
Attending to attention:
Reverse correlation reveals subtle cues to attentiveness in others' faces

Clara Colombatto¹ (0000-0003-3293-8741)
Brian J. Scholl² (0000-0003-3610-0890)

¹Department of Psychology, University of Waterloo

²Department of Psychology and Wu-Tsai Institute, Yale University

Running Head : Attending to attention
Address for Correspondence : Clara Colombatto
Department of Psychology
University of Waterloo
200 University Avenue West
Waterloo, ON, Canada N2L 3G1
Email : clara.colombatto@uwaterloo.ca
Keywords : Social Perception; Face Perception; Reverse Correlation;
Attentiveness
Preregistration : https://aspredicted.org/K8Q_2M1
Data & Materials : https://osf.io/2pt7y/?view_only=1cd6095a5fef43faaee5807a6eedf6b
Word Count : 2090 (Introduction, Results, Discussion) + 1457 (Methods)
Version : 6/29/25 — In press, *Psychonomic Bulletin & Review*

Abstract (221 words)

Some of the most foundational properties we can perceive from others' faces involve cognitive states, such as how *attentive* (vs. distracted) they seem — an important ability, since the likelihood of someone in our local environment affecting our fitness is enhanced when they are attentive. But how can we tell whether another person is attentive? This study reveals that the way in which we perceive attentiveness in others' faces is straightforward in some ways, but deeply counterintuitive in others. We explored this using reverse correlation, a data-driven approach that can reveal the nature of internal representations without prior assumptions. In two online studies (n=200 each), observers viewed pairs of faces created by adding randomly generated noise (across many spatial frequencies) to a constant base face, and had to select which appeared to be most attentive. Analyses of automatically extracted facial landmarks from the resulting 'classification images' revealed the determinants of perceived attentiveness. Some cues were straightforward: attentive faces indeed had more direct eye gaze, and larger pupils. But other novel and equally robust cues were subtle and surprising; for example, attentive faces reliably had darker (as if more flared, or retroussé) nostrils. These powerful and subtle effects of facial cues on impressions of attentiveness highlight the importance of attention not just as a perceptual process, but as an object of perception itself.

Introduction

When looking at other people's faces, we reflexively perceive a variety of social properties — from stable traits (such as trustworthiness or extraversion) to more transient states (such as surprise or anger). But while a great deal of research has explored such properties (Todorov, 2017), there may be an even more foundational set of features that we readily perceive from others' faces: cognitive states, such as how *attentive* people seem. The ability to perceive others' attentiveness seems especially important since someone who is actively attending to us may be more likely to engage in actions which affect us in the coming moments — more so than someone who is distracted or otherwise inattentive.

Some cues to others' attention seem straightforward, such as eye gaze — since people typically look in the direction in which they are attending. For example, when choosing among multiple options, people tend to look towards the objects they desire most (Shimojo et al., 2003); when making a sandwich, people look towards the ingredients they will grab next (Land & Hayhoe, 2001); and when walking along a trail, people look where they will soon step (Matthis et al., 2018).

But there is also clearly more to attentiveness than just gaze direction: in some situations we may look in one direction while attending elsewhere (Colombatto, Chen, & Scholl, 2020), or while being lost in thought altogether (Smallwood & Schooler, 2013). In fact, experience-sampling studies have shown how people spend almost half of their waking life lost in thought (Killingsworth & Gilbert, 2010; Seli et al., 2018). In these cases, the direction of gaze may not correspond to any underlying object of attention, and is thus relatively uninformative with respect to others' underlying mental states. And even when we are in an overall attentive state, we may nevertheless not be attending to the same point at which we are gazing: we often fail to notice objects that are right in front of us, as in the phenomenon of inattentional blindness — which is actually (and perhaps counterintuitively) greatest for stimuli that appear directly at

fixation (Mack & Rock, 1998). So while research in social perception has typically focused on the *direction* of attention (as signaled by overt gaze), it also seems important for us to perceive the *degree* to which people are attentive in the first place.

But just how can we tell whether another person is attentive (vs. distracted)? Here we employ data-driven techniques to explore internal representations of attentiveness. In particular, we employed reverse correlation — a technique that can reveal internal representations without the need to have prior assumptions or to test specific hypotheses (Dotsch & Todorov, 2012; Gosselin & Schyns, 2003; Mangini & Biederman, 2004). In such studies, observers typically view images created by superimposing random noise patterns (with independent random variation across many spatial frequencies) onto faces, and must then select which of two images best captures a particular trait (e.g. looking more trustworthy). After repeating such judgments many times (always with different random noise patterns, but the same underlying faces), the selected images can then be averaged to create ‘classification images’ that represent the visual features associated with trait or state of interest. In this way, reverse correlation can recover observers’ internal visual representations, and can identify specific subtle properties that observers associate with the trait in question. For example, past work employing reverse correlation has identified specific visual cues underlying the perception of demographic characteristics (e.g., ethnicity; Dotsch et al. 2008), social traits (e.g., trustworthiness; Dotsch & Todorov, 2012), emotions (e.g., surprise; Jack et al., 2012), and even aspects of our self-images (Maister et al., 2021).

