

FISH: Face Intensity-Shape Histogram Representation for Automatic Face Splicing Detection

Marco Fanfani^{a,*}, Fabio Bellavia^a, Massimo Iuliani^{a,b},
Alessandro Piva^{a,b}, Carlo Colombo^a

^a*Dept. of Information Engineering, University of Florence, Florence, Italy*
^b*FORLAB Multimedia Forensics Laboratory, University of Florence, Prato, Italy*

Abstract

Tampered images spread nowadays over any visual media influencing our judgement in many aspects of our life. This is particularly critical for face splicing manipulations, where recognizable identities are put out of context. To contrast these activities on a large scale, automatic detectors are required.

In this paper, we present a novel method for automatic face splicing detection, based on computer vision, that exploits inconsistencies in the lighting environment estimated from different faces in the scene. Differently from previous approaches, we do not rely on an ideal mathematical model of the lighting environment. Instead, our solution, built upon the concept of histogram-based features, is able to statistically represent the current interaction of faces with light, untied from the actual and unknown reflectance model. Results show the effectiveness of our solution, that outperforms existing approaches on real-world images, being more robust to face shape inaccuracies.

Keywords: Image Forensics, Scene level analysis, Geometric Constraints, Lighting environment, Face splicing detection.

1. Introduction

2 Manipulated images are becoming ubiquitous in everyday life. Thanks to
3 the advancement of photo-editing software, highly realistic tampering can be

*Corresponding author

Email address: marco.fanfani@unifi.it (Marco Fanfani)

4 produced even by non-expert users, with deep social impact and critical conse-
5 quences in our perception of reality. In order to detect and contrast the spread
6 of these fake images, Image Forensics has developed several solutions [1] aimed
7 at determining if an image is pristine or tampered according to the presence,
8 absence or inconsistency between the traces left by operations such as image
9 acquisition, compression and other editing processes.

10 Face splicing, achieved by inserting into an original image a human face
11 retrieved from a different photo, is one of the most critical tampering since
12 it deals with people identity and can be used to produce images where spe-
13 cific subjects are inserted into an inconvenient and awkward context. *Signal*
14 *level* traces found as invisible footprints into the signal statistics, such as de-
15 mosaicing [2] or compression [3, 4] artifacts, or noise [5, 6], can be employed to
16 detect face splicing. Unluckily, these solutions have a limited applicability, since
17 the abovementioned traces may be partially or completely spoiled by common
18 operations on images, such as resizing, compression, etc. [7]. More recently, al-
19 ternative methods based on deep-learning [8] or exploiting the inconsistencies at
20 the *physical level* of the scene represented in the image have arisen, considering
21 shadows [9], perspective [10], or lighting [11, 12] incongruities.

22 In this paper we present a novel technique to detect face splicing based on
23 physical-level analysis of the imaged scene. Previous works exploiting physical
24 traces in the image try to directly extract and estimate the lighting parameters
25 (i.e., the light source position, color and intensity) on each single face in the
26 image, from which to detect inconsistencies indicating possible tampering. The
27 major novelties of our approach are:

- 28 • Instead of a complex and partially incomplete ideal model characterizing
29 the interaction of light with faces, we propose to employ histogram-based
30 features. Histograms have proved to be very effective in many computer
31 vision tasks [13] and, to the best of our knowledge, were never employed
32 for face splicing detection;
- 33 • Novel ad-hoc metrics to compute distances between FISH features have

34 also been designed, taking into account pixel saturation and albedo differ-
35 ences, so as to further improve the accuracy of our face splicing detector;

- 36 • Since our face features depend only on some image data statistics, without
37 focusing on a particular mathematical model, on real images they outper-
38 form the state-of-the-art approach of [14] due to their higher robustness
39 against image noise and face shape estimation inaccuracies;
- 40 • Finally, our approach is computationally more efficient, since it relies only
41 on histogram computation, while the state-of-the-art requires complex face
42 and lighting renderings.

43 The paper is organized as follows: In the next Section, a brief overview
44 of the state-of-the-art methods is presented. The proposed histogram-based
45 representation is described in Sect. 3, and used as the main building block
46 for the fully automatic pipeline of Sect. 4. An experimental evaluation of our
47 approach is reported in Sect. 5, and conclusions are finally drawn in Sect. 6.

48 **2. State of the Art**

49 Estimating the light source parameters of a real scene is quite a challenging
50 task [15] which can prove extremely useful for detecting tampered images. In
51 the recent literature on image forensics, some methods aim to detect image
52 inconsistencies by estimating the color of the light source (i.e. the illuminant),
53 while others focus on fitting a parametric model describing the interaction of the
54 light source with the environment, for which the light source location/direction
55 is usually the most relevant parameter.

56 The estimation of the light source color is strictly connected with the *colour*
57 *constancy* problem [16], that requires to subtract the real light color from the
58 input image in order to make the scene appear as it was acquired under a white
59 illuminant. In the case of forensic applications, features related to light color are
60 extracted on several patches of the images using the Gray-World assumption [17,
61 18], or physical-based solutions like the Inverse Intensity-Chromaticity [19] and

62 compared across the image looking for anomalies. In particular, in [11] a SVM
 63 classifier is trained on features extracted from an illuminant map (i.e., a super-
 64 pixel tessellation of the image, associating each patch to its illuminant color)
 65 computed by solving the color constancy problem.

