Face Likeness Affects Face Learning

Benjamin Balas¹ and Claire Bailly²

Authors' Note:

1. Psychology Department, North Dakota State University. Fargo, ND, USA.
2. North Dakota Governor's School, Fargo, ND. USA

Correspondence regarding this article should be addressed to Benjamin Balas,
Department of Psychology, North Dakota State University, Fargo, ND 58102

Contact: Benjamin.balas@ndsu.edu

Word Count = 6400 words.

Abstract

Learning to recognize new faces depends critically on learning to cope with the substantial appearance variability that faces are subject to: The same individual can look very different under different viewing conditions. When attempting to match or discriminate between unfamiliar faces, observers tend to misinterpret instances of within-person variability as different exemplars. Exposure to high variability during face learning tends to support better performance in subsequent recognition tasks, and idiosyncratic face variability becomes a part of representations that support accurate face recognition. Though it is clear that the *amount* of variability is important for learning, it remains an open question whether or not different types of variability support more robust recognition abilities after training. In the present study, we addressed this question by asking observers to complete an incidental learning task in which they were either presented with good likenesses or poor likenesses of target faces. Following this training task, all participants were tested with a novel set of face images including new exemplars of the individuals presented during training. We found that the quality of face likenesses affected performance during the training task, but did not lead to different outcomes at test. Both kinds of training improved different aspects of performance relative to a control group that received no training at all, however, suggesting that both good and poor examples of an individual's appearance are useful for establishing representations for recognition.

Introduction

If you want to be good at recognizing faces, you must learn to cope with variability. Any individual face can look very different from moment to moment as a function of lighting, expression, viewpoint, and changes in hairstyle, makeup, or accessories like sunglasses or hats. This *intra-personal variability* across different images of the same face is one key difficulty in achieving accurate face recognition. The other half of the problem is the varying amounts of *extra-personal variability* across different individuals. Some pictures of different people look very similar, while in other cases different people look quite different. Successfully recognizing faces requires distinguishing between these two kinds of variability. Formalizing face recognition along these lines is useful in machine vision applications (Moghaddam, Jebara & Pentland, 2000), and also helps explain some aspects of human face perception (Balas, 2012; 2013).

Making the distinction between these two sources of appearance variability also helps to highlight some of the key limitations of human face recognition. In particular, a number of recent studies have demonstrated that observers find it difficult to identify instances of intra-personal variability when asked to recognize unfamiliar faces (Jenkins et al., 2011; Andrews et al., 2015; Balas & Pearson, 2017;). Specifically, observers tend to assign different labels to images of the same person, while rarely making the opposite error of assigning the same label to images of different people (Jenkins et al., 2011). This pattern of errors has been described as a failure to “tell people together” (Andrews et al. 2015), which nicely emphasizes the key problem: Variation across different images of the same unfamiliar person is often interpreted as evidence of a change in identity.

This pattern of errors depends on the degree of familiarity observers have with the faces under consideration. The well-known “other-race effect” is frequently described in terms of overall poorer recognition of other-race faces relative to own-race faces (Meissner & Brigham, 2001), and with regard to appearance variability the aforementioned pattern of errors is exacerbated for other-race faces. When attempting to group other-race faces according to identity, observers have an increased tendency to incorrectly classify images of the same person as belonging to different people (Laurence, Zhou & Mondloch, 2015). The effects of limited experience with faces are not restricted to category-specific deficits in recognition, however. Observers with globally limited exposure to faces also have poorer abilities to distinguish between intra- and extra-personal appearance variability. This includes individuals from depopulated regions (Balas & Saville, 2017), who have reduced category-selective neural responses to faces (Balas & Saville, 2015) and homeschooled individuals (Short, Balas & Wilson, 2017). However, the effects of limited face exposure can be dramatically mitigated by familiarity with specific individuals’ faces. Regarding the other-race effect, using familiar other-race faces reduces observers’ recognition errors regarding both types of appearance variability (Zhou & Mondloch, 2016). Face familiarity also improves performance for own-race faces (Jenkins et al., 2011), suggesting that face recognition depends on individual-specific representations of appearance that are learned through extended exposure and feedback.

