ARE EXPECTATIONS OF OBSTRUCTED FACIAL FEATURES ACCURATE?

A Thesis

Submitted to the Department of Psychology

In Partial Fulfillment of the Requirements

For the Degree of

Bachelor of Arts Honours

in

Psychology

University of Regina

By

Magdalyn Derow

Regina, Saskatchewan

April, 2023

Abstract

To fill in missing facial information of partially obstructed, unfamiliar faces, it is believed that people form an accurate holistic expectation. In support of this claim, Winand (2022) demonstrated that participants could correctly match the bottom half of a face to its top half. Yet, the study is limited by the fact that participants may have been able to match the face halves using superficial characteristics such as shading and texture rather than the shapes and sizes of features. Therefore, the purpose of this study was to replicate Winand's (2022) findings with a task in which such superficial matching strategies cannot be used. This was achieved by morphing images together to blur out such inconsistencies in the stimuli. Undergraduates (N=122) were shown the top and bottom halves of a face that belonged to either a single-identity (two photos of the same identity morphed together) or dual-identity (two different identities morphed together). Participants toggled between two randomly chosen bottom halves that belonged to either the same people shown in the top half or different people, and chose the bottom half that best matched the top half. Overall, accuracy was well above chance, but highest when choosing the best single-identity bottom half for a single-identity top. Thus, although incorporating another identity decreases accuracy, people are generally able to find similarities among the top and bottom half identities without the aid of superficial characteristics. This suggests that people accurately form a holistic expectation based on the available top features.

Acknowledgments

This study was supported by the University of Regina's department of Psychology. I would like to extend my sincere appreciation towards my supervisors Dr. Chris Oriet and Yaren Koca for their guidance, support, patience, and, most importantly, their belief in me. I will apply the knowledge you have given me regarding research, academia, and life in general throughout my future endeavors. I am honoured to have had the experience of having two supervisors and could not have asked for anyone better. I would also like to thank Amanpreet Sidhu and Logan Ueckert for helping Yaren prepare the stimuli for the experiment, your assistance helped the process run smoother. I would also like to thank my fellow honours classmates for their support as well as the motivation they brought me by observing their passion and determination to succeed. I cannot wait to see where your journeys will take you all. I would also like to thank my professors for encouraging me to apply to the honours program and providing me with the skill set to succeed. I would also like to thank all the C5 members for providing me with the opportunity to present to them and for offering me constructive criticism. Last but not least, I would like to thank my family and friends for lending me their ears, when I needed to run some ideas past them, and for their never-ending support throughout my thesis journey.

Are Expectations of Obstructed Facial Features Accurate?

In response to COVID-19, the use of masks has become a regular occurrence within society. In turn a new phenomenon known as "mask-fishing," when an individual is shocked at how someone appears once they remove their mask, has emerged (Robles, 2022). In order for mask-fishing to occur, an expectation has to be formed or else one would not be shocked. This suggests that people have a holistic tendency to fill in the missing pieces of facial information. Although one may assume these expectations are not accurate due to masks interfering with visual perception, Winand's (2022) study demonstrates when given the top half of a face participants can correctly match it with its bottom half. However, the study is limited by superficial characteristics like alignment and tonal inconsistencies, since they may have been used as cues to find the answer. What happens when superficial characteristics are removed by being blurred over through morphing images together? Are people still accurate? Therefore, the aim of this study is to examine whether people can correctly predict occluded facial features from the top available features when they cannot rely on superficial characteristics.

Young et al. (1987) discovered the composite face effect which is thought to arise from the holistic approach people take to processing faces. Through conducting four experiments, the authors explored the influence of facial alignment and orientation on face processing. The first experiment presented composite faces in which the top and bottom halves came from different famous people. In one condition the top and bottom halves of two different faces were aligned, while in the other condition the top and bottom halves of two different faces were misaligned. The time required to identify the name of the top or bottom half of each composite was measured. The reaction times were longer for aligned composites, suggesting that one half of the face interfered with identifying the other half only when the halves made a complete whole.

