Psychonomic Bulletin & Review (2021) 28:2027–2034
https://doi.org/10.3758/s13423-021-01965-2

BRIEF REPORT

And like that, they were gone: A failure to remember recently attended unique faces

Joyce Tam[1] · Michael K. Mugno[1] · Ryan E. O'Donnell[1] · Brad Wyble[1]

Accepted: 27 May 2021 / Published online: 8 July 2021

Abstract

Attribute amnesia (AA) is a phenomenon in which participants have difficulty answering an unexpected question about an attended attribute of the most recent target stimulus. A similar situation can occur in cases of real-life eyewitness identification when the eyewitness did not explicitly try to remember the alleged perpetrator's face despite having attended to it. We found that AA is generalizable to novel faces, such that when participants were unexpectedly asked to identify a face, performance was poor, even though they had just attended to that face seconds ago (N = 40 each in an initial experiment and its replication). This finding shows that unexpected face identification is inaccurate even when the face had just been attended to and suffered minimal decay and interference, implying that AA can explain some cases of failure of eyewitness identification that cannot be attributed to a lack of attention or post-event interference.

Keywords Attribute amnesia · Eyewitness identification · Face recognition · Working memory

Introduction

Attribute amnesia (AA) is a phenomenon in which people are poor at reporting an unexpectedly probed attribute, even if it had just been attended to seconds ago (H. Chen & Wyble, 2015a, 2016).[1] Generalizing this finding to novel faces can help us understand face identification failure in real-life situations, particularly in eyewitness identification, where witnesses sometimes do not realize they will be asked to identify the potential perpetrator at the time of witnessing the crime. In the current study, we found AA with a set of unique and novel face stimuli, such that having attended to a face did not guarantee the formation of a reportable memory representation of its identity, even if the report was requested only several seconds later. This observation suggests that attention to facial attributes does not necessarily lead to accurate face identification, and AA may partly explain inaccuracies in eyewitness identification.

In a typical AA experiment, an attribute, such as the identity of a letter, must be attended in order to detect a target and report its location, but participants do not expect to report the attribute itself. At some point in the experiment, they are given a surprise question about the attribute. For example, H. Chen and Wyble (2015a) asked participants to report the location of a letter presented among three digits. Participants completed 155 of these location trials before unexpectedly being probed to report the identity of the letter they had just seen on the 156th trial. Although their performance during the location trials was near ceiling, participants were generally unable to report the letter's identity on the surprise trial. This is surprising according to conventional understandings of working memory and attention because being able to detect the letter requires attentional selection based on its orthography. Such attentional focus is often linked to the creation of a working-memory representation (e.g., Chun, 2011; Cowan, 2001; Kiyonaga & Egner, 2013; Postle, 2006; Roelfsema, 2005). In addition, AA has also been generalized to pre-surprise tasks other than location report (e.g., H. Chen et al., 2016; H. Chen & Wyble, 2016). For example, H. Chen and Wyble (2016) similarly asked participants to find a letter among three digits

[1] The use of the term "amnesia" in the current paper and in previous AA studies reference Wolfe's (1999) proposal of "inattentional amnesia," which describes when unattended visual information is "seen, but will be instantly forgotten" (p.73) because vision does not necessarily start the memory-encoding process. AA suggests that even with attention, vision does not necessarily start the memory-encoding process either.

✉ Joyce Tam
joyce.m.y.tam@gmail.com

[1] Department of Psychology, Pennsylvania State University, University Park, PA 16801, USA

but instructed them to report the letter's color in the pre-surprise trials. However, a surprise trial of the letter's identity was still near the chance level. Taken together, AA demonstrates that selection based on a certain attribute is not sufficient for the information of that attribute to be successfully remembered.

AA has been replicated and extended in multiple studies (e.g., Born et al., 2020; H. Chen et al., 2016; H. Chen & Wyble, 2015b, 2016; H. Chen et al., 2019; W. Chen & Howe, 2017; Y. V. Jiang et al., 2016; Swan et al., 2017; Wyble et al., 2019), but most studies used simple, repetitive stimuli (typically letters) as targets. We studied whether AA is generalizable to a set of unique face stimuli, which have substantial ecological, cultural, and societal relevance. An example of where a phenomenon such as AA could occur in a real-life example is when an eyewitness encountered a perpetrator but did not realize in that moment that a crime was occurring or was about to occur. In such a case, an eyewitness might have attended to a perpetrator's face without the intention to encode it for future identification. Previous studies have shown that various attentional failures reduced the chance of noticing a crime and thus impaired accuracy of subsequent reports about that event's details (see Cutler et al., 1987; Hyman et al., 2018; Loftus, 2005). For example, participants might have been focused on another demanding task (inattentional blindness; Chabris et al., 2011; Rivardo et al., 2011), faced sudden visual interruptions (change blindness; Davis et al., 2008), or been distracted by a salient item like a weapon (the weapon focus effect; Fawcett et al., 2011; Kocab & Sporer, 2016; Steblay, 1992). These studies show that attention is necessary for a reportable memory to be formed. AA addresses the other side of the relationship between attention and memory and suggests that having paid attention is not sufficient for accurate eyewitness memory even in the short term.

