

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

The Role of Serial Dependence in Visual Perception

Permalink

<https://escholarship.org/uc/item/3f62297z>

Author

Liberman, Alina

Publication Date

2016

Peer reviewed|Thesis/dissertation

The Role of Serial Dependence in Visual Perception

By

Alina Liberman

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Neuroscience

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Whitney, Chair

Professor Sonia Bishop

Professor Michael Silver

Professor Bruno Olshausen

Professor Stephen Palmer

Summer 2016

Abstract

The Role of Serial Dependence in Visual Perception

by

Alina Liberman

Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor David Whitney, Chair

From moment to moment, we perceive objects in the world as continuous despite fluctuations in their image properties due to factors like occlusion, visual noise, and eye movements. The mechanism by which the visual system accomplishes this object continuity remains elusive. Recent results demonstrate that the perception of low-level stimulus features such as orientation and numerosity is systematically biased (i.e. pulled) towards visual input from the recent past. The spatial region over which current orientations are pulled by previous orientations is known as the continuity field, which is temporally tuned for the past 10-15 seconds. This perceptual pull could contribute to the visual stability of low-level features over short time periods, but it does not address how visual stability occurs at the level of complex objects or during visual discontinuities. Here, we examined whether the visual system facilitates stable perception by biasing current perception of objects towards recently seen objects, and whether it operates across disruptions in visibility. First, we used psychophysics to show that the perception of face identity is systematically biased towards identities seen up to several seconds prior, even across changes in viewpoint. This perceptual bias did not depend on subjects' prior responses or on the method used to measure identity perception. Next, we tested whether this serial dependence, or positive perceptual pull, helps stabilize perceived expression. To test this, observers judged random facial expressions (ranging from happy to sad to angry). We found a pull in perceived expression toward previously seen expressions, but only when the 1-back and current face shared the same identity. Finally, we examined whether the continuity field helps maintain perceived object identity during occlusion. Specifically, we found that the perception of an oriented Gabor that emerged from behind an occluder was significantly pulled towards the random (and unrelated) orientation of the Gabor that was seen entering the occluder. Importantly, this serial dependence was stronger for predictable, continuously moving trajectories, compared to unpredictable ones or static displacements. While this serial dependence in object perception manifests as a misperception, it is adaptive: visual processing echoes the stability of objects in the world to create perceptual continuity. These results provide further evidence for the existence of an object-selective continuity field that helps maintain perceived object stability over time and across visual noise or disruptions.

Acknowledgments

This work would not have been possible without the generosity, support and assistance of many people.

I would first like to thank the members of the Whitney Lab over the last five years: Alice Albrecht, Yang Bai, Wesley Chaney, Jason Fischer, Steve Haroz, Anna Kosovicheva, Gerrit Maus, Allison Yamanashi Leib, Mauro Manassi, Timothy Sweeny, Santani Teng, Benjamin Wolfe, Nicole Wurnitsch, Ye Xia, and Kathy Zhang. It has been a great pleasure to work with and learn from such a passionate and pleasant group of scientists.

I would like to give a huge thank you to all of my scientific mentors. I would first like to thank David Whitney for accepting me into his lab and helping me grow as a scientist. I am very grateful for his encouragement, support, and all of the time spent discussing experiments and research ideas. It is a great privilege to learn from someone who is so passionate and knowledgeable about vision science. I would also like to thank Jack Gallant and Michael Silver for taking me on as a rotation student during my first year and giving me the opportunity to learn from them and their lab members, especially Rachel Denison, Caterina Gratton, Alex Huth, and Natalia Bilenko. I would like to thank Bill Prinzmetal for his valuable feedback on my work and for sharing his expertise with me. I would also like to thank Kalanit Grill-Spector, who first sparked my excitement and curiosity about perception and the brain at Stanford University.

I would also like to thank my thesis committee, Sonia Bishop, Bruno Olshausen, Steve Palmer, and Michael Silver, for their great research advice and encouragement. Our meetings and discussions made my research significantly better.

I greatly enjoyed being a member of the Helen Wills Neuroscience community. I would like to thank Tony Leonard and Kati Markowitz and the rest of the amazing Helen Wills staff for their help and advice over the years. I would also like to thank the neuroscience graduate students for creating such an inclusive and engaging academic environment.

I would like to thank my family and friends for their support throughout graduate school. I would especially like to thank my husband Andrew Sinkov. He has been incredibly supportive and encouraging through all of the highs and lows of graduate school and often had more faith in me than I did. This experience would have been much more difficult without him by my side, and my papers would contain significantly more scientific jargon.

This research was supported by National Eye Institute grant R01EY018216, an NRSA predoctoral fellowship, and an NSF graduate research fellowship.

Table of Contents

Chapter 1: Introduction	1
Chapter 2: Serial dependence in the perception of faces.....	4
Chapter 3: Serial dependence promotes the stability of perceived emotional expression in an identity-specific manner.....	21
Chapter 4: Serial dependence promotes object stability during occlusion.....	32
Chapter 5: Conclusion	43
References	44
Appendix A: Supplemental Figures for Chapter 2	51

Chapter 1: Introduction

We perceive objects and people in the world as stable and continuous even though their image properties frequently change due to factors like occlusion, visual noise, and eye movements. One important question in neuroscience is how the visual system promotes the perception of object continuity and discounts these changes over time. An intriguing possibility is that the visual system leverages the temporal persistence and regularity of objects in the physical world to stabilize perception over time. A perceptual bias towards recently seen objects would promote identity stability and counteract the often distorted, noisy and discontinuous retinal images of those objects.

Current studies show that the perception of low-level stimulus features such as orientation and numerosity is indeed serially dependent—systematically biased (i.e. pulled) towards visual input from the recent past (Cicchini, Anobile, & Burr, 2014; Corbett, Fischer, & Whitney, 2011; Fischer & Whitney, 2014). The spatial region over which current orientations are pulled by previous orientations is known as the continuity field, which is temporally tuned for the past 10-15 seconds and requires attention to the perceived orientations (Fischer & Whitney, 2014). Chapter 2 examines the object selectivity of the continuity field by testing whether the perception of face identity is systematically biased towards identities seen up to several seconds prior, even across changes in viewpoint (Lieberman, Fischer, & Whitney, 2014). Following lengthy exposure to faces, many researchers have reported negative aftereffects of facial identity or emotional expression (repulsion away from the adapting stimulus) (Carbon & Leder, 2005; Fox & Barton, 2007; Fox, Oruc, & Barton, 2008; Leopold, Rhodes, Müller, & Jeffery, 2005; Rhodes, Jeffery, Clifford, & Leopold, 2007; Tillman & Webster, 2012; Webster & Maclin, 1999; Webster, Kaping, Mizokami, & Duhamel, 2004). In Chapter 2, we test whether there is a perceptual pull towards recently seen face identities, a previously unknown form of visual adaptation. Even though a positive perceptual bias in perceived identity manifests as a misperception, it is adaptive: visual processing echoes the stability of face identity in the world to create perceptual continuity.

The object selective continuity field makes two additional predictions that I test in Chapters 3 and 4. If the continuity field is about object-level information, it predicts that there should be serial dependence not just in identity, but also in facial expression (Chapter 3), even though the underlying mechanisms for identity and expression are thought to differ (Bruce & Young, 1986; Haxby, Hoffman, & Gobbini, 2000). If the continuity field facilitates the perception of stable object identity despite occlusions, then there should be stronger serial dependence for objects that pass behind occluders and reemerge at expected versus unexpected locations (like cars through a tunnel; Chapter 4).

The visual system is very sensitive to emotional expression; observers are able to make above chance discrimination of expressions shown as briefly as 30-50 ms (Calvo & Esteves, 2005; Kirouac & Doré, 1984; Milders, Sahraie, & Logan, 2008). However, an individual's emotional expressions are not constantly or spontaneously changing. It is therefore important to balance the need to detect new facial expressions in others with maintaining perceived stability of an individual's emotional state. In Chapter 3, we examine whether there is serial dependence for emotional expressions for the purpose of perceived expression stability. To test this, we used a behavioral psychophysical task that examines whether the perception of emotional expression is biased towards recently seen expressions. Facial expression perception has also exhibited

identity-dependent adaptation and processing whereas facial identity recognition may be independent of changes in expression (Fox et al., 2008; Fox & Barton, 2007; Schweinberger & Soukup, 1998; Schweinberger, Burton, & Kelly, 1999; Wild-Wall, Dimigen, & Sommer, 2008). We therefore also examine whether the amplitude of serial dependence for emotional expression is dependent on the identity of the 1-back face. We predict that there should be a larger perceptual pull from previously seen faces if the previous face has the same identity as the current face compared to when the previous face has a different identity.

A person can easily recognize a friend's face under difficult or partial viewing conditions, such as occlusion of the face by a scarf. In Chapter 4, we test the prediction that if the object-selective continuity field maintains an object's identity or appearance in the presence of noise or occlusion, then we should find serial dependence in the perception of an object that reemerges from behind an occluder: the object's identity should appear pulled by the identity of the object that entered the occlusion. Previous research shows that observers see an object moving continuously through an occluder as the trajectory of one object, even if that object looks different before and after emerging from the occluder (Burke, 1952). Additionally, change detection of the color of an object after occlusion improves when the object is perceived as moving on a continuous trajectory (Flombaum & Scholl, 2006), indicating that we perceive objects as persisting even when we cannot see them. Furthermore, there is increased object-based attention to occluded objects compared to when those objects were visible, suggesting an increase in attentional resources with occlusion (Flombaum, Scholl, & Pylyshyn, 2008). This knowledge of object continuity and persistence under occlusion is learned at a very young age and is present early in infancy (Carey & Xu, 2001; Spelke, 1994; Spelke, Kestenbaum, Simons, & Wein, 1995). However, no study has systematically addressed how an object *appears* as it reemerges from behind an occluder. Since the object-selective continuity field promotes temporal stability of object identity (Lieberman et al., 2014), it could play a role in stabilizing the appearance of an object after it is obscured. In Chapter 4, we examine whether the perception of an object as it moves from behind an occluder is pulled or biased towards the identity of the object that entered the occluder. To test this, we measured the perceptual pull of an oriented Gabor as it exits from an occluder towards the Gabor orientation that subjects saw entering the occluder as it travels along an expected or unexpected trajectory.

There are many prior studies demonstrating that perceptual history can shape current perception, but none of them predict the existence of a perceptual bias towards recently seen identities or expressions, or a bias that promotes object continuity under occlusion. Priming and negative aftereffects show a type of perceptual dependence on the recent past, yet are distinct from our predictions. Adaptation studies show that prior exposure to a variety of stimulus features (Anstis, Verstraten, & Mather, 1998; Campbell & Maffei, 1971; Webster, 2011; Webster et al., 2004; Webster & Mollon, 1991) results in a stimulus-specific negative aftereffect or perceptual repulsion away from the adapting stimulus (for reviews see (Kohn, 2007; Thompson & Burr, 2009; Webster, 2012)). However, our findings demonstrate a positive perceptual pull towards the recent past and are therefore not a result of known forms of adaptation. Our experiments may be related to perceptual priming effects (Faivre & Kouider, 2011; Kristjánsson, Ingvarsdóttir, & Teitsdóttir, 2008; Kristjánsson, Bjarnason, Hjaltason, & Stefánsdóttir, 2009; Maljkovic & Nakayama, 1994; 1996; 2000), but there are important differences. Priming generally occurs for reaction time (Maljkovic & Nakayama, 1994; 1996; 2000) and, where relevant, can improve discriminability of primed stimuli (Sigurdardóttir, Kristjánsson, & Driver, 2008); serial dependence does not impact reaction time and is a

reduction in the discriminability of objects (Fischer & Whitney, 2014; Liberman et al., 2014) for the sake of perceptual stability. The continuity field will therefore produce a systematic bias of perception towards previously seen identities, emotions or the object perceived before occlusion—perceptual serial dependence, and this would be adaptive, at least on relatively short time scales.

Chapter 2: Serial dependence in the perception of faces

The research in this chapter has been published in the following article and is included with permission:

Liberman A, Fischer J, and Whitney D. "Serial Dependence in the Perception of Faces", *Current Biology* 24, 2569–2574, November 3, 2014.

Introduction

One of vision's most important functions is to maintain the stable perception of objects so that they *look* the same from moment to moment. Object motion, occlusion, changes in lighting, and eye movements all contribute to noisy visual input, so how is it that we still perceive objects as the same from one moment to the next? We hypothesized that the visual system might leverage the statistical regularities of objects in the physical world to stabilize perception over time. Objects present a second ago usually continue to exist at this moment; the visual system might harness this temporal predictability by biasing the perception of an object in the present moment toward the appearance of that object in the recent past. Such serially dependent perception would carry the benefit of making our visual experience more fluid and continuous, but its existence is not a foregone conclusion. Following exposure to stimuli that range from basic features to faces, there are usually negative aftereffects (repulsion away from what was recently perceived) (Campbell & Maffei, 1971; Clifford et al., 2007; Rhodes et al., 2004; Webster et al., 2004). In addition, maximally independent perception from moment-to-moment would provide the most unbiased estimates of object identity. Nevertheless, a systematic bias of perception towards what was previously seen—perceptual serial dependence—could be adaptive, at least on relatively short time scales. This bias would help us perceive objects as stable over space and time, even though the retinal images of those objects are often distorted, noisy, and discontinuous.

Previous research has shown that the perception of orientation at any given moment is systematically biased towards recently seen orientations, contributing to the visual stability of low-level features over short periods of time (Fischer & Whitney, 2014). The spatial region over which current orientations are pulled by previous orientations is known as the continuity field, which is temporally tuned for the past 10-15 seconds. However, an orientation-selective continuity field does not explain how perceptual stability arises for complex objects. Our goal is therefore to determine whether serial dependence is fundamental to the stable perception of objects in the world.

In this series of experiments, we examined whether the perception of a face is biased towards previously seen face identities. To test this, we presented a random series of morphed faces (**Figure 2.1**) several seconds apart and measured the perceived face identity of each face morph. We found that perceived identity at a given moment was pulled by the identity presented 5 seconds ago or more, even when the previous identity was viewed from a different angle. Our results reveal the object-selective continuity field (CF) as a mechanism that promotes perceptual stability by biasing perceived object identity toward recently seen identities.

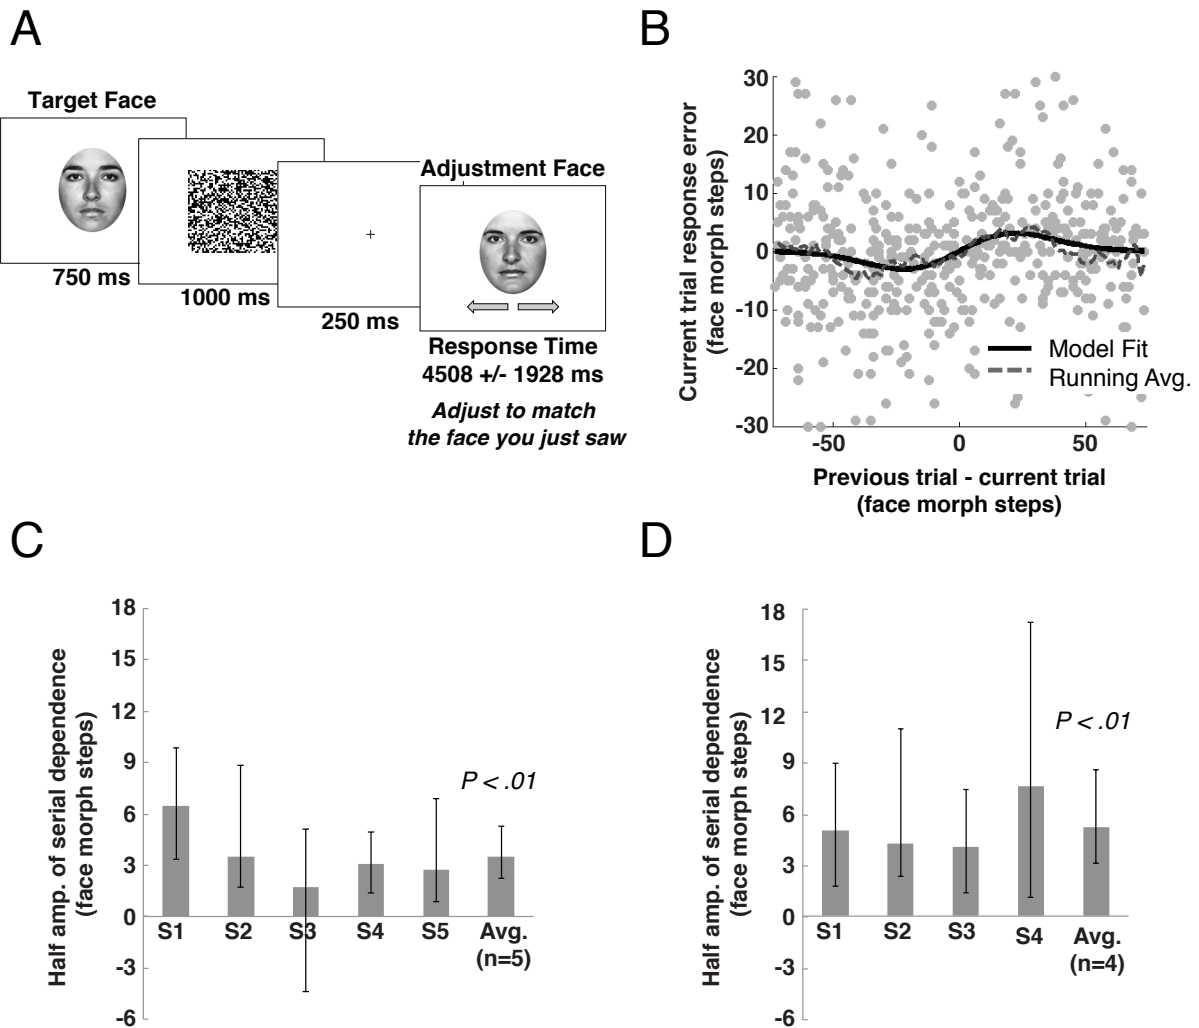


Figure 2.1. Experiment 1 Trial Sequence and Results. **(A)** Trial sequence for the method of adjustment task in Experiment 1. On each trial, a randomly selected target face was presented for 750 ms, followed by a 1000 ms noise mask of black and white pixels to reduce afterimages, and a 250 ms fixation cross. Subjects then saw a test screen containing a random adjustment face, which they modified by scrolling through the continuous identity wheel to match the target face (see Figure S1A). **(B)** Example data from subject 4, with each data point showing performance on one trial. In units of face morph steps, the x-axis is the shortest distance along the morph wheel between the current and 1-back target face (1-back target face – current target face), and the y-axis is the shortest distance along the morph wheel between the selected match face and target face (match face – current target face). Positive x-axis values indicate that the 1-back target face was clockwise on the face morph wheel relative to the current target face, and positive y-axis values indicate that the current match face was also clockwise relative to the current target face (Figure S1A). The running average (dashed line) reveals a clear trend in the data, which followed a derivative-of-Gaussian shape (model fit depicted as solid line). **(C)** Half-amplitude of the serial dependence for each subject in Experiment 1a. Error bars are bootstrapped 95% confidence intervals and p-value is based on group permuted null distribution. Additional experiments show that memory confusions cannot fully explain our pattern of results (see Figure S2). **(D)** Half-amplitude of the serial dependence for each subject in Experiment 1b for trials with no 1-back response. Sequential effects have been known to influence subjects' responses by introducing dependencies between current and previous

trial decisions (Luce & Green, 1974; Tanner, Rauk, & Atkinson, 1970; Wiegiersma, 1982a; 1982b). However, these results are not due entirely to sequential decision biases since we observed serially dependent perception without a 1-back response. Error bars are bootstrapped 95% confidence intervals and p-value is based on group permuted null distribution.

Results

Experiment 1a: Serial Dependence of Perceived Face Identity

We presented a random series of faces drawn from an identity morph continuum (**Figure S1A**) and measured the perceived identity of each face using a method of adjustment (MOA) task. On each trial, a random target face was presented followed by a test screen containing a random adjustment face, which subjects matched to the target face using the continuous identity morph wheel (**Figure 2.1A** and **Figure S1A**). “Target face” denotes the face that subjects tried to match, “adjustment face” denotes the randomly selected starting point for matching the target, and “match face” denotes the face that subjects selected as most similar to the target face.

Perceptual error was calculated as the shortest distance along the morph wheel between the match and target faces. Each subject’s error on the current trial was compared to the difference in target face identities between the current and previous trial (**Figure 2.1B**). We fit a simplified Gaussian derivative (DoG) to each subject’s data and calculated p-values using permutation analysis (**Figure 2.1B**; see **Experimental Procedures**). All subjects displayed a positive DoG half-amplitude, indicating that perceived identity on a given trial was significantly pulled in the direction of identities presented in the previous trial ($p < 0.01$, $n=5$, group permuted null) (**Figure 2.1C**). The largest attraction of perceived identity occurred when the 1-back target face was, on average, ± 24.5 morph frames away from the current target face, which resulted in an average perceptual pull towards the 1-back face of ± 3.5 face morph frames. The full amplitude of the effect was therefore 7 face morph steps, indicating that the current face appeared pulled toward the previous face by over 1.5 times the just noticeable difference (JND; see **Experimental Procedures**). No subject showed a significant influence of faces seen two trials back, which may reflect a narrow temporal window over which identity serial dependence occurs. Average response time (RT) across subjects was 4250 ± 2168 ms; the 1-back face occurred, on average, ~ 7500 ms prior to the current trial face. Perceived face identity was therefore strongly attracted toward the identity of a random target face seen more than 7 seconds prior.

Experiment 1b: Serial Dependence Without Previous Responses

To control for potential biases caused by prior responses or adjustment stimuli, four subjects completed a variation of Experiment 1a where they did not respond on a randomly selected 50% of trials. During these surprise “no response” trials, subjects saw the target face followed by a 2000 ms blank screen before beginning the next trial. We analyzed the subset of trials in Experiment 1b with no response on the 1-back trial, and fit a DoG to each subject’s data. All subjects had a positive DoG half-amplitude ($p < 0.01$, $n=4$, group permuted null) (**Figure 2.1D**), indicating that subject responses, *per se*, are not necessary for the serial dependence of perceived identity. Nonetheless, the process of responding and attending to the stimulus may play an important role in serial dependence. Two participants showed a 2-back serial dependence effect ($p < 0.05$, $n=2$, permuted null), potentially due to the shorter time between current and 1-

back target faces with no response. The 1-back target face occurred 3250 ms prior to the current target face when no response was required on the 1-back trial.

Experiment 2: Serial Dependence of Face Perception Using Constant Stimuli

Experiments 1a and 1b demonstrated that perception of face identity is pulled towards recently seen identities, a misperception that could facilitate stable face identity perception over time. In Experiment 2, we used a two-interval forced choice (2IFC) design to determine whether serial dependence altered perception independent of response method. This experiment also had the benefit of reducing response time and the number of intervening stimuli seen during the response period.

