

Specularity, the Zeta-image, and Information-Theoretic Illuminant Estimation

Mark S. Drew¹, Hamid Reza Vaezi Joze¹, and Graham D. Finlayson²

¹ School of Computing Science,
Simon Fraser University
{mark,hrv1}@cs.sfu.ca

² School of Computing Sciences,
University of East Anglia
graham@cmp.uea.ac.uk

Abstract. Identification of illumination is an important problem in imaging. In this paper we present a new and effective physics-based colour constancy algorithm which makes use of a novel log-relative-chromaticity planar constraint. We call the new feature the Zeta-image. We show that this new feature is tied to a novel application of the Kullback-Leibler Divergence, here applied to chromaticity values instead of probabilities. The new method requires no training data or tunable parameters. Moreover it is simple to implement and very fast. Our experimental results across datasets of real images show the proposed method significantly outperforms other unsupervised methods while its estimation accuracy is comparable with more complex, supervised, methods. As well, the new planar constraint can be used as a post-processing stage for any candidate colour constancy method in order to improve its accuracy.

1 Introduction

Identification of illumination is an important problem in image processing for digital cameras, for both still images and video. In a scene consisting of dielectric materials (e.g., plastics, and indeed most non-metals) there is typically substantive specular content. This does not necessarily mean extremely bright mirror-like reflection, but can consist for example of the glint of light reflected from blades of grass, or the sheen of light reflected from a desk surface. For non-metals, this very common specular content is an important indicator of the colour of the lighting in a scene.

Many colour constancy algorithms have been proposed (see [1, 2] for an overview). The foundational colour constancy method, the so-called White-Patch or Max-RGB method, estimates the light source colour from the maximum response of the different colour channels. Another well-known colour constancy method is based on the Grey-World hypothesis, which assumes that the average reflectance in the scene is achromatic. Grey-Edge is a recent version of the Grey-World hypothesis that says: the average of the reflectance differences in a

scene is achromatic [3]. The Gamut Mapping algorithm, a more complex and more accurate algorithm, was introduced by Forsyth [4]. It is based on the assumption that in real-world images, for a given illuminant one observes only a limited number of colours. Several extensions have been proposed.

Lee [5] proposed a method which uses specularly to compute illumination by using the fact that in the CIE chromaticity diagram the coordinates of the colours from different points from the same surface will fall on a straight line connected to the specular point. This is the case when the light reflected from a uniform surface is an additive mixture of the specular component and the diffuse component. This seminal work initiated a substantial body of work on identifying specular pixels and using these to attempt to discover the illuminant (e.g. [6, 7]). Although these works are theoretically strong, none of them reports performance over real world datasets of images with and without specularities.

In this paper, we set out a new discovery, consisting of a planar constraint that must be obeyed, in a certain standard model of reflectance, by specular or near-specular pixels in an image. The new feature involved we call the Zeta-image,¹ and below we show that this feature is tied to the information-theoretic concept of applying one distribution to generate bitstring codes for another; here we view chromaticity components, which add to 1, in the role of probabilities. We present a novel physics-based colour constancy algorithm based on a log-relative-chromaticity planar constraint, which requires no training data or tunable parameters. It is easy to implement and very fast compared to more complex colour constancy methods such as gamut mapping. Our experimental results over three large datasets of both laboratory and real world images show that the proposed method significantly outperforms other unsupervised methods while its accuracy of estimation is comparable with more complex methods that need training data and tunable parameters.

2 Relative Chromaticity Near Specular Point

2.1 Image Formation Model and Relative Chromaticity

Let the RGB 3-vector at each pixel be denoted \mathbf{R} , with components $R_k, k = 1..3$. Let the RGB 3-vector for the light itself as seen by the camera be e_k , and let the 3-vector for the reflectance at a pixel as seen by the camera, under equi-energy white light, be s_k . Now in a product-of-vectors simple model [8] we approximately have the matte (“body”, i.e., non-specular) RGB value at that pixel equal to

$$R_k \simeq \sigma s_k e_k / q_k \tag{1}$$

where σ is shading. In the standard Lambertian model for matte shading, σ equals lighting-direction dotted into surface normal. Here, q_k is a triple giving the overall (integrated) camera sensor strength [9].

