Consciousness and Cognition 85 (2020) 103005

Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

Getting to the heart of it: Multi-method exploration of nonconscious prioritization processes

Yaniv Abir[a], Ran R. Hassin[b,*]

[a] Psychology Department, Columbia University, USA
[b] Department of Psychology and The Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel

ARTICLE INFO

Keywords:
Consciousness
Face perception
Continuous flash suppression
Prioritization

ABSTRACT

Understanding the determinants of consciousness is crucial for theories that see it as functionally adaptive, and for explaining how consciousness affects higher-level cognition. The invention of continuous flash suppression (CFS), a long-duration suppression technique, resulted in a proliferation of research into the process of prioritization for consciousness. We developed a new technique, repeated masked suppression (RMS), that facilitates the measurement of long suppression times, but relies on different visual principles. RMS enables a theoretical leap: It allows scientists to examine the central process of prioritization across different suppression methods. In five experiments (n = 282) we collected chronometric RMS and CFS data, finding that the previously reported face inversion effect and the face priority-dimension generalize beyond CFS. Our results validate the use of multi-method designs in the study of prioritization for consciousness. Furthermore, we show how RMS could be used online to reach diverse samples, previously beyond the reach of consciousness science.

1. Introduction

The cognitive sciences have long been engaged with the hard problem of consciousness: making it measurable, tractable, and explainable (Chalmers, 1995). A useful prism into the workings of conscious and non-conscious processing has been to ask what type of information enters our conscious experience in the first place (e.g. Blake & Logothetis, 2002; Dubois & Faivre, 2014; Jiang, Costello, & He, 2007). Since only a narrow subset of mental content is prioritized for conscious experience, prioritization and selection have long been regarded as crucial for the study of consciousness (Baars, 1997; Cohen, Dennett, & Kanwisher, 2016; Mack & Rock, 1998).

Understanding the process of prioritization is crucial for theories that see consciousness as functionally adaptive. It is by the process of prioritization that information is either made available to the computational faculties of consciousness or left for the purview of non-conscious processing. According to the global workspace (Baars, 1997; Dehaene & Naccache, 2001) and information integration theories (IIT; Tononi, Boly, Massimini, & Koch, 2016), this difference is substantial: The combined, emergent processing power of disparate brain regions is bestowed only upon the contents of consciousness (or uniquely characterizes it, according to IIT). In line with these theoretical perspectives, scientists have identified different neuronal and behavioral signatures for content that has crossed the threshold of consciousness in comparison to content that has not (e.g. Dehaene, Changeux, Naccache, & Sergent, 2006; Koch, Massimini, Boly, & Tononi, 2016; Melloni et al., 2007). These differences inform our theories of human cognition and behavior, positing consciousness as an important moderator of information processing.

* Corresponding author.
E-mail address: ran.hassin@mail.huji.ac.il (R.R. Hassin).

https://doi.org/10.1016/j.concog.2020.103005
Received 17 February 2020; Received in revised form 9 July 2020; Accepted 9 August 2020
Available online 22 September 2020
1053-8100/

Furthermore, crossing the threshold of consciousness may have far-reaching real-world repercussions: While nonconscious stop signals cause a slowing of potent responses and sometimes prevent them, conscious signals reliably induce strong inhibition (Van Gaal, Ridderinkhof, Scholte, & Lamme, 2010). This example formalizes the difference between a driver becoming consciously aware of a child crossing the street, and mere non-conscious processing of this pertinent cue. Thus, a robust science of prioritization for consciousness is of great interest.

When Continuous Flash Suppression (CFS; Tsuchiya & Koch, 2005), a visual presentation paradigm that facilitates the measurement of prioritization for consciousness, was invented over a decade ago, a paradigmatic shift in the study of consciousness ensued. Unlike any previous method, with CFS visual stimuli can be presented for prolonged durations below the threshold of awareness. In an influential variant of this technique presentation is maintained until the target stimulus breaks through CFS (bCFS; Jiang et al., 2007) and becomes visible. Among other advances, the ability to prolong and measure the time it takes for a stimulus to reach awareness led to the creation of a robust science of the prioritization for consciousness.

Researchers in this novel field have documented the various low-level visual factors that drive prioritization for consciousness (e.g. Chung & Khuu, 2014; Li & Li, 2015). Furthermore, they have demonstrated how stimuli with semantic meaning that is chronically important, such as social stimuli, valenced stimuli or readable text, enjoy privileged access to awareness (e.g. Abir, Sklar, Dotsch, Todorov, & Hassin, 2018; Chen & Yeh, 2012; Jiang et al., 2007; Yang, Zald, & Blake, 2007; but see Rabagliati, Robertson, & Carmel, 2018). Lastly, the prioritization system was shown to be adaptive on short time scales. It is, for example, reliably biased by the contents of working memory, or by the convergence of multi-modal stimulation (e.g. Gayet, Paffen, & der Stigchel, 2013; Zhou, Jiang, He, & Chen, 2010). The use of bCFS has also given rise to an enlightening discussion about the distinction between low- and high-level vision in consciousness research (Hassin, 2013; Hesselmann & Moors, 2015), about integration in conscious vs. nonconscious vision (Moors, Hesselmann, Wagemans, & van Ee, 2017; Mudrik, Faivre, & Koch, 2014; Sklar, Deouell, & Hassin, 2018), and about robust experimental methodology (e.g. Dubois & Faivre, 2014; Yang, Brascamp, Kang, & Blake, 2014).

But the theoretical and empirical developments achieved with CFS come with a serious qualification, that has heretofore been largely ignored. The vast majority of results concerning prioritization were derived from only one method of masking stimuli for long durations, that is, CFS. Consequentially, the theoretical construct of prioritization for consciousness has mostly been operationalized only as overcoming the suppression elicited by continuous monocular flashes of high contrast stimuli. It has been speculated that such results should generalize beyond the specific suppression mechanisms (e.g. Keysers & Perrett, 2002), but an empirical investigation telling apart results that are idiosyncratic features of bCFS from those that are central characteristics of prioritization is lacking.

In this paper, we begin to address this serious caveat. We developed a new long-duration masking paradigm - breaking repeated masking suppression (bRMS) - and used it to examine the generality of findings from the bCFS literature (Abir et al., 2018; Jiang et al., 2007). In five experiments with 282 participants we collected psychometrically rich bRMS data in conjunction with bCFS data, establishing convergent validity for bRMS and quantifying shared variance between the two paradigms. This shared variance is theoretically crucial, being evidence for common processes that operate across different suppression methods. Thus, for the first time we are able to quantify the workings of a central prioritization mechanism.

On top of the theoretical advances our paradigm facilitates, our new technique has another advantage: It can be easily implemented online, thus expanding our ability to examine prioritization to populations that were historically very difficult to reach (for example, participants of varied age; 20–69 years in this dataset).

The core idea behind bCFS that has made it so successful is to utilize the properties of binocular vision in order to render participants initially unaware of otherwise perfectly visible stimuli (Tsuchiya & Koch, 2005). This is done by presenting target stimuli only to one eye of an observer, while a dynamically changing masking stimulus is presented to the other eye. When the target eventually breaks into awareness, participants are asked to report its location on screen with a speeded response. Differences in participants' response times - or breaking times (BT) - are assumed to be equivalent to differences in prioritization for consciousness (Dubois & Faivre, 2014; Yang et al., 2014).[1]

Assuming isomorphism between bCFS BTs and time to visual awareness has been crucial in facilitating the prolific and internally consistent literature mentioned above. This assumption, however, has not yet been systematically validated using other suppression techniques.

In bRMS, we chose to utilize backward and forward masking (Enns & Di Lollo, 2000) in order to render visual suppression. Thus, rather than separate stimulus and mask onto two different ocular streams as in bCFS, in bRMS they are separated in time. We presented participants with masks interleaved with a target stimulus appearing at a lower contrast level. The masking stimulus is presented for a duration of 67 ms each time, while the target is presented for a duration of only 34 ms (see Fig. 1).[2] A similar repetitive masking technique has been long known to generate stable invisibility of simple line targets using metacontrast masking (The standing wave of invisibility effect; Macknik & Livingstone, 1998; Werner, 1935). A variant of this method has been previously used in order to present more complex stimuli below the threshold of awareness (Kawakami & Yoshida, 2015).

Here, we take a similar approach to bCFS, and present stimuli with RMS long enough for the observer to form a conscious percept of them. As in bCFS, participants indicate the target's location and the main variable of interest is BT. Due to the similar mode of response, bRMS BTs enjoy the same psychometric properties that made bCFS into such a successful paradigm.

[1] Since the contrast of the target stimulus is increased during the first second of presentation, BTs in paradigms of this sort confound the necessary time period and sufficient contrast for conscious awareness. Disentangling these two sources of stimulus energy is beyond the scope of this paper.

