Cognition 238 (2023) 105529

Contents lists available at ScienceDirect

Cognition

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

Experiencing without knowing? Empirical evidence for phenomenal consciousness without access

Yoni Zion Amir[a], Yaniv Assaf[a,b], Yossi Yovel[a,c], Liad Mudrik[a,d,*]

[a] Sagol School for Neuroscience, Tel-Aviv University, Tel Aviv 69978, Israel
[b] School of neurobiology, biochemistry and biophysics, Faculty of Life Sciences, Tel Aviv University, Tel Aviv 69978, Israel
[c] School of Zoology, Faculty of Life Sciences, Tel Aviv University, Tel Aviv, 69978, Israel
[d] School of Psychological Sciences, Tel-Aviv University, Tel Aviv 69978, Israel

ARTICLE INFO

Keywords:
Consciousness
Phenomenal consciousness
Access consciousness
Qualia
Unconscious processing
Cognition

ABSTRACT

Can one have a phenomenal experience to which one does not have access? That is, can you experience something without knowing? The dissociation between phenomenal (P) and access (A) consciousness is widely debated. A major challenge to the supporters of this dissociation is the apparent inability to experimentally demonstrate that P-without-A consciousness exists; once participants report having a P-experience, they already have access to it. Thus, all previous empirical support for this dissociation is indirect. Here, using a novel paradigm, we create a situation where participants (Experiment 1, $N = 40$) lack online access to the stimulus yet are nevertheless able to retrospectively form judgements on its phenomenal, qualitative aspects. We further show that their performance cannot be fully explained by unconscious processing or by a response to stimulus offset (Experiment 2, $N = 40$). This suggests that P and A consciousness are not only conceptually distinct, but might also be teased apart empirically.

Statement of relevance: A critical question in the scientific quest towards solving the problem of consciousness focuses on the ability to isolate conscious experiences at their purity, without any accompanying cognitive processes. This challenge has been augmented by a highly influential – yet controversial – dissociation suggested by the philosopher Ned Block between ***Phenomenal consciousness***, or the "what it is like" to have an experience, and ***Access consciousness***, indexing the ability to report that one has that experience. Critically, these two types of consciousness most typically go together, making it highly difficult – if not impossible – to isolate Phenomenal consciousness. Our work shows that the dissociation between phenomenal and access consciousness is not merely conceptual, but can also be empirically demonstrated. It further opens the gate to future studies pinpointing the neural correlates of the two types of consciousness.

1. Introduction

Consciousness is one of the most perplexing phenomena in nature: although there is no other phenomenon with which we are so intimately familiar, scholars have spent centuries trying to understand exactly what it is and how it operates, and have yet to arrive to an agreed upon explanation (Yaron, Melloni, Pitts, & Mudrik, 2021). Among different accounts, a highly influential yet controversial suggestion by Ned Block dissociated between two types of consciousness: phenomenal consciousness (*P-consciousness*) and access consciousness (*A-consciousness*) (Block, 1995). The former refers to the purely subjective aspects of experience (i.e., "what it is like" Nagel, 1974), also known as qualia (Tye, 1997). A-consciousness, on the other hand, refers to the cognitive access one has to the experience, translated into the ability to report it or perform some mental operations on it. Notably, although these two types of consciousness typically go together, Block argued that they do not always co-occur: we can have P-consciousness without A-consciousness (*P-without-A*), suggesting that we do not have access to all the contents we experience (Block, 1995, 2005, 2007).

This dissociation evoked philosophical discussions (Block, 2007; Cohen & Dennett, 2011; Naccache, 2018) and has inspired neuroscientific and cognitive work: The recurrent processing theory (Lamme & Roelfsema, 2000) adopted it, suggesting that P-consciousness is subserved by recurrent processing within sensory areas, while attention and

* Corresponding author at: Tel Aviv University, PO Box 39040, Ramat Aviv, Tel Aviv 69978, Israel.
E-mail address: mudrikli@tauex.tau.ac.il (L. Mudrik).

https://doi.org/10.1016/j.cognition.2023.105529
Received 22 March 2023; Accepted 16 June 2023
Available online 30 June 2023
0010-0277/

fronto-parietal activations are needed for A-consciousness (Lamme, 2004). Other studies of processing without attention found different neural activations for Kanizsa shapes and non-Kanizsa shapes during attentional blink (Fahrenfort, Van Leeuwen, Olivers, & Hogendoorn, 2017) and inattentional blindness (Vandenbroucke, Fahrenfort, Sligte, & Lamme, 2014). The authors then claimed that these represent forms of phenomenal experiences that are not accessed (though notably, these could also simply be unconscious processes). Finally, some studies focused on the claim that perception overflows access, or report, to substantiate the argument that A-consciousness is different from P-consciousness: the former has a limited capacity while the latter is rich and includes information we cannot report (Block, 2011). An ongoing debate surrounds this point (Block, 2011, 2014; Bronfman, Brezis, Jacobson, & Usher, 2014; Cohen, Dennett, & Kanwisher, 2016; De Gardelle, Sackur, & Kouider, 2009; Kouider, De Gardelle, Sackur, & Dupoux, 2010; Sperling, 1960).

Yet, beyond the theoretical argument (Block, 1995) and the above-mentioned studies whose results have been interpreted in light of this dissociation, thus far there has been *no direct empirical evidence for P-without-A consciousness*. Such evidence would have to meet three conditions: (a) participants should have no access to the stimulus during its presentation, (b) they should nevertheless demonstrate that they had some qualitative, phenomenal experience of the stimulus, and (c) it should be established that (b) does not simply reflect unconscious processing (Cohen & Dennett, 2011; Kouider, Sackur, & de Gardelle, 2012; Phillips, 2011). Notably, satisfying (a) and (b) is not an easy feat, as it requires the isolation of P-consciousness from its cognitive representation (Block, 2007; Kouider et al., 2010).