There are reasons whereby novel faces might be expected to be immune to AA. Multiple lines of research suggest that humans preferentially process faces over other types of stimuli (see Palermo & Rhodes, 2007). For example, upright faces required less time to be perceived through continuous flash suppression (CFS) than inverted faces, suggesting that naturally oriented faces reach awareness more readily (Y. Jiang et al., 2007; Stein et al., 2012), and identity-specific face representations can be formed in as short as 75 ms (Tanaka, 2001). Task-irrelevant faces can also capture attention and impair task performance (Bindemann et al., 2005; Lavie et al., 2003). Moreover, researchers comparing short-term memory for faces and other real-world objects estimated that faces enjoy higher capacity limit (Curby & Gauthier, 2007) or memory resolution (Scolari et al., 2008). It would thus be instructive to determine whether a "face advantage" can be observed in an AA experiment, i.e., whether face identification accuracy in the surprise trial would be as high as the following control trial.

Another reason to doubt that AA would be observed for novel faces is that some classes of novel stimuli had been shown to be immune from AA (H. Chen et al., 2019; W. Chen & Howe, 2017). For example, when H. Chen et al. (2019) used a set of unique, non-repeating line-drawing stimuli as targets, the surprise trial accuracy was as high as the following control trials (H. Chen et al., 2019). This moderation effect from target repetition with similar pictorial stimuli was also demonstrated earlier by W. Chen and Howe (2017) with a set of colored pictures of real-life objects. In contrast to both studies, letter targets exhibit AA whether they repeat across the experiment or not (H. Chen & Wyble, 2016). These results together show that target repetition and the type of target stimuli interactively affect AA. Thus, the use of unique faces in the current study also tests whether the absence of AA with unique targets is extendable to another set of stimuli.

Methods

In the current study, we investigated AA with photorealistic, artificially generated faces from the NVIDIA StyleGAN (Karras et al., 2019). The faces from this GAN were sampled with algorithms developed by Baylies (2020) and refined by Mugno et al. (2020), which can be downloaded from https://osf.io/5cqw8/. This face generator allowed us to build a big enough set of face stimuli such that each face would appear only once in the experiment. The paradigm was also modified relative to the original AA study (H. Chen & Wyble, 2015a) to reduce the difficulty of the surprise test by using a longer and unmasked search array (see Methods), which should allow information to be encoded effectively (Vogel et al., 2006; Woodman & Vogel, 2008). The surprise question was also shorter, reducing the interference from extra information at retrieval (Makovski et al., 2008; Souza et al., 2016). These modifications made it unlikely that report failures on the surprise trial are due to capacity limitations in encoding and retrieval.

To foreshadow our key result, despite these modifications to make the surprise test simpler, AA was present for unique faces. Only about half of the participants were able to report the identity of the target face on the surprise trial, even though they were able to locate the target at near ceiling levels in the pre-surprise trials and could easily report the target identity in the control trials after the surprise.

Participants

The initial experiment and the replication had separate sample sizes of 40 participants each, doubling that of H. Chen and Wyble's (2015a) original study. The initial sample size was set to detect a minimum of 30% accuracy increase from the surprise trial to the first control trial (see Design).[2] Note that the findings were replicated using the same sample size. Participants completed the experiment on their personal computers as the experiment was hosted online. For the initial experiment, 41 participants were recruited from the Penn State University subject pool in exchange for course credits. Forty-five additional participants were recruited from Prolific.co for the replication, and each was compensated with US$1.11. After excluding participants who performed near chance (≤ 25%) in either the pre-surprise location trials or the control trials, data from 40 participants in each experiment were analyzed (initial experiment: M_{Age} = 19.3 years, SD = 1.57 years, 26 female; replication experiment: M_{Age} = 24.8 years, SD = 6.07 years, 19 female). The experiments were approved by the institutional review board at Penn State University.

Stimuli

The faces used in the study were generated from a modified version of StyleGAN, a generative adversarial network developed by Karras et al. (2019). Each face is completely artificial and uniquely generated within a space of 512 possible dimensions using the algorithms developed by Mugno et al. (2020). Faces were further selected such that they do not have highly distinguishing features like glasses. Each face was given a centered circular crop (diameter = 250 px) with a Gaussian blur radius of 10 px. The circular crop removes backgrounds and other outstanding attributes such as the hair and the shirt collar. The resultant images were then resized for the experiment (see below). There were 47 young faces and 96 adult faces selected, with the young serving as targets and the adult as distractors; no face was repeated in the experiment (Fig. 1).

Design

The trial sequence is shown in Fig. 2. Each trial began with a black fixation cross (height = 30 px) in the center of a white background, displayed for a random interval between 800 and 1,800 ms. This was followed by the search array lasting 2,000 ms, where one young face and three adult faces (diameter = 225 px) were placed in a square arrangement, each centered 198 px from the fixation cross. Note that the exact stimulus sizes and positions depend on the resolution of the participant's personal computers. The target and distractor locations randomly varied from trial to trial. This was followed by a blank delay of 500 ms. In the first 27 trials, a location screen followed the delay and assigned numbers (1, 2, 3, 4) that correspond to the locations of the faces in the search array. Participants were asked to input the number corresponding to the location of the target young face and were promptly given feedback.

