This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

Is the phenomenological overflow argument really supported by subjective reports?

Florian Cova

Department of Philosophy & Swiss Center for Affective Sciences, University of Geneva

Maxence Gaillard

Interfaculty Initiative in Information Studies, The University of Tokyo

François Kammerer

Institut Supérieur de Philosophie, Université catholique de Louvain, FNRS

Correspondence

Florian Cova, Université de Genève, Département de philosophie, 5 rue De-Candolle, CH-1211 Genève 4, Switzerland

Email: florian.cova@gmail.com

Does phenomenal consciousness overflow access-consciousness? Some researchers have claimed that it does, relying on interpretations of various psychological experiments such as Sperling's or Landman's, and crucially using alleged subjective reports from participants to argue in favor of these interpretations. However, systematic empirical investigations of participants' subjective reports are scarce. To fill this gap, we reproduced Sperling's and Landman's experiments, and carefully collected reports made by subjects about their own experiences, using questionnaires and interviews. We found that participants' subjective reports do not support the overflow thesis, but rather suggest alternative interpretations of the experiments.

KEYWORDS

phenomenal consciousness; access-consciousness; overflow; subjective data; introspection; Sperling

1. INTRODUCTION

One important current debate in the science of consciousness bears on whether we should embrace the distinction between "phenomenal consciousness" (P-consciousness) and "access consciousness" (A-consciousness). A number of researchers in psychology and neuroscience

(Lamme, 2003; Milner & Goodale, 2008; Zeki, 2003), some of them explicitly following Block (1995), have accepted such a distinction. Other researchers, such as those working within the *global workspace* framework (Baars, 1988; Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Dehaene & Naccache, 2001), deny that we should ultimately distinguish between these two forms of consciousness.

One of the most crucial arguments in favor of this distinction is the *overflow argument* (Block, 2007, 2011), according to which we have to make the distinction because the content of P-consciousness is richer than the content of A-consciousness. Proponents of the overflow argument usually consider that it is supported by various psychological experiments, such as Sperling's famous experiment on iconic memory (Sperling, 1960), or more recent experiments on change detection (Landman, Spekreijse, & Lamme, 2003). However, they admit that the objective data provided by such experiments are compatible with other views—views that do not recognize the reality of phenomenal overflow. In order to strengthen their case, they appeal to *subjective reports* from participants—subjects' reports about their own experiences.

Opponents of the overflow thesis have often tried to dismiss the use of such subjective reports, stating for example that subjects might be victims of an illusion regarding their own experience—notably when it comes to the *richness* of their experience—so that we should not take subjective reports at face value (Kouider, de Gardelle, Sackur, & Dupoux, 2010). Others have argued that data provided by participants' subjective reports are too vague to provide direct support for the overflow thesis and that they are also compatible with other, competing views (Phillips, 2011).

Here, we track down the origin of these claims about subjective data: They stem merely from informal, imprecise and under-determined reports. In order to know whether the overflow argument is really supported by participants' subjective reports, we present the results of two studies in which we systematically collected subjective data. Using experimental paradigms from Sperling's and Landman's experiments, we probed participants' subjective experience, using quantitative questionnaires and qualitative interviews. We found that, when asked about their subjective experiences, participants give diverse reports, most of which do *not* support the overflow interpretation of these experiments.

2. PHENOMENAL CONSCIOUSNESS AND ACCESS CONSCIOUSNESS

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

Ned Block first introduced the distinction between phenomenal consciousness and access consciousness as a *conceptual* distinction (Block, 1995). He suggested to distinguish—at least conceptually and maybe empirically—between two forms of consciousness: phenomenal consciousness ("P-consciousness") and access consciousness ("A-consciousness").

What does this distinction amount to? As Block puts it, "P-consciousness is experience" (Block, 1995, p. 230). P-conscious states are states such that there is something it is like for the subject to be in them. Their content is inherently *phenomenal* (Block, 1995, p. 232). P-consciousness is *not* defined in functional terms (in terms of the role played by P-conscious states in a cognitive system), even though P-conscious states *may play* certain functional roles. The paradigmatic P-conscious states are sensations.

A-consciousness, on the other hand, is defined as follows:

> [A] state is access-conscious (A-conscious) if, in virtue of one's having the state, a representation of its content is (1) … poised for use as a premise in reasoning, (2) poised for use for rational control of action, (3) poised for rational control of speech. (Block, 1995, p. 231).

The content of A-conscious states is *representational,* and such states are defined *functionally,* as states playing a certain kind of functional role in a cognitive system.

Block's original paper presented and defended this distinction as a distinction in our *concepts* of consciousness. Most (but not all) of his arguments in defense of this conceptual distinction were based on thought experiments: descriptions of *merely possible* situations in which A-consciousness and P-consciousness come apart. He left open the possibility that A-consciousness and P-consciousness, although distinct from a conceptual point of view, could empirically be shown to be *identical*: "[P]erhaps P-consciousness and A-consciousness amount to much the same thing empirically even though they differ conceptually" (Block, 1995, p. 242).

Block's distinction has been extremely influential. As a conceptual distinction, is has been widely accepted in the philosophical and the scientific study of consciousness – at least as a starting point.[1] Things are different when it comes to the *empirical* understanding of the distinction. Major contemporary theories of consciousness differ on whether P-consciousness

[1] Some critics of this distinction have argued that the concept of phenomenal consciousness (conceived of as a non-functionally defined form of consciousness) is flawed and that it is an obstacle to scientific research on consciousness (Cohen & Dennett, 2011, p. 362–363; Kouider, de Gardelle, Sackur, & Dupoux, 2010, p. 304). However, most researchers accept the distinction between phenomenal consciousness and access consciousness as prima facie convincing from a conceptual point of view, even though scientific inquiry might lead us to reject it.

and A-consciousness come apart empirically. Some theorists, sometimes explicitly following Block, claim that we should empirically distinguish between access consciousness and phenomenal consciousness. According to these "dissociative" (Cohen & Dennett, 2011, p. 359) views of consciousness, the two forms of consciousness can come apart, and they do not correspond to the same kind of brain states (Lamme, 2003, 2006; Milner & Goodale, 2008; Zeki, 2003).

Other researchers have defended an "integrative" view of consciousness. For them, phenomenal consciousness and access consciousness really amount to the same thing, and phenomenal consciousness typically cannot come apart from a certain form of access consciousness. Researchers working within the *global workspace* framework (Baars, 1988; Dehaene et al., 2006; Dehaene & Naccache, 2001; Kouider et al., 2010) have typically defended this kind of view – although they are not alone in this (Cohen & Dennett, 2011). The debate between these two conceptions has been one of the most heated and fundamental debates within the field of the scientific and philosophical study of consciousness.

The main argument in favor of not only a *conceptual* but also *empirical* distinction between phenomenal consciousness and access consciousness has come progressively to be known as the "overflow argument". A first version (Block, 1995, pp. 234, 244) of this argument can be found in Block's seminal paper, "On a confusion about a function of consciousness."[2] The importance of the argument increased later as Block reinforced and systematized this line of thought (Block, 2007, 2011). The overflow argument in itself is quite simple and goes as follows. (1) *Some contents are phenomenally conscious but not access conscious*. This is the *overflow thesis*, which amounts to saying that the content of phenomenal consciousness *overflows* the content of access consciousness. However, (2) *if phenomenal consciousness were not empirically distinct from access consciousness, no content would be phenomenally conscious without being access conscious*. Therefore, (3) *phenomenal consciousness is empirically distinct from access consciousness*. The argument is convincing only if one accepts its crucial and controversial premise: the overflow thesis. Consequently, a substantial part of the debate between dissociative and integrated conceptions of consciousness focused on the overflow thesis.

[2] In this part of the paper, Block goes as far as claiming that there is an empirical, and not only conceptual, distinction between P-consciousness and A-consciousness.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

3. THE OVERFLOW THESIS AND THE USE OF SUBJECTIVE REPORTS

The most robust defense of the overflow thesis appeals to the interpretation of the results of some classic psychological experiments: first and foremost, Sperling's 1960 series of experiments (Sperling, 1960), and, to a certain extent, other experiments, such as Landman's 2003 experiment (Landman et al., 2003)—hereafter respectively Sperling's experiment and Landman's experiment.

In Sperling's experiment (Block, 2007, p. 487; Sperling, 1960), arrays of characters (6 to 12 letters) were flashed to participants for a short amount of time (between 15 and 500ms). When participants were asked to report the letters they had seen, they were able to report only about four of the letters. Sperling then modified the experiment, by playing a tone after the array of characters was replaced by a blank. Participants were told to report the top row if the tone was high, the middle row is the tone was intermediate, or the bottom row is the tone was low. In this version, participants were able to report *almost all* the characters in the relevant row.

According to Block, the best interpretation of this experiment is the following: participants were *phenomenally conscious* of *all the letters* in a detailed and determinate way, but the limited capacity of their working memory made it impossible for all of the specific information about the letters to become access conscious at the same time. What was *access conscious* was only some specific information about four letters, while the access conscious information about the rest of the letters was merely of a generic kind (they were represented as "letters" and not as such and such particular letters).

In Landman's experiment (Block, 2007, p. 488; Landman et al., 2003), participants were flashed for half a second eight rectangles arranged around a dot located in the middle of the screen. Participants were supposed to keep looking at the dot. After the flash, the rectangles were replaced with a blank. Then, another array appeared with eight rectangles (at the same location as previously) and a cue pointing at one of the rectangles. Participants had to say whether the indicated rectangle had changed orientation. Typically, participants were able to track the orientation of about four items (out of eight). However, a second version featured a slight variation: the cue indicating the rectangle about which subjects have to make their report was shown *before* the blank, when the eight rectangles were shown for the first time. Participants then unsurprisingly gave correct answers most of the time. Finally, in a third version, the line was shown *during* the blank (after the disappearance of the eight initial rectangles, but before the apparition of the second group of rectangles). Participants' reports were then almost as accurate as in the second version.

According to Block, such results are best explained in the same way as the results of Sperling's experiment. Participants were phenomenally conscious of all the rectangles in a detailed way (with their specific orientation) from the start, but the limited capacity of their working memory made it impossible for them to consciously *access* the specific orientation of each of the rectangles. However, when participants were shown a cue (even during the blank), this made the relevant information (already phenomenally conscious) access conscious, allowing participants to report accurately the change in orientation. Landman's experiment, like Sperling's experiment, supposedly shows that the content of phenomenal consciousness overflows the content of access consciousness.

Many researchers noted that this interpretation is *not* the only interpretation compatible with the experimental data (Block, 2011; Brown, 2012; Cohen & Dennett, 2011; Kouider et al., 2010; Phillips, 2011; Stazicker, 2011). There are at least two competing interpretations of the data: the fragmentary phenomenology hypothesis and the generic phenomenology hypothesis. Both hypotheses deny the reality of phenomenal overflow. The fact that participants are able to make correct reports about a greater number of items when properly cued (whether in Sperling's or in Landman's experiment) is accounted for by the idea that information concerning the totality of the items is indeed stored somewhere in the participants' cognitive system, but in an unconscious way - in some sort of detailed unconscious iconic memory. Such unconscious information only becomes conscious *after* the cue, when subjects access it, which means there is no phenomenal overflow.

So, what is conscious (both in the sense of P-consciousness and A-consciousness as, according to these theorists, the two cannot come apart) *before* the cue appears? According to the fragmentary phenomenology hypothesis, participants are then only conscious of mere fragments of items (Kouider et al., 2010) – for example, fragments of letters in Sperling's case. According to the generic phenomenology hypothesis (Cohen & Dennett, 2011; Stazicker, 2011), participants are then merely conscious of *generic* items: for example, in Sperling's case, they are conscious of *letters*, or of letter-shaped forms, but not of *specific and determinate letters* (at least not for *all* the letters). Both hypotheses crucially deny that, prior to the cue, subjects are conscious of all items in a detailed and determinate way.

