

Kent Academic Repository

Full text document (pdf)

Citation for published version

Demanet, Jelle and Dhont, Kristof and Notebaert, Lies and Pattyn, Sven and Vandierendonck, Andre (2007) Pixelating Familiar People in the Media: Should Masking Be Taken at Face Value? *Psychologica Belgica*, 47 (4). pp. 261-276. ISSN 0033-2879.

DOI

<https://doi.org/10.5334/pb-47-4-261>

Link to record in KAR

<https://kar.kent.ac.uk/42831/>

Document Version

UNSPECIFIED

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

PIXELATING FAMILIAR PEOPLE IN THE MEDIA: SHOULD MASKING BE TAKEN AT FACE VALUE?

Jelle DEMANET, Kristof DHONT, Lies NOTEBAERT, Sven PATTYN, &
André VANDIERENDONCK
Ghent University

This study questions the effectiveness of masking faces by means of pixelation on television or in newspapers. Previous studies have shown that masking just the face leads to unacceptably high recognition levels, making it likely that participants also use other cues for recognition, such as hairstyle or clothes. In the current study we investigate this possibility by means of an identification task in which participants had to identify (partially) masked images of familiar people. To demonstrate that non-facial cues become increasingly important for recognition as faces are masked more strongly, we manipulated the size of the masked area and the degree of pixelation. Confirming our expectations, increasing the size of masked area or its level of deterioration led to lower recognition rates. More importantly, also an interaction effect between the two variables emerged, showing that additional visual information partly compensates the downswing in recognition when masking becomes stronger. Although in some conditions low recognition rates were found, masking was never a hundred percent effective, making it clear that the media should approach this issue with care. Implications of our findings and future directions are considered.

When we see the face of a person we are familiar with, it is easy to identify that person. Variations in luminance, face mimics, and even special characteristics such as beards, hoods and glasses do not cause many problems for identification. Moreover, familiar faces can be recognised even when the quality of the presented images is very poor, which is, for instance, often the case for security surveillance videos (Bruce, Henderson, Newman, & Burton, 2001; Burton, Wilson, Cowan, & Bruce, 1999). Apparently, we are doing very well at recognising people by their face.

But what about masked faces? Are we capable of recognising a familiar person in a newspaper or on television when some sort of masking technique

The first four authors have contributed equally to the manuscript. Jelle Demanet, Kristof Dhont, Lies Notebaert, Sven Pattyn, and André Vandierendonck are affiliated to Ghent University; Jelle Demanet and André Vandierendonck to the Department of Experimental Psychology; Kristof Dhont and Sven Pattyn to the Department of Developmental, Personality, and Social Psychology, and Lies Notebaert to the Department of Experimental Clinical and Health Psychology.

Correspondence concerning this article should be addressed to Jelle Demanet, Department of Experimental Psychology, Ghent University, Henri Dunantlaan 2, 9000, Ghent. E-mail: Jelle.Demanet@UGent.be or to Kristof Dhont, Department of Developmental, Personality and Social Psychology, Ghent University, Henri Dunantlaan 2, 9000, Ghent. E-mail: Kristof.Dhont@UGent.be

is applied? Given the fact that vague pictures of familiar people lead to high levels of identification (see Bruce et al., 2001), it is very relevant to question the effectiveness of deliberate attempts to make faces unrecognisable. Therefore, the goal of the present study is to search for conditions that can serve as guidelines for everyday journalistic practice. Focusing on the technique of pixelation, we want to examine which areas of a displayed person should be masked and how severely the quality of the masked area should be deteriorated to truly render the depicted person unrecognisable. These are important questions, since there is no consensus on this subject today.

Masking in the media

Although journalists may appeal on the freedom of press and their professional duty to inform the public, under certain circumstances several laws prohibit them to portray people in a recognisable manner. Indeed, national laws protecting the right on anonymity or the portrait right, a recent recommendation of the Committee of Ministers of the Council of Europe (2003), and deontological codes for journalists all describe situations in which the freedom of press is subordinate to the private rights of the depicted person.

Therefore, numerous techniques are used in the media to obscure the identity of a person. Examples can be found daily in newspapers and on television: a black rectangle is added in front of a person's eyes, faces are blurred or pixelated, or only the back of a person is shown. However, journalists often rely on their own, subjective impression to decide whether a person is unrecognisable; a procedure that might be problematic as journalists are usually not familiar with the depicted person. Therefore, they are not in the position to assess whether people who are familiar with the person are unable to recognise him or her as well.