Familiarity results in improved abilities to distinguish between intra- and extra-personal face variability, but how do observers learn correct patterns of generalization over different images that depict the same familiar individual? One important by-product of familiarity is a greater reliance on internal, rather than external, features (Ellis, Shepherd & Davies, 1979;

Young et al. 1985; Bruce et al., 1999). Limiting observers' access to these cues confirms that internal features carry more valuable information for coping with appearance variability correctly (Longmore, Liu & Young, 2014; Kemp et al., 2016; Kramer et al., 2018). Still, what kind of experience with faces best facilitates this shift in face processing and supports the establishment of a robust representation of the appearance of a familiar individual's face? Several recent results highlight the importance of high variability across images of an unfamiliar individual during learning. Murphy et al. (2015), for example, found that face exemplar variation during training led to measurable advantages at test compared to repeated exposure to a limited set of exemplars. Along similar lines, Ritchie & Burton (2019) demonstrated that even when training conditions involve a matched number of unique exemplars, high variability learning supports better performance than low-variability learning. Learning is also positively affected by separating exposure to new exemplars across multiple instances (Menon, Kemp & White, 2018) suggesting that variability is integrated effectively across separate encounters with a new face. Importantly, within-person variability is not just noise that observers must discount, but serves as a cue to identity. Expression variability, for example, is built into representations of face identity (Redfern & Benton, 2019) and computationally, idiosyncratic variability provides diagnostic information for individuation (Burton et al., 2016). Developmentally, high-variability training is particularly important; Children who learn an unfamiliar face following high- or low-variability training only benefit from high-variability training, in contrast to adults who benefit from both kinds of exposure (Baker, Laurence, & Mondloch, 2017).

Variability is not enough to ensure robust learning, however. Kramer et al. (2017) demonstrated that natural variability, which they operationalized by imposing face inversion or contrast negation, is necessary for learning new faces. Ritchie, Mireku & Kramer (2019) also found that presenting observers with multiple images vs. single images did not affect performance in a live face-matching task. Thus, not all forms of variability are effective for learning new face images, and some tasks require differing magnitudes or types of face variability.

Presently, we explored how the *type* of face variability, rather than the amount, affects observers' ability to learn new identities. Specifically, we examined how exposure to training images that are "good likenesses" affects subsequent recognition of those individuals. Observers show matching advantages for faces that they endorse as good likenesses (Ritchie, Kramer & Burton, 2018) even though these likeness judgments vary across observers. Do good likenesses of a face work better as a training set than poor likenesses? If robust face representations depend on a good estimate of an individual's average appearance (Kramer, Ritchie & Burton, 2015), good likenesses may be more useful. However, poor likenesses may also serve a useful purpose for face learning: non-representative images may provide useful information regarding worst-case scenarios for recognition, leading to representations that easily handle simpler types of variability. If learning to recognize an individual depends on determining what part of face-space their images occupy, good and poor likenesses thus may provide distinct information about the shape of this region: Good likenesses may provide information about it's center, while poor likenesses may provide information about it's boundary. The question we address here is whether or not these two types of appearance variability differ in terms of their impact on face learning.

We examined this question through the application of a card-sorting task that has been widely used in recent studies (Jenkins et al., 2011; Andrews et al., 2015; Balas & Pearson, 2017; Laurence, Zhou & Mondloch, 2016). Participants in this task sort a large set of images into groups based on identity such that each group of images contains all and only images of a single unique individual. While participants usually overestimate the number of unique faces in the set when the true number is unknown, providing them with that number leads to a substantial improvement in performance (Jenkins et al., 2011). This cued version of the sorting task has been used to support incidental learning of the faces depicted in the stimulus set (Andrews et al., 2015), and we use a similar approach here. Participants in our task sorted images of two unfamiliar individuals into two groups based on identity. Half of these participants sorted images that were rated by independent observers as being good likenesses, while the other half of our participants sorted images that were rated as being poor likenesses. Following this cued sorting task, all participants completed an unconstrained version of the task in which they were presented with novel images of the two individuals presented during training, and a third individual who had not been included. We predicted that performance with poor likenesses would be poorer during the training phase, but may lead to superior performance at test due to prior exposure to especially challenging images. Briefly, while we did find that performance during the training phase differed in the predicted direction, test performance did not. However, compared to participants who received no training, we observed effects of both training conditions on different performance outcomes. We discuss these results in the context of emerging theories regarding how variability is incorporated into representations for face recognition and practical applications for face learning.

