Illusions of size

Sami Ryan Yousif

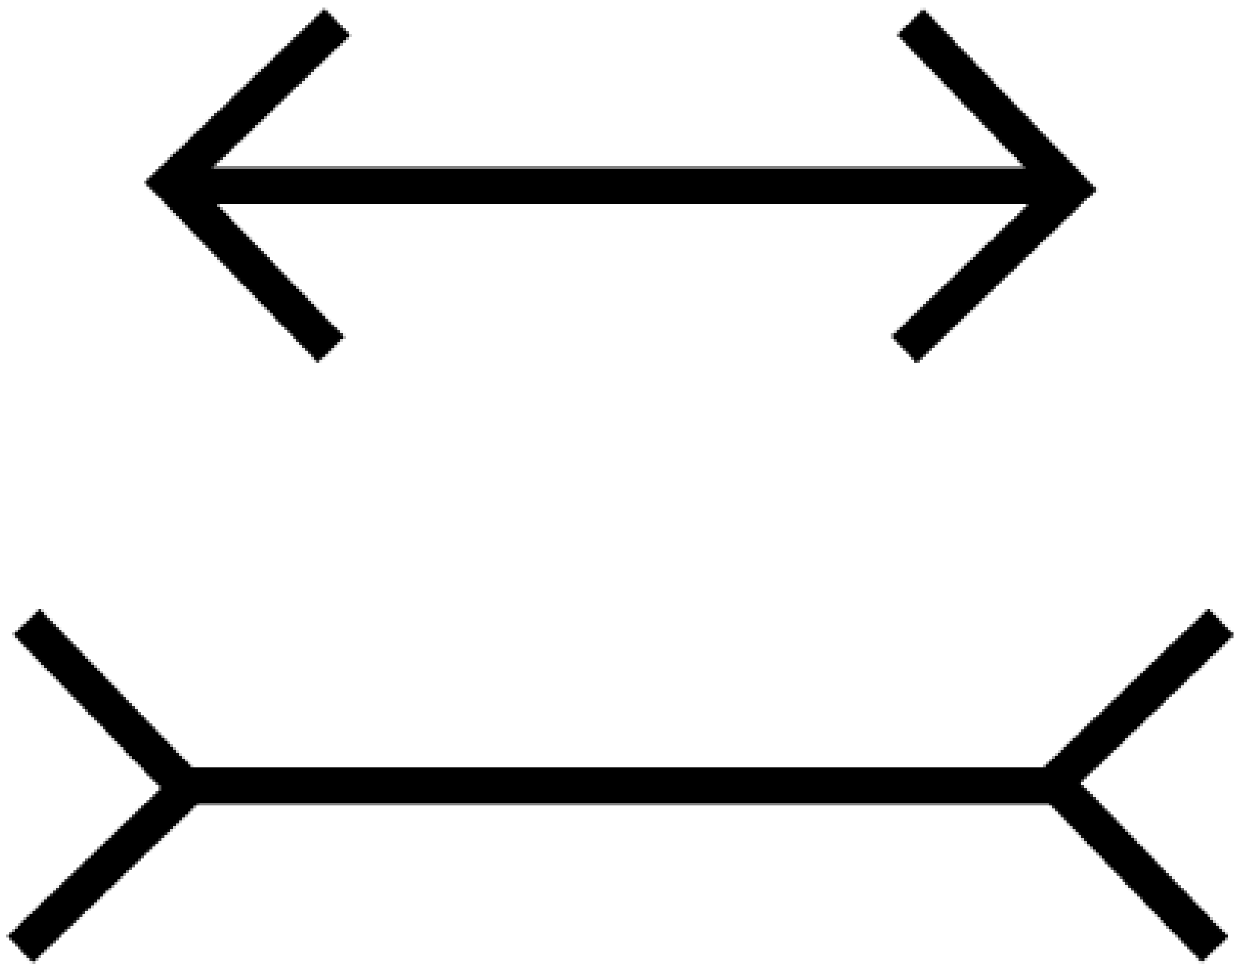

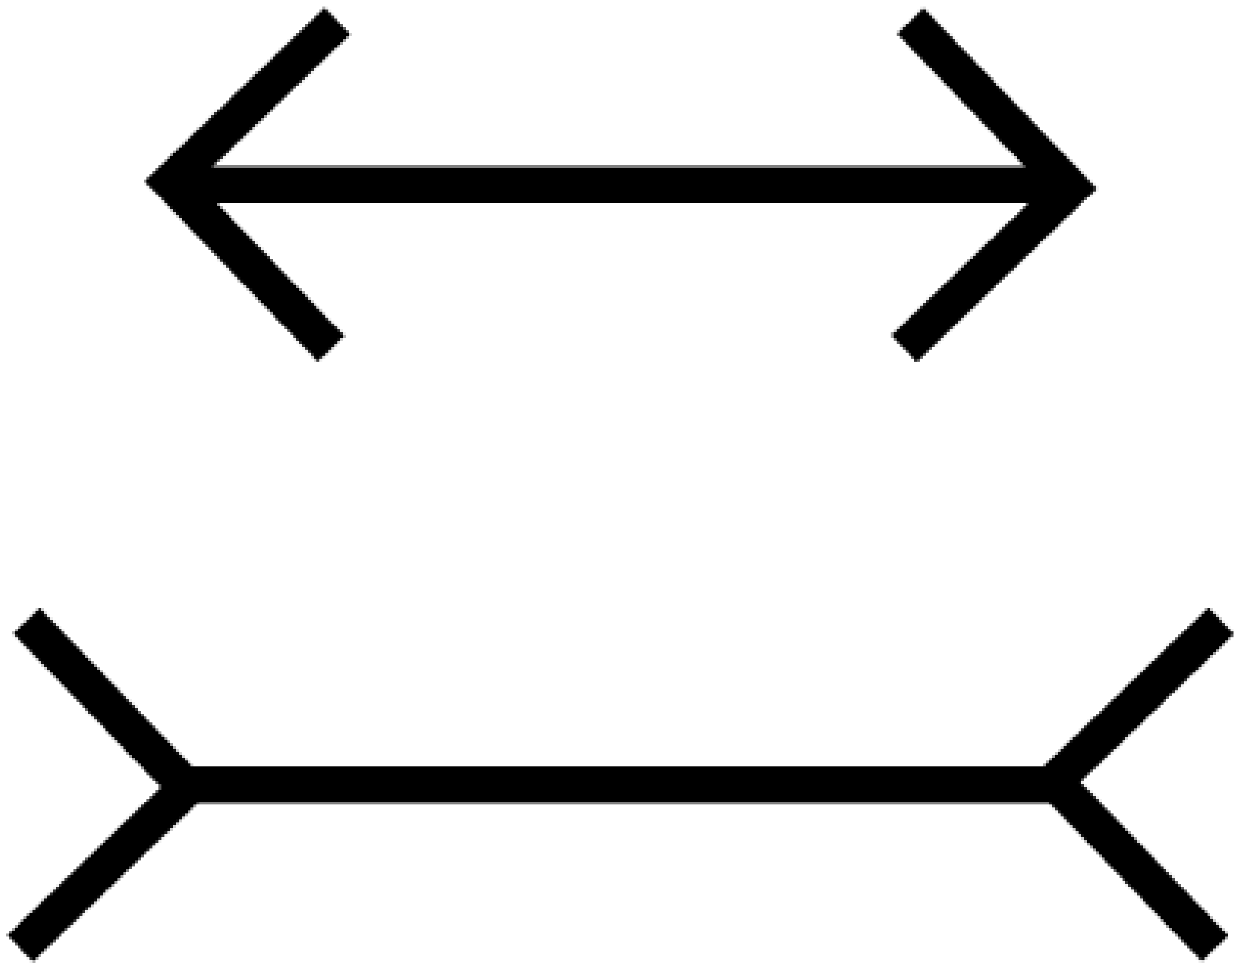

The Muller-Lyer illusion !!

Is it about **depth perception?**

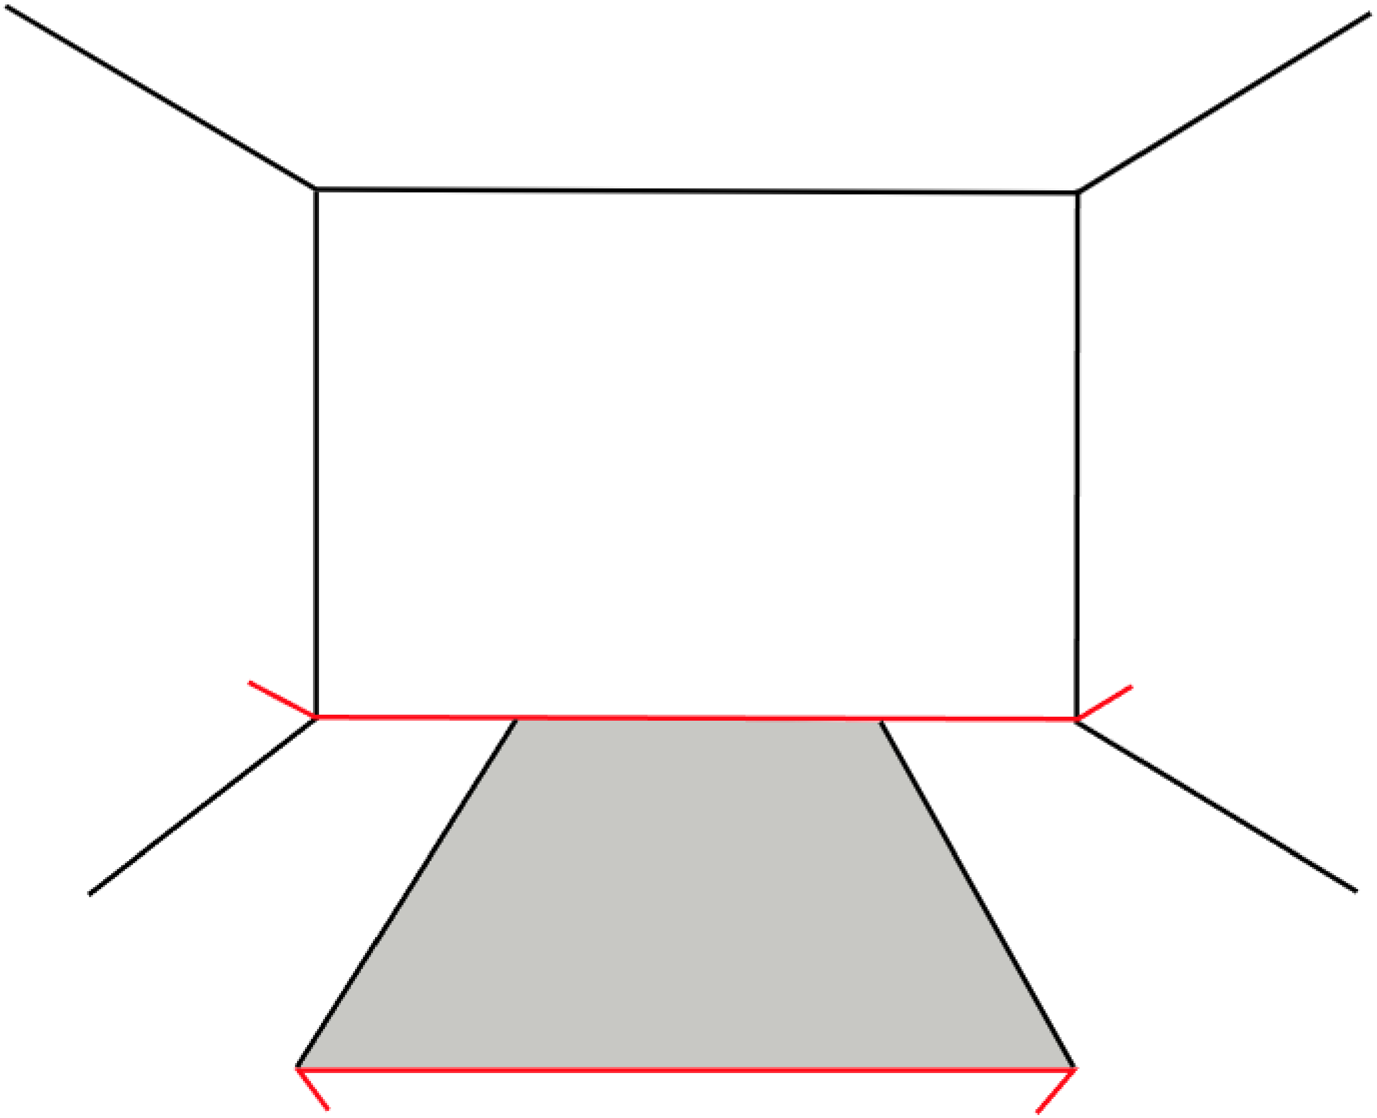

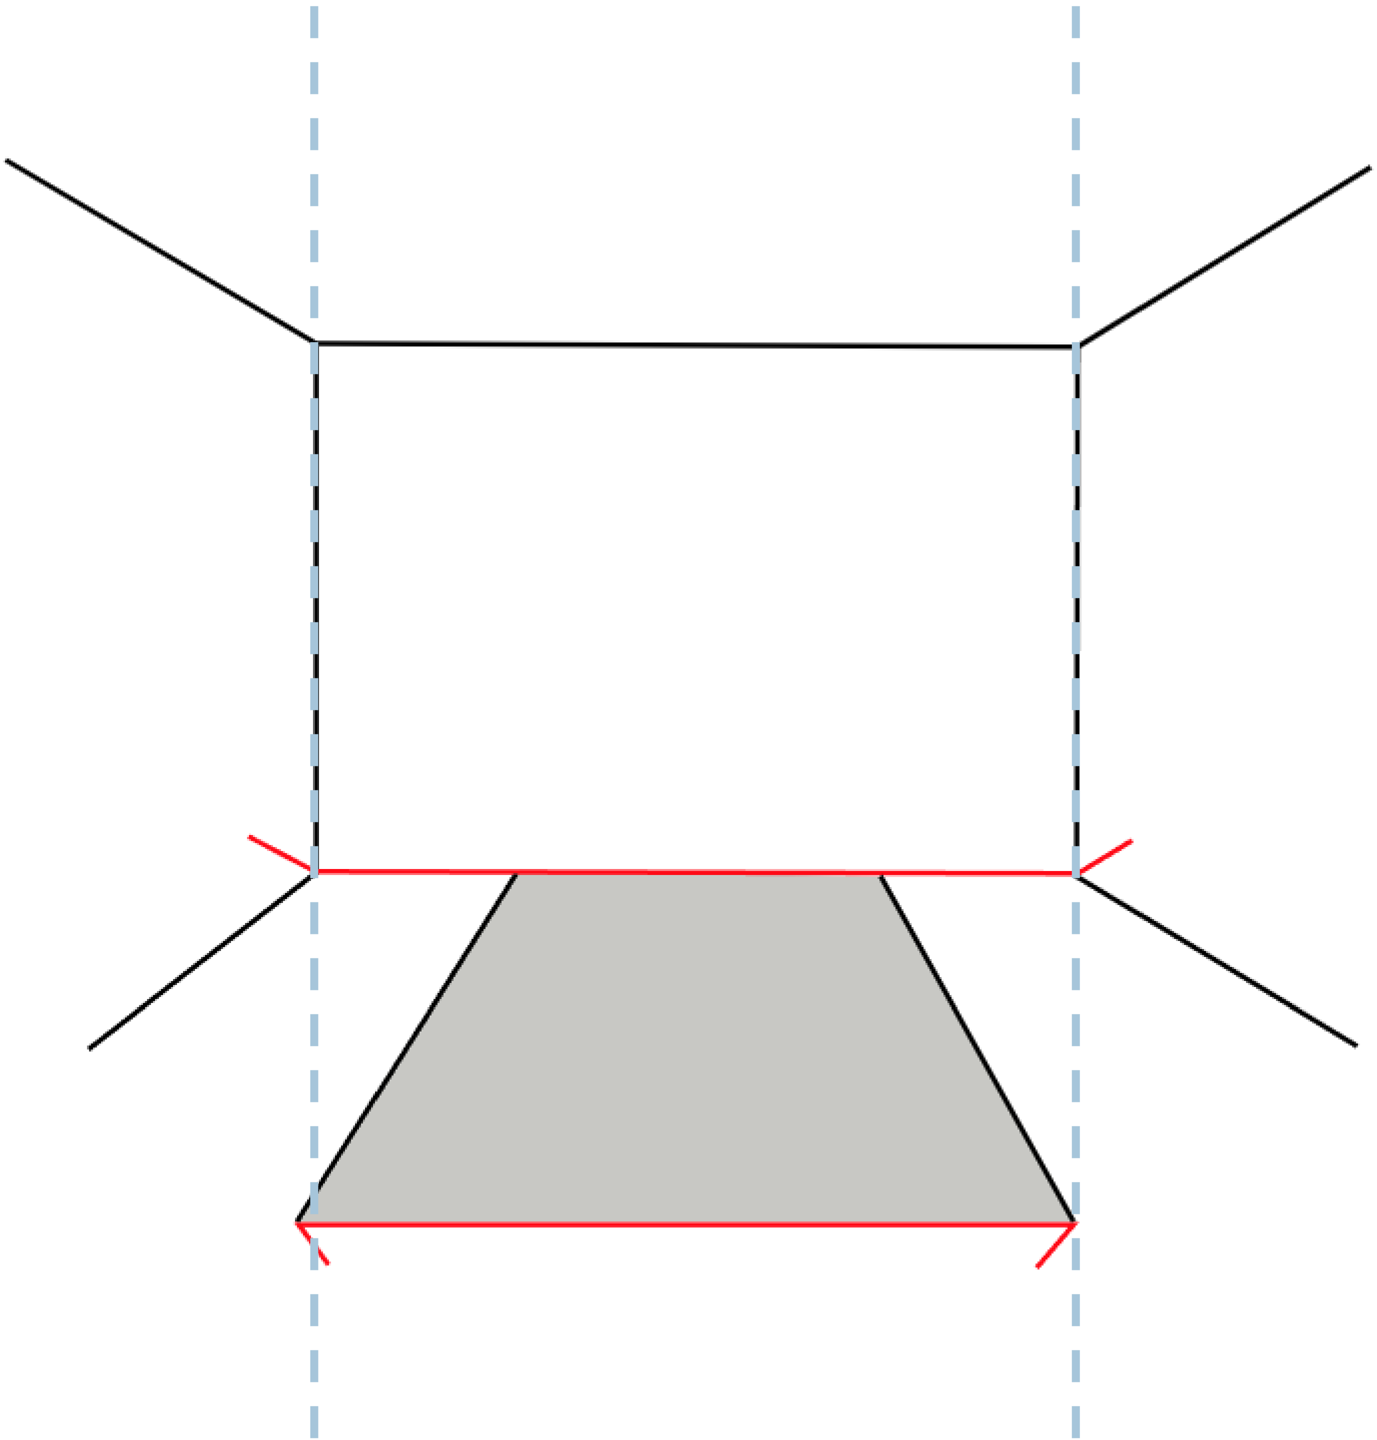

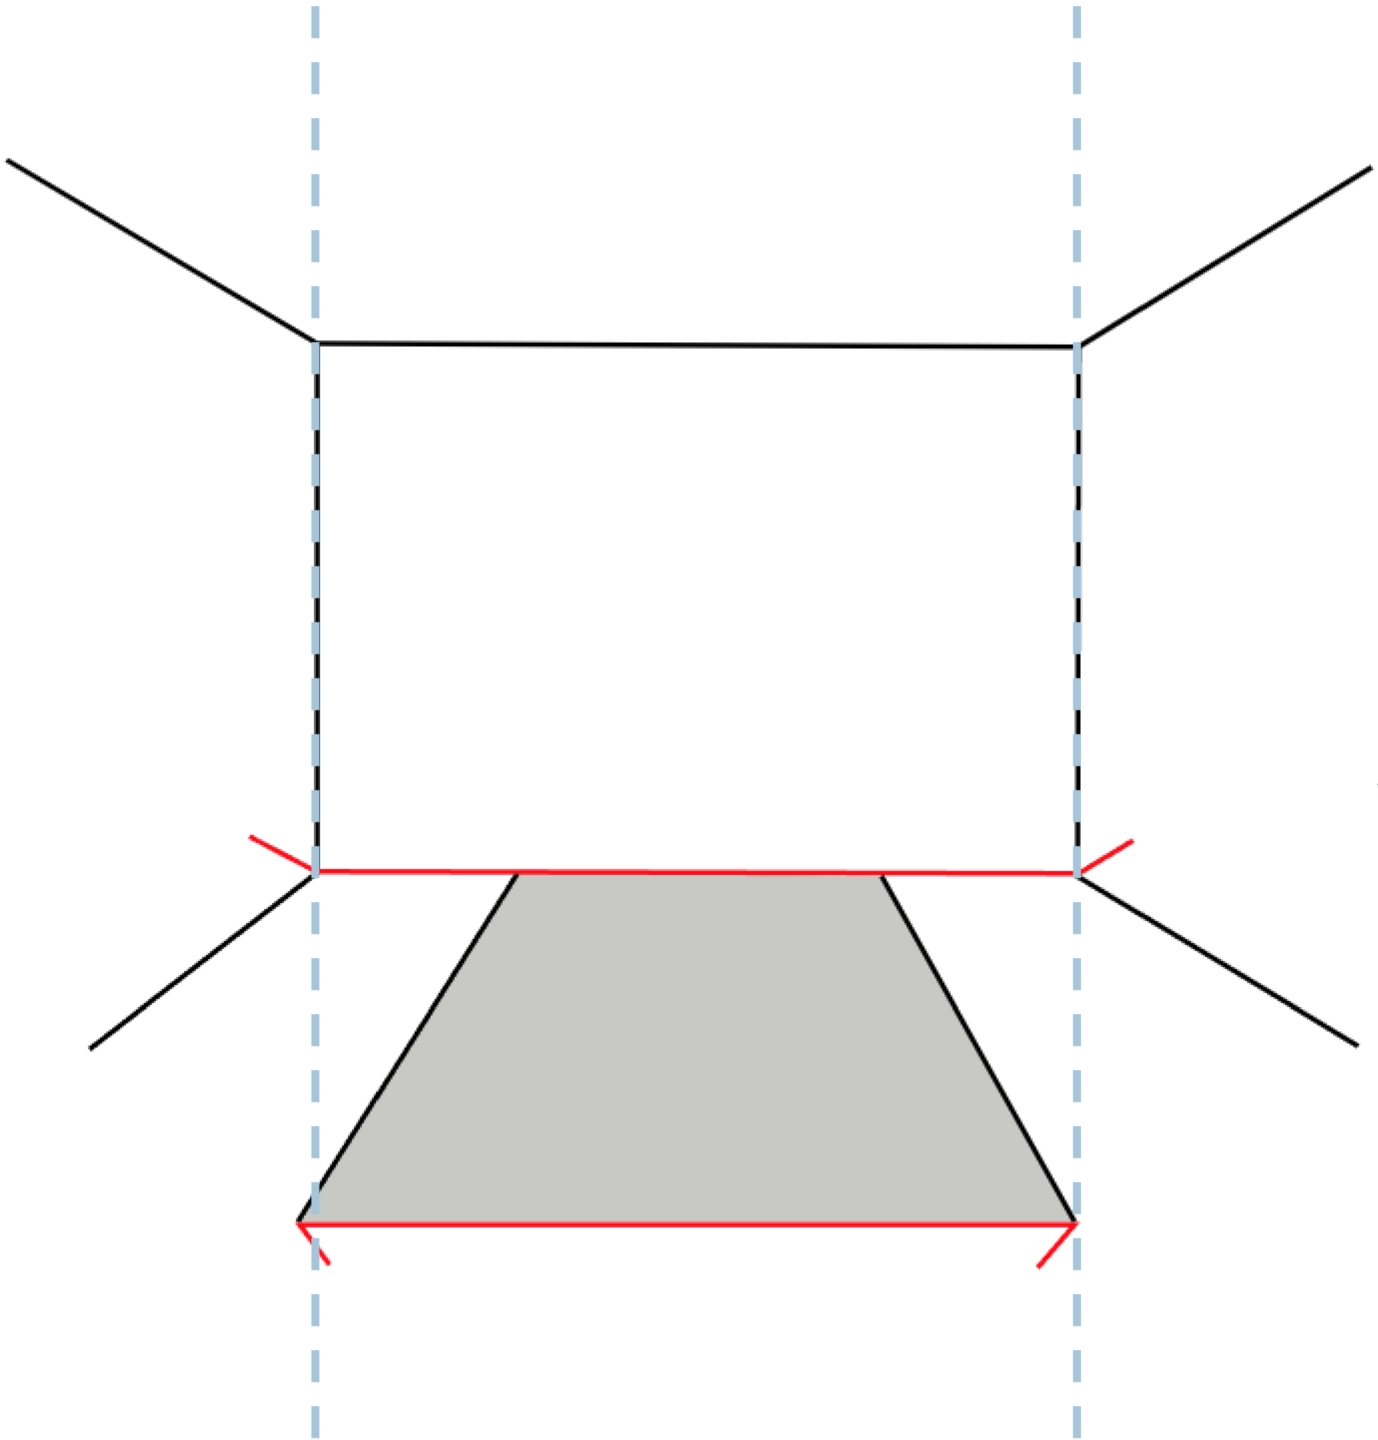

The **"Carpentered world hypothesis"**

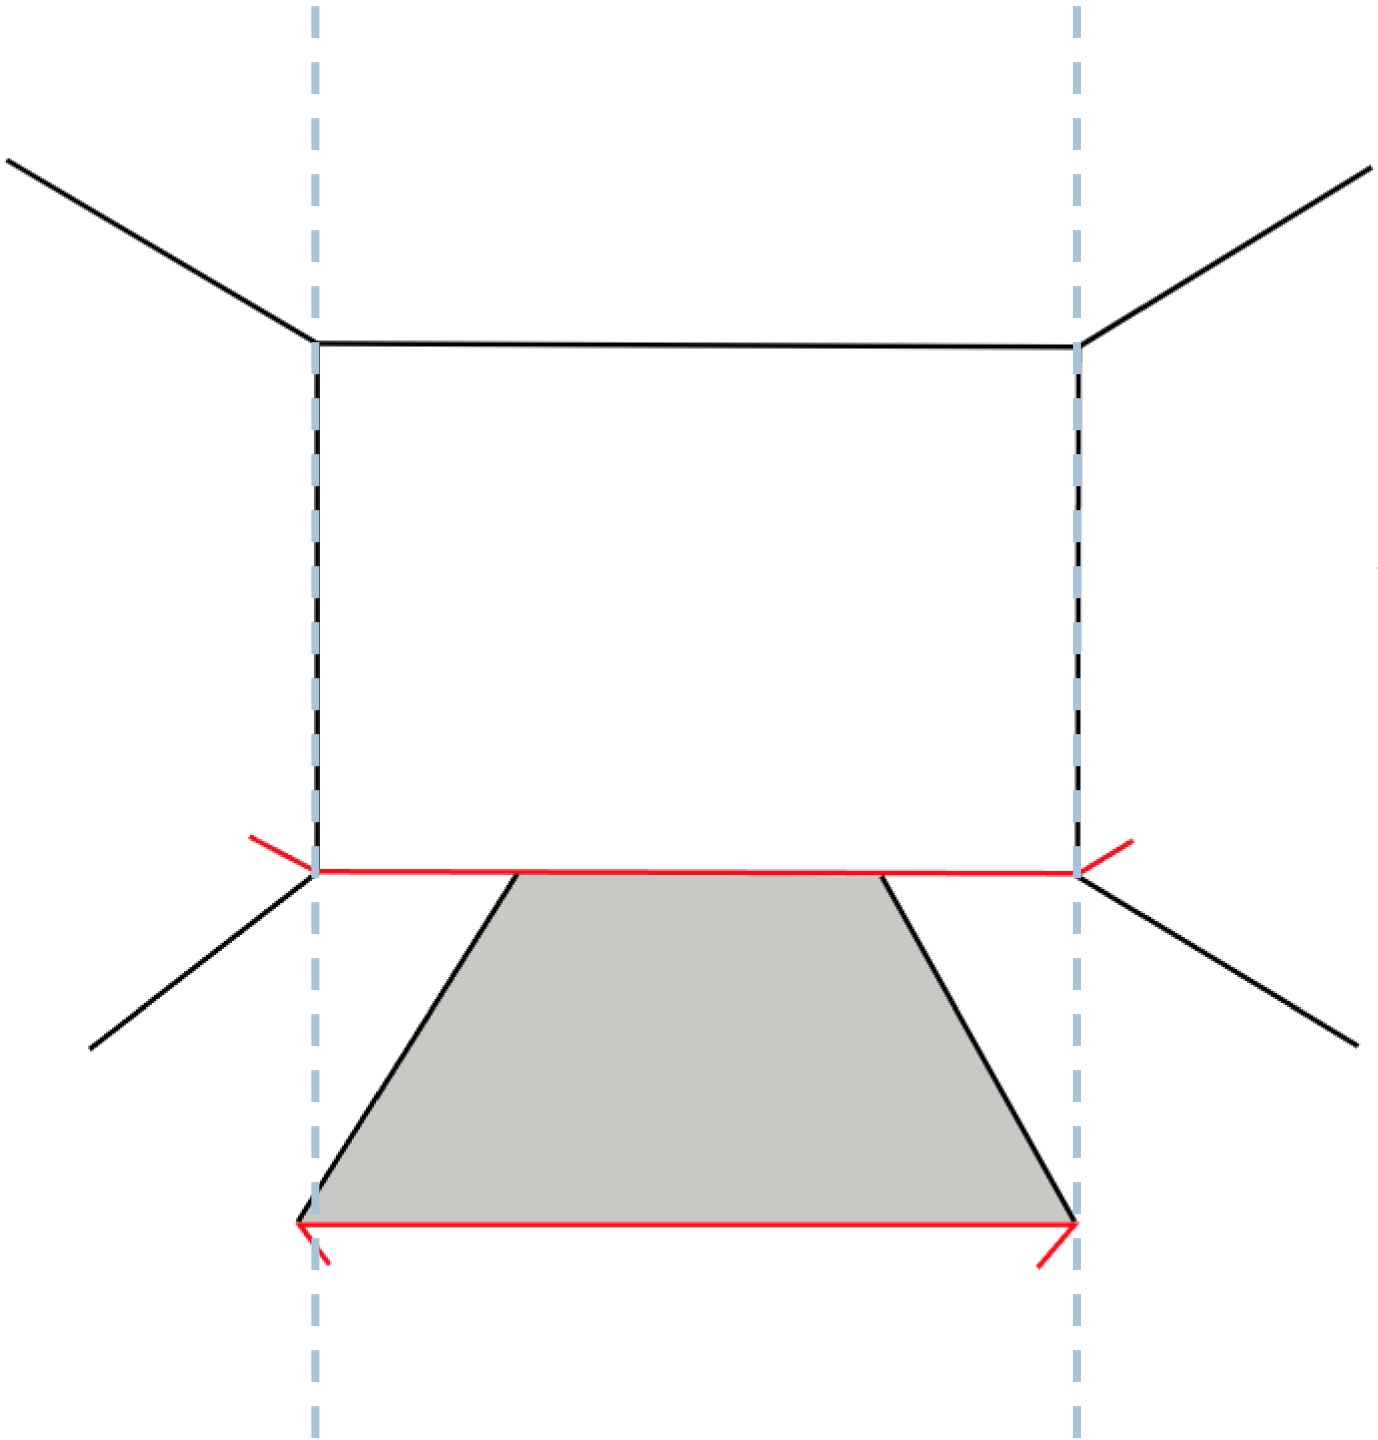

!!

The "Carpentered world hypothesis"

How can you test it?

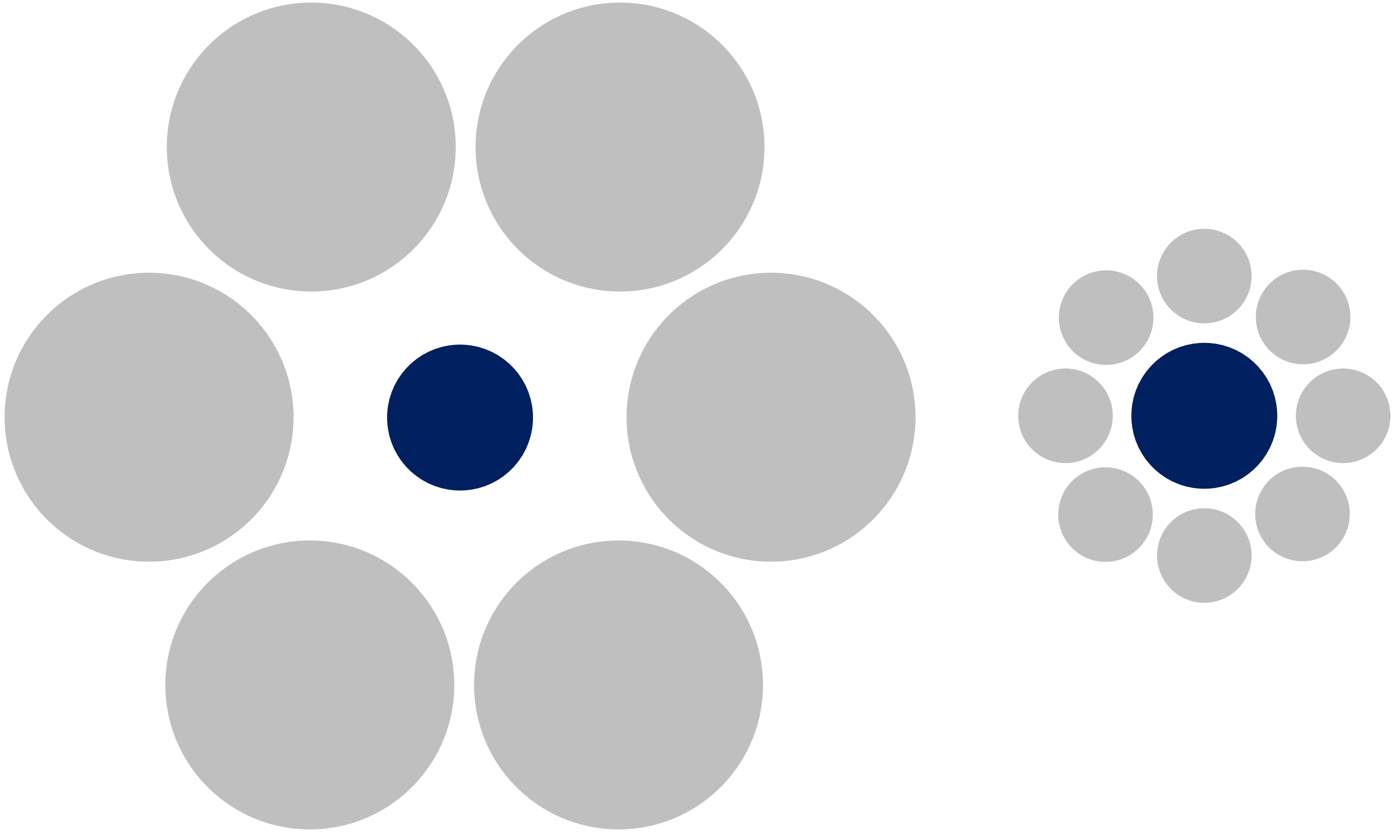

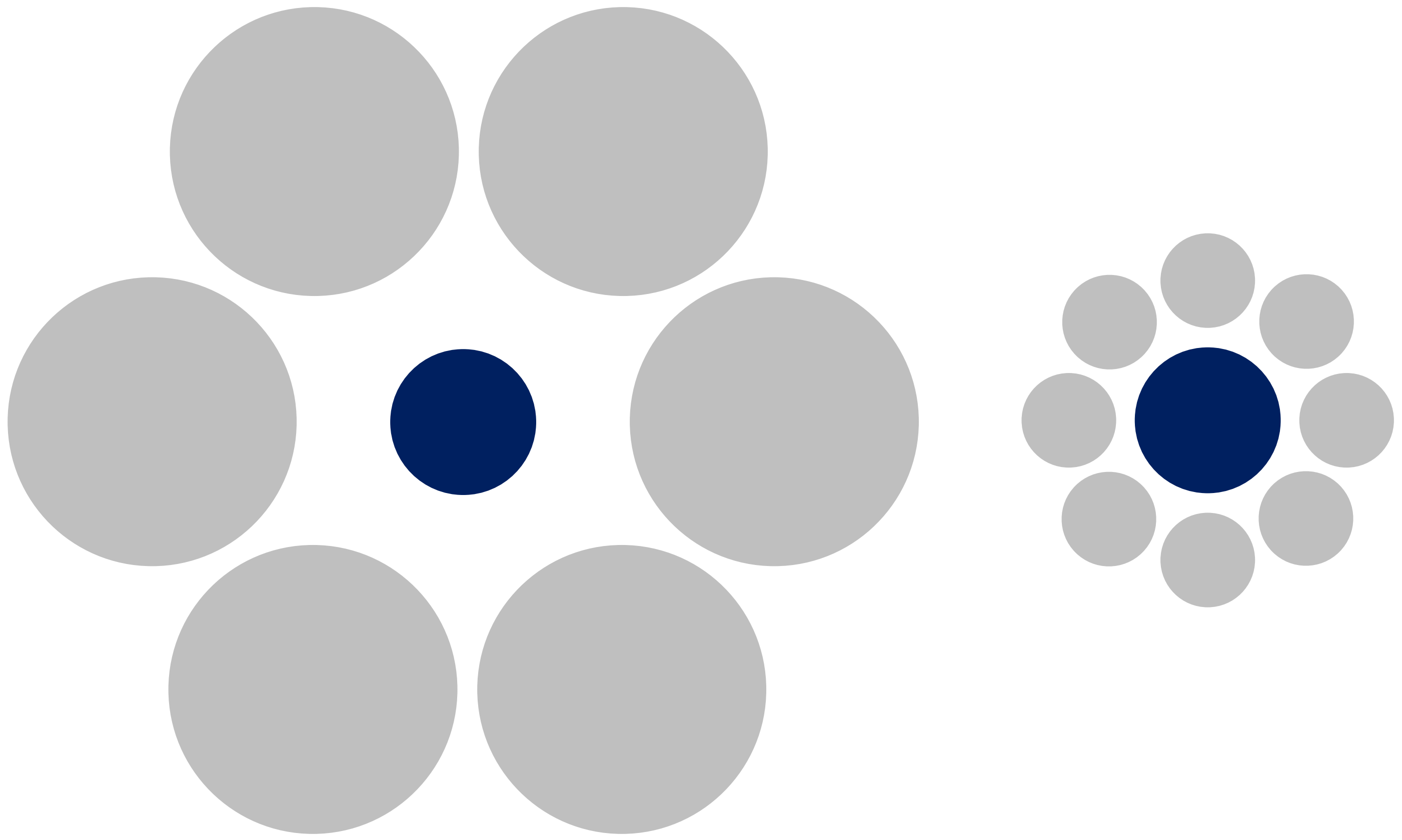

The Ebbinghaus illusion!!

The Ponzo illusion !!

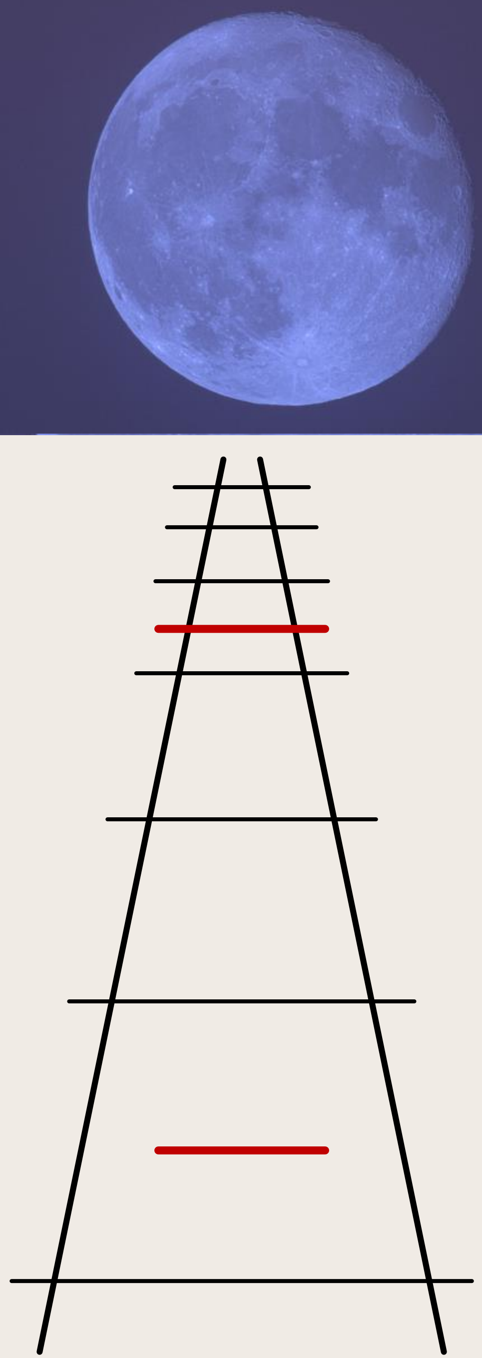

Speaking of the moon....


The moon illusion ‼

Let's talk about illusions you've not seen before.

Which two dots are farther apart?

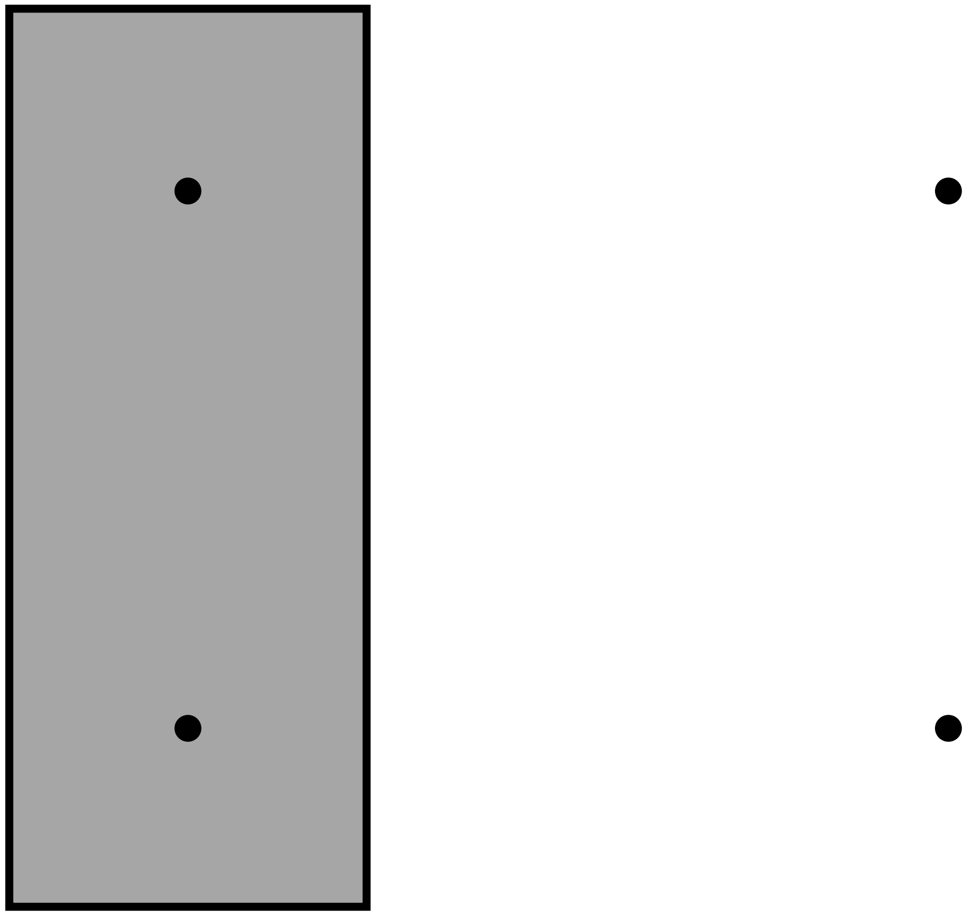

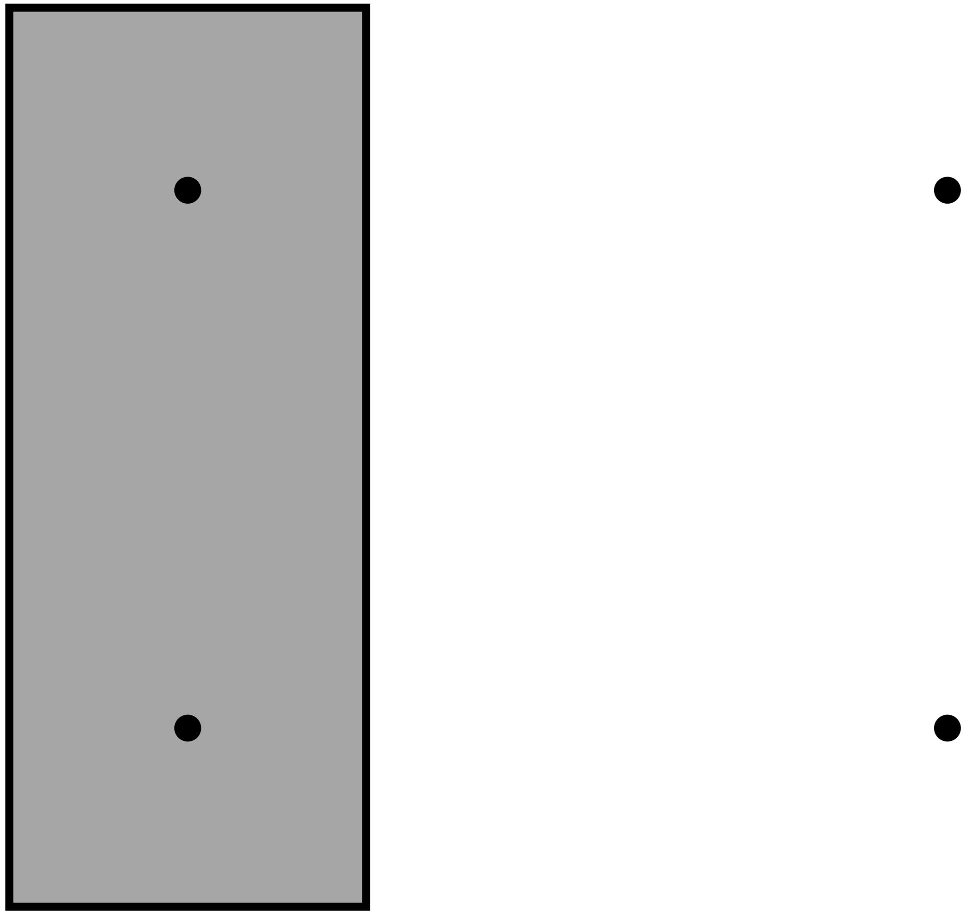

Object-based warping

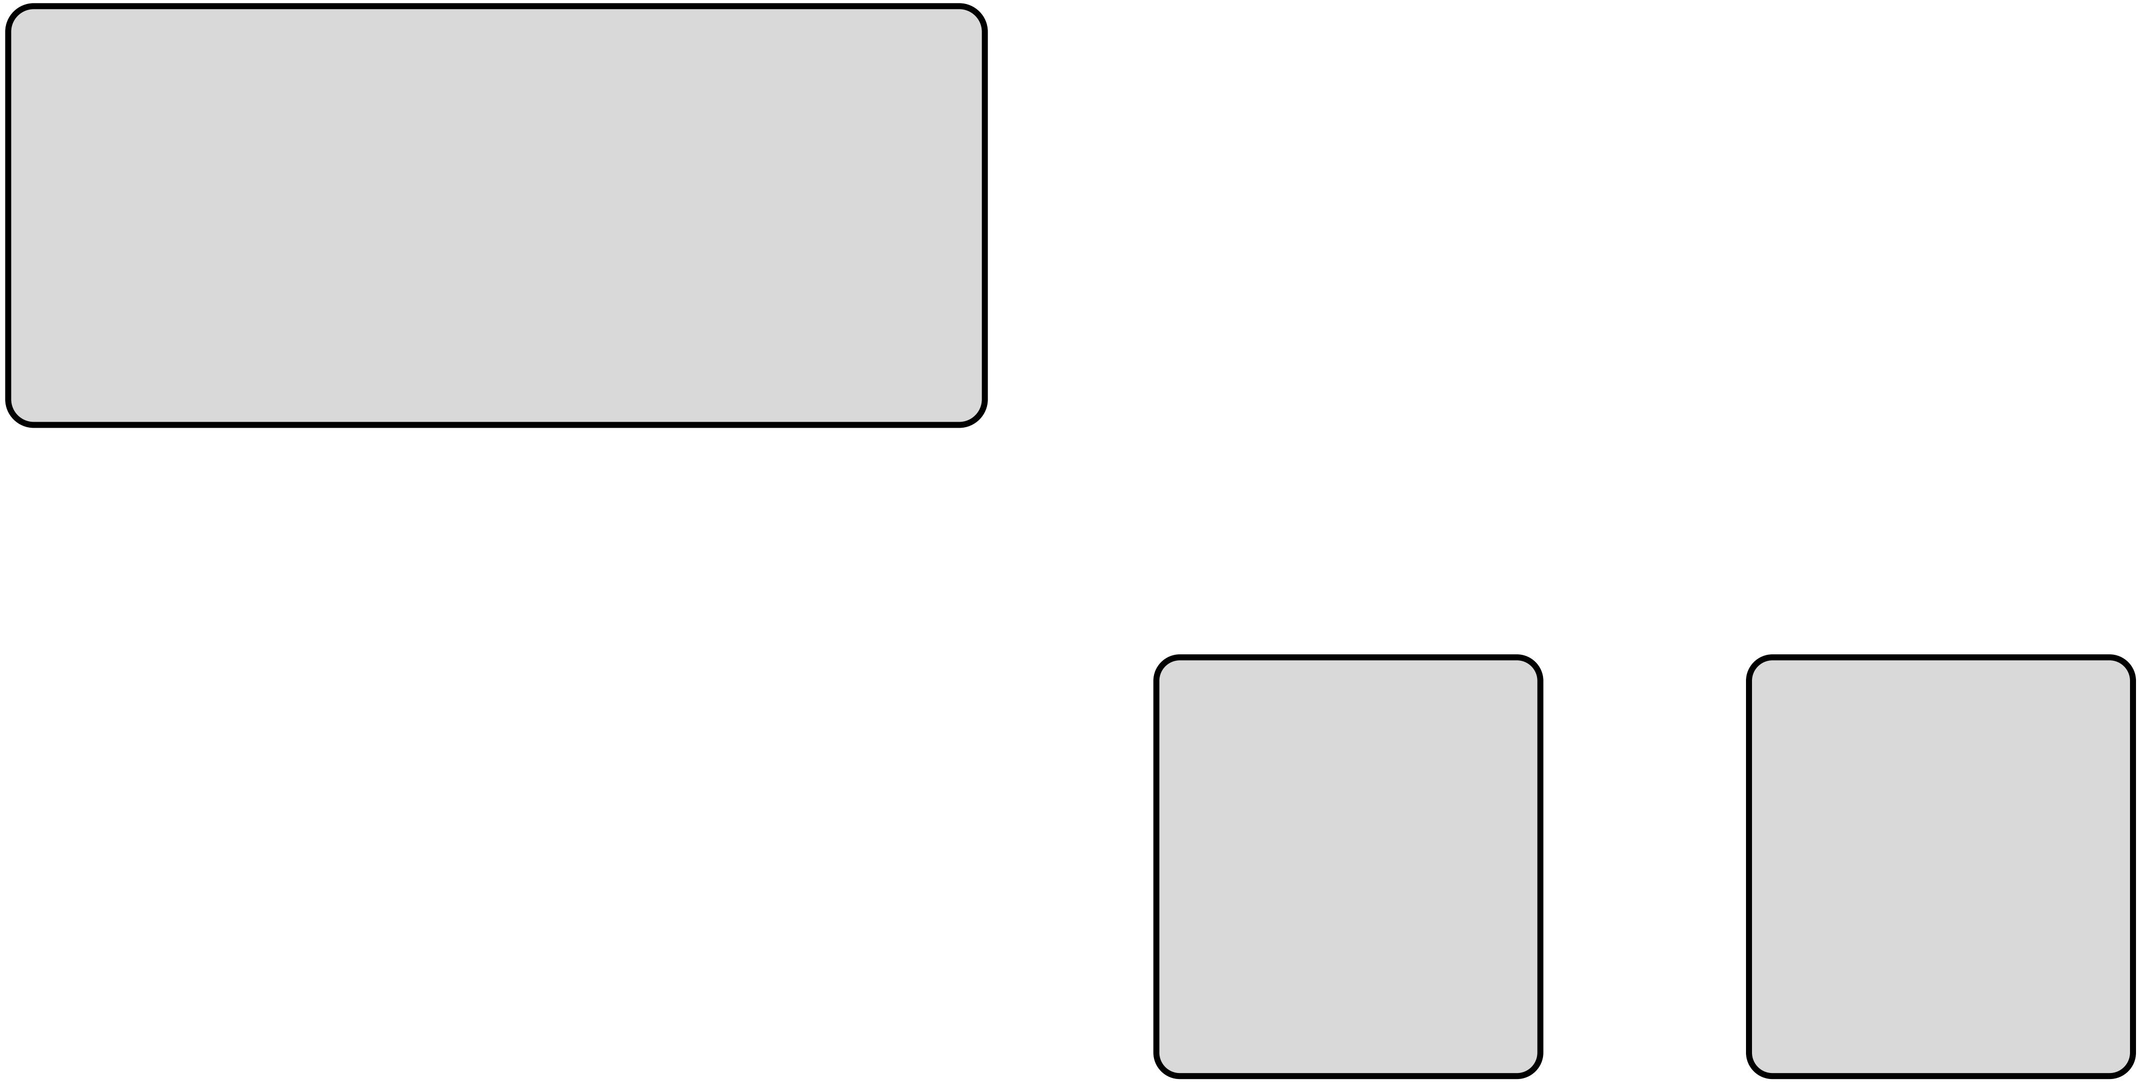

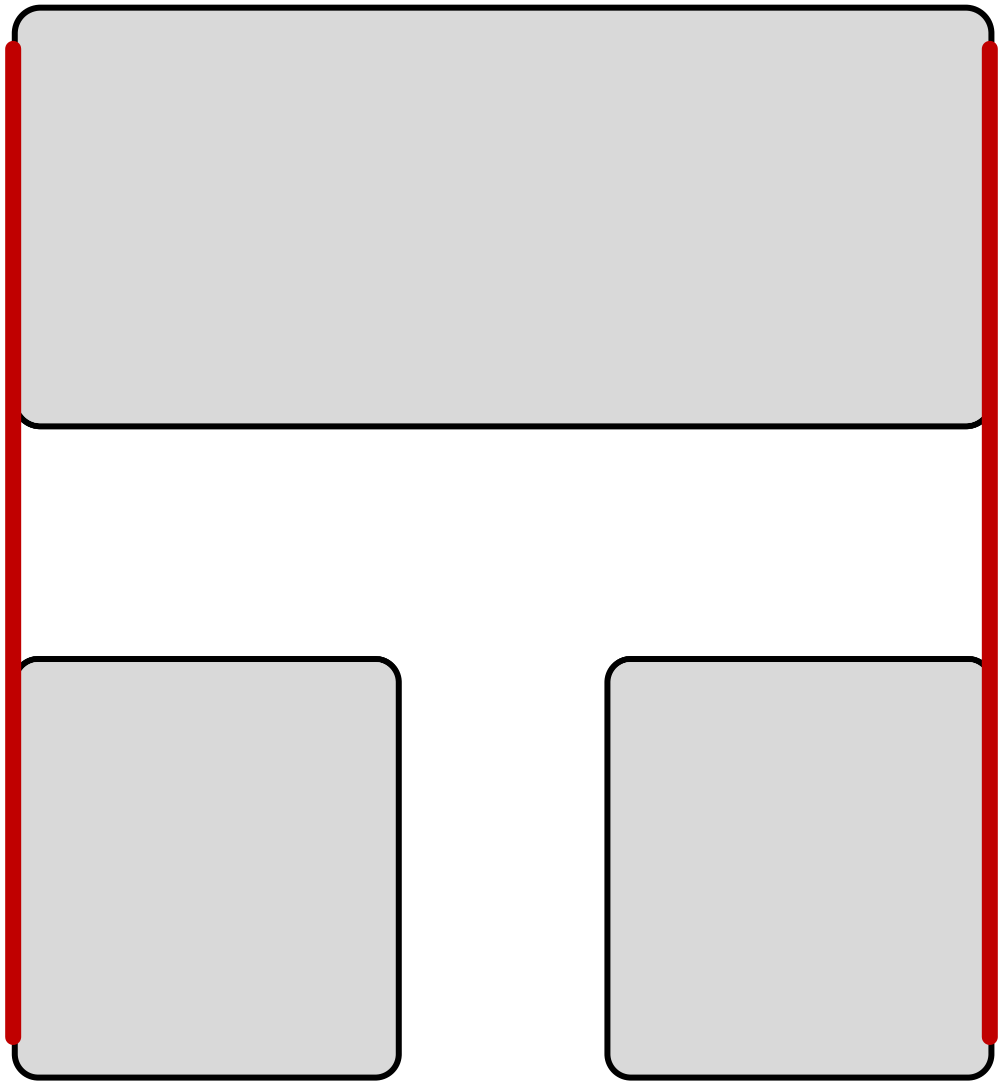

The one-is-more illusion

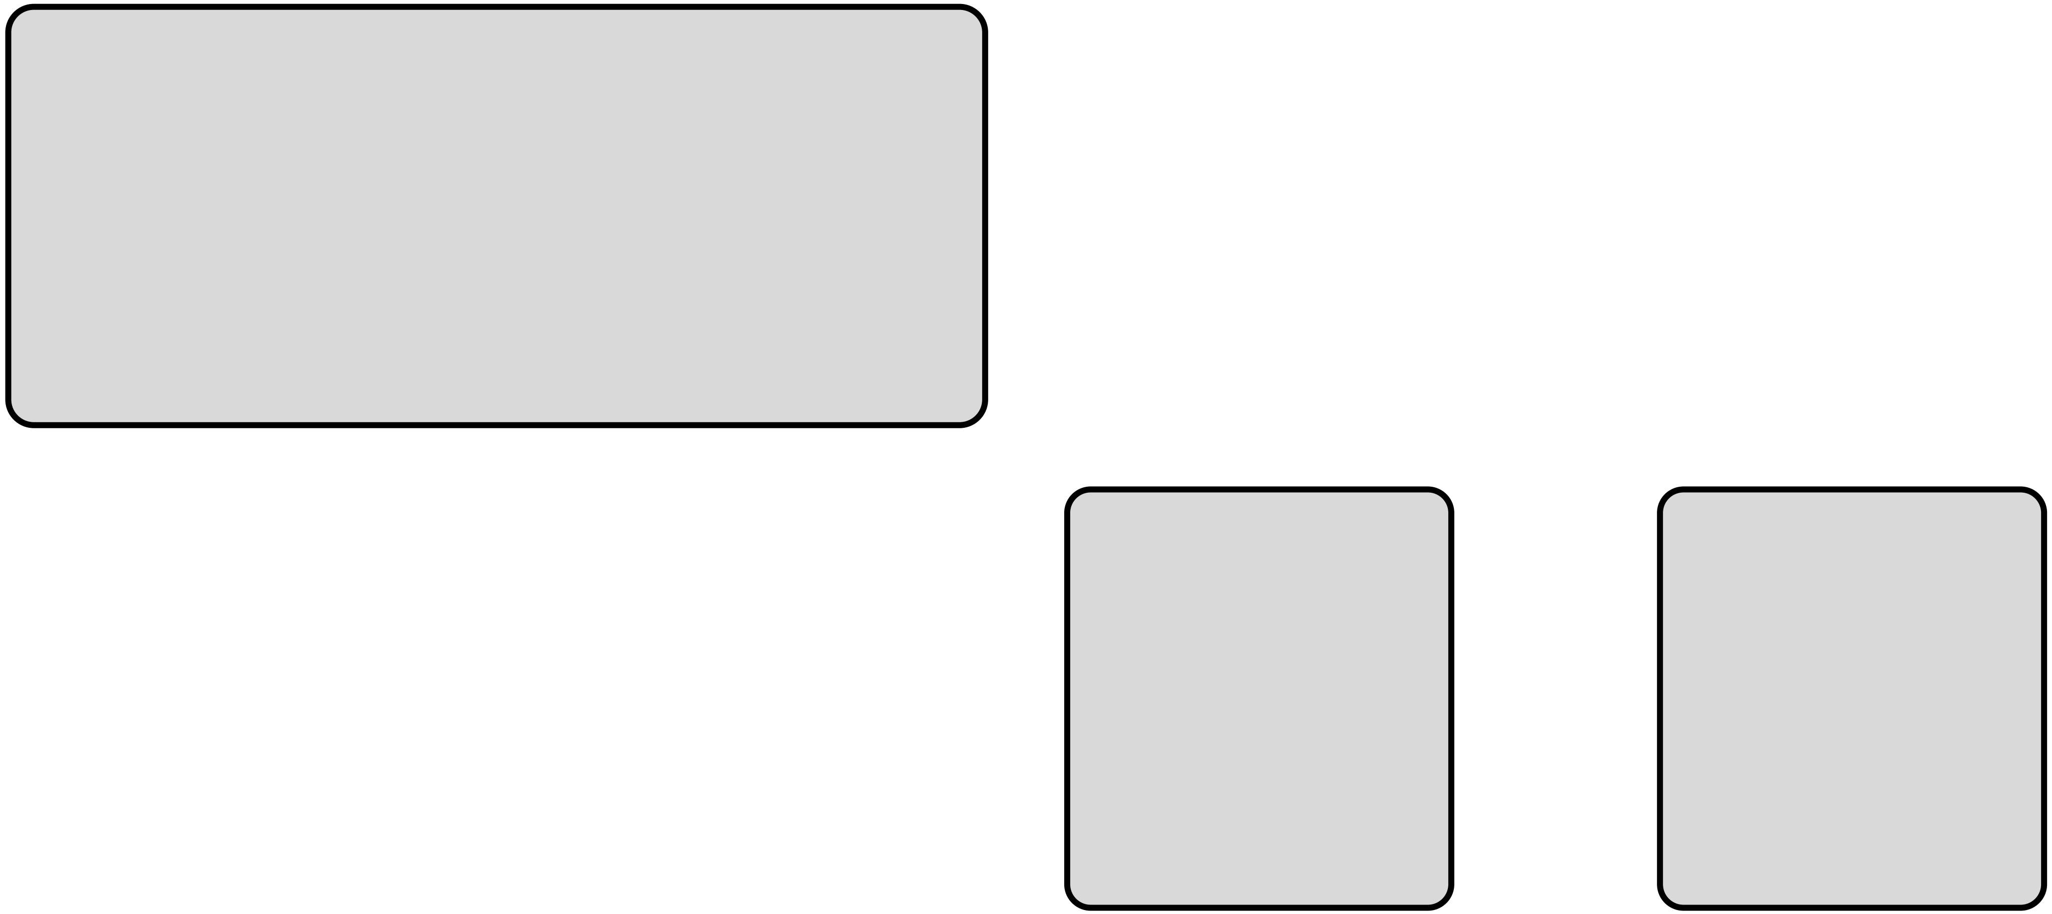

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Cognition

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

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

ARTICLE INFO

Keywords:
Spatial perception
Time perception
Segmentation
Object-based attention

ABSTRACT

We distinguish between discrete objects and continuous entities in categorization and language, but might we actually *see* such stimuli differently? Here we report the *one-is-more illusion*, wherein 'objecthood' changes what we perceive in an unexpected way. Across many variations and tasks, observers perceived a single continuous object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Cognition

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

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

ARTICLE INFO

Keywords:
Spatial perception
Time perception
Segmentation
Object-based attention

ABSTRACT

We distinguish between discrete objects and continuous entities in categorization and language, but might we actually *see* such stimuli differently? Here we report the *one-is-more illusion*, wherein 'objecthood' changes what we perceive in an unexpected way. Across many variations and tasks, observers perceived a single continuous object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

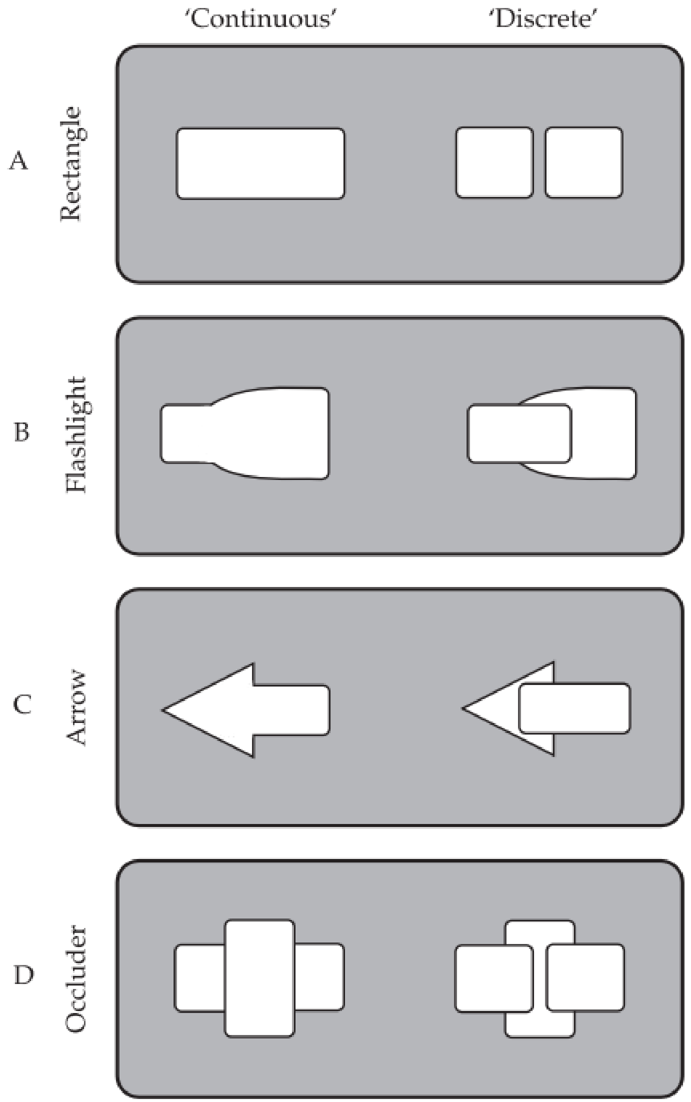

Fig. 2. Depiction of the four shape contrasts tested in Experiments 1a and 1b. In each case, the entities in each row are equally wide.

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Cognition

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

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

ARTICLE INFO

Keywords:
Spatial perception
Time perception
Segmentation
Object-based attention

ABSTRACT

We distinguish between discrete objects and continuous entities in categorization and language, but might we actually *see* such stimuli differently? Here we report the *one-is-more illusion*, wherein 'objecthood' changes what we perceive in an unexpected way. Across many variations and tasks, observers perceived a single continuous object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

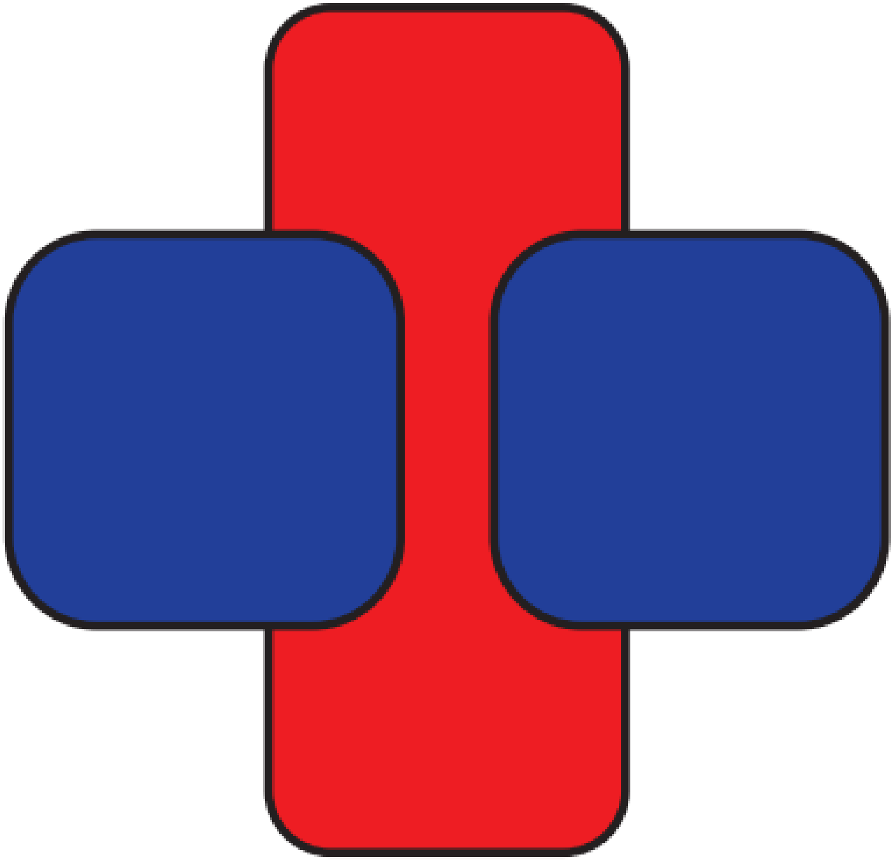

Illusions of **size** are more common than you think…

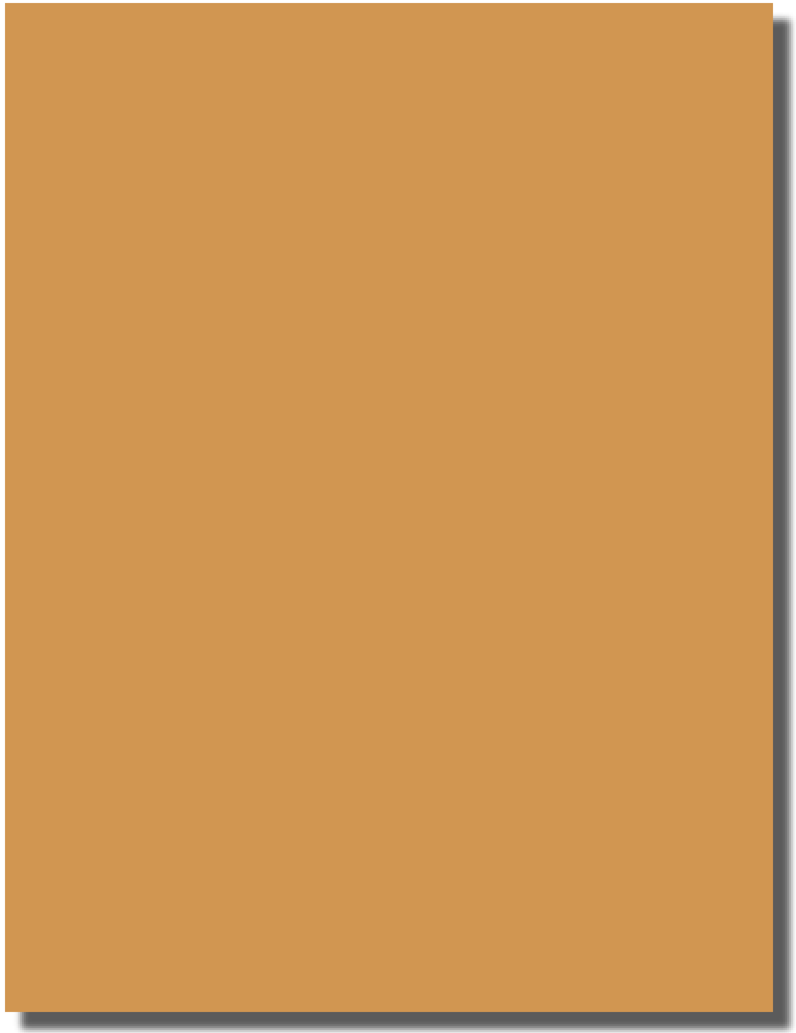

The left page is larger than the right page, **by how much**?

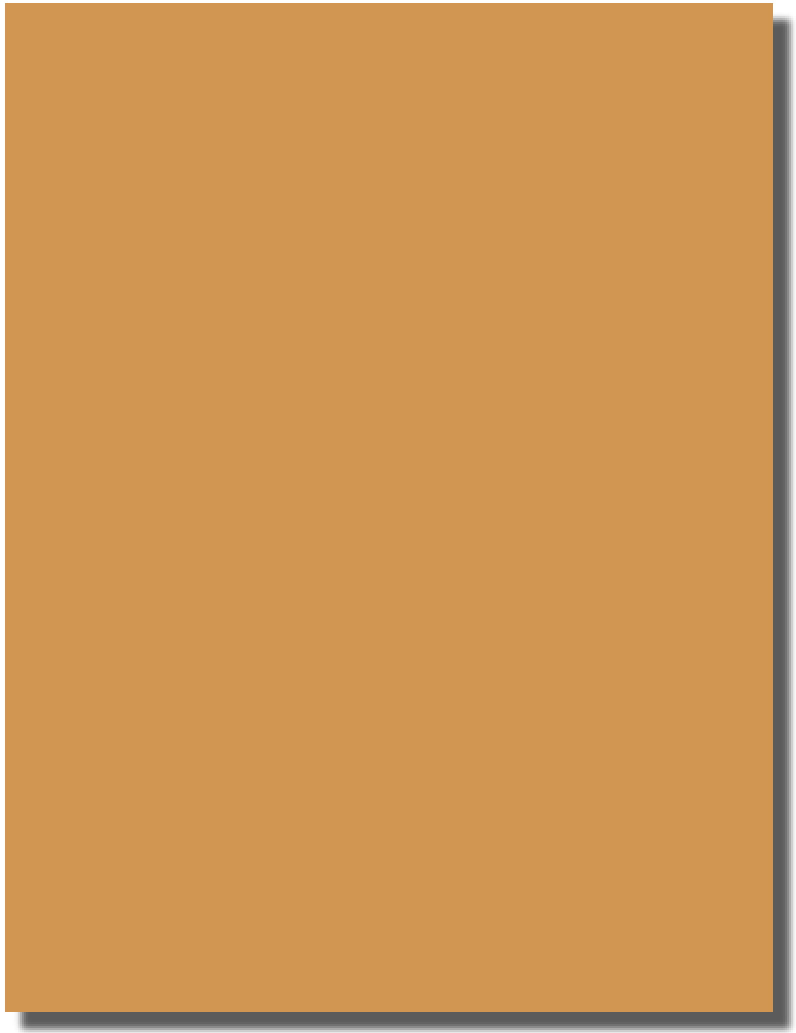

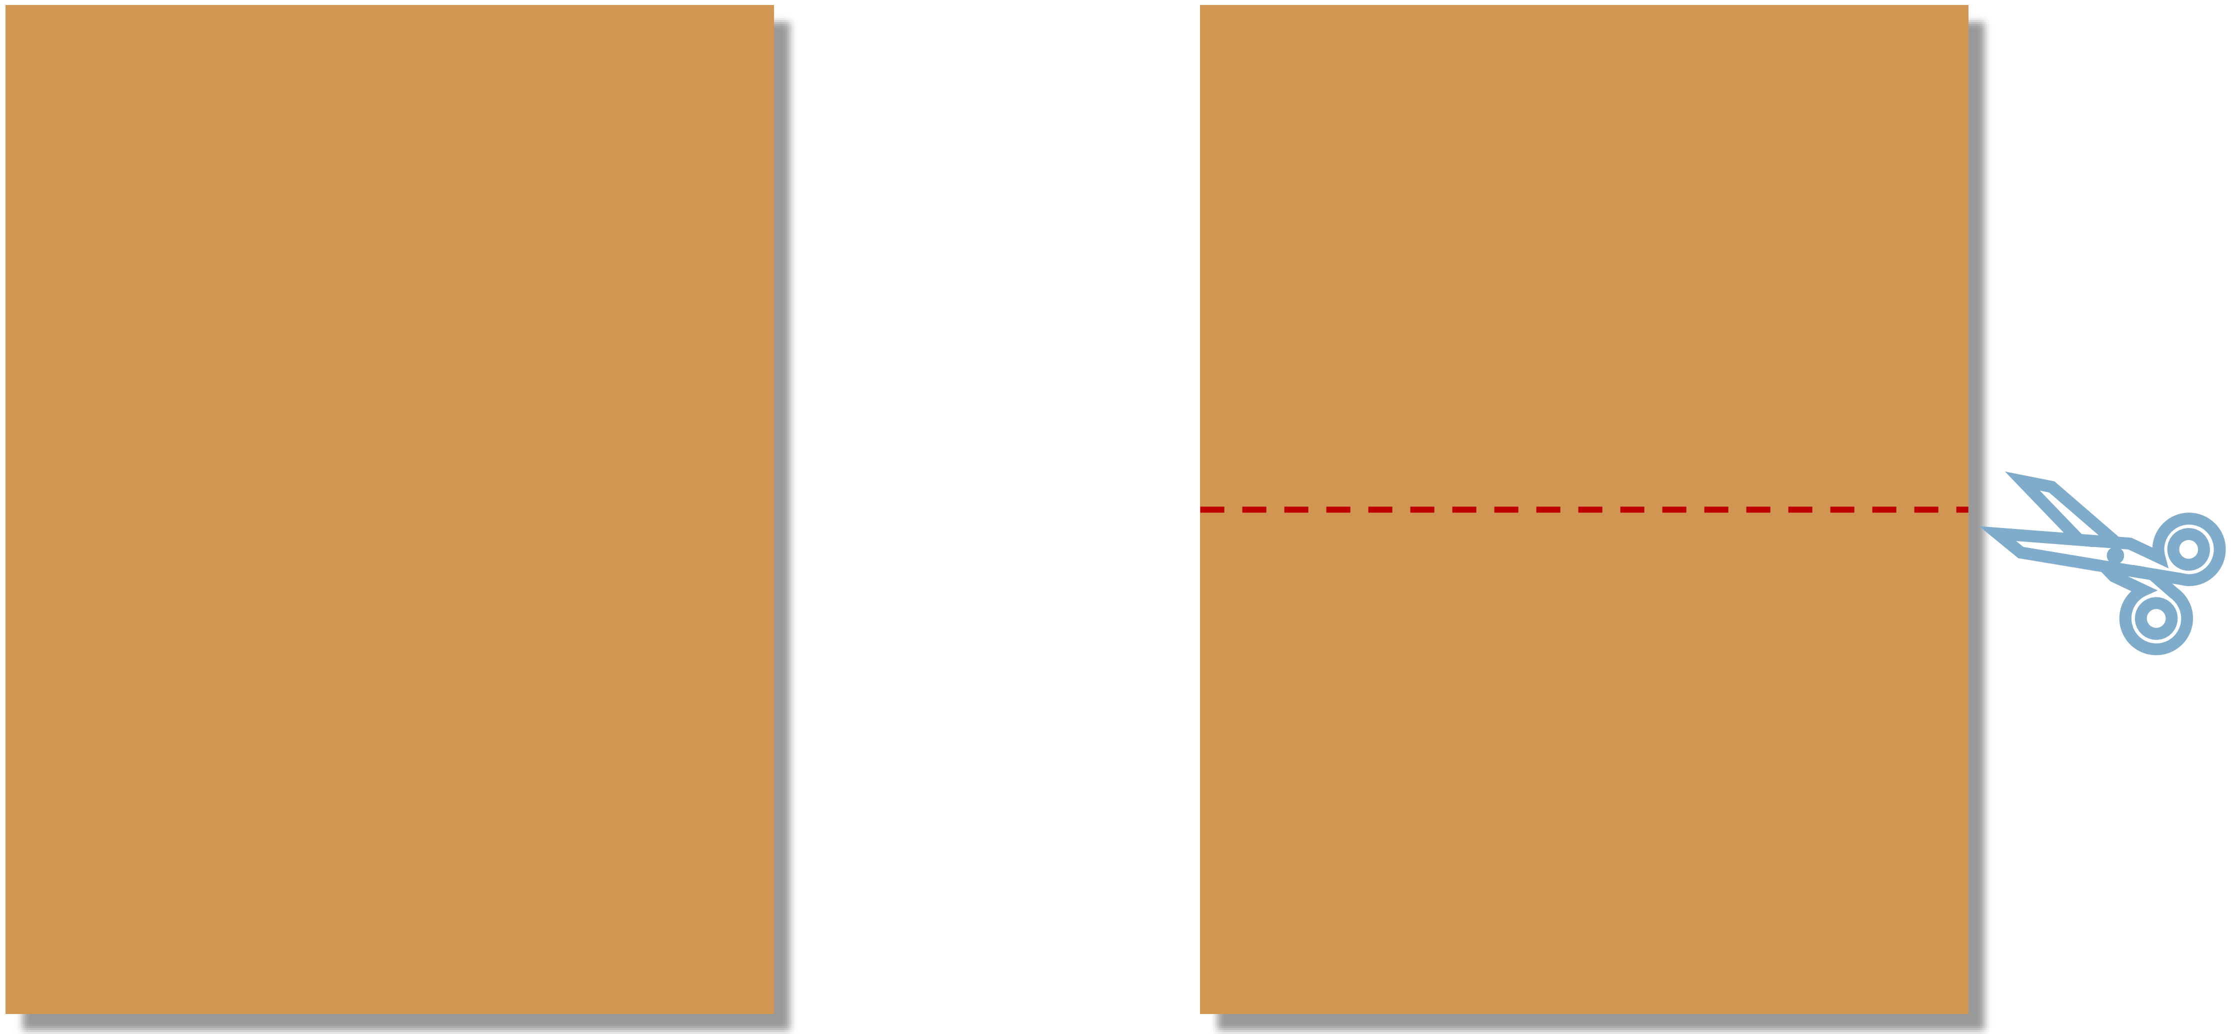

The left page is **twice as large** as the right page…

…or is it?

"A potential explanation might be the **general incapability of accurately comparing more than one geometrical dimension at once** — in everyday life, we solve this perceptual-cognitive bottleneck by reducing the complexity of such a task via aligning parts with same lengths."

!!

– Carbon, 2016

Which has more cumulative area?

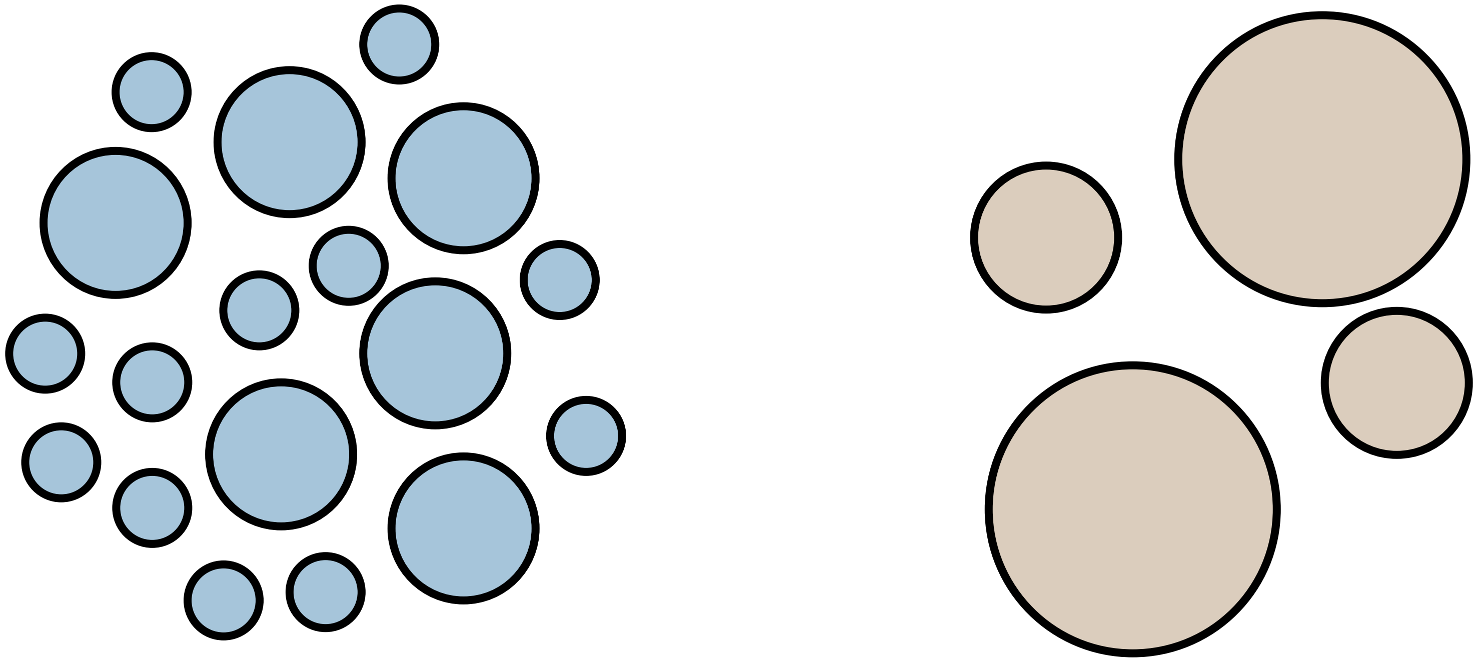

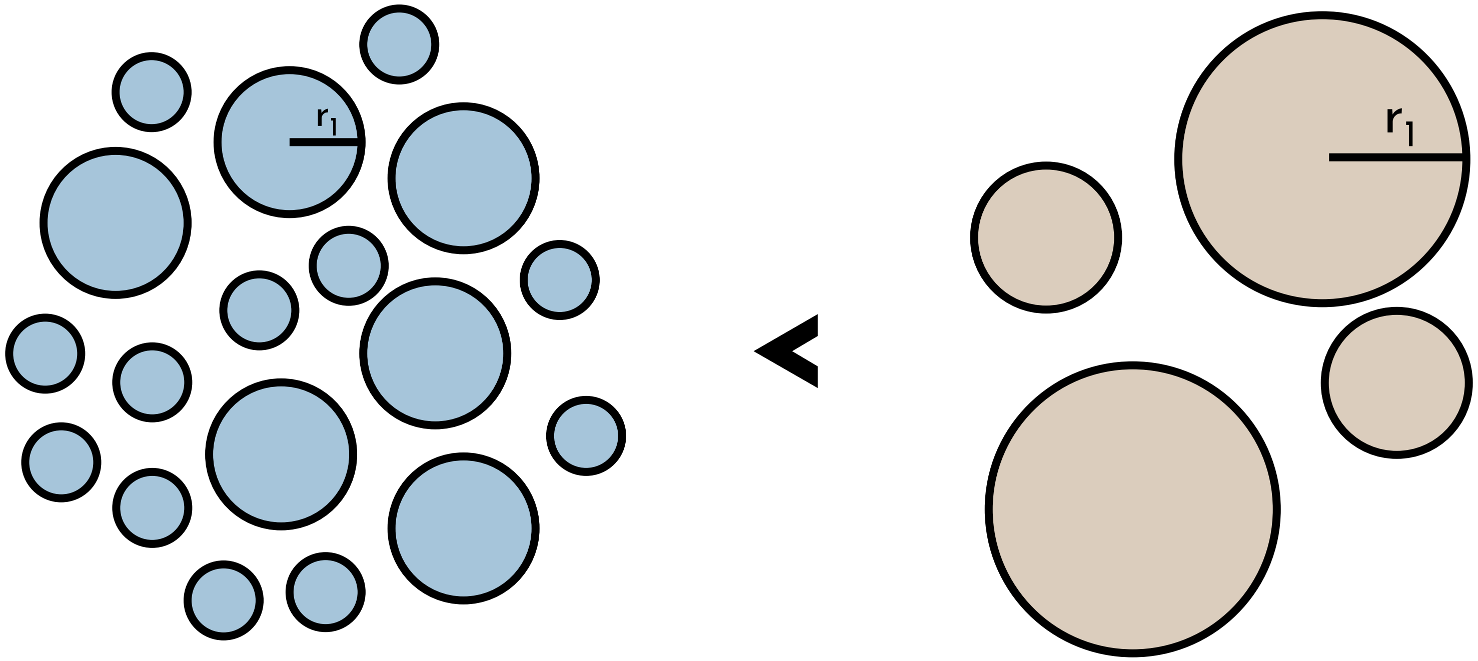

Area = $\pi r_1^2 + \pi r_2^2 + \ldots + \pi r_n^2$

Area = $\pi r_1^2 + \pi r_2^2 + \pi r_3^2 + \pi r_4^2$

What about perceived area?

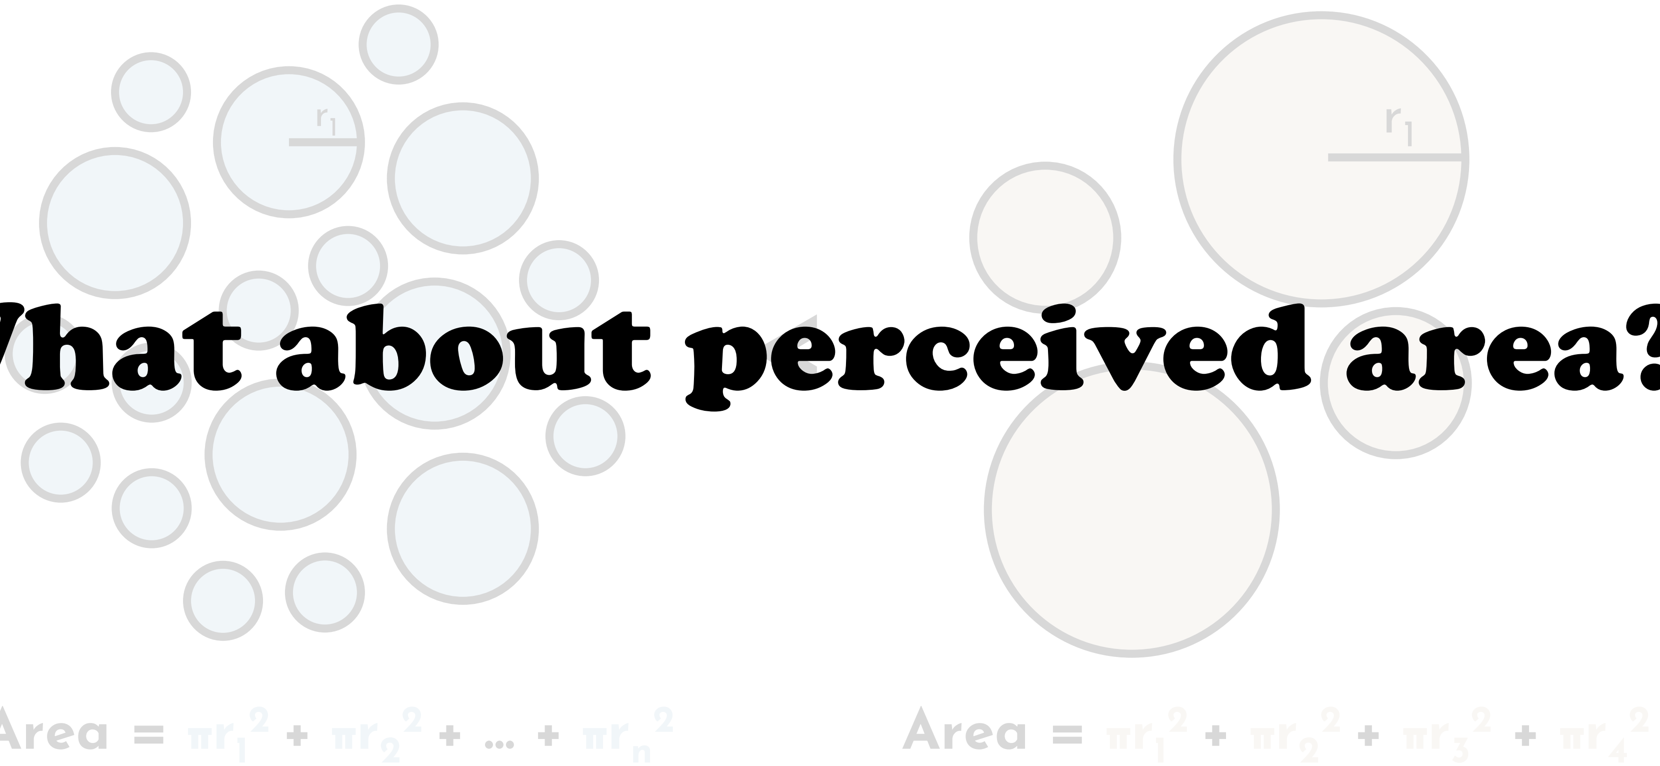

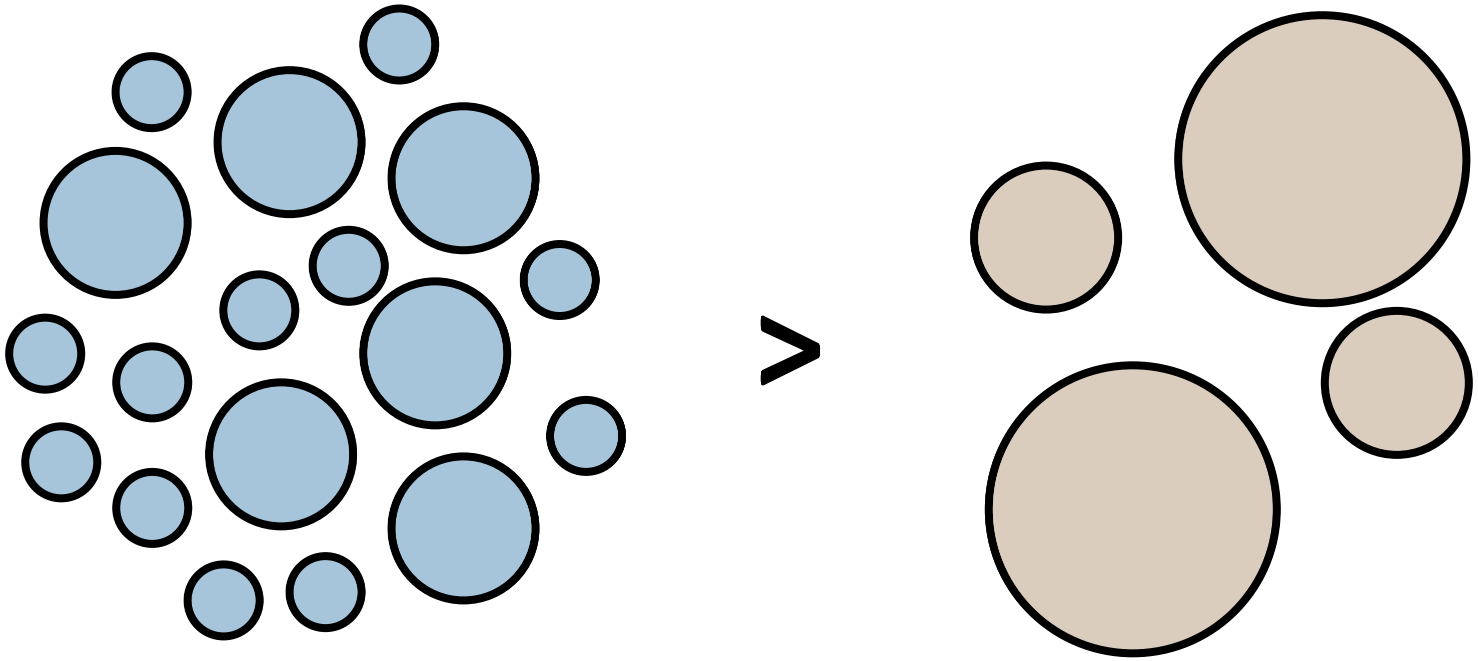

~~Area = $\pi r_1^2 + \pi r_2^2 + \ldots + \pi r_n^2$~~

Area = $l_1+w_1 + l_2+w_2 + \ldots + l_n+w_n$

~~Area = $\pi r_1^2 + \pi r_2^2 + \pi r_3^2 + \pi r_4^2$~~

Area = $l_1+w_1 + l_2+w_2 + l_3+w_3 + l_4+w_4$

The Additive-Area Heuristic: An Efficient but Illusory Means of Visual Area Approximation

Psychological Science
1–9

DOI: 10.1177/0956797619831617
www.psychologicalscience.org/PS
SAGE

Sami R. Yousif and Frank C. Keil
Department of Psychology, Yale University

Abstract

How do we determine how much of something is present? A large body of research has investigated the mechanisms and consequences of number estimation, yet surprisingly little work has investigated area estimation. Indeed, area is often treated as a pesky confound in the study of number. Here, we describe the *additive-area heuristic*, a means of rapidly estimating visual area that results in substantial distortions of perceived area in many contexts, visible even in simple demonstrations. We show that when we controlled for additive area, observers were unable to discriminate on the basis of true area, per se, and that these results could not be explained by other spatial dimensions. These findings reflect a powerful perceptual illusion in their own right but also have implications for other work, namely, that which relies on area controls to support claims about number estimation. We discuss several areas of research potentially affected by these findings.

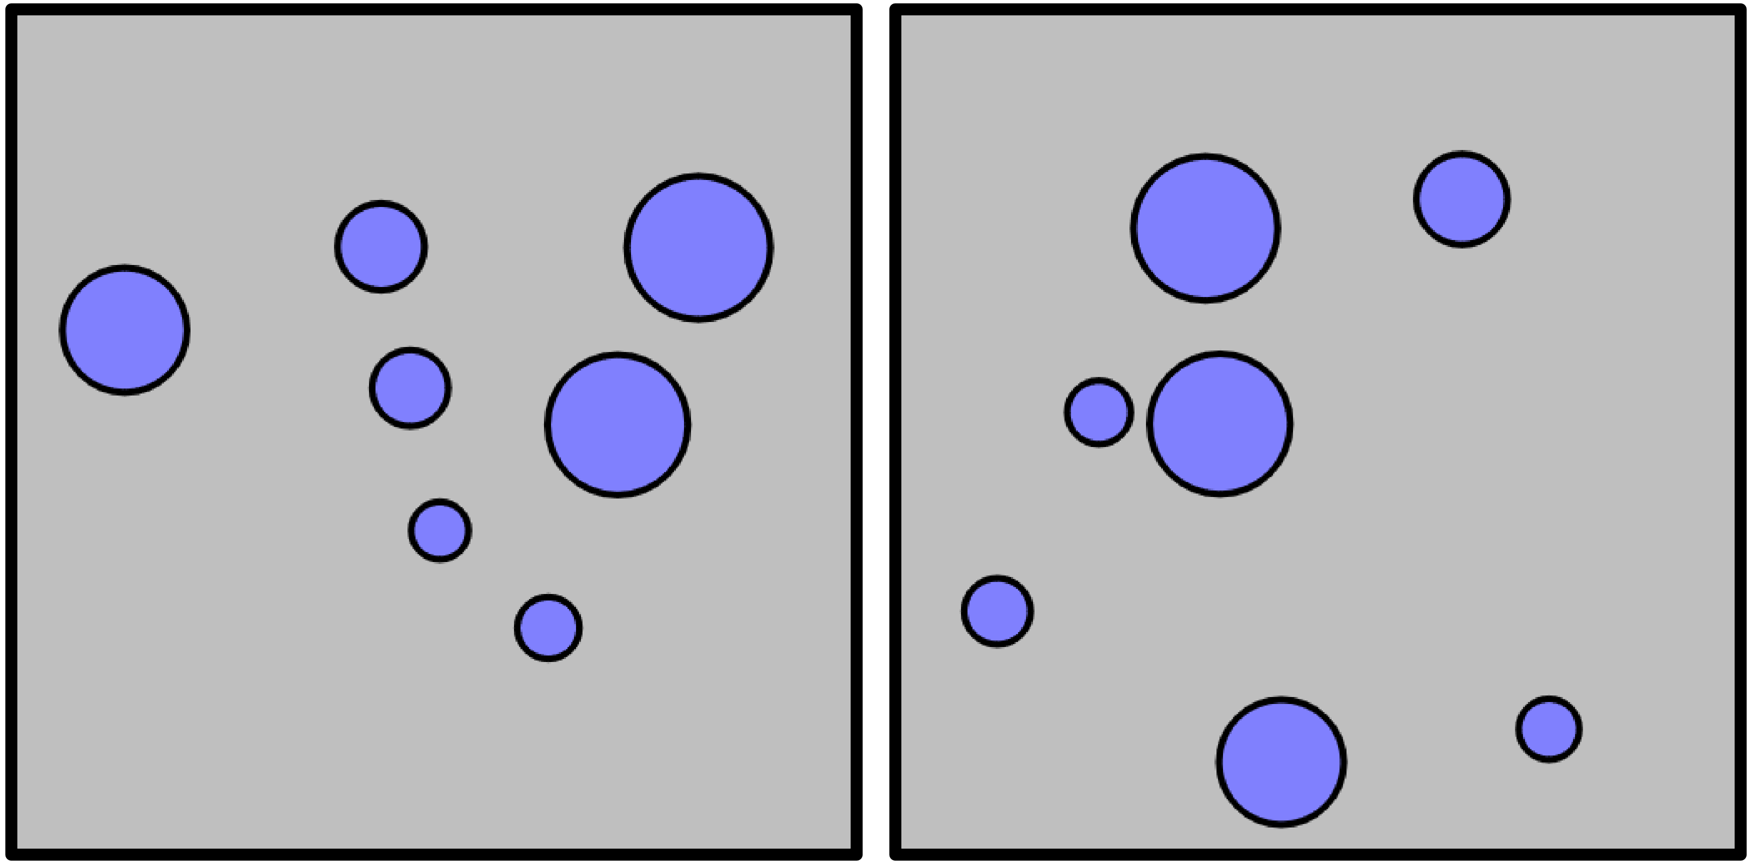

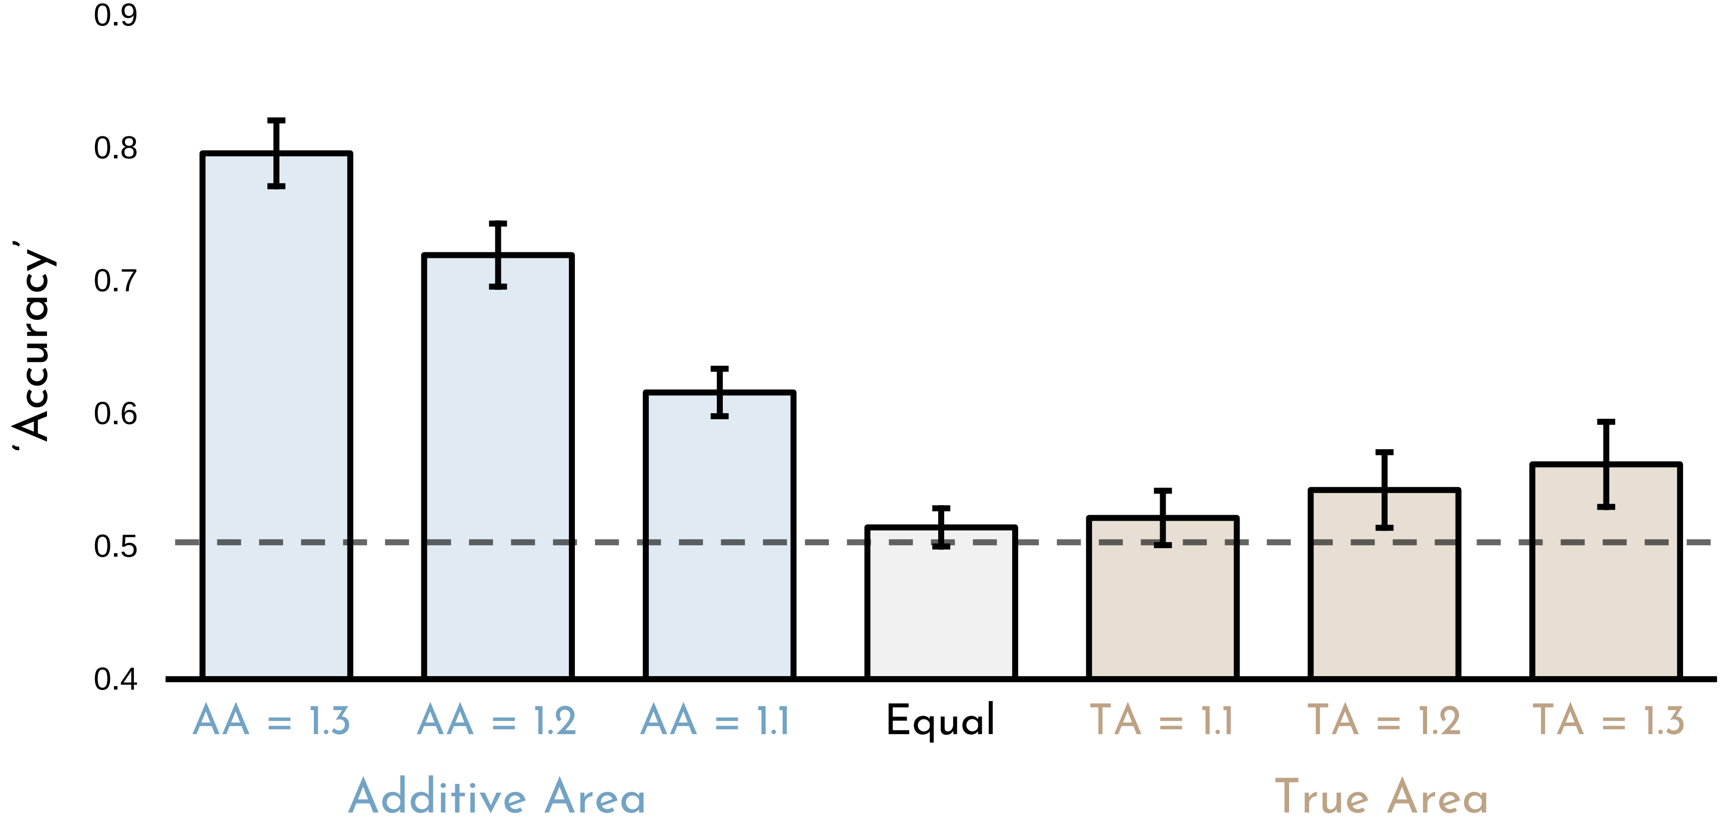

'Accuracy'
0.9
0.8
0.7
0.6
0.5
0.4
AA = 1.3
AA = 1.2
AA = 1.1
Equal
TA = 1.1
TA = 1.2
TA = 1.3
Additive Area
True Area

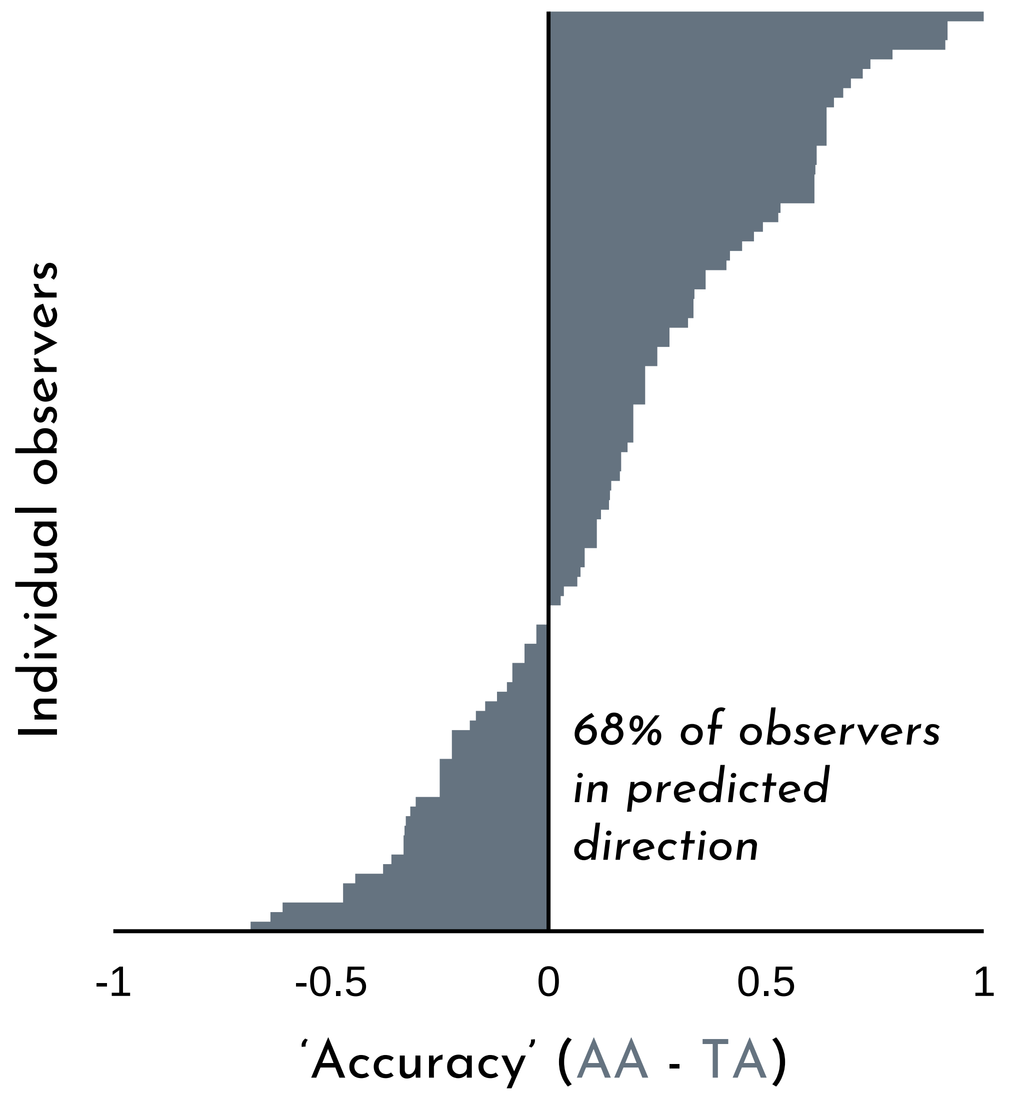
Individual observers
68% of observers
in predicted
direction
-1
-0.5
0
0.5
1
'Accuracy' (AA - TA)

The same rules apply in a 3D world!

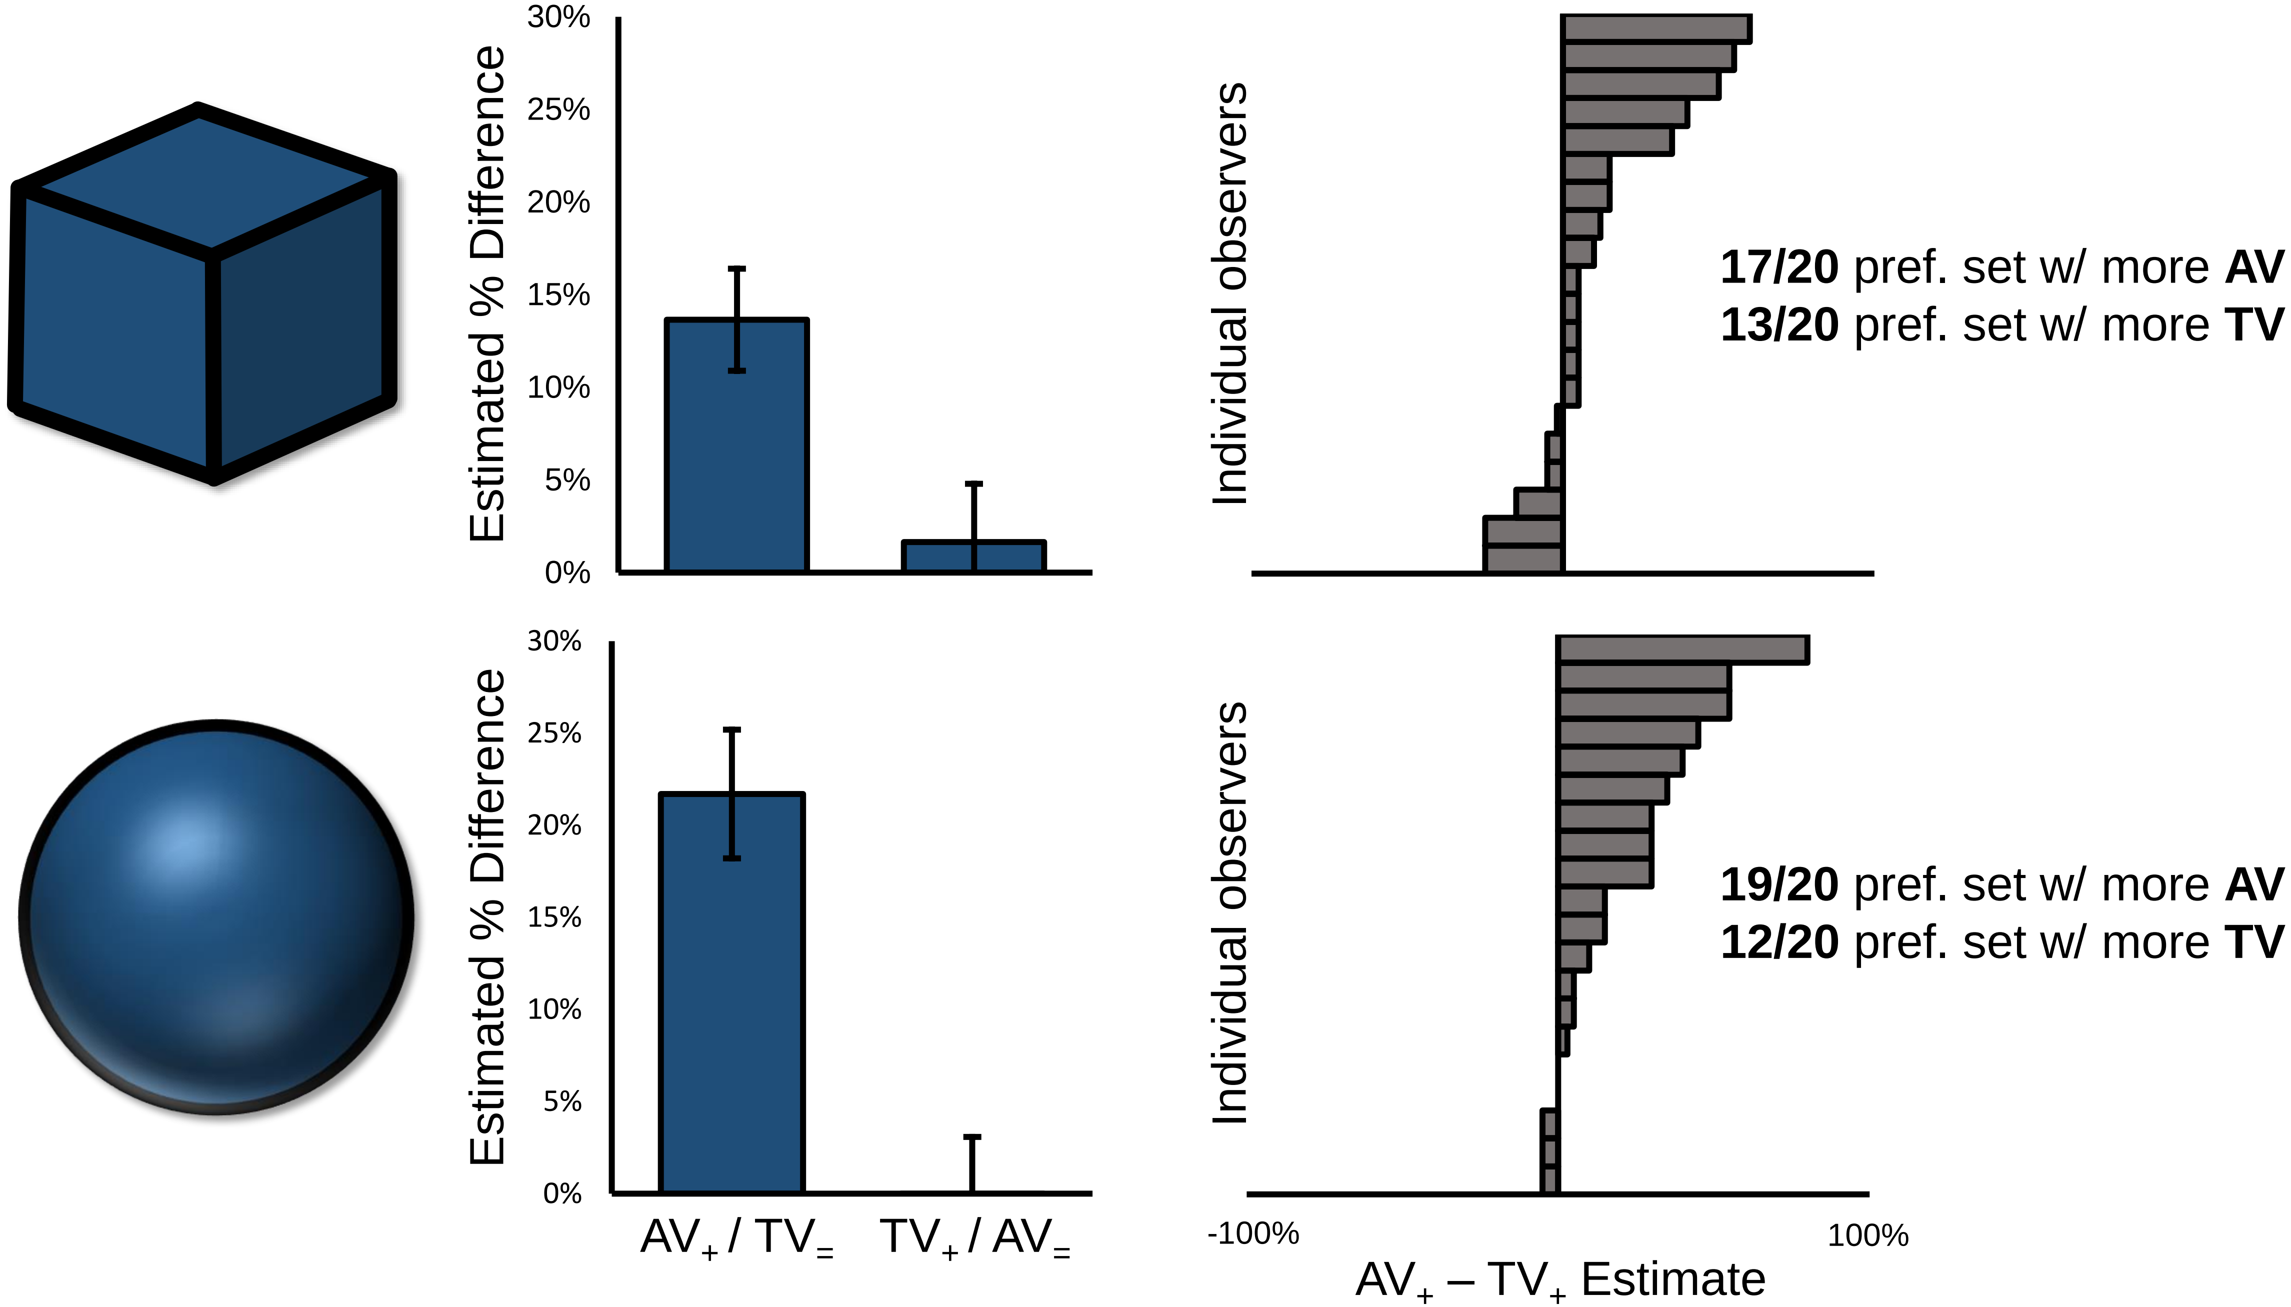
30%
25%
20%
15%
10%
5%
0%
Estimated % Difference
Individual observers
17/20 pref. set w/ more AV
13/20 pref. set w/ more TV
19/20 pref. set w/ more AV
12/20 pref. set w/ more TV
AV+ / TV=
TV+ / AV=
-100%
100%
AV+ – TV+ Estimate

Short Report

PERCEPTION

Perception
2021, Vol. 50(5) 462–469

DOI: 10.1177/03010066211003746
journals.sagepub.com/home/pec
SAGE

A Ubiquitous Illusion of Volume: Are Impressions of 3D Volume Captured by an "Additive Heuristic"?

Elizabeth Bennette, Frank C. Keil and Sami R. Yousif
Yale University, United States

Abstract

Several empirical approaches have attempted to explain perception of 2D and 3D size. While these approaches have documented interesting perceptual effects, they fail to offer a compelling, general explanation of everyday size perception. Here, we offer one. Building on prior work documenting an "Additive Area Heuristic" by which observers estimate perceived area by summing objects' dimensions, we show that this same principle—an "additive heuristic"—explains impressions of 3D volume. Observers consistently discriminate sets that vary in "additive volume," even when there is no true difference; they also *fail* to discriminate sets that truly differ (even by amounts as much as 30%) when they are equated in "additive volume." These results suggest a failure to properly integrate multiple spatial dimensions, and frequent reliance on a perceptual heuristic instead.

Keywords

3D perception, perception, perceptual organization, spatial cognition, spatial vision

Date Received: 13 September 2020; accepted: 26 February 2021

In a classic demonstration, Piaget presented children with two identical glasses containing equal amounts of water. Water from one glass was then poured into another taller, skinnier glass, and children were asked which glass held more water. Famously, children select the taller glass—and this is seen as evidence that children fail to understand the conservation of volume between the two containers (Piaget, 1952). But what is the nature of this bias? Why do children perceive the taller glass as having more in the first place?

Corresponding author:
Sami R. Yousif, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, United States.
Email: sami.yousif@yale.edu

Illusions of color

Sami Ryan Yousif

the dress


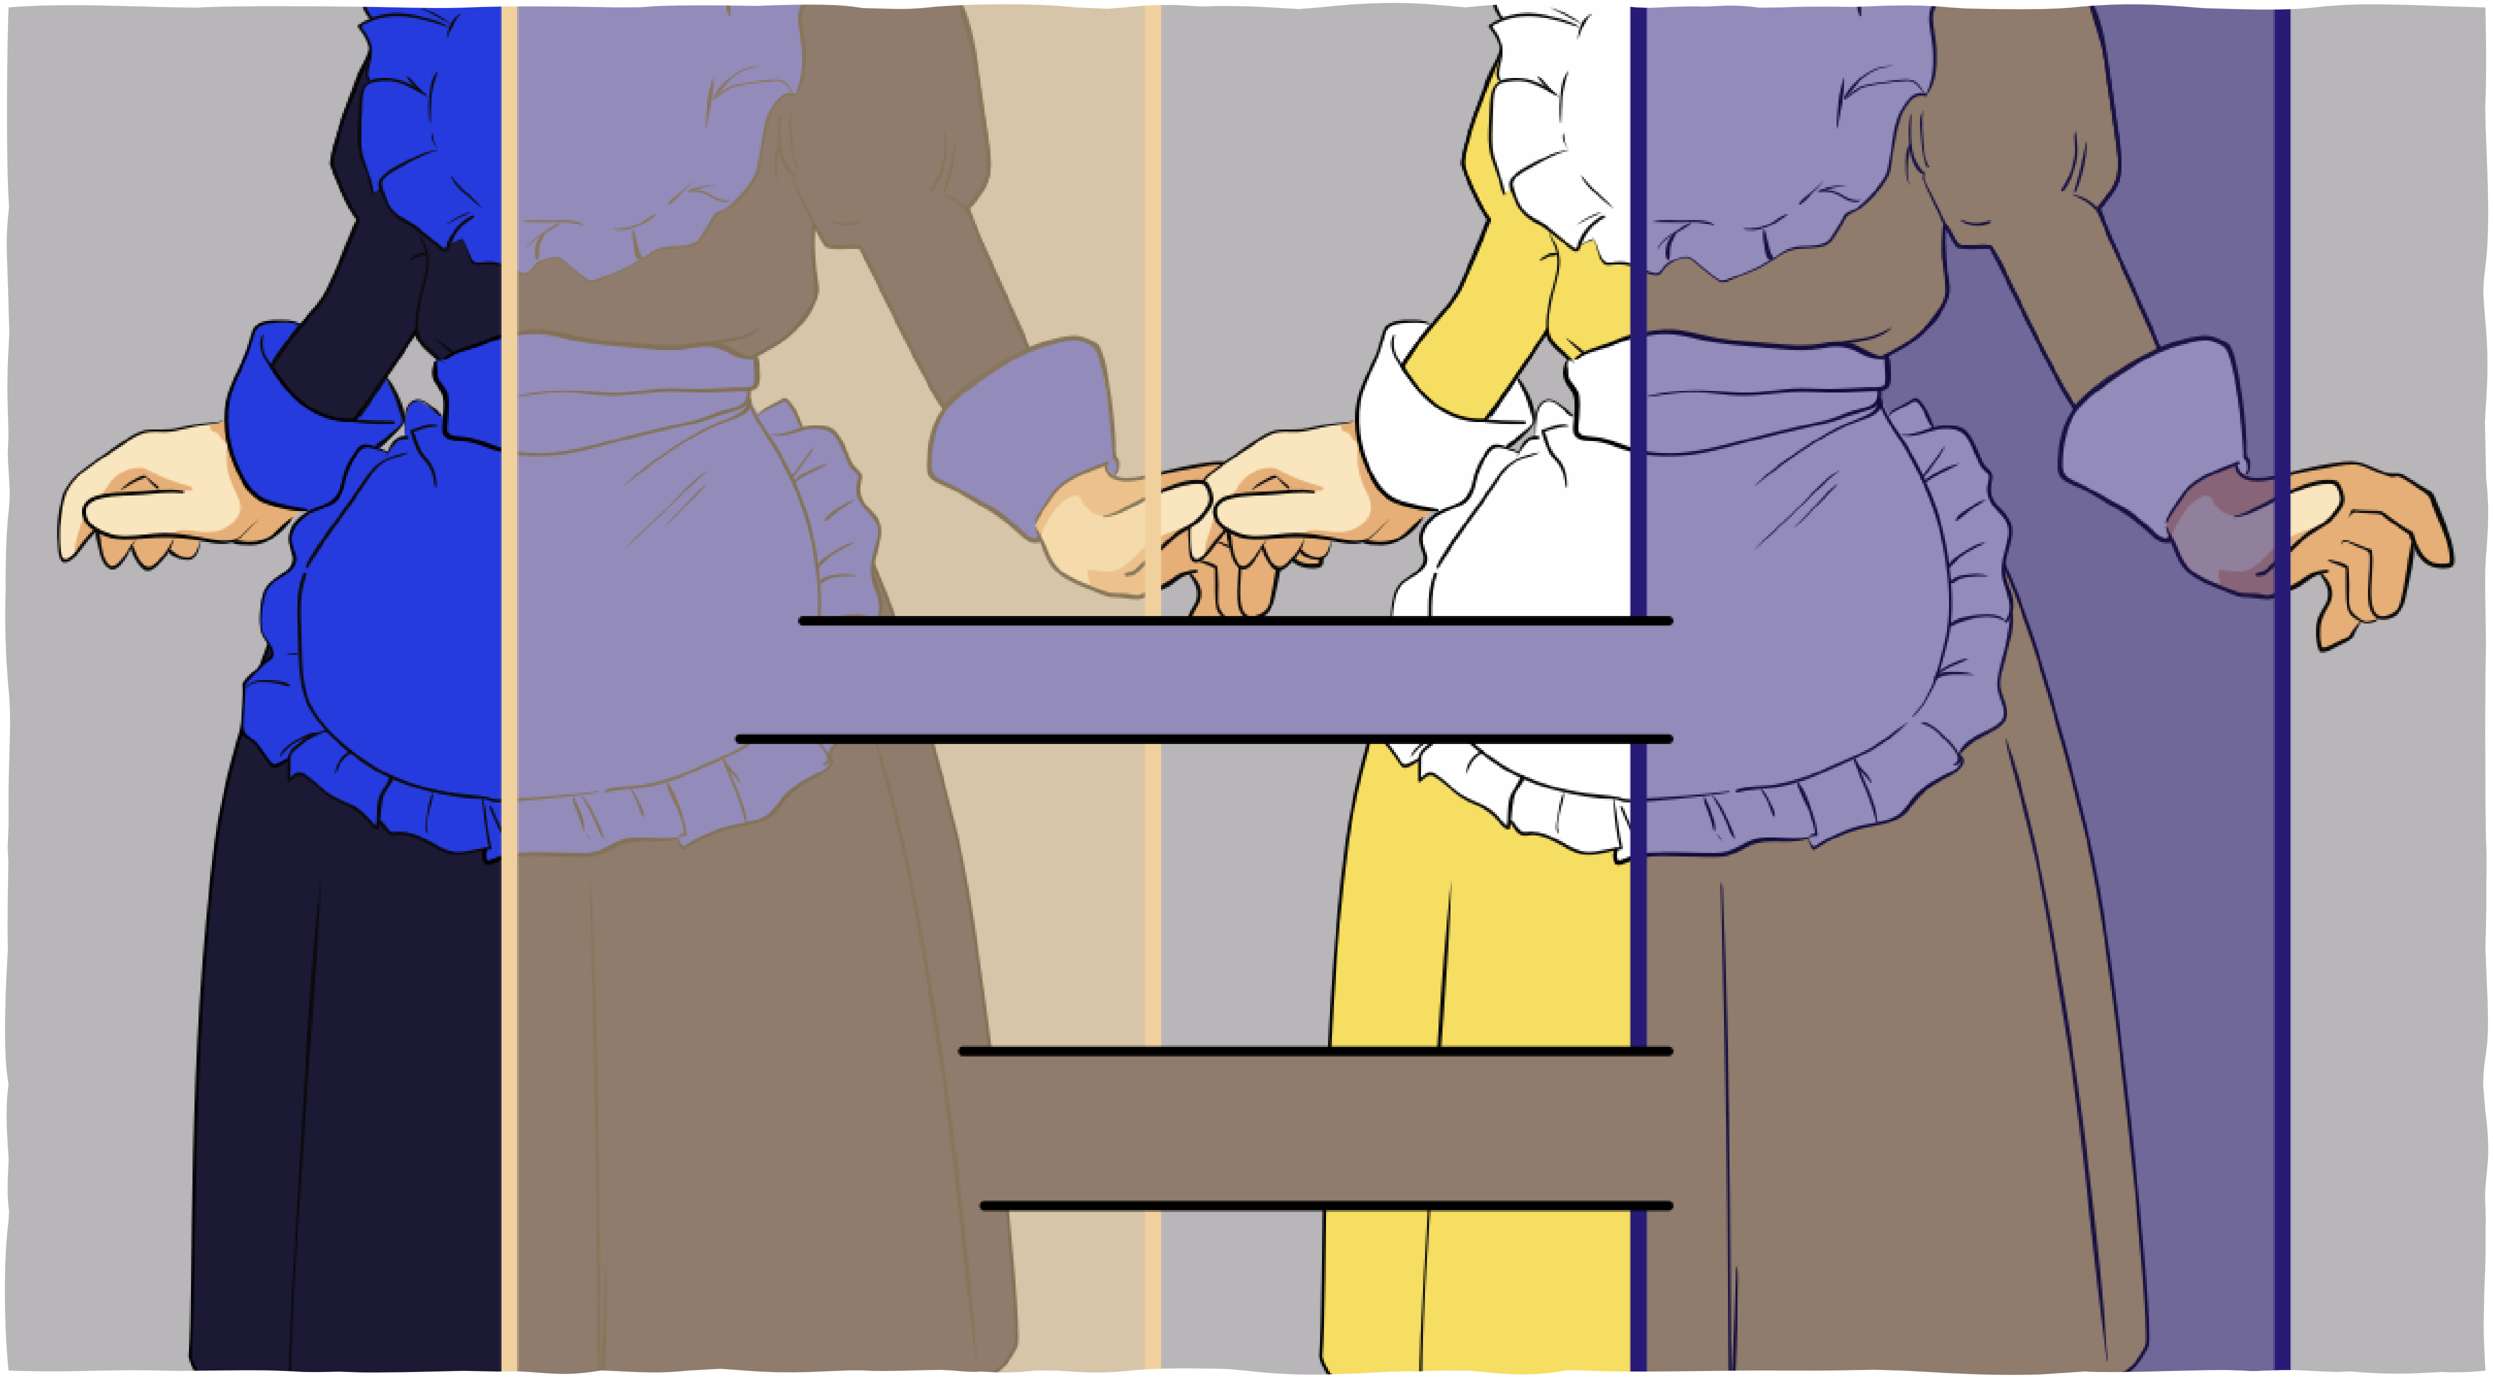

The visual system distinguishes between *light* and *color*.

unconscious inference

(Daddy) Hermann von Helmholtz

[Visual perception] is the result of an ***unconscious*** *and involuntary activity; and for this very reason it strikes our consciousness as a foreign and overpowering force of nature.*

(1856/1925, p. 117)

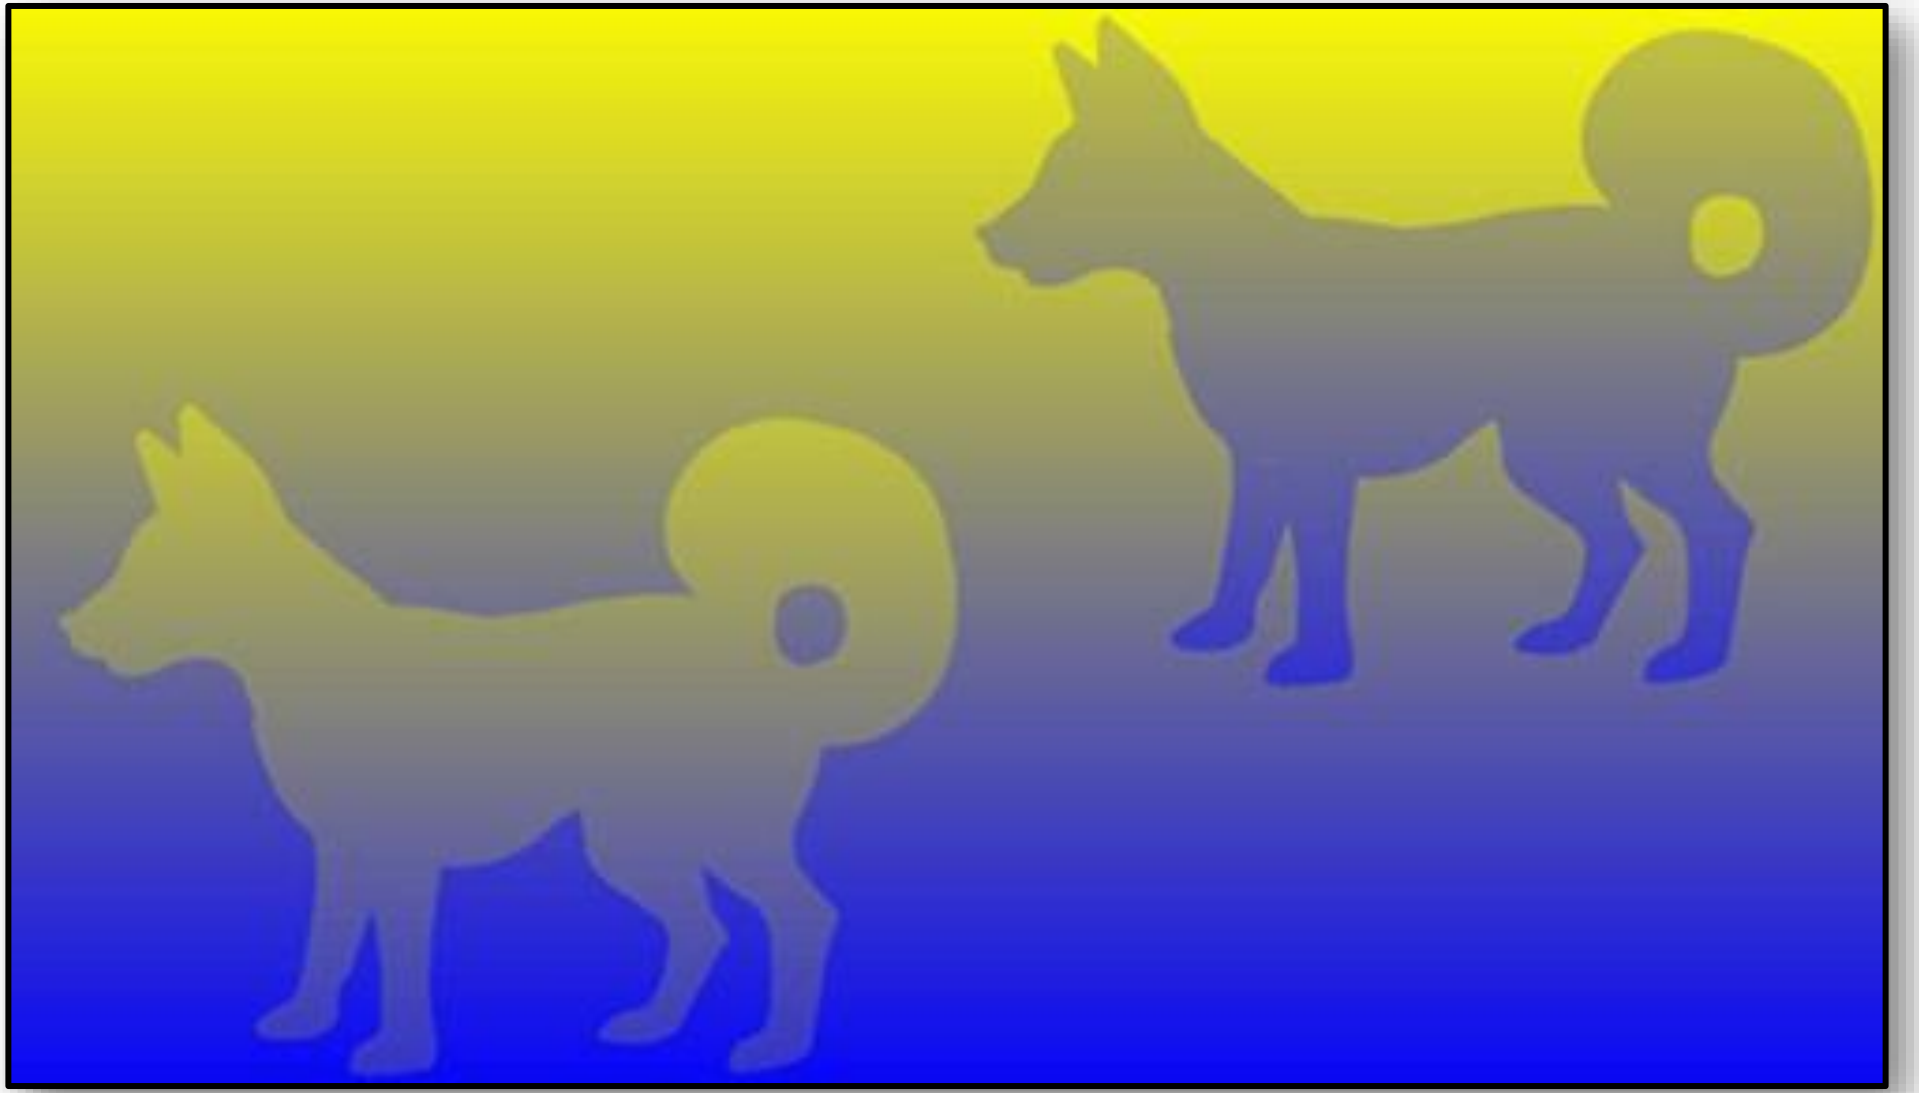

A simple inference...

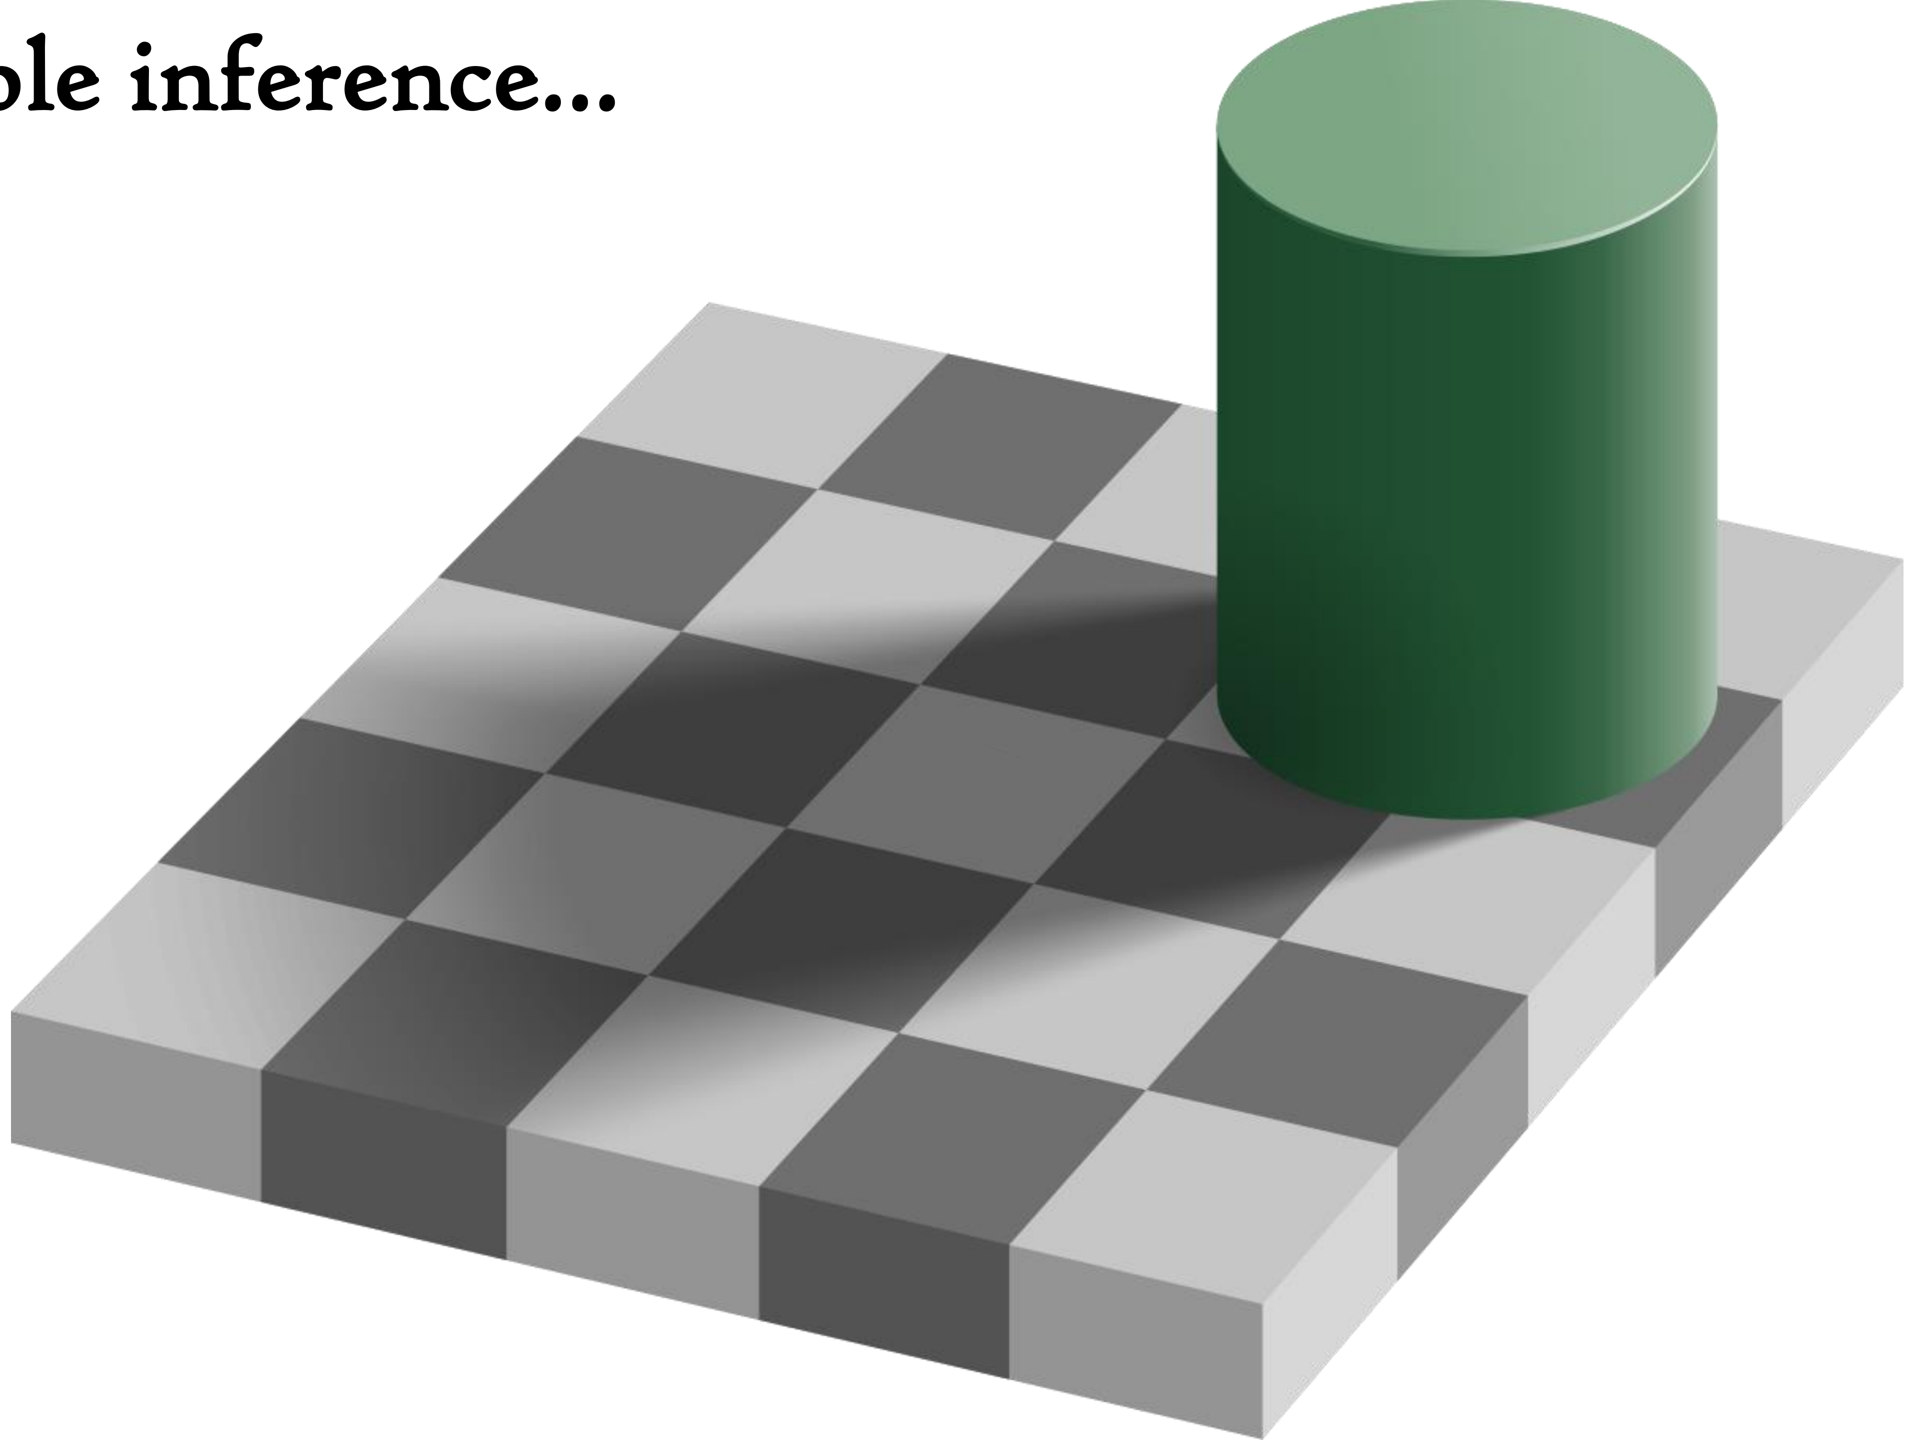

!!

Color constancy

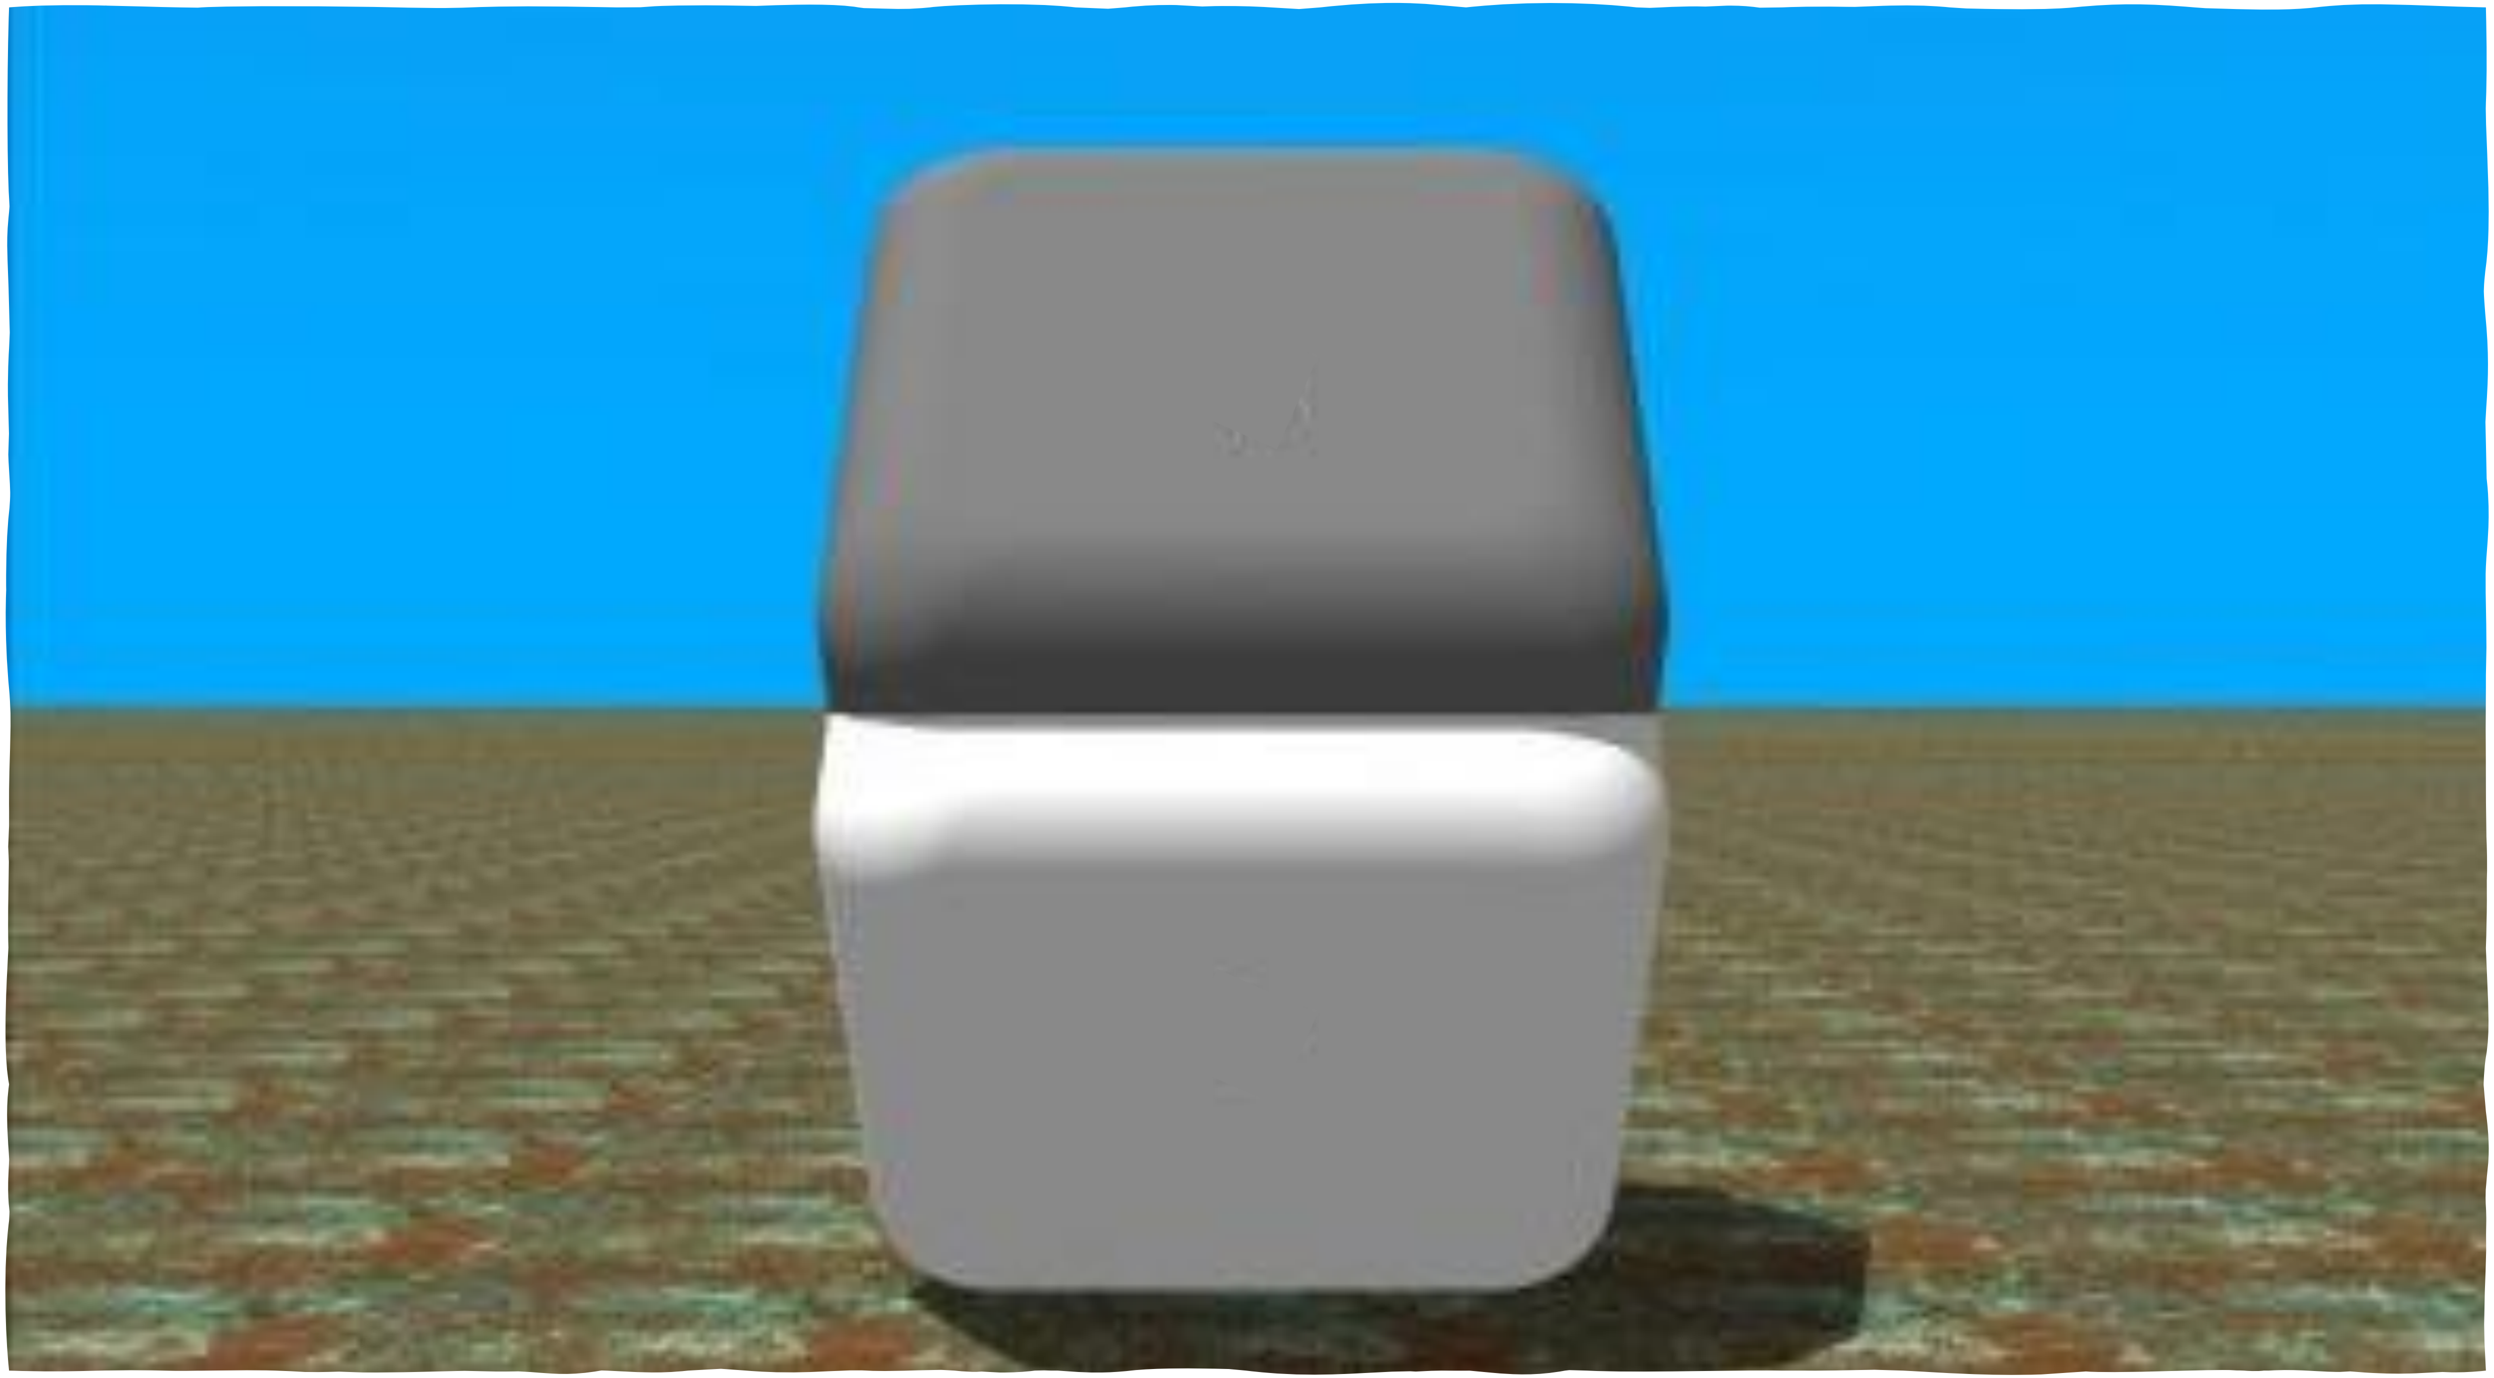

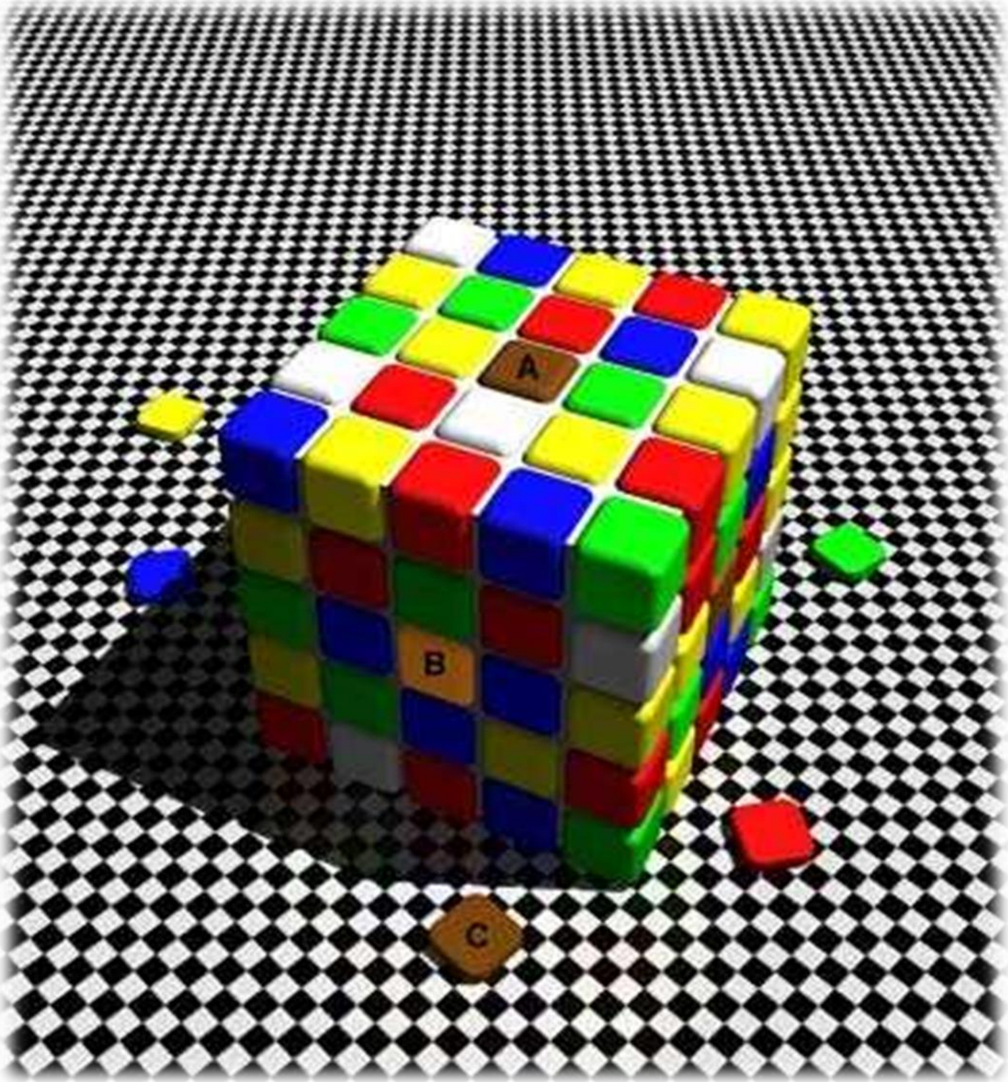
A
B
C

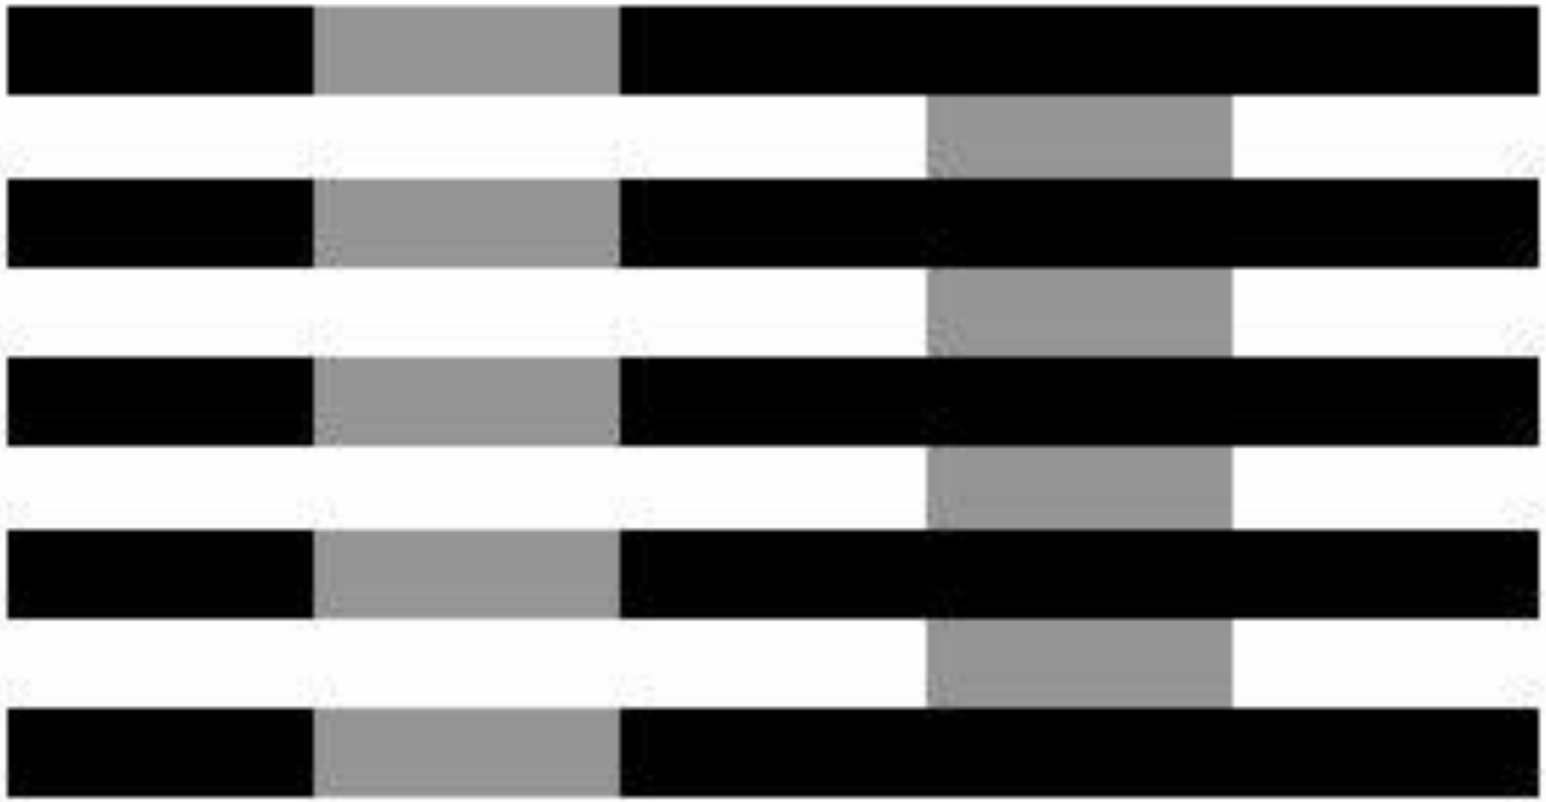

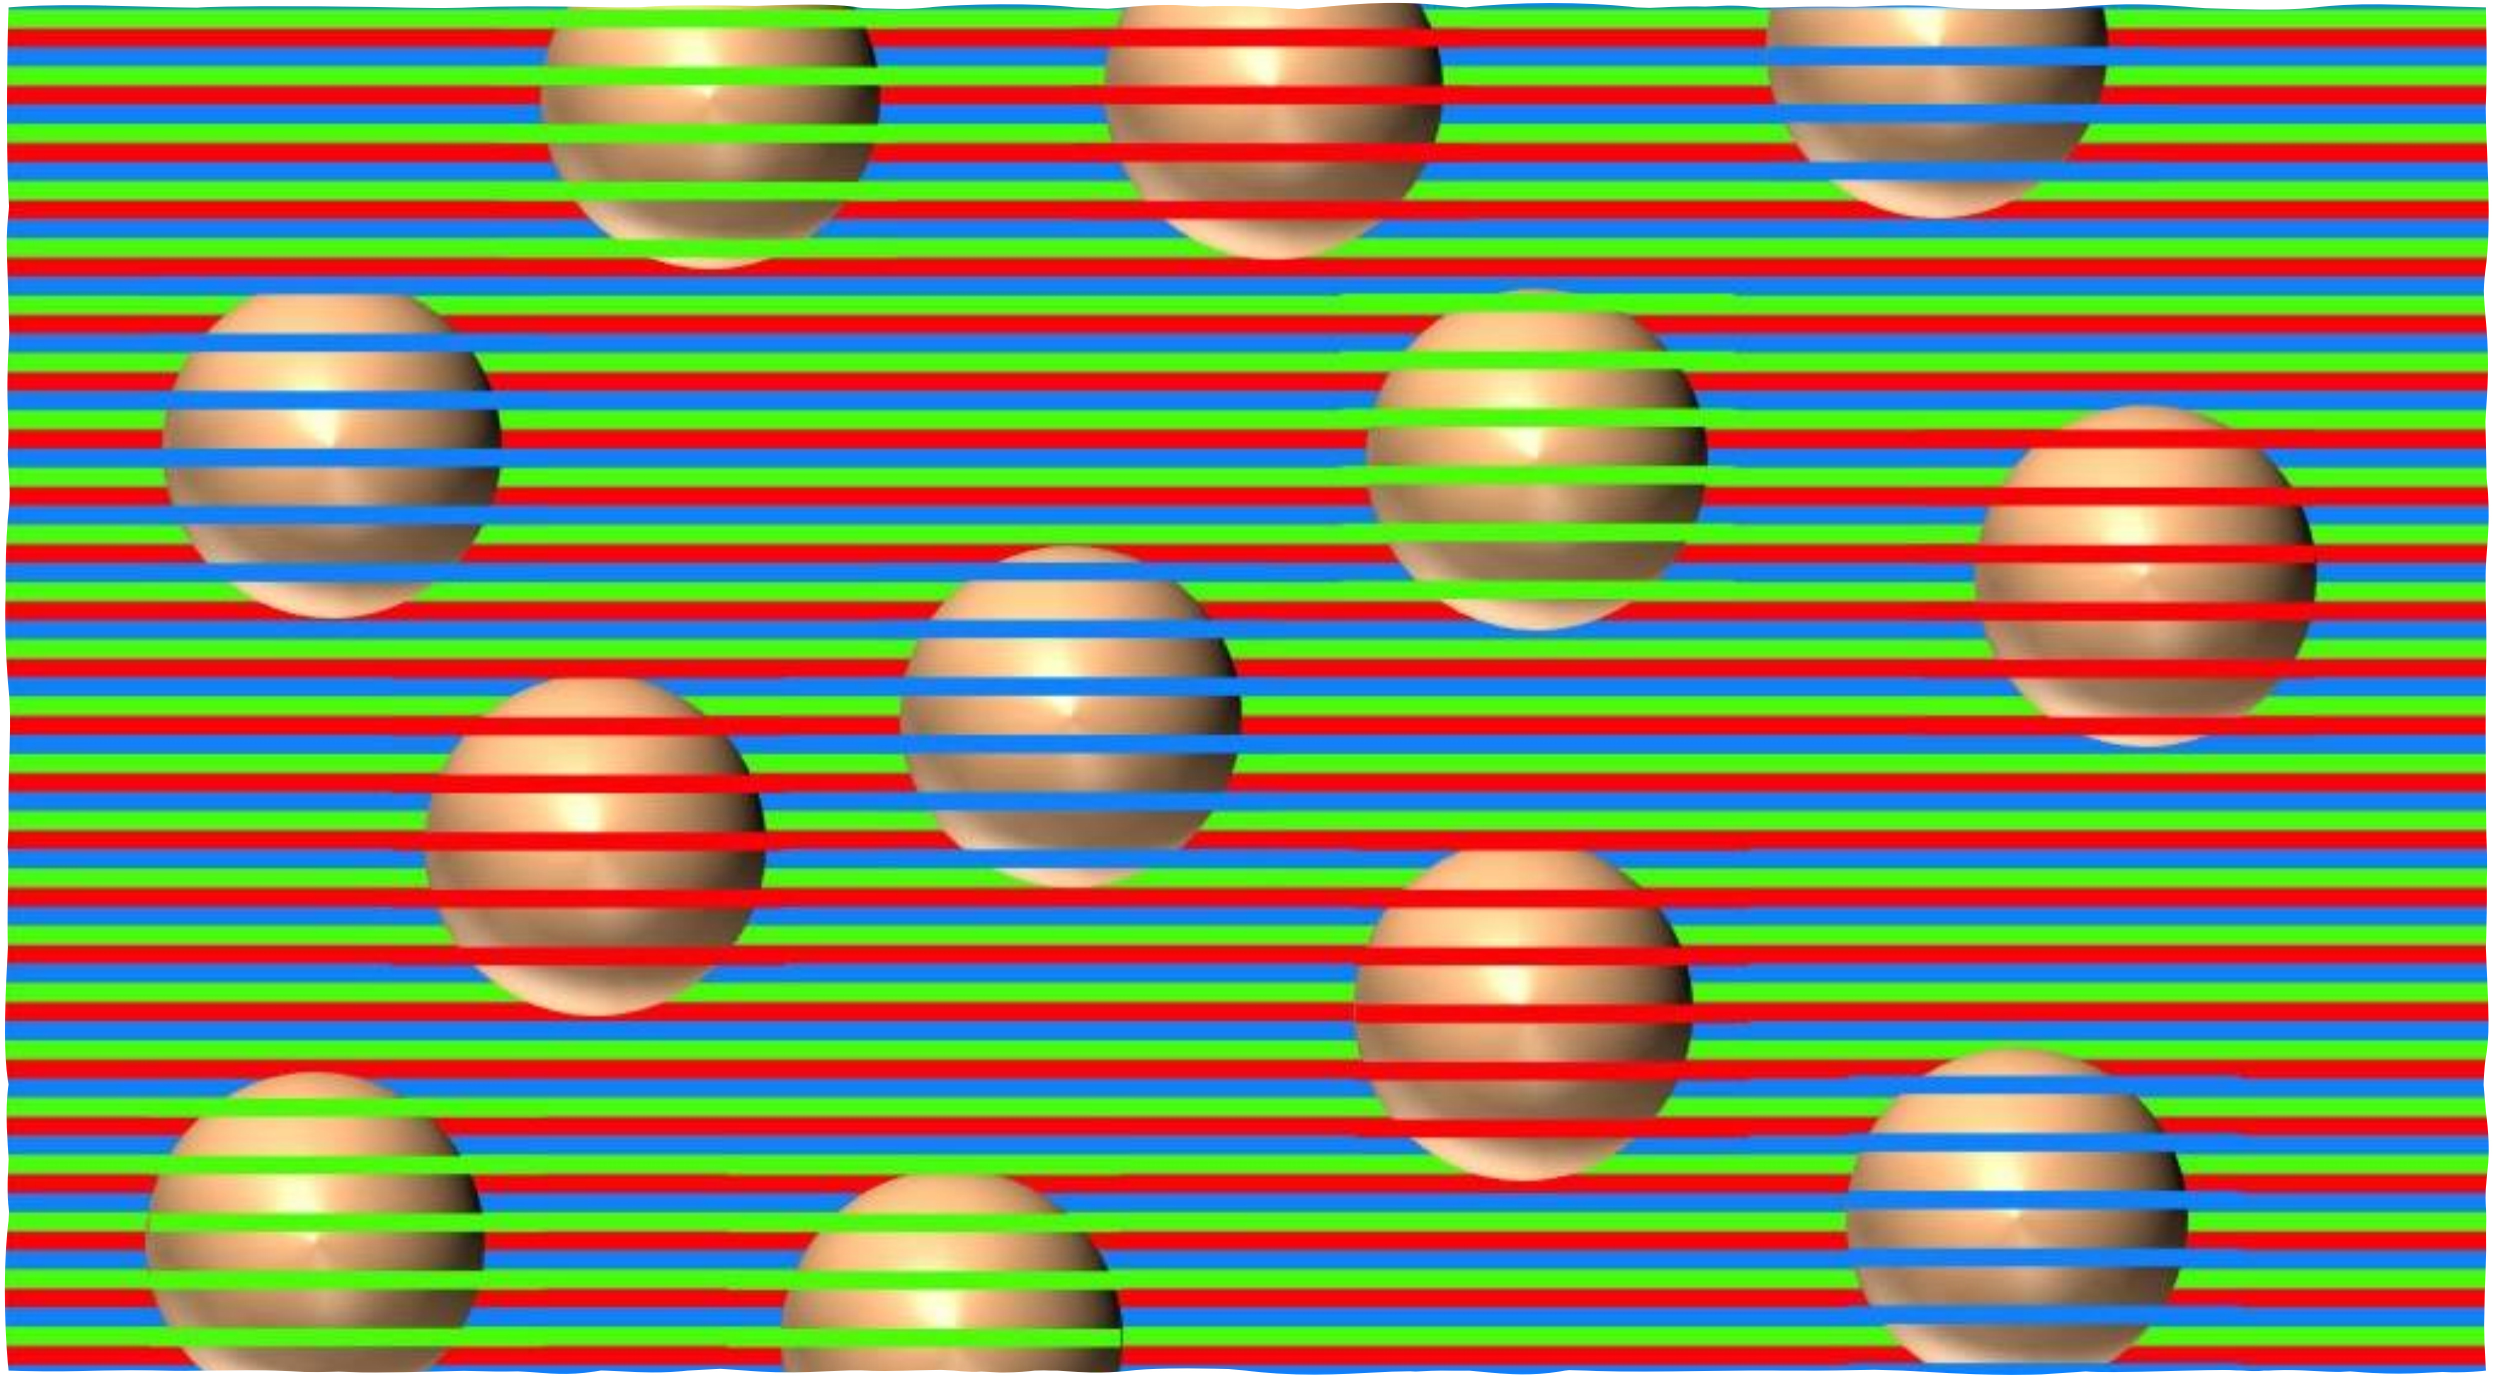

Color constancy is everywhere!

