

An analysis of monochrome conversions and normalizations on the Local Binary Patterns texture descriptors

Navid Nourani-Vatani¹, Mark De Deuge¹, Bertrand Douillard², Stefan B. Williams¹

¹Australian Centre for Field Robotics

The University of Sydney

Darlington NSW 2050, Australia

f.lastname@acfr.usyd.edu.au

²Jet Propulsion Laboratory

M/S 198-235

4800 Oak Grove Dr.

Pasadena, CA 91109, USA

Abstract

While the importance of the choice of color space for color descriptors has been studied extensively, a similar study for image texture descriptors is missing. This publication investigates the effect of color-to-monochrome conversions, image normalization, and metrics on the discriminative power of texture descriptors. The measure of the discriminative power of a feature is formulated as supervised spectral feature analysis. This analysis allows to measure the relative performance of a feature under varying conditions as long as the feature dimension is maintained. Feature discrimination evaluation is applied to Local Binary Patterns texture descriptors and it is shown how the proposed metric directly maps to classification performance. Based on this metric, we demonstrate that the choice of color-to-monochrome conversion and normalization can have a significant effect on the performance of the LBP descriptors.

1. Introduction

Given that visual features are central to image classification, metrics to evaluate the discriminative power of these features are essential for building accurate classifiers. This publication investigates the effect of monochrome transformations, image normalizations, and metrics for feature discrimination in the context of texture features. In particular, the focus is on variations of the popular Local Binary Pattern (LBP) texture descriptors [2, 17], although the presented methodology is equally applicable to other features.

Texture descriptors are normally extracted from monochrome images [2, 10, 17, 20]. Various authors have proposed so-called color-texture descriptors by extracting texture information within and between each color channel [5, 6, 19] or by combining texture descriptors with color descriptors [15, 22, 23, 24]. However, color information—

and obviously the image luminosity—is illumination dependent. Because of this, image processing in outdoor environments is particularly difficult due to changing illumination, strong sun, shadows, etc. As such, much research has focused on stabilizing the color information by making it more robust to changes in illumination [4, 8]. Image normalization can also provide some gray-scale invariance and remove some of the global gray-scale properties and is frequently used in the literature [8, 13, 17, 18].

From the color/color-texture literature, it is obvious that it is important to investigate the properties of the color space. We are, however, unaware of a study on the impact of color-to-monochrome conversion on descriptors. Many descriptors, e.g. LBP, SIFT [14] and HOG [7], are after all in their basic forms computed from the luminosity of the images. As we will show in Section 3, each color space has its own definition of luminosity, which are lossy and can involve non-linear conversions of the RGB components.

To analyze the effect of color-to-monochrome conversions and normalization, in this paper we perform supervised spectral feature analysis [26] on the extracted descriptors. Our contributions are threefold. 1) We introduce a novel metric based on spectral feature analysis for estimating the discriminatory power of (texture) descriptors. The advantage of this metric is that it allows a *classifier free* evaluation of a feature's discriminative power. 2) We validate the use of this metric as a measure of feature discrimination. The obtained values of the discrimination measures are shown to map monotonically to actual classification results obtained with Support Vector Machines (SVMs). 3) We employ this metric to compare the discriminative power of features formed from six different monochrome conversion, and in which various image normalization strategies are used. For each conversion and normalization tested, this analysis produces metrics quantifying the separation in feature space of various classes and the invariance of the texture descriptor with respect to illumination.

The remainder of this paper is organized as follow. Next section briefly describes the various LBP descriptors employed in this paper. Section 3 introduces the tested monochrome conversion and normalization methods. In Section 4 we introduce our feature analysis methodology. Spectral feature analysis and classification results are presented in Section 5. Finally, in Section 6 we present our concluding remarks.

2. Texture descriptors

A large number of features have been developed to describe the texture in an image [2, 10, 16, 20]. In various studies [17, 20] these descriptors have been compared and the Local Binary Patterns (LBP) descriptors have been shown to often outperform many of the other descriptors. Because of the calculation simplicity and robustness, these descriptor are often the texture descriptor of choice [1, 9, 15] and thus we focus on the LBP descriptors. However, the approach for quantifying the separation of the texture classes in feature space introduced later can be equally applied to other feature descriptors.

