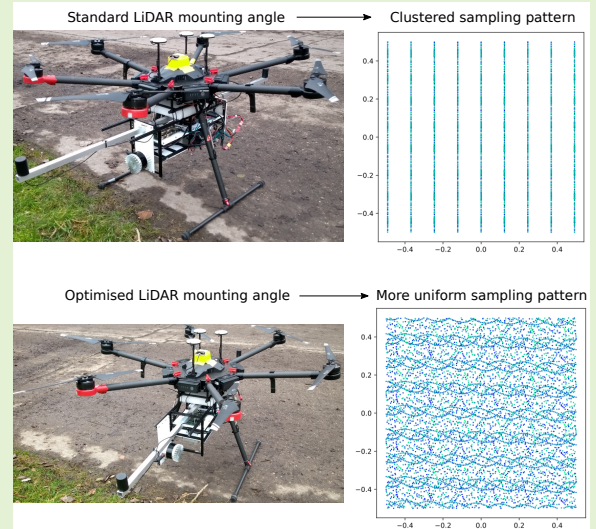


On the optimal mounting angle for a spinning LiDAR on a UAV

Laurens Diels, Michiel Vlamincx, Bart De Wit, Wilfried Philips *Senior Member, IEEE*, and Hiep Luong

Abstract—UAV-based 3D LiDAR mapping has received increasing interest from both the research community and industrial operations for the 3D reconstruction of large and dangerous environments. When a LiDAR sensor is rigidly mounted to a UAV, one usually orients it such that the LiDAR’s upward direction coincides with the UAV’s flight direction. For a spinning LiDAR this fully specifies the mounting orientation because of circular symmetry. However, we argue in this paper that this configuration is suboptimal, especially when the LiDAR’s full horizontal field of view is not important, either because of a limited field of view of other sensors in case of sensor fusion, or because the scene to (swath) map via LiDAR is less wide than the LiDAR’s field of view. Indeed, in this situation many of the points the LiDAR samples are not of interest. Additionally, in the usual configuration the UAV’s movement can compensate for the limited elevation resolution of the LiDAR, but much less so for the azimuth resolution. For a simplified theoretical model we introduce a quality metric for mounting angles. Using a simulation we then derive the optimal mounting angle, which we validate via real world flights. We suggest to rotate this standard mounting orientation over an angle of -84.5° over the UAV’s vertical direction.

Index Terms—LiDAR, Remote sensing, Swath mapping, UAV



I. INTRODUCTION

UNMANNED aerial vehicles (UAVs), colloquially known as drones, have become ever more important in the field of remote sensing. They allow cost-effective yet sufficiently detailed surveying of scenes on the scale of small forests, fields and buildings. In addition they are ideally placed for infrastructure inspection of hard to reach places [1]–[6].

LiDAR (Light Detection And Ranging) systems measure distances and can be used for 3D mapping purposes. Commonly used are spinning LiDARs, which contain a number of rotating lasers to determine these distances over multiple azimuth and elevation angles (because of the rotation and presence of multiple lasers, respectively). For example, the Ouster OS1-128 LiDAR has 128 lasers, and hence 128 corresponding elevation angles. Multiple settings are possible, but for the one we will focus on the most, it samples during each 360° sweep for each elevation at 1024 azimuths, and this at 20 sweeps per second. Thus every 50 ms we obtain $128 \cdot 1024 = 131\,072$ distances, which can be mapped into 3D

Euclidean space. By combining different sweeps we can get a very dense point cloud. Compared to other 3D reconstruction approaches such as photogrammetry, LiDAR point clouds are not only denser, but typically more accurate and much less time consuming to compute [7]–[9]. A downside is they do not contain RGB information. If this is desired, sensor fusion with an RGB camera will be needed.

It is natural to equip UAVs with LiDARs for airborne 3D mapping. This mounting can occur via a gimbal, where the LiDAR can rotate independently of the UAV, or via a rigid construction. In this paper we will only consider the latter case, although the conclusions are still useful for the gimbal situation. If we introduce right-handed reference frames for the UAV and the LiDAR, where the UAV’s $x^{(U)}$ -direction points forward and the $z^{(U)}$ -direction points upwards, and the LiDAR’s $z^{(L)}$ -direction points upwards (and the $x^{(L)}$ – and $y^{(L)}$ -directions are of less importance due to circular symmetry), the usual way to rigidly mount the LiDAR is to make the LiDAR’s $z^{(L)}$ -direction coincide with the UAV’s $x^{(U)}$ -direction, as illustrated in Figure 1. Indeed, this is the configuration used in e.g. [9]–[12]. To the best of our knowledge, this is the first paper examining other nadir configurations.

A downside to this standard configuration is that while we can compensate for the LiDAR’s limited elevation resolution and vertical field of view in a controlled manner by

LD, MV, WP and HL are with the imec research group of Image Processing and Interpretation (IPI) at the department of Telecommunication and Information Processing (TELIN) of Ghent University.

BDW is with the department of Geography of Ghent University.

All authors are with the UAV Research Centre (URC) of Ghent University.




Fig. 1. The UAV ($\cdot^{(U)}$), LiDAR ($\cdot^{(L)}$) and ground ($\cdot^{(G)}$) reference systems, in blue, yellow and cyan, respectively. The image shows our DJI M600 UAV with payload, containing in particular an Ouster OS1-128 spinning LiDAR, mounted in the standard orientation. Note that in our mathematical description below we will ignore the spatial extent of the UAV and LiDAR and let the origins of their reference systems coincide. Also the origin of the ground reference system will be placed differently.

adjusting the UAV's speed as the UAV's movement direction corresponds to the LiDAR's elevation direction, the same cannot be said for the LiDAR's azimuth resolution. The UAV's instability in flight will help to some extent, but this mostly lies beyond the pilot's control.