Color constancy is everywhere!

For context, a crash course…

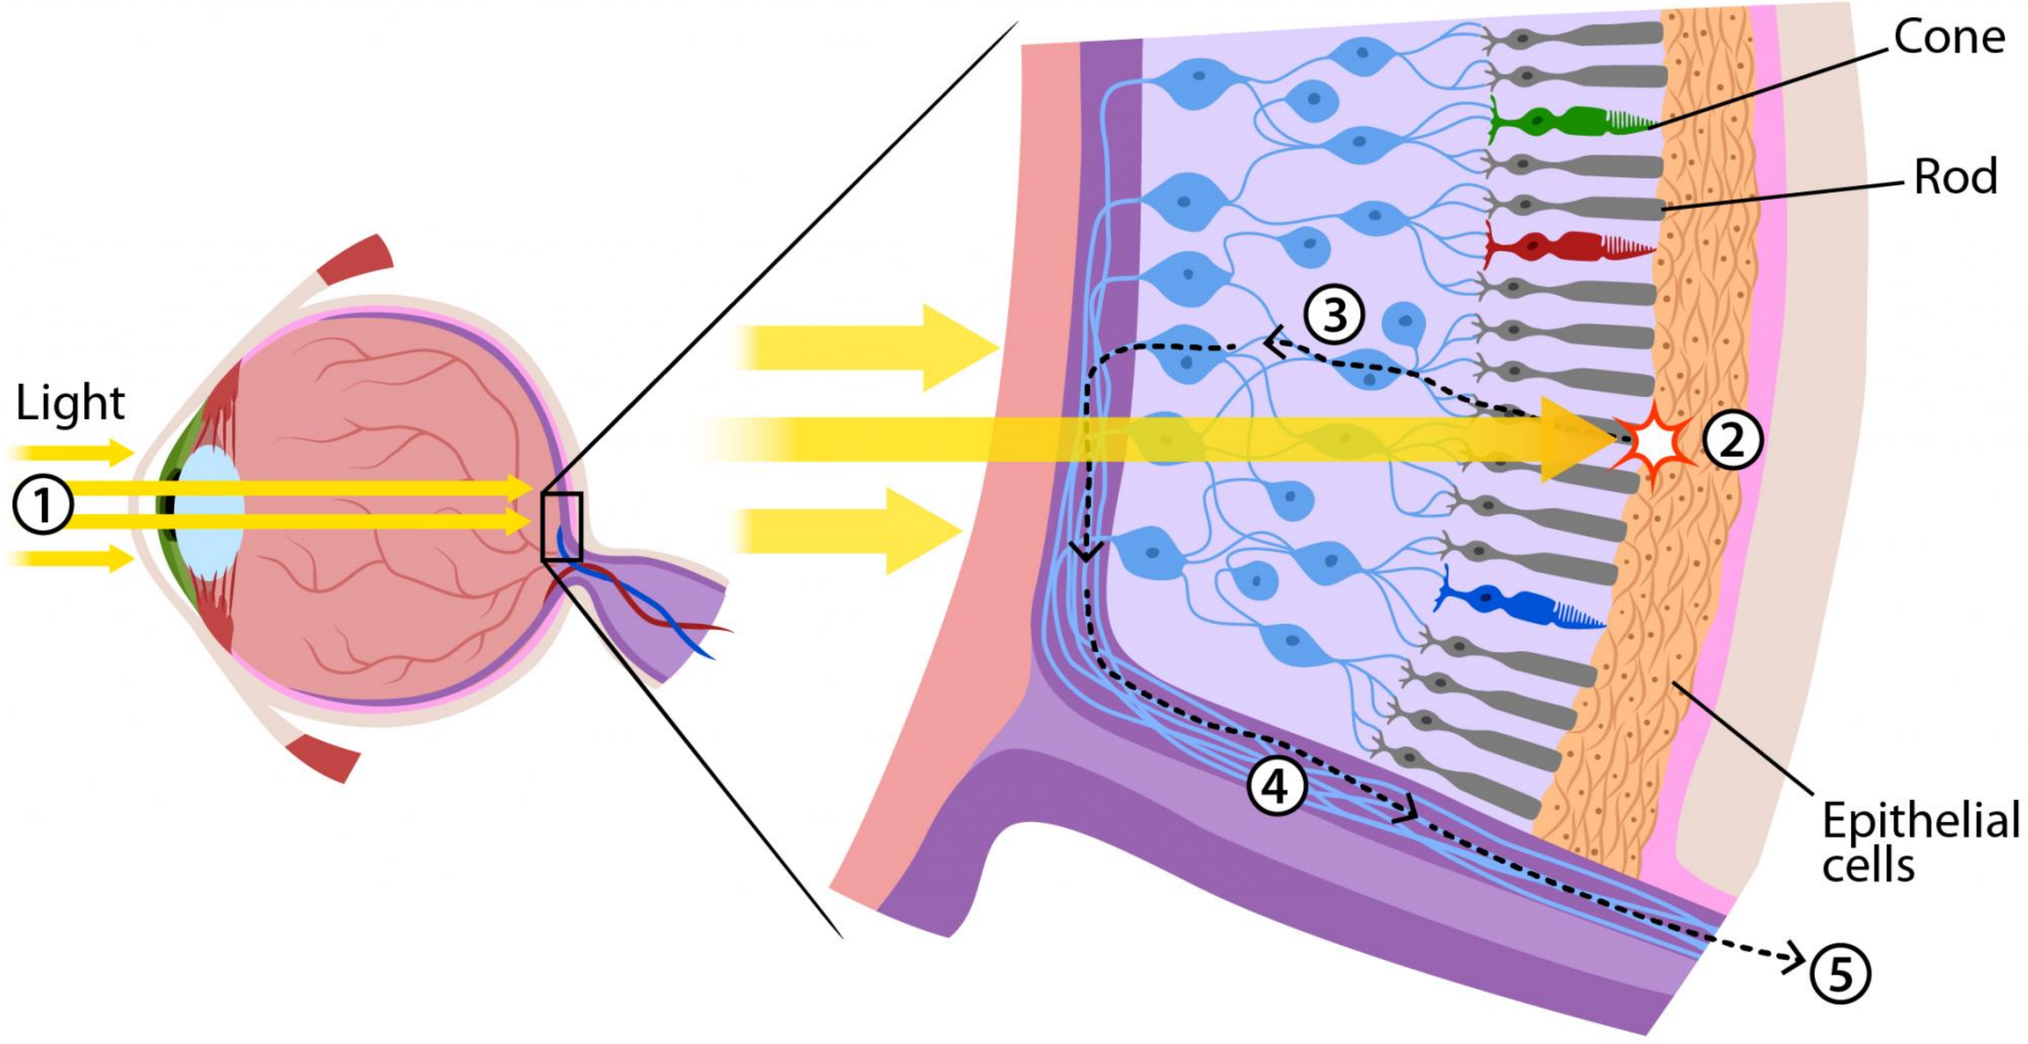
Light
1
2
3
4
5
Cone
Rod
Epithelial cells

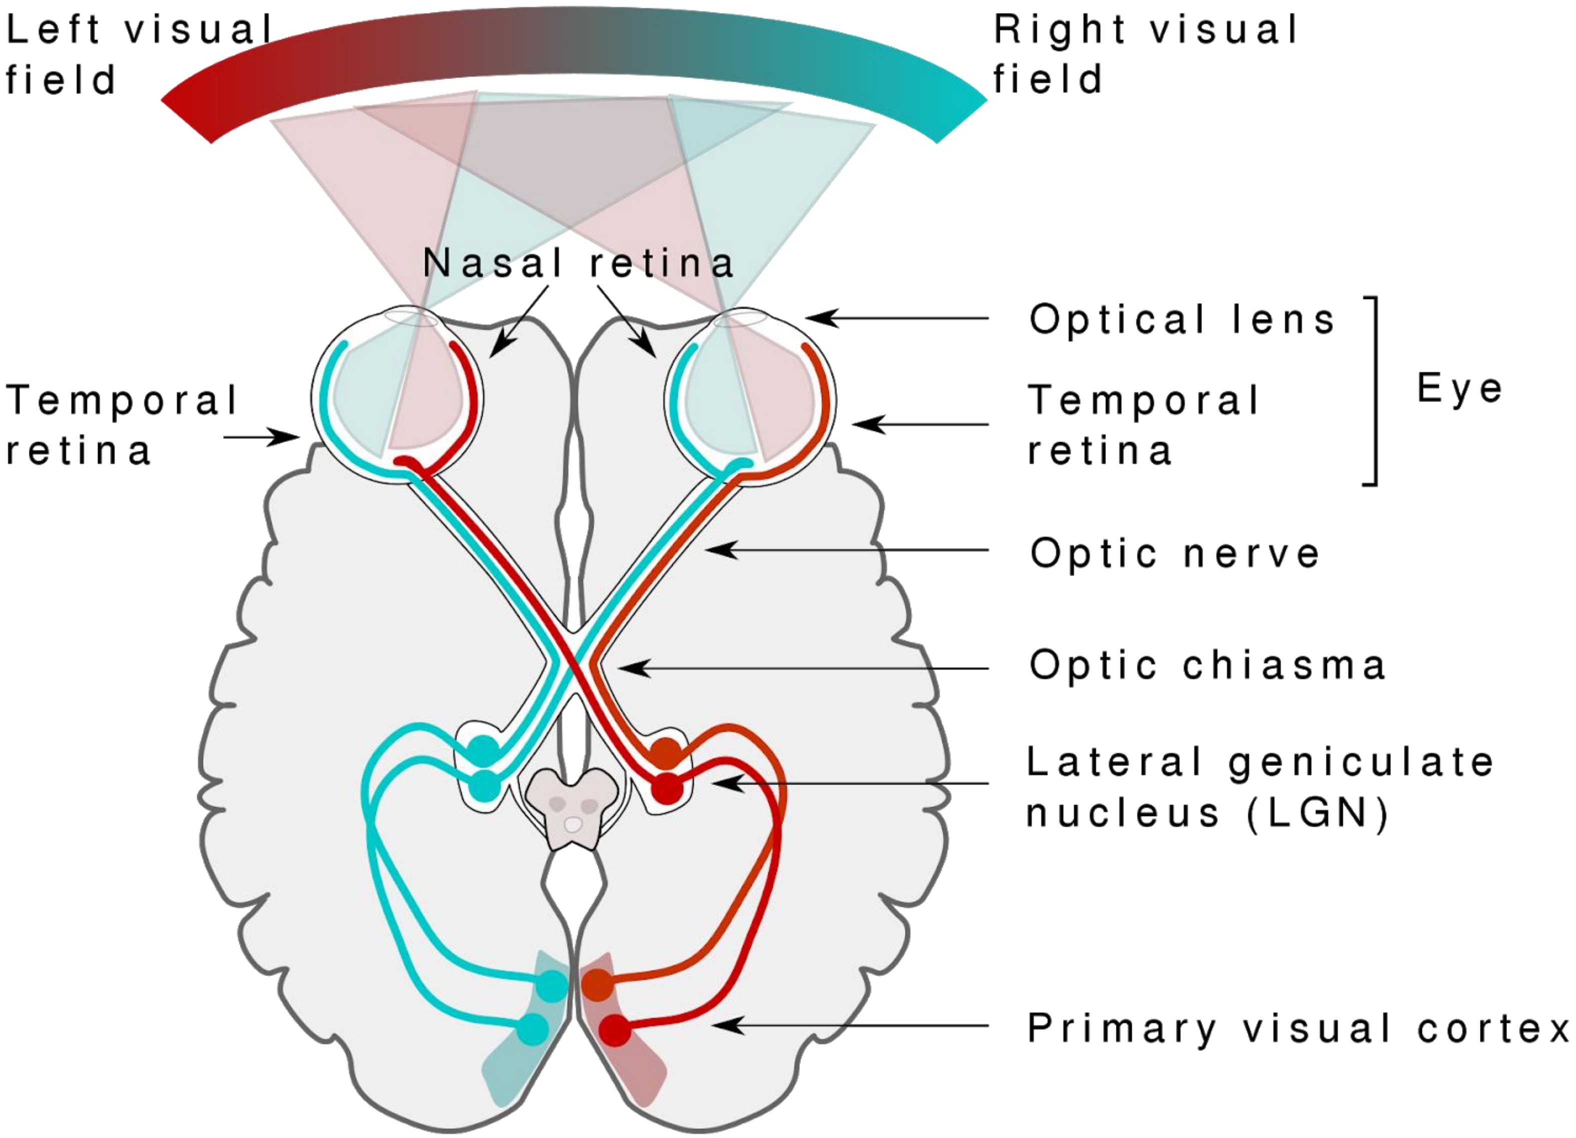
Left visual field
Right visual field
Nasal retina
Optical lens
Temporal retina
Temporal retina
Eye
Optic nerve
Optic chiasma
Lateral geniculate nucleus (LGN)
Primary visual cortex

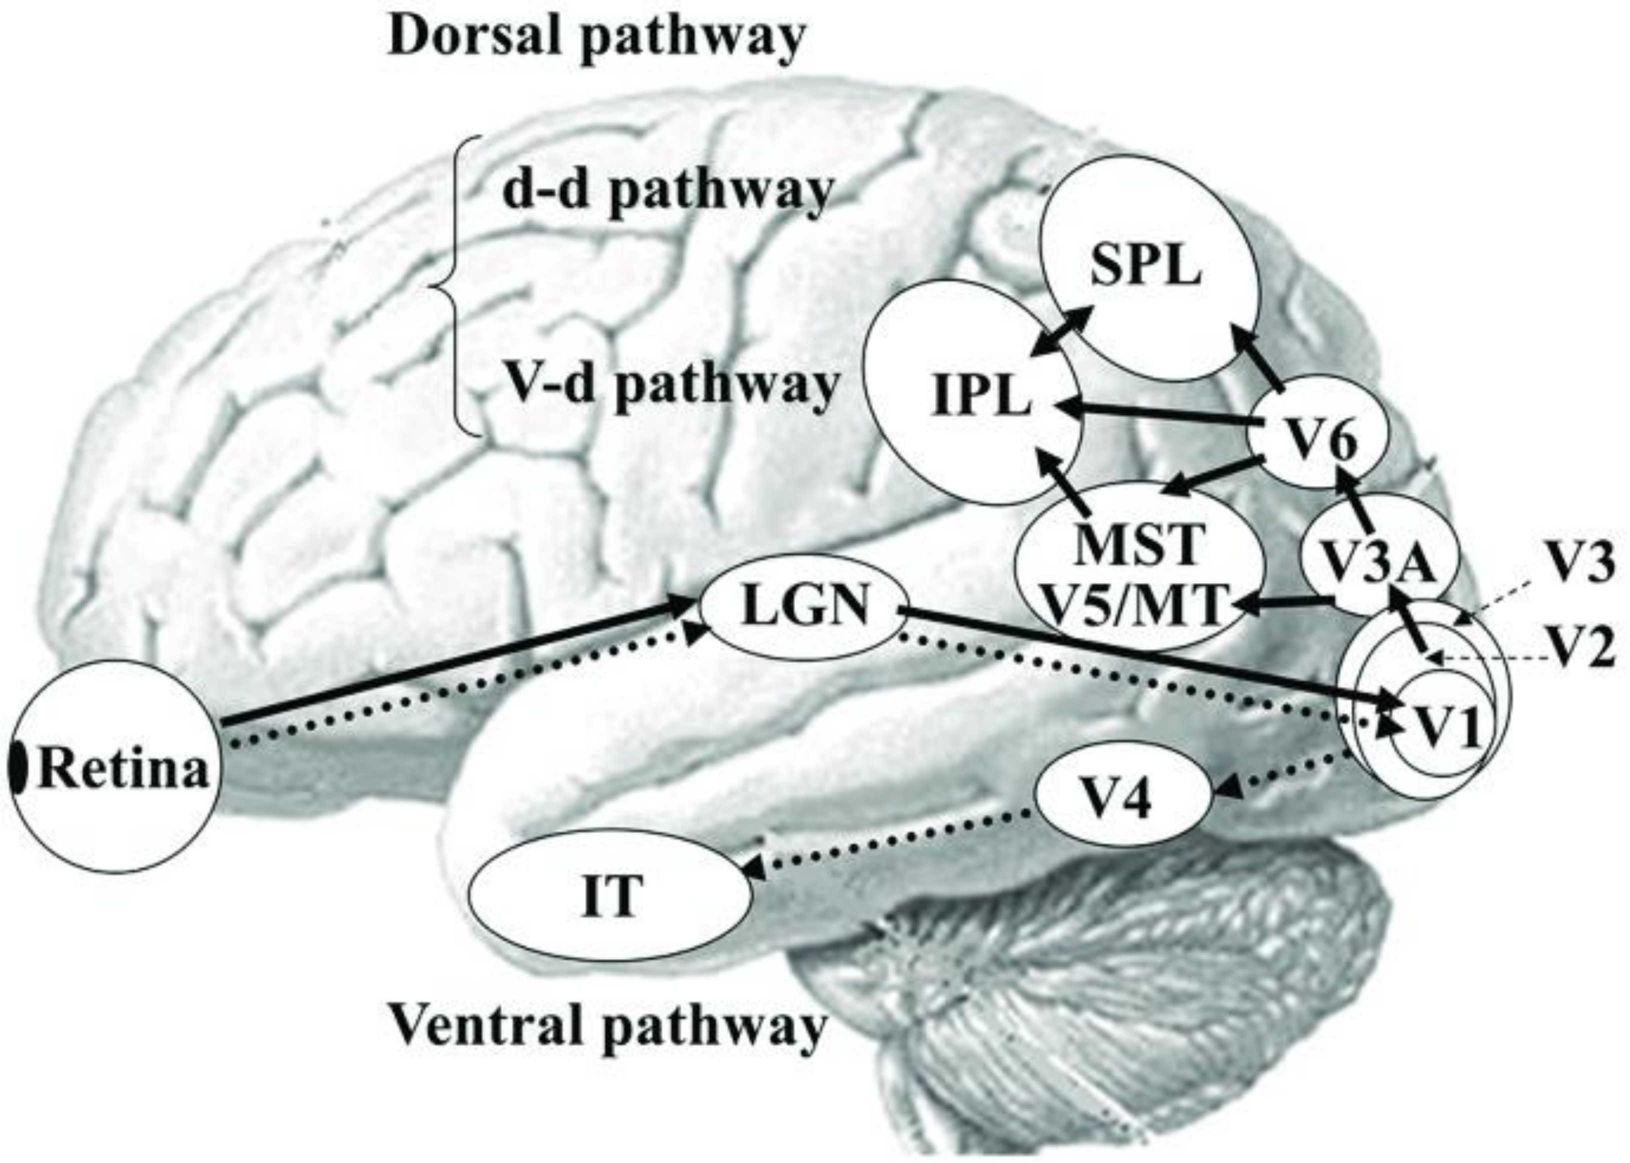
Dorsal pathway
d-d pathway
V-d pathway
SPL
IPL
V6
MST
V5/MT
V3A
V3
V2
LGN
V1
Retina
V4
IT
Ventral pathway

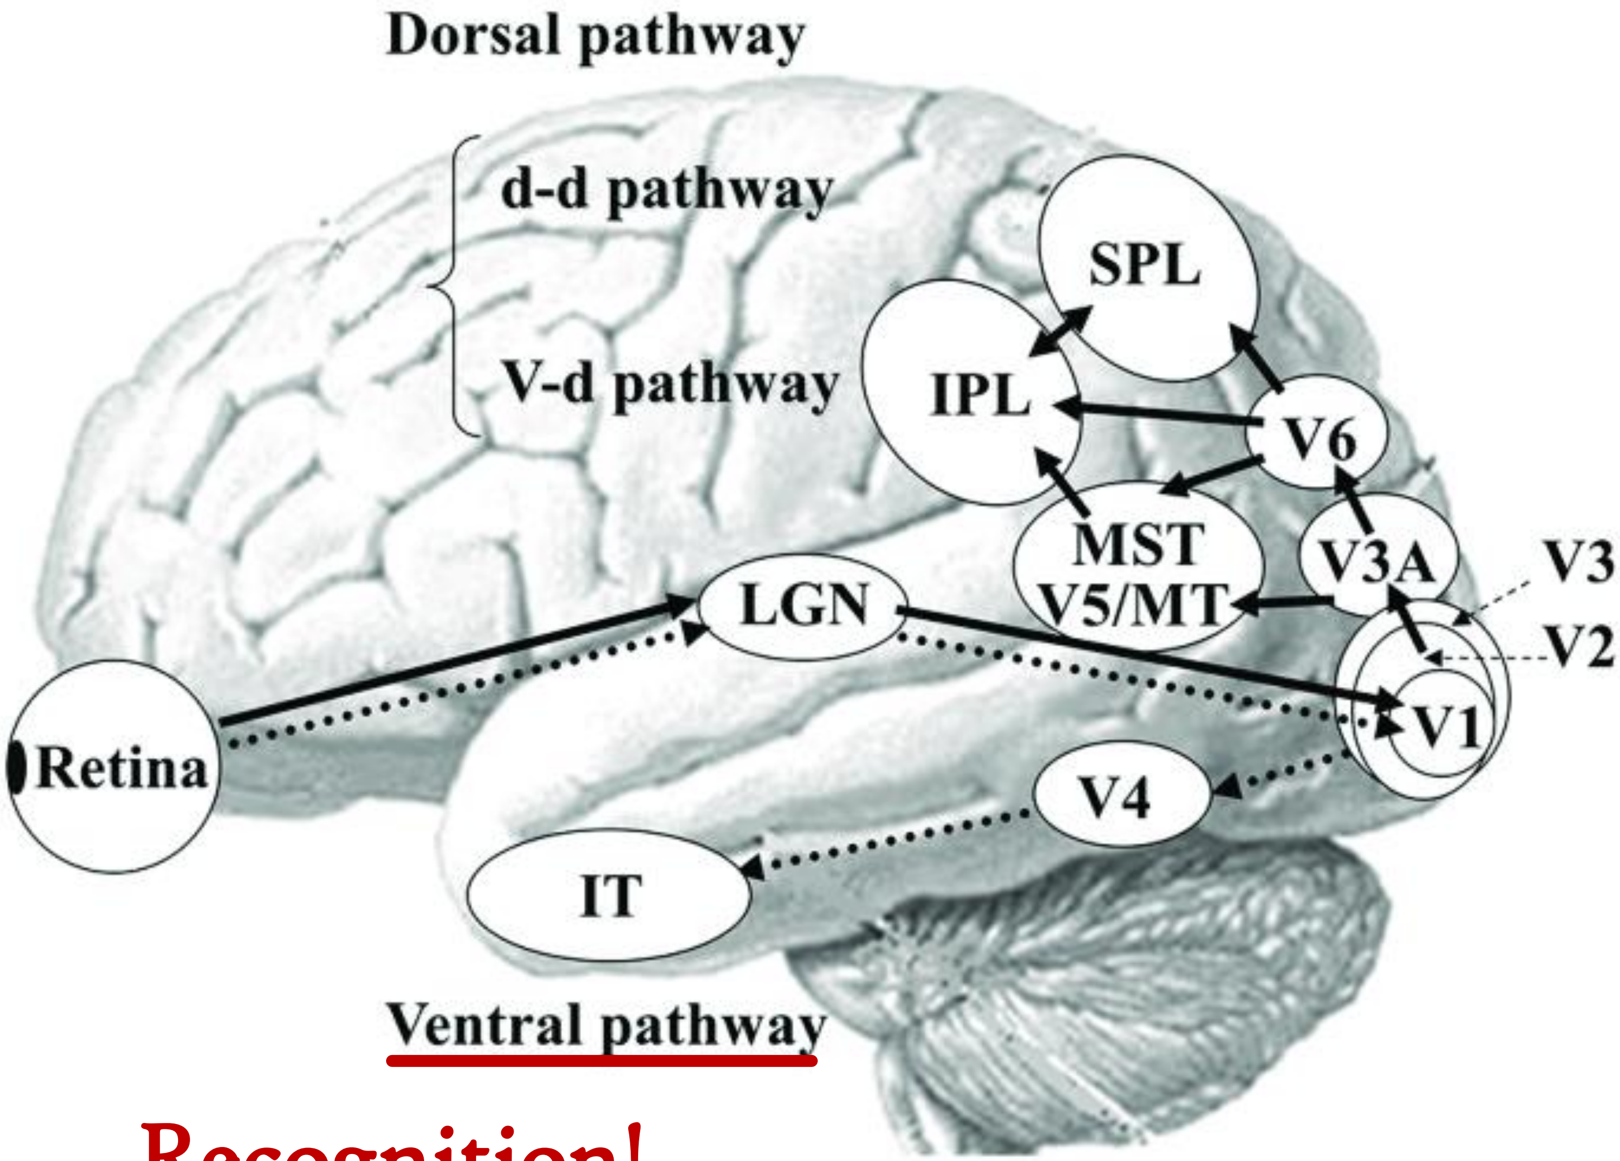

Recognition!

!!

Action!

!!

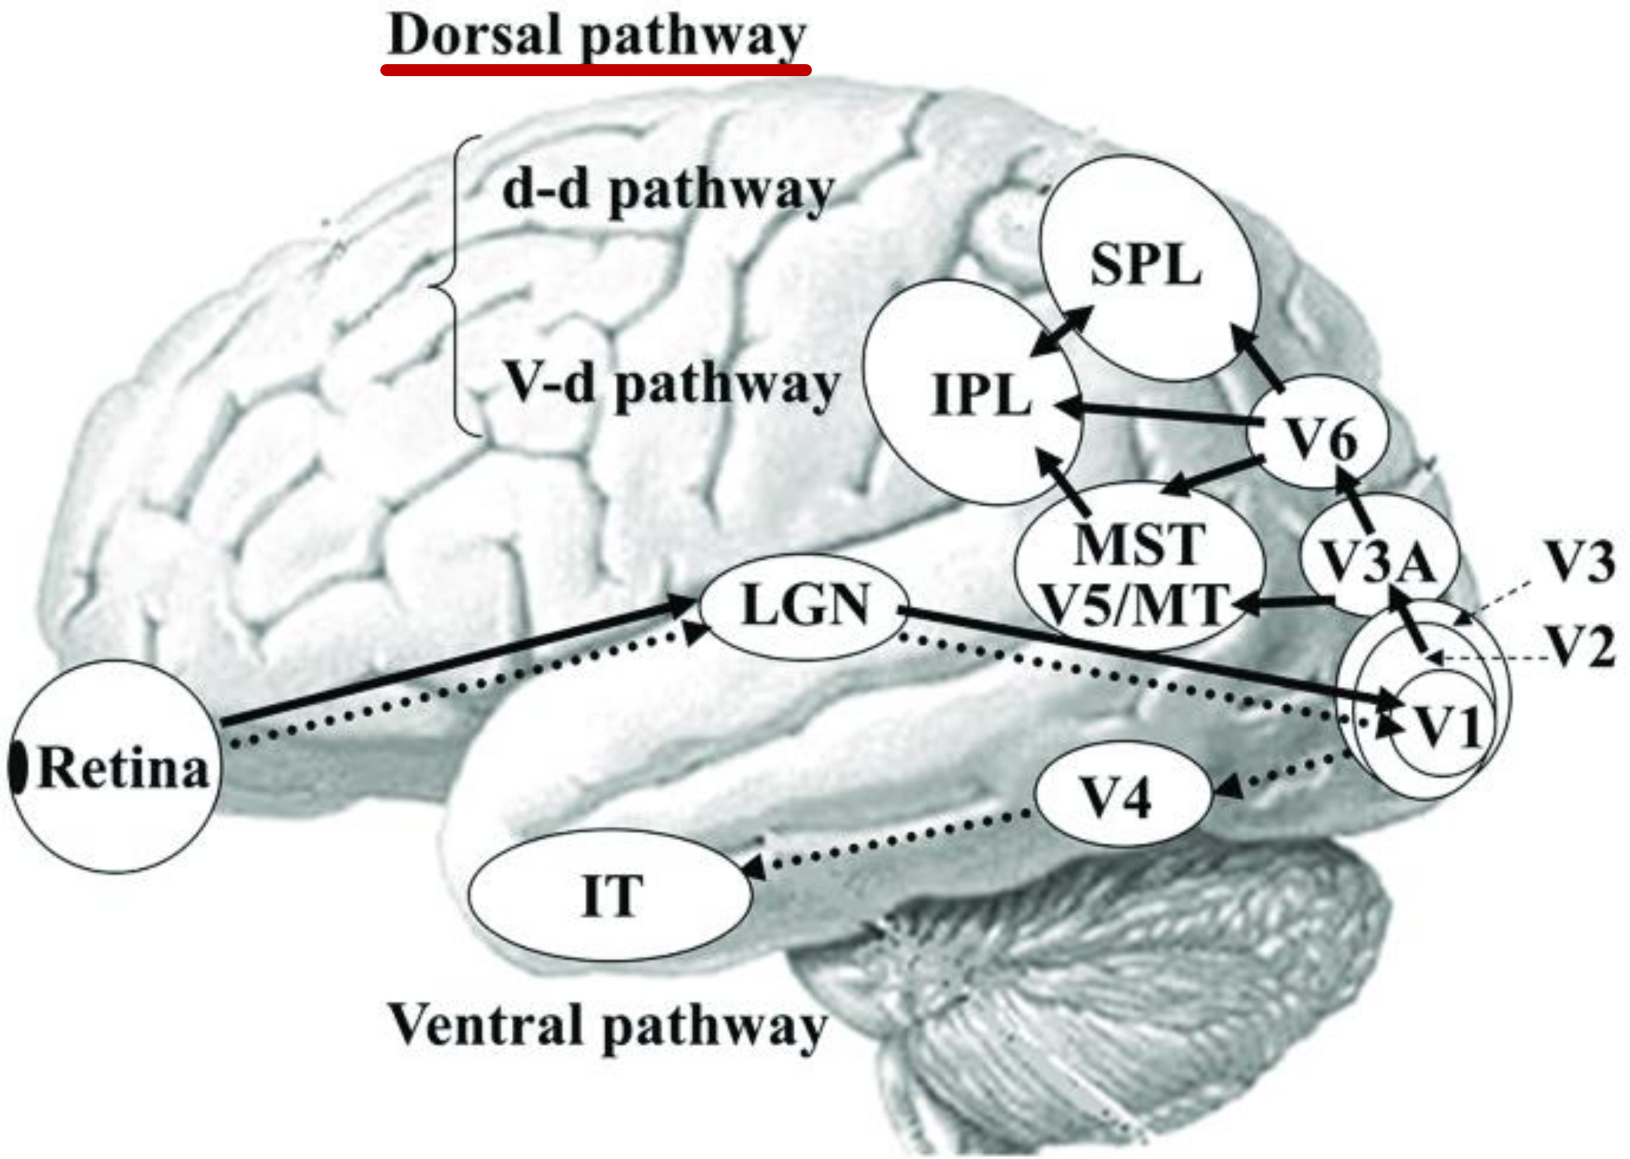
Dorsal pathway
d-d pathway
V-d pathway
SPL
IPL
V6
MST
V5/MT
V3A
V3
V2
V1
LGN
Retina
V4
IT
Ventral pathway

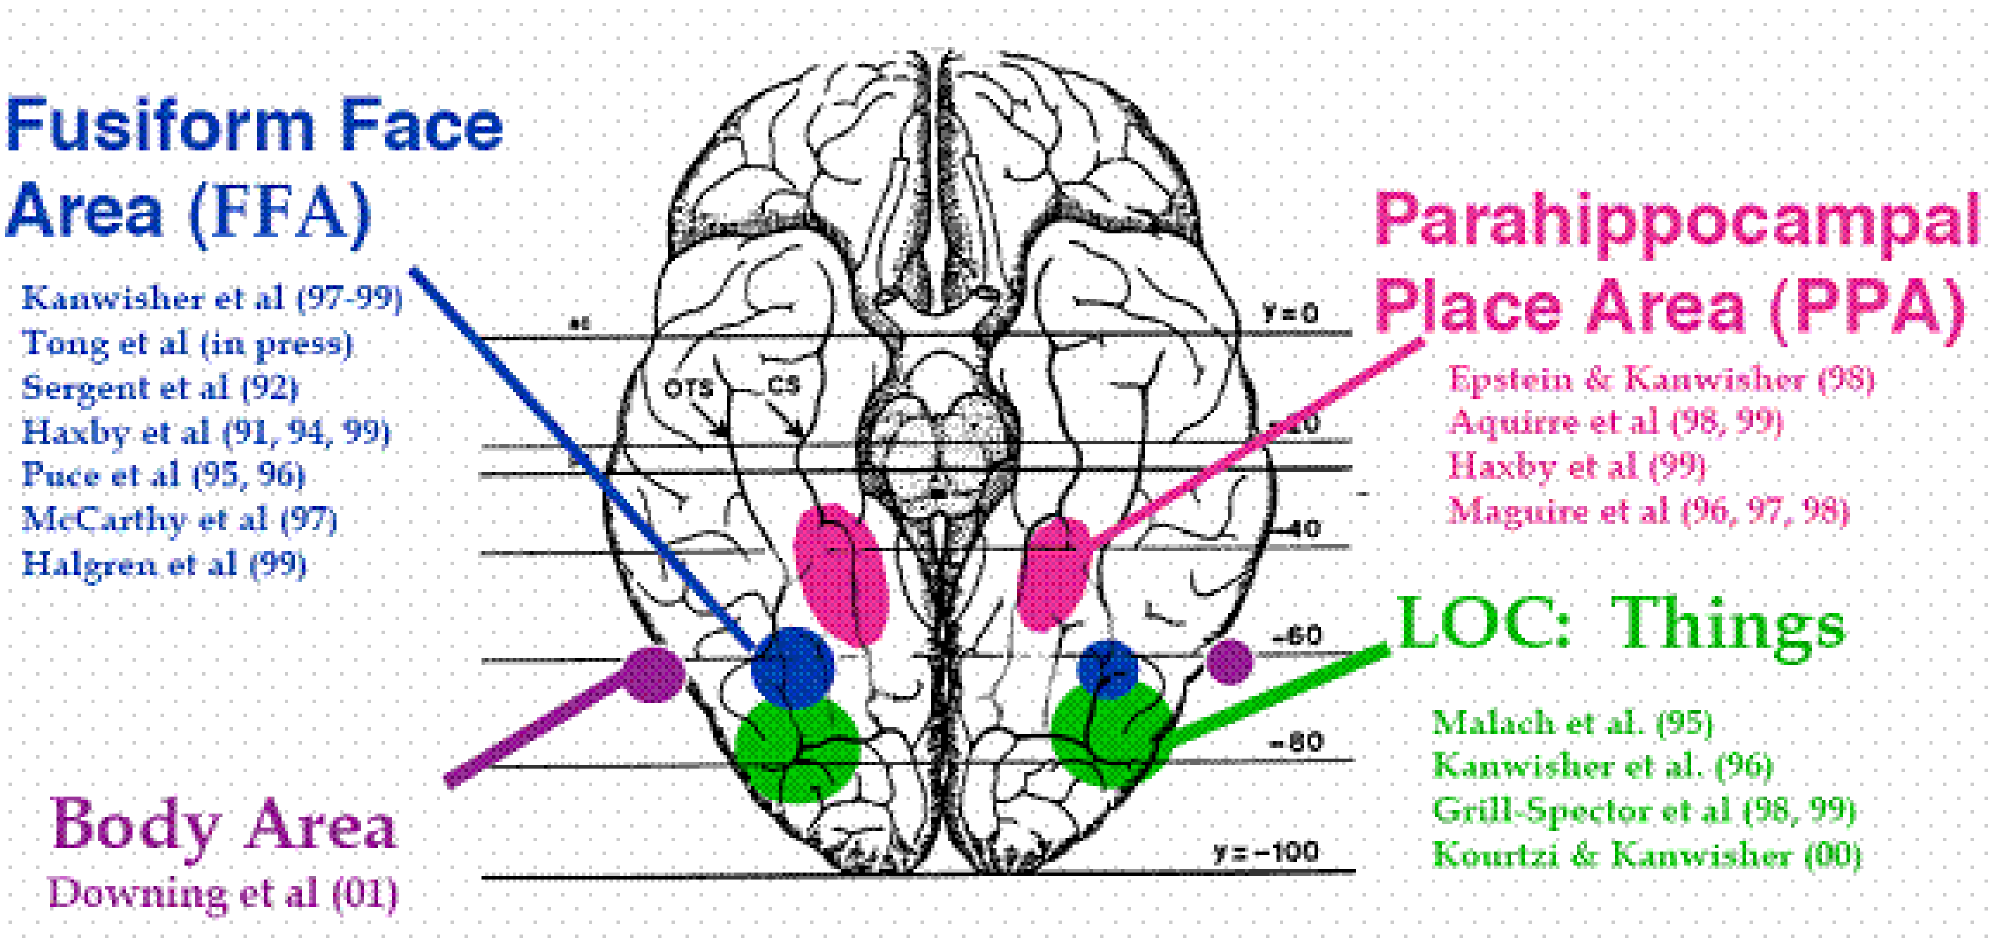

Fusiform Face Area (FFA)
Kanwisher et al (97-99)
Tong et al (in press)
Sergent et al (92)
Haxby et al (91, 94, 99)
Puce et al (95, 96)
McCarthy et al (97)
Halgren et al (99)
Parahippocampal Place Area (PPA)
Epstein & Kanwisher (98)
Aquirre et al (98, 99)
Haxby et al (99)
Maguire et al (96, 97, 98)
LOC: Things
Malach et al. (95)
Kanwisher et al. (96)
Grill-Spector et al (98, 99)
Kourtzi & Kanwisher (00)
Body Area
Downing et al (01)
OTS
CS
y=0
-40
-60
-80
y=-100

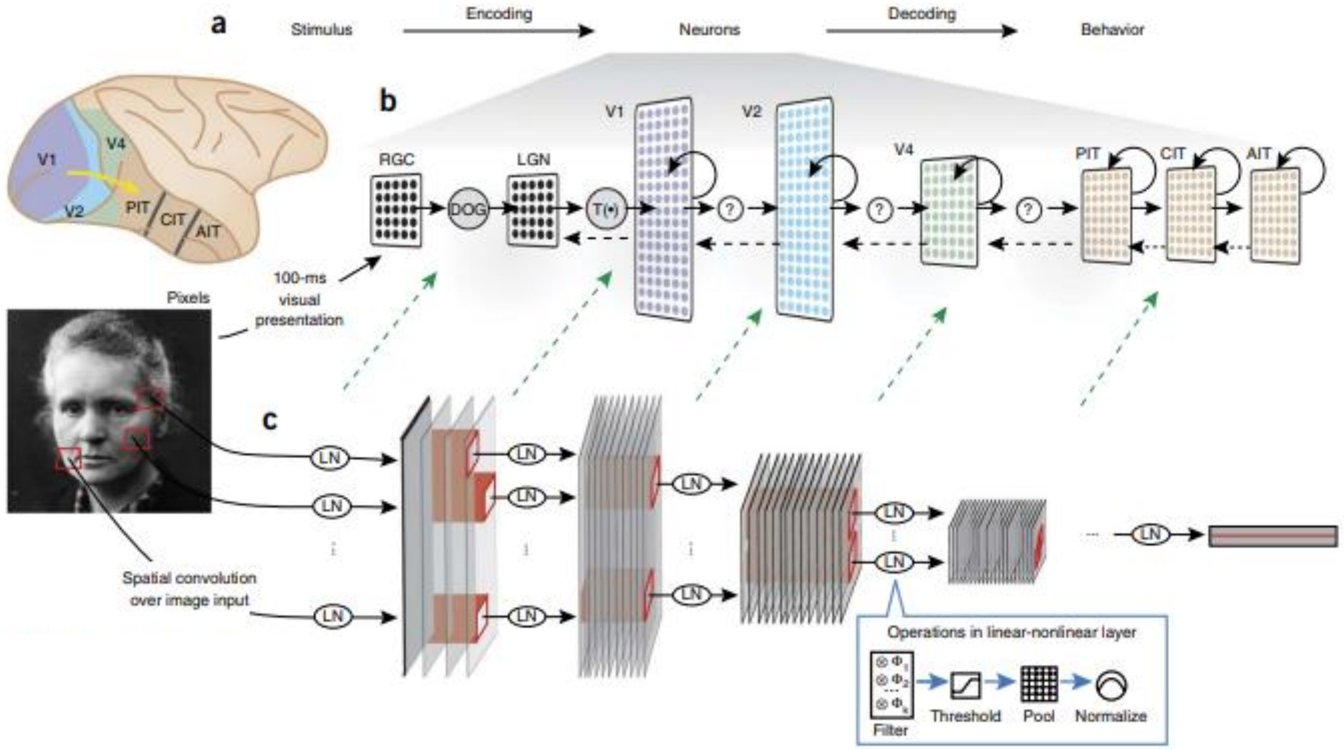
a
Stimulus
Encoding
Neurons
Decoding
Behavior
b
V1
V4
V2
PIT
CIT
AIT
RGC
DOG
LGN
T(•)
?
Pixels
100-ms
visual
presentation
c
LN
Spatial convolution
over image input
Operations in linear-nonlinear layer
Filter
Threshold
Pool
Normalize

Big question: *Where do these illusions occur?*


Color
adaptation
!!

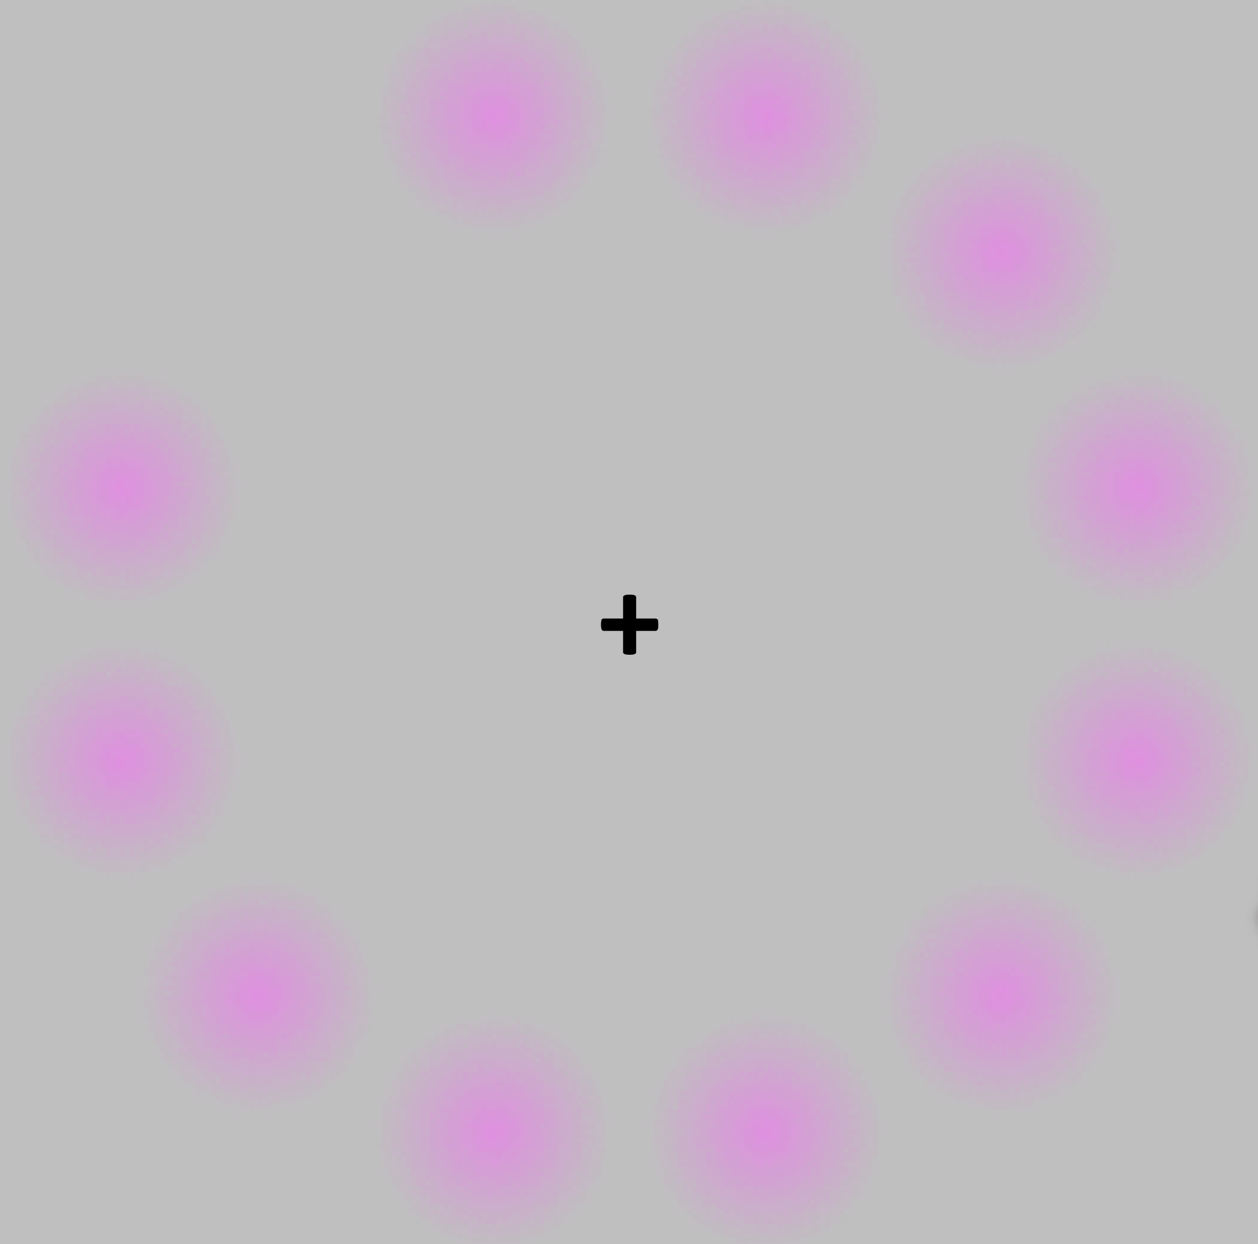

Ignaz Troxler (1780-1866)

ILLUSIONS OF MOTION

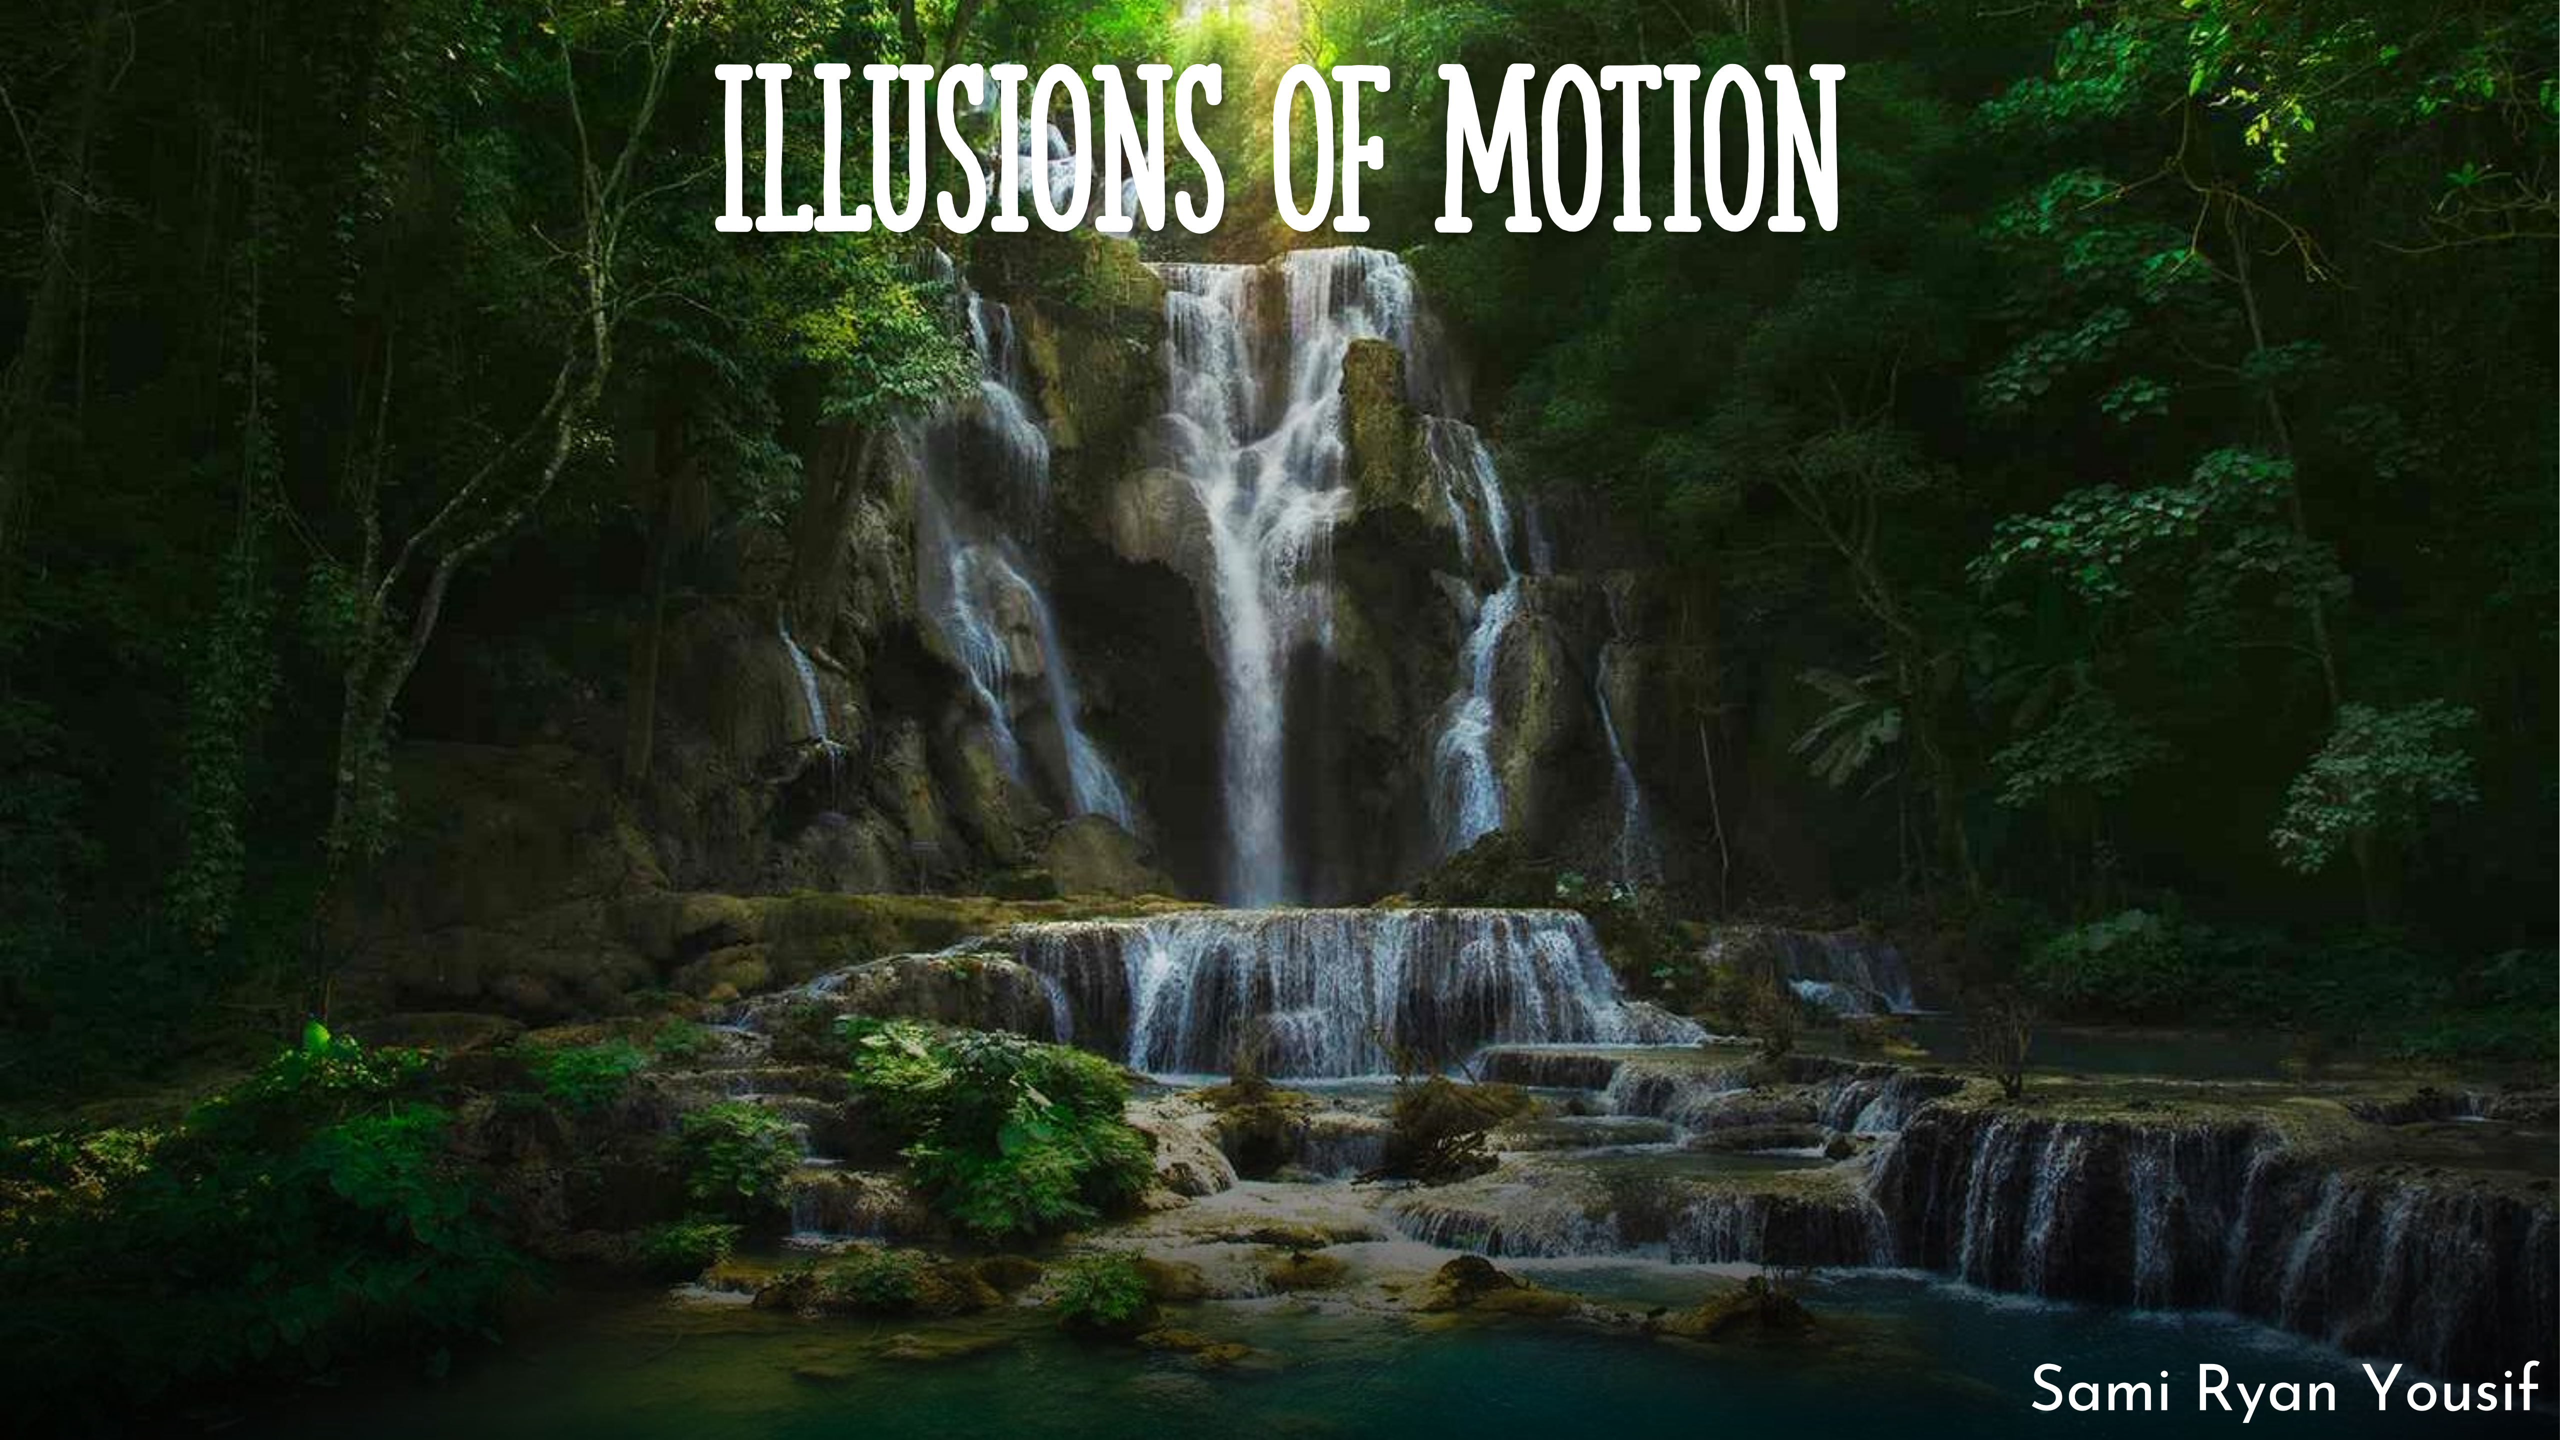

Sami Ryan Yousif

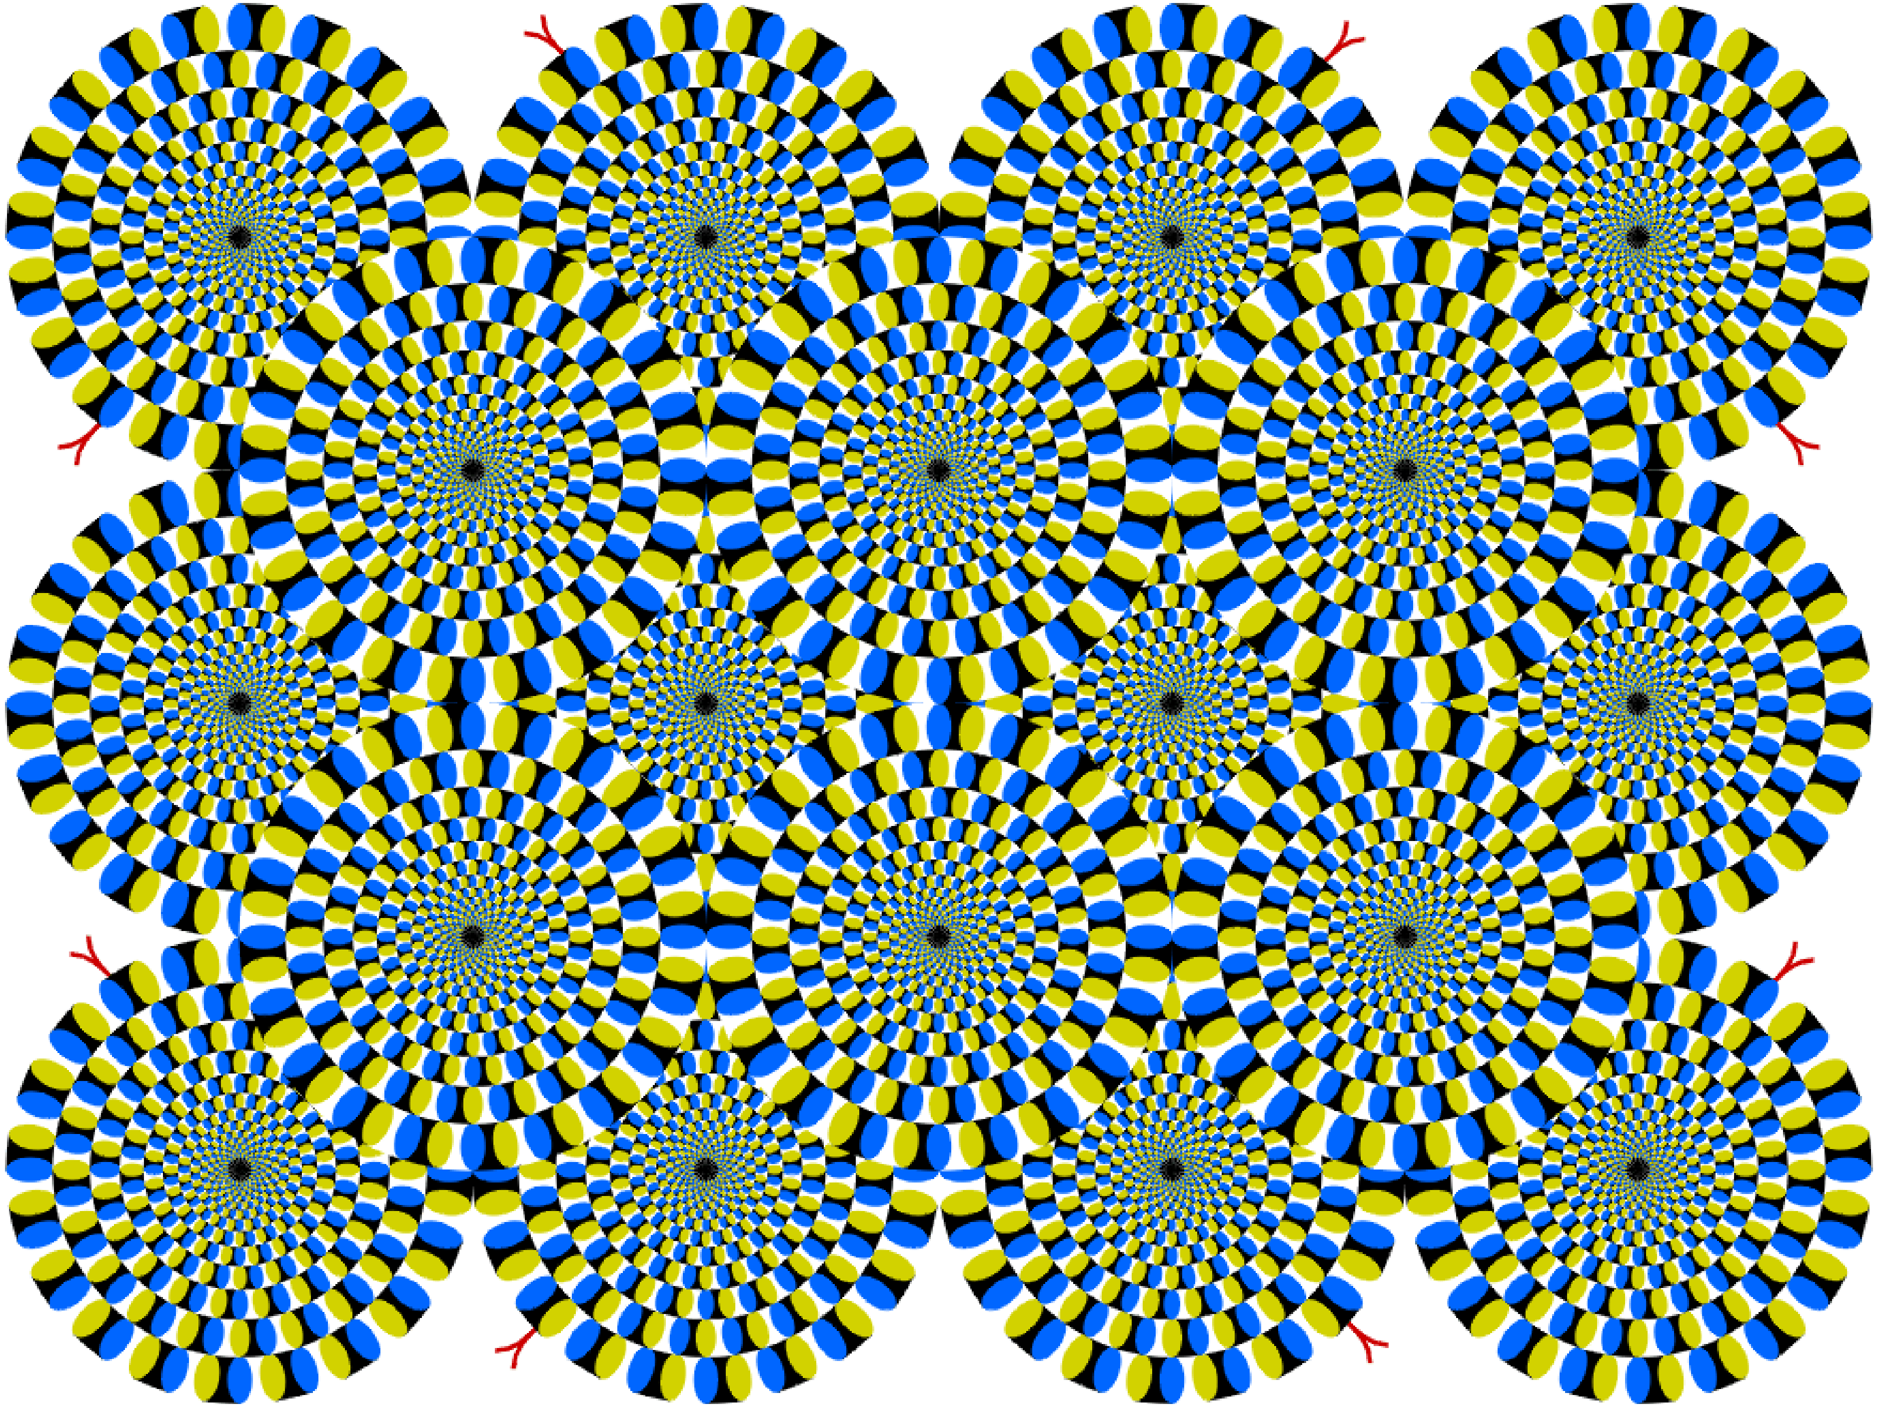

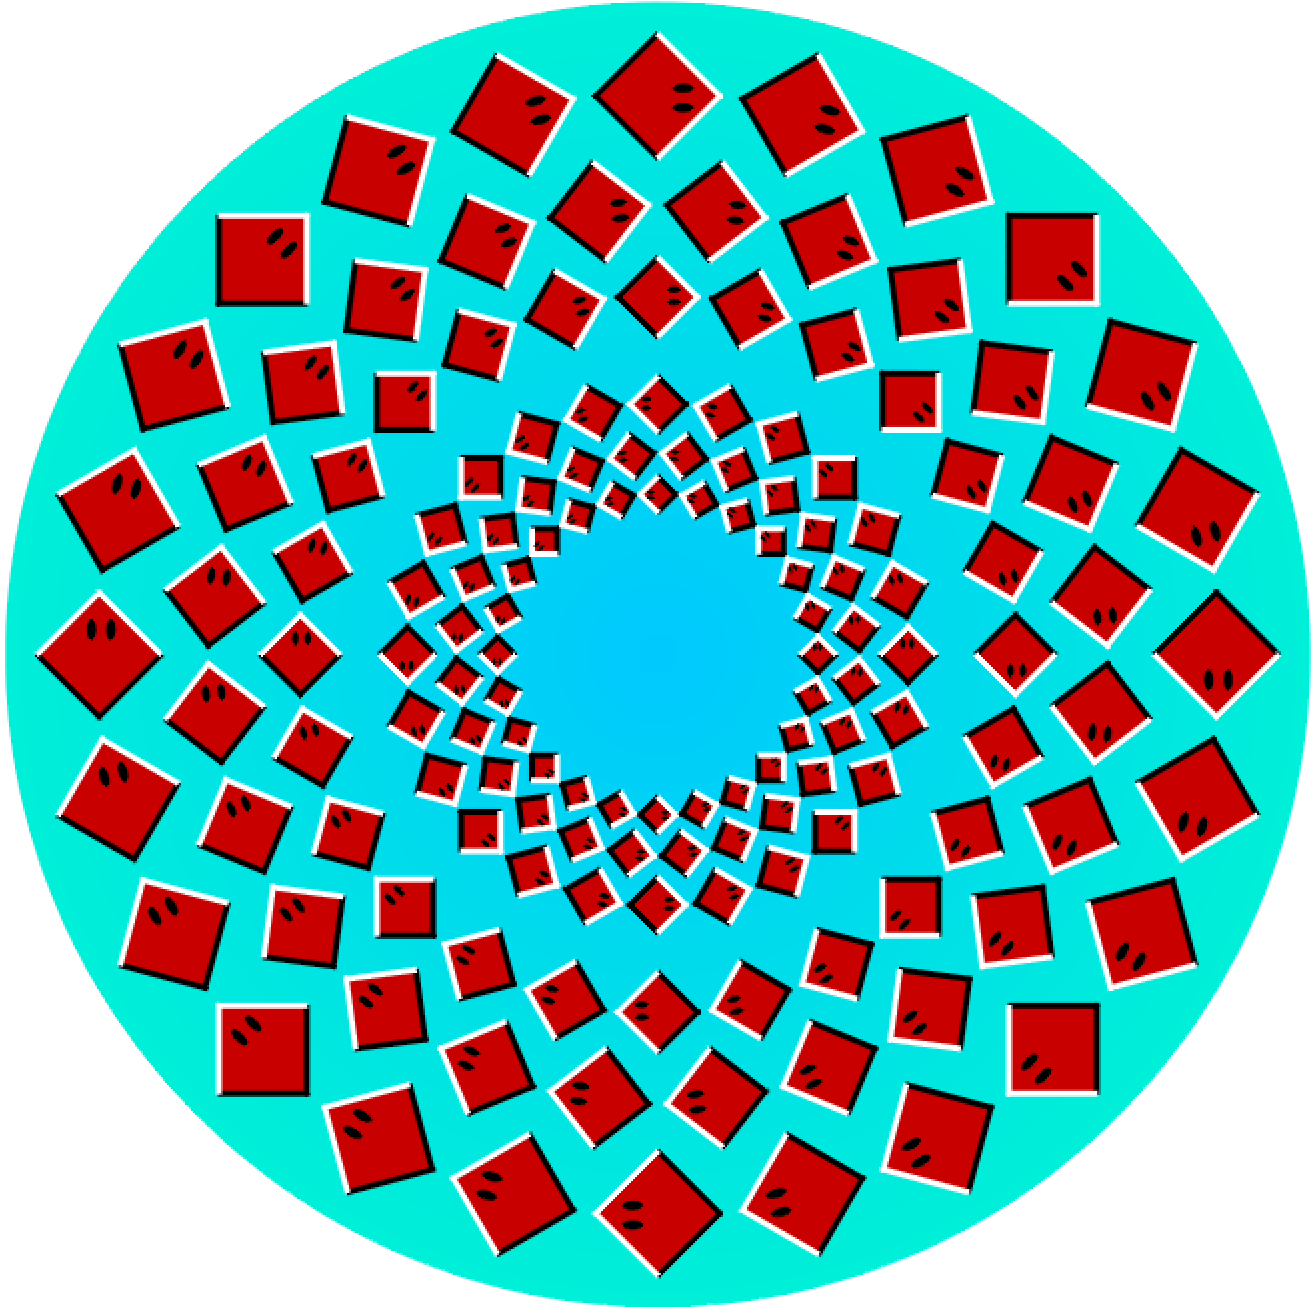

Why?

‼

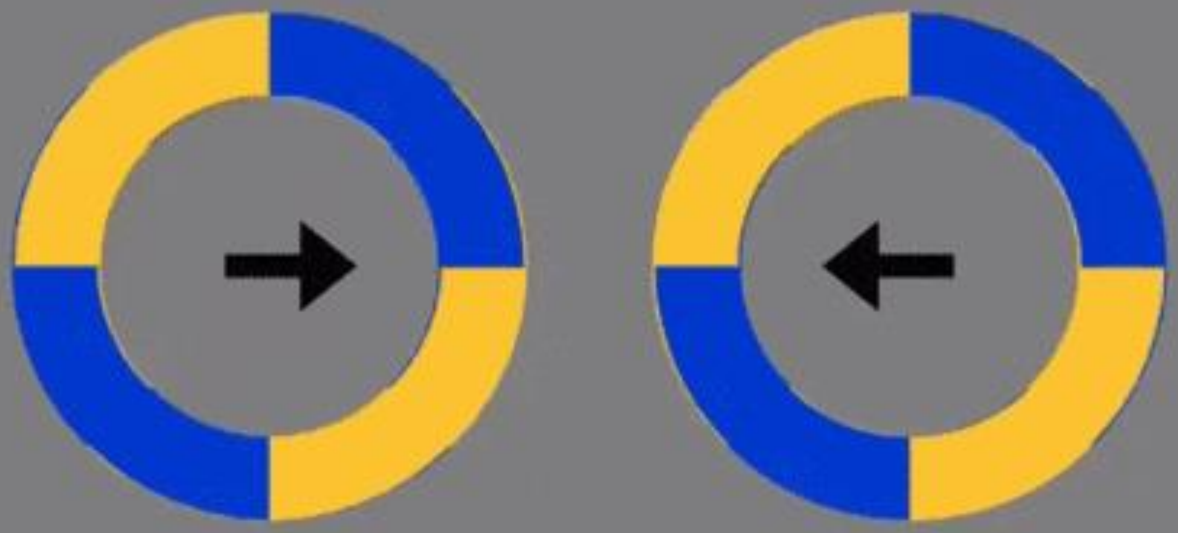

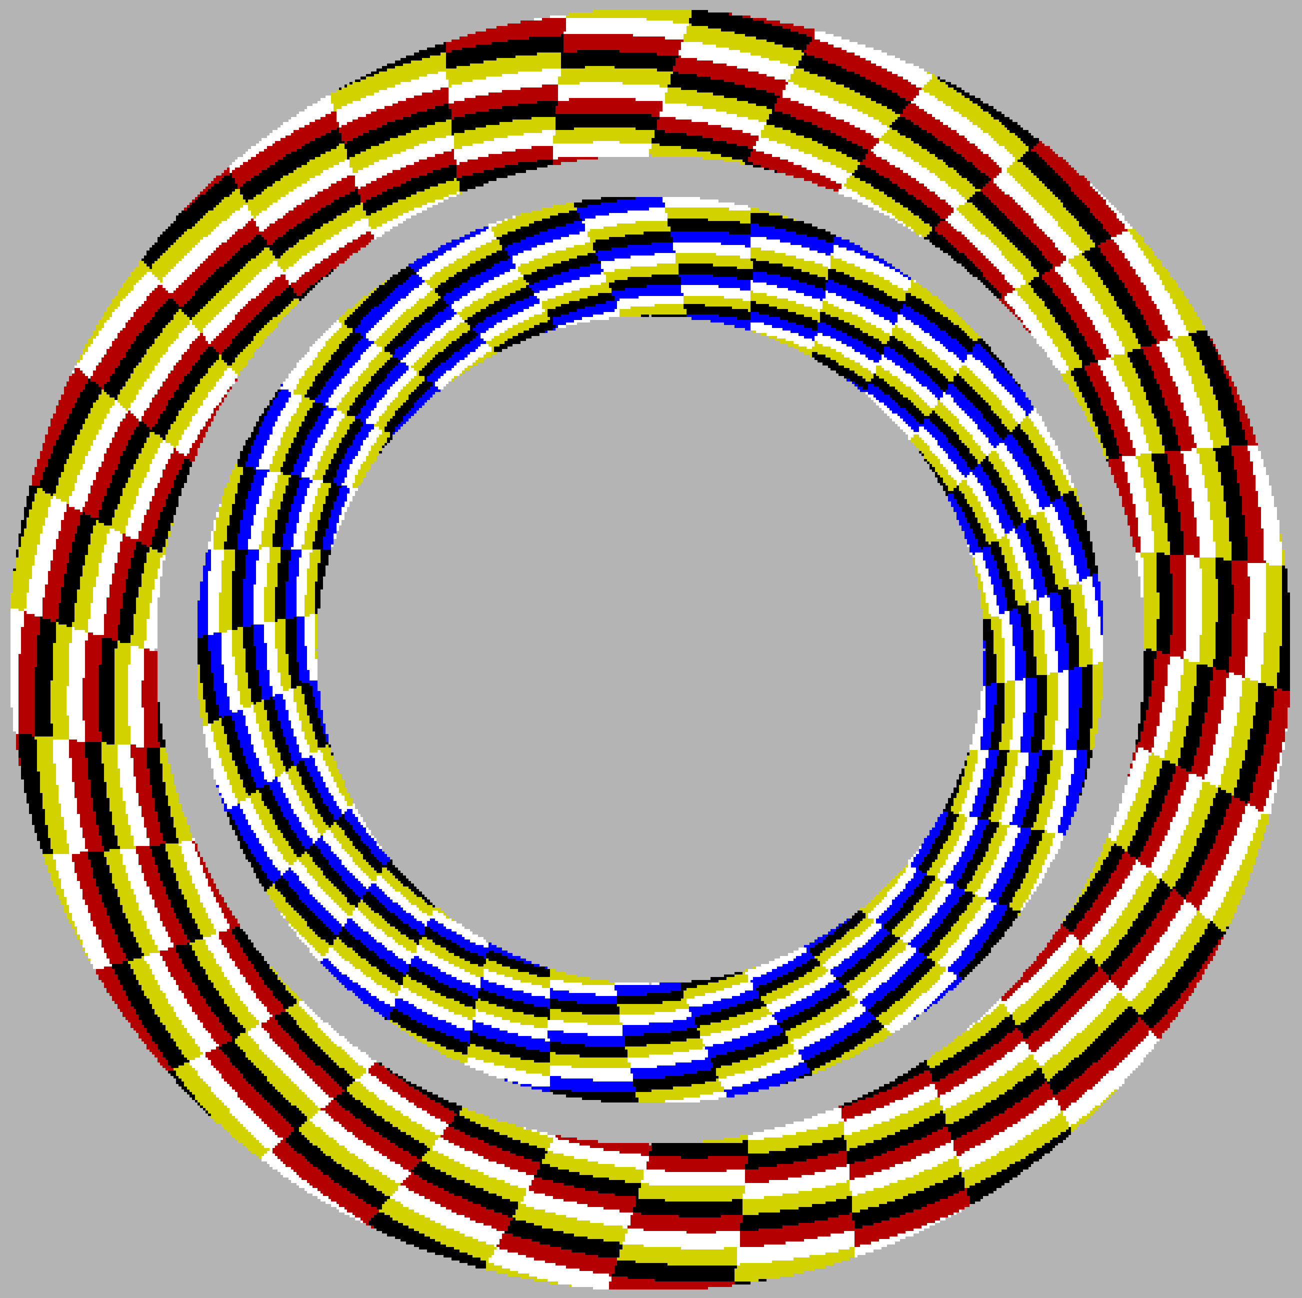

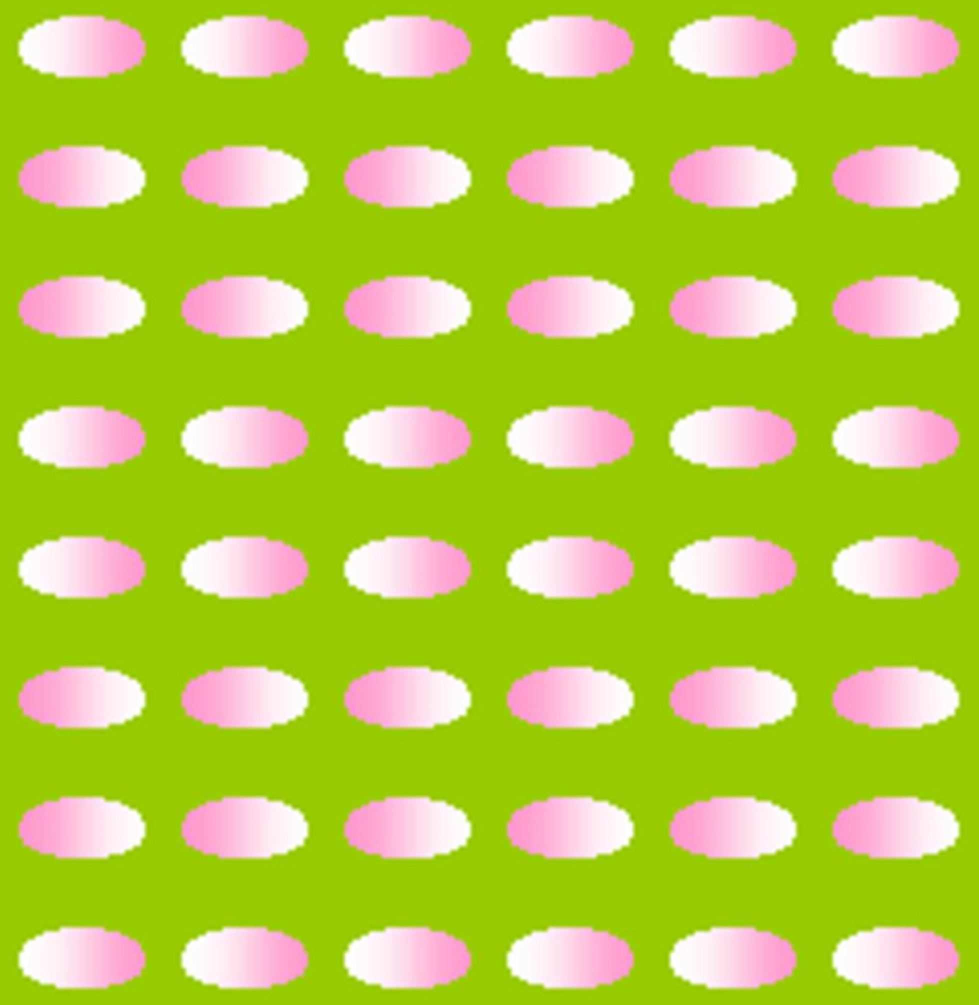

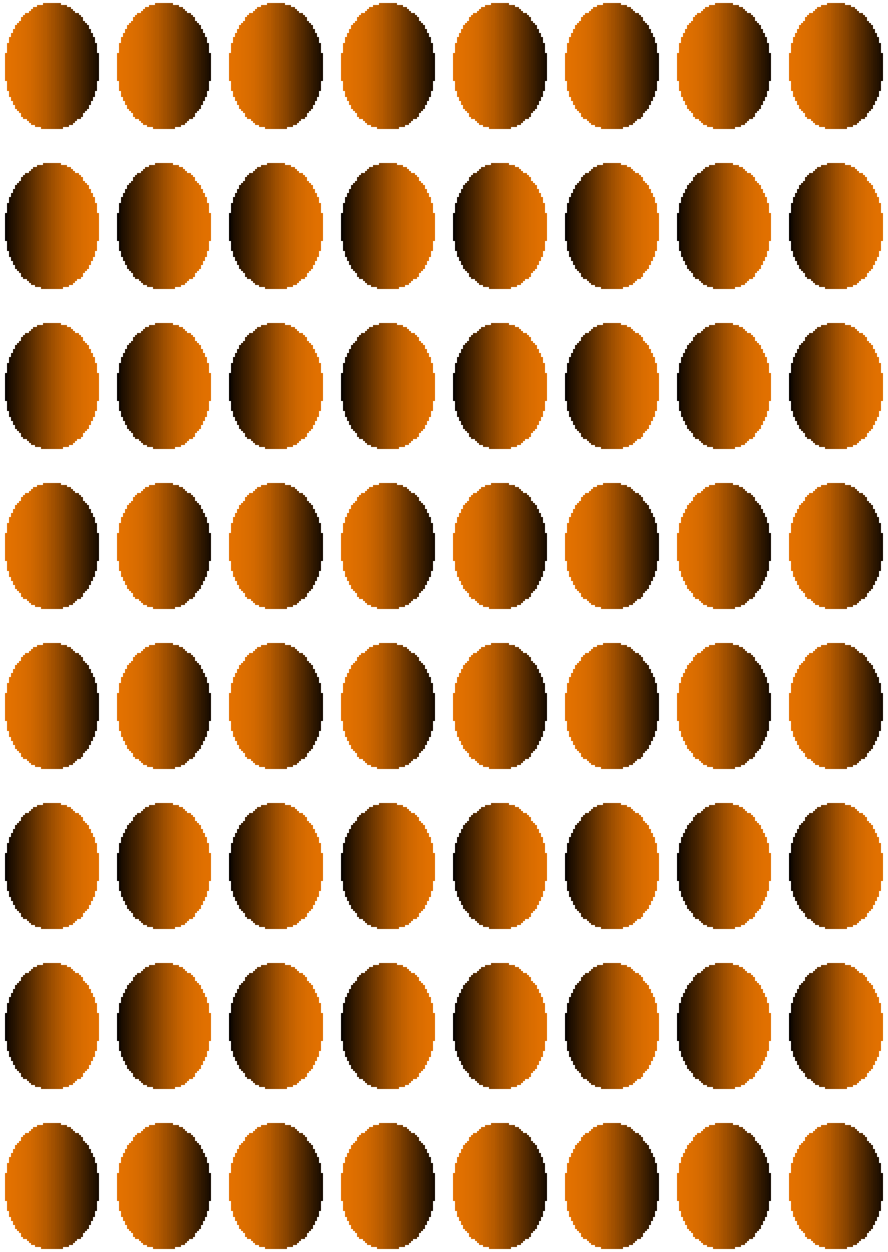

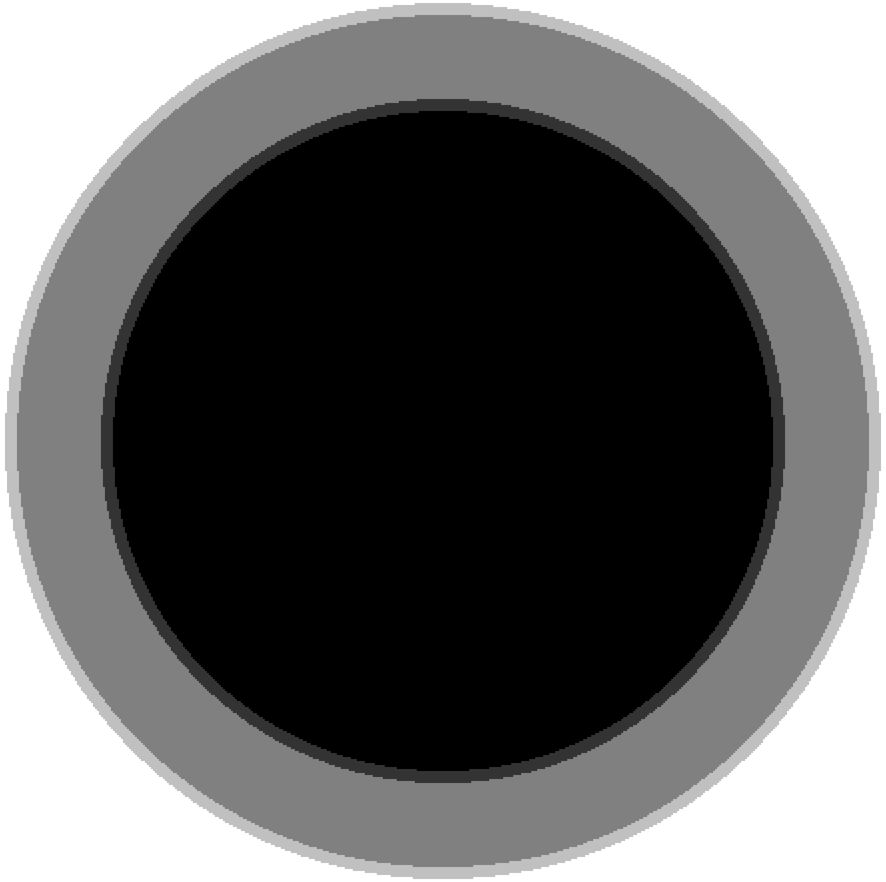

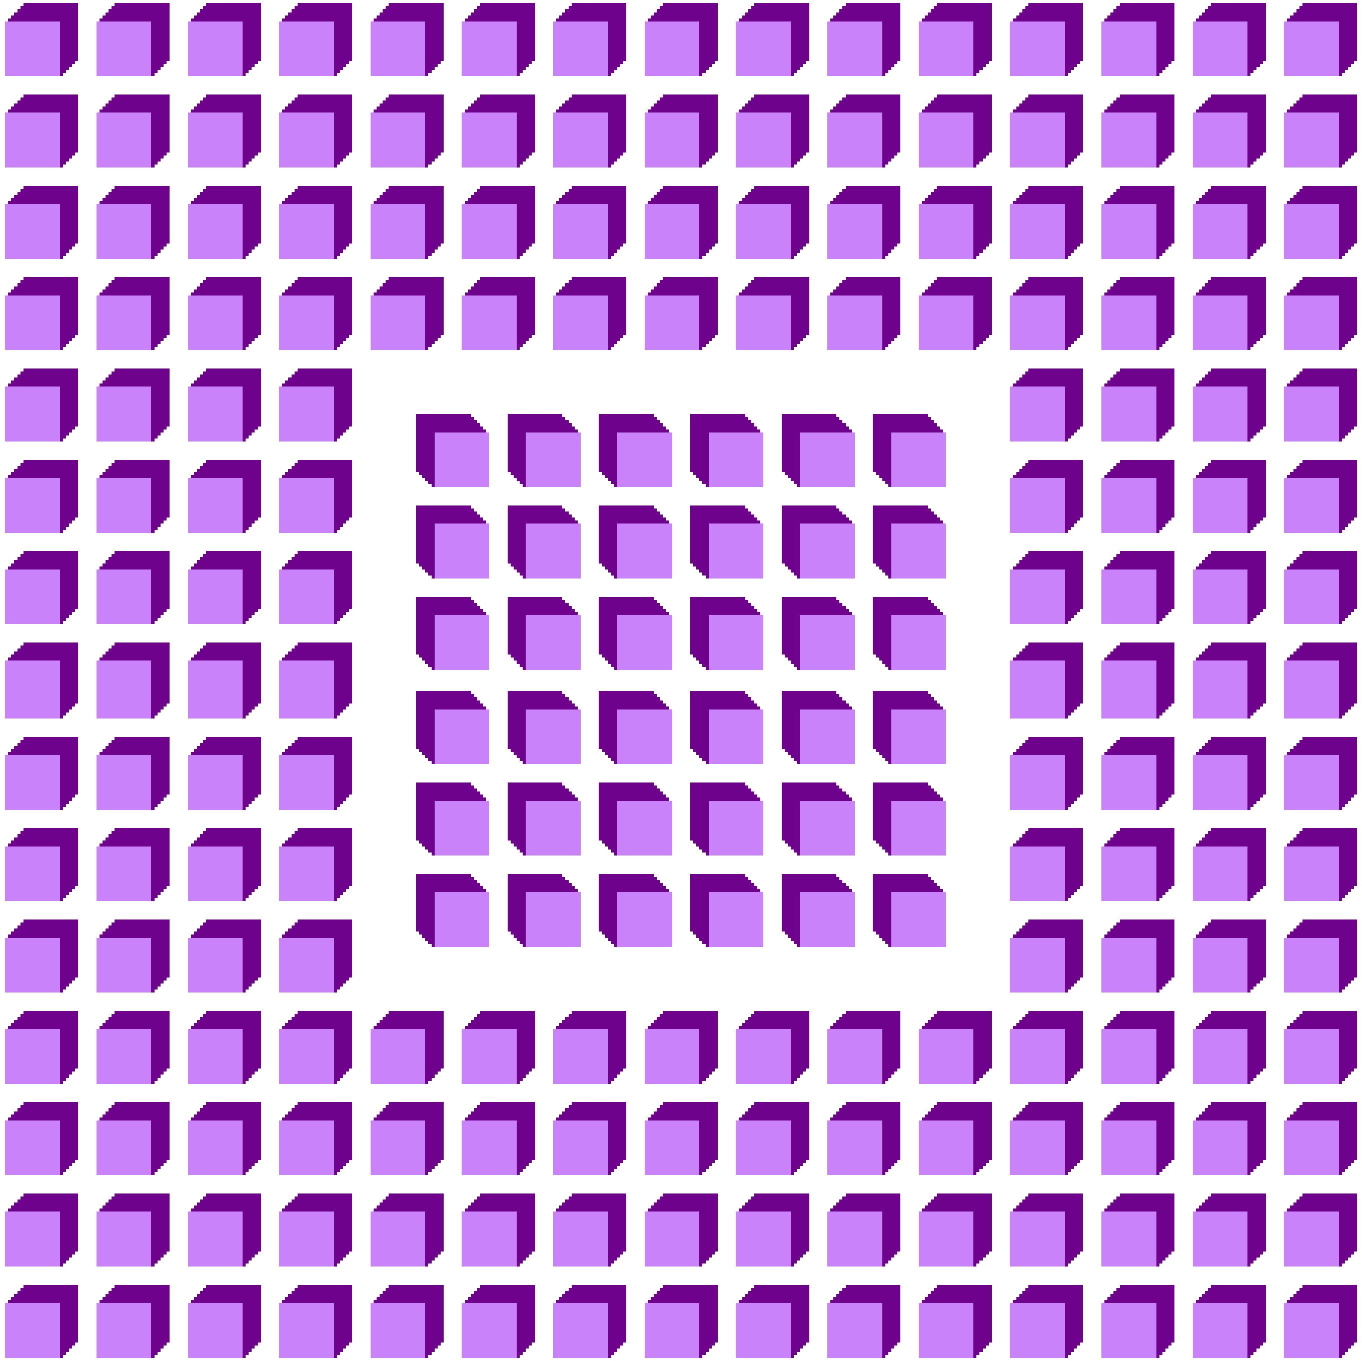

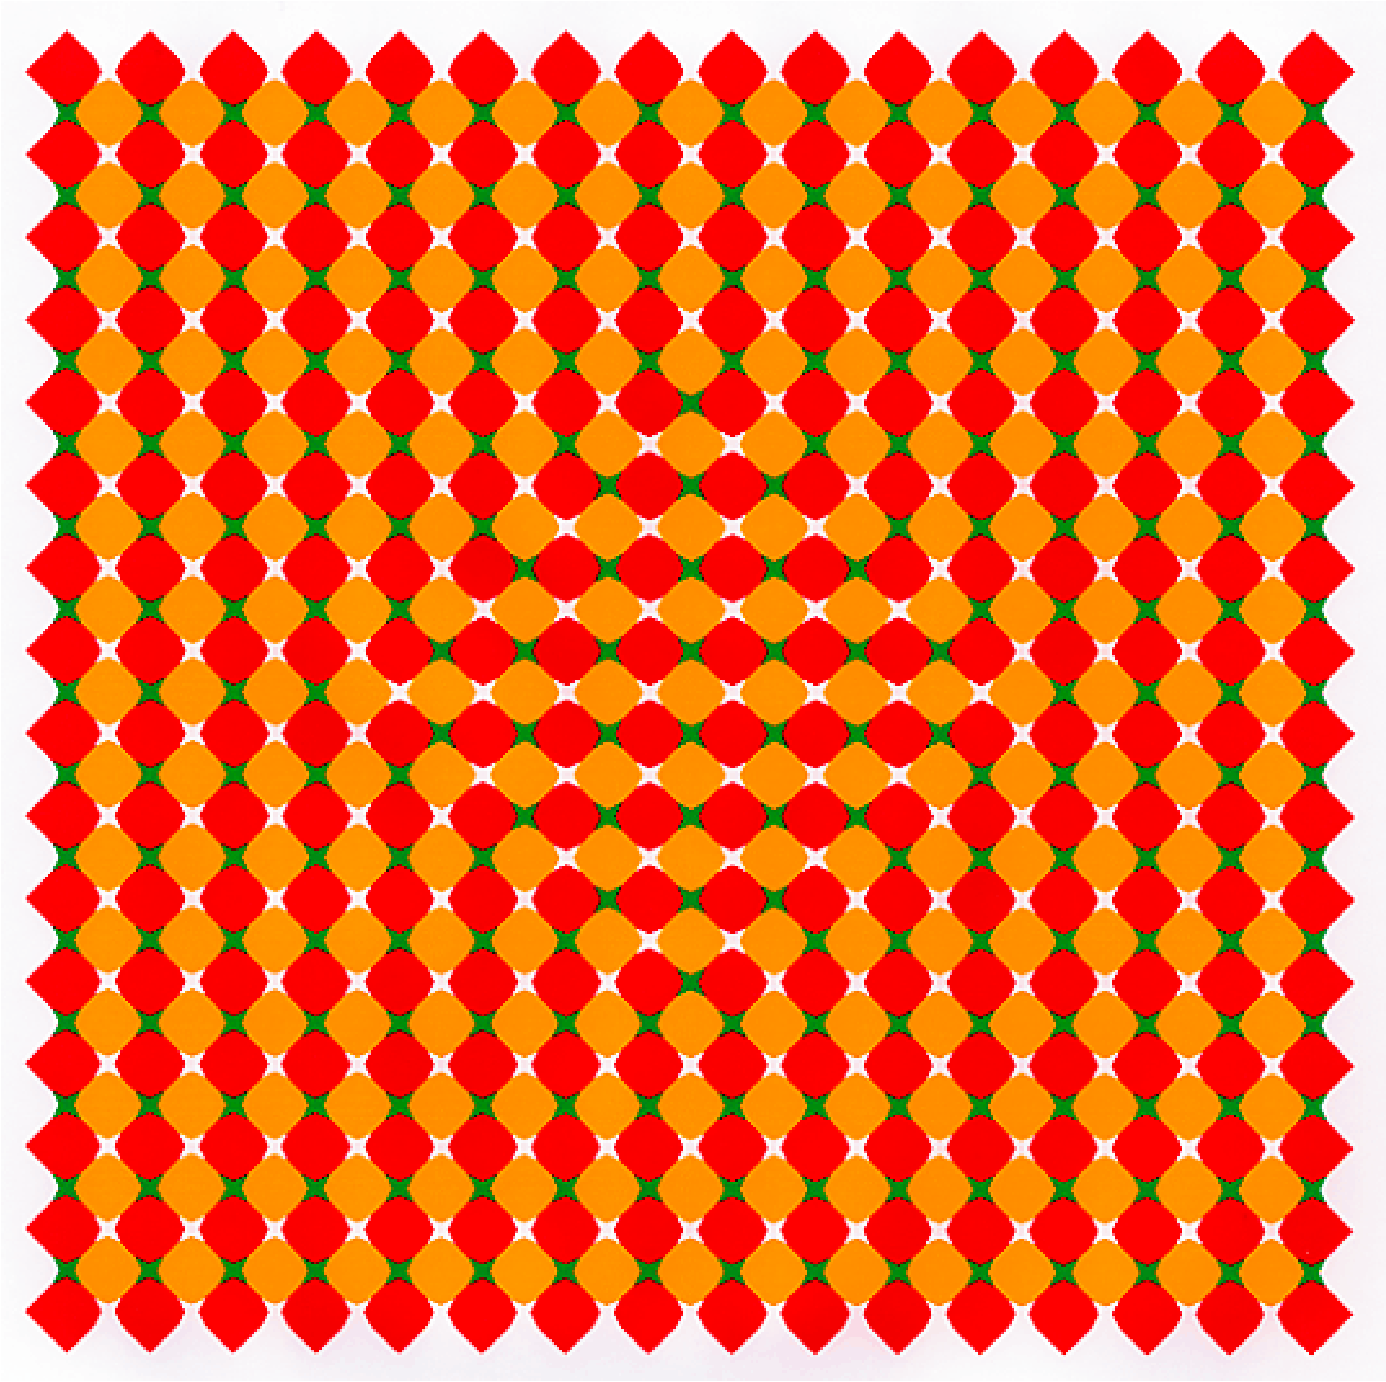

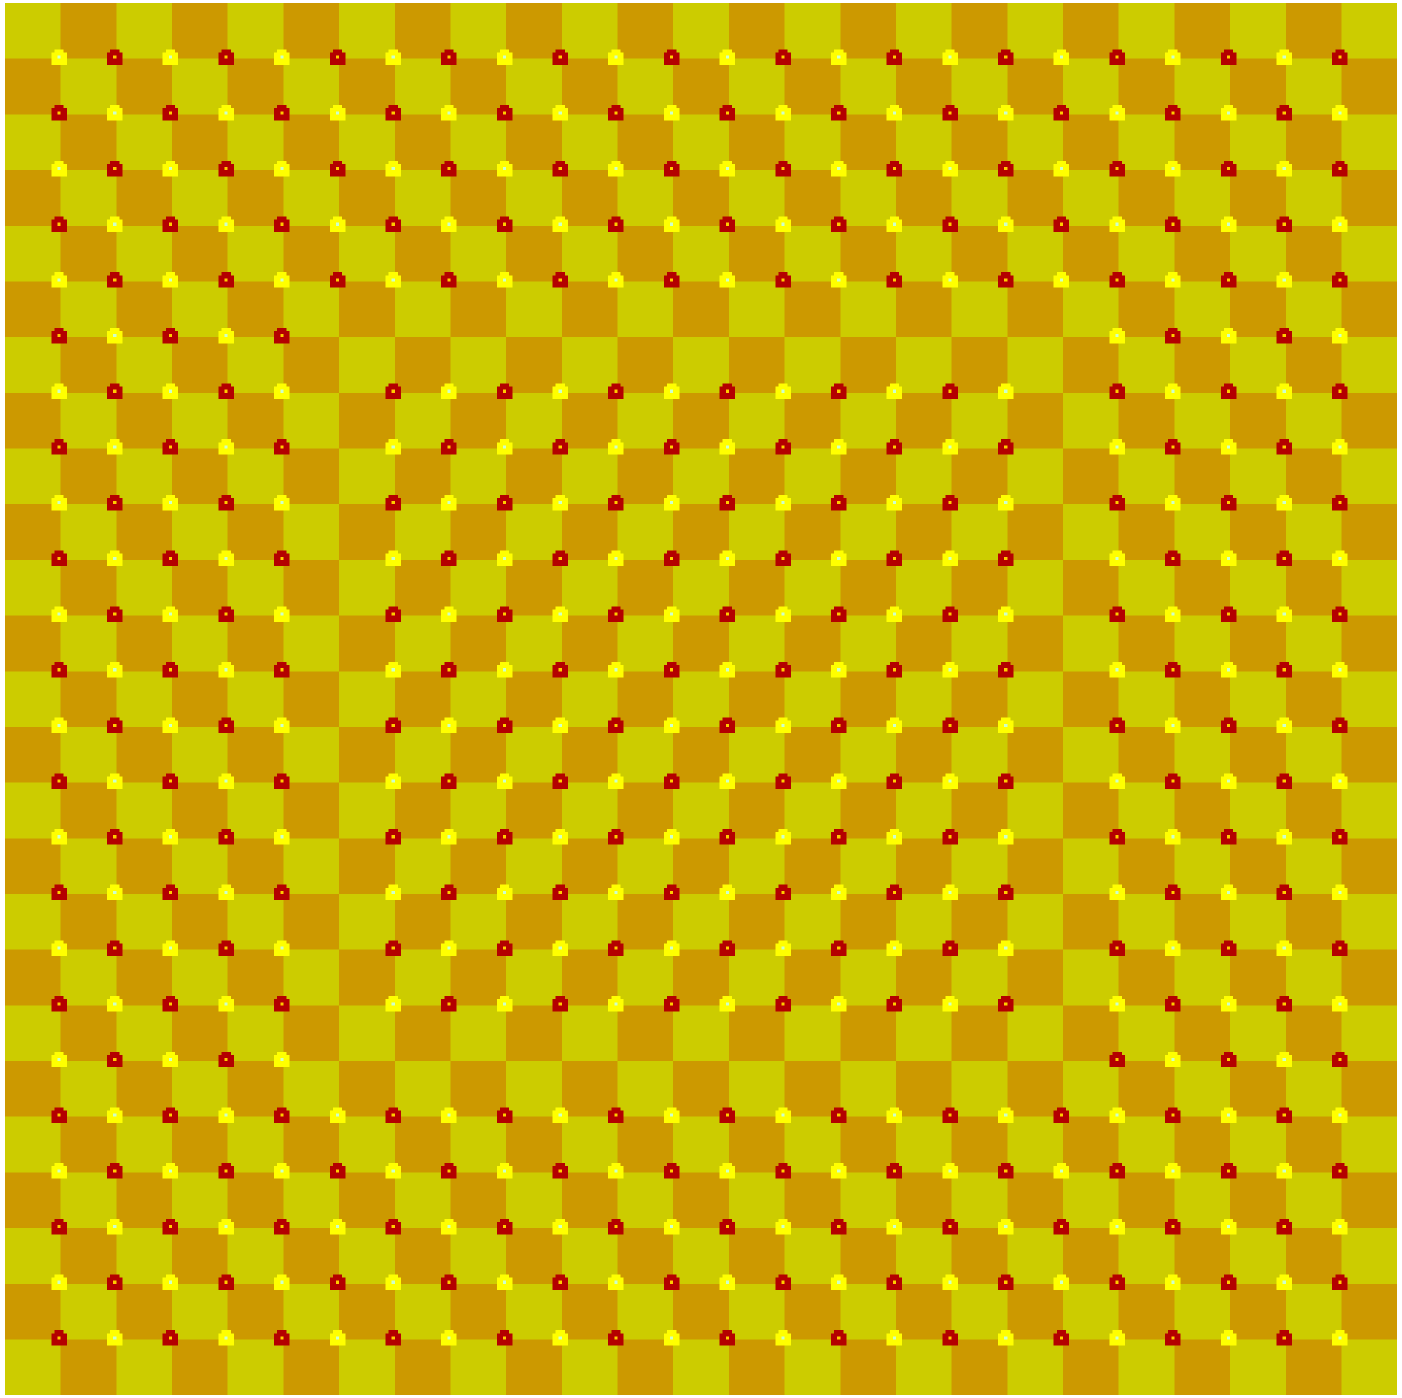

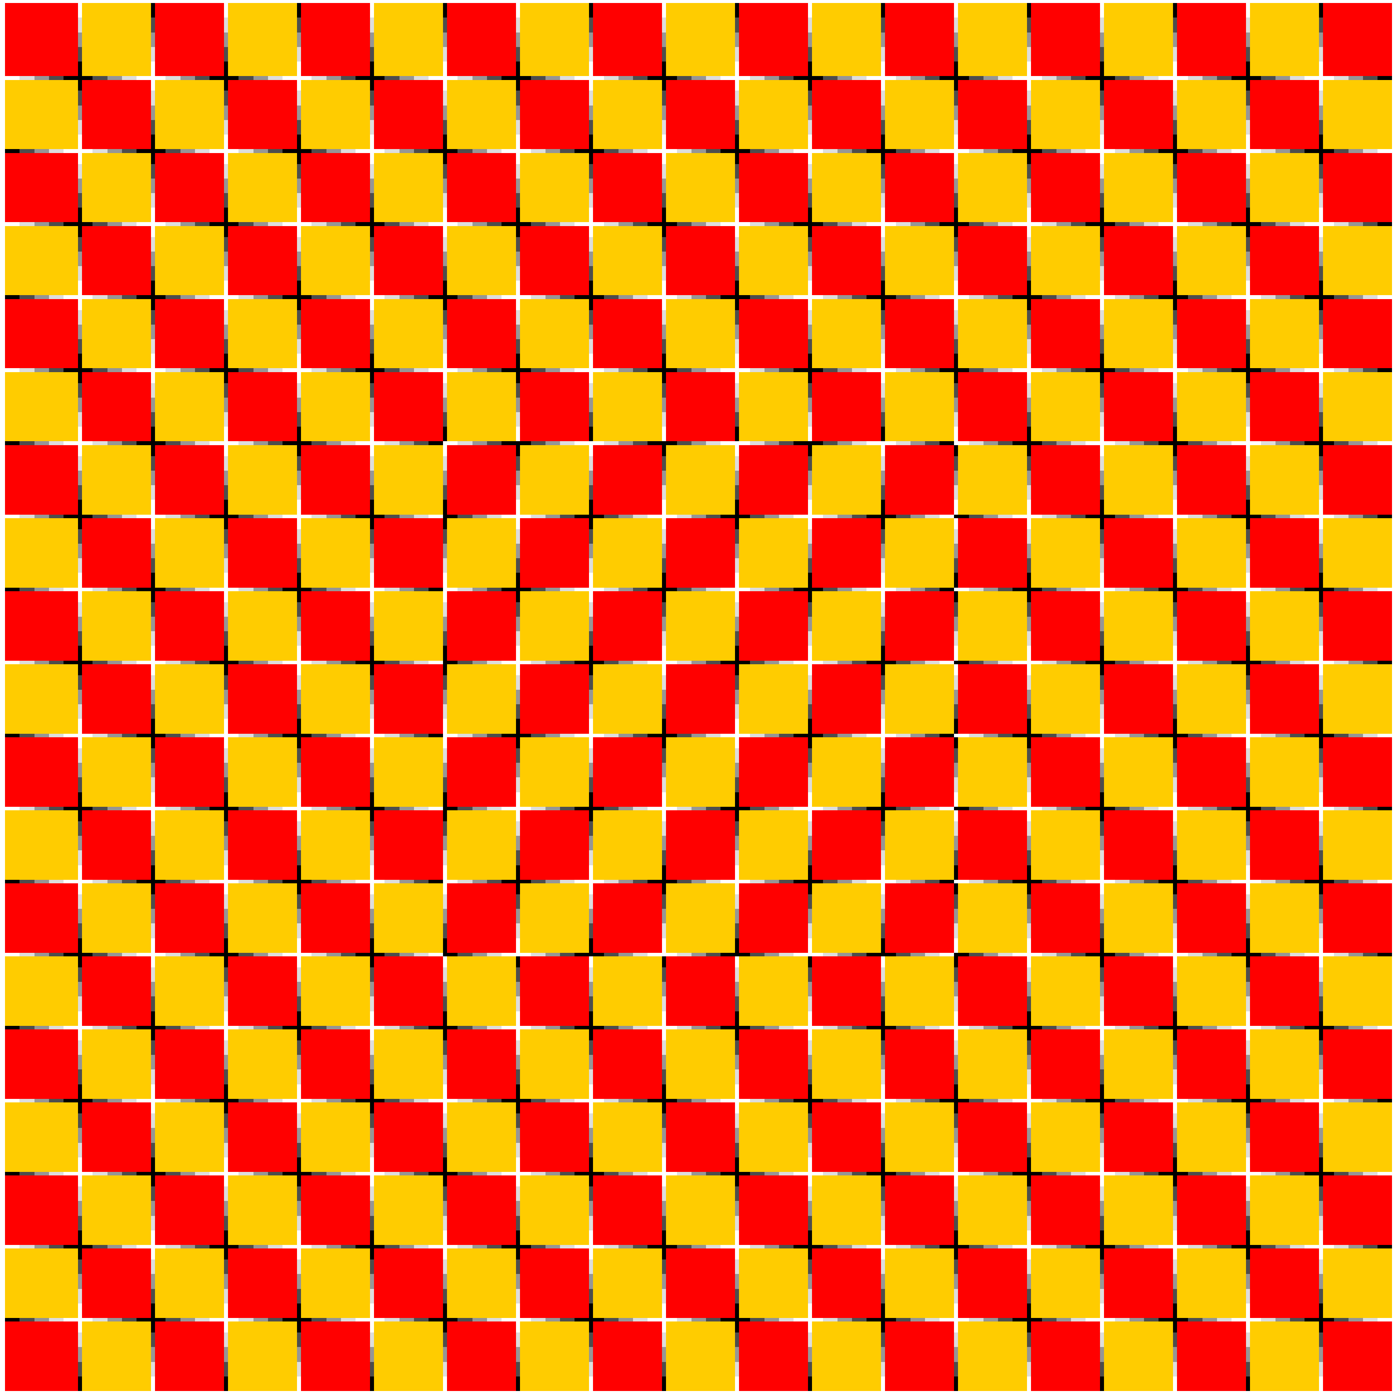

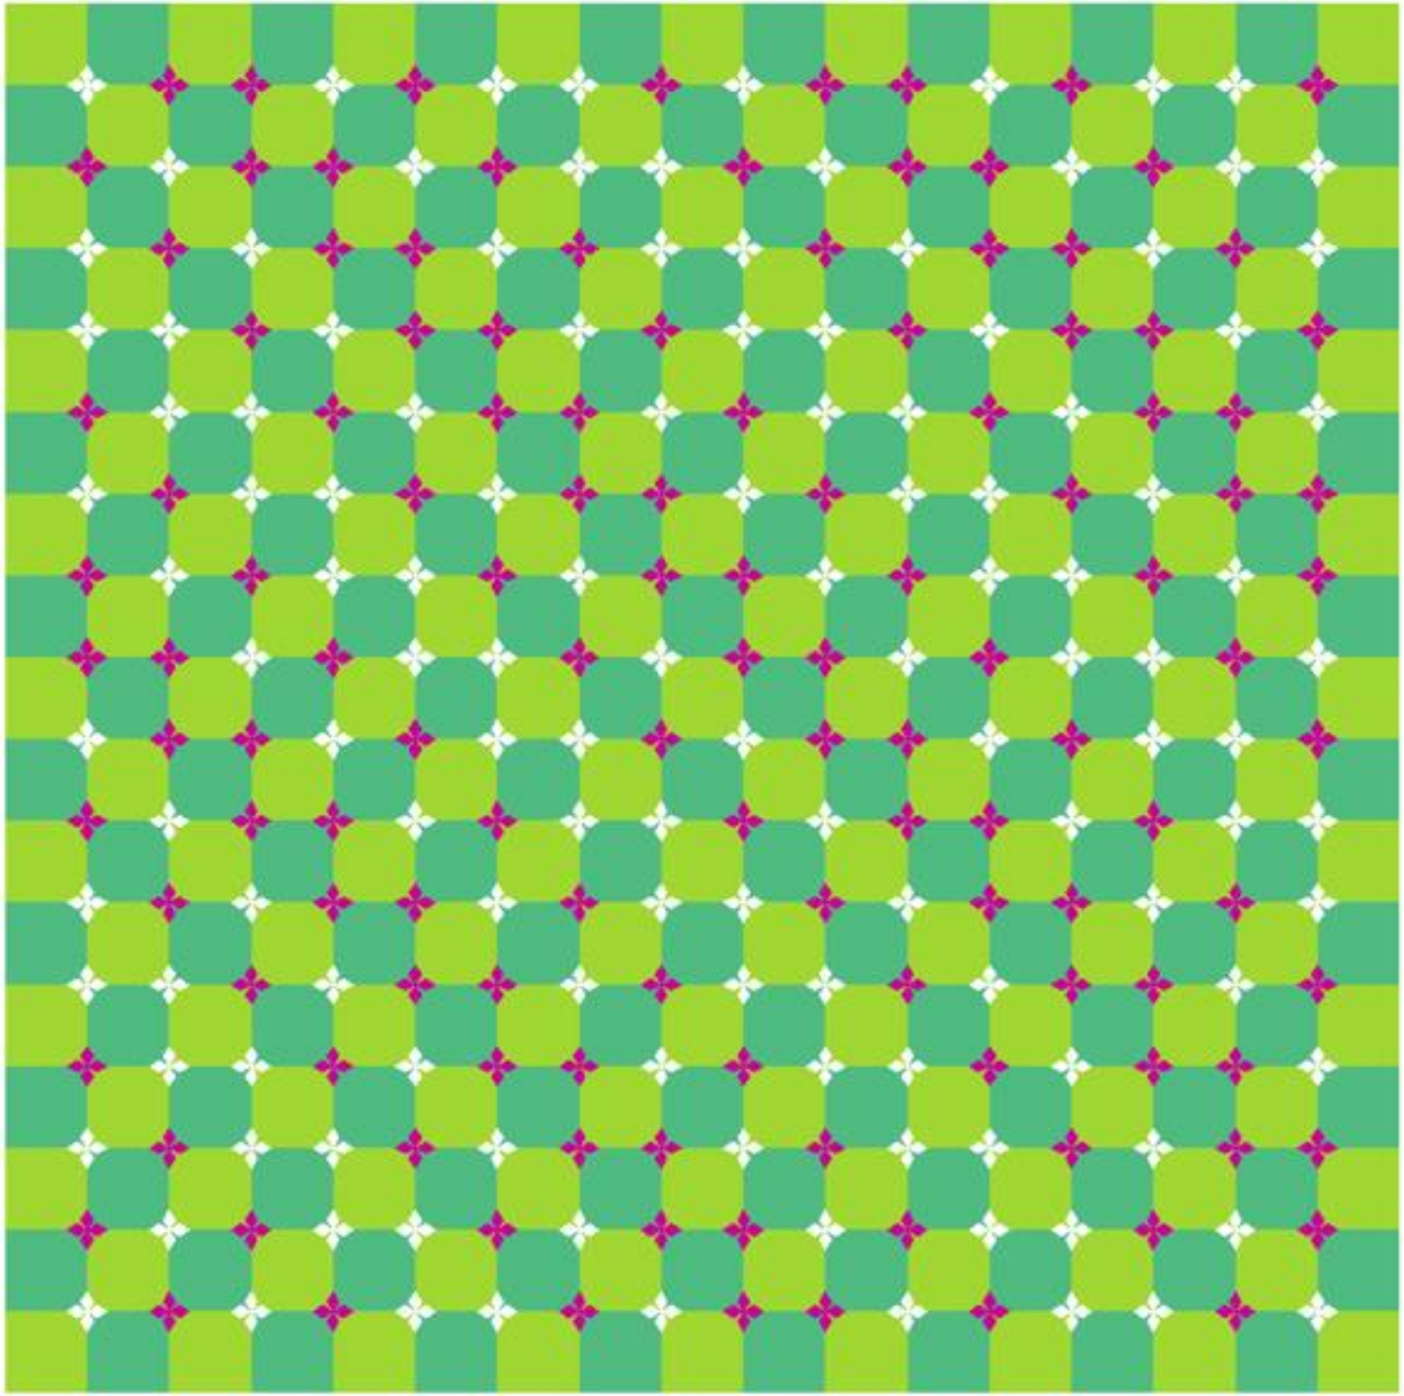

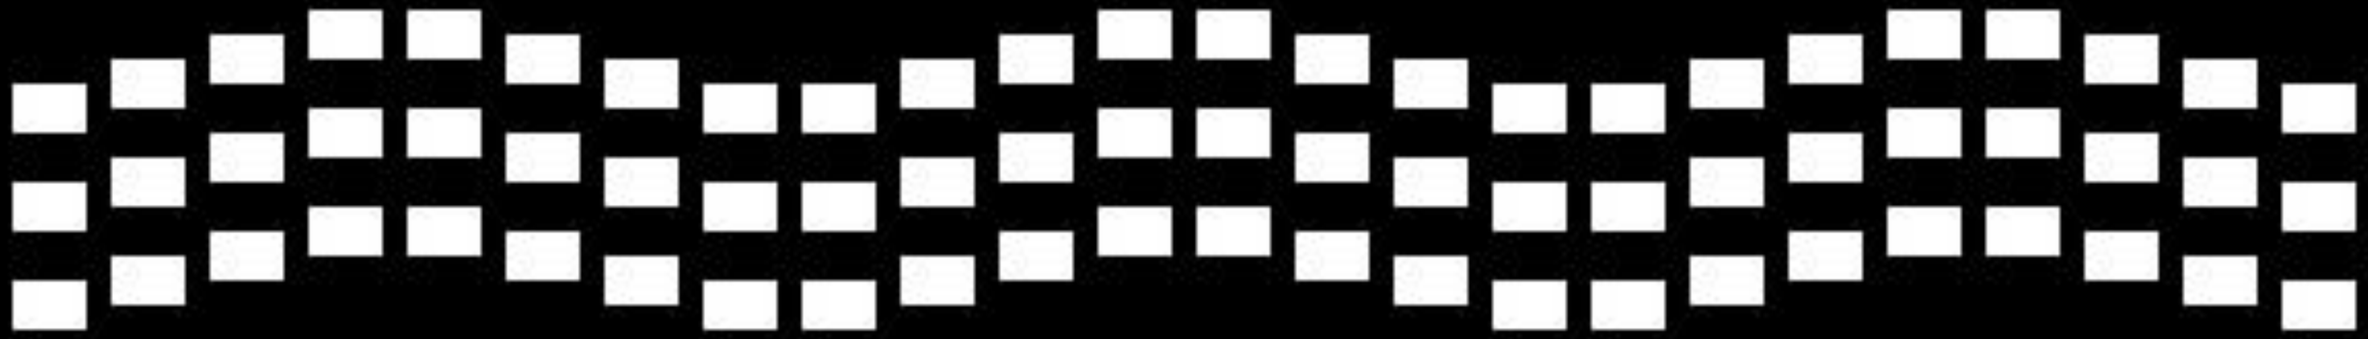

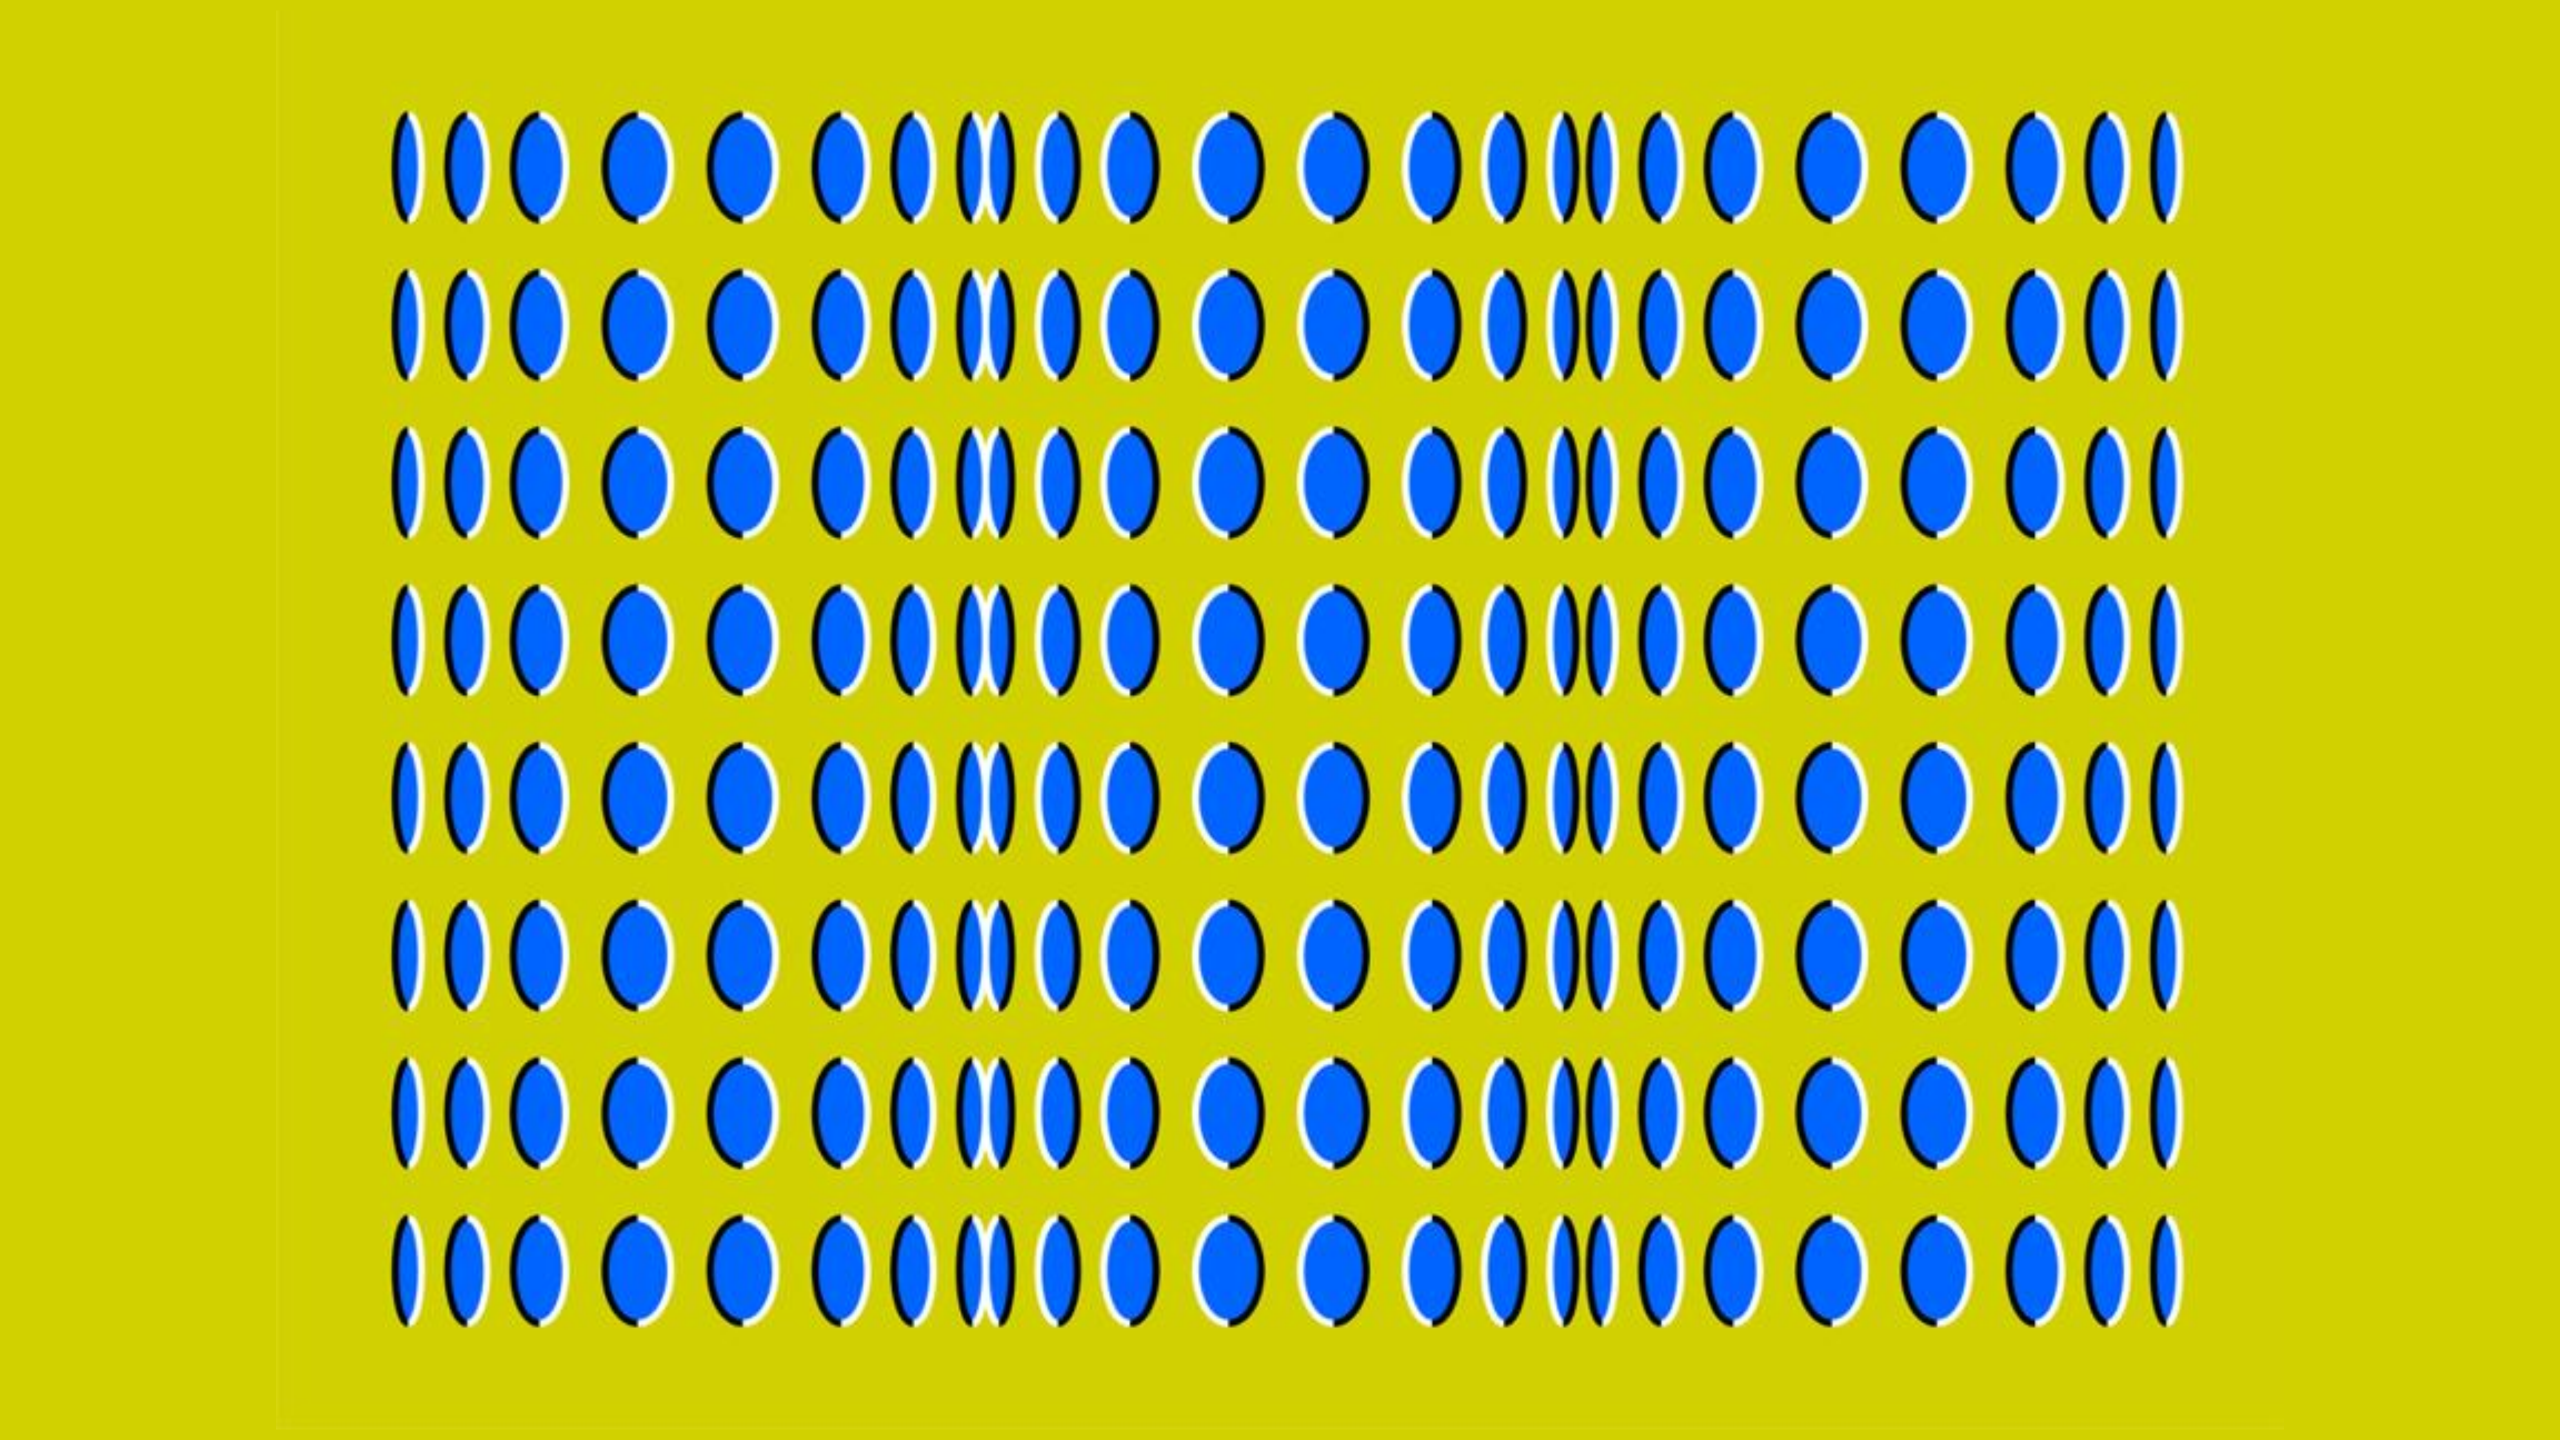

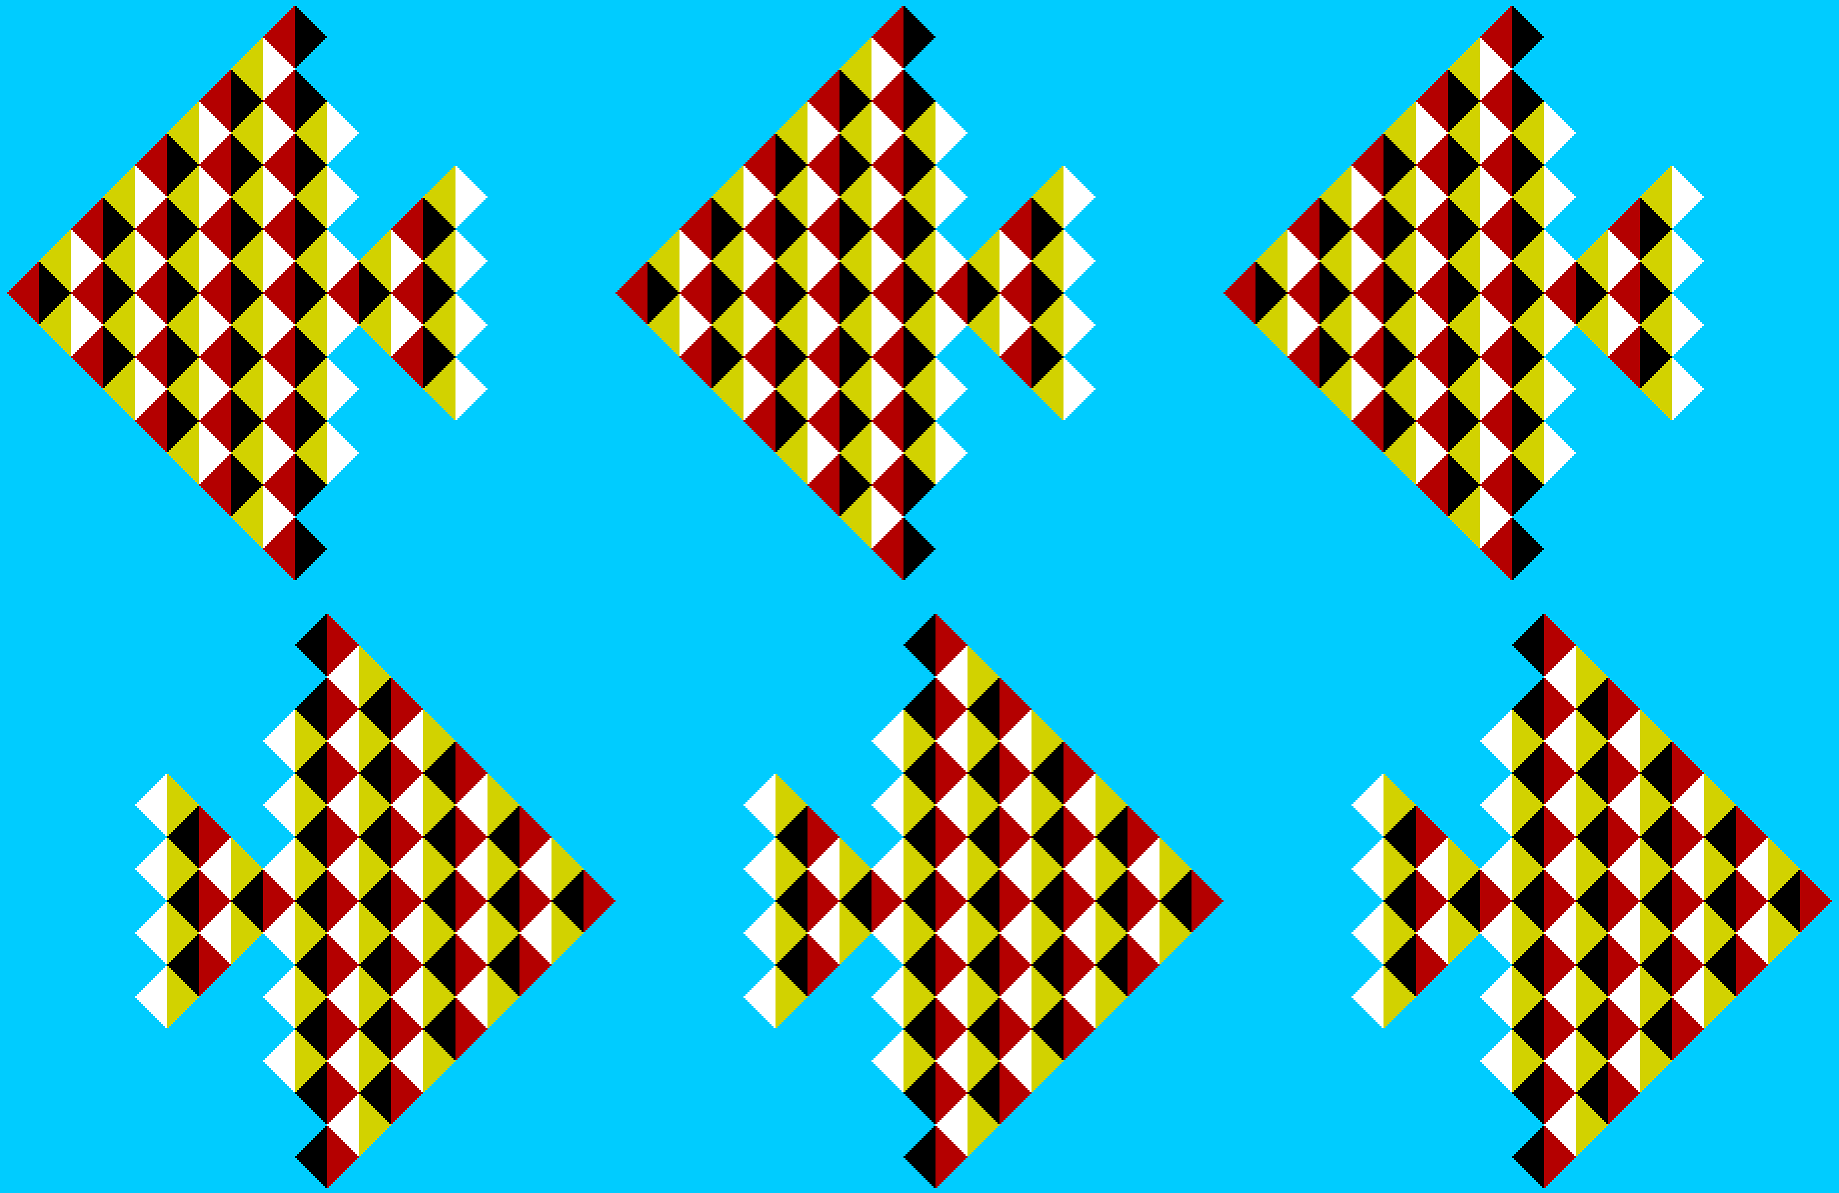

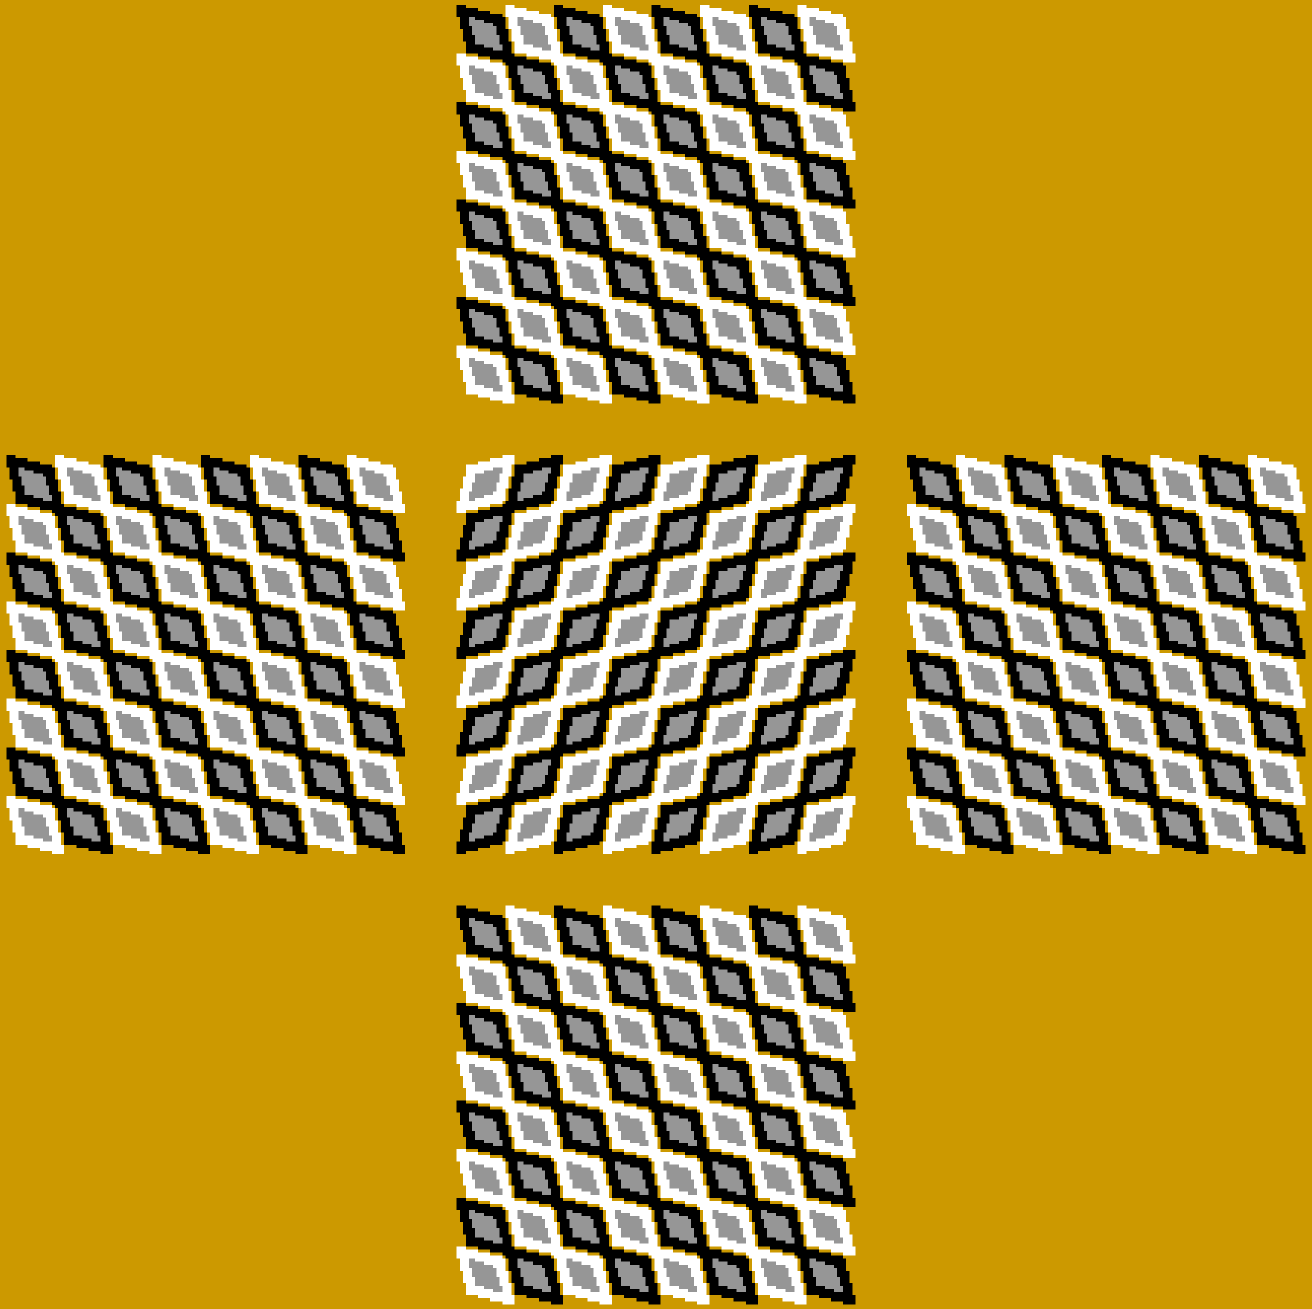

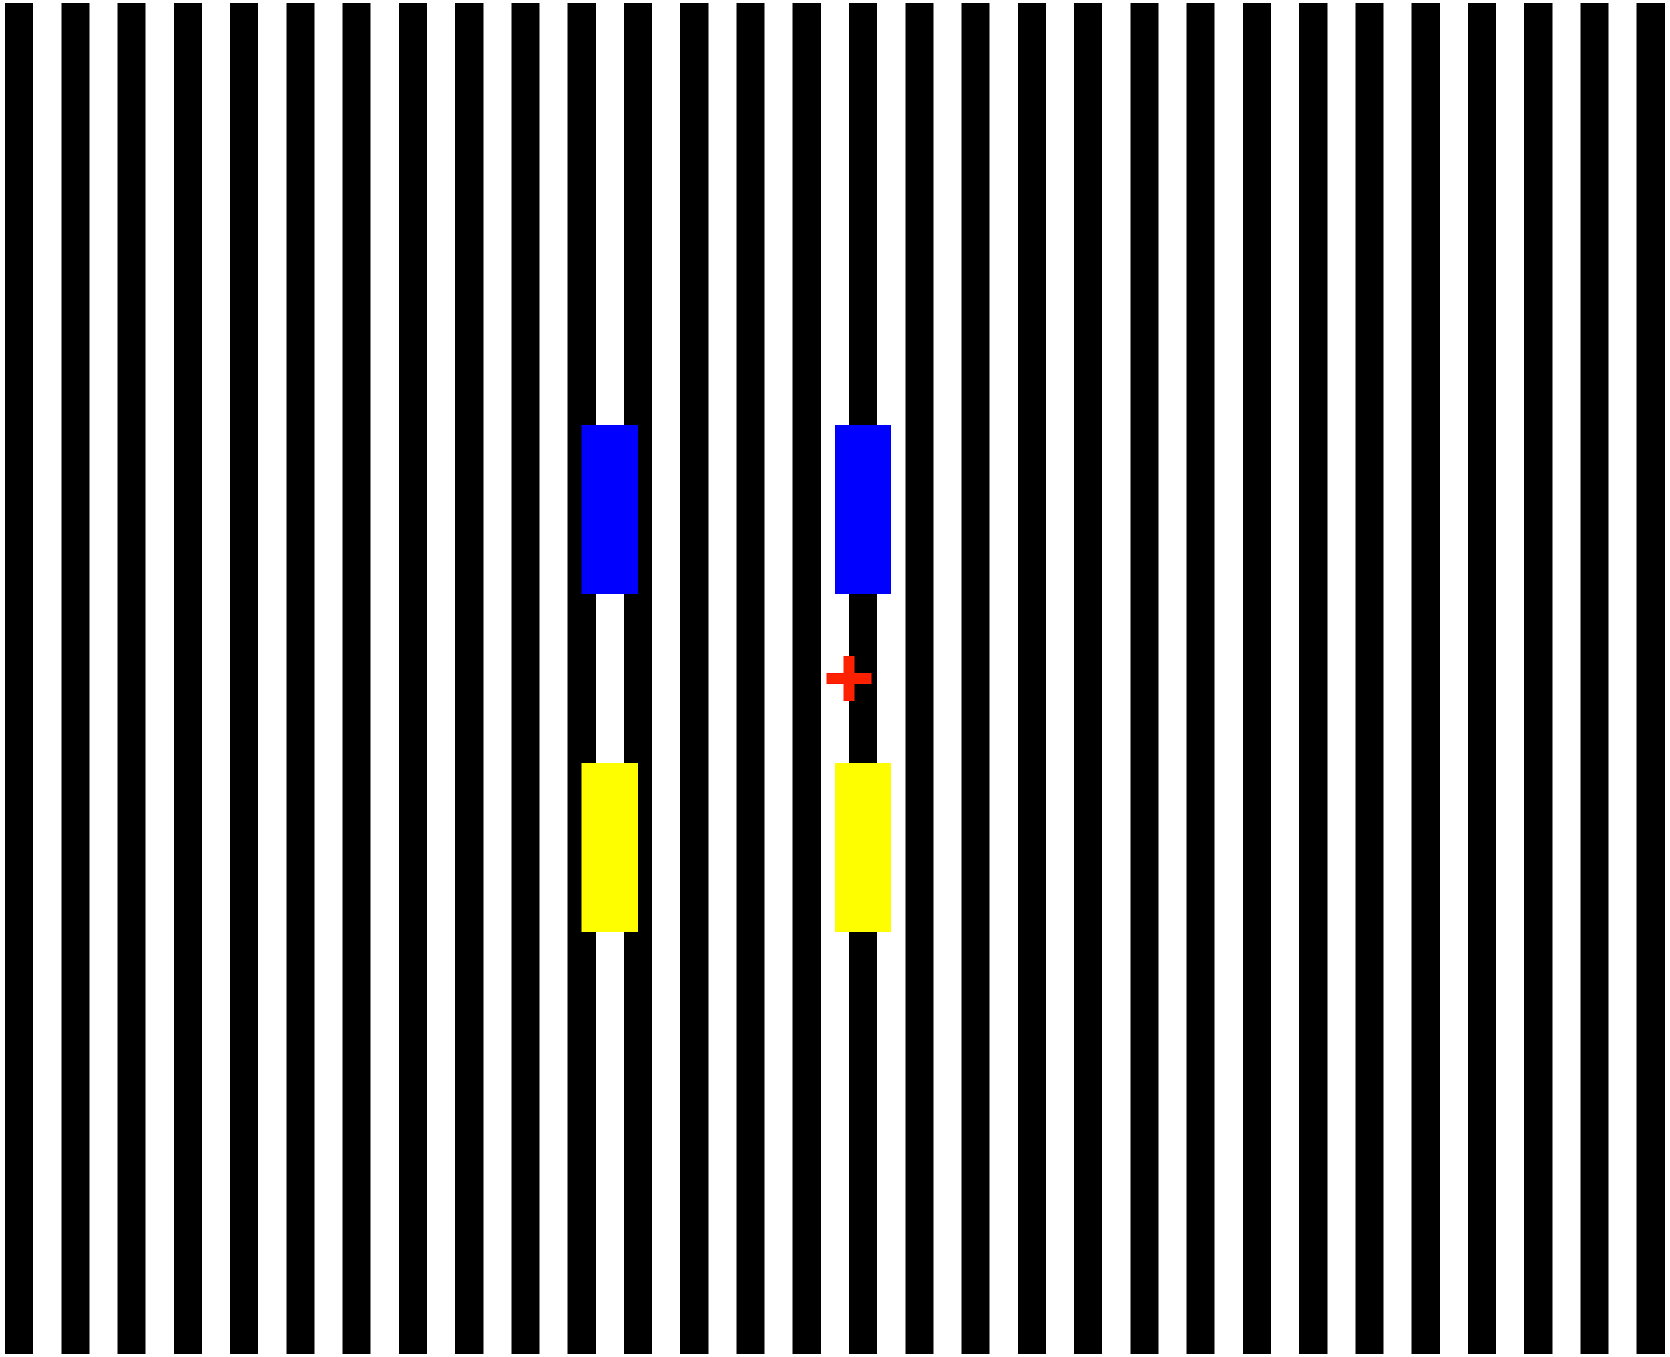

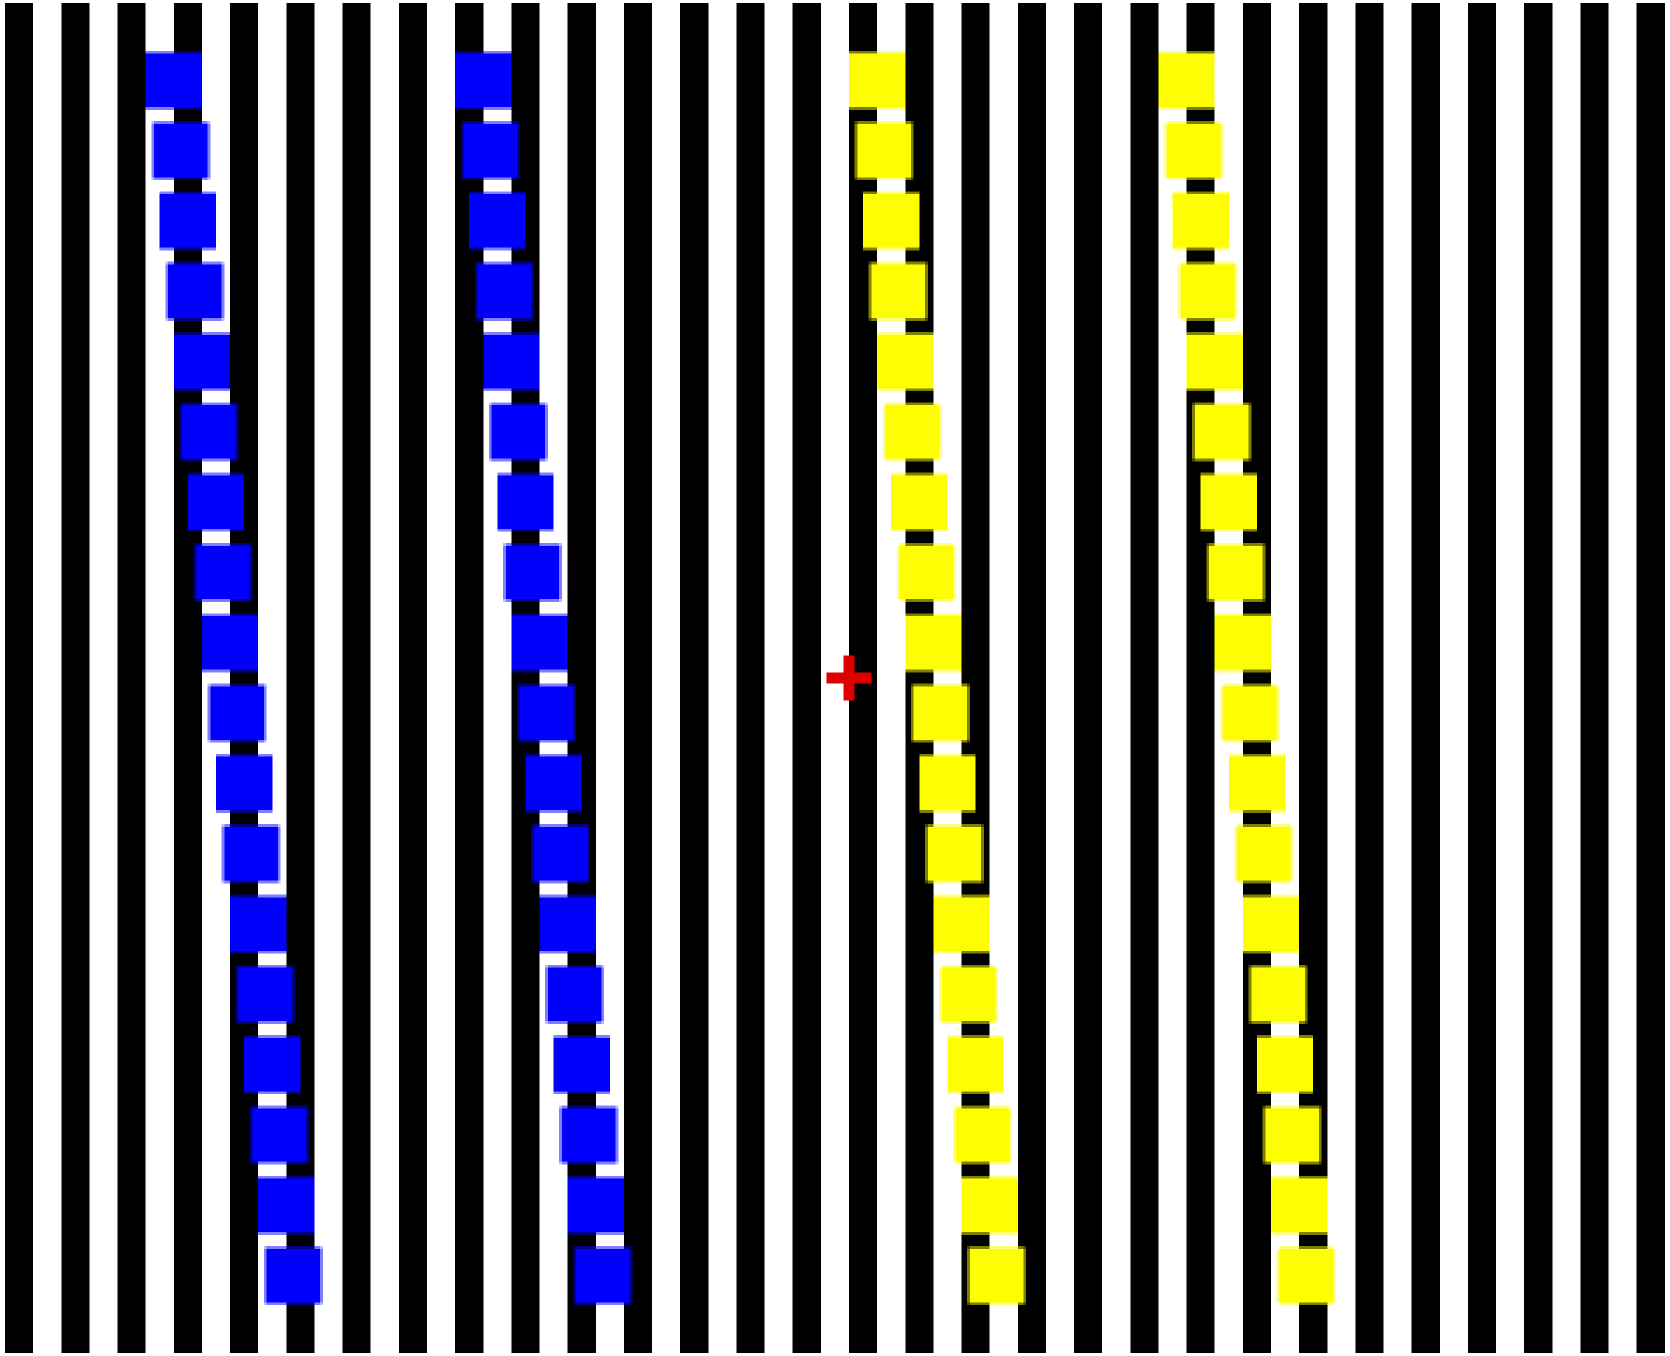

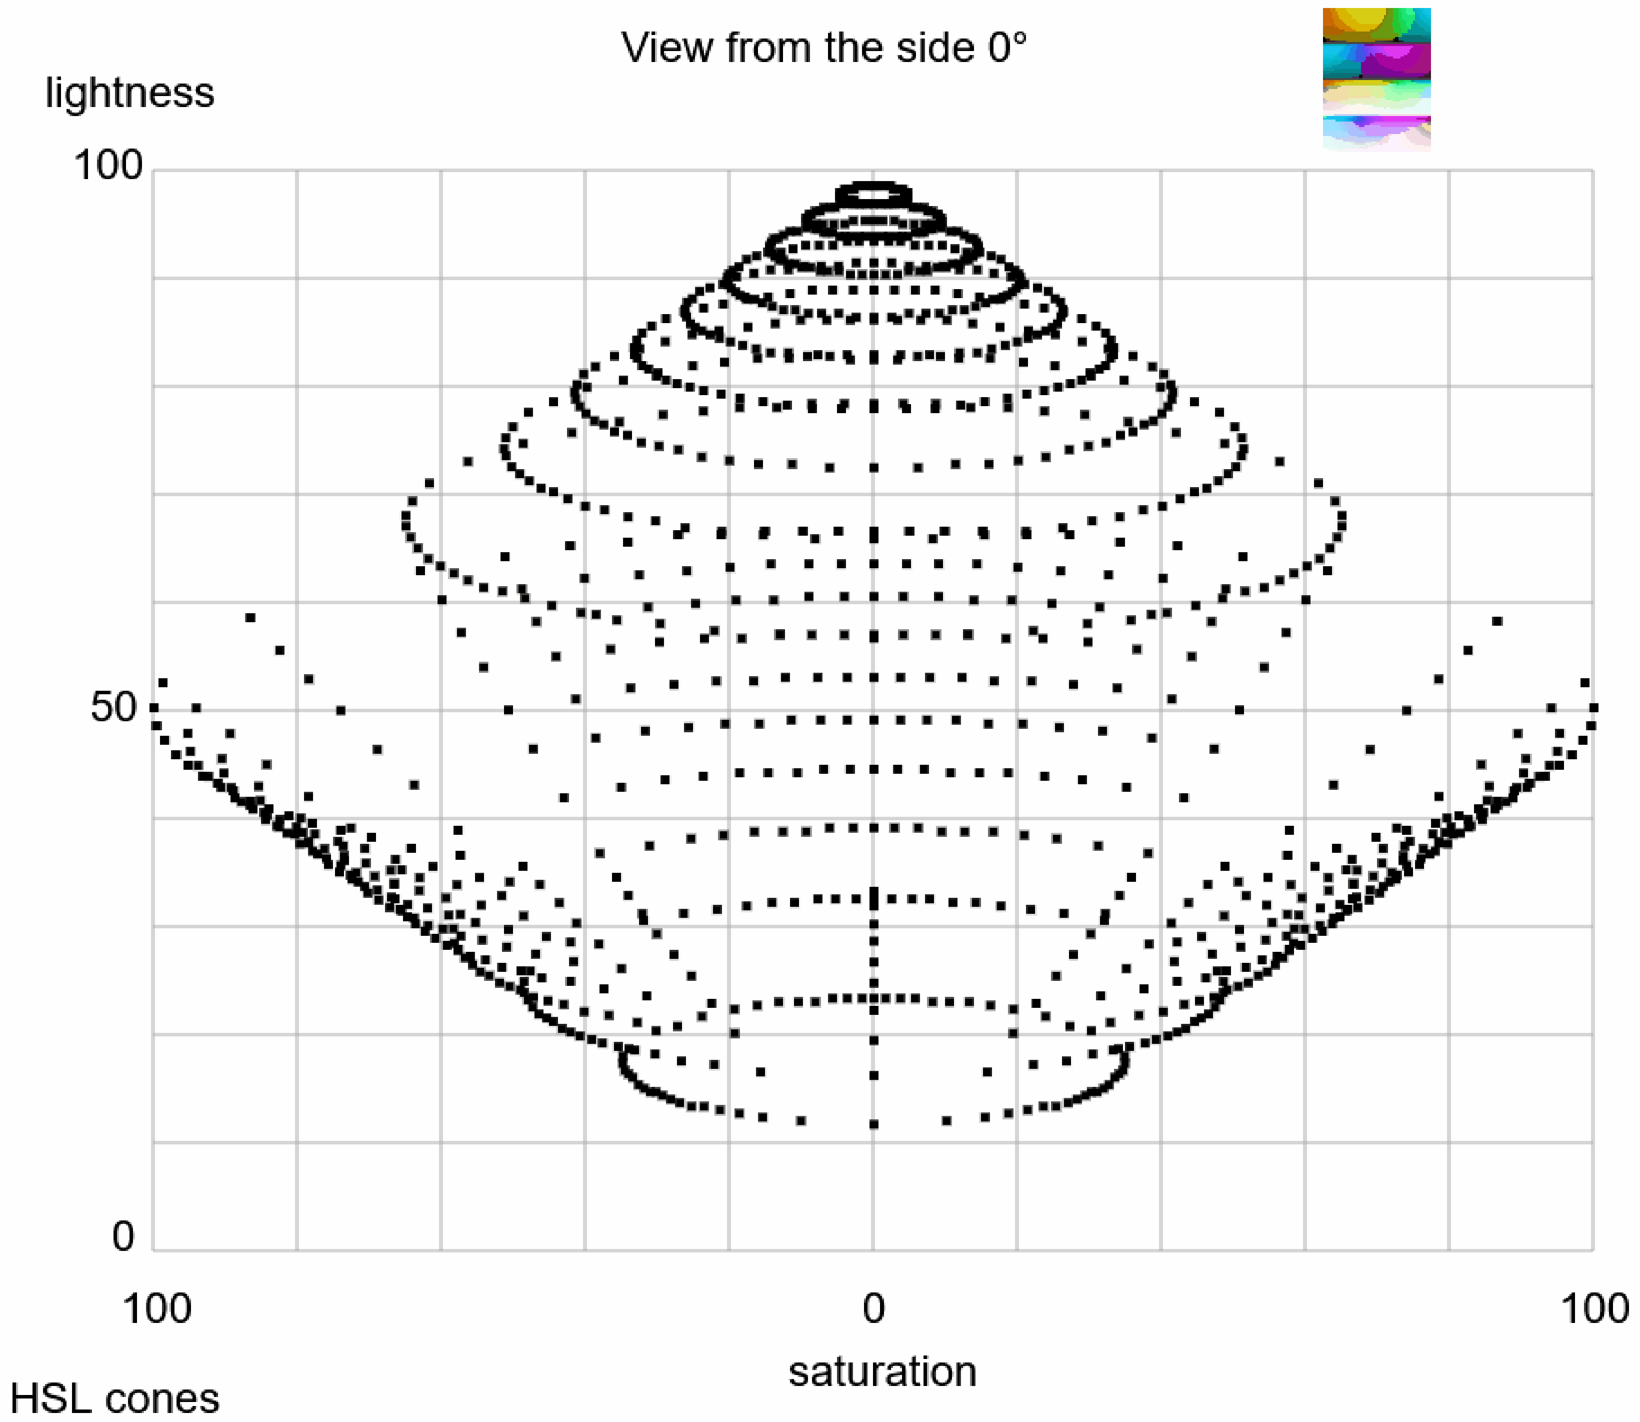
View from the side 0°
lightness
100
50
0
100
0
100
saturation
HSL cones

Illusions of motion also occur naturally!

