Diagnostic Features for Human Categorisation of Adult and Child Faces

Simon Faghel-Soubeyrand^{1,2}, Juliane A. Kloess², Frédéric Gosselin¹, Ian Charest^{1,2}, and
Jessica Woodhams²

¹Département de Psychologie, Université de Montréal

²Centre for Applied Psychology, School of Psychology, University of Birmingham

Author Note

Simon Faghel-Soubeyrand <https://orcid.org/0000-0002-7867-8589>

Juliane A. Kloess <https://orcid.org/0000-0002-8342-7880>

Frédéric Gosselin <https://orcid.org/0000-0002-3797-4744>

Ian Charest <https://orcid.org/0000-0002-3939-3003>

Jessica Woodhams <https://orcid.org/0000-0002-9674-6653>

The authors have no known conflict of interest to disclose.

The authors would like to express their gratitude and appreciation to West Midlands Police for their assistance, time and effort in supporting the study undertaken and presented here, as well as the Institute for Global Innovation for financially supporting this study.

Correspondence concerning this article should be addressed to Jessica Woodhams, School of Psychology, University of Birmingham, 52 Pritchatts Road, Edgbaston, Birmingham, B15 2SA, UK. E-mail: J.Woodhams@bham.ac.uk

Word count : 4,444

Abstract

Knowing how humans differentiate children from adults has useful implications in many areas of both forensic and cognitive psychology. Yet, how we extract age from faces has been surprisingly underexplored in both disciplines. Here, we used a novel data-driven experimental technique to objectively measure the facial features human observers use to categorise child and adult faces. Relying on more than 35,000 trials, we used a reverse correlation technique that enabled us to reveal how specific features which are known to be important in face-perception—position, spatial-frequency, and orientation—are associated with accurate child and adult discrimination. This showed that human observers relied on evidence in the nasal bone and eyebrow area for accurate adult categorisation, while they relied on the eye and jawline area to accurately categorise child faces. For orientation structure, only facial information of vertical orientation was linked to face-adult categorisation, while features of horizontal and, to a lesser extent oblique orientations, were more diagnostic of a child face. Finally, we found that spatial-frequency diagnosticity showed a U-shaped pattern for face-age categorisation, with information in low and high spatial frequencies being diagnostic of child faces, and mid spatial frequencies being diagnostic of adult faces. Through this first characterisation of the facial features of face-age categorisation, we show that important information found in psychophysical studies of face-perception in general (i.e. the eye area, horizontals, and mid-level spatial-frequencies) is crucial to the practical context of face-age categorisation, and present data-driven procedures through which face-age classification training could be implemented for real-world challenges.

Introduction

The amount of imagery depicting child sexual abuse (referred to here as indecent images of children, IIOC¹; Edwards, 2013) in circulation has dramatically increased in the last 25 years, from estimates of thousands of such images in the late 1990s to millions or tens of millions nowadays (Home Office, 2017; Krasodonski-Jones et al., 2017). In the United Kingdom, the task of assessing whether a suspect's digital material includes IIOC is conducted by digital forensics analysts who are members of specialist teams in police forces across the UK. This task involves determining whether a given image a) depicts a child, and b) if it is of an indecent nature (Kloess et al., 2019). In other countries, professionals undertaking the same task have the added challenge of distinguishing between children within specific age bands (e.g., in Germany, distinctions are made regarding IIOC that depict children under 18 years and children under 14 years; Ratnayake et al., 2014).

Identifying and classifying IIOC is a time-consuming task which is also stressful and emotionally challenging (Franqueira et al., 2018; Ratnayake et al., 2014). In addition, some imagery presents particular difficulties for the human classifier (Kloess et al., 2019; Ratnayake et al., 2014). Unsurprisingly, the potential to (semi-)automate the identification and classification of IIOC using algorithms is currently being explored (e.g., Gao & Ai, 2009; Home Office, 2017; Sae-Bae et al., 2014). Developing our understanding of what makes children's faces distinct from adult faces, in particular the features used by human observers to discriminate them, is therefore an important step in developing software and advancing training for digital forensics analysts (and allied professionals). This study builds on existing qualitative research that has explored the aspects and attributes within imagery that digital forensics analysts report drawing on in order to inform their decision-making in the process of identifying and classifying IIOC, including specific facial and bodily features of children

¹ This material is also referred to as child sexual exploitation material (Merdian et al., 2021), child pornography and sexually exploitative material of children (Franqueira et al., 2018).

(Kloess et al., 2019, 2021; Michalski et al., 2019), by applying a new (data-driven) reverse-correlation technique relying on Gabor wavelets to this problem area. In doing so, it fills important gaps in the literature, given i) the scarcity of studies that have examined facial differences associated with age (Gao & Ai, 2009) and, most importantly, ii) the absence of studies having revealed the specific facial features human observers use for age categorisation.

