As if the Sun Went Round the Earth: A Systematic Review of Empirical Evidence Concerning Perceptual Richness

---

A Thesis
Presented to
The Division of Philosophy, Religion, Psychology, and Linguistics
Reed College

In Partial Fulfillment
of the Requirements for the Degree
Bachelor of Arts

Sophronia Neff Barone
May 2022

Approved for the Division

(Psychology)

Michael Pitts

Acknowledgements

First, I would like to thank Michael, my adviser, for believing in me and giving me the opportunity to work on this project. Thank you to all the people that have supported me during my time at Reed. Words cannot do justice. Mom, thank you for being one of my first teachers and for always being there when I need someone. Dad, thank you for always encouraging my curiosity and passions, and for inspiring me to think deeply about the world. Sydney, I am so lucky to have you as my sister. Bella, thank you for being your cosmic self and always making everything an adventure. Maddi, thank you for being like a sister to me and encouraging me to believe in all of the magic. Kevin, thank you for helping me remember everything. Kaitlin, thank you for all the pure joy you have brought to my life, I cannot wait to see the amazing things you are going to do in this world. Clay, Max, and Elle, thank you for making science fun and for being such wonderful friends. Ryan, Caroline, Charles, Cleo, and Rosalia, thank you for brightening my days, and for giving me homes away from home.

List of Abbreviations

AST	Attention Schema Theory
GWT	Global Workspace Theory
IAB	Inattentional blindness
IIT	Integrated Information Theory

Table of Contents

List of Figures

Abstract

Is conscious perception rich or sparse? A key ongoing debate in consciousness research concerns the bandwidth of perceptual experience. The current study involved a PRISMA literature review of relevant empirical papers. Twenty-six studies were included that explicitly addressed the rich vs. sparse debate. The methods, interpretations, and underlying assumptions of the included studies were critically compared to find opportunities for experimental and theoretical progress. Primary themes and topics of these 26 studies included change blindness, sensory memory, and ensemble perception. Key micro-debates were identified including whether sensory memory is an unconscious process, whether attention is necessary for representation in sensory memory, how exactly element perception and ensemble perception interact, and the nature of representation in instances of change blindness. Possible directions for future experiments aimed at answering some of these questions are discussed. The overall pattern of results was also considered in the context of popular theories of consciousness, including sensorimotor theories of perception, predictive processing, and attention schema theory.

Dedicated to my Dad; it all started with a bead craft….

Introduction

> "He [Ludwig Wittgenstein] once greeted me with the question: 'Why do people say that it was natural to think that the sun went round the earth rather than that the earth turned on its axis?' I replied: 'I suppose, because it looked as if the sun went round the earth.' 'Well,' he asked, 'what would it have looked like if it had *looked* as if the earth turned on its axis?'"
> – Elizabeth Anscombe (1959, p. 151)

Defining Consciousness

Consciousness is an ambiguous term with a variety of meanings. It is commonly used to mean the opposite of whatever happens during death, dreamless sleep, and states of medically determined unconsciousness such as general anesthesia or coma--i.e., "she lost consciousness." What, exactly, did she lose? This understanding reflects a colloquial conceptualization of consciousness as a state of alertness necessary to process and respond to information from the environment (or in the case of dreams, processing internally stored information and simulating responses). If this was the only understanding of consciousness, the problem would be quite simple. Machines do this, as well as newborn infants, all animals, and even simple organisms, yet whether these entities are "conscious" is quite controversial. Why is this? And what can this controversy tell us about how we understand consciousness?

Consciousness as subjective experience

Consciousness is currently defined in scientific research as subjective, first-person experience (Graziano, 2019; Koch, 2019). Terms such as "consciousness," "subjective experience," and "first-person experience" will be used interchangeably throughout the rest of this thesis. But what is meant by them? A calculator can enter a state of "alertness" (turned on) where it can process information, but does it have a subjective *experience* of doing so? This element of *experience* is what makes consciousness more complex than mere alertness. Consider the philosopher's zombie (Campbell 1970; Kirk 1974). This thought experiment asks us to imagine a being that is indistinguishable from a real person, yet has no inner experience. The being may laugh, but have no *experience* of joy. It may cry, but have no *experience* of sadness. For familiarity's sake, this same reasoning is applied often in animal ethics. Despite showing outward signs of distress, many are doubtful as to whether animals *experience* suffering.

Subjectivity: specificity & privateness

What does it mean for experience to be *subjective*? Are all experiences subjective? In a famous 1974 paper, philosopher Thomas Nagel asked: "what is it like to be a bat?" He argued that for an organism to have conscious experience means that "...there is something that it is like to *be* that organism--something it is like *for* the organism" (p. 436, emphasis in original). It is important to understand that what is *for* the organism cannot be for anything else--it is that highly *specific.*

To begin to understand the specificity of subjectivity, it is helpful to first understand the *privateness* of subjectivity. Privateness is readily apparent in the history and etymology of the word *consciousness,* or "...the only space in which you could be together with God even before death" (Metzinger, 2009, p. 26). In 1700, John Locke's *Essay Concerning Human Understanding* defined consciousness as "the perception of what passes in a man's own mind" (p. 48), conceptualizing consciousness as a form of self-awareness. Etymologically related is the word "conscience," the inner knowing of right and wrong. Consciousness enables conscience--prior deliberation on a course of action, a private space to "prepare." Outside of God, if one so believes, the conscious experience is so private that it is only able to be known by the experiencer. Does contemporary scientific research threaten this privateness?

Phenomenal experience

Hypothetically, if technology were to advance to a point where someone's neural activity could be finely monitored in real time, and scientists understood what all of that activity meant, privateness may be threatened in the sense that one could no longer lie. One could no longer suppress their darker thoughts, only revealing the version of themselves they wish the world to see. But this is beside the point. Peering into an individual's mind, onlookers would be met with patterns of neural activity, which they would then interpret. This is different from having the *experience* produced by the mind observed. While the patterns of activity can exist for (be observed by) many, the *experience* can only exist for one.

Nagel asks us to imagine what it is like to be a bat, but then almost immediately admits that it is impossible for us to do so successfully--we can never know what it is like to be a bat! This is because, he argues, "our experience provides the basic material for our imagination, whose range is therefore limited. It will not help to try to imagine that one has webbing on one's arms…" (p. 439). Additionally, this limitation "...exists between one person and another" (p. 440). Because we have never had the *experience* of being another, we cannot successfully imagine *what it is like* to be another. We can try our best to theorize the mental states of others, and often do so quite well, but we will never have *full* knowledge, because that knowledge only exists for a single experiencer.

The word *phenomena* is derived from the Greek *phainomenon*, meaning "appearance" (Smith & Zalta, 2018). In Western philosophy, *phenomenology* is the study of consciousness, or, the study of appearances. Metzinger aptly and poetically defined consciousness as "...the appearance of a world" (2009, p. 15). The phenomenal nature of subjectivity can be understood as the limitation Nagel describes as existing "...between one person and another" (p. 440). The phenomenal appearance is for a specific subject, from a specific point of view, occurring at a specific moment.

Physicalism (discussed in more detail on page 5) asserts that for two instances of identical physical facts, the social, psychological, etc. aspects of those instances will also be identical. But can there truly be two physically identical instances? We have discussed here the importance of specificity to phenomenal experience–specific moment, specific point-of-view. If two instances were to be truly physically identical, they would have to overlap in both time and space. The curiosity of the possibility of such an occurrence (referred to in philosophy as the problem of coinciding objects, or, colocation–can two objects be in the same place at the same time?) highlights the importance of these factors–can two phenomenal experiences be identical if they occur in a different place, at a different time, for a different person?

Subjectivity & objectivity: irresolvable tension?

To summarize, Nagel highlights a tension between subjectivity and empirical reductionism in the form of an epistemic conflict. Reductionism employs *objectivity*, striving towards a non-specific point of view to generate a non-specific set of information. While subjective experience contains evidence that is private and can only be known in the same way by one, objective knowledge is public and can be known in the same way by many. Because of this, systems and theories grounded in objectivity are potentially incapable of thoroughly explaining all aspects of subjectivity. The public aspects of subjectivity such as neural activity can be explained, and an individual may describe to others what something is "like," however the fleeting, firsthand *experience* associated with any given neural activity and accompanying report can only be known by her who has it. Nagel argues: "The reason is that every subjective phenomenon is *essentially* connected with a single point of view, and it seems inevitable that an objective, physical theory will abandon that point of view" (p. 437, emphasis added).

The scientific study of consciousness

Perspective & theory

The predictive power enjoyed by the physical sciences has led to an increase in the popularity of *physicalism*. Physicalism, also known as materialism, assumes that everything that exists is physical (or, material) in nature. This is opposed to *dualism,* which posits a mind or soul separate from the physical world. Leading physicalist attempts at reducing consciousness claim that it is the product of brains and neural activity. While a close relationship between brains and consciousness is unanimously accepted--we almost all agree that sleep, death, brain injury, and psychoactive substances impact states of consciousness--the idea that the physical properties and processes of the brain are both necessary and sufficient for consciousness is what sets physicalist understandings apart from those that rely on dualism. Is it possible for this objective perspective to describe consciousness without, as Nagel warns, abandoning the specificity of subjectivity?

Physicalism & satisfaction

Truthfully, neuroscience is a relatively young and developing field with much progress still to be made. With that being said, we are still far off from being able to conclude that any failure of the physical sciences to satisfyingly explain consciousness is due to an epistemic incompatibility–such "failures" may simply reflect the need for a more comprehensive understanding of the brain and what it does. In philosophy of consciousness, *the explanatory gap* refers to the intuitive difficulty faced when attempting to conceptualize conscious experience as the product of physical activity. However, difficulty does not mean impossibility. Associating specific patterns of neural activity with consciousness may someday be as accepted as associating H_2O with water--while both terms uncontroversially refer to the same thing, they are used in different contexts for different purposes (Lewis, 1972). This idea is known in philosophy as *identity* *theory*--certain mental states are identical with certain brain states. One may reject identity theory on the grounds that, as discussed earlier, brain activity cannot be truly *identical* with any mental state, as the mental state is *experienced* as something quite different than the firing of neurons that is *observed.* However, just because something (neural activity) appears different from different perspectives (first vs. third person) doesn't mean existing correspondences are not worth investigating. As we learn more about the brain, such correspondences may become more "intuitive."

Another possibility is that dissatisfaction with physical reductions of consciousness may boil down to personal predisposition and preference. While this thesis so far has paid special attention to points of view that are particularly dissatisfied with physicalism, it is important to acknowledge that many academics and non-academics alike are especially satisfied by physicalism, so much so that the scientific study of consciousness developed as a field in the first place! Because there are so many varying perspectives, it would not be fair to give up on psychophysical reductionism until the connection between consciousness and the brain has been solved or at least exhausted.

Perhaps satisfaction will ultimately come as the result of compromise—what would happen if the sciences displayed a greater and more consistent understanding of the tensions and difficulties at play? Acknowledging diversity of perspective and the epistemic challenges faced by empiricism when dealing with this topic is an important part of the approach of this thesis.

Method

Aspirations & difficulties

One of the main goals of the scientific study of consciousness is to isolate which specific brain regions, mechanisms, and processes are necessary and sufficient for consciousness (Crick & Koch, 2003; Koch et al., 2016; Koch, 2019; Dehaene, 2014). This is obviously not without its challenges. As discussed, conscious experience is private. This makes it a particularly difficult phenomenon for outside observers, such as researchers, to assess. How can researchers connect patterns of neural activity with specific mental states if the mental state can only be known by one person? The standard attempt to resolve this challenge is heavy reliance on participants to report and describe their conscious experiences. For example, visibility of a stimulus may be manipulated so that brain activity associated with unconscious vs. conscious visual perception can be compared. To determine when the stimulus is consciously perceived so that corresponding brain activity can be isolated, participants are often instructed to make some type of report when they see the stimulus, like a simple keypress or verbal description.

Critiques

Reliance on participant report, however, has not been without critique. In a 1995 paper, philosopher Ned Block outlined a need to reassess science's conceptual understanding of consciousness, lest we fail to avoid a potentially catastrophic methodological and conceptual error. He proposed a distinction between *access consciousness* and *phenomenal consciousness*. *Access consciousness* refers specifically to "...availability [of conscious contents] for use in reasoning and rationally guiding speech and action," whereas *phenomenal consciousness* refers to the first-person, subjective, what-it-is-like experience that we have been discussing thus far (Block, 1995). Block's access-phenomenal distinction was a response to evidence for a possible discrepancy found between what participants were able to *report* experiencing and what they intuitively *felt* they had experienced. An iconic example of this discrepancy is provided by the classic 1960 experiment developed by George Sperling. In this experiment, participants were briefly shown a 3x4 array of letters and were asked to list as many letters as possible from memory. Despite claiming they had seen it all, participants could accurately report only 3-4 out of 12 letters. Intriguingly, participants were able to report 3-4 letters from any row that they were cued to report (even if the cue came after the briefly flashed letters were gone), which may suggest that a detailed and accurate representation of the letters had at some point been constructed, and participants' intuition was correct: they may have consciously perceived all of the letters in detail even if only for a fleeting instant (Block, 2011; Sperling, 1960). These results led Block to argue that "perceptual consciousness overflows access" (Block, 2011). Because participant report is currently the most reliable way to confirm states of consciousness, Block argued that much of consciousness research only examines conscious access, potentially missing the what-it-is-like that consciousness science seeks and claims to address (1995).

The current study

The current study directly addresses Block's assertion that phenomenal consciousness overflows access through a systematic literature review of empirical papers centered around the question: is perception rich or sparse? This topic was chosen because *perceptual* overflow can be seen as a manifestation of *phenomenal* overflow, and research on conscious visual perception is abundant.

What does it mean for perception to be "rich" or "sparse"? Based on introspection, most individuals report perceptually rich visual experiences—meaning they are full of *detail* throughout the visual field. However, as shown by the Sperling paradigm (1960), this same detail is often un-reportable. In other words, we seem to be able to report *that* our perceptual experiences are rich in detail, but we are often unable to report the details themselves. Does this mean that the rich detail is not experienced at all, and our impression of richness is illusory (perception is sparse), or that the details are experienced but cannot be accessed in many cases due to limitations in cognitive processing (perception is rich)?

Perceptual richness has become a key debate in consciousness research, with a myriad of complicated data generated by carefully designed laboratory experiments. Confusingly, many studies that seek to examine this issue provide data that can be interpreted in favor of either side of the debate, constantly calling for rigorous refinement of both method and interpretation. Consider two competing interpretations of the results from Sperling's 1960 study: in favor of "team-rich," participants may have had a detailed yet fleeting conscious experience of all of the letters in the array, however their attention was diverted by the task of remembering and reporting the letters and thus the post-cue was useful for helping them to reorient, access, and report a specific subset of the contents of their experience. In favor of "team-sparse," participants may have initially represented 3-4 letters in high detail (wherever they happened to be focusing their attention), and the peripheral letters as a low-definition "gist" or "average." While this may have given them the impression of rich detail throughout the array, it was only once the post-cue manipulated their attention that the full, detailed representation of a specific subset was formed, allowing them to accurately report these contents.

Despite the interesting evidence and interpretations on each side of the debate, determining conclusively whether perception is rich or sparse is perhaps a simplification of the issue. For example, if laboratory evidence were to one day fully conclude that perception is sparse, that does not address the question: why does it *feel* rich? Because no matter what the objective evidence shows, we still subjectively experience the appearance of richness—and isn't that subjective appearance what consciousness research claims to be all about?

This issue can be understood with the following analogy. For much of human history, we assumed the sun revolved around the earth, because of how it appeared. Copernicus then famously discovered that this assumption based on appearance was incorrect, and that the earth in fact revolves around the sun, spinning on its axis. Following this discovery, the appearance did not change: to this very day, it still looks to all of us as if the sun moves across the sky around the earth, despite collectively knowing better. The scientific study of consciousness is often approached as if there is only one question: does the sun revolve around the earth or not? In other words, are things as they appear to be? While this is certainly a necessary question, what do we do if things aren't as they appear to be? Do we throw up our hands and yell "We've done it! We've proved that consciousness is an illusion!"? Seeing as consciousness *is* the appearance, if science wishes to satisfactorily describe consciousness, the least it can do is not forsake its accountability to the subjective *appearance* of things. Of course, this analogy is not perfect, and there are key issues related to consciousness that it does not address. However, the hope is that conscientious adoption of this approach will result in a discussion of perceptual richness that is more thorough and satisfying than typically encountered.

A convenient outcome of this review would be to find sufficient empirical evidence to conclude that perception is rich—things are as they appear to be. However, because there is so much conflicting evidence, it is wise to prepare for this not to be the case. Keeping this in mind, I prepared to orient this review not around just one question (is perception rich?) but two questions, with the second question being: why does perception subjectively seem rich? To avoid presumption, the precise way in which this question was addressed was not finalized until completion of the review corresponding to the first question. Results of the initial review informed the direction taken to address the second question.

Relevance & significance

Block's phenomenal/access distinction is still debated, yet it has not failed to shape a considerable portion of the theoretical conversation around consciousness research. Its influence has resulted in major theories of consciousness taking up the task of explicitly explaining phenomenal consciousness--lest they lose face by repeating the mistake of failing to consider it. Because phenomenal consciousness is considered the more "pure" or "authentic" form of consciousness--i.e. what science is *really* getting at when it says "first-person, subjective experience" --there is significant pressure in consciousness research to clearly account for it. Thus, a review of empirical research on issues specific to the challenges posed by phenomenal consciousness as a concept (perceptual overflow, in this case) that considers the relevant tensions and challenges is extremely pertinent.

Theories of consciousness

Additionally, because of the popularity of the access/phenomenal distinction, many leading scientific theories of consciousness have at least commented on if not made clear predictions regarding phenomenal overflow and perceptual richness. These ideas and predictions are considered throughout the review to increase its relevance to contemporary theory and research.

Global Workspace Theory

Global Workspace Theory (GWT) is a theory of access consciousness (Newman & Baars, 1993). Depending on who you ask, this may be a strength or a weakness. To those who believe there is a phenomenal consciousness that overflows access, GWT is an insufficient model. GWT defines what is consciously experienced as that which has been broadcast in the global workspace (Dehaene & Naccache, 2001). According to GWT, when information is broadcast, it becomes available to local processing by cognitive systems involved in memory, language, and decision-making. This definition of consciousness as the availability of information is almost identical to Block's definition of access consciousness (Mashour, 2020; Block, 1995). If we only experience what is available to cognitive systems such as language, memory, and decision-making, then we only experience what is accessed, and phenomenal consciousness is illusory. This has resulted in criticism of the theory, claiming that it is insufficient for addressing key subjective experiences such as phenomenal overflow and perceptual richness. In defense of GWT's ability to explain such experiences, Shanahan & Baars (2007) argue that access is often incorrectly conflated with working memory (a system that can preserve information temporarily if it is needed for use in an ongoing task). For example, they argue that information in the workspace is inherently conscious, whereas information stored in working memory can be unconscious. Someone who is told to quickly memorize a string of numbers will encode as many numbers as possible into working memory. Until asked to recall the numbers, they are not conscious of them (unless they are rehearsing the numbers over in their head). While the numbers remain available in working memory for the duration of the task, the numbers are not consciously perceived (broadcast) for the entire duration. Further, Shanahan & Baars (2007) argue that in order to enter working memory, information must first pass through the workspace, and thus consciousness. There is a possibility that richly detailed information may be globally broadcast to local processors at some point, but by the time this information has been reported, additional details may have been lost, never making it from short-term to working-memory (Shanahan, 2005 in Shahanan & Baars, 2007). Despite this possibility of a fleeting yet detailed experience, GWT is clearly on the sparse side of the rich vs. sparse debate. A version of the general GWT, the global neuronal workspace (GNWT) hypothesis specifically claims that we can only consciously experience one (or at best a few) things in any given moment (Dehaene, 2014). Importantly, GNWT proposes extensive unconscious processing (including rich and detailed unconscious perception), accompanied by much more limited conscious processing. In some sense, this access-only theory (Naccache, 2018) rejects Block's notion of a rich but fleeting phenomenal consciousness and replaces it with a rich but fleeting unconscious.

Integrated Information Theory

GWT and Integrated Information Theory (IIT) are perhaps the two most well-known and hotly debated theories of consciousness. Notably, IIT is not a functional theory of consciousness. Instead of explaining *what* the brain does, IIT places importance on *how* the brain does things, theorizing that consciousness is not a result of any one process, but rather a phenomenon inherently linked to how the brain operates. Further, IIT can be seen as an identity theory, as it postulates that consciousness is identical to *phi*–a term referring to maximum integrated information. Unlike GWT, IIT is known for being a theory of phenomenal consciousness. Proponents of IIT believe perceptual richness is real and non-illusory, advocating for paradigms that don't rely on the use of report (access) so that consciousness can be more accurately studied (Tsuchiya et al., 2015; Haun et al., 2017;). Additionally, introspection is valued in IIT as a valid tool for learning truths about consciousness, including that conscious perception is rich (Haun et al., 2017; Koch, 2019). While not emphasized by this theory explicitly, IIT's adherence to the rich view of conscious perception necessarily implies that it considers the extent of unconscious processing to be much more limited than competing theories on the sparse side of the debate.