‼

“Phi phenomenon / apparent motion”

The visual system distinguishes between *light* and *color*.

unconscious inference

But motion illusions are so much more than this!


Color
adaptation

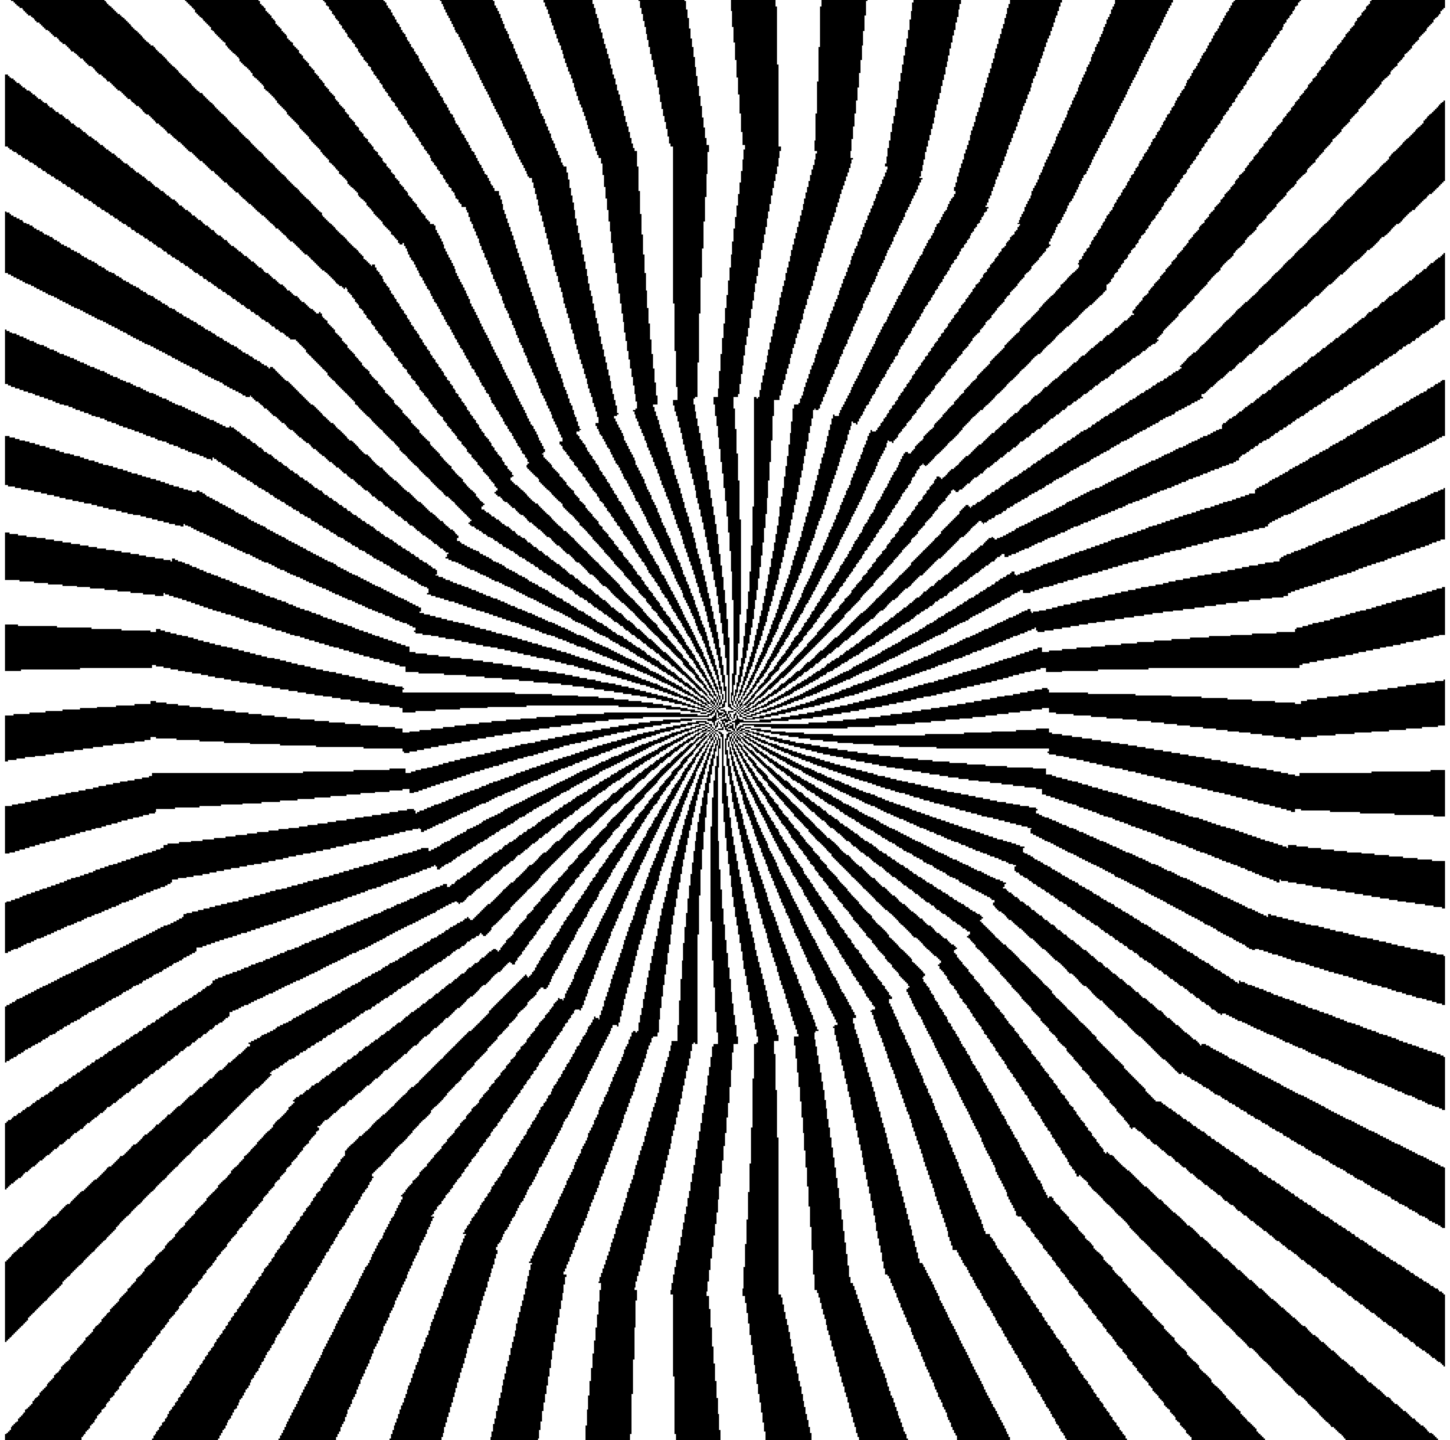

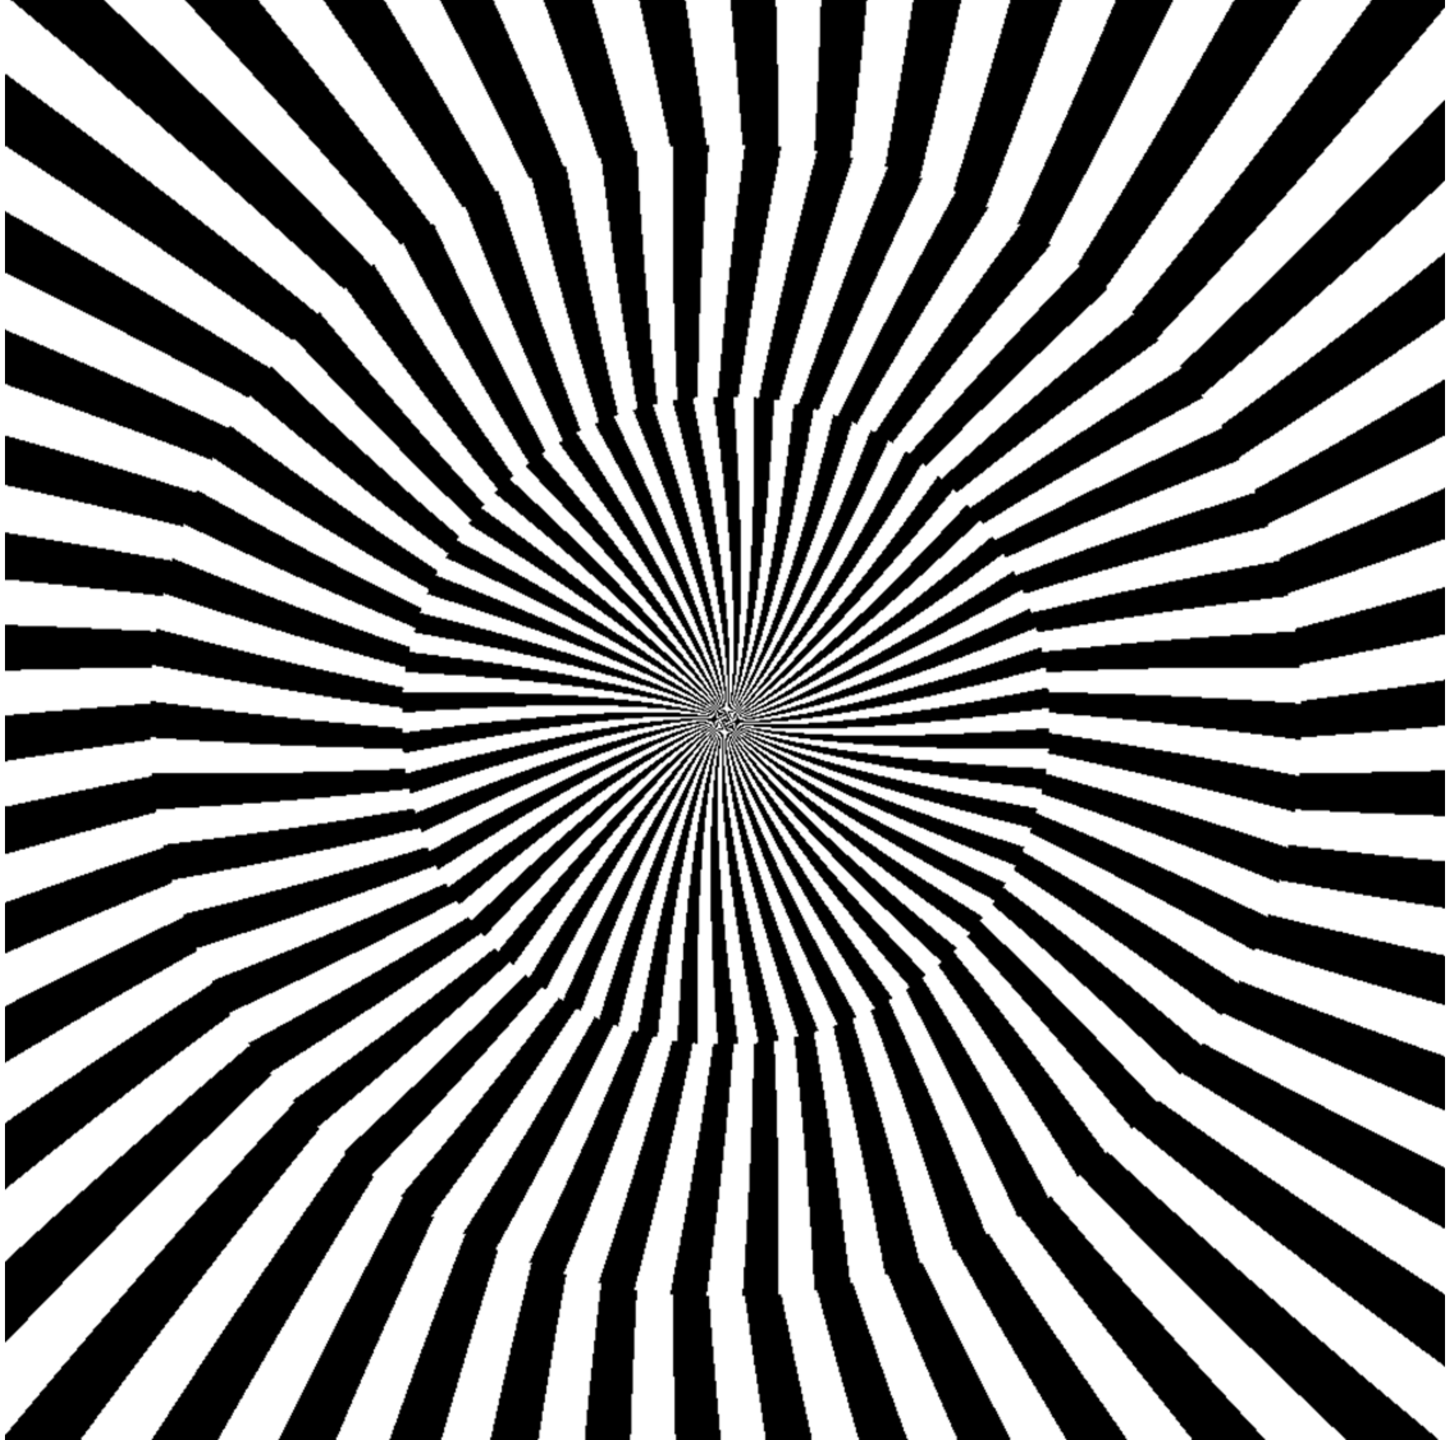

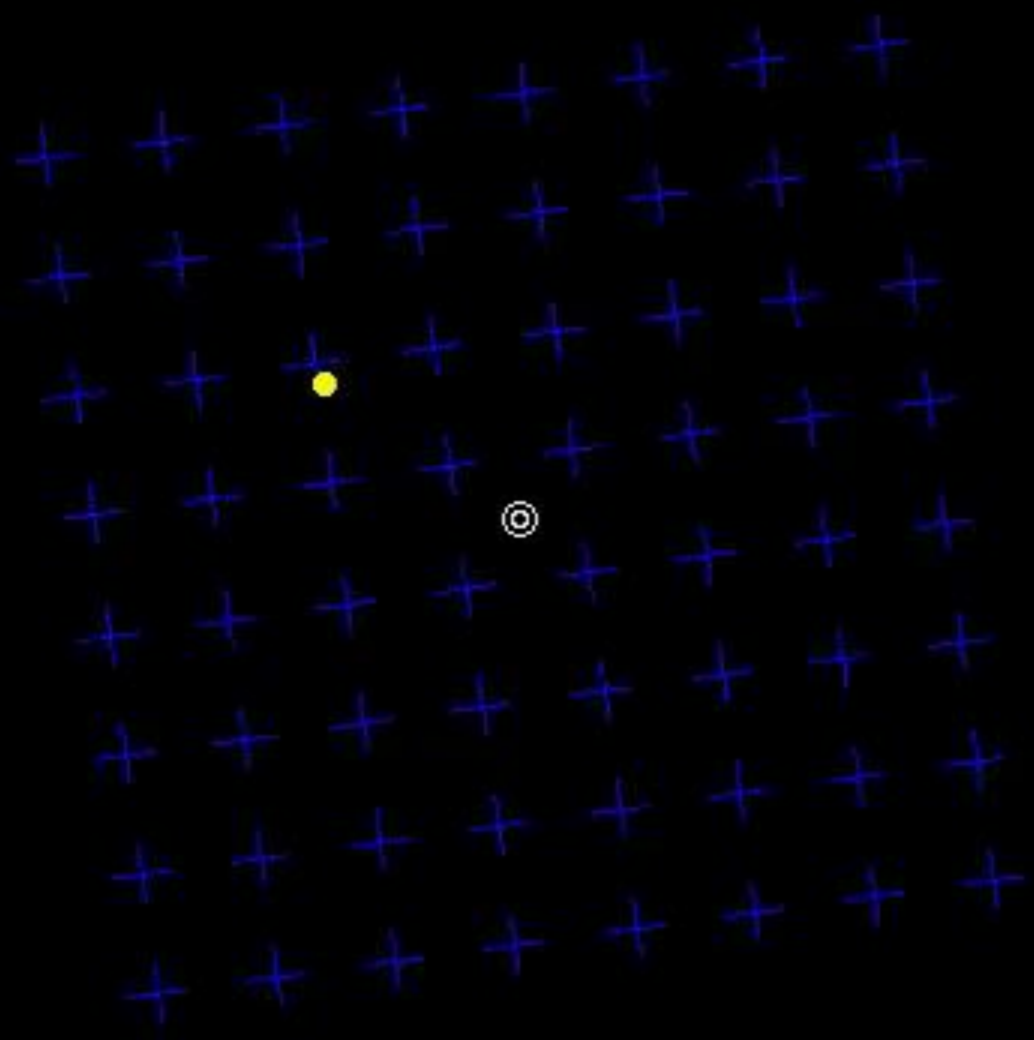

Motion-induced blindness!

Motion-induced blindness!

A perceptual scotoma?

‼

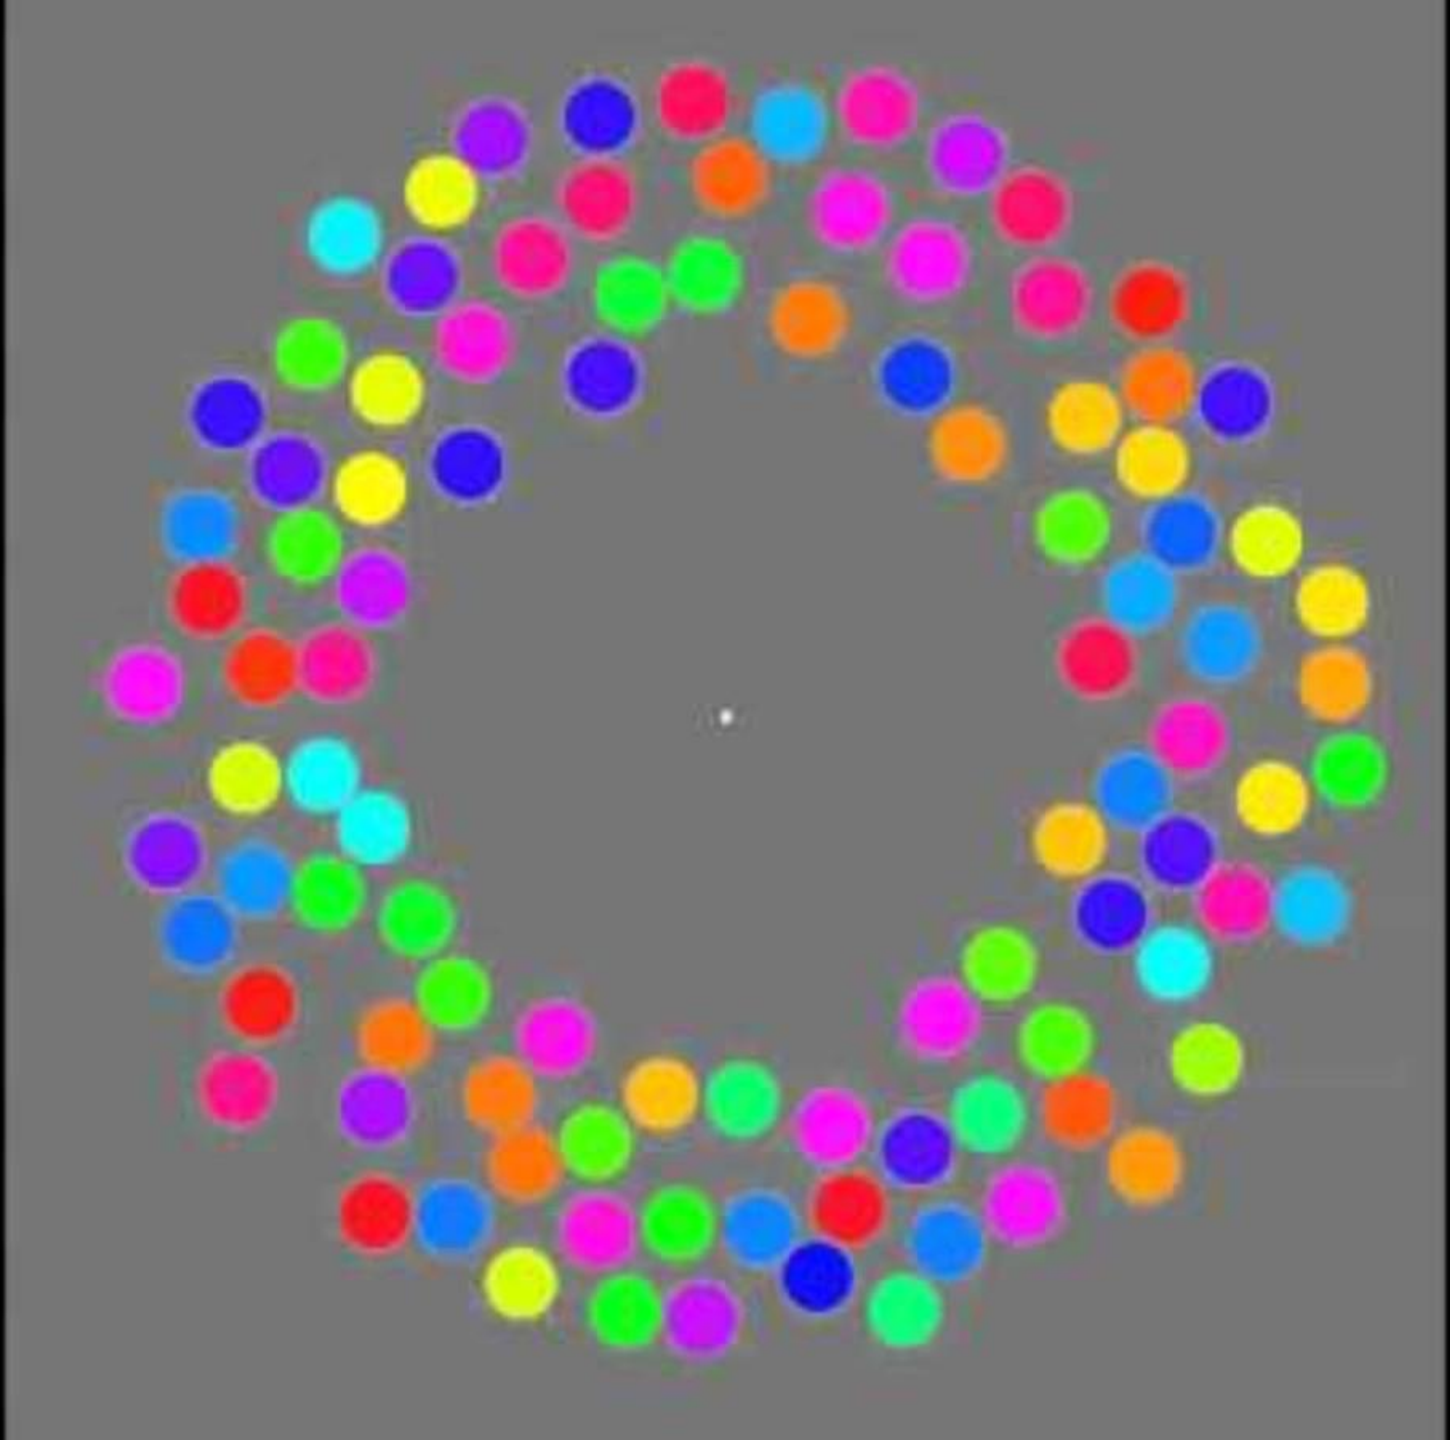

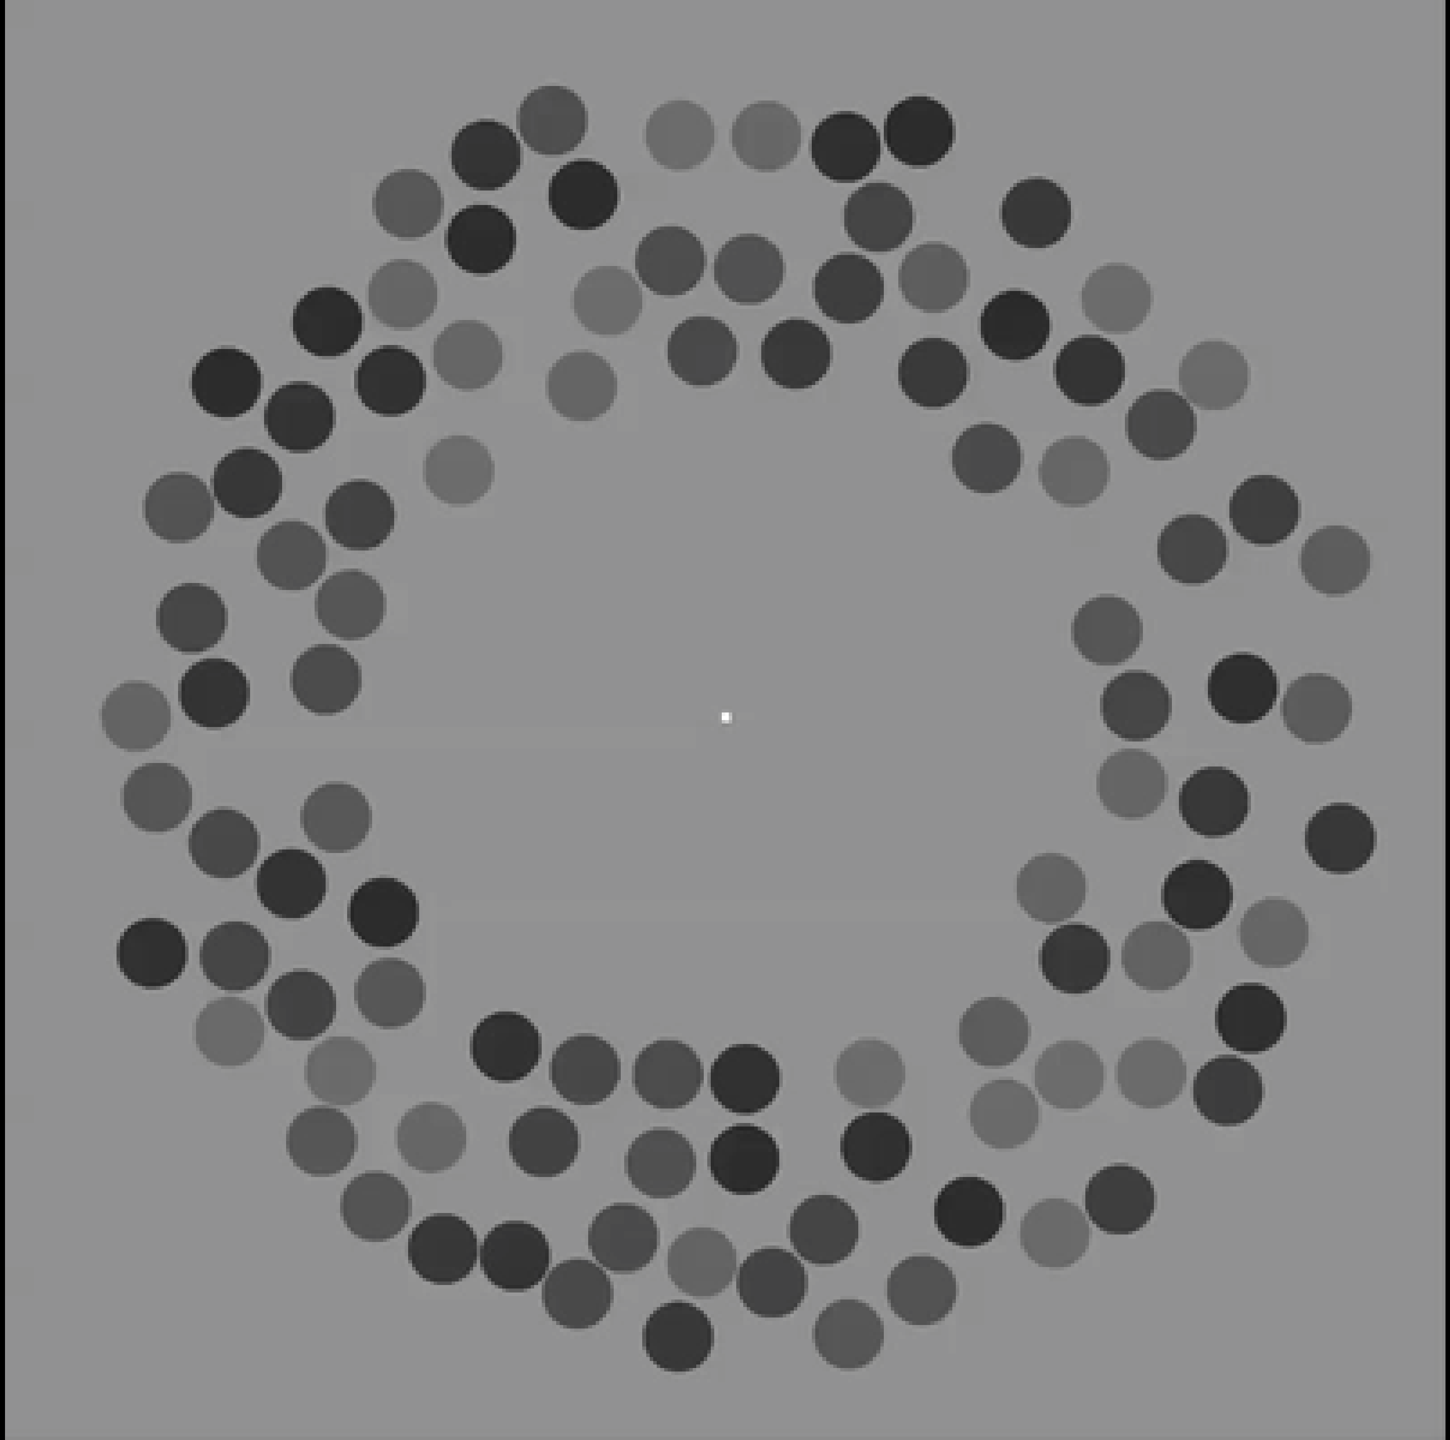

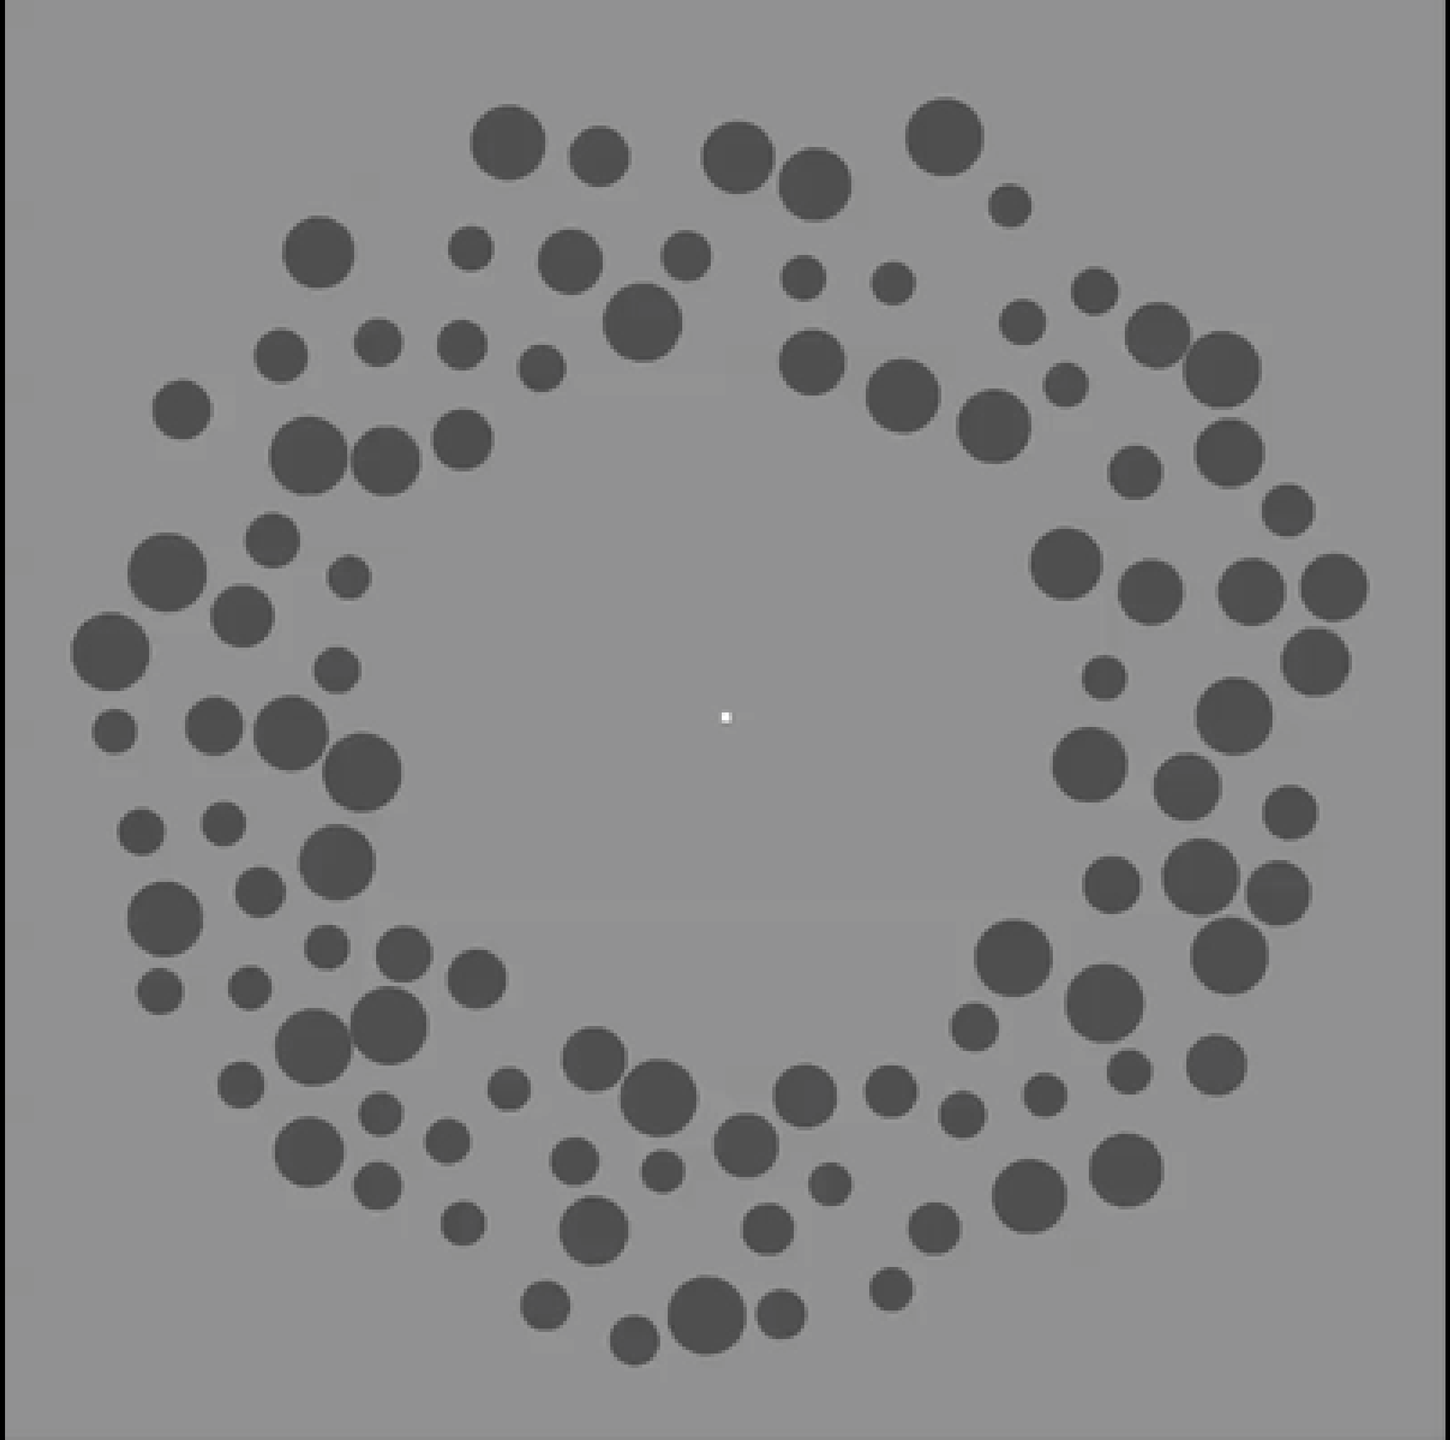

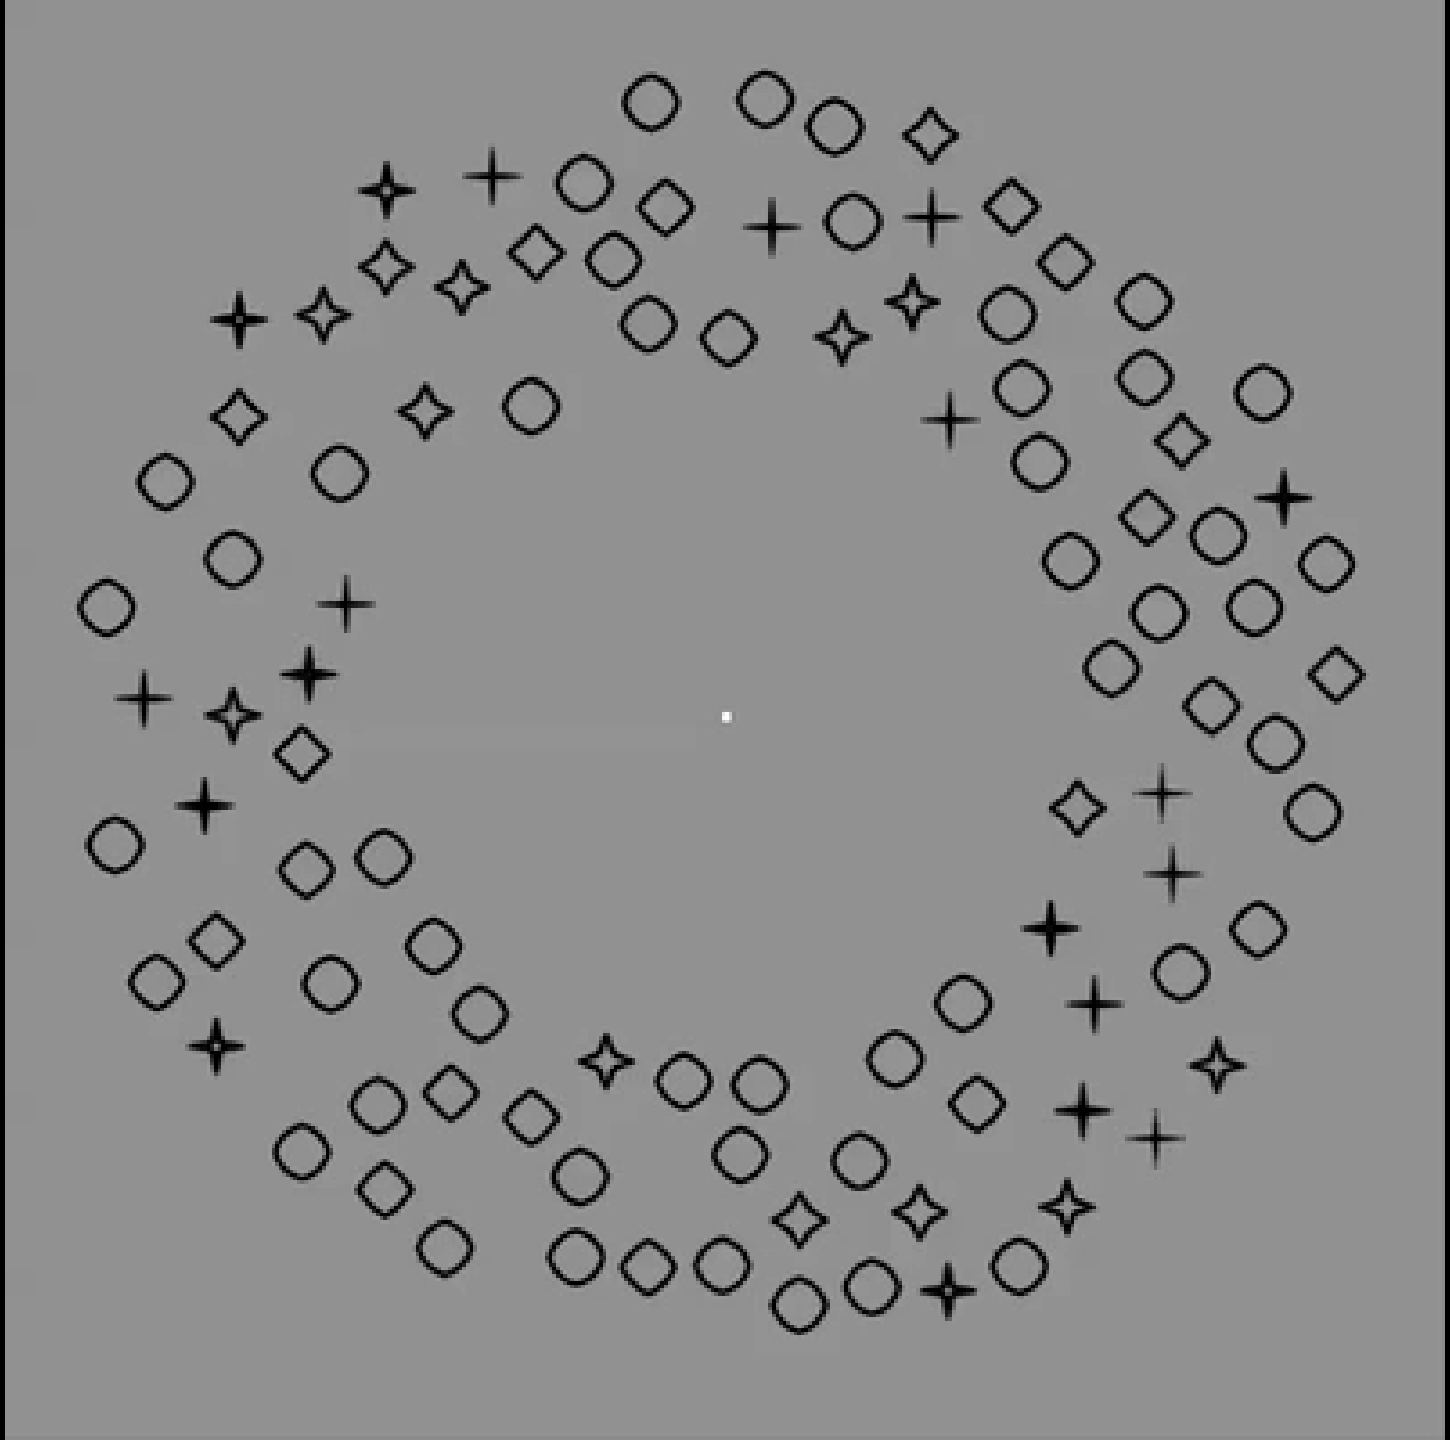

Current Biology 21, 140–143, January 25, 2011 DOI 10.1016/j.cub.2010.12.019

Report

Motion Silences Awareness of Visual Change

Jordan W. Suchow[1,*] and George A. Alvarez[1]
[1]Department of Psychology, Harvard University, Cambridge, MA 02138, USA

Summary

Loud bangs, bright flashes, and intense shocks capture attention, but other changes—even those of similar magnitude—can go unnoticed. Demonstrations of change blindness have shown that observers fail to detect substantial alterations to a scene when distracted by an irrelevant flash, or when the alterations happen gradually [1–5]. Here, we show that objects changing in hue, luminance, size, or shape appear to stop changing when they move. This motion-induced failure to detect change, silencing, persists even though the observer attends to the objects, knows that they are changing, and can make veridical judgments about their current state. Silencing demonstrates the tight coupling of motion and object appearance.

Results

We created a series of movies in which 100 dots were arranged in a ring around a central fixation mark (Figure 1A). Each movie alternated between two phases, stationary and moving. During the stationary phase, the dots changed rapidly in hue, luminance, size, or shape. During the moving phase, the dots continued to change at the same rate while the entire ring rotated about its center. Observers were instructed to adjust the rate of change during the stationary phase to match the apparent rate of change in the moving phase. The results revealed a graded effect: the faster the ring rotated, the slower the dots seemed to change (Figure 1B). The fastest rotation (0.33 Hz) produced nearly complete silencing. Several visual demonstrations can be found at http://visionlab.harvard.edu/silencing/ and in the Supplemental Information available online (Movie S1, Movie S2, Movie S3, and Movie S4).

Determining the Perceived State

During silencing, rapidly changing objects appear nearly static, which raises an immediate question: What is the perceived state (e.g., red, bright, big, round) at any given moment? To illustrate, consider an observer who fails to notice an object change gradually from yellow to red. One possibility is that the observer always sees yellow, never updating his percept to incorporate the new hue—this is freezing, erroneously keeping hold of an outdated state [6]. Another possibility is that he always sees the current hue (e.g., yellow, orange, then red) but is unaware of the transition from one to the next—this is implicit updating [4].

Both accounts are plausible. Temporal freezing, filling-in, and illusory color-shape conjunction are three known phenomena in which the visual system paints a percept that differs from reality, either by retaining an outdated version of a changing stimulus or by inferring its current or future state [6–8]. Alternatively, in continuous change-blindness, part of a scene changes gradually, and though oblivious to the change, the observer perceives its current state veridically [3, 4].

To distinguish these two accounts of silencing—freezing and implicit updating—we created a change-detection task that generalizes Hollingworth and Henderson's reversion test [4]. In that study, observers viewed a picture of a room while, unbeknownst to them, the camera angle gradually shifted. After some time, the camera angle suddenly reverted to its original state. Observers pressed a button if they saw the picture change. The two accounts make different predictions as to whether the observers noticed the reversion: implicit updating predicts success, whereas freezing predicts failure. In fact, the reversion was obvious, ruling against freezing and in favor of implicit updating [4]. Here, instead of performing a single test in which the dots flip to their original state (i.e., their hue at the onset of motion), we performed a separate test for each state in the dots' history—past, present, and future. This generalized reversion test affords greater sensitivity in determining the perceived state. The two accounts both predict that observers will notice some reversions while failing to notice others but differ as to which reversions they predict will go unnoticed (Figure 2; red segments in "predictions" panel at top).

We found that observers noticed flips to the past and future, but not to the present (Figure 2; bottom panel); this occurred regardless of whether the objects stopped, continued to move, or were masked at the time of the reversion. The average magnitude of an unnoticed flip was $-14° \pm 12°$ (mean ± standard error of the mean [SEM]) when the objects stopped moving, $-8° \pm 10°$ when they continued, and $-14° \pm 11°$ when they were masked. Though each of these values is slightly negative, none are significantly different from 0° (one-sample test for mean angle of circular data, $p = 0.23$, $p = 0.43$, and $p = 0.20$, respectively), and all are reliably different from 180° ($p < 0.001$ for each). Importantly, each distribution is markedly nonuniform, which implies that observers were able to make a judgment that depended on the objects' state (Rayleigh test for uniformity of circular data, $p < 0.001$ for each). Silenced changes are updated implicitly—the observer sees the current state.

Incidentally, freezing of stationary color changes has been found to last for ≈200 ms [6], which corresponds to a $-10°$ change in hue in our reversion test. Though the data rule out the possibility that temporal freezing explains silencing, they leave open the possibility that freezing persists within a local window, such that the perceived color consistently lags a bit behind the actual color; this would explain the observed, though not statistically significant, lag.

Motion in Space versus on the Retina

When an object moves but the observer's gaze does not—as in the movies presented here—two types of motion occur simultaneously: the object moves in space, and its image moves on the retina. Which causes silencing? We created four variants of the original movie that together dissociate the two types of motion. In the first variant, the object moves while the observer's gaze remains fixed, producing motion both in space and on the retina. In the second, the object moves

*Correspondence: suchow@fas.harvard.edu

Illusions of number

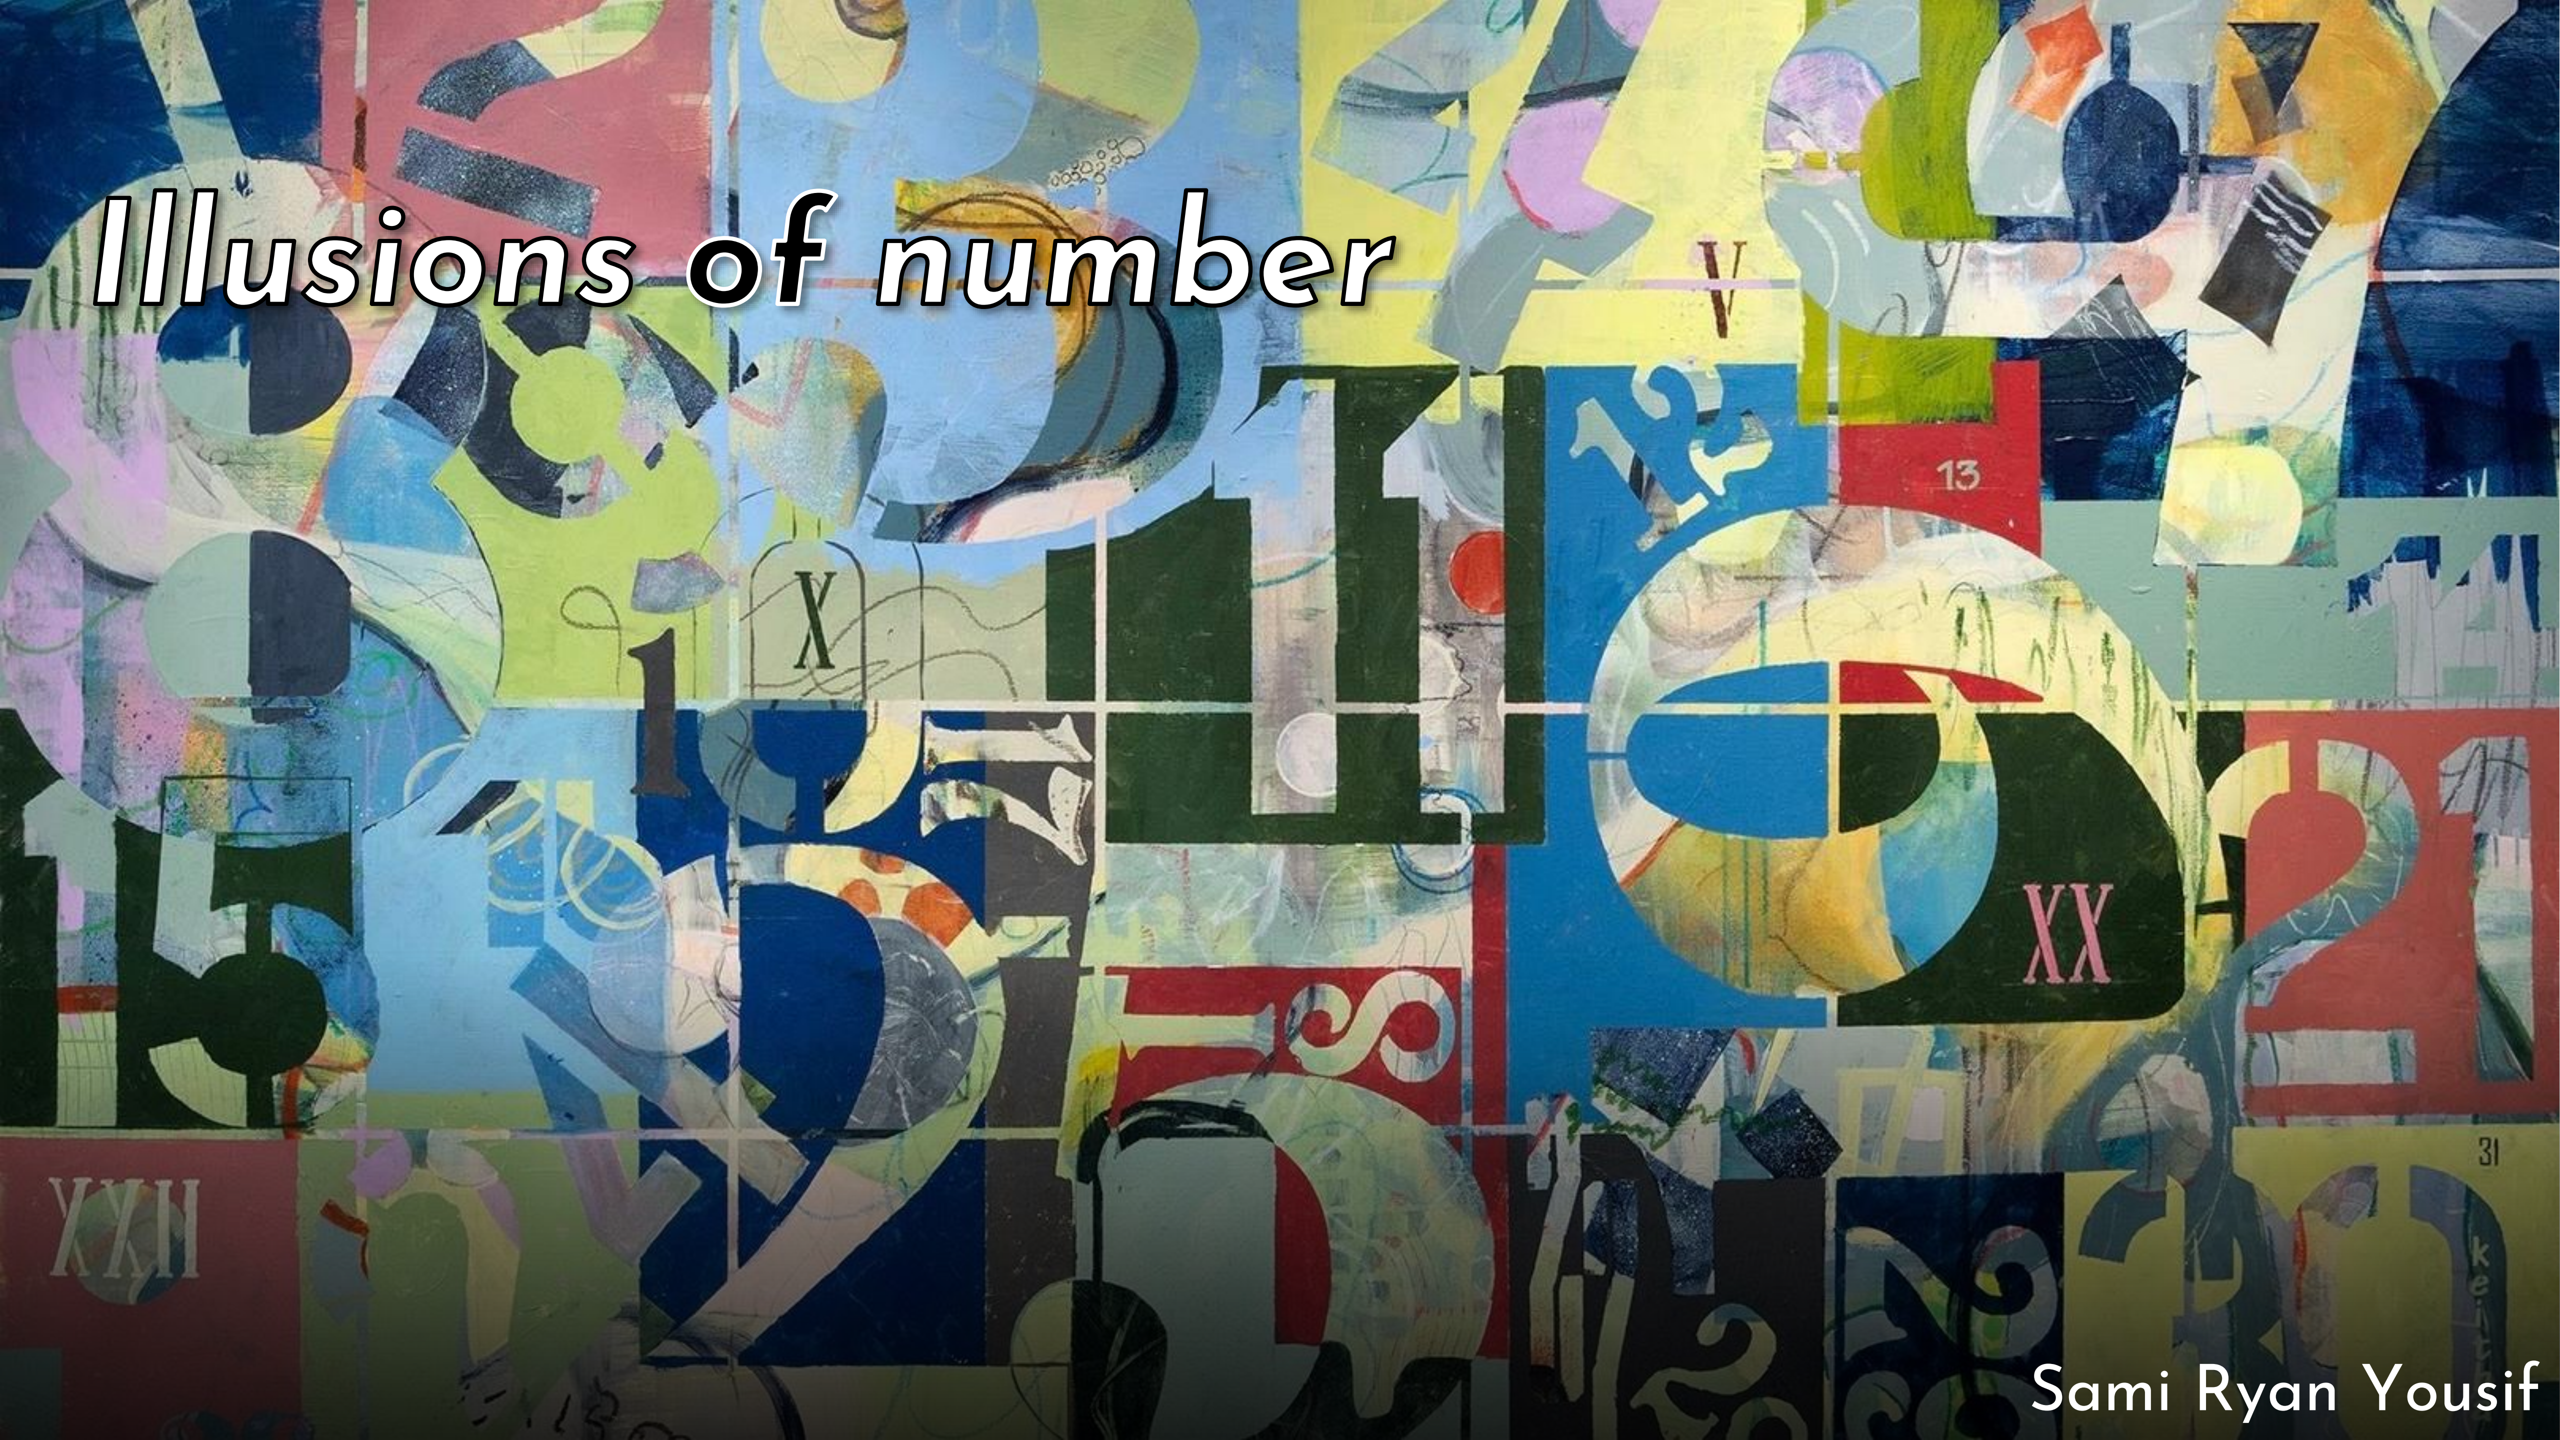

Sami Ryan Yousif

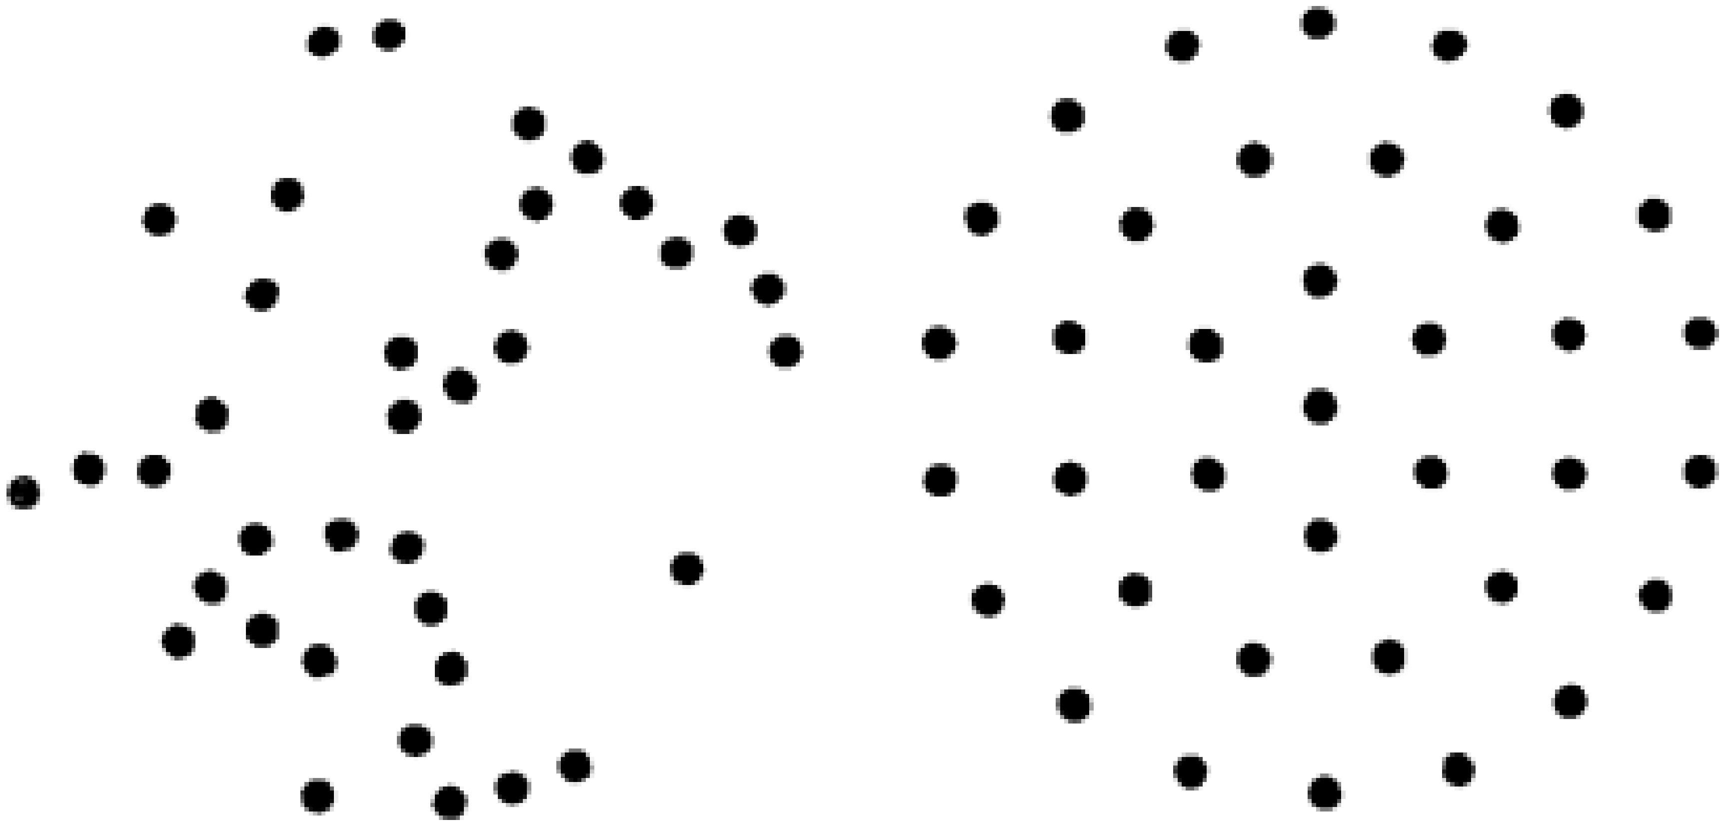

The **regular-random illusion**

!!

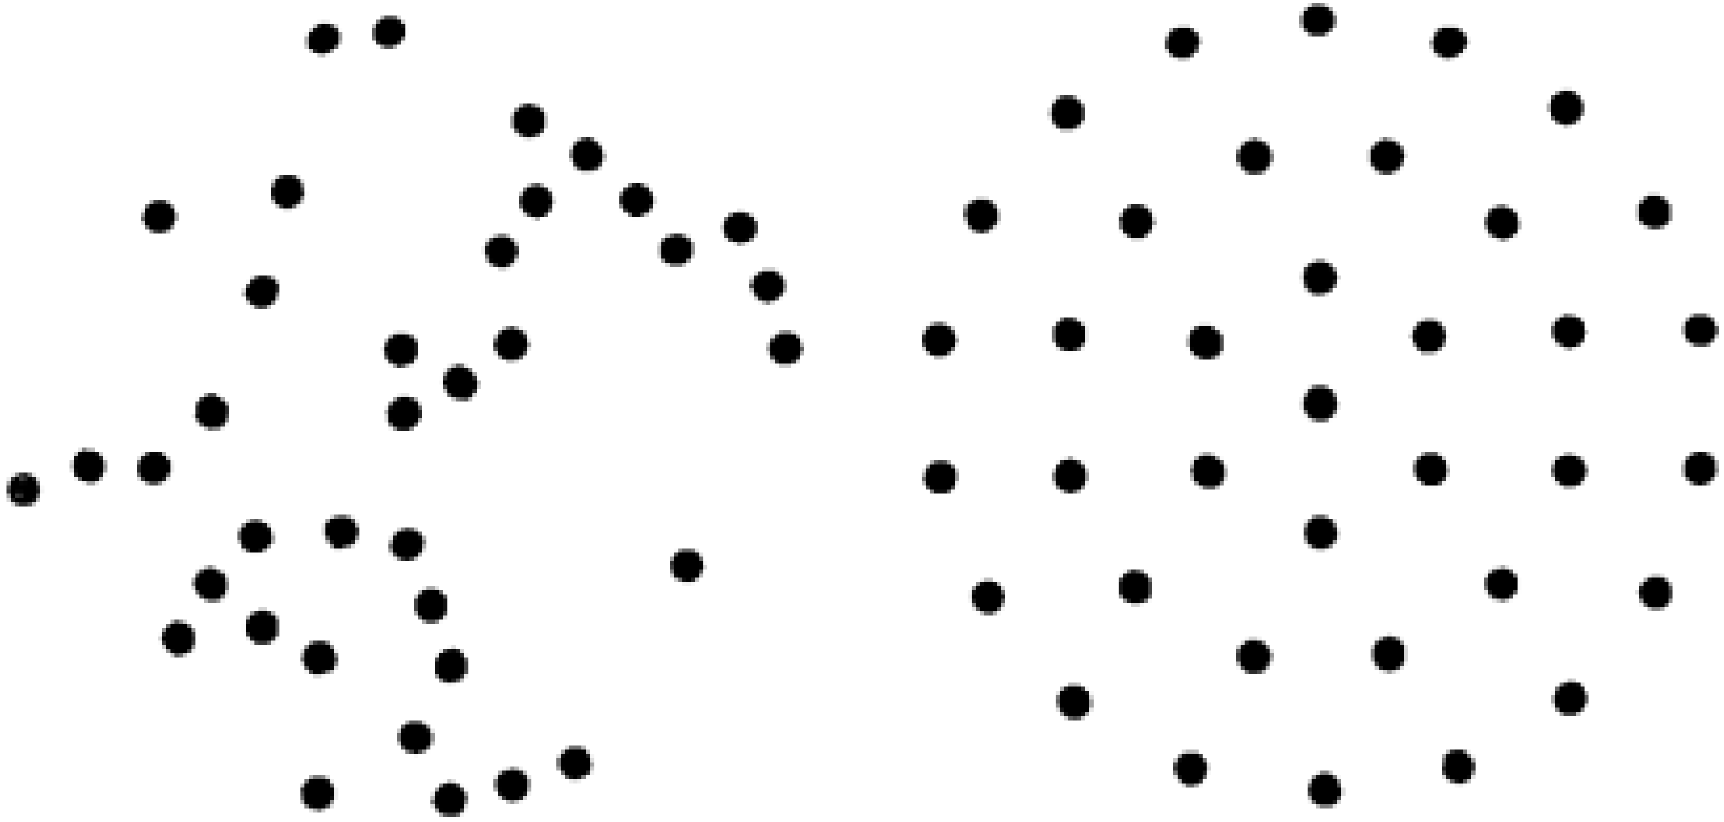

Ginsburg, 1976

Perception & Psychophysics
1991, 49 (4), 303-314

Occupancy model of perceived numerosity

JÜRI ALLIK and TIIA TUULMETS
University of Tartu, Tartu, Estonia

Observers saw 234 different pairs of stochastically organized dot patterns and indicated which of the two patterns appeared to be more numerous. All of the data can be accounted for by supposing that the choice of the more numerous pattern is based on the determination of the occupancy indices of both patterns. Each dot is posited to have an impact upon its neighborhood in a constant occupancy radius R. The area of the stimulus plane occupied collectively by all dots provides a basis for judging relative numerosity; the pattern with the larger occupancy value is chosen as more numerous. The occupancy model, besides providing a general explanation of known numerosity illusions in strictly quantitative terms, can explain some puzzling aspects of numerosity perception.

Quantification is one of the most impressive acts of the human mind. On many occasions, however, the direct one-by-one counting of items is impossible: the number of objects is too large, the viewing time is too limited, the separation of already-counted objects from not-yet-counted ones is too difficult, and so forth. Nevertheless, in all such situations, the observer is able to estimate the approximate number of items on the basis of an instantaneous impression of numerosity. In a typical numerosity discrimination experiment, the observer indicates which of the two presented random-dot patterns appears to be more numerous. It is intuitively compelling to think that the observer's decision is based on an internal representation of numerosity—that the pattern producing the greater subjective magnitude of an internal process is chosen as being more numerous. All currently constructed psychophysical scales describe perceived numerosity as a power function of the objective number of items in the stimulus (Indow & Ida, 1977; Krueger, 1972, 1984). All these numerosity scales are very tentative, however, for they fail to take into account the well-documented dependence of perceived numerosity on the *spatial configuration* of dots. The same numbers of dots distributed differently in space may appear to be very different in the apparent number of their elements. For example, objects occupying a more extended area on the display usually appear to contain more numerous elements (Bevan, Maier, & Helson, 1963; Binet, 1890; Ponzo, 1928). Many other configurations of dots have been found to increase or decrease in their apparent visual number, relative to the same number of randomly distributed dots (Frith & Frith, 1972; Ginsburg, 1976; Ginsburg & Goldstein, 1987; Krueger, 1972; Taves, 1941; Vos, van Oeffelen, Tibosch, & Allik, 1988). Such results indicate that the perceptual system is not able to abstract the number per se from all the other stimulus attributes (see Allik, 1989).

Gestalt categories of perceptual organization, such as spatial proximity, provide more realistic candidates for the stimulus properties that serve as bases for numerosity judgments. Unfortunately, most of these principles have not been formulated in quantitatively measurable terms, and, as a rule, they can be communicated only through graphic examples. One of the few attempts to explain perceived numerosity in strictly formal terms was undertaken by Vos et al. (1988). The basic idea is that the perceived numerosity depends not on the number of dots as such but on a more complex spatial property of the dot pattern—namely, the total area of the plane apparently filled with dots. The impact each dot has upon its neighborhood is portrayed as a monotonically decreasing spread (dispersion) function. The regions in the image where the sum of all (potentially overlapping) individual spread functions exceeds a pre-established threshold value are regarded as being filled with dots.

A formal description of the CODE (COntour DEtector) model, including the selection of optimal parameters (form and width of the spread function, threshold value), was provided by van Oeffelen and Vos in 1983. According to the CODE model, the width of the spread function depends on the distance to its nearest neighbor. The filled area index predicted the sign or direction of many known numerosity illusions (Vos et al. 1988). In addition, five specially constructed 36-dot patterns filling approximately 33% of the total area were rated as being more numerous than five 36-dot patterns filling only 15% of the total area. However, the model has not been tested in a more demanding manner. It is also worth noting that a general concern of the Vos et al. study, as with other numerosity studies in general, was to predict the *sign* of the numerosity illusion, not its *magnitude*.

Although the basic idea of the model—that the perceived numerosity can be identified with the filled area—seems

The authors thank Norman Ginsburg, Tarow Indow, and an anonymous reviewer for helpful comments on an earlier version of this article. The discussion of this research topic with Piet Vos has been very stimulating for many years. We are grateful to Enn Veldi for help in revising our English. We are extremely obliged to Lester Krueger for his valuable suggestions and editorial help. Correspondence concerning this article may be sent to Jüri Allik, Department of Psychology, University of Tartu, 78 Tiigi Street, Tartu, Estonia 202400 (U.S.S.R.).

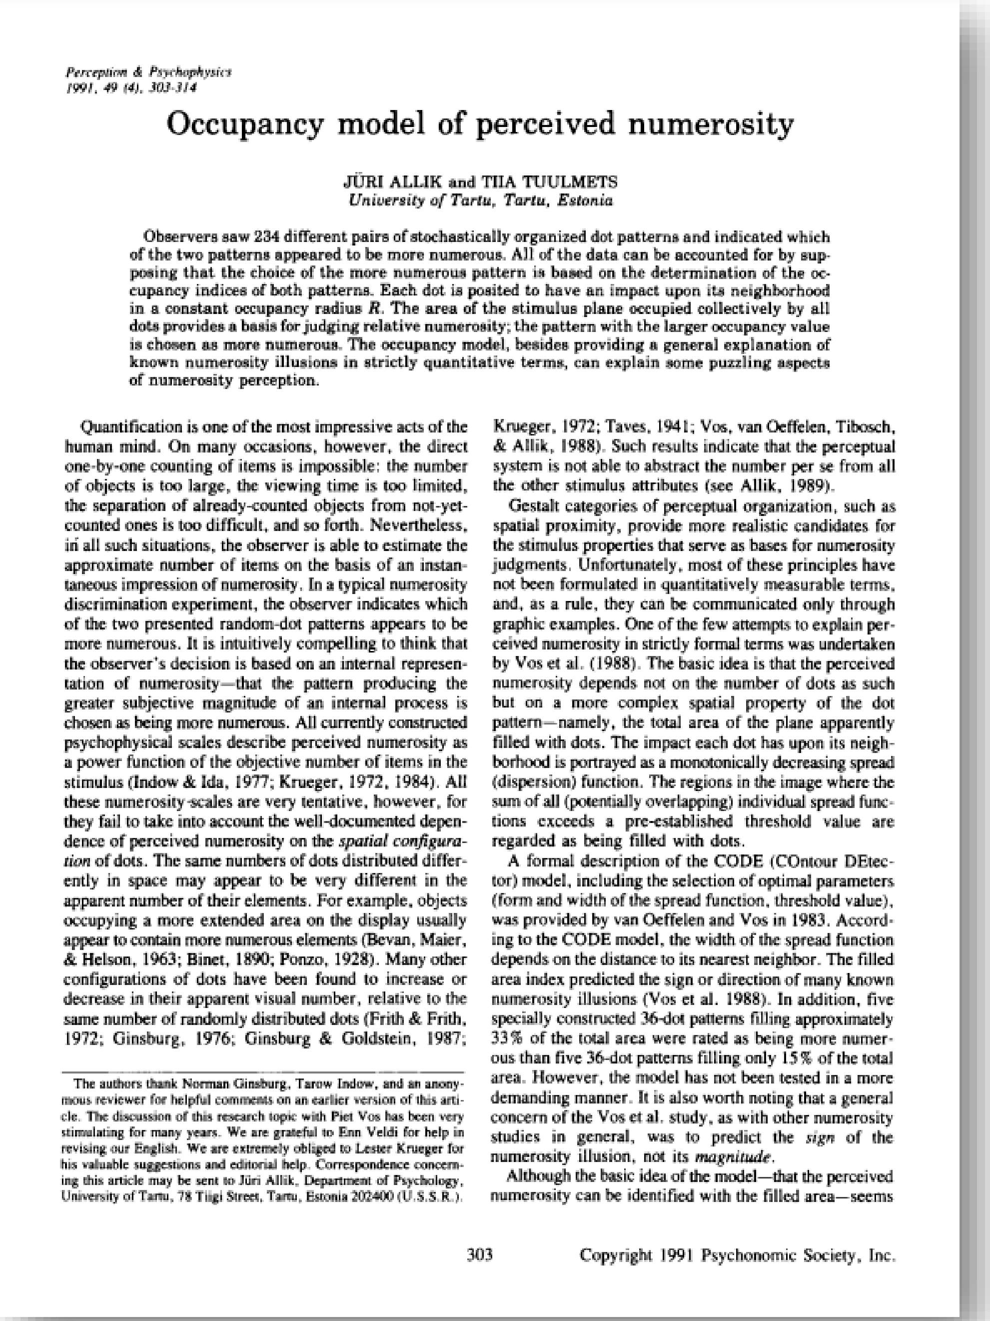
Perception & Psychophysics
1991, 49 (4), 303-314

Occupancy model of perceived numerosity

JÜRI ALLIK and TIIA TUULMETS
University of Tartu, Tartu, Estonia

Observers saw 234 different pairs of stochastically organized dot patterns and indicated which of the two patterns appeared to be more numerous. All of the data can be accounted for by supposing that the choice of the more numerous pattern is based on the determination of the occupancy indices of both patterns. Each dot is posited to have an impact upon its neighborhood in a constant occupancy radius *R*. The area of the stimulus plane occupied collectively by all dots provides a basis for judging relative numerosity; the pattern with the larger occupancy value is chosen as more numerous. The occupancy model, besides providing a general explanation of known numerosity illusions in strictly quantitative terms, can explain some puzzling aspects of numerosity perception.

Quantification is one of the most impressive acts of the human mind. On many occasions, however, the direct one-by-one counting of items is impossible: the number of objects is too large, the viewing time is too limited, the separation of already-counted objects from not-yet-counted ones is too difficult, and so forth. Nevertheless, in all such situations, the observer is able to estimate the approximate number of items on the basis of an instantaneous impression of numerosity. In a typical numerosity discrimination experiment, the observer indicates which of the two presented random-dot patterns appears to be more numerous. It is intuitively compelling to think that the observer's decision is based on an internal representation of numerosity—that the pattern producing the greater subjective magnitude of an internal process is chosen as being more numerous. All currently constructed psychophysical scales describe perceived numerosity as a power function of the objective number of items in the stimulus (Indow & Ida, 1977; Krueger, 1972, 1984). All these numerosity scales are very tentative, however, for they fail to take into account the well-documented dependence of perceived numerosity on the *spatial configuration* of dots. The same numbers of dots distributed differently in space may appear to be very different in the apparent number of their elements. For example, objects occupying a more extended area on the display usually appear to contain more numerous elements (Bevan, Maier, & Helson, 1963; Binet, 1890; Ponzo, 1928). Many other configurations of dots have been found to increase or decrease in their apparent visual number, relative to the same number of randomly distributed dots (Frith & Frith, 1972; Ginsburg, 1976; Ginsburg & Goldstein, 1987; Krueger, 1972; Taves, 1941; Vos, van Oeffelen, Tibosch, & Allik, 1988). Such results indicate that the perceptual system is not able to abstract the number per se from all the other stimulus attributes (see Allik, 1989).

Gestalt categories of perceptual organization, such as spatial proximity, provide more realistic candidates for the stimulus properties that serve as bases for numerosity judgments. Unfortunately, most of these principles have not been formulated in quantitatively measurable terms, and, as a rule, they can be communicated only through graphic examples. One of the few attempts to explain perceived numerosity in strictly formal terms was undertaken by Vos et al. (1988). The basic idea is that the perceived numerosity depends not on the number of dots as such but on a more complex spatial property of the dot pattern—namely, the total area of the plane apparently filled with dots. The impact each dot has upon its neighborhood is portrayed as a monotonically decreasing spread (dispersion) function. The regions in the image where the sum of all (potentially overlapping) individual spread functions exceeds a pre-established threshold value are regarded as being filled with dots.

A formal description of the CODE (COntour DEtector) model, including the selection of optimal parameters (form and width of the spread function, threshold value), was provided by van Oeffelen and Vos in 1983. According to the CODE model, the width of the spread function depends on the distance to its nearest neighbor. The filled area index predicted the sign or direction of many known numerosity illusions (Vos et al. 1988). In addition, five specially constructed 36-dot patterns filling approximately 33% of the total area were rated as being more numerous than five 36-dot patterns filling only 15% of the total area. However, the model has not been tested in a more demanding manner. It is also worth noting that a general concern of the Vos et al. study, as with other numerosity studies in general, was to predict the *sign* of the numerosity illusion, not its *magnitude*.

Although the basic idea of the model—that the perceived numerosity can be identified with the filled area—seems

The authors thank Norman Ginsburg, Tarow Indow, and an anonymous reviewer for helpful comments on an earlier version of this article. The discussion of this research topic with Piet Vos has been very stimulating for many years. We are grateful to Enn Veldi for help in revising our English. We are extremely obliged to Lester Krueger for his valuable suggestions and editorial help. Correspondence concerning this article may be sent to Jüri Allik, Department of Psychology, University of Tartu, 78 Tiigi Street, Tartu, Estonia 202400 (U.S.S.R.).

!!

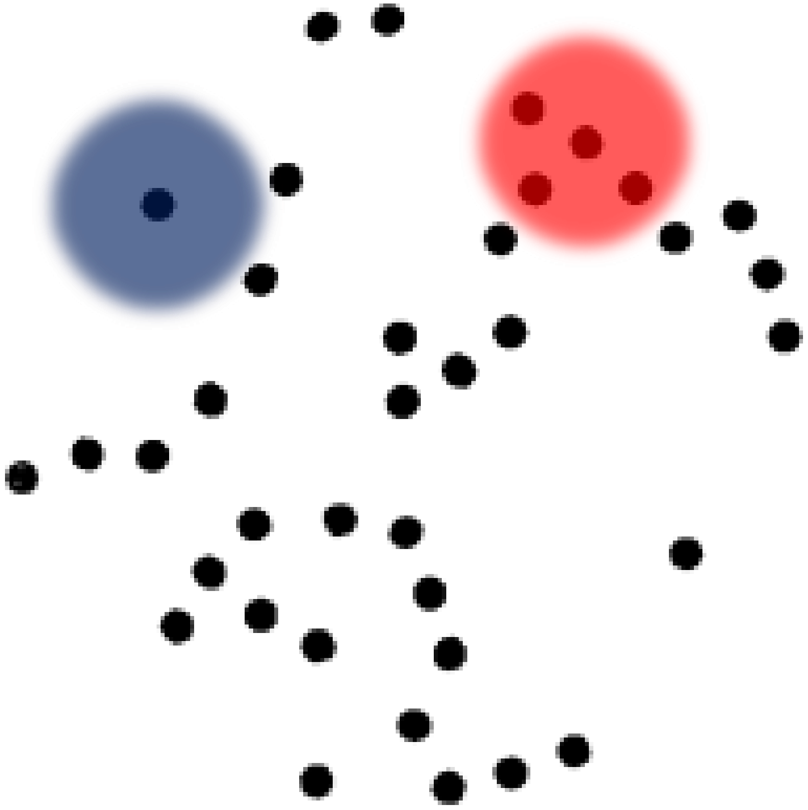

“receptive field?”

The coherence illusion !!

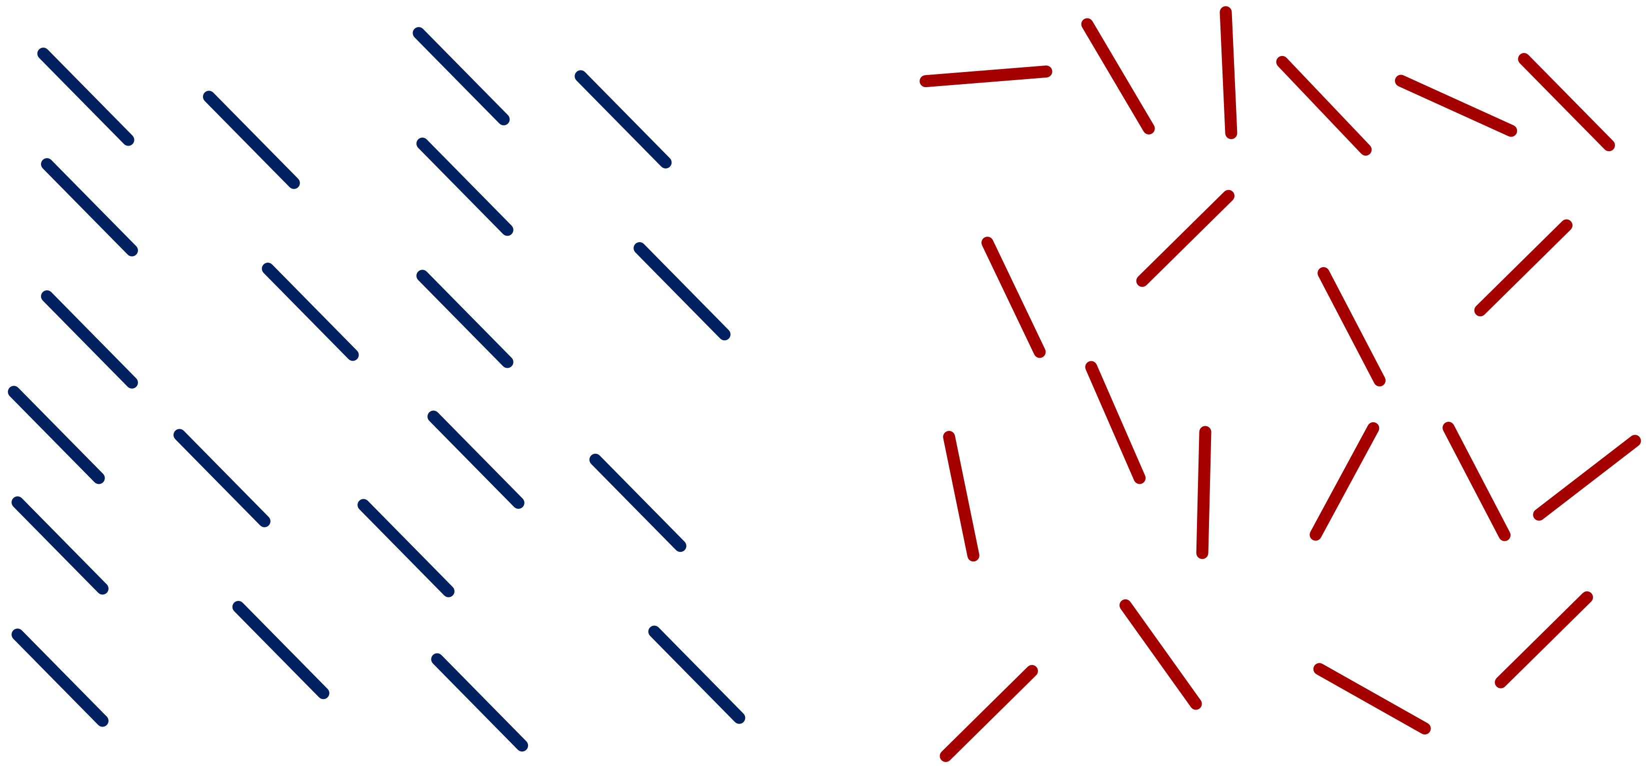

DeWind et al., 2020

Similarly oriented objects appear more numerous

Nicholas K. DeWind Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Michael F. Bonner Department of Cognitive Science, Johns Hopkins University, Baltimore, MD, USA

Elizabeth M. Brannon Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Several non-numerical factors influence the numerical estimation of visual arrays, including the spacing of items and whether they are arranged randomly or symmetrically. Here we report a novel numerosity illusion we term the *coherence illusion*. When items in an array have a coherent orientation (all pointing in the same direction) they seem to be more numerous than when items are oriented randomly. Participants show parametric effects of orientation coherence in three distinct numerical judgment tasks. These findings are not predicted by any current model of numerical estimation. We discuss array entropy as a possible framework for explaining both the coherence illusion and the previously reported regular-random illusion.

Introduction

Educated adults can precisely quantify a set by counting and using number words to denote cardinality. But people also directly perceive the approximate number of objects in a set (e.g., Dehaene, 1997). Educated adults share this number sense with primates (e.g., Brannon & Terrace, 1998), rodents (e.g., Meck & Church, 1983), birds (e.g., Honig & Stewart, 1989), and a wide variety of other animals (for review see Merritt, DeWind, & Brannon, 2012). The approximate number sense emerges early in human development (e.g., Izard, Sann, Spelke, & Streri, 2009) and is present in adults from societies that lack a verbal counting system (Pica, Lemer, Izard, & Dehaene, 2004). The approximate number sense is theorized to provide a foundation for symbolic mathematics (Feigenson, Dehaene, & Spelke, 2004), and the precision of approximate numerical discrimination has been found to be correlated with mathematical achievement in children (for review see Chen & Li, 2014).

Several non-numerical attributes of arrays have been noted to affect their perceived numerosity. For example, connecting elements in an array with a thin line decreases their perceived numerosity (Franconeri, Bemis, & Alvarez, 2009; He, Zhang, Zhou, & Chen, 2009). These authors propose that the perceived numerical decrease is due to object segmentation. Items that are part of a contiguous portion of space may be partially viewed as a single object. Two objects connected by a line are contiguously connected, and so an array composed of connected pairs is viewed as less numerous than the same array without the connections. Interestingly, this effect does not require an explicit spatial connection. Illusory contours connecting objects as well as statistical regularity in the color of neighboring objects is sufficient to create a more abstract "connection" between the objects and reduce perceived numerosity (Kirjakovski & Matsumoto, 2016; Zhao & Yu, 2016). Connectedness affects numerosity representations in extrastriate cortical areas and operates on signals associated with perceived numerosity rather than lower level representations in primary visual cortex. For example, connectedness affects numerosity adaptation and is detectable in neural responses to stimulus arrays after 150 ms in visual area V3, but not in lower cortical areas (Fornaciai, Cicchini, & Burr, 2016; Fornaciai & Park, 2018). Connectedness and grouping of objects also affects numerosity signals in the intraparietal areas (He, Zhou, Zhou, He, & Chen, 2015).

Another effect of array configuration on perceived numerosity is spacing; greater average spacing between the elements of a set increases their perceived number (DeWind, Adams, Platt, & Brannon, 2015; Gebuis & Reynvoet, 2012a), and regularly spaced elements seem to be more numerous than randomly spaced elements (Ginsburg, 1976). These effects may be two examples of the same phenomenon: that greater distance between

Citation: DeWind, N. K., Bonner, M. F., & Brannon, E. M. (2020). Similarly oriented objects appear more numerous. *Journal of Vision*, *20*(4):4, 1–11, https://doi.org/10.1167/jov.20.4.4.

Similarly oriented objects appear more numerous

Nicholas K. DeWind
Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Several non-numerical factors influence the numerical estimation of visual arrays, including the spacing of items and whether they are arranged randomly or symmetrically. Here we report a novel numerosity illusion we term the *coherence illusion*. When items in an array have a coherent orientation (all pointing in the same direction) they seem to be more numerous than when items are oriented randomly. Participants show parametric effects of orientation coherence in three distinct numerical judgment tasks. These findings are not predicted by any current model of numerical estimation. We discuss array entropy as a possible framework for explaining both the coherence illusion and the previously reported regular-random illusion.

number of objects in a set (e.g., Dehaene, 1997). Educated adults share this number sense with primates (e.g., Brannon & Terrace, 1998), rodents (e.g., Meck & Church, 1983), birds (e.g., Honig & Stewart, 1989), and a wide variety of other animals (for review see Merritt, DeWind, & Brannon, 2012). The approximate number sense emerges early in human development (e.g., Izard, Sann, Spelke, & Streri, 2009) and is present in adults from societies that lack a verbal counting system (Pica, Lemer, Izard, & Dehaene, 2004). The approximate number sense is theorized to provide a foundation for symbolic mathematics (Feigenson, Dehaene, & Spelke, 2004), and the precision of approximate numerical discrimination has been found to be correlated with mathematical achievement in children (for review see Chen & Li, 2014).

primary visual cortex. For example, connectedness affects numerosity adaptation and is detectable in neural responses to stimulus arrays after 150 ms in visual area V3, but not in lower cortical areas (Fornaciai, Cicchini, & Burr, 2016; Fornaciai & Park, 2018). Connectedness and grouping of objects also affects numerosity signals in the intraparietal areas (He, Zhou, Zhou, He, & Chen, 2015).

Another effect of array configuration on perceived numerosity is spacing; greater average spacing between the elements of a set increases their perceived number (DeWind, Adams, Platt, & Brannon, 2015; Gebuis & Reynvoet, 2012a), and regularly spaced elements seem to be more numerous than randomly spaced elements (Ginsburg, 1976). These effects may be two examples of the same phenomenon: that greater distance between

Citation: DeWind, N. K., Bonner, M. F., & Brannon, E. M. (2020). Similarly oriented objects appear more numerous. *Journal of Vision*, *20*(4):4, 1–11, https://doi.org/10.1167/jov.20.4.4.

How is number perceived?

Number is perceived **directly**

It is processed **early in cortex**

Acts as a **perceptual primitive**

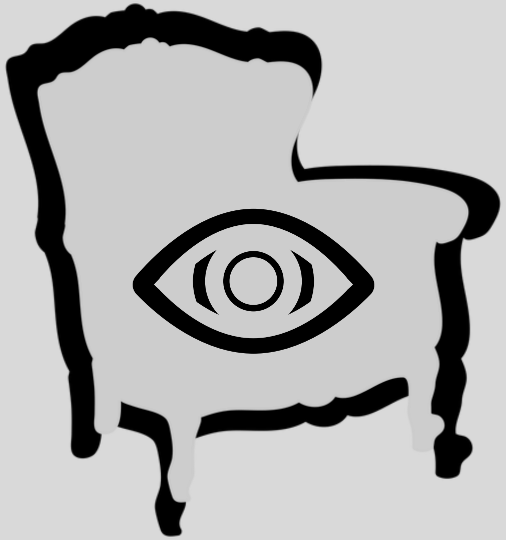

'Domain-specific view'

'Domain-general view'

Number is perceived **indirectly**

It is **inferred or constructed**

Not a primitive, but an **output**

American Psychological Association

Journal of Experimental Psychology: General

0096-3445/19/$12.00

http://dx.doi.org/10.1037/xge0000553

Visual Illusions Help Reveal the Primitives of Number Perception

Edwina Picon, Denitza Dramkin, and Darko Odic
University of British Columbia

The human perceptual system is responsive to numerical information within visual and auditory scenes. For example, when shown 2 displays of dots, observers can instantly, albeit approximately, identify the set that is more numerous. Theories in perceptual and cognitive psychology have focused on 2 mechanisms for how vision accomplishes such a feat: Under the domain-specific encoding theory, number is represented as a primary visual feature of perception, much like motion or color, while under the domain-general theory, the visual system represents number indirectly, through a complex combination of features such as the size of the dots, their total cluster, and so forth. Evidence for the latter theory often comes from "congruency effects:" the finding that participants frequently select the side where the dots on the screen are denser, larger, or brighter, rather than the side that is actually more numerous. However, such effects could also stem from response conflicts between otherwise independent dimensions. Here, we test these 2 competing accounts by embedding numerical displays within visual illusions that create large conflicts between number and other non-numeric dimensions—including contour length, convex hull, and density—and contrast participants' performance on a number discrimination task (i.e., "Which side has more dots?") against a number estimation task (i.e., "How many dots are there?"), which should eliminate response conflicts. Across 3 experiments, we find that while contour length illusions only affect number perception in discrimination tasks, the influences of convex hull and density on number perception persist in both discrimination *and* estimation tasks, supporting a more domain-general account of number encoding.

Keywords: approximate number system, number perception, number sense, visual illusions

Supplemental materials: http://dx.doi.org/10.1037/xge0000553.supp

Perception gets a lot from very little. Even just a cursory glance at Figure 1 yields an instant and automatic sense of number: Without counting you can easily decide whether there are more blue dots than yellow dots. Decades of work have shown that the visual and auditory systems of human newborns and nonhuman animals are sensitive to changes in number (Izard, Sann, Spelke, & Streri, 2009; for review of the nonhuman animal literature, see Vallortigara, 2017), and that this "number sense" contributes to an assortment of other cognitive abilities, including our understanding of currency (Marques & Dehaene, 2004) and foraging behavior in nonhuman primates (Piantadosi & Cantlon, 2017). As a result, the perceptual sense of number has been of great interest to cognitive, computational, developmental, and comparative psychologists.

Our sensitivity to visual number information simultaneously showcases the efficiency and the mystery of perception: Because number cannot be extracted from any single objective feature, such as wavelengths of light (for color), salt concentration (for taste), pressure on skin (for touch), and so forth, how do our perceptual systems encode and represent number? To date, two types of theories have been put forward to answer this question. Under the first, the *domain-specific encoding theory*, number is represented relatively early in sensory processing by dedicated and specialized neurons, thus constituting a primitive of perception, much like color, orientation, and motion. For example, in the model of Dehaene and Changeux (1993), low-level neurons instantiate a two-dimensional object map whose total activity corresponds to the number of objects in the scene: the more objects, the more populated the map, and the higher the representation of number (see also Stoianov & Zorzi, 2012). Consistent with this domain-specific encoding account, Burr and Ross (2008; Ross & Burr, 2010, 2012) have repeatedly demonstrated that we can perceptually adapt to number. In the same way that staring at a green square subsequently produces an illusory percept of a red square, staring at a display of 100 dots subsequently produces an illusory diminished sense of number compared with staring at a display of 50 dots (see Burr & Ross, 2008 for demonstration). Because adaptation effects are most often caused by the fatiguing of low-level neurons due to repeated exposure, Burr and Ross (2008; Ross & Burr, 2010, 2012) have concluded that number must therefore be a low-level feature akin to motion, color, orientation, and so forth (for other arguments in favor of number as a primary visual feature, see Anobile, Cicchini, & Burr, 2016).

Edwina Picon, Denitza Dramkin, and Darko Odic, Department of Psychology, University of British Columbia.

This research was supported by an NSERC DG provided to Darko Odic. We thank Catalina Canon for help in data collection and Justin Halberda for various conversations about the data and the article.

Correspondence concerning this article should be addressed to Darko Odic, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, British Columbia V6T 1Z4. E-mail: darko.odic@psych.ubc.ca

American Psychological Association
Journal of Experimental Psychology: General

0096-3445/19/$12.00
http://dx.doi.org/10.1037/xge0000553

Visual Illusions Help Reveal the Primitives of Number Perception

Edwina Picon, Denitza Dramkin, and Darko Odic
University of British Columbia

The human perceptual system is responsive to numerical information within visual and auditory scenes. For example, when shown 2 displays of dots, observers can instantly, albeit approximately, identify the set that is more numerous. Theories in perceptual and cognitive psychology have focused on 2 mechanisms for how vision accomplishes such a feat: Under the domain-specific encoding theory, number is represented as a primary visual feature of perception, much like motion or color, while under the domain-general theory, the visual system represents number indirectly, through a complex combination of features such as the size of the dots, their total cluster, and so forth. Evidence for the latter theory often comes from "congruency effects:" the finding that participants frequently select the side where the dots on the screen are denser, larger, or brighter, rather than the side that is actually more numerous. However, such effects could also stem from response conflicts between otherwise independent dimensions. Here, we test these 2 competing accounts by embedding numerical displays within visual illusions that create large conflicts between number and other non-numeric dimensions—including contour length, convex hull, and density—and contrast participants' performance on a number discrimination task (i.e., "Which side has more dots?") against a number estimation task (i.e., "How many dots are there?"), which should eliminate response conflicts. Across 3 experiments, we find that while contour length illusions only affect number perception in discrimination tasks, the influences of convex hull and density on number perception persist in both discrimination *and* estimation tasks, supporting a more domain-general account of number encoding.

Keywords: approximate number system, number perception, number sense, visual illusions

Supplemental materials: http://dx.doi.org/10.1037/xge0000553.supp

Perception gets a lot from very little. Even just a cursory glance at Figure 1 yields an instant and automatic sense of number: Without counting you can easily decide whether there are more blue dots than yellow dots. Decades of work have shown that the visual and auditory systems of human newborns and nonhuman animals are sensitive to changes in number (Izard, Sann, Spelke, & Streri, 2009; for review of the nonhuman animal literature, see Vallortigara, 2017), and that this "number sense" contributes to an assortment of other cognitive abilities, including our understanding of currency (Marques & Dehaene, 2004) and foraging behavior in nonhuman primates (Piantadosi & Cantlon, 2017). As a result, the perceptual sense of number has been of great interest to cognitive, computational, developmental, and comparative psychologists.

Our sensitivity to visual number information simultaneously showcases the efficiency and the mystery of perception: Because number cannot be extracted from any single objective feature, such as wavelengths of light (for color), salt concentration (for taste), pressure on skin (for touch), and so forth, how do our perceptual systems encode and represent number? To date, two types of theories have been put forward to answer this question. Under the first, the *domain-specific encoding theory*, number is represented relatively early in sensory processing by dedicated and specialized neurons, thus constituting a primitive of perception, much like color, orientation, and motion. For example, in the model of Dehaene and Changeux (1993), low-level neurons instantiate a two-dimensional object map whose total activity corresponds to the number of objects in the scene: the more objects, the more populated the map, and the higher the representation of number (see also Stoianov & Zorzi, 2012). Consistent with this domain-specific encoding account, Burr and Ross (2008; Ross & Burr, 2010, 2012) have repeatedly demonstrated that we can perceptually adapt to number. In the same way that staring at a green square subsequently produces an illusory percept of a red square, staring at a display of 100 dots subsequently produces an illusory diminished sense of number compared with staring at a display of 50 dots (see Burr & Ross, 2008 for demonstration). Because adaptation effects are most often caused by the fatiguing of low-level neurons due to repeated exposure, Burr and Ross (2008; Ross & Burr, 2010, 2012) have concluded that number must therefore be a low-level feature akin to motion, color, orientation, and so forth (for other arguments in favor of number as a primary visual feature, see Anobile, Cicchini, & Burr, 2016).

Edwina Picon, Denitza Dramkin, and Darko Odic, Department of Psychology, University of British Columbia.

This research was supported by an NSERC DG provided to Darko Odic. We thank Catalina Canon for help in data collection and Justin Halberda for various conversations about the data and the article.

Correspondence concerning this article should be addressed to Darko Odic, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, British Columbia V6T 1Z4. E-mail: darko.odic@psych.ubc.ca

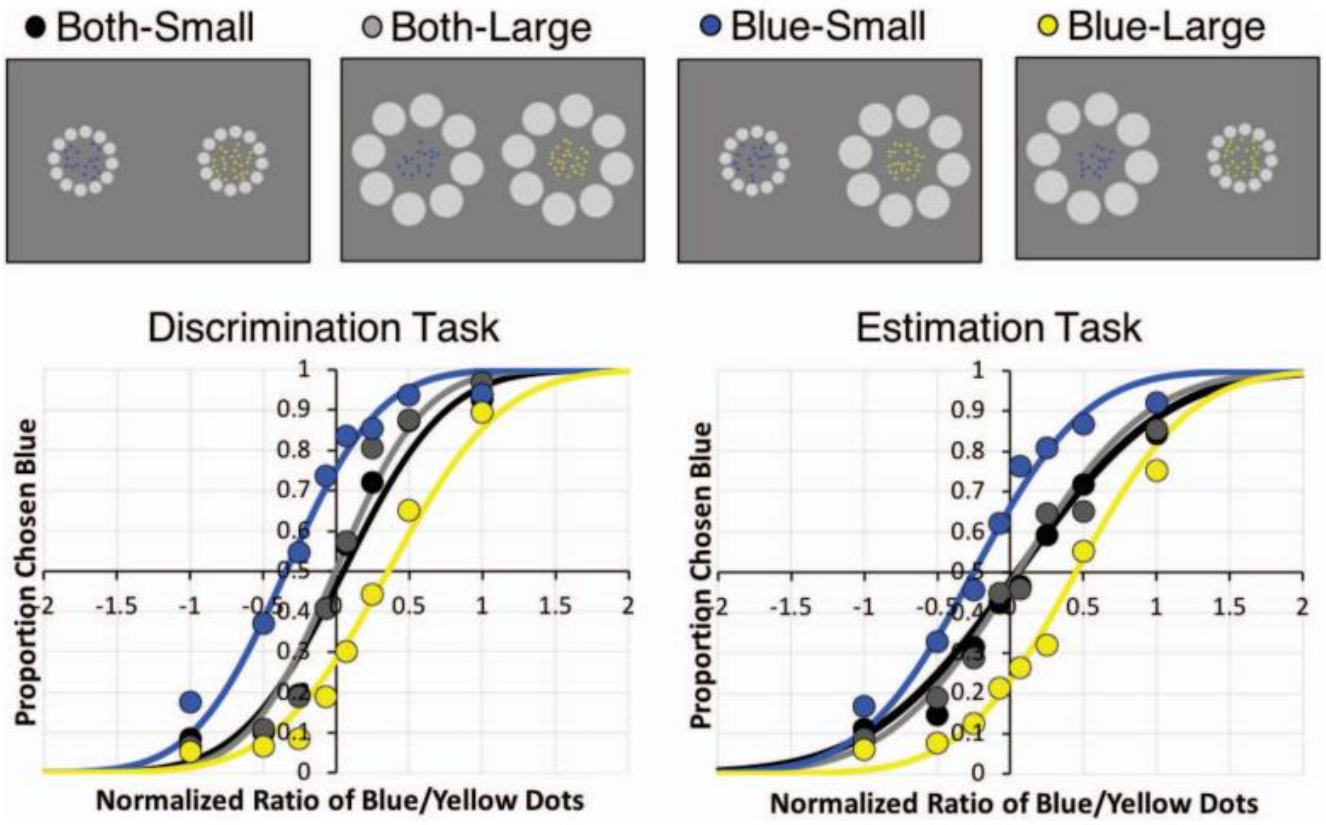

Figure 3. Experiment 2 stimuli and results. The top panel illustrates four example trials, one from each of the four conditions. The bottom panel shows the data from the Experiment 2 discrimination and estimation conditions; the lines are the best-fit psychophysical cumulative normal model. See the online article for the color version of this figure.

American Psychological Association

Journal of Experimental Psychology: General

0096-3445/19/$12.00

http://dx.doi.org/10.1037/xge0000553

Visual Illusions Help Reveal the Primitives of Number Perception

Edwina Picon, Denitza Dramkin, and Darko Odic
University of British Columbia

The human perceptual system is responsive to numerical information within visual and auditory scenes. For example, when shown 2 displays of dots, observers can instantly, albeit approximately, identify the set that is more numerous. Theories in perceptual and cognitive psychology have focused on 2 mechanisms for how vision accomplishes such a feat: Under the domain-specific encoding theory, number is represented as a primary visual feature of perception, much like motion or color, while under the domain-general theory, the visual system represents number indirectly, through a complex combination of features such as the size of the dots, their total cluster, and so forth. Evidence for the latter theory often comes from "congruency effects:" the finding that participants frequently select the side where the dots on the screen are denser, larger, or brighter, rather than the side that is actually more numerous. However, such effects could also stem from response conflicts between otherwise independent dimensions. Here, we test these 2 competing accounts by embedding numerical displays within visual illusions that create large conflicts between number and other non-numeric dimensions—including contour length, convex hull, and density—and contrast participants' performance on a number discrimination task (i.e., "Which side has more dots?") against a number estimation task (i.e., "How many dots are there?"), which should eliminate response conflicts. Across 3 experiments, we find that while contour length illusions only affect number perception in discrimination tasks, the influences of convex hull and density on number perception persist in both discrimination *and* estimation tasks, supporting a more domain-general account of number encoding.

Keywords: approximate number system, number perception, number sense, visual illusions

Supplemental materials: http://dx.doi.org/10.1037/xge0000553.supp

Perception gets a lot from very little. Even just a cursory glance at Figure 1 yields an instant and automatic sense of number: Without counting you can easily decide whether there are more blue dots than yellow dots. Decades of work have shown that the visual and auditory systems of human newborns and nonhuman animals are sensitive to changes in number (Izard, Sann, Spelke, & Streri, 2009; for review of the nonhuman animal literature, see Vallortigara, 2017), and that this "number sense" contributes to an assortment of other cognitive abilities, including our understanding of currency (Marques & Dehaene, 2004) and foraging behavior in nonhuman primates (Piantadosi & Cantlon, 2017). As a result, the perceptual sense of number has been of great interest to cognitive, computational, developmental, and comparative psychologists.

Our sensitivity to visual number information simultaneously showcases the efficiency and the mystery of perception: Because number cannot be extracted from any single objective feature, such as wavelengths of light (for color), salt concentration (for taste), pressure on skin (for touch), and so forth, how do our perceptual systems encode and represent number? To date, two types of theories have been put forward to answer this question. Under the first, the *domain-specific encoding theory*, number is represented relatively early in sensory processing by dedicated and specialized neurons, thus constituting a primitive of perception, much like color, orientation, and motion. For example, in the model of Dehaene and Changeux (1993), low-level neurons instantiate a two-dimensional object map whose total activity corresponds to the number of objects in the scene: the more objects, the more populated the map, and the higher the representation of number (see also Stoianov & Zorzi, 2012). Consistent with this domain-specific encoding account, Burr and Ross (2008; Ross & Burr, 2010, 2012) have repeatedly demonstrated that we can perceptually adapt to number. In the same way that staring at a green square subsequently produces an illusory percept of a red square, staring at a display of 100 dots subsequently produces an illusory diminished sense of number compared with staring at a display of 50 dots (see Burr & Ross, 2008 for demonstration). Because adaptation effects are most often caused by the fatiguing of low-level neurons due to repeated exposure, Burr and Ross (2008; Ross & Burr, 2010, 2012) have concluded that number must therefore be a low-level feature akin to motion, color, orientation, and so forth (for other arguments in favor of number as a primary visual feature, see Anobile, Cicchini, & Burr, 2016).

Edwina Picon, Denitza Dramkin, and Darko Odic, Department of Psychology, University of British Columbia.

This research was supported by an NSERC DG provided to Darko Odic. We thank Catalina Canon for help in data collection and Justin Halberda for various conversations about the data and the article.

Correspondence concerning this article should be addressed to Darko Odic, Department of Psychology, University of British Columbia, 2136 West Mall, Vancouver, British Columbia V6T 1Z4. E-mail: darko.odic@psych.ubc.ca

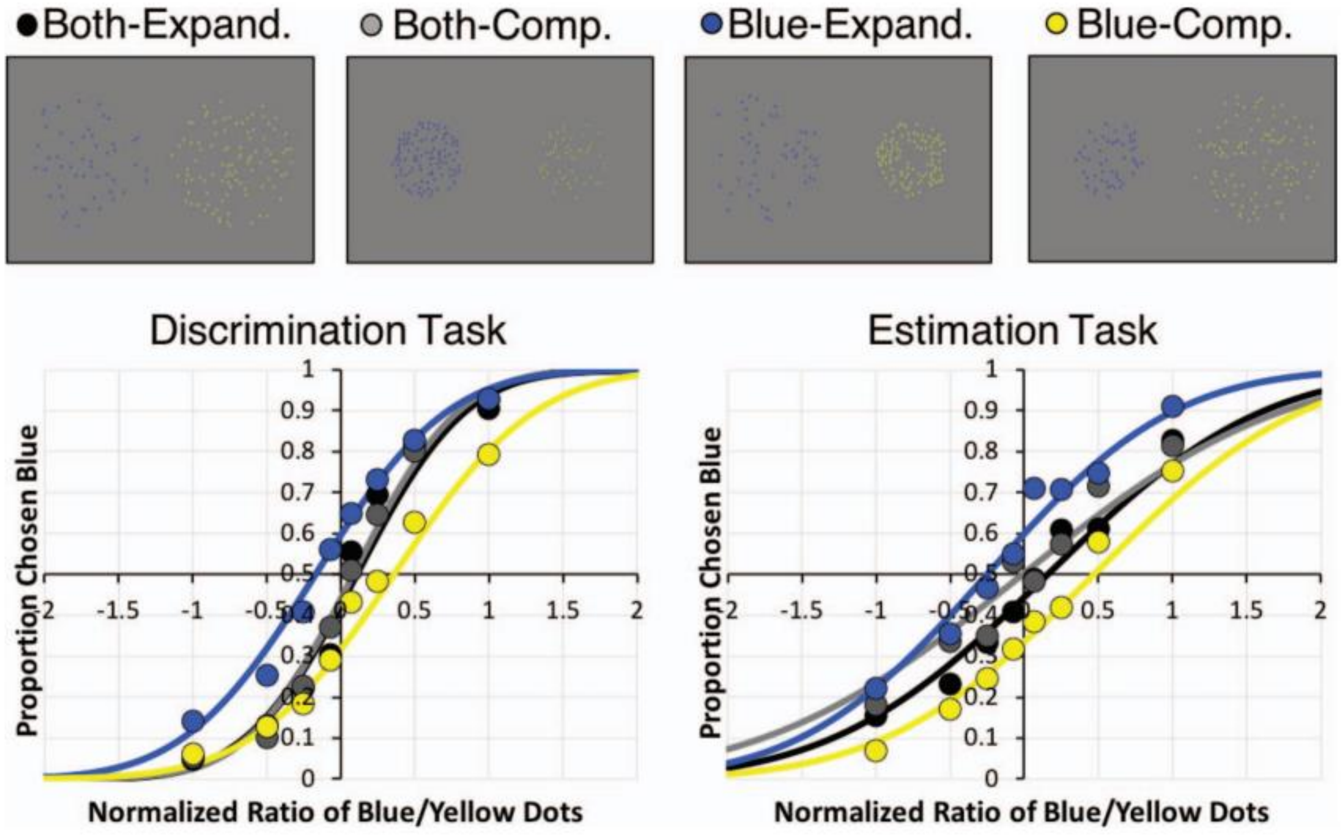

Figure 5. Experiment 3 stimuli and results. The top panel illustrates four example trials, one from each of the four conditions. The bottom panel shows the data from the Experiment 3 discrimination and estimation conditions; the lines are the best-fit psychophysical cumulative normal model. See the online article for the color version of this figure.

“Congruency effects”

!!

"General Magnitude Theory" / "A Theory of Magnitude" (ATOM)

Stroop effect

Red Green Purple
Brown Blue Red

Purple Red Brown
Red Green Blue

Numerical Stroop effect

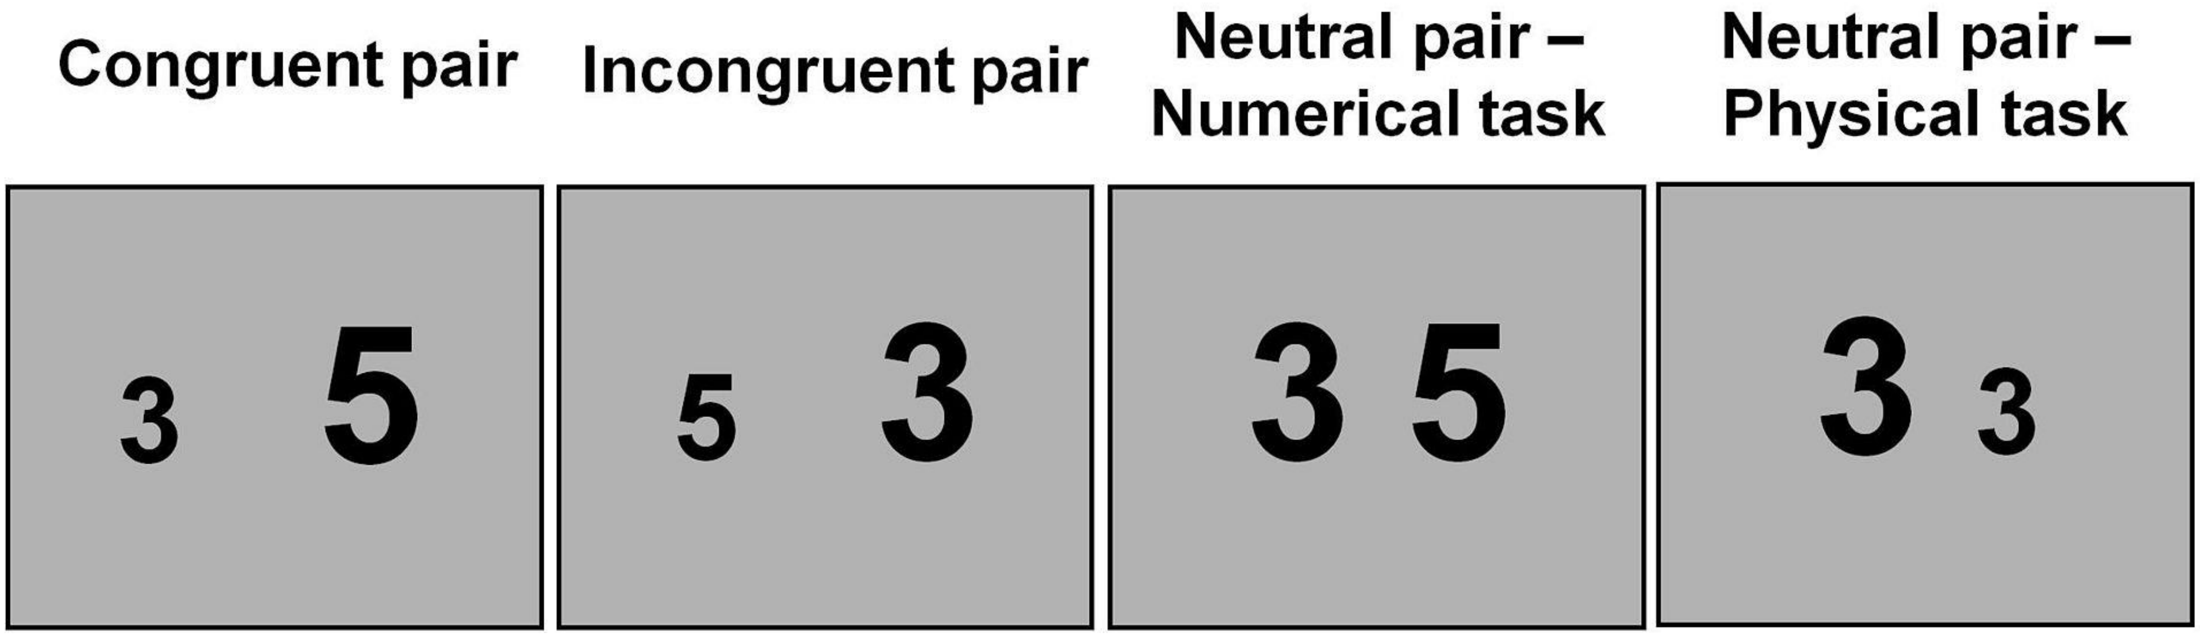

The crowd size illusion

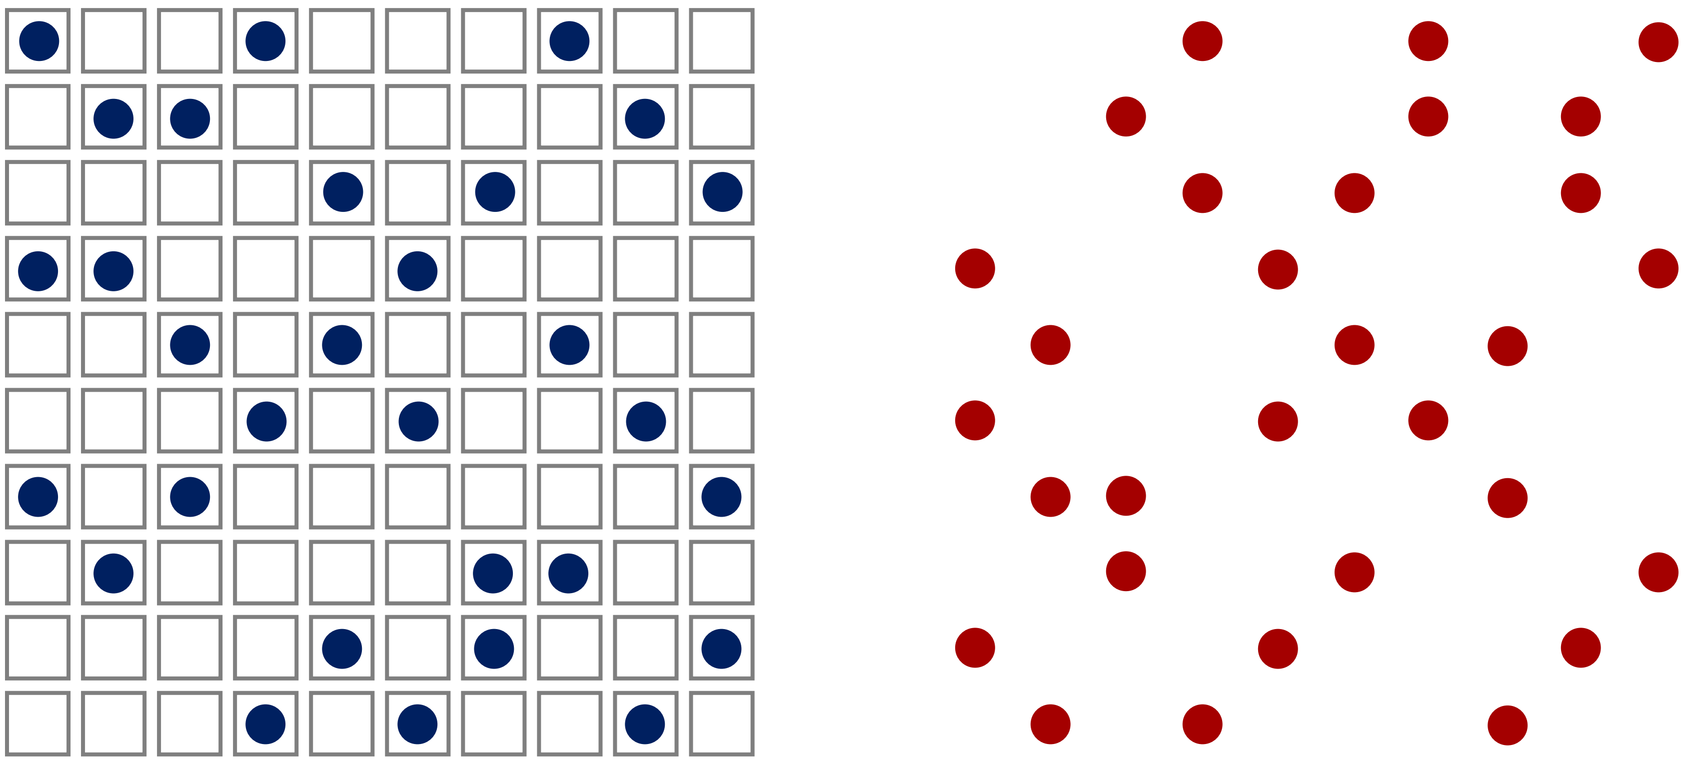

Waterhouse & Yousif, 2025

The crowd size illusion

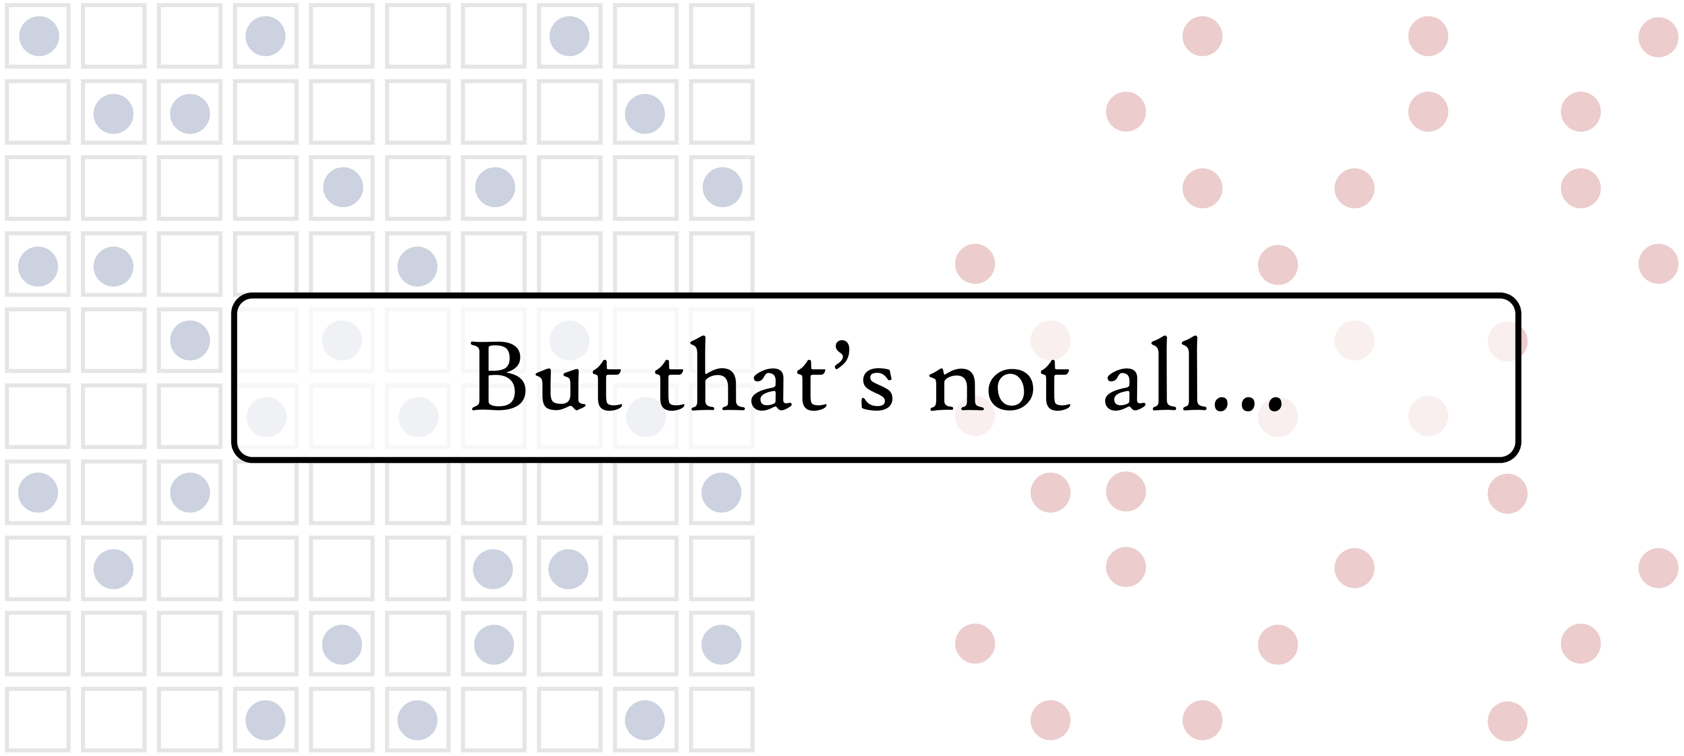

But that's not all...

Waterhouse & Yousif, 2025

The crowd size illusion

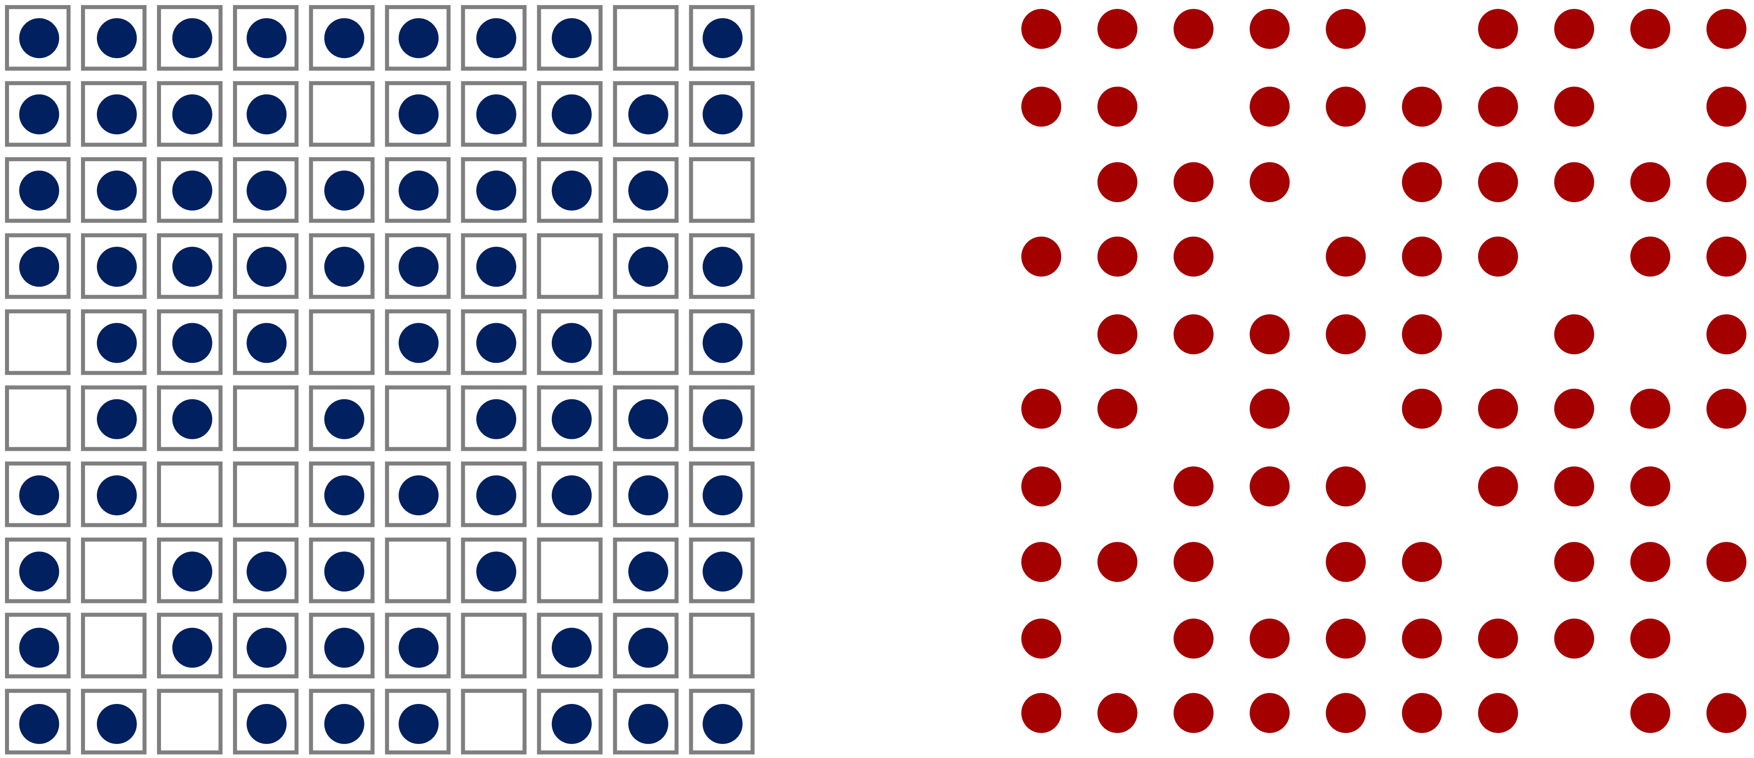

Waterhouse & Yousif, 2025

The connectedness illusion

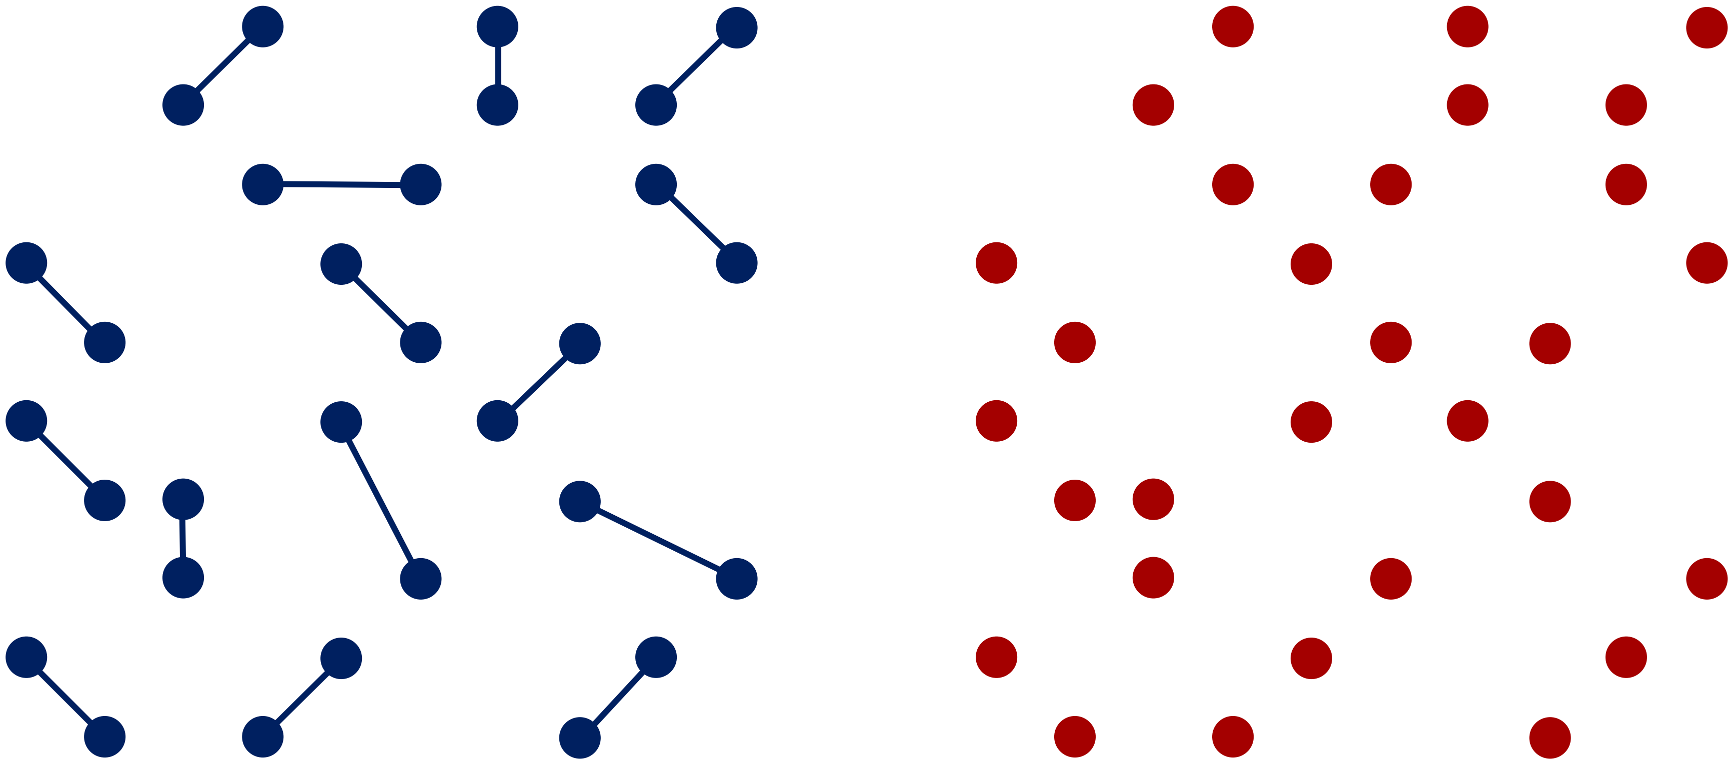

Franconeri et al., 2009;
He et al., 2009

‼

Let's talk about **objects.**

(and attention)

!!

object-based attention

object-based attention

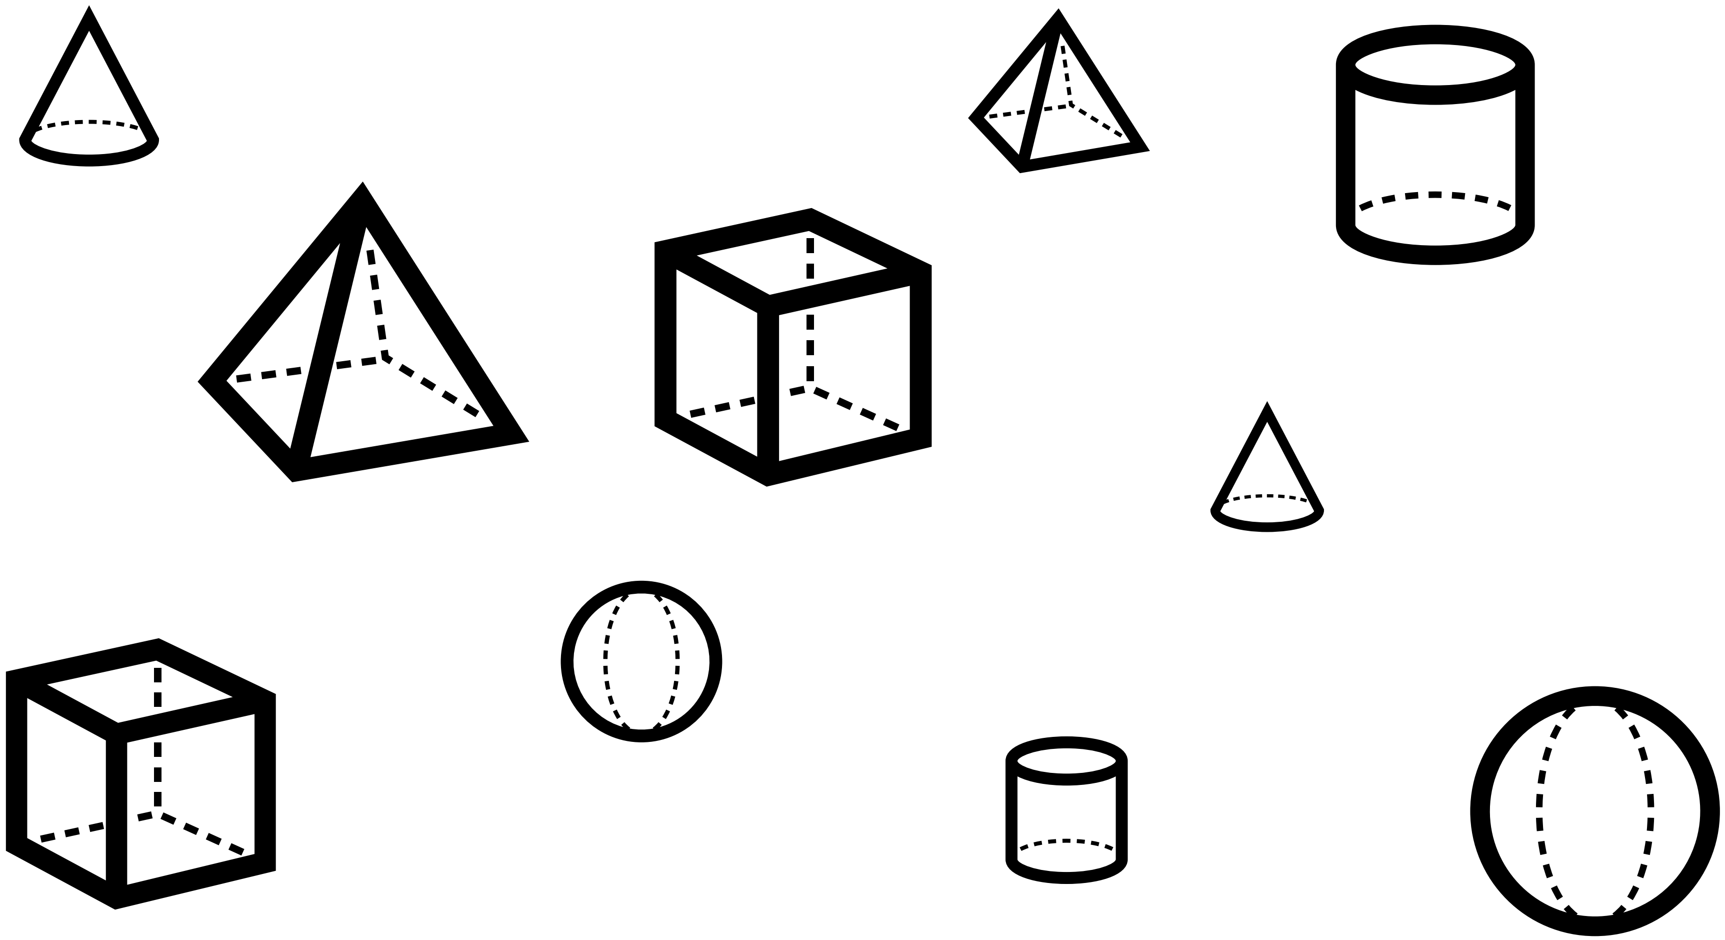

object-based attention

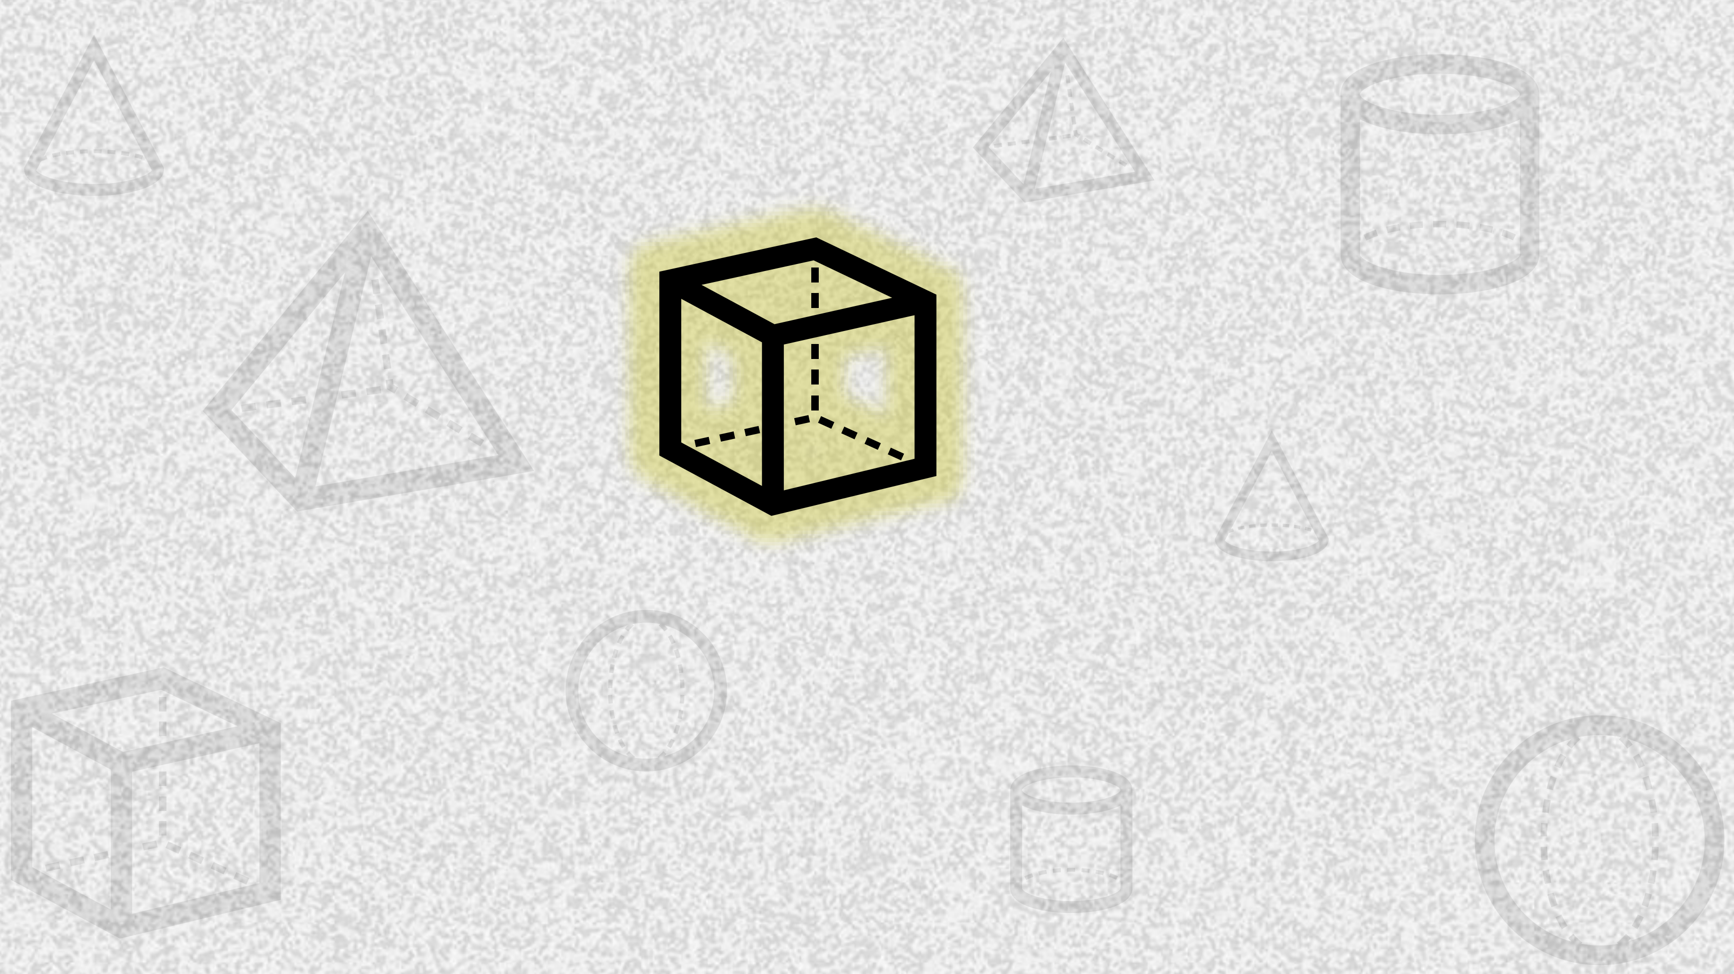

object-based attention

object-based attention

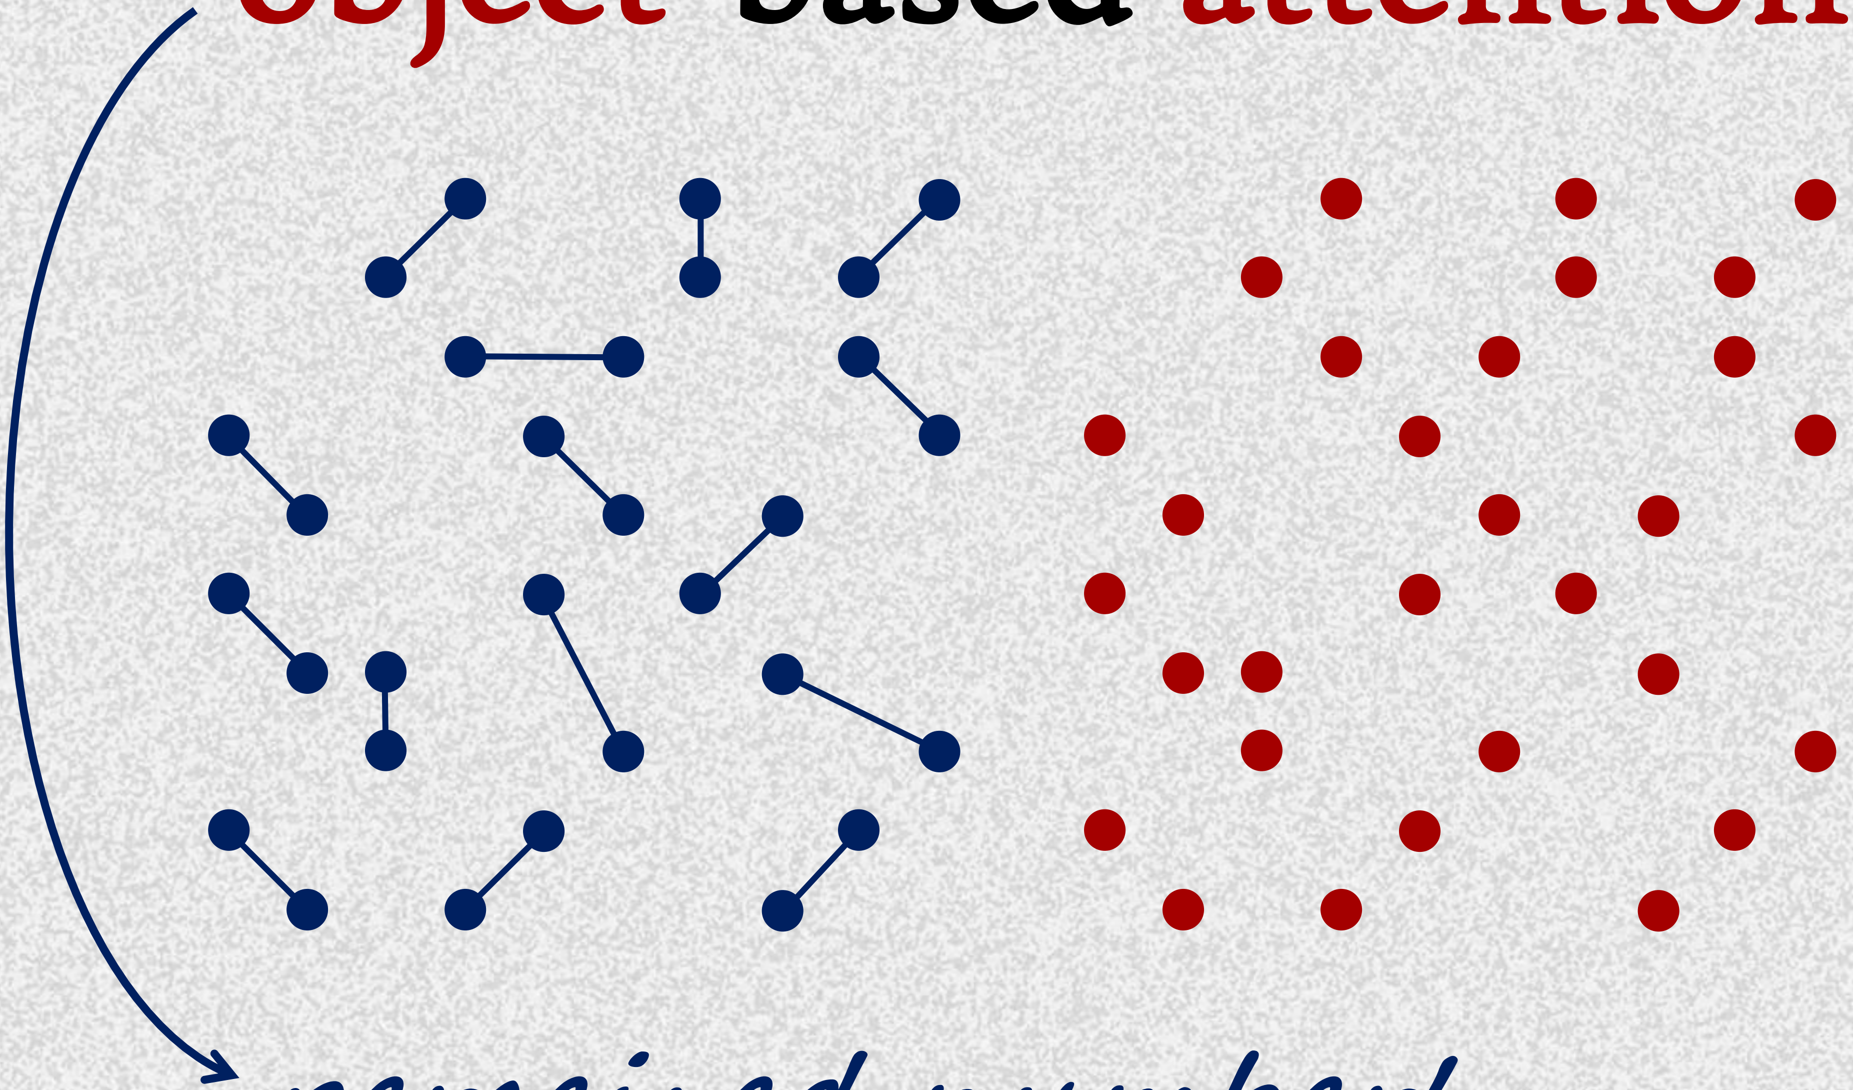

!!

perceived number!