66 Parametric models describing the interaction between light and the environ-
 67 ment are based on the spherical harmonics representation [20, 21]. In particular,
 68 under the assumption of convex Lambertian surfaces with uniform albedo, lin-
 69 ear camera response and distant light sources, for each color channel the light
 70 intensity $I(\mathbf{x}_k)$ measured at pixel \mathbf{x}_k can be modeled as a linear combination
 71 of the spherical harmonics $\{Y_{n,m}(N(\mathbf{X}_k))\}$. Up to the second order, these are
 72 evaluated as

$$\begin{aligned}
 Y_{0,0}(N(\mathbf{X}_k)) &= \frac{1}{\sqrt{4\pi}} & Y_{1,-1}(N(\mathbf{X}_k)) &= \sqrt{\frac{3}{4\pi}} y_k \\
 Y_{1,0}(N(\mathbf{X}_k)) &= \sqrt{\frac{3}{4\pi}} z_k & Y_{1,1}(N(\mathbf{X}_k)) &= \sqrt{\frac{3}{4\pi}} x_k \\
 Y_{2,-2}(N(\mathbf{X}_k)) &= 3\sqrt{\frac{5}{12\pi}} x_k y_k & Y_{2,-1}(N(\mathbf{X}_k)) &= 3\sqrt{\frac{5}{12\pi}} y_k z_k \\
 Y_{2,0}(N(\mathbf{X}_k)) &= \frac{1}{2}\sqrt{\frac{5}{4\pi}} (3z_k^2 - 1) & Y_{2,1}(N(\mathbf{X}_k)) &= 3\sqrt{\frac{5}{12\pi}} x_k z_k \\
 Y_{2,2}(N(\mathbf{X}_k)) &= \frac{3}{2}\sqrt{\frac{5}{12\pi}} (x_k^2 - y_k^2) & &
 \end{aligned} \tag{1}$$

73 In the above formulation, the pixel \mathbf{x}_k is the projection of a surface 3D point
 74 \mathbf{X}_k , with normal $N(\mathbf{X}_k) = [x_k, y_k, z_k]$. The coefficients up to the second order
 75 of the spherical harmonics, i.e. $\ell_{n,m}$ with $n = \{0, 1, 2\}$ and $m = \{-n, n\}$, almost
 76 uniquely identify the lighting environment. In order to estimate them, the linear

77 system $M\ell = \mathbf{I}$, or explicitly

$$\begin{bmatrix} r_{0,0}(N(\mathbf{X}_1)) & r_{1,-1}(N(\mathbf{X}_1)) & \dots & r_{2,2}(N(\mathbf{X}_1)) \\ r_{0,0}(N(\mathbf{X}_2)) & r_{1,-1}(N(\mathbf{X}_2)) & \dots & r_{2,2}(N(\mathbf{X}_2)) \\ \vdots & \vdots & \ddots & \vdots \\ r_{0,0}(N(\mathbf{X}_K)) & r_{1,-1}(N(\mathbf{X}_K)) & \dots & r_{2,2}(N(\mathbf{X}_K)) \end{bmatrix} \begin{bmatrix} \ell_{0,0} \\ \ell_{1,-1} \\ \vdots \\ \ell_{2,2} \end{bmatrix} = \begin{bmatrix} I(\mathbf{x}_1) \\ I(\mathbf{x}_2) \\ \vdots \\ I(\mathbf{x}_K) \end{bmatrix} \quad (2)$$

78 is solved, where $r_{0,0}(N(\mathbf{X}_k)) = \pi Y_{0,0}(N(\mathbf{X}_k))$, $r_{1,m}(N(\mathbf{X}_k)) = \frac{2\pi}{3} Y_{1,m}(N(\mathbf{X}_k))$,
 79 $r_{2,m}(N(\mathbf{X}_k)) = \frac{\pi}{4} Y_{2,m}(N(\mathbf{X}_k))$, and $K \geq 9$ pixel sampling locations \mathbf{x}_k are used.

80 A possible splice is noticed when, in the same image, lighting coefficients
 81 relative to different parts of the scene exhibit relevant differences. In particular,
 82 lighting coefficients are estimated from occluding boundaries in [12], and from
 83 human faces in [22, 23, 14], after retrieving their 3D shape. To the best of
 84 our knowledge, the complex model described in [14], enriched to overcome the
 85 strict assumptions behind the spherical harmonics representation given above,
 86 is the current state-of-the-art in face splicing based on lighting observations.
 87 However, it still shows the main drawbacks inherent in retrieving the spherical
 88 lighting coefficients. More specifically, light estimation is very sensitive to the
 89 shape accuracy of the object upon which the matrix M is computed, i.e., the
 90 normals of the sampled points. This makes the solution very unstable, as can
 91 be noted by the performance degradation from synthetically rendered faces to
 92 real faces [14], for which the 3D shape is usually obtained automatically using
 93 morphable models [24, 25] or, more recently, deep learning [26]. Furthermore,
 94 still in the case of faces from real images, the advantages of using complex
 95 lighting models over simple ones are quite negligible.

96 According to these observations, and considering the difficulty in obtaining
 97 more accurate 3D models, in this paper we propose a different approach to
 98 face splicing based on an indirect estimation of the lighting map. In particular,
 99 instead of computing analytically the lighting coefficients, we build histograms
 100 relating surface normals with their intensity values, by statistically modelling the

101 interaction map between light and the surface. The resulting descriptor design is
 102 inspired by histogram-based keypoint descriptors [13] employed in robust image
 103 matching. Indeed, the histograms associated to different faces are stable and
 104 robust to shape variations, and can be successfully used to indirectly measure
 105 lighting inconsistencies between spliced and pristine faces.

106 3. Face Intensity-Shape Histogram (FISH)

107 Under the assumption of convex and Lambertian surfaces with fixed albedo
 108 and distant light sources, the image intensity values of points in the scene only
 109 depend on their associated surface normals. In the case of faces, the resulting
 110 channel-wise mapping function $L : \mathbb{R}^3 \rightarrow \mathbb{R}$ from normals $\mathbf{n} = [x \ y \ z]^T$, $z > 0$ to a
 111 color channel intensity of the image $I = L(\mathbf{n})$ can be statistically modelled using
 112 a histogram-based representation, referred to as *Face Intensity-Shape Histogram*
 113 (FISH), computed as follows.

114 Given a face in the image and its associated 3D shape model (see Fig. 1a
 115 and 1b, respectively), we first pre-process the model so as to remove face regions
 116 strongly violating the assumptions above (see Fig. 1c). These regions include
 117 neck and ears (that yield poorly estimated normals), mouth, eyes and eyebrows
 118 (that have a different albedo and reflectance with respect to face skin), and
 119 saturated areas (i.e., pixels with maximum intensity among all channels out of
 120 the range [15, 240] for 8-bit RGB images).




