UNIVERSITY OF AMSTERDAM

UvA-DARE (Digital Academic Repository)

Generalized gamut mapping using image derivative structures for color constancy

Gijsenij, A.; Gevers, T.; van de Weijer, J.

DOI
10.1007/s11263-008-0171-3

Publication date
2010

Document Version
Final published version

Published in
International Journal of Computer Vision

Link to publication

Citation for published version (APA):
Gijsenij, A., Gevers, T., & van de Weijer, J. (2010). Generalized gamut mapping using image derivative structures for color constancy. *International Journal of Computer Vision*, *86*(2-3), 127-139. https://doi.org/10.1007/s11263-008-0171-3

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, P.O. Box 19185, 1000 GD Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (https://dare.uva.nl)

Download date:14 Nov 2026

Int J Comput Vis (2010) 86: 127–139
DOI 10.1007/s11263-008-0171-3

Generalized Gamut Mapping using Image Derivative Structures for Color Constancy

Arjan Gijsenij · Theo Gevers · Joost van de Weijer

Received: 11 February 2008 / Accepted: 6 August 2008 / Published online: 1 October 2008
© The Author(s) 2008. This article is published with open access at Springerlink.com

Abstract The gamut mapping algorithm is one of the most promising methods to achieve computational color constancy. However, so far, gamut mapping algorithms are restricted to the use of pixel values to estimate the illuminant.

Therefore, in this paper, gamut mapping is extended to incorporate the statistical nature of images. It is analytically shown that the proposed gamut mapping framework is able to include any linear filter output. The main focus is on the local n-jet describing the derivative structure of an image. It is shown that derivatives have the advantage over pixel values to be invariant to disturbing effects (i.e. deviations of the diagonal model) such as saturated colors and diffuse light. Further, as the n-jet based gamut mapping has the ability to use more information than pixel values alone, the combination of these algorithms are more stable than the regular gamut mapping algorithm. Different methods of combining are proposed.

Based on theoretical and experimental results conducted on large scale data sets of hyperspectral, laboratory and real-world scenes, it can be derived that (1) in case of deviations of the diagonal model, the derivative-based approach outperforms the pixel-based gamut mapping, (2) state-of-the-art algorithms are outperformed by the n-jet based gamut mapping, (3) the combination of the different n-jet based gamut mappings provide more stable solutions, and (4) the fusion strategy based on the intersection of feasible sets provides better color constancy results than the union of the feasible sets.

Keywords Color Constancy · Illuminant estimation · Gamut mapping · n-jet

A. Gijsenij (✉) · T. Gevers
University of Amsterdam, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
e-mail: gijsenij@science.uva.nl

T. Gevers
e-mail: gevers@science.uva.nl

J. van de Weijer
Computer Vision Center (CVC), Campus Universitat Autonoma de Barcelona, 08193 Bellaterra, Spain
e-mail: joost@cvc.uab.es

1 Introduction

The color of a light source has a significant influence on the measurements of object colors in a scene. Fortunately, humans have, to some extent, the ability of color constancy: they perceive the same color of an object despite large differences in illumination (Delahunt and Brainard 2004; Foster et al. 2006). Various computer vision related topics like human-computer interaction (Yang et al. 1998), color appearance models (Fairchild 2005) and color feature extraction (Gevers and Smeulders 2000) would benefit from a similar ability.

Many color constancy algorithms have been proposed, see (Hordley 2006) for a recent overview. Two widely used algorithms make use of simple statistics of the image to estimate the color of the illuminant. One is based on the assumption that the average color in a scene is achromatic, called the *Grey-World assumption* (Buchsbaum 1980), while the other assumes that the maximum response in a scene is caused by a perfect reflectance, called the *White-Patch* assumption (Land 1977). Similar methods are the Shades of Grey algorithm (Finlayson and Trezzi 2004), which actually embodies the Grey-World and the White-Patch algorithms as special cases, and the Grey-Edge algorithm (van de Weijer et al. 2007a), which extends these algorithms to include higher-order statistics (e.g. edges). Other examples of color

constancy algorithms include methods that use knowledge acquired in a learning phase (Brainard and Freeman 1997; D'Zmura et al. 1995; Gehler et al. 2008; Gijsenij and Gevers 2007; van de Weijer et al. 2007b) and gamut-based methods (Finlayson et al. 2001; Finlayson et al. 2006; Forsyth 1990; Gijsenij et al. 2007). Gamut-based methods are very promising algorithms to achieve color constancy with high accuracy. Therefore, in this paper, we focus on this type of algorithm.

Gamut mapping algorithms are restricted to the use of pixel values to estimate the illuminant. They are based on the simplifying assumption that in real-world images, under a given light source, only a limited number of colors are observed. However, because the focus is solely on pixel values, additional information that is present in higher-order structures is ignored. Recent work by van de Weijer et al. (2007a) has shown that image derivatives can be used to improve color constancy. The Grey-World algorithm is extended to incorporate derivative information, resulting in the Grey-Edge algorithm.

In this paper, the aim is to provide a gamut mapping framework by extending the original gamut mapping approach to enable the incorporation of any linear filter output. The main focus will be on the local n-jet (i.e. derivative information). It has been shown that the n-jet completely describes the derivative structure of an image (Kass and Witkin 1987; Koenderink and van Doorn 1987). Hence, instead of using pixel values only, edges and higher-order structures are included in the n-jet description. As the n-jet based gamut mapping framework generates a number of algorithms, more stable solutions can be found by using the additional information simultaneously. Consequently, different ways of combining these algorithms are proposed to estimate the illuminant. The first approach is by using the feasible sets obtained by the different algorithms, and combining these sets into a new feasible set. The same estimator used for the original gamut mapping algorithm is applied to these new feasible sets which results in the final estimate. The second fusion approach considers the different gamut mapping algorithms as independent algorithms and the fusion is at the level of the outputs of the algorithms.

The paper is organized as follows. In Sects. 2 and 3, color constancy is defined and the original gamut mapping is explained, respectively. In Sect. 4, the gamut mapping is extended to incorporate the derivative information in terms of the n-jet. In Sect. 5, the different combination strategies are discussed. Finally, in Sects. 6 and 7, the experimental results and the conclusions are presented.

2 Color Constancy

In this paper, the aim of color constancy is to estimate the color of the light source. This estimate is used to transform the original image, so that it appears if it was taken under a canonical (often white) light source. To model this process, in this section, image formation and transformation are considered. Therefore, in Sect. 2.1, the image formation model is considered first. Then, in Sect. 2.2, the model to transform images from one light source to another is presented.

2.1 Reflection Model

An image $\mathbf{f}$ can be modeled under the assumption of Lambertian reflectance as follows:

$$\mathbf{f}(\mathbf{x}) = \int_\omega e(\lambda)\rho_k(\lambda)s(\mathbf{x},\lambda)d\lambda, \tag{1}$$

where $e(\lambda)$ is the color of the light source, $s(\mathbf{x},\lambda)$ is de surface reflectance and $\rho_k(\lambda)$ is the camera sensitivity function ($k \in \{R,G,B\}$). Further, ω and $\mathbf{x}$ are the visible spectrum and the spatial coordinates respectively.

In order to create a more realistic model, Shafer (Shafer 1985) proposes to add a "diffuse" light term to the model of (1). The diffuse light is considered to have low intensity and to be coming from all directions in an equal amount:

$$\mathbf{f}(\mathbf{x}) = \int_\omega e(\lambda)\rho_k(\lambda)s(\mathbf{x},\lambda)d\lambda + \int_\omega a(\lambda)\rho_k(\lambda), \tag{2}$$

where $a(\lambda)$ is the term that models the diffuse light. Using this equation, objects under daylight can be modeled more accurately, since daylight consists of both a point source (the sun) and diffuse light coming from the sky. However, the assumption that diffuse light is equal in all directions does not often hold in practice. A more realistic approximation is to consider the diffuse light to be dependent on the position in the image, according to:

$$\mathbf{f}(\mathbf{x}) = \int_\omega e(\lambda)\rho_k(\lambda)s(\mathbf{x},\lambda)d\lambda + \int_\omega \overline{a(\mathbf{x},\lambda)}\rho_k(\lambda), \tag{3}$$

where we assume the dependence of the position to be low-frequent, which is as indicated by the overline.

