Version of Record: https://www.sciencedirect.com/science/article/pii/S0042698922000207
Manuscript_4943b77ff483492baaab338d0c64b2b2

Structure of visual biases revealed by individual differences

Mark Wexler [a,b,1], *Pascal Mamassian* [a,c], *Alexander C. Schütz* [d]
[a]*CNRS,* [b]*Université de Paris,* [c]*École Normale Supérieure, PSL University,* [d]*University of Marburg*

Multiple studies have shown that certain visual stimuli are perceived in accordance with strong biases that are both robust within individuals and highly variable from one individual to the next. These biases undergo small changes over time that demonstrate that they constitute latent states of the visual system. The literature to date indicates that the individual biases for different stimulus classes are independent of each other. Here we asked whether some of these biases are nonetheless related to one another. We measured individual biases for five classes of stimuli in 1000 participants. The stimuli were two different versions of two-dimensional apparent motion, smooth motion in Glass patterns, and two different structure-from-motion stimuli. There were pronounced individual biases in all stimuli, and these biases varied in direction and strength across individuals. Some biases were not independent: the two biases for apparent motion direction were most strongly correlated, and they were both correlated, but less strongly, to the bias direction for smooth motion. While all other pairs of biases had unrelated directions, the strengths of all biases were correlated. The correlation of bias strengths may be due to either a common factor across the stimulus types, or an attentional effect. Only a tiny fraction of the between-participant variance can be explained by age and gender. These results show that latent states of the visual system that we measure as individual biases are organized in a structured way, and call out for further study of this under-explored aspect of visual perception.

Keywords: vision, perceptual bias, individual differences, motion perception, ambiguous motion

Introduction

Recently it has become clear that the appearance of many supra-threshold visual stimuli is highly observer-dependent, as revealed both by formal experiments (Afraz et al., 2010; Boeykens et al., 2019; Carter & Cavanagh, 2007; Cretenoud et al., 2021; Drissi-Daoudi et al., 2020; Finlayson et al., 2020; Goutcher, 2016; Greenwood et al., 2017; Grzeczkowski et al., 2017; Houlsby et al., 2013; Hwang & Schütz, 2020; Kosovicheva & Whitney, 2017; Mamassian & Wallace, 2010; Moutsiana et al., 2016; Schütz, 2014; Schütz & Mamassian, 2016; Schwarzkopf et al., 2011; Schwarzkopf & Rees, 2013; Wang et al., 2020; Wexler, 2018; Wexler et al., 2015; Witzel et al., 2017) as well as by popular phenomena such as #TheDress and #TheShoe. Similar idiosyncratic phenomena exist in auditory perception (Deutsch, 1986; Pressnitzer et al., 2018), in temporal processing (Grabot & Kayser, 2020; Grabot & van Wassenhove, 2017), or in motor control (Schütz, 2014; Lebovich et al., 2019). Robust idiosyncratic effects seem to occur nearly everywhere we look for them. Individual differences are often described in

[1] Corresponding author. Email: mark.wexler@u-paris.fr. Postal address: Integrative Neuroscience and Cognition Center, CNRS UMR 8002, Université de Paris, 45 rue des Saints-Pères, 75006 Paris, France.

terms of perceptual biases where the same physical stimulus is consistently perceived by each individual, but in different ways by different individuals.

Are individual differences in perception permanent, or can they change over time? The answer to this question depends on their neural substrate, which can be more permanent or transitory. On the most permanent end of the spectrum, individual differences could be due to interindividual variations in brain anatomy such as differences in size or structure of cortical areas (Llera et al., 2019). Anatomical variations have been associated with some idiosyncratic perceptual effects (Moutsiana et al., 2016; Schwarzkopf et al., 2011; Schwarzkopf & Rees, 2013). Less permanent are interindividual differences in brain connectivity (Finn et al., 2015), which may undergo long-term changes due to learning and perceptual experience. Connectivity differences have been tied to interindividual perceptual differences (Genç et al., 2011). Finally, the brain, as a complex system, can obviously be in one of many transitory states (Bassett & Gazzaniga, 2011; Sporns, 2010; Lebovich et al., 2019; Turkheimer et al., 2021), as manifested through spiking patterns on single-neuron or population levels, or through synaptic weights. Although not as commonly evoked as the more permanent factors, interindividual differences in transitory brain states may also underlie interindividual differences in perception. Obviously, the more transitory the neural substrate, the more individual perceptual factors or biases can change over time.

Past studies that have measured individual biases in a few sessions over time have noted that biases generally seem stable (e.g., Carter & Cavanagh, 2007; Drissi-Daoudi et al., 2020). In our own previous work, we have shown that idiosyncratic differences—while remaining quite stable in many cases—do undergo change, and aspects of this change are systematic. Using large datasets with measurements on time scales ranging from seconds to months, we have shown that idiosyncratic perceptual biases undergo changes that obey certain regularities. For example, idiosyncratic biases can be momentarily deviated from their stable values, but tend to 'decay' or return to their original values over several tens of seconds (Wexler et al., 2015). When measured over time scales of hours (Wexler, 2018) or months (Wexler et al., 2015), idiosyncratic biases tend to undergo random walks or drifts. More important than the concrete dynamics is the fact the biases *preserve state*: they do not assume new, independent values every time they are measured. Instead, when they evolve, their new values tend to 'remember' past values, while adding small changes to these past values. Thus, at least certain kinds of idiosyncratic perceptual phenomena, which we measure and quantify as perceptual biases, do undergo change, with dynamics that are stateful, i.e., remembering their past values. These observations imply that at least some idiosyncratic perceptual phenomena are not due to permanent differences between observers' brains, but to more transitory differences in connectivity, or differences in the state of neural activations and synaptic weights.

Therefore idiosyncratic biases can be thought of as internal states of the brain, or memories. This leads us to ask an obvious question: how many of these internal states or memories exist? In our past work (Wexler et al., 2015) we studied two families of stimuli, which we called SFM (structure-from-motion) and TFM (transparency-from-motion). For each of the two families of stimuli, we measured an idiosyncratic bias, which was a two-dimensional vector. Measuring these two vectors in a large sample of observers, we found no correlation between them. In our other past work, we similarly found no correlations between idiosyncratic biases in transparency-from-motion and binocular rivalry (Hwang & Schütz, 2020) or between different motor responses to the same stimuli (Schütz, 2014). Thus, at least for these cases, it seems that there is an independent bias for each type of stimulus. In the present study, we ask whether there is an independent bias or internal state for *every* type of stimulus (and if so, what constitutes a stimulus type). If not, can we find stimuli whose biases are related?

In the current study we showed five types of stimuli to a large number (1000) of participants. Participants showed idiosyncratic differences in the perception of these stimuli, which we quantified by calculating a bias for each of the five stimulus types. All five biases were two-dimensional vectors. We computed between-participant correlations between these vectors in order to check if any of the biases were related to each other. In fact, we measured each of the biases twice, which allowed us to compare between-stimulus correlations to within-stimulus ones.

Methods

Participants

Participants were recruited using the crowdsourcing platform Prolific.co, where they had previously signed up to participate in internet-based studies. All demographic data was self-reported and gathered on the crowdsourcing platform prior to signing up for our study. Participants were required to have reported being fluent in English, and to have no mental health conditions that are uncontrolled by medication and that impact their daily lives. Participants were also required to have a desktop or laptop computer available to perform the study (rather than a tablet or telephone). 1000 participants completed the study, and were paid £2.10 (the median duration of the study was 16.4 minutes). According to self-reports, 58.3% of the participants were men, 41.2% were women, and 0.5% did not report their gender. Of the 97.4% of the participants who reported their age, the youngest and oldest participants were 18 and 84 years old, with 25%, 50% and 75% percentiles being 24, 32, and 44 years old. 84.8% of the participants reported being right-handed, 11.5% left-handed, and 3.6% ambidextrous. 91.2% reported that English was their first language; 48.4% reported that they currently resided in the UK, 45.6 in the USA, and the rest in other countries or did not report residence.

This study was approved by the local ethics committee (CER Université Paris Descartes). Participants were deemed to have given informed consent by proceeding to the on-line experiment after having been informed of the nature of the study.