It's not just about (overt) connections!

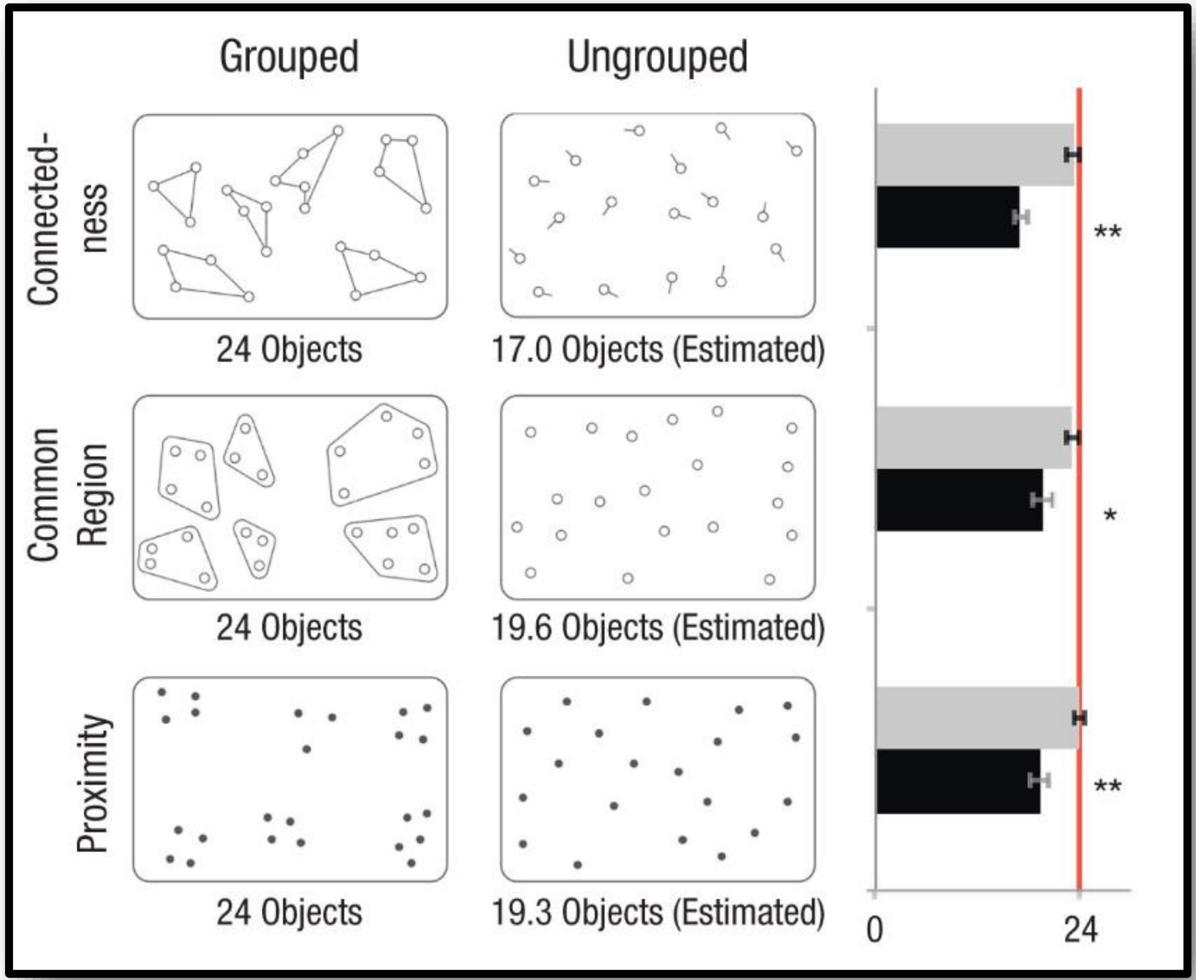

Yu et al., 2019

We've encountered object-based attention before!

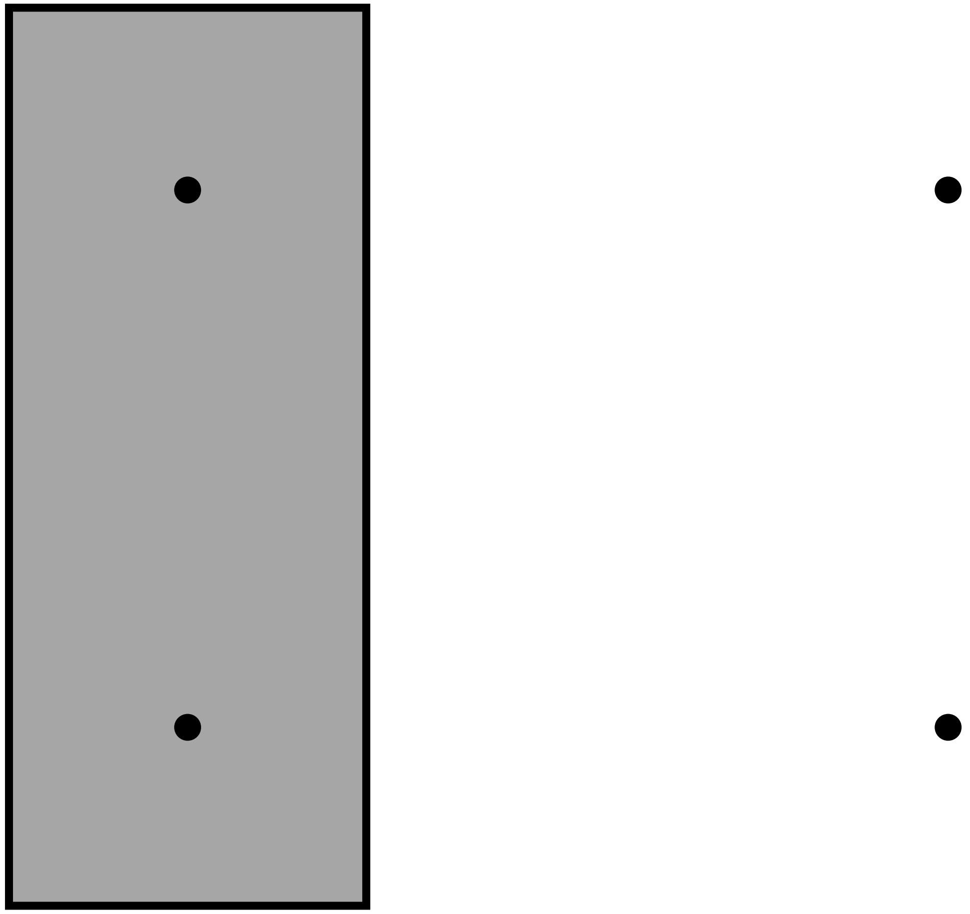

Object-based warping

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Cognition

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

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

ARTICLE INFO

Keywords:
Spatial perception
Time perception
Segmentation
Object-based attention

ABSTRACT

We distinguish between discrete objects and continuous entities in categorization and language, but might we actually *see* such stimuli differently? Here we report the *one-is-more illusion*, wherein 'objecthood' changes what we perceive in an unexpected way. Across many variations and tasks, observers perceived a single continuous object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

The one-is-more illusion

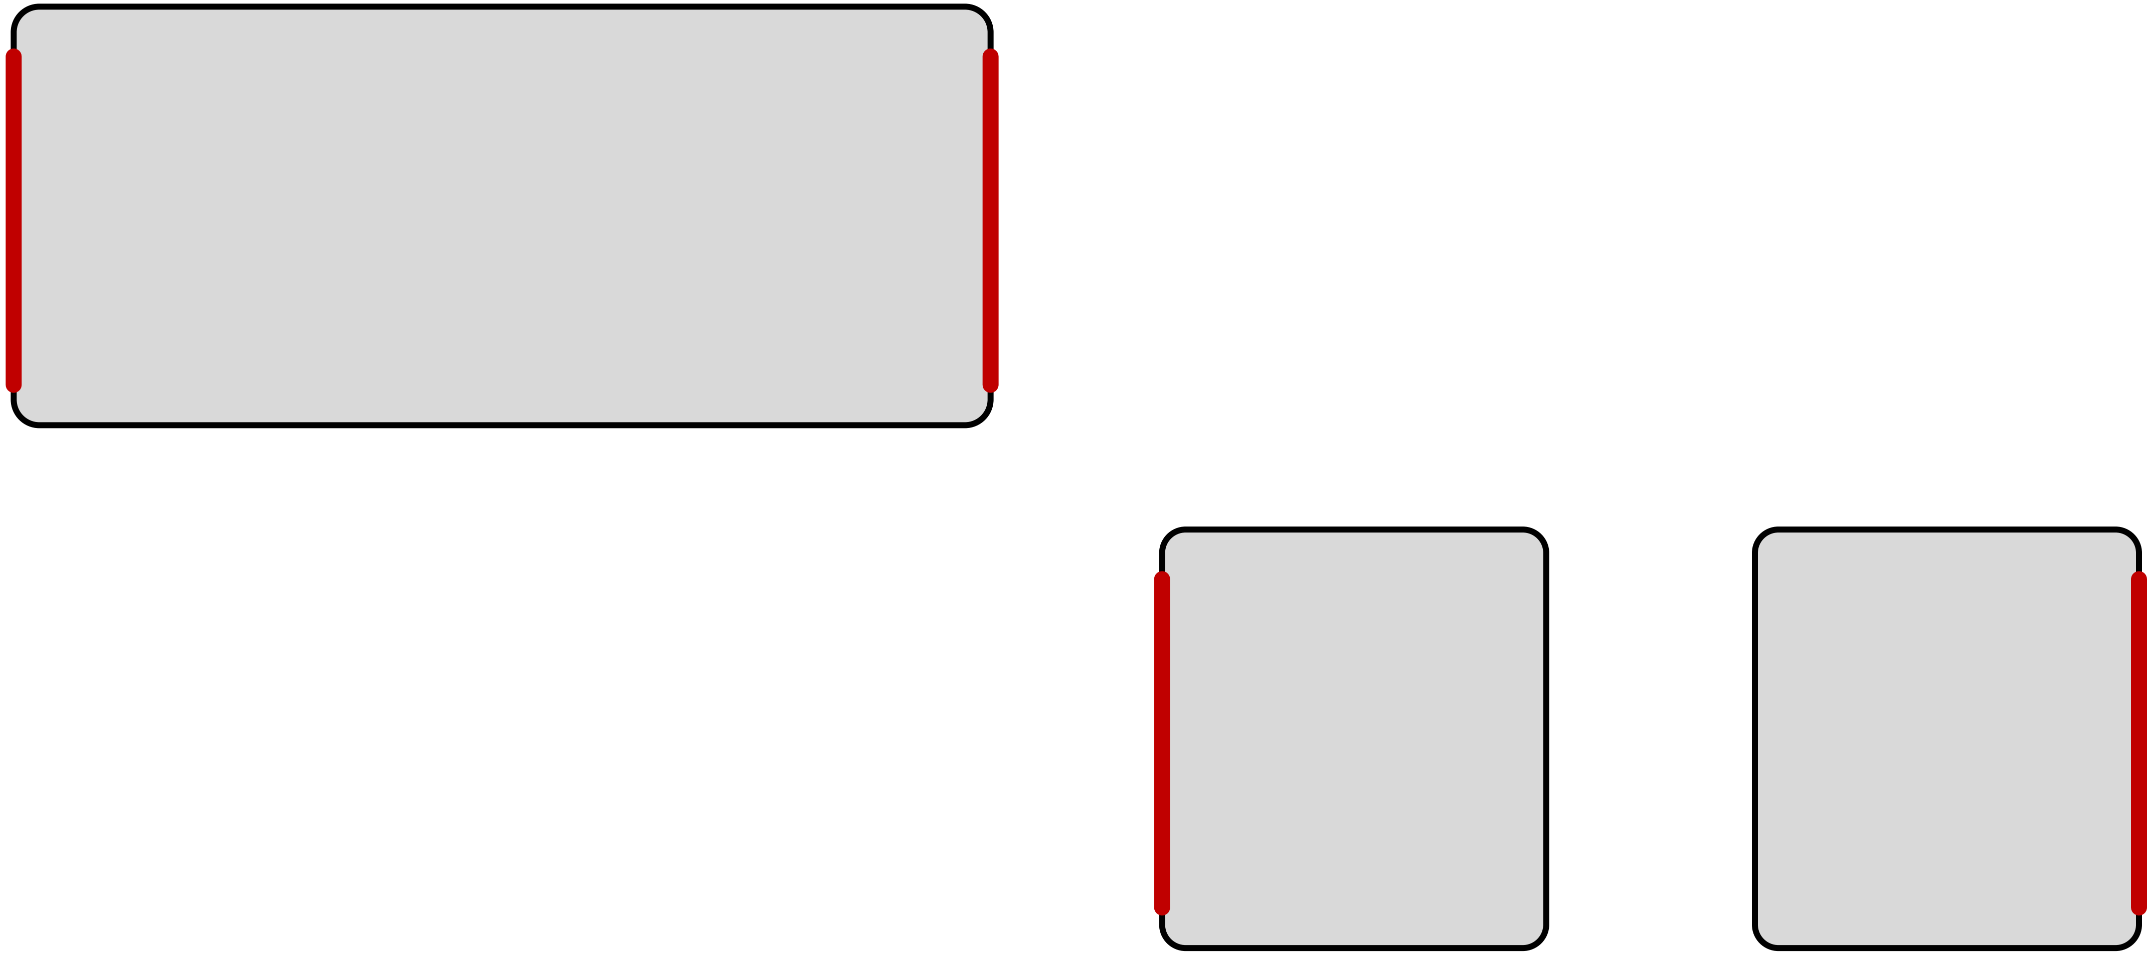

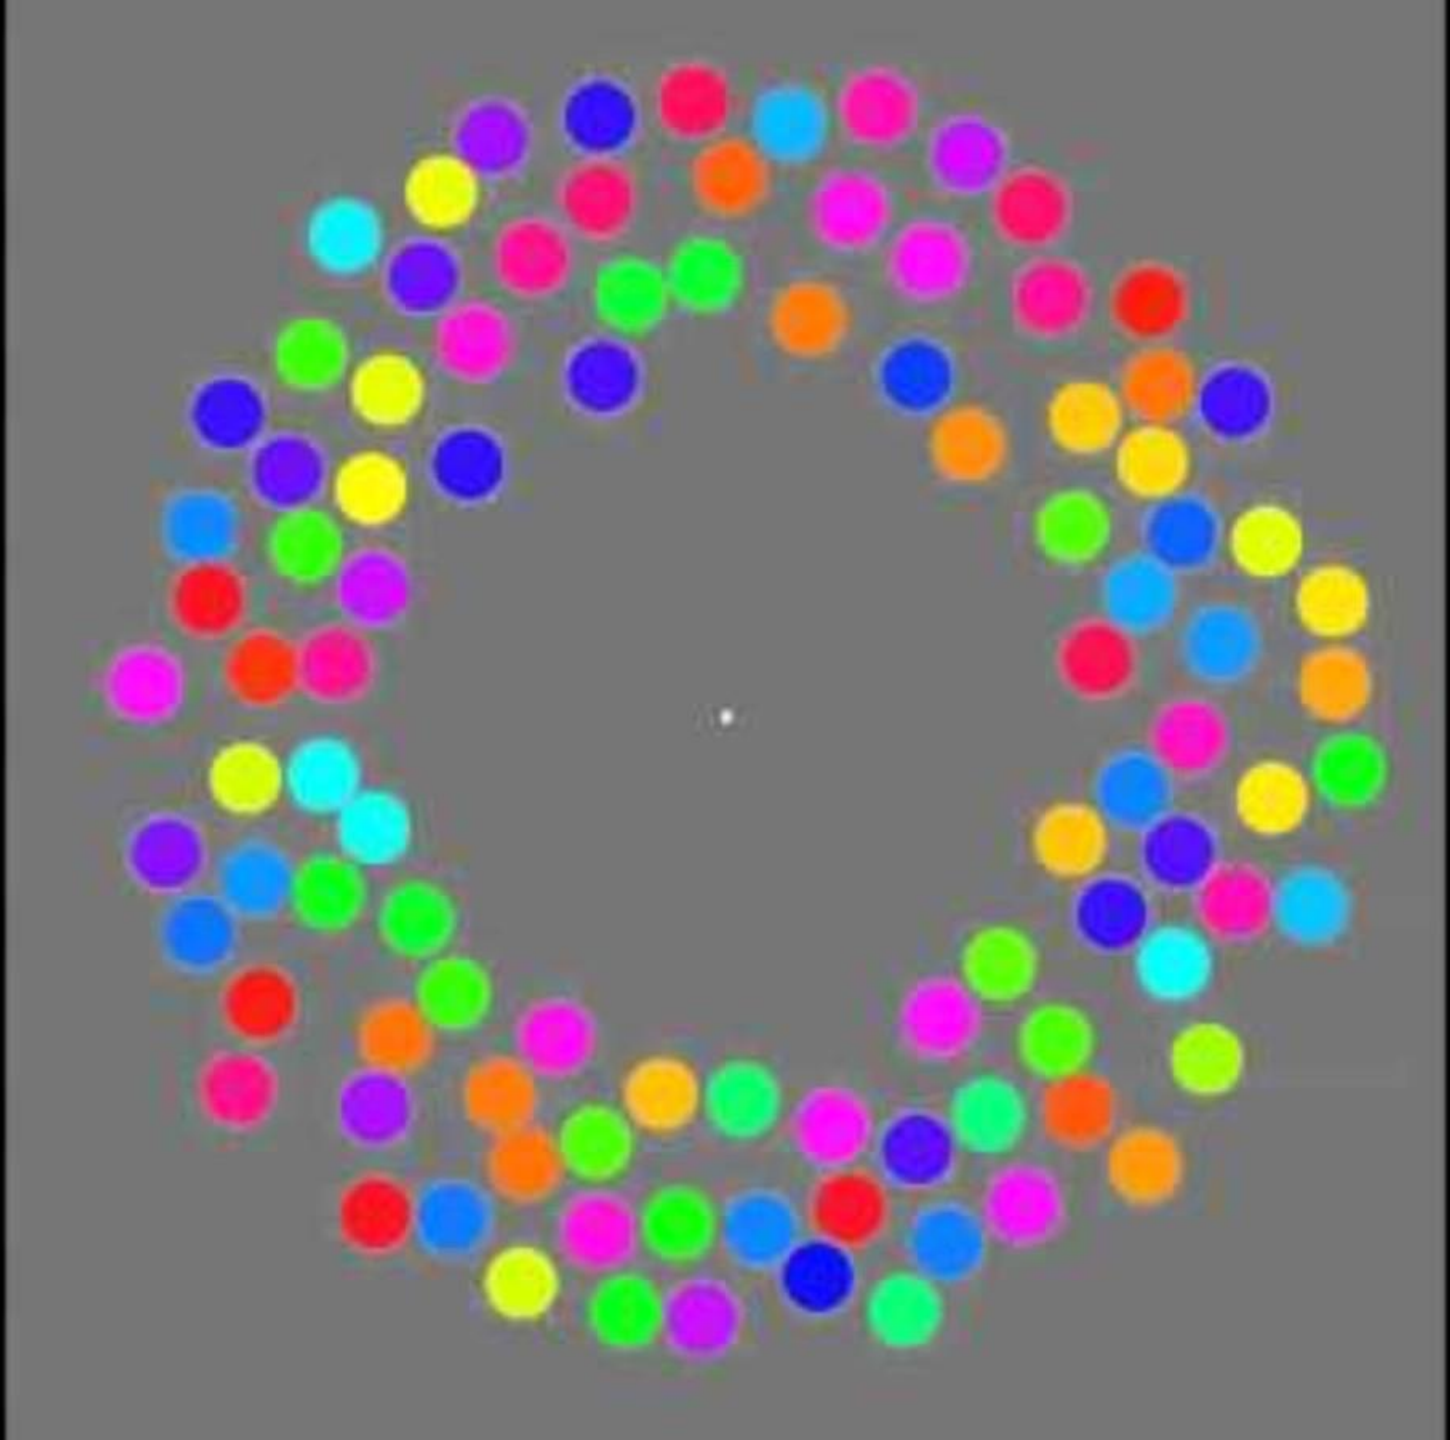

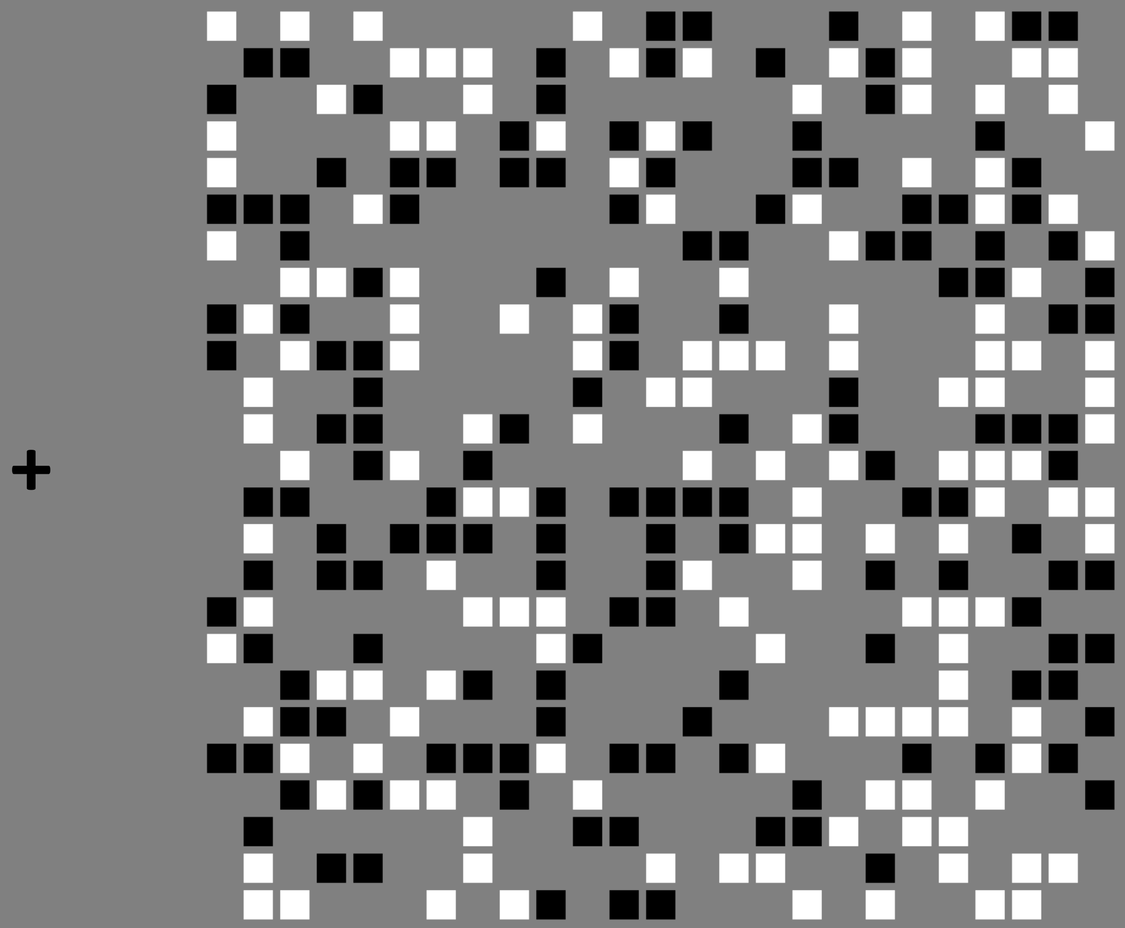

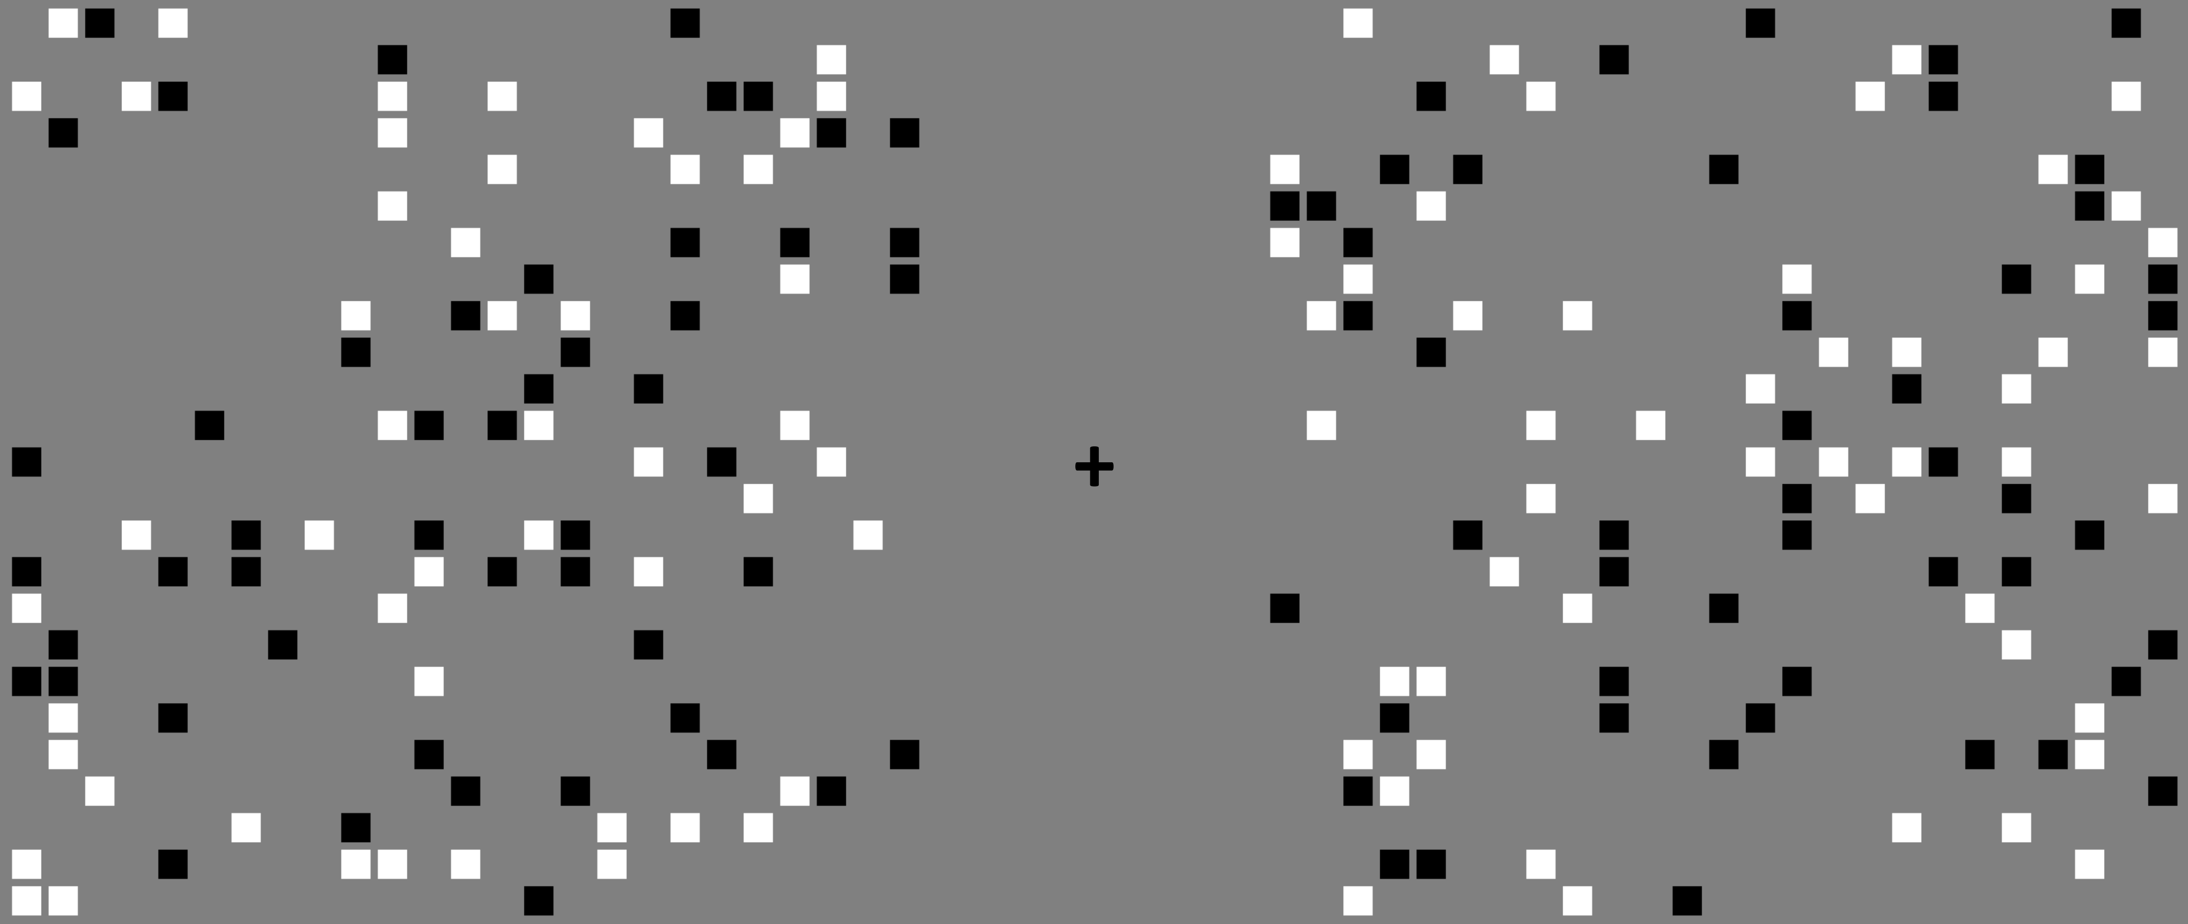


Number
adaptation?!

Current Biology 18, 425–428, March 25, 2008 DOI 10.1016/j.cub.2008.02.052

Report

A Visual Sense of Number

David Burr[1,2,*] and John Ross[2]
[1]Dipartimento di Psicologia
Università Degli Studi di Firenze
Via S. Nicolò 89
Florence 50125
Italy
[2]School of Psychology
University of Western Australia
Perth WA 6009
Australia

Summary

Evidence exists for a nonverbal capacity for the apprehension of number, in humans [1] (including infants [2, 3]) and in other primates [4–6]. Here, we show that perceived numerosity is susceptible to adaptation, like primary visual properties of a scene, such as color, contrast, size, and speed. Apparent numerosity was decreased by adaptation to large numbers of dots and increased by adaptation to small numbers, the effect depending entirely on the numerosity of the adaptor, not on contrast, size, orientation, or pixel density, and occurring with very low adaptor contrasts. We suggest that the visual system has the capacity to estimate numerosity and that it is an independent primary visual property, not reducible to others like spatial frequency or density of texture [7].

Results and Discussion

Jevons, a 19[th] century economist, rather than counting beans, assessed his accuracy in estimating the number of beans in a box at a single glance [8]. He made no errors at four or below but became increasingly inaccurate as the number of beans increased beyond four. Subsequent studies have confirmed his findings and the lack of errors below five has led to the concept of *subitizing* [9–12], usually presumed to be a separate process allowing rapid apprehension of the numerosity of collections containing fewer than five objects. The perception of larger numbers is usually assumed to require slower and more cognitive processes, like counting.

All primary visual properties are susceptible to *adaptation*, sometimes giving rise to dramatic aftereffects, like the waterfall illusion [13], and changes in color, size, distance, spatial frequency, and orientation. If numerosity was a primary visual property, like color or motion, it too should be prone to adaptation. The online demonstration shows that it is. After 30 s adaptation to the two different adaptor patches, the two subsequent patches appear to differ considerably in numerosity (whereas inspection after adaptation wears off, or counting, shows that they both number 30 dots). We quantified adaptation effects by asking subjects whether a *test* stimulus (of variable numerosity), presented to the region that had been adapted, appeared more or less numerous than a *probe* stimulus (of fixed numerosity), presented to a different unadapted position a little later. The proportion of trials where the test appeared more numerous than the probe was plotted against test numerosity and fitted with cumulative Gaussian functions whose mean estimates the point of subjective equality (PSE) between test and probe, and standard deviation the threshold for discriminating between the two (the just-noticeable difference [jnd]). Figure 1A shows sample psychometric functions for a 30 element probe, with and without adaptation to a 400 element stimulus. The ratio of the matched test to probe increases from unity (30 dots) with no adaptation to more than 3 (100 dots) after adaptation (we increased the test number to compensate for the reduction in its apparent numerosity). Note also that that after adaptation the psychometric function is steeper (on logarithmic coordinates), implying a smaller jnd.

We first measured the effect of adapting to a large number (400) of dots as a function of number of dots in the probe (Figure 1B). The amount of adaptation was fairly constant with probe numerosity down to about 12 dots and then decreased as the probe approached the subitizing range. The precision of the match, given by the jnd or Weber fraction (jnd expressed as a fraction of dot number), did not deteriorate during adaptation, the average percentage Weber fractions for unadapted and adapted conditions being 28% for unadapted and 26% for the adapted conditions.

We next investigated whether adaptation to small numbers can cause an increase in apparent numerosity. The red circles of Figure 2 show that adaptation occurred in both directions: Adaptation to small numbers increased apparent numerosity (so the matched number decreased), and adaptation to large numbers decreased apparent numerosity. Adaptation to 50 dots (the number of the probe) had no effect, with the amount of adaptation increasing with the difference between adapt and probe number. The curves of both subjects were well fit by linear regression on log coordinates, with a slope around 0.25.

In order to test whether adaptation depends on numerosity per se or is derived from other factors, like texture density [7], we performed a number of controls. We first varied the size of the adaptor and test dots, in order to vary pixel density. In the above-described study (red circles of Figure 2), both adaptor and test dots were circles of 6 pixel (20 arcmin) diameter (28 pixel area). We repeated the experiment with square adaptor stimuli of 8 × 8 pixels (64 pixels) and test stimuli of 3 × 3 pixels (9 pixels, 1/7 as many as the adaptor). If pixel density were the relevant attribute, the curves of Figure 2 should shift leftwards by a factor of 7, so the null point occurs when adaptor and test pixel density are matched (for adaptation dot number of 7). This clearly does not occur. For naive observer PB, the curves remain superimposed; for DB, there is a slight shift in the opposite direction.

We also examined the effect of adaptor contrast. As Figure 2C shows, contrast of adaptor dots had little effect on the magnitude of adaptation. At contrasts as low as 12%, the adaptation effect is still nearly 2-fold, dropping only near detection threshold. It appears that the only factor that affects adaptation is numerosity, not density, orientation, or contrast.

*Correspondence: dave@in.cnr.it

But wait!

‘Old news’ hypothesis

‘*Old news*’ hypothesis

Two quick proofs of concept.

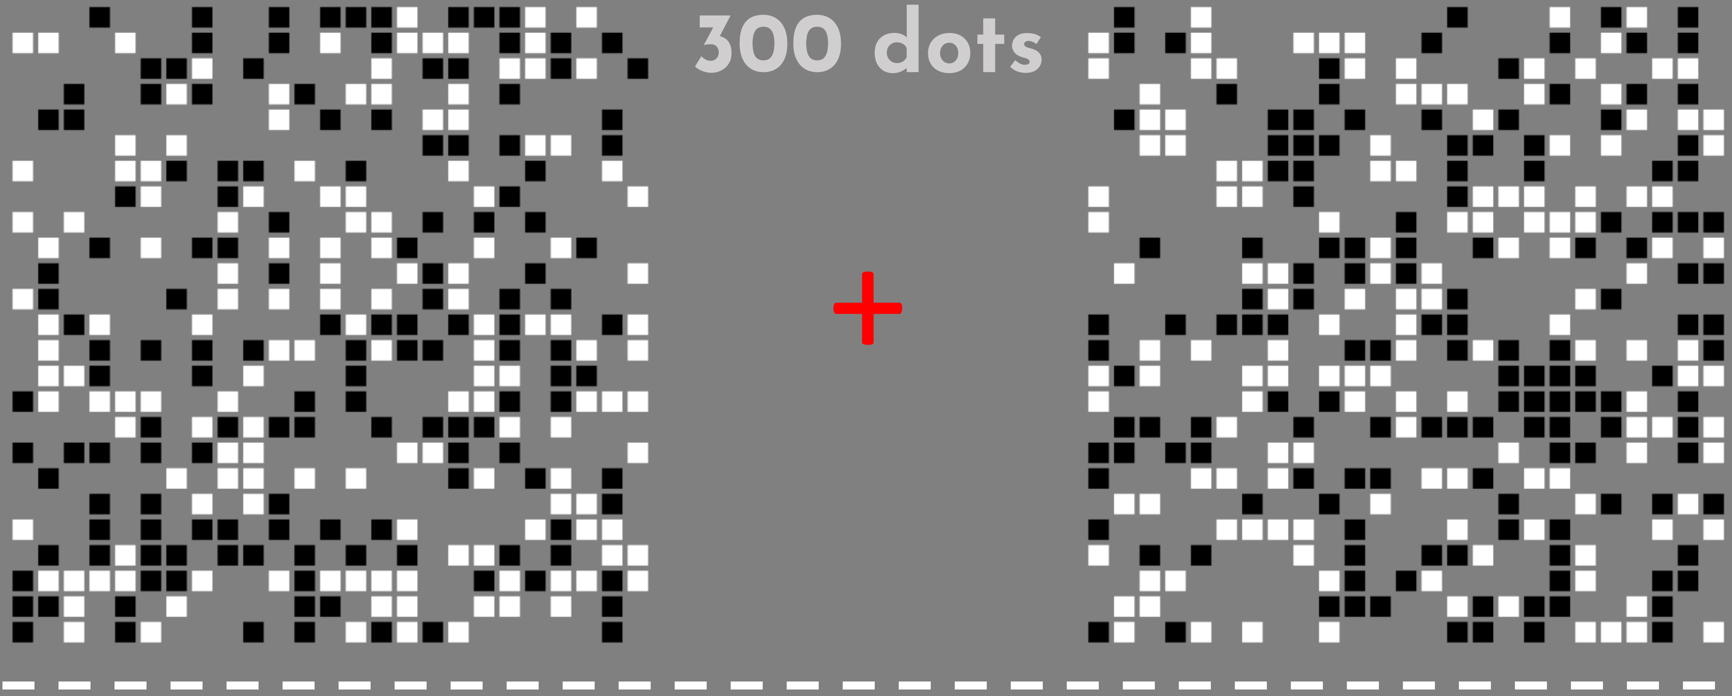
300 dots

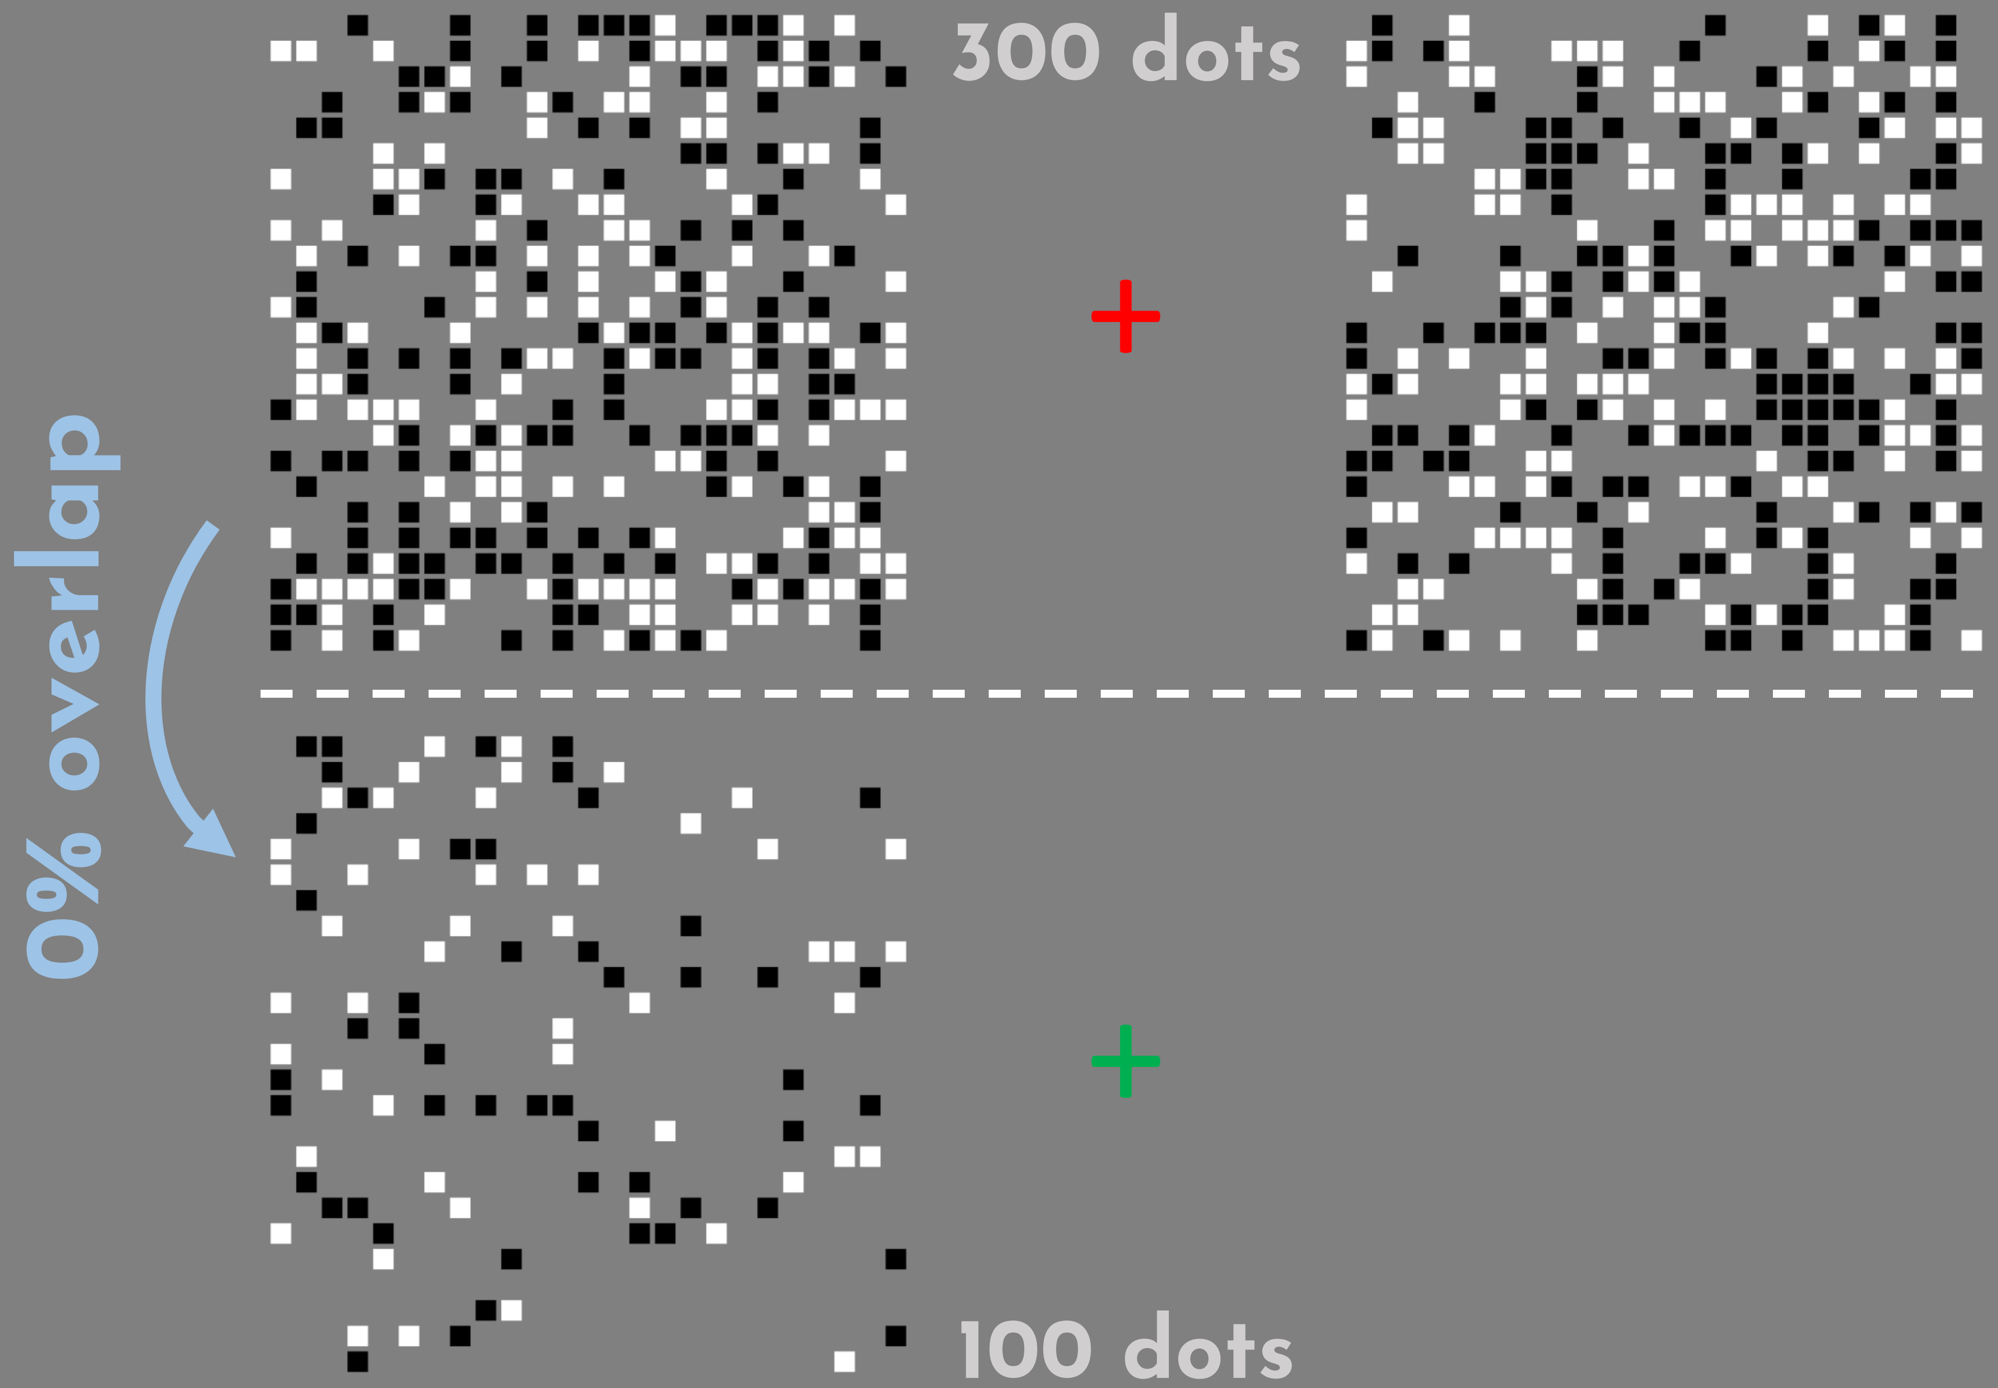
300 dots
0% overlap
100 dots

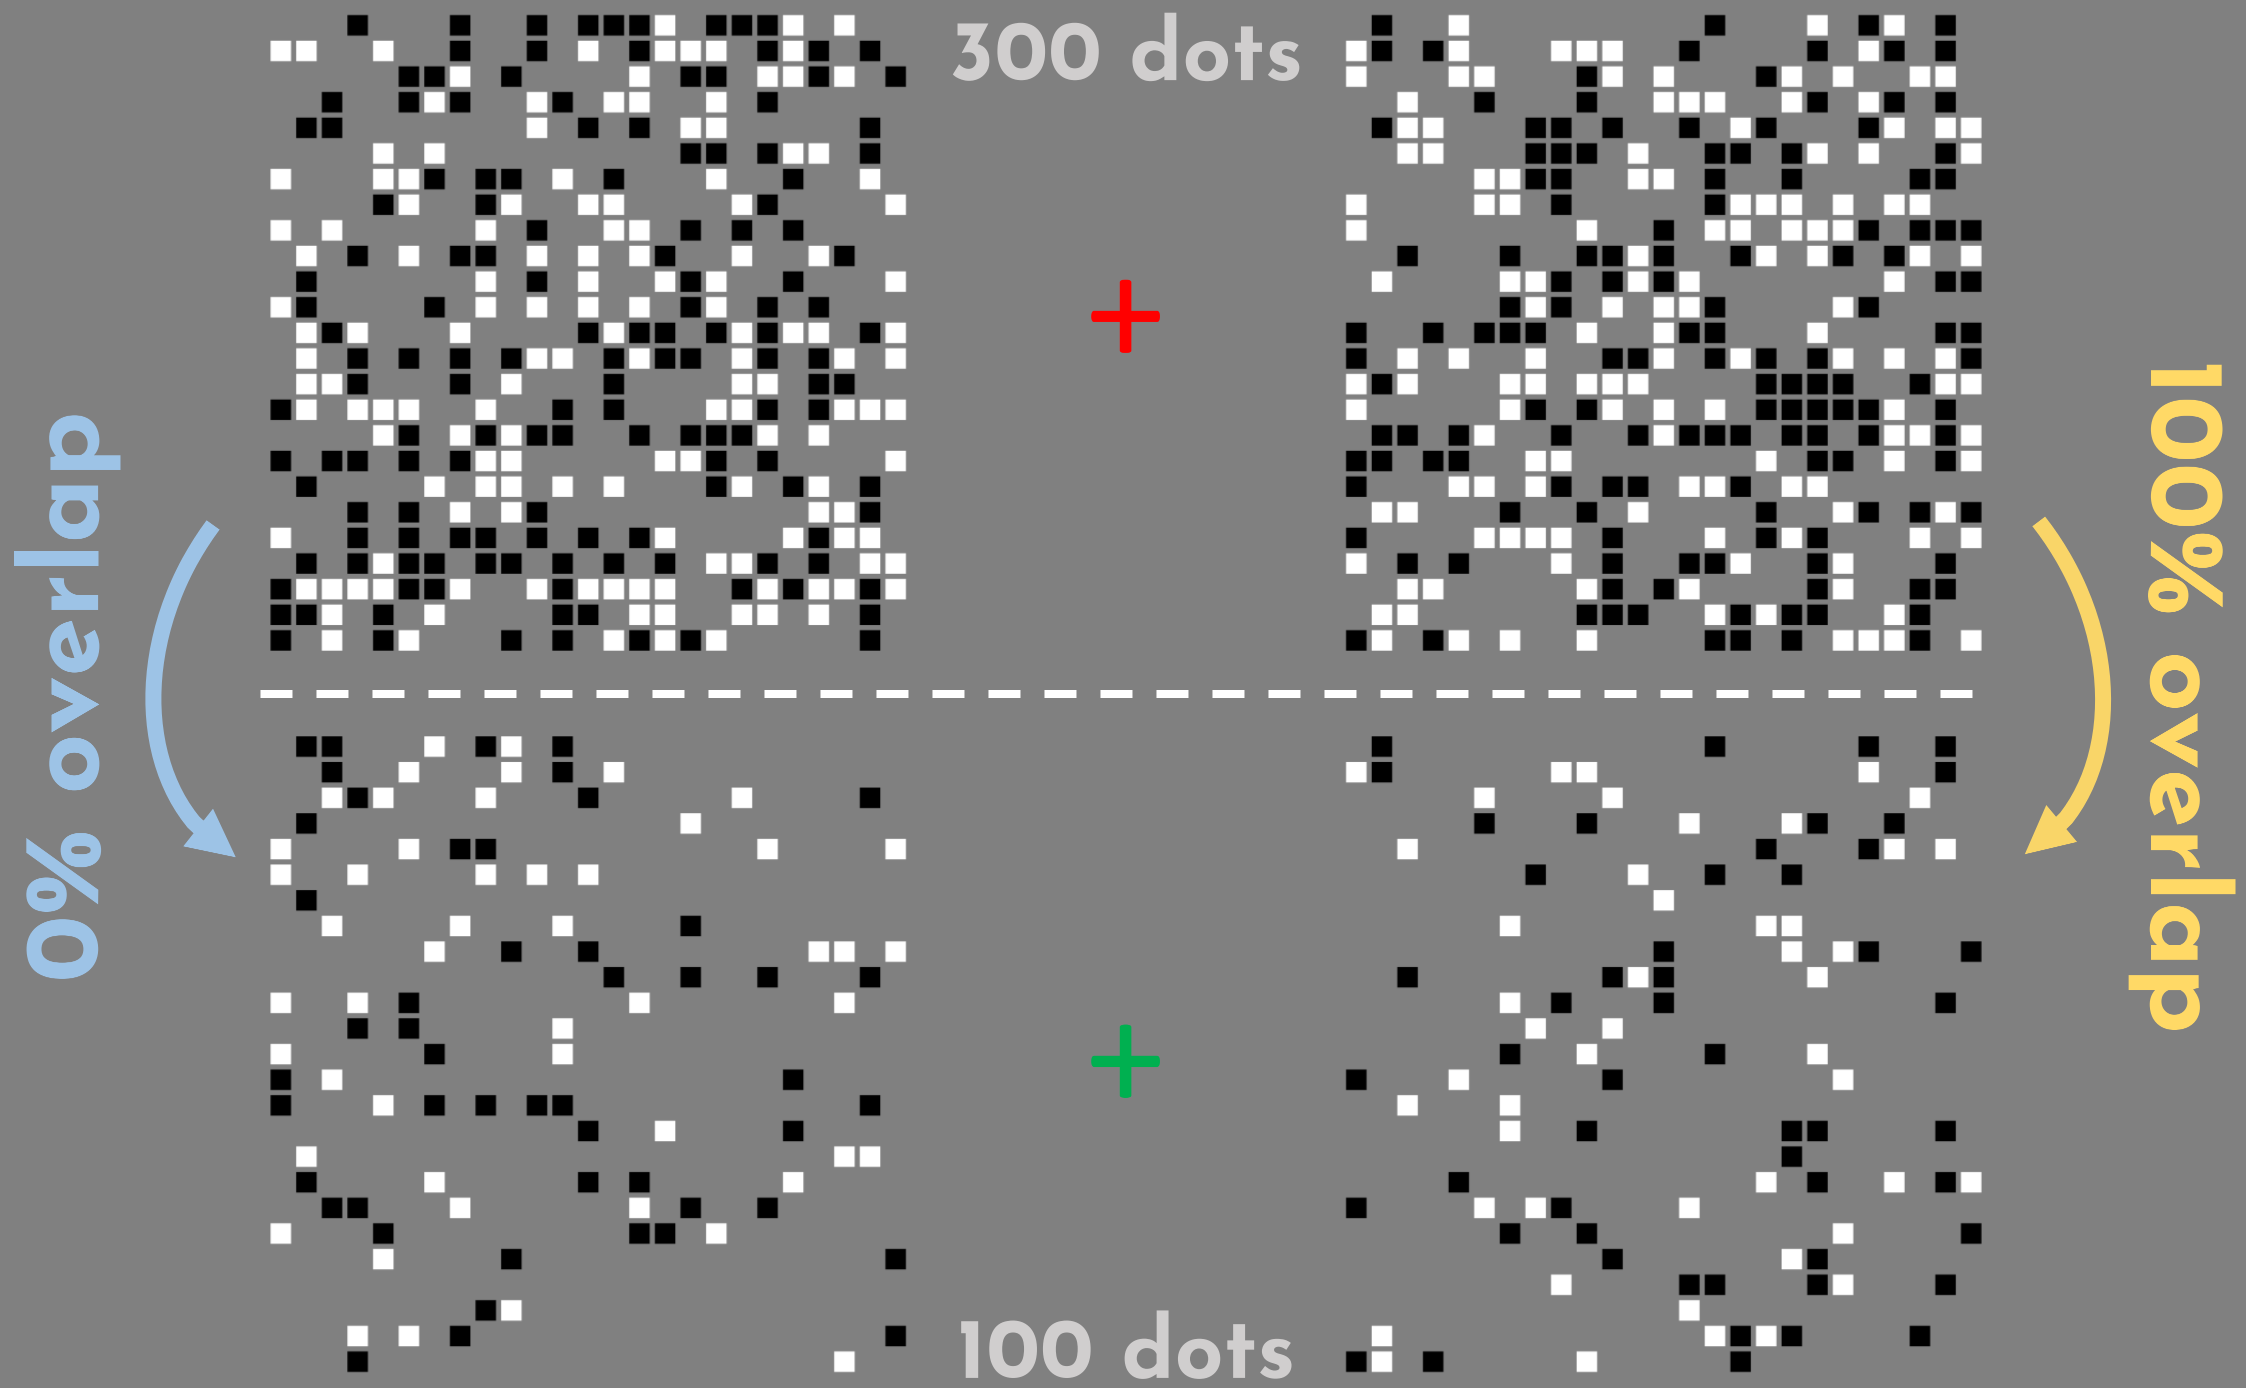
300 dots
0% overlap
100% overlap
100 dots

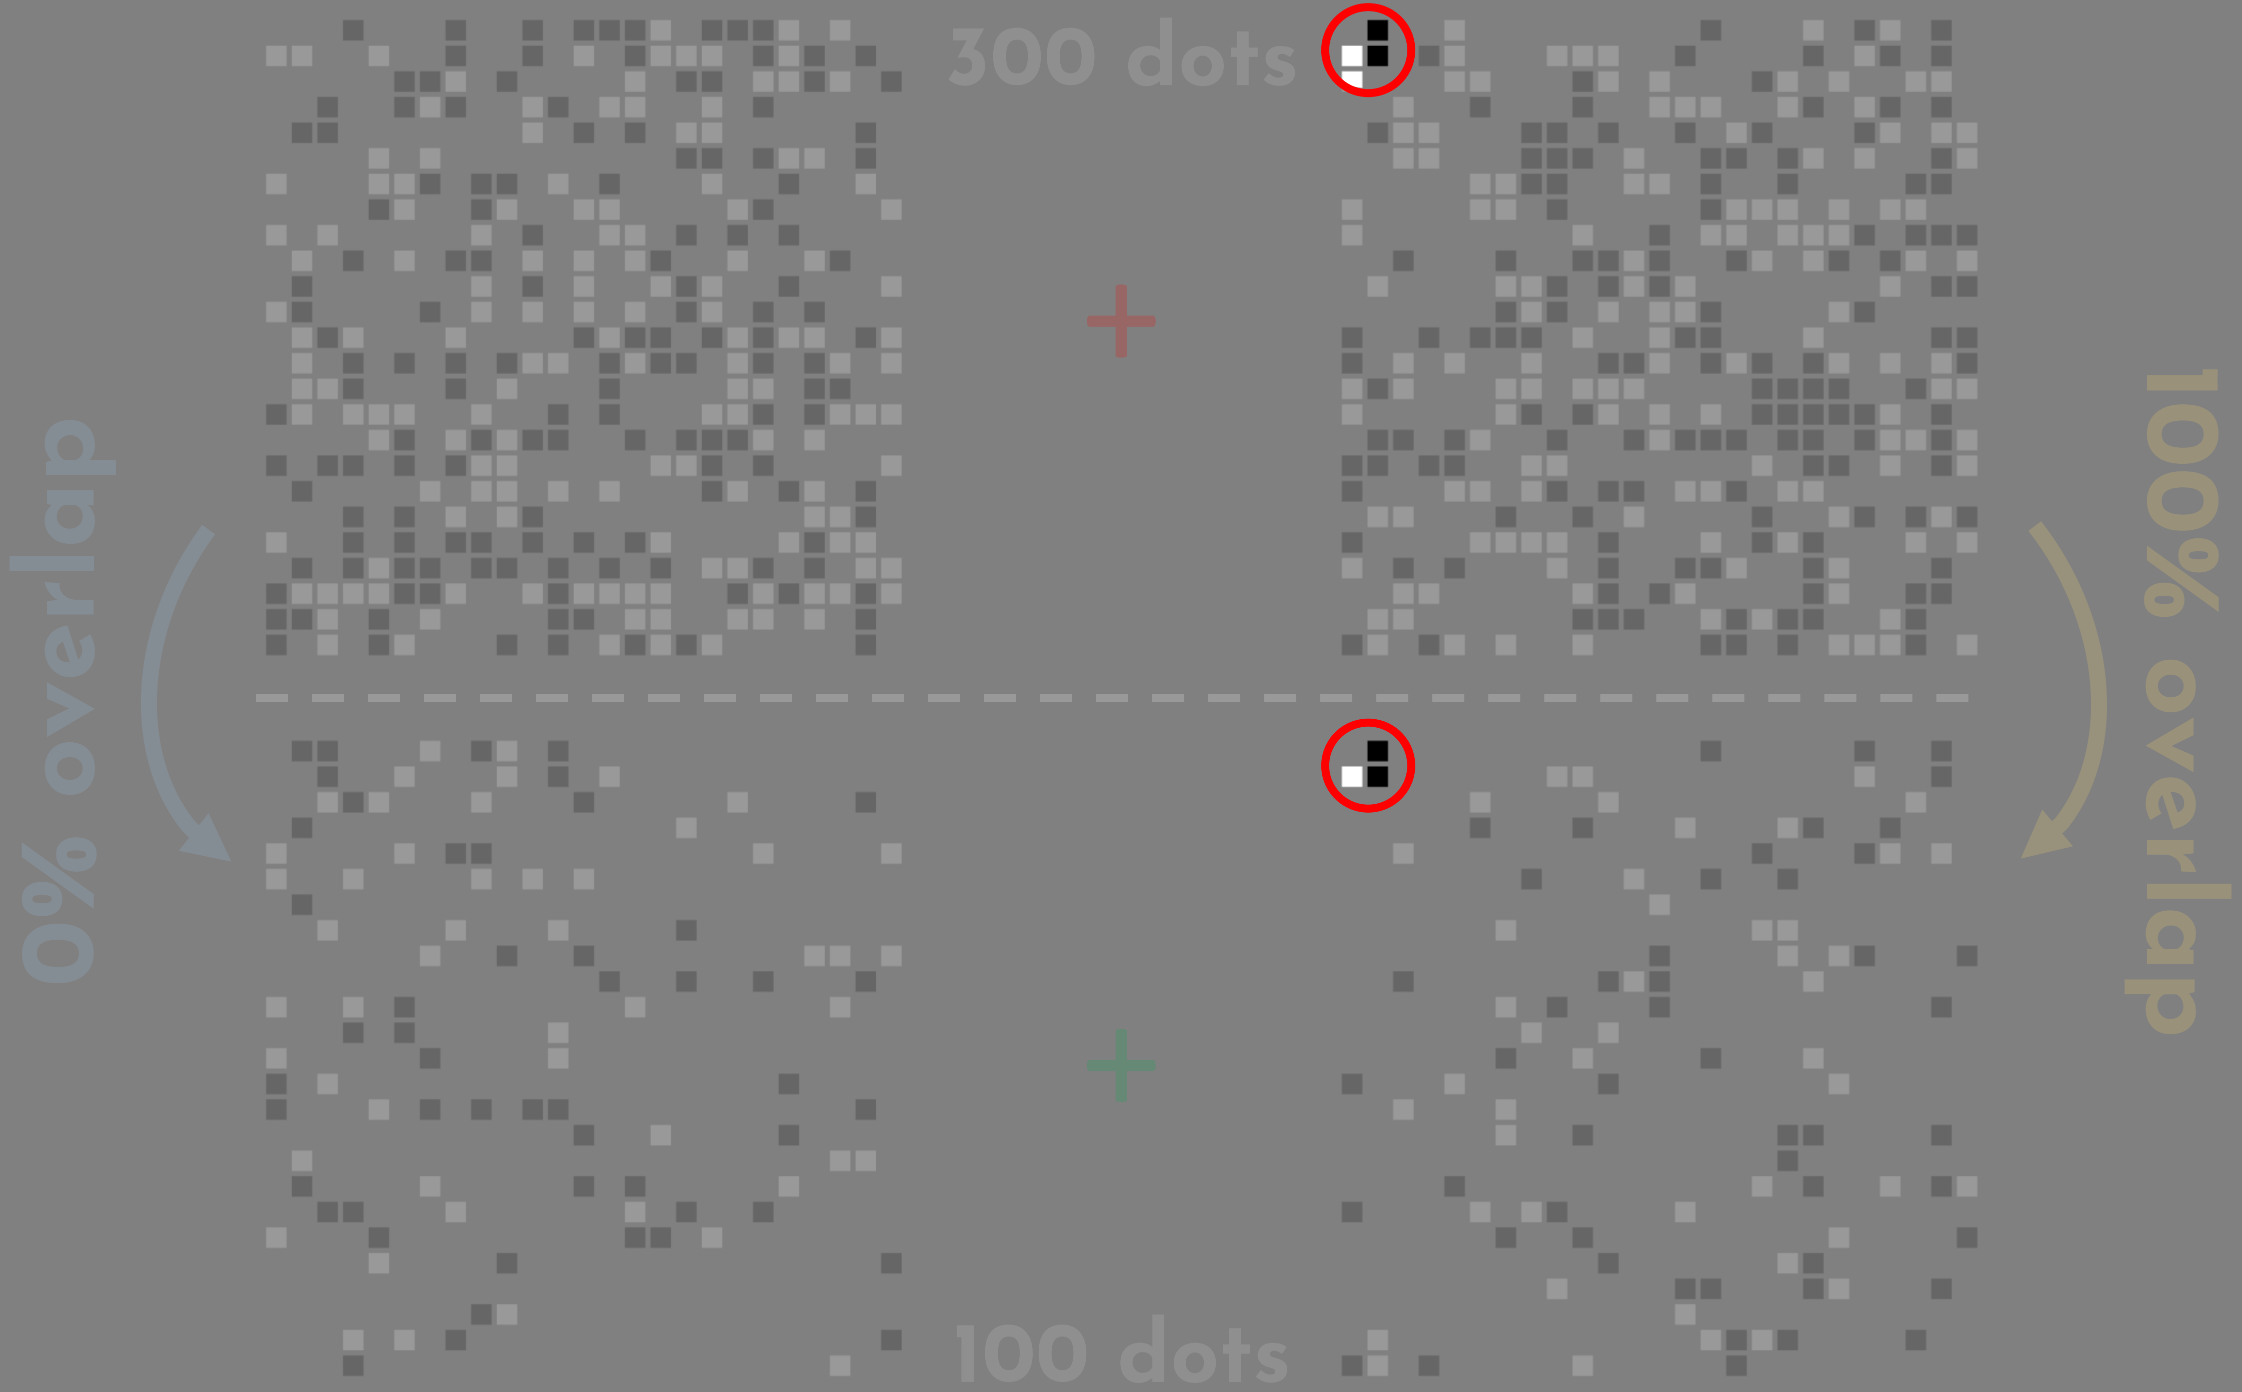
300 dots
100 dots
0% overlap
100% overlap

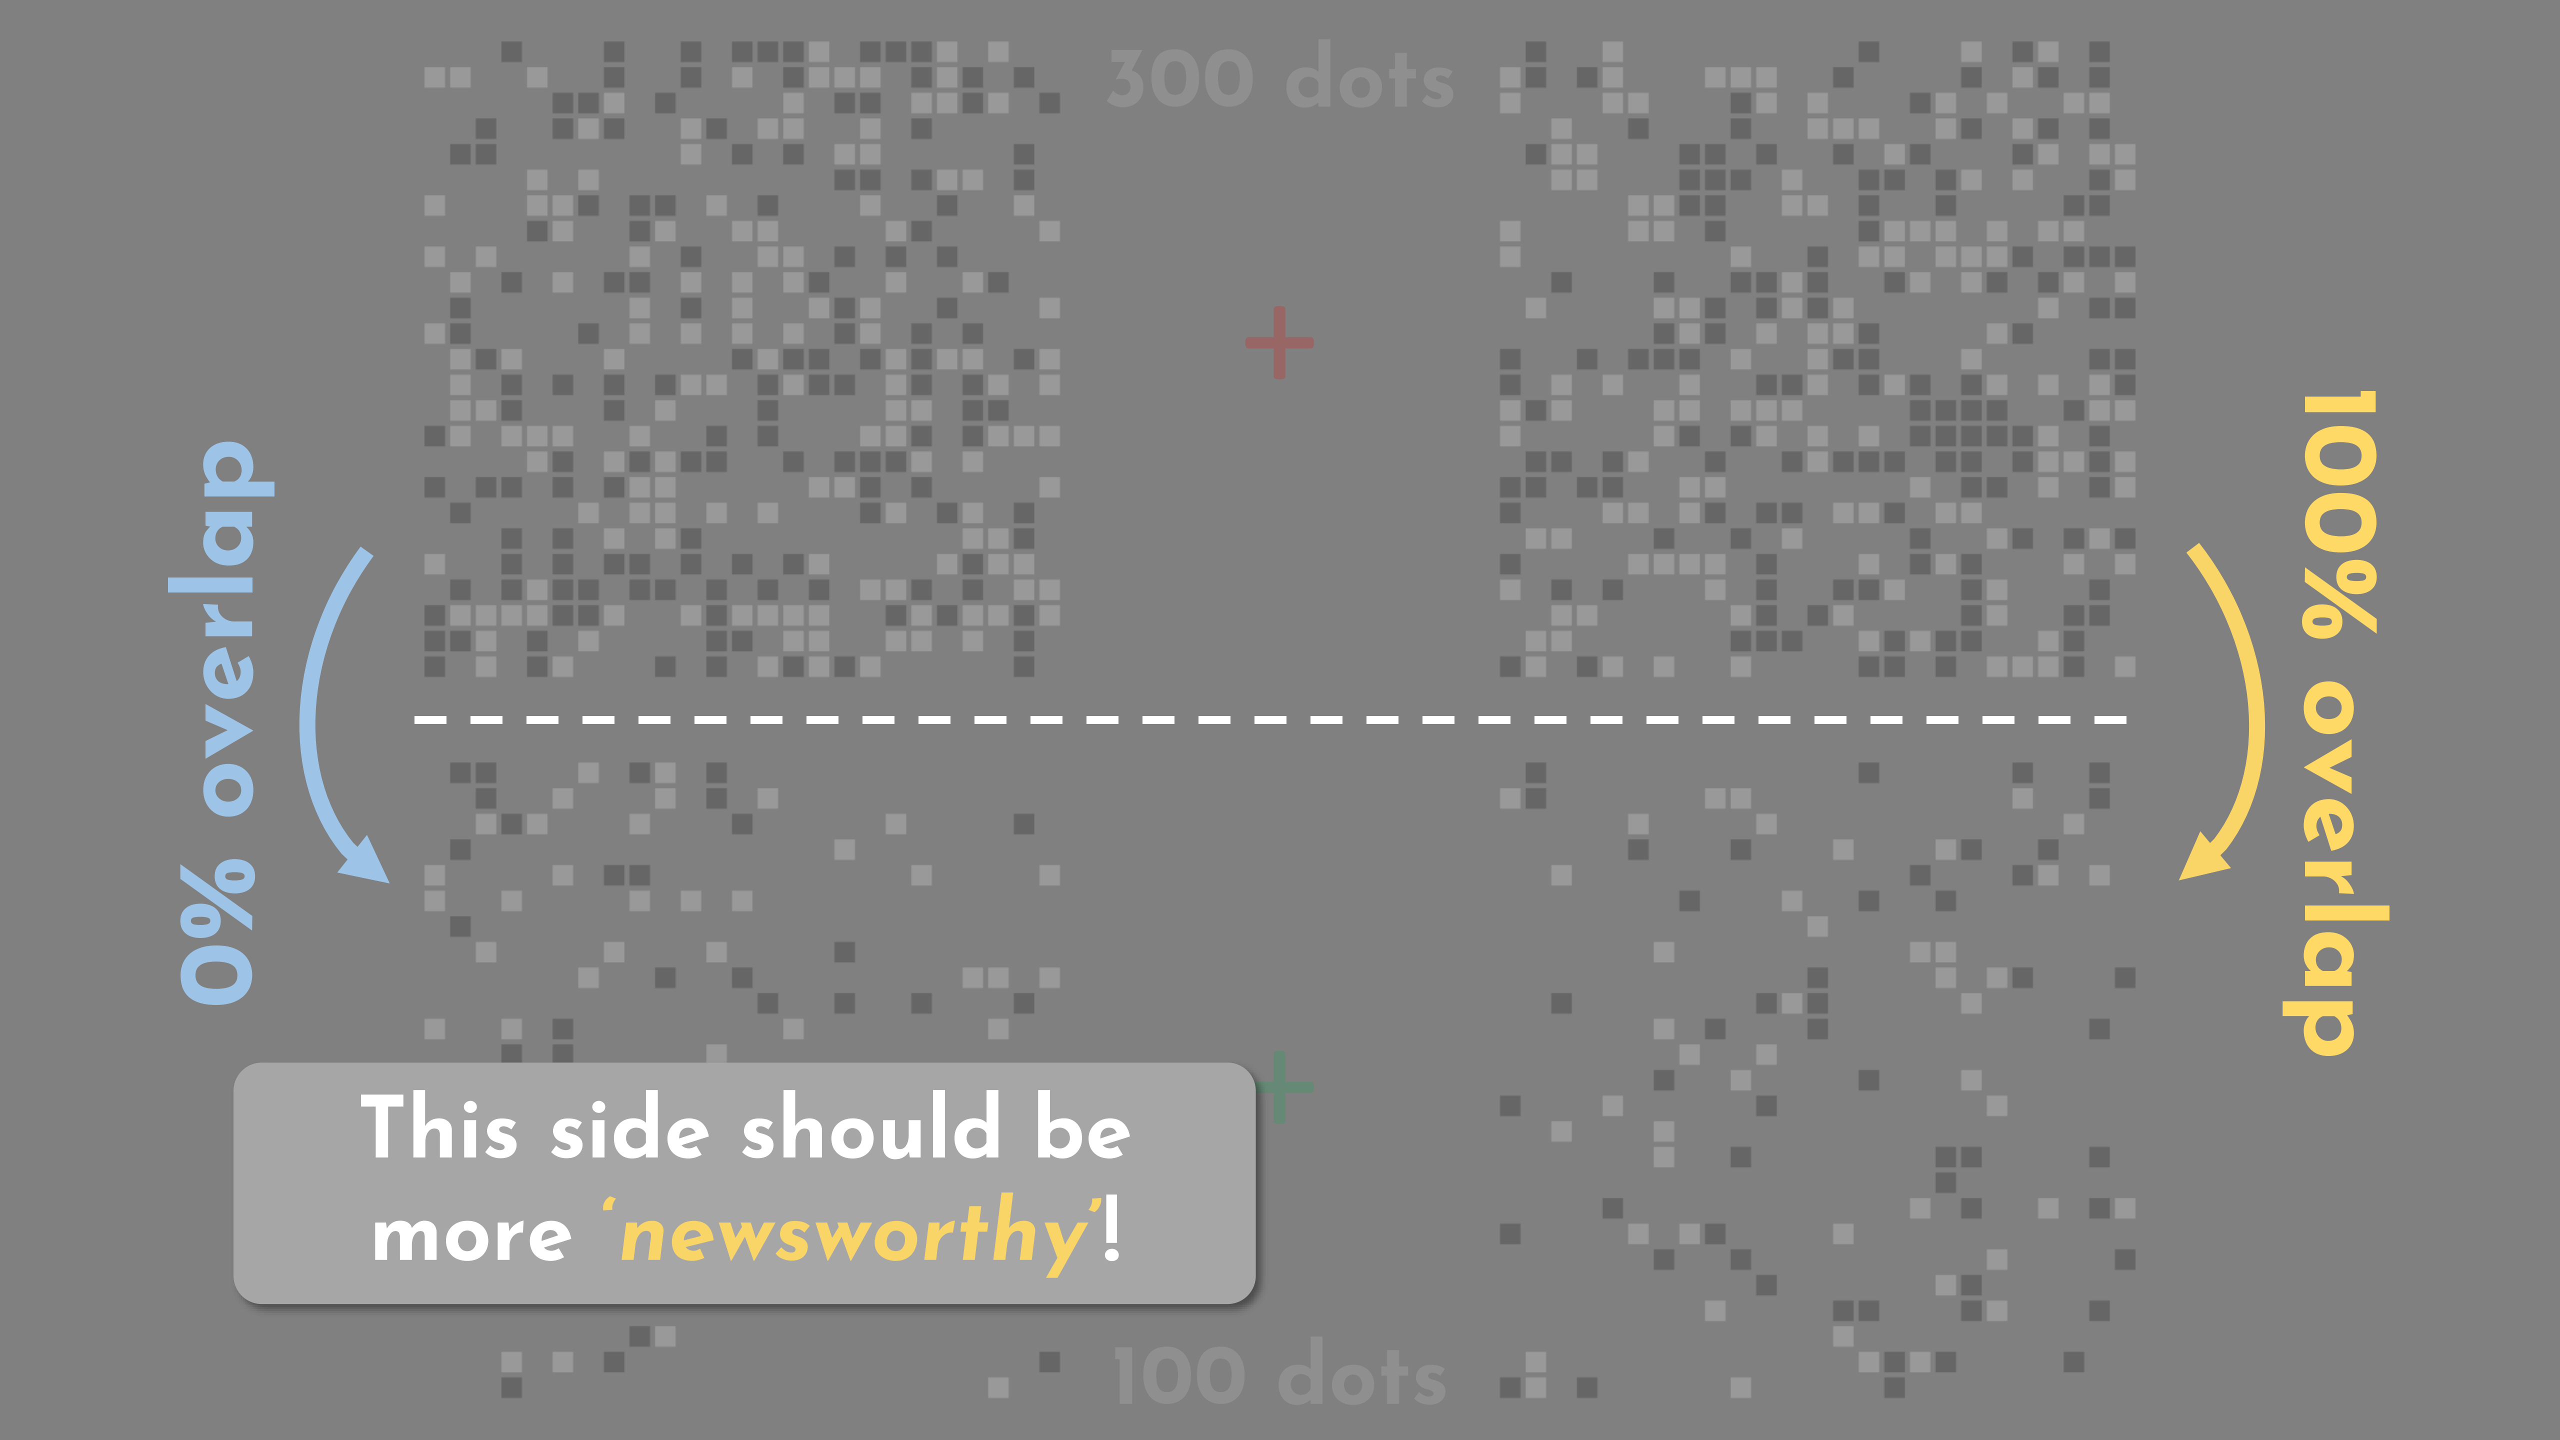
300 dots
0% overlap
100% overlap
100 dots
This side should be more 'newsworthy'!

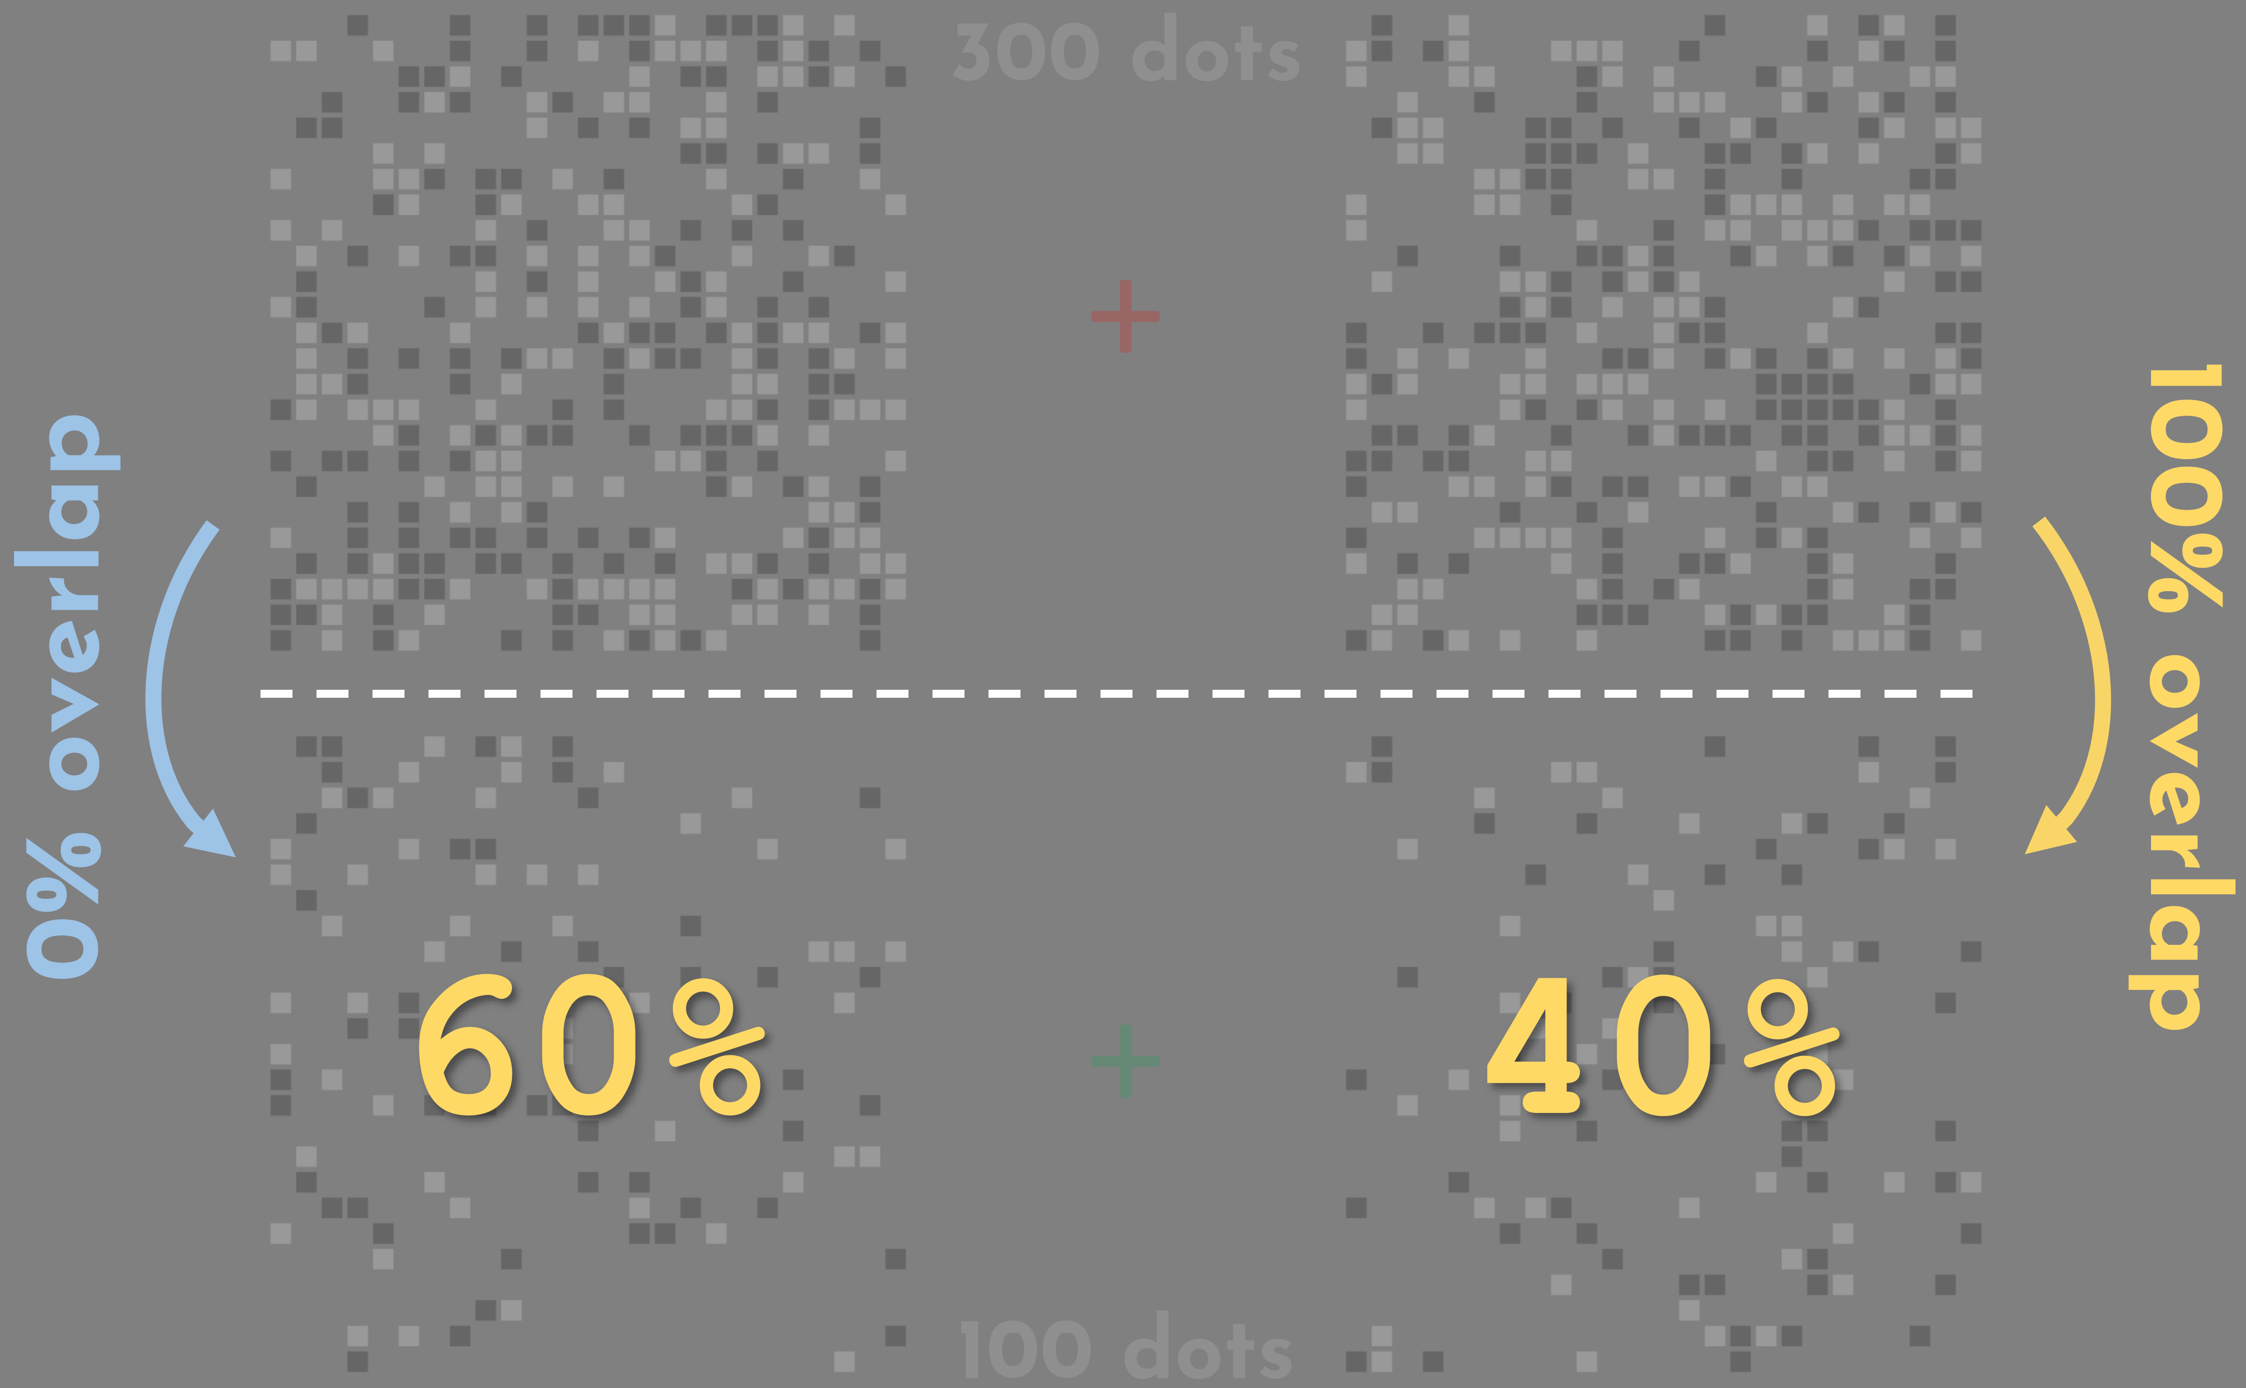
300 dots
0% overlap
100% overlap
60%
40%
100 dots

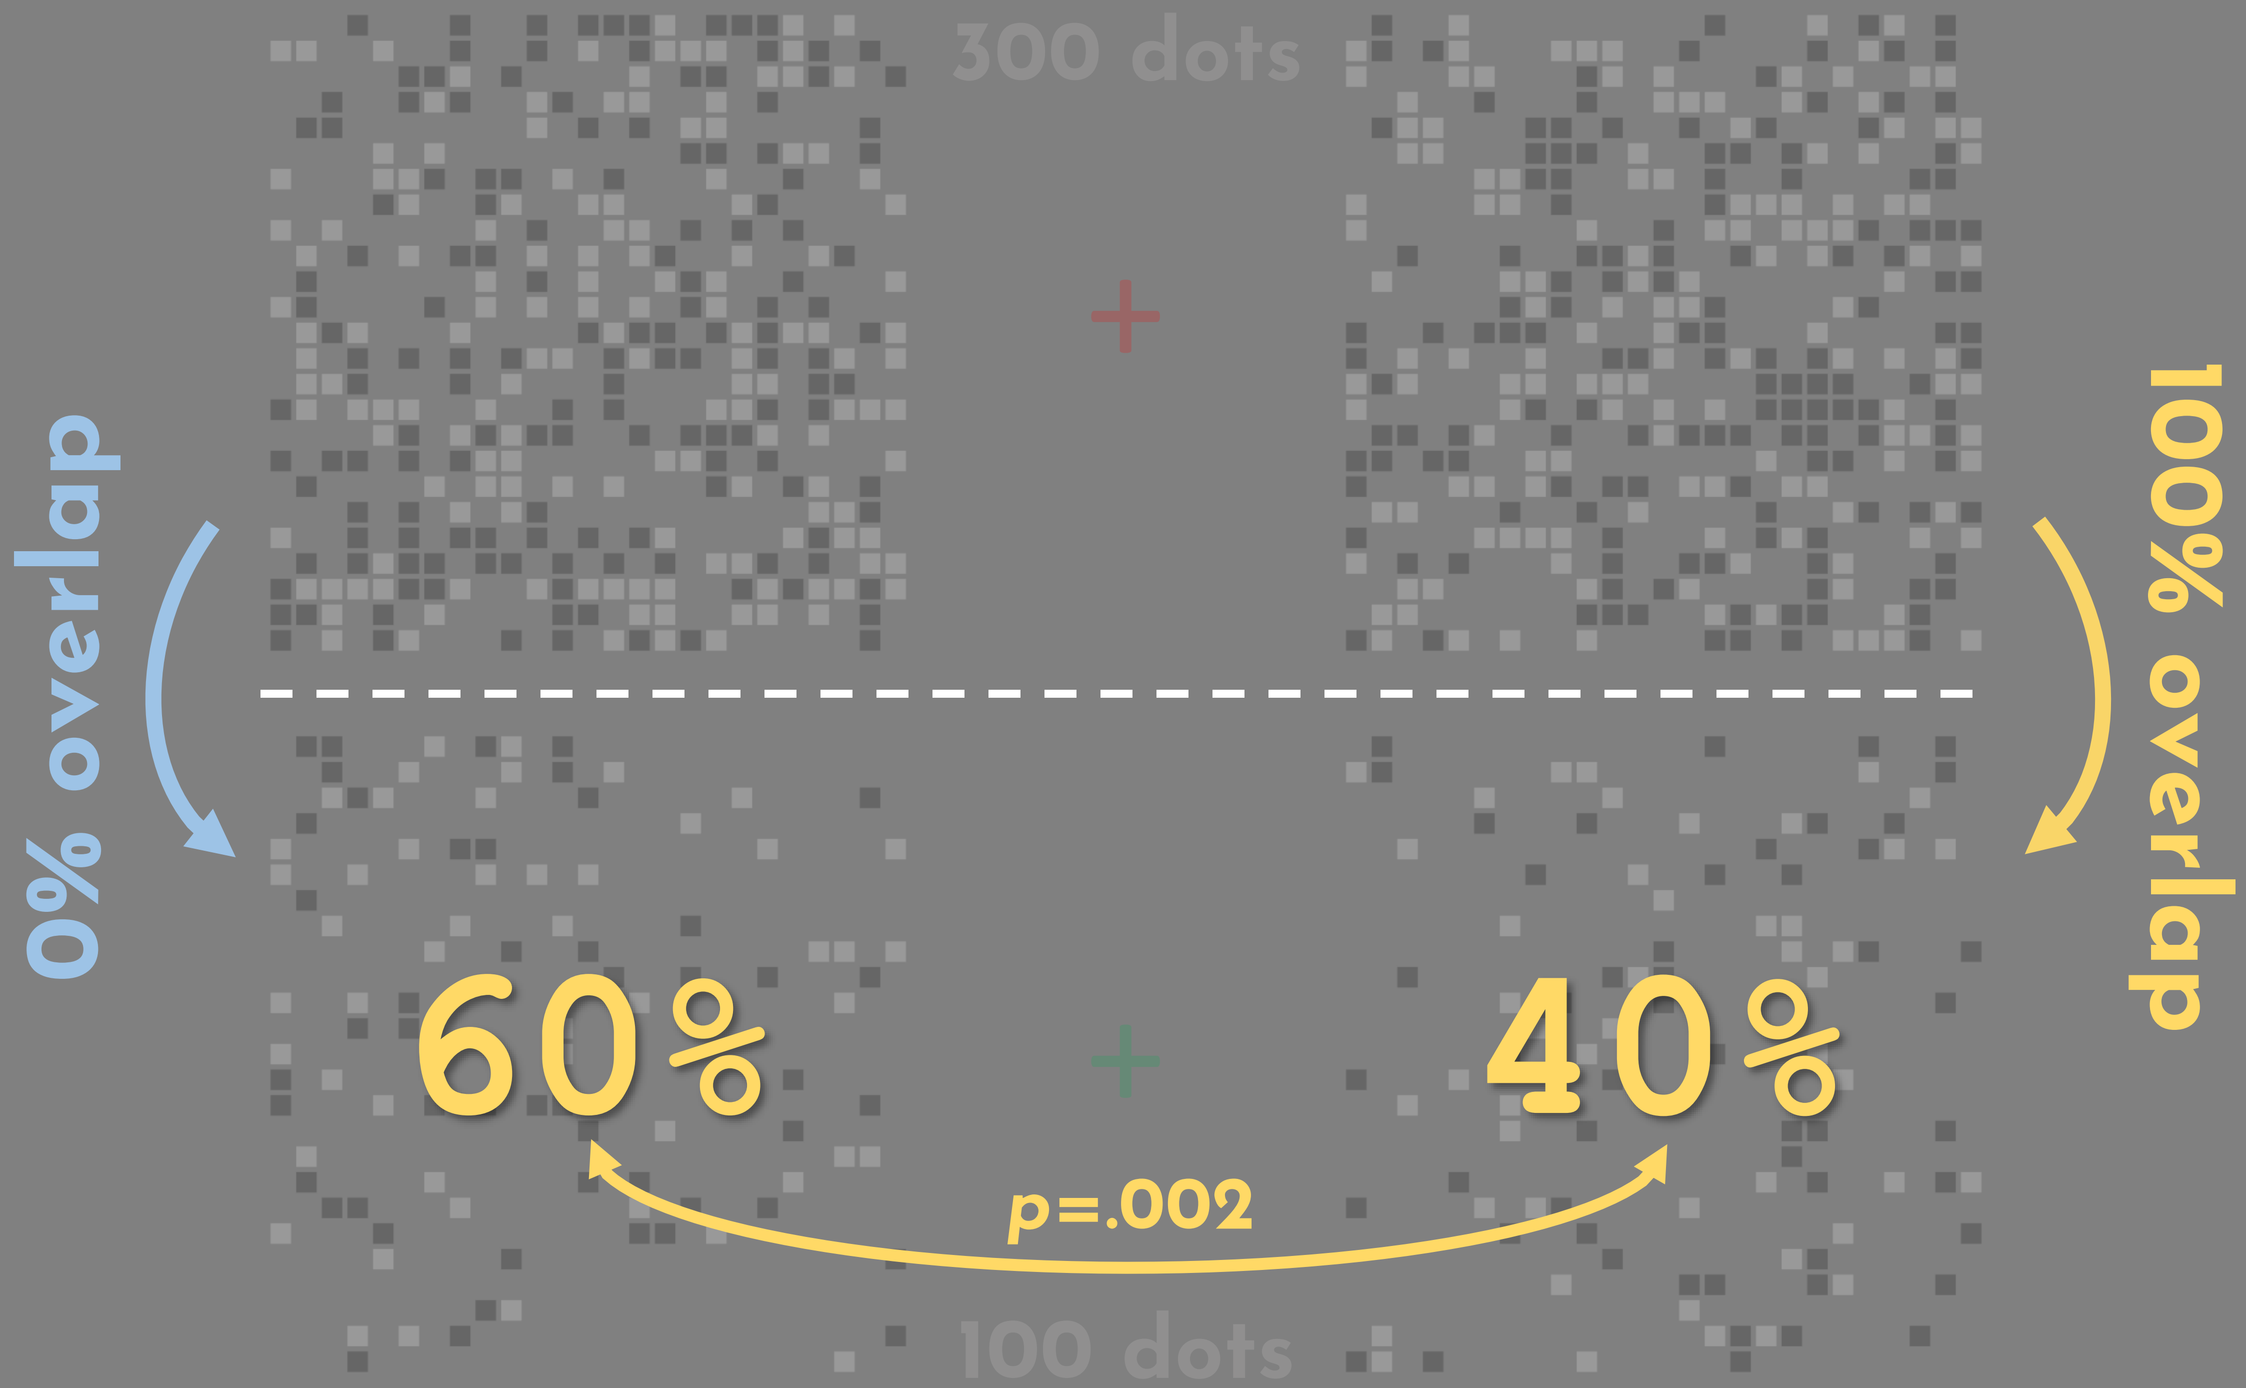
300 dots
0% overlap
100% overlap
60%
40%
p=.002
100 dots

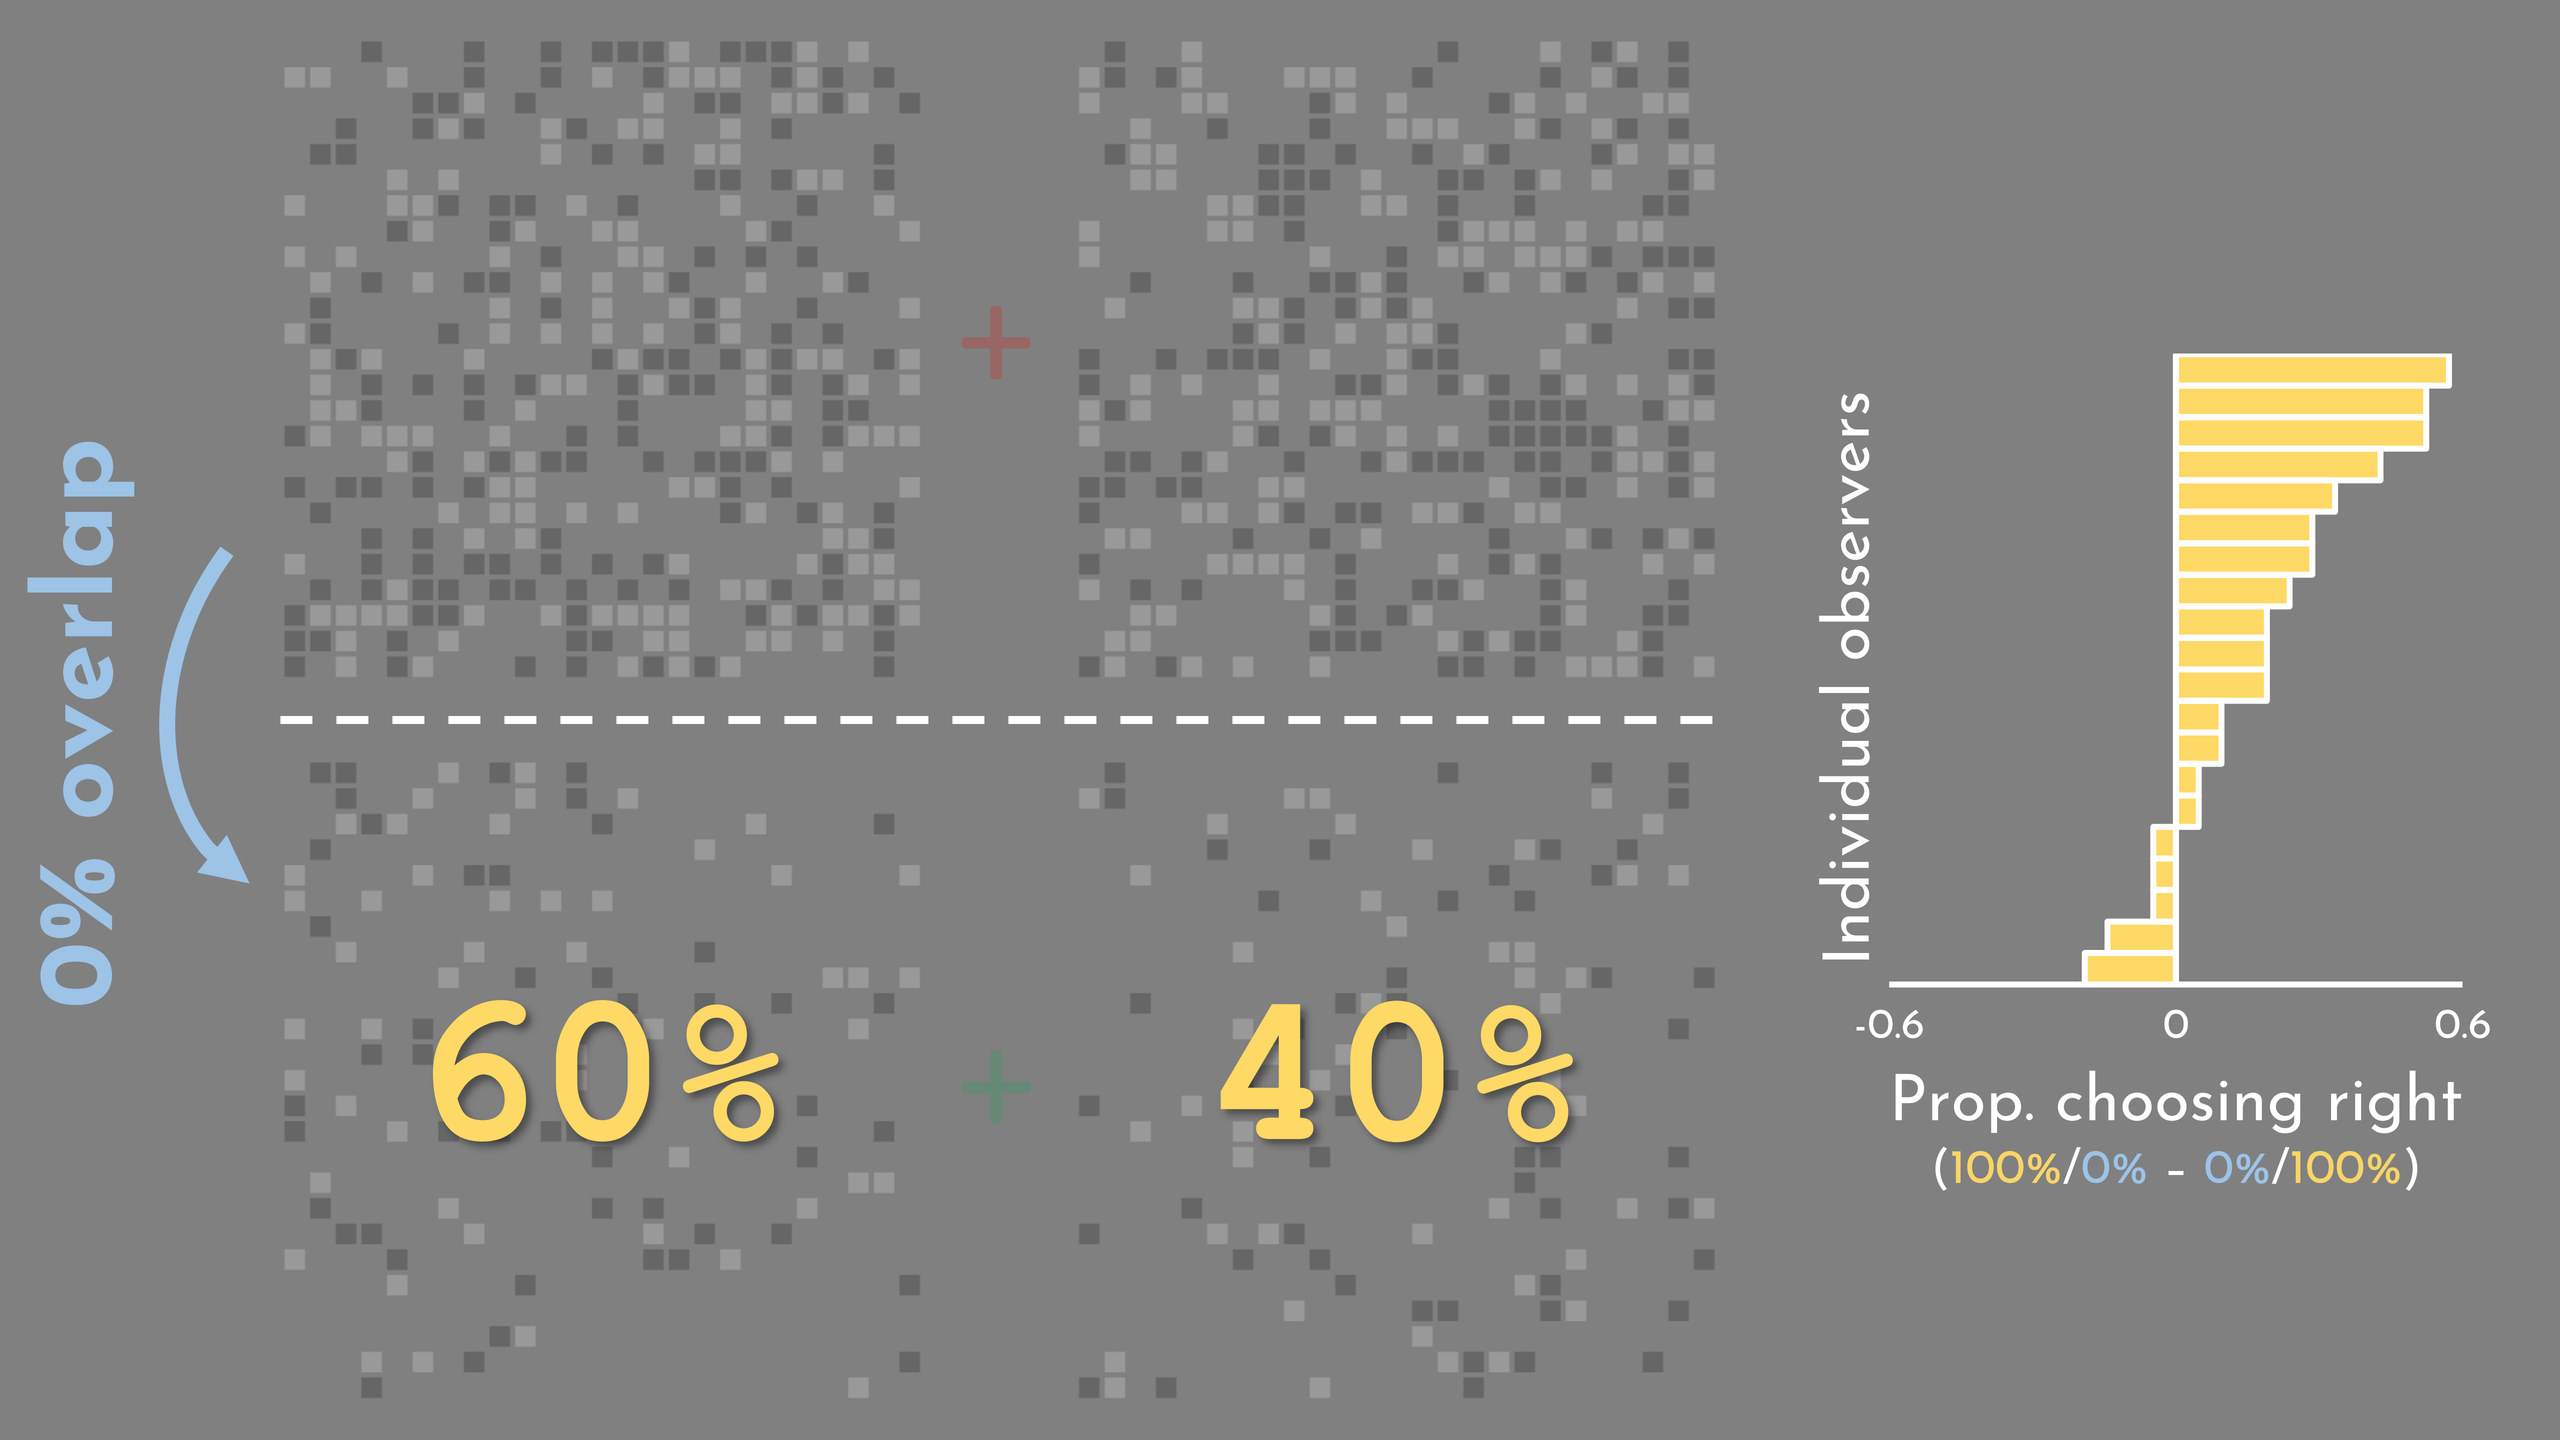
0% overlap
60%
40%
Individual observers
-0.6
0
0.6
Prop. choosing right
(100%/0% - 0%/100%)

What if *locations* are constant…

…but **colors** change?

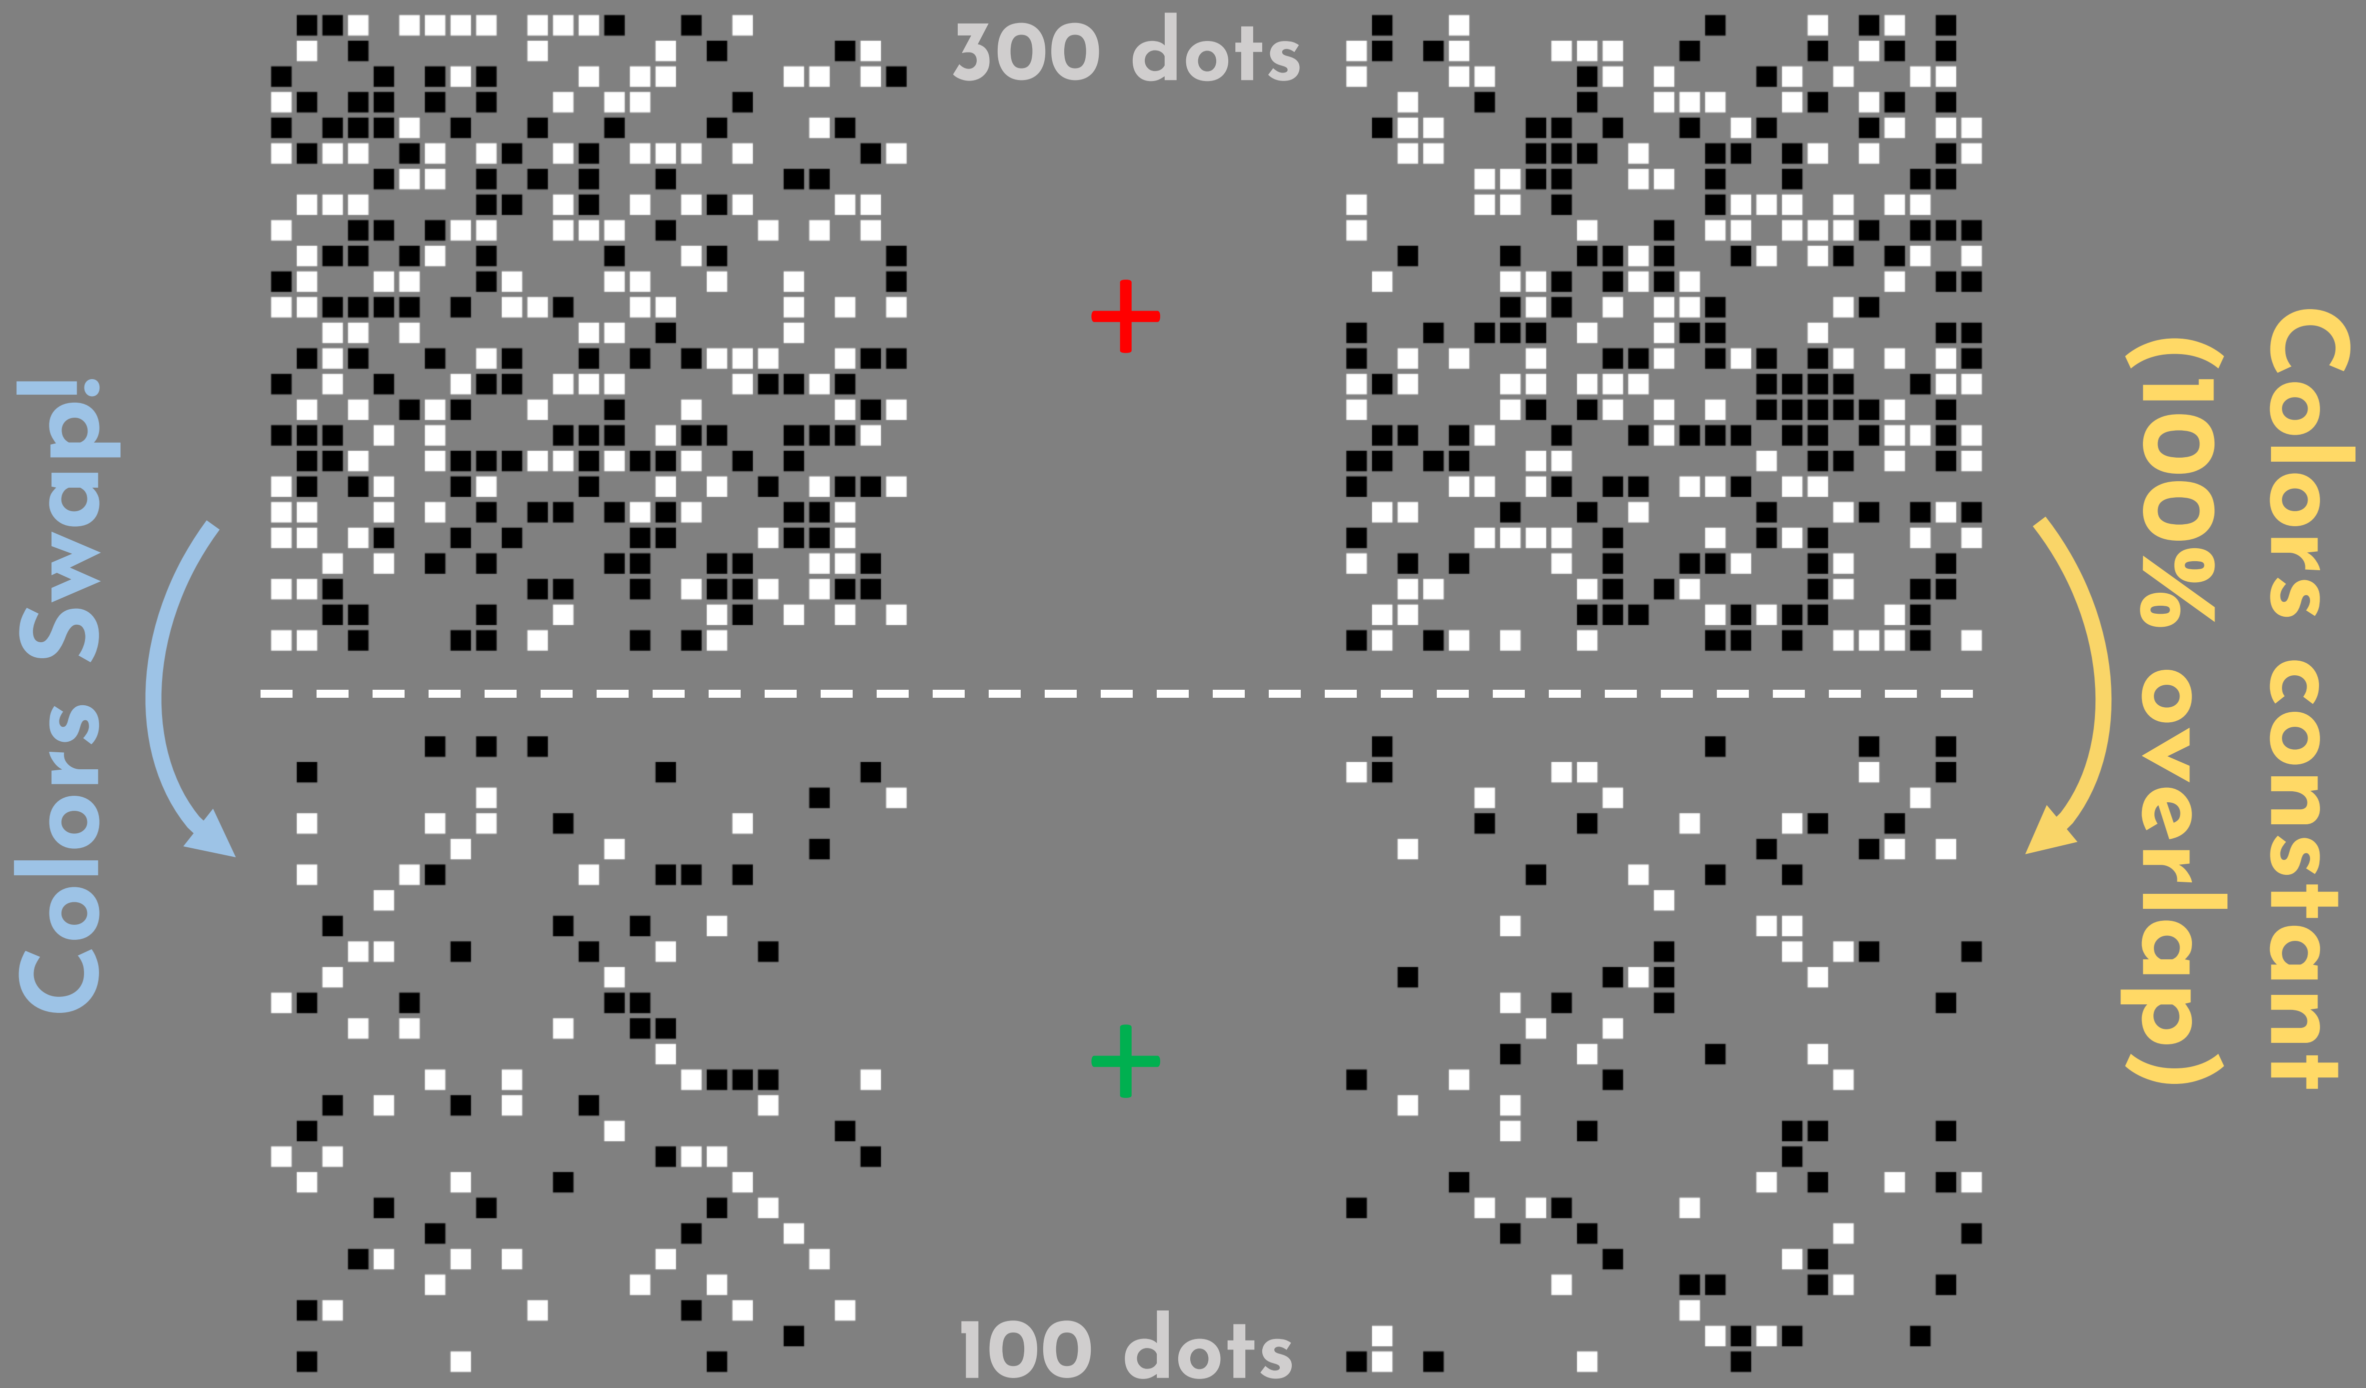
300 dots
100 dots
Colors Swap!
Colors constant
(100% overlap)

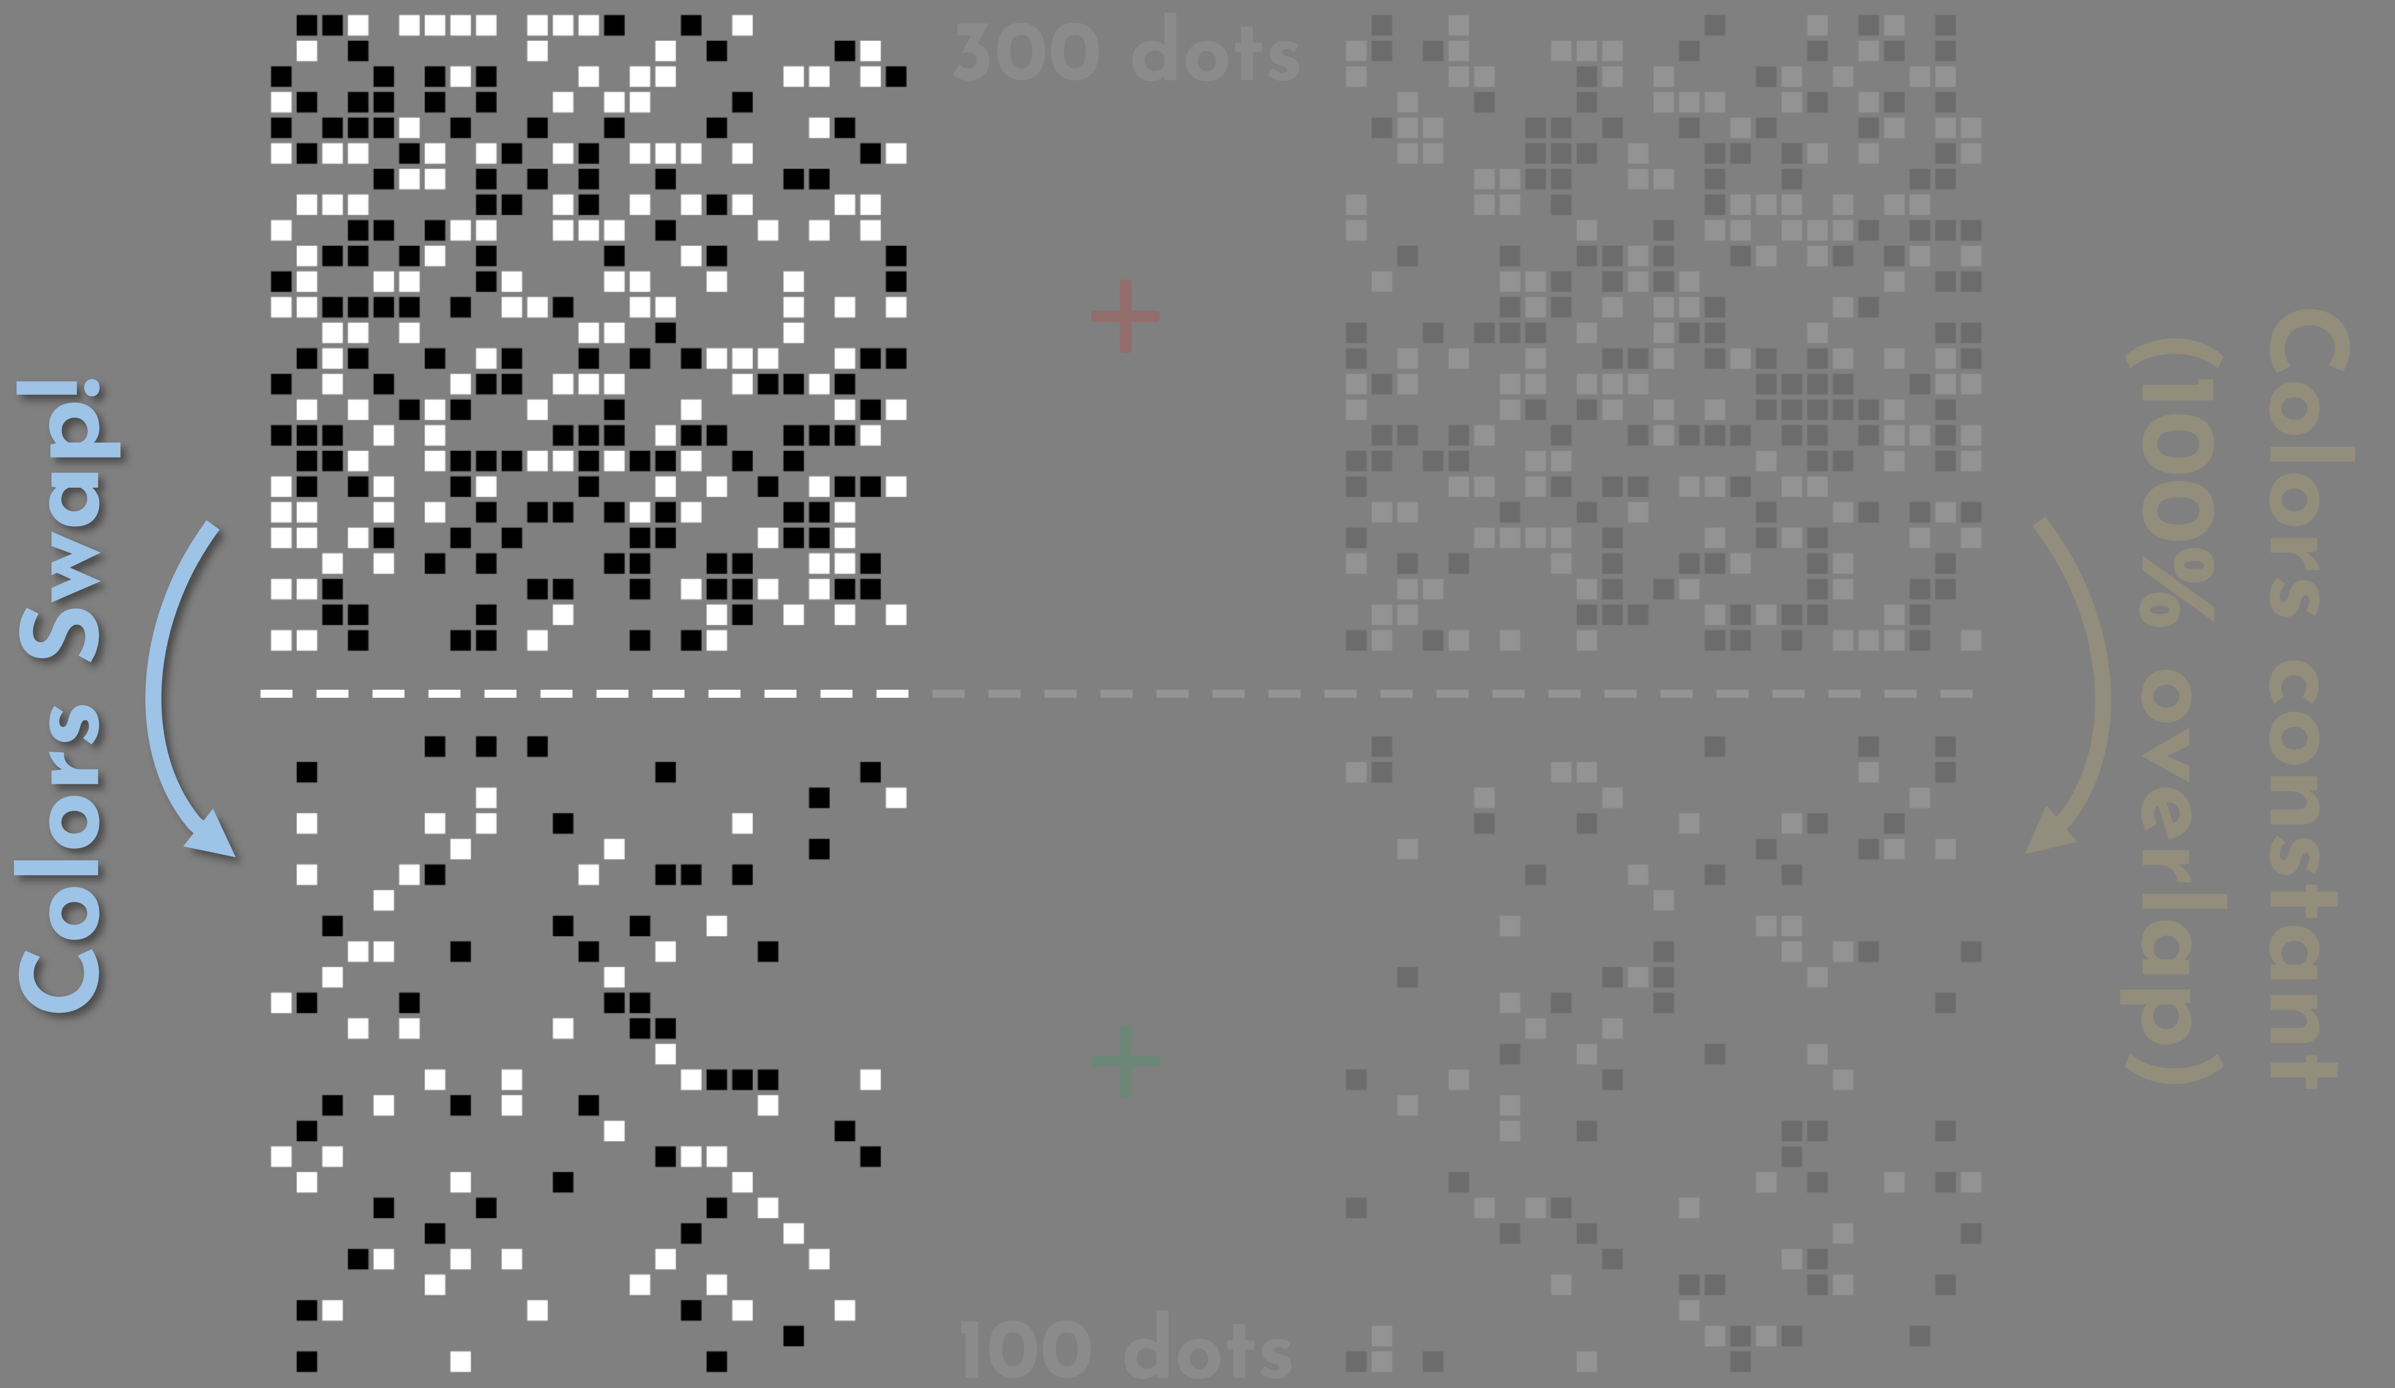
300 dots
100 dots
Colors Swap!
Colors constant (100% overlap)

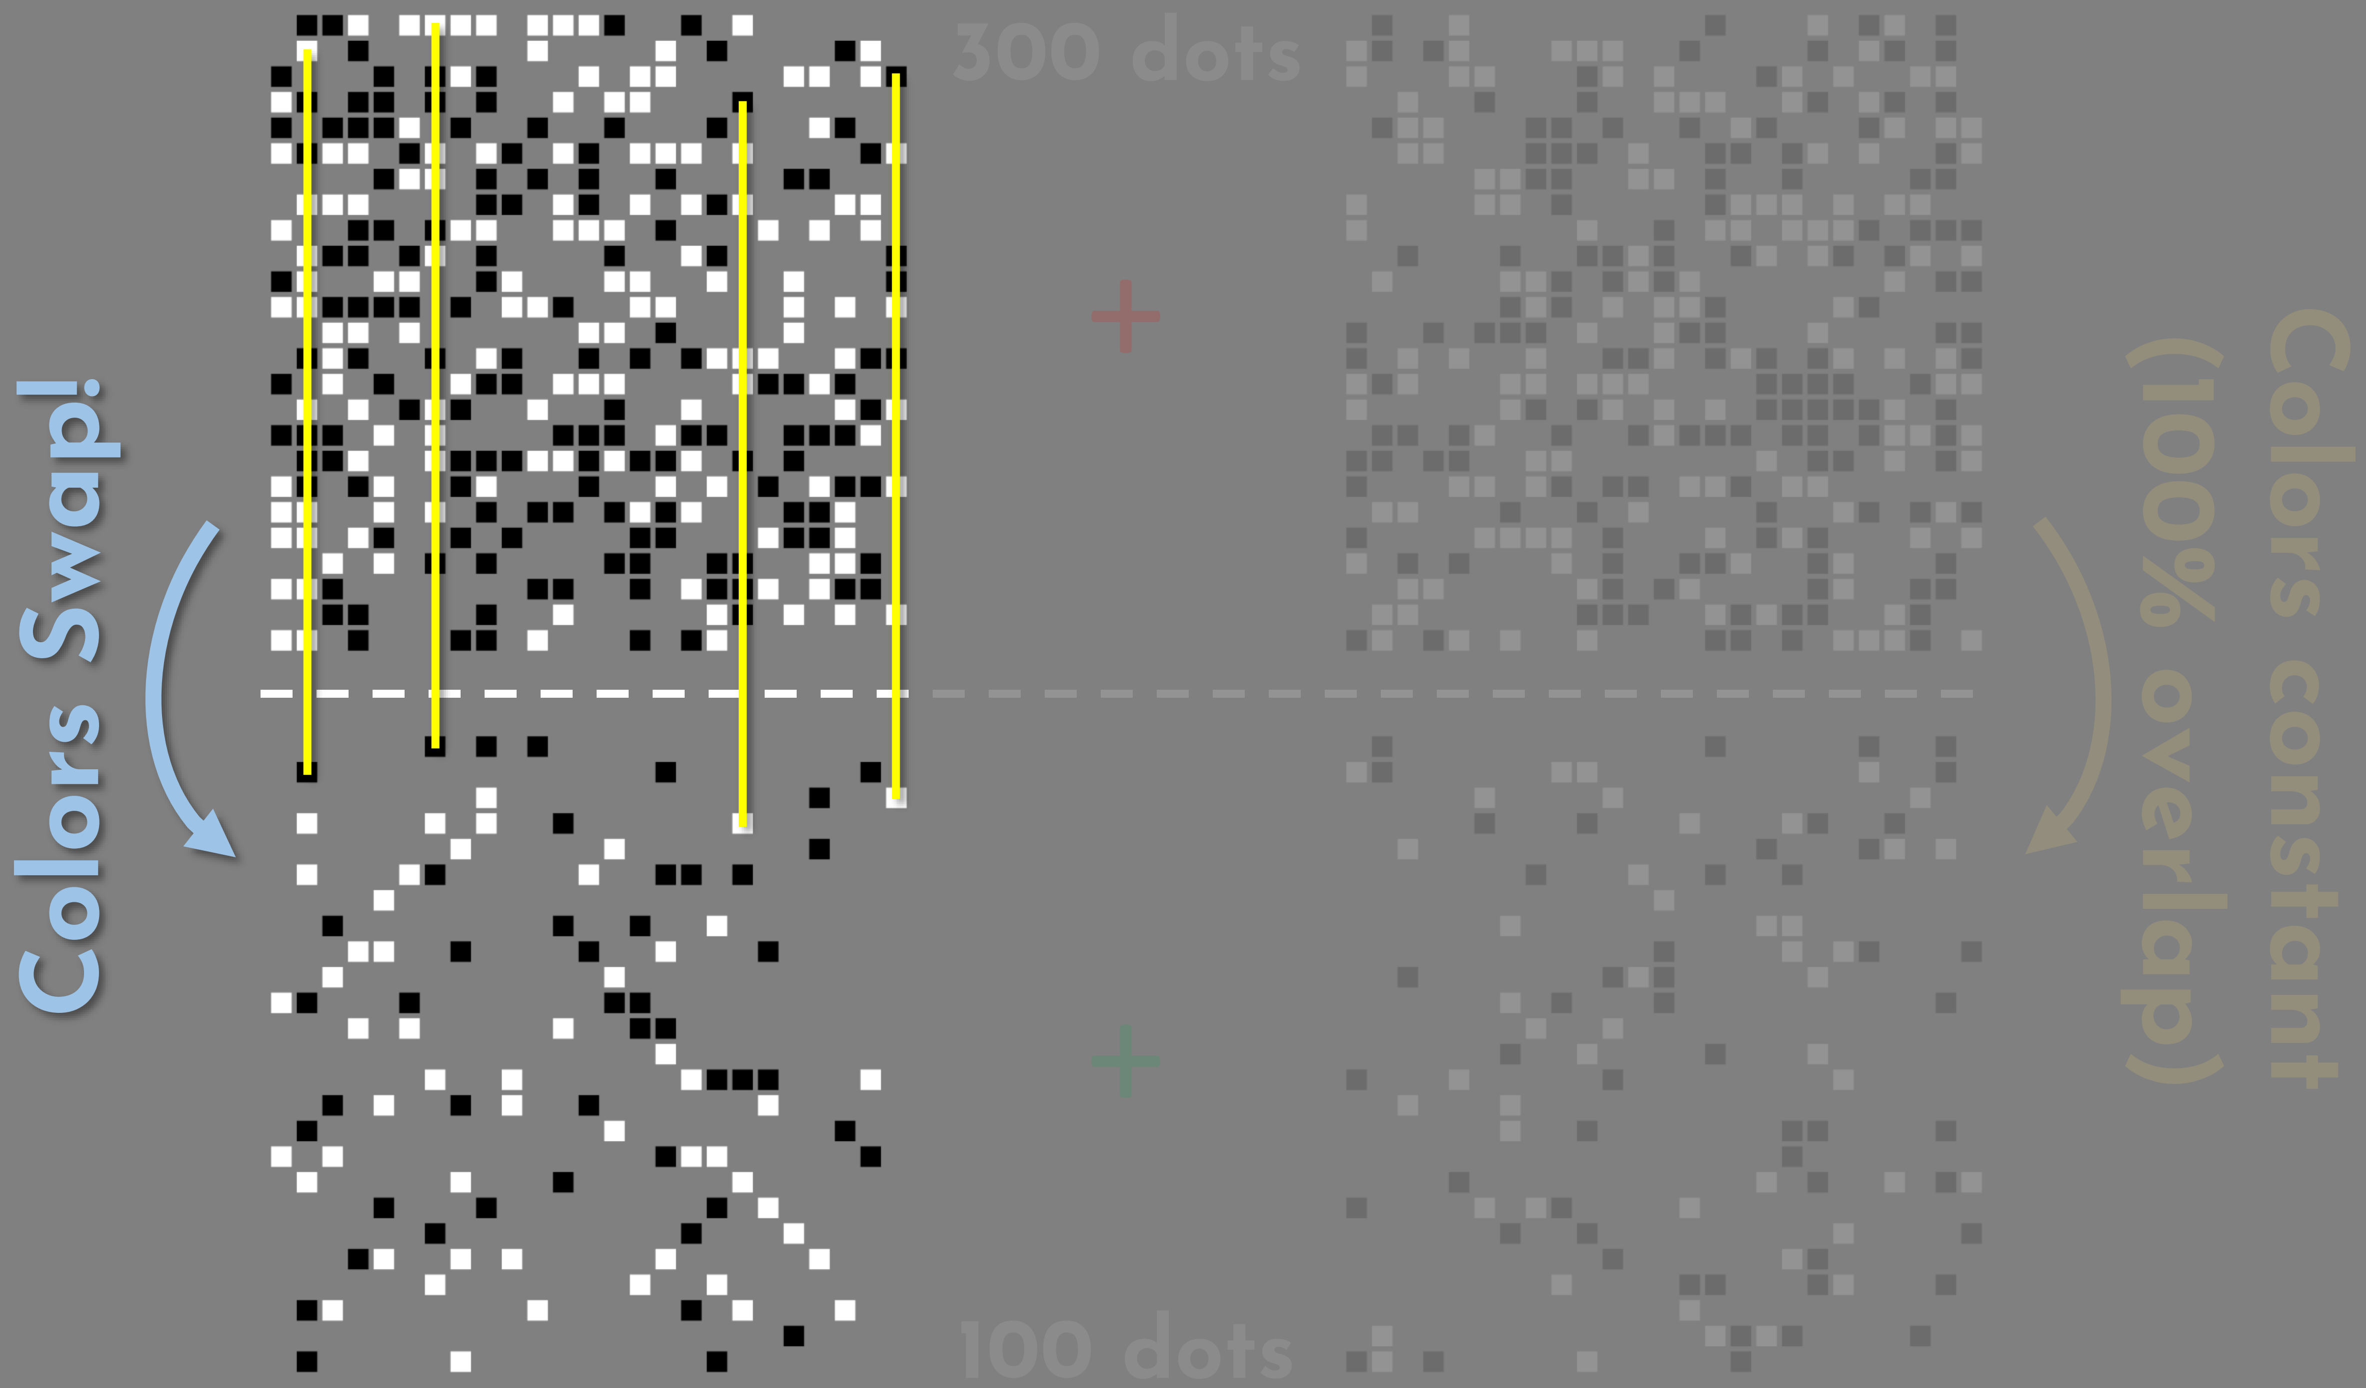

Colors Swap!
300 dots
100 dots
Colors constant (100% overlap)

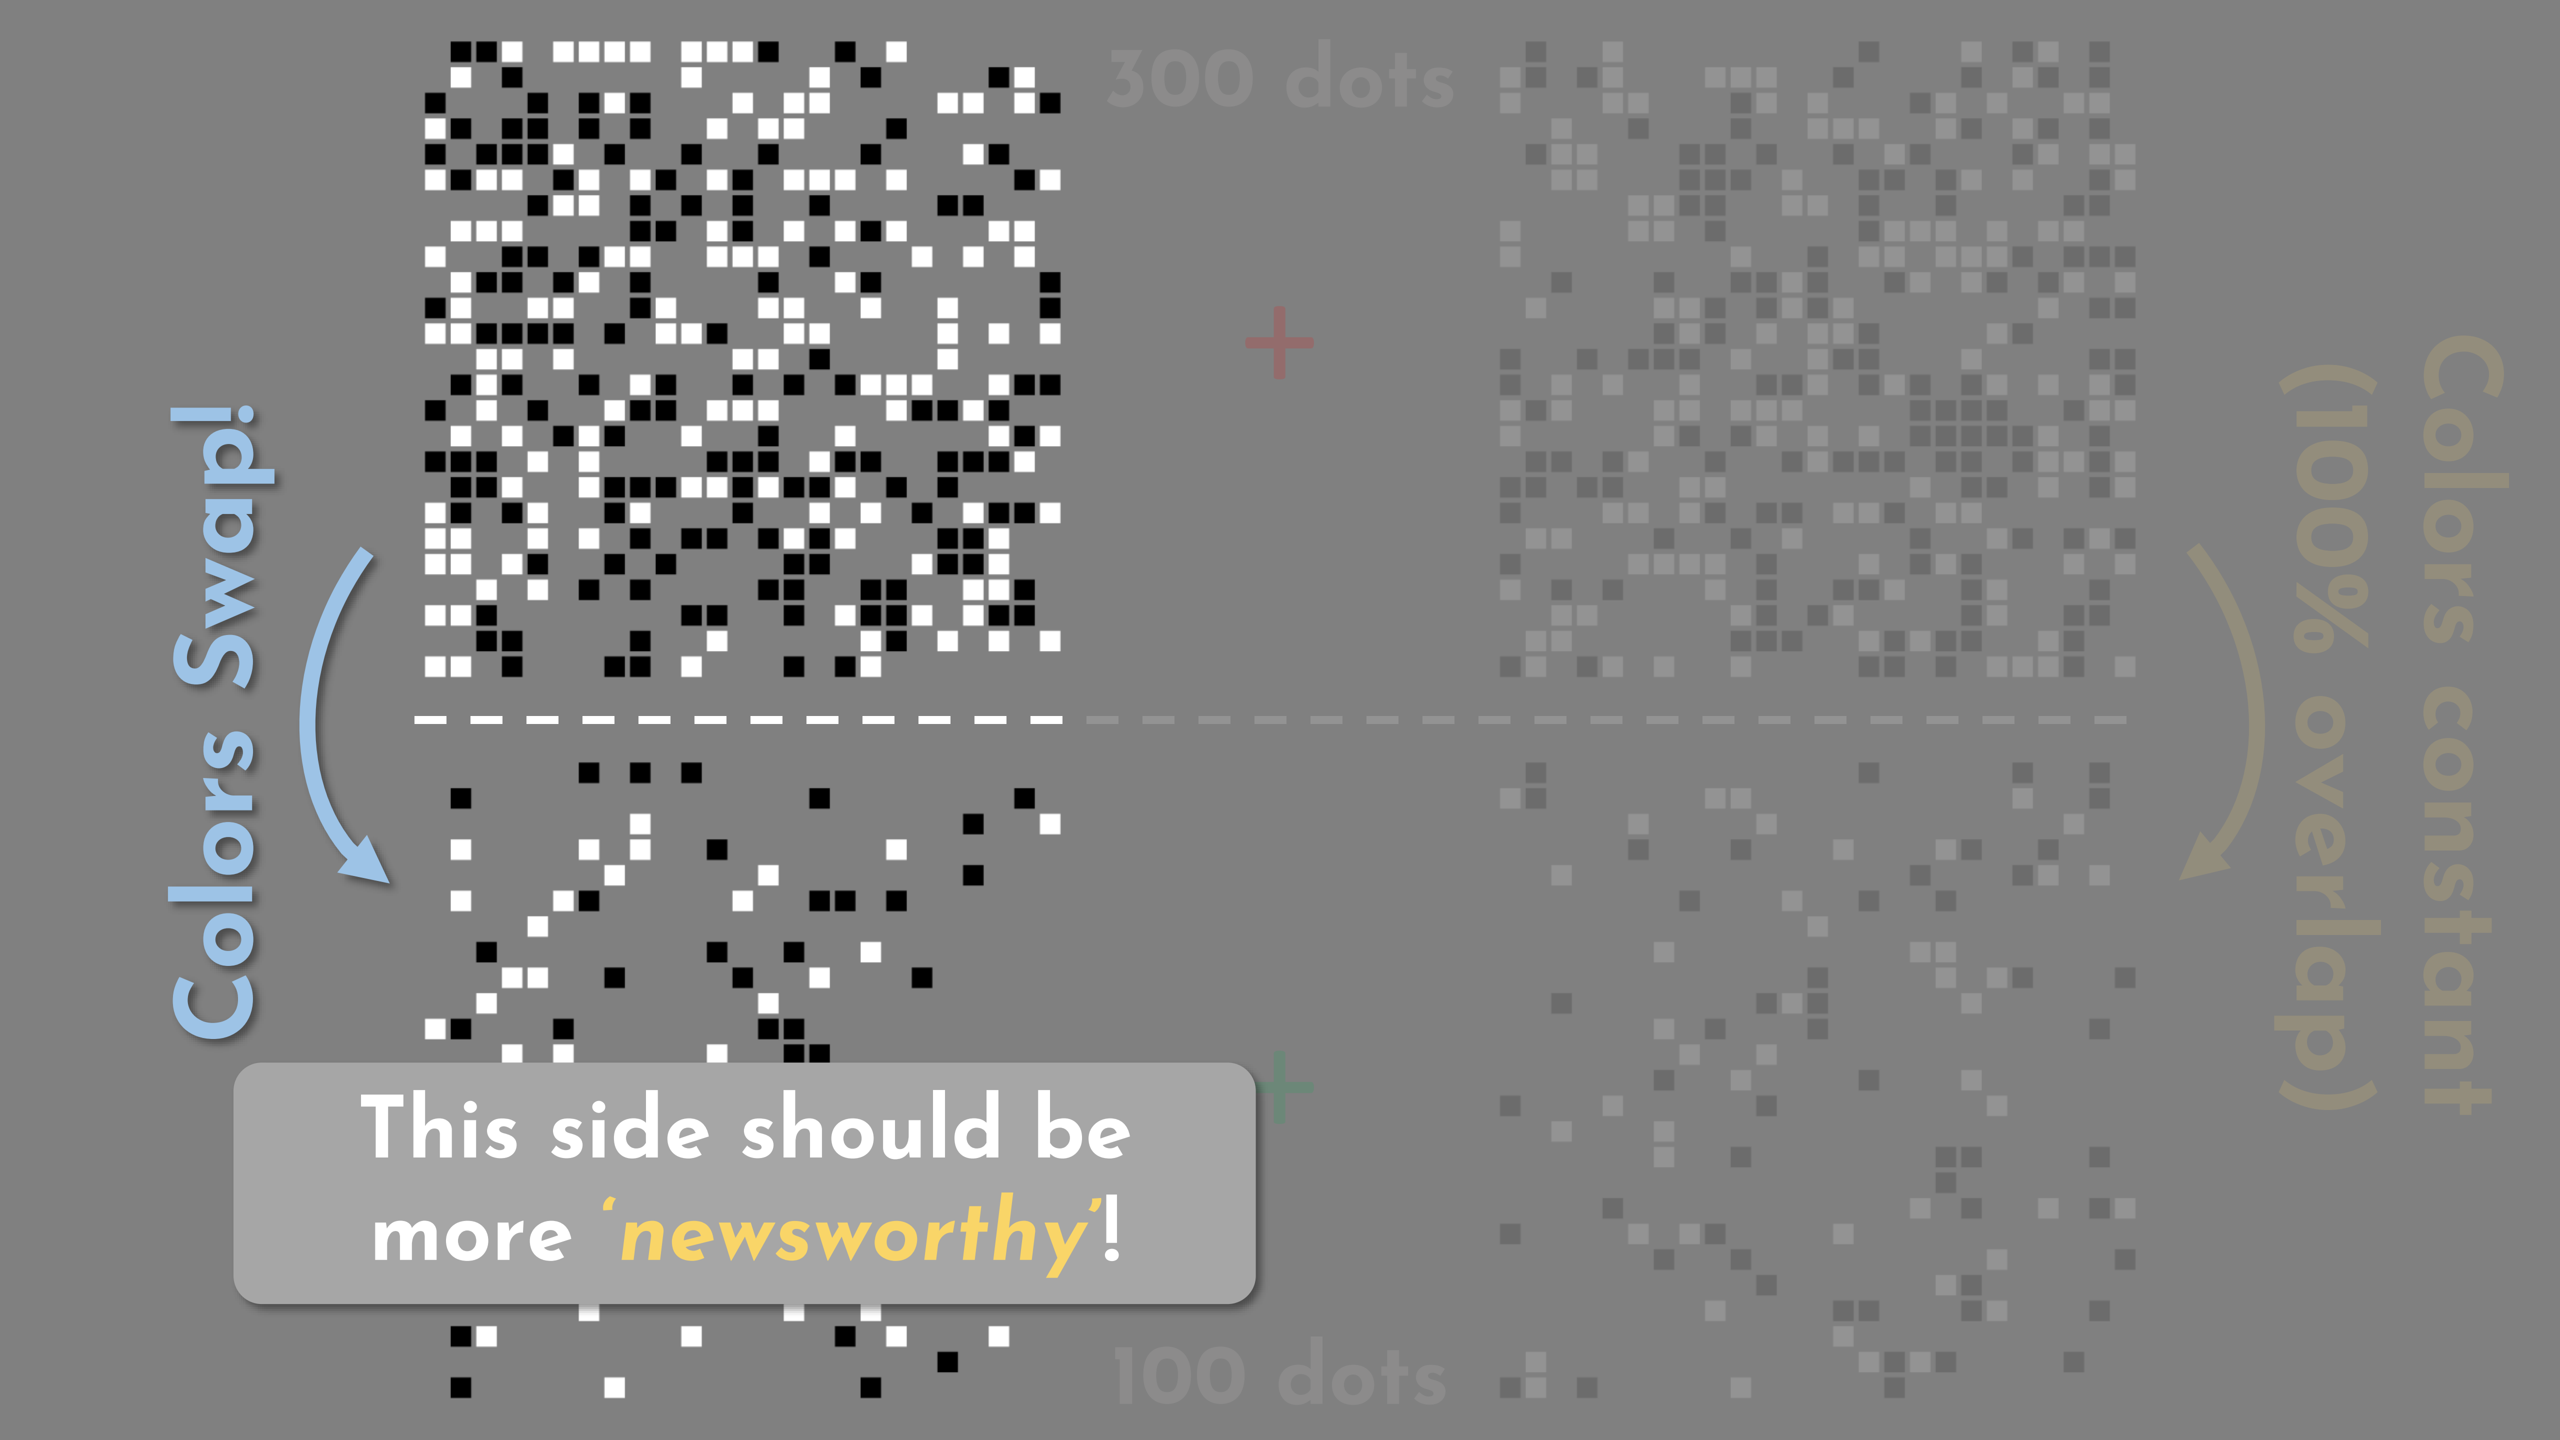
300 dots
Colors Swap!
Colors constant (100% overlap)
This side should be more ‘newsworthy’!
100 dots