We explored the determinants of perceived attentiveness in two preregistered online experiments. Observers viewed pairs of faces created by adding randomly generated noise to a constant base face, and had to select which appeared to be most attentive — as in the example depicted in Figure 1. Automated computer vision tools were then used to measure facial landmarks in the resulting ‘classification images’, and to explore how such landmarks differed for attentive vs. distracted faces — with a focus on both seemingly straightforward properties

(such as eye size and gaze direction) and those that seem less intuitive (such as the subtle characteristics of mouths and nostrils). We report the results of both a primary experiment and a direct replication, which was run to ensure the reliability of the results and generalizability to a new sample.



Figure 1. Depiction of the key task: on each trial, observers viewed two images (employing the same base face, but different patterns of random noise) and simply reported which looked more attentive.

Method

Observers

For each experiment, 200 observers (Experiment 1: 83 females, 115 males, 2 non-binary, $M_{\text{age}}=28.94$, age range=18–70; Experiment 2: 83 females, 113 males, 3 non-binary, 1 fluid, $M_{\text{age}}=30.34$, age range=18–75) were recruited through Prolific Academic (prolific.com; Palan & Schitter, 2018), and each completed a 20-minute session in exchange for monetary

compensation according to the standard suggested rate on Prolific. All observers reported normal or corrected-to-normal visual acuity, and completed the experiment on a desktop or laptop computer (as opposed to a phone or tablet). There were no specific restrictions with respect to age, residence, or language; the only additional inclusion criterion was that observers had not previously taken part in previous studies from our laboratory involving perceived attentiveness (e.g., on pupil dilation; see Colombatto & Scholl, 2022). This preregistered sample size was chosen for Experiment 1 before data collection began based on pilot data, and was fixed to be identical in Experiment 2.

Apparatus

Observers were redirected to a website where stimulus presentation and data collection were controlled by custom software written in HTML, JavaScript, CSS, and PHP. Since the experiment was rendered on observers' own web browsers, viewing distance, screen size, and display resolutions could vary dramatically, but stimulus size was standardized to each person's screen using the Virtual Chin Rest (Li et al., 2020). Observers' browser windows were automatically put in full-screen mode at the beginning of the experiment.

Stimuli and Design

Observers viewed pairs of faces created by superimposing random noise patterns on a constant base face (a grayscale average of all male faces in the Karolinska Face Database; Lundqvist et al., 1998). There were 1200 noise patterns, each subtending 512×512 px. Each noise pattern consisted of 60 superimposed sinusoidal patterns at five spatial frequency scales (1, 2, 4, 8, and 16 cycles per image), six orientations (0° , 30° , 60° , 90° , 120° , and 150°), two phases (0, $\pi/2$) and random contrast (Dotsch et al., 2008; Dotsch & Todorov, 2012). Each of these 1200 patterns was then added to the base face ('original noise') and subtracted from the base face ('inverted'), for a total of 2400 images, or 1200 pairs.

On each trial, observers viewed a pair of faces in noise and were asked to select which of the two appeared more attentive (as detailed below, and as depicted in Figure 1). Each observer

viewed 400 pairs of faces (one on each trial), randomly selected among the 1200 pairs. The faces each subtended 10° and were centered on the screen, 5° apart horizontally. The position of the ‘original noise’ stimulus within each pair (i.e., with added vs. subtracted noise) was counterbalanced, such that it appeared on the left on half the trials, and on the right on the remaining half.

This forced selection between pairs of stimuli (with added vs. subtracted noise) is but one way to measure observers’ impressions in reverse correlation paradigms. Such impressions can also be obtained via other means, such as continuous ratings of single stimuli (e.g., Oosterhof & Todorov, 2008; Jack et al., 2012). Forced selections between complementary stimuli, however, allow for greater discriminability, and a greater signal as measured in fewer trials (for a discussion, see Brinkman et al., 2017). Perhaps thanks to such advantages, this form of reverse correlation has been commonly used in face perception research (e.g., Dotsch et al., 2008; Dotsch & Todorov, 2012).

Procedure

After agreeing to a consent form and completing the virtual chinrest procedure, observers were told that they would view pairs of faces, and their task would be to “select which of the two people looks more attentive” — defined as “someone who seems focused and engaged with what is going on around them, as opposed to someone who seems distracted or lost in thought”. We also emphasized that they should report their first impressions rather than calculated judgments: “Of course, this task might seem a bit odd, since you don't know these people and we’re asking you to determine who looks more attentive by just looking at their faces. Nonetheless, you might find that you sometimes have a gut intuition about who looks more attentive — and that’s what we’re after. So when making your responses, don’t think about it too much: we’re really just interested in your gut reaction.”

Observers then completed 400 trials each, with self-paced breaks every 25% of the way. Each trial began with a 2° white fixation cross on a gray background (HEX #404040) for 1s,

followed by the stimuli, until a keypress. Observers responded by pressing the “left” or “right” arrow keys to select the face on the left or the right, respectively. The next trial began immediately after a response.

