COGNITIVE NEUROSCIENCE

Illusion of visual stability through active perceptual serial dependence

Mauro Manassi[1,2]* and David Whitney[1,3,4]

Despite a noisy and ever-changing visual world, our perceptual experience seems remarkably stable over time. How does our visual system achieve this apparent stability? Here, we introduce a previously unknown visual illusion that shows direct evidence for an online mechanism continuously smoothing our percepts over time. As a result, a continuously seen physically changing object can be misperceived as unchanging. We find that online object appearance is captured by past visual experience up to 15 seconds ago. We propose that, because of an underlying active mechanism of serial dependence, the representation of the object is continuously merged over time, and the consequence is an illusory stability in which object appearance is biased toward the past. Our results provide a direct demonstration of the link between serial dependence in visual representations and perceived visual stability in everyday life.

INTRODUCTION

Why do objects in the world appear to be so stable despite constant changes in their retinal images? Retinal images continuously fluctuate because of many sources of internal and external noise ranging from retinal image motion, occlusions and discontinuities, lighting changes, and perspective changes, among many other sources of noise. However, the objects do not appear to jitter, fluctuate, or change identity from moment to moment. This problem—why the world seems unchanging over time—is decades, if not centuries, old (1, 2).

The most common modern explanations for the appearance of object stability revolve around some form of change blindness (3–5) or inattentional blindness (6, 7), in which the capacity limits of visual short-term memory (8, 9) prevent us from being aware of things that change. This class of explanation revolves around the limits of perceptual, decisional, or memory processing: Any fluctuations in the retinal images of objects do not cause the objects to appear to change identity because we simply do not notice those changes. None of these explanations of perceptual stability make any predictions about the online appearance of objects at any particular moment in time. According to these accounts, one may not recall whether an object at this moment has changed identity, but the object at this moment is not misperceived as being some other thing.

While the limits of perceptual and cognitive processing certainly define the envelope of our awareness, there is an alternative but complementary explanation for why objects in the world appear stable. The hypothesis is that the visual system uses an active stabilization mechanism, a pull toward the past, which manifests as serial dependence in perceptual judgments. Serial dependence causes objects at any moment to be misperceived as being more similar to those in the recent past. This serial dependence has been reported in the appearance of things (10–13), perceptual decisions about things (14–16), and the memory for things (17–19). In all of these examples, serial dependence is found for random or unpredictable sequential images.

If serial dependence serves to stabilize perception, then it should generate an online illusion of stability in which a single physically changing object can be misperceived as unchanging. For example, an object that smoothly changes identity could appear "frozen" and unchanging; change blindness in this case would arise from continuously biasing our representations toward the past. The key difference between the predictions from an active serial dependence and that of the passive change blindness–based explanations (3–5) is that serial dependence predicts that the online perceptual appearance of attended things can be altered and incorrect. Change blindness explanations make no predictions about an ongoing misrecognition of attended objects at each moment in time. Here, we show that the online appearance of a single physically changing object can be made to seem stable through an active mechanism of perceptual serial dependence.

RESULTS

In experiment 1, two separate groups of 44 and 45 participants rated on Mechanical Turk the age of a young or old static face embedded in a blue frame (13 and 25.5 years, respectively; Fig. 1A, white square and circle). A third group of 47 independent participants were presented with a movie of a face that morphed gradually, aging from young to old. These observers then rated the age of the old face embedded in a blue frame (20.2 years; Fig. 1A, blue circle). For simplicity, we will refer to the rating of the static face alone as the "reference face" (white data points) and to ratings of the static face preceded by the movie as the "test face" (blue circle). Participants were made aware of the age rating task only after the task-irrelevant movie, when the static face and blue frame appeared; they were not warned, cued, or asked to make any explicit decision about the movie while watching it. In addition, each participant completed one, and only one, single trial across all our experiments. Thus, every experiment consisted of one-shot independent trials, from independent subjects who had no prior knowledge or experience with the stimulus or task.

We compared the age ratings between physically identical static faces, either alone (reference face) or with a preceding video (test face). The last frame of the video was identical to the test/reference face. Although the two faces (test and reference) were identical, the

[1]Department of Psychology, University of California, Berkeley, CA, USA. [2]School of Psychology, University of Aberdeen, King's College, Aberdeen, UK. [3]Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA. [4]Vision Science Group, University of California, Berkeley, CA, USA.
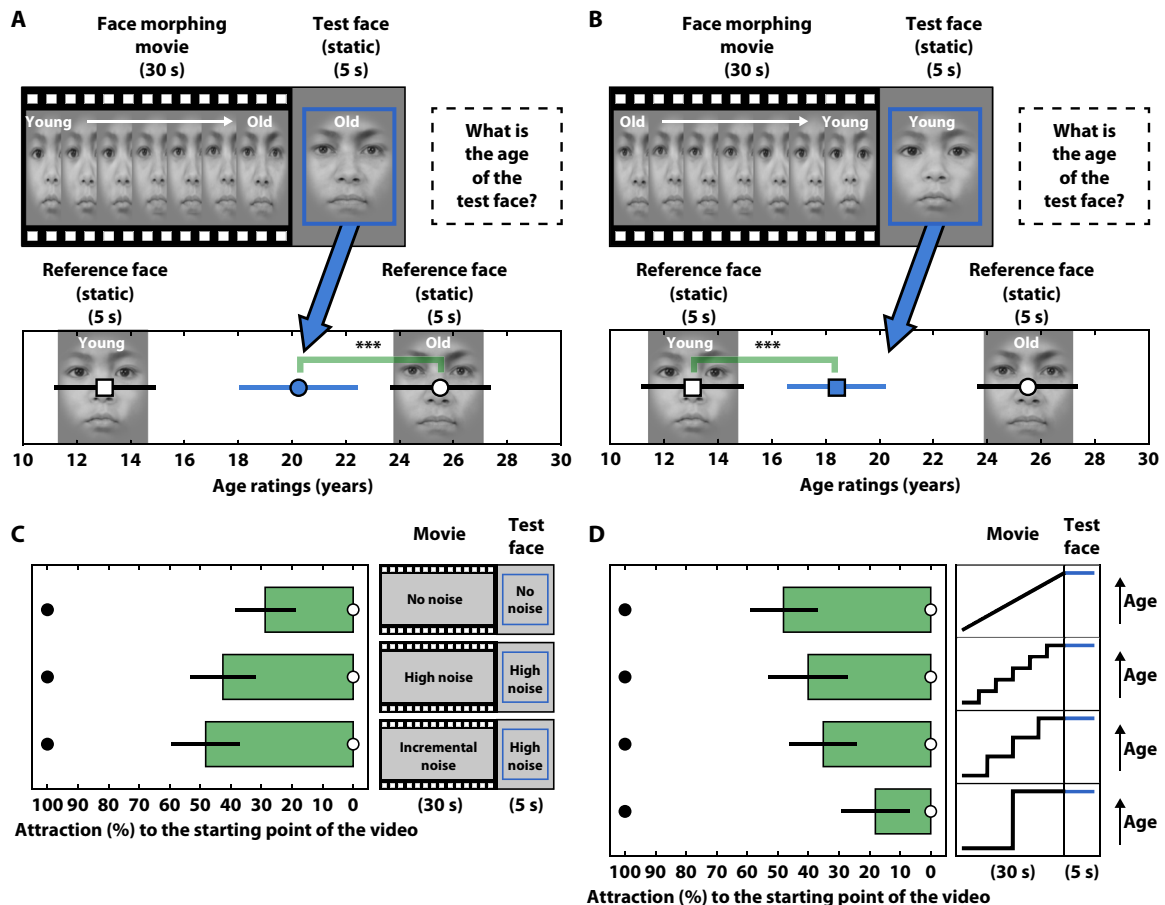*Corresponding author. Email: mauro.manassi@abdn.ac.uk

