Achieving Turbidity Robustness on Underwater Images Local Feature Detection

Felipe Codevilla felipe.codevilla@furg.br Joel De O. Gaya joelfelipe@furg.br Nelson Duarte Filho dmtnldf@furg.br Silvia S. C. Botelho silviacb@furg.br Center of Computational Sciences (C3) Federal University of Rio Grande (FURG) Rio Grande, Brazil

Abstract

Methods to detect local features have been made to be invariant to many transformations. So far, the vast majority of feature detectors consider robustness just to over-land effects. However, when capturing pictures in underwater environments, there are media specific properties that can degrade the visual quality the captured images. Little work has been made in order to study the robustness that the popular feature detectors have to underwater environment image conditions. We develop a new dataset, called TURBID, where we produced real seabed images with different amounts of degradation. On this dataset, we search over multiple feature detectors from the literature to indicate the ones with more robust properties.

We concluded that scale-invariant detectors are more robust to degradation of underwater images. Finally, we elected *Center Surround Extremas*, *KAZE*, *Difference of Gaussians* and the *Hessian-Laplace* as the best detectors for this environment on all tested scenes.

1 Introduction

The understanding and development of techniques to find local features has been largely improved for over-land images in the past decades. Local features are important due to their robustness and generality, *i.e.*, they can be used in many applications [$\[mathbb{23}\]$]. In the early stages, the objective was to find local features on regions that are invariant to rotation and translation, *e.g.* the *Harris* corner detector [$\[mathbb{12}\]$] or the Hessian blob detector [$\[mathbb{6}\]$]. Invariance to scale was also incorporated to feature points detectors after the development of the *scale space theory* for computer vision [$\[mathbb{13}\]$]. Many popular detectors like SIFT (*DoG*) [$\[mathbb{16}\]$], or SURF (*Fast-Hessian*) [$\[mathbb{12}\]$], have been developed with this approach. Several studies evaluated feature detectors, with respect to their invariance to general transformations [$\[mathbb{12}\]$] [$\[mathbb{2}\]$] [$\[mathbb{2}\]$] [$\[mathbb{2}\]$] [$\[mathbb{2}\]$]. Nevertheless, there is a clear lack of studies concerning robustness towards specific noise conditions.

We are here interested in the specific conditions that exist in the underwater environment. Its photometric properties demand a special treatment in order to overcome the image degradation caused by underwater media properties, here called turbidity for simplicity. The image acquired by visual sensors is highly degraded by the light scattering process on underwater images. The degradation is especially true for higher distances, where the scene could be considered as planar. One could use restoration methods [23] or even try to directly recover the color or contrast properties of the captured image. However, finding local features on an image with an enhanced quality, does not implicate into finding features with higher robustness to the underwater scattering-based degradation.

Today, with the advent of ROVs (Remotely Operated Vehicles) and AUVs (Autonomous Underwater Vehicles) several images are collected from the underwater environment favouring new computer vision applications. With this, applications benefited by finding descriptive feature points in underwater environments grows every year. They are essential for many applications like 3D reconstruction $[\square]$, visual odometry $[\square]$ and ocean floor mapping $[\square]$.

However, most of these applications rely on the best over-land feature detectors, without considering the water photometric properties. It is likely that some algorithms have a better behaviour than others when applied on images degraded by specific underwater conditions. Hence, it is fundamental that feature detectors are tested under different turbidity levels. This is important since, on common ocean survey missions, images are captured under different levels of degradation (Figure 1).



Figure 1: Different turbidity levels of a seabed scene captured on a AUV survey mission [23].

In this context, to evaluate feature detectors, we propose a new dataset called *TURBID*. This dataset is based on real underwater scenes photographs. The pictures are placed on the bottom of a tank filled with a milk-water solution and then are re-photographed with the degradation controlled by the amount of milk. This dataset is an improvement in terms of visual diversity when compared to previous efforts [11] and is one of the main contributions of this work.

