

Face Phylogeny Tree Using Basis Functions

Sudipta Banerjee, *Member, IEEE* and Arun Ross, *Senior Member, IEEE*

Abstract—Photometric transformations, such as brightness and contrast adjustment, can be applied to a face image repeatedly creating a set of near-duplicate images. Identifying the original image from a set of such near-duplicates and deducing the relationship between them are important in the context of digital image forensics. This is commonly done by generating an image phylogeny tree — a hierarchical structure depicting the relationship between a set of near-duplicate images. In this work, we utilize three different families of basis functions to model pairwise relationships between near-duplicate images. The basis functions used in this work are orthogonal polynomials, wavelet basis functions and radial basis functions. We perform extensive experiments to assess the performance of the proposed method across three different modalities, namely, face, fingerprint and iris images; across different image phylogeny tree configurations; and across different types of photometric transformations. We also utilize the same basis functions to model geometric transformations and deep-learning based transformations. We also perform extensive analysis of each basis function with respect to its ability to model arbitrary transformations and to distinguish between the original and the transformed images. Finally, we utilize the concept of approximate von Neumann graph entropy to explain the success and failure cases of the proposed IPT generation algorithm. Experiments indicate that the proposed algorithm generalizes well across different scenarios thereby suggesting the merits of using basis functions to model the relationship between photometrically and geometrically modified images.

1 INTRODUCTION

IN many applications, the face image of an individual may be subjected to photometric transformations such as brightness adjustment, histogram equalization, gamma correction, etc. These photometric transformations may be applied in a sequential fashion, resulting in an array of near-duplicate face images (see Figure 1). While some of these transformations can be used to improve face recognition [1], others may be maliciously used for image ‘tampering’ [2]. The availability of inexpensive photo editing tools and the ability to rapidly share images across social media platforms, has resulted in the posting of a large number of near-duplicates on the internet. Identification of the original image from a set of such near-duplicates is important in the context of digital image forensics [3]. Further, since these images are created by the successive application of photometric transformations, inferring the *order of evolution* is a challenging but interesting problem [4]. The order of evolution can be represented as an Image Phylogeny Tree (IPT), that indicates the relationship between the root node (original image) and the child nodes (transformed images) via directed links as illustrated in Figure 2. Deducing the IPT from the set of near-duplicates in an automated fashion has the following advantages:

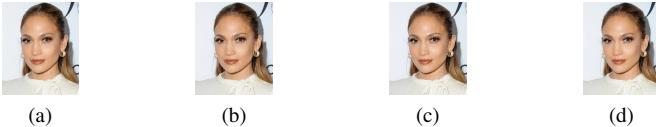


Fig. 1. Examples of variations of the same image uploaded on multiple websites with subtle modifications making them appear almost identical.

1. Indication of image tampering: An image can be tampered for a number of reasons. It can be done to airbrush celebrity faces on magazine covers,¹ or to depict fake situations to garner

• Both the authors are with the Department of Computer Science and Engineering, Michigan State University, East Lansing, MI, 48824.
E-mail: {banerj24, rossarun} @ cse.msu.edu

1. <https://www.cbsnews.com/news/uk-curb-airbrushed-images-keep-bodies-real/>

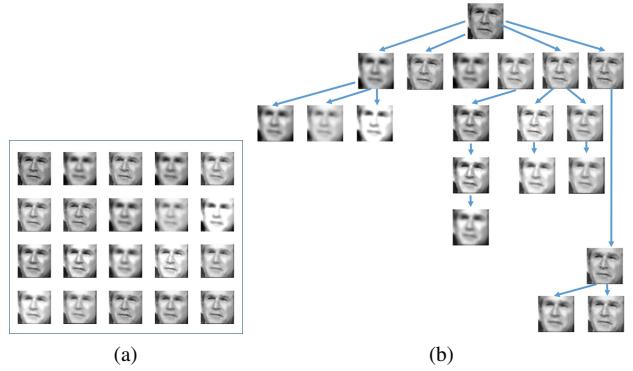


Fig. 2. Examples of photometrically related near-duplicate face images. (a) A set of 20 related images and (b) their corresponding Image Phylogeny Tree (IPT). In our work (a) is the input and (b) is the output.

political attention. In either case, tampered images convey false information. An IPT has directed links, and therefore, can be used to trace an image back to its origin, *i.e.*, the root node that denotes the original image.

2. Preserving chain of custody: Face images can be produced as culpable biometric evidence in legal proceedings [5]. The admissibility of such evidence is contingent on its integrity, *i.e.*, it should not have been tampered with. The chain of custody [6] of the digital evidence can be established via (i) hardware identification (the device used to acquire the biometric sample) [7], and (ii) analysis at the image level. IPT construction involves image-level analysis and can be leveraged to determine the authenticity between a pair of biometric samples *i.e.*, the original versus the tampered.

In this work, we tackle this challenging problem in the context of different biometric modalities (face, fingerprint and iris images) subjected to different types of photometric and geometric transformations resulting in IPTs of different configurations.

The remainder of the paper is organized as follows. Section 2 discusses the related work in the context of image phylogeny

literature. Section 3 describes the proposed method used to construct the IPT. Section 4 describes the dataset and the experiments conducted to evaluate the proposed method. Section 5 reports the results and the performance of the proposed algorithm. Section 6 describes the graph entropy based measure to assess the reconstructed IPTs. Finally, Section 7 concludes the paper.

2 RELATED WORK

The task of near-duplicate detection and retrieval (NDDR) is closely related to this work. NDDR is a well-researched topic which involves finding semantically related images for a given query image. In some cases, a query image may be a composite of multiple donor images, *i.e.*, portions of different images are *donated* to produce the composite. Provenance analysis methods [8] have been used to identify the donor images of a composite. In such cases, undirected phylogeny trees [9] are suitable, since deducing the links is more critical than determining the direction of the link. Image phylogeny trees (IPTs), on the other hand, deal with the task of determining ancestral relationships, *i.e.*, explicitly determining the parent and child node as well as the predecessors. In the case of IPT construction, the typical assumption is that there is a single root node and all the images are related such that there is no isolated node [10], [11], [12], [13], [14], [15]. We make the same assumption in this paper. Other works exist that consider the presence of multiple root nodes resulting in Image Phylogeny Forests [16], [17], [18].

The construction of an IPT commonly comprises of two major steps:

- Firstly, a pairwise asymmetric (dis)similarity measure is used to determine the degree of (dis)similarity between a pair of images and to discriminate successfully between the forward and reverse transformation directions. For an image pair $(\mathbf{I}_i, \mathbf{I}_j)$, where, \mathbf{I}_i is the original image and \mathbf{I}_j is the transformed image, the forward transformation direction refers to $\mathbf{I}_i \rightarrow \mathbf{I}_j$ and the reverse direction refers to $\mathbf{I}_j \rightarrow \mathbf{I}_i$. For a set of n near-duplicate images, this pairwise asymmetric measure is used to populate an $n \times n$ matrix.
- Secondly, a tree spanning algorithm utilizes the asymmetric matrix computed in the first step to determine the relationship between nodes and deduce a hierarchical tree structure.

In [10], the IPT is considered to be a minimal spanning tree (MST). In this work, we consider the ancestral links also as correct edges. Therefore, we interpret IPT as a directed acyclic graph (DAG) while relaxing the MST criterion. Majority of the literature focuses on different kinds of geometric transformations, *e.g.*, cropping, scaling, rotation; and pixel intensity-based transformations, *e.g.*, brightness and contrast adjustment, gamma transformation, and compression operations. Thus, traditionally, the first step of asymmetric measure computation involves geometric registration, followed by color channel normalization and compression matching [10]. Other works focus on using wavelet-based denoising technique [19], and a combination of gradient estimation and mutual information techniques [20] to derive an improved asymmetric measure.

The emphasis in the aforementioned works are on images depicting natural scenes. In this work, we focus on biometric images, *viz.*, face, fingerprints and iris. Recently, work has been done in modeling photometric transformations for iris images [21] using simple linear and quadratic functions. In [22], Legendre polynomials and Gaussian Radial Basis Functions were employed

to model photometric transformations for face images. The basis functions play an important role in not only reliably modeling the transformations, but also, in successfully discriminating between the forward and reverse directions. Therefore, the selection of a suitable set of basis functions is necessary to model the transformations and to distinguish between forward and reverse transformations. In the literature, there are a number of basis functions from the polynomial, wavelet and radial basis families. Even within a family, the basis functions can differ from each other, in terms of their ranges of orthogonality, associated weights, and other properties. In this paper, we expand on the work done in [22], by considering a larger number of potential basis functions. Our work *differs* from the previous work [22] as follows:

- 1) We consider three different families of basis functions for modeling photometric and geometric transformations: (i) Orthogonal polynomial family (Legendre and Chebyshev), (ii) Wavelet family (Gabor), and (iii) Radial Basis family (Gaussian and Bump).
- 2) We perform cross-modality testing, *i.e.*, learning the parameters of the basis functions using *face* images, and testing it on near-infrared *iris* images and optical sensor *fingerprint* images.
- 3) We test on multiple IPT configurations to evaluate the robustness of the proposed method. Also, we assess our method's robustness to unseen photometric and geometric transformations (*i.e.*, transformations not used during the training phase) accomplished using deep learning-based schemes, as well as open-source and commercial software. Furthermore, we have performed qualitative assessment of the IPTs reconstructed using the proposed method on near-duplicates downloaded from the internet.
- 4) We visualize the results using t -distributed stochastic neighbor embedding (t -SNE) to better understand the ability of the basis functions in modeling the transformations and discriminating between forward and reverse transformation directions.
- 5) We employ von Neumann directed graph entropy to better understand and evaluate the reconstructed IPTs.

3 PROPOSED METHOD

A photometrically related image pair $(\mathbf{I}_i, \mathbf{I}_j)$ can be generated by applying a single transformation or a sequence of transformations to one image resulting in the other image. However, to construct the IPT we require to differentiate between the original image and the transformed image. Say, if \mathbf{I}_i is the original image and \mathbf{I}_j is the transformed image, then the IPT should have a directed link as follows: $\mathbf{I}_i \rightarrow \mathbf{I}_j$. Applying this same principle to a set of near-duplicate photometrically related images, we need two sets: the first set denoting the *parent* nodes and the second set denoting the *child* nodes. These two sets are then used to construct the IPT (*parent* \rightarrow *child*). So the first step is to identify the sets of parent and child nodes from an array of near-duplicates.

We proceed to identify the parent and the child node for each pair of images by first modeling the transformation that relates the two images. We use parameterized basis functions to model the transformations in both directions ($\mathbf{I}_i \rightarrow \mathbf{I}_j$ and $\mathbf{I}_i \leftarrow \mathbf{I}_j$). But modeling the transformation does not indicate which is the parent node and which is the child node. To accomplish this, we require an asymmetric measure to distinguish between the forward and reverse directions. We pose the asymmetric measure computation

as the *likelihood* ratio problem [22]. To compute the likelihood ratio, we adopt a supervised framework with a training phase and a testing phase. In the *training* phase, we model numerous transformations for a large number of near-duplicate image pairs in both directions (in this phase we know the original and the transformed images *apriori*). This results in two sets of parameter distributions, one for the forward transformation and the other for the reverse transformation. In the *testing* phase, for a given near-duplicate pair, we first model the transformations in both directions. Next, we use the estimated parameters to determine how *likely* they are to originate from the forward parameter distribution as opposed to the reverse parameter distribution. This step leads to the computation of the asymmetric similarity measure. We repeat this step for all image pairs in the near-duplicate set. Upon pairwise modeling of all the near-duplicate images in the set, we perform thresholding to identify *related* image pairs. The similarity measure is then utilized to identify which image from the related pair is the parent, thus, making the other image its child. The sets of parent and child nodes are ultimately used to generate the IPT.