Apparatus

The experiment was performed within an internet browser window on each participant's own computer (either desktop or laptop, but neither tablet nor telephone). From data reported by the participant's internet browser, we determined that 73.7% of the participants had Windows as their operating system, 19.2% Mac OS, 4.1% Chrome OS, and 3% other types of Linux. The most common browsers were Chrome (73.7%), Firefox (14%), Edge (6.1%), Safari (4.4%), and Opera (1.5%) (Internet Explorer was excluded because it was incompatible with the experiment software). The most common monitor resolutions were 1920x1080 (23.9%), 1366x768 (22.7%), 1440x900 (12%), 1536x864 (11.3%), and 1280x720 (5.8%). The experimental software automatically switched to full-screen mode whenever possible (and otherwise requested that the participant do so manually). We found that 90.9% of our participants were in full-screen mode; in the remainder, the browser window filled a median of 84.6% of the area of the monitor. The experimental software attempted to draw the visual stimulus on every monitor refresh. We calculated an approximate monitor refresh rate by taking the mean of the actual number of refreshes per second during all stimuli. We found that 90.9% of the participants had refresh rates between 50 and 70 Hz, 2.3% below 50 Hz, and 6.8% above 70 Hz.[2]

[2] We attempted to draw the stimulus on every monitor refresh using the requestAnimationFrame() function in Javascript. Low refresh rates were likely due to underpowered or overloaded computers, while high rates were

Stimuli

There were five types of stimuli. They have in common that each stimulus has an ambiguous orientation (direction of motion or tilt), which we will call □, and is compatible with two orientations, □ and □ + 180 deg.[3] Each block of 16 trials had equally spaced values of □ = 5.625, 16.875, …, 163.125, 174.375 deg, in random order. The duration of all stimuli was 900 ms. Examples of stimuli are shown in Figure 1. In the descriptions of the stimuli below and throughout the rest of this article, we will describe lengths and positions using approximate lengths on the monitor.[4]

AM1 (apparent motion of type 1: see http://lab-perception.org/demo/penta?exp=am1). A grating in which the luminance varied sinusoidally in one dimension, with spatial period 1.06 cm, tilted counterclockwise from the vertical by angle □. Two alternating images were shown: a first image with a random phase, and a second one in which the first grating was displaced in the direction θ by half a cycle, i.e., by 0.53 cm. Each image was shown for 150 ms, so the temporal frequency was 3.33 Hz. The direction of motion was therefore ambiguous, either □ or □ + 180 deg—the directions of the two smallest displacements compatible with the stimulus (Ullman, 1979). Participants reported perceived motion direction, by selecting one of two arrows.

AM2 (apparent motion of type 2: see http://lab-perception.org/demo/penta?exp=am2). Similar to AM1, except that the sinusoidal grating was replaced by a two-dimensional square grid of circular dots (radius 0.1 cm). Initial offset or phase in both directions was random. The stimulus is compatible with motions in the directions □ and □ + 180 deg, for the same reason as AM1. Participants reported perceived motion direction, by selecting one of two arrows. Stimuli similar to AM1 and AM2 have been used before to study idiosyncratic biases in apparent motion perception (Laubrock et al., 2008; Morikawa & McBeath, 1992; Schütz, 2014; Williams et al., 2003).

GLA (Glass motion: see http://lab-perception.org/demo/penta?exp=gla). Static images in which pairs of dots are separated by the same displacement vector (or rotation), known as Glass patterns, create a vague impression of motion along the axis of the vector, but whose direction is obviously ambiguous (Glass, 1969). Rapid sequences, in which successive frames are uncorrelated and therefore carry no coherent motion signals, create a much more vivid perception of motion along the axis of the displacement vector (Ross et al., 2000; Donato, Pavan & Campana, 2020). Each frame of this stimulus consisted of 100 randomly-placed pairs of circular dots (radius 0.05 cm, separation between centers 0.2 cm, some dots were occluded by the frame), separated by an axis oriented at angle □. On each stimulus frame (usually every 16.7 ms) a new image with different, uncorrelated dot pairs was displayed. Because this stimulus specifies only an axis and not a direction, it is compatible with motion in directions □ and □ + 180 deg. Participants reported perceived motion direction, by selecting one of

likely due to high-performance or variable refresh rate monitors. Regardless of monitor refresh rates, we attempted to draw stimuli at rates as close as possible to 60 frames per second.

[3] To describe the geometry of the stimuli, we will use a reference frame with origin at the center of the browser window, positive x axis pointing to the participant's right, y axis upwards, and z axis towards the participant. Angles are zero in the direction of the positive x axis, positive for the participant's counter-clockwise. For SFM stimuli, surface slant and tilt are defined by the normal vector (direction towards the participant): the angle between the normal vector and the z axis is the slant, the angle of the projection of the normal vector onto the xy plane the tilt.

[4] To draw items of approximately the same physical size on devices that have different pixel densities, browsers use an abstract unit of length called the 'CSS pixel,' defined as being *approximately* equal to 1/96 of an inch (World Wide Web Consortium, 2019). Our experimental software uses these units because it is the best automatic way to ensure roughly consistent lengths. For convenience, these units have been converted to centimeters. It is important to keep in mind that CSS lengths can differ from physical lengths by 10% or more.

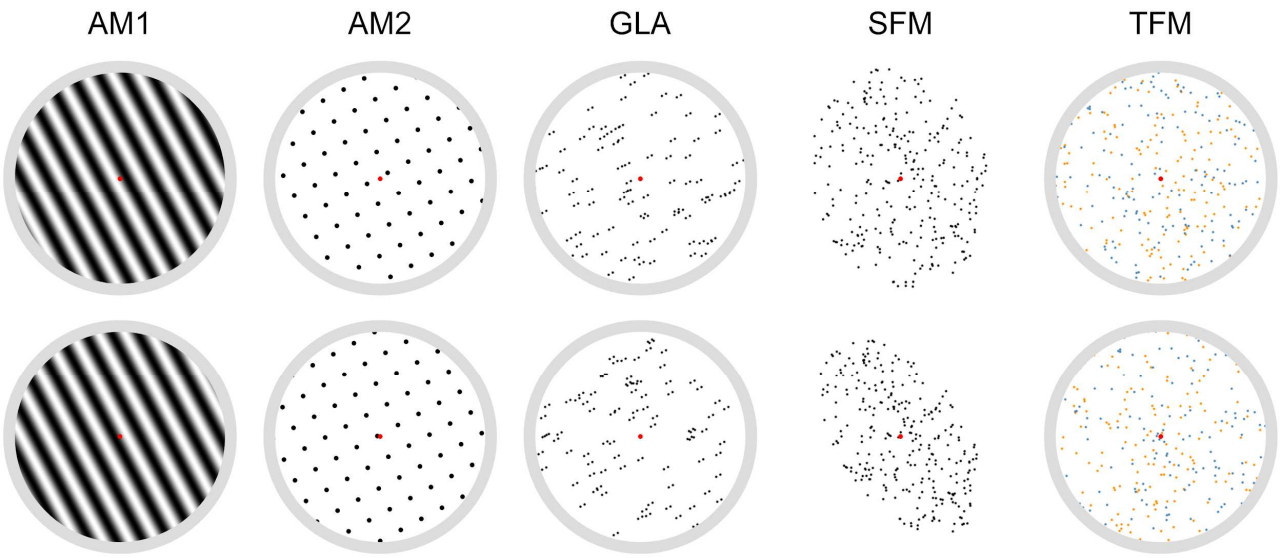

Figure 1. Illustrations of two frames of each of the five stimulus types. In this figure, all of the directions or tilts are compatible with 28.125 deg, or the opposite direction, 208.125 deg. Animated examples can be seen at the web URLs given in the text. The images shown are screenshots from the actual experiment. **AM1.** A one-dimensional sinusoidal grating in luminance. The two frames differ by phase 180 deg, making the direction of apparent motion ambiguous. **AM2.** A two-dimensional array of dots with the same spatial frequency as AM1. As for AM1, the phase difference is 180 deg, resulting in ambiguous apparent motion. **GLA.** Each frame consists of 'Glass patterns' of pairs of dots, separated by the same 2d vector. A sequence of such Glass patterns with the same separation vector was shown (but with the frames independent of one another), resulting in perception of very fast motion along the axis of the separation vector, but with ambiguous direction. **SFM.** A set of dots simulating a plane slanted in depth, undergoing a rotation in depth. (The frames shown are the starting point and endpoint of the rotation.) Because we used a parallel projection, the tilt (the direction of the plane's norm, as projected onto the image) is ambiguous up to 180 deg. **TFM.** Two sets of dots move in opposite directions, without any depth cues. (For clarity the two sets are shown here in different colors; in the actual stimulus all dots were black. The frames shown are the starting point and endpoint of the motion.) When viewing this sequence, most observers perceive one of the two sets of dots as transparent and closer to them than the other set. Which of the two directions of motion is perceived as transparent and closer constitutes the 180-deg ambiguity.

two arrows. Although it is known that kinetic Glass patterns lead to the perception of fast motion (Ross et al., 2000), it has not been previously reported that the direction of this motion is subject to individual bias.

SFM (structure from motion: see http://lab-perception.org/demo/penta?exp=sfm). The stimulus consisted of 284 dots, placed randomly and uniformly on a disk of radius 5.0 cm. In order to arrive at its central orientation, notionally this disk was first rotated in depth by 45 deg about an axis pointing upwards, and then rotated about an axis pointing out of the monitor by the stimulus tilt, equal to angle . This surface therefore had slant of 45 deg, and a tilt that varied from trial to trial. During the trial, this virtual surface was rotated in depth about axis (cos , sin , 0) where = + 45 deg, in order to obtain an intermediate level of shear (Cornilleau-Pérès et al., 2002; van Boxtel et al., 2003). The rotation speed was 30 deg/s and the rotation was symmetric about the central orientation, so that over the 900 ms duration of the stimulus, the orientation of the surface varied by a rotation between -13.5 and +13.5 deg, about the above-mentioned axis, from the central orientation. On each frame, the positions of the dots were projected using parallel projection (i.e., a virtual dot having 3D coordinates (x, y, z) was drawn at coordinates (x, y) on the monitor), using circular dots of radius 0.05 cm. Because of the parallel

projection, this stimulus is compatible with tilts and + 180 deg. Participants reported perceived tilt by selecting one of two icons, depicting the possible tilts unambiguously.[5]

TFM (transparency from motion: see http://lab-perception.org/demo/penta?exp=tfm). Two sets of dots, without any depth cues, moved in opposite directions and + 180 deg. Each set of dots moved at speed 1.76 cm/s, and the average number of dots visible in each set was 127. Each dot was drawn as a circle of radius 0.05 cm. Although no explicit depth cues are present, for most observers the two sets of dots are seen at different depths, with one set of dots perceived as transparent and closer to the observer, and often as more salient, than the other (Gibson et al., 1959; see references in Mamassian & Wallace, 2010). Participants reported which of the two directions of motion, or + 180 deg, corresponded to the closer or more salient set.

