

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Confidence in absence as confidence in counterfactual visibility

Permalink

<https://escholarship.org/uc/item/2sk2d257>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Mazor, Matan

Schipper, Maya

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Confidence in absence as confidence in counterfactual visibility

Maya Schipper & Matan Mazor

University of Oxford

Abstract

When things are perceived clearly they can be detected with confidence. But under what conditions can one be confident that something is absent? Here we use a meta-perceptual illusion to show that confidence in absence scales not with visibility itself, but with the subjective belief that a stimulus would have been visible, if present. In two pre-registered experiments, participants detected the presence or absence of letters in frames of dynamic noise, and rated their decision confidence. Across trials, stimuli could appear bigger or smaller. Critically, while perceptual sensitivity was increased for smaller stimuli, participants' meta-perceptual beliefs (measured with post-experiment debriefing and prospective confidence ratings) were that larger letters were easier to detect. Accordingly, while confidence in presence scaled with objective visibility (and was therefore higher for smaller stimuli), confidence in absence scaled with beliefs about counterfactual visibility (and was therefore higher for bigger stimuli). This dissociation between the effect of stimulus size on confidence in presence and absence diminished as the experiment progressed: a sign of meta-perceptual learning. Furthermore, the effect of size on confidence in absence, but not in presence, correlated with a meta-perceptual parameter from an ideal observer model of perceptual detection, fitted to decision and response time data alone. Overall, we conclude that confidence in absence closely tracked participants' model-derived expectations about the visibility of counterfactual stimuli.

Keywords: perception, metacognition, counterfactual reasoning, absence

Introduction

Perceptual decisions vary not only in content, but also in levels of subjective persuasion: while some decisions are made with confidence and precision, others are accompanied by a subjective feeling of doubt and uncertainty. This subjective sense of confidence carries significance: in most settings, decisions that are made with higher levels of confidence are more likely to be correct (Fleming, 2024; Henmon, 1911; Nelson, 1990). In recent years, much research has been devoted to identifying the computational underpinning of perceptual confidence. According to inferential, or Bayesian, accounts, confidence reflects the estimated probability of a decision to be correct, as extracted based on prior knowledge and perceptual evidence, sometimes subject to computational constraints (Adler & Ma, 2018; Aitchison et al., 2015; Meyniel et al., 2015). Alternative, heuristic accounts, argue that confidence reflects a simple readout of the magnitude, or precision, or perceptual evidence (Calder-Travis et al., 2024; Xue et al., 2024). According to these accounts, confidence in a

perceptual decision tracks its probability of being correct only by virtue of tracking the amount of perceptual evidence in its support.

Critically, computational accounts of perceptual confidence—both inferential and heuristic—have been primarily designed with discrimination tasks in mind, in which decisions are made about the category of a presented stimulus (Calder-Travis et al., 2024; Rausch et al., 2018; Shekhar & Rahnev, 2024; Xue et al., 2024). Critically, however, the two model families are most clearly dissociated in detection tasks, in which decisions are made about the presence or absence of stimuli. Previous work (Mazor et al., 2025) showed that detection decisions are based not only on perceptual evidence (visibility), but also the believed probability of obtaining such evidence (beliefs about visibility)—a model-derived quantity that cannot be derived from the strength of perceptual evidence. The latter is especially salient in perceptual decisions in the absence of a stimulus (counterfactual visibility: “*I would have seen the target if it was present*”). To illustrate their point, Mazor and colleagues had participants perform a hard detection task of partly occluded stimuli. Decision confidence was similarly affected by occlusion when a target was present or absent. According to their model, despite this superficial similarity between the effects of occlusion on decisions about presence and absence, these two effects had different origins: occlusion affected confidence in presence because it reduced the quality of perceptual evidence, but it affected confidence in absence because participants believed that it reduced the quality of perceptual evidence, making them suspect that they might have missed the target.