Here, we strive at meeting all three conditions using a novel paradigm that mimics an everyday phenomenon that is often given as an example for P-without-A (Block, 1995): imagine you are working from home, completely absorbed in your work; you probably would not notice the constant noise emitted by the refrigerator compressor. If asked about it, you would most likely say you do not hear a thing. Yet when the refrigerator turns off, you notice that something has changed in your experience, and you are even able to tell, in retrospect, what that noise sounded like, despite denying hearing it in real time. Such a retrospective experience is familiar to us, yet it has never been quantified or demonstrated in a multi-trial experiment, when participants know that they will be asked about their experience (so their lack of real-time access cannot be explained by not paying attention to the probed stimulus, not knowing how it should sound like, nor not understanding the question).

In a preregistered experiment (https://osf.io/yhav8), participants heard a mixture of sounds sometimes including one of six pink noise signals, and were asked to report if they can hear certain sounds. The sounds were gradually turned off, until in some trials, only the pink noise remained. At that point, participants were asked whether they heard something. Then, if the noise was still on, it was abruptly stopped and participants were asked if they have heard the change. This yielded three types of trials: "Access trials" (A-trials), where participants had cognitive access to the pink noise while it was playing and noticed that it had stopped, "Phenomenal trials" (P-trials), where participants denied online access to the stimulus, but did notice it after it stopped, and "No-consciousness trials" (NC-trials), where neither took place. Lastly, participants were asked to perform a 2-Alternatives-Forced-Choice (2AFC) discrimination task where two (different) pink-noise signals were presented. To demonstrate P-without-A, we asked if participants can retroactively discriminate the qualitative aspects of the stimulus they denied hearing in P-trials, satisfying the first two conditions described above. The comparison with NC-trials tested the contribution of unconscious processing to performance, and an additional control experiment further excluded the option that the results could be explained merely by a retrospective response to the offset of the noise.

2. Methods

Participants: Overall, 143 participants took part in Experiments 1 and 2. Sample size was determined based on a pilot experiment we conducted ($N = 14$). There, participants answered the forced choice task correctly in 66 out of 130 P-trials in a 3-AFC discrimination task ($\chi2(1) = 4.90$, $p = 0.027$). To calculate power, we ran a simulation analysis with 1000 iterations, so that in each iteration we randomly picked five P-trials from each participant in the pilot experiment (N trials = 70) and ran a one-sided binomial test, with chance performance defined as 33.33% (since in the pilot experiment, there were three optional sounds rather than two, as was the case for the current experiments). The resulting values were entered into the binom.power function included in the R binom package, which provided the required sample size for obtaining a significant effect based on these results with 95% power. This generated a distribution of required sample sizes over the 1000 iterations. The median of that distribution was chosen to serve as the current sample size, which was 40 participants. Notably, this refers to participants who had at least four P-trials; thus, given the great variability in participants' tendency to have P-trials, 103 participants (57 women, mean age = 26.8, SD = 4.3) were run in Experiment 1 until this sample size was completed. Forty participants were included in Experiment 2 (32 women, mean age = 22.65, SD = 1.66). Participants were awarded either course credit or a financial compensation (~$15) for their participation. All participants had normal or corrected-to-normal vision and hearing. People with past neurological, or mental disorders, or taking psychiatric medicines, were not included. The experiments were approved by the Tel Aviv University ethics committee. All participants signed a consent form and were explained that they could leave the experiment at any stage if they wish to do so.

2.1. Exclusion criteria

None of the participants failed to meet the first exclusion criteria, as all participants had >50% success in the discrimination test during the practice task. As described above, 62 participants were not included in the final analysis of Experiment 1 for not having more than three P-trials. Additionally, one participant who did have more than three P-trials was also excluded since she had no A-trials at all, suggesting she probably did not understand the instructions, and was clearly an outlier with respect to the other participants.

2.2. Apparatus

Participants sat in a dimly lit and sound attenuated room. To minimize all noise, only the 24″ VG248 screen (100-Hz refresh rate, located 60 cm from the participant) and the Creative A5 speakers (and not the computer) were positioned in the experimental room. All stimuli were presented using Matlab and Psychtoolbox 3 (Brainard, 1997).

2.3. Stimuli

The sound stimuli included six pink noises (low-pass filtered signals) with different upper cutoff frequencies (2 kHz–7 kHz, in jumps of 1 kHz), with peak intensity of 32 dB SPL at the ear. All pink noises were generated using the Avisoft Saslab Lite software. Pink noises were found to be distinguishable in a pretest ($N = 8$, M = 77.1%, SD = 17.1). Twelve semantically meaningful sounds were downloaded from internet sources: an ambulance siren, babies' laughter, bells rings, cat meows, cricket chirps, crowd clapping, dog barks, drums' noise, frog croaks, sheep bleats, a phone ring and snores. These sounds were further edited in Audacity to generate two additional versions (±18% pitch), to be used in the practice session. Two pure tones (600 Hz, 1000 Hz,) and a compound one (3000 Hz–15,000 Hz) were also downloaded from internet resources. The intensities of the semantic sounds and the tones varied (M = 44 dB SPL at the ear, min = 32 dB SPL).

2.4. Procedure

Trial procedure is presented in Fig. 1. All experiments started with a discrimination practice task ($N_{trials} = 20$), where a target sound (semantic/pink noise) was presented for 1 s. Two sounds were sequentially presented in a random order for 1 s each: the target sound and one of its versions (higher/lower pitch, or a different pink noise). Participants were asked to identify the target sound (Fig. 1A). Then, the main task started. In Experiment 1, each participant went through 48 trials. Each trial began with the presentation of a randomly selected mixture of sounds, played continuously: two semantic sounds and one tone (Fig. 1B). In 2/3 of the trials, these sounds were accompanied with an ongoing pink noise (pink-noise presentation was pseudorandomly intermixed, with the constraint that the first trial included a pink noise and the second trial did not). Then, eleven questions were sequentially presented at random order (RGB: 127, 127, 127, 5.24° X 7.62°). Nine questions prompted participants to report if they were hearing a certain sound (e.g., 'do you hear cats?'). In another question, participants were asked "Do you hear *something*?", so a positive response was expected in case they heard any sound. Finally, another question was a confidence rating one with respect to the previous answer. The sounds were randomly turned off during the first eleven steps, with the tone being last to remain, besides the pink noise (if presented). When the tone was also turned off (and the pink noise, if presented, was still on), participants were asked again to report if they hear something and rate their confidence. Then, the screen went blank, the pink noise was turned off (again, if presented), and a new question appeared ("Did you notice any change?"). Two seconds after participants' response, a discrimination task was presented, akin to the practice task (difference between the two alternatives was 2 kHz). Participants were asked to choose which sound was presented during the trial, or guess if they did not hear any noise. Finally, at the end of each trial, a memory question was presented to assess engagement level l ("Were you asked about hearing dogs in the last trial?").