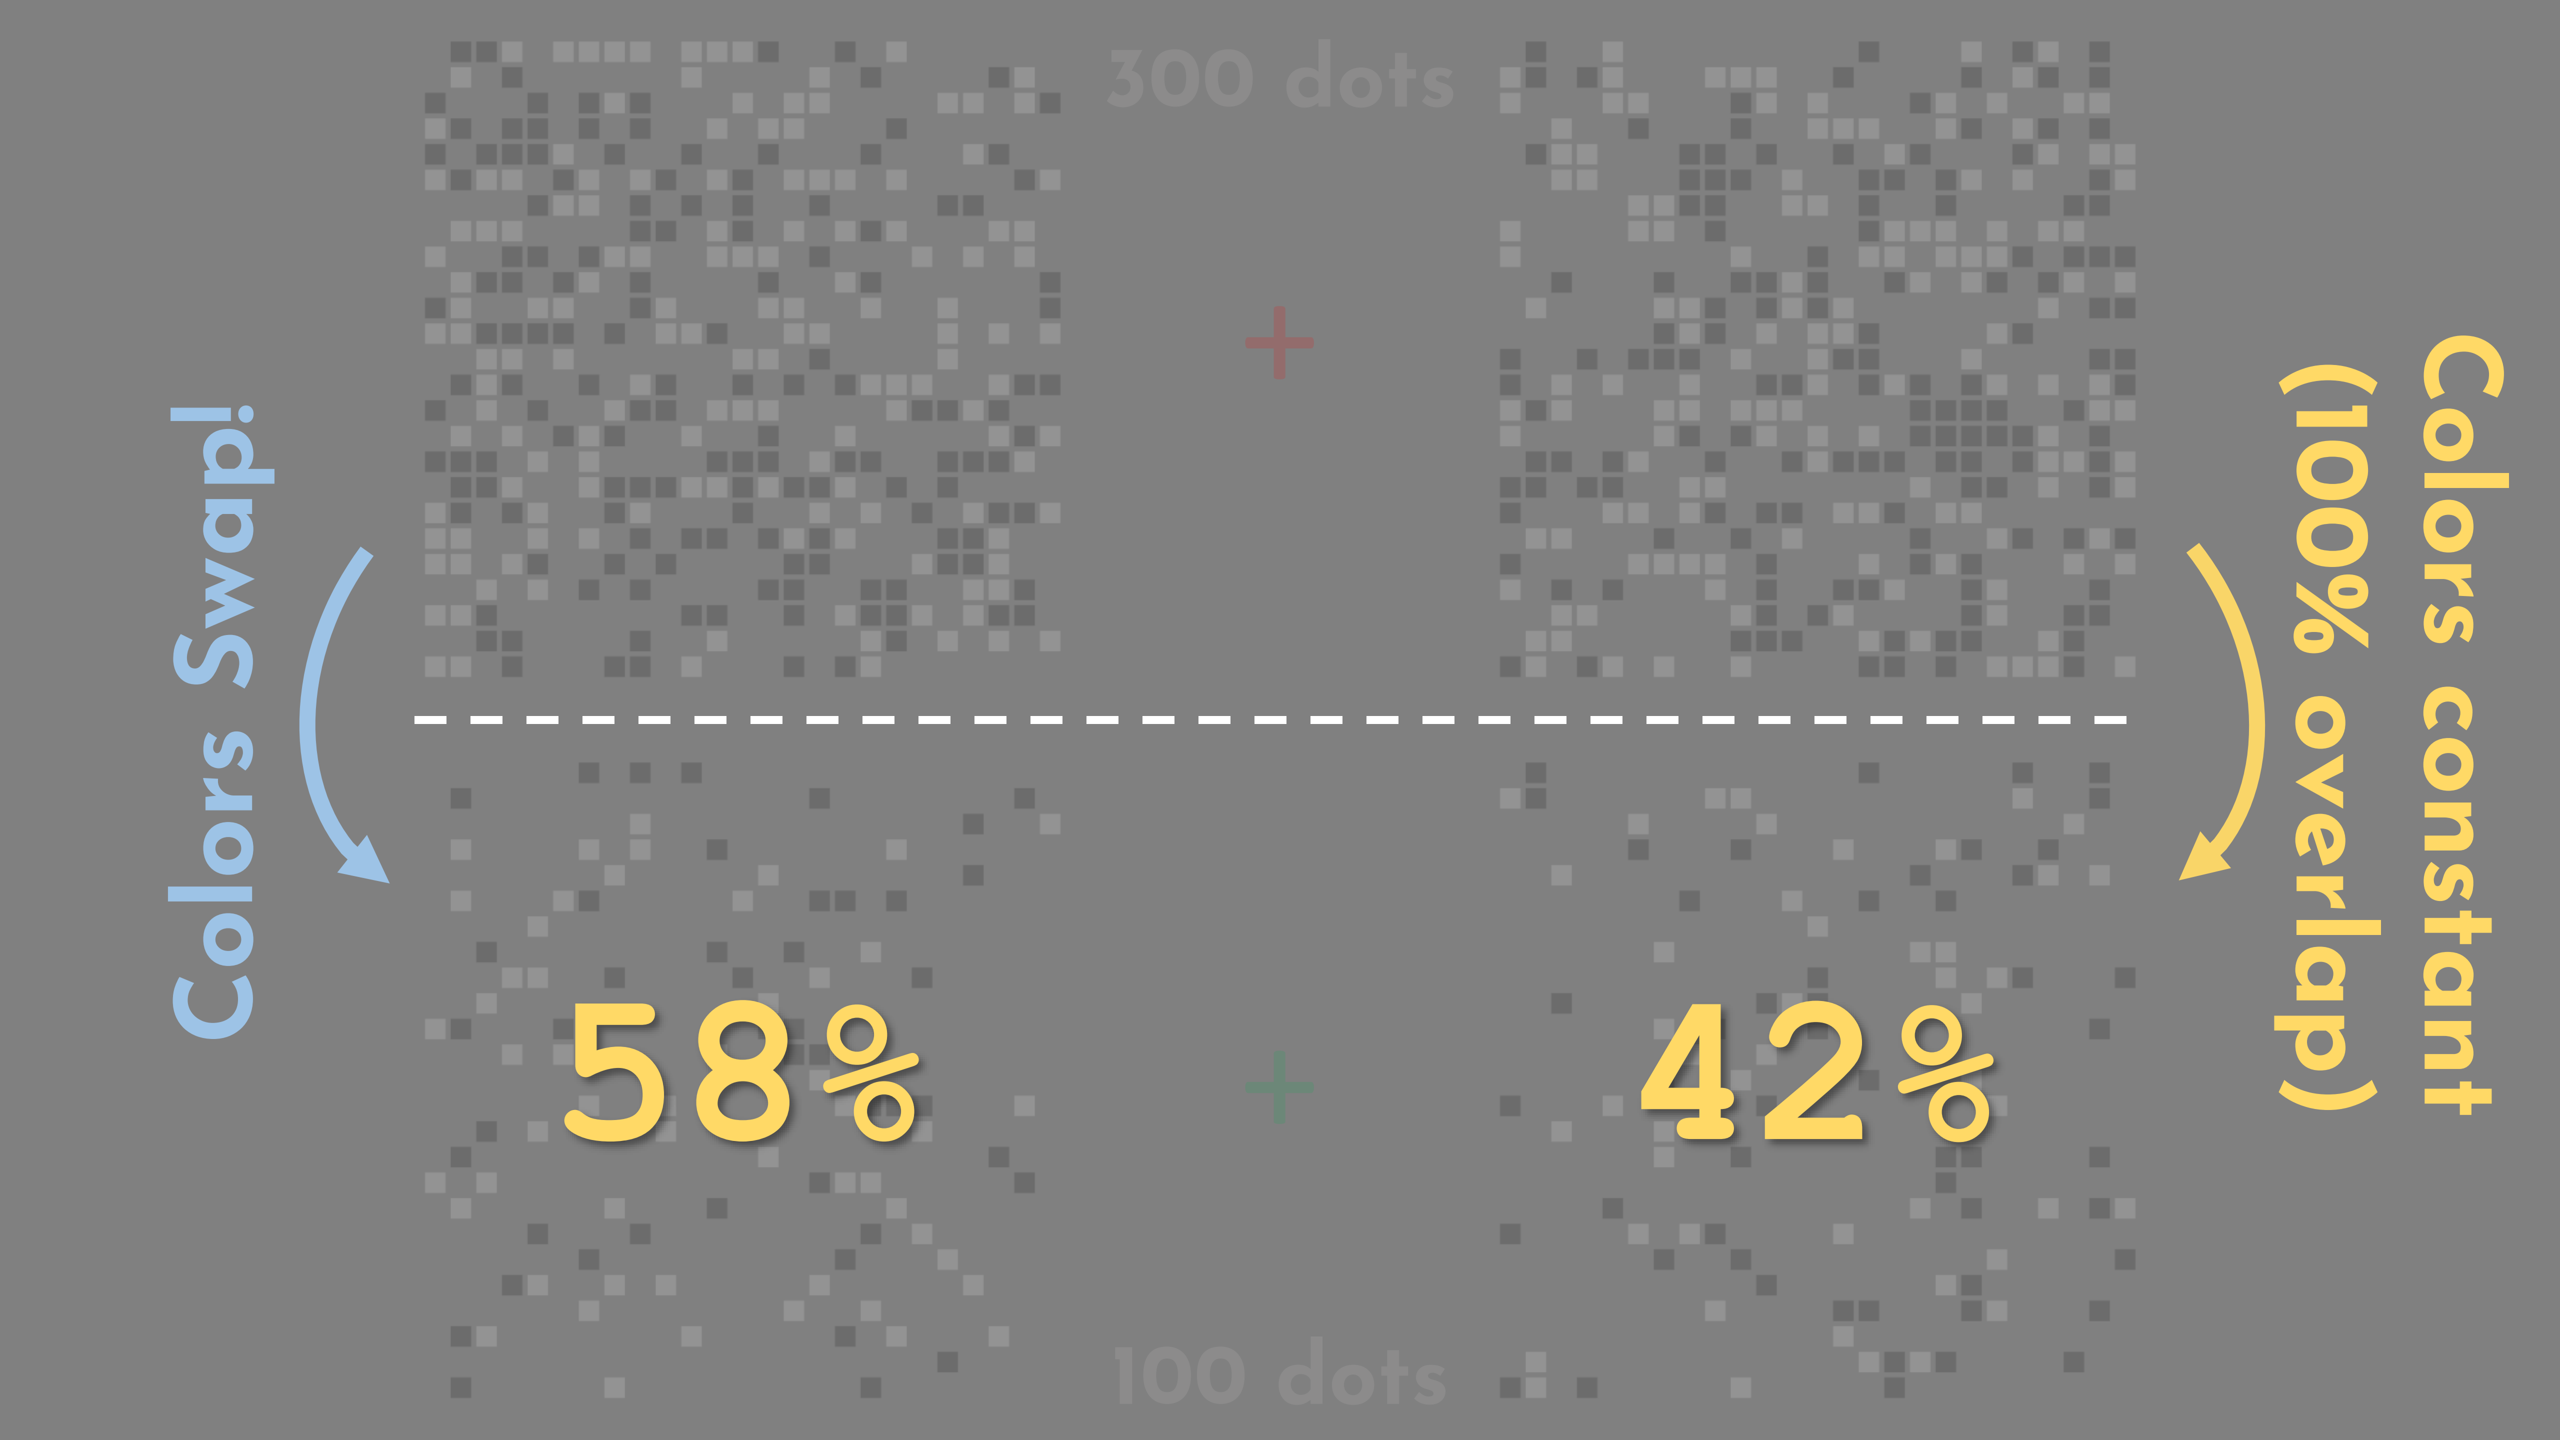
300 dots
Colors Swap!
Colors constant
(100% overlap)
58%
42%
100 dots

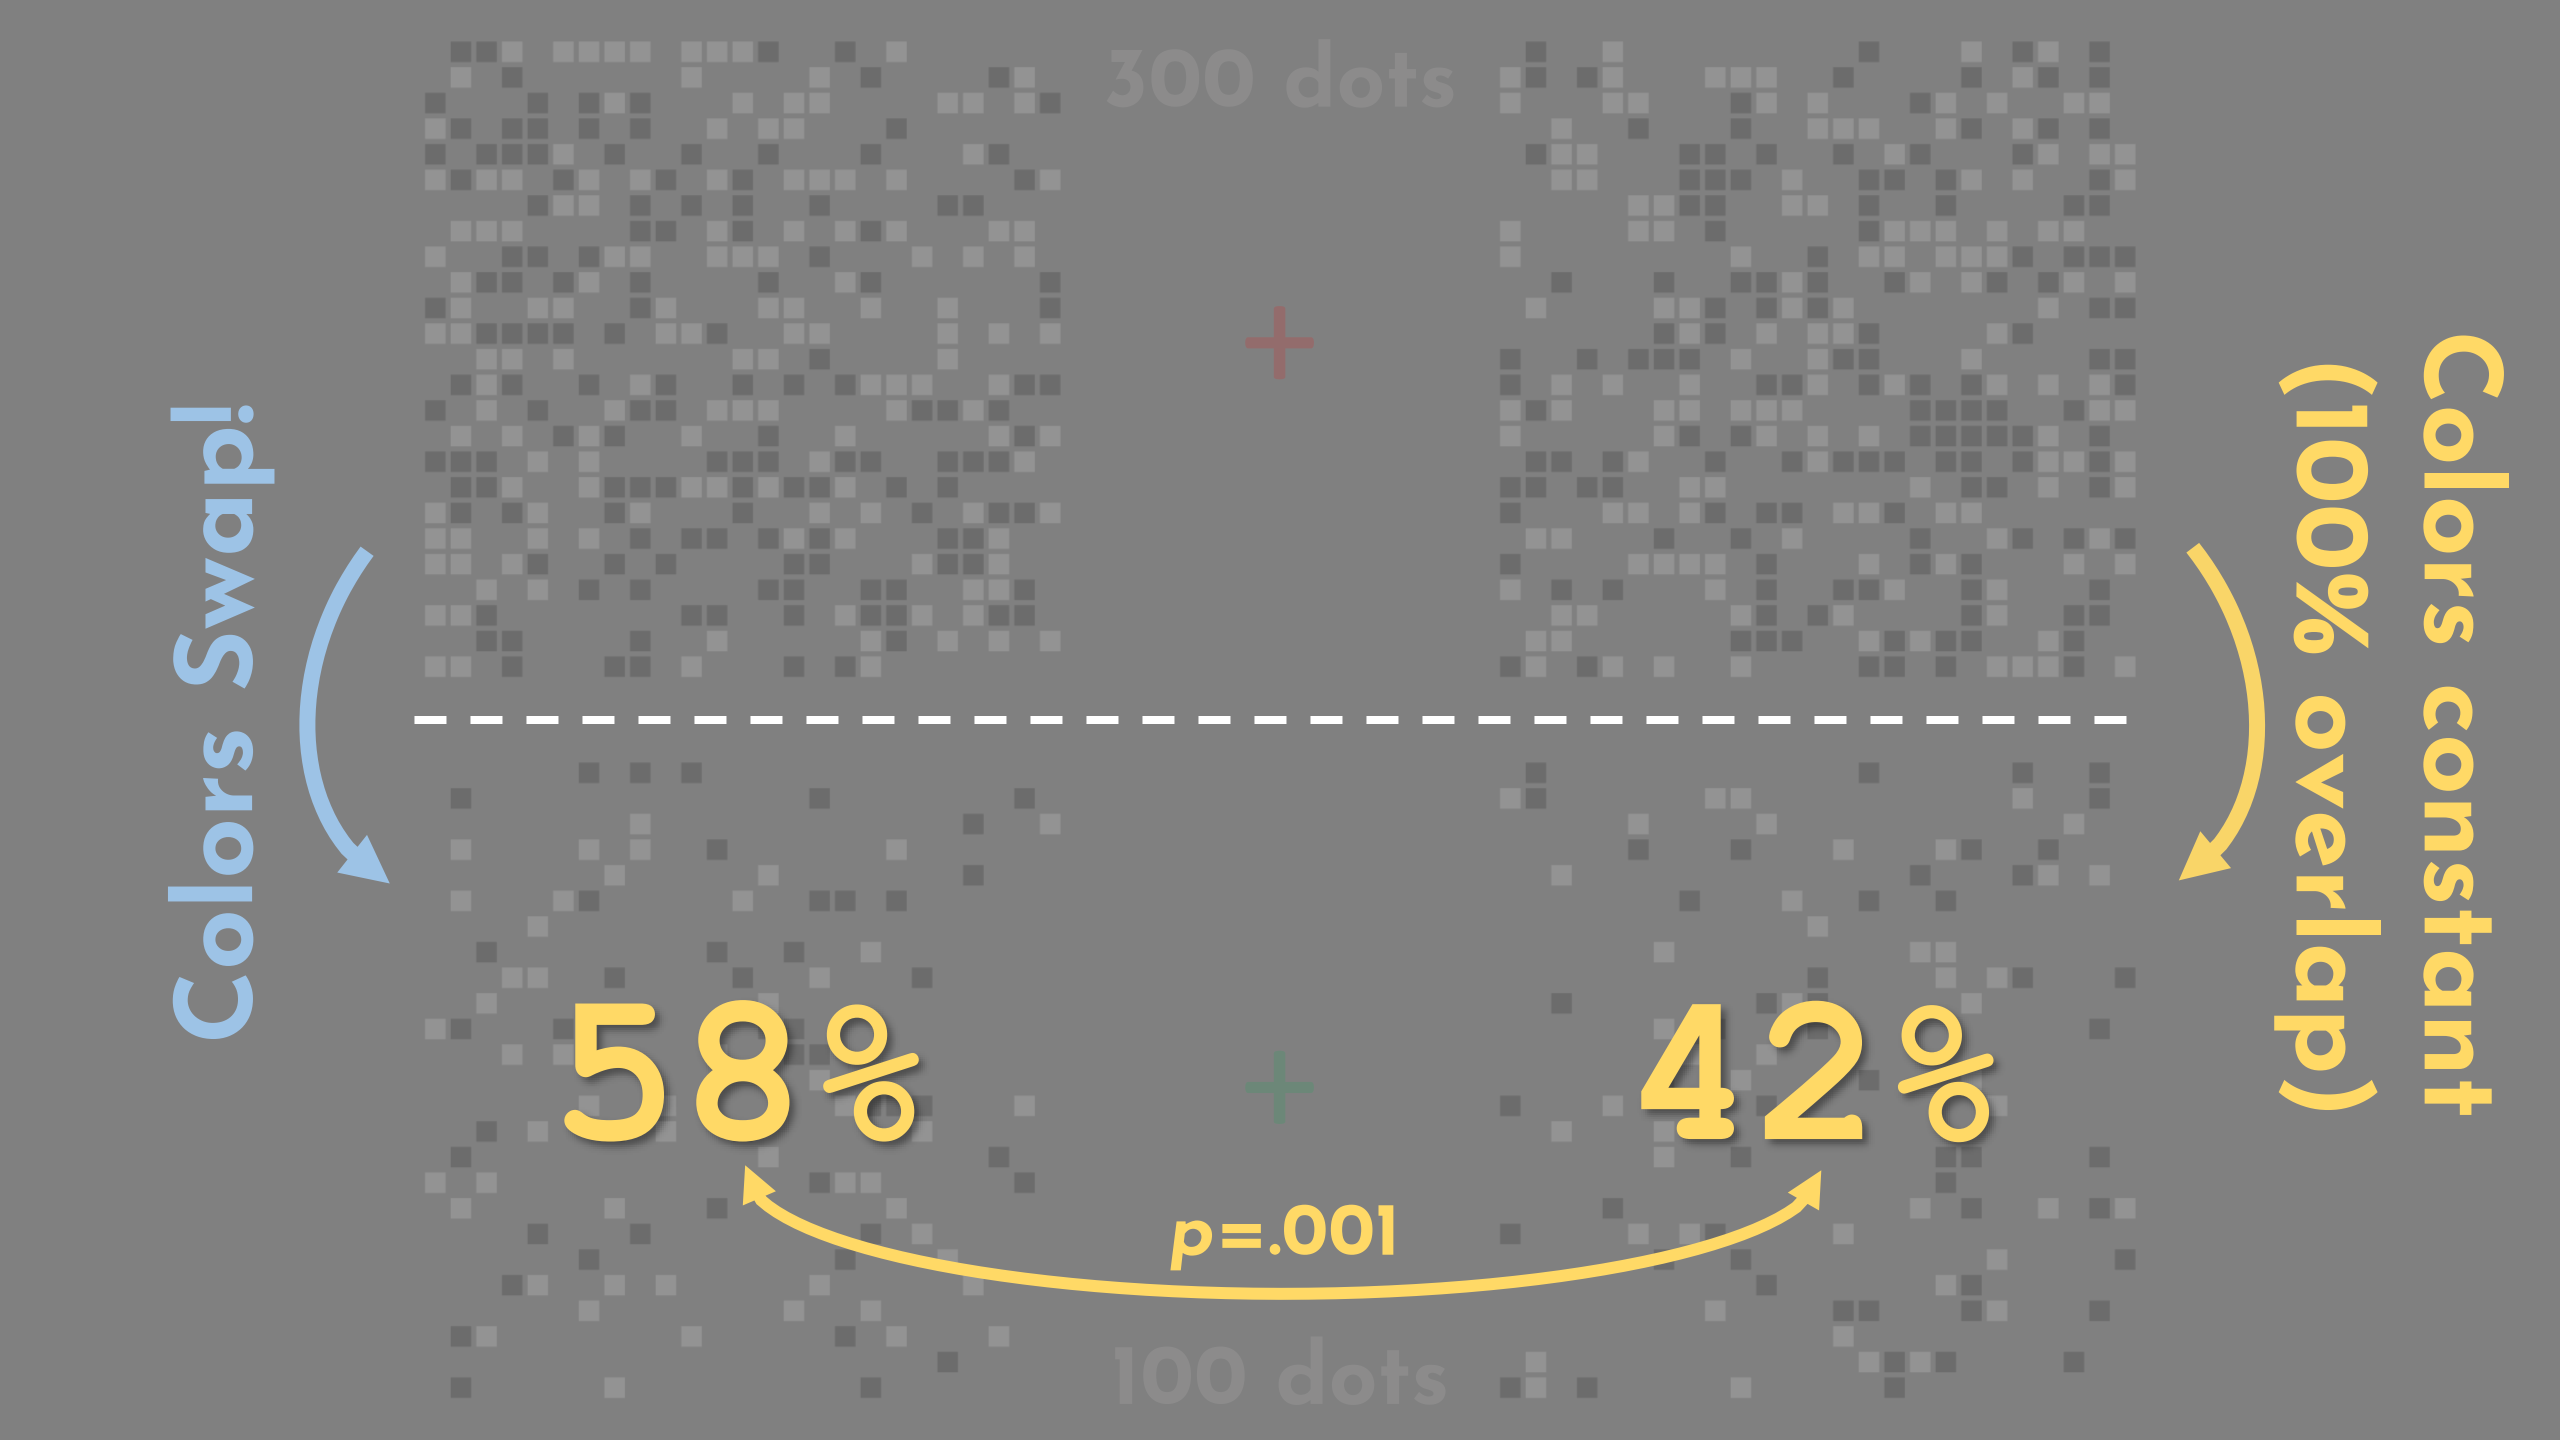
300 dots
100 dots
Colors Swap!
Colors constant (100% overlap)
58%
42%
p=.001

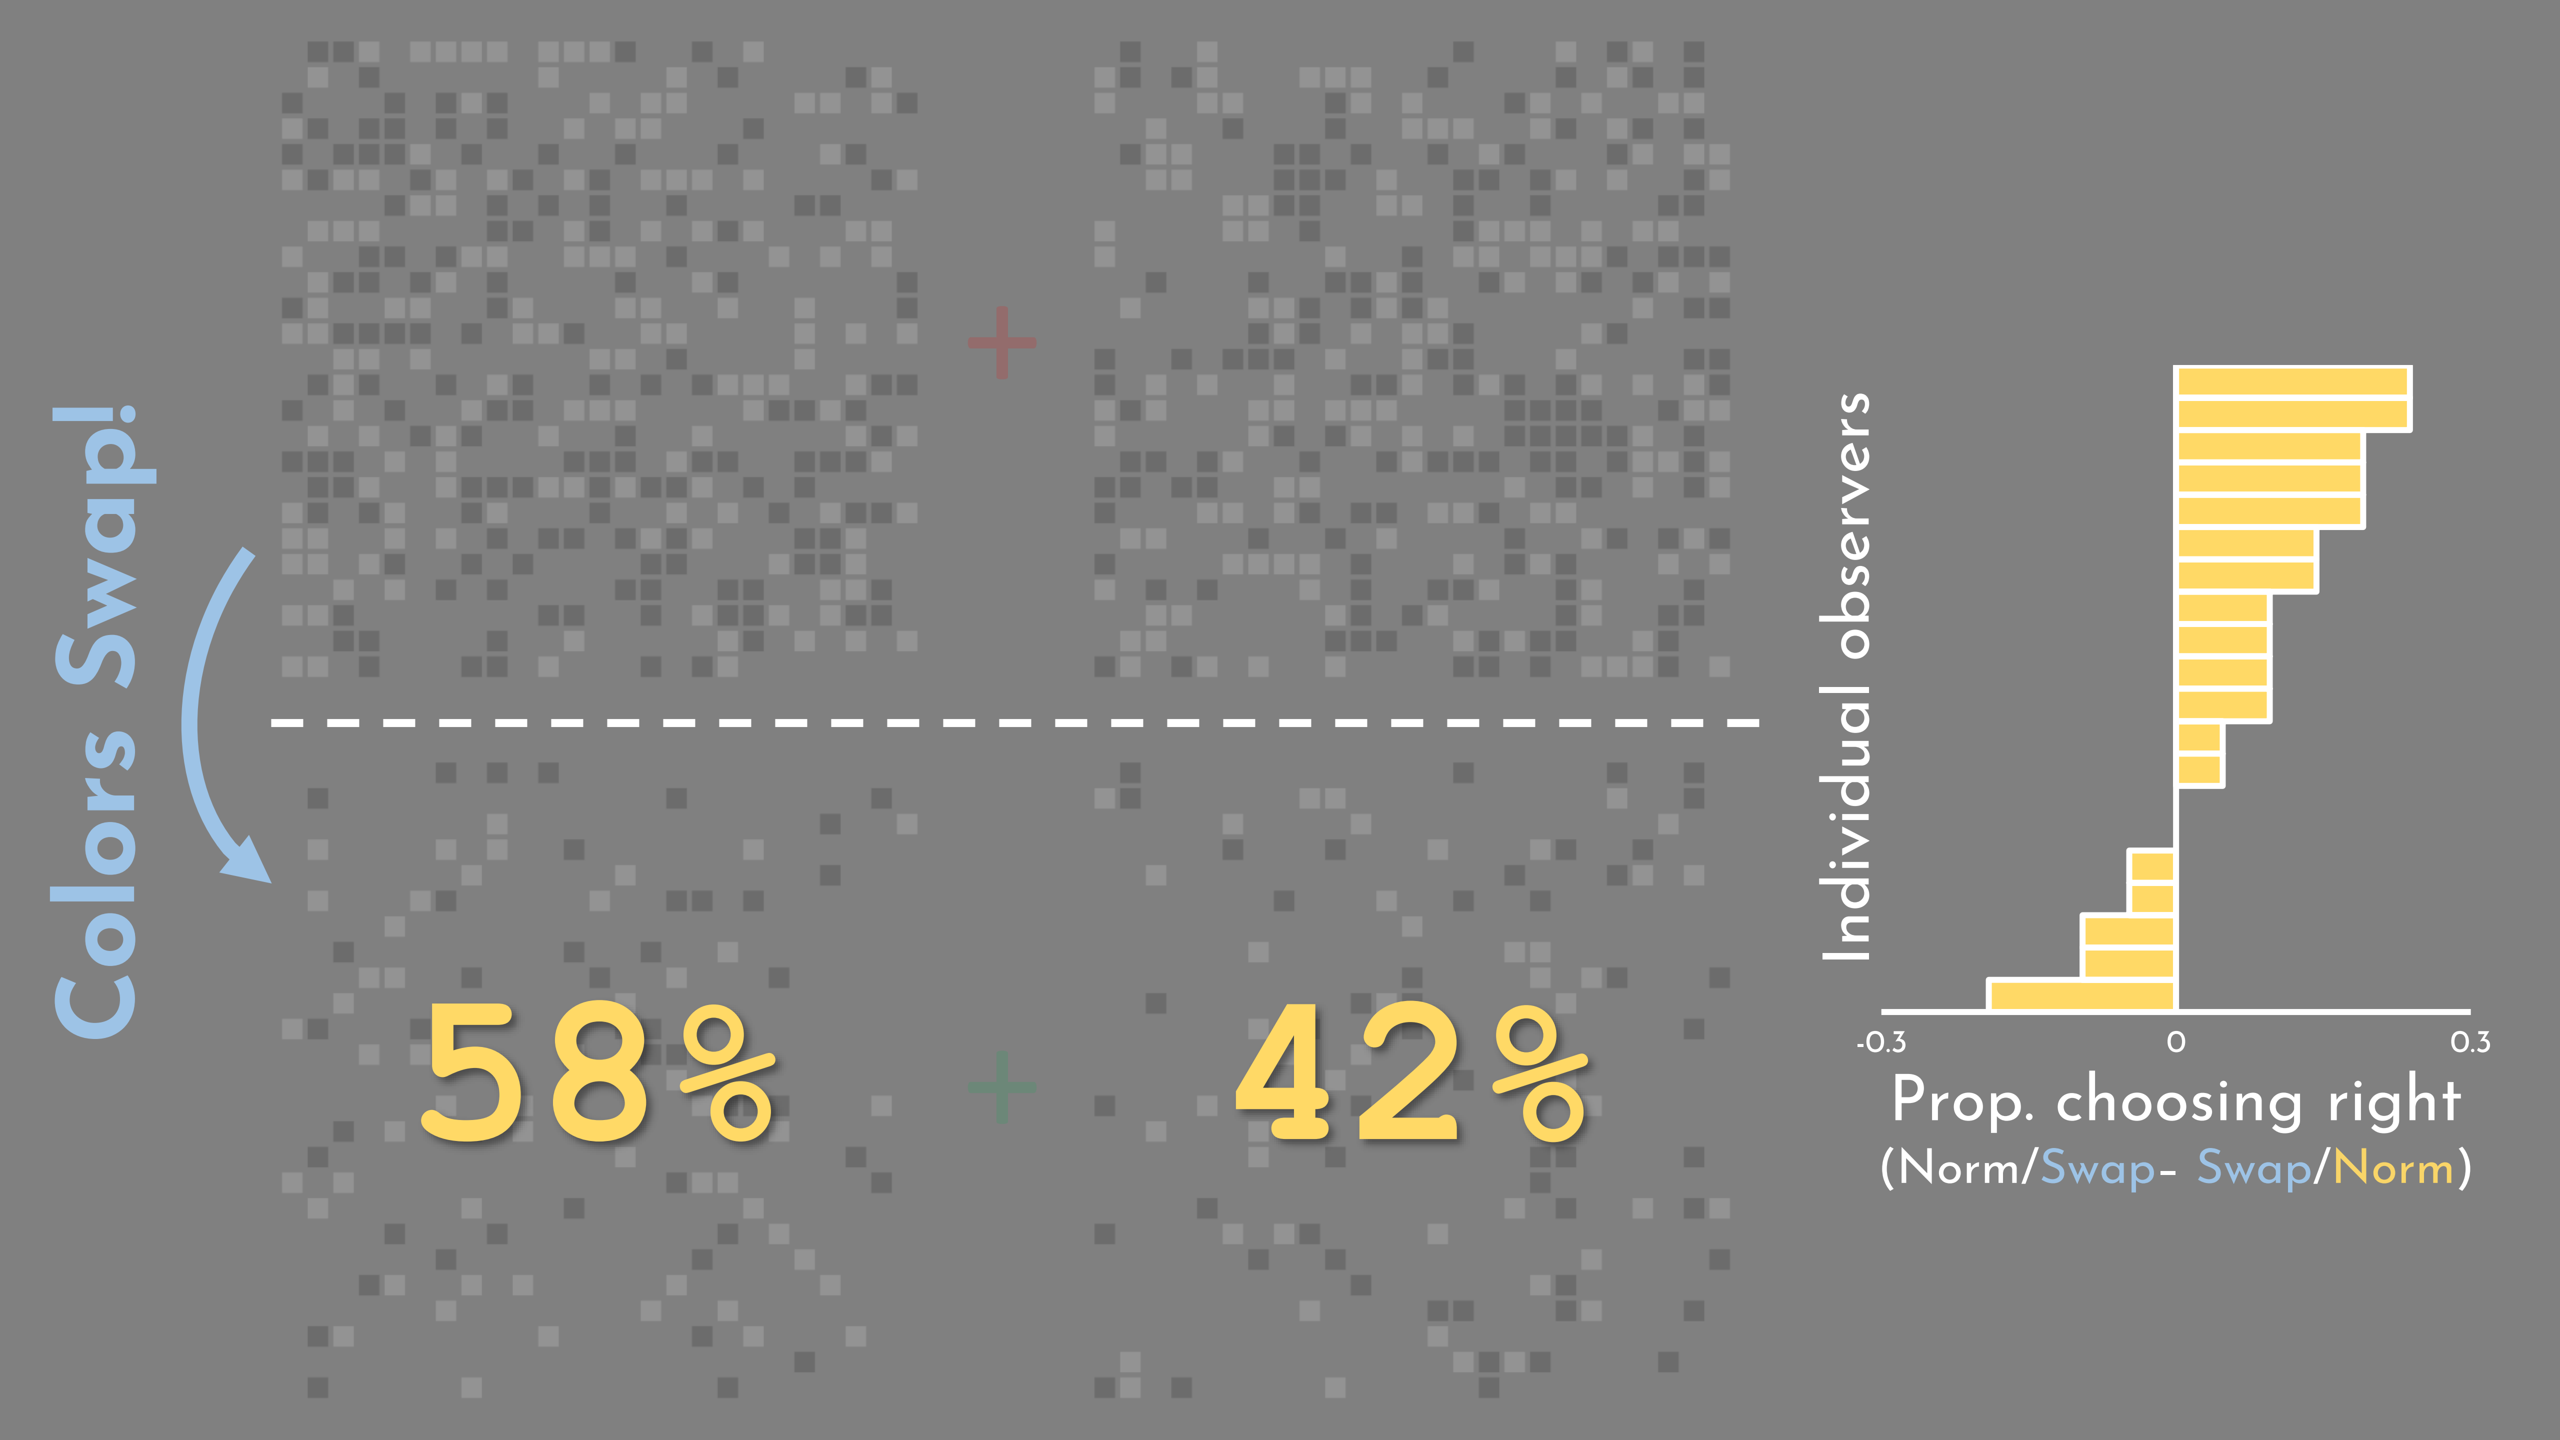
Colors Swap!
58%
42%
Individual observers
-0.3
0
0.3
Prop. choosing right
(Norm/Swap- Swap/Norm)

The visual system is sensitive to the '*newsworthiness*' of information!

What about reverse adaptation?

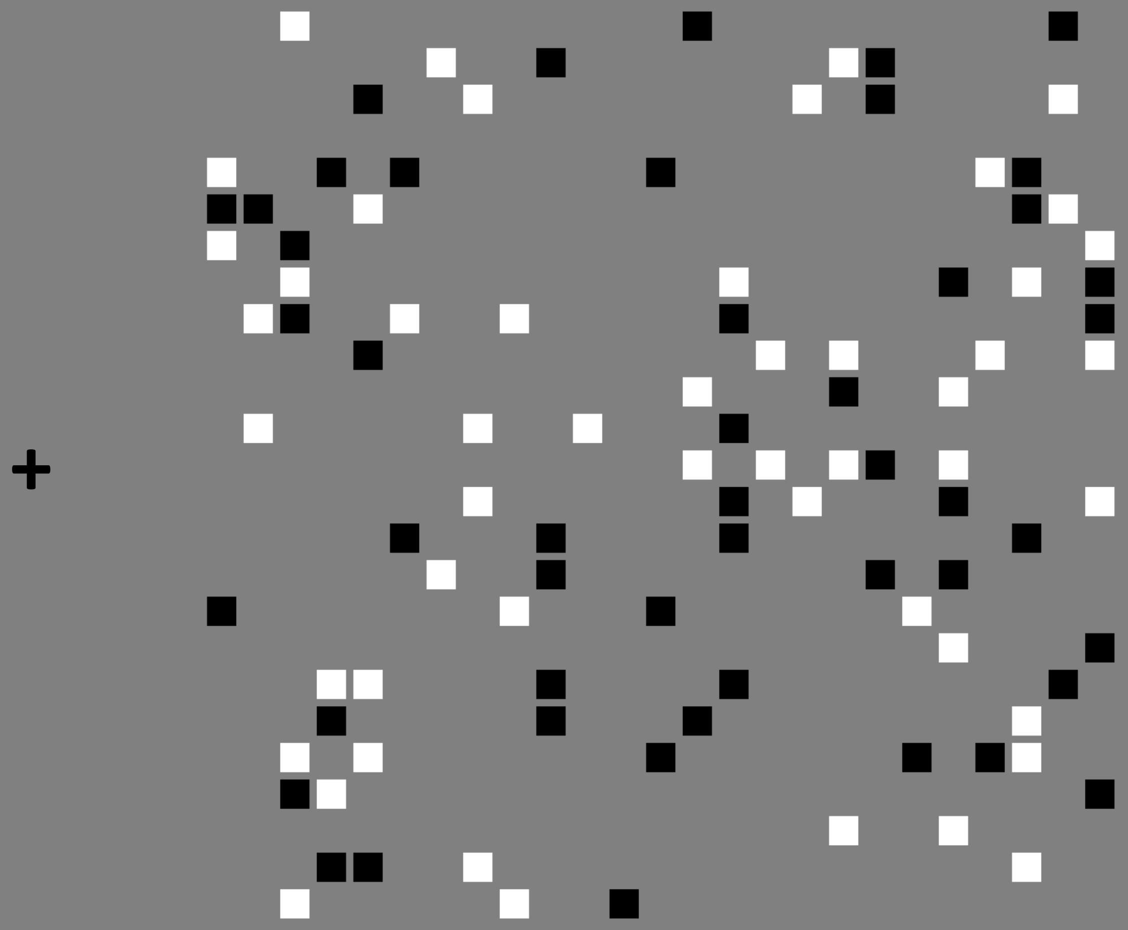

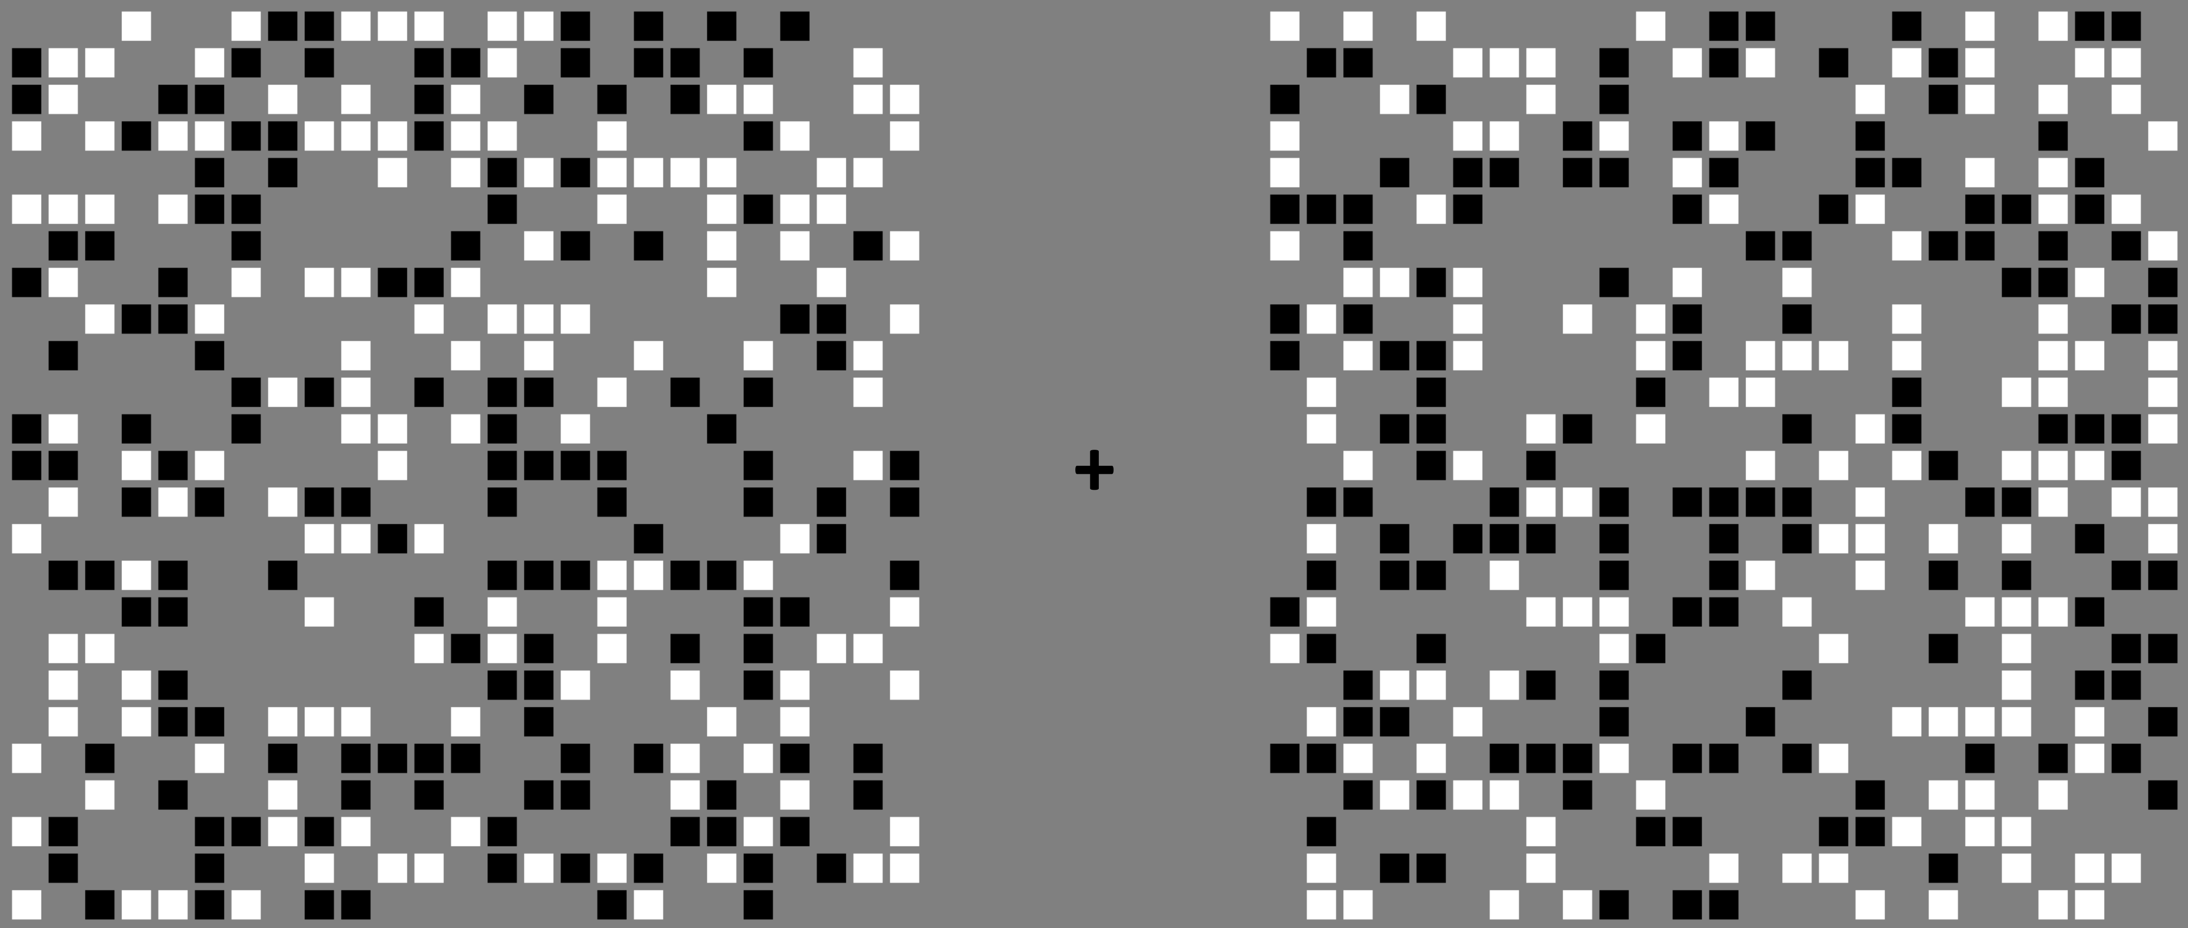


Cognition 249 (2024) 105813

Contents lists available at ScienceDirect

Cognition

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

Number adaptation: A critical look

Sami R. Yousif[a,*], Sam Clarke[b], Elizabeth M. Brannon[a]

[a] Department of Psychology, University of Pennsylvania, USA
[b] Department of Philosophy, University of Southern California, USA

ARTICLE INFO

Keywords:
Number
Perception
Adaptation
Old news

ABSTRACT

It is often assumed that adaptation — a temporary change in sensitivity to a perceptual dimension following exposure to that dimension — is a litmus test for what is and is not a "primary visual attribute". Thus, papers purporting to find evidence of number adaptation motivate a claim of great significance: That number is something that can be seen in much the way that canonical visual features, like color, contrast, size, and speed, can. Fifteen years after its reported discovery, number adaptation's existence seems to be nearly undisputed, with dozens of papers documenting support for the phenomenon. The aim of this paper is to offer a counterweight — to critically assess the evidence for and against number adaptation. After surveying the many reasons for thinking that number adaptation exists, we introduce several lesser-known reasons to be skeptical. We then advance an alternative account — the old news hypothesis — which can accommodate previously published findings while explaining various (otherwise unexplained) anomalies in the existing literature. Next, we describe the results of eight pre-registered experiments which pit our novel old news hypothesis against the received number adaptation hypothesis. Collectively, the results of these experiments undermine the number adaptation hypothesis on several fronts, whilst consistently supporting the old news hypothesis. More broadly our work raises questions about the status of adaptation itself as a means of discerning what is and is not a visual attribute.

1. Introduction

It is sometimes joked that vision science primarily serves to catalogue phenomena long known by magicians, cinematographers, and petty thieves. Occasionally, however, its discoveries offer to profoundly transform our understanding of what it means to see. Take the reported discovery of *visual number adaptation*. Since the pioneering work of Burr and Ross (2008) it has become widely accepted that observers visually adapt to the number of items in a seen collection, much as we adapt to other visible properties, like color, size, and motion. The claim is that prolonged exposure to a large number of seen items causes a middling number of items in that region to appear less numerous than they otherwise would. Conversely, prolonged exposure to a small number of items reportedly causes a middling number of items in that region to appear more numerous than they otherwise would.

These are stunning results. In canonical examples of visual number adaptation, observers enjoy obvious and phenomenologically striking aftereffects. If you adapt to 300 dots in a left-hand region of visual space, a test display containing 100 dots in that region will look remarkably sparse when compared to an otherwise identical collection of 100 dots in an un-adapted region (see Burr & Ross, 2008; see also Demo #1 in the supplemental materials on our OSF page). Since researchers have taken steps to rule out simpler explanations (e.g., by controlling for the total brightness and/or surface area of collections), received wisdom is that these results reflect adaptation to *the number* of items in seen collections. And because adaptation effects of this sort have been deemed rare or absent from thought and post-perceptual cognition (Block, 2022; Webster, 2015; c.f. Phillips & Firestone, 2022), number adaptation has been taken to suggest that number is a "primary" visual attribute, on a par with color and other low-level visual properties (Anobile, Cicchini, & Burr, 2016; c.f., Smortchkova, 2020). So, while numbers are abstract objects, located outside of space or time, number adaptation has been taken to establish that numbers nevertheless feature in the contents of human vision and visual experience; that, strange as it sounds, we literally *see* number.

Given the practical, philosophical, and theoretical implications of these claims, it is perhaps surprising that the existence of visual number adaptation has gone largely unchallenged (but see Dakin, Tibber,

* Corresponding author at: Department of Psychology, University of Pennsylvania, 425 S. University Ave, Stephen A. Levin Bldg., Philadelphia, PA 19104-6241, USA.
E-mail address: sryousif@sas.upenn.edu (S.R. Yousif).

https://doi.org/10.1016/j.cognition.2024.105813
Received 18 February 2024; Received in revised form 7 May 2024; Accepted 13 May 2024
Available online 30 May 2024

To *see* more: cogdevlab.org/number-adaptation/

Cognition 254 (2025) 105939

Contents lists available at ScienceDirect

Cognition

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

ELSEVIER

COGNITION

Discussion

Seven reasons to (still) doubt the existence of number adaptation: A rebuttal to Burr et al. and Durgin

Sami R. Yousif [a,*], Sam Clarke [b], Elizabeth M. Brannon [c]

[a] *Department of Psychology and Neuroscience, University of North Carolina, Chapel Hill, Chapel Hill, NC, USA*
[b] *Department of Philosophy, University of Southern California, Los Angeles, CA, USA*
[c] *Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA*

A B S T R A C T

Does the visual system adapt to number? For more than fifteen years, most have assumed that the answer is an unambiguous "yes". Against this prevailing orthodoxy, we recently took a critical look at the phenomenon, questioning its existence on both empirical and theoretical grounds, and providing an alternative explanation for extant results (*the old news hypothesis*). We subsequently received two critical responses. Burr, Anobile, and Arrighi rejected our critiques wholesale, arguing that the evidence for number adaptation remains overwhelming. Durgin questioned our old news hypothesis — preferring instead a theory about density adaptation he has championed for decades — but also highlighted several ways in which our arguments do pose serious challenges for proponents of number adaptation. Here, we reply to both. We first clarify our position regarding number adaptation. Then, we respond to our critics' concerns, highlighting seven reasons why we remain skeptical about number adaptation. We conclude with some thoughts about where the debate may head from here.

The phenomenon of number adaptation occupies a funny place in our hearts. All three of us were once convinced of the existence of number adaptation and cited it accordingly (see, e.g., Clarke & Beck, 2021; Fornaciai et al., 2017; Yousif & Keil, 2020). This is the lens through which our recent critique of number adaptation (Yousif et al., 2024) should be viewed. We are not protecting a theoretical position that we've wedded ourselves to. We are taking a critical look at number adaptation *despite* our prior theoretical commitments.

Our aim for this work was not to challenge the evidence that number is directly perceived (see Section #6.2 of our original paper) but instead to voice our growing concern that what looks like number adaptation is actually a more basic form of item-level adaptation. Here, we respond to two critiques of this suggestion. Burr, Arrighi and Anobile reject our concerns with number adaptation, saying that "adaptation to numerosity…has been far from refuted" (p. 1), and that "there exists overwhelming evidence for numerosity adaptation" (p. 6). Burr and colleagues contend that we *ought* to believe in number adaptation in principle — because adaptation has been observed for other perceptual attributes. "If [number] did not adapt," they write, "it would be unique amongst perceptual attributes, worthy of very special attention" (p. 5). Ironically, we worry that many other forms of high-level adaptation have become widely accepted precisely because of the path that number adaptation paved. This is what makes the stakes in this debate so high. Perhaps, influenced by the fuzzy criteria that have been used to support number adaptation, other cases of high-level adaptation have been similarly misunderstood (see, e.g., Storrs, 2015; Yousif & Clarke, 2024). If one domino falls, others may follow. As foreshadowed in the final sentence of our original paper, this debate "raises questions about the nature and meaning of adaptation itself" (p. 14).

In contrast, Durgin (2024) suggests that our critique of number adaptation brings into focus "how fragile some findings may be" (p. 1) and applauds our "simpler, but not simplistic, approach to understanding a complex phenomenon" (p. 6). He explains that our work "both illustrates and encourages critical thinking about the nature of adaptation" (p. 6). Nevertheless, he argues that our theory is incomplete and favors his long-established view that density is the mechanism underlying putative number adaptation (see Durgin, 1995, 2008).

In what follows we first clarify our position and then respond to our commentators' critiques, offering seven reasons why we continue to think that the evidence for number adaptation is unconvincing — and why we continue to think that our 'old news hypothesis' explains extant results better than the existing alternatives. We conclude with some thoughts about where this debate may head from here.

1. What we are arguing

Our claim is that apparent cases of number adaptation are unlikely to be driven by adaptation to number *per se*. We argue that (1) The

* Corresponding author.
E-mail address: sami.yousif@unc.edu (S.R. Yousif).

To *see* more: cogdevlab.org/number-adaptation/

Food for thought:

Food for thought:

1. I mentioned that the domain-specific number encoding theory is dominant despite the evidence against it. **Why?**

Food for thought:

1. I mentioned that the domain-specific number encoding theory is dominant despite the evidence against it. **Why?**

2. Should we believe that number adaptation is genuine and/or like other kinds of adaptation? **Why or why not?**

Food for thought:

1. I mentioned that the domain-specific number encoding theory is dominant despite the evidence against it. **Why?**

2. Should we believe that number adaptation is genuine and/or like other kinds of adaptation? **Why or why not?**

3. **What makes *number* worth talking about?**


ILLUSIONS
OF FACES
Sami Ryan Yousif

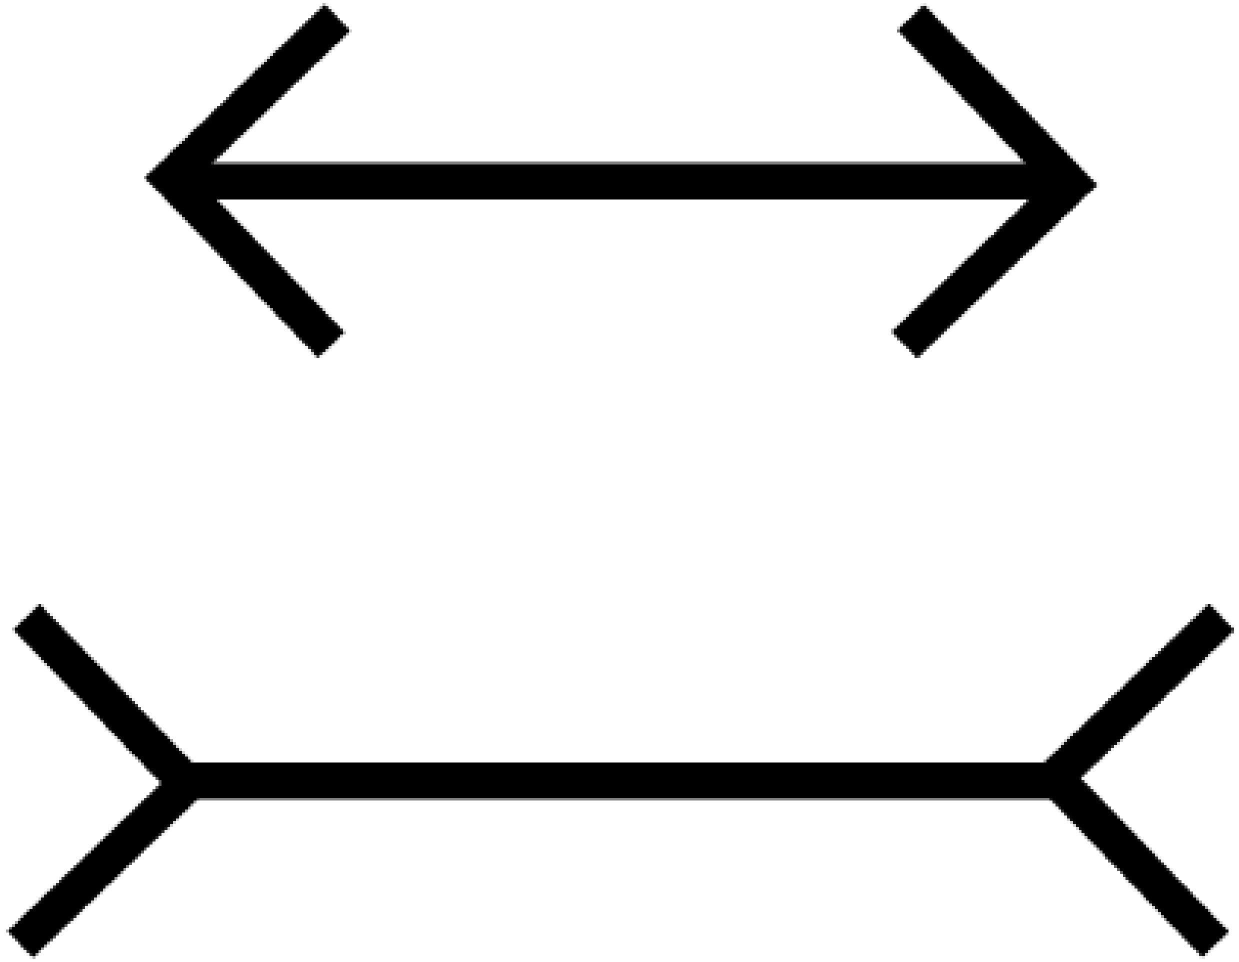

The **Muller-Lyer illusion** ‼

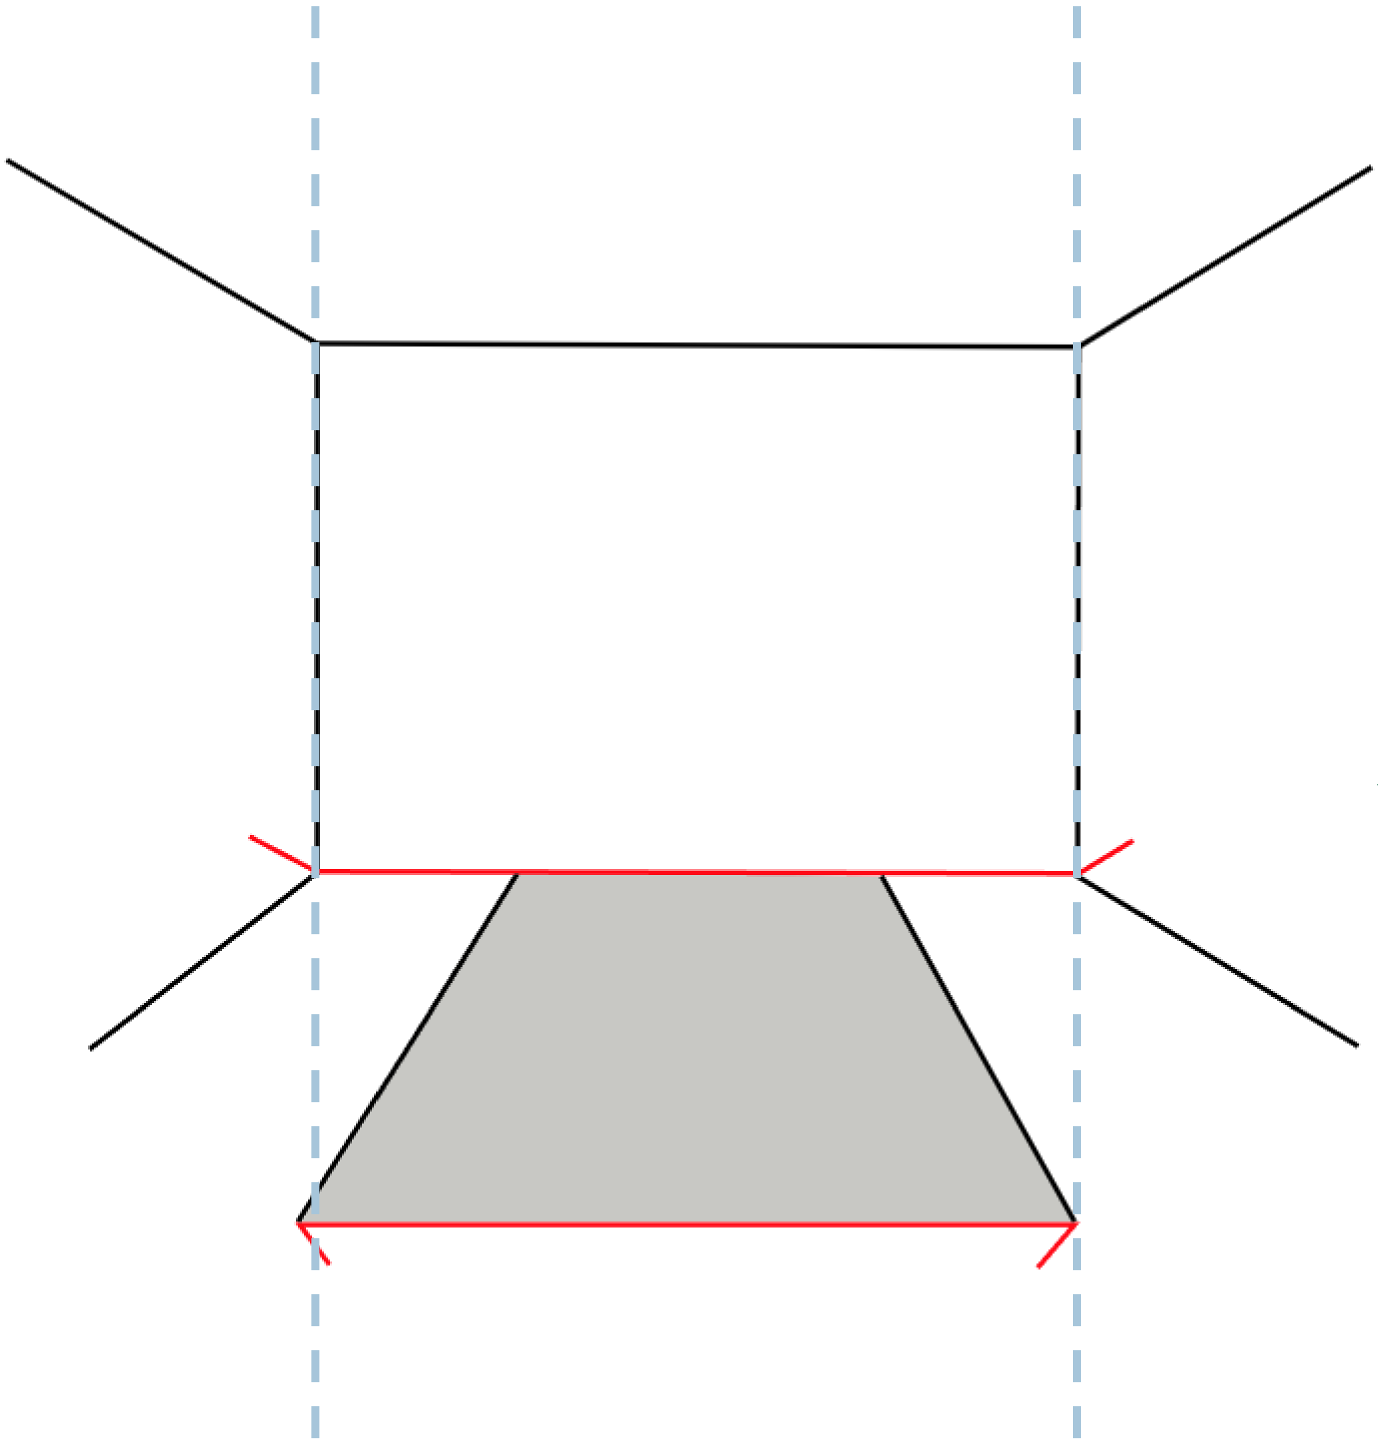

The "Carpentered world hypothesis"

Is visual perception WEIRD?
The Müller-Lyer illusion and the Cultural Byproduct Hypothesis

Dorsa Amir[1] & Chaz Firestone[2]

[1]Department of Psychology and Neuroscience, Duke University
[2]Department of Psychological and Brain Sciences, Johns Hopkins University

Correspondence: Dorsa Amir (dorsa.amir@duke.edu), Chaz Firestone (chaz@jhu.edu)

Preprint, 1/24/25 — Forthcoming in ***Psychological Review***

Is visual perception WEIRD?
The Müller-Lyer illusion and the Cultural Byproduct Hypothesis

Dorsa Amir[1] & Chaz Firestone[2]

[1]Department of Psychology and Neuroscience, Duke University
[2]Department of Psychological and Brain Sciences, Johns Hopkins University

Correspondence: Dorsa Amir (dorsa.amir@duke.edu), Chaz Firestone (chaz@jhu.edu)

Preprint, 1/24/25 — Forthcoming in ***Psychological Review***

Abstract

A fundamental question in the psychological sciences is the degree to which culture shapes core cognitive processes — perhaps none more foundational than how we perceive the world around us. A dramatic and oft-cited "case study" of culture's power in this regard is the Müller-Lyer illusion, which depicts two lines of equal length but with arrowheads pointing either inward or outward, creating the illusion that one line is longer than the other. According to a line of research stretching back over a century, depending on the society you were raised in (and how much carpentry you were exposed to), you may not see the illusion at all — an ambitious and influential research program motivating claims that seemingly basic aspects of visual processing may actually be "culturally evolved byproducts". This *Cultural Byproduct Hypothesis* bears on foundational issues in the science, philosophy, and sociology of psychology, and remains popular today. Yet, here we argue that it is almost certainly false. We synthesize evidence from diverse fields which demonstrate that: (1) the illusion is not limited to humans, appearing in non-human animals from diverse ecologies; (2) the statistics of natural scenes are sufficient to capture the illusion; (3) the illusion does not require straight lines typical of carpentry (nor even any lines at all); (4) the illusion arises in sense modalities other than vision; and (5) the illusion arises even in congenitally blind subjects. Moreover, by reexamining historical data and ethnographic descriptions from the original case studies, we show that the evidence for cultural variation and its correlation with key cultural variables is in fact highly inconsistent, beset by questionable research practices, and misreported by later discussions. Together, these considerations undermine the most popular and dramatic example of cultural influence on perception. We further extend our case beyond this phenomenon, showing that many of these considerations apply to other visual illusions as well, including similarly implicated visual phenomena such as the Ebbinghaus, Ponzo, Poggendorf, and Horizontal-Vertical illusions. We conclude by outlining future approaches to cross-cultural research on perception, and we also point to other potential sources of cultural variation in visual processing.

Is visual perception WEIRD?

The Müller-Lyer illusion and the Cultural Byproduct Hypothesis

Dorsa Amir[1] & Chaz Firestone[2]

[1]Department of Psychology and Neuroscience, Duke University
[2]Department of Psychological and Brain Sciences, Johns Hopkins University

Correspondence: Dorsa Amir (dorsa.amir@duke.edu), Chaz Firestone (chaz@jhu.edu)

Preprint, 1/24/25 — Forthcoming in ***Psychological Review***

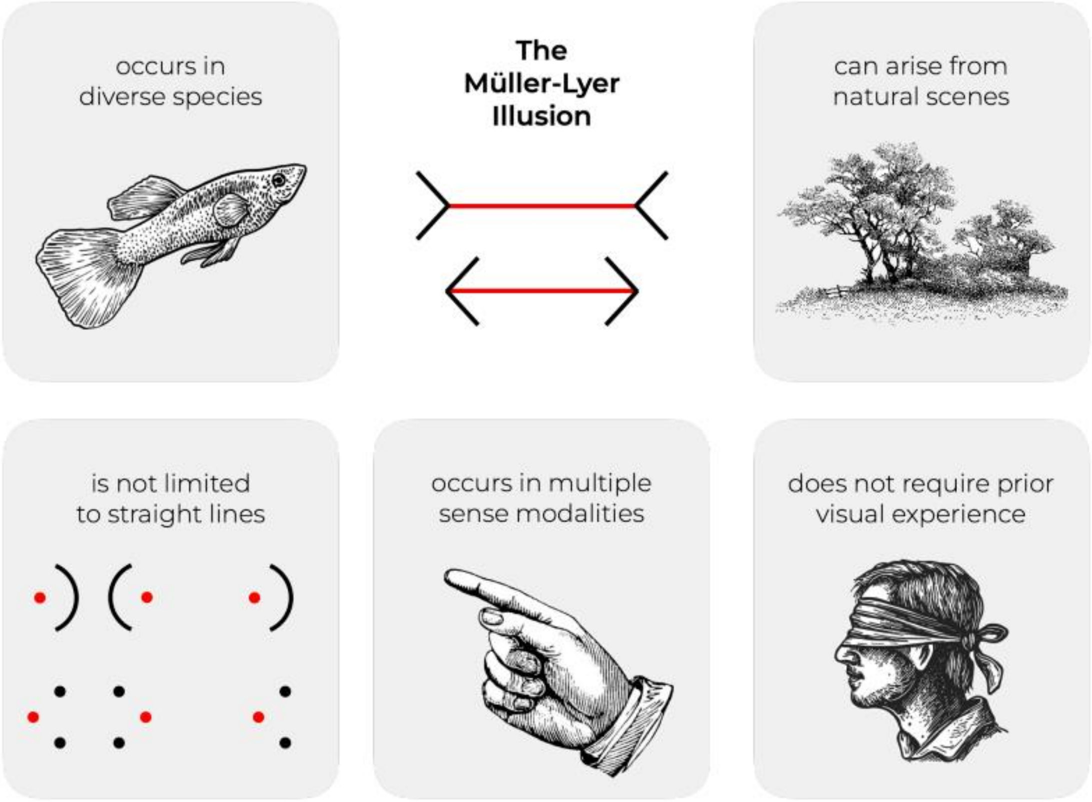

Figure 2: A visual summary of empirical challenges to the Cultural Byproduct Hypothesis of visual illusion, using the Müller-Lyer illusion as a case study. The illusion arises in species from diverse ecologies (**Section 4.1**), can be captured by the statistic of natural scenes (**Section 4.2**), occurs with stimuli that do not have straight lines (nor even any lines at all; **Section 4.3**), arises in senses other than vision (**Section 4.4**), and is experienced even by congenitally blind subjects whose sight is restored (**Section 4.5**).

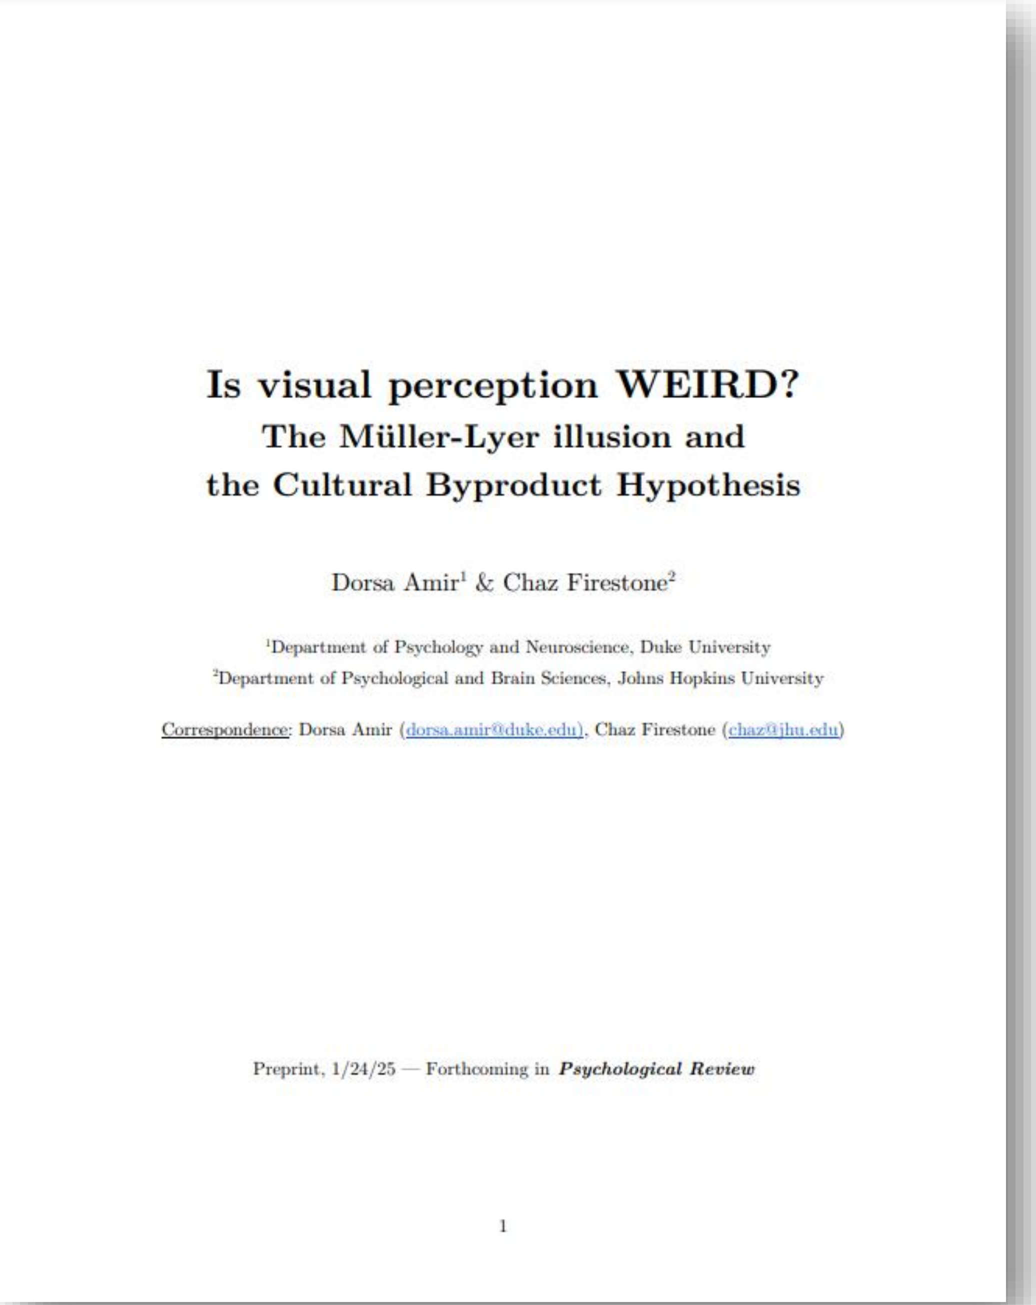

Is visual perception WEIRD?
The Müller-Lyer illusion and the Cultural Byproduct Hypothesis

Dorsa Amir[1] & Chaz Firestone[2]

[1]Department of Psychology and Neuroscience, Duke University
[2]Department of Psychological and Brain Sciences, Johns Hopkins University

Correspondence: Dorsa Amir (dorsa.amir@duke.edu), Chaz Firestone (chaz@jhu.edu)

Preprint, 1/24/25 — Forthcoming in ***Psychological Review***

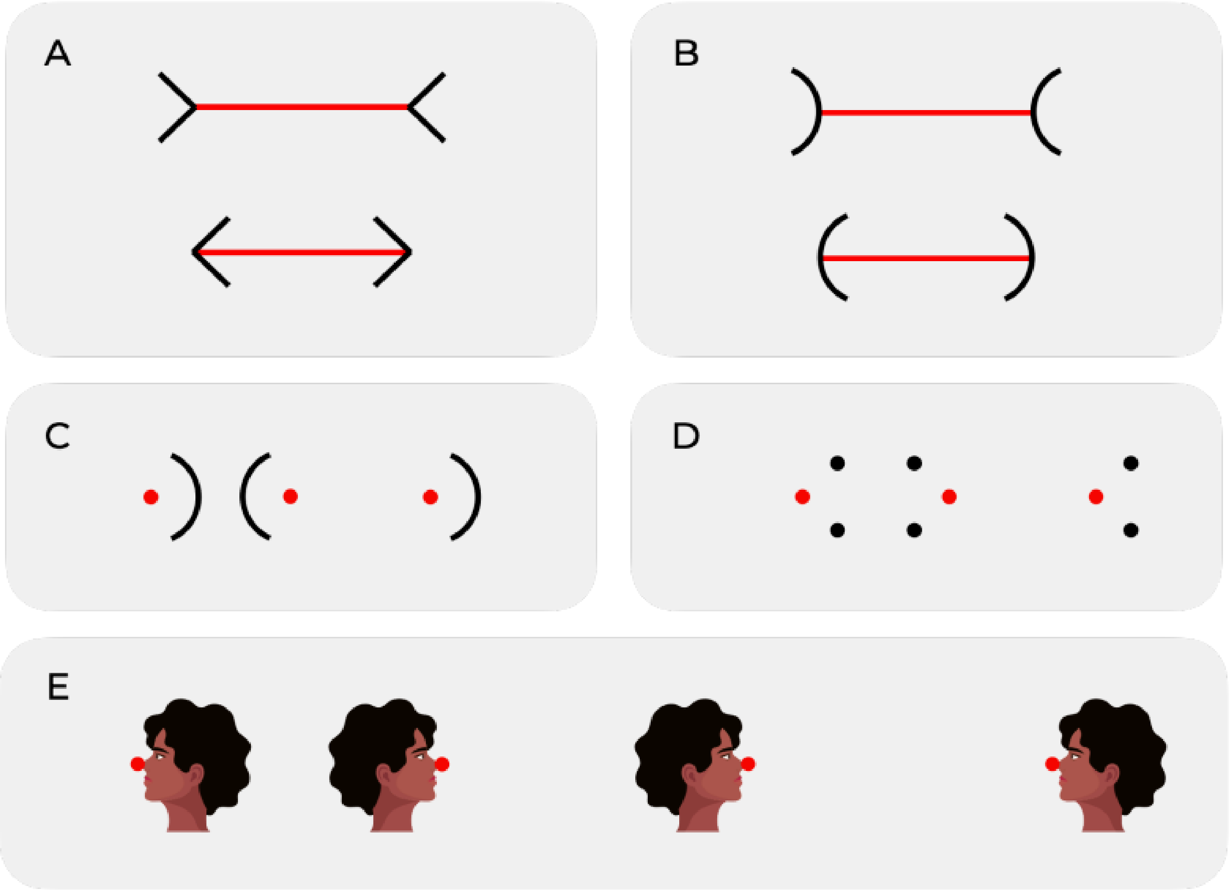

Figure 5: (A) The traditional Müller-Lyer illusion, along with variants based on (B) semicircles with lines, (C) semicircles with dots, (D) dots only, and (E) faces. These additional variants are increasingly distant from the straight lines and sharp angles present in the original stimulus.

Is visual perception WEIRD?

The Müller-Lyer illusion and the Cultural Byproduct Hypothesis

Dorsa Amir[1] & Chaz Firestone[2]

[1]Department of Psychology and Neuroscience, Duke University

[2]Department of Psychological and Brain Sciences, Johns Hopkins University

Correspondence: Dorsa Amir (dorsa.amir@duke.edu), Chaz Firestone (chaz@jhu.edu)

Preprint, 1/24/25 — Forthcoming in ***Psychological Review***

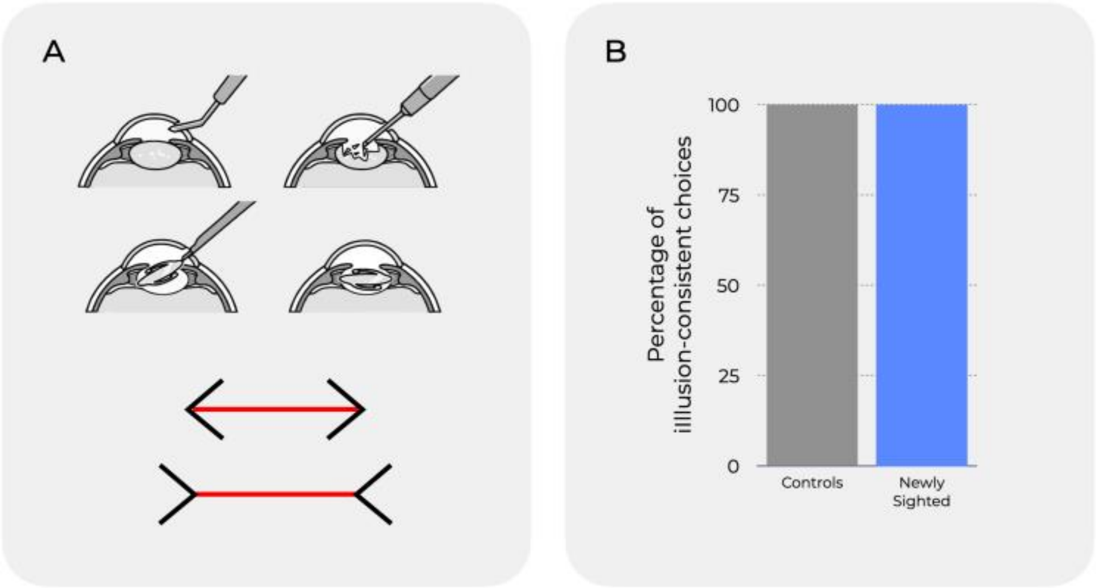

Figure 7: (A) Congenital cataracts can be replaced by a procedure in which the cataract — the clouded lens — is emulsified; an intraocular lens is then put in its place, effectively restoring vision. (B) Immediately upon the restoration of sight, newly sighted subjects are susceptible to the Müller-Lyer illusion (data drawn from Gandhi et al., 2015).

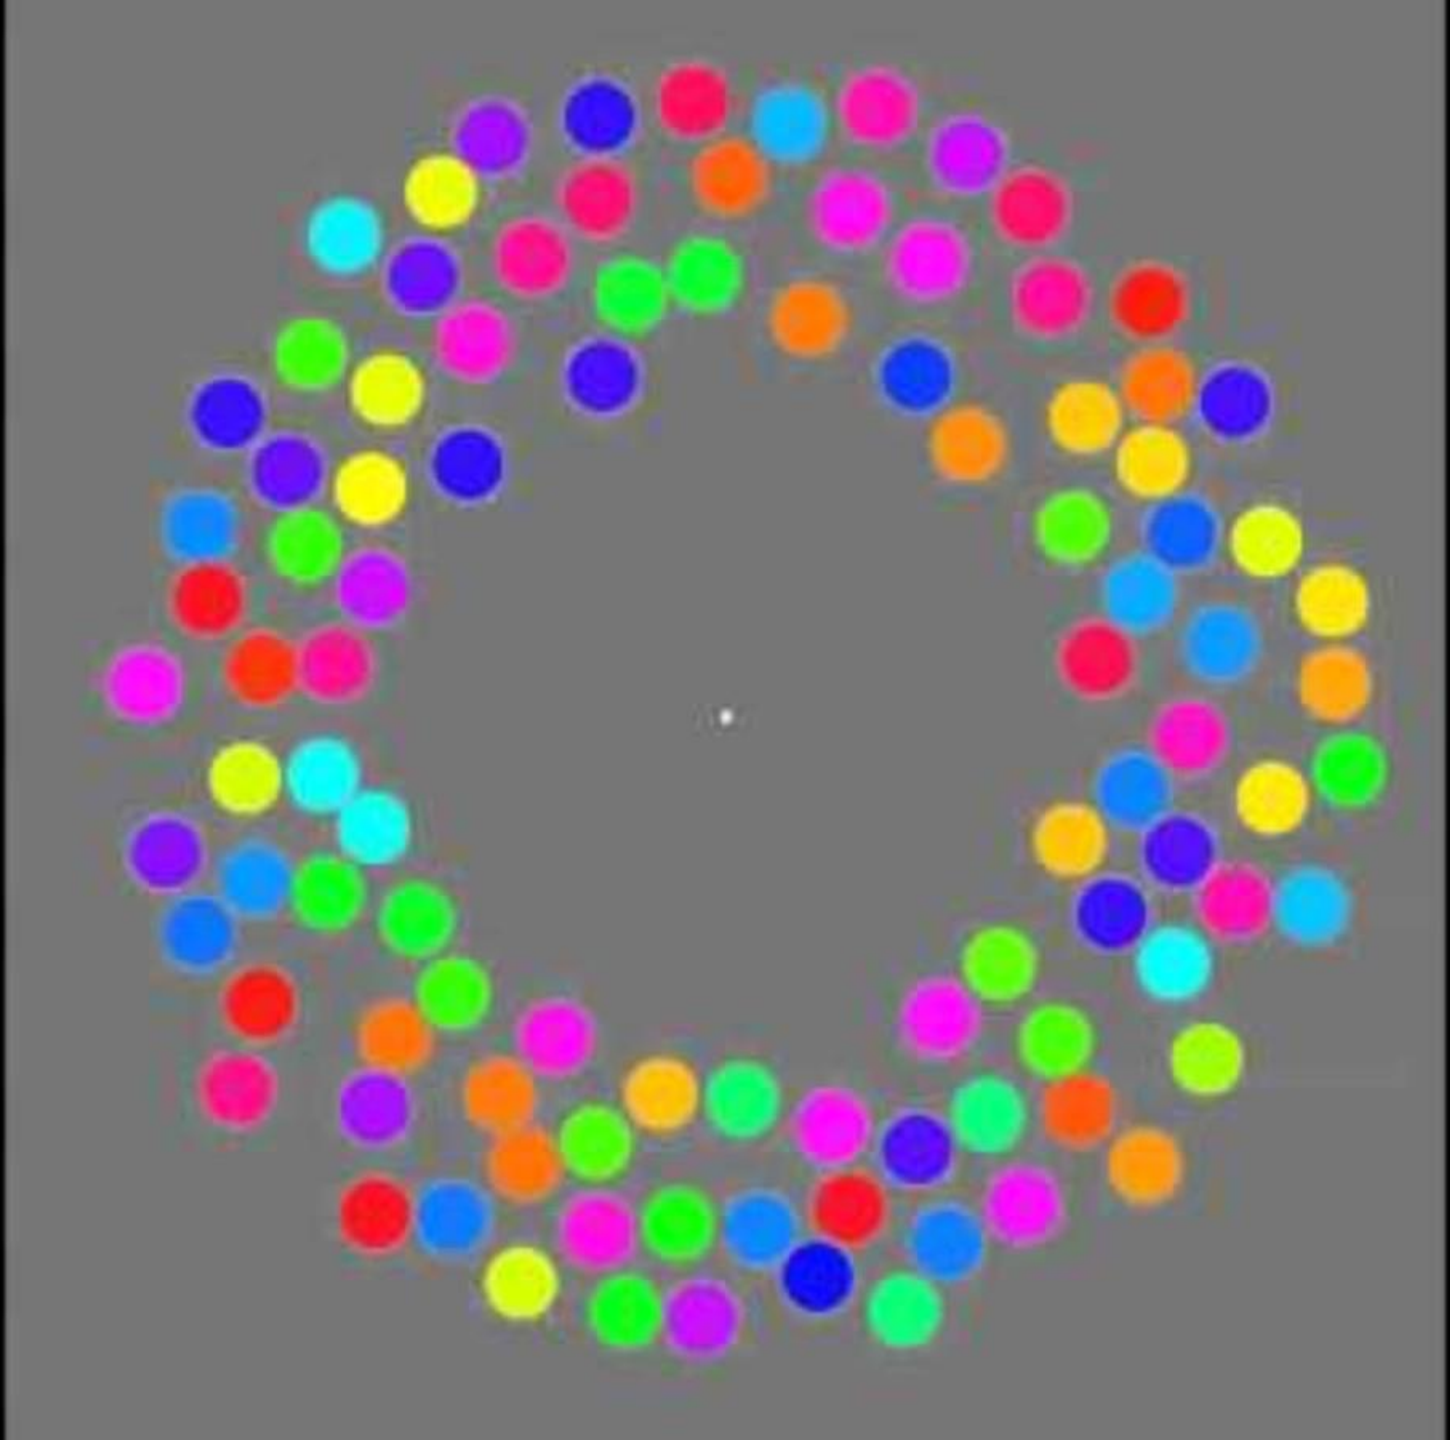

Research Article

Psychological Science
2024, Vol. 35(5) 504–516

DOI: 10.1177/09567976241235104
www.psychologicalscience.org/PS

The Motion-Silencing Illusion Depends on Object-Centered Representation

Qihan Wu and Jonathan I. Flombaum
Department of Psychological and Brain Sciences, Johns Hopkins University

Abstract

Motion silencing is a striking and unexplained visual illusion wherein changes that are otherwise salient become difficult to perceive when the changing elements also move. We develop a new method for quantifying illusion strength (Experiments 1a and 1b), and we demonstrate a privileged role for rotational motion on illusion strength compared with highly controlled stimuli that lack rotation (Experiments 2a to 3b). These contrasts make it difficult to explain the illusion in terms of lower-level detection limits. Instead, we explain the illusion as a failure to attribute changes to locations. Rotation exacerbates the illusion because its perception relies upon structured object representations. This aggravates the difficulty of attributing changes by demanding that locations are referenced relative to both an object-internal frame and an external frame. Two final experiments (4a and 4b) add support to this account by employing a synchronously rotating external frame of reference that diminishes otherwise strong motion silencing. All participants were Johns Hopkins University undergraduates.

Keywords

vision, object perception, change detection, illusions, preregistered

Received 8/22/23; Revision accepted 1/11/24

Motion silencing (MS) is a striking illusion wherein changes that are otherwise salient become difficult to perceive when the changing elements move (Suchow & Alvarez, 2011). The typical stimulus includes 100 dots arranged to form a ring. Each dot starts with a random color, changing color by continuously cycling through color space. The changing colors are highly noticeable when the ring is stationary but barely so when the ring rotates. Now the color changes become difficult to detect, silenced, an illusion that gains strength as rotation speed increases.

The prevailing explanation for the illusion appears to be that local change detectors have small receptive fields (Suchow & Alvarez, 2011). Fast rotation means that color changes happen too slowly to register. Effectively, if an item remains unchanged as it passes through a detector, the change cannot be caught, even if the item has changed quite a bit when you compare it after and before.

The present study sought to explore MS in two ways. First, we sought to investigate the effects of silencing on objective discrimination. A changing but silenced stimulus should be difficult to distinguish from an unchanging stimulus. Second, we sought evidence to support a higher-level explanation for the illusion: that it arises when color changes are detected but difficult to attribute to specific locations. Toward this end, we examine whether rotation amplifies silencing, compared with similar but nonrotational motion. We reason that rotation complicates the process of attributing changes to locations because it increases the demand for labeling locations in multiple reference frames at once.

We were motivated to explore whether rotation increases silencing for two related reasons. First, we observed that nearly all the examples of silencing involve rotation. But the two exceptions we are aware of support the view that silencing is an attribution error. Peirce (2013) demonstrated that motion itself can be silenced in the presence of global changes—such as

Corresponding Author:
Jonathan I. Flombaum, Department of Psychological and Brain Sciences, Johns Hopkins University
Email: flombaum@jhu.edu

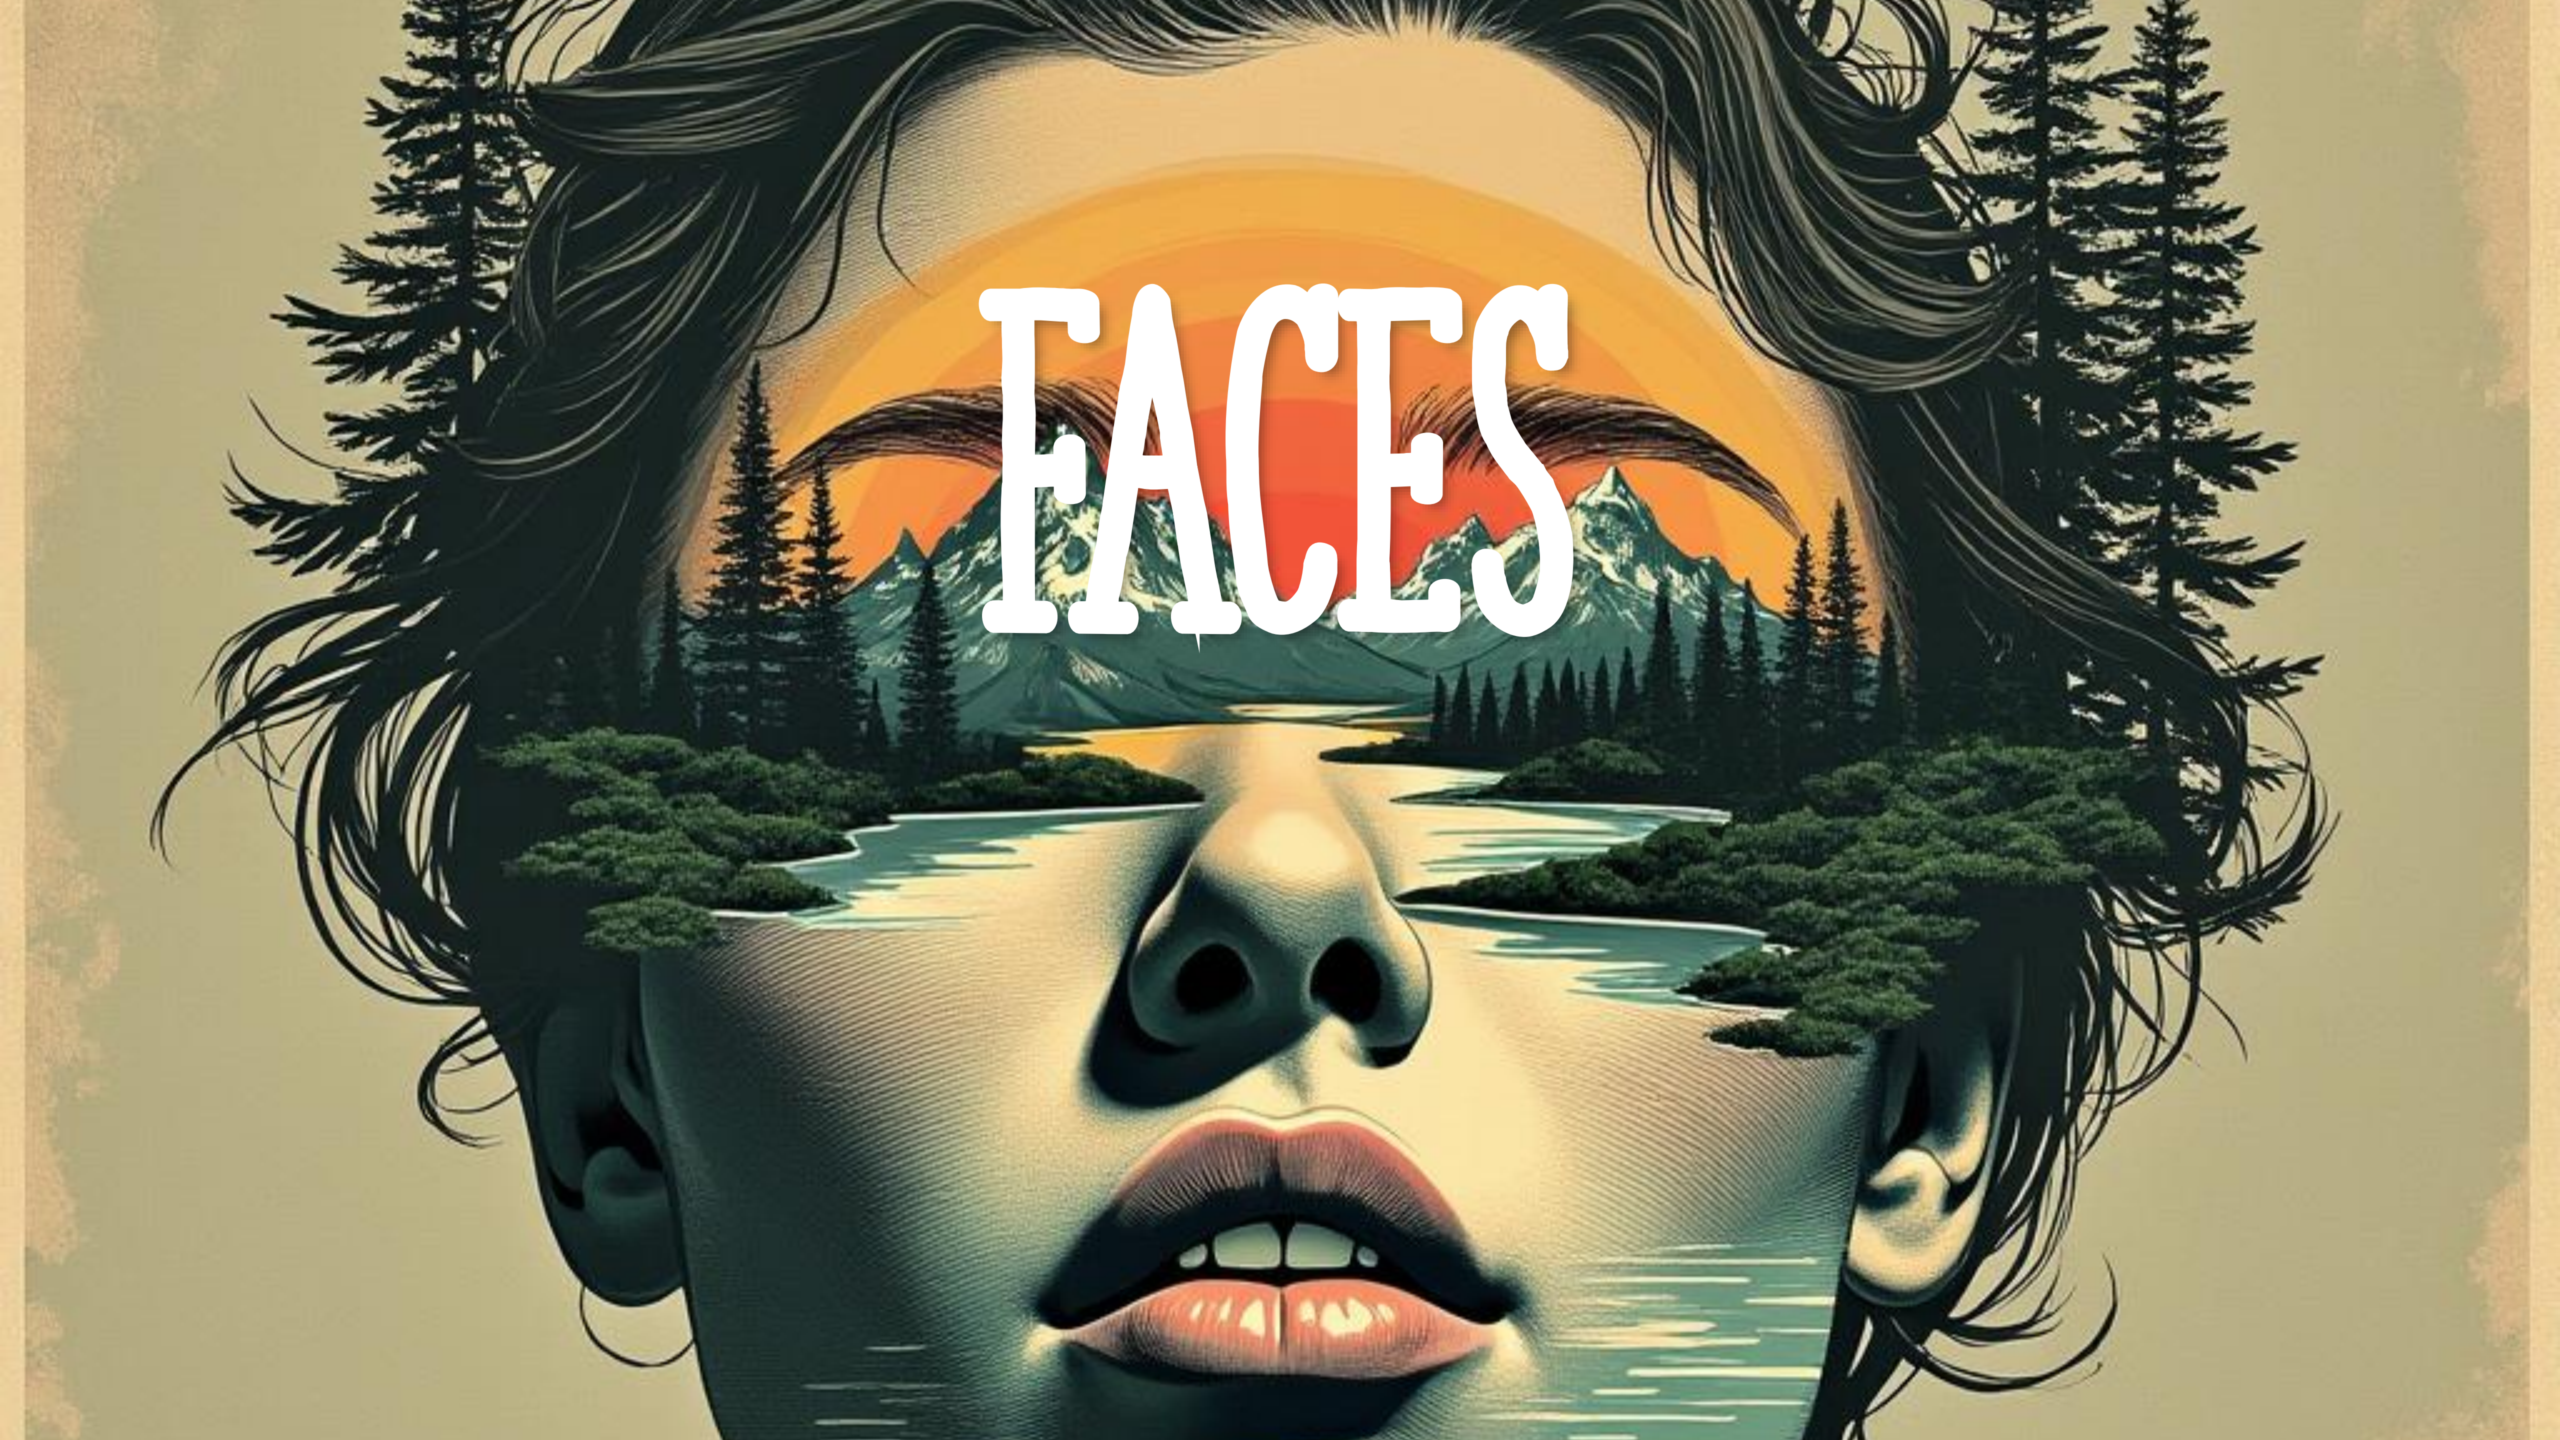
FACES

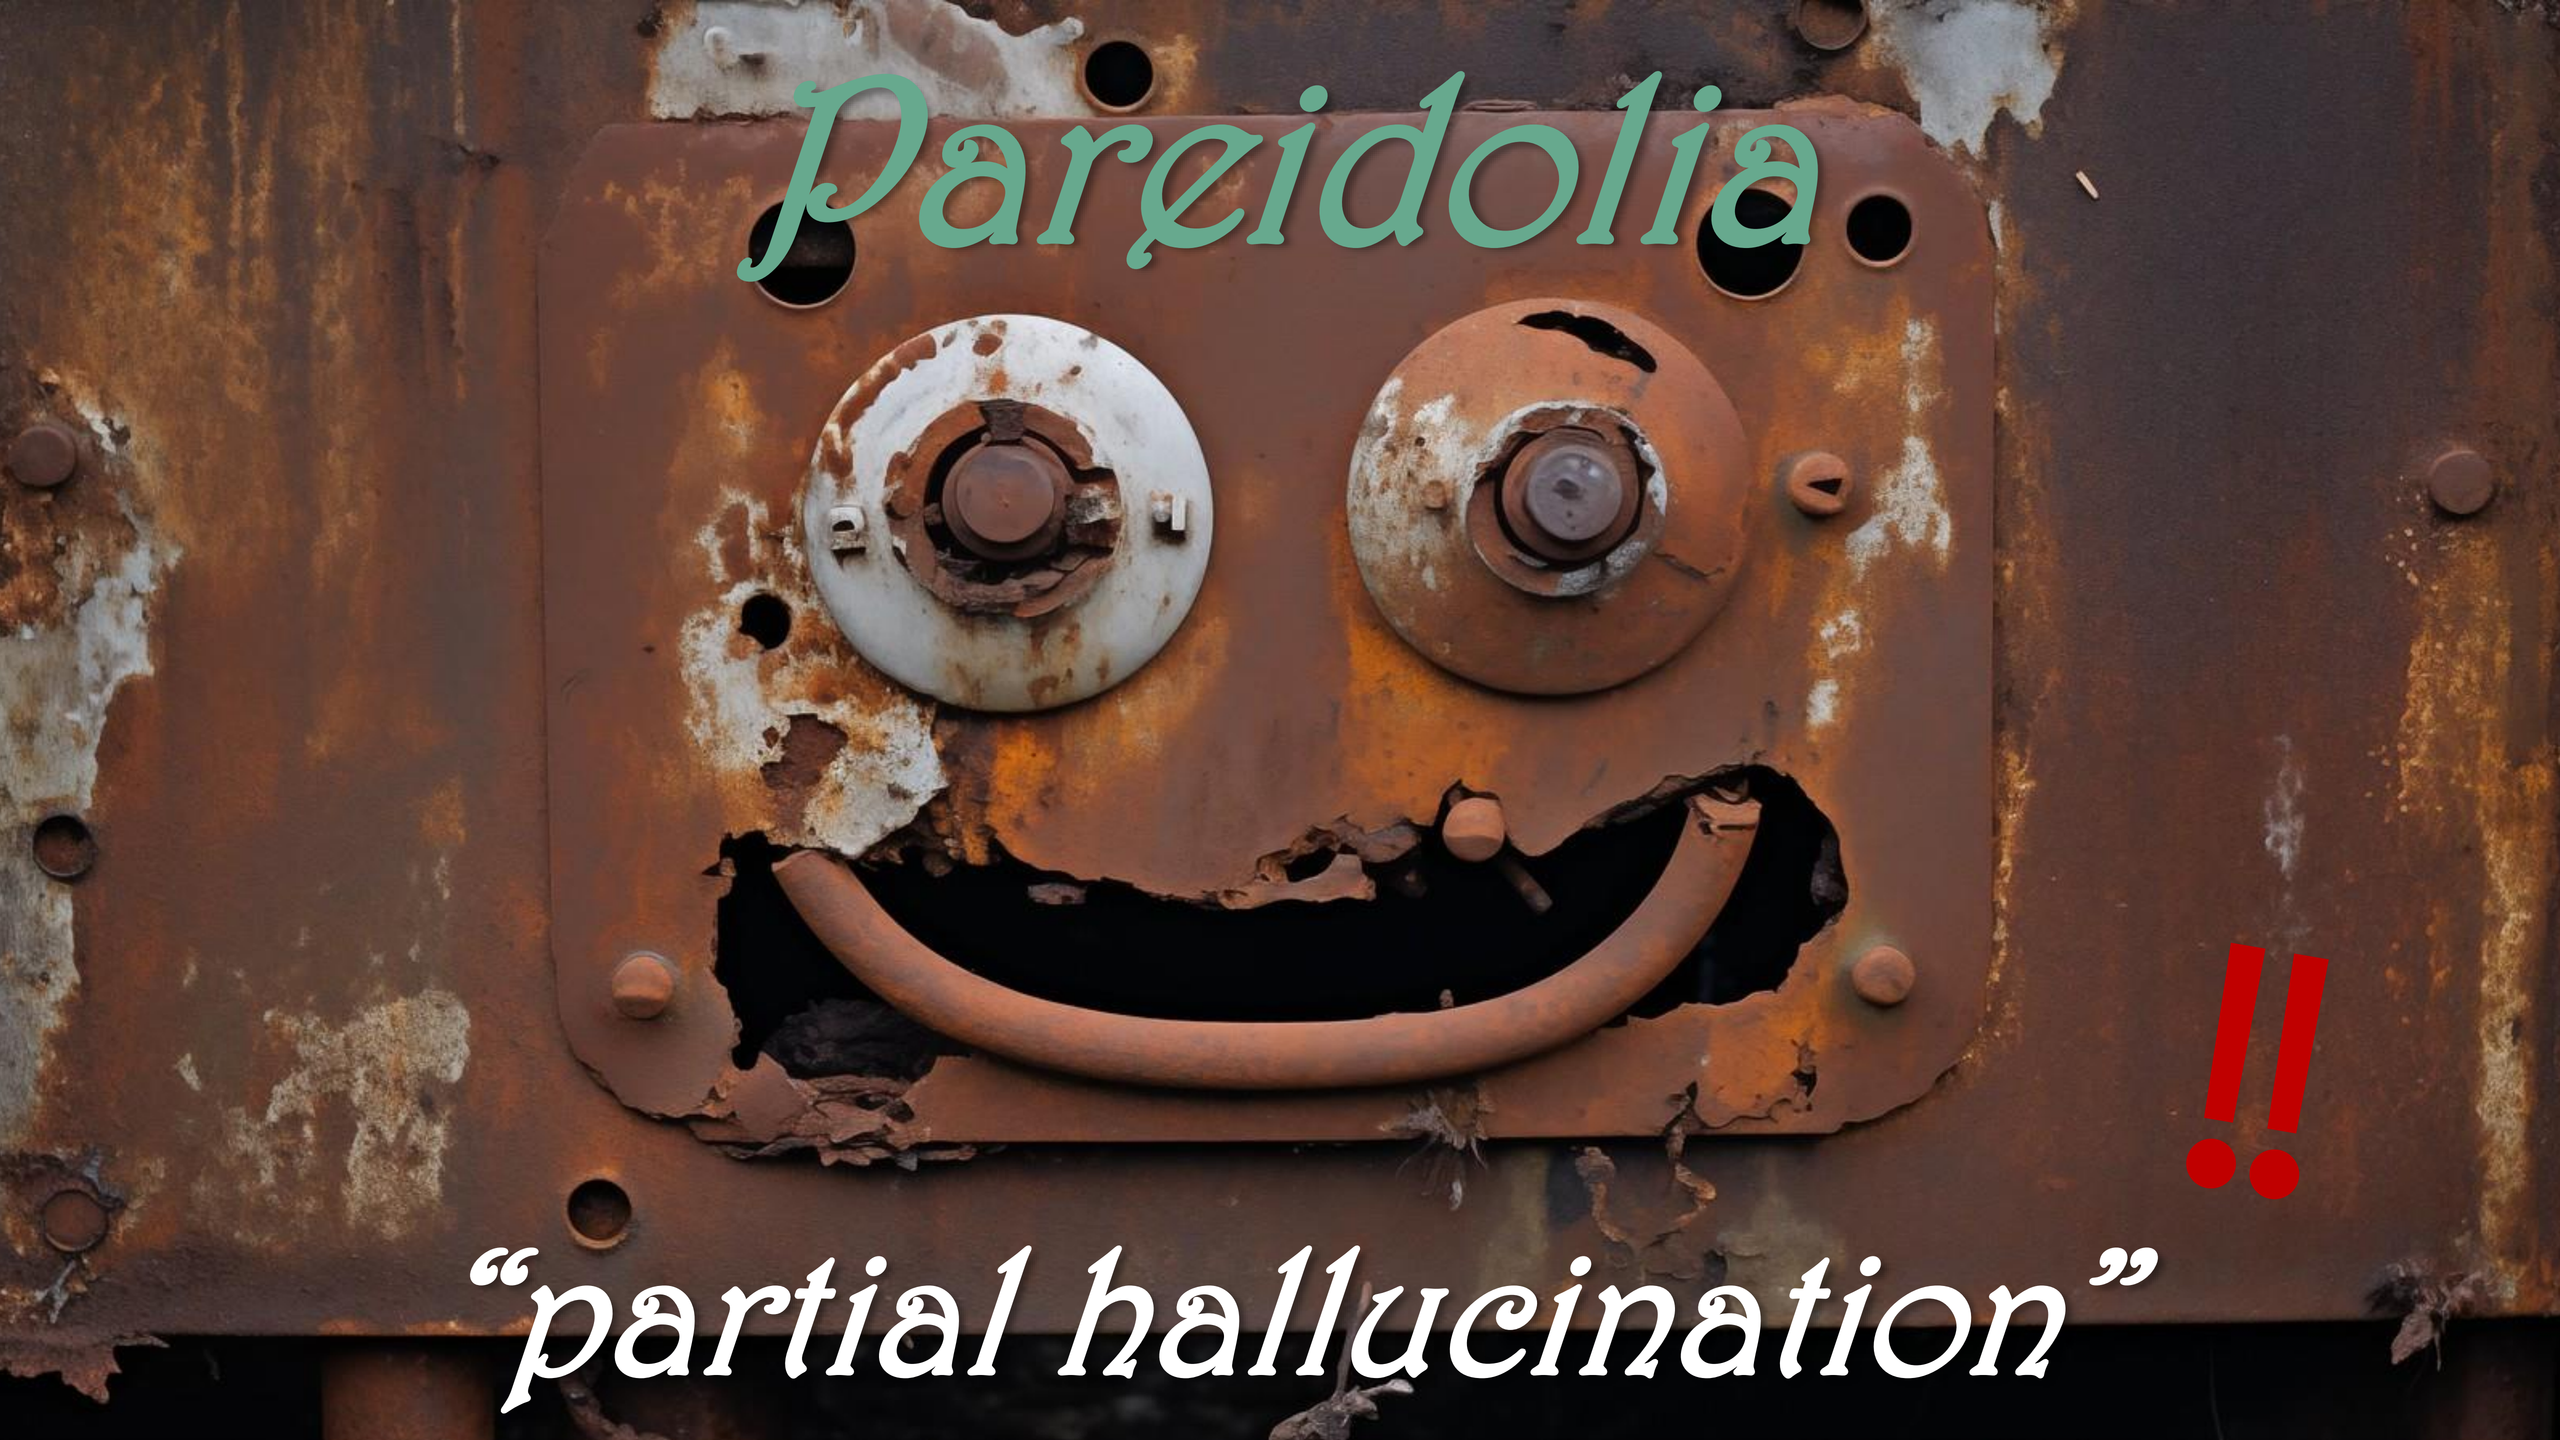
Pareidolia
“partial hallucination”
‼

Faces are perhaps *the most important* stimulus we regularly encounter.

Which of the following faces is more trustworthy?


PSYCHOLOGICAL SCIENCE

Research Article

First Impressions

Making Up Your Mind After a 100-Ms Exposure to a Face

Janine Willis and Alexander Todorov

Princeton University

ABSTRACT—***People often draw trait inferences from the facial appearance of other people. We investigated the minimal conditions under which people make such inferences. In five experiments, each focusing on a specific trait judgment, we manipulated the exposure time of unfamiliar faces. Judgments made after a 100-ms exposure correlated highly with judgments made in the absence of time constraints, suggesting that this exposure time was sufficient for participants to form an impression. In fact, for all judgments—attractiveness, likeability, trustworthiness, competence, and aggressiveness—increased exposure time did not significantly increase the correlations. When exposure time increased from 100 to 500 ms, participants' judgments became more negative, response times for judgments decreased, and confidence in judgments increased. When exposure time increased from 500 to 1,000 ms, trait judgments and response times did not change significantly (with one exception), but confidence increased for some of the judgments; this result suggests that additional time may simply boost confidence in judgments. However, increased exposure time led to more differentiated person impressions.***

Lavater's (1772/1880) *Essays on Physiognomy*, which was written in 1772 and reprinted in more than 150 editions by 1940, described in minute detail how to relate facial features to personality traits (e.g., "the nearer the eyebrows are to the eyes, the more earnest, deep, and firm the character," p. 59). Although these ideas strike most people today as ludicrous and bring to mind phrenology, empirical evidence shows that the effects of facial appearance on social outcomes are pervasive. In almost every significant domain of life, attractive people get better outcomes than unattractive people (Hamermesh & Biddle, 1994; Zebrowitz, 1999). The effects of baby-faced appearance are as pervasive as are the effects of attractiveness (Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, baby-faced individuals are less likely to receive severe judicial outcomes than mature-faced individuals (Zebrowitz & McDonald, 1991).

From the structure of the face, people form not only global impressions, but also specific trait impressions (Hassin & Trope, 2000). For example, we showed that inferences of competence, based solely on facial appearance, predicted the outcomes of U.S. congressional elections in 2000, 2002, and 2004 (Todorov, Mandisodza, Goren, & Hall, 2005). Although we measured impressions on a variety of traits, including attractiveness, trustworthiness, and likeability, the trait inference that predicted the election outcomes was competence. Competence was also rated as the most important attribute for a person running for a public office. This finding suggests that person attributes that are important for specific decisions are inferred from facial appearance and influence these decisions.

From both the standard-intuition and the rational-actor points of view, trait inferences from facial appearance should not influence important deliberate decisions. However, to the extent that these inferences occur rapidly and effortlessly, their effects on decisions may be subtle and not subjectively recognized. Using the terms of dual-process theories (Chaiken & Trope, 1999; Kahneman, 2003), we have argued that trait inferences from faces can be characterized as fast, intuitive, unreflective System 1 processes that contrast with slow, effortful, and deliberate System 2 processes (Todorov et al., 2005). We provided preliminary evidence for this proposal by showing that inferences of competence based on 1-s exposure to the faces of the winners and the runners-up for the Senate races sufficed to predict the election outcomes.

In this article, we report a series of studies in which we systematically manipulated the exposure time of faces to further explore the minimal conditions under which people make trait inferences from facial appearance. Research on visual processing has shown that high-level object representations can be constructed very rapidly from visual scenes (Grill-Spector & Kanwisher, 2005; Rousselet, Fabre-Thorpe, & Thorpe, 2002; Thorpe, Fize, & Marlot, 1996). It is possible that inferences about socially significant attributes are also rapidly extracted

Address correspondence to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, e-mail: atodorov@princeton.edu.

Research Article

First Impressions

Making Up Your Mind After a 100-Ms Exposure to a Face

Janine Willis and Alexander Todorov

Princeton University

ABSTRACT—*People often draw trait inferences from the facial appearance of other people. We investigated the minimal conditions under which people make such inferences. In five experiments, each focusing on a specific trait judgment, we manipulated the exposure time of unfamiliar faces. Judgments made after a 100-ms exposure correlated highly with judgments made in the absence of time constraints, suggesting that this exposure time was sufficient for participants to form an impression. In fact, for all judgments—attractiveness, likeability, trustworthiness, competence, and aggressiveness—increased exposure time did not significantly increase the correlations. When exposure time increased from 100 to 500 ms, participants' judgments became more negative, response times for judgments decreased, and confidence in judgments increased. When exposure time increased from 500 to 1,000 ms, trait judgments and response times did not change significantly (with one exception), but confidence increased for some of the judgments; this result suggests that additional time may simply boost confidence in judgments. However, increased exposure time led to more differentiated person impressions.*

Lavater's (1772/1880) *Essays on Physiognomy*, which was written in 1772 and reprinted in more than 150 editions by 1940, described in minute detail how to relate facial features to personality traits (e.g., "the nearer the eyebrows are to the eyes, the more earnest, deep, and firm the character," p. 59). Although these ideas strike most people today as ludicrous and bring to mind phrenology, empirical evidence shows that the effects of facial appearance on social outcomes are pervasive. In almost every significant domain of life, attractive people get better outcomes than unattractive people (Hamermesh & Biddle, 1994; Zebrowitz, 1999). The effects of baby-faced appearance are as pervasive as are the effects of attractiveness (Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, baby-faced individuals are less likely to receive severe judicial outcomes than mature-faced individuals (Zebrowitz & McDonald, 1991).

From the structure of the face, people form not only global impressions, but also specific trait impressions (Hassin & Trope, 2000). For example, we showed that inferences of competence, based solely on facial appearance, predicted the outcomes of U.S. congressional elections in 2000, 2002, and 2004 (Todorov, Mandisodza, Goren, & Hall, 2005). Although we measured impressions on a variety of traits, including attractiveness, trustworthiness, and likeability, the trait inference that predicted the election outcomes was competence. Competence was also rated as the most important attribute for a person running for a public office. This finding suggests that person attributes that are important for specific decisions are inferred from facial appearance and influence these decisions.

From both the standard-intuition and the rational-actor points of view, trait inferences from facial appearance should not influence important deliberate decisions. However, to the extent that these inferences occur rapidly and effortlessly, their effects on decisions may be subtle and not subjectively recognized. Using the terms of dual-process theories (Chaiken & Trope, 1999; Kahneman, 2003), we have argued that trait inferences from faces can be characterized as fast, intuitive, unreflective System 1 processes that contrast with slow, effortful, and deliberate System 2 processes (Todorov et al., 2005). We provided preliminary evidence for this proposal by showing that inferences of competence based on 1-s exposure to the faces of the winners and the runners-up for the Senate races sufficed to predict the election outcomes.

In this article, we report a series of studies in which we systematically manipulated the exposure time of faces to further explore the minimal conditions under which people make trait inferences from facial appearance. Research on visual processing has shown that high-level object representations can be constructed very rapidly from visual scenes (Grill-Spector & Kanwisher, 2005; Rousselet, Fabre-Thorpe, & Thorpe, 2002; Thorpe, Fize, & Marlot, 1996). It is possible that inferences about socially significant attributes are also rapidly extracted

Address correspondence to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, e-mail: atodorov@princeton.edu.

!!

attractiveness

likeability

trustworthiness

competence

aggressiveness


PSYCHOLOGICAL SCIENCE

Research Article

First Impressions

Making Up Your Mind After a 100-Ms Exposure to a Face

Janine Willis and Alexander Todorov

Princeton University

ABSTRACT—*People often draw trait inferences from the facial appearance of other people. We investigated the minimal conditions under which people make such inferences. In five experiments, each focusing on a specific trait judgment, we manipulated the exposure time of unfamiliar faces. Judgments made after a 100-ms exposure correlated highly with judgments made in the absence of time constraints, suggesting that this exposure time was sufficient for participants to form an impression. In fact, for all judgments—attractiveness, likeability, trustworthiness, competence, and aggressiveness—increased exposure time did not significantly increase the correlations. When exposure time increased from 100 to 500 ms, participants' judgments became more negative, response times for judgments decreased, and confidence in judgments increased. When exposure time increased from 500 to 1,000 ms, trait judgments and response times did not change significantly (with one exception), but confidence increased for some of the judgments; this result suggests that additional time may simply boost confidence in judgments. However, increased exposure time led to more differentiated person impressions.*

Lavater's (1772/1880) *Essays on Physiognomy*, which was written in 1772 and reprinted in more than 150 editions by 1940, described in minute detail how to relate facial features to personality traits (e.g., "the nearer the eyebrows are to the eyes, the more earnest, deep, and firm the character," p. 59). Although these ideas strike most people today as ludicrous and bring to mind phrenology, empirical evidence shows that the effects of facial appearance on social outcomes are pervasive. In almost every significant domain of life, attractive people get better outcomes than unattractive people (Hamermesh & Biddle, 1994; Zebrowitz, 1999). The effects of baby-faced appearance are as pervasive as are the effects of attractiveness (Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, baby-faced individuals are less likely to receive severe judicial outcomes than mature-faced individuals (Zebrowitz & McDonald, 1991).

From the structure of the face, people form not only global impressions, but also specific trait impressions (Hassin & Trope, 2000). For example, we showed that inferences of competence, based solely on facial appearance, predicted the outcomes of U.S. congressional elections in 2000, 2002, and 2004 (Todorov, Mandisodza, Goren, & Hall, 2005). Although we measured impressions on a variety of traits, including attractiveness, trustworthiness, and likeability, the trait inference that predicted the election outcomes was competence. Competence was also rated as the most important attribute for a person running for a public office. This finding suggests that person attributes that are important for specific decisions are inferred from facial appearance and influence these decisions.

From both the standard-intuition and the rational-actor points of view, trait inferences from facial appearance should not influence important deliberate decisions. However, to the extent that these inferences occur rapidly and effortlessly, their effects on decisions may be subtle and not subjectively recognized. Using the terms of dual-process theories (Chaiken & Trope, 1999; Kahneman, 2003), we have argued that trait inferences from faces can be characterized as fast, intuitive, unreflective System 1 processes that contrast with slow, effortful, and deliberate System 2 processes (Todorov et al., 2005). We provided preliminary evidence for this proposal by showing that inferences of competence based on 1-s exposure to the faces of the winners and the runners-up for the Senate races sufficed to predict the election outcomes.

In this article, we report a series of studies in which we systematically manipulated the exposure time of faces to further explore the minimal conditions under which people make trait inferences from facial appearance. Research on visual processing has shown that high-level object representations can be constructed very rapidly from visual scenes (Grill-Spector & Kanwisher, 2005; Rousselet, Fabre-Thorpe, & Thorpe, 2002; Thorpe, Fize, & Marlot, 1996). It is possible that inferences about socially significant attributes are also rapidly extracted

Address correspondence to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, e-mail: atodorov@princeton.edu.

"Thin slicing"


PSYCHOLOGICAL SCIENCE

Research Article

First Impressions

Making Up Your Mind After a 100-Ms Exposure to a Face

Janine Willis and Alexander Todorov

Princeton University

ABSTRACT—*People often draw trait inferences from the facial appearance of other people. We investigated the minimal conditions under which people make such inferences. In five experiments, each focusing on a specific trait judgment, we manipulated the exposure time of unfamiliar faces. Judgments made after a 100-ms exposure correlated highly with judgments made in the absence of time constraints, suggesting that this exposure time was sufficient for participants to form an impression. In fact, for all judgments—attractiveness, likeability, trustworthiness, competence, and aggressiveness—increased exposure time did not significantly increase the correlations. When exposure time increased from 100 to 500 ms, participants' judgments became more negative, response times for judgments decreased, and confidence in judgments increased. When exposure time increased from 500 to 1,000 ms, trait judgments and response times did not change significantly (with one exception), but confidence increased for some of the judgments; this result suggests that additional time may simply boost confidence in judgments. However, increased exposure time led to more differentiated person impressions.*

Lavater's (1772/1880) *Essays on Physiognomy*, which was written in 1772 and reprinted in more than 150 editions by 1940, described in minute detail how to relate facial features to personality traits (e.g., "the nearer the eyebrows are to the eyes, the more earnest, deep, and firm the character," p. 59). Although these ideas strike most people today as ludicrous and bring to mind phrenology, empirical evidence shows that the effects of facial appearance on social outcomes are pervasive. In almost every significant domain of life, attractive people get better outcomes than unattractive people (Hamermesh & Biddle, 1994; Zebrowitz, 1999). The effects of baby-faced appearance are as pervasive as are the effects of attractiveness (Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, baby-faced individuals are less likely to receive severe judicial outcomes than mature-faced individuals (Zebrowitz & McDonald, 1991).

From the structure of the face, people form not only global impressions, but also specific trait impressions (Hassin & Trope, 2000). For example, we showed that inferences of competence, based solely on facial appearance, predicted the outcomes of U.S. congressional elections in 2000, 2002, and 2004 (Todorov, Mandisodza, Goren, & Hall, 2005). Although we measured impressions on a variety of traits, including attractiveness, trustworthiness, and likeability, the trait inference that predicted the election outcomes was competence. Competence was also rated as the most important attribute for a person running for a public office. This finding suggests that person attributes that are important for specific decisions are inferred from facial appearance and influence these decisions.

From both the standard-intuition and the rational-actor points of view, trait inferences from facial appearance should not influence important deliberate decisions. However, to the extent that these inferences occur rapidly and effortlessly, their effects on decisions may be subtle and not subjectively recognized. Using the terms of dual-process theories (Chaiken & Trope, 1999; Kahneman, 2003), we have argued that trait inferences from faces can be characterized as fast, intuitive, unreflective System 1 processes that contrast with slow, effortful, and deliberate System 2 processes (Todorov et al., 2005). We provided preliminary evidence for this proposal by showing that inferences of competence based on 1-s exposure to the faces of the winners and the runners-up for the Senate races sufficed to predict the election outcomes.

In this article, we report a series of studies in which we systematically manipulated the exposure time of faces to further explore the minimal conditions under which people make trait inferences from facial appearance. Research on visual processing has shown that high-level object representations can be constructed very rapidly from visual scenes (Grill-Spector & Kanwisher, 2005; Rousselet, Fabre-Thorpe, & Thorpe, 2002; Thorpe, Fize, & Marlot, 1996). It is possible that inferences about socially significant attributes are also rapidly extracted

Address correspondence to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, e-mail: atodorov@princeton.edu.

592 Copyright © 2006 Association for Psychological Science Volume 17—Number 7

"Thin slicing"

This is more **consequential** than you might think!

Inferences of Competence from Faces Predict Election Outcomes

Alexander Todorov,[1,2]* Anesu N. Mandisodza,[1]† Amir Goren,[1] Crystal C. Hall[1]

We show that inferences of competence based solely on facial appearance predicted the outcomes of U.S. congressional elections better than chance (e.g., 68.8% of the Senate races in 2004) and also were linearly related to the margin of victory. These inferences were specific to competence and occurred within a 1-second exposure to the faces of the candidates. The findings suggest that rapid, unreflective trait inferences can contribute to voting choices, which are widely assumed to be based primarily on rational and deliberative considerations.

Inferences of Competence from Faces Predict Election Outcomes

Alexander Todorov,[1,2*] Anesu N. Mandisodza,[1†] Amir Goren,[1] Crystal C. Hall[1]

We show that inferences of competence based solely on facial appearance predicted the outcomes of U.S. congressional elections better than chance (e.g., 68.8% of the Senate races in 2004) and also were linearly related to the margin of victory. These inferences were specific to competence and occurred within a 1-second exposure to the faces of the candidates. The findings suggest that rapid, unreflective trait inferences can contribute to voting choices, which are widely assumed to be based primarily on rational and deliberative considerations.

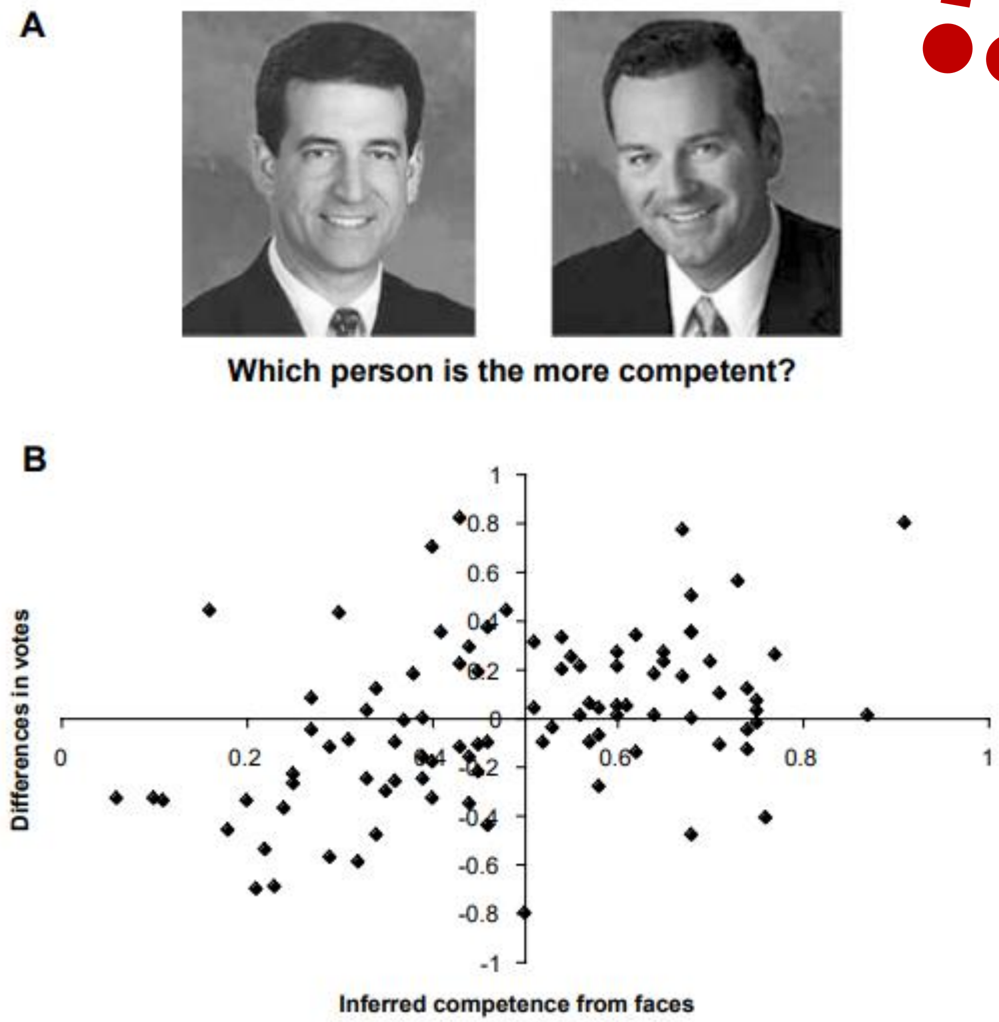

Which of the following faces is more attractive?

Peshek et al., 2011

Axelsson et al., 2018

Figure 3. Averaged images of 16 individuals (eight women) photographed twice in a cross-over design, during experimentally induced (*a*) acute sickness and (*b*) placebo. Images made by Audrey Henderson, MSc, St Andrews University, using Psychomorph. Here, 184 facial landmarks were placed on each image before composites displaying the average shape, colour and texture were created [20].

Axelsson et al., 2018

Also, faces are weird.

Face-inversion effect

Thatcher effect

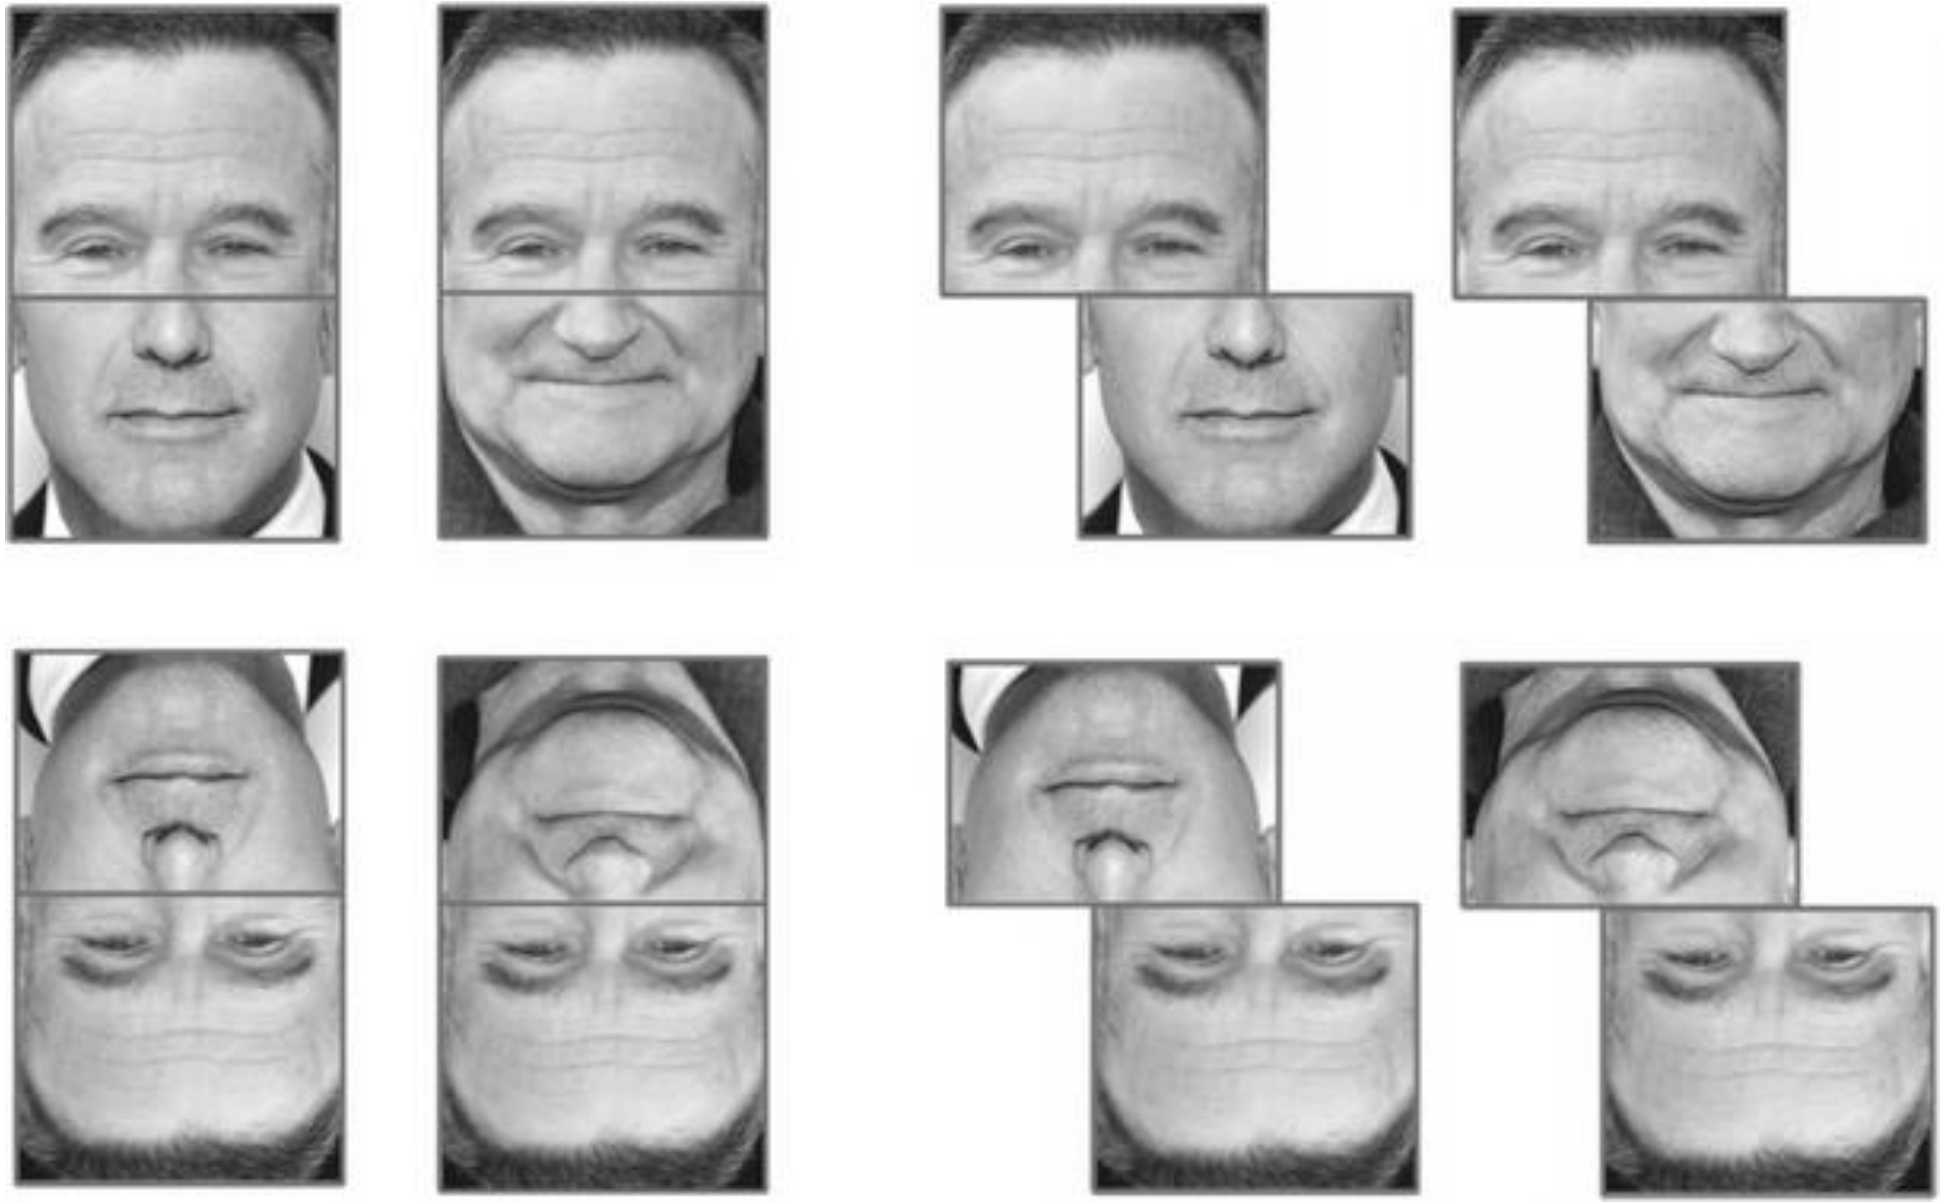

!!

Composite face illusion

Prosopagnosia

Psychonomic Bulletin & Review
2009, 16 (2), 252-257
doi:10.3758/PBR.16.2.252

BRIEF REPORTS

Super-recognizers: People with extraordinary face recognition ability

RICHARD RUSSELL
Harvard University, Cambridge, Massachusetts

BRAD DUCHAINE
University College London, London, England

AND

KEN NAKAYAMA
Harvard University, Cambridge, Massachusetts

We tested 4 people who claimed to have significantly better than ordinary face recognition ability. Exceptional ability was confirmed in each case. On two very different tests of face recognition, all 4 experimental subjects performed beyond the range of control subject performance. They also scored significantly better than average on a perceptual discrimination test with faces. This effect was larger with upright than with inverted faces, and the 4 subjects showed a larger "inversion effect" than did control subjects, who in turn showed a larger inversion effect than did developmental prosopagnosics. This result indicates an association between face recognition ability and the magnitude of the inversion effect. Overall, these "super-recognizers" are about as good at face recognition and perception as developmental prosopagnosics are bad. Our findings demonstrate the existence of people with exceptionally good face recognition ability and show that the range of face recognition and face perception ability is wider than has been previously acknowledged.

Developmental prosopagnosia is a condition marked by exceptionally poor face recognition ability despite normal vision and absence of brain damage or other cognitive deficits. The first case of developmental prosopagnosia (alternately called *congenital* or *hereditary* prosopagnosia) was described in 1976 (McConachie, 1976), and the condition was thought to be very rare (Kress & Daum, 2003). However, recent years have witnessed estimates suggesting that 2%–2.5% of the general population has the condition (Kennerknecht et al., 2006; Nakayama et al., 2006) and rapid growth in the number of cases reported (Barton, Cherkasova, Press, Intriligator, & O'Connor, 2003; Behrmann & Avidan, 2005; Duchaine & Nakayama, 2006b; Le Grand et al., 2006).

The widespread use of terms such as "condition," "disorder," and "impaired" to describe developmental prosopagnosia indicates a prevailing notion of face recognition ability as being either normal (i.e., roughly average) or pathological. The prevalence of this notion may be due in part to the apparent lack of people who are as far above average at face recognition as developmental prosopagnosics are below average. Finding such people would support an alternate notion of a broad distribution of face recognition ability, with (at least some cases of) developmental prosopagnosia representing the lower tail of the distribution.

We have been contacted by several people who, with telling anecdotes, have made persuasive claims to having superior face recognition abilities. The goal of the present study was to evaluate these claims through objective testing. Specifically, we sought to determine whether these people are indeed much better than normal at recognizing faces. A secondary goal was to determine whether they are also better than average at perceiving subtle differences among simultaneously presented faces and to measure their face inversion effect.

Following recent media coverage of our work on developmental prosopagnosia, several people contacted us because they believe that they are the opposite of prosopagnosic—that they are exceptionally *good* at recognizing faces. We have tested 3 of these people: C.S., a 26-year-old female PhD student; C.L., a 40-year-old female homemaker; and J.J., a 36-year-old female municipal employee. We also tested M.R., a 31-year-old male computer programmer. M.R. was referred by his roommate (a developmental prosopagnosic, whom we tested in our laboratory),

R. Russell, rrussell@wjh.harvard.edu

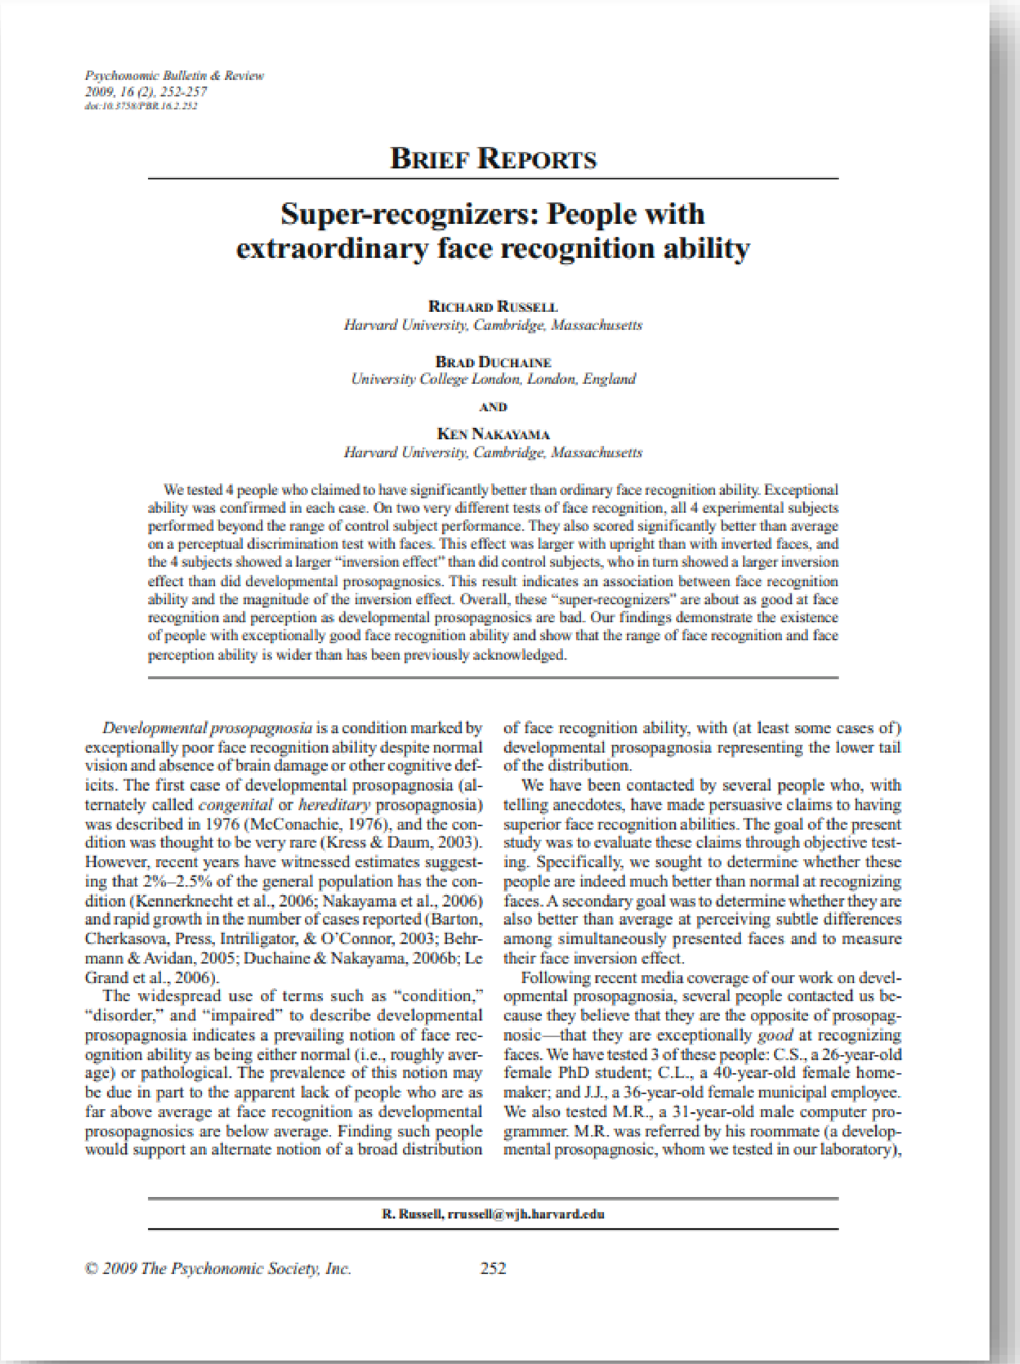

Psychonomic Bulletin & Review
2009, 16 (2), 252-257
doi:10.3758/PBR.16.2.252

Brief Reports

Super-recognizers: People with extraordinary face recognition ability

Richard Russell
Harvard University, Cambridge, Massachusetts

Brad Duchaine
University College London, London, England

and

Ken Nakayama
Harvard University, Cambridge, Massachusetts

We tested 4 people who claimed to have significantly better than ordinary face recognition ability. Exceptional ability was confirmed in each case. On two very different tests of face recognition, all 4 experimental subjects performed beyond the range of control subject performance. They also scored significantly better than average on a perceptual discrimination test with faces. This effect was larger with upright than with inverted faces, and the 4 subjects showed a larger "inversion effect" than did control subjects, who in turn showed a larger inversion effect than did developmental prosopagnosics. This result indicates an association between face recognition ability and the magnitude of the inversion effect. Overall, these "super-recognizers" are about as good at face recognition and perception as developmental prosopagnosics are bad. Our findings demonstrate the existence of people with exceptionally good face recognition ability and show that the range of face recognition and face perception ability is wider than has been previously acknowledged.

Developmental prosopagnosia is a condition marked by exceptionally poor face recognition ability despite normal vision and absence of brain damage or other cognitive deficits. The first case of developmental prosopagnosia (alternately called *congenital* or *hereditary* prosopagnosia) was described in 1976 (McConachie, 1976), and the condition was thought to be very rare (Kress & Daum, 2003). However, recent years have witnessed estimates suggesting that 2%–2.5% of the general population has the condition (Kennerknecht et al., 2006; Nakayama et al., 2006) and rapid growth in the number of cases reported (Barton, Cherkasova, Press, Intriligator, & O'Connor, 2003; Behrmann & Avidan, 2005; Duchaine & Nakayama, 2006b; Le Grand et al., 2006).

The widespread use of terms such as "condition," "disorder," and "impaired" to describe developmental prosopagnosia indicates a prevailing notion of face recognition ability as being either normal (i.e., roughly average) or pathological. The prevalence of this notion may be due in part to the apparent lack of people who are as far above average at face recognition as developmental prosopagnosics are below average. Finding such people would support an alternate notion of a broad distribution of face recognition ability, with (at least some cases of) developmental prosopagnosia representing the lower tail of the distribution.

We have been contacted by several people who, with telling anecdotes, have made persuasive claims to having superior face recognition abilities. The goal of the present study was to evaluate these claims through objective testing. Specifically, we sought to determine whether these people are indeed much better than normal at recognizing faces. A secondary goal was to determine whether they are also better than average at perceiving subtle differences among simultaneously presented faces and to measure their face inversion effect.

Following recent media coverage of our work on developmental prosopagnosia, several people contacted us because they believe that they are the opposite of prosopagnosic—that they are exceptionally *good* at recognizing faces. We have tested 3 of these people: C.S., a 26-year-old female PhD student; C.L., a 40-year-old female homemaker; and J.J., a 36-year-old female municipal employee. We also tested M.R., a 31-year-old male computer programmer. M.R. was referred by his roommate (a developmental prosopagnosic, whom we tested in our laboratory),

R. Russell, rrussell@wjh.harvard.edu

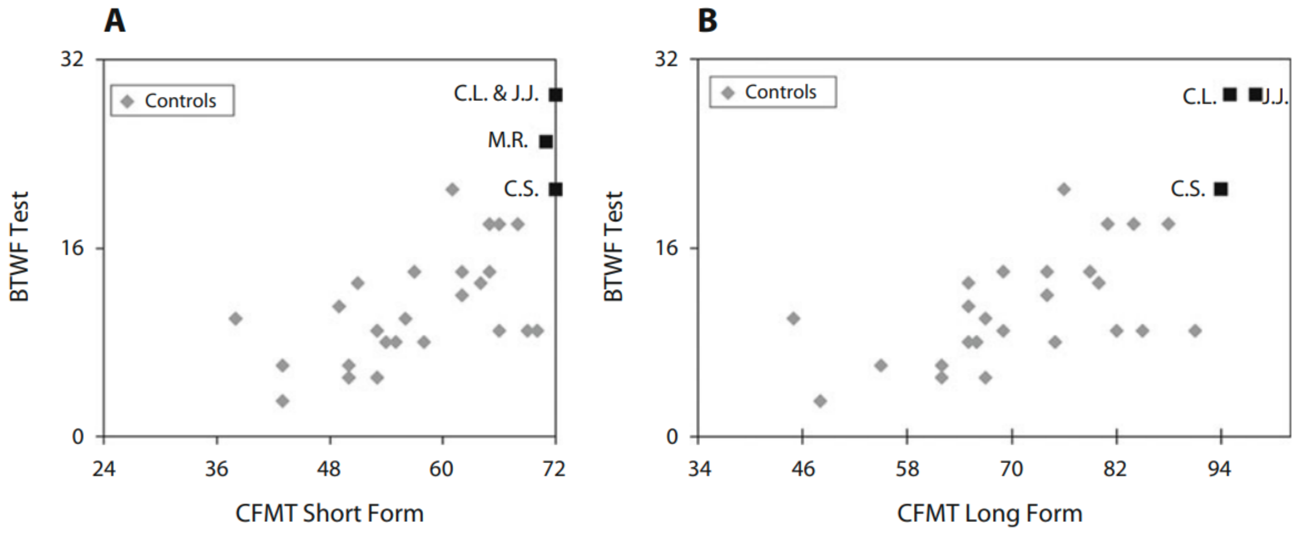

Figure 3. (A) Performance on the Before They Were Famous (BTWF) test plotted against the Cambridge Face Memory Test (CFMT) short form. (B) Performance on the BTWF test plotted against the CFMT long form. The dependent variable in both tests is the number correct. Minimum values on each axis represent chance performance on the given task; maximum performance is 56 on the BTWF test, 72 on the CFMT short form, and 102 on the CFMT long form.

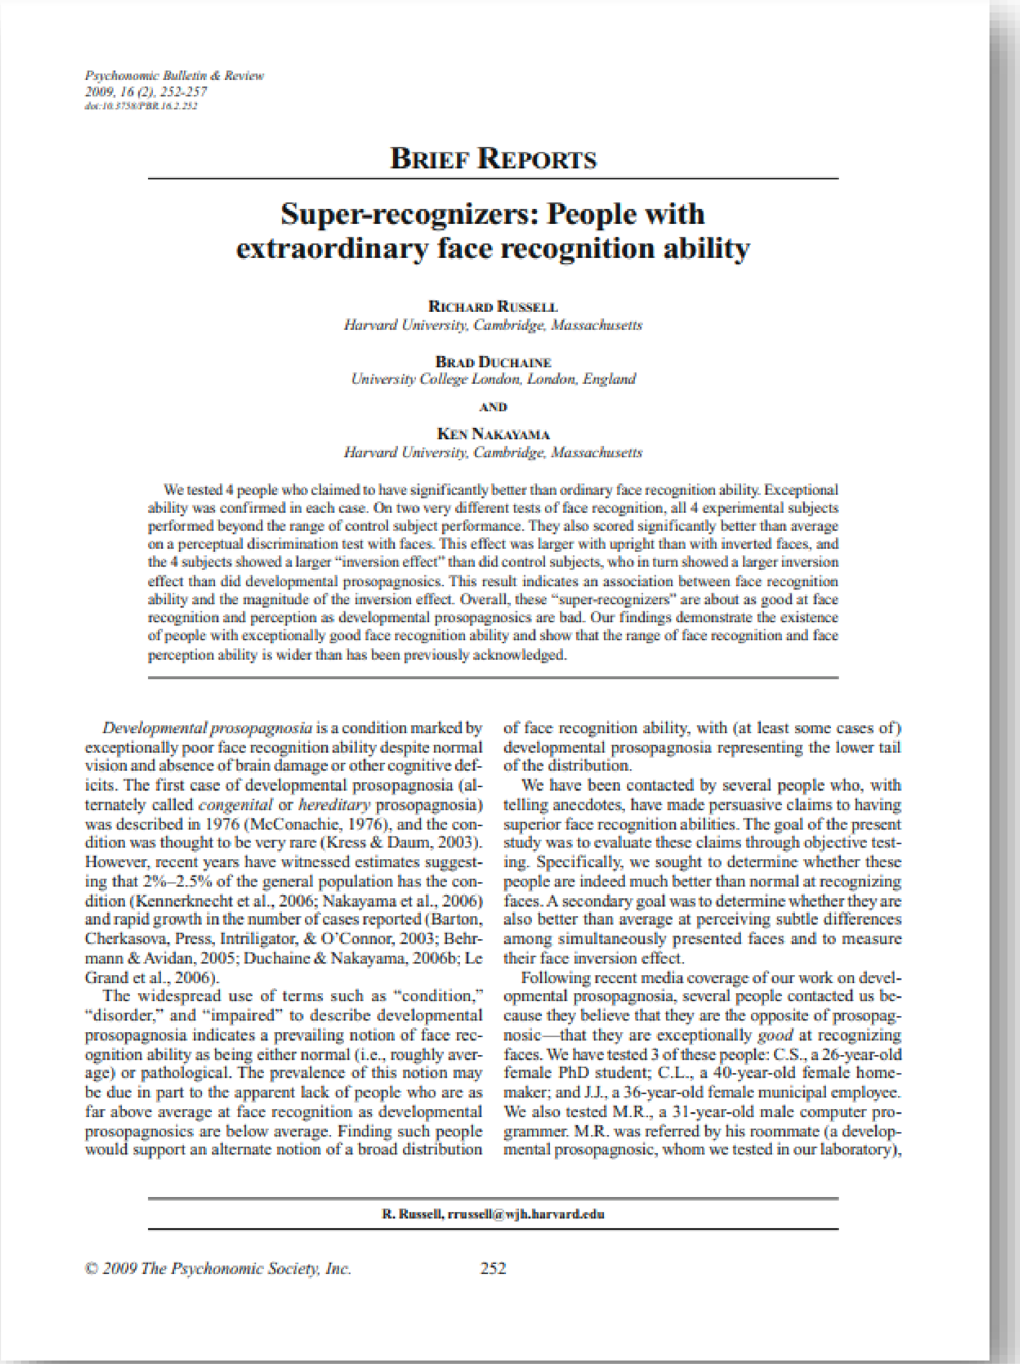

Psychonomic Bulletin & Review
2009, 16 (2), 252-257
doi:10.3758/PBR.16.2.252

BRIEF REPORTS

Super-recognizers: People with extraordinary face recognition ability

RICHARD RUSSELL
Harvard University, Cambridge, Massachusetts

BRAD DUCHAINE
University College London, London, England

AND

KEN NAKAYAMA
Harvard University, Cambridge, Massachusetts

We tested 4 people who claimed to have significantly better than ordinary face recognition ability. Exceptional ability was confirmed in each case. On two very different tests of face recognition, all 4 experimental subjects performed beyond the range of control subject performance. They also scored significantly better than average on a perceptual discrimination test with faces. This effect was larger with upright than with inverted faces, and the 4 subjects showed a larger "inversion effect" than did control subjects, who in turn showed a larger inversion effect than did developmental prosopagnosics. This result indicates an association between face recognition ability and the magnitude of the inversion effect. Overall, these "super-recognizers" are about as good at face recognition and perception as developmental prosopagnosics are bad. Our findings demonstrate the existence of people with exceptionally good face recognition ability and show that the range of face recognition and face perception ability is wider than has been previously acknowledged.

Developmental prosopagnosia is a condition marked by exceptionally poor face recognition ability despite normal vision and absence of brain damage or other cognitive deficits. The first case of developmental prosopagnosia (alternately called *congenital* or *hereditary* prosopagnosia) was described in 1976 (McConachie, 1976), and the condition was thought to be very rare (Kress & Daum, 2003). However, recent years have witnessed estimates suggesting that 2%–2.5% of the general population has the condition (Kennerknecht et al., 2006; Nakayama et al., 2006) and rapid growth in the number of cases reported (Barton, Cherkasova, Press, Intriligator, & O'Connor, 2003; Behrmann & Avidan, 2005; Duchaine & Nakayama, 2006b; Le Grand et al., 2006).

The widespread use of terms such as "condition," "disorder," and "impaired" to describe developmental prosopagnosia indicates a prevailing notion of face recognition ability as being either normal (i.e., roughly average) or pathological. The prevalence of this notion may be due in part to the apparent lack of people who are as far above average at face recognition as developmental prosopagnosics are below average. Finding such people would support an alternate notion of a broad distribution of face recognition ability, with (at least some cases of) developmental prosopagnosia representing the lower tail of the distribution.

We have been contacted by several people who, with telling anecdotes, have made persuasive claims to having superior face recognition abilities. The goal of the present study was to evaluate these claims through objective testing. Specifically, we sought to determine whether these people are indeed much better than normal at recognizing faces. A secondary goal was to determine whether they are also better than average at perceiving subtle differences among simultaneously presented faces and to measure their face inversion effect.

Following recent media coverage of our work on developmental prosopagnosia, several people contacted us because they believe that they are the opposite of prosopagnosic—that they are exceptionally *good* at recognizing faces. We have tested 3 of these people: C.S., a 26-year-old female PhD student; C.L., a 40-year-old female homemaker; and J.J., a 36-year-old female municipal employee. We also tested M.R., a 31-year-old male computer programmer. M.R. was referred by his roommate (a developmental prosopagnosic, whom we tested in our laboratory),

R. Russell, rrussell@wjh.harvard.edu

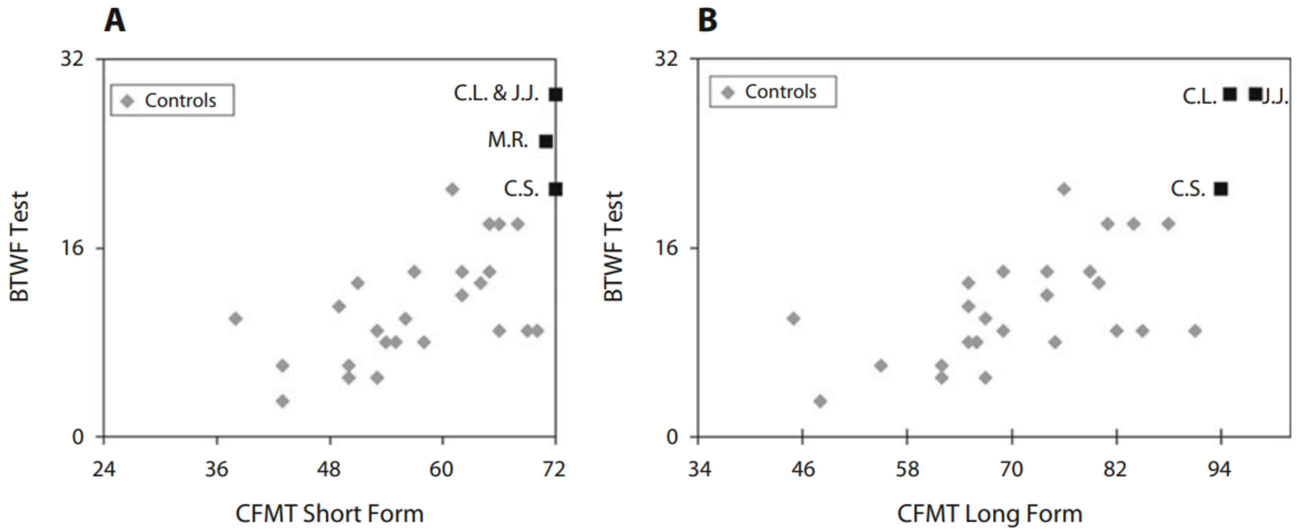

Figure 3. (A) Performance on the Before They Were Famous (BTWF) test plotted against the Cambridge Face Memory Test (CFMT) short form. (B) Performance on the BTWF test plotted against the CFMT long form. The dependent variable in both tests is the number correct. Minimum values on each axis represent chance performance on the given task; maximum performance is 56 on the BTWF test, 72 on the CFMT short form, and 102 on the CFMT long form.

Continuous variation!

!!

Perception, 2003, volume 32, pages 285 – 293

DOI:10.1068/p5027

The role of eyebrows in face recognition

Javid Sadr¶, Izzat Jarudi, Pawan Sinha¶
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology,
45 Carleton Street, E25-201, Cambridge, MA 02142, USA; e-mail: sadr@mit.edu; sinha@ai.mit.edu
Received 23 April 2002, in revised form 4 November 2002; published online 7 March 2003

Abstract. A fundamental challenge in face recognition lies in determining which facial characteristics are important in the identification of faces. Several studies have indicated the significance of certain facial features in this regard, particularly internal ones such as the eyes and mouth. Surprisingly, however, one rather prominent facial feature has received little attention in this domain: the eyebrows. Past work has examined the role of eyebrows in emotional expression and nonverbal communication, as well as in facial aesthetics and sexual dimorphism. However, it has not been made clear whether the eyebrows play an important role in the identification of faces. Here, we report experimental results which suggest that for face recognition the eyebrows may be at least as influential as the eyes. Specifically, we find that the absence of eyebrows in familiar faces leads to a very large and significant disruption in recognition performance. In fact, a significantly greater decrement in face recognition is observed in the absence of eyebrows than in the absence of eyes. These results may have important implications for our understanding of the mechanisms of face recognition in humans as well as for the development of artificial face-recognition systems.

1 Introduction

Evolutionary history has seen a considerable reduction in the amount of hair on the human face (McNeill 2000; figure 1). As such, the presence of the eyebrow might seem a curiosity. Do the eyebrows in fact serve a useful purpose or are they merely an evolutionary vestige? While it is perhaps true that eyebrows may provide the eyes modest protection against such things as rain and perspiration, it is perhaps more relevant that the eyebrows also appear to serve a number of functions that are more visual in nature. Eyebrows may serve as high-contrast lines that give the appearance of the brow greater clarity and emphasis, and their associated musculature allows for sophisticated, often involuntary gestures that may be discerned from a relatively large distance. As such, the eyebrows appear to play an important role in the expression of emotions and in the production of other social signals, and they may also contribute to the sexual dimorphism

Figure 1. Overall, humans have relatively little facial hair as compared to other primates; the conspicuous presence of the hair forming the eyebrows is somewhat intriguing. A number of selection pressures may be responsible for the development and persistence of the eyebrows over the course of primate and especially human evolution. A color version of this figure can be seen at http://www.perceptionweb.com/misc/p5027/.

¶ Authors to whom correspondence and requests for materials should be addressed.

Perception, 2003, volume 32, pages 285–293

DOI:10.1068/p5027

The role of eyebrows in face recognition

Javid Sadr¶, Izzat Jarudi, Pawan Sinha¶
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology,
45 Carleton Street, E25-201, Cambridge, MA 02142, USA; e-mail: sadr@mit.edu; sinha@ai.mit.edu
Received 23 April 2002, in revised form 4 November 2002; published online 7 March 2003

Abstract. A fundamental challenge in face recognition lies in determining which facial characteristics are important in the identification of faces. Several studies have indicated the significance of certain facial features in this regard, particularly internal ones such as the eyes and mouth. Surprisingly, however, one rather prominent facial feature has received little attention in this domain: the eyebrows. Past work has examined the role of eyebrows in emotional expression and nonverbal communication, as well as in facial aesthetics and sexual dimorphism. However, it has not been made clear whether the eyebrows play an important role in the identification of faces. Here, we report experimental results which suggest that for face recognition the eyebrows may be at least as influential as the eyes. Specifically, we find that the absence of eyebrows in familiar faces leads to a very large and significant disruption in recognition performance. In fact, a significantly greater decrement in face recognition is observed in the absence of eyebrows than in the absence of eyes. These results may have important implications for our understanding of the mechanisms of face recognition in humans as well as for the development of artificial face-recognition systems.