Methods

We pre-registered our experimental design via the Open Science Framework (<https://osf.io/v5ngc>), including a description of our sample size, experimental manipulations, and analysis plan. After completion of this design and analysis of the data, we recruited an additional participant group to support additional exploratory analyses. We describe these two phases of the experiment separately to emphasize which portions of the study were completed in accordance with our pre-registration, and which were completed afterwards.

Participants

According to our pre-registered design, we aimed to recruit a sample of 72 participants. We ultimately recruited a total of 78 participants, exceeding our initial target due to incidental scheduling errors. We discarded data from 4 participants who did not follow task instructions: These participants either responded that test images all depicted one person, or that each image depicted a different person without inspecting the full stimulus set. Our final sample was thus comprised of 74 participants, approximately half of which (N=38, 20 female) were randomly assigned to the “Good Likenesses” training condition, and the remaining participants (N=36, 22 female) were assigned to the “Poor Likenesses” condition.

After completing the pre-registered design and analyzing the data collected from these participants, we collected data from a third group of participants who were only completed the

test phase of our task. This sample was comprised of 36 participants (19 female), none of whom were included in the initial sample.

All participants self-reported normal or corrected-to-normal vision and were naïve to the purpose of the study. We recruited participants from the NDSU Undergraduate Psychology Study Pool, and all participants received course credit for volunteering. We obtained written informed consent from each participant prior to the beginning of the testing session.

Stimuli

Our stimulus set was comprised of 112 unique images of 3 individuals, all of whom were female Dutch models that our participants were not familiar with. Two individuals (Sylvie Van der Vaart and Doutzen Kroes) were depicted in 48 images, while the third individual (Dewi Driegen) was depicted in 16 images. These images were selected to include a wide range of ambient appearance variation and were obtained via Google Image searches using each individual's name as the search term. Each face was segmented from its parent image (including as much of the external contour as possible), resized for printing on a 3" x 2" card, and printed in full color.

For each image of Sylvie Van der Vaart and Doutzen Kroes, we obtained Likeness ratings from naive observers using the Amazon Mechanical Turk (AMT). To do so, we created average images of both individuals, and presented each original image alongside the corresponding average face. AMT participants were asked to provide a 1-7 rating regarding the quality of the likeness between the original image and the average. We obtained 20 unique judgments per image and calculated the mean rating per image. For each model, we sorted the images according to average likeness ratings and divided the stimuli into thirds: the bottom third comprised our "Poor Likeness" training set, the top third comprised our "Good Likeness" training set, and the remainder were used during the test phase of our task, alongside images of our third individual.

Procedure.

Individuals in our initial sample completed both a training and test phase of the task. During training, participants received either the "Good Likeness" or "Poor Likeness" images, each containing 16 images each of Sylvie Van der Vaart and Doutzen Kroes. We asked these participants to sort the cards into two groups based on identity. Next, all participants received the same test set of images, containing 16 novel images of the individuals presented during training, as well as 16 images of the third individual. Participants sorted these cards according to identity as well, but were not told how many unique people were depicted in the stimulus set. Participants were not told that the individuals presented during training were included in this stimulus set, and were not given any further guidance regarding how to complete the task.

Individuals in our second sample only completed the test phase of the task. Immediately after providing written consent to participate, they received the test set of stimuli and completed the sorting task according to the procedures described above. All participants completed their experimental session within approximately 30 minutes.

Results

Pre-registered analyses

We analyzed the data from each participant by computing signal detection descriptors of sorting performance according to the method outlined in Balas & Pearson (2017). This approach considers the disposition of each *pair* of images in the full stimulus set. Each image pair either depicts the same individual or two different individuals, and so should either be sorted together or apart. Each type of image pair can be sorted incorrectly, leading to *same-person/different-group* errors or *different-person/same-group* errors. By counting these errors in a particular sorting solution, we can calculate a hit rate and false alarm rate for each participant, which further allows us to use d' (sensitivity) and c (response criterion) values to characterize performance. We carried out these computations for each participant's card-sorting solution in the *training* and *test* sessions separately. We also counted the number of groups each participant made in the test phase of the experiment.

