

Max-Correlation Embedding Computation

Antonio Robles-Kelly

NICTA*, Locked Bag 8001, Canberra ACT 2601, Australia
Research School of Eng., ANU, Canberra ACT 0200, Australia

Abstract. In this paper, we present a method to compute an embedding matrix which maximises the dependence of the embedding space upon the graph-vertex coordinates and the incidence mapping of the graph. This treatment leads to a convex cost function which, by construction, attains its maximum at the leading singular value of a matrix whose columns are given by the incidence mapping and the embedded vertex coordinates. This, in turn, maximises the correlation between the spaces in which the embedding and the graph vertex coordinates are defined. It also maximises the dependence between the embedding and the incidence mapping of the graph. We illustrate the utility of the method for purposes of approximating the colour sensitivity functions of a set of over 20 commercially available digital cameras using a library of spectral reflectance measurements.

1 Introduction

In the pattern analysis community, there has recently been renewed interest in the embedding methods motivated by graph theory. One of the best known of these is ISOMAP [1]. Related algorithms include locally linear embedding which is a variant of PCA that restricts the complexity of the input data using a nearest neighbor graph [2], and the Laplacian eigenmap that constructs an adjacency weight matrix for the data-points and projects the data onto the principal eigenvectors of the associated Laplacian matrix [3]. Collectively, these methods are sometimes referred to as manifold learning theory.

Embedding methods can also be used to transform the graph-matching problem into one of point-pattern alignment. The problem is to find matches between pairs of point sets when there is noise, geometric distortion and structural corruption. There is a considerable literature on the problem and many contrasting approaches, including relaxation [4] and optimisation [5], have been attempted. However, the main challenge in graph matching is how to deal with differences in node and edge structure. One of the most elegant recent approaches to the graph matching problem has been to use graph-spectral methods [6], and exploit information conveyed by the eigenvalues and eigenvectors of the adjacency

* NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

matrix. For instance, Umeyama [7] has developed a method for finding the permutation matrix which best matches pairs of weighted graphs of the same size, by using a singular value decomposition of the adjacency matrices. Scott and Longuet-Higgins [8], on the other hand, align point-sets by performing singular value decomposition on a point association weight matrix. Shapiro and Brady [9] have reported a correspondence method which relies on measuring the similarity of the eigenvectors of a Gaussian point-proximity matrix. Kosinov and Caelli [10] have improved this method by allowing for scaling in the eigenspace. More recently, Sebastian and Kimia [11] have used a distance metric analogous to the string edit distance to perform object recognition from a dataset of shock graphs.

The main argument levelled against the techniques mentioned above is that they adopt a heuristic approach to the relational matching problem by using a goal-directed graph similarity measure. To overcome this problem, several authors have proposed more general approaches using ideas from information and probability theory. For instance, Wong and You [12] defined an entropic graph-distance for structural graph matching. Christmas, Kittler and Petrou [4] have shown how a relaxation labeling approach can be employed to perform matching using pairwise attributes whose distribution is modelled by a Gaussian. Wilson and Hancock [13] have used a MAP (maximum *a posteriori*) estimation framework to accomplish purely structural graph matching. Recently, Caetano *et al.* have proposed a method to estimate the compatibility functions for purposes of learning graph matching [14].

Here, we focus on the recovery of an embedding matrix based upon the graph and the embedding itself. We do this by maximising the correlation for both, the node-set for the graph and the metric space in which the embedding is defined. To this end, we depart from a cost function which aims at minimising the matrix norm between the embedding and the incidence mapping of the graph. We then rewrite the cost function so as to involve the eigenfunctions of two matrices of inner products. We show the utility of the method presented here for purposes of approximating the spectral sensitivity function of a set of over 20 digital cameras using a library of reflectances of a calibration target, *i.e.* an X-Rite ColorChecker chart.