1 Introduction

Evolutionary history has seen a considerable reduction in the amount of hair on the human face (McNeill 2000; figure 1). As such, the presence of the eyebrow might seem a curiosity. Do the eyebrows in fact serve a useful purpose or are they merely an evolutionary vestige? While it is perhaps true that eyebrows may provide the eyes modest protection against such things as rain and perspiration, it is perhaps more relevant that the eyebrows also appear to serve a number of functions that are more visual in nature. Eyebrows may serve as high-contrast lines that give the appearance of the brow greater clarity and emphasis, and their associated musculature allows for sophisticated, often involuntary gestures that may be discerned from a relatively large distance. As such, the eyebrows appear to play an important role in the expression of emotions and in the production of other social signals, and they may also contribute to the sexual dimorphism

Figure 1. Overall, humans have relatively little facial hair as compared to other primates; the conspicuous presence of the hair forming the eyebrows is somewhat intriguing. A number of selection pressures may be responsible for the development and persistence of the eyebrows over the course of primate and especially human evolution. A color version of this figure can be seen at http://www.perceptionweb.com/misc/p5027/.

¶ Authors to whom correspondence and requests for materials should be addressed.

Perception, 2003, volume 32, pages 285 – 293

DOI:10.1068/p5027

The role of eyebrows in face recognition

Javid Sadr¶, Izzat Jarudi, Pawan Sinha¶
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 45 Carleton Street, E25-201, Cambridge, MA 02142, USA; e-mail: sadr@mit.edu; sinha@ai.mit.edu
Received 23 April 2002, in revised form 4 November 2002; published online 7 March 2003

Abstract. A fundamental challenge in face recognition lies in determining which facial characteristics are important in the identification of faces. Several studies have indicated the significance of certain facial features in this regard, particularly internal ones such as the eyes and mouth. Surprisingly, however, one rather prominent facial feature has received little attention in this domain: the eyebrows. Past work has examined the role of eyebrows in emotional expression and nonverbal communication, as well as in facial aesthetics and sexual dimorphism. However, it has not been made clear whether the eyebrows play an important role in the identification of faces. Here, we report experimental results which suggest that for face recognition the eyebrows may be at least as influential as the eyes. Specifically, we find that the absence of eyebrows in familiar faces leads to a very large and significant disruption in recognition performance. In fact, a significantly greater decrement in face recognition is observed in the absence of eyebrows than in the absence of eyes. These results may have important implications for our understanding of the mechanisms of face recognition in humans as well as for the development of artificial face-recognition systems.

1 Introduction

Evolutionary history has seen a considerable reduction in the amount of hair on the human face (McNeill 2000; figure 1). As such, the presence of the eyebrow might seem a curiosity. Do the eyebrows in fact serve a useful purpose or are they merely an evolutionary vestige? While it is perhaps true that eyebrows may provide the eyes modest protection against such things as rain and perspiration, it is perhaps more relevant that the eyebrows also appear to serve a number of functions that are more visual in nature. Eyebrows may serve as high-contrast lines that give the appearance of the brow greater clarity and emphasis, and their associated musculature allows for sophisticated, often involuntary gestures that may be discerned from a relatively large distance. As such, the eyebrows appear to play an important role in the expression of emotions and in the production of other social signals, and they may also contribute to the sexual dimorphism

Figure 1. Overall, humans have relatively little facial hair as compared to other primates; the conspicuous presence of the hair forming the eyebrows is somewhat intriguing. A number of selection pressures may be responsible for the development and persistence of the eyebrows over the course of primate and especially human evolution. A color version of this figure can be seen at http://www.perceptionweb.com/misc/p5027/.

¶ Authors to whom correspondence and requests for materials should be addressed.

Perception, 2003, volume 32, pages 285 – 293

DOI:10.1068/p5027

The role of eyebrows in face recognition

Javid Sadr¶, Izzat Jarudi, Pawan Sinha¶
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology,
45 Carleton Street, E25-201, Cambridge, MA 02142, USA; e-mail: sadr@mit.edu; sinha@ai.mit.edu
Received 23 April 2002, in revised form 4 November 2002; published online 7 March 2003

Abstract. A fundamental challenge in face recognition lies in determining which facial characteristics are important in the identification of faces. Several studies have indicated the significance of certain facial features in this regard, particularly internal ones such as the eyes and mouth. Surprisingly, however, one rather prominent facial feature has received little attention in this domain: the eyebrows. Past work has examined the role of eyebrows in emotional expression and nonverbal communication, as well as in facial aesthetics and sexual dimorphism. However, it has not been made clear whether the eyebrows play an important role in the identification of faces. Here, we report experimental results which suggest that for face recognition the eyebrows may be at least as influential as the eyes. Specifically, we find that the absence of eyebrows in familiar faces leads to a very large and significant disruption in recognition performance. In fact, a significantly greater decrement in face recognition is observed in the absence of eyebrows than in the absence of eyes. These results may have important implications for our understanding of the mechanisms of face recognition in humans as well as for the development of artificial face-recognition systems.

1 Introduction

Evolutionary history has seen a considerable reduction in the amount of hair on the human face (McNeill 2000; figure 1). As such, the presence of the eyebrow might seem a curiosity. Do the eyebrows in fact serve a useful purpose or are they merely an evolutionary vestige? While it is perhaps true that eyebrows may provide the eyes modest protection against such things as rain and perspiration, it is perhaps more relevant that the eyebrows also appear to serve a number of functions that are more visual in nature. Eyebrows may serve as high-contrast lines that give the appearance of the brow greater clarity and emphasis, and their associated musculature allows for sophisticated, often involuntary gestures that may be discerned from a relatively large distance. As such, the eyebrows appear to play an important role in the expression of emotions and in the production of other social signals, and they may also contribute to the sexual dimorphism

Figure 1. Overall, humans have relatively little facial hair as compared to other primates; the conspicuous presence of the hair forming the eyebrows is somewhat intriguing. A number of selection pressures may be responsible for the development and persistence of the eyebrows over the course of primate and especially human evolution. A color version of this figure can be seen at http://www.perceptionweb.com/misc/p5027/.

¶ Authors to whom correspondence and requests for materials should be addressed.

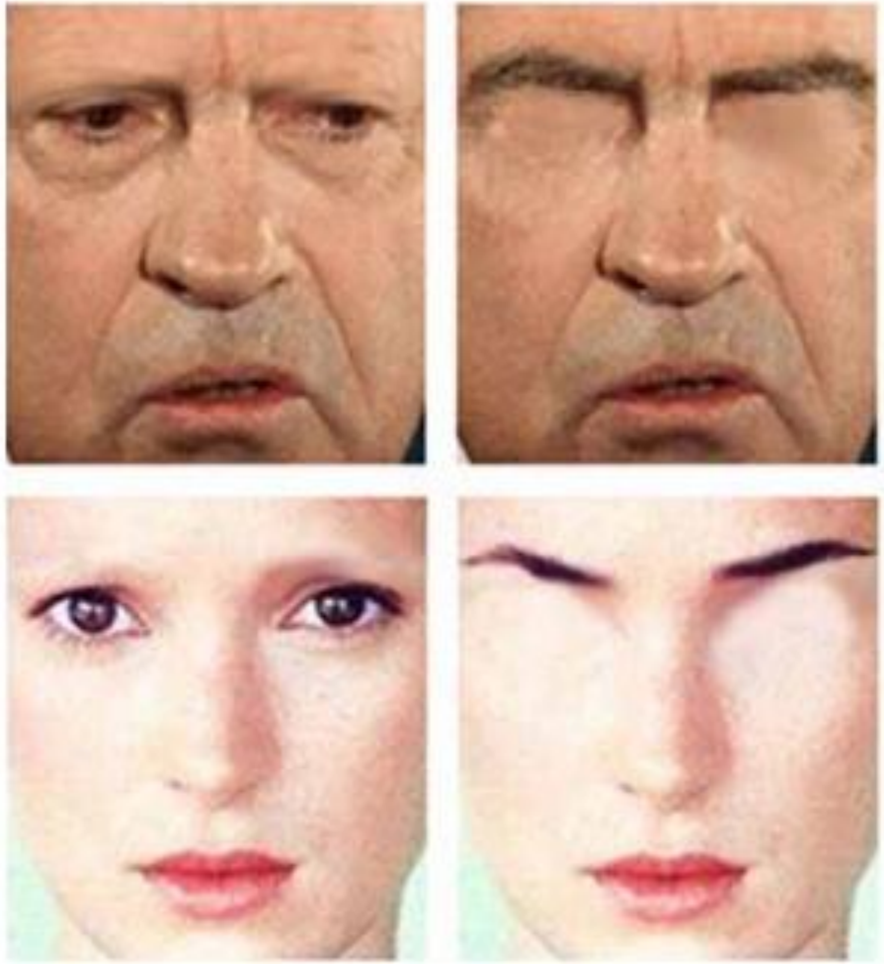

Perception, 2003, volume 32, pages 285 – 293

DOI:10.1068/p5027

The role of eyebrows in face recognition

Javid Sadr¶, Izzat Jarudi, Pawan Sinha¶
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology,
45 Carleton Street, E25-201, Cambridge, MA 02142, USA; e-mail: sadr@mit.edu; sinha@ai.mit.edu
Received 23 April 2002, in revised form 4 November 2002; published online 7 March 2003

Abstract. A fundamental challenge in face recognition lies in determining which facial characteristics are important in the identification of faces. Several studies have indicated the significance of certain facial features in this regard, particularly internal ones such as the eyes and mouth. Surprisingly, however, one rather prominent facial feature has received little attention in this domain: the eyebrows. Past work has examined the role of eyebrows in emotional expression and nonverbal communication, as well as in facial aesthetics and sexual dimorphism. However, it has not been made clear whether the eyebrows play an important role in the identification of faces. Here, we report experimental results which suggest that for face recognition the eyebrows may be at least as influential as the eyes. Specifically, we find that the absence of eyebrows in familiar faces leads to a very large and significant disruption in recognition performance. In fact, a significantly greater decrement in face recognition is observed in the absence of eyebrows than in the absence of eyes. These results may have important implications for our understanding of the mechanisms of face recognition in humans as well as for the development of artificial face-recognition systems.

1 Introduction

Evolutionary history has seen a considerable reduction in the amount of hair on the human face (McNeill 2000; figure 1). As such, the presence of the eyebrow might seem a curiosity. Do the eyebrows in fact serve a useful purpose or are they merely an evolutionary vestige? While it is perhaps true that eyebrows may provide the eyes modest protection against such things as rain and perspiration, it is perhaps more relevant that the eyebrows also appear to serve a number of functions that are more visual in nature. Eyebrows may serve as high-contrast lines that give the appearance of the brow greater clarity and emphasis, and their associated musculature allows for sophisticated, often involuntary gestures that may be discerned from a relatively large distance. As such, the eyebrows appear to play an important role in the expression of emotions and in the production of other social signals, and they may also contribute to the sexual dimorphism

Figure 1. Overall, humans have relatively little facial hair as compared to other primates; the conspicuous presence of the hair forming the eyebrows is somewhat intriguing. A number of selection pressures may be responsible for the development and persistence of the eyebrows over the course of primate and especially human evolution. A color version of this figure can be seen at http://www.perceptionweb.com/misc/p5027/.

¶ Authors to whom correspondence and requests for materials should be addressed.

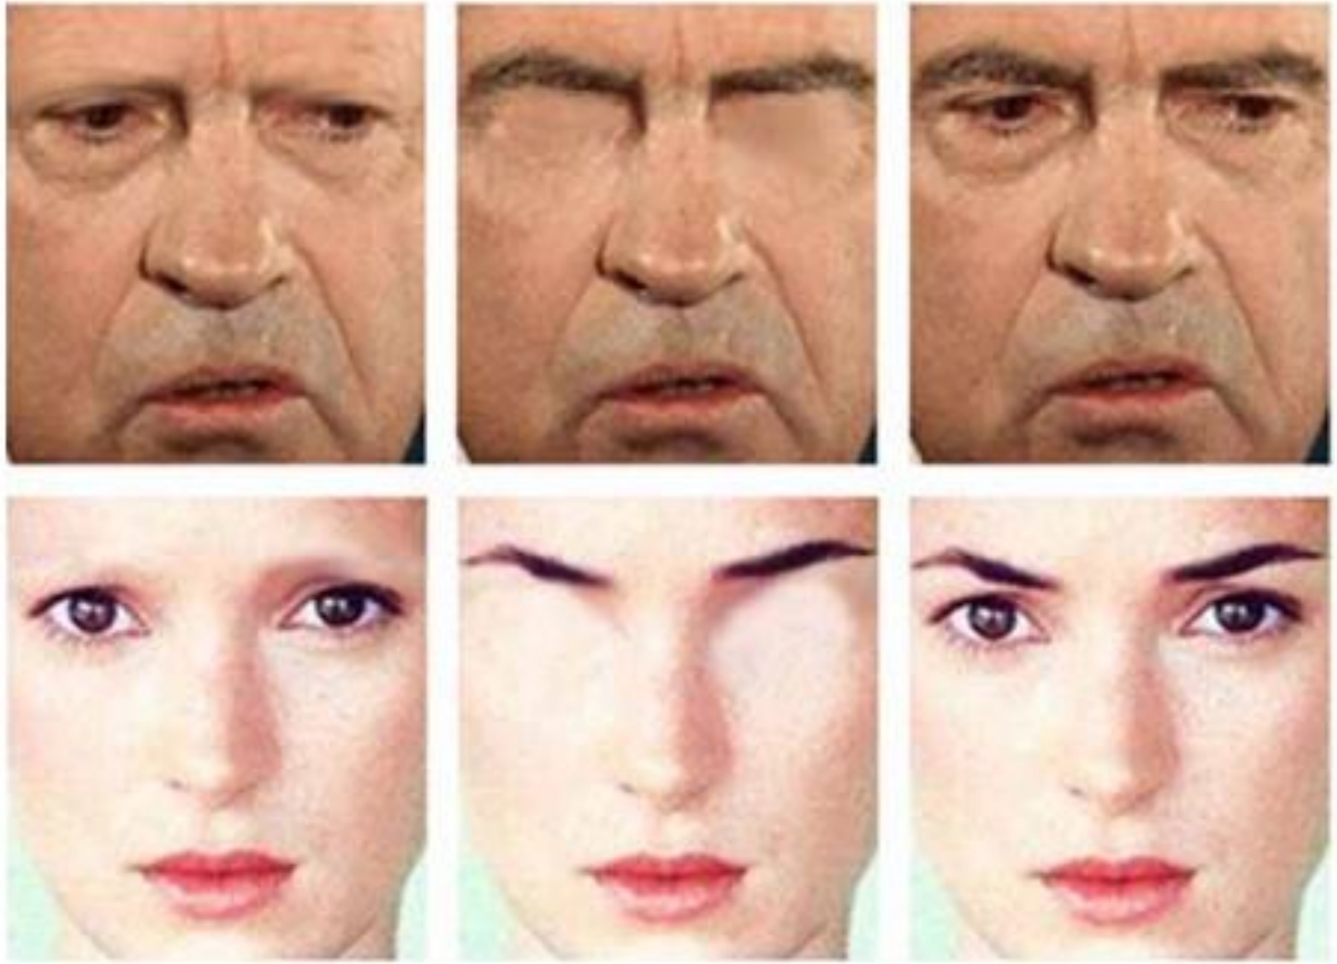

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B

Phil. Trans. R. Soc. B (2011) **366**, 1702–1725
doi:10.1098/rstb.2010.0360

Review

Visual adaptation and face perception

Michael A. Webster[1,*] and Donald I. A. MacLeod[2]

[1]*Department of Psychology, University of Nevada, Reno, NV 89557, USA*
[2]*Department of Psychology, University of California, San Diego, La Jolla, CA 92093, USA*

The appearance of faces can be strongly affected by the characteristics of faces viewed previously. These perceptual after-effects reflect processes of sensory adaptation that are found throughout the visual system, but which have been considered only relatively recently in the context of higher level perceptual judgements. In this review, we explore the consequences of adaptation for human face perception, and the implications of adaptation for understanding the neural-coding schemes underlying the visual representation of faces. The properties of face after-effects suggest that they, in part, reflect response changes at high and possibly face-specific levels of visual processing. Yet, the form of the after-effects and the norm-based codes that they point to show many parallels with the adaptations and functional organization that are thought to underlie the encoding of perceptual attributes like colour. The nature and basis for human colour vision have been studied extensively, and we draw on ideas and principles that have been developed to account for norms and normalization in colour vision to consider potential similarities and differences in the representation and adaptation of faces.

Keywords: face perception; adaptation; colour vision; after-effects; neural coding

1. VISUAL ADAPTATION AND THE PERCEPTION OF FACES

Stare carefully at the cross in the image in figure 1 for a minute or so, and then close your eyes. A clear image of a face will appear after a few seconds. (A popular version of this illusion uses an image of Christ, so we have tried it with a devil.) The phantom image is a negative afterimage of the original picture, and arises because each part of the retina adjusts its sensitivity to the local light level in the original picture. Thus, when a uniform field of dim light is transmitted by the closed eyelids, cells that were previously exposed to dark (or light) regions will respond more (or less). Note that the afterimage is also much more recognizable as a face compared with the original, in part because it has the correct brightness polarity (e.g. dark eyes). This polarity-specific difference [1] is one of many examples that have been used to argue that face perception may depend on specialized processes that are distinct from the mechanisms mediating object recognition [2]. However, while the visual coding of faces may depend on face-specific pathways, the principles underlying how these processes are organized and calibrated may be very general. In particular, the principles of sensitivity regulation that give rise to light adaptation in the retina are likely to be manifest at all stages of visual coding, and may be fundamentally important to all perceptual analyses. Thus, just as adapting to different light levels can affect the perception of brightness, adaptation to different faces can affect the appearance of facial attributes, and in both cases these sensitivity changes may shape the nature of visual coding in functionally similar ways.

In this review, we consider the properties and implications of adaptation effects in face perception. Although it may not be obvious, the appearance of faces does depend strongly on the viewing context, and thus the same face can look dramatically different depending on which faces an observer is previously exposed to. These face after-effects offer a window into the processes and dynamics of face perception and have now been studied extensively. Here, we consider what these studies have revealed about human face perception. But before examining this, it is worth emphasizing that it is important to understand how face perception and recognition are affected by perceptual adaptation for a number of reasons. First, faces are arguably the most important social stimuli for humans and the primary means by which we perceive the identity and state of conspecifics. Any process that modulates that perception thus has profound implications for understanding human perception and behaviour, and these consequences are important regardless of the basis of the adaptation. That is, how our judgements of such fundamental attributes as identity, expression, fitness or beauty are affected by the diet of faces we encounter is essential to understanding the psychology of face perception, whether those context effects result from changes in sensitivity or criterion and whether they reflect high-level and face-specific representations or low-level and thus generic response changes in the visual system. Second, a central focus of research in face perception has been to understand how information

* Author for correspondence (mwebster@unr.edu).

One contribution of 10 to a Theme Issue 'Face perception: social, neuropsychological and comparative perspectives'.


PSYCHOLOGICAL SCIENCE

Research Article

First Impressions

Making Up Your Mind After a 100-Ms Exposure to a Face

Janine Willis and Alexander Todorov

Princeton University

ABSTRACT—*People often draw trait inferences from the facial appearance of other people. We investigated the minimal conditions under which people make such inferences. In five experiments, each focusing on a specific trait judgment, we manipulated the exposure time of unfamiliar faces. Judgments made after a 100-ms exposure correlated highly with judgments made in the absence of time constraints, suggesting that this exposure time was sufficient for participants to form an impression. In fact, for all judgments—attractiveness, likeability, trustworthiness, competence, and aggressiveness—increased exposure time did not significantly increase the correlations. When exposure time increased from 100 to 500 ms, participants' judgments became more negative, response times for judgments decreased, and confidence in judgments increased. When exposure time increased from 500 to 1,000 ms, trait judgments and response times did not change significantly (with one exception), but confidence increased for some of the judgments; this result suggests that additional time may simply boost confidence in judgments. However, increased exposure time led to more differentiated person impressions.*

Lavater's (1772/1880) *Essays on Physiognomy*, which was written in 1772 and reprinted in more than 150 editions by 1940, described in minute detail how to relate facial features to personality traits (e.g., "the nearer the eyebrows are to the eyes, the more earnest, deep, and firm the character," p. 59). Although these ideas strike most people today as ludicrous and bring to mind phrenology, empirical evidence shows that the effects of facial appearance on social outcomes are pervasive. In almost every significant domain of life, attractive people get better outcomes than unattractive people (Hamermesh & Biddle, 1994; Zebrowitz, 1999). The effects of baby-faced appearance are as pervasive as are the effects of attractiveness (Montepare & Zebrowitz, 1998; Zebrowitz, 1999). For example, baby-faced individuals are less likely to receive severe judicial outcomes than mature-faced individuals (Zebrowitz & McDonald, 1991).

From the structure of the face, people form not only global impressions, but also specific trait impressions (Hassin & Trope, 2000). For example, we showed that inferences of competence, based solely on facial appearance, predicted the outcomes of U.S. congressional elections in 2000, 2002, and 2004 (Todorov, Mandisodza, Goren, & Hall, 2005). Although we measured impressions on a variety of traits, including attractiveness, trustworthiness, and likeability, the trait inference that predicted the election outcomes was competence. Competence was also rated as the most important attribute for a person running for a public office. This finding suggests that person attributes that are important for specific decisions are inferred from facial appearance and influence these decisions.

From both the standard-intuition and the rational-actor points of view, trait inferences from facial appearance should not influence important deliberate decisions. However, to the extent that these inferences occur rapidly and effortlessly, their effects on decisions may be subtle and not subjectively recognized. Using the terms of dual-process theories (Chaiken & Trope, 1999; Kahneman, 2003), we have argued that trait inferences from faces can be characterized as fast, intuitive, unreflective System 1 processes that contrast with slow, effortful, and deliberate System 2 processes (Todorov et al., 2005). We provided preliminary evidence for this proposal by showing that inferences of competence based on 1-s exposure to the faces of the winners and the runners-up for the Senate races sufficed to predict the election outcomes.

In this article, we report a series of studies in which we systematically manipulated the exposure time of faces to further explore the minimal conditions under which people make trait inferences from facial appearance. Research on visual processing has shown that high-level object representations can be constructed very rapidly from visual scenes (Grill-Spector & Kanwisher, 2005; Rousselet, Fabre-Thorpe, & Thorpe, 2002; Thorpe, Fize, & Marlot, 1996). It is possible that inferences about socially significant attributes are also rapidly extracted

Address correspondence to Alexander Todorov, Department of Psychology, Green Hall, Princeton University, Princeton, NJ 08544-1010, e-mail: atodorov@princeton.edu.

attractiveness

likeability

trustworthiness

competence

aggressiveness

One more thing:
Watch carefully.

Hollow face illusion

Food for thought:

Food for thought:

1. For number we talked about domain-specific vs. domain-general coding. What about faces?

2. Should we believe that face adaptation is genuine and/or like other kinds of adaptation? Why or why not?

3. Are our judgments about faces accurate?


ILLUSIONS OF SOUND

Illusions of faces: Addendum

PLOS COMPUTATIONAL BIOLOGY

RESEARCH ARTICLE

Human-like face pareidolia emerges in deep neural networks optimized for face and object recognition

Pranjul Gupta[1*], **Katharina Dobs**[1,2]

1 Department of Experimental Psychology, Justus Liebig University Giessen, Giessen, Germany, 2 Center for Mind, Brain, and Behavior, Universities of Marburg, Giessen and Darmstadt, Marburg, Germany

* pgupta013@gmail.com

Citation: Gupta P, Dobs K (2025) Human-like face pareidolia emerges in deep neural networks optimized for face and object recognition. PLoS Comput Biol 21(1): e1012751. https://doi.org/10.1371/journal.pcbi.1012751

Editor: Chris I. Baker, National Institute of Mental Health and Neuro Sciences, UNITED STATES OF AMERICA

Received: July 03, 2023

Accepted: December 24, 2024

Published: January 27, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: https://doi.org/10.1371/journal.pcbi.1012751

Data availability statement: The experimental stimuli set and MEG RDMs used in this study

Abstract

The human visual system possesses a remarkable ability to detect and process faces across diverse contexts, including the phenomenon of face pareidolia—seeing faces in inanimate objects. Despite extensive research, it remains unclear why the visual system employs such broadly tuned face detection capabilities. We hypothesized that face pareidolia results from the visual system's optimization for recognizing both faces and objects. To test this hypothesis, we used task-optimized deep convolutional neural networks (CNNs) and evaluated their alignment with human behavioral signatures and neural responses, measured via magnetoencephalography (MEG), related to pareidolia processing. Specifically, we trained CNNs on tasks involving combinations of face identification, face detection, object categorization, and object detection. Using representational similarity analysis, we found that CNNs that included object categorization in their training tasks represented pareidolia faces, real faces, and matched objects more similarly to neural responses than those that did not. Although these CNNs showed similar overall alignment with neural data, a closer examination of their internal representations revealed that specific training tasks had distinct effects on how pareidolia faces were represented across layers. Finally, interpretability methods revealed that only a CNN trained for both face identification and object categorization relied on face-like features—such as 'eyes'—to classify pareidolia stimuli as faces, mirroring findings in human perception. Our results suggest that human-like face pareidolia may emerge from the visual system's optimization for face identification within the context of generalized object categorization.

Author summary

Have you ever seen a face in an inanimate object, like the foam in your coffee? This phenomenon, known as face pareidolia, demonstrates our brain's remarkable ability to detect faces even when they're not there. But why does this happen? Our study explores

Illusory faces are more likely to be perceived as male than female

Susan G. Wardle[a,1,2], Sanika Paranjape[a,1], Jessica Taubert[b], and Chris I. Baker[a]

[a]Section on Learning and Plasticity, Laboratory of Brain & Cognition, National Institute of Mental Health, Bethesda, MD 20814; and [b]Section on Neurocircuitry, Laboratory of Brain & Cognition, National Institute of Mental Health, Bethesda, MD 20814

Edited by Susan Fiske, Psychology Department, Princeton University, Princeton, NJ; received October 1, 2021; accepted November 17, 2021

Despite our fluency in reading human faces, sometimes we mistakenly perceive illusory faces in objects, a phenomenon known as face pareidolia. Although illusory faces share some neural mechanisms with real faces, it is unknown to what degree pareidolia engages higher-level social perception beyond the detection of a face. In a series of large-scale behavioral experiments (n_{total} = 3,815 adults), we found that illusory faces in inanimate objects are readily perceived to have a specific emotional expression, age, and gender. Most strikingly, we observed a strong bias to perceive illusory faces as male rather than female. This male bias could not be explained by preexisting semantic or visual gender associations with the objects, or by visual features in the images. Rather, this robust bias in the perception of gender for illusory faces reveals a cognitive bias arising from a broadly tuned face evaluation system in which minimally viable face percepts are more likely to be perceived as male.

face perception | gender | bias | pareidolia | face evaluation

Human faces convey a rich amount of social information beyond their identity (1–3). We are able to rapidly evaluate the age (4), gender (5, 6), and emotional expression (7) of the faces of individuals, even if they are not known to us, in addition to more abstract traits, such as trustworthiness and aggressiveness (8, 9). Although these judgements are based on visual information, biases have been identified that suggest that both perceptual and cognitive factors are involved in face evaluation (10–13). For example, people tend to judge faces as closer to their own age (10, 13), and damage to the amygdala is associated with perceiving unfamiliar faces as more trustworthy and approachable (12). Biases in face perception have important implications for understanding the neural processing of faces and their role in complex social behaviors (3). However, it is still unknown to what extent these behavioral biases arise from the tuning of the underlying face-processing mechanisms or, alternatively, from the nature of the experimental stimuli and task (10, 11). Here we approach this question from a new angle by examining face evaluation for a different class of faces: illusory faces in inanimate objects.

Face pareidolia is the spontaneous perception of illusory facial features in inanimate objects (Fig. 1), and can be thought of as a natural error of our face detection system (14–18). It has recently been shown that nonhuman primates also experience face pareidolia (14, 15), and that illusory faces engage similar neural mechanisms to real faces in the human brain (18). However, it is unclear to what degree higher-level social perception beyond the detection of a face occurs in pareidolia. Investigation of face evaluation in illusory faces has the potential to reveal new insight into the underlying mechanisms of face perception. A key feature of face pareidolia is that it involves the spontaneous perception of a face in an inanimate object, and consequently it is an example of face perception that is divorced from many characteristics that typically accompany the faces of living organisms, such as the motion of facial muscles (e.g., to form emotional expressions), chronological age, and biological sex. The primary question we address here is whether illusory faces are perceived to have these traits even in the absence of their biological specification. As there is no a priori reason why an illusory face should be perceived to have a specific age, gender or expression, any reliable perception of these attributes would be informative about inherent properties of the underlying system.

Studies using human faces have suggested potential biases in the perceived characteristics of human faces along dimensions such as age (10, 13) and gender (10, 11, 19) under conditions of visual uncertainty. However, determining the potential origin and generality of these biases has proven difficult and highlights the fundamental challenges inherent in understanding how the perception of specific traits is linked to face processing. Human faces are visually complex, and our brains are incredibly well-adapted to processing faces as a cohesive whole (20). Consequently, it is challenging to empirically isolate particular aspects of a human face (e.g., biological sex) from other interdependencies (e.g., identity). Additionally, since human faces have a biologically specified age and gender, it is necessary to introduce uncertainty via deliberate experimental manipulation of the stimuli. Studies of human faces have used various forms of image manipulation, including removing hair (21, 22), showing silhouettes of faces in profile (23), adding visual noise (24), and synthetically generating faces by morphing along stimulus dimensions, such as gender (10, 11, 19). A critical advantage of using pareidolia to probe the tuning of the face-processing system is that no decisions about stimulus manipulation need to

Significance

Face pareidolia is the phenomenon of perceiving illusory faces in inanimate objects. Here we show that illusory faces engage social perception beyond the detection of a face: they have a perceived age, gender, and emotional expression. Additionally, we report a striking bias in gender perception, with many more illusory faces perceived as male than female. As illusory faces do not have a biological sex, this bias is significant in revealing an asymmetry in our face evaluation system given minimal information. Our result demonstrates that the visual features that are sufficient for face detection are not generally sufficient for the perception of female. Instead, the perception of a nonhuman face as female requires additional features beyond that required for face detection.

Author contributions: S.G.W., S.P., J.T., and C.I.B. designed research; S.G.W. and S.P. performed research; S.G.W. and S.P. contributed new reagents/analytic tools; S.G.W., S.P., J.T., and C.I.B. analyzed data; and S.G.W., S.P., J.T., and C.I.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]S.G.W. and S.P. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: susan.wardle@nih.gov.

This article contains supporting information online at http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2117413119/-/DCSupplemental.

Published January 24, 2022.

Illusory faces are more likely to be perceived as male than female

Susan G. Wardle[a,1,2], Sanika Paranjape[a,1], Jessica Taubert[b], and Chris I. Baker[a]

[a]Section on Learning and Plasticity, Laboratory of Brain & Cognition, National Institute of Mental Health, Bethesda, MD 20814; and [b]Section on Neurocircuitry, Laboratory of Brain & Cognition, National Institute of Mental Health, Bethesda, MD 20814

Edited by Susan Fiske, Psychology Department, Princeton University, Princeton, NJ; received October 1, 2021; accepted November 17, 2021

Despite our fluency in reading human faces, sometimes we mistakenly perceive illusory faces in objects, a phenomenon known as face pareidolia. Although illusory faces share some neural mechanisms with real faces, it is unknown to what degree pareidolia engages higher-level social perception beyond the detection of a face. In a series of large-scale behavioral experiments (n_{total} = 3,815 adults), we found that illusory faces in inanimate objects are readily perceived to have a specific emotional expression, age, and gender. Most strikingly, we observed a strong bias to perceive illusory faces as male rather than female. This male bias could not be explained by preexisting semantic or visual gender associations with the objects, or by visual features in the images. Rather, this robust bias in the perception of gender for illusory faces reveals a cognitive bias arising from a broadly tuned face evaluation system in which minimally viable face percepts are more likely to be perceived as male.

face perception | gender | bias | pareidolia | face evaluation

Human faces convey a rich amount of social information beyond their identity (1–3). We are able to rapidly evaluate the age (4), gender (5, 6), and emotional expression (7) of the faces of individuals, even if they are not known to us, in addition to more abstract traits, such as trustworthiness and aggressiveness (8, 9). Although these judgements are based on visual information, biases have been identified that suggest that both perceptual and cognitive factors are involved in face evaluation (10–13). For example, people tend to judge faces as closer to their own age (10, 13), and damage to the amygdala is associated with perceiving unfamiliar faces as more trustworthy and approachable (12). Biases in face perception have important implications for understanding the neural processing of faces and their role in complex social behaviors (3). However, it is still unknown to what extent these behavioral biases arise from the tuning of the underlying face-processing mechanisms or, alternatively, from the nature of the experimental stimuli and task (10, 11). Here we approach this question from a new angle by examining face evaluation for a different class of faces: illusory faces in inanimate objects.

Face pareidolia is the spontaneous perception of illusory facial features in inanimate objects (Fig. 1), and can be thought of as a natural error of our face detection system (14–18). It has recently been shown that nonhuman primates also experience face pareidolia (14, 15), and that illusory faces engage similar neural mechanisms to real faces in the human brain (18). However, it is unclear to what degree higher-level social perception beyond the detection of a face occurs in pareidolia. Investigation of face evaluation in illusory faces has the potential to reveal new insight into the underlying mechanisms of face perception. A key feature of face pareidolia is that it involves the spontaneous perception of a face in an inanimate object, and consequently it is an example of face perception that is divorced from many characteristics that typically accompany the faces of living organisms, such as the motion of facial muscles (e.g., to form emotional expressions), chronological age, and biological sex. The primary question we address here is whether illusory faces are perceived to have these traits even in the absence of their biological specification. As there is no a priori reason why an illusory face should be perceived to have a specific age, gender or expression, any reliable perception of these attributes would be informative about inherent properties of the underlying system.

Studies using human faces have suggested potential biases in the perceived characteristics of human faces along dimensions such as age (10, 13) and gender (10, 11, 19) under conditions of visual uncertainty. However, determining the potential origin and generality of these biases has proven difficult and highlights the fundamental challenges inherent in understanding how the perception of specific traits is linked to face processing. Human faces are visually complex, and our brains are incredibly well-adapted to processing faces as a cohesive whole (20). Consequently, it is challenging to empirically isolate particular aspects of a human face (e.g., biological sex) from other interdependencies (e.g., identity). Additionally, since human faces have a biologically specified age and gender, it is necessary to introduce uncertainty via deliberate experimental manipulation of the stimuli. Studies of human faces have used various forms of image manipulation, including removing hair (21, 22), showing silhouettes of faces in profile (23), adding visual noise (24), and synthetically generating faces by morphing along stimulus dimensions, such as gender (10, 11, 19). A critical advantage of using pareidolia to probe the tuning of the face-processing system is that no decisions about stimulus manipulation need to

Significance

Face pareidolia is the phenomenon of perceiving illusory faces in inanimate objects. Here we show that illusory faces engage social perception beyond the detection of a face: they have a perceived age, gender, and emotional expression. Additionally, we report a striking bias in gender perception, with many more illusory faces perceived as male than female. As illusory faces do not have a biological sex, this bias is significant in revealing an asymmetry in our face evaluation system given minimal information. Our result demonstrates that the visual features that are sufficient for face detection are not generally sufficient for the perception of female. Instead, the perception of a nonhuman face as female requires additional features beyond that required for face detection.

Author contributions: S.G.W., S.P., J.T., and C.I.B. designed research; S.G.W. and S.P. performed research; S.G.W. and S.P. contributed new reagents/analytic tools; S.G.W., S.P., J.T., and C.I.B. analyzed data; and S.G.W., S.P., J.T., and C.I.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]S.G.W. and S.P. contributed equally to this work.

[2]To whom correspondence may be addressed. Email: susan.wardle@nih.gov.

This article contains supporting information online at http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2117413119/-/DCSupplemental.

Published January 24, 2022.

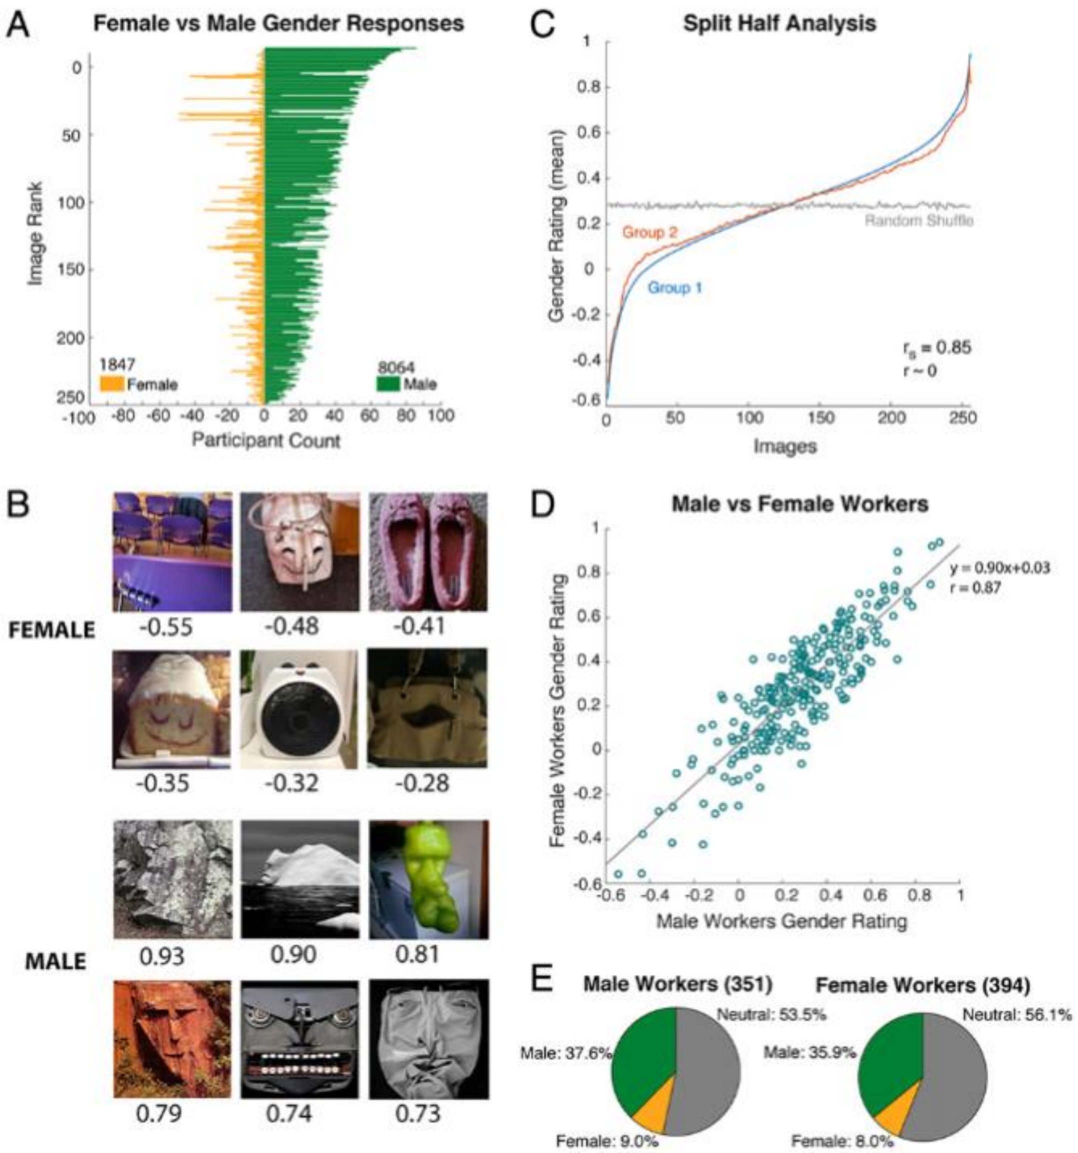


Vox

Shepard tone

Risset Rhythm

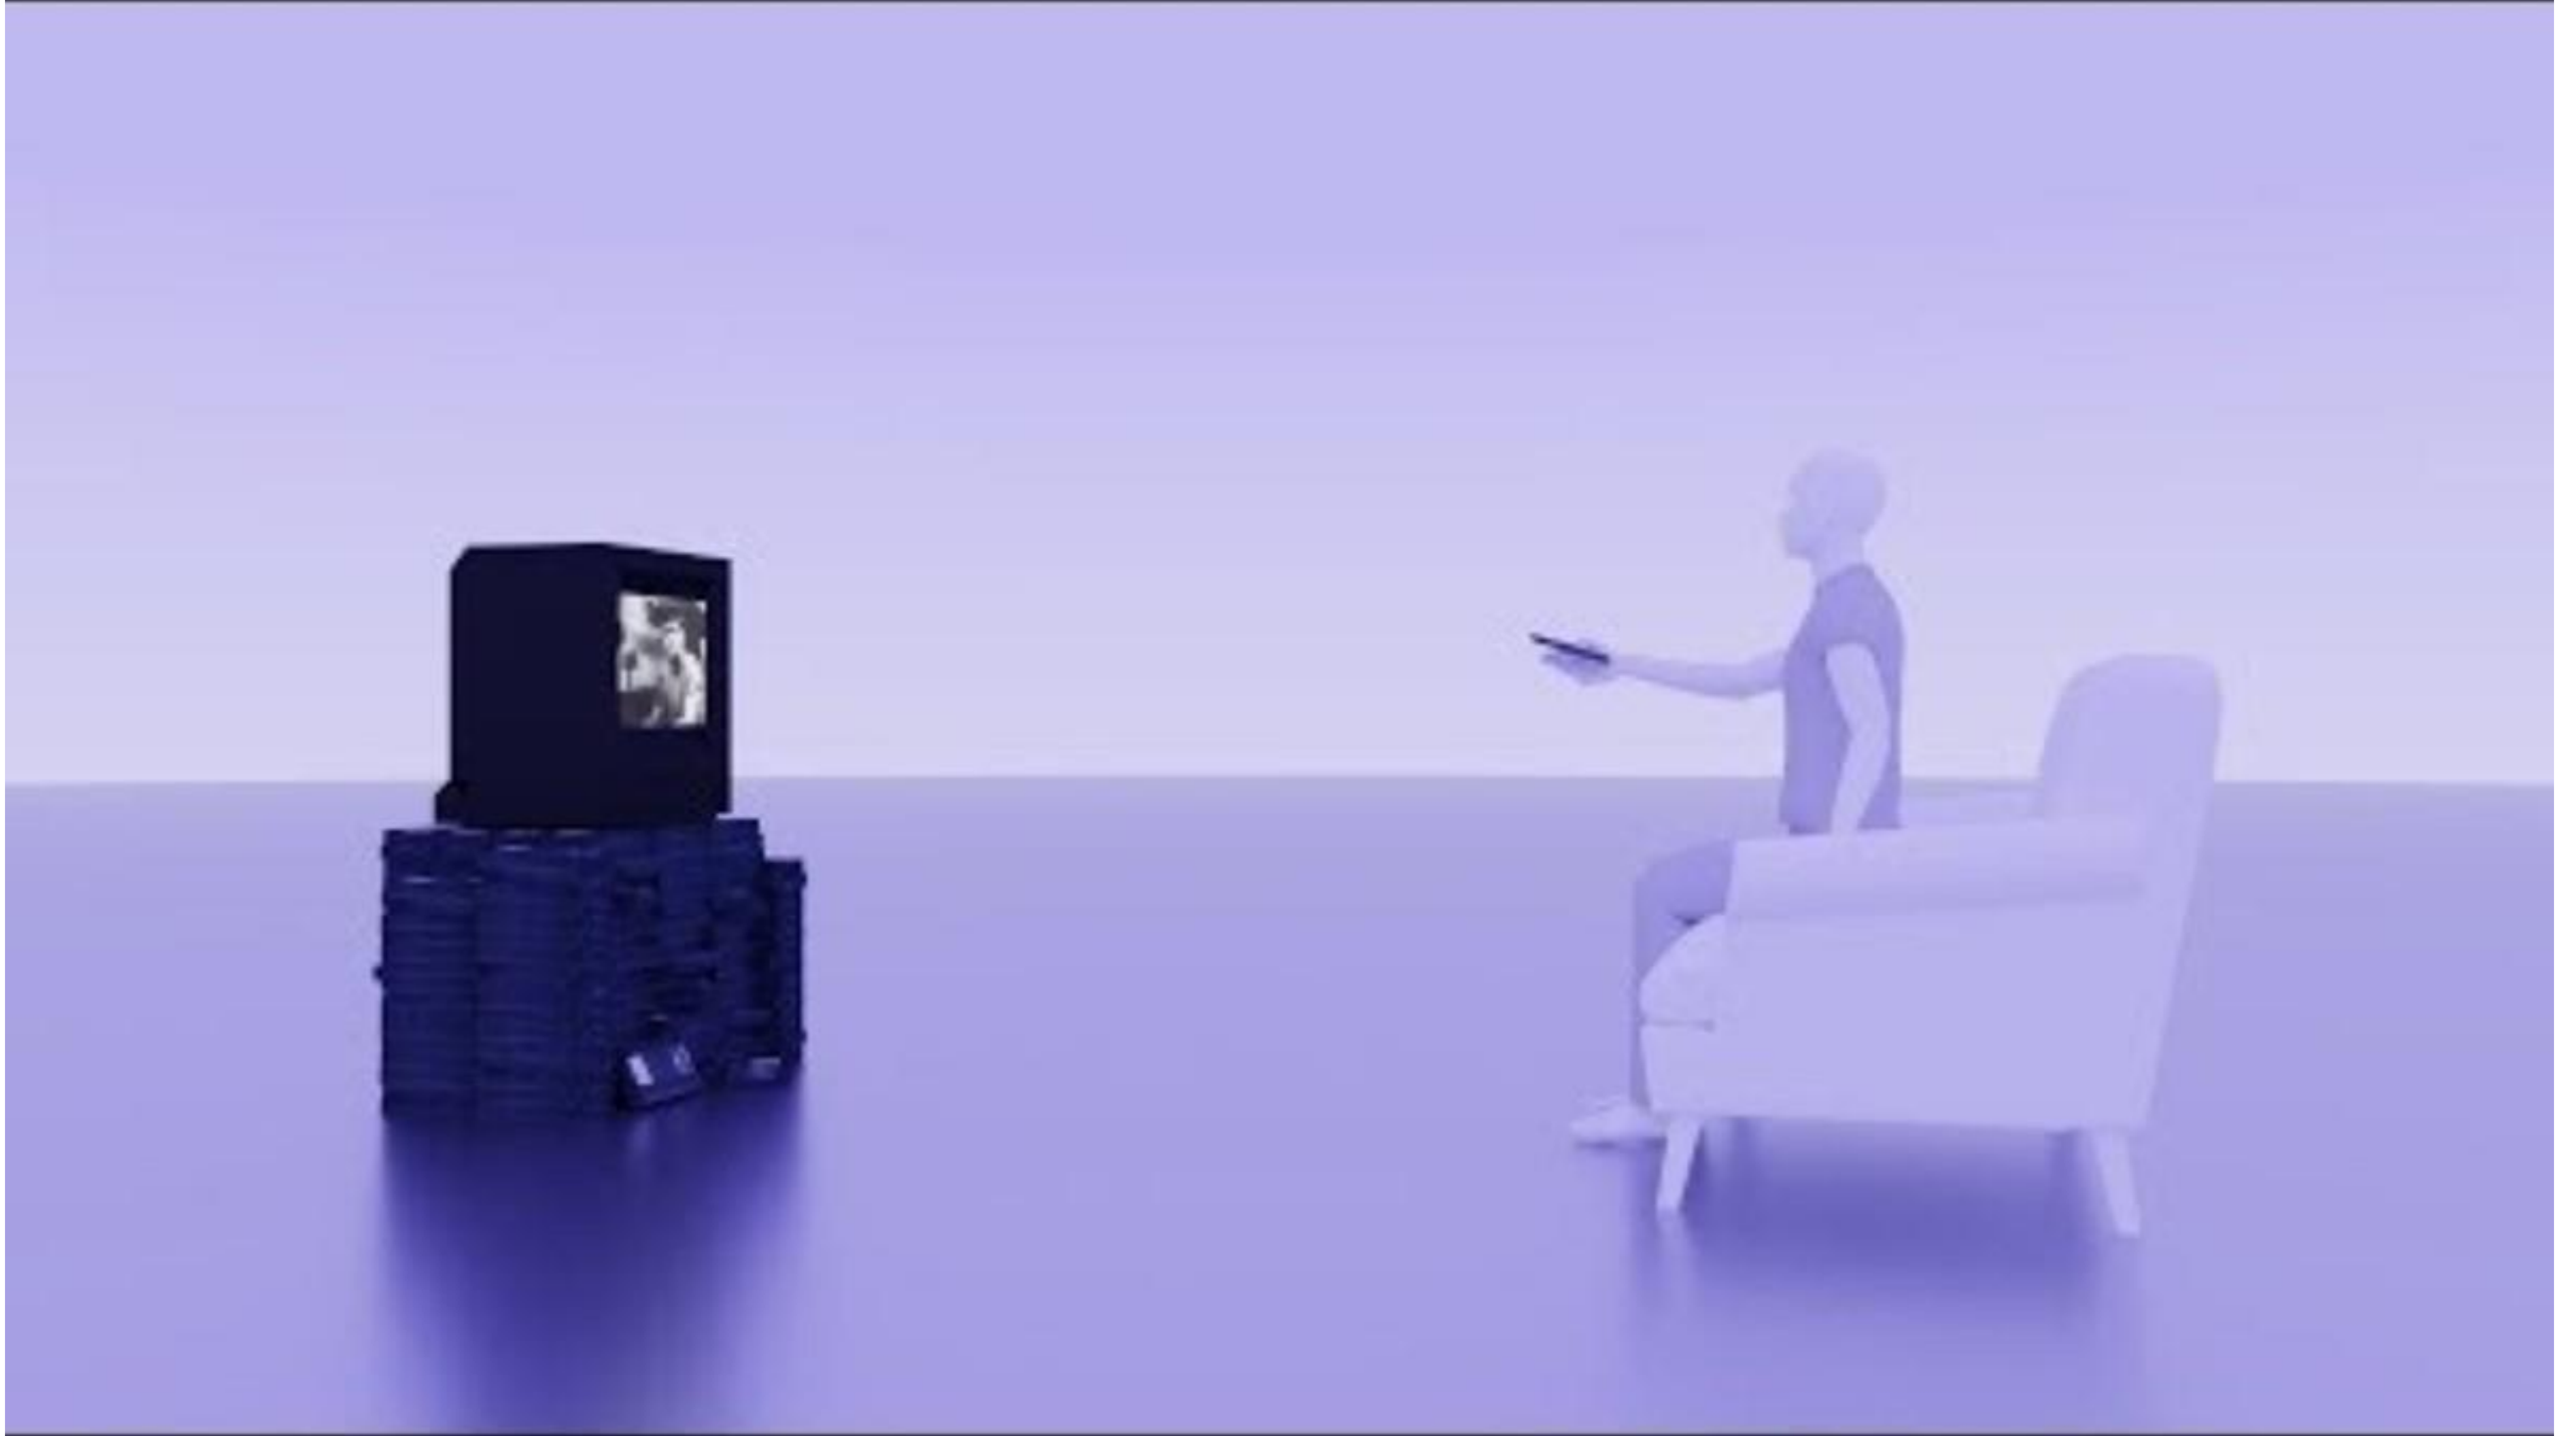

Speech-to-song illusion

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

Segmentation
Object-based attention

object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

Auditory Illusion

4

Mystery Melody


ILLUSIONS
OF FACES
Sami Ryan Yousif

McGurk Effect

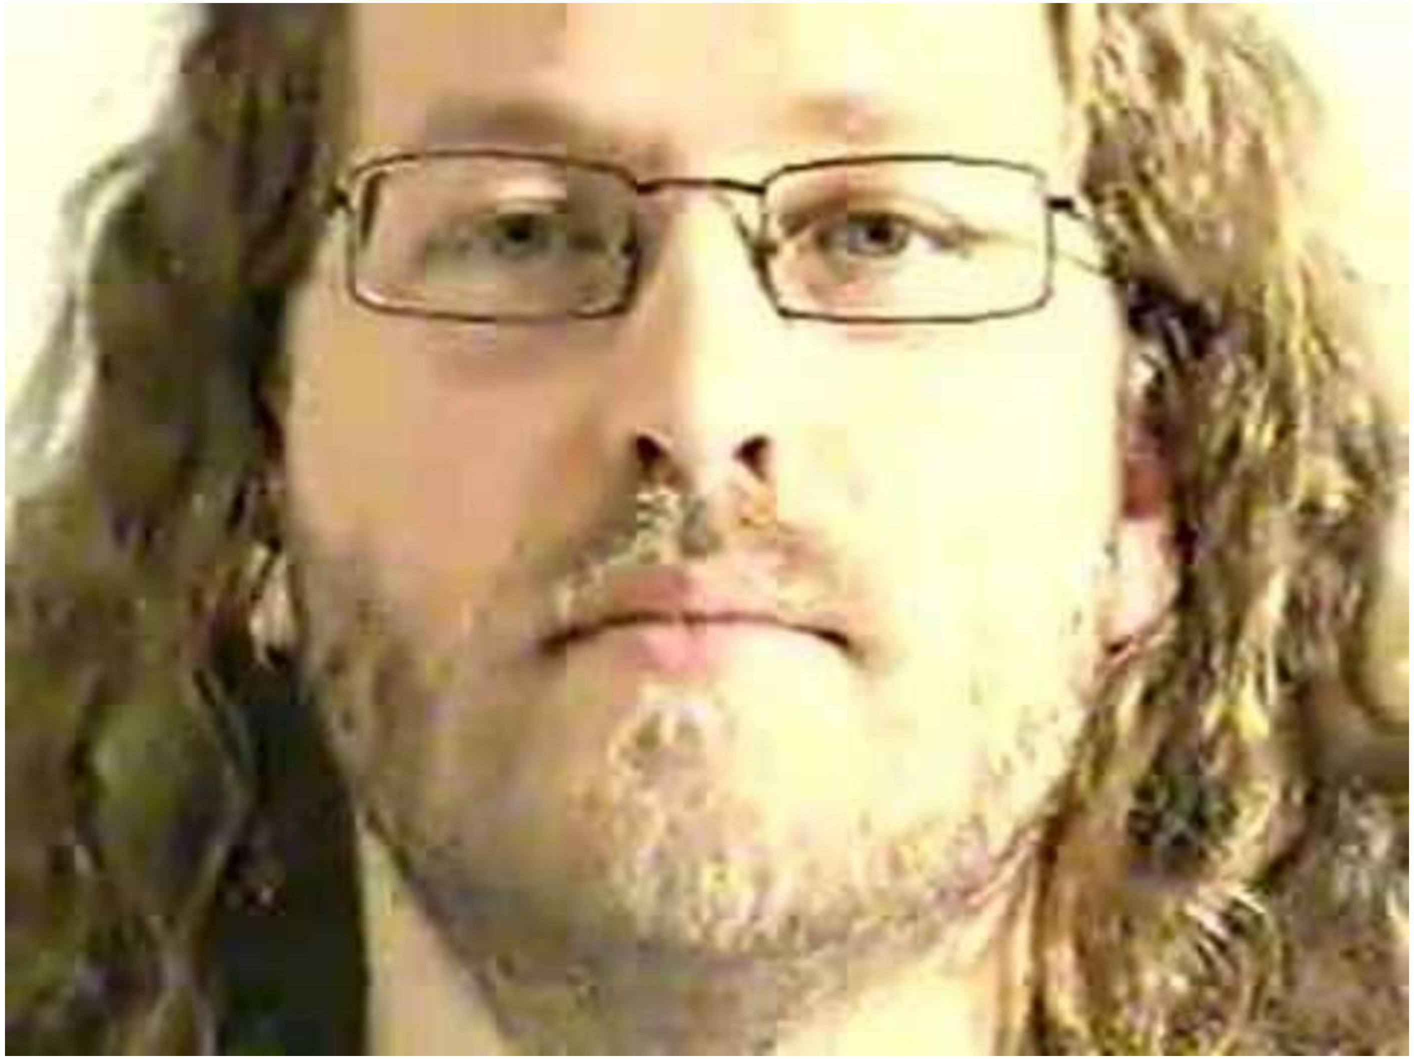

McGurk Effect

‼

A *cross-modal* *illusion!*

A related phenomenon (thanks TikTok!)

Auditory Illusion

Phantom Words

The opposite is also true:
***Semantic satiation*.**

*Spoken language is much more **ambiguous** than we realize!*

Cognition 235 (2023) 105401

Contents lists available at ScienceDirect

Cognition

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

Consequences of phonological variation for algorithmic word segmentation

Caroline Beech [*], Daniel Swingley

Department of Psychology, University of Pennsylvania, 425 S University Ave, Philadelphia, PA 19104, USA

ARTICLE INFO

Keywords:
Language acquisition
Computational modeling
Word segmentation
Phonological variation

ABSTRACT

Over the first year, infants begin to learn the words of their language. Previous work suggests that certain statistical regularities in speech could help infants segment the speech stream into words, thereby forming a proto-lexicon that could support learning of the eventual vocabulary. However, computational models of word segmentation have typically been tested using language input that is much less variable than actual speech is. We show that using actual, transcribed pronunciations rather than dictionary pronunciations of the same speech leads to worse segmentation performance across models. We also find that phonologically variable input poses serious problems for lexicon building, because even correctly segmented word forms exhibit a complex, many-to-many relationship with speakers' intended words. Many phonologically distinct word forms were actually the same intended word, and many identical transcriptions came from different intended words. The fact that previous models appear to have substantially overestimated the utility of simple statistical heuristics suggests a need to consider the formation of the lexicon in infancy differently.

1. Introduction

Although infants are born knowing little about their native language, they quickly learn a great deal from the speech they hear. Within months, they become familiar with their native language's sound categories (Werker & Tees, 1984), as well as the relative frequency of different sequences of speech sounds (Archer, Czarnecki, & Curtin, 2021; Jusczyk, Friederici, Wessels, Svenkerud, & Jusczyk, 1993). Beyond learning about their language's phonology, infants also begin to learn words. Months before their first birthday, they recognize the meanings of some common words, including both concrete nouns (Bergelson & Swingley, 2012) and a little later, more abstract words (Bergelson & Swingley, 2013), and by the second half of the first year, they recognize the spoken form of a variety of words familiar from home experience or laboratory exposure (e.g., Hallé & Boysson-Bardies, 1994; Jusczyk & Aslin, 1995; Jusczyk & Hohne, 1997; Schreiner, Altvater-Mackensen, & Mani, 2016; Swingley, 2005a; Vihman, Nakai, DePaolis, & Hallé, 2004).

An important step in the process of language learning is word segmentation, or pulling out words from the continuous stream of speech. It is easy to understand that this is a difficult problem—one only needs to listen to a parent speaking to an infant in an unfamiliar language to recognize that it is quite hard to infer where one word ends and the next begins. This problem is difficult for infants too, which is why infants learn words more easily when they are presented in one-word utterances than when they are embedded in longer utterances (Brent & Siskind, 2001; Keren-Portnoy, Vihman, & Fisher, 2019; Swingley & Humphrey, 2018). Yet infants do manage to break utterances into parts. Laboratory studies demonstrate that infants can extract words from their phonetic contexts (e.g., Jusczyk & Aslin, 1995), and infants have some knowledge of grammatical words that never appear in isolation (e.g., Shi & Lepage, 2008).

Research into infants' early discovery of words has taken two forms: experiments that present continuous speech to infants and test which elements they retain, and computational models that evaluate what infants might learn were they to parse and retain speech sequences according to a particular set of computable heuristics. The present paper continues the latter line, but differs from most prior work in examining the consequences of normal phonological variability. When words are realized in more than one way, does the phonological structure of the lexicon still permit simple probabilistic heuristics to succeed in producing the foundation of the early vocabulary?

In principle, there are several cues that could be helpful in word segmentation, once the infant has some familiarity with phonological regularities present in the lexicon. For example, in English, strong syllables tend to coincide with word onsets, suggesting that English speakers could learn to use stress patterns or vowel-reduction patterns to detect where an unknown word begins (Cutler & Norris, 1988).

* Corresponding author.
E-mail addresses: cbeech@sas.upenn.edu (C. Beech), swingley@psych.upenn.edu (D. Swingley).

https://doi.org/10.1016/j.cognition.2023.105401
Received 15 July 2022; Received in revised form 27 December 2022; Accepted 6 February 2023
Available online 12 February 2023

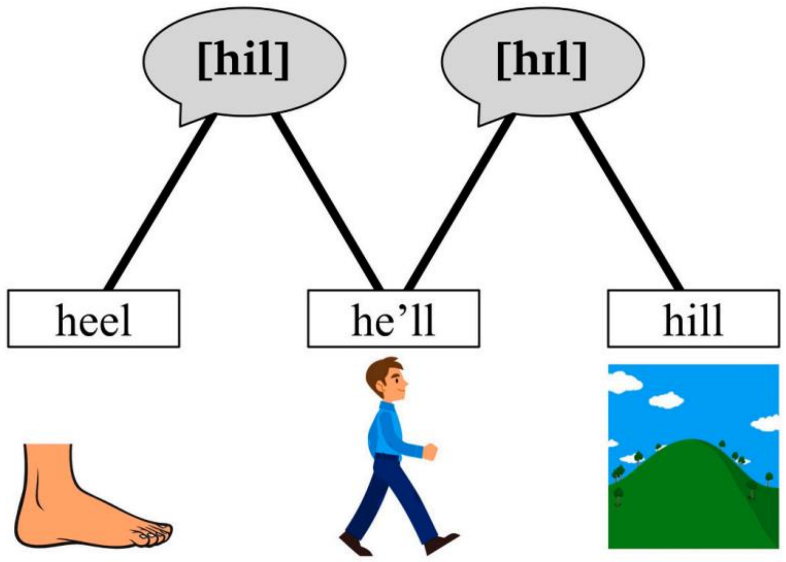

Fig. 2. Example many-to-many mapping. Edges between phonological word forms (shaded) and orthographic words (white) represent attested pronunciations. In this network, the same orthographic word ("he'll") can have multiple pronunciations ([hil] and [hɪl]), and a single phonological word form can map onto multiple orthographic words and thus meanings (e.g., [hil] maps onto both "heel" and "he'll").

We [fɪl] in the gaps with our prior knowledge

PNAS | RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES
SOCIAL SCIENCES

The perception of silence

Rui Zhe Goh[a,b,1], Ian B. Phillips[a,b,1,2], and Chaz Firestone[a,b,1,2]

Edited by Steven Luck, University of California, Davis, CA; received February 8, 2023; accepted May 10, 2023 by Editorial Board Member Michael S. Gazzaniga

Auditory perception is traditionally conceived as the perception of sounds—a friend's voice, a clap of thunder, a minor chord. However, daily life also seems to present us with experiences characterized by the absence of sound—a moment of silence, a gap between thunderclaps, the hush after a musical performance. In these cases, do we positively *hear* silence? Or do we just *fail to hear*, and merely judge or infer that it is silent? This longstanding question remains controversial in both the philosophy and science of perception, with prominent theories holding that sounds are the only objects of auditory experience and thus that our encounter with silence is cognitive, not perceptual. However, this debate has largely remained theoretical, without a key empirical test. Here, we introduce an empirical approach to this theoretical dispute, presenting experimental evidence that silence can be genuinely perceived (not just cognitively inferred). We ask whether silences can "substitute" for sounds in event-based auditory illusions—empirical signatures of auditory event representation in which auditory events distort perceived duration. Seven experiments introduce three "silence illusions"—the one-silence-is-more illusion, silence-based warping, and the oddball-silence illusion—each adapted from a prominent perceptual illusion previously thought to arise only from sounds. Subjects were immersed in ambient noise interrupted by silences structurally identical to the sounds in the original illusions. In all cases, silences elicited temporal distortions perfectly analogous to the illusions produced by sounds. Our results suggest that silence is truly heard, not merely inferred, introducing a general approach for studying the perception of absence.

absence perception | silence | event representation | temporal illusions

Significance

Do we only hear sounds? Or can we also hear silence? These questions are the subject of a centuries-old philosophical debate between two camps: the perceptual view (we literally hear silence), and the cognitive view (we only judge or infer silence). Here, we take an empirical approach to resolve this theoretical controversy. We show that silences can "substitute" for sounds in event-based auditory illusions. Seven experiments introduce three "silence illusions," adapted from perceptual illusions previously thought to arise only with sounds. In all cases, silences elicited temporal distortions perfectly analogous to their sound-based counterparts, suggesting that auditory processing treats moments of silence the way it treats sounds. Silence is truly perceived, not merely inferred.

What do we hear? The canonical answer is that auditory perception is the perception of *sounds* and their properties—the pitch of a friend's voice, the loudness of a thunderclap, the timbre of a minor chord. This traditional view has considerable pedigree, with influential historical sources holding that sounds are the sole objects of auditory perception (1, cf.2). It is also the answer favored in contemporary scholarship: Prominent scientific accounts conceive the fundamental units of auditory perception as sounds (or auditory streams comprised of sounds; ref. 3 and 4), and many philosophical theories agree, holding that "all auditory perception involves the perception of sound" (5) and that "if anything at all is heard, what is heard is necessarily a sound" (6) (see also refs. 7 and 8). The pervasiveness of this canonical view about the contents of auditory perception might seem unsurprising—what else might we hear, if not sound?

However, there has long been a stubborn and intuitive counterexample: experiences of *silence*, which are characterized by the absence of sound. Silence confronts us throughout our daily lives—consider an awkward pause in a conversation, a suspenseful gap between thunderclaps, or the hush at the end of a musical performance. What is the nature of these experiences?

Silence: Heard or Inferred?

One possibility is that experiences of silence are simply cases in which we *fail to hear*, and then use our faculties of reasoning and judgment to *infer* that it is silent. This interpretation is offered by those who defend the traditional sound-only view of audition, holding that an experience of silence is merely the "cognitive accompaniment of an absence of experience" and "is itself no form of hearing" (9). This cognitive view may be motivated by a deeper assumption about perception, namely that we can genuinely perceive only what is present in the world, not what is absent (9, 10). After all, one might think, absences are nonentities—they do not exist—and so can hardly impinge on our sensory apparatus.

However, an alternative possibility which arguably does more justice to our phenomenology is that we literally perceive silences. This interpretation has recently received

Author affiliations: [a]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD 21218; and [b]Department of Philosophy, Johns Hopkins University, Baltimore, MD 21218

Author contributions: R.Z.G., I.B.P., and C.F. designed research; R.Z.G. performed research; R.Z.G. analyzed data; and R.Z.G., I.B.P., and C.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. S.L. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: rgoh1@jhu.edu, ianbphillips@jhu.edu, or chaz@jhu.edu.

[2]I.B.P. and C.F. contributed equally to this work.

Published July 10, 2023.

PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES
SOCIAL SCIENCES

The perception of silence

Rui Zhe Goh[a,b,1], Ian B. Phillips[a,b,1,2], and Chaz Firestone[a,b,1,2]

Edited by Steven Luck, University of California, Davis, CA; received February 8, 2023; accepted May 10, 2023 by Editorial Board Member Michael S. Gazzaniga

Auditory perception is traditionally conceived as the perception of sounds—a friend's voice, a clap of thunder, a minor chord. However, daily life also seems to present us with experiences characterized by the absence of sound—a moment of silence, a gap between thunderclaps, the hush after a musical performance. In these cases, do we positively *hear* silence? Or do we just *fail to hear*, and merely judge or infer that it is silent? This longstanding question remains controversial in both the philosophy and science of perception, with prominent theories holding that sounds are the only objects of auditory experience and thus that our encounter with silence is cognitive, not perceptual. However, this debate has largely remained theoretical, without a key empirical test. Here, we introduce an empirical approach to this theoretical dispute, presenting experimental evidence that silence can be genuinely perceived (not just cognitively inferred). We ask whether silences can "substitute" for sounds in event-based auditory illusions—empirical signatures of auditory event representation in which auditory events distort perceived duration. Seven experiments introduce three "silence illusions"—the one-silence-is-more illusion, silence-based warping, and the oddball-silence illusion—each adapted from a prominent perceptual illusion previously thought to arise only from sounds. Subjects were immersed in ambient noise interrupted by silences structurally identical to the sounds in the original illusions. In all cases, silences elicited temporal distortions perfectly analogous to the illusions produced by sounds. Our results suggest that silence is truly heard, not merely inferred, introducing a general approach for studying the perception of absence.

absence perception | silence | event representation | temporal illusions

Significance

Do we only hear sounds? Or can we also hear silence? These questions are the subject of a centuries-old philosophical debate between two camps: the perceptual view (we literally hear silence), and the cognitive view (we only judge or infer silence). Here, we take an empirical approach to resolve this theoretical controversy. We show that silences can "substitute" for sounds in event-based auditory illusions. Seven experiments introduce three "silence illusions," adapted from perceptual illusions previously thought to arise only with sounds. In all cases, silences elicited temporal distortions perfectly analogous to their sound-based counterparts, suggesting that auditory processing treats moments of silence the way it treats sounds. Silence is truly perceived, not merely inferred.

What do we hear? The canonical answer is that auditory perception is the perception of *sounds* and their properties—the pitch of a friend's voice, the loudness of a thunderclap, the timbre of a minor chord. This traditional view has considerable pedigree, with influential historical sources holding that sounds are the sole objects of auditory perception (1, cf.2). It is also the answer favored in contemporary scholarship: Prominent scientific accounts conceive the fundamental units of auditory perception as sounds (or auditory streams comprised of sounds; ref. 3 and 4), and many philosophical theories agree, holding that "all auditory perception involves the perception of sound" (5) and that "if anything at all is heard, what is heard is necessarily a sound" (6) (see also refs. 7 and 8). The pervasiveness of this canonical view of the contents of auditory perception might seem unsurprising—what else might we hear, if not sound?

However, there has long been a stubborn and intuitive counterexample: experiences of *silence*, which are characterized by the absence of sound. Silence confronts us throughout our daily lives—consider an awkward pause in a conversation, a suspenseful gap between thunderclaps, or the hush at the end of a musical performance. What is the nature of these experiences?

Silence: Heard or Inferred?

One possibility is that experiences of silence are simply cases in which we *fail to hear*, and then use our faculties of reasoning and judgment to *infer* that it is silent. This interpretation is offered by those who defend the traditional sound-only view of audition, holding that an experience of silence is merely the "cognitive accompaniment of an absence of experience" and "is itself no form of hearing" (9). This cognitive view may be motivated by a deeper assumption about perception, namely that we can genuinely perceive only what is present in the world, not what is absent (9, 10). After all, one might think, absences are nonentities—they do not exist—and so can hardly impinge on our sensory apparatus.

However, an alternative possibility which arguably does more justice to our phenomenology is that we literally perceive silences. This interpretation has recently received

Author affiliations: [a]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD 21218; and [b]Department of Philosophy, Johns Hopkins University, Baltimore, MD 21218

Author contributions: R.Z.G., I.B.P., and C.F. designed research; R.Z.G. performed research; R.Z.G. analyzed data; and R.Z.G., I.B.P., and C.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. S.L. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: rgoh1@jhu.edu, ianbphillips@jhu.edu, or chaz@jhu.edu.

[2]I.B.P. and C.F. contributed equally to this work.

Published July 10, 2023.

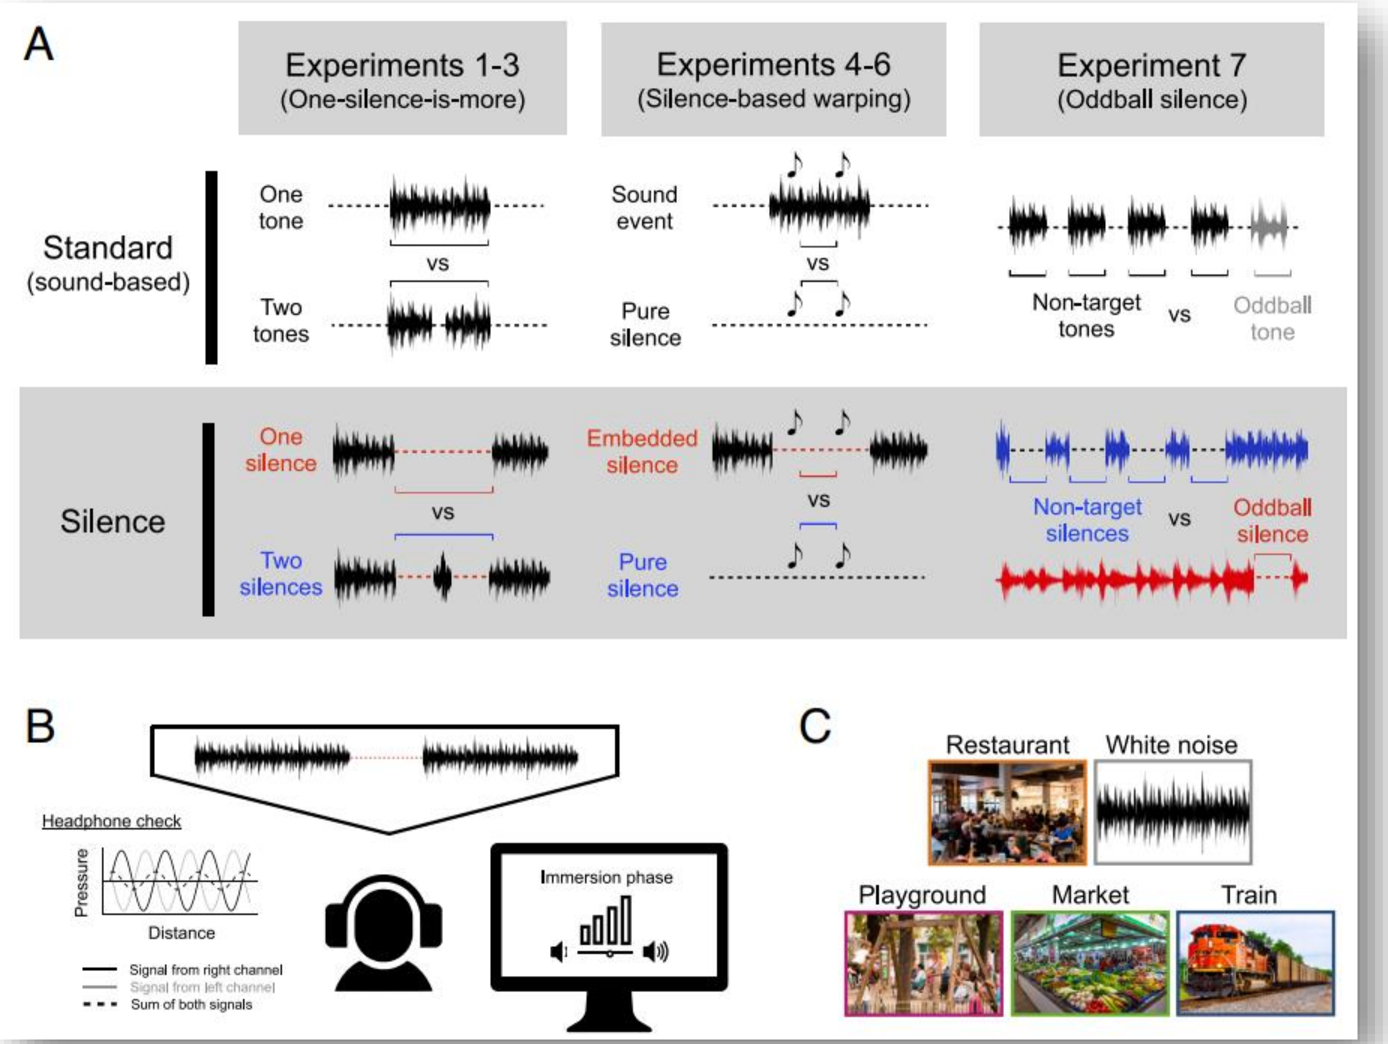

Food for thought:

1. We are mostly unaware of all the ambiguities in sound that we just discussed. **How/why?**

2. **Are there any really obvious ways in which sound perception is not analogous to visual perception?**

3. **What is the deal with music?**


ILLUSIONS OF TIME
Sami Ryan Yousif

The one-is-more illusion

Cognition 185 (2019) 121–130

Contents lists available at ScienceDirect

Original Articles

The one-is-more illusion: Sets of discrete objects appear less extended than equivalent continuous entities in both space and time

Check for updates

Sami R. Yousif*, Brian J. Scholl*

Yale University, USA

Segmentation
Object-based attention

object (e.g. a rectangle) as longer than an equated set of multiple discrete objects (e.g. two shorter rectangles separated by a gap). This illusion is phenomenologically compelling, exceptionally reliable, and it extends beyond space, to time: a single continuous tone is perceived to last longer than an equated set of multiple discrete tones. Previous work has emphasized the importance of objecthood for processes such as attention and visual working memory, but these results typically require careful analyses of subtle effects. In contrast, we provide striking demonstrations of how perceived objecthood changes the perception of other properties in a way that you can readily see (and hear!) with your own eyes (and ears!).

1. Introduction

One of the most fundamental and pervasive distinctions in cognitive science is that between the continuous and the discrete. Indeed, one of the key insights of the cognitive revolution was that intelligent behavior could be explained in part by appeal to discrete symbolic representations, even when the neural implementations of those discrete symbols might themselves be continuous (for seminal reviews see Newell, 1980; Pylyshyn, 1984). In cognitive psychology, this distinction has inspired spirited debate about the mechanisms of learning — where continuous, gradual processes (such as long-term potentiation) are contrasted with approaches that rely on storing and updating the values of discrete variables (e.g. Gallistel, 2000). And in developmental psychology, language researchers have sought to understand how the child's mind turns a continuous stream of syllables into representations of discrete words (e.g. Saffran, Aslin, & Newport, 1996).

Perhaps nowhere, though, has the distinction between the continuous and the discrete been more salient in cognitive science than in the study of perception. Sometimes this distinction is drawn explicitly, for example when asking about the temporal resolution of perception (e.g. VanRullen & Koch, 2003; see also Asplund, Fougnie, Zughni, Martin, & Marois, 2014). In other cases, the distinction is just as fundamental, but more implicit. For example, arguably the two most active areas in the study of visual cognition over the past two decades have been visual working memory and attention — and in both of these areas, this distinction has been central. The underlying units of visual working memory, for example, have been characterized as both discrete (limited by the number of 'slots' corresponding to encoded objects, regardless of their features; e.g. Luck & Vogel, 1997) and as continuous (limited by the overall amount of encoded information, regardless of how that information is distributed among objects; e.g. Alvarez & Cavanagh, 2004), and this remains an area of active debate (for a review, see Suchow, Fougnie, Brady, & Alvarez, 2014). And visual selective attention has similarly been characterized as both continuous (operating akin to a spotlight that selects undifferentiated spatial regions of the visual field; for a review, see Cave & Bichot, 1999) and discrete (selecting and shifting among individual objects rather than spatial regions; for a review, see Scholl, 2001).

One seemingly awkward aspect of these various theories, however, is that despite being theories of perception (and thus of seeing), the relevant effects cannot typically be seen. Instead, these effects (e.g. 'same-object-advantages' in object-based attention; e.g. Egly, Driver, & Rafal, 1994) are often relatively small, and only come out in the statistical wash. As such, the present experiments asked (for the first time, to our knowledge): does this sort of 'objecthood', beyond influencing attention and memory, also affect what we see in the first place?

* Corresponding authors at: Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA.
E-mail addresses: sami.yousif@yale.edu (S.R. Yousif), brian.scholl@yale.edu (B.J. Scholl).

https://doi.org/10.1016/j.cognition.2018.10.002
Received 18 April 2018; Received in revised form 5 October 2018; Accepted 5 October 2018
Available online 22 January 2019

Journal of Experimental Psychology: General
2014, Vol. 143, No. 1, 71–76

0096-3445/14/$12.00 DOI: 10.1037/a0032296

BRIEF REPORT

Attentional Rhythm: A Temporal Analogue of Object-Based Attention

Julian De Freitas, Brandon M. Liverence, and Brian J. Scholl
Yale University

The underlying units of attention are often discrete visual objects. Perhaps the clearest form of evidence for this is the *same-object advantage*: Following a spatial cue, responses are faster to probes occurring on the same object than they are to probes occurring on other objects, while equating brute distance. Is this a fundamentally spatial effect, or can same-object advantages also occur in time? We explored this question using independently normed rhythmic temporal sequences, structured into phrases and presented either visually or auditorily. Detection was speeded when cues and probes both lay within the same rhythmic phrase, compared to when they spanned a phrase boundary, while equating brute duration. This same-phrase advantage suggests that object-based attention is a more general phenomenon than has been previously suspected: Perceptual structure constrains attention, in both space and time, and in both vision and audition.

Keywords: object-based attention, rhythm, music perception, auditory perception

Supplemental materials: http://dx.doi.org/10.1037/a0032296.supp

The most fundamental feature of our visual experience may be space itself, since it appears to be the medium in which all other visual representations exist. Perhaps for this reason, much of the classic work on visual attention was grounded in spatial metaphors (Fernandez-Duque & Johnson, 1999), likening attention to a spotlight or zoom lens (for a review, see Cave & Bichot, 1999). However, more recent research in cognitive science has revealed that a wide range of mental processes, including visual attention, operate over units that are fundamentally discrete.

Object-Based Attention

A perfect example of such processing is *object-based attention*, a class of effects in which discrete visual objects act as units of selection, constraining otherwise equated shifts of spatial attention (for reviews, see Chen, 2012; Scholl, 2001). Perhaps the most direct evidence for this comes from demonstrations of *same-object advantages* in spatial shifts of attention. In a classic demonstration of this effect (Egly, Driver, & Rafal, 1994), observers viewed two vertically oriented rectangles, with their attention cued by a luminance change to an end of one rectangle (see Figure 1A). After a brief delay, a probe appeared at an end of one of the rectangles, and observers pressed a key in response. The cue validly predicted the location of the probe on most trials, but on invalid trials the probe occurred at either the opposite end of the cued rectangle (within object) or at an equidistant point on the neighboring rectangle (between object). There was a same-object advantage: On invalidly cued trials, observers responded faster to within-object than between-object probes. The mechanisms underlying such effects have been explored and debated in dozens of subsequent studies (Chen, 2012; Scholl, 2001).

Structure in Space

Despite the name *object-based attention*, several results suggest that objects per se are not required. This type of effect has also been demonstrated with groups (e.g., Dodd & Pratt, 2005), parts (e.g., Barenholtz & Feldman, 2003; Vecera, Behrmann, & McGoldrick, 2000), surfaces (e.g., He & Nakayama, 1995), and texture flows (e.g., Ben-Shahar, Scholl, & Zucker, 2007)—surely reflecting the same underlying general influence of structure on attention. Similarly, studies of individual cues to objecthood have revealed same-"object" advantages even to types of structure that lack key intuitive features of objects, such as closure (Avrahami, 1999; Marino & Scholl, 2005) and connectedness (Ben-Shahar et al., 2007; Feldman, 2007). We may conclude that "object-based attention" is really a more general phenomenon, in which spatial attention is influenced by visual structure (of many kinds).

Structure in Time

But how general? Object-based attention has traditionally been conceptualized in terms of a specific modality (*visual* structure) and dimension (influencing *spatial* attention). Here we explore whether object-based attention may reflect an even more abstract

This article was published Online First April 15, 2013.
Julian De Freitas, Brandon M. Liverence, and Brian J. Scholl, Department of Psychology, Yale University.
For helpful comments and conversation we thank Mike Dodd, Chaz Firestone, Alexandria Marino, Greg McCarthy, and Ève Poudrier. Readers may visit the authors' website at http://www.yale.edu/perception/
Correspondence concerning this article should be addressed to Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: brian.scholl@yale.edu

Effects like the one-is-more illusion suggest deep connections between ***vision*** *and* ***audition****, between the processing of* ***space*** *and* ***time***

Event-Based Warping: An Illusory Distortion of Time Within Events

Rui Zhe Goh[1,2,*], Hanbei Zhou[1,*], Chaz Firestone[1,2,†], Ian Phillips[1,2,†]

[1]Department of Psychological & Brain Sciences, Johns Hopkins University
[2]Department of Philosophy, Johns Hopkins University
*These authors should be considered joint first author
†These authors should be considered joint senior author
Correspondence: rgoh1@jhu.edu, hzhou43@jhu.edu, chaz@jhu.edu, ianbphillips@jhu.edu

Abstract

Objects and events are fundamental units of perception: Objects structure our experience of space, and events structure our experience of time. A striking and counterintuitive finding about object representation is that it can warp perceived space, such that stimuli within an object appear farther apart than stimuli in empty space. Might events distort perceived time in the same way objects distort perceived space? Here, four experiments (N=400 adults) show that they do: Just as stimuli within an *object* are perceived as farther apart in *space*, stimuli within an *event* are perceived as further apart in *time*. Such “Event-based Warping” is elicited both by events characterized by sound (E1), as well as events characterized by silence (E2). Moreover, these effects cannot be explained by surprise, distraction or attentional-cueing (E3), and also arise cross-modally (from audition to vision; E4). We suggest that object-based warping and event-based warping are both instances of a more general phenomenon in which representations of *structure* — whether in space or in time — generate powerful and analogous perceptual distortions.

Public significance statement

Perception segments sensory input received across continuous time into discrete events. For instance, when we listen to a musical piece, our auditory system segments continuous sound waves into the discrete musical notes, phrases, and motifs that we hear. Whereas much previous work has focused on the effects of event segmentation on downstream processes that depend on it (such as attention and memory), here we show that event segmentation can produce an *upstream* effect, distorting our experience of time itself — the very time that was segmented into events in the first place.

Keywords

event segmentation | time perception | multisensory processing

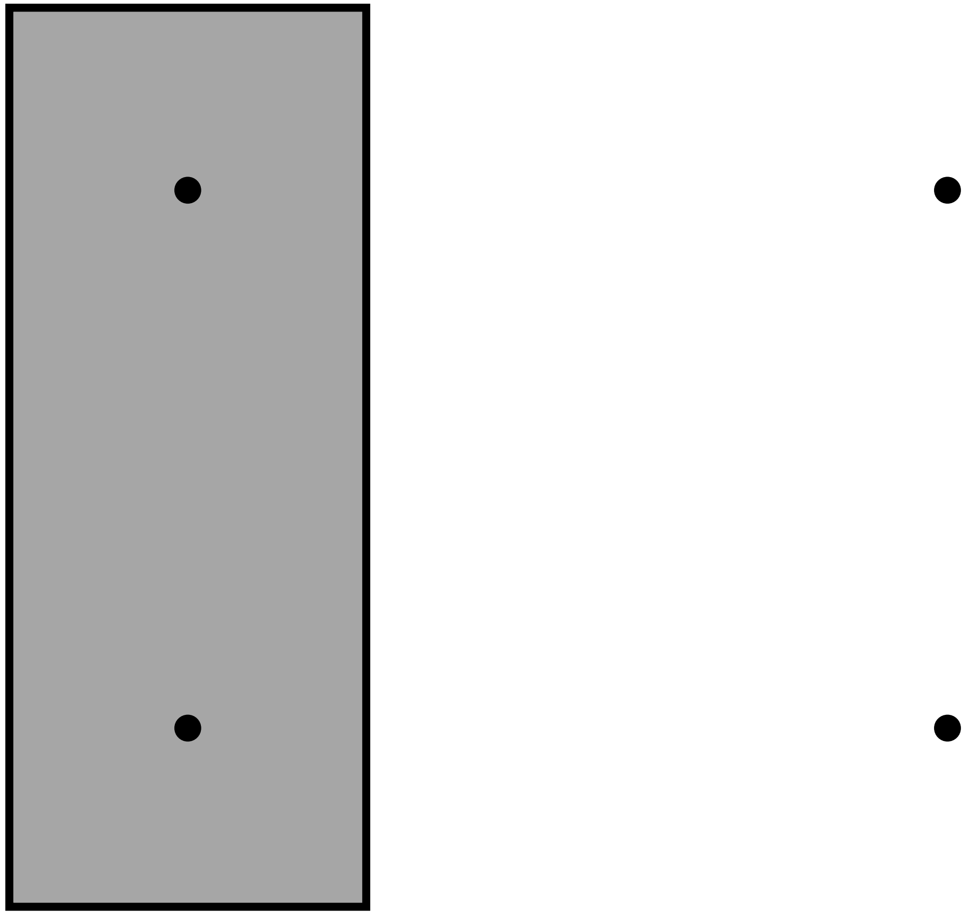

Object-based warping

Event-Based Warping: An Illusory Distortion of Time Within Events

Rui Zhe Goh[1,2,*], Hanbei Zhou[1,*], Chaz Firestone[1,2,†], Ian Phillips[1,2,†]

[1]Department of Psychological & Brain Sciences, Johns Hopkins University
[2]Department of Philosophy, Johns Hopkins University
*These authors should be considered joint first author
†These authors should be considered joint senior author
Correspondence: rgoh1@jhu.edu, hzhou43@jhu.edu, chaz@jhu.edu, ianbphillips@jhu.edu

Abstract
Objects and events are fundamental units of perception: Objects structure our experience of space, and events structure our experience of time. A striking and counterintuitive finding about object representation is that it can warp perceived space, such that stimuli within an object appear farther apart than stimuli in empty space. Might events distort perceived time in the same way objects distort perceived space? Here, four experiments (N=400 adults) show that they do: Just as stimuli within an *object* are perceived as farther apart in *space*, stimuli within an *event* are perceived as further apart in *time*. Such "Event-based Warping" is elicited both by events characterized by sound (E1), as well as events characterized by silence (E2). Moreover, these effects cannot be explained by surprise, distraction or attentional-cueing (E3), and also arise cross-modally (from audition to vision; E4). We suggest that object-based warping and event-based warping are both instances of a more general phenomenon in which representations of *structure* — whether in space or in time — generate powerful and analogous perceptual distortions.

Public significance statement
Perception segments sensory input received across continuous time into discrete events. For instance, when we listen to a musical piece, our auditory system segments continuous sound waves into the discrete musical notes, phrases, and motifs that we hear. Whereas much previous work has focused on the effects of event segmentation on downstream processes that depend on it (such as attention and memory), here we show that event segmentation can produce an *upstream* effect, distorting our experience of time itself — the very time that was segmented into events in the first place.

Keywords
event segmentation | time perception | multisensory processing

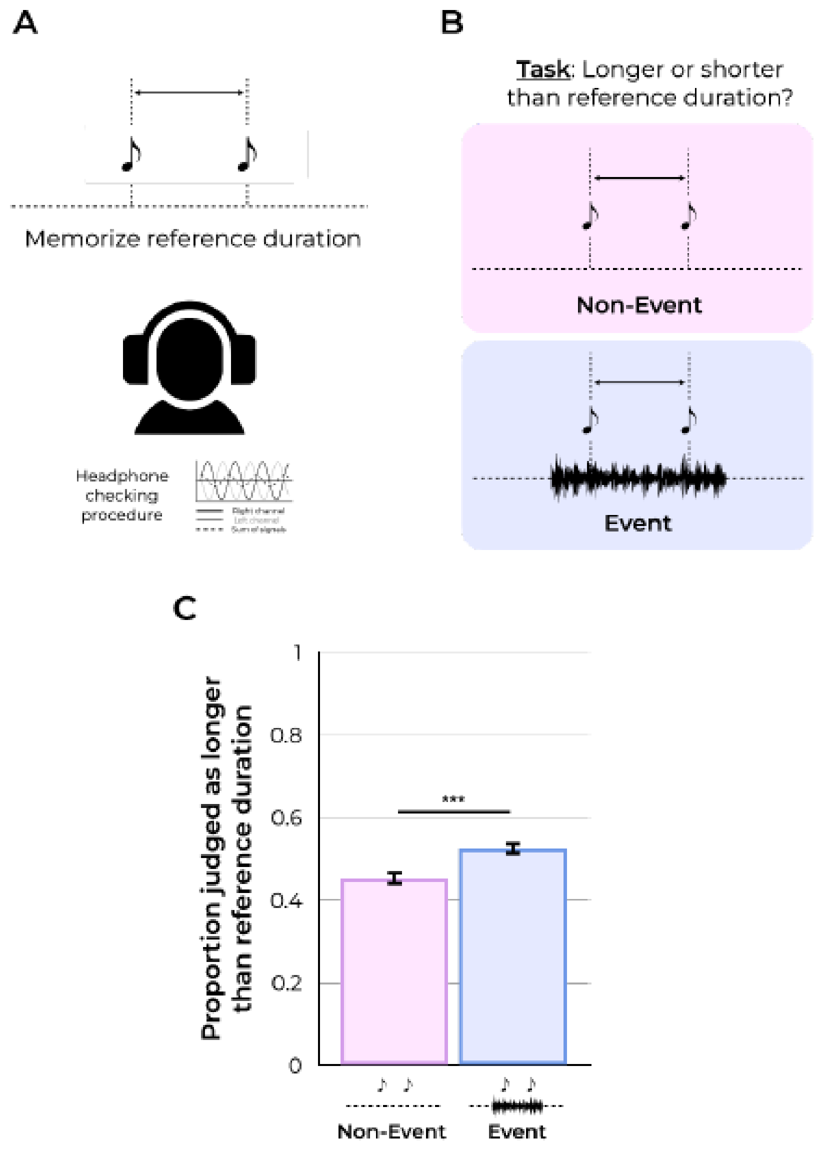

PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES SOCIAL SCIENCES

The perception of silence

Rui Zhe Goh[a,b,1], Ian B. Phillips[a,b,1,2], and Chaz Firestone[a,b,1,2]

Edited by Steven Luck, University of California, Davis, CA; received February 8, 2023; accepted May 10, 2023 by Editorial Board Member Michael S. Gazzaniga

Auditory perception is traditionally conceived as the perception of sounds—a friend's voice, a clap of thunder, a minor chord. However, daily life also seems to present us with experiences characterized by the absence of sound—a moment of silence, a gap between thunderclaps, the hush after a musical performance. In these cases, do we positively *hear* silence? Or do we just *fail to hear*, and merely judge or infer that it is silent? This longstanding question remains controversial in both the philosophy and science of perception, with prominent theories holding that sounds are the only objects of auditory experience and thus that our encounter with silence is cognitive, not perceptual. However, this debate has largely remained theoretical, without a key empirical test. Here, we introduce an empirical approach to this theoretical dispute, presenting experimental evidence that silence can be genuinely perceived (not just cognitively inferred). We ask whether silences can "substitute" for sounds in event-based auditory illusions—empirical signatures of auditory event representation in which auditory events distort perceived duration. Seven experiments introduce three "silence illusions"—the one-silence-is-more illusion, silence-based warping, and the oddball-silence illusion—each adapted from a prominent perceptual illusion previously thought to arise only from sounds. Subjects were immersed in ambient noise interrupted by silences structurally identical to the sounds in the original illusions. In all cases, silences elicited temporal distortions perfectly analogous to the illusions produced by sounds. Our results suggest that silence is truly heard, not merely inferred, introducing a general approach for studying the perception of absence.

absence perception | silence | event representation | temporal illusions

Significance

Do we only hear sounds? Or can we also hear silence? These questions are the subject of a centuries-old philosophical debate between two camps: the perceptual view (we literally hear silence), and the cognitive view (we only judge or infer silence). Here, we take an empirical approach to resolve this theoretical controversy. We show that silences can "substitute" for sounds in event-based auditory illusions. Seven experiments introduce three "silence illusions," adapted from perceptual illusions previously thought to arise only with sounds. In all cases, silences elicited temporal distortions perfectly analogous to their sound-based counterparts, suggesting that auditory processing treats moments of silence the way it treats sounds. Silence is truly perceived, not merely inferred.

What do we hear? The canonical answer is that auditory perception is the perception of *sounds* and their properties—the pitch of a friend's voice, the loudness of a thunderclap, the timbre of a minor chord. This traditional view has considerable pedigree, with influential historical sources holding that sounds are the sole objects of auditory perception (1, cf.2). It is also the answer favored in contemporary scholarship: Prominent scientific accounts conceive the fundamental units of auditory perception as sounds (or auditory streams comprised of sounds; ref. 3 and 4), and many philosophical theories agree, holding that "all auditory perception involves the perception of sound" (5) and that "if anything at all is heard, what is heard is necessarily a sound" (6) (see also refs. 7 and 8). The pervasiveness of this canonical view about the contents of auditory perception might seem unsurprising—what else might we hear, if not sound?

However, there has long been a stubborn and intuitive counterexample: experiences of *silence*, which are characterized by the absence of sound. Silence confronts us throughout our daily lives—consider an awkward pause in a conversation, a suspenseful gap between thunderclaps, or the hush at the end of a musical performance. What is the nature of these experiences?

Silence: Heard or Inferred?

One possibility is that experiences of silence are simply cases in which we *fail to hear*, and then use our faculties of reasoning and judgment to *infer* that it is silent. This interpretation is offered by those who defend the traditional sound-only view of audition, holding that an experience of silence is merely the "cognitive accompaniment of an absence of experience" and "is itself no form of hearing" (9). This cognitive view may be motivated by a deeper assumption about perception, namely that we can genuinely perceive only what is present in the world, not what is absent (9, 10). After all, one might think, absences are nonentities—they do not exist—and so can hardly impinge on our sensory apparatus.

However, an alternative possibility which arguably does more justice to our phenomenology is that we literally perceive silences. This interpretation has recently received

Author affiliations: [a]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD 21218; and [b]Department of Philosophy, Johns Hopkins University, Baltimore, MD 21218

Author contributions: R.Z.G., I.B.P., and C.F. designed research; R.Z.G. performed research; R.Z.G. analyzed data; and R.Z.G., I.B.P., and C.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. S.L. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: rgoh1@jhu.edu, ianbphillips@jhu.edu, or chaz@jhu.edu.

[2]I.B.P. and C.F. contributed equally to this work.

Published July 10, 2023.

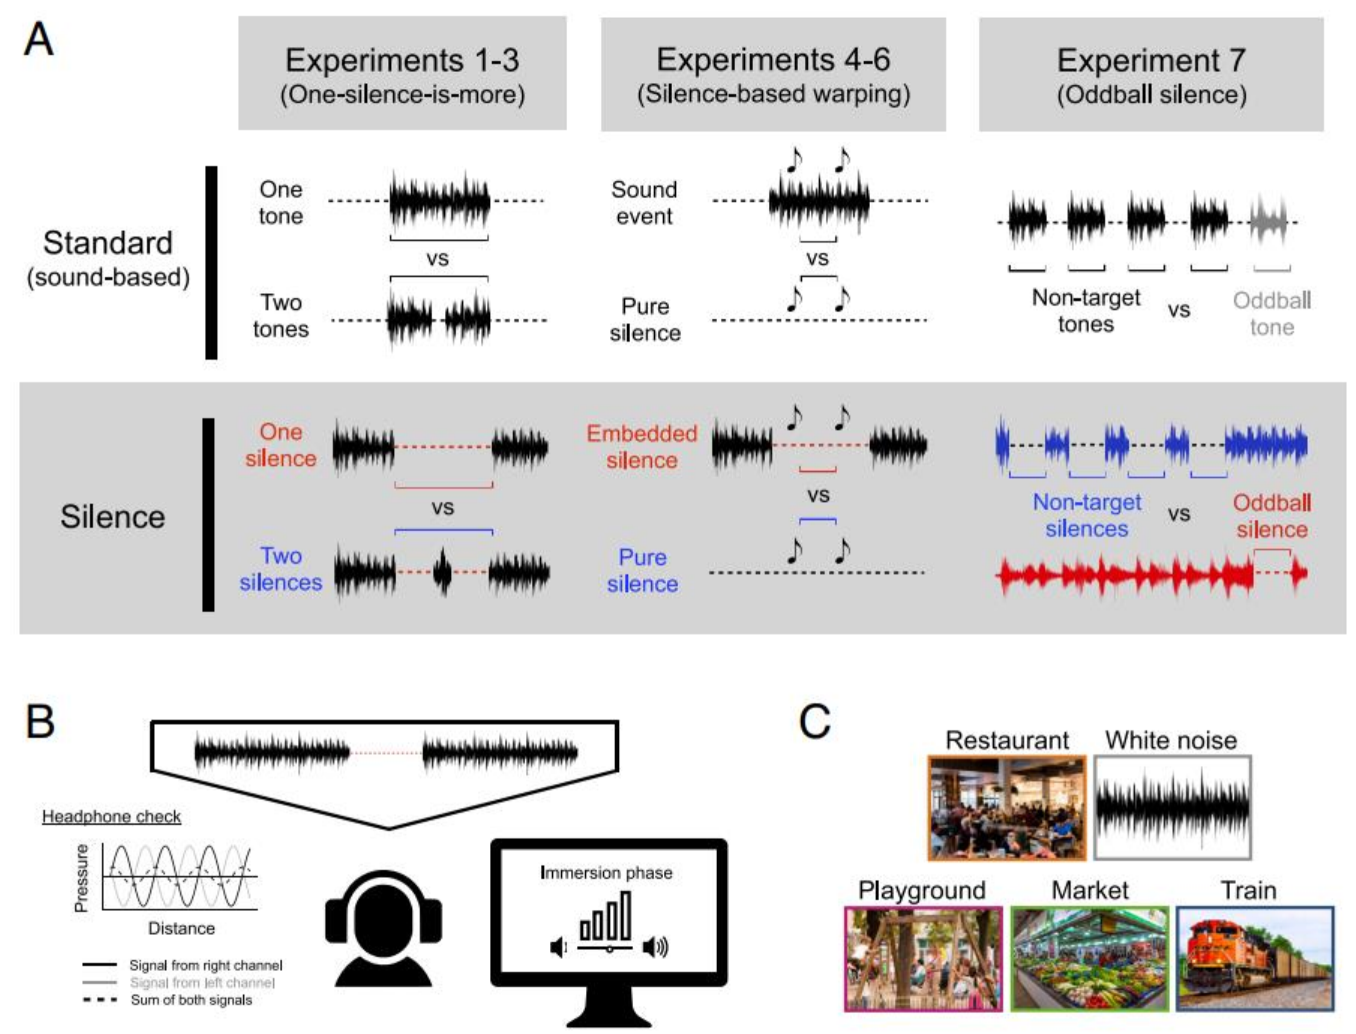

Other interactions between space and time.

Which is further? Green -> Blue or Blue -> Red?

The inverse...

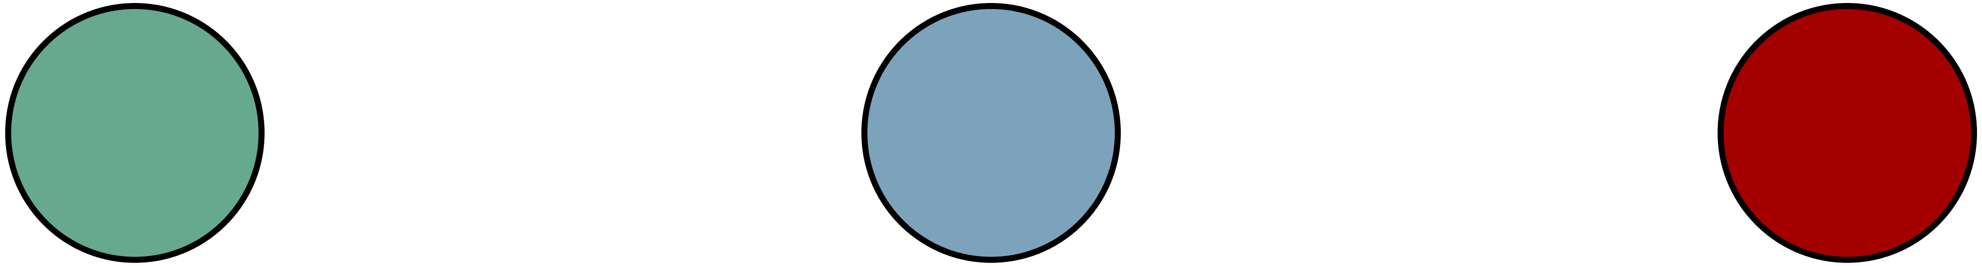

Interactions between time and memory

Research Article

ASSOCIATION FOR PSYCHOLOGICAL SCIENCE

Psychological Science
1–17

DOI: 10.1177/09567976221129533
www.psychologicalscience.org/PS
SAGE

Mnemonic Content and Hippocampal Patterns Shape Judgments of Time

Brynn E. Sherman[1], Sarah DuBrow[2], Jonathan Winawer[3], and Lila Davachi[4,5]

[1]Department of Psychology, University of Pennsylvania; [2]Department of Psychology, University of Oregon; [3]Department of Psychology and Center for Neural Science, New York University; [4]Department of Psychology, Columbia University; and [5]Department of Clinical Research, Nathan Kline Institute for Psychiatric Research

Abstract

Our experience of time can feel dilated or compressed, rather than reflecting true "clock time." Although many contextual factors influence the subjective perception of time, it is unclear how memory accessibility plays a role in constructing our experience of and memory for time. Here, we used a combination of behavioral and functional MRI measures in healthy young adults ($N = 147$) to ask the question of how memory is incorporated into temporal duration judgments. Behaviorally, we found that event boundaries, which have been shown to disrupt ongoing memory integration processes, result in the temporal compression of duration judgments. Additionally, using a multivoxel pattern similarity analysis of functional MRI data, we found that greater temporal pattern change in the left hippocampus within individual trials was associated with longer duration judgments. Together, these data suggest that mnemonic processes play a role in constructing representations of time.

Keywords

time estimation, memory, event boundaries, neuroimaging

Received 8/3/21; Revision accepted 9/9/22

> An hour, once it lodges in the queer element of the human spirit, may be stretched to fifty or a hundred times its clock length; on the other hand, an hour may be accurately represented on the timepiece of the mind by one second.
>
> —Virginia Woolf (1928)

On a busy vacation, time may escape you—by the time you go to the museum, grab lunch in the park, shop for souvenirs, and visit a historical site, the day may seem to have flown by. Yet when recalling the trip to a friend, that same day may feel like a week; all of those events could not have possibly occurred within the same few hours. This puzzle raises a fundamental question of how the structure of experience can paradoxically influence subjective impressions of time in experience and in reflection.

A great deal of work has focused on the latter: how the structure of experience influences how we remember elapsed time. In particular, abrupt shifts in context, or event boundaries, influence memory for time. Memory for the temporal order of events is disrupted across event boundaries (DuBrow & Davachi, 2013; Heusser et al., 2018; Horner et al., 2016). Further, intervals that contain an event boundary are remembered as longer than equivalently timed intervals without a boundary (Clewett et al., 2020; Ezzyat & Davachi, 2014), and mnemonic duration judgments scale with the number of events (Faber & Gennari, 2015, 2017; Lositsky et al., 2016). Such findings converge with the intuition that busy days feel long in memory: Events may serve to dilate time in memory. How, though, do events also result in time feeling subjectively shorter in the moment?

Event boundaries also influence memory on the more immediate time scale of working memory by reducing

Corresponding Author:
Brynn E. Sherman, University of Pennsylvania, Department of Psychology
Email: brynns@sas.upenn.edu

Research Article

Psychological Science
1–17

Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/09567976221129533
www.psychologicalscience.org/PS
SAGE

Mnemonic Content and Hippocampal Patterns Shape Judgments of Time

Brynn E. Sherman[1], Sarah DuBrow[2], Jonathan Winawer[3], and Lila Davachi[4,5]
[1]Department of Psychology, University of Pennsylvania; [2]Department of Psychology, University of Oregon; [3]Department of Psychology and Center for Neural Science, New York University; [4]Department of Psychology, Columbia University; and [5]Department of Clinical Research, Nathan Kline Institute for Psychiatric Research

Abstract
Our experience of time can feel dilated or compressed, rather than reflecting true "clock time." Although many contextual factors influence the subjective perception of time, it is unclear how memory accessibility plays a role in constructing our experience of and memory for time. Here, we used a combination of behavioral and functional MRI measures in healthy young adults (N = 147) to ask the question of how memory is incorporated into temporal duration judgments. Behaviorally, we found that event boundaries, which have been shown to disrupt ongoing memory integration processes, result in the temporal compression of duration judgments. Additionally, using a multivoxel pattern similarity analysis of functional MRI data, we found that greater temporal pattern change in the left hippocampus within individual trials was associated with longer duration judgments. Together, these data suggest that mnemonic processes play a role in constructing representations of time.

Keywords
time estimation, memory, event boundaries, neuroimaging

Received 8/3/21; Revision accepted 9/9/22

> An hour, once it lodges in the queer element of the human spirit, may be stretched to fifty or a hundred times its clock length; on the other hand, an hour may be accurately represented on the timepiece of the mind by one second.
>
> —Virginia Woolf (1928)

On a busy vacation, time may escape you—by the time you go to the museum, grab lunch in the park, shop for souvenirs, and visit a historical site, the day may seem to have flown by. Yet when recalling the trip to a friend, that same day may feel like a week; all of those events could not have possibly occurred within the same few hours. This puzzle raises a fundamental question of how the structure of experience can paradoxically influence subjective impressions of time in experience and in reflection.

A great deal of work has focused on the latter: how the structure of experience influences how we remember elapsed time. In particular, abrupt shifts in context, or event boundaries, influence memory for time. Memory for the temporal order of events is disrupted across event boundaries (DuBrow & Davachi, 2013; Heusser et al., 2018; Horner et al., 2016). Further, intervals that contain an event boundary are remembered as longer than equivalently timed intervals without a boundary (Clewett et al., 2020; Ezzyat & Davachi, 2014), and mnemonic duration judgments scale with the number of events (Faber & Gennari, 2015, 2017; Lositsky et al., 2016). Such findings converge with the intuition that busy days feel long in memory: Events may serve to dilate time in memory. How, though, do events also result in time feeling subjectively shorter in the moment?

Event boundaries also influence memory on the more immediate time scale of working memory by reducing

Corresponding Author:
Brynn E. Sherman, University of Pennsylvania, Department of Psychology
Email: brynns@sas.upenn.edu

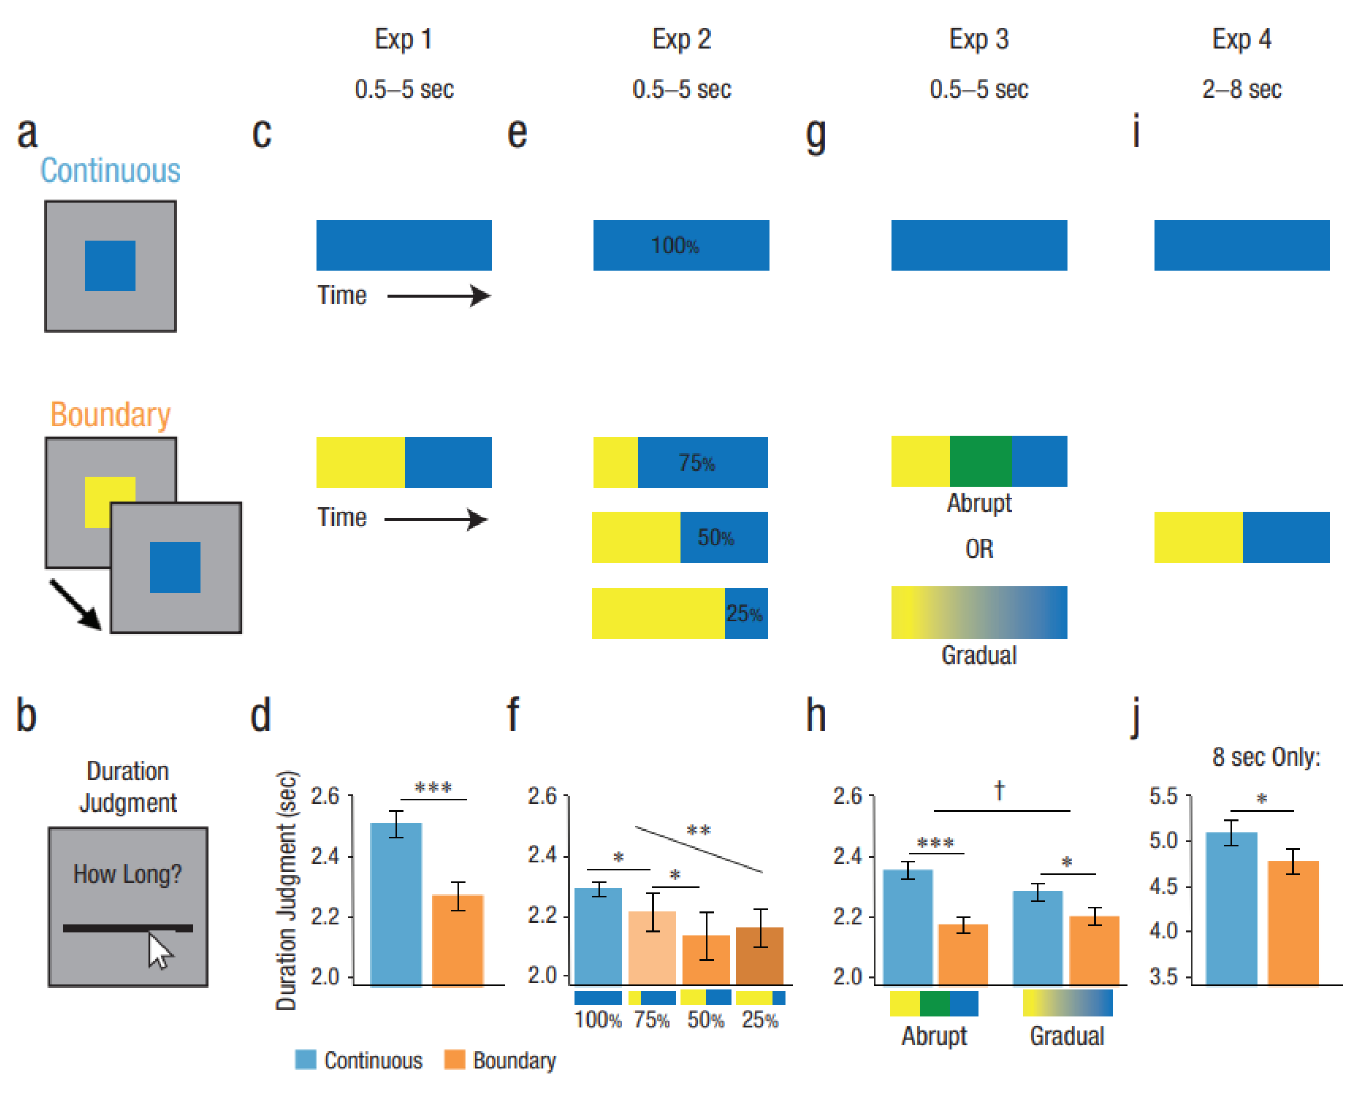

Research Article

Mnemonic Content and Hippocampal Patterns Shape Judgments of Time

Psychological Science
1–17

Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/09567976221129533
www.psychologicalscience.org/PS
SAGE

Brynn E. Sherman[1], Sarah DuBrow[2], Jonathan Winawer[3], and Lila Davachi[4,5]
[1]Department of Psychology, University of Pennsylvania; [2]Department of Psychology, University of Oregon; [3]Department of Psychology and Center for Neural Science, New York University; [4]Department of Psychology, Columbia University; and [5]Department of Clinical Research, Nathan Kline Institute for Psychiatric Research

Abstract

Our experience of time can feel dilated or compressed, rather than reflecting true "clock time." Although many contextual factors influence the subjective perception of time, it is unclear how memory accessibility plays a role in constructing our experience of and memory for time. Here, we used a combination of behavioral and functional MRI measures in healthy young adults ($N = 147$) to ask the question of how memory is incorporated into temporal duration judgments. Behaviorally, we found that event boundaries, which have been shown to disrupt ongoing memory integration processes, result in the temporal compression of duration judgments. Additionally, using a multivoxel pattern similarity analysis of functional MRI data, we found that greater temporal pattern change in the left hippocampus within individual trials was associated with longer duration judgments. Together, these data suggest that mnemonic processes play a role in constructing representations of time.

Keywords
time estimation, memory, event boundaries, neuroimaging

Received 8/3/21; Revision accepted 9/9/22

> An hour, once it lodges in the queer element of the human spirit, may be stretched to fifty or a hundred times its clock length; on the other hand, an hour may be accurately represented on the timepiece of the mind by one second.
>
> —Virginia Woolf (1928)

On a busy vacation, time may escape you—by the time you go to the museum, grab lunch in the park, shop for souvenirs, and visit a historical site, the day may seem to have flown by. Yet when recalling the trip to a friend, that same day may feel like a week; all of those events could not have possibly occurred within the same few hours. This puzzle raises a fundamental question of how the structure of experience can paradoxically influence subjective impressions of time in experience and in reflection.

A great deal of work has focused on the latter: how the structure of experience influences how we remember elapsed time. In particular, abrupt shifts in context, or event boundaries, influence memory for time. Memory for the temporal order of events is disrupted across event boundaries (DuBrow & Davachi, 2013; Heusser et al., 2018; Horner et al., 2016). Further, intervals that contain an event boundary are remembered as longer than equivalently timed intervals without a boundary (Clewett et al., 2020; Ezzyat & Davachi, 2014), and mnemonic duration judgments scale with the number of events (Faber & Gennari, 2015, 2017; Lositsky et al., 2016). Such findings converge with the intuition that busy days feel long in memory: Events may serve to dilate time in memory. How, though, do events also result in time feeling subjectively shorter in the moment?

Event boundaries also influence memory on the more immediate time scale of working memory by reducing

Corresponding Author:
Brynn E. Sherman, University of Pennsylvania, Department of Psychology
Email: brynns@sas.upenn.edu

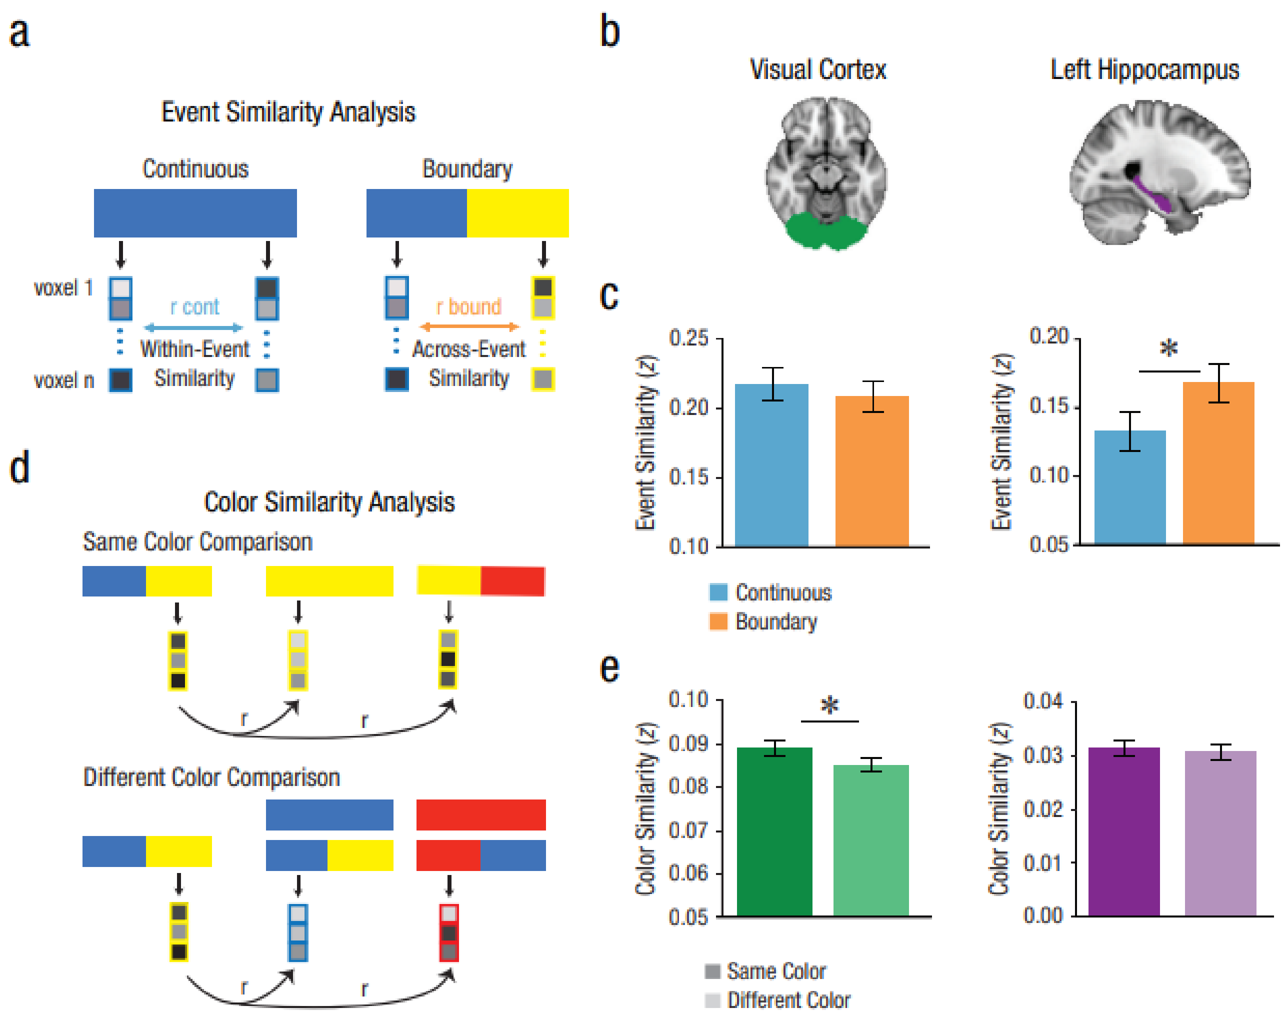

Research Article

Mnemonic Content and Hippocampal Patterns Shape Judgments of Time

"Our everyday experiences convey a powerful truth: Our perception of time often diverges from the reality of time. When enjoying an active vacation with family, time moves quickly: Hours go by in minutes. When sitting through an unnecessary meeting, time moves slowly: Minutes go by in hours. What is the origin of these phenomenologically compelling illusions of time perception? Past research has examined how a range of specific factors, from emotions to blinking, contributes to the distortion of time. Here, in contrast, we evaluated how the content and accessibility of our memories shape time perception. We show that context shifts, known to disrupt memory processing, also lead to robust contractions of perceived time."

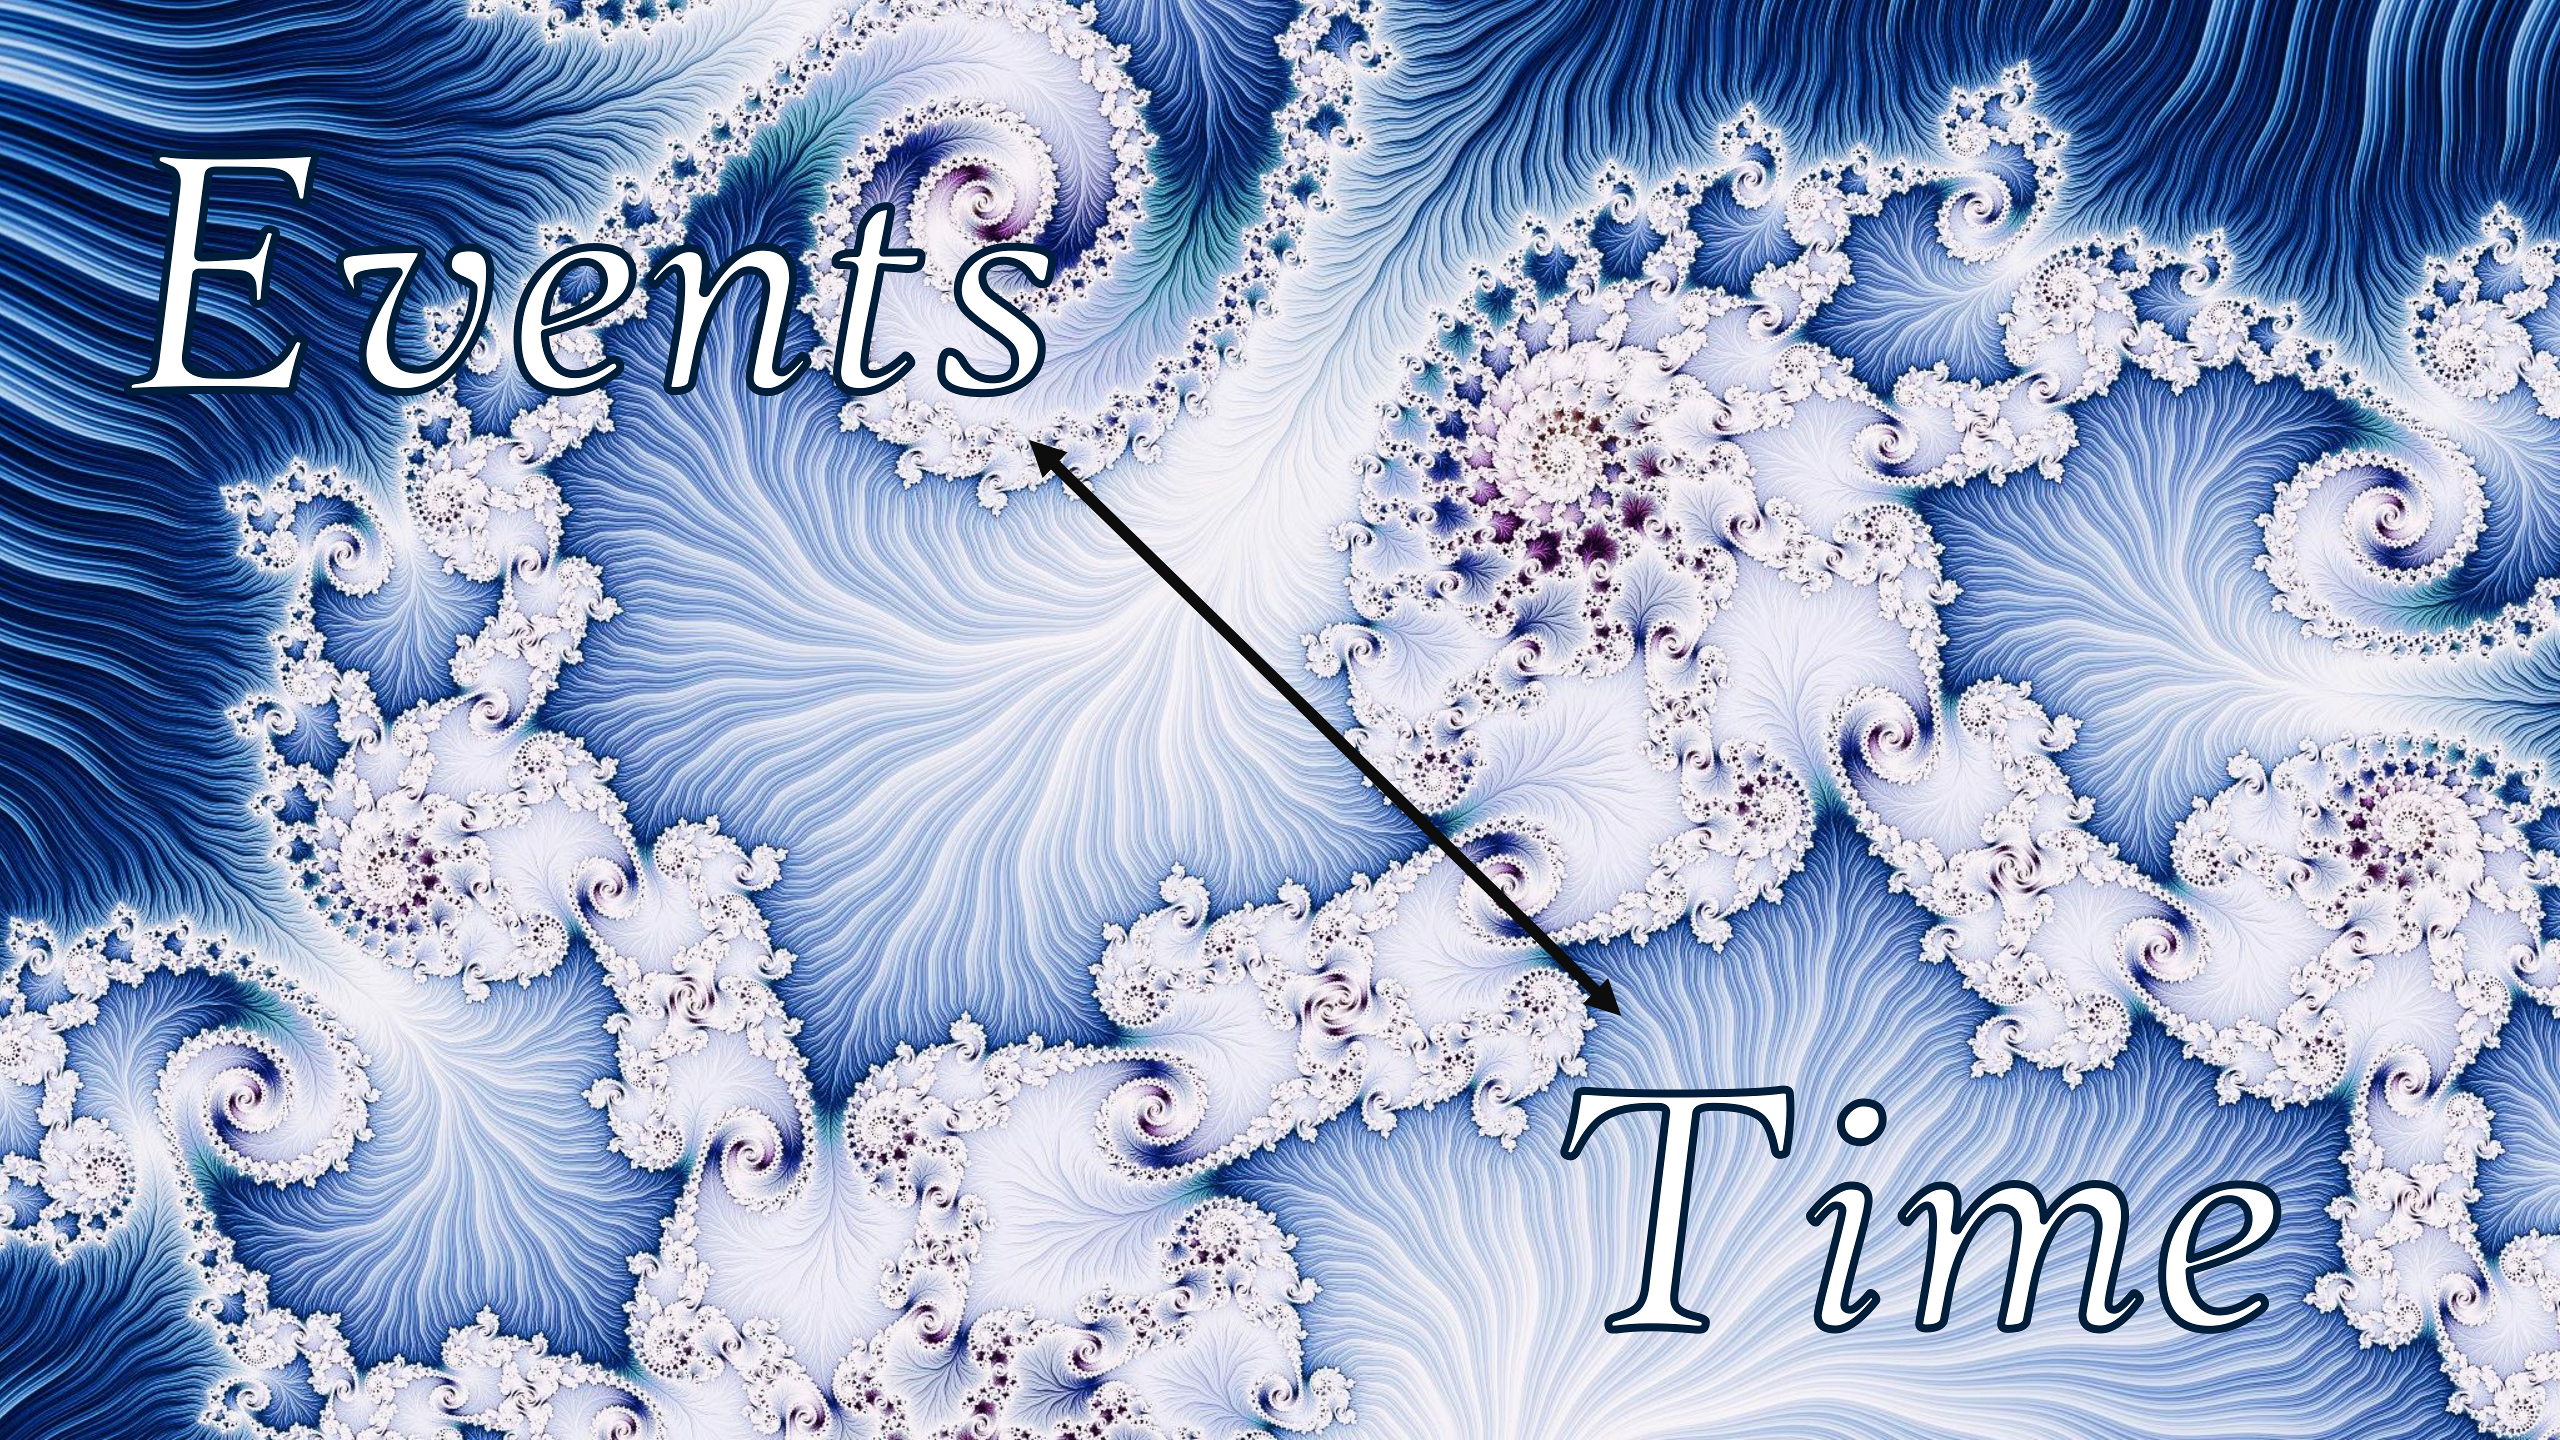
Events
Time

Psychonomic Bulletin & Review
https://doi.org/10.3758/s13423-023-02311-4

THEORETICAL/REVIEW

More than a moment: What does it mean to call something an 'event'?

Tristan S. Yates[1] · Brynn E. Sherman[2] · Sami R. Yousif[2]

Accepted: 14 May 2023

Abstract
Experiences are stored in the mind as discrete mental units, or 'events,' which influence—and are influenced by—attention, learning, and memory. In this way, the notion of an 'event' is foundational to cognitive science. However, despite tremendous progress in understanding the behavioral and neural signatures of events, there is no agreed-upon definition of an event. Here, we discuss different theoretical frameworks of event perception and memory, noting what they can and cannot account for in the literature. We then highlight key aspects of events that we believe should be accounted for in theories of event processing—in particular, we argue that the structure and substance of events should be better reflected in our theories and paradigms. Finally, we discuss empirical gaps in the event cognition literature and what the future of event cognition research may look like.

Keywords Event cognition · Memory · Attention · Prediction error · Perception · Event boundaries

Time is divided into years, which are divided into months, days, hours, and minutes. However, our experience is not represented in these arbitrary units; rather, our lives are divided into *events* which may span moments, months, or decades. This is to say that the units of *time* (months, days, hours, and minutes) are not the same as the units of *experience* ('events').

The notion that experience is subdivided into events has become a foundational idea in cognitive science. It offers a way of describing how mental representations of experience (which are often discrete) differ from reality (which is often more continuous). Once we label those mental representations (i.e., as events), we can search for them in the mind and brain (Baldassano et al., 2017; Ezzyat & Davachi, 2011; Heusser et al., 2016) and make predictions about how people will attend to and/or remember certain experiences.

But what is an event exactly? Events have been defined based on subjective ratings (e.g., Newtson, 1973), neural representational similarity measures (Baldassano et al., 2017; Geerligs et al., 2021), and by measuring the influence of temporal structure on memory (Clewett & Davachi, 2017) or perception (Liverence & Scholl, 2012; Meyerhoff et al., 2015; Sherman, DuBrow et al., 2023; Yousif & Scholl, 2019). Events have been described as discrete moments in time lasting as short as hundreds of milliseconds (Michotte, 1963) as well as extended periods of time lasting as long as centuries (Teigen et al., 2017; see also Sastre et al., 2022). They have been likened to visual objecthood and attention (see, e.g., Casati & Varzi, 2008; De Freitas et al., 2014; Ji & Papafragou, 2022; Tversky et al., 2008; Zacks & Tversky, 2001), and they have also been argued to reflect inferences about the causal structure of the world (Shin & DuBrow, 2021; see also Radvansky, 2012). They have been studied in vision (e.g., Tauzin, 2015), in memory (e.g., DuBrow & Davachi, 2013), in language (e.g., Ünal et al., 2021), and in more naturalistic scenarios (e.g., Sastre et al., 2022; Swallow et al., 2018). Boundaries between them can be triggered by anything from movement through a doorway (Radvansky & Copeland, 2006; Radvansky et al., 2010, 2011; see also Radvansky, 2012), to a change in background color (Heusser et al., 2018), to the movement of dots (Ongchoco & Scholl, 2019). This complexity has led some to avoid concrete definitions altogether; Schwartz (2008) presciently wrote that events are merely "what we make of them" (p. 1).

In some ways, it may be easy to say what an event *is*: It is any discrete experience. However, it is much harder to say what an event *is not*. If something as simple as walking

✉ Tristan S. Yates
tristan.yates@yale.edu

✉ Brynn E. Sherman
brynns@sas.upenn.edu

✉ Sami R. Yousif
sryousif@sas.upenn.edu

1 Department of Psychology, Yale University, New Haven, CT, USA

2 Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Published online: 05 July 2023

Psychonomic Bulletin & Review
https://doi.org/10.3758/s13423-023-02311-4

THEORETICAL/REVIEW

More than a moment: What does it mean to call something an 'event'?

Tristan S. Yates[1] · Brynn E. Sherman[2] · Sami R. Yousif[2]

Accepted: 14 May 2023

Abstract

Experiences are stored in the mind as discrete mental unit
learning, and memory. In this way, the notion of an 'even
progress in understanding the behavioral and neural si
we discuss different theoretical frameworks of event p
literature. We then highlight key aspects of events th
particular, we argue that the structure and substance
we discuss empirical gaps in the event cognition li

Keywords Event cognition · Memory · Attention

Time is divided into years, which are divided int
days, hours, and minutes. However, our experience i
resented in these arbitrary units; rather, our lives are
into *events* which may span moments, months, or de
This is to say that the units of *time* (months, days, hours,
minutes) are not the same as the units of *experience* ('events

The notion that experience is subdivided into events has become a foundational idea in cognitive science. It offers a way of describing how mental representations of experience (which are often discrete) differ from reality (which is often more continuous). Once we label those mental representations (i.e., as events), we can search for them in the mind and brain (Baldassano et al., 2017; Ezzyat & Davachi, 2011; Heusser et al., 2016) and make predictions about how people will attend to and/or remember certain experiences.

But what is an event exactly? Events have been defined based on subjective ratings (e.g., Newtson, 1973), neural

017;
see … visual
objecthood … 2008; De
Freitas et al., 2014; … 2022; Tversky et al., 2008; Zacks & Tversky, 2001), and they have also been argued to reflect inferences about the causal structure of the world (Shin & DuBrow, 2021; see also Radvansky, 2012). They have been studied in vision (e.g., Tauzin, 2015), in memory (e.g., DuBrow & Davachi, 2013), in language (e.g., Ünal et al., 2021), and in more naturalistic scenarios (e.g., Sastre et al., 2022; Swallow et al., 2018). Boundaries between them can be triggered by anything from movement through a doorway (Radvansky & Copeland, 2006; Radvansky et al., 2010, 2011; see also Radvansky, 2012), to a change in background color (Heusser et al., 2018), to the movement of dots (Ongchoco & Scholl, 2019). This complexity has led some to avoid concrete definitions altogether; Schwartz (2008) presciently wrote that events are merely "what we make of them" (p. 1).

In some ways, it may be easy to say what an event *is*: It is any discrete experience. However, it is much harder to say what an event *is not*. If something as simple as walking

✉ Tristan S. Yates
tristan.yates@yale.edu

✉ Brynn E. Sherman
brynns@sas.upenn.edu

✉ Sami R. Yousif
sryousif@sas.upenn.edu

[1] Department of Psychology, Yale University, New Haven, CT, USA

[2] Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

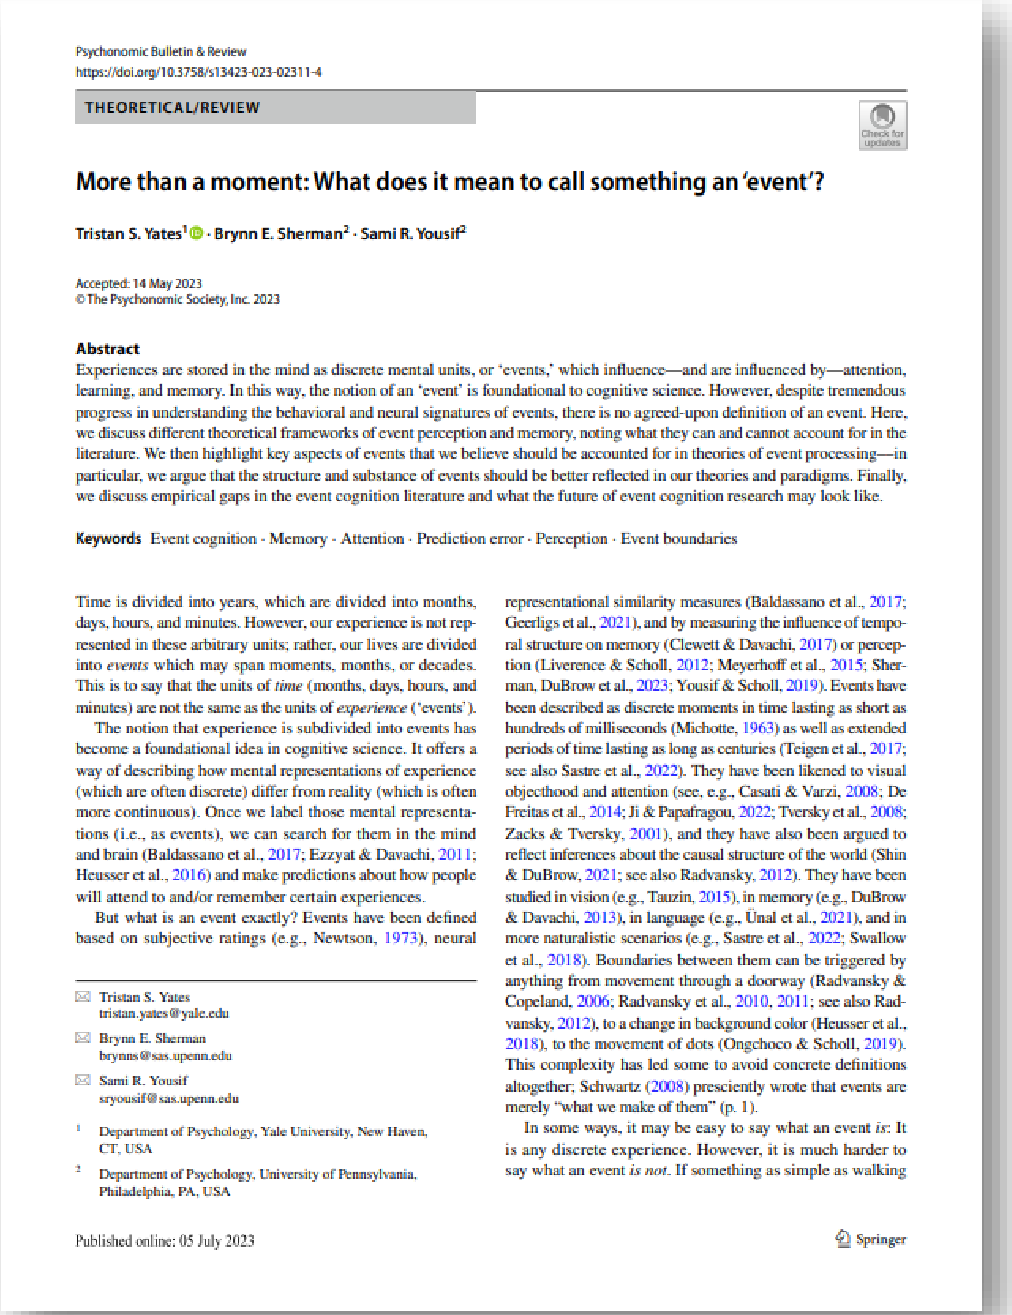
Psychonomic Bulletin & Review
https://doi.org/10.3758/s13423-023-02311-4

THEORETICAL/REVIEW

More than a moment: What does it mean to call something an 'event'?

Tristan S. Yates[1] · Brynn E. Sherman[2] · Sami R. Yousif[2]

Accepted: 14 May 2023

Abstract
Experiences are stored in the mind as discrete mental units, or 'events,' which influence—and are influenced by—attention, learning, and memory. In this way, the notion of an 'event' is foundational to cognitive science. However, despite tremendous progress in understanding the behavioral and neural signatures of events, there is no agreed-upon definition of an event. Here, we discuss different theoretical frameworks of event perception and memory, noting what they can and cannot account for in the literature. We then highlight key aspects of events that we believe should be accounted for in theories of event processing—in particular, we argue that the structure and substance of events should be better reflected in our theories and paradigms. Finally, we discuss empirical gaps in the event cognition literature and what the future of event cognition research may look like.

Keywords Event cognition · Memory · Attention · Prediction error · Perception · Event boundaries

Time is divided into years, which are divided into months, days, hours, and minutes. However, our experience is not represented in these arbitrary units; rather, our lives are divided into *events* which may span moments, months, or decades. This is to say that the units of *time* (months, days, hours, and minutes) are not the same as the units of *experience* ('events').

The notion that experience is subdivided into events has become a foundational idea in cognitive science. It offers a way of describing how mental representations of experience (which are often discrete) differ from reality (which is often more continuous). Once we label those mental representations (i.e., as events), we can search for them in the mind and brain (Baldassano et al., 2017; Ezzyat & Davachi, 2011; Heusser et al., 2016) and make predictions about how people will attend to and/or remember certain experiences.

But what is an event exactly? Events have been defined based on subjective ratings (e.g., Newtson, 1973), neural representational similarity measures (Baldassano et al., 2017; Geerligs et al., 2021), and by measuring the influence of temporal structure on memory (Clewett & Davachi, 2017) or perception (Liverence & Scholl, 2012; Meyerhoff et al., 2015; Sherman, DuBrow et al., 2023; Yousif & Scholl, 2019). Events have been described as discrete moments in time lasting as short as hundreds of milliseconds (Michotte, 1963) as well as extended periods of time lasting as long as centuries (Teigen et al., 2017; see also Sastre et al., 2022). They have been likened to visual objecthood and attention (see, e.g., Casati & Varzi, 2008; De Freitas et al., 2014; Ji & Papafragou, 2022; Tversky et al., 2008; Zacks & Tversky, 2001), and they have also been argued to reflect inferences about the causal structure of the world (Shin & DuBrow, 2021; see also Radvansky, 2012). They have been studied in vision (e.g., Tauzin, 2015), in memory (e.g., DuBrow & Davachi, 2013), in language (e.g., Ünal et al., 2021), and in more naturalistic scenarios (e.g., Sastre et al., 2022; Swallow et al., 2018). Boundaries between them can be triggered by anything from movement through a doorway (Radvansky & Copeland, 2006; Radvansky et al., 2010, 2011; see also Radvansky, 2012), to a change in background color (Heusser et al., 2018), to the movement of dots (Ongchoco & Scholl, 2019). This complexity has led some to avoid concrete definitions altogether; Schwartz (2008) presciently wrote that events are merely "what we make of them" (p. 1).

In some ways, it may be easy to say what an event *is*: It is any discrete experience. However, it is much harder to say what an event *is not*. If something as simple as walking

✉ Tristan S. Yates
tristan.yates@yale.edu

✉ Brynn E. Sherman
brynns@sas.upenn.edu

✉ Sami R. Yousif
sryousif@sas.upenn.edu

1 Department of Psychology, Yale University, New Haven, CT, USA

2 Department of Psychology, University of Pennsylvania, Philadelphia, PA, USA

Published online: 05 July 2023

Springer

!!

Events : time :: objects : space

Attentional rhythm

Object-based attention

Journal of Experimental Psychology:
Human Perception and Performance
2012, Vol. 38, No. 3, 549–554

0096-1523/12/$12.00 DOI: 10.1037/a0027228

OBSERVATION

Discrete Events as Units of Perceived Time

Brandon M. Liverence and Brian J. Scholl
Yale University

In visual images, we perceive both space (as a continuous visual medium) and objects (that inhabit space). Similarly, in dynamic visual experience, we perceive both continuous time and discrete events. What is the relationship between these units of experience? The most intuitive answer may be similar to the spatial case: time is perceived as an underlying medium, which is later segmented into discrete event representations. Here we explore the opposite possibility—that our subjective experience of time itself can be influenced by how durations are temporally segmented, beyond more general effects of change and complexity. We show that the way in which a continuous dynamic display is segmented into discrete units (via a *path shuffling* manipulation) greatly influences duration judgments, independent of psychophysical factors previously implicated in time perception, such as overall stimulus energy, attention and predictability. It seems that we may use the passage of discrete events—and the boundaries between them—in our subjective experience as part of the raw material for inferring the strength of the underlying "current" of time.

Keywords: time dilation, event perception, segmentation, object persistence

"Time is a sort of river of passing events, and strong is its current."

—Marcus Aurelius (1909–1914, [translated version] *Meditations, IV, 43*)

In the physical world, time seems to flow inexorably forward, and always at the same objective pace (see Carroll, 2010). In our subjective experience, in contrast, the pace at which time seems to flow can speed up or slow down, as a function of lower-level variables such as stimulus energy (e.g., Brown, 1995), and higher-level factors such as attention (e.g., Tse, Intriligator, Rivest, & Cavanagh, 2004) and predictability (e.g., Pariyadath & Eagleman, 2007). Continuous visual input is also carved up in our subjective experience into discrete units, so that we perceive durations as being filled by particular events. There is considerable agreement between observers on where such units begin and end across a wide range of stimuli (e.g., Zacks, 2004; Zacks, Speer, & Reynolds, 2009), and discrete event representations also influence attention and memory for dynamic scenes (e.g., Levin & Varakin, 2004; Swallow, Zacks, & Abrams, 2009).

How do these two primary features of our dynamic perceptual experience—time and events—relate to each other? Intuitively, time is the underlying medium for event segmentation—the "river" on which events flow. Here, in contrast, we explore the possibility that subjective time is influenced by how temporal experience is segmented into discrete events. This study is inspired by analogous explorations of space and object perception: though space seems like the underlying medium for object segregation, the way in which perceptual experience is segmented into discrete objects influences spatial perception (e.g., Vickery & Chun, 2010), attention (e.g., Scholl, 2001), and memory (e.g., Luck & Vogel, 1997). Might the underlying "units" of time perception be similarly discrete?

Experiment 1: Path Shuffling

Here we explore whether temporal judgments are influenced by incidental visual event segmentation. To isolate a role for segmentation, we used animations of simple shapes that could be carefully controlled, and segmentation cues that relied on simple and universal visual properties such as objecthood and spatiotemporal continuity. In particular, an object moving along a continuous spatiotemporal path is perceived to be a single persisting individual, whereas disruptions to perceived spatiotemporal continuity result in the percept of multiple independent objects (for a review, see Scholl, 2007). This kind of cue—the perceived disappearance of old objects followed by the appearance of new objects—has been found in previous work to induce event boundaries (e.g., Zacks et al., 2009).

To induce event segmentation in the current study, we introduced stark spatiotemporal gaps into brief animations of a single moving object (e.g., Figure 1). When the resulting segments are played "in order" (on "Forwards" trials), observers see a single

This article was published Online First February 27, 2012.

Brandon M. Liverence and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and/or comments on earlier drafts, we thank Marc Buehner, Todd Horowitz, Josh New, Barbara Tversky, John Wearden, and the members of the Scholl and Chun labs at Yale University.

Correspondence concerning this article should be addressed to Brandon M. Liverence or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: brandon.liverence@yale.edu, brian.scholl@yale.edu

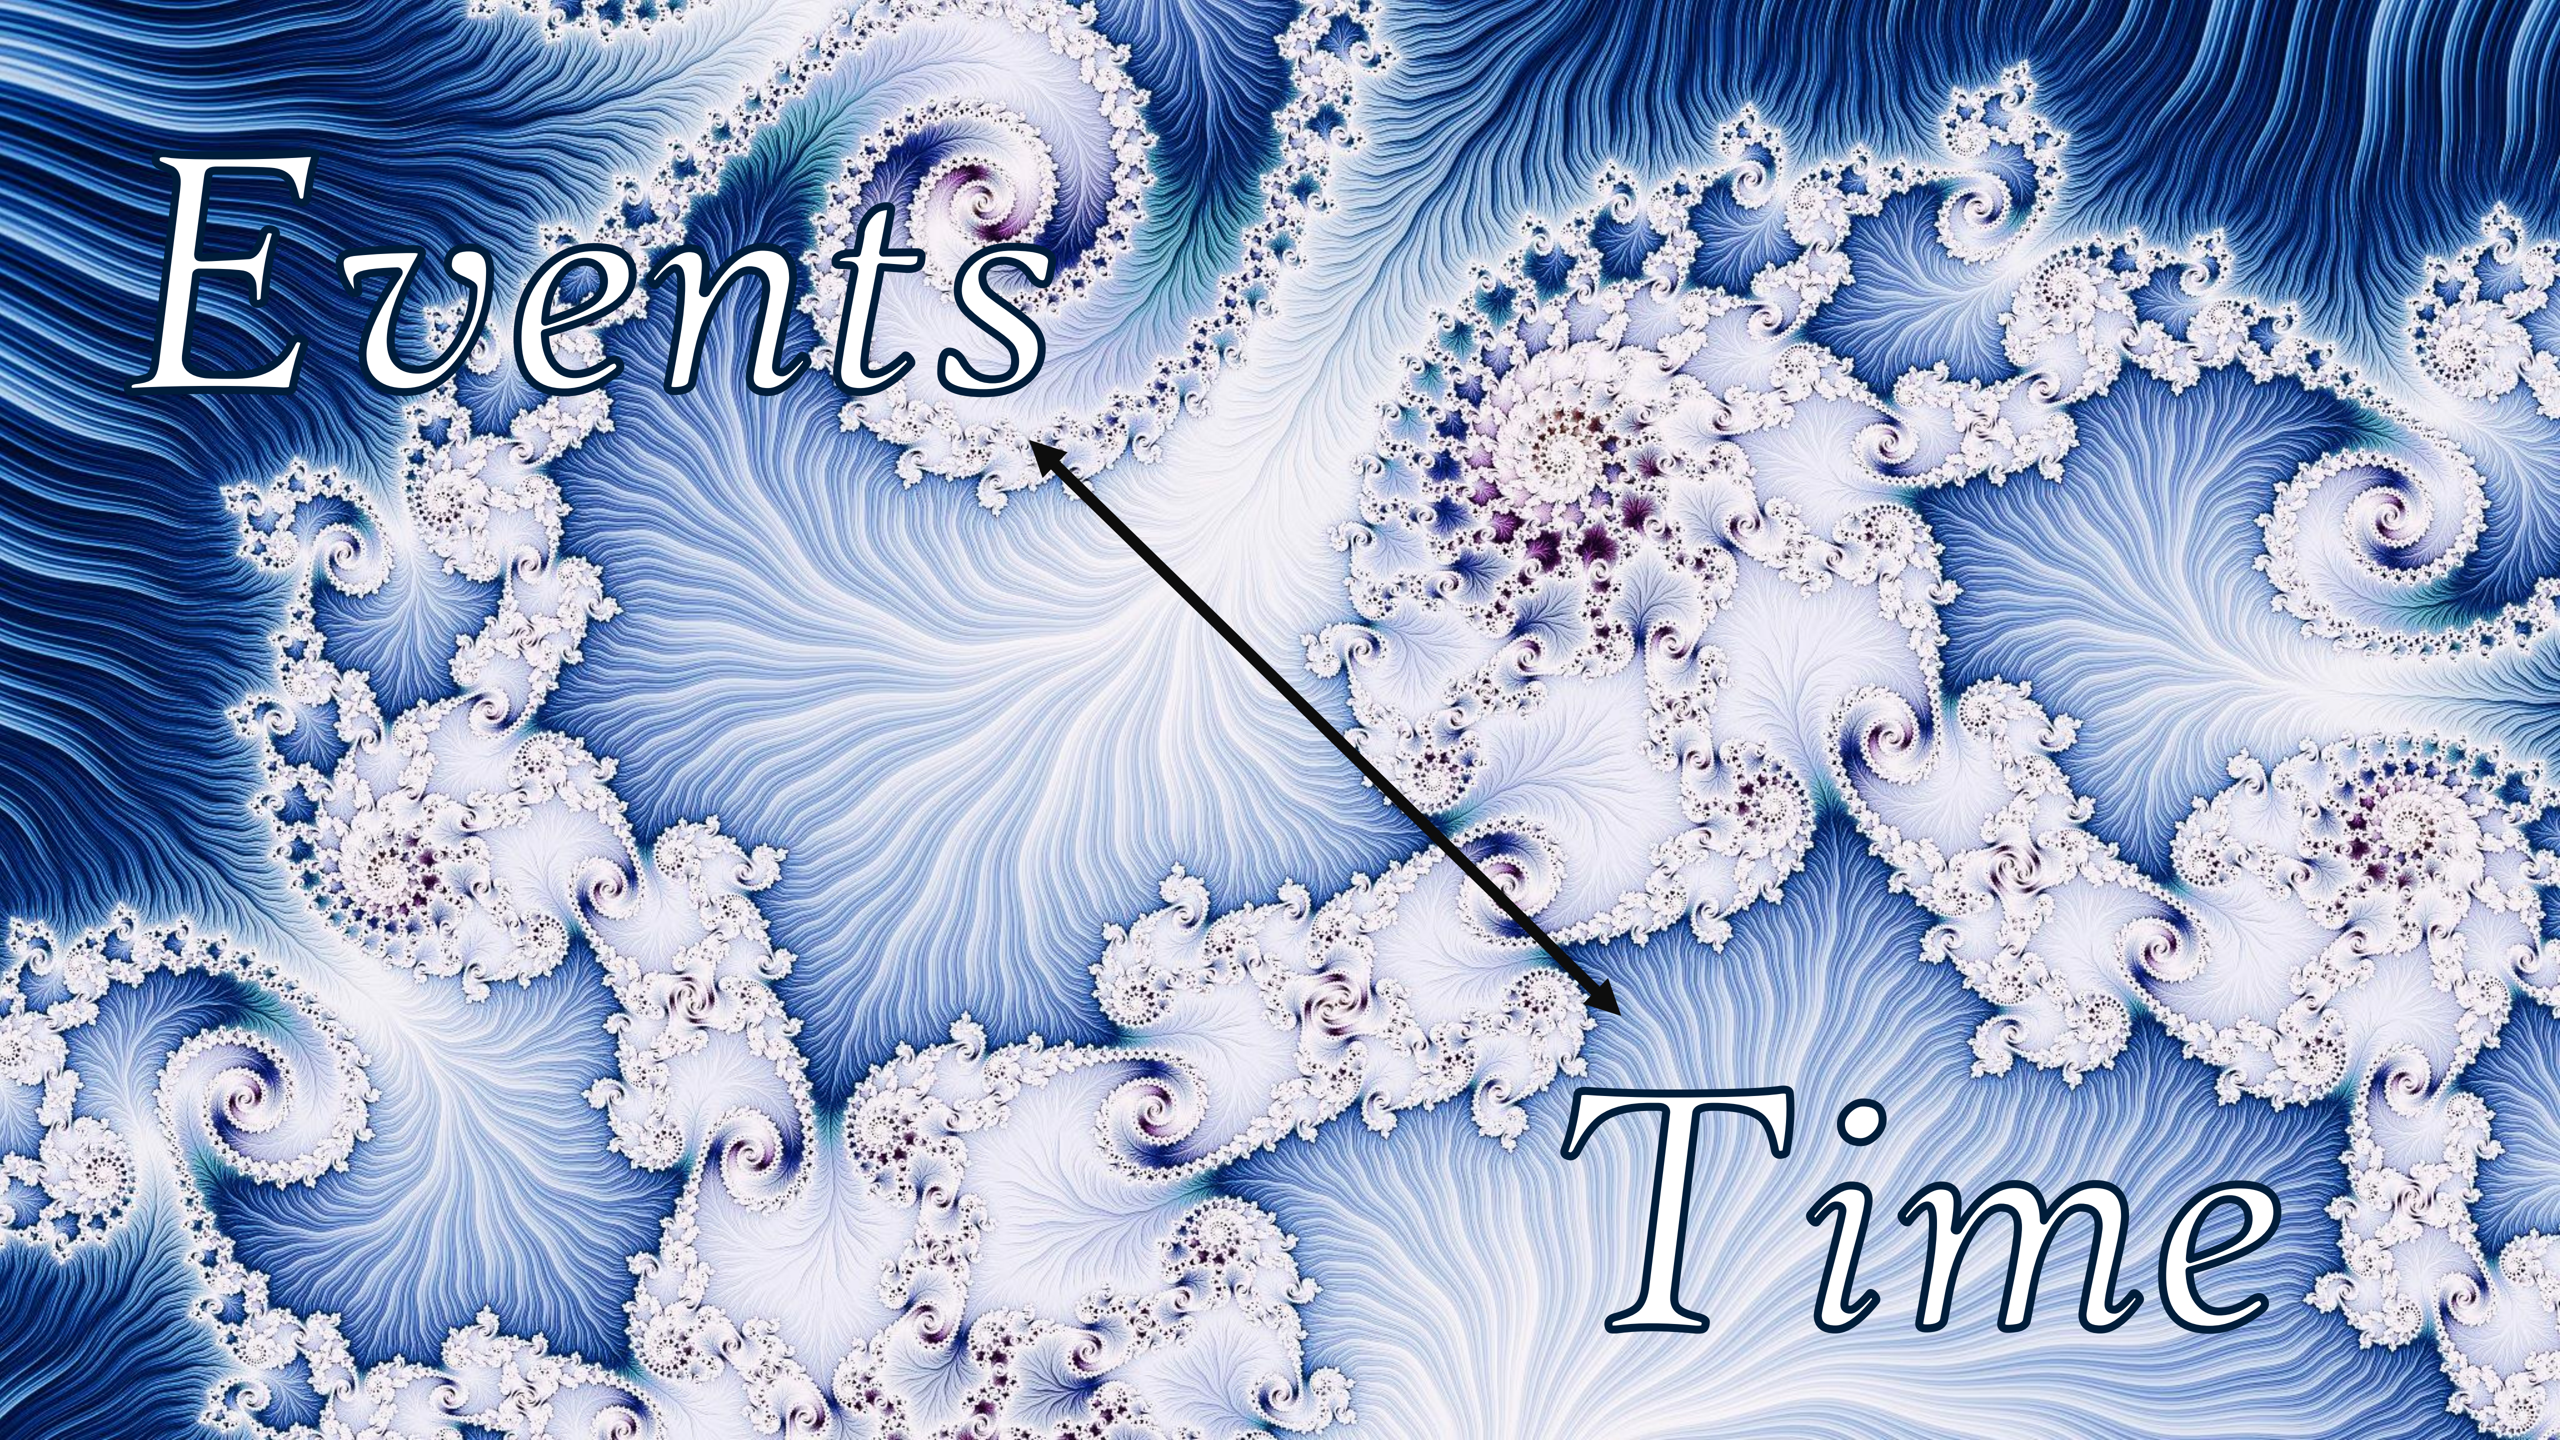
Events
Time

Cognition 249 (2024) 105833

Contents lists available at ScienceDirect

Cognition

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

Event representation at the scale of ordinary experience

Sami R. Yousif [a,*], Sarah Hye-yeon Lee [b], Brynn E. Sherman [a], Anna Papafragou [b]

[a] University of Pennsylvania, Department of Psychology, United States of America
[b] University of Pennsylvania, Department of Linguistics, United States of America

ARTICLE INFO

Keywords:
Event representation
Episodic memory
Temporal memory

ABSTRACT

Weeks are divided into weekdays and weekends; years into semesters and seasons; lives into stages like childhood, adulthood, and adolescence. How does the structure of experience shape memory? Though much work has examined event representation in human cognition, little work has explored event representation at the scale of ordinary experience. Here, we use shared experiences — in the form of popular television shows — to explore how memories are shaped by event structure at a large scale. We find that memories for events in these shows exhibit several hallmarks of event cognition. Namely, we find that memories are organized with respect to their event structure (boundaries), and that beginnings and endings are better remembered at multiple levels of the event hierarchy simultaneously. These patterns seem to be partially, but not fully, explained by the perceived story-relevance of events. Lastly, using a longitudinal design, we also show how event representations evolve over periods of several months. These results offer an understanding of event cognition at the scale of ordinary human lives.

