}_\alpha$ ,  $\tilde{D}$  for  $\alpha = 0$ ) with changing viewing angle  $\alpha$ .

the same  $D_f$ , however, requires the same scanning time for the same  $v_f$ , as an identical distance must be flown. Scanning at lower altitudes ( $h_1 < h_2$ ) benefits from a  $h_2/h_1$  times higher spatial sampling resolution compared to scanning from higher altitudes ( $h_2 > h_1$ ), and from  $h_2/h_1$  times more intermediate classification results for the same scanning distance, as  $h_2/h_1$  times more single images are captured.

As explained in Ref.<sup>25</sup>, the orthographically projected occlusion density  $\tilde{D}$  in single images depends statistically on the density and size of the occluders (i.e., density of forest and size of branches, trunks, leaves, etc.) and the height of the occlusion volume (i.e., the height of trees). Here, the assumption of an orthographic projection implies an orthographic viewing angle  $\alpha = 0$  with respect to the synthetic aperture (SA) plane (i.e., looking straight down at the forest ground). More oblique viewing angles ( $\alpha > 0$ ), however, result in an increase in projected occlusion density due to the longer imaging distance from the SA plane through the occlusion volume to the forest ground (see proof in Sect. S1 of the Supplementary Material):

$$\tilde{D}_\alpha = 1 - (1 - \tilde{D})^{\frac{1}{\cos(\alpha)}}. \quad (7)$$

As illustrated in Fig. 3, large viewing angles (and consequently cameras with a large  $FOV = 2 \cdot \alpha$ ) are inefficient for occlusion removal. In the case of forests, oblique viewing angles would additionally cause larger projected occluder sizes, since side views of tree trunks project to larger footprints than top-down views. Larger occluder sizes, however, would require even larger sampling distances for efficient occlusion removal<sup>25</sup>.

Considering all the above sampling parameters, the following conclusions can be drawn:

1. A larger  $FOV$  is beneficial only up to a certain degree (when occlusion removal becomes seriously inefficient because viewing angles are too oblique).
2. For a given flying speed, higher flight altitudes have no effect on scanning time (i.e., the time needed to cover a certain forest range), but lead to lower spatial sampling resolution on the ground. Flight altitudes should therefore be chosen to be as low as possible.
3. A higher flying speed decreases not only scanning time but also the integral image overlap (and thus the number of times a person can be detected correctly). A suitable trade-off between flying speed and detection probability must be chosen. This might depend on the amount of occlusion (i.e., slower flights for denser forests).


Faster imaging speeds and shorter processing times are always beneficial, as both increase integral image overlap and allow faster flying speeds.

### Combined classification

In Ref.<sup>31</sup>, a fully autonomous and classification driven drone was developed and deployed to carry out wilderness search and rescue operations without human intervention. The system utilizes YOLO<sup>55</sup> (a common and widespread object detection model<sup>56,57</sup> which provides real-time performance and consumes low enough power<sup>58</sup> to be deployed on mobile processors) to achieve state-of-the-art results for person detection in the wilderness. However, the system performs classification on discrete (non-overlapping) integral images (single integral images over a 30 m flight distance and after 30 s of flight) and thus tends the probability of detecting persons on a single classification chance. Here we demonstrate how the continuous computation of integral images and combined classifications within them increase the chance that a person is detected correctly by a factor of  $o_f$  compared to single classifications of non-overlapping integral images, as in Refs.<sup>30,31</sup>.

Many real-time object-detection algorithms, such as YOLO<sup>55</sup>, output classification results as axis-aligned bounding boxes (AABBs) together with detection confidence scores. For classification combination, we project all AABBs of all individual integral images to a common coordinate system that is defined by the digital elevation model (DEM) of the ground surface. Thus, for a single discrete point on the ground surface, we collect