At the end of the experiment, observers answered a series of questions that allowed us to exclude those (per the preregistered criteria) who reported having interrupted the survey (Expt 1 $n=24$; Expt 2 $n=32$); reported past participation in a similar study ($ns=10, 13$); reported not paying attention (by answering less than 50 on a 1-100 scale, from “I was very distracted” to “I was very focused”; $ns=14, 24$); encountered problems ($ns=6, 5$); or who failed to answer our questions sensibly ($ns=2, 2$). In addition, we excluded observers whose responses were unreliable, potentially due to inattentive or random responding. To quantify reliability, we computed classification images for each observer separately for the first and second half of the experiment (i.e., in trials 1-200 vs. 201-400), and we excluded observers who had a negative correlation ($ns=104, 131$).¹ Observers across these criteria ($ns=134, 172$, some of whom triggered multiple criteria) were excluded and replaced without ever analyzing their data.

Analyses

As per our preregistered plan, we first computed classification images by averaging the selected noise patterns, and anti-classification images by averaging the non-selected noise patterns (Dotsch, 2015). Next, we extracted landmark position and action unit activation via OpenFace (Baltrušaitis et al., 2013; Baltrušaitis et al., 2015; Wood et al., 2015), and quantified specific features from each observer’s attentive and non-attentive classification images, according to the following procedure:

Primary analyses. (1) The absolute deviation of gaze direction in radians, derived from the gaze direction angle as computed via OpenFace, and averaged across the x and y directions; (2) The area of the pupils in pixels, i.e. the area of the elliptical region bounded by the pupil

¹ This exclusion rate seemed large, though not unusually so for online data collection on such populations. In additional analyses, we confirmed that all of the patterns reported here remain qualitatively identical when including all subjects who were excluded due to this criterion.

landmarks as computed via OpenFace (landmark indices: 21, 23, 25, 27 for the left pupil; 49, 51, 53, 55 for the right pupil), and averaged across the left and right pupils; (3) The vertical position of the labial commissure as a percentage of the height of the mouth, i.e. the height of the region bounded by the mouth landmarks in 2D as computed via OpenFace (landmark indices: 48, 54 for the labial commissure; 50, 52, 57 for the mouth height), and averaged across the left and right commissure; (4) The mean luminance of the nostril region, i.e. the rectangular region bounded by the nose landmarks in 2D as computed via OpenFace (landmark indices: 30, 31, 33, 35).

Control analyses. (5) The absolute deviation of head rotation in radians, derived from the head pose rotation as computed via OpenFace, and averaged across the x and y axes; (6) The area of the chin in pixels, i.e. the area of the elliptical region with the height determined by the chin landmarks in 2D as computed via OpenFace (landmark indices: 57, 8), and the width determined by the mouth landmarks in 2D as computed via OpenFace (landmark indices: 58, 56); (7) The vertical position of the mouth as a percentage of the height of the lower half of the face, i.e. the height of the region bounded by the face landmarks in 2D as computed via OpenFace (landmark indices: 50, 52, 57 for the mouth; 33, 8 for the lower half of the face); (8) The mean luminance of a rectangular region on the cheek of width, height, and vertical position equal to the nostril region (see #4), but horizontally centered at the halfway point of the cheek region bounded by the face landmarks in 2D as computed via OpenFace (landmark indices: 1, 31 for the left cheek; 35, 15 for the right cheek), and averaged across the left and right cheek.

Exploratory analyses. (9) The sum of raisers in the brow region (AU1, AU2) minus the brow lowerer (AU4); (10) The sum of raisers in the eye region (AU5, AU6, AU7); (11) The difference between the lip corner puller (AU12) and depressor (AU15); (12) The intensity of blinks (AU45).

Results

The top row of Figure 2 presents the resulting ‘classification images’ — the averages of those random noise patterns that were collectively perceived to produce attentive vs. inattentive faces (with the images from Experiment 1 depicted in Figure 2a, and those from Experiment 2 depicted in Figure 2b). When overlaid back onto the base face (as in the middle row of Figure 2), these images reveal the noticeably distinct prototypical visual representations of attentiveness and inattentiveness. Unsurprisingly, attentive faces differed in the eye region, but they also differed in more unexpected and even counterintuitive ways — e.g. in the curvature of the labial commissure, or the salience and contrast of the nostrils.