By assuming that the color of the light source $\mathbf{e}$ depends on the color of the light source $e(\lambda)$ as well as the camera sensitivity function $\rho_k(\lambda)$, color constancy is then equivalent to the estimation of $\mathbf{e}$:

$$\mathbf{e} = \int_\omega e(\lambda)\rho_k(\lambda)d\lambda. \tag{4}$$

Since, in general, only the image values of $\mathbf{f}$ are known, this is an under-constrained problem, and it therefore cannot be solved without further assumptions.

2.2 Diagonal Model

In this paper, color constancy is achieved by determining the color of the light source of an input image. However, in

many cases the color of the light source is of less importance than the appearance of the input image under a reference light. Therefore, the aim of most of the color constancy methods is to transform all colors of the input image, taken under an unknown light source, to colors as they appear under a reference light source. This transformation can be modeled by a diagonal mapping or *von Kries Model* (von Kries 1970). The diagonal mapping is given as follows:

$$\mathbf{f}^c = \mathcal{D}^{u,c}\mathbf{f}^u, \tag{5}$$

where $\mathbf{f}^u$ is the image taken under an unknown light source, $\mathbf{f}^c$ is the same image transformed, so it appears if it was taken under the reference light (called canonical illuminant), and $\mathcal{D}^{u,c}$ is a diagonal matrix which maps colors that are taken under an unknown light source u to their corresponding colors under the canonical illuminant c:

$$\begin{pmatrix} R^c \\ G^c \\ B^c \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^u \\ G^u \\ B^u \end{pmatrix}. \tag{6}$$

However, under some conditions, the diagonal model is too strict, and no solutions are found (this situation is called the null solution problem). This could be caused by saturated colors, the presence of surfaces that were not represented in the canonical gamut or scattering in the lens (veiling illumination), for instance. To overcome this, Finlayson et al. (2005) accounted for this shortcoming by adding an offset term to the diagonal model, resulting in the diagonal-offset model:

$$\begin{pmatrix} R^c \\ G^c \\ B^c \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^u \\ G^u \\ B^u \end{pmatrix} + \begin{pmatrix} o_1 \\ o_2 \\ o_3 \end{pmatrix}. \tag{7}$$

Deviations from the diagonal model are reflected in the offset term $(o_1, o_2, o_3)^T$. Ideally, this term will be zero, which is the case when the diagonal model is valid.

Interestingly, by means of the offset, the diagonal model also takes diffuse lighting into account as approximated by (2). To obtain position dependent diffuse lighting of (3), the following model is proposed, called local-diagonal-offset model:

$$\begin{pmatrix} R^c \\ G^c \\ B^c \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^u \\ G^u \\ B^u \end{pmatrix} + \begin{pmatrix} \overline{o_1(\mathbf{x})} \\ \overline{o_2(\mathbf{x})} \\ \overline{o_3(\mathbf{x})} \end{pmatrix}. \tag{8}$$

In conclusion, (8) is more robust against deviations from the diagonal model (e.g. saturated colors), diffuse light (assuming that the dependence of the position is low-frequent) and veiling illumination and hence will be used in the remainder of this paper.

3 Gamut Mapping

The gamut mapping algorithm has been introduced by Forsyth (1990). It is based on the assumption, that *in real-world images, for a given illuminant, one observes only a limited number of colors*. Consequently, any variations in the colors of an image (i.e. colors that are different from the colors that can be observed under a given illuminant) are caused by a deviation in the color of the light source. This limited set of colors that can occur under a given illuminant is called the *canonical gamut* $\mathcal{C}$, and it is found in a training phase by observing as many surfaces under one known light source (called the *canonical illuminant*) as possible.

In general, a gamut mapping algorithm takes as input an image taken under an unknown light source (i.e. an image of which the illuminant is to be estimated), along with the precomputed canonical gamut. Next, the algorithm consists of three important steps:

1. Estimate the gamut of the unknown light source by assuming that the colors in the input image are representative for the gamut of the unknown light source. So, all colors of the input image are collected in the input gamut $\mathcal{I}$.
2. Determine the set of *feasible mappings* $\mathcal{M}$, i.e. all mappings that can be applied to the gamut of the input image and that result in a gamut that lies completely within the canonical gamut. Under the assumption of the diagonal mapping, a unique mapping exists that converts the gamut of the unknown light source to the canonical gamut. However, since the gamut of the unknown light source is simply estimated by using the gamut of *one* input image, in practice several mappings are obtained. Every mapping i in the set $\mathcal{M}$ should take the input gamut completely inside the canonical gamut:

$$\mathcal{M}_i\mathcal{I} \in \mathcal{C}. \tag{9}$$

3. Apply an estimator to select one mapping from the set of feasible mappings. The selected mapping can be applied to the canonical illuminant to obtain an estimate of the unknown illuminant. The original method (Forsyth 1990) used the heuristic that the mapping resulting in the most colorful scene, i.e. the diagonal matrix with the largest trace, is the most suitable mapping. Alternatives are the average of the feasible set or a weighted average (Barnard 2000).

These are the basic steps of gamut mapping algorithms. Several extensions have been proposed. Difficulties in implementation are addressed in (Finalyson 1996; Finlayson and Hordley 2000), where it is shown that the gamut mapping algorithm can also be computed in chromaticity space $(\frac{R}{B}, \frac{G}{B})$. However, the performance of this $2D$ approach is a slightly lower than the performance of the $3D$ approach.

Dependency on the diagonal model is addressed in (Barnard 2000), where the canonical gamut is systematically enlarged by accounting for deviation of the diagonal model. In (Finlayson et al. 2005), the diagonal-offset model is introduced to account for diffuse light. Further, in (Finlayson et al. 2006), the problem of illuminant estimation is effectively reduced to the problem of *illuminant classification*, resulting in the gamut-constrained illuminant estimation algorithm. Finally, in (Finlayson and Xu 2003) an efficient implementation is introduced using convex programming.

4 Gamut Mapping using Derivative Structures

As discussed above, gamut mapping is based on the assumption that only a limited set of colors is observed under a certain illuminant. Multiple phenomena in nature (e.g. blurring) and imaging conditions (e.g. scaling) can cause the mixture of colors. Therefore, if two colors are observed under a certain illuminant, then also all colors in between could be observed under this illuminant, since the set of all possible colors which can be seen under a certain illuminant form a convex hull (i.e. gamut). In this paper, the gamut theory is extended by proving that the above is not only true for *image values* but also for every *linear combination of image values*. Hence, the correct estimate of an illuminant will also map every gamut which is constructed by a linear combination of image values back into the canonical gamut constructed with the same linear operation.

4.1 Gamut Mapping of Linear Combinations of Pixel Values

In the previous section, we have shown that the image values form a gamut, and that the transformations of the gamuts under illuminant changes follows the models give in (5)–(8). Here we will look at the image gamuts which are formed by a linear combination of image values.

Consider a set of image values:

$$\mathbf{F} = \{\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_n\} \tag{10}$$

were $\mathbf{f} = \{R, G, B\}$, and a image feature g which is a linear combination of image values $\mathbf{g} = \mathbf{w}^T\mathbf{F}$.

If we consider the *von Kries Model* the relation between the image values of an object taken under two different light sources is modeled by the diagonal model $\mathbf{f} = \mathbf{D}\mathbf{f}'$. Then for the feature $\mathbf{g}$ the following holds:

$$\begin{aligned}\mathbf{g} &= \mathbf{w}^T\mathbf{F} = w_1\mathbf{f}_1 + w_2\mathbf{f}_2 + \cdots + w_n\mathbf{f}_n \\ &= w_1\mathbf{D}\mathbf{f}'_1 + w_2\mathbf{D}\mathbf{f}'_2 + \cdots + w_n\mathbf{D}\mathbf{f}'_n \\ &= \mathbf{D}\big(w_1\mathbf{f}'_1 + w_2\mathbf{f}'_2 + \cdots + w_n\mathbf{f}'_n\big) \\ &= \mathbf{D}\big(\mathbf{w}^T\mathbf{F}'\big) = \mathbf{D}\mathbf{g}',\end{aligned} \tag{11}$$

proving that also for measurements $\mathbf{g}$ the diagonal models holds. The above is of importance, because it shows that gamut mapping can also be performed on all measurements $\mathbf{g}$ which are a linear combination of the image values $\mathbf{f}$.