Sensorimotor/enactive theory of perception

Sensorimotor (and enactive) theories of perception offer a novel and compelling explanation for why perception *seems* rich. Such theories propose that the conscious perception of physical stimuli involve *sensorimotor contingencies*—that is, information regarding how the visual experience of a given percept will change as a result of different movements (head, eye, body, etc.) (Noë, 2004). Sensorimotor theory also acknowledges the complexities of intuition and introspection. They acknowledge that we understand that we don't experience as much detail in our periphery; "We don't take ourselves to experience all environmental detail in consciousness all at once. Rather, we take ourselves to be situated in an environment, to have access to environmental detail as needed by turns of the eyes and head and by repositioning the body" (Noë, 2004, p. 59). Sensorimotor theories argue that this is why it *seems* as if perception is rich—we know we will find richness if we look for it, and this contingency is directly tied to what we perceive.

Attention Schema Theory

To understand Attention Schema Theory (AST), we must first clarify the distinction between attention and awareness. Attention is generally considered an unconscious process, whereas awareness is conscious. Attention can simply be defined as the selection of certain sensory, cognitive, etc. information over other information to undergo further processing. However, this processing does not necessarily result in consciousness (an example of this would be subliminal priming, where participants are presented with stimuli so briefly that it does not reach consciousness, yet results show that performance is still affected).

Throughout the rest of this paper, this is the distinction between attention and awareness that will be used.

AST claims that consciousness is the result of an *attention schema*–a loose, rough internal model of the attentional system itself. This sort of model can be understood as it relates to the *body schema*. For example, Graziano et al. (2020) ask readers to consider the human arm. If the arm is amputated, some patients experience what is known as a "phantom limb" where they continue to feel pain and sensation as if their arm were still physically there. Such an experience is able to occur thanks to the body schema–a model of the body and how it works that is represented by the brain in order to more finely control movement and positioning. AST claims that attention is modeled in the same way, and that our feeling of "consciousness" is a byproduct or artifact of this model. Just like the body schema (a rough caricature of the real body) is useful for controlling the body, the attention schema (a rough caricature of the real attention mechanisms) is useful for controlling attention. What AST contributes to the discussion of perceptual richness has less to do with the theory explaining why perception subjectively feels rich and more to do with why we believe we have a subjective experience that *feels* like anything at all. The attention schema contains information that allows the system to claim that a specific experience has "…a subjective, what-it-feels-like component" (Graziano, 2020). AST cuts underneath issues of phenomenality and perceptual overflow and asks not what is subjectively experienced or why, but rather *how* can a machine like the brain come to know or believe it has such experiences in the first place?

Methods

Protocol

This review followed the Preferred Reporting Items for Systematic reviews and Meta-Analyses extension for Scoping Reviews (PRISMA - ScR) guidelines (Tricco et al., 2018).

Eligibility criteria

Included publications must have been peer-reviewed, available in English, and published between 1995 and 2021 to ensure up-to-date work on this topic while still allowing room for some foundational research from the mid-late '90s when the issue of perceptual overflow was first introduced (Block, 1995). Studies on non-human animal models, children (<18) or the elderly, that examine psychopathologies, and/or that involve psychoactive substances were excluded due to inability to reliably compare results across studies. Because the issue of perceptual richness is so heavily informed by information gained via introspection, the risk of including studies with participants who may show systematic differences in introspection compared to adult human primates was sought to be avoided. A full list of inclusion and exclusion criteria can be found in Appendix A.

Information sources

Searches were conducted on the Scopus and PsychInfo databases. Results were downloaded to Zotero, where duplicates were automatically identified and manually removed.

Search method

Final search terms were chosen according to the idea that it is better to discard many irrelevant papers than to miss out on relevant papers. Preliminary searches were conducted on Scopus, PsychInfo and Google Scholar to identify key relevant papers. These papers provided a way to measure the effectiveness of different search term combinations (i.e., how many of the key papers were present in the results using *xyz* terms?). The same groups of final search terms were used on both Scopus and Web of Science, an example is given below:

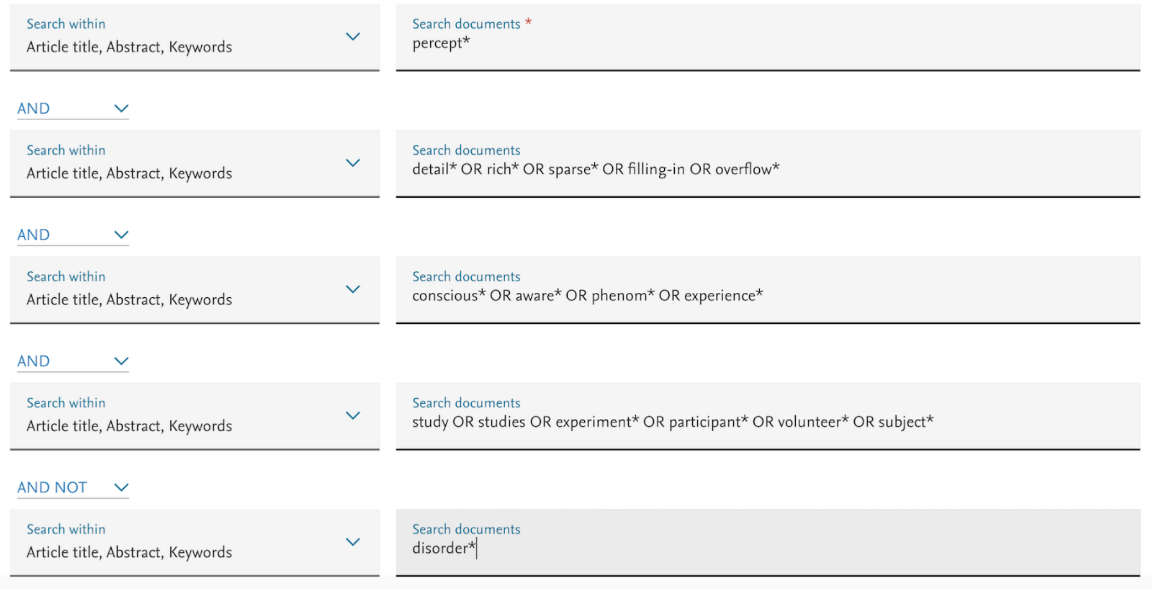

Search within
Article title, Abstract, Keywords
Search documents *
percept*

AND

Search within
Article title, Abstract, Keywords
Search documents
detail* OR rich* OR sparse* OR filling-in OR overflow*

AND

Search within
Article title, Abstract, Keywords
Search documents
conscious* OR aware* OR phenom* OR experience*

AND

Search within
Article title, Abstract, Keywords
Search documents
study OR studies OR experiment* OR participant* OR volunteer* OR subject*

AND NOT

Search within
Article title, Abstract, Keywords
Search documents
disorder*

Figure 1: Scopus search terms
Shows search terms as entered into the Scopus database. The asterisk "*" can be replaced by any letter or series of letter, so "percept*" can result in "perception," "detail*" in "details" or "detailedness," etc.

The rationale for each term is as follows:

1. Percept*: Because the topic of the current study is perception, relevant studies should contain some variation of the term perception in the abstract, title, or keywords.
2. Detail* OR rich OR sparse OR filling-in OR overflow*: Because perception is a common word used in many contexts, this string of terms aims to narrow the scope of the results to those dealing specifically with perceptual overflow.
3. Conscious* OR aware* OR phenom* OR experience: We included these terms because we are interested in perceptual overflow specifically as a topic in the broader field of consciousness research. These terms are used interchangeable with "consciousness" in the literature.
4. Study OR studies OR experiment* OR participant* OR volunteer* OR subject*: Because we want to review empirical papers involving human subjects only, these terms were chosen with the aim of excluding papers not fitting that criteria (i.e., reviews, animal models, etc.)
5. (And not) disorder*: In pilot searches, many papers were returned that regarded various health and psychological disorders (i.e., disorders of perception, patients' perceptions of specific disorders). This keyword was excluded to filter out such papers.

Once the above terms were entered into the database, additional filters were applied. These filters limited the results to papers relevant to the fields of psychology and neuroscience, papers available in English, papers that are peer-reviewed, and papers published from 1995-2021. The application of these filters can be seen in the figure below:

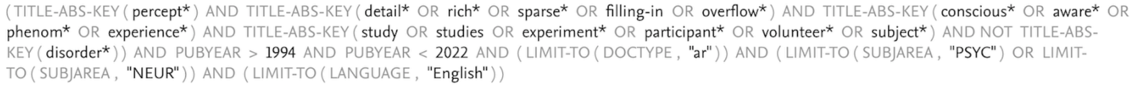

```
( TITLE-ABS-KEY ( percept* ) AND TITLE-ABS-KEY ( detail* OR rich* OR sparse* OR filling-in OR overflow* ) AND TITLE-ABS-KEY ( conscious* OR aware* OR
phenom* OR experience* ) AND TITLE-ABS-KEY ( study OR studies OR experiment* OR participant* OR volunteer* OR subject* ) AND NOT TITLE-ABS-
KEY ( disorder* ) ) AND PUBYEAR > 1994 AND PUBYEAR < 2022 AND ( LIMIT-TO ( DOCTYPE , "ar" ) ) AND ( LIMIT-TO ( SUBJAREA , "PSYC" ) OR LIMIT-
TO ( SUBJAREA , "NEUR" ) ) AND ( LIMIT-TO ( LANGUAGE , "English" ) )
```

Figure 2: Scopus search terms & filters

Depicts all search terms in initial search and any additional filters applied

The final searches were conducted on January 8, 2022. The Scopus search resulted in 1,395 papers, 1,387 of which were downloadable to Zotero. The PsychInfo search resulted in 2,344 papers, for a total of 3,731 papers downloaded to Zotero between the two databases combined. After removing duplicates, 3,016 papers were screened for relevance. Relevance was defined as specifically referencing the rich vs. sparse debate in the title and/or abstract. Twenty-two papers fit this relevance criteria. An additional 4 papers were cross-referenced from the relevant papers, for a total of 26 papers included in this review. See Appendix B for a full bibliography of the 26 included papers.

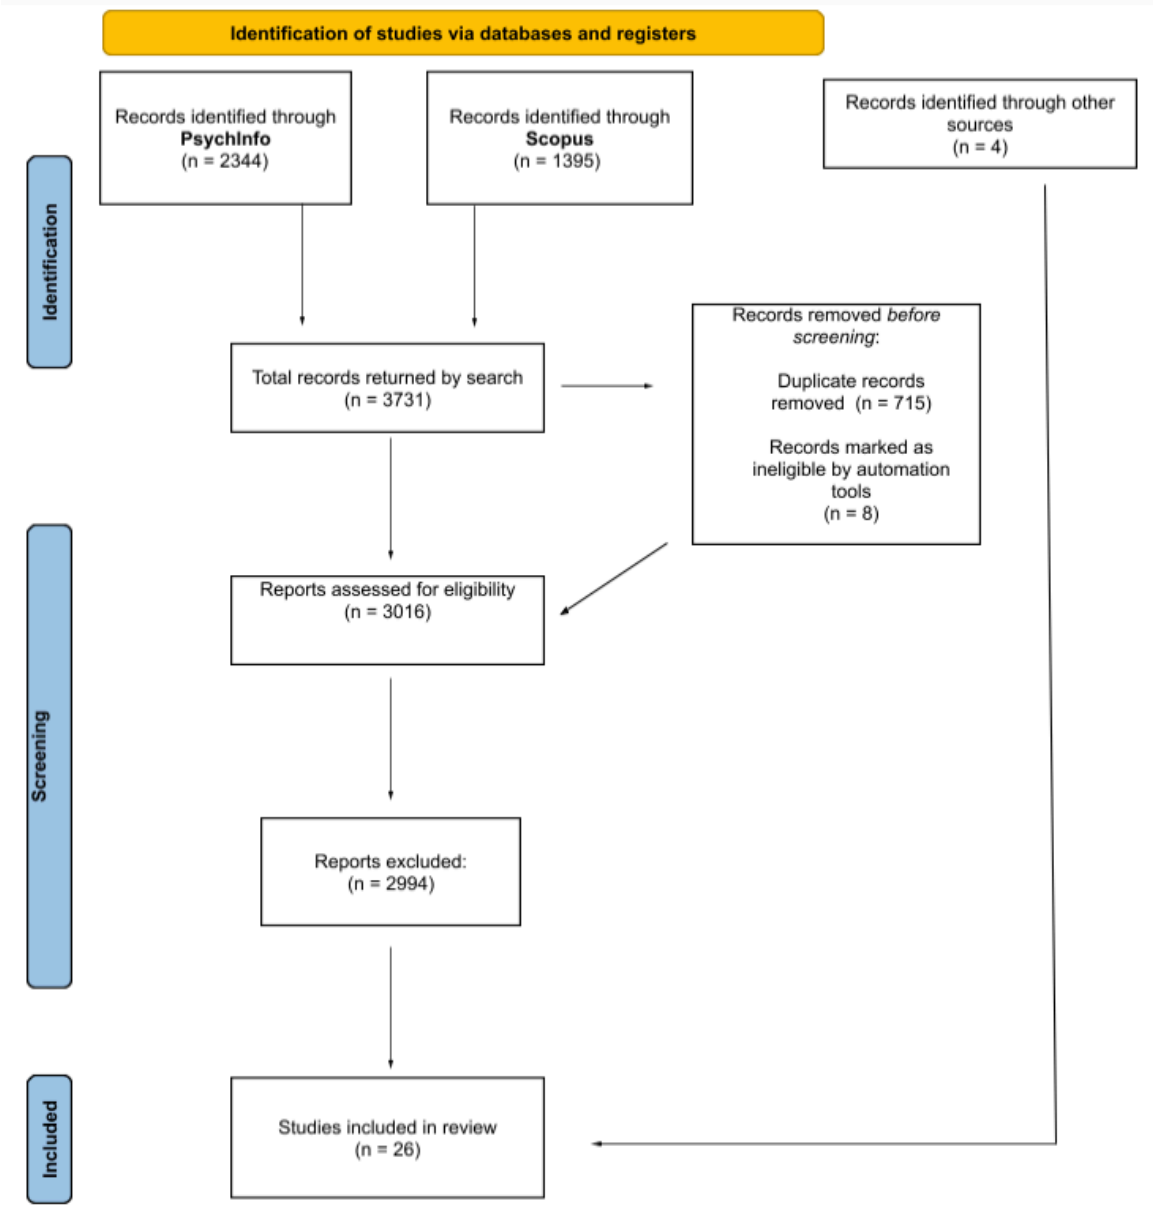

Figure 3: Search flow chart

A diagram depicting the selection process of the papers included in this review

Results

Color perception & summary statistics

Color perception has lent insight to the study of subjective richness. A common example is our perception of a rich, blue sky. The retinal periphery is composed of mostly rods (as opposed to cones), which do not differentiate color. Despite this, when we look at the cloudless sky, we see an endless expanse of blue that does not decrease in vibrancy further away from the center point of fixation. How are we able to perceive rich blue in our periphery despite the lack of color-differentiating photoreceptors?

Cohen et al. (2020a) used virtual reality (VR) environments to investigate peripheral color perception in real world settings. As participants freely explored VR environments with eye- and head-movements, the color of the periphery was desaturated such that color only remained within a 10°, 17.5°, 25°, or 32.5° radius from the center of fixation (determined via eye-tracking). In other words, a small portion in the center of the visual field would remain in-color, while the periphery would fade to black and white. By integrating online eye-tracking with the VR environment, wherever participants looked, the virtual world was colored, while large portions of the periphery were rendered colorless. An image has been included on the next page to provide a sense of what a significant portion of the visual field was desaturated.

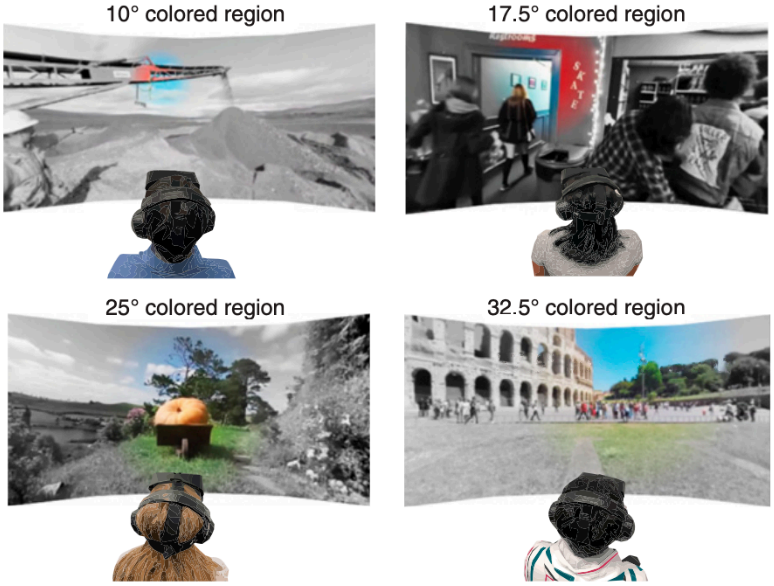

Figure 4: Stimuli from Cohen et al. (2020a)
This image segment reproduced from Cohen et al. (2020a) depicts the four different desaturation conditions participants experienced in the VR study.

Whether participants noticed the absence of color varied across conditions of how much of the visual field was desaturated. When 32.5° remained in color, 83% of participants failed to notice the desaturated periphery. When 25° remained, this number dropped to 65%. In the condition of 17.5°, 50% of participants failed to notice. Even when a mere 10° remained in color, a considerable number of participants (33%) failed to notice the lack of peripheral color. This finding was interpreted by Cohen et al. (2020) as evidence for sparse perception because of how many participants failed to notice the striking change.

The inability of subjects to detect achromatic regions has been demonstrated before. Similar in concept to Cohen et al. (2020a), Balas & Sinha (2007) presented participants with "color chimeras," created by applying a color saturation mask to natural or artificial scenes (example pictured on the net page). Color chimeras are designed to help determine how chromatic "filling-in" responds to different manipulations of a scene's summary statistics, by varying the saturation in different visual regions. Such manipulations can help researchers determine whether filling-in occurs via foveal/central-outward or peripheral-inward processes.

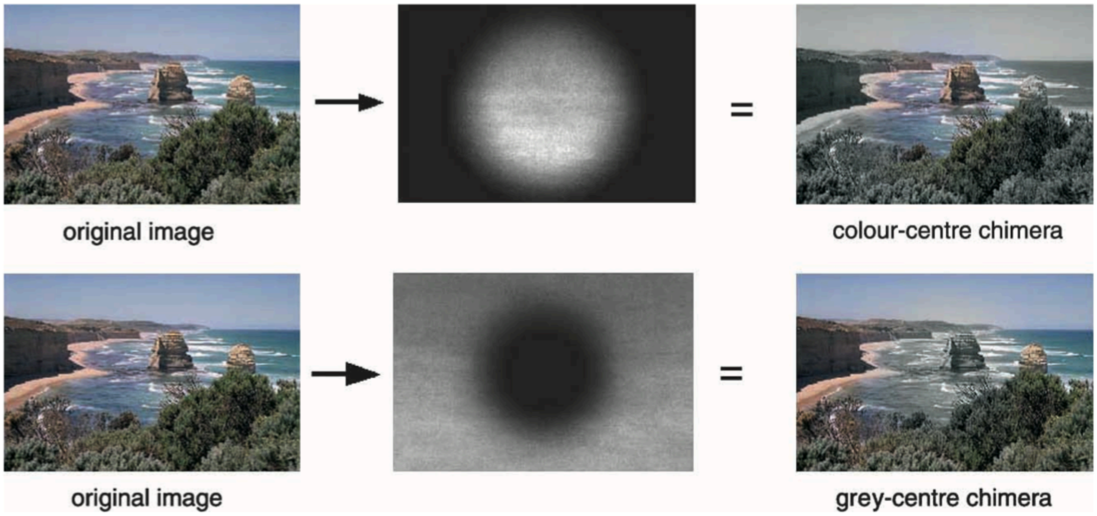

Figure 5: Color chimera stimuli
Stimuli reproduced from Balas & Sinha (2007). The saturation masks (center column) are applied to the original images (left column). The resulting image has a different saturation distribution.

Participants were presented with full-color, greyscale (black-and-white), and “mixed” (saturation masked) images, and were asked to respond via keypress whether they perceived the image as full-color, greyscale, or mixed. Results showed that participants tended to report center-color chimeras as full-color. This effect was associated with mask width (i.e., a larger colored center region resulted in more reports of a full-color image), with full-color reports reaching “nearly 50%” for the largest width (Balas & Sinha, 2007). Results were similar for grey-center images, and a similar main effect of mask size was found for grey-center images, the smaller the mask the more likely participants were to report a full-color image. Balas & Sinha (2007) saw this as suggesting color filling-in can be both a center-outward and peripheral-inward process. For grey-center images, there were also significant effects of “color” (mask) and “texture” (natural scene structure vs. artificial) manipulations. For example, falsely perceiving images as fully black-and-white did not occur as often as perceiving them as fully in color. Based on their results, Balas & Sinha (2007) conclude that color-spreading is likely to rely on *summary statistics*—an "average" or "gist" of the contents of a scene that the brain computes to help make perception more efficient. Summary statistics—also referred to as ensemble perception--will be discussed in more depth below.

