Exploring Colour Constancy Solutions

F. Tous, M. Vanrell and R. Baldrich
Dept. d'Informàtica – Computer Vision Center, Universitat Autònoma de Barcelona
Bellaterra, Barcelona, Spain

Abstract

The aim of this paper is to explore the space of diagonal colour constancy solutions. In gamut mapping approaches, finding the illuminant of a scene implies to find the set of feasible maps and afterwards to apply certain decision criterion to select a proper solution. This last step has been usually based on a heuristic computation over the feasible set. However an analysis on how are the solutions of this feasible set is not known by the authors. This is the essential contribution of this paper, since we explore on a reduced version of the feasible set some specific properties of the solutions. Criteria such as, maximum volume, feasible set average, maximum area on chromaticity plane or grey world solutions have been explored, and this works conclude that this usual criteria do not always assure finding optimal solutions, and therefore, further work remains to be done in this sense. Finally, we outline that some criteria related to the position of the optimal mapped image on the chromaticity plane should be taken into account.

Introduction

Colour constancy is the ability of the human visual system to build internal colour representations without the effects of the scene illuminant. A complete answer on how human visual system reaches this capability has not given yet. However, there are several theories on how it would be.[1-3]

In computer vision, colour constancy is a main focus of interest, since a good efficiency on automatic image understanding requires a stable, i.e. illuminant invariant, representation of colour images. Hence, a wide range of colour constancy methods have been proposed in the literature.[5-11] The performance ratios of these methods vary depending on how it is measured,[4] but still does not exist a colour constancy method that performs exactly on all sort of images and under weak assumptions. In this framework, methods based on the gamut mapping approach[5] can be regarded as the best in some of the performance rankings. In this work we will focus on the properties of the solutions given by this approach. The analysis will be done on synthetic images since error measurements are straightforward and free of the influence of sensor errors.

Gamut mapping methods do not estimate the scene illuminant of an image, but they directly estimate the transform from this unknown illuminant to a canonical one. The transform of an illuminant change is usually modelled by a linear diagonal model,[12]

$$(R^C, G^C, B^C) = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} R^U \\ G^U \\ B^U \end{pmatrix} \tag{1}$$

Then, to solve the colour constancy problem implies to estimate three parameters, named α, β and γ, which are the components of the diagonal matrix that maps a colour under an unknown illuminant, (R^U G^U B^U), to a colour under a canonical illuminant, (R^C G^C B^C). It is obvious to note that we are dealing with an under-constrained problem since we only know (R^U G^U B^U) from every pixel of a given image. To deal with this problem, gamut mapping methods usually act in two steps, firstly, computing the set of all feasible maps, and secondly, selecting the estimated map.

The first step finds the limits, within the space of all possible maps for the set of feasible solutions. That is the space of all possible (α, β, γ) solutions. We will call it the space of maps. These are the solutions that assure to map the image gamut inside the canonical gamut, henceforth, the feasible set.

The second step involves applying a heuristic criterion to select one of the feasible solutions. Different heuristics have been proposed for this selecting criterion, such as the map which maximizes the volume of the mapped image, the average of the feasible set, or the maximum of different norms: L_1 or L_2.[13]

Whilst we can find several works expending efforts to improve the first step of gamut mapping approaches, not so many efforts have been expended in discussing about the second step. The performance of different selection criteria has been recently explored in Ref. [13]. Because we think it is worth to go further on this point, we propose an empirical exploration of the feasible set of solutions given by a gamut mapping approach. It will allow exploring not the performance but the adequacy of these selection criteria.

The empirical approach we propose in this paper is based on a sampling process of the continuous set of feasible solutions. By computing the error for every specific solution we try to extrapolate some conclusions about the properties of the optimal solutions. In this work we will regard the minimum angle error as the optimality criterion for the solutions. Therefore the norm of the best solutions will not be considered. To explore the feasible set we have two options: to define a sampling procedure on the (α, β, γ) space, in order to get a subset of the feasible solutions obtained by a general gamut mapping algorithm[5]; or directly to build a heuristically reduced set of feasible solutions. We have worked on the second option and we have built a constraint-satisfaction search tree based on a specific assumption about the solutions.

We have organized the paper as follows. Firstly, we explain the search method we have used to build a sampled version of the feasible set. Secondly, we explore within the sampled version of the feasible set showing the error solution with respect to different criteria. Finally we derive some conclusions.

Sampling the Feasible Set

The goal of this section is to define a procedure to directly build a reduced set of feasible solutions. To do this, we propose to work on a search tree that gives a subset of maps accomplishing the constraint of being within the feasible set.

