

# Contour-as-Face Framework: A Method to Preserve Privacy and Perception

Journal of Marketing Research  
 2020, Vol. 57(4) 617-639  
 © American Marketing Association 2020  
 Article reuse guidelines:  
[sagepub.com/journals-permissions](http://sagepub.com/journals-permissions)  
 DOI: 10.1177/0022243720920256  
[journals.sagepub.com/home/mrj](http://journals.sagepub.com/home/mrj)



Yinghui Zhou, Shasha Lu, and Min Ding

## Abstract

Consumers and marketers use facial information to make important inferences about others in many business contexts. However, consumers and firms are increasingly concerned about privacy and discrimination. To address privacy–perception trade-offs, the authors propose a novel contour-as-face (CaF) framework that transforms face images into contour images incorporating both the nonoutline and outline features of facial parts. In three empirical studies, the authors (1) compare human perceptions of face and contour images along 15 dimensions commonly assessed in marketing contexts; (2) investigate the effectiveness of contour images for protecting anonymity related to identity, age, and gender; and (3) implement the CaF framework in a real-life online dating context. Results show that the CaF framework effectively resolves privacy–perception trade-off problems by preserving the information that is useful for humans to make inferences about many relevant perceptual dimensions in marketing while making it virtually impossible for humans to infer identity and very difficult to infer age and gender accurately—two critical discrimination factors. Results from the field implementation demonstrate the feasibility and value of using the CaF framework for real-life decision making.

## Keywords

facial contour, face perception, facial privacy, Fourier transformation, online dating

Online supplement: <https://doi.org/10.1177/0022243720920256>

The face plays an important role in decision-making processes in various contexts, including politics (e.g., Todorov et al. 2005), law and justice (e.g., Zebrowitz and McDonald 1991), leadership (e.g., Rule and Ambady 2008), the labor market (e.g., Bóo, Rossi, and Urzúa 2013; Ruffle and Shtudiner 2014), and marketing (e.g., DeShields, Kara, and Kaynak 1996; Kahle and Homer 1985; Keh et al. 2013; Valentine et al. 2014; Xiao and Ding 2014). In the social networking era, the amount of facial information available online has increased dramatically. In 2013, Facebook revealed that its users had uploaded 250 billion photos and were uploading 350 million new photos every day (Smith 2019).

This spread of facial images creates a “face dilemma” for consumers. On the one hand, the online presence of facial information helps people make decisions when using social networking sites such as Facebook, Match.com, and LinkedIn; on the other hand, it may increase the risk of misuse, exploitation, and other improper disclosures of identity information (Cobb and Kohno 2017; Senior and Pankanti 2011) as well as face-based discrimination (e.g., Ruffle and Shtudiner 2014). The European Union’s General Data Protection Regulation reflects the public’s heightened awareness of the need to protect personal information, including facial data.

The face dilemma also affects marketing decisions in various contexts, including service employee selection/allocation, professional service evaluation, and customer relationship management (e.g., Gomulya et al. 2017; Keh et al. 2013). Customers use facial information to assess an employee’s emotional state, attractiveness, and trustworthiness. These assessments in turn affect customer satisfaction, loyalty, and service quality evaluations, especially in the online environment (e.g., Rezlescu et al. 2012). Therefore, it is important for firms to consider facial information when allocating employees to interact with different customers in different contexts. HireVue, a hiring intelligence firm, uses face perception modeling to determine the fitness of potential employees (often those who will be interacting with customers) for companies such as IBM, Unilever, and Hilton. Yet firms must avoid discrimination

Yinghui Zhou is Assistant Professor of Marketing, College of Management, Shenzhen University, and Research Associate at the Institute for Sustainable Innovation and Growth (iSIG), School of Management, Fudan University, China (email: yhzhou@szu.edu.cn). Shasha Lu is Lecturer (Assistant Professor) of Marketing, Cambridge Judge Business School, University of Cambridge, UK (email: s.lu@jbs.cam.ac.uk). Min Ding is Bard Professor of Marketing, Smeal College of Business, and Affiliate Professor, College of Information Sciences and Technology, Pennsylvania State University, USA (email: minding@psu.edu).

based on demographics that can be inferred from an employee's face, such as gender and age (Altonji and Blank 1999). In many places, laws have been passed to prohibit recruitment and wage discrimination based on facial appearance (Ellis and Watson 2012).

Resolving the face dilemma thus requires protecting different facial privacy data in two types of applications: in Type I applications (e.g., mostly related to consumers), facial identities must be protected to reduce the risk of violating individuals' privacy, and in Type II applications (e.g., employee allocation, professional service evaluation), demographic traits such as gender and age must be protected to prevent face-based discrimination against individuals from certain groups. In both types of applications, it is important to preserve facial cues that help people make useful inferences when making decisions. We propose a novel solution to address the face dilemma: the contour-as-face (CaF) framework. This framework asks a human to make judgments and decisions using face contours as a proxy for faces and then directly models the judgments and decisions as a function of face contours using Fourier transformation. We show that the CaF framework can anonymize face images for Type I and Type II applications without (substantially) compromising the utility of facial information in decision making.

In the sections that follow, we first review the existing literature and describe the motivations for our work by discussing five marketing domains related to the face dilemma. We then introduce the CaF framework as a way to resolve the face dilemmas in Type I and Type II applications. We present three empirical studies in which we investigated the effectiveness of the CaF framework in addressing the face dilemma and CaF's usefulness in real-life business applications. We conclude with a general discussion of the use of the CaF framework for face-related business contexts, as well as future research directions.

## Literature Review and Motivation

When seeing a face, the human brain instantly processes facial information and makes inferences about the person along various dimensions such as identity, demographics, physical state, emotional state, and social traits (Todorov 2017). Table 1 summarizes the types of face-based perceptions and application contexts commonly studied in the literature.

The literature on face-based perceptions is too vast to be covered in this article, so we focus our discussion on the widespread usage of facial data in marketing and business contexts. We first provide a brief review of previous research on the use of face perceptions in business-related contexts; we subsequently discuss privacy issues associated with using facial data. We then describe the motivation for our work: to resolve the two types of privacy–perception trade-offs associated with the use of facial data in various marketing domains. Finally, we discuss the challenges and limitations of existing solutions in addressing the trade-off problems.

## Face-Based Perceptions

Face-based perceptions have been widely used in business-related contexts. First, studies have consistently shown that face-based perceptions and preferences can affect individuals' evaluations of others in various marketing contexts, including advertising, personal selling, service encounters, and networking/dating (e.g., Small and Verrochi 2009; Valentine et al. 2014; Xiao and Ding 2014). Customers demonstrate higher purchase intentions or satisfaction levels when spokespersons, salespeople, and service providers are perceived to be more attractive than their counterparts or demonstrate positive emotional states (e.g., DeShields, Kara, and Kaynak 1996; Hennig-Thurau et al. 2006; Kahle and Homer 1985; Wang et al. 2017). As a result, social inferences from job candidates' facial information are widely used in the employee selection process (e.g., Caers and Castelyns 2011). For example, face-based perceptions such as maturity, dominance, and competence have been shown to influence the selection of employees (e.g., Gorn, Jiang, and Johar 2008; Graham, Harvey, and Puri 2016; Keh et al. 2013). Second, perceptions from customer facial data can be used to improve targeting and segmentation, firms' understanding of customer preferences, and marketing effectiveness (Xiao, Kim, and Ding 2013). For example, Lu, Xiao, and Ding (2016) analyzed customers' facial expressions in videos to infer the customers' product preferences. Affectiva, an emotion analytics company, helps brands measure customer responses to digital content through facial frames of customers captured while they are viewing digital content online. Lapetus Solutions, an artificial intelligence company, infers various types of information related to customers' health conditions from customers' facial images, which has been used by insurance companies for segmentation and targeting.

## Facial Privacy

The use of facial data in business contexts is a double-edged sword. As the main source of personal biometric information, the face can be used to identify a person's demographics or identity. The prevalence of publicly available facial information gives rise to risks and concerns about facial privacy, including face-based discrimination and identity privacy violations.

*Face-based discrimination.* Face perceptions of demographic information can bias human judgments and decisions (for a review, see Olivola, Funk, and Todorov [2014]) in business contexts. Studies have shown that job allocation decisions are affected by perceptions of employees' gender and sexual dimorphism (e.g., Altonji and Blank 1999; Landau 1995). For example, Acquisti and Fong (2020) found that access to information about candidates' sexual dimorphism and ethnicity via online social network profiles could bias employers' callback decisions. The use of face information in automated hiring software is facing widespread public backlash, with fears that it may lead to more biased hiring decisions (Martin 2018).

**Table 1.** Relevant Literature on Face-Based Perception.

Perceptual Domain	Quantitative Facial Feature Measurements							Representative References	
	Dimension	Qualitative Description	Texture	Size	Spatial Relation	Angle, Curvature	Contour/Outline		Application Context(s)
Identification	Personal identity	N	Y	Y	Y	Y	N	Video game, ID registration, video surveillance, suspect tracking	Li and Jain (2011); Zhao et al. (2003)
	Age	N	Y	Y	Y	Y	N	Identification, service eligibility	Berry and McArthur (1986); Rhodes (2009)
Demographics/ background information	Gender	N	Y	Y	Y	Y	N	Human-computer interaction	Yang, Li, and Ai (2006)
	Sexual dimorphism	Y	Y	N	N	N	N	First impression, dating, health	Perrett et al. (1998)
	Sexual orientation	N	N	Y	Y	Y	N	Employee selection	Wang and Kosinski (2018)
	Height	Y	N	N	N	N	N	Leadership selection	Re et al. (2013)
	Ethnicity	Y	N	N	N	N	N	Human-computer interaction, surveillance	Fu, He, and Hou (2014)
Emotional states	Social status	Y	N	N	N	N	N	Employee selection	Bjornsdottir and Rule (2017)
	Occupation	Y	N	N	N	N	N	Leadership selection	Olivola, Eubanks, and Lovelace (2014)
	Facial expressions	Y	Y	Y	Y	Y	N	Market research, robotics	Fasel and Luetttin (2003)
Physical states	Arousal	N	N	N	N	Y	N	Driver fatigue detection, human-computer interaction, interactive gaming	Gu, Ji, and Zhu (2002)
	Physical strength	Y	N	N	N	N	N	First impression	Todorov et al. (2008)
Social traits	Fatness	Y	N	N	N	N	N	Health	Coetzee, Perrett, and Stephen (2009)
	Health condition	Y	Y	N	N	N	N	Dating, employee selection	Henderson et al. (2016)
	Pleasantness	Y	N	N	N	N	N	First impression	Willis and Todorov (2006)
	Attractiveness	N	N	N	Y	N	N	Networking and dating	Cunningham (1986); Cunningham et al. (1990)
	Approachability	N	Y	Y	Y	Y	N	First impression	Vernon et al. (2014)
	Confidence	Y	Y	Y	N	N	N	First impression	Vernon et al. (2014)
	Intelligence	Y	N	N	N	N	N	Dating, employee selection, leadership selection	Oosterhof and Todorov (2008); Zebrowitz et al. (2002)
Trustworthiness	Trustworthiness	Y	N	N	Y	N	N	Lie detection, first impression	Oosterhof and Todorov (2008); Todorov et al. (2008)
	Aggressiveness	N	N	Y	Y	N	N	Leadership selection, criminal profiling	Carré, McCormick, and Mondloch (2009)
	Dominance	Y	Y	Y	Y	Y	N	Leadership selection	Oosterhof and Todorov (2008); Vernon et al. (2014); Geniole et al. (2015)
Baby-facedness	Baby-facedness	Y	Y	Y	Y	N	N	First impression	Berry and McArthur (1986); Zebrowitz and McDonald (1991)
	Competence	Y	N	N	N	N	N	Leadership selection, employee selection	Todorov et al. (2005); Rule and Ambady (2008)

Notes: The complete references of Table 1 can be found in Web Appendix D. Y = the corresponding aspect has been studied in the representative references; N = the corresponding aspect has not been studied in the representative references.

**Table 2.** Application Domains for Privacy–Perception Trade-Off.

Dimension	Online Dating	Social Networking	Professional Service Evaluation	Employee Selection/Allocation	Customer Relationship Management
Identification	N	N	—	—	N
Age	Y	N	N	N	Y
Sexual dimorphism	Y	N	N	N	Y
Emotional state (e.g., smile)	Y	—	Y	—	—
Health	Y	—	—	—	—
Fatness	Y	—	—	—	—
Arousal	Y	—	—	—	—
Aggressiveness	Y	Y	—	—	—
Approachability	Y	Y	Y	Y	Y
Attractiveness	Y	—	Y	Y	—
Baby-facedness	Y	Y	Y	—	—
Confidence	Y	Y	Y	Y	—
Dominance	Y	Y	Y	Y	Y
Intelligence	Y	Y	Y	Y	—
Pleasantness	Y	Y	Y	Y	Y
Trustworthiness	Y	Y	Y	Y	—

Notes: Y = perception of this dimension is desired in the corresponding application context; N = perception of this dimension is not desired in the corresponding application context; — = perception of this dimension is irrelevant or its relation to the application context is not clear.

Despite a large body of literature on the subject, to the best of our knowledge, scholars have yet to propose an effective method to protect disadvantaged groups from face-based discrimination, especially in the online recruitment and job allocation contexts.

**Identity privacy.** As the use of visual data in business contexts has increased, so too have customers' concerns about identity privacy when sharing images and videos (Cobb and Kohno 2017; Gross and Acquisti 2005). In a recent survey, more than 75% of the respondents reported that they did not feel comfortable with firms' use of facial data in either a commercial or human resource context (Ada Lovelace Institute 2019). There are many face images that circulate daily on social media platforms (e.g., Facebook, LinkedIn), photo-sharing apps (e.g., Instagram), and online dating websites. It is likely that a human may look at a photo on an "anonymous" website and recognize the pictured person from memory. Likewise, algorithms can match a photo on an anonymous website, such as a dating site, to a photo on a nonanonymous website such as LinkedIn or Facebook, thus revealing the identity of the person on the anonymous website (Cobb and Kohno 2017).

From a firm's perspective, the level of protection of customers' personal data affects a firm's marketing and financial performance. Privacy has been found to be a key driver of online trust (Martin, Borah, and Palmatier 2017). Higher levels of data security and privacy protection, as perceived by customers, can lead to higher levels of trust and increased purchase intentions, effective purchase behavior, and loyalty to a firm/brand (Martin, Borah, and Palmatier 2017). As a result, protecting customer privacy while gaining insights from customer data has been recognized as the top challenge for chief marketing officers of large firms such as HSBC, Yahoo, and Target, all of

which have experienced data breaches in recent years (Benes 2018).