In this work, we seek to model an arbitrary transformation using a set of parametric functions, that we refer to as basis functions. Such an approach is needed since the space of photometric transformations is very vast; further, each of these transformations has a large number of parameter values. For example, a simple brightness adjustment can be accomplished using a large number of brightness values. The use of a fixed set of parametric functions to approximate a transformation reduces the otherwise complex task of modeling the photometric transformation. Thus, the task of approximating the transformations involves learning the *parameters* of the basis functions, subject to a criterion. In our case, the criterion or the objective function is the minimization of the photometric error between a near-duplicate image pair. This is formulated as below:

$$PE(\mathbf{I}_i, \mathbf{I}_j) = \min_{\alpha} \sum_p \|\mathbf{I}_i(p) - \mathcal{T}[\mathbf{I}_j(p)|\alpha]\|_2^2. \quad (1)$$

Here, $\mathcal{T}[\cdot|\alpha]$ denotes the photometric transformation. We model the transformation using the basis function as $\mathbf{I}_i(p) \approx \mathcal{T}[\mathbf{I}_j(p)|\alpha] \approx \sum_{h=1}^m \alpha_h \mathbb{B}_h[\mathbf{I}_j(p)]$, where the transformation is applied to each pixel p . $\alpha = [\alpha_1, \dots, \alpha_m]^T$ is the parameter vector to be estimated and m is the number of basis functions. In this work, we have five different types of basis functions, so the value of m depends upon the choice of the basis function. Next, we describe the process of modeling the transformations using the basis functions and the parameter estimation routines.

3.1 Parameter Estimation of Basis Functions

3.1.1 Orthogonal Polynomial Basis Functions

- 1) **Legendre polynomials** are a class of orthogonal polynomials defined in the interval $[-1, 1]$. The Legendre polynomial of degree n computed at x is denoted as $P_n(x)$ and is written as follows:

$$L_n(x) = 2^n \sum_{k=0}^n x^k \binom{n}{k} \binom{\frac{n+k-1}{2}}{n}. \quad (2)$$

Legendre polynomials have been successfully used for image template matching [23], and image reconstruction and compression [24]. Note that Eqn.(2) simplifies to a linear function for $n = 1$ and a quadratic polynomial for $n = 2$.

- 2) **Chebyshev polynomials** are a special case of Jacobi polynomials defined in the interval $[-1, 1]$. There are two kinds of Chebyshev polynomials, here we are interested in Chebyshev polynomials of first kind which have been extensively used for approximating complex functions such as graph convolution [25]. The explicit representation is presented below:

$$C_n(x) = x^n \sum_{k=0}^{\lfloor \frac{n}{2} \rfloor} \binom{n}{2k} (1 - x^{-2})^k. \quad (3)$$

In the notation, $\sum_{h=1}^m \alpha_h \mathbb{B}_h[\mathbf{I}_j(p)]$, $\mathbb{B}_h[\cdot]$ equals $L_h(\cdot)$ if Legendre polynomial is used and $C_h(\cdot)$ if Chebyshev polynomial is employed, and $m = n + 1$. Next, we solve the objective function in Eqn.(1) using the inverse compositional estimation (ICE) algorithm [26], [27]. The IC update rule expresses the updated transformation as a composition of the current transformation and the inverse of the incremental transformation. See [26] for a detailed derivation of the IC update rule. The parameter α is computed as $\alpha \leftarrow \frac{\alpha_{old}}{1 + \Delta\alpha}$. The IC update rule is an iterative optimization algorithm and updates the new α using the previous value, α_{old} , and the incremental $\Delta\alpha$. The incremental parameter vector is computed as,

$$\Delta\alpha = (\mathbf{J}_S \mathbf{J}_S^T + \lambda \mathbf{Id})^{-1} \mathbf{J}_S \mathbf{E}. \quad (4)$$

Here, the term \mathbf{J}_S is known as the Jacobian of the source image (\mathbf{I}_i), and the term $(\mathbf{J}_S \mathbf{J}_S^T)$ is known as the approximate Hessian matrix. We applied L_2 regularization to the Hessian matrix. Here, λ denotes the regularization parameter, \mathbf{Id} denotes the identity matrix, and \mathbf{E} denotes the error image computed between the source image (\mathbf{I}_i) and the modeled target image ($\mathcal{T}[\mathbf{I}_j|\alpha_{old}]$). In this work, α is a 6-dimensional vector for both Legendre and Chebyshev polynomials.

3.1.2 Wavelet Basis Functions

Gabor wavelets are used for extracting texture information from images [28] and has been selected as one of the basis functions for modeling the transformations. We employed a bank of Gabor filters parameterized with the wavelength and orientation. A set of four discrete wavelengths $\{2, 3, 4, 5\}$ and four orientations $\{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$ are selected. Each wavelength corresponds to a single filter scale that treats the image at a different resolution. Thus, we have a bank of sixteen Gabor filters. We filtered the image with the Gabor bank and we obtained 16 filtered responses. However, in our case we combined the orientation responses for each wavelength, thus reducing the total number of responses from 16 to 4. Finally, we use ICE to estimate the 4-dimensional parameter vector α ($m = 4$).

3.1.3 Radial Basis Functions

The polynomial and wavelet basis functions model the transformations at pixel level. In pixel-level modeling, the photometric error between each pixel of the original and transformed image pair is minimized using a weighted linear combination of basis functions. However, local filtering operations such as median filtering are applied in a patch-wise manner. Therefore, we used the family of radial basis functions that possesses nice smoothing properties to model transformations at the patch level. In patch-level modeling, the photometric error between two patches, one patch belonging to the original image and the second patch belonging to the transformed image, is minimized using a weighted linear combination

of basis functions. Therefore, patch-level modeling considers all the pixels within a patch for minimizing the photometric error. This resolves the spatial dependencies observed in patch-based photometric transformations.

- 1) **Gaussian radial kernel** is the first type of smoothing functions considered in the work. Gaussian RBF computed at x is denoted as $K(x)$ and is written as $K(x) = \exp \|x - \mu\|^2$. Consider an image pair $(\mathbf{I}_i, \mathbf{I}_j)$ each of which is tessellated into N_P non-overlapping blocks of size 16×16 . For the block q , where $q = [1, \dots, N_P]$, let $\mathbf{I}_j^q(p) = \mathcal{T}[\mathbf{I}_i^q(p) | \boldsymbol{\alpha}_q] \approx \sum_p \alpha_{p,q} K[\mathbf{I}_i^q(p)]$. Here, p denotes the pixel intensity value within the q^{th} block and μ indicates the average of the pixel intensity values within that block. For Gaussian RBF, $m = p$, *i.e.*, for a 16×16 block, the local least squares estimation yields $\boldsymbol{\alpha}_q \approx \boldsymbol{\alpha}_q^T \mathbb{B}[\mathbf{I}_i^q]$, where $\mathbb{B}[\mathbf{I}_i^q] = K[\mathbf{I}_i^q(p)]$, $\forall p$. The local least squares method is used to estimate the coefficient vector $\boldsymbol{\alpha}_q$ for each block. The final $\boldsymbol{\alpha}$ is a 256-dimensional vector obtained by computing the average of all $\boldsymbol{\alpha}_q$ s. See [22] for detailed derivation.
- 2) **Bump RBF** is a smooth compact function which can be interpreted as a Gaussian function scaled to lie on a disc of unity radius. It is not analytic unlike Gaussian RBFs, but can be used as generalized functions which are essential in converting discontinuous functions to smooth functions [29]. In this case, $K(x) = \exp \left(-\frac{1}{1 - x^2} \right)$ for $x \in [-1, 1]$. Here, x is mean-centered. Using least squares estimation, we obtain $\boldsymbol{\alpha}$ (256-dimensional).

3.2 Asymmetric Measure Computation and IPT Construction

The asymmetric measure can be in the form of pairwise similarity or dissimilarity, but in this work, we adopt a similarity-based asymmetric measure. The similarity measure computed between a pair of images determines whether an image pair is photometrically related or not, *i.e.*, whether a link should exist between a pair of nodes (images) in the IPT; it also helps determine the direction of the link by identifying the parent node and the child node. The parameters estimated for modeling the transformations are utilized to compute this similarity measure as described below.

3.2.1 Likelihood ratio for computing the asymmetric similarity measure

Given a pair of images, $(\mathbf{I}_i, \mathbf{I}_j)$, we first estimate the parameter vectors $\boldsymbol{\alpha}_{ij}$ and $\boldsymbol{\alpha}_{ji}$ in both directions ($\mathbf{I}_i \rightarrow \mathbf{I}_j$ and $\mathbf{I}_j \rightarrow \mathbf{I}_i$). The parameter vectors are necessary but not sufficient for constructing the IPT. We compute the *likelihood ratio* from the estimated parameters to yield a similarity score which can discriminate between the forward and reverse directions, and can thus be used to construct the IPT. To compute the likelihood ratio, we need the probability distribution of the parameter vectors — $p_f(\boldsymbol{\alpha})$ and $p_b(\boldsymbol{\alpha})$ corresponding to forward and reverse directions for a large number of training images. The probability distributions are generated in a supervised fashion, where we assume that we know for an image pair from the training set $(\mathbf{I}_r, \mathbf{I}_s)$, \mathbf{I}_r is the original image and \mathbf{I}_s is the transformed image. Then the forward transformation refers to $(\mathbf{I}_r \rightarrow \mathbf{I}_s)$, and the reverse transformation refers to $(\mathbf{I}_s \rightarrow \mathbf{I}_r)$. The set of $\boldsymbol{\alpha}_{rs}$ vectors computed for a large number of image pairs are used to estimate $p_f(\boldsymbol{\alpha})$. Similarly, the

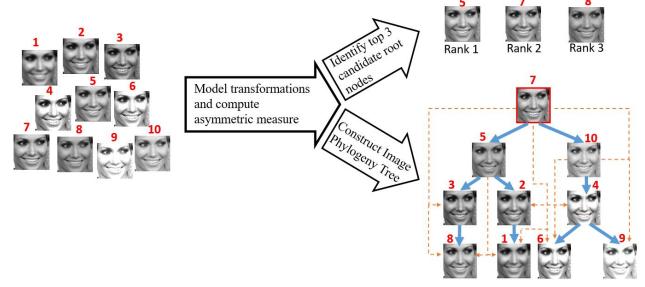


Fig. 3. The outline of the proposed method. The proposed method first models the photometric transformations between every image pair and then computes the asymmetric measure. Given a set of near-duplicate images as input (on the left) the two objectives are: (i) to determine the candidate set of root nodes, and (ii) to construct the IPT when the root image is known. The dashed arrows indicate ancestral links and the bold arrows indicate immediate links between parent and child nodes.

set of $\boldsymbol{\alpha}_{sr}$ parameter vectors are used to determine $p_b(\boldsymbol{\alpha})$. We utilized Parzen window based non-parametric density estimation scheme [30] to obtain $p_f(\boldsymbol{\alpha})$ and $p_b(\boldsymbol{\alpha})$.

Upon obtaining the forward and the reverse parameter distributions, we now compute the likelihood ratios as follows: $\Lambda_{ij} = \frac{p_f(\boldsymbol{\alpha}_{ij})}{p_b(\boldsymbol{\alpha}_{ij})}$. Similarly, $\Lambda_{ji} = \frac{p_f(\boldsymbol{\alpha}_{ji})}{p_b(\boldsymbol{\alpha}_{ji})}$. Our intuition is that we will observe a higher value of Λ_{ij} compared to Λ_{ji} , if \mathbf{I}_i is the original image and \mathbf{I}_j is the transformed image. In this case, $\boldsymbol{\alpha}_{ij}$ belongs to the forward distribution and should result in a higher value of Λ_{ij} . Conversely, $\boldsymbol{\alpha}_{ji}$ belongs to reverse distribution, resulting in a lower value of Λ_{ij} . The likelihood ratios are further used to populate the similarity matrix.

