

Water Detection From Downwash-induced Optical Flow for a Multirotor UAV

Ricardo Pombeiro, Ricardo Mendonça, Paulo Rodrigues, Francisco Marques,
André Lourenço, Eduardo Pinto, Pedro Santana and José Barata

Abstract—Terrains of different nature exhibit distinct structural dynamics when exposed to wind due to their own intrinsic material composition. This translates into peculiar optical flow patterns that can be used to identify the terrain type. In this sense, this paper proposes an active vision-based water detection model that exploits the predictable optical flow patterns induced by the downwash effect of vertical take-off and landing Unmanned Aerial Vehicles (UAV). To determine whether a water surface is below the UAV, the system tracks the optical flow in the video feed captured by a downward-looking camera. Then, an histogram of optical flow orientation is built and compared against a model histogram, given the expected effect induced by the downwash effect. The histograms are compared and similarity between both histograms is used as the likelihood of the terrain as being covered by water. The resulting classification can be used to guide the landing of the UAV or to produce cost maps supporting ground vehicles' safe navigation. The model was successfully validated on 20 videos acquired with an hexacopter while hovering above sandy, grassy, and water-covered terrains.

I. INTRODUCTION

Water detection is a crucial capability for a wide range of autonomous vehicles, either for ground vehicles to avoid water bodies, aerial vehicles to determine suitable landing areas, or surface vehicles to detect safe passageways.

Laser scanners have proven to be useful to detect water as no-return situations usually hint the presence of a water body [1], [2], [3]. However, down to a certain depth, like on small puddles or shallow waters, some light may be reflected off the underwater ground surface. This creates a false indication that a dry and solid surface may be present where a potentially hazardous water body may be, rendering the approach, by itself, insufficient.

Several vision-based solutions have been proposed as well. Stereo vision methods for water detection rely to a large extent on the fact that object reflections are reconstructed as 3D points below the surface level [4], [5], [6]. However, this technique is limited to the near-field and suffers from common stereo-vision pitfalls, such as sensitivity to lighting and weather conditions. Monocular vision cues have also been proposed for water detection. For instance, symmetry operators can be used to detect reflections on the water surface [7], [8]. Textureless and high brightness image patches could also be indicative of water presence [5]. Detecting sky reflections

R. Pombeiro, R. Mendonça, P. Rodrigues, F. Marques, A. Lourenço, E. Pinto, and J. Barata are with CTS-UNINOVA, Universidade Nova de Lisboa (UNL), Portugal. E-mails: {rjrp, rmm, pmr, fam, afl, emp, jab}@uninova.pt

P. Santana is with ISCTE - Instituto Universitário de Lisboa (ISCTE-IUL), Portugal and Instituto de Telecomunicações, Portugal. E-mail: pedro.santana@iscte.pt




Fig. 1. The UAV's propulsion system used as an active perception mechanism to identify water bodies from their susceptibility to the vehicle's propeller whirlwind. (a) Downwash effect caused by the UAV hovering above a water body. (b) Video frame captured by the UAV's downward-looking camera. (c) Optical flow induced by the UAV's downwash effect (depicted by the blue lines on the captured image).

[9] and colour variations as the vehicle moves [10] are also interesting cues. These cues rely on the assumption that the camera is roughly pointing towards the horizon line, which is considerably different from the bird's-eye view usually available to aerial vehicles.

Rather than relying on a set of known cues, complex sets of perceptual features can also be employed by learning image classifiers from offline datasets [1], [8], [11]. However, reflections, foliage or other objects will impair the classifiers' accuracy. To mitigate this problem, we have proposed a method for a surface vehicle to supervise, based on its own sensors, the online learning of water/land aerial image classifiers to be used by an Unmanned Aerial Vehicle (UAV) teammate [12]. However, the aerial vehicle's dependency on a surface robotic teammate renders this approach mostly limited to marsupial robotic systems [13].