In order to have a better comprehension of the map set, instead of directly work on the space of all (α , β , γ) maps, we will work on the space of all possible correspondences between image surfaces and canonical surfaces, i.e., the latter corresponds to a set of selected surfaces that sample a wide range of the gamut of a canonical illuminant. Hence, for a given correspondence we compute a diagonal matrix that approximates the transform of the selected image surfaces to the canonical surfaces.

Thus, from a given image, *I*, acquired under an unknown illuminant, *U*, we extract the (R, G, B) values of *n* surfaces,

$$S^U(I) = \{S_1^U, S_2^U, \ldots, S_n^U\},$$

and given a canonical set of surfaces, denoted as

$$S^C = \{S_1^C, S_2^C, \ldots, S_k^C\},$$

we can build the set of all possible correspondences between S^U and S^C, which can be expressed as

$$Corr = \{S_1^U = S_{p_1}^C, S_2^U = S_{p_2}^C, \ldots, S_n^U = S_{p_n}^C;\ p_i = 1,..,k\}$$

where $\#Corr = k^n$, and means that we presume that surface S_i^U corresponds to surface

$$S_{p_i}^C$$

when it is seen under the canonical illuminant. Then, for a given set of matchings we can compute *n* diagonal transforms, one for each pair of two surfaces, and their mean can be used to estimate a diagonal map for this correspondence. On this space of all possible pairs between image and canonical surfaces we have k^n possible solutions, which implies an exponential complexity and therefore it is hard to be computed.

Building the set of all these solutions implies to do a search on a tree with *n* levels of depth, being *k* the branching factor, and where each complete path (from the root to a leaf node) of the tree represents a concrete solution. This will be a constraint satisfaction search algorithm since all the maps being out of the feasible set will be removed. Although, this constraint is sometimes very effective, still the number of solutions can be too large; e.g. if we work with images of 10 different coloured surfaces and with about 35 canonical surfaces we need to explore a tree of 35^{10} paths.

In order to reduce the size of this tree, we need to reduce the size of the $S^U(I)$ and S^C sets. To this end, we propose some assumptions on these two sets:
Assumption 1 (*existence of significant surfaces): An image can be approximately represented by a fixed number of significant surfaces, namely r, whose gamut nearly covers the entire image gamut.*

Assumption 2 *(relaxed grey-world): The average of the significant surfaces of an image is close to grey.*

Assumption 3 *(colour structure invariance): The most likely correspondences between image surfaces and canonical surfaces are those which maintain the colour structure of the image gamut. That is, relative positions of colour surfaces within the gamut should not change.*

Apart from reducing the set of canonical and image surfaces, an important goal of these assumptions is also to ensure that an important amount of optimal maps will be selected with the sampling process, since the final goal is to explore the nature of these optimal solutions.

Above assumptions introduce important changes on the sampling process. The branching factor is reduced thanks to assumption 3, and assumption 1 reduces the depth of the tree to *r*. These assumptions are introduced into the sampling algorithm with the following steps:

1. Selection of the *r* significant surfaces from $S^U(I)$ to $S^{Sig}(I)$.
2. Mapping surfaces of $S^{Sig}(I)$ with a grey world transform, it is denoted as $S^{GW}(I)$.
3. For each surface, $i: 1..r$, of $S^{GW}(I)$ we select the *m* nearest neighbours surfaces from the canonical surfaces, S^C, denoted as,

   $$S^{NN_i}$$

   (see figure 1).
4. Computing the set of all possible correspondences between each

   $$S_i^{Sig}(I) \text{ with its corresponding } S^{NN_i},$$

it is given by

$$RCorr =$$

$$S_1^{Sig} = S_{p_1}^{NN_1}, S_2^{Sig} = S_{p_2}^{NN_2}, \ldots, S_r^{Sig} = S_{p_r}^{NN_r};\ p_i = 1,..,m$$

where $\#RCorr = m^r$ that can be significantly reduced depending on the number of selected values.

5. For each element of *RCorr*, the corresponding (α , β , γ) map is computed.

Some considerations have to be explained regarding the step 1 of the given algorithm. The goal of this step is to reduce the number of image surfaces. Based on assumption 1 we do the following process:

1. Computing the extreme point of the convex hull of the image surfaces, (solid points of figure 2).

2. For each point of the previous step, the sum of distance to the rest of points is computed.
3. The r points with greatest values of the previous point are selected as significant surfaces, (Significant image surfaces in figure 2), obtaining $S^{Sig}(I)$.