### *Motivation: Addressing Privacy–Perception Trade-Offs*

These conflicting objectives must be addressed to provide value for stakeholders while reducing (or even eliminating) privacy infringement and discrimination. We formally define two major types of privacy–perception trade-offs when using facial data. A Type I identity protection–perception trade-off occurs when an optimizer (e.g., firm) attempts to preserve facial cues that are useful for human (e.g., customer) judgments or decisions while reducing facial cues that contain personally identifiable information, thus protecting facial identities. A Type II discrimination prevention–perception trade-off occurs when an optimizer attempts to preserve facial cues that aid decision making while reducing facial cues that may facilitate appearance-based discrimination.

Based on the literature, we identified five marketing domains (i.e., dating, networking, professional service evaluation, employee selection/allocation, and customer relationship management) that have pressing needs for effective tools to solve Type I (identity protection–perception) and/or Type II (discrimination prevention–perception) trade-off problems. As Table 2 shows, faces are used to achieve different objectives in each of these contexts.

Examples of Type I contexts may include dating and customer relationship management, where firms need to represent the target faces (the faces that are being observed) in a way that can serve as a basis for forming a "first impression" but prevent identification by perceivers (customers or persons who observe the face information). For example, in the online dating context, a face is typically linked to a user's dating profile.

**Table 3.** A Comparison of Existing Methods for Addressing Privacy–Perception Trade-Off.

Application Domains	Key Aspects in Addressing Privacy–Perception Trade-Off	Face Anonymization Methods		Face Perception Modeling Methods		This Research (CaF)
		Masking, Pixelization, Blurring	Face Replacement	PS	EF	
Face representation presented to human	Anonymization of identity information	High	High	—	—	High
	Anonymization of demographic information	Medium	High	—	—	High
	Preservation of information useful for face perceptions	Low	Low	—	—	High
Face perception modeling	Accuracy in capturing original face information	—	—	Medium	High	High
	Robustness to face image quality or scale change	—	—	Medium	Low	High
	Computational efficiency	—	—	High	Low	High

Notes: High = the performance of the method on a given dimension is high; Medium = the performance of the method on a given dimension is medium; Low = the performance of the method on a given dimension is low; — = does not apply.

Users may want to conceal their facial identities while retaining the ability to rely on face perceptions (including emotional and physical states and social traits) when screening potential dates. In the context of customer relationship management, a face is typically linked with a target customer. Employees may want to provide customized customer service based on information inferred from the customer's face (e.g., approachability) without invading the customer's identity privacy.

Examples of Type II contexts may include employee selection/allocation and professional service evaluation, where firms need to represent the target faces (e.g., service providers' faces) in a way that can serve as a basis for matching the perceivers' (e.g., customers') preferences without inducing face-based discrimination. For example, in the employee selection/allocation context, a face is typically linked with a (potential) employee. Firms may want to improve the target customer's experience, satisfaction, and purchase intention by selecting/allocating employees whose facial features match the target customer's preferences on certain perceptual dimensions (e.g., pleasantness, trustworthiness) without discriminating employees in terms of demographics (e.g., gender, age). In the professional service evaluation context, a face is typically linked with a professional service provider (e.g., accountant, banker) who may wish to reveal their facial information on dimensions that would help customers evaluate their service quality (e.g., approachability, confidence, trustworthiness, intelligence) while preventing demographics-based discrimination in customers' judgments or decisions.

Both types of trade-off problems are potentially relevant to users in the social networking context (e.g., Facebook, LinkedIn), where a face is usually linked to a user's personal profile. In addition to protecting their own identity privacy, users may want to avoid attracting interest or experiencing discrimination from others based on their gender or age.

### Challenges in Resolving the Privacy–Perception Trade-Offs

Despite the large body of literature on face perceptions and facial privacy, to the best of our knowledge, scholars have not

studied how to effectively address the privacy–perception trade-offs. The most relevant research domains are face anonymization, which focuses on anonymizing face identity information, and face perception modeling, which focuses on understanding the face information used in various face perceptions. We reviewed the existing methods in both domains based on their abilities to represent and preserve face information for human perception and to mask identity and demographic information. Table 3 provides a comparison of these methods.

**Face anonymization.** To reduce the risk of identification in the course of facial data processing, considerable research in computer science has been directed toward the anonymization of faces—that is, the representation of original faces with modified or concealed face information. One way to conceal facial identity is by reducing the amount of identifiable information contained in a face image. Naive methods such as masking, pixelization, and blurring are commonly used, all of which involve removing critical information about important facial parts (e.g., eyes, nose) from the face (Newton, Sweeney, and Malin 2005). By doing so, however, these naive methods can quickly reduce the reliability (i.e., the extent of agreement with perceptions based on the original face image) of human perception on privacy-insensitive dimensions of the face, including emotional states and social traits. Another way to anonymize facial identity information is by replacing the original face image with a surrogate face that has been constructed from a cluster of selected faces (Newton, Sweeney, and Malin 2005). This face replacement method can be modified to preserve the original face's demographic information (e.g., gender, age) by selecting surrogate faces with the same demographic characteristics (Du et al. 2014). However, it is unable to preserve facial information used for face perceptions on other dimensions, such as emotional states and social traits.

**Face perception modeling.** Unlike face anonymization research, face perception modeling focuses on identifying key information in an original face image that influences human

perception. Previous studies have proposed two main approaches in this domain.

One is the physiognomic (PS) approach, whereby the target face is represented by PS features derived from the coordinates of facial landmarks on key facial parts (i.e., eyebrows, eyes, nose, mouth, and face). The PS approach has been a popular method of investigating the drivers of various face perceptions (Todorov et al. 2008; Vernon et al. 2014). However, it suffers from several inherent problems, including arbitrary facial landmark placement, arbitrary facial feature selection, and a lack of information for the regions between landmarks (Xiao and Ding 2014). Moreover, typical PS measures fail to accurately capture the subtle variations among faces (e.g., shape of face; Xiao and Ding 2014). Because the face image is reduced to a set of points in the PS approach (see Web Appendix C for an example), it is difficult for a human being to make meaningful inferences from it.

The other approach is the eigenface (EF) approach, whereby the target face is represented by a linear function of eigenvectors (referred to as “eigenfaces”) derived from the covariance matrix of a set of face images (for details, see Turk and Pentland [1991]). It was widely used for automatically detecting and matching faces from images but has also been used as a modeling method to understand customers’ preferences for faces (for details, see Xiao and Ding [2014]). There are several limitations of the EF approach: its effectiveness relies heavily on the texture information contained in the target face image and the training images, making it sensitive to changes in the image quality, illumination, scale, and facial expression (Chellappa, Wilson, and Sirohey 1995) that are prevalent in natural face images. It is computationally time-consuming to calculate the eigenvectors when new faces are added to the training database. More importantly, it is not feasible for humans to make any types of inferences about the target face from the function of eigenfaces.

From our review, neither the face anonymization nor face perception modeling methods can effectively represent a face in a way that masks privacy information (e.g., identity, age, gender) while allowing a human to make reasonably reliable perceptions. In the next section, we propose a CaF framework to address this gap in the literature and practice.

## The CaF Framework

To address the two types of trade-off problems in the five marketing domains, it is important that firms represent face information in a way that can effectively (1) conceal the facial cues that lead to undesired perceptions (e.g., identity, age, gender), thus reducing the privacy risk and encouraging the sharing of face information; (2) preserve facial cues that lead to desired perceptions (e.g., emotional states, social traits), thus aiding human judgment and decision making; (3) allow for direct modeling of relations between the face information and the perceiver’s judgments/decisions, thus improving business practices related to faces; and (4) encompass diverse real-life business contexts.


In this section, we begin by examining the visual cues driving different types of face perceptions to identify a solution that can satisfy these four criteria. We then propose a CaF framework that is grounded in social psychology, neuropsychology, and computer vision research as a solution. The CaF framework has two main components: the representation component presents facial contours instead of face images to the perceivers, and the modeling component quantifies the relationship between facial features and perceptions using techniques from the face perception and computer vision literature streams.

### Visual Cues Driving Face Perceptions

Humans use two main types of visual cues in face perception: texture cues and structural cues (Meinhardt-Injac, Persike, and Meinhardt 2013). Texture cues reflect variations in intensity and color on the facial surface. Evidence shows that texture cues provide critical information regarding a face’s identity, gender, and age and are thus essential in human face recognition, gender classification, and age perception (e.g., Bruce and Langton 1994; Johnston, Hill, and Carman 1992; Porcheron, Mauger, and Russell 2013). For example, a human’s face recognition performance drops by as much as 20%–30% when pigmentation and shading information are eliminated (Johnston, Hill, and Carman 1992). Bruce and Langton (1994) found that skin texture information plays an important role in human gender classification. Porcheron, Mauger, and Russell (2013) found that skin texture and facial contrast are essential for individuals to estimate age with reasonable accuracy. These findings indicate that removing texture cues may effectively mask the facial information that enables identification or induces gender/age-based discrimination.

Structural cues consist of the shapes of facial parts and their spatial relationships. The shapes of facial parts are largely determined by their contours (i.e., boundary outlines), which contain much of the biological information that is fundamental in face perception (Lestrel 2008). Considerable converging evidence shows that the contours (used interchangeably with “shape” in prior studies) of facial parts as well as their spatial configurations can effectively explain human face perception (Tanaka and Farah 1993) along various dimensions, including attractiveness, approachability, confidence, trustworthiness, dominance, baby-facedness, and emotional state (e.g., Todorov et al. 2008; Torrance et al. 2014; Vernon et al. 2014). Although evidence shows that structural cues also contribute to perceptions of identity, age, and gender when combined with texture cues (e.g., Bruce and Langton 1994; George and Hole 2000), the independent role of structural cues in face recognition or age/gender classification remains unclear.

These findings suggest that the trade-off problems may be resolved by removing texture cues from the face while preserving structural cues as captured in facial contours. Drawing on these findings, we propose a radical approach that uses facial contours as a solution satisfying the four criteria delineated at the beginning of this section. Next, we describe the two components of the CaF framework.



**Figure 1.** An example of the CaF representation.

### CaF Representation

In previous studies, face representation involved presenting a face image to a viewer. The CaF representation component presents a facial contour image instead of a face image as the visual stimulus. This study marks the first time anyone has adopted this representation technique in academia or practice. In line with the literature and practice, the proposed CaF representation component preserves complete contour information about five key facial parts—the eyebrows, eyes, nose, mouth, and overall face—as well as their spatial relations. Figure 1 provides an example of the CaF representation. Additional facial parts (e.g., ears) and three-dimensional contour information may also be used when appropriate.

Compared with the existing face representation approaches used in face anonymization (see Table 3), the CaF representation preserves more information related to privacy-insensitive face perceptions. It can thus be used to elicit human perceptions for decision making in many business contexts. This is impossible using modified or concealed face representations and is not feasible using PS or EF representations.

### CaF Modeling

Contour-as-face modeling involves two key steps: extracting facial features from the CaF representation and quantifying the effect of the facial features on face perceptions, usually through statistical models or machine learning approaches. The CaF modeling component extracts two types of information from the CaF representation: outline features (i.e., features capturing the holistic contour information of the five key facial parts) and nonoutline features (i.e., features capturing the configural information of facial parts, including the sizes and angles of the five facial parts and their spatial arrangements).

**Outline feature extraction.** To extract the outline features, the CaF modeling component applies the centroid distance–Fourier descriptor (CD-FD) method, which uses discrete Fourier transformation to convert the contours to mathematically precise Fourier coefficients. These Fourier coefficients are then converted into input variables (referred to as “Fourier descriptors”) for statistical or machine learning models. The CD-FD method is a popular technique used to accurately represent target contour information with features that remain stable through various transformational changes. The following

steps briefly describe the method (for a detailed description of the method, see Web Appendix A).

**Step 1:** Acquire complete contour boundary outline information from the target object image (e.g., key facial parts such as the eyebrows, eyes, nose, mouth, and overall face). The contour boundary information can be extracted either manually or by applying contour/edge detection to the target image. After aligning the starting point and the nominal orientation of the contours, the complete contour boundary outline can be represented by a set of coordinates of contour boundary points, given by  $\{(x_t, y_t), t = 0, 1, \dots, N - 1\}$ , where  $N$  is the number of boundary points.

**Step 2:** Derive the centroid distance function from the coordinates of the contour boundary points. The centroid distance function,  $r(t)$ , is given by the distance of the boundary points from the centroid of the contour, denoted as  $\{x_{\text{centroid}}, y_{\text{centroid}}\}$ , and is represented by

$$r(t) = \sqrt{(x_t - x_{\text{centroid}})^2 + (y_t - y_{\text{centroid}})^2}, \quad (1)$$

$$t = 0, 1, \dots, N - 1$$

where  $x_{\text{centroid}} = (\sum_{t=0}^{N-1} x_t) / N$ , and  $y_{\text{centroid}} = (\sum_{t=0}^{N-1} y_t) / N$ . The use of centroid distance function ensures that the Fourier descriptors are invariant to translation changes due to the subtraction of the centroid.

**Step 3:** Obtain Fourier coefficients by applying discrete Fourier transformation to the centroid distance function  $r(t)$ . The Fourier coefficients are given by

$$a_n = \frac{1}{N} \sum_{t=0}^{N-1} r(t) \exp\left(-i2\pi \times \frac{nt}{N}\right), \quad n = 0, 1, \dots, N - 1, \quad (2)$$

where  $i$  is the imaginary unit. Each Fourier coefficient  $a_n$  is a unique complex number consisting of a magnitude component, denoted  $|a_n|$ , which captures the scale change in the contour at a certain frequency, and a phase component, denoted  $\phi_n$ , which captures the rotational change in the contour at the frequency. Both the magnitude and phase components of Fourier coefficients contain important information about the complete contour. The lower-frequency coefficients contain information about the global contour, and the higher-frequency coefficients contain information about the finer local details of the contour. It is worth noting that the impact of phase component on contour depends on the magnitude component (Fieguth, Bloore, and Domsa 2004); the effects of phase components thus can be captured by the products of the phase angles and their corresponding magnitudes.