There are many physical changes associated with the maturation of faces. During childhood, craniofacial growth alters the shape and size of a child's face, as well as the relative positions of different facial features. These changes are due to the development of permanent teeth and puberty. The greatest changes are seen between birth and the age of five years, with the rate of change being non-linear (Ricanek et al., 2015). Compared to adult faces, children's faces are typically characterized by a "protruding forehead, large head, round face, big eyes and a small nose or mouth" (Komori & Nittono, 2013, p. 285).

Findings from research on the border-control task of the facial matching of children demonstrate that differentiating between faces of similar-aged children is difficult for non-expert (Kramer et al., 2018) and expert (i.e., nursery workers and super-recognizers; Bate et al., 2020; Belanova et al., 2018) observers, and that this task is harder the younger the child who is depicted (Michalski et al., 2019). In combination with research regarding how child and adult faces differ physically, this suggests that, particularly with younger children, their faces are defined by a common set of qualities which differ to the faces of adults. While these qualities may make the task of face-matching children more difficult, they should make the task of discriminating between child and adult faces easier, assuming that those who are engaged in the discrimination task attend to valid features for age categorisation.

In relation to this, Kloess et al. (2019) assessed the inter-rater agreement of law enforcement personnel experienced at identifying and classifying IIOC, finding levels of

agreement that were not always adequate. In subsequent focus groups, the officers reported that features they found more indicative of a younger age included large eyes relative to the rest of the face, the presence of milk teeth, or the eruption of adult teeth². They also reported that the classification task was more challenging when there was a mismatch between the apparent maturity of the depicted person's face and their body (or vice-versa), and where the depicted person was made to look younger or older (e.g., by means of make-up and clothing). More recently, Kloess et al. (2021) followed up on their 2019 study by asking digital forensics analysts to describe the key facial attributes they use in determining the presence of a child in imagery. These included: (i) large eyes, (ii) a small nose, (iii) a round facial shape, (iv) an absence of cheekbones, (v) early signs of teeth development, and (vi) smooth skin.

Attending to the face (rather than other cues) may thus improve the accuracy of the identification and classification of IIOC. It is therefore important to further investigate (i) whether participants do indeed rely on the features previously identified by police staff when discriminating between child and adult faces, and, relatedly, (ii) whether attending to these features/facial areas does indeed lead to more accurate decision-making. In this matter, psychophysical studies on face perception have outlined the behavioural relevance of important features of facial images, such as orientation structure (e.g. Dakin & Watt, 2009; Duncan et al., 2017; Goffaux et al., 2016; Goffaux & Greenwood, 2016; Pachai, Sekuler & Bennett, 2013), spatial location (e.g. Adolphs et al., 2005; Faghel-Soubeyrand et al., 2019; Gosselin & Schyns, 2001; Tardif et al., 2019), colour (Benitez-Quiroz et al., 2018; Dupuis-Roy et al., 2019; Nestor & Tarr, 2008; Young et al., 2013), and spatial frequency (Estéphan et al., 2018; Faghel-Soubeyrand et al., 2020; Tardif et al., 2017; Willenbockel et al., 2010; Gaspar, Sekuler & Bennett, 2008). However, while the use of these physical properties have been revealed for various face-processing tasks (e.g. face-identification,

² The use of dentition developmental stage in making accurate age assessments has been questioned, and may therefore only be useful in images where the depicted person's mouth is open (see Ratnayake et al., 2014), which is not always the case in the context of IIOC.

face-expression & face-gender recognition in the above-cited studies), the diagnostic features for face-age recognition have never been explored.

The present study tackles this question by employing a recently developed reverse-correlation technique, called Diagnostic Feature Mapping (DFM; Alink & Charest, 2020), in order to precisely quantify the use of three core physical properties of facial images of children and adults, namely location, orientation structure and spatial frequency. These properties are randomly sampled in a similar way to Bubbles (Gosselin & Schyns, 2001). A multiple linear regression applied to these samples and response accuracy can reveal the diagnosticity of behaviourally relevant features of images which are known to be crucial to the human visual system (i.e., different spatial-frequencies, orientations, spatial coordinates and sizes in the striate cortex; Hubel & Wiesel, 1968). Contrasting with previous psychophysical studies, however, is the capacity of DFM to reveal these three core features in a single experimental paradigm, permitting a thorough description of the facial information relied for categorization tasks. Using such Gabor features was also motivated by studies in computer vision (Gao & Ai, 2009), suggesting that models using Gabors as priors perform well in explaining differences between facial images of children and adults. DFM therefore provided an informed, comprehensive and entirely data-driven way to reveal the specific facial features associated with face-age categorisation.

