I Know That Person! Why Can't I Remember?

Logan Ueckert

Department of Psychology, University of Regina

Honours Thesis

Supervisors: Dr. Chris Oriet and Yaren Koca

April 24, 2024

Acknowledgement

I would like to thank the University of Regina and the Psychology Department for the opportunity to conduct this research. I would especially like to extend a heartfelt thank you to Dr. Chris Oriet and Yaren Koca for their supervision throughout this project. I am appreciative of the time that each invested in this work as without them it would not have been possible.

Abstract

The butcher-on-the-bus phenomenon refers to repeatedly encountering a person in a specific context and later failing to recognize them in a novel context. Despite the failure to recognize the person, a strong sense of familiarity occurs. Although this phenomenon commonly occurs, little is known about its underlying cause. One possibility is that the informativeness of the context increases the strength of the association between context and identity. This association may then interfere with recognizing the same person in a novel context. This study was designed to examine the following question: Is recognition of faces in previously-unseen contexts influenced by the allocation of attention to previously-seen contexts? Attention to context was manipulated by altering whether the context provided identity-specific information such as the target's occupation or hobbies, and whether the context was consistent (mimicking the experience of repeatedly encountering a face in the same context) or varied. Unexpectedly, I found that identities first encountered in an informative context were later recognized better in a new context than those encountered in uninformative contexts. This suggests that the informativeness of context provides a powerful cue that aids in the learning of a new face, surpassing the benefit of exposure to within-person variability.

I Know That Person! Why Can't I Remember?

In 1980 Mandler used the term butcher-on-the-bus to describe a common phenomenon related to identity recognition. While the phrase "the-butcher-on-the-bus" may not be familiar to everyone, many people can relate to the phenomenon that it illustrates. This phenomenon describes repeatedly encountering a person in a specific context and later failing to recognize them in a novel context. Continuing with the example from Mandler (1980), one can become familiar with the butcher in the market, but later when the same person is seen on the bus one cannot recognize them as the butcher, yet they still get the sense that they are familiar. I set out to investigate this phenomenon and proposed a potential explanation for it.

In order to understand the butcher-on-the-bus phenomenon it is important to first understand the perceptual and cognitive processes that are involved in recognizing a face. There are multiple cues that can be used to identify a person but the face is the most commonly used identifier. Seeing a person's face provides a lot of information. Bruce and Young (1986) divided the information obtained from viewing a face into the following seven categories: pictorial, structural, visually derived semantic, identity-specific semantic, name, expression, and facial speech codes. Pictorial information constitutes the basic shapes and colours present in a *particular* view of a face. In contrast, structural information is that information that remains constant across *different* views of a face. Visually derived semantic information includes information regarding a person's social category such as gender and age. The identity-specific semantic category includes information that cannot be obtained from unfamiliar faces such as their occupation because they require time to become associated with a person. Name information, although identity-specific, is arbitrary and does not guide social interactions.

Expression information includes all of the combinations of facial actions to convey emotion. Facial speech codes include the information derived from the mouth region that facilitates speech comprehension. Unlike the previous types of information, Bruce and Young argue that expression and facial speech codes are not directly involved in the process of face recognition, yet are still important to the process of distinguishing faces from non-faces.

Bruce and Young (1986) describe a number of interrelated components involved in the processing of the different types of facial information outlined in Figure 1. The inputs are structurally encoded into features that are view-independent (i.e., those that do not change from one view to another such as the metric distance between facial features) and view-dependent (i.e., those that do change across views such as expression). This allows the information to be used by the other components. Identity recognition is supported by two of these: face recognition units and person identity nodes. A face recognition unit responds to any view of a person's face; their voice or name will not evoke a response. However, a person identity node can be accessed via their face, voice, name, gait, or even clothes. Bruce and Young’s model illustrates that information from sources beyond an individual’s face can be used to identify a person.