We will restrict ourselves to situations where in the default mounting configuration the horizontal field of view is not fully utilised, as we are limited by the field of view of other sensors, or as we are only interested in mapping a narrow strip. Examples of the latter situation include the mapping of (rail)roads, rivers, aqueducts and power lines. In both cases we can still make use of the large horizontal field of view by rotating the LiDAR $\pm 90^\circ$ about the $z^{(U)}$ -axis. However, this will also have a similar problem as before, where we can now only compensate for the azimuth resolution. Therefore, we intuitively expect a rotation with an angle 'close to', but still somewhat short of $\pm 90^\circ$ to work best. This intuition will be rigorously investigated in this paper.

Our main contributions are as follows.

- 1) We introduce a quality metric to measure the performance of each LiDAR mounting angle, at least in a simplified mathematical model.
- 2) By optimising the mounting angle we find that angles close to -84.5° up to sign are ideal.
- 3) We validate these results using real-world flights, which confirm the superiority of the suggested mounting angles over the standard one, in the case of single-pass flights (e.g. for swath mapping).

This paper is structured as follows. In section II we will properly define the mounting angle and derive mathematically how it impacts the LiDAR samples, at least in an idealised flight model. We will also see that we should avoid certain mounting angles such as the default one in Figure 1. In section III we will define a cost function, so that we can detect problematic angles. This of course also allows us to minimise the cost to find optimal angles. This is done in section IV. Having a (range of) optimal mounting angle(s) we then put




Fig. 2. The definition of the mounting angle α . The figure shows the LiDAR (grey rectangle) when looking along the UAV's negative $z^{(U)}$ -direction (i.e. downwards from the UAV's point of view). Note that from this view point the cylindrical LiDAR would indeed appear as a rectangle. The LiDAR's $y^{(L)}$ - and $z^{(L)}$ -directions are shown in grey. The UAV's $x^{(U)}$ - and $y^{(U)}$ -directions are drawn in black and placed to intersect at the LiDAR's centre. Not shown are the LiDAR's $x^{(L)}$ -direction, which points in to the page (away from the reader), and the aforementioned UAV's $z^{(U)}$ -direction, which points out of the page (towards the reader). Note that we use signed angles, and that in this illustration we have $\alpha < 0$.

the theory to the test in section V where we describe the results from actual flights performed with the recommended and default mounting angles to see if we indeed notice an improvement in practice. We finish in section VI with our conclusion and further research directions.

II. THE MOUNTING ANGLE

For land surveying one can usually fly directly over the area of interest, so it makes sense to mount the LiDAR at nadir, i.e. the LiDAR's $z^{(L)}$ -axis should be perpendicular to the UAV's $z^{(U)}$ -axis. But we can still rotate the LiDAR about the UAV's $z^{(U)}$ -axis. We define the mounting angle α as the oriented angle from the UAV's $x^{(U)}$ -direction to the LiDAR's $z^{(L)}$ -direction as shown in Figure 2. Thus, an angle of $\alpha = 0$ corresponds to the standard orientation.

Before demonstrating the effect of changing the mounting angle, we first completely fix the LiDAR reference frame by declaring that the $x^{(L)}$ -axis is the negative UAV $z^{(U)}$ -direction, i.e. points towards the ground during stable flight. The LiDAR reference frame's origin is naturally at the centre of the LiDAR, where we also place the origin of the UAV reference frame (as was already the case in Figure 2). To model the effect of the mounting angle we make the following simplifying assumptions.

- The LiDAR's $n \in \mathbb{N}$ lasers (e.g. $n = 128$) are uniformly spaced in terms of elevation angle (in the LiDAR reference system), within the range of $[-\beta_v, \beta_v]$, where $\beta_v \in (0, \pi/2)$ hence is half the vertical field of view (e.g. $\beta_v = \frac{\pi}{8} = 22.5^\circ$).
- The LiDAR samples $m \in \mathbb{N}$ points per laser per full 360° sweep (e.g. $m = 1024$), at a constant rate so that also the corresponding azimuth angles are uniformly spaced. The azimuth at time $t = 0$ is $\theta = 0$.
- The LiDAR system rotates at a frequency of $f \in \mathbb{R}_{>0}$ sweeps per second (e.g. $f = 20$ Hz), in the clockwise

direction when viewed along the negative $z^{(L)}$ -direction, i.e. from above in the LiDAR frame. This spinning direction is chosen to be consistent with the Ouster OS1-128.

- The UAV flies at a constant speed $s \in \mathbb{R}_{>0}$ in its $x^{(U)}$ -direction, at a constant height of $h \in \mathbb{R}_{>0}$, in a completely stable fashion over a perfectly flat plane.