Method

Participants

A screening questionnaire was used to screen for conditions that would prevent a volunteer being able to take part in the experiment. This resulted in one volunteer being excluded. Of the 16 participants who took part in the experiment (male = 9, female = 7, right-handed = 14, left-handed = 2, White British = 12, and Asian-Indian = 4, aged between 21-47 years ($M = 38.44$, $SD = 8.33$)), seven were students or staff at the University of

Birmingham, and nine were digital forensics analysts recruited from West Midlands Police. All participants had no history of psychiatric diagnosis and all had normal or corrected-to-normal vision. Full ethical approval for the study was granted by the Science, Technology, Engineering and Mathematics Ethical Review Committee at the University of Birmingham (ERN_15-1374AP10A). In addition, the researchers adhered to the British Psychological Society's Code of Ethics and Conduct (2018) throughout the study, and the experiment was conducted in accordance with the Declaration of Helsinki. Participants were recruited through the University of Birmingham's Research Participation Scheme, via the last author's professional network at the University of Birmingham, and by means of an advertisement that was circulated via internal channels within West Midlands Police to officers who were working with child sexual exploitation materials on a daily basis.

Stimuli

A total of 12 images depicting faces of either children or adults (6 adults, 6 males; all with neutral expressions) were selected from the Radboud Face dataset (Langner et al., 2010), and converted into 250 x 250 pixels grayscale images. Each image was aligned based on twenty landmarks (averaged to six mean coordinates for left and right eyes, left and right eyebrows, nose and mouth) using Procrustes transformations (rotation, scaling and translation), and revealed through an elliptical mask that excluded non-facial cues such as hair, ears and neck. The luminance profile of the resulting face images was equalized across images using the SHINE toolbox (Willenbockel et al., 2010).

Gabors Wavelet Decomposition

A more in-depth description of the Gabors wavelet decomposition algorithm has been published elsewhere (Alink & Charest, 2020). Briefly, we used a custom-made Matlab program which aimed to reduce each face image to a subset of the most important 2200 Gabor wavelets. To do so, we considered wavelets of 20 different spatial frequencies (SFs),

exponentially increasing between 2.4 and 87 cycles per image (

$SF = \frac{10}{45} \times 1.08^n$; $n = [1, 2, \dots, 20]$). They had 12 orientations between 15° to 180° in

constant steps of 15° , and were centred on each possible pixel of the input grayscale image.

The final sets of Gabor features were selected based on their best fit (least-square correlation) to the original grayscale images. Amplitudes were set to an equal value for all wavelets.

Partial reconstructions of the images at every given trial were created by randomly selecting a subset of the 2200 Gabor wavelets, and summing them (see Figure 1a). The resulting range of pixel values was modified linearly so that all stimuli covered the full 0-255 range.

Apparatus

The experimental program ran on Windows computers in the MATLAB environment (The Mathworks, Natick, MA), using functions from the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were shown on 22-in. monitors (1920 x 1080 pixels at 60Hz). Participants performed the experiment in a dimly lit room, and viewing distance was maintained at ~ 76 cm.

Experimental procedure

The experiment took place in the laboratories of the School of Psychology at the University of Birmingham. Participants were first given standardized instructions about the experimental task. Each trial began with a grey screen and a central fixation point (250 ms), which participants were instructed to gaze at. The partial face reconstruction spanning 8 deg of visual angle was then shown in the central area of the screen, and remained visible until the participant's response. Participants were asked to categorise the faces depicted in the stimuli as either children or adults as fast and as accurately as possible. The stimulus presentation order was randomized for each participant. All 16 participants completed 2400 trials of this adult vs. child face discrimination task, with short breaks every 200 trials

(approximately 10 minutes). The task was completed in one session that lasted between 1.5-2 hours. Overall, a total of 38,400 trials were recorded for this study.

The quantity of face information (i.e., the number of Gabor Wavelet features) necessary to maintain an accuracy of 75% was adjusted on a trial-by-trial basis with the QUEST algorithm (Watson & Pelli, 1983). Royer and colleagues (2015) showed that such a threshold (measured using a similar method) is strongly correlated with three commonly used face recognition ability tests ($r = -.79$ with the mean of the Cambridge Face Memory Test +, the Cambridge Face Perception Test, and the Glasgow Face Matching Test short version). We thus used the number of Gabors required to maintain an accuracy of 75% as our individual index of face-age discrimination performance.

After the experiment, participants were verbally debriefed and were given a debriefing sheet to take with them. Participants recruited through the Research Participation Scheme were compensated with course credits and participants from West Midlands Police were reimbursed for their travel expenses and time.

Figure 1. Image feature decomposition and reverse-correlation procedure. A reconstructed image (B, lower image) was created using a subset of 2200 Gabors wavelets varying in spatial-frequency (from 2.4 to 87 cpi), orientation (10° to 180°), x-y coordinates and size. These 2200 wavelets, the most correlated with the original image (A), are referred to as Gabor features (see also Alink & Charest, 2020 for further details). (C) At any given trial, the stimulus was created by randomly sampling the 2200 Gabor features to reveal parts of the reconstructed image. Participants were asked to indicate whether the resulting sampled face was a child or an adult. Each participant was submitted to 2.4K trials, for a grand total of >38k trials. (D) We obtained the diagnostic features by weighing and summing the randomly sampled feature vectors across trials with z-scored behavioural accuracies.

