

Detection of Interest Points in Turbid Underwater Images

Rafael Garcia, Nuno Gracias
Computer Vision and Robotics Group
University of Girona, Girona, 17001 Spain
Email: {rafa, ngracias}@eia.udg.edu

Abstract—Our research is motivated by an evident lack of evaluation of recent image matching techniques for applications in underwater vision. This paper is a first step in this direction. This work compares the performance of popular salient keypoint detectors on images degraded by turbidity. We show that, as opposed to over-land, on images acquired in water medium, Hessian-based approaches outperform their Laplacian and Harris counterparts. Fast Hessian, the detector of the Speeded Up Robust Features (SURF) matching technique, is recognized to be the best method for scale-invariant detection. Conversely, when invariance to scale is not required, a combination of standard Hessian and Harris with sub-pixel accuracy and non-maxima suppression is more accurate. The objective of our work was also to create and distribute a reference set of turbid images, which can be used to evaluate processing, detection, description and matching techniques for underwater applications. We present a collection of 36 images acquired by a specially designed trinocular system under 12 gradually increasing turbidity levels. We also draw attention to image quality assessment method called SSIM, Structural SIMilarity index, which reliably quantifies degradation of image quality caused by turbidity. As a whole, the major goal of this paper is to provide an updated reference for researchers dealing with keypoint detection in underwater imaging.

I. INTRODUCTION

Detection of interest points in images has been an area of active research for decades. Local features proved to be suitable for matching and recognition in a diversity of applications. Existing extraction algorithms range from classic Harris corner detector [1], through invariant to scale Laplacian-of-Gaussian [2] up to intensity region stability measures [3]. Development of new techniques and constant upgrade on the classical ones aim at greater invariance to geometric and photometric deformations of images [4]. Precise and accurate detection of the same salient points contributes to the success of image matching, especially when the appropriate descriptor is used [5]. The demand of detectors to tolerate changes in translation, rotation, scale, camera viewpoint, illumination and noise gives rise to the need of their evaluation.

The goal of our work is to compare the most popular and recent detection techniques and their original ancestors on images acquired under water. Working in the area of seabed mosaicing we encountered a complete absence of any works concerning evaluation of matching techniques for underwater imagery. However, the number of applications where underwater images should be searched for matchable features grows larger and larger each year. This process is stimulated

by the evolution in the area of underwater robotics. Navigation of underwater vehicles [6], construction of the sea bottom visual maps [7], building panoramas of ocean sceneries [8] (see Fig.1) and 3-dimensional reconstruction of the underwater terrain [9] include image matching as an essential component. In most cases, keypoint detectors and descriptors are chosen in underwater applications on basis of their performance on over-land images.

From the year 2000 prominent works on detector evaluation keep bringing to the forefront best of the contemporary to them techniques. Thus, earlier comparisons [10], which investigated the “non-invariant to scale detectors”, concentrated on and promoted the classic Harris corner detector [1], which made it the most popular single-scale detector both over-land and under water. Its close counterparts like Hessian and Laplacian were left behind until the introduction of scale and affine invariant methods [4]. Recently, the majority of popular detection techniques [5] are based on Harris, Hessian and Laplacian, and are variations of the original methods. Recent comparisons focus on evaluation of invariance under deformations [4]. At present, Difference-of-Gaussian (DoG) and affine invariant Harris are claimed to be generally more robust, hence being the most widely used techniques. This coincides with preferences seen in recent underwater applications, where the most popular detectors are various versions of Harris ([6], [11], [12]) and DoG ([13], [14]). We selected 11 popular recent techniques for evaluation: standard Harris, Hessian and Laplacian [15], scale and affine invariant Harris and Hessian in different combinations [4], DoG, FastHessian [16] and the scale invariant edge-based technique SedgeLap [17]. The performance of all the detectors is tested for the special case of underwater images degraded by turbidity under controlled conditions.

