



Deep Convolutional Gaussian Mixture Model for Stain-Color Normalization of Histopathological Images

Farhad Ghazvinian Zanjani^(✉), Svitlana Zinger, and Peter H. N. de With

Eindhoven University of Technology, 5612 AJ Eindhoven, The Netherlands
f.ghazvinian.zanjani@tue.nl

Abstract. Automated microscopic analysis of stained histopathological images is degraded by the amount of color and intensity variations in data. This paper presents a novel unsupervised probabilistic approach by integrating a convolutional neural network (CNN) and the Gaussian mixture model (GMM) in a unified framework, which jointly optimizes the modeling and normalizing the color and intensity of hematoxylin- and eosin-stained (H&E) histological images. In contrast to conventional GMM-based methods that are applied only on the color distribution of data for stain color normalization, our proposal learns how to cluster the tissue structures according to their shape and appearance and simultaneously fits a multivariate GMM to the data. This approach is more robust than standard GMM in the presence of strong staining variations because fitting the GMM is conditioned on the appearance of tissue structures in the density channel of an image. Performing a gradient descent optimization in an end-to-end learning, the network learns to maximize the log-likelihood of data given estimated parameters of multivariate Gaussian distributions. Our method does not need ground truth, shape and color assumptions of image contents or manual tuning of parameters and thresholds which makes it applicable to a wide range of histopathological images. Experiments show that our proposed method outperforms the state-of-the-art algorithms in terms of achieving a higher color constancy.

Keywords: Computational pathology · Stain-color normalization
Gaussian mixture model (GMM)
Convolutional neural network (CNN)

1 Introduction

Computational histopathology involves computer-aided diagnosis (CAD) for microscopic analysis of stained histopathological slides to study presence, localization or grading of disease. Manual staining by adding contrasting dyes prior to microscopic imaging is a common clinical practice. Such a non-quantified manual procedure causes significant variations in tissue intensity and color, which also

originate from other factors such as imaging device characteristics. Such undesirable effects become more problematic when different laboratories share digital images [1]. A well-known approach for compensating the color variations is applying color normalization techniques. Recent studies show that color normalization as preprocessing stage leads to a higher performance in a CAD system [2]. To do so, several methods are investigated [1, 3–9]. The proposed approaches can be divided into two categories: *stain-color deconvolution* and *template color matching*, which are first briefly explained below, and then we describe our contributions.

Stain-color deconvolution [10] methods are considering prior knowledge of the reference stain vector for every dye and split an input RGB image into three stain channels, each representing the actual stain color. Ruifrok *et al.* [10] introduce this prior knowledge by manually selecting pixels that represent a specific stain class and then compute the color deconvolution vector. Because of the semi-automatic nature, several studies are performed later for automated extraction of stains by e.g. using the singular value decomposition (SVD) technique [6], probabilistic Gaussian mixture model (GMM) [8], using a prior for stain matrix estimation [3] and stain color descriptions along with training a supervised relevance vector machine [4]. Although these solutions lead to a better stain estimation, they solely limited to image color information, while the spatial dependency among tissue structures has been ignored [1]. This causes shortcomings for stain deconvolution approaches when severe staining variations occur in the data.

Template color matching methods proposed by Reinhard *et al.* [5] rely on aligning the statistical color distribution (e.g. mean and standard deviation) of a *source* image with a *template* image. The authors used a set of linear transformations for assigning a unimodal distribution to each channel of the CIELAB color model. Afterwards, each channel was treated independently for alignment. Since there is a dependency between the color channels due to dye contribution, this approach has drawbacks which have been addressed in [1, 4]. For solving this problem, separate transformations are performed on stain classes [8], or on tissue classes [1, 9]. For avoiding artifacts at the border of different classes under different transformations, a weighted contribution of these transformations in the final color image is considered. Two proposed solutions consist of estimating weights of the GMM [8] and training a naive Bayesian classifier [1]. This solution introduce multiple parameters and thresholds which cannot be optimally applied to a new dataset or even fail if the tissue type changes.