Figure 1: (a) Detected face; (b) Registered 3D shape (using 3DMM); (c) Masked 3D shape;
 (d) FISH (best viewed in color).

121 FISH bins $i = 0, \dots, \mathfrak{B}$ are sampled according to the vertices of a semi-

122 icosphere, that approximates a semi-sphere limited to the positive z -axis by a
 123 simplicial polyhedron at subdivision level 3 (i.e. an icosphere). Since an ico-
 124 sphere has 642 vertexes, of which only 305 with strictly positive z coordinate, it
 125 holds $\mathfrak{B} = 304$. Each bin corresponds to a distinct *quantized* surface normal \mathbf{n}_i
 126 (see Fig. 1d). FISH bin values $I_i = L(\mathbf{n}_i)$ for each color channel are computed
 127 via Gaussian kernel density estimation as explained hereafter. Let $\hat{\mathbf{n}}_k = N(\mathbf{X}_k)$
 128 and $\hat{I}_k = I(\mathbf{x}_k)$ be respectively the 3D shape normal vector of \mathbf{X}_k and the in-
 129 tensity value of a pixel \mathbf{x}_k , which is the projection of \mathbf{X}_k as in Eq. 2. (Notice
 130 that index i refer to bins, while index k to pixels/normals sampled on the face.)
 131 Then

$$I_i = \sum_k \frac{w_{ik}}{w_i} \hat{I}_k \quad (3)$$

132 where the sum is over the masked face pixels, with weights

$$w_i = \sum_k w_{ik} \quad (4)$$

133 computed from the Gaussian distribution

$$z_{ik} = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2} \left(\frac{\arccos(\mathbf{n}_i \cdot \hat{\mathbf{n}}_k)}{2\sigma} \right)^2} \quad (5)$$

134 subject to a influence cutoff threshold τ_k :

$$w_{ik} = \begin{cases} z_{ik} & \text{if } z_{ik} > \tau_k \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

135 The value of τ_k corresponds to the 2.5th percentile of the distribution of the
 136 z_{ik} , for $i = [0, \dots, \mathfrak{B}]$. In this way, weights associated to normals $\hat{\mathbf{n}}_k$ that are
 137 too far from the i^{th} bin representative \mathbf{n}_i are forced to zero. The standard
 138 deviation σ used to define the kernel bandwidth in Eq. 5 is equal to 3/8 times
 139 the average angular distance between two adjacent vertexes of the icosphere.

140 By concatenating the bin values for each channel, i.e.,

$$\mathbf{I} = \mathbf{L}(\mathbf{n}_i) = [L_R(\mathbf{n}_i) L_G(\mathbf{n}_i) L_B(\mathbf{n}_i)]^T \quad (7)$$

141 the final FISH descriptor \mathbf{L} is obtained.

142 FISH descriptors can be used to compare faces in a probe image. The more
 143 two FISH descriptors are similar, the more the corresponding faces are likely to
 144 be exposed to the same lighting conditions. A possible definition of the distance
 145 $\mathcal{D}(a, b)$ between two FISH descriptors \mathbf{L}^a and \mathbf{L}^b associated to faces a and b is

$$\mathcal{D}(a, b) = \left(\sum_{\substack{i=0, \dots, \mathfrak{B} \\ (w_i^a > 0) \wedge (w_i^b > 0)}} \|\mathbf{I}_i^a - \mathbf{I}_i^b\|^2 \right)^{\frac{1}{2}} \quad (8)$$

146 where $\mathbf{I}_i^a = \mathbf{L}^a(\mathbf{n}_i)$, $\mathbf{I}_i^b = \mathbf{L}^b(\mathbf{n}_i)$, $\|\cdot\|$ is the Euclidean norm L_2 —chosen
 147 experimentally, as it gives the best results among L_1 , L_2 , Wave edges, Canberra,
 148 Correlation, Bhattacharyya and Kullback Leibler—and w_i^a , w_i^b are defined as in
 149 Eq. 4. Notice that the above definition of $\mathcal{D}(a, b)$ takes explicitly into account
 150 the presence of empty histogram bins.

151 As shown in Fig. 2, unhandled skin albedo would result in an incorrect
 152 FISH-based face matching.




Figure 2: (b),(c): In the absence of skin tone normalization, the FISH descriptors for two faces in a pristine image (a) look different to each other, while they should not. (Best viewed in color.)

153 In order to remove skin color effects when comparing two FISH descriptors
 154 \mathbf{L}^a and \mathbf{L}^b , we developed and tested two normalization strategies. The first

155 strategy consists of simply pre-normalizing \mathbf{L} by the mean RGB value $\boldsymbol{\mu}$ of the
 156 associated masked face, under the common assumption that albedo is a scale
 157 factor, i.e.,

$$\mathring{I}_i = \mathring{\mathbf{L}}(\mathbf{n}_i) = \mathbf{L}(\mathbf{n}_i)/\boldsymbol{\mu} \quad (9)$$

158 channel-wise, so that

$$\mathcal{D}'(a, b) = \mathcal{D}(\mathring{\mathbf{L}}^a, \mathring{\mathbf{L}}^b) \quad . \quad (10)$$

159 In the second strategy, color saturation is taken into account. In detail, the
 160 FISH descriptor \mathbf{L}^a is normalized with respect to its albedo μ_a , then the albedo
 161 μ_b of \mathbf{L}^b is applied, clipping saturated values, i.e.,

$$I_i^{a \rightarrow b} = \mathbf{L}^{a \rightarrow b}(\mathbf{n}_i) = \min(255, \mathbf{L}^a(\mathbf{n}_i) \frac{\mu_b}{\mu_a}) \quad (11)$$