The efficacy of masking techniques

The majority of research on masking has focused on the techniques of blurring and pixelation (also known as 'coarse quantisation' or 'blocking'). The blurring technique removes details from an image by transforming each pixel into a weighted average of itself and its surrounding pixels. Thereby some of the high spatial frequencies of the image, which are known to aid recognition (see Fiorentini, Maffei, & Sandini, 1983), are removed. The weighting procedure that is used most often is based on a Gaussian distribution (i.e., Gaussian blurring).

The pixelation technique, on the other hand, decreases the number of pixels in (a part of) an image, by replacing the original, smaller pixels with a limited number of relatively large pixels. The luminance of each new pixel is

equal to the average luminance of the original pixels that it replaces. This technique removes high spatial frequencies from the image and, at the same time, introduces spurious frequency noise.

Harmon and Julesz (1973) were the first to demonstrate that pixelation reduces people's ability to recognise an image. Ever since, this finding has been replicated and extended by many others (e.g., Bachmann, 1991; Costen, Parker, & Craw, 1994, 1996; Lander, Bruce, & Hill, 2001; Morrone, Burr, & Ross, 1983). Bachmann (1991) for instance, found evidence for a critical amount of pixels that is sufficient for recognising a face (i.e., 18 pixels horizontally and vertically; see also Costen et al., 1994, 1996).

From a practical point of view however, it is not yet clear whether these findings can be applied in everyday journalistic activities. As recently noted by Lander et al. (2001), the stimuli used by Bachmann (1991) contained pictures that only showed the face area, i.e., without the hair or any other body part. Moreover, the depicted people did not have any special characteristics like beards or glasses. Although the author had good reasons to do this in light of his research goals, it undermines the applicability of his findings in more realistic settings.

To clarify this issue, Lander et al. (2001) conducted two experiments to investigate the effectiveness of pixelation (Experiment 1) and blurring (Experiment 2). In both experiments, participants saw television images of familiar (famous) and unfamiliar people who had to be identified. These images showed at least the head and shoulders of the portrayed people, all from a frontal viewpoint. In their pixelation experiment, masking was achieved by reducing the resolution of the appropriate face area to 10 or 20 pixels width-wise. The second experiment was almost identical, except that the faces were now blurred instead of pixelated.

The study of Lander et al. (2001) revealed two interesting findings concerning the present research goals. First, the results indicated that the overall performance level of identification was considerably high. For instance, in the first experiment, participants identified on average 65% of the familiar people. Second, as the authors expected, identification significantly ameliorated as the images were less deteriorated. For instance, with 10 pixels/face the mean identification rate was 55% whereas with 20 pixels/face this was 78%. Both findings clearly show that despite the degraded quality of the images, participants were still able to recognise a considerably high percentage of the familiar faces. Hence, as suggested by Lander et al. (2001, p. 109), "the higher recognition levels shown in this experiment are likely to be a reflection of the additional cues to identity available in both our moving and static images." Indeed, some of the people in the pictures wore glasses or had a remarkable haircut which could have facilitated identification.

Recognition cues

Past research on familiar face recognition elucidated that internal facial features, such as mouth, nose, and eyes, are more important face recognition cues than external features like the hair or the chin (Ellis, Shepherd, & Davies, 1979). Ellis et al. (1979) suggested that the construction of cognitive representations of faces is based on the internal features, which are relatively stable over time and more informative. Several other studies corroborated the internal cue superiority and have led to the hypothesis that during familiar face recognition people rely on the least changeable characteristics (Bruce & Young, 1998). For instance, Young, Hay, McWeeny, Flude, and Ellis (1985) found that participants were faster in matching pictures of complete faces with pictures of only the inner-face parts than with pictures of only the outer-face parts. However, concerning the recognition of unfamiliar faces, it appears that outer-face parts, compared to inner-face parts, are equally important (Ellis et al., 1979; Young et al., 1985) or even more important (Bruce, Henderson, Greenwood, Hancock, Burton, & Miller, 1999; Shepherd, Davies, & Ellis, 1981) cues for recognition.

Based on these results, we can conclude that people rely heavily on internal features for familiar face recognition. However, because these studies only presented stimuli consisting of specific facial parts, the question remains whether these internal features are equally important in a situation in which the masked images show the complete person. Clearly, Lander et al. (2001) already provided a provisory answer to this question. Masking the face, and as such also the internal features, appears to be necessary, but insufficient. Even when seriously degraded, participants still recognised at least 35% of the static familiar faces. Furthermore, Lander et al. (2001) suggested that the recognition rate of the pictures of low quality was still considerably high because in this condition participants were also relying on specific features that were not part of the inner-face, like the haircut or glasses. In sum, the effectiveness of masking is at least disputable and raises the question of how such masking techniques can be applied more successfully.