Like occlusion, manipulations that affect the visibility of stimuli normally have a similar effect on beliefs about visibility. But to show that counterfactual visibility uniquely determines confidence in absence, we needed to identify a manipulation that dissociates counterfactual visibility from true visibility. That is, we need a manipulation that affects visibility in a way that systematically deviates from people's intuitive beliefs about their own perception. Specifically, in the present study, we have chosen to manipulate stimulus size. Evidence for a size-related bias in perception has been found in, for example, studies of the font size effect (e.g., Price et al., 2016; Rhodes & Castel, 2008), which suggest that larger cues can influence judgments of learning, either through greater encoding fluency (Koriat, 1997), or *a priori* beliefs (Yang et al., 2018).

Thus, participants may erroneously believe that stimulus visibility scales with stimulus size even when this

is not the case. Now, if confidence in absence is indeed uniquely determined by such beliefs about visibility, then we should find that it tracks this bias and decreases with smaller stimuli, and that this does not happen for confidence in presence.

Results

All code, anonymized data, pre-registrations, and live demos of the experiments can be accessed at github.com/self-model/confidenceInAbsenceSize.

217 English-speaking adult participants were recruited via Prolific.com to take part in Exp. 1 (pre-registration: osf.io/r62nm), in which they detected the presence or absence of a target letter (A or S, on different trials) in dynamic patches of visual noise. The target letter was present on 50% of the trials, and in the remaining 50% target-absent trials stimuli consisted of pure noise, generated by sampling, with replacement, pixels from the target image. The stimulus appeared on the screen, refreshing at 15 frames per second, until a response was recorded, at which point the participant rated their decision confidence on an analog scale. After applying our pre-registered exclusion criteria (above-chance accuracy, and no more than 25% of trials slower than 5 seconds or faster than 100 milliseconds), 214 participants were included in the analysis. On average, participants responded correctly on 0.93 ($SD=0.08$) of trials, and had a slight but significant bias to report target absence (proportion of “target present” responses: $M=0.48$, 95% CI [0.48,0.49], $t(213)=-3.69$, $p<.001$).




Figure 1: Experiment 1. A) Stimuli. Left: example frames from target-present (blue) and target-absent (red) trials. Right: on different trials, stimuli appeared big or small. B) trial structure. Participants performed two 16-trial detection blocks in which the target was the letter S and two blocks in which the target was the letter A. Stimulus presentation was preceded by a frame, and followed by an analog confidence rating scale. The order of the two letters was randomised between participants.

Our critical manipulation was of stimulus size. Stimuli were 18-by-12 pixels in size, but on different trials

individual pixels appeared large (5-by-5 screen pixels per stimulus pixel) or small (3-by-3 screen pixels per stimulus pixel). Letters were easier to detect when stimuli appeared smaller (d' for small stimuli: 3.05; d' for big stimuli: 2.88; difference: $t(213)=3.88$, $p<.001$), perhaps due to the spatial low frequency, pixelated nature of stimuli (Schyns & Oliva, 1994). Accordingly, participants were faster to detect letters when they appeared small (mean median reaction time in correct letter detections in seconds: small: 2.03 seconds; big: 2.15 seconds; $t(213)=6.04$, $p<.001$). They were also more confident in making decisions about presence when stimuli appeared small (mean confidence in correct letter detections on a 0-1 scale: small: 0.83; big: 0.80; $t(213)=-6.13$, $p<.001$), in line with confidence in presence tracking the objective visibility of stimuli. Finally, the three effects were correlated: participants whose perceptual sensitivity benefited more when stimuli appeared smaller were also relatively more confident ($r=.26$, 95% CI [.13,.38], $t(212)=3.92$, $p<.001$) and faster ($r=-.15$, 95% CI [-.28,-.02], $t(212)=-2.19$, $p=.029$) in their correct letter detections in small compared to big displays.