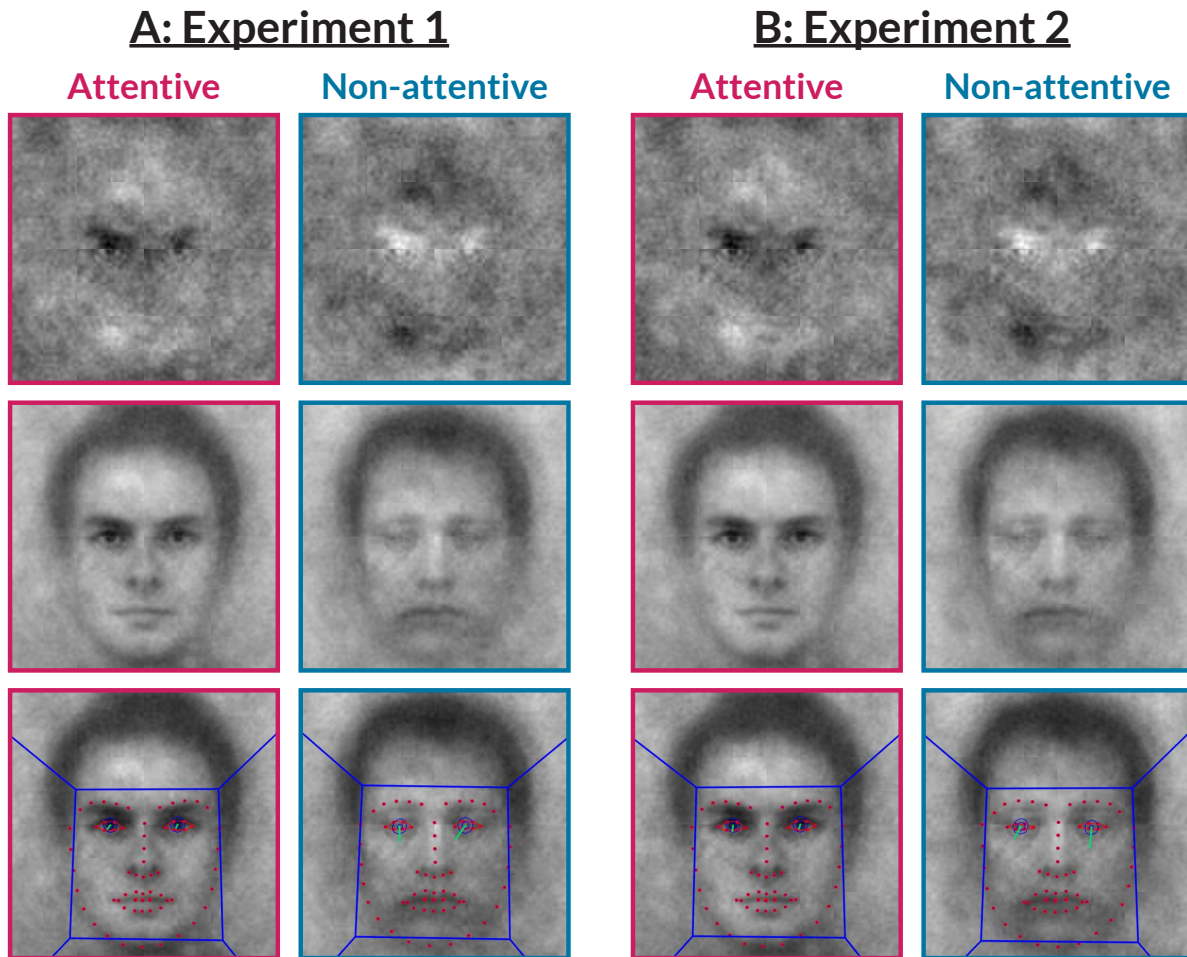


Figure 2. Determinants of perceived attentiveness. Classification images for attentive and non-attentive faces in Experiment 1 (A) and Experiment 2 (B), presented in isolation (top row), overlaid over the base face (middle row), and with the facial landmarks as extracted via OpenPose (bottom row).

Primary analyses. To quantify these differences, we extracted facial landmarks via OpenFace, a toolkit for automated facial behavior analysis, including landmark detection and gaze estimation (Baltrušaitis et al., 2018; Figure 2, bottom row; for details on feature extraction, see Method). Figure 3 depicts landmark data for each observer in Experiment 1 (first column) and Experiment 2 (second column), as well as those resulting faces with the minimum and maximum values for each feature across both experiments (third and fourth columns). A series of pre-registered paired-sample t-tests confirmed that compared to the resulting non-attentive faces, the resulting attentive faces had more direct eye gaze (Expt 1: $t(199)=5.38, p<.001, d=0.38, CI=[0.24, 0.52]$; Expt 2: $t(199)=4.40, p<.001, d=0.31 [0.17, 0.45]$; Figure 3A) and also more dilated pupils (Expt 1: $t(199)=7.11, p<.001, d=0.50 [0.36, 0.65]$; Expt 2: $t(199)=8.13, p<.001, d=0.57 [0.42, 0.72]$; Figure 3B). Attentive faces also had higher labial commissures relative to their mouth position — as if they were smiling (Expt 1: $t(199)=2.67, p=.008, d=0.19 [0.05, 0.33]$; Expt 2: $t(199)=3.79, p<.001, d=0.27 [0.13, 0.41]$; Figure 3C). Finally, and perhaps most surprisingly, attentive faces reliably had darker and more salient nostrils — as if more flared, or retroussé (Expt 1: $t(199)=2.88, p=.004, d=0.20 [0.06, 0.34]$; Expt 2: $t(199)=3.30, p=.001, d=0.23 [0.09, 0.37]$; Figure 3D). (See Table 1 for all means and standard deviations.)