All stimuli were displayed for 900 ms, at a frame rate as close as possible to 60 Hz. All but SFM were enclosed inside a circular frame, a light gray annulus with inside radius of 4.8 cm and outside radius of 5.3 cm. All stimuli were black (lowest available luminance level), with the exception of AM1 which was a sinusoidal luminance grating. The background was white (highest available luminance). A red fixation dot with radius 0.1 cm was shown in the center of the window throughout the stimulus, and for 750 ms immediately prior.

Procedure

After being recruited to take part in the study, participants' browsers were redirected to our own server which hosted the experiment. The experiment was programmed in Javascript, and made use of the HTML5 canvas element for high-performance display functions. Participants were asked to close other applications and browser tabs on their computers as much as possible, and to run the experiment in a calm environment where they would not be disturbed, and to sit upright throughout. The experimental software verified that participants were using a desktop or laptop computer, and a browser that was compatible with the software (a vast majority were). The browser window was then put into full-screen mode (if this was not possible, participants were asked to switch to full-screen mode manually).

The experiment consisted of 10 blocks, two of each of the five stimulus types. Each block consisted of 16 trials. The 16 trials in each block were in pseudorandom order, but the order of the trials was the same for the two blocks of the same stimulus type. In one group of participants (N = 105), the two blocks of each stimulus type immediately succeeded one another. In a second group (N = 895), each stimulus type occurred once in the first five blocks, and once again and in the same order in the second five blocks. In both groups, the first three stimulus types were AM1, AM2, and GLA (in random order), followed by SFM and TFM (in random order). This order of conditions was chosen to simplify instructions that were similar for the first three stimulus types.

Prior to the first block of each stimulus type, the participant was led through step-by-step instructions for performing the task, on two trials (which could be repeated if requested by the participant). They were then given five practice trials (which could also be repeated), after which the main 16 trials of the block were performed. When the two blocks of each stimulus type did not immediately follow one another, at the start of the second block two practice trials were performed.

[5] According to the standard definition of tilt, described above, surfaces seen from above (e.g., the floor) have upward tilt (90 deg), and ones seen from below have downward tilt (270 deg). Because this is counter-intuitive, the tilts reported in the Results section have been additionally rotated by 180 deg.

Each trial began with a fixation point shown by itself for 750 ms, followed by the stimulus whose duration was 900 ms. After this the stimulus and fixation point were replaced by two icons to the left and to the right of the center of the window, one of which the participant selected with the mouse cursor. For the AM1, AM2, GLA and TFM stimuli the icons were arrows pointing in opposite directions; for SFM, they were surfaces with opposite tilts. After guiding the cursor to one of the icons (which changed color when the cursor approached) and clicking a mouse button to confirm the response, participants were directed to return the cursor to the center of the window in order to begin the next trial. When a participant either failed to respond or move the cursor to the center of the window for more than 3 s, a textual reminder was displayed.

Results

Raw data is available at https://osf.io/qb4my/.

Bias distributions of individual stimuli

We computed two-dimensional bias vectors for each participant and each stimulus family as follows:

$$\mathbf{B} = \frac{1}{C_n N} \sum_{i=1}^{N} (\cos\theta_i, \sin\theta_i)$$

where θ_i is the reported motion direction for AM1, AM2, GLA, and TFM stimulus families, and the reported tilt for SFM; N is the number of trials. The normalization factor

$$C_n = \left(n \sin\frac{\pi}{2n} \right)^{-1}$$

is included in order for the maximum length of the bias vector to be equal to 1 (here n is the number of unique directions or tilts). The interpretation of the bias vector **B** is as follows. Its length (from 0 to 1) is a measure of bias strength. The maximum value, 1, corresponds to all reported directions lying in one 180 deg semicircle. The direction of the vector corresponds to the bias direction: the preferred direction for AM1, AM2, GLA, and TFM, and the preferred tilt for SFM.

Distributions of the bias vectors in our population are shown in Figure 2 for the five stimulus families. The figure shows the population distributions of the 2D vectors (inside unit disks because their maximum length is 1), as well as the distributions of their lengths (bias strengths) and angles (bias directions). As can be seen in the second row of Figure 2, the biases are clearly unevenly distributed. More precisely, bias strengths (shown as orange bars) are sharply peaked at their maximum, 1. For comparison, the blue bars in the second row show the distribution of bias strengths if there were no biases—if all responses were independent random choices (obtained by simulating random, independent responses on every trial). The differences between the empirical and null bias strength distributions are so large—given the N = 1000 sample size—that Kolmogorov-Smirnov p-values were too small to compute. Bias directions are also non-uniformly distributed around the circle: the Rayleigh test for uniformity yields $p < 10^{-40}$ for all five stimulus families (Mardia & Jupp, 1999).

An obvious pattern in the distributions shown in Figure 2 is the peaks in the cardinal directions. However, the four peaks in the cardinal directions tend to have unequal sizes. For all five stimuli, the