2 Graph Theory and Spectral Geometry

As mentioned above, we aim at computing a linear mapping that can be used to embed the graph-vertices into a space of finite dimensionality based upon a known transformation to a subspace constrained by the edge space. In this manner, the embedding will reflect the structure of the edge-space of the graph while being based upon a known relationship between the graph vertex-set and the target space Ω . This has two main advantages. Firstly, the target space for the recovered mapping can be used to constrain the embedding. Secondly, note that the mapping sought here embeds the graph vertices using a linear operator drawn from spectral geometry. This is not only practically useful but

theoretically important since it provides a link between the spectra of graphs and linear operators.

2.1 On the Incidence Mapping of Graphs

Here, we aim at recovering a mapping \mathcal{T} which is a matrix whose dimensionality is $\Omega \times |\mathcal{V}|$. In other words, we aim at recovering an operator which can embed the nodes of a graph G into a space \mathfrak{R}^Ω . To commence, we require some formalism. Let $G = (\mathcal{V}, \mathcal{E}, \mathcal{A})$ denote the graph with node-set $\mathcal{V}_i = \{V_1, \dots, V_{|\mathcal{V}|}\}$, edge-set $\mathcal{E} = \{e | V_a, V_c \in \mathcal{V}\}$ and attribute-set $\mathcal{A} = \{A_1, \dots, A_{|\mathcal{V}_i|}\}$.

Here, we view, in general, the vertex-attributes $\mathbf{A}(a)$ as vectors, where each of these has a one-to-one correspondence to a graph vertex. This also permits the computation of the weight matrix \mathcal{W} with elements $\mathcal{W}(a, c)$ for the graph G . The weight matrix \mathcal{W} can be related to the un-normalised Laplacian through the relationship $\mathcal{L} = \mathbf{D} - \mathcal{W}$, where \mathbf{D} is a diagonal matrix such that $\mathbf{D} = \text{diag}(\text{deg}(1), \text{deg}(2), \dots, \text{deg}(|\mathcal{V}|))$ and $\text{deg}(c) = \sum_{a=1}^{|\mathcal{V}|} \mathcal{W}(a, c)$ is the degree of the node indexed c in the graph [6].

The use of the graph Laplacian is important, since it permits the use of the incidence mapping. Note that the incidence mapping \mathcal{I} is independent of the orientation of the edges in \mathcal{E} . Moreover, it is an operator independent of the vertex-basis, i.e. its permutation invariant [15], which can be recovered via a Young-Householder [16] decomposition on the graph Laplacian such that $\mathcal{L} = \mathcal{I}\mathcal{I}^T$.

2.2 Embedding Computation

With these ingredients, we can formalise the problem as that of recovering the linear mapping \mathcal{T} such that

$$\min_{\mathcal{T}} \left\{ f(\mathcal{T}) \right\} = \min_{\mathcal{T}} \left\{ \|\mathcal{Y} - \mathcal{T}\mathcal{I}\|^2 \right\} \quad (1)$$

given the embedding $\mathcal{Y} \in \mathfrak{R}^{\Omega \times |\mathcal{V}|}$ of \mathcal{V} in Ω and $\mathcal{I} \in \mathfrak{R}^{\Gamma \times |\mathcal{V}|}$, as before, is the incidence mapping of G .

It is worth noting in passing that this is akin to point pattern matching settings where the problem is that of finding a transformation which can be used to map the data points onto their counterparts in the model point-set. Nonetheless its similarities, the main difference is that, here, given the coordinates $\mathcal{I}_{v,\gamma}$ and $\mathcal{Y}_{v,\omega}$ of the embeddings and incidence mappings for the node $v \in \mathcal{V}$ in the dimensions $\gamma \in \Gamma$ and $\omega \in \Omega$, we aim at recovering the entries $\phi_{i,j}$ of the matrix \mathcal{T} such that

$$\min_{\mathcal{T}} \left\{ f(\mathcal{T}) \right\} = \min_{\phi_{\omega,\gamma} \in \mathcal{T}} \left\{ \sum_{v \in \mathcal{V}} \left(\|\mathcal{Y}_{v,\omega} - \sum_{\gamma \in \Gamma} \phi_{\omega,\gamma} \mathcal{I}_{v,\gamma}\|^2 \right) \right\} \quad (2)$$

rather than the corresponding permutation and rotation matrices.