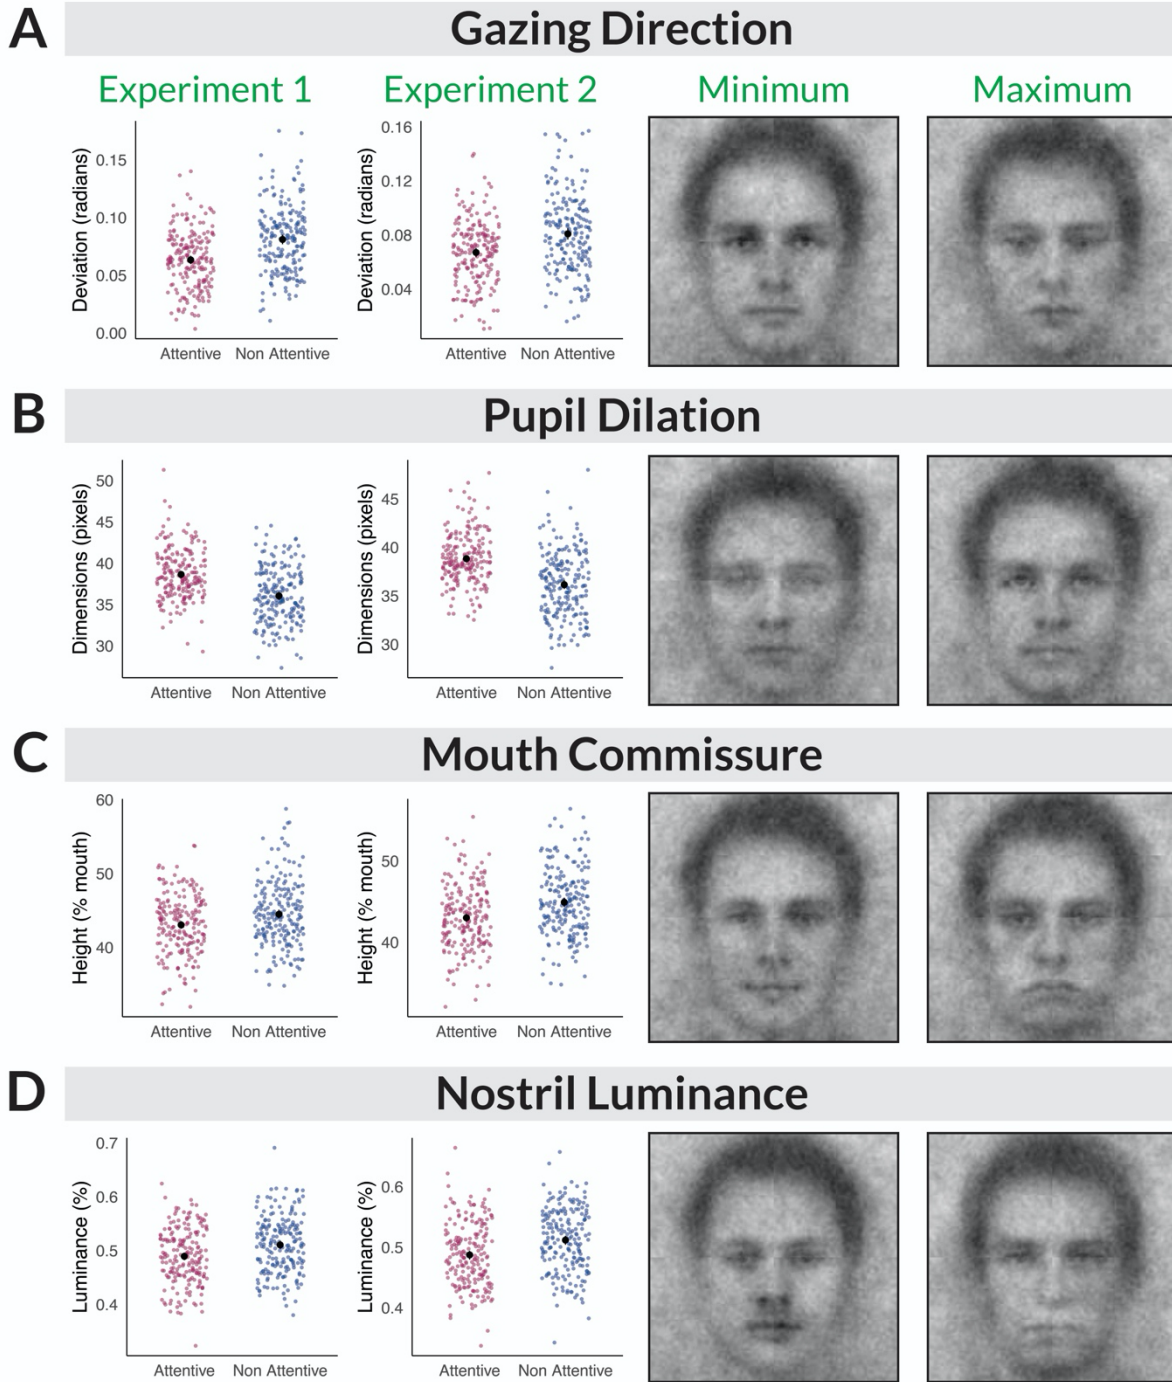


Figure 3. Analyses of facial features for attentive and inattentive faces. For each facial landmark (rows), we depict the comparison between attentive and non-attentive classification images overlaid onto the base face for Experiment 1 (first column) and Experiment 2 (second column), and the resulting faces with highest and lowest values across both experiments (third and fourth columns). Individual points represent values for the CI of each observer – jittered horizontally to avoid excessive overlap. Error bars (barely visible) represent 95% confidence intervals.