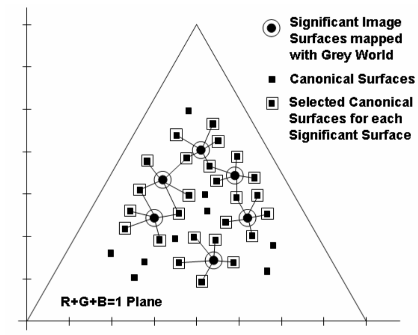

Figure 1. r-nearest neighbours of canonical surfaces for significant surfaces (m=5, r=6).

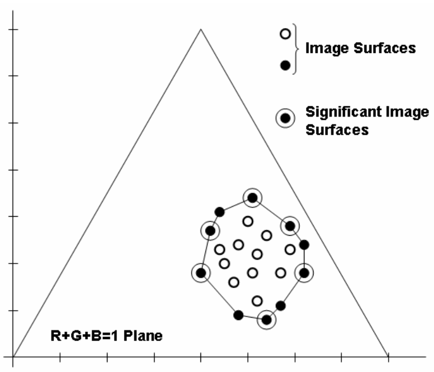

Figure 2. Selected Significant Image Surfaces as the farthest within the surfaces in the convex hull of the image surfaces (r=6).

The introduction of assumptions helps us to reduce the number of samples of the feasible set. In the next section we compute the angular error obtained when applying each one of the computed sampling, and subsequently, these solutions will be explored to analyse the behaviour of those with minimum error.

Exploration

To explore how the best solutions act within the feasible set, we empirically work on their observation. Our work has been carried out with synthesized data, creating a set of canonical surfaces and images consisting of random Munsell chips. For the canonical surfaces we have selected 35 representative surfaces from the Munsell chips and the macbeth color chart. As canonical illuminant we have synthesized a planckian illuminant with CCT=6500K, and a gausian narrow-band sensor has been built, with centers in 450, 540 and 610 nm. With this configuration, the RGB values of the canonical surfaces have been built, i.e., we have 35 RGB values representing the canonical surfaces under the canonical illuminant.

We have also generated several synthetic images, which are a set of randomly selected Munsell surfaces, built from their reflectances under a random illuminant, chosen from a set of 11 different illuminants. For these images we will compute the colour constancy error for each of the solutions generated with the reduced matching tree. In our experiments this tree has been computed with parameters, $m=5$ an $r=6$, that means we have got 6 significant surfaces of an image and these have been put in correspondence with the 5 nearest surfaces of the canonical gamut for each image surface. This implies to explore a set of $m^r=5^6=15625$ possible maps.

In this work, we have used as optimality criterion the angular error between the recovered RGB color vectors for an estimated solution and the true RGB color vector under the canonical illuminant, considering we know the image illuminant, the RMS error of the angular error for the r surfaces is used.[4]

Since the angle between two RGB color vectors only rate the error estimating the chromaticity of the surfaces, we will only focus on the recovery of the chromaticity properties of the image surfaces as they could be seen under the canonical illuminant, and not their intensity. The analysis of this error over a large range of solutions within a large set of images, has taken us to conclude some interesting properties on the map selection criteria.

In some gamut mapping methods, once the feasible set is computed, the map that gives the gamut with the largest volume is usually selected as the optimal map, i.e., the volume is used as an estimator for the best map. Because in our study we only work with the angular error, we have used the area of the convex hull of the mapped image in chromaticity coordinates as a possible measure of optimality.

In figure 3 we can see some plots of the angular error versus the area of the mapped image gamut for three different images. The coordinates of each point in this plot represent the properties of a specific map of the generated feasible set, built with the proposed sampling algorithm.

Observing these plots it can be inferred that the optimal solution has a fixed area, i.e., there is only one area that gives the best solution, however, for the cases we have studied, it rarely is the maximum area. This leads us to consider other measures apart from the area.

In figure 4 we see the scatter plot corresponding to the same images of figure 3, presenting in this case the angular error versus the volume of images transformed by the generated maps. Again, in this case the explored images rarely present a maximum volume for the optimal solutions. In the same plots, bright dots represent the average map computed from these feasible sets, in these three cases the average map presents an approximate value of 6.