2.1. LBP overview

Ojala et al. [17] introduced the shift and rotation invariant Local Binary Patterns (LBP) texture descriptor. The LBP descriptor is calculated for every pixel by comparing the value of a center pixel with the values of P circularly-symmetric surrounding pixels at a radius R from this center pixel, $LBP_{P,R}$. By taking the sign of the difference between the center and surrounding pixels a binary descriptor is generated, which is also scale invariant.

More recently, Ahonen et al. [2] introduced an LBP descriptor (LBP-HF) that locally preserves the rotational variations of the surrounding pixels (by not performing the aforementioned rotation) while being rotationally invariant at the patch level. This is achieved by representing the LBP descriptor histograms in the Fourier space.

An extension to the rotation-invariant descriptor, $LBP_{P,R}^i$, the rotation-invariant *uniform* descriptor, $LBP_{P,R}^{riu2}$, has also been developed [17] after realizing that a certain few patterns occur in more than 90% of the observed textures. These “fundamental” patterns are more uniform and contain very few spatial transitions ($0 \rightarrow 1$ and $1 \rightarrow 0$ transitions). Here, *u2* corresponds to uniform patterns with less than 2 transitions.

LBPs allow for detecting patterns in circular neighborhoods of any quantization of the angular space and at any spatial resolution by employing interpolation. Multiple LBP operators can also be combined (e.g. concatenated) to generate a multi-resolution descriptor.




Figure 1: The effect of calculating the $LBP_{8,1}^{riu2}$ histogram using different patch sizes N of a uniformly textured image. a) The textured image. b) The variance in the LBP histograms.

2.2. Selected LBPs

In this paper we analyze the performance of the most popular Local Binary Pattern descriptors, the rotation-invariant uniform feature ($LBP_{P,R}^{riu2}$) [17] and the newer Local Binary Pattern Histogram Fourier feature (LBP-HF) [2]¹. In our feature analysis we compare the performance of both $LBP_{8,1}$ and $LBP_{24,3}$. As such, we are comparing the effect of color-to-monochrome conversion and image normalization on four descriptors: $LBP_{8,1}^{riu2}$, $LBP_{24,3}^{riu2}$, $LBP\text{-HF}_{8,1}$, and $LBP\text{-HF}_{24,3}$.

Figure 1 shows the effect of varying the patch size when calculating the histogram for $LBP_{8,1}^{riu2}$. Here the descriptor is calculated on a single image with a repeating pattern and the histograms are extracted from non-overlapping square patches at sizes $N = [7, 15, 31, 63, 125]$ ². The boxplots show the variation of each dimension of the descriptor for each patch size. It is observed that the median of the values across the different patch sizes are similar while the variation decreases with the patch size. In the remainder of this document a patch size of $N = 31$ is used, which exhibits an acceptable tradeoff between precision, number of samples and computation time.

3. Color models

In this study, we analyse the effect of various color-to-monochrome transformation on the LBP features extracted from the converted monochrome image. Furthermore, we investigate how the change in the source illumination affects the descriptor. This is of particular interest in outdoor scenes, where both the intensity and chromaticity change continuously due to the change in illumination (e.g. motion of the sun in the sky or wavelength dependent attenuation of light in water).

The next section reviews some of the commonly employed color-to-monochrome conversions from various

¹We employ the C++ implementation freely available from <https://github.com/nourani/LBP>.

²The effect is similar for the other descriptors. $LBP_{8,1}^{riu2}$ was chosen as it has the smallest feature dimension $n = 10$ making the illustration more clear.

color spaces while section 3.2 looks at the image normalization methods.

3.1. Monochrome conversions

The conversion to monochrome is lossy and the choice of how this compression is done can have significant effect on the discriminative power of a texture descriptor. In this study we will demonstrate the effect of extracting Local Binary Pattern descriptors from the luminosity channels of these six color spaces: XYZ, Lab, HSV, HSL, HSI, and YUV₆₀₁.