**Figure 4.** Test sites, flight directions, and coverages (left) at constant flying speeds of  $v_f = 4, 6,$  and  $10$  m/s over unoccluded (open field) terrain (top rows), over conifer forest (middle rows), and over broadleaf forest (bottom rows). Resulting integral images (right) reveal the appearance of the same person (bounding boxes indicate the manually labeled ground-truth appearances) of  $o_f = 13.8$  for  $4$  m/s,  $9.2$  for  $6$  m/s, and  $5.5$  for  $10$  m/s. Note that the shape of the ground-truth bounding boxes varies slightly due to misregistration in the integral images. The GPS coordinates of the test sites are:  $48^\circ 20' 08.4''$  N,  $14^\circ 19' 34.6''$  E (open field),  $48^\circ 19' 58.1''$  N,  $14^\circ 19' 48.1''$  E (conifer forest),  $48^\circ 19' 59.8''$  N,  $14^\circ 19' 52.2''$  E (broadleaf forest). An enlarged view is provided in Sect. S2 of the Supplementary Material for better readability.

a maximum of  $o_f$  overlapping AABBs and combine their individual confidence scores to a single score. This results in a DEM-aligned confidence map (initially filled with zero values). To distinguish between false and true detections on a per-pixel basis, each entry in this confidence map can be thresholded after  $o_f$  new image samples past its first appearance within the drone's FOV. Note that, if a DEM is not available, the ground surface can be automatically approximated by a simple shape<sup>28</sup> (e.g., a plane).


Our hypothesis is that true detection often (but not always) project higher scores to the same spatial location, while false detections project predominantly (but not exclusively) lower scores to more randomly distributed locations. Thus, combining the projected scores should emphasize true and suppress false detection scores, and consequently separate the score ranges of these two groups more clearly than in single integral image classifications. This will in turn allow better thresholding and thus improve overall classification performance.

Confidence scores of multiple integral images can be combined by statistical or mathematical methods. We evaluated three different methods for combining the confidence scores of integral images (median, maximum, and the product of median and maximum), and present them along with confidence scores obtained in single integral images by using our previous approach<sup>31</sup> in the “Results” section.

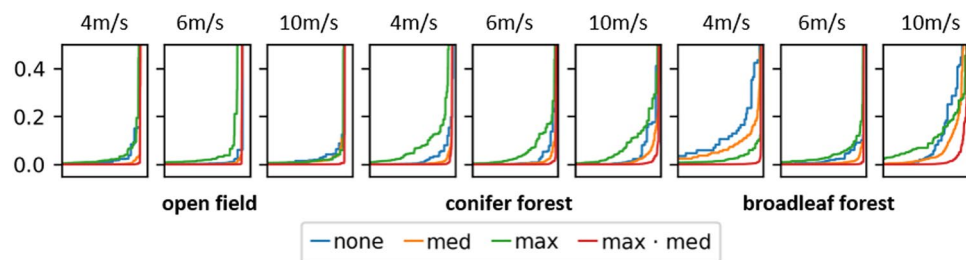
## Results

To evaluate our combination strategies, we performed 1D SA test flights at various constant speeds ( $v_f = 4, 6,$  and  $10$  m/s) from  $30$  m AGL above unoccluded (open field) terrain, and  $37$  m above dense forests (conifer and broadleaf forest). Integral images were computed from  $D_f = N = 30$  single images recorded at a video rate of  $t_i = 0.33$  s (30 fps) and a FOV of  $43.10^\circ$ . The processing speed achieved was around  $t_p = 0.5$  s. Details on the hard- and software used, the new AOS image acquisition that achieves fast flying speeds and large integral image overlaps, the improved integral image computation that applies deferred rendering to reduce processing time, and the implementation of the person classification are provided in the “Methods” section.

Figure 4 illustrates the test sites and computed single integral images along the 1D SA flight paths together with the ground-truth labels of persons on the ground. As predicted by (3), the same person appears in  $o_f$  integral images for different  $v_f$  ( $14$  vs.  $13.8$  for  $4$  m/s,  $9-10$  vs.  $9.2$  for  $6$  m/s, and  $5-6$  vs.  $5.2$  for  $10$  m/s). Note that slight variations are caused by different transition angles through the FOV and the fact that the FOV differs slightly for horizontal/vertical and diagonal image axes. Figure 5 shows and compares, for each flight, the probability maps of single integral images (utilizing the same YOLO model as in Ref.<sup>31</sup>, but additionally trained with unoccluded targets, as described in the “Methods” section) and combined classification results. For the open-field flights and single integral images, high confidence scores were obtained at the ground-truth positions, while low-score detections could easily be filtered out by a distinguished confidence threshold. This, however, was not the case for occlusion (especially during faster flights), where classification performance dropped significantly, and finding a proper confidence threshold to distinguish between false and true detections became increasingly difficult.