Block does not seem to deny that these competing hypotheses are compatible with the objective data provided by Sperling's and Landman's experiments (for example, with the performance of participants). He rather argues for the superiority of the overflow thesis by way of an inference to the best explanation (Block, 2007, p. 486): he claims that we should choose the overflow thesis because it is the one which is best supported by our overall data. Indeed, according to Block, everything being equal, we should choose the hypothesis which fits best participants' reports about their own experience. We should not endorse a hypothesis implying that participants' reports about their own experiences are *false*, and that subjects are victims of *illusions*, unless we have some compelling evidence for it. However, according to Block, participants in these studies typically report *seeing all letters or rectangles clearly*. This gives us a reason to reject the alternatives to the overflow thesis.[3]

This methodological *desideratum* is what allows Block to reject competing interpretations of Sperling's and Landman's experiments. Block repeatedly appeals to the idea that, in this kind of experiments, participants claim that, prior to the cue, they can see all or almost all items (Block, 1995, p. 244, 2007, p. 494). He explicitly argues against the fragmentary phenomenology hypothesis and the generic phenomenology hypothesis by pointing out that they contradict what participants say about their own experiences (Block, 2011, p. 570). According to him, these hypotheses imply that participants are victims of "illusions" concerning their own phenomenology. He claims that, according to the generic phenomenology hypothesis, subjects are victims of a "generic illusion": they confuse their real phenomenology (a phenomenology of "letter-likeness") with "specific letter-shape representations that are already specified in consciousness". Block also claims that, according to the fragmentary phenomenology hypothesis, subjects are victims of a "fragmentary illusion", in which they have the impression to see "a grid of specific letters" when what they really see are "sparse fragments" (Block, 2011, p. 568). The idea that endorsing these competing hypotheses implies denying the truth of participants' reports, so that the defenders of these hypotheses are committed to the view that subjects are victims of illusions concerning their own phenomenology, is crucial for Block's rejection of these views, as well as for his argument for the overflow thesis.[4]

[3] "This supports what the subjects say … I am taking what subjects say at face value (though of course I am prepared to reject what subjects say if there is evidence to that effect)" (Block, 2007, p. 488).

[4] Defenders of the overflow thesis have also appealed to considerations independent of subjective reports about phenomenology. Block (2007), for example, mentions data on visual neglect/extinction and suggests that it provides some indirect support for the idea that consciousness might occur without access. Others (Bronfman, Brezis, Jacobson, & Usher, 2014; Usher, Bronfman, Talmor, Jacobson, & Eitam, 2018) tried to build a case for the overflow thesis on the basis of performance on color-diversity tasks (see below for a discussion), or on a study of an irrelevance-blindness paradigm (Usher et al., 2018, Section 5). However, the appeal to subjective reports remains crucial for the defenders of the overflow thesis, notably given that these independent considerations are

This is a postprint. Please refer to the published version, to be found here: https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

There are many ways to oppose Block's argument. One could state that it is acceptable to endorse a view in which participants are victims of illusions about their phenomenology, notably because this illusion (the "rich phenomenology illusion") is predicted and explained by the view (Kouider et al., 2010). Block (2007, pp. 491–494, 2011, pp. 568–574) seems to suppose that the debate lies just there: to what extent should we trust people about their own phenomenology? Does available evidence justify the claim that participants in such studies are victims of *illusions*? Here, we focus on a different question: is it really the case that participants' reports *support* the overflow thesis rather than competing hypotheses? Do participants *really* make reports about their own experiences that contradict these competing hypotheses?

4. ARE THERE ANY DATA ABOUT PARTICIPANTS' SUBJECTIVE REPORTS?

4.1. The quest for origins

Where do the subjective reports emphasized by Block and others come from? Whether they take these subjective reports as empirical evidence in favor of the dissociative view (Block, 2007; 2011), or whether they dismiss them as stemming from some kind of illusion (De Gardelle et al., 2009; Kouider et al., 2010), all researchers mentioned above rely on two kinds of sources: (1) anecdotal reports, including reports about researchers' own personal experience, and (2) scientific publications. Interestingly, scientific publications are scarce: ultimately, when one tracks down the origin of these claims by looking at citations given in the relevant articles, one always ends up stumbling on the same three groups of references: (i) Baars' *A Cognitive Theory of Consciousness* (Baars, 1988), (ii) two different papers by Sperling (1960, 1983), and (iii) a series of articles dating back from the very beginnings of scientific psychology. A 'genealogical' citation tree for the claim that participants report seeing all items can be found in Figure 1.

susceptible to various interpretations (see below our discussion of Bronfman et al., 2014; Ward, Bear, & Scholl, 2016). This justifies our focus on subjective reports.

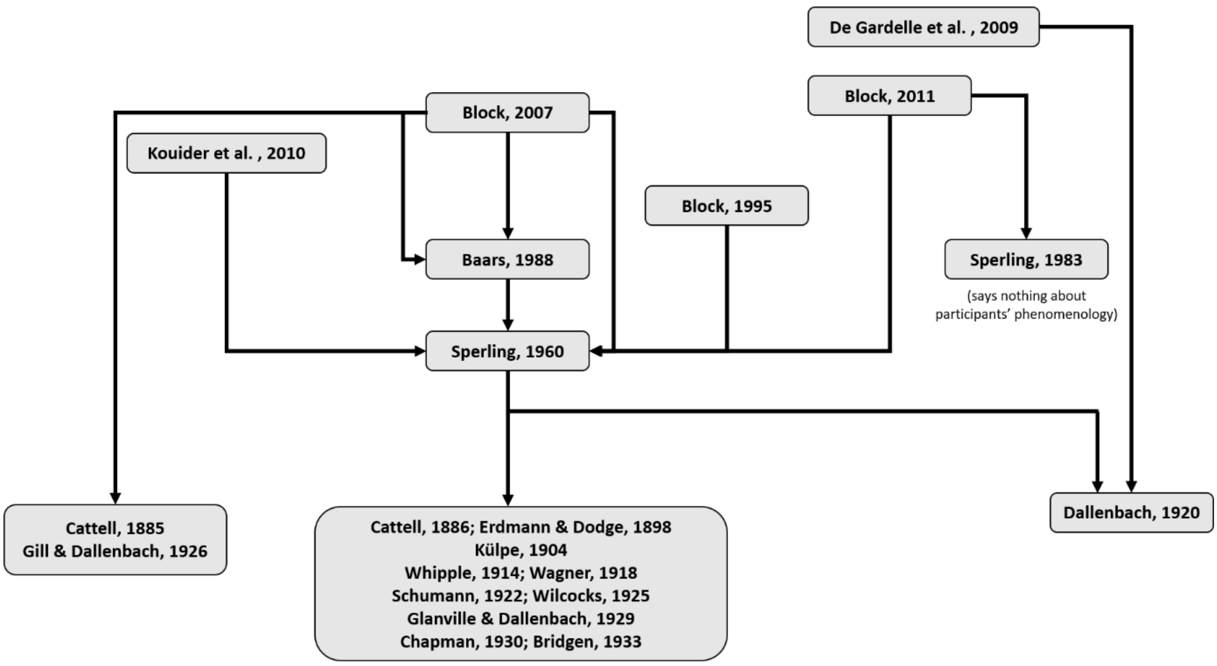

FIGURE 1 Tracking the source of the claim according to which participants report seeing most letters in Sperling-like experiments. Every time a paper does not provide original empirical data in favor of this claim but refers to another paper, an arrow indicates the paper that is referred to.

Baars (1988) is an important node in this citation tree. Looking for an operational definition of consciousness, Baars states in his introduction: "In the course of this book, we will often appeal to the reader's personal experience, but only for the sake of illustration. From a scientific point of view, all evidence can be stated in entirely objective terms". He then provides a criterion indicating that people are conscious of an event: "if (1) they can say immediately afterwards that they were conscious of it and (2) we can independently verify the accuracy of their report" (Baars, 1988, p. 15). Baars then refers to Sperling's experiment, in which "observers typically claim that they can see all the letters, but they can only recall three or four of them," as a case of an experiment that perfectly illustrates the discrepancy between what subjects can say immediately about their feeling and what can be objectively reported or measured by the experimenter: "It is troubling that subjects—and experimenters serving as subjects—continue to insist that they are momentarily conscious of *all* the elements in the array" (ibid., italics in the original). Baars gives no empirical evidence of his own for this claim, which means we are ultimately sent back to Sperling's papers.

Sperling (1983) is just a brief comment that adds nothing new and only repeats in passing that "subjects in tachistoscopic experiments frequently make the curious observation that they saw more than they remembered". This leaves us with Sperling's 1960 seminal paper (Sperling, 1960). Overall, five subjects took part in seven different experiments, and they came to the laboratory three times a week for one month, performing 12 sessions each in total.

So, what does Sperling have to say about participants' subjective experience? Here are the relevant passages:

> "Observers enigmatically insist that they have seen more than they can remember afterwards" (Introduction)
>
> "The answers proposed are a systematic elaboration of an observation that is readily made by most viewers of the actual tachistoscopic presentation. They report that the stimulus field appears to be still readable at the time a tone is heard which follows the termination of the stimulus by 150 msec. In other words, the subjective image or sensation induced by the light flash outlasts the physical stimulus at least until the tone is heard." (Discussion)
>
> "Observers commonly assert that they can *see* more than they can *report"* (Conclusion, italics in the original)[5]

Nothing refers to a proper methodology to collect subjective reports, beyond informal observations about what subjects "*insist"* on or "*assert."* In the context of Sperling's demonstration, subjective reports are not a central piece of the argument but only one element in a body of evidence suggesting the reality of a very short-term memory. That subjective reports appeared consistent with the participants' performance when properly cued was arguably reassuring and probably did not incline Sperling to investigate further. However, if these subjective reports are to be taken as empirical evidence about what participants saw and felt, then the methodology for collecting subjective data is clearly lacking. What did the participants precisely say that they were seeing? How did "most viewers" formulate their observations? This seminal paper did not even try to provide answers to these questions.

This leaves us with a third group of studies: psychological studies from the end of the 19th century to the 1930s—when introspective methods in psychology were in common use. Of course, because these studies predate Sperling or Landman and colleagues' studies, they cannot directly inform us about participants' reports during these experiments. However, one could argue that, to the extent that these studies investigate the phenomenology of participants

[5] This last sentence has been later exposed by Block as an important piece of empirical evidence in favor of the overflow thesis, and is arguably the source of the widespread assumption according to which "subjects see more than they can report (or remember)".

briefly presented with wide arrays of shapes and letters, their conclusion (that participants report seeing most of these shapes in detail) can be justifiably extended to what participants report (or would report) in Sperling's or Landman's experiment. For example, the three studies by Dallenbach and colleagues (see Figure 1) use tachistoscopic presentations to distinguish between attributive and cognitive "clearness" of experience, which somehow overlaps with the phenomenal and access consciousness distinction. Yet there are reasons to be wary of an inference from past studies to current debates:

1. These studies use *very small samples*: Three participants in each study, and six different participants overall (because of multiple participations).
2. Participants are typically *not naïve:* In each study, one of the three participants is the experimenter himself. Others are often psychology professors with prior theoretical commitments.
3. Participants are *highly trained:* In each study, participants typically practice the task for several months. This raises two worries. The first is that their prior knowledge of the task might have an effect on both their performance and phenomenology. The second is that this makes it doubtful that their experience can be taken as representative of the experience of typical participants in Sperling and Landman-style participants, on which most philosophical discussions have focused.
4. Results are *hard to interpret:* Even if studies by Dallenbach and his colleagues are cited in defense of the overflow thesis (because participants typically report seeing all items at the same level of clarity), they could also be cited to support the generic phenomenology view, as participants typically report missing out on certain properties of shapes but not others (Dallenbach, 1920) or being able to grasp certain properties of shapes, such as color, without really perceiving their form (Glanville & Dallenbach, 1929).