Critically, however, this negative effect of stimulus size on target visibility went against participants' reported beliefs. When asked at the end of the experiment—that is, after having experienced that small stimuli are easier to detect—whether stimulus size made any difference to task difficulty, only a small minority of 38/218 participants reported that bigger letters were harder to detect, 85 participants reported no difference, and the remaining 92 reported that smaller stimuli were harder to detect. Participants' beliefs were not linked to the objective effect of size on their perceptual sensitivity ($F(2,211)=0.52$, $MSE=0.42$, $p=.595$). Exp. 2 below provides additional evidence that participants erroneously believed that bigger stimuli were easier to detect, using prospective confidence ratings. Together, our design produced a dissociation between objective and perceived difficulty: while detection was objectively easier when stimuli appeared small, participants believed the opposite was the case.

Confidence in absence tracks believed, not objective, visibility

With this dissociation, we returned to decisions about absence. As per our logic, these decisions are made not upon the accumulation of sufficient perceptual evidence for absence, but based on a failure to accumulate evidence for presence, together with a belief that such evidence would have been available if a target was present (Mazor, et al., 2025). In line with this, we replicated detection asymmetries in reaction time (decisions about absence were on average slower by 0.35 seconds; $t(213)=11.82$, $p<.001$) and confidence (confidence in absence was on average lower by 0.05 on a 0-1 scale; $t(213)=6.75$, $p<.001$).

Importantly, if decisions about absence are made based on a belief about visibility, rather than visibility itself, confidence in those trials should be lower not when visibility was low (big stimuli), but when visibility was

believed to be low in participants' internal model of their own perception (small stimuli; "*It would have been hard to see the target in this small display, so I might have missed it*"). Indeed, unlike confidence in presence, which was higher for smaller stimuli, confidence in absence was higher for bigger stimuli (mean confidence in correct rejections: small: 0.76; big: 0.78; $t(213)=4.68$, $p<.001$). Furthermore, unlike the effect of stimulus size on confidence in presence, here we observed no correlation with the effect of stimulus size on perceptual sensitivity ($r=.03$, 95% CI $[-.11,.16]$, $t(212)=0.38$, $p=.703$). Participants' confidence in absence was determined by their beliefs about visibility, rather than by visibility itself.




Figure 2: A) visual sensitivity (measured as d') as a function of stimulus size. B) mean confidence in correct responses as a function of target presence and size. Error bars represent the standard error; ***: $p<.001$. C) correlations between the effect of size on perceptual sensitivity and on confidence, separately for target-present and target-absent decisions.

An ideal observer model of perceptual detection dissociates between true and believed visibility

As shown above, our paradigm successfully dissociated between true and believed visibility. A recent model of perceptual detection allows to formalise this dissociation, describing perceptual detection as a series of momentary decisions made by an ideal observer in a partially observed environment (for a more technical description of the model, see Mazor et al., 2025). On each trial, the agent is in one of two world states: a target letter is either objectively present

or absent. The agent does not have direct access to the world: it only has access to the noisy all-or-none activations of a sensor, which activates with probability θ_{absent} when a target is absent, and with a higher probability $\theta_{present}$ when a target is present. The agent observes a sequence of activations and inactivations at a rate of 20 per second, and, based on this information, decides whether to commit to a "target present" or a "target absent" decision, or, crucially, to accumulate more evidence. The trial ends once the agent commits to a decision, and the agent is internally rewarded only if its decision is correct.

Critically, the agent does not know the true values of θ_{absent} and $\theta_{present}$. Instead, it has subjective beliefs about these values $\bar{\theta}_{absent}$ and $\bar{\theta}_{present}$, which may be more or less accurate. This way, the model can describe a setting where perceptual information is of low quality ($\theta_{absent} \approx \theta_{present}$), but the agent believes it is of high quality ($\bar{\theta}_{absent} < \bar{\theta}_{present}$) or vice versa. Finally, in deciding how to act, the agent follows the optimal policy to maximise its long-term reward, based on its beliefs about visibility and subjective temporal discounting of value (which is assumed to take an exponential form). We derive the optimal policy by solving the Bellman equation using backward induction. As a consequence, this model does not have free parameters that correspond to decision biases, criteria or boundaries. Instead, model parameters correspond to properties of the sensor, the agents' beliefs about the sensor, and the agents' value function and decision noise. The model is fit to decisions and decision times, but qualitative predictions about confidence can be made by taking the model-derived subjective probability of being correct at the time of making the decision.