In Experiment 2, the procedure was similar to the practice trials in Experiment 1. In each trial ($N = 108$), one of the three tones was presented for 1 s with a gradual onset of 200 ms and gradual offset of 10 ms, in which one of the pink noises was gradually introduced had a 10 ms, remaining for additional 30/40/50 ms. Participants were first asked to rate the clarity of their experience using the 4-points Perceptual Awareness Scale (1-did not hear at all, 2-vague perception; 3-parly heard; 4-heard clearly (Ramsøy & Overgaard, 2004)), and then performed a 2AFC discrimination task (identical to Experiment 1, with the probe noises presented for 2 s each; Fig. 1C).

2.5. Statistical analysis

Logistic mixed models (Agresti, 2003) were used to determine if the accuracy in the discrimination tasks was above chance, with participants defined as random effect. A positive intercept, estimated using Wald's test, indicated above-chance performance. To assess differences between trial types, the same models were run, but with trial type as a fixed effect

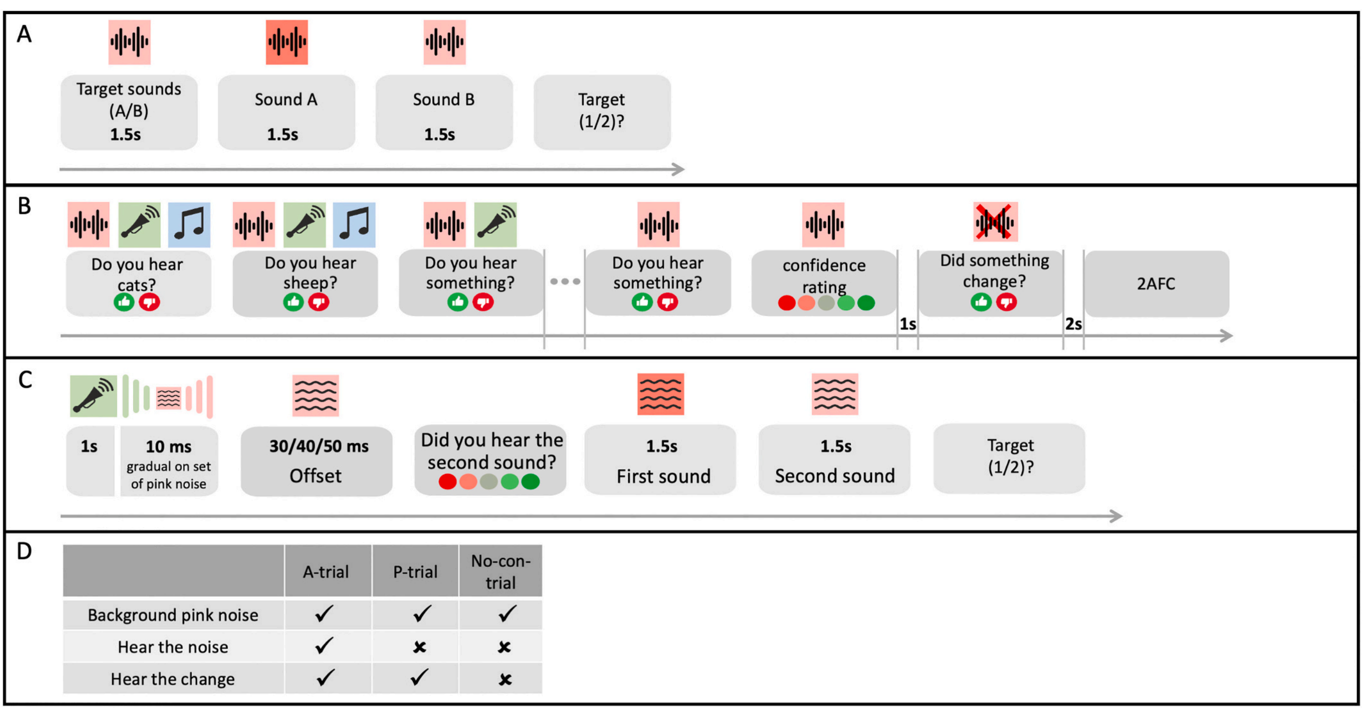

Fig. 1. Experiment 1 procedures and trial types: A. An example trial in the discrimination practice phase, in which a target sound was played ("original sound"), followed by two sounds at different pitches, one of which was the target sound (order counterbalanced). Then, the participant was asked to determine if the target sound appeared first or second in the sequence. **B.** An example trial in Experiment 1, where in the first ten steps random questions about hearing specific sounds (Do you hear cats?) or hearing something (Do you hear something?) were presented. During these ten steps, the sounds were gradually and randomly turned off as the participant advanced in the trial, until in 2/3 of the trials only the pink noise remained (in the other 1/3, no sound was presented at this point). Then, the participant was asked "do you hear something?", and asked to rate her confidence in the response. At that point, the pink noise was also turned off, and the participant was asked if she heard any change. This was followed by a discrimination task (similar to panel A, but without the target sound being presented again). **C.** An example trial in Experiment 2. Each trial started with a tone playing for 1 s with 200 ms gradual onset. During the last 10 ms, the tone was gradually turned off, while a pink noise started to play, again gradually introduced (within those 10 ms). Then, the pink noise remained for a short duration (30/40/50 ms) until its abrupt offset. Immediately afterwards, participants were asked to rate to what degree they have heard the second sound (e.g., the pink noise). The trial ended with a discrimination task identical to Experiment 1. **D.** Summary of trial types. All trials have background pink noise. In A-trials (Access trials), participants heard the pink noise while it was playing and heard the change when it stopped. In P-trials (phenomenal trials) and NC-trials (no consciousness trials) they did not hear the ongoing pink noise. In P-trials they heard the change after the sound stopped, while in NC-trials they did not. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and participants as a random effect. First we checked if the latter model is preferred over a depleted model containing only the participants as a random effect, indicating a main effect of trial type. Then we preformed post-hoc tests with one of the trial types defined as the intercept, to look for specific differences between trial types.