Next, if we consider the diagonal-offset model given by $\mathbf{f} = \mathbf{D}\mathbf{f}' + \mathbf{o}$ then,

$$\begin{aligned}\mathbf{g} &= \mathbf{w}^T\mathbf{F} = w_1\mathbf{f}_1 + w_2\mathbf{f}_2 + \cdots + w_n\mathbf{f}_n \\ &= w_1\big(\mathbf{D}\mathbf{f}'_1 + \mathbf{o}\big) + \cdots + w_n\big(\mathbf{D}\mathbf{f}'_n + \mathbf{o}\big) \\ &= \mathbf{D}\big(\mathbf{w}^T\mathbf{F}'\big) + \left(\sum_{i=1}^{n} w_i\right)\mathbf{o} = \mathbf{D}\mathbf{g}' + \left(\sum_{i=1}^{n} w_i\right)\mathbf{o}.\end{aligned} \tag{12}$$

Hence, to estimate the illuminant change between $\mathbf{g}'$ an $\mathbf{g}$ we have to estimate both the diagonal matrix $\mathbf{D}$ and the offset $\mathbf{o}$. However, in the special case that $\sum_{i=1}^{n} w_i = 0$, the offset term $\mathbf{o}$ cancels out.

A similar reasoning can be applied to the local-diagonal-offset model of (8). In this case we have to ensure that all image values $\mathbf{f}_n$ which are linearly combined in $\mathbf{g}$ are taken from a local neighborhood where the offset $\overline{\mathbf{o}}$ can be considered constant. Hence, to perform gamut mapping under the local-diagonal-offset model the linear combination g has to satisfy two restrictions: (1) the weights w should sum up to zero, (2) the values $\mathbf{f}_n$ should come from a local neighborhood. Both these restrictions are satisfied by image derivative filters: the sum over the weights of filter is equal to zero, and since it is filter the values are taken from a local neighborhood. This makes image derivatives especially attractive for gamut mapping since, contrary to zero-order image value gamuts, they allow to estimate illuminant models under the more accurate local-diagonal-offset model.

In this paper, we investigate gamut mapping based on the statistical nature of images in terms of their derivative structure. The derivative structure of an image is described (in a complete sense) by means of the n-jet, see (Koenderink and van Doorn 1987) and (Kass and Witkin 1987). In this paper, we consider gamuts up to the second order structure, which is given by

$$\{\mathbf{f}, \mathbf{f}_x, \mathbf{f}_y, \mathbf{f}_{xx}, \mathbf{f}_{xy}, \mathbf{f}_{yy}\}, \tag{13}$$

where the derivatives are computed for image $\mathbf{f}$ by a convolution with a Gaussian at the scale of the derivative filter,

$$\mathbf{f} \otimes \frac{\partial}{\partial x} G^{\sigma} = \frac{\partial}{\partial x}\big(\mathbf{f} \otimes G^{\sigma}\big). \tag{14}$$

Since these derivative filters are all linear filters, it follows from (11) that the gamuts of the n-jet behave similarly under illuminant variations as a normal zero-order gamut.

4.2 n-jet Gamuts

The basic steps of the gamut mapping algorithm are identical when using derivative (n-jet) images. However, when us-

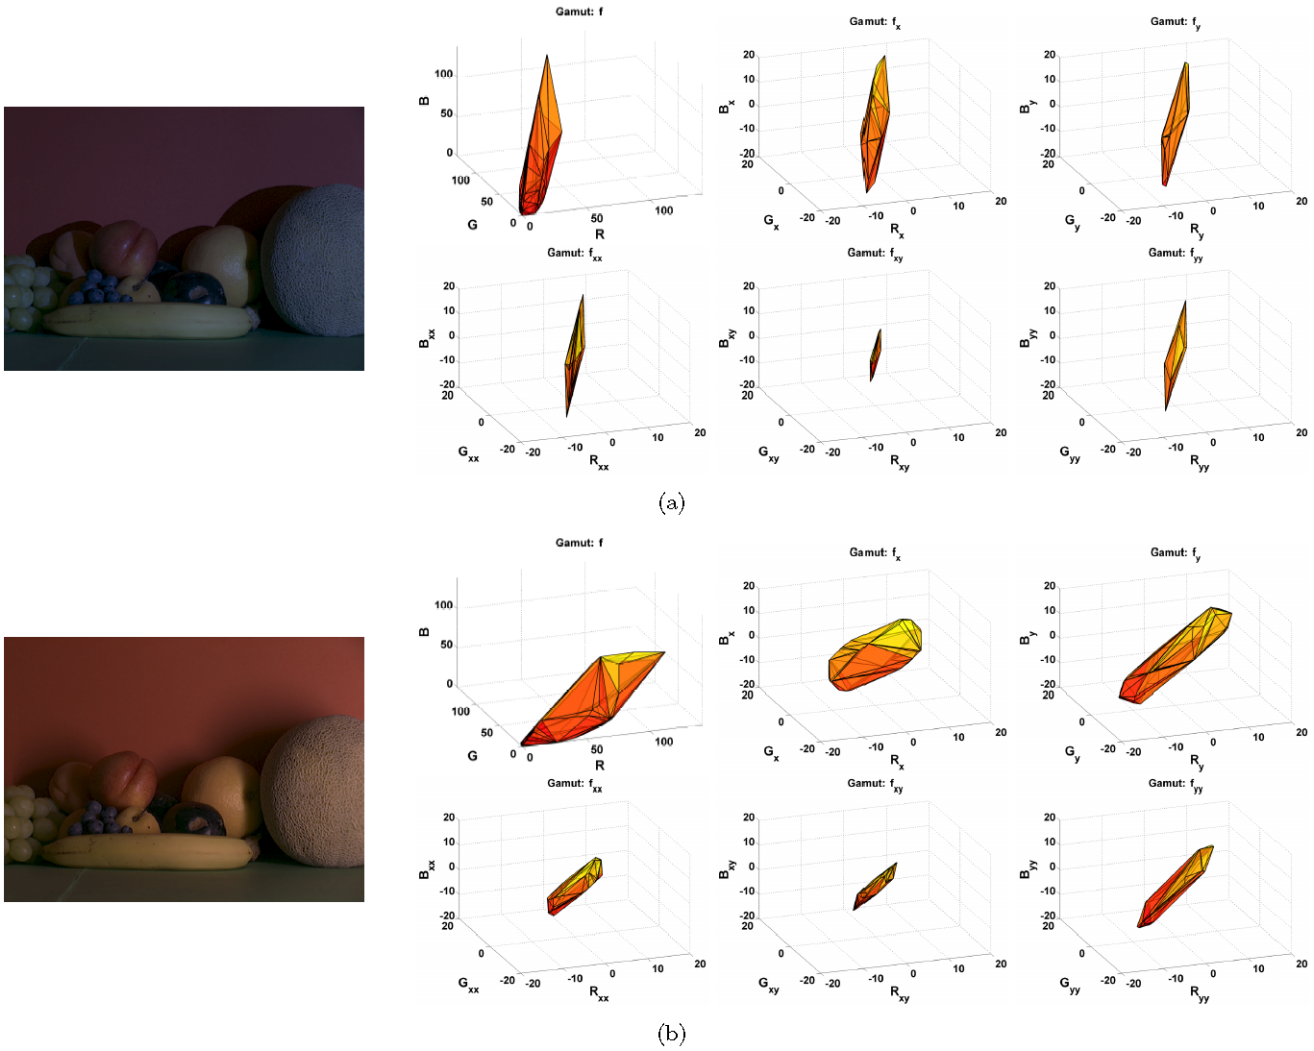

Fig. 1 Examples of the gamuts of the different n-jet images for a scene taken under two different light sources (images from Barnard et al. 2002). What is shown is the gamut of the corresponding image, using information that is present in either pixel values (**f**), edges ($\mathbf{f}_x$ and $\mathbf{f}_y$) or higher-order statistics ($\mathbf{f}_{xx}$, $\mathbf{f}_{xy}$ and $\mathbf{f}_{yy}$). Comparing the gamuts of the two images for one type of information (e.g. $\mathbf{f}_x$ of image (**a**) with $\mathbf{f}_x$ of image (**b**)) clearly shows the discriminative power of the different n-jets

ing derivatives, during the construction of the gamuts (both the canonical gamut and the input gamut), the values that are captured in the gamut are symmetric (e.g. if a transition from surface a to surface b is present, then the transition from surface b to surface a should also be included in the gamut). Further, note that the diagonal model can consist of strictly positive elements only. For the pixel-based gamut mapping this restriction is imposed naturally, but the first and second-order gamuts can contain negative as well as positive values. Hence, during implementation one should make sure that the diagonal mappings that are found contain strictly positive elements only. Further note that the complexity of the algorithm based on pixel and derivative information remains the same (and hence the difference in runtime can be neglected).