The present study

The current work focuses on the technique of pixelation. As mentioned earlier, compared to blurring, pixelation not only removes high spatial frequencies from an image, but also introduces spurious frequency noise. Thus, from a practical point of view, by applying the pixelation technique we are more likely to succeed in finding those conditions that can truly prevent identification (see Costen et al., 1994, 1996).

Because our main goal concerns the practical implications of applying

masking techniques, we are not primarily interested in the individual contribution of specific external and internal cues to person identification. Instead, we try to evaluate the pixelation technique as it is commonly used in the media. More concretely, when people are pixelated in the media, only the face itself is masked consistently, while hair, clothes, and body are masked less often.

Therefore, the current study examines the role of the size of the masked surface area. Thereby, we can inspect whether masking external cues, such as the hair and body, on top of masking the face results in lower recognition rates. The level of masking is also manipulated by using 3 levels of pixelation, namely 8, 14, and 20 pixels/face (measured horizontally from ear to ear). By crossing both variables in a within-subject design, we want to determine how recognisable familiar people are when different degrees of pixelation are applied to different sizes of masked surface.

Consistent with previous research, we expect that an increase of the size of the masked area as well as its level of deterioration will lead to lower recognition rates (see e.g., Bachmann, 1991; Lander et al., 2001). More specifically, while very high recognition rates are expected for a masking of 20 pixels/face, a masking of 14 and 8 pixels/face should lead to a decrease in recognition. Analogously, a decrease in recognition rate is expected as the masked area expands from covering only the face to covering the complete head and body of a person.

More importantly, we predict a significant interaction between the size of the masked area and level of pixelation. Specifically, we expect a negligible effect of the size of the masked area for a masking of 20 pixels/face, while for a masking of 14 and 8 pixels/face, significant lower recognition rates are expected for a larger masked area. This effect should be more apparent when pixelation is coarser (8 pixels/face compared to 14 pixels/face). This hypothesis stems from the finding that external cues (e.g., movement; Lander et al., 2001) have already proven to interact with level of pixelation. Hence, we expect a similar interaction with masked area which confines the accessibility to cues such as hairstyle or clothing.

Finally, we want to stress that a masking technique is not very efficient if, in spite of a drastic decrease in recognition, more than half of the stimuli can still be identified. Therefore, we explicitly want to point out the necessity of considering the absolute values of the recognition rates and thereby predict that both a high degree of deterioration and a rather large masked area will prove indispensable when sufficiently low recognition rates are desired.

Method

Participants

Thirty-seven volunteers participated in the experiment, of which 16 were male and 21 female. They were aged between 18 and 53 years (mean age 23 years) and had normal or corrected-to-normal vision.

Stimuli, materials and design

A set of 180 digital colour photos was selected comprising 90 familiar (famous) and 90 unknown people. All pictures displayed the bust of the people from a frontal viewpoint on a white background (see Figure 1). The pictures were presented on a 15 inch colour monitor, connected to an IBM compatible personal computer. Viewing distance was 60 cm and the depicted faces were approximately 4 cm wide and presented centrally on the screen.

The familiar people were selected on the basis of their presumed familiarity. The set of 90 stimuli comprised four equally large categories; politicians, musicians, Hollywood actors, and a residual category consisting of sportsmen and -women and local television personalities. As in Lander et al. (2001), pictures of unknown people were added to see how often participants would mistakenly guess the identity of an unfamiliar person. These pictures were chosen to approximately reflect the general appearance of the familiar people. All pictures were found on the Internet and were of high quality. The entire set of both familiar and unfamiliar people consisted of an equal amount of males and females.

The identity of the people was masked by applying the pixelation technique (see Figure 1). Moreover, we manipulated both the size of the masked area and the severity of masking. The masked area contained either the face area, the area covering both face and hair, or the area containing the entire person (i.e., the bust). The degree of pixelation was measured by counting the number of pixels in width from ear to ear (see Lander et al., 2001). Three levels of pixelation were applied: 8 pixels/face, 14 pixels/face, and 20 pixels/face.

The two factors, Masked Area and Level of Pixelation, were manipulated within-subjects and crossed. Both factors had three levels (i.e., face, head, or bust for Masked Area and 8 pixels/face, 14 pixels/face, or 20 pixels/face for Level of Pixelation). Figure 1 shows a particular familiar person in each of the nine possible masking conditions.