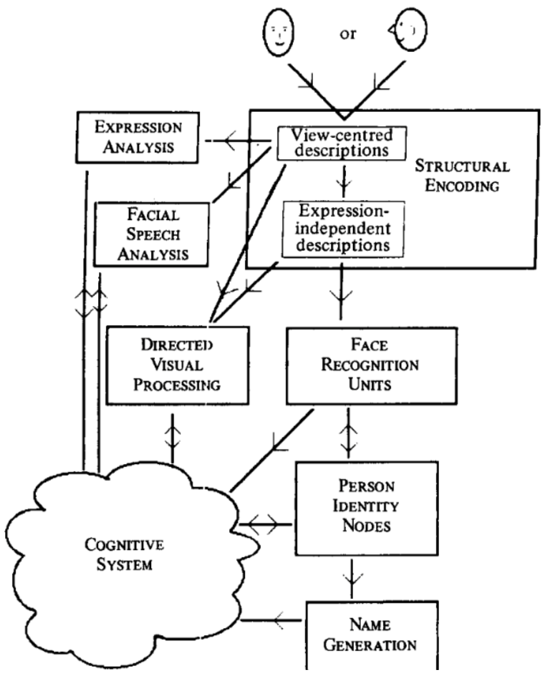

Figure 1. A functional model of face recognition taken from Bruce and Young (1986).

Bruce and Young (1986) discuss two common errors in face recognition and address them in terms of their proposed model for face recognition. People frequently know that a face is familiar to them but are unable to recall any information about the person. In their model this is explained as an activation of the face recognition unit giving rise to the feeling of familiarity without access to the person identity node. The other error involves the “tip-of-the-tongue” phenomenon where a person is recognized as familiar but their name cannot be retrieved. The authors explain this using their model by proposing the error results from activation of both the face recognition units and the person identity nodes in the absence of activation of the name generation node. Their model suggests that the butcher-on-the-bus phenomenon may be the result of activation of the face recognition units, in the absence of activation of both the person identity nodes, and the name generation node.

It is important to understand how a newly encountered face becomes familiar because there is a large discrepancy in the recognition of familiar and unfamiliar faces. It may be the case that structural information is not encoded as effectively for unfamiliar faces. Alternatively, the face recognition units may not be as refined for unfamiliar faces as they are for familiar ones. While judging two different images of an unfamiliar face depicting the same person is highly error prone, the same task is relatively effortless for familiar faces (e.g., Bruce et al., 1999). Recent literature shows that this is because each face varies in a unique way across different viewing conditions such as lighting, age, angle, and emotional expression. This variability leads people to mistakenly identify two images of the same unfamiliar person as different people, a mistake that does not occur for familiar faces (Jenkins et al. 2011).

Idiosyncratic variability has been neglected or intentionally excluded as noise in earlier face recognition research through the use of images that vary systematically (e.g., two images of the same face that only vary in emotional expression). However, recently there has been an increased interest in using ambient images that vary in unsystematic ways where viewing conditions are not controlled. Ritchie and Burton (2017) argued that within the context of face recognition unsystematic variability is crucial to the process of learning faces because for a face to become familiar the unique ways in which that face varies need to be learned. In line with this, Ritchie and Burton (2017) found that having greater exposure to a single face's range of variability enhances the ability to recognize that face. In their experiment subjects completed a learning phase where they were shown a series of photographs of each identity paired with that person's name. During the learning phase subjects learned half of the identities from high-variability images (i.e., taken from several different uncontrolled viewing conditions), and the remaining half were learned from low-variability images (i.e., screenshots taken from interviews). After training, subjects' memory for these faces was tested using new images. The authors found that faces which were learned from high variability images were recognized more accurately (Experiment 1A), and were recognized faster even when the test image closely resembled the low variability set (Experiment 1B). Additionally, when subjects were given a matching task where they were asked to indicate whether two images depict the same person, subjects were more accurate in matching two photos of a face learned from high variability images (Experiment 2). These findings demonstrate that exposure to high within-person variability leads to better learning of a newly encountered face.

The role of *context* in becoming familiar with a new face through exposure to within-person variability is not well established. However, there is existing research concerning the role of context in memory for words. Godden and Baddeley conducted two experiments examining the effect of context on memory. Although neither of these experiments involved face recognition, they both provide valuable insight into the role of context in recall and recognition.

Godden and Baddeley (1975) had divers learn lists of words either on land or underwater, and then asked them to recall the words in the same learning and recall environment (congruent condition) or a different learning and recall environment (incongruent condition). The authors found that recall performance was better in the congruent condition than in the incongruent condition suggesting that memory recall is context-dependent.