**Figure 5.** Probability maps of single integral images (first 16 columns) and combined classification results (last three columns: maximum, median, and maximum·median) for all test flights. Detections are indicated with AABBs, and confidence scores are color coded (see the logarithmic color bar on the right). An enlarged view is provided in Sect. S2 of the Supplementary Material for better readability.



**Figure 6.** Plotted confidence scores (y-axis) of all detections (x-axis) in low-to-high order (from left to right). The steepest gradient, which supports the best defined confidence-score thresholding, was always achieved with maximum·median combination. While the difference to other combination methods (and to no combination (none), that is, single integral image classification) was low without occlusion (open field), it was significant in the presence of occlusion (conifer and broadleaf forest).

The results of maximum, median, and maximum·median combined classifications are shown in the last three columns of Fig. 5. While the maximum emphasized true but also aggregated false detections, and the median suppressed false but also reduced true detections, the maximum·median best suppressed false and emphasized true detections. That the score ranges between true and false detections can be separated best by a threshold with a maximum·median combination is also illustrated by the plot of confidence-score-ordered detections shown in Fig. 6. A steep gradient supports distinct confidence-score thresholds. Here, maximum·median outperformed a maximum and median combination as well as single integral image classification (without combination). Table 1 compares the maximum confidence scores of false detections with the minimum confidence scores of true detections. Their ratio represents the degree of separation between both groups (and how well confidence thresholds can be chosen). A value below 1 indicates overlapping scores for false and true detections. Thresholding will consequently lead to false classifications (or missing true classifications), which was always the case for our forest flights and classifications in single integral images (results of our previous approach<sup>31</sup>). All combination methods increased the separation between false detections and true detection scores (ratio > 1), while the maximum·median method performed best. This implies that (in contrast to single integral image classification) a distinct and robust threshold can be found which leads to no false but all true detections.

## Methods

We utilized an octocopter (MikroKopter OktoXL 6S12, two LiPo 4500 mAh batteries, 4.5 kg) that carried the following payload (cf. Fig. 1): a FLIR Vue Pro thermal camera (9 mm fixed focal length lens, 7.5  $\mu$ m to 13.5  $\mu$ m spectral band, 14-bit non-radiometric, 118 g, 1.2 W), a Flir HDMI and power module providing HDMI video output from the camera (640  $\times$  480 @30 Hz; 15 g), a video capture card (Basetech, 640  $\times$  480 @30 Hz, 22 g), a

		Single	Max	Median	Max-median
Open field	4 m/s	0.274/0.119	0.617/0.119	0.380/0.015	0.235/0.001
		2.302	5.184	25.333	235.0
	6 m/s	0.249/0.099	0.471/0.099	0.365/0.015	0.172/0.0002
		2.515	4.757	24.333	860.0
	10 m/s	0.241/0.033	0.565/0.033	0.313/0.033	0.177/0.001
		7.303	17.121	9.484	177.0
Conifer forest	4 m/s	<b>0.122/0.147</b>	0.305/0.147	0.221/0.092	<i>0.067/0.008</i>
		<b>0.829</b>	2.074	2.402	8.375
	6 m/s	<b>0.084/0.157</b>	0.235/0.157	0.144/0.027	<i>0.033/0.002</i>
		<b>0.535</b>	1.496	5.333	16.5
	10 m/s	<b>0.107/0.140</b>	0.276/0.140	0.204/0.041	<i>0.056/0.006</i>
		<b>0.764</b>	1.971	4.975	9.333
Broadleaf forest	4 m/s	<b>0.002/0.049</b>	0.406/0.049	0.007/0.004	<i>0.003/0.0001</i>
		<b>0.041</b>	8.285	1.75	30.0
	6 m/s	<b>0.037/0.098</b>	0.397/0.098	0.188/0.017	<i>0.075/0.0008</i>
		<b>0.377</b>	4.051	11.058	93.75
	10 m/s	<b>0.008/0.224</b>	0.302/0.224	0.054/0.054	<i>0.016/0.004</i>
		<b>0.035</b>	1.348	1.0	4.0