The faces used in this experiment were a subset of those used in Experiment 1, including original Face A (#1), original Face B (#50), and the 48 face morphs in between (**Figure S1A**). Before the experiment, subjects were trained to recognize Faces A and B (see **Experimental Procedures**). Immediately after training, participants were shown sequences of two faces per trial and judged which of the two faces looked more similar to face A (**Figure 2.2A**). The initial face presented in each trial, “first face,” was presented for 1000 ms. The following face, “second face,” was presented for 500 ms and differed randomly from the first face by ± 12 , ± 6 , or 0 face morphs steps (**Figure S1B**). Since subjects saw the first face for twice as long as the second face, we expected the 1-back first face to have a stronger pull on the perception of the subsequent trial’s first face.

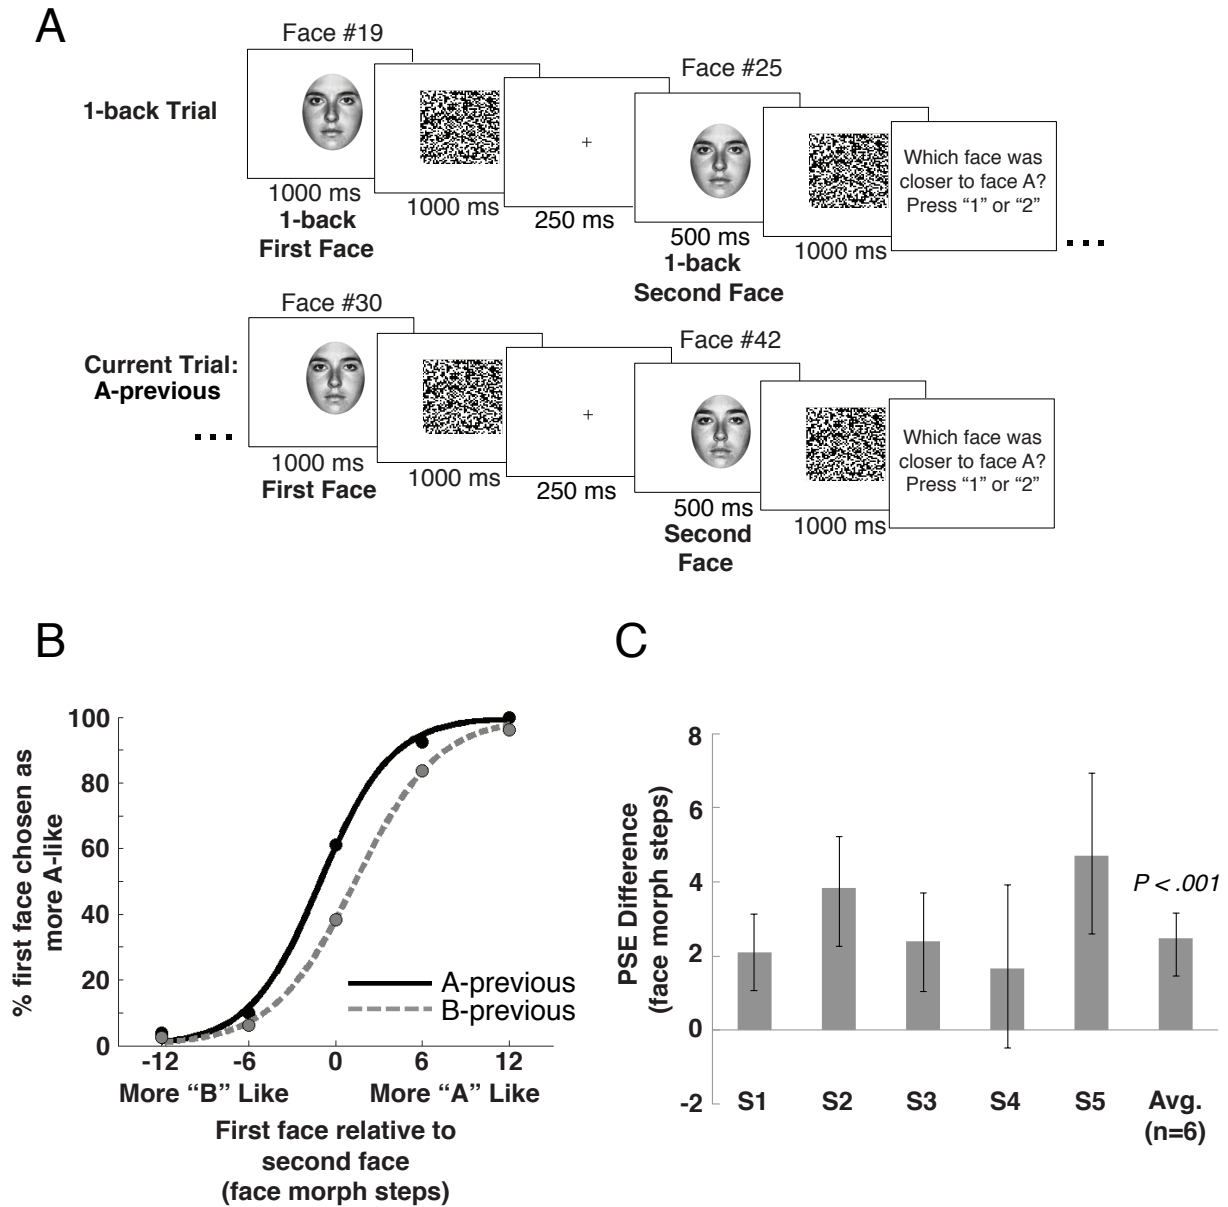


Figure 2.2. Experiment 2 Trial Sequence and Results. **(A)** Trial sequence for 2IFC task in Experiment 2. For each trial, the first face was presented for 1000 ms followed by a 1000 ms noise mask and 250 ms fixation cross. Subjects saw the second face for 500 ms and judged whether the first face (press “1”) or second face (press “2”) looked more like original Face A. Trial type was determined by comparing the position in the morph continuum of the current trial first face to that of the 1-back trial first face. If the 1-back first face was closer to original Face A along the morph continuum, the trial was labeled an “A-previous” trial. Faces are shown here without added noise. **(B)** Example psychometric functions for subject 3. The abscissa shows the difference in the identity of the first face relative to the second face in the current trial. Trials that fell in bins -12 and -6 had a first face that was more B-like relative to the second face, trials in the 0 bin had identical first and second faces, and trials in the +6 and +12 bin had a first face that was more A-like relative to the second face. The ordinate shows the proportion of first faces on the current trial that were chosen as being more A-like. The black curve consists of all trials with 1-back first faces that were more “A”-like and the gray dotted curve consists of all trials with 1-back first faces that were more “B”-like. If the 1-back first face positively pulled subjects’ perception of face

identity, then there should be a leftward horizontal displacement of the black curve relative to the gray dotted curve, which is what we found for all subjects (see **Figure S3A**). (C) PSE difference between the black and gray dotted curve for each subject. Error bars are bootstrapped 95% confidence intervals and p-value is based on group permuted null distribution.

Trials for which the 1-back first face was comparatively closer to Face A along the morph continuum were labeled “A-previous” and trials for which the 1-back first face was comparatively more similar to Face B were labeled “B-previous.” Each subject saw an equal number of A-previous and B-previous trials, but presentation order was shuffled. We fit separate logistic functions to A-previous and B-previous trials and calculated the slope and point of subject equality (PSE) for each logistic curve fit (**Figure 2.2B & Figure S3A**; also see **Experimental Procedures**). We also fit several lagged logistic regression models to the data to sequentially examine the influence of each previously seen face.

If the 1-back first face pulled perception of the current trial’s first face more than the second face, then there should be a leftward displacement of the A-previous logistic curve relative to the B-previous curve (a significant difference in the PSE (*b*) parameter). A leftward shift in PSE would indicate that the presentation of a relatively more A-like 1-back first face altered subjects’ perception such that the current first face actually appeared more A-like. However, if subjects’ identity perception was repelled or not influenced by previously viewed faces, we would observe a rightward displacement of the A-previous curve relative to the B-previous curve, or no displacement at all. We assessed the significance of each subject’s PSE shift using a permutation test to calculate a null distribution of PSE differences (see **Experimental Procedures**).

We found a significant leftward shift of the A-previous curve ($p < 0.001$, permuted null distribution, $n=6$) (**Figure 2.2C**), with 4 of the 5 subjects showing a significant shift (each $p < 0.001$, permuted null distribution; see **Experimental Procedures**) (**Figure S3A**). For A-previous trials, subjects were more likely to perceive the current trial first face (than second face) as A-like since the current and 1-back first face were closer together in time and presented for twice as long. Average response time across all subjects was 435 ± 205 ms ($n=6$); the 1-back first face occurred ~ 5685 ms prior to the current first face. We found no consistent slope differences between A-previous and B-previous trials, indicating no difference in sensitivity (permutation test; see **Experimental Procedures**).

To determine whether the 1-back second face also influenced subjects’ perception, we fit several lagged logistic regression models to each subject’s data and determined which model best predicted responses. Each successive model tested whether considering another face further back in the past explained significantly more variance in responses compared to a model without that face (see **Experimental Procedures**). For 4 of the 5 subjects, the best-fitting lagged logistic regression model included both the 1-back first and second face (least significant subject: $F(2, 816) = 11.682$, $p < 0.0001$), indicating that both faces presented on the previous trial (more than 5 seconds ago) significantly biased perception in the present trial.

Experiment 3: Serial Dependence Occurs Across Face Viewpoints

In Experiment 3, our goal was to determine whether perceived identity is serially dependent across different face viewing angles, where basic image features change but identity remains stable. The procedure for Experiment 3 was identical to that of Experiment 2, except

subjects were trained on two original neutral male face identities, Face A and Face B, within each of three possible viewpoints (frontal, left, right) (**Figure S1C**). Importantly, no two sequential trials contained the same viewpoint, but the first and second face within a trial were always viewed from the same angle (**Figure 2.3A**).

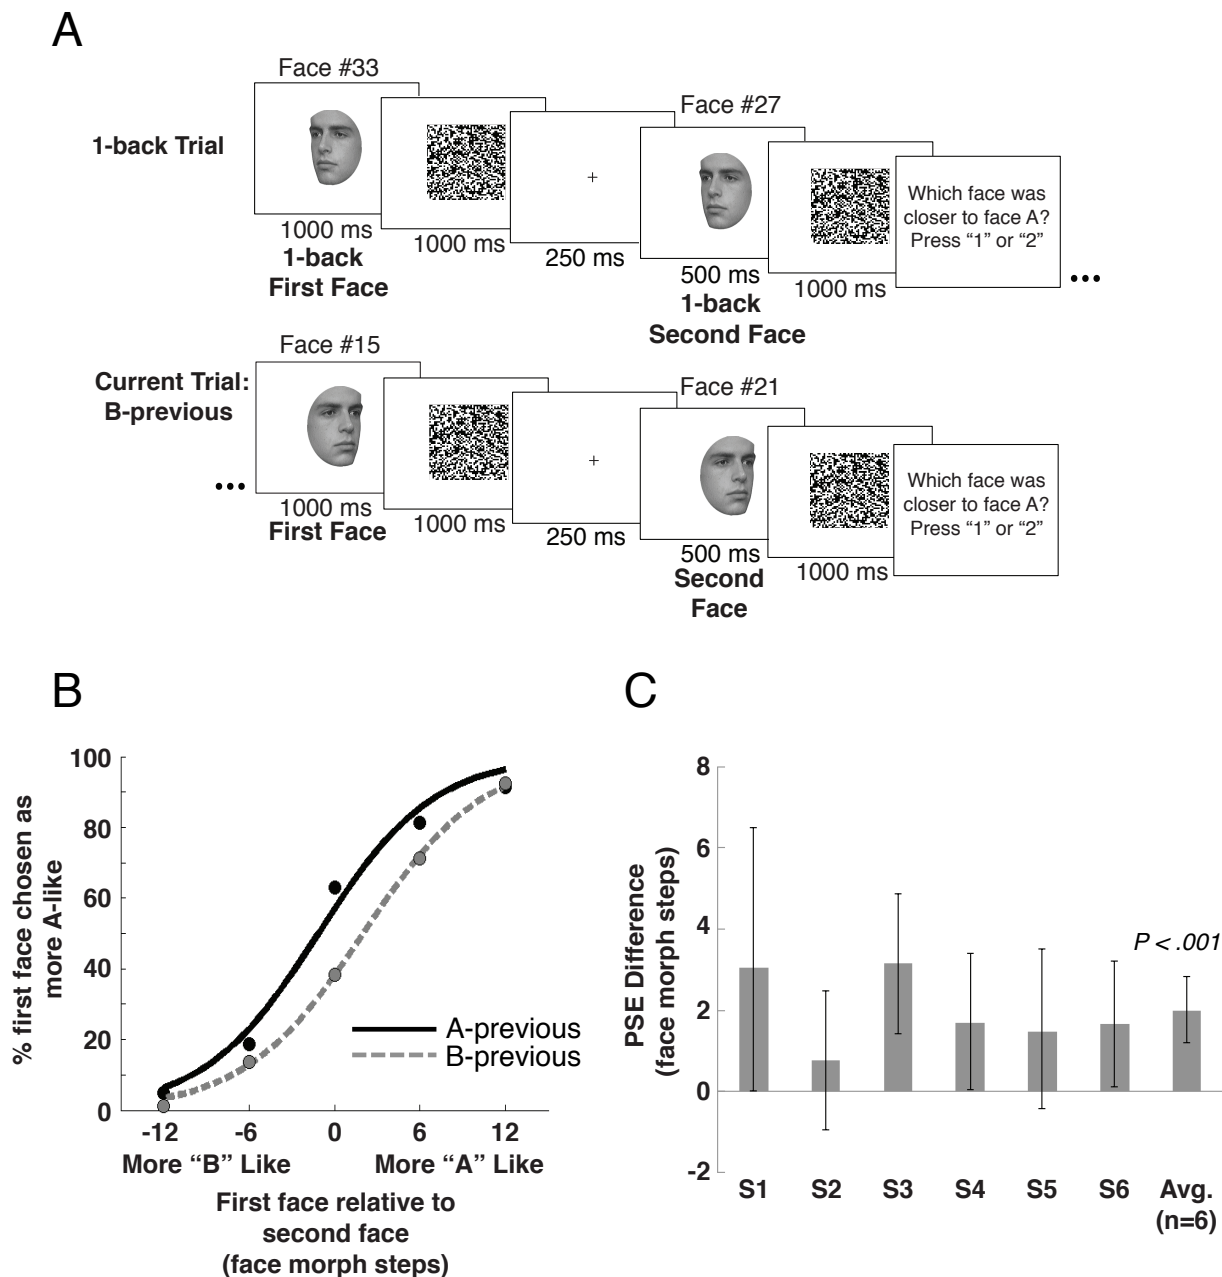


Figure 2.3. Experiment 3 Trial Sequence and Results. **(A)** Trial sequence for the 2IFC task in Experiment 3. We used grayscale image morphs based on 2 original neutral male faces across three different viewpoints (frontal, left, right), cropped by an oval to remove hairline (see Figure S1C). The trial sequence was identical to that of Experiment 2, except both 1-back trial faces were always of a different viewpoint relative to current trial faces. Faces in the figure are shown without added noise. **(B)** Example

data from subject 3. The black curve consists of all trials with 1-back first faces that were more “A”-like and the gray dotted curve consists of all trials with 1-back first faces that were more “B”-like. If the 1-back first face positively pulls subjects’ perception of face identity, then there should be a leftward horizontal displacement of the black curve relative to the gray dotted curve, which is what we report for all subjects (see Figure S3B). (C) PSE difference for each subject. Error bars are bootstrapped 95% confidence intervals and p-value is based on group permuted null distribution.

Even across different viewpoints, subjects were more likely to perceive the current target face as A-like if the 1-back target face was more A-like (**Figure 2.3B and Figure S3B**). All subjects showed a leftward displacement of the A-previous curve relative to the B-previous curve ($p < 0.001$, permuted null distribution, $n=6$), with 4 of the 6 subjects showing a significant shift (each $p < 0.05$, permuted null distribution) (**Figure 2.3C and Figure S3B**). The average RT across subjects was 327 ± 174 ms; the 1-back target face was seen ~ 5577 ms prior to the current target face. There were no significant slope differences between A-previous and B-previous psychometric functions. We also simulated the effect of response hysteresis (responding the same way on successive trials) using the exact trial sequences presented and determined that response hysteresis could not cause this serial dependence (**Figure S3C**). These results show that serial dependence can operate on high-level identity representations rather than simply on low-level features.

Experiment 4: Serial Dependence Across Face Viewpoints: Method of Adjustment

To extend the results of Experiment 3, we presented sequential faces from different viewpoints and measured perceived face identity using an MOA task identical to that in Experiment 1. We used a new set of female faces with two possible viewpoints (right or left-facing profiles), drawn from a circular identity continuum to avoid any edge effects that might have been present in the 1-dimensional stimulus arrays in Experiment 3 (**Figure S1D**). Within a trial, the target and adjustment faces were always shown in the same viewpoint, but the viewpoint differed from one trial to the next (**Figure 2.4A**; see **Experimental Procedures**).

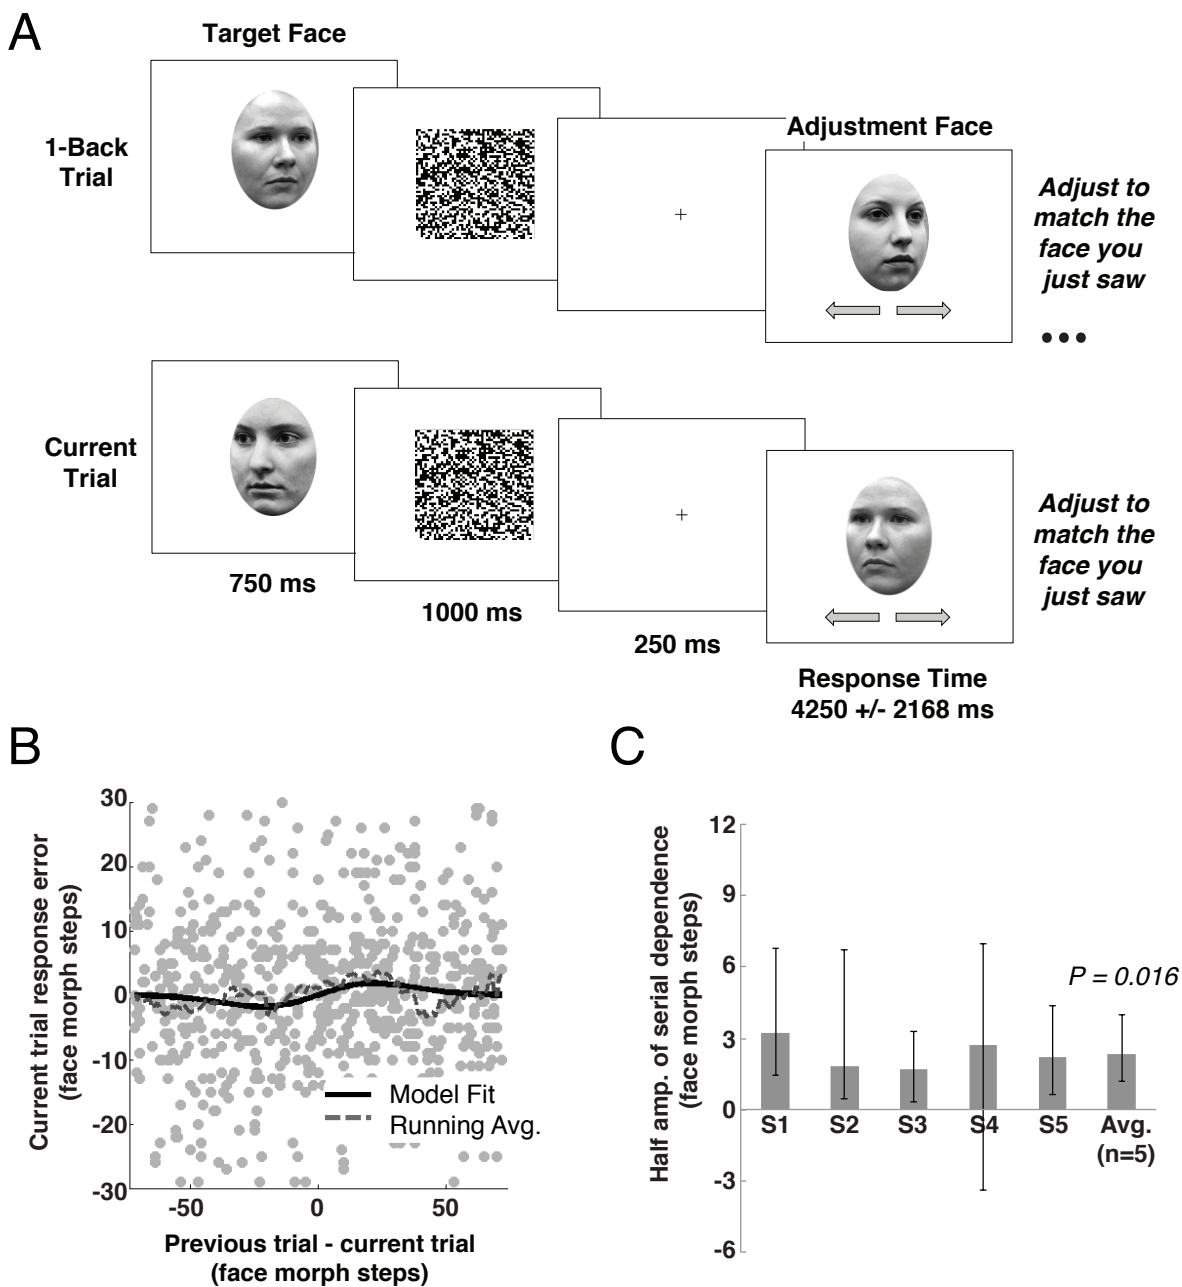


Figure 2.4. Experiment 4 Trial Sequence and Results. **(A)** Trial sequence for the method of adjustment task in experiment 4. A circular morphed continuum of female face identities was created with two possible viewpoints (right or left-facing profiles) (see Figure S1D). The identities shown here are similar to those in the experiment, with permission obtained for reproduction purposes. Randomly selected target faces were presented for 750 ms followed by a 1000 ms mask of black and white pixels. Subjects responded by matching the adjustment face to the target face. The 1-back trial target face was always from a different viewpoint relative to the current trial target face, but the target and adjustment face had the same viewpoint. **(B)** Example data from subject 2, with each dot showing performance on one trial. In units of face morph steps, the x-axis is the shortest distance along the morph wheel between the current

and 1-back target face, and the y-axis is the shortest distance along the morph wheel between the selected match face and target face. The DoG model fit is depicted as a solid line and the running average is depicted as a dashed line. (C) Half-amplitude of the serial dependence for each subject in Experiment 4. Error bars are bootstrapped 95% confidence intervals and p-value is based on group permuted null distribution.