Finally, increasing the size of stimuli scales the probability of sensor activation by a factor of α , corresponding to the true perceptual effect of size on visibility. Whenever $\alpha > 1$, increasing stimulus size makes stimuli more visible, and whenever $\alpha < 1$ it makes them less visible. The agent holds a subjective belief about this value, $\bar{\alpha}$, which it uses to flexibly interpret evidence and set its subjective decision boundary as a function of stimulus size.

Ten model parameters were fitted to the data of individual subjects, with the parameters of interest being α and $\bar{\alpha}$, respectively corresponding to the true and believed effects of size on stimulus visibility. These were allowed to vary between 0.58 to 2.22. To account for noise in the decision-making process, action selection followed a softmax distribution. A previous study that used this model showed good parameter recovery (Mazor et al., 2025).

To examine the model's success in capturing qualitative patterns found in behavioural data, we simulated data from the fitted model parameters. The simulated data accurately matched participants' actual response time and error rates, predicting an advantage for smaller stimuli when a target is

present, but not when a target is absent (see Fig. 3B, left and middle panels). In contrast, while the model correctly predicted the increase in confidence for smaller stimuli in target-present trials, it failed to predict that the pattern should be opposite in target-absent trials (Fig. 3B)




Figure 3: A) A schematic description of the ideal observer model. B) Comparison of human and simulated data for Exp. 1, for response time, error rate, and confidence. Lines and points represent human data; semi-transparent rectangles represent data simulated from model parameters, which were fitted to reaction time and accuracy (but not confidence) data of individual participants. Rectangles are centred at their mean value, and their height is twice the standard error.

We next examined the fitted parameters. Across participants, α was 0.87 (SD=0.37) and significantly lower than 1 (a t-test on $\log(\alpha)$ against 0: $t(213)=-7.90$, $p<.001$), meaning that, according to the model, big stimuli were objectively less visible than small stimuli. Surprisingly, α was also slightly lower than 1 (mean=0.99, SD=0.39; a t-test on $\log(\alpha)$ against 0: $t(213)=-3.29$, $p=.001$), meaning that, according to the model, big stimuli were also believed to be less visible. This goes against participants' self-reported difficulty ratings and confidence in absence in Exp. 1. Critically however, the model correctly produced a significant difference between α and $\bar{\alpha}$ across individuals ($t(213)=-7.11$, $p<.001$), thereby capturing the fact that participants systematically misrepresented the effect of stimulus size on stimulus visibility.

Some support for our interpretation of these model parameters as reflecting true and believed visibility comes from their correlations with confidence effects. α , which reflects the effect of size on stimulus visibility, was strongly correlated with the effect of size on confidence in target-presence ($r_s=-.37$, $p<.001$), but not at all in target-absence ($r_s=.03$, $p=.662$; Fig. 4, first row).

Conversely, $\bar{\alpha}$, which reflects participants' beliefs about the effect of size on stimulus visibility, was not correlated with the effect of size on confidence in presence ($r_s=-.08$, $p=.275$), but it correlated with the effect of size on confidence in absence ($r_s=-.17$, $p=.014$). Notably, these are correlations between model parameters and out-of-sample data which was not available to the model during the parameter fitting process (the model was fitted to decisions and decision times, not confidence ratings). This double dissociation supports our psychological interpretation of these two model parameters as reflecting true and believed effects of size on visibility, as well as our broader theoretical claim that confidence in presence reflects true visibility and confidence in absence reflects believed, or counterfactual, visibility.




Figure 4: A) correlations between parameter α and the effect of size on confidence, separately for target-present and target-absent decisions. B) correlations between parameter $\bar{\alpha}$ and the effect of size on confidence, separately for target-present and target-absent decisions. ***: $p<.001$; *: $p<.05$.