We now consider the sampling pattern we obtain on this ground plane, for which we introduce a third (and final) reference frame, where the $x^{(G)}$ -direction corresponds to the UAV's negative $y^{(U)}$, and the $y^{(G)}$ -direction to the UAV's $x^{(U)}$. This is also shown in Figure 1. In Figure 2 the ground plane's $x^{(G)}$ - and $y^{(G)}$ -directions would thus be in the usual orientation: $x^{(G)}$ points to the right of the page and $y^{(G)}$ towards the top. The origin of the ground plane reference system is the location on the ground beneath the UAV at time $t = 0$. Note that this does not indicate the start of the flight, and that we allow negative times and indices, as this will prove more convenient mathematically. Under these assumptions and conventions one can show that at the k -th ($k \in \mathbb{Z}$) discrete laser firing time step the LiDAR samples the n points $\mathbf{Q}(k, l)$ with ground plane coordinates

$$\begin{aligned} \mathbf{Q}(k, l)_{x^{(G)}} &= \frac{h}{\cos\left(\frac{2k}{m}\pi\right)} \left(\cos\alpha \sin\left(\frac{2k}{m}\pi\right) \right. \\ &\quad \left. - \sin\alpha \tan\left(-\beta_v + \frac{2\beta_v}{n-1}l\right) \right), \\ \mathbf{Q}(k, l)_{y^{(G)}} &= \frac{h}{\cos\left(\frac{2k}{m}\pi\right)} \left(\sin\alpha \sin\left(\frac{2k}{m}\pi\right) \right. \\ &\quad \left. + \cos\alpha \tan\left(-\beta_v + \frac{2\beta_v}{n-1}l\right) \right) \\ &\quad + \frac{sk}{mf} \\ \mathbf{Q}(k, l)_{z^{(G)}} &= 0 \end{aligned} \quad (1)$$

for $l \in [0, n-1]_{\mathbb{Z}}$ (where this notation denotes the interval of integers). Indeed, these equations are derived as follows. Consider a LiDAR azimuth-elevation pair (θ, φ) at a specific time t . The LiDAR fires a laser in this direction, which hits a ground point a distance $r \in \mathbb{R}_{\geq 0}$ away. Transforming the spherical coordinates (r, θ, φ) into Euclidean coordinates $(x^{(L)}, y^{(L)}, z^{(L)})$ can be done using the standard formulas and by observing that $x^{(L)} = h$, which will allow us to express r in terms of h , θ and φ . This way we obtain

$$\begin{aligned} x^{(L)} &= r \cos\theta \cos\varphi = h \\ y^{(L)} &= r \sin\theta \cos\varphi = h \tan\theta \\ z^{(L)} &= r \sin\varphi = h \frac{\tan\varphi}{\cos\theta}. \end{aligned}$$

To switch over to the ground reference system we use that the directions $x^{(G)}$, $y^{(G)}$ and $z^{(G)}$ correspond to $-y^{(L)}$, $z^{(L)}$ and $-x^{(L)}$ when ignoring the mounting angle, i.e. if $\alpha = 0$ (see Figure 1), and that the mounting angle induces a $z^{(G)}$ -rotation. We also need to add the movement term and origin

offset which combined is simply $(0, st, h)$. This results in

$$\begin{aligned} x^{(G)} &= -y^{(L)} \cos\alpha - z^{(L)} \sin\alpha \\ &= -\frac{h}{\cos\theta} (\cos\alpha \sin\theta + \sin\alpha \tan\varphi) \\ y^{(G)} &= -y^{(L)} \sin\alpha + z^{(L)} \cos\alpha + st \\ &= -\frac{h}{\cos\theta} (\sin\alpha \sin\theta - \cos\alpha \tan\varphi) + st \\ z^{(G)} &= -x^{(L)} + h \\ &= 0. \end{aligned}$$

Finally, firing step $k \in \mathbb{Z}$ corresponds to an azimuth $\theta_k = -\frac{2\pi}{m}k$ as the LiDAR spins clockwise, and to a time step $t_k = \frac{k}{mf}$ (so that it takes a time of $1/f$ for m firings and azimuth angles, i.e. a for a full sweep). Since the l -th elevation angle ($l \in [0, n-1]_{\mathbb{Z}}$) is $\varphi_l = -\beta_v + \frac{2\beta_v}{n-1}l \in [-\beta_v, \beta_v]$ (which includes the distinct boundary angles), we then obtain (1).

However, this obviously is only correct when a laser pulse actually reaches the ground plane, i.e. has azimuth and elevation indices k and l such that the corresponding azimuth θ_k and elevation φ_l lie in $(-\pi/2, \pi/2)$. For the elevation this is guaranteed by the limited vertical field of view of the LiDAR system. But for the azimuths we need to artificially restrict the horizontal field of view.

We could use a horizontal field of view of $(-\pi/2, \pi/2)$, but azimuths close to $\pm\pi/2$ will travel a long time, which means that the returns from the laser pulses might not be registered by the LiDAR system. Additionally we have in mind flights where the horizontal field of view of the LiDAR is not important since the scene is small (e.g. in case of swath mapping), or the LiDAR data is to be fused with for example a camera with a smaller field of view. Therefore, we artificially introduce a smaller effective LiDAR horizontal field of view $[-\beta_h, \beta_h]$ (which corresponds to a swath width of $2h \tan\beta_h$). To be more precise, we specify some $\beta_h \in \mathbb{R}_{>0}$ and only consider time and azimuth indices

$$k \in K := \{k_0 + jm \mid j, k_0 \in \mathbb{Z}, -k_{0,\max} \leq k_0 \leq k_{0,\max}\}$$

where $k_{0,\max} := \left\lfloor \frac{m\widetilde{\beta}_h}{2\pi} \right\rfloor$ is the maximal azimuth index in the sweep around time $t = 0$ ($|\theta_k| \leq \widetilde{\beta}_h$). Now the effective half horizontal field of view β_h is the discretised version of $\widetilde{\beta}_h$, namely the azimuth corresponding to $k_{0,\max}$:

$$\beta_h = |\theta_{k_{0,\max}}| = \frac{2\pi}{m} \left\lfloor \frac{m}{2\pi} \widetilde{\beta}_h \right\rfloor.$$

In principle we allow $\widetilde{\beta}_h$ to be any value in $(0, \pi/2)$. But in practice it depends on the intended swath width, or the angle of view in case of sensor fusion with a camera.