Indeed, the cost function above could be tackled using a least squares solution. This naturally leads to a solution akin to a linear regressor whereby $\phi_{\omega,\gamma}$ can be

viewed as the slope of the lines $\mathcal{Y}_{v,\omega} = \phi_{\omega,\gamma} \sum_{\gamma \in \Gamma} \mathcal{I}_{v,\gamma}$. This can be viewed as a minimisation on the distance about the $\mathcal{Y}_{v,\omega}$ variables [17]. Note that it would be more desirable to use the distance, i.e. norm, spanned by both, the embedding and the incidence mapping. Thus, we rewrite the cost function above making use of the matrix $\mathbf{M}_{\omega,\gamma} = [\mathcal{Y}_{\cdot,\omega} | \mathcal{I}_{\cdot,\gamma}]$, where $\mathcal{Y}_{\cdot,\omega} = [\mathcal{Y}_{1,\omega}, \mathcal{Y}_{2,\omega}, \dots, \mathcal{Y}_{|\mathcal{V}|,\omega}]^T$ and $\mathcal{I}_{\cdot,\gamma} = [\mathcal{I}_{1,\gamma}, \mathcal{I}_{2,\gamma}, \dots, \mathcal{I}_{|\mathcal{V}|,\gamma}]^T$ as follows

$$\min_{\mathcal{T}} \left\{ f(\mathcal{T}) \right\} = \min_{\substack{\xi \in \Omega \\ \psi \in \Gamma}} \left\{ \sum_{v \in \mathcal{V}} \left(\left\| \mathcal{Y}_{v,\omega} - \sum_{\gamma \in \Gamma} \frac{\xi^T \mathbf{M}_{\omega,\gamma} \psi}{\xi^T \psi} \mathcal{I}_{v,\gamma} \right\|^2 \right) \right\} \quad (3)$$

where, by construction, ξ^T and ψ are the eigenvectors of $\mathbf{M}_{\omega,\gamma}^T \mathbf{M}_{\omega,\gamma}$ and $\mathbf{M}_{\omega,\gamma} \mathbf{M}_{\omega,\gamma}^T$, respectively [18].

The advantage of Equation 3 resides in the fact that, as we will see in the following section, the term

$$\phi_{\omega,\gamma} = \max_{\substack{\xi \in \Omega \\ \psi \in \Gamma}} \left\{ \frac{\xi^T \mathbf{M}_{\omega,\gamma} \psi}{\xi^T \psi} \right\}$$

maximises both, the correlation, in the geometric sense, of both, the pairs $\mathcal{Y}_{\cdot,\omega}$, $\mathcal{I}_{\cdot,\gamma}$ and $\mathcal{Y}_{v,\cdot}$, $\mathcal{I}_{v,\cdot}$. This is, it maximises the dependence of the recovered embedding upon the incidence mapping and that of the target space on the vertex coordinates. Moreover, as an added advantage, the computation of $\phi_{\omega,\gamma}$ can be done in a straightforward manner via the application of Singular Value Decomposition (SVD) to the matrix $\mathbf{M}_{\omega,\gamma}$ [18], i.e. $\phi_{\omega,\gamma}$ is the leading singular value of $\mathbf{M}_{\omega,\gamma}$.

2.3 Max-Correlation

Now we examine the link between the cost function above and the eigenvectors ξ and ψ . To this end, we make use of the matrix of scalar products $\mathbf{H} = \mathbf{M}_{\omega,\gamma} \mathbf{M}_{\omega,\gamma}^T$. Note that, since the developments here apply equally to the $\mathbf{M}_{\omega,\gamma}^T \mathbf{M}_{\omega,\gamma}$, we focus on \mathbf{H} throughout the section.