All subjects displayed a positive DoG half-amplitude, indicating that perceived face identity was significantly pulled toward faces presented one trial ago, even though the 1-back face was always from a different viewpoint ($p < 0.02$, $n=5$, group permuted null; **Figure 2.4C**). The largest attraction of perceived identity occurred when the 1-back target face was, on average, ± 23.1 morph frames away from the current target face, which resulted in an average perceptual pull towards the 1-back face of ± 2.34 morph frames. Average response time (RT) across subjects was 4508 ± 1928 ms; the 1-back face occurred ~ 7758 ms prior to the current trial face. There was no significant difference in serial dependence amplitude between Experiment 4 and Experiment 1. The cross-viewpoint effects in Experiment 4 therefore indicate that serial dependence occurs at the level of object-centered perceptual representations and does not depend on low-level stimulus features.

Discussion

Our experiments demonstrated that perceived face identity was pulled by identities encountered 5 or more seconds ago. This effect did not depend on subjects responding on the previous trial, and there was no perceptual pull on the current face if the previously seen face was sufficiently different. To determine whether serial dependence operates at the level of identity, our final two experiments manipulated the viewpoint of the sequentially presented faces. We found that identity perception is serially dependent across different face viewpoints, even without any prior associative perceptual training (Wallis & Bühlhoff, 2001). Some existing models, including Bayesian models of perceptual dependencies (Cicchini et al., 2014) and physiologically-motivated population coding models (Fischer & Whitney, 2014; Maunsell & Treue, 2006), can produce serially-dependent effects, but, importantly, our cross-viewpoint results take these further by demonstrating that the perceptual continuity field (Fischer & Whitney, 2014) can operate at the level of object-centered perceptual representations.

History Effects In Perception

Prior studies have shown that perceptual history can shape current perception, but none predict the existence of a continuity field (Fischer & Whitney, 2014) tuned for face identity. Priming, negative aftereffects, hysteresis, and other phenomena show a type of perceptual dependence on the recent past, yet are distinct from our findings. Adaptation studies show that prior exposure to a variety of stimulus features (Anstis et al., 1998; Campbell & Maffei, 1971; Webster et al., 2004; Webster & Mollon, 1991) results in a stimulus-specific negative aftereffect or perceptual repulsion away from the adapting stimulus (for reviews see (Kohn, 2007; Thompson & Burr, 2009; Webster, 2012)). Our results, however, indicate a positive perceptual pull towards the recent past and are therefore not a result of known forms of adaptation. Since the serial dependence effect is restricted to faces seen 3-10 seconds ago, our experiments do not show a long-term positive aftereffect, as described in (Chopin & Mamassian, 2012). There is potentially a small contribution of memory confusion between the current and 1-back face, if subjects sometimes mistakenly reported the 1-back face in Experiments 1 or 4, but additional

control experiments show that memory confusions are unlikely and cannot account for our pattern of results (see **Figure S2**).

Furthermore, our results are not due to typical hysteresis of near-threshold stimuli (Eagleman, Jacobson, & Sejnowski, 2004; Preminger, Sagi, & Tsodyks, 2007; Williams, Phillips, & Sekuler, 1986) or stabilization of bistable stimuli (Hock, Kelso, & Schöner, 1993; Leopold, Wilke, Maier, & Logothetis, 2002; Maloney, Martello, Sahm, & Spillmann, 2005; Pearson & Brascamp, 2008; Sterzer & Rees, 2008), since our stimuli are randomly presented (thus disrupting hysteresis (Preminger et al., 2007)) and are not bistable. Our results may be related to perceptual priming effects (Kristjánsson et al., 2008; Kristjánsson et al., 2009; Maljkovic & Nakayama, 1994; 1996; 2000), but there are important differences. Priming generally occurs for reaction time (Maljkovic & Nakayama, 1994; 1996; 2000) and, where relevant, can improve discriminability of primed stimuli (Sigurdardottir et al., 2008). Our results reveal a counterintuitive reduction in the discriminability of sequential stimuli due to serial dependence (**Figure S3D**). Nevertheless, the possible interaction between priming and serial dependence remains an open question.

Is Serial Dependence in Perception Inevitable?

There are three possible perceptual consequences of prior exposure to a visual stimulus: a negative aftereffect, complete independence in sequential perception, or positive serial dependence. Although negative face aftereffects can emerge following brief adaptation (Leopold et al., 2005; Rhodes et al., 2007), like they sometimes do for basic features (Glasser, Tsui, Pack, & Tadin, 2011; Kosovicheva et al., 2012), serial dependence on the timescale of our experiments trumped any potential negative aftereffects.

Complete independence in moment-to-moment perceptual judgments would theoretically be the most bias-free and accurate. Given independent temporal noise and estimates of object identity, a temporal integration mechanism without serial dependence could better estimate the instantaneous state of the world compared to one that introduces a sequential dependence. Our results suggest that the visual system instead favors perceptual continuity over short periods of time, at the cost of introducing potential perceptual biases.

Conclusion

Our results are consistent with a continuity field (Fischer & Whitney, 2014), yet go significantly beyond related findings on perceptual history effects (Cicchini et al., 2014; Corbett et al., 2011; Fischer & Whitney, 2014) by showing that the perception of faces, and not just features, is serially dependent. Thus, the continuity field is object-selective, surviving changes in viewpoint, and reflects a mechanism that produces serially-dependent perception of objects for the purpose of visual stability. By recycling previously perceived identities, the object-selective continuity field decreases the neural computations necessary for the identification of perceptually similar objects over time.

Experimental Procedures

General Methods

For all experiments, faces were centered on a white background and overlaid with a central

fixation cross. Subjects viewed stimuli at a distance of 56 cm on a monitor with a resolution of 1024 x 768 and a refresh rate of 100 Hz. Subjects used a keyboard or mouse for all responses.

All experimental procedures were approved by the UC Berkeley Institutional Review Board. Participants were affiliates of UC Berkeley and provided written informed consent before participation. All participants had normal or corrected-to-normal vision, and all except one were naïve to the purpose of the experiment.

Experiment 1a & 1b

Subjects. Five subjects (4 female; age = 24-32 years) participated in Experiment 1a and four subjects (4 female; age = 18-31) participated in Experiment 1b. One of the subjects in Experiment 1b was not naïve to the experiment, and one of the subjects participated in both Experiment 1a and 1b.

Stimuli and procedure. We used a set of 147 Caucasian female faces with neutral expressions (**Figure S1A**), which were generated from three original Ekman identities (Ekman & Friesen, 1976a) using Morph 2.5 (Gryphon Software). Each presented face subtended 5.9 x 7.3 degrees of visual angle. During the experiment, subjects were tested on their ability to identify randomly chosen target faces with a method of adjustment (MOA) task. We measured subjects' identification errors on the MOA task to determine whether a subject's perception of each target identity was influenced by previously seen target identities.

Experiment 1a

On each trial, a random target face was presented for 750 ms, followed by a 1000 ms noise mask of randomly shuffled black and white pixels, to reduce afterimages, and then a 250 ms fixation cross prior to the response (**Figure 2.1A**). Subjects then saw a test screen containing a random adjustment face, which they adjusted to match the target face. After picking a match face, subjects saw a 1000 ms noise mask followed by a 1000 ms fixation cross before the next trial began. Here, we use the terms “target face” to mean the face that subjects tried to match, “adjustment face” to denote the randomly-selected face used as the starting point for matching the target, and “match face” for the face that subjects selected as most similar to the target face. The experiment was self-paced and subjects were allowed to take as much time as necessary to respond. We recorded responses based on the numerical value of the match face along the morph continuum, with possible values ranging from 1 to 147. Four subjects each completed 540 trials, and one subject completed 624 trials.

Experiment 1b

In order to rule out any potential biases due to previous motor responses, four additional subjects completed a version of the experiment where half of the trials, selected randomly, did not require a response. Subjects instead saw a surprise blank screen for 2000 ms during the response period, followed by the next trial. Each of the subjects in the no-response condition completed 2382 total trials over 5-6 sessions.

Memory confusion experiments

2AFC

To determine whether subjects might be experiencing memory confusion, and whether the probability of confusing the target and 1-back face changes with increasing similarity between those faces, we ran a control experiment that measured how often subjects mistakenly reported the 1-back target face rather than the current target face (**Figure S2A & B**). The stimulus set, trial sequence, and timing were similar to that of Experiment 1b: on each trial, subjects saw a random face for 750 ms, followed by a 1000 ms noise mask and 250 ms fixation cross. Subjects then saw a blank screen for 2000 ms, followed by a screen displaying two faces—a target and a lure—one was randomly assigned to the left and the other to the right of fixation. Subjects indicated which face they had last seen (L or R) using the keyboard arrow keys. One of the comparison faces was always the current face (the target, and thus, a correct response), and the second comparison face (the lure) was either a random face (50% of trials) or the 1-back face (50% of trials; picking this 1-back face would constitute a memory confusion). The random lure faces served to establish a baseline for how often subjects made discrimination errors or lapses. We tested three subjects in this 2AFC experiment (all of whom also participated in Experiment 1), and collapsed their data for subsequent analyses. Two subjects completed 500 trials over two sessions, and one subject completed 300 trials in one session.

MOA

To measure the precision of subjects' memory for the current and 1-back faces, we ran the same three subjects from the 2AFC experiment above (and Experiment 1) in a modified version of Experiment 1a (**Figure S2C & D**). The procedure was identical to that of Experiment 1a, except that on a random and unpredictable 25% of trials, subjects were asked to match the adjustment face to the 1-back target face rather than the current target face. The histogram in Figure S2C shows the error distribution (collapsed across subjects) for judgments of the current target face (75% of all responses), and the histogram in Figure S2D shows the collapsed error distribution for judgments of the 1-back target face (25% of all trials). Two subjects completed 250 trials in one session, and one subject completed 300 trials over two sessions.

Analysis. Identification error was computed as the shortest distance along the morph wheel between the match face and the target face. Identification error was compared to the difference in target face identities between the current and previous trial, computed as the shortest distance along the morph wheel between the previous target face (1-back) and the current target face. Trials were considered lapses and excluded if errors exceeded ± 60 morph units (3.5 standard deviations from mean on average, less than 5% of data excluded) or if the response time was longer than 10 seconds. We fit a simplified Gaussian derivative (DoG) to each subject's data of the form:

$$y = abcxe^{-(bx)^2}$$

where parameter y is identification error on each trial (match face – current target face), x is the difference along the wheel between the current and 1-back target face (1-back target face – current target face), a is half the peak-to-trough amplitude of the derivative-of-Gaussian, b scales the width of the Gaussian derivative, and c is a constant, $\sqrt{2}/e^{-0.5}$, which scales the curve to make the a parameter equal to the peak amplitude (**Figure 2.1B**). We fit the Gaussian derivative using constrained nonlinear minimization of the residual sum of squares.

For each subject's data, we generated confidence intervals by calculating a bootstrapped distribution of the model-fitting parameter values by resampling the data with replacement 10,000 times (Efron & Tibshirani, 1993). On each iteration, we fit a new DoG to obtain a bootstrapped half-amplitude and width for each subject. We used the half amplitude of the DoG, the a parameter in the above equation, to measure the degree to which subjects' reports of face identity were pulled in the direction of n-back face identities. If subjects' perception of face identity was repelled by the 1-back face (e.g., because of a negative aftereffect; (Clifford et al., 2007; Webster & MacLeod, 2011)) or not influenced by the 1-back face (because of independent, bias-free perception on each trial), then the half-amplitude of the DoG should be negative or close to zero, respectively.

In order to calculate significance, we also generated a null distribution of half amplitude (a) values for each subject using a permutation analysis. We randomly shuffled each subject's response errors relative to the difference between the current and 1-back target face and recalculated the DoG fit for each iteration of the shuffled data. We ran this procedure for 10,000 iterations in order to generate a within-subject null distribution of half amplitude values. P-values were calculated by computing the proportion of half amplitudes in each subject's null distribution that were greater than or equal to the observed half amplitude. To test significance at the group level, we chose a random a parameter value index (without replacement) from each subject's null distribution and averaged those values across all five subjects. We repeated this procedure for 10,000 iterations in order to generate a group null distribution of average half amplitude values, and calculated the p-value as described above.

Experiment 2

Subjects. Six subjects (3 female; age = 23-30 years) participated in the experiment. One of the subjects was not naïve to the experiment, and two of the subjects had participated in Experiment 1. One subject became unavailable after one run of the experiment and was only included in the group analysis.

Stimuli and procedure. We used a subset of 50 female faces from the 147 female face morphs used in Experiment 1. This subset of faces consisted of two original female identities, Face A (#1) and Face B (#50), and the 48 morphs between them (**Figure S1A**). Each face was presented in an oval aperture to mask out the hairline and subtended 5.9 x 7.3 degrees of visual angle. Noise was added to the faces to increase difficulty by randomly replacing 5-10% of the pixels in the entire image with black or white pixels.

Before beginning the experiment, subjects were trained to recognize Face A and Face B. During training, subjects were initially shown each face, with a label, and then tested on their recall through 30 randomized trials (15 trials per face). During the randomized trials, subjects were shown either Face A or Face B and had to respond by correctly identifying which face they were shown. If subjects did not get at least 90% of trials correct (27/30), they had to repeat the training phase.

Immediately after training, subjects began the main experiment. Participants were shown a sequence of two faces in each trial, and they had to decide which of the two looked more similar

to Face A (much like a two-interval-forced-choice [2IFC] task). The initial face presented in each trial sequence, “first face,” was drawn from a subset of 26 faces from the 50 faces used in this experiment. These 26 faces were taken from the center of the morph continuum and could range from face morph #13 to face #38. The following face in the trial sequence, “second face,” could differ from the first face by ± 12 , ± 6 , or 0 face morphs. Trials that fell in bins -12 and -6 had a first face that was more B-like relative to the second face, trials that fell in the 0 bin had identical first and second faces, and trials that fell in the +6 and +12 bin had a first face that was more A-like relative to the second face (**Figure S1B**). The second face could range anywhere along the morph continuum between Face A and Face B, while the first face was limited to the center of the continuum. Within a trial, the likelihood of the first face being A- or B-like relative to the second face was randomized.

The first face was presented for 1000 ms, followed by a 1000 ms noise mask and 250 ms fixation cross. The second face was presented for 500 ms followed by a 1000 ms noise mask (**Figure 2.2A**). Subjects responded by identifying which of the two faces looked more similar to Face A by pressing “1” for the first face or “2” for the second face, followed by a 1500 ms fixation cross before the next trial. Four subjects completed 820 trials over two runs, one subject completed 1230 trials over three runs, and one subject became unavailable after one run of 410 trials, but the inclusion of their data had no impact on the group effects.

Analysis. Trials were sorted into one of two groups: B-previous or A-previous. Group membership was determined by comparing the position in the morph continuum of the current trial first face to that of the 1-back first face (**Figure 2A**). Trials for which the 1-back first face was closer to Face A along the morph continuum were labeled as “A-previous” trials and trials for which the 1-back first face was more similar to Face B were labeled as “B-previous” trials. Each subject saw an equal number of A-previous and B-previous trials, but presentation order was shuffled. Once we separated a subject’s data into two groups, we fit a separate psychometric function to A-previous and B-previous trials using the following logistic equation:

$$P(\text{respond A on trial } t) = \frac{1}{1 + e^{-a(x_t - b)}}$$

where x_t is the difference between the first and second face for trial t (i.e., -12, -6, 0, 6, or 12), parameter a scales with the slope, and b is the point of subjective equality (PSE).

We generated confidence intervals by calculating a bootstrapped distribution of model-fitting parameter values. Within each trial type and bin, we resampled the data with replacement for 10,000 iterations and fit a new psychometric function on each iteration (Efron & Tibshirani, 1993). We then computed the difference between “A-previous” and “B-previous” bootstrapped b (PSE) values in order to generate a distribution of PSE differences. To test for significance, we ran a permutation analysis where we shuffled the ‘A-previous’ and ‘B-previous’ labels within each of the five bins. We then recalculated logistic curve fits for the new, randomly assigned A-previous and B-previous trials and computed the difference in PSE between the new parameters. We ran this procedure for 10,000 iterations in order to generate a within-subject null distribution of difference scores. We calculated a p-value by computing the proportion of difference values in each subject’s null distribution that were greater than or equal to the observed difference between curves.

To calculate the just noticeable difference (JND) for this set of face morphs, we collapsed each subject's A-previous and B-previous psychometric functions (**Figure S3A**) into one set of data and fit that data with a single logistic function. We then found the x values at which the logistic curve passed through 25% and 75% on the y-axis, and calculated the JND by taking half of the absolute difference between these two x values. The resulting JND was about 4.5 face morph steps between these female faces.

In order to determine whether the 1-back second face also pulled subjects' perception, we fit several lagged logistic regression models to each subject's data and determined which model best predicted subjects' responses. Each successive model tested whether considering another face further back in the past explained significantly more variance in subjects' responses compared to a model without that face. Each subject's data was fit using the following logistic function:

$$P(\text{respond A on trial } t) = \text{Logit}^{-1}(\alpha + \beta_0 \times \text{Offset}_t + \sum_{i=1}^m \beta_i (X_{t-i} - X_t))$$

where α is a constant, X_t is the position on the linear face continuum of the first face of trial t , $X_{(t-i)}$ is the position on the face continuum of the i th-back first or second face, and Offset_t is the difference between the first and second face for trial t (i.e., -12, -6, 0, 6, or 12). Each lag was calculated as the difference between the current trial first face and the i th-back face, i.e. Lag 1: (1-back second face – current first face); Lag 2: (1-back first face – current first face); Lag 3: (2-back second face – current face), and added to the model in a stepwise manner. A negative lag value indicated that the i th-back face was closer to Face A (#1) compared to the current first face. To select the best model for each subject's data, we implemented a stepwise procedure where we tested significance of each additional lag using AIC (Akaike, 1974). Once the highest order lag no longer decreased the model AIC, we stopped adding additional i th-back face differences.

We ran an F-test to determine whether the model with the highest order significant lag had a significantly better fit to the data compared to a model with no lags added (i.e. only including the difference between the first and second face on the current trial as a predictor).

Experiment 3

Subjects. Six subjects (3 female; age = 24-36 years) participated in the experiment. One of the subjects was not naïve to the experiment, and three of the subjects had participated in Experiment 2.

Stimuli and procedure. We used grayscale image morphs based on 2 original neutral male faces across three different viewpoints (frontal, left, right), cropped by an oval to remove the hairline (**Figure S1C**). Each presented face subtended 5.64 x 7.47 degrees of visual angle. We randomly replaced 5% of the pixels in each image with black or white pixels in order to increase difficulty by reducing small features that might be diagnostic markers of a given identity.

The procedure for Experiment 3 was identical to that of Experiment 2, except subjects were trained on the two original male face identities, Face A and Face B, within each of the three possible viewpoints (**Figure S1C**).

During the main experiment, subjects were shown the first face for 1000 ms, followed by a 1000 ms noise mask and a 250 ms fixation cross. They then saw the second face for 500 ms followed by a 1000 ms noise mask (**Figure 2.3A**). Subjects had to indicate which of the two faces looked more similar to male Face A, after which they saw a fixation dot for 1500 ms. No two sequential trials contained the same viewpoint, but the target and comparison face (within a single trial) were always viewed from the same angle. Six subjects completed 820 trials over two runs.

Analysis. Experiment 3 had identical logistic equation fitting, bootstrap, and permutation analysis as in Experiment 2.

Experiment 4

Subjects. Five subjects (4 female; age = 18-37 years) participated in the experiment. One of the subjects was not naïve to the experiment, and three of the subjects had participated in Experiment 1.

Stimuli and procedure. We used a continuum of 144 Caucasian female faces with neutral expressions (**Figure S1D**), which were generated from three original identities across two different viewpoints (left- and right-facing profile), cropped by an oval to remove the hairline. Each presented face subtended 5.9 x 7.3 degrees of visual angle. During the experiment, subjects were tested on their ability to identify randomly chosen target faces with a method of adjustment (MOA) task. The procedure for Experiment 4 was identical to that of Experiment 1, except subjects were trained on the three original female face identities within each of the two possible viewpoints (**Figure S1D**).

During training, subjects were familiarized with Face A, B, and C in each of the two viewpoints. Subjects were initially shown each face turned to the right, with a label, and then tested on recall through 30 randomized trials (10 trials per face). During the randomized trials, subjects were shown one of the three faces and had to respond by correctly identifying which face they were viewing. If subjects did not get at least 90% of trials correct (27/30), they had to repeat the training phase. The same training procedure was then repeated for the left-facing identities.

During the main experiment, a random target face was presented for 750 ms, followed by a 1000 ms noise mask of randomly shuffled black and white pixels, to reduce afterimages, and then a 250 ms fixation cross prior to the response (**Figure 2.4A**). Subjects then saw a test screen containing a random adjustment face, which they adjusted to match the target face. Importantly, no two sequential trials contained the same viewpoint, but the target and adjustment face (within a single trial) were always viewed from the same angle.

Analysis. Experiment 4 had identical DoG fitting, bootstrap, and permutation analysis as in Experiment 1.

Chapter 3: Serial dependence promotes the stability of perceived emotional expression in an identity-specific manner

Introduction

The perception of emotional expression is fundamental for successful social interactions, personal emotion regulation, the experience of empathy, and many other vital activities (Salovey & Mayer, 1990). Individuals with Parkinson's disease, schizophrenia, traumatic brain injury, and other cognitive deficits have impairments in recognizing facial affect, which may have deleterious effects on their personal and social interactions (Crocker & McDonald, 2009; Jacobs, Shuren, Bowers, & Heilman, 1995; Martin, Baudouin, Tiberghien, & Franck, 2005). Most research on emotion perception focuses on fast emotion categorization or recognition speed and accuracy (Edwards, 1998; Ekman & Friesen, 1971; Kirouac & Doré, 1984; Stel & van Knippenberg, 2008; Tracy & Randles, 2011; Tracy & Robins, 2008). The visual system is very sensitive to emotional expression; observers are able to make above chance discrimination of expressions shown as briefly as 30-50 ms (Calvo & Esteves, 2005; Kirouac & Doré, 1984; Milders et al., 2008). Yet, emotional expressions do not constantly or spontaneously change. Thus, it is important to balance the ability to detect new facial expressions with the need to maintain perceived *stability* of an individual's emotional state. However, no study has addressed how the visual system promotes the perception of expression stability over time.

From moment to moment, we perceive the identities of objects and people in the world as stable and continuous even though their image properties frequently change due to factors like occlusion, visual noise, changes in viewpoint, and eye movements. Previous studies have shown that the perception of orientation, numerosity, and other low-level stimulus features is serially dependent—systematically biased (i.e. pulled) towards similar visual input from the recent past (Cicchini et al., 2014; Corbett et al., 2011; Fischer & Whitney, 2014). This serial dependence is tuned over visual space and time, as well as in feature space (object similarity; Fischer & Whitney, 2014; Liberman et al., 2014). The spatio-temporal region over which current object features, such as orientation, are pulled by previously seen features is known as the Continuity Field (CF).