Poynton [21] recommended to computed the brightness component as a properly-weighted sum of red, green, and blue, e.g.:

$$Y_{601} = 0.299 \cdot R + 0.587 \cdot G + 0.114 \cdot B. \quad (1)$$

Similarly the luminance in the XYZ color space can be calculated from RGB, using D50 white-point, by:

$$Y_{XYZ} = 0.228 \cdot R + 0.737 \cdot G + 0.03 \cdot B. \quad (2)$$

The perceptually uniform lightness value of the Lab color space can be calculated from (2):

$$L_{LAB} = \begin{cases} 116(\frac{1}{3}\frac{29}{6}^2 Y + \frac{4}{29}) - 16, & Y \leq \frac{6}{29} \\ 116Y^{1/3} - 16, & \text{otherwise.} \end{cases} \quad (3)$$

In the formulation for HSV/HSL/HSI color spaces, the brightness is defined as the largest component or an average of the color components:

$$V = \max(R, G, B), \quad (4)$$

$$L = \text{mean}(\min(R, G, B), \max(R, G, B)), \quad (5)$$

$$I = \text{mean}(R, G, B). \quad (6)$$

These representations are highly nonlinear, which also introduces spokes (i.e. discontinuities) into the hue circle [21].

3.2. Image normalization

Image normalization can provide some gray-scale invariance and remove some of the global gray-scale properties. We investigate the effect of image normalization, before descriptor calculation, through *histogram stretching* and *histogram averaging* applied globally on the whole image or locally on the image patches:

Global averaging: Ensure the image has an average of 128 and a standard deviation of 20.

Global stretch: Stretch the image histogram to span the range [0, 255]. This was employed by Kim et al. [13].

Local averaging: Perform averaging but on the 31×31 image patches. This was employed by Ojala et al. [17] but with patch size of 128×128 .

Local stretch: Perform histogram stretching on the 31×31 image patches. This was employed by Osuna et al. [18].

None: No normalization is performed. This is our baseline for the performance of the normalization.

The next section introduces our feature analysis methodology.

4. Feature Analysis

We introduce a metric for evaluating the discriminative power of a descriptor, which is here applied to texture descriptors. This metric is used to evaluate the sensitivity of the descriptor to illumination changes and color space. Rather than evaluating the features on various classifiers, we seek to capture a measure of the expected performance of the feature space without evaluating classifiers and thus performing a feature evaluation which does not factor in the specifics of the classifier.

For classification purposes the features of the same class are desired to be close (even under illumination change), and features of other classes to be well separated. This desired feature space structure can be encoded in an affinity matrix S_{ij} , where the entry (i, j) encodes the connectivity of data points i and j . The discriminative power of a feature can be evaluated by comparing the separation it provides in feature space against the desired separation encoded by S_{ij} . This evaluation can be performed with supervised spectral feature analysis [26].

In our case, we would like to know what choice of color space and normalization results in the best separation of the descriptor features for different textures, while remaining invariant to changes in illumination for the same texture.

4.1. Supervised Spectral Feature Analysis

Mappings to the feature space should in principle keep feature responses of the same class close under different illuminations. This desired intra class proximity can be encoded using the Laplacian score [11]. This score evaluates the r^{th} feature f_r with the following measure:

$$L_r = \frac{\sum_{ij} (f_{r_i} - f_{r_j})^2 S_{ij}}{\text{Var}(f_r)}, \quad (7)$$

where S_{ij} is the encoding of affinity of the data points i and j encoded using feature f_r . For example, f_{r_i} is the patch i encoded using feature f_r . In the supervised setting S_{ij}^{intra} is defined as:

$$S_{ij}^{\text{intra}} = \begin{cases} \frac{1}{n^2}, & \text{if } C_i = C_j \wedge i \neq j \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

where C is the class and n in the number of elements in that class.




Figure 2: Intra-class and inter-class separation. For the sake of clarity only one set of intra-class and inter-class connectivities are shown.

We extend the intra-class Laplacian score, to an inter-class compactness score S_{ij}^{inter} in order to include an assessment of the separability of classes in the feature spaces:

$$S_{ij}^{inter} = \begin{cases} \frac{1}{n \cdot m}, & \text{if } C_i \neq C_j \\ 0, & \text{otherwise,} \end{cases} \quad (9)$$

where n and m are the number of elements in the classes C_i and C_j , respectively. This allows us to assess if features from different classes, which we desire to remain separate, are overlapping. An example of the intra-class similarity and the inter-class distances can be seen in Figure 2. Here solid lines represent the intra-class feature affinity S_{ij}^{intra} that is calculated for a single feature in each class. The dashed lines represent the inter-class distances from a single feature in one class to the features of another class. Thus we use these metrics to assess the agreement between our feature space and color space combination and our desired (separable) feature space.