3.2.2 IPT Construction

The similarity matrix \mathbf{S} of size $n \times n$ is populated by the likelihood ratio values as follows: $\mathbf{S}_{ij} = \Lambda_{ij}; i, j = [1, \dots, n]$ and $i \neq j$. The diagonal elements of the similarity matrix are ignored as we do not consider self-loops in an IPT. The similarity matrix is then employed for (i) identifying the candidate root nodes and (ii) constructing the IPT. The steps are described below.

- 1) **Indicator matrix representation:** This step helps in pruning the outliers which can be falsely identified as root nodes. The indicator matrix is constructed by thresholding the similarity matrix against a suitable threshold, which results in a binary matrix. The details of the threshold selection are described in [22]. The indicator matrix serves as an adjacency matrix of a coarse directed acyclic graph that is further refined for constructing the IPT. The reason it is referred to as “coarse” is that it may contain some spurious edges.
- 2) **Candidate root nodes identification:** In this work, we consider each IPT to have a single root node. The authors in [12], [16] considered each node as a potential root, one at a time, and then computed a cost function for each IPT constructed using the potential root. The IPT resulting in the least cost function value was selected, and its root node was used for the final evaluation. This involves $O(n^3)$ (can be optimized to $O(n^2)$) complexity as reported in [12]. In contrast, the method proposed here computes a set of three candidate root nodes, which corresponds to the top 3 choices for the root node out of n nodes. This requires finding the nodes having the highest number of 1’s in the indicator matrix (we consider ancestral edges as correct edges). The

entire process requires summing each row of the indicator matrix followed by sorting and this results in $O(n \log n)$ computational complexity.

3) **IPT generation:** We construct the IPT as described in [22] using a depth-first search-based tree spanning technique. The choice of depth-first search (DFS) over breadth-first search (BFS) is motivated by the fact that DFS has a linear memory requirement with respect to the nodes and results in a faster search ($O(n)$) and is therefore used for topological sorting. The total computational complexity for the IPT construction using the proposed method is $O(n \log n) + O(n) \approx O(n \log n)$.

The outline of the proposed method for constructing the IPT is illustrated in Figure 3.

TABLE 1
Description of the datasets used in this work.

Modality	Name of the Dataset	Dataset Identifier	No. of subjects	No. of images
Face	LFW	Partial Set	391	12,290
		Full set	468	27,270
Iris	CASIA-IrisV2	—	37	7,260
	Device2	—	525	5,005
Fingerprint	CASIA-IrisV4	—	—	—
	Thousand	—	—	—
Fingerprint	FVC 2000	Config I	110	8,800
	DB3	Config II	90	7,200

4 EXPERIMENTS

In this section, we describe the datasets employed, the experiments conducted and finally report the results.

4.1 Datasets

We used four datasets belonging to three different modalities to conduct experiments. For the face modality, we used images from the Labeled Faces in the Wild (LFW) dataset [31]. For the iris modality, we used near-infrared iris images from the CASIA-IrisV2 Device2 subset [32] and CASIA-IrisV4 Thousand subset [33]. For the fingerprint modality, we used images from the FVC2000 DB3 dataset [34]. The description of the datasets is provided in Table 1. We selected four photometric transformations, *viz.*, Brightness adjustment, Median filtering filtering, Gaussian smoothing, and Gamma transformation as used in [22] to test the proposed IPT construction algorithm. The parameter range for each of the transformations is described in Table 2.

4.2 Experimental Methodology

We performed seven experiments which are described below.

4.2.1 Experiment 1: Efficacy of basis functions

In this experiment, we evaluate the ability of the basis functions to (i) model the photometric transformations and (ii) discriminate between the forward and reverse directions. To accomplish the first task, we perform two evaluation methods. The first evaluation involves *deterministically* selected parameters, whereas the second evaluation involves *randomly* selected transformation parameters. For the first evaluation, we select a face image, \mathbf{I} and subject it to a single transformation, *e.g.*, gamma adjustment, parameterized with a specific γ value, resulting in \mathbf{I}' . We repeat this process 200 times, each time we use an incrementally modified γ_{new}

($\gamma_{new} = \gamma_{old} + \Delta\gamma$), thus, resulting in 200 near-duplicate image pairs. Furthermore, we repeat this process for 5 images corresponding to 5 different subjects. Therefore, we have a total of 1,000 photometrically related image pairs for a single transformation. We conduct this process for each of the four transformations indicated in Table 2. Next, we use the basis functions to model the transformation *only* in the forward direction in this experiment. Then, we use *t*-SNE [35] to reduce the dimensionality of the estimated vectors and project them onto 3-dimensions. This experiment is conducted to assess the ability of the basis functions in modeling the transformations and we visually interpret the results. The second protocol involves modeling 2,000 image pairs in the forward direction using the five basis functions, where each of the 500 images were subjected to each of the four photometric transformations using randomly selected parameters. We further computed the residual photometric error (PE), for each image pair which is the difference in the pixel intensity between the actual target image and the output modeled using the basis functions. We then computed the mean over all 2,000 image pairs. The mean of the residual PE evaluates the ability of the basis functions to accurately model the transformations.

To evaluate the second task, *i.e.*, discriminating between the forward and reverse directions, we selected 400 image pairs (100 image pairs corresponding to each of the four transformations) from the 2,000 pairs, that were generated using randomly selected parameters. Then, we modeled the transformations in *both* forward and reverse directions, and estimated the parameters. We further used *t*-SNE to obtain a 2-dimensional embedding of the estimated parameters, and visualized the projections to analyze the performance of the basis functions.

4.2.2 Experiment 2: IPT Reconstruction

In this experiment, we evaluated the proposed approach in terms of (i) Root identification and (ii) IPT Reconstruction accuracy metrics as used in [22]. We followed the same experimental protocol in [22], and assessed the performance of the basis functions on both a partial set and a full-set of face images from the LFW dataset.² We used the IPT configuration presented in Figure 4(a) for this experiment.

4.2.3 Experiment 3: Cross-modality testing on multiple configurations

We tested the proposed approach on iris images and fingerprint images. This experiment is intended to demonstrate the generalizability of the proposed method across modalities.

1) **Iris Images** —We applied a random sequence of photometric transformations on near-infrared iris images. We evaluated the root identification and IPT reconstruction accuracies for 726 IPTs using the same procedure as described in Section 3.2. We used the same IPT configuration as the one used for face images (see Figure 4(b)). Note, the parameter probability distributions in the forward and reverse directions are computed using a training set comprising of *face* images; the test images are *iris* images. The test iris images are acquired in the near-infrared spectrum, in contrast to the training face images that are acquired in the visible spectrum. As a result, this experiment can also be treated as an assessment of the basis functions for cross-spectral modeling.

2. [http://iprobe.cse.msu.edu/dataset_detail.php?id=1&title=Near-Duplicate_Face_Images_\(NDFI\)](http://iprobe.cse.msu.edu/dataset_detail.php?id=1&title=Near-Duplicate_Face_Images_(NDFI))

TABLE 2

Photometric transformations and the range of the corresponding parameters used in the training and testing experiments. The transformed images are scaled to [0, 255]. Note that experiments were also conducted using other complex photometric transformations besides the ones listed here.

Photometric Transformations	Level of Operation	Parameters	Range
Brightness adjustment	Global	[a,b]	$a \in [0.9, 1.5]$, $b \in [-30, 30]$
Median filtering	Local	size of window [m,n]	$m \in [2, 6]$, $n \in [2, 6]$
Gaussian smoothing	Global	standard deviation	$stddev \in [1, 3]$
Gamma transformation	Global	gamma	$gamma \in [0.5, 1.5]$

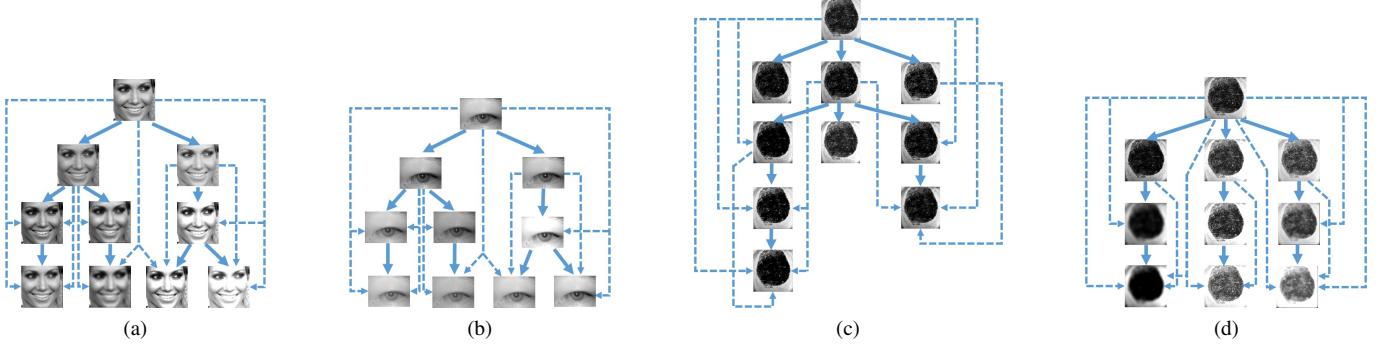


Fig. 4. IPT configurations used in Experiments 2 and 3 for the face, iris and fingerprint modalities. Note that the same configuration was tested across two modalities (Face and Iris) while, two different configurations were tested for the same modality (Finger). The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

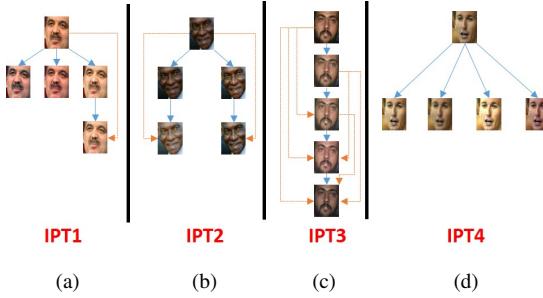


Fig. 5. IPT configurations used in Experiment 4. The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

2) **Fingerprint Images** — Two different IPT configurations are tested as depicted in Figures 4(c) and 4(d). This experiment tests the generalizability of the proposed approach as a function of the breadth and depth of the IPT. We refer to Figure 4(c) as Config I and Figure 4(d) as Config II. The IPT configuration used for testing the face and iris images is more balanced (similar distribution of nodes on the left and the right sides of the root) compared to the Config I structure which has more depth than breadth, whereas Config II has the same breadth at successive depths.

We also performed an intra-modality experiment which serves as the baseline experiment to compare against the performance of the cross-modality experiment. The train and test partitions for the intra-modality experiments are as follows:

- Iris images — We used 5,005 images from the CASIA-IrisV4 Thousand subset belonging to 525 subjects to learn the parameter distributions in the forward and the reverse directions. We then tested it on 726 IPTs (same configuration as in Figure 4(b)) constructed from the CASIA-IrisV2 Device2 subset. This experiment can also be considered as a cross-

dataset experiment, due to the use of two different datasets in the training and testing phases.

- Fingerprint images — We used the same dataset (*intra-dataset*) in training and testing but we strictly followed a subject disjoint protocol. This, however, resulted in a lesser number of training images. 560 images from 70 subjects were used for creating parameter distributions and then tested on 3,200 images from 40 subjects. We used the same configurations as depicted in Figures 4(c) and 4(d).

4.2.4 Experiment 4: Robustness to unseen photometric transformations

We considered a closed set of 4 transformations in the training stage. However, a gamut of image and video editing tools such as Photoshop, GIMP and Snapchat filters exist which can be used for image manipulation, particularly for face images. In this context, we constructed a small test set of images transformed using Photoshop operations (Hue and Saturation adjustment, Curve transformation, Color balance, and Blur filters) to create 35 IPTs corresponding to 5 subjects. The IPT configurations are selected such that they cover diverse breadth and depth values possible for an IPT with 5 nodes. See Figure 5. The trained parameter distributions did not encounter instances of Photoshopped images; hence, this experiment will demonstrate the robustness of the basis functions in handling unseen transformations.