Fig. 1. Experiments 1 to 3. (**A**) Experiment 1. Two groups of observers were asked to rate the age of a young or old static face embedded in a blue frame (white square and circle; reference faces). A third group was presented with a face morphing movie gradually aging from young to old and was then asked to rate the age of the old face embedded in a blue frame (blue circle, test face). Although the test face and reference faces are identical, the old test face was rated as much younger than what it actually was (green bracket). White text is shown for illustration. (**B**) A fourth group was presented with a face morphing movie gradually rejuvenating from old to young and was then asked to rate the age of the young face embedded in a blue frame (blue square; test face). The young test face was rated as much older than what it actually was (green bracket). ***P < 0.0001. (**C**) Experiment 2. Attraction percentage was computed as age difference between reference faces (e.g., reference face: old) and test faces (e.g., movie: young to old; test face: old) divided by the total age range (e.g., old reference face − young reference face). Increasing (A) and decreasing (B) age directions were equally balanced. White and black circles indicate zero (0%) and full (100%) attraction toward the beginning of the movie. When both movie and test face were presented without noise, test face age ratings were attracted toward 28% of the movie. When both movie and test face were presented with high constant dynamic noise, attraction was around 42% (A and B). When the movie was presented with high dynamic incremental noise and the test face with high noise, attraction was around 48%. (**D**) Experiment 3. When the age in the face morphing movie increased in gradual steps of 6, 4, and 2, attraction gradually decreased. Incremental noise and high noise were added to the movie and static face (reference or test faces), respectively. Error bars are bootstrapped 95% confidence intervals. Photo credit: Anthony Cerniello. Computer-generated face images were slightly modified for visualization purposes.

old test face, seen after the video, was rated as 5 years younger than the old reference face, seen without the video (20.2 versus 25.5 years; *P* < 0.001; Fig. 1A, green bracket). We propose that, because of serial dependence, the identity of the face is continuously merged over time through the movie, and hence, observers perceive a slower age change (movie S1). As a result, the static test face is misperceived as biased toward the content of the movie seen 12 to 15 seconds ago. This is an online illusion of stability, and it shows that a single physically changing object can be misperceived as unchanging. This serial dependence effect occurs on a perceptual level, as observers were asked to judge the age of a static test face without any prior instructions, no prior decision about the movie, and no memory load.

To test whether the stability illusion is due to a simple unidirectional bias in age ratings, a fourth group of 45 new participants watched a movie of a rejuvenating face that gradually morphed from old to young. Following the movie, observers rated the age of a young static test face embedded in a blue frame (Fig. 1B, blue square). The young face was rated as 5 years older than its actual age (18.4 versus 13 years; *P* < 0.001; Fig. 1B, green bracket). This confirms that the stability illusion can cause faces to appear younger or older depending on the previously seen faces.

Serial dependence in perception is modulated by uncertainty, such that increased noise increases serial dependence (*11*, *20*). To test whether noise plays a role in the illusion of stability here, we presented the stimuli with and without noise to separate groups of observers. As a measure of the stability illusion strength, we calculated the bias in age ratings toward the beginning of the movie as an attraction index. We computed the absolute difference between test

and reference faces (e.g., old reference age – old test face) and divided that by the total age range (e.g., old reference face – young reference face; white circles in Fig. 1, A and B). Age directions were equally balanced (Fig. 1, A and B). When the movie and test face were presented alone (movie S2) or with superimposed dynamic noise, the static test face ratings were attracted by 28 and 42% of the movie, respectively [Fig. 1C; Fig. 1 (A and B) refers to the high noise condition; see Materials and Methods]. When the movie was presented with increasing dynamic noise and test face with high noise,