Regarding to [8, 9], we propose a stain-color normalization method based on GMM but with a different approach. Our contribution is threefold. First, we introduce a new unsupervised stain-color normalization method based on deep convolutional Gaussian mixture models (DCGMM). Our proposal benefits from a convolutional neural network (CNN) for performing soft-assignment clustering of tissue structures in an unsupervised manner. Second, in contrast to the previous GMM-based color normalization methods that work on color point-clouds by considering each pixel independently in input space, our method fits a GMM to an input color image with processing and involving the visual contents, such as the appearance and shape of regions in the image intensity (density) channel by

means of the CNN. Such an approach outperforms the previous methods specifically when strong staining variations appear in the data (see Fig. 1). This outcome occurs because our method is independent from the chromatic information for tissue class assignment and consequent distribution modeling. Finally, instead of using a common expectation-maximization (EM) algorithm for optimizing the GMM, we introduce an end-to-end learning procedure for optimizing the parameters of the CNN and the GMM together. To our knowledge, it is the first time that this approach is used for training a DCGMM.




Fig. 1. Example of standard GMM failure in tissue clustering using HSD color space; (a) RGB H&E image; (b) 3-class standard GMM; (c) 3-class DCGMM.

2 Methods

We first provide a brief overview of the standard GMM method and introduce the notation used in the paper. Afterwards, we introduce our DCGMM method.

Gaussian Mixture Model of data (\mathbf{x}), can be presented as a linear superposition of K Gaussian mixtures in terms of discrete *latent* variables (\mathbf{z}), in the form of

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \quad (1)$$

The K -dimensional binary random variable \mathbf{z} has one-hot encoding ($z_k \in \{0, 1\}$; $\sum_{k=1}^K z_k = 1$), which represents the tissue class in our study. In Eq. (1), the mixing coefficients π_k must satisfy $0 \leq \pi_k \leq 1$ together with $\sum_{k=1}^K \pi_k = 1$, in order to fulfill a valid probability definition [11]. Here, \mathcal{N} stands for a multivariate normal distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix $\boldsymbol{\Sigma}_k$. If we consider π_k as prior probability of class z_k , its posterior probability called *responsibility* can be written as follows [11]:

$$\gamma(z_k) = p(z_k = 1 | \mathbf{x}) = \frac{\pi_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}. \quad (2)$$

According to Eq. (1), the (natural) log-likelihood function for an image ($\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$) with total number of pixels (observations) equal to N is

$$\ln p(\mathbf{X} | \boldsymbol{\pi}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \sum_{n=1}^N \ln \left\{ \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right\}. \quad (3)$$

Given the GMM, the goal is to maximize the likelihood function (Eq. (3)) with respect to the parameters $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k, \pi_k)$. The common approach for this is using the EM algorithm by iteratively evaluating the responsibilities (Eq. (2)) and re-estimating the parameters.

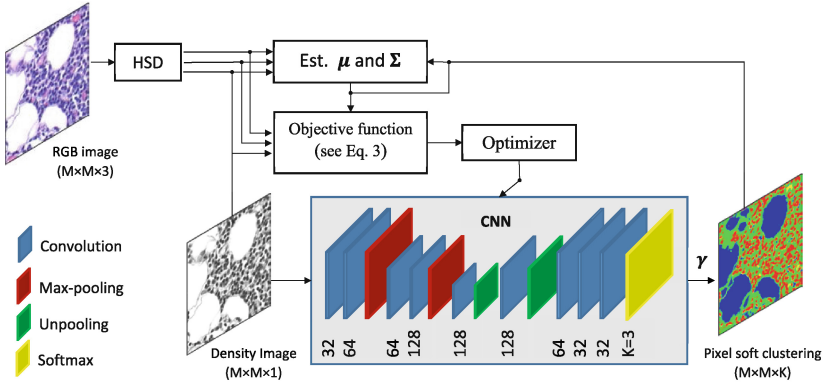


Fig. 2. Block diagram of DCGMM in training phase.