On this dataset, we test feature detectors, considering different detection approaches, with respect to their robustness to the degradation caused by turbidity. The objective is to find which are the detectors that are more robust to the general structural degradation of underwater environments, and which are the properties that these detectors hold to sustain this robustness. Finally, we also will present the applications that could be benefited by this study.

2 Underwater Noise Properties

When light propagates in an underwater environment, it interacts with the suspended particles in the medium being both *scattered* and *absorbed*. With this, the amount of light that is scattered or absorbed instead of transmitted on a straight line is defined as *turbidity*.

The *turbidity* culminate into certain results when imaging an underwater scene. Besides the shortening of information imaged, a corruption of information also happens. *Backscattering* is when reflected sources from outside the captured scene are scattered in a wide angle eventually reaching the image plane. This effect creates a characteristic veil on the image that reduces contrast and suppress fine structures on the image. *Forward-scattering* is when light is spread on a short angle, blurring the information from the image scene and creating a contrast reduction.

Schechner and Kapel [22] showed that *most* of the degradation on underwater images happens due to the *backscattering* effect. Thus, for this work, we discarded the *forward scattering* component. We also assumed that the illumination is homogeneous. With this, a common simplification of the underwater propagation model is presented as:

$$T_i(x) = J(x)e^{-c_i d(x)} + A(x)(1 - e^{-c_i d(x)}),$$
(1)

where $T_i(x)$ is an image with a certain *i* amount of turbidity, in a image position *x*. J(x) is the clear image, $e^{-c_i d(x)}$ is the *turbidity* portion that determines the amount of degradation. A(x) is a function that represents the contribution from the *backscattering* effect. The c_i is called the media attenuation constant and is proportional to the amount of floating particles in the media.

Given an image $T_i(x)$, it is possible to measure the amount of the degradation the scene J(x) has suffered. This is proportional to the amount of turbidity from the environment. The turbidity is dependent of the quantity and the type of floating sediment that scatters the incident light. However, to estimate degradation in a image we must consider that the illuminated volume and the intensity of light are also applicable [1].

Image signals are usually highly structured, their pixels show a dependency that carries most of the information from the scene [\square]. When the image signal J(x) is attenuated and the *backscattering* effect dominates, the original image (J(x)) structure is naturally lost. Under these considerations, *Garcia and Gracias* [\square] proposed to use a variation of the *Structural Similarity Index* (*SSIM*) [\square] to evaluate image degradation due to turbidity called *Structural Degradation Index* (*SDI*). The *SSIM* measures the structural similarity between image. The *SDI* is a simple variation to SSIM computed as:

$$SDI = 100(1 - SSIM). \tag{2}$$

The *SDI* performed fairly well to measure image quality. However, to help the comparison between different scenes we used a normalized version of the *SDI* (*NSDI*).

Let an image T_N as being one image completely formed by the *backscattering*, A(x), of a given scene. We measure the *NSDI*_i of an arbitrary turbid image T_i as a ratio between the degradation from the pure backscattering *SDI*_N and the degradation *SDI*_i from an arbitrary image, hence:

$$NSDI_i = SDI_i / SDI_N. \tag{3}$$

3 Achieving Turbidity Robustness in Local Feature Detection

A local feature is usually defined as a small region that differs from its immediate neighbourhood [23]. A local feature detector normally search for invariant patterns on images. These patterns can be invariant to rotation, translation, scale, and be robust to many noise sources. A scale invariant pattern is an special class, where it is detected based on simulating a 3D representation $L(x, \sigma)$, of the image J(x) with a scale parameter σ . Based on the diffusive properties of scale, the scale can be simulated by the application of a smoothing function on the original image:

$$L(x,\sigma) = J(x) * d(x,\sigma), \tag{4}$$

where $d(x, y, \sigma)$ is a function to that is a solution to the diffusion equation that models the scale space [13]. The maxima of a gradient function, such as Laplacian, at the $L(x, \sigma)$ space, is used to find the good scale invariant points. After, a blob or corner detector is used to improve the detection quality.