the attraction was around 48% (experiment 2; Fig. 1C and movie S3). The results are consistent with the finding that serial dependence in perception increases with noise and uncertainty (*11*, *20*). As the latter noise condition yielded the strongest illusory effect, it was used across subsequent experiments (Figs. 2 to 4). Video discriminability or noise cannot be considered as valid explanations for our illusion; our illusion was 1.5 times the Just Noticeable Difference (JND) of the video (minimum video duration for which observers could correctly discriminate the age; see experiment 12).
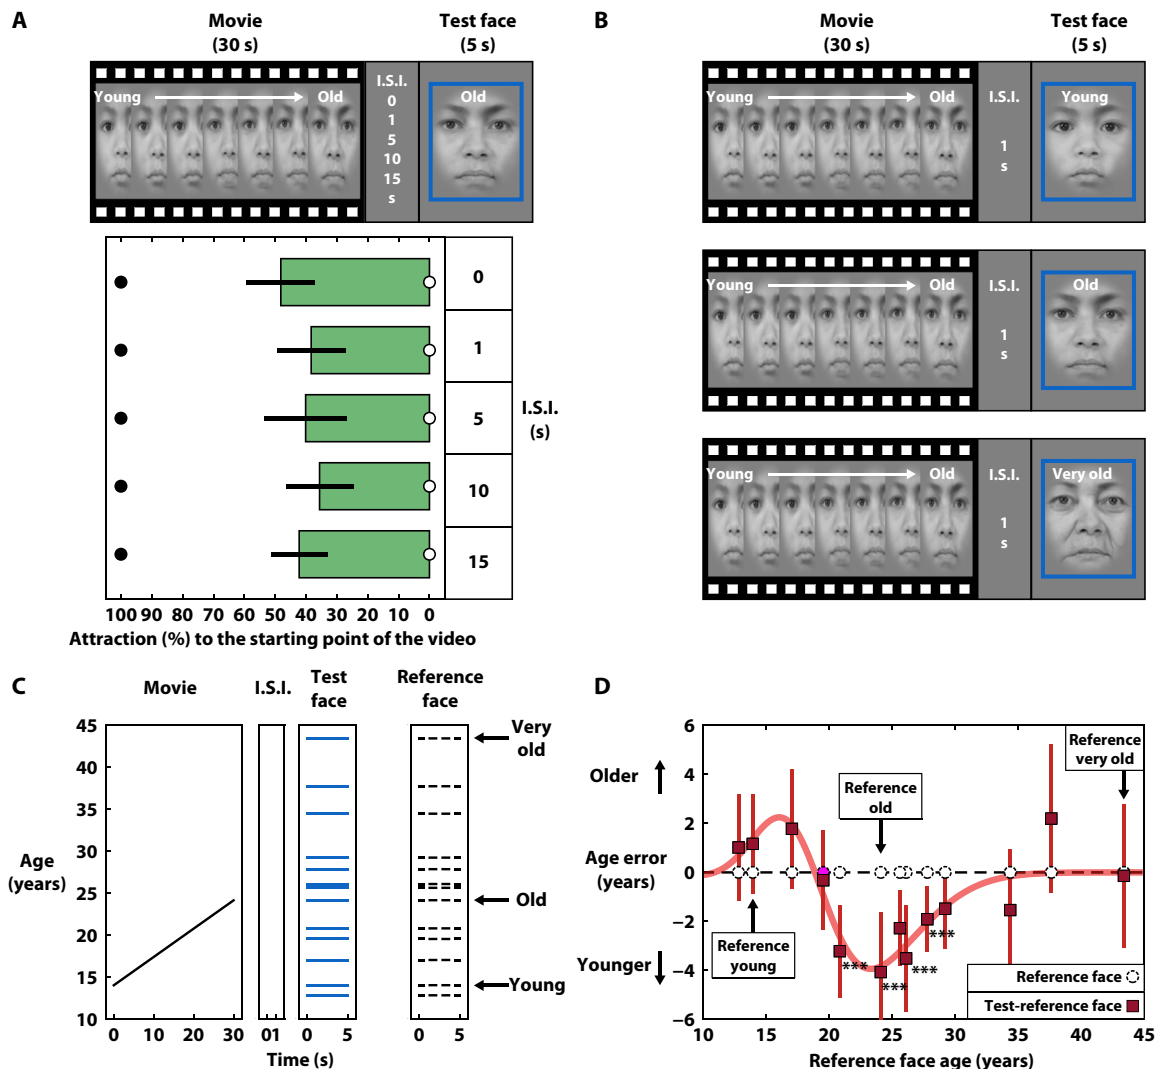