Moreover, beyond methodological issues concerning such and such particular studies, there is also the larger issue that this literature is unfamiliar to most present-day researchers. Thus, in absence of a systematic review of this literature, it is hard to determine whether the studies chosen to show that participants have the impression to see all letters in detail are representative of this psychological literature or whether a more systematic and comprehensive review of the literature would reveal serious disagreements between researchers. Thus, we can reasonably conclude from our search that we have no precise idea of what participants in Sperling and Landman and colleagues' studies had the impression to see.

4.2. Addressing the lack of subjective data

Having found little relevant data in scientific publications, we are left with anecdotal remarks, personal comments or general claims which do not seem to be supported by empirical investigation. For instance, Block claims that "subjects (including myself) in overflow experiments often testify that their responses are based on specific phenomenology that was there all along" (2007, responses to objections R2.2). According to Kouider and colleagues (2010), "the overflow argument is rooted in the intuition that we are conscious of much more than we can describe and manipulate. For instance, when observing a complex visual scene, we feel that we have a rich visual experience even if we can report only a few elements". The precise delineation of this "intuition" and of how much it is grounded in the reports of participants in Sperling's experiment is far from clear. Block himself noticed that the debate was lacking solid grounds when he wrote (1995, p. 244): "I am P-conscious of all the letters at once, that is, jointly, and not just as blurry or vague letters, but as specific letters (or at least specific shapes), but I don't have access to all of them jointly at once (I would like to know whether others describe what it is like in this way, but the prejudice against introspection tends to keep answers to such questions from the journals)".

This is precisely the kind of question the current paper aims to answer. Because the importance given to participants' subjective reports in the overflow debate sharply contrasts with our lack of empirical data about these reports, we decided to collect participants' reports in a systematic, non-anecdotal way. More precisely, our goal was to collect subjects' reports about the visual experiences they have when they take part in the kind of experiments typically cited in the literature (Sperling's experiments for Study 1, Landman et al.'s experiments for Study 2). To our knowledge, this is the first project of its kind, as there has been so far no other attempt at systematically collecting detailed participants' reports about their subjective experiences in the Sperling and Landman paradigms.[6]

[6] De Gardelle and colleagues (2009) used what they called a "free subjective report" procedure in which, after being presented with arrays of letters and symbols, participants were presented with 3–8 items and instructed to select those of the items which were displayed in the previous array according to their subjective feeling. However, though useful to determine whether participants had the feeling to see a particular symbol that was present in the original array (which was de Gardelle and colleagues' goal), this procedure does not give us direct insight in the aspects of participants' phenomenology that are relevant to our goal: it does not help distinguish between *specific* and *generic* phenomenology, and it cannot directly inform us on whether participants had the impression to see *all* letters or only *part* of them).

However, recent debates on the nature of consciousness have highlighted the need for a more systematic recording of subjective reports in Sperling-type experiments. For instance, Bronfman, Brezis, Jacobson, and Usher (2014) presented participants with arrays of 24 letters (4 rows) for 300ms in which some rows contained letters of a wide variety of colors (from blue to red), while others contained letters with only slight differences in color (e.g., different shades of blue). While the main task was to report a given letter in a row indicated by a visual cue, participants were also asked as a secondary task to estimate the diversity of color for one row chosen at random (either cued or uncued). Participants correctly reported color diversity for uncued rows, on a par with reports for cued rows, which suggests that participants had access to some information about the uncued rows. Was this information conscious or unconscious?

To answer this question, the authors conducted another series of research in which participants were asked to assess the color diversity of the whole array. Experimenters systematically collected subjective reports through a forced-choice procedure in which participants had to indicate after each trial whether they (i) “did not see the colors”, (ii) “saw the colors partially” or (iii) “saw the colors well”. Researchers observed that participants who gave answers (ii) and (iii) reported color diversity with a greater accuracy than those who gave answer (i), which were at chance level, suggesting, in line with the overflow thesis, that the information participants reported about uncued rows in previous studies was conscious.

Note, however, that these studies only investigate participants’ reports about ensemble properties: the average color diversity of a row or array of letters, rather than individual colors of single letters. Could it be that participants are aware of such ensemble properties without being conscious of individual properties? To explore this possibility, Ward and colleagues (2016) conducted a similar study, but with a more fine-grained measure of participants’ subjective experience. Participants were asked to choose between the four following options: (i) “I had no sense that any of the letters had any color at all”, (ii) “I had a vague sense that the letters were colored in general, but I didn’t clearly perceive the individual colors of individual letters”, (iii) “I had a clear sense that the letters were colored in general, but I didn’t clearly perceive the individual colors of individual letters”, or (iv) “I had a clear sense that the letters were colored in general, and I could also clearly perceive the individual colors of individual letters”. Overall, participants very rarely claimed that they saw individual colors of individual letters and preferred in-between answers, suggesting that participants’ access to individual letters is better described by the generic phenomenology hypothesis.

As an independent confirmation of this conclusion, a trick was then introduced in the experiment: colors of individual letters shifted during the task, but color diversity for the whole rows was kept constant. The shift was unnoticed by participants until debriefing, which suggests they were not aware of the color of particular letters. Ward and colleagues thus concluded that participants experience the statistical diversity of colors without experiencing individual colored letters. One can be conscious of some aspects of a visual scene and extract some general properties of it without being aware of each element of the scene, according to the statistical perception interpretation (Cohen, Dennett, & Kanwisher, 2016).

This debate illustrates three crucial points. First, the systematic collection and study of participants' reports can provide valuable insight in what participants experience during such experiments. Second, participants' subjective reports are in line with more objective measures: they correlate with their performance in report tasks. Third, merely asking participants whether they saw all colors (or letters) is too vague: a positive answer to such a question can be interpreted both as support for the overflow thesis ("I saw all colors in detail") and as support for the generic phenomenology thesis ("I saw all colors, but not in a detailed way"). This is why more fine-grained ways of probing participants' phenomenology are required. Based on these insights, we decided to empirically explore participants' reports about their own phenomenology during the kind of studies put forward by proponents of the overflow thesis, beginning with Sperling's 1960 experiment.

5. STUDY 1

In Study 1, we investigated participants' reports about their phenomenology in a Sperling-like design. Because participants' reports were collected through a comprehensive qualitative interview taking place at the end of the study, we decided to focus on a single experimental condition, so that participants did not have to carry out the difficult task of keeping track of their phenomenology through different conditions. Given recent worries that post-cues might affect and change the phenomenology (see for example Phillips, 2011; Sergent et al., 2013; Gross & Flombaum, 2017), we decided to use the full report condition.

We also chose to recruit naïve participants, with no prior training in introspection and no prior knowledge of the task. There were two reasons for this choice. The first was practical: training a sufficient number of participants would have taken more resources and expertise than we had. The second was theoretical: as most of the current philosophical debate is premised on subjective reports made by psychology students without prior training in introspection, it seemed adequate to recruit similar participants for our study.

Participants were asked to report as many letters as they could of an array of 12 letters that would be briefly displayed on a computer screen (see Figure 2A). The whole experiment was divided into three blocks, and participants had to fill a pen-and-paper questionnaire at the end of each block. At the end of the study, participants engaged in a brief interview with the experimenter. A full description of the methodology is available in Supporting Information (see Appendix A, "Study 1, Methods").

5.1. Results: Main task

Participants' average hit rate (i.e. number of letters correctly reported) was 4.2, which is line with the existing literature: participants tended to correctly report 4 letters out of 12 (see Figure 2B). Participants' false alarm rate (i.e. number of letters incorrectly reported) was 0.3, which suggests that participants did not tend to hallucinate letters or try reporting letters at random.

5.2. Results: Questionnaire

Question 1. We asked participants: "In general, did you feel like you saw all the letters displayed on the screen?". Participants were offered three possible answers ("YES"; "NO"; "I don't know"). They were also asked to indicate how certain they were of their answer on a (0-100) % scale. Participants' answers to Question 1 are summarized in Table 1, which shows that answers depend on the block. Overall, most participants start by denying that they see all the letters, but YES and NO answers are equally distributed by the end of the third block.

TABLE 1 Participants' answers to Question 1 in Study 1 (whether they felt like they saw all the letters displayed on the screen). Right column indicates mean (and standard deviation) for participants' confidence in their answer (on a scale from 0 to 100).

	% of YES/NO answers	**Confidence**
Block 1	YES: 22% NO: 78%	YES: 70.8 (23.2) NO: 74.6 (22.0)
Block 2	YES: 34% NO: 64%	YES: 76.4 (18.4) NO: 76.7 (18.0)
Block 3	YES: 50%	YES: 75.4 (19.5)

	NO: 42%	NO: 81.7 (16.1)

Question 2. We asked participants: "Think about what you experienced during the block that just ended. Which of the following statements best describe what you experienced when the letters were displayed on the screen?". We offered six possible answers. The first and second answers correspond to different versions of the overflow hypothesis (1:"I felt like I saw all the letters in detail and was able to identify them at this time"; 2: "I felt like I saw all the letters in detail, but without necessarily being able to identify them and tell which letters they were at this time"). The third answer corresponds to the generic phenomenology hypothesis (3: "I felt like I saw all the letters, but not in detail: I just saw where they were, and that they were letters"). The fourth and fifth answers correspond to the partial phenomenology hypothesis (4: "I felt like I saw most letters, but not all"; 5: "I felt like I saw only a small part of the letters"), while the sixth answer is a "strange" answer used to test subjects' understanding of the task (6: "I felt like I saw nothing"). Participants were asked to indicate how certain they were of their answer on a (0-100) % scale. Participants' answers to Question 2 are presented in Table 2.

TABLE 2 Participants' answers to Question 2 in Study 1 (best description of what they experienced when the letters were displayed on screen). Numbers in parentheses indicate participants' confidence in their answer (on a scale from 0 to 100).

	Block 1	**Block 2**	**Block 3**
All letters in detail and identified	8% (81.3)	4% (95.0)	6% (88.3)
All letters in detail	8% (82.5)	12% (79.2)	14% (77.1)
All letters, but not in detail	34% (79.1)	36% (79.2)	38% (80.0)
Most letters	36% (82.8)	30% (76.0)	26% (81.2)
Small part of the letters	12% (74.2)	12% (80.8)	14% (82.9)
Nothing	0%	2% (40.0)	2% (10.0)

5.3. Results: Interview

Participants' interviews were independently coded by the first and second author. Inter-rater agreement was assessed using joint probability of agreement.

(A) Free description. (Overall inter-rater agreement: 90.8%) We began our interview by asking participants to simply describe, in their own words, what they had the impression to see when the letters were displayed on the screen. This was an open-ended question, the content of which we then proceeded to code. When participants only stated the obvious, i.e. that they had seen letters, we coded their answers as falling in the *Letters* category (agreement: 88.0%). If participants stated that they generally saw some of the letters, but only part of them, we coded their answers as falling in the *Partial* category (agreement: 92.0%). If they made a distinction between letters they saw clearly and others they saw in a more vague and indistinct way, their answer fell in the *Generic* category (agreement: 82.0%). If they stated that they clearly saw all the letters, their answer fell in the *Overflow* category (agreement: 96.0%). Finally, if their answer was absolutely uninformative (e.g. "I felt stress"), it was coded as *Uncategorized* (agreement: 96.0%). The same answer could fall into multiple categories, either because some categories were not exclusive (such as the *Partial* and *Generic* ones), or because participants reported that their experience evolved from one block to another.

In the end, 62% [52%][7] of answers fell into the *Generic* category, 20% [28%] in the *Letters* category, 16% [20%] in the *Overflow* category, 12% [12%] in the *Partial* category, and 0% [4%] were *Uncategorized*.

(C) Hypothesis matching. (Overall inter-rater agreement: 100.0%) Then, we presented participants with two scientific hypotheses about participants' experience: one that sounded like the overflow hypothesis, and one that sounded like a disjunction of the generic and partial phenomenology hypotheses. Participants were asked to indicate which one best matched their own experience. 22% [22%] chose the overflow hypothesis, 72% [72%] chose the generic/partial phenomenology hypothesis, and 6% [6%] remained undecided.