All *P*-values were corrected according to the Tree-FDR method (Bogomolov, Peterson, Benjamini, & Sabatti, 2021) (see supplementary information).This method corrects for all comparisons that would have made this work publishable, while taking into account families of comparisons and nested analyses.

3. Results

3.1. Experiment 1

3.1.1. Lack of cognitive access to the stimulus in real time

We first asked if P-trials can even be found; and the answer was positive (Fig. 2A; for the effects over time, see Fig. 2B), though with high between-individuals variability: across 103 participants, in 12.3% (SD = 12.6%) of the trials, participants denied hearing the stimulus when presented, but noticed the change when it stopped. 67.7% of the trials (SD = 32.2%) were A-trials, and NC-trials constituted 20.0% (SD = 29.3%). In line with the predefined exclusion criteria, only 40 participants who had at least four P-trials were included for further analysis (Fig. 2C). For those participants, 24.0% of the trials were P-trials (SD = 12.8%), 57.2% of the trials were A trials (SD = 25.7%), and 18.8% of the trials were NC trials (SD = 19.0%). Participants were also fairly confident in their responses about hearing or not-hearing the noise, with M = 3.88 (SD = 0.74), 4.20 (0.63) and 4.19 (0.61) for P, A and NC trials, respectively (on a 1–5 scale, where 5 denotes "highly confident").

To better characterize the behavior of these participants, we examined their performance in all other questions that appeared in the main task (e.g., "Do you hear dogs?"; "During this trial, were you asked if you heard cats?"; See Methods). Reassuringly, performance was very high both for the within-trial questions (M = 95.90%, SD = 2.73%) and, to a lesser degree, for the post-trial memory question (M = 75.10%, SD = 8.43%).

3.1.2. Phenomenal experience of the stimulus

We then tested if participants also had retrospective access to the *qualitative nature* of the stimulus. Here and elsewhere, we used logistic mixed models (for details, see Statistical Analysis) to test (a) if performance is above chance; or (b) if trial types differ. The estimated intercept in P-trials was above zero (M = 0.88, 95% CI = [0.63,1.19], z(284) = 6.20, $p < 0.001$, OR = 2.42), suggesting that participants were indeed above chance in identifying the pink noise that they heard during the trial (Fig. 3A). As a sanity check, we also tested if A-trials are above chance, which was indeed the case (M = 1.50, 95% CI = [1.24,1.79], z (693) = 11.10, p < 0.001, OR = 4.47. Fig. 3A).

3.1.3. Performance cannot be explained by unconscious processing

A critical question in this study is if the above-chance performance in the discrimination task should be ascribed to P-consciousness (i.e., having retrospective access also to the qualitative nature of the sound that was not detected while playing), or simply to unconscious processing (Cohen & Dennett, 2011; Kouider et al., 2012). According to the latter account, the unconscious processing of the pink noise might have pushed participants towards the correct answer without them having any conscious experience of the noise. This interpretation can be examined using the NC-trials: if above chance performance will be found on those trials, it would support the interpretation that unconscious processing might suffice for successfully discriminating between the