Motionless water surfaces do not exhibit a detectable texture if there is no natural source of disturbance. However, having the UAV approach the water body its rotor blades' whirl

provoke a significant agitation on its surface (see Fig. 1(a)). This paper builds on this to devise a strategy that actively induces a somewhat predictable movement on the water surface through the vehicle's downwash, facilitating the water/land classification. In previous work [14], we exploited the chaotic nature of wind-rippled water bodies to segment images into water and land categories, showing good results even beyond the near-field. In particular, to detect the chaotic pattern, the optical flow field computed over several input video frames is analysed through an entropy-based heuristic. Herein we study how to enhance the uniqueness of the optical flow signature on water surfaces by actively altering their motion field with the UAV's downwash effect (see Fig. 1(c)).

This paper is organized as follows. Section II describes the proposed model and its experimental validation is presented of Section III. Finally, conclusions and future work directions are provided in Section IV.

II. WATER DETECTION METHOD

The downwash effect induced by each rotor of the aerial vehicle turns the chaotic trajectory of the water particles into more predictable motion patterns. Concretely, the air turbulence induces a quasi-radial dispersion of water particles outwards from a pivot point that vertically matches the rotor's position. Such phenomena is easily imprinted in the optical flow captured by a downward-looking camera, as depicted in Fig. 1(c). Ripples induced by wind currents have a signature that differs considerably from the ones induced by the downwash effect and, so, can be distinguished with the method herein proposed. On solid ground, covered by sand or dirt, the rotor's downwash influence is comparatively lower. Such variation on the resulting optical flow is a pivotal visual cue towards a robust self-supervised water/land segmentation for autonomous aerial vehicles.

A. Method Overview

The method assumes the UAV hovers at an altitude roughly of 2 m, as static as possible, above a target surface. Specifically, the UAV needs to actively move towards the region, lowering down to an altitude to which the downwash effect is felt on the putative water surface (see Fig. 2). Remaining stable at that position, the UAV acquires a short video with its downward-looking camera, to look for a water-characteristic signature of the downwash effect on the surface so to classify it as either water or non-water. Once a classification is issued by the proposed method, the UAV takes the decision of landing or pursuing the next sampling point, depending on the mission specification.

The proposed water detector starts by computing the sparse optical flow on a set of n consecutive frames captured by the UAV's downward-looking camera. The sparse optical flow is computed with a pyramidal implementation of the Lucas-Kanade method. The optical flow trackers are initialised on the vertices of a $w \times h$ regular grid, covering the whole image space. Then, the trackers follow the motion field in the subsequent $n - 1$ frames (see Fig. 1(c)). If the surface below




Fig. 2. A rendering of a typical mission (a) UAV getting an aerial panorama of the environment. (b) UAV receives a waypoint to access (represented as the red blob). (c) The UAV downwash generating turbulence on the water's surface. (d) The UAV's FOV viewing the expected quasi-radial effect on the water's surface.

the UAV is covered by water, the trackers are expected to move radially with respect to the image centre. Therefore, to check for the presence of water, the method seeks for this signature in the trackers' paths in the image space. Following the approach used in our previous work [14], fake *weak* features are overlaid in each frame to prevent trackers from drifting away on homogeneous and textureless terrain regions. To filter out deceptive motion bias caused by sporadic external winds (i.e., not intentionally induced), the optical flow computation is repeated over r iterations of consecutive video segments. The actual number of iterations depends on the estimated classification accuracy gained so far (see Section II-D).

B. Optical Flow Descriptor

The underlying assumption of the proposed method is that the optical flow induced by the UAV's downwash on a water surface is distinguishable from any other surface's. To classify the optical flow, a simple global descriptor is computed. A global descriptor is image-wise, which means that the method classifies the image as a whole and not its parts. As the UAV is assumed to be hovering at a very low altitude, the area of the surface being imaged is quite small and, thus, an image-wise classification suffices in most situations.

The image descriptor is a unidimensional histogram of the optical flow directions over the whole set of $w \cdot h$ trackers. The histogram is uniformly spaced over the interval $[0^\circ, 360^\circ]$ with a set of 16 bins, normalised so as to ensure that each bin exhibits a frequency within the interval $[0, 1]$. This motion field's orientation quantisation provides an adequate speed-accuracy trade-off for the tested dataset.

The orientation of a given optical flow tracker is defined as the angle between the line segment connecting the first and last position of the tracker and the x-axis at the centre of the first input image of the video segment. Fig. 3 illustrates the orientation computation process.