4.2.5 Experiment 5: Ability to handle geometric transformations

We designed this experiment to assess the ability of basis functions in modeling geometric transformations. We selected some well-known geometric transformations such as sampling using linear interpolation and affine transformations that include translation, scaling and rotation. We have selected these particular transformations as they have also been utilized in [10] for creating near-duplicates. The details about the geometric transformations and

TABLE 3

Experiment 5: Geometric transformations and their parameter ranges used in this work.

Geometric transformations	Parameters
Re-sampling	[90%, 110%]
Generic Affine	Rotation Translation Scaling
	[-5°, 5°] [5, 20] [90%, 110%]

their respective parameter ranges are described in Table 3. We randomly selected 500 images belonging to 97 subjects from the Labeled Faces in the wild (LFW) dataset [31]. We then applied four geometric modifications (see Table 3) on each of these images in a random sequence with random parameter values to create 500 image phylogeny trees (IPTs). Each IPT contains 10 images so we have a total of 5,000 images. An example IPT consisting of geometrically modified images is presented in Figure 6. Note that the IPT configuration is the same as the one used to evaluate photometrically modified images. We have conducted the experiment using the following two protocols, and evaluated the performance using root identification accuracy at Ranks 1, 2 and 3 and IPT reconstruction accuracy.

1. The first protocol involves training on the photometrically modified images, while testing on geometrically modified images. In this protocol, the training set did not include any geometrically modified images, so it assesses the robustness of the basis functions on different classes of transformations (*i.e.*, photometric versus geometric). We have *not* modified the asymmetric measure computation method or the tree-spanning method used in constructing the IPT.

2. The second protocol involves training and testing on geometrically modified images. To accomplish this task, we created a new training set of 5,865 pairs of original and geometrically transformed images using the LFW dataset. The objective is to evaluate the performance of the basis functions when trained and tested on geometric modifications, unlike in the first protocol.

We compared the performance of the proposed method with a baseline algorithm described in [10]. The baseline algorithm uses Speeded-Up Robust Features (SURF) and RANSAC algorithm for the task of geometric registration, followed by color channel normalization, and uses the residual photometric error as the asymmetric measure. The Oriented Kruskal algorithm is used for spanning the IPT which is a minimal spanning tree in their case. We implemented the baseline algorithm in two ways.

(a) Firstly, we used SURF and M-SAC (M-estimator sample and consensus scheme which is an improved variant of RANSAC) for geometric registration. We did not perform color channel normalization since we used gray-scale images. We then rescaled pixel intensities in the original and modified images to [0, 255] prior to evaluation. We used the Oriented Kruskal algorithm for constructing the IPT.

(b) Secondly, we used the best performing basis function to compute the likelihood ratio to be used as the asymmetric measure and then employed the Oriented Kruskal for spanning the IPT.

It is important to note that, unlike the baseline method, the proposed method does not require any separate geometric registration for modeling the geometric transformations.

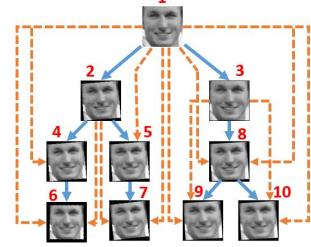


Fig. 6. Experiment 5: An example IPT generated using geometrically modified near-duplicate images. The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

4.2.6 Experiment 6: Ability to handle near-duplicates available online

Near-duplicate images of celebrities and political figures are widely circulated on the internet. The actual sequence of generation of such near-duplicates may be unknown, but these images represent pragmatic scenarios where the ground truth may not be always available. In this experiment, we analyze how the proposed asymmetric measure and IPT construction methods can handle such images. To this end, we followed the suggestion presented in [36] and used Google image search to download 40 near-duplicates retrieved from the following 5 queries: *Angelina Jolie*, *Kate Winslet*, *Superman*, *Britney Spears* and *Bob Marley*. We used training parameter distributions learnt from both geometrically and photometrically modified images for each of the five basis functions used in this work. We then used the proposed asymmetric measure computation method to identify top 3 candidate root nodes. For each of the candidate root node we then reconstructed an IPT. Due to unavailability of ground truth, we could not evaluate the accuracy of the reconstructed IPTs, but we present qualitative assessment of the reconstructed IPTs.

4.2.7 Experiment 7: Ability to handle deep learning-based transformations and image augmentation schemes

Several deep learning-based transformations and image augmentation packages are available that can be used for applying sophisticated transformations to images in an automated fashion generating a large number of near-duplicates. In this experiment, we used images generated using a deep learning-based autoencoder [37] and open source image augmentation packages [38]. We conducted the experiment using two protocols.

1. The first protocol involves a deep convolutional autoencoder [37]. The autoencoder was trained on $\sim 19,000$ images from the CelebA dataset [39] to generate 80 near-duplicate images belonging to 16 subjects. The resultant IPT configuration is depicted in Figure 7(a). The convolutional autoencoder comprises of an encoder block that consists of five convolutional layers, followed by ReLU after each convolutional layer, and the decoder block comprises of traditional convolutional layers and nearest-neighbor based upsampling.³ We did not use de-convolution or transposed convolution layers, as they can lead to checkerboard artifacts. The intuition behind using an autoencoder for generating near-duplicate images is to leverage its ability to perform high fidelity reconstruction of the original input images. This fits the definition of ‘near-duplicates’ in our image phylogeny task and has, therefore, been used in this experiment. We apply the original

3. <https://sebastianraschka.com/deep-learning-resources.html>

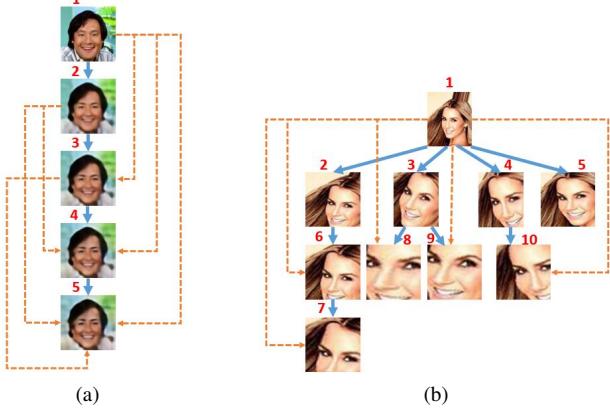


Fig. 7. Experiment 7: (Left) IPT test configuration used for evaluating proposed method on autoencoder generated near-duplicates. (Right) IPT test configuration used for evaluating proposed method on images generated by open source image augmentation package. The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

image as an input to the autoencoder to generate the first set of near-duplicates at depth=1. This first set of reconstructed images are again fed as input to the same autoencoder to generate near-duplicates at depth=2, and so on until we generate near-duplicates at depth=5.

2. The second protocol involves Augmentor [38], a data augmentation tool used when training deep neural networks. We used this tool, which is an open source package in Python, to apply random distortions such as zoom, cropping, rotation, resampling and elastic deformations on an image. See Figure 7(b). Some of these image transformations and their diverse parameter ranges (training involved rotation values in the interval $[-5^\circ, 5^\circ]$, whereas, testing using Augmentor involved rotation values in the interval $[-10^\circ, 10^\circ]$) are not encountered during the training stage. We randomly selected 100 images belonging to 100 subjects from the CelebA dataset. We applied the Augmentor on each of these 100 images to create 100 IPTs. Each IPT contains 10 images. So we tested on a total of 1,000 near-duplicate images.

In addition, we also used some images synthesized using a deep learning-based generative network known as Beauty-Glow [40]. The generative network performs a style transfer on the makeup of the individual in face images, resulting in near-duplicates as shown in Figure 8(a). Images are generated by sequentially increasing the magnification value of the makeup, highlighting the intensity of the makeup. We used 13 IPTs (each IPT contains 7 images), resulting in a total of 91 images.

5 RESULTS AND ANALYSIS

In this section, we first report the results observed for the experiments described in the previous section, and we further present our insights into the findings.

5.1 Results of Experiment 1

The 3D projected vectors obtained using *t*-SNE are illustrated in Figure 9 for each of the transformations modeled using the five basis functions. Each column denotes a photometric transformation, and each row denotes a basis function. As evident from the projections, the basis functions can model the majority of the

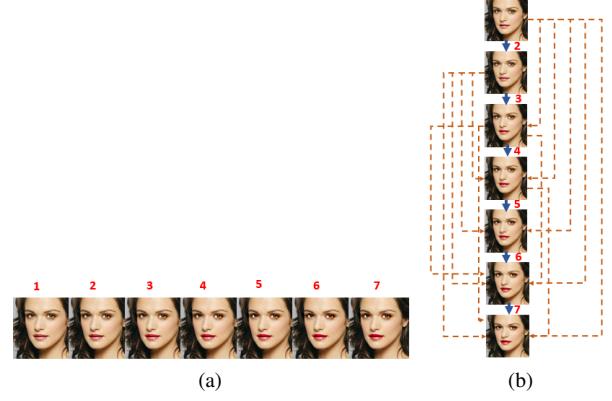


Fig. 8. Experiment 7: (Left) Near-duplicates generated using Beauty-Glow generative network. (Right) IPT constructed using Chebyshev polynomials for the near-duplicates on the left. The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

transformations fairly well. The parameters governing each transformation are incrementally modified and, hence, their projections should ideally span a continuous trajectory. Also, we expect to observe this behavior irrespective of the transformations used or the identity of the subject. We indeed observe such a behavior for most of the cases, except for Gamma adjustment, where the polynomials and wavelet functions flounder. Note that median filtering requires integer parameter values (height and width of window). Therefore, in the *t*-SNE results (last column in Figure 9) we observe small clusters, depicting accurate modeling of discrete parameterized transformations. Out of all the basis functions, the radial basis functions seem to model the transformations the best. Figure 10 further substantiates that the RBFs are best at modeling transformations while Gabor wavelets perform relatively poorly. The RBFs result in the lowest mean residual photometric error, suggesting a more accurate modeling of the photometric transformations.

In terms of discriminability between the forward and reverse directions, the projections are almost indistinguishable in the two directions for the wavelet functions, but they are relatively better for polynomial functions and the RBFs, as evidenced in Figure 11. The polynomials have fairly well-separated projections, indicating their ability to discriminate between the forward and reverse directions. We anticipate that this ability will be reflected in the IPT reconstruction experiments.

5.2 Results of Experiment 2

The results of root identification and IPT reconstruction are presented in Tables 4 and 5. Results indicate that polynomials (Legendre and Chebyshev) perform the best in a majority of cases among the set of five basis functions selected in this work. The results are consistent with the observations reported in Figure 11, which indicates sufficient discriminability offered by the polynomials. For the partial set, Legendre polynomials perform best both in terms of root identification (89.91%) and IPT reconstruction (70.61%) accuracies, closely followed by Chebyshev polynomials. For the full set, Gaussian RBF performs the best in terms of root identification accuracy (80.85%) while Chebyshev polynomials perform the best in terms of IPT reconstruction accuracy (66.54%, a small improvement of $\approx 1.5\%$ is observed compared to the results in [22]).