Fig. 2. Experiments 4 to 6. (**A**) Experiment 4. Experimental design was identical to Fig. 1 (A and B), except for an I.S.I. of 0, 1, 5, 10, or 15 s between the movie and static face. Incremental noise and high noise were added to the movie and static face (reference or test faces), respectively. Attraction percentage was computed with equally balanced increasing and decreasing age directions in the face morphing movie (Fig. 1, A and B). Static face age rating was attracted toward the movie at all tested I.S.I. (**B** to **D**) Experiment 5. (B and C) Thirteen groups of observers were presented with a movie with a face gradually aging from young to old (black line) and, after an I.S.I. of 1 s (gray line), were asked to rate the age of the static face (test faces; blue lines). The test face was randomly chosen by random sampling the movie. As a control, the other 13 observers' groups were asked to rate the age of a static face with the same ages (reference faces; dashed lines). (D) Standardization of previous graph in terms of the distance of test face rating from the reference baseline (white dots, dashed line). For each test face age, age error in test faces was computed as the difference between reference and test face (red squares). Negative and positive values indicate that the static face was rated as younger and older than what it actually was, respectively. As in Fig. 1A, when the test face was old, it was rated as much younger than what it actually was (white circles in the center). When the test face was even older (white circles on the right side), attraction gradually decreased. When the static face was younger (white circles on the left side), attraction gradually decreased and flipped to the opposite direction, i.e., a young test face was rated as older than what is actually was. Error bars are bootstrapped 95% confidence intervals. ***P < 0.0001. Photo credit: Anthony Cerniello. Computer-generated face images were slightly modified for visualization purposes.
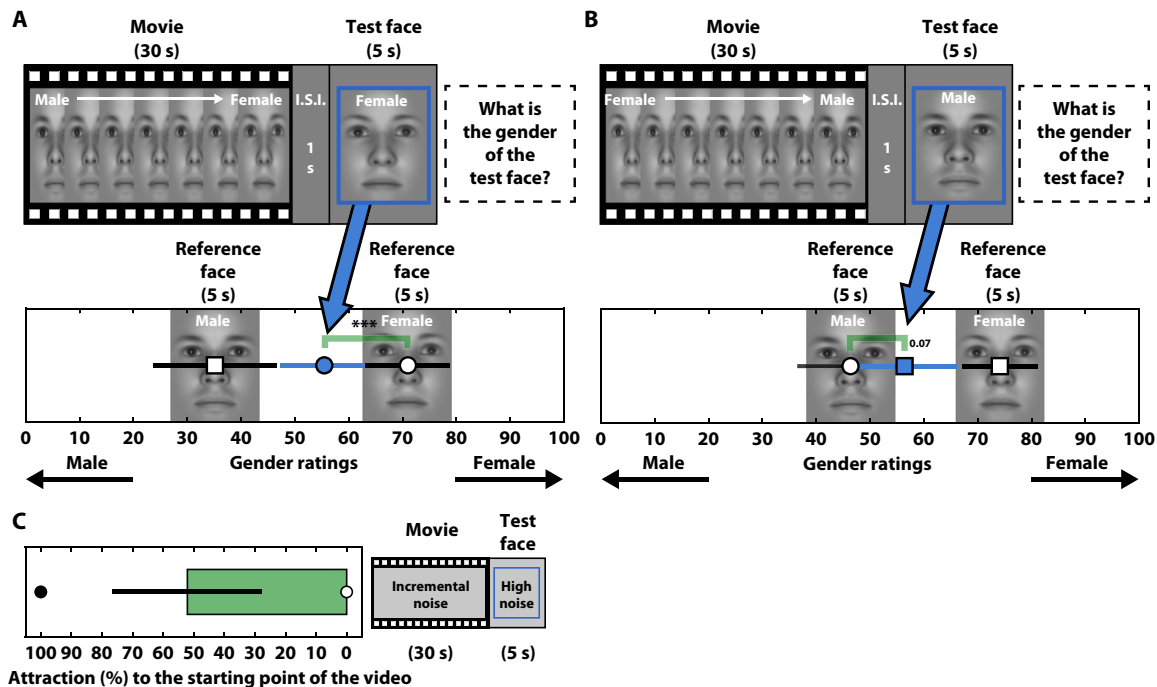
Fig. 3. Experiment 6. (**A**) Two groups of observers were asked to rate the gender of a male or female static reference face embedded in a blue frame (white square and circle; reference faces). A third group was presented with a face morphing movie gradually changing gender from male to female and was then asked to rate the gender of the test female face embedded in a blue frame (blue circle; test face). The female face was rated as much more masculine than what it actually was (green bracket). (**B**) A fourth group was presented with a face morphing movie gradually changing gender from female to male and was then asked to rate the gender of the male face embedded in a blue frame (blue square; test face). The male face was rated as much more feminine than what it actually was (green bracket). Different static faces ratings between (A) and (B) are due to incremental noise [(A) low-high; (B) high-low]. (**C**) Test face gender ratings were attracted 52% of the way toward the starting point of the movie. Photo credit: Mauro Manassi.

It might be argued that our results are due to a central tendency bias, i.e., the tendency to rate test faces as being close to middle age, independent of movie content. To test this hypothesis, experiment 3 replicated the same conditions as in Fig. 1 (A and B), but instead of a linear increase or decrease in the age of the face, the age of the face was morphed in different staircase functions (Fig. 1D). This left intact the starting and ending points of the movies (young and old). Attraction gradually decreased with decreasing the number of age steps in the movie, thus showing that our illusion is not only due to a simple response or central tendency bias but also it strongly depends on the whole content of the face morphing movie (Fig. 1D). As a further proof that integration occurs across the entire movie range, we computed the attraction with the last 6, 18, and 30 seconds of the video preceding the test face. Attraction linearly increased with increasing video duration, thus showing that the attraction effect involves all parts of the preceding video (experiment 11; fig. S3A).

If our illusion is due to an active mechanism of perceptual serial dependence, then it should occur on a broad temporal range in accordance with previous literature (*10*, *12*). Accordingly, in experiment 4, we measured the temporal strength of our illusion with an interstimulus interval (I.S.I.) of 0, 1, 5, 10, and 15 seconds between the movie and test face (movie S4). Test face age ratings were attracted toward the movie at all intervals, thus showing that our illusion extends across a large period of time ($P < 0.001$ in all conditions; Fig. 2A). These results further show that, without intervening trials, serial dependence magnitude extends over a larger period of time than previously shown (*10*).

If consistent with previous serial dependence literature on face stimuli, our illusion should be determined by face feature similarity [face identity: (*21*, *22*); face expression: (*23*, *24*)], and hence, it should occur only when the face morphing movie and test face are similar. Crucially, and unlike previous passive change blindness–based explanations (*3*–*5*), any modulation of the illusion by feature similarity would be consistent with serial dependence and would allow us to make predictions about the apparent age of the test face.

In experiment 5, we presented a movie of a face that morphed from young to old (as in Fig. 1A), and after an interval of 1 second, we varied the age of the static test face by making it younger or older than the original test old face (Fig. 2, B and C, and movie S5). On the basis of the known tuning of serial dependence for face similarity (*23*), we formulated three predictions. First, our illusion should occur only with faces similar in age (and thus identity) to the test face and not between dissimilar faces. We found that the old test face was rated as younger (attraction effect) only for a few similar identities that were most similar to the old face; the attraction disappeared for more dissimilar identities ($P < 0.01$; Fig. 2D, red squares with asterisks). Second, as the "old" test face was perceived as being ~20 years old after watching the movie (Fig. 1A), we predict that, when a reference face that is 20 years old is used as a test face after the movie, the degree of attraction for that face should be zero. We found no attraction for a test face of 20 years of age, meaning that the 20-year-old face is the age actually perceived at the end of the movie (Fig. 2D, purple dot). Third, test faces younger than ~20 years old should be perceived as older, because the movie content contains
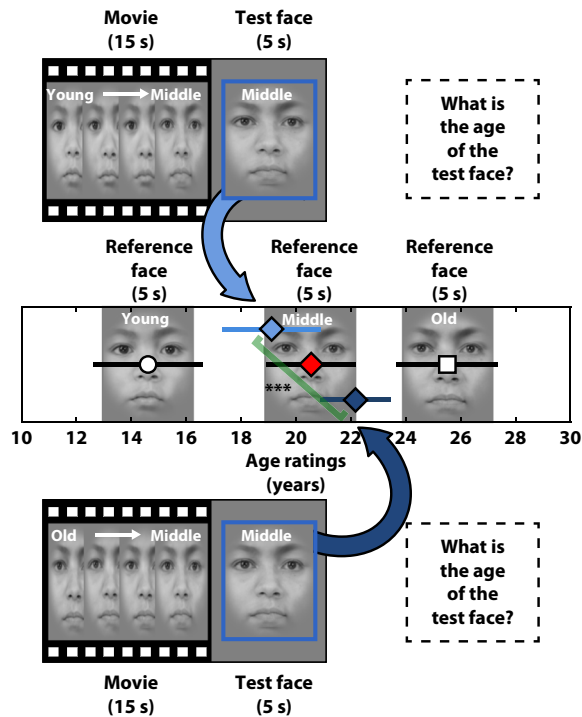
Fig. 4. Experiment 7. Three groups of observers were asked to rate the age of a young, middle, or old face embedded in a blue frame (white circle, red diamond, and white square; reference faces). A fourth group was presented a movie with a face gradually aging from young to middle age and were then asked to rate the age of the test middle face embedded in a blue frame (light blue diamond; test face). The middle test face was rated as younger than it actually was. A fifth group was presented a movie with a face gradually rejuvenating from old to middle age and were then asked to rate the age of the test middle face embedded in a blue frame (dark blue diamond; test face). The middle test face was rated as older than it actually was. ***$P < 0.0001$. Photo credit: Anthony Cerniello. Computer-generated face images were slightly modified for visualization purposes.