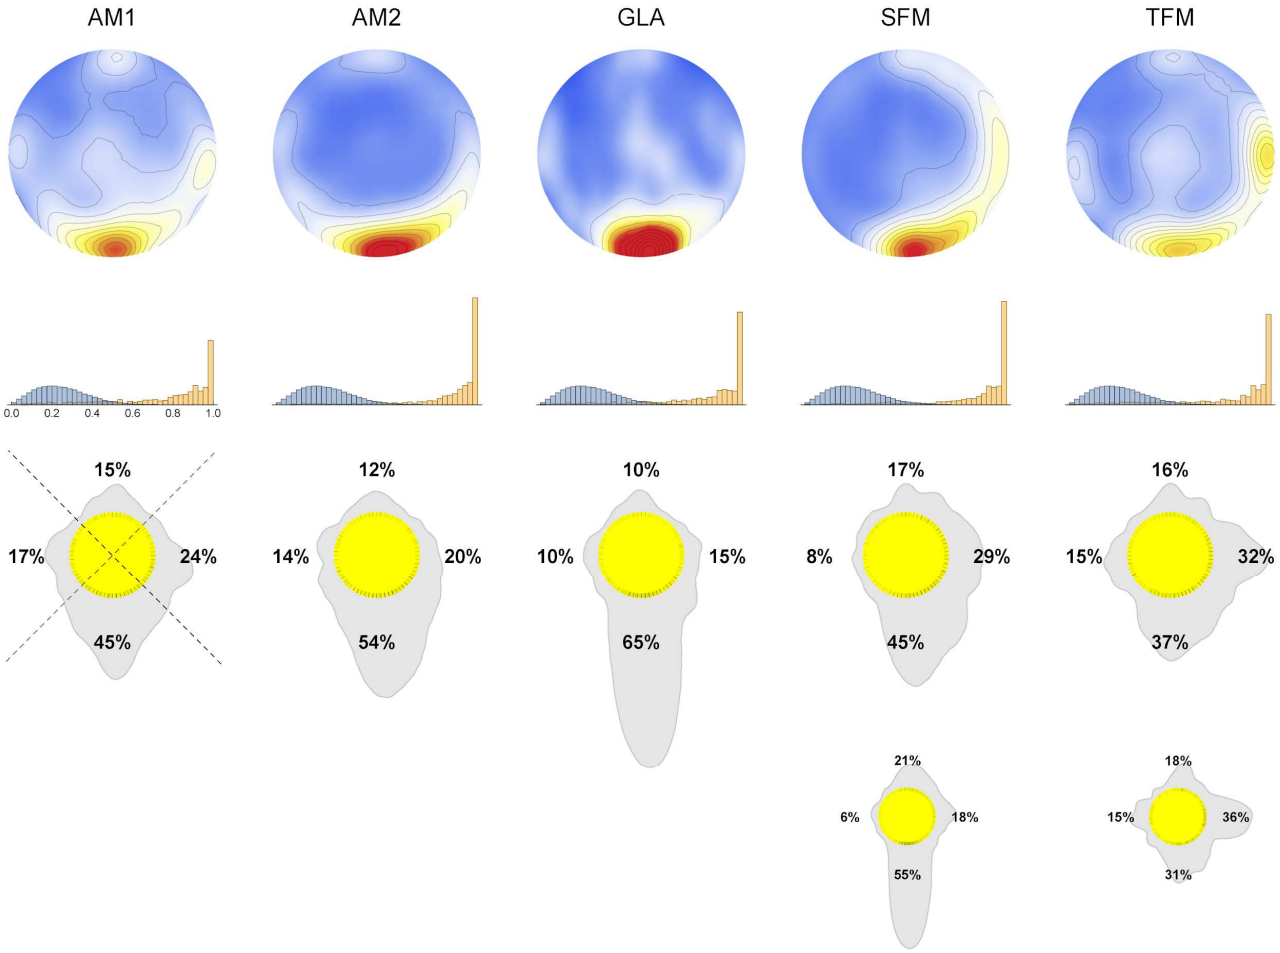

Figure 2. The distributions of biases for each of the five stimulus families. The top row shows the density of the 2d bias vectors in the population. The length of the bias vector, between 0 and 1, is a measure of bias strength; its direction is the orientation of the bias. Higher densities are shown in red, lower ones in blue. The middle row shows probability distribution functions of the bias strengths alone (orange histograms). The blue histogram in the left-hand graph shows the null-hypothesis distribution of bias strengths (if the response on each trial were independent and random). The bottom row shows PDFs for bias directions alone, as a polar plot: at any given polar angle, the distance between the outer curve and the rim of the yellow circle is proportional to the PDF of bias direction at that angle (individual bias directions are shown as strokes on the rim of the yellow circle). The four numbers show the percentages of bias directions lying in 90° sectors centered on the four cardinal directions. The two smaller polar histograms beneath the SFM and TFM plots show the population distributions of bias directions for nearly identical stimuli from Wexler et al. (2015).

peak to the bottom is stronger than the peak to the top, and the peak to the right stronger than the peak to the left (see the percentages on the bottom row of Figure 2). A bootstrap analysis corrected for false discoveries from multiple tests revealed that all bottom-top and right-left asymmetries were statistically significant, with the highest p value equal to 0.0011 (Efron & Tibshirani, 1993; Benjamini & Hochberg, 1995).[6]

[6] The tests were performed by resampling participants' biases, drawing randomly with replacement the same number of biases as the actual number of participants. For each resample we calculated an asymmetry index: for example, for a right-left asymmetry we calculated the number of biases falling within 45° of the rightward direction, and subtracted the number of biases falling within 45° of the leftward direction. We drew 10^5 such bootstrap resamples, and calculated the empirical distribution of the resulting asymmetry indices. The fraction of resamples with negative indices was an estimate for the probability of the null hypothesis (Efron & Tibshirani,

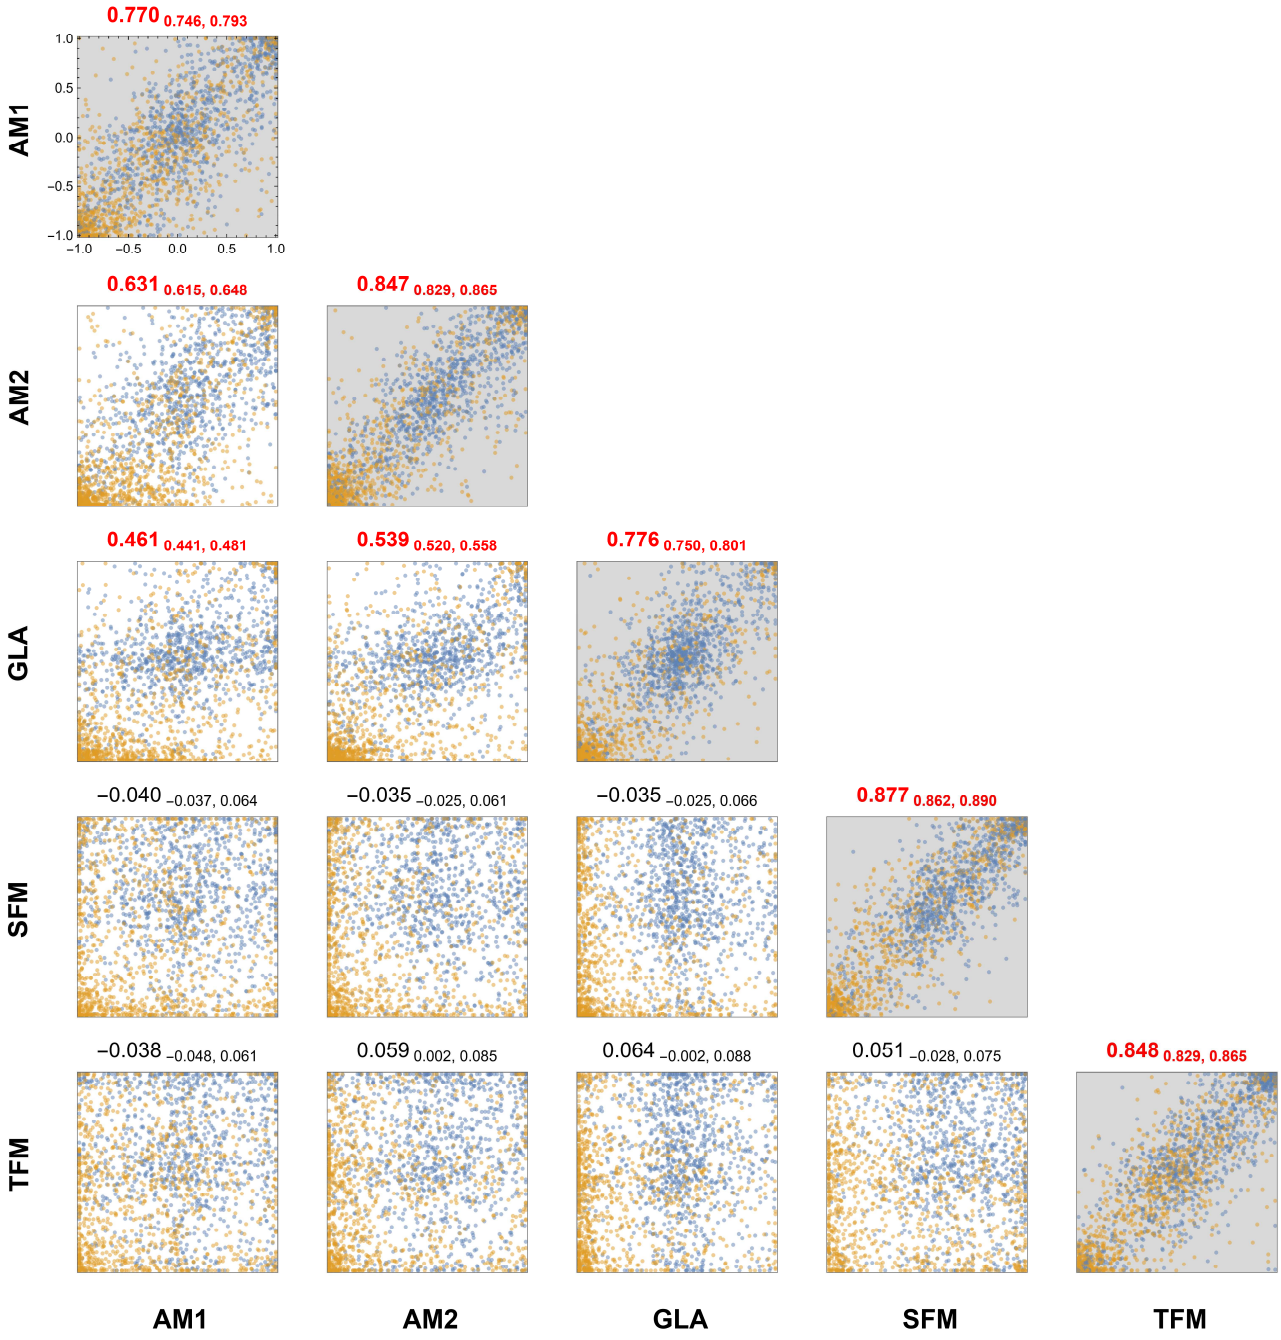

Figure 3. Between-participant correlations of the five biases. Each *off-diagonal* scatterplot shows the correlation between two different biases. The biases are 2-dimensional vectors, so their horizontal components are represented in blue, and their vertical components in orange. Each dot represents either the horizontal or vertical components of bias vectors for two different (see labels at bottom and left) in the same participant. Scatterplots on the *diagonal* represent correlations between the two repetitions of the same stimulus. To make the relationships between the dots easier to see, we jittered the dots by random numbers between -0.1 and +0.1 in both coordinates. The number above each scatterplot is the vector correlation coefficient, with uncorrected individual 95% confidence interval shown in the subscript (bootstrap with 10^4 samples). Coefficients in red are significantly greater than zero, as revealed by the bootstrap test with correction for false discovery (FDR 5%).

1993). We corrected the probability values for false discoveries due to multiple tests using the Benjamini-Hochberg (1995) procedure, with the false discovery rate set to 5%.

It is worth pointing out that we have previously measured the population distributions of the SFM and TFM biases in a large population that are very similar to those found here: these distributions are shown in the smaller comparison figures at the bottom of Fig. 2 (Wexler, Duyck & Mamassian, 2015).