Related to the phenomenon of "filling-in" is the *uniformity illusion*–the ability of central stimuli to cause peripheral stimuli to be perceived inaccurately. Suárez-Panilla et al. (2018) presented participants with a grid of Gabor patches (pictured below) and instructed participants to press a key when the grid appeared uniform. During the control session, the presented stimuli were actually uniform, so both subjective and objective uniformity occurred. Meanwhile, in the illusion session, only subjective uniformity occurred. The illusion was fairly successful, with participants subjectively experiencing uniformity in 27% of trials.

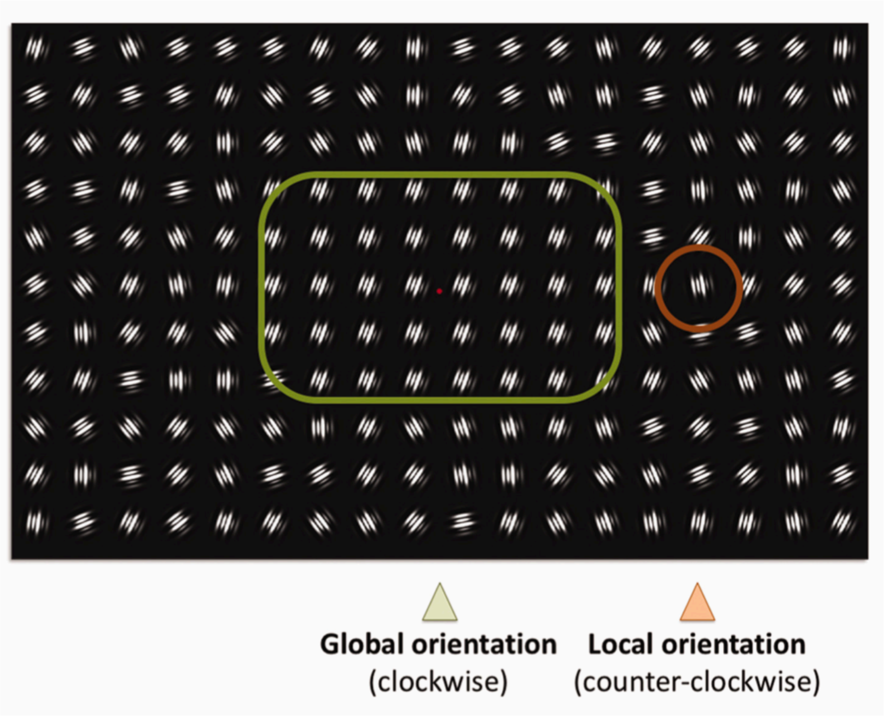

Figure 6: Array of Gabor patches
A segment of an image reproduced from Suárez-Panilla et al. (2018) displaying an example of their stimuli. Under the uniformity illusion, peripheral patches (those outside the green rectangle) would be perceived as matching those within the rectangle.

After viewing either the objectively or subjectively uniform array, participants were presented with a single Gabor patch from the preceding array and were asked to report its orientation. The location of the patch didn't change from one presentation to the next, but the orientation of the single test patch was changed from its original to intermediary between global (orientation of the illusory patches, consistent with those in the green rectangle) or local (the actual orientation of the single patch as it was in the previous array). Suárez-Panilla et al. (2018) used this measurement to assess whether the illusion was driven by higher- or lower-level processing. Participants exhibited a bias in reporting test patch orientation as local (consistent with their actual orientation, not their illusory orientation). Suárez-Panilla et al. (2018) interpreted this result as evidence that the illusion is driven by higher-order (more categorical and conceptual) rather than lower-order (more brute detail-ordered, lines, contours, etc.) processing. If the illusion was driven by lower-level visual areas associated with processing detail such as orientation, Suárez-Panilla et al. (2018) argue that participants would report the orientation of the test patch as consistent with the illusion. Because participants were still biased to report the objective orientation of the test patch, it seems likely that the illusion does not rely on and affect lower-level visual processing.

Otten et al. (2017) tested the applicability of the uniformity illusion to a wide array of visual features such as shape, luminance, motion, orientation, pattern, and identity. Examples are pictured below. Participants were initially presented with the center "patch," and the periphery faded in over the course of 2 seconds. The image remained on the screen until the participant reported uniformity (or would expire after 10 seconds). Similar to Suárez-Panilla et al. (2018), some trials involved objective uniformity. In such trials, the peripheral content would slowly become objectively uniform to central content, mimicking the effects of the illusion. For identity and pattern uniformity, there was an added trial condition where objective uniformity was present at and maintained from initial presentation. This allowed Otten et al. (2017) to establish a baseline measurement of how long it took participants to detect uniformity in a more general sense.

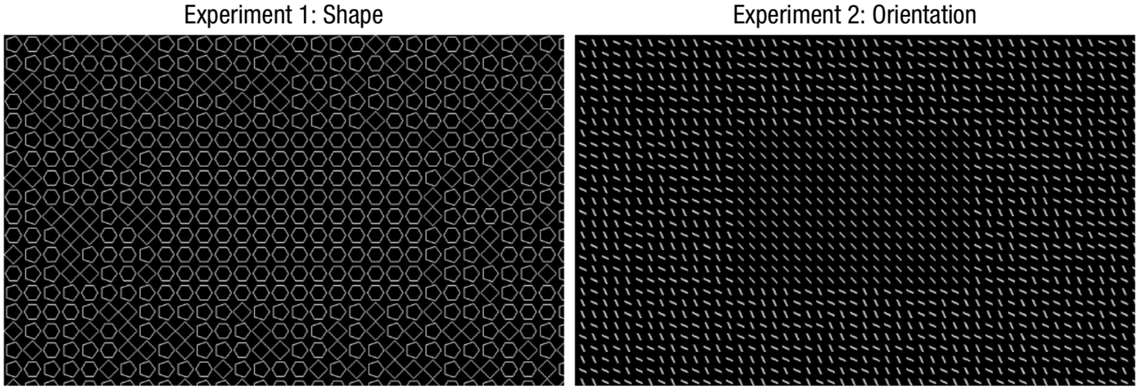

Figure 7: The uniformity illusion
Examples of stimuli used in Otten et al. (2017). Try it yourself! First, note that the pattern of shapes differs in the center vs. the periphery. Then, fix your gaze on the center of the image until you notice the pattern become uniform across the entire image.

The illusion was successful for all stimulus types, with at least 2 seconds needed for participants to experience the illusion once they began looking at the image. One exception to this was for stimuli in the identity condition, for which subjective uniformity occurred almost immediately. In the experiment using identity stimuli (letters), participants' confidence in the illusion was also measured. Otten et al. (2017) found that confidence was similar for both real and illusory uniformity, which they interpreted in support of the strength of the illusion as it suggests that the illusory effects are almost as strong as perceiving the real thing.

The studies discussed so far emphasize participants' inability to see things "correctly," or their tendency to perceive things that are not there. This highlights what could be seen as perceptual deficits. However, additional research has been done on peripheral color perception that claim to find evidence for rich and accurate, yet fleeting, perception of peripheral content. Bronfman et al. (2014) examined a summary statistic known as color diversity. Like the Sperling (1960) paradigm, participants were briefly presented arrays of random letters and were tasked with reporting letters from a specific "cued" row. In this version of the Sperling paradigm, the relevant row was always *pre-cued* (i.e., indicated before array onset), and the letters were colored, either with high color diversity (color of each letter sampled from all over the color wheel) or low color diversity (color of each letter sampled from a narrow sub-region of the color wheel). An example of the color-diversity variations in stimuli is pictured below. In addition to reporting letters from the cued-row, on some trials, participants were also asked to perform a secondary task of reporting the color diversity (high or low) of either the cued or a non-cued row. Replicating Sperling's (1960) results, participants could accurately report ~2-4 letters from the cued row, but in Experiments 1-4, participants also showed above-chance accuracy in judging the color-diversity of letters, both in the cued and non-cued rows. Accuracy of color-diversity judgments increased when participants were instructed to report color diversity before reporting the letter (Experiment 2), which Bronfman et al. (2014) interpreted as evidence that endogenous attention increased perception of color diversity.

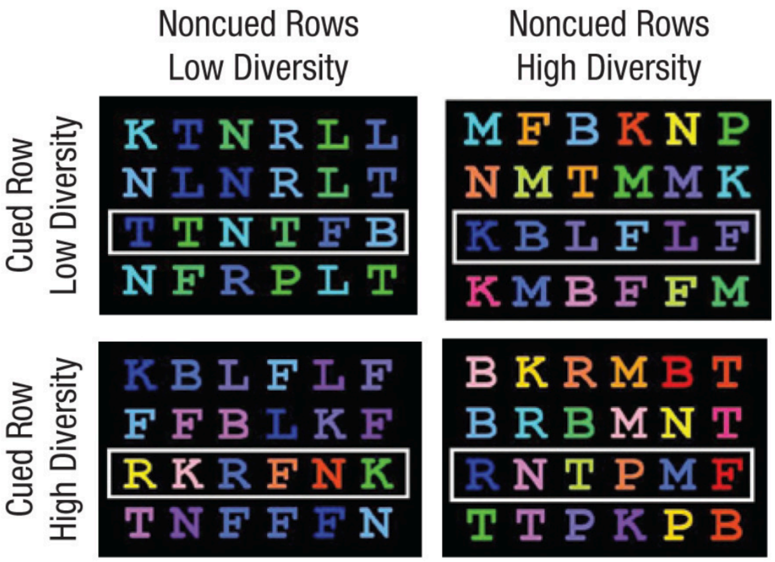

Figure 8: Color diversity of stimuli in Bronfman et al. (2014)
An image segment reproduced from Bronfman et al. (2014) depicting their stimuli. As shown, stimuli could have high color-diversity overall (bottom right), low color-diversity overall (top left), or differences in color diversity between the cued and non-cued rows (bottom left and top right).

To further ensure that this performance could not be due to unconscious processing, Experiment 5 replicated the procedure with masking and brief presentation (16.7ms). Participants were also asked to rank their visibility on a scale: 1 did not see the colors, 2 partially saw the colors, or 3 saw the colors well. Bronfman et al. (2014) found that performance on trials where participants chose 1 was statistically at-chance (47%), but significantly above chance on trials where 2 (68%) or 3 (84%) were chosen. This was taken to suggest that above chance performance in the earlier experiments cannot be explained by unconscious processing alone.

Bronfman et al. (2014) saw their study as evidence for rich perception because it revealed an ability to accurately estimate color-diversity (even in less-attended, peripheral regions of the display), which they saw as a remnant of a fleeting experience of all (or most of) the individual original colors. This interpretation hinges on the assumption that the perception of summary statistics relies on perception of individual elements. If this assumption is correct, then the results of Bronfman et al. (2014) may indeed support rich perception. In a follow-up study, however, Ward et al. (2016) investigated the validity of this assumption. Experiment 1 replicated the procedure of Bronfman et al. (2014) with a modification to the 1-3 scale the original study used in their fifth experiment. The questions posed by Ward et al. (2016) more directly probed participants' perceptions of individual colors, asking:

1) I had no sense that any of the letters had any color at all
2) I had a vague sense that the letters were colored in general, but I didn't clearly perceive the individual colors of individual letters.
3) I had a clear sense that the letters were colored in general, but I didn't clearly perceive the individual colors of individual letters.
4) I had a clear sense that the letters were colored in general, and I could also clearly perceive the individual colors of individual letters.

Results showed that participants chose option 1 rarely, and options 2-4 about equally. In replication of Bronfman et al. (2014), on trials when participants reported no sense of the colors at all (option 1), performance on the color-diversity judgment task was at chance. In addition, Ward et al. (2016) found no significant difference between performance on trials where participants chose option 2 (63.29%) vs option 4 (73.53%).

In a second experiment, Ward et al. (2016) used change blindness to test performance on a task where perception of individual colors was indeed required. On half of the trials, the colors of the unattended letters were reshuffled during the 17ms blank period. This caused a change to the individual elements from prechange (150ms) to postchange (150ms), however because no new colors were added or subtracted to the array, the overall diversity statistic was preserved. Participants were then asked if they noticed the change. None of the 12 participants noticed the change. Performance on the main task did not suffer (accuracy at estimating color diversity remained above-chance). A third and final experiment replicated the procedure in Experiment 2, however this time the pre-change display duration was increased from 150ms to 650ms. Above-chance diversity estimation was again replicated, however no participants noticed the change, despite the longer exposure.

Ward et al. (2016) interpreted their results as a challenge to the fundamental assumption of Bronfman et al. (2014): that perception of summary statistics implies the perception of individual elements. The conclusions Bronfman et al. (2014) drew in support of rich perception rest on this assumption, but Ward et al.'s (2016) experiments showed that summary statistics (at least the color diversity statistic) can be accurately judged (at above-chance levels) despite complete change blindness of the individual elements.

Change blindness

Researchers often reference *inattentional blindness* (IAB) and *change blindness* (CB) as examples of failure to notice details we intuitively feel we perceive. IAB and CB occur both in- and outside of laboratory settings, strengthening the argument for sparse perception. *Inattentional blindness* refers to the failure to notice salient, central stimuli because of a lack of or misallocation of attention. Simon & Chabris (1999) used a highly creative paradigm to demonstrate IAB for complex, realistic stimuli. Participants were shown a video in which two teams (3 players each) passed a basketball back and forth to each other. One team wore black shirts, and the other team wore white shirts, and participants were instructed to count the number of times players on the black or white team (each participant assigned to either "black" or "white" condition) passed the ball. This task was the easy condition. The hard condition required participants to count how many times the assigned team performed specific types of passes (bounce, aerial, etc.). There were two scenarios researchers were curious whether people would notice while performing their assigned task: 1) a woman holding an open umbrella would walk from left to right across the scene 2) a shorter woman in a full-body gorilla costume would perform the same action. Videos were presented in conditions of transparent and opaque. Across all conditions, a shocking 46% of participants failed to notice the key scenario. Participants were more likely, however, to notice the event while performing the easy task than the hard task. The woman holding the umbrella was noticed more often than the gorilla. There was no difference in color condition on noticing the woman with the umbrella, but participants were more likely to notice the gorilla when paying attention to the black team than when paying attention to the white team.

In a final, additional scenario (perhaps the most well-known), participants were asked to perform the easy task on the white team while the woman in the gorilla costume walked into the center of the players, turned to face the camera, and banged her fists against her chest before walking off. The entirety of this critical event lasted a full 9 seconds. Shockingly, 50% of participants failed to notice the gorilla. To watch the video shown in the experiment, visit this link: https://www.youtube.com/watch?v=vJG698U2Mvo

CB is similar to IAB, but instead of failing to notice the appearance of an unexpected object or event, CB involves a failure to notice a change, even if one is "on the look-out" for potential changes. In CB, such failures occur primarily due to a visual disruption (such as an eye-movement, blink, or sudden disruption to the scene itself). Early research on both change and inattentional blindness was in part motivated by practical applications such as improving cancer screening, TSA baggage screening, and prevention of traffic accidents. O'Regan et al. (1999) conducted a study that sought to simulate mud splashing against the windshield of a car, this simulation is pictured on the following page. Participants were tasked with detecting changes between pairs of natural scenes, however a "mudsplash" would be superimposed over parts of the image precisely at the time of the change (but never covering the location of the change). Changes were categorized into "central interest" and "marginal interest," but were controlled for equal size and salience. Results showed that participants almost always detected central-interest changes on the first round, but marginal interest-changes were detected at the second round or later. In a follow-up experiment, O'Regan et al. (1999) modified the stimuli so that instead of multiple "mudsplashes," there was only one, and it covered the location of the change. This modification was to ensure that the masking rectangle was not disrupting representation of the scene. Participants were once again able to immediately detect central-interest change, however marginal interest changes went undetected.

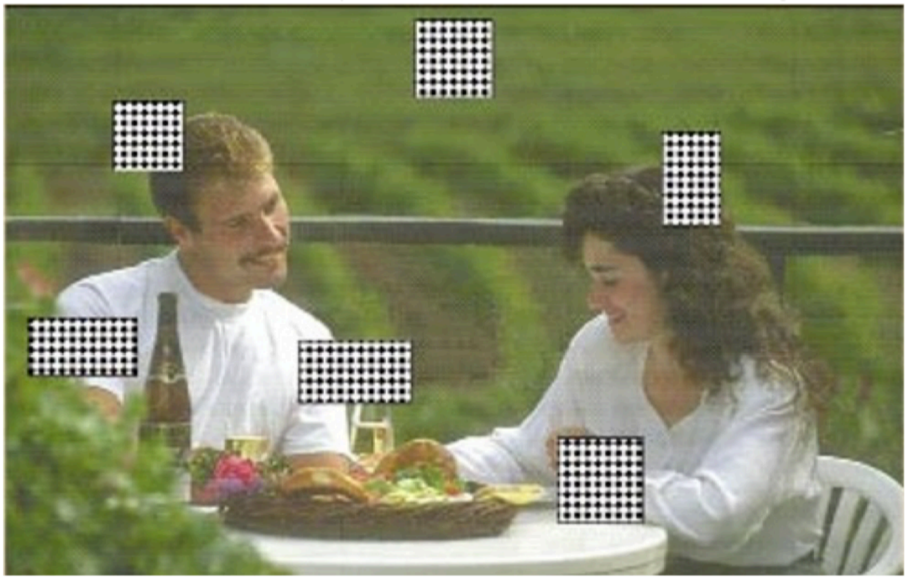

Figure 9: Change blindness & "mudsplashes"
An example of the stimuli used in O'Regan et al. (1999). The sudden appearance of black-and-white checkered patches aimed to simulate mud suddenly splashing onto the windshield of a car.

O'Regan et al. (1999) interpreted their findings as evidence for a sparse representation of the external world, because of such a striking failure to notice environmental detail. They further hypothesize that the external world may act as a sort of "outside memory," and that our feeling of richness comes from knowing we can immediately access rich detail through series of eye and head-movements so effortless that we are often not even consciously aware of making them.

While CB has been widely cited as evidence for sparse perception, there are other interpretations of the phenomenon. Mitroff et al. (2004) cite three common explanations for CB. The first is a failure to represent (encode) the pre-change (original) image in any way. For example, a participant fails to detect a change because there was no representation of (certain details of) the first image to access and use to compare the second image against. The second explanation is that the pre-change image is sufficiently represented, but then is quickly forgotten or overwritten by successive information. In this case, a participant views and represents the prechange image, but then the appearance of the postchange (changed) image interrupts this fragile representation, rendering the participant unable to accurately detect a change. The final explanation Mitroff et al. (2004) refer to is a mere failure of comparison. A participant accurately represents both the prechange and post-change image and indeed compares them, however they do so ineffectually, failing to detect the change in time.

Determining which explanation is accurate would provide crucial insight into the nature of mental representations, including their richness or sparseness. If the first explanation is true, perception is sparse because not much detail is represented. If the second is true, then perception is rich but representations are extremely fragile, explaining why people feel as if they experience more than they can report. If the third is true, then perception is rich and failures of access are due to some other cause than fragile representations.

Mitroff et al. (2004) attempted in four experiments to determine if representations were still preserved in instances of failure to detect a change. If this were the case, then it may be concluded that CB is a failure of comparison, because a representation of the pre-change image is indeed preserved. To assess this, Mitroff et al. (2004) compared performance against two different boundaries of chance: lower-chance (50%, standard) and upper-chance (a more conservative measure calculated individually for each participant). The upper-chance boundary was intended to reflect the maximum level of accuracy that could be obtained *without* sustaining representations of both the pre- and post-change displays. Thus, Mitroff et al. (2004) interpreted performance above this boundary as evidence for the existence of both representations and thus comparison.

Experiment 1 consisted of a standard change detection task using an array of line drawings. On 90% of trials, one of the drawings was replaced (change trials). To assess quality and content of representation on every trial (even when a participant failed to detect a change), participants were asked a series of three questions after each trial. All three questions presented participants with two objects and asked them which had been present. The first question presented the prechange object and a novel object, the second presented the postchange object and a novel object, and the third was a novel object and a randomly chosen object from the tested array not implicated in the change. Results of the first experiment revealed participants were fairly adept at detecting changes (65.88%)

and were able to perform better than both the lower boundary of chance and the upper boundary of chance. Mitroff et al. (2004) concluded from this above upper-chance performance that participants were able to access representations of both pre- and post-change information, even on trials where they failed to detect a change, suggesting that CB may indeed result from a comparison failure.

Experiment 2 followed the same procedure as Experiment 1 with the following exception: the third follow-up question was removed, and participants were explicitly informed that all follow-up questions would involve objects involved in the change (recall that the third question in Experiment 1 did not involve any changed objects, only a novel object and a random unchanged object). Results replicated Experiment 1 (no significant difference in detection, performance above upper boundary). These results ensured that participants were not adopting a strategy of ignoring the change detection task in lieu of actively trying to encode as many items in the array as possible in order to perform well on the questions. If the only items questioned were those changed, participants would have no motive to focus on any other task. Experiment 3 required participants to answer the change detection question *before* the memory question. Results revealed no significant difference in detection compared to the prior two experiments, however accuracy for the prechange question decreased in unaware trials, suggesting that while the prechange information may be represented, the stability and the robustness of the representation can reasonably be called into question. This is apparent due to performance suffering due to the interruption of being asked the detection question first. Finally, Experiment 4 simply assessed whether the memory questions swayed performance on the change detection task by removing the memory questions altogether. Results showed no significant difference in accuracy, confirming that the probes had no effect on the main task. Overall, these follow-up assessments were taken by Mitroff et al. (2004) to confirm the reliability of their main finding, which they claim is that CB is the result of a comparison failure (Experiment 1, above paragraph).