Considering the composites interfered with participants` ability to identify constituent parts, this experiment suggests that faces are processed holistically rather than by individual features. Experiment 2 compared the effects of upright aligned and misaligned composites to inverted aligned and misaligned composites using the same procedure as Experiment 1. The goal was to investigate whether disrupting holistic processing with inversion (Yin, 1969) would reduce the composite face effect. The results in the upright condition replicated the findings of Experiment 1, in which longer reaction times were observed for aligned stimuli than for misaligned. On the contrary, there was no difference between inverted aligned and misaligned reaction times, suggesting that the interference from the non-target face half originates from holistic processing. In Experiment 3, participants studied only the top halves of unfamiliar faces followed by a recognition test in which these top halves were paired with different people`s bottom halves. Since unfamiliar and familiar faces are processed differently (Bruce & Young, 1986), the authors wished to confirm that the results of Experiment 1 generalize to unfamiliar faces. As hypothesized, the reaction times for aligned composites were still longer than for misaligned composites, but the difference was smaller than in the two previous experiments as reaction times to unfamiliar faces were increased overall. The authors concluded that bottom halves disrupt the recognition of unfamiliar top halves, providing further evidence of holistic processing of faces. Lastly, in Experiment 4, since internal and external features influence recognition differently (Ellis et al., 1979), the authors created composites by combining internal and external features from different individuals' faces rather than different top and bottom halves to test the generality of the composite face effect. In the aligned condition both the internal and external features were upright. In the misaligned condition one set of features was upright while the other set was inverted. Participants were instructed to name the owner of the internal or external

features as quickly as possible. Longer reaction times were observed for aligned faces than for misaligned faces. Collectively, all four experiments demonstrate that faces are processed holistically such that perception of a face as an integrated whole interferes with the processing of its constituent parts, creating the perception of an emergent novel identity.

Due to COVID-19, surgical masks are mandated in certain settings as a first line of defense; thus, mask use is prevalent within daily social interactions. Consequently, research demonstrates that surgical masks impede unfamiliar identity recognition (Carragher & Hancock, 2020; Freud et al., 2020), suggesting the processing of masked faces relies less on holistic information compared to non-obstructed faces. Specifically, Freud et al. (2020) hypothesized that face masks interfere with perception of faces and consequently alter how faces are processed (Freud et al., 2020). In order to explore this, Freud et al. (2020) conducted two experiments where participants' perceptual abilities were tested either through a non-masked (the control group) or a masked face memory test (experimental group). The tests had two phases: Phase 1 had six targets for participants to familiarize themselves with, then Phase 2 had three faces, one of which was a target face that they were asked to identify (Duchaine & Nakayama, 2006). Participants also completed the test a second time when faces were inverted. Within Experiment 2, the same procedures were in place except that in the experimental group faces were only masked in one of the two phases randomly (Freud et al., 2020). The results of both experiments indicated that masks impair recognition accuracy. In Experiment 1, although participants performed above chance in the masked condition, their accuracy declined by 15% compared to non-masked performance. Despite faces only being masked in one of the two phases in Experiment 2, a similar deficit to Experiment 1 was replicated. In terms of holistic processing, Freud et al. (2020) found that participants' accuracy in the control group (non-masked) dropped

by 26% when faces were inverted as opposed to upright, showing evidence for the face inversion effect. For participants in the experimental group (masked), accuracy dropped by only 16%, suggesting inversion has less of an impact on masked faces. Interestingly, the inverted face effect is thought to occur because it reduces configural processing, the ability to form relationships between two or more facial features at once (Freud et al. 2020; Piepers & Robbins, 2012). Since holistic processing is essentially configural processing but when all facial features are processed together (Piepers & Robbins, 2012), Freud et al.'s (2020) results suggest that mask wearing decreases holistic processing, which in turn decreases perception accuracy.

Concurrently, Carragher and Hancock (2020) analyzed the effects of surgical masks on human face recognition performance, for both familiar and unfamiliar faces and compared the results to a deep neural network (DNN) simulation designed to emulate human face recognition performance. To accomplish this, they had participants view pairs of images of people, some of whom were wearing surgical masks. Participants were asked to indicate whether the images depict the same person or two different individuals. They tested participants in one of three conditions: (1) control (no masks), (2) mixed (only one face wearing a mask), and (3) masked (both faces wearing masks). As predicted they found that masks impact sensitivity, the ability to accurately differentiate matches from mismatches. Specifically, they observed that sensitivity for the control stimuli was significantly higher than sensitivity for mixed and masked stimuli, suggesting that presence of masks increases the difficulty of face matching. Further, masks elicited a response bias. When familiar faces were masked, participants more often than not considered them to be matches as indicated by results on one test. Conversely, unfamiliar faces were more likely to be reported as mismatches when measured using a different test. This tendency can be seen in both the mixed and masked conditions. The findings confirm that

humans process faces holistically. Accordingly, their findings demonstrate that surgical masks interfere with a person's ability to tell whether two photos depict the same person.