162 $\mathbf{L}^{a \rightarrow b}$ is then compared with \mathbf{L}^b . The final distance is made symmetric by also
 163 considering the case in which the μ_a is applied to \mathbf{L}^b , so that

$$\mathcal{D}''(a, b) = \min(\mathcal{D}(\mathbf{L}^a, \mathbf{L}^{b \rightarrow a}), \mathcal{D}(\mathbf{L}^b, \mathbf{L}^{a \rightarrow b})) \quad (12)$$

164 Referring to Fig. 3 we present an example of both normalization strategies for
 165 the faces of the pristine image in Fig. 2. Fig. 3a and 3d show the FISHs \mathbf{L}^a and
 166 \mathbf{L}^b without any normalization: Their comparison produces a distance of 65.42.
 167 In this case, the effect of the skin color strongly affects the distance, introducing
 168 a bias related to the face albedo. This can be suppressed by normalizing each
 169 descriptor with its mean RGB value, thus obtaining the FISHs $\mathring{\mathbf{L}}^a$ and $\mathring{\mathbf{L}}^b$,
 170 shown in Fig. 3b and 3e. Comparing these normalized descriptors yields a
 171 distance of 25.07. However, $\mathring{\mathbf{L}}^a$ and $\mathring{\mathbf{L}}^b$ cannot take into account saturated
 172 values that go outside the range $[0, 255]$. In this case, using the FISH descriptors
 173 $\mathbf{L}^{a \rightarrow b}$ (Fig. 3c) and $\mathbf{L}^{b \rightarrow a}$ (Fig. 3f) can handle this saturation side-effects. In
 174 particular, to compute $\mathcal{D}''(a, b)$, we first evaluate the distance between \mathbf{L}^b and
 175 $\mathbf{L}^{a \rightarrow b}$ (i.e. Fig. 3d and 3c), and between \mathbf{L}^a and $\mathbf{L}^{b \rightarrow a}$ (i.e. Fig. 3a and 3f), and




Figure 3: Normalized descriptor obtained from the pristine image of Fig. 2. While the distance \mathcal{D} without any normalizations ((a) and (d)) obtains a score of 65.42, \mathcal{D}' ((b) and (e)) lowers the score to 25.07. Finally, \mathcal{D}'' obtains 7.02 as the minimum between 13.45 (from (a) and (f)) and 7.02 (from (c) and (d)). (Best viewed in color.)

176 then we select the minimum among the two distances, that in this case is 7.02.

177 Figure 4 shows an example face, together with results synthesized from the
 178 inverse mapping of the FISH model and from the spherical harmonics coefficients
 179 obtained as described in [27]. Since the FISH model preserves better shading
 180 details than the spherical harmonics model, FISH fits better real data, which also
 181 implies an implicit relaxation of the strict assumptions defining the interaction
 182 of light with the environment.




Figure 4: Examples of inverse synthesized face. (a) Original image; (b) Masked face; (c) FISH reverse mapping synthesis; (d) spherical harmonics synthesis. (Best viewed in color.)




Figure 5: Pipeline for automatic face splicing detection using FISH descriptors.

183 4. Automatic face splicing detection pipeline

184 We employed the FISH descriptor to develop a fully automated pipeline for
 185 face splicing detection, that can be divided into the following three steps (see
 186 Fig. 5):

- 187 • **Face detection.** The method proposed in [28] is used, which exploits
 188 general Deformable Part Models trained to specifically detect faces. Sub-
 189 parts of the object are detected by taking into account the deformation
 190 with respect to a mean shape (detection threshold is set to 0.3). From
 191 each detected face region, 68 face landmarks are successively localized
 192 according to the face alignment algorithm of [29], based on Supervised
 193 Descent Method, used with the default parameters.
- 194 • **Face shape and normals estimation.** Face landmarks computed at
 195 the previous step are used to register a *3D Morphable Model* (3DMM)
 196 and to obtain an estimate of the face shape. In particular, we adopted the
 197 solution presented in [30], combining the *Basel Face Model* [24] and the
 198 *Face Warehouse* model [25] in order to be able to adapt the model to both
 199 identity and expression. As an alternative approach, we also tested the
 200 recent method proposed in [26] based on convolutional neural networks.
- 201 • **FISH descriptors extraction and comparison.** See Sect. 3.

202 Note that, since our method, as well as [14] and [27], compares lighting esti-
 203 mates to detect discrepancies, at least two faces are required. Moreover, in the

204 case that only two faces are detected, the pipeline can detect the occurrence
205 of tampering, but is unable to indicate which of the two is the tampered face,
206 while, if more than two faces are found, the spliced face can be localized as the
207 one with the greatest distance in terms of FISH descriptors from the other faces.
208 Notice also that it is assumed that all the subjects under analysis are subjected
209 to the same lighting environment.

210 5. Experimental evaluation

211 In order to gain a deep insight into FISH performance, several comparative
212 tests were carried out using different datasets that cover increasing levels of
213 complexity, from a fully synthetic setup (Sect. 5.1), through a controlled face
214 acquisition setup with manual 3D model estimation (Sect. 5.2), to a real-world,
215 unconstrained scenario (Sect. 5.3).

216 5.1. Synthetically generated faces

217 This evaluation employs the Syn1 and Syn2 datasets, presented in [14], where
218 two sets of 3D synthetic faces have been rendered with known random lights.
219 Since FISH does not compute spherical harmonics, a direct estimation of the
220 error in terms of lighting coefficients as in [14] cannot be done. Nevertheless,
221 a higher distance between the related FISH descriptors must be expected as
222 the discrepancy in two lighting environments increases. Under this observation,
223 the correlation between the difference of two ground-truth spherical harmonics
224 vectors, corresponding to the two faces to be checked, and the distance of the
225 related FISH descriptors, provides a good indicator of the method accuracy.
226 For this scope, we created *virtually* spliced probes by considering two faces with
227 different lighting, and evaluated the correlation between the scores obtained
228 by FISH and the ground-truth values in terms of Spearman’s rank correlation
229 coefficient (SROCC). Additionally, in order to evaluate the method robustness
230 w.r.t. noise in the images and in the 3D shape estimates, the evaluation was
231 repeated by injecting Gaussian noise with zero mean and variable standard

232 deviation σ . In particular, a Gaussian noise with $\sigma_{\text{RGB}} = \{5, 7\}$ was added
 233 to each RGB channel independently, and similarly a Gaussian noise with $\sigma_{\text{N}} =$
 234 $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ was added to each normal vector dimension independently.