During the experiment, each participant was shown all 180 familiar and unfamiliar people in one of the nine possible masking conditions. Because each stimulus was presented only once, our data are not affected by any rep-




Figure 1.

All nine possible masking conditions based on the original picture of 1 of the 90 familiar stimuli, namely Leonardo DiCaprio. Level of Pixelation varied from left (20 pixels/face) to middle (14 pixels/face) to right (8 pixels/face). The levels of Masked Area varied from top (*face*) to middle (*head*) to bottom (*bust*).

etition effects (i.e., identification facilitation because of former successful identifications of the same stimulus). Furthermore, the experiment was set up as to ensure that each participant was shown 20 stimuli (10 familiar and 10 unfamiliar) in each masking condition. For this purpose, nine lists were created in which both the order of the stimuli and the masking conditions were randomized. This randomization however, had both a within-list restriction resulting in a design similar to a Latin Square and a between-list restriction. The within-list restriction guaranteed that each masking condition occurred equally often in each list (i.e., 10 familiar and 10 unfamiliar stimuli per condition per list). The between-list restriction assured that each masking condition of each picture was presented once across the nine lists. Hence, it was ensured that each participant provided enough data for each masking condition, given a sufficient amount of controlled randomization.

Procedure

Each participant was randomly assigned to one of the nine experimental lists and was tested individually. They were asked to verbally identify the portrayed person. As in Lander et al. (2001), both the name of the person and unambiguous semantic information about the person (e.g., the name of an actor's character or ministerial function) were accepted as correct answers. General information (e.g., 'actor' or 'politician') was considered insufficient for a correct response.

The pictures were shown for three seconds, after which a question appeared on screen asking whether the participant had recognised the person. Participants had six seconds to respond before the next picture appeared, which proved amply sufficient. The experimenter wrote down every answer but did not give any feedback. Afterwards participants had to identify all supposedly familiar people once again in original, unmasked form to control whether participants really were familiar with these people.

Data analysis

For the purpose of data-analysis, the raw data were transformed into measures based on Signal Detection Theory (SDT). First, mean Hit and False Alarm Rates were calculated for every participant for each of the nine experimental conditions. The Hit Rates were calculated by dividing the number of correct recognitions of familiar stimuli by the total number of known stimuli for that participant in that condition and the False Alarm Rates by dividing the number of incorrect identifications of unfamiliar stimuli by the total number of unknown stimuli.

Next, the signal detection measures d' and c were calculated to determine

whether variations in the Hit Rate should be addressed to variations in sensitivity or response bias, respectively¹. The discrimination or sensitivity index d' measures the extent to which signal trials (i.e., familiar stimuli) can be distinguished from noise trials (i.e., unfamiliar stimuli). More specifically, the value of the sensitivity parameter d' denotes the distance between the distributions of the signal and noise trials and is scaled by the standard deviation of these distributions. A d' -value of 0 indicates that the participant is unable to distinct signal from noise trials. Positive values point to an increase in participant's ability to make such distinction, whereas negative values are generally caused by sampling errors and mistakes (e.g., when a participant repeatedly says 'yes' while he meant to say 'no' in a yes/no task). For the current experiment we expect higher d' -values in those conditions in which masking was less efficient.

The parameter c is a response bias measure. For participants who tend to respond positively on both signal and noise trials (e.g., by wrongfully identifying unfamiliar people), the value of this parameter is negative (i.e., a liberal criterion). Positive estimates of this parameter point to a tendency to respond negatively (i.e., a conservative criterion). Values close to 0 indicate that participants applied a neutral criterion.

Because d' can in principle adopt values ranging from $-\infty$ to $+\infty$, the log-linear approach was applied to handle extreme Hit and False Alarm Rates. This approach consists of adding 0.5 to each number of Hits and False Alarms and adding 1 to the total number of signal and noise trials. Based on these modifications, new estimates can be acquired for both the Hit Rate and False Alarm Rate, which now cannot take on the extreme values of 0 or 1. Hence, also the parameters d' and c can be easily calculated without any occurrence of extreme values (Hautus, 1995).