older identities across the duration of the morph movie and, hence, should bias test face perception toward older ages. When the test face was younger, it was rated as older than it actually was (Fig. 2D, white circles on the left side). Our results and predictions were very well captured by a two-parameter derivative of Gaussian model, in accordance with previous results (10, 18, 19, 25), and ideal observer models proposed in the serial dependence literature (20).

Our illusion is not restricted to the specific identity we tested. We replicated the effect with a different set of identities (see fig. S1). As a further confirmation of the featural tuning of our illusion, we found that attraction toward the past was higher when the movie and static test face had the same identity compared to when they had different ones (see fig. S1D).

It might be argued that the illusion of stability here is restricted to judgments of age. Conversely, if serial dependence induces apparent stability more generally, then the effect should hold true across different stimuli and tasks. To address this, in experiment 6, we tested whether our stability effect generalizes to a new set of stimuli and a gender rating task (Fig. 3 and movie S6). Four separate groups of ~50 participants (see Materials and Method for more details) rated on Mechanical Turk the gender (male/female) of static reference faces with low and high dynamic noise, embedded in a blue frame. Ratings were on a scale from 0 to 100% (percent female). Male faces were rated as 46 to 35% (low and high noise, respectively), and female faces were rated as 74 to 71% (low and high noise, respectively; Fig. 3, A and B, white squares and circles; reference faces). A fifth group of 53 participants viewed a face morphing movie that gradually changed gender from male to female. After this, they were asked to rate the gender of the female test face embedded in a blue frame (Fig. 3A, blue circle; test face). A sixth group of 50 participants watched a face morphing movie that gradually changed gender from female to male and then rated the gender of the male test face embedded in a blue frame (Fig. 3B, blue square; test face). The static test face was pulled toward the previously viewed movie and rated as more masculine (55.4% versus 71%; $P = 0.002$; Fig. 4A, green brackets) or feminine (56% versus 46%; $P = 0.07$; Fig. 4B; green brackets) than it actually was. Collapsed across both conditions, static test face ratings were biased more than halfway toward the starting point of the movie ($P < 0.001$; Fig. 3C). Hence, our effect is not limited to age judgments or particular stimuli.

A prediction of the serial dependence–induced illusion of stability is that a face should not only appear stable and biased toward the past but also the same face should look different depending on one's previous perceptual experience. To test this hypothesis, in a baseline condition of experiment 7, three separate groups of observers rated the age of a young, middle, or old reference face, respectively (young: 14.6 years; middle: 20.5 years; old: 25.5; Fig. 4, white and red shapes). In the first condition, we presented a movie of a face morphing from young to middle age, followed by a middle-aged test face and age rating task (Fig. 4, light blue diamond). In the second condition, we presented a movie of a face morphing from old to middle age, followed by a middle-aged test face and age rating task (Fig. 4, dark blue diamond). Note that the middle-aged test face in both conditions was identical. Whereas the reference middle-aged static face was rated as 20.5 years old (red diamond), it was rated as younger (19.1 years; light blue diamond) or older (22.1 years; dark blue diamond) depending on the content of the previous movie (blue diamonds; $P < 0.001$; movie S7). We replicated the same result using the gender morph stimuli from Fig. 3 (fig. S2). This experiment rules out any explanations based on a central tendency bias in responses or perception (26, 27), because the middle face was pulled away from the center and toward either end of the morph continuum, depending entirely on the direction of the previously viewed movie.

DISCUSSION

Our results reveal a novel perceptual illusion: A continuous physically changing object is misperceived as unchanging because of a perceptual pull toward our past visual experience (Fig. 1). Our stability illusion cannot be explained by passive change blindness (3–5); the underlying mechanism of serial dependence actively modifies the perceptual appearance of objects, making them appear as they were several seconds ago (Figs. 2 to 4).

Visual perception was previously proposed to be determined by continuity fields: spatiotemporal integration mechanisms that continuously bias object perception toward our past visual experience (10, 12, 28). The mechanism(s) of continuity fields serve two main purposes. First, they reduce the number of potential neural computations across time for each perceived object, by recycling previously perceived features and objects. In this sense, misperception is an efficient strategy because the visual system exploits to its advantage natural temporal redundancies (11, 20). Equally important, continuity

fields promote a continuous and stable representation of the world. Continuity fields could help explain the perception of stability in many visual illusions (29–32), including the sequential dependencies in perception, decision, and memory (10, 17). However, all serial dependence literature has so far focused on random or unpredictable sequential images, mostly using unnatural successions of static images with a trial-by-trial analysis. Hence, the actual link between serial dependence and visual stability has yet to be demonstrated. In this regard, our illusion represents the first direct demonstration that inducing online serial dependence can cause apparent perceptual stability in an otherwise continuously changing object.