The local extrema detected by the maxima of gradient functions usually have more invariance to structural deformation due to scale [I]. As stated before, underwater phenomena also create structural degradation. Further, it tends to eliminate all the *finer scale structures*, which is equivalent to the scale phenomena. This can be observed on Figure 2 where we show a comparison between scale generate by a Gaussian blur and a certain level of turbidity with a similar computed SDI.



Consequently, we can assume the invariant points detected by some scale invariant detector can have also a good robustness to turbidity. For that, we propose to analyze the robustness of different scale-space generations functions $d(x, \sigma)$ and maxima detection approaches.

4 The TURBID Dataset

Multiple datasets exist to isolate some special properties that are interesting to evaluate algorithms [1]. For the underwater properties, it is quite a challenge to reproduce in a controlled manner since it is hard to obtain untouched seabed structures.

Feature detectors are highly dependent on the image texture, and shape properties. In order to isolate the turbidity properties, here we propose an dataset that have the real underwater scenes, with all different visual aspects, but having a controlled produced turbidity.



Here we propose the TURBID dataset ¹. We photograph scenes compressing three different high quality printed pictures of real underwater seabed scenes previously photographed at the Bahamas. The printed scenes were selected as the ones with lesser prior degradation possible (Fig. 4 first column). The scenes contains the actual structures of the underwater floor, plus some human made structures, with a natural drawback that is the change of resolution due to printing and re-photographing.²

The pictures were photographed inside a 1000 litres tank, uniformly illuminated with two 30 watts fluorescent light strips (Figure 3) and completely filled with fresh water. As an image capture device, we used a static Go Pro Hero3 Black Edition with a resolution of 3000x4000 pixels. The fluorescent light strips were placed outside the water in order to simulate natural lightning. We capture multiple images where the only changing properties between scenes is the image degradation due to turbidity. We did it in order to minimize any other possible phenomena on the scene.



Figure 3: The experimental setup created in order to evaluate feature detection algorithms. It is composed by a camera, fluorescent lights and underwater seabed printed photos.

We first photograph a image, T_0 , on a totally clean fresh water. After that, the amount of degradation is controlled by successively adding whole milk in water in the tank. That process was repeated 19 times producing the images $T_1...T_{19}$. We choose to use whole milk since it has a higher size particle that induces a lot of wide angle scattering, increasing the *backscattering* effect [20], the main source of underwater degradation [22].

The amount of milk added is also successively increased as show on Table 1. Note that for the last image we added 60 ml of milk to obtain a image with approximately just the backscattering effect.

On the Figure 4, we show the images obtained by the experiment. Each row represents a different printed picture with different levels of turbidity on each column.

This dataset differs from its previous [\square] by the following advantages: *i*: it uses real and varied image visual aspect; *ii*: more turbidity levels, containing a level where the original aspects are completely attenuated (T_{19}) and; *iii*: higher image resolution, allowing a more precise accuracy analysis.

Apparently, the considered scenes are not extensive to evaluate general 3D scenes. However, we endorse that the produced illuminated volume is 3D, and yet there are no particular differences in depths, the formed turbid phenomena is real. Also, the visual structures previously photographed are maintained.

¹ The link to download the dataset can be found at

https://mega.co.nz/#!h05FQJpI!8u5BcyeewSn6mVfc6vBwXjvaTZtzelnDFLnHeKk3hDE

 $^{^{2}}$ There was mainly a change from 20 *pixels/mm*² to 4 *pixels/mm*² of resolution and some additional noise from printing issues.

Image (T_i)	Amount of Whole Milk	Milk Added
T1	5 ml	5 ml
T2	10 ml	5 ml
T3	15 ml	5 ml
T4	20 ml	5 ml
T5	25 ml	5 ml
T6	30 ml	5 ml
T7	36 ml	6 ml
T8	42 ml	6 ml
T9	50 ml	8 ml
T10	58 ml	8 ml
T11	66 ml	8 ml
T12	74 ml	8 ml
T13	82 ml	8 ml
T14	90 ml	8 ml
T15	100 ml	10 ml
T16	110 ml	10 ml
T17	120 ml	10 ml
T18	130 ml	10 ml
T19	190 ml	60 ml

Table 1: Amounts of milk added for each obtained image on the experiment.