Results

Data of one participant were left out of all further analyses, because the familiarity check revealed that for some masking conditions, the participant was unfamiliar with 7 or 8 of the 10 familiar people in those conditions. As a consequence, each of the experimental lists was completed by exactly four

¹These parametric measures were preferred to their non-parametric equivalents 'A' en 'B'' because both types of measures assume some form of underlying distribution for the Hits and False Alarms anyhow and because the parametric measures are generally preferable in most contexts (Pastore, Crawley, Berens, & Skelly, 2003; for an overview on SDT, see Stanislaw & Todorov, 1999; for in depth accounts on different bias or discrimination indices, see MacMillan & Creelman, 1990; Swets, 1986, respectively).

participants. The 36 participants recognised on average 77 of the 90 familiar persons. Data of familiar persons which were not recognised by a participant in unmasked form were left out of the analysis for that participant. The mean overall Hit and False Alarm Rate was 55.4% and 7.8%, respectively.

Sensitivity parameter d'

A 3 (Masked Area) \times 3 (Level of Pixelation) Repeated Measures Multivariate Analysis of Variance (MANOVA) was conducted on the sensitivity parameter d'^2 . The main effects of Masked Area and Level of Pixelation both proved significant, respectively, $F(2, 34) = 13.06, p < .001$ and $F(2, 34) = 128.56, p < .001$, showing that the more a picture was deteriorated by either of both variables, the more effective masking was. Furthermore, also the interaction effect between both variables was significant; $F(4, 32) = 5.00, p < .005$. Figure 2 shows the mean value and standard deviation of d' for each experimental condition.

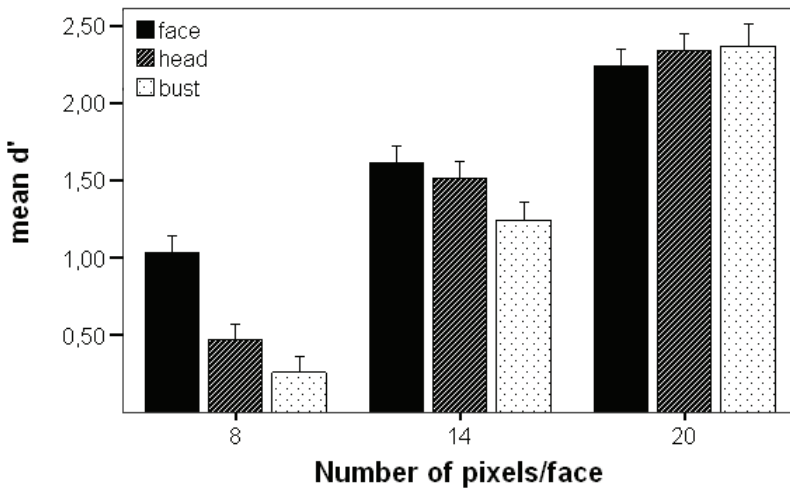


Figure 2.

The mean and standard error of d' for each experimental condition based on 'Level of Pixelation' and 'Masked Area'.

²Generally, when repeated measures ANOVA is used, the violation of the sphericity assumption makes the regular, univariate F -values inaccurate. In the present study for example, the analysis of d' Level of Pixelation failed Mauchly's test of sphericity; approximate $\chi^2(2) = 10.08, p < .01$. Several approaches have been developed that correct for the inaccuracy of such univariate test results (e.g., Greenhouse-Geisser, Huynh-Feldt). However, O'Brien and Kaiser (1985) stated that a multivariate approach (i.e., Repeated Measures MANOVA), which does not assume sphericity, is more appropriate for analysing such data, even though such tests are less powerful.

Figure 2 also gives a hint as to how to interpret the interaction effect between Masked Area and Level of Pixelation. Apparently, when masking is rather superficial (i.e., 20 pixels/face), there appears to be no difference between the different levels of Masked Area, but as masking gets more severe, a difference does appear between those levels. This trend corresponds to our predictions and was statistically tested by executing six planned comparisons; one for each pair of two consecutive levels of Masked Area in each level of Level of Pixelation separately. As expected, the levels 'face' and 'head' and the levels 'head' and 'bust' did not differ significantly when the masking consisted of 20 pixels/face; F 's < 1 . For 14 pixels/face, the levels 'face' and 'head' did not differ significantly; $F < 1$, but the levels 'head' and 'bust' did; $F(1, 35) = 5.55, p < .05$. Finally, for the highest level of pixelation (i.e., 8 pixels/face), both the levels 'face' and 'head' and the levels 'head' and 'bust' differed significantly; respectively, $F(1, 35) = 14.75, p < .001$ and $F(1, 35) = 4.56, p < .05$.