[2] The resulting long-term suppression is distinct from RSVP streams, where the target image is presented only once and only briefly (Lawrence, 1971; Raymond, Shapiro, & Arnell, 1992)

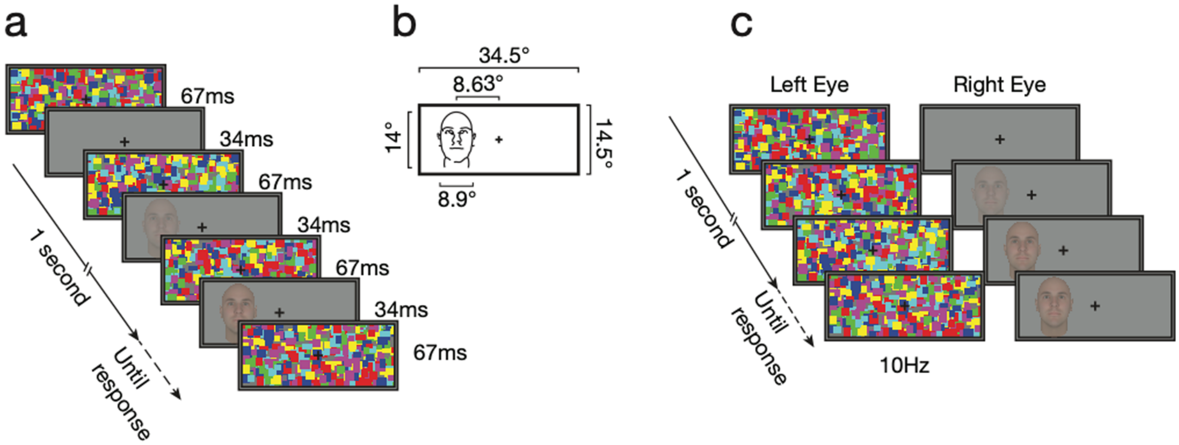

Fig. 1. Stimuli and presentation procedures. In breaking repeated masking suppression (bRMS) high-contrast suppressors are presented before and after each short flash of the target face stimulus. An example trial is schematically depicted in panel (a). Panel (b) describes the presentation sizes for each frame in all lab experiments, and the relative presentation proportions for online experiments. Panel (c) depicts the breaking continuous flash suppression paradigm (bCFS), where suppressor stimuli and target faces are presented to different eyes.

After developing bRMS we validated that it measures prioritization for consciousness by replicating two robust effects described in the bCFS literature – the face inversion effect, and the face priority-dimension. We chose these two phenomena as test cases as they have been demonstrated to be replicable and reliably different from post-conscious biases (Abir et al., 2018; Jiang et al., 2007). By replicating these effects we demonstrated that they are indeed general properties of the prioritization process, rather than specific characteristics of the bCFS technique. Furthermore, we emphasized comparing not only population-level effects, but also the spread of individual differences in BTs, as robust differences are a hallmark of the prioritization for consciousness process (Sklar et al., 2020).

Experiments 1a, 1b and 2 constitute successful replications of the face inversion effect with bRMS. It is well established that bCFS BTs for faces in their cardinal orientation are shorter than for inverted faces (Jiang et al., 2007). This effect is commonly ascribed to the privileged processing afforded to facial stimuli in the visual system. Replicating this effect in bRMS is thus a first step in establishing the convergent validity of bRMS, and a crucial demonstration of the generalizability of a basic pattern in prioritization of faces for visual awareness

Experiment 3a examines the face priority-dimension finding previously reported in bCFS. Following the literature in social trait judgment (Oosterhof & Todorov, 2008), reverse correlation was used to estimate the combination of facial characteristics that best predict an increase in bCFS BTs (Abir et al., 2018). The resultant priority-dimension was found to be correlated with judgments of dominance, such that more dominant faces break earlier into consciousness. The priority-dimension, being a high-dimensional description of BTs, can be used as a strong test of convergence between bCFS and bRMS. Furthermore, it has previously been shown to be a sensitive measure for distinguishing between different cognitive tasks that share psychometric attributes (Abir et al., 2018)[3]. Thus, in Experiment 3b we measured this dimension for bCFS and bRMS within the same participant group, testing for the correlation and convergence between the two paradigms.

2. Experiments 1a and 1b

In experiment 1a we test the effect of face inversion in bRMS. In experiment 1b we replicate the results of experiment 1a and characterize the inversion effect in bRMS relative to bCFS within the same group of participants.

2.1. Methods

The protocol for all experiments was approved by the Institutional Review Board of The Hebrew University. Informed consent was obtained from all participants.

2.1.1. Participants

32 Hebrew University students (12 male, age: M = 24.59, SD = 2.63, range 20–32) participated in experiment 1a in return for course credit (n = 4) or payment of 30 NIS (n = 28). We based our participant exclusion criteria and desired sample size on previous studies (Abir et al., 2018). Furthermore, the original bCFS study reported an inversion effect size of Cohen's d = 0.92 (Jiang et al., 2007), which requires 26 observations to achieve 0.90 power to detect an effect with a *t*-test. We thus aimed for an n = 30 sample after exclusions.

[3] In addition, it has been shown that the face priority-dimension cannot be fully explained by low-level visual differences between face images (Abir et al., 2018; cf. Stein, 2019). The question of low-level vs. high-level vision is orthogonal to the cardinal discussion in this paper.

Another 65 Hebrew University students (20 male, age: M = 23.29, SD = 2.44, range 18–33) participated in experiment1b for course credit (n = 35) or payment of 40 NIS (n = 30). We aimed to have 50 participants after applying exclusion criteria, since we wanted additional power to conduct correlational analyses between bRMS and bCFS data from this experiment. A sample of this size would allow us to detect a positive correlation as strong as r = 0.4 with 0.90 power.

2.1.2. Stimuli and apparatus

Stimuli consisted of multicolour suppressors (“Mondrians” – random amalgams of rectangles of varying sizes and colors) and target stimuli. The target pictures depicted 55 faces taken from a set of 300 faces randomly generated using FaceGen 3.1 (Oosterhof & Todorov, 2008). Faces were presented either in their cardinal orientation or rotated by 180°. In experiment 1a a scrambled face condition was also included for exploratory purposes. We report results pertaining to this condition in the Appendix A. Stimuli were presented within a black frame with a cross in the middle, centered on a grey background. Fig. 1 depicts stimuli and their presented size.

Stimuli were presented on a Samsung SyncMaster SA950 LCD monitor (refresh rate 60 Hz), using PsychToolbox (Brainard, 1997) for MATLAB. Participants sat approximately 60 cm from the screen and used a chin rest. In experiment 1b, participants wore compatible 3d glasses, allowing for separate presentation to each eye.

2.1.3. Procedure

On each trial of the bRMS task Mondrians and target faces were presented in rapid alternation. Mondrians were presented for 66.67 ms and target faces for 33.34 ms each time, creating a 10 Hz rhythm. On each trial of the bCFS task in experiment 1b participants were presented with Mondrians changing at 10 Hz to one eye, and one face stimulus to the other eye.

Common to both paradigms, the contrast of the target faces was ramped up during the first second of presentation to one of six levels of maximum contrast, ranging between 30% and 90% (defined as the rgba alpha channel value). Contrast for the Mondrians decreased to zero starting at 7 s and up to 10 s into the trial.

On each trial participants had to indicate by pressing one of two keys if the face (or shape, in experiment 1a) appeared on the left or right side of the screen. They were instructed to respond as soon as they became aware of the location of the face. This reaction-time, relative to trial onset, serves as our dependent measure of BT. If no response was given, the trial ended after 10 s. Stimuli order and presentation location were randomized.

Participants first read the instructions for the experiment, and then completed 25 practice trials with the experimenter present in the room. The experimenter then left for the remainder of the experimental block. Each face of the 30 faces was presented to participants at each level of maximum contrast, in each of the presentation conditions (Normal/Inverted/Scrambled for experiment 1a, Normal/Inverted for experiment 1b). Breaks were given every 100 trials.

Experiment 1a consisted only of a bRMS block. Experiment 1b consisted of both a bRMS and a bCFS block in counterbalanced order.

2.2. Analysis and results

Analysis was carried out using the R statistical environment 3.6.0 in conjunction with the Stan probabilistic programming language 2.19.2.

2.2.1. Data preparation

The data preparation procedure was based on previous bCFS studies (Abir et al., 2018). Participants who reported shutting one eye during the bCFS block were excluded from analysis (Exp. 1b: n = 12). Mean accuracy of responses for the top three contrast levels was calculated per participant. Participants' whose accuracy scores were more than 3SD from the group mean were excluded (Exp. 1a: n = 1; Exp. 1b: n = 3). Error trials were excluded (Exp. 1a: n = 740, 4.42%; Exp. 1b bRMS: n = 1453, 8.08%; Exp. 1b bCFS: n = 1035, 5.76%), and so were trials with BTs of less than 200 ms (Exp. 1a: n = 2, 0.01%; Exp. 1b bRMS: n = 8, 0.04%; Exp. 1b bCFS: n = 3, 0.02%). Trials with BTs more than 3SD from the mean of each participant in each contrast level and each paradigm were further removed from analysis (Exp. 1a: n = 261, 1.56%; Exp. 1b bRMS: n = 295, 1.64%; Exp. 1b bCFS: n = 333, 1.85%). In experiment 1b, a software error occurred during 25 bRMS trials, which were excluded from analysis (1.85%).