However, in another study, Godden and Baddeley (1980) had a group of divers again learn lists of words. On a given day groups of divers learned on land or underwater, and then were tested on land or underwater. Each diver eventually participated in each condition. The four conditions were: learn dry-recognize dry, learn dry-recognize wet, learn wet-recognize dry, and learn wet-recognize wet. The test was a recognition task where the participants were given a list of words and had to indicate which words they had seen in the previous list. The authors found that only a main effect of the original learning environment was significant. More words were recognized when learning took place on dry land as opposed to when learning occurred underwater. This was the case regardless of the test environment. The authors concluded that their results showed no evidence of a context-dependent effect. For this reason their results are consistent with the findings of previous research and support the notion that recognition memory is resistant to extrinsic context dependency. Based on their findings the authors suggest that intrinsic context, which refers to aspects of a stimulus that must be processed in order to comprehend it, affects recognition memory. This is because it directly determines what material is learned, and guides the participant back to the interpretation of the stimulus that occurred during acquisition. On the other hand, the relationship between extrinsic context and material learned is completely arbitrary. Therefore, the extrinsic context does not determine the interpretation of the material, and cannot add anything to more powerful cues presented by the physical presence of the words to be remembered.

While there have been experiments conducted to examine individual aspects of the butcher-on-the-bus phenomenon, to date there has been little research conducted with the sole purpose to better understand the phenomenon. One of the few studies related to this topic was conducted by Gruppuso et al. (2007). These authors sought to determine the relative contribution of conscious recollection and unconscious influences to the recognition of faces in an unfamiliar context. To accomplish this they presented participants with face-context picture pairs (such as that in Figure 2), allowed participants to study the pairings, and then had them complete a memory test.

Figure 2. An example of a stimulus used by Gruppuso et al. (2007).

During the study period participants were instructed to rate the likelihood that the person in the photo would be associated with the context. They had no knowledge that a memory test would follow. At test, participants were presented with faces (both studied and unstudied). Each face was paired with a context photo. In the *congruent context*, the context presented at test was the context paired with the face at study. In the *incongruent context*, the context presented at test was a context previously-paired with a different previously seen face. In the *new context*, a previously unseen context was presented alongside the face at test. During the test participants were asked to only judge whether the faces were old or new. If the face was judged to be old, then they proceeded to give a remember or know judgement. They were not asked about the context. Gruppuso et al. (2007) distinguish between knowing (i.e., the sense of familiarity with the stimulus) and remembering (i.e., the ability to retrieve the source alongside the information that is retrieved). The butcher-on-the-bus phenomenon is an example of knowing in the absence of remembering.

Gruppuso et al. (2007) found that studied faces were more likely to be remembered if they were encountered in the congruent context versus the incongruent context or the new context. It was the combination of a particular face being learned in a particular context that resulted in the highest recollection accuracy. It was not the case that reinstatement of any previously seen context increased recognition. Rather, they found that the effect of reinstating an incongruent context was detrimental to remembering, but had little effect on knowing. Reinstatement of an incongruent context was worse than presenting a new context for remembering. However, reinstatement of an incongruent context was no worse than presenting a new context for knowing. Their results suggest that a person’s response to seeing the same person in two different contexts is dependent on whether sufficient source-specifying information required to make the identification can be recalled. People are less likely to recollect source-specifying information when the butcher is seen on the bus (an atypical context), yet they are no less likely to experience a sense of familiarity.

The experiment conducted by Gruppuso and colleagues has several characteristics that limit its ability to adequately explain the butcher-on-the-bus phenomenon. First, it is not clear that the butcher-on-the-bus phenomenon can be studied using context-face pairings as done in the previous study. In all cases the relationship between the context and paired face was entirely arbitrary. In contrast, when the butcher-on-the-bus phenomenon occurs naturally, the relationship between the person and the learning context is not arbitrary. Secondly, a key feature of this phenomenon is that the target identity is learned within a context, and using a context-face pairing may not be sufficient to truly replicate the natural conditions of the phenomenon. Finally, participants in the aforementioned study were not exposed to within person variability which has been found to be an essential aspect of learning a new face (Ritchie & Burton, 2017).