In Fig. 1, a few examples of gamuts of the different n-jet images are shown. From these images, it can be derived that the pixel-based gamut (i.e. the gamut of **f**), the edge-based gamuts (i.e. the gamuts of $\mathbf{f}_x$ and $\mathbf{f}_y$), as well as the gamuts using higher-order statistics (i.e. the gamuts of $\mathbf{f}_{xx}$, $\mathbf{f}_{xy}$ and $\mathbf{f}_{yy}$) are considerably different although they were computed from the same scene where the only difference is a change in the color of the light source.

5 Combination of Gamut Mapping Algorithms

It is beneficial to incorporate additional information into illuminant estimation (Hordley 2006). This can either be done by means of supplemental algorithms (Bianco et al. 2007; Schaefer et al. 2005) or by using higher-order statistics in combination with pixel values (Gijsenij and Gevers 2007). In this section, the goal is to exploit these two different ways

of combining derivative-based gamut mapping algorithms to provide additional information to estimate the illuminant.

The use of additional information introduces two mutually exclusive opportunities to increase the performance. First, the uncertainty of the estimates of the gamut mapping algorithm can be reduced. Second, the probability of finding the correct illuminant estimate can be increased. In general, the gamut mapping algorithm produces a set of illuminant estimates, called the *feasible set*. From this feasible set, one final illuminant estimate is selected using some method. If the size of the feasible set is large, then the possibility of selecting the wrong estimate is relatively large, i.e. the uncertainty of the final estimate is relatively high. On the other hand, a smaller feasible set results in a lower probability that the correct illuminant is contained inside this set. If multiple feasible sets can be found, which are all different from one another, by using the different n-jet images, then we can choose to either increase or decrease the size of the final feasible set. Intuitively, a smaller feasible set results in a more accurate illuminant estimate than a larger feasible set.

The first approach, proposed in Sect. 5.1, is by combining the feasible sets obtained by the different algorithms. The second fusion method, described in Sect. 5.2, considers the different gamut mapping algorithms as separate algorithms and combines the final estimates of the algorithms.

5.1 Combining Feasible Sets

Each gamut mapping algorithm produces a *feasible set* which contains all diagonal mappings that map the gamut of the input image inside the canonical gamut. Hence, the feasible set is a set of possible light sources. Since all gamut mapping algorithms produce such a set, these sets can be used for the combination, instead of selecting only one mapping per algorithm. Since each feasible set represents all illuminant estimates that are considered possible, a natural approach of combining the feasible sets is to consider only those estimates that are present in all feasible sets i.e. an intersection of the feasible sets. Another approach of combining the feasible sets is to consider every estimate that is present in all feasible sets, i.e. the union of the feasible sets:

$$\hat{\mathcal{M}}_{\text{intersect}} = \bigcap_i \mathcal{M}_i, \tag{15}$$

$$\hat{\mathcal{M}}_{\text{union}} = \bigcup_i \mathcal{M}_i, \tag{16}$$

where $\hat{\mathcal{M}}_{\text{intersect}}$ is the intersection of all feasible sets, $\hat{\mathcal{M}}_{\text{union}}$ is the union, and $\mathcal{M}_i$ is the feasible set produced by algorithm i. Then, on these combined feasible sets, an estimator is applied similar to step three of the gamut mapping algorithm.

5.2 Combining Algorithm Outputs

As a second possibility, the use of additional information is used in a later stage. Several methods can be considered. Bianco et al. (2007) propose a number of alternatives, of which a regular average of the outputs is the simplest combination strategy and the *No-N-Max*-method is the most effective. The latter method is a simple average of the outputs, excluding the N estimates that have the largest distance from the other estimates, where N is an adjustable parameter:

$$\hat{\mathbf{e}}_N = \frac{\sum_{i=1,\ldots,(N-n)} \mathbf{e}_i}{n - N}, \tag{17}$$

where $\mathbf{e}_i$ is the ith estimate, ranked according to the average distance to all other estimates. Further, $\hat{\mathbf{e}}$ is the result of the combination of the n algorithms, and N is the number of estimates that are excluded. Hence, $N = 0$ is equal to a simple average of all estimates.

6 Experiments

This section presents an empirical evaluation of the proposed algorithms. The first experiments are conducted on hyperspectral data, resulting in a systematic analysis of the derivative-based gamut mapping. Then, experiments on images taken under laboratory settings are conducted. Finally, the proposed methods are evaluated on real-world images. For all the experiments, we follow the implementation of (Finlayson and Xu 2003; Finlayson et al. 2005). We use the L_1 norm maximization.

6.1 Performance Measure

For all images in the data set (both hyperspectral and real-world images), the correct color of the light source $\mathbf{e}_l$ is known *a priori*. To measure how close the estimated illuminant resembles the true color of the light source, the angular error ϵ is used:

$$\epsilon = \cos^{-1}(\hat{\mathbf{e}}_l \cdot \hat{\mathbf{e}}_e), \tag{18}$$

where $\hat{\mathbf{e}}_l \cdot \hat{\mathbf{e}}_e$ is the dot product of the two normalized vectors representing the true color of the light source $\mathbf{e}_l$ and the estimated color of the light source $\mathbf{e}_e$. To provide more insight in the evaluation, the median as well as the mean angular error will be reported.

6.2 Hyperspectral Data

The first data set that is used consists of 1995 surface reflectance spectra, taken from several different sources, see Barnard et al. (2002). These surface reflectance spectra can

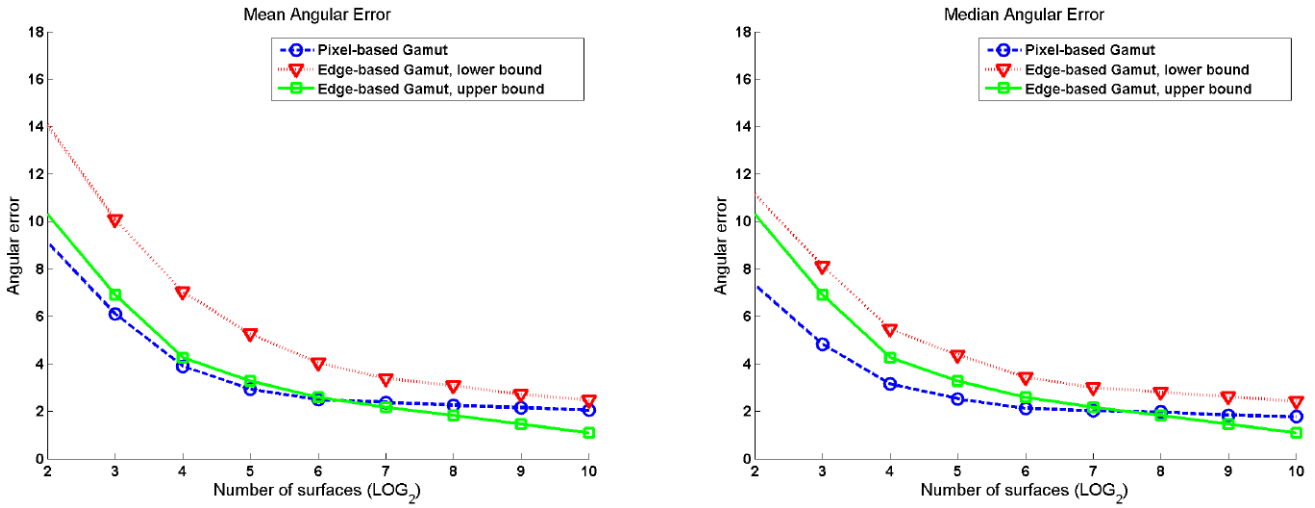

Fig. 2 Results of experiment 1. The *first figure* shows the mean angular error over 1000 scenes for every number of surfaces, the *second figure* shows the median angular error

be converted into (R, G, B)-values by combining them with an illuminant that is selected from a set of 287 illuminant spectra (also from Barnard et al. 2002), using (1).

Using this data set, a number of experiments are conducted. The first experiment is concerned with the performance of the gamut mapping algorithm as function of the number of surfaces (edges). The second experiment evaluates the robustness against color clipping, and the third experiment focuses on the robustness of the methods against deviations from the diagonal model as modeled by the *constant* offset in the diagonal model.