Taken together, these results suggest that most participants did not think they saw all the letters in detail, either because they thought they saw only some letters in detail and the others in a 'vaguer', more generic way, or because they thought they only saw part of the letters (or both). This clearly goes against the idea that subjective reports support the overflow thesis. However, as we realized this through our interviews, we introduced additional questions to

[7] The first number reports results following the first author's coding, while the number between brackets indicate what the results would be following the second author's coding.

collect more detailed reports on participants' impressions. At the 12th participant, we introduced questions (E) and (I), to which only 39 participants answered.

(E) "Distinct" versus "blurry letters". (Overall inter-rater agreement: 97.4%) We asked participants whether it was correct that they saw certain letters in detail but other in a less precise way. 37 [36] participants out of 39 answered that this was correct.

(F) Letters seen in detail but not reported. (Overall inter-rater agreement: 96.0%) Then, we asked participants whether they sometimes felt like they saw a letter in detail but then failed to report it. 33 [32] participants answered that this was indeed the case, 15 [15] answered that this was not the case, and 1 [3] did not answer. This means that a majority of participants (67.3% [64.0%]) had the impression to see certain letters distinctly, but to be unable to report them.

(I) Number of letters seen in detail. (Overall inter-rater agreement: 97.4%) These results were corroborated by participants' answers to our last question, where we asked them how many letters, on average, they had the impression to have seen in detail.[8] The mean answer was 5.29 [5.31] (*SD* = 1.93 [1.93], Median = 4.5 [4.5]), which is superior to the average hit rate for these 39 participants ($M = 4.22$, $SD = 0.54$): $t(38) = 3.43$, $p = 0.001$ (see Figure 2B for the distribution of both answers).

[8] For participants who answered by a range of numbers (e.g. "4 to 5"), we took the median as the answer (e.g. 4.5).

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

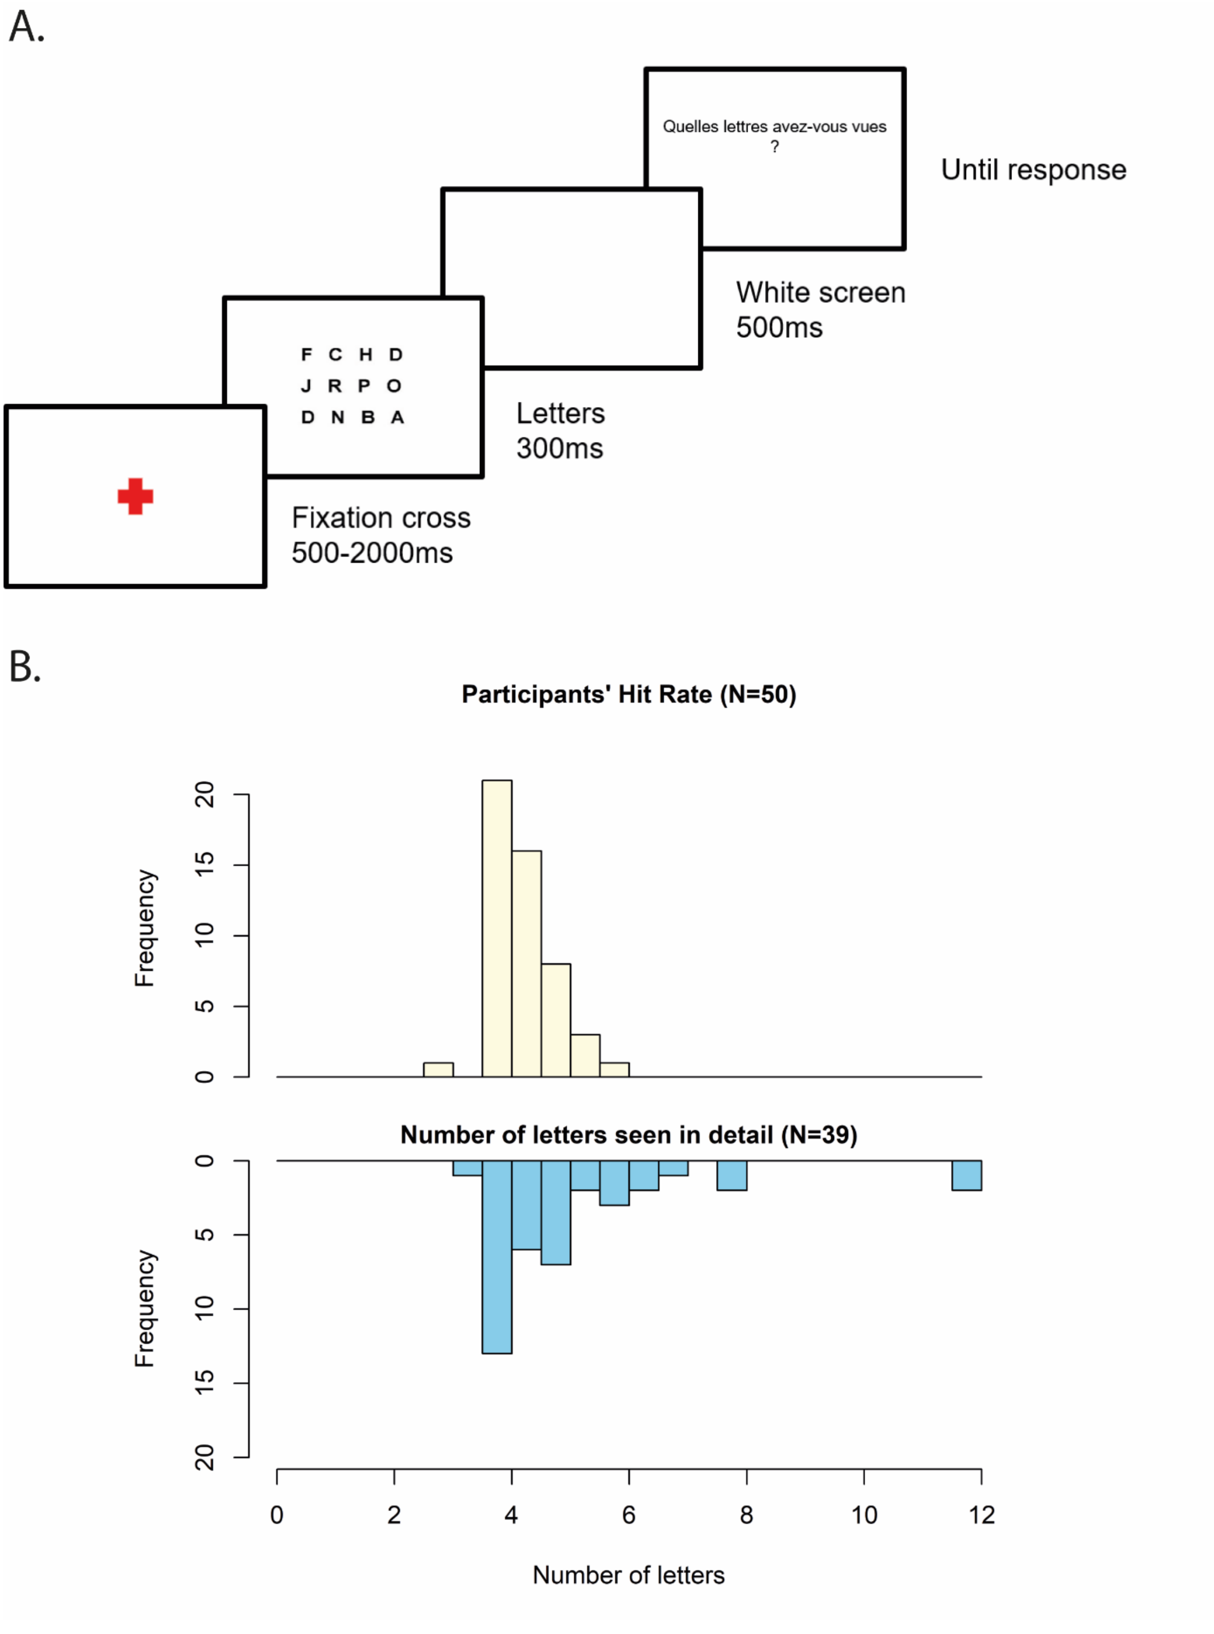

FIGURE 2 (A) A standard trial in Study 1. (B) Distribution of participants' hit rate (for the 50 participants of Study 1), and numbers of letters reported as seen in detail (for the last 39 participants of Study 1).

(G & H) Explanation for letters seen in detail but not reported. (Overall inter-rater agreement: 96.0%) We interrogated participants about the reasons that led them not to report some of the letters they did see in detail. We proposed two non-mutually exclusive explanations, though they were free to propose alternate ones. The first was that, though they saw these letters in detail, they did not have time to categorize them as being such or such letters. The second was that, though they read the letters and stored them into memory, they forgot them before they could report it. Overall, out of 33 [32] participants who claimed not to have reported all the letters they had seen in detail, 9 [9] agreed with the first explanation and 26 [26] agreed with the second explanation (meaning that 4 participants agreed with both explanations). Most participants actually put forward the second explanation even before we proposed it, and reported having memorized some letters but forgetting them as they stored new letters into memory.

5.4. Discussion

Overall, our results failed to support the claim that most participants report seeing all the letters in detail. To our questionnaire's first question (whether participants "felt like they saw all the letters," participants' answers were mixed with very few "YES" answers in the first two blocks (22% and 34%), and only half (50%) of participants selecting this answer in the third block.

Now, does this steady increase of participants' "YES" answers reflect a change in phenomenology or a progressive reinterpretation of an obviously ambiguous question (from "did you see all letters in detail?" to 'did you see all letters, even if not in detail?')? Participants' answers to the second question (choosing the best description of what they experienced when the letters were displayed on the screen) seem to favor the latter explanation, as they display little changes throughout all three blocks. However, we do not have an explanation for why this change in interpretation occurred mostly between the second and third blocks rather than between the first and the second.

Most participants' answers were in line either with the generic phenomenology hypothesis (34 to 38% depending on the block), or the partial phenomenology hypothesis (40to 48% depending on the block). Answers in line with the overflow hypothesis were consistently in the minority (16 to 20% depending on the block).

One might wonder whether these results are not simply due to participants interpreting response options in a peculiar way (for example, interpreting a question about the number of letters they saw in detail as asking them how many letters they were able to keep in memory and report). However, one strength of our study was the debriefing interview, that allowed us to probe thoroughly participants' reports. Participants' recollections during the interview are in line with their answers to the questionnaire. Most participants (more than 90%) reported not seeing all letters in detail and seeing certain letters only in a vague, generic way. Most of them (around 2/3) did report seeing more letters than they could report, but the average difference between the number of letters reported as seen in detail and the number of letters actually correctly reported was low (around 1). Moreover, most of these participants (more than 3/4) explained this gap in terms of having categorized and memorized (and thus accessed) the additional letter but being unable to retain it in memory long enough. Taking their reports at face value, it seems that participants only had the impression to see some of the letters in detail (4 to 6), and the others in a more generic way. Generally, letters seen in detail were those of the line they focused on and a few (1 or 2) letters from another line. Participants tried to memorize letters in a sequential way, but this led them to forget some of the letters they accessed.

Still, one might wonder about an apparent discrepancy between participants' answers to the questionnaire and their answers during the interview: while participants' answers during the questionnaire seemed to favor the partial phenomenology hypothesis, their answers during the interview seemed to favor the generic phenomenology hypothesis. However, this discrepancy is only apparent. Indeed, categorization of participants' answers as "generic", "partial", or "both" in the interview is based on their open, non-guided initial description of their experience. As such, we were only able to categorize as 'partial' participants that took the initiative to report that they had the feeling that they did not see some of the letters. This does not mean that those who did not (and thus were not coded as giving a "partial" answer) had necessarily seen all the letters. Moreover, the fact that we asked participants to describe "what they had the impression to see" might have motivated them to avoid dwelling on what they did *not* have the impression to see.