If we also consider an additional specular component, this equation becomes

$$R_k \simeq \sigma s_k e_k / q_k + \beta e_k \tag{2}$$

¹ Patent applied for.

where β represents the amount of specular component at that pixel. The value of β for a pixel will depend upon the lighting direction, the surface normal, and the viewing geometry. Here, the specular component βe_k is simply assumed to be the same colour as the light itself, in a Neutral Interface Model [5] for dielectrics. For purposes of discovering properties of this equation, let us assume for the time being that β is simply a constant — in actuality it will be a scalar property of each pixel and this issue is further discussed below in §2.3. Let us lump values $\sigma s_k/q_k$ into a single quantity and for convenience call this simply s_k . Now we have

$$R_k = s_k e_k + \beta e_k \tag{3}$$

The chromaticity ρ_k is colour without magnitude, in an L_1 norm: $\rho = \{R, G, B\}/(R + G + B)$, so here we have

$$\rho_k = \frac{s_k e_k + \beta e_k}{\sum_{j=1}^3 (s_j e_j + \beta e_j)} \tag{4}$$

Let the chromaticity of the light itself be denoted $\rho_k^e = e_k/\sum_j e_j$. Now here we wish to examine the properties of the Relative Chromaticity, which we define to be the chromaticity divided by the chromaticity of the light, ρ_k^e . Let us call this quotient chromaticity χ_k , so that

$$\chi_k = \frac{\rho_k}{\rho_k^e} = \frac{s_k e_k + \beta e_k}{\sum_{j=1}^3 (s_j e_j) + \beta \sum_{j=1}^3 e_j} \cdot \frac{\sum_{j=1}^3 e_j}{e_k} \tag{5}$$

where all divisions are taken to be component-wise. Dividing by the light chromaticity is the main innovative step in this paper: it is an *ansatz* that we claim will bear fruit by generating a constraint on the illuminant colour.

For convenience, let $E \equiv \sum_{j=1}^3 e_j = |e|$ where $|\cdot|$ is the L_1 norm. Then we arrive at

$$\chi_k = \frac{s_k + \beta}{\frac{(\sum_j s_j e_j)}{E} + \beta} \tag{6}$$

So, for a pixel with no matte component s_k but only a purely specular component, we would have $\chi_k \equiv 1$ for all 3-vector elements $k = 1..3$.

2.2 Log-Relative-Chromaticity and Planar Constraint

Next we show that in fact log values are preferable, in that a simple planarity constraint falls out of the formulation once we move to the log domain.

Let us define a new quantity, ψ , which is the logarithm of the ratio χ defined above: we call this the Log-Relative-Chromaticity:

$$\psi_k = \log(\chi_k) = \log(\rho_k/\rho_k^e) \tag{7}$$

Now *near* a specular point, we can take the limit as $(1/\beta) \rightarrow 0$. Let $\alpha = 1/\beta$. Then in the limit as specularity increases, ψ goes to

$$\psi_k = \lim_{\alpha \rightarrow 0} \log \left\{ (\alpha s_k + 1) / \left(\alpha \sum_j (s_j e_j) / E + 1 \right) \right\}$$

Using a Maclaurin series,

$$\psi_k = \alpha \left(s_k - \frac{\sum_j s_j e_j}{E} \right) + O(\alpha^2) \quad (8)$$

Omitting $O(\alpha^2)$, we note by inspection that the quantity ψ_k is *orthogonal to the illuminant vector*:

$$\sum_k \psi_k e_k \equiv 0, \text{ so also } \sum_k \psi_k \rho_k^e = 0 \quad (9)$$

Therefore we have a planar constraint on image pixels that are near-specular:

Planar Constraint: *For near-specular pixels, Log-Relative-Chromaticity values are orthogonal to the light chromaticity.*

Note that in eq. (8) above we have expressed this orthogonality in a different way than the usual, Euclidean-norm based calculation of the part of the vector \mathbf{s} that is orthogonal to vector \mathbf{e} , viz. $(s_k - \mathbf{s} \cdot \hat{\mathbf{e}} \hat{e}_k)$ for normalized light vector $\hat{\mathbf{e}}$. Nevertheless, it is easy to verify that eq. (9) does indeed hold. Eq. (8) means an L_1 -based projection onto the plane orthogonal to the light.

2.3 Varying Specular Factor β

The fact that specular scalar factor β is not a constant makes no difference to the argument: for near-specular pixels any value of β still leads to quantity ψ_k lying in the plane orthogonal to the illuminant.