Let ξ_l be the l^{th} eigenvector of \mathbf{H} scaled so its sum of squares is equal to the corresponding eigenvalue τ_l . Since $\mathbf{H}\xi_l = \tau_l \xi_l$ and $(\mathbf{J}\mathbf{J}^T)\xi_l = \mathbf{H}\xi_l$, it follows that the squared distance between a pair of entries in the matrix \mathbf{H} can be written as

$$\| \eta_i - \eta_j \|^2 = \sum_{l=1}^N \tau_l (\xi_l(i) - \xi_l(j))^2 = \mathbf{H}(i,i) + \mathbf{H}(j,j) - 2\mathbf{H}(i,j) \quad (4)$$

where η_i and η_j are coordinates in the embedding space such that their inner product corresponds to the entry indexed i, j of \mathbf{H} and N is its rank.

With these ingredients, we can recover the variables ξ_i for the vertices in the graph such that the weighted correlations between their embedding vectors are maximum or minimum by extremising the quantity

$$\epsilon = \sum_{i,j} \left\| \xi_i \eta_i - \xi_j \eta_j \right\|^2 \quad (5)$$

To take our analysis further, we use Equation 4 and, after some algebra, write

$$\epsilon = \sum_{i,j} (\xi_i^2 \mathbf{H}(i,i) + \xi_j^2 \mathbf{H}(j,j) - 2\xi_i \xi_j \mathbf{H}(i,j)) \quad (6)$$

Note that, Equation 6 can be divided into two sets of terms, one for the diagonal and the other for the off-diagonal elements of \mathbf{H} as follows

$$\epsilon = 2M \sum_i \xi_i^2 \mathbf{H}(i,i) - \sum_{i,j} 2\xi_i \xi_j \mathbf{H}(i,j) \quad (7)$$

where M is the order of \mathbf{H} and we have used the fact that

$$\sum_{i,j} \xi_i^2 \mathbf{H}(i,i) = N \sum_i \xi_i^2 \mathbf{H}(i,i) \quad (8)$$

and

$$\sum_{i,j} \xi_i^2 \mathbf{H}(i,i) = \sum_{i,j} \xi_j^2 \mathbf{H}(j,j) \quad (9)$$

Note that maximising the first term in the right-hand side of Equation 7 implies minimising the second one and vice versa. The proof of this hinges in the properties of spectral radii of symmetric matrices [19, 20]. This is also consistent with the work of Chung on isoperimetric inequalities [21]. Thus, we can focus on the term

$$\hat{\epsilon} = - \sum_{\substack{i,j \\ i \neq j}} 2\xi_i \xi_j \mathbf{H}(i,j) \quad (10)$$

Furthermore, to write Equation 10 in compact form, we can define a matrix $\hat{\mathbf{H}}$ which comprises the off-diagonal elements of \mathbf{H} as follows

$$\hat{\mathbf{H}}(i,j) = \begin{cases} \mathbf{H}(i,j) & \text{if } i \neq j \\ 0 & \text{otherwise} \end{cases} \quad (11)$$

and write

$$\hat{\epsilon} = -2\boldsymbol{\Pi}^T \hat{\mathbf{H}} \boldsymbol{\Pi} \quad (12)$$

where $\boldsymbol{\xi} = [\xi_1, \xi_2, \dots, \xi_M]^T$ is a column vector of order M whose i^{th} element is given by ξ_i . Note that the expression above is a Rayleigh quotient. Thus, maximising ϵ is equivalent to minimising $\boldsymbol{\xi}^T \hat{\mathbf{H}} \boldsymbol{\xi}$, which implies that $\boldsymbol{\xi}$ is given by the eigenvector of $\hat{\mathbf{H}}$ which corresponds to the eigenvalue whose rank is the smallest. In this case, $\boldsymbol{\xi}$ is the maximiser of the correlation between the vectors η_i and η_j .

3 Recovering Camera Spectral Sensitivity Functions

In computer vision, video and graphics, we rely upon cameras and rendering contexts to capture and reproduce colour information. Moreover, the accurate

capture and reproduction of colours as acquired by digital camera sensors is an active area of research which has applications in colour correction [22–24], camera simulation [25] and sensor design [26].