More important, as previously showed, when capturing scenes underwater, usually the turbidity phenomena is just significant on larger distances. For this case the depths differences on the scene are not meaningful, especially when turbidity gets higher. With this, we can easily approximate the scenes with 2D photos. This can be observed comparing the real and simulated photos from, respectively, Figures 1 and 4.

Finally, the proposed dataset and methodology is not only useful for feature detection evaluation, but also can be used to test many underwater computer vision algorithms, specially for an accurate *ground truth* on image restoration algorithms.

5 Evaluating Local Features

On this section we present how to evaluate the robustness of local features towards water turbidity.

We use the repeatability criteria, which is associated with the possibility to find a feature point in the same spot after an image transformation. We defend that, if a point detected is repeated under different levels of turbidity, it means that this feature detector is able to find good points under image degradation due to turbidity. Also, the repeatability is closely correlated to feature detector quality [22]. Related to this, the *Localization Accuracy* is also accounted to measure how close the points are repeated.

To compute the *Repeatability* criteria, we first compute the feature points for each detector on the clean image T_0 and also for each turbid image ($T_1...T_{19}$). Feature points tend to concentrate on higher saliency regions, especially on underwater images where the illumination is less uniform. Regarding to this issue, we applied a *non-maxima-suppression*



Figure 4: The images captured over different levels of degradation due to turbidity controlled by milk addition. We photographed three different printed pictures, *Photo 1* (first row), *Photo 2* (second row) and *Photo 3* (third row). On the first column we show the clean image (no milk) for each captured photography. The second column represents a *Low Turbidity* degradation range with around 15 ml of milk (T_4). The *Medium Turbidity* level range is show on the third column and contains around 50 ml of milk (T_{10}). Finally the last (fourth) column is the *High Turbidity* level range with around 100 ml of milk (T_{16}).

technique that selects only the best on each local neighbourhood. We determine the parameter of the *non-maxima-suppression* as δ . This parameter determines the minimum distance each feature will have from others. After this, the *N* best features are selected based on each feature detector own criteria.

The number N of detected features should be sufficiently large, such that a considerable number of features are detected on every object. However higher quantity of features usually induces the problem of less robust features being found. To compute the *NSDI* index we used the image T_{19} as the backscatter image for normalization.

Having the features computed, the repeatability criteria is calculated as following: on the clean image (T_0) , we take each feature and evaluate if it is still resistant with the presence of turbidity. For the feature to be considered resistant, on a level *NSDI*_i of degradation, for an image T_i , it has to be detected on the T_i and not move more than a predefined location error factor *e* pixels with respect to T_0 . Subsequently, the number of resistant features found, from T_1 to T_{19} are counted. Considering this procedure, the repeatability is measured by:

$$R_i = \frac{N_i}{N_0},\tag{5}$$

where N_0 is the number of features on the clean image, and N_i is the number of features *repeated* on the image T_i .

The ε parameter is associated with the localization accuracy. The higher is ε , higher is the tolerance on localization accuracy when evaluating the feature detectors. Further, for lower values of ε we evaluate just the feature points that can be very accurately located.



Figure 5: Results for $N = 1500 \varepsilon = 5$. Testing the repeatability with respect to the ranging of the NSDI degradation index. The best results were obtained by DoG [16].



Figure 6: Results for N = 100 and $\varepsilon = 5$, showing the repeatability with respect to the ranging of the NSDI degradation index. For this case the *CenSurStar* [I] and KAZE [I] got the best results

6 Results

6.1 Tested Feature Detectors

As single scale detectors we evaluate, *Harris* $[\square]$, *Hessian* $[\square]$ and *Laplacian* $[\square]$. All the three implementations were done in Matlab $[\square]$.

We analyzed the scale-invariant *Harris-Laplace* [II], *Hessian-Laplace* [II] and the *Difference of Gaussians* (*DoG*) 2004 methods using the *VL FEAT* implementation [II].

Also with scale invariance we evaluate the popular *FastHessian* from SURF $[\square]$ with the MATLAB implementation.