On trial 28 (the surprise trial) an unexpected task was presented following the delay. The target face along with three unseen young faces from the same pool were randomly arranged in a horizontal line, with corresponding numbers (5, 6, 7, 8) below them. Participants were asked to input the number that corresponds to the target face they had just seen on the previous search array. After this surprise question, they were asked to report the location using the same prompt that they were accustomed to. Trials 29–32 repeated this new pattern, acting as controls. Note that the exact wording of the questions can also be found in Fig. 2. Due to experimenter error, one adult face was included with the young face set, meaning that there was a 2.1% (1/47) chance that this adult face was the target face in the surprise trial and a 6.4% chance (3/47) that this adult face was one of the foils in the surprise identity question. Note that these probabilities also apply to the first control trial, and thus the validity of comparing the surprise and first control accuracy was not compromised. In the replication experiment, this adult face was replaced with a young face. This adult face was never included in the adult face set.

Results

The principal results of the initial experiment are summarized in Table 1A. The pre-surprise accuracy was 92%, showing that participants were able to identify and locate the youngest face easily. However, the accuracy of the surprise identity trial was 50%, meaning that only 20 out of the 40 participants answered it correctly. In contrast, the accuracy in the first control trial was 90%. To test whether accuracy improved from the surprise to the first control trial, we employed the pairwise trinomial sign test[3] (Bian et al., 2009) which signified a highly significant result, n_+ (number of participants who answered the surprise

[2] In detail, we ran the trinomial test (Bian et al., 2009; see Results for the parameters it considers) with simulated data of sample sizes starting from N = 10, with 10-sample increments. For each sample size, data were simulated to generate the most conservative estimate of statistical significance such that N_d/N = 0.3 and n_- approaches its maximum possible value. At N = 40, with N_d = 12, n_+ = 26, n_- = 14, simulated p = .040.

[3] Previous AA studies primarily used the chi-square test to assess the significance of the difference between single-trial accuracies. However, in the surprise versus first control accuracy comparison, each participant contributes to two cells of the chi-square contingency table, which violates the independence assumption of the test (McHugh, 2013). Such paired observations are better assessed with a sign test. We used the trinomial test, which is an iteration of the binomial test that takes into consideration the presence of ties (Bian et al., 2009). The test statistics, N_d is given by $|n_+ - n_-|$ and accompanied by a bootstrapped 95% CI. The decision to use the trinomial test was made without first performing a chi-square test.