2.2.2. Face inversion effect

Average BTs are consistently longer for inverted faces relative to faces in their cardinal orientation (Fig. 2). This pattern is evident in experiment 1a and in both paradigms of experiment 1b. We corroborated this observation by fitting the data with a lognormal regression model, chosen to capture to the strong skewness of reaction-time data (Ratcliff, 1993). We used the log contrast level, face orientation, presentation paradigm and their interactions to predict BTs. All terms were allowed to vary by participant and stimulus. Recommended priors for reaction-time data were used for facilitating and regularizing the fitting process (Schad, Betancourt, & Vasishth, 2019; Table A.1).

Crucially, we found a robust face inversion effect in experiment 1a b = 0.06 95% PI [0.03, 0.09]. This finding was replicated in experiment 1b for bRMS b = 0.05 95% PI [0.02, 0.08], and for bCFS b = 0.08 95% PI [0.04, 0.11]. As can be seen in Fig. 2, these effects correspond to a 5.13–8.26% rise in mean BTs when a face is inverted, and the 95% posterior interval for both paradigms and experiments did not include zero. The inversion effect did not differ considerably between the two paradigms b = 0.03 95% PI [−0.01, 0.07]. These analyses demonstrate a similar face inversion effect both in bCFS and bRMS.

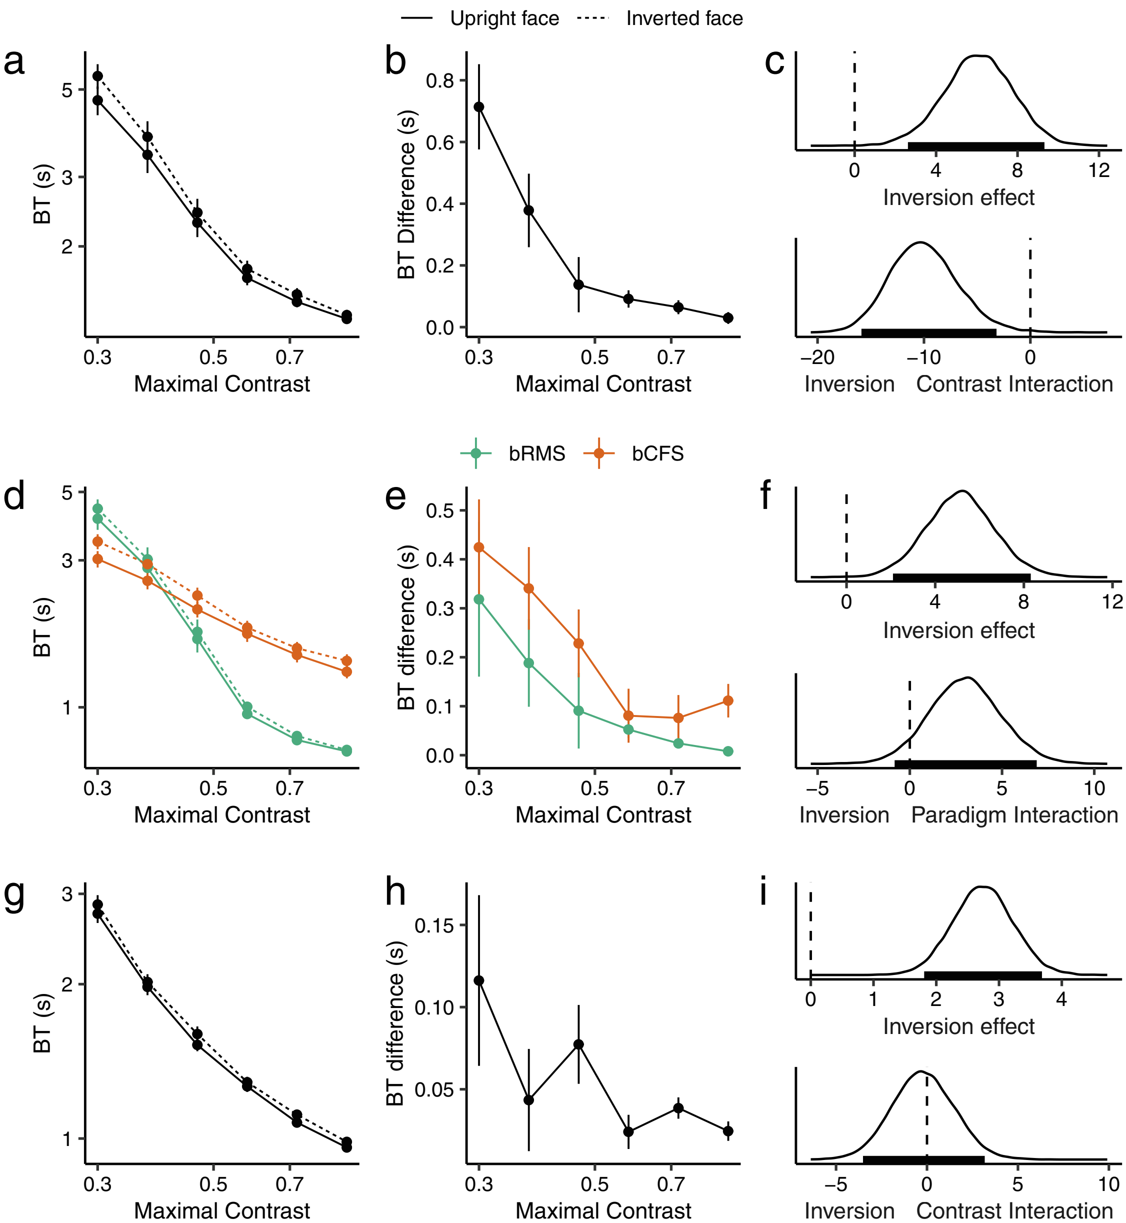

Fig. 2. Results of face inversion experiments. Results from Experiment 1a are presented on the top row. Mean BTs (a) and the difference between BTs for inverted and upright faces (b) are plotted as a function of presentation contrast. Mean BTs and contrast are plotted on a log scale. Error bars denote 1 SE. Posterior distributions of the coefficients from a multilevel lognormal regression are plotted in (c). On the top is the coefficient for the inversion effect, on the bottom the interaction term for face inversion and contrast level. Thick black lines denote the 95% posterior interval. Equivalent results from Experiment 1b, split by paradigm, are plotted in the middle row. The posterior for the overall inversion effect across paradigms is plotted in the top of (f), with the interaction term of face orientation with paradigm plotted on the bottom. Results from the online sample of Experiment 2 are plotted on the bottom row, mirroring the plots from Experiment 1a on the top row.

We also found a significant interaction between face orientation and log contrast in experiment 1a $b = -0.11$ 95% PI $[-0.17, -0.03]$ and a similar, but non-significant, trend in experiment 1b $b = -0.07$ 95% PI $[-0.15, 0.00]$. This effect was not hypothesized, and therefore needs to be cautiously addressed. We suspect it might stem from the deviation from linearity in the relation between log contrast and log-BTs.

2.2.3. Individual differences

Using the same model, we find robust individual differences in both bRMS and bCFS BTs (see Fig. 3). The population of