We tested Center Surround Extremas (*CenSurStar*) detector $[\square]$ using a eight sided star shaped polygon. The implementation is from *OpenCV* $[\square]$ adapted to *MATLAB*.

Finally, we analysed the the anisotropic diffusion based *KAZE* detector $[\Box]$ with different kernels(*KAZE G1, KAZE G2, KAZE G3*) and using the original author's implementation. For all methods the *N* best points were selected based on the peak scores returned by the methods.

6.2 Robustness to Turbidity

Figures 5 and 6 shows the obtained results for repeatability for a *High Quantity*, N = 1500, and *High Response*, N = 100, feature situations. The two situations are a trade-off between high dispersion of features (N = 1500) or computing just the most accurate features (N = 100).

On each case we show the plots for the value of repeatability from the three printed photos testing multiple detectors in function of the *NSDI* index. For both cases a $\delta = 30$ was used, this assures that the points are not overly located on salient regions. We set an error to $\varepsilon = 5$.

CODEVILLA ET AL. : ACHIEVING TURBIDITY ROBUSTNESS



Figure 7: Plots showing the average localization error considering a maximum search of $\varepsilon = 15$. There is a clear better result for scale invariant methods, with emphasis to *DoG*.

The curves from Figs. 5 and 6 confirm our previous hypothesis as there is a clear tendency of robustness from the scale invariant methods. The best detectors were in all cases Hessian-Laplace, *DoG*, *KAZE* and *CenSurStar*.

The *Harris-Laplace* obtained poor results, mainly because the know bad compatibility between corners and scale space generation. The Box-Filter approximations of the *FastHessian* introduce noise in the scale-space, making the repeatability of detection decreased, for some features.

CenSurStar obtained good results specially for the most accurate features (Fig. 6), which was also true for *KAZE*. This helps to conclude about the importance to consider non-Gaussian scale-space representations. However, note that *CenSurStar* had bad repeatability for a higher quantity of features (Figure 5).

The results obtained by scale invariant detector are also able to find more stable points. Figure 7 shows the average localization error ε for each level of degradation with a N = 1500 and a maximum error of 15. We can perceive again a clear advantage of the *DoG* and the *KAZE* methods.

The obtained results contrasted with the results obtained by *Garcia and Gracias* [1]. This difference is mainly because our test scenes had more diverse visual aspects. More importantly, the difference happens due to fine parameter tuning explained on the next section.

6.3 Parameter Selection

All the detectors parameters were tuned to optimize their results for TURBID. We did that since it is remarkable the sensibility of the methods over implementation details [[[]], for instance, a change in the initial image smoothing can lead into a high change of repeatability. With this, all the methods had an initial Gaussian blur of $\sigma = 1.6$.

Also, the number of scales searched is also important for scale invariant methods. For the *Harris-Laplace*, *Hessian-Laplace* and *DoG*, a set of scales with 8 octaves and 3 scales per octave was set. For the *FastHessian* a box filter size to from 9 to 383 was used. For *CenSurStar* we ranged the star box filter sizes from 3-283. Each *KAZE* kernel was made with 4 octaves and 4 scale per octave.

6.4 Summary and Applications

Harris [12] *Hessian* [6] and *Laplacian*[23] approaches performed poorer than the scale invariant methods. *Harris* is generally used as very precise detector and is used in underwater tracking applications [1]. However, we show that, on the present scenario, the use of scale is useful also for precision. We present below some underwater vision applications, and the recommended feature detectors.

Camera Calibration and Tracking: These kind of applications demand precision, whereas the feature point must be located on the same image spots. These applications have been largely used on underwater environments $[\Box 2]$ $[\Box]$. For this case, we recommend the use of $DoG[\Box 3]$ or Hessian-Laplace $[\Box 3]$. Both had the best results for localization accuracy (Figure 7).

