

Learning Artistic Lighting Template from Portrait Photographs

Xin Jin¹, Mingtian Zhao^{2,3}, Xiaowu Chen^{1,*},
Qinping Zhao¹, and Song-Chun Zhu^{2,3}

¹ State Key Laboratory of Virtual Reality Technology and Systems,
School of Computer Science and Engineering, Beihang University, Beijing, China
jinxin@vrlab.buaa.edu.cn,
chen@buaa.edu.cn

² Lotus Hill Institute, Ezhou, China

³ Department of Statistics, University of California, Los Angeles, USA
{mtzhao, sczhu}@stat.ucla.edu

Abstract. This paper presents a method for learning artistic portrait lighting template from a dataset of artistic and daily portrait photographs. The learned template can be used for (1) classification of artistic and daily portrait photographs, and (2) numerical aesthetic quality assessment of these photographs in lighting usage. For learning the template, we adopt Haar-like local lighting contrast features, which are then extracted from pre-defined areas on frontal faces, and selected to form a log-linear model using a stepwise feature pursuit algorithm. Our learned template corresponds well to some typical studio styles of portrait photography. With the template, the classification and assessment tasks are achieved under probability ratio test formulations. On our dataset composed of 350 artistic and 500 daily photographs, we achieve a 89.5% classification accuracy in cross-validated tests, and the assessment model assigns reasonable numerical scores based on portraits' aesthetic quality in lighting.

1 Introduction

The word *photography*, first used in 1839 by Sir John Herschel, came from two Greek words *photos* (light) and *graphé* (drawing) [1]. Just like brushes and pigments of painters, light is the major tool of photographers for creating beautiful pictures. The art of lighting in photography has been so heavily influencing the aesthetics of photographs, that learning to manipulate lighting is the doorway to high-quality photography. Especially, understanding how light works and having an appreciation of good lighting is at the root position of photographer training [2]. The goal of this paper is to enable computers to *appreciate* good lighting, namely, to distinguish artistic photographs with outstanding lighting composition from daily (commonplace) ones (as shown in Fig.1), and to quantitatively assess their visual aesthetics quality in lighting usage.

* Corresponding author.




Fig. 1. Two portrait photographs (a) is usually considered as an artistic photograph, while (b) is often considered as a daily one

In recent literature, classification of images according to the quality of visual aesthetics, including photographs [3,4,5,6,7], paintings [8], etc., and numerical quality assessment of them [3,5], have attracted increasing interest. Most existing work towards this interest follows three steps: (1) collecting an image dataset according to specific objectives, and separating it into two subsets containing the “good” and “bad” images, respectively, (2) designing various features and extracting them from these collected images, and (3) training a classifier to automatically judge a test image’s quality class, or fitting a model to data with consensus scores of aesthetics quality from various training sources in order to assess the test image numerically. In early work [3,4] only global features were used, then Datta et al. [5] considered local features within image regions, and a few later studies [6,7,8] also included features extracted between regions. However, most of them did not use content-oriented features (e.g., features specially designed for portrait, landscape), which limited their performance, because aesthetic metrics usually vary a lot over diverse image contents [2,9,10,11].

As an attempt for content-oriented design towards the problem, this paper focuses on analyzing the lighting of portrait photography. Although factors such as pose, personality, expression, etc., no doubt affect the aesthetics, the use of lighting itself plays a key role. According to professional studies on portrait photography [2,9,10,11], lights and shadows on the face, with their relative locations, their area ratios, etc., are the dominant facts of attraction, which also make the main contribution to 3D perception of the 2D image on the photograph. Fig.2 includes four typical lighting styles of artistic portrait photographs [2,9,10,11]. The *Rembrandt* style is usually implemented with sidelight. It is featured by a triangular highlight below one of the eyes, with its surrounding areas in dark shadows. The *Paramount* style has a butterfly shape for the shadow between nose and mouth, which is achieved with a key light in front of and above the model’s face. The *Loop* style is named after the loop-shaped shadow below the nose. Light position for this style should be between those of Rembrandt and Paramount styles. And with a more extreme sidelight than that of the Rembrandt style, the *Split* style has half of the face in shadow. Such categorization of styles inspired us that there probably exists an *artistic lighting template* for each style, specified by the local contrast within the parts of the face. Thus we divide the area of a frontal face into 16 rectangular parts as shown in Fig.3. Choosing rectangles to represent facial parts is




Fig. 2. Artistic lighting templates. Top row: four typical artistic lighting styles of portrait photography. Middle rows: Haar-like lighting contrast filters for measuring the local contrast. Large local contrast can usually be captured by one of the filters, for example, the Rembrandt style can be captured by the center-vs-periphery filter applied below one of the eyes. Bottom row: example lighting templates for the styles, designed manually for illustration.

for computational efficiency. We then adopt Haar-like features within each rectangle to capture its local lighting contrast, which are extracted by a bank of filters of different contrast types, target channels and statistics. The last row of Fig.2 shows the corresponding example lighting templates with the most distinctive features for the above artistic lighting styles (note: features are selected manually for illustration).