Beyond orientation and other basic features, serial dependence occurs at higher levels of perception as well. We have recently demonstrated that the continuity field is object-selective by showing that the perception of face identity is systematically biased towards identities seen up to several seconds prior, even across changes in viewpoint (Liberman et al., 2014). If the continuity field promotes the perceived stability of emotional expression as well as identity, then there should be serial dependence not just in identity, but also in facial expression. We therefore predicted that perceived emotional expression would be biased towards recently seen emotional expressions. Here, we tested this using a psychophysical task, and we also determined whether this serial dependence in perceived emotional expression depended on the similarity of sequential facial identities. Because expressions within an individual may be more auto-correlated than across individuals, we expected that there should be a larger perceptual pull from previously seen faces if the previous face had the same identity as the current face.

Experiment 1: Serial dependence of perceived emotional expression

If the continuity field (CF) facilitates the perceptual stability of perceived emotional expression, then there should be a measurable serial dependence in judged expression; the

perception of a facial expression at one moment in time should be pulled towards recently seen expressions. To test this, we presented a series of random facial expressions drawn from an expression morph continuum and had observers report the facial expression that they last saw through a method of adjustment task (**Figure 3.1A & B**). The question was whether the perceived expression at a given moment was serially dependent on the expressions of the faces seen several seconds previously.

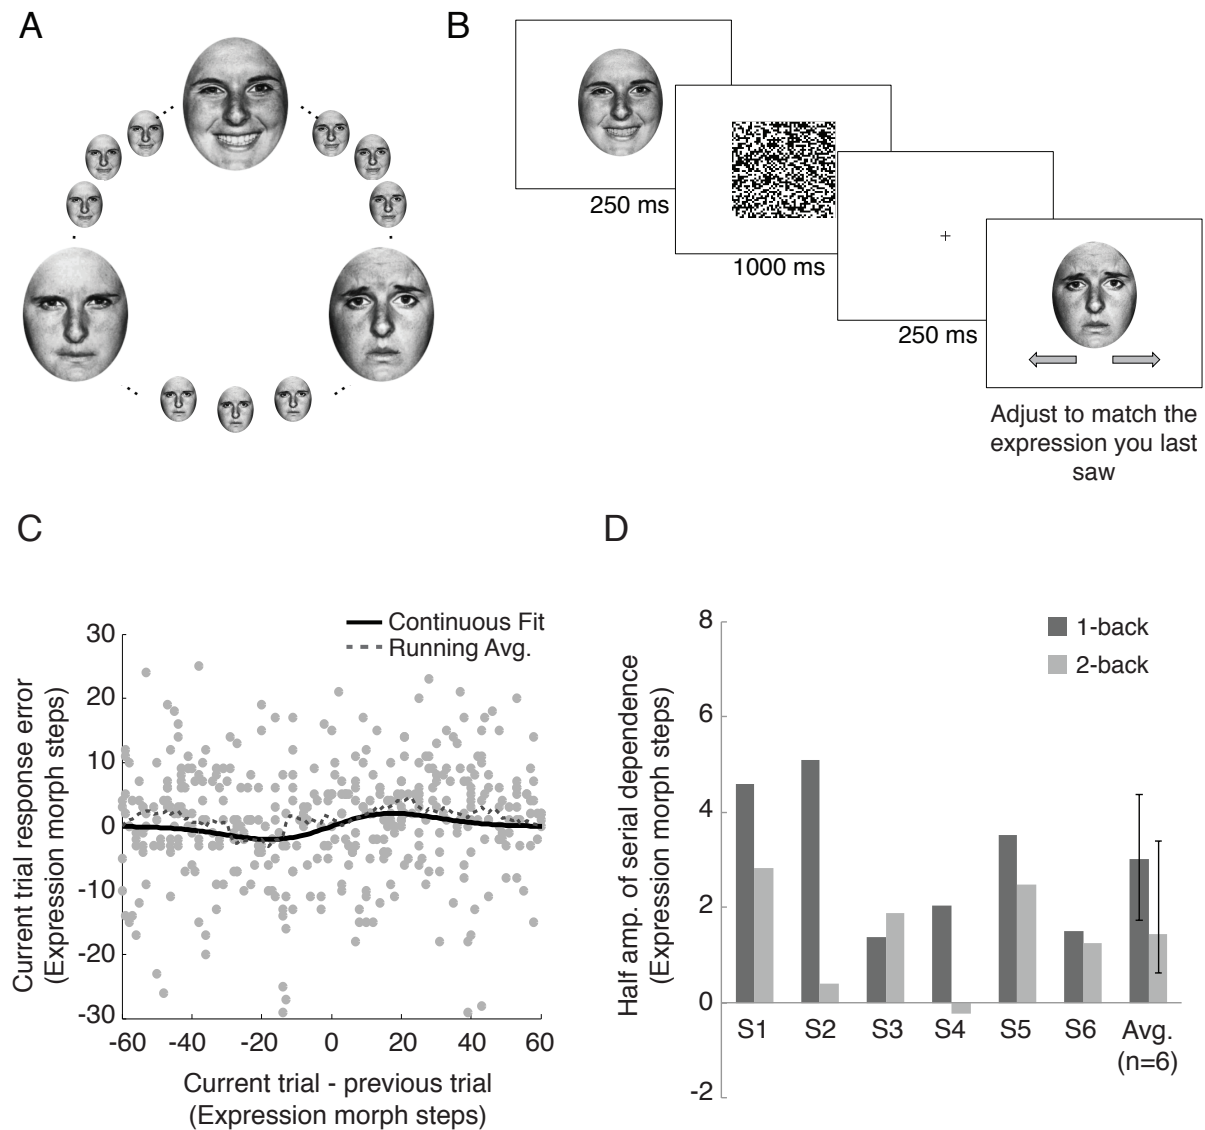


Figure 3.1. Experiment 1 Trial Sequence and Results. **(A)** Face morphs used in Experiments 1 and 2. These morphs were based on one original female Ekman identity displaying a happy, sad, or angry face (Ekman & Friesen, 1976b). A set of 48 morphs was created between these expressions, resulting in a face morph continuum of 147 faces. **(B)** Trial sequence for the method of adjustment task in Experiment 1. After selecting an expression that was matched to the previous trial, subjects saw a 1000 ms noise mask followed by a 1000 ms fixation cross before the next trial began. **(C)** Example data from subject 4, with each data point showing performance on one trial. The x-axis is the difference between the current and 1-back target expression (1-back target expression – current target expression, in units of expression morph

steps), and the y-axis is the difference between the selected match expression and target expression (match expression – current target expression). **(D)** Half-amplitude of the serial dependence for each subject in Experiment 1. On average, subjects had a significant, positive perceptual pull of the current facial expression towards the expression seen one or two trials previously (both $p < .05$, permuted null distribution). Error bars are bootstrapped 95% confidence intervals.

Methods

Participants

Six subjects (3 female) ranging in age from 19 to 33 years ($M = 26.7$, $SD = 5.5$ years) participated in Experiment 1. One of the subjects in Experiment 1 was not naïve to the experiment. All experimental procedures were approved by the UC Berkeley Institutional Review Board. Participants were affiliates of UC Berkeley and provided written informed consent before participation. All participants had normal or corrected-to-normal vision.

Stimuli and procedure

We used a set of 147 Caucasian female face morphs with different expressions (**Figure 3.1A**), which were generated using Morph 2.5 (Gryphon Software) from one original Ekman identity displaying a happy, sad, or angry face (Ekman & Friesen, 1976b), cropped by an oval aperture to remove the hairline. Each presented face subtended 5.9×7.3 degrees of visual angle. During the experiment, subjects were tested on their ability to match randomly chosen target expressions with a method of adjustment (MOA) task. We measured subjects' errors on the MOA task to determine whether a subject's perception of each target expression was influenced by previously seen target expressions. For all experiments, faces were centered on a white background and overlaid with a central fixation cross. All experiments were programmed in MATLAB (The MathWorks, Natick, MA) using Psychophysics Toolbox (Brainard, 1997). Subjects viewed stimuli at a distance of 56 cm on a monitor with a resolution of 1024×768 and a refresh rate of 100 Hz. Subjects used a keyboard or mouse for all responses.

On each trial, a random target expression was presented for 250 ms, followed by a 1000 ms noise mask of random black and white pixels (to reduce afterimages) and then a 250 ms fixation cross prior to the response (**Figure 3.1B**). Subjects then saw a test screen containing a random adjustment expression, which they adjusted to match the target facial expression. After picking a match expression, subjects saw a 1000 ms noise mask followed by a 1000 ms fixation cross before the next trial began. Here, we use the terms “target expression” to mean the face that subjects tried to match, “adjustment expression” to denote the randomly-selected face used as the starting point for matching the target, and “match expression” for the facial expression that subjects selected as most similar to the target expression. The experiment was self-paced and subjects were allowed to take as much time as necessary to respond. We recorded responses based on the numerical value of the match expression along the morph continuum, with possible values ranging from 1 to 147. Six subjects each completed 500 trials.

Analysis

Response (perceptual) error was computed as the shortest distance along the morph wheel between the match expression and the target expression (y-axis). This response error was then compared to the difference in expressions between the current and previous trial (x-axis), computed as the shortest distance along the morph wheel between the previous target expression

(1-back) and the current target expression. For each subject's data, trials were considered lapses and excluded if error exceeded 3 standard deviations from the mean or if the response time was longer than 10 seconds (less than 5% of data excluded on average). We fit a simplified Gaussian derivative (DoG) to each subject's data of the form:

$$y = abcxe^{-(bx)^2}$$

where parameter y is identification error on each trial (match expression – current target expression), x is the difference along the wheel between the current and 1-back target expression (1-back target expression – current target expression), a is half the peak-to-trough amplitude of the derivative-of-Gaussian, b scales the width of the Gaussian derivative, and c is a constant, $\sqrt{2}/e^{-0.5}$, which scales the curve to make the a parameter equal to the peak amplitude (**Figure 3.1C**). We fit the Gaussian derivative using constrained nonlinear minimization of the residual sum of squares.

For each subject's data, we generated confidence intervals by calculating a bootstrapped distribution of the model-fitting parameter values by resampling the data with replacement 5,000 times (Efron & Tibshirani, 1993). On each iteration, we fit a new DoG to obtain a bootstrapped half-amplitude and width for each subject. We used the half amplitude of the DoG—the a parameter in the above equation—to measure the degree to which subjects' reports of facial expression were pulled in the direction of n-back expressions. If subjects' perception of facial expression was repelled by the 1-back expression (e.g., because of a negative aftereffect; (Clifford et al., 2007; Webster et al., 2004)) or not influenced by the 1-back expression (because of independent, bias-free perception on each trial), then the half-amplitude of the DoG should be negative or close to zero, respectively.

In order to calculate significance, we also generated a null distribution of half amplitude (a) values for each subject using a permutation analysis. We randomly shuffled each subject's response errors relative to the difference between the current and 1-back target expression and recalculated the DoG fit for each iteration of the shuffled data. We ran this procedure for 5,000 iterations to generate a within-subject null distribution of half amplitude values. P-values were calculated by computing the proportion of half amplitudes in each subject's null distribution that were greater than or equal to the observed half amplitude. To test significance at the group level, we chose a random a parameter value index (without replacement) from each subject's null distribution and averaged those values across all five subjects. We repeated this procedure for 5,000 iterations in order to generate a group null distribution of average half amplitude values, and calculated the p-values as described above.

Results


All subjects displayed a positive DoG half-amplitude, indicating that perceived expression on a given trial was significantly pulled in the direction of expressions presented in the previous trial ($p < 0.001$, $n=6$, group permuted null, **Figure 3.1D**), with 4 of the 6 subjects showing significant serial dependence ($p < .05$, permuted null). The largest attraction of perceived expression occurred when the 1-back target expression was, on average, ± 23.1 morph frames away from the current target expression, which resulted in an average perceptual pull towards the 1-back face of ± 3.0 face morph frames. Most subjects also showed an influence of expressions seen two trials back ($p < .05$, $n=6$, group permuted null). Average response time (RT) across subjects was 2695 ms ($SD = 1523$ ms), so the 1-back face occurred, on average, 6195 ms prior to the current trial face, and the 2-back face occurred, on average, 12390 ms prior to the

current face. Perceived facial expression was therefore pulled toward the expression of a random target expression seen more than 12 seconds prior.

Experiment 2: Is serial dependence for emotional expression identity dependent?

In Experiment 1, we found serial dependence in perceived emotional expression. This is consistent with the idea that the continuity field facilitates the stability of perceived expression by echoing the physical autocorrelations of facial expressions in natural environments (Lieberman et al., 2014). Individual faces convey expressions that vary over time, but these expressions do not randomly or suddenly change. Therefore, emotional expressions might be more physically auto correlated within an individual face than across different identities. If the visual system mirrors this, we would expect stronger serial dependence in perceived expression within a given identity, than across identities. In Experiment 2, we tested whether the amplitude of serial dependence for facial expression was modulated by the similarity between the current and the 1-back face identity.

A



B

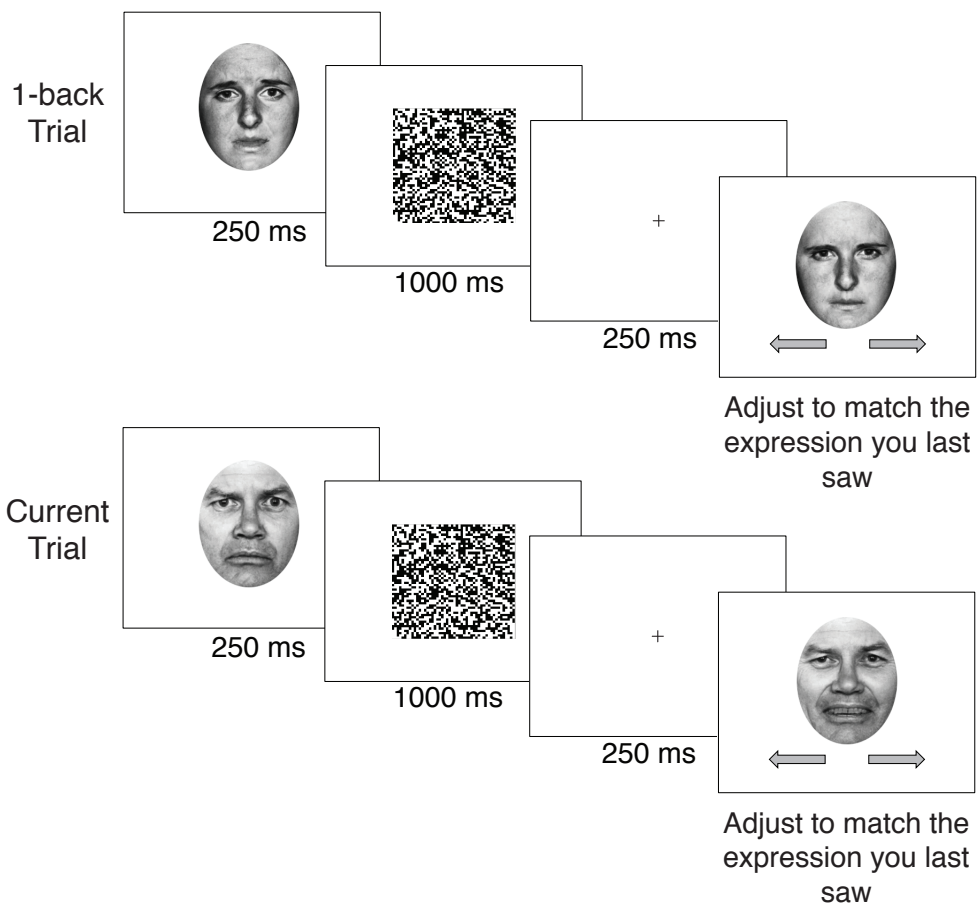


Figure 3.2. Experiment 2 Stimuli and Trial Sequence. **(A)** In addition to the female morphs (**Figure 3.1A**), a set of male face morphs was also used in Experiment 2. These stimuli consisted of grayscale image morphs based on an original male identity (Ekman & Friesen, 1976b) with happy, sad, and angry expressions. **(B)** Trial sequence for the method of adjustment task in Experiment 2. On each trial, a randomly selected target expression of either gender was presented for 250 ms, followed by a 1000 ms noise mask of black and white pixels to reduce afterimages, and a 250 ms fixation cross. Subjects then saw a test screen containing a random adjustment face, which they modified by scrolling through the continuous expression wheel to match the target expression. After picking a match expression, subjects saw a 1000 ms noise mask followed by a 1000 ms fixation cross before the next trial began.

Methods

Participants

Seven subjects (4 female) ranging in age from 20 to 37 years ($M = 29.1$, $SD = 5.5$ years) participated in Experiment 2. One subject was excluded because their response error SD was more than two SD s away from the other subjects, but the inclusion of their data did not change the pattern or significance of the results. One of the subjects in Experiment 2 was not naïve to the experiment, and five of the subjects also participated in Experiment 1. All experimental procedures were approved by UC Berkeley Institutional Review Board. Participants were affiliates of UC Berkeley and provided written informed consent before participation. All participants had normal or corrected-to-normal vision.

Stimuli and procedure

The faces used in Experiment 2 consisted of two emotional morph continua: the first morph continuum was the set of female faces used in Experiment 1 (**Figure 3.1A**), and the second set was based on a male face (**Figure 3.2A**) (Ekman & Friesen, 1976b). We created the male face morph continuum from three facial expressions (happy, sad, and angry) using the same morph procedures as described in Experiment 1 (**Figure 3.2A**).

During the experiment, subjects were tested on their ability to match randomly chosen target emotions with a MOA matching task, similar to the task in Experiment 1. However, on each trial, participants now saw a randomly chosen male or female target expression for 250 ms, followed by a 1000 ms noise mask of randomly shuffled black and white pixels, and then a 250 ms fixation cross prior to the response (**Figure 3.2B**). Subjects then saw a test screen containing a random adjustment expression with the same gender as the target expression, which they adjusted to match the last expression they saw. After picking a match expression, subjects saw a 1000 ms noise mask followed by a 1000 ms fixation cross before the next trial began. Six subjects each completed 400 trials. All other experiment procedures were identical to Experiment 1.

Analysis

Similar to Experiment 1, response error was computed as the shortest distance along the morph wheel between the match expression and the target expression (x-axis) and compared to the difference in expressions between the current and previous trial (y-axis). Trials where the 1-back target expression and the current target expression shared the same identity were labeled “Same 1-back Identity,” and trials where the 1-back target expression and the current target expression had different identities were labeled “Different 1-back Identity.” For each subject, we fit a separate simplified DoG function to Same 1-back and Different 1-back trials, according to

the fitting and significance testing procedures described in Experiment 1 (**Figure 3.3A**). We then determined whether there was a significant difference in the serial dependence amplitude between same 1-back identity and different 1-back identity trials using a permutation analysis.

For the permutation analysis, we shuffled the same 1-back and different 1-back trial labels, recalculated DoG fits for the new, randomly assigned trial types, and took the difference between same and different 1-back DoG amplitude. We ran this procedure for 5,000 iterations in order to generate a within-subject null distribution of difference scores. We calculated a p-value by computing the proportion of difference values in each subject's null distribution that were greater than or equal to the observed difference between amplitudes. To test significance at the group level, we chose a random α -difference parameter value index (without replacement) from each subject's null distribution of differences and averaged those values across all five subjects. We repeated this procedure for 5,000 iterations in order to generate a group null distribution of average difference values, and calculated the p-value as described above.