Robot Localization and Object Classification: Feature detectors have been explored on these applications underwater for loop closing on SLAM [III] [I] or visual vocabulary building [III]. For this case, the general repeatability is more important than a precise localization accuracy. We would not recommend the *FastHessian* algorithm, commonly used in some applications [I]. On the other hand *DoG* [III] or Hessian-Laplace [III] are recommended in general. However, it is useful to also use *CenSurStar* [II] or *KAZE G2* [I] specially when less features are expected to be found.

7 Conclusions

This paper evaluated multiple feature detectors in underwater environments in order to conclude about the characteristics that produce robustness to turbidity.

We proposed a completely new and open dataset (*TURBID*), containing several photographed real seabed printed pictures that had different amounts of turbidity degradation. Besides the use for local feature evaluations, the *TURBID* dataset can be used for multiple applications, with emphasis on serving as a ground truth for underwater restoration applications.

We measured the structural degradation (SDI) of our produced image and evaluated the feature detectors repeatability and localization accuracy towards this degradation.

We found that, finding scale invariant points is a useful way to find structural degradation robust points. We elected KAZE [D], *Center Surround Extremas* [D], *Difference of Gaussians* [D] and *Fast Hessian* [D] as good feature points detectors for underwater environments in all tested scenes. This is important, specially since detectors without scale invariance are normally recommended to be used in situations where scale changing is low. We identified some specific applications and recommended the best feature detectors for each application.

As a future work, we endorse that scale-space generation functions that consider the turbidity based image degradation would obtain even more robust features.

8 Acknowledgements

The authors would like to thank to the Brazilian National Agency of Petroleum, Natural Gas and Biofuels (ANP), to the Funding Authority for Studies and Projects (FINEP) and to Ministry of Science and Technology (MCT) for their financial support through the Human Resources Program of ANP to the Petroleum and Gas Sector - PRH-ANP/MCT. This paper is also a contribution of the Brazilian National Institute of Science and Technology - INCT-Mar COI funded by CNPq Grant Number 610012/2011-8.

References

- Motilal Agrawal, Kurt Konolige, and Morten Rufus Blas. Censure: Center surround extremas for realtime feature detection and matching. In *Computer Vision–ECCV 2008*, pages 102–115. Springer, 2008.
- [2] Pablo Fernández Alcantarilla, Adrien Bartoli, and Andrew J Davison. Kaze features. In *Computer Vision–ECCV 2012*, pages 214–227. Springer, 2012.
- [3] Josep Aulinas, Marc Carreras, Xavier Llado, Joaquim Salvi, Rafael Garcia, Ricard Prados, and Yvan R Petillot. Feature extraction for underwater visual slam. In OCEANS, 2011 IEEE-Spain, pages 1–7. IEEE, 2011.
- [4] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110, 2008.
- [5] C. Beall, B.J. Lawrence, V. Ila, and F. Dellaert. 3d reconstruction of underwater structures. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4418–4423, 2010.
- [6] P. R. Beaudet. Rotationally invariant image operators. In Proceedings of the 4th International Joint Conference on Pattern Recognition, pages 579–583, Kyoto, Japan, November 1978.
- [7] S. S. C. Botelho, Paulo Drews-Jr, M. Figueiredo, C. Da Rocha, and G. L. Oliveira. Appearance-based odometry and mapping with feature descriptors for underwater robots. *Journal of the Brazilian Computer Society*, 15:47 – 54, 09 2009. ISSN 0104-6500.
- [8] G. Bradski. The opency library. Dr. Dobb's Journal of Software Tools, 2000.
- [9] P. Corke, C. Detweiler, M. Dunbabin, M. Hamilton, D. Rus, and I. Vasilescu. Experiments with underwater robot localization and tracking. In *Robotics and Automation*, 2007 *IEEE International Conference on*, pages 4556–4561, 2007. doi: 10.1109/ROBOT.2007.364181.
- [10] P Drews, Silvia Botelho, and Sebastião Gomes. Slam in underwater environment using sift and topologic maps. In *Robotic Symposium*, 2008. LARS'08. IEEE Latin American, pages 91–96. IEEE, 2008.
- [11] Rafael Garcia and Nuno Gracias. Detection of interest points in turbid underwater images. In OCEANS, 2011 IEEE-Spain, pages 1–9. IEEE, 2011.
- [12] Chris Harris and Mike Stephens. A combined corner and edge detector. In Alvey vision conference, volume 15, page 50. Manchester, UK, 1988.
- [13] Tony Lindeberg. Scale-space theory: A basic tool for analyzing structures at different scales. *Journal of applied statistics*, 21:225–270, 1994.
- [14] Tony Lindeberg. Feature detection with automatic scale selection. *International journal of computer vision*, 30(2):79–116, 1998.