To check to what extent participants' data showed biases at the individual level, we analyzed the distributions of individual responses for each stimulus type. We applied the Rayleigh test for anisotropy or non-uniformity (Mardia & Jupp, 1999), corrected by the Benjamini-Hochberg control for false discoveries, with the false discovery rate set to 5%. For the five stimulus types we found that the fraction of participants having significantly anisotropic responses ranged from 85% to 95% (85% for AM1, 86% for TFM, 90% for GLA, 92% for SFM, 95% for AM2).

Correlations between different biases

The major goal of this study was to check whether the five biases that we measured were independent, or whether they co-varied to some degree. In other words, if you know the bias for one stimulus type for a given participant, to what extent can you predict his or her bias for another stimulus type? Because biases are two-dimensional vectors, to carry out these analyses we used a measure of vector correlation, generalizing Pearson correlation coefficients to vectors (Hanson et al., 1992; Buneo, 2011). As with the Pearson correlation, a vector correlation coefficient of +1 denotes two sets of 2D vectors identical up to scale and translation, but also rotation; a coefficient of -1 denotes a similar relation, but with an additional reflection; and a coefficient of 0 denotes no particular relation of these kinds between the sets of vectors.

Figure 3 shows both the raw data of the bias vectors for each pair of stimulus conditions, and the correlation coefficients. Correlations between two sets of two-dimensional vectors are difficult to represent visually, because a scatterplot would require 4 dimensions. To simplify, we plot two-dimensional scatterplots, but with the two bias vectors for each participant represented as two dots: a blue dot representing the horizontal components of the two vectors (its horizontal position representing the value for one stimulus type, the vertical position for the other), and an orange dot representing the vertical components of the two vectors. This is for the off-diagonal scatterplots in Figure 3. For the scatterplots on the main diagonal—the ones with the gray background—we represent the same thing, but for the correlation between the two independent estimates for each bias, arising from the two repetitions for each stimulus type.[7]

We calculated vector correlations between each pair of biases. The values are given above each scatterplot in Figure 3. We tested whether the correlations were significantly different from zero using a bootstrap over participants, corrected by the Benjamini-Hochberg procedure (10^4 bootstrap samples, false discovery rate 5%). The correlation coefficients significantly different from zero are shown in red in Figure 3. Thus, the only significant correlations between *different* stimuli were between AM1 and AM2, between AM2 and GLA, and between AM1 and GLA (all $p < 10^{-4}$). Additionally, the correlations between the repetitions of each of the five stimuli, which we will call "self-correlations," were all significantly different from zero (all $p < 10^{-5}$, 10^5 bootstrap samples).[8] Further bootstrap tests on differences between correlation coefficients revealed that the AM1-AM2 coefficient is significantly

[7] To make the comparison of between-stimulus and within-stimulus correlation fair, for every pair of different stimuli we actually calculated 4 correlation coefficients, between each of the two repetitions of the first stimulus and each of the two repetitions of the second. We then took the mean of the four correlations using the z-transformation.

[8] We used different numbers of bootstrap samples in different analyses because some of the underlying measures were more time-consuming to compute.

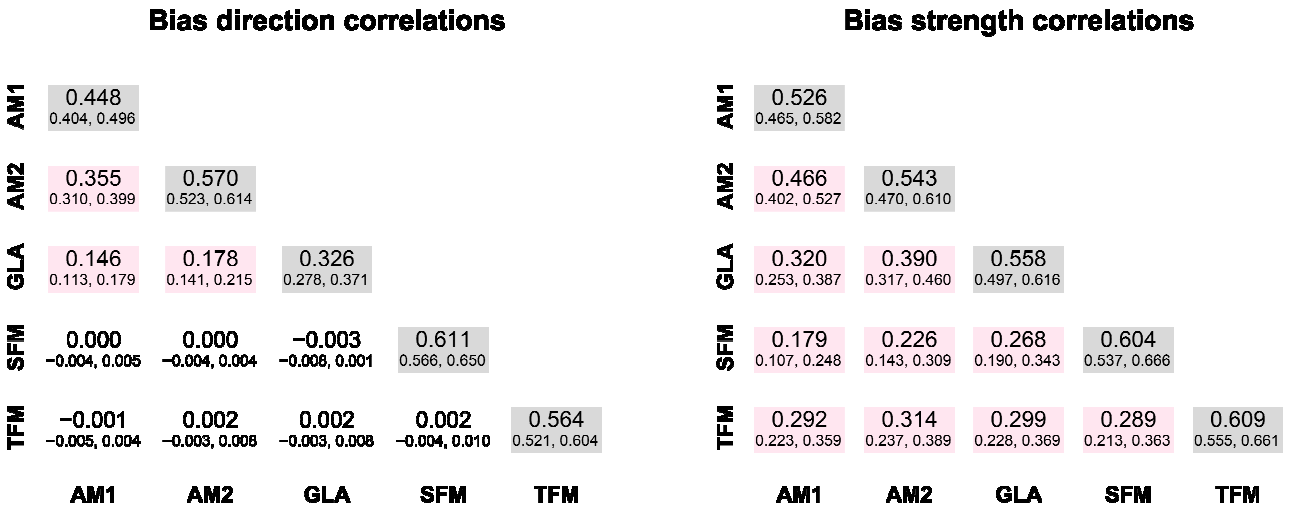

Table 1. Separate correlations between the directions and strengths of the biases. Left: Fisher-Lee Π circular correlations between bias directions (orientations of the bias vectors). Right: Spearman R correlations between bias strengths (Euclidean lengths of the bias vectors). The two smaller numbers below each correlation give the 95% confidence interval (derived from the bootstrap, 10^3 resamples for directions, 10^4 for strengths). Gray background: correlations between the biases from the two blocks of each stimulus. Pink and gray backgrounds: correlations significantly different from zero.

greater than the AM1-GLA and AM2-GLA coefficients, and that the AM2-GLA coefficient is significantly greater than the AM1-GLA one (all $p < 10^{-4}$, 10^4 bootstrap samples). Furthermore, similar tests revealed that each of the AM1-AM2, AM1-GLA, and AM2-GLA correlations was weaker (quantitatively by about 20-40%) than the self-correlations between the repetitions of two underlying biases (all $p < 10^{-4}$, 10^4 bootstrap samples). In other words, for each pair of different biases X and Y that had a significant correlation, this correlation was weaker than the self-correlations of X and Y with themselves. Thus, the temporal variations of the differences between stimulus types were stronger than the temporal variations of the biases themselves.

Since biases are two-dimensional vectors, we wondered whether the correlations that we have found apply to the bias strengths (the vectors' lengths), bias directions (the vectors' orientations), or both. We therefore calculated separate correlations between the strengths (using ordinary Pearson correlations) and the directions (using Fisher-Lee correlations of circular variables; see, e.g., Mardia & Jupp, 2000). The results, shown in Table 1, are in partial agreement with the ones obtained through vector correlation. The structure of direction correlations (left side of Table 1) mirrors the vector correlations: the largest correlation is between AM1 and AM2, second-largest AM1-GLA and AM2-GLA, and all other correlations not significant (tested using the bootstrap with 10^3 samples). The bias strength correlations, however, yield a surprise: *all* pairs of biases have significantly positive correlations—even though, as for the other correlation measures, the AM1-AM2 correlation is the highest, and AM1-GLA and AM2-GLA second-highest (bootstrap with 10^4 samples). Given our large sample size, the absence of correlations between bias directions is most likely not an issue of power but indeed reflects differences between bias direction and strength. Thus, some biases are unrelated in direction, but nonetheless correlated in their strengths, with correlation coefficient around $R \approx 0.25$.[9]

[9] The reader may have noticed that the separate correlations of directions and bias strengths (Table 1) are generally smaller than the vector correlations (Fig. 3). For the case of bias directions, this is due to the fact that the Fisher-Lee Π correlations are less robust under noise than vector correlations, which we have verified using simulations. In the case of bias strengths, the smaller correlation coefficients are due to ceiling effects, because most bias strengths are close to the maximum, 1 (see Fig. 2, second line).

	AM1	AM2	GLA	SFM	TFM
AM1					
AM2	0.165 0.108, 0.225				
GLA	0.066 −0.086, 0.119	0.086 −0.101, 0.139			
SFM	−0.054 −0.084, 0.108	−0.021 −0.070, 0.077	−0.050 −0.082, 0.102		
TFM	−0.034 −0.073, 0.089	0.061 −0.091, 0.111	−0.028 −0.079, 0.086	0.059 −0.069, 0.115	

Table 2. Vector correlations between changes in different biases across the two repetitions of each stimulus type. The small numbers below each correlation coefficient show the 95% confidence interval (bootstrap with 10^5 samples). The only significant correlation is shown with a pink background (corrected for false discovery at 5%).