Time-resolved analysis reveals meta-perceptual learning within the experiment

If participants started the experiment expecting bigger letters to be easier to detect, by the end of it they had experienced 15 minutes of feedback to the contrary, such that small letters were systematically easier to detect. In our next analysis, we asked whether experience with the task resulted in learning about the true effect of size on perceptual sensitivity, and whether this was evident in confidence ratings. We therefore computed the difference in confidence in correct responses as a function of stimulus

size, block number (each block consisted of 16 trials) and target presence (see Fig. 5).

A clear pattern emerged: the effect of size on confidence systematically decreased as a function of block number, both when a target was absent ($F(1,637.92)=13.86$, $p<.001$ for the fixed effect in the model $\Delta\text{confidence} \sim \text{block} + (1|\text{subject})$), and, surprisingly, when a target was present (although this decrease was not as steep; $F(1,635.17)=5.80$, $p=.016$). While our theoretical model can account for the target-absent effect as reflecting meta-perceptual learning, the target-present effect is harder to explain, especially given that the true effect of stimulus size on visibility, measured with detection sensitivity, remained similar between blocks ($F(1,641)=0.59$, $p=.443$). Tellingly, the two confidence effects were statistically independent: participants who showed the steepest drop in effect size across blocks in target-absent trials did not necessarily show a similar drop in target-present trials (correlation between slopes from the model $\Delta\text{confidence} \sim \text{block}$ fitted to target-present and target-absent trials of individual subjects: $r=-.04$, 95%CI $[-.18, .09]$, $t(212)=-0.63$, $p=.527$), suggestive of a different mechanism behind the two effects. We return to this point in the Discussion.




Figure 5: The effect of stimulus size on decision confidence in target-present trials (in blue) and in target-absent trials (in red) as a function of block number. Error bars represent the standard error; *, $p<.05$, **, $p<.01$, ***, $p<.001$

Prospective confidence ratings follow beliefs about visibility

We reasoned that if confidence in absence reflects participants' beliefs about perception, rather than anything about the perceptual input itself, it should align with participants' confidence prior to seeing the stimulus at all. In Exp. 2, we tested this hypothesis.

147 English-speaking adult participants were recruited via Prolific.com to participate in Exp. 2 (pre-registration: osf.io/z652f). After applying our pre-registered exclusion criteria (above-chance accuracy, and no more than 25% of trials slower than 5 seconds or faster than 100 milliseconds), 138 participants were included in the analysis.

Similar to Exp.1, the critical manipulation was one of stimulus size, with stimuli appearing either large or small on

different trials. However, this time, in the first two blocks of the experiment we asked participants to rate their confidence *before* deciding whether the target is present or absent. To this end, participants were presented with an empty black frame indicating the size of the stimulus they were about to observe and asked to rate how confident they are that they will be correct in the upcoming decision. Only afterwards were they presented with the noisy stimulus and asked to detect the presence or absence of the target (Fig. 6A, first row). Blocks 3 and 4 followed a similar trial structure to that of Exp. 1 (Fig. 6A, second row). Since our time-resolved analysis from Exp. 1 revealed rapid meta-perceptual learning within 64 experimental trials, and since we were most interested in participants' miscalibrated beliefs about their perception, Exp. 1 included only 32 trials: 16 with prospective and 16 with post-decisional confidence ratings.




Figure 6: Experiment 2. A) Trial structure in the main blocks of the experiment, in prospective and post-decisional confidence blocks (trials were preceded by a fixation cross and followed by a feedback screen, as in Exp. 1). B) visual sensitivity (measured as d') as a function of stimulus size. C) The effect of frame size on prospective decision confidence. D) mean post-decisional confidence in correct responses as a function of target presence and size. Error bars represent the standard error; *, $p<.05$, ***, $p<.001$.

Here, too, letters were easier to detect when the stimuli appeared smaller (d' for small stimuli: 2.41, d' for big stimuli: 2.13; $t(137)=4.52$, $p<.001$; Fig. 6B). Critically, prospective confidence ratings showed the opposite pattern, with participants being significantly more confident that they will be correct in their upcoming decision when they expected the stimulus to be big, compared to small (mean confidence in upcoming letter detections on a 0-1 scale: small: 0.73; big: 0.78; $t(137)=5.75$, $p<.001$; Fig. 6C). This is

consistent with a prior meta-perceptual belief that bigger stimuli are more visible.