Training phase. We analyzed participants' d' and c values during the training phase using two-tailed independent-samples Bayesian t-tests carried out in JASP (JASP, 2018). For our d' values, this analysis yielded a Bayes factor of 6.6 in favor of the hypothesis that our two participant groups had unequal means, meaning that the data we observed were approximately 6 times more likely under this model than a null model in which the mean d' values were equal. Specifically, participants in the Good Likeness condition had higher d' values ($M=1.03$, $SD=1.23$) than participants in the Poor Likeness condition ($M=0.33$, $SD=0.85$). For our values of the response criterion, c , the same analysis yielded a Bayes factor of 0.25. This means that the data we observed in this case was approximately 4 times more likely under the null model in which the two groups had equal means. Participants in the Good Likeness ($M=0.012$, $SD=0.11$) and the Poor Likeness condition ($M=0.007$, $SD=0.08$) did not differ in their tendency to split cards into separate groups, which is of course constrained in this phase of the task.

Test phase. As above, we analyzed the d' and c values obtained from each participant using independent-sample Bayesian t-tests. For d' values, this yielded a Bayes Factor of 0.61, indicating that the data we observed were approximately 1.6 times more likely under a null model in which the two groups of participants had the same mean sensitivity. Thus sensitivity in the Good Likeness condition ($M=0.93$, $SD=0.53$) did not clearly differ from sensitivity in the Poor Likeness condition ($M=0.77$, $SD=0.35$).

For our response criterion values, the same analysis yielded a Bayes factor of 0.32, indicating that the data we observed were approximately 3 times more likely under the null model. Response criterion values in the Good Likeness condition ($M=-1.26$, $SD=0.54$) thus do not appear to differ from those in the Poor Likeness condition ($M=-1.16$, $SD=0.47$). Likewise, the comparison between the number of groups made by observers in the Good Likeness condition ($M=13.6$, $SD=8.7$) and Poor Likeness condition ($M=11.6$, $SD=6.9$) yielded a Bayes Factor of 0.36, indicating nearly 3 times more support for the null model. Overall, these results suggest that the manipulation we imposed regarding the quality of likenesses during training had little to no effect on performance at test.

Exploratory analyses

Following this analysis, we recruited an additional group of participants who completed the test phase without any form of prior incidental learning. We calculated d' and c values for this participant group as well, and compared the results from this sample with each of the previous groups using Bayesian Independent-Samples t -tests.

With regard to our d' values, we found that the comparison between performance in the “No training” group ($M=0.63$, $SD=0.29$) and the “Good likenesses” group ($M=0.93$, $SD=0.53$) yielded a Bayes Factor of 8.4, indicating strong evidence in favor the alternative hypothesis. The comparison between the “No training” group and the “Poor likenesses” group yielded a Bayes Factor of 0.96, indicating relatively equal evidence for the null and alternative hypotheses. Thus, training with Good Likenesses appeared to confer sensitivity benefits relative to no exposure to training images, while Poor Likenesses did not. The data from all three conditions is displayed in Figure 1.

Figure 1 - Boxplots and jitter elements depicting the d' values observed in all three training conditions. There is strong evidence for a difference between the Good Likeness and No Training conditions, but little evidence for other differences across conditions.

For response criterion values, the comparison between the “No training” group ($M=-1.45$, $SD=0.37$) and the “Good likenesses” group ($M=-1.26$, $SD=0.54$) yielded a Bayes Factor of 0.99, indicating equal amounts of evidence for the null and alternative hypothesis. The comparison between the “No training” group and the “Poor likenesses” group ($M=-1.13$, $SD=0.51$) yielded a Bayes Factor of 8.4, indicating strong support for the alternative hypothesis. Training with poor

likenesses appears to reduce the tendency to assign images to different groups relative to receiving no training. The data from all three of these conditions is depicted in Figure 2.

Figure 2 - Boxplots and jitter elements depicting the response criterion values observed in all three training conditions. There is strong evidence for a difference between the Poor Likeness and No Training conditions, but little evidence for other differences across conditions.

Finally, we compared performance across conditions using the number of groups participants made in each condition. Comparing the average number of groups made in the No Training condition ($M=16.1$, $SD=6.9$) to the number of groups made in the Poor Likeness condition led to a Bayes Factor of 4.7, indicating strong support that these means differ. The same comparison between the No Training condition and the Good Likeness condition ($M=13.6$, $SD=8.7$) led to a Bayes Factor of 0.54, indicating anecdotal evidence in favor of the null hypothesis. The data from all three of these conditions is depicted in Figure 3.