If holistic processing plays a significant role in identity recognition and surgical masks disrupt unfamiliar face recognition, is it possible the unseen portion is filled in with an expectation to form a holistic representation? The phenomenon of "mask-fishing," suggests this may be the case, since an expectation has to be formed for one to be shocked. Recently, Kramer and Jones (2022) proposed that when people view partially obscured faces they imagine that the hidden features are similar to the corresponding features of the prototypical (i.e., average) face. Participants were shown either the top or the bottom half of a face. The other half of the face was a morph created by averaging the actual target`s missing half with the corresponding half of the prototypical face in varying proportions, yielding a gradient ranging from 0% target/100% prototype to 100% target/0% prototype. Their task was to imagine what the missing half looks like and adjust the relative proportion of the target and prototype in the morph to achieve the best match. Kramer and Jones (2022) observed that participants preferred an option that more closely resembled the prototypical face to either the actual face or a version that exaggerates the person's distinctive features. The results provide evidence that the unseen portion of the face can be filled in based on prior experience.

Additionally, Winand (2022) conducted a study in which participants were shown the top half of a face with either the corresponding bottom half or a different person`s bottom half in an effort to provide insights into how holistic processing can aid in the formation of expectations. Each participant was asked to identify which bottom half out of two faces best paired with the top half. Notably, the participants were led to believe that there were no correct answers to discourage participants from employing strategies to match the top- and the bottom-half based on

superficial, image-specific similarities (e.g., similarities in shading or lighting). It is important to control for such strategies because participants could perform better than expected by guessing even if they are unable to match the top and bottom halves of faces based on features. Winand (2022) hypothesized that, by using information in the top half of a face to compute what the bottom half should look like, participants could generate a holistic representation of a face that would allow them to recognize when the correct bottom half is paired with the top half. Participants` mean accuracy was above chance at 69%, suggesting they were not randomly guessing. Therefore, the findings suggest that when given a forced choice between a correct and incorrect bottom half of a face, people can reliably choose the correct bottom half.

Notably, participants were not given the opportunity to form an expectation of how the bottom half appears because they were only shown complete faces. Further, tonal and alignment inconsistencies between the top and bottom halves may have influenced participants` choices despite the instruction to ignore such differences. However, in Experiment 2 the top half was presented by itself to give participants the opportunity to form their own expectations without the influence of the bottom half. Participants could view the top half for as long as they wished. They then could press the spacebar to reveal either the target's or a different individual's bottom half and were asked to determine whether the features in the bottom half look like they match the top half's features. Although participants had more difficulty with the task of determining which pairs were a match than in Experiment 1, they still performed above chance (60% accuracy). The results of Experiment 2 suggest that participants generate an expectation of what the complete face looks like based on the top half. Like Experiment 1, Experiment 2 suggests people are relatively accurate at matching the top half of a face with its corresponding bottom half.

Based on the above research, it is evident that masks have a negative effect on human

perceptual abilities, yet participants in Freud et al.'s (2020) and Winand's (2022) studies were still able to perform above chance. Although Winand's study explores to what extent these expectations are accurate, participants may have been able to get the correct answer by attending to superficial characteristics, such as tonal and alignment inconsistencies, rather than facial features. This may explain why participants performed above chance, even if they did not have an expectation of what the missing half of the face looked like. This is a possibility as in Young et al.'s (1987) study participants had faster reaction times for misaligned composites. However, I believe that through morphing identities together these superficial characteristics can be mitigated. Therefore, the aim of this study is to examine whether people can correctly predict occluded facial features from the visibly available features when they cannot rely on matching superficial characteristics. I hypothesize that the mismatched bottoms will appear 'incorrect' to the participants without them relying on similarities in superficial characteristics.

Methods

Participants

Based on G*Power for a factorial ANOVA (alpha = 0.05, Cohen's d = 0.25, power = 0.80), I determined a minimum of 128 participants were needed. Simultaneously, I had a stopping rule of BF 30 (strong evidence), since my planned analyses included an Open-Ended Sequential Bayes Factor Analysis. Thus, if I obtained a BF 30 before 128 participants, the plan was to stop collecting data and vice versa. Ultimately, 122 undergraduate students, who were enrolled at the University of Regina, were recruited through the Psychology Participant Pool. This number of participants was sufficient as indicated in the results section. Participants were predominantly female (two non-binary, 10 non-disclosed, and 48 male), right-handed (one ambidextrous, seven left-handed, nine non-disclosed), and ranged from ages 18 to 47. Notably,

60 participants had a visual impairment but were wearing glasses or contacts during the task. A 1% credit towards any chosen 100- or 200-level psychology course taken by each participant was granted as compensation. Prior to conducting the experiment, the University of Regina's Research Ethics Board gave their approval.