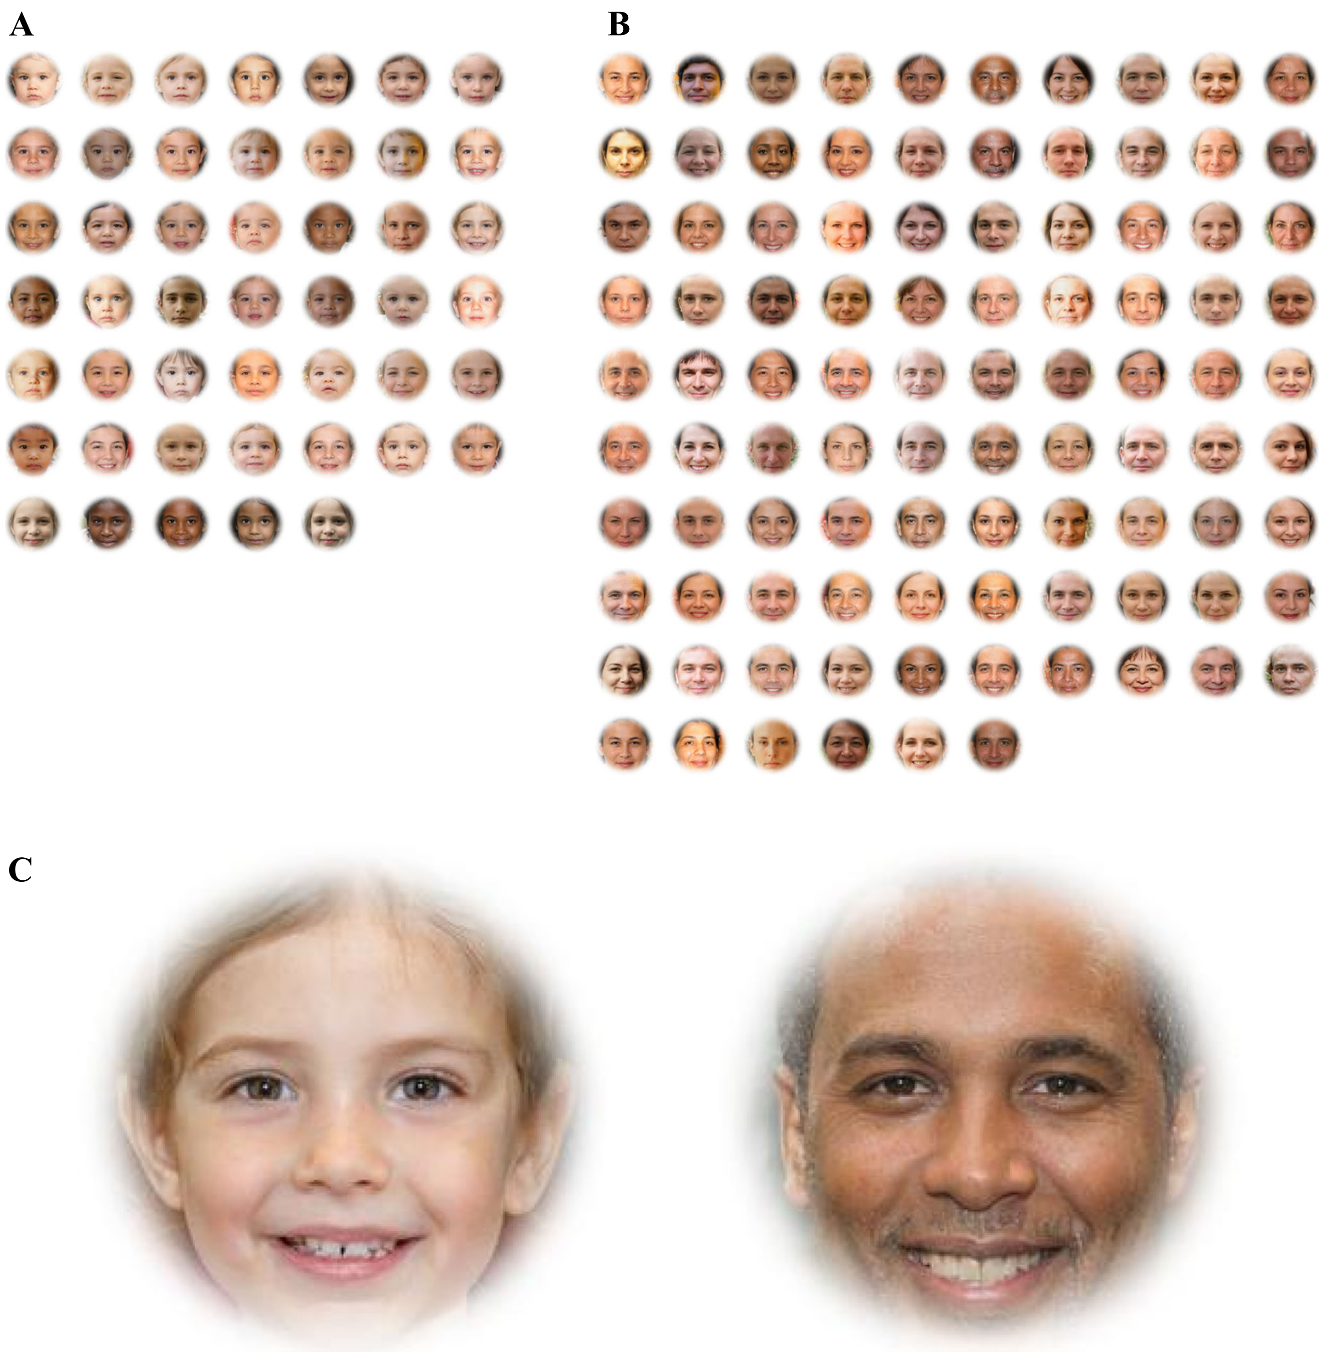

Fig. 1 Stimuli used in the experiments. Each face appeared only once in each experiment. **a** Young faces used as targets in all trials and foil options in the identity task in trials 28–32. The set after replacing the erroneously included adult face is shown (see main text). **b** Adult faces used as distractors in the search arrays. **c** Enlarged figure of one of the young faces (left) and one of the adult faces (right). All of these faces are generated artificially using an adversarial network

question incorrectly but the first control question correctly) = 17, n_- (number of participants who answered the surprise question correctly but then answered the first control question incorrectly) = 1, n_0 (number of participants who answered both questions correctly or both incorrectly) = 22, N_d = 16, $p < .001$, 95% CI [9, 23].

Results from the replication study agreed (Table 1B). Pre-surprise accuracy was 95%. Only 24 out of the 40 participants answered the surprise identity question correctly (60%), which was significantly lower than the number of people who got the first control question correct (83%), n_+ = 12, n_- = 3, n_0 = 25, N_d = 9, p = .0138, 95% CI [2, 16].