Thus, our results do not seem to support the claim that most participants report seeing all letters (even less all letters in detail). However, our study suffers from two shortcomings. First, it requires subjects to perform two tasks: reporting letters, and introspecting one's own phenomenology. Even if participants did understand that the two questions (corresponding to the two tasks) were different (since we took care to emphasize this difference), it might still be that the former task influences participants' performance in the latter. We designed a follow-up to Study 1 to address this problem. We instructed participants to "simply look at the screen" while they were shown 10 groups of 12 letters (300 ms each). They were then given the same questionnaire as in Study 1 (Question 1 and Question 2), with an additional Question 3, asking them: "On average, how many letters did you have the impression to see in detail?". The methodology and results of this follow-up are detailed in Supporting Information (see Appendix B, "Follow-up to Study 1").

Second, Study 1 only uses the full report paradigm, while discussions of the overflow argument are often based on the partial report paradigm. We addressed this limitation in Study 2.

6. STUDY 2

As mentioned earlier, Sperling's experiments are not the only ones supposed to show a gap between what people have the impression to see and what people are able to report: Block also appeals to more recent experiments by Landman and colleagues on iconic memory. In this second study, we investigated participants' phenomenology in this category of studies.

Additionally, one might worry that, because Study 1 and its follow-up focused on the full report condition, results might not generalize to the partial report (cued) condition. We took advantage of this study to correct for this shortcoming, by presenting participants with three different conditions. Consequently, we replaced the qualitative interview at the end of the study with a single simple question at the end of each trial, together with shorter pen-and-pencil questionnaires that participants had to fill at the end of each experimental block. This also presented the advantage of allowing us to collect participants' phenomenological reports during the task itself and for each single trial, rather than relying on a post-study interview asking participants for their average phenomenology.

In each trial, the screen would display for a very brief moment a circle composed of 8 or 12 rectangles, each being horizontally or vertically oriented. The circle would then disappear for a moment, then reappear. This second circle of rectangles could either be identical to the first, or different to the extent that the orientation of one rectangle at most changed. Participants were then instructed to report whether the figure had changed or not. Within each trial, a cue (a yellow arrow) would indicate which of the 8 or 12 rectangles was susceptible to change its orientation. However, this cue could appear at three different moments, depending on the trial: during the appearance of the first figure (*before-cue* condition), after the disappearance of the first figure but before the appearance of the second (*between-cue* condition), or during the appearance of the second figure (*after-cue* condition). At the end of each trial, participants were invited to indicate what they had the feeling to see, by clicking on one of the three following answers (see Figure 3A):

A. I felt like I saw all rectangles in detail (orientation included)
B. I felt like I saw all rectangles, but not necessary their details (orientation included)
C. I felt like I did not see some of the rectangles.

The study is divided into six parts. At the end of each of part, participants were instructed to take a break from the computer task to fill a pen-and-paper questionnaire containing four questions. In Question 1, subjects were asked to say when did the cue appear in this block (three possible answers were offered: “Before rectangles disappeared”; “Between their disappearance and reappearance”; “After rectangles reappared”). Question 2 asked: “On average, how many rectangles did you have the impression to see in detail (orientation included)?”. Question 3 asked: “On average, how many rectangles did you have the impression to *just* see (that is: even in a vague way, without necessarily seeing their orientation?)”. Question 4 asked: “Was there times when you saw a rectangle without seeing the direction in which it was oriented?” (possible answers offered: “YES”; “NO”; “I don’t know”). For Question 2, 3 and 4, participants were asked to indicate how certain they were of their answer on a (0-100) % scale. A full presentation of the methods and the questionnaire of Study 2 is available in Supporting Information (see Appendix C, “Study 2, Methods”).

6.1. Results: Main task

For each position of the cue (*before, between, after*) and each number of rectangles (8 or 12), we computed an estimate of participants’ capacity (i.e. number of rectangles subjects have available for comparison) using the same formula as Landman, Spekreijse & Lamme (2003):

*capacity = (hit rate * number of rectangles – number of rectangles * false alarm rate)/(1 – false alarm rate)*

Figure 3B shows our estimation of participants’ capacities for each cue condition and each number of rectangles. An ANOVA with capacity as a dependent variable and number of rectangles (8, 12) and position of the cue (before, between, after) as two within-subjects independent factors revealed a significant effect of the number of rectangles ($F(1,28)=24.4$, $p < .001$), a significant effect of the position of the cue ($F(2,56) = 62.4$, $p < .001$), and a significant interaction effect ($F(2,56) = 13.8$, $p < .001$).[9] Crucially, participants' capacity was higher in the *between-cue* condition compared to the *after-cue* condition, both for eight rectangles arrays (M = 5.76 vs. 3.90; $t(28) = 5.95$, $p < .001$) and 12 rectangles arrays (M = 6.60 versus 4.44; $t(28) = 2.81$, $p = .009$).

Participants' capacity in the *after-cue* condition can be considered an estimate of the number of rectangles' orientations they would normally be able to report (but see Gross & Flombaum, 2017 for a critique of this assumption). Results displayed in Figure 3B indicate that this estimate lies between four and five rectangles.

[9] The interaction is trivially driven by the fact that participants' capacity in the *before-cue* condition is almost perfect, yielding a difference of four rectangles between eight and 12 rectangles arrays in the *before-cue* condition, while this difference is roughly one rectangle in the *between-cue* and *after-cue* conditions.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

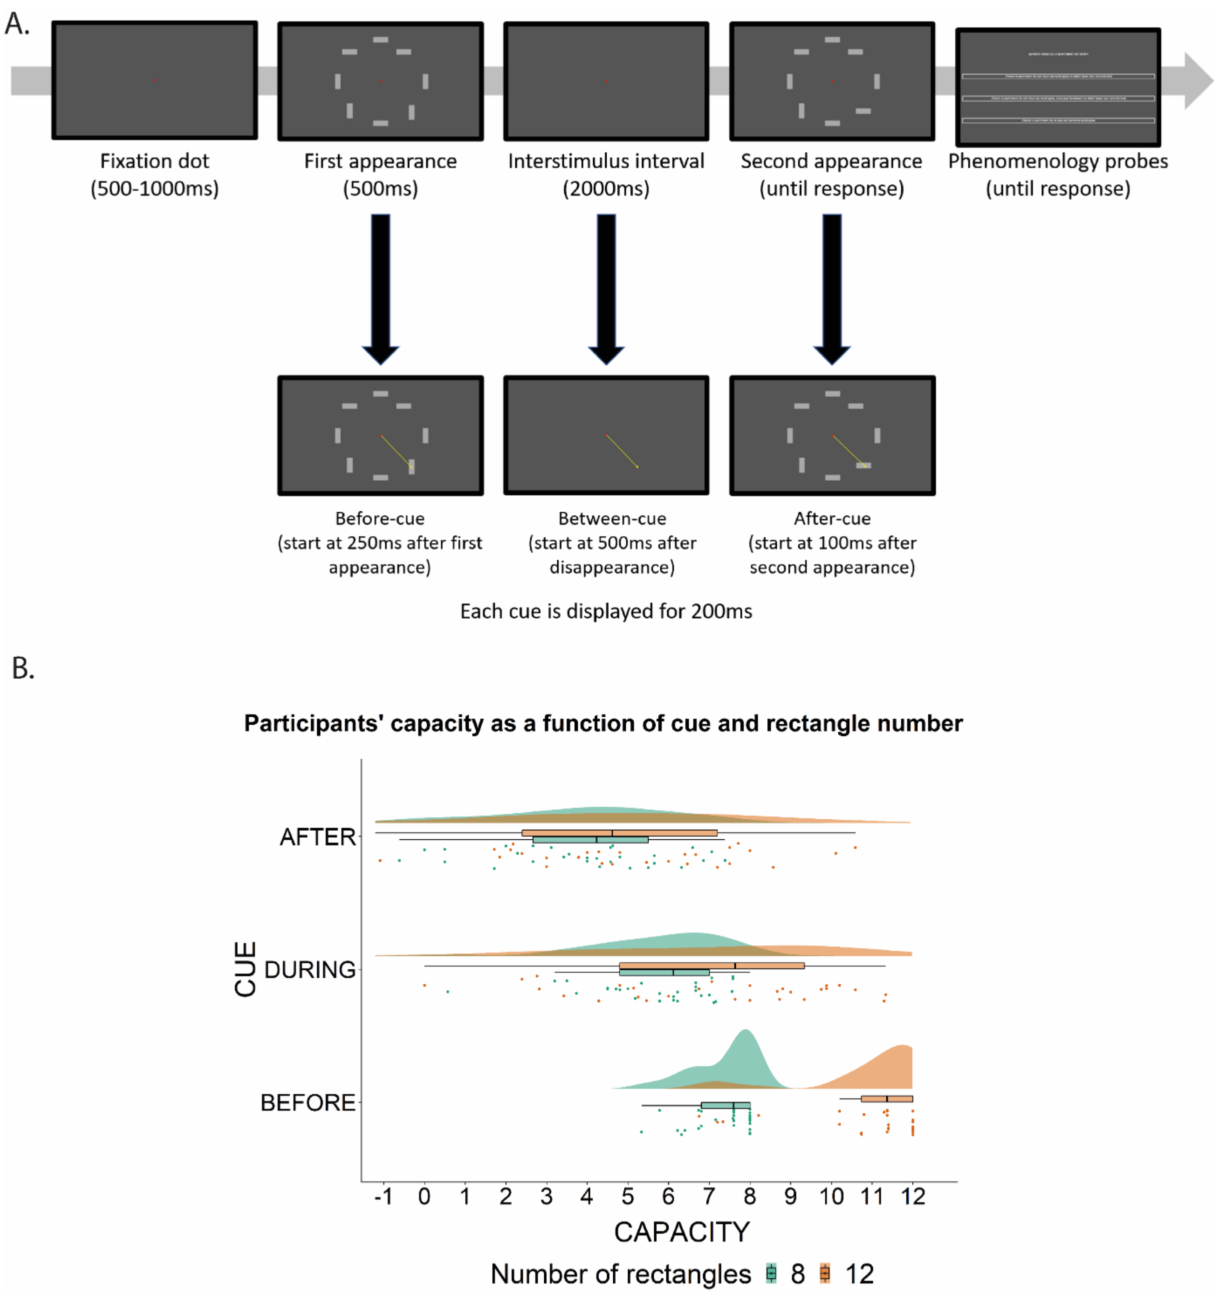

FIGURE 3 (A) A standard trial in Study 2. (B) Distribution of participants' capacity (in numbers of rectangles) for each number of rectangles (8 or 12) and each position of the cue (before, during, after). 3 data points do not appear on the Figure: "-6" in the 12 rectangles, between-cue condition, and "-9" and "-1.5" in the 12 rectangles, after-cue condition. Lines within boxplots indicate the median.

6.2. Results: In-task phenomenology probes

We first analyzed participants' answers to the multiple choices questions asked at the end of every trial during the main task. Participants' answers are summarized in Figure 4.

FIGURE 4 Overall proportion of participants' answers to in-task phenomenology probes as a function of number of rectangles (8 or 12) and position of the cue (before, between, after). Proportions are calculated on a total of 1160 answers per condition.