3 Planar Constraint Method

Here we begin construction of an algorithm by considering first a simple search method as motivation, and then stating an analytic solution.

3.1 Global Search

The planar constraint suggests that the dot product for near-light-colour (e.g., specular) pixels is minimized for the correct illuminant. This points to a useful descriptor for finding the specular point.

Suppose we were to assume that for any candidate illuminant the lowest 10-percentile, say, of dot-product values (9) could be near-specular pixels. Now, to find the correct illuminant, we need to minimize dot-product values (9) over candidate illuminants for those lowest 10-percentile pixels. Thus an optimization can be stated as follows:

Define the **Zeta-image** ζ as the dot-product of the log-relative-chromaticity $\boldsymbol{\psi}$, eq. (7), with a putative light direction:

$$\zeta = -\boldsymbol{\psi} \cdot \boldsymbol{\rho}^e = -\log(\boldsymbol{\rho} / \boldsymbol{\rho}^e) \cdot \boldsymbol{\rho}^e \quad (10)$$

$$\begin{aligned}
 \text{Minimize :} \quad & \min_{\rho^e} \mathcal{I} = \sum_{\psi \in \Psi_0} |\zeta| \\
 \text{subject to} \quad & \sum_{k=1}^3 \rho_k^e = 1, 0 < \rho_k^e < 1, k = 1..3
 \end{aligned} \tag{11}$$

where Ψ_0 is the set of pixel dot-product values (9) with the candidate illuminant chromaticity ρ^e that are in the lowest 10-percentile.

The meaning of eq. (11) is that we first carry out a search, over possible illuminant chromaticities ρ^e . This can be phrased as either an optimization-based approach or, as here, a simple hierarchical grid search. Then we adopt a heuristic that says that the lowest 10-percentile of values of dot-products with the candidate illuminant could be specular or in general illuminant-coloured. For these pixels we calculate the sum of absolute values of dot-products and take as the best candidate light that which delivers the minimum sum.

Fig. 1(a) shows an input image, and Fig. 1(b) shows a boolean map of the lowest 10-percentile of dot-product values (9) with the correct illuminant chromaticity of that image. In contrast, if we show the lowest 10-percentile of dot products with the chromaticity of an incorrect light, where ψ values are constructed using that incorrect light, the boolean map identifying putative specular/light-coloured pixels is as displayed in Fig. 1(c). We see that using the correct light produces a much more plausible map of possible pixels that will help identify the light.

The float-valued Zeta-image is displayed in Fig. 1(d). We show next that we can directly use the Zeta-image to analytically find the correct illuminant.

3.2 Analytic Solution

Having motivated the method, we now state an analytic solution that in fact produces excellent results and is very simple and fast. Suppose we identify a possible set Ψ_0 of specular pixels by any convenient method — e.g., we could simply take the top 5% of brightness.

Let the number of bright pixels be N . Then our analytic solution is as follows:

Theorem: Up to an L_1 normalization, the light colour is given by the geometric mean

$$\rho_k^e = g(\rho_k) \equiv \left(\prod_{i=1}^N \rho_k^i \right)^{1/N}, k = 1..3 \tag{12}$$



Fig. 1. (a): Image taken under measured illuminant. (b): Light-coloured pixels identified using planar constraint, when correct illuminant is chosen. (c): Putative light-coloured pixels for incorrect illuminant. (d): Analytic Zeta-image (float, 1-intensity).

To prove this, let us first solve an auxiliary optimization replacing (11). For physical lights, we expect ζ from eq. (10) to be non-negative, so [with no absolute-value bars as in \mathcal{I} in eq. (11)] we first solve

$$\min_{\boldsymbol{\rho}^e} \sum_i \zeta_i + \lambda (\sum_{k=1}^3 \rho_k^e - 1) \quad (13)$$

where $i = 1..N$ ranges over the N pixels in $\boldsymbol{\psi} \in \Psi_0$, and λ is a Lagrange multiplier enforcing that $\boldsymbol{\rho}^e$ is a chromaticity. Taking partial derivatives with respect to ρ_k^e we have the normal equations

$$-\sum_i [\log(\rho_k^i) - \log(\rho_k^e) - 1] + \lambda = 0, \text{ with solution } \rho_k^e = g(\rho_k) \cdot \exp(-(N+\lambda)/N).$$

The meaning of (13) then is: **The planarity constraint yields the geometric mean of the chromaticities as the solution for the light**, up to trivial scaling of the L_1 -norm.