An additional claim for the existence of preserved representations in CB comes from Simons et al. (2002). In a real-world scenario, each participant would walk down a path on Harvard campus, where a woman in athletic clothes would approach them and ask for directions. A group of eight confederates (undercover researchers) would pass in-between, talking amongst themselves, and either take from the woman (removal condition) or give to the woman (addition condition) a standard orange basketball. Once the participant finished providing directions, they were immediately debriefed as to the nature of the study and asked if they had noticed what had occurred. In the addition condition, only 1 out of 6 participants noticed the change upon first questioning, and an additional participant reported noticing after initial questioning. In the removal condition, 2 out of 5 participants noticed the change upon first questioning, 3 noticed only when asked about the basketball specifically. Like Mitroff et al. (2004), Simons et al. (2002) took this to mean that participants may "show" CB even when they have a representation of the scene. When asked specifically, participants did recall the basketball, even though earlier, when asked indirectly, they did not indicate such.

A second experiment replaced the basketball with a novel red and white striped ball, allowing researchers the opportunity to ask participants to recall specific details about the changed item. Simons et al. (2002) also added control conditions where no change occurred (the ball was either present or absent throughout the exchange) and control probes (participants were asked about changes that never occurred) to establish a baseline sense of how trustworthy participants' reports were more generally. Only 2 out of 14 participants in the removal condition and 1 out of 13 in the addition condition failed to report the change at any point during the questioning. While the results are unclear for report of detail, Simons et al. (2002) report that "over half" of the participants in the removal condition were able to mention distinguishing features of the novel ball. 69% of participants across conditions were able to "'discover' memory" of the ball, after being cued by researchers.

A third and final experiment got rid of the addition condition. The woman held either a soccer ball or a stuffed bunny holding a get-well sign and asked for directions to the hospital. The woman was accompanied by another confederate who held objects either consistent or inconsistent: soccer cleats and an ice pack (consistent with ball, inconsistent with bunny) or a plant (inconsistent with the soccer ball, consistent with the bunny). Once again, a group of confederates would pass by and remove the objects held by the two women. 39% of participants failed to report the change in the consistent condition, and 31% failed in the inconsistent condition. Across both conditions, 35% never reported the change.

While many participants failed to notice the change at all, a considerable number failed to notice only initially and required further, more specific questioning. Can this behavior be explained not by the preservation of a once-conscious representation, but rather by unconscious processes such as *priming* (the ability of sensory, emotional, and/or linguistic input to subconsciously bias behavior)? Were the details of the change processed unconsciously, allowing the participants to make highly accurate guesses despite never having consciously experienced the details?

In a 2002 paper, Mitroff et al. attempted to rule out this explanation. The first experiment sought to determine whether the presence of a change can affect performance despite unawareness of the change. Such an effect would imply that changes can undergo unconscious processing, and prime participants as explained above. Participants completed a change detection task on computerized stimuli such as gratings and numbers, and then were asked a series of additional questions regarding their *confidence* in their response. Mitroff et al. (2002) found a correlation between response time and confidence. Participants responded more quickly when they were more confident and low confidence accompanied incorrect responses. Confidence was found to be a better predictor of response variability than implicit sensitivity. But could implicit detection be responsible for feelings of confidence? Mitroff et al. (2002) say no, arguing that the overall bias found to report no change reflects a conservatism consistent with *explicit* processes. In other words, participants tended to err on the side of caution, and this was interpreted by Mitroff et al. (2002) as evidence that participants were consciously and intentionally considering their responses, as opposed to responding reflexively.

Two additional experiments examined *implicit localization.* Experiment 2 asked whether participants could correctly guess the location of the change despite being unaware of it, assuming an ability to do so is proof of some extent of unconscious processing. Results showed a marked increase in detection over time, and participants were most likely to correctly guess the location of the change on trials where they reported awareness of the change. Because participants performed better on trials where they reported awareness, it could not be concluded that unconscious processing played a primary role in localization. If that were the case, performance would have been similar on both aware and unaware trials. Experiment 3 assessed whether attention was drawn to the location of the change, despite being unaware of it, to determine if awareness or attention better predicts performance. In the first part of Experiment 3, participants performed a change detection task that contained distractor stimuli. In part two, a different group of participants performed the same task, however this time they were forced to report whether or not they saw the change after only one cycle (prechange, blank, postchange) whereas in part one the cycle repeated until a participant reported a change. Performance was not found to be significantly different in parts one and two. Part three of Experiment 3 tested the hypothesis that performance in parts one and two was due to accumulation. Mitroff et al. (2002) argued that if repeated viewing of a change results in the accumulation of localized information, then enhanced performance should only occur when the location of the change is maintained across exposures. If it is not maintained, then there is nothing to learn/accumulate and benefit from. The procedure was identical to part two, however this time the location of the change on each trial was randomized. Results showed performance identical to part two in regards to error rates, response time, and the pattern of slope difference, which Mitroff et al. (2002) took as evidence that localization cannot be explained by implicit accumulation.

While Experiments 1-3 assessed implicit registration and localization of a change, a fourth and final experiment tested *implicit identification.* Participants were presented with an array of rectangles. The critical change (⅔ of trials) involved a single rectangle changing orientation from either vertical to horizontal or horizontal to vertical. An unchanged rectangle was then cued, and participants were asked to report the orientation of the cued rectangle. Mitroff et al. (2002) reasoned that if implicit identification occurred, participants would respond significantly faster when the orientation of the changed and cued rectangles were congruent. In other words, they would be primed by unconscious information and able to make a more accurate guess. Results showed this not to be the case. Participants were not faster to respond to congruent trials, regardless of whether or not they were aware of the change. However, participants did make fewer errors on congruent trials, once again whether or not they noticed the change. A follow-up study confirmed that these results were not due to implicit learning of spatial contingencies. Overall, Mitroff et al. (2002) interpret their results as evidence for "...explicit comparison mechanisms" (as opposed to implicit, automatic, unconscious detection) underlying CB.

Attempts have been made to integrate both the "impoverished representation" (representations are either non-existent or fragile & fleeting) and "rich representation" (representations are rich, but they are not compared effectively) views of CB. Specifically, Varakin & Levin (2010) highlight situations where there is both poor change detection and good long-term visual recognition of analogous stimuli, suggesting that the nature of representations is potentially not best captured by examining how the visual system processes unexpected changes. Two tasks were implemented. The "object task" required participants to respond to a cued object but not to a cued empty location. The "upright task" required participants to respond to cued upright but not inverted objects. These tasks were intended to test pre-attentive and attention-requiring processes, respectively. The first task is more similar to mere detection: something is present or it is not. The second task is more similar to identification and requires more processing: don't just report the presence of something, but tell us a detail about it (whether it is upright). Stimuli consisted of four line drawings and four blank squares (object condition) or four upright line drawings and four inverted line drawings (upright condition). Half of the trials contained changes. Participants were asked if they detected a change (yes or no) and to localize/identify the change via the object's name. Upon completion of the 16 change detection trials, participants' long-term memories were tested. They were presented with two arrays used in no-change trials that occurred earlier in the experiment. Both arrays were identical with one exception: one of the arrays contained a novel object that hadn't been presented in any of the trials. Participants had to correctly identify which array they had previously seen in the experiment. This memory test served to determine whether participants could remember details similar to those that they failed to report in the change detection task.

Change detection analysis involved only the first and last trials. This is because researchers could be sure that on the first trial, participants were not intentionally looking for a change, and on the last trial, they were intentionally looking for a change. For the first trial, significantly more participants in the upright condition responded “yes” when asked if they noticed a change than participants in the object condition. This means that participants detected more changes from inverted to upright (and vice versa) than from object present to absent (and vice versa). Localization performance was also higher in the upright condition than in the object condition. Interestingly, localization performance in the object condition was not above chance, however, for the upright condition, performance was above chance even in trials where participants selected “no” regarding whether they detected a change. This implies that participants were able to recall details about the objects (such as where the changed object was located) despite not detecting a change. For the last trial, where researchers could be sure that participants were intentionally looking for a change, detection was roughly equivalent in both conditions, and localization was more accurate in the object condition. Regarding long-term memory, participants displayed above-chance (chance = 50%) recognition for both the upright (67%) and object (69%) conditions. Varakin & Levin (2010) take this result to suggest that CB can occur under conditions that still lead to functional visual long-term memory. Because of this, they argue that change detection may not be an accurate, functional measurement of the richness/sparseness of mental representations.

A second experiment involved three upright stimuli, three inverted stimuli, and three blank squares in both the object and upright conditions. This change was intended to result in a more complex task. Results showed that long-term memory remained above chance despite an inability to consistently detect unexpected changes. A third and final experiment extended the paradigm to natural scenes. Additionally, instead of a yes/no binary, a 6-point scale was used to report detection. Results showed that subjects’ confidence that something changed did not differ in situations where a change did and did not occur. This was true in both the intentional and the non-intentional trial. In both the intentional and unintentional trials, localization was not above chance. Long-term recognition, however, was once again well above chance (50%) in both conditions (62% object, 75% upright), which Varakin & Levin (2010) took to further support their hypothesis that "visual representations get rich and act poor."

Sensory memory

Another proposed explanation of CB suggests that CB occurs because experimenters do not probe early stages of visual short-term memory that have larger capacities than those stages traditionally probed (working memory, for example). To best understand this, we will take a brief detour to discuss the different stages of visual memory.

The first stage of visual memory is called iconic memory. Iconic memory refers to the automatic processing and brief maintenance of information sensed by our visual system. Everything that is being picked up by your visual field right now is in your iconic memory, at least very briefly (until you make another eye movement). If certain information picked up by the visual system is deemed important, then that information will be transferred to visual short-term memory. Information in visual short-term memory that needs to be actively manipulated to complete some kind of task is then transferred to working memory. For example, the last few words you just read, and the more important themes of this paper are currently stored in your working memory. Finally, with repeated exposure or intentional memorization, items in working memory can enter long-term memory, where things like telephone numbers, anniversary dates, and important life memories are stored.

In 2010, Sligte et al. found that when they tested early stages of visual memory, performance on a change detection task was higher than when later stages were tested. Participants were shown an array of eight objects and were instructed to memorize as many as possible. On each trial, the location of one of the eight items was "cued" (a cue refers to a very brief stimulus that is presented in such a way that it orients attention to a more important stimulus). Critically, the cue latencies varied across three conditions: 10ms after offset of initial display (early retro-cue), 1000ms after offset of initial display (late retro-cue), or 1000ms after offset of initial display but 100ms after onset of test display (post-change cue; the interval between initial and test display was shortened from 2,000 to 900ms in the last condition to allow for this overlap). These three intervals were intended to correspond to *iconic memory*, *fragile visual short term memory*, and *visual working memory,* respectively. After a brief interval where nothing was presented, participants were asked if they had noticed a change. If the trial did contain a change, after responding (regardless of accuracy) participants were presented with the pre-change object and three novel distractor objects that had been in neither array. Participants were then asked to identify the changed item. Sligte et al. (2010) argued that a more detailed representation would enable participants to both detect *and* identify a change, rather than either detect *or* identify changes alone. Identification is traditionally considered to engage higher-level processing than mere detection because it relies on perception and categorization of certain details of an object. For example, I might see a cat dart out of the corner of my eye. If asked if I noticed anything, I could say yes. But if asked *what* I noticed, I might not know. I might say “some type of animal?” The example of the darting cat illustrates why identification (I saw a cat) may signify more thorough processing than detection (I saw something) alone.

Sligte et al. (2010) found that performance was highest in the early retro-cue condition and significantly decreased over the remaining conditions. These results were taken to indicate noteworthy capacity differences among the three memory systems manipulated, with earlier memory stages having higher capacities. To determine the amount of *highly detailed* representations, specifically, authors multiplied change detection accuracy with identification accuracy. Participants displayed the highest number of *detailed* representations in the early retro-cue condition (iconic memory). Additionally, iconic memory and fragile VSTM were found to contain *primarily* highly detailed representations, whereas visual working memory contained one highly detailed and one sparsely detailed representation.

Sligte et al. (2010) interpret this increased performance as suggesting that we may very well have rich, detailed perception like we intuitively feel we experience. In other words, they are saying that our intuitions about the world seeming rich are accurate. They argue that the reason other experiments cited as evidence for “sparse” perception fail to capture this is because they are testing too late of a stage of memory. It can help to imagine memory as a bottle neck, with iconic memory as the body of the bottle containing the vast richness of our immediate experience. However, as time progresses, the capacity of our memory gets more and more narrow, resulting in an inability of later memory stages to do justice to the richness of experience justice.

Sligte et al. (2010) falls victim to a common criticism of many experimental paradigms: can results obtained using such simple stimuli be generalized to the complex real-world settings we find ourselves in during everyday life? Clarke & Mack (2015) sought to answer this question. Participants performed a deletion change detection task (where they were asked to detect a feature that is removed from the scene), however this time the stimuli were natural scenes. The prechange scene was presented (500ms), and, in conditions with a cue, a red arrow cue at 0, 300, or 1000ms was presented after the pre-change image disappeared. The three time intervals were once again intended to correspond with the three relevant stages of memory: iconic, fragile visual short term, and working memory. In another condition, the cue was presented at the same time as the postchange scene. Results were consistent with Sligte et al. (2010), showing that the cue improved performance and that this improvement declined as ISI between pre-change scene and cue increased (in other words, performance was better when early stages of memory were tested). Further analysis revealed that the cue enhanced change detection only at the 0ms and 300ms cue conditions, and change identification at the 0, 300, and 1000ms cue conditions. Comparing the cue conditions directly to each other, there was a significant difference in *identification* between iconic memory (0ms) and working memory (1000ms), but not between iconic memory (0ms) and fragile short-term memory (300ms), or between fragile short-term memory (300ms) or working memory (1000ms), which Clarke & Mack (2015) took to imply a *gradual* fade in memory. Regarding *detection*, there was a significant difference between iconic memory (0ms) and working memory (1000ms), fragile short-term memory (300ms) and working memory (1000ms), but not between fragile short-term memory (300ms) and iconic memory (0ms). The authors interpreted these differences as evidence that the early cues enhance performance by directing participants' attention to a representation of the scene stored in iconic memory. This representation fades over time, which is why there is no effect of the 1000ms and postchange cues on performance.

Sligte et al. (2010) imply that more CB studies would replicate their results if early stages of memory were tested instead of working memory alone. One example they cite is a study by de Gardelle et al. (2009) that used a modified Sperling paradigm. In this study, participants were tasked with reporting as many letters as possible from a cued row. In the first experimental phase, the arrays contained letters. In the second, they contained some catch symbols. The third phase introduced pseudo-letters. In phases 2 and 3, on 1/3 of trials participants were shown 3-8 randomly chosen items and asked to choose which had been part of the previous array, according to their "subjective feeling." Participants had the option of responding that none of the items pictured had been included. While results showed that participants correctly identified catch symbols, participants' sensitivity to pseudo-letters was not significantly greater than 0. Interestingly, the presence of pseudo-letters significantly increased participants' sensitivity to their real-letter counterparts, which de Gardelle et al. (2009) took as suggesting that participants subjectively perceive the pseudo-letters as real letters.

A second experiment increased the saturation of the background and foreground colors and made the auditory cues more distinct. There was also the addition of an "only letters" option in the report, to help determine if participants were indeed subjectively perceiving pseudo-letters as real letters (for example, if pseudo-letters are shown and a participant selects "only letters," it can be reasonably inferred that they are subjectively perceiving pseudo-letters as real letters). Participants took the experiment and were called back for a second session a week later. This time, they were explicitly instructed that there were pseudo-letters they may not have seen last time. Participants were given a cheat sheet of all possible stimuli (letters, catch symbols, pseudo-letters) and were told to study the key for five minutes before beginning the experiment. The only significantly different result from Experiment 1 was that participants showed equal sensitivity to pseudo-letters as they did their real-letter counterparts. de Gardelle et al. (2009) interpret this result as evidence that pseudo-letters are indeed subjectively perceived as real letters, and go on to explain potential mechanisms, such as Bayesian predictive coding (where real letters are a prior resulting in perception of "predicted" real letters even when they are not there). This conclusion supports a sparse interpretation of perception. Similar to perceiving a black-and-white periphery as being full of rich color, incorrectly perceiving pseudo-letters as real letters is interpreted by some researches as evidence of a shocking deficit in the perceptual system, and how can a system with such deficits be considered to perceive things "richly"? Iconic memory researchers would push back against this conclusion, arguing that de Gardelle et al. (2009) may have had different results if they had probed earlier memory stages. Specifically, they argue that the presentation of the mask after the memory array may have disrupted the representation of items in iconic memory, causing them to not make it through the bottleneck into working memory (Vandenbroucke et al., 2012).

Perceptual reentry–the (re)appearance of a stimulus in perception after its physical disappearance–has been cited as a threat to sensory memory's (a term used from here on out to refer to combined iconic and fragile visual short-term memory) explanation of perceptual richness (Ward et al., 2016). In 2004, Mitroff & Scholl demonstrated perceptual-reentry using a *motion-induced blindness* (MIB) paradigm (explained in the image below).

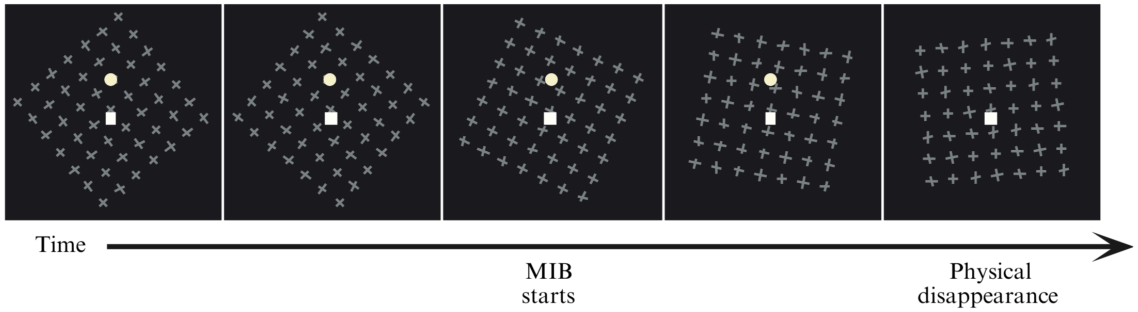

Figure 10: Motion-induced blindness
An example of the procedure reproduced from Mitroff & Scholl (2004). Participants were asked to fixate on the center square while the pattern in the background would rotate around. Eventually, participants would become "blind" to the white circle above the square, even though the circle was still present. For a better example of this illusion, see the following video: https://www.youtube.com/watch?v=Hfrb94mKCJw focus on the flashing green dot in the center, and eventually the three yellow dots on the edges will disappear.

Participants were presented with the stimuli pictured above and were instructed to notify experimenters when MIB had occurred (when they could no longer perceive the white circle). Then, the circle would be removed from the screen. Shockingly, participants almost always saw the circle disappear–95% of the time, even when they were blind to the circle beforehand. This is compared to 98% when they were aware of the circle the whole time. Mitroff & Scholl (2004) argue that such a finding is evidence that the brain maintains a representation of the circle even when participants are blind to it. A second experiment sought to determine if this representation was formed initially and then maintained over time, or, alternatively, was constantly re-formed and refreshed during MIB. MIB was once again induced in participants, but this time, after participants reported a loss of awareness, the stimulus (a line in experiment 2 instead of a circle) would rotate to the left or right before disappearing. Participants were able to accurately report the orientation of the stimulus when it disappeared, which Mitroff & Scholl (2004) interpreted in favor of the refresh hypothesis–if participants simply maintained an initial representation, how could they accurately report the change in orientation that occurred when they could not subjectively see the line? Ward et al. (2016) argues that MIB–as displayed in this study–challenges sensory memory because it could offer an alternative explanation of why participants performed so well when early memory was probed in Sligte et al. (2010). In Mitroff & Scholl (2004), participants reported subjective unawareness of a stimulus. When the stimulus disappeared–i.e., a change was made to it–subjective awareness suddenly occurred. Could the same process be involved in Sligte et al. (2010)? Is it possible that participants were unaware of the stimuli prior to the moment of change, similar to Mitroff & Scholl (2004)?

It has also been suggested (Ward et al., 2016) that *late visual reactivation* could also explain the results of Sligte et al. (2010). Like MIB, this phenomenon involves the perception of a no-longer present stimulus, induced by post-cueing attention. An example of this is illustrated in Sergent et al. (2013). Participants were presented with two circles to the right and left of a centrally located fixation cross. Their task was to report the orientation of a single Gabor patch that would randomly flash in one of the circles. Attention was cued via dimming of one or both circles and could appear either before or after the target stimulus. Sergent et al. (2013) found an increase in accuracy when the cue and the patch appeared in the same circle, even when the cue was presented 400ms after the target first appeared. A second experiment was conducted to ensure that participants were subjectively unaware of the target prior to the post-cue, as the results are especially notable in this case. To determine this, subjects reported their subjective awareness of the target on a scale. Results identified two key types of trials: 1) near-chance performance and a subjective visibility rating of 0% and 2) high performance and high visibility. This categorization may initially seem to support the results of Sligte et al. (2010), because participants performed better when they reported that they were aware of the target prior to the cue. This would imply that the cue did not retroactively bring the target into awareness, but that it aided an already existing awareness. However, Sergent et al. (2013) argue that attention, not awareness, is still predominantly responsible for performance and for the *balance* between these two types of trials. Their evidence for this comes from a model of the subjectivity ratings, which significantly fit the data. Notably, the estimate of “aware” trials provided by this model was

"...strongly influenced by both pre- and post-cued attention and could account for the effect observed on detection sensitivity" (Sergent et al. 2013), suggesting that attention was, after all, a better predictor of performance than awareness.