We conclude that the points sampled on the ground plane have coordinates given by (1), for time and azimuth indices $k \in K$ and elevation indices $l \in [0, n-1]_{\mathbb{Z}}$. We denote the full point cloud, i.e. the set of all such points $\mathbf{Q}(k, l)$ by $L(\alpha)$. From now on we will work exclusively in 2D ground plane coordinates, and thus we will no longer include the $z^{(G)}$ -component, nor write the superscripts $\cdot^{(G)}$. Figure 3 shows an example of the 2D point clouds from a single LiDAR




Fig. 3. Sampled LiDAR point clouds from a single sweep (a) or a number of consecutive sweeps (b). More precisely we consider $k \in [-k_{0,\max}, k_{0,\max}] + jm$ for $j = 0$ in (a) and $-7 \leq j \leq 7$ in (b). The points are coloured according to the scan index j , with the first scan ($j = -7$) being dark blue and the last scan ($j = 7$) bright green. One can observe that the full point cloud consists of shifted copies of the point cloud from a single sweep, which also follows from (1). The parameters used here are $\alpha = \pi/6$, $\beta_h = \pi/4$ ($= \beta_v$), $\beta_v = \pi/8$, $m = 256$, $n = 16$, $f = 10$ Hz, $h = 5$ m and $s = 15$ m/s. These values are chosen for illustrative purposes. Note that the absolute size of the point cloud is not important here (or in Figure 4 below): we are only interested in its shape.




Fig. 4. Part of two combined point clouds at a high zoom level. In (a) we use $\alpha = 0$ and observe line patterns, in (b) we use $\alpha = \pi/4 = 45^\circ$ and get a more uniform sampling pattern. The parameters used here are realistic: $\beta_h = \beta_v = \pi/4$, $\beta_v = \pi/8$, $m = 1024$, $n = 128$, $f = 20$ Hz, $h = 20$ m, $s = 5$ m/s. The colours again indicate sweep number, in the Winter (blue to green) colour map.

sweep, and from 15 consecutive sweeps, for a mounting angle of $\alpha = \pi/6 = 30^\circ$.

Although (1) in principle shows the effect of the mounting angle α , the issue with $\alpha = 0$ is easily shown by examining a part of the full sampled point cloud up close, as in Figure 4a. The problem is that in the $\alpha = 0$ case the UAV's movement can compensate for the LiDAR's limited vertical field of view, but the resolution along scan lines (i.e. lines of constant elevation (index l)) cannot be improved. Using $\alpha = \pm\pi/2$ would roughly reverse the situation: we can get excellent sampling across scan lines, but poor sampling in between them. Other angles allow for more uniform sampling, as shown in Figure 4b where $\alpha = \pi/4$.

III. QUALITY OF A MOUNTING ANGLE

Figure 4 shows that certain angles are clearly suboptimal for most use cases. This raises the question which angle would be best, which in turns begs the question how we would measure the quality of an angle. Although multiple criteria are possible, we will only present a single one. We will focus on a smaller region of interest of the entire point cloud, our choice of which will be explained in the second subsection. For now, simply assume it is a bounded representative region of the part of the ground plane relevant for the full (infinite) point cloud.

1) Quality metrics: To ensure uniform density, we want every (not necessarily LiDAR-sampled) point \mathbf{S} in the region of interest to be close to a LiDAR point, i.e. the distance to \mathbf{S} 's nearest neighbour in the LiDAR point cloud should be minimised. This ensures we avoid gaps such as in Figure 4a.

We now take $p \in \mathbb{N}$ uniformly random samples in the region of interest R , yielding the set $S_p \subseteq R$. Then we choose some aggregation function $\mathcal{A}_s : \mathcal{P}_f(\mathbb{R}) \rightarrow \mathbb{R}$, i.e. a function that summarises a finite set of numbers into a single one. Here $\mathcal{P}_f(\mathbb{R})$ denotes the set of finite subsets of \mathbb{R} . Possible examples are the mean or a quantile. Then our cost for mounting angle α is

$$c_{\mathcal{A}_s}(\alpha; p) = \mathcal{A}_s(\{d(\mathbf{S}, L(\alpha)) \mid \mathbf{S} \in S_p\}), \quad (2)$$

where

$$d(\mathbf{S}, L(\alpha)) = \min \{\|\mathbf{S} - \mathbf{Q}\|_2 \mid \mathbf{Q} \in L(\alpha)\}$$

and we switched from quality to cost terminology to indicate we want to minimise $c_{\mathcal{A}_s}$.

Note that this is how we would calculate this in practice (given R), but from a theoretical point of view we can also use random variables. For example, if we use the mean \mathbb{E} for \mathcal{A}_s , then we get

$$c_{\mathbb{E}}(\alpha) = \mathbb{E}_{\mathbf{S} \sim \mathcal{U}(R)} [d(\mathbf{S}, L(\alpha))],$$

where $\mathcal{U}(R)$ denotes the uniform distribution on R . By the law of large numbers, $c_{\mathbb{E}}(\alpha; p) \rightarrow c_{\mathbb{E}}(\alpha)$ as $p \rightarrow \infty$. However, such a probabilistic analysis is beyond the scope of this paper and could form the basis for future work.