Experiences are complex and varied, but our memories of them are structured in the form of discrete events. We remember our high school years as distinct from our college years and those years as distinct from everything after. We remember our summer months as distinct from the grind of the semester. We remember our weekdays as distinct from our weekends.

A significant body of work in psychology and cognitive science has been interested in the nature of event representation — how it is that we carve *continuous* experience up into *discrete* 'units' (for reviews, see Radvansky & Zacks, 2011; Shipley & Zacks, 2008; Yates, Sherman, & Yousif, 2023; Zacks, 2020). Most of this work, unsurprisingly, has focused on events that can be easily studied in the laboratory: events with a simple linear structure, typically lasting seconds, minutes, or hours, at most. Yet our lives unfold not on the timescale of moments or minutes, but weeks and months and years. And any one 'moment' may be part of a representation at multiple scales (e.g., on a Friday afternoon, on the 9th of October, in the year 2020). Moreover, our lives are not one single narrative: There are overlapping plots, a revolving door of characters, and unexpected twists (see Conway, 1996). What does event representation look like at this scale — for events that share the 'texture' of ordinary experience?

1. Event representation at shorter vs. longer timescales

In everyday language as well as in cognitive science, an "event" may refer to many different things (see Yates et al., 2023). Some work emphasizes *event boundaries* (see, e.g., Radvansky & Copeland, 2006; Radvansky, Krawietz, & Tamplin, 2011; Radvansky, Tamplin, & Krawietz, 2010; see also Radvansky, 2012). Other work subtly manipulates context (e.g., a change from an orange to a purple background; DuBrow & Davachi, 2013; DuBrow & Davachi, 2016; Ezzyat & Davachi, 2014; Heusser, Ezzyat, Shiff, & Davachi, 2018). Yet other work relies on more naturalistic events, in the form of written stories (Copeland, Radvansky, & Goodwin, 2009; Doolen & Radvansky, 2021, 2022; Speer, Zacks, & Reynolds, 2007) or video media (e.g., Baldassano et al., 2017; Yates et al., 2022; Zacks, Speer, Swallow, & Maley, 2010). While these studies certainly reveal something about the nature of event processing, they do not fully reflect the richness of ordinary experience (but see Teigen, Böhm, Bruckmüller, Hegarty, & Luminet, 2017; Rouhani et al., 2023).

Current theories of event cognition point to two influential ideas about event representation. The first is the notion that event *boundaries* influence memory in a variety of ways. For instance, boundaries influence associative memory: The temporal order of items across events is more easily confused than temporal order within events, suggesting that temporal memory is organized around event structure to some extent

* Corresponding author.
E-mail address: sryousif@sas.upenn.edu (S.R. Yousif).

https://doi.org/10.1016/j.cognition.2024.105833
Received 26 September 2023; Received in revised form 15 May 2024; Accepted 22 May 2024
Available online 3 June 2024

Cognition 249 (2024) 105833

Contents lists available at ScienceDirect

Cognition

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

Event representation at the scale of ordinary experience

Sami R. Yousif[a,*], Sarah Hye-yeon Lee[b], Brynn E. Sherman[a], Anna Papafragou[b]

[a] University of Pennsylvania, Department of Psychology, United States of America
[b] University of Pennsylvania, Department of Linguistics, United States of America

ARTICLE INFO

Keywords:
Event representation
Episodic memory
Temporal memory

ABSTRACT

Weeks are divided into weekdays and weekends; years into semesters and seasons; lives into stages like childhood, adulthood, and adolescence. How does the structure of experience shape memory? Though much work has examined event representation in human cognition, little work has explored event representation at the scale of ordinary experience. Here, we use shared experiences — in the form of popular television shows — to explore how memories are shaped by event structure at a large scale. We find that memories for events in these shows exhibit several hallmarks of event cognition. Namely, we find that memories are organized with respect to their event structure (boundaries), and that beginnings and endings are better remembered at multiple levels of the event hierarchy simultaneously. These patterns seem to be partially, but not fully, explained by the perceived story-relevance of events. Lastly, using a longitudinal design, we also show how event representations evolve over periods of several months. These results offer an understanding of event cognition at the scale of ordinary human lives.

Experiences are complex and varied, but our memories of them are structured in the form of discrete events. We remember our high school years as distinct from our college years and those years as distinct from everything after. We remember our summer months as distinct from the grind of the semester. We remember our weekdays as distinct from our weekends.

A significant body of work in psychology and cognitive science has been interested in the nature of event representation — how it is that we carve *continuous* experience up into *discrete* 'units' (for reviews, see Radvansky & Zacks, 2011; Shipley & Zacks, 2008; Yates, Sherman, & Yousif, 2023; Zacks, 2020). Most of this work, unsurprisingly, has focused on events that can be easily studied in the laboratory: events with a simple linear structure, typically lasting seconds, minutes, or hours, at most. Yet our lives unfold not on the timescale of moments or minutes, but weeks and months and years. And any one 'moment' may be part of a representation at multiple scales (e.g., on a Friday afternoon, on the 9th of October, in the year 2020). Moreover, our lives are not one single narrative: There are overlapping plots, a revolving door of characters, and unexpected twists (see Conway, 1996). What does event representation look like at this scale — for events that share the 'texture' of ordinary experience?

1. Event representation at shorter vs. longer timescales

In everyday language as well as in cognitive science, an "event" may refer to many different things (see Yates et al., 2023). Some work emphasizes *event boundaries* (see, e.g., Radvansky & Copeland, 2006; Radvansky, Krawietz, & Tamplin, 2011; Radvansky, Tamplin, & Krawietz, 2010; see also Radvansky, 2012). Other work subtly manipulates context (e.g., a change from an orange to a purple background; DuBrow & Davachi, 2013; DuBrow & Davachi, 2016; Ezzyat & Davachi, 2014; Heusser, Ezzyat, Shiff, & Davachi, 2018). Yet other work relies on more naturalistic events, in the form of written stories (Copeland, Radvansky, & Goodwin, 2009; Doolen & Radvansky, 2021, 2022; Speer, Zacks, & Reynolds, 2007) or video media (e.g., Baldassano et al., 2017; Yates et al., 2022; Zacks, Speer, Swallow, & Maley, 2010). While these studies certainly reveal something about the nature of event processing, they do not fully reflect the richness of ordinary experience (but see Teigen, Böhm, Bruckmüller, Hegarty, & Luminet, 2017; Rouhani et al., 2023).

Current theories of event cognition point to two influential ideas about event representation. The first is the notion that event *boundaries* influence memory in a variety of ways. For instance, boundaries influence associative memory: The temporal order of items across events is more easily confused than temporal order within events, suggesting that temporal memory is organized around event structure to some extent

* Corresponding author.
E-mail address: sryousif@sas.upenn.edu (S.R. Yousif).

https://doi.org/10.1016/j.cognition.2024.105833
Received 26 September 2023; Received in revised form 15 May 2024; Accepted 22 May 2024
Available online 3 June 2024

GAME OF THRONES™

"The North secedes from the Seven Kingdoms and proclaims Robb as king."

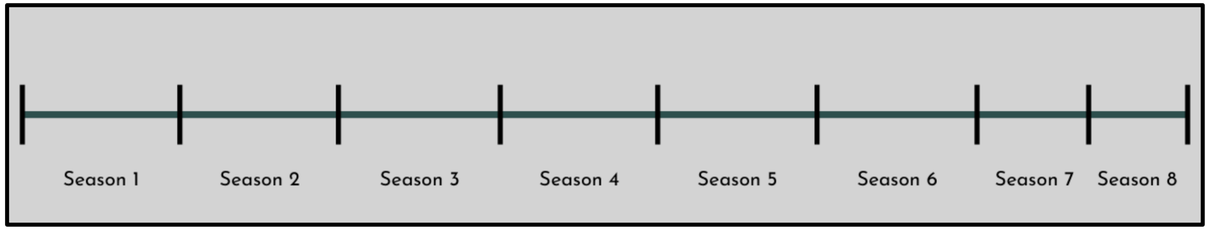

Cognition 249 (2024) 105833

Contents lists available at ScienceDirect

Cognition

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

Event representation at the scale of ordinary experience

Sami R. Yousif[a,*], Sarah Hye-yeon Lee[b], Brynn E. Sherman[a], Anna Papafragou[b]

[a] University of Pennsylvania, Department of Psychology, United States of America
[b] University of Pennsylvania, Department of Linguistics, United States of America

ARTICLE INFO

Keywords:
Event representation
Episodic memory
Temporal memory

ABSTRACT

Weeks are divided into weekdays and weekends; years into semesters and seasons; lives into stages like childhood, adulthood, and adolescence. How does the structure of experience shape memory? Though much work has examined event representation in human cognition, little work has explored event representation at the scale of ordinary experience. Here, we use shared experiences — in the form of popular television shows — to explore how memories are shaped by event structure at a large scale. We find that memories for events in these shows exhibit several hallmarks of event cognition. Namely, we find that memories are organized with respect to their event structure (boundaries), and that beginnings and endings are better remembered at multiple levels of the event hierarchy simultaneously. These patterns seem to be partially, but not fully, explained by the perceived story-relevance of events. Lastly, using a longitudinal design, we also show how event representations evolve over periods of several months. These results offer an understanding of event cognition at the scale of ordinary human lives.

Experiences are complex and varied, but our memories of them are structured in the form of discrete events. We remember our high school years as distinct from our college years and those years as distinct from everything after. We remember our summer months as distinct from the grind of the semester. We remember our weekdays as distinct from our weekends.

A significant body of work in psychology and cognitive science has been interested in the nature of event representation — how it is that we carve *continuous* experience up into *discrete* 'units' (for reviews, see Radvansky & Zacks, 2011; Shipley & Zacks, 2008; Yates, Sherman, & Yousif, 2023; Zacks, 2020). Most of this work, unsurprisingly, has focused on events that can be easily studied in the laboratory: events with a simple linear structure, typically lasting seconds, minutes, or hours, at most. Yet our lives unfold not on the timescale of moments or minutes, but weeks and months and years. And any one 'moment' may be part of a representation at multiple scales (e.g., on a Friday afternoon, on the 9th of October, in the year 2020). Moreover, our lives are not one single narrative: There are overlapping plots, a revolving door of characters, and unexpected twists (see Conway, 1996). What does event representation look like at this scale — for events that share the 'texture' of ordinary experience?

1. Event representation at shorter vs. longer timescales

In everyday language as well as in cognitive science, an "event" may refer to many different things (see Yates et al., 2023). Some work emphasizes *event boundaries* (see, e.g., Radvansky & Copeland, 2006; Radvansky, Krawietz, & Tamplin, 2011; Radvansky, Tamplin, & Krawietz, 2010; see also Radvansky, 2012). Other work subtly manipulates context (e.g., a change from an orange to a purple background; DuBrow & Davachi, 2013; DuBrow & Davachi, 2016; Ezzyat & Davachi, 2014; Heusser, Ezzyat, Shiff, & Davachi, 2018). Yet other work relies on more naturalistic events, in the form of written stories (Copeland, Radvansky, & Goodwin, 2009; Doolen & Radvansky, 2021, 2022; Speer, Zacks, & Reynolds, 2007) or video media (e.g., Baldassano et al., 2017; Yates et al., 2022; Zacks, Speer, Swallow, & Maley, 2010). While these studies certainly reveal something about the nature of event processing, they do not fully reflect the richness of ordinary experience (but see Teigen, Böhm, Bruckmüller, Hegarty, & Luminet, 2017; Rouhani et al., 2023).

Current theories of event cognition point to two influential ideas about event representation. The first is the notion that event *boundaries* influence memory in a variety of ways. For instance, boundaries influence associative memory: The temporal order of items across events is more easily confused than temporal order within events, suggesting that temporal memory is organized around event structure to some extent

* Corresponding author.
E-mail address: sryousif@sas.upenn.edu (S.R. Yousif).

https://doi.org/10.1016/j.cognition.2024.105833
Received 26 September 2023; Received in revised form 15 May 2024; Accepted 22 May 2024
Available online 3 June 2024

Game of Thrones

!!

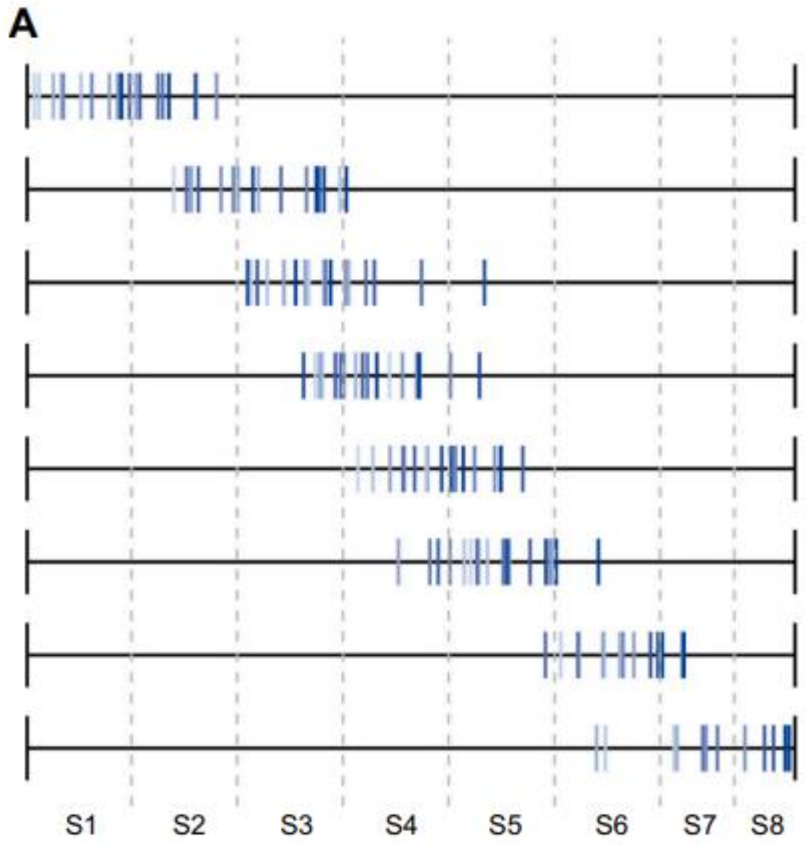

Robust temporal expansion

OPEN ACCESS Freely available online PLoS one

The Effect of Predictability on Subjective Duration

Vani Pariyadath[1], David Eagleman[1,2]*

1 Department of Neuroscience, Baylor College of Medicine, Houston, Texas, United States of America, 2 Department of Psychiatry, Baylor College of Medicine, Houston, Texas, United States of America

Events can sometimes appear longer or shorter in duration than other events of equal length. For example, in a repeated presentation of auditory or visual stimuli, an unexpected object of equivalent duration appears to last longer. Illusions of duration distortion beg an important question of time representation: when durations dilate or contract, does time in general slow down or speed up during that moment? In other words, what entailments do duration distortions have with respect to other timing judgments? We here show that when a sound or visual flicker is presented in conjunction with an unexpected visual stimulus, neither the pitch of the sound nor the frequency of the flicker is affected by the apparent duration dilation. This demonstrates that subjective time in general is not slowed; instead, duration judgments can be manipulated with no concurrent impact on other temporal judgments. Like spatial vision, time perception appears to be underpinned by a collaboration of separate neural mechanisms that usually work in concert but are separable. We further show that the duration dilation of an unexpected stimulus is not enhanced by increasing its saliency, suggesting that the effect is more closely related to prediction violation than enhanced attention. Finally, duration distortions induced by violations of progressive number sequences implicate the involvement of high-level predictability, suggesting the involvement of areas higher than primary visual cortex. We suggest that duration distortions can be understood in terms of repetition suppression, in which neural responses to repeated stimuli are diminished.

Citation: Pariyadath V, Eagleman D (2007) The Effect of Predictability on Subjective Duration. PLoS ONE 2(11): e1264. doi:10.1371/journal.pone.0001264

INTRODUCTION

Time is commonly thought to fluctuate in its subjective rate of passage. For example, upon first glance, the second hand of a clock sometimes seems to be frozen in position momentarily before it continues to tick at a normal pace [1,2]. Perceived duration can be warped by saccades [3,4], flicker [5], and life-threatening events, which are sometimes anecdotally reported to unfold in slow motion [6,7]. The neural basis of such distortions remains unknown.

To gain traction on time representation and its plasticity, we turn to a duration distortion that is easily reproduced in the laboratory. Specifically, the first stimulus in a train of repeated presentations is often perceived to have a longer duration than successive stimuli. Participants report duration dilations of as much as 50% in trains of visual stimuli [8,9], and as much as 15% in trains of auditory stimuli [10]. The above studies proposed that the illusion is a consequence of increased arousal at the first appearance of the stimulus.

Similarly, when an oddball stimulus appears midstream in a repeated presentation of stimuli (auditory or visual), the judged duration of the oddball is overestimated by up to 50% [11,12]. Tse and his co-authors (2004) proposed an attentional explanation–specifically, that the duration dilation results from an increase in information processed at the time of the oddball due to the deployment of attentional resources. Tse et al (2004) refer to the duration dilation as 'time's subjective expansion'.

But what does it mean to say that subjective time expands? We here set out to distinguish two hypotheses. In the first, perception works like a movie camera: when one aspect of the scene slows down, everything is slowed down. Thus, if a police car launching off a ramp were filmed using slow-motion photography, it would not only have a longer duration in the air, but also its sirens would blare in a lower pitch, and its lights would blink at a lower temporal frequency. In this case, duration, sound pitch and visual flicker all change hand-in-hand. The second hypothesis, in contrast, supposes that different temporal judgments are generated by different neural mechanisms–and while they often align, they are not required to. Thus, the police car may be judged to have a longer duration in the air, even while the frequencies of its sounds and flickering lights remain unchanged. In this paper, we distinguish these two hypotheses by testing the specific entailments of duration distortions, and in this way are able to directly address the notion of "time's" subjective expansion.

MATERIALS AND METHODS

Participants sat 59 cm from a computer monitor and fixated at the center of the screen, and made responses using the keyboard. All participants had normal or corrected-to-normal vision, and were consented according to the procedures of the Institutional Review Board at Baylor College of Medicine.

RESULTS

Experiment 1

We began by quantifying the midstream oddball illusion [11]. Six subjects ran 84 trials in which they watched 9 repeated presentations of a photograph with an oddball photograph randomly embedded between the 5th and 8th presentation. Photographs subtended $3.1 \times 3.1°$ of visual angle and were repeatedly presented at fixation for 500 ms with ISIs of 300 ms. The duration of the oddball varied between 300–700 ms (Figure 1a). After each trial,

Academic Editor: David Burr, Istituto di Neurofisiologia, Italy

Received October 17, 2007; **Accepted** November 1, 2007; **Published** November 28, 2007

Funding: RO1 NIH NS053960 (DME)

Competing Interests: The authors have declared that no competing interests exist.

* **To whom correspondence should be addressed.** E-mail: eagleman@bcm.edu

Psychological Research (2006) 70: 77–87
DOI 10.1007/s00426-004-0195-4

ORIGINAL ARTICLE

Rolf Ulrich · Judith Nitschke · Thomas Rammsayer

Perceived duration of expected and unexpected stimuli

Received: 26 April 2004 / Accepted: 3 September 2004 / Published online: 18 December 2004

Abstract Three experiments assessed whether perceived stimulus duration depends on whether participants process an expected or an unexpected visual stimulus. Participants compared the duration of a constant standard stimulus with a variable comparison stimulus. Changes in expectancy were induced by presenting one type of comparison more frequently than another type. Experiment 1 used standard durations of 100 and 400 ms, and Experiments 2 and 3 durations of 400 and 800 ms. Stimulus frequency did not affect perceived duration in Experiment 1. In Experiments 2 and 3, however, frequent comparisons were perceived as shorter than infrequent ones, and discrimination performance was better for infrequent comparisons. Overall, this study supports the notion that infrequent stimuli increase the speed of an internal pacemaker.

Introduction

Temporal information processing in humans and animals is often explained in terms of a hypothetical internal pacemaker counter mechanism (e.g., Allan, Kristofferson, & Wiens, 1971; Church & Gibbon, 1982; Creelman, 1962; Gibbon, 1977; Killeen & Fetterman, 1988; Rammsayer & Ulrich, 2001; Treisman, 1963; Treisman, Faulkner, Naish, & Brogan, 1990; Treisman, Faulkner, & Naish, 1992; Wearden, 1991). For example, the scalar timing theory (Gibbon, 1991) assumes among other things that an internal neural pacemaker generates successive pulses according to a Poisson process and that a further device accumulates these emitted pulses. According to this theory, then, the number of pulses accumulated during a given time interval is the internal representation of this interval. Although the scalar timing theory was formerly developed as a model for animal timing (e.g., Gibbon, 1977), it has also been successfully applied to human timing (e.g., Droit-Volet, 2002; Fortin, 2003; Malapani & Fairhurst, 2002; McCormack, Brown, Maylor, Richardson, & Darby, 2002; Penney, Gibbon, & Meck, 2000; Wearden, 1991, 1995; for a review see Grondin, 2001a).

More recently, several studies on human timing have focused on the issue of how experimental conditions (such as stimulus modality, e.g., Wearden, Edwards, Fakhri, & Percival, 1998), neurochemical substances (e.g., Rammsayer, 1999), and cognitive factors (e.g., Brown & Boltz, 2002) influence the speed of the internal pacemaker and, thus, perception of time. A number of studies provided convincing evidence that cognitive factors such as the amount of attention devoted to time could greatly alter the perceived duration of an interval (e.g., Casini & Macar, 1999; Chen & O'Neill, 2001; Enns, Brehaut, & Shore, 1999; Hemmes, Brown, & Kladopoulos, 2004; Mattes & Ulrich, 1998; Zakay & Block, 1997; for a review see Lejeune, 1998). The most typical findings were obtained in dual-task conditions, where participants have to divide attention between time perception and a nontemporal task. The major result is that allocation of attention to the nontemporal task shortens the perceived duration of a simultaneously occurring event (e.g., Brown, 1995, 1997). According to Thomas and Weaver's (1975) attenuation hypothesis, attentional resources have to be divided between the internal pacemaker counter timing mechanism and a nontemporal stimulus processor. It is assumed that when attention is directed to the stimulus processor, a certain number of pulses remain unregistered. Consequently, the less attention is directed to time, the shorter the time experienced will be.

An alternative strategy for investigating attentional mechanisms represents the cuing paradigm (Posner,

R. Ulrich (✉) · J. Nitschke
Psychologisches Institut, Universität Tübingen,
Friedrichstrasse 21, 72072 Tübingen, Germany
E-mail: ulrich@uni-tuebingen.de
Tel.: +49-7071-2972410
Fax: +49-7071-292410

T. Rammsayer
Universität Göttingen, Göttingen, Germany

OPEN ACCESS Freely available online PLoS one

The Effect of Predictability on Subjective Duration

Vani Pariyadath[1], David Eagleman[1,2]*

1 Department of Neuroscience, Baylor College of Medicine, Houston, Texas, United States of America, 2 Department of Psychiatry, Baylor College of Medicine, Houston, Texas, United States of America

Events can sometimes appear longer or shorter in duration than other events of equal length. For example, in a repeated presentation of auditory or visual stimuli, an unexpected object of equivalent duration appears to last longer. Illusions of duration distortion beg an important question of time representation: when durations dilate or contract, does time in general slow down or speed up during that moment? In other words, what entailments do duration distortions have with respect to other timing judgments? We here show that when a sound or visual flicker is presented in conjunction with an unexpected visual stimulus, neither the pitch of the sound nor the frequency of the flicker is affected by the apparent duration dilation. This demonstrates that subjective time in general is not slowed; instead, duration judgments can be manipulated with no concurrent impact on other temporal judgments. Like spatial vision, time perception appears to be underpinned by a collaboration of separate neural mechanisms that usually work in concert but are separable. We further show that the duration dilation of an unexpected stimulus is not enhanced by increasing its saliency, suggesting that the effect is more closely related to prediction violation than enhanced attention. Finally, duration distortions induced by violations of progressive number sequences implicate the involvement of high-level predictability, suggesting the involvement of areas higher than primary visual cortex. We suggest that duration distortions can be understood in terms of repetition suppression, in which neural responses to repeated stimuli are diminished.

Citation: Pariyadath V, Eagleman D (2007) The Effect of Predictability on Subjective Duration. PLoS ONE 2(11): e1264. doi:10.1371/journal.pone.0001264

INTRODUCTION

Time is commonly thought to fluctuate in its subjective rate of passage. For example, upon first glance, the second hand of a clock sometimes seems to be frozen in position momentarily before it continues to tick at a normal pace [1,2]. Perceived duration can be warped by saccades [3,4], flicker [5], and life-threatening events, which are sometimes anecdotally reported to unfold in slow motion [6,7]. The neural basis of such distortions remains unknown.

To gain traction on time representation and its plasticity, we turn to a duration distortion that is easily reproduced in the laboratory. Specifically, the first stimulus in a train of repeated presentations is often perceived to have a longer duration than successive stimuli. Participants report duration dilations of as much as 50% in trains of visual stimuli [8,9], and as much as 15% in trains of auditory stimuli [10]. The above studies proposed that the illusion is a consequence of increased arousal at the first appearance of the stimulus.

Similarly, when an oddball stimulus appears midstream in a repeated presentation of stimuli (auditory or visual), the judged duration of the oddball is overestimated by up to 50% [11,12]. Tse and his co-authors (2004) proposed an attentional explanation–specifically, that the duration dilation results from an increase in information processed at the time of the oddball due to the deployment of attentional resources. Tse et al (2004) refer to the duration dilation as 'time's subjective expansion'.

But what does it mean to say that subjective time expands? We here set out to distinguish two hypotheses. In the first, perception works like a movie camera: when one aspect of the scene slows down, everything is slowed down. Thus, if a police car launching off a ramp were filmed using slow-motion photography, it would not only have a longer duration in the air, but also its sirens would blare in a lower pitch, and its lights would blink at a lower temporal frequency. In this case, duration, sound pitch and visual flicker all change hand-in-hand. The second hypothesis, in contrast, supposes that different temporal judgments are generated by different neural mechanisms–and while they often align, they are not required to. Thus, the police car may be judged to have a longer duration in the air, even while the frequencies of its sounds and flickering lights remain unchanged. In this paper, we distinguish these two hypotheses by testing the specific entailments of duration distortions, and in this way are able to directly address the notion of "time's" subjective expansion.

MATERIALS AND METHODS

Participants sat 59 cm from a computer monitor and fixated at the center of the screen, and made responses using the keyboard. All participants had normal or corrected-to-normal vision, and were consented according to the procedures of the Institutional Review Board at Baylor College of Medicine.

RESULTS

Experiment 1

We began by quantifying the midstream oddball illusion [11]. Six subjects ran 84 trials in which they watched 9 repeated presentations of a photograph with an oddball photograph randomly embedded between the 5th and 8th presentation. Photographs subtended 3.1×3.1° of visual angle and were repeatedly presented at fixation for 500 ms with ISIs of 300 ms. The duration of the oddball varied between 300–700 ms (Figure 1a). After each trial,

Academic Editor: David Burr, Istituto di Neurofisiologia, Italy

Received October 17, 2007; Accepted November 1, 2007; Published November 28, 2007

Funding: RO1 NIH NS053960 (DME)

Competing Interests: The authors have declared that no competing interests exist.

* To whom correspondence should be addressed. E-mail: eagleman@bcm.edu

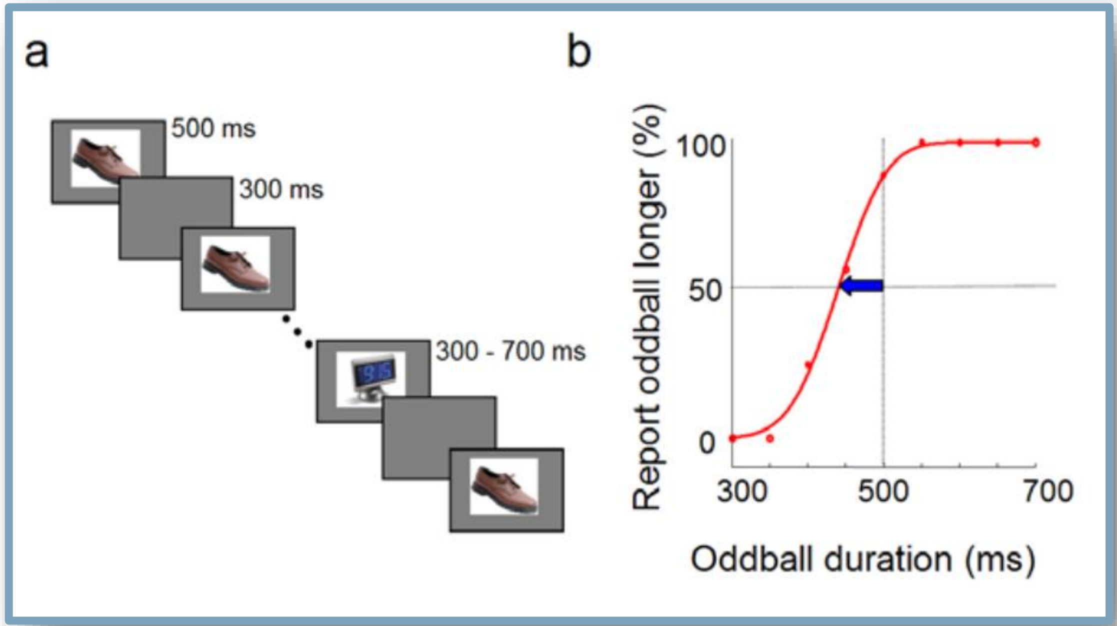

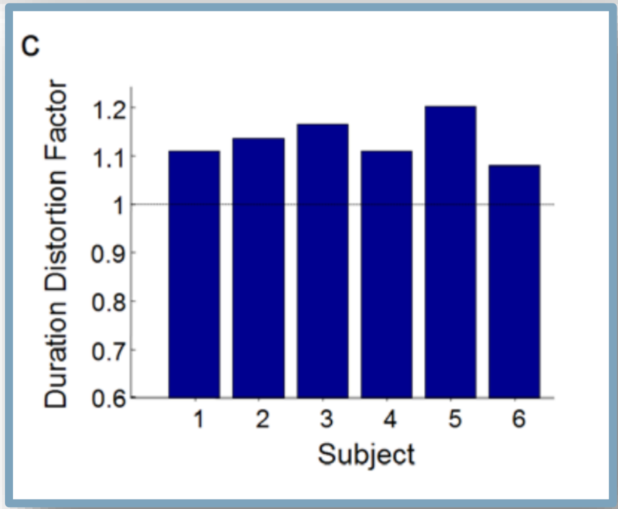

PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES
SOCIAL SCIENCES

The perception of silence

Rui Zhe Goh[a,b,1], Ian B. Phillips[a,b,1,2], and Chaz Firestone[a,b,1,2]

Edited by Steven Luck, University of California, Davis, CA; received February 8, 2023; accepted May 10, 2023 by Editorial Board Member Michael S. Gazzaniga

Auditory perception is traditionally conceived as the perception of sounds—a friend's voice, a clap of thunder, a minor chord. However, daily life also seems to present us with experiences characterized by the absence of sound—a moment of silence, a gap between thunderclaps, the hush after a musical performance. In these cases, do we positively *hear* silence? Or do we just *fail to hear*, and merely judge or infer that it is silent? This longstanding question remains controversial in both the philosophy and science of perception, with prominent theories holding that sounds are the only objects of auditory experience and thus that our encounter with silence is cognitive, not perceptual. However, this debate has largely remained theoretical, without a key empirical test. Here, we introduce an empirical approach to this theoretical dispute, presenting experimental evidence that silence can be genuinely perceived (not just cognitively inferred). We ask whether silences can "substitute" for sounds in event-based auditory illusions—empirical signatures of auditory event representation in which auditory events distort perceived duration. Seven experiments introduce three "silence illusions"—the one-silence-is-more illusion, silence-based warping, and the oddball-silence illusion—each adapted from a prominent perceptual illusion previously thought to arise only from sounds. Subjects were immersed in ambient noise interrupted by silences structurally identical to the sounds in the original illusions. In all cases, silences elicited temporal distortions perfectly analogous to the illusions produced by sounds. Our results suggest that silence is truly heard, not merely inferred, introducing a general approach for studying the perception of absence.

absence perception | silence | event representation | temporal illusions

Significance

Do we only hear sounds? Or can we also hear silence? These questions are the subject of a centuries-old philosophical debate between two camps: the perceptual view (we literally hear silence), and the cognitive view (we only judge or infer silence). Here, we take an empirical approach to resolve this theoretical controversy. We show that silences can "substitute" for sounds in event-based auditory illusions. Seven experiments introduce three "silence illusions," adapted from perceptual illusions previously thought to arise only with sounds. In all cases, silences elicited temporal distortions perfectly analogous to their sound-based counterparts, suggesting that auditory processing treats moments of silence the way it treats sounds. Silence is truly perceived, not merely inferred.

What do we hear? The canonical answer is that auditory perception is the perception of *sounds* and their properties—the pitch of a friend's voice, the loudness of a thunderclap, the timbre of a minor chord. This traditional view has considerable pedigree, with influential historical sources holding that sounds are the sole objects of auditory perception (1, cf.2). It is also the answer favored in contemporary scholarship: Prominent scientific accounts conceive the fundamental units of auditory perception as sounds (or auditory streams comprised of sounds; ref. 3 and 4), and many philosophical theories agree, holding that "all auditory perception involves the perception of sound" (5) and that "if anything at all is heard, what is heard is necessarily a sound" (6) (see also refs. 7 and 8). The pervasiveness of this canonical view about the contents of auditory perception might seem unsurprising—what else might we hear, if not sound?

However, there has long been a stubborn and intuitive counterexample: experiences of *silence*, which are characterized by the absence of sound. Silence confronts us throughout our daily lives—consider an awkward pause in a conversation, a suspenseful gap between thunderclaps, or the hush at the end of a musical performance. What is the nature of these experiences?

Silence: Heard or Inferred?

One possibility is that experiences of silence are simply cases in which we *fail to hear*, and then use our faculties of reasoning and judgment to *infer* that it is silent. This interpretation is offered by those who defend the traditional sound-only view of audition, holding that an experience of silence is merely the "cognitive accompaniment of an absence of experience" and "is itself no form of hearing" (9). This cognitive view may be motivated by a deeper assumption about perception, namely that we can genuinely perceive only what is present in the world, not what is absent (9, 10). After all, one might think, absences are nonentities—they do not exist—and so can hardly impinge on our sensory apparatus.

However, an alternative possibility which arguably does more justice to our phenomenology is that we literally perceive silences. This interpretation has recently received

Author affiliations: [a]Department of Psychological and Brain Sciences, Johns Hopkins University, Baltimore, MD 21218; and [b]Department of Philosophy, Johns Hopkins University, Baltimore, MD 21218

Author contributions: R.Z.G., I.B.P., and C.F. designed research; R.Z.G. performed research; R.Z.G. analyzed data; and R.Z.G., I.B.P., and C.F. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. S.L. is a guest editor invited by the Editorial Board.

[1]To whom correspondence may be addressed. Email: rgoh1@jhu.edu, ianbphillips@jhu.edu, or chaz@jhu.edu.

[2]I.B.P. and C.F. contributed equally to this work.

Published July 10, 2023.

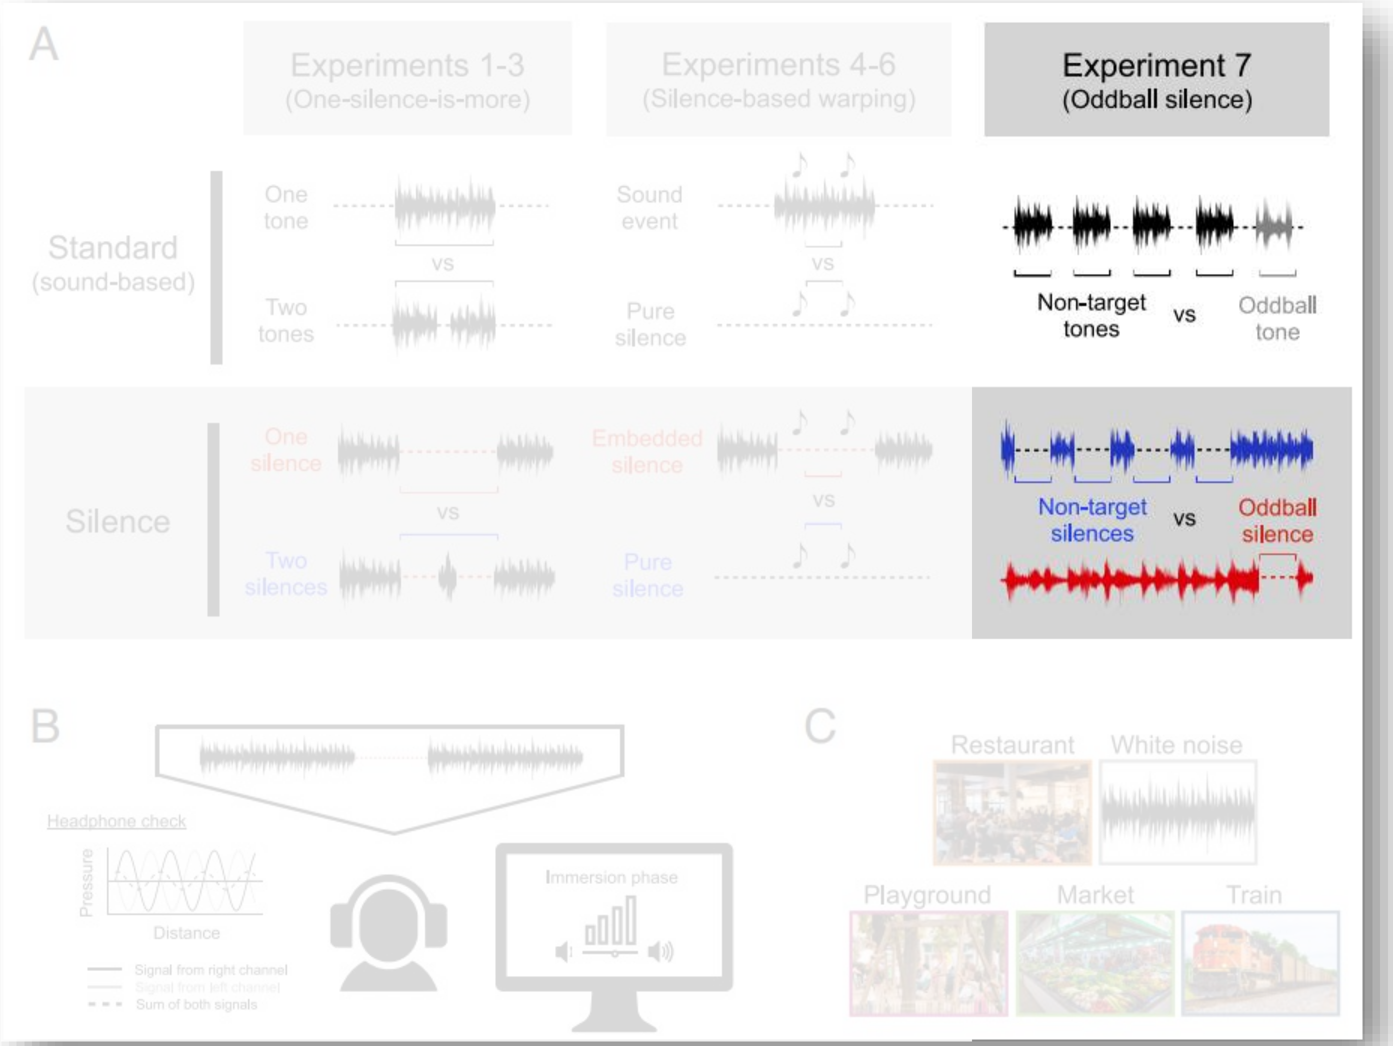

Memory & Cognition
1974, Vol. 2, No. 1A, 153-160

Memory and the experience of duration in retrospect*

RICHARD A. BLOCK†
University of Oregon, Eugene, Oregon 97403

Two experiments investigated the relationship between long-term memory for events occurring during an interval and the experience of duration of the interval in retrospect. In both, Ss attended to a sequence consisting of a standard, an experimental, and a second standard interval. Then unexpected comparative duration and memory judgments were requested. In Experiment I, either 30 or 60 unrelated words occurred during the 180-sec experimental interval. When more words had occurred, judgments of duration of the experimental interval, judgments of number of words presented, and number of words recognized all increased, but free recall of words was unaffected. In Experiment II, 80 categorized words occurred during the 160-sec experimental interval, with category instances in either blocked or random order. When words were blocked by category, judgments of duration of the experimental interval, free recall, and recognition all increased, but judgments of number of words were unaffected. Results were discussed in terms of Ornstein's (1969) "storage size" hypothesis.

Some of the first philosophers and psychologists who discussed the experience of duration realized that it must be intimately related to memory processes. Aristotle thought that "only those animals which perceive time remember, and the organ whereby they perceive time is also that whereby they remember [McKeon, 1941, pp. 607-608]." More recently, a number of psychologists have proposed a distinction between the experience of duration in passing and in retrospect, postulating separate mechanisms to account for the two types of temporal experience. As the term is typically used, "the experience of duration in passing" refers to the awareness of the apparent length of an interval between a past event and the psychological present. On the other hand, "the experience of duration in retrospect" typically refers to the awareness of the apparent length of an interval between two past events.

There is little agreement among theorists concerning the processes underlying the experience of duration in passing. However, contemporary theorists, along with Aristotle, generally agree that memory plays a central role in the experience of duration in retrospect (e.g., Frankenhaeuser, 1959; Ornstein, 1969; Michon, 1970). Since memory is assumed to be involved, the present paper explores the role of long-term memory in determining the retrospective experience of duration of relatively long intervals.

Although there has been much theoretical speculation about the role of memory in the experience of duration, only a few investigators have attempted to determine experimentally which aspects of memory might be involved. Frankenhaeuser (1959, Experiments 3 and 4) studied the relationship between memory for the number of events that occurred during an interval and retrospective estimates of duration of the interval. In these experiments, Ss read random digits until stopped. Then they estimated either the duration of the interval, the number of digits read, or both. There was a close correspondence between the two types of estimates over intervals ranging in duration from 4 to 53 sec. The interpretation of these results, however, is complicated by the fact that Ss were tested repeatedly and were aware that duration judgments would be requested at the end of some of the intervals. Under these circumstances, Ss might have adopted a special strategy for making judgments of duration. Thus, the degree of correspondence between the two types of judgments might have been somewhat artifactual and possibly not generalizable to situations in which Ss are not aware that duration judgments would be requested at the end of the interval.

Using a paradigm in which Ss were not aware that duration judgments would be requested, Ornstein (1969) conducted several experiments in which both memory for the contents of an interval and judgments of duration were measured. In one of them (Experiment VII), Ss learned either of two different paired-associate lists for 6 min, and they were asked to make a comparative magnitude judgment of duration of the learning task after either no delay or a 2-week delay. Duration judgments were closely paralleled by a memory

*This article is based on a PhD dissertation submitted to the University of Oregon. The author wishes to thank members of the dissertation committee–Douglas L. Hintzman, Michael I. Posner, Wayne A. Wickelgren, and Don S. Levi–for criticism and advice. Thanks are also due Jean F. Block, David R. Gostnell, and A. John Wisdom for comments on an earlier version of the manuscript. The research was supported by Predoctoral Fellowship 2-F01-MH48356-02 from the National Institute of Mental Health, U.S. Public Health Service, and by a grant to Douglas L. Hintzman from the Office of Education, U.S. Department of Health, Education, and Welfare. Contractors undertaking such projects under government sponsorship are encouraged to express freely their professional judgment in the conduct of the project. Points of view or opinions stated do not, therefore, necessarily represent official Office of Education position or policy.

†Requests for reprints should be sent to Richard A. Block, Department of Psychology, State University of New York, Plattsburgh, New York 12901.

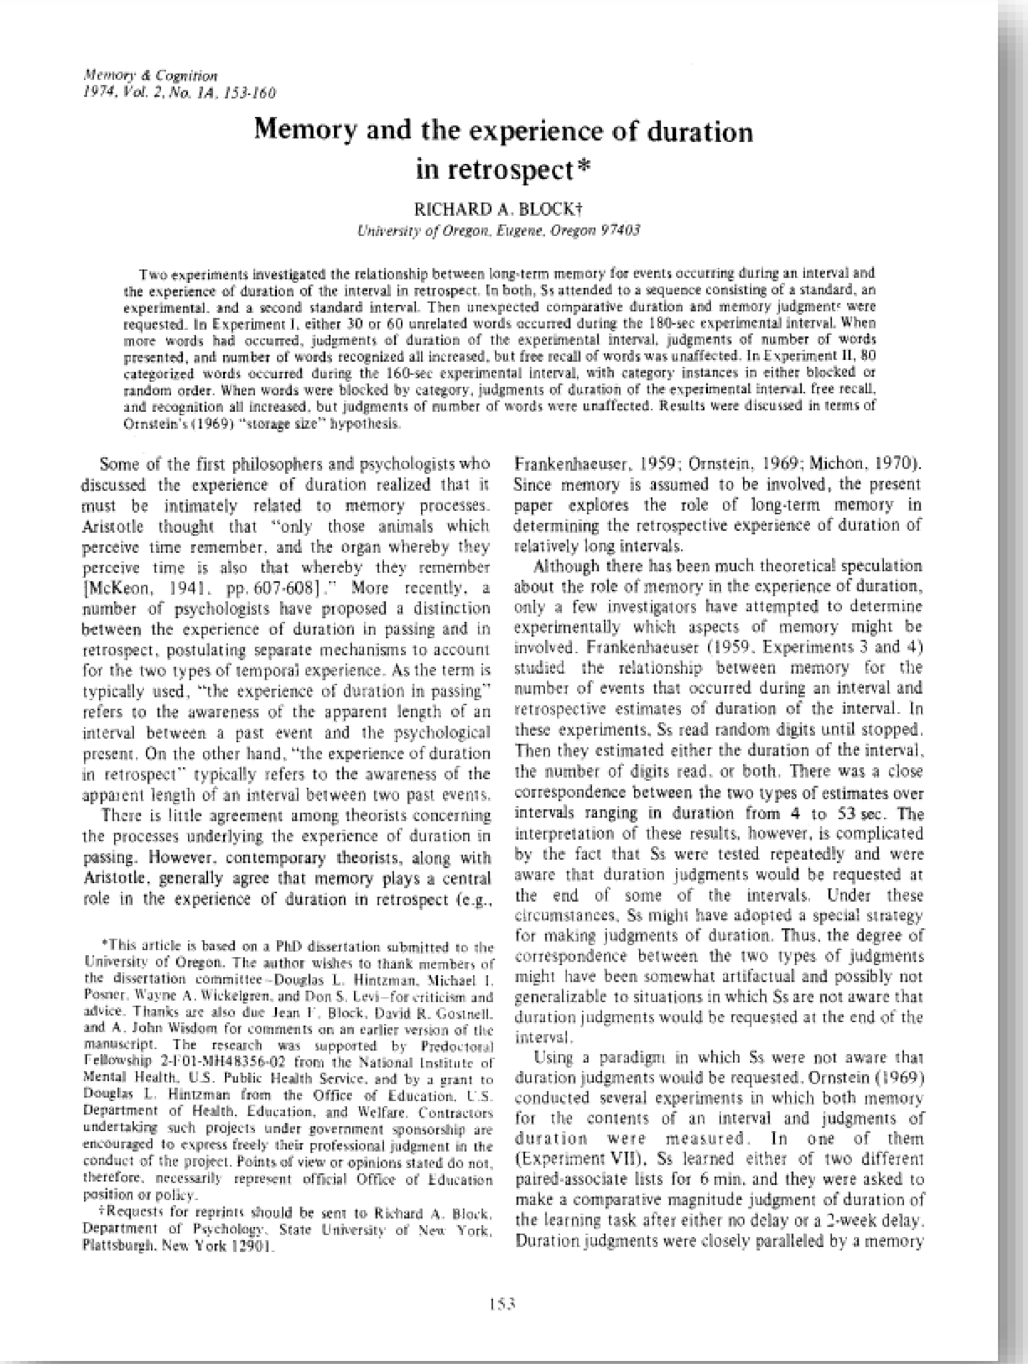

Memory & Cognition
1974, Vol. 2, No. 1A, 153-160

Memory and the experience of duration in retrospect*

RICHARD A. BLOCK†
University of Oregon, Eugene, Oregon 97403

Two experiments investigated the relationship between long-term memory for events occurring during an interval and the experience of duration of the interval in retrospect. In both, Ss attended to a sequence consisting of a standard, an experimental, and a second standard interval. Then unexpected comparative duration and memory judgments were requested. In Experiment I, either 30 or 60 unrelated words occurred during the 180-sec experimental interval. When more words had occurred, judgments of duration of the experimental interval, judgments of number of words presented, and number of words recognized all increased, but free recall of words was unaffected. In Experiment II, 80 categorized words occurred during the 160-sec experimental interval, with category instances in either blocked or random order. When words were blocked by category, judgments of duration of the experimental interval, free recall, and recognition all increased, but judgments of number of words were unaffected. Results were discussed in terms of Ornstein's (1969) "storage size" hypothesis.

Some of the first philosophers and psychologists who discussed the experience of duration realized that it must be intimately related to memory processes. Aristotle thought that "only those animals which perceive time remember, and the organ whereby they perceive time is also that whereby they remember [McKeon, 1941, pp. 607-608]." More recently, a number of psychologists have proposed a distinction between the experience of duration in passing and in retrospect, postulating separate mechanisms to account for the two types of temporal experience. As the term is typically used, "the experience of duration in passing" refers to the awareness of the apparent length of an interval between a past event and the psychological present. On the other hand, "the experience of duration in retrospect" typically refers to the awareness of the apparent length of an interval between two past events.

There is little agreement among theorists concerning the processes underlying the experience of duration in passing. However, contemporary theorists, along with Aristotle, generally agree that memory plays a central role in the experience of duration in retrospect (e.g., Frankenhaeuser, 1959; Ornstein, 1969; Michon, 1970). Since memory is assumed to be involved, the present paper explores the role of long-term memory in determining the retrospective experience of duration of relatively long intervals.

Although there has been much theoretical speculation about the role of memory in the experience of duration, only a few investigators have attempted to determine experimentally which aspects of memory might be involved. Frankenhaeuser (1959, Experiments 3 and 4) studied the relationship between memory for the number of events that occurred during an interval and retrospective estimates of duration of the interval. In these experiments, Ss read random digits until stopped. Then they estimated either the duration of the interval, the number of digits read, or both. There was a close correspondence between the two types of estimates over intervals ranging in duration from 4 to 53 sec. The interpretation of these results, however, is complicated by the fact that Ss were tested repeatedly and were aware that duration judgments would be requested at the end of some of the intervals. Under these circumstances, Ss might have adopted a special strategy for making judgments of duration. Thus, the degree of correspondence between the two types of judgments might have been somewhat artifactual and possibly not generalizable to situations in which Ss are not aware that duration judgments would be requested at the end of the interval.

Using a paradigm in which Ss were not aware that duration judgments would be requested, Ornstein (1969) conducted several experiments in which both memory for the contents of an interval and judgments of duration were measured. In one of them (Experiment VII), Ss learned either of two different paired-associate lists for 6 min, and they were asked to make a comparative magnitude judgment of duration of the learning task after either no delay or a 2-week delay. Duration judgments were closely paralleled by a memory

*This article is based on a PhD dissertation submitted to the University of Oregon. The author wishes to thank members of the dissertation committee–Douglas L. Hintzman, Michael I. Posner, Wayne A. Wickelgren, and Don S. Levi–for criticism and advice. Thanks are also due Jean F. Block, David R. Gostnell, and A. John Wisdom for comments on an earlier version of the manuscript. The research was supported by Predoctoral Fellowship 2-F01-MH48356-02 from the National Institute of Mental Health, U.S. Public Health Service, and by a grant to Douglas L. Hintzman from the Office of Education, U.S. Department of Health, Education, and Welfare. Contractors undertaking such projects under government sponsorship are encouraged to express freely their professional judgment in the conduct of the project. Points of view or opinions stated do not, therefore, necessarily represent official Office of Education position or policy.

†Requests for reprints should be sent to Richard A. Block, Department of Psychology, State University of New York, Plattsburgh, New York 12901.

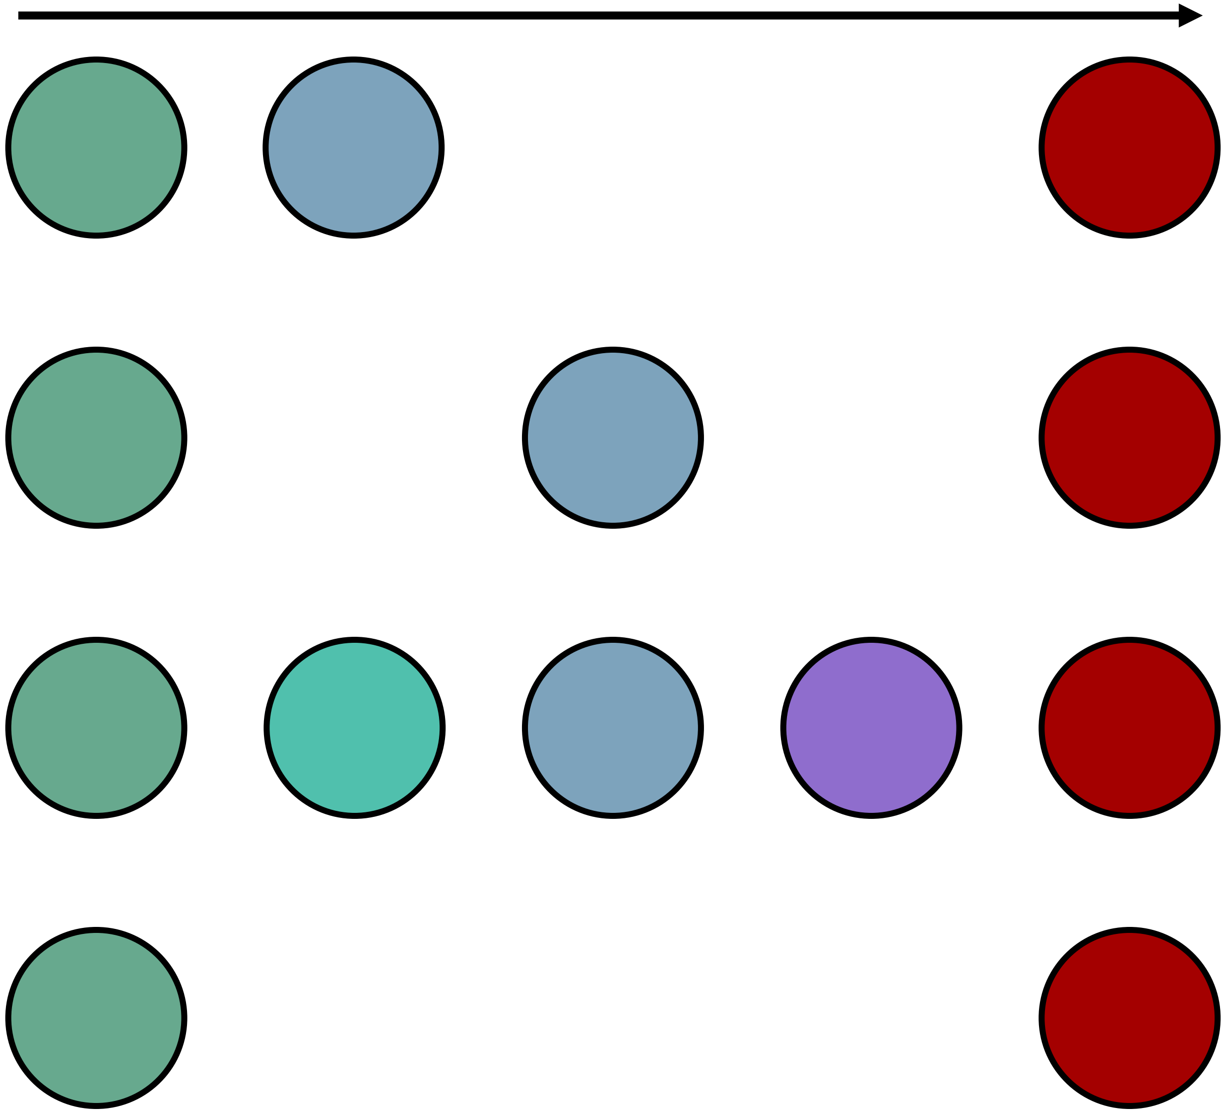

"storage size" hypothesis !!

An illusion of time caused by repeated experience

Brynn E. Sherman[1] & Sami R. Yousif[2]

[1] University of Pennsylvania

[2] University of North Carolina, Chapel Hill

Running Head : An illusion of time
Word Count : 2022 (Excluding Methods and Results)
Data : https://osf.io/nx5t7/?view_only=358f4e72d98b41cd875c19be0c97c090
Version : Accepted

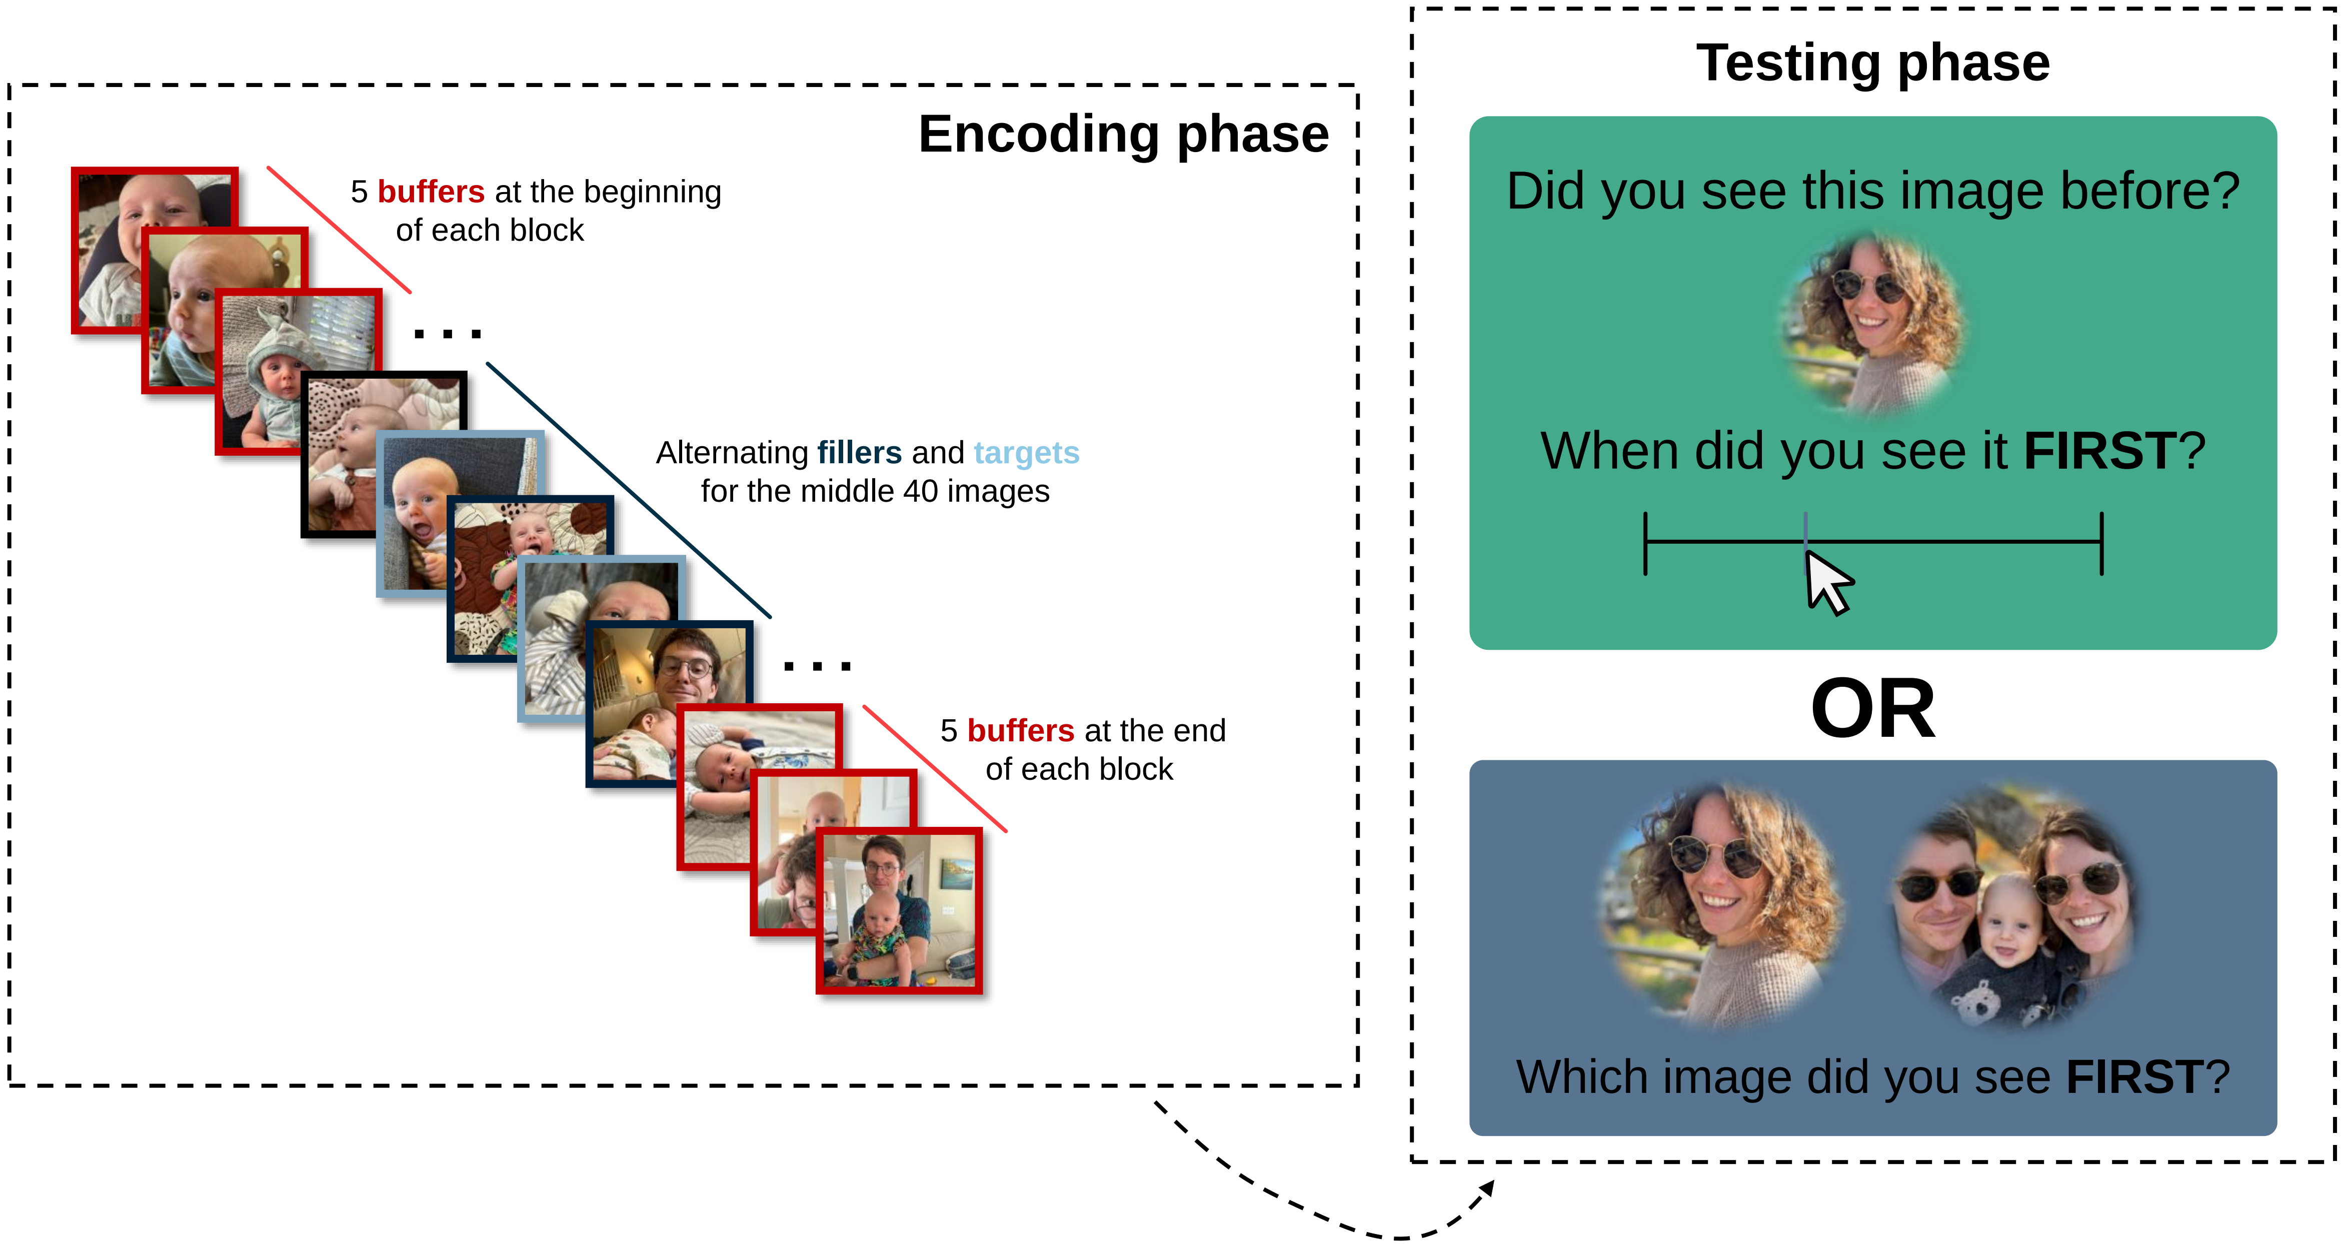

Encoding phase
5 buffers at the beginning of each block
Alternating fillers and targets for the middle 40 images
5 buffers at the end of each block
Testing phase
Did you see this image before?
When did you see it FIRST?
OR
Which image did you see FIRST?

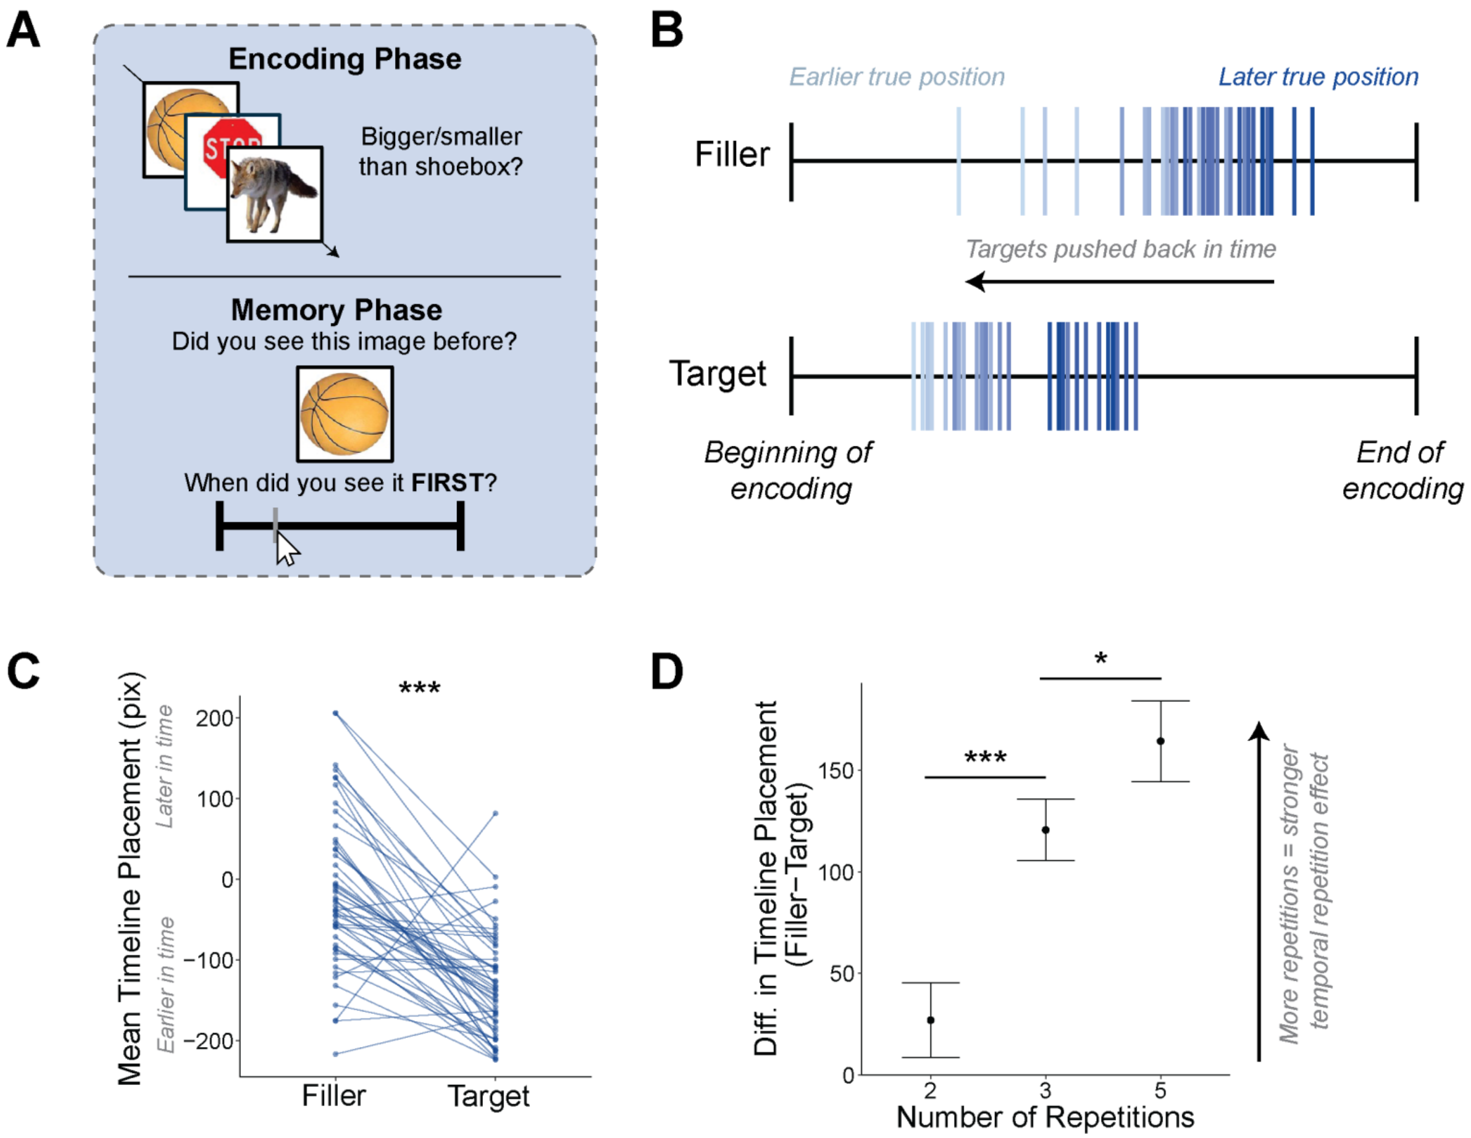
A
Encoding Phase
Bigger/smaller than shoebox?
Memory Phase
Did you see this image before?
When did you see it FIRST?
B
Earlier true position
Later true position
Filler
Targets pushed back in time
Target
Beginning of encoding
End of encoding
C
***
Mean Timeline Placement (pix)
Later in time
Earlier in time
200
100
0
−100
−200
Filler
Target
D
*
***
Diff. in Timeline Placement (Filler−Target)
150
100
50
0
2
3
5
Number of Repetitions
More repetitions = stronger temporal repetition effect

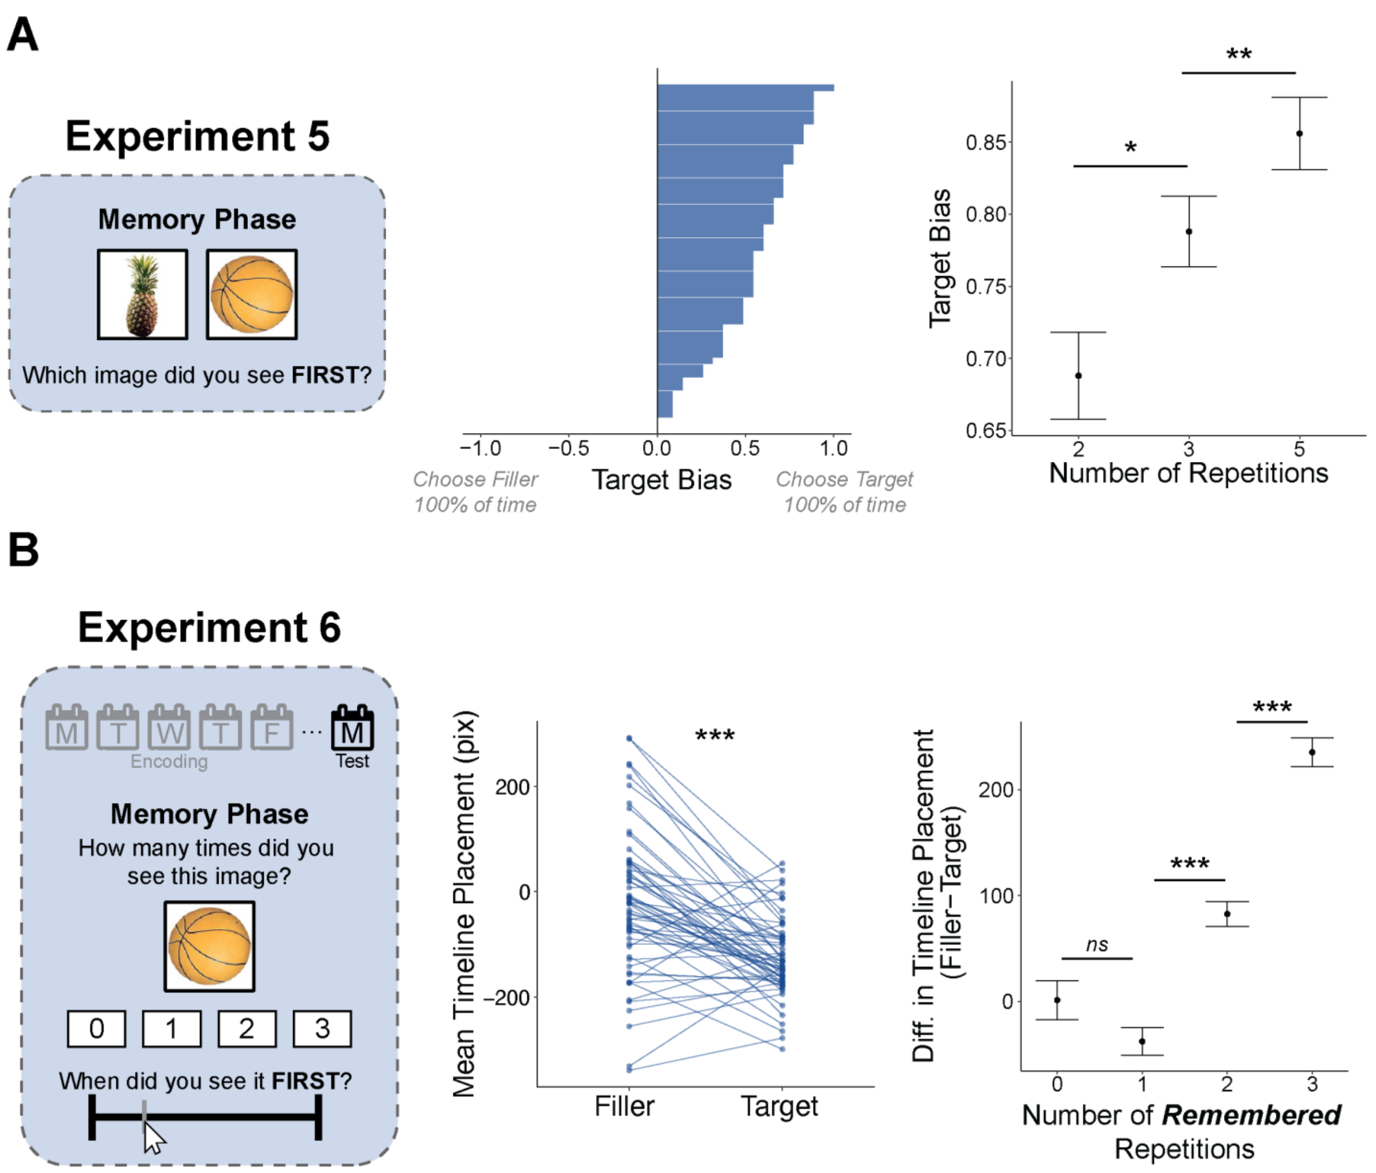

A
Experiment 5
Memory Phase
Which image did you see FIRST?
−1.0
−0.5
0.0
0.5
1.0
Choose Filler 100% of time
Target Bias
Choose Target 100% of time
Target Bias
0.65
0.70
0.75
0.80
0.85
2
3
5
Number of Repetitions
*
**
B
Experiment 6
M T W T F ... M
Encoding
Test
Memory Phase
How many times did you see this image?
0 1 2 3
When did you see it FIRST?
Mean Timeline Placement (pix)
200
0
−200
Filler
Target
***
Diff. in Timeline Placement (Filler−Target)
0
100
200
0
1
2
3
Number of Remembered Repetitions
ns
***
***

An illusion of time caused by repeated experience

Brynn E. Sherman[1] & Sami R. Yousif[2]

[1] University of Pennsylvania

[2] University of North Carolina, Chapel Hill

Running Head	:	An illusion of time
Word Count	:	2022 (Excluding Methods and Results)
Data	:	https://osf.io/nx5t7/?view_only=358f4e72d98b41cd875c19be0c97c090
Version	:	Accepted

!!

Food for thought:

1. For number and faces we talked about domain-specific vs. domain-general coding. **What about time?**

2. Time perception seems really important! **Why are we so bad at it?**

3. **What other factors influence perceived time?**

ILLUSIONS OF EVERYTHING (PART I)

Sami Ryan Yousif

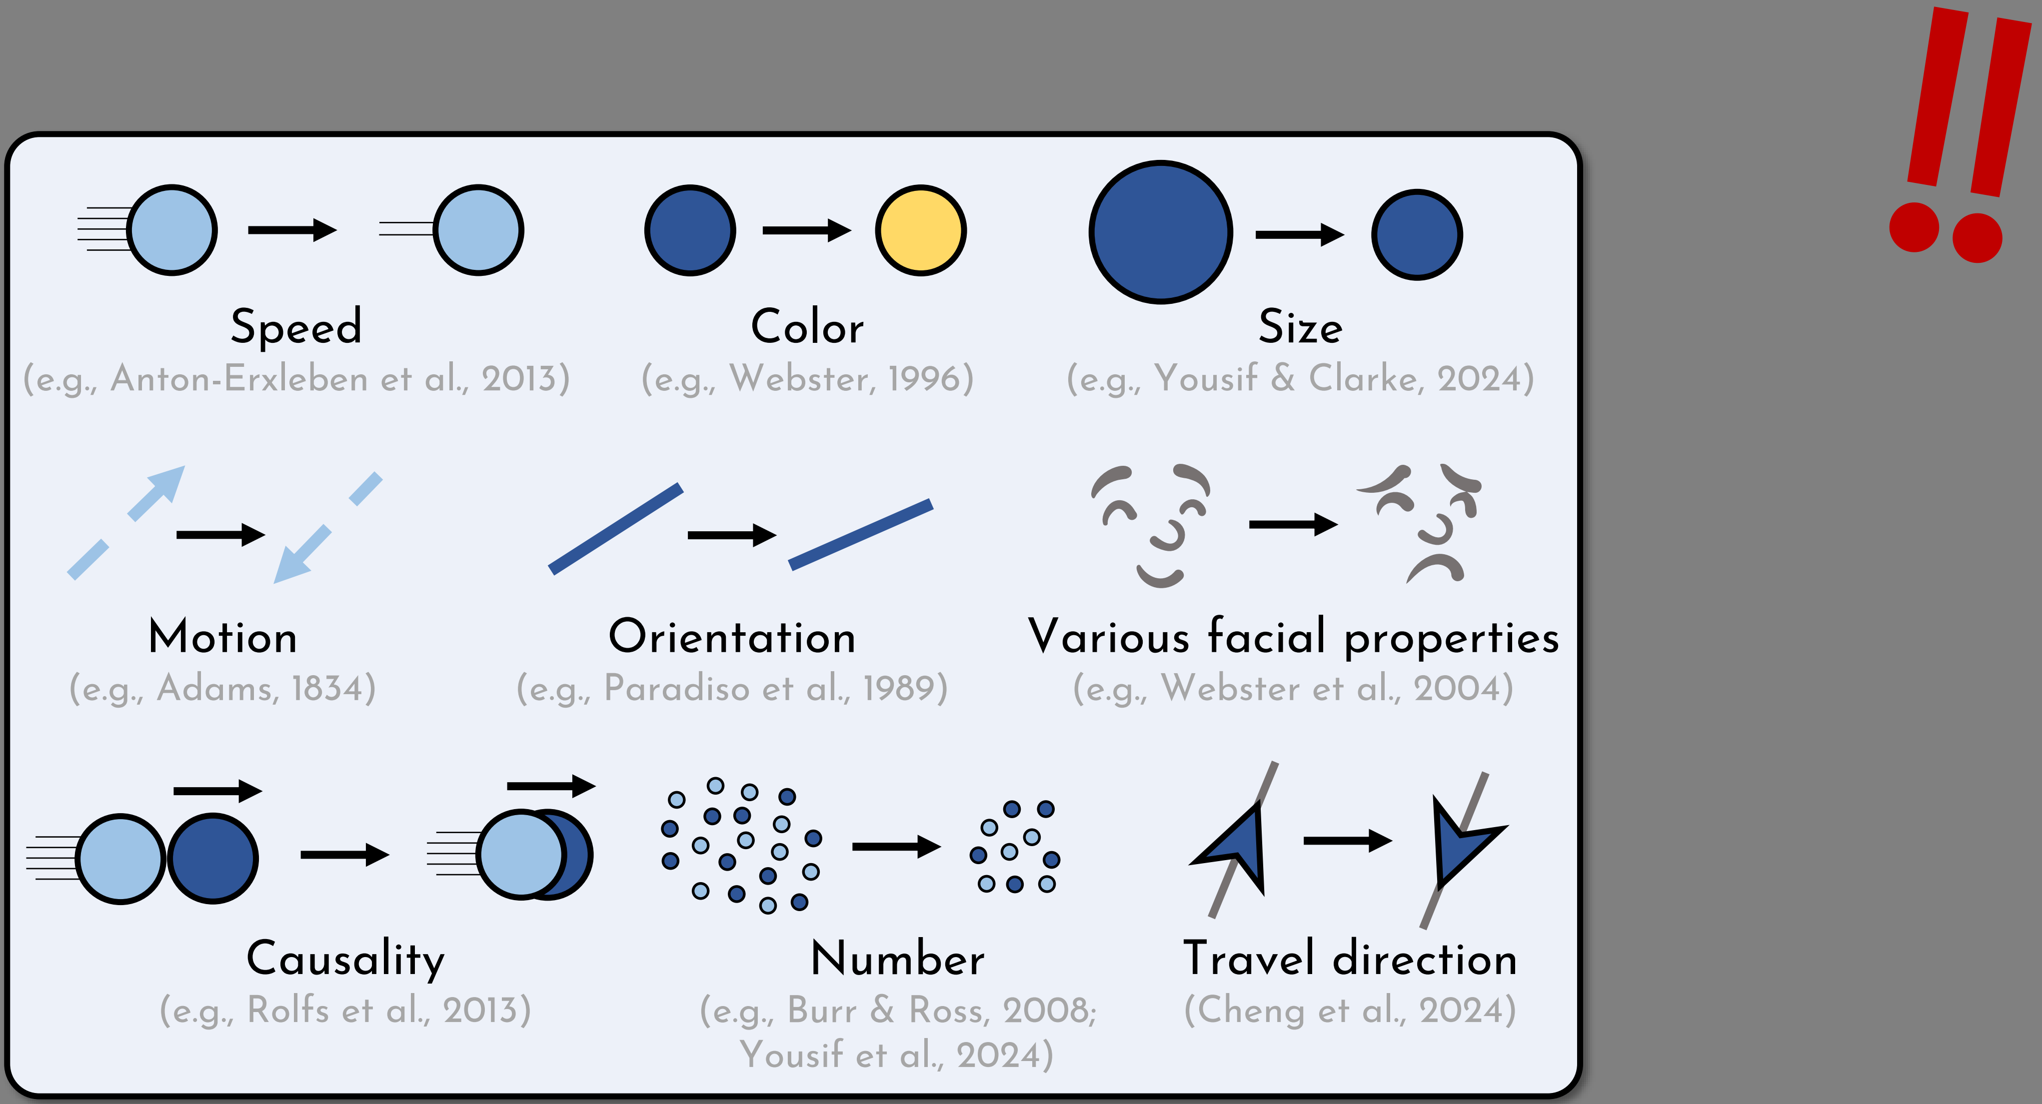

We "adapt" to all sorts of things!

But **adaptation** is more than a ~~visual~~ curiosity!

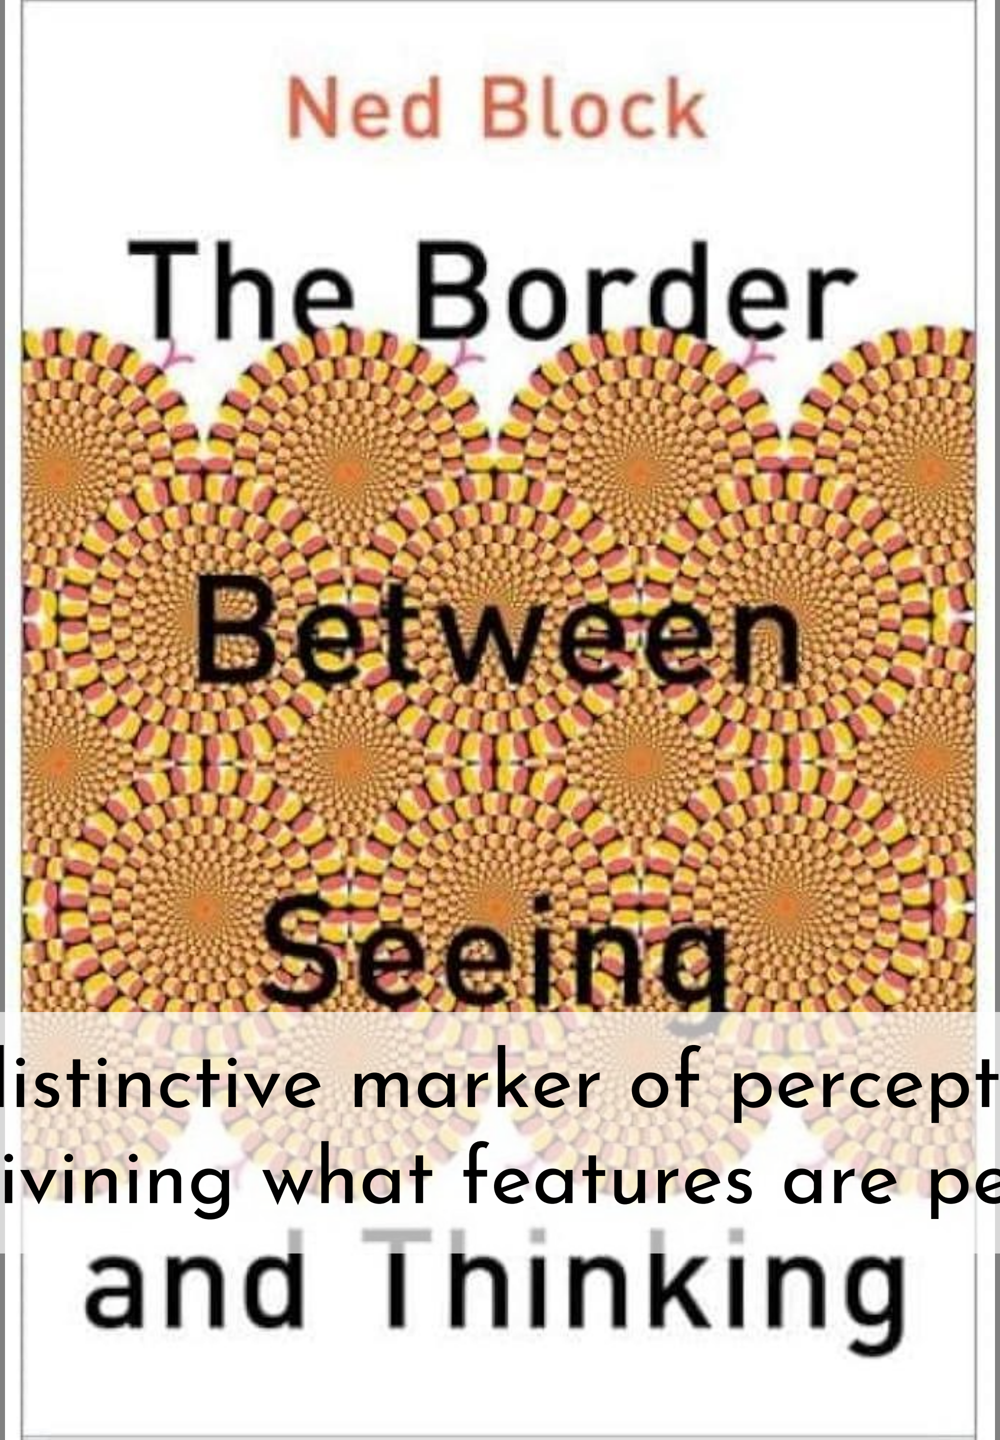

Adaptation is a distinctive marker of perceptual content – thus a **litmus test** for divining what features are perceptual in nature.

PSYCHOLOGY

Prevalence-induced concept change in human judgment

David E. Levari[1], Daniel T. Gilbert[1*], Timothy D. Wilson[2], Beau Sievers[3], David M. Amodio[4], Thalia Wheatley[3]

Why do some social problems seem so intractable? In a series of experiments, we show that people often respond to decreases in the prevalence of a stimulus by expanding their concept of it. When blue dots became rare, participants began to see purple dots as blue; when threatening faces became rare, participants began to see neutral faces as threatening; and when unethical requests became rare, participants began to see innocuous requests as unethical. This "prevalence-induced concept change" occurred even when participants were forewarned about it and even when they were instructed and paid to resist it. Social problems may seem intractable in part because reductions in their prevalence lead people to see more of them.

Adaptation-like effects are said to pervade human judgment, with weighty **social consequences**.

fMR-adaptation: a tool for studying the functional properties of human cortical neurons

Kalanit Grill-Spector [a,*], Rafael Malach [b]

[a] *Department of Brain and Cognitive Sciences, MIT, NE20-444, 77 Mass. Ave., Cambridge, MA 02139, USA*

[b] *Neurobiology Department, Weizmann Institute of Science, Rehovot 76100, Israel*

Received 20 March 2000; received in revised form 21 November 2000; accepted 22 November 2000

Abstract

The invariant properties of human cortical neurons cannot be studied directly by fMRI due to its limited spatial resolution. One voxel obtained from a fMRI scan contains several hundred thousands neurons. Therefore, the fMRI signal may average out a heterogeneous group of highly selective neurons. Here, we present a novel experimental paradigm for fMRI, functional magnetic resonance-adaptation (fMR-A), that enables to tag specific neuronal populations within an area and investigate their functional properties. This approach contrasts with conventional mapping methods that measure the averaged activity of a region. The application of fMR-A to study the functional properties of cortical neurons proceeds in two stages: First, the neuronal population is adapted by repeated presentation of a single stimulus. Second, some property of the stimulus is varied and the recovery from adaptation is assessed. If the signal remains adapted, it will indicate that the neurons are invariant to that attribute. However, if the fMRI signal will recover from the adapted state it would imply that the neurons are sensitive to the property that was varied. Here, an application of fMR-A for studying the invariant properties of high-order object areas (lateral occipital complex – LOC) to changes in object size, position, illumination and rotation is presented. The results show that LOC is less sensitive to changes in object size and position compared to changes of illumination and viewpoint. fMR-A can be extended to other neuronal systems in which adaptation is manifested and can be used with event-related paradigms as well. By manipulating experimental parameters and testing recovery from adaptation it should be possible to gain

Adaptation is a fundamental property of biological systems, and, therefore, a critical **tool** for studying the brain.

Thus: claims about **adaptation** are consequential!

Current Biology 18, 425–428, March 25, 2008 DOI 10.1016/j.cub.2008.02.052

Report

A Visual Sense of Number

David Burr[1,2,*] and John Ross[2]
[1]Dipartimento di Psicologia
Università Degli Studi di Firenze
Via S. Nicolò 89
Florence 50125
Italy
[2]School of Psychology
University of Western Australia
Perth WA 6009
Australia

Summary

Evidence exists for a nonverbal capacity for the apprehension of number, in humans [1] (including infants [2, 3]) and in other primates [4–6]. Here, we show that perceived numerosity is susceptible to adaptation, like primary visual properties of a scene, such as color, contrast, size, and speed. Apparent numerosity was decreased by adaptation to large numbers of dots and increased by adaptation to small numbers, the effect depending entirely on the numerosity of the adaptor, not on contrast, size, orientation, or pixel density, and occurring with very low adaptor contrasts. We suggest that the visual system has the capacity to estimate numerosity and that it is an independent primary visual property, not reducible to others like spatial frequency or density of texture [7].

Results and Discussion

Jevons, a 19th century economist, rather than counting beans, assessed his accuracy in estimating the number of beans in a box at a single glance [8]. He made no errors at four or below but became increasingly inaccurate as the number of beans increased beyond four. Subsequent studies have confirmed his findings and the lack of errors below five has led to the concept of *subitizing* [9–12], usually presumed to be a separate process allowing rapid apprehension of the numerosity of collections containing fewer than five objects. The perception of larger numbers is usually assumed to require slower and more cognitive processes, like counting.

All primary visual properties are susceptible to *adaptation*, sometimes giving rise to dramatic aftereffects, like the waterfall illusion [13], and changes in color, size, distance, spatial frequency, and orientation. If numerosity was a primary visual property, like color or motion, it too should be prone to adaptation. The online demonstration shows that it is. After 30 s adaptation to the two different adaptor patches, the two subsequent patches appear to differ considerably in numerosity (whereas inspection after adaptation wears off, or counting, shows that they both number 30 dots). We quantified adaptation effects by asking subjects whether a *test* stimulus (of variable numerosity), presented to the region that had been adapted, appeared more or less numerous than a *probe* stimulus (of fixed numerosity), presented to a different unadapted position a little later. The proportion of trials where the test appeared more numerous than the probe was plotted against test numerosity and fitted with cumulative Gaussian functions whose mean estimates the point of subjective equality (PSE) between test and probe, and standard deviation the threshold for discriminating between the two (the just-noticeable difference [jnd]). Figure 1A shows sample psychometric functions for a 30 element probe, with and without adaptation to a 400 element stimulus. The ratio of the matched test to probe increases from unity (30 dots) with no adaptation to more than 3 (100 dots) after adaptation (we increased the test number to compensate for the reduction in its apparent numerosity). Note also that that after adaptation the psychometric function is steeper (on logarithmic coordinates), implying a smaller jnd.

We first measured the effect of adapting to a large number (400) of dots as a function of number of dots in the probe (Figure 1B). The amount of adaptation was fairly constant with probe numerosity down to about 12 dots and then decreased as the probe approached the subitizing range. The precision of the match, given by the jnd or Weber fraction (jnd expressed as a fraction of dot number), did not deteriorate during adaptation, the average percentage Weber fractions for unadapted and adapted conditions being 28% for unadapted and 26% for the adapted conditions.

We next investigated whether adaptation to small numbers can cause an increase in apparent numerosity. The red circles of Figure 2 show that adaptation occurred in both directions: Adaptation to small numbers increased apparent numerosity (so the matched number decreased), and adaptation to large numbers decreased apparent numerosity. Adaptation to 50 dots (the number of the probe) had no effect, with the amount of adaptation increasing with the difference between adapt and probe number. The curves of both subjects were well fit by linear regression on log coordinates, with a slope around 0.25.

In order to test whether adaptation depends on numerosity per se or is derived from other factors, like texture density [7], we performed a number of controls. We first varied the size of the adaptor and test dots, in order to vary pixel density. In the above-described study (red circles of Figure 2), both adaptor and test dots were circles of 6 pixel (20 arcmin) diameter (28 pixel area). We repeated the experiment with square adaptor stimuli of 8 × 8 pixels (64 pixels) and test stimuli of 3 × 3 pixels (9 pixels, 1/7 as many as the adaptor). If pixel density were the relevant attribute, the curves of Figure 2 should shift leftwards by a factor of 7, so the null point occurs when adaptor and test pixel density are matched (for adaptation dot number of 7). This clearly does not occur. For naive observer PB, the curves remain superimposed; for DB, there is a slight shift in the opposite direction.

We also examined the effect of adaptor contrast. As Figure 2C shows, contrast of adaptor dots had little effect on the magnitude of adaptation. At contrasts as low as 12%, the adaptation effect is still nearly 2-fold, dropping only near detection threshold. It appears that the only factor that affects adaptation is numerosity, not density, orientation, or contrast.

*Correspondence: dave@in.cnr.it

But I didn’t tell you everything!

PROCEEDINGS OF THE ROYAL SOCIETY B

rspb.royalsocietypublishing.org

Research

A generalized sense of number

Roberto Arrighi[1], Irene Togoli[1] and David C. Burr[1,2]

[1]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, via San Salvi 12, Florence 50135, Italy
[2]Institute of Neuroscience CNR, via Moruzzi 1, Pisa 56124, Italy

Cite this article: Arrighi R, Togoli I, Burr DC. 2014 A generalized sense of number. *Proc. R. Soc. B* **281**: 20141791.
http://dx.doi.org/10.1098/rspb.2014.1791

Received: 18 July 2014
Accepted: 7 October 2014

Much evidence has accumulated to suggest that many animals, including young human infants, possess an abstract sense of approximate quantity, a *number sense*. Most research has concentrated on apparent numerosity of spatial arrays of dots or other objects, but a truly abstract sense of number should be capable of encoding the numerosity of any set of discrete elements, however displayed and in whatever sensory modality. Here, we use the psychophysical technique of *adaptation* to study the sense of number for serially presented items. We show that numerosity of both auditory and visual sequences is greatly affected by prior adaptation to slow or rapid sequences of events. The adaptation to visual stimuli was spatially selective (in external, not retinal coordinates), pointing to a sensory rather than cognitive process. However, adaptation generalized across modalities, from auditory to visual and vice versa. Adaptation also generalized across *formats*: adapting to sequential streams of flashes affected the perceived numerosity of spatial arrays. All these results point to a perceptual system that transcends vision and audition to encode an abstract sense of number in space and in time.

PROCEEDINGS OF THE ROYAL SOCIETY B

rspb.royalsocietypublishing.org

Research

Cite this article: Arrighi R, Togoli I, Burr DC. 2014 A generalized sense of number. *Proc. R. Soc. B* **281**: 20141791.
http://dx.doi.org/10.1098/rspb.2014.1791

Received: 18 July 2014
Accepted: 7 October 2014

A generalized sense of number

Roberto Arrighi[1], Irene Togoli[1] and David C. Burr[1,2]

[1]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, via San Salvi 12, Florence 50135, Italy
[2]Institute of Neuroscience CNR, via Moruzzi 1, Pisa 56124, Italy

Much evidence has accumulated to suggest that many animals, including young human infants, possess an abstract sense of approximate quantity, a *number sense*. Most research has concentrated on apparent numerosity of spatial arrays of dots or other objects, but a truly abstract sense of number should be capable of encoding the numerosity of any set of discrete elements, however displayed and in whatever sensory modality. Here, we use the psychophysical technique of *adaptation* to study the sense of number for serially presented items. We show that numerosity of both auditory and visual sequences is greatly affected by prior adaptation to slow or rapid sequences of events. The adaptation to visual stimuli was spatially selective (in external, not retinal coordinates), pointing to a sensory rather than cognitive process. However, adaptation generalized across modalities, from auditory to visual and vice versa. Adaptation also generalized across *formats*: adapting to sequential streams of flashes affected the perceived numerosity of spatial arrays. All these results point to a perceptual system that transcends vision and audition to encode an abstract sense of number in space and in time.

eLIFE elifesciences.org

RESEARCH ARTICLE

A shared numerical representation for action and perception

Giovanni Anobile[1,2†], Roberto Arrighi[1†], Irene Togoli[1], David Charles Burr[1,3,4*]

[1]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, Florence, Italy; [2]Department of Developmental Neuroscience, Stella Maris Scientific Institute, Pisa, Italy; [3]Institute of Neuroscience, National Research Council, Pisa, Italy; [4]School of Psychology, University of Western Australia, Perth, Australia

Abstract Humans and other species have perceptual mechanisms dedicated to estimating approximate quantity: a *sense of number*. Here we show a clear interaction between self-produced actions and the perceived numerosity of subsequent visual stimuli. A short period of rapid finger-tapping (without sensory feedback) caused subjects to underestimate the number of visual stimuli presented near the tapping region; and a period of slow tapping caused overestimation. The distortions occurred both for stimuli presented sequentially (series of flashes) and simultaneously (clouds of dots); both for magnitude estimation and forced-choice comparison. The adaptation was spatially selective, primarily in external, real-world coordinates. Our results sit well with studies reporting links between perception and action, showing that vision and action share mechanisms that encode numbers: a generalized *number sense*, which estimates the number of self-generated as well as external events.

DOI: 10.7554/eLife.16161.001

***For correspondence:** dave@in.cnr.it

†These authors contributed equally to this work

Competing interests: The authors declare that no competing interests exist.

Funding: See page 12

Received: 17 March 2016
Accepted: 18 July 2016
Published: 09 August 2016

Reviewing editor: Richard Ivry, University of California, Berkeley, United States

Introduction

Animals, including humans, estimate spontaneously and reasonably accurately the approximate quantity of arrays of objects, without recourse to other forms of representation, such as density (*Cicchini et al., 2016*). Even newborn infants of less than 3 days show selective habituation to number (*Izard et al., 2009*). There is now very good evidence in both human and non-human primates that number is encoded by intraparietal and prefrontal cortex (*Castelli et al., 2006*; *Dehaene et al., 2003*; *Harvey et al., 2013*; *Nieder, 2005, 2012, 2016*; *Nieder et al., 2006*; *Nieder and Miller, 2004*; *Piazza and Eger, 2016*; *Piazza et al., 2004, 2007*), even in numerically naive monkeys (*Viswanathan and Nieder, 2013*). All these studies point to the existence of a *visual sense of number* within a parietal–frontal network (*Dehaene, 2011*).

A truly abstract sense of number should be capable of encoding the numerosity of any set of discrete elements, displayed simultaneously or sequentially, in whatever sensory modality. Some evidence exists for such a generalized number sense. Neurons in the lateral prefrontal cortex (lPFC) of behaving monkeys encode numerosity for both auditory and visual sensory modalities, suggesting supra-modal numerosity processing (*Nieder, 2012*). Another study reported separate populations of neurons in the intraparietal sulcus (IPS) responding selectively to sequential or simultaneous numerical displays, while a third set of neurons showed numerosity selectivity for both simultaneous and sequential presentations, suggesting that the information about spatial and temporal numerosity converges to a more abstract representation (*Nieder et al., 2006*). There is also evidence from functional imaging in humans for a right lateralized fronto-parietal circuit activated by both auditory and visual number sequences, and that right IPS is involved in processing both sequential and simultaneous numerosity formats (*Castelli et al., 2006*; *Piazza et al., 2006*).

eLIFE elifesciences.org

RESEARCH ARTICLE

A shared numerical representation for action and perception

Giovanni Anobile[1,2†], Roberto Arrighi[1†], Irene Togoli[1], David Charles Burr[1,3,4*]

[1]Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, Florence, Italy; [2]Department of Developmental Neuroscience, Stella Maris Scientific Institute, Pisa, Italy; [3]Institute of Neuroscience, National Research Council, Pisa, Italy; [4]School of Psychology, University of Western Australia, Perth, Australia

Abstract Humans and other species have perceptual mechanisms dedicated to estimating approximate quantity: a *sense of number*. Here we show a clear interaction between self-produced actions and the perceived numerosity of subsequent visual stimuli. A short period of rapid finger-tapping (without sensory feedback) caused subjects to underestimate the number of visual stimuli presented near the tapping region; and a period of slow tapping caused overestimation. The distortions occurred both for stimuli presented sequentially (series of flashes) and simultaneously (clouds of dots); both for magnitude estimation and forced-choice comparison. The adaptation was spatially selective, primarily in external, real-world coordinates. Our results sit well with studies reporting links between perception and action, showing that vision and action share mechanisms that encode numbers: a generalized *number sense*, which estimates the number of self-generated as well as external events.
DOI: 10.7554/eLife.16161.001

*For correspondence: dave@in.cnr.it

†These authors contributed equally to this work

Competing interests: The authors declare that no competing interests exist.

Funding: See page 12

Received: 17 March 2016
Accepted: 18 July 2016
Published: 09 August 2016

Reviewing editor: Richard Ivry, University of California, Berkeley, United States

Introduction

Animals, including humans, estimate spontaneously and reasonably accurately the approximate quantity of arrays of objects, without recourse to other forms of representation, such as density (*Cicchini et al., 2016*). Even newborn infants of less than 3 days show selective habituation to number (*Izard et al., 2009*). There is now very good evidence in both human and non-human primates that number is encoded by intraparietal and prefrontal cortex (*Castelli et al., 2006; Dehaene et al., 2003; Harvey et al., 2013; Nieder, 2005, 2012, 2016; Nieder et al., 2006; Nieder and Miller, 2004; Piazza and Eger, 2016; Piazza et al., 2004, 2007*), even in numerically naive monkeys (*Viswanathan and Nieder, 2013*). All these studies point to the existence of a *visual sense of number* within a parietal–frontal network (*Dehaene, 2011*).

A truly abstract sense of number should be capable of encoding the numerosity of any set of discrete elements, displayed simultaneously or sequentially, in whatever sensory modality. Some evidence exists for such a generalized number sense. Neurons in the lateral prefrontal cortex (lPFC) of behaving monkeys encode numerosity for both auditory and visual sensory modalities, suggesting supra-modal numerosity processing (*Nieder, 2012*). Another study reported separate populations of neurons in the intraparietal sulcus (IPS) responding selectively to sequential or simultaneous numerical displays, while a third set of neurons showed numerosity selectivity for both simultaneous and sequential presentations, suggesting that the information about spatial and temporal numerosity converges to a more abstract representation (*Nieder et al., 2006*). There is also evidence from functional imaging in humans for a right lateralized fronto-parietal circuit activated by both auditory and visual number sequences, and that right IPS is involved in processing both sequential and simultaneous numerosity formats (*Castelli et al., 2006; Piazza et al., 2006*).

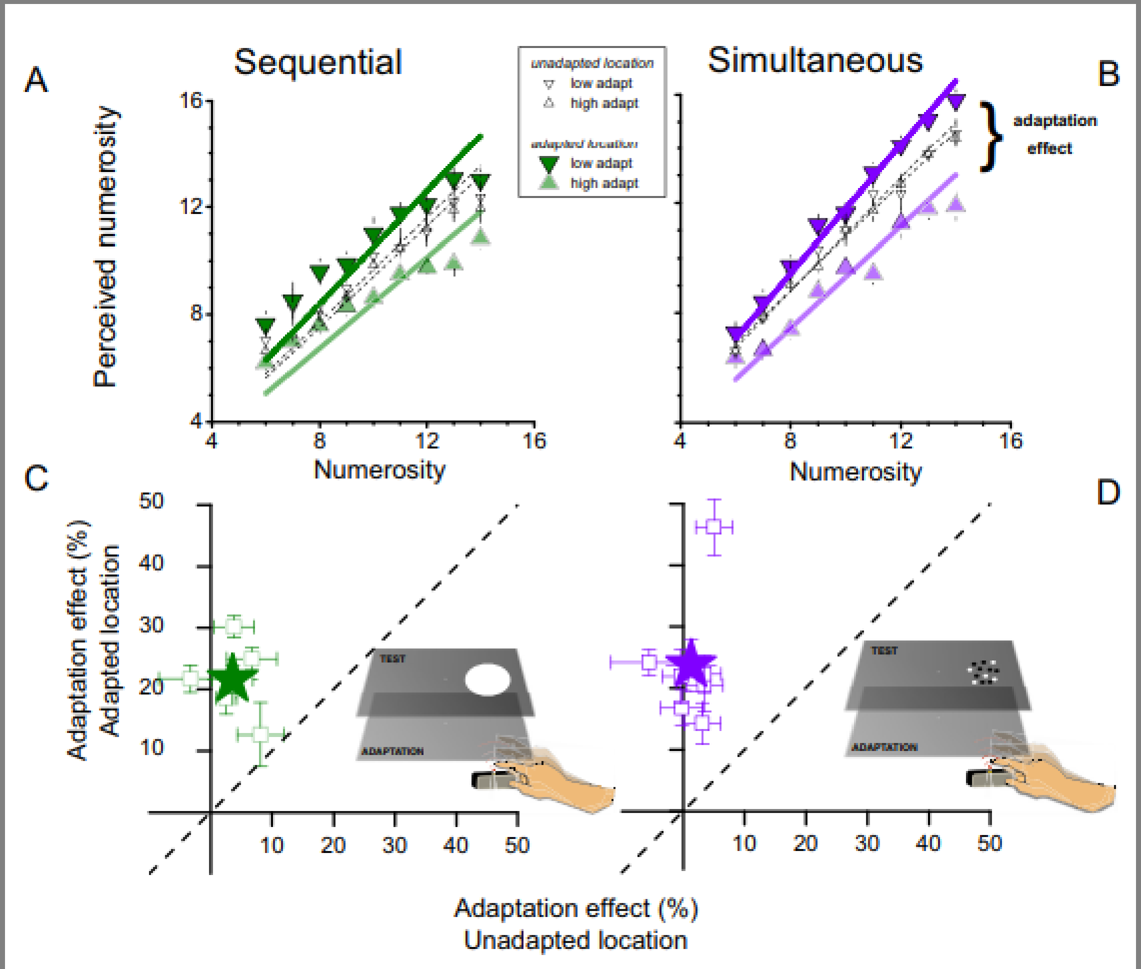

This is weird, because other times, adaptation is highly fragile!

iScience

CellPress
OPEN ACCESS

Article

Numerosity perception is tuned to salient environmental features

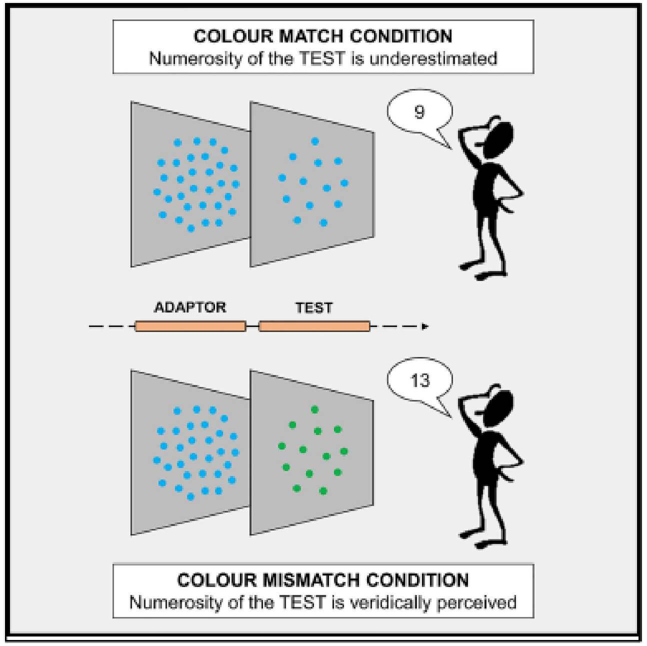

Paolo Antonino Grasso, Giovanni Anobile, Roberto Arrighi, David Charles Burr, Guido Marco Cicchini

paolo.grasso@unifi.it

Highlights

Numerosity perception is highly selective to salient non-numerical features

Visual numerosity adaptation requires similar perceived test and adapter color

Auditory numerosity adaptation requires similar test and adapter pitch

Grasso et al., iScience 25, 104104
April 15, 2022
https://doi.org/10.1016/j.isci.2022.104104

Cognition 249 (2024) 105813

Contents lists available at ScienceDirect

Cognition

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

Number adaptation: A critical look

Sami R. Yousif[a,*], Sam Clarke[b], Elizabeth M. Brannon[a]

[a] *Department of Psychology, University of Pennsylvania, USA*
[b] *Department of Philosophy, University of Southern California, USA*

ARTICLE INFO

Keywords:
Number
Perception
Adaptation
Old news

ABSTRACT

It is often assumed that adaptation — a temporary change in sensitivity to a perceptual dimension following exposure to that dimension — is a litmus test for what is and is not a "primary visual attribute". Thus, papers purporting to find evidence of number adaptation motivate a claim of great significance: That number is something that can be seen in much the way that canonical visual features, like color, contrast, size, and speed, can. Fifteen years after its reported discovery, number adaptation's existence seems to be nearly undisputed, with dozens of papers documenting support for the phenomenon. The aim of this paper is to offer a counterweight — to critically assess the evidence for and against number adaptation. After surveying the many reasons for thinking that number adaptation exists, we introduce several lesser-known reasons to be skeptical. We then advance an alternative account — the old news hypothesis — which can accommodate previously published findings while explaining various (otherwise unexplained) anomalies in the existing literature. Next, we describe the results of eight pre-registered experiments which pit our novel old news hypothesis against the received number adaptation hypothesis. Collectively, the results of these experiments undermine the number adaptation hypothesis on several fronts, whilst consistently supporting the old news hypothesis. More broadly our work raises questions about the status of adaptation itself as a means of discerning what is and is not a visual attribute.

1. Introduction

It is sometimes joked that vision science primarily serves to catalogue phenomena long known by magicians, cinematographers, and petty thieves. Occasionally, however, its discoveries offer to profoundly transform our understanding of what it means to see. Take the reported discovery of *visual number adaptation*. Since the pioneering work of Burr and Ross (2008) it has become widely accepted that observers visually adapt to the number of items in a seen collection, much as we adapt to other visible properties, like color, size, and motion. The claim is that prolonged exposure to a large number of seen items causes a middling number of items in that region to appear less numerous than they otherwise would. Conversely, prolonged exposure to a small number of items reportedly causes a middling number of items in that region to appear more numerous than they otherwise would.

These are stunning results. In canonical examples of visual number adaptation, observers enjoy obvious and phenomenologically striking aftereffects. If you adapt to 300 dots in a left-hand region of visual space, a test display containing 100 dots in that region will look remarkably sparse when compared to an otherwise identical collection of 100 dots in an un-adapted region (see Burr & Ross, 2008; see also Demo #1 in the supplemental materials on our OSF page). Since researchers have taken steps to rule out simpler explanations (e.g., by controlling for the total brightness and/or surface area of collections), received wisdom is that these results reflect adaptation to *the number* of items in seen collections. And because adaptation effects of this sort have been deemed rare or absent from thought and post-perceptual cognition (Block, 2022; Webster, 2015; c.f. Phillips & Firestone, 2022), number adaptation has been taken to suggest that number is a "primary" visual attribute, on a par with color and other low-level visual properties (Anobile, Cicchini, & Burr, 2016; c.f., Smortchkova, 2020). So, while numbers are abstract objects, located outside of space or time, number adaptation has been taken to establish that numbers nevertheless feature in the contents of human vision and visual experience; that, strange as it sounds, we literally *see* number.

Given the practical, philosophical, and theoretical implications of these claims, it is perhaps surprising that the existence of visual number adaptation has gone largely unchallenged (but see Dakin, Tibber,

* Corresponding author at: Department of Psychology, University of Pennsylvania, 425 S. University Ave, Stephen A. Levin Bldg., Philadelphia, PA 19104-6241, USA.

E-mail address: sryousif@sas.upenn.edu (S.R. Yousif).

https://doi.org/10.1016/j.cognition.2024.105813
Received 18 February 2024; Received in revised form 7 May 2024; Accepted 13 May 2024
Available online 30 May 2024

To *see* more: cogdevlab.org/number-adaptation/

	Repulsivity?	Phenomenology?	Retinotopy?	Spatiotopy?	Bidirectionality?	Attention?	Cross-modal?	Cross-stimulus?
Color	✓	✓	✓	✓	✓	–	–	–
Motion	✓	✓	✓	✓	✓	–	–	–
Orientation	✓	?	✓	✓	✓	?	–	–
Speed	✓	?	✓	✓	?	✓	–	–
Facial features	✓	X	✓	✓	✓	–	–	–
Causality	✓	X	✓*	✓	✓*	–	–	–
Number	✓	?	X	?	X	?	?	✓
Size	✓	X	X	?	?	?	–	–
Variance	✓	X	–	✓	✓	–	–	✓
Repetition suppression	✓	X	X	X	–	–	✓	✓
PICC	✓	X	–	–	✓	–	–	–
Serial depend.	X	X	✓	✓	✓	✓	X	–
Rand. # gen.	✓	X	–	–	✓	–	✓	–

"**Adaptation** is not a unified kind".

Adaptation to value?

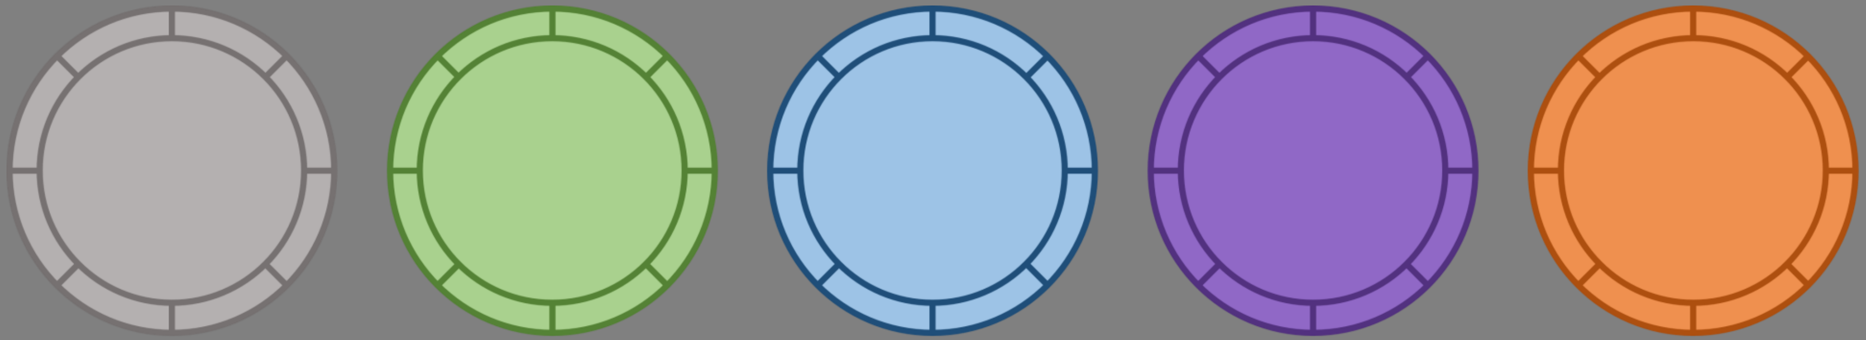

1 point 2 points 3 points 4 points 5 points

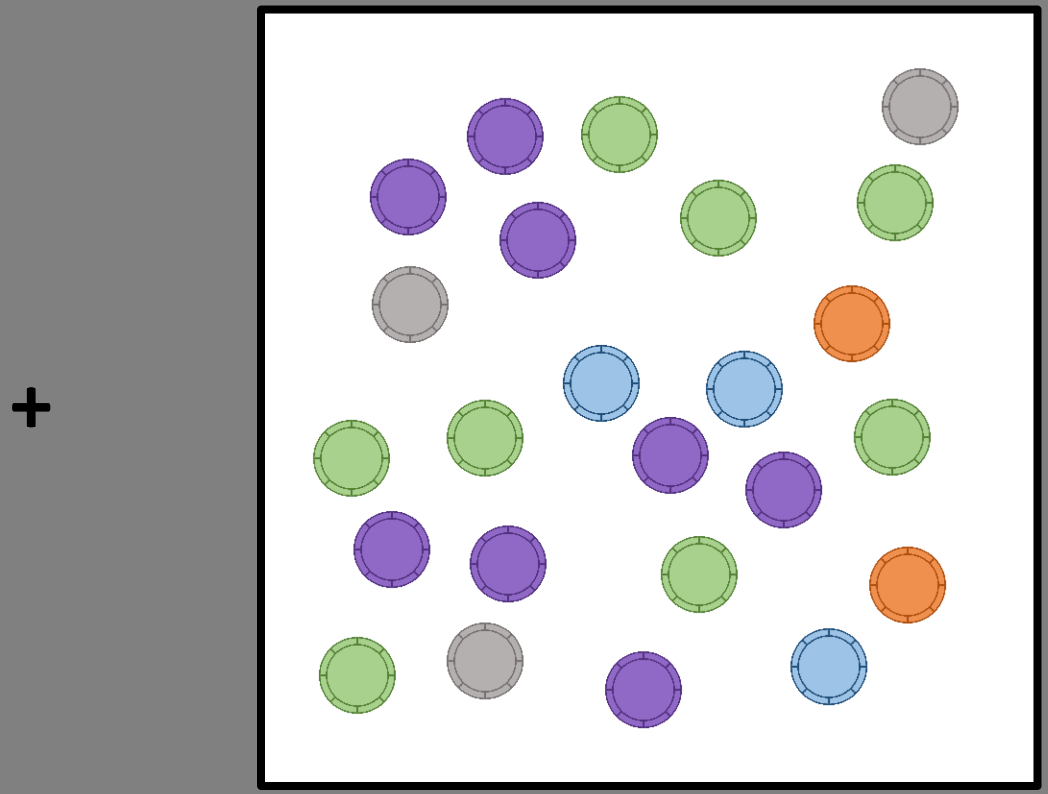

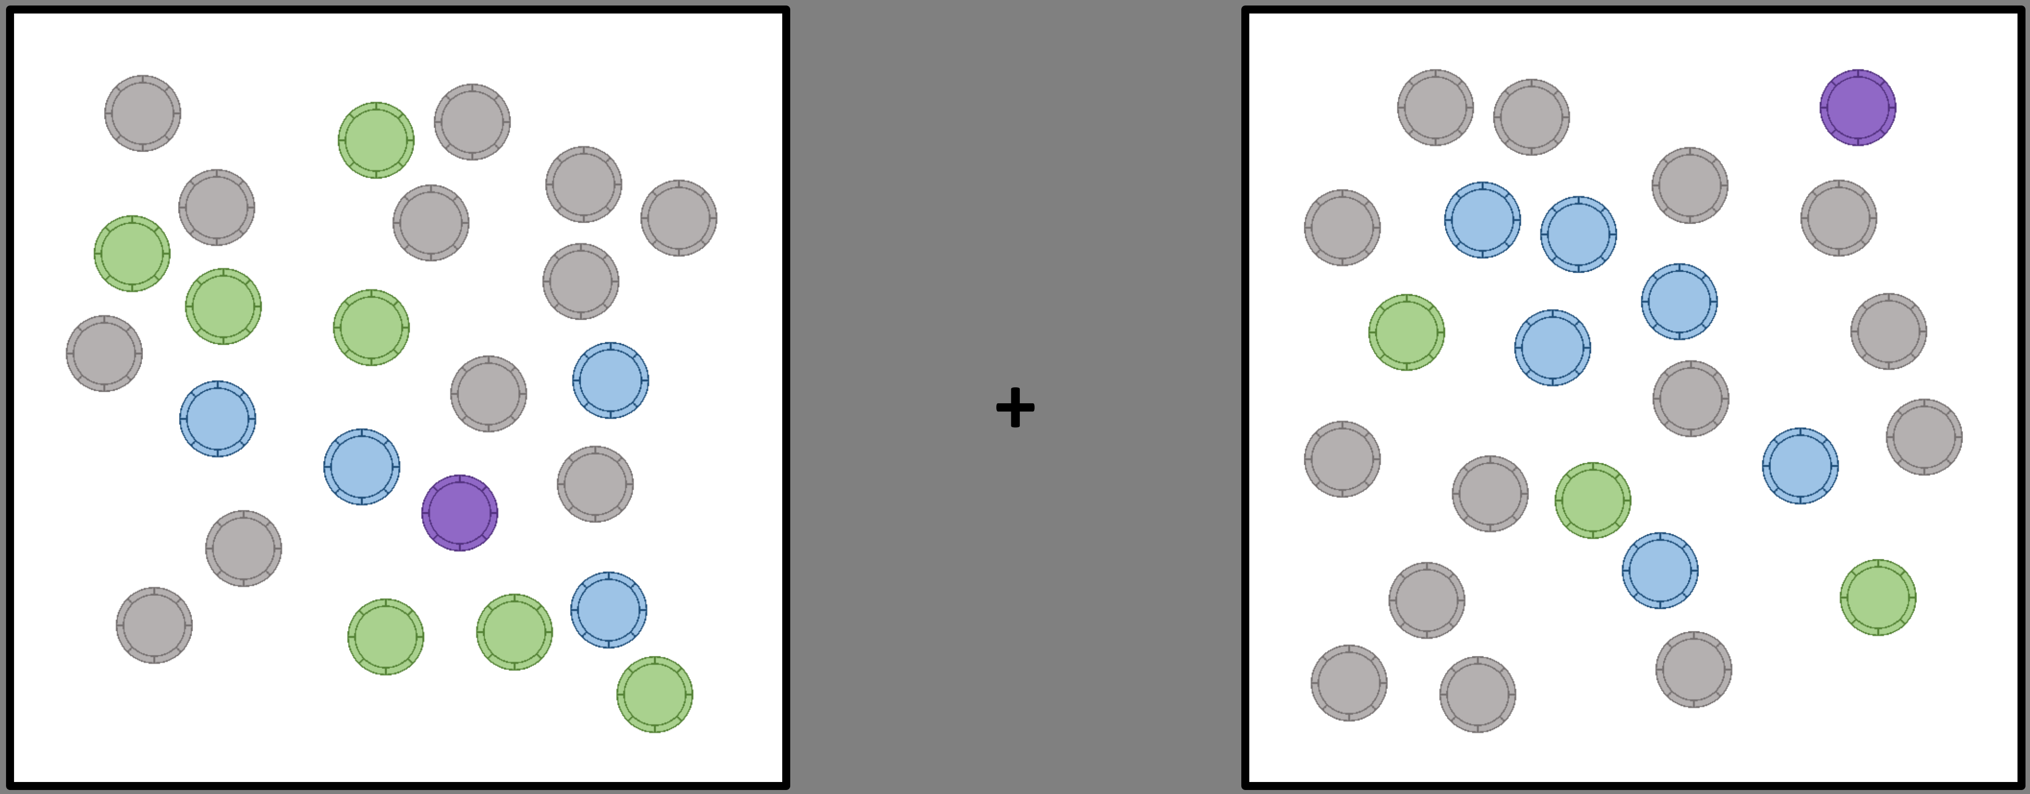

The idea is that **perceptual adaptation** to value isn't possible.

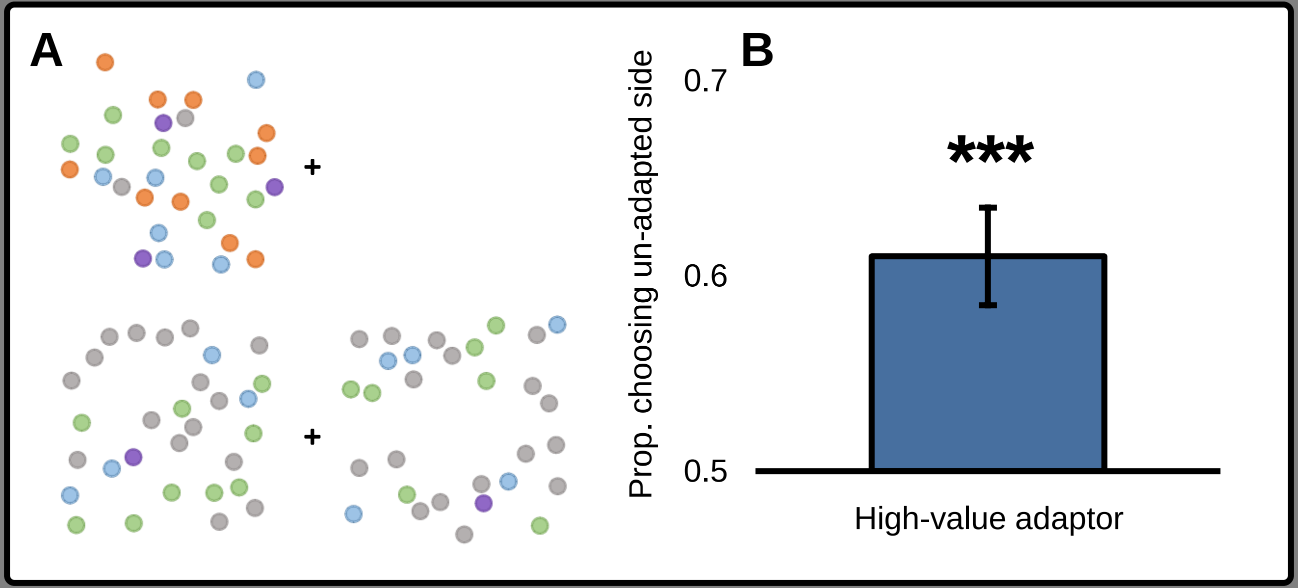
A
B
Prop. choosing un-adapted side
0.7
0.6
0.5
***
High-value adaptor

How far does this go?

Helton (2016)

Is adaptation really just a perceptual phenomenon?

ILLUSIONS OF EVERYTHING (PART II)

Sami Ryan Yousif

Top-down effects

Top-down effects

A **top-down effect** is any case where what you **think** affects what you **perceive**.

A famous example…

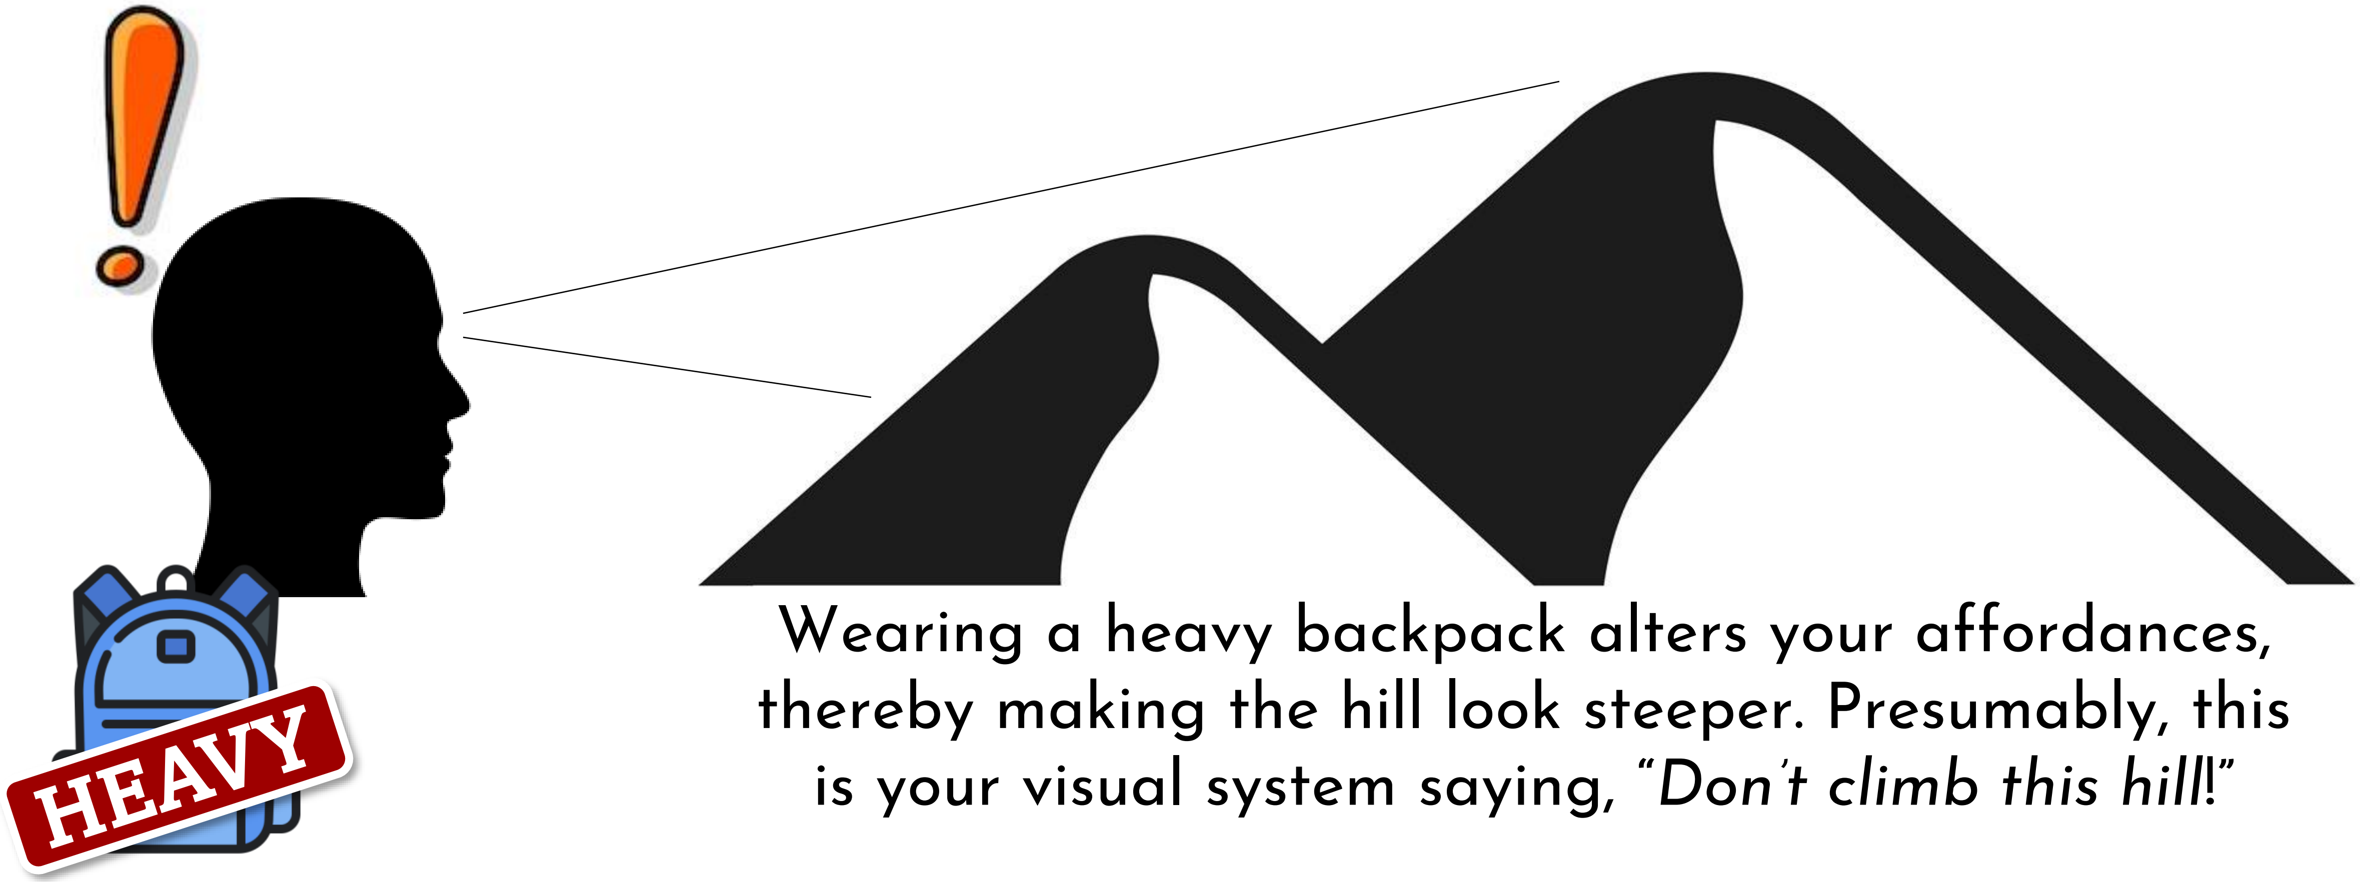

Wearing a heavy backpack alters your affordances, thereby making the hill look steeper. Presumably, this is your visual system saying, *“Don’t climb this hill!”*

Journal of Experimental Psychology:
Human Perception and Performance
1999, Vol. 25, No. 4, 1076–1096

0096-1523/99/$3.00

Visual–Motor Recalibration in Geographical Slant Perception

Mukul Bhalla
Loyola University New Orleans

Dennis R. Proffitt
University of Virginia

In 4 experiments, it was shown that hills appear steeper to people who are encumbered by wearing a heavy backpack (Experiment 1), are fatigued (Experiment 2), are of low physical fitness (Experiment 3), or are elderly and/or in declining health (Experiment 4). Visually guided actions are unaffected by these manipulations of physiological potential. Although dissociable, the awareness and action systems were also shown to be interconnected. Recalibration of the transformation relating awareness and actions was found to occur over long-term changes in physiological potential (fitness level, age, and health) but not with transitory changes (fatigue and load). Findings are discussed in terms of a time-dependent coordination between the separate systems that control explicit visual awareness and visually guided action.

In conscious awareness, the apparent slant of hills is greatly exaggerated. For example, 5° hills appear to be about 20°, and 10° ones look to be about 30° (Proffitt, Bhalla, Gossweiler, & Midgett, 1995). Be that as it may, people are not especially prone to stumble whenever the terrain over which they walk changes in slant. When ascending a 5° hill, people appropriately raise their feet to accommodate this incline, not a 20° one.

In this and our previous article (Proffitt et al., 1995), we suggest that the exaggeration of slant in conscious awareness promotes the function of relating distal inclines to one's physiological potential. Given gravity and one's physiology, a long 5° hill is actually rather difficult to ascend, and consequently it appears to be quite steep. Given this proposal, we predict that geographical slant will change with changes in physiological potential.

The first purpose of this article is to provide support for this prediction. In four experiments, we show that hills appear steeper when people are (a) encumbered by wearing a heavy backpack, (b) fatigued after a long run, (c) of low physical fitness, and (d) elderly or in poor health. None of these manipulations influenced a measure of visually guided actions directed at geographical slants.

This article's second purpose is to show that the visual systems that inform conscious awareness of slant and visually guided actions are dissociable yet are also transformationally connected. The evidence for dissociation derives from the fact that manipulations of physiological potential evoke changes in conscious awareness without any concomitant changes in visually guided actions. The evidence for transformational connectedness is seen in the finding that the guidance system can be indirectly informed by conscious representations. This is illustrated in the following example: Without looking at a hill, a person can be asked to make a motor adjustment conforming to some verbally given slant angle, say 20°. The response to such an instruction will be a motor response of 5°. Note from the earlier example that this would be the appropriate motor accommodation to a 5° hill that, in consciousness, appeared to be 20°. Conscious representations and motor adjustments are not the same, but they are internally consistent. All four of the studies reported herein investigated the time course over which this internal consistency is maintained. We found that it is not maintained over transitory changes in physiological potential lasting an hour or less; however, it is maintained over longer periods of months and years.

Conscious Slant Perception

In our previous article (Proffitt et al., 1995), we proposed that the discrepancy between conscious slant perception and visually guided actions is functionally advantageous given the different goals that these different systems subserve. Conscious slant perception informs the planning of relatively long-term molar behaviors such as selecting and modulating gait style, whereas the visual guidance system informs the execution of specific behaviors in the immediate action space. A goal of gait style selection is to maintain an acceptable rate of energy expenditure. Whenever terrain slant changes, gait must also change, or one's aerobic state will fluctuate outside of desired values. Gait selection is future oriented in that it requires the regulation of behaviors

Mukul Bhalla, Department of Psychology, Loyola University New Orleans; Dennis R. Proffitt, Department of Psychology, University of Virginia.

The experiments reported in this article were part of Mukul Bhalla's dissertation research conducted at the University of Virginia. This research was supported by National Institute of Mental Health Grant MH52640 and National Aeronautics and Space Administration Grant NCC-2-925.

We thank the students of the University of Virginia and the residents and members of The Collonades retirement community and the Charlottesville Senior Center for participating in these experiments. We would also like to thank Jill Seaks and Marie Anderson for help in running experiments. Finally, we gratefully acknowledge the helpful comments and advice of Marco Bertamini, Bennett Bertenthal, Linda Bunker, Sarah Creem, Glen Gaesser, Michael Kubovy, and Tyrone Yang.

Correspondence concerning this article should be addressed to Mukul Bhalla, Department of Psychology, Box 194, Loyola University, 6363 St. Charles Avenue, New Orleans, Louisiana 70118. Electronic mail may be sent to bhalla@nadal.loyno.edu.

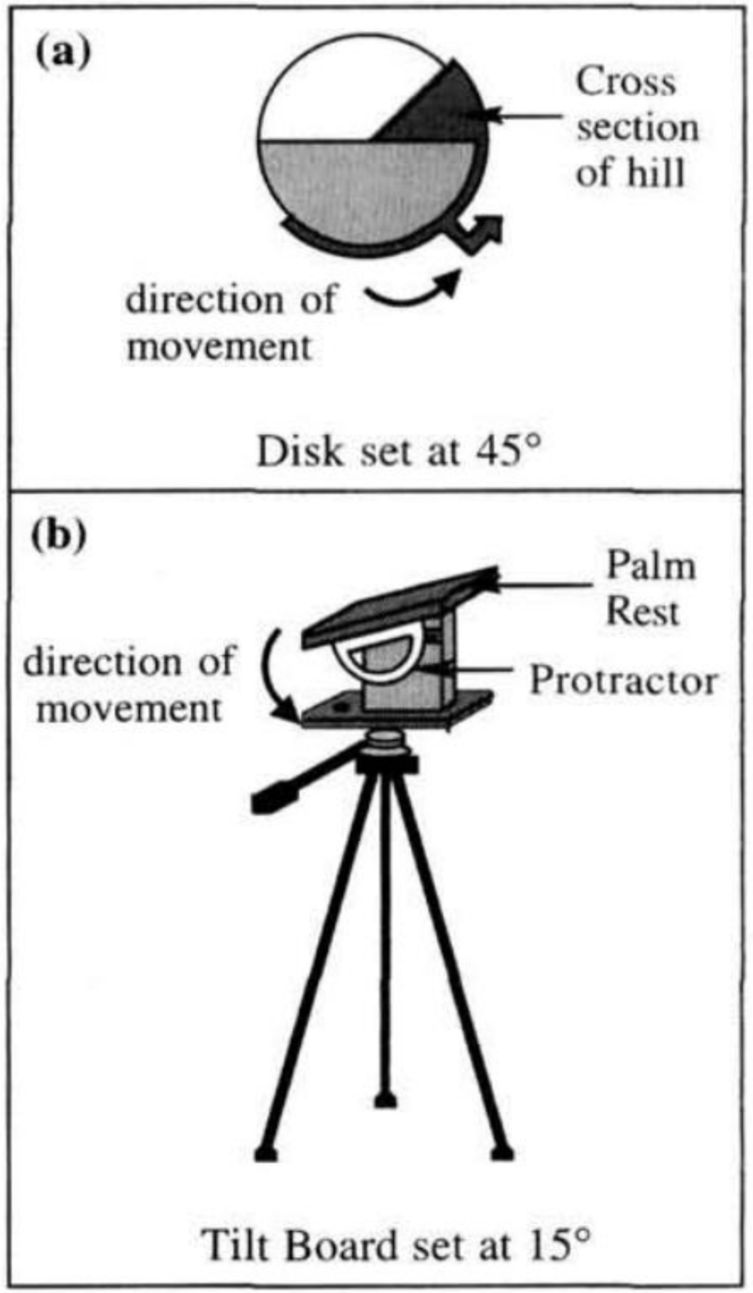

But it isn't just hills and steepness...

- Body size influences size perception
- Emotional arousal influences height perception
- Field goal kicking performance influences size perception
- Sentences about faces affects face perception
- Glucose levels influence slant perception
- Social closeness influence spatial perception
- Parkour ability influences height perception
- Tool use influences distance perception
- Object knowledge influences color perception
- Threat influences distance perception

But it isn't just hills and steepness…

- Imagining danger influences distance perception
- Ease of blocking a ball influences speed perception
- Social comparison influences face perception
- "Illusory ownership of a virtual child body" influences size perception (???)
- "The presence of comrades decreases… formidability of an opponent" (???)
- Age influences distance perception (of walks)
- Morality influences perceptual pop-out
- Political leaning influences the perception of Obama's skin color
- Social power influences weight perception
- Self-affirmation influences distance perception
- Revealing secrets influences slant perception

To make sense of top-down effects, we need to understand modularity.

!!

In the broadest sense, modularity refers to the simple fact that the mind has parts, not unlike your computer has different apps and programs (Fodor, 1983)

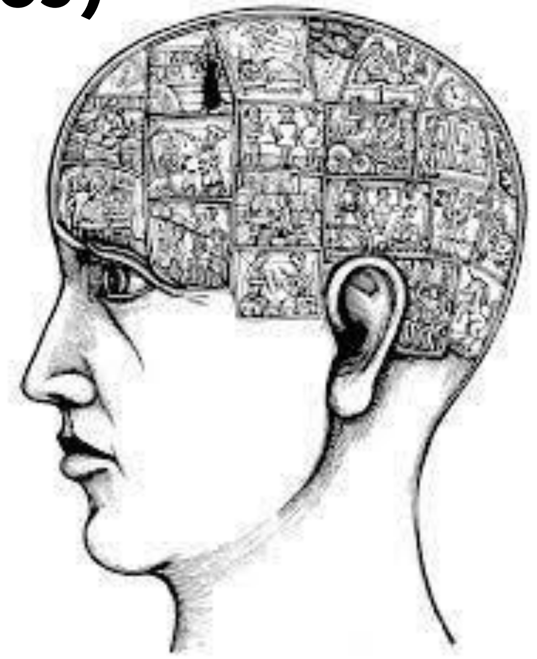

Nine characteristics of modularity…

1. Domain specificity
2. Mandatory operation
3. Limited central accessibility
4. Fast processing
5. **Informational encapsulation***
6. ‘Shallow’ outputs
7. Fixed neural architecture
8. Characteristic and specific breakdown patterns
9. Characteristic ontogenetic pace and sequencing

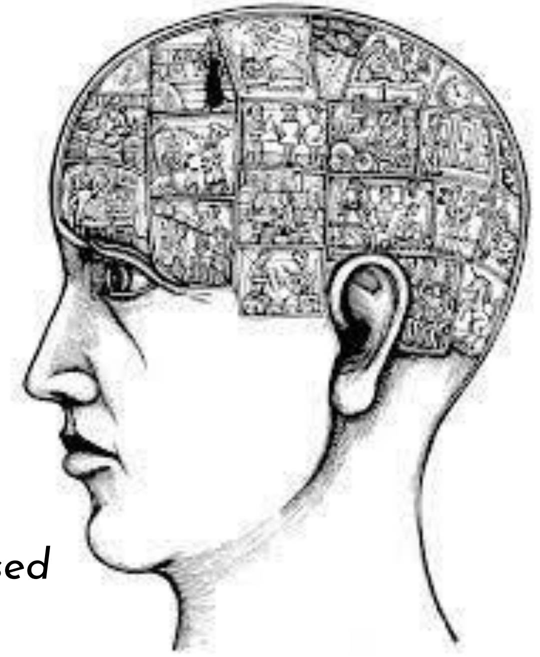

**Oft discussed as the most important feature of modularity, and sometimes discussed in terms of “(im)penetrability” rather than “encapsulation” (see Pylyshyn, 1999)*

Informational encapsulation*

!!

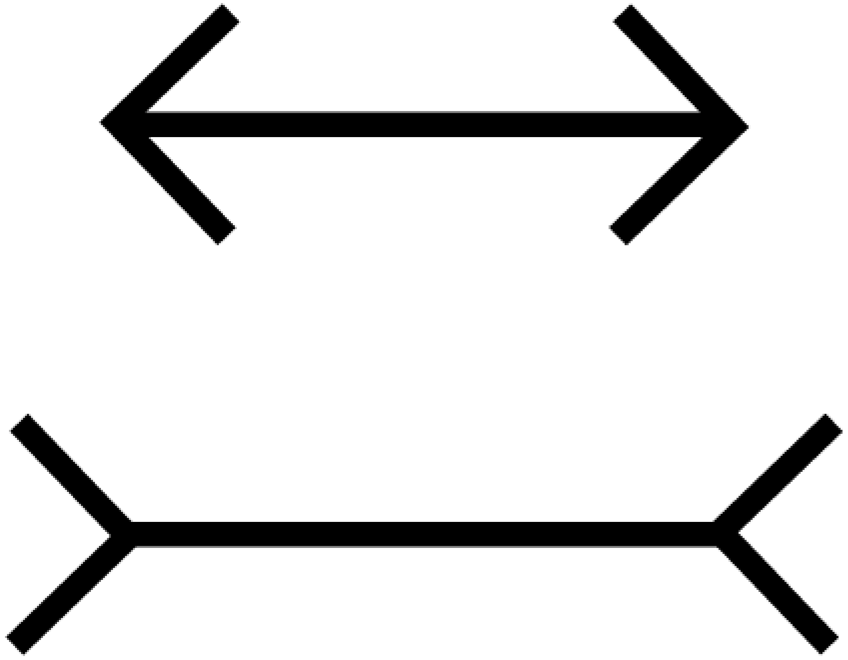

"**Cognitive impenetrability** of visual perception!"

– Pylyshyn, 1999

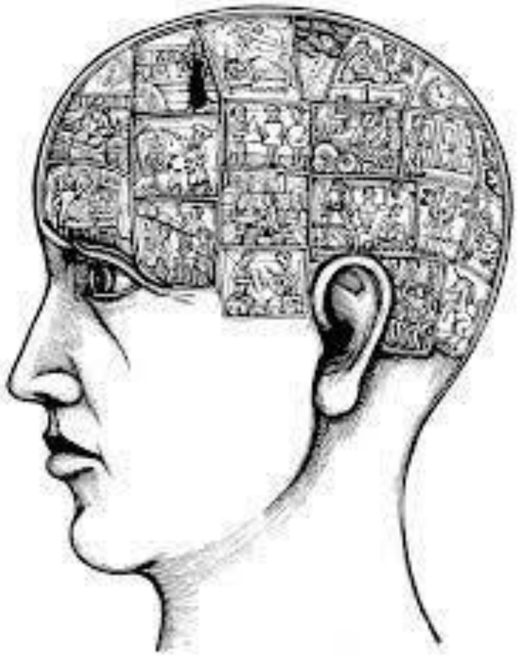

BEHAVIORAL AND BRAIN SCIENCES (2016), Page 1 of 77
doi:10.1017/S0140525X15000965, e229

Cognition does not affect perception: Evaluating the evidence for "top-down" effects

Chaz Firestone
Department of Psychology, Yale University, New Haven, CT 06520-8205
chaz.firestone@yale.edu

Brian J. Scholl
Department of Psychology, Yale University, New Haven, CT 06520-8205
brian.scholl@yale.edu

Abstract: What determines what we see? In contrast to the traditional "modular" understanding of perception, according to which visual processing is encapsulated from higher-level cognition, a tidal wave of recent research alleges that states such as beliefs, desires, emotions, motivations, intentions, and linguistic representations exert direct, top-down influences on what we see. There is a growing consensus that such effects are ubiquitous, and that the distinction between perception and cognition may itself be unsustainable. We argue otherwise: None of these hundreds of studies – either individually or collectively – provides compelling evidence for true top-down effects on perception, or "cognitive penetrability." In particular, and despite their variety, we suggest that these studies all fall prey to only a handful of pitfalls. And whereas abstract theoretical challenges have failed to resolve this debate in the past, our presentation of these pitfalls is empirically anchored: In each case, we show not only how certain studies *could* be susceptible to the pitfall (in principle), but also how several alleged top-down effects *actually are* explained by the pitfall (in practice). Moreover, these pitfalls are perfectly general, with each applying to dozens of other top-down effects. We conclude by extracting the lessons provided by these pitfalls into a checklist that future work could use to convincingly demonstrate top-down effects on visual perception. The discovery of substantive top-down effects of cognition on perception would revolutionize our understanding of how the mind is organized; but without addressing these pitfalls, no such empirical report will license such exciting conclusions.


BEHAVIORAL AND BRAIN SCIENCES (2016), Page 1 of 77
doi:10.1017/S0140525X15000965, e229

Cognition does not affect perception: Evaluating the evidence for "top-down" effects

Chaz Firestone
Department of Psychology, Yale University, New Haven, CT 06520-8205
chaz.firestone@yale.edu

Brian J. Scholl
Department of Psychology, Yale University, New Haven, CT 06520-8205
brian.scholl@yale.edu

Abstract: What determines what we see? In contrast to the traditional "modular" understanding of perception, according to which visual processing is encapsulated from higher-level cognition, a tidal wave of recent research alleges that states such as beliefs, desires, emotions, motivations, intentions, and linguistic representations exert direct, top-down influences on what we see. There is a growing consensus that such effects are ubiquitous, and that the distinction between perception and cognition may itself be unsustainable. We argue otherwise: None of these hundreds of studies – either individually or collectively – provides compelling evidence for true top-down effects on perception, or "cognitive penetrability." In particular, and despite their variety, we suggest that these studies all fall prey to only a handful of pitfalls. And whereas abstract theoretical challenges have failed to resolve this debate in the past, our presentation of these pitfalls is empirically anchored: In each case, we show not only how certain studies *could* be susceptible to the pitfall (in principle), but also how several alleged top-down effects *actually are* explained by the pitfall (in practice). Moreover, these pitfalls are perfectly general, with each applying to dozens of other top-down effects. We conclude by extracting the lessons provided by these pitfalls into a checklist that future work could use to convincingly demonstrate top-down effects on visual perception. The discovery of substantive top-down effects of cognition on perception would revolutionize our understanding of how the mind is organized; but without addressing these pitfalls, no such empirical report will license such exciting conclusions.

A case study of the theoretical value modularity.

perception vs. cognition

seeing vs. *thinking*

Six pitfalls:

!!

- An overly confirmatory research strategy
- *Perception vs. judgment*
- *Demand and response bias*
- *Low-level differences*
- *Peripheral attentional effects*
- *Memory and recognition*

Food for thought:

- Is there a clear divide between perception / cognition?
- Are there any counterexamples to modularity?
- Are there other clear *pieces of mind*?
- Are other parts of the mind modular? Is any part?
- Hallucination??


ILLUSIONS OF MEMORY
Sami Ryan Yousif

Timothy Bates @timothycbates

Another one bites the dust! So exciting to see this since @JoHenrich included the idea that the Müller-Leyer illusion is dependent on culture in the incredibly influential (and mostly lost in translation) WEIRD thesis. I felt it was unlikely. Looked up the data and it was all ancient, tiny, weak studies. But how to test it?.... Well, @DorsaAmir @chazfirestone just did the hard work and Bingo: Turns out it was yet another "important if true", "*Too big to test (TM)*" Shibboleth. Nice work!

> **Dorsa Amir** @DorsaAmir · Jan 25
>
> Does the culture you grow up in shape the way you see the world? In a new Psych Review paper, @chazfirestone & I tackle this centuries-old question using the Müller-Lyer illusion as a case study. Come think through one of history's mysteries with us🧵 (1/13):

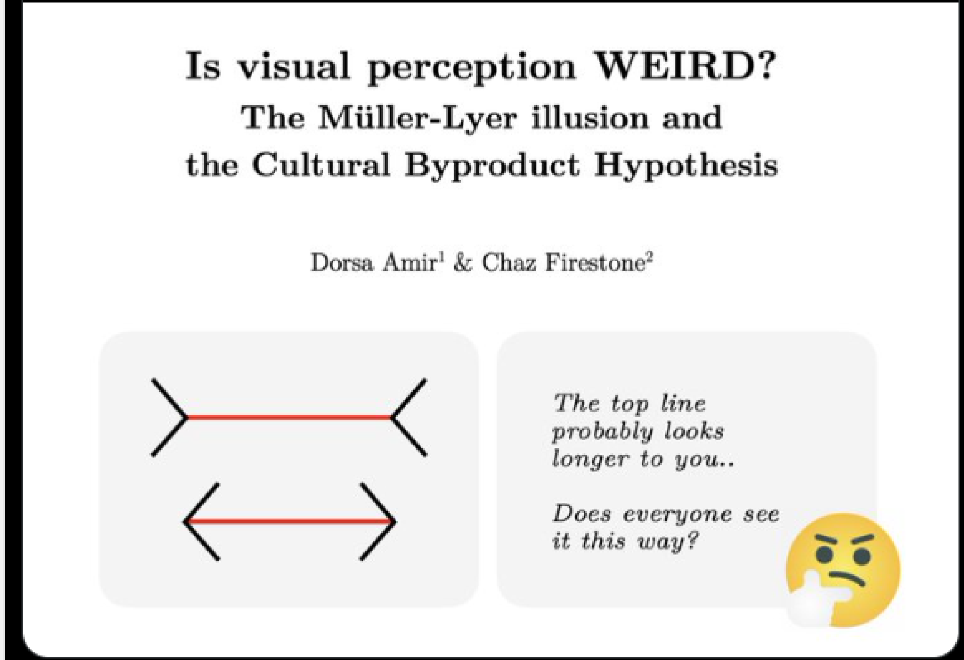

5:48 PM · Jan 26, 2025 · **54.9K** Views

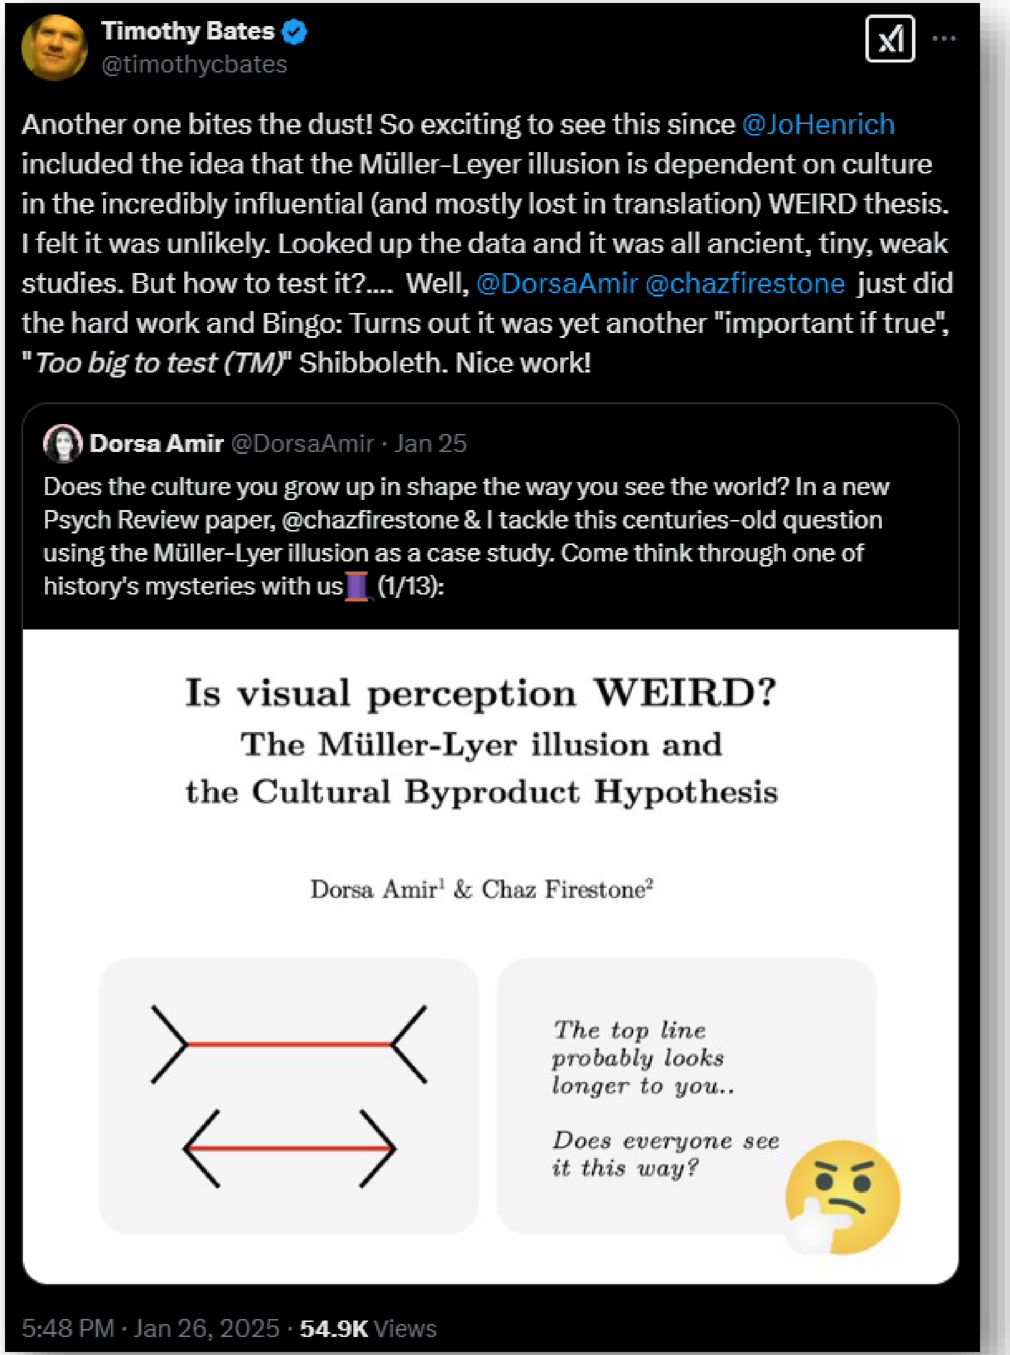

Timothy Bates @timothycbates

Another one bites the dust! So exciting to see this since @JoHenrich included the idea that the Müller-Leyer illusion is dependent on culture in the incredibly influential (and mostly lost in translation) WEIRD thesis. I felt it was unlikely. Looked up the data and it was all ancient, tiny, weak studies. But how to test it?.... Well, @DorsaAmir @chazfirestone just did the hard work and Bingo: Turns out it was yet another "important if true", "*Too big to test (TM)*" Shibboleth. Nice work!