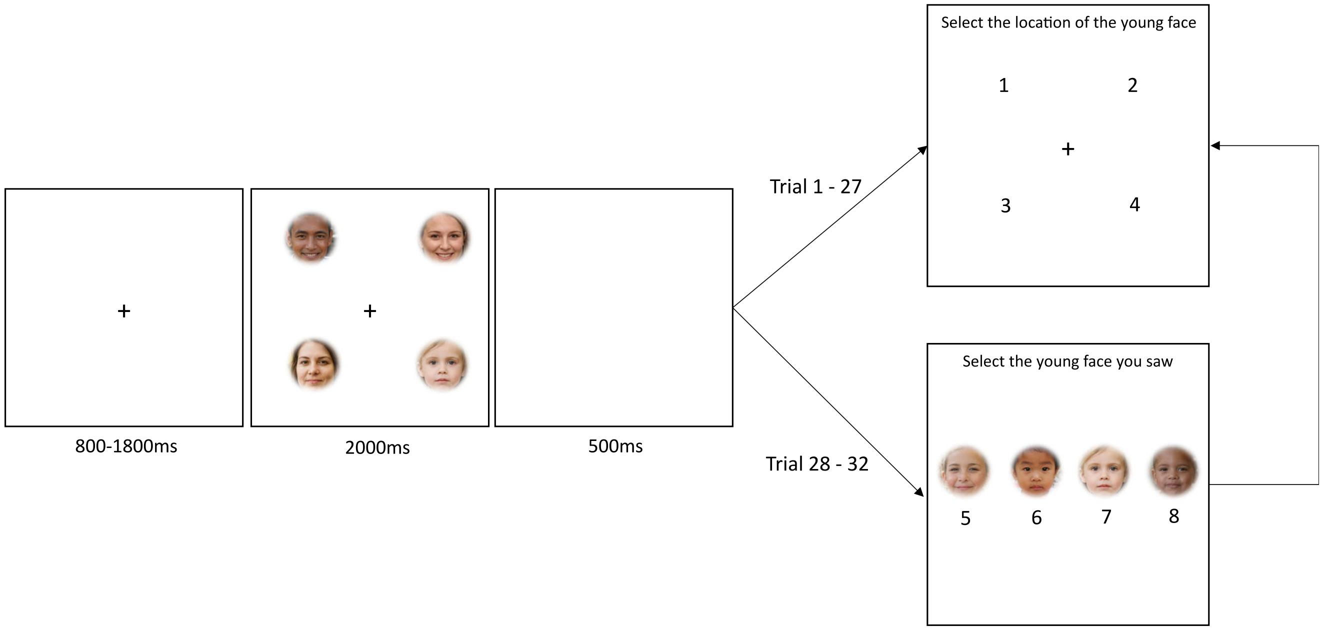

Fig. 2 Schematic representation of the experiment. After a fixation with a random duration between 800 and 1,800 ms, participants were presented with four faces for 2,000 ms, during which they were to memorize the location of the youngest face (bottom right in the current example). After a delay of 500 ms, they would see the text “Select the location of the young face” and report the location by pressing a number key. On the surprise trial and the control trials, they were probed with the text “Select the young face you saw” to select the target face they had just seen among four young faces before completing the location task. The figure is not drawn to scale

Discussion

We reported and replicated the first observation of AA with unique faces. Participants performed near ceiling in the pre-surprise trials during which they reported the location of the youngest face among distractor adult faces. However, when they were unexpectedly asked to report the identity of the young target face on a surprise trial, accuracy was 50% and 60%, respectively, suggesting that having attended to and selected the face did not produce a reportable memory representation in the surprise trial, even if the surprise question was presented only several seconds after the search array. Importantly, the identity question accuracy increased in the first subsequent control trial, showing that participants could memorize the face identity accurately if they expected its report.

Table 1 Summary of results. The reaction time (RT) of the surprise trial is reported separately for participants who answered the question accurately (A) and inaccurately (IA). Otherwise, only reaction times of correct trials are shown. (**A**) Results from the initial experiment. (**B**) Results from the replication

		Pre-surprise	Surprise	Control 1	Control 2	Control 3	Control 4
A							
Location	Accuracy	92%	78%	92%	93%	95%	88%
	RT (s)	1.00	3.90	1.47	0.97	1.13	0.85
Identity	Accuracy	NA	50%	90%	95%	93%	90%
	RT (s)	NA	6.72 (A) 6.70 (IA)	2.40	2.74	2.22	2.37
B							
Location	Accuracy	95%	60%	95%	95%	83%	100%
	RT (s)	0.83	3.95	1.47	1.21	0.90	0.93
Identity	Accuracy	NA	60%	83%	88%	93%	98%
	RT (s)	NA	6.96 (A) 10.97 (IA)	2.90	2.67	2.14	2.07

NA not applicable

AA with face targets

The current study shows that AA can occur with face stimuli that have been theorized to be preferentially detected, identified, and remembered (Curby & Gauthier, 2007; Y. Jiang et al., 2007; Scolari et al., 2008; Stein et al., 2012; Tanaka, 2001), and possibly capture attention (Bindemann et al., 2005; Lavie et al., 2003). Furthermore, the results have potential bearings on failures of eyewitness identification. Witnesses who attempt to identify the perpetrator from a line-up may not have expected to be required to do so at the time of witnessing the crime, analogous to the surprise identity task in the current study. Eyewitness identification is deemed more believable if the witnesses report that they paid attention during the crime (Bradfield & Wells, 2000; Palmer et al., 2013), and previous research has indeed shown that attentional failures impede awareness of and knowledge about the crime (see Cutler et al., 1987; Hyman et al., 2018; Loftus, 2005). Our data further show that having paid attention is not always sufficient for identification of a face, even at a short encoding-test latency: an unexpected identification task yielded poor accuracy for a face that was attended several seconds ago. Thus, the expectation of memory report should be explicitly controlled or manipulated in related studies (which has been shown to not be the case; Baldassari et al., 2021), as encoding cannot be assumed from participants having paid attention in such studies. It is also important to note that this gap between attention and intention of encoding is more observable when participants are entrenched in a cover task, such as in the current design where participants completed 27 pre-surprise location report trials before the surprise identity report trial. Report of unexpected attributes tends to be better, though still impaired at the start of an experiment (Eitam et al., 2013) and develops over the first few dozen trials (Wyble et al., 2019). Future research will also be required to study AA under more realistic, crime-scene-like contexts.