Confidence in absence showed a similar pattern to prospective confidence (in full dissociation from objective accuracy), and was higher when stimuli were expected to appear big (mean confidence in correct rejections: small: 0.79, big: 0.81; $t(136) = 2.73$, $p = .007$). Intriguingly, confidence in presence was unaffected by the size manipulation (mean confidence in correct letter detections: small: 0.83, big: 0.83; $t(134) = -0.49$, $p = .626$), and so were target-present reaction times (mean median reaction time in correct letter detections in seconds: small: 2.28, big: 2.31, $t(137) = 0.64$, $p = .522$). Finally, across subjects, the effect of stimulus size on post-decisional confidence ratings was uncorrelated with the effect of stimulus size on perceptual sensitivity or with the effect of stimulus size on prospective confidence, both in target-present (all p 's > .30) and in target-absent trials (all p 's > .60).

Discussion

Like decisions about the presence of objects, decisions about absence vary in subjective confidence. Previous work argued that, despite apparent similarities, confidence in presence and in absence reflect two distinct quantities: the first is based on perceptual evidence, and the second on meta-perceptual beliefs ("*I would have seen it if it was there*"; Ghetti, 2003; Kanai et al., 2010; Mazor et al., 2020, 2025; Mazor & Fleming, 2022; Meuwese et al., 2014). For example, Kanai et al. showed that confidence in absence tracks objective accuracy when sensitivity is manipulated in ways that can be monitored by participants (such as attentional manipulations) but not when manipulated in ways that cannot be monitored (such as perceptual manipulations). Only attentional manipulations allowed participants to monitor the likelihood with which a target would have been detected, if present. Outside perception, participants were more confident that a region was mine-free when the overall density of mines in the surrounding area was high (thereby making the absence of evidence for mines in the region more informative; Hsu et al., 2017). Similarly, people are more confident that an event did not occur (for example, that a name was not on a previously learned list) when they believe they would have remembered if it did (for example, because they would have remembered their own name; Glanzer & Adams, 1985).

Here we provide additional support for this process dissociation, using a novel meta-perceptual illusion. Across two experiments, we show that confidence in presence scales with perceptual sensitivity, but that confidence in absence scales with participants' miscalibrated beliefs about visibility. This finding has theoretical significance for two timely debates in cognitive science. First, it lends support to inferential accounts of confidence formation, in which confidence—or at least confidence in absence—is not a mere perceptual read-out of stimulus strength (Shekhar & Rahnev, 2024; Xue et al., 2024) but the integration of perceptual input with prior expectations about the world

and, crucially, the self too (Katyal & Fleming, 2024; Olawole-Scott & Yon, 2024; Seow et al., 2021). And second, it aligns with decisions about absence involving sophisticated, metacognitive processes (De Cornulier, 1988; Levesque, 1986), making them a useful window into people's models of their own minds (Mazor, 2025).

While our proposed theoretical model accounts for much of the data, some patterns remain unexplained. First, a block-wise reduction in effect size of stimulus size on confidence in absence was mirrored by a similar (albeit weaker) reduction in the effect of size on confidence in presence (see Fig. 5). We interpreted the target-absent effect as the manifestation of meta-perceptual learning, but the target-present effect is harder to explain. One possibility is that a reduction in the effect of size on confidence in presence reflects perceptual learning, whereby participants learned to adjust their focus as a function of stimulus size (for example by squeezing their eyes, as some participants indicated in the open debrief). Crucially, however, participants did not get significantly better at the task as a function of block number, and the difference in perceptual sensitivity between small and big stimuli remained similar across blocks.