**Step 4:** Obtain the scale-invariant Fourier coefficients by normalizing the magnitudes of the Fourier coefficients by the zero-frequency component, or  $|a_0| = |\sum_{t=0}^{N-1} r(t) / N|$ . The magnitude component of a scale-invariant Fourier coefficient is given by

$$|b_n| = \frac{|a_n|}{|a_0|}, \quad (3)$$

where  $n = 1, \dots, N - 1$ , and  $|b_0| = 1$ .

Step 5: Keep the first  $K$  scale-invariant Fourier coefficients that are needed to recover the original contour information with high accuracy. Although discrete Fourier transformation often results in many Fourier coefficients, only a small set (e.g., 10–15) is generally needed to capture the majority of variation among the contours and describe the original contour information with high accuracy (Zhang and Lu 2002). These Fourier coefficients are kept for further analysis.

Step 6: Construct the Fourier descriptors of the contour. Each Fourier descriptor is a vector derived from the corresponding scale-invariant Fourier coefficient. For the purpose of statistical modeling, each Fourier descriptor consists of the normalized magnitude  $|b_n|$ , and its product with the corresponding phase angle; that is,  $|b_n| \times \phi_n$ , denoted by  $d_n = \{|b_n|, |b_n| \times \phi_n\}$ ,  $n = 1, \dots, K$ . The  $K$  Fourier descriptors are then used to represent the outline features of the contour, denoted by  $D = \{d_1, \dots, d_K\}$ .

**Nonoutline feature extraction.** The nonoutline features are extracted following the face perception literature (e.g., Vernon et al. 2014). First, align the contour image so that the centroids of the left and right eyes are level. Second, calculate five size variables (i.e., the sizes of the eyebrows, eyes, nose, mouth, and overall face)<sup>1</sup>; four angle variables (i.e., the angles of the eyebrows, eyes, mouth, and overall face); and six spatial placement variables (i.e., eyebrow separation, eye separation, and the vertical placements of the eyebrows, eyes, nose, and mouth) from the aligned contour image. Finally, normalize the extracted features based on face size (where applicable) to obtain scale-invariant features, resulting in 14 variables that are then used to represent the nonoutline features of the CaF representation, denoted by  $X$ .

**The CaF model specification.** The CaF model quantifying the effect of facial contours on face perceptions thus can be given as

$$Y = f(Z, X), \quad (4)$$

where  $Y$  represents the perception of a facial contour image and  $X$  represents the nonoutline features of the contour image.  $Z = \{D^1, \dots, D^5\}$  represents the outline features of the five key facial parts, where  $D^m = \{d_1^m, \dots, d_{K_m}^m\}$ ,  $m = 1, \dots, 5$ ;  $K_m$  is the number of Fourier coefficients chosen to describe the contour information of facial part  $m$ ; and  $d_n^m = \{|b_n^m|, |b_n^m| \times \phi_n^m\}$ ,  $n = 1, \dots, K_m$ .

We note that our proposed CaF modeling method is a general tool that can be applied to face perception modeling regardless of whether the perception is based on face images or facial contour images (as used in the CaF representation). It can also be used for a wide variety of face perception or face

preference modeling applications (e.g., face-related recommendation systems).

Compared with the existing methods used in face perception modeling (see Table 3), the CaF modeling method has several advantages. First, unlike the PS approach, which only captures a small portion of the contour outline information of facial parts, the Fourier descriptors serve as a holistic measure that captures the original facial contour with little loss of information. Second, the CaF method is an objective and robust measure of facial features. Unlike the PS approach, the CaF method does not make a priori assumptions about the importance of specific facial information (e.g., facial landmarks). Unlike the texture features used in the EF approach, the contour information used in the CaF approach is largely invariant to changes in image quality and background. Moreover, the Fourier descriptors are invariant to translation and scale change (Zhang and Lu 2002). Third, the CaF method is easy to implement and computationally simple, which is especially helpful when working with large face data sets (e.g., images or videos on social media). Finally, because of the CaF method's ability to explicitly describe the contour at different levels of detail (using an additive, orthogonal Fourier series), the researcher can control the contour information entering the model through the number of Fourier coefficients, thus enabling understanding of the effects of different types of facial contour information on face perceptions.

In the following sections, we present analyses from three empirical studies that examine the effectiveness of the proposed CaF framework in addressing the face dilemma. Study 1 investigates the effectiveness of the CaF framework (CaF representation and CaF modeling) in capturing and preserving face information for human perceptions on commonly assessed dimensions and the underlying mechanisms. Study 2 examines the effectiveness of using the CaF representation for masking private information (e.g., identity, age, gender) in real life. We implemented the CaF framework in a real-life online dating context in Study 3 to investigate the feasibility of using the CaF framework for decision making in actual business contexts.

## Study 1: The Effectiveness of CaF for Preserving Perception

We aimed to answer the following research questions in Study 1: (1) Do contours of facial parts contain important information that people use in perceptions of faces? (2) If the facial contour information does affect face perceptions, can people make reliable inferences from the CaF representation on the dimensions that are commonly assessed in marketing contexts? (3) What aspects of the facial contour image drive the contour-based perceptions across different perceptual dimensions? To answer the three research questions, we first characterized the relationship between facial contour information and human perceptions of faces on various dimensions using the CaF modeling method and compared the model performance to that of two benchmarks that model face perception. Then, we compared

<sup>1</sup> Because the left and right eyebrows/eyes are generally symmetrical, here we focus only on the left eyebrow and left eye contour information for model simplicity.



human perceptions on various dimensions after exposure to two types of face representations (facial contour images vs. face images). Finally, we analyzed which aspects of facial contour images influence humans' perceptual inferences across different perceptual dimensions using variations of the CaF model.

## Study Design

**Stimulus faces and contour images.** To elicit real-life face perception experiences, we used images of real faces from publicly available sources, including university faculty webpages, company employee webpages, and job search sites. We selected 216 stimulus faces (107 men, 109 women) that satisfied the following criteria: (1) not a celebrity (and no resemblance to a celebrity), (2) reasonable resolution (photo size larger than 100 KB), (3) frontal face view, and (4) contours of all facial parts (eyebrows, eyes, nose, mouth, overall face) visible and clear. This ensured that inferences were not affected by poor image quality or nonvisual factors that are irrelevant to face or contour information (e.g., using faces of well-known people). To generate the face image stimuli, we cropped the face photos to include only the face and neck area, thereby reducing the effects from visual factors other than face. To generate the contour image stimuli, we manually extracted the contours of five key facial parts for each face image: the eyebrows, eyes, nose, mouth, and overall face shape (for an example, see Figure 1). The 216 face images and corresponding facial contour images were used as visual stimuli for the study.

**Data collection.** Following the face perception literature (e.g., Todorov et al. 2008; Vernon et al. 2014), we used 15 perceptual dimensions that have proven to be critical in forming a first impression. These dimensions can be categorized as follows: demographics (e.g., age, sexual dimorphism); emotional state (e.g., smile); apparent physical state (e.g., health, fatness); and social traits (e.g., aggressiveness, approachability, arousal, attractiveness, baby-facedness, confidence, dominance, intelligence, pleasantness, trustworthiness). We applied a between-subjects design to eliminate any potential confusion (and interaction) caused by asking the same person to provide both face-based and contour-based perceptions.

Three hundred fifty-four participants (98 men, 256 women) were in the face image group. Each participant was randomly presented with a (male or female) face image and asked to provide ratings on the 15 dimensions using a seven-point scale. The participants were asked to provide ratings only if they did not know the individuals depicted in the images; they were also allowed to decide how many images to rate. To eliminate potential own-gender perception bias (Hills et al. 2018), we used only those ratings for face images of the opposite gender (about half of the total ratings collected) in the subsequent data analysis. Our conclusions were not qualitatively different when all ratings were included. On average, each participant rated 9.3 face images of the opposite gender.

Five hundred seventy participants (158 men, 412 women) were in the contour image group. We used the same procedure as for the face image group but substituted contour images for the face images. Because this was the first time that contours were used as stimuli for face perceptions, we wanted to avoid introducing any artificial effects related to asking participants to rate dimensions that could be difficult to evaluate from contours. Therefore, we asked participants to provide ratings for any of the 15 dimensions they found interesting and were able to rate using facial contours.<sup>2</sup> The participants in the contour group were shown only contour images of the opposite gender to avoid own-gender perception bias (Hills et al. 2018). On average, each participant rated 25.9 contour images of the opposite gender.

**Measures.** For each face/contour image in the database, we obtained three sets of information: (1) nonoutline features, (2) outline features, and (3) the participants' average ratings of the face image and corresponding contour image on the 15 dimensions. We derived 14 nonoutline features for each face/contour image following the procedure described in the CaF modeling subsection. To facilitate statistical analysis, we standardized the 14 nonoutline features with a mean of 0 and variance of 1, yielding 14 nonoutline variables. The outline features were obtained using the CD-FD method described in the CaF modeling subsection. We retained the Fourier coefficients that can reconstruct the contour information with a root mean square error (RMSE) ratio below 5%.<sup>3</sup> The number of Fourier coefficients retained for the left eyebrow, left eye, nose, mouth, and overall face were 5, 5, 12, 9, and 9, respectively. A higher number of Fourier coefficients are needed to capture features with subtle but important differences, such as the nostril area on the nose contour, which typically has substantial variation in the details that convey different perceptual information. A total of 40 Fourier descriptors (i.e., 80 outline features) were retained for the five facial parts. We standardized the Fourier descriptors with a mean of 0 and variance of 1 to facilitate statistical analysis, resulting in 80 outline variables. Overall, we collected 48,013 valid face-based ratings and 101,630 valid contour-based ratings on the 15 dimensions. On average, each face image received 15 valid ratings on each perception dimension, and each contour image received 31 valid ratings on each perception dimension.

<sup>2</sup> An assessment of the potential selection bias reveals that there are no significant differences in the choices of most perceptual dimensions among people with different age, gender, height, or education level. It seems, however, that women are more likely to choose confidence when evaluating male contour images, and people with higher education level are more likely to choose aggressiveness. In addition, people who are taller are less likely to choose trustworthiness.

<sup>3</sup> We calculated the RMSE by comparing the reconstructed contour using Fourier coefficients with the original contour. We calculated the RMSE ratio by dividing the RMSE of reconstructed contour using Fourier coefficients by the RMSE of reconstructed contour using the zero-frequency component.

**Analysis and Results**

*The role of contour information in perceptions of faces.* We examine the role of facial contour information in the perceptions of faces by comparing the model fit and prediction performance of the CaF method with that of two existing face perception models; namely, the PS model (Vernon et al. 2014) and the EF model (Xiao and Ding 2014). We begin the analysis by describing the specifications of the three models.

As discussed previously, the PS model extracts PS measures from the coordinates of salient landmarks on human faces. Following the state-of-the-art work in Vernon et al. (2014), we manually labeled 179 facial landmarks on the 216 face images and derived 47 PS measures based on the placements of the landmarks, 37 of which were related to the five key facial parts in CaF representation and the remaining 10 were related to other facial regions such as the ear, cheekbone, and cheek. We compared the CaF model with a PS model using the 37 PS measures as independent variables (see Web Appendix C for a list of the variables used in the PS model). We also compared the CaF model with a PS model using all 47 PS measures. The conclusions were not qualitatively different (see Web Appendix C).

The EF model extracts facial texture features through a function of eigenvectors derived from the covariance matrix of training face images (for details, see Turk and Pentland [1991]). Xiao and Ding (2014) applied this method in a marketing context to model the preferences of faces in print ads. Following their procedure, we first normalized the 216 face images and removed the noise related to the angle and size of the face. Then we obtained the eigenvectors and eigenvalues using the principle component analysis method. We retained 57 eigenvectors that explained 95% of the variance among the face images. Finally, we calculated the loading vector (57 × 1) for each face. These EF loadings were used as independent variables in the EF model.

For all three models, the mean perception ratings of each face image on the 15 dimensions were used as dependent variables, denoted as  $Y^j, j = 1, \dots, 15$ . Following the face perception modeling literature (Vernon et al. 2014; Xiao and Ding 2014), we applied a linear regression model to characterize the relationships between facial information and face perceptions. Letting a  $37 \times 1$  vector  $V = (v_1, \dots, v_{37})'$  denote the independent variables used in the PS model; a  $57 \times 1$  vector  $W = (w_1, \dots, w_{57})'$  denote the independent variables used in the EF model; a  $14 \times 1$  vector  $X = (x_1, \dots, x_{14})'$  denote the nonoutline variables; and a  $80 \times 1$  vector  $Z = (z_1, \dots, z_{80})'$  denote the outline variables, we specify the PS, EF, and CaF models as follows:

$$PS : Y^j = \alpha_{PS}^j + \beta_{PS}^j V + \epsilon_{PS}^j$$

$$EF : Y^j = \alpha_{EF}^j + \beta_{EF}^j W + \epsilon_{EF}^j$$

$$CaF : Y^j = \alpha_{CaF}^j + \beta_{CaF}^j X + \gamma_{CaF}^j Z + \epsilon_{CaF}^j$$

**Table 4.** Model Comparison: CaF, PS, and EF.

Dimension	R <sup>2</sup>			Adjusted R <sup>2</sup>			Leave-One-Out Prediction (MSE)		
	CaF	PS	EF	CaF	PS	EF	CaF	PS	EF
Age	<b>.72</b>	.48	.51	<b>.50</b>	.37	.33	<b>.29</b>	.33	.34
Aggressiveness	<b>.66</b>	.41	.35	<b>.40</b>	.29	.12	<b>.47</b>	.51	.62
Approachability	<b>.66</b>	.53	.41	.39	<b>.44</b>	.20	.40	<b>.38</b>	.57
Arousal	<b>.87</b>	.74	.71	<b>.76</b>	.68	.61	<b>.31</b>	.35	.50
Attractiveness	<b>.80</b>	.68	.69	<b>.65</b>	.61	.58	<b>.47</b>	.52	.59
Baby-facedness	<b>.60</b>	.23	.33	<b>.29</b>	.07	.09	<b>.28</b>	.33	.34
Confidence	<b>.82</b>	.67	.69	<b>.69</b>	.61	.58	<b>.30</b>	.34	.41
Dominance	<b>.59</b>	.30	.32	<b>.26</b>	.15	.07	<b>.32</b>	.34	.37
Fatness	<b>.62</b>	.39	.41	<b>.33</b>	.27	.19	.50	<b>.48</b>	.59
Health	<b>.71</b>	.52	.62	<b>.49</b>	.42	<b>.49</b>	<b>.23</b>	.26	.24
Intelligence	<b>.73</b>	.47	.61	<b>.53</b>	.37	.47	<b>.25</b>	.31	.26
Pleasantness	<b>.78</b>	.69	.62	.61	<b>.62</b>	.49	.41	<b>.40</b>	.56
Sexual dimorphism	<b>.91</b>	.81	.85	<b>.84</b>	.77	.79	<b>.61</b>	.86	.80
Smile	<b>.89</b>	.83	.74	<b>.80</b>	.79	.64	<b>.32</b>	<b>.32</b>	.66
Trustworthiness	<b>.74</b>	.63	.58	.54	<b>.55</b>	.42	.29	<b>.26</b>	.36

Notes: The best models are highlighted in bold.