However so far we have omitted absolute value bars, with a full optimization minimizing $\mathcal{I} = \sum_i |\zeta_i|$. We now observe that the form (10) formally has the structure of the Kullback-Leibler Divergence from information theory, in that chromaticities for image, ρ_k^i , and light, ρ_k^e all add to unity: $\sum_{k=1}^3 \rho_k = 1$. We are minimizing $\sum_{i=1}^N [\sum_{k=1}^3 -\rho_k^e \log(\rho_k^i/\rho_k^e)]$, which has the K-L structure except that instead of summing over probabilities for symbols we are here summing over colour-channels. Thus each dot product (the sum over k at each pixel) is necessarily nonnegative since it represents the extra bits required to code samples from ρ_k^e when using a code based on ρ_k^i . Hence we can simply consider the minimization (13), with solution (12) up to scaling. \square

As a final step, we calculate a final value for ρ_k^e by trimming pixels to the least-10% values of the Zeta-image ζ and recalculating the geometric mean (12).

Areas of images that are specular tend to be bright. According to our theory the geometric mean of pixels in these bright (generally, specular) regions is the optimal estimate for the illuminant. Our insight is in contradistinction to the work of Choudhury and Medioni [10] and Funt and Shi [11] which proposed finding the max after a local mean calculation (e.g. after local blurring). In the presence of specular highlights we instead claim that illuminant estimation is the mean of the max, not the converse. Indeed, using the correct ordering is crucial (a fact borne out by our experiments reported below).

3.3 2nd Algorithm: Planar Constraint Applied as Post-processing

Suppose we have an estimate $\boldsymbol{\rho}^{e*}$ of the correct illuminant, from any colour constancy algorithm. If our estimate is indeed near the correct illuminant we can then identify as near-specular pixels those whose absolute value of dot-product eq. (9) with the candidate light chromaticity falls in the lowest 10-percentile, say (i.e., nearest to zero). Forming a Singular Value Decomposition of $\boldsymbol{\psi}$ values for those pixels determines the best-fit plane. Within the model presented here that normal should be close to the illuminant chromaticity in direction. Because of the additional evidence brought to bear by eq. (9) we expect the estimate to

improve. If instead the illuminant estimate is wrong, then we have found that the above SVD step will almost always not change it much and no harm is done by carrying out this post-processing step. We carry out

$$\begin{aligned} \psi_k &= \log\left(\frac{\rho_k}{\rho_k^{e^*}}\right), \quad k = 1..3; \quad \Psi = \psi^T \rho^{e^*}; \\ \psi(\Psi_0) &= \text{SVD}(\psi(\Psi_0)) = U \text{diag}(\mathbf{d}) V^T; \\ \rho^e &= \mathbf{V}_3 / |\mathbf{V}_3|; \quad \text{success} = (d_3 \text{ small}) \& (\rho^e \simeq \rho^{e^*}) \end{aligned} \quad (14)$$

where ρ^{e^*} is the estimate of the illumination provided by a colour constancy method, Ψ_0 is the lowest 10-percentile of Ψ , and ρ^e is the estimate of the illumination based on the planar constraint.

The meaning of eq. (14) is that, for any estimate ρ^{e^*} of the light chromaticity ρ^e , if the model (9) is obeyed around the light point then SVD should produce an estimate of the light that agrees with ρ^{e^*} .

We demonstrate below that this planar constraint does indeed improve the estimate of ρ^e , verifying the suitability of the plane constraint applied as a post-processing step, for any candidate colour constancy algorithm. In the next section we will demonstrate the substantial improvement delivered by this simple planar constraint when added to each of several well-known colour constancy algorithms as a post-processing step.

4 Experiment Results

4.1 Datasets

We apply our proposed method to three different real-image datasets [12–14] and compare our results to other colour constancy algorithms.

Our first experiment uses the Barnard dataset [12], denoted the SFU Laboratory dataset, which contains 321 measured images under 11 different measured illuminants. The scenes are divided into two sets as follows: minimal specularities (22 scenes, 223 images – i.e., 19 missing images); and non-negligible dielectric specularities (9 scenes, 98 images – 1 illuminant is missing for 1 scene).