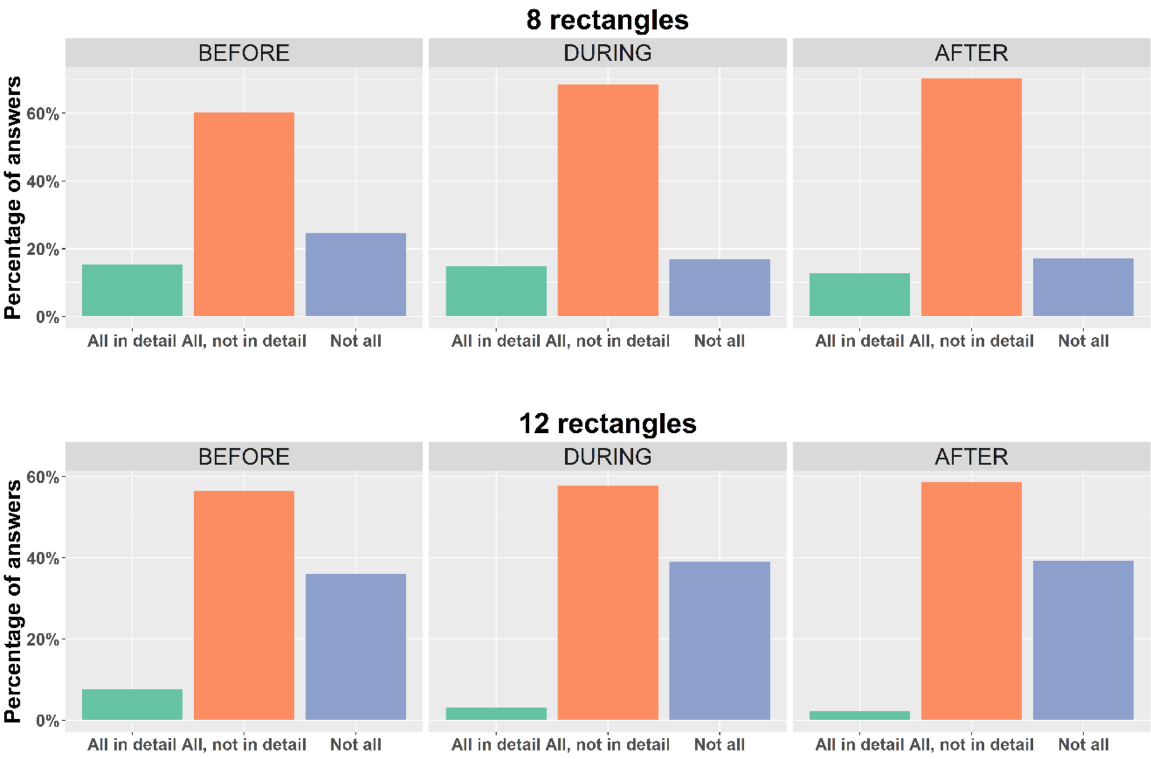

In all conditions, the report according to which participants saw all rectangles in detail (orientation included) was the least chosen (2 to 15% of answers). On the contrary, the report according to which participants saw all rectangles, but not in detail, was always the preferred one (56 to 70% of answers). The third answer, according to which participants did not see all rectangles, always came second. It was more likely to be chosen when there were 12 rectangles.[10]

Thus, it seems that participants' reports about their phenomenology are consistent with a blend of partial (not all rectangles) and generic phenomenology (all rectangles, but not in detail). One interesting question is whether this blend reflects between-participants or within-participants variability. Figure 5 presents the distribution of participants' answers for each of the 29 participants for each number of rectangles and position of the cue, and suggests the

[10] Following one reviewer's suggestion, we converted participants' answers in numeric data (All in detail = 2, All but not in detail = 1, Not all = 0) and conducted a repeated-measure ANOVA with these numeric data as dependent variable and position of the cue and number of rectangles as within-subject factors. We found a significant effect of number of rectangles ($F(1,28) = 16.15$, $p < .001$), but no significant effect of cue ($F(2,56) = 0.07$, $p = .94$), and only a marginally significant interaction effect ($F(1,28) = 2.52$, $p = .09$).

existence of important between-subjects variability: a few participants consistently report seeing all rectangles in detail in certain conditions, while others consistently report seeing all rectangles but not in detail, and others consistently report not seeing all rectangles.

FIGURE 5 Distribution of each individual participants' answers to in-task phenomenology probes as a function of number of rectangles (8 or 12) and position of the cue (before, between, after). Distributions are calculated on a total of 40 answers per participant and per condition. Each single vertical bar represents one single participant's answers to the 40 in-task probes, and how many of these answers fall into each category (black = all rectangles in detail, dark gray = all rectangles but not in detail, pale gray = not all rectangles).

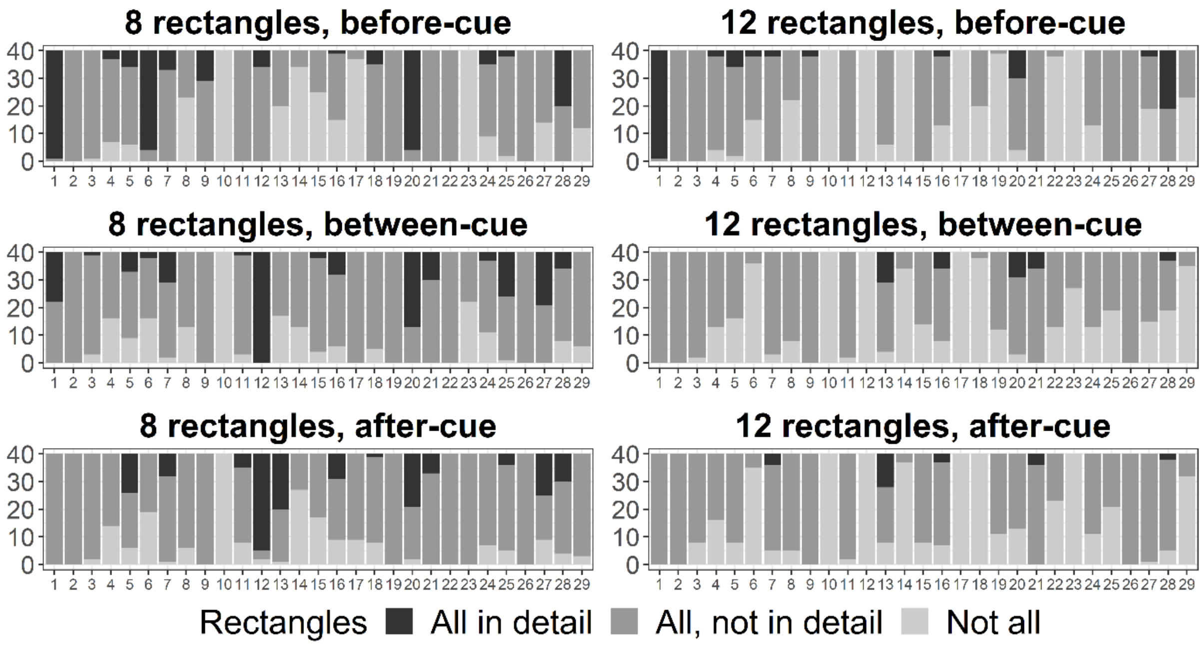

6.3. Results: Numbers of rectangles seen in detail and simpliciter

So, most participants reported not seeing all rectangles in detail, but to what extent was the number of rectangles they reported seeing in detail higher than the number of rectangles they were able to report? We began by analyzing participants' answers to Questions 2 and 3 of our questionnaire. Question 2 asked them how many rectangles they had the impression to see in detail, while Question 3 asked them how many rectangles they had the impression to see *simpliciter* (that is: not necessarily in detail). Participants' answers to these questions for each number of rectangles and position of the cue are summarized in Figure 6.

FIGURE 6. Distribution of participants' answers to Question 2 (number of rectangles seen in detail) and Question 3 (number of rectangles seen overall) for each number of rectangles and position of the cue.

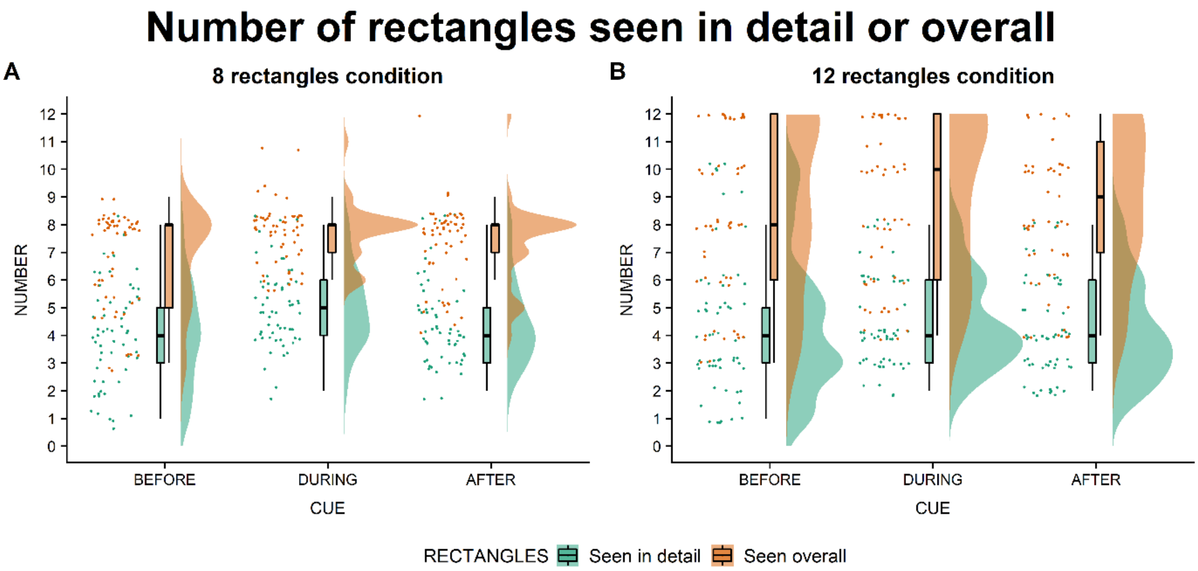

We conducted an ANOVA with number of rectangles as a dependent variable and number of rectangles (8, 12), position of the cue (before, between, after) and type of question (in detail, not in detail) as three within-subjects independent factors. It revealed a significant effect of number of rectangles ($F(1,324) = 4.4$, $p = .04$), a significant effect of question type ($F(1,324) = 59.24$, $p < .001$), and a significant interaction between rectangles and question type ($F(1,324) = 13.8$, $p = .03$). All other effects were non-significant. Overall, participants tended to report seeing more rectangles when there were more rectangles, and when they were asked about rectangles seen simpliciter rather than rectangles seen in detail. This latter difference increases when there are 12 rectangles rather than eight, hence the interaction effect. Interestingly, there was no significant effect of the cue's position on phenomenological reports.

That participants report seeing certain rectangles without seeing them in detail is corroborated by participants' answers to Question 4, in which participants were asked whether there were times when they saw a rectangle without seeing the direction in which it was oriented. Answers to this question are presented in Figure 7 and shows that most participants claim having seen certain rectangles without seeing their orientation.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

FIGURE 7 Percentage of participants' answers to Question 4 (whether they happened to see rectangles without seeing their orientation) as a function of number of rectangles (8 or 12) and position of the cue (before, between, after). Numbers at the top of bars indicate participants' average confidence in their answer (on a scale from 0 to 100).

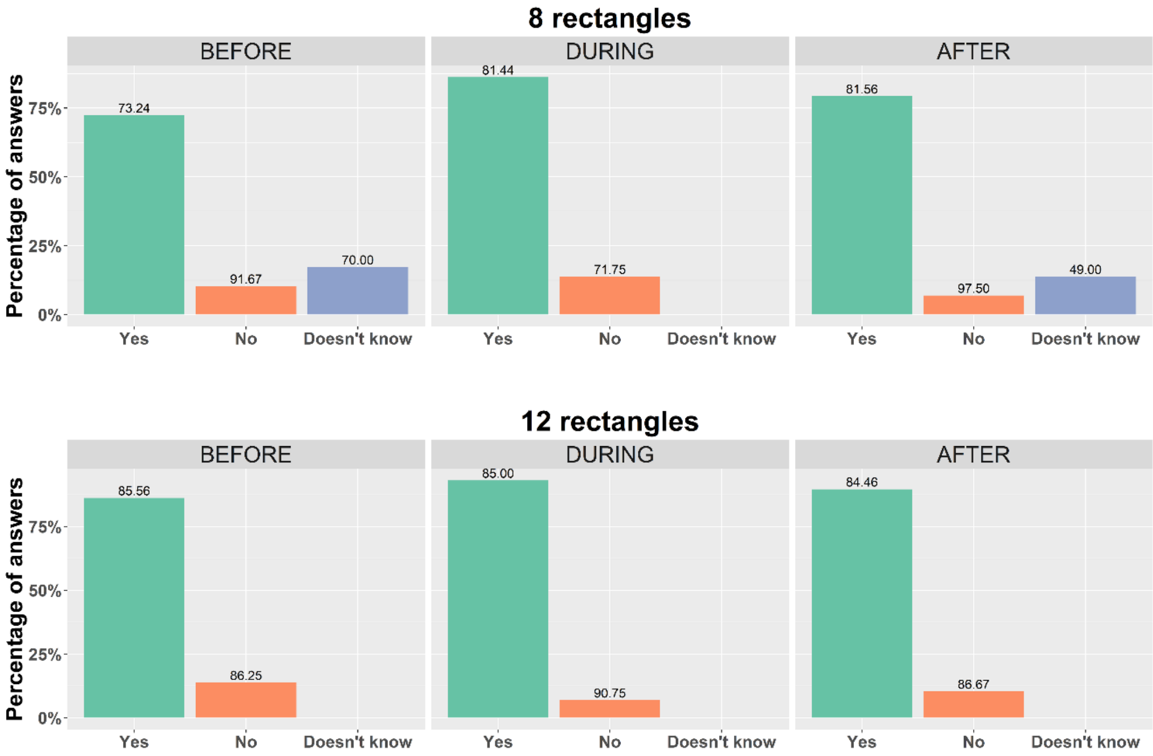

Overall, participants reported being able to see between four and five rectangles in detail.