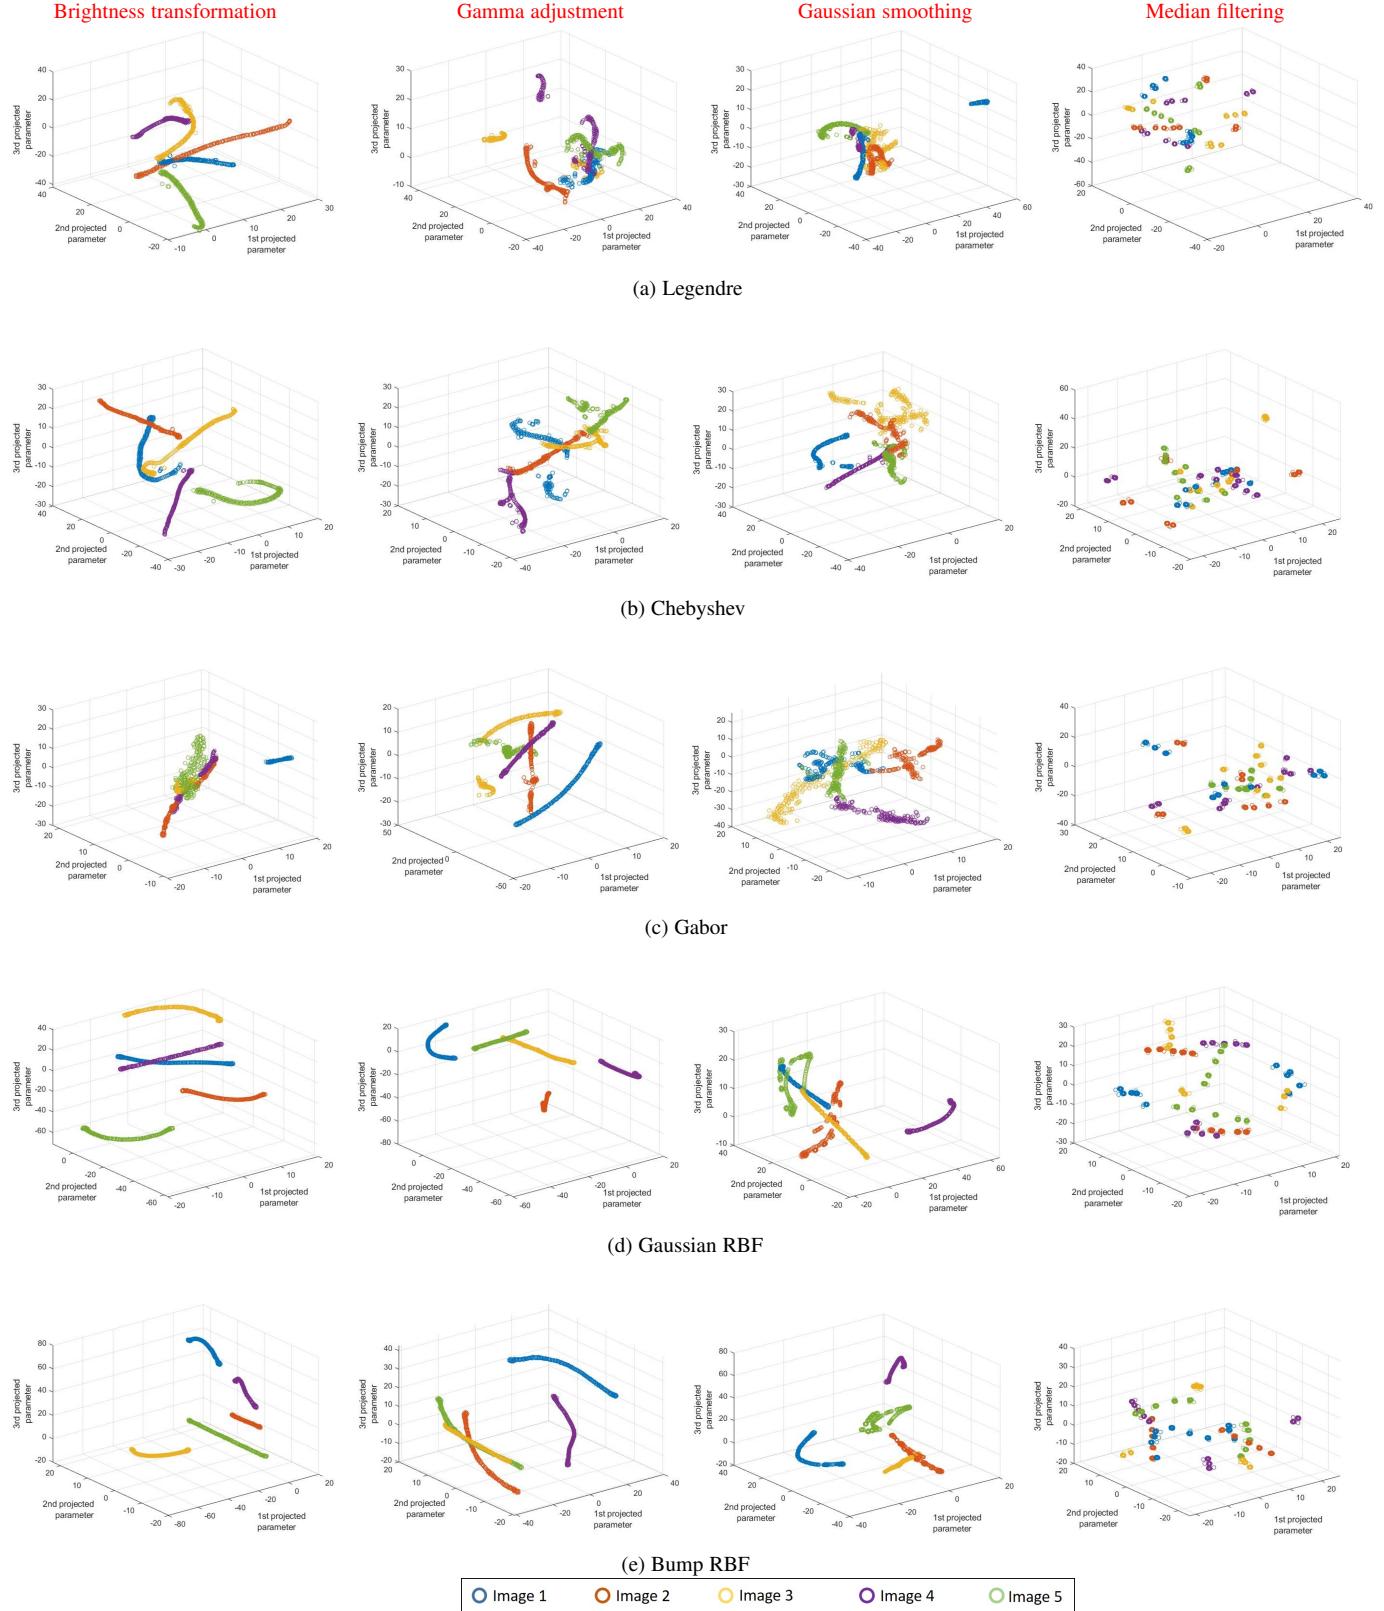


Fig. 9. Experiment 1: 3D projected parameters using t-SNE corresponding to each photometric transformation (column) modeled using each basis function (row). Each color represents a single image. A total of 5 images were modeled. Gaussian and Bump RBFs model majority of the transformations reasonably well as indicated by the last two rows. The Brightness transformation was easiest to model as the parameters of the basis functions follow a continuous path.

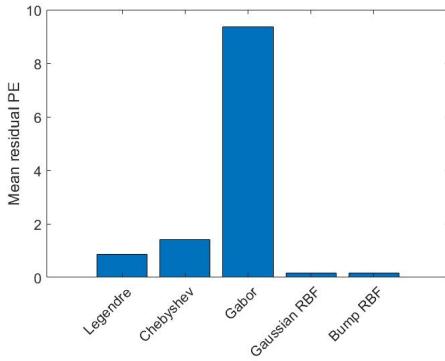


Fig. 10. Experiment 1: The photometric error between the actual output and the the output modeled using the basis functions is denoted as residual photometric error (PE). The mean of the residual PE is demonstrated for 2,000 image pairs modeled in both forward and reverse directions using the five basis functions. Gabor resulted in the highest residual PE, and the RBFs yield the lowest residual PE demonstrating their efficacy in reliably modeling the transformations.

5.3 Results of Experiment 3

The results for the cross-modality experiments are presented in Table 6 for iris images, and in Tables 7 and 8 for fingerprint images. The purpose of this set of experiments is to assess how the proposed method performs for (i) the same IPT configuration but used across two different modalities, and (ii) the same modality but tested on different IPT configurations. Note that the training modality is different than the test modality in both cases.

The results in Table 6 indicate that Chebyshev polynomials obtain 94.90% root identification accuracy at Rank 3 and an IPT reconstruction accuracy of 67.90% for iris images. It is closely followed by Legendre polynomials. However, other basis functions perform poorly, specifically, the Gabor wavelets. Gabor wavelets are good texture descriptors, *i.e.*, they can extract the high-frequency features reliably from an image. In the case of photometric transformations, the pixel intensity gradient, which contributes toward high-frequency features are not significantly affected and, thus, the wavelet fails to correctly model the transformation between an image pair.

Image phylogeny on near-duplicate fingerprint images is extremely hard, as evident from the results in Tables 7 and 8. Visual inspection reveals that the set of near-duplicate fingerprint images appear to be black blobs on a white background, with no discernible differences between the set. Therefore, the root identification and IPT reconstruction accuracies are worse compared to the face and iris modalities: the best root identification performance is 65.28% and IPT reconstruction accuracy is 70.59%. This experiment also shows that the performance varies across configurations. Specifically, symmetric configurations (Config-II) can be more difficult to reconstruct than asymmetric configurations (Config-I).

Next, we compare the results of the cross-modality experiments with the baseline, which is the intra-modality experiment described in Section 4.2.3. The results of the baseline experiments are reported in Table 9. The results indicate that cross-modality performance is commensurate with the intra-modality performance. For example, for iris images, the intra-modality experiment obtains the best root identification accuracy of 94.63% at Rank 3, while the cross-modality experiment obtains the best root identification accuracy of 94.90% at Rank 3. Furthermore, the intra-modality experiment obtains the highest IPT reconstruction

TABLE 4
Experiment 2: Root identification and IPT reconstruction accuracies for face images (Partial Set).

Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
Legendre	65.90/82.18/89.91	70.61
Chebyshev	53.62/74.69/85.52	70.53
Gabor	27.66/41.74/56.06	55.54
Gaussian RBF	65.25/81.04/87.79	66.15
Bump RBF	63.79/80.07/86.41	66.52

TABLE 5
Experiment 2: Root identification and IPT reconstruction accuracies for face images (Full set).

Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
Legendre	50.45/66.74/75.68	65.05
Chebyshev	45.18/65.13/76.86	66.54
Gabor	29.48/44.77/58.01	55.46
Gaussian RBF	56.44/71.87/80.85	63.84
Bump RBF	55.34/70.85/80.09	64.27

accuracy of 68.62%, while the cross-modality experiment obtains the highest IPT reconstruction accuracy of 67.90% in the case of iris images.

5.4 Results of Experiment 4

The training set comprised of images modified using 4 rudimentary photometric transformations. In the real world, a plethora of image editing applications exist, thereby making image phylogeny for face images a generally difficult problem. We hypothesize that by creating a training dataset through random parameters on simple transformations, the unseen transformations can be reliably modeled. Results reported in Table 10 indicate that unseen transformations were modeled fairly well. Legendre polynomials performed the best in terms of root identification accuracy with 76.47% averaged across the four IPT configurations (see Figure 5). Chebyshev polynomials performed the best in terms of IPT reconstruction accuracy with 76.25% averaged across the four IPT configurations (a small improvement of $\approx 1.67\%$ is observed compared to the results in [22]).

5.5 Results of Experiment 5

We present examples of geometric transformation modeling using the basis functions in Figure 12. We observe that Gaussian RBFs outperform Legendre polynomials at modeling the geometric transformations (see Figure 12) due to two reasons - (i) the radial basis functions can potentially span infinite range of values as opposed to the polynomials which can span values within a finite

TABLE 6
Experiment 3A: Root identification and IPT reconstruction accuracies for iris images in the cross-modality setting.

Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
Legendre	56.75/76.58/87.88	67.53
Chebyshev	72.59/88.29/94.90	67.90
Gabor	5.79/12.40/19.70	51.23
Gaussian RBF	40.08/63.64/76.45	66.74
Bump RBF	39.67/60.88/76.03	66

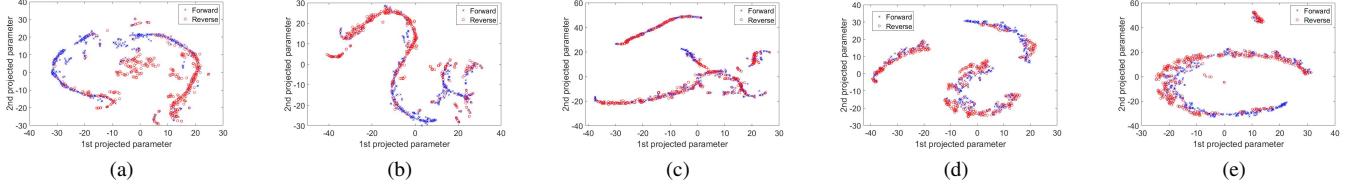


Fig. 11. Experiment 1: 2D projected parameters using t-SNE in forward and reverse directions, corresponding to all 4 transformations modeled using each basis function: (a) Legendre, (b) Chebyshev, (c) Gabor, (d) Gaussian RBF and (e) Bump RBF. Legendre and Chebyshev polynomials can better discriminate between forward and reverse directions as indicated by the relatively better separated parameter distributions compared to the remaining basis functions.

TABLE 7

Experiment 3B: Root identification and IPT reconstruction accuracies for fingerprint images (Config -I) in the cross-modality setting.

Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
Legendre	29.66/44.32/56.82	68.99
Chebyshev	31.93/46.14/57.39	70.59
Gabor	22.50/37.27/51.36	68.08
Gaussian RBF	30.80/50.34/61.14	68.98
Bump RBF	31.59/ 51.70/62.50	68.51

TABLE 8

Experiment 3B: Root identification and IPT reconstruction accuracies for fingerprint images (Config -II) in the cross-modality setting.

Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
Legendre	34.58/51.11/59.31	65.82
Chebyshev	35/55.28/65.28	65.93
Gabor	31.39/48.75/63.33	60.76
Gaussian RBF	14.58/23.75/36.81	59.29
Bump RBF	17.50/28.33/40.42	59.96

interval, and (ii) the RBFs did patch-level modeling compared to the pixel-level modeling done by the polynomials. The results in Table 11 indicate that the basis functions perform significantly better when trained on geometrically modified images (second protocol) compared to when trained on photometrically modified images (first protocol). As anticipated, if the class of transformations are the same in both training and testing set, the results are better, but surprisingly, even with photometrically modified

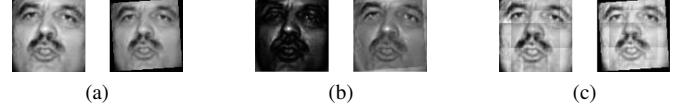


Fig. 12. Experiment 5: Example of geometric transformation (rotation) modeling using basis functions. (a) Original image (on the left) and the transformed image (on the right). (b) Modeled image pair using Legendre polynomials (modeled original image is on the left and modeled transformed image is on the right). (c) Modeled image pair using Gaussian RBF (modeled original image is on the left and modeled transformed image is on the right).

training images, the basis functions are able to reliably handle geometric transformations. The basis functions outperform the baseline (see first row in Table 11) by $\sim 50\%$ in terms of root identification accuracy and $\sim 56\%$ in terms of IPT reconstruction accuracy. In the case of substituting the asymmetric measure in the baseline with the proposed asymmetric measure (we used Gaussian RBF), while retaining the tree spanning algorithm (see second row in Table 11), an improvement of $\sim 32\%$ in terms of root identification accuracy and an improvement of $\sim 52\%$ in terms of IPT reconstruction accuracy is observed.

We also observe that the IPT reconstruction accuracy is lower for geometrically modified images compared to photometrically modified images. We tried to further analyze this difference in performance. Visual inspection revealed that the geometrically modified images appeared ‘more similar’ to the original image compared to the photometrically altered images. This can be attributed to the restrictive parameter range used in geometric transformations. A restrictive parameter range ensures that the images are indeed near-duplicates. A wide variation in the parameter values may result in highly dissimilar images, thereby, destroying the notion of near-duplicates. To quantify the degree of similarity between the original images and the modified images, we performed face recognition using a commercial face matcher. In the face recognition experiment, the original images served as the gallery and the geometrically modified images served as the probe samples. We repeat this same process for photometrically modified images belonging to Set I of Experiment 2. ROC curves presented in Figure 13 indicate a true match rate of 96.27% at a false match rate of 0.01% for the geometrically altered images, and a true match rate of 91.87% at a false match rate of 0.01% for the photometrically altered images. Note, the probe and gallery sizes for the photometrically modified images are four times more than that for the geometrically modified images. Nonetheless, the basis functions can better handle the identification of the original image (root node) and reconstruction of the IPT compared to the existing method.

TABLE 9

Experiment 3: Baseline performance of basis functions in terms of root identification and IPT reconstruction accuracies in the intra-modality setting.

Modality & Configuration	Basis Function	Root identification (%) Rank 1/2/3	IPT Reconstruction (%)
IRIS	Legendre	64.46/83.75/91.46	68.62
	Chebyshev	76.31/89.53/94.63	66.62
	Gabor	8.54/18.60/27.55	54.33
	Gaussian RBF	29.89/49.31/66.25	60.77
	Bump RBF	25.21/38.15/48.76	58.06
FINGERPRINT Config-I	Legendre	34.69/48.13/57.81	71.92
	Chebyshev	38.75/58.44/67.55	71.46
	Gabor	5.94/14.06/23.75	64.37
	Gaussian RBF	28.13/48.13/59.69	66.96
	Bump RBF	39.06/56.56/66.56	68.35
FINGERPRINT Config-II	Legendre	35.63/50.31/63.12	66.22
	Chebyshev	42.50/56.56/69.69	65.52
	Gabor	5.31/9.69/17.81	55.23
	Gaussian RBF	26.56/42.19/54.37	59.91
	Bump RBF	31.56/48.44/60.31	63.44

TABLE 10
Experiment 4: Root identification and IPT reconstruction accuracies for unseen photometric transformations.

Basis Functions	IPT Configurations							
	IPT 1		IPT 2		IPT 3		IPT 4	
	Root identification Rank 3 (%)	IPT Reconstruction (%)		Root identification Rank 3 (%)	IPT Reconstruction (%)		Root identification Rank 3 (%)	IPT Reconstruction (%)
Legendre	66.67	82.22	90	71.67	77.78	44.44	71.43	100
Chebyshev	44.44	84.44	60	75	44.44	45.56	71.43	100
Gabor	44.44	82.22	80	71.67	77.78	48.89	71.43	100
Gaussian RBF	66.67	84.44	40	68.33	44.44	43.33	71.43	100
Bump RBF	66.67	84.44	60	68.33	33.33	44.44	71.43	100

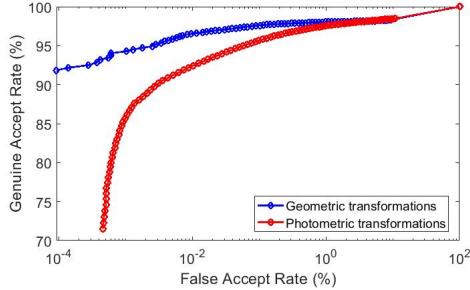


Fig. 13. Experiment 5: ROC curves for recognition of the original images with the photometrically and geometrically modified images using a COTS face matcher. The recognition performance is higher for geometrically altered images compared to photometrically modified images indicating high degree of similarity with the original images.

TABLE 11

Experiment 5: Root identification and IPT reconstruction accuracies for geometric transformations. The top two rows indicate the baseline algorithms. The baselines yield only one root node as output so results are reported only at Rank 1 and the remaining ranks are indicated Not Applicable (NA). In this experiment, the testing (TE) is always done on geometrically modified images (indicated by TE-GM) but the training (TR) can be done using either photometrically modified images (indicated by TR-PM) or geometrically modified images (TR-GM). Results indicate training on geometrically modified images yield best performance when tested on geometric transformations.

Method	Protocol	Root identification accuracy at Ranks 1/2/3 (%)	IPT reconstruction accuracy (%)
Baseline	(SURF + MSAC) + Oriented Kruskal Gaussian RBF + Oriented Kruskal	8.00 / NA / NA	3.62
		27.20 / NA / NA	7.24
Legendre	TR-PM, TE-GM TR-GM, TE-GM	14.20 / 24.40 / 37.00 23.20 / 39.60 / 52.80	52.78 54.30
Chebyshev	TR-PM, TE-GM TR-GM, TE-GM	7.00 / 13.40 / 19.60 23.20 / 35.60 / 49.20	51.40 59.81
Gabor	TR-PM, TE-GM TR-GM, TE-GM	25.80 / 41.20 / 52.60 25.60 / 40.00 / 54.60	57.75 58.49
Gaussian RBF	TR-PM, TE-GM TR-GM, TE-GM	25.40 / 42.40 / 57.00 58.60 / 75.80 / 86.00	55.67 51.40
Bump RBF	TR-PM, TE-GM TR-GM, TE-GM	7.80 / 17.60 / 31.00 46.60 / 65.80 / 77.00	49.82 56.32

5.6 Results of Experiment 6

Examples of the near-duplicates retrieved from the internet and their corresponding IPT reconstructions are presented in Figure 14. Qualitative analysis indicates that the reconstructed IPTs can depict the relationship between the near-duplicates reasonably well. For example, the images 3, 4 and 1 for the *Bob Marley* images (see Figure 14(c)) should ideally follow the sequence as indicated by the IPT. Similarly, image 2 appears to be a cropped version of image 3 which is correctly constructed by the proposed method.

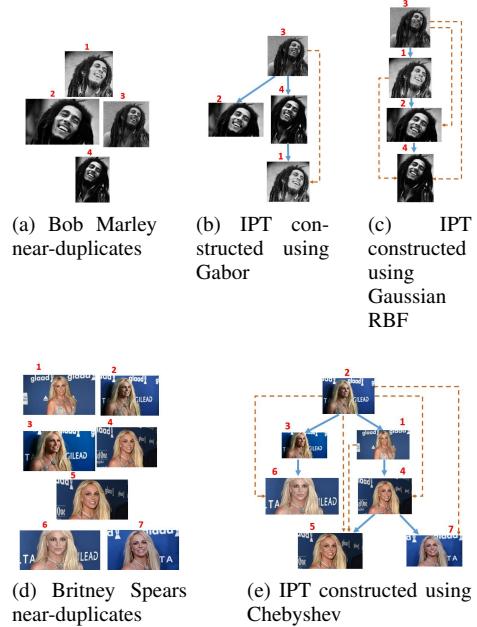


Fig. 14. Experiment 6: Examples of near-duplicates available online and their corresponding IPTs constructed using the proposed method. The first row corresponds to (a) 4 near-duplicates retrieved using the query *Bob Marley*, (b) IPT constructed using Gabor trained on photometric distribution (the top 3 candidate root nodes are 2,3,1) and (c) IPT constructed using Gaussian RBF trained on geometric distribution (the top 3 candidate root nodes are 3,2,1). The second row corresponds to (d) 7 near-duplicates retrieved using the query *Britney Spears* and (e) IPT constructed using Chebyshev trained on photometric distribution (the top 3 candidate root nodes are 2,4,5). The bold arrows indicate immediate links and the dashed arrows indicate ancestral links.

5.7 Results of Experiment 7

We used parameter distributions learnt from both photometrically modified images and geometrically modified images for IPT reconstruction. The results are reported in Table 12 in terms of root identification and IPT reconstruction accuracies.

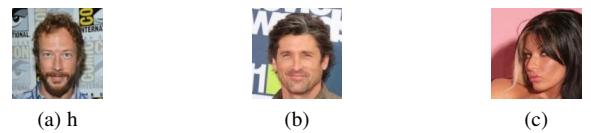


Fig. 15. Example images from the CelebA dataset containing prominent background details in the face images.