235 Table 1 reports the results obtained by FISH and the baseline method of [27].
 236 For our pipeline using FISH descriptors, the superscript ‘†’ (i.e. FISH[†]) indicates
 237 that no mask is applied to the saturated pixels.

Table 1: SROCC on Syn1 Syn2 (best results in bold)

Method	Original	Image noise		Shape noise				
		$\sigma_{\text{RGB}} = 5$	$\sigma_{\text{RGB}} = 7$	$\sigma_{\text{N}} = 0.1$	$\sigma_{\text{N}} = 0.2$	$\sigma_{\text{N}} = 0.3$	$\sigma_{\text{N}} = 0.4$	$\sigma_{\text{N}} = 0.5$
FISH [†] with \mathcal{D}	0.7639	0.7639	0.7636	0.7670	0.7492	0.7170	0.6738	0.6191
FISH [†] with \mathcal{D}'	0.8625	0.8626	0.8620	0.8608	0.8457	0.8278	0.8057	0.7941
FISH [†] with \mathcal{D}''	0.8544	0.8545	0.8538	0.8484	0.8288	0.8077	0.7846	0.7673
FISH with \mathcal{D}	0.7639	0.7639	0.7636	0.7671	0.7491	0.7170	0.6738	0.6192
FISH with \mathcal{D}'	0.8627	0.8628	0.8621	0.8609	0.8459	0.8278	0.8059	0.7940
FISH with \mathcal{D}''	0.8543	0.8545	0.8539	0.8485	0.8289	0.8078	0.7846	0.7672
Kee & Farid [27]	0.8131	0.8135	0.8137	0.8183	0.8127	0.7896	0.7557	0.7365

238 As shown in the table, FISH correlation with light coefficients is high, in
 239 particular using the distance normalization schemes \mathcal{D}' , and \mathcal{D}'' . FISH with
 240 distance normalizations has better correlation than the baseline spherical har-
 241 monics estimation method of [27] also when noise is added. Note that FISH
 242 and FISH[†] obtain very close results, since for these images no saturated pix-
 243 els are present (i.e. there are not highlights or strong shadows). Results with
 244 the method of [14] are not reported in Table 1 since nothing can actually be
 245 said about the behavior of this approach in the presence of noise. Indeed, this
 246 method does not use the normal vectors directly: It requires to render the
 247 face 3D model on 42 images with different lightings and estimate the optimized
 248 transfer coefficients. This can only be done with the knowledge of additional
 249 data, unavailable to us. If no noise is present, the solution of [14] obtains a very
 250 high correlation value (0.9592), thanks to the availability of the original true 3D
 251 face model for the rendering process, which actually is an unrealistic scenario
 252 in practical situations.

253 *5.2. Real faces in a controlled acquisition setup*

254 For this test, the Yale Face Database B (YaleB) [31] was used, that includes
 255 a set of images obtained from 10 distinct faces captured in different poses under

256 49 different lighting conditions. Following [14], we focused on frontal faces, thus
 257 reducing the dataset to 490 test images. Analogously to the previous experimen-
 258 tal evaluation on Syn1 and Syn2, a *virtually* spliced dataset was generated by
 259 considering for the negative (pristine) set all the face pairs of different identities
 260 with the same lighting, obtaining $(49 \times 10 \times 9)/2 = 2205$ pristine images. On
 261 the other hand, there are $(49 \times 10 \times 48 \times 9)/2 = 105840$ tampered probes, from
 262 which the positive (spliced) set was generated by randomly sampling a number
 263 of examples equal to that of the negative class. (A similar experiment was car-
 264 ried out in [14], where the authors randomly sampled 10000 probes for both the
 265 negative and positive classes, thus introducing repetitions in the negative class.
 266 Hence the slight discrepancies between our results and theirs.)




Figure 6: ROC curves for the *virtual* splicing test on YaleB: (a) FISH[†], (b) FISH, where the three distance \mathcal{D} , \mathcal{D}' , \mathcal{D}'' and reported respectively in red, green and blue. In (c) ROCs for [14] and [27]. (best viewed in color)

267 Figure 6 reports the Receiver Operating Characteristic (ROC) plots for our
 268 FISH and FISH[†], using all the distances \mathcal{D} , \mathcal{D}' , and \mathcal{D}'' , together with results
 269 from [14] and [27], obtained by using the code available online. The Area Under
 270 the Curve (AUC) is reported in Table 2 for completeness, together with the
 271 True Positive Rate (TPR) at 0.01, 0.05, and 0.10 False Positive Rate (FPR). For
 272 this controlled acquisition setup on real face images, all the methods obtained
 273 comparable results. Notice that for this test, high-quality 3D face shapes were

274 computed using Face Gen¹, which requires several input images from different
 275 views for face, and manually annotated landmarks. It is worth remarking that
 276 this is still an unrealistic application scenario for us, as we target to work with
 277 real and noisy images on an automatic pipeline.