**Table 1.** Minimum confidence scores of true detections/maximum confidence scores of false detections (average overall detections and all integral images) and the ratio of those scores. Values below 1 (bold) indicate false classifications. The maximum-median combination performed best. In the case of occlusion (conifer and broadleaf forest), it separated false from true confidence scores by a factor of 4–93 (italic).

single-board system-on-chip computer, SoCC (RaspberryPi 4B, 8 GB RAM, 65 g, 6.4 W), an LTE communication hat (Sixfab 3G/4G & LTE base hat and a SIM card, 35 g), and a Vision Processing Unit, VPU (Intel Neural Compute Stick 2, 30 g, 1.5 W). The payload (350 g in total) was mounted on a rotatable gimbal and was positioned to keep the camera pointing downwards during flight.

Our software was implemented in Python and runs on the SoCC, where various sub-processes, such as drone communication, imaging acquisition, and image processing (including integral image computation and classification), run in parallel to make use of the SoCC's multi-core capabilities. Inter-process queues are applied for efficient and secure communication between sub-processes. The drone communication sub-process interacts with the drone, utilizing a MikroKopter-customized serial protocol to receive IMU/GPS positions and send waypoint instructions which include GPS location, orientation, and speed. Received GPS positions and orientations are time-stamped and communicated to the image-processing sub-process. Note that all experiments were carried out in accordance with relevant guidelines and regulations. The study protocols were approved by ethics committee of Upper Austrian government.

**Image acquisition.** For imaging, we grabbed digital video frames from the video capture card connected to the thermal camera using OpenCV's video capturing module. The camera was set to SAR mode (mode suited to search-and-rescue operations in the wilderness, using 100% field of view), providing a high tonal range for higher temperatures and fewer gray levels for colder temperatures. The images were time-stamped to assign individual GPS coordinates, preprocessed using OpenCV's pinhole camera model to remove the lens distortion, and cropped to a field of view of 43.10° and a resolution of 512 px × 512 px.

To match the slow and asynchronous GPS signal rate (5 Hz in our case) to the faster video rate (30 Hz in our case), we applied time-based linear interpolation (assuming constant flying speeds). We thus interpolated GPS coordinates for each video frame grabbed. This, however, led to an interpolation error that depends on flying speed and imaging time:


Given the drone's constant flying speed  $v_f$  [m/s] and the camera's imaging time  $t_i$  [s], the maximum interpolation error caused by an unknown delay between capturing and measuring the capturing timestamp (which includes the transmission time of the image from the camera to the processor) is:

$$E_{I_{max}} = (v_f \cdot t_i) / 2, \quad (8)$$

and adds to the GPS error. Thus, for  $t_i = 0.033$  s (30 fps), and  $v_f = 1$  m/s, 4 m/s, 6 m/s, and 10 m/s,  $E_{I_{max}} = 1.67$  cm, 6.67 cm, 10 cm, and 16.67 cm, for example. Note that assuming a constant drone speed results in  $E_{I_{max}}$  being independent of the speed of the GPS sensor, as GPS positions of recorded camera frames can be linearly interpolated if they cannot be measured sufficiently fast. Only for non-constant speed segments (e.g., during acceleration and deceleration) are fast GPS samples beneficial for more precise piecewise linear or non-linear interpolations.

**Deferred integral imaging computation.** In principle, the preprocessed and GPS-assigned single thermal images are projected onto and averaged at a digital elevation model (DEM) using their individual poses and the camera's fixed intrinsic parameters. The DEM is a triangular mesh compatible with most standard graphics





**Figure 7.** Integration times [ms] for various  $N$  (number of images integrated) and a range of DEM sizes (vertex counts in millions). The new deferred rendering (solid lines) is compared with the old classical rendering (dashed lines). Classical and deferred renderings have similar computation times for small numbers of vertices. However, unlike for deferred rendering, rendering times drastically increase with vertex count for classical rendering. The plotted curves are averaged over 100 integral computations.

pipelines, as explained in Ref.<sup>31</sup>. For an integral image, the DEM with all projections is finally rendered from the center perspective of all integrated poses.