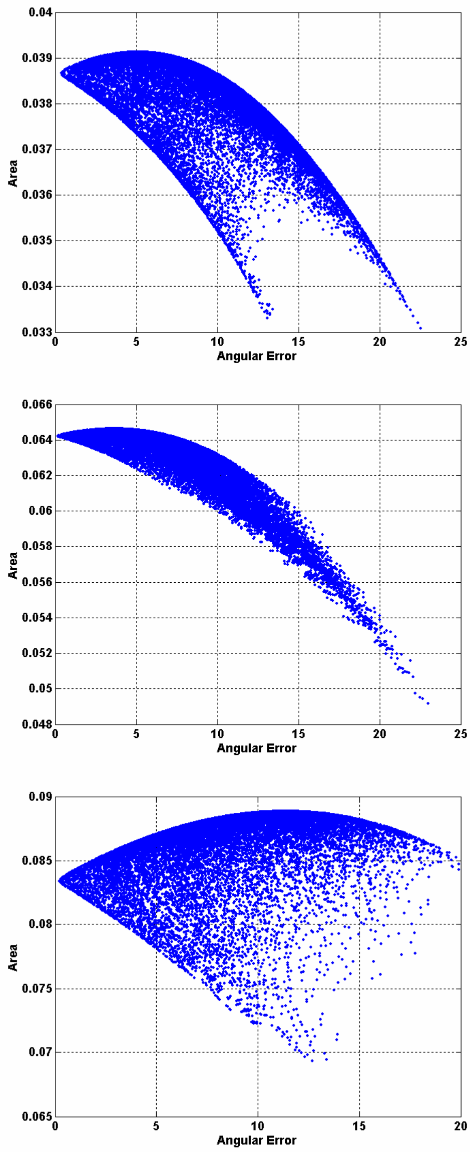

Figure 3. Scatter plots of the angular error versus area of the mapped images.

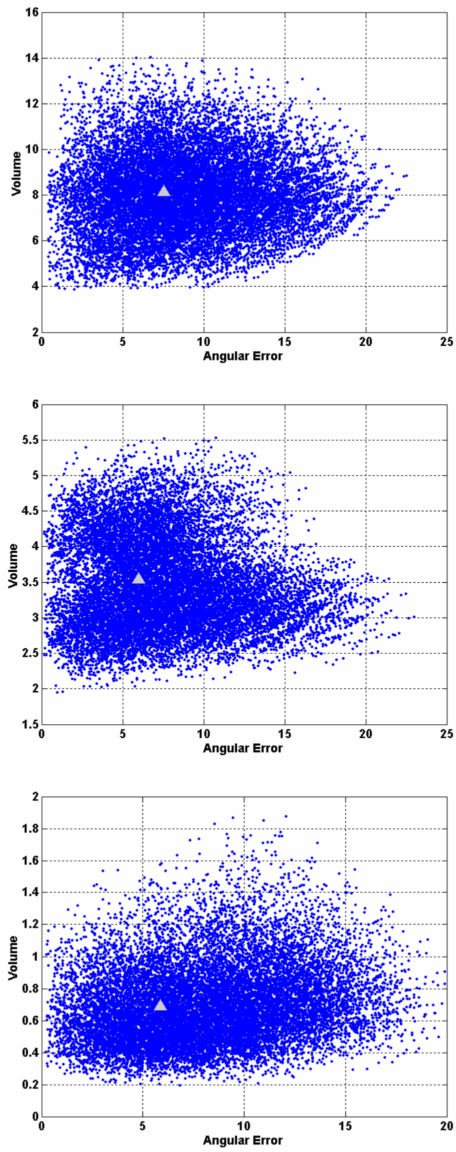

Figure 4. Scatter plots of the angular error versus volume of the mapped images.

Another way to characterize a solution map is by considering the position of the mapped image within the canonical gamut. This position can be given on the coordinate system with origin on the center of the canonical gamut. Transforming the center of the gamut of the mapped image to its polar coordinates on the specified system, we obtain a distance parameter and an angle parameter. The behavior of the mapped images versus these two parameters is given in figures 5 and 6 respectively.

As in the case of the area, we can not take the minimum value of the distance to select the optimal solution (what would mean selecting the grey world solution), and the problem cannot reduced to a maximisation or minimisation problem neither.

From figures 4, 5 and 6 we have deduced that optimal solutions converge on a unique area, a unique distance and a unique angle. By crossing the information from the values of the optimal maps, we can see in figure 7, that these optimal maps present a clear intersection between maps of optimal angle, maps of optimal distance and some maps of optimal area.