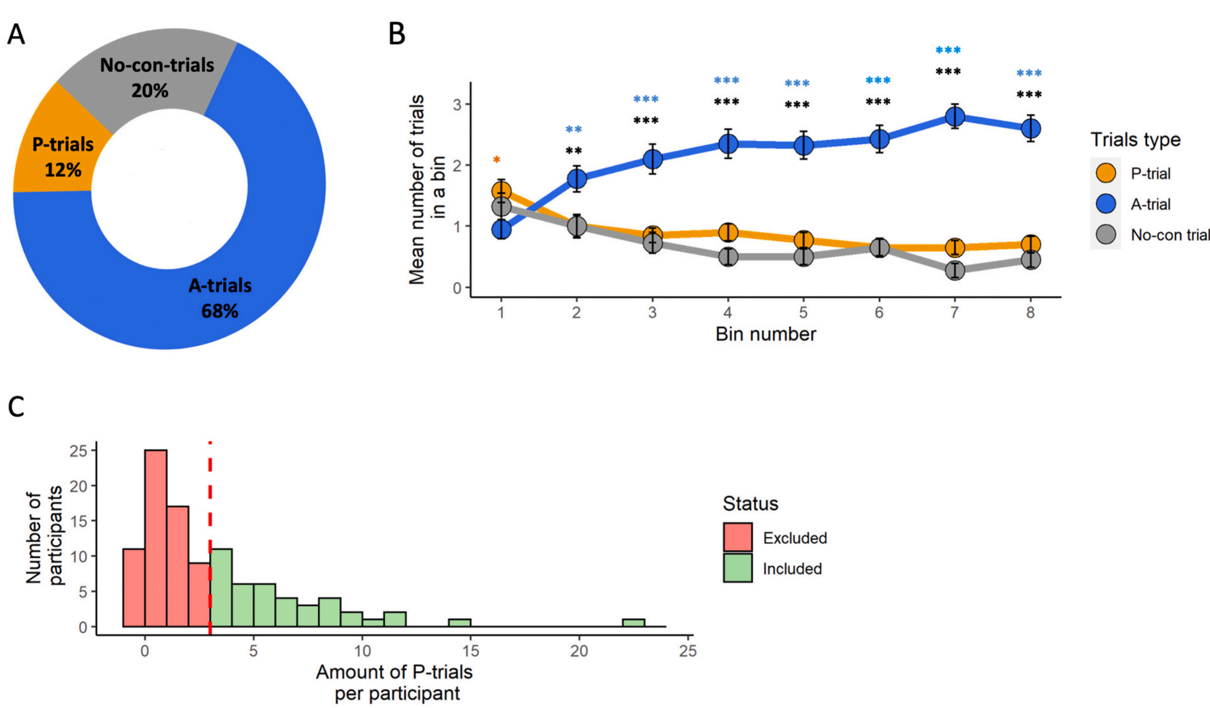

Fig. 2. Trial frequencies in Experiment 1. A. Trials distribution across participants. **B.** Number of P, A and NC-trials throughout the experiment. Trials with background pink-noise were divided into eight bins of four trials. The dots represent the average number of P, A and NC-trials, and the black lines represent the standard error in each bin tested with binomial test. In the first bin there were more P-trials than A-trials, and in all other bins the number of A-trials was higher than both the number of P-trials and the number of NC-trials. One asterisk (*) represents significance below 0.05, two asterisks (**) represent significance below 0.01, and three asterisks (***) represent significance below 0.001. Orange asterisks represent more P-trials than A-trials, while blue and black asterisks represent more A-trials than P-trials and NC-trials, respectively. **C.** Distribution of the number of P-trials per participant. The Y-axis represents the number of participants who have a certain number of P-trials. Participants who have less than four P-trials (left) were excluded from further statistical analysis. The dotted red line represents the minimum limit of four relevant trials. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

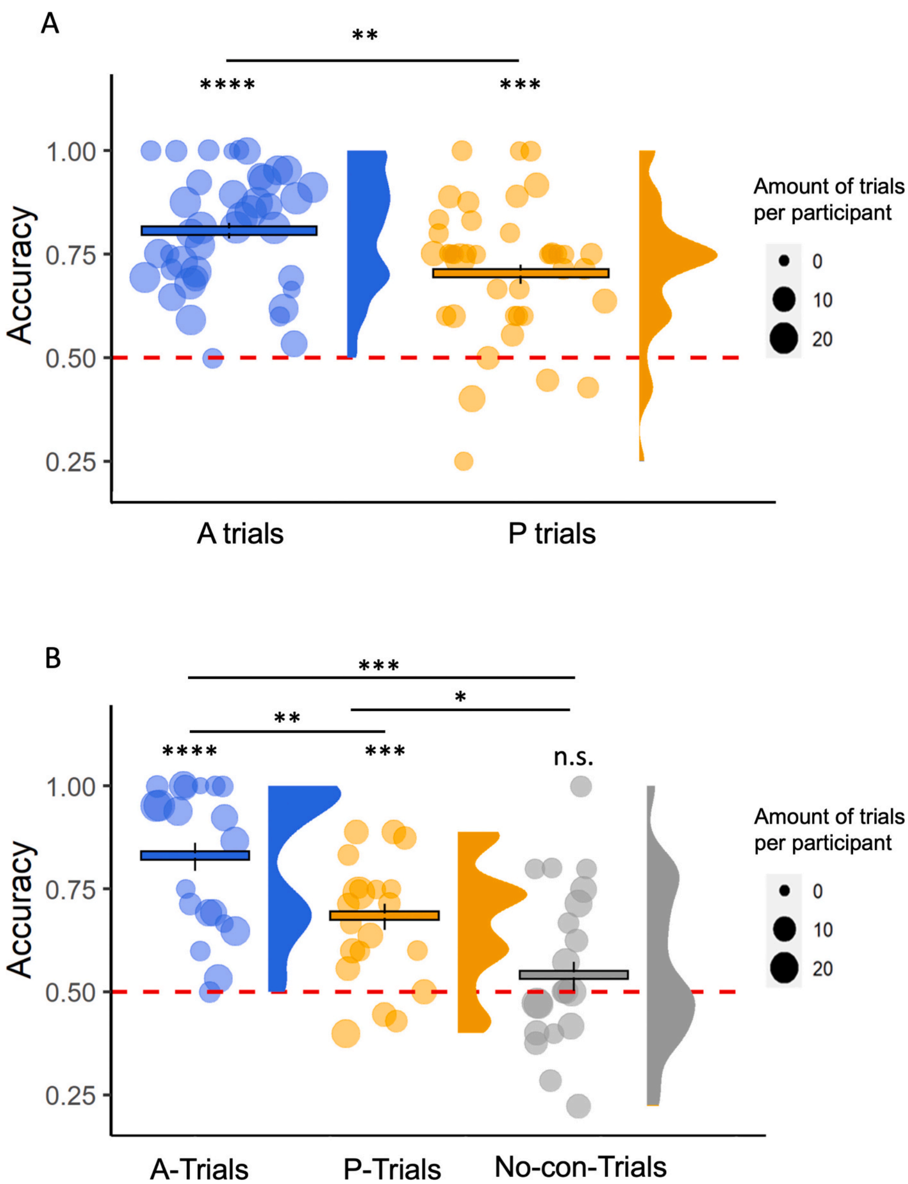

Fig. 3. Performance in discrimination tests for A-trials, P-trials and NC-trials. Each dot represents performance of an individual participant and its diameter represents the number of trials this participant had. The horizontal lines depict the weighted-average accuracy across all participants, with vertical lines marking the standard error. The distribution of these data points is presented to the right of each plot. **A.** Performance in P (left) and A (right) trials ($N = 40$). In both trial types, discrimination was higher than chance. Yet performance was overall higher in A-trials than in P-trials. **B.** Performances in discrimination test in all trial types, only for participants who had more than three NC-trials ($N = 20$). Performance in A-trials is higher than performance in P-trials ($p < 0.001$) and in NC-trials ($p < 0.001$). Performance in P-trials is also higher than the performance in NC-trials ($p = 0.048$). Finally, in this model too both A and P trials performance is above chance ($p < 0.001$ for both), but not in NC-trials. In all panels, one asterisk (*) represents significance below 0.05, two asterisks (**) represent significance below 0.01, three asterisks (***) represent significance below 0.001, and four asterisks (****) represents infinitesimal *p* value.

sounds, at least to some degree. Notably, the pink noise in these trials is exactly the same as the one presented in the P trials, so there is no reason to expect a difference in the level of unconscious processing in both cases. Yet performance in NC-trials (for participants who had at least four such trials; $N = 20$), suggested otherwise, and was not greater than chance (M = 0.16, 95% CI = [−0.12,0.49], z(195) = 1.14, $p = 0.254$, OR = 1.18).