Feature Diagnosticity Index

We quantified how the Gabor features (the Gabor wavelets) were differently associated with face representations by computing a Feature Diagnosticity index (FDi; see (Alink & Charest, 2020)) for each of the features composing our face images, for every participant, and stimulus independently. This was done by weighing the features presented across trials by the z-scored accuracy values, providing us with a three-dimensional matrix of

FDI with the dimensions number of participants ($n = 16$), number of images (12) and number of features (2200).

Results

Accuracy was close to the target 75% ($M_{acc} = .7577$, $SD_{acc} = .0443$). One participant was excluded from further analysis because of poor performance (number of Gabors threshold was greater than 3 standard deviations above the mean of all participants). The average number of Gabors required to attain 75% accuracy ranged from a low of 156.52 (i.e. best performance) to a high of 443.75 (i.e. lowest performance; $M_{nbgabors} = 275.52$, $SD_{nbgabors} = 73.86$).

Facial areas for child and adult classification

To reveal the facial attributes used to categorise faces as child vs. adult, we reprojected each Gabors FDI to their specific x-y image position, smoothed (with a Gaussian kernel with a standard deviation of 10 pixels) the resulting image regression coefficients for each stimulus, and compared them between children and adults faces with unpaired t-tests. These comparisons of regression coefficients are displayed on Figure 2a, with hot colors (positive t-statistics) indicating evidence for accurate adult classification, and cold colors (negative values) indicating child classification evidence. Significant facial features ($ps < .05$) for adult classification evidence were found in an area around the nasal bone and eyebrows structure. In contrast, we found that the eyes, nose, and jawline were used as evidence for child face categorisation. The significant regions have been highlighted on an average of all face stimuli on Figure 2b for child evidence (left bottom image) and adult evidence (right bottom image). The average regression coefficient image (irrespective of category type) is shown in supplementary Figure 1A.

Figure 2. **Facial features for child and adult categorisation.** Regression coefficients for 2D feature positions (for accurate classification of child and adult images) were spatially smoothed, and compared between child and adult faces with unpaired t-tests. On the first panel, hot colors (positive t-statistics) indicate evidence for accurate adult classification, and cold colors (negative values) indicate evidence for accurate child classification. The two bottom panels show pixels (i.e., regions of faces) significantly linked to accurate child and adult categorisation, respectively.

Spatial-frequencies for child and adult classification

To determine the spatial frequencies (or information granularity) used to accurately categorise child and adult faces, we correlated the spatial-frequency parameters of the presented Gabor features with accurate responses for child and adult faces across trials for all participants. The regression coefficients were compared between child and adult faces with unpaired t-tests. The resulting evidence magnitude for the 20 SFs bands are shown in Figure

3, with significant SF bands (t-stats with $p < .05$) highlighted. Overall, this revealed that low (coarse details of 5.1 cpi) and high SFs (fine details of 49.5 cpi) were used as evidence of child faces, while mid-SFs were relied on by our participants as evidence for adult classification. The average SFs regression coefficient (irrespective of category type) are shown in supplementary Figure 1B.

Figure 3. Spatial frequency evidence for categorisation of child and adult faces. Correlation coefficients for feature granularity (for accurate classification of child and adult images) were compared between child and adult faces for each one of the 20 SF bins (2.4-87.3 cpi) with unpaired t-tests. For visualisation purposes, the lower panel images display the facial information for each of the SF bands for child and adult faces. Here, positive t-statistics shows SF evidence for child-like faces, while negative values shows SF evidence for adult-like faces. Significant SF bands —low SF of 5 cpi and high SF of 50 cpi for child faces, mid-SF of 22 cpi for adult faces—are outlined with an asterisk.

Orientation of facial traits for child and adult classification

To determine the type of orientation structure relied upon for accurate categorisation of child and adult faces, we correlated the orientation parameter of the presented Gabor features with accurate responses for child and adult faces across trials for all participants. The regression coefficients were then compared between child and adult faces with unpaired t-tests. The corresponding t-statistics are shown in Figure 4, with significant orientation bands (t-stats with $p < .05$) highlighted. Overall, this revealed that horizontal (90°) and oblique facial structures (60° and 150°) were used as evidence for classification of a face as a child. For adult face categorisation, however, vertical features (180°) were linked to accurate responses in our participants. The average orientation regression coefficient (irrespective of category type) is shown in supplementary Figure 1A.

Figure 4. What orientation structure do we rely on to categorise child and adult faces? Correlations coefficients for feature orientation (for accurate classification of child and adult images) were compared between child and adult faces for each one of the 12 orientation bands (10° to 180° in 15° steps) with unpaired t-tests. For visualization purposes, the lower panel images display the facial information for each of the 12 orientation bands for child and

adult faces. Positive t-statistics shows orientation evidence for child-like faces, while negative values shows orientation evidence for adult-like faces. Significant orientation bands—horizontal (90°) and oblique (60° & 150°) information for child faces, vertical information (180°) for adult faces—are outlined with an asterisk.