Can sensory memory benefits be explained by unconscious processes such as attention? Vandenbroucke et al. (2012) sought to address this issue. Their study used the *Kanizsa illusion* (pictured on the following page) to assess the characteristics of early (iconic) and late (fragile) sensory memory. If the Kanizsa illusion can be processed by sensory memory, then sensory memory must involve high-level processing such as perceptual organization and perceptual inference (Vandenbroucke et al., 2012 cite Harris et. al., 2011 as evidence that induction of the illusion relies on these processes). Vandenbroucke et al. (2012) further argue that because these processes cannot be supported by unconscious processing alone, it would stand to reason that items subject to these processes must not be entirely unconscious.

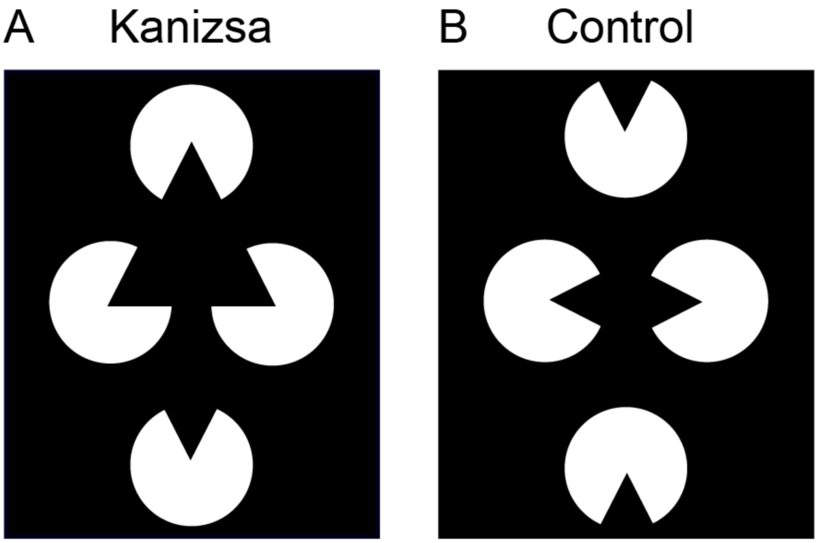

Figure 11: The Kanizsa illusion
An example of the stimuli used in Vandenbroucke et al. (2012). On the left is the Kanizsa illusion. Note that there is no actual triangle drawn in the image. Instead, viewers are able to infer a triangle. Because there are no actual contour lines hitting the retina, it is argued that the illusion relies on higher-level processing involved in inferring image details.

Like Sligte et al. (2010), Vandenbroucke et al. (2012) used a change detection paradigm with three conditions of retro-cueing to manipulate early sensory memory (iconic memory), late sensory memory (fragile visual short-term memory), and working memory. In the first experiment, the key change was made to the figure via the inducers (the white circles that sort of look like Pac-Man that help cause the illusion to occur). In the Kanizsa condition, the illusory triangle formed by the inducers acted as an aid to participants by creating a more noticeable change. The early and late sensory memory conditions were identical except for the duration of the blank occurring between the memory display and the cue (33ms early sensory, 1000ms late sensory). Unlike these two conditions, the working memory condition contained a 900ms post-memory blank and the test display was presented before the cue. Performance was higher in the early sensory memory condition than the late sensory memory condition, and performance in the late sensory memory condition was higher than in the working memory condition. Similar to Sligte et al. (2010) and Clarke & Mack (2015), the authors interpreted this as evidence that sensory memory has a larger capacity than working memory. Performance on trials containing the illusion was higher than control trials, meaning that the illusion was successful in creating a more noticeable change that boosted performance. Notably, this performance benefit was found even in sensory memory conditions, meaning that participants' sensory memory must have processed the illusion in order to receive the benefits. Vandenbroucke et al. (2012) take this as suggesting that sensory memory involves higher order processing, because the illusion was effective in sensory memory. Following the authors' argument, this would mean that items in sensory memory are consciously perceived.

In a second experiment, the inducers were made isoluminant to the background to decrease the strength of the illusion. Authors argue that this allows for the illusion to be *cognitively inferred*, rather than directly perceptually inferred. This distinction was compared to "...whether you decide there should be a triangle compared to actually perceiving a triangle…" (Vandenbroucke et al., 2012). Replicating the first experiment, performance in the perceptual-inference condition was found to be better in the sensory memory conditions than in the working memory condition, and performance was also enhanced in Kanizsa trials compared to control trials. Performance was also enhanced in Kanizsa compared to control trials in the isoluminant/cognitive-inference condition, however this effect was the same across all three memory conditions. Vandenbroucke et al. (2012) interpreted the disappearance of the performance difference between sensory memory and working memory in the cognitive-inference condition as evidence that the benefits of the illusion must be due to perceptual inference. Further, they argue that this supports their hypothesis that items in sensory memory do undergo high-level processing, and that these representations are phenomenally experienced (as evidenced by their being perceptually (directly) rather than cognitively (indirectly) inferred).

To assess brain regions associated with representation in sensory memory, Vandenbroucke et al. (2014a) used fMRI to measure brain activity in response to presentation of the Kanizsa illusion in different conditions of inattentional blindness. Also investigated in this study was whether perceptual inference is purely perceptual vs. relying (at least to some extent) on cognitive access.

Vandenbroucke et al. (2014a) used four main stimuli: a Kanizsa illusion, a Kanizsa control image, a line image, and a line control image (pictured on the following page). These stimuli were embedded in a two-back task where participants were instructed to report a repetition of a letter that was presented two trials (two "letters") earlier. The cognitive demands of this task were intended to induce inattentional blindness, providing researchers with the ability to compare brain activity correlated with perception of the illusion in sensory memory (inattentional blindness) vs. access of the illusion (no inattentional blindness). Because of the temporal imprecision of fMRI, Vandenbroucke et al. (2014a) used inattentional blindness vs. no inattentional blindness to compare sensory vs. working memory respectively. To avoid confusion, it should be noted that the authors relied on an uncommon conceptualization of inattentional blindness. Usually, inattentional blindness is interpreted as subjective unawareness of a stimulus (hence the "blindness" part of the term "inattentional blindness"). However, Vandenbroucke et al. (2014a) interpret inattentional blindness as involving perceptual awareness, along with a failure of cognitive access (despite being aware, participants do not "know" or report their awareness, processes the researchers believe involve cognitive access).

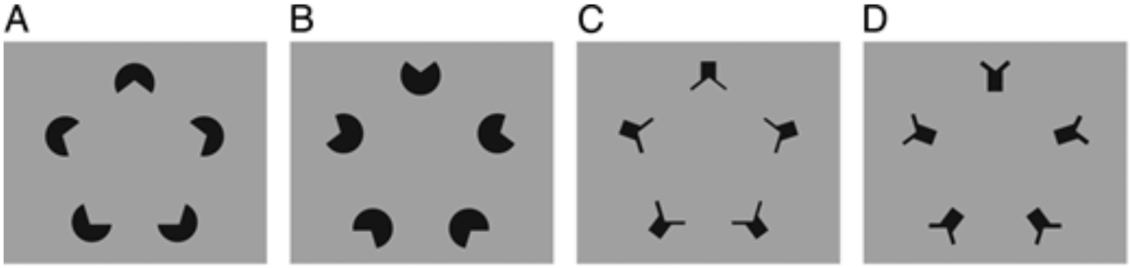

Figure 12: Kanizsa illusion II
Segment of an image from Vandenbroucke et al. (2014a) depicting their four main stimuli. A) the Kanizsa illusion, where a pentagon is perceptually inferred despite there being no real lines drawn B) Kanizsa control C) line image of a cognitively inferable pentagon D) line control image

fMRI analyses focused on the following regions: V1, V2, V3 V3AB, V4, and LOC. Heightened activity was found in both low- and high-level visual areas for the Kanizsa figure compared to its control. This difference was not found between the line figure and its control, which the authors suggest implies that the cognitive inference involved in the line figure is not accompanied by any unique visual/neural signal, providing additional evidence that perception of the illusion relies solely on perceptual inference and cognitive inference. Univariate analysis revealed no difference in patterns of neural activity in conditions of inattentional blindness vs. no inattentional blindness, which Vandenbroucke et al. (2014a) claim as evidence that the illusion is processed even in the absence of access/report. If participants are aware of the stimulus when they report it, then the brain activity associated with reporting should also be associated with awareness. Additionally, multivariate analysis showed that the neural signature of the non-reported illusion could be used to classify the signature of the reported illusion. In other words, the patterns of brain activity associated with the illusion when participants were inattentionally blind was so similar to the activity when they were not that the patterns could be used to identify and predict each other.

Finally, research on sensory memory has also investigated how confident participants feel in their performance across the three stages of memory. Vandenbroucke et al. (2014b) presented participants with white rectangles of four possible orientations on a black background. Then, a cue (always valid) was presented at three conditions of visual memory (iconic/early-sensory, 50ms ISI; fragile/late-sensory, 1000ms ISI; working, 900ms ISI) with the cue occurring at the end of the listed ISI in the first two conditions, and after a brief replication of the memory array (100ms) for 500ms in the last condition. Participants performed a change detection task on the array. After reporting whether they detected a change (*objective* measure), participants indicated their *subjective* confidence in their answer: sure, doubt, or guess. Results replicated much of the iconic memory research discussed, showing a higher detection sensitivity (d′) for iconic memory than fragile and working memory. Once again, iconic memory capacity was highest, followed by fragile and then working memory capacity. Participants were biased more liberally (false alarms) during the sensory memory conditions and more conservatively (misses) during the working memory condition. Because of this bias, Vandenbroucke et al. (2014b) ran an additional experiment as they worried that the bias could have confounded participants' perceptions of their choices (confidence). For example, a liberal bias may reasonably be associated with higher confidence, while a more conservative bias may be associated with lower confidence. To assess for bias, the change detection task was replaced with a change identification task. Instead of rectangles, the array contained arrows and cued arrows would change orientation. Participants were asked to report the specific change in orientation (clockwise or counterclockwise). Because this task relied on identification and not a more automatic process like detection, Vandenbroucke et al. (2014b) predicted that it would result in a response bias closer to 0. Sensitivity did not differ significantly from Experiment 1 and response bias was insignificant. *Metacognition*–defined by the researchers as the relationship between confidence and accuracy, such that when correct responses are accompanied by high confidence and incorrect by low, a participant can be said to have *high metacognition*–was calculated for each of the three memory conditions. Metacognition was lower for iconic memory than for working memory, but not significantly different between fragile and working memory. Vandenbroucke et al. (2014b) interpreted these results as suggesting that fragile memory may involve explicit, conscious processing of information, similar to that which occurs in working memory. On the other hand, because of iconic memory's lower metacognition, the authors suggest that "Possibly, the mechanisms underlying iconic memory are partly implicit," or, partly unconscious rather than fully conscious (Vandenbroucke et al., 2014b).

Miscellaneous

Some relevant articles did not directly address the major topics discussed above: CB, sensory memory, color perception, or summary statistics. These papers will be discussed below.

Research on perceptual richness predominantly involves brief presentations as it aims to study momentary, immediate experience. However, this is not what things are like in the "real world." In everyday life, we experience long-term immersion in stable environments that we can explore with eye and head movements. Studies involving subjective detailedness in such stable situations have been conducted. In 2006, David Melcher presented 20 participants with natural scenes for 5, 10, or 20 seconds, a set of durations chosen to best replicate scene-scanning as it occurs outside of the laboratory. Some images were immediately followed by a memory probe, where participants were asked to recall objects from the scenes. Other images were not immediately followed by the probe, but rather the probe occurred after 4 to 6 intermediary images. In a second experiment, scenes were presented for 1 or 10 seconds, with some conditions containing a 10 or 60 second delay between a 10 second and 1 second presentation. The delay period contained either a reading task intended to strain working memory, or a visual short-term memory task where participants performed a change detection task. Results showed that performance on the memory probes increased with duration of the delay, repetition, and for central stimuli. This effect was not impacted by either the working- or visual short-term memory tasks. Melcher (2006) took these results to suggest that our experience of rich detailedness in the environment comes from the accumulation of detail in memory over time.

In an additional study, Melcher (2010) once again examined the accumulation of details for natural scenes, however this time looking at the effects of the emotional charge of the scenes and individual differences in the participants. In Experiment 1, participants performed a color change detection task, a digit span task, and a memory task like that used in Melcher (2006) where participants were required to identify objects from a natural scene. Similar to the previous study, performance at the scene memory task improved as duration increased. No significant individual differences among participants were found in the color change detection task, however a significant correlation was found between performance on the change detection and scene memory tasks. Another significant correlation was found between performance on the digit span and change detection tasks, revealing that there may be individual abilities that contribute to and impact performance on the kinds of tasks involved in this research.

In Experiment 2, participants were shown neutral, highly negative, or highly positive emotional images. ⅓ of the images did not have an immediate memory probe, and participants were called back to the lab to take the probe 1 week later. Results showed that memory was impaired for negative images, such that participants' performance did not improve over time as was normally found. For positive scenes, however, participants performed better than they did on trials involving either neutral or negative scenes. Performance again improved at the 1-week mark, however at this point no effect of emotion was found.

The results of these two studies contribute to our understanding of perceptual richness by providing data related to factors such as time, memory, emotional content, and individual difference. However, their relevance to the rich vs. sparse debate as discussed in this thesis is lacking simply because they do not manipulate the type of momentary, immediate perception the debate is centered around. While the 2006 paper claims results as evidence against perceptual sparseness, such a claim is not particularly relevant because the study simply does not manipulate the same kind of perception the many other studies involved in this debate do. In that sense, it is almost as if these two studies address a different topic. Yet on the other hand, these two papers do highlight a gap in the current debate–namely, its applicability to real-world situations.

In a separate study, Heinrich & Bach (2010) found that participants perceived higher subjective detailedness in a smaller version of two otherwise identical images. Participants were presented with two versions of the same image (either a natural scene or a random pattern) and were asked to report via keypress which image they thought was more detailed. Critically, one image was scaled down to ⅓ the original size. Results showed that in order for both images to receive the same subjective detailedness ranking, the *objective detailedness* of the smaller image had to be reduced, implying that participants perceived smaller images as more detailed than larger images. Objective detailedness was defined in this study as the number of spatial frequency cycles in the image. A spatial frequency cycle is a change in contrast, as pictured on the following page:

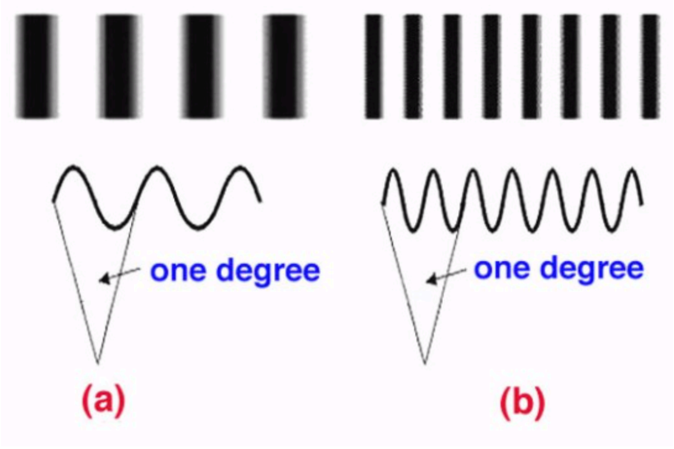

Figure 13: Spatial frequency cycles
This image (reproduced from Kalloniatis & Luu, 2007) shows the amount of spatial frequency cycles per degree. The array on the right (b) has a higher number of spatial frequency cycles per degree. This is comparable to how Heinrich & Bach (2010) calculated the number of spatial frequency cycles per image.

Heinrich & Bach (2010) propose their *visuo-cognitive hypothesis* to explain their results. This hypothesis posits a visuo-cognitive *schema* for assessing image quality, including not only a schematic representation of the physical characteristics of the image, but also an intermediate level representation of the properties of the visual system itself. Because this schema is aware of the gross limitations of the visual system (such as blind spots, interruptions due to blinks and eye-movements, etc.), it is accustomed to predicting more detail than is immediately available. However, Heinrich & Bach (2010) mention that for lower-frequency images, the schema will assume that a lack of detail is due to the physical properties of the image and not the deficits of the visual system itself. The opposite is true of higher-frequency images, with such stimuli causing the schema to overpredict detail. Henrich & Bach (2010) argue that such a schema may explain the discrepancies found in subjective detailedness reported for images of different sizes and frequencies. Small images are perceived as more detailed because the schema interprets the lack of information due to size as not due to the image, but due to a failure of the visual system. Thus, it compensates by predicting more detail, giving the viewer an inaccurate subjective sense of the image.

Park et al. (2010) argue that *refreshing* may allow us to subjectively feel as if we perceive more detail in a scene than is immediately visually available. *Refreshing* refers to a cyclical, post-perceptual process where part(s) of a scene is/are represented despite not being physically seen anymore. Participants were presented with panoramic scenes for 1.5 seconds while undergoing fMRI imaging. On trials in the refresh condition, participants were instructed to reimagine (refresh) a cued part of the scene (in this case, either the left or right side of the scene). On trials in the non-refresh condition, instead of cues, half of the immediately preceding scene was displayed, so that participants directly viewed half of the scene instead of mentally refreshing it. Congruent cues would cue the same side that would later be refreshed or perceived, and non-congruent cues would cue the alternative side. Participants did not make any overt responses on a trial-by-trial basis, but rather answered surprise memory questions at the end of the experiment. fMRI regions of interest included the parahippocampal place area (PPA)--a region associated with scene perception– and retrosplenial cortex (RSC)--the area of cortex where the hippocampus is located.

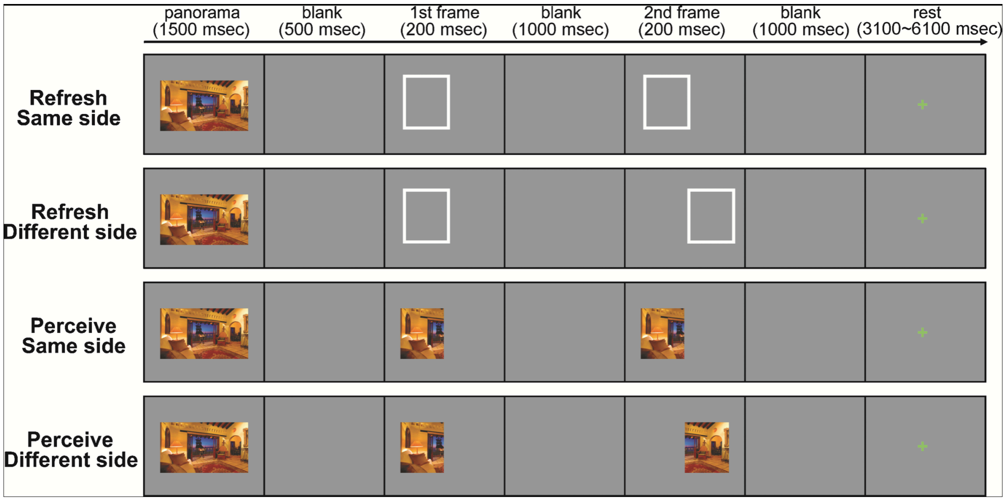

Figure 14: Park et al. (2010) procedure
Visual diagram of the procedure used in Park et al. (2010) (reproduced). Note that the "different side" conditions refer to an incongruency in the cue and the side probed. The opposite is true for the "same side" condition.

Activation in the PPA was attenuated for same-side perception but not different-side. The same was found to be true for refreshing. Park et al. (2010) claim this indicates that the same neuronal populations are activated when viewing a scene for the second time that are activated upon initial viewing, and that refreshing a scene in large part replicated this process. In other words, activity in the PPA when refreshing a scene is similar to that associated with actually viewing the scene. In accordance with this, there was no significant difference found between activity in the PPA between refreshing and perceiving. Similar to the PPA, activity in the RSC was attenuated during the same-side perceive condition, meaning once again that activity was similar across multiple viewings of the same image. Unlike in the PPA, in the RSC there was a difference found between perceiving and refreshing: a significant decrease in activity when the same-side was refreshed compared to when it was merely perceived. The authors interpret these results as indicating that the PPA is involved in refreshing information that is more specific and localized within a larger scene, whereas the RSC can refresh information that is not restricted to such a limited region within the scene. In this sense, "...the PPA and RSC may represent different levels of scene information" (Park et al., 2010). Additionally, whole-brain analysis showed that when participants refreshed scenes (both same and different sides), heightened activity occurred in frontal areas (superior frontal gyrus, middle frontal gyrus, inferior frontal gyrus).