We then compared performance between the three trial types (A, P and NC). An effect of trial type was found ($\chi^2(2) = 38$, $p < 0.001$; See Methods for details about the model), with better discrimination performance for A-trials than for both P-trials (β = −0.98, CI = [−1.49,-0.49], z(576) = −3.85, $p < 0.00$, OR = 0.38) and NC-trials (β = −1.42, 95% CI = [−1.91,-0.96], z(576) = −5.91, $p < 0.001$, OR = 0.24), and – importantly - better performance for P-trials compared with NC-trials (β = −0.44, 95% CI = [−0.88,-0.01], z(576) = −1.98, $p = 0.048$, OR = 0.64; see Fig. 3B). The same differences between A and P trials were found in an additional analysis, where all 40 participants were included (see supplementary information).

3.2. Experiment 2

3.2.1. The offset response

Experiment 1 suggested that P-trials performance cannot be solely explained by unconscious processing. However, it might still be claimed that performance is driven by the offset response, where hearing the offset alone allows the reconstruction of the sound due to cognitive inferring (Kouider et al., 2010), which can then be used to perform the task. This explanation was examined in Experiment 2; following piloting (see supplementary information), in Experiment 2 the pink-noise stimulus was presented for different short durations (30/40/50 ms) when fully attended, without additional sounds besides a 1 s (3000 kHz) pure tone that preceded it. The tone was designed to mask the onset of pink noise, which started while the tone was still being played, akin to Experiment 1. Critically however, here the pink noise was not presented throughout the entire stimulation (but only at the end), so P-without-A consciousness could not be obtained, and its duration was shortened so to minimize as much as possible the chances of hearing it in real time, as it was played, while still hearing its offset (see Methods). Participants were asked to rate the clarity of the experience of the pink noise using a variant of the Perceptual Awareness Scale (Ramsøy & Overgaard, 2004) (with 1 indexing “I did not hear anything”, 2 representing “I barely heard something”, 3 meaning “I partly heard the sound” and 4 denoting “I heard the entire sound clearly”) (see Methods) and to perform a 2AFC discrimination task (identical to Experiment 1, see supplementary materials). Only trials receiving perceptual rating >1 were included (N = 1124 across participants). That is, in the included trials, participants declared hearing the sound, despite the short duration, indicating that they have indeed heard either its offset, or the stimulus itself while it was playing.

Performance in the discrimination task in these trials was assessed for each stimulus duration using a logistic mixed model. Above chance performance was not found for the 30 ms and 40 ms stimuli (M = 0.13, 95% CI = [−0.10, 0.35], z(347) = 1.18, p = 0.238, OR = 1.14; M = 0.13, 95% CI = [−0.12, 0.38], z(389) = 1.04, p = 0.297, OR = 1.16, respectively), yet it was found for the 50 ms ones (M = 0.33, 95% CI = [0.08, 0.57], z(385) = 2.73, p < 0.001, OR = 1.39). Importantly, even in these longer duration trials, performance was not as good as in the P-trials (See analysis below). Thus, even when the offset of the sound was clearly audiable (as indicated by participants’ ratings), the information conveyed in it was not sufficient for reconstructing the qualitative aspect of the sound in retrospect to obtain a similar performance as the one found in P-trials. This conclusion is supported by a direct comparison of performance in the P-trials in Experiment 1, showing that performance in P-trials was higher than in 30, 40, and, importantly – also in the 50 ms trials (β = −0.75, 95% CI = [−1.11,-0.41], z(1410) = −4.22, p < 0.001, OR = 0.47; β = −0.76, 95% CI = [−1.11,-0.42], z(1410) = −4.36, p < 0.001, OR = 0.47; β = −0.54, 95% CI = [−0.89,-0.20], z(1410) = −3.08, p = 0.002, OR = 0.58 respectively; Fig. 4). Note that when only focusing on trials in which participants declared not hearing the stimulus, an effect of unconscious processing for the 50 ms stimuli was found, as well as a marginally significant effect for the 40 ms stimuli (M = 0.17, 95% CI = [0.05,0.30], z(1053) = 2.72, p = 0.007, OR = 1.19; M = 0.13, 95% CI = [−0.01, 0.26], z(389) = 1.91, p = 0.056, OR = 1.13. respectively). Critically though, this above chance performance was still much lower than performance in P-trials (β = −0.82, 95% CI = [−1.11, −0.53], z (3479) = −5.52, p < 0.001, OR = 0.44; β = −0.74, 95% CI = [−1.04, −0.46], z(3479) = −5.01, p < 0.001, OR = 0.48; β = −0.7, 95% CI = [−0.99, −0.41], z(3479) = −4.7, p < 0.001, OR = 0.5, respectively for 30, 40, 50 ms stimuli).

4. Discussion

Can we experimentally study the phenomenal aspects of experience at their purity, and even separate them from access-related processes? This question has been a matter of ongoing debate (Aru, Bachmann, Singer, & Melloni, 2012; Block, 1995), considered by some as one of the greatest challenges in the neuroscientific study of consciousness (Block, 2019). Others still deny that phenomenology can and should be separated from access (Cohen & Dennett, 2011), stressing that access – and its associated cognitive functions – covary with subjective experience (Dehaene, Lau, & Kouider, 2021). Indeed, thus far there has been no convincing demonstration that the two can be empirically dissociated. Here, we show that this can, in principle, be done, though quite rarely (i.e., only in few trials and in a subsample of the participants) and in very specific conditions.

We created a unique situation that meets the three conditions necessary to demonstrate P-without-A consciousness (Block, 1995): (a) participants did not have online access to the stimulus while it was presented. Yet, (b) their performance in the discrimination task suggested they had retrospective access to the phenomenal aspects of the stimulus. Note that participants’ reports of not hearing the stimulus occurred despite repeated exposure to the stimulus and the task. Thus, failing to hear the stimulus did not stem from it being presented below the threshold of perception, or from misunderstanding the task. Importantly, the analysis of the NC-trials suggested that (c) unconscious processing of the stimulus cannot explain retrospective performance (though see further discussion below).