In Ref.<sup>31</sup>, the above process was implemented with classical rendering, for which the entire DEM had to be processed for each projection. This did not achieve practical computation times for large numbers of projections and/or large DEMs (cf. Figure 7). To decouple image projection from the processing of the DEM's geometry, we now utilize a rendering technique known as deferred rendering (DR). This requires the DEM's geometry to be processed only once for each integral image and consequently speeds up computation time. With DR, the processed DEM geometry is preserved in the graphics memory and can be reused for all projections as long as the rendering perspective does not change. As shown in Fig. 7, compared to classical rendering, this leads to a significant decrease in computation time for large DEMs and large numbers of projections. Storing the DEM's processed geometry requires hardware support for floating-point precision buffers, which is supported by our SoCC.

Integral image computation was implemented using C, C++, and OpenGL, and integrated into the main Python program using Cython. Preprocessing and integration of 30 thermal images require around 99 ms and 115 ms, respectively, for DEMs with 34k–2.6m vertices.

**Classification.** Person detection in the integral images was performed with the YOLOv4-tiny network architecture<sup>59</sup> on the VPU. More details on (pre-)training, parameters, and the selection of the best training weights can be found in Ref.<sup>31</sup>. Classifications are achieved in 84 ms and are optionally transmitted to a remote mobile device using the LTE communication hat connected to the SoCC.

For classifier training (using the Darknet software), we used 17 previously recorded flights<sup>30</sup> (F0–F11 and O1–O6; excluding F7). The data contains 7 flights over forests with persons on the ground (F0–F6), 4 flights over forests without persons (F8–F11), and 6 flights over a meadow with persons but without occlusions (O1–O6). From these flights F1, F8, and O6 were used for validation, while the remaining 14 flights were used for training. For our experiments, the recordings were resampled from 2D synthetic apertures to 1D synthetic aperture lines, pose data was computed from GPS readings, and persons were labeled, as explained in Ref.<sup>31</sup>.

Note that manual compass correction was applied to each flight. To compute the integral images of the training and validation datasets, we applied the following augmentations: We varied the synthetic aperture size  $N$  in 7 steps: 1 (pinhole), 5, 10, 15, 20, 25, and 30. Note that, due to resampling, some 1D apertures had fewer single images (e.g., the longest lines were only in the range  $20 < N \leq 25$  for some scenes). Additionally, we applied 10 random image rotations by varying the up vector of the integral's virtual camera. The digital elevation model was translated up and down by 3 m in steps of 1 m to simulate defocus. Furthermore, we computed an additional integral image with a random compass error (rotation around each single-image camera's forward axis by  $\pm 15$  degrees). This led to a total of 980 variations for each of the 179 resampled 1D apertures.

The trained network achieved an average precision score of 88.7% (without combined classification) on test data presented in Table 2 of Ref.<sup>31</sup> (previously 86.2% in Ref.<sup>31</sup>) and was also able to detect unoccluded persons in open terrain (e.g., meadows, fields), since it was additionally trained with the unoccluded person data from Refs.<sup>30,60</sup>. Note, that for combined classification, a precision score cannot be determined as pixel-precise ground truth labels are not available.

**Ethics declarations.** The ethics committee of the Upper Austrian government approved the study, and participants provided written informed consent.



## Discussion

We have shown that false detections can be significantly suppressed and true detections significantly boosted by classifying based on multiple combined AOS—rather than single—integral images. This improves classification rates especially in the presence of occlusion. To make this possible, we modified the AOS imaging process to support large overlaps between subsequent integrals, thus enabling real-time and on-board scanning and processing (processing speed of 0.5 s at 9.1 W) for groundspeeds of up to 10 m/s (while previously the maximum was 1 m/s<sup>30,31</sup>). One of the major limitations of utilizing a deep learning-based classifier is its non-deterministic nature which (in contrast to the derivable dependency between target visibility in integral images and occlusion density<sup>25,28</sup>) prevents us from establishing a distinct mathematical relationship between person detection rate and occlusion. Empirical evidence<sup>31</sup> only suggest a strong correlation between classifier performance and visibility while the relation between visibility and occlusion density has already been described for randomly distributed occluder models<sup>25</sup>. Furthermore, any simulation-based correlation investigation between occlusion density and classifier performance will be heavily dependent on the classifier configuration and the training data.