Deep Convolutional GMM (DCGMM): The recent development of deep generative models has invoked some extensions to the standard GMM [12, 13]. Two proposed approaches are (1) constructing a stack of multiple GMM layers on top of each other in a hierarchical architecture [13] that is optimized by EM-based algorithm and (2) using auto-encoder neural networks while applying the GMM on their low-dimensional representations [12] that has been studied for unsupervised anomaly detection. This paper presents a different extension to the standard GMM by introducing the DCGMM that is a fully-convolutional CNN of which the parameters are optimized to fit a GMM to the input image, in an end-to-end learning, using gradient descent and back-propagation algorithms.

We aim to fit a GMM to the pixel-color distribution conditioned on tissue classes. For processing the image and detecting the tissue classes, the high capability of the CNN has been exploited. To do so, estimating the responsibility coefficients (Eq. (2)) is performed by a CNN. For a better understanding, one can consider that the E-step in an EM-based optimization is replaced by a CNN. However, all parameters of the GMM and CNN are jointly optimized by the gradient descent algorithm. In the DCGMM, the negative log-likelihood (maximizing Eq. (3)) is used as the loss function.

Our color normalization algorithm can be split into two phases, training the DCGMM and the color transformation (inference). In the training phase, we fit a GMM to the data. After the training stage is finished and in inference mode, the template image and source image are separately supplied to the model. Consequently, the parameters of the fitted Gaussian distributions and their mixture coefficients ($\boldsymbol{\pi}$) are computed in those two images. Afterwards, the multivariate

Gaussian distributions of the source image are transformed (aligned) to have similar parameters of distributions as the template image, while $\boldsymbol{\pi}$ is kept unchanged. In the rest of this section, we explain these two phases in detail.

Fitting Gaussian Distribution to Data: First the RGB color values are transferred to the hue-saturation-density (HSD) color system [14]. In a HSD model of histopathological images, the density channel (D) is linearly related to the amount of stain while the other two chromatic channels are independent. This property suits well to the analysis of transmitted light microscopy, compared to the alternative color spaces [1]. We only use the normalized to zero-mean (centered) density channel as the input to the network. Ignoring the chromatic information and only clustering the tissue structures according to their appearance (normalized density channel), alleviating the effect of strong staining variations in images. The proposed CNN has a fully-convolutional architecture, consisting of several convolutional layers, rectified linear units (ReLU) as nonlinearity functions and (un)pooling operators. The reduced image size after applying two stages of max-pooling returns back to its original size by applying un-pooling operations. After the last convolutional layer, there is a softmax layer. The network aims to estimate the responsibility values (see Eq. (2)) for each pixel in the input image. Using the softmax layer at the output of the network guarantees the constraint of $\sum_{k=1}^K \gamma_k = 1$. Since we are working on H&E histopathological images, each pixel in the image mostly belongs to one out of three clusters ($K = 3$): hematoxylin, eosin and background (not stained). Because the biological composition of tissue related to each pixel in the image leads to a varying stain absorption ratio between pixels, the color of each pixel can be presented by a weighted sum of the different stains used. This property can be reflected in the responsibility coefficient (γ) of a GMM, which is estimated in the softmax layer of the network in our proposed model. The calculation of the required partial derivatives of negative log-likelihood (loss function) with respect to its parameters ($\boldsymbol{\pi}$, $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$) for performing a gradient descent algorithm can be found in [15, p. 45].