Fig. 3. (a) A demonstration of professional photographers’ common practice to assess portrait photographs’ lighting quality. The face is coarsely divided into 3×5 grids according to the positions of hair, eyes, eyebrows, etc. Lights and shadows in these grids are the criteria for lighting usage. Inspired by professional photographers, we do face alignment (b) to detect feature points, then divide the face into 16 rectangular parts (c) corresponding to the nose, mouth, eyes, etc.

We learn an artistic portrait template for these styles composed of local lighting contrast features using a stepwise feature pursuit algorithm. The learning algorithm is built on a log-linear model of the probability distributions of artistic portraits, with the feature responses as factors. The learned template can be used for classifying artistic and daily portrait photographs in terms of lighting usage. Furthermore, we use human experiments to obtain consensus scores of portrait lighting to which we fit a regression model with our selected features for predicting the aesthetic quality of a portrait photograph’s lighting.

The contributions of this paper include: (1) a learned common artistic lighting template of portrait photographs, (2) a method for modeling the lighting usage of portrait photographs, with content-oriented designs of features and template, and (3) evaluations of the strengths and weaknesses of our methods and designs by applying them on a dataset containing 350 artistic and 500 daily photographs.

2 Feature Design

We adopt Haar-like features to capture the lighting characteristics on important parts of the face. Each feature F consists of four components.

1. Spatial contrast type T , which can be left-vs-right, top-vs-bottom, or center-vs-periphery, as shown in Fig.2 (above the bottom row).
2. Rectangular region R on the image lattice. According to the type T , region R is divided into two equally sized subregions R_0 and R_1 as shown in black and white areas in Fig.2.
3. Target channel C of the image. Our adopted channels include the graylevel, the three channels L , a and b in the CIE 1976 (L^*, a^*, b^*) color space (note:

L differs slightly from graylevel although the two channels are heavily correlated), and the hue and saturation channels in the HSV color space. In order to capture the effect of staggered highlight under sidelight, we also involve the channel of graylevel gradients, as well as an edge channel (see component 4 below) including all edge pixels generated by a Canny edge detector [12].

4. Target statistic S of the image. We consider 3 types of statistics: (1) mean value μ , (2) histogram h , and (3) density ρ which was specially designed for the edge channel (i.e., the proportion of edge pixels).

With the above feature design, we define the local contrast r between the statistics μ , h and ρ of the two subregions R_0 and R_1 as

$$r_\mu = |\mu_1 - \mu_0|, \quad r_h = \text{JS}(h_1||h_0), \quad r_\rho = |\rho_1 - \rho_0|, \quad (1)$$

where $\text{JS}(\cdot||\cdot)$ denotes the discrete Jensen-Shannon divergence [13]. We use r as the feature response of F in our model.

3 Template Learning

Let Ω_A denote the set of artistic portrait photographs, and Ω_D denote the set of daily ones, we would like to build a template for Ω_A against Ω_D , as well as a probability distribution upon this template. A template is a group of features which characterize the artistic photographs. Suppose the template is composed of a set of features $\{F_1, \dots, F_K\}$, a probabilistic model for each image $\mathbf{I} \in \Omega_A$ can be defined in a log-linear form [14,15] as

$$p(\mathbf{I}) = q(\mathbf{I}) \prod_{k=1}^K \frac{1}{z_k} \exp\{\lambda_k r_k(\mathbf{I})\}, \quad (2)$$

in which $q(\mathbf{I})$ is the null distribution of \mathbf{I} without any knowledge of the feature responses r_k , λ_k are weight parameters, and z_k are normalizing constants for the factors. For our case, we use the template to describe the new information of photographs in Ω_A compared with those in Ω_D , and leave the rest information in $q(\mathbf{I})$. In this way, $q(\mathbf{I})$ models the daily photographs.