Our approach is similar to marginal Fisher analysis (MFA) [25]. The goal of MFA is to find a linear projection that preserves the desired intra-class structure, while increasing the inter-class distances [25]. However, our approach uses the intra and inter class Laplacians to assess these characteristics in the given feature space.

It should be noted that the Laplacian score cannot be used to compare the performances of features with different dimensions due to the non-linear relationship between the distributions of L_2 distance and the dimension of the feature space. Hence, the goal of this analysis is not to evaluate which of the LBP descriptors perform best but to realize if the dependencies on monochrome conversions and normalizations is applicable to all the descriptors.

5. Results

5.1. Feature analysis experiment

To evaluate the proposed concept we have carried out an experiment using a standard texture data set. We present the results from our feature analysis and validate it by showing results from a linear SVM classifier.

5.1.1 Data set

The *Outex_TC_00012* image data set³ contains 24 textures at 9 orientations and three illuminations: 2300K horizon sunlight denoted as ‘horizon’, 2856K incandescent CIE A denoted as ‘inca’, and 4000K fluorescent TL84 denoted as ‘TL84’. The rotation-invariance of the LBP descriptors on this data set has been shown before [17, 2] and is not of interest in the context of our analysis. We are interested in the stability of the features under varying illumination. Figure 1 shows an example image from this data set.

5.1.2 Feature analysis

Each image in the data set is converted to the six color spaces and the luminosity channel extracted. Patches of size 31×31 are extracted to be evaluated with each of the LBP descriptors. Local normalization is performed on these patches prior to LBP calculation while global normalization is applied before to the patch extraction. Note that only one form of normalization is carried out, not both global and local. Finally, all features are normalized to have unit norm. This process results in a total of 30 experiments (6 monochrome conversions \times 5 normalization methods) where each experiment has a data set of 32400 patches (450 samples/class \times 24 classes \times 3 illuminations).

Since our ideal feature space increases the margins between different classes and maintains intra-class locality we can define the affinity matrix S_{ij}^{intra} (8) and our desired inter-class dissimilarity matrix S_{ij}^{inter} (9) with $n = m = 450$. In each experiment we evaluate the ability of the feature space to enhance separation and maintain locality with the change in illumination.

Figure 3 shows the spectral analysis results. The results for the six monochrome conversion are shown in six different colors while the results for employing normalization methods are shown with different markers. The results are presented using a normalized score since the Laplacian score is unit-less and only the relative score can be evaluated. Given a luminosity-normalization pair, the further to the left a marker is positioned the better *intra-class* closeness is achieved. The further towards the top, the better *inter-class* separation. Hence, markers towards the top left are performing better than markers towards the bottom right.

We can observe that the descriptors extracted from the luminance channel of the XYZ color space (shown in red) generally demonstrate higher intra-class closeness; i.e. the red markers are furthest towards the left. For the other conversion, the descriptors extracted from the luminosity channels of the Lab, YUV and HSV spaces perform generally better than those from the HSI and HSL color spaces.

³Data sets available from <http://www.outex.oulu.fi>.

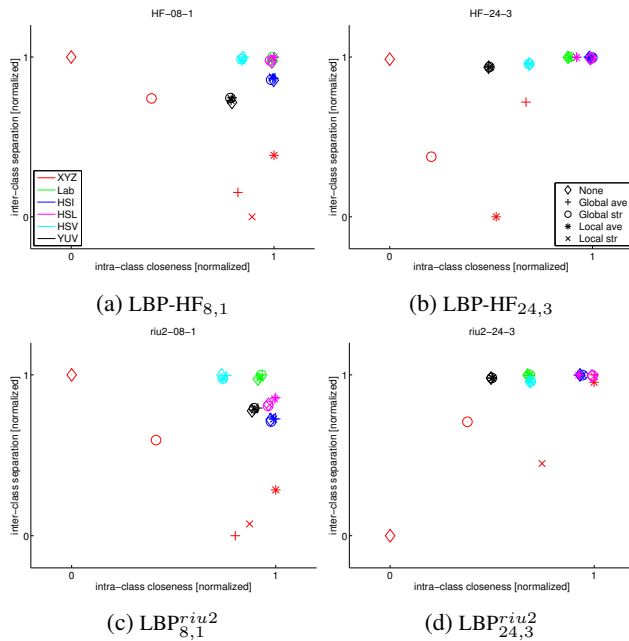


Figure 3: Intra-class closeness and inter-class separation of the features in any color space-normalization pair. Markers towards the top left are performing better than markers towards the bottom right. See text for detailed analysis.