Algorithm 1. DCGMM training (left) and inference (right) algorithm

$\boldsymbol{\theta}_{net} \leftarrow$ random CNN parameters repeat $\mathbf{X} \leftarrow$ input image $\mathbf{X}_h, \mathbf{X}_s, \mathbf{X}_d \leftarrow$ HSD(\mathbf{X}) $\boldsymbol{\gamma} \leftarrow f_{net}(\bar{\mathbf{X}}_d, \boldsymbol{\theta}_{net})$ $N_k \leftarrow \sum_{n=1}^N \gamma(z_{nk})$ $\boldsymbol{\mu}_k \leftarrow \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) \mathbf{x}_n$ $\boldsymbol{\Sigma}_k \leftarrow \frac{1}{N_k} \sum_{n=1}^N \gamma(z_{nk}) (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T$ $\boldsymbol{\pi}_k \leftarrow \frac{N_k}{N}$ $\mathcal{L} \leftarrow -\sum_{n=1}^N \ln \{ \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \}$ $\boldsymbol{\theta}_{net} \stackrel{+}{\leftarrow} -\nabla_{\boldsymbol{\theta}_{net}}(\mathcal{L})$ until stopping criterion	$\mathbf{X}_{h,s,d}^t \leftarrow$ HSD Template image $\boldsymbol{\gamma}^t, \boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t, \boldsymbol{\pi}^t \leftarrow$ Ops.(1-5) $\mathbf{X}_{h,s,d}^s \leftarrow$ HSD Source image $\boldsymbol{\gamma}^s, \boldsymbol{\mu}^s, \boldsymbol{\Sigma}^s, \boldsymbol{\pi}^s \leftarrow$ Ops.(1-5) for k:=1 to K step 1 do $\mathbf{Y} \leftarrow \mathbf{X}_k^s - \boldsymbol{\mu}_k^s$ \triangleright Centering $\boldsymbol{\Phi}^s \boldsymbol{\Lambda}^s \boldsymbol{\Phi}^{s-1} \leftarrow \boldsymbol{\Sigma}_k^s$ \triangleright SVD $\mathbf{Z} \leftarrow \boldsymbol{\Lambda}^{s-\frac{1}{2}} \boldsymbol{\Phi}^{sT} \mathbf{Y}$ \triangleright Whitening $\boldsymbol{\Phi}^t \boldsymbol{\Lambda}^t \boldsymbol{\Phi}^{t-1} \leftarrow \boldsymbol{\Sigma}_k^t$ \triangleright SVD $\mathbf{X}_k^{new} \leftarrow \boldsymbol{\Phi}^t \boldsymbol{\Lambda}^{t\frac{1}{2}} \mathbf{Z} + \boldsymbol{\mu}_k^t$ end for $\mathbf{X}^{new} \leftarrow \sum_{k=1}^K [\boldsymbol{\gamma}_k^s \circ \mathbf{X}_k^{new}]$
---	---

Table 1. Standard Deviation (SD) and Coefficient of Variation (CV) of NMI for all five laboratories for hematoxylin and eosin dyes.

Method	Hematoxylin												Eosin	
	Lab 1		Lab 2		Lab 3		Lab 4		Lab 5		Average		SD	CV
	SD	CV	SD	CV	SD	CV	SD	CV	SD	CV	SD	CV		
Original	0.033	0.065	0.031	0.060	0.037	0.078	0.029	0.049	0.028	0.051	0.032	0.060	0.0563	0.0748
Macenko [6]	0.029	0.052	0.026	0.046	0.020	0.037	0.025	0.044	0.020	0.035	0.024	0.043	0.0362	0.0439
Reinhard [5]	0.032	0.058	0.025	0.044	0.020	0.035	0.030	0.052	0.029	0.049	0.027	0.047	0.0386	0.0494
Khan [4]	0.066	0.156	0.067	0.155	0.085	0.158	0.054	0.110	0.049	0.093	0.064	0.135	0.0434	0.0555
Vahadane [7]	0.036	0.065	0.032	0.058	0.024	0.046	0.023	0.042	0.020	0.038	0.027	0.050	0.034	0.041
Bejnordi [1]	0.016	0.029	0.015	0.027	0.018	0.034	0.029	0.055	0.024	0.044	0.021	0.038	0.0191	0.0220
Ours	0.022	0.045	0.017	0.034	0.017	0.036	0.014	0.030	0.017	0.035	0.017	0.036	0.0188	0.0209