Finally, to verify whether these results are common for all stimuli and are not caused by only a small number of pictures, an additional analysis was conducted across the familiar stimuli instead of across participants. Familiar faces which were not known by at least two of four participants in each experimental condition were left out of this analysis. This way, 61 of the 90 original familiar stimuli entered the analysis. Similar to the analysis across participants, a 3 x 3 Repeated Measures MANOVA was conducted for the Hit Rates³. The results completely paralleled those of our previous analyses, showing a main effect of Masked Area; $F(2, 59) = 15.74, p < .001$, a main effect of Level of Pixelation; $F(2, 59) = 262.07, p < .001$, and an interaction effect between both variables; $F(4, 57) = 11.98, p < .001$.

Bias parameter c

A 3 (Masked Area) x 3 (Level of Pixelation) Repeated Measures MANOVA was conducted on the bias parameter c . The main effects of Masked Area and Level of Pixelation, as well as their interaction were significant; $F(2, 34) = 20.02, p < .001$, $F(2, 34) = 144.24, p < .001$, and $F(4, 32) = 8.67, p < .001$, respectively. Figure 3 shows the mean value of c for each experimental condition. The positive values of c in each of the experimental conditions point to a rather conservative disposition of the participants.

In line with the analyses of d' , 6 planned comparisons were conducted to further investigate the interaction effect. This analysis revealed that at a pixelation level of 20 pixels/face, the levels 'face' and 'head' and the levels 'head' and 'bust' did not differ significantly; $F(1, 35) = 3.04, ns$, and $F < 1$,

³ d' measures could not be calculated because there were no noise trials for this analysis.

respectively. For 14 pixels/face, the levels 'face' and 'head' and the levels 'head' and 'bust' did not differ significantly either; $F(1, 35) = 3.04$, *ns*, and $F < 1$, respectively. For the most severe level of pixelation (i.e., 8 pixels/face), the levels 'face' and 'head' did differ significantly; $F(1, 35) = 44.37$, $p < .001$, while the levels 'head' and 'bust' did not; $F < 1$.

A comparison of Figures 2 and 3 indicates that the bias is inversely related to the signal detection measure d' . This was statistically confirmed since, on the mean condition level, the c parameter correlated almost perfectly with d' , $r = -.96$, $p < .001$. This means that the lower d' was, the higher the bias was to deny recognition of the stimulus.

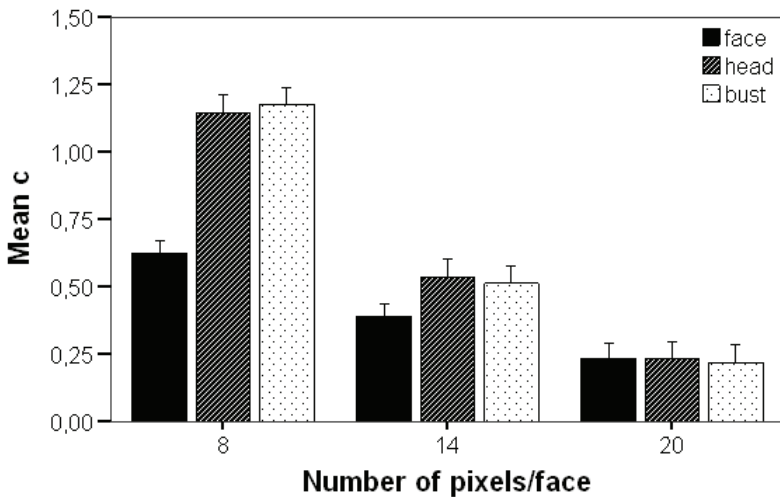


Figure 3.

The mean and standard error of the bias parameter c for each experimental condition based on 'Level of Pixelation' and 'Masked Area'.

Discussion

The media often have to obscure a person's identity on legal grounds and try to achieve this by applying masking techniques. However, since journalists determine in a subjective manner to which degree a presented person is still recognisable, it is not at all clear whether they truly succeed in masking the identity of that person. In the current study, we investigated the effectiveness of such masking techniques. In order to be effective, a masked person should be truly unidentifiable, i.e., the lower the identification rate, the better. The first aim of this study was to examine the impact of both the size of the masked area as well as the degree of pixelation on identification and

whether the effect of external cues, such as hair, shoulders, and clothes varies across different degrees of pixelation. The second aim was to find a masking condition in which recognition is sufficiently hampered to be of any practical use.

Regarding our first aim, we found a general trend concerning the level of pixelation. Indeed, supporting evidence was found for the intuitive proposition that the chance to identify a person decreases as pixelation becomes more severe, as was already reported previously (e.g., Bachmann, 1991; Costen et al., 1994, 1996; Lander et al., 2001; Morrone et al., 1983). Besides the effect of degree of pixelation, we also found a significant main effect of the size of the masked area; the larger the masked area, the more difficult it was to recognise the depicted person.