Second, while our model accounted for response and response time patterns as well as for the negative effect of stimulus size on confidence in presence, it incorrectly predicted that confidence in absence should follow a similar trend and be higher for smaller stimuli. Despite this failure to predict the overall effect of stimulus size on confidence in absence, the meta-perceptual parameter $\bar{\alpha}$ was correlated with the magnitude of this effect across participants (Fig. 4B). Together, it seems that beliefs about visibility may contribute to confidence in absence in ways that are not captured by our ideal observer model. This may happen if, for example, confidence integrates post-decisional evidence (Petrusic & Baranski, 2003; Yeung & Summerfield, 2012), or weighs task difficulty as an independent factor (Katyal & Fleming, 2024; Olawole-Scott & Yon, 2024; Seow et al., 2021).

Finally, while Exp. 2 confirmed that confidence in absence is affected by stimulus size in a way that parallels the effect of size on prospective confidence, it also produced a series of disappointing null findings, with no significant effect of size on confidence in presence and no correlations between size effects on post-decisional confidence and on prospective confidence or perceptual sensitivity. Disappointed by these null findings, we reminded ourselves that Exp. 2 had a smaller sample size and only 8 trials per block (half the number of trials in Exp. 1). As such, we are not confident that these correlations are truly absent — we cannot be certain that we would have detected them if they were present.

Acknowledgements

MM is supported by a post doctoral research fellowship from All Souls College at the University of Oxford. The

authors would like to acknowledge the use of the University of Oxford Advanced Research Computing facility in carrying out this work.