Note that the first data set consists of (R, G, B)-values, i.e. surface reflectance spectra in combination with an illuminant spectrum, instead of images. For these experiments, an edge is defined as the difference between two (R, G, B)-values and there is no distinction between edges in the x-direction and edges in the y-direction. 2nd-order statistics are not considered in this experiment. Also note that a scene with n different surfaces may contain several edges. However, the lower bound on the number of different edges in a scene with n surfaces is $n - 1$ edges, since every surface in a scene connects to at least one other surface (e.g. the background, which also is a surface in a scene). The upper bound on the number of edges is $\frac{1}{2}n(n-1)$ different edges, when each surface in the scene connects to every other surface. A transition from surface a to surface b is considered to be identical to a transitions from surface b to surface a.

Experiment 1. The first experiment is to compute the color constancy performance as a function of the number of surface (or edges). Surfaces are randomly selected from the database of surface reflectance spectra. To compare the performance of the pixel-based gamut mapping with the edge-based gamut mapping algorithm, two different numbers of edges are created for every scene: the lower bound and the upper bound. Along with a randomly selected illuminant spectrum, a data set is created containing n (R, G, B)-values (i.e. surfaces). Further, $n - 1$ (i.e. lower bound) and $\frac{1}{2}n(n-1)$ (i.e. upper bound) transitions are generated, and this process is repeated 1000 times, for $n = \{4, 8, 16, 32, 64, 128, 256, 512, 1024\}$.

In Fig. 2, results are shown for the pixel-based gamut mapping algorithm and the edge-based gamut mapping algorithm. The median and the mean angular error as function of the number of surfaces are shown. These graphs show that the pixel-based gamut mapping slightly outperforms the edge-based gamut mapping when the number of surfaces is small ($\log_2(n) < \pm 6$). However, for a larger number of surfaces (and hence a larger number of edges), the edge-based gamut mapping outperforms the pixel-based gamut mapping.

Experiment 2. The second experiment involves the simulated effect of color clipping. Would a real-world image be created using the generated (R, G, B)-values, then the pixel values are often bound to a maximum. This effect is called color clipping, and to analyze how this affects the performance, an experiment is performed. For this experiment, the number of surfaces is kept fixed throughout the test. For each iteration, $n = 8$ surfaces are generated by randomly selecting 8 reflectance spectra and 1 illuminant spectrum. These 8 surfaces are used to generate a scene with 8 edges, which is near the lower-bound on the number of edges. Since all used illuminants are normalized so that perfect white under that illuminant has a maximum value of 255 among the three channels, no values higher than 255 can occur. To simulate the effect of color clipping on the performance of the gamut mapping algorithms, several different clipping levels are simulated, discarding pixels with an R, G or B-value higher than that level. If a surface is

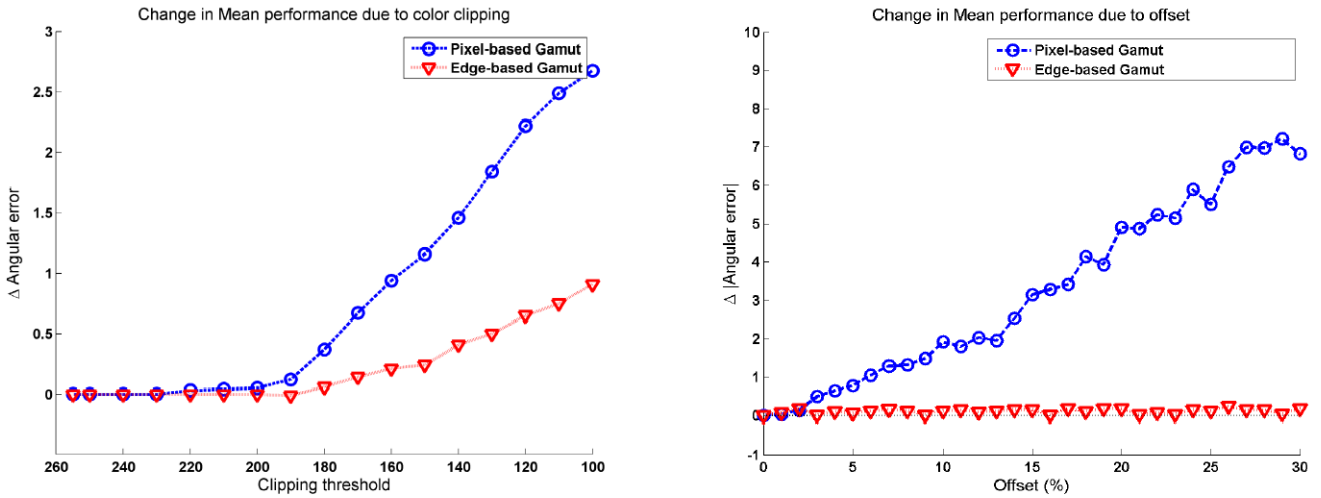

(a) Experiment 2 (b) Experiment 3

Fig. 3 Results of experiment 2 and 3. The figures show the change in mean performance of the two algorithms. In figure (**a**), it can be seen that the pixel-based gamut mapping is more affected by color clipping than the edge-based gamut mapping. Further, in figure (**b**), it can be seen that the error of the pixel-based gamut mapping algorithm increases as the offset increases, while the edge-based gamut mapping is not affected

clipped (and consequently discarded from the computation), then the corresponding edge is discarded as well.

Both methods suffer from the effect of color clipping, see Fig. 3(a). However, the error of the pixel-based gamut mapping increases faster than the error of the edge-based gamut mapping.

Experiment 3. The third experiment involves the robustness against deviations from the diagonal model, which is modeled by simulating an offset of the diagonal model, (7). Again, the number of surfaces is kept fixed throughout the test. For each iteration, $n = 8$ surfaces are generated by randomly selecting 8 reflectance spectra. These 8 surfaces are used to generate a scene with 8 edges, which is near the lower-bound on the number of edges. Next, instead of randomly selecting one illuminant spectra like the previous experiment, these spectra are combined with the canonical illuminant to form (R, G, B)-values. The diagonal-offset model is used to create the (R, G, B)-values under a different illuminant. The color of this illuminant is determined by applying the diagonal model to the canonical illuminant. The values of the elements of the diagonal matrix are randomly selected from the range $[0.5 \ldots 1.5]$, and the offset is gradually increased, ranging from 0% of the average pixel value in the scene (i.e. no offset) to 30% of the average pixel value. So, a new $(R_{\text{new}}, G_{\text{new}}, B_{\text{new}})$-value is generated by randomly selecting a value in the range of:

$$R_{\text{new}} = \left[\left(R - \frac{p}{100}x\right) \ldots \left(R + \frac{p}{100}x\right)\right], \tag{19}$$

$$G_{\text{new}} = \left[\left(G - \frac{p}{100}y\right) \ldots \left(G + \frac{p}{100}y\right)\right], \tag{20}$$

$$B_{\text{new}} = \left[\left(B - \frac{p}{100}z\right) \ldots \left(B + \frac{p}{100}z\right)\right], \tag{21}$$

where (x, y, z) is the average pixel value in the scene (i.e. the average (R, G, B)-values of the 8 surfaces) and the offset is $p\%$.

In Fig. 3(b), the results are shown, relative to the performance without the simulation of the disturbing effects. From this experiment, it can be observed that by adding an offset to the diagonal model, the performance of the edge-based gamut mapping algorithm is not affected, whereas the performance of the regular gamut mapping drops dramatically. The error increases linearly with the increasing offset.

To conclude, pixel-based gamut mapping performs slightly better than edge-based gamut mapping when the number of surfaces (and edges) is low. When the number of surfaces and edges start to increase, edge-based gamut mapping starts to perform better than pixel-based gamut mapping. However, in case of color clipping or deviations of the diagonal model, e.g. diffuse light or object reflections, the edge-based method outperforms the pixel-based gamut mapping. This indicates that the information that is used by the edge-based method is supplementary to the pixel values, and that for real-world images, where the shortcoming of the diagonal model is often the case, the additional edge-information could result in an increased performance of the gamut mapping algorithm.