Stimuli and Apparatus

300 unfamiliar face identities (100 for the top halves and 200 for the bottom halves) were gathered from the Glasgow University Face Database (GUFD; Burton et al., 2010). This number is necessary as identities should not be repeated to prevent implicit learning from previous exposure. The photos consisted of Caucasian faces of people ages 18 to 30 with neutral expression at a front view angle with no obstructions, such as glasses. Two images of each identity were collected. To mitigate the influence of tonal and alignment inconsistencies between the top and bottom halves, single-identity and dual-identity composites were created by morphing two images together using InterFace, a software for morphing faces and were presented in grayscale (Kramer et al., 2017). Single-identity composites were created by morphing two images of the same identity (A & A`). Dual-identity composites were created by morphing two images of different identities (A & B). Each resulting composite were divided into its top and bottom half; therefore, the experiment has two independent variables (1) Top-half condition: Single Identity Tops versus Dual Identity Tops (sID tops v. dID tops) and (2) Bottom-half condition: Single Identity Bottoms versus Dual Identity Bottoms (sID bottoms v. dID bottoms). When the two independent variables are crossed, 4 conditions are formed (1) Single ID top half - Single ID bottom half (Figure 1A), Single ID top half - Dual ID bottom half (Figure 1B), Dual ID top half - Single Id bottom half (Figure 1C), and Dual ID top half - Dual ID bottom half (Figure 1D). Notably, the Dual ID top half - Single ID bottom half (Figure 1C)

constituted as the control condition, due to the fact that each bottom option is a morph of one of the two top identities. Therefore, neither choice is more correct than the other. The experiment was programmed using jsPsych (de Leeuw, 2015) and run on Pavlovia.org. The experiment was conducted on participants' personal computers.