However, as expected, interpretation of these effects should take into account the significant interaction between both variables. More specifically, when a picture was masked slightly (20 pixels/face), the size of the masked area had no influence. For more deteriorated pictures (14 and 8 pixels/face) recognition rates between the different sizes of masked area did differ. This kind of interaction seems to indicate that the more severe the masking the more people appeal to other recognition cues.

For the current study, this boils down to the fact that people make use of cues such as hairstyle and clothing to recognise moderately (14 pixels/face) and severely (8 pixels/face) masked individuals. In fact, because our study was conducted under laboratory conditions, *it might even underestimate* the potential use of bodily features as recognition cues in a more realistic setting.

First, only the bust of the stimuli was depicted while in everyday media people are often depicted head-to-toe, thus increasing the potential relevance of bodily characteristics as recognition cues. Second, the familiar stimuli in the current study only consisted of celebrities. It is well possible that clothing is more distinctive for personally familiar persons than it is for celebrities.

Regarding the practical implications of this study, we hypothesised that both a large masked area and a high level of pixelation of this masked area are needed to reach reasonably low recognition rates. It is clear that masking with 14 and 20 pixels/face did not sufficiently limit recognition. The d' values for these conditions indicated that the distance between the signal and noise distributions varied between 1.25 and 2.37 times the standard deviation, meaning that participants were well able to distinguish familiar from unfamiliar stimuli. The question whether a masking of 8 pixels/face is sufficient, depends on which standard is used. When only complete unrecognisability is an acceptable goal (i.e., a d' -value of 0), then this condition will not satisfy either. If a balance should be reached between the unrecognisability of the person and the quality of the picture, then a condition in which d' -values are lower than 0.5 may be acceptable. However, whatever standard is

used, our study made clear that simply masking the face is insufficient, because, no matter how severely the person was masked, d' was always larger than 1, with a Hit Rate of at least 45%.

Finally, we found significant effects concerning the bias index. The positive values of c in our study point to a tendency for participants to reply negatively when asked whether they recognised the depicted person. Furthermore, the parameter c varied significantly over conditions. For the interpretation of these effects, the corresponding values of the recognition measures can be taken into account. The highly negative correlations between c and d' indicate that the tendency to respond negatively increased as the ability to identify the depicted person decreased. The fact that participants were inclined to deny knowing the presented stimulus when a large part of the picture is severely masked, is of course a very intuitive finding.

Although the current study can prove very useful in some contexts, we do want to highlight some restrictions of our design and give several directions for future research. A first restriction concerns our use of static images as stimuli. Indeed, it is a well established fact that movement has an impact on person identification (e.g., Knappmeyer, Thornton, & Bühlhoff, 2003; Knight & Johnston, 1997; Lander et al., 2001; Lander, Christie, & Bruce, 1999; Thornton & Kourtzi, 2002). Lander et al. (2001), for example, showed that the addition of natural movement leads to a general increase in the ability to identify masked people. It would be interesting to investigate whether movement ameliorates recognition equally under all circumstances or whether it interacts with variables such as the size of the masked area or the level of pixelation. This question has important practical implications, since in audiovisual media people are usually moving.

A second restriction has already been reported above. Because our stimuli only portrayed the bust of the depicted persons, the actual usefulness of bodily cues such as the clothing or body shape might be wrongfully minimised. Presenting pictures that show the target person head-to-toe, instead of only showing the bust, may clarify this issue. Moreover, presenting the complete person truly addresses the question of the efficacy of masking on person identification, whereas most previous research focused on the effect of masking on face recognition.

Finally, the effect of external, non-visual cues, such as the voice can be explored further. The voice could be gradually transformed, as is often done in the media, to examine its effect on recognition in addition to the manipulations of the current study.

Conclusion

Although the media often struggle to find the right balance between the quality of an image and the necessity of masking someone's identity, there appear to be no guidelines at all as to how to apply such masking techniques successfully. We realise that we, as scientists, are not in the position to impose standards on the media. Nevertheless, we tried to clarify what conditions must be fulfilled to attain sufficiently low recognition rates. From our research, we can conclude that it is advisable for the media to take caution when masking people who do not want to be recognised. We showed that masking just the face leads to an unacceptably high degree of recognition, independent of which level of pixelation was used. Under the current conditions (i.e., presenting only the bust of the person), it seems that especially the hair needs to be masked, although the clothing and body also appear to play a role in recognition. In everyday media then, other cues (e.g., 'background story', complete outfit, movement, or voice) are often available as well and might facilitate identification even more. Therefore, we advise a total and severe masking in order to make people truly unrecognisable.