To better understand the relation between the spectral radiance and the colour output of digital cameras, recall that we can express the colour output of the detector at pixel u as follows

$$I_k(u) = g(u)\mathbf{S}(u)^T \text{diag}(\mathbf{Q}_k)\mathbf{L}, \quad (13)$$

where $I_k(u)$ is the image radiance for any of the three colour channels $k = \{R, G, B\}$ at the pixel u . $\mathbf{S}(u)$ is a vector indexed to wavelength whose entries are given by the surface reflectance $S(\lambda, u)$ at the wavelength λ . \mathbf{L} is the power spectrum of the light with the elements $L(\lambda)$ corresponding to the spectral power at the wavelength λ . \mathbf{Q}_k is a vector whose element $Q_k(\lambda)$ corresponds to the spectral sensitivities of the k^{th} colour sensor at the wavelength λ . When dealing with flat surfaces such as colour charts, we can assume that $g(u) = 1$. This expression has been used widely in the literature [27] and is consistent with reflectance models in computer vision, such as that in [28].

By inspection, it is straightforward to note that, in Equation 13, if the object reflectance and illuminant power spectrum are known, the camera spectral sensitivity functions are, indeed, a linear mapping which “embeds” the product of the reflectance and the illuminant into the colour space. Further, we can view the product of the reflectance and the illuminant as the incidence mapping as presented previously and the ensuing colour triples as the embedding \mathcal{Y} . As a result, the matrices $\mathbf{M}_{\omega, \gamma}$ are defined in the colour and wavelength spaces. This is, Ω corresponds to the colour and Γ to the wavelength domain.

In the following experiments, we employ the dataset presented in [29]. This is one of the most complete studies on commercial digital camera spectral responses comprising 28 commercial models¹ Note that the dataset presented in [29] does not contain colour imagery, but rather the sensitivity functions themselves. Thus, for the dataset in [29], we have used the ground-truth power spectrum of the illuminants and the reflectance for each of the colour tiles in a semi-gloss (SG) X-Rite ColorChecker target with 140 colour patches. This is straightforward since the ColorChecker is a flat surface whose mean-scattered power can be easily computed.

3.1 Illuminants

Throughout the section, we use two standard calibrated light sources. This is in line with the standard illuminants defined by the CIE [30]. Our calibrated light sources correspond to the A and D series of illuminants. For our A series illuminant, we have used a tungsten-filament light with a correlated colour temperature (CCT) [31] of 2700°K. Our D series light is an artificial sunlight with a CCT of 6500°K (D65).

¹ These can be downloaded from <http://www.cis.rit.edu/jwgu/research/camspec/db.php>

It is worth noting in passing that the use of these two light sources is also aimed at spanning across a wide variety of real-world settings. This is as the A series illuminant correspond to the incandescent filament lights widely used in households and street lighting, whereas the D series illuminant accounts for outdoor environments.

3.2 Reflectance Library

Recall that we also require the spectral reflectance of the color tiles in our X-Rite target so as to compute the covariance matrices used by our method. To this end, we have acquired the reflectance of each colour tile in the X-Rite charts using a StellarNet Bluewave Spectrometer. The spectrometer delivers a spectrum of 1716 samples per tile over the visible and near infrared range. The measurements have been effected using a two-way integrating sphere and a halogen-Deuterium calibrated light source in the $[200nm - 1700nm]$ range.

Note that, for our reflectance library, we have followed the ISO standard for the visible spectrum and archived the reflectance in the range $[400nm, 780nm]$, which yields a total of 599 samples per tile over the 164 colours in the two charts. It is worth stressing in passing that we have opted for the ISO standard over the CIE since the latter is a subset of the former (the CIE standard dictates the visible range is given by the interval $[400nm, 700nm]$) [32]. As mentioned earlier, for our reflectance, we the 140 tiles of 100 different colours including the white ones. The inclusion of the white tile is important since this allows for the illuminant power spectrum computation as required in Equation 13.

3.3 Experiments

With the spectral and colour data in hand, we proceed to provide a quantitative analysis regarding the approximation yielded by our method. To this end, we have used the Euclidean deviation, in degrees, between the sensitivity functions approximated by our method and the corresponding ground truth.