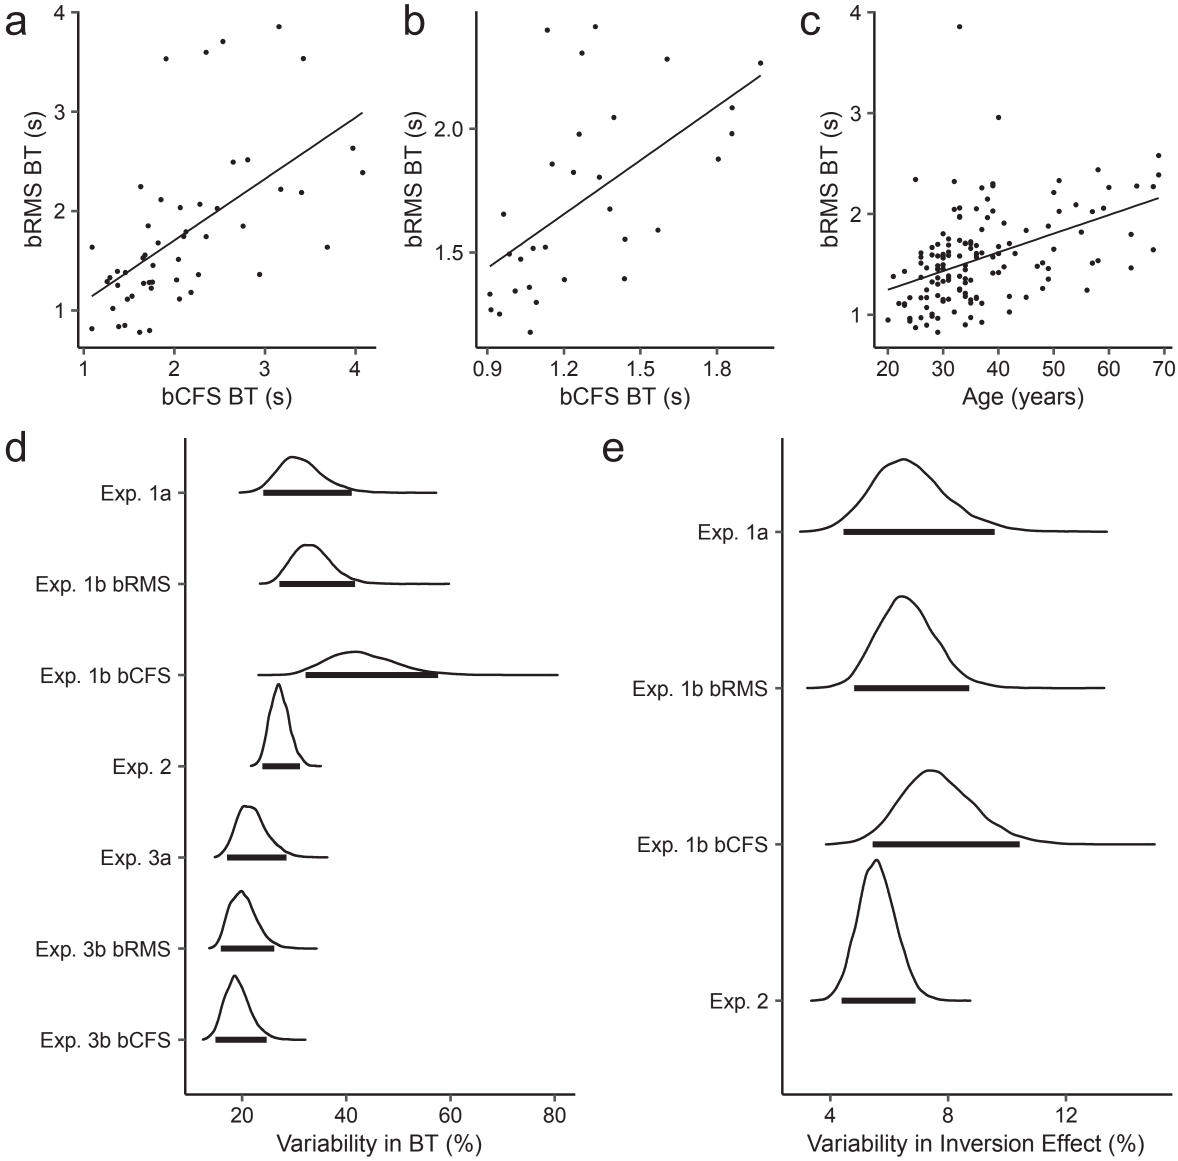

Fig. 3. Correlations and magnitude of individual differences in bRMS and bCFS. Scatter plots of mean BTs for each participant in the two paradigms are plotted in (a) for Experiment 1b, and in (b) for Experiment 3b. The scatter plot of age and bRMS BTs is depicted in (c). Solid lines are best linear fits. The bottom row depicts the posterior distribution of the scale of individual differences in mean BT (d) and the inversion effect (e). These can be read as the expected percentage difference, relative to mean BT, between the average participant and a participant removed from the population mean by 1 SD. Posterior estimates were derived by multilevel lognormal regression.

participants is estimated to have an SD of 30.89–33.17% around the mean BT. Experiment 1b allowed us to directly compare individual differences across paradigms. While the scale of individual differences in bCFS tended to be larger than bRMS (median posterior difference 9.25%, 95% PI [0.80%, 20.60%]), individual average BTs in the two paradigms were strongly correlated $r = 0.66$, 95% PI [0.63, 0.68] Fig. 3. Furthermore, the individual differences in the inversion effect in the two paradigms are also reliably correlated $r = 0.33$ 95% PI [0.02, 0.64]. This correspondence in individual differences is strong evidence for shared processing in bRMS and bCFS.

3. Experiment 2

bRMS uniquely affords the opportunity to perform experiments online, gathering data from diverse samples that have never been part of the consciousness literature (non-WEIRD participants; Henrich, Heine, & Norenzayan, 2010). Experiment 2 demonstrates this approach by documenting the face inversion effect and individual differences in BTs in a varied sample collected through Amazon

Mechanical Turk (MTurk).

3.1. Methods

3.1.1. Participants

242 participants recruited through MTurk completed the task (104 male, age: M = 38.10, SD = 11.59, range 20–69) in return for payment of $4. Data was collected on weekdays between 7AM and 2PM ET. Only participants with an approval rating of at least 95% qualified for the experiment.

Sample size was chosen to be larger than in experiments 1a and 1b since we expected greater variability in an online sample but did not have an estimate of effect size to use for power analysis. We aimed for 120 participants in our analysed sample, and so doubled that number to allow for potentially high exclusion rates.

3.1.2. Stimuli and procedure

Stimuli were the same as for Experiment 1a. The experiment code was ported to JavaScript and the jsPsych package (De Leeuw, 2015), and run online using the psiTurk software (Gureckis et al., 2016). High performance animation for bRMS trials was generated using the GSAP animation package.

The experimental procedure was identical to Experiment 1a with necessary adaptations for online administration that are described in the supplemental material.

3.2. Analysis and results

3.2.1. Data preparation

The data preparation procedure was similar to that used in Experiment 1a, with adaptations for online data. Three participants were excluded based on their questionnaire responses: one reported presentation problems, two participants responded with multiple expletives. Next, trials on which participants were engaged with another application on their computer were excluded from analysis (n = 380, 0.47%). Eighty-two participants experienced multiple problems with timely visual presentation as recorded by the animation code and were excluded from analysis. Specific trials during which presentation problems were detected for the remaining participants were further excluded (n = 405, 0.78%). These high exclusion rates are due to insufficient available computational power on participants computers, either due to dated hardware or concurrently running applications. This disadvantage was partially mitigated in the most current version of the online bRMS code, which tests participants' hardware before initiating the experiment.

Five participants with accuracy scores for the three top contrast levels more than 3SD from the group mean were excluded from analysis. Error trials were excluded (n = 2736, 5.29%), and so were trials with BTs of less than 200 ms (n = 60, 0.12%), trials in which the response was given after the 15 s response deadline (n = 1052, 2.04%), and trials with BTs more than 3SD from the mean of each participant in each contrast level (n = 866, 1.68%). Finally, participants for whom the remaining data set included less than two-thirds of the original trial count were excluded from analysis (n = 3), leaving the data of 142 in the analyzed dataset (71 male, age: M = 36.74, SD = 11.18, range 20–69).

3.2.2. Face inversion effect

Group average BTs were consistently longer for inverted faces relative to faces in their cardinal orientation (Fig. 2g, h). This is corroborated by a lognormal regression model like the one used for Experiment 1a, in which a robust inversion effect is evident b = 0.03, 95% PI [0.02, 0.04]. Using this model, we find no substantial difference in the inversion effect across contrast levels b = 0.00, 95% PI [−0.03, 0.02].

3.2.3. Individual differences

Robust individual differences are evident in this online sample and were estimated as part of fitting the above-mentioned model. The population of participants is estimated to have an SD of 27.21% (95% PI [23.91, 31.10]) around the mean BT.

Having access to a large online sample affords the opportunity to explore the correlates of these robust differences. Our sample presents a wide range of ages, and indeed we find that this variability is correlated with BTs such that older participants present longer BTs $r = 0.45$, $p < .001$ (see Fig. 3c). While this result is not adjusted for general reaction-time variability (cf. Sklar et al., 2020), it demonstrates the research potential of measuring prioritization for consciousness in online samples.

4. Experiments 3a and 3b

In these experiments we tested whether the characteristics of a target face predict bRMS BTs similarly to the way they predict bCFS BTs. We have previously established that bCFS BTs are well predicted by a combination of facial features termed the priority-dimension, that is negatively correlated with perceived facial dominance (Abir et al., 2018). In Experiment 3a, we first identified such a priority-dimension in bRMS alone, and then in Experiment 3b used both paradigms on the same group of participants, allowing us to directly estimate the similarity of the dimensions.

4.1. Methods

4.1.1. Participants

36 Hebrew University students (13 male, age: M = 28.08, SD = 2.24, range 19–29) participated in Experiment 3a for course credit (n = 9), or payment of 30 NIS. Another 42 Hebrew University students (7 male, age: M = 24.07, SD = 3.02, range 20–33) participated in experiment 3b for payment of 40 NIS.

As in our previous work with the priority-dimension (Abir et al., 2018), we aimed to have data from 30 participants after exclusions. We previously found that with n = 14 participants one would have power of 0.95 to detect a priority-dimension with the previously reported magnitude (Abir et al., 2018). Since we aimed for a more elaborate analysis, taking into account uncertainty due to differences between participants, we chose a larger sample size.

4.1.2. Stimuli and procedure

The procedure for Experiment 3a was similar to Experiment 1a, and the procedure for Experiment 3b was similar to Experiment 1b. The important difference in Experiments 3a and 3b is our use of the full stimulus set of 300 randomly generated faces, presented twice for each paradigm. This set of faces were generated by drawing values from a standard normal distribution on each of the 50 parameters defining the shape of each face (Oosterhof & Todorov, 2008). Presenting a large randomly generated stimulus set allows us to compute an unbiased estimate of the combined contribution of the shape parameters to variations in BTs. All faces were presented in their cardinal orientation at the same contrast level of 70%.