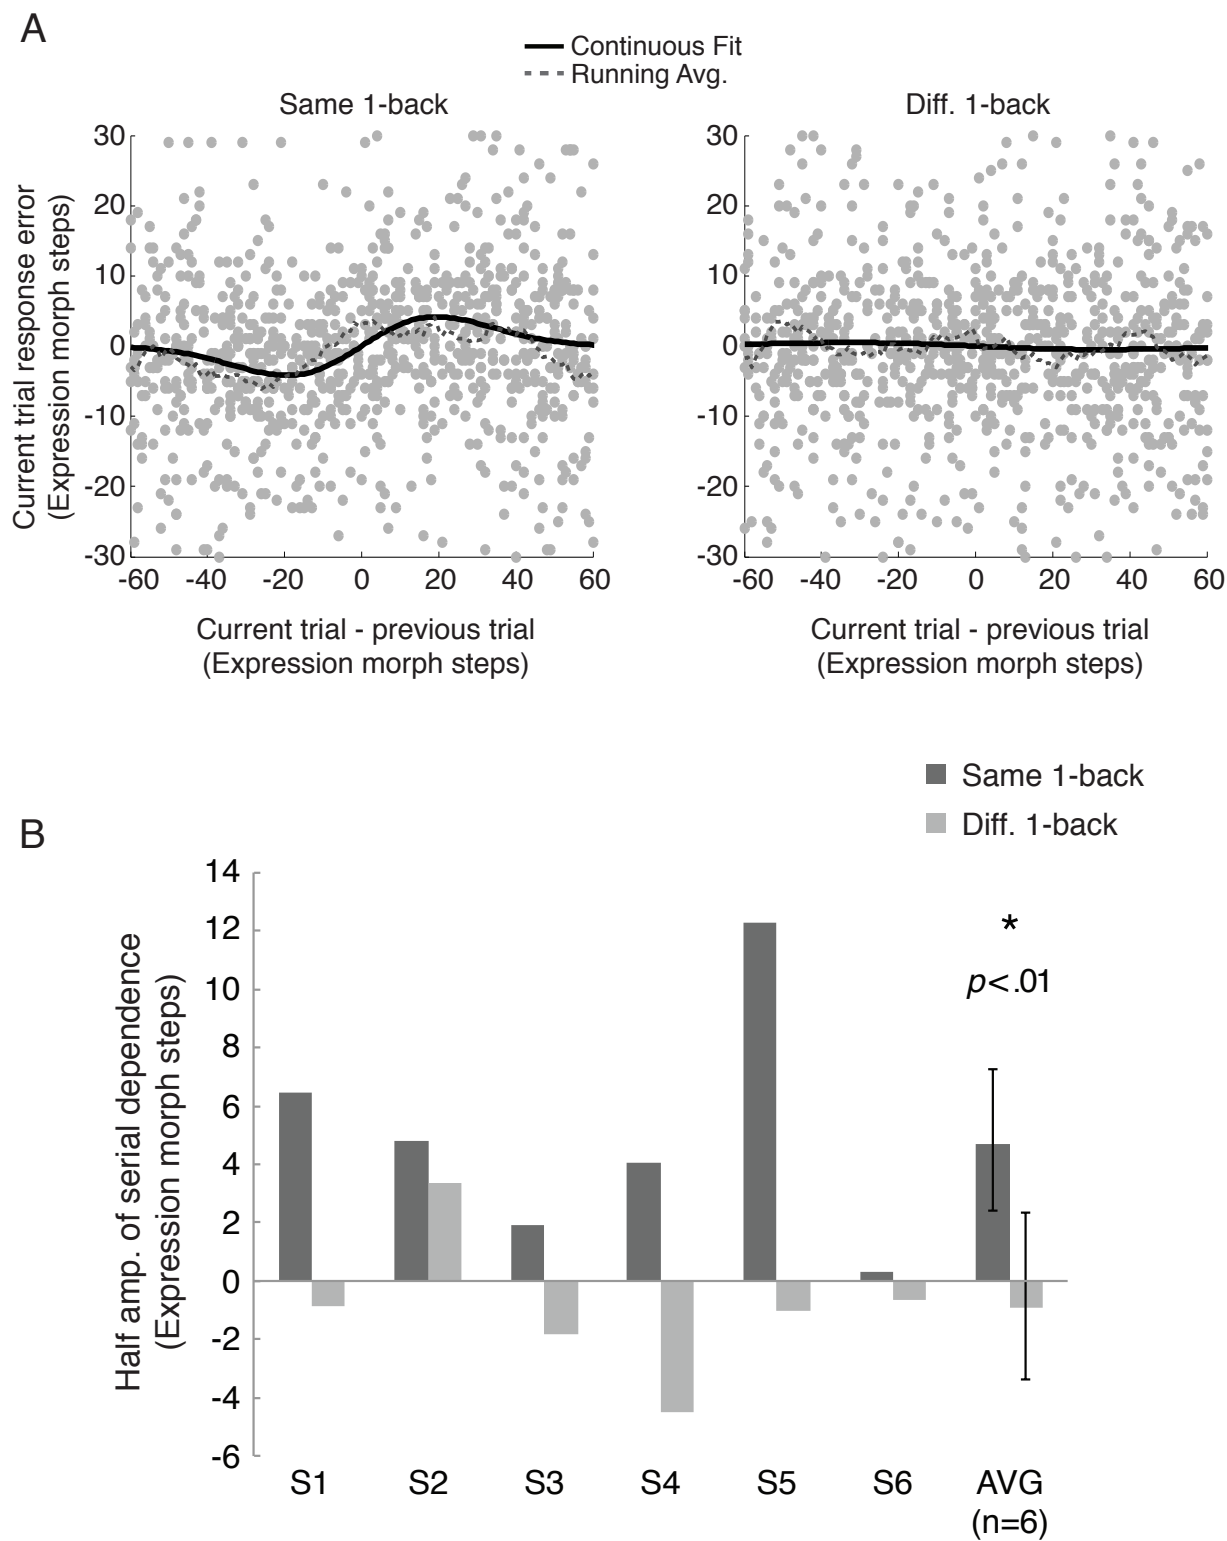


Figure 3.3. Experiment 2 Results. **(A)** Collapsed data from all subjects for trials with the Same 1-back identity (left), and collapsed data from all subjects for all trials with a different 1-back identity (right). Each data point shows performance on one trial. **(B)** Half-amplitude of the serial dependence for each

individual subject in Experiment 2 for same 1-back and different 1-back trials. On average, there was a significantly larger amplitude of serial dependence for Same 1-back trials compared to Different 1-back trials ($p < .001$, permuted null distribution). Error bars are bootstrapped 95% confidence intervals and p -value is based on group permuted null distribution. Subject 2 - Subject 6 also participated in Experiment 1, and their data can be seen in the Fig. 3.1D.

Results

The half-amplitude of serial dependence was significantly larger for the Same 1-back trials compared to the Different 1-back trials ($p < 0.001$, $n=6$, group permuted null, **Figure 3.3B**), with 3 of the 6 subjects showing significantly larger serial dependence ($p < .05$, permuted null). Additionally, the Same 1-back trials showed an overall positive serial dependence effect, replicating the results from Experiment 1 ($p < 0.001$, $n=6$, group permuted null, **Figure 3.3B**). The Different 1-back trials showed no overall serial dependence effect for perceived expression ($p=.8$, $n=6$, group permuted null, **Figure 3.3B**). Therefore, perceived facial expression was pulled towards the facial expression seen one trial ago, but only if the current and previous facial expression came from similar identities. Average response time (RT) across subjects was 2874 ms ($SD = 1553$ ms) for this experiment. This result suggests that the object-selective continuity field maintains the stability of perceived facial expressions over time in an identity-dependent manner. These results also demonstrate that the positive serial dependence found in the same 1-back trials is not entirely due to previous motor responses or general response biases (Luce & Green, 1974; Tanner et al., 1970; Wiegersma, 1982a; 1982b), since responding on the previous trial did not elicit a serial dependence effect in both conditions.

Discussion

Our experiments demonstrated that perceived emotional expression was perceptually pulled by expressions seen up to two trials previously (6 to 12 seconds ago). Furthermore, this serial dependence effect was selective to the identity of the previously seen faces. We saw no perceptual pull on the current expression if the previously seen facial expression had a dissimilar identity. The continuity field is therefore a mechanism that helps maintain the perceptual stability of emotional expression, in addition to facial identity and low level features (Fischer & Whitney, 2014; Liberman et al., 2014).

Several alternative explanations for our results can be ruled out. A generalized response bias, or motor serial dependence would not predict the serial dependence we report (Luce & Green, 1974; Tanner et al., 1970; Wiegersma, 1982a; 1982b), since subjects did not show expression serial dependence when they responded to a 1-back face from a different identity. Furthermore, adaptation and associated negative aftereffects, priming, and other phenomena show a type of perceptual dependence on the recent past, yet remain distinct from serial dependence and the CF. Adaptation studies show that prior exposure to a variety of stimulus features (Anstis et al., 1998; Campbell & Maffei, 1971; Webster et al., 2004; Webster & Mollon, 1991) results in a stimulus-specific negative aftereffect, or perceptual repulsion, away from the adapting stimulus (for reviews see (Thompson & Burr, 2009; Webster, 2012)). Additionally, both emotional expression and facial identity have previously been reported to exhibit negative aftereffects (Carbon & Leder, 2005; Fox et al., 2008; Fox & Barton, 2007; Leopold et al., 2005; Rhodes et al., 2007; Tillman & Webster, 2012; Webster et al., 2004; Webster & Maclin, 1999).

However, our experiments show a positive perceptual pull towards the recent past and are therefore not a result of known forms of adaptation. The reason we find serial dependence rather

than a negative aftereffect is likely because of (1) the brief exposure duration in our experiments (adaptation studies generally expose observers to an image for many seconds or even minutes), and (2) the long interstimulus intervals in our experiments (aftereffects have been observed to decline after the adapting stimulus is removed), as well as the fact that (3) each trial had a random expression, which would tend to wash out adaptation and reduce negative aftereffects. With that in mind, both adaptation and serial dependence are likely operating here and in previous studies, albeit with different time courses, as found in the orientation domain (Fischer & Whitney, 2014). Furthermore, for adaptation, emotional expression is also identity centered—aftereffects are tuned to identity (Fox et al., 2008; Fox & Barton, 2007; Schweinberger et al., 1999; Schweinberger & Soukup, 1998; Wild-Wall et al., 2008). Although serial dependence and adaptation for facial expression show different time scales and perceptual effects, they do both seem to conform to the identity of the face.

The perceptual serial dependence we report may be related to priming effects (Faivre & Kouider, 2011; Kristjánsson et al., 2008; Kristjánsson et al., 2009; Maljkovic & Nakayama, 1994; 1996), but there are important differences. Priming generally manifests in reaction time (Kahneman, Treisman, & Gibbs, 1992; Maljkovic & Nakayama, 1994; 1996) and, where relevant, can improve discriminability of primed stimuli (Sigurdardottir et al., 2008); serial dependence does not impact reaction time and is a reduction in the discriminability of objects (Fischer & Whitney, 2014; Liberman et al., 2014), for the sake of perceptual stability. The CF is a spatiotemporal operator that affects appearance: it makes (even slightly different) objects look the same over time. The CF is one mechanism (of potentially many) that could generate effects that may appear similar to priming, but is in fact distinct from previously reported priming effects. This is not to say that priming, adaptation, and serial dependence are unrelated, as they may play complementary roles in helping to establish and/or maintain perceptual stability.

In summary, our results demonstrate a serial dependence in perceived emotional expression that is identity selective. A continuity field may therefore operate on perceived expression of faces for the purpose of enhancing perceptual stability in an identity-dependent manner. By recycling previously perceived identities and expressions, the object-selective CF decreases the neural computations necessary for the identification of similar objects over time.

Chapter 4: Serial dependence promotes object stability during occlusion

From moment to moment, we perceive the identities of objects and people as stable even though their image properties frequently change due to factors like occlusion, changes in viewpoint, and eye movements. Whether and how the visual system promotes perceived object continuity over time remains an important question. Recent experiments have demonstrated a serial dependence in perception: a bias in the perceived identity of objects toward similar objects seen in the last few seconds (Cicchini et al., 2014; Fischer & Whitney, 2014; Liberman et al., 2014). Intriguingly, visual processing may mirror the temporal regularity of objects in the physical world through a perceptual bias that promotes stability of identity perception.

Previous studies revealed that the perception of orientation, numerosity, and other low-level features is indeed serially dependent—systematically biased towards similar visual input from the recent past (Cicchini et al., 2014; Corbett et al., 2011; Fischer & Whitney, 2014). Beyond these basic features, face identity perception is also systematically biased towards identities seen up to several seconds prior, even across viewpoint changes (Liberman et al., 2014). This serial dependence is tuned in feature space (object similarity; (Fischer & Whitney, 2014; Liberman et al., 2014)) as well as in space and time. The spatio-temporal region over which current object features, such as orientation, are pulled by previously seen features is known as the Continuity Field (CF). For orientation, the spatial tuning of the CF extends over 20 or more degrees of visual space and has a spatiotopic component (Fischer & Whitney, 2014). The CF is temporally tuned over the past 5-15 seconds; moreover, the CF is gated by attention, meaning that the perception of orientation is only serially dependent on previously attended orientations (Fischer & Whitney, 2014).

Although serially dependent perception manifests as a misperception, it could be adaptive: visual processing echoes the stability of objects in the world to create perceptual continuity. This would be especially helpful in cases where there is noise, occlusion, or discontinuities in the retinal image. For example, how similar does an object look after it is temporarily obscured and then reappears (e.g., because it moves behind an occluder, because of motion parallax, etc.)? If the object-selective continuity field facilitates the stability of an object's appearance in the presence of noise or occlusion, then we should find serial dependence in the perception of an object that reemerges from an occluder.

A great deal of research has investigated dynamic occlusion and the processes that support tracking moving objects behind occluders, often referred to by names such as object permanence, the tunnel effect, amodal integration, etc. (Baillargeon, Spelke, & Wasserman, 1985; Burke, 1952; Flombaum & Scholl, 2006; Michotte, Thines, & Crabbé, 1991). However, previous work does not make a clear prediction about whether there is serial dependence in the perception of objects that enter and reemerge from behind occluders. It is well known that observers perceive objects moving continuously behind an occluder as following a single trajectory, even if those objects look different before entering and after exiting from behind the occluder (Burke, 1952). Additionally, detecting the change in a object's color after occlusion improves when the object is perceived as moving along a continuous trajectory (Flombaum & Scholl, 2006), indicating that we perceive objects as persisting even when we cannot see them. Furthermore, object-based attention increases for occluded objects compared to when those same objects are visible, suggesting that attentional resources increase during object occlusion (Flombaum et al., 2008). This knowledge of object continuity and persistence during occlusion is

learned at a young age and is present early in infancy (Carey & Xu, 2001; Spelke et al., 1995). Despite a great deal of research, it still remains to be explored how an object *appears* as it reemerges from behind an occluder.

If the continuity field helps maintain perceived object identities during visual discontinuities, then we should be able to measure serial dependence through occlusion—sequential objects should look more similar than they actually are. Moreover, we should find stronger serial dependence for continuous trajectories that are expected and natural, as opposed to trajectories that are unpredictable or unexpected. Note that the change detection experiments mentioned earlier (Flombaum & Scholl, 2006) do not anticipate these results: serial dependence makes similar objects look identical, which should weaken change detection and increase change blindness. The seeming contradiction in these predictions is easily reconciled because these change detection experiments tested sequential objects that were easily discriminable (Flombaum & Scholl, 2006). Thus, no serial dependence would be expected in those cases, because serially dependent perception is tuned to feature and object similarity (Fischer & Whitney, 2014; Liberman et al., 2014). In the following series of experiments, we examined whether perceptual serial dependence facilitates the appearance of stable object identities despite occlusion.

Experiment 1: Does serial dependence stabilize perceived object identity during interrupted visibility?

If the object-selective continuity field facilitates the stability of an object's appearance in the presence of noise or occlusion, then we should find serial dependence when that object emerges from behind an occluder. The object's identity should appear to be pulled by the identity of the object that first entered the occluder, especially when it travels behind the occluder in a predictable or expected trajectory. To test this hypothesis, we presented a series of randomly oriented Gabor patches that traveled behind an occluder along either a continuous or discontinuous trajectory (**Figure 4.1**). Our question was whether the perceived orientation of the Gabor that emerged from the occluder was serially dependent on the orientation of the Gabor that entered the occluder.

Methods

Stimuli

Experiments were conducted in a darkened experimental booth. Subjects viewed stimuli on a CRT monitor (1024 x 768, 100Hz, Dell Trinitron) at a distance of 56 cm. All experiments were programmed in MATLAB (The MathWorks, Natick, MA) using Psychophysics Toolbox (Brainard, 1997). The stimulus consisted of drifting Gabor patches. Each Gabor patch was a sine wave grating (carrier) 4 cycle/deg, 29% Michelson contrast, with a Gaussian contrast envelope (SD of 4 deg) and Brownian noise ($1/f^2$ spatial noise). In Experiment 1, subjects saw a Gabor patch traveling behind an occluder (26 degrees high, 13.25 degrees wide; **Figure 4.1A**). The Gabor always appeared on the left side of the screen and traveled towards the right side of the screen. The orientation of each Gabor that appeared (both before entering the occluder as well as after exiting the occluder) was randomly selected from a uniform distribution, with possible orientations of 1 to 180 degrees. Thus, the orientation of the Gabor patch could be very similar or very different when it re-emerged from the occluder. The Gabor patch traveled across the screen at a speed of 19.9 degrees per second.

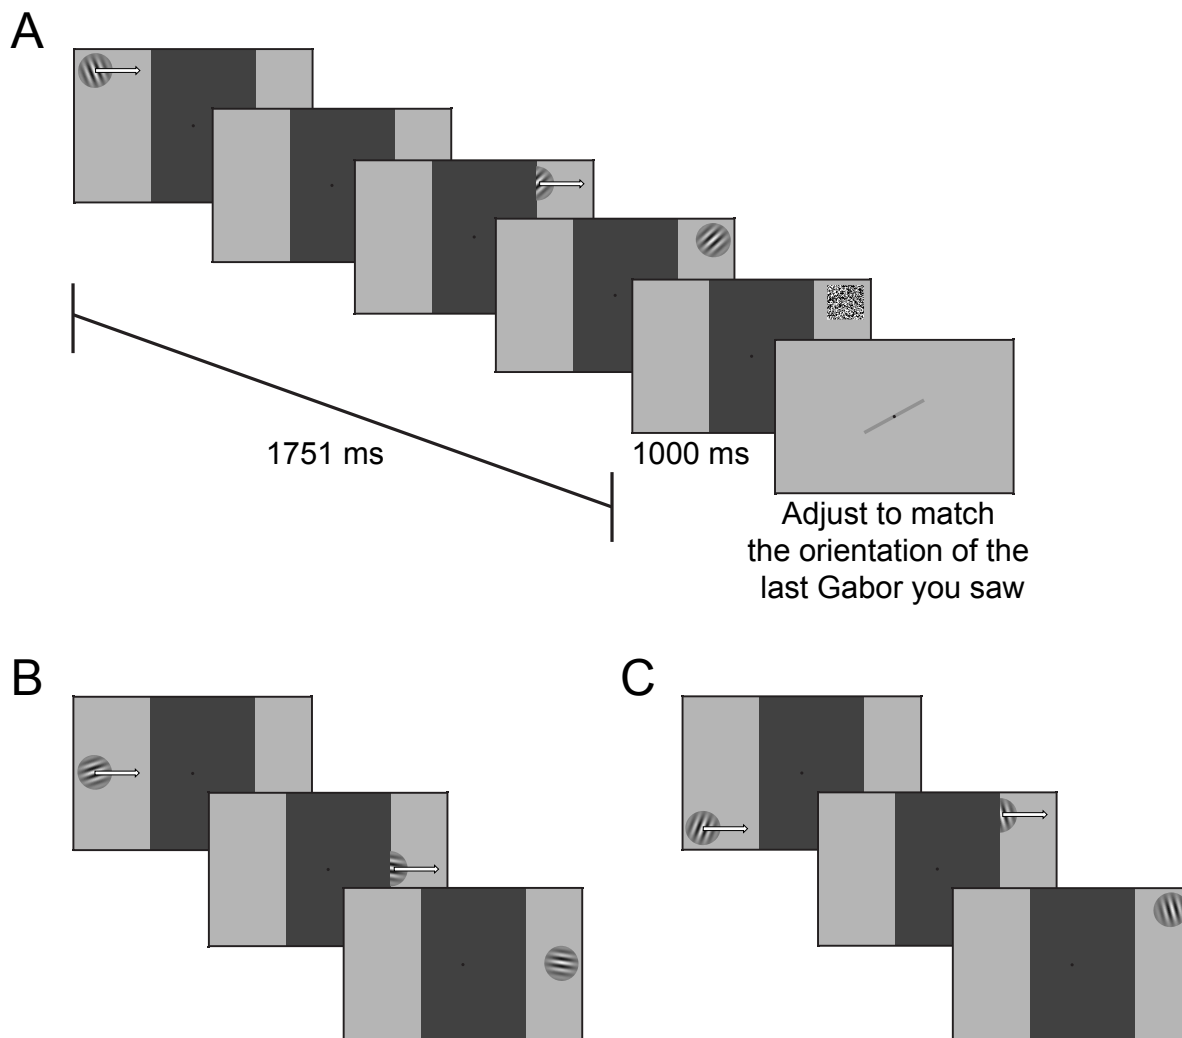


Figure 4.1. Experiment 1 stimuli and procedure. **(A)** Example continuous moving trajectory trial sequence for a Gabor with a starting position at the top left corner of the screen. Subjects responded by adjusting a rectangular bar (.24 x 4 degrees) at fixation to match the perceived orientation of the second, exiting Gabor patch. **(B)** Example continuous moving trajectory path for a Gabor with a starting position in the middle of the screen. **(C)** Example discontinuous moving trajectory trial for a Gabor with a starting position at the bottom left corner of the screen.

Participants

All experimental procedures were approved by the UC Berkeley Institutional Review Board and were in accordance with the Declaration of Helsinki. Participants were affiliates of UC Berkeley and provided written informed consent before participation. A total of 13 subjects (6 female) participated in this experiment, ranging in age from 19 to 37 years ($M = 27.5$, $SD = 5.4$). The partially collected data from two participants was discarded because they were not available to complete the full experiment, resulting in a total of 11 subjects in Experiment 1. All participants

had normal or corrected-to-normal vision, and all except one were naïve to the purpose of the experiment. To determine our sample sizes, we estimated the anticipated effect size from previously reported data (Fischer & Whitney, 2014) for spatially and temporally separated Gabors. For these experiments, our anticipated effect size (Cohen's d) for our Experiment 1 trial conditions was $d=3$, which required a minimum of 6 subjects for a power of .8 and probability level of $p=.05$ (2-tailed). We stopped data collection once we reached twice the minimum number of subjects (12), but collected data from one additional subject due to subject exclusions.

Procedures

In this experiment, subjects maintained fixation on a central point while their head was stabilized in a chin rest as they viewed Gabor stimuli. They saw a Gabor patch appear either in the upper left hand corner, lower left hand corner, or middle left edge of the screen (**Figure 4.1**). This first Gabor patch (entering Gabor) had a random orientation (1-180 degrees) and traveled across the screen towards the edge of an occluder for 610 ms at a fixed speed. The Gabor was no longer visible once it traveled behind the occluder for 721 ms. A second Gabor patch (exiting Gabor) with a new, random orientation then exited from the opposite side of the occluder and traveled for 420 ms to the right edge of the screen, where it stopped and disappeared. It was replaced by a 1000 ms random noise patch to reduce any afterimages. The exiting Gabor moved away from the occluder in a trajectory that was either consistent with the path of the entering Gabor (continuous straight moving trajectory, **Figure 4.1B**), or along a path that was not consistent with the entering Gabor (discontinuous moving trajectory, **Figure 4.1C**). Subjects saw a total of 3 possible continuous trajectories: the Gabor moved horizontally across the screen and behind an occluder along the top, middle, or bottom of the screen. We generated the discontinuous trajectory by randomizing the location of the exiting Gabor to be one of the possible alternate trajectories that were not continuous with the initial entering Gabor. For example, a (discontinuous) Gabor whose initial position on the screen was in the top left corner would exit the occluder from either the middle or bottom of the screen.

There were three randomized trial types in this experiment: continuous moving trials (40% of trials), discontinuous moving trials (40% of trials), and catch moving trials (20% of trials). For each non-catch trial, subjects responded by using the keyboard to adjust a randomly oriented rectangular bar (.24 by 4 degrees, 10.2 cd/m²) at fixation to match the perceived orientation of the exiting Gabor patch. After adjustment, the rectangular bar was replaced by a 1500 ms noise patch to reduce any afterimages. Participants then saw a fixation dot for 1000 ms before beginning the next trial. During the catch moving trials, subjects only saw the entering Gabor and had to respond to its orientation, without ever seeing an exiting Gabor. The purpose of these surprise catch trials was to make sure subjects attended to the orientation of both the entering and exiting Gabors, since serial dependence has been found to be gated by attention (Fischer & Whitney, 2014). The entering Gabor was on the screen longer than the exiting Gabor to further facilitate subjects' attention to that Gabor. Subjects completed 2 sessions of 350 trials each. We predicted that the amplitude of the serial dependence effect should be higher for continuous versus discontinuous trials.

Data Analysis

For both continuous and discontinuous moving trials, response error was computed as the difference between the response orientation and the physical orientation of the exiting Gabor. Response error was then compared to the difference between the exiting and entering Gabor

orientations. Positive values on the abscissa indicate that the entering Gabor was more clockwise than the exiting Gabor, and positive errors on the ordinate indicate that the reported orientation was more clockwise than the actual orientation of the exiting Gabor. For each subject's data, trials were considered lapses and excluded if error exceeded 3 standard deviations from the mean (less than 2% of data excluded on average) or if the response time was longer than 10 seconds.