Although the responsibility term is estimated by the CNN from only the density channel of the image, the mean and covariance matrices (μ_k and Σ_k) are estimated from all three channels of the HSD color image. The randomly initialized parameters of the CNN are updated by ADAM gradient-based optimization with a fixed learning rate of 10^{-3} . The scheme of our model is depicted in Fig. 2.

Transformation of Multivariate Gaussian Distributions: By training the model in the color normalization task, two GMMs are fitted to the source and template images, individually. Afterwards, a set of transformations are applied to align the multivariate Gaussian distributions between the source and the template. These transformations consist of three operations: *mean centering*, *whitening* and *coloring* transformation. Let us assume that (μ^s, Σ^s) and (μ^t, Σ^t) are the estimated parameters of the two distributions in source and template image, respectively. By shifting the mean of source image to the origin (*centering*) and then whitening which involves the SVD algorithm, the source image will have a zero mean and an identity covariance matrix. Consequently, by applying a coloring transformation which is the inverse of whitening but after replacing the eigenvalues (Λ) and eigenvectors (Φ) of source distribution with the template distribution, the whitened Gaussian distribution of source image scales and rotates to have the same covariance matrix as the template image (Σ^t). Finally, the distribution is shifted to obtain the same mean as the template distribution. For clarifying this procedure, pseudocode in Algorithm 1 shows these steps in detail.

3 Experimental Results

Histopathology Image Dataset: We focus on inter-laboratory variations of the H&E staining, as it is a major concern in large-scale application of CAD in pathology. For better comparison with recent studies, we use the same dataset as has been introduced in [1]. The dataset contains 625 images (each 1388×1040 pixels) from 125 digitized H&E stained WSIs of lymph nodes from 3 patients and was collected from five Dutch pathology laboratories, each using their own

routine staining protocols (more details can be found in [1]). Our model is trained on randomly cropped patches (576×576 pixels) and evaluated on full-sized images by using leave-one-out cross-validation based on the above laboratories.

Results: We trained the DCGMM on the dataset. The model easily converges in a few minutes. The inference computation time for each image in its original size (1388×1040 pixels) is about 0.6 s, implemented in the TensorFlow library and running on a TITAN Xp GPU. The performance of our method is compared to that of five competing state-of-the-art algorithms: linear appearance normalization by Macenko *et al.* [6], statistical color properties alignment by Reinhard *et al.* [5], nonlinear mapping for stain normalization by Khan *et al.* [4], sparse non-negative matrix factorization by Vahadane *et al.* [7] and WSI color standardizer by Bejnordi *et al.* [1]. The normalized median intensity (NMI) measure [1, 9] is used to evaluate color constancy of normalized images. Quantitative analysis is based on independently evaluating the color constancy in the regions that show mostly absorbed hematoxylin or eosin. Since nuclei mostly absorb hematoxylin, they first are detected automatically by using a fast radial symmetry transform and a marker-controlled watershed algorithm [1]. Since our generated hematoxylin masks slightly differ from what were used in [1], the obtained NMI scores in our benchmark are not exactly the same as [1]. However, the results are in agreement with [1]. For evaluation of the eosin analysis, several regions are manually annotated for 25 images. The evaluation results of different methods for assessing hematoxylin and eosin regions are shown in Table 1. The results clearly indicate that our proposed method results in the lowest variation in color after normalization of the images and it outperforms competing state-of-the-art methods. Figure 3 illustrates an example of the template image, a source image and the outcomes of color normalization by the different methods.