We compared the three models on both goodness-of-fit and leave-one-out prediction performance.<sup>4</sup> The results appear in Table 4. The CaF model performs best overall in explaining the variances in face-based perception ratings (average R<sup>2</sup> = .74; average adjusted R<sup>2</sup> = .54) and prediction performance (average mean square error [MSE] = .36) when compared with the PS model (average R<sup>2</sup> = .56; average adjusted R<sup>2</sup> = .47; average MSE = .40) and the EF model (average R<sup>2</sup> = .56; average adjusted R<sup>2</sup> = .40; average MSE = .48). On most dimensions, the CaF model performs the best, or equally as well, in predicting face-based perception ratings and in explaining the variances in face-based perception ratings. These results suggest that the contour features can improve the explanatory and predictive power of face perception models across various dimensions, relative to the PS and EF features.

Compared with the PS model, the CaF model is able to increase the proportion of explained variance (after adjusting for the number of independent variables) on most dimensions such as baby-facedness (by 22%), intelligence (by 16%), age (by 13%), dominance (by 11%), aggressiveness (by 11%), arousal (by 8%), and confidence (by 8%). This suggests that the contours of the intervening regions among landmarks contain important information people need to make useful inferences on these dimensions. The performance of the PS model and CaF model are close on the smile, trustworthiness, pleasantness, attractiveness, and approachability dimensions, all of which are correlated with the perception of valence (Oosterhof and Todorov 2008). This suggests consistency between the contour features and PS measures in capturing the facial information that is important for valence perceptions along these

<sup>4</sup> We used Lasso (Tibshirani 1996) to alleviate potential overfitting problems.

dimensions. The EF model appears to perform the worst on most dimensions except health, which is largely affected by skin texture information (Henderson et al. 2016). This provides evidence that the structural cues as captured by the contour features play a more important role in most face perceptions than the texture cues captured by eigenvectors.

These results demonstrate that (1) the complete contours of facial parts contain important information people need for face perceptions, and (2) compared with the two state-of-the-art benchmark models, the CaF model is generally more effective in capturing the key facial information driving face perceptions on most dimensions. These findings provide strong evidence for the use of contour information in modeling face perceptions and justify the use of the CaF framework in this context.

*Comparing perceptions of contour images versus face images.* To examine the effectiveness of the CaF representation in preserving information for commonly assessed face perceptions, we compared human perceptions on the 15 dimensions after exposure to two types of face representations (facial contour images vs. face images). The perceptions based on face images were used as benchmarks to evaluate the reliability of contour-based perceptions. We evaluated the effectiveness of the CaF representation from two aspects: (1) whether and to what extent contour-based perceptions relate to perceptions based on original face images and (2) how contour-based perceptions differ from face-based perceptions across the different perceptual dimensions.

For each face image and the corresponding contour image, we calculated the mean of perception ratings for each dimension. Figure 2 displays the distributions of the mean perception ratings for the 216 contour images and 216 face images on all 15 dimensions. To evaluate the relationship between contour-based and face-based perceptions and the reliability of contour-based perceptions, we calculated the Pearson correlation and the intra-class correlation (ICC, a measure widely used in reliability analysis; for details of ICC criteria, see Hallgren [2012]), of mean perception ratings for all 15 dimensions. The average correlation between contour-based and face-based mean perception ratings for all 15 dimensions was .41, and the average ICC value was .51, indicating that people can make reasonably reliable perceptions using contour images.

A closer examination reveals that the reliability of contour-based inferences varies across perceptual dimensions: it is much easier for humans to make reliable inferences from face contours about some dimensions (e.g., smile, approachability) than others (e.g., gender, age).

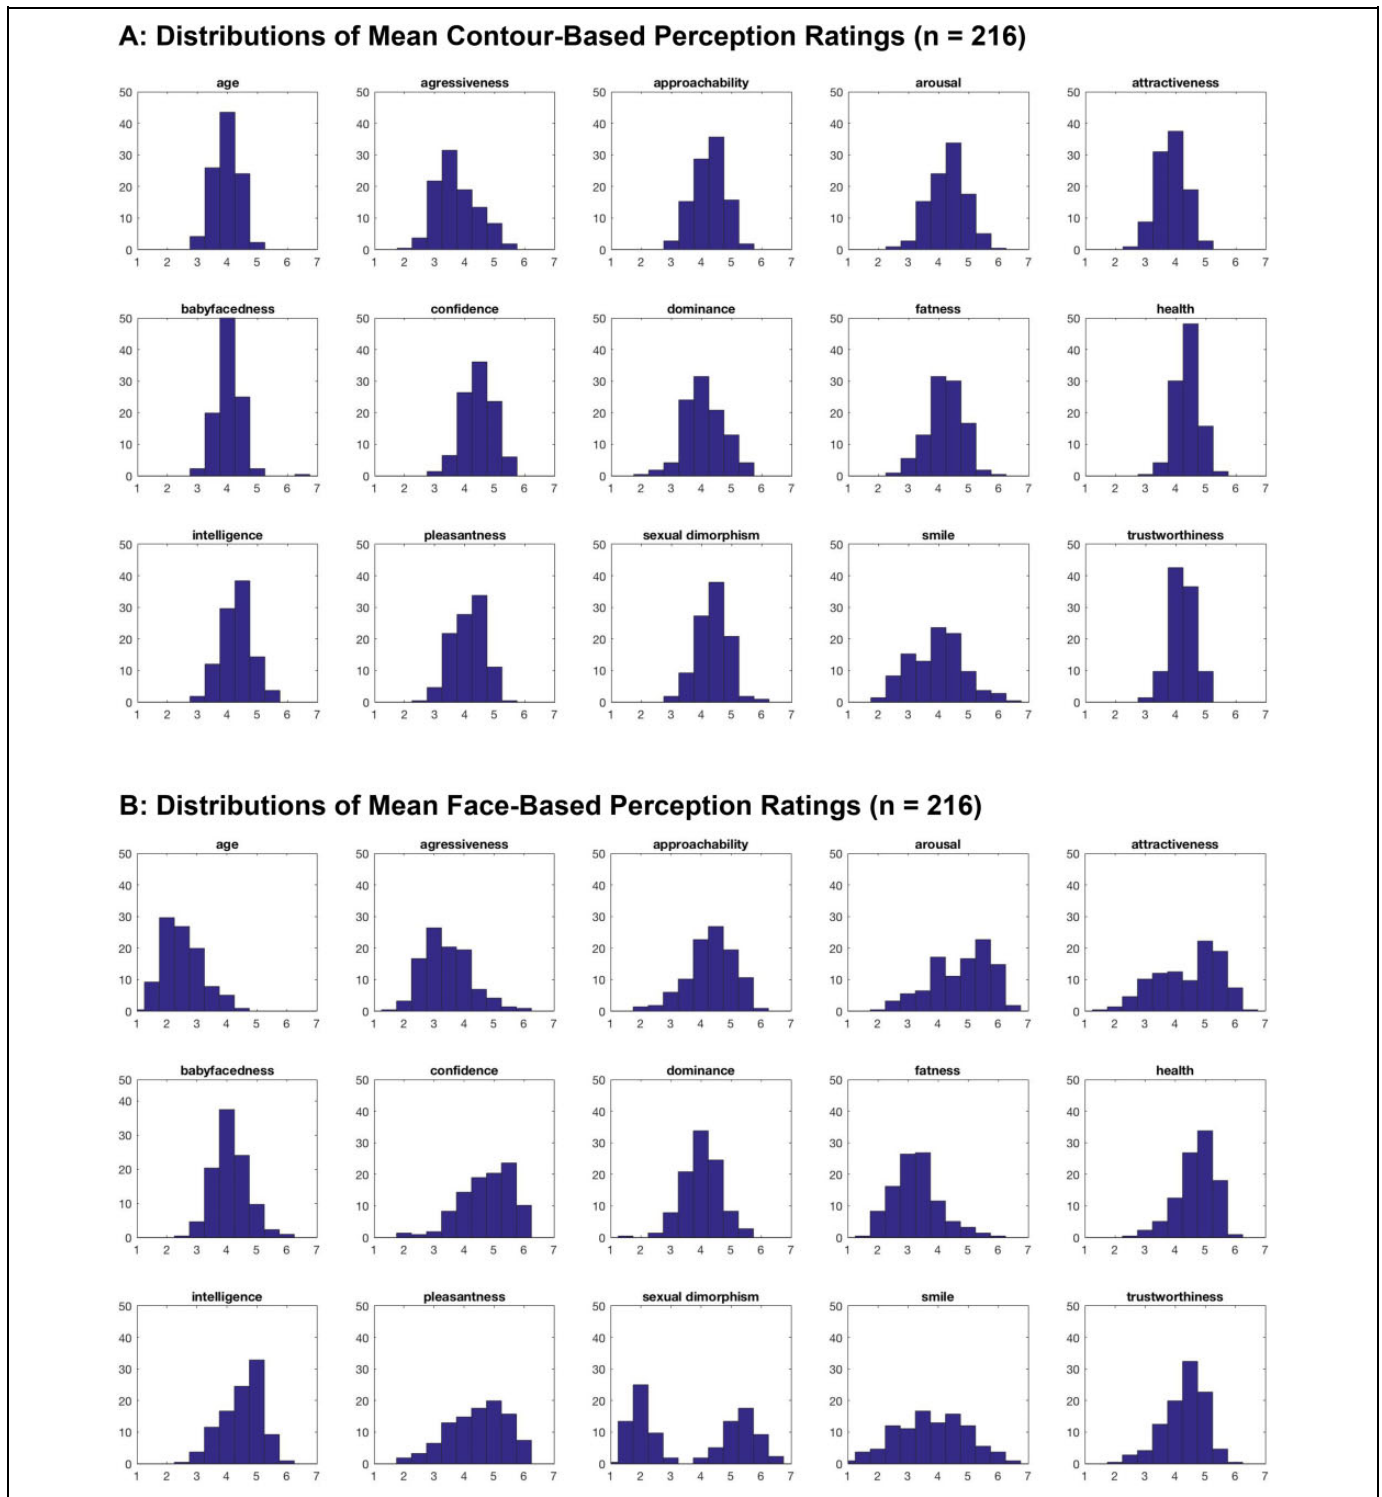
First, people can infer a positive emotional state and correlated dimensions such as pleasantness, attractiveness, approachability, arousal, and confidence (referred to as “valence” dimensions) with considerably high reliability based on contours. Specifically, the contour-based face perceptual inferences were quite reliable for seven dimensions: smile ( $r_{\text{smile}} = .81, p < .01$ ;  $ICC_{\text{smile}} = .88$ ), pleasantness ( $r_{\text{pleasantness}} = .62, p < .01$ ;  $ICC_{\text{pleasantness}} = .70$ ), attractiveness ( $r_{\text{attractiveness}} = .62, p < .01$ ;  $ICC_{\text{attractiveness}} = .63$ ), approachability ( $r_{\text{approachability}} = .51, p < .01$ ;  $ICC_{\text{approachability}} = .64$ ), arousal ( $r_{\text{arousal}} = .55, p < .01$ ;

$ICC_{\text{arousal}} = .66$ ), and confidence ( $r_{\text{confidence}} = .49, p < .01$ ;  $ICC_{\text{confidence}} = .60$ ). These findings are consistent with prior research indicating that valence perceptions are highly correlated with each other and are driven by expression signaling such as a smile (Oosterhof and Todorov 2008; Todorov et al. 2008; Vernon et al. 2014). A further analysis of the discrepancies between the mean contour-based and face-based perceptions on the valence dimensions seems to indicate that the perception of valence is slightly lower when using contour images (e.g.,  $M_{\text{attractiveness}} = 3.87$  vs. 4.43,  $p < .01$ ;  $M_{\text{pleasantness}} = 4.14$  vs. 4.46,  $p < .01$ ;  $M_{\text{arousal}} = 4.31$  vs. 4.76,  $p < .01$ ;  $M_{\text{confidence}} = 4.46$  vs. 4.76,  $p < .01$ ). Interestingly, it seems that faces are perceived to be more homogeneous on these dimensions when using contour images than when using face images. For example, the contour-based perceptions of smile, approachability, attractiveness, pleasantness, arousal, and confidence seem to vary less than face-based perceptions in our data set ( $SD_{\text{smile}} = .89$  vs. 1.14,  $p < .01$ ;  $SD_{\text{approachability}} = .53$  vs. .77,  $p < .01$ ;  $SD_{\text{attractiveness}} = .47$  vs. 1.06,  $p < .01$ ;  $SD_{\text{pleasantness}} = .53$  vs. .94,  $p < .01$ ;  $SD_{\text{arousal}} = .59$  vs. .98,  $p < .01$ ;  $SD_{\text{confidence}} = .51$  vs. .86,  $p < .01$ ), indicating that people can make more differentiating judgments on these dimensions if given more facial information than the facial contours.

Compared with face-based perceptions, contour-based perceptions on the fatness dimension were fairly reliable ( $r_{\text{fatness}} = .48, p < .01$ ;  $ICC_{\text{fatness}} = .64$ ), which is consistent with research suggesting that facial contour information is predictive of fatness (Mayer et al. 2017). In general, it seems that the faces in our database tended to be perceived as fatter when using contour images than when using face images ( $M_{\text{fatness}} = 4.21$  vs. 3.27,  $p < .01$ ;  $SD_{\text{fatness}} = .59$  vs. .77,  $p < .01$ ).