I have attempted to study the underlying cognitive mechanisms of the butcher-on-the-bus phenomenon in a way that more closely resembles the way that it naturally occurs. In order to accomplish this, I have addressed the limitations of Gruppuso et al.'s (2007) study. Instead of presenting pairings of face and context, images depicting a person within a context will be displayed. In addition, the informativeness of the context has been manipulated. This allows for the examination of person recognition not only when there is an arbitrary relationship between person and context, but also when there is a meaningful relationship between person and context. Importantly, because of my research design participants will see multiple images depicting the same identity during training and thus will have exposure to within person variability.

This study was designed to examine the following question: Is recognition of faces in previously-unseen contexts influenced by the allocation of attention to previously-seen contexts? Attention to the context was manipulated through altering whether the context is informative or not. I hypothesized that if attention to previously-seen contexts interferes with recognition of a face in a new context then *remember* responses should vary across conditions. Specifically, I expected the single informative context to yield the lowest proportion of "remember" responses because the target should be strongly associated with the previous context, resulting in participants struggling to remember the target in a novel context; this is the typical situation that yields the 'butcher-on-the-bus' phenomenon. The presence of an informative learning context should strengthen the association between identity and context, and thus make remembering the same identity in a novel context more difficult. In contrast, the single uninformative context will yield an intermediate number of "remember" responses because the identity will not be strongly associated with the context. Finally, I expect the multiple uninformative context condition to yield the highest number of "remember" responses because no specific context has been associated with the identity. Additionally, if familiarity is unaffected by previously seen contexts then the proportion of *know* responses will not vary across conditions. This is because familiarity is automatic, as a result *know* responses should not be affected by the amount of attention allocated to the learning images.

Method

Participants

I required 29 participants based on a power analysis using custom software to detect a medium sized effect (Cohen's $f = .25$), 80% power, and an alpha of .05 in a one-way within-subject analysis of variance (ANOVA) with three levels. I recruited 36 University of Regina undergraduate students with normal or corrected-to-normal vision through the participant pool. All procedures were carried out in accordance with the Canadian Tri-Council Policy Statement on the ethical treatment of research participants and received approval from the University of Regina Research Ethics Board prior to its commencement. All participants signed a consent form prior to their participation, and received 1% bonus credit toward a psychology class as compensation. Data collected from all 36 participants were included for analysis (32 female, 4 male; age: $M = 22$ years, $SD = 5.51$).

Stimuli and Apparatus

There were 18 target identities included with 6 photos each, as well as 18 foil identities with one photo each. All photos were ambient images downloaded from the internet in which lighting, angle, and expression were not controlled. The width of the images was standardized to 500 pixels however the height of each image varied. There is no established precedent for the size of images being used as stimuli when the background information is important, as context has not typically been included in face research stimuli. For the purposes of my study it was necessary to crop the images in such a way that enough of the context was present, while other people besides the intended identity were not included. Further efforts to make all images the same physical size would have required breaking the aspect ratio for many of the images. This would have led to stretching or distorting of the faces and their features. Ultimately, I deemed maintaining the aspect ratio to be of greater importance than having each photo be the same size. This decision was due to the fact that face processing is a high-level task and therefore not heavily dependent on the size of the retinal image. The identities used in the study came from social media and were chosen to be unfamiliar to the Canadian participants. The experiment was programmed using jsPsych (de Leeuw, 2015) and ran on Pavlovia.org, which is a platform for running psychology experiments online. The data were analyzed using JASP (Version 0.14.1; JASP Team, 2023).

Procedure

Each participant completed a training phase followed by a test phase. During the training phase participants were shown a random sequence of 90 target images divided into five blocks of 18 images with a break between blocks. Participants were asked to rate the distinctiveness of the face shown in each photo using a slider placed at the bottom of the screen to ensure that attention was being devoted to each face. This task also served to ensure that participants did not intentionally try to memorize the identities they were shown during the training phase. Distinctiveness was rated from very typical to very distinctive.

The three within-subjects conditions used were as follows: single informative context, single uninformative context, or multiple uninformative context (control). In the single informative context each target person was presented in a single context that could provide identity-specific information such as their occupation or hobbies. I expected the informativeness of the context to increase the amount of attention allotted by participants. In the single uninformative context each target person was presented in a single context that does not provide any identity-specific information. I expected this condition to not draw participants' attention to the context because of the lack of information conveyed by it. In the multiple uninformative context each target person was presented in five different contexts. Similarly, I expected this condition to not draw participants' attention to the context due to the lack of informativeness. Six targets were assigned to each of the three conditions.