Table 1 Angular errors for several versions of the gamut mapping on the SFU data set, containing 321 images. The images used for the computation of the canonical gamut are left out, leaving 290 images in the test set. When no solution is found, the white illuminant is taken as the estimate

Method	Mean ϵ	Median ϵ
$\mathcal{G}^{\sigma=3}(\mathbf{f})$	4.3°	2.6°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_x)$	5.1°	3.6°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_y)$	5.5°	3.5°
$\mathcal{G}^{\sigma=3}(\nabla\mathbf{f})$	4.5°	3.0°
$\mathcal{G}^{\sigma=5}(\mathbf{f}_{xx})$	5.5°	4.2°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_{xy})$	5.0°	3.2°
$\mathcal{G}^{\sigma=5}(\mathbf{f}_{yy})$	6.0°	3.5°
$\mathcal{G}^{\sigma=3}(\nabla\nabla\mathbf{f})$	5.1°	3.7°

6.3 Images under Laboratory Settings

The next experiments will be conducted on a set of real images, taken under laboratory settings (Barnard et al. 2002). The images in this data set are all indoor scenes. There are 31 scenes, taken under 11 different light sources, resulting in a total of 321 images. For more information on this set, we refer to (Barnard et al. 2002). The canonical gamuts were computed by taking one image from every scene. These images were then left out of the test set, so the algorithms are tested on 290 images. As a preprocessing step for $\mathcal{G}(\mathbf{f})$, i.e. the regular gamut mapping algorithm, a Gaussian averaging is used with the same scale that is used to compute the derivatives of $\mathbf{f}$. This procedure was found to improve the pixel-based gamut results. Further, to suppress high frequency noise amplification from the derivative operation, the image derivatives are computed with Gaussian derivative filters.

n-jet The results of the gamut mapping algorithm on the different n-jet images are shown in Table 1. The gradient of the image was computed by taking the correlation of the color channels into account, see (Di Zenzo 1986). It can be seen that the regular gamut mapping performs best on this data set, with a mean and median angular error of 4.3° and 2.6°, respectively. Using higher-order statistics, the best performance is obtained using the gradient of the image.

Combination As explained in Sect. 5, one of the main advantages of using the n-jet in combination with the gamut mapping is the possibility of information fusion. In Table 2, the results of different combination strategies are shown. In this table, the results for the Leave-N-out algorithm are obtained using the chromaticity-values of the estimates of the different gamut-mapping algorithms. Further, the combination of $\mathbf{f}$, $\mathbf{f}_x$, $\mathbf{f}_y$ and $\nabla\mathbf{f}$ denotes the 1-jet, and the 1-jet in combination with $\mathbf{f}_{xx}$, $\mathbf{f}_{xy}$, $\mathbf{f}_{yy}$ and $\nabla\nabla\mathbf{f}$ denotes the 2-jet.

Table 2 Angular errors for several fusion strategies using the different versions of the gamut mapping on the SFU data set, containing 321 images. The images used for the computation of the canonical gamut are left out, leaving 290 images in the test set. When no solution is found, the white illuminant is taken as the estimate

Method	Mean ϵ	Median ϵ
$\mathcal{G}^{\sigma=3}(\mathbf{f})$	4.3°	2.6°
Leave-N-out, $N=0$	4.0°	2.5°
Leave-N-out, $N=1$	3.9°	2.4°
Leave-N-out, $N=2$	3.9°	2.3°
Leave-N-out, $N=3$	3.9°	2.3°
Intersection: Complete 1-jet	4.2°	2.1°
Intersection: Complete 2-jet	4.2°	2.2°
Intersection: Only 1-order	5.0°	2.9°
Intersection: Only 2-order	4.7°	2.8°
Union: Complete 1-jet	4.8°	2.9°
Union: Complete 2-jet	5.0°	3.0°
Union: Only 1st-order	4.6°	3.4°
Union: Only 2nd-order	6.1°	4.6°

Finally, the combination of $\mathbf{f}_x$, $\mathbf{f}_y$ and $\nabla\mathbf{f}$ expresses the first-order, and $\mathbf{f}_{xx}$, $\mathbf{f}_{xy}$, $\mathbf{f}_{yy}$ and $\nabla\nabla\mathbf{f}$ the second-order versions.

Using the intersection to combine the feasible sets results in more accurate estimations that when using the union of the feasible sets. The best results are obtained when intersecting the feasible sets of the 1-jet, but results with the 2-jet are small. Notice that error when using only the 1st or 2nd-order are higher, indicating that the pixel values depict valuable information. Moreover, the median error of the combination of the 1-jet is much lower than when only pixel values are used, which indicates that the stability of the results is increased by using the combination instead of one single source of information.

The results for combination using the outputs, i.e. leave-N-out as shown in Table 2, are comparable to combination using the intersection of the feasible sets. Improved results are obtained when the 2 or 3 most distant estimates are excluded. In this case, the mean angular error is slightly lower than when using the intersection of the feasible sets, while the median angular error is slightly higher.

Comparison to the State-of-the-Art Now, the n-jet based gamut mapping algorithm is compared to other algorithms, which can be considered state-of-the-art. The results of the gamut mapping with a diagonal-offset model (Finlayson et al. 2005) are obtained by using the L_1 norm maximization of the diagonal transform, while minimizing the L_1 norm of the offset. In Table 3, the results are shown. It can be derived that when using the intersection of the feasible sets of $\mathbf{f}$ and $\mathbf{f}_x$ this will outperforms all other methods. The

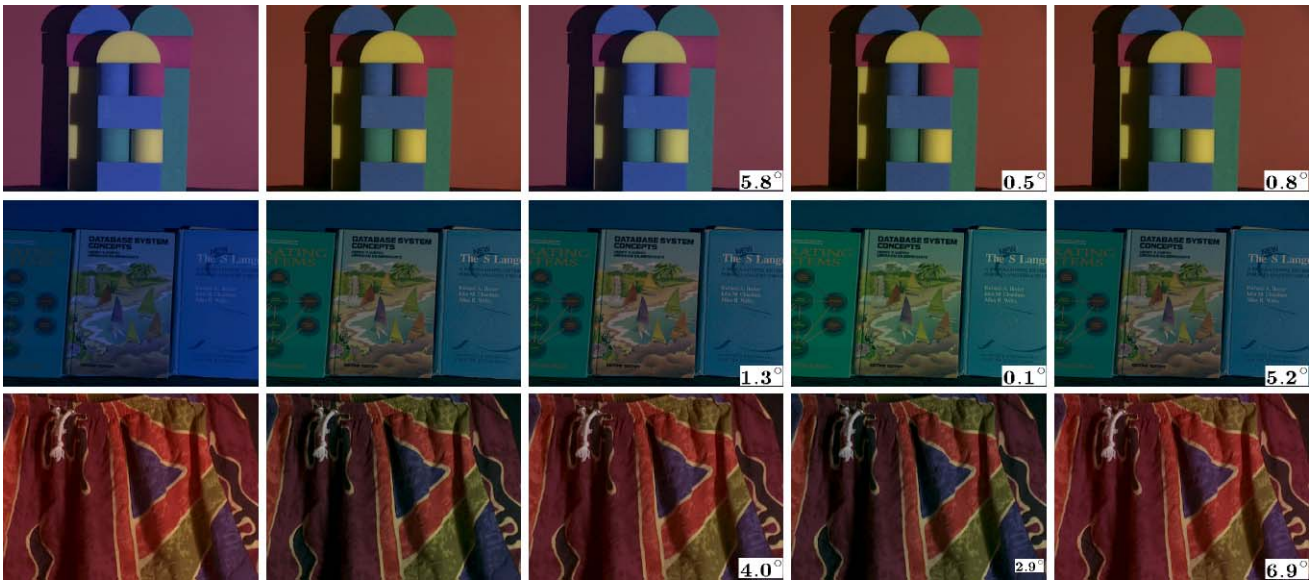

Fig. 4 Examples of the results of the pixel-based gamut mapping and the edge-based gamut mapping. The original image is shown on the *left*, followed by the result of perfect correction (using the ground truth), the result of the pixel-based gamut mapping (gamut mapping on $\mathbf{f}$), the result of the edge-based gamut mapping (gamut mapping on $\mathbf{f}_x$) and the result of the intersection of $\mathbf{f}$ and $\mathbf{f}_x$, respectively. The angular errors are displayed in the *bottom right corner* of the images

Table 3 Angular errors for several color constancy algorithms on the SFU data set. The images used for the computation of the canonical gamut are left out, leaving 290 images in the test set

Method	Mean ϵ	Median ϵ
Intersection: Complete 1-jet	4.2°	2.1°
Regular gamut with offset-model	4.7°	3.1°
Grey-World	9.8°	7.0°
White-Patch	9.2°	6.5°
Shades of Grey	6.3°	3.9°
Color by Correlation	9.9°	6.8°

method nearest in performance is the gamut mapping with a diagonal-offset model, with a mean and median angular error of 4.7° and 3.1°, respectively. Note that the results of the Grey-World,[1] White-Patch,[1] Shades-of-Grey[1] and Color by Correlation[2] are obtained using the software that is available on the web.