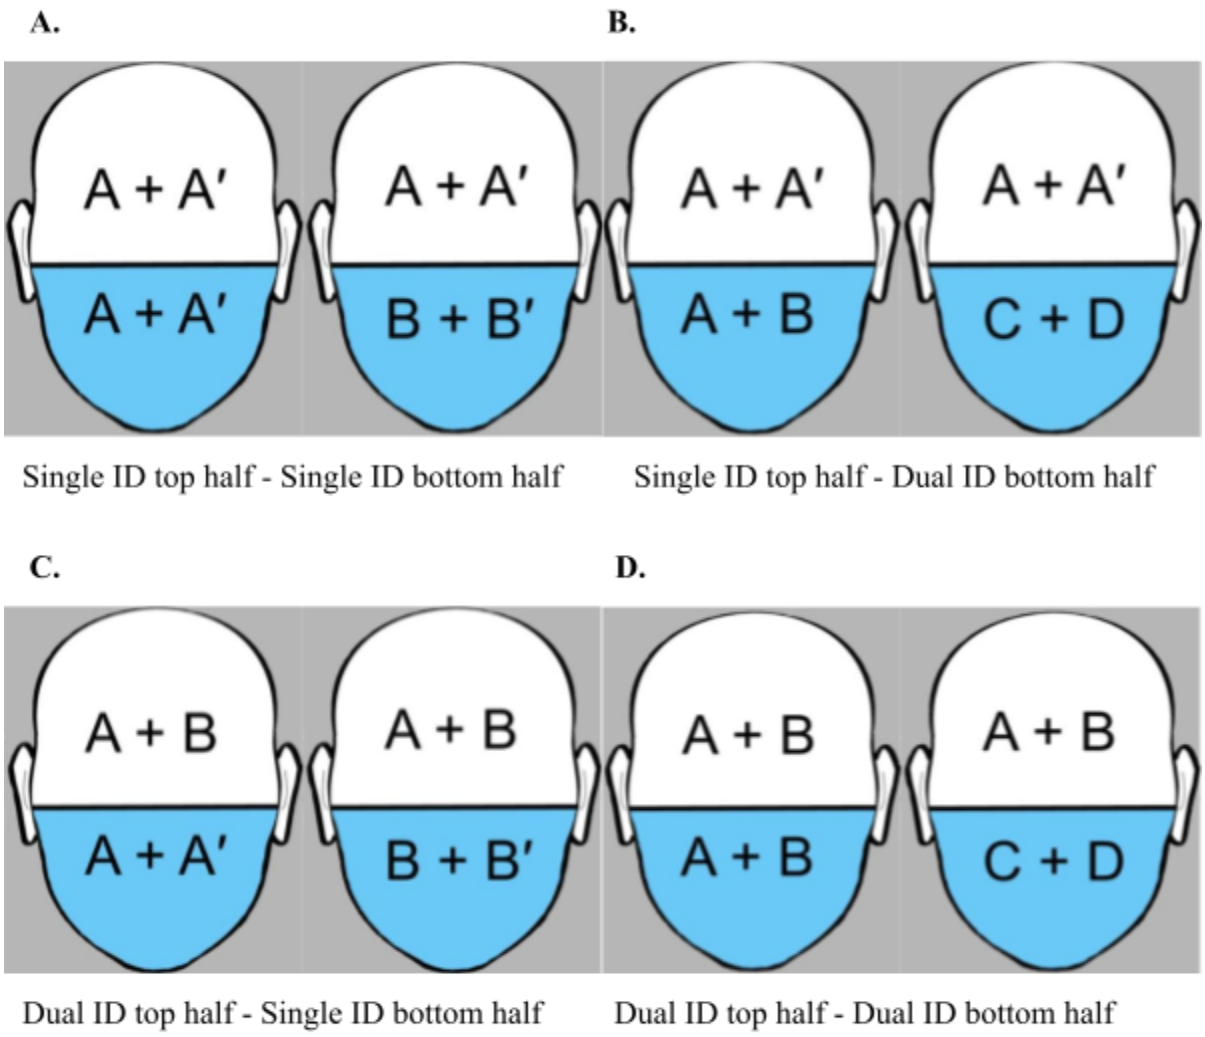

Figure 1. Visuals of the four conditions that will be used. Note: There will only be one top half presented and the participants will toggle between the two bottom half options.

Procedure.

Once the participants completed the consent and demographic forms administered through Qualtrics software, they were redirected to the experiment on the Pavlovia website. Each participant was randomly assigned to one of the four conditions, making the experiment between-subjects. Prior to beginning the experiment, participants were instructed to view the top half shown and toggle between two possible bottom halves. Their task was to indicate which bottom half best matches the top half. Specifically, they were instructed to rely on feature consistency between the top and bottom halves for making their decisions. The goal was to have people base their decisions on the shapes and sizes of features rather than superficial similarities like lighting and quality of the alignment of the top and bottom halves.

Participants completed 50 unique randomized trials. Each identity was only shown once, so participants could not familiarize themselves with an identity. Participants could toggle between two bottom halves as many times as they wished by pressing the "j" key on their keyboard without any time constraints. Following an examination of each composite, participants locked in their answer by pressing the "enter" key on their keyboard. Once the "enter" key was tapped the participants could no longer return to the previous trial. Immediately after, a plus sign "+" appeared at the center of their screen in place of the previous stimuli and disappeared on its own to signal the beginning of a new trial.

Analysis & Results

To compare the mean scores of subject accuracy under the four different conditions two factorial ANOVAs, a null hypothesis statistical test and a Open-Ended Sequential Bayes Factor Analysis, with participants' accuracy as the dependent variable, were conducted to see which hypothesis had supporting evidence- either the null or alternative hypothesis. The null hypothesis is that participants are guessing when matching the two halves when superficial characteristics

are removed. The alternative hypothesis is participants are accurate in matching the two halves without relying on superficial characteristics. JASP (JASP Team, 2022) was the software used to perform the tests. I had expected that participants would perform above chance for all conditions besides our control condition (Dual ID top half - Single ID bottom half), which should be at chance. However, I did expect participants in Single ID top half - Dual ID bottom (Figure 1B) to have less accuracy than participants of conditions Single ID top half - Single ID bottom half (Figure 1A) & Dual ID top half - Dual ID bottom half (Figure 1D), since there is no correct answer but a more correct answer. There is no perfectly correct answer because an additional identity is presented with the top's identity blended into the bottom option, so it partially matches the top's identity. If mismatched bottoms appear incorrect to the participants, then it will suggest they are able to form expectations of how a face should look as a whole without relying on similarities in superficial characteristics.