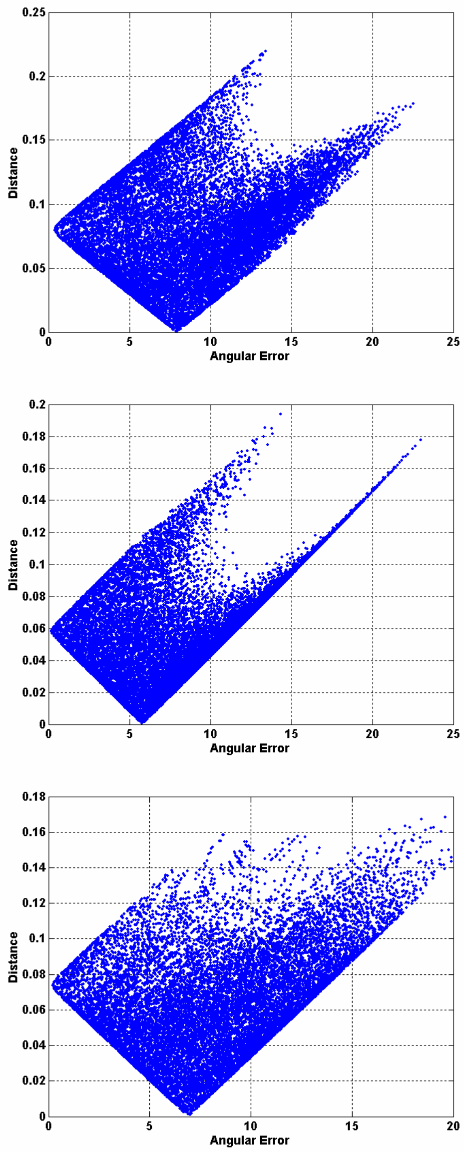

Figure 5. Scatter plots of the angular error versus position distance of the mapped images.

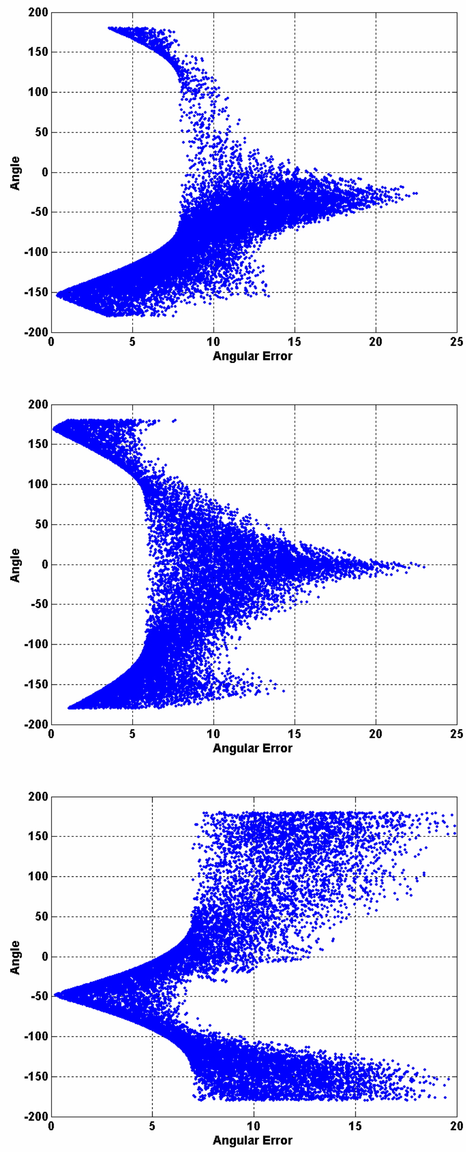

Figure 6. Scatter plots of the angular error versus position angle of the mapped images.

Results and Conclusions

In this paper we have explained some conclusions derived from an empirical work on colour constancy solutions. Methods based on gamut mapping approaches present a decision step that selects the optimal map from the set of all feasible maps computed at a former step. Different criteria have been proposed for this decision step, however, an exploration on how are the solutions within this feasible set have not been explored. This is the essential contribution of this paper.

We propose a method to sample the feasible set of maps based on some important assumptions on the image content. Considering that in an empirical work we can synthesise the optimal solution, we can compute the angular error between each sample solution and the optimal solution. This allows us to plot the relationship between the angular error of every sample with respect to specific properties of this sample.