6.4. Comparison between capacity and number of rectangles seen in detail

Finally, we compared our estimation of participants' capacity (see Figure 3B) to the number of rectangles they reported seeing in detail (see Figure 6). We conducted an ANOVA with number of rectangles/capacity as a dependent variable and number of rectangles (8, 12), position of the cue (before, between, after) and measure (capacity vs. number of rectangles seen in detail) as three within-subjects independent factors. It revealed a significant effect of number of rectangles ($F(1,324) = 4.7$, $p = .03$), a significant effect of position of the cue ($F(2,324) = 8.4$, $p < .001$), and a significant effect of question type ($F(1,324) = 24.9$, $p < .001$). There was also a significant interaction effect between rectangles and question ($F(1,324) = 8.4$, $p = .03$) and between cue and question ($F(2,324) = 6.3$, $p = .002$).

Most of these effects were again driven by a simple fact: participants' reports of rectangles seen in detail tended to be stable across conditions (position of cue and number of rectangles), while estimates of participants' capacity tended to be very high in the *before-cue* condition, and more so when they were 12 rectangles rather than eight. To determine whether the *before-cue* condition was the only condition driving the difference between estimate of participants' capacity and the number of rectangles reported as seen in detail, we compared both measures separately for each condition. When there were eight rectangles, we found a significant difference in the *before-cue* condition ($t(28) = 9.08$, $p < .001$), a significant difference in the *between-cue* condition ($t(28) = 2.44$, $p = .02$), and no significant difference in the *after-cue* condition ($t(28) = -1.38$, $p < .001$). When there were 12 rectangles, we found a significant difference in the *before-cue* condition ($t(28) = 11.21$, $p < .001$), a significant difference in the *between-cue* condition ($t(28) = 3.16$, $p = .004$), and no significant difference in the *after-cue* condition ($t(28) = 0.33$, $p = .75$). For all four significant differences, the number of rectangles seen in detail was lower than estimates of participants' capacity.

6.5. Discussion

The results of our study did not seem to support the claim that most participants report seeing all items rectangles in detail. Answers to in-task phenomenology probes yielded few answers in this direction (ranging from 2% to 15%; see Figure 4). Moreover, the number of participants reporting seeing all rectangles in a consistent way (at least $3/4^{th}$ of the time) ranged from only 0 to 3 depending on the condition (see Figure 5).

Pen-and-pencil questionnaires administered at the end of each block yielded similar results. The average number of rectangles participants reported seeing in detail was quite low, ranging from 3.97 to 5.00 (see Figure 6), and was not higher than participants' estimated capacity (see Figure 3B), suggesting that the number of items reported as seen in detail did not really "overflow" participants' capacity.

Moreover, participants' answers to the questionnaire also provided tentative evidence for the existence of generic and partial phenomenology. On the generic side, the number of rectangles reported as seen *simpliciter* was consistently higher than the number of rectangles reported as seen in detail (see Figure 6), and most participants reported seeing certain rectangles without seeing their orientation (see Figure 7). On the partial side, less than half of participants reported seeing 12 rectangles overall in the 12 rectangles condition (see Figure 6).

Interestingly, the position of the cue did not seem to significantly impact participants' phenomenological reports, while it significantly affected participants' capacity. This suggests that our participants' phenomenological reports are somewhat robust, and can thus be generalized to wider contexts.

7. GENERAL DISCUSSION

In this paper, our goal was to examine the empirical foundations of Ned Block's overflow argument. Indeed, one crucial premise of the overflow argument, as it is currently defended in the literature, is the overflow thesis, which is itself notably supported by the claim that participants in studies such as Sperling's or Landman et al.'s experiments *typically* report seeing much more letters than they can recall. However, having found not much empirical support for this claim in the existing literature, we decided to put it directly to test. What can then be concluded from our results with respect to the overflow thesis? Do most participants *really* report seeing all (or, at least, most) letters without enough details, even though they are not able to report them?

7.1. Implications for the support of the overflow thesis

As it turns out, our results showed that most subjects reported a phenomenology distinct from the one reported by Block. Indeed, here is how Block (1995, p. 244) reported his own phenomenology: "I am P-conscious of all the letters at once, that is, jointly, and not just as blurry or vague letters, but as specific letters (or at least specific shapes)". However, in direct opposition with this description, a majority of our participants reported either (i) seeing only part of the letters, or (ii) seeing all letters, but only some in a detailed way, while the others were just blurry or vague.

In our first study, answers to the questionnaire suggested that participants have a preference for descriptions of their phenomenology that are either partial (seeing only part of the letters) or generic (seeing only certain letters in detail). During the interviews, a majority of participants unequivocally claimed that they had seen some letters more clearly than others, and more than half of them spontaneously described their experience in ways that were very evocative of a generic phenomenology (with only certain letters seen in detail). Overall, this suggests that, when probed about their visual experience, most participants report something akin to a generic phenomenology.

Similarly, in our second study, about two thirds of participants answered that they "saw all rectangles, but not in detail," while one third answered they "did not see some of the rectangles". When asked to estimate the number of rectangles they could see in detail, participants' answers in the post-cue condition matched their actual performance, i.e. the number of rectangles that they could report. This suggests that participants are phenomenologically conscious of the details of a part of the stimulus only, and that this is precisely the part of the stimulus that they can access. Moreover, in support of the generic phenomenology view, a large majority of participants declared seeing some rectangles without seeing the direction in which they were oriented. This idea is consistent with another report from the second study: participants report seeing more rectangles *simpliciter* than they see rectangles in detail. Again, this suggests that the generic phenomenology is a better fit of what most participants report than the overflow hypothesis.[11]

In light of these results, it is thus tempting to reject Block's claim that participants in Sperling's and Landman's experiments typically report seeing all the letters in enough detail to identify them and to conclude that, in its current state, one of the crucial pieces of argumentation in favor of the overflow thesis lacks serious empirical foundations.

7.2. Objections and responses

However, proponents of the overflow thesis might reject our conclusion on several grounds. A first possibility is to argue that some of our results are compatible with the overflow thesis or might even be used to support it.

---

[11] Note that we deliberately used the same set of stimuli as the studies we replicated. However, recent studies on overflow (e.g., Bronfman, Ward) use generally more complex stimuli, introducing colors or larger arrays of letters. The nature of the stimulus should be taken into consideration when trying to generalize our results or compare them to other studies. The more complex the stimuli, the more likely perception will not be able to capture the details of the scene. The most complex visual stimuli are natural scenes, e.g., in the context of the laboratory, drawings, photographs as stimuli. While only a very short flash is needed to categorize a scene (Thorpe, Fize, & Marlot, 1996) and a short presentation time (e.g., 200ms) allow participants to capture the "schema" of a scene (Biederman, 1981), the exact nature of the "schema" or "gist" of the scene is a debated issue. Furthermore, if most experiments in gist perception have been conducted with forced choices questions, Fei-Fei and colleagues (2007) have shown that under a free report paradigm participants are able to describe many details of a scene that has been presented only under 500ms. These results, considered as a testimony to the richness of visual experience, have been sometimes interpreted as supporting the overflow thesis (Haun et al., 2017). In our studies, stimuli are much simpler than a natural scene and in a sense, the gist of the scene does not change: participants still face an array of letters or rectangles at each trial.

One might try such a move on the ground that, in Study 1, we actually observe that participants report seeing clearly and in detail more letters than they can actually report. Actually, the difference is roughly one letter (M = 5.29 vs. 4.22), which is far less impressive than the common claim that participants have the impression of seeing *much more* than they can report, but proponents of the overflow thesis do not necessarily have to show that participants see all the items in a given scene: it is enough to prove their point that participants report having seen more letters than they could actually access.

So, do the results of Study 1 support the overflow hypothesis? Not necessarily. Although most participants (around 65%) declared seeing in detail more letters than they could actually report, this does not mean that they declared seeing in detail more letters than they have *accessed*. Indeed, participants might sometimes have *accessed* certain letters while being unable to report them because they were not able to keep them in memory. When asked about the reasons why they could not report some of the letters they saw in detail, a wide majority of these participants (around 80%) pointed to the limit of their working memory: they saw the letters clearly and memorized them, but failed to retain them in memory long enough to report them. More interestingly, during the interview, many of these participants informally reported having accessed letters by pronouncing them mentally, that is, reading letters and putting them "in memory". In addition, very few of these participants (around 28% - so, 18% of the total number of participants) reported that the letters they saw in detail but failed to report were letters they did not have the time to identify and categorize at the time they appeared. Taken together, these data suggest that, for most of our participants, the letters they reported seeing in detail but did not report were at some point accessed (since categorizing letters presupposes accessing them in some fashion). Thus, though our results support the claim that, sometimes, most participants report seeing a bit more than they can report, our results cannot be reasonably used to argue that participants report seeing a bit more than they can actually access. At best, our results are compatible with the claim that in Study 1 (but not in Study 2) a small subset of participants seems to report seeing in detail on average one more letter than they can actually access. We doubt that this kind of claim is likely to impress proponents of the generic or partial phenomenology hypothesis.

A second way to reject our conclusion is to question our choice of participants. Indeed, one could object that our participants had no training in introspection, and that as a consequence they were not able to deliver reliable and clear introspective reports.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

One way to answer this objection would be to side with opponents of introspection, who claim that trained introspection is not a reliable process because introspection is permeable to theoretical commitments. The history of the use of introspection in psychology has shown that the content of observations made by highly trained subjects-experimenters is correlated with the theoretical commitments of these experimenters (Lyons 1986, p. 21). Thus, one might argue that the fact that participants in our study were not professional practitioners of introspection and had no prior expectations regarding the task they were performing, are the very reasons why we should consider their reports "trustworthy".

However, for the purpose of the present research, we need not enter in such debates about the reliability of introspection, nor reject the idea that reports of those trained in introspection are more "trustworthy". In fact, what we need to do is to distinguish two senses in which subjective reports can be trustworthy. In a strong sense, subjective reports are trustworthy when they accurately reflect participants' own phenomenology. In a weaker sense, subjective reports are trustworthy when they accurately reflect what participants *think* their own phenomenology to be (i.e. when participants are not lying or hiding something about their own phenomenology). Though considering our participants' reports "trustworthy" in the strong sense would be a bold supposition, it seems perfectly reasonable to consider them "trustworthy" in the second, weaker sense. This is all we need for our purpose.