We have used these two metrics so as to account for both, variations in the “shape” of the colour sensitivity functions and power spectrum of the illuminant with respect to the ground truth as well as colour variations induced by the approximation presented here. For our dataset, we have used the colours extracted from the trichromatic imagery of the X-Rite ColourChart as acquired by each camera. For both datasets, we have compared these ground truth data to the colours computed using the sensitivity functions and illuminant power spectrum approximated by our method when applied to the colour checker reflectance library.

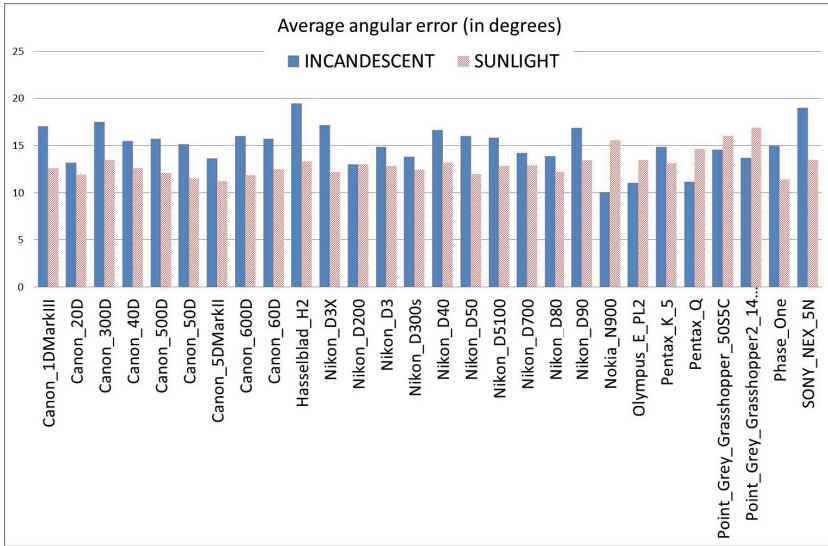


Fig. 1. The average angular error on the spectral responses across three channels for each camera in the dataset presented in [29]




Fig. 2. Colour sensitivity functions for a sample camera in our dataset. The two panels show the ground truth and approximated sensitivity functions for the Nikon D80 camera in [29].

Figure 1 shows the Euclidean angular errors, in degrees, for the sensitivity functions corresponding to the cameras in the dataset. Note that the Euclidean angular errors are often in the order of 12 degrees for the dataset in [29].

To illustrate the quality of the approximation in a qualitative manner, in Figure 2 we show the ground truth and approximated colour sensitivity functions a sample camera in the dataset. Note the close accordance of the colour sensitivity functions approximated by our method with respect to the ground truth. Note, however, that even for 16.94 degrees average error yielded by the method presented here, the overall shape approximated by our approach is in good accordance with the ground truth (in the first panel).

4 Conclusions

In this paper, we present a method to compute an embedding which maximises the dependence of the embedding matrix upon the graph-vertex coordinates and that of the target space on the incidence mapping. This treatment leads to a convex cost function whose optimum is attained by the leading singular value of a matrix whose columns are given by the incidence mapping and the embedded vertex coordinates. We illustrate the utility of the method for purposes of approximating the colour sensitivity functions of a set of over 20 commercially available digital cameras from a single image of a colour calibration target. We do this by using a set of spectral reflectance measurements. Thus, our spectral sensitivity recovery via the computation of the corresponding embedding can be viewed as the result of maximising the relationship between the colour values yielded by the camera and the spectra in the reflectance library.