Discussion

What specific visual features drive our accurate categorisation of faces as either a child or an adult? One practical implication of answering this question is the improvement of the identification and classification of indecent images of children (IIOC; Edwards, 2013) which involves determining whether a given image a) depicts a child, and b) if it is of an indecent nature (Home Office, 2017; Kloess et al., 2019). Here, we used a data-driven approach—Diagnostic Feature Mapping (DFM; Alink & Charest, 2020; see also Gosselin & Schyns, 2001 for a similar approach) to reveal the specific facial features human participants draw on in order to distinguish child vs. adult faces. Specifically, we revealed the use of three types of features important to the visual system—spatial-frequency, position, and orientation—that human observers extract while they complete child vs. adult categorisations, yielding a precise and comprehensive characterization of the relevant facial information for face-age categorisation.

The participants in the present study used information located around the nasal bone and eyebrows (supraorbital ridge) to classify adult faces. There was some evidence that they also used an area of the cheek, which may correspond with the flattening of the cheeks that occurs during the transition from childhood to adulthood (Kau & Richmond, 2008). In contrast, the area around the eyes, nose and jawline were more indicative of a child's face. Indeed, this matches with physical differences in faces of children: large eyes relative to the rest of the face, a small, wider, up-turned nose with a concave bridge, and a rounder chin, are all physiological characteristics that have been reported (Liggett, 1975; Sforza et al., 2011). Our findings are also in line with some of the facial features reported by digital forensics analysts in Kloess et al. (2021), who explicitly described the following facial attributes as

indicative of a younger age: (i) large eyes, (ii) a small nose, (iii) a round facial shape, and (iv) an absence of cheekbones. The present study corroborated these findings by revealing, with implicit and data-driven methods, the facial attributes linked to accurate child vs. adult categorisation. But our DFM paradigm enabled us to go further than revealing facial areas for accurate face-age categorisation from faces.

Indeed, DFM's reliance on Gabor wavelets parameters enabled us to show the specific level of detail (spatial-frequencies, SF) and orientation structure humans extract from images of faces in order to determine the age from a face. We found that the spatial-frequencies associated with age classification manifested a U pattern from low to high SF, with low SF (coarse features) and high SF (fine features) being associated with child faces, and mid-level SF being associated with adult faces. The fact that coarse facial features are associated with younger, child-like faces aligns with the layman notion that some facial features of children are larger, such as the forehead, the head, and the eyes (Komori & Nittono, 2013). Interestingly, the mid-level SFs (10 cpi) found to be indicative of adult faces matches the level of information required for (adult) facial identification found in previous studies (e.g., Tardif et al., 2017, 2019; Willenbockel et al., 2010). This might explain, in turn, why identification of children's faces has been reported to be so difficult in lab conditions (Bate et al., 2020; Belanova et al., 2018; Kramer et al., 2018), as well as in real-world conditions (e.g. Kloess et al., 2019, 2021); the level of detail of facial information for our classification of a face as a child (i.e. low & high-SFs) does not match the facial information we normally use for adult face-identification (i.e. mid-SFs).

For orientation information, we found that vertical features of a face, presumably around the nasal bone, is indicative of an adult's face. Indeed, as shown in Figure 2, the eyebrow region as well as the nasal bone were associated with accurate identification of stimuli as an adult. These features coincide with areas of the skull that reliably differentiate

the sexes in adult skulls, namely the glabella and the supraorbital ridge (Graw et al., 1999; Tanner & Tanner, 1990). Between puberty and adulthood, men develop a prominent supraorbital ridge and a larger glabella; in younger males, these features are not as apparent leading to uncertainty when classifying younger skulls as male or female (Graw et al., 1999). Our participants may therefore be using features reliably associated with adult male faces to differentiate child from adult faces in our sample.

In contrast, oblique features (i.e., specifically information at angles of 60° and 150°) were more indicative of a child's face, presumably from the rounder jawline and eyes of children's faces. In addition, we found that horizontal information—which is paramount for face identification and the categorization of facial expressions of emotions (Duncan et al., 2017; Goffaux et al., 2016; Goffaux & Dakin, 2010; Goffaux & Greenwood, 2016; Pachai et al., 2013)—was significantly associated with child categorisation. In fact, information for face-age categorisation peaked around this horizontal band (specifically at 75° when averaged across both child and adults). This finding is yet another piece of evidence showing that horizontal information is at the core of face processing tasks, and extends the importance of this information for the first time to the challenging task of face-age categorisation. Our findings also go a step further than previous psychophysical studies of face perception by showing that it is the horizontal information in the eye region that is the most relevant in the context of this face processing task. To our knowledge, this is the first time horizontal information is linked in a direct manner to eye information in a face processing task (e.g., see Duncan et al., 2017 for an indirect link).

