Does the surveillance paradigm provide evidence for unconscious evaluative conditioning?
A Bayesian perspective

Benedek Kurdi and Melissa J. Ferguson
Cornell University, Ithaca, New York

Author Note

Benedek Kurdi and Melissa J. Ferguson, Department of Psychology, Cornell University.

Correspondence concerning this article should be addressed to Benedek Kurdi, Department of Psychology, Cornell University, Ithaca, NY 14850, email: bk493@cornell.edu.

Abstract

In evaluative conditioning (EC), a neutral conditioned stimulus acquires the valence of an intrinsically valenced unconditioned stimulus as a result of repeated pairings between the two. Evidence for EC in the absence of awareness of the pairings has been provided using the surveillance paradigm (Olson & Fazio, 2001). In this commentary, we report a Bayesian reanalysis of data from a multi-lab replication of this paradigm (Moran et al., 2020). Unlike the original frequentist analyses, we conclude that the replication does not provide convincing evidence in favor of unconscious EC under any of the exclusion criteria, including the ones originally used by Olson and Fazio (2001). Moreover, when informative priors relying on the totality of meta-analytic evidence are used, three of four analyses provide evidence against an effect. This commentary highlights the benefits of using Bayesian analytic approaches and cautions against inferring the presence of unconscious EC in the surveillance paradigm.

Keywords: Bayesian analyses; Bayesian equivalence testing; evaluative conditioning; incidental learning; unconscious cognition

Statement of Relevance

Our likes and dislikes are a powerful determinant of behaviors and outcomes across a multitude of domains. One important source of such preferences is what goes together with what in the environment: Initially neutral objects that co-occur with positive objects tend to become positive and initially neutral objects that co-occur with negative objects tend to become negative. This phenomenon is called evaluative conditioning (EC). Whether EC can occur incidentally (that is, without awareness of the pairings) has been a matter of considerable theoretical debate. We report new analyses of a large-scale replication project to show that Bayesian models do not provide convincing support for, and often even provide evidence against, unaware EC effects in the so-called surveillance paradigm (Olson & Fazio, 2001). Likes and dislikes may very well change incidentally through EC; however, alternative procedures will be needed to conclusively demonstrate such effects.

Does the surveillance paradigm provide evidence for unconscious evaluative conditioning?

A Bayesian perspective

Moran et al. (2020) recently conducted a large-scale registered replication of the so-called surveillance paradigm (Olson & Fazio, 2001) to investigate whether evaluative conditioning (EC) effects can emerge in the absence of conscious awareness of CS–US pairings. In the paradigm, participants are surreptitiously exposed to pairings of one Pokémon character (CS+) with positive words and images (US+) and a second Pokémon character (CS-) with negative words and images (CS-). The presence of unconscious EC effects is inferred if participants show a preference for the CS+ over the CS- but are unable to report CS–US contingencies on a retrospective memory test.

The multi-lab replication effort yielded mixed results. A small but statistically significant unconscious EC effect emerged when the original exclusion criteria by Olson and Fazio (2001) were used to exclude aware participants from analyses [R1]. However, when alternative exclusion criteria, including modified criteria suggested by Olson and Fazio and criteria suggested by Bar-Anan, De Houwer, and Nosek (2010) were used, the effect disappeared [R2]. Nevertheless, ambiguity remains given that the significant effect obtained under the original criteria and the nonsignificant effects obtained under the alternative criteria did not significantly differ from each other [R3]. Moran et al. also reported a meta-analytic estimate of all previous results obtained using the surveillance paradigm. This analysis produced a small but statistically significant effect when the authors did not correct for publication bias [R4]; however, correcting for publication bias reduced the effect to nonsignificance [R5].

Given that this replication project involved extensive resources, including 1478 participants recruited by twelve labs in ten countries, any additional insight on the collected data would be especially helpful. Here we report preregistered secondary analyses of the data reported by

Moran et al. (2020) in a Bayesian equivalence testing framework (Kruschke, 2018) in the hope of creating more clarity about the strength of the evidence in favor of unconscious EC effects in the surveillance paradigm.