2) Region of interest: As region of interest we will take a rectangle with width w_R and height h_R centred at the origin of the ground plane:

$$R = \left[-\frac{w_R}{2}, \frac{w_R}{2}\right] \times \left[-\frac{h_R}{2}, \frac{h_R}{2}\right].$$

For the height we observe that the point cloud is periodic along the y -direction. Indeed this follows from (1) which also shows that the period is s/f , as can also be observed in Figure 4b where $s/f = 0.25$ m. Therefore we can take $h_R = s/f$ (or any multiple thereof). Next, we want to avoid that our region of interest has 'empty' regions at the sides (extreme x -values), devoid of LiDAR points for certain angles α , as this would mean that such angles automatically get high costs. In other words, we want to avoid penalising angles closer to $\pm\pi/2$ for the resulting reduced field of view. This can be achieved by taking $w_R = 2h \tan \beta_v$, which is the minimal (x -)width of the

point cloud of a single sweep for $\alpha = \pi/2$. In conclusion the region of interest is

$$R = [-h \tan \beta_v, h \tan \beta_v] \times \left[-\frac{s}{2f}, \frac{s}{2f} \right].$$

IV. SIMULATION

If we fix a LiDAR sensor (and its settings), we fix β_v , m , n and f . We can also fix $\widetilde{\beta}_h$, say to $\pi/4$. The remaining parameters are h and s , which depend on the UAV flight. The LiDAR sensor mounted on the UAV will normally not change, but the UAV will be used for different areas, requiring different flight heights and speeds. Therefore, we will determine the optimal mounting angle α as a function of h and s . But note that scaling h and s by the same factor will result in a point cloud scaled by that factor (i.e. $\mathbf{Q}(k, l; \lambda h, \lambda s) = \lambda \mathbf{Q}(k, l; h, s)$ in (1), for all $\lambda \in \mathbb{R}_{>0}$). Similarly, also the region of interest will be scaled with the same factor, and finally the same holds for the cost. This implies that the optimal mounting angle α only depends on the height-speed ratio h/s .¹

We then computed the cost of all angles between -90° and 90° (inclusive) using a step size of 0.25° for a total of 721 angles, and all height-speed ratios between 0.1 and 50 using a step size of 0.1 for a total of 500 height-speed ratios. For the other LiDAR system related parameters we had an Ouster OS1-128 in mind with $m = 1024$, $n = 128$, $f = 20$ Hz, $\beta_v = \pi/8$ and used $\widetilde{\beta}_h = \pi/4$. We used $p = 10^5$ uniform samples and aggregated them using the mean, median (50%-quantile), 95%-quantile and maximum as aggregation function \mathcal{A}_s . For each height-speed ratio, the optimal angle is then the one minimising this aggregate. However, we will focus on the results for the 95%-quantile, as this is a more robust version of the worst-case analysis (the maximum).

These computed costs have a huge peak for $\alpha = 0$, with values on average about 11 times larger than the median cost at a specific height-speed ratio. To avoid saturating our figure with this peak, we removed it by clipping the costs. These clipped costs are presented in Figure 5. We observe some complicated patterns, but in general see, consistent with our intuition from the introduction, that the minimal costs for each height-speed ratio are reached for angles near, but not at $\pm 90^\circ$. This becomes clearer in Figure 6 where we present the optimal mounting angle for each height-speed ratio.

The results for the other aggregation functions are quite similar. However, we do want to mention that when using the maximum function for sample aggregator the local maxima for $\alpha = \pm 90^\circ$ are almost as pronounced as the peak for $\alpha = 0$.

Since in practice the height-speed ratio is not completely constant during a flight, and might be slightly different than the planned value, it is more useful to provide a single mounting angle that works well for all height-speed ratios. We approach this by defining another optimisation function. For each height-speed ratio $r := h/s$ we can find the optimal mounting




Fig. 5. The 95%-quantile sample costs in metres, clipped to remove the peak at $\alpha = 0$, presented as a heat map. Note that we work with the height h at a fixed speed $s = 1$ m/s, instead of the height-speed ratio, as this allows interpreting the cost as a physical distance.

Optimal LiDAR mounting angles based on minimal 95%-quantile distances from samples




Fig. 6. The optimal mounting angle for each height-speed ratio, according to the 95%-quantile sample cost. To avoid the incorrect interpretation that the height-speed ratio represents time, we differentiate between metres in the y - and z -directions, yielding the unit $m_z/(m_y/s) = m_z/m_y \cdot s$ instead of just seconds.

angle $\alpha_{\text{opt}}(r)$ (as in Figure 6). The relative performance $q_{\mathcal{A}_s}(\alpha, r)$ of a mounting angle α at height-speed ratio r is then (well-)defined as the inverse quotient of α 's cost $c_{\mathcal{A}_s}(\alpha; h, s)$ at any height h and speed s with $h/s = r$, to the optimal (angle's) cost $c_{\mathcal{A}_s}(\alpha_{\text{opt}}(r); h, s) = \min_{\alpha'} c_{\mathcal{A}_s}(\alpha'; h, s)$:

$$q_{\mathcal{A}_s}(\alpha, r) = \frac{c_{\mathcal{A}_s}(\alpha_{\text{opt}}(r); h, s)}{c_{\mathcal{A}_s}(\alpha; h, s)} \in [0, 1]$$