4.2. Analysis and results

4.2.1. Data preparation

The data preparation procedure was identical to that used previously for the priority-dimension (Abir et al., 2018). Participants who reported looking only at one side of the screen or shutting one eye were excluded from analysis (Exp. 3a: n = 5, Exp. 3b: n = 7). Participants whose accuracy rate was less than 90% were also excluded from analysis (Exp. 3a: n = 1, Exp. 3b: n = 0). Next, error trials were excluded (Exp. 3a: n = 401, 2.23%; Exp. 3b bRMS: n = 399, 1.11%, bCFS: n = 446, 1.24%), as were trials with BTs shorter than 200 ms (Exp. 3a: n = 0; Exp. 3b bRMS: n = 11, 0.03%, bCFS: n = 1, 0.00%), and trials with BTs more than 3SD from the each participant's mean (Exp. 3a n = 253, 1.41%; Exp. 3b bRMS: n = 313, 0.87%, bCFS: n = 336, 0.93%).

4.2.2. Priority-dimension

We fit data from each paradigm with a multiple regression model, predicting BTs from all 50 shape parameters defining the appearance of a face stimulus. Building on a multitude of previous work interested in modelling the facial characteristics that drive a psychological outcome, we were interested in the overall pattern of regression coefficients, rather than the contribution of each one separately. As in any regression model, the resultant face shape coefficients can be used to generated predicted BTs for each face stimulus. Furthermore, changing a face stimulus along this vector of coefficients, which we termed the priority-dimension, is equivalent to changing it in the direction that causes the largest changes in BTs, and thus prioritization for consciousness.

As in Experiments 1a, 1b and 2, we used a lognormal likelihood for the multiple regression model. The intercept and the 50 face-shape coefficients were allowed to vary by participant (see Fig. A.2). For Experiment 3b we used a multivariate model: regressions were fit simultaneously for both bRMS and bCFS. The correlation between the 50 face-shape coefficients for bRMS and bCFS was directly estimated in the model. Additionally, this model included a term for the correlation in individual differences in general BTs between the two paradigms. The same priors used for Experiments 1a, 1b and 2 were utilized in these models. See the Appendix A for a technical description of the model, and validation with posterior-predictive checks.

The face-shape regression model explained a consistent proportion of BT variability. As a measure of effect size, we compare the predicted BT for the average face stimulus, with the predicted BT for a face removed in face-space from the average by 1 SD along the priority-dimension (Todorov, Dotsch, Wigboldus, & Said, 2011). Based on the data from Experiment 3a, bRMS BTs are expected to change by 2.39% (95% PI [1.66%, 3.18%]) for such a change in the presented face. In Experiment 3b, bRMS BTs are expected to change by 3.61% (95% PI [2.70%, 4.58%]). This is in line with the change observed in bCFS data from Experiment 3b (2.17%, 95% PI [1.16%, 3.13%]) and for the originally published priority dimension (2.95%, 95% PI [2.35%, 3.60%]; see Fig. 5).

Importantly, we compared the face-shape coefficients for bRMS and bCFS in experiment 3b and found a strong correlation between the bRMS and bCFS priority-dimensions in this participant sample: median posterior r = 0.62, 95% PI [0.17, 0.88].

4.2.2.1. Correlations with previously published data. The correlation with the original priority-dimension (Abir et al., 2018) is strong for the face-shape coefficients predicting bRMS data in experiment 3a (median posterior r = 0.50, 95% PI [0.29, 0.66]) and experiment 3b (median posterior r = 0.45, 95% PI [0.25, 0.61]).

Fig. 4 depicts the full correlation structure between the priority-dimensions reported here for bRMS and bCFS, while Fig. 5 depicts the correlation of these dimensions with the face-shape coefficients that predict judgments of trustworthiness and dominance, as measured by Oosterhof and Todorov (2008). Oosterhof and Todorov showed in their seminal paper that the space of facial social trait judgments reduces well to judgments of trustworthiness and dominance.

4.2.3. Individual differences

A strong correlation between average bRMS and bCFS BTs across participants is replicated in Experiment 3b median r = 0.54,

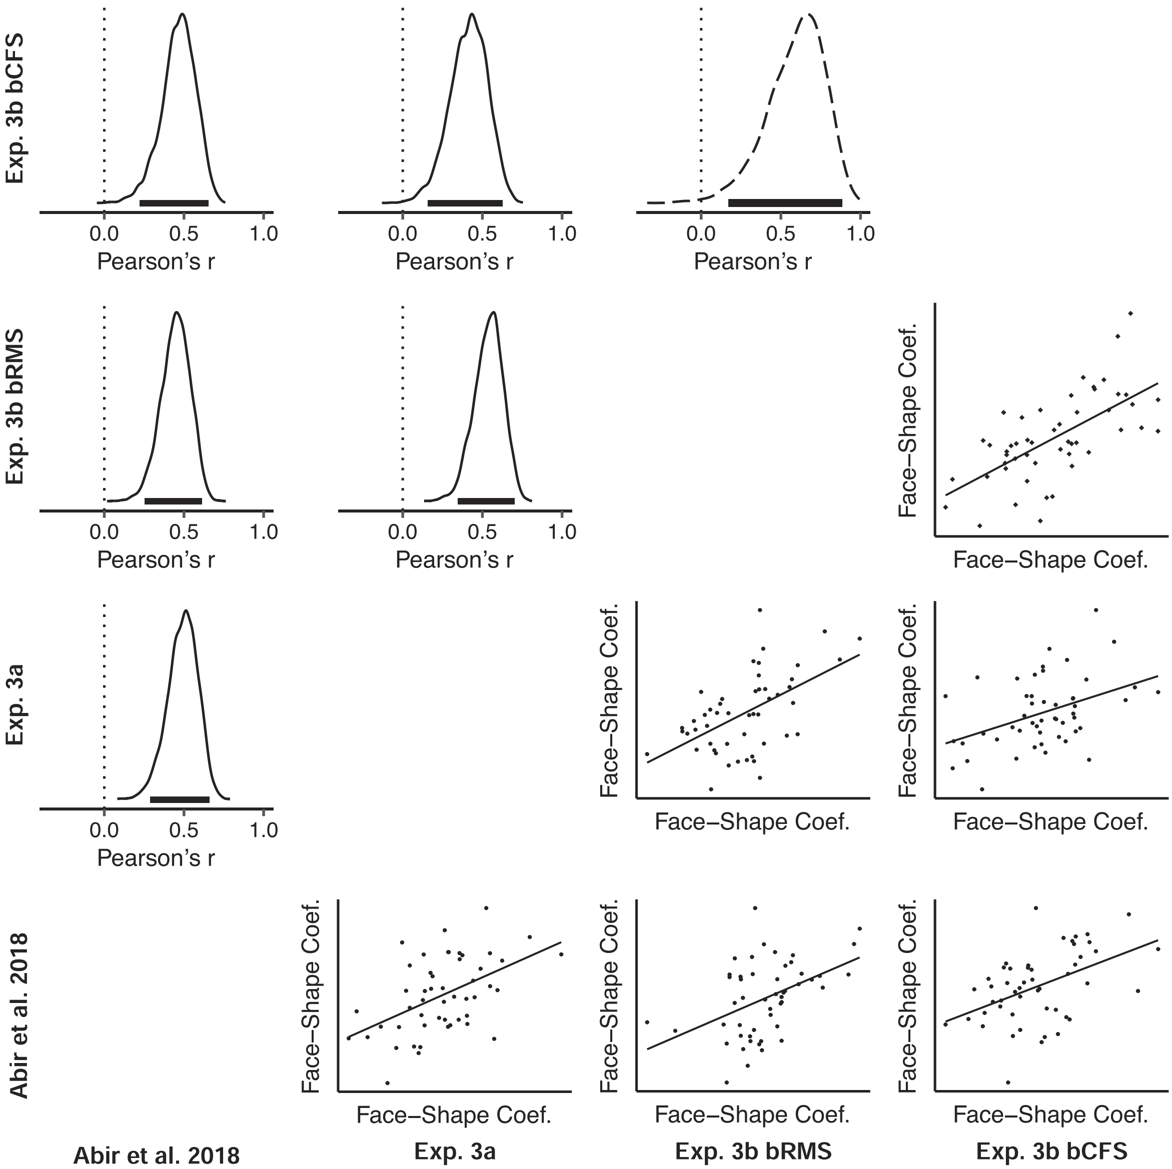

Fig. 4. Strong positive correlations between face-shape coefficients in Experiments 3a, 3b and data previously published in (Abir et al., 2018). Bivariate correlations and scatter plots are organized in a matrix, where each row and each column correspond to a different dataset, labeled in bold. Scatter plots depicting the correspondence between face-shape coefficients in each of the experiments are presented below the diagonal. Lines are derived by linear regression for plotting purposes. Above the diagonal are posterior distributions for the correlation between face-shape coefficients in the two datasets. Thick lines denote 95% posterior intervals. Correlations between bRMS and bCFS BTs measured within the same participant sample are depicted with diamonds and a dashed line, correlations for BTs measured in two different samples in circles and a solid line.

95% CI [0.25, 0.73] Fig. 3b.