Experiment 2 further ruled out hearing the offset of the sound signal as an alternative explanation, as discrimination in trials where only the offset was heard was either at chance, or lower than in P-trials. This finding mitigates the concern that hearing only the abrupt stimulus offset is sufficient for discriminating the pink noises due to cognitive inference. Notably, the results are also not compatible with simple echoic memory (e.g., when not paying attention to your teacher talking, and then being asked to repeat what the teacher said, you can reconstruct it). This is because, as opposed to the teacher example, in our experiments participants denied hearing the sound in real time, despite being directly asked about it and trying to hear it. This also differentiates our results from ones showing the effects of a post-cue on perception (Sergent et al., 2013); here, the abrupt offset (possibly serving as a post-cue) does not lead to experiencing the stimulus as taking place now, but rather to it having been experienced in the past, while already being over now. Thus, the offset cue in our case does not lead to experiencing the stimulus itself, but to the experience of realizing it had been experienced before. In the post-cue work, on the other hand, the post-cue led to the stimulus being experienced as occurring at that time.

Arguably, this supports the proposed P vs. A dissociation, which has so far been mainly a theoretical one (Block, 1995). It also paves the way

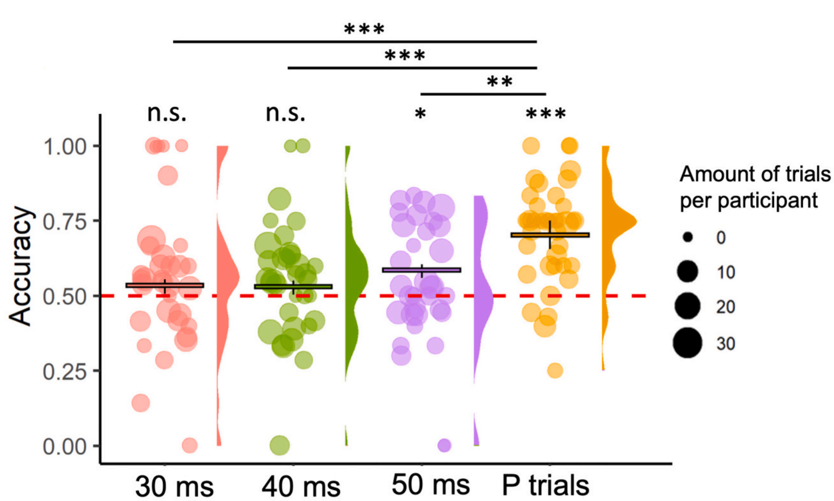

Fig. 4. Experiment 2 results compared to P-trials results. Note that the three left columns depict the results from Experiment 2 (N = 40), while the rightmost column presents the results from Experiment 1 (N = 40, presented also in Fig. 3) to facilitate their comparison. The same annotations used in Fig. 3 apply here.

to studying the neural correlates of P-without-A-consciousness, previously held to be an unsurmountable challenge (Block, 1995; Chalmers, 1995). Even the more recent suggestion of a ‘no-cognition’ paradigm (Block, 2019) concerns a method where P-consciousness is dissociated from attention (for discussion see Block, 2020; Phillips & Morales, 2020), but not from A-consciousness itself. Past claims for the neural correlates of P-without-A consciousness (Fahrenfort et al., 2017; Vandenbroucke et al., 2014) lack a key ingredient that is present here; the retrospective access participants have to the qualitative nature of the stimulus. That is, in those studies, participants deny having *any access* to the stimuli, and the authors then interpret neural activations that are content-specific as correlates of P-consciousness. But what guarantees that this is indeed the case, and that these activations do not simply reflect unconscious processing? The rationale there is that the probed processes involve illusory perceptual integration (in the form of Kanizsa shapes, for example), and that such integration is a hallmark of phenomenal experience. This is an assumption held by some theories (e.g. Lamme, 2020), but not all of them (e.g. Lau & Rosenthal, 2011), and it has yet to be convincingly demonstrated (for a study claiming for unconscious processing of the Kanizsa illusion see Wang, Weng, & He, 2012, and for reviews on other types of integration without awareness – or there lack of see Hirschhorn, Kahane, Gur-Arie, Faivre, & Mudrik, 2021; Mudrik, Faivre, & Koch, 2014). In our case, conversely, participants directly reported that something has changed once the stimulus had stopped, indexing in retrospect that they indeed had a previous experience to which they had no access in real time. Their performance in the objective task further confirmed that this experience also had a clear qualitative character, as they could differentiate it from other experiences. Comparing neural activations in such trials with those during NC trials could uniquely isolate the underpinnings of P-without-A consciousness (though, notably, the paradigm would have to be tweaked to yield more P-trials).

Yet though our findings strengthen the P vs. A dissociation, they also highlight its limitations. First, they demonstrate how difficult it is to find cases where P and A indeed dissociate, showcasing that this is the exception, rather than the rule. Second, they demonstrate one of the greatest challenges this dissociation faces: differentiating between P-consciousness and unconscious processing (Phillips, 2011), since the way to do so either relies on report (A-consciousness), or requires additional assumptions (e.g., about unconscious information being un-integrated or not allowing grouping) that have yet to be proven. Here too, one could always claim that we did not fully rule out unconscious processing as an explanation. Indeed, this might be an impossible task, since there is no agreed upon marker that could differentiate between the two (especially because so far, the dissociation has been mostly theoretical). Our attempt to do so relies on two pieces of evidence: at the subjective level, participants directly state that they experience themselves as having had an experience to which they did not have access in real time. They say they ‘realize’ that they actually did hear the sound before it disappeared, despite reporting otherwise. At the objective level, the difference in performance between NC-trials and P-trials, where unconscious processing should in principle be the same, provides further support to this claim. Further research might dig deeper into this question, to investigate the difference between retrospectively noticing the offset of a prolonged sound to which one has adapted, and a shortly-presented stimulus.