The performance in underwater images should be investigated separately, since photometric artifacts induced by the water medium cannot be encountered over-land (apart from instances with fog [18] or atmospheric haze [19]). These artifacts come from an interaction of light with the water. The attenuation of the light as it travels in water consists of absorption and scattering. If scattering can be neglected, the absorption effect causes loss of contrast and domination of green (moderate depth and big particles) or blue (deep ocean and minor particles) wavelengths. Such image degradations can be successfully eliminated by a preprocessing based on, e.g., CLAHE [20]. However, in turbid waters, with constantly




Fig. 1. Application of matching techniques to underwater images: detail of a ship wreck mosaic.

moving particles of different sizes, the scattering effect dominates, causing significant loss in the visual image quality.

An important outcome of our work is a reference set of turbid images acquired using special experimental setup. Our setup includes a trinocular underwater camera system, a water tank with interior mate black painting, a natural scenery built from real sand, rocks and shells, an artificial underwater light source and skimmed-milk as turbidifier. A second contribution is a proposal of a methodology to assess the quality of turbid images, based on the Structural SIMilarity Index, introduced by Wang [21]. We demonstrate that the widely used MSE (mean squared error) and PSNR (peak signal-to-noise ratio) fail against turbidity-induced image degradation. Measures such as the concentration of suspended solids or nephelometric turbidity cannot be directly related to visual quality of images, since the latter depends also on the scattering volume, i.e., the intersection of the illuminated underwater volume and the cone defined by the camera's field of view [22]. Finally we evaluate the selected detection techniques calculating the percentage of repeatedly detected keypoints in turbid images versus *ideal* features detected in the *clear water* image, provided some preset localization error. The repeatability results are demonstrated for different values of localization error, non-maxima suppression radii and varying maximum number of keypoints to detect. Our results clearly point out the detectors more robust to turbidity. In conclusion we discuss reasons of better or poorer performance of different techniques and advice methods more appropriate for keypoints detection in the special case of underwater imagery.

II. DETECTORS

Descriptions of all the selected detection techniques can be found in many available sources. Hence, we present each method very briefly, concentrating instead on its performance as reported in previous comparative studies.

The *Harris* feature detector, presented by Harris and Stephens in 1988 [1], is based on the auto-correlation matrix. If both eigenvalues of the matrix for a given point are large, distinct, positive values, this indicates a corner. To localize Harris keypoints with *sub-pixel* accuracy, we use a 2D fit of a paraboloid to the keypoint neighborhood and take its maximum as an accurate position. Features detected by Hessian and Laplacian are refined in the same way. Harris feature detector responds to corners and highly textured points, it is invariant to

rotation, translation and linear changes in lighting. However, the classic version of Harris is not invariant neither to scale nor to affine transformations [23]. Still, Harris demonstrated an excellent performance in the evaluation of single-scale interest point detectors presented by Schmid et al. in 2000 [10].

Hessian is one of the first corner detectors, developed in 1978 by Beaudet [15], who proposed a rotationally invariant measure of saliency given by the determinant of the Hessian matrix. A local maximum of the determinant indicates a feature. Laplacian differs from Hessian in that it calculates the trace of the same matrix. The second derivatives used in Hessian matrix give strong responses on blobs and ridges. Hessian keypoints are considered to be more stable [16] than their Laplacian counterparts. The determinant formula includes mixed second-order derivatives and does not give large values where signal changes only in one direction, i.e. on edges, while the formula of trace accepts such features.