In conclusion, using the combined n-jet approach improves the performance of single gamut mapping algorithms on this data set, especially in terms of stability. Further, current state-of-the-art algorithms are outperformed using the n-jet based gamut mapping. Finally, examples of results are shown in Fig. 4. The first image shows the original unprocessed image. The second image shows how the image where the illuminant is perfectly estimated (for this purpose, the ground truth is used). The third, fourth and fifth image depict the result of the pixel-based gamut mapping (i.e. gamut mapping using $\mathbf{f}$), edge-based gamut mapping (i.g. gamut mapping using $\mathbf{f}_x$) and the combination of the two methods (i.e. $\mathcal{G}(\mathbf{f}) \cap \mathcal{G}(\mathbf{f}_x)$). The combination of the two methods results in an estimate that is almost close to the best performance that is obtained by either of the two algorithms. Obviously, the combination of the two methods does not improve results for every image, as can be derived from by example at the bottom.

6.4 Real-World Images

The last experiment is performed on a set of real-world images (Ciurea and Funt 2003). This set consists of images that are captured using a camera with a grey sphere mounted on top. There are 15 different scenes, both indoor and outdoor, resulting in a total of 11346 image (see Ciurea and Funt 2003 for more information). Note that the images in this set are not linear, as they were gamma-corrected (with unknown value for gamma) after capturing. Since the value of gamma is unknown, we did not correct for this before estimating the illuminants.

For this set, the canonical gamuts were computed using 14 of the 15 categories, and the algorithms are tested on the remaining category. This procedure was repeated 15 times, leaving every category out for testing once. Same as for the previous data set, for the regular gamut mapping

[1] http://lear.inrialpes.fr/people/vandeweijer/software.

[2] http://kobus.ca/research/programs/colour_constancy.

Table 4 Angular errors for the real-world data set containing 11346 images, both indoor and outdoor. Note that when no solution could be found, then a white illuminant was taken as estimate

Method	Mean ϵ	Median ϵ
$\mathcal{G}^{\sigma=5}(\mathbf{f})$	7.1°	5.8°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_x)$	7.1°	5.9°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_y)$	7.0°	6.0°
$\mathcal{G}^{\sigma=3}(\nabla\mathbf{f})$	6.8°	5.8°
$\mathcal{G}^{\sigma=3}(\mathbf{f}_{xx})$	7.6°	6.1°
$\mathcal{G}^{\sigma=1}(\mathbf{f}_{xy})$	7.5°	5.8°
$\mathcal{G}^{\sigma=5}(\mathbf{f}_{yy})$	7.0°	6.1°
$\mathcal{G}^{\sigma=5}(\nabla\nabla\mathbf{f})$	7.0°	6.3°

algorithm, i.e. $\mathcal{G}(\mathbf{f})$, a Gaussian averaging is used as pre-processing step, and the derivative operation is computed with Gaussian derivative filters.

n-jet The results of the gamut mapping algorithm on the different n-jet images for this data set are shown in Table 4. It can be seen that the median angular error of the gamut mapping applied to 0th-order data (i.e. pixel values) is similar to the performance of the gamut mapping applied to 1st-order data (i.e. edges). However, the mean angular error when using edges is lower than when using pixel values. This indicates that it is better to use edges than pixel values for real-world images. The performance of the gamut mapping algorithm using 2$^{\text{nd}}$-order information is significantly worse that both pixel values and edges.

Combination Combining the results of the n-jet approach slightly improves the performance, see Table 5.

For this data set, the same trend as for the previous data set is observed: the combination of the different methods results in more stable solutions than when using a single algorithm, and the intersection of the feasible sets is better than the union. The best performance, in terms of mean as well as median angular error, is obtained by simply averaging the estimates of the different n-jet based methods.

Comparison Finally, results of a few other color constancy algorithms are shown in Table 6. The n-jet based gamut mapping outperforms the regular gamut mapping with the diagonal-offset model on this data set, even when the n-jet information is not used. This indicates that the diagonal model does not fail often for the images that are in this data set.

For this real-world data set, the same trend as for the images under laboratory settings can be observed, namely that the additional information that is provided by using the n-jet improves the performance and stability of the gamut

Table 5 Angular errors for several combination strategies using the different versions of the gamut mapping on the real-world data set containing 11346 images, both indoor and outdoor. Note that when no solution could be found, then a white illuminant was taken as estimate

Method	Mean ϵ	Median ϵ
$\mathcal{G}^{\sigma=5}(\mathbf{f}_y)$	6.9°	5.6°
Leave-N-out, $N=0$	6.5°	5.5°
Leave-N-out, $N=1$	6.5°	5.5°
Leave-N-out, $N=2$	6.5°	5.5°
Leave-N-out, $N=3$	6.5°	5.6°
Intersection: Complete 1-jet	6.9°	5.8°
Intersection: Complete 2-jet	6.9°	5.9°
Intersection: Only 1st-order	6.8°	5.7°
Intersection: Only 2nd-order	7.1°	6.2°
Union: Complete 1-jet	7.1°	5.8°
Union: Complete 2-jet	7.3°	5.9°
Union: Only 1st-order	6.9°	5.8°
Union: Only 2nd-order	7.7°	6.5°

Table 6 Angular errors for several algorithms on the real-world data set

Method	Mean	Median
Leave-N-out	6.5°	5.5°
Regular gamut with offset-model	7.2°	5.7°
Grey-World	7.9°	7.0°
White-Patch	7.1°	6.7°
Color-by-Correlation	8.1°	6.5°

mapping algorithm. The intersection of the feasible sets provides better results than the union of the feasible sets. Further, results, for this data set, are also better than when using the Leave-N-out strategy. Finally, a few example images of this data set are shown in Fig. 5. Again the unprocessed and the perfectly corrected images are show first, but now they are followed by the result of the pixel-based gamut mapping (i.e. gamut mapping using $\mathbf{f}$), the result of the intersection of the gamut mapping on the edges in the x and y-direction (i.e. $\mathcal{G}(\mathbf{f}_x) \cap \mathcal{G}(\mathbf{f}_y)$) and the result of the union of the same two algorithms (i.e. $\mathcal{G}(\mathbf{f}_x) \cup \mathcal{G}(\mathbf{f}_y)$). On average, the intersection provides the best results, but for some images the union provided better results. In general, it can be seen that the combination of the different algorithms result in more stable solutions, i.e. the variation of the performance between the images in one data set is less for the combination of algorithms than for the single algorithms.