In the test phase subjects were shown a target or a foil identity in a novel context and instructed to make a *remember*, *know*, or *new* judgement. Subjects were asked to respond "remember" if they had a specific recollection of the person and to respond "know" if the target felt familiar, but they did not have a specific recollection of the person. Additionally, subjects were asked to respond "new" if they were sure that the person was not shown during training. After giving their responses subjects were asked which of the previously seen contexts they most strongly associated with the identity. During the test phase photos of the 18 target identities were interspersed among photos of 18 foil identities, with one photo per identity. These 36 photos were presented in a random order. All photos presented during the test phase were previously unseen and depicted the identities in a novel context.

Results

The analysis consisted of both classic and Bayesian ANOVA. The Bayes factor was used to assess how well the observed results were predicted by both the null and alternative hypothesis. A $BF_{10} = j$ would indicate that a given set of results are j times more likely to occur under the alternative hypothesis (Schönbrodt & Wagenmakers, 2018). Conversely, a $BF_{01} = k$ indicates that the observed results are k times more likely to occur under the null hypothesis.

Three *d'* scores, a measure of discriminative power, were calculated for each subject, one score for each of the three conditions using the Hautus (1995) correction. The Hautus correction was used to avoid incidents of a false alarm rate of zero leading to an undefined *d'* score. These scores were then analyzed using a repeated measures ANOVA to determine if there was a significant effect of condition on *d'*. This test indicated a significant difference between at least two conditions, $F(2,70) = 13.826$, $MSE = 0.451$, $p < 0.001$, $\eta_p^2 = 0.283$, with substantial evidence in favour of this effect, $BF_{10} = 2617.171$. The descriptive statistics regarding this analysis are displayed in Table 1 and the average *d'* score by condition is illustrated in Figure 3.

Table 1

d' Descriptive Statistics

Condition	N	Mean	SD	SE	Coefficient of Variation
Multiple Uninformative	36	1.089	0.928	0.155	0.852
Single Informative	36	1.717	0.914	0.152	0.532
Single Uninformative	36	0.930	0.726	0.121	0.781

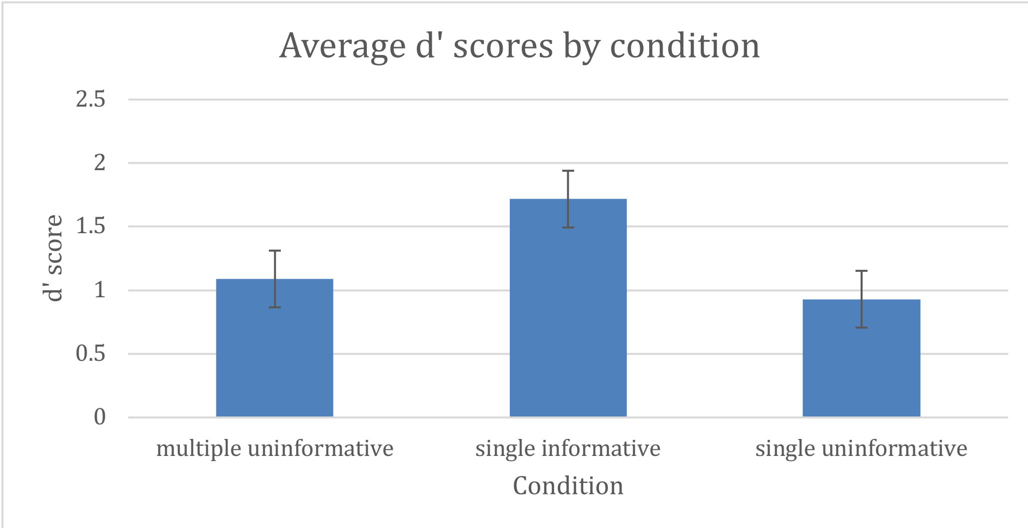

Figure 3. Average *d'* scores by conditions. Error bars represent 95% confidence interval (Masson & Loftus, 2003)

A follow-up test using Bonferroni correction as a pairwise comparison was conducted with the results shown in Table 2. This indicated a significant difference between the multiple uninformative condition and single informative condition $M_{difference} = -0.628$, $t = -3.967$, $p < 0.001$. As well, there was a significant difference between the single informative condition and single uninformative condition, $M_{difference} = 0.787$, $t = 4.973$, $p < 0.001$.