Table 2: Tests on YaleB (best results in bold)

Method	AUCs	TPR @ 0.01 FPR	TPR @ 0.05 FPR	TPR @ 0.10 FPR
FISH [†] with \mathcal{D}	0.9360	0.5066	0.7315	0.8295
FISH [†] with \mathcal{D}'	0.9439	0.6390	0.7864	0.8671
FISH [†] with \mathcal{D}''	0.9653	0.7950	0.8739	0.9161
FISH with \mathcal{D}	0.9049	0.1887	0.5633	0.7592
FISH with \mathcal{D}'	0.9719	0.8127	0.9034	0.9356
FISH with \mathcal{D}''	0.9611	0.7923	0.8739	0.9120
Peng et al. [14]	0.9754	0.8345	0.8961	0.9311
Kee & Farid [27]	0.9531	0.7120	0.8082	0.8680

278 5.3. Real faces in the wild

279 Tests with a fully unconstrained scenario were carried out by evaluating
 280 our automated pipeline on the DSO-1 dataset [11] containing real images. The
 281 DSO-1 dataset includes 100 pristine and 100 spliced images, with challenging
 282 manipulations. The dataset shows high variation of people poses and expres-
 283 sions, captured in indoor and outdoor scenarios under uncontrolled lighting
 284 conditions. Occlusions caused by other faces or objects (like glasses or hair) are
 285 also present. To the best of our knowledge, DSO-1 is the only freely available
 286 real-world face splicing database.

287 In order to compare our results with those reported in [14], we strictly fol-
 288 lowed their protocol², by excluding some DSO-1 images and by limiting the
 289 comparison to face pairs.

290 Table 3 reports the AUC of the ROC curve for different versions of our
 291 method and the current state-of-the-art methods. For our pipeline using FISH
 292 descriptors, the superscript ‘ \star ’ is applied when the recent CCN method described
 293 in [26] is employed to compute the 3D face model instead of the standard 3DMM.
 294 Figure 7 also reports ROC curves for our pipelines.

¹<https://facegen.com/modeller.htm>

²https://github.com/bomb2peng/CASIA_3Dlighting/tree/master/datasets/DSO-1

Table 3: Face splicing detection in terms of AUC on the DSO-1 dataset (best results in bold). Results for the state-of-the-art methods have been retrieved from [14]

Method	AUC
FISH [†] with \mathcal{D}	0.5454
FISH [†] with \mathcal{D}'	0.5462
FISH [†] with \mathcal{D}''	0.5962
FISH with \mathcal{D}	0.5374
FISH with \mathcal{D}'	0.5588
FISH with \mathcal{D}''	0.6135
FISH* with \mathcal{D}	0.5376
FISH* with \mathcal{D}'	0.5672
FISH* with \mathcal{D}''	0.6169
Peng et al. [14]	0.5795
Kee & Farid [27]	0.5715
Fan et al. [32]	0.5633

295 Results show that all the methods based on FISH obtain a better AUC with
 296 respect to the state-of-the-art in combination with the \mathcal{D}'' distance, demon-
 297 strating the effectiveness of the proposed solution. Exclusion of saturated pixels
 298 produce an additional improvement, while the albedo handling mechanism is
 299 very critical, as shown by the changes of performance when employing \mathcal{D} , \mathcal{D}'
 300 and \mathcal{D}'' . Moreover, while FISH* does not considerably improve the results with
 301 respect to the other FISH variants, as it lowers the False Positive Rate (FPR)
 302 but also slightly decreases the True Positive Rate (TPR), nevertheless it benefits
 303 from a minor computational complexity and code management over FISH. In
 304 addition, the FISH descriptor can better handle errors on the 3D shape cluster-
 305 ing and in weighting the contributions of similar normal vectors, thus reducing
 306 the impact of incorrectly estimated normals. For this reason, FISH can be more
 307 reliable in a fully automatic scenario, where the accuracy of the 3D face model
 308 is lower than in a synthetic scenario.

309 5.4. Distance normalization on FISH

310 As it can be noticed from experiments reported in Sec. 5.1 and Sec. 5.2, in
 311 all the tests on the Syn1 and Syn2 and using FISH[†] on the Yale database, the
 312 best results are achieved with the \mathcal{D}' distance, while using FISH on Yale and




Figure 7: ROC curves on DSO-1 with respectively (a) FISH[†], (b) FISH, and (c) FISH^{*}. For each version, the three distance \mathcal{D} , \mathcal{D}' , \mathcal{D}'' and reported respectively in red, green and blue (best viewed in color).

313 in all cases on the DSO-1 dataset, it is \mathcal{D}'' that obtains the best scores. This
 314 behavior is reasonably due to the different ranges of RGB values that can be
 315 found in the images. Table 4 reports for each dataset the standard deviation of
 316 the average RGB color of the related faces with and without saturated values.
 317 The standard deviation values are computed over the mean RGB value of each
 318 face, considering all the pixels exploited to compute the FISH descriptor (i.e. all
 319 pixels that are projection of a 3D vertex of the face model).

Table 4: Standard deviation of the average RGB color of the faces. Note that for YaleB only gray-scale images are provided

With saturated pixels				Without saturated pixels			
Dataset	STD			Dataset	STD		
	R	G	B		R	G	B
Syn1	12.99	9.91	8.92	Syn1	12.99	9.91	8.92
Syn2	12.99	9.91	8.92	Syn2	12.99	9.91	8.92
YaleB	23.93			YaleB	16.02		
DSO-1	29.99	29.54	28.58	DSO-1	25.21	23.79	23.13

320 According to the table, \mathcal{D}' gives better results in the case of low variance (e.g.
 321 inferior to 20), while \mathcal{D}'' obtains better results for data with higher variance.
 322 Notice also that no saturated pixels are found in the synthetic datasets, which
 323 confirms their limits in simulating a real scenario.

324 *5.5. Computational complexity*

325 Both FISH and the methods of [14] and [27] share the initial steps of the
326 pipeline (i.e., face detection and alignment, and 3D shape estimation). These
327 steps take most of the time spent in computation, that in our Matlab imple-
328 mentation correspond respectively on about 9 seconds for face detection on each
329 image, plus 0.15 and 0.08 seconds for face alignment and 3DMM fitting for each
330 single face detected.