Looking at the normalization methodologies, we observe that the best intra-class closeness is achieved by not performing any normalization (the \diamond in the figure). We see no obvious answer to which normalization methodology results in the best inter-class separation, however, we do observe that the XYZ color space is showing less inter-class separation when applying normalization.

5.1.3 SVM evaluation

To verify the spectral feature analysis results presented above, we have performed classification employing a linear SVM [13]. We performed 3-fold cross validation by training on one illumination, e.g. horizon, and testing on the other two, e.g. inca and TL84⁴. The cost parameter, C , was found by performing an exponential grid search in the range $[10^{-3} - 10^3]$ [12]. The training set consisted of 10800 patches (450 patches/image \times 24 classes) and the test set of 21600 patches (450 patches/image \times 24 classes \times 2 illuminations).

The results of the classifications are shown in Figure 4 and presented in terms of the average F_1 score of the three training/test sets. We observe that the best classification results are achieved in the XYZ, Lab and YUV color space,

⁴We observed that when trained on two illuminations and tested on the third, the SVM classified the test samples with 100% accuracy regardless of monochrome conversion and normalization method. We are therefore presenting the more challenging problem of training on a single illumination and testing on two other.

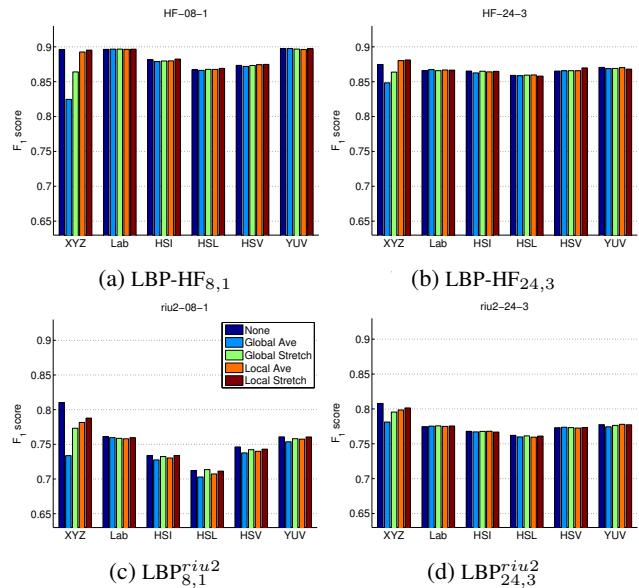


Figure 4: SVM classification results. The results are shown in terms of the mean, best and worst F_1 scores. See text for detailed analysis.

while the descriptors extracted from the luminosity channels of the HSI/HSL/HSV perform worst. This was expected looking at the spectral analysis results in Figure 3. Also here, we do not observe any noticeable performance difference between the normalization methods, except in the XYZ color space.

These findings are significant since they provide an experimental validation of the use of spectral feature analysis as a general tool for measuring separation in feature space. However, it is important to notice that the presented methodology cannot be employed to measure the difference between the different features, i.e. we could not predict the LBP-HF_{8,1} descriptor to be performing better than the LBP_{8,1}^{riu2} or any of the other descriptors. Developing feature analysis methodologies that can detect this is an ongoing research topic and, as supported by these results, of great importance.

5.2. Real world data set validation

To validate the results attained in the previous section we have carried out another experiment where we perform benthic habitat classification using LBP descriptors.

5.2.1 Tasmania data set

The data set consists of 1258 images from 14 different Autonomous Underwater Vehicle (AUV) dives captured of the east coast of Tasmania, Australia. Due to the nature of the underwater images they exhibit large variability in illumination, attenuation and color correctness. The images have been labeled by expert marine scientists, with each image




Figure 5: Example images from the *Tasmania* data set.

containing 50 labels, defining a taxonomical hierarchy of 19 species and substrate forms [3]. Figure 5 illustrates a few example images from the data set⁵.

5.2.2 Results

Multi-resolution LBP features have previously been shown [17, 2] to perform superior to single-resolution descriptors. In this more challenging data set, we therefore employed a multi-resolution LBP-HF descriptor, by concatenating a LBP-HF_{8,1} and a LBP-HF_{8,3} descriptor. Two-thirds of the data set was used for training and one third for validation. Again we employed three-fold cross validation and exponential grid search for the cost parameter C .