We fit a simplified Gaussian derivative using constrained nonlinear minimization (Nelder-Mead) to each subject's data for continuous and discontinuous trials separately using the following equation:

$$y = abcxe^{-(bx)^2}$$

Parameter y is response error on each trial (response Gabor orientation – exiting Gabor orientation), x is the difference between the entering and exiting Gabor orientation, a is half the peak-to-trough amplitude of the derivative-of-Gaussian, b scales the width of the Gaussian derivative, and c is a constant, $\sqrt{2}/e^{-0.5}$, which scales the curve to make the a parameter equal to the peak amplitude. We used the a parameter in the above equation from each subject's data as a metric of serial dependence—the degree to which subjects' reports of orientation of the exiting Gabor were pulled in the direction of the entering Gabor (**Figure 4.2A**). If a subject's perception of orientation were repelled by the entering Gabor (e.g., because of a negative tilt aftereffect (Campbell & Maffei, 1971)) or not influenced by the entering Gabor, then the half-amplitude of the best fitting Gaussian derivative should be negative or close to zero, respectively. We then ran a 2-tailed paired Student t-test to compare the amplitude of serial dependence for continuous (**Figure 4.2A**) and discontinuous (**Figure 4.2B**) moving trials across all subjects. Effect sizes for within-subject comparisons were calculated using Cohen's d_{av} (Lakens, 2013). We used a false discovery rate (FDR) procedure to correct for multiple comparisons (Benjamini & Hochberg, 1995).

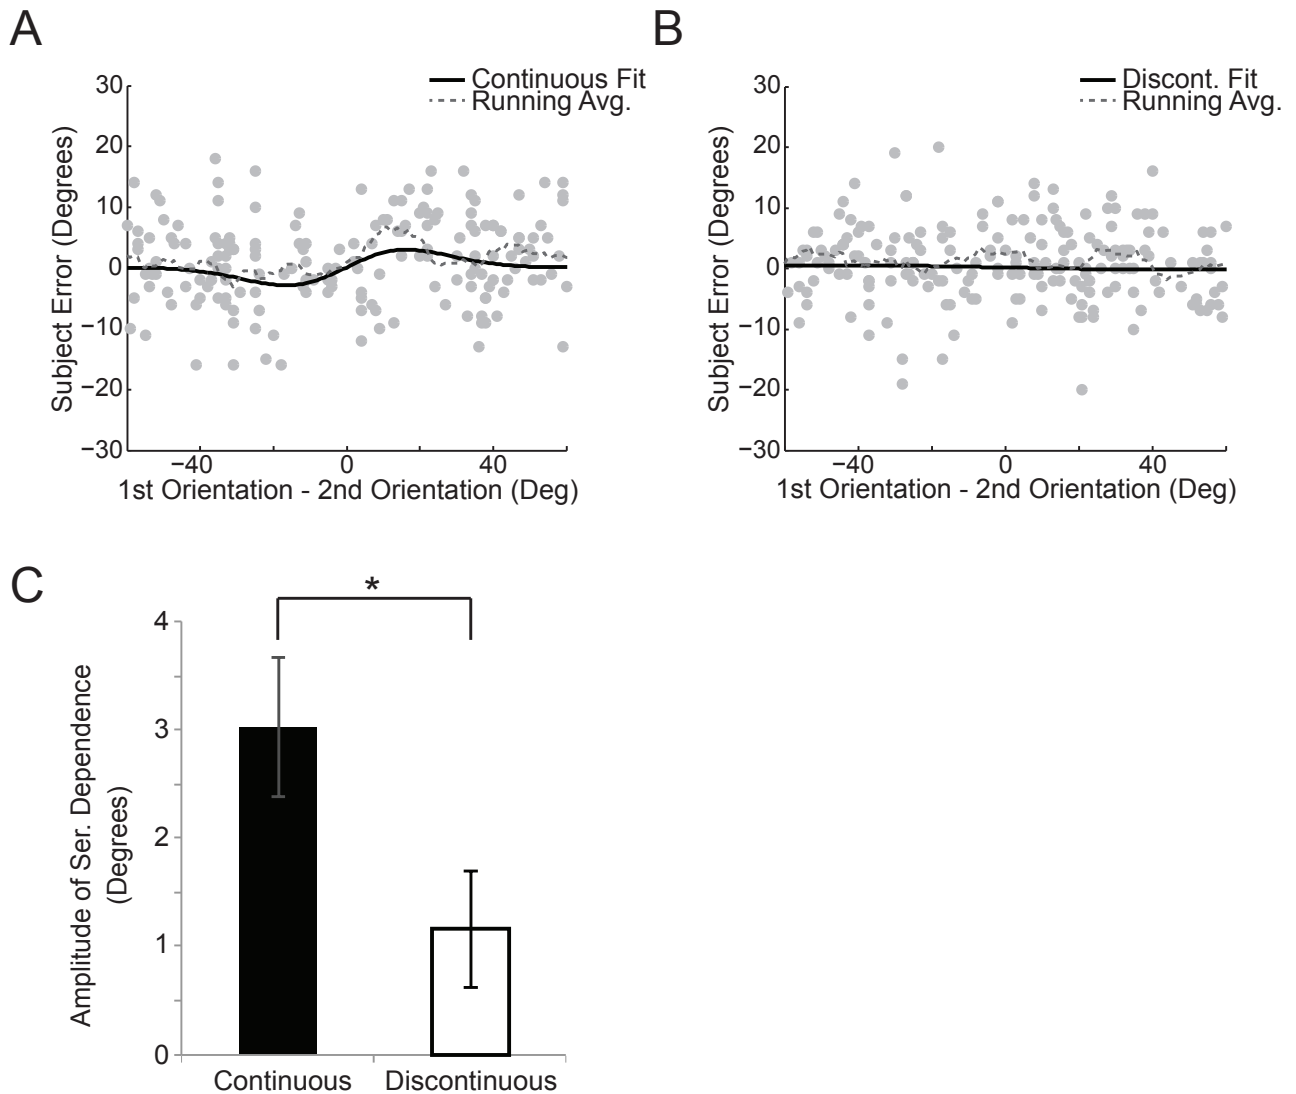


Figure 4.2. Experiment 1 results. **(A)** Example data from a representative subject for all continuous moving trajectory trials. The derivative of a Gaussian (solid black line) was fit to the entire range of the data. Here, for visualization, the x-axis is truncated slightly. **(B)** Example data from the same representative subject for all discontinuous moving trajectory trials. **(C)** Average amplitude of serial dependence across 11 subjects for continuous and discontinuous trials. Error bars are standard error of the mean (SEM). Serial dependence was significantly stronger when the object moved along a continuous versus discontinuous trajectory behind the occluder.

Results

The half-amplitude of serial dependence was significantly larger for the continuous compared to the discontinuous moving trials (**Figure 4.2C**; $t(10) = 4.69$, $p = .003$, $d_{av} = .92$, FDR corrected 2-tailed paired t-test). Additionally, both the continuous and discontinuous trials

showed an overall positive serial dependence effect ($t(10) = 4.69$, $p = .003$, $d = 1.41$, continuous trials; $t(10) = 2.05$, $p = .11$, $d = .62$, discontinuous trials; FDR corrected 2-tailed t-test), as predicted by the temporal and spatial tuning of the continuity field (Fischer & Whitney, 2014). Average subject response time across all trials was 1828 ms ($SD = 910$ ms). The perceived orientation of the exiting Gabor was pulled toward the orientation of the entering Gabor, and this effect was significantly stronger when the object moved along a continuous trajectory behind the occluder. This result suggests that object continuity modulates serial dependence in orientation perception: the perceptual pull from previous stimuli was enhanced when the drifting Gabor was perceived as traveling along a continuous, predictable path.

Experiment 2: Is motion necessary for perceived perceptual continuity?

In Experiment 1, we found that serial dependence helps promote spatiotemporal continuity by making objects that enter and exit from behind an occluder in a predictable way appear more similar than they truly are. Because serial dependence also pulls the apparent orientation of sequentially presented static objects (Fischer & Whitney, 2014), we asked whether motion was, in fact, important for the exiting Gabor's orientation to be perceptually pulled towards the entering Gabor. To test this, we presented subjects with a series of static, randomly oriented Gabor patches on either side of an occluder (**Figure 4.3**). Since there were less continuity cues for the static Gabors, we predicted that there would be significantly less serial dependence when objects did not move along a trajectory.

Methods

Stimuli

In Experiment 2, subjects saw two stationary Gabors appear on either side of an occluder. The Gabor stimuli had all of the same properties as the Gabor stimuli in Experiment 1, but they were stationary. The occluder was the same size and in the same location as Experiment 1 (**Figure 4.3A**). Subjects reported the perceived orientation of the last Gabor that they saw by adjusting the orientation of a response bar shown at a central fixation point (as in Experiment 1). Catch trials were also included: on 80% of the trials, the last seen Gabor was the nominally "exiting" Gabor (though it did not move), while on the other 20% of trials, randomly interleaved, the last seen Gabor was the nominally "entering" Gabor (though it did not move). As in Experiment 1, the orientation of the response bar was initially random on each trial and subjects used a keyboard for all responses.

Participants

All experimental procedures were approved by the UC Berkeley Institutional Review Board and were in accordance with the Declaration of Helsinki. Participants were affiliates of UC Berkeley and provided written informed consent before participation. A total of 12 subjects (6 female) participated in this experiments, ranging in age from 19 to 37 years ($M = 25.9$, $SD = 4.9$). 11 of the subjects also participated in Experiment 1, and one subject's data was discarded because the subject was not available to complete the full experiment. All participants had normal or corrected-to-normal vision, and all except one were naïve to the purpose of the experiment.

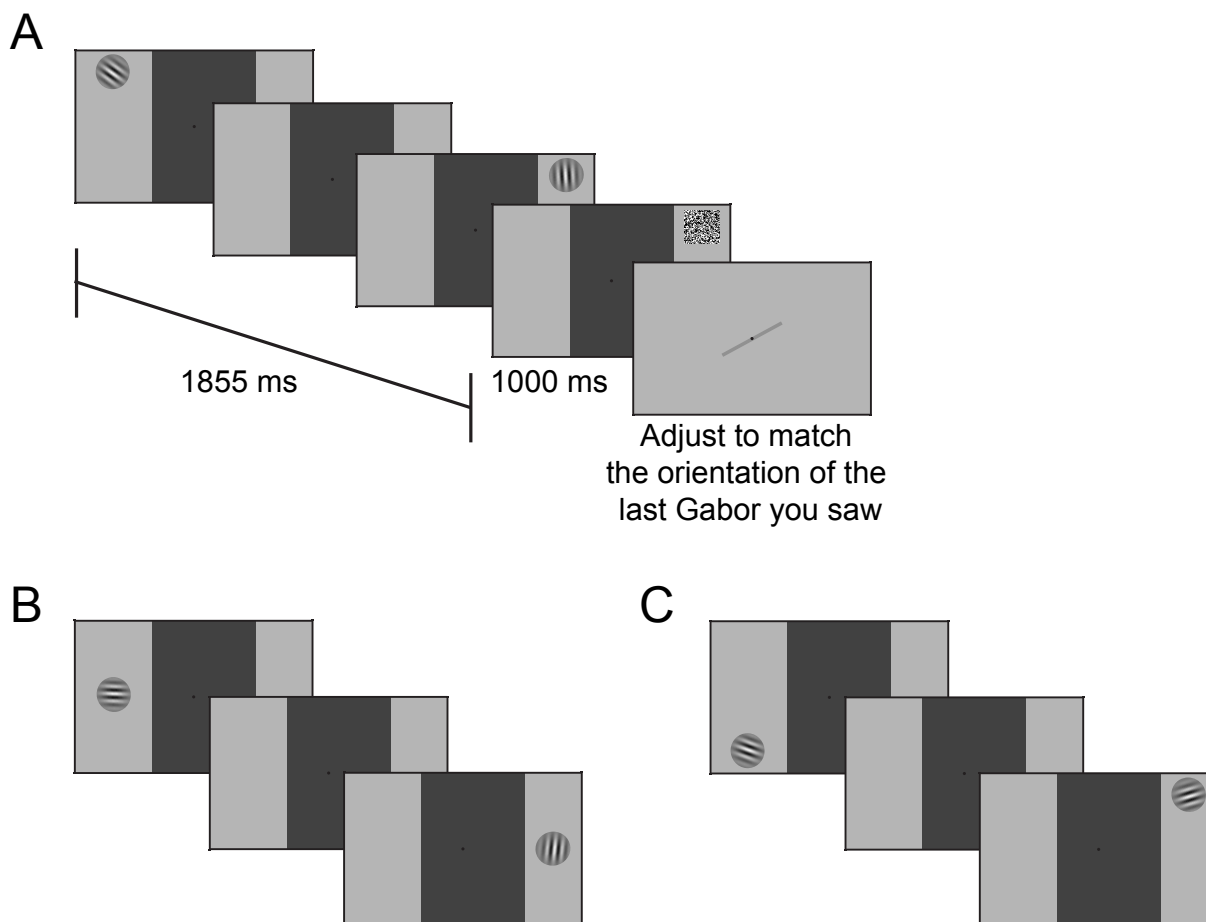


Figure 4.3. Experiment 2 stimuli and procedure. **(A)** Example continuous static trajectory trial sequence for a Gabor with a starting position at the top left of the screen. Subjects responded by adjusting a rectangular bar (.24 x 4 degrees) at fixation to match the perceived orientation of the exiting Gabor patch. **(B)** Example aligned static trial for a Gabor with a starting position in the middle of the screen. **(C)** Example misaligned static trial for a Gabor with a starting position in the bottom left of the screen.

Procedures

The stimuli and procedures were identical to those in Experiment 1, except that the Gabors remained stationary throughout all trials. The Gabor first appeared centered halfway between the left side of the screen and the closest edge of the occluder to try to match the average location of the moving Gabor. It stayed in the same initial location without moving for 610 ms (for consistency we will still refer to this as the “entering” Gabor, although it did not actually move or display any accretion cues). It then disappeared for an 885 ms interstimulus interval. Another static Gabor (akin to the “exiting” Gabor) then appeared centered between the right edge of the occluder and the right side of the screen for 360 ms before being replaced by a 1000 ms noise patch (**Figure 4.3A**). The possible locations for the Gabor patches were based on the locations used for the drifting patches in Experiment 1; the entering Gabor could be in any one of the three horizontal locations as the entering Gabor in Experiment 1 and the exiting Gabor could be in any one of three horizontal locations of the exiting Gabor in Experiment 1. The

“exiting” Gabor patch in this experiment was shown in a location that was either analogous to a continuous trajectory in Experiment 1 (aligned static trials; **Figure 4.3B**), or analogous to a discontinuous trajectory in Experiment 1 (misaligned static trials; **Figure 4.3C**). Subjects completed 2 sessions of 350 trials each.

Data Analysis

All data analysis procedures for Experiment 2 were identical to Experiment 1. We compared response error to the difference between the exiting and entering Gabor orientations. We used the same curve fitting method as in Experiment 1 to determine whether the exiting Gabor was perceptually pulled toward the entering Gabor’s orientation while they both remained stationary.

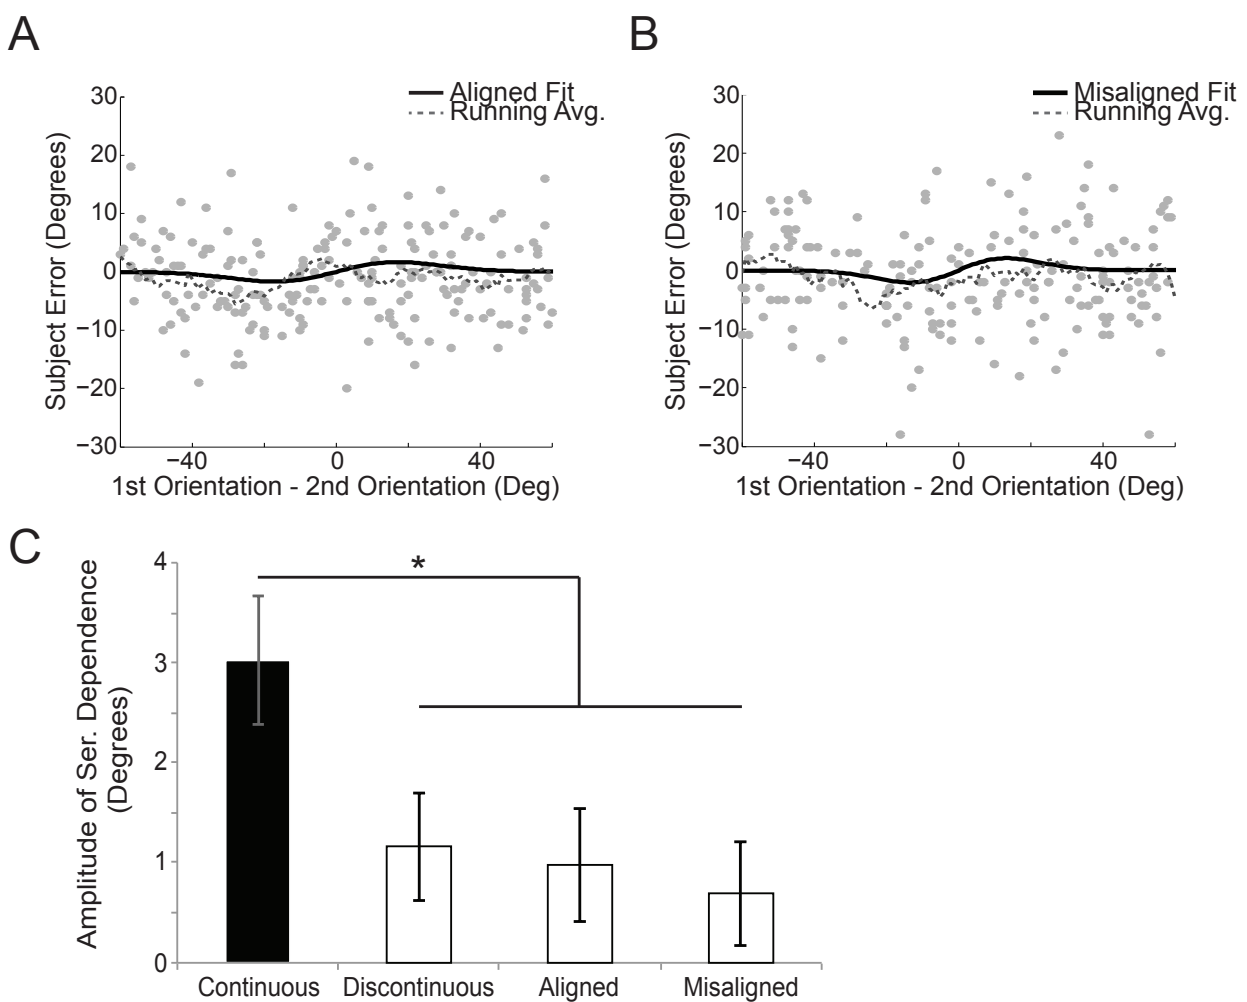


Figure 4.4. Experiment 2 results. (A) Example data from a representative subject for all aligned static trajectory trials. (B) Example data from the same representative subject for all misaligned static trajectory trials. (C) Amplitude of serial dependence compared across both experiments. Error bars are SEM. Serial dependence was significantly stronger when the object moved along a perceived continuous trajectory, versus a discontinuous trajectory or static presentation ($t(10) = 4.69$, $p = .003$, $d_{av} = .92$, continuous vs.

discontinuous; $t(10) = 2.68$, $p = .046$, $d_{av} = 1.07$, continuous vs. aligned static; $t(10) = 2.88$, $p = .044$, $d_{av} = 1.2$, continuous vs. misaligned static; FDR corrected 2-tailed paired t-test).

Results

We compared the half-amplitude of serial dependence for trials when subjects saw the second static Gabor in a horizontally aligned versus misaligned location on the screen (**Figure 4.4**). Average subject response time across all trials was 1802 ms ($SD = 959$ ms) for this experiment. The half-amplitude of serial dependence was not significantly different for the aligned compared to the misaligned static trials (**Figure 4.4C**; $t(10) = .36$, $p > .250$, $d_{av} = .17$, FDR corrected 2-tailed paired t-test). Neither the aligned nor misaligned trials showed a significant serial dependence effect ($t(10) = 1.93$, $p = .11$, $d = .58$, aligned trials; $t(10) = 1.33$, $p = .24$, $d = .4$, misaligned trials; FDR corrected 2-tailed t-test). However, they did show an obviously positive trend. More importantly, the amplitude of serial dependence was significantly stronger in the continuous moving trials (Experiment 1) compared to either the aligned static or misaligned static trials (**Figure 4.4C**; $t(10) = 2.68$, $p = .046$, $d_{av} = 1.07$, continuous vs. aligned static; $t(10) = 2.88$, $p = .044$, $d_{av} = 1.2$, continuous vs. misaligned static; FDR corrected 2-tailed paired t-test).

General Discussion

The Continuity Field is a mechanism proposed to facilitate the perceptual stability of objects (Fischer & Whitney, 2014; Liberman et al., 2014). It does so through serial dependence, which is a systematic bias in the appearance of an object's features and identity toward similar objects seen in the last several seconds. Our results demonstrate that serial dependence in orientation perception operates across interruptions in visibility. Importantly, this serial dependence was stronger for continuously moving trajectories, compared to unexpected trajectories or static displacements, reflecting expectations of a stable and predictable world.

Several alternative explanations for our results can be ruled out. A generalized response bias, or motor serial dependence would not predict the spatiotemporal tuning of the serial dependence we report (Luce & Green, 1974; Tanner et al., 1970; Wiegersma, 1982a; 1982b), since subjects only responded to one of the two Gabors shown on each trial. Additionally, static displacements produced little serial dependence in orientation perception; continuous and predictably moving objects resulted in significantly stronger serial dependence. The relatively small serial dependence in the static conditions is consistent with previous results that characterized the spatial tuning of the CF (Fischer & Whitney, 2014); orientation serial dependence decreases as the spatial distance between Gabors increases. Our results are therefore consistent with previous findings and inconsistent with motor and generalized response biases. Furthermore, adaptation and associated negative aftereffects, priming, and other phenomena show a type of perceptual dependence on the recent past, yet remain distinct from serial dependence and the CF. Adaptation studies show that prior exposure to a variety of stimulus features (Anstis et al., 1998; Campbell & Maffei, 1971; Webster et al., 2004; Webster & Mollon, 1991) results in a stimulus-specific negative aftereffect, or perceptual repulsion, away from the adapting stimulus (for reviews see (Thompson & Burr, 2009; Webster, 2012)). However, our experiments show a positive perceptual pull towards the recent past and are therefore not a result of known forms of adaptation. Our proposed experiments may be related to perceptual priming effects (Faivre & Kouider, 2011; Kristjánsson et al., 2008; Kristjánsson et al., 2009; Maljkovic & Nakayama, 1994; 1996), but there are important differences. Priming generally manifests in reaction time (Maljkovic & Nakayama, 1994; 1996) and, where relevant, can improve

discriminability of primed stimuli (Sigurdardottir et al., 2008); serial dependence does not impact reaction time and is a reduction in the discriminability of objects for perceptual stability (Fischer & Whitney, 2014; Liberman et al., 2014). The CF is a spatiotemporal operator that affects appearance: it makes (even slightly different) objects look the same over time. The CF is one mechanism (of potentially many) that could generate effects that may appear similar to priming, but is in fact distinct from previously reported priming effects. This is not to say that priming, adaptation, and serial dependence are unrelated, but they are distinguishable and the complementary roles they play in perceptual stability remain a rich area of investigation.