Harris Laplace and Hessian Laplace (*HarLap* and *HesLap*) [23] are adaptations of Harris and Hessian detectors respectively, to scale changes via incorporating the Laplacian-of-Gaussian (LoG) scale-space [2]. HarHes detector [24] is a combination of HarLap and HesLap. Harris Affine and Hessian Affine [23] (*HarAff* and *HesAff*) are HarLap and HesLap respectively, upgraded by affine normalization, that can significantly shift the original points. The DoG detector of SIFT is a close approximation to the LoG via replacing the second-order derivation by subtraction inside the Gaussian scale space. To detect features, each sampled point in the DoG space is compared to its $3 \times 3 \times 3$ scale-space neighborhood. A keypoint is selected only if an extremum is detected. Once a candidate keypoint has been found, a 3D quadratic function is fitted to the nearby data to determine the precise location of the maximum. Finally, the ratio of principal curvatures is checked via Hessian to discard points on edges [25].

The comparison of Mikolajczyk et al. from 2004 [23] shows that HarLap and HesLap are better than HarAff and DoG over scale change. In a more recent review by Moreles and Perona [5], where HarLap, HesLap, HarAff, HesAff and DoG are compared in combination with various description techniques on human-made objects, HesAff and DoG perform consistently better with all the descriptors.

SURF (speeded up robust features) was presented by Bay et al. [16] in 2006. The detector of SURF, *FastHessian*, uses the determinant of Hessian matrix for the localization of keypoints

both in space and scale. To speed up the computations, the second-order derivatives of Hessian are approximated using box filters. SURF scale space is constructed by up-scaling the filter size instead of iteratively reducing the image size. The localization of interest points is identical to the DoG approach. When comparing FastHessian to DoG, HarLap and HesLap versus change of viewpoint, scale and decreasing light, FastHessian demonstrates similar or better detection repeatability in standard over-land scenes¹ [16]. SURF is also known to be significantly faster than other techniques. For instance, Murillo et al. [26], give preference to SURF over SIFT for in-door robot localization, as it often gives better localization accuracy while being faster.

The edge-based detector *SedgeLap* was proposed by Mikolajczyk et al. [17] in 2003. *SedgeLap* uses a multi-scale Canny edge detector and a Laplacian to find and verify features in space and scale.

III. OPTICAL PHENOMENA AFFECTING UNDERWATER IMAGES

Attenuation of the light in water is a process by which light energy is “lost” from the beam because of absorption and scattering. *Absorption* is the conversion of light energy into a different form (normally heat) and its consequent removal from the image-forming process. *Backscatter* is the portion of source illumination reflected back, while *forward scattering* comes from the small-angle deviations of light interacting with suspended sediments. Forward scattering heavily degrades the sharpness of the imaging system [22]. The amount of light scattering produced by matter suspended in water is measured by *nephelometers*. The value of *suspended sediment concentration* (SSC), which gives a mass of suspended matter per volumetric unit of water does not provide any information about the optical visibility [27]. The actual *visual water quality* characteristic has been traditionally measured using a *Secchi disk*, a white or black-and-white disk that is lowered into the water until the image is judged to disappear from view. However, none of these measures can reliably quantify optical quality of images, which depends on three factors: SSC, intensity of light and scattering volume. *Scattering volume* is the intersection of the camera field of view and the cone of light generated by the lamp. Since scattering is accumulated, the bigger the volume is, the heavier its impact on the visual quality of images [28]. A schematic concept of the scattering volume is depicted in Fig.2(b).

We consider *turbidity* as the scattering and attenuation of light, which causes loss of water clarity and transparency. In this situation the scattering effect cannot be neglected and absorption is also present. We define *turbid images* as images where the visibility of the desired scene is degraded by turbidity.

IV. EXPERIMENTAL SETUP: ACQUISITION OF TURBID IMAGES

The experimental setup comprises an underwater video system containing three cameras, a point-like light source, tank with water and skimmed milk (see Fig.2). The trinocular system is composed of three Rolera-XR cameras acquiring images of 696×520 pixels at 12 bit per pixel. These cameras possess sensitive CCDs with large pixels of $13.7\mu\text{m} \times 13.7\mu\text{m}$, that allow image acquisition under low illumination. The distance between the optical centers of cameras is 30 cm. The distance between each camera and the bottom of the tank is 65 cm. The light source is a halogen lamp of 12 V, placed into the underwater housing and connected to a stable power supply to have constant light intensity during the whole experiment. The light source was fixed in the water, 5 cm above the plane of camera centers, at distance of 20, 25 and 45 cm from cameras respectively. The inside walls of the water tank are mate black to prevent inter-reflections. The scene was set using natural sand, rocks and shells. At the start of the experiments the tank was filled with 400 liters of fresh water.