In terms of the real-world implications of our findings, there appears to be a common set of facial features that is associated with children's faces, which may explain why differentiating between faces of different children can present such a challenge (e.g., the task facing border control agents; Bate et al., 2020; Belanova et al., 2018; Kramer et al., 2018;

Michalski et al., 2019). Our findings are also relevant to the task of identifying IIOC, whether by a human or an AI algorithm, and provides an evidence base for the training of digital forensics analysts (and applied professionals) who have to differentiate between faces of children from those of adults. It is reassuring that our findings coincide with some of the features described in accounts given by experts in this task as to the features to which they attend (Kloess et al., 2021). However, by incorporating our findings into training, this knowledge can be imparted to new recruits, ensuring that they are using the features empirically demonstrated to be associated with accurate differentiation from the start. Our findings are also relevant to discussions about verifying children online for safeguarding (GCHQ, 2020), and to technology companies who design age assurance algorithms that interface with a webcam (e.g., Yoti).

Finally, it is important to acknowledge the limitations of our study. The facial stimuli represented 12 individuals and all were of a White ethnicity. Three key implications stem from this: (i) our method needs replication with a larger set of faces; (ii) it should be extended to the study of faces of other ethnicities (Gao & Ai, 2009); and (iii) the interaction between the ethnicity of the viewer and the ethnicity of the stimuli should be evaluated. According to the Global Threat Assessment completed by the WeProtect Global Alliance (2019), children from lower income countries are at particular risk of sexual exploitation and abuse in order to provide family income given the high level of poverty in some of these areas. This also results in an increase of non-White children being depicted in IIOC, which the analysts in the Kloess et al. (2021) study reported as presenting an additional challenge due to the lack of familiarity with those of a different ethnicity to their own, and the fact that different ethnicities can represent great variation in terms of the stages and nature of physical development.

In summary, we used a comprehensive data-driven technique to reveal how three major types of facial information are used by human observers to make face-age categorisations. Notably, the sampling of all three visual features—position, spatial-frequency, orientation—simultaneously in the DFM procedure permitted us to interpret and integrate these findings in a single experiment/data set, relating spatial information (e.g., features of the jawline) with orientation information (e.g., oblique information). The features of the face our participants relied on to differentiate child from adult faces align with what is known about changes in facial physiology from childhood to adulthood, and with some of what digital forensics analysts have self-reported about the features to which they attend when trying to classify IIOC. However, our research goes beyond extending the relevance of these features to providing evidence that attending to these features (over others) is associated with *accurate* face-age classification. This not only advances our knowledge of the psychophysical correlates for face-age categorisation in human observers for the first time, but also constitutes a crucial step towards implementing evidence-based face-age classification training (see Adolphs et al., 2005; Faghel-Soubeyrand et al., 2019 for similar procedures).

References

- Adolphs, R., Gosselin, F., Buchanan, T. W., Tranel, D., Schyns, P., & Damasio, A. R. (2005). A mechanism for impaired fear recognition after amygdala damage. *Nature*, 433(7021), 68–72. <https://doi.org/10.1038/nature03086>
- Alink, A., & Charest, I. (2020). Clinically relevant autistic traits predict greater reliance on detail for image recognition. *Scientific Reports*, 10(1), 14239. <https://doi.org/10.1038/s41598-020-70953-8>
- Alliance, W. G. (2020). *Global Threat Assessment 2019: Working together to end the sexual exploitation of children online*. <https://apo.org.au/node/278176>
- Bate, S., Bennetts, R., Murray, E., & Portch, E. (2020). Enhanced matching of children’s faces in “super-recognisers” but not high-contact controls. *I-Perception*, 11(4), 2041669520944420. <https://doi.org/10.1177/2041669520944420>
- Belanova, E., Davis, J. P., & Thompson, T. (2018). Cognitive and neural markers of super-recognisers’ face processing superiority and enhanced cross-age effect. In *Cortex* (Vol. 108, pp. 92–111). <https://doi.org/10.1016/j.cortex.2018.07.008>
- Benitez-Quiroz, C. F., Srinivasan, R., & Martinez, A. M. (2018). Facial color is an efficient mechanism to visually transmit emotion. *Proceedings of the National Academy of Sciences of the United States of America*, 115(14), 3581–3586. <https://doi.org/10.1073/pnas.1716084115>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. <https://www.ncbi.nlm.nih.gov/pubmed/9176952>
- Duncan, J., Gosselin, F., Cobarro, C., Dugas, G., Blais, C., & Fiset, D. (2017). Orientations for the successful categorization of facial expressions and their link with facial features. *Journal of Vision*, 17(14), 7. <https://doi.org/10.1167/17.14.7>
- Dupuis-Roy, N., Faghel-Soubeyrand, S., & Gosselin, F. (2019). Time course of the use of