Factors other than attentional failures also affect the accuracy of eyewitness identification (Wells, 2002). For example, memory might be biased by misleading information encountered between the time of memory formation and report (Loftus, 2005). However, such post-event misinformation is absent in our case as the identity report was presented after a 500-ms blank delay that followed the search array. Unconscious transference (falsely identifying an innocent bystander; Loftus, 1976) should also play a minimal role in the current study, as the foil options were always novel faces that had not appeared in the experiment before. However, we cannot rule out the possibility that pre-event misinformation contributed to the identification failure (Eakin et al., 2003), for example, via erroneous conjunctions of features from different faces seen before the surprise trial (Reinitz et al., 1992). In addition, our subjects were provided with a forced-choice task and had no opportunity to opt out of choosing a face. Therefore, our findings do not distinguish between failures to retrieve any information about the facial identity (i.e., low confidence, or guessing), or retrieving an inaccurate memory with high confidence.

The current AA effect was also observed with a paradigm that minimized the chance that report failures on the surprise trial were due to capacity limitations of encoding or retrieval as the surprise question was made easier compared to the original study by H. Chen and Wyble (2015a). In addition to the search array being presented for 2,000 ms (compared to 150 ms in the original), the current paradigm eliminated masks after the search array. If participants were encoding the target face identity, they should have been able to do so, as the presentation time was well above the time estimated for detecting an identity-specific face (Tanaka, 2001) and remembering at least two different faces (Curby & Gauthier, 2007). Also, a much shorter, five-word surprise question was used compared to the 26-word question in the original AA paper, thus interference at test should have been reduced (Makovski et al., 2008; Souza et al., 2016; Tabi et al., 2019). Despite these modifications to the procedure, AA was still present. The inability to report the face's identity despite these simplifications reinforces the likelihood that AA reflects a lack of memory-encoding effort towards an attended stimulus, rather than capacity limitations in encoding or retrieval due to the task design. This is consistent with previous evidence suggesting that AA reflects a reduced effort in memory encoding despite attention (H. Chen & Wyble, 2016).

AA with unique targets

In previous work, when target stimuli were reused across trials, AA was present regardless of the target type (H. Chen et al., 2019). However, AA was observed with unique letter targets (H. Chen & Wyble, 2016; H. Chen et al., 2019) but not unique pictorial targets (H. Chen et al., 2019; W. Chen & Howe, 2017). Therefore, it seems that target repetition and stimulus characteristics of the target set interactively affect the presence of AA. In the current study using unique face stimuli as targets, AA was present, demonstrating that the target repetition boundary condition is not extendable to our stimuli set.

Stimulus type aside, multiple other factors have been identified to affect whether/how much AA can be observed. That set includes the type of the surprise question (Swan et al., 2017), the surprise attribute (H. Chen & Wyble, 2015b), and the number of pre-surprise trials (Wyble et al., 2019). However, these factors are unlikely to explain the poor accuracy on the surprise trial in this study. The current surprise question type (four-alternative force choice) and surprise attribute (identity) were similar to that in H. Chen et al. (2019) and W. Chen and Howe (2017). While AA has been found to be attenuated with fewer pre-surprise trials (Wyble et al., 2019), the number of pre-surprise trials in the current study (27) is the same as in H. Chen et al. (2019); moreover, W. Chen and Howe (2017) failed to find AA with unique pictorial targets even with 155 pre-surprise trials. Further studies would be required to identify critical stimulus characteristics of a target set that influence whether unique targets would show AA. Another point of interest is to investigate how the surprise trial accuracy for faces would alter with the use of different criteria for target selection (e.g., locating a female face among male faces in the pre-surprise trials). A selection criterion that encourages more elaborate processing of the face identity might improve surprise identification performance.

Conclusion

We found AA with unique face targets. It shows that people have difficulty recognizing a face when they did not expect to have to recognize it, even if they just attended to and selected the face seconds ago. Our finding provides an important bridge for the AA effect to be understood in a more applied context, i.e., eyewitness identification. In addition, our work expands a boundary condition of AA by showing that the occurrence of AA with unique targets is not restricted to the use of simple stimuli such as letters.

Author Note This study is supported by funding from NSF grant 1734220 awarded to author B. W. We thank Dr. Ira Hyman, an anonymous reviewer, and members of the WybleLab for helpful comments and insights. Stimuli, experiment script, and data are available at http://osf.io/b3asj/

References

Baldassari, M. J., Moore, K. N., Hyman, I., Hope, L., Mah, E., & Lindsay, D. S. (2021). *The Effect of Pre-Event Instructions on Eyewitness Identification Stage 1 Registered Report Manuscript. OSF Preprints.* https://doi.org/10.31219/osf.io/sk5bc

Baylies, P. (2020). *Stylegan-encoder* [Github code]. Retrieved May 17, 2021, from https://github.com/pbaylies/stylegan-encoder

Bian, G., McAleer, M., & Wong, W.-K. (2009). A Trinomial Test for Paired Data When There are Many Ties. *Mathematics and Computers in Simulation*, *81*(6), 1153–1160. https://doi.org/10.2139/ssrn.1410589

Bindemann, M., Burton, A. M., Hooge, I. T. C., Jenkins, R., & De Haan, E. H. F. (2005). Faces retain attention. *Psychonomic Bulletin & Review*, *12*(6), 1048–1053.

