Edge Classification using Photo-Geometric features

Josep M. Gonfaus[1] Theo Gevers [1,2] Arjan Gijsenij [2] F. Xavier Roca[1] Jordi Gonzàlez[1]

[1]*Centre de Visió per Computador*
Universitat Autònoma de Barcelona

[2]*Faculty of Science*
Universiteit van Amsterdam

Abstract

Edges are caused by several imaging cues such as shadow, material and illumination transitions. Classification methods have been proposed which are solely based on photometric information, ignoring geometry to classify the physical nature of edges in images.

In this paper, the aim is to present a novel strategy to handle both photometric and geometric information for edge classification. Photometric information is obtained through the use of quasi-invariants while geometric information is derived from the orientation and contrast of edges. Different combination frameworks are compared with a new principled approach that captures both information into the same descriptor.

From large scale experiments on different datasets, it is shown that, in addition to photometric information, the geometry of edges is an important visual cue to distinguish between different edge types. It is concluded that by combining both cues the performance improves by more than 7% for shadows and highlights.

Figure 1. Shadow-Edge Detection. Photometric information (left) is enhanced when it is combined with geometric information (right).

1. Introduction

Edges are fundamental visual cues which are at the basis of many image understanding and computer vision methods. Edges are caused by a large variety of imaging variables such as shadows, highlights, illumination and material changes. The classification of edges by their physical origin is useful for image understanding, where corresponding edge types (e.g. material edges) are considered for a specific task at hand while discounting other accidental and disturbing edge types (such as shadows and highlight edges). In this paper, we consider the problem of discriminating shadow and specular edge types based on local surface properties.

In general, edge classification is commonly based on photometric information only. For instance, Gevers and Stokman [3] distinguish between shadow-geometry, highlight and material edges by using a rule-based approach. They compute image-derivatives in different (invariant) color spaces, and then assign class labels to edges based on the whether they are present or not in different (invariant) color spaces. Van de Weijer et al. [14] propose a slightly different scheme, called *quasi-invariants*, but is still based on similar invariant principles. Further, Finlayson et al. [1] propose a transformation that results in an image that is invariant to the light intensity and color. This approach distinguished between shadow and material edges. Shadow edges are then detected by subtracting the derivatives of this invariant image from the derivatives of the original image. Since only material edges are present in both images, shadow edges remain when subtracting the two images.

The above methods use photometric information for edge classification ignoring geometric information completely. However, geometric information may contain valuable information for edge classification which has not been used so far. Taking the example presented in Fig. 1, on the left we only use photometric information to detect edges caused by shadows. On top, high probable regions are isolated, and on the bottom the im-

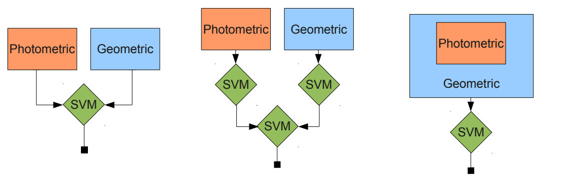

(a) Early Fusion (b) Late Fusion (c) Low-Level Integration

Figure 2. Illustration of the different combination frameworks to combine photometric and geometric features.

age without shadow edges is shown. On the right, we use the geometry on top of the photometric space to detect the same type of edges. Notice how shadow-edges are more localized, for example the surrounding area of the mirrors is now properly discovered.

Therefore, in this paper, the aim is to use both photometric *and* geometric information for edge classification and evaluate how they can be combined. It is shown that the local geometry of edges is as important as photometric information to distinguish between different types of edges. For instance, in [12], the use of monochromatic cues provide a more accurate perceptual recovering of intrinsic images. However, strong illumination constraints difficult its applicability to natural images. In [4], several local photometric and geometric descriptors are evaluated to only detect shadow edges. However, the coarse combination of all the features leads to a very costly approach.

2 Combination Framework

To take advantage of both photometric and geometric information for the classification of the physical nature of edges, both types of features need to be combined. Three approaches are compared in this paper. First, two widely known data-driven frameworks are detailed and finally, a more fundamental method is formulated to integrate the photometric color variants with the geometric features. Fig. 2 illustrates the differences between the approaches here described.

Early fusion (EF). This approach is based on concatenating the different image descriptors into one large feature vector, which contains the information of all used features [11]. The method combines specific properties of the descriptors into a larger feature space. This increases the dimensionality of the feature space significantly, but it is also able to distinguish whether geometric information (e.g. vertical or horizontal edges) is relevant or not.

One of the main drawbacks of this approach is related to feature space encoding. First, each new feature that is added requires learning the whole model again. Secondly, when features with different characteristics are concatenated (like number of dimensions or different bin distributions), then a tedious cross-validation procedure is required to find the proper weighting for the different features. Further, more examples are required to properly model the increase of dimensionality.