Second, it appears to be much harder or even impossible for humans to make reliable inferences on the two demographic dimensions of gender and age and their correlated dimensions. For example, Figure 2, Panels A and B, show that participants can perceive gender reasonably well from face images (Figure 2, Panel B, sexual\_dimorphism) but fail to identify gender-specific facial features using contours (Figure 2, Panel A, sexual\_dimorphism). This indicates that masculine and feminine facial cues are less distinguishable when using contour images ( $r_{\text{sexual_dimorphism}} = .26, p < .01$ ;  $ICC_{\text{sexual_dimorphism}} = .24$ ). This finding is consistent with prior studies showing that texture cues such as skin color and contrast are important for gender perception (Bruce and Langton 1994). Contour-based perceptions for the second demographic dimension, age, were also unreliable ( $r_{\text{age}} = .31, p < .01$ ;  $ICC_{\text{age}} = .44$ ). A comparison between Panel A and Panel B on the age dimension shows that perceptions of age based on contour images deviate significantly from perceptions based on the original face images ( $M_{\text{age}} = 3.98$  vs. 2.52,  $p < .01$ ;  $SD_{\text{age}} = .42$  vs. .67,  $p < .01$ ).<sup>5</sup> Moreover, no significant correlation exists between contour-based and face-

<sup>5</sup> An alternative interpretation of the results is that inferred ages based on contours were older because the overwhelming majority of the 216 photos were of younger people. It is possible that inferences of contour-based ages would be younger than face-based ages if the face images depicted older people.




**Figure 2.** Distributions of mean contour-based versus face-based perception ratings, by dimension.

Notes: In all graphs, the x-axis is the mean perception rating, and the y-axis is the percentage of total images. For sexual dimorphism, 1 = “very feminine,” and 7 = “very masculine.”

based perceptions on another dimension related to age, babyfacedness ( $r_{\text{baby-facedness}} = .05, p > .10$ ;  $\text{ICC}_{\text{baby-facedness}} = .09$ ), indicating that participants cannot effectively infer babyfacedness from contour images. Our analyses seem to support

prior findings that perceptions of age and age-related dimensions are largely affected by texture cues such as skin contrast and skin pigmentation (Porcheron, Mauger, and Russell 2013), which are absent from contour images.



**Figure 3.** Mean relative difference in perception ratings across contour clusters.

Contour-based inferences on the three dimensions that are correlated with gender and age, dominance, aggressiveness, and health (Henderson et al. 2016; Oosterhof and Todorov 2008) were also not reliable ( $r_{\text{dominance}} = .14, p < .05$ ;  $\text{ICC}_{\text{dominance}} = .25$ ;  $r_{\text{aggressiveness}} = .25, p < .01$ ;  $\text{ICC}_{\text{aggressiveness}} = .40$ ;  $r_{\text{health}} = .35, p < .01$ ;  $\text{ICC}_{\text{health}} = .47$ ). This supports prior findings showing that perceptions of dominance, aggressiveness, and health are sensitive to signals of physical strength related to sexual dimorphism (gender) and maturity (age) (Henderson et al. 2016; Oosterhof and Todorov 2008) and provides additional evidence that physical strength signals are contained primarily in the texture information rather than in the contour information. Contour-based perceptions on the two dimensions that have proven to be related to both valence dimensions and gender/age, intelligence, and trustworthiness (Oosterhof and Todorov 2008) are also shown to be less reliable ( $r_{\text{intelligence}} = .35, p < .01$ ;  $\text{ICC}_{\text{intelligence}} = .50$ ;  $r_{\text{trustworthiness}} = .34, p < .01$ ;  $\text{ICC}_{\text{trustworthiness}} = .45$ ). This indicates that both facial contour and texture information are needed for making reliable perceptions on the two dimensions.

Overall, the evidence suggests that humans can make effective inferences from facial contours on perceptual dimensions other than sexual dimorphism (gender), age, and the correlated perceptual dimensions. This points to the great potential of the CaF representation in masking age and gender-related facial cues to protect privacy and eliminate discrimination while preserving the information necessary for people to make relevant inferences about other dimensions such as smile, pleasantness, attractiveness, approachability, arousal, and confidence.

*Heterogeneity among faces and the reliability of contour-based perception.* Figure 2 indicates that another factor that potentially

accounts for the variations in the reliability of contour-based perception is the heterogeneity among faces. We explored this through a cluster analysis using the relative differences between the means of contour-based and face-based perception ratings on all 15 dimensions. We used the K-means clustering and tested for  $K = 1, 2, 3, 4, 5, 6$ . The four-cluster result is the best based on the gap statistic (Tibshirani, Walther, and Hastie 2001). The four clusters include 69, 36, 77, and 34 images, respectively.

Figure 3 shows comparisons among the four clusters on the reliability of contour-based perceptions, along with the average face-based ratings on all 15 dimensions for each cluster. The results show that there are substantial differences in the reliability of contour-based perceptions among different face groups. In Clusters 1 and 2, estimates of (masculine) sexual dimorphism are more pronounced, but most social trait perceptions (e.g., arousal, attractiveness) tend to be underestimated. By contrast, in Cluster 4, most social trait perceptions (e.g., arousal, attractiveness) tend to be overestimated, but (masculine) sexual dimorphism tends to be less pronounced. Because most faces in Clusters 1 and 2 are female, whereas most faces in Clusters 3 and 4 are male, the overestimation of masculinity for female faces and the underestimation of masculinity for male faces may be attributable to the fact that facial features related to sexual dimorphism are masked in contour images. It is worth noting that the impact of facial contours on perceptions varies significantly even within the same gender. For example, faces in Cluster 4 tend to be overestimated on approachability, whereas contour-based ratings for Cluster 3 are quite close to face-based ratings on the same dimension.

**Table 5.** Contribution of Nonoutline and Outline Variables to Contour-Based Perceptions.

Dimension	R <sup>2</sup>			Adjusted R <sup>2</sup>			F-Test (p-Value)	
	M <sub>0</sub>	M <sub>1</sub>	M <sub>2</sub>	M <sub>0</sub>	M <sub>1</sub>	M <sub>2</sub>	M <sub>2</sub> vs. M <sub>0</sub>	M <sub>2</sub> vs. M <sub>1</sub>
Age	.18	.52	.60	.12	.24	.30	<b>.01</b>	.05
Aggressiveness	.39	.54	.66	.35	.27	.39	.23	<b>.00</b>
Approachability	.32	.66	.74	.27	.46	.54	<b>.00</b>	<b>.00</b>
Arousal	.47	.77	.83	.43	.63	.70	<b>.00</b>	<b>.00</b>
Attractiveness	.33	.68	.74	.28	.50	.54	<b>.00</b>	<b>.04</b>
Baby-facedness	.09	.40	.44	.03	.04	.01	.59	.76
Confidence	.41	.68	.75	.37	.48	.56	<b>.00</b>	<b>.00</b>
Dominance	.33	.53	.62	.28	.26	.32	.24	<b>.03</b>
Fatness	.32	.81	.84	.28	.69	.72	<b>.00</b>	<b>.02</b>
Health	.23	.57	.60	.18	.32	.29	.05	.86
Intelligence	.32	.57	.66	.27	.32	.39	<b>.03</b>	<b>.01</b>
Pleasantness	.26	.69	.73	.21	.50	.51	<b>.00</b>	.22
Sexual dimorphism	.32	.57	.69	.27	.32	.44	<b>.00</b>	<b>.00</b>
Smile	.41	.83	.86	.37	.73	.76	<b>.00</b>	<b>.03</b>
Trustworthiness	.20	.44	.51	.15	.11	.14	.54	.21

Notes: M<sub>0</sub> is the nonoutline model with the 14 nonoutline variables. M<sub>1</sub> is the outline model with the 80 outline variables. M<sub>2</sub> is the full CaF model including the 14 nonoutline variables and 80 outline variables. Significant F-test results ( $p < .05$ ) are highlighted in boldface.

*Drivers of contour-based perceptions.* To investigate the information in contour images that drives contour-based perceptions, we quantified the relationships between contour facial features and contour-based perception ratings. Following the prior studies on face perceptions (Todorov 2017), we analyzed the contributions of the two types of contour features (nonoutline and outline) on contour-based face perceptions.

To discern the effect of the two types of contour features, we compared the CaF model (M<sub>2</sub>) with two variations: a nonoutline model (M<sub>0</sub>) with 14 nonoutline variables as independent variables and an outline model (M<sub>1</sub>) with 80 outline variables as independent variables. For all three models, the mean perception ratings of each contour image on the 15 dimensions were used as dependent variables, denoted as  $U^j$ ,  $j = 1, \dots, 15$ . We specify M<sub>0</sub>, M<sub>1</sub>, and M<sub>2</sub> as follows:

$$M_0 : U^j = \alpha_0^j + \beta_0^j X + \varepsilon_0^j$$

$$M_1 : U^j = \alpha_1^j + \gamma_1^j Z + \varepsilon_1^j$$

$$M_2 : U^j = \alpha_2^j + \beta_2^j X + \gamma_2^j Z + \varepsilon_2^j$$

We performed standard linear regression analysis. The results appear in Table 5. The results confirmed our intuition that both outline and nonoutline variables contribute to contour-based face perceptions. Overall, the full model (M<sub>2</sub>) explains a substantial proportion of the variances in perceptions (average  $R^2 = .68$ , average adjusted  $R^2 = .44$ ) compared with the outline model (M<sub>1</sub>) (average  $R^2 = .62$ , average adjusted  $R^2 = .39$ ) and the nonoutline model (M<sub>0</sub>) (average  $R^2 = .31$ ,

average adjusted  $R^2 = .26$ ). A comparison between the two variation models and the full model at the dimension level shows that both the outline and nonoutline variables significantly affect the contour-based perceptions on most of the dimensions. These results suggest that researchers should use the full CaF model with both outline and nonoutline variables when modeling face perceptions based on CaF representation.

As we expected, the effects of nonoutline and outline variables vary greatly across different perceptual dimensions. First, for the dimensions on which it is easy for humans to make reliable inferences from the contour images—such as smile, pleasantness, approachability, attractiveness, arousal, confidence, and fatness—human perceptions seem to be driven mainly by the outline variables. The full CaF model is able to capture substantial proportions of the variances on these dimensions. Second, for the dimensions on which it is difficult for humans to make reliable inferences from the contour images, the main drivers of people's perceptions vary. For example, the results indicate that outline information plays a dominant role in people's contour-based perceptions of age, sexual dimorphism, and health, while nonoutline information plays a dominant role in people's perceptions of dominance and aggressiveness. Neither outline nor nonoutline variables could effectively explain perceptions of baby-facedness, which is consistent with our findings that there is no significant correlation between contour-based and face-based perceptions on the baby-facedness dimension.

Here, we make two additional notes regarding the results. First, ideally, latent class analysis (or other methods that account for heterogeneity) would be used to further examine the heterogeneity among faces. Such an analysis is not meaningful with our limited face-contour data and large number of variables; however, it should be explored in the future when more data become available. Second, we assessed the out-of-sample prediction performance of our models and found that M<sub>2</sub> was only marginally better than M<sub>0</sub> in a few cases and equivalent in others. Given the generally superior fit of M<sub>2</sub>, we hypothesize that this could be due to overfitting and substantial facial heterogeneity. Creating a model that accounts for heterogeneity when a much larger set of face contours and ratings becomes available might help solve this problem and enable exploration of the heterogeneity in contour-based perceptions of different faces.

To further explore the contributions of outline variables associated with individual facial parts, we compared five variations of the CaF model, each using the outline features of a facial part in addition to the 14 nonoutline features as independent variables (see Table 6). The results show that the contributions of outline variables to perceptions of different dimensions vary greatly across facial parts.<sup>6</sup> In general, the outlines of the eyebrows and eyes do not significantly affect most perceptions. The outline of the mouth, however, explains as

<sup>6</sup> We also analyzed the contributions of nonoutline variables to perceptions of different dimensions, and the results show that they also vary across facial parts.

**Table 6.** Effects of Outline Variables of Different Facial Parts on Contour-Based Perceptions.

Dimension	M <sub>0</sub>			M <sub>0</sub> + Eyebrow Contour			M <sub>0</sub> + Eye Contour			M <sub>0</sub> + Nose Contour			M <sub>0</sub> + Mouth Contour			M <sub>0</sub> + Face Contour		
	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )	R <sup>2</sup>	Adj. R <sup>2</sup>	F-Test (vs. M <sub>0</sub> )
Age	.18	.12	.21	.12	.55	.22	.12	.38	.38	.25	<b>.00</b>	.33	.21	<b>.00</b>	.37	.26	<b>.00</b>	
Aggressiveness	.39	.35	.43	.35	.32	.44	.37	.12	.46	.35	.52	.49	.40	<b>.01</b>	.43	.34	.72	
Approachability	.32	.27	.35	.27	.50	.36	.28	.23	.41	.29	.26	.61	.54	<b>.00</b>	.39	.29	.23	
Arousal	.47	.43	.49	.42	.77	.49	.42	.83	.56	.46	.09	.69	.63	<b>.00</b>	.56	.48	<b>.01</b>	
Attractiveness	.33	.28	.36	.27	.59	.35	.27	.68	.47	.35	<b>.01</b>	.54	.45	<b>.00</b>	.53	.45	<b>.00</b>	
Baby-facedness	.09	.03	.12	.01	.89	.12	.01	.81	.18	.00	.78	.20	.06	.18	.20	.06	.13	
Confidence	.41	.37	.43	.35	.90	.43	.36	.87	.49	.39	.25	.60	.53	<b>.00</b>	.49	.40	.08	
Dominance	.33	.28	.38	.30	.11	.37	.29	.25	.45	.33	.05	.40	.29	.28	.39	.29	.39	
Fatness	.32	.28	.38	.30	.08	.36	.28	.37	.52	.42	<b>.00</b>	.48	.39	<b>.00</b>	.78	.74	<b>.00</b>	
Health	.23	.18	.25	.15	.96	.26	.17	.66	.33	.19	.38	.42	.32	<b>.00</b>	.36	.25	<b>.01</b>	
Intelligence	.32	.27	.35	.27	.56	.37	.29	.13	.46	.34	<b>.01</b>	.44	.34	<b>.01</b>	.43	.34	<b>.01</b>	
Pleasantness	.26	.21	.29	.20	.69	.32	.23	.14	.35	.21	.41	.61	.54	<b>.00</b>	.35	.24	.15	
Sexual dimorphism	.32	.27	.36	.28	.43	.35	.27	.50	.48	.37	<b>.00</b>	.49	.40	<b>.00</b>	.43	.33	<b>.02</b>	
Smile	.41	.37	.44	.37	.52	.45	.38	.21	.52	.42	<b>.04</b>	.79	.75	<b>.00</b>	.49	.40	.08	
Trustworthiness	.20	.15	.23	.13	.68	.25	.15	.33	.28	.13	.72	.31	.19	<b>.05</b>	.27	.14	.56	

Notes: Significant F-test results ( $p < .05$ ) are highlighted in bold.

much as 79% (adjusted  $R^2 = .75$ ) of the variance in contour-based perceptions for smile, indicating that the mouth outline drives the perceptions of smiles in contour images. Similarly, the mouth outline drives perceptions of pleasantness, confidence, arousal, and approachability, which supports findings from prior studies that these perceptions tend to be loaded on the same dimension and are sensitive to emotional cues (Todorov et al. 2008). Including face outline information can explain as much as 78% (adjusted  $R^2 = .74$ ) of the variance in contour-based perceptions of fatness, indicating that the face outline drives perceptions of fatness in contour images. Both the face and mouth outlines drive perceptions of attractiveness. Compared with outlines of the mouth and/or face, the nose outline plays a relatively less important role on most dimensions.