For a more real-world (out of the laboratory) image experiment we used the re-processed version of the Gehler colour constancy dataset [15], denoted the ColorChecker dataset [13]. This dataset consists of 568 images, both indoor and outdoor. The illuminant ground truth for these images is known because each image has a Macbeth ColorChecker placed in the scene. The ColorChecker is masked off in tests.

Ciurea and Funt [14] introduced the GreyBall dataset, which contains 11346 images extracted from video recorded under a wide variety of imaging conditions. The images are divided into 15 different clips taken at different locations. The ground truth was acquired by attaching a grey sphere to the camera, displayed in the corner of the image. This grey sphere must be masked during experiments.

4.2 Previous Methods

To compare, we use the standard well-known colour constancy methods: White-Patch, Grey-World, and Grey-Edge implemented by [3]. For Grey-Edge we use optimal settings, which differ per dataset [16] ($p = 7$, $\sigma = 4$ for the SFU Laboratory dataset and $p = 1$, $\sigma = 6$ for the ColorChecker dataset). We also use the result provided by Gijsenij and et al. [17] for pixel-based gamut mapping, using the best general gamut mapping setting, which is for *1st-jet* as reported in [17] (although we could not precisely match their exact results using the code they released). For other methods we use results as provided by Gijsenij [2, 16].

For methods which need training data, such as the gamut mapping methods, in the SFU Laboratory dataset 31 images (all images recorded under the *syl-50MR16Q*-illuminant) were used for computation of the canonical gamut, and subsequently these were omitted from the test set. For the ColorChecker dataset, three-fold cross-validation was used to learn the canonical gamut (with the folds as well as the ground truth supplied with the original dataset). Testing for supervised methods is as described in [2], §VII-A.

4.3 Post-processing

Table 1 shows the accuracy of the plane constraint eq. (14) in §3.3 as a post-processing step applied to the results of each of well-known colour constancy algorithms, in order to improve the estimate.

Table 1. Median of angular errors for well-known colour constancy algorithms for the SFU Laboratory [12] dataset and ColorChecker dataset [13], plus result after post-processing with planar constraint eq. (14) for each colour constancy algorithm

Method	SFU Lab.	ColorChecker
White-Patch	6.5°	5.7°
White-Patch + Planar Con.	5.1°	4.4°
Grey-World	7.0°	6.3°
Grey-World + Planar Con.	5.0°	4.3°
Grey-Edge	3.2°	4.3°
Grey-Edge + Planar Con.	2.7°	3.8°

4.4 Global Search and Analytic Solution Experiment

Table 2 indicates the accuracy of the proposed methods for the SFU Laboratory dataset [12], the ColorChecker dataset [13] and the GreyBall dataset [14], in terms of the mean and median of angular errors, for several colour constancy algorithms applied to these datasets. For those methods which need data-dependent tunable parameters, we utilize optimal parameters for their dataset. For an overview of results of different algorithms on these datasets refer to [2, 16].

To our knowledge, for the SFU Laboratory dataset the Planar Constraint Search eq. (11) does best in terms of median angular error compared to any reported colour constancy method, even those needing training data. We do note that for this dataset the Planar Constraint Search eq. (11) is not the best for the ColorChecker dataset, with Gamut Mapping methods performing better. However, both Planar Search and the Analytic method of §3.2 (Geomean) do as well or better than the other relatively fast methods for the GreyBall dataset, and are only bested by the much more complex method [18].

Run-times average 5.2s for Planar-Constraint Search and 415ms for the Analytic method, compared to 617ms for the GreyEdge algorithm and 63.2s for 1st-Jet Gamut Mapping, operating on the SFU Laboratory dataset using (un-optimized) Matlab.