Primary Analyses					
Measures		Experiment 1		Experiment 2	
Feature	Unit	Attentive	Non-Attentive	Attentive	Non-Attentive
Gaze deviation	Radians	0.06 (0.02)	0.08 (0.03)	0.07 (0.02)	0.08 (0.03)
Pupil area	Pixels	38.63 (3.01)	36.03 (3.33)	38.82 (2.62)	36.14 (3.32)
Labial Comm. height	%	43.04 (4.01)	44.47 (4.40)	43.00 (3.88)	44.90 (4.08)
Nostril luminance	%	0.49 (0.05)	0.51 (0.05)	0.49 (0.05)	0.51 (0.05)

Control Analyses					
Measures		Experiment 1		Experiment 2	
Feature	Unit	Attentive	Non-Attentive	Attentive	Non-Attentive
Head deviation	Radians	0.05 (0.01)	0.05 (0.02)	0.05 (0.02)	0.05 (0.02)
Chin area	Pixels	157.08 (4.89)	158.44 (5.77)	157.79 (5.68)	158.60 (6.27)
Mouth height	%	35.24 (0.73)	35.24 (0.80)	35.23 (0.89)	35.26 (0.94)
Cheek luminance	%	0.51 (0.03)	0.49 (0.03)	0.51 (0.03)	0.49 (0.03)

Exploratory Analyses					
Measures		Experiment 1		Experiment 2	
Feature	Unit	Attentive	Non-Attentive	Attentive	Non-Attentive
Brow raisers	A.U.	-0.20 (0.62)	-0.36 (0.68)	-0.16 (0.64)	-0.36 (0.65)
Eye raisers	A.U.	0.53 (0.51)	0.51 (0.55)	0.55 (0.57)	0.64 (0.62)
Lip raisers	A.U.	-0.11 (0.40)	-0.39 (0.45)	-0.09 (0.41)	-0.34 (0.40)
Blink intensity	A.U.	0.68 (0.27)	0.89 (0.28)	0.68 (0.27)	0.91 (0.33)

Table 1. Means and standard deviations (in parentheses) for primary, control, and exploratory analyses, for both experiments. ‘Comm’ = Commissure.

Control analyses. In additional preregistered control analyses, we verified that such differences were specific to the highlighted features of interest, and not to other facial features: the resulting attentive faces had more direct eye gaze, but if anything a larger head rotation (Expt 1: $t(199)=1.99, p=.048, d=0.14$ [0.00, 0.28]; Expt 2: $t(199)=1.60, p=.112, d=0.11$ [0.03, 0.25]). (In additional analyses not reported here, we also confirmed that attentive and non-attentive faces did not differ systematically when only considering tilt along only the X axis, i.e. the pitch.) Attentive faces had larger pupils, but if anything smaller chin areas (Expt 1: $t(199)=2.13, p=.034, d=0.15$ [0.01, 0.29]; Expt 2: $t(199)=1.11, p=.270, d=0.08$ [-0.06, 0.22]). Even within the mouth region, attentive faces had a higher labial commissure, but not a higher mouth overall (Expt 1: $t(199)=0.05, p=.960, d=0.00$ [-0.14, 0.14]; Expt 2: $t(199)=0.28, p=.781,$

$d=0.02$ [-0.12, 0.16]). And while attentive faces had darker nostrils, they did not have darker cheeks, which if anything were brighter (Expt 1: $t(199)=6.79$, $p<.001$, $d=0.48$ [0.63, 0.33]; Expt 2: $t(199)=3.29$, $p=.001$, $d=0.23$ [0.09, 0.37]).

Exploratory analyses. Beyond individual facial features, we also conducted an additional preregistered secondary analysis of action unit activation. Attentive faces had more activations in the raisers in the brow region (Expt 1: $t(199)=2.35$, $p=.020$, $d=0.17$ [0.03, 0.31]; Expt 2: $t(199)=2.97$, $p=.003$, $d=0.21$ [0.07, 0.35]) as well as in the lip region (Expt 1: $t(199)=5.29$, $p<.001$, $d=0.37$ [0.23, 0.52]; Expt 2: $t(199)=5.01$, $p<.001$, $d=0.35$ [0.21, 0.50]). There were no differences in the raisers for the eye region (Expt 1: $t(199)=0.23$, $p=.818$, $d=0.02$ [-0.12, 0.15]; Expt 2: $t(199)=1.39$, $p=.165$, $d=0.10$ [-0.04, 0.24]), although attentive classification images had lower intensity in the blink muscles (Expt 1: $t(199)=6.65$, $p<.001$, $d=0.47$ [0.32, 0.62]; Expt 2: $t(199)=6.52$, $p<.001$, $d=0.46$ [0.31, 0.61]).

Discussion

This study explored for the first time how we perceive attentiveness in others' faces — and the answers were straightforward in some ways, yet deeply counterintuitive in others. On one hand, attentive faces had wider pupils and more direct gaze. This is relatively unsurprising in both lay and scientific terms. After all, a common synonym for attentive is 'watchful'. And both features have been linked to underlying mental states (Benedek et al., 2018; Colombatto & Scholl, 2022; Emery, 2000; Ziman et al., 2023), and are frequently measured by scientists studying attention (Sirois & Brisson, 2014). On the other hand, attentive faces also had higher lip curvature and nostrils that more prominently flared, or *retroussé*. This is a novel and counterintuitive pattern, to say the least: there are thousands of studies of the influence of eyes on social perception (Emery, 2000), but almost none about the importance of nostrils. (The last 10 years of presentations at the *Vision Sciences Society* [a leading venue for perception

research] featured over 15000 abstracts, only 4 of which even mentioned nostrils, and none in the context of attention.)