In our experiment we emulated turbidity using skimmed milk. Skimmed milk molecules cover the diameter range from 10 to 600 nm [29]. Particles about 10 nm in size scatter equal amount of light forward and backward. The forward scattering starts dominating for particles about 100 nm and, close to 1000 nm, there is strong small-angle forward scattering and weak backscattering [30]. Hence, using skimmed milk as turbidifier allowed us to have all types of scattering.

Image acquisition was performed using only the artificial illumination (Fig.2). First, 50 images were recorded in clear water. Then the turbidity was gradually increased by adding fixed amounts of milk, starting at 100 ml with increments of 50 ml. This procedure was repeated for 10 times (see Table I). For each milk concentration, 50 images were acquired with 5 seconds delay between shots, to avoid noise correlation. The temporal variance of pixel intensities is therefore small, with a maximum value of 180 for raw 12-bit images and 15 once converted to 8-bit. Conversion of images was necessary since all the detectors work only with 8-bit data. Apart from conversion, histogram stretching was performed to enhance contrast of the final images (see Fig.3).

V. IMAGE QUALITY ASSESSMENT

The first problem when working with turbid images is the absence of a standard technique to evaluate their quality. The quality of the turbid image is understood as the visibility of the target scene and we denote it as *visual clarity*. To quantify loss of visual clarity we tried standard indices: mean squared error (MSE) and peak signal-to-noise ratio (PSNR). MSE was calculated for 8-bit images. PSNR, which was originally designed for wide dynamic range, was calculated for the raw 12-bit data. PSNR estimates the power of corrupting noise that affects the accuracy of original signal and is usually expressed in terms of the logarithmic decibel scale. All the calculated indices were averaged through 50 images for each turbidity level. Typical values for PSNR in lossy image and

¹Dataset available at: <http://www.robots.ox.ac.uk/vgg/data/data-aff.html>




Fig. 2. Experimental Setup. From left to right: trinocular underwater camera system, definition of the scattering volume, light scattering in turbid water.




Fig. 3. Sequence of turbid images acquired by trinocular system. From left to right: camera 0, 1 and 2. Up-down: clear water, 200 ml and 400 ml of milk.

video compression are between 30 and 50 dB, where higher values indicate better quality. The results displayed in Table 1 show that both indices fail to provide adequate measures of visual clarity for underwater turbid images. PSNR values are around 40 dB for all the turbidity levels, while MSE is surprisingly small. MSE and PSNR are easy to calculate, but have no useful application in turbid images.

$$MSE(I_{turbid}) = \frac{1}{mn} \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} ||I_{clear}(i, j) - I_{turbid}(i, j)||^2 \quad (1)$$

$$PSNR = 10 \log_{10} \frac{4095^2}{MSE_{12bits}} \quad (2)$$

To quantify the visual clarity of turbid images we propose to use the Structural SIMilarity Index (SSIM) [21]. SSIM measures similarity of structure, contrast and luminance between reference and distorted images. The structural information is considered to be independent of scene luminance and contrast and permanent, while contrast can vary. For our images, the visual structure of the scene is first blurred and then lost completely with increase in milk concentration. The contrast of the scene decreases due to absorption, and finally stabilizes for the last solutions, where the only perceptible object is the white suspended matter of the turbidifier. The luminance of the scene is kept constant due to the artificial illumination of constant intensity. However, SSIM estimates luminance from the mean of the signal. With increase of milk concentration, the amount of white matter in the water increases, consequently increasing the mean. This produces a slight inaccuracy of the SSIM measure, which will be discussed later. In general, however, SSIM decreases in full accordance with the loss of visual image clarity and quantifies it quite reliably. The SSIM index ranges from 1 in clear water to 0.3 for the last solution images (see Table I). Taking 100 as a measure of full similarity, we introduce a new index, called *Structural Degradation Index* (SDI) defined from the SSIM as