Figure 3 - Boxplots and jitter elements depicting the number of groups made by participants in all three training conditions. There is strong evidence for a difference between the Poor Likeness and No Training conditions, but little evidence for other differences across conditions.

Though these analyses are exploratory and include comparisons between data collected as part of our pre-registered design and data collected after our initial analysis, they suggest that the incidental learning provided during our training period did affect subsequent performance. Specifically, the lack of a training session reduced sensitivity relative to the “Good likenesses” condition, and led the response criterion to be more negative relative to the “Poor likenesses” condition. The former effect suggests that training with “Good likenesses” tended to improve the ability to distinguish between intra- and extra-personal face variation, while the latter suggests that participants’ tendency to split cards into more groups is also affected by training.

Discussion

Our results reveal that face likeness - whether an image subjectively looks representative of a person’s appearance - affects observers’ ability to correctly evaluate intra- and extra-personal variability. In our training phase, for example, we found that faces rated as poor likenesses were more difficult to correctly sort by identity compared to faces rated as good likenesses, as quantified by lower d' values in the former case. Critically, both of these sets of stimuli were comprised of ambient images that varied substantially in appearance across exemplars. This result constitutes a manipulation check with regard to the composition of the two training decks we created to examine face learning, which is important to establish given recently published evidence suggesting that the perception of face likeness is idiosyncratic (Ritchie, Kramer & Burton, 2018). Besides verifying that we split our stimulus set into meaningfully different subsets, this outcome also indicates that not all types of image variability are created equal, and

that face likeness is a useful means of examining this issue. Whatever properties lead a face to be rated as a poor likeness also lead to poor sorting performance for unfamiliar, even when observers know the number of unique identities in the stimulus set. In previous reports (Jenkins et al., 2011; Andrews et al., 2015), this cue frequently leads to high accuracy (though see Balas & Pearson, 2017 for results with a 4-category sorting task), so it is noteworthy that poor likenesses lead to such low sensitivity. Besides this result, the lack of a difference in response criterion values is expected due to the instruction to sort the cards presented during the training phase into two piles. We have reported elsewhere that constraining the number of groups participants can make essentially freezes their response criterion within a signal detection framework (Balas & Saville, 2017; Balas & Pearson, 2017), so it is not surprising that imposing a '2-sort' task on participants fixes their response criterion near zero.

Our key research question was whether these different incidental learning experiences (Good vs. Poor likenesses) affected performance at test. Considering the results from the pre-registered portion of our study, likeness quality does not have a large effect on subsequent learning. Our sensitivity values were slightly more likely under the null hypothesis (though a Bayes Factor of this magnitude represents weak evidence for either model), and our response criterion values were three times more likely under the null hypothesis. The impact of incidental learning thus does not appear to depend on the quality of the likenesses provided as training exemplars, which is consistent with previous results demonstrating that high variability leads to better learning outcomes. If we are correct in asserting that both of our training sets had high levels of appearance variability, then our results suggest that it isn't terribly important what kind of variability observers are exposed to. Despite the fact that poor likenesses were difficult to assign to identities in a constrained sorting task, the exposure to images with high variability proved as useful in that case as the alternative condition.

There is an alternate interpretation of these results, however, which is why we included the additional sample of participants after completing our pre-registered analyses. Perhaps neither training condition did anything to change how participants were going to perform in the test phase of the experiment. Rather than good and poor likenesses being equally good for learning to recognize new faces, it could also be the case that this was too little exposure to new identities to have an effect on subsequent performance. We argue that this is why the comparisons between both of these training groups and the participants who completed the task without training are crucial. These results demonstrate that performance in the test phase was better in distinct ways when each kind of training was provided relative to no training. Good likenesses led to higher sensitivity values than no training, and poor likenesses led to less negative response criterion values, which indicates a reduced tendency to split the stimulus set into too many groups (as reflected in our comparison of group numerosity). The latter effect on response criterion could conceivably be non-perceptual: Providing the information that the correct number of groups is 2 in the training phase could lead to an anchoring effect that makes lower group numerosity more plausible. However, this can't entirely account for the difference in sensitivity values between Good Likeness training and no training. In this case, it's not enough to have a small number of groups, they have to have the correct composition. One interpretation of these results would be that good likenesses are in fact better for face learning. In terms of how we envisioned these aspects of facial appearance reflecting properties of a face-space representation, this could indicate that it is better to establish a useful estimate of typical

appearance for a new individual than to learn about worst-case scenarios. The utility of modal face appearance has been demonstrated in a number of studies, including both computational (Burton et al., 2005; Burton et al., 2016) and behavioral (Kramer, Ritchie & Burton, 2015) assessments, so our data potentially resonate with these prior reports. While our data don't allow us to draw this conclusion unequivocally (the Good Likeness and Poor Likeness conditions didn't differ from one another in terms of sensitivity after all), we think that this conclusion is consistent with our results and multiple prior studies.