Table 2

Pairwise Comparison of d' Scores Using Bonferroni Correction

Post Hoc Comparisons - Conditions						
		Mean Difference	SE	t	Cohen's d	P_{bonf}
Multiple Uninformative	Single Informative	-0.628	0.158	-3.967	-0.729	< 0.001
	Single Uninformative	0.159	0.158	1.006	0.185	0.953
Single Informative	Single Uninformative	0.787	0.158	4.973	0.914	< 0.001

Note. p-value adjusted for comparing a family of 3.

Distinctiveness ratings were analyzed as a way to determine whether there was a potential confound of identity. It has previously been found that distinctive faces tend to be more memorable (Valentine, 1991). Due to different identities being assigned to the various conditions it was necessary to check whether identities assigned to a given condition were consistently rated as being more distinctive. Distinctiveness ratings corresponding to the target identities were averaged across conditions and then analyzed via repeated measures ANOVA. Mauchly's test of sphericity indicated that the assumption of sphericity was violated ($\varepsilon = 0.701$). As a result, the Greenhouse-Geisser (1959) correction was applied (Girden, 1992). This analysis did not show a significant difference between the average distinctiveness ratings across conditions $F(1.402, 49.079) = 0.651$, $MSE = 154.921$, $p = 0.473$, $\eta_p^2 = 0.018$, with evidence supporting the absence of an effect $BF_{01} = 6.884$. The descriptive statistics for distinctiveness are shown in Table 3.

Table 3

Distinctiveness Descriptive Statistics

Condition	N	Mean	SD	SE	Coefficient of Variation
Multiple Uninformative	36	48.331	12.862	2.144	0.266
Single Informative	36	48.139	20.397	3.399	0.424
Single Uninformative	36	45.812	17.397	2.900	0.380

A repeated measures ANOVA was performed to compare the effect of condition on the proportion of "know" responses. There were no significant differences in the proportion of "know" responses among conditions $F(2,70) = 1.340$, $MSE = 0.016$, $p = 0.268$, $\eta_p^2 = 0.037$, $BF_{01} = 3.898$. Descriptive statistics for this analysis are shown in Table 4.

Table 4

Proportion "know" Descriptive Statistics

Condition	N	Mean	SD	SE	Coefficient of Variation
Multiple Uninformative	36	0.116	0.153	0.026	1.325
Single Informative	36	0.074	0.146	0.024	1.972
Single Uninformative	36	0.116	0.173	0.029	1.493

A repeated measures ANOVA was performed to compare the effect of condition on the proportion of “remember” responses. There was a significant difference among the conditions on the proportion of “remember” responses $F(2,70) = 17.587$, $MSE = 0.031$, $p < 0.001$, $\eta_p^2 = 0.334$, $BF_{10} = 21644.132$. Table 5 displays the associated descriptive statistics.

Table 5

Proportion “remember” Descriptive Statistics

Condition	N	Mean	SD	SE	Coefficient of Variation
Multiple Uninformative	36	0.556	0.316	0.053	0.569
Single Informative	36	0.694	0.297	0.049	0.427
Single Uninformative	36	0.449	0.306	0.051	0.680

The Bonferroni correction was then used as a follow-up pairwise comparison and indicated a significant difference between all three conditions at the level of α = 0.05. The mean difference between each comparison and the associated significance is shown in Table 6.

Table 6

Pairwise Comparison of Proportion "remember" Using Bonferroni Correction

Post Hoc Comparisons - Condition

		Mean Difference	SE	t	P_{bonf}
Multiple Uninformative	Single Informative	-0.139	0.041	-3.347	0.004
	Single Uninformative	0.106	0.041	2.566	0.037
Single Informative	Single Uninformative	0.245	0.041	5.914	< 0.001

Note. P-value adjusted for comparing a family of 3.

The proportion of "know" and "remember" responses by condition were then examined in relation to the overall proportion of hits in each condition as displayed in Figure 4.