Previous researchers have suggested that object files (Kahneman et al., 1992; Treisman, 1988) or object tokens (Chun & Cavanagh, 1997; Kanwisher, 1987; Kanwisher & Driver, 1992) maintain perceptual stability by using the perceived spatiotemporal continuity of objects as they move, change or have disruptions in visibility. However, object tokens or files are a conceptual and not mechanistic description of the ability to track objects during occlusion. They may contribute to or even generate object permanence (Kahneman et al., 1992), but they do not make a clear prediction about what objects should *look like* after emerging from dynamic occlusion, or whether there should or should not be serial dependence.

Furthermore, object file effects are mainly found in RT improvements, similar to priming (Kahneman et al., 1992). We did not find any significant RT effects in our data, which we tested by fitting both a Gaussian and a straight line to normalized trial RTs (y-axis) relative to the difference between the first and second Gabor orientation (x-axis) for continuous moving trajectory trials. We tested model significance using AIC (Akaike, 1974), and found that the straight line had the lowest AIC in 9/11 subjects, indicating no RT improvements for trials with the largest serial dependence effects.

Object tokens are also often invoked in studies of change detection through occlusion (Flombaum et al., 2008; Flombaum & Scholl, 2006). Our measured serial dependence appears to be inconsistent with those predictions (or findings) of change detection experiments (Flombaum et al., 2008; Flombaum & Scholl, 2006), at least until one considers the feature tuning of the CF (Fischer & Whitney, 2014; Liberman et al., 2014). Experiments that employ change blindness as a tool typically use sequential objects that are very distinct (Flombaum & Scholl, 2006), where little or no serial dependence would be expected. In contrast, serial dependence is strongest (represented as a ratio of the difference between the sequential objects) for similar objects. Serial dependence should therefore increase change blindness, but only for similar objects. Moreover, change detection experiments require a comparison over time—is the first object similar to the second- and therefore involve working memory. In contrast, our task does not require an explicit comparison between the first and second displays; it does not tax memory, and thus may not involve the same working memory demands.

Conclusion

Our results suggest that the CF is tuned to the spatiotemporal predictability of dynamic objects for the purpose of perceptual stability. The CF is highly tuned to continuous and expected object trajectories, mirroring the predictability of objects in the world and reflecting expectations of stability. The CF therefore provides a key mechanism that could support functions of perceptual stability during object tracking and while objects are occluded.

Chapter 5: Conclusion

The continuity field is a mechanism proposed to facilitate the perceptual stability of objects across time (Fischer & Whitney, 2014; Liberman et al., 2014). It does so through serial dependence, which is a systematic bias in the appearance of an object's features and identity toward similar objects seen in the last several seconds. In this set of experiments, we have demonstrated that the object-selective continuity field helps maintain the stability of perceived facial identity, perceived facial expression of similar identities, and perceived orientation across interruptions in visibility. These results bring us closer to answering how the visual system pieces together a coherent and stable perception of the world.

In Chapter 2, we found that people's perception of facial identity was significantly influenced by identities encountered 5 or more seconds prior. This was not a result of response bias since a response on the previous trial was not required. In addition, this serial dependence was perceptually tuned to face identity such that there was no perceptual pull on the current face if the previously seen face had a sufficiently different identity. Importantly, we found that the continuity field for facial identity operates across different face viewpoints, indicating that this serial dependence occurs at the level of object-centered perceptual representations and does not depend on low-level stimulus features.

In Chapter 3, we found that the continuity field stabilizes the expression attributed to an individual face by making it perceptually biased towards recently seen expressions. Our experiments showed that subjects made consistent errors when reporting the expression of the target face, seeing it as more similar to the facial expressions presented on the previous two trials (up to 12 seconds previously). Additionally, we found stronger serial dependence in expression when the 1-back identity was the same as the current identity. Overall, these results indicate that perceived facial expression at any one moment is biased towards recently seen expressions, especially within the same identity. The continuity field may therefore support apparent stability of emotional expressions by introducing perceived serial dependence in an identity-selective manner.

Finally, in Chapter 4, our results demonstrated that serial dependence in orientation perception operates across interruptions in visibility. Importantly, this serial dependence was stronger for continuously moving trajectories, compared to unexpected trajectories or static displacements, reflecting expectations of a stable and predictable world. Overall, our findings suggest that the visual system takes advantage of prior visual information when representing objects in the world, even across interruptions in visibility. These findings further our knowledge about object perception and suggest new possibilities for neural and psychophysical research in object processing, highlighting the importance of temporal integration of visual information.

References

- Akaike, H. (1974). A new look at the statistical identification model. *IEEE Trans. Auto. Control*, *19*, 716–723.
- Anstis, S., Verstraten, F. A. J., & Mather, G. (1998). The motion aftereffect. *Trends in Cognitive Sciences*, *2*(3), 111–117.
- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*.
- Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B (Methodological)*, *57*(1), 289–300.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, *77*(3), 305–327.
- Burke, L. (1952). On the tunnel effect. *Quarterly Journal of Experimental Psychology*, *4*(3), 121–138. <http://doi.org/10.1080/17470215208416611>
- Calvo, M., & Esteves, F. (2005). Detection of emotional faces: low perceptual threshold and wide attentional span. *Visual Cognition*, *12*(1), 13–27. <http://doi.org/10.1080/13506280444000094>
- Campbell, F. W., & Maffei, L. (1971). The tilt after-effect: a fresh look. *Vision Research*, *11*(8), 833–840.
- Carbon, C.-C., & Leder, H. (2005). Face adaptation: Changing stable representations of familiar faces within minutes? *Advances in Cognitive Psychology*, *1*(1), 1–7.
- Carey, S., & Xu, F. (2001). Infants' knowledge of objects: Beyond object files and object tracking. *Cognition*, *80*(1), 179–213.
- Chopin, A., & Mamassian, P. (2012). Predictive Properties of Visual Adaptation. *Current Biology*, *22*(7), 622–626. <http://doi.org/10.1016/j.cub.2012.02.021>
- Chun, M. M., & Cavanagh, P. (1997). Seeing Two as One: Linking Apparent Motion and Repetition Blindness. *Psychological Science*, *8*(2), 74–79. <http://doi.org/10.2307/40063149?ref=search-gateway:0a1a4a3f613e3d0fcb9654e6bbb86a19>
- Cicchini, G. M., Anobile, G., & Burr, D. C. (2014). Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proceedings of the*

National Academy of Sciences of the United States of America.
<http://doi.org/10.1073/pnas.1402785111>

- Clifford, C. W., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., & Schwartz, O. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research*, *47*(25), 3125–3131.
- Corbett, J. E., Fischer, J., & Whitney, D. (2011). Facilitating Stable Representations: Serial Dependence in Vision. *PloS One*, *6*(1), e16701.
- Crocker, V., & McDonald, S. (2009). Recognition of emotion from facial expression following traumatic brain injury. *Brain Injury*, *19*(10), 787–799.
<http://doi.org/10.1080/02699050500110033>
- Eagleman, D. M., Jacobson, J. E., & Sejnowski, T. J. (2004). Perceived luminance depends on temporal context. *Nature*, *428*(6985), 854–856. <http://doi.org/10.1038/nature02467>
- Edwards, K. (1998). The face of time: Temporal cues in facial expressions of emotion. *Psychological Science*.
- Efron, B., & Tibshirani, R. (1993). An introduction to the bootstrap.
- Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology*, *17*(2), 124–129.
- Ekman, P., & Friesen, W. V. (1976a). Measuring facial movement. *Environmental Psychology and Nonverbal Behavior*, *1*(1), 56–75. <http://doi.org/10.1007/BF01115465>
- Ekman, P., & Friesen, W. V. (1976b). Pictures of facial affect. *Consulting Psychologists, Palo Alto, CA*.
- Faivre, N., & Kouider, S. (2011). Increased sensory evidence reverses nonconscious priming during crowding. *Journal of Vision*, *11*(13), 16–16. <http://doi.org/10.1167/11.13.16>
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 1–9. <http://doi.org/10.1038/nn.3689>
- Flombaum, J. I., & Scholl, B. J. (2006). A temporal same-object advantage in the tunnel effect: Facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(4), 840–853. <http://doi.org/10.1037/0096-1523.32.4.840>
- Flombaum, J. I., Scholl, B. J., & Pylyshyn, Z. W. (2008). Attentional resources in visual tracking through occlusion: The high-beams effect. *Cognition*, *107*(3), 904–931.
<http://doi.org/10.1016/j.cognition.2007.12.015>

- Fox, C. J., & Barton, J. J. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research, 1127*, 80–89.
- Fox, C. J., Oruc, I., & Barton, J. J. S. (2008). It doesn't matter how you feel. The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision, 8*(3), 11–11. <http://doi.org/10.1167/8.3.11>
- Glasser, D. M., Tsui, J. M. G., Pack, C. C., & Tadin, D. (2011). Perceptual and neural consequences of rapid motion adaptation. *Proceedings of the National Academy of Sciences of the United States of America, 108*(45), E1080–8. <http://doi.org/10.1073/pnas.1101141108>
- Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2000). The distributed human neural system for face perception. *Trends in Cognitive Sciences, 4*(6), 223–233. [http://doi.org/10.1016/S1364-6613\(00\)01482-0](http://doi.org/10.1016/S1364-6613(00)01482-0)
- Hock, H. S., Kelso, J. S., & Schöner, G. (1993). Bistability and hysteresis in the organization of apparent motion patterns. *Journal of Experimental Psychology: Human Perception and Performance, 19*(1), 63.
- Jacobs, D. H., Shuren, J., Bowers, D., & Heilman, K. M. (1995). Emotional facial imagery, perception, and expression in Parkinson's disease. *Neurology*.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*.
- Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*.
- Kanwisher, N., & Driver, J. (1992). Objects, Attributes, and Visual Attention: Which, What, and Where. *Current Directions in Psychological Science, 1*(1), 26–31. <http://doi.org/10.2307/20182120?ref=search-gateway:f6b1fc579d31775500487387647d8364>
- Kirouac, G., & Doré, F. Y. (1984). Judgment of facial expressions of emotion as a function of exposure time. *Perceptual and Motor Skills, 59*(1), 147–150.
- Kohn, A. (2007). Visual Adaptation: Physiology, Mechanisms, and Functional Benefits. *Journal of Neurophysiology, 97*(5), 3155–3164. <http://doi.org/10.1152/jn.00086.2007>
- Kosovicheva, A. A., Maus, G. W., Anstis, S., Cavanagh, P., Tse, P. U., & Whitney, D. (2012). The motion-induced shift in the perceived location of a grating also shifts its aftereffect. *Journal of Vision, 12*(8), 7–7. <http://doi.org/10.1167/12.8.7>
- Kristjansson, A., Ingvarsdottir, A., & Teitsdottir, U. D. (2008). Object- and feature-based priming in visual search. *Psychonomic Bulletin & Review, 15*(2), 378–384. <http://doi.org/10.3758/PBR.15.2.378>

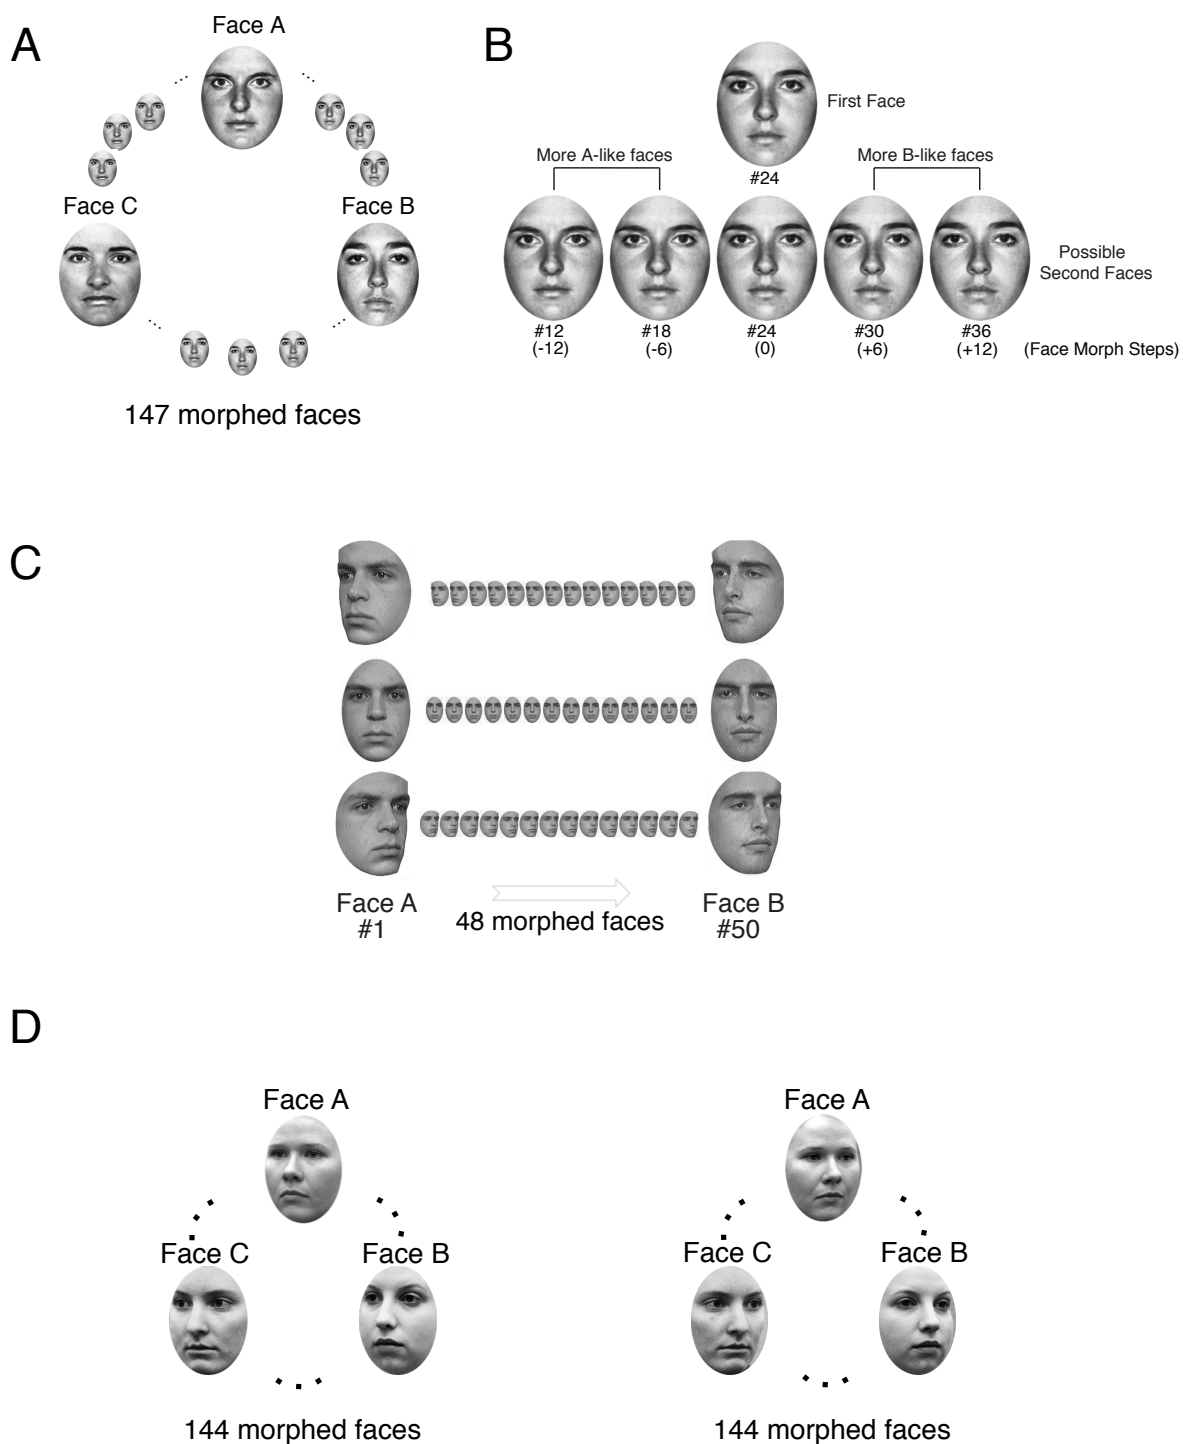
- Kristjánsson, Á., Bjarnason, A., Hjaltason, Á. B., & Stefánsdóttir, B. G. (2009). Priming of luminance-defined motion direction in visual search. *Attention, Perception, & Psychophysics*, *71*(5), 1027–1041. <http://doi.org/10.3758/APP.71.5.1027>
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*. <http://doi.org/10.3389/fpsyg.2013.00863/abstract>
- Leopold, D. A., Rhodes, G., Müller, K.-M., & Jeffery, L. (2005). The dynamics of visual adaptation to faces. *Proceedings of the Royal Society B: Biological Sciences*, *272*(1566), 897–904.
- Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience*, *5*(6), 605–609. <http://doi.org/10.1038/nn851>
- Lieberman, A., Fischer, J., & Whitney, D. (2014). Serial Dependence in the Perception of Faces. *Current Biology*, *24*(21), 2569–2574. <http://doi.org/10.1016/j.cub.2014.09.025>
- Luce, R. D., & Green, D. M. (1974). Detection, discrimination, and recognition. In E. Carterette & M. P. Friedman (Eds.), *Handbook of perception: Vol. 2, Psychophysical judgment and measurement* (Vol. 2, pp. 299–342). ACADEMIC PRESS INC.
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, *22*(6), 657–672.
- Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Attention, Perception, & Psychophysics*, *58*(7), 977–991.
- Maljkovic, V., & Nakayama, K. (2000). Priming of popout: III. A short-term implicit memory system beneficial for rapid target selection. *Visual Cognition*, *7*(5), 571–595.
- Maloney, L., Martello, M. F. D., Sahm, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences*, *102*(8), 3164.
- Martin, F., Baudouin, J.-Y., Tiberghien, G., & Franck, N. (2005). Processing emotional expression and facial identity in schizophrenia. *Psychiatry Research*, *134*(1), 43–53. <http://doi.org/10.1016/j.psychres.2003.12.031>
- Maunsell, J., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*.
- Michotte, A., Thines, G., & Crabbé, G. (1991). Amodal completion of perceptual structures. *Michotte's Experimental Phenomenology of Perception*, 140–167.

- Milders, M., Sahraie, A., & Logan, S. (2008). Minimum presentation time for masked facial expression discrimination. *Cognition & Emotion*, *22*(1), 63–82. <http://doi.org/10.1080/02699930701273849>
- Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, *12*(9), 334–341. <http://doi.org/10.1016/j.tics.2008.05.006>
- Preminger, S., Sagi, D., & Tsodyks, M. (2007). The effects of perceptual history on memory of visual objects. *Vision Research*, *47*(7), 965–973. <http://doi.org/10.1016/j.visres.2007.01.003>
- Rhodes, G., Jeffery, L., Clifford, C. W. G., & Leopold, D. A. (2007). The timecourse of higher-level face aftereffects. *Vision Research*, *47*(17), 2291–2296. <http://doi.org/10.1016/j.visres.2007.05.012>
- Rhodes, G., Jeffery, L., Watson, T. L., Jaquet, E., Winkler, C., & Clifford, C. W. G. (2004). Orientation-Contingent Face Aftereffects and Implications for Face-Coding Mechanisms. *Current Biology*, *14*(23), 2119–2123. <http://doi.org/10.1016/j.cub.2004.11.053>
- Salovey, P., & Mayer, J. D. (1990). Emotional intelligence. *Imagination, Cognition and Personality*, *9*(3), 185–211.
- Schweinberger, S. R., & Soukup, G. R. (1998). Asymmetric relationships among perceptions of facial identity, emotion, and facial speech. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(6), 1748–1765. <http://doi.org/10.1037/0096-1523.24.6.1748>
- Schweinberger, S. R., Burton, A. M., & Kelly, S. W. (1999). Asymmetric dependencies in perceiving identity and emotion: experiments with morphed faces. *Perception & Psychophysics*, *61*(6), 1102–1115.
- Sigurdardottir, H. M., Kristjánsson, Á., & Driver, J. (2008). Repetition streaks increase perceptual sensitivity in visual search of brief displays. *Visual Cognition*, *16*(5), 643–658. <http://doi.org/10.1080/13506280701218364>
- Spelke, E. (1994). Initial knowledge: Six suggestions. *Cognition*, *50*(1), 431–445.
- Spelke, E. S., Kestenbaum, R., Simons, D. J., & Wein, D. (1995). Spatiotemporal continuity, smoothness of motion and object identity in infancy. *British Journal of Developmental Psychology*, *13*(2), 113–142.
- Stel, M., & van Knippenberg, A. (2008). The role of facial mimicry in the recognition of affect. *Psychological Science*.
- Sterzer, P., & Rees, G. (2008). A neural basis for percept stabilization in binocular rivalry. *Journal of Cognitive Neuroscience*, *20*(3), 389–399. <http://doi.org/10.1162/jocn.2008.20039>

- Tanner, T. A., Rauk, J. A., & Atkinson, R. C. (1970). Signal recognition as influenced by information feedback. *Journal of Mathematical Psychology*, 7(2), 259–274.
- Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology : CB*, 19(1), R11–4. <http://doi.org/10.1016/j.cub.2008.10.014>
- Tillman, M. A., & Webster, M. A. (2012). Selectivity of face distortion aftereffects for differences in expression or gender. *Frontiers in Psychology*, 3.
- Tracy, J. L., & Randles, D. (2011). Four Models of Basic Emotions: A Review of Ekman and Cordaro, Izard, Levenson, and Panksepp and Watt. *Emotion Review*, 3(4), 397–405. <http://doi.org/10.1177/1754073911410747>
- Tracy, J. L., & Robins, R. W. (2008). The automaticity of emotion recognition. *Emotion*, 8(1), 81–95. <http://doi.org/10.1037/1528-3542.8.1.81>
- Treisman, A. (1988). Features and objects: The fourteenth bartlett memorial lecture. *The Quarterly Journal of Experimental Psychology Section A*, 40(2), 201–237. <http://doi.org/10.1080/02724988843000104>
- Wallis, G., & Bühlhoff, H. H. (2001). Effects of temporal association on recognition memory. *Proceedings of the National Academy of Sciences*, 98(8), 4800–4804.
- Webster, M. (2012). Evolving concepts of sensory adaptation. *F1000 Biology Reports*, 4. <http://doi.org/10.3410/B4-21>
- Webster, M. A. (2011). Adaptation and visual coding. *Journal of Vision*, 11(5). <http://doi.org/10.1167/11.5.3>
- Webster, M. A., & MacLeod, D. I. A. (2011). Visual adaptation and face perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1571), 1702–1725. <http://doi.org/10.1098/rsbl.2006.0440>
- Webster, M. A., & Maclin, O. H. (1999). Figural aftereffects in the perception of faces. *Psychonomic Bulletin & Review*, 6(4), 647–653.
- Webster, M. A., & Mollon, J. D. (1991). Changes in colour appearance following post-receptoral adaptation. *Nature*.
- Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, 428(6982), 557–561.
- Wiegersma, S. (1982a). A control theory of sequential response production. *Psychological Research*, 44(2), 175–188. <http://doi.org/10.1007/BF00308449>

- Wiegersma, S. (1982b). Sequential response bias in randomized response sequences: A computer simulation. *Acta Psychologica*, *52*(3), 249–256.
- Wild-Wall, N., Dimigen, O., & Sommer, W. (2008). Interaction of facial expressions and familiarity: ERP evidence. *Biological Psychology*, *77*(2), 138–149.
<http://doi.org/10.1016/j.biopsycho.2007.10.001>
- Williams, D., Phillips, G., & Sekuler, R. (1986). Hysteresis in the perception of motion direction as evidence for neural cooperativity. *Nature*, *324*(6094), 253–255.
<http://doi.org/10.1038/324253a0>

Appendix A: Supplemental Figures for Chapter 2



2

Figure S1. (Related to Figure 2.1, 2.2, 2.3, and 2.4) (A) Face morphs used in Experiments 1 and 2. The stimuli consisted of grayscale image morphs based on 3 original female Ekman faces (Ekman &

Friesen, 1976a) with neutral expressions, cropped by an oval aperture to remove the hairline. A set of 48 morphs was created between these identities, resulting in a face morph continuum of 147 faces. **(B)** Experiment 2 trial structure. The faces used in this experiment were a subset of those in panel (A), including original face A (#1), original face B (#50), and the 48 face morphs in between. The first face presented in each trial sequence was drawn from a subset of 26 faces taken from the center of the morph continuum and could range from face morph #13 to face #38. The second face in the sequence could differ from the first face by ± 12 , ± 6 , or 0 face morph steps. Trials that fell in bins -12 and -6 had a first face that was more B-like relative to the second face, trials that fell in the 0 bin had identical first and second faces, and trials that fell in the +6 and +12 bin had a first face that was more A-like relative to the second face. **(C)** Face morphs used in Experiment 3. We used grayscale image morphs based on 2 original neutral male faces across three different viewpoints (frontal, left, right), cropped by an oval to remove the hairline. A set of 48 morphs was created within each of the viewpoints, resulting in three sets of 50 face morphs. **(D)** Experiment 4 stimuli. Grayscale image morphs were created from 3 original neutral female faces across two viewpoints (left and right), cropped by an oval to remove hairline. The identities shown here are similar to those actually used in the experiment, with permission obtained for reprint purposes. A set of 47 morphs was created within each of the viewpoints, resulting in two continuous sets of 144 face morphs.