*Self-revealed decision rules in contour-based perceptions.* To understand the decision rules that people use to make inferences from contour images, we asked another group of participants to explicitly describe the facial cues they used to make inferences based on contour images. We collected responses from a total of 76 participants (19 men, 57 women), each of whom was asked to rate at least five (randomly assigned) contours on five dimensions they had chosen from the 15 dimensions. Two research assistants coded the responses independently and categorized the self-revealed decision rules into different types of contour features. Tables 7 and 8 present the results and sample decision rules. In line with our previous analysis, we analyzed the self-revealed decision rules related to the types of contour features, as well as the key facial parts.

Overall, the results from Table 7 demonstrate great consistency with the model-based findings. The self-revealed decision rules show that people rely on both outline and nonoutline

information in their judgments of facial contours. In addition, compared with the nonoutline features, people generally use more outline features in their judgments. On average, 85% of participants mentioned the use of outline features in their judgments across each of the 15 dimensions, while 54% of participants mentioned the use of nonoutline features.

Consistent with our model-based findings, the outline information of the mouth (51%<sup>7</sup>) and the face (42%) are generally used more often in people's perceptions than the outline information of other facial parts such as the eyebrows (16%), eyes (15%), and nose (7%). At the dimension level (see Tables 7 and 8), the outline information of the mouth is used most often among all facial parts in perceptions of smile (100%), pleasantness (90%), confidence (68%), approachability (71%), and arousal (55%), which is consistent with the model-based findings. The outline information of the face is used most often in perceptions of fatness (90%), health (64%), and age (54%), which supports our findings that face outline variables have significant effects on contour-based perceptions of these dimensions. These findings provide strong evidence that the proposed CaF model is effective in capturing the information people use in contour-based perceptions.

In contrast, the nonoutline information of mouth (3%) and face (5%) is used less often in people's judgments than the nonoutline information related to other facial parts, such as eyebrow (34%), eye (27%), and nose (15%). For example, people used nonoutline information, including the size, angle, and spatial placement, of the eyebrows and eyes in their judgments of arousal (59% and 62%, respectively), approachability

<sup>7</sup> On average, 51% of people stated that they used the contour of mouth in their judgments of the contour images.

**Table 7.** Self-Revealed Decision Rules for Contour-Based Perceptions.

Dimension	Percentage of Participants Mentioning Outline Features in Self-Revealed Decision Rules (%)					Overall (at least one outline feature is mentioned)	Percentage of Participants Mentioning Nonoutline Features in Self-Revealed Decision Rules (%)					Overall (at least one nonoutline feature is mentioned)
	Eyebrow	Eye	Nose	Mouth	Face		Eyebrow	Eye	Nose	Mouth	Face	
Age	9	6	6	28	<b>54</b>	78	<b>22</b>	15	6	4	11	43
Aggressiveness	25	8	0	<b>58</b>	33	92	<b>50</b>	33	17	0	0	58
Approachability	27	18	0	<b>71</b>	18	90	<b>55</b>	31	10	6	2	67
Arousal	17	21	7	<b>55</b>	31	76	59	<b>62</b>	3	3	0	86
Attractiveness	3	9	31	43	<b>49</b>	86	23	<b>43</b>	29	3	6	69
Baby-facedness	13	31	13	19	<b>56</b>	88	6	6	13	0	6	25
Confidence	14	14	7	<b>68</b>	4	82	<b>43</b>	25	7	7	0	71
Dominance	15	31	0	<b>62</b>	<b>62</b>	92	<b>46</b>	31	15	0	0	54
Fatness	0	3	8	13	<b>90</b>	90	0	18	15	8	<b>36</b>	44
Health	0	9	0	36	<b>64</b>	91	9	9	9	0	0	18
Intelligence	18	0	18	27	18	55	27	<b>55</b>	18	9	0	73
Pleasantness	33	24	5	<b>90</b>	29	100	<b>43</b>	24	5	0	0	48
Sexual dimorphism	14	5	9	36	<b>55</b>	77	<b>68</b>	23	55	5	9	77
Smile	27	32	0	<b>100</b>	14	100	18	14	5	5	0	27
Trustworthiness	22	17	0	56	56	83	<b>33</b>	22	17	0	0	56

Notes: The percentages shown in the table are calculated by (Number of participants who mentioned the feature for the selected dimension)/(Number of participants who selected the dimension). Two research assistants were hired to code the participants' responses independently. The Cohen's kappa of two independent coders is  $>.90$ . The facial part that was mentioned most for each perceptual dimension is highlighted in bold.

**Table 8.** Examples of Self-Revealed Decision Rules for Contour-Based Perceptions.

Dimension	Examples of Self-Revealed Decision Rules
Age	"The person is young if the face is rounder, and the person is middle-aged if the overall face contour has sharp corners; If the face contour is neither round nor has sharp corners, the person is old; The person is old if the corners of mouth are curving down."
Aggressiveness	"The person is more aggressive if the face contour has more sharp corners; The person is more aggressive if the corners of mouth are curving down."
Approachability	"The person is more approachable if the mouth has a smiling curve and the corners of mouth are curving up, and if the person has flat eyebrows with an elongated downward curve."
Arousal	"The level of arousal is higher if the mouth has a smiling curve, and if the contour has raised eyebrows."
Attractiveness	"The person is more attractive if the face is rounder and has fewer sharp corners, and if the contour has smaller nose and bigger eyes."
Baby-facedness	"The person is less mature if the contours of face, eyebrows, eyes and mouth have fewer sharp corners."
Confidence	"The person is more confident if the mouth has a smiling curve and the eyebrows arch down."
Dominance	"The person is more dominant if the face has more sharp corners and higher eyebrow ridges."
Fatness	"The person is fatter if the face and chin are rounder; the person is thin if the chin is pointy."
Health	"The person is healthy if the face is round, but not too round."
Intelligence	"The person is more intelligent if the nostril ridges are smaller, and if the contour has smaller nose, and thinner mouth."
Pleasantness	"The person is more pleasant if the mouth has a smiling curve and the eyebrow-eye distance is larger; The person is less pleasant if the mouth-chin distance is larger and the corners of mouth are turning down."
Sexual dimorphism	"The person is more masculine if the face is more squared and the chin has sharp corners, and if the face has thicker and longer eyebrows and larger nose."
Smile	"The person is smiling if the corners of mouth are upward, and the eyebrows and eyes bent down."
Trustworthiness	"The person is more trustworthy if the mouth has a smiling curve, and if the face is broader and has raised eyebrows."

(55% and 31%, respectively), aggressiveness (50% and 33%, respectively), and dominance (46% and 31%, respectively). The nonoutline information of the eyebrows (68%) and nose (55%) seems to be used most often in people's contour-based perceptions on sexual dimorphism.

In summary, the results in Study 1 demonstrate that the CaF framework is effective as both a stimulus and modeling method, and humans can make meaningful inferences from contour images that are quite consistent with inferences made from the corresponding face images. In addition, our findings



show that contours can significantly reduce the accuracy of people's perceptions of age and gender, two perceptual dimensions that contours are intended to mask to help prevent discrimination.

## Study 2: The Effectiveness of CaF for Protecting Privacy

A face's level of anonymity typically equates to the extent to which face recognition is prevented (Newton, Sweeney, and Malin 2005). In Study 2, we aimed to assess the CaF representation's ability to mask private information (e.g., identity, age, gender) from human recognition. We focused on human recognition for two reasons. First, marketing contexts in which the CaF framework is most likely to be used (Table 2) typically require privacy protection from other humans, not sophisticated algorithms. Second, it is difficult for algorithms to unmask identities using face contours derived from different photos of the same person, because the contours may differ greatly due to the face angle, head orientation, emotional state (e.g., smiling changes the contours of mouth and face) and the apparent physical state (e.g., the eyes and mouth may be open to different degrees).

### Study Design

A typical face recognition task requires participants to identify a target face that matches a stimulus face from a designated database. To test how well contour images are able to anonymize one's identity and demographics to prevent human recognition/perception, we asked participants to view contour images and identify the corresponding faces from a database. To mimic the use of contour images in real-life applications, we first constructed a database of 260 faces (184 men, 76 women) from the faculty pages at a university. We randomly chose 10 individuals from the database as the target people and searched publicly available sources for different face photos of the target people. We then manually generated the facial contours using these face photos. If a different face photo was not publicly available, we randomly chose another individual from the database as a target individual.

A total of 138 students (47 men, 91 women) from a different university participated in the study and received a fixed payment for participating. The participants were presented with 10 contour images as well as the 260 faces in the database and were asked to identify the faces corresponding to each of the 10 contour images. Each participant was permitted to make up to three guesses per contour image with a decreasing order of confidence score (e.g., "How likely do you think the guessed face and the person in the contour image are the same person?").

For each contour image, we offered a cash prize (about US\$75) to be split evenly among all participants who correctly identified the target face in the contour image as their top choice. If no one correctly identified the target face in their top choice, the prize would be split evenly by all participants who

correctly identified the target face in their second choice, and so on. If no one correctly identified the target face in any of their three guesses, the cash prize would be split evenly by all participants.

### Analysis and Results

The participants submitted a total of 4,089 guesses about the identities of the people represented by the contour images, of which 1,380 were top choices. We present our analysis of the participants' top choices here; conclusions from the analyses of the second and third choices were not qualitatively different. The results appear in Table 9.

*Effect of CaF on anonymizing identity.* Overall, the results show that it is very difficult for humans to discern a person's identity based on contour images. As shown in Table 9, 37 out of the 1,380 guesses were correct, making the overall recognition rate 2.68%.

A further analysis of the face identification results reveals four key findings. First, the ease of recognition seems to vary across different faces, which is consistent with the findings from Study 1. For example, the recognition rates of the 10 contour images ranged from 0% to 10.87%; the overwhelming majority of contour images were not easily identified, with the exception of one (Contour Image 7) whose target face was also the most often guessed face (recognition rate of 10.87%). However, this rate is not significantly larger than the guess rate of the second-most-often-guessed face (5.80%).<sup>8</sup> The verbal feedback collected from participants suggests that the high recognition rate was due to the unique mouth and face contour of the target face. Second, in general, a single contour image can be matched with multiple faces. For example, among the 260 faces in the database, the number of unique guesses for a given contour image ranged from 58 to 77. In fact, out of the nine target faces that were correctly identified, seven were not among the five most-often-guessed faces for the corresponding contours. This indicates that contour images effectively conceal identifiable information. Third, statistical tests show that for the majority of faces, the correct recognition rates were not statistically different from the random guesses among the unique faces guessed by all participants. While the correct identification rate for Contour Image 5 (5.07%) was higher than random guess rate among unique faces guessed by all participants ( $p < .01$ ), it is not significantly different from the guess rates of the other four most-often-guessed faces. In addition, the correct face was only the third most-often-guessed face. Similarly, for Contour Image 7, although the correct identification rate was higher than random guesses among unique faces guessed by all participants ( $p < .01$ ), there were no significant differences among the guess rates of the top three most-often-guessed faces. This indicates that although people

<sup>8</sup>  $p = .10$  for hypothesis testing  $H_0: p_1 \leq p_2$  versus  $H_1: p_1 > p_2$  using proportion test, where  $p_1$  and  $p_2$  are proportions of the first- and second-most-often-guessed faces, respectively.

**Table 9.** Results of Human Identification Using Contour Images.

Contour Image	Correct Identification Rate (%)	$p$ -Value <sup>a</sup>	Five Faces Guessed Most Often (%)					Total Unique Faces Guessed by All Participants	Age Ratio	Gender Ratio
			1st	2nd	3rd	4th	5th			
1	.72	.90	12.32	7.25	6.52	6.52	4.35	60	.87	.40
2	1.45	.58	8.70	4.35	4.35	4.35	3.62	71	.97	.97
3	1.45	.60	7.25	5.07	3.62	3.62	3.62	69	1.20	.10
4	.72	.89	12.32	7.25	6.52	5.80	3.62	63	1.36	.73
5	5.07	.01	7.97	6.52	5.07 <sup>b</sup>	4.35	4.35	62	.68	.15
6	2.90	.15	7.97	4.35	4.35	4.35	3.62	68	.62	.13
7	10.87	.00	10.87 <sup>b</sup>	5.80	5.07	4.35	3.62	71	1.07	.57
8	2.90	.22	11.59	8.70	7.97	4.35	3.62	58	.72	.43
9	.72	.88	8.70	7.25	5.07	4.35	3.62	66	.78	.12
10	.00	1.00	7.25	5.07	4.35	3.62	2.90	77	.92	.62

<sup>a</sup> $p$ -value of the hypothesis testing  $H_0: p \leq p_0$  versus  $H_1: p > p_0$  using the exact binomial test, where  $p$  and  $p_0$  are the correct identification rate and  $1/I$  (total unique faces guessed by all participants), respectively.

<sup>b</sup>Target face matches the contour image.

may use some information from the contour image to narrow down the choice set, it is very difficult for them to identify the correct face beyond random guesses within the choice set. Fourth, even when people correctly identified the faces from the contour images, they were not confident about their judgments based on the contour images. For example, the average confidence score for all 37 correct guesses was 48.6% (SD = 13.6%), which is not significantly larger than the average confidence score for the incorrectly identified faces within individuals<sup>9</sup> (M = 44.6%, SD = 6.6%) or across individuals<sup>10</sup> (M = 44.5%, SD = 8.5%).