$$SDI = 100(1 - SSIM) \quad (3)$$

This index provides a decreasing integer scale for image degradation, which is more intuitive and easier to interpret. For this reason, we use SDI instead of SSIM in the abscissa axis of our comparative plots.

VI. RESULTS

Keypoints coming from all the detectors were spread using *non-maxima suppression*. In images of unstructured natural scenes and especially underwater images, features tend to agglomerate on the most salient (in shallow water) or the most illuminated (in deep-ocean surveys) objects. Non-maxima suppression takes the keypoint with the highest value of saliency and eliminates all other keypoints around it which are closer than some predefined radius (*non-maxima suppression radius*). The procedure is repeated until all the originally detected

points are searched through or a required number of features is selected.

A. Accuracy Criterion

The detectors were compared by evaluating how many accurate points are detected by each technique in a turbid image. To be considered as “accurate”, the turbid point is allowed to move not more than some predefined shift error ϵ from the reference keypoint, detected in the clear image. All the detectors were used with the default parameters set by the authors, apart from the saliency threshold. To promote a fair comparison we set the saliency threshold to 0 for all the methods, then we apply the non-maxima suppression with the same radius R and finally we take fixed maximum number N of the first best keypoints. For each camera, the reference “clear image” was produced by averaging 50 images acquired in clear water. For each turbidity level, the accuracy of the detector was calculated on 50 sample images, as defined by Eq. 4, and then averaged. N_{turbid}^ϵ denotes the number of ϵ -accurate keypoints detected in the turbid image.

$$Accuracy(\epsilon, R, N) = \frac{N_{turbid}^\epsilon}{N} \quad (4)$$

B. Comparison of Detectors

First, we display the results acquired with our default values of $\epsilon = 0.5$ pixel, $R = 5$ pixel and $N = 1000$ keypoints. A value of 0.5 pixel is a demanding shift error for images of size 696×520 , leading to only the very accurate points being preserved. Non-maxima suppression radius of 5 pixels spreads enough even keypoints from the single-scale detectors, while 1000 is half the average total number of keypoints detected. Accuracy of all the detectors is plotted versus visual image quality expressed by the SDI index.

Fig.4 shows that detectors are clearly divided into two groups for all the cameras. Harris, Hessian, Laplacian, HesLap, DoG, FastHessian and HesAff perform acceptably, while the remaining show poor results. With the decrease in the scattering volume, the slope of all the curves becomes smaller, and the portion of accurate keypoints detected for the last levels of turbidity is significantly larger. In our experimental setup camera 0 has the biggest scattering volume followed by camera 1, and camera 2 has half scattering volume than camera 0.

Among the single-scale detectors, Hessian demonstrates the best performance. It significantly outperforms Harris, because corners are less common for underwater scenes than blobs. Even human-made objects, being for a long time in the water, loose sharp edges and get smoother outlines, which are better recognized by blob detectors. Blobs are more stable to clutter introduced by turbidity than corners, hence the difference in accuracy between Hessian and Harris increases for higher turbidity levels. Hessian slightly outperforms Laplacian as well, since the former penalizes points on edges, which are also sensitive to noise. The single-scale detectors are easy to implement, they are fast and, when supplemented with sub-pixel accuracy, also highly accurate. However, they are only applicable when the change in scale between images is very




Fig. 4. Accuracy of detectors versus turbidity.