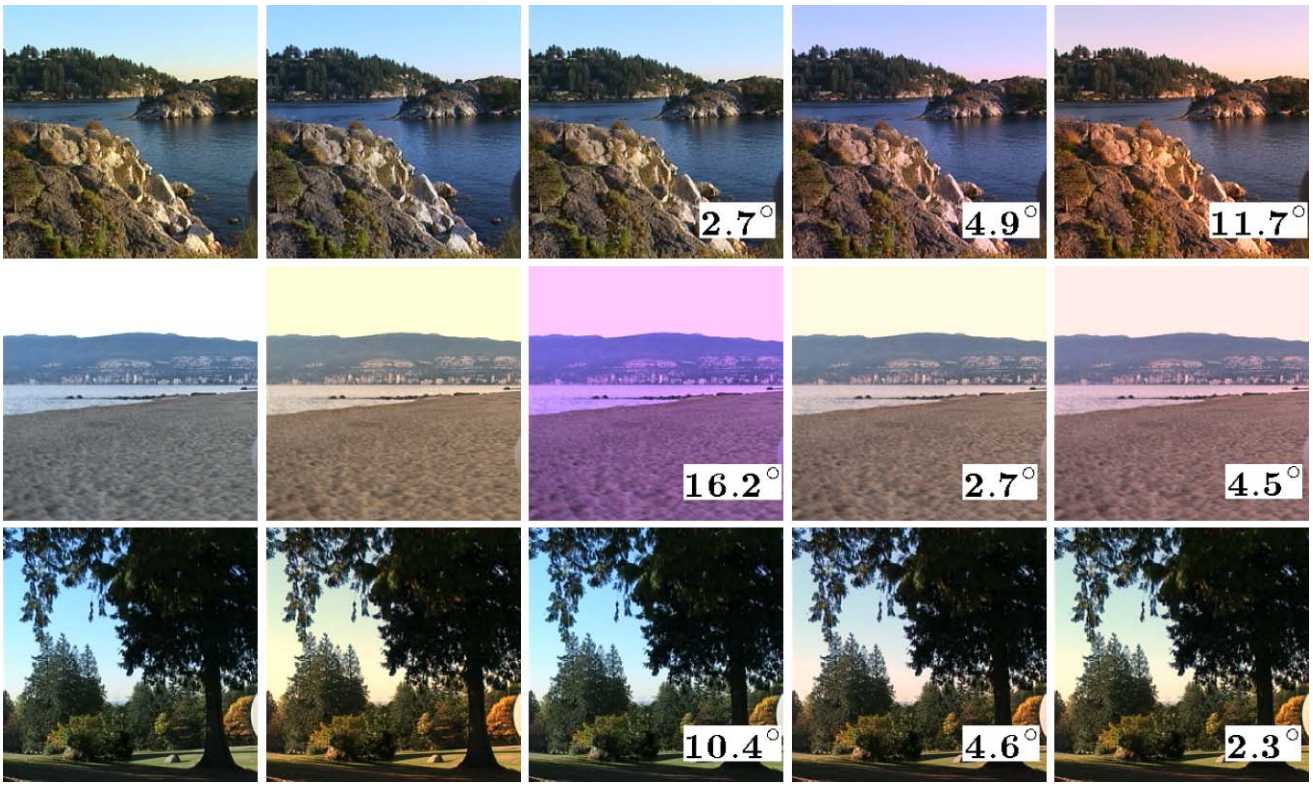

Fig. 5 Examples of the results of the pixel-based gamut mapping and the edge-based gamut mapping. The original image is shown on the *left*, followed by the result of perfect correction (using the ground truth), the result of the pixel-based gamut mapping (gamut mapping on $\mathbf{f}_x$), the result of the intersection of $\mathbf{f}_x$ and $\mathbf{f}_y$ and the result of the union of $\mathbf{f}_x$ and $\mathbf{f}_y$, respectively. The angular errors are displayed in the *bottom right corner* of the images

7 Conclusion

In this paper, gamut mapping has been extended to incorporate the statistical nature of images. It has been analytically shown that the proposed gamut mapping framework is able to include any linear filter output. However, the main focus was on the local n-jet describing the derivative structure of an image. As the n-jet based gamut mapping framework generates a number of algorithms, more information becomes available and different ways of combining are proposed to increase the performance and stability of the gamut mapping algorithm.

Experiments on large scale data sets of hyperspectral, laboratory and real-world scenes showed that (1) in case of deviations of the diagonal model, e.g. diffuse light or object reflections, the edge-based method outperformed the pixel-based gamut mapping, (2) state-of-the-art algorithms are outperformed by the n-jet based gamut mapping, (3) the combination of the different n-jet based gamut mappings provide more stable solutions, and (4) the fusion strategy based on intersection provides better color constancy results than the strategy based on the union of the feasible sets.

Acknowledgements We thank the anonymous reviewers for their valuable comments. This work is partially supported by the Consolider-Ingenio 2010 CSD2007-00018 of Spanish Ministery of Science and the Ramon y Cajal Program.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

Barnard, K. (2000). Improvements to gamut mapping colour constancy algorithms. In *European conference on computer vision* (pp. 390–403).

Barnard, K., Martin, L., Funt, B., & Coath, A. (2002). A data set for color research. *Color Research and Application*, *27*(3), 147–151.

Bianco, S., Gasparini, F., & Schettini, R. (2007). Combining strategies for white balance. In *Digital photography III*. IS&T.

Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America A*, *14*, 1393–1411.

Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of the Franklin Institute*, *310*(1), 1–26.

Ciurea, F., & Funt, B. V. (2003). A large image database for color constancy research. In *Proceedings of the eleventh color imaging conference* (pp. 160–164). IS&T—The Society for Imaging Science and Technology.

Delahunt, P. B., & Brainard, D. H. (2004). Does human color constancy incorporate the statistical regularity of natural daylight? *Journal of Vision*, *4*(2), 57–81.

Di Zenzo, S. (1986). A note on the gradient of a multi-image. *Computer Vision, Graphics and Image Processing*, *33*(1), 116–125.

D'Zmura, M., Iverson, G., & Singer, B. (1995). Probabilistic color constancy. In *Geometric representations of perceptual phenomena* (pp. 187–202). Lawrence Erlbaum Associates.

Fairchild, M. D. (2005). *Wiley-IS&T series in imaging science and technology*. *Color appearance models* (2nd ed.). Chichester: Wiley. ISBN 0-470-01216-1.

Finalyson, G. D. (1996). Color in perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *18*(10), 1034–1038.

Finlayson, G. D., & Hordley, S. D. (2000). Improving gamut mapping color constancy. *IEEE Transactions on Image Processing*, *9*(10), 1774–1783.

Finlayson, G. D., & Trezzi, E. (2004). Shades of gray and colour constancy. In *Proceedings of the twelfth color imaging conference* (pp. 37–41). IS&T—The Society for Imaging Science and Technology.

Finlayson, G. D., & Xu, R. (2003). Convex programming color constancy. In *IEEE workshop on color and photometric methods in computer vision, in conjunction with ICCV'03* (pp. 1–8).

Finlayson, G. D., Hordley, S. D., & Hubel, P. M. (2001). Color by correlation: a simple, unifying framework for color constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *23*(11), 1209–1221.

Finlayson, G. D., Hordley, S. D., & Xu, R. (2005). Convex programming colour constancy with a diagonal-offset model. In *Proceedings of IEEE international conference on image processing* (pp. 948–951).

Finlayson, G. D., Hordley, S. D., & Tastl, I. (2006). Gamut constrained illuminant estimation. *International Journal of Computer Vision*, *67*(1), 93–109.

Forsyth, D. A. (1990). A novel algorithm for color constancy. *International Journal of Computer Vision*, *5*(1), 5–36.

Foster, D. H., Amano, K., & Nascimento, S. M. C. (2006). Color constancy in natural scenes explained by global image statistics. *Visual Neuroscience*, *23*(3–4), 341–349.

Gehler, P. V., Rother, C., Blake, A., Minka, T., & Sharp, T. (2008). Bayesian color constancy revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1–8).

Gevers, T., & Smeulders, A. W. M. (2000). Pictoseek: combining color and shape invariant features for image retrieval. *IEEE Transactions on Image Processing*, *9*(1), 102–119.

Gijsenij, A., & Gevers, T. (2007). Color constancy using natural image statistics. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1–8).

Gijsenij, A., Gevers, T., & van de Weijer, J. (2007). Color constancy by derivative-based gamut mapping. In *Photometric analysis for computer vision (PACV'07), in conjunction with ICCV'07*, Rio de Janeiro, Brazil, October 2007.

Hordley, S. D. (2006). Scene illuminant estimation: past, present, and future. *Color Research and Application*, *31*(4), 303–314.

Kass, M., & Witkin, A. (1987). Analyzing oriented patterns. *Computer Vision, Graphics and Image Processing*, *37*(3), 362–385.

Koenderink, J. J., & van Doorn, A. J. (1987). Representation of local geometry in the visual system. *Biological Cybernetics*, *55*(6), 367–375.

Land, E. H. (1977). The retinex theory of color vision. *Scientific American*, *237*(6), 108–128.

Schaefer, G., Hordley, S., & Finlayson, G. (2005). A combined physical and statistical approach to colour constancy. In *Proceedings of the IEEE computer society conference on computer vision and pattern recognition* (pp. 148–153). Washington: IEEE Computer Society.

Shafer, S. A. (1985). Using color to separate reflection components. *Color Research and Application*, *10*(4), 210–218.

van de Weijer, J., Gevers, T., & Gijsenij, A. (2007a). Edge-based color constancy. *IEEE Transactions on Image Processing*, *16*(9), 2207–2214.

van de Weijer, J., Schmid, C., & Verbeek, J. J. (2007b). Using high-level visual information for color constancy. In *Proceedings of the international conference on computer vision*, Rio de Janeiro, Brazil.

von Kries, J. (1970). Influence of adaptation on the effects produced by luminous stimuli. In D. L. MacAdam (Ed.), *Sources of color vision* (pp. 109–119). Cambridge: MIT Press.

Yang, J., Stiefelhagen, R., Meier, U., & Waibel, A. (1998). Visual tracking for multimodal human computer interaction. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 140–147).