5. Discussion

We presented bRMS, a new paradigm for measuring prioritization for consciousness, and demonstrated its convergence with bCFS, the established method of measurement in the field. We examined both the face inversion effect and the face priority-dimension in bRMS in a series of experiments using robust experimental and statistical methodology. Our results support the attribution of these effects to differences in prioritization for consciousness, rather than merely differences in binocular suppression.

Allowing for the successful and simple collection of data online, bRMS promises to greatly expand the volume of data collected in this field of research, and the generalizability of inferences made across varied populations and variable suppression methods.

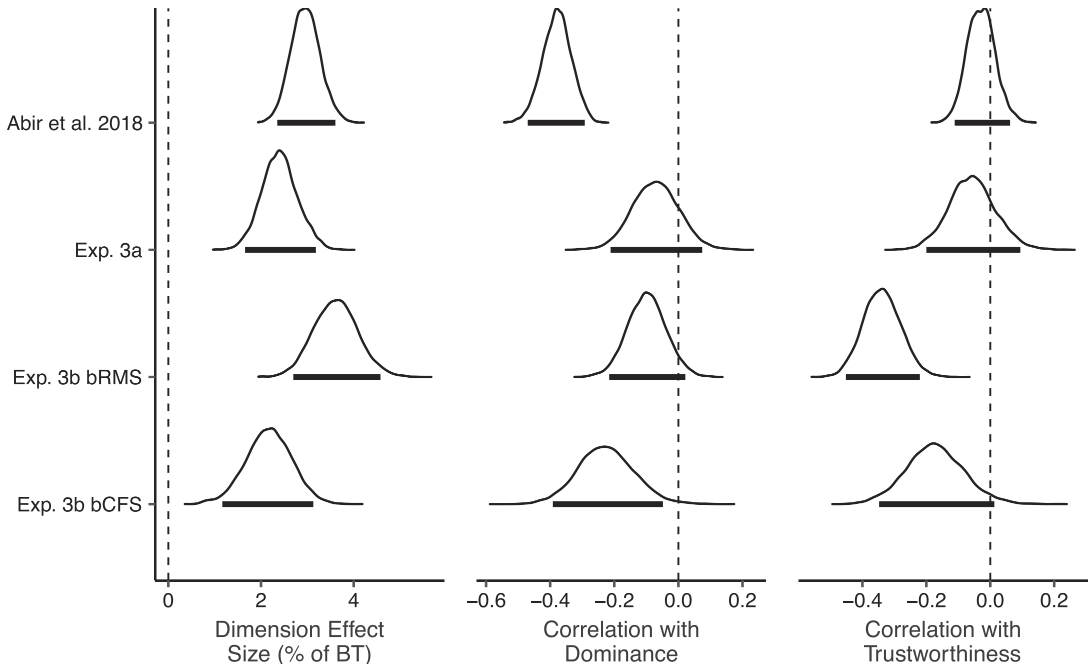

Fig. 5. Face dimensions in bRMS and bCFS. Effect sizes for the face dimensions in Experiments 3a and 3b, as well as for data previously published in Abir et al. (2018), are depicted on the left. These can be read as the expected percentage difference, relative to mean BT, between the average face and a face removed from the average face by 1 SD along the priority-dimension. Posterior distributions for the correlation between the BT face dimensions and social trait dimensions reported by Oosterhof and Todorov (2008) are shown in the middle column (for dominance) and on the right (for trustworthiness). Thick lines denote the 95% posterior interval.

Tapping online participant pools supports research into questions as of yet unexplored, such as the covariance of individual differences in BTs with various cultural and biological attributes. bRMS, then, opens new horizons for the study of consciousness.

Along with the advantages of running prioritization experiments online, our experiments also demonstrate some of the challenges. We find shorter BTs online relative to those measured in lab, and a smaller inversion effect size. These can be attributed to a considerable source of additional variance in online experiments relative to lab experiments: variance in presentation hardware and environmental conditions. It is important to remember, though, that these differences may also reflect the reality of taking our effects outside of our laboratories and into the homes of subjects, thereby reaching a much more diverse population.

Our intention was to develop and test a tool for investigating long-duration prioritization and allow for inferences that go beyond one paradigm. Yet, the comparison with previously published findings is also interesting. First, regarding the priority-dimension, the negative correlation of the priority-dimension with facial dominance reported here and previously, such that more dominant faces are more highly prioritized for consciousness, does not agree with pervious findings reported with different task designs. Stewart et al. (2012) used a small set of computer-generated faces that varied only on the dominance and trustworthiness social trait dimensions. They find longer bCFS BTs for more dominant faces. It remains an open question whether this discrepancy is due to the restricted sampling of stimuli, which result in distorted inference across the space of possible faces, to the difference in context induced by the two different tasks, engendering differing cognitive responses, or to the difference in participant populations – American college students in Stewarts and collogues' data, and Israeli and Palestinian university students in our data.

Previous work suggests that different methods of suppressing stimuli from consciousness impact different levels of processing in the cognitive system (e.g. Breitmeyer, 2015), and might interact with attention in differing ways (Breitmeyer & Ogmen, 2000; Tong, Meng, & Blake, 2006). In spite of these differences in suppression mechanism, our studies revealed similar prioritization processes operating under both paradigms (Keysers & Perrett, 2002). We did not find systematic differences in the face inversion and priority-dimension effect between bRMS and bCFS. A trend towards weaker effects in bRMS relative to bCFS appears in some experiments and comparisons, but not in others. Further research, perhaps employing simpler stimuli, is needed in order to find the differences between the two paradigms that may exist on top of the great similarity described here. The existence of similarities and differences would enrich our theory of prioritization, as well as our understanding of specific ways in which stimuli become conscious.

As test cases for this new multi-method design, we chose two effects concerning face stimuli. We knew these effects to be robust, from our own work and from the work of many others, and so were good necessary tests of convergence between the paradigms. Moving forward, it will be interesting to revisit the many factors reported in the literature as predictors of prioritization. With bRMS and bCFS, it should be possible to reevaluate the replicability and generalizability of this literature.

Our results are a mere demonstration of the potential promise of running bRMS experiments online. An interesting extension to the age effect reported here would be to measure individual differences in both BTs and general reaction-times simultaneously (Sklar et al., 2020). Independently, being able to characterize the face priority-dimension in different cultures (cf. Jack, Garrod, Yu, Caldara, & Schyns, 2012) opens an exciting research avenue.

For these reasons, we make our experimental code for running bRMS experiments in the lab and online available to the scientific community. We hope these will be useful tools for researchers of consciousness.

5.1. Conclusions

Research into prioritization for consciousness has made great advancements in the past decade but has consistently focused on one operationalization of prioritization: bCFS. Using bCFS and bRMS in tandem, we uncover a central prioritization process, that is sensitive to facial features in a similar manner regardless of suppression method. bRMS shows further promise as it allows for prioritization experiments to be conducted in online diverse samples, considerably expanding the possibilities of prioritization for consciousness research.

CRediT authorship contribution statement

Yaniv Abir: Conceptualization, Methodology, Formal analysis, Investigation. **Ran R. Hassin:** Conceptualization, Methodology, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank Nitai Kerem, Ella Margolin and Yoav Resnik for their help with data collection.

Data statement

All data and code are available at https://doi.org/10.24433/CO.2463446.v1. Materials and scripts for running the experiments are available at https://git.io/JeFfi.

Appendix A

A.1. Model fitting

All models in this paper were implemented in the Stan probabilistic programming language. For Experiments 1a, 1b and 2, the brms R package was used to interface with Stan. For Experiments 3a and 3b we used the rstan package. In all cases we ran 5–15 chains of 1000–1200 samples each, collecting a total of 7500–12,000 samples post warmup (according to available hardware). These values were titrated up where necessary to ensure $\widehat{R}$ metrics close to one, and large effective sample size metrics.

For all models, we used priors recommended by Schad et al. (2019) for reaction-time data (Table A.1). These priors enabled and sped-up the process of model fitting. In frequentist terms applying the priors on the effects of interest is equivalent to fitting a ridge-regression model (Gelman et al., 2013).

Table A.1
Choice of priors for regularized lognormal regression models. Where the recommendation by Schad et al. (2019) was for using informative priors, the comments column describes how these were chosen.