Fig. 3. Calculation of the optical flow (filled line) orientation when computed over n consecutive frames. The dashed line connects the initial (filled circle) and final positions (arrow head) of the optical flow tracker. The optical flow orientation, θ , is computed as the angle between the dashed line and the x-axis of the image plane.




Fig. 4. Optical flow trackers' orientation histogram obtained from observing the downwash effect on water-covered terrain. This histogram was computed for the video input depicted in Fig. 1.

C. Water Signature Detection

The perceived surface's motion field is influenced not only by its materials' intrinsic dynamics (e.g., waving of grasses, water ripples) but also by the camera's relative motion, either caused by an external disturbance or the UAV's inaccurate motion control. For instance, when the aerial vehicle hovers above ground surfaces while gently moving linearly, the orientation histogram exhibits a higher frequency on the orientation that matches the linear motion.

Ideally, when hovering above water surfaces subject to the downwash effect, the optical flow orientation histogram is expected to approximate a uniform distribution due to the radial dispersion of trackers. However, quantisation biases and the influence noise and inevitable UAV displacements have on the optical flow result in non perfectly uniform histograms

(see Fig. 4). Still, as results will show, the likelihood that the surface below the UAV is covered by water can be obtained by checking how similar the normalised orientation histogram, computed from the video segment, is from the ideal normalised uniform histogram. Similarity between both histograms is computed as $1 - d$, where d is the *Bhattacharyya* distance. Hence, similarity is 0 for a perfect mismatch and 1 for a perfect match.

Changes in altitude induce a radial optical flow - the expected water signature - even on surfaces not covered by water. Therefore, it is mandatory that the UAV remains at the same altitude during the video acquisition process. Nevertheless, as optical flow trackers are computed only over a few frames, the system is not sensitive to the slow altitude changes naturally present when hovering outdoors.

D. Terrain Labelling

The presence of multiple potential terrain classes on the video input, the noise impinging the sensor, the unexpected UAV motion, the weak downwash effect, and the ripples induced by wind modulate the orientation histogram in unpredictable ways. As a result, the classification of the surface based on the n frames may be inconclusive and the whole process is repeated for another r consecutive video segments of n frames. That is, r video segments of n frames are captured, the orientation histogram of each video segment is computed and its similarity with the ideal uniform histogram determined. Finally, an average similarity is computed over all video segments.

Intermediate average similarity values, empirically defined as encompassed by the interval $[0.45, 0.55]$, are taken as non-confident. Low (< 0.45) and high (> 0.55) average similarity values correspond to confident classifications, either non-water or water, respectively.

When the detector is unable to produce a confident output, the detector will repeat the process until a decision is taken. This process aims at maintaining computation low and responsiveness high for most cases by processing only a reduced number of frames. Computation is increased and responsiveness lowered only when facing low confidence situations.

III. EXPERIMENTAL RESULTS

A. Data Gathering

To validate the proposed water detector, the aerial vehicle *Vigil R6* (see Fig. 5) was used. Developed for the RIVER-WATCH experiment [12], this UAV uses the VRBRAIN control from Virtual Robotix for low level processing, which encompasses an onboard Inertial Measurement Unit (IMU) and a Ublox Global Position System (GPS) for pose estimation. An Hardkernell's Odroid-XU device is used as the high level processing unit. This unit is equipped with an Exynos octa-core ARM-based 1.8 GHz CPU, and it runs the Indigo distribution of the Robot Operating System (ROS) [15] on the Linux Xubuntu 14.04 distribution. OpenCV [16] was used for low-level image processing routines. The UAV is equipped with an SJ4000 HD camera mounted on an active gimbal to




Fig. 5. Vigil R6 during the field trials.

ensure that its optical axis is parallel to the gravity vector. The water detector was empirically parametrised as follows: $w = 640$, $h = 480$, $n = 5$, and $r = 7$. $n = 5$ was chosen to represent approximately 1/3 of a second, considering the input video's frame rate.

A set of 20 videos was acquired by flying the UAV over diverse locations under various wind conditions (see Fig. 7). This set includes video sequences with the presence of water, grass, sand, and dirt (see Fig. 7). Videos were acquired at 17 frames per second with an image resolution of 640×480 . With this resolution, the system takes on average 3.2s to reach a decision on an external i5-based dual-core 1.8GHz CPU. A GPU-based optical flow implementation would reduce considerably the computation time.