Our results cannot be explained by a unidirectional bias in age ratings (Fig. 1, A versus B, and Fig. 3, A versus B) or a strictly age-related illusion (Figs. 3 and 4; same illusion with gender judgments). Furthermore, explanations based on central tendency (27), cognitive anchoring (33), or summary statistics representation computed in a given temporal window (34–37) are not consistent with our results (Figs. 1D and 2D and fig. S1D) (34, 38); the illusion does not simply depend on the starting point of the video, as the gradual change of the entire video is necessary (experiments 3 and 11), as well as similarity between video and test face (experiments 5 and 9). In addition, note that explanations based on negative aftereffects resulting from face adaptation (39, 40) would predict opposite results to the ones we reported: Subsequent faces would be repelled away instead of attracted. Last, our illusion cannot be explained by poor sensitivity to change in the video, as the JND for video discrimination is lower than the illusion itself (experiment 12); thus, the impression of an unchanging face, and the bias toward previously seen faces, cannot be driven by poor discriminability alone.

There is currently an intense debate on the underlying mechanisms of serial dependence and the neural origin of continuity fields: Among others, serial dependence was proposed to occur at the level of perception (10, 11, 20, 41, 42), decision (14, 15, 25), and memory (17, 43). To address this issue, we tested our illusion with a single-trial design, in which participants only experience one trial and, therefore, (i) make no prior decision about the movie, (ii) have no prior instructions about the aging/gender rating task until they are presented with the test face, and, thus, (iii) have no memory load across the experiment. Consequently, we propose that our illusion occurs on a truly perceptual level, as the reader can directly experience in movie S8 and at www.mauromanassi.com/stability-illusion. It is important to mention that serial dependence on a perceptual level does not contradict or deny the existence of serial dependence on other levels, such as decision or memory (14, 16, 17, 19).

This illusion demonstrates an active mechanism of perception stabilization in real time. The perceptual outcome of this active mechanism is blindness to temporal changes in the environment by actively smoothing the appearance of objects over a relatively long time scale. Crucially, whereas standard change blindness accounts assume a passive loss of information because of capacity constraints, we show evidence for a complementary active mechanism that achieves the same ends (change blindness) with a purposeful computational goal of stabilizing the appearance of the world.

MATERIALS AND METHODS
Online recruitment and selection
All experimental procedures were approved by and conducted in accordance with the guidelines and regulations of the University of California (UC) Berkeley Institutional Review Board (IRB). Participants were recruited online on Mechanical Turk and provided informed consent in accordance with the IRB guidelines of UC Berkeley.

To avoid any cue to the task, the title of the Mechanical Turk survey was "Very very very short survey." Each observer was paid 20 cents for participating in the study. The experiment lasted ~1 min (35-s reference face, 65-s movie and test face), and I.S.I. (0 to 25 s) and response time. The maximum time allotted for completing the experiment was 10 min. Turk workers' qualifications were defined as (i) only within the United States, (ii) with more than 50 hits already approved, and (iii) not having already completed the survey. To make sure each participant completed only one survey, we (i) selected "having already completed the experiment" as a qualification, (ii) used TurkGate (44) to exclude participants and avoid survey previews, and (iii) removed participants with the same IP (Internet Protocol) address across all the experiment datasets, considering only the first participation as valid. From Mechanical Turk, participants were redirected to a Qualtrics study through a weblink. Experiments 11 and 12 were run through PsychoPy with the same design (45), and experiment 12 was run locally.

Stimuli
Face stimuli for videos in experiments 1 to 5, 7, and 9 were extracted and modified from a clip by A. Cerniello (https://vimeo.com/74033442). Face stimuli for videos in experiments 6 and 10 (gender stimuli) and experiments 8 and 9 (aging stimuli) were made in Daz 3D. Face morphing movies were created by morphing between pairs of faces (male/female and young/old). Each face was cropped within an oval Gaussian mask and shown in the center of the screen with a gray background. Test and reference faces lasted 5 s, and the face morphing movie lasted 30 s.

In the noise conditions (Fig. 1C), a dynamic pink rectangular noise mask was superimposed on the face with different transparency levels. In the identity in Figs. 1, 2, and 4, transparency in the low and high noise condition was 25 and 55%, respectively. Transparency gradually increased from 25 to 55% in the incremental noise condition. Middle face transparency was 40% in Fig. 4. In the identity in fig. S1, transparency in the low and high noise condition was 15 and 40%, respectively. Transparency gradually increased from 15 to 40% in the incremental noise condition. In the identity in Fig. 3 and fig. S2, transparency in the low and high noise condition was 25 and 40%, respectively. Transparency gradually increased from 25 to 40% in the incremental noise condition. Middle face transparency was 32.5% in fig. S2.

Survey
In Qualtrics, participants were presented with a video frame (width, 560; height, 315 pixels) and preliminary instructions on how to initiate the one-trial video: "Please click on this video once to start it and double click to activate Full Screen mode. Follow the instructions of the video and insert the response in the interface below at the end:". When participants clicked on the video, survey instructions were shown for 30 seconds: "In this survey you will see a face for a random period that goes from 5 to 35 seconds. Please stare at the face as much as you can during the entire period. A blue frame will randomly appear during the survey, and you will be asked a question about what you see inside the frame. After giving the answer with the interface below, please click on the blue arrow button. Important: you are allowed to do this survey only once, without going back to the beginning or stopping it."

After the survey instructions, a reference face alone or a test face preceded by the face morphing movie appeared for 5 s. At the same time, task instructions appeared. Aging task instructions were as follows: "What is the age of **this** face? Please type the age number in the interface below." Participants were asked to type the age number in an interface below. Gender task instructions were as follows: "What is the gender of **this** face? Adjust the dot in the reference below to rate the gender of the scale (MALE=0 FEMALE=100)." Participants were asked to adjust a rating bar from 0 to 100% (bar starting point was 50%). After 5 s, the test/reference face and task disappeared, and participants were asked to "give the answer using the interface below and click on the blue arrow to continue," thus ending the survey. Task instructions were shown for only 5 s when the reference or test face appeared, and participants were unable to move the video backward or forward (video controls were disabled). In experiments 6 and 10, as a confirmation that the gender task was fully understood, participants were additionally asked at the end of the survey to indicate the previous question among 10 possible alternatives: age, attractiveness, shape, size, image ratio, image duration, gender, image visibility, race, and empathy. Only participants responding "Gender" were considered for the study. For each condition (reference or test face), 70 participants were recruited.