331 Additionally, FISH and [27] just require to estimate the normal vectors of
332 the face shape, which takes about 10 seconds on average on our Matlab non
333 optimized implementation, while [14] exploits 3D information to synthesize 42
334 images of the face under different known illuminations in order to estimate the
335 *transfer coefficients* that are exploited to retrieve the lighting vector. Although
336 we cannot effectively verify the computational time spent by [14] as we lack data
337 to replicate this step, it would reasonably be equal or surpass the time spent
338 by FISH, since rendering software typically has to estimate the shape normal
339 vectors in addition to other steps. Moreover, [14] also requires to solve N 42x9
340 linear systems (i.e., 42 images per 9 lighting transfer functions, for each of the
341 sampling points).

342 For the final step, both methods in [27] and [14] solve a linear system with
343 N equations, that in our implementation takes about 5 milliseconds. On the
344 other hand, the FISH histogram has a computational complexity of $O(N\mathfrak{B})$,
345 that in our non-optimized implementation takes about 80 milliseconds.

346 Considering the whole pipeline, FISH running times are comparable to those
347 of [27], since most of the time is spent in the first step of the pipeline, while [14]
348 should spend more time for the computation of the *transfer coefficients*.

349 Notice that the distance computation is slightly slower for our solution, due
350 to the higher dimension of the histogram w.r.t. the lighting vector, but this has
351 a negligible impact over the computation time for the whole pipeline.

352 *5.6. Limitations*

353 FISH splicing detection, similarly to [14] and [27], relies on the comparison
354 of physical lighting environments from distinct faces, and requires at least two
355 faces in a probe image. Additionally, this kind of approach would not work if
356 the scene strongly violates the assumption of Lambertian surfaces illuminated
357 by distant lights, such in the case when objects in the scene cast strong shadows
358 over one of the faces under inspection. Finally, image resolution should be
359 sufficiently high to allow accurate face alignment and sampling of light color
360 intensity data.

361 **6. Conclusion**

362 This paper presented a novel approach to face splicing detection based on
363 light analysis. The proposed FISH descriptor is designed according to a sta-
364 tistical representation based on histograms, implicitly estimating the mapping
365 between image intensities and 3D normal vectors. FISH can alleviate the im-
366 pact of the low accuracy of the 3D face model, which typically strongly affects
367 the methods based on spherical harmonics. The effectiveness and robustness of
368 our solution has been demonstrated on three different datasets: While in the
369 controlled scenarios of Syn1/Syn2 and YaleB FISH obtains results comparable
370 to the state-of-the-art, on images acquired on real scenarios with unconstrained
371 lighting conditions, such those of the DSO-1 dataset, it outperform all the ex-
372 isting face splicing detectors based on lighting analysis.

373 **Acknowledgments**

374 This material is based on research sponsored by the Air Force Research Lab-
375 oratory and the Defense Advanced Research Projects Agency under agreement
376 number FA8750-16-2-0188. The U.S. Government is authorized to reproduce
377 and distribute reprints for Governmental purposes notwithstanding any copy-
378 right notation thereon.

379 **References**

- 380 [1] A. Piva, An overview on image forensics, ISRN Signal Processing 2013
381 (2013) Article ID 496701, 22 pages.
- 382 [2] P. Ferrara, T. Bianchi, A. De Rosa, A. Piva, Image forgery localization via
383 fine-grained analysis of cfa artifacts, Information Forensics and Security,
384 IEEE Transactions on 7 (5) (2012) 1566–1577. doi:10.1109/TIFS.2012.
385 2202227.
- 386 [3] B. Li, T. Ng, X. Li, S. Tan, J. Huang, Revealing the trace of high-
387 quality JPEG compression through quantization noise analysis, Informa-
388 tion Forensics and Security, IEEE Transactions on 10 (3) (2015) 558–573.
389 doi:10.1109/TIFS.2015.2389148.
- 390 [4] T. Bianchi, A. Piva, Image forgery localization via block-grained analysis
391 of jpeg artifacts, Information Forensics and Security, IEEE Transactions on
392 7 (3) (2012) 1003–1017. doi:10.1109/TIFS.2012.2187516.
- 393 [5] M. Chen, J. Fridrich, M. Goljan, J. Lukas, Determining image origin and in-
394 tegrity using sensor noise, Information Forensics and Security, IEEE Trans-
395 actions on 3 (1) (2008) 74–90. doi:10.1109/TIFS.2007.916285.
- 396 [6] B. Mahdian, S. Saic, Using noise inconsistencies for blind image forensics,
397 Image and Vision Computing 27 (10) (2009) 1497 – 1503.
- 398 [7] M. Zampoglou, S. Papadopoulos, Y. Kompatsiaris, Detecting image splic-
399 ing in the wild (web), in: Proc. IEEE Int Multimedia & Expo Workshops
400 (ICMEW) Conf, 2015, pp. 1–6.
- 401 [8] P. Zhou, X. Han, V. I. Morariu, L. S. Davis, Two-stream neural networks
402 for tampered face detection, in: 2017 IEEE Conference on Computer Vision
403 and Pattern Recognition Workshops (CVPRW), 2017, pp. 1831–1839.
- 404 [9] E. Kee, J. F. O’Brien, H. Farid, Exposing photo manipulation with in-
405 consistent shadows, ACM Trans. Graph. 32 (3) (2013) 28:1–28:12. doi:
406 10.1145/2487228.2487236.