These investigations raise several interesting questions about the origins of internal representations of attentiveness and the relative contribution of specific features to global impressions of attentiveness. For example, we may speculate about the roles that the counterintuitive features — especially those involving the nostrils and mouths — may have played. Why might more prominently flared (or retroussé) nostrils be associated with perceived attentiveness? This could be related to other properties such as head orientation or eye widening. When a person is attentive, they might subtly tilt or raise their head, which could make the nostrils more visible — though this was not supported by the analyses of head deviation. Or they might raise their eyebrows or widen their eyes, making the nostrils more visible for reasons related to facial musculature. Or, of course, nostril flaring could just be an intrinsic cue to attentiveness, without arising indirectly from other factors. (Though nostrils have almost never been explored in face perception, one study suggests that ‘super-recognisers’ who excel in face recognition tend to focus more on the nose region when viewing faces; Bobak et al., 2017. And a recent study of the ‘uncanny valley’ in face perception showed that, relative to clearly human or clearly avatar faces, human-avatar ambiguous faces elicit greater fixations on the nose area; Grebot et al., 2022.) And why might variations in the lip region also be associated with perceived attentiveness? This could relate to possible links between perceived attentiveness and properties such as extraversion and trustworthiness. Indeed, in social contexts attentiveness often seems to be a positive trait, indicating engagement and social connection. And as such, attentiveness may come to be associated with cues such as smiling given their relation to these other properties. Future work could thus explore how such visual

cues to attentiveness may reflect its social significance.²

Beyond the main preregistered analyses involving gaze, pupil size, nostrils, and labial commissure, additional control analyses revealed that attentive faces also had smaller chin areas and brighter cheeks. These additional differences seem consistent with the key findings reported here. For example, brighter cheeks produce higher contrast with the darker pupils and nostrils, and smaller chin areas are consistent with a smiling expression that pulls the musculature upward. And of course, these patterns of additional results also serve to highlight how internal representations of attentiveness involve several subtle aspects of facial appearance, including broad changes in facial musculature and expression.

Overall, these powerful and consistent effects of facial cues on perceived attentiveness (a) contribute to the growing study of naive theories of attention (Elekes & Király, 2021), (b) highlight how perceived attentiveness transcends the perceived *direction* of attention (e.g. via gaze), and (c) underscore the importance of attention not just as a perceptual process, but as an object of perception itself.

² Noticing such subtle differences in everyday life might require a situation wherein someone is continually looking toward you while initially not paying attention. This may occur more often in the modern era, due to videoconferencing — as when one suddenly realizes during a call that their videocamera is actually on. Anecdotal evidence from one of the authors suggests that in such situations one may immediately change one's facial expression from a resting pose that is simultaneously neutral, distracted, and unflattering (cf. 'perceived resting negative emotion'; Hester, 2019) to one that is abruptly much different and highly attentive.

Open Practices Statement

The hypotheses, methods, and analysis plan were preregistered prior to data collection. Preregistration for Experiment 1 can be found at https://aspredicted.org/K8Q_2M1 . There were no deviations from the preregistration. Experiment 2 was a direct replication of Experiment 1. All study materials, primary data, and analysis scripts are publicly available at https://osf.io/2pt7y/?view_only=1cd6095a5fef43faaaee5807a6eedf6b .

Declaration

Acknowledgments: For helpful conversations and/or comments on previous drafts, we thank the members of the Yale Perception & Cognition Laboratory.

Funding: This project was funded by ONR MURI #N00014-16-1-2007 awarded to BJS.

Conflicts of interest/Competing interests: The authors declare no conflicts of interest.

Ethics approval: These experiments were approved by an internal IRB panel at Yale University. The procedures adhere to the tenets of the Declaration of Helsinki.

Consent to participate: All participants provided informed consent prior to participating in the study.

Consent for publication: Not applicable.

Availability of data and materials: All primary data and study materials are publicly available at https://osf.io/2pt7y/?view_only=1cd6095a5fef43faaaee5807a6eedf6b .

Code availability: Analysis scripts are publicly available at https://osf.io/2pt7y/?view_only=1cd6095a5fef43faaaee5807a6eedf6b .

Authors' contributions: **Clara Colombatto**: Conceptualization, Methodology, Software, Investigation, Formal analysis, Validation, Data curation, Visualization, Writing - Original Draft, Writing - Review & Editing. **Brian J. Scholl**: Conceptualization, Methodology, Visualization, Writing - Original Draft, Writing - Review & Editing, Supervision, Funding acquisition.