Late fusion (LF). The second data-driven fusion approach avoids the problem of merging unbalanced or very different types of feature vectors by learning each of the descriptors independently, and after combining the posterior probabilities of each of the classifiers [6].

The advantage of this approach is that the learning stage becomes faster, since feature vectors are much smaller. The output of each of the single classifiers is used to learn a new linear SVM, which learn the importance of each one of the features, and it allows to avoid the raw high-dimensional features. However, this approach fails when specific relations between the features are important, because it will be unable to detect them. For example, the combined occurrence of a specific geometric response (e.g. diagonal edges) and a high contrasted patch response can not be learned because one classifier can only process one type of information simultaneously.

Low-level integration (LLI). In contrast to the previous approaches, photometric features are used as a pre-processing step for the geometric features. This implies that this low-level integration approach can also be combined with the previous data-driven approaches.

To describe a patch, the image is converted into the desired photometric invariant space, and the geometric descriptor is applied on each one of their channels independently. The main difference with conventional geometric features applied only on the illuminant space, is that material edges are similarly encoded than any other type of edges. With this approach, by definition, the patch is described in two clear directions: the variant and the invariant. Therefore, both types of edges will be easily differentiable. The final descriptors need to be normalized by the magnitude of the body reflectance in order to be comparable among images [14].

This low-level integration approach has not been used before to classify the different types of edges. Other works have explored the invariances of several color spaces to improve the discriminative power of the descriptors [13]. Their aim is to obtain feature descriptions that can be easily found under different illumination conditions. In contrast, in our case, we are looking for the opposite, discover the edges that are variant to certain characteristics (e.g. shadows or highlights).

3 Experiments

To validate our method, different experiments are performed. First, a shadow edge classifier is presented. Next, we extend shadow edge classification to other types (specular edges in particular) by using the same approach. Finally, the proposed method is used to estimate the illuminant direction in a scene.

For extracting the geometric information in the images, we use the SIFT descriptor [9] and the Weibull descriptor [2]. To obtain the photometric information we use the quasi-invariants approach [14].

3.1 Edge Classification

Shadow-Edge Classification. The aim of this experiment is to derive the common characteristics of shadow edges. To this end, we use the annotations from [4], which contains 7047 patches extracted from 3699 images from outdoor and indoor scenarios [8]. All patches are 19×19 pixels, and split into *shadow patches*, which contains at least one clear shadow edge (in any position of the patch) and *non-shadow patches*, which corresponds to patches without any shadow edge.

Each patch is described by one of the previous discussed descriptors. Then, a SVM-classifier is used to evaluate a 10-fold cross-validation. Single features are learnt with linear and RBF-kernel, and the best results are kept. Parameters are tuned by cross-validation.

Results are summarized in Table 1, where we report the Area under the ROC curve (AUC). This metric is invariant to an unbalanced number of examples per class, which is desirable to detect image effects that only appear from time to time (e.g.. shadow or specular edges, reflections). It can be observed that none of the single features is able to exceed 0.77. The combination by means of Early (EF) or Late (LF) Fusion ends up with a similar performance of 0.83. However, when combining the integration of the new low-level features, performance is still improved to 0.84, outperforming previously reported methods.

Shadows in natural images can take any contrast or shape. However, our improvement can be explained by the fact that shadows tend to be mainly diagonal oriented, or in a patch, all of them usually follow the same direction, and this information was lost in the previous approaches.

Highlight-Edge Classification. As an illustration of the generalization of our method, we extended our edge classification technique to highlights, since specularities are valuable clues for color constancy [5]. Given the photometric invariance that the quasi-invariants provide, we focus now on the opponent color space, which

Features	Shadow	Highlight
Quasi-invariants (QI)	0.77	0.78
SIFT	0.77	0.82
Weibull	0.77	0.76
LLI (SIFT + QI)	0.79 (1)	0.83
LLI (Weibull + QI)	0.80 (2)	0.85
EF (SIFT + Weibull + QI)	0.83	0.82
LF (SIFT + Weibull)	0.81	0.83
LF (SIFT + QI)	0.82	0.83
LF (Weibull + QI)	0.79	0.80
LF (SIFT + Weibull + QI)	0.83	0.83
EF((1) + (2))	0.83	0.85
LF((1) + (2))	**0.84**	**0.87**
Best previously reported [4]	0.82	-

Table 1. Results for edge classification. The results denote the area under the corresponding ROC-curve.

accompanies the specular variant and invariant space. For more details on quasi-invariants, we refer to [14]. 400 patches of highlights and 400 of non-highlights are extracted from a car exposition dataset [10].