References

- Adler, W. T., & Ma, W. J. (2018). *Comparing Bayesian and non-Bayesian accounts of human confidence reports* | *PLOS Computational Biology*.
<https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1006572>
- Aitchison, L., Bang, D., Bahrami, B., & Latham, P. E. (2015). Doubly Bayesian Analysis of Confidence in Perceptual Decision-Making. *PLOS Computational Biology*, 11(10), e1004519.
<https://doi.org/10.1371/journal.pcbi.1004519>
- Calder-Travis, J., Charles, L., Yeung, N., & Bogacz, R. (2024). *Bayesian confidence in optimal decisions*.
<https://doi.org/10.17605/OSF.IO/QPSEM>
- De Cornulier, B. (1988). “Knowing Whether”, “Knowing Who”, and Epistemic Closure. In M. Meyer (Ed.), *Questions and Questioning* (pp. 182–192). De Gruyter.
<https://doi.org/10.1515/9783110864205.182>
- Fleming, S. M. (2024). Metacognition and Confidence: A Review and Synthesis. *Annual Review of Psychology*, 75(Volume 75, 2024), 241–268.
<https://doi.org/10.1146/annurev-psych-022423-032425>
- Ghetti, S. (2003). Memory for nonoccurrences: The role of metacognition. *Journal of Memory and Language*, 48(4), 722–739.
- Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, 13(1), 8–20. <https://doi.org/10.3758/BF03198438>
- Henmon, V. A. C. (1911). The relation of the time of a judgment to its accuracy. *Psychological Review*, 18(3), 186–201. <https://doi.org/10.1037/h0074579>
- Hsu, A. S., Horng, A., Griffiths, T. L., & Chater, N. (2017). When Absence of Evidence Is Evidence of Absence: Rational Inferences From Absent Data. *Cognitive Science*, 41(S5), 1155–1167.
<https://doi.org/10.1111/cogs.12356>
- Kanai, R., Walsh, V., & Tseng, C. (2010). Subjective discriminability of invisibility: A framework for distinguishing perceptual and attentional failures of awareness. *Consciousness and Cognition*, 19(4), 1045–1057.
<https://doi.org/10.1016/j.concog.2010.06.003>
- Katyal, S., & Fleming, S. (2024). *Gender and anxiety reveal distinct computational sources of underconfidence*. OSF. <https://doi.org/10.31234/osf.io/5w2vd>
- Koriat, A. (1997). Monitoring one’s own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of Experimental Psychology: General*, 126(4), 349–370.
<https://doi.org/10.1037/0096-3445.126.4.349>
- Levesque, H. J. (1986). Making believers out of computers. *Artificial Intelligence*, 30(1), 81–108.
[https://doi.org/10.1016/0004-3702\(86\)90068-8](https://doi.org/10.1016/0004-3702(86)90068-8)
- Mazor, M. (2025). *Inference about Absence as a Window into the Mental Self-Model*. *Open Mind*
- Mazor, M., & Fleming, S. M. (2022). Efficient search termination without task experience. *Journal of Experimental Psychology: General*, 151(10), 2494–2510. <https://doi.org/10.1037/xge0001188>
- Mazor, M., Friston, K. J., & Fleming, S. M. (2020). Distinct neural contributions to metacognition for detecting, but not discriminating visual stimuli. *eLife*, 9, e53900.
<https://doi.org/10.7554/eLife.53900>
- Mazor, M., Moran, R., & Press, C. (2025). Beliefs about perception shape perceptual inference: An ideal observer model of detection. *Psychological Review*.
- Meuwese, J. D., van Loon, A. M., Lamme, V. A., & Fahrenfort, J. J. (2014). The subjective experience of object recognition: Comparing metacognition for object detection and object categorization. *Attention, Perception, & Psychophysics*, 76, 1057–1068.
- Meyniel, F., Sigman, M., & Mainen, Z. F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, 88(1), 78–92.
<https://doi.org/10.1016/j.neuron.2015.09.039>
- Nelson, T. O. (1990). Metamemory: A Theoretical Framework and New Findings. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 26, pp. 125–173). Academic Press.
[https://doi.org/10.1016/S0079-7421\(08\)60053-5](https://doi.org/10.1016/S0079-7421(08)60053-5)
- Olawole-Scott, H., & Yon, D. (2024). Expectations About Precision Bias Metacognition and Awareness. *Psychological Review*, 128(1)(45).
<https://doi.org/10.1037/xge0001371>
- Petrusic, W. M., & Baranski, J. V. (2003). Judging confidence influences decision processing in comparative judgments. *Psychonomic Bulletin & Review*, 10(1), 177–183.
<https://doi.org/10.3758/BF03196482>
- Price, J., McElroy, K., & Martin, N. J. (2016). The role of font size and font style in younger and older adults’ predicted and actual recall performance. *Aging, Neuropsychology, and Cognition*, 23(3), 366–388.
<https://doi.org/10.1080/13825585.2015.1102194>
- Rausch, M., Hellmann, S., & Zehetleitner, M. (2018). Confidence in masked orientation judgments is informed by both evidence and visibility. *Attention, Perception, & Psychophysics*, 80(1), 134–154.
<https://doi.org/10.3758/s13414-017-1431-5>
- Rhodes, M. G., & Castel, A. D. (2008). Memory predictions are influenced by perceptual information: Evidence for metacognitive illusions. *Journal of Experimental Psychology: General*, 137(4), 615–625.
<https://doi.org/10.1037/a0013684>
- Schyns, P. G., & Oliva, A. (1994). From Blobs to Boundary Edges: Evidence for Time- and Spatial-Scale-Dependent Scene Recognition. *Psychological Science*, 5(4), 195–200.
<https://doi.org/10.1111/j.1467-9280.1994.tb00500.x>

- Seow, T. X. F., Rouault, M., Gillan, C. M., & Fleming, S. M. (2021). How Local and Global Metacognition Shape Mental Health. *Biological Psychiatry*, 90(7), 436–446. <https://doi.org/10.1016/j.biopsych.2021.05.013>
- Shekhar, M., & Rahnev, D. (2024). How do humans give confidence? A comprehensive comparison of process models of perceptual metacognition. *Journal of Experimental Psychology: General*, 153(3), 656–688. <https://doi.org/10.1037/xge0001524>
- Xue, K., Shekhar, M., & Rahnev, D. (2024). *A novel behavioral paradigm reveals the nature of confidence computation in perceptual decision making* | JOV | ARVO Journals. <https://jov.arvojournals.org/article.aspx?articleid=2801303>
- Yang, C., Huang, T. S.-T., & Shanks, D. R. (2018). Perceptual fluency affects judgments of learning: The font size effect. *Journal of Memory and Language*, 99, 99–110. <https://doi.org/10.1016/j.jml.2017.11.005>
- Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: Confidence and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594), 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>