- 407 [10] M. Iuliani, M. Fanfani, C. Colombo, A. Piva, Reliability assessment of
408 principal point estimates for forensic applications, *J. Visual Communica-*
409 *tion and Image Representation* 42 (2017) 65–77.
- 410 [11] T. Carvalho, C. Riess, E. Angelopoulou, H. Pedrini, A. de Rezende Rocha,
411 Exposing digital image forgeries by illumination color classification, *IEEE*
412 *Transactions on Information Forensics and Security* (2013) 1182–1194.
- 413 [12] M. Johnson, H. Farid, Exposing digital forgeries in complex lighting envi-
414 ronments, *Information Forensics and Security, IEEE Transactions on* 2 (3)
415 (2007) 450–461. doi:10.1109/TIFS.2007.903848.
- 416 [13] F. Bellavia, C. Colombo, Rethinking the sGLOH descriptor, *IEEE Transac-*
417 *tions on Pattern Analysis and Machine Intelligence* 40 (4) (2018) 931–944.
- 418 [14] B. Peng, W. Wang, J. Dong, T. Tan, Optimized 3d lighting environment
419 estimation for image forgery detection, *IEEE Transactions on Information*
420 *Forensics and Security* 12 (2) (2017) 479–494. doi:10.1109/TIFS.2016.
421 2623589.
- 422 [15] W. Zhou, C. Kambhamettu, A unified framework for scene illuminant es-
423 timation, *Image and Vision Computing* 26 (3) (2008) 415 – 429.
- 424 [16] A. Gijsenij, T. Gevers, J. van de Weijer, Computational color constancy:
425 Survey and experiments, *IEEE Transactions on Image Processing* 20 (9)
426 (2011) 2475–2489. doi:10.1109/TIP.2011.2118224.
- 427 [17] G. Buchsbaum, A spatial processor model for object colour per-
428 ception, *Journal of the Franklin Institute* 310 (1) (1980) 1 – 26.
429 doi:https://doi.org/10.1016/0016-0032(80)90058-7.
430 URL [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/0016003280900587)
431 [0016003280900587](http://www.sciencedirect.com/science/article/pii/0016003280900587)
- 432 [18] J. van de Weijer, T. Gevers, A. Gijsenij, Edge-based color constancy, *IEEE*
433 *Transactions on Image Processing* 16 (9) (2007) 2207–2214. doi:10.1109/
434 *TIP.2007.901808.*

- 435 [19] R. T. Tan, K. Nishino, K. Ikeuchi, Color constancy through inverse-
436 intensity chromaticity space, *J. Opt. Soc. Am. A* 21 (3) (2004) 321–334.
437 doi:10.1364/JOSAA.21.000321.
- 438 [20] R. Ramamoorthi, P. Hanrahan, On the relationship between radiance and
439 irradiance: determining the illumination from images of a convex lam-
440 bertian object, *J. Opt. Soc. Am. A* 18 (10) (2001) 2448–2459. doi:
441 10.1364/JOSAA.18.002448.
442 URL <http://josaa.osa.org/abstract.cfm?URI=josaa-18-10-2448>
- 443 [21] R. Basri, D. W. Jacobs, Lambertian reflectance and linear subspaces, *IEEE*
444 *Transactions on Pattern Analysis and Machine Intelligence* 25 (2) (2003)
445 218–233. doi:10.1109/TPAMI.2003.1177153.
- 446 [22] B. Peng, W. Wang, J. Dong, T. Tan, Improved 3d lighting environ-
447 ment estimation for image forgery detection, in: 2015 IEEE International
448 Workshop on Information Forensics and Security (WIFS), 2015, pp. 1–6.
449 doi:10.1109/WIFS.2015.7368587.
- 450 [23] B. Peng, W. Wang, J. Dong, T. Tan, Automatic detection of 3d lighting in-
451 consistencies via a facial landmark based morphable model, in: 2016 IEEE
452 International Conference on Image Processing (ICIP), 2016, pp. 3932–3936.
- 453 [24] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, T. Vetter, A 3d face model
454 for pose and illumination invariant face recognition, in: 2009 Sixth IEEE In-
455 ternational Conference on Advanced Video and Signal Based Surveillance,
456 2009, pp. 296–301.
- 457 [25] C. Cao, Y. Weng, S. Zhou, Y. Tong, K. Zhou, Facewarehouse: A 3d facial
458 expression database for visual computing, *IEEE Transactions on Visualiza-*
459 *tion and Computer Graphics* 20 (3) (2014) 413–425. doi:10.1109/TVCG.
460 2013.249.
461 URL <http://dx.doi.org/10.1109/TVCG.2013.249>

- 462 [26] G. Trigeorgis, P. Snape, I. Kokkinos, S. Zafeiriou, Face normals in-the-wild
463 using fully convolutional networks, in: 2017 IEEE Conference on Computer
464 Vision and Pattern Recognition (CVPR), 2017, pp. 340–349. doi:10.1109/
465 CVPR.2017.44.
- 466 [27] E. Kee, H. Farid, Exposing digital forgeries from 3-d lighting environments,
467 in: 2010 IEEE International Workshop on Information Forensics and Secu-
468 rity, 2010, pp. 1–6. doi:10.1109/WIFS.2010.5711437.
- 469 [28] M. Mathias, R. Benenson, M. Pedersoli, L. Van Gool, Face detection with-
470 out bells and whistles, in: Computer Vision – ECCV 2014: 13th Eu-
471 ropean Conference, Zurich, Switzerland, September 6-12, 2014, Proceed-
472 ings, Part IV, Springer International Publishing, Cham, 2014, pp. 720–735.
473 doi:10.1007/978-3-319-10593-2_47.
474 URL https://doi.org/10.1007/978-3-319-10593-2_47
- 475 [29] X. Xiong, F. D. la Torre, Supervised descent method and its applications to
476 face alignment, in: 2013 IEEE Conference on Computer Vision and Pattern
477 Recognition, 2013, pp. 532–539. doi:10.1109/CVPR.2013.75.
- 478 [30] X. Zhu, Z. Lei, J. Yan, D. Yi, S. Z. Li, High-fidelity pose and expression
479 normalization for face recognition in the wild, in: 2015 IEEE Conference
480 on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 787–796.
481 doi:10.1109/CVPR.2015.7298679.
- 482 [31] A. S. Georghiades, P. N. Belhumeur, D. J. Kriegman, From few to many:
483 illumination cone models for face recognition under variable lighting and
484 pose, IEEE Transactions on Pattern Analysis and Machine Intelligence
485 23 (6) (2001) 643–660.
- 486 [32] W. Fan, K. Wang, F. Cayre, Z. Xiong, 3d lighting-based image forgery
487 detection using shape-from-shading, in: 2012 Proceedings of the 20th Eu-
488 ropean Signal Processing Conference (EUSIPCO), 2012, pp. 1777–1781.