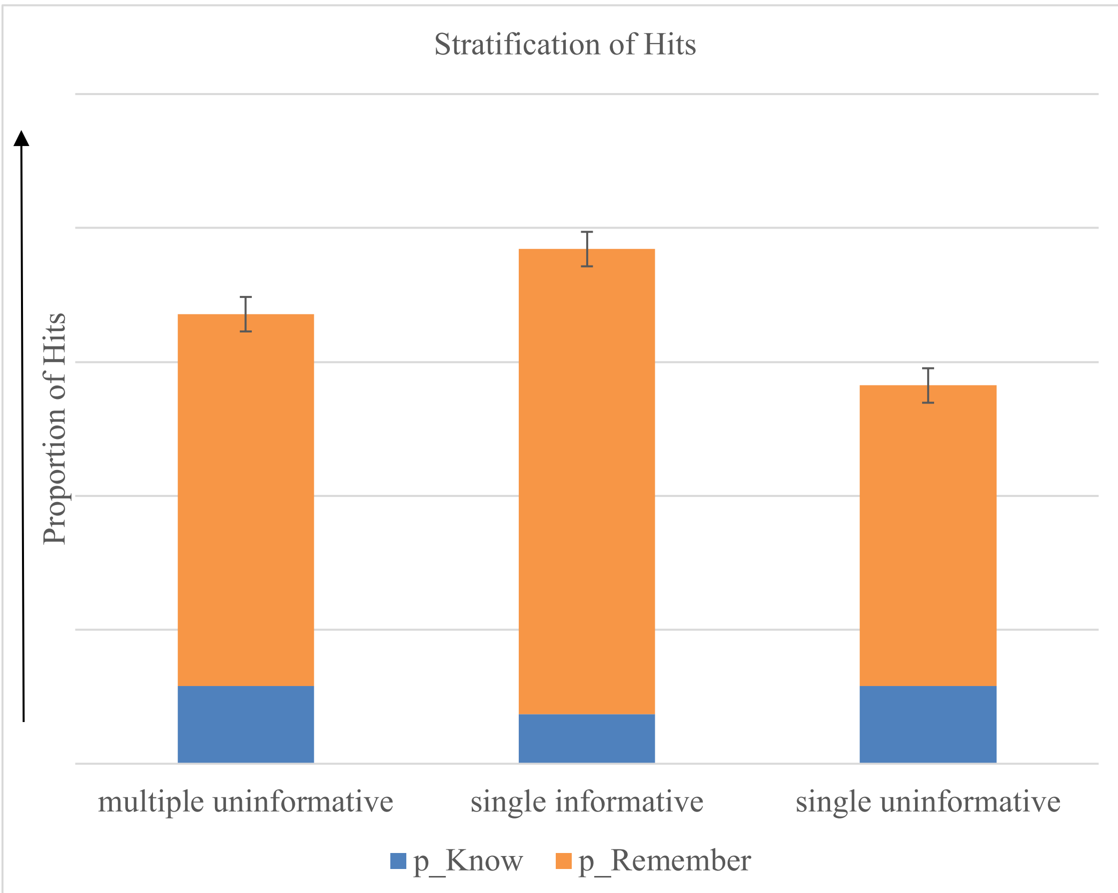

Figure 4. Shows the overall proportion of hits by category as well as the stratification of hits into "remember" and "know" responses. Error bars represent 95% confidence interval for the overall proportion of hits (Masson & Loftus, 2003).

Discussion

The observed proportion of "know" responses was in line with my hypothesis that "know" responses would not significantly differ across conditions due to the automaticity of familiarity. However, the results concerning the proportion of "remember" responses were not as previously expected. While it was true that the proportion of "remember" responses varied by condition, the variation was contrary to my expectation. The highest proportion of "remember" responses was observed in the single informative condition, which I previously hypothesized would have the lowest proportion. I made this prediction under the expectation that repeated exposure to an identity in the same informative context (e.g., butcher in the market) would interfere with the ability to remember the same identity in a novel context (e.g., butcher on the bus). Interestingly my results suggest that repeated exposure to an identity in a single informative context actually improves a person's ability to recognize the same identity in a novel context.