Born, S., Jordan, D., & Kerzel, D. (2020). Attribute amnesia can be modulated by foveal presentation and the pre-allocation of endogenous spatial attention. *Attention, Perception, and Psychophysics*, *82*(5), 2302–2314. https://doi.org/10.3758/s13414-020-01983-7

Bradfield, A. L., & Wells, G. L. (2000). The perceived validity of eyewitness identification testimony: A test of the five Biggers criteria. *Law and Human Behavior, 24*(5), 581–594. https://doi.org/10.1023/A:1005523129437

Chabris, C. F., Weinberger, A., Fontaine, M., & Simons, D. J. (2011). You do not talk about Fight Club if you do not notice Fight Club: Inattentional blindness for a simulated real-world assault. *i-Perception, 2*, 150–153.

Chen, H., Swan, G., & Wyble, B. (2016). Prolonged focal attention without binding: Tracking a ball for half a minute without remembering its color. *Cognition*, *147*, 144–148. https://doi.org/10.1016/j.cognition.2015.11.014

Chen, H., & Wyble, B. (2015a). Amnesia for Object Attributes: Failure to Report Attended Information That Had Just Reached Conscious Awareness. *Psychological Science*, *26*(2), 203–210. https://doi.org/10.1177/0956797614560648

Chen, H., & Wyble, B. (2015b). The location but not the attributes of visual cues are automatically encoded into working memory. *Vision Research*, *107*, 76–85. https://doi.org/10.1016/j.visres.2014.11.010

Chen, H., & Wyble, B. (2016). Attribute Amnesia Reflects a Lack of Memory Consolidation for Attended Information. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(2), 225–234. https://doi.org/10.1037/xhp0000133

Chen, H., Yu, J., Fu, Y., Zhu, P., Li, W., Zhou, J., & Shen, M. (2019). Does attribute amnesia occur with the presentation of complex, meaningful stimuli? The answer is, "it depends". *Memory and Cognition*, *47*(6), 1133–1144. https://doi.org/10.3758/s13421-019-00923-7

Chen, W., & Howe, P. D. L. (2017). Attribute amnesia is greatly reduced with novel stimuli. *PeerJ, 5*(2), e4016. https://doi.org/10.7717/peerj.4016

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. https://doi.org/10.1016/j.neuropsychologia.2011.01.029

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*(1), 87–114. https://doi.org/10.1017/S0140525X01003922

Curby, K. M., & Gauthier, I. (2007). A visual short-term memory advantage for faces. *Psychonomic Bulletin and Review*, *14*(4), 620–628. https://doi.org/10.3758/BF03196811

Cutler, B. L., Penrod, S. D., & Martens, T. K. (1987). The reliability of eyewitness identification - The role of system and estimator variables. *Law and Human Behavior*, *11*(3), 233–258. https://doi.org/10.1007/BF01044644

Davis, D., Loftus, E. F., Vanous, S., & Cucciare, M. (2008). "Unconscious transference" can be an instance of "change blindness". *Applied Cognitive Psychology*, *22*(5), 605–623. https://doi.org/10.1002/acp.1395

Eakin, D. K., Schreiber, T. A., & Sergent-Marshall, S. (2003). Misinformation Effects in Eyewitness Memory: The Presence and Absence of Memory Impairment as a Function of Warning and Misinformation Accessibility. *Journal of Experimental Psychology: Learning Memory and Cognition, 29*(5), 813–825. https://doi.org/10.1037/0278-7393.29.5.813

Eitam, B., Yeshurun, Y., & Hassan, K. (2013). Blinded by irrelevance: Pure irrelevance induced "blindness". *Journal of Experimental Psychology: Human Perception and Performance, 39*(3), 611–615. https://doi.org/10.1037/a0032269

Fawcett, J. M., Russell, E. J., Peace, K. A., & Christie, J. (2011). Of guns and geese: a meta-analytic review of the "weapon focus" literature. *Psychology, Crime and Law*, *19*(1), 1–32. https://doi.org/10.1080/1068316X.2011.599325

Hyman, I. E., Wulff, A. N., & Thomas, A. K. (2018). Crime Blindness: How Selective Attention and Inattentional Blindness Can Disrupt Eyewitness Awareness and Memory. *Policy Insights from the Behavioral and Brain Sciences*, *5*(2), 202–208. https://doi.org/10.1177/2372732218786749

Jiang, Y., Costello, P., & He, S. (2007). Processing of Invisible Stimuli: Advantage of Upright Faces and Recognizable Words in Overcoming Interocular Suppression. *Psychological Science*, *18*(4), 349–356.