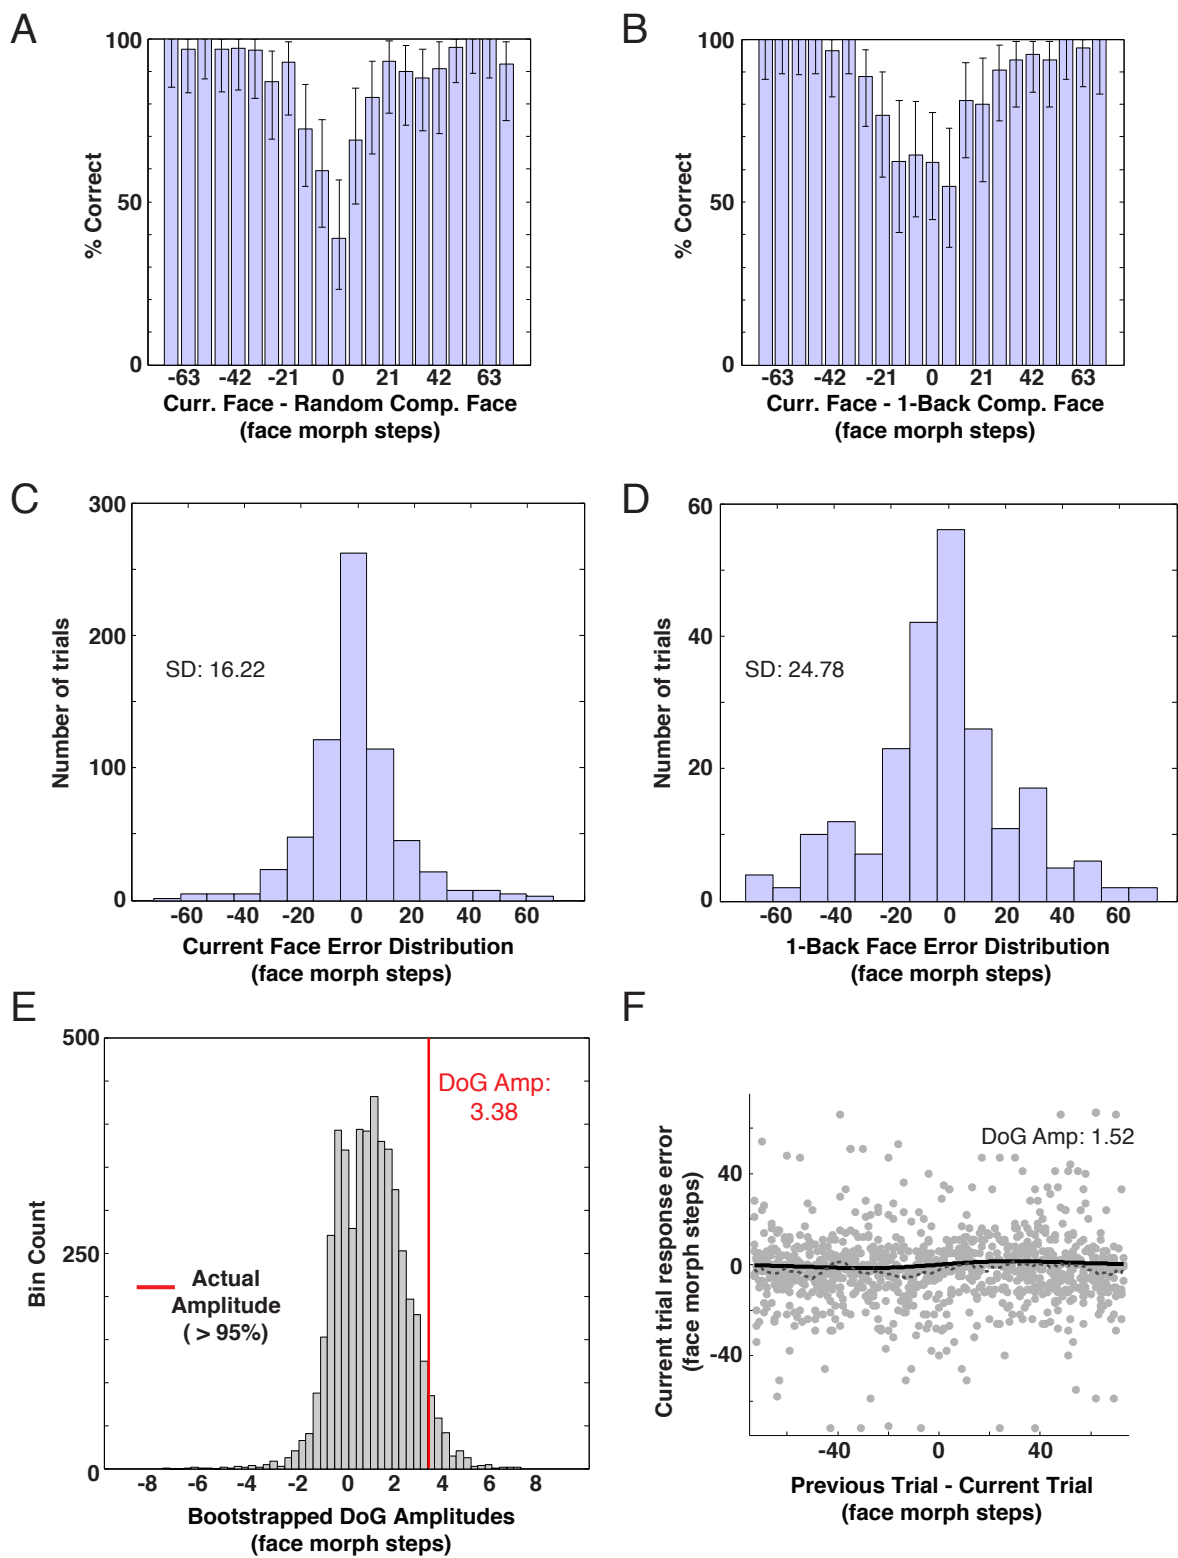


Figure S2. (Related to Figure 2.1). (A-B) 2AFC experiment testing for memory confusion. It is possible that memory confusion, mistakenly reporting the 1-back face rather than the current face, could

have contributed to the pattern of results in Experiment 1 (e.g., Figure 2.1). To determine whether subjects might have experienced memory confusion, and whether the probability of confusing the target and 1-back face changed with increasing similarity between those faces, we ran a control experiment that measured how often subjects mistakenly reported the 1-back target face rather than the current target face. The stimulus set, trial sequence, and timing were similar to that of Experiment 1b: on each trial, subjects saw a random face for 750 ms, followed by a 1000 ms noise mask and 250 ms fixation cross. Subjects then saw a blank screen for 2000 ms, followed by a screen displaying two faces—a target and a lure—one was randomly assigned to the left and the other to the right of fixation. Subjects indicated which face they had last seen (L or R) using the keyboard arrow keys. One of the comparison faces was always the current face (the target, a correct response), and the second comparison face (the lure) was either a random face (50% of trials) or the 1-back face (50% of trials; picking this 1-back face would constitute a memory confusion). The random lure faces served to establish a baseline for how often subjects made discrimination errors or made lapses. A subject may have mistakenly picked the 1-back lure face on a given trial (rather than the current target face) due to a discrimination error rather than a memory confusion; thus, it is important to have a baseline measure of discrimination error. We tested three subjects in this 2AFC experiment (all of whom had also participated in Experiment 1), and collapsed their data for subsequent analyses. Panel A shows percent correct (with 95% binomial confidence intervals) for trials with a random lure face, binned by the morph step difference between the target and lure faces (discrimination errors and lapses). Panel B shows percent correct for trials with a 1-back lure face, binned by the morph step difference between the target and lure faces (memory confusions). There was no significant difference in the distribution of errors obtained for trials with a random lure versus trials with a 1-back lure (panels A versus B; $P = 0.25$, permutation test), indicating that memory confusions per se were not common and that most errors were attributable to discrimination errors or lapses. (C-D) The precision of memory for current and 1-back faces. To measure the precision of subjects' memory, we ran the same three subjects from the 2AFC experiment above (and Experiment 1) in a modified version of Experiment 1a. The procedure was identical to that of Experiment 1a, except that on a random and unpredictable 25% of trials, subjects were asked to match the adjustment face to the 1-back target face rather than the current target face. The histogram in panel C shows the error distribution (collapsed across subjects) for judgments of the current target face (75% of all responses), and the histogram in panel D shows the collapsed error distribution for judgments of the 1-back target face (25% of all trials). The standard deviation of the 1-back face error distribution is significantly larger than the standard deviation of the current face error distribution, indicating that subjects were less precise in their overall recollection of the 1-back faces ($P < 0.001$, permuted null distribution and Mann-Whitney U test). One might argue that, since subjects had to report the 1-back face on only 25% of the trials, the task became more difficult and resulted in broader error distributions for reporting the 1-back face. However, it is critical to note that in Experiment 1 and 4 (Figure 2.1 and 2.4), subjects were not required to recall or report the 1-back face at all. Thus, by having subjects report the 1-back face on 25% of trials, the 1-back error distribution we measure here is conservatively narrow—likely more precise than the error distribution in Experiments 1 and 4 for the 1-back face. (E) Memory confusion model. Although there was no significant bias to picking the 1-back face over a random face (no overall difference between panels A and B), we examined whether memory confusion could contribute to the pattern of results in Figure 1. For any bin in which subjects were more likely to choose a 1-back lure than a random lure (higher error rate in panel B than panel A), we attributed the excess errors to memory confusions (a lenient criterion for what counts as memory confusion), and used those values to constrain the memory model. Based on the difference between panels A and B, a total of 2.9% of responses could potentially be classified as memory confusions. To directly model the potential influence of memory confusions, we ran 5000 bootstrapped iterations where we randomly chose 1000 target faces per iteration and simulated method of adjustment responses to each target face based on the empirical error distribution in panel C. For the same 1000 random target face trials, we also simulated responses to each 1-back target face based on the 1-back error distribution in panel D (simulating responses subjects would

give if they had experienced memory confusions). We then used this model to generate response data in which subjects experienced memory confusions on a portion of the trials. To do this, we took the 1000 simulated current target face responses and replaced a percentage of those with 1-back target face responses, based on the frequency of memory confusions per bin in the 2AFC experiment. That is, we used the memory confusion error rate to choose a corresponding proportion of randomly sampled trials from the 1-back error distribution. We then treated the resulting simulated data just as we had treated the empirical data from Figure 1, fitting a derivative of a Gaussian (DoG) curve to the combined simulation data and estimating the amplitude of the DoG curve. We repeated this fitting procedure for all iterations of the memory confusion model, with the amplitude estimate from each iteration reflecting the apparent serial dependence that might arise as a result of memory confusions. The histogram in panel E shows DoG amplitudes for 5000 simulation iterations. Without constraining the width of the DoG fits to that of the empirical fits from Experiment 1, 95% of the simulated DoG amplitudes were smaller than the empirical amplitude (red line). Constraining the width of the DoG simulations to within +/- 5 face morph steps of peak serial dependence (testing whether memory confusions could produce serial dependence of the same width and amplitude as we observed in Experiment 1), over 99% of the simulated amplitudes were lower than the empirical amplitude. Finally, even when an unrealistically large proportion of errors are classified as memory confusions (50% of the errors in panel B), over 97% of the simulated amplitudes were smaller than the empirical amplitude. Thus, although there could be a small contribution of memory confusion to our results, the constellation of results here shows that memory confusion cannot completely account for the serial dependence of face perception. (F) One example memory confusion simulation. One example simulated data set, with a DoG amplitude of ~1.5 - significantly less than the empirical amplitude. This simulated amplitude is one of the 5000 iterations in panel E.

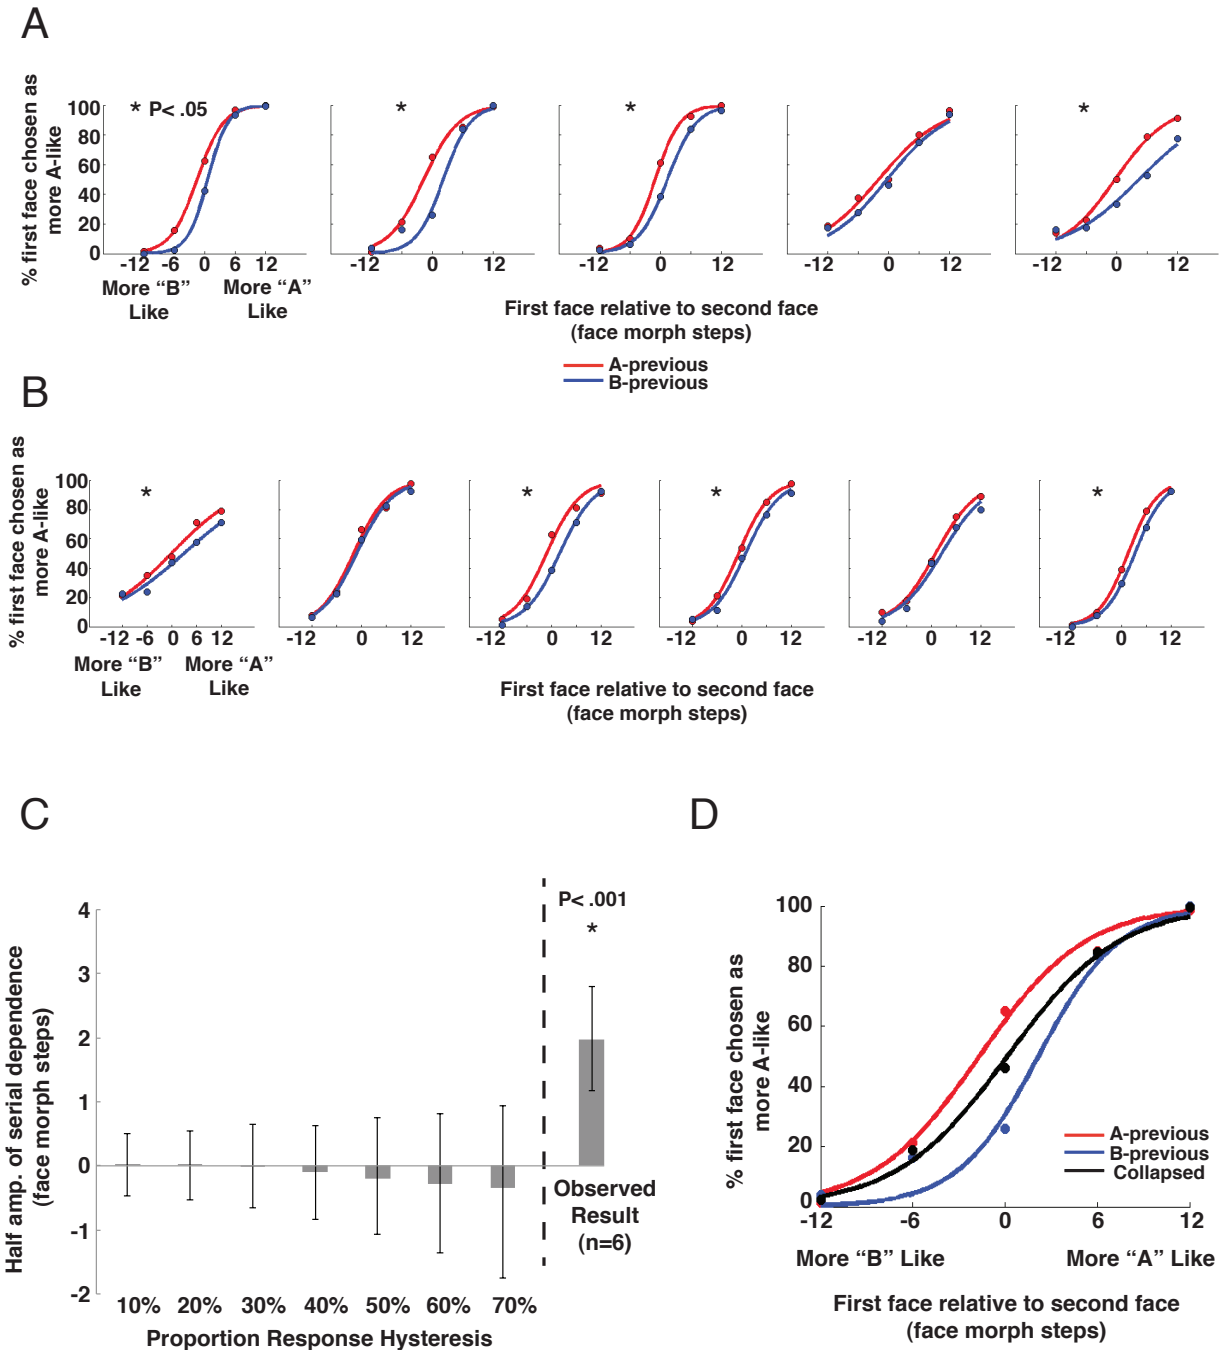


Figure S3. (Related to Figure 2.2 and 2.3) (A) Data from each subject for Experiment 2. Abscissa shows the identity of the first face relative to the second face. Trials that fell in bins -12 and -6 on the x-axis had a first face that was more B-like relative to the second face, trials that fell in the 0 bin had identical first and second faces, and trials that fell in the +6 and +12 bins had a first face that was more A-like relative to the second face. The ordinate shows the proportion of first faces that were chosen as being more A-like. The red data consists of all trials with 1-back first faces that were more "A"-like and the blue data consists of all trials with 1-back first faces that were more "B"-like. Asterisk indicates significance at the 0.05 level; P-values are based on each subject's permuted null distribution. (B) Data for each subject from Experiment 3. The format of the graphs is the same as in panel (A). Experiment 3

was identical to Experiment 2, except sequential trials always contained a different viewpoint. The first and second face within a trial were always viewed from the same angle. Four of six individual subjects showed a significant shift in their psychometric functions based on the identity shown in the previous trial, with the remaining 2 subjects showing a trend in the same direction. **(C)** Response hysteresis simulation for Experiment 3. The 2IFC design of Experiments 2 and 3 was designed to disentangle perceptual serial dependence from hysteresis in subjects' responses. Nonetheless, we conducted a simulation to test whether repetition of responses could have produced the serial dependence we observed in Experiments 2 and 3. For each subject's trial sequence presented in Experiment 3, we simulated responses with various proportions of response hysteresis. For example, to generate 10% response hysteresis, 90% of trials (randomly chosen) had correct responses and the remaining 10% of trials repeated the 1-back response. For correct trials where subjects saw the same first and second face, we assigned a response at random. We then computed the amplitude of serial dependence within the simulated response sequence, and repeated this analysis 1000 times in order to generate bootstrapped confidence intervals at each proportion of hysteresis. At 80% and greater response hysteresis, the simulated responses became too noisy to reliably fit with psychometric functions. No amount of response hysteresis produced significant serial dependence, indicating that response hysteresis could not be responsible for the perceptual serial dependence we observed in Experiments 2 and 3. **(D)** Perceptual serial dependence reduces sensitivity. Here we show example data from subject 2 in Experiment 2 with A-previous and B-previous psychometric functions collapsed (black curve). By definition, when the psychometric functions in Experiments 2 and 3 that are separated by trial type (e.g., blue and red functions) are collapsed, they yield a shallower single function compared to the average slope of the two logistic functions separately. That is, subjects' sensitivity in Experiments 2 and 3 was higher when computed separately for trials with different 1-back face identities versus when computed on collapsed data from all trial types ($p < 0.001$, permuted group null, $n=12$). This consistent shift in PSE when separating responses based on trial-type indicates that the reduced sensitivity is an important perceptual consequence of serial dependence. While perceptual priming leads to an improvement in sensitivity and performance, our results actually show the opposite effect: the perceived identity of a face was misperceived as being more similar to a previous face, which actually reduced overall sensitivity to stimulus difference.