TABLE I

STATISTICS FOR THE ACQUIRED TURBID IMAGES. MSE: MEAN SQUARED ERROR FOR 8-BIT IMAGES. PSNR: PEAK SIGNAL-TO-NOISE RATIO FOR 12-BIT IMAGES. SSIM: STRUCTURAL SIMILARITY INDEX FOR 8-BIT IMAGES. SDI: STRUCTURAL DEGRADATION INDEX FOR 8-BIT IMAGES.

Solution	clear	1	2	3	4	5	6	7	8	9	10	11
Milk (ml)	0	100	150	200	250	300	350	400	450	500	550	600
MSE	cam0	36	7	3	3	2	2	2	3	3	3	4
	cam1	50	52	48	8	8	5	4	3	2	3	3
	cam2	21	8	4	3	2	3	4	5	4	4	5
PSNR	cam0	48	37	38	39	40	42	43	44	45	46	42
	cam1	48	33	33	33	35	36	37	43	44	45	41
	cam2	48	29	28	28	28	49	47	46	46	43	44
SSIM	cam0	0.99	0.87	0.78	0.67	0.59	0.51	0.46	0.43	0.4	0.38	0.37
	cam1	0.99	0.85	0.76	0.68	0.62	0.56	0.49	0.44	0.4	0.37	0.35
	cam2	0.1	0.93	0.88	0.79	0.7	0.64	0.57	0.53	0.48	0.44	0.38
SDI	cam0	1	13	22	33	41	49	54	57	60	62	65
	cam1	1	15	24	32	38	44	51	56	60	63	66
	cam2	0	7	12	21	30	36	43	47	52	56	62

small. In this case, if two detectors can be used, Hessian and Harris are complementary. If only one detector should be selected, Hessian is the best choice for underwater images.

Among the scale-invariant detectors FastHessian clearly outperforms the others. The difference increases with turbidity until it is reduced at the last solutions. FastHessian uses box-filters to approximate Hessian for detection in scale and space. This approach appears to be more appropriate against turbidity than classic LoG (HesLap) or its approximation (DoG). Altogether, detectors based on approximations perform additional smoothing of the noise and clutter introduced by scattering, hence being more stable against them. The Hessian-based techniques stand out clearly in comparison with others. HarLap performs poorly, while the instability of corners with increase of turbidity gets worse, due to the non-maxima suppression in scale. Suppression in scale is inherent to all the scale-invariant detectors, it is implemented by original authors and performed during detection. Corners detected at upper scale-space levels, i.e. in highly smoothed images, are projected to the original image, where their location has a much larger uncertainty in terms of the original image resolution. This is the reason why single-scale detectors outperform their scale-invariant counterparts (Hessian and FastHessian, Laplacian and DoG, Harris and HarLap). Fig.4 shows, that HesAff and HarAff ideally coincide with HesLap and HarLap. Hence, we conclude that affine normalization does not introduce inaccuracy in central point localization. The worst performing detector is the edge-based SedgeLap. Edges are blurred, occluded and heavily shifted because of scattering, hence keypoints detected on their basis are highly unstable and inaccurate.

On the whole, the best choice to detect keypoints in underwater images is FastHessian when dealing with change in scale. DoG fires on the same geometrical structures (blobs), so their joining will not improve the results. HarLap detects corners but gives many inaccurate points.

C. Influence of the selected Parameters

In this section we selected the four best detectors (Harris, Hessian, DoG and FastHessian), to study their performance under different values of shift error, non-maxima suppression radius and maximum number of keypoints. Fig.5 (a) shows that regardless of whether the best 500 or 2000 keypoints are selected, the relative order of the detectors remains the same. Accuracy clearly decreases with increase of N only for DoG. It is evident, that with increase of ϵ , accuracy increases. However, Fig.5 (b) additionally verifies that the relative order of detectors remains unchanged. At the same time the influence of the non-maxima suppression radius is more complicated (see Fig.5 (c)). With increase of radius R , for bigger scattering volume (camera 0), accuracy of single-scale detectors increases, while it decreases for DoG and FastHessian. If we have a smaller scattering volume, the accuracy of Hessian, DoG and FastHessian decreases, while Harris does not change its performance.