- chromatic and achromatic facial information for sex categorization. *Vision Research*, 157, 36–43. <https://doi.org/10.1016/j.visres.2018.08.004>
- Edwards, I. (2013). Victims, Sentencing Guidelines, and the Sentencing Council. In *Sentencing Guidelines* (pp. 71–85). <https://doi.org/10.1093/acprof:oso/9780199684571.003.0005>
- Estéphan, A., Fiset, D., Saumure, C., Plouffe-Demers, M.-P., Zhang, Y., Sun, D., & Blais, C. (2018). Time Course of Cultural Differences in Spatial Frequency Use for Face Identification. *Scientific Reports*, 8(1), 1816. <https://doi.org/10.1038/s41598-018-19971-1>
- Faghel-Soubeyrand, S., Dupuis-Roy, N., & Gosselin, F. (2019). Inducing the use of right eye enhances face-sex categorization performance. *Journal of Experimental Psychology. General*, 148(10), 1834–1841. <https://doi.org/10.1037/xge0000542>
- Faghel-Soubeyrand, S., Lecomte, T., Bravo, M. A., Lepage, M., Potvin, S., Abdel-Baki, A., Villeneuve, M., & Gosselin, F. (2020). Abnormal visual representations associated with confusion of perceived facial expression in schizophrenia with social anxiety disorder. *NPJ Schizophrenia*, 6(1), 28. <https://doi.org/10.1038/s41537-020-00116-1>
- Franqueira, V. N. L., Bryce, J., Al Mutawa, N., & Marrington, A. (2018). Investigation of Indecent Images of Children cases: Challenges and suggestions collected from the trenches. *Digital Investigation*, 24, 95–105. <https://doi.org/10.1016/j.diin.2017.11.002>
- Gao, F., & Ai, H. (2009). Face Age Classification on Consumer Images with Gabor Feature and Fuzzy LDA Method. *Advances in Biometrics*, 132–141. https://doi.org/10.1007/978-3-642-01793-3_14
- Goffaux, V., & Dakin, S. C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*, 1, 143. <https://doi.org/10.3389/fpsyg.2010.00143>

- Goffaux, V., Duecker, F., Hausfeld, L., Schiltz, C., & Goebel, R. (2016). Horizontal tuning for faces originates in high-level Fusiform Face Area. *Neuropsychologia*, *81*, 1–11.
<https://doi.org/10.1016/j.neuropsychologia.2015.12.004>
- Goffaux, V., & Greenwood, J. A. (2016). The orientation selectivity of face identification. *Scientific Reports*, *6*, 34204. <https://doi.org/10.1038/srep34204>
- Gosselin, F., & Schyns, P. G. (2001). Bubbles: a technique to reveal the use of information in recognition tasks. *Vision Research*, *41*(17), 2261–2271.
<https://www.ncbi.nlm.nih.gov/pubmed/11448718>
- Graw, M., Czarnetzki, A., & Haffner, H. T. (1999). The form of the supraorbital margin as a criterion in identification of sex from the skull: investigations based on modern human skulls. *American Journal of Physical Anthropology*, *108*(1), 91–96.
[https://doi.org/10.1002/\(SICI\)1096-8644\(199901\)108:1<91::AID-AJPA5>3.0.CO;2-X](https://doi.org/10.1002/(SICI)1096-8644(199901)108:1<91::AID-AJPA5>3.0.CO;2-X)
- Home Office. (2017, October 3). *Home Office to crack down on online child sexual abuse with new cutting-edge technology*. GOV.UK.
<https://www.gov.uk/government/news/home-office-to-crack-down-on-online-child-sexual-abuse-with-new-cutting-edge-technology>
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215–243.
<https://doi.org/10.1113/jphysiol.1968.sp008455>
- Kau, C. H., & Richmond, S. (2008). Three-dimensional analysis of facial morphology surface changes in untreated children from 12 to 14 years of age. *American Journal of Orthodontics and Dentofacial Orthopedics: Official Publication of the American Association of Orthodontists, Its Constituent Societies, and the American Board of Orthodontics*, *134*(6), 751–760. <https://doi.org/10.1016/j.ajodo.2007.01.037>
- Kloess, J. A., Woodhams, J., & Hamilton-Giachritsis, C. E. (2021). The challenges of