Indeed, let us remind the reader that our goal in this paper is to test the main empirical premise for Ned Block's "overflow argument". And this premise, the overflow thesis, is precisely supposed to be supported by the alleged fact that participants *report* seeing in detail most, if not all of the items in Sperling's and Landman's experiment. This last fact is not a fact about the accuracy of participants' subjective reports, but about their content – about what people *think* they see. As such, testing it does not require one to assume that participants' reports are accurate, only that they are honest. To put it otherwise: because Ned Block's argument is that we should endorse the overflow hypothesis as the *simplest* and thus *best* explanation of the data, which corresponds to what participants think of their own phenomenology, it is sufficient, in order to undermine his argument, to show that participants *do not* have such thoughts about their own phenomenology, given that they do not actually make the corresponding reports (we presuppose here that participants are sincere – their reports express what they think). No assumption about the accuracy of these reports is required to reject Block's argument.

At this point, proponents of the overflow hypothesis might feel some relief: to the extent that one cannot assume our participants' reports to be introspectively accurate ("trustworthy" in the *strong sense*), our results do not refute the overflow hypothesis. Our results only allow us to reject one particular argument for this thesis, Block's argument, which crucially requires this thesis to be supported by participants' reports. This might lead some readers to think that we did not accomplish much. After all, all the overflow proponent need to do is to replace Block's argument by another, alternate argument grounded not in any participants' subjective reports, but, say, in *expert* participants' reports (that is: reports from participants with training in introspection). Such an alternate argument would claim that we should endorse the overflow hypothesis because it fits the opinion of expert participants about their phenomenology.

However, Block's argument cannot simply be replaced by such an argument. Indeed, to our knowledge, no study has ever systematically collected expert participants' reports during an experimental paradigm similar to the ones used by Sperling and Landman and colleagues. Barring such data, this kind of argument is baseless. So, until such data are collected (and show that expert participants' phenomenological reports actually matches the overflow hypothesis), there is no replacement for Block's argument, and our results actually achieve something substantive by depriving the overflow thesis of one of its main bases.

Moreover, we also think that our results dramatically shift the dialectic of the debate, by displacing the burden of proof. Indeed, let us remind the methodological principle put forward by Ned Block himself: we should choose the best and simplest explanation for participants' results and reports in experiments such as Sperling and Landman's. However, because all major accounts can provide an explanation for these results and reports, the best explanation is the one that assumes that participants are not victims of an illusion and are not deeply mistaken about their own phenomenology. This logic led Block to conclude that the overflow hypothesis is the best hypothesis. However, the same logic, applied to our results, should lead us to conclude that the generic and partial phenomenology hypotheses are in a much better position here than the overflow hypothesis. Indeed, arguing against our results that we should not take naïve participants' reports seriously and that we should rather trust reports from trained participants amounts precisely to claiming that most participants are *mistaken* about their phenomenology – thus undermining Block's original argument and turning it upon its head. If defenders of the overflow hypothesis want to dismiss naïve participants' reports as inaccurate to salvage their hypothesis, they must do exactly what Block required from the generic and partial phenomenology theorists: they owe us an explanation. They must claim that naïve subjects are victims of an illusion regarding their own phenomenology, and they must explain this illusion. The burden of proof has thus shifted.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

Thus, for our results to refute Block's argument in favor of the overflow thesis, we do not need to suppose that our participants' reports are accurate (or more accurate than the one of trained participants). Now, it is also true that, from a more general perspective, it would be interesting (if not crucial) to also systematically gather reports from trained participants. However, this goes beyond the reach of the present research.

8.3. Explaining the disagreement

Finally, one interesting observation that we can draw from our results is that there seems to be no obvious and unanimous answer among our participants regarding the content of their phenomenal experience. Some participants consistently report seeing all items presented during the task, though not in detail; others consistently report seeing only part of them, while a small minority report seeing all items in detail (see Figure 5). In addition, some participants consistently shift from one description to another depending on the trial, adding intra-observer variability to the inter-observer variations. Though such results might not be particularly surprising, the possibility of wide variations in participants' subjective reports has never really been discussed with respect to the overflow debate, as almost all the contenders took the fact that subjects report a detailed phenomenology for granted.

Our studies do not allow us to determine exactly what is the source of these variations. Of course, the most fascinating hypothesis is that differences in reports reflect *actual differences* in phenomenology, but there are more trivial, less exciting possibilities. For example, different participants might simply interpret our response options in different ways (again, let us remind that "seeing" lends itself to many different interpretations). Or, less trivially, different participants might rely on different introspective processes. Still, the possibility that participants might have different phenomenologies is interesting, as it could provide an explanation for the current philosophical disagreement on what participants in Sperling-like studies really see. In other words, the existence of diverging theoretical options in the overflow debate might correspond to the fact that different persons have different experiences. Block's claim that he could see all the letters in Sperling's experiment is, after all, a subjective report as well. Although a minority, some of our participants reported experiences precisely in line with the overflow thesis. Block and other scholars sympathetic to the overflow thesis might belong to a subgroup of persons with a rich phenomenology, while other persons would have a generic, or partial, phenomenology.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

The idea that different participants might have substantively distinct phenomenologies deserves further investigation. The study of consciousness rests on the idea that human beings share a common frame of experience. Of course, all researchers agree that consciousness can have various shapes inside of this common frame, but what if there is an irreducible variety of forms of consciousness, even among human beings?

8.4 Conclusion

By running partial replications of Sperling's 1960 experiments and Landman and colleagues' 2003 experiments, and collecting participants' subjective reports in a systematic way, we have shown that, contrary to what has been claimed (notably by Ned Block), these reports do not support the overflow thesis. Our results rather suggest that participants report a mix of partial and generic phenomenology, thus undercutting the main premise for Block's argument in favor of the overflow thesis.

More generally, our results can be seen as a reminder that philosophers and scientists should avoid relying on "subjective data," when these data come merely from informal reports or personal experience. A systematic collection of detailed and precise reports is the best way to ensure the reliability of subjective reports. In other words, a systematic collection of reports must be recognized as legitimate by anyone who thinks that subjective reports must be taken at face value, at least *prima facie* (e.g., Block). Even researchers doubting the direct use of subjective reports in the study of consciousness may want to study them as indicating something about introspection.

SUPPLEMENTARY MATERIALS

Materials and data for all studies are accessible at https://osf.io/5gn23/ This includes program used to display stimuli, that exclusively run on open sources software (Python, Pygame), allowing the reader to run his own version of the studies to check the generalizability of our results. If help is needed in setting up the experiment, do not hesitate to contact the corresponding author.

ACKNOWLEDGMENTS

This project originates in the CNRS "Methods and Foundations of science" Summer School 2014. The authors would like to thank the audience at the Consciousness Club at The University of Tokyo, and the Thumos Research Group at University of Geneva. Figures 2 and 3 were produced using the R scripts and examples provided by Allen et al. (2019).

REFERENCES

Allen, M., Poggiali, D., Whitaker, K., Marshall, T. R., & Kievit, R. A. (2019). Raincloud plots: A multi-platform tool for robust data visualization. *Wellcome Open Research*, *4*, 63. doi:10.12688/wellcomeopenres.15191.1

Baars, B. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Biederman, I. (1981). On the Semantics of a Glance at a Scene. In Kubovy, M., & Pomerantz, J. (Eds.), *Perceptual Organization* (pp. 213–53). London: Routledge.

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

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*, 481-548. doi:10.1017/S0140525X07002786

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

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*, 1394-1403. doi:10.1177/0956797614532656

Brigden, R. L. (1933). A tachistoscopic study of the differentiation of perception. *Psychological Monographs*, *44*(1), 153-166.

Brown, R. (2012). The myth of phenomenological overflow. *Consciousness and Cognition*, *21*(2), 599-604. doi:10.1016/j.concog.2011.06.005

Cattell, J. M. (1885) The inertia of the eye and brain. *Brain*, *8*(3), 295–312.

Cattell, J. M. (1886). Über die Trägheit der Netzhaut und des Sehcentrums. *Philosophische Studien*, (*3*), 94–127.

Chapman, D. W. (1930). *The Comparative Effects of Determinate and Indeterminate Aufgaben Upon Certain Types of Attentional Report*. Doctoral dissertation, Harvard University.

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, *15*(8), 358-364. doi:10.1016/j.tics.2011.06.008

Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the Bandwidth of Perceptual Experience? *Trends in Cognitive Sciences*, *20*(5), 324-335. doi:10.1016/j.tics.2016.03.006

Dallenbach, K. M. (1920). Attributive vs. cognitive clearness. *Journal of Experimental Psychology, 3*(3), 183-230.

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

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, *10*(5), 204-211. doi:10.1016/j.tics.2006.03.007

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*, 1-37. doi:10.1016/s0010-0277(00)00123-2

Erdmann, B., & Dodge, R. (1898). *Psychologische Untersuchungen über das Lesen auf experimenteller Grundlage*. Halle: Niemeyer.

Fei-Fei, L., Lyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision, 7*(10), 1-29. doi:10.1167/7.1.10

Gill, N. F., & Dallenbach, K. M. (1926) A preliminary study of the range of attention. *The American Journal of Psychology, 37*(2), 247–56

Glanville, A. D., & Dallenbach, K. M. (1929). The range of attention. *The American Journal of Psychology*, *41*(2), 207-236.

Gross, S., & Flombaum, J. (2017). Does perceptual consciousness overflow cognitive access? The challenge from probabilistic, hierarchical processes. *Mind & Language*, *32*(3), 358-391. doi:10.1111/mila.12144

Haun, A. M., Tononi, G., Koch, C., & Tsuchiya, N. (2017). Are we underestimating the richness of visual experience? *Neuroscience of Consciousness*, *3*(1). doi:10.1093/nc/niw023

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*(7), 301-307. doi:10.1016/j.tics.2010.04.006

Külpe O. (1904). Versuche über Abstraktion [Experiments on abstraction]. *Bericht über den Kongress für experimentelle Psychologie*, *1*, 56–68.

Lamme, V. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, *7*, 12-18. doi:10.1016/s1364-6613(02)00013-x

Lamme, V. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, *10*(11), 494-501. doi:10.1016/j.tics.2006.09.001

Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*(2), 149-164. doi:10.1016/s0042-6989(02)00402-9

Lyons, W. (1986). *The Disappearance of Introspection*. Cambridge: Bradford Book.

Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, *46*(3), 774-785. doi:10.1016/j.neuropsychologia.2007.10.005

Phillips, I. (2011). Perception and Iconic Memory: What Sperling Doesn't Show. *Mind and Language*, *26*(4), 381-411. doi:10.1111/j.1468-0017.2011.01422.x

Schumann, F. (1922). Das Erkennungsurteil. *Zeitschrift für Psychologie*, *88*, 205-224.

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. doi:10.1016/j.cub.2012.11.047

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1-29.

Sperling, G. (1983). Why we need iconic memory. *Behavioral and Brain Sciences*, *6*(1), 37-39. doi:10.1017/S0140525X00014564

Stazicker, J. (2011). Attention, Visual Consciousness and Indeterminacy. *Mind and Language*, *26*(2), 156-184. doi:10.1111/j.1468-0017.2011.01414.x

Thorpe, S., D. Fize, & Marlot, C.. (1996). Speed of Processing in the Human Visual System. *Nature*, *381*(6582), 520–22. doi:10.1038/381520a0.

Usher, M., Bronfman, Z. Z., Talmor, S., Jacobson, H., & Eitam, B. (2018). Consciousness without report: insights from summary statistics and inattention "blindness". *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *373*(1755). doi:10.1098/rstb.2017.0354

Wagner, J. (1918). Experimentelle Beiträge zur Psychologie des Lesens. *Zeitschrift für Psychologie*, *80*, 1-75.

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. doi:10.1016/j.cognition.2016.01.010

Wilcocks, R. W. (1925). An examination of Külpe's experiments on abstraction. *The American Journal of Psychology*, *36*(3) 324-341.

Whipple, G. M. (1914). *Manual of Mental and Physical Tests, Part I, Simpler Processes*. Baltimore: Warwick and York.

This is a postprint. Please refer to the published version, to be found here:
https://onlinelibrary.wiley.com/doi/abs/10.1111/mila.12291

Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Sciences*, *7*(5), 214-218. doi:10.1016/S1364-6613(03)00081-0