1. For the near-duplicates generated using autoencoder, the Gabor model performs the best in terms of root identification accuracy (Rank 1 accuracy of 81.25%) and the Chebyshev polynomials perform the best in terms of IPT reconstruction accuracy

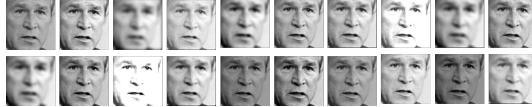


Fig. 16. An example of a set of photometrically modified images used for comparison between the proposed algorithm and human assessment. 20 such sets were simulated and each user was asked to identify their top three choices for the original image (root node).

(64.37%). Gabor outperforms the remaining basis functions due to the characteristics of the CelebA dataset that includes face images along with background details (see Figure 15). Gabor as a texture descriptor is able to accurately model the change in the texture of the background regions in the outputs reconstructed using the autoencoder.

2. For the near-duplicates generated using the image augmentation package, maximum root identification accuracy of 50% at Rank 3 is obtained by Bump RBF while Chebyshev obtains an IPT reconstruction accuracy of 75.20%. The image augmentation package includes elastic distortions that are complex non-linear deformations, that the basis functions did not model accurately.

3. For the near-duplicates generated using the BeautyGlow network, we only performed qualitative evaluation due to the small size of the test set. We observe that the IPT constructed in Figure 8(b) indicates a deep tree with gradual increase in the intensity of the make-up which is anticipated.

We observed that the basis functions perform better when trained on photometrically modified images compared to when they are trained on geometrically modified images, in the case of both autoencoder and Augmentor. This is perhaps due to the fact that the autoencoder does not introduce geometric modifications; the modifications are restricted to structural and textural details. On the other hand, the image augmentation library produces random geometric modifications such as rotation and re-sampling, but uses more sophisticated techniques to remove some of the artifacts associated with such operations (*e.g.*, removes the black padding near the borders of a rotated image). In such cases, we speculate that the basis functions view the geometrically altered image as a photometrically modified image. This explains the better performance of the basis functions when trained on photometrically modified images compared to geometrically modified images.

In addition to the seven experiments discussed above, we performed another small experiment to compare human ability through visual inspection with the performance of the proposed algorithm. To accomplish this task, we generated 20 sets of photometrically modified images, where each set contained 20 images subjected to a random sequence of the 4 transformations described in Table 2. An example of such a set is depicted in Figure 16. Next, we asked each of 12 evaluators to indicate their top 3 preferences for the original image (root node) of each set.

We compared the root identification accuracy as assessed by human evaluators with the output generated by the proposed algorithm for all the five basis functions employed in this work. The bar plots presented in Figure 17 indicate that the proposed algorithm performs comparably with manual evaluation. At Rank 1 the best performing algorithm (Chebyshev) is outperformed by humans by $\approx 5\%$, however, at Ranks 2 and 3, the Chebyshev polynomials outperform human evaluation by $\approx 10\%$.

Next, we highlight the main take-away points from all the

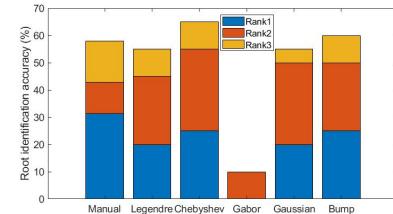


Fig. 17. Bar plots comparing the performance of human evaluators with the proposed algorithm in terms of root identification for 20 sets of images, each set containing 20 near-duplicate face images (see Figure 16).

experiments.

- 1) Radial basis functions performed best at modeling the baseline photometric transformations (Brightness adjustment, Gamma transformation, Gaussian smoothing and Median filtering) and also geometric transformations. They resulted in the lowest residual error when used for modeling the transformations.
- 2) Orthogonal polynomials performed best at reliably discriminating between the forward and reverse directions. They resulted in the highest root identification and IPT reconstruction accuracies in a majority of the cases involving photometric and geometric modifications.
- 3) The proposed approach of utilizing “likelihood ratio” generalizes well across multiple IPT configurations and different biometric modalities.
- 4) The proposed approach is capable of handling different classes of transformations both photometric and geometric. In addition, they generalize well over previously unseen transformations based on deep learning tools and image editing software.

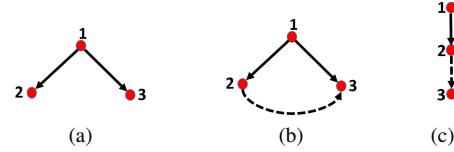


Fig. 18. Toy example demonstrating the effect of inserting a spurious edge on the von Neumann entropy. (a) Groundtruth IPT (b) Correctly reconstructed IPT with spurious edge (c) Incorrectly reconstructed IPT with spurious edge. Note, the spurious edge is indicated by a dashed line.

6 EXPLANATORY MODEL

Finally, in this section, we present a graph entropy-based explanatory model which analyzes the failure cases of the IPTs reconstructed using the proposed method. A failed IPT reconstruction can involve (i) missing edges, and (ii) spurious edges. Such failure cases can be quantified using the approximate von Neumann entropy measure for directed graphs. Graph entropy is computed in terms of its in-degree and out-degree [41]. Consider a directed graph $G(V, E)$ with a set of nodes denoted by V and the set of edges denoted by E . The adjacency matrix of such a graph is defined as,

$$A_{uv} = \begin{cases} 1 & \text{if } (u, v) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

TABLE 12

Experiment 7: Root identification and IPT reconstruction accuracies for deep learning-based transformations. The near duplicates are generated either using the autoencoder (left block) or image augmentation schemes for training deep neural networks (right block).

Basis functions	For near-duplicates generated using autoencoder				For near-duplicates generated using image augmentation schemes			
	Performance when trained on photometric transformations		Performance when trained on geometric transformations		Performance when trained on photometric transformations		Performance when trained on geometric transformations	
	Root identification accuracy (%) @ Ranks 1/2/3	IPT reconstruction accuracy (%)	Root identification accuracy (%) @ Ranks 1/2/3	IPT reconstruction accuracy (%)	Root identification accuracy (%) @ Ranks 1/2/3	IPT reconstruction accuracy (%)	Root identification accuracy (%) @ Ranks 1/2/3	IPT reconstruction accuracy (%)
Legendre	25.00 / 50.00 / 68.75	50.63	25.00 / 50.00 / 62.50	62.50	16.00 / 27.00 / 40.00	74.00	16.00 / 30.00 / 42.00	71.27
Chebyshev	25.00 / 56.25 / 62.50	64.37	37.50 / 43.75 / 68.75	46.88	13.00 / 24.00 / 39.00	75.20	18.00 / 29.00 / 43.00	72.00
Gabor	81.25 / 100 / 100	55.00	31.25 / 43.75 / 56.25	49.38	12.00 / 25.00 / 36.00	71.93	18.00 / 30.00 / 40.00	70.27
Gaussian RBF	87.50 / 93.75 / 93.75	55.00	43.75 / 56.25 / 68.75	50.62	17.00 / 30.00 / 43.00	65.20	19.00 / 30.11 / 41.00	67.67
Bump RBF	56.25 / 68.75 / 87.50	57.50	25.00 / 37.50 / 43.75	55.00	17.00 / 32.00 / 50.00	66.20	12.00 / 26.00 / 44.00	67.40

TABLE 13

Approximate von Neumann entropy for analysis of spurious edges. The mean and the standard deviation of the differences between the entropy of the ground truth and the reconstructions are reported. Low values indicate accurate reconstructions and smaller number of spurious as well as missing edges.

Basis Function	Entropy (Mean and standard deviation)
Legendre6	-0.0009 \pm 0.0045
Chebyshev	-0.0000 \pm 0.0044
Gabor	-0.0036 \pm 0.0039
Gaussian RBF	-0.0035 \pm 0.0028
Bump	-0.0034 \pm 0.0030

The in-degree and out-degree of node u are presented as $d_u^{in} = \sum_{v \in V} A_{vu}$; $d_u^{out} = \sum_{v \in V} A_{uv}$. The von Neumann entropy for the directed graph G can be written as:

$$H = 1 - \frac{1}{|V|} - \frac{1}{2|V|^2} \left[\sum_{(u,v) \in E} \frac{d_u^{in}}{d_v^{in} d_u^{out}} + \sum_{(u,v) \in E_2} \frac{1}{d_u^{out} d_v^{out}} \right]. \quad (5)$$

Here, $E_1 = \{(u,v) | (u,v) \in E \text{ and } (v,u) \notin E\}$; $E_2 = \{(u,v) | (u,v) \in E \text{ and } (v,u) \in E\}$, such that, $E = E_1 \cup E_2$ and $E_1 \cap E_2 = \emptyset$. The maximum value of entropy for a directed graph is equal to $1 - \frac{1}{|V|}$ for a star graph where all the nodes in the graph have either incoming links or outgoing links but not both. The minimum entropy is obtained for a cyclic graph, implying all nodes are fully connected, and the value of minimum entropy is $1 - \frac{1}{|V|} - \frac{1}{2|V|^2}|V| = 1 - \frac{1}{|V|} - \frac{1}{2|V|} = 1 - \frac{1.5}{|V|}$. In our work, we want to move toward maximum entropy since minimum entropy results in the worst case scenario of reconstruction where the result is a cyclic graph. Thus, Eqn. (5) can be simplified as follows.

$$H = 1 - \frac{1}{|V|} - \frac{1}{2|V|^2} \left[\sum_{(u,v) \in E_1} \frac{d_u^{in}}{d_v^{in} d_u^{out}} \right]. \quad (6)$$

The above equation is applicable in our work since $E_2 = \emptyset$, otherwise, we will end up having a cycle in the reconstructed IPT. We hypothesize that spurious edges decrease the entropy of the IPT, which is further reduced if the IPT misses correct edges. This can be illustrated using a toy example presented in Figure 18. The reconstruction yields a 100% IPT reconstruction accuracy for the first reconstructed IPT (Figure 18(b)), which has no missing edge but one spurious edge, and 50% for the second reconstructed IPT (Figure 18(c)), which has one missing edge and one spurious edge. The von Neumann entropy as computed from Eqn. (6) for the ground truth configuration (Figure 18(a)) is, $H(GT) = 1 - \frac{1}{3} - \frac{1}{18} \left[\frac{d_1^{in}}{d_2^{in} d_1^{out}} + \frac{d_1^{in}}{d_3^{in} d_1^{out}} \right] = 1 - \frac{1}{3} - \frac{1}{18} \left[0 \right] = 0.67$.

Note that the in-degree of node 1 is 0. This value also corresponds to the maximum entropy of a directed graph with 3 nodes. The entropy for the first reconstructed IPT (Figure 18(b)) is, $H(IPT1) = 1 - \frac{1}{3} - \frac{1}{18} \left[\frac{d_1^{in}}{d_2^{in} d_1^{out}} + \frac{d_1^{in}}{d_3^{in} d_1^{out}} + \frac{d_2^{in}}{d_3^{in} d_2^{out}} \right] = 1 - \frac{1}{3} - \frac{1}{18} \left[0 + 0 + \frac{1}{2 \times (1)^2} \right] = 0.64$. The entropy for the second reconstructed IPT (Figure 18(c)) is, $H(IPT2) = 1 - \frac{1}{3} - \frac{1}{18} \left[\frac{d_1^{in}}{d_2^{in} d_1^{out}} + \frac{d_2^{in}}{d_3^{in} d_2^{out}} \right] = 1 - \frac{1}{3} - \frac{1}{18} \left[0 + \frac{1}{1 \times (1)^2} \right] = 0.61$. Thus, $H(GT) > H(IPT1) > H(IPT2)$. This demonstrates that inaccurate IPTs (missing edges and spurious edges) can reduce the entropy from the ground truth entropy, and this property can be leveraged for evaluating the goodness of the reconstructed IPTs. Reconstructed IPTs with missing edges and spurious edges will tend to have lower entropy.