**Effect of CaF on anonymizing age.** To assess whether the CaF representations effectively anonymize age, we hired three research assistants to estimate the ages of all faces in the databases. We used the average perceived age for each face as its estimated age. We first calculated the absolute difference in estimated ages between a face a participant guessed and the actual target face corresponding to a contour, then averaged the values of all participants' guesses.

The results seem to suggest that age can be completely anonymized using contour images. For example, the difference between the target face and a guessed face ranged from 4.5 years to 9.0 years across the ten contour images (M = 6.23, SD = 1.47), which is not statistically different from a random

guess (M = 6.93, SD = 1.23).<sup>11</sup> To quantify the effectiveness of contour images in age anonymization, we defined an age ratio using the estimated ages of the target faces as ground truth. We calculated the age ratio by dividing the mean absolute error of age guess by the mean absolute error of random guess (i.e., if participants randomly pick a person from the database). Higher age ratios indicate better anonymization performance (where 1 represents a random guess). Table 9 shows that, in general, it is very difficult to infer age from contour images.

**Effect of CaF on anonymizing gender.** To assess whether contours can effectively anonymize gender, we calculated the percentage of guesses that were the incorrect gender, which ranged from 2.9% to 68.4% across all ten contour images (M = 23.5%; SD = 23.3%). Statistical tests show that incorrect identification rates for gender are significantly smaller than a random guess assumption (M = 45.8%, SD = 21.5%),<sup>12</sup> indicating that people can infer gender information from a contour image to some extent. However, further analysis of the confidence scores of the guesses within individual revealed that the average confidence level of correct gender guesses was 44.5% (SD = 7.8%), which is not significantly different from the average confidence score of incorrect gender guesses (M = 45.1%, SD = 9.2%). The results suggest that people are not confident about their gender guesses based on contour images even when they are correct.

To quantify the effectiveness of contour images in gender anonymization, we calculated a gender ratio in a similar way as the age ratio, by dividing the mean absolute error of gender guess by that of random guess. Higher gender ratios indicate better anonymization performance (where 1 represents a random guess). Table 9 shows that the ease of gender recognition

<sup>9</sup>  $p = .09$  for hypothesis testing  $H_0: m_1 \leq m_2$  versus  $H_1: m_1 > m_2$  using Wilcoxon signed rank test, where  $m_1$  and  $m_2$  are the average confidence scores of correctly identified faces and not correctly identified faces for each participant. We used the Wilcoxon signed rank test because the Shapiro-Wilk and Anderson-Darling normality tests show the distributions of the confidence scores of matched faces and mismatched faces are far from normal distributions.

<sup>10</sup>  $p = .25$  for hypothesis testing  $H_0: m_1 \leq m_2$  versus  $H_1: m_1 > m_2$  using Wilcoxon signed rank test, where  $m_1$  and  $m_2$  are the average confidence scores of correctly identified faces and not correctly identified faces across participants.

<sup>11</sup> The  $p$ -value of the two-sided paired t-test is .23 and that of the Wilcoxon signed rank test is also .23.

<sup>12</sup> The  $p$ -value of the one-sided paired t-test is less than .01, and that of the Wilcoxon signed rank test is also less than .01.

varies greatly across faces. Note this is not a direct inference of sexual dimorphism based on face contours, as assessed in Study 1. Because we allowed participants to compare the contour images with face images, the relatively higher accuracy may be due to a correlation in facial contour similarity.

Overall, our analysis provides convincing evidence that it is not easy for a human to guess the identity, age, or gender associated with a contour image even if the corresponding database of face images includes just 260 photos. This suggests that contour images can anonymize one's identity, as well as age and gender—two critical factors associated with potential discrimination in many real-life contexts—to prevent human recognition/perception.

### Study 3: A Field Implementation of CaF in an Online Dating Context

In Study 3, we aimed to investigate the feasibility of applying the CaF framework for real-life decision making that is based on or affected by face perceptions while protecting the privacy of individuals. We chose to implement the CaF framework in the online dating context for three reasons. First, face and face perceptions play an important role in the decision-making process in an online dating context. In a field experiment run by OkCupid, profile pictures accounted for over 90% of users' profile ratings, while text information accounted for less than 10% (Hern 2014). Second, from the customers' perspective, revealing one's facial identity on an online dating website/app poses inherent privacy risks, as an online dating profile usually contains more sensitive information about the user and is usually accessible to a wider audience. As a result, users often experience negative feelings because of the exposure of their faces (Cobb and Kohno 2017). Third, online dating has grown into a billion-dollar industry with tremendous commercial, personal, and social value. In the United States, statistics show that there are over 2,100 dating apps or websites and the online dating service industry is estimated to be worth \$3 billion (IBISWorld 2019). A proper solution to the face dilemma in the online dating context can substantially improve customer experience and recommendation systems, thus increasing customer trust and retention, frequency of use, and communication (the presentation of solutions to customers), among other aspects, which are also present in other industries. Next, we describe the field study design, then present analysis of the customers' responses.

#### The Field Study Design

We constructed an online dating site (denoted as "MKQ"), similar to Tinder (<https://tinder.com>), which relies on facial information and minimal demographics. Unlike Tinder users, MKQ users do not search for potential partners on the site; instead, MKQ matches users automatically based on their facial contour preferences. After registering, users are first asked to provide a facial contour image (or submit a face photo to be converted to facial contour) and basic information such as

gender, age, height, and university. To enable automatic matching, users are asked to complete a calibration task by rating randomly assigned facial contours taken from people of the opposite gender using a seven-point scale (everything else being equal, 1 represented a low likelihood of befriending a person with that contour and 7 represented a high likelihood of befriending a person with that contour). Users decide how many contours they would like to rate, knowing that the more they rate, the more accurate the matching will be. Once a user receives a notification about a match made by the site, they can go to the website to view details about the matched partner. The user may click "yes" to accept the match and initiate contact or click "no" to reject the match and wait for the next round of matching. If both parties click "yes," the two matched users can start to communicate in a private chatroom. We provide details about the matching procedure in Web Appendix B.

MKQ is implemented using PHP and MySQL on a cloud server rented from 1&1 IONOS, a cloud and hosting firm (<https://www.ionos.com/>). Due to limited resources, we did not promote MKQ commercially. All MKQ users needed a university-based email address to be eligible.

#### Analysis and Results

By the end of the first month after the launch of MKQ, 385 users (106 men, 217 women, 62 undisclosed) had registered using their personal university email accounts. The users understood that MKQ is a real service that matches individuals with potentially desirable partners of the opposite gender. There is no fee to join, and members do not receive any compensation. By the end of the first month, 198 (70 men, 128 women) out of the 385 members had completed 13,311 contour evaluations with ratings covering the entire range of the scale ( $M = 3.54$ ,  $SD = 1.47$ ). Evidently, individuals can make strong negative and positive inferences from contour images in an online dating context, based on the number of contour ratings at the extreme ends of the scale (ratings of 1 = 1,339; ratings of 2 = 1,977; ratings of 3 = 3,102; ratings of 4 = 3,522; ratings of 5 = 2,136; ratings of 6 = 957; ratings of 7 = 278). The MKQ members appear to have had no problem making discriminant inferences based on contour images to facilitate dating decisions.

*Users' dating-related inferences from CaF representations.* To investigate whether the CaF representation can elicit users' dating-related inferences and their drivers, we compared the CaF model with two variations: a nonoutline model and an outline model, using site members' average ratings for each contour image in the calibration task as the dependent variable. An assessment of model fit shows that the full CaF model explains 76% (adjusted  $R^2 = .57$ ) of the variance in users' contour ratings, the nonoutline model explains 40% of the variance (adjusted  $R^2 = .36$ ), and the outline model explains 69% of the variance (adjusted  $R^2 = .51$ ). The results showed that the CaF representation is a feasible way of eliciting customers' dating-related decisions and the contour outline

information is critical for making dating-related inferences from contour images. To investigate the contributions of each facial part's outline information in driving contour-based dating decisions, we compared five models using the 14 nonoutline variables and the outline variables of each facial part as independent variables. The results revealed that including outline information for the nose, mouth, or face can significantly improve the model fit from 40% to 55% (adjusted  $R^2 = .45$ ), 54% (adjusted  $R^2 = .46$ ), and 59% (adjusted  $R^2 = .52$ ), respectively, whereas the outlines of eyebrows and eyes did not have a significant impact on contour-based dating decisions. The results suggest that, in general, the outlines of nose, mouth, and face play a more important role than those of eyebrows and eyes in people's dating-related inferences based on contour images.

**Users' responses to CaF-based matching.** To investigate whether customers are willing to adopt the CaF-based matching recommendations in the online dating context, we analyzed users' responses to the CaF-based matching. During the first month of operation, we made six rounds of matchings for a total of 175 matched pairs among 171 (81 men and 90 women) unique users (some of whom were matched more than once). For ease of explanation, we refer to each time a user receives a matching recommendation as one matching event. In 123 (35.1%) out of the 350 matching events, users returned to the website and viewed the detailed information about the match after receiving the recommendation. Out of these 123 returning users, 114 (92.7%) responded by clicking either "yes" (103; 83.7%), indicating their willingness to initiate contact, or "no" (11; 8.9%), indicating their refusal to initiate contact. The remaining 9 users (7.3%) did not take any action after viewing the information. Out of the 103 matching events when users clicked "yes," the matched partners clicked "yes" in 52 events (50.5%), leading to 26 mutually interested pairs. The remaining 51 events when users clicked "yes" (49.5%) involved matched partners who either never replied or clicked "no." By the end of the first month, we observed that 22 (84.6%) out of the 26 mutually interested pairs engaged in active discussions beyond exchanges of greetings. Nine pairs (18 users) eventually felt comfortable enough to exchange personal contact information to continue communication off-site, representing successful initial matching by MKQ.

Messages in these chatrooms reveal that many users were happy with the experience and were comfortable evaluating contour images and providing contour images of themselves for dating purposes (e.g., "It's amazing how it can match based on my preference on facial contours!"; "I found the facial contour preference calibration tests super fun!"). The users generally appreciated the benefits of MKQ in terms of protecting their privacy and avoiding appearance-based interests (e.g., "This is so simple and clean, much better than the normal (networking) sites, where all my information is exposed"; "I like the fact that we don't see each other's photos in here"). In addition, those who did not want their dating chances to be influenced by their appearance commented very positively on

the contour-based dating site. Those who normally are considered beautiful/handsome did not want to date people who were simply attracted to their looks, and those who normally are considered less desirable, appearance-wise, did not want to be eliminated from consideration without having had a chance to show who they truly are. The use of CaF addresses these concerns while still allowing users to make relevant, face-based perceptual inferences.

Although our field implementation was small in scale, it provides evidence that the CaF framework can be successfully implemented and embraced by consumers in real life. Results from the field implementation showed that consumers are willing to accept and use contour images to make real-life decisions and appreciate contour images for their ability to protect privacy while preserving valuable information for perception.

## Discussion

This research provides a means by which researchers and practitioners can best represent and analyze face information to elicit desired face perceptions and mask the undesired face perceptions for perceivers. We make several contributions to the marketing literature. First, we proposed a rigorous method to model face data that has never been used in the business, social science, or computer science disciplines. Second, for the first time in the literature and practice, face contours were used to elicit perceptions from humans on 15 fundamental dimensions, providing a comprehensive understanding of the trade-offs between desirable (e.g., emotional state) and undesirable (e.g., age, gender) facial inferences. Third, the CaF framework has proven to effectively address the privacy-perception trade-offs associated with using facial data. Finally, we demonstrated the feasibility of the CaF framework by implementing it in a real-life online dating context. The evidence from the field study supports the use of the CaF framework in a recommendation system to assist users in making face-based decisions.

In practice, this research has several implications for business practices related to facial data, especially in contexts where face perceptions play an important role in decision making, such as dating, social networking, retailing, advertising, employee allocation, and customer relationship management. First, the CaF provides a viable solution for marketers to resolve the face dilemma, allowing business to protect privacy while preserving the value of face perceptions when managing customers' facial data. For example, social media firms can use the CaF representation as a way to help prevent demographic-based discrimination while preserving users' ability to infer useful information for social networking. In retailing contexts where customers' facial data are collected and used to support marketing decisions, firms can use the CaF representation to protect customers' identity privacy while maintaining the ability of employees to infer useful information for improving marketing effectiveness, including segmentation, targeting, customer relationship management, and product/service personalization. Second, firms can use CaF to systematically model customers' face perceptions or face-based preferences to

provide better employee allocation while masking sensitive information (e.g., age, gender) of the target face to avoid discrimination. This would be helpful in application contexts where customers' evaluations of products and services are affected by the facial features of service providers, such as professional service evaluation and personal selling. Finally, as a novel method of modeling face information, CaF can be incorporated into face-related recommendation systems even when face images are used instead of contours (as used in the CaF representation).

There are three caveats in implementing the CaF framework in practice. First, firms need to evaluate customers' preferences regarding the privacy–perception trade-offs before adopting the CaF framework in their operations. Firms (e.g., online dating sites) can also customize the CaF representation to meet the different needs of their users. As an example, for users who do not have substantial privacy concerns, firms can add more details of the face, such as wrinkles, or other facial parts, such as ears, to the CaF representation to improve face perceptions. Second, the variation in the reliability of contour-based perceptions across different perceptual dimensions indicates that marketers should have a good understanding of perceivers' (e.g., customers') desired and undesired face perceptions before applying the CaF framework to specific contexts. Third, it should be noted that a person with a unique contour for their facial features (as compared with other faces in the database) can be more easily recognized. For example, based on the feedback from study participants, the face with the highest recognition rate in Study 2 was recognized due to the uniqueness of its mouth and face contours.

The CaF belongs to the research domain called “artificial empathy,” defined as “the ability of nonhuman models to predict a person's internal states (e.g., cognitive, affective, physical) given the signals he or she emits (e.g., facial expression, voice, gesture) or to predict a person's reaction (including, but not limited to internal states) when he or she is exposed to a given set of stimuli (e.g., facial expression, voice, gesture, graphics, music, etc.)” (Xiao, Kim, and Ding 2013, p. 244). The CaF framework enables us to predict how a person will react when they see someone's face, without revealing the face itself. We hope this research helps stimulate further discussion about how to address privacy-related problems associated with visual data in various marketing contexts. Here, we list several promising extensions. First, it would be worthwhile to explore the use of more realistic-looking CaF representation, such as 3-D CaF representation or a representation that includes other facial parts (e.g., ears). Second, given the significant heterogeneity in human faces, it would be worthwhile to test the CaF framework in a larger-scale study with more diverse faces/contours. Researchers might also explore the heterogeneity in how people respond to specific features in contour images of different faces using a larger data set. Finally, researchers might explore the use of the CaF modeling method to model other types of visual information in related marketing contexts, such as perceptions of brands' visual representations (e.g., BMW's iconic grill).