Data analysis

Exclusion criteria were the following. First, repeated IP address from the participant. Second, unrelated response. In the age rating task, responses were excluded if they did not include a number between 1 and 100. In the gender rating task, responses were excluded if gender was not chosen in the additional question at the end of the survey. Third, outlier responses. Responses were excluded if age or gender ratings exceeded 2 standard deviations from the mean on each condition. On average, when all the criteria were considered, 20% of the data were removed for each condition, thus leaving an average number of ~50 participants for each condition.

Sample sizes for each condition in each experiment. Each number here refers to the number of independent individual observers in each condition. Each observer only participated in a single trial. Experiment 1 (incremental noise): Reference Young (low noise): 44; Reference Young (high noise): 45; Reference Old (low noise): 44; Reference Old (high noise): 44; Young-to-Old: 47; Old-to-Young: 45. Experiment 2 (no noise): Reference Young: 45; Reference Old: 44; Young-to-Old: 47; Old-to-Young: 45. Experiment 2 (stable high noise): Reference Young: 45; Reference Old: 44; Young-to-Old: 46; Old-to-Young: 62. Experiment 3: Continuous video (see experiment 1, incremental noise). Six steps: Young-to-Old: 46; Old-to-Young: 42. Four steps: Young-to-Old: 45; Old-to-Young: 46. Two steps: Young-to-Old: 43; Old-to-Young: 48. Experiment 4: ISI 0: Young-to-Old: 47; Old-to-Young: 45. ISI 1: Young-to-Old: 43; Old-to-Young: 49. ISI 5: Young-to-Old: 46; Old-to-Young: 47. ISI 10: Young-to-Old: 53; Old-to-Young: 51. ISI 15: Young-to-Old: 58; Old-to-Young: 51. Experiment 5: Test Frame 2: 44; Test Frame 450: 47; Test Frame 1050: 44; Test Frame 1350: 72; Test Frame 1450: 47; Test Frame 1850: 51; Test Frame 2250: 45; Test Frame 2650: 45; Test Frame 3850: 44; Test Frame 4650: 45; Test Frame 5450: 42; Test Frame 5850: 44; Test Frame 7050: 53; Video Frame 2: 48; Video Frame 450: 45; Video Frame 1050: 45; Video Frame 1350: 49; Video Frame 1450: 48; Video Frame 1850: 45; Video Frame 2250: 43; Video Frame 2650: 56; Video Frame 3850: 58; Video Frame 4650: 54; Video Frame 5450: 37; Video Frame 5850: 46; Video Frame 7050: 55. Experiment 6: Reference Male (low noise): 48; Reference

Male (high noise): 50; Reference Female (low noise): 44; Reference Female (high noise): 45; Male-to-Female: 53; Female-to-Male: 50. Experiment 7: Reference Young (low noise): 44; Reference Young (high noise): 45; Reference Old (low noise): 44; Reference Old (high noise): 44; Reference Middle: 71; Young-to-Middle: 54; Old-to-Middle: 50. Experiment 8: Reference Young (low noise): 49; Reference Young (high noise): 46; Reference Old (low noise): 54; Reference Old (high noise): 46; Young-to-Old: 56; Old-to-Young: 54. Experiment 9: 1: Identity Experiment 1; 2: Identity Experiment 8. Young 1–to–Old 2: 89; Young 2–to–Old 1: 53; Old 1–to–Young 2: 59; Old 2–to–Young 1: 48. Experiment 10: Male-to-Middle: 50; Female-to-Middle: 54. Experiment 11: Six seconds: Young-to-Old: 44; Old-to-Young: 36. Eighteen seconds: Young-to-Old: 40; Old-to-Young: 42. Experiment 12: 11 participants.

For each condition, we generated mean and confidence intervals by calculating a bootstrapped distribution of the average values by resampling the data with replacement 5000 times. Attraction index was computed by calculating the ratio between the test rating (e.g., old test face) and the absolute difference between the two reference face ratings (e.g., young and old). In the incremental noise conditions, absolute difference was computed by calculating the difference between reference faces with low noise (starting point) and high noise (ending point). Hence, four reference ratings (e.g., young/old and low/high noise) were computed for each comparison with two test faces (e.g., young-to-old, test old; old-to-young, test young). Increasing and decreasing age directions were equally balanced.

For each video condition, we generated confidence intervals by calculating a bootstrapped distribution of the values by resampling the data with replacement 5000 times. Attraction indexes reported in the figures are mean bootstrapped averages.

REFERENCES AND NOTES

1. I. Al-Haytham, *The optics of Ibn Al-Haytham (Al Sabra, Trans.)* (Warburg Institute, 1083).
2. H. Von Helmholtz, *Handbuch der Physiologischen Optik [Handbook of physiological optics]* (Voss, 1866).
3. D. J. Simons, C. F. Chabris, Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception* **28**, 1059–1074 (1999).
4. D. J. Simons, R. A. Rensink, Change blindness: Past, present, and future. *Trends Cogn. Sci.* **9**, 16–20 (2005).
5. R. A. Rensink, Change detection. *Annu. Rev. Psychol.* **53**, 245–277 (2002).
6. A. Mack, I. Rock, *Inattentional Blindness* (MIT Press, 1998), vol. 33.
7. U. Neisser, The control of information pickup in selective looking, in *Perception and Its Development: A Tribute to Eleanor J. Gibson* (Psychology Press, 1979), pp. 201–219.
8. A. Hollingworth, A. M. Richard, S. J. Luck, Understanding the function of visual short-term memory: Transsaccadic memory, object correspondence, and gaze correction. *J. Exp. Psychol. Gen.* **137**, 163–181 (2008).
9. S. J. Luck, E. K. Vogel, The capacity of visual working memory for features and conjunctions. *Nature* **390**, 279–281 (1997).
10. J. Fischer, D. Whitney, Serial dependence in visual perception. *Nat. Neurosci.* **17**, 738–743 (2014).
11. G. M. Cicchini, K. Mikellidou, D. Burr, Serial dependencies act directly on perception. *J. Vis.* **17**, 6 (2017).
12. M. Manassi, A. Liberman, A. Kosovicheva, K. Zhang, D. Whitney, Serial dependence in position occurs at the time of perception. *Psychon. Bull. Rev.* **25**, 2245–2253 (2018).
13. J. Taubert, D. Alais, D. Burr, Different coding strategies for the perception of stable and changeable facial attributes. *Sci. Rep.* **6**, 32239 (2016).
14. A. Abrahamyan, L. L. Silva, S. C. Dakin, M. Carandini, J. L. Gardner, Adaptable history biases in human perceptual decisions. *Proc. Natl. Acad. Sci. U.S.A.* **113**, E3548–E3557 (2016).
15. J.-M. Lueckmann, J. H. Macke, H. Nienborg, Can serial dependencies in choices and neural activity explain choice probabilities? *J. Neurosci.* **38**, 3495–3506 (2018).