Due to the slow sampling rate of standard GPS, only constant-speed segments are currently supported. For acceleration and deceleration segments, faster and more precise Real-Time Kinematic (RTK) devices are beneficial. Non-linear interpolation and prediction models will be investigated in the future to enable better mapping of faster imaging rates to slower pose sampling. In future, we will also investigate additional improvements, such as adaptive combination techniques, online adaptive thresholding, and options for extending flight endurance to enable fully autonomous beyond-visual-line-of-sight (BVLOS) search-and-rescue missions.

## Data availability

The datasets generated during and/or analyzed during the current study are available in the Combined Person Classification with AOS data repository on Zenodo: <https://doi.org/10.5281/zenodo.5013639>.

Received: 8 November 2021; Accepted: 21 February 2022

Published online: 09 March 2022

## References

1. Moreira, A. *et al.* A tutorial on synthetic aperture radar. *IEEE Geosci. Remote Sens. Magn.* **1**, 6–43 (2013).
2. Li, C. J. & Ling, H. Synthetic aperture radar imaging using a small consumer drone. In *Proc. IEEE Int. Symp. Antennas Propag. USNC-URSI Nat. Radio Sci. Meeting*, 685–686 (2015).
3. Rosen, P. A. *et al.* Synthetic aperture radar interferometry. *Proc. IEEE* **88**, 333–382 (2000).
4. Levanda, R. & Leshem, A. Synthetic aperture radio telescopes. *Signal Process. Mag.* **27**, 14–29 (2010).
5. Dravins, D., Lagadec, T. & Nuez, P. D. Optical aperture synthesis with electronically connected telescopes. *Nat. Commun.* **6**, 6852 (2015).
6. Ralston, T. S. *et al.* Interferometric synthetic aperture microscopy (ISAM). *Nat. Phys.* **3**, 129–134 (2007).
7. Hansen, R. E. *Introduction to Synthetic Aperture Sonar* 3–28 (IntechOpen, 2011).
8. Hayes, M. P. & Gough, P. T. Synthetic aperture sonar: A review of current status. *IEEE J. Ocean. Eng.* **34**, 207–224 (2009).
9. Hansen, R. E. *et al.* Challenges in seafloor imaging and mapping with synthetic aperture sonar. *IEEE Trans. Geosci. Remote Sens.* **49**, 3677–3687 (2011).
10. Bülow, H. & Birk, A. Synthetic aperture sonar (SAS) without navigation: Scan registration as basis for near field synthetic imaging in 2D. *Sensors* **20**, 4440 (2020).
11. Jensen, J. A. *et al.* Synthetic aperture ultrasound imaging. *Ultrasonics* **44**, 5–15 (2006).
12. Zhang, H. K. *et al.* Synthetic tracked aperture ultrasound imaging: Design, simulation, and experimental evaluation. *J. Med. Imaging* **3**, 027001 (2016).
13. Barber, Z. W. & Dahl, J. R. Synthetic aperture ladar imaging demonstrations and information at very low return levels. *Appl. Opt.* **53**, 5531–5537 (2014).
14. Turbide, S. *et al.* Synthetic aperture lidar as a future tool for earth observation. *Proc. SPIE* **10563**, 1115–1122 (2017).