Correlations between changes in biases

In most of our participants (N = 895) we first measured each of the five biases once in the first half of the experiment, and then we measured them again in the same order in the second half. A typical order of conditions could have been, for example, $AM2_{(1)}$, $GLA_{(1)}$, $AM1_{(1)}$, $TFM_{(1)}$, $SFM_{(1)}$, $AM2_{(2)}$, $GLA_{(2)}$, $AM1_{(2)}$, $TFM_{(2)}$, $SFM_{(2)}$. Thus, for any two conditions—say AM2 and AM1—there was an overlap of the period between the two measures of the first bias—$AM2_{(1)}$ and $AM2_{(2)}$ in our example—and the two measures of the second—$AM1_{(1)}$ and $AM1_{(2)}$. If the internal representations of two biases are wholly or partly shared, then the *change* in one of the biases across its two measurements may be correlated with the change in the other. (In the other 105 participants, the second measurement of each bias immediately succeeded the first, so there was no such overlap—which is why we exclude them from the following analysis.) The vector correlations of bias changes across the two repetitions of each bias measurement are shown in Table 2. A bootstrap test with Benjamini-Hochberg correction revealed that only one correlation, between AM1 and AM2, was significantly different from zero, with a vector correlation coefficient of 0.165 ($p = 0.004$, 10^5 bootstrap samples, false discovery rate 5%).

Effects of demographic variables

We wondered whether the population distributions of biases are fixed, or whether they vary across populations. From our crowdsourcing platform, we were able to obtain the values of four parameters that participants had previously reported in the platform's surveys:

- Age: ranging from 18 to 84, when reported
- Gender: the platform does not have a non-binary option, but we removed from the demographic analysis the small number of participants who did not report their gender. We wound up with 58% men and 42% women.
- Manual preference: we removed a small number of participants who reported that they were ambidextrous, leaving 88% right-handers and 12% left-handers.
- Country of residence: by far the two largest countries were the UK and USA, representing together 94% of our participants. We removed the other countries from the demographic

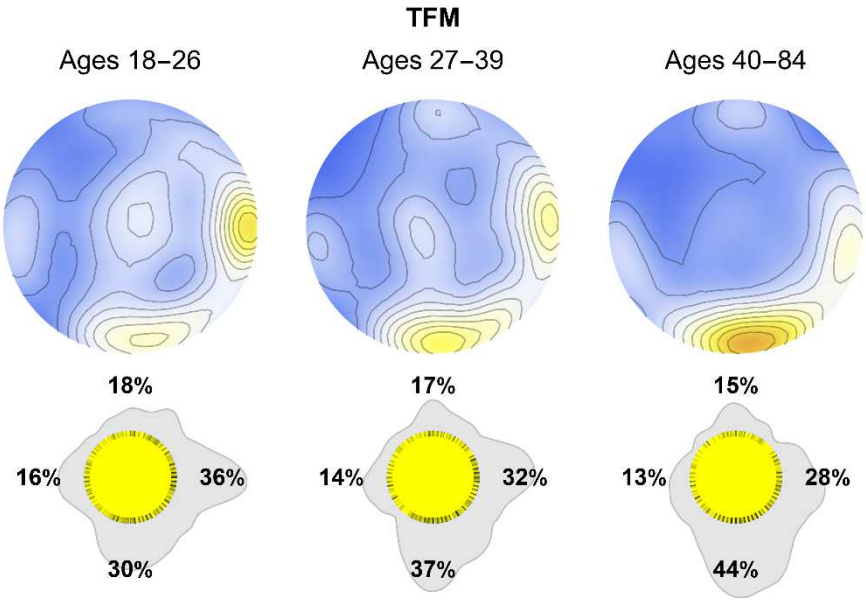

Figure 4. An example of a demographic effect on a population distribution of biases. Here, the effect of age on the distribution of TFM biases. The upper graphs show the distributions of the 2d bias vectors; the lower graphs are polar histograms representing distributions of bias directions. With age, the downward peak seems to grow, while the other peaks seem to shrink. Only the growth in the downward peak attains statistical significance.

analysis, leaving 53% UK and 47% USA.
Keeping only participants for whom we had a value for each of the four parameters left us with 879 participants out of 1000 for the demographic analysis.

An example of a demographic effect is shown in Figure 4. Here, age seems to have an effect on the shape of the TFM bias distribution. To visualize the effect of age on the bias distribution, we categorized age into 3 groups, dividing at approximate 1/3 and 2/3 quantiles. The downward peak appears to grow with age, while the other peaks appear to shrink. In the youngest group the rightward peak is larger than the downward one, while for the oldest participants the reverse is true. In the statistical tests that we describe in the next paragraph, we found that the effect of age on the vertical component of bias attained significance, among several other such effects.

In order to statistically test for the effects of demographic variables on biases, we performed linear regressions of the horizontal and vertical components of the biases on the four variables listed above.
A difficulty with our analysis is that the groups are unbalanced: for example, the gender ratio in older participants is not the same as in younger ones. An advantage of the linear regression is that it corrects for lack of balance. In order to test the significance of the regression coefficients, we performed a bootstrap analysis (10^4 repetitions), and corrected for false discoveries using the Benjamini-Hochberg correction. Because we had not planned to perform this analysis, we felt that a particularly stringent correction was in order, and therefore put the false discovery rate to 1%.

We found five significant effects of demographic variables on the biases, all of them on the vertical component (maximum $p = 0.0008$). There were three significant effects of age: older age was associated with lower (i.e., more negative) vertical components of GLA and TFM biases, and higher (more positive) vertical components of SFM biases. There were also two significant effects of gender: women had more negative vertical components of AM1 and GLA biases than men. The TFM effect is

illustrated in Figure 4, while the other effects are shown in Fig. S1 in the Supplementary Materials. (As cross-validation, the same five effects were the only ones found to have non-zero coefficients using a lasso regression with lasso parameter set to 0.1 (James et al., 2021).) However, it should be noted that these demographic effects, although significant because of our large sample size, are tiny in magnitude as compared to between-individual variations: on the average the 10 linear models only explain 1.7% of the variance of the bias components, with the most 'successful' model—the one for vertical component of the GLA bias—accounting for 3.4% of the variance. The details of the fits, including the coefficients and R^2 values, are shown in Table S1 in the Supplementary Materials.

Discussion

We found that each of the five stimuli that we studied admitted significant idiosyncratic differences, which we quantified as biases. The population distributions of the biases were non-uniform: they had peaks in the cardinal directions, and these peaks have left-right and up-down asymmetries (as we have previously shown for SFM and TFM stimuli: Wexler et al., 2015). The five biases were not independent: using vector correlations, we found that AM1 and AM2 were most closely related, and both were correlated with the GLA bias, though slightly less strongly. The SFM and TFM biases were independent of one another (as we had shown previously: Wexler et al., 2015), and uncorrelated with the other three biases. Additionally, changes in the AM1 and AM2 biases across the two measurements were significantly correlated.

We also separately calculated the correlations between the directions of the five biases, and between their strengths (the polar and radial components of the bias vectors). We found that the bias directions mirrored the vector correlations: AM1 and AM2 were most strongly correlated, both had weaker but significant correlations with GLA, and none of the other pairs were correlated. In contrast, we found that *all* five bias strengths had significant positive correlations (albeit with the AM1-AM2, AM1-GLA and AM2-GLA more strongly correlated than others). Therefore, we have found several pairs of biases whose strengths are correlated (with a Spearman $R \approx 0.25$, highly significant given our sample size) but whose directions seem to be entirely unrelated. Two scenarios seem possible. First, we may have discovered a non-specific tendency in certain people than in others to have stronger biases, across a variety of visual stimuli. However, we cannot exclude another scenario: some participants may have paid less attention than others across all tasks; less attention would have led to greater noise, which in turn would have lowered all biases. Although there is some indirect evidence against this second scenario—there were very few participants with weak biases across all stimuli (for example, only 30/1000 had all five biases with strength below 0.75), we think that it is not decisive. Another possible cause of spurious correlations could have been the large differences in apparatus between participants (lighting, quality of the display, distance to the monitor) but not between blocks within each participant. A lower-quality display, for instance, by also effectively injecting noise into the stimuli, could have lowered bias strengths across the board. In spite of the significant positive correlations between all bias strengths, therefore, we cannot be certain of the existence of a common factor across all five biases without further studies.

The observed differences between age and gender groups are difficult to interpret. The effect of age might be mediated by many other variables. There are many heterogeneous results about the effects of age on motion perception (reviewed by Billino & Pilz, 2019) highlighting the complexity of the involved processing stages and the differential effects of ageing. A specific issue of our internet-based study might be that participants were not tested under standardized conditions and that there might have been systematic differences in the testing conditions or equipment between different age groups. In

addition, despite the large sample size, our sample is not representative for any specific age group. Gender differences in perception have been reported for instance in color preferences (Bonnardel et al., 2018; Hurlbert & Ling, 2007; Sorokowski et al., 2014) and the weighting of different cues for establishing correspondence in apparent motion stimuli (Shechter et al., 1991), as well as in the performance of some visual tasks (Shaqiri, 2018). To our knowledge there are no reports of gender differences that might explain the observed difference in the vertical bias. In any case, the observed effects of age and gender in our study should not be overinterpreted because they explained only a small amount of variance.