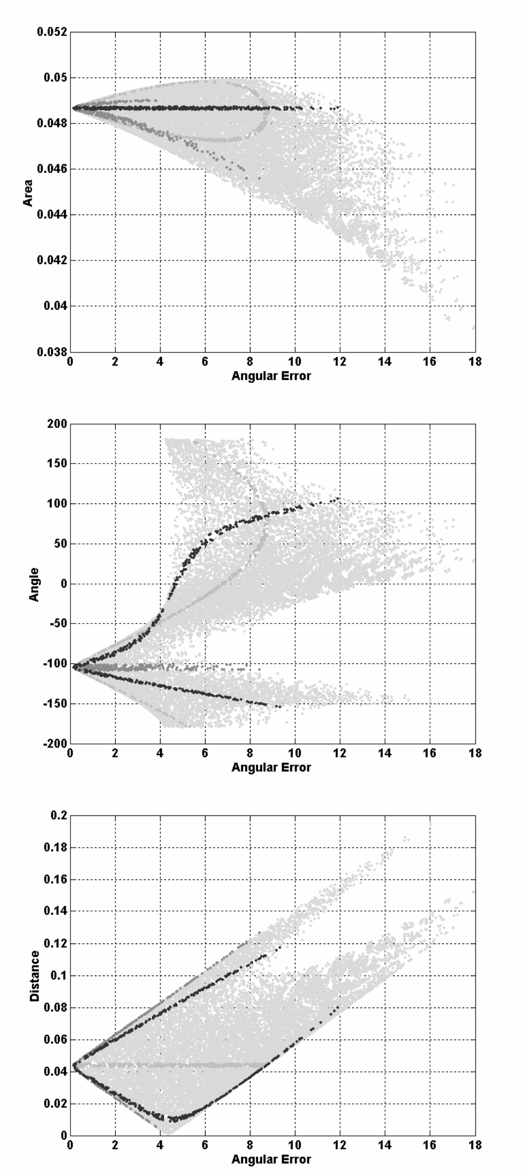

Figure 7. Scatter plots of the angular error versus area, angle and distance. Solutions grouped by optimal area (darkest), by optimal angle and by optimal distance (brightest).

Explored properties are: the volume of the image mapped with a specific solution, the mean map of the feasible set, the area of the chromaticity gamut of the mapped image and the position of this mapped image within the chromaticity gamut of the canonical surfaces. For the set of samples we have built, we can state that in general, the map with maximum volume, or the map which is the mean of the feasible set, or even the map with maximum chromaticity area do not present optimal errors for the analysed images.

On the other hand, we have explored other properties such as the position of the mapped image on the chromaticity plane. Since we have worked on a sampled set of maps, we can not build a general solution to the colour constancy problem, but we can conclude that some other criteria should be defined in order to approximate the optimal solutions.

Further work needs to be done in order to exploit these results for the definition of effective selection criteria. More properties should be explored and other sampling procedures could be defined.

References

1. B.A. Wandell, Foundations of Vision, Sinauer Associates, Inc, 1995.
2. Edwin H. Land, Scientific American, 237, 23 (1977).
3. H. Spitzer and S. Semo, Pattern Recognition, 35, 8 (2002).
4. K. Barnard, et al., IEEE Trans. on IP, 11, 9 (2002).
5. D.A. Forsyth., I. Journal of Computer Vision, 5, 1 (1990).
6. L.T. Maloney and B.A. Wandell, J. of the Opt. Soc. of America, 3, 1 (1986).
7. G. Sapiro. IEEE Trans. On PAMI, 21, 11 (1999).
8. Brian Funt, Florian Ciurea and John McCann. Retinex in matlab. 8th SID, pg. 112. (2000).
9. G. D. Finlayson, S. Hordley and P.M. Hubel, IEEE Trans. on PAMI, 23, 11 (2001).
10. G.D. Finlayson. IEEE Trans. on PAMI, 18, 10 (1996).
11. B. Funt, K. Barnard and L. Martin. Is colour constancy good enough? In 5th ECCV, pg. 445. (1998).
12. G.D. Finlayson, M.S. Drew and B.V. Funt. J. Opt. Soc. Am. A, 11 (1994).
13. G.D. Finlayson and R.Xu, Convex Programming Color Constancy, IEEE Workshop on Color and Photometric Methods in Computer Vision, (2003).

Biography

Francesc Tous is a research fellow at the Computer Vision Center at the Universitat Autònoma de Barcelona. In the year 2000 he acquired the degree of engineer in Computer Science by the Universitat Autònoma de Barcelona, and in 2002 the MSc in Computer Science by the Universitat Autònoma de Barcelona (area of Computer Vision). Currently he is a PhD Student at the Computer Vision Center and his research is focused on colour in computer vision, especially in the colour constancy problem and colour invariant representations.