15. Vaish, V. *et al.* Using plane + parallax for calibrating dense camera arrays. In *Proc. CVPR*, 2–9 (2004).
16. Vaish, V. *et al.* Reconstructing occluded surfaces using synthetic apertures: Stereo, focus and robust measures. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, 2331–2338 (2006).
17. Zhang, H., Jin, X. & Dai, Q. Synthetic aperture based on plenoptic camera for seeing through occlusions. In *Advances in Multimedia Information Processing—PCM* (eds Hong, R. *et al.*) 158–167 (Springer, 2018).
18. Yang, T. *et al.* Kinect based real-time synthetic aperture imaging through occlusion. *Multimedia Tools Appl.* **75**(12), 6925–6943 (2016).
19. Joshi, N. *et al.* Synthetic aperture tracking: Tracking through occlusions. In *2007 IEEE 11th International Conference on Computer Vision*, 1–8 (2007).
20. Pei, Z. *et al.* Occluded object 3d reconstruction using camera array synthetic aperture imaging. *Sensors* **19**, 607 (2019).
21. Yang, T. *et al.* All-in-focus synthetic aperture imaging. In *Computer Vision—ECCV 2014* (eds Fleet, D. *et al.*) 1–15 (Springer, 2014).
22. Pei, Z. *et al.* Synthetic aperture imaging using pixel labeling via energy minimization. *Pattern Recogn.* **46**(1), 174–187 (2013).
23. Kurmi, I., Schedl, D. C. & Bimber, O. Airborne optical sectioning. *J. Imaging* **4**, 102 (2018).
24. Bimber, O., Kurmi, I. & Schedl, D. C. Synthetic aperture imaging with drones. *IEEE Comput. Graph. Appl.* **39**, 8–15 (2019).
25. Kurmi, I., Schedl, D. C. & Bimber, O. A statistical view on synthetic aperture imaging for occlusion removal. *IEEE Sens. J.* **19**, 9374–9383 (2019).
26. Kurmi, I., Schedl, D. C. & Bimber, O. Thermal airborne optical sectioning. *Remote Sensing* **11**, 1668 (2019).
27. Schedl, D. C., Kurmi, I. & Bimber, O. Airborne optical sectioning for nesting observation. *Sci. Rep.* **10**, 7254 (2020).
28. Kurmi, I., Schedl, D. C. & Bimber, O. Fast automatic visibility optimization for thermal synthetic aperture visualization. *IEEE Geosci. Remote Sens. Lett.* **18**, 836–840 (2021).
29. Kurmi, I., Schedl, D. C. & Bimber, O. Pose error reduction for focus enhancement in thermal synthetic aperture visualization. *IEEE Geosci. Remote Sens. Lett.* **19**, 1–5 (2021).
30. Schedl, D. C., Kurmi, I. & Bimber, O. Search and rescue with airborne optical sectioning. *Nat. Mach. Intell.* **2**, 783–790 (2020).
31. Schedl, D. C., Kurmi, I. & Bimber, O. An autonomous drone for search and rescue in forests using airborne optical sectioning. *Sci. Robot.* **6**(55), 1188 (2021).
32. Syngé, E. H. A method of investigating the higher atmosphere. *Lond. Edinb. Dublin Philos. Mag. J. Sci.* **9**, 1014–1020 (1930).
33. Molebny, V. *et al.* Laser radar: Historical prospective—from the east to the west. *Opt. Eng.* **56**, 31220 (2016).