Results

Participants' average accuracy was above chance in all four conditions: Single ID top half - Single ID bottom half (Figure 1A; $M = 76.3$), Single ID top half - Dual ID bottom half (Figure 1B; $M = 62.6$), Dual ID top half - Single Id bottom half (Figure 1C; $M = 58.2$), and Dual ID top half - Dual ID bottom half (Figure 1D; $M = 67.9$). The 2 x 2 factorial ANOVA revealed that there is a statistically significant main effect for the top condition, $F(1,117) = 19.512$, $p < .001$, $\eta_p^2 = .143$, no main effect for the bottom condition, $F(1,1) = 1.892$, $p < .172$, $\eta_p^2 = .016$, and a significant interaction of top condition by bottom condition, $F(1,117) = 66.574$, $p < .001$, $\eta_p^2 = .016$. Specifically, a Tukey post-hoc test with alpha = .05 revealed that there is a difference between conditions Dual ID top half - Dual ID bottom half and Single ID top half - Dual ID bottom half $t(1) = 2.63$, $p < .047$, Dual ID top half - Dual ID bottom half and Dual ID top half -

Single Id bottom half $t(1) = 4.94$, $p < .001$, Dual ID top half - Dual ID bottom half and 1 $t(1) = 4.12$, $p < .001$, Single ID top half - Dual ID bottom half and Single ID top half - Single ID bottom half $t(1) = 6.56$, $p < .001$, and Dual ID top half - Single ID bottom half and Single ID top half - Single ID bottom half $t(1) = 8.93$, $p < .001$. In contrast, the difference between Single ID top half - Dual ID bottom half and Dual ID top half - Single Id bottom half $t(1) = 2.141$, $p < .146$ is not statistically significant. The alternative hypothesis is also supported by the results of the Open-Ended Sequential Bayes Analysis $BF_{10} = 1.82 \text{ X } 10^{11}$). Taken together, the results are consistent with the hypothesis that participants can accurately match the top and bottom half of a face even when superficial characteristics cannot be used to determine the correct pairing.

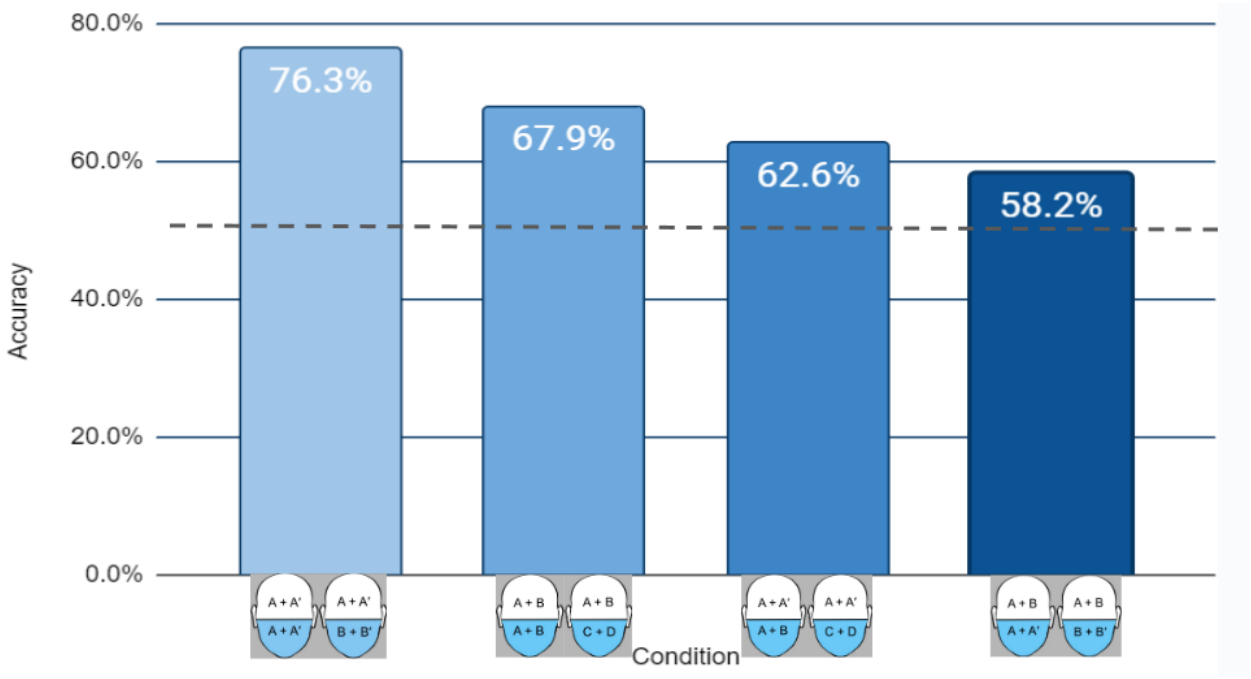

Figure 2. A visual of the participants' mean accuracy in all four conditions: starting from the left to the right Single ID top half - Single ID bottom half; Dual ID top half - Dual ID bottom half; Single ID top half - Dual ID bottom half; Dual ID top half - Single Id bottom half.