Table 2. Angular errors for several colour constancy algorithms for SFU Laboratory dataset [12], ColorChecker dataset [13] and GreyBall dataset [14]

Dataset Methods	SFU Laboratory		Color Checker		Gray Ball	
	Median	Mean	Median	Mean	Median	Mean
White Patch	6.5°	9.1°	5.7°	7.4°	5.3°	6.8°
Gray World	7.0°	9.8°	6.3°	6.4°	7.0°	7.9°
Gray Edge	3.2°	5.6°	4.5°	5.3°	4.7°	5.9°
Bayesian [15]	-	-	3.5°	4.8°	-	-
Gamut Mapping	2.3°	3.7°	2.5°	4.1°	5.8°	7.1°
Gamut Mapping 1jet [17]	2.1°	3.6°	2.5°	4.1°	5.8°	6.9°
Natural Image Statistics	-	-	3.1°	4.2°	3.9°	5.2°
Planar Constraint Search	1.9°	4.3°	2.8°	4.1°	4.6°	5.9°
Geomean	2.1°	6.2°	2.7°	4.2°	4.7°	5.8°

5 Conclusions

In this paper we present a novel physics-based insight regarding a plane constraint that obtains for log-relative-chromaticity values near the illuminant point (for white surfaces, or specularities in the neutral-interface model). This insight provides a useful and very simple method for identifying the illuminant chromaticity. Experiment results over datasets consisting of laboratory images and of real-world images demonstrate that the proposed method significantly outperforms other unsupervised methods while its accuracy of illuminant estimation is comparable with the best (supervised) methods but much faster. As well, the plane constraint can also be brought to bear to improve estimates provided by other illuminant estimation algorithms, in a post-processing step; experimental results indicate that estimate errors can be reduced by some 15 percent by this simple and very fast mechanism.

References

1. Hordley, S.D.: Scene illuminant estimation: past, present, and future. *Color Research and Application* 31, 303–314 (2006)
2. Gijsenij, A., Gevers, T., van de Weijer, J.: Computational color constancy: Survey and experiments. *IEEE Trans. on Image Proc.* 20 (2011)
3. van de Weijer, J., Gevers, T.: Color constancy based on the grey-edge hypothesis. In: *Int. Conf. on Image Proc.*, pp. II:722–II:725 (2005)
4. Forsyth, D.: A novel approach to color constancy. In: *Int. Conf. on Computer Vision*, pp. 9–18 (1988)
5. Lee, H.: Method for computing the scene-illuminant chromaticity from specular highlights. *J. Opt. Soc. Am. A* 3, 1694–1699 (1986)
6. Lehmann, T.M., Palm, C.: Color line search for illuminant estimation in real-world scenes. *J. Opt. Soc. Amer. A* 18, 2679–2691 (2001)
7. Tan, R., Nishino, K., Ikeuchi, K.: Color constancy through inverse-intensity chromaticity space. *J. Opt. Soc. Am. A* 21, 321–334 (2004)
8. Borges, C.: Trichromatic approximation method for surface illumination. *J. Opt. Soc. Am. A* 8, 1319–1323 (1991)
9. Drew, M., Finlayson, G.: Multispectral processing without spectra. *J. Opt. Soc. Am. A* 20, 1181–1193 (2003)
10. Choudhury, A., Medioni, G.: Perceptually motivated automatic color contrast enhancement. In: *Color and Reflec. in Imaging and Comp. Vis. Workshop*, pp. 1893–1900 (2009)
11. Funt, B.V., Shi, L.: The rehabilitation of MaxRGB. In: *18th Color Imaging Conf.*, pp. 256–259 (2010)
12. Barnard, K., Martin, L., Funt, B.V., Coath, A.: A data set for colour research. *Color Res. and Applic.* 27, 147–151 (2002)
13. Shi, L., Funt, B.V.: Re-processed version of the Gehler color constancy dataset of 568 images (2010), <http://www.cs.sfu.ca/~colour/data/>
14. Ciurea, F., Funt, B.V.: A large image database for color constancy research. In: *Color Imag. Conf.*, pp. 160–164 (2003)
15. Gehler, P., Rother, C., Blake, A., Minka, T., Sharp, T.: Bayesian color constancy revisited. In: *Comp. Vis. and Patt. Rec.* (2008)
16. Gijsenij, A.: Color constancy: Research website on illuminant estimation (2012), <http://staff.science.uva.nl/~gijsenij/colorconstancy/>
17. Gijsenij, A., Gevers, T., Weijer, J.: Generalized gamut mapping using image derivative structures for color constancy. *Int. J. of Comp. Vis.* 86, 127–139 (2008)
18. Gijsenij, A., Gevers, T.: Color constancy using natural image statistics and scene semantics. *Trans. on Patt. Anal. and Mach. Intell.* 33, 687–698 (2011)