The influence of the scattering volume is well pronounced in all the previously displayed plots. For bigger scattering volumes SDI is stretched in the beginning and compressed at the end of the scale, while for smaller scattering volumes it is more evenly distributed. This is a matter of discussion, since if SDI would express all the impact of the turbidity on image quality (scattering volume included), the accuracy values for the same SDI should be the same. Consequently, we consider the development and tuning of Structural Degradation Index a subject of future work.

VII. CONCLUSIONS AND FUTURE WORK

As result of the performed experiments, a reference set of turbid underwater images was acquired. The recorded data can be used to test processing, detection, description and matching techniques for applications dealing with underwater imagery.

The most popular recent keypoint detectors were compared in terms of their accuracy when detecting salient points in




Fig. 5. Accuracy of detectors versus turbidity for different values of N , ϵ and R .

turbid images. A combination of classic Hessian and Harris (or Hessian only) supplemented by sub-pixel accuracy was recognized as the best non-invariant-to-scale detector. Since classic methods do not incorporate image scale space, they perform fast and they are easy to implement. Among the scale-invariant detectors, FastHessian, the detector of SURF, demonstrated the best performance in all tests. It outperformed DoG, Hessian Laplace and Harris Laplace. SURF is also known to be very fast.

The final result of our work is the answer to the problem of quality assessment of turbid images. We have demonstrated that the use of the SSIM index is a reliable measure to quantify the degradation of images due to scattering and attenuation. The slight inaccuracy of the index, that became apparent during comparisons, highlighted at the same time a possible method to tune it.

The main direction of the future work consists in the evaluation of recent description and matching techniques on turbid underwater images, taking into account parallax and non-planarity of natural underwater scenes.

ACKNOWLEDGEMENTS

This work has been partially funded through the Spanish Ministry of Science and Innovation under grant CTM2010-15216 and the EU under grant FP7-ICT-2009-248497. Nuno Gracias is funded by the Spanish MCINN under the Ramon y Cajal program.