We are interested in selecting an informative subset of the features for the template, rather than using the whole feature set, for two reasons: (1) features tend to be correlated, and (2) we prefer a simple template for both capability of generalization and computational efficiency. Since selecting an optimal subset of features at once is a non-trivial task, we adopt a stepwise feature pursuit algorithm [16,15] to select one feature at a time to construct $p(\mathbf{I})$.

Given a candidate feature set, we begin with an empty template corresponding to the daily photo distribution $p_0(\mathbf{I}) = q(\mathbf{I})$ at step 0. Then at each step t , we choose the max-gain feature

$$\begin{aligned} F_{(t)} &= \arg \max_{F_k} \text{KL}(p_t(r_k(\mathbf{I}))||p_{t-1}(r_k(\mathbf{I}))) \\ &\approx \arg \max_{F_k} \text{KL}(p_t(r_k(\mathbf{I}))||q(r_k(\mathbf{I}))) \\ &\approx \arg \max_{F_k} \text{KL}(h_A(r_k(\mathbf{I}))||h_D(r_k(\mathbf{I}))) \end{aligned} \quad (3)$$




Fig. 4. Template learning. For each feature F_k in the candidate feature set, its responses $r_k(\mathbf{I})$ is calculated for all images in Ω_A and Ω_D , and two histograms $h_A(r_k(\mathbf{I}))$ and $h_D(r_k(\mathbf{I}))$ are obtained for the artistic and daily photographs, respectively. Then the KL divergence between the two histograms can be computed, as an approximation to the information gain of F_k on our dataset.

on our dataset, where $\text{KL}(\cdot || \cdot)$ is the Kullback-Leibler divergence, and $h_A(\cdot)$ and $h_D(\cdot)$ are the histograms over Ω_A and Ω_D , respectively, as shown in Fig.4. In Eq.(3), the second approximation applies empirical estimates of marginal probabilities with instances in the dataset. In order for the first approximation, where we assume $p_{t-1}(r_k(\mathbf{I})) \approx q(r_k(\mathbf{I}))$, to be feasible, we apply local inhibition in the template learning process to reduce the correlations among the selected features in the pursuit steps. Noticing the fact that the rectangular regions have no or very little overlaps with each other, we simply assume independence of features in different regions, and use the inhibition strategy that selects only one feature from each rectangular part of the face. Meanwhile, for step t , the parameters $\lambda_{(t)}$ and $z_{(t)}$ can be computed by solving the system

$$\begin{aligned} \mathbb{E}_q \left[\frac{1}{z_{(t)}} \exp\{\lambda_{(t)} r_{(t)}(\mathbf{I})\} r_{(t)}(\mathbf{I}) \right] &= \mathbb{E}_{p_t}[r_{(t)}(\mathbf{I})] = \mathbb{E}_f[r_{(t)}(\mathbf{I})] \\ \mathbb{E}_q[\exp\{\lambda_{(t)} r_{(t)}(\mathbf{I})\}] &= z_{(t)} \end{aligned} \quad (4)$$

with $\mathbb{E}_q[\cdot] \approx \text{Mean}_{\Omega_D}(\cdot)$ and $\mathbb{E}_f[\cdot] \approx \text{Mean}_{\Omega_A}(\cdot)$ as empirical estimates according to our dataset. The feature pursuit process stops when the Bayesian Information Criterion (BIC) of the model reaches its minimum [17].

4 Inference for Classification and Assessment

4.1 Classification

Since the template is built for artistic photographs against daily ones, besides modeling the former, it can also be used for classifying these two types. Given a test portrait photo \mathbf{I} , we can calculate a template matching score with

$$\text{MatchingScore}(\mathbf{I}) = \log \frac{p(\mathbf{I})}{q(\mathbf{I})} = \sum_{k=1}^K (\lambda_k r_k(\mathbf{I}) - \log z_k) . \quad (5)$$

This follows a probability ratio test formulation. If the matching score is greater than a learned threshold corresponding to a certain significance level (e.g., the equal error rate (EER) threshold), the test photo is classified as an art photo, otherwise it is classified as a daily one.

This classification algorithm differs from AdaBoost [18] although the two have similar formulations. In AdaBoost, each $r_k(\mathbf{I})$ is a binary classifier, while in the above method it is a continuous feature response. Using a continuous $r_k(\mathbf{I})$ augmented from binary classifier has two advantages: (1) it can usually lead to a model with relatively smaller number of features since $r_k(\mathbf{I})$ becomes more informative, and (2) the model can be extended to predict a meaningful continuous score (see the next section).