References

- Baltrušaitis, T., Robinson, P., & Morency, L. P. (2013). Constrained local neural fields for robust facial landmark detection in the wild. In *Proceedings of the IEEE International Conference on Computer Vision*, 354-361.
- Baltrušaitis, T., Mahmoud, M., & Robinson, P. (2015, May). Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *2015 11th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG 2015)*, 6, 1-6.
- Baltrušaitis, T., Zadeh, A., Lim, Y. C., & Morency, L. P. (2018). Openface 2.0: Facial behavior analysis toolkit. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, 59-66.
- Benedek, M., Daxberger, D., Annerer-Walcher, S., & Smallwood, J. (2018). Are you with me? Probing the human capacity to recognize external/internal attention in others' faces. *Visual Cognition*, 26, 511-517.
- Bobak, A. K., Parris, B. A., Gregory, N. J., Bennetts, R. J., & Bate, S. (2017). Eye-movement strategies in developmental prosopagnosia and "super" face recognition. *Quarterly Journal of Experimental Psychology*, 70, 201-217.
- Brinkman, L., Todorov, A., & Dotsch, R. (2017). Visualising mental representations: A primer on noise-based reverse correlation in social psychology. *European Review of Social Psychology*, 28, 333-361.
- Colombatto, C., Chen, Y. C., & Scholl, B. J. (2020). Gaze deflection reveals how gaze cueing is tuned to extract the mind behind the eyes. *Proceedings of the National Academy of Sciences*, 117, 19825-19829.

- Colombatto, C., & Scholl, B. J. (2022). Unconscious pupillometry: An effect of “attentional contagion” in the absence of visual awareness. *Journal of Experimental Psychology: General*, *151*, 302-308.
- Dotsch, R., Wigboldus, D. H., Langner, O., & Van Knippenberg, A. (2008). Ethnic out-group faces are biased in the prejudiced mind. *Psychological Science*, *19*, 978-980.
- Dotsch, R., & Todorov, A. (2012). Reverse correlating social face perception. *Social Psychological and Personality Science*, *3*, 562-571.
- Dotsch, R. (2015). rcicr: Reverse-correlation image-classification toolbox [R package Version 0.3. 0].
- Elekes, F., & Király, I. (2021). Attention in naïve psychology. *Cognition*, *206*, Article 104480.
- Emery, N. J. (2000). The eyes have it: The neuroethology, function and evolution of social gaze. *Neuroscience & Biobehavioral Reviews*, *24*, 581-604.
- Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, *14*, 505-509.
- Grebot, I. B. D. F., Cintra, P. H. P., de Lima, E. F. F., & de Castro, M. V. (2022). Uncanny valley hypothesis and hierarchy of facial features in the human likeness continua: An eye-tracking approach. *Psychology & Neuroscience*, *15*, 28-42.
- Hester, N. (2019). Perceived negative emotion in neutral faces: Gender-dependent effects on attractiveness and threat. *Emotion*, *19*, 1490-1494.
- Jack, R. E., Garrod, O. G., Yu, H., Caldara, R., & Schyns, P. G. (2012). Facial expressions of emotion are not culturally universal. *Proceedings of the National Academy of Sciences*, *109*, 7241-7244.
- Land, M. F., & Hayhoe, M. M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research*, *41*, 3559-3565.

- Li, Q., Joo, S. J., Yeatman, J. D., & Reinecke, K. (2020). Controlling for participants' viewing distance in large-scale, psychophysical online experiments using a virtual chinrest. *Scientific Reports, 10*, 904.
- Lundqvist, D., Flykt, A., & Öhman, A. (1998). Karolinska directed emotional faces. *PsycTESTS Dataset, 91*, 630.
- Mack, A., & Rock, I. (1998). *Inattentional Blindness*. Cambridge, MA: MIT Press.: Perception without attention. In R. D. Wright (Ed.), *Visual attention* (pp. 55-76). Oxford University Press.
- Maister, L., De Beukelaer, S., Longo, M. R., & Tsakiris, M. (2021). The self in the mind's eye: Revealing how we truly see ourselves through reverse correlation. *Psychological Science, 32*, 1965-1978.
- Mangini, M. C., & Biederman, I. (2004). Making the ineffable explicit: Estimating the information employed for face classifications. *Cognitive Science, 28*, 209-226.
- Matthis, J. S., Yates, J. L., & Hayhoe, M. M. (2018). Gaze and the control of foot placement when walking in natural terrain. *Current Biology, 28*, 1224-1233.
- Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*, 11087-11092.
- Palan, S., & Schitter, C. (2018). Prolific. ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance, 17*, 22-27.
- Seli, P., Beaty, R. E., Cheyne, J. A., Smilek, D., Oakman, J., & Schacter, D. L. (2018). How pervasive is mind wandering, really? *Consciousness and Cognition, 66*, 74-78.
- Shimojo, S., Simion, C., Shimojo, E., & Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nature Neuroscience, 6*, 1317-1322.
- Sirois, S., & Brisson, J. (2014). Pupillometry. *Wiley Interdisciplinary Reviews: Cognitive Science, 5*, 679-692.

Smallwood, J., & Schooler, J. W. (2013). The restless mind. *Psychology of Consciousness: Theory, Research, and Practice*, 1, 130-149.

Todorov, A. (2017). *Face value: The irresistible influence of first impressions*. Princeton University Press.

Wood, E., Baltrušaitis, T., Zhang, X., Sugano, Y., Robinson, P., & Bulling, A. (2015). Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, 3756-3764.

Zadeh, A., Chong Lim, Y., Baltrusaitis, T., & Morency, L. P. (2017). Convolutional experts constrained local model for 3d facial landmark detection. In *Proceedings of the IEEE International Conference on Computer Vision Workshops* (pp. 2519-2528).

Ziman, K., Kimmel, S. C., Farrell, K. T., & Graziano, M. S. (2023). Predicting the attention of others. *Proceedings of the National Academy of Sciences*, 120, e2307584120.