Discussion

The aim of the present study was to replicate Winand's (2022) results to address whether participants can accurately fill in missing facial information of partially obstructed faces. To achieve this, I conducted the same matching task as Winand but attempted to remove her study's limitation of superficial stimulus characteristics, like discolouration, with Photoshop and morphing together photos. In accordance with Winand's findings, the present study demonstrates that participants can accurately identify which bottom half best suits the top half of a face. Using morphing I demonstrated that participants are able to accurately determine which bottom half appears "incorrect" without relying on similarities in superficial characteristics. Within all four conditions participants performed well above chance.The decrease in accuracy from Single ID top half - Single ID bottom half (Figure 1A) to Dual ID top half - Dual ID bottom half (Figure 1D), as seen in Figure 2, may be explained by Kramer and Jones' (2022) claim that people assume hidden facial features to be similar to "average" faces in their environment. Including an extra identity (dual-identity) into morphs may eliminate some identity cues that people use to match the bottom half by making the face more typical/average, whereas single-identity morphs may have a more obvious answer due to their own individualistic characteristics being morphed together with themselves.

Although there is support for the composite face effect (Young et al., 1987) and for surgical masks impeding unfamiliar face recognition (Carragher and Hancock, 2020; Freud, 2020), the question still remains whether the expectation of the bottom half is accurate when a face is partially obstructed.The present study does not directly address this, considering participants never had the opportunity to form their own expectation of what the top halves' bottom halves would look like, as they were always shown one of the two bottom options paired with the top. Replicating Winand and Oriet's (2022) Experiment 2 also with morphed stimuli

could be a future direction to consider, as it may open up an avenue to understand whether the generation of a holistic representation for partially obstructed faces occurs.

Like Winand's (2022) study, the present study has its own limitations. Mainly, morphing cannot remove superficial characteristics entirely. There is also the potential to combine differing superficial characteristics when morphing together two separate identities. This means that although our goal was to blur out superficial characteristics, if they are different ones, they may not even each other out and accidentally make the other option more obvious. Additionally, the generalizability of our findings may be limited by the fact that all of our participants are Canadian undergraduate students. Replicating our findings with a more representative sample would be ideal for future studies. Further, allowing participants to toggle as many times as they wish does not necessarily simulate a natural response of seeing an unfamiliar face in real time unless one is analyzing an individual, such as security protocols in relation to photo IDs and/or an eye-witness identifying a potential suspect in a line-up. Our results may suggest that people can tell whether the individual behind a mask is the person in the photo on a piece of ID, which is of value, since some individuals are immunocompromised and should not have to jeopardize their health by having to remove their mask to be identified. However, an experiment directly addressing this is necessary. As for eye-witness line-ups, a few areas should be explored- children's perception and the impact of the other-race effect on accuracy performance. Developmental psychology research has shown that holistic processing begins around 7-8 months (Kanazawa et al., 2018); however, children's accuracy may be inhibited, depending on their age, as a result of not being exposed to enough faces during their short life span. Additionally, being around mask wearers early in infancy may have the potential to slow visual development. Similarly, the other-race effect, people's facial recognition memory of other-race

faces, is poorer than for their own race (Cross et al., 1971), may be the result of not being exposed to enough other-race faces. Potentially, a lack of exposure in both areas may prevent the formation of accurate holistic expectations. Not to mention, with marginalized groups being more likely to be wrongly convicted (Itskovich et al., 2023), exploring the other-race effect with this paradigm is of importance. Further, in live eye-witness line-ups the eye-witness compares faces right beside each other. Therefore, having participants do a matching task, where they cannot toggle between options and are rather presented side-by-side may be beneficial. Notably, it would be important to switch the order of the "correct" bottom half as some individuals may have a response bias that leads to them favouring one side over the other irrespective of what is shown there. Beyond faces, this experiment may pique one's interest to explore the accuracy of holistically formed expectations of objects missing visual information (Gestalt), considering humans possess the ability to recognize objects' by filling in missing pieces of visual information and that the relationship between the manifestation of holistic processing for faces versus objects is ambiguous (Curby & Moerel, 2019).