(where compared to (2) we made the dependence on the number of samples p implicit, and the dependence on the height and speed explicit). For any aggregation function \mathcal{A}_r we can then compute the global relative performance

$$\bar{q}_{\mathcal{A}_r; \mathcal{A}_s}(\alpha) = \mathcal{A}_r(\{q_{\mathcal{A}_s}(\alpha, r) \mid r\})$$

by aggregating over all considered height-speed ratios r . The globally best mounting angle α_{opt} , independent of the height-speed ratio, is then obtained by maximising this relative performance:

$$\alpha_{\text{opt}} = \arg \max_{\alpha} \bar{q}_{\mathcal{A}_r; \mathcal{A}_s}(\alpha).$$

For similar reasons as mentioned before, we prefer to use the 5%-quantile for \mathcal{A}_r . In combination with the 95%-quantile for \mathcal{A}_s we then obtained $\alpha_{\text{opt}} = -84.5^\circ$, with relative performance $\bar{q}_{\mathcal{A}_r; \mathcal{A}_s}(\alpha_{\text{opt}}) = 95.12\%$. In other words, we do not lose much performance by switching to a fixed mounting

¹This of course also speeds up the optimisation considerably. Although this is not the focus of this paper, we mention that we also decreased the execution time by exploiting the symmetry of the point cloud, and by reasoning about upper bounds of the nearest neighbour distances in the cost function. The fact that such reasoning was already quite complicated, leads us to believe that analytical optimisation of the cost function is intractable.

TABLE I

THE OPTIMAL MOUNTING ANGLE α_{opt} (IN DEGREES) AND ITS RELATIVE PERFORMANCE $\bar{q}_{\mathcal{A}_r; \mathcal{A}_s}(\alpha_{\text{opt}}) = \max_{\alpha} \bar{q}_{\mathcal{A}_r; \mathcal{A}_s}(\alpha)$ (IN PERCENT) FOR THE DIFFERENT AGGREGATION FUNCTIONS \mathcal{A}_s AND \mathcal{A}_r

| \mathcal{A}_s | \mathcal{A}_r | | | |
|-----------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | mean | median | 5%-quantile | minimum |
| mean | -85.75° 98.84% | 86.75° 99.04% | -85.00° 97.40% | 85.50° 95.60% |
| median | 86.75° 99.15% | -86.75° 99.23% | -85.50° 98.21% | 85.50° 96.86% |
| 95%-quantile | -85.50° 97.84% | 86.75° 98.21% | -84.50° 95.12% | -84.50° 92.14% |
| maximum | -85.50° 90.20% | 86.75° 90.80% | -85.50° 81.56% | 83.50° 70.74% |

angle for all height-speed ratios. For different aggregation functions \mathcal{A}_r and \mathcal{A}_s we obtain quite similar optimal angles, at least up to the sign. The exact numbers can be found in Table I. So although we advise to use $\alpha = -84.5^\circ$ when possible, all $\alpha \in \pm[82^\circ, 88^\circ]$ (or so) should also work well.

V. EMPIRICAL VALIDATION

To experimentally validate our theoretically optimal (range of) angle(s), we equipped a DJI M600 UAV with an Ouster OS1-128 LiDAR and an SBG Quanta GNSS and IMU, used to register the different LiDAR scans. We then performed the same test flight twice, once with the LiDAR mounted at 0° (see Figure 1), and once with a mounting angle of $\alpha = 88^\circ$ (see Figure 7). In an older version of the simulation this angle was optimal. Its global relative performance in the current version is 92.15% for 95%-quantile sample and 5%-quantile height-speed ratio aggregation. The relative performance at the relevant height-speed ratio 6 (see below) is 96.64% for the 95%-quantile sample aggregator.

The parameters for the Ouster OS1-128 were $m = 1024$, $n = 128$, and $\beta_v = \pi/8$, but we now used the default frequency $f = 10$ Hz, as opposed to the 20 Hz used in the simulation. We also did not artificially limit the horizontal field of view of the LiDAR, but the sensor's maximum range of some 120 m will still implicitly limit it. The flight parameters were set as $h = 30$ m and $s = 5$ m/s, although of course in practice these fluctuate a bit. This should not be a problem as our mounting angle is designed to work well for all height-speed ratios. However, the yaw, pitch and roll are not constant as opposed to in the simplified model we used in the analysis and simulation. For the purpose of this validation we restricted ourselves to a single flight line (Figure 8). The yaw, pitch and roll for the 0° flight are shown in Figure 9.

Figure 10 shows the registered scans using our GPS and IMU. To measure the performance we opted to take a look at a 97 cm by 96 cm calibration plate in the centre of these point clouds. It is pictured in Figure 11 and indicated by the red arrows in Figure 10. The LiDAR sampled points on this plate are shown in Figure 12. There are no clear line patterns (which would go from the bottom-right to the top-left) visible in



Fig. 7. The LiDAR mounted at $\alpha = 88^\circ$. This was achieved by 3D printing with an inclination of 2° .



Fig. 8. The flight line over UGent campus Sterre, shown in a Google Maps orthophoto. Time is colour coded using the Turbo colour map, where blue indicates the start of the flight line and red the end.

Figure 12a, showing that this problem from Figure 4a is largely eliminated by the slight fluctuations in orientation. These will also affect the cost landscape. In particular the cost for $\alpha = 0^\circ$ will be noticeable lower than for perfectly stable flights. However, as flight instability is unpredictable, modelling it stochastically will also lead to a stochastic cost landscape, and stochastic optimal angles, significantly complicating the analysis. Therefore we leave dealing with this instability for future work.