Priors for reaction time data		
Type of coefficient	Prior	Comment
Intercept (grand mean)	Normal(7,0.6)	Reflects our expectations for mean RTs between 600 and 2000 ms, which we based on prior studies (Abir et al., 2018) and on piloting. This informed choice for the prior speeds up model-fitting considerably but does not change estimates for any of the effects of interest.
Effects of predictors	Normal(0, 0.05)	
Scale of individual differences	Normal(0, 0.1)	
Scale of residual error	Normal(0, 0.5)	

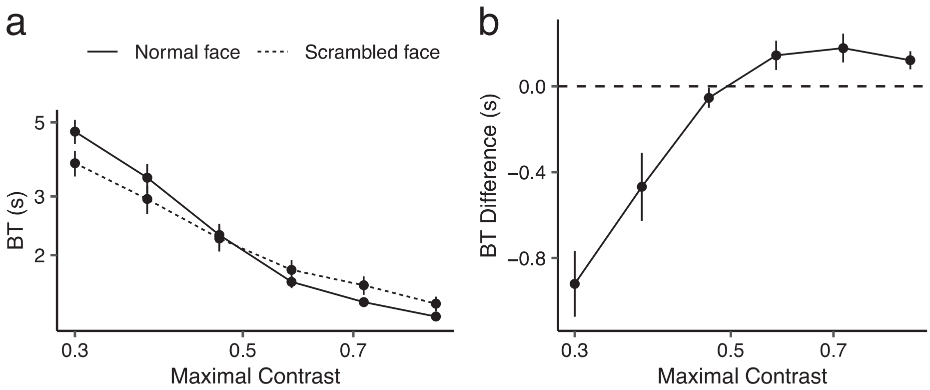

Fig. A.1. Supplementary results from Exp. 1a. Mean BTs (a) and the difference between BTs for scrambled and normally presented faces (b) are plotted as a function of presentation contrast. Mean BTs and contrast are plotted on a logarithmic scale. Error bars denote 1 SE.

A.2. Experiment 1a

Fig. A.1 depicts BTs for diffeomorphically scrambled (Stojanoski & Cusack, 2014) face images relative to normally presented images. Diffeomorphically scrambled faces have previously been used as a control condition in bCFS studies, as they preserve many of the low-level visual features of the face image, but inhibit its identifiability as a face (Abir et al., 2018). To our knowledge, the comparison of BTs for normal and scrambled faces has not been documented in the bCFS literature. We find that the difference in BT is strongly modulated by presentation contrast, underlining the importance of sampling along the contrast continuum when comparing BTs for different class of stimuli in prioritization experiments (Holland, 1986; Waskom, Okazawa, & Kiani, 2019).

A.3. Experiment 2

Before starting the experiment, participants were asked to adjust the size of a circle displayed on screen, so that it matches the size of a US coin of their choice held against the screen. This procedure yielded an estimate of the physical resolution of each participant's monitor, used to adjust display size of all stimuli so that their physical size was uniform across participants. In order to ensure participants' full engagement with the task, no time limit was set on trials, with stimuli remaining on screen until response. Fifteen seconds into the trial the presentation contrast for Mondrians was linearly decreased to 0% over the course of 5 s, rendering the stimulus perfectly visible. The longer possible duration of the trial was chosen to discourage participants from not responding.

Additionally, participants' performance was continuously monitored. Participants who did not achieve 85% accuracy on the training block, were asked to repeat it. Failing to achieve criterion accuracy on the second attempt resulted in the termination of the experiment. During the experimental block, if accuracy during the last 20 trials of the top three contrast levels dipped below 80%, a message was displayed asking participants to respond more accurately. This message was displayed also if five out of the previous eight trials had reaction times of less than 250 ms.

A.4. Experiments 3a and 3b

The multiple regression model predicting BTs from face-shape parameters can be described by the following formula in R syntax:

$$\log(BT)\ 1 + fs_1 + fs_2 + \cdots + fs_{50} + (1 + fs_1 + \cdots + fs_{50}|Participant)$$

where each fs_i is one face-shape parameter. For fitting this regression formula we utilized the same regularizing priors described in Table A.1.

After fitting the multiple regression model to the data, we performed several posterior predictive checks (Gelman et al., 2013) as a validation of the adequacy of this model as a descriptor of the data. We generated simulated BTs using multiple draws from the posterior of the model and compared summary statistics of these simulated data with the summary statistics of the observed data. Fig. A.2 depicts BT histograms for a random sample of participants, along with the distribution of histograms generated from model draws. Similarly to previous formulations of lognormal models for reaction time data (Schad et al., 2019) the model seems to capture most of the features of the histograms.

In previous work of similar vein, correlations between face dimensions were evaluated by averaging the data across participants and using Pearson correlation to evaluate relationships between different datasets, ignoring the hierarchical structure of the data and the non-normal likelihood of reaction time data (e.g. Abir et al., 2018). To compare these methods' performance on our data, we computed the above metrics on the simulated BTs from our model. One thousand simulated BT datasets were generated from the posterior draw that had a correlation parameter for bRMS and bCFS dimensions closest to the median posterior value. The correlation between bRMS and bCFS average BTs for each of the 300 faces, and the correlation between bRMS and bCFS coefficients derived by classical reverse correlation from these average BTs was computed for each of the simulated datasets. These distributions are depicted

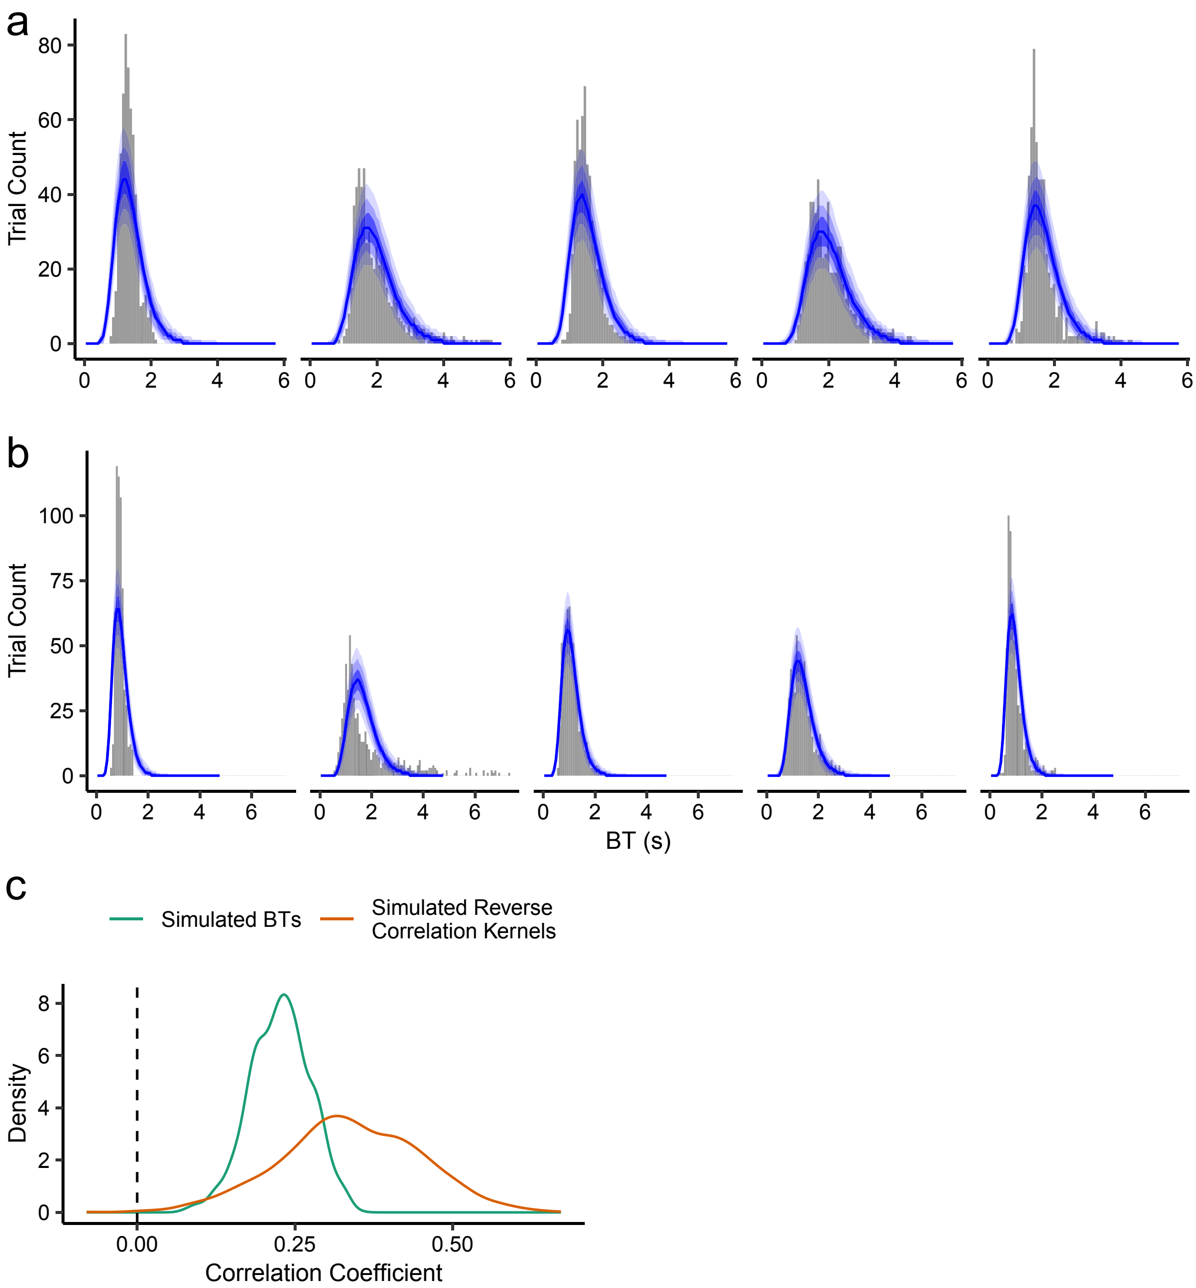

Fig. A.2. Posterior predictive checks for the face-shape regression model. BT histograms for five randomly picked participants are shown in (a) for bRMS data, and (b) for bCFS data. Overlaid in blue are the median histograms for simulated data for these same participants. Ribbons denote 50%, 80% and 95% posterior intervals (from darkest to lightest ribbon). Plot (c) depicts the posterior predictive distribution of correlations between mean bRMS and bCFS BTs for the 300 face stimuli (in green) and of correlations between reverse correlation kernel coefficients for bRMS and bCFS (in orange). These simulated datasets were generated from a posterior draw with median correlation between the latent dimensions for the t wo paradigms. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

in Fig. A.2c. As is evident these distributions are wider and centered at a lower value than our estimate of the correlation reported in the main text. This demonstrates the increased power to detect a correlation between the face-shape coefficients of the two paradigms afforded by the hierarchical regression model in comparison to a simple reverse-correlation approach.