The results are shown in Figure 6. Figure 6a shows the spectral analysis result, while Figures 6b-6c show the precision-recall and the F_1 -scores, respectively. From the spectral analysis results we notice that the normalization is having a very large impact on this data set; all the \diamond are in the bottom right. We can see that this is clearly reflected in the precision-recall plot of the classification results in Figure 6b. From the spectral analysis, the global methods are producing higher inter-class separation while the local methods are producing higher intra-class closeness. In the SVM classification results, the global normalization methods produce slightly better results than the local normalization methods. Hence, we can observe that the additional inter-class separation gained by employing global normalization was more significant than the added intra-class closeness in the local normalization methods. For this data set, this difference is possibly due to excessive noise amplification in the small patches. Finally, from the F_1 scores we can see that again that the XYZ, YUV and Lab conversions result in the highest classification results. The best classification is achieved by employing XYZ conversion and global image stretching, which results in an $F_1 = 75.3$. The bottom line is that an increase of almost 40% is attained by choosing a more appropriate monochrome conversion and normalization method.

⁵The data set can be downloaded from <http://marine.acfr.usyd.edu.au/datasets>.

6. Conclusion

This paper has introduced a *classifier-free* metric to analyze the performance of descriptors in terms of their intra-class compactness and inter-class separability under a varying condition. We demonstrated the effect of color-to-monochrome conversions and image normalization methods on the popular Local Binary Pattern texture descriptors. The results attained with the proposed metric were validated extensively using standard classification techniques.

The outcome of our analysis was that the choice of color-to-monochrome conversion does indeed matter; this should not come as a surprise as these conversions are lossy and can be non-linear. The best performance was consistently achieved by employing a linear conversion from RGB: XYZ, YUV, Lab, and HSV. In particular the conversion to XYZ resulted in best performance in all tests.

We also demonstrated that the spectral analysis can capture the effects of image normalization and found that there can be noticeable improvement when the illumination changes across the images are extreme. This is in light that the LBP features are themselves designed to be robust against changes in intensity.

While LBPs were used in this study the proposed method can be equally applied to other descriptors and other test scenarios. In future studies, it would be interesting to investigate the effect of color-to-monochrome conversion on other popular descriptors such as SIFT, HOG etc, and to define metrics to compare descriptors of various dimensions. One limitation of the presented method is that it cannot be used to compare descriptors with different dimensionality. The ability to do so is very important and is an active research topic.

Acknowledgement

This work was supported by the Australian Government through the SIEF program. The authors would like to thank their colleagues at the Australian Centre for Field Robotics, in particular Mr Daniel Steinberg for his valuable comments. The authors would also like to acknowledge the Australian National Research Program (NERP) Marine Biodiversity Hub for the taxonomical labeling.

References

- [1] T. Ahonen, A. Hadid, and M. Pietikainen. Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(12):2037–2041, 2006.
- [2] T. Ahonen, J. Matas, C. He, and M. Pietikainen. Rotation invariant image description with local binary pattern histogram Fourier features. *Image Analysis*, pages 61–70, 2009.
- [3] M. Bewley, B. Douillard, N. Nourani-Vatani, A. Friedman, O. Pizarro, and S. Williams. Automated species detection: An experimental approach to kelp detection from sea-floor