Results shown in Table 1 report a similar behaviour as the classification of shadow-edges. The most remarkable fact is that the low-level integration framework is the one performing better. This can be easily explained by the fact that the specularites are characterized as spots, i.e. the edges around of the highlight follow a circular pattern. In this case, our new approach, is the only one able to detect such patterns in the photometric and geometric space at the same time. The combination improves over a single feature by more than 7%.

3.2 Illumination Direction

To demonstrate the possibilities of using shadow edge classification in different applications, we use our method to estimate the direction of the illumination source. In contrast to [7] (use multiple cues), we only focus on shadow edges to estimate the sun position. Given a single image, the LLI (SIFT + QI) descriptor is extracted every 9 pixels. Then, by using the previously learned classifier from Section 3.1. The classifier responses are used to weight the edge orientations of the descriptor. After being accumulated into an histogram of 8 bin orientations, a SVM-regression is learnt to recover the sun position.

We use a subset of real still images extracted from outdoor webcams, which were kindly provided by [7]. It consists of 13 different scenarios containing 391 im-

Figure 3. Qualitative results. Shadow edge detection and sun position estimation.

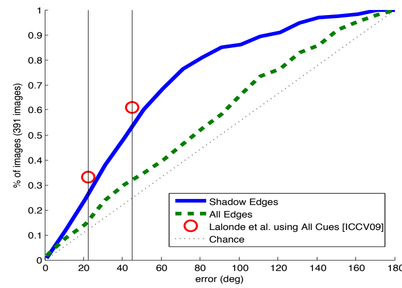

Figure 4. Quantitative results. The percentatge of correctly estimated images taking into account the error (in degrees).

ages in total. The evaluation methodology is based on cross-validation of sequences, leaving each sequence out for testing once. Some scenarios with the estimated (gray) and the real (blue) shadows are shown in Fig. 3.

Fig. 4 reports the cumulative histogram of errors in sun position estimation. A large improvement is observed when using shadow-edges or all the edges. For example, by assuming at most 45 degrees of error, more than 54% of the images are well classified, with respect to the 31% achieved by counting all the edges. By using only the shadow-edges with a very naive scheme, the method achieves comparable results to [7].

4 Conclusions

In this paper, we present the idea to improve edge classification by enriching the widely used photometric information with several geometric features. A principled approach has been proposed to obtain an integrated combination of photometric and geometric information. With this approach, the combined representation performs the best among the single features.

From the experiments on different datasets, it is shown that the addition of the geometry of edges is an important visual cue to distinguish between different edge types. It is shown that by combining both cues improves over using only one by 7% for shadows and highlights. This remarks the fact that it not only works for shadows. Other types of edges can be also identified by using the appropriate photometric color space as a basis of the framework.

Acknowledgements

The authors acknowledge the support of the Spanish Research Programs Consolider-Ingenio 2010: MIPRCV (CSD200700018); Avanza I+D ViCoMo (TSI-020400-2009-133) and DiCoMa (TSI-020400-2011-55); along with the Spanish projects TIN2009-14501-C02-01 and TIN2009-14501-C02-02.

References

[1] G. Finlayson and S. Hordley. Color constancy at a pixel. *JOSA*, 2001.

[2] J. M. Geusebroek and A. W. M. Smeulders. A six-stimulus theory for stochastic texture. *IJCV*, 2005.

[3] T. Gevers and H. Stokman. Classifying color edges in video into shadow-geometry, highlight, or material transitions. *IEEE Trans. Multimedia*, 2003.

[4] A. Gijsenij and T. Gevers. Shadow edge detection using geometric and photometric features. In *Proc. ICIP*, 2009.

[5] A. Gijsenij, T. Gevers, and J. van de Weijer. Physics-based edge evaluation for improved color constancy. In *Proc. CVPR*, Miami Beach, Florida, USA, 2009.

[6] J. Kittler, M. Hatef, R. Duin, and J. Matas. On combining classifiers. *PAMI.*, 1998.

[7] J.-F. Lalonde, A. A. Efros, and S. G. Narasimhan. Estimating natural illumination from a single outdoor image. In *Proc. ICCV*, 2009.

[8] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *Proc. CVPR*, 2006.

[9] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 2004.

[10] M. Ozuysal, V. Lepetit, and P.Fua. Pose estimation for category specific multiview object localization. In *Proc. CVPR*, 2009.

[11] C. G. M. Snoek, M. Worring, and A. W. Smeulders. Early versus late fusion in semantic video analysis. In *ACM Multimedia*, 2005.

[12] M. F. Tappen, W. T. Freeman, and E. H. Adelson. Recovering intrinsic images from a single image. *PAMI.*, 2005.

[13] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek. Evaluating color descriptors for object and scene recognition. *PAMI.*, 2010.

[14] J. van de Weijer, T. Gevers, and J.-M. Geusebroek. Edge and corner detection by photometric quasi-invariants. *PAMI.*, 2005.