4.2 Numerical Assessment

In addition to classification, it is usually more useful to have a reasonable numerical assessment of an input photo on its aesthetic quality. This can be achieved by extending our learning algorithm to a regression framework.

In an empirical manner, we define the quality of a portrait photograph \mathbf{I} as the probability p that it is better than another random chosen portrait photograph \mathbf{J} , namely,

$$\text{QualityScore}(\mathbf{I}) = p = E_{f(\mathbf{J})} [\mathbf{1}(\mathbf{I} \text{ wins against } \mathbf{J} \text{ in quality})] , \quad (6)$$

where f is the distribution of all portrait photographs (either artistic or daily), and $\mathbf{1}(\cdot)$ is the indicator function. For predicting the score p , we randomly choose a fixed number n of photographs to compare with \mathbf{I} , then the number of wins of \mathbf{I} should follow a binomial distribution $\text{binom}(n, p)$. We do such comparison experiments on many test images and obtain their scores, then estimate the effects of the features on the score using logistic regression [19], by fitting the model

$$\log \frac{p}{1-p} = \sum_{k=1}^K (\lambda_k r_k(\mathbf{I}) - \log z_k) = \lambda_0 + \sum_{k=1}^K \lambda_k r_k(\mathbf{I}) . \quad (7)$$

This model is able to output a score $p \in (0, 1)$ for the quality of test photographs.

5 Experiments

5.1 Data Collection

Prior to experiments on learning and evaluating the template, we collect a image dataset to work on. We collect 350 artistic portrait photographs from 3 sources: (1) masterpieces of portrait photography from famous photographers (e.g., Yousuf Karsh, Arnold Newman), (2) collections from professional photography websites (e.g., photo.net, portrait-photos.org), and (3) scanned copies from professional portrait photography books focusing on lighting [2,9,10,11].

The 500 daily photographs for comparison are obtained from 2 main sources: (1) popular photo hosting websites (e.g., flickr.com), and (2) image search engines’ results for the keywords “face” and “daily life” (e.g., images.google.com).

The collected dataset tries to randomize irrelevant factors by spanning over various poses, ethnic groups, etc. All images in our dataset are aligned using AAM [20] to a standard frontal face before entering the algorithms. Sometimes manual corrections are needed for better accuracy after automatic alignment due to shadow effects.

5.2 The Learned Template

Fig.5 displays the set of candidate features and a learned template composed of the selected 12 most informative features. The 12 rectangular parts cover most of the area on the face. Most of the selected features are of the left-vs-right spatial type. This matches the common lighting strategies for portrait photography in studios: photographers usually change the directions of light in the horizontal dimension. The parts for nose and the area between mouth and nose are the top significant areas, which makes sense as these two parts are relatively complex in geometry and have various appearances under different illumination conditions. It is worth noticing that all color features (except one on saturation) are ignored by the template learning process, which in turn proves the conjecture that lighting is the major factor of artistic portrait photography.

5.3 Classification Results

Fig.6 and Fig.7 shows the performance of our method in classifying artistic and daily portrait photographs. Running 5-fold cross-validation, an average performance is displayed with the ROC curve in Fig.6. The average classification accuracy is 89.5% in these experiments using the EER threshold, which is quite good considering the relatively small training sample size. Look at the failure cases in Fig.7—some are indeed hard to classify unless information such as poses and expressions is involved.

5.4 Numerical Assessment Results

For the numerical assessment task, training data with the quality of photographs are necessary for fitting the regression model (see Section 4.2). We obtain the




Fig. 5. (a) The feature set with 224 candidate features: for each of the 16 rectangular parts of the faces, we have 14 local contrast filters. (b) A learned common artistic portrait lighting template composed of the 12 most informative features, shown with ranks marked by colored boundaries corresponding to the legend. The target channel of the selected features are also marked. As for target statistics, all selected features are on means.




Fig. 6. The classification performance of the log-linear model with a fixed number of 12 features. The ROC curve displays the average performance in the 5-fold cross-validated tests.