Bayesian analyses offer considerable benefits over the frequentist analyses used by Moran et al. First, with large samples, frequentist analyses can yield statistically significant results even when the null hypothesis is more likely to be true than the alternative hypothesis (Benjamin et al., 2018). As such, the present analysis gives us the opportunity to reexamine the robustness of the replication project's main result [R1] to a different analytic framework. Nonsignificant findings, in turn, are inherently ambiguous in a frequentist framework (Dienes, 2014): They could indicate either lack of adequate power or the genuine absence of an effect. Therefore, it is important to probe whether alternative participant exclusions provide evidence in favor of the null hypothesis [R2] and whether different exclusion criteria, in fact, yield qualitatively similar results [R3]. Finally, Bayesian analyses allow for the incorporation of prior beliefs into models: Here we compared results under a default (uninformative) prior, the unadjusted meta-analytic estimate [R4], and the adjusted meta-analytic estimate [R5].

We fit intercept-only mixed effects models to the data, with standardized preference for the CS+ over the CS- as the dependent variable and random intercepts for data collection sites. In model 1, we placed a default (uninformative) prior on the intercept; in model 2, we used the unadjusted meta-analytic effect size reported by Moran et al. (2020) as an informative prior; and in model 3, we used the publication bias-adjusted meta-analytic effect size for the same purpose. We preregistered $\beta = 0.1$ as the smallest effect consistent with the alternative hypothesis (H_1); the area of the posterior distribution below this value was seen as consistent with the null hypothesis (H_0). The main quantity of interest, upon which we base our inferences, is the proportion of the

posterior distribution consistent with H_1 versus H_0 . If more than 95 percent of the posterior are consistent with H_1 or H_0 , we see this as decisive evidence in favor of the corresponding hypothesis; other values represent varying degrees of mixed evidence.

The outcome of the Bayesian reanalysis using default priors is shown in the top row of Figure 1. From such reanalysis, it is clear that the original Olson and Fazio exclusion criteria provide the relatively strongest evidence in favor of unconscious EC effects. However, with over 23 percent of the posterior distribution favoring H_0 , the evidence for H_1 is not decisive. As such, the present analysis strongly qualifies the main conclusion of the replication project [R1]: Although a statistically significant result may have been obtained in the original analysis, the Bayesian reanalysis shows that considerable portions of the posterior distribution around the estimate are actually consistent with a diminishingly small, or even negative, effect.

When alternative exclusion criteria were used, the data provided convincing evidence for H_0 , with proportions of the posterior distribution favoring no effect falling between 93 and 96 percent. As such, unlike the original frequentist analysis [R2], the current Bayesian analysis warrants the conclusion that the nonsignificant results obtained using alternative exclusion criteria did not emerge due to lack of statistical power; rather, the present models positively suggest an absence of unaware EC effects. In contrast, the present analyses indicate that the lack of significant differences between exclusion criteria [R3] was most likely due to lack of statistical power: Comparing different exclusion criteria in a Bayesian framework (top row, Figure 1) makes it clear that the original exclusion criteria and alternative exclusion criteria result in qualitatively different conclusions from the same data. Specifically, whereas the former yield some (limited) support for H_1 over H_0 , the latter provide clear evidence in favor of H_0 over H_1 .

Finally, we fit the same models using the unadjusted meta-analytic effect size (row 2, Figure 1) and the adjusted meta-analytic effect size (row 3, Figure 1) as informative priors, thus explicitly incorporating the results of previous work relying on the surveillance paradigm into our analyses [R4–R5]. The most important takeaway from these analyses is that, presumably given the large sample used in the replication project, the results seem quite robust to the choice of prior. Specifically, even the most lenient analysis (using the unadjusted meta-analytic effect size as the informative prior and the original exclusion criteria) does not provide convincing evidence for H_1 over H_0 ; under alternative exclusion criteria, the data provide consistent evidence for H_0 over H_1 .