We compute this entropy-based measure to analyze the accuracy of the IPTs reconstructed using the proposed method for the face images (full set with 2,727 IPTs). The maximum and minimum entropy for a graph containing 10 nodes is 0.85 and 0.90, respectively. The von Neumann entropy as computed from Eqn. (6) for the ground truth IPT (see Figure 4(a)) is 0.89. We then compute the entropy corresponding to the reconstructed IPTs. Finally, we compute the difference between the entropy of the ground truth configurations and the reconstructed configurations. We report the mean and the standard deviation of the differences in the entropy in Table 13. The results indicate that the differences between the ground truth entropy and the entropy of the reconstructed IPTs are the smallest when the polynomials are utilized as basis functions. This further corroborates our findings that the polynomials result in the best IPT reconstruction accuracies (see Tables 4 - 9).

7 CONCLUSION

In this work, we presented a method to model pairwise photometric transformations between a set of near-duplicate biometric images using basis functions. The modeling of the photometric transformations is used in conjunction with the likelihood ratio based asymmetric measure to identify the original image and deduce how the images are related to each other. We performed comprehensive experiments using face, iris and fingerprint images under different experimental settings including intra-modality, cross-modality, and cross-dataset scenarios. Results indicate promising performance with a root identification accuracy of $\approx 94\%$ and an IPT reconstruction accuracy of $\approx 72\%$. We also analyzed the performance of the basis functions using t-SNE and observed that some of the photometric transformations are easier to model than

others. We also successfully modeled geometric transformations using the basis functions. We observed that the proposed scheme can be generalized to previously unseen transformations accomplished using image editing software and deep learning-based methods. Finally, we utilized graph entropy measure to evaluate the reconstructed IPTs. The deviations between the ground truth and reconstructed IPTs are minimal indicating that the proposed reconstruction algorithm performs reasonably well. Future work will focus on performing a global analysis that will simultaneously consider all the nodes to derive accurate IPT reconstructions.

8 ACKNOWLEDGMENTS

This material is based upon the work supported by the National Science Foundation under Grant Number 1618518.

REFERENCES

- [1] V. Štruc and N. Pavesic, “Photometric normalization techniques for illumination invariance,” in *Advances in Face Image Analysis: Techniques and Technologies*. IGI-Global, 01 2011, pp. 279–300.
- [2] A. Agarwal, A. Sehwag, R. Singh, and M. Vatsa, “Deceiving face presentation attack detection via image transformations,” *IEEE International Conference on Multimedia Big Data*, 08 2019.
- [3] H. Farid, “Digital image forensics,” *Scientific American*, vol. 298, no. 6, pp. 66–71, 2008.
- [4] A. Ross, S. Banerjee, C. Chen, A. Chowdhury, V. Mirjalili, R. Sharma, T. Swearingen, and S. Yadav, “Some research problems in biometrics: The future beckons,” in *12th IAPR International Conference on Biometrics (ICB)*, Crete, Greece, June 2019.
- [5] N. A. Spaun, “Forensic biometrics from images and video at the federal bureau of investigation,” in *IEEE International Conference on Biometrics: Theory, Applications, and Systems (BTAS)*, Sep. 2007, pp. 1–3.
- [6] J. Čosić and M. Baća, “(Im)proving chain of custody and digital evidence integrity with time stamp,” in *The 33rd International Convention MIPRO*, May 2010, pp. 1226–1230.
- [7] N. Bartlow, N. Kalka, B. Cukic, and A. Ross, “Identifying sensors from fingerprint images,” in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, June 2009, pp. 78–84.
- [8] D. Moreira, A. Bharati, J. Brogan, A. da Silva Pinto, M. Parowski, K. W. Bowyer, P. J. Flynn, A. Rocha, and W. J. Scheirer, “Image provenance analysis at scale,” *IEEE Transactions on Image Processing (T-IP)*, vol. 27, no. 12, 2018.
- [9] A. Bharati, D. Moreira, A. Pinto, J. Brogan, K. Bowyer, P. Flynn, W. Scheirer, and A. Rocha, “U-phylogeny: Undirected provenance graph construction in the wild,” in *IEEE International Conference on Image Processing (ICIP)*, 05 2017.
- [10] Z. Dias, A. Rocha, and S. Goldenstein, “Image phylogeny by minimal spanning trees,” *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 2, pp. 774–788, April 2012.
- [11] Z. Dias, S. Goldenstein, and A. Rocha, “Large-scale image phylogeny: Tracing image ancestral relationships,” *IEEE MultiMedia*, vol. 20, no. 3, pp. 58–70, July 2013.
- [12] ———, “Exploring heuristic and optimum branching algorithms for image phylogeny,” *Journal of Visual Communication and Image Representation*, vol. 24, no. 7, pp. 1124 – 1134, 2013.
- [13] P. Bestagini, M. Tagliasacchi, and S. Tubaro, “Image phylogeny tree reconstruction based on region selection,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 2059–2063.
- [14] N. L. Philippe, W. Puech, and C. Fiorio, “Phylogeny of JPEG images by ancestor estimation using missing markers on image pairs,” in *Sixth International Conference on Image Processing Theory, Tools and Applications (IPTA)*, Dec 2016, pp. 1–6.
- [15] S. Milani, M. Fontana, P. Bestagini, and S. Tubaro, “Phylogenetic analysis of near-duplicate images using processing age metrics,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 2054–2058.
- [16] Z. Dias, S. Goldenstein, and A. Rocha, “Toward image phylogeny forests: Automatically recovering semantically similar image relationships,” *Forensic Science International*, vol. 231, no. 13, pp. 178 – 189, 2013.
- [17] A. Oliveira, P. Ferrara, A. D. Rosa, A. Piva, M. Barni, S. Goldenstein, Z. Dias, and A. Rocha, “Multiple parenting identification in image phylogeny,” in *IEEE International Conference on Image Processing (ICIP)*, Oct 2014, pp. 5347–5351.
- [18] M. A. Oikawa, Z. Dias, A. de Rezende Rocha, and S. Goldenstein, “Manifold learning and spectral clustering for image phylogeny forests,” *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 1, pp. 5–18, Jan 2016.
- [19] A. Melloni, P. Bestagini, S. Milani, M. Tagliasacchi, A. Rocha, and S. Tubaro, “Image phylogeny through dissimilarity metrics fusion,” in *Fifth European Workshop on Visual Information Processing (EUVIP)*, Dec 2014, pp. 1–6.
- [20] F. Costa, A. Oliveira, P. Ferrara, Z. Dias, S. Goldenstein, and A. Rocha, “New dissimilarity measures for image phylogeny reconstruction,” *Pattern Analysis and Applications*, vol. 20, 03 2017.
- [21] S. Banerjee and A. Ross, “Computing an image phylogeny tree from photometrically modified iris images,” in *Proc. of 3rd International Joint Conference on Biometrics (IJCB)*, October 2017.
- [22] ———, “Face phylogeny tree: Deducing relationships between near-duplicate face images using legendre polynomials and radial basis functions,” in *10th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS)*, Tampa, Florida, September 2019.
- [23] S. Omachi and M. Omachi, “Fast template matching with polynomials,” *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2139–2149, Aug 2007.
- [24] G. Li and C. Wen, “Legendre polynomials in signal reconstruction and compression,” in *5th IEEE Conference on Industrial Electronics and Applications*, June 2010, pp. 1636–1640.
- [25] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 3844–3852.
- [26] A. Bartoli, “Groupwise geometric and photometric direct image registration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 12, pp. 2098–2108, Dec 2008.
- [27] S. Baker and I. Matthews, “Lucas-Kanade 20 years on: A unifying framework,” *International Journal of Computer Vision*, vol. 56, no. 3, pp. 221–255, 2004.
- [28] J. G. Daugman, “Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters,” in *Journal of Optical Society of America A*, vol. 2, no. 7, July 1985, pp. 1160–1169.
- [29] L. W. Tu, “Bump functions and partitions of unity,” in *An Introduction to Manifolds*. New York, NY: Springer New York, 2008, pp. 127–134.
- [30] C. Sammut, “Density estimation,” in *Encyclopedia of Machine Learning and Data Mining*, C. Sammut and G. I. Webb, Eds. Boston, MA: Springer US, 2017, pp. 348–349.
- [31] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, “Labeled faces in the wild: A database for studying face recognition in unconstrained environments,” University of Massachusetts, Amherst, Tech. Rep. 07-49, October 2007.
- [32] “CASIA Iris Database Version 2,” <http://biometrics.idealtest.org/dbDetailForUser.do?id=2>, [Online accessed: 12th April 2019].
- [33] “CASIA Iris Database Version 4,” <http://biometrics.idealtest.org/dbDetailForUser.do?id=4>, [Online accessed: 30th August 2019].
- [34] D. Maio, D. Maltoni, R. Cappelli, J. L. Wayman, and A. K. Jain, “FVC2000: fingerprint verification competition,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, no. 3, pp. 402–412, March 2002.
- [35] L. van der Maaten and G. E. Hinton, “Visualizing data using t-SNE,” *Journal of Machine Learning Research*, vol. 9, pp. 2431–2556, November 2008.
- [36] J. J. Foo, J. Zobel, and R. Sinha, “Clustering near-duplicate images in large collections,” in *Proceedings of the International Workshop on Workshop on Multimedia Information Retrieval*. New York, NY, USA: Association for Computing Machinery, 2007, p. 2130.
- [37] X. Guo, X. Liu, E. Zhu, and J. Yin, “Deep clustering with convolutional autoencoders,” in *Neural Information Processing*, D. Liu, S. Xie, Y. Li, D. Zhao, and E.-S. M. El-Alfy, Eds. Cham: Springer International Publishing, 2017, pp. 373–382.
- [38] “Augmentor: Image augmentation library in python for machine learning,” <https://github.com/maclboice/Augmentor>, [Online accessed: 3rd January 2020].
- [39] “CelebA dataset,” <http://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>, [Online accessed: 3rd January 2020].
- [40] H.-J. Chen, K.-M. Hui, S.-Y. Wang, L.-W. Tsao, H.-H. Shuai, and W.-H. Cheng, “BeautyGlow: On-demand makeup transfer framework with reversible generative network,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[41] C. Ye, R. C. Wilson, C. H. Comin, L. da F. Costa, and E. R. Hancock, “Entropy and heterogeneity measures for directed graphs,” in *Similarity-Based Pattern Recognition*, E. Hancock and M. Pelillo, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 219–234.



Sudipta Banerjee received the B.Tech. degree in Electronics and Communication Engineering from West Bengal University of Technology, India, in 2011, and the M.E. degree in Electronics and Tele-Communication Engineering with specialization in Control Engineering from Jadavpur University, India, in 2014. She is currently pursuing her Ph.D. degree in Computer Science and Engineering at Michigan State University under the supervision of Dr. Arun Ross in the iProBe research laboratory. Her research interests include image forensics in the context of digitally altered biometric images.



Arun Ross is the John and Eva Cillag Endowed Chair in the College of Engineering and a Professor in the Department of Computer Science and Engineering at Michigan State University, United States. He received the B.E. (Hons.) degree in Computer Science from BITS Pilani, India, and the M.S. and PhD degrees in Computer Science and Engineering from Michigan State University, United States. He served as the faculty at West Virginia University between 2003 and 2012 where he received the Benedum

Distinguished Scholar Award for excellence in creative research and the WVU Foundation Outstanding Teaching Award. Ross was a recipient of the NSF CAREER Award and was designated a Kavli Fellow by the US National Academy of Sciences in 2006. He received the JK Aggarwal Prize in 2014 and the Young Biometrics Investigator Award in 2013 from the International Association of Pattern Recognition. He is the co-author of the textbook, *Introduction to Biometrics* and the monograph, *Handbook of Multibiometrics*.