The pattern of bias correlations that we have obtained has a certain logic to it. The AM1, AM2, and GLA biases govern the perception of two-dimensional motion: slow apparent motion for AM1 and AM2, fast streaming motion for GLA. The TFM bias, though also measured using reports of two-dimensional motion direction, involves perceiving motion in both the reported and the opposite directions, and has more to do with the way that motion directions are segmented in depth. It is therefore likely to be a *three-dimensional* motion bias, related to the preferred axis of depth rotation of the ambiguous sphere or cylinder rotation stimuli commonly used to study bistability (Andersen & Bradley, 1998; Leopold et al., 2002). Finally, the SFM bias is a preference for three-dimensional surface orientation. If this reasoning is correct, then our results indicate that all three of our biases for two-dimensional motion are correlated—with the two apparent-motion stimuli correlated to each other most strongly—while there are no significant correlations between two-dimensional motion, three-dimensional motion, and three-dimensional structure biases. Thus, our results seem to point to internal biases for broad features rather than specific stimuli.

A possible objection to the correlations between the AM1, AM2 and GLA biases is that they were due not to perception, idiosyncratic or not, but rather to response biases. Indeed, for these stimuli, the response consisted in choosing one of two arrows. Perhaps certain participants simply preferred clicking on, say, downward arrows, while others preferred rightward arrows, and so on? Such response biases could conceivably lead to correlations such as we have found. However, we should note that a fourth stimulus, TFM, had the exact same probe as AM1, AM2, and GLA: an arrow. If the correlations were due to response bias, we would expect similar correlations with TFM. The fact that we find no correlations between the TFM and any of the others argues strongly against a response-bias explanation.

A closely related debate has concerned common factors in switching between different interpretations of bistable stimuli. It has long been observed that subjects undergo perceptual switches at very different rates. For example, Carter and Pettigrew (2003) found that switching rates for both binocular rivalry and motion-induced blindness each varied by about a factor of 10 in their subject population. They also found that the two switching rates were significantly correlated to one another, concluding that this was evidence for a common oscillator governing both types of bistability. However, this and other results have been criticized on methodological grounds, namely that any correlations are actually due to decisional biases as to what constitutes a switch (Gallagher & Arnold, 2014). The correlations have sometimes been found difficult to replicate, and a meta-analysis has shown them to be quite weak (Brascamp et al., 2018). Nevertheless, there is evidence for some correlations (Brascamp et al., 2019; Cao et al., 2018; Pastukhov et al., 2019; Steinwurzel et al., 2020). Consistent with our results, correlations in the dynamics of ambiguous perception seem to be stronger for more similar stimuli.

More generally, in the past decade there have been a number of studies that have leveraged individual differences to look for a putative 'common factor' in vision (see Tulver, 2019 for a review). For the most part, these studies have found correlations only between very similar stimuli and tasks, and little

evidence for a general common factor in vision. For instance Grzeczkowski et al. (2017) measured the susceptibility of a pool of participants to a set of illusions, and found that only the Ponzo and Ebbinghaus illusions—the only two size illusions tested—were significantly correlated. Follow-up studies showed that variations on the same illusion were correlated but different illusions were not, and that the individual susceptibilities were stable in participants across time (Cretenoud et al., 2019, 2021). Our present finding of correlations between biases related to two-dimensional motion, but not three-dimensional motion or three-dimensional shape, echo this general tendency of multiple local, but no global, factors.

Ultimately the usefulness of our results lies in the information that they provide about the neural mechanisms underlying visual perception. We have shown that the perception of the five types of visual stimuli in this study is mediated by five latent variables that can only be manifestations of underlying brain states. These variables are not independent, but present a non-trivial set of correlations. In the past decade neuroscientists have become more interested in individual differences in human brains, both for their own sake and for what they have to tell about general brain function (Dubois & Adolphs, 2016). For instance, Charest el al. (2014) have shown that individual variations in infero-temporal activations predict individual variations in perceived similarities between objects. Psychophysical data about individual biases and their correlations may guide and broaden these techniques in neuroscience.

Acknowledgements

We are grateful to Luc Tamisier for his kind help in setting up and managing the web server used for this study. We would also like to thank an anonymous reviewer for comments that were extremely useful.

Data availability statement

Anonymized raw data, demographic information, and computed biases are available from the repository https://osf.io/qb4my/.

References

Afraz, A., Pashkam, M. V., & Cavanagh, P. (2010). Spatial heterogeneity in the perception of face and form attributes. *Current Biology*, *20*(23), 2112–2116. https://doi.org/10.1016/j.cub.2010.11.017

Andersen, R. A., & Bradley, D. C. (1998). Perception of three-dimensional structure from motion. *Trends in Cognitive Sciences*, *2*(6), 222–228. https://doi.org/10.1016/S1364-6613(98)01181-4

Bassett, D. S., & Gazzaniga, M. S. (2011). Understanding complexity in the human brain. *Trends in Cognitive Sciences*, *15*(5), 200–209. https://doi.org/10.1016/j.tics.2011.03.006

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.

Billino, J., & Pilz, K. S. (2019). Motion perception as a model for perceptual aging. *Journal of Vision*, *19*(4), 3. https://doi.org/10.1167/19.4.3

Boeykens, C., Wagemans, J., & Moors, P. (2019). Biases in the perception of the ambiguous motion quartet across spatial scale. *Journal of Vision*, *19*(10), 152a.

https://doi.org/10.1167/19.10.152a
Bonnardel, V., Beniwal, S., Dubey, N., Pande, M., & Bimler, D. (2018). Gender difference in color preference across cultures: An archetypal pattern modulated by a female cultural stereotype. *Color Research & Application*, *43*(2), 209–223. https://doi.org/10.1002/col.22188
Brascamp, J. W., Becker, M. W., & Hambrick, D. Z. (2018). Revisiting individual differences in the time course of binocular rivalry. *Journal of Vision*, *18*(7), 3. https://doi.org/10.1167/18.7.3
Brascamp, J. W., Qian, C. S., Hambrick, D. Z., & Becker, M. W. (2019). Individual differences point to two separate processes involved in the resolution of binocular rivalry. *Journal of Vision*, *19*(12), 15. https://doi.org/10.1167/19.12.15
Cao, T., Wang, L., Sun, Z., Engel, S. A., & He, S. (2018). The Independent and Shared Mechanisms of Intrinsic Brain Dynamics: Insights From Bistable Perception. *Frontiers in Psychology*, *9*, 589. https://doi.org/10.3389/fpsyg.2018.00589
Carter, O. L., & Cavanagh, P. (2007). Onset rivalry: Brief presentation isolates an early independent phase of perceptual competition. *PloS One*, *2*(4), e343. https://doi.org/10.1371/journal.pone.0000343
Carter, O. L., & Pettigrew, J. D. (2003). A common oscillator for perceptual rivalries? *Perception*, *32*(3), 295–305. https://doi.org/10.1068/p3472
Charest, I., Kievit, R. A., Schmitz, T. W., Deca, D., & Kriegeskorte, N. (2014). Unique semantic space in the brain of each beholder predicts perceived similarity. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(40), 14565–14570. https://doi.org/10.1073/pnas.1402594111
Cornilleau-Pérès, V., Wexler, M., Droulez, J., Marin, E., Miège, C., & Bourdoncle, B. (2002). Visual perception of planar orientation: Dominance of static depth cues over motion cues. *Vision Research*, *42*(11), 1403–1412.
Cretenoud, A. F., Grzeczkowski, L., Kunchulia, M., & Herzog, M. H. (2021). Individual differences in the perception of visual illusions are stable across eyes, time, and measurement methods. *Journal of Vision*, *21*(5), 26. https://doi.org/10.1167/jov.21.5.26
Cretenoud, A. F., Karimpur, H., Grzeczkowski, L., Francis, G., Hamburger, K., & Herzog, M. H. (2019). Factors underlying visual illusions are illusion-specific but not feature-specific. *Journal of Vision*, *19*(14), 12. https://doi.org/10.1167/19.14.12
Deutsch, D. (1986). A musical paradox. *Music Perception*, *3*, 275–280.
Donato, R., Pavan, A., & Campana, G. (2020). Investigating the Interaction Between Form and Motion Processing: A Review of Basic Research and Clinical Evidence. *Frontiers in Psychology*, *11*, 2867. https://doi.org/10.3389/fpsyg.2020.566848
Drissi-Daoudi, L., Doerig, A., Parkosadze, K., Kunchulia, M., & Herzog, M. H. (2020). How stable is perception in #TheDress and #TheShoe? *Vision Research*, *169*, 1–5. https://doi.org/10.1016/j.visres.2020.01.007
Dubois, J., & Adolphs, R. (2016). Building a Science of Individual Differences from fMRI. *Trends in Cognitive Sciences*, *20*(6), 425–443. https://doi.org/10.1016/j.tics.2016.03.014
Efron, B., & Tibshirani, R. J. (1993). *An Introduction to the Bootstrap.* Chapman & Hall/CRC.
Finlayson, N. J., Neacsu, V., & Schwarzkopf, D. S. (2020). Spatial Heterogeneity in Bistable Figure-Ground Perception. *I-Perception*, *11*(5), 2041669520961120. https://doi.org/10.1177/2041669520961120
Finn, E. S., Shen, X., Scheinost, D., Rosenberg, M. D., Huang, J., Chun, M. M., Papademetris, X., & Constable, R. T. (2015). Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. *Nature Neuroscience*, *18*(11), 1664–1671. https://doi.org/10.1038/nn.4135
Gallagher, R. M., & Arnold, D. H. (2014). Interpreting the temporal dynamics of perceptual rivalries. *Perception*, *43*(11), 1239–1248. https://doi.org/10.1068/p7648