Fig. 3. Performance of different stain color-normalization methods on an *H&E* image. (a) template image (b) original images, (c) Macenko *et al.*[6], (d) Vahadane *et al.*[7], (e) Reinhard *et al.*[5], (f) Bejnordi *et al.*[1], (g) Khan *et al.* [4], (h) DCGMM.

4 Conclusions

We have introduced a framework for modeling and normalizing the colors in H&E histopathological images. Our model jointly optimizes the parameters of a CNN and the parameters of multivariate GMM in an end-to-end learning framework. By minimizing the negative log-likelihood loss function, the CNN learns how to cluster the image structures for optimally fitting the GMM on the data. Our proposal takes only one assumption on the number of clusters that is evidently chosen to three ($K = 3$) for H&E images. It does not need manual tuning of parameters and thresholds which makes it applicable to a wide range of histopathological images collected from different organs. Since our method processes and fits the GMM conditioned on the tissue classes, in comparison with previous methods, it is more robust in presence of strong stain variations.

References

1. Bejnordi, B., et al.: Stain specific standardization of whole-slide histopathological images. *IEEE Trans. Med. Imaging* **35**(2), 404–415 (2016)
2. Ciompi, F., Geessink, O., Bejnordi, B.E., et al.: The importance of stain normalization in colorectal tissue classification with convolutional networks. In: 2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017), pp. 160–163. IEEE (2017)
3. Niethammer, M., Borland, D., Marron, J.S., Woosley, J., Thomas, N.E.: Appearance normalization of histology slides. In: Wang, F., Yan, P., Suzuki, K., Shen, D. (eds.) *MLMI 2010*. LNCS, vol. 6357, pp. 58–66. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15948-0_8
4. Khan, A., Rajpoot, N., Treanor, D., Magee, D.: A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution. *IEEE Trans. Biomed. Eng.* **61**(6), 1729–1738 (2014)
5. Reinhard, E., Adhikhmin, M., Gooch, B., Shirley, P.: Color transfer between images. *IEEE Comput. Graph. Appl.* **21**(5), 34–41 (2001)
6. Macenko, M., et al.: A method for normalizing histology slides for quantitative analysis. In: *Proceedings of IEEE International Symposium on Biomedical Imaging Nano Macro*, pp. 1107–1110 (2009)
7. Vahadane, A., et al.: Structure-preserving color normalization and sparse stain separation for histological images. *IEEE Trans. Med. Imaging* **35**(8), 1962–1971 (2016)
8. Magee, D., et al.: Color normalization in digital histopathology images. In: *Proceedings of Optical Tissue Image Analysis Microscopy, Histopathology Endoscopy*, pp. 100–111 (2009)
9. Ajay Basavanahally, A.M.: Em-based segmentation-driven color standardization of digitized histopathology. *Proc. SPIE* **12**, 8676–8676-12 (2013)
10. Ruifrok, A.C., Johnston, D.A.: Quantification of histochemical staining by color deconvolution. *Analyt. Quant. Cytol. Histol. Int. Acad. Cytol. Am. Soc. Cytol.* **23**(4), 291–299 (2001)
11. Christopher, M.B.: *Pattern Recognition and Machine learning*. Springer, New York (2016)

12. Zong, B., et al.: Deep autoencoding gaussian mixture model for unsupervised anomaly detection. In: International Conference on Learning Representations (2018)
13. van den Oord, A., Schrauwen, B.: Factoring variations in natural images with deep gaussian mixture models. In: Advances in Neural Information Processing Systems, vol. 27, pp. 3518–3526. Curran Associates, Inc. (2014)
14. van der Laak, J.A., Pahlplatz, M.M., Hanselaar, A.G., de Wilde, P.C.: Hue-saturation-density (hsd) model for stain recognition in digital images from transmitted light microscopy. *Cytometry* **39**(4), 275–284 (2000)
15. Petersen, K.B., Pedersen, M.S.: The Matrix Cookbook. Technical University of Denmark, version 20121115, November 2012