Even though the optimal angles might shift somewhat by flight instability, we believe our conclusions will still, by and large, remain valid under general flight conditions. Indeed, in our test flight the number of points sampled on the calibration plate was significantly lower for $\alpha = 0^\circ$ than for $\alpha = 88^\circ$: 1518 versus 2453. This experiment thus demonstrates that we indeed get a significantly higher sampling density for $\alpha = 88^\circ$.

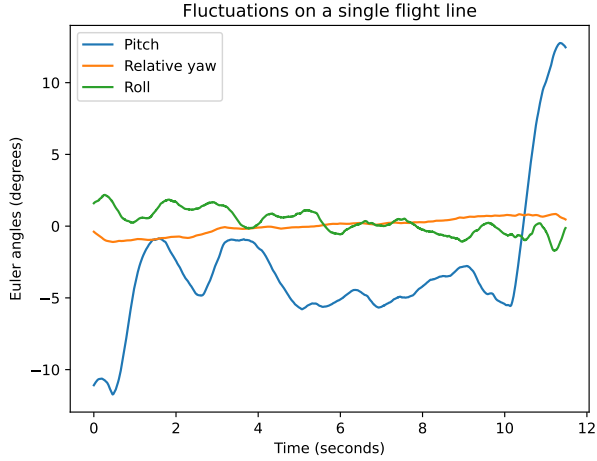


Fig. 9. The fluctuation of yaw, pitch and roll on the flight line of Figure 8. In the plot the yaw has been zero-centred by subtracting its mean of -42.3° with respect to true north. For comparison, the angular resolutions of the LiDAR are $\frac{2\pi}{m} \approx 0.352^\circ$ for azimuths and $\frac{2\beta_v}{n-1} \approx 0.354^\circ$ for elevations.




Fig. 10. The LiDAR point cloud from the flight line, for both mounting configurations: $\alpha = 0^\circ$ (a) and $\alpha = 88^\circ$ (b). The colour indicates the $z^{(G)}$ -value, using the Jet colour map. Both images are captured from the same virtual camera pose, clearly demonstrating the difference in field of view between the configurations. The red arrows point to the calibration plate.

VI. CONCLUSION AND FUTURE WORK

In this paper we investigated the optimal way to mount a spinning LiDAR rigidly to a UAV for use in land surveying of areas where the horizontal field of view of the LiDAR is of less importance, either because the scene itself is narrow, i.e. for swath mapping, or because of limitations of other sensors. We found that using the default configuration of 0° has a number of flaws. The obvious one is that many of the LiDAR samples




Fig. 11. Our calibration plate, placed on a small table to be able to detect it easily in the LiDAR data.




Fig. 12. The (zero-centre shifted WGS 84 / UTM 31N easting (x) and northing (y) coordinates of) points on the calibration plate, for mounting angles $\alpha = 0^\circ$ (a) and $\alpha = 88^\circ$ (b). Note that there is some noise (the borders of the rectangle are unclear), presumably due to imperfect calibration. The plate (on the small table) remains clearly distinguishable from the ground, however, so that the number of points should be unaffected.

lie outside of the region of interest. A second flaw is that when ignoring flight instability, we noticed very localised sampling patterns.

To get rid of both problems the best we can, we found that a mounting angle of -84.5° is optimal, although the sign is less important and similar angles will also perform well. This was corroborated by our real-world validation where we did indeed observe some advantages of angles closer to $\pm 90^\circ$ compared to 0° . In particular the sampling density in the region of interest was significantly higher. To get the same density at 0° one would need to fly over the area multiple times, resulting in a longer data acquisition time, especially if the UAV's batteries need to be recharged. As can be seen in Figure 12(a), the theoretical issue of the line patterns from Figure 4a is at least partially overcome by the fluctuations in orientation resulting from the UAV's flight, which in practice is not completely steady.

When considering wide scenes where the LiDAR is used in isolation, the horizontal field of view will matter more. When using $\alpha = 0^\circ$ we will need fewer flight lines to cover the entire scene, compared to when we would use $\alpha = -84.5^\circ$. But to reach the same sampling density as with our rotated LiDAR perhaps we might need to perform more flight lines for 0° regardless. Additionally the sampled points at higher absolute azimuths for $\alpha = 0^\circ$ will still have a number of

issues. For example, if we want to map an orchard or a forest consisting of multiple rows of trees, we will have occlusion problems when only performing a single flight line. Therefore, we would be forced to carry out more flight lines in any case, just like in the case of $\alpha = -84.5^\circ$. Next, assume we have multiple flight lines and want to perform sensor fusion with a camera with a more narrow field of view. In this case a LiDAR point which was not in the camera view at a certain time, might be encountered again later, on a different flight line. This can compensate for the lower density of $\alpha = 0^\circ$ on a single flight line. However, such LiDAR samples are actually still of limited use, because of the temporal difference and possible occlusion issues due to the different poses of the LiDAR and camera at the time of capturing the point. Therefore it is not clear which mounting angle should be preferred in these situations and a future extension of this work would be to more thoroughly examine the situation of flight trajectories with multiple (parallel) flight lines.

Another future extension of this work is to validate the results for various practical use-cases: for instance check that $\alpha = -84.5^\circ$ results in higher quality digital elevation models or in more reliable object detection than $\alpha = 0^\circ$.

ACKNOWLEDGEMENT

This work was financially supported by the following project: COMP4DRONES ECSEL Joint Undertaking (JU) under grant agreement No 826610.

The authors are grateful to anonymous reviewers for their comments and suggestions.