- [15] Tony Lindeberg and Jan-Olof Eklundh. On the computation of a scale-space primal sketch. *Journal of Visual Communication and Image Representation*, 2(1):55 – 78, 1991. ISSN 1047-3203. doi: http://dx.doi.org/10.1016/1047-3203(91)90035-E.
- [16] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [17] MATLAB. version 7.10.0 (R2010a). The MathWorks Inc., Natick, Massachusetts, 2010.
- [18] Krystian Mikolajczyk and Cordelia Schmid. Scale & affine invariant interest point detectors. *International journal of computer vision*, 60(1):63–86, 2004.
- [19] Krystian Mikolajczyk, Tinne Tuytelaars, Cordelia Schmid, Andrew Zisserman, Jiri Matas, Frederik Schaffalitzky, Timor Kadir, and Luc Van Gool. A comparison of affine region detectors. *International journal of computer vision*, 65(1-2):43–72, 2005.
- [20] Srinivasa G Narasimhan, Mohit Gupta, Craig Donner, Ravi Ramamoorthi, Shree K Nayar, and Henrik Wann Jensen. Acquiring scattering properties of participating media by dilution. ACM Transactions on Graphics (TOG), 25(3):1003–1012, 2006.
- [21] T. Nicosevici, N. Gracias, S. Negahdaripour, and R. Garcia. Efficient three-dimensional scene modeling and mosaicing. *Journal of Field Robotics*, 26:759–788, 2009.
- [22] O. Pizarro, R. Eustice, and H. Singh. Large area 3d reconstructions from underwater surveys. In OCEANS '04. MTTS/IEEE TECHNO-OCEAN '04, volume 2, pages 678– 687 Vol.2, 2004. doi: 10.1109/OCEANS.2004.1405509.
- [23] O. Pizarro, S.B. Williams, M.V. Jakuba, M. Johnson-Roberson, I. Mahon, M. Bryson, D. Steinberg, A. Friedman, D. Dansereau, N. Nourani-Vatani, D. Bongiorno, M. Bewley, A. Bender, N. Ashan, and B. Douillard. Benthic monitoring with robotic platforms x2014; the experience of australia. In *Underwater Technology Symposium (UT), 2013 IEEE International*, pages 1–10, March 2013. doi: 10.1109/UT.2013.6519909.
- [24] Oscar Pizarro, Paul Rigby, Matthew Johnson-Roberson, Stefan B Williams, and Jamie Colquhoun. Towards image-based marine habitat classification. In OCEANS 2008, pages 1–7. IEEE, 2008.
- [25] Raimondo Schettini and Silvia Corchs. Underwater image processing: state of the art of restoration and image enhancement methods. *EURASIP Journal on Advances in Signal Processing*, 2010:14, 2010.
- [26] Cordelia Schmid, Roger Mohr, and Christian Bauckhage. Evaluation of interest point detectors. *International Journal of computer vision*, 37, 2000.
- [27] Tali Treibitz and Yoav Y. Schechner. Instant 3descatter. In Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition -Volume 2, CVPR '06, pages 1861–1868, 2006. doi: 10.1109/CVPR.2006.155.
- [28] Tinne Tuytelaars and Krystian Mikolajczyk. Local invariant feature detectors: a survey. *Foundations and Trends* (R) *in Computer Graphics and Vision*, 3(3):177–280, 2008.

- [29] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. http://www.vlfeat.org/, 2008.
- [30] Zhou Wang, Alan C Bovik, Hamid Rahim Sheikh, and Eero P Simoncelli. Image quality assessment: From error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13(4):600–612, 2004.