Overall, the results of these Bayesian models not only fail to provide evidence for unconscious EC; rather, they overwhelmingly provide evidence against this possibility. However, it should be noted that unconscious EC has recently been observed in paradigms quite different from the surveillance paradigm investigated here (Greenwald & De Houwer, 2017; Jurchis, Costea, Dienes, Miclea, & Opre, in press). More generally, given the ubiquity of incidental learning across myriad domains of human cognition (e.g., Knowlton, Ramus, & Squire, 1992; Saffran & Kirkham, 2018; Schapiro, Rogers, Cordova, Turk-Browne, & Botvinick, 2013), it seems improbable that all evaluative learning requires conscious awareness. Nonetheless, what paradigm will be best suited to robustly demonstrate unconscious EC effects remains an open question. And, as the present analysis suggests, the surveillance paradigm most likely does not offer a sufficiently sensitive option.

References

- Bar-Anan, Y., De Houwer, J., & Nosek, B. A. (2010). Evaluative conditioning and conscious knowledge of contingencies: A correlational investigation with large samples. *Quarterly Journal of Experimental Psychology*, 63(12), 2313–2335.
<http://doi.org/10.1080/17470211003802442>
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., et al. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2(1), 6–10.
<http://doi.org/10.1038/s41562-017-0189-z>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5(e33400), 1507–17. <http://doi.org/10.3389/fpsyg.2014.00781>
- Greenwald, A. G., & De Houwer, J. (2017). Unconscious conditioning: Demonstration of existence and difference from conscious conditioning. *Journal of Experimental Psychology: General*, 146(12), 1705–1721. <http://doi.org/10.1037/xge0000371>
- Jurchis, R., Costea, A., Dienes, Z., Miclea, M., & Opre, A. (in press). Evaluative conditioning of artificial grammars: Evidence that non-conscious structures bias affective evaluations of novel stimuli. *Journal of Experimental Psychology: General*.
- Knowlton, B. J., Ramus, S. J., & Squire, L. R. (1992). Intact artificial grammar learning in amnesia: Dissociation of classification learning and explicit memory for specific instances. *Psychological Science*, 3(3), 172–179. <http://doi.org/10.1111/j.1467-9280.1992.tb00021.x>
- Kruschke, J. K. (2018). Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science*, 1(2), 270–280.
<http://doi.org/10.1177/2515245918771304>

- Moran, T., Hughes, S., Hussey, I., Vadillo, M. A., Olson, M. A., Aust, F., et al. (2020). *Incidental attitude formation via the surveillance task: A pre-registered replication of Olson and Fazio (2001)*. Manuscript submitted for publication.
- Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, 12(5), 413–417. <http://doi.org/10.1111/1467-9280.00376>
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69(1), 181–203. <http://doi.org/10.1146/annurev-psych-122216-011805>
- Schapiro, A. C., Rogers, T. T., Cordova, N. I., Turk-Browne, N. B., & Botvinick, M. M. (2013). Neural representations of events arise from temporal community structure. *Nature Neuroscience* 16(4), 486–492. <http://doi.org/10.1038/nn.3331>

Author Contributions

B. Kurdi conceived and performed the data analysis and interpretation. B. Kurdi and M. J. Ferguson drafted the manuscript. Both authors approved the final version of the manuscript for submission.

Open Practices Statement

The analyses reported in this article were formally preregistered; the preregistration can be accessed at <https://osf.io/sfh4v/>. The data analysis scripts are posted at <https://osf.io/u4mv8/> and the data and materials are posted at <https://osf.io/hs32y/>.

Figure 1. Posterior distribution of the unconscious evaluative conditioning (EC) effect under different priors and exclusion criteria. Positive scores correspond to the theoretically expected preference for CS+ over CS-. The dashed vertical line shows the posterior mean and the solid vertical line shows the smallest effect size of interest ($\beta = 0.1$). Areas displayed in red are consistent with the null hypothesis $H_0: B < 0.1$ and areas displayed in light blue are consistent with the directional alternative hypothesis $H_1: B \geq 0.1$. Percentages denote the proportion of the posterior distribution consistent with H_0 and H_1 , respectively.