- identifying and classifying child sexual exploitation material: Moving towards a more ecologically valid pilot study with digital forensics analysts. *Child Abuse & Neglect* (Vol. 118, p. 105166). <https://doi.org/10.1016/j.chiabu.2021.105166>
- Kloess, J. A., Woodhams, J., Whittle, H., Grant, T., & Hamilton-Giachritsis, C. E. (2019). The Challenges of Identifying and Classifying Child Sexual Abuse Material. *Sexual Abuse* (Vol. 31, Issue 2, pp. 173–196). <https://doi.org/10.1177/1079063217724768>
- Komori, M., & Nittono, H. (2013). Influence of Age-independent Facial Traits on Adult Judgments of Cuteness and Infantility of a Child's Face. *Procedia - Social and Behavioral Sciences*, 97, 285–291. <https://doi.org/10.1016/j.sbspro.2013.10.235>
- Kramer, R. S. S., Mulgrew, J., & Reynolds, M. G. (2018). Unfamiliar face matching with photographs of infants and children. *PeerJ*, 6, e5010. <https://doi.org/10.7717/peerj.5010>
- Krasodonski-Jones, A., White, C., Eccleston, D., & Marsh, O. (2017). *Online child sexual abuse imagery*. <https://dera.ioe.ac.uk/id/eprint/30926>
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S. T., & van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition and Emotion*, 24(8), 1377–1388. <https://doi.org/10.1080/02699930903485076>
- Liggett, J. (1974). *The human face*. Stein and Day.
- Merdian, H. L., Perkins, D., McCashin, D., & Stevanovic, J. (2021). Integrating structured individual offending pathway analysis into group treatment for individuals who have accessed, shared, and/or distributed child sexual exploitation material: a feasibility study and preliminary outcome evaluation. In *Psychology, Crime & Law* (Vol. 27, Issue 6, pp. 579–605). <https://doi.org/10.1080/1068316x.2020.1849690>
- Michalski, D., Heyer, R., & Semmler, C. (2019). The performance of practitioners conducting facial comparisons on images of children across age. *PloS One*, 14(11), e0225298. <https://doi.org/10.1371/journal.pone.0225298>

- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013). Sensitivity to Information Conveyed by Horizontal Contours is Correlated with Face Identification Accuracy. *Frontiers in Psychology*, 4, 74. <https://doi.org/10.3389/fpsyg.2013.00074>
- Ratnayake, M., Obertová, Z., Dose, M., Gabriel, P., Bröker, H. M., Brauckmann, M., Barkus, A., Rizgeliene, R., Tutkuvienė, J., Ritz-Timme, S., Marasciuolo, L., Gibelli, D., & Cattaneo, C. (2014). The juvenile face as a suitable age indicator in child pornography cases: a pilot study on the reliability of automated and visual estimation approaches. *International Journal of Legal Medicine*, 128(5), 803–808. <https://doi.org/10.1007/s00414-013-0875-y>
- Ricanek, K., Bhardwaj, S., & Sodomsky, M. (2015). A review of face recognition against longitudinal child faces. *BIOSIG 2015*. <https://dl.gi.de/handle/20.500.12116/2272>
- Royer, J., Blais, C., Gosselin, F., Duncan, J., & Fiset, D. (2015). When less is more: Impact of face processing ability on recognition of visually degraded faces. *Journal of Experimental Psychology. Human Perception and Performance*, 41(5), 1179–1183. <https://doi.org/10.1037/xhp0000095>
- Sae-Bae, N., Sun, X., Sencar, H. T., & Memon, N. D. (2014). Towards automatic detection of child pornography. *2014 IEEE International Conference on Image Processing (ICIP)*, 5332–5336. <https://doi.org/10.1109/ICIP.2014.7026079>
- Sforza, C., Grandi, G., De Menezes, M., Tartaglia, G. M., & Ferrario, V. F. (2011). Age- and sex-related changes in the normal human external nose. In *Forensic Science International* (Vol. 204, Issues 1-3, pp. 205.e1–e205.e9). <https://doi.org/10.1016/j.forsciint.2010.07.027>
- Tanner, J. M., & Tanner, J. M. (1990). *Foetus Into Man: Physical Growth from Conception to Maturity*. Harvard University Press. <https://play.google.com/store/books/details?id=YxpimctaWd4C>

- Tardif, J., Fiset, D., Zhang, Y., Estéphan, A., Cai, Q., Luo, C., Sun, D., Gosselin, F., & Blais, C. (2017). Culture shapes spatial frequency tuning for face identification. *Journal of Experimental Psychology. Human Perception and Performance*, 43(2), 294–306.
<https://doi.org/10.1037/xhp0000288>
- Tardif, J., Morin Duchesne, X., Cohan, S., Royer, J., Blais, C., Fiset, D., Duchaine, B., & Gosselin, F. (2019). Use of face information varies systematically from developmental prosopagnosics to super-recognizers. *Psychological Science*, 30(2), 300–308.
<https://journals.sagepub.com/doi/abs/10.1177/0956797618811338>
- Watson, A. B., & Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Perception & Psychophysics*, 33(2), 113–120. <https://doi.org/10.3758/bf03202828>
- Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: the SHINE toolbox. *Behavior Research Methods*, 42(3), 671–684. <https://doi.org/10.3758/BRM.42.3.671>
- Young, S. G., Elliot, A. J., Feltman, R., & Ambady, N. (2013). Red enhances the processing of facial expressions of anger. *Emotion*, 13(3), 380–384.
<https://doi.org/10.1037/a0032471>

supplementary figure 1. Features for accurate classification of child vs adult faces, irrespective of the class.

Correlation between accurate responses and position (A), spatial-frequencies (B) and orientations (C) have been averaged irrespective of the stimulus class. Significance was assessed using permutation testing : a null distribution was created by regressing the presence of each feature type across trials with randomly shuffled the accuracy vector of each participant (repeated 1000 times).