Lastly, we report the effect size for the face-shape multiple regression in the main text. To derive it, we calculated the predicted BT for the average face stimulus (for which all face-shape parameters are zero) and for a face removed from it by 1 SD in face-space along the priority-dimension. Since the units of the face-shape parameters are face-space SDs, moving 1 SD from the average face in face space is equivalent to moving along a normalized vector. Hence, if b is the vector of face-shape coefficients found by multiple regression, we need to predict BT for a face defined by the vector $\hat{x} = \frac{b}{||b||}$:

$$\log(BT) = a + b\hat{x} + e = a + b\cdot\frac{b}{||b||} + e = a + \frac{||b||^2}{||b||} + e = a + ||b|| + e$$

where a is a general intercept, and e is an error term. We find that the expected difference in log BT is the norm of the coefficient vector b. That is the quantity we report for effect size.

Appendix B. Supplementary material

Supplementary data to this article can be found online at https://doi.org/10.1016/j.concog.2020.103005.

References

Abir, Y., Sklar, A. Y., Dotsch, R., Todorov, A., & Hassin, R. R. (2018). The determinants of consciousness of human faces. *Nature Human Behaviour, 2*(3), 194–199. https://doi.org/10.1038/s41562-017-0266-3.

Baars, B. J. (1997). In the theatre of consciousness. Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies, 4*(4), 292–309.

Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews Neuroscience, 3*(1), 13–21. https://doi.org/10.1038/nrn701.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.

Breitmeyer, B. G. (2015). Psychophysical "blinding" methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition, 35*, 234–250.

Breitmeyer, B. G., & Ogmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics, 62*(8), 1572–1595.

Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200–219.

Chen, Y.-C., & Yeh, S.-L. (2012). Look into my eyes and I will see you: Unconscious processing of human gaze. *Consciousness and Cognition, 21*(4), 1703–1710. https://doi.org/10.1016/j.concog.2012.10.001.

Chung, C., & Khuu, S. K. (2014). The processing of coherent global form and motion patterns without visual awareness. *Frontiers in Psychology, 5*, 195.

Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences, 20*(5), 324–335. https://doi.org/10.1016/j.tics.2016.03.006.

De Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods, 47*(1), 1–12.

Dehaene, S., Changeux, J., Naccache, L., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: A testable taxonomy. *Trends in Cognitive Sciences, 10*(5), 204–211.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*(1–2), 1–37.

Dubois, J., & Faivre, N. (2014). Invisible, but how? The depth of unconscious processing as inferred from different suppression techniques. *Frontiers in Psychology, 5*, 1117.

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences, 4*(9), 345–352.

Gayet, S., Paffen, C. L. E., & der Stigchel, S. (2013). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science, 24*(12), 2472–2480.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis.* Chapman and Hall/CRC.

Gureckis, T. M., Martin, J., McDonnell, J., Rich, A. S., Markant, D., Coenen, A., ... Chan, P. (2016). psiTurk: An open-source framework for conducting replicable behavioral experiments online. *Behavior Research Methods, 48*(3), 829–842.

Hassin, R. R. (2013). Yes it can: On the functional abilities of the human unconscious. *Perspectives on Psychological Science, 8*(2), 195–207. https://doi.org/10.1177/1745691612460684.

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature, 466*(7302), 29.

Hesselmann, G., & Moors, P. (2015). Definitely maybe: Can unconscious processes perform the same functions as conscious processes? *Frontiers in Psychology, 6*, 584.

Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association, 81*(396), 945–960.

Jack, R. E., Garrod, O. G. B., Yu, H., Caldara, R., & Schyns, P. G. (2012). Facial expressions of emotion are not culturally universal. *Proceedings of the National Academy of Sciences of the United States of America, 109*(19), 7241–7244. https://doi.org/10.1073/pnas.1200155109.

Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science, 18*(4), 349–355. https://doi.org/10.1111/j.1467-9280.2007.01902.x.

Kawakami, N., & Yoshida, F. (2015). Perceiving a story outside of conscious awareness: When we infer narrative attributes from subliminal sequential stimuli. *Consciousness and Cognition, 33*, 53–66.

Keysers, C., & Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences, 6*(3), 120–125.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience, 17*(5), 307.

Lawrence, D. H. (1971). Two studies of visual search for word targets with controlled rates of presentation. *Perception & Psychophysics, 10*(2), 85–89.

Li, Y., & Li, S. (2015). Contour integration, attentional cuing, and conscious awareness: An investigation on the processing of collinear and orthogonal contours. *Journal of Vision, 15*(16), 10.

Mack, A., & Rock, I. (1998). *Inattentional blindness*, *Vol. 33*. Cambridge, MA: MIT Press.

Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nature Neuroscience, 1*(2), 144–149.

Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., & Rodriguez, E. (2007). Synchronization of neural activity across cortical areas correlates with conscious perception. *Journal of Neuroscience, 27*(11), 2858–2865.

Moors, P., Hesselmann, G., Wagemans, J., & van Ee, R. (2017). Continuous flash suppression: Stimulus fractionation rather than integration. *Trends in Cognitive Sciences, 21*(10), 719–721.

Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences, 18*(9), 488–496.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*(32), 11087–11092.

Rabagliati, H., Robertson, A., & Carmel, D. (2018). The importance of awareness for understanding language. *Journal of Experimental Psychology: General, 147*(2), 190–208. https://doi.org/10.1037/xge0000348.

Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin, 114*(3), 510.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18*(3), 849.

Schad, D. J., Betancourt, M., & Vasishth, S. (2019). Toward a principled Bayesian workflow in cognitive science. ArXiv Preprint ArXiv:1904.12765.

Sklar, A. Y., Deouell, L. Y., & Hassin, R. R. (2018). Integration despite fractionation: Continuous flash suppression. *Trends in Cognitive Sciences, 22*(11), 956–957.

Sklar, A. Y., Goldstein, A. Y., Abir, Y., Dotsch, R., Todorov, A., & Hassin, R. R. (2020). Did you see it? Robust individual variance in the prioritization of contents to conscious awareness. PsyArXiv, January 20. https://doi.org/10.31234/osf.io/hp7we.

Stein, T. (2019). The breaking continuous flash suppression paradigm: Review, evaluation, and outlook. In G. Hesselmann (Ed.). *Transitions between Consciousness and*

Unconsciousness. London: Routledge.
Stewart, L. H., Ajina, S., Getov, S., Bahrami, B., Todorov, A., & Rees, G. (2012). Unconscious evaluation of faces on social dimensions. *Journal of Experimental Psychology: General*. https://doi.org/10.1037/a0027950.
Stojanoski, B., & Cusack, R. (2014). Time to wave good-bye to phase scrambling: Creating controlled scrambled images using diffeomorphic transformations. *Journal of Vision, 14*(12), 6.
Todorov, A., Dotsch, R., Wigboldus, D. H. J., & Said, C. P. (2011). Data-driven methods for modeling social perception. *Social and Personality Psychology Compass, 5*(10), 775–791.
Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences, 10*(11), 502–511. https://doi.org/10.1016/j.tics.2006.09.003.
Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: From consciousness to its physical substrate. *Nature Reviews Neuroscience, 17*(7), 450.
Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience, 8*(8), 1096–1101.
Van Gaal, S., Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. A. F. (2010). Unconscious activation of the prefrontal no-go network. *Journal of Neuroscience, 30*(11), 4143–4150.
Waskom, M. L., Okazawa, G., & Kiani, R. (2019). Designing and interpreting psychophysical investigations of cognition. *Neuron, 104*(1), 100–112.
Werner, H. (1935). Studies on contour: I. Qualitative analyses. *The American Journal of Psychology, 47*(1), 40–64.
Yang, E., Brascamp, J., Kang, M.-S., & Blake, R. (2014). On the use of continuous flash suppression for the study of visual processing outside of awareness. *Frontiers in Psychology, 5*, 724.
Yang, E., Zald, D. H., & Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion (Washington, D.C.), 7*(4), 882–886. https://doi.org/10.1037/1528-3542.7.4.882.
Zhou, W., Jiang, Y., He, S., & Chen, D. (2010). Olfaction modulates visual perception in binocular rivalry. *Current Biology, 20*(15), 1356–1358.