5. Conclusions

In this study, we report results that get us closer than ever to providing direct evidence for phenomenal consciousness without access. While also highlighting the limitations of this dissociation (most importantly, the difficulty to differentiate between phenomenal consciousness and unconscious processing), our results do show that phenomenal consciousness can be empirically distilled both from access consciousness, and possibly also from unconscious processing. With the novel paradigm we developed, future research could explore these different types of processing further, and even track the neural correlates of each type.

Data access

The preregistration for Experiment 1 can be accessed at: https://doi.org/10.17605/OSF.IO/QRU3B.

Experiment 2 was not preregistered.

Deidentified data for both experiments along with a codebook and the data-analysis scripts are posted at:

https://osf.io/gbqnk/?view_only=0ca93a6300cc44a894af467f929dc21d;

The materials used in these studies are also available there.

Declaration of Competing Interest

The authors declare no competing interests.

Data availability

Deidentified data for all experiments along with a codebook and the data-analysis scripts are shared in the file "open practice".

Acknowledgments

L.M. is CIFAR Tanenbaum Fellow in the Brain, Mind, and Consciousness program.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cognition.2023.105529.

References

Agresti, A. (2003). *Categorical data analysis (3rd ed.)*.
Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience and Biobehavioral Reviews, 36*(2), 737–746. https://doi.org/10.1016/j.neubiorev.2011.12.003
Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences, 18*(2), 227–247. https://doi.org/10.1017/S0140525X00038188
Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences, 9*(2), 46–52. https://doi.org/10.1016/j.tics.2004.12.006
Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences, 30*(5–6), 481–548. https://doi.org/10.1017/S0140525X07002786
Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences, 15*(12), 567–575. https://doi.org/10.1016/j.tics.2011.11.001
Block, N. (2014). Rich conscious perception outside focal attention. *Trends in Cognitive Sciences, 18*(9), 445–447. https://doi.org/10.1016/j.tics.2014.05.007
Block, N. (2019). What is wrong with the no-report paradigm and how to fix it. *Trends in Cognitive Sciences, 23*(12), 1003–1013. https://doi.org/10.1016/j.tics.2019.10.001
Block, N. (2020). Letter finessing the bored monkey problem. *Trends in Cognitive Sciences, 24*(1), 1. https://doi.org/10.1016/j.tics.2019.12.012
Bogomolov, M., Peterson, C. B., Benjamini, Y., & Sabatti, C. (2021). Hypotheses on a tree: New error rates and testing strategies. *Biometrika, 108*(3), 575–590.
Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*(4), 433–436. https://doi.org/10.1163/156856897X00357
Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We see more than we can report: “Cost free” color phenomenality outside focal attention. *Psychological Science, 25*(7), 1394–1403. https://doi.org/10.1177/0956797614532656
Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies, 2*(3), 200–219. https://doi.org/10.4324/9780203826430-11
Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences, 15*(8), 358–364. https://doi.org/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. https://doi.org/10.1016/j.tics.2016.03.006
De Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition, 18*(3), 569–577. https://doi.org/10.1016/j.concog.2009.03.002
Dehaene, S., Lau, H., & Kouider, S. (2021). What is consciousness, and could machines have it? *Robotics, AI, and Humanity*, 43–56.

Fahrenfort, J. J., Van Leeuwen, J., Olivers, C. N. L., & Hogendoorn, H. (2017). Perceptual integration without conscious access. *Proceedings of the National Academy of Sciences, 114*(14), 3744–3749. https://doi.org/10.1073/pnas.1617268114

Hirschhorn, R., Kahane, O., Gur-Arie, I., Faivre, N., & Mudrik, L. (2021). Windows of integration hypothesis revisited. *Frontiers in Human Neuroscience, 14*(January), 1–9. https://doi.org/10.3389/fnhum.2020.617187

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. https://doi.org/10.1016/j.tics.2010.04.006

Kouider, S., Sackur, J., & de Gardelle, V. (2012). Do we still need phenomenal consciousness? Comment on Block. *Trends in Cognitive Sciences, 16*(3), 140–141. https://doi.org/10.1016/j.tics.2012.01.003

Lamme, V. A. F. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks, 17*(5–6), 861–872. https://doi.org/10.1016/j.neunet.2004.02.005

Lamme, V. A. F. (2020). Visual functions generating conscious seeing. *Frontiers in Psychology, 11*(February), 83. https://doi.org/10.3389/fpsyg.2020.00083

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23*(11), 571–579.

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*(8), 365–373. https://doi.org/10.1016/j.tics.2011.05.009

Mudrik, L., Faivre, N., & Koch, C. (2014). Information integration without awareness. *Trends in Cognitive Sciences, 18*(9), 488–496. https://doi.org/10.1016/j.tics.2014.04.009

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

Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review, 83*(4), 435–450.

Phillips, I. B. (2011). Perception and iconic memory: What Sperling doesn't show. *Mind & Language, 26*(4), 381–411. https://doi.org/10.1111/j.1468-0017.2011.01422.x

Phillips, I. B., & Morales, J. (2020). The fundamental problem with no-cognition paradigms. *Trends in Cognitive Sciences, 24*(3), 165–167. https://doi.org/10.1016/j.tics.2019.11.010

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences, 3*(1), 1–23.

Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Current Biology, 23*(2), 150–155. https://doi.org/10.1016/j.cub.2012.11.047

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74*(11), 1–29. https://doi.org/10.1167/13.15.22

Tye, M. (1997). Qualia. In *In Stanford encyclopedia of philosophy*.

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

Wang, L., Weng, X., & He, S. (2012). Perceptual grouping without awareness: Superiority of Kanizsa triangle in breaking interocular suppression. *PLoS One, 7*(6), Article e40106.

Yaron, I., Melloni, L., Pitts, M., & Mudrik, L. (2021). The consciousness theories studies (ConTraSt) database: Analyzing and comparing empirical studies of consciousness theories. *BioRxiv.*. https://doi.org/10.1101/2021.06.10.447863