Fig. 7. Examples of classification results

Table 1. Logistic Regression Coefficients

Feature	Est.	Std.Err.	z -score	p -value	Feature	Est.	Std.Err.	z -score	p -value
(I)	-2.04	0.14	-14.82	< 0.01	F_7	1.99	0.46	4.27	< 0.01
F_1	8.39	1.29	6.51	< 0.01	F_8	4.45	0.69	6.45	< 0.01
F_2	0.05	0.74	0.07	0.95	F_9	2.06	0.74	2.78	0.01
F_3	-3.47	1.98	-1.75	0.08	F_{10}	-1.88	0.66	-2.84	< 0.01
F_4	-0.09	0.86	-0.11	0.91	F_{11}	-0.22	0.33	-0.66	0.51
F_5	1.57	0.61	2.57	0.01	F_{12}	-4.04	0.77	-5.26	< 0.01
F_6	0.32	0.80	0.40	0.69					

consensus quality scores (i.e., winning probabilities against randomly chosen images) with human experiments.

Human Experiments. We randomly choose 50 photographs (either artistic or daily) from our dataset as training examples. We also choose 10 graduate students of various majors as test subjects to do the comparisons between photographs.

For each \mathbf{I} in the 50 training images, another 50 random images $\mathbf{J}_1, \dots, \mathbf{J}_{50}$ from the rest of the dataset are sampled with replacement. Each of the 50 pairs $(\mathbf{I}, \mathbf{J}_1), \dots, (\mathbf{I}, \mathbf{J}_{50})$ is displayed to a random test subject, who will compare the two images and report their relative rank order in the quality of lighting. In this way, a total of 2500 comparisons are performed, and the numbers of wins and losses for the 50 training images are obtained for fitting the prediction model in Section 4.2. Here the replacement ensures the constant probability condition for assuming a binomial distribution [19].

Regression. Most features have favorable small p -values in the logistic regression (see Table 1). Fig.8 plots the fitted scores vs. the scores obtained from human experiments. Due to the small size of the training set, and the small




Fig. 8. Goodness-of-fit visualization of the logistic regression for quality assessment




Fig. 9. Examples of numerical assessment: the top-10, middle-10, and bottom-10 photographs are displayed with predicted quality scores

number of Bernoulli trials, there are still a considerable residual deviance in the fitting (null deviance: 936.85 on 49 degrees of freedom, residual deviance: 398.50 on 37 degrees of freedom). But since no heteroskedasticity or non-normal effects are noticed [19], we believe that our empirical definition for the aesthetic quality and the experimental design should make good sense.

Fig.9 shows examples of prediction, which includes the top-10, middle-10, and bottom-10 photographs predicted by our regression model, for the rest of the photographs in our dataset besides the training examples. For the two photographs in Fig.1, the predicted scores are (a) 0.781 and (b) 0.157, respectively.

6 Discussions

In previous sections, we have demonstrated the capabilities of our template-based model for classification and assessment despite its simplicity. For further diagnosis of the method and potential improvements, a few important aspects need to be considered.

Template and Feature Capabilities. Our common template can be understood as an “average” template of various art lighting styles. In Fig.7, the results show that the learned common template is not quite suitable for the Paramount style, which however can deal with Rembrandt, Split and Loop styles very well. The reason for this might be that the learned features’ contrast types in our template are mostly left-vs-right. In practice, photographers use a key light above and in front of the model to form a butterfly shaped shadow between mouth and nose. This can usually be better captured by top-vs-down features. But in the training process, this effect is often submerged by the other styles. This problem can be partially solved by learning a generative template for each style, in other words, we believe it is necessary to learn a mixture model for artistic portraits, which is more effective and efficient. Besides, if we trade computational efficiency for more complex feature shapes (e.g., triangle, ellipse), we can expect to model the styles like Rembrandt and Paramount better.

Local Inhibition Strategy. We apply local inhibition for reducing the correlations among selected features, which is necessary for the stepwise pursuit algorithm. But the inhibition strategy that allows only one feature in each rectangular part is imperfect, for example, many informative features might be missed during the process. If it can be done with an acceptable computational cost, a pre-computation for the feature correlations on our dataset is helpful. But the scalability of this method is limited by frequent modifications of the dataset and candidate features.