As the first attempt to address the emerging concerns among customers and marketers over facial data protection, the CaF framework may yield benefits such as increased customer trust and retention and better service recommendations. We hope this new framework will become a valuable tool for businesses using faces or face representations and help improve consumer welfare in marketing.

### Associate Editor

Vikas Mittal

### Acknowledgments

The preliminary work in this research was completed while the first author was a postdoctoral scholar at the Institute for Sustainable Innovation and Growth (iSIG) at the School of Management, Fudan University (FDSM), where the third author serves as an advisory professor. The authors thank the *JMR* review team for their insightful suggestions. They are grateful to Tom Hartley (University of York) for providing the list of social traits used in Vernon et al. (2014). They also thank Vithala R. Rao (Cornell University), Li Xiao (Fudan University), Songting Dong (University of New South Wales), David Stillwell (University of Cambridge), Vincent Mak (University of Cambridge), Yi Wang (Central University of Finance and Economics), Jie Xu (Fudan University), Ming Ji (House of Enlightened Needs) and Fulin Zhang (iSIG) for their valuable comments as well as the seminar and conference participants at Shanghai Jiao Tong University, the 38th ISMS Marketing Science Conference, and the 2016 Annual Conference of China Marketing Science for their feedback.

### Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported by the National Natural Science Foundation of China [Grant 71301099], the Institute for Sustainable Innovation and Growth (iSIG) at the School of Management at Fudan University, Shenzhen University, and the Cambridge Judge Business School at University of Cambridge.

### References

- Acquisti, Alessandro and Christina Fong (2020), “An Experiment in Hiring Discrimination via Online Social Networks,” *Management Science*, 66 (3), 1005–24.
- Ada Lovelace Institute (2019), “Beyond Face Value: Public Attitudes to Facial Recognition Technology” (September 2), <https://www.adalovelaceinstitute.org/beyond-face-value-public-attitudes-to-facial-recognition-technology/>.
- Altonji, Joseph G. and Rebecca M. Blank (1999), “Race and Gender in the Labor Market,” *Handbook of Labor Economics*, 3 (Part C), 3143–3259.
- Benes, Ross (2018), “Protecting Customer Data Is a Top Challenge for CMOs” *eMarketer*, (November 13), <https://www.emarketer.com/content/protecting-customer-data-is-a-top-challenge-for-cmos>.

- Bóo, Florencia L., Martín A. Rossi, and Sergio S. Urzúa (2013), "The Labor Market Return to an Attractive Face: Evidence from a Field Experiment," *Economics Letters*, 118 (1), 170–72.
- Bruce, Vicki and Steve Langton (1994), "The Use of Pigmentation and Shading Information in Recognizing the Sex and Identities of Faces," *Perception*, 23 (7), 803–22.
- Caers, Ralf and Vanessa Castelyns (2011), "LinkedIn and Facebook in Belgium: The Influences and Biases of Social Network Sites in Recruitment and Selection Procedures," *Social Science Computer Review*, 29 (4), 437–48.
- Chellappa, Rama, Charles L. Wilson, and Saad Sirohey (1995), "Human and Machine Recognition of Faces: A Survey," *Proceedings of the IEEE*, 83 (5), 705–41.
- Cobb, Camille and Tadayoshi Kohno (2017), "How Public is My Private Life? Privacy in Online Dating," in *Proceedings of the 26th International Conference on World Wide Web*. New York: Association for Computing Machinery, 1231–40.
- DeShields, Oscar W., Jr., Ali Kara, and Erdener Kaynak (1996), "Source Effects in Purchase Decisions: The Impact of Physical Attractiveness and Accent of Salesperson," *International Journal of Research in Marketing*, 13 (1), 89–101.
- Du, Liang, Meng Yi, Erik Blasch, and Haibin Ling (2014), "GARP-Face: Balancing Privacy Protection and Utility Preservation in Face De-Identification," in *IEEE International Joint Conference on Biometrics*. Clearwater, FL: Institute of Electrical and Electronics Engineers 1–8.
- Ellis, Evelyn and Philippa Watson (2012), *EU Anti-Discrimination Law*. Oxford, UK: Oxford University Press.
- Fieguth, Paul, Paul Bloore, and Adrian Domsa (2004), "Phase-Based Methods for Fourier Shape Matching," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*. Montreal: Institute of Electrical and Electronics Engineers, 425–28.
- George, Patricia A. and Graham J. Hole (2000), "The Role of Spatial and Surface Cues in the Age-Processing of Unfamiliar Faces," *Visual Cognition*, 7 (4), 485–509.
- Gomulya, David, Elaine M. Wong, Margaret E. Ormiston, and Warren Boeker (2017), "The Role of Facial Appearance on CEO Selection After Firm Misconduct," *Journal of Applied Psychology*, 102 (4), 617–35.
- Gorn, Gerald J., Yuwei Jiang, and Gita Venkataramani Johar (2008), "Babyfaces, Trait Inferences, and Company Evaluations in a Public Relations Crisis," *Journal of Consumer Research*, 35 (1), 36–49.
- Graham, John R., Campbell R. Harvey, and Manju Puri (2016), "A Corporate Beauty Contest," *Management Science*, 63 (9), 3044–56.
- Gross, Ralph and Alessandro Acquisti (2005), "Information Revelation and Privacy in Online Social Networks," in *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*. New York: Association for Computing Machinery, 71–80.
- Hallgren, Kevin A. (2012), "Computing Inter-Rater Reliability for Observational Data: An Overview and Tutorial," *Tutorials in Quantitative Methods for Psychology*, 8 (1), 23–34.
- Henderson, Audrey J., Iris J. Holzleitner, Sean N. Talamas, and David I. Perrett (2016), "Perception of Health from Facial Cues," *Philosophical Transactions of the Royal Society B*, 371 (1693), 20150380.
- Hennig-Thurau, Thorsten, Markus Groth, Michael Paul, and Dwayne D. Gremler (2006), "Are All Smiles Created Equal? How Emotional Contagion and Emotional Labor Affect Service Relationships," *Journal of Marketing*, 70 (3), 58–73.
- Hern, Alex (2014), "OKCupid: We Experiment on Users. Everyone Does," *The Guardian* (July 29), <https://www.theguardian.com/technology/2014/jul/29/okcupid-experiment-human-beings-dating>.
- Hills, Peter J., J. Michael Pake, Jack R. Dempsey, and Michael B. Lewis (2018), "Exploring the Contribution of Motivation and Experience in the Postpubescent Own-Gender Bias in Face Recognition," *Journal of Experimental Psychology: Human Perception and Performance*, 44 (9), 1426–46.
- IBISWorld (2019), "Dating Services Industry in the US—Market Research Report," (December), <https://www.ibisworld.com/united-states/market-research-reports/dating-services-industry/>.
- Johnston, Alan, Harold Hill, and Nicole Carman (1992), "Recognising Faces: Effects of Lighting Direction, Inversion, and Brightness Reversal," *Perception*, 42 (11), 365–75.
- Kahle, Lynn R. and Pamela M. Homer (1985), "Physical Attractiveness of the Celebrity Endorser: A Social Adaptation Perspective," *Journal of Consumer Research*, 11 (4), 954–61.
- Keh, Hean T., Run Ren, Sally R. Hill, and Xuan Li (2013), "The Beautiful, the Cheerful, and the Helpful: The Effects of Service Employee Attributes on Customer Satisfaction," *Psychology and Marketing*, 30 (3), 211–26.
- Landau, Jacqueline (1995), "The Relationship of Race and Gender to Managers' Ratings of Promotion Potential," *Journal of Organizational Behaviour*, 16 (4), 391–400.
- Lestrel, Pete E. (2008), *Fourier Descriptors and Their Applications in Biology*. Cambridge, UK: Cambridge University Press.
- Lu, Shasha, Li Xiao, and Min Ding (2016), "A Video-Based Automated Recommender (VAR) System for Garments," *Marketing Science*, 35 (3), 484–510.
- Martin, Kelly D., Abhishek Borah, and Robert W. Palmatier (2017), "Data Privacy: Effects on Customer and Firm Performance," *Journal of Marketing*, 81 (1), 36–58.
- Martin, Nicole (2018), "Are AI Hiring Programs Eliminating Bias or Making It Worse?" *Forbes* (December 13), <https://www.forbes.com/sites/nicolemartin1/2018/12/13/are-ai-hiring-programs-eliminating-bias-or-making-it-worse>.
- Mayer, Christine, Sonja Windhager, Katrin Schaefer, and Philipp Mitteroecker (2017), "BMI and WHR are Reflected in Female Facial Shape and Texture: A Geometric Morphometric Image Analysis," *PLoS One*, 12 (1), e0169336.
- Meinhardt-Injac, Bozana, Malte Persike, and Günter Meinhardt (2013), "Holistic Face Processing Is Induced by Shape and Texture," *Perception*, 42 (7), 716–32.
- Newton, Elaine M., Latanya Sweeney, and Bradley Malin (2005), "Preserving Privacy by De-Identifying Face Images," *IEEE Transactions on Knowledge and Data Engineering*, 17 (2), 232–43.
- Olivola, Christopher Y., Friederike Funk, and Alexander Todorov (2014), "Social Attributions from Faces Bias Human Choices," *Trends in Cognitive Sciences*, 18 (11), 566–70.

- Oosterhof, Nikolaas N. and Alexander Todorov (2008), "The Functional Basis of Face Evaluation," *Proceedings of the National Academy of Sciences*, 105 (32), 11087–92.
- Porcheron, Aurélie, Emmanuelle Mauger, and Richard Russell (2013), "Aspects of Facial Contrast Decrease with Age and Are Cues for Age Perception," *PLoS One*, 8 (3), e57985.
- Rezlescu, Constantin, Brad Duchaine, Christopher Y. Olivola, and Nick Chater (2012), "Unfakeable Facial Configurations Affect Strategic Choices in Trust Games With or Without Information About Past Behavior," *PloS One*, 7 (3), e34293.
- Ruffle, Bradley J. and Ze'ev Shtudiner (2014), "Are Good-Looking People More Employable?" *Management Science*, 61 (8), 1760–76.
- Rule, Nicholas O. and Nalini Ambady (2008), "The Face of Success: Inferences from Chief Executive Officers' Appearance Predict Company Profits," *Psychological Science*, 19 (2), 109–11.
- Senior, Andrew W. and Sharathchandra Pankanti (2011), "Privacy Protection and Face Recognition," in *Handbook of Face Recognition*, Stan Z. Li and Anil K. Jain, eds. London: Springer, 671–91.
- Small, Deborah A. and Nicole M. Verrochi (2009), "The Face of Need: Facial Emotion Expression on Charity Advertisements," *Journal of Marketing Research*, 46 (6), 777–87.
- Smith, Kit (2019), "53 Incredible Facebook Statistics and Facts," Brandwatch (June 1), <https://www.brandwatch.com/blog/face-book-statistics/>.
- Tanaka, James W. and Martha J. Farah (1993), "Parts and Wholes in Face Recognition," *Quarterly Journal of Experimental Psychology A*, 46 (2), 225–45.
- Tibshirani, Robert (1996), "Regression Shrinkage and Selection via The Lasso," *Journal of the Royal Statistical Society: Series B*, 58 (1), 267–88.
- Tibshirani, Robert, Guenther Walther, and Trevor Hastie (2001), "Estimating the Number of Clusters in a Data Set via the Gap Statistic," *Journal of the Royal Statistical Society: Series B*, 63 (2), 411–23.
- Todorov, Alexander (2017), *Face Value: The Irresistible Influence of First Impressions*. Princeton, NJ: Princeton University Press.
- Todorov, Alexander, Anesu N. Mandisodza, Amir Goren, and Crystal C. Hall (2005), "Inferences of Competence from Faces Predict Election Outcomes," *Science*, 308 (5728), 1623–26.
- Todorov, Alexander, Chris P. Said, Andrew D. Engell, and Nikolaas N. Oosterhof (2008), "Understanding Evaluation of Faces on Social Dimensions," *Trends in Cognitive Sciences*, 12 (12), 455–60.
- Torrance, Jaimie S., Joanna Wincenciak, Amanda C. Hahn, Lisa M. DeBruine, and Benedict C. Jones (2014), "The Relative Contributions of Facial Shape and Surface Information to Perceptions of Attractiveness and Dominance," *PLoS One*, 9 (10), e104415.
- Turk, Matthew and Alex Pentland (1991), "Eigenfaces for Recognition," *Journal of Cognitive Neuroscience*, 3 (1), 71–86.
- Valentine, Katherine A., Norman P. Li, Lars Penke, and David I. Perrett (2014), "Judging a Man by the Width of His Face: The Role of Facial Ratios and Dominance in Mate Choice at Speed-Dating Events," *Psychological Science*, 25 (3), 806–11.
- Vernon, Richard J.W., Clare A.M. Sutherland, Andrew W. Young, and Tom Hartley (2014), "Modeling First Impressions from Highly Variable Facial Images," *Proceedings of the National Academy of Sciences*, 111 (32), 3353–61.
- Wang, Ze, Surendra N. Singh, Yexin Jessica Li, Sanjay Mishra, Maureen Ambrose, and Monica Biernat (2017), "Effects of Employees' Positive Affective Displays on Customer Loyalty Intentions: An Emotions-as-Social-Information Perspective," *Academy of Management Journal*, 60 (1), 109–29.
- Xiao, Li, Hye-jin Kim, and Min Ding (2013), "An Introduction to Audio and Visual Research and Applications in Marketing," in *Review of Marketing Research*, Naresh Malhotra, ed. Bingley, UK: Emerald Group Publishing, 213–253.
- Xiao, Li and Min Ding (2014), "Just the Faces: Exploring the Effects of Facial Features in Print Advertising," *Marketing Science*, 33 (3), 338–52.
- Zebrowitz, Leslie A. and Susan M. McDonald (1991), "The Impact of Litigants' Baby-Facedness and Attractiveness on Adjudications in Small Claims Courts," *Law and Human Behavior*, 15 (6), 603–23.
- Zhang, Dengsheng and Guojun Lu (2002), "Shape-Based Image Retrieval Using Generic Fourier Descriptor," *Signal Processing: Image Communication*, 17 (10), 825–48.