16. A. Braun, A. E. Urai, T. H. Donner, Adaptive history biases result from confidence-weighted accumulation of past choices. *J. Neurosci.* **38**, 2418–2429 (2018).

17. A. Kiyonaga, J. M. Scimeca, D. P. Bliss, D. Whitney, Serial dependence across perception, attention, and memory. *Trends Cogn. Sci.* **21**, 493–497 (2017).

18. D. P. Bliss, J. J. Sun, M. D'Esposito, Serial dependence is absent at the time of perception but increases in visual working memory. *Sci. Rep.* **7**, 14739 (2017).

19. M. Fritsche, P. Mostert, F. P. de Lange, Opposite effects of recent history on perception and decision. *Curr. Biol.* **27**, 590–595 (2017).

20. G. M. Cicchini, K. Mikellidou, D. C. Burr, The functional role of serial dependence. *Proc. Biol. Sci.* **285**, 20181722 (2018).

21. A. Liberman, J. Fischer, D. Whitney, Serial dependence in the perception of faces. *Curr. Biol.* **24**, 2569–2574 (2014).

22. K. Turbett, R. Palermo, J. Bell, D. A. Hanran-Smith, L. Jeffery, Serial dependence of facial identity reflects high-level face coding. *Vision Res.* **182**, 9–19 (2021).

23. A. Liberman, M. Manassi, D. Whitney, Serial dependence promotes the stability of perceived emotional expression depending on face similarity. *Atten. Percept. Psychophys.* **80**, 1461–1473 (2018).

24. G. Mei, S. Chen, B. Dong, Working memory maintenance modulates serial dependence effects of perceived emotional expression. *Front. Psychol.* **10**, 1610 (2019).

25. D. Pascucci, G. Mancuso, E. Santandrea, C. Della Libera, G. Plomp, L. Chelazzi, Laws of concatenated perception: Vision goes for novelty, decisions for perseverance. *PLoS Biol.* **17**, e3000144 (2019).

26. E. C. Poulton, S. Poulton, *Bias in Quantifying Judgements* (Taylor & Francis, 1989).

27. H. L. Hollingworth, The central tendency of judgment. *J. Philos. Psychol. Sci. Methods* **7**, 461–469 (1910).

28. G. M. Cicchini, G. Anobile, D. C. Burr, Compressive mapping of number to space reflects dynamic encoding mechanisms, not static logarithmic transform. *Proc. Natl. Acad. Sci.* **111**, 7867–7872 (2014).

29. K. Yarrow, P. Haggard, R. Heal, P. Brown, J. C. Rothwell, Illusory perceptions of space and time preserve cross-saccadic perceptual continuity. *Nature* **414**, 302–305 (2001).

30. P. Tse, P.-J. Hsieh, The infinite regress illusion reveals faulty integration of local and global motion signals. *Vision Res.* **46**, 3881–3885 (2006).

31. Z.-L. Lu, L. A. Lesmes, G. Sperling, Perceptual motion standstill in rapidly moving chromatic displays. *Proc. Natl. Acad. Sci.* **96**, 15374–15379 (1999).

32. M. Lisi, P. Cavanagh, Dissociation between the perceptual and saccadic localization of moving objects. *Curr. Biol.* **25**, 2535–2540 (2015).

33. A. Tversky, D. Kahneman, Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131 (1974).

34. J. H. McDermott, M. Schemitsch, E. P. Simoncelli, Summary statistics in auditory perception. *Nat. Neurosci.* **16**, 493–498 (2013).

35. A. R. Albrecht, B. J. Scholl, Perceptually averaging in a continuous visual world: Extracting statistical summary representations over time. *Psychol. Sci.* **21**, 560–567 (2010).

36. S. C. Chong, A. Treisman, Representation of statistical properties. *Vision Res.* **43**, 393–404 (2003).

37. J. Haberman, T. Harp, D. Whitney, Averaging facial expression over time. *J. Vis.* **9**, 1–113 (2009).

38. D. Whitney, A. Yamanashi Leib, Ensemble perception. *Annu. Rev. Psychol.* **69**, 105–129 (2018).

39. S. R. Schweinberger, R. Zäske, C. Walther, J. Golle, G. Kovács, H. Wiese, Young without plastic surgery: Perceptual adaptation to the age of female and male faces. *Vision Res.* **50**, 2570–2576 (2010).

40. S. F. O'Neil, M. A. Webster, Adaptation and the perception of facial age. *Vis. Cognit.* **19**, 534–550 (2011).

41. M. Fornaciai, J. Park, Spontaneous repulsive adaptation in the absence of attractive serial dependence. *J. Vis.* **19**, 21–21 (2019).

42. Y. Murai, D. Whitney, Serial dependence revealed in history-dependent perceptual templates. *Curr. Biol.* **31**, 3185–3191.e3 (2021).

43. J. Barbosa, H. Stein, R. L. Martinez, A. Galan-Gadea, S. Li, J. Dalmau, K. C. S. Adam, J. Valls-Solé, C. Constantinidis, A. Compte, Interplay between persistent activity and activity-silent dynamics in the prefrontal cortex underlies serial biases in working memory. *Nat. Neurosci.* **23**, 1016–1024 (2020).

44. G. Goldin, A. Darlow. (2013), vol. TurkGate (Version 0.4.0) [Software].

45. J. W. Peirce, Generating stimuli for neuroscience using PsychoPy. *Front. Neuroinform.* **2**, 10 (2009).