34. Du, B. *et al.* High-speed photon-counting laser ranging for broad range of distances. *Sci. Rep.* **8**, 4198 (2018).
35. Mohandes, M., Deriche, M. & Aliyu, S. O. Classifiers combination techniques: A comprehensive review. *IEEE Access* **6**, 19626–19639 (2018).
36. Chitroub, S. Classifier combination and score level fusion: Concepts and practical aspects. *Int. J. Image Data Fusion* **1**(2), 113–135 (2010).
37. Nweke, H. F. *et al.* Al-garadi, Data fusion and multiple classifier systems for human activity detection and health monitoring: Review and open research directions. *Inf. Fusion* **46**, 147–170 (2019).
38. Xu, L., Krzyzak, A. & Suen, C. Y. Methods of combining multiple classifiers and their applications to handwriting recognition. *Trans. Syst. Man* **22**(3), 418–435 (1992).
39. Kittler, J. *et al.* On combining classifiers. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 226–239 (1998).
40. Kuncheva, L. I. A theoretical study on six classifier fusion strategies. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(2), 281–286 (2002).
41. Tulyakov, S. *et al.* Review of classifier combination methods. In *Machine Learning in Document Analysis and Recognition* (eds Marinai, S. & Fujisawa, H.) 361–386 (Springer, 2008).
42. Castanedo, F. A review of data fusion techniques. *Sci. World J.* **2013**, 1–19 (2013).
43. Okan, K., Kose, N. & Rigoll, G. Motion fused frames: Data level fusion strategy for hand gesture recognition. In *Proc. CVPR IEEE Workshops*, 2103–2111 (2018).
44. Bloch, I. Information combination operators for data fusion: A comparative review with classification. *IEEE Syst. Trans. Man Cybern. A Syst. Hum.* **26**(1), 52–67 (1996).
45. Gosa, M. *et al.* A survey of decision fusion and feature fusion strategies for pattern classification. *IETE Tech. Rev.* **27**(4), 293–307 (2010).
46. Yang, J. *et al.* Feature fusion: Parallel strategy vs serial strategy. *Pattern Recogn.* **36**(6), 1369–1381 (2003).
47. Zadrozny, B. & Elkan, C. Transforming classifier scores into accurate multiclass probability estimates. In *Proc. Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 694–699 (2002).
48. Amin, J. *et al.* A new approach for brain tumor segmentation and classification based on score level fusion using transfer learning. *J. Med. Syst.* **43**, 326–326 (2019).
49. Chen, M. *et al.* Image retrieval based on multi-feature similarity score fusion using genetic algorithm. In *The 2nd International Conference on Computer and Automation Engineering (ICCAE)*, 45–49 (2010).
50. Vatsa, M., Singh, R. & Noore, A. Improving iris recognition performance using segmentation, quality enhancement, match score fusion, and indexing. *IEEE Trans. Syst. Man Cybern.* **38**, 1021–1035 (2008).
51. Tax, D. M. J., Duin, R. P. W. & Van Breukelen, M. Comparison between product and mean classifier combination rules. In *Proc. Workshop on Statistical Pattern Recognition* (1997).
52. Hashem, S. Optimal linear combinations of neural networks. *Neural Netw.* **10**(4), 599–614 (1997).
53. Wetzstein, G. *et al.* Computational plenoptic imaging. *Comput. Graph. Forum* **30**, 2397–2426 (2011).
54. Wu, G. Light field image processing: An overview. *IEEE J. Sel. Top. Signal Process* **11**(7), 926–954 (2017).
55. Redmon, J. *et al.* You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788 (2016).
56. Haris, M. & Glowacz, A. Road object detection: A comparative study of deep learning-based algorithms. *Electronics* **10**(16), 1932 (2021).
57. Ahmed, M. *et al.* Survey and performance analysis of deep learning based object detection in challenging environments. *Sensors* **21**(15), 5116 (2021).
58. Nvidia Co. *Jetson Benchmarks* (2021). <https://developer.nvidia.com/embedded/jetson-benchmarks>. Accessed 28 February 2022.
59. Bochkovskiy, A., Wang, C. Y. & Liao, H. Y. M. YOLOv4: Optimal speed and accuracy of object detection. <http://arxiv.org/abs/2004.10934> (2020).
60. Schedl, D. C., Kurmi, I. & Bimber, O. Data: Search and rescue with airborne optical sectioning. *Nat. Mach. Intell.* **2**(12), 783–790. <https://doi.org/10.5281/zenodo.3894773> (2020).

## Acknowledgements

This research was funded by the Austrian Science Fund (FWF) under Grant Number P 32185-NBL and by the State of Upper Austria and the Austrian Federal Ministry of Education, Science, and Research via the LIT—Linz Institute of Technology under Grant Number LIT-2019-8-SEE-114. D.C.S. was with Johannes Kepler University Linz, during the study, and is now at the University of Applied Sciences Upper Austria.

## Author contributions

O.B. originated the concept. O.B., D.C.S., and I.K. conceived and designed the experiments. D.C.S. and I.K. performed the experiments. O.B., D.C.S., and I.K. analyzed the data. D.C.S. and I.K. contributed materials/analysis tools. O.B., D.C.S., and I.K. wrote the paper.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-07733-z>.

**Correspondence** and requests for materials should be addressed to O.B.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022