Genç, E., Bergmann, J., Singer, W., & Kohler, A. (2011). Interhemispheric connections shape subjective experience of bistable motion. *Current Biology*, *21*(17), 1494–1499. https://doi.org/10.1016/j.cub.2011.08.003

Gibson, E. J., Gibson, J. J., Smith, O. W., & Flock, H. (1959). Motion parallax as a determinant of perceived depth. *Journal of Experimental Psychology*, *58*(1), 40–51. https://doi.org/10.1037/h0043883

Goutcher, R. (2016). Motion direction influences surface segmentation in stereo transparency. *Journal of Vision*, *16*(15), 17. https://doi.org/10.1167/16.15.17

Grabot, L., & Kayser, C. (2020). Alpha Activity Reflects the Magnitude of an Individual Bias in Human Perception. *Journal of Neuroscience*, *40*(17), 3443–3454. https://doi.org/10.1523/JNEUROSCI.2359-19.2020

Grabot, L., & van Wassenhove, V. (2017). Time Order as Psychological Bias. *Psychological Science*, *28*(5), 670–678. https://doi.org/10.1177/0956797616689369

Greenwood, J. A., Szinte, M., Sayim, B., & Cavanagh, P. (2017). Variations in crowding, saccadic precision, and spatial localization reveal the shared topology of spatial vision. *Proceedings of the National Academy of Sciences*, *114*(17), E3573–E3582. https://doi.org/10.1073/pnas.1615504114

Grzeczkowski, L., Clarke, A. M., Francis, G., Mast, F. W., & Herzog, M. H. (2017). About individual differences in vision. *Vision Research*, *141*, 282–292. https://doi.org/10.1016/j.visres.2016.10.006

Houlsby, N. M. T., Huszár, F., Ghassemi, M. M., Orbán, G., Wolpert, D. M., & Lengyel, M. (2013). Cognitive tomography reveals complex, task-independent mental representations. *Current Biology*, *23*(21), 2169–2175.

Hurlbert, A. C., & Ling, Y. (2007). Biological components of sex differences in color preference. *Current Biology*, *17*(16), R623–R625. https://doi.org/10.1016/j.cub.2007.06.022

Hwang, B.-W., & Schütz, A. C. (2020). Idiosyncratic preferences in transparent motion and binocular rivalry are dissociable. *Journal of Vision*, *20*(12), 3. https://doi.org/10.1167/jov.20.12.3

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An Introduction to Statistical Learning: With Applications in R* (2nd ed.). Springer.

Kosovicheva, A., & Whitney, D. (2017). Stable individual signatures in object localization. *Current Biology*, *27*(14), R700–R701. https://doi.org/10.1016/j.cub.2017.06.001

Laubrock, J., Engbert, R., & Kliegl, R. (2008). Fixational eye movements predict the perceived direction of ambiguous apparent motion. *Journal of Vision*, *8*(14), 13.1-17. https://doi.org/10.1167/8.14.13

Lebovich, L., Darshan, R., Lavi, Y., Hansel, D., & Loewenstein, Y. (2019). Idiosyncratic choice bias naturally emerges from intrinsic stochasticity in neuronal dynamics. *Nature Human Behaviour*, *3*(11), 1190–1202. https://doi.org/10.1038/s41562-019-0682-7

Leopold, D. A., Wilke, M., Maier, A., & Logothetis, N. K. (2002). Stable perception of visually ambiguous patterns. *Nature Neuroscience*, *5*(6), 605–609. https://doi.org/10.1038/nn851

Llera, A., Wolfers, T., Mulders, P., & Beckmann, C. F. (2019). Inter-individual differences in human brain structure and morphology link to variation in demographics and behavior. *ELife*, *8*, e44443. https://doi.org/10.7554/eLife.44443

Mamassian, P., & Wallace, J. M. (2010). Sustained directional biases in motion transparency. *Journal of Vision*, *10*(13), 23. https://doi.org/10.1167/10.13.23

Mardia, K. V., & Jupp, P. E. (2000). *Directional Statistics*. Wiley.

Morikawa, K., & McBeath, M. K. (1992). Lateral motion bias associated with reading direction. *Vision Research*, *32*(6), 1137–1141.

Moutsiana, C., de Haas, B., Papageorgiou, A., van Dijk, J. A., Balraj, A., Greenwood, J. A., & Schwarzkopf, D. S. (2016). Cortical idiosyncrasies predict the perception of object size. *Nature*

Communications, *7*, 12110. https://doi.org/10.1038/ncomms12110
Pastukhov, A., Kastrup, P., Abs, I. F., & Carbon, C.-C. (2019). Switch rates for orthogonally oriented kinetic-depth displays are correlated across observers. *Journal of Vision*, *19*(6), 1. https://doi.org/10.1167/19.6.1
Pressnitzer, D., Graves, J., Chambers, C., de Gardelle, V., & Egré, P. (2018). Auditory Perception: Laurel and Yanny Together at Last. *Current Biology*, *28*(13), R739–R741. https://doi.org/10.1016/j.cub.2018.06.002
Ross, J., Badcock, D. R., & Hayes, A. (2000). Coherent global motion in the absence of coherent velocity signals. *Current Biology*, *10*(11), 679–682. https://doi.org/10.1016/s0960-9822(00)00524-8
Schütz, A. C. (2014). Inter-individual differences in preferred directions of perceptual and motor decisions. *Journal of Vision*, *14*(12), 16.1-17.
Schütz, A. C., & Mamassian, P. (2016). Early, local motion signals generate directional preferences in depth ordering of transparent motion. *Journal of Vision*, *16*(10), 24. https://doi.org/10.1167/16.10.24
Schwarzkopf, D. S., & Rees, G. (2013). Subjective size perception depends on central visual cortical magnification in human v1. *PloS One*, *8*(3), e60550. https://doi.org/10.1371/journal.pone.0060550
Schwarzkopf, D. S., Song, C., & Rees, G. (2011). The surface area of human V1 predicts the subjective experience of object size. *Nature Neuroscience*, *14*(1), 28–30. https://doi.org/10.1038/nn.2706
Shaqiri, A., Roinishvili, M., Grzeczkowski, L., Chkonia, E., Pilz, K., Mohr, C., Brand, A., Kunchulia, M., & Herzog, M. H. (2018). Sex-related differences in vision are heterogeneous. *Scientific Reports*, *8*(1), 7521. https://doi.org/10.1038/s41598-018-25298-8
Shechter, S., Hillman, P., Hochstein, S., & Shapley, R. M. (1991). Gender differences in apparent motion perception. *Perception*, *20*(3), 307–314. https://doi.org/10.1068/p200307
Sorokowski, P., Sorokowska, A., & Witzel, C. (2014). Sex differences in color preferences transcend extreme differences in culture and ecology. *Psychonomic Bulletin & Review*, *21*(5), 1195–1201. https://doi.org/10.3758/s13423-014-0591-8
Sporns, O. (2010). *Networks of the Brain*. MIT Press.
Steinwurzel, C., Animali, S., Cicchini, G. M., Morrone, M. C., & Binda, P. (2020). Using psychophysical performance to predict short-term ocular dominance plasticity in human adults. *Journal of Vision*, *20*(7), 6. https://doi.org/10.1167/jov.20.7.6
Tulver, K. (2019). The factorial structure of individual differences in visual perception. *Consciousness and Cognition*, *73*, 102762. https://doi.org/10.1016/j.concog.2019.102762
Turkheimer, F. E., Rosas, F. E., Dipasquale, O., Martins, D., Fagerholm, E. D., Expert, P., Váša, F., Lord, L.-D., & Leech, R. (2021). A Complex Systems Perspective on Neuroimaging Studies of Behavior and Its Disorders. *The Neuroscientist*, 1073858421994784. https://doi.org/10.1177/1073858421994784
Ullman, S. (1979). *The interpretation of visual motion*. MIT Press.
van Boxtel, J. J. A., Wexler, M., & Droulez, J. (2003). Perception of plane orientation from self-generated and passively observed optic flow. *Journal of Vision*, *3*(5), 318–332.
Wang, Z., Murai, Y., & Whitney, D. (2020). Idiosyncratic perception: A link between acuity, perceived position and apparent size. *Proceedings of the Royal Society B: Biological Sciences*, *287*(1930), 20200825. https://doi.org/10.1098/rspb.2020.0825
Wexler, M. (2018). Multidimensional internal dynamics underlying the perception of motion. *Journal of Vision*, *18*(5), 7. https://doi.org/10.1167/18.5.7
Wexler, M., Duyck, M., & Mamassian, P. (2015). Persistent states in vision break universality and time invariance. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(48), 14990–14995. https://doi.org/10.1073/pnas.1508847112

Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J., & Assad, J. A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. *Nature Neuroscience*, *6*(6), 616–623. https://doi.org/10.1038/nn1055

Witzel, C., O'Regan, J. K., & Hansmann-Roth, S. (2017). The dress and individual differences in the perception of surface properties. *Vision Research*. https://doi.org/10.1016/j.visres.2017.07.015

World Wide Web Consortium. (2019). *W3C Candidate Recommendation*. https://www.w3.org/TR/css-values-3/