B. Results

Fig. 6 presents the optical flow and orientation histograms of representative videos depicting surfaces covered by water (row 1), tall and sparse vegetation (row 2), small and dense vegetation (row 3), and bare surfaces (rows 4 and 5). As expected, the figure shows that the absence of water renders the orientation histograms mostly unimodal, even in the presence of a motion field induced by waving vegetation and UAV slight motion. Contrarily, the orientations' frequency is spreader in the water surface histograms, closer to the expected uniformity (see also Fig. 4). However, the histograms obtained with water-covered surfaces are not perfectly uniform, which owes mostly to the difficulty of tracking the radial water ripples throughout their lifetime and also to the inability of the UAV to maintain a static hovering.

Despite the difficulties in obtaining perfectly uniform optical flow orientation histograms in water-covered surfaces, the proposed method was capable of correctly classifying all tested videos, except video *Grass C*, in which ambiguous results were obtained. In this case, the presence of shadowed tall dense vegetation induced a motion field that partially fooled the system. Increasing the image resolution and sampling rate,




Fig. 6. Data set optical flow (left column) and corresponding orientation histograms (right column). Video segments from top to bottom: Water G, Grass A, Grass F, Dirt A, Dirt B.

enlarging the number of frames processed by the method, filtering out the motion field induced by UAV's ego-motion, and automatically learning the average similarity thresholds from datasets are possible developments to improve the accuracy of the system. The average similarity values used for the final classification of the surface are presented in Fig. 7.

The system exhibited different levels of confidence on its final classification across the test dataset. Recall that confident outputs are those generated from extreme average similarity values. On water environments, videos *Water A*, *Water B*, *Water C*, *Water E*, *Water H*, and *Water I* were correctly classified as water with good confidence. Videos *Water D*,

Water G, and *Water J* also got correctly classified, but with scores closer to the indecision region, particularly the video *Water D*. This was caused by the presence of wind that partially disrupted the downwash effect on the water surface. On grassy environments, the system produced confident correct classifications for videos *Grass B*, *Grass D*, and *Grass E*. Videos *Grass A* and *Grass F* were also correctly classified but with lower confidence values. This is due to the fact the optical flow on the high grass blades affected by the downwash effect partially resembles the one induced by water under the same conditions. On bear ground surfaces, videos *Dirt A*, and *Dirt C* were confidently classified correctly as non-water. Conversely, videos *Dirt B* and *Sand A* generated correct yet closer to indecision results. This is explained by not so gentle altitude changes during the video acquisition, which caused momentary radial optical flow.

These results confirm that the downwash effect induced by the UAV's rotors shapes the surface's motion field in a way that it facilitates the separation of both water and non-water classes in the robot's visual input.

IV. CONCLUSIONS

A vision-based water detection model that exploits the predictable optical flow patterns induced by the downwash effect of vertical take-off and landing UAVs was presented.

A set of experimental results show the ability of the method to discriminate water surfaces from other typical surfaces found in natural environments, despite the continuous motion exhibited by the UAV while hovering. The proposed method enhances the autonomy level of aerial vehicles by giving them the ability to better evaluate whereto land safely. The method relies on the expected radial optical flow pattern induced by the downwash effect, thus discarding the need for knowing beforehand the hard-to-predict appearance of water or the need for learning image classifiers, which can easily get outdated. Moreover, by relying solely on visual input, the method is well suited for small sized aerial vehicles that do not afford enough payload capacity for additional sensors.

As future work we expect to implement explicitly active perception [17] policies so as to better shape the visual input and, consequently, increase the discriminatory power of the proposed method. We would also like to improve ground mobility by integrating the active perception concepts herein studied with the ones developed in our previous work on haptic-based active perception for vegetation detection in ground vehicles [18].

ACKNOWLEDGEMENT

This work was partially funded by the CTS multi-annual funding, through the PIDDAC Program funds.