REFERENCES

- [1] C. Harris and M. Stephens. A combined corner and edge detector. In *Proceedings of Alvey Conference*, pages 189–192, Manchester, UK, 1988.
- [2] T. Lindeberg. *Discrete Scale-Space Theory and the Scale-Space Primal Sketch*. PhD thesis, Royal Institute of Technology, CVAP, Stockholm, Sweden, 1991.
- [3] J. Matas, O. Chum, M. Urban, and T. Pajdla. Robust Wide Baseline Stereo from Maximally Stable Extremal Regions. In *Proceedings of the 13th BMVC*, pages 384–393, 2002.
- [4] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. Van Gool. A comparison of affine region detectors. *International Journal of Computer Vision*, 65:43–72, 2005.
- [5] P. Moreels and P. Perona. Evaluation of Features Detectors and Descriptors Based on 3d Objects. *IJCV*, 73:263–284, 2007.
- [6] P. Corke, C. Detwiler, M. Dunbabin, M. Hamilton, D. Rus, and I. Vasilescu. Experiments with Underwater Robot Localization and Tracking. In *Proceedings of International Conference on Robotics and Automation*, pages 4556–4561, 2007.
- [7] I. Mahon and S. Williams. SLAM using Natural Features in an Underwater Environment. In *Proceedings of the 8th International Conference on Control, Automation, Robotics and Vision*, pages 2076–2081, Kunming, China, 2004.
- [8] S.B. Williams and I. Mahon. Design of an Unmanned Underwater Vehicle for Reef Surveying. In *3rd IFAC Symposium on Mechatronic Systems*, Australia, 2004.
- [9] O. Pizarro, R. Eustice, and H. Singh. Large area 3d reconstructions from underwater surveys. *OCEANS 2004*, 2:678–687, 2004.
- [10] C. Schmid, R. Mohr, and C. Bauckhage. Evaluation of Interest Point Detectors. *IJCV*, 37:151–172, 2000.
- [11] O. Pizarro and H. Singh. Towards large-area underwater mosaicing for scientific applications. *IEEE Journal of Oceanic Engineering*, 28:651–672, 2003.
- [12] M. Dunbabin, K. Usher, and P. Corke. Visual motion estimation for an autonomous underwater reef monitoring robot. In *Proceedings of the International Conference on Field and Service Robotics*, pages 57–68, Port Douglas, Australia, 2005.
- [13] L. Liu, F. Peng, Y. Tian, Y. Xu, and K. Zhao. Fast image matching for localization in deep-sea based on the simplified SIFT (scale invariant feature transform) algorithm. In *SPIE*, volume 6795, 2007.
- [14] V. Brandou, A.G. Allais, M. Perrier, E. Malis, P. Rives, J. Sarrazin, and P.M. Sarradin. 3D Reconstruction of Natural Underwater Scenes Using the Stereovision System IRIS. *OCEANS - Europe*, pages 1–6, 2007.
- [15] P. Beaudet. Rotationally invariant image operators. In *Proceedings of the 4th ICPR: International Joint Conference on Pattern Recognition*, pages 579–583, 1978.
- [16] H. Bay, T. Tuytelaars, and L.J.V. Gool. SURF: Speeded Up Robust Features. In *Proceedings of the ECCV*, pages 404–417, Australia, 2006.
- [17] K. Mikolajczyk, A. Zisserman, and C. Schmid. Shape recognition with edge-based features. In *Proceedings of the BMVC*, volume 2, pages 779–788, 2003.
- [18] S.K. Nayar and S.G. Narasimhan. Vision in bad weather. In *Proceedings of ICCV*, page 820, Washington, DC, USA, 1999. IEEE Computer Society.
- [19] Y. Schechner, S. Narasimhan, and S. Nayar. Polarization-Based Vision through Haze. *Applied Optics, Special issue*, 42:511–525, 2003.
- [20] K. Zuiderveld. Contrast Limited Adaptive Histogram Equalization. *Pearson Education, Inc.*, 1994.
- [21] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13:600–612, April 2004.
- [22] C.J. Funk, S.B. Bryant, and P.J. Heckman, Jr. Handbook of Underwater Imaging System Design. *Ocean Technology Department*, 1972.
- [23] K. Mikolajczyk and C. Schmid. Scale And Affine Invariant Interest Point Detectors. In *IJCV*, volume 1, pages 63–86, 2004.
- [24] A. Ramisa, S. Vasudevan, D. Scharamuzza, R.L. de Mántaras, and R. Siegwart. A tale of two object recognition methods for mobile robots. In *Proceedings of ICVS*, Greece, 2008.
- [25] D. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *IJCV*, 60:91–110, 2004.
- [26] A.C. Murillo, J.J. Guerrero, and C. Sagüés. SURF features for efficient robot localization with omnidirectional images. In *Proceedings of ICRA*, pages 3901–3907, Italy, 2007.
- [27] D. Pavanelli and A. Bigi. Indirect Methods to Estimate Suspended Sediment Concentration: Reliability and Relationship of Turbidity and Settleable Solids. *Biosystems Engineering*, 90:75–83, 2005.
- [28] T. Treibitz and Y.Y. Schechner. Instant 3Descatter. In *CVPR(2)*, pages 1861–1868. IEEE Computer Society, 2006.
- [29] P. Udabage, R. Sharma, D. Murphy, I. McKinnon, and R. Beckett. Size distribution of reconstituted skim milk using field-flow fractionation. *Journal of Microcolumn Separations*, 9:557–563, 1998.
- [30] National Aeronautics and Space Administration. Scattering: Seeing the Microscopic Among the Giants. *Educational Brief*, 2005.