These unexpected results cannot be explained by a systematic difference in the distinctiveness of identities across conditions. A possible explanation for this unexpected finding is that the informativeness of the context is drawing more attention to the identity within the context. As a result, the identity is then better encoded in memory. The presence of an informative context may provide a powerful top-down cue which increases the efficiency with which exposures to the same face can be aggregated. The benefit of this top-down cue appears to surpass the benefit of exposure to within-person variability. This notion is supported by lower sensitivity in the multiple uninformative condition which included high within-person variability.

There is also another possible explanation for why I was unable to induce the butcher-on-the-bus phenomenon and instead found that the consistency and informativeness of the context reduced its presence. It may be the case that in order for the butcher-on-the-bus phenomenon to occur attention needs to not be paid to the face while still being devoted to the context. Due to the training task participants were given it was necessary for them to pay attention to each face. My intention was to create a strong sensation of the butcher-on-the-bus phenomenon. However, it is possible that the task I presented participants with actually eliminated the butcher-on-the-bus phenomenon.

It would be beneficial for a future study to employ a training task that does not directly draw attention to the face. Additionally, future studies should aim to incorporate a rotation of identities between conditions such that targets would be randomly divided into three sets. For one subject, each set of targets would be assigned to one condition. For another subject, a different set of targets would be assigned to one condition. This assignment of targets to conditions would then rotate through the remaining subjects such that, across all subjects, each target set would be assigned to each condition an equal number of times. This would ensure that there is no confound between target identity and condition. Originally, I planned to implement such a rotation of identities, although ultimately due to time constraints for collecting stimuli the planned rotation of target identities within the three conditions was not possible.

The importance of context in face recognition research has typically been neglected and as a result the role of context has been under-explored. Although my results were not as expected they still demonstrate that context has an important role in face recognition. As a result, it is important to begin introducing context into research regarding face recognition, as in the real world it is seldom the case that face recognition occurs without context. Furthermore, it is important to understand the conditions in which recognition of people seen on multiple occasions is accurate. There are possible implications extending to suspect identification in legal proceedings and careers that involve recognizing unfamiliar faces (e.g., police, security, border services).

References

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

Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J., Burton, A. M., & Miller, P. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied*, *5*(4), 339. https://doi.org/10.1037/1076-898X.5.4.339

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

Girden, E. R. (1992). *ANOVA: Repeated measures* (No. 84). sage.

Godden, D. R., & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and underwater. *British Journal of Psychology*, *66*(3), 325-331. https://doi.org/10.1111/j.2044-8295.1975.tb01468.x

Godden, D., & Baddeley, A. (1980). When does context influence recognition memory?. *British journal of Psychology*, *71*(1), 99-104. https://doi.org/10.1111/j.2044-8295.1980.tb02735.x

Greenhouse, S.W., Geisser, S. On methods in the analysis of profile data.*Psychometrika* **24**, 95–112 (1959). https://doi-org.libproxy.uregina.ca/10.1007/BF02289823

Gruppuso, V., Lindsay, D. S., & Masson, M. E. (2007). I'd know that face anywhere!. *Psychonomic Bulletin & Review*, *14*(6), 1085-1089. https://doi.org/10.3758/BF03193095

Hautus, M. J. (1995). Corrections for extreme proportions and their biasing effects on estimated values of d′. *Behavior Research Methods, Instruments, & Computers*, *27*, 46-51. https://doi.org/10.3758/BF03203619

JASP Team (2023). JASP (Version 0.17.3)[Computer software].

Jenkins, R., White, D., Van Montfort, X., & Burton, A. M. (2011). Variability in photos of the same face. *Cognition*, *121*(3), 313-323. https://doi.org/10.1016/j.cognition.2011.08.001

Mandler, G. (1980). Recognizing: The judgment of previous occurrence. *Psychological Review*, *87*(3), 252–271. https://doi.org/10.1037/0033-295X.87.3.252

Masson, M. E., Loftus, G. R. (2003). Using confidence intervals for graphically based data interpretation. *Canadian Journal of Experimental Psychology/Canadian Journal of Experimental Psychology*, *57*(3), 203. https://doi.org/10.1037/h0087426

Schönbrodt, F. D., & Wagenmakers, E.-J. (2018). Bayes factor design analysis: Planning for compelling evidence. *Psychonomic Bulletin & Review*, *25*(1), 128–142. https://doi.org/10.3758/s13423-017-1230-y

Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology*, *43*(2), 161-204. https://doi.org/10.1080/14640749108400966