Conclusion

To be shocked at how someone appears when they remove their mask, one has to form an expectation. To evaluate the accuracy of these expectations, Winand (2022) had participants choose the correct bottom half of a face out of two options for its corresponding top half. Although participants were above chance in accuracy, the study was limited by superficial characteristics like skin texture. I conducted the same experiment but removed the superficial characteristics with morphing. Just like Winand, I found that participants' accuracy was well above chance without the help of superficial characteristics. Evidently, participants are able to match the correct bottom half of a face to its top half, indicating they could see the similarities

between halves. However, future research is needed to directly address whether people do form a holistic expectation when the bottom half facial features are hidden.

References

Burton, A.M., White, D., & McNeill, A. (2010). The Glasgow face matching test. *Behavior Research Methods, 42*, 286–291. https://doi.org/10.3758/BRM.42.1.286

Bruce, V., & Young, A. (1986). Understanding face recognition. *The British Journal of Psychology*, *77*(3), 305–327. https://doi.org/10.1111/j.2044-8295.1986.tb02199.x

Carragher, & Hancock, P. J. B. (2020). Surgical face masks impair human face matching performance for familiar and unfamiliar faces. *Cognitive Research: Principles and Implications*, *5*(1), 59–59. https://doi.org/10.1186/s41235-020-00258-x

Cross, J.F., Cross, J. & Daly, J. (1971). Sex, race, age, and beauty as factors in recognition of faces. *Perception & Psychophysics 10*, 393–396. https://doi.org/10.3758/BF03210319

Curby, K.M., & Moerel, D. (2019). Behind the face of holistic perception: Holistic processing of Gestalt stimuli and faces recruit overlapping perceptual mechanisms. *Atten Percept Psychophys 81*, 2873–2880.. https://doi.org/10.3758/s13414-019-01749-w

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a web browser. *Behavior Research Methods*, *47*(1), 1–12. https://doi.org/10.3758/s13428-014-0458-y

Duchaine, B., & Nakayama, K. (2006). The Cambridge Face Memory Test: Results for neurologically intact individuals and an investigation of its validity using inverted face stimuli and prosopagnosic participants. *Neuropsychologia*, *44*(4), 576–585. https://doi.org/10.1016/j.neuropsychologia.2005.07.001

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*(4), 431–439. https://doi.org/10.1068/p080431

Freud, E., Stajduhar, A., & Rosenbaum, R.S. (2020). The COVID-19 pandemic masks the way people perceive faces. *Sci Rep, 10*, 22344. https://doi.org/10.1038/s41598-020-78986-9

Itskovich, E., Factor, R., & Ohana, D. (2023). Haven't they suffered enough? Time to exoneration following wrongful conviction of racially marginalized minority- vs. majority-group members. *Punishment & Society*, 146247452211483–. https://doi.org/10.1177/14624745221148318

JASP Team. (2022) JASP (Version 0.16.4)[Computer software].https://jasp-stats.org/

Kanazawa, S., Nakato, E., & Yamaguchi, M. (2018). Holistic processing in mother's face perception for infants. *Infant Behavior & Development, 50*, 257-263. https://doi.org/10.1016/j.infbeh.2018.01.007

Kramer, R. S .S., Jenkins, R., & Burton, A. M. (2017). InterFace: A software package for face image warping, averaging, and principal components analysis. *Behav Res 49*, 2002–2011. https://doi.org/10.3758/s13428-016-0837-7

Kramer, R. S. S., & Jones, A. L. (2022). Incomplete faces are completed using a more average face. *Cognitive Research: Principles and Implications*, *7*(1), 1–12. https://doi.org/10.1186/s41235-022-00429-y

Piepers, D. W., & Robbins, R. A. (2012). A review and clarification of the terms "holistic," "configural," and "relational" in the face perception literature. *Frontiers in Psychology*, *3*, 559–559. https://doi.org/10.3389/fpsyg.2012.00559

Robles, J. (2022). Masks are gone on campus, 'mask fishing' is officially over. *University Wire*. https://login.libproxy.uregina.ca:8443/login?url=https://www.proquest.com/wire-feeds/masks-are-gone-on-campus-mask-fishing-is/docview/2637751649/se-2

Winand, C. (2022). Examining accuracy in a facial feature matching task. *Faculty of Arts, University of Regina*. http://hdl.handle.net/10294/14908

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*(1), 141–145. https://doi.org/10.1037/h0027474

Young, Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception (London)*, *42*(11), 1166–1178. https://doi.org/10.1068/p160747n