There are several important limitations of our study that we think are worth considering in future work. First, unlike previous studies that have differentiated between training conditions using the magnitude of variability (Baker, Laurence & Mondloch, 2017), we don't have an easy means of assessing the variability of appearance in our two sets of training stimuli. In previous studies this manipulation has been operationalized through intuitive manipulations (e.g. images taken from a single video session rather than across multiple sessions) rather than quantitative estimates of variability. Still, these intuitive manipulations of appearance variability are probably accurate enough for the purposes of distinguishing between high and low variability. In our case, we just don't know what's in our two training sets in terms of this variable. This is potentially problematic, but could also be interesting to examine closely: What if "Poor Likeness" means "more variable," for example? This should support better learning, but perhaps the influence of appearance variability on face learning is U-shaped or otherwise non-linear. A second limitation of our study is the lack of feedback during our training phase, which may have reduced the effects of incidental learning on subsequent recognition. Providing such feedback during training, perhaps using a perceptual matching task similar to the one developed by Matthews, Davis & Mondloch (2018), could be a useful means of increasing the likelihood that training will have an effect on test performance. Finally, examining recognition at test in a broader set of contexts could also be a useful way to examine the effects of different kinds of face variability presented during training. Imposing a constrained sort at training and test may remove or reduce the subjective aspects of the sorting task used here, for example. Evaluating performance in a pairwise same-different task or a match-to-sample experiment could also reduce the amount of individual variation we observe across participants in the sorting paradigm. All of these approaches are potentially useful ways to extend and refine the results presented here. provide intriguing support that the nature of face variability observed during exposure affects the accuracy of subsequent recognition.

Overall, our results provide intriguing support that the nature of face variability observed during exposure affects the accuracy of subsequent recognition. Face learning depends not only on seeing many different exemplars of a face, but also on seeing representative exemplars of the individuals you hope to recognize later. While there remain many open questions regarding how to make likeness more concrete, and how to appropriately quantify the amount of variability within a set of faces, our results demonstrate that considering these difficult questions more closely may lead to new insights into how representations for face recognition are built from visual experience.

Acknowledgements

The authors would like to thank Dr. Jeff Boyer for his work with the ND Governor's School program to facilitate a pilot version of this study. Thanks also to Sophie Wamre, Trent Johnson, and Amanda Auen for assistance with participant recruitment, testing, and data transcription.

References

- Andrews, S., Jenkins, R., Cursiter, H. & Burton, A. M. (2015) Telling faces together: Learning new faces through exposure to multiple instances, *The Quarterly Journal of Experimental Psychology*, 68, 2041-2050.
- Andrews, S., Jenkins, R., Cursiter, H. & Burton, A.M. (2015) Telling faces together: Learning new faces through exposure to multiple instances. *Quarterly Journal of Experimental Psychology*, 68, 2041-2050.
- Baker, K.A., Laurence, S. & Mondloch, C.J. (2017) How does a newly encountered face become familiar? The effect of within-person variability on adults' and children's perception of identity.
- Balas, B. (2012) Bayesian Face Recognition and Perceptual Narrowing in Face-Space. *Developmental Science*, 15, 579-588.
- Balas, B. (2013) Developing race categories in infancy via Bayesian face recognition. *Visual Cognition*, 21, 1138-1164.
- Balas, B. & Pearson, H. (2017) Intra- and extra-personal variation in person recognition. *Visual Cognition*, 25, 456-469.
- Balas, B. & Saville, A. (2015) N170 face specificity and face memory depend on hometown size. *Vision Research*, 69, 211-217.
- Balas, B. & Saville, A. (2017) Hometown size affects the processing of naturalistic face variability. *Vision Research*, 141, 228-236.
- Bruce, V., Henderson, Z., 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(4), 339-360.
- Burton, A.M. (2013) Why has research in face recognition progressed so slowly? *Quarterly Journal of Experimental Psychology*, 66, 1467-1485.
- Burton, A.M. Jenkins, R., Hancock, P.J. & White, D. (2005) Robust representations for face recognition: the power of averages. *Cognitive Psychology*, 51, 256-284.