Wang et al. (2021) conducted a study that sought to probe into the nature of reportability. Namely, can participants fail to report fully attended items? Participants were shown one of eight possible numbers which was presented in either red, blue, yellow, or purple against a grey background. Participants' task was to report whether the number was odd or even as quickly as possible. After the 28th trial, participants were asked a "surprise" question regarding the identity or color of the number. Questions were forced choice, and all four options were of the same parity (all were odd or all were even). Participants were quite successful at the parity task (reporting oddness or evenness of the numbers), however performance on the surprise questions was not as high (35% correct for color; 55% correct for identity). Wang et al. (2021) describe their results as evidence that *attribute amnesia*–failure to report a task-irrelevant attribute of a stimulus due to the cognitive demand of the task–occurs even in a task with a very low cognitive load. It appeared as if because participants were tasked with focusing on one feature of the numbers (their parity), they failed to represent and recall features of the numbers irrelevant to the task (their color or even the name/identity of the number), even though the task was not very difficult and all features involved very obvious and seemingly hard to miss. Experiment 2 replicated the procedure of Experiment 1 but used Chinese characters (more complex, meaning-imbued stimuli) instead of numbers. Participants were asked to respond to whether the character corresponded with a type of furniture. The surprise question probed the direct identity of the character. Results were similar to Experiment 1, with only 50% of participants answering the surprise question correctly. This seems to imply that attribute amnesia is not reduced even when stimuli are more generally meaningful.

Experiment 3 tested if the results would replicate in a more real-world context. Participants were given 32 cards and asked to sort the cards by parity as quickly as possible. Upon completion of the task, participants were asked to report the color and identity of the last card (forced choice). The results showed that attribute amnesia can still occur in more realistic settings, with correct responses a mere 35% for identity and 40% for color.

Wang et al. (2021) claim that because their task had such a low cognitive load, it is likely that all attributes of the stimuli were not only attended, but also consciously perceived (using Kouider & Dehaene (2007) criteria for access consciousness: sufficient strength of input stimulus and sufficient top-down attention; i.e., in other words, their argument is that it is unreasonable to assume that participants did not consciously perceive such noticeable information). They claim, because of this argument, that their findings are evidence that perceptual consciousness overflows report—participants failed to accurately report information they obviously perceived, so therefore the capacity of perception must be greater than the capacity of access. Further, they argue that a unique strength of their study is that it involved stimuli that were fully attended, unlike other studies on perceptual overflow which predominantly investigate stimuli outside of attention.

Finally, Solovey et al. (2015) investigated a phenomenon known as *subjective inflation*–an overestimation of the reliability of perceived detail. Specifically, they hypothesized that, under a lack of attention, participants may employ a more "liberal detection criterion" for peripheral stimuli compared to central stimuli. Stimuli were sometimes images containing variably oriented gratings embedded in noise, and other times just images of noise. Participants were tasked with reporting the presence or absence of gratings at cued (arrow) locations (peripheral or central). Participants reported subjectively perceiving the targets more often at peripheral locations, even when the targets were not actually present. However, the difference in overall detection sensitivity between the two locations was not significantly different. Despite this, Solovey et al. (2015) interpret these results as confirming their hypothesis that more liberal detection criteria are applied to peripheral content.

A second experiment sought to assess the impact of manipulating the size of peripheral content. Will the subjective impression of peripheral and central stimuli be more alike if peripheral stimuli are enlarged to compensate for the fact that they are not centrally located? Solovey et al. (2015) argued that such an effect would imply that subjective inflation may in part be due to the number of neurons recruited for peripheral vs. central processing. Because enlarging peripheral stimuli would recruit more neurons to process the stimuli, it may mimic central processing on a neuronal level. Results confirmed researchers' hypothesis, showing that the detection bias disappeared.

A third and final experiment examined the possibility that inflation is a cognitive strategy by providing trial-by-trial feedback. If feedback enhances performance, Wang et al. (2015) claim it could be reasoned that subjective inflation is a strategy that can be learned from experience. Results showed that the addition of feedback did not significantly affect either sensitivity or bias compared to no-feedback. Solovey et al. (2015) take this result in support of their hypothesis that subjective inflation is not a cognitive strategy, but rather perceptual in nature.

Discussion

Sensory memory: benefits & challenges

Studies that claim to find evidence for a large capacity sensory memory system offer a captivating explanation of perceptual richness because they put a satisfying end to the rich vs. sparse debate. In other words, high-capacity sensory memory does not leave one asking: well, then, why does it *seem* as if the sun revolves around the earth? If the results are to be believed, high-capacity sensory memory ends the discussion with perception simply being rich, full stop. Things are as they seem. We indeed perceive (via sensory memory) more than we can report (via working memory), and thus our intuition is accurate. Perception overflows access.

However, research on high-capacity sensory memory is still far from reaching such a firm conclusion. There is still disagreement in the field as to whether iconic memory (a component of sensory memory) involves primarily conscious or unconscious processing. If items in iconic memory are processed unconsciously, then the compilation of results supporting a high-capacity sensory memory system cannot be considered proof of rich conscious perception. One curious result included in this review is from Vandenbroucke et al. (2014b). To reiterate, this study found that metacognition–a process associated with explicit, conscious processing–was lower for iconic memory than fragile memory, and not significantly different between fragile and working memory. The authors then go on to discuss sensory memory (iconic + fragile memory) as involving the same explicit, conscious processing as does working memory, evidenced by similar metacognition. However, they never clearly justify the jump from iconic memory having lower metacognition to (iconic memory + fragile memory) having metacognition roughly equal to that of working memory. Because of iconic memory's lower metacognition, the authors themselves even hesitantly admit: "Possibly, the mechanisms underlying iconic memory are partly implicit" (Vandenbroucke et al., 2014b).

Access & report

The question of whether iconic memory is primarily an implicit, unconscious process highlights an important conceptual issue relating back to the original phenomenal/access distinction. In the study involving fMRI imaging and perception of the Kanizsa illusion, Vandenbroucke et al. (2014a) used an unusual interpretation of inattentional blindness (IAB). Traditionally, IAB is associated with a lack of awareness—if a participant is inattentionally blind to something, then they are unaware of it. However, Vandenbroucke et al. (2014a) take the stance that participants are not unaware of stimuli subject to IAB, they merely haven't *accessed* it. This is conceptually consistent with iconic memory's explanation for CB–participants are not actually *blind* to the change, rather the window of access has been missed by testing late-stage, lower-capacity memory (working memory). Vandenbroucke et al. (2014a) cite their finding of no significant difference in brain activity between instances of IAB and no IAB as evidence for awareness of stimuli during IAB—if activity did not differ, and brain activity is responsible for awareness, then how could level of awareness differ? However, the study failed to include imaging of frontal areas (especially dorsolateral and inferior prefrontal cortex), which have been strongly associated with conscious access, and in some cases with conscious perception, even in "no-report" situations (Del Cul et al., 2009; Whyte et al., 2022; Kapoor et al., 2019). The failure to include such a highly relevant and implicated brain region severely limits Vandenbroucke et al.'s (2014a) claim, as a difference may have been found if this crucial brain region were imaged. However, even if they did include analyses of frontal areas and their data did indeed show differences in prefrontal cortex for stimuli in IAB vs. non-IAB conditions, the authors would have likely still interpreted such differences as being associated with access and report, which they argue is possible in non-IAB but not IAB conditions.

It is important to note that sensory memory paradigms rely on report, and thus involve access. Since Ned Block outlined the access/phenomenal distinction in 1995, consciousness research has gradually moved toward mitigating the confounding impacts of report (Tsuchiya et al., 2015). Because report is a process that inherently relies on access (recall that access is defined as the availability of information via cognitive systems for use in processes such as report), researchers have sought ways to isolate pure, phenomenal experience, untainted by the neural activity associated with access. This has led to a rise in popularity of no-report paradigms that do not probe awareness on a trial-by-trial basis, so that the effects of report do not taint time-sensitive neural activity and behavioral effects (Pitts et al., 2014; Cohen et al., 2020b). If items in sensory memory are indeed consciously experienced, and the results of studies such as that led by Sligte et al. (2010) are to be taken as evidence for perceptual richness, then the attempts of no-report paradigms have been in vain, because access itself is not the problem, only when access occurs. Block's distinction (1995) was inspired by concerning results of studies such as those using the Sperling (1960) paradigm: why is performance inconsistent with what people claim to experience? Is science missing something? High-capacity sensory memory, however, explains the results of the original studies that were so concerning to Block: researchers simply failed to probe the correct stage of memory. The sensory memory paradigm prompts us to consider if access is the real problem, or simply *when* information is accessed. Block (2011) argues for the latter, saying that the issue in the original Sperling (1960) paradigm arose not because the phenomenal contents were entirely *inaccessible*, but merely because they were *unaccessed.* The sensory memory paradigm resolves Block's concerns by enabling participants to access what was formerly unaccessed.

Despite this, access is often discussed in the field as a problem to work around. That is, access is an issue not because things are either inaccessible or unaccessed, but because access itself is a confounding factor. No matter when things are accessed, access may still pose an issue. This is the interpretation of access that has inspired the development of no-report paradigms. Additionally, access is understood to involve different brain processes than those involved in pure experience itself (Cohen et al., 2020b). In this sense, the issue of access ultimately still remains. Are there still contents of experience that are unaccessed in the sensory memory paradigm? Do we experience things that we cannot access in any circumstances, even through introspection (ironically, a form of access, through which we have discovered the issue of access)? Are no-report paradigms created in vain if high-capacity sensory memory resolves Block's concerns? Is there merit to isolating brain activity associated with experience vs. report & access?

Relevance of results to popular theories of consciousness

The studies included in this review were considered in the context of popular scientific theories of consciousness. Out of this consideration, Attention Schema Theory (AST), predictive processing, and the sensorimotor theory of perception emerged as being especially relevant. These theories can offer a compelling explanation of results by accounting for why perception *seems* rich, even in instances when objective measures suggest a much sparser perception.

Attention Schema Theory

The relevance of AST to the results discussed in this review regards how a discrepancy between subjective report and behavioral performance may occur. Research predominantly asks why people don't perform in accordance with their reports (why does behavior exhibit sparseness) instead of: why do people report something so different than what their behavior exhibits? The former places more trust in the behavior, while the latter places more trust in the subjective experience. Pursuing the former will lead to more attempts to create paradigms that capture behaviorally what participants claim to experience. Pursuing the latter accepts both the existing behavioral evidence for sparse perception and the subjective evidence for perception feeling rich, focusing efforts on understanding what might cause such a difference.

Without taking a clear side on the issue, AST explains how a discrepancy like this—between behavior and report—could happen. Recall that the attention schema is what causes the "sense," so to speak, that we are conscious. Just like the body schema provides information about the arm, the attention schema provides information about attention. However, the arm schema is not the arm (remember the example of the phantom limb?). Neither is the attention schema attention. The schema is a *loose, rough* model. This allows for efficiency, but it also means that some detail gets left out. The information provided by the attention schema is, according to Graziano (2020), precisely what allows us to say that we are conscious, to talk about consciousness, and to experience consciousness via our ability to "know" that we are conscious. Similarly, the information provided by the body schema allows us to move the arm, hand, and fingers intentionally and with finesse. However, the looseness of the attention schema can result in some missed detail—this is what allows for a discrepancy between what is shown by behavior and what is experienced by the behaver. I can describe throwing a ball, but I do not do so in terms of exact weight, force, etc. Similarly, I can tell you about how I perceive a scene, but I may not be able to do so in terms of every precise detail.

Additionally, the idea of the attention schema is similar to Heinrich & Bach's (2010) visuo-cognitive hypothesis (see Results: Miscellaneous).

Participants rated smaller images as more detailed than their larger counterparts, and Heinrich & Bach (2010) hypothesized that this was due to a schema of the visual system that predicted the amount of detail in an image based on both information from the image and information from the schema about the visual system itself. For the purposes of current research on perceptual richness, which mainly examines visual perception, the attention schema and the visual schema described by Heinrich & Bach (2010) have considerable overlap. Could it be that our intuitive sense of a rich world is based on predictions made by an efficient & functional yet imperfect schema?

Predictive processing & controlled hallucinations

Models as reality

The major themes identified in the results also showed particular relevance to predictive processing, along with Anil Seth's idea of conscious perception as a "controlled hallucination" (2021). It is important to note that Seth points out a key difference between predictive processing and controlled hallucinations, specifically that predictive processing is a theory describing the mechanisms of the brain, whereas the idea that reality is a controlled hallucination builds off of predictive processing and seeks to explain phenomenal experience in ways predictive processing alone fails to (Seth, 2021). The term "controlled hallucination" refers to how, according to Seth's theory (2021), we do not perceive physical reality directly. Rather, what we perceive is a prediction about what the system detects in physical reality. In that sense, what we see can be considered a sort of hallucination. However, the hallucination is "controlled" by the error minimization process, which will be discussed in more detail below. Going forward, "predictive processing" as a term will refer to an idea of consciousness that includes not only the original tenets of predictive processing, but also Seth's additional insights into phenomenal experience as a controlled hallucination.

The idea that our perception is not a direct perception of physical reality overlaps conceptually with AST. AST posits that we are able to believe we are conscious thanks to a schematic model of attention. AST also posits the existence of visual models, similar in kind to the attention schema (Graziano, 2020). Both theories hypothesize that reality is perceived indirectly through models.

The predictive process

Predictive processing hypothesizes that the brain acts like a prediction machine, constantly making top-down predictions about environmental content and refining these predictions based on bottom-up sensory input. This process is referred to as *error minimization* and can be understood as a form of Bayesian probability. Bayesian probability is a form of probability that directly accounts for knowledge of factors relevant to the predicted event, in addition to how *reliable* this knowledge is. This is contrasted with other conceptions and calculations of probability that emphasize the mere frequency of events or the inherent nature of certain things/systems to behave in particular ways. The practice of updating relevant information weighted for reliability is how, according to predictive processing, the brain makes predictions about the environment.

For example, imagine you are trying to make a prediction about what to wear—will the weather today be warmer or colder? You look out your window, and you see the sun shining. You know that sunshine in your geographic region is associated with warmer weather, so you predict that the weather might be warmer. In Bayesian statistics, this kind of information based on past experience is known as a *prior*. In this scenario, you also have a birdfeeder outside your bedroom window. Over time, you've noticed a trend regarding the kinds of birds that visit your birdfeeder: on warm mornings, more red birds come, whereas on colder mornings, more brown birds come. Soon after waking up, you notice some brown birds fly over and perch on your birdfeeder. Because you have past reason to reliably associate brown birds with colder weather, the introduction of this additional prior to the calculation will decrease the probability that the morning is warm. Finally, you walk downstairs and step outside. You feel that, although the sun is out, the day is quite brisk. This information is again factored into the calculation. Because of its immediacy in answering the question, it is considered the most reliable of all the factors, and the prediction updates to "cold outside."

However, to really familiarize ourselves with the nuance of Bayesian probability, especially in regards to *reliability,* let us consider one last variable. Suppose when you stepped outside, your skin and hair were wet—you had just taken a shower, and you did not bother to dry off before stepping out. Because you know that being wet can cause you to experience air temperature differently (prior), this would reduce the *reliability* of the information you gained from stepping outside and feeling the air temperature on your skin. So, although the prediction overall may still be that the morning is cold, the probability may not be quite as high as if you were to step outside with dry skin and hair. This process sounds incredibly familiar to the visuo-cognitive hypothesis proposed by Heinrich & Bach (2010) that was discussed above in relation to AST. The visual schema weighted the *reliability* of the larger and smaller image differently based on knowledge about the visual system and past experience with similar sensory information. This caused the predicted (experienced) amount of detailedness to differ significantly between the two images.

Seth's idea of consciousness as a controlled hallucination emphasizes that what we experience *is* the prediction (2021). We do not experience predictions *about* what we see. Rather, what we see *is* the prediction itself! When less detail is predicted, we experience less detail. When more detail is predicted, we experience more detail.

Inattentional blindness: a case of bad predictions

Seth (2021) uses the example of an animal approaching. Despite the moving object appearing dark and furry, it is probably unlikely to be a gorilla, because, let's say, in the hypothetical region that you find yourself in, you have not had many interactions with gorillas. This means there is no good reason for "gorilla" as a prior to be factored into the prediction, so the end prediction is unlikely to be "gorilla." In the process of concluding that it is not a gorilla, the brain uses attention to test updated predictions and new incoming sensory information against each other as the creature approaches. For example, maybe there really is a gorilla approaching! As the creature approaches, more clearly visible sensory input is detected. When the brain tests the prediction "not gorilla" against this new sensory information, an error signal is triggered, and the prediction may update to "gorilla."

Seth (2021) argues that the involvement of attention in this process is what is responsible for inattentional blindness (IAB). Once again using gorillas as an example, Seth references the classic study by Simons & Chabris (1999) (the same study discussed in our results section). According to Seth, participants failed to notice the gorilla because their attention was engaged in the task of counting the ball passes. While this is the typical explanation given by researchers, Seth (2021) explains *why* the diversion of attention results in IAB, not merely that it does. Participants missed the gorilla because attention was diverted, and attention is needed for an effective error minimization process. Participants did not predict a gorilla, so neither did they see one. As soon as researchers told participants about the existence of a gorilla in the video, "gorilla" became a prior, and attention acted to search for and more heavily-weight related sensory information, causing an update to the prediction which enabled subjects to consciously perceive the gorilla.

The example of the gorilla serves to show how, according to Seth's (2021) theory, we cannot see that which we do not predict. Rather, what we see is that which we predict. Predictive processing can help explain why we have an intuitive sense of peripheral detail–it is because we predict it will be there based on out past experience (prior) of detail when we go to look for it.

Sensorimotor theory of perception

Perception is an active, exploratory process

Actively looking for and encountering detail is hypothesized by some to be an important part of perception (Noë, 2004; O'Regan, 2011; Seth, 2021). Seth (2021) argues that actively engaging the environment is an important way of encountering new sensory information that can aid the error minimization process. This frames perceptual consciousness as an active process. Instead of being passive, where some type of final image is presented to consciousness, what is perceived relies on constant interaction with the environment. This constant interaction provides access to new sensory information, including ways of testing and determining the reliability of certain sensory information, which helps the ever-ongoing process of making better predictions.

The sensorimotor theory of perception is another theory that emphasizes active engagement with the environment. Similar to predictive processing, the sensorimotor theory of perception explains why perception *seems* rich. It does so while also accounting for instances of change and inattentional blindness. J. Kevin O'Regan (2011, p.23) provides a thought experiment ("the hand-in-the-bag game") that helps show how the sensorimotor theory can account for poor performance in experiments involving brief visual presentations (such as change and inattentional blindness paradigms). O'Regan asks us to imagine that someone has put an object–let's say a banana, for our example–into a bag. You are then asked to put your hand in the bag, and to guess the identity of the object. However, the catch is that you cannot move your hand around inside the bag, feeling the object in different ways. You must make your guess based only on the initial point of contact. O'Regan highlights how this seemingly small restriction can immensely impact the difficulty of the task. If you were able to freely explore the item, it would not be long at all until you figured out that it was a banana. However, only given one vantage point, who knows when the guessing will end!

O'Regan argues that vision works much like touch in this way. In order to piece a scene together–to form concrete, identifiable objects out of various combinations of sizes, orientations, and colors–we must *explore* the scene visually (with eye and head movements, for example), actively *enacting* what we see, similar to how we have to explore the object in the bag in order to perceive it as something identifiable. The reason why phenomena such as CB occur, according to O'Regan, is in large part because participants in CB studies are not afforded the opportunity to explore the scenes they are so briefly presented with.

Sensorimotor theory and predictive processing are highly compatible as theories of perceptual consciousness. According to Seth's (2021) interpretation of predictive processing, if we are conscious of predictions, then we are necessarily only conscious of what we are actively engaging in the error minimization process. According to sensorimotor perception, we only perceive that which we are actively manipulating (exploring). In this way, the two theories can be combined: we are conscious of predictions, and these predictions are formed via a process of exploring the environment as described by sensorimotor theory. Additionally, our experience of exploring the environment results in the formation of *sensorimotor contingencies* (Noë, 2004). Sensorimotor contingencies can be understood as knowledge of how motion impacts sensory experience (for example, if I walk closer to this building, it will appear taller). Such information based on past experience is crucial to the error minimization process, as it contributes to calculations of the reliability of priors and sensory information. In this way, sensorimotor contingencies can be seen as priors.

But there remains the question: why, if I can only "see" what I am actively manipulating, does it feel as if I am perceiving my surroundings in rich detail? Not only the unattended periphery, but how can I also feel a sense of awareness of the space behind me, above me, and beneath me? For example, as I write this, I am sitting at a table next to a window. Outside of the window, far outside the scope of even my periphery, is a grove of trees. I know this because I sit here quite often. How am I able to experience the *perceptual presence* of these trees, despite not "actively manipulating" them?

Philosopher Alva Noë refers to the *visual potential* of a space, object, etc. He defines visual potential as "...the way its [object's, space's] aspect changes as a result of movement…" (2004, p.77). Visual potential can be interpreted as a sort of prediction about an object or space, like that discussed by predictive processing. From sitting here before, I know that there are trees off my left shoulder, so when I turn my *attention* (not necessarily my head or eyes) to that region, I engage it in active manipulation and a prediction is made based on my past experiences of seeing trees there. This is also considered a sensorimotor contingency. For example, clearly seeing (the sensory component) the trees outside in full detail is *contingent* on my moving my head and eyes (the motor component). O'Regan sums this up by saying: "Because the information is available on demand, I have the impression of seeing everything" (2011, p.28). This provides an explanation of why perception *seems* rich despite an inability to report detail. I may not be able to tell you how many branches are on the tree, but I feel as if I could, because all I'd have to do is look!