REFERENCES

- [1] Matthies, L. H., Bellutta, P., and McHenry, M. (2003). Detecting water hazards for autonomous off-road navigation. In SPIE Proceedings, Vol. 5083, pp. 231-242. DOI: 10.1117/12.496942. International Society for Optics and Photonics.
- [2] Hong, T. H., Legowik, S., and Nashman, M. (1998). Obstacle detection and mapping system (No. NISTIR-6213). US Department of Commerce, Technology Administration, National Institute of Standards and Technology.
- [3] Elkins, L., Sellers, D., and Monach, W. R. (2010). The Autonomous Maritime Navigation (AMN) project: Field tests, autonomous and co-operative behaviors, data fusion, sensors, and vehicles. *Journal of Field Robotics*, 27(6), pp. 790-818.
- [4] Rankin, A., Huertas, A., and Matthies, L. (2005). Evaluation of stereo vision obstacle detection algorithms for off-road autonomous navigation. Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space Administration.
- [5] Yao, T., Xiang, Z., Liu, J., and Xu, D. (2007). Multi-feature fusion based outdoor water hazards detection. In Proceedings of the IEEE International Conference on Mechatronics and Automation (ICMA), pp. 652-656. IEEE.
- [6] Murray, D., and Little, J. J. (2000). Using real-time stereo vision for mobile robot navigation. *Autonomous Robots*, 8(2), pp. 161-171.
- [7] Subramanian, A., Gong, X., Riggins, J. N., Stilwell, D. J., and Wyatt, C. L. (2006). Shoreline mapping using an omni-directional camera for autonomous surface vehicle applications. In Proceedings of the International Conference IEEE/MTS OCEANS. IEEE.
- [8] Gong, X., Subramanian, A., and Wyatt, C. L. (2007). A two-stage algorithm for shoreline detection. In Proceedings of the IEEE Workshop on Applications of Computer Vision (WACV). IEEE.
- [9] Rankin, A. L., Matthies, L. H., and Bellutta, P. (2011). Daytime water detection based on sky reflections. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pp. 5329-5336. IEEE.
- [10] Rankin, A., and Matthies, L. (2010). Daytime water detection based on colour variation. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 215-221. IEEE.
- [11] Xu, A., and Dudek, G. (2010). A vision-based boundary following framework for aerial vehicles. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 81-86. IEEE.
- [12] Pinto, E., Marques, F., Mendona, R., Loureno, A., Santana, P., and Barata, J. (2014). An Autonomous Surface-Aerial Marsupial Robotic Team for Riverine Environmental Monitoring: Benefiting from Coordinated Aerial, Underwater, and Surface Level Perception. In Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO). IEEE.
- [13] Marques, F., Loureno, A., Mendona, R., Pinto, E., Rodrigues, P., Santana, P., and Barata, J. (2015). A Critical Survey On Marsupial Robotic Teams for Environmental Monitoring of Water Bodies. In Proceedings of the International Conference IEEE/MTS OCEANS. IEEE Press.
- [14] Santana, P., Mendonca, R., and Barata, J. (2012). Water detection with segmentation guided dynamic texture recognition. In Proceedings of the IEEE International Conference on Robotics and Biomimetics (ROBIO), pp. 1836-1841. IEEE.
- [15] Quigley, M., Conley, K., Gerkey, B., Faust, J., Foote, T., Leibs, J., Wheeler, R., and Ng, A. Y. (2009). ROS: an open-source Robot Operating System. In Proceedings of the ICRA Workshop on Open Source Software, Vol. 3, No. 3.2.
- [16] Bradski, G., and Kaehler, A. (2008). *Learning OpenCV: Computer vision with the OpenCV library*. O'Reilly Media, Inc..
- [17] Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE* 76(8), pp. 996-1005.
- [18] Baleia, J., Santana, P. and Barata, J. (2015). On Exploiting Haptic Cues for Self-Supervised Learning of Depth-Based Robot Navigation Affordances. *Journal of Intelligent & Robotic Systems*. DOI: 10.1007/s10846-015-0184-4. Springer.




Fig. 7. Representative frames of the test dataset acquired with the UAV hovering at 2 m above the surface under analysis. The Average Similarity (AS) computed per video so as to generate a water/non-water classification is also represented. Note that water-covered surfaces are the ones producing higher AS values.