References

1. Tenenbaum, J.B., de Silva, V., Langford, J.C.: A global geometric framework for nonlinear dimensionality reduction. *Science* 290(5500), 2319–2323 (2000)
2. Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* 290, 2323–2326 (2000)
3. Belkin, M., Niyogi, P.: Laplacian eigenmaps and spectral techniques for embedding and clustering. *Neural Information Processing Systems 14*, 634–640 (2002)
4. Christmas, W.J., Kittler, J., Petrou, M.: Structural matching in computer vision using probabilistic relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17(8), 749–764 (1995)
5. Gold, S., Rangarajan, A.: A graduated assignment algorithm for graph matching. *PAMI* 18(4), 377–388 (1996)
6. Chung, F.: *Spectral Graph Theory*. American Mathematical Society (1997)
7. Umeyama, S.: An eigen decomposition approach to weighted graph matching problems. *PAMI* 10(5), 695–703 (1988)
8. Scott, G., Longuet-Higgins, H.: An algorithm for associating the features of two images. *Proceedings of the Royal Society of London* 244(B), 21–26 (1991)
9. Shapiro, L., Brady, J.M.: Feature-based correspondance - an eigenvector approach. *Image and Vision Computing* 10, 283–288 (1992)
10. Caelli, T., Kosinov, S.: An eigenspace projection clustering method for inexact graph matching. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 26(4), 515–519 (2004)
11. Sebastian, T.B., Klein, P.N., Kimia, B.B.: Shock-based indexing into large shape databases. In: *European Conference on Computer Vision*, vol. 3, pp. 731–746 (2002)
12. Wong, A.K.C., You, M.: Entropy and distance of random graphs with application to structural pattern recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 7, 599–609 (1985)
13. Wilson, R., Hancock, E.R.: Structural matching by discrete relaxation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 19(6), 634–648 (1997)
14. Caetano, T., Cheng, L., Le, Q., Smola, A.: Learning graph matching. In: *Proceedings of the 11th International Conference on Computer Vision*, pp. 14–21 (2007)
15. Biggs, N.L.: *Algebraic Graph Theory*. Cambridge University Press (1993)

16. Young, G., Householder, A.S.: Discussion of a set of points in terms of their mutual distances. *Psychometrika* 3, 19–22 (1938)
17. Björck, A.: Numerical methods for least squares problems. SIAM (1996)
18. Golub, G.H., Loan, C.F.V.: *Matrix Computations*. The Johns Hopkins Press (1996)
19. Torgerson, W.S.: Multidimensional scaling I: Theory and method. *Psychometrika* 17, 401–419 (1952)
20. Varga, R.S.: *Matrix Iterative Analysis*, 2nd edn. Springer (2000)
21. Chung, F.: Discrete Isoperimetric Inequalities. *Surveys in Differential Geometry IX* (2004)
22. Wandell, B.A.: The synthesis and analysis of color images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 9(1), 2–13 (1987)
23. Brainard, D.H., Stockman, A.: *Colorimetry*. McGraw-Hill (1995)
24. Finlayson, G.D., Drew, M.S.: The maximum ignorance assumption with positivity. In: *Proceedings of the IS&T/SID 4th Color Imaging Conference*, pp. 202–204 (1996)
25. Longere, P., Brainard, D.H.: Simulation of digital camera images from hyperspectral input. In: van den Branden Lambrecht, C. (ed.) *Vision Models and Applications to Image and Video Processing*, pp. 123–150. Kluwer (2001)
26. Ejaz, T., Horiuchi, T., Ohashi, G., Shimodaira, Y.: Development of a camera system for the acquisition of high-fidelity colors. *IEICE Transactions on Electronics E89-C(10)*, 1441–1447 (2006)
27. Kimmel, R., Elad, M., Shaked, D., Keshet, R., Sobel, I.: A variational framework for retinex. *International Journal of Computer Vision* 52(1), 7–23 (2003)
28. Finlayson, G.D., Schaefer, G.: Solving for colour constancy using a constrained dichromatic reflection model. *International Journal of Computer Vision* 42(3), 127–144 (2001)
29. Jiang, J., Liu, D., Gu, J., Süssstrunk, S.: What is the space of spectral sensitivity functions for digital color cameras? In: *Workshop on Applications of Computer Vision*, pp. 168–179 (2013)
30. Wyszecki, G., Stiles, W.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*. Wiley (2000)
31. Judd, D.B., Macadam, D.L., Wyszecki, G., Budde, H.W., Condit, H.R., Henderson, S.T., Simonds, J.L.: Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America* 54(8), 1031–1036 (1964)
32. Robles-Kelly, A., Huynh, C.P.: *Imaging Spectroscopy for Scene Analysis*. Springer (2013)