Contingencies & predictions in the laboratory

The idea of perception as an active process that provides us with a *prediction of richness* (despite not yet having the ability to report the specific details of that richness) can explain a significant portion of the results of papers included in this review. However, this idea is not without its challenges, which will also be discussed.

Predicting color & pattern

Sensorimotor predictions can explain filling-in and the uniformity illusion. An iconic example of this is the study led by Cohen et al. (2020a), where participants perceived color in VR environments despite a large part of the periphery being fully desaturated. Participants began the experiment exploring their environments over a period of up to 7 seconds before the periphery was desaturated. This allowed them to gain sensorimotor experience such that exploratory eye and head movements resulted in color. Additionally, this experience mirrors the daily experience all participants have—exploring the real world full of color. By the time the periphery was desaturated, participants already had all the information needed to make predictions about peripheral content. Because peripheral color was predicted, peripheral color was experienced. The VR environments held up their end of the contingency as well—wherever a participant looked, they saw color in the center of their visual field. Throughout the experiment, there was no reason to ever doubt the contingency. Thus, in Cohen et al.'s (2020a) experiment, the fact that many subjects did not notice the removal of large swaths of peripheral color is not really evidence that they didn't experience color in the periphery in the first place (due to perception being sparse), but instead it may be evidence for a failure to update perceptual priors of a richly colored world (which is what we consciously experience). In other words, subjects in this study may have consciously experienced a richly colored complete visual world that did not match the reality once Cohen et al. (2020a) physically removed the color from the periphery.

Predictions about peripheral content can partially explain the results of Balas & Sinha (2007), where participants perceived images in full color even when parts of the images were desaturated. This was true of both types of desaturated images—when the center was in color and periphery black-and-white, and when the center was black-and-white and the periphery color. For color-center images, the same reasoning can be applied as applied above to Cohen et al. (2020a). Because we have so much experience of encountering color when we investigate the periphery, it makes sense that participants perceived the grayscale periphery of center color images as colored. However, how can sensorimotor predictions explain why participants also perceived grayscale center images in full color? It could be that because color is such a reliable prior (meaning, we have so little evidence to doubt its presence when detected), it trumps the reliability of what is perceived at direct fixation. For example, consider a situation in which these participants were presented with full grayscale images. It would be highly unlikely for any to perceive those in full color. It seems as if even a relatively small amount of color detected in the periphery is, in many cases, strong enough to impact perception at fixation. After all, most of us with healthy vision do not have much experience with a black-and-white world.

This reasoning can be applied to visual elements other than color. Consider the patterns used in the studies that examined the uniformity illusion (Suárez-Panilla et al., 2018; Otten et al., 2017). In these cases, fixation was the more reliable prior, and the central patterns resulted in predictions that the periphery would match. This is consistent with the findings of Solovey et al. (2015), where participants employed a more liberal detection criterion in the periphery than in the center of the visual field. This is evidence that the brain is aware of peripheral shortcomings, and, like Heinrich & Bach's (2010) visuo-cognitive hypothesis, takes this into account during the error minimization process. The brain "knows" that sensory detection is weaker at the periphery, so peripheral content is a less reliable prior than central content. This could be why central content determines how peripheral content appears in situations such as the uniformity illusion.

Finally, back to where it all began, consider the Sperling (1960) paradigm. Participants felt that they saw all the letters, despite not being able to report their detail. This situation is a perfect embodiment of O'Regan's (2011) statement: "Because the information is available on demand, I have the impression of seeing everything" (p.28). Participants saw letters at fixation and used this information to predict that letters would also be in the periphery. Because they predicted letters, they saw the letters. If what we see is a prediction and not reality directly, this could explain how we are able to see detail that isn't "there" (we can't report, in the case of Sperling [1960], or in reality isn't there, like the color in Cohen et al., [2020a]).

The world as outside memory

Briefly introduced when discussing the results of O'Regan et al. (1999) is the idea of *the world as outside memory.* O'Regan discusses this idea in more depth in his 2011 book, using the example of imagining his grandmother's face:

> “The mental image does not light up simultaneously as when you switch on your television….The details of her face are not all ‘painted in’ simultaneously to some kind of inner picture that contains these details. The details remain implicit until I ask myself about each particular item: Is she wearing her gold or her silver earrings? Is she wearing her glasses? The more I consider her face, the more parts of it I can detail….Finally when I’m in a state of mind in which I know all these details are immediately mentally available to myself, I can say I am having a mental image of my grandmother.” (30)

Rather than an entire scene in full detail appearing to the perceiver, the perceiver is required to actively piece together the details of the scene. This is why some details can be missed–a scene was presented too briefly for the perceiver to actively probe and construct that element of the scene. Why that same perceiver feels a sense of richness in day-to-day life, however, is because they know that detail is readily available to them if they look for it in the outside world.

Is this reasoning compatible with evidence for preserved representations in CB, described in Mitroff et al. (2004) and Simons et al. (2002)? An important detail to consider from Simons et al. (2002) is that, unlike other CB studies (including those papers discussed in the sensory memory section), participants were exposed to the key stimulus (the ball held by the woman) for an extended period of time. If this allowed them to represent details from the scene, different mechanisms may have been at play than in instances of brief presentation. Interestingly, a fair number of participants were able to answer questions about the ball, however recall that in most instances participants were only able to report details when asked specific, prompting questions by researchers. The ability to answer questions about the ball in the absence of the ball would not necessarily challenge the idea of world as outside memory, as both cases could be true: the ability to form and preserve detailed representations of some things, while offloading information about other things to the external hard drive of the environment. It is possible that those participants who failed to report details about the ball offloaded that information, while those who did report details represented them.

The results of Varakin & Levin (2010) pose an interesting challenge to the idea of outside memory. Varakin & Levin (2010) found that, under analogous circumstances, participants exhibited CB to some stimuli yet were able to recall other stimuli from long-term memory. This led Varakin & Levin (2010) to conclude that sufficiently detailed representations of stimuli in CB experiments are formed, otherwise how could they be encoded in long-term memory? They argue that because CB still occurred, change detection tasks might not best capture the nature and function of these representations, and CB may lead researchers to form false conclusions that the representations formed during CB are sparse. Is it possible for both explanations of CB—Varakin & Levin's (formation of a rich representation that is for some reason not recalled) and O'Regan's (the world as outside memory)—to be true? Can something be represented both internally and externally? Could both be true, just in different circumstances?

Representation & the Cartesian theater

The proposed existence of internal, mental representations often runs up against a conceptual wall. Considering the folk idea of the *conscience* as an example, we often think about consciousness as some sort of abstraction of our mental life. A sort of inner eye that oversees our thoughts, feelings, and experiences (see Graziano [2020] for a compelling explanation of how the attention schema could give rise to this intuitive conclusion). The idea of things being represented "in" the mind has is linked to such a belief in an inner eye. For example, consider the discussion of the quality of representations involved in CB: who or what is this scene represented for?

This question was taken up by philosopher Daniel Dennett (1991) in his discussion of the *Cartesian theater* of consciousness. The Cartesian theater can be thought of as a sort of stage, or screen, in which certain thoughts, sights, sounds, etc. appear "in" consciousness. This conceptualizes consciousness as a container:

that which is inside is conscious, that which is outside is unconscious. This passive idea of consciousness as something that is presented *for* is in stark contrast to sensorimotor theories that conceptualize consciousness as an active process that is constructed *by.* Mechanistically, the idea of consciousness as something constructed *by* is easier to parse than something that is constructed *for,* because the former does not have to explain: for who/what, precisely? This makes the world as outside memory an attractive idea, because instead of solving the problem of representations and the Cartesian theater, it overwrites it entirely. Would it be beneficial to disregard the concept of internal representation altogether? Representations have long been a heavily researched topic in cognitive science, and their dismissal would not come easily. Sensorimotor theories that propose the existence of external representations should be considered in future research and conversations regarding mental representations, in hopes that eventually the "for who?" problem of the container will be solved, by a more comprehensive understanding of representations or/and overwhelming evidence for external rather than internal representation (as is the case if the world is outside memory).

Concluding remarks

Ecological validity

Two papers (Melcher 2006 & Melcher 2010) included in the results stood out among the rest by studying not brief, momentary awareness, but rather the accumulation of detail in memory over time. For this reason, the results of these two papers cannot be directly compared to the results of the other papers included in this review, as they examine quite different phenomena. However, Melcher's (2006 & 2010) design highlights an aspect of real-world awareness that other paradigms failed to address.

More often than not, we find ourselves in stable and familiar environments: our bedrooms, offices, cars, etc. We spend prolonged periods of time in these places and return to them frequently. This allows us to, in Melcher's words, "accumulate" information about these scenes. As I type this at my desk, I can see my bedroom door in my peripheral vision. I know there is artwork hung on my door, despite not being able to actually see it. Still, it phenomenally feels as if it were there—I can sense its *perceptual presence*. Additionally, I could, if asked, report the content and placement of most of the artwork. It makes sense that someone in a situation such as mine would report having an intuitively rich experience of their environment, even though they are also aware that they cannot, without turning their head, actually *see* the very content that contributes to their experience of richness. This type of awareness is similar to the example of the grove of trees discussed on page 68.

Are we looking in the wrong places by studying awareness of brief presentations of stimuli? What real-world situations do these paradigms seek to emulate? While we most often find ourselves in familiar environments, this is not always the case. Imagine driving 60mph in a car, down a road you have never traveled before. Trees, signs, and other vehicles whir past. Maybe there are some clouds, or some mountains, further away from the road that you can focus on for a sustained period of time, but the majority of stimuli in this situation are incredibly fleeting. Considering this example, it makes sense that change and inattentional blindness have been studied in some applied research settings due to their relevance for car accidents and road safety. Similarly, can you think back to the day you moved into your current residence? Can you remember what it was like to experience your now familiar surroundings before they were so deeply encoded in your memory? The study of brief visual presentations may help us understand these kinds of quickly changing or uniquely novel situations; however research on perceptual richness/sparseness must also make room for the effects of long-term memory and accumulation on phenomenal experience.

Additionally, according to a predictive processing framework, even situations like a car ride or a new house rely on accumulation in the form of priors. For example, maybe it is your first time in a specific house, but it is not your first time ever stepping foot into a house. Likewise, maybe you have never driven down this particular road, but you have certainly driven down others. These past experiences develop priors that assist your perception in these novel situations. The development of such priors is another good reason to examine the effects of detail accumulation and memory on subjective richness.

Intuition, introspection, & language

As we have seen, the way we think and talk about perception and consciousness can impact what is researched and how that research is approached. This is evident in how Block's (1995) access-phenomenal distinction has shaped the field since its introduction. Yet on an even more fundamental level, the topic of perceptual richness rests on the idea that we all intuitively feel that perception is rich. Why is this an assumption, and why should it be questioned?

It is generally accepted that most of us feel as if the world is rich with detail. There is a great deal of truth to this. However, it is also true that most humans with healthy vision, if asked, would admit that their vision is not as detailed in the periphery as it is at fixation. Maybe not all of us realize exactly *how* poor peripheral vision is (for example, consider the exercise below, which certainly surprised me the first time I tried it), but we aren't entirely naïve to the limitations of our visual systems.

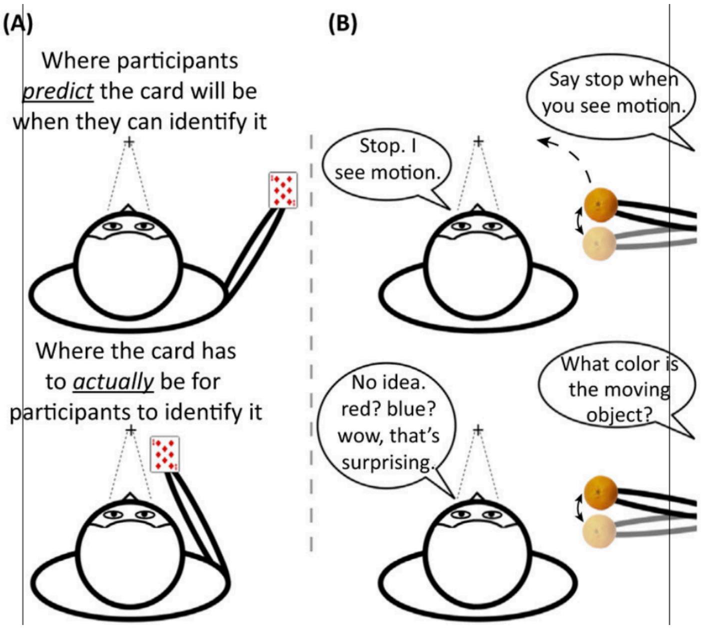

Figure 15: The limits of peripheral vision
Find a friend and try the two exercises pictured above. For exercise (A), predict where you think the playing card will be when you can identify it. Then, have your friend move the card closer in-front of your face. Note the difference between your prediction and where the card is actually. For exercise (B), have your friend find a brightly colored object, but do not let them reveal the color to you. Have your friend move the object around in your periphery, stopping as soon as they get close enough that you see motion. Once you can see motion, your friend should stop moving the object. Can you guess what color it is? Reproduced from Cohen et al., 2016.

How might consciousness research change if the field as a whole placed more weight on the accuracy of our introspections? Certainly, intuition can lead us astray. Think of how surprised participants must have been to realize they missed the gorilla (Simons & Chabris, 1999) or that all but a small portion of their visual field had grown black-and-white (Cohen et al., 2020a). However, also note how simply asking participants a more complex set of questions resulted in wildly different results (Bronfman et al., 2014; Ward et al., 2016), revealing that we may know more about our own visual perception than is often claimed by research.

Future research could benefit from asking participants more complex questions. By obtaining a more thorough understanding of the subjective experiences occurring during perceptual experiments, researchers will be able to study the correlations between subjective experience and objective behavior more specifically.

Future directions

What should be done next based on the theories and experiments discussed so far? Outlined below is a series of open questions that, if answered, could result in considerable progress in the rich vs. sparse debate.

Sensory memory: conscious or unconscious processing?

The first question regards the nature of the contents of sensory memory—namely, are they consciously perceived? Or are they only subject to unconscious processing? The approach taken by Vandenbroucke et al. (2012) is quite clever, using a stimulus believed to depend on conscious perception (the Kanizsa illusion). However, whether the Kanizsa illusion depends on consciousness is not unanimously agreed upon. There is research that shows unconscious processing of the Kanizsa illusion (Persuh et al., 2016). Despite this, most research agrees with the stance taken by Vandenbroucke et al. (2012), with many studies also relying on the assumption that the Kanisza illusion involves explicit, conscious processing.

One way to assess whether the contents of sensory memory are consciously perceived is to determine whether attention is required for sensory memory. A study by Mack et al. (2016) claims that attention is indeed required. Participants in this study had to perform a dual task, and Mack et al. (2016) found that iconic memory benefits disappeared when the dual task diverted attention away from the critical stimuli. If iconic memory has a bandwidth as large as Sligte et al. (2010) claim, and if attention is not required for representation in iconic memory, then why didn't participants in Mack et al. (2016) exhibit sensory memory benefits? One argument may be that the cue in Mack et al. (2016) came a bit later (250ms) than in Sligte et al. (2010) (10ms for iconic; 1000ms for fragile short-term). However, participants in Mack et al. (2016) did not even show the benefits that participants in Sligte et al. (2010) showed in the 1000ms cue condition (longer than 250ms).

It could be argued that participants perceived all the stimuli in Mack et al. (2016), however they weren't able to *access* it. But this implies that sensory memory benefits are benefits to access and falls into the access or no-access trap discussed on pages 57-58.

The results of Mack et al. (2016) seem to support Ward et al.'s (2016) interpretation of Sergent et al. (2013)—that the performance benefits found in Sligte et al. (2010) are due to cueing *attention,* not awareness, implying that items in sensory memory are not consciously perceived but undergo implicit processing. Because sensory memory would offer such a clear and concise explanation of subjective richness, strong evidence on one side (conscious or unconscious) would considerably advance the debate.

Filling-in the predictions?

Another productive question to address concerns the relationship between filling-in (as exhibited in studies examining the uniformity illusion [Otten et al., 2017; Suárez-Panilla et al., 2018] and color spreading [Balas & Sinha, 2007]) and the predictive process. What overlap is there between filling-in and the predictive process? Levinson & Baillet (2022) highlight some important distinctions. Because filling-in has offered a widely accepted explanation for phenomena such as the uniformity illusion, it should be seriously considered as a hypothesis before handing the wheel over to predictive processing.

Levinson & Baillet (2022) discuss the possibility of both filling-in and priors to explain the uniformity illusion. They note that the two explanations are not necessarily compatible, questioning the reliability of central stimuli as a prior. For example, they discuss how real-world environments are not uniform. If our priors are based on accumulated real-world experiences, why would they result in predictions of uniformity in experimental settings, when such uniformity is not experienced in the real world? This is an insightful criticism.

The entire scene of the room I am in right now is not uniform with what I am currently fixating on (my computer). However, an image is not an immersive, panoramic environment. It is often the case that patches of motifs/patterns are uniform in ways the real-world is not. The same is true for text on a page (Sperling, 1960). When discussing stimuli such as this, central stimuli may indeed be a reliable prior. In real-world settings, details in the center of the visual field may act as a prior by predicting that the periphery also probably contains richly detailed content. Specific details themselves are not predicted.

To compare filling-in and predictive processing, future research should carefully manipulate stimuli such that the strength of different priors (such as contrast, uniformity, central or peripheral location, frequency, color, etc.) can be analyzed. How do these priors influence perception of different stimuli? Can this perception be primed by experience with visual surfaces such as text? Can uniformity and filling-in occur in instances where it does not make sense to consider dominant stimuli a prior? Answering these questions can help shed light on the relationship between priors and processes such as filling-in.

The problem of representations

One last question to examine concerns the nature of CB. To reiterate the question posed by Mitroff et al. (2004), does CB occur because: 1) none of the pre-change stimuli were represented in any way 2) they were represented, but the representation was fragile and easily overwritten 3) they were represented in full detail, but the system did not engage the representation effectively to notice the change. The results of Varakin & Levin (2010) are convincing evidence in favor of the last option. Participants displayed long-term memory for the same kinds of stimuli they also exhibited CB for, implying that participants may have represented the stimuli but for some reason unrelated to the quality of the representation failed the change detection task.

The sensorimotor explanation of CB offers yet another competing explanation: participants rely on their ability to access detail in the external world, so there is no need for an internal representation. Participants exhibit CB because they did not engage in active manipulation of the critical portion of the scene.

Future research should seek to determine which one of the following two options best explains CB: the world as outside memory, or a comparison failure? Perhaps both explanations are valid, and they account for CB in different circumstances.

In some sense, both explanations could be attempts at describing the same phenomenon, but, because of different conceptual approaches to the issue of representations, it seems as if they are discussing different mechanisms. In both instances, participants fail to engage a representation—whether it is internal, or in the outside world. From this angle, it could be argued that Varakin & Levin (2010) and O'Regan (2011) are describing the same explanation for why CB occurs, just using two different conceptualizations of representation.

Another approach may be to abandon CB altogether. The brief flashes of images used in change and inattentional blindness paradigms do not at all replicate everyday perception. And this everyday perception is what we base our intuitive claim of rich perception off of. Such an approach would not only be more ecologically valid, but also pay more respect to introspection and the subjective experience itself.

Additionally, future research could more deeply examine the error minimization process in instances of both blindness and "noticing." Instead of the traditional approach of trying to determine the conditions that best capture attention (size, contrast, relevance, etc.) and lead to "noticing," research should try determining the conditions that lead to the most accuracy in the error minimization process. What relative differences between physical stimuli and subjectively perceived prediction are needed for the error signal to be triggered, resulting in an "accurate" perception of the stimulus? Under what conditions does this occur more quickly?

Final remarks on future directions

Settling the questions outlined above would result in considerable progress on the rich vs. sparse debate. Determining whether attention is necessary for representation in sensory memory, whether sensory memory is an unconscious process, the nature of representations in CB, and how priors relate to processes such as filling-in show promise as future directions in this debate. Additionally, one general comment that applies to all topics is that more consideration should be paid to how above- and below-chance performance is interpreted in research on perceptual richness. While some studies cite behavior above chance as strong evidence for perceptual richness (such as 67% and 69%, Varakin & Levin (2010), and 65.88%, Mitroff et al. (2004)), other studies interpret similar behavior as strong evidence for sparseness (such as in Cohen et al., 2020a, participants in the 10 condition failed to notice 33% of the time, which means they *did notice* 66% of the time--why wasn't this above chance performance interpreted as evidence for perceptual richness?).

Returning to the issue of satisfaction

This review began with a discussion of the tensions that arise when trying to merge the epistemic systems of subjectivity and objectivity. Commonly felt dissatisfactions were raised, and ways to mitigate these feelings were proposed. How have we done so far?

Instead of simplifying conscious perception to rich vs. sparse, attempts were made to account for the *subjective* richness that will persist no matter which side the empirical evidence eventually settles on. In the introduction, the analogy of Copernicus's discovery was used to describe the complexity of the debate. This analogy helped distill two questions: is perception rich? if not, why does it seem rich? Two answers were proposed based on the results of this review. Sensory memory provides an affirmative answer to the first question. Sensorimotor predictions account for the "if not." Which of these two answers is correct relies on continued research.