References

- Bachmann, T. (1991). Identification of spatially quantised tachistoscopic images of faces: How many pixels does it take to carry identity? *European Journal of Cognitive Psychology*, *3*, 85-103.
- Bruce, V., Henderson, V., Greenwood, K., Hancock, P.J.B., Burton, A.M., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, *5*, 339-360.
- Bruce, V., Henderson, V., Newman, C., & Burton A.M. (2001). Matching identities of familiar and unfamiliar faces caught on CCTV images. *Journal of Experimental Psychology: Applied*, *7*, 207-218.
- Bruce, V., & Young, A. (1998). *In the eye of the beholder: The science of face perception*. Oxford: Oxford University Press.
- Burton, A.M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science*, *10*, 243-248.
- Costen, N.P., Parker, D.M., & Craw, I. (1994). Spatial content and spatial quantization effects in face recognition. *Perception*, *23*, 129-146.
- Costen, N.P., Parker, D.M., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception and Psychophysics*, *58*, 602-612.
- Ellis, H.D., Shepherd, J.W., & Davies, G.M. (1979). Identification of familiar and unfamiliar faces from internal and external features – Some implications for theories of face recognition. *Perception*, *8*, 431-439.
- Fiorentini, A., Maffei, L., & Sandini, G. (1983). The role of high spatial frequencies

- in face perception. *Perception*, *12*, 195-201.
- Harmon, L.D., & Julesz, B. (1973). Masking in visual recognition – Effects of 2-dimensional filtered noise. *Science*, *180*, 1194-1197.
- Hautus, M. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d' . *Behavior Research Methods, Instruments, & Computers*, *27*, 46-51.
- Knappmeyer, B., Thornton, I.M., & Bülhoff, H.H. (2003). The use of facial motion and facial form during the processing of identity. *Vision Research*, *43*, 1921-1936.
- Knight, B., & Johnston, A. (1997). The role of movement in face recognition. *Visual Cognition*, *4*, 265-273.
- Lander, K., Bruce, V., & Hill, H. (2001). Evaluating the effectiveness of pixelation and blurring on masking the identity of familiar faces. *Applied Cognitive Psychology*, *15*, 101-116.
- Lander, K., Christie, F., & Bruce, V. (1999). The role of movement in the recognition of famous faces. *Memory and Cognition*, *27*, 974-985.
- MacMillan, N.A., & Creelman, C.D. (1990). Response bias: Characteristics of detection theory, threshold theory, and “nonparametric” indexes. *Psychological Bulletin*, *107*, 401-413.
- Morrone, M.C., Burr, D.C., & Ross, J. (1983). Added noise restores the recognizability of coarse quantized images. *Nature*, *305*, 226-228.
- O'Brien, R.G., & Kaiser, M.K. (1985). MANOVA method for analyzing repeated measures designs: An extensive primer. *Psychological Bulletin*, *97*, 316-333.
- Pastore, R.E., Crawley, E.J., Berens, M.S., & Skelly, M.A. (2003). “Nonparametric” A' and other modern misconceptions about signal detection theory. *Psychonomic Bulletin and Review*, *10*, 556-569.
- Recommendation Rec(2003)13 of the Committee of Ministers to member states on the provision of information through the media in relation to criminal proceedings.* (2003). Retrieved November 10, 2006, from Council of Europe Website: <https://wcd.coe.int/ViewDoc.jsp?id=51365>.
- Shepherd, J., Davies, G., & Ellis, H. (1981). Studies of cue saliency. In G. Davis, H. Ellis, & J. Shepherd (Eds.). *Perceiving and remembering faces* (pp. 105-131). London: Academic Press.
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, *31*, 137-149.
- Swets, J.A. (1986). Indices of discrimination or diagnostic accuracy: Their ROCs and implied models. *Psychological Bulletin*, *99*, 100-117.
- Thornton, I.M., & Kourtzi, Z. (2002). A matching advantage for dynamic human faces. *Perception*, *31*, 113-132.
- Young, A.W., Hay, D.C., McWeeny, K.H., Flude, B.M., & Ellis, A.W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception*, *14*, 737-746.

Received May 25, 2007

Revision received May 8, 2008

Accepted May 24, 2008