Jiang, Y. V., Shupe, J. M., Swallow, K. M., & Tan, D. H. (2016). Memory for recently accessed visual attributes. *Journal of Experimental Psychology: Learning Memory and Cognition*, *42*(8), 1331–1337. https://doi.org/10.1037/xlm0000231

Karras, T., Laine, S., & Aila, T. (2019). A style-based generator architecture for generative adversarial networks. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, *2019-June*, 4396–4405. https://doi.org/10.1109/CVPR.2019.00453

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin and Review*, *20*(2), 228–242. https://doi.org/10.3758/s13423-012-0359-y

Kocab, K., & Sporer, S. L. (2016). The weapon focus effect for person identifications and descriptions: A meta-analysis. In M. Miller & B. H. Bornstein (Eds.), *Advances in psychology and law* (Vol. 1, pp. 71–117). Springer. https://doi.org/10.1007/978-3-319-29406-3_3

Lavie, N., Ro, T., & Russell, C. (2003). The role of perceptual load in processing distractor faces. *Psychological Science*, *14*(5), 510–515. https://doi.org/10.1111/1467-9280.03453

Loftus, E. F. (1976). Unconscious Transference in Eyewitness Identification. *Law and Psychology Review, 2*, 93-98.

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning and Memory*, *12*(4), 361–366. https://doi.org/10.1101/lm.94705

Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting Attention in Visual Working Memory Reduces Interference From Memory Probes. *Journal of Experimental Psychology: Learning Memory and Cognition*, *34*(2), 369–380. https://doi.org/10.1037/0278-7393.34.2.369

McHugh, M. L. (2013). The Chi-square test of independence Lessons in biostatistics. *Biochemia Medica*, *23*(2), 143–149. https://doi.org/10.11613/BM.2013.018

Mugno, M., Wyble, B., & Hudson, R. (2020). Face Orbitals: Adapting Deep Learning Models to Generate Artificial Faces. Poster presented at Penn State annual Undergraduate Exhibition, University Park, PA. https://doi.org/10.17605/OSF.IO/5CQW8

Palermo, R., & Rhodes, G. (2007). Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia*, *45*(1), 75–92. https://doi.org/10.1016/j.neuropsychologia.2006.04.025

Palmer, M. A., Brewer, N., Weber, N., & Nagesh, A. (2013). The confidence-accuracy relationship for eyewitness identification decisions: Effects of exposure duration, retention interval, and divided attention. *Journal of Experimental Psychology: Applied*, *19*(1), 55–71. https://doi.org/10.1037/a0031602

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*(1), 23–38. https://doi.org/10.1016/j.neuroscience.2005.06.005

Reinitz, M. T., Lammers, W. J., & Cochran, A. B. P. S. (1992). Memory-conjunction errors: Miscombination of stored stimulus features can produce illusions of memory. *Memory & Cognition, 20*(1), 1–11. https://doi.org/10.3758/BF03208247

Rivardo, M. G., Brown, K. A., Rodgers, A. D., Maurer, S. V., Camaione, T. C., Minjock, R. M., & Gowen, G. M. (2011). Integrating Inattentional Blindness and Eyewitness Memory. *North American Journal of Psychology*, *13*(3), 519–538.

Roelfsema, P. R. (2005). Elemental operations in vision. *Trends in Cognitive Sciences*, *9*(5), 226–233. https://doi.org/10.1016/j.tics.2005.03.012

Scolari, M., Vogel, E. K., & Awh, E. (2008). Perceptual expertise enhances the resolution but not the number of representations in working memory. *Psychonomic Bulletin and Review*, *15*(1), 215–222. https://doi.org/10.3758/PBR.15.1.215

Souza, A. S., Rerko, L., & Oberauer, K. (2016). Getting More From Visual Working Memory : Retro-Cues Enhance Retrieval and Protect From Visual Interference. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(6), 890–910.

Steblay, N. (1992). A Meta-Analytic Review of the Weapon Focus Effect. *Law and Human Behaviour*, *16*(4), 413–424.

Stein, T., Sterzer, P., & Peelen, M. V. (2012). Privileged detection of conspecifics: Evidence from inversion effects during continuous flash suppression. *Cognition*, *125*(1), 64–79. https://doi.org/10.1016/j.cognition.2012.06.005

Swan, G., Wyble, B., & Chen, H. (2017). Working memory representations persist in the face of unexpected task alterations. *Attention, Perception, & Psychophysics, 79*(5), 1408-1414.

Tabi, Y. A., Husain, M., & Manohar, S. G. (2019). Recall cues interfere with retrieval from visuospatial working memory. *British Journal of Psychology*, *110*(2), 288–305. https://doi.org/10.1111/bjop.12374

Tanaka, J. W. (2001). The Entry Point of Face Recognition: Evidence for Face Expertise. *Journal of Experimental Psychology: General*, *130*(3), 534–543.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1436–1451. https://doi.org/10.1037/0096-1523.32.6.1436

Wells, G. L. (2002). Eyewitness testimony. In *Encyclopedia of Crime and Punishment* (pp. 663–668). Retrieved February 16, 2021, from https://lib.dr.iastate.edu/psychology_pubs/75

Wolfe, J. M. (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting Memories* (pp.71-94). MIT Press.

Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin and Review*, *15*(1), 223–229. https://doi.org/10.3758/PBR.15.1.223

Wyble, B., Hess, M., O'Donnell, R. E., Chen, H., & Eitam, B. (2019). Learning how to exploit sources of information. *Memory and Cognition*, *47*(4), 696–705. https://doi.org/10.3758/s13421-018-0881-x

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.