How do the results of this review relate back to the original, broader conversation of phenomenal consciousness and the "what it's like" component of subjective experience? Phenomenal consciousness is thought to be the aspect of consciousness that cannot be studied empirically, due to issues of privacy and reliance on report. Unless some technology is someday invented that allows researchers and participants to share consciousness, the following will always be true: the conscious experience itself can never be observed empirically, only the associated neural activity. This is not enough to satisfy some. Others find this approach completely satisfactory. New insight from sensorimotor predictions may help make this gap more palatable. Like the problem of representations, the "what it is like" component of experience can be thought of not as a final product to locate in the brain, but as an active process. We enact our experience. Still, the experience of this experience will be different from first- and third-person points of view, thus, from an epistemological standpoint, there will be aspects of consciousness that empiricism can never "know." As science continues to make progress on this final frontier, hopefully discussing consciousness in empirical terms will become more familiar. While it will never serve as a satisfactory *replacement* to all individuals—with some people preferring religious explanations, phenomenological explanations based off their own immediate experience, etc.—there is no reason why it should not have a seat at the table.

Satisfaction may ultimately be personal and cultural. The satisfactoriness of an empirical account of consciousness may depend less on the explanation itself and more on the cultural predispositions of populations to find empirical explanations in general satisfactory. There is sometimes an attitude in consciousness research; an emphasis that nothing is at it seems, and that we all go about our lives victim to some great illusion. Those who believe in their intuitive experience of rich perception despite phenomena such as change and inattentional blindness are held in disdain, seen as entertaining a childish, naïve belief. Very often, little recognition is given to the fact that when all is said and done, the sun will still appear to rise and set, moving in an arc across the sky and around the earth.

Hopefully this review has paid more respect to subjective phenomenological evidence by seriously considering both what things *are* and how things *seem.* The scientific study of consciousness is, after all, the study of subjective appearance. Despite this, many have followed the siren song of figuring out how things *actually* are, forsaking the very subjectivity the field claims to study. While both avenues—the actuality and appearance—are equally deserving of inquiry, a comprehensive scientific explanation of consciousness should account for both how things are and how things appear.

Appendix A: Inclusion & Exclusion Criteria

Inclusion criteria

Included papers must:

- Be available in English
- Present original empirical data (no reviews)
- Be published between 1995 and 2021. This is to ensure relevance, including early responses to the access/phenomenal distinction (Block, 1995).
- Be peer-reviewed
- Reference the rich vs. sparse debate in the title or abstract. Because perception is such a commonly used word, this criterion was intended to ensure relevance. A reference could be anywhere from an explicit mention of the debate to simply using terms such as "phenomenal overflow," "rich perception," "access or phenomenal," "phenomenal consciousness," etc.

Exclusion criteria

Excluded papers:

- Were pilot studies
- Were case studies
- Involved the use or administration of psychoactive drugs
- Examined non-human animal models
- Included persons under the age of 18 in their samples
- Included persons with psychiatric conditions in their samples
- Included senior citizens in their samples

Appendix B: Articles Included in the Review

Balas, B., & Sinha, P. (2007). “Filling-in” colour in natural scenes. *Visual Cognition*, *15*(7), 765–778. Scopus. https://doi.org/10.1080/13506280701295453

Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We see more than we can report: “Cost free” color phenomenality outside focal attention. *Psychological Science*, *25*(7), 1394–1403. https://doi.org/10.1177/0956797614532656

Clarke, J., & Mack, A. (2015). Iconic memory for natural scenes: Evidence using a modified change-detection procedure. *Visual Cognition*, *23*(7), 917–938. https://doi.org/10.1080/13506285.2015.1103826

Cohen, M. A., Botch, T. L., & Robertson, C. E. (2020). The limits of color awareness during active, real-world vision. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *117*(24), 13821–13827. https://doi.org/10.1073/pnas.1922294117

de Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, *18*(3), 569–577. https://doi.org/10.1016/j.concog.2009.03.002

Heinrich, S. P., & Bach, M. (2010). Less is more: Subjective detailedness depends on stimulus size. *Journal of Vision*, *10*(10). https://doi.org/10.1167/10.10.2

Melcher, D. (2006). Accumulation and persistence of memory for natural scenes. *Journal of Vision*, *6*(1), 8–17. https://doi.org/10.1167/6.1.2

Melcher, D. (2010). Accumulating and remembering the details of neutral and emotional natural scenes. *Perception*, *39*(8), 1011–1025. https://doi.org/10.1068/p6670

Mitroff, S. R., & Scholl, B. J. (2004). Seeing the Disappearance of Unseen Objects. *Perception*, *33*(10), 1267–1273. https://doi.org/10.1068/p5341no

Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(4), 798–815. https://doi.org/10.1037/0096-1523.28.4.798

Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, *66*(8), 1268–1281. https://doi.org/10.3758/BF03194997

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of "mudsplashes." *Nature*, *398*(6722), 34–34. https://doi.org/10.1038/17953

Otten, M., Pinto, Y., Paffen, C. L. E., Seth, A. K., & Kanai, R. (2017). The uniformity illusion: Central stimuli can determine peripheral perception. *Psychological Science*, *28*(1), 56–68. https://doi.org/10.1177/0956797616672270

Park, S., Chun, M. M., & Johnson, M. K. (2010). Refreshing and integrating visual scenes in scene-selective cortex. *Journal of Cognitive Neuroscience*, *22*(12), 2813–2822. https://doi.org/10.1162/jocn.2009.21406

Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing Attention after the Stimulus Is Gone Can Retrospectively Trigger Conscious Perception. *Current Biology*, *23*(2), 150–155. https://doi.org/10.1016/j.cub.2012.11.047

Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, *28*(9), 1059–1074. https://doi.org/10.1068/p281059

Simons, D. J., Chabris, C. F., Schnur, T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition: An International Journal*, *11*(1), 78–97. https://doi.org/10.1006/ccog.2001.0533

Sligte, I., Vandenbroucke, A., Scholte, H. S., & Lamme, V. (2010). Detailed Sensory Memory, Sloppy Working Memory. *Frontiers in Psychology*, *1*. https://www.frontiersin.org/article/10.3389/fpsyg.2010.00175

Solovey, G., Graney, G. G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception, & Psychophysics*, *77*(1), 258–271. https://doi.org/10.3758/s13414-014-0769-1

Suárez-Pinilla, M., Seth, A. K., & Roseboom, W. (2018). The illusion of uniformity does not depend on the primary visual cortex: Evidence from sensory adaptation. *I-Perception*, *9*(5). https://doi.org/10.1177/2041669518800728

Vandenbroucke, A. R. E., Fahrenfort, J. J., Sligte, I. G., & Lamme, V. A. F. (2014). Seeing without knowing: Neural signatures of perceptual inference in the absence of report. *Journal of Cognitive Neuroscience*, *26*(5), 955–969. https://doi.org/10.1162/jocn_a_00530

Vandenbroucke, A. R. E., Sligte, I. G., Barrett, A. B., Seth, A. K., Fahrenfort, J. J., & Lamme, V. A. F. (2014). Accurate Metacognition for Visual Sensory Memory Representations. *Psychological Science*, *25*(4), 861–873. https://doi.org/10.1177/0956797613516146

Vandenbroucke, A. R. E., Sligte, I. G., Fahrenfort, J. J., Ambroziak, K. B., & Lamme, V. A. F. (2012). Non-attended representations are perceptual rather than unconscious in nature. *PLoS ONE*, *7*(11). https://doi.org/10.1371/journal.pone.0050042

Varakin, D. A., & Levin, D. T. (2006). Change blindness and visual memory: Visual representations get rich and act poor. *British Journal of Psychology*, *97*(1), 51–77. https://doi.org/10.1348/000712605X68906

Wang, R., Fu, Y., Chen, L., Chen, Y., Zhou, J., & Chen, H. (2021). Consciousness can overflow report: Novel evidence from attribute amnesia of a single stimulus. *Consciousness and Cognition*, *87*, 103052. https://doi.org/10.1016/j.concog.2020.103052

Ward, E. J., Bear, A., & Scholl, B. J. (2016). Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining

whether awareness overflows access. *Cognition*, *152*, 78–86.
https://doi.org/10.1016/j.cognition.2016.01.010

Bibliography

Anscombe, G. E. M. (1959). *An Introduction to Wittgenstein's Tractatus*. Hutchinson University Library.

Balas, B., & Sinha, P. (2007). "Filling-in" colour in natural scenes. Visual Cognition, 15(7), 765–778. Scopus. https://doi.org/10.1080/13506280701295453

Block, N. (1995). On a confusion about a function of consciousness. Behavioral and Brain Sciences, 18(2), 227–247. https://doi.org/10.1017/S0140525X00038188

Block, N. (2011). Perceptual consciousness overflows cognitive access. Trends in Cognitive Sciences, 15(12), 567–575. https://doi.org/10.1016/j.tics.2011.11.001

Bourget, D., & Chalmers, D. J. (2014). What Do Philosophers Believe? Philosophical Studies, 170(3), 465–500. https://doi.org/10.1007/s11098-013-0259-7

Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We see more than we can report: "Cost free" color phenomenality outside focal attention. *Psychological Science*, *25*(7), 1394–1403. https://doi.org/10.1177/0956797614532656

Campbell, K. K. (1970). *Body and Mind*. Doubleday.

Clarke, J., & Mack, A. (2015). Iconic memory for natural scenes: Evidence using a modified change-detection procedure. *Visual Cognition*, *23*(7), 917–938. https://doi.org/10.1080/13506285.2015.1103826

Cohen, M. A., Botch, T. L., & Robertson, C. E. (2020a). The limits of color awareness during active, real-world vision. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *117*(24), 13821–13827. https://doi.org/10.1073/pnas.1922294117

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

Cohen, M. A., Ortego, K., Kyroudis, A., & Pitts, M. (2020b). Distinguishing the Neural Correlates of Perceptual Awareness and Postperceptual Processing. *Journal of Neuroscience*, *40*(25), 4925–4935. https://doi.org/10.1523/JNEUROSCI.0120-20.2020

Crick, F., & Koch, C. (2003). A framework for consciousness. *Nature Neuroscience*, *6*(2), 119–126. https://doi.org/10.1038/nn0203-119

de Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, *18*(3), 569–577. https://doi.org/10.1016/j.concog.2009.03.002

Dehaene, S. (2014). *Consciousness and the Brain: Deciphering How the Brain Codes Our Thoughts*. Penguin.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*(1), 1–37. https://doi.org/10.1016/S0010-0277(00)00123-2

Del Cul, A., Dehaene, S., Reyes, P., Bravo, E., & Slachevsky, A. (2009). Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain*, *132*(9), 2531–2540. https://doi.org/10.1093/brain/awp111

Dennett, D. C. (2017). *Consciousness Explained*. Little, Brown.

Graziano, M. S. A. (2019). *Rethinking Consciousness: A Scientific Theory of Subjective Experience*. W. W. Norton & Company.

Graziano, M. S. A. (2020). Consciousness and the attention schema: Why it has to be right. *Cognitive Neuropsychology*, *37*(3–4), 224–233. https://doi.org/10.1080/02643294.2020.1761782

Graziano, M. S. A., Guterstam, A., Bio, B. J., & Wilterson, A. I. (2020). Toward a standard model of consciousness: Reconciling the attention schema, global workspace, higher-order thought, and illusionist theories. *Cognitive*

Neuropsychology, *37*(3–4), 155–172. https://doi.org/10.1080/02643294.2019.1670630

Harris, J. J., Schwarzkopf, D. S., Song, C., Bahrami, B., & Rees, G. (2011). Contextual Illusions Reveal the Limit of Unconscious Visual Processing. *Psychological Science*, *22*(3), 399–405. https://doi.org/10.1177/0956797611399293

Haun, A. M., Tononi, G., Koch, C., & Tsuchiya, N. (2017). Are we underestimating the richness of visual experience? *Neuroscience of Consciousness*, *2017*(1), niw023. https://doi.org/10.1093/nc/niw023

Heinrich, S. P., & Bach, M. (2010). Less is more: Subjective detailedness depends on stimulus size. *Journal of Vision*, *10*(10). https://doi.org/10.1167/10.10.2

Kalloniatis, M., & Luu, C. (2007, June 5). *Figure 22, [Spatial frequency is a measure...].* [Text]. University of Utah Health Sciences Center. https://www.ncbi.nlm.nih.gov/books/NBK11509/figure/ch25kallspatial.F22/

Kapoor, V., Dwarakanath, A., Safavi, S., Werner, J., Panagiotaropoulos, T. I., & Logothetis, N. K. (2019). *Neuronal discharges in the prefrontal cortex reflect changes in conscious perception during a no report binocular rivalry paradigm*. https://pure.mpg.de/pubman/faces/ViewItemOverviewPage.jsp?itemId=item_3164050

Kirk, R. (1974). Sentience and Behaviour. *Mind*, *83*(329), 43–60.

Koch, C. (2019). *The Feeling of Life Itself: Why Consciousness Is Widespread but Can't Be Computed*. MIT Press.

Koch, C., Massimini, M., Boly, M., & Tononi, G. (2016). Neural correlates of consciousness: Progress and problems. *Nature Reviews Neuroscience*, *17*(5), 307–321. https://doi.org/10.1038/nrn.2016.22

Levinson, M., & Baillet, S. (2022). Perceptual filling-in dispels the veridicality problem of conscious perception research. *Consciousness and Cognition*, *100*, 103316. https://doi.org/10.1016/j.concog.2022.103316

Lewis, D. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, *50*(3), 249–258. https://doi.org/10.1080/00048407212341301

Locke, J. (1700). *An Essay Concerning Human Understanding*.

Mack, A., Erol, M., Clarke, J., & Bert, J. (2016). No iconic memory without attention. *Consciousness and Cognition*, *40*, 1–8. https://doi.org/10.1016/j.concog.2015.12.006

Mashour, G. A., Roelfsema, P., Changeux, J.-P., & Dehaene, S. (2020). Conscious Processing and the Global Neuronal Workspace Hypothesis. *Neuron*, *105*(5), 776–798. https://doi.org/10.1016/j.neuron.2020.01.026

Melcher, D. (2006). Accumulation and persistence of memory for natural scenes. *Journal of Vision*, *6*(1), 8–17. https://doi.org/10.1167/6.1.2

Melcher, D. (2010). Accumulating and remembering the details of neutral and emotional natural scenes. *Perception*, *39*(8), 1011–1025. https://doi.org/10.1068/p6670

Metzinger, T. (2010). *The Ego Tunnel: The Science of the Mind and the Myth of the Self*. Basic Books.

Mitroff, S. R., & Scholl, B. J. (2004). Seeing the Disappearance of Unseen Objects. *Perception*, *33*(10), 1267–1273. https://doi.org/10.1068/p5341no

Mitroff, S. R., Simons, D. J., & Franconeri, S. L. (2002). The siren song of implicit change detection. *Journal of Experimental Psychology: Human Perception and Performance*, *28*(4), 798–815. https://doi.org/10.1037/0096-1523.28.4.798

Mitroff, S. R., Simons, D. J., & Levin, D. T. (2004). Nothing compares 2 views: Change blindness can occur despite preserved access to the changed information. *Perception & Psychophysics*, *66*(8), 1268–1281. https://doi.org/10.3758/BF03194997

Naccache, L. (2018). Why and how access consciousness can account for phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1755), 20170357. https://doi.org/10.1098/rstb.2017.0357

Nagel, T. (1974). What Is It Like to Be a Bat? *The Philosophical Review*, *83*(4), 435–450. https://doi.org/10.2307/2183914

Newman, J., & Baars, B. J. (1993, December). *A Neural Attentional Model for Access to Consciousness: A Global Workspace Perspective* (Journal (Paginated) No. 2). Concepts in Neuroscience. https://web-archive.southampton.ac.uk/cogprints.org/73/

Noe, A. (2004). *Action in Perception*. MIT Press.

O'Regan, J. K. (2011). *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*. Oxford University Press, USA.

O'Regan, J. K., Rensink, R. A., & Clark, J. J. (1999). Change-blindness as a result of "mudsplashes." *Nature*, *398*(6722), 34–34. https://doi.org/10.1038/17953

Otten, M., Pinto, Y., Paffen, C. L. E., Seth, A. K., & Kanai, R. (2017). The uniformity illusion: Central stimuli can determine peripheral perception. *Psychological Science*, *28*(1), 56–68. https://doi.org/10.1177/0956797616672270

Park, S., Chun, M. M., & Johnson, M. K. (2010). Refreshing and integrating visual scenes in scene-selective cortex. *Journal of Cognitive Neuroscience*, *22*(12), 2813–2822. https://doi.org/10.1162/jocn.2009.21406

Persuh, M., Emmanouil, T. A., & Ro, T. (2016). Perceptual overloading reveals illusory contour perception without awareness of the inducers. *Attention, Perception, & Psychophysics*, *78*(6), 1692–1701. https://doi.org/10.3758/s13414-016-1146-z

Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S. A. (2014). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, *101*, 337–350. https://doi.org/10.1016/j.neuroimage.2014.07.024

Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing Attention after the Stimulus Is Gone Can

Retrospectively Trigger Conscious Perception. *Current Biology*, *23*(2), 150–155. https://doi.org/10.1016/j.cub.2012.11.047

Seth, A. (2021). *Being You: A New Science of Consciousness*. Penguin.

Shanahan, M., & Baars, B. (2007). Global workspace theory emerges unscathed. *Behavioral and Brain Sciences*, *30*(5–6), 524–525. https://doi.org/10.1017/S0140525X07003056

Simons, D. J., & Chabris, C. F. (1999). Gorillas in Our Midst: Sustained Inattentional Blindness for Dynamic Events. *Perception*, *28*(9), 1059–1074. https://doi.org/10.1068/p281059

Simons, D. J., Chabris, C. F., Schnur, T., & Levin, D. T. (2002). Evidence for preserved representations in change blindness. *Consciousness and Cognition: An International Journal*, *11*(1), 78–97. https://doi.org/10.1006/ccog.2001.0533

Sligte, I., Vandenbroucke, A., Scholte, H. S., & Lamme, V. (2010). Detailed Sensory Memory, Sloppy Working Memory. *Frontiers in Psychology*, *1*, 175. https://doi.org/10.3389/fpsyg.2010.00175

Smith, D. W. (2018). Phenomenology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2018). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2018/entries/phenomenology/

Solovey, G., Graney, G. G., & Lau, H. (2015). A decisional account of subjective inflation of visual perception at the periphery. *Attention, Perception, & Psychophysics*, *77*(1), 258–271. https://doi.org/10.3758/s13414-014-0769-1

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1–29. https://doi.org/10.1037/h0093759

Suárez-Pinilla, M., Seth, A. K., & Roseboom, W. (2018). The illusion of uniformity does not depend on the primary visual cortex: Evidence from sensory adaptation. *I-Perception*, *9*(5). https://doi.org/10.1177/2041669518800728

Tricco, A. C., Lillie, E., Zarin, W., O'Brien, K. K., Colquhoun, H., Levac, D., Moher, D., Peters, M. D. J., Horsley, T., Weeks, L., Hempel, S., Akl, E. A., Chang, C., McGowan, J., Stewart, L., Hartling, L., Aldcroft, A., Wilson, M. G., Garritty, C., … Straus, S. E. (2018). PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Annals of Internal Medicine*, *169*(7), 467–473. https://doi.org/10.7326/M18-0850

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. F. (2015). No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences*, *19*(12), 757–770. https://doi.org/10.1016/j.tics.2015.10.002

Vandenbroucke, A. R. E., Fahrenfort, J. J., Sligte, I. G., & Lamme, V. A. F. (2014a). Seeing without knowing: Neural signatures of perceptual inference in the absence of report. *Journal of Cognitive Neuroscience*, *26*(5), 955–969. https://doi.org/10.1162/jocn_a_00530

Vandenbroucke, A. R. E., Sligte, I. G., Barrett, A. B., Seth, A. K., Fahrenfort, J. J., & Lamme, V. A. F. (2014b). Accurate Metacognition for Visual Sensory Memory Representations. *Psychological Science*, *25*(4), 861–873. https://doi.org/10.1177/0956797613516146

Vandenbroucke, A. R. E., Sligte, I. G., Fahrenfort, J. J., Ambroziak, K. B., & Lamme, V. A. F. (2012). Non-attended representations are perceptual rather than unconscious in nature. *PLoS ONE*, *7*(11). https://doi.org/10.1371/journal.pone.0050042

Varakin, D. A., & Levin, D. T. (2006). Change blindness and visual memory: Visual representations get rich and act poor. *British Journal of Psychology*, *97*(1), 51–77. https://doi.org/10.1348/000712605X68906

Wang, R., Fu, Y., Chen, L., Chen, Y., Zhou, J., & Chen, H. (2021). Consciousness can overflow report: Novel evidence from attribute amnesia of a single stimulus. *Consciousness and Cognition*, *87*, 103052. https://doi.org/10.1016/j.concog.2020.103052

Ward, E. J., Bear, A., & Scholl, B. J. (2016). Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining whether awareness overflows access. *Cognition*, *152*, 78–86. https://doi.org/10.1016/j.cognition.2016.01.010

Whyte, C. J., Hohwy, J., & Smith, R. (2022). An active inference model of conscious access: How cognitive action selection reconciles the results of report and No-Report paradigms. *Current Research in Neurobiology*, 100036. https://doi.org/10.1016/j.crneur.2022.100036