REFERENCES

- [1] N. Metni and T. Hamel, "A UAV for bridge inspection: Visual servoing control law with orientation limits," *Automation in construction*, vol. 17, no. 1, pp. 3–10, 2007.
- [2] J. Nikolic, M. Burri, J. Rehder, S. Leutenegger, C. Huerzeler, and R. Siegwart, "A UAV system for inspection of industrial facilities," in *2013 IEEE Aerospace Conference*. IEEE, 2013, pp. 1–8.
- [3] G. Morgenthal and N. Hallermann, "Quality assessment of unmanned aerial vehicle (UAV) based visual inspection of structures," *Advances in Structural Engineering*, vol. 17, no. 3, pp. 289–302, 2014.
- [4] C. A. F. Ezequiel, M. Cua, N. C. Libatique, G. L. Tangonan, R. Alampay, R. T. Labuguen, C. M. Favila, J. L. E. Honrado, V. Canos, C. Devaney *et al.*, "UAV aerial imaging applications for post-disaster assessment, environmental management and infrastructure development," in *2014 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 2014, pp. 274–283.
- [5] F. Outay, H. A. Mengash, and M. Adnan, "Applications of unmanned aerial vehicle (UAV) in road safety, traffic and highway infrastructure management: Recent advances and challenges," *Transportation research part A: policy and practice*, vol. 141, pp. 116–129, 2020.
- [6] G. Luzardo, M. Vlaminck, D. Lefkaditis, W. Philips, and H. Luong, "GPS-assisted feature matching in aerial images with highly repetitive patterns," in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 2021, pp. 8348–8351.
- [7] S. Esposito, P. Fallavollita, W. Wahbeh, C. Nardinocchic, and M. Balsia, "Performance evaluation of UAV photogrammetric 3D reconstruction," in *2014 IEEE Geoscience and Remote Sensing Symposium*. IEEE, 2014, pp. 4788–4791.
- [8] M. Vlaminck, H. Luong, and W. Philips, "Multi-resolution ICP for the efficient registration of point clouds based on octrees," in *2017 Fifteenth IAPR International Conference on Machine Vision Applications (MVA)*. IEEE, 2017, pp. 334–337.
- [9] J. Sofonia, Y. Shendryk, S. Phinn, C. Roelfsema, F. Kendoul, and D. Skocaj, "Monitoring sugarcane growth response to varying nitrogen application rates: A comparison of UAV SLAM LiDAR and photogrammetry," *International Journal of Applied Earth Observation and Geoinformation*, vol. 82, p. 101878, 2019.
- [10] K.-W. Chiang, G.-J. Tsai, Y.-H. Li, and N. El-Sheimy, "Development of LiDAR-based UAV system for environment reconstruction," *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 10, pp. 1790–1794, 2017.
- [11] K. Liu, X. Shen, L. Cao, G. Wang, and F. Cao, "Estimating forest structural attributes using UAV-LiDAR data in Ginkgo plantations," *ISPRS journal of photogrammetry and remote sensing*, vol. 146, pp. 465–482, 2018.
- [12] Y.-C. Lin, Y.-T. Cheng, T. Zhou, R. Ravi, S. M. Hasheminasab, J. E. Flatt, C. Troy, and A. Habib, "Evaluation of UAV LiDAR for mapping coastal environments," *Remote Sensing*, vol. 11, no. 24, p. 2893, 2019.

Laurens Diels received the master's degrees in Mathematics and Artificial Intelligence from KU Leuven in 2017 and 2020, respectively. He is currently pursuing the PhD degree in Computer Science Engineering at the Image Processing and Interpretation research group (IPI) and UAV Research Centre (URC) of Ghent University. His research focuses on the processing of point clouds captured from UAV-mounted LiDARs and the fusion with other sensors.

Michiel Vlaminck received his PhD degree in Computer Science Engineering from Ghent University in 2020. Since then, he is working as a postdoctoral researcher at the UAV Research Centre of Ghent University. He is currently working on the topic of 3D scene reconstruction using active depth sensors. His research focuses on applications in the domains of augmented reality, autonomous robots and UAV.

Bart De Wit received his master's degree in Geography in 2008, and works since then at the department of Geography at Ghent University. He participated in several projects that are geomatics related, including GIS, cartography, web based GIS, land surveying, mobile mapping, photogrammetry and remote sensing. Bart has 9 years of experience in flying with drones, and processing data for geomatics related projects: creating 3D models, working with multispectral cameras and LiDAR.

Wilfried Philips (Senior Member, IEEE) was born in Aalst, Belgium, in 1966. He received the Diploma degree in Electrical Engineering and the PhD degree in Applied Sciences from Ghent University, Ghent, Belgium, in 1989 and 1993, respectively. He is a Senior Full Professor with the Department of Telecommunications and Information Processing, Ghent University, where he heads the Image Processing and Interpretation Research Group. He also leads the activities in image processing and sensor fusion within the research institute IMEC. His main research interests include image and video quality improvement and estimation, real-time computer vision, and sensor data processing. He is also a Co-Founder of the Senso2Me company, which provides Internet of Things solutions for elderly care.

Hiep Luong received the PhD degree in Computer Science Engineering from Ghent University, in 2009. He has been working as a Postdoctoral Researcher and the Project Manager of Image Processing and Interpretation (IPI), an imec Research Group, Ghent University. He is currently a Professor at IPI and leads the UAV Research Centre of Ghent University. His research interests include image and real-time video processing for various fields, such as HDR imaging, (bio) medical imaging, depth and multi-view processing, and multi-sensor fusion for UAV and AR applications.