Dorsa Amir @DorsaAmir · Jan 25

Does the culture you grow up in shape the way you see the world? In a new Psych Review paper, @chazfirestone & I tackle this centuries-old question using the Müller-Lyer illusion as a case study. Come think through one of history's mysteries with us🧵 (1/13):

5:48 PM · Jan 26, 2025 · **54.9K** Views

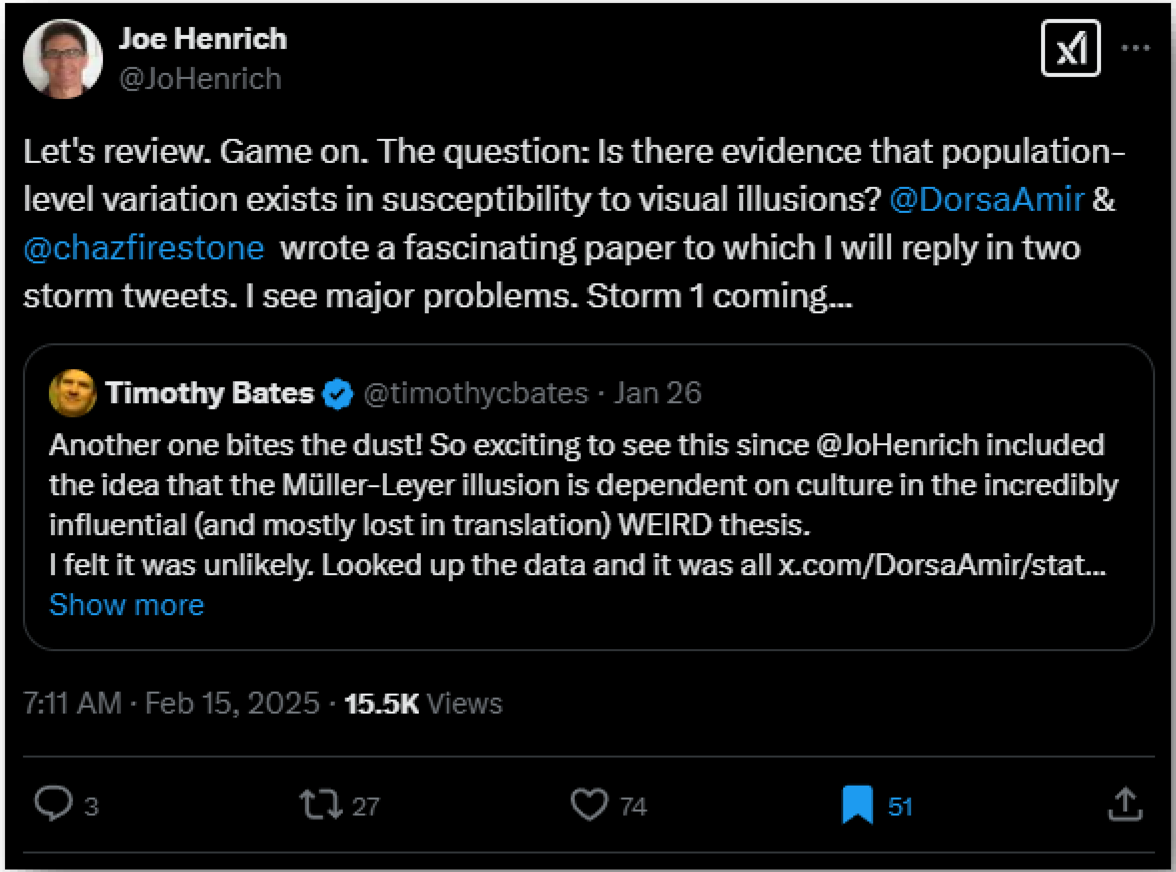

Joe Henrich @JoHenrich

Let's review. Game on. The question: Is there evidence that population-level variation exists in susceptibility to visual illusions? @DorsaAmir & @chazfirestone wrote a fascinating paper to which I will reply in two storm tweets. I see major problems. Storm 1 coming...

Timothy Bates @timothycbates · Jan 26

Another one bites the dust! So exciting to see this since @JoHenrich included the idea that the Müller-Leyer illusion is dependent on culture in the incredibly influential (and mostly lost in translation) WEIRD thesis.
I felt it was unlikely. Looked up the data and it was all x.com/DorsaAmir/stat...
Show more

7:11 AM · Feb 15, 2025 · **15.5K** Views

3 · 27 · 74 · 51


Joe Henrich
@JoHenrich
Let's review. Game on. The question: Is there evidence that population-level variation exists in susceptibility to visual illusions? @DorsaAmir & @chazfirestone wrote a fascinating paper to which I will reply in two storm tweets. I see major problems. Storm 1 coming...
Timothy Bates @timothycbates · Jan 26
Another one bites the dust! So exciting to see this since @JoHenrich included the idea that the Müller-Leyer illusion is dependent on culture in the incredibly influential (and mostly lost in translation) WEIRD thesis.
I felt it was unlikely. Looked up the data and it was all x.com/DorsaAmir/stat...
Show more
7:11 AM · Feb 15, 2025 · 15.5K Views
3
27
74
51

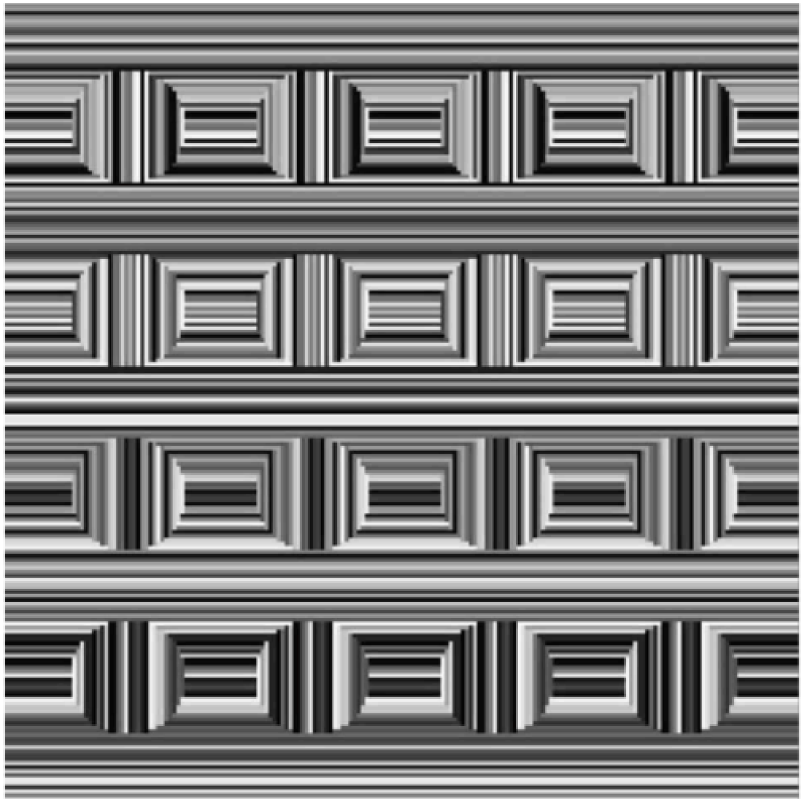


Joe Henrich
@JoHenrich
Let's review. Game on. The question: Is there evidence that population-level variation exists in susceptibility to visual illusions? @DorsaAmir & @chazfirestone wrote a fascinating paper to which I will reply in two storm tweets. I see major problems. Storm 1 coming...
Timothy Bates @timothycbates · Jan 26
Another one bites the dust! So exciting to see this since @JoHenrich included the idea that the Müller-Leyer illusion is dependent on culture in the incredibly influential (and mostly lost in translation) WEIRD thesis.
I felt it was unlikely. Looked up the data and it was all x.com/DorsaAmir/stat...
Show more
7:11 AM · Feb 15, 2025 · 15.5K Views
3
27
74
51

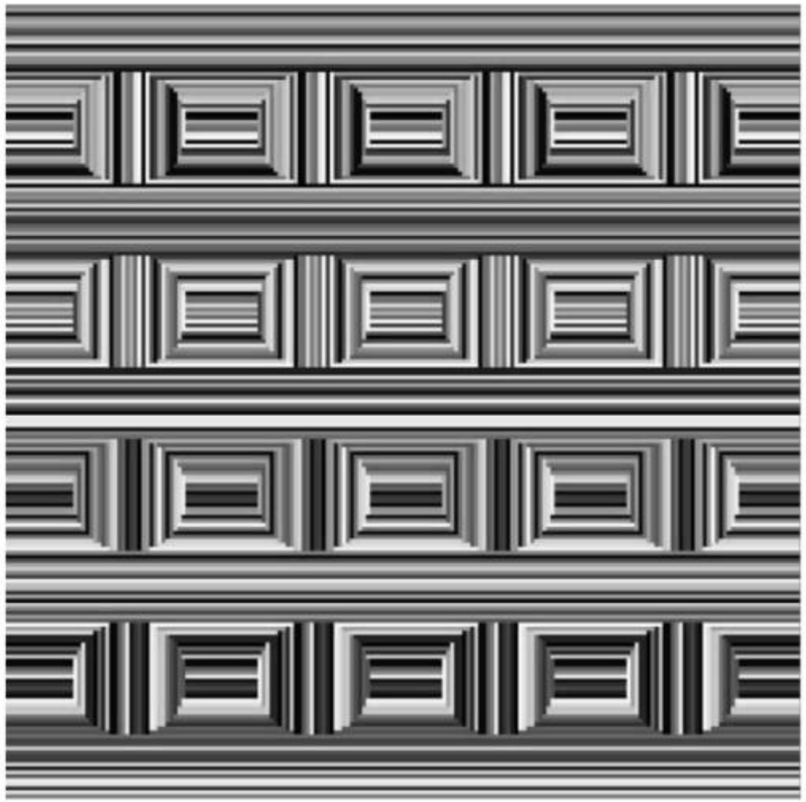

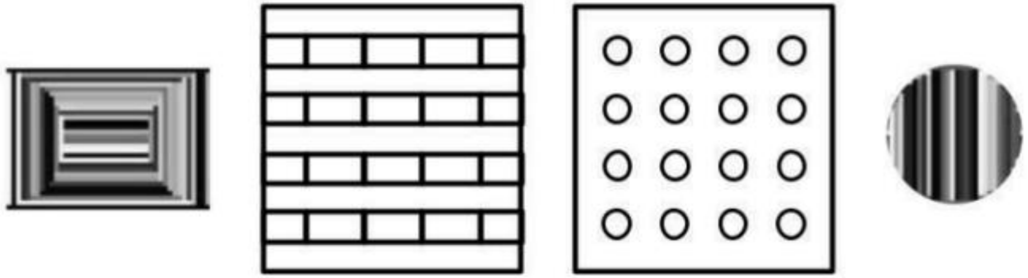

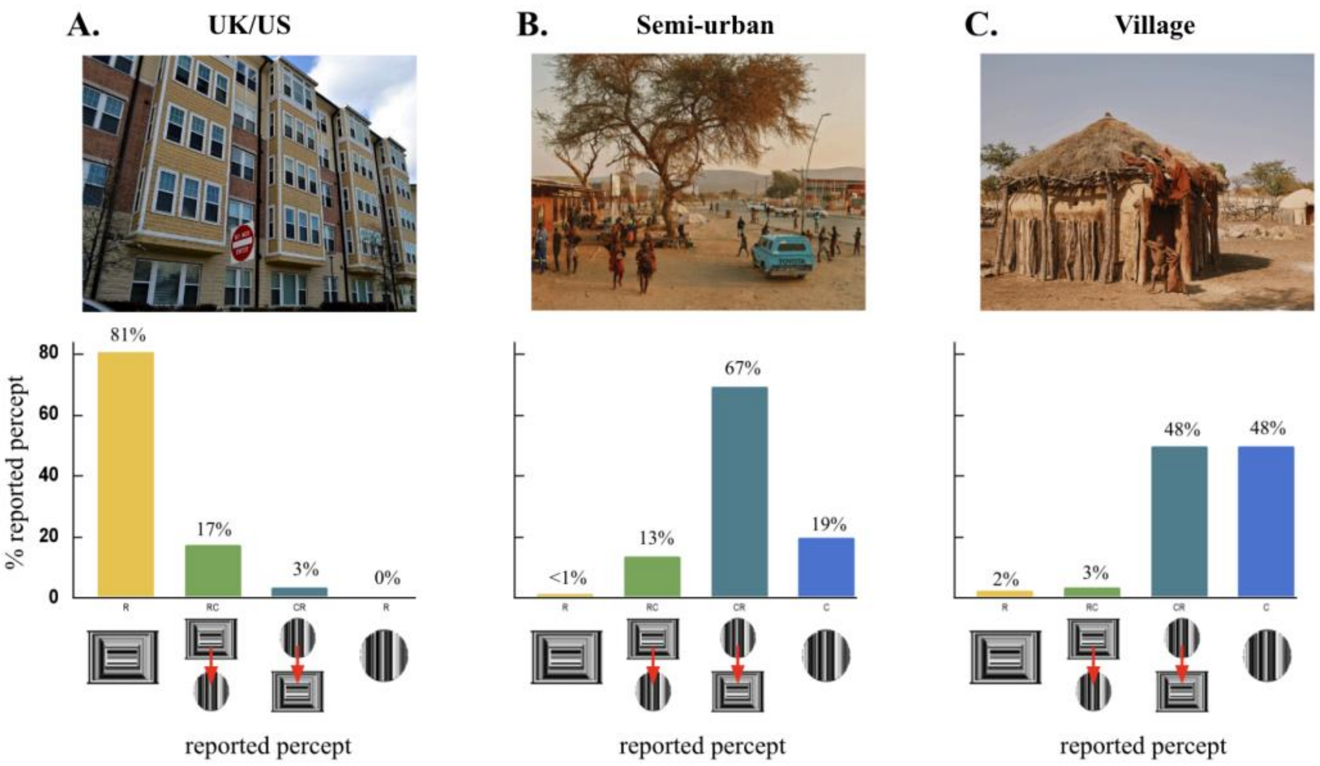
A.
UK/US
B.
Semi-urban
C.
Village
% reported percept
80
60
40
20
0
81%
17%
3%
0%
<1%
13%
67%
19%
2%
3%
48%
48%
R
RC
CR
C
reported percept
reported percept
reported percept

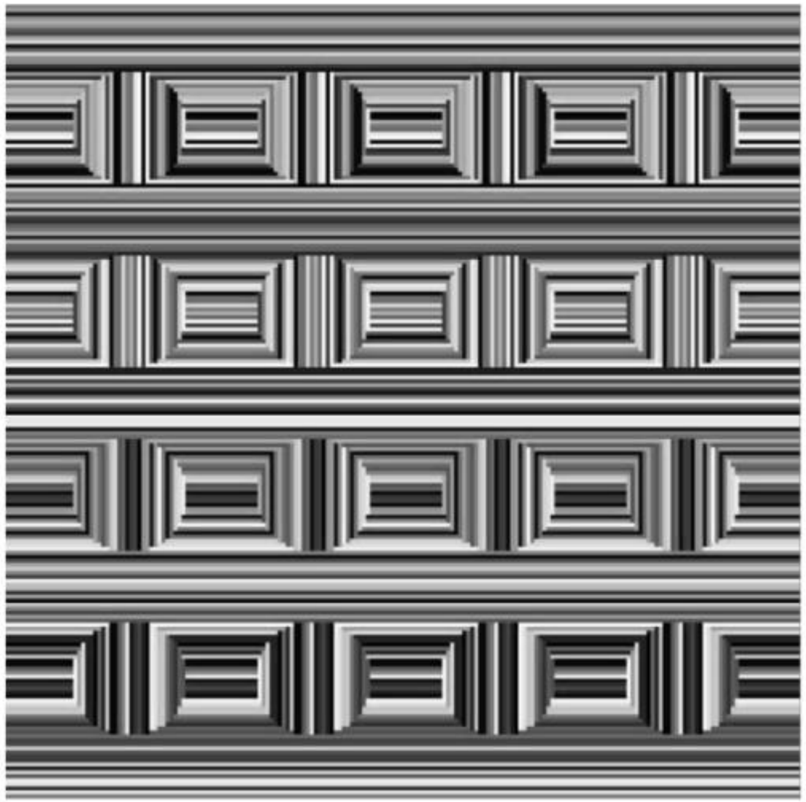

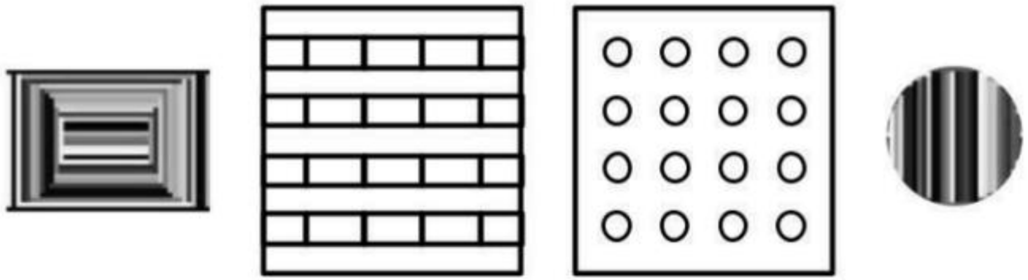

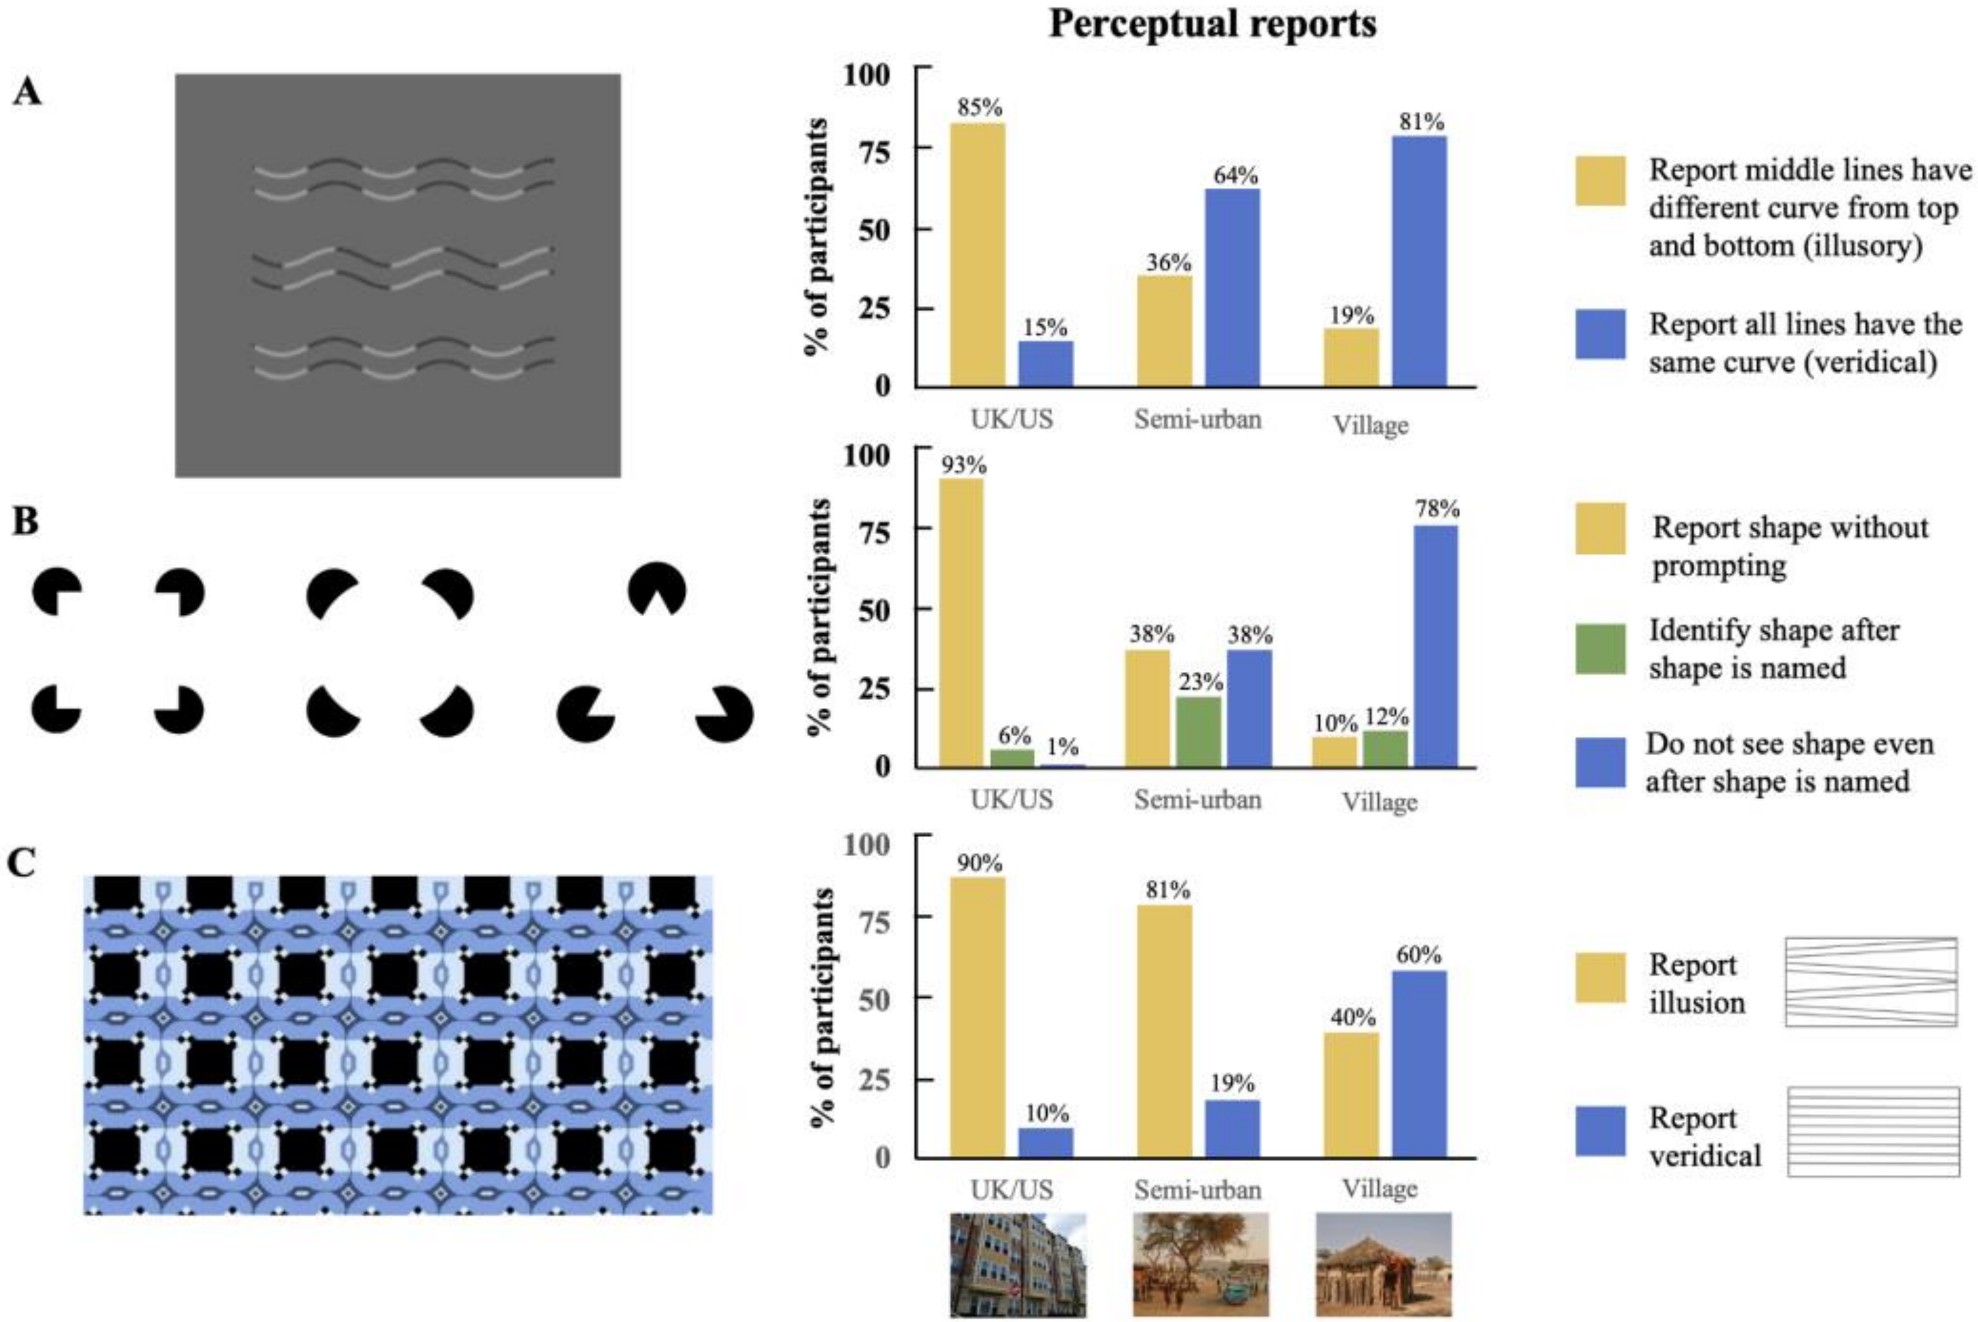

Perceptual reports
A
B
C
% of participants
100
75
50
25
0
85%
15%
36%
64%
19%
81%
UK/US
Semi-urban
Village
Report middle lines have different curve from top and bottom (illusory)
Report all lines have the same curve (veridical)
93%
6%
1%
38%
23%
38%
10%
12%
78%
Report shape without prompting
Identify shape after shape is named
Do not see shape even after shape is named
90%
10%
81%
19%
40%
60%
Report illusion
Report veridical


ILLUSIONS OF MEMORY


WHAT'S YOUR
EARLIEST MEMORY?


SUPER
MARIO WORLD
Nintendo
Explore nine worlds and 96 levels of non-stop action!
SUPER NINTENDO
ENTERTAINMENT SYSTEM


00
000
05
A

Psychonomic Bulletin & Review (2021) 28:928–936
https://doi.org/10.3758/s13423-020-01829-1

BRIEF REPORT

Hi-def memories of lo-def scenes

Jose Rivera-Aparicio[1] · Qian Yu[1] · Chaz Firestone[1]

Accepted: 6 October 2020 / Published online: 14 January 2021

Abstract
The study of visual memory is typically concerned with an image's *content:* How well, and with what precision, we can recall which objects, people, or features we have seen in the past. But images also vary in their *quality*: The same object or scene may appear in an image that is sharp and highly resolved, or it may appear in an image that is blurry and faded. How do we remember those properties? Here six experiments demonstrate a new phenomenon of "vividness extension": a tendency to (mis)remember images as though they are "enhanced" versions of themselves — that is, sharper and higher quality than they actually appeared at the time of encoding. Subjects briefly saw images of scenes that varied in how blurry they were, and then adjusted a new image to be as blurry as the original. Unlike an old photograph that fades and blurs, subjects misremembered scenes as more vivid (i.e., less blurry) than those scenes had actually appeared moments earlier. Follow-up experiments extended this phenomenon to saturation and pixelation – with subjects recalling scenes as more colorful and resolved – and ruled out various forms of response bias. We suggest that memory misrepresents the quality of what we have seen, such that the world is remembered as more vivid than it is.

Keywords Scene perception · Memory · Vividness · Boundary extension

Introduction

What's in a picture? The images we see have *content*, such as objects, places, events, and people; a photograph, for example, might show two friends strolling on a beach. However, images vary not only in their content, but also in their *quality*: A photograph of the very same beach-walk may be sharp, resolved, and saturated; or it could be blurry, grainy, or faded – and this too is a property of images that we may see and encode. Research on visual memory is typically concerned with memory for the former kinds of details — the capacity and precision of our ability to recall contentful information such as which people, objects, events, places, colors, or shapes were present in an image. But what about the latter kinds of details? How do we remember the quality of the images we see?

Suppose, for example, that you see a blurry photograph, grainy television broadcast, or faded piece of artwork, and you attempt to recall not what the images were about, but instead how blurry, grainy, or faded they were, regardless of their content. Is memory for such properties accurate? Or might such memories show biases, misrepresenting the quality of the images we see? Here, we explore this question by investigating memory for image quality.

Are memories better or worse than the real thing?

What kind of bias might there be for memories of image quality? On one hand, memories decay (in that stored information becomes harder to retrieve over time), and this might produce a corresponding "decay" in the quality attributed to those features. In that case, as memories of a photograph or broadcast fade into the past, we might remember the images themselves as having *been* grainier and more faded. Indeed, it has recently been suggested that memories "literally fade," such that images may be estimated as less bright or salient when they are recalled from memory than when they are viewed online (Cooper, Kensinger, & Ritchey, 2019). This prediction not only seems subjectively intuitive, but it could also naturally emerge from sensory recruitment models of working memory, on which such memories are essentially reactivations of sensory states (Harrison & Tong, 2009;

✉ Chaz Firestone
chaz@jhu.edu

[1] Department of Psychological and Brain Sciences, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA

Psychonomic Bulletin & Review (2021) 28:928–936
https://doi.org/10.3758/s13423-020-01829-1

BRIEF REPORT

Hi-def memories of lo-def scenes

Jose Rivera-Aparicio[1] · Qian Yu[1] · Chaz Firestone[1]

Accepted: 6 October 2020 / Published online: 14 January 2021

Abstract

The study of visual memory is typically concerned with an image's *content:* How well, and with what precision, we can recall which objects, people, or features we have seen in the past. But images also vary in their *quality*: The same object or scene may appear in an image that is sharp and highly resolved, or it may appear in an image that is blurry and faded. How do we remember those properties? Here six experiments demonstrate a new phenomenon of "vividness extension": a tendency to (mis)remember images as though they are "enhanced" versions of themselves — that is, sharper and higher quality than they actually appeared at the time of encoding. Subjects briefly saw images of scenes that varied in how blurry they were, and then adjusted a new image to be as blurry as the original. Unlike an old photograph that fades and blurs, subjects misremembered scenes as more vivid (i.e., less blurry) than those scenes had actually appeared moments earlier. Follow-up experiments extended this phenomenon to saturation and pixelation – with subjects recalling scenes as more colorful and resolved – and ruled out various forms of response bias. We suggest that memory misrepresents the quality of what we have seen, such that the world is remembered as more vivid than it is.

Keywords Scene perception · Memory · Vividness · Boundary extension

Introduction

What's in a picture? The images we see have *content*, such as objects, places, events, and people; a photograph, for example, might show two friends strolling on a beach. However, images vary not only in their content, but also in their *quality*: A photograph of the very same beach-walk may be sharp, resolved, and saturated; or it could be blurry, grainy, or faded – and this too is a property of images that we may see and encode. Research on visual memory is typically concerned with memory for the former kinds of details — the capacity and precision of our ability to recall contentful information such as which people, objects, events, places, colors, or shapes were present in an image. But what about the latter kinds of details? How do we remember the quality of the images we see?

Suppose, for example, that you see a blurry photograph, grainy television broadcast, or faded piece of artwork, and you attempt to recall not what the images were about, but instead how blurry, grainy, or faded they were, regardless of their content. Is memory for such properties accurate? Or might such memories show biases, misrepresenting the quality of the images we see? Here, we explore this question by investigating memory for image quality.

Are memories better or worse than the real thing?

What kind of bias might there be for memories of image quality? On one hand, memories decay (in that stored information becomes harder to retrieve over time), and this might produce a corresponding "decay" in the quality attributed to those features. In that case, as memories of a photograph or broadcast fade into the past, we might remember the images themselves as having *been* grainier and more faded. Indeed, it has recently been suggested that memories "literally fade," such that images may be estimated as less bright or salient when they are recalled from memory than when they are viewed online (Cooper, Kensinger, & Ritchey, 2019). This prediction not only seems subjectively intuitive, but it could also naturally emerge from sensory recruitment models of working memory, on which such memories are essentially reactivations of sensory states (Harrison & Tong, 2009;

✉ Chaz Firestone
chaz@jhu.edu

[1] Department of Psychological and Brain Sciences, Johns Hopkins University, 3400 N Charles St, Baltimore, MD 21218, USA

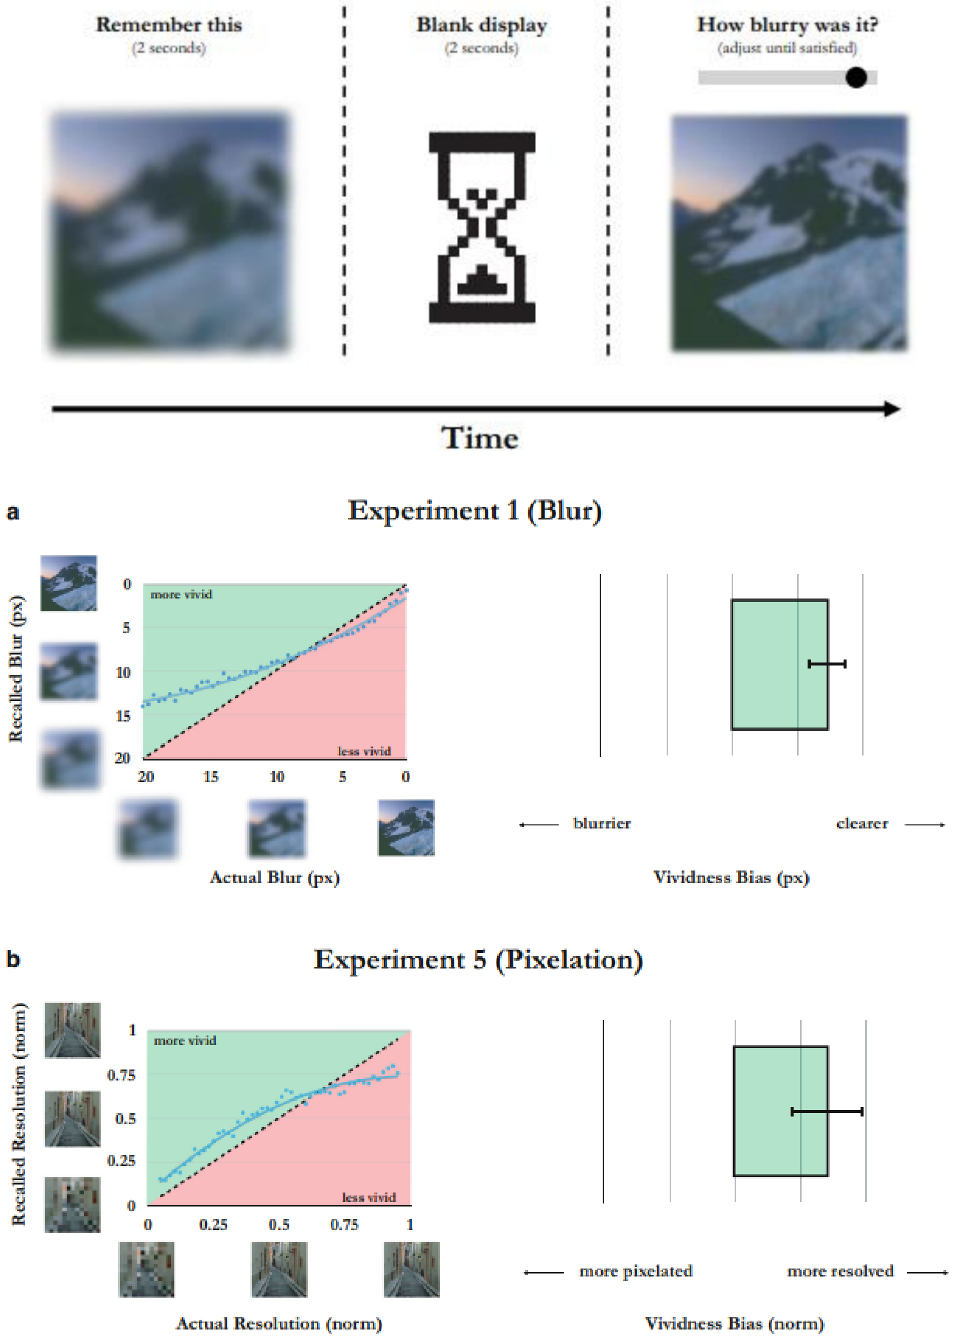

It isn't just our visual memories that are "hi-def"!

!!

"Rosy Retrospection"

“Rosy Retrospection” ‼

Temporal Adjustments in the Evaluation of Events: The “Rosy View”

In a series of three investigations we examined people's anticipation of, actual experiences in, and subsequent recollection of meaningful life events: a trip to Europe, a Thanksgiving vacation, and a 3-week bicycle trip in California. The results of all three studies supported the hypothesis that people's expectations of personal events are more positive than their actual experience during the event itself, and their subsequent recollection of that event is more positive than the actual experience. The “rosy view” phenomenon is associated with an increase in the number of negative thoughts during the event which seem to be caused by distractions, disappointment, and a less positive view of the self. However, these effects are short-lived; within days after the event, people have much more positive evaluations of the event. We discuss alternative interpretations for our findings and implications for group and organizational settings.

Mitchell et al., 1997

Our memories aren't nearly as reliable as we think!

For example: Where was the **mushroom**?

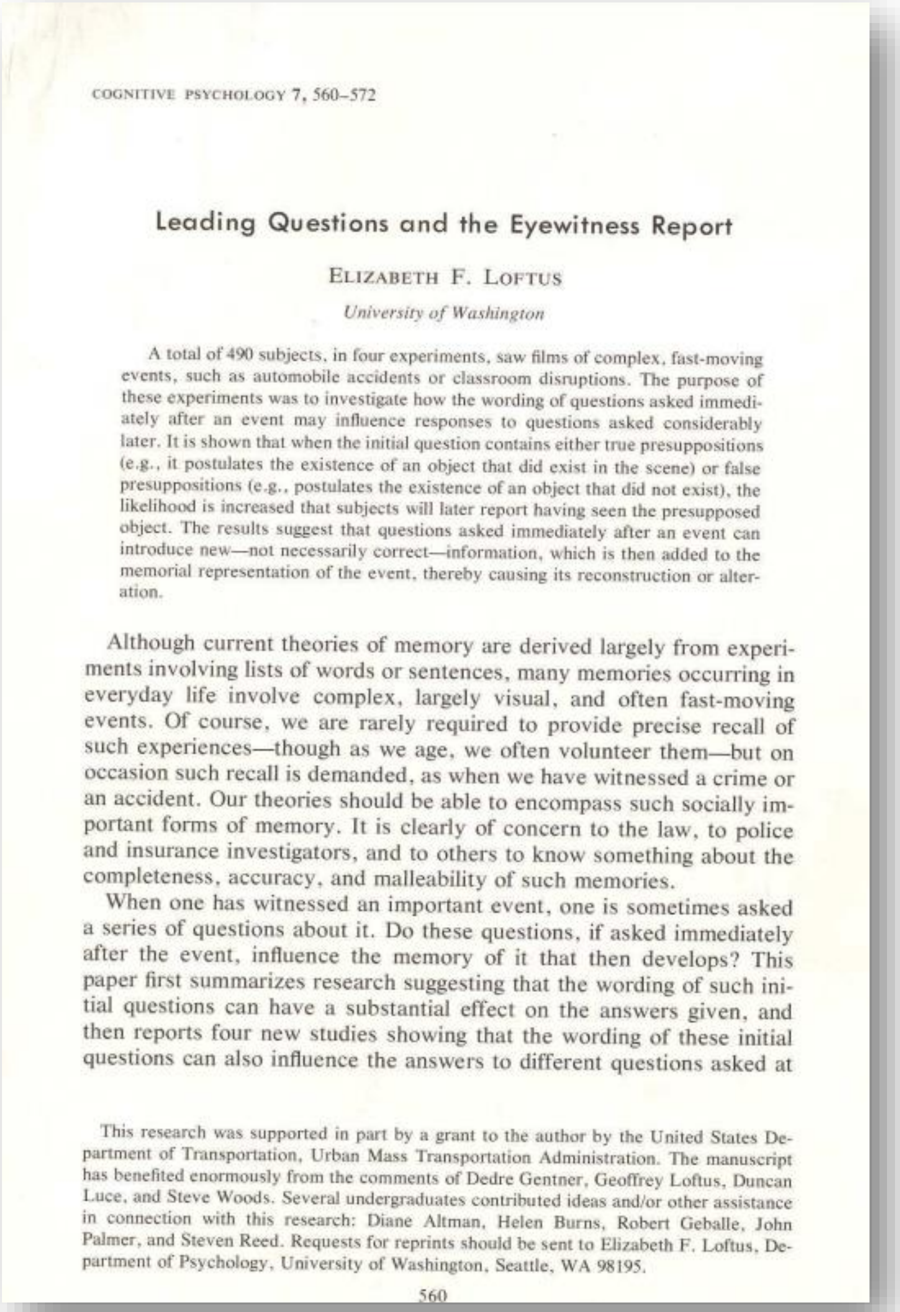
COGNITIVE PSYCHOLOGY 7, 560–572

Leading Questions and the Eyewitness Report

ELIZABETH F. LOFTUS

University of Washington

A total of 490 subjects, in four experiments, saw films of complex, fast-moving events, such as automobile accidents or classroom disruptions. The purpose of these experiments was to investigate how the wording of questions asked immediately after an event may influence responses to questions asked considerably later. It is shown that when the initial question contains either true presuppositions (e.g., it postulates the existence of an object that did exist in the scene) or false presuppositions (e.g., postulates the existence of an object that did not exist), the likelihood is increased that subjects will later report having seen the presupposed object. The results suggest that questions asked immediately after an event can introduce new—not necessarily correct—information, which is then added to the memorial representation of the event, thereby causing its reconstruction or alteration.

Although current theories of memory are derived largely from experiments involving lists of words or sentences, many memories occurring in everyday life involve complex, largely visual, and often fast-moving events. Of course, we are rarely required to provide precise recall of such experiences—though as we age, we often volunteer them—but on occasion such recall is demanded, as when we have witnessed a crime or an accident. Our theories should be able to encompass such socially important forms of memory. It is clearly of concern to the law, to police and insurance investigators, and to others to know something about the completeness, accuracy, and malleability of such memories.

When one has witnessed an important event, one is sometimes asked a series of questions about it. Do these questions, if asked immediately after the event, influence the memory of it that then develops? This paper first summarizes research suggesting that the wording of such initial questions can have a substantial effect on the answers given, and then reports four new studies showing that the wording of these initial questions can also influence the answers to different questions asked at

This research was supported in part by a grant to the author by the United States Department of Transportation, Urban Mass Transportation Administration. The manuscript has benefited enormously from the comments of Dedre Gentner, Geoffrey Loftus, Duncan Luce, and Steve Woods. Several undergraduates contributed ideas and/or other assistance in connection with this research: Diane Altman, Helen Burns, Robert Geballe, John Palmer, and Steven Reed. Requests for reprints should be sent to Elizabeth F. Loftus, Department of Psychology, University of Washington, Seattle, WA 98195.

Where was the mushroom?

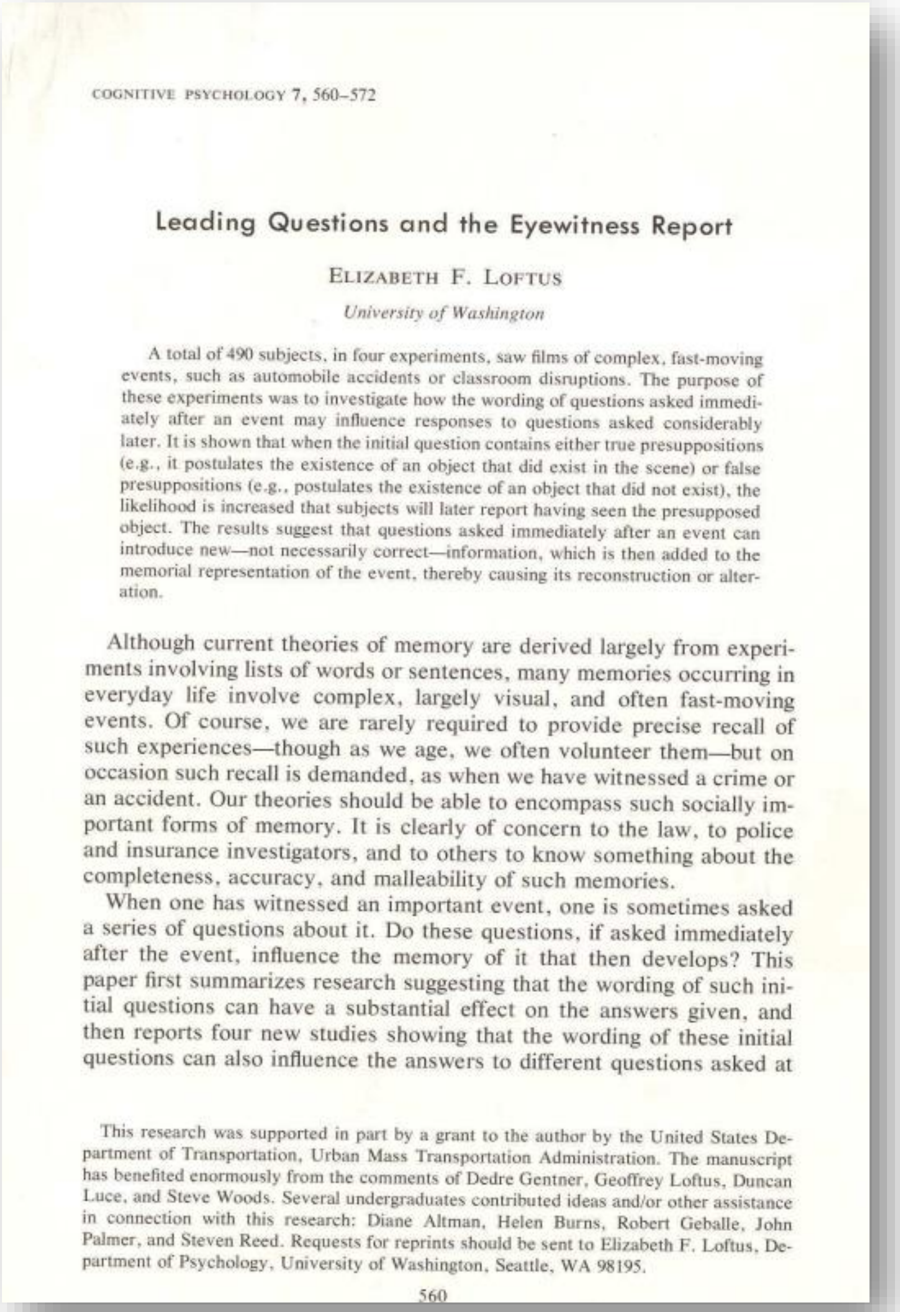
COGNITIVE PSYCHOLOGY 7, 560–572

Leading Questions and the Eyewitness Report

ELIZABETH F. LOFTUS

University of Washington

A total of 490 subjects, in four experiments, saw films of complex, fast-moving events, such as automobile accidents or classroom disruptions. The purpose of these experiments was to investigate how the wording of questions asked immediately after an event may influence responses to questions asked considerably later. It is shown that when the initial question contains either true presuppositions (e.g., it postulates the existence of an object that did exist in the scene) or false presuppositions (e.g., postulates the existence of an object that did not exist), the likelihood is increased that subjects will later report having seen the presupposed object. The results suggest that questions asked immediately after an event can introduce new—not necessarily correct—information, which is then added to the memorial representation of the event, thereby causing its reconstruction or alteration.

Although current theories of memory are derived largely from experiments involving lists of words or sentences, many memories occurring in everyday life involve complex, largely visual, and often fast-moving events. Of course, we are rarely required to provide precise recall of such experiences—though as we age, we often volunteer them—but on occasion such recall is demanded, as when we have witnessed a crime or an accident. Our theories should be able to encompass such socially important forms of memory. It is clearly of concern to the law, to police and insurance investigators, and to others to know something about the completeness, accuracy, and malleability of such memories.

When one has witnessed an important event, one is sometimes asked a series of questions about it. Do these questions, if asked immediately after the event, influence the memory of it that then develops? This paper first summarizes research suggesting that the wording of such initial questions can have a substantial effect on the answers given, and then reports four new studies showing that the wording of these initial questions can also influence the answers to different questions asked at

This research was supported in part by a grant to the author by the United States Department of Transportation, Urban Mass Transportation Administration. The manuscript has benefited enormously from the comments of Dedre Gentner, Geoffrey Loftus, Duncan Luce, and Steve Woods. Several undergraduates contributed ideas and/or other assistance in connection with this research: Diane Altman, Helen Burns, Robert Geballe, John Palmer, and Steven Reed. Requests for reprints should be sent to Elizabeth F. Loftus, Department of Psychology, University of Washington, Seattle, WA 98195.

!!

Where was the mushroom?

vs.

What did you see?

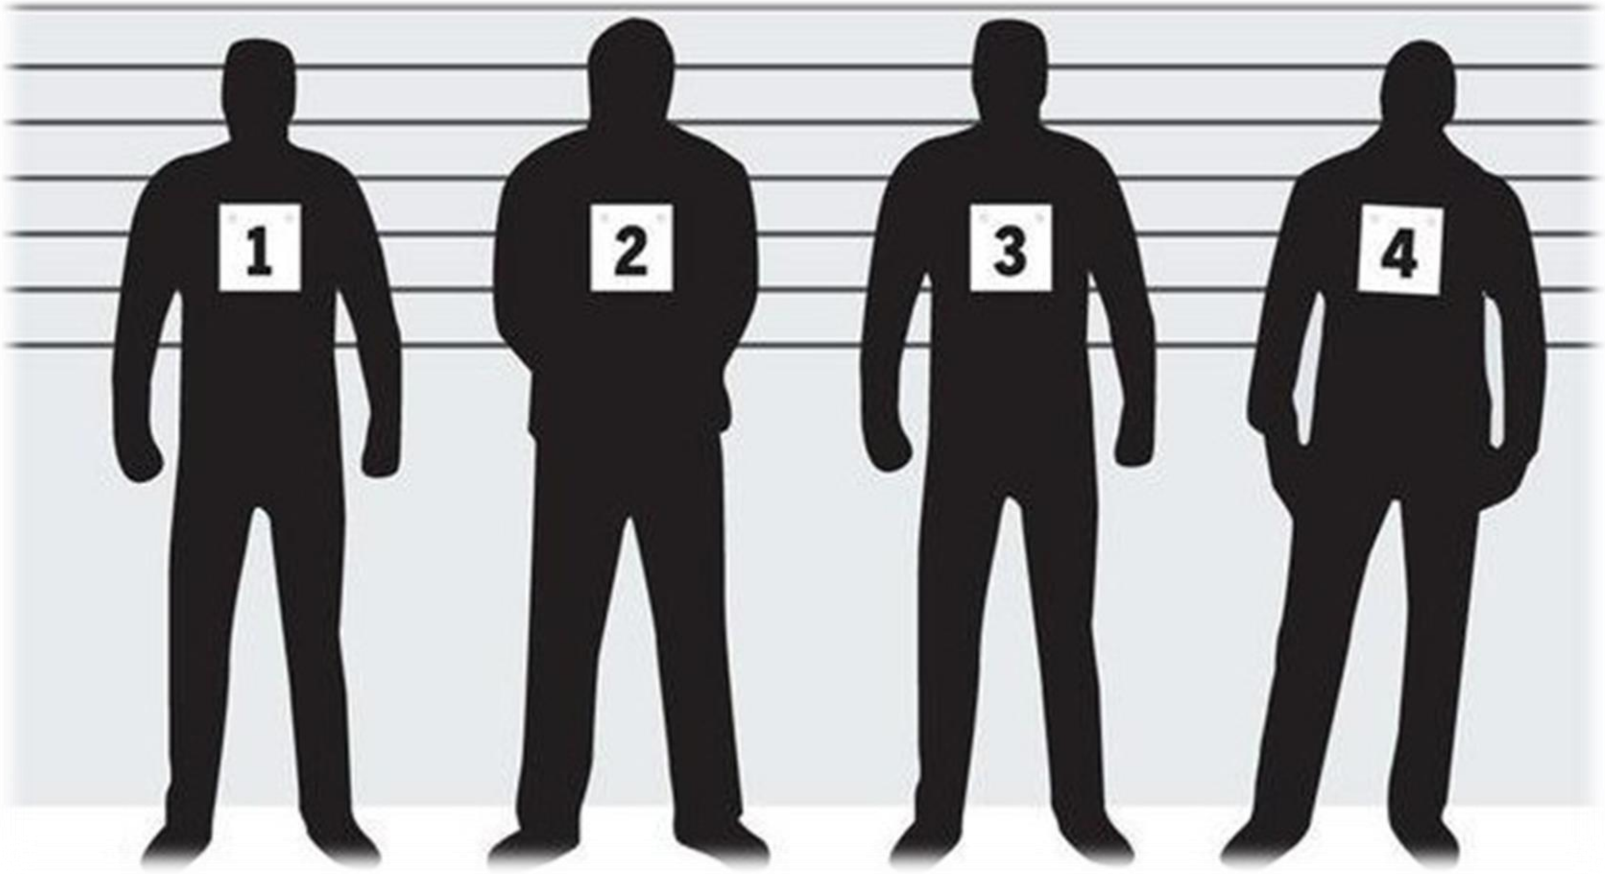

1
2
3
4

Unconscious Transference in Eyewitness Identification

Elizabeth F. Loftus, Ph.D*

I. Introduction

A ticket agent in a railroad station was held up at gunpoint and subsequently recognized a sailor in a lineup as the culprit. The sailor had an ironclad alibi, however, and was eventually released from custody. The ticket agent was later interviewed in an attempt to determine why he had misidentified the sailor; he said that when he saw the sailor in the lineup his face looked familiar. As it happened, the sailor's base was near the railroad station, and on three occasions prior to the robbery he had purchased tickets from this agent. It appears then that the ticket agent mistakenly assumed that the familiarity of the face related back to the robbery, when it undoubtedly related back to the three times the sailor had bought train tickets.[1]

In the ticket agent case, a person seen in one situation (buying tickets) was confused with or recalled as the person seen in a second situation (committing armed robbery). Glanville Williams[2] has termed this phenomenon "unconscious transference". In his discussion of an English murder case which may have resulted in the execution of an innocent man, Williams points out that one of the witnesses who identified the defendant had seen him briefly prior to the crime and may have unconsciously effected a transference.[3]

There are many cases in which one might infer from the statements of facts that a transference had occurred, but cases in which the transference problem has been specifically discussed are virtually nonexistent.[4] Rather than rely on individual cases, we shall examine the psychological literature to confirm the occurrence of unconscious transference.

In a recent experiment by Buckhout[5] 141 students witnessed a

---

* Professor of Psychology, University of Washington. Requests for reprints should be sent to Professor Elizabeth Loftus, Department of Psychology, University of Washington, Seattle, Washington 98195.

1. P. Wall, Eye-Witness Identification in Criminal Cases (1965).
2. G. Williams, The Proof of Guilt (1963).
3. *Id.*
4. Wall, *supra* note 1.
5. Buckhout, *Eyewitness Testimony*, 231 Scientific Am., Dec., 1974, at 23.

Did you ever get lost in the mall?

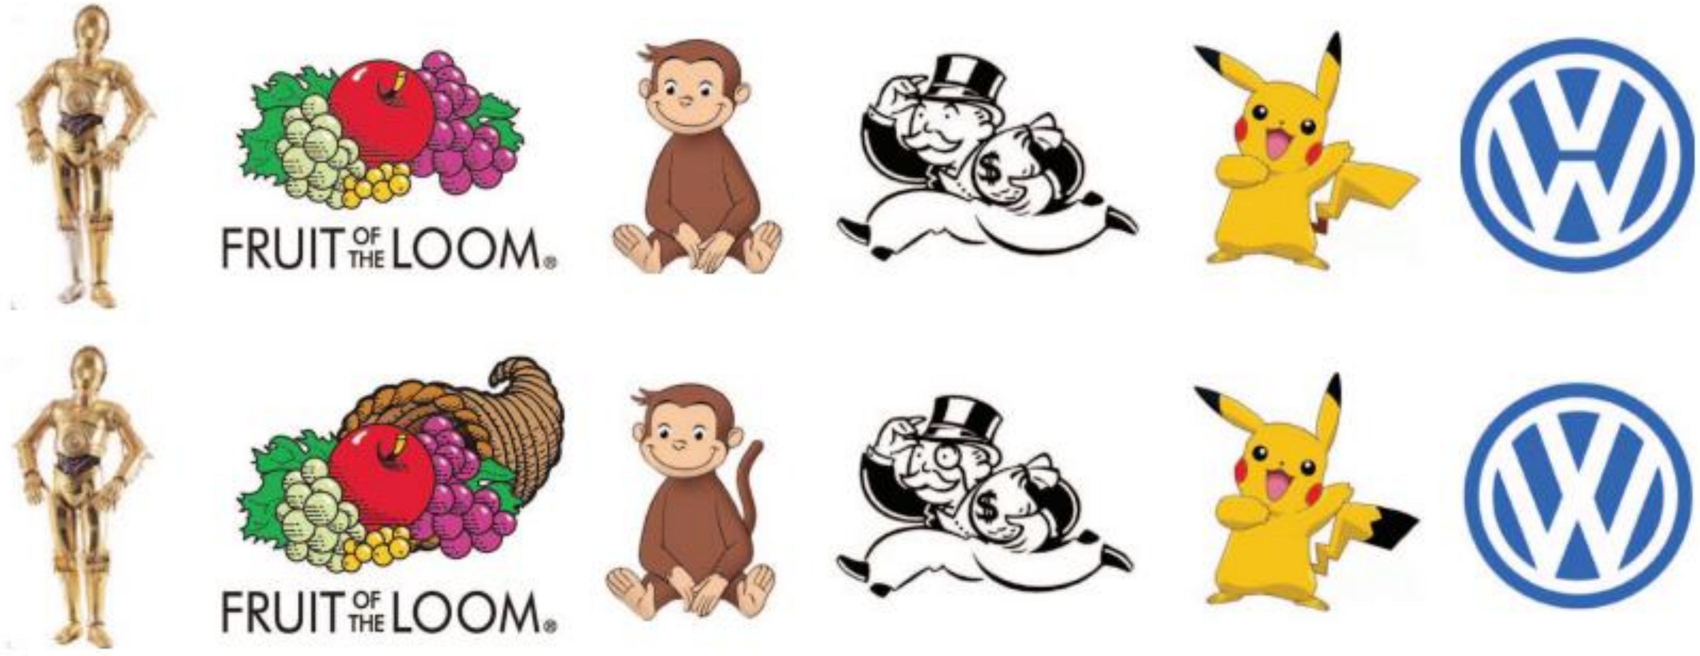
FRUIT OF THE LOOM
FRUIT OF THE LOOM

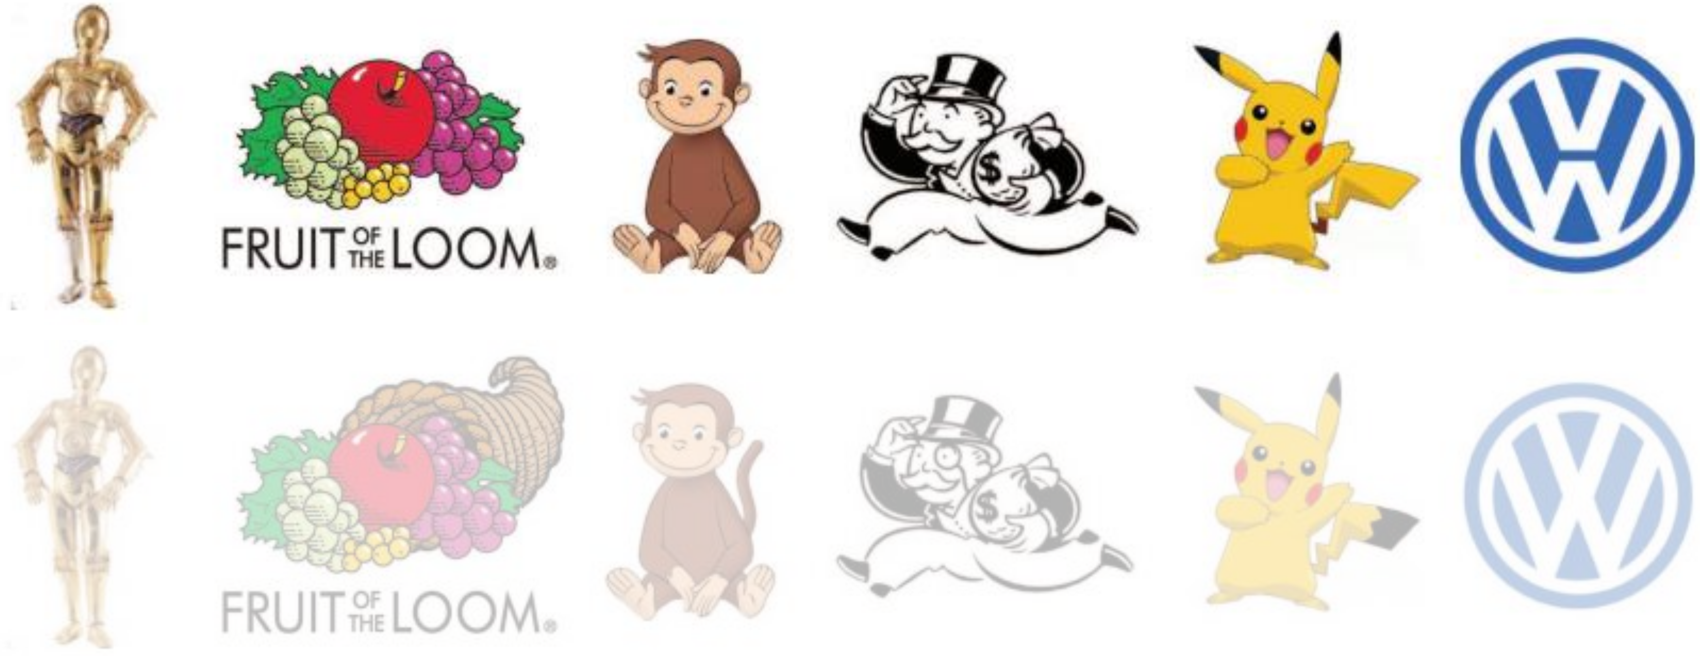
FRUIT OF THE LOOM
FRUIT OF THE LOOM

Research Article

Psychological Science
2022, Vol. 33(12) 1971–1988

DOI: 10.1177/09567976221108944
www.psychologicalscience.org/PS
SAGE

The Visual Mandela Effect as Evidence for Shared and Specific False Memories Across People

Deepasri Prasad and Wilma A. Bainbridge
Department of Psychology, The University of Chicago

Abstract

The Mandela effect is an Internet phenomenon describing shared and consistent false memories for specific icons in popular culture. The visual Mandela effect is a Mandela effect specific to visual icons (e.g., the Monopoly Man is falsely remembered as having a monocle) and has not yet been empirically quantified or tested. In Experiment 1 ($N = 100$ adults), we demonstrated that certain images from popular iconography elicit consistent, specific false memories. In Experiment 2 ($N = 60$ adults), using eye-tracking-like methods, we found no attentional or visual differences that drive this phenomenon. There is no clear difference in the natural visual experience of these images (Experiment 3), and these errors also occur spontaneously during recall (Experiment 4; $N = 50$ adults). These results demonstrate that there are certain images for which people consistently make the same false-memory error, despite the majority of visual experience being the canonical image.

Keywords

visual memory, recognition, visual recall, memory errors, drawing paradigm, open data, open materials

Received 7/7/21; Revision accepted 4/10/22

Popular icons, such as characters or logos, are intentionally designed to be eye catching and memorable. Unless the design changes, people are repeatedly exposed to the same canonical icon (i.e., the official, in-use design), which creates a strong sense of familiarity. However, there are some icons for which many people report strongly remembering a different version, one that is not the canonical icon. Interestingly, this incorrect version is the same across those individuals with this false memory. This phenomenon of specific and consistent visual false memories for certain images in popular culture is called the visual Mandela effect (VME). If such specific shared false memory exists, it suggests commonalities across our experiences of these images or a role of properties intrinsic to these images on false memory.

The term "Mandela effect" was coined by Fiona Broome, a paranormal researcher, to describe her false memory of Nelson Mandela dying in prison in the 1980s (Broome, 2010). She claimed that other people also had the same false memory. The term has since propagated on the Internet to describe instances in which many people share highly specific false memories for names, events, or images. For example, people report having a strong recollection that the Monopoly Man, mascot for the board game Monopoly, wears a monocle. To their surprise, he does not nor has he ever (Fig. 1). Broome and other individuals on the Internet have interpreted this shared false experience as evidence of alternate dimensions. Thus, most of the existing literature on the Mandela effect uses it as an example of a conspiracy theory (Maswood & Rajaram, 2019). However, examining the Mandela effect as a psychological phenomenon could shed light on the nature of memory representations and how false memories form. In this study, we provide a comprehensive examination of the VME, a specific type of Mandela effect in the visual modality, and investigate its consistency across people.

Corresponding Author:
Wilma A. Bainbridge, The University of Chicago, Department of Psychology
Email: wilma@uchicago.edu

Mandela effect:

Vivid, shared false memories.

Synthese (2014) 191:155–185
DOI 10.1007/s11229-013-0247-7

Is memory for remembering? Recollection as a form of episodic hypothetical thinking

Felipe De Brigard

Received: 16 November 2012 / Accepted: 16 January 2013 / Published online: 5 February 2013

Abstract Misremembering is a systematic and ordinary occurrence in our daily lives. Since it is commonly assumed that the function of memory is to remember the past, misremembering is typically thought to happen because our memory system malfunctions. In this paper I argue that not all cases of misremembering are due to failures in our memory system. In particular, I argue that many ordinary cases of misremembering should not be seen as instances of memory's malfunction, but rather as the normal result of a larger cognitive system that performs a different function, and for which remembering is just one operation. Building upon extant psychological and neuroscientific evidence, I offer a picture of memory as an integral part of a larger system that supports not only thinking of what was the case and what potentially could be the case, but also what could have been the case. More precisely, I claim that remembering is a particular operation of a cognitive system that permits the flexible recombination of different components of encoded traces into representations of possible past events that might or might not have occurred, in the service of constructing mental simulations of possible future events.

Keywords Memory · Cognitive function · Remembering · Hypothetical thinking · Core brain network · Episodic future thinking · Episodic counterfactual thinking

So that imagination and memory are but one thing, which for diverse considerations hath diverse names.

Thomas Hobbes, Leviathan 1.2.

F. De Brigard (✉)
Department of Psychology, Harvard University, William James Hall, 860, Cambridge, MA 02139, USA
e-mail: brigard@wjh.harvard.edu

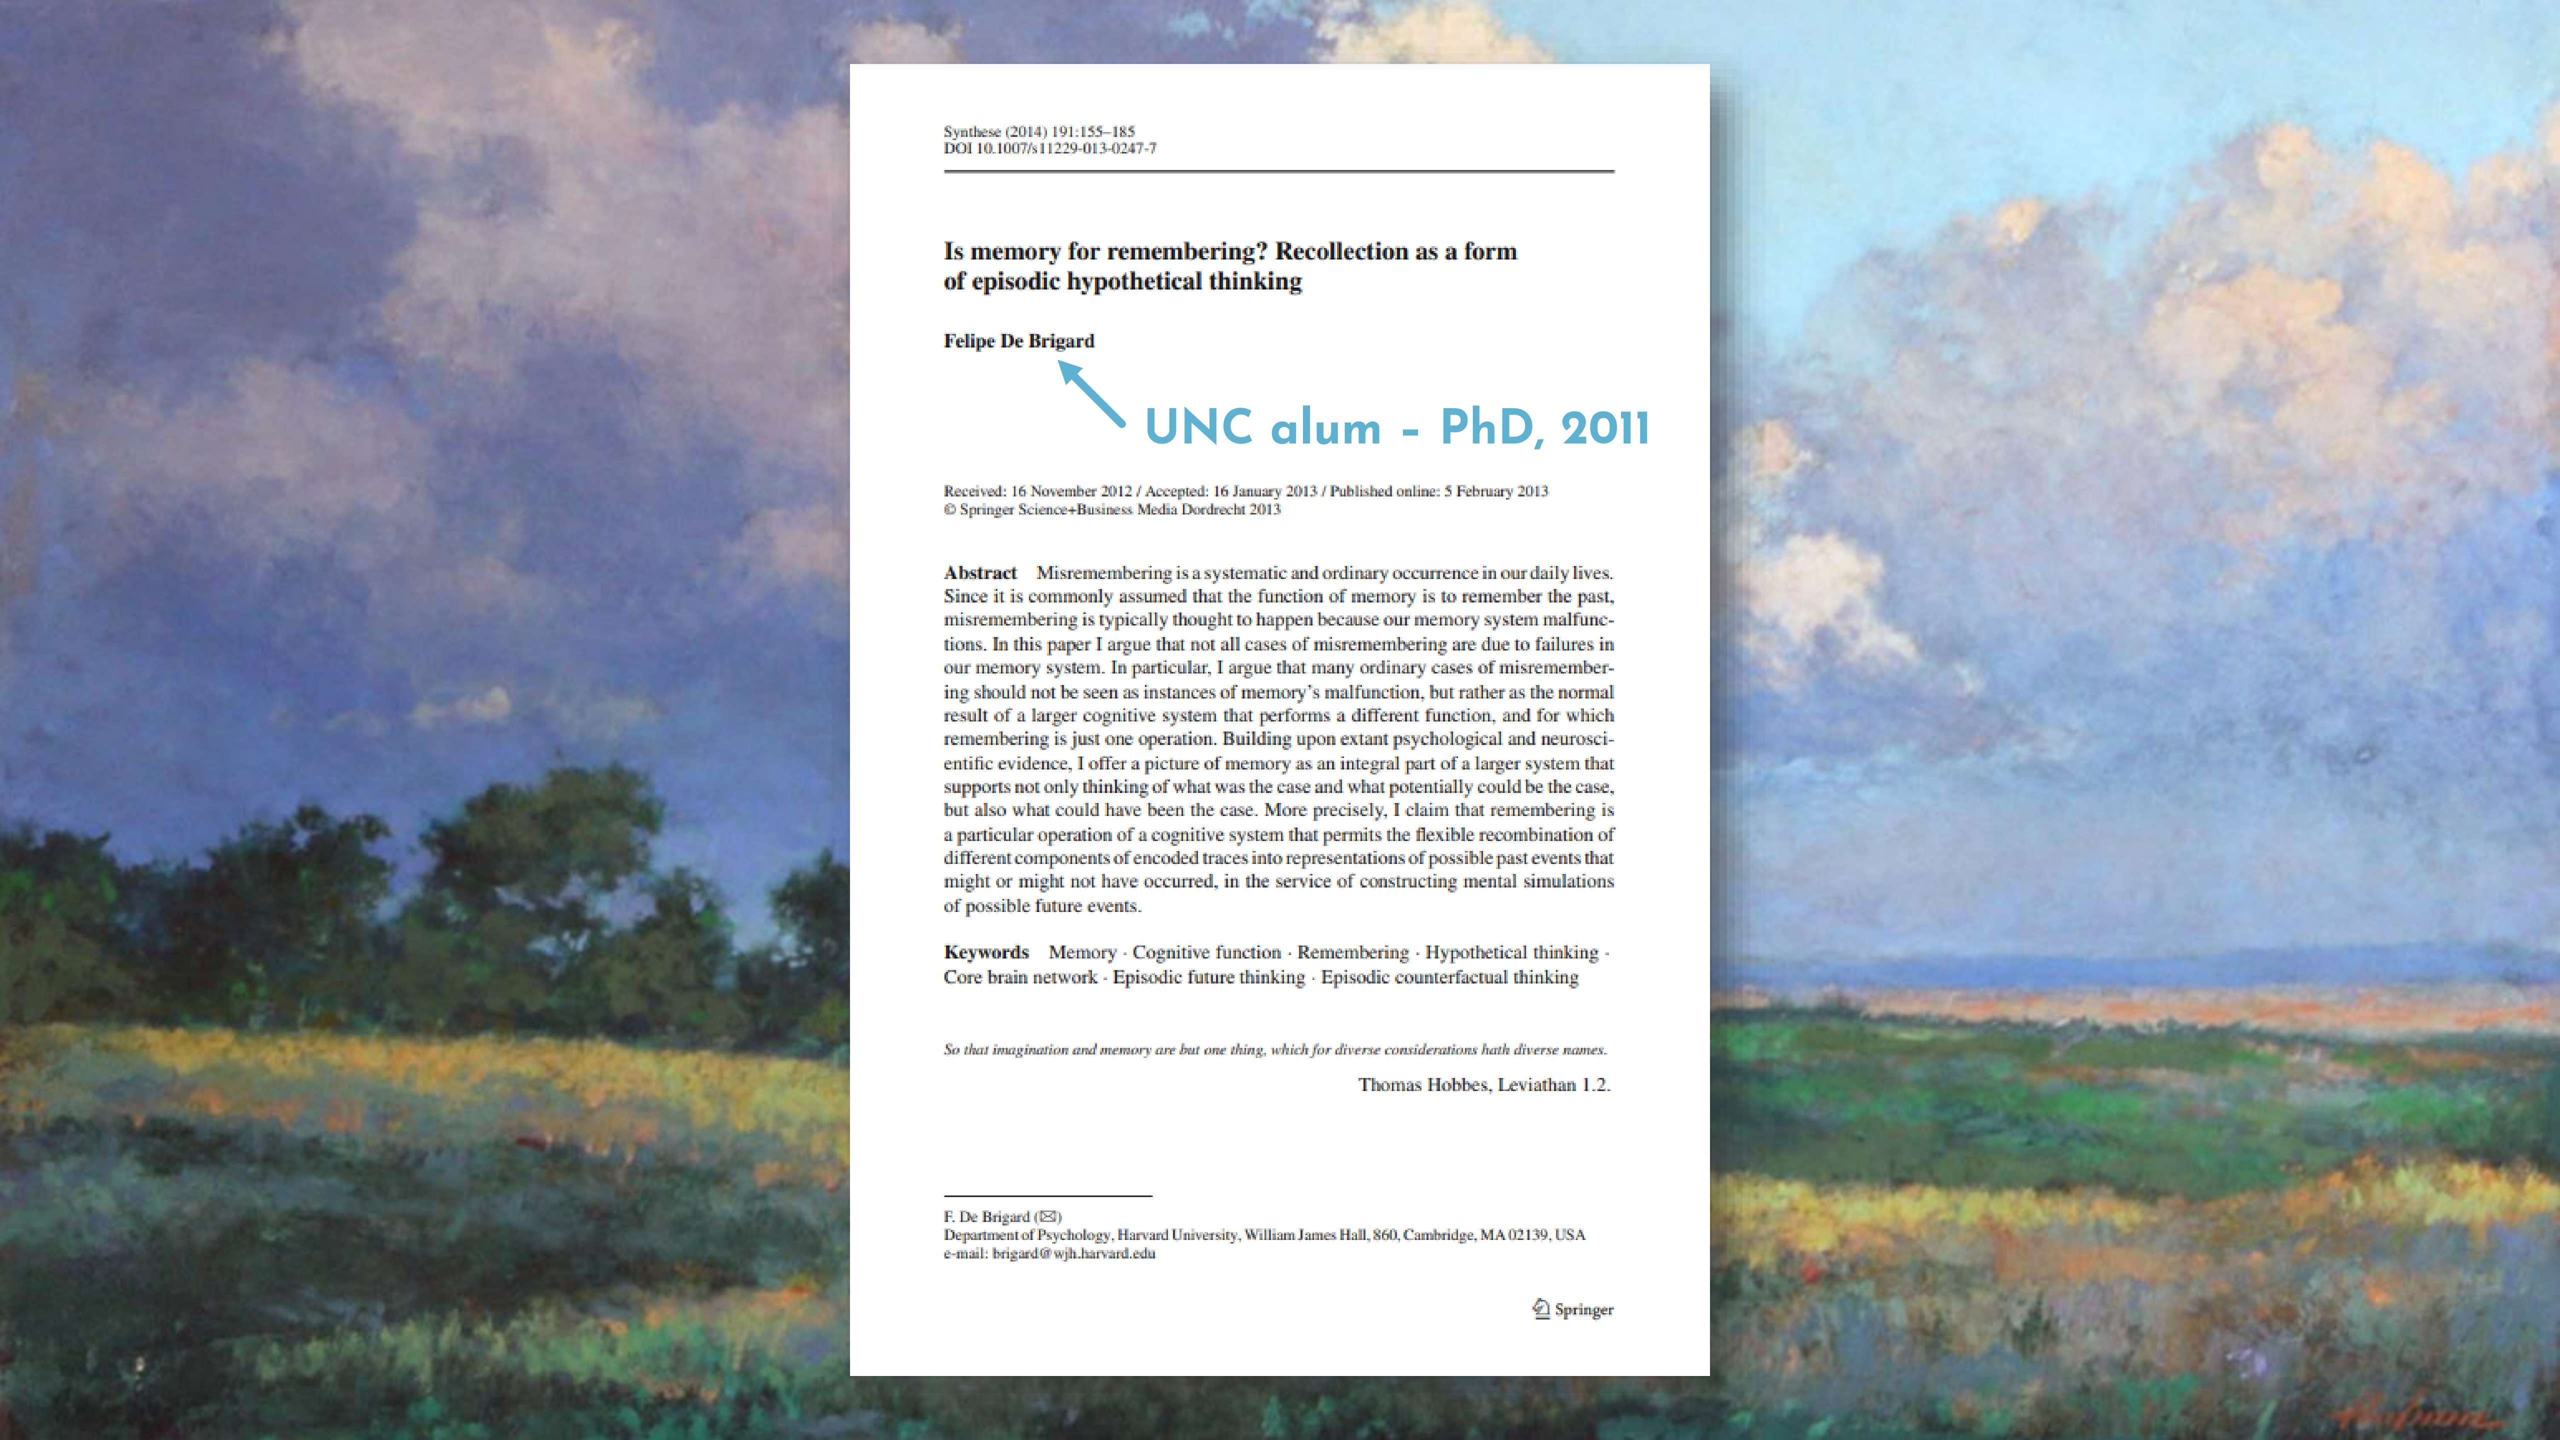

Synthese (2014) 191:155–185
DOI 10.1007/s11229-013-0247-7

Is memory for remembering? Recollection as a form of episodic hypothetical thinking

Felipe De Brigard

Received: 16 November 2012 / Accepted: 16 January 2013 / Published online: 5 February 2013

Abstract Misremembering is a systematic and ordinary occurrence in our daily lives. Since it is commonly assumed that the function of memory is to remember the past, misremembering is typically thought to happen because our memory system malfunctions. In this paper I argue that not all cases of misremembering are due to failures in our memory system. In particular, I argue that many ordinary cases of misremembering should not be seen as instances of memory's malfunction, but rather as the normal result of a larger cognitive system that performs a different function, and for which remembering is just one operation. Building upon extant psychological and neuroscientific evidence, I offer a picture of memory as an integral part of a larger system that supports not only thinking of what was the case and what potentially could be the case, but also what could have been the case. More precisely, I claim that remembering is a particular operation of a cognitive system that permits the flexible recombination of different components of encoded traces into representations of possible past events that might or might not have occurred, in the service of constructing mental simulations of possible future events.

Keywords Memory · Cognitive function · Remembering · Hypothetical thinking · Core brain network · Episodic future thinking · Episodic counterfactual thinking

So that imagination and memory are but one thing, which for diverse considerations hath diverse names.

Thomas Hobbes, Leviathan 1.2.

F. De Brigard (✉)
Department of Psychology, Harvard University, William James Hall, 860, Cambridge, MA 02139, USA
e-mail: brigard@wjh.harvard.edu

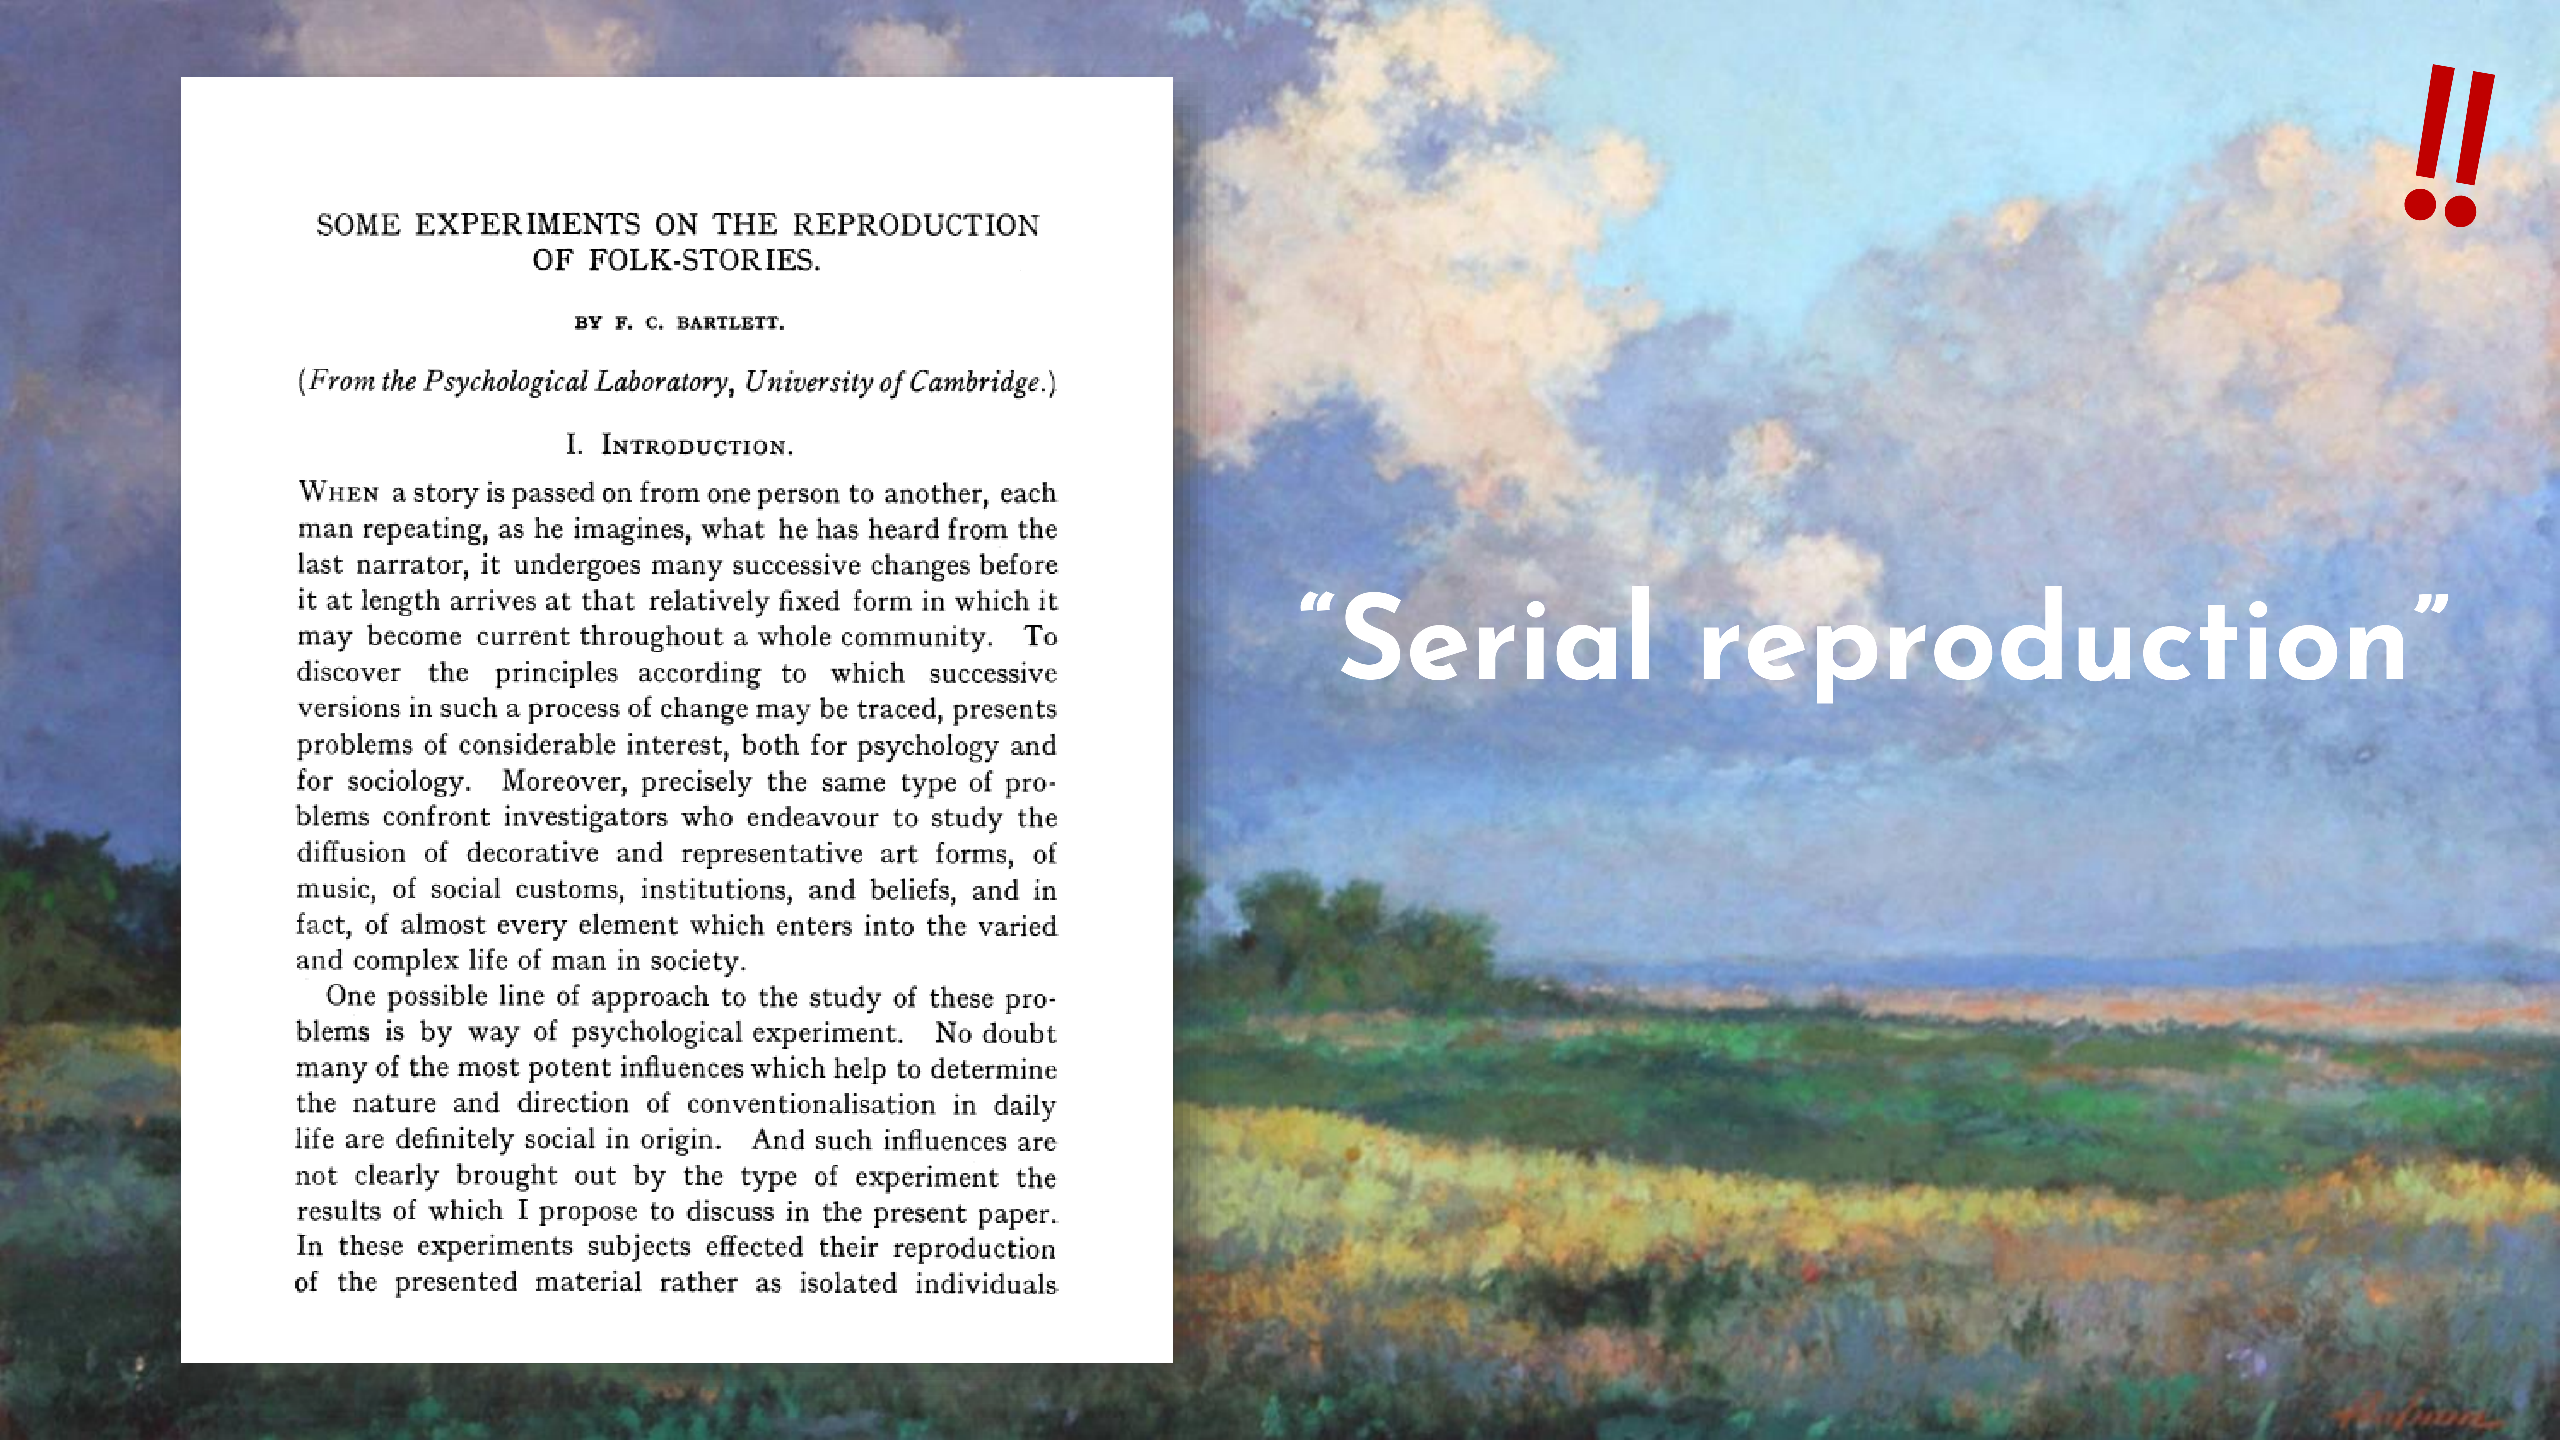

SOME EXPERIMENTS ON THE REPRODUCTION OF FOLK-STORIES.

BY F. C. BARTLETT.

(*From the Psychological Laboratory, University of Cambridge.*)

I. INTRODUCTION.

WHEN a story is passed on from one person to another, each man repeating, as he imagines, what he has heard from the last narrator, it undergoes many successive changes before it at length arrives at that relatively fixed form in which it may become current throughout a whole community. To discover the principles according to which successive versions in such a process of change may be traced, presents problems of considerable interest, both for psychology and for sociology. Moreover, precisely the same type of problems confront investigators who endeavour to study the diffusion of decorative and representative art forms, of music, of social customs, institutions, and beliefs, and in fact, of almost every element which enters into the varied and complex life of man in society.

One possible line of approach to the study of these problems is by way of psychological experiment. No doubt many of the most potent influences which help to determine the nature and direction of conventionalisation in daily life are definitely social in origin. And such influences are not clearly brought out by the type of experiment the results of which I propose to discuss in the present paper. In these experiments subjects effected their reproduction of the presented material rather as isolated individuals

American Psychological Association

Journal of Experimental Psychology: General

0096-3445/18/$12.00

2018, Vol. 147, No. 10, 1466–1487
http://dx.doi.org/10.1037/xge0000446

TeleFace: Serial Reproduction of Faces Reveals a Whiteward Bias in Race Memory

Stefan Uddenberg and Brian J. Scholl
Yale University

How is race encoded into memory when viewing faces? Here we demonstrate a novel systematic bias in which our memories of faces converge on certain prioritized regions in our underlying "face space," as they relate to perceived race. This convergence was made especially salient using a new visual variant of the method of serial reproduction: "TeleFace." A single face was briefly presented, with its race selected from a smooth continuum between White and Black (matched for mean luminance). The observer then reproduced that face, using a slider to morph a test face along this continuum. Their response was then used as the face initially presented to the next observer, and so on down the line in each reproduction chain. White observers' chains consistently and steadily converged onto faces significantly Whiter than they had initially encountered—Whiter than both the original face in the chain and the continuum's midpoint—regardless of where chains began. Indeed, even chains beginning near the Black end of the continuum inevitably ended up well into White space. Very different patterns resulted when the same method was applied to other arbitrary face stimuli. These results highlight a systematic bias in memory for race in White observers, perhaps contributing to the more general notion in social cognition research of a 'White default.'

Keywords: face memory, face perception, race perception, serial reproduction

Faces, Spaces, and Races

Although the underlying biological reality of race is disputed, its psychological and perceptual reality is not (Cosmides, Tooby, & Kurzban, 2003). As such, a key task for research on social perception is to determine how race is encoded into visual memory—especially in the context of what are surely the most ubiquitous and salient social stimuli of all: faces. And as is so often the case, insights into this process may come from an exploration of how it can go awry. Visual memory for race in faces, like any form of memory, is imperfect. This can arise from noisy representations, but some of this imperfection may also arise because of *biases* in memory.

In what may be the most popular framework for discussing face perception, race is characterized as a (set of) dimension(s) within a multidimensional 'face space.' In this framework, a bias in visual memory may arise when some particular region(s) of the space are prioritized—such that a face representation is in effect pulled from one region in the space toward some other region. Here we explore the possibility that visual memory for faces—using race as a case study—is in effect biased toward our representation of a 'default face.' We aim to document the existence of such a bias, to characterize its nature and extent, and to highlight how it conflicts in a striking way with most past research on memory for race. In doing so, we also aim to demonstrate how a variant of the method of serial reproduction—here applied to visual stimuli (in what we call the 'TeleFace' task)—may be especially useful for revealing the nature of such 'default' representations.

Representing Faces

Human faces are among the most salient and important visual stimuli we encounter in our everyday lives. Yet the differences between faces are subtle, as they all share the same parts and global configuration. Great strides have been made in recent years in understanding how we perceive and remember various aspects of such stimuli in terms of an underlying multidimensional "face space" (for recent reviews see Todorov, Olivola, Dotsch, & Mende-Siedlecki, 2015; Valentine, Lewis, & Hills, 2016). In this framework, faces are represented not as arbitrarily varying visual stimuli, but rather in terms of a specific set of continuous dimensions that collectively comprise a psychological similarity space. As such, faces that are similar to one another on a number of dimensions are located close to one another in the 'space,' whereas faces with dissimilar values on many dimensions are further apart. Some such dimensions may reflect relatively simple geometric properties (e.g., nose size or forehead height; e.g., Hurlbert, 2001), while others may reflect more holistic and complex patterns (e.g.,

This article was published Online First July 16, 2018.

Stefan Uddenberg and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and/or comments on earlier drafts, we thank Marvin Chun, Chaz Firestone, Greg McCarthy, and the members of the Yale Perception and Cognition Laboratory. For help developing the paradigm for online testing, and for general technical wizardry, we thank Emily Ward. The data presented in this article were also presented as talks at the 15th and 17th annual meetings of the Vision Sciences Society.

Correspondence concerning this article should be addressed to Stefan Uddenberg or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: stefan.uddenberg@yale.edu or brian.scholl@yale.edu

American Psychological Association

Journal of Experimental Psychology: General

0096-3445/18/$12.00

2018, Vol. 147, No. 10, 1466–1487
http://dx.doi.org/10.1037/xge0000446

TeleFace: Serial Reproduction of Faces Reveals a Whiteward Bias in Race Memory

Stefan Uddenberg and Brian J. Scholl
Yale University

How is race encoded into memory when viewing faces? Here we demonstrate a novel systematic bias in which our memories of faces converge on certain prioritized regions in our underlying "face space," as they relate to perceived race. This convergence was made especially salient using a new visual variant of the method of serial reproduction: "TeleFace." A single face was briefly presented, with its race selected from a smooth continuum between White and Black (matched for mean luminance). The observer then reproduced that face, using a slider to morph a test face along this continuum. Their response was then used as the face initially presented to the next observer, and so on down the line in each reproduction chain. White observers' chains consistently and steadily converged onto faces significantly Whiter than they had initially encountered—Whiter than both the original face in the chain and the continuum's midpoint—regardless of where chains began. Indeed, even chains beginning near the Black end of the continuum inevitably ended up well into White space. Very different patterns resulted when the same method was applied to other arbitrary face stimuli. These results highlight a systematic bias in memory for race in White observers, perhaps contributing to the more general notion in social cognition research of a 'White default.'

Keywords: face memory, face perception, race perception, serial reproduction

Faces, Spaces, and Races

Although the underlying biological reality of race is disputed, its psychological and perceptual reality is not (Cosmides, Tooby, & Kurzban, 2003). As such, a key task for research on social perception is to determine how race is encoded into visual memory—especially in the context of what are surely the most ubiquitous and salient social stimuli of all: faces. And as is so often the case, insights into this process may come from an exploration of how it can go awry. Visual memory for race in faces, like any form of memory, is imperfect. This can arise from noisy representations, but some of this imperfection may also arise because of *biases* in memory.

In what may be the most popular framework for discussing face perception, race is characterized as a (set of) dimension(s) within a multidimensional 'face space.' In this framework, a bias in visual memory may arise when some particular region(s) of the space are prioritized—such that a face representation is in effect pulled from one region in the space toward some other region. Here we explore the possibility that visual memory for faces—using race as a case study—is in effect biased toward our representation of a 'default face.' We aim to document the existence of such a bias, to characterize its nature and extent, and to highlight how it conflicts in a striking way with most past research on memory for race. In doing so, we also aim to demonstrate how a variant of the method of serial reproduction—here applied to visual stimuli (in what we call the 'TeleFace' task)—may be especially useful for revealing the nature of such 'default' representations.

Representing Faces

Human faces are among the most salient and important visual stimuli we encounter in our everyday lives. Yet the differences between faces are subtle, as they all share the same parts and global configuration. Great strides have been made in recent years in understanding how we perceive and remember various aspects of such stimuli in terms of an underlying multidimensional "face space" (for recent reviews see Todorov, Olivola, Dotsch, & Mende-Siedlecki, 2015; Valentine, Lewis, & Hills, 2016). In this framework, faces are represented not as arbitrarily varying visual stimuli, but rather in terms of a specific set of continuous dimensions that collectively comprise a psychological similarity space. As such, faces that are similar to one another on a number of dimensions are located close to one another in the 'space,' whereas faces with dissimilar values on many dimensions are further apart. Some such dimensions may reflect relatively simple geometric properties (e.g., nose size or forehead height; e.g., Hurlbert, 2001), while others may reflect more holistic and complex patterns (e.g.,

This article was published Online First July 16, 2018.

Stefan Uddenberg and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and/or comments on earlier drafts, we thank Marvin Chun, Chaz Firestone, Greg McCarthy, and the members of the Yale Perception and Cognition Laboratory. For help developing the paradigm for online testing, and for general technical wizardry, we thank Emily Ward. The data presented in this article were also presented as talks at the 15th and 17th annual meetings of the Vision Sciences Society.

Correspondence concerning this article should be addressed to Stefan Uddenberg or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: stefan.uddenberg@yale.edu or brian.scholl@yale.edu

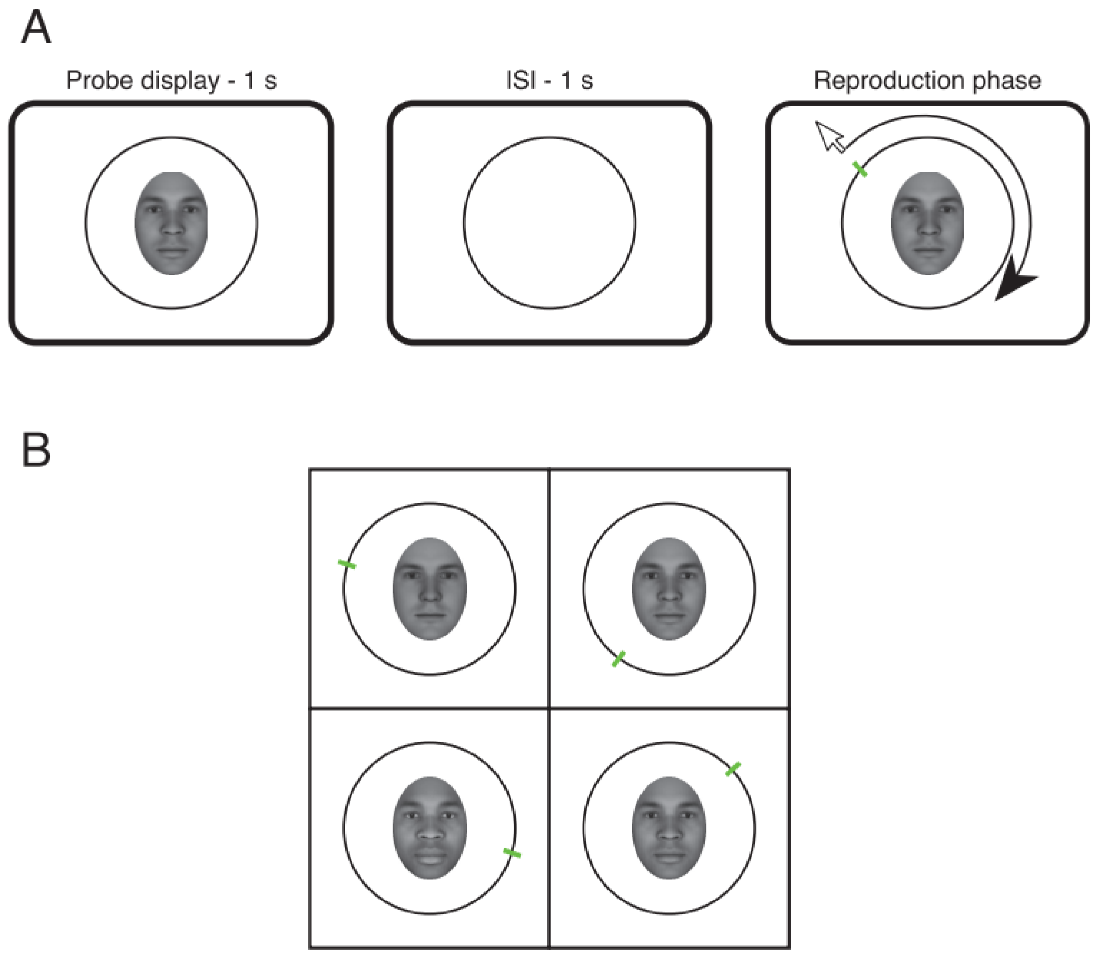

American Psychological Association

Journal of Experimental Psychology: General

0096-3445/18/$12.00

2018, Vol. 147, No. 10, 1466–1487
http://dx.doi.org/10.1037/xge0000446

TeleFace: Serial Reproduction of Faces Reveals a Whiteward Bias in Race Memory

Stefan Uddenberg and Brian J. Scholl
Yale University

How is race encoded into memory when viewing faces? Here we demonstrate a novel systematic bias in which our memories of faces converge on certain prioritized regions in our underlying "face space," as they relate to perceived race. This convergence was made especially salient using a new visual variant of the method of serial reproduction: "TeleFace." A single face was briefly presented, with its race selected from a smooth continuum between White and Black (matched for mean luminance). The observer then reproduced that face, using a slider to morph a test face along this continuum. Their response was then used as the face initially presented to the next observer, and so on down the line in each reproduction chain. White observers' chains consistently and steadily converged onto faces significantly Whiter than they had initially encountered—Whiter than both the original face in the chain and the continuum's midpoint—regardless of where chains began. Indeed, even chains beginning near the Black end of the continuum inevitably ended up well into White space. Very different patterns resulted when the same method was applied to other arbitrary face stimuli. These results highlight a systematic bias in memory for race in White observers, perhaps contributing to the more general notion in social cognition research of a 'White default.'

Keywords: face memory, face perception, race perception, serial reproduction

Faces, Spaces, and Races

Although the underlying biological reality of race is disputed, its psychological and perceptual reality is not (Cosmides, Tooby, & Kurzban, 2003). As such, a key task for research on social perception is to determine how race is encoded into visual memory—especially in the context of what are surely the most ubiquitous and salient social stimuli of all: faces. And as is so often the case, insights into this process may come from an exploration of how it can go awry. Visual memory for race in faces, like any form of memory, is imperfect. This can arise from noisy representations, but some of this imperfection may also arise because of *biases* in memory.

In what may be the most popular framework for discussing face perception, race is characterized as a (set of) dimension(s) within a multidimensional 'face space.' In this framework, a bias in visual memory may arise when some particular region(s) of the space are prioritized—such that a face representation is in effect pulled from one region in the space toward some other region. Here we explore the possibility that visual memory for faces—using race as a case study—is in effect biased toward our representation of a 'default face.' We aim to document the existence of such a bias, to characterize its nature and extent, and to highlight how it conflicts in a striking way with most past research on memory for race. In doing so, we also aim to demonstrate how a variant of the method of serial reproduction—here applied to visual stimuli (in what we call the 'TeleFace' task)—may be especially useful for revealing the nature of such 'default' representations.

Representing Faces

Human faces are among the most salient and important visual stimuli we encounter in our everyday lives. Yet the differences between faces are subtle, as they all share the same parts and global configuration. Great strides have been made in recent years in understanding how we perceive and remember various aspects of such stimuli in terms of an underlying multidimensional "face space" (for recent reviews see Todorov, Olivola, Dotsch, & Mende-Siedlecki, 2015; Valentine, Lewis, & Hills, 2016). In this framework, faces are represented not as arbitrarily varying visual stimuli, but rather in terms of a specific set of continuous dimensions that collectively comprise a psychological similarity space. As such, faces that are similar to one another on a number of dimensions are located close to one another in the 'space,' whereas faces with dissimilar values on many dimensions are further apart. Some such dimensions may reflect relatively simple geometric properties (e.g., nose size or forehead height; e.g., Hurlbert, 2001), while others may reflect more holistic and complex patterns (e.g.,

This article was published Online First July 16, 2018.

Stefan Uddenberg and Brian J. Scholl, Department of Psychology, Yale University.

For helpful conversation and/or comments on earlier drafts, we thank Marvin Chun, Chaz Firestone, Greg McCarthy, and the members of the Yale Perception and Cognition Laboratory. For help developing the paradigm for online testing, and for general technical wizardry, we thank Emily Ward. The data presented in this article were also presented as talks at the 15th and 17th annual meetings of the Vision Sciences Society.

Correspondence concerning this article should be addressed to Stefan Uddenberg or Brian J. Scholl, Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205. E-mail: stefan.uddenberg@yale.edu or brian.scholl@yale.edu

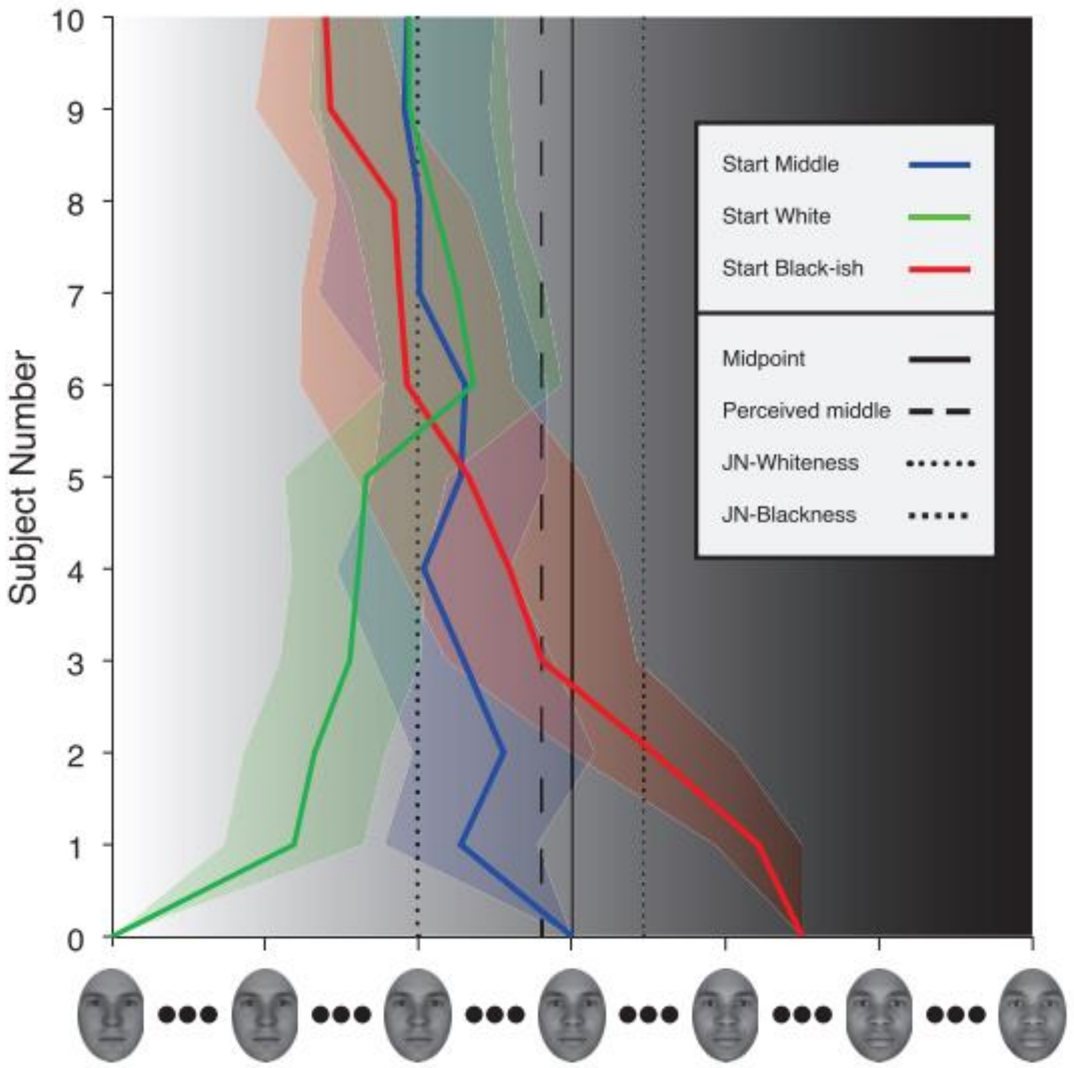

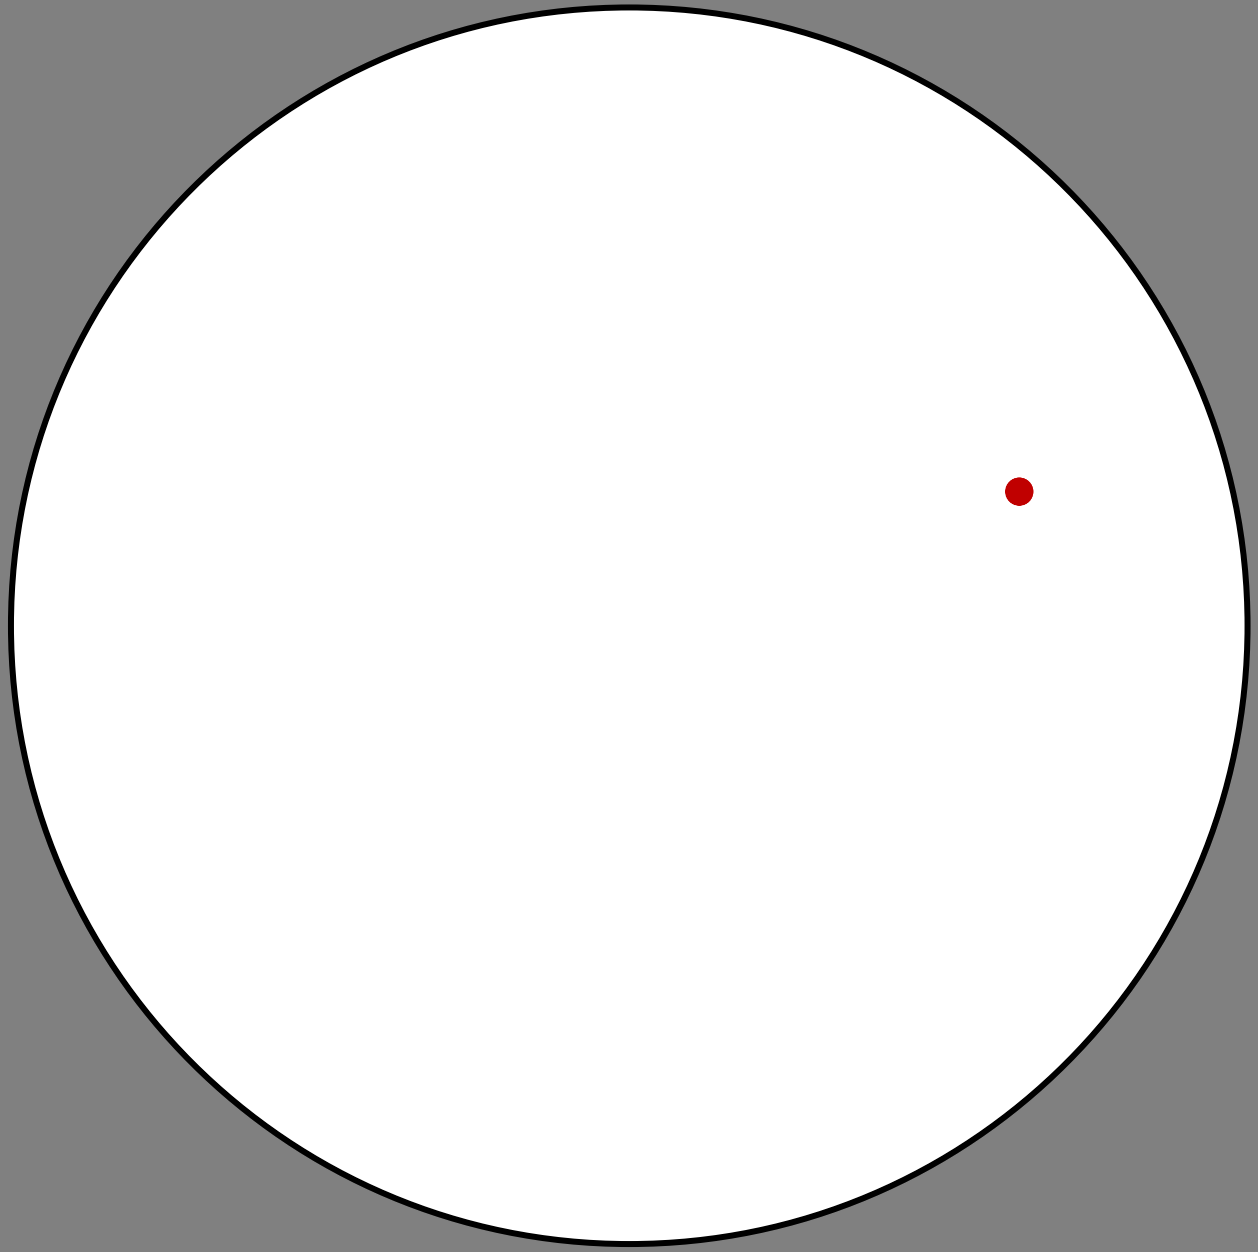

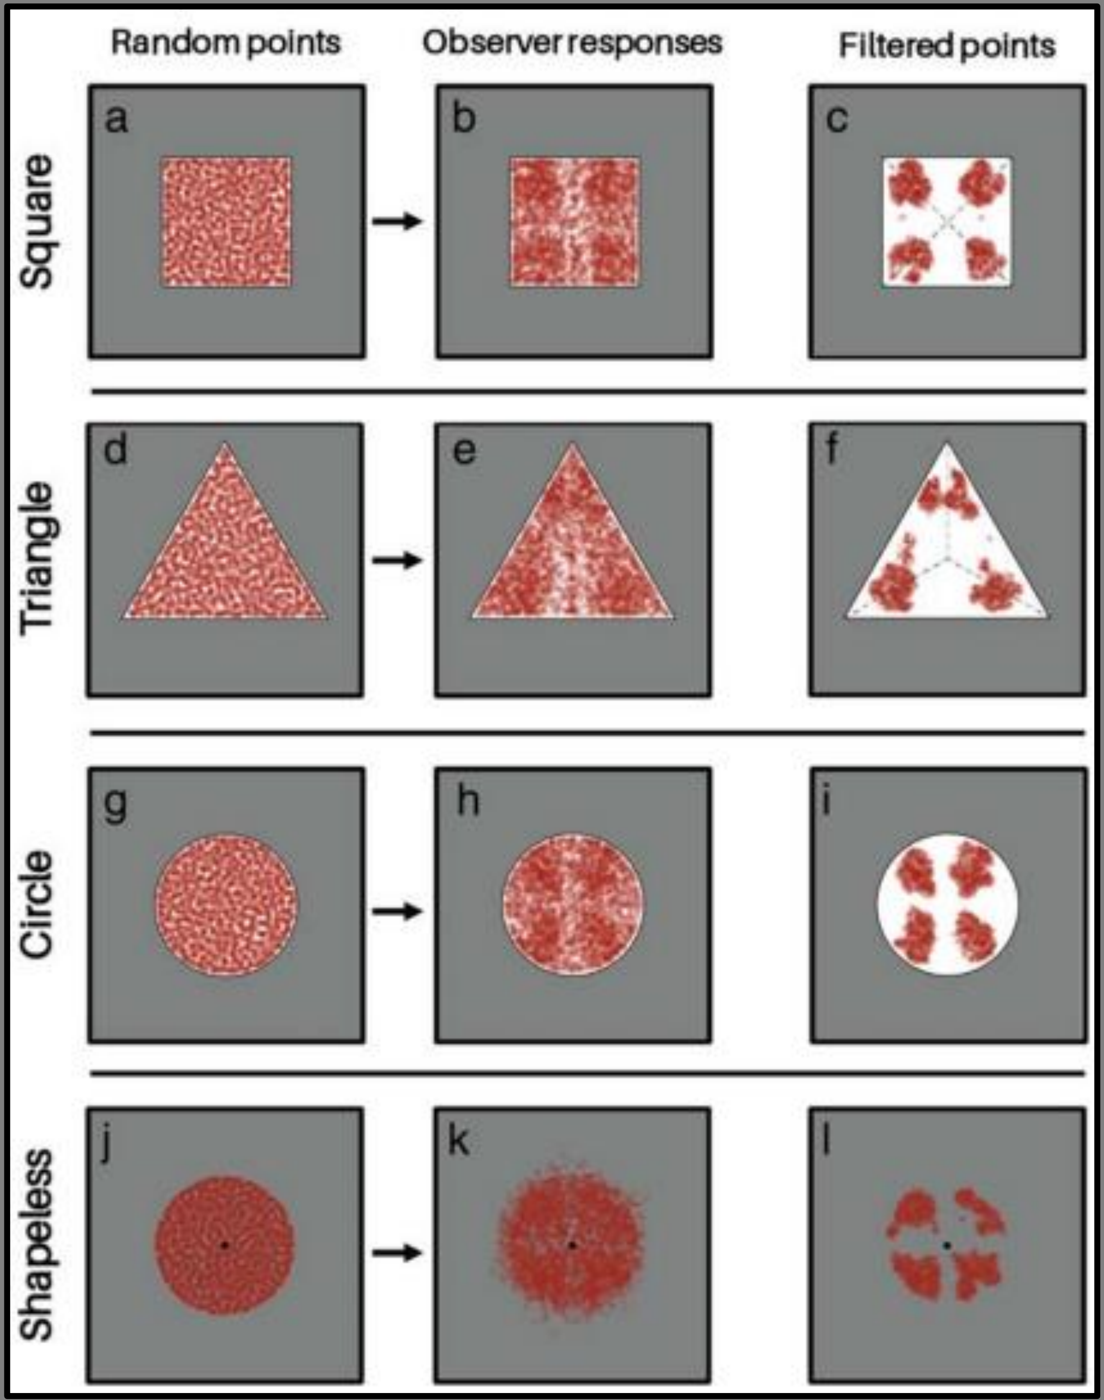

Random points
Observer responses
Filtered points
Square
Triangle
Circle
Shapeless
a
b
c
d
e
f
g
h
i
j
k
l

Attention, Perception, & Psychophysics (2020) 82:3124–3143
https://doi.org/10.3758/s13414-020-02029-8

Systematic angular biases in the representation of visual space

Sami R. Yousif[1] · Yi-Chia Chen[2] · Brian J. Scholl[1]

Published online: 29 April 2020

Abstract

Representing spatial information is one of our most foundational abilities. Yet in the present work we find that even the simplest possible spatial tasks reveal surprising, systematic *mis*representations of space—such as biases wherein objects are perceived and remembered as being nearer to the centers of their surrounding quadrants. We employed both a placement task (in which observers see two differently sized shapes, one of which has a dot in it, and then must place a second dot in the other shape so that their relative locations are equated) and a matching task (in which observers see two dots, each inside a separate shape, and must simply report whether their relative locations are matched). Some of the resulting biases were shape specific. For example, when dots appeared in a triangle during the placement task, the dots placed by observers were biased away from certain parts of the symmetry axes. But other systematic biases were not shape specific, and seemed instead to reflect differences in the grain of resolution for different regions of space. For example, with both a circle and even a shapeless configuration (with only a central landmark) in the matching task, observers were better at discriminating angular differences (when a dot changed positions around the circle, as opposed to inward/outward changes) in cardinal versus oblique sectors. These data reveal a powerful angular spatial bias, and highlight how the resolution of spatial representation differs for different regions and dimensions of space itself.

Keywords Spatial perception · Shape perception · Spatial biases

The ability to accurately perceive and represent space is vital to our success as a species. This is intuitively obvious in the context of skills such as navigation—because even getting to and from your home each day would be impossible without spatial representations. But, in addition, spatial representations are also thought to lie at the foundation of many other cognitive processes, from object representation (e.g., Driver, Davis, Russell, Turatto, & Freeman, 2001; Kahneman, Treisman, & Gibbs, 1992), to numerical processing (e.g., Dehaene, Bossini, & Giraux, 1993; Zorzi, Priftis, & Umiltà, 2002), to reasoning about social relationships (e.g., Parkinson & Wheatley, 2013).

✉ Sami R. Yousif
sami.yousif@yale.edu

1 Department of Psychology, Yale University, Box 208205, New Haven, CT 06520-8205, USA

2 Harvard University, Cambridge, MA, USA

Spatial representation and spatial biases

Ever since classic work on the 'cognitive map' (Tolman, 1948), a great deal of research has been devoted to explaining the nature of spatial representations. Some of this work appeals to an underlying coordinate system, in which objects can be represented in absolute terms—as in proposals for a Euclidean map that represents locations and spatial relationships in a common coordinate system (e.g., Gallistel, 1990; O'Keefe & Nadel, 1978). (Such representations are often thought to be supported by 'place cells' that fire selectively to specific locations; e.g., O'Keefe & Dostrovsky, 1971.) Other work appeals to more relative sorts of representations, in which locations are represented via their angle and distance relationships to other objects (e.g., Kuipers, Tecuci, & Stankiewicz, 2003; Werner, Krieg-Brückner, & Herrmann, 2000), and some animals navigate via representations of egocentric paths, based, for example, on counting the number of steps taken (e.g., McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Müller & Wehner, 1988). (And these sorts of representations are sometimes thought to be supported by 'grid cells' that fire selectively to periodic regions of space with a particular hexagonal

Psychological Review
1991, Vol. 98, No. 3, 352–376

Categories and Particulars: Prototype Effects in Estimating Spatial Location

Janellen Huttenlocher, Larry V. Hedges, and Susan Duncan
University of Chicago

A model of category effects on reports from memory is presented. The model holds that stimuli are represented at 2 levels of detail: a fine-grain value and a category. When memory is inexact but people must report an exact value, they use estimation processes that combine the remembered stimulus value with category information. The proposed estimation processes include truncation at category boundaries and weighting with a central (prototypic) category value. These processes introduce bias in reporting even when memory is unbiased, but nevertheless may improve overall accuracy (by decreasing the variability of reports). Four experiments are presented in which people report the location of a dot in a circle. Subjects spontaneously impose horizontal and vertical boundaries that divide the circle into quadrants. They misplace dots toward a central (prototypic) location in each quadrant, as predicted by the model. The proposed model has broad implications; notably, it has the potential to explain biases of the sort described in psychophysics (contraction bias and the bias captured by Weber's law) as well as asymmetries in similarity judgments, without positing distorted representations of physical scales.

In this article we propose a model of category effects found in reports from episodic memory, that is, reports of the what, when, and where of particular experiences. For example, a person may try to remember the particular properties of an object (e.g., its size and color) or where an object was located. When memory is inexact, people's reports are reconstructions, influenced by schematic or category information (cf. Bartlett, 1932; Brewer & Nakamura, 1984). If information is simply forgotten, a default value may be reported (e.g., the usual color of that sort of object, or a location central to the area where the object could be). If information is remembered, but inexactly, reports may be blends, intermediate between an actual stimulus value and a category value (cf. Belli, 1988).

At present, precise models of category effects on reports of particular experiences are lacking. In proposing such a model here, we begin with stimulus domains based on continuous physical dimensions; object height, temporal or spatial location, and so on. The assumptions of the model are the familiar ones implicit in the previous examples: that memory is hierarchically organized and inexact, and that, in reporting, people may draw from information at two levels (a particular value and a category). The model is novel in that it posits that reports of particular stimulus values are based on estimation procedures that take account of prior (category) information. One of these estimation processes, truncation resulting from category boundaries, was described earlier in Huttenlocher, Hedges, and Prohaska (1988). A second estimation process, weighting with a prototype (a central value in the category), is the focus in the present article.

The proposed uses of category information in estimation introduce systematic biases in reporting even when memory, although inexact, is not itself biased. While introducing bias, these uses of category information nevertheless may be rational; that is, they may improve the overall accuracy of reports by decreasing their variability. This function of categories—the adjustment of inexactly represented stimulus values in a way that may potentially yield more accurate estimates—has not, thus far, been explored.

In the next section, we describe the general form of the proposed model. (The mathematical formulation is presented in the Appendix.) Then we apply the model to the estimation of spatial location. Four experiments are presented in which people reproduce the location of a dot in a circle: The observed patterns of bias are those predicted by the model. The use of category information, in our experiments, improves the overall accuracy of reporting. After explicating the model and producing the evidence for it in the case of the representation of spatial location, we consider the application of the model to more general issues. Notably, the model has implications for claims of systematic distortion in the mental representation of values along physically measurable dimensions, including spatial location, based on biases in reporting (e.g., asymmetries in distance judgments, the biases described in psychophysics). The model shows that, at least in some cases, such biases in reporting can be explained without positing biases in the representation of physical stimulus values. In addition, the model has implications for arguments concerning the representation of category information in memory. In particular, at least for the purpose of estimation, models that posit explicit representation of category information (boundaries and prototypes) may yield a more natural explanation than do models that posit the implicit representation of categories as sets of exemplars.

This work was supported in part by Air Force Office of Scientific Research, Program in Cognition, Grant AFOSR-88-0215. In addition to overall coauthorship, Larry V. Hedges developed the formal presentation that makes up the Appendix.

We thank Ray Adams and Jack Vevea for their help in both data analyses and statistical modeling and Miriam Bassok for helpful comments on an earlier draft of this article.

Correspondence concerning this article should be addressed to Janellen Huttenlocher, Department of Psychology, University of Chicago, 5848 South University Avenue, Chicago, Illinois 60637.


Memory errors
are a tool!

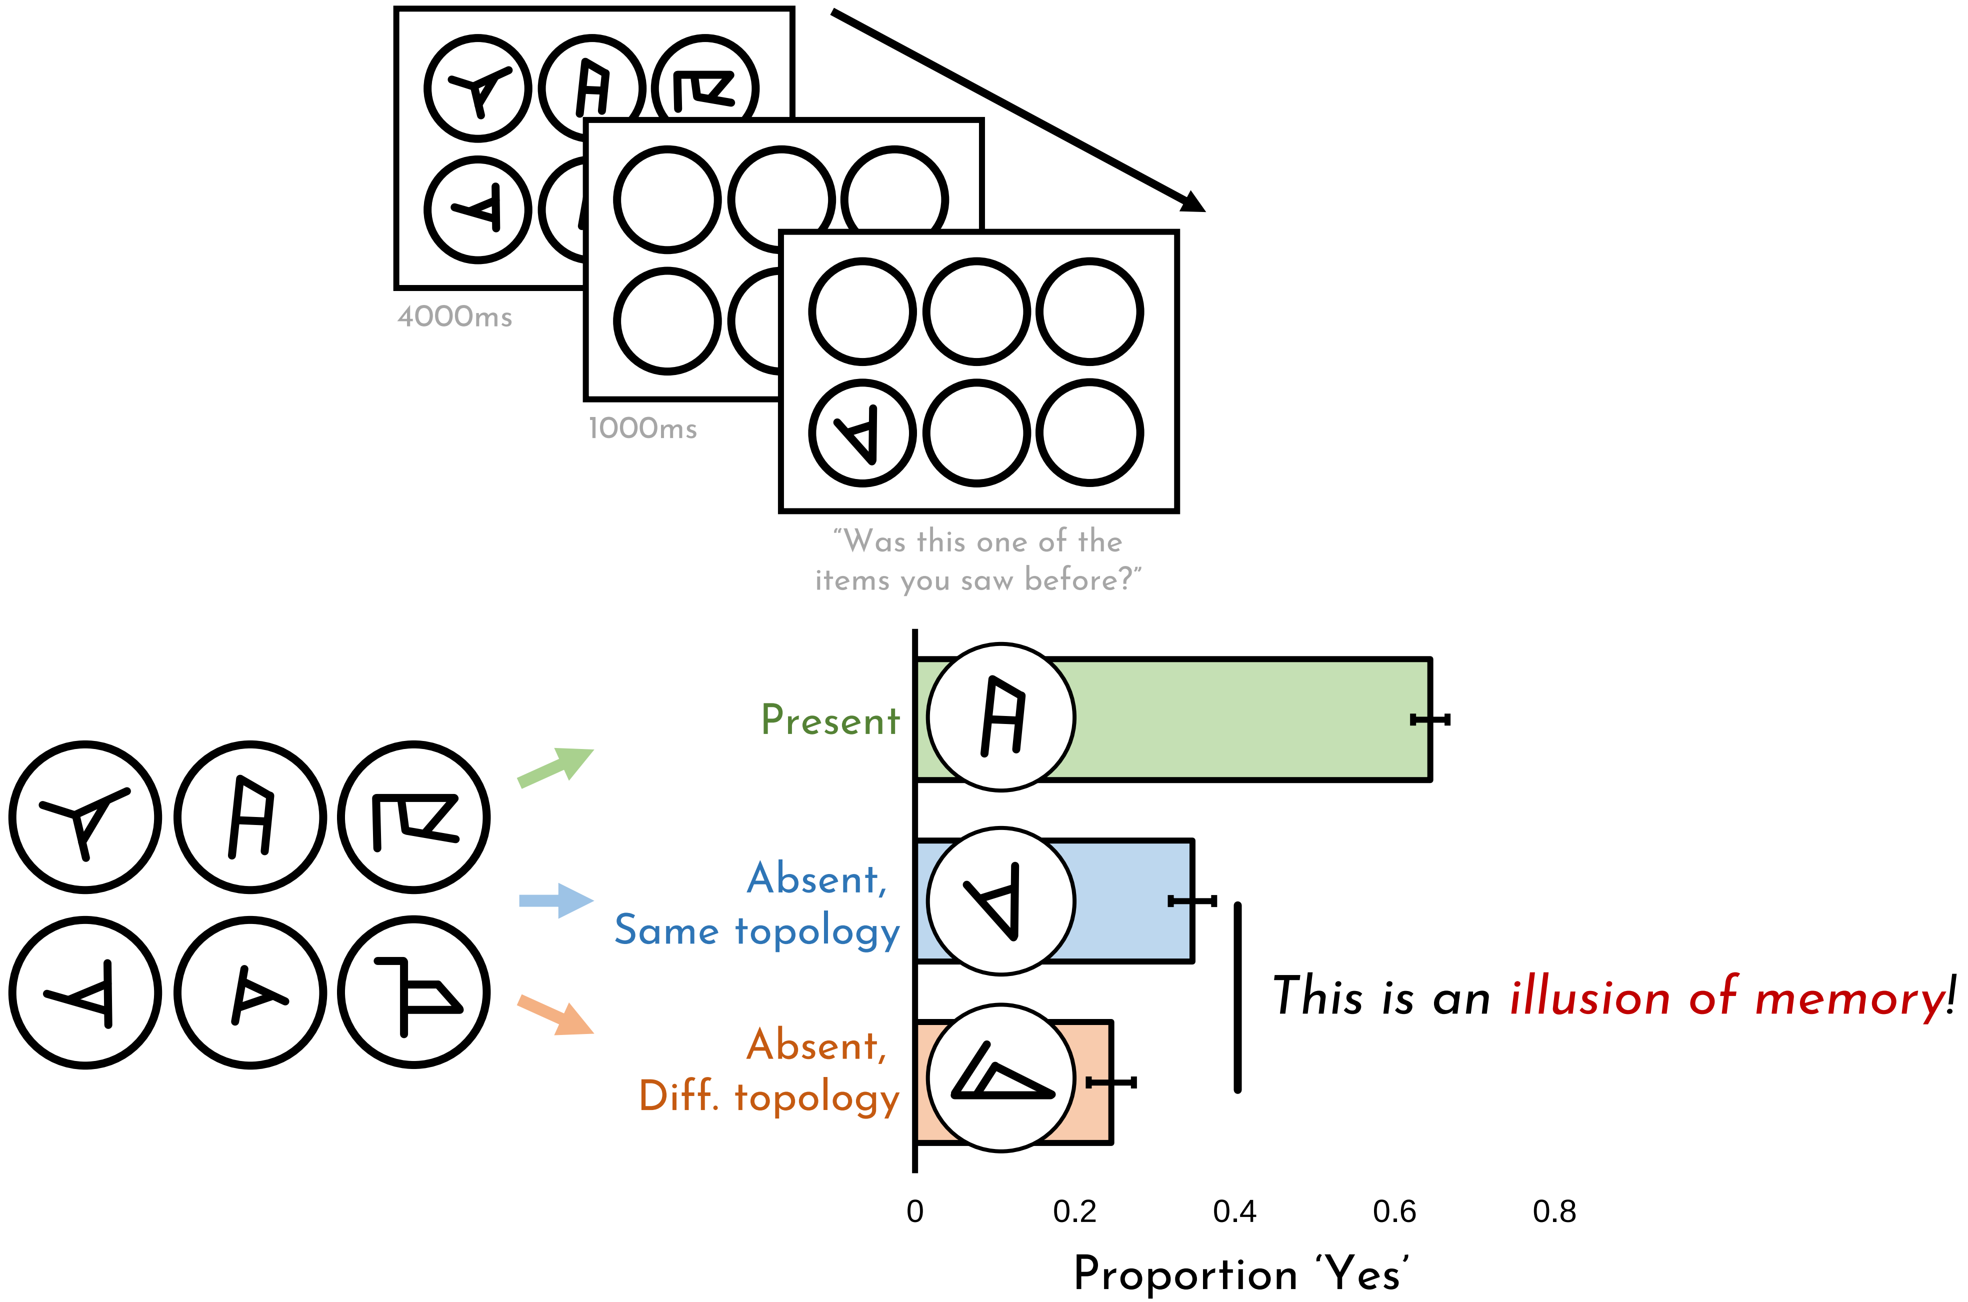

4000ms
1000ms
"Was this one of the items you saw before?"
Present
Absent, Same topology
Absent, Diff. topology
This is an illusion of memory!
0
0.2
0.4
0.6
0.8
Proportion 'Yes'

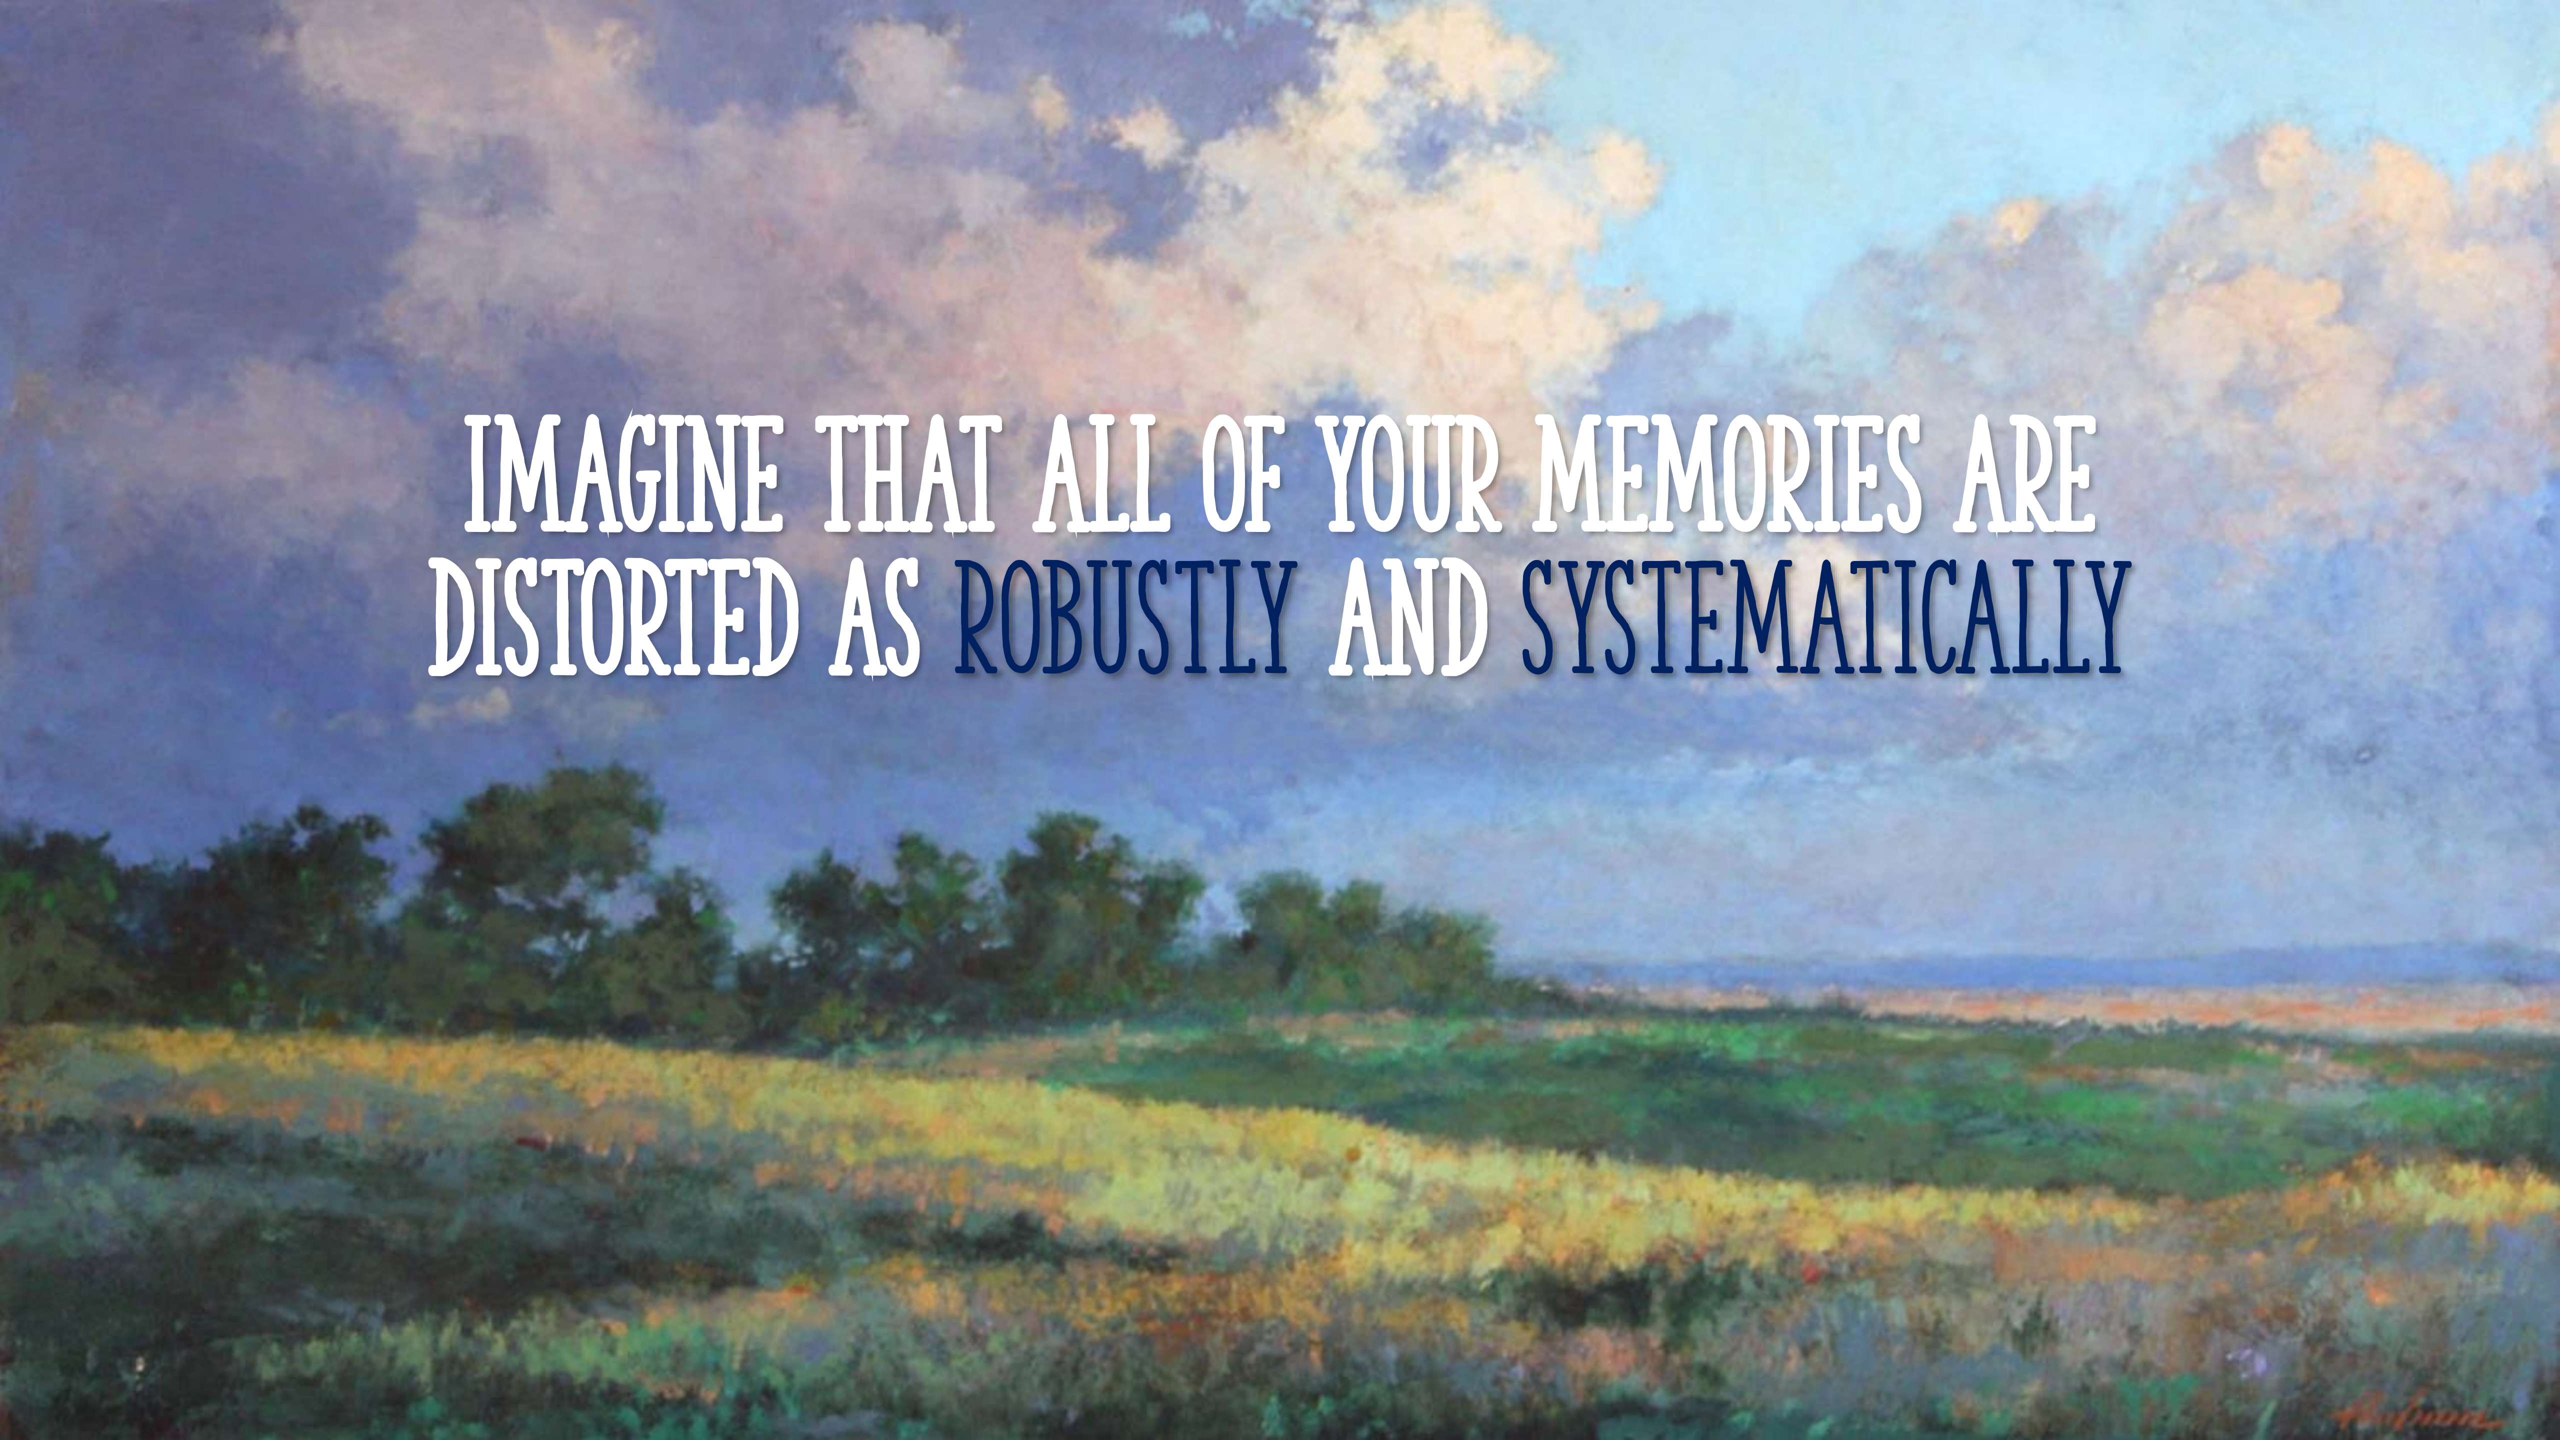
IMAGINE THAT ALL OF YOUR MEMORIES ARE
DISTORTED AS ROBUSTLY AND SYSTEMATICALLY

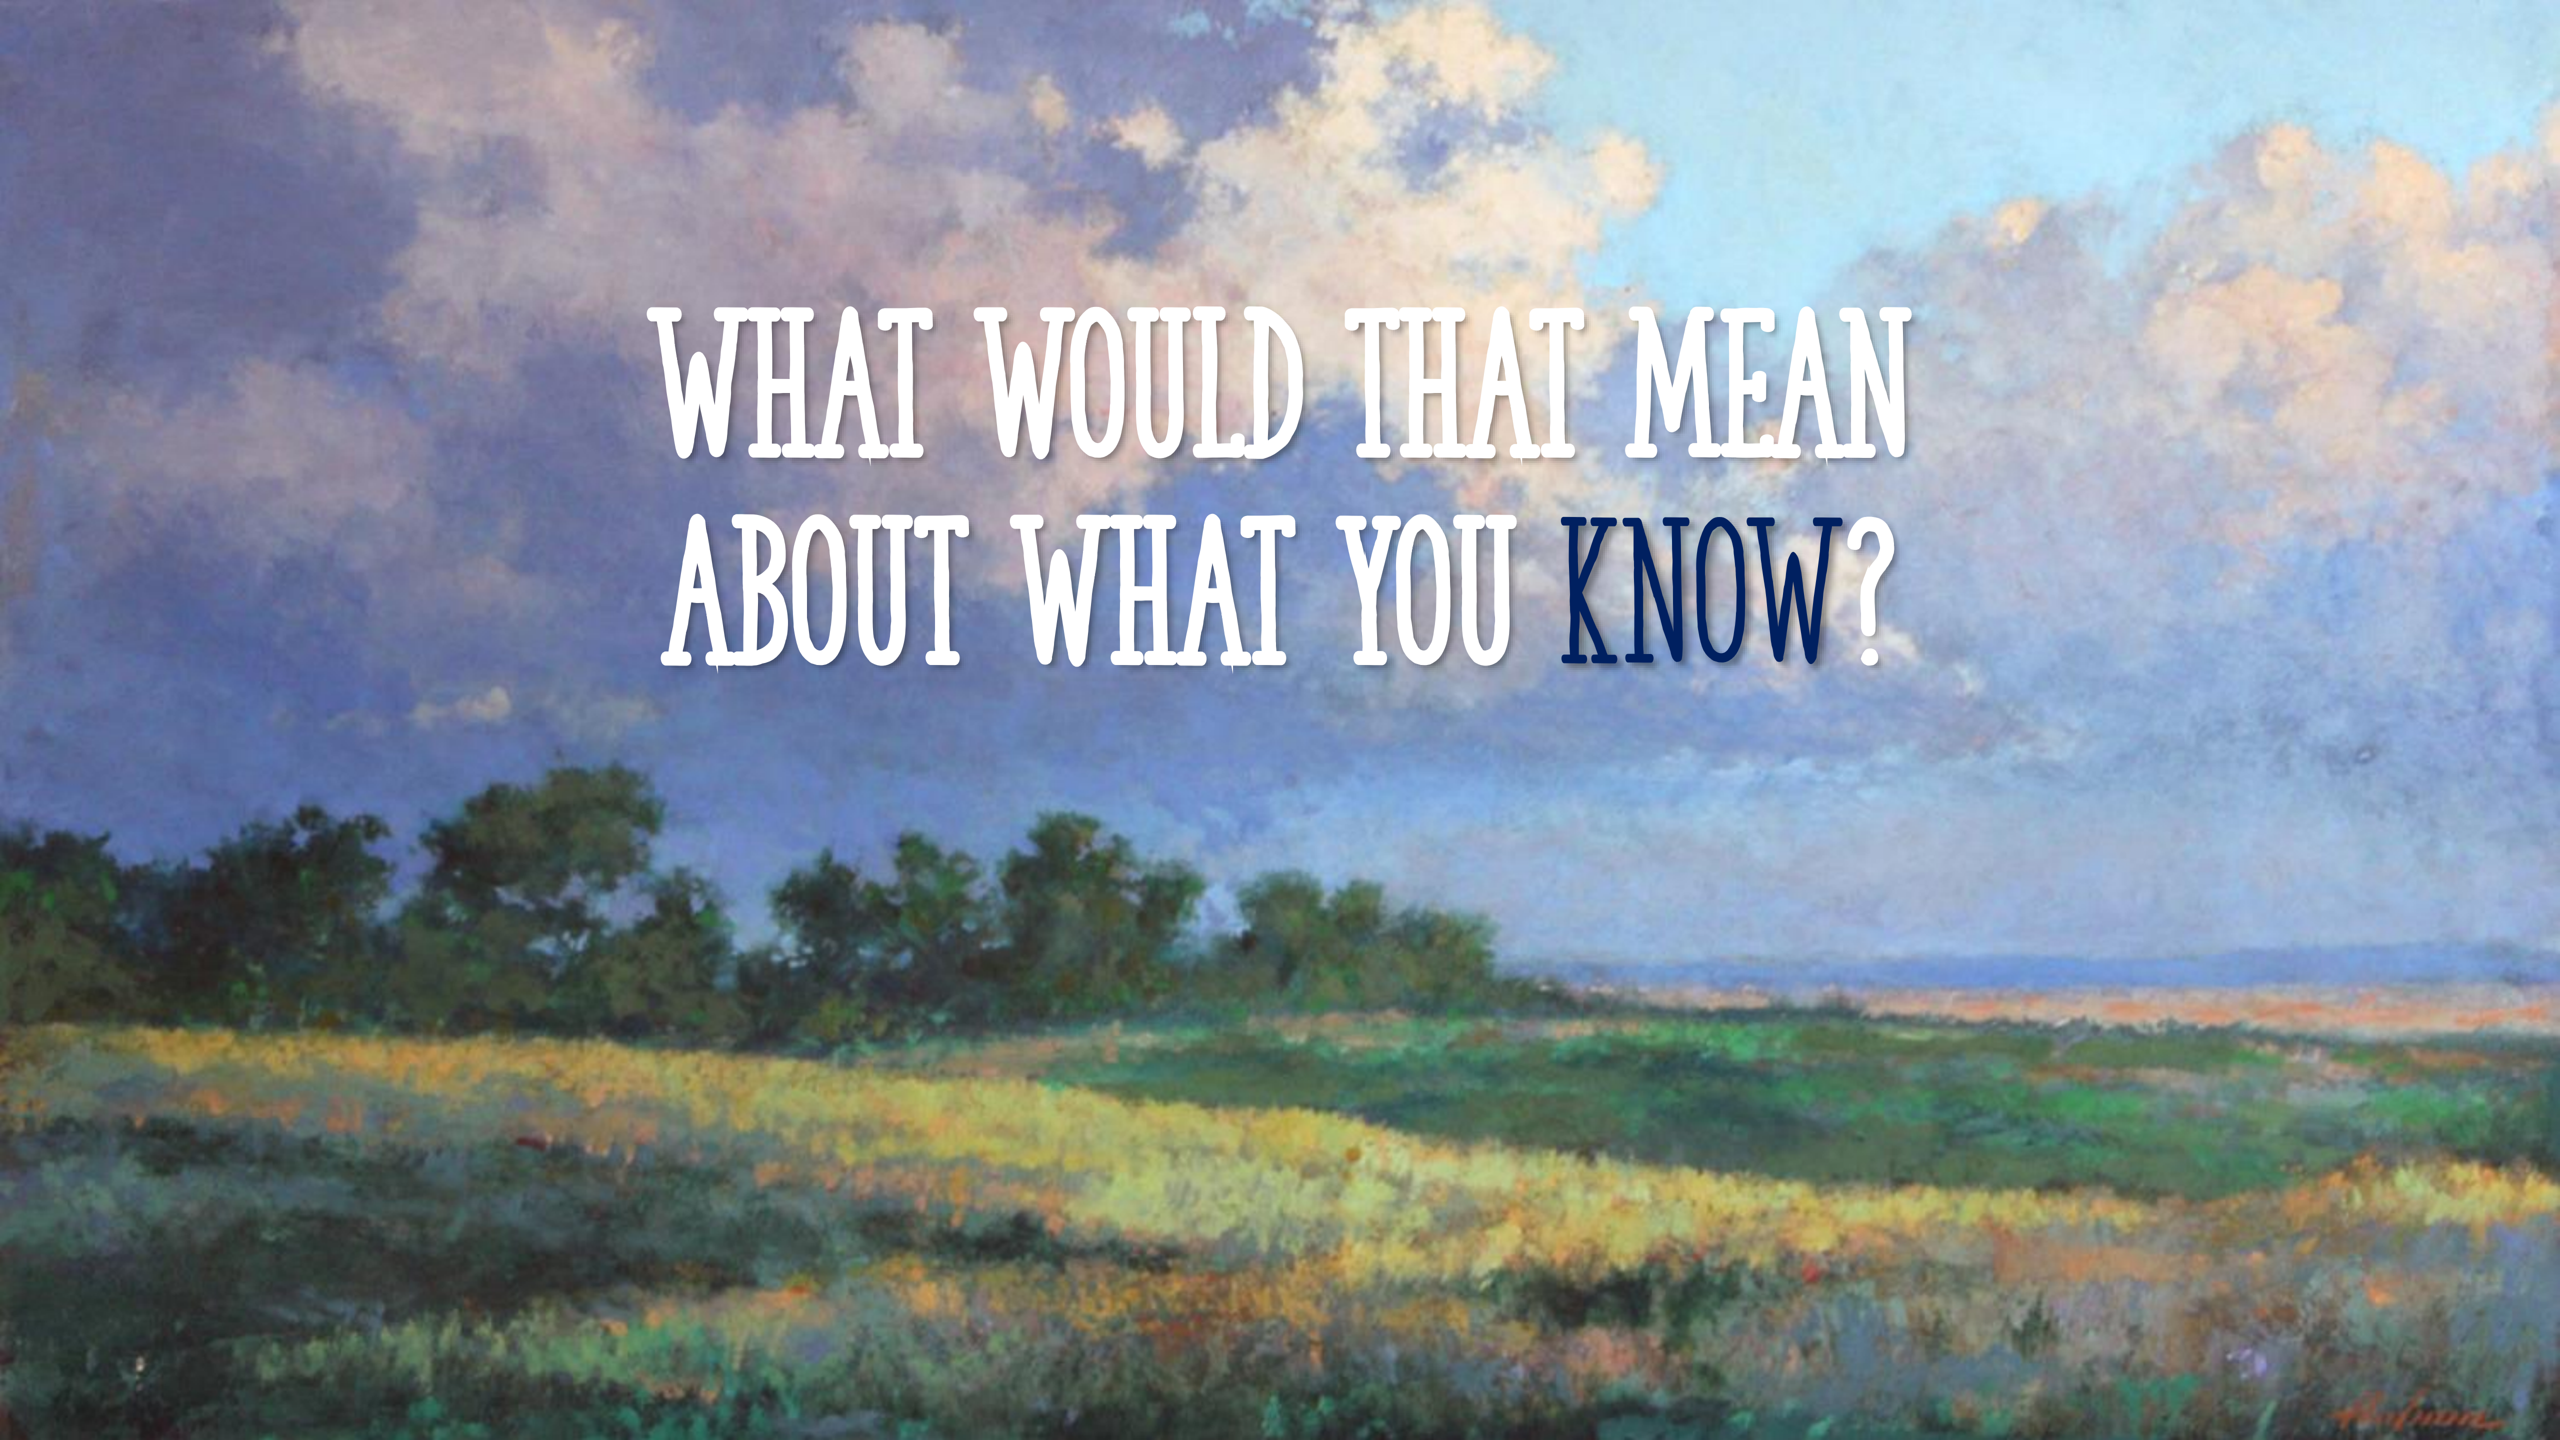
WHAT WOULD THAT MEAN
ABOUT WHAT YOU KNOW?


WE MUST BE EPISTEMICALLY VIGILANT.