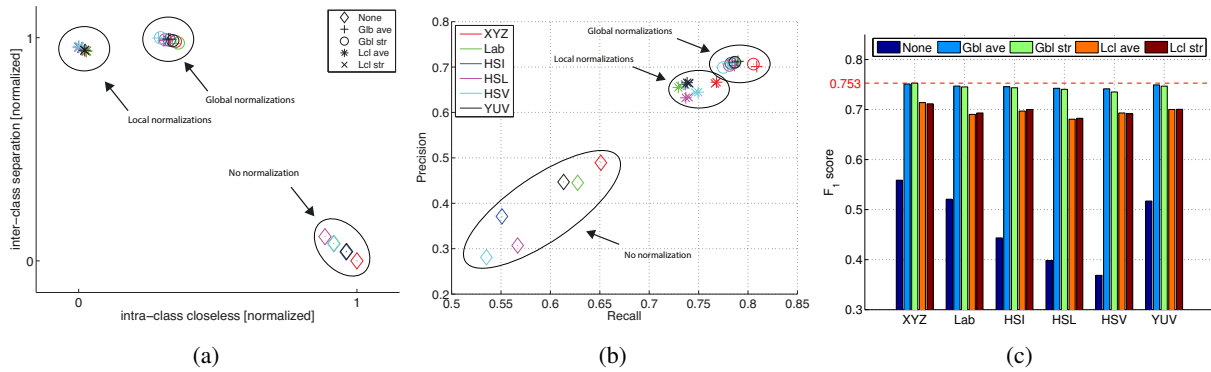


Figure 6: a) Spectral analysis result, b) Precision-recall results and c) F_1 -scores for the *Tasmania* data set.

- auv images. In *Australasian Conference on Robotics and Automation*, 2012.
- [4] M. Bryson, M. Johnson-Roberson, O. Pizarro, and S. Williams. Colour-Consistent Structure-from-Motion Models using Underwater Imagery. In *Robotics: Science and Systems (RSS)*, Sydney, 2012.
- [5] C.-H. Chan, J. Kittler, and K. Messer. Multi-scale local binary pattern histograms for face recognition. *Advances in biometrics*, pages 809–818, 2007.
- [6] C.-H. Chan, J. Kittler, and K. Messer. Multispectral local binary pattern histogram for component-based color face verification. In *IEEE International Conference on Biometrics: Theory, Applications, and Systems*, pages 1–7, 2007.
- [7] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 886–893, 2005.
- [8] G. Finlayson, S. Hordley, G. Schaefer, and G. Yun Tian. Illuminant and device invariant colour using histogram equalisation. *Pattern Recognition*, 38(2):179–190, 2005.
- [9] H. Grabner and H. Bischof. On-line boosting and vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 260–267, 2006.
- [10] R. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Transactions on Systems, Man and Cybernetics*, 3(6):610–621, 1973.
- [11] X. He, D. Cai, and P. Niyogi. Laplacian score for feature selection. *Advances in Neural Information Processing Systems*, 18:507, 2006.
- [12] C. Hsu, C. Chang, C. Lin, et al. A practical guide to support vector classification. <https://www.cs.sfu.ca/people/Faculty/teaching/726/spring11, 2003>.
- [13] K. Kim, K. Jung, S. Park, and H. Kim. Support vector machines for texture classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(11):1542–1550, 2002.
- [14] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004.
- [15] M. Marcos, M. Soriano, and C. Saloma. Classification of coral reef images from underwater video using neural networks. *Optics Express*, 13(22):8766–8771, 2005.
- [16] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on featured distributions. *Pattern recognition*, 29(1):51–59, 1996.
- [17] T. Ojala, M. Pietikäinen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, 2002.
- [18] E. Osuna, R. Freund, and F. Girosit. Training support vector machines: an application to face detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 130–136, 1997.
- [19] M. Pietikäinen, T. Mäenpää, and J. Viertola. Color texture classification with color histograms and local binary patterns. In *Workshop on Texture Analysis in Machine Vision*, 2002.
- [20] R. Porter and N. Canagarajah. Robust rotation-invariant texture classification: wavelet, Gabor filter and GMRF based schemes. In *IEE Proceedings of Vision, Image and Signal Processing*, pages 180–188. IET, 1997.
- [21] C. Poynton. Frequently Asked Questions about Gamma. <http://www.poynton.com>, 1997.
- [22] M. Soriano, S. Marcos, C. Saloma, M. Quibilan, and P. Alino. Image classification of coral reef components from underwater color video. In *OCEANS*, pages 1008–1013, 2001.
- [23] K. E. van de Sande, T. Gevers, and C. G. Snoek. Evaluating color descriptors for object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1582–1596, 2010.
- [24] J. Van De Weijer and C. Schmid. Coloring local feature extraction. *European Conference on Computer Vision*, pages 334–348, 2006.
- [25] S. Yan, D. Xu, B. Zhang, H. Zhang, Q. Yang, and S. Lin. Graph embedding and extensions: A general framework for dimensionality reduction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(1):40–51, 2007.
- [26] Z. Zhao and H. Liu. Spectral feature selection for supervised and unsupervised learning. In *International conference on Machine Learning*, pages 1151–1157. ACM, 2007.