Burton, A.M., Kramer, R.S.S., Ritchie, K.L. & Jenkins, R. (2016) Identity from variation: Representations of faces derived from multiple instances. *Cognitive Science*, 40, 202-223.

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.

JASP Team (2018). JASP (Version 0.8.6)[Computer software].

Jenkins, R., White, D. van Montford, X. & Burton, A.M. (2011) Variability in photos of the same face. *Cognition*, 121, 313-323.

Kemp, R.I., Caon, A., Howard, M. & Brooks, K.R. (2016) Improving unfamiliar face matching by masking the external facial features. *Applied Cognitive Psychology*, 30, 622-627.

Kramer, R.S.S., Jenkins, R., Young, A.W. & Burton, A.M. (2017) Natural variability is essential to learning new faces. *Visual Cognition*, 25, 470-476.

Kramer, R.S.S., Manesi, Z., Towler, A., Reynolds, M.G. & Burton, A.M. (2018) Familiarity and within-person facial variability: The importance of the internal and external features. *Perception*, 47, 3-15.

Kramer, R.S.S., Ritchie, K.L. & Burton, A.M. (20125) Viewers extract the mean from images of the same person: A route to face learning. *Journal of Vision*, 15:4, 1-9.

Kramer, R.S.S., Young, A.W. & Burton, A.M. (2018) Understanding face familiarity, *Cognition*, 172, 46-58.

Laurence, S., Zhou, X. & Mondloch, C.J. (2016) The flip side of the other-race coin: They all look different to me. *British Journal of Psychology*, 107, 374-388.

Longmore, C.A., Liu, C.H & Young, A.W. (2014) The importance of internal facial features in learning new faces. *Quarterly Journal of Experimental Psychology*, 68, 249-260.

Matthews, C.M., Davis, E.E. & Mondloch, C.J. (2018) Getting to know you: The development of mechanisms underlying face learning. *Journal of Experimental Child Psychology*, 167, 295-313.

Meissner, C. A., &Brigham, C. (2001). Thirty years of investigating the own-race bias in memory for faces: A meta-analytic review. *Psychology, Public Policy, & Law*,7, 3–35.

Menon, N., Kemp, R.I. & White, D. (2018) More than a sum of parts: robust face recognition by integrating variation. *Royal Society Open Science*, 5, 172381.

- Menon, N., White, D. & Kemp, R.I. (2015) Variation in photos of the same face drives improvements in identity verification. *Perception*, 44, 1332-1341.
- Moghaddam, B., Jebara, T. & Pentland, A. (2000) Bayesian Face Recognition. *Pattern Recognition*, 33, 1771-1782.
- Murphy, J., Ipser, A., Gaigg, S.B. & Cook, R. (2015) Exemplar variance supports robust learning of facial identity. *J. Exp. Psychol: Hum. Percept. Perform.*, 41, 577-581.
- Ritchie, K.L. & Burton, A.M. (2019) Learning faces from variability. *Quarterly Journal of Experimental Psychology*.
- Ritchie, K.L., Kramer, R.S.S., & Burton, A.M. (2018) What makes a face photo a 'good likeness'? *Cognition*, 170, 1-8.
- Ritchie, K.L., Mireku, M.O. & Kramer, R.S.S. (2019) Face averages and multiple images in a live matching task. *British Journal of Psychology*. Doi: 10.1111/bjop.12388.
- Short, L.A., Balas, B. & Wilson, C. (2017) The effect of educational environment on identity recognition and perceptions of within-person variability. *Visual Cognition*, 25, 940-948.
- Young, A.W. & Burton, A.M. (2018) Are we face experts? *Trends in Cognitive Sciences*, 22, 100-110.
- 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.
- Zhou, X. & Mondloch, C.J. (2016) Recognizing "Bella Swan" and "Hermione Granger": No Own-Race Advantage in Recognizing Photos of Famous Faces. *Perception*, 45, 1426-1429.