Interpreting the Regression Model. Look at the regression coefficients in Table 1. The coefficient for F_1 is approximately 8.39 with a standard error of 1.29. For this coefficient, the underlying meaning of the numbers is that the larger the local contrast is, the better the photo quality becomes. This does not make much sense as it contradicts with some photographs in our dataset, for example, the demo photo for the Paramount style. To better understand the contribution of each feature, Fig.10 displays part of the estimated coefficients (for the intercept and first two features only) using linear quantile regressions [21] which fit the response $\log(\#_{\text{win}}/\#_{\text{loss}})$, from 5% to 95% quantiles with a step of 5%. The plots reflect the non-linear contributions of the features. For example, F_2 has a more obvious positive contribution thus is more important for low-quantile (daily) photographs than for high-quantile (artistic) ones, although its average effect over various images is small (see the point estimate of its coefficient in Table 1). Similar non-linear effects for the other features are also observed.




Fig. 10. Coefficient plots of linear quantile regressions, in which the predictors are feature values, and the response is $\log(\#\text{win}/\#\text{loss})$. The horizontal axis is the response quantile, and the vertical axis is the coefficient of each feature. The black curves are quantile regression estimates, and the dashed red lines are for ordinary least-squares regression estimates.

7 Conclusion

In this paper, we learn an artistic lighting template for classification and assessment of artistic and daily portrait photographs in lighting. Using Haar-like features applied on rectangular areas of the face, our method generates a template with justifiable interpretations. We use a log-linear probabilistic model to characterize artistic portrait photographs against daily ones, which gives satisfactory performance in both of the two tasks. Potential improvements of this method include more expressive features, better feature selection strategy with reduced assumptions, and more powerful statistical models, etc.

Acknowledgments

This work was partially supported by National Natural Science Foundation of China (90818003&60933006), National High-Tech (863) and Key Technologies R&D Programs of China (2009AA01Z331&2008BAH29B02), National Grand Fundamental Research (973) Program of China (2006CB303007), and Specialized Research Fund for the Doctoral Program of Higher Education (20091102110019).

References

1. Wikipedia, <http://en.wikipedia.org/wiki/photography>
2. Hurter, B.: The best of photographic lighting — techniques and images for digital photographers, 2nd edn. Amherst Media (2007)
3. Tong, H., Li, M., Zhang, H., He, J., Zhang, C.: Classification of digital photos taken by photographers or home users. *PCM* (1), 198–205 (2004)
4. Ke, Y., Tang, X., Jing, F.: The design of high-level features for photo quality assessment. In: *CVPR*, pp. 419–426 (2006)

5. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Studying aesthetics in photographic images using a computational approach. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) ECCV 2006. LNCS, Part 3, vol. 3953, pp. 288–301. Springer, Heidelberg (2006)
6. Luo, Y., Tang, X.: Photo and video quality evaluation: Focusing on the subject. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part III. LNCS, vol. 5304, pp. 386–399. Springer, Heidelberg (2008)
7. Wong, L.K., Low, K.L.: Saliency-enhanced image aesthetics class prediction. In: ICIP (2009)
8. Li, C., Chen, T.: Aesthetic visual quality assessment of paintings. *IEEE Journal of Selected Topics in Signal Processing* 3, 236–252 (2009)
9. Hunter, F., Biver, S., Fuqua, P.: *Light: Science and Magic: An Introduction to Photographic Lighting*, 3rd edn. Focal Press (2007)
10. Grey, C.: *Master Lighting Guide for Portrait Photographers*. Amherst Media (2004)
11. Praker, D.: *Basics Photography: Lighting*. AVA Publishing (2007)
12. Canny, J.F.: A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 8, 679–714 (1986)
13. Lin, J.: Divergence measures based on the shannon entropy. *IEEE Trans. Info. Theory* 37 (1991)
14. Della Pietra, S., Della Pietra, V., Lafferty, J.: Inducing features of random fields. *IEEE Trans. Pattern Anal. Mach. Intell.* 19, 380–393 (1997)
15. Si, Z., Gong, H., Wu, Y.N., Zhu, S.C.: Learning mixed templates for object recognition. In: CVPR, pp. 272–279 (2009)
16. Friedman, J.H.: Exploratory projection pursuit. *Journal of American Stat. Assoc.* 82, 249–266 (1987)
17. Schwarz, G.: Estimating the dimension of a model. *Ann. Statist.* 6 (1978)
18. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55, 119–139 (1997)
19. Faraway, J.J.: *Extending the Linear Model with R*. Taylor & Francis Group, Abington (2006)
20. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. In: Burkhhardt, H., Neumann, B. (eds.) ECCV 1998. LNCS, Part 2, vol. 1407, pp. 484–498. Springer, Heidelberg (1998)
21. Koenker, R.: *Quantile Regression*. Cambridge University Press, Cambridge (2005)