

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Novel Target Detection Task Using Artificial Stimuli: The Effect of Familiarity.

Permalink

<https://escholarship.org/uc/item/16t1p38x>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 42(0)

Authors

McLaren, Rossy

Civile, Ciro

Cooke, Anna

et al.

Publication Date

2020

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

A Novel Target Detection Task Using Artificial Stimuli: The Effect of Familiarity.

Rossy McLaren (R.P.McLaren@exeter.ac.uk), Ciro Civile (C.Civile@exeter.ac.uk),
Anna Cooke (ac798@exeter.ac.uk) and Ian P.L. McLaren (I.P.L.McLaren@exeter.ac.uk)

School of Psychology, University of Exeter, UK.

Abstract

In this paper we demonstrate that a target detection task is facilitated when the background on which the target is presented is a familiar one, even though the target appears at a random location. We compare performance in that condition with one where the background is randomly generated and establish a significant difference between these two versions of the task in terms of both d' and criterion, C . We also go on to look at the effect of a tDCS procedure that we know to affect discrimination performance on this difference, discovering that it seems to reduce or reverse the difference in criterion for these two conditions. We ascribe this effect to the neurostimulation manipulation shifting the distribution of information used to reach a decision.

Keywords: Associative learning; Perceptual Learning; Error-based Modulation of Salience, Target Detection

Introduction

In this paper we look at the issue of target detection in humans, and, in particular, how target detection is influenced by the familiarity of the background against which it is detected, even when the target is not in any way predicted by that background. In other words, the target appears at a random location and there is a 50:50 chance it will occur on any trial. The basic paradigm we will use employs checkerboards, a 16 x 16 arrangement of black and white squares, that we can control to be either randomly determined from trial to trial or a fixed arrangement (see Figure 1b for an example). The target is itself a regular 5 x 5 arrangement of black and white squares in the shape of an H that can occur at any permissible location, with the only constraint being that the target must fit onto the checkerboard background. The participants' task is to signal, by pressing the appropriate key, whether the target is present or not on a given trial.

We've described the basic paradigm at this point because it helps to place this experiment in the context of previous literature. It is very reminiscent of the work done on search image (Tinbergen, 1960; Pietrewicz and Kamil, 1979; Plaisted and Mackintosh, 1995), in that it involves detection of a target against a "cryptic" background, and that would be a fair characterization of our paradigm. In those experiments, however, the manipulation typically

concerns familiarity with the targets rather than the background against which they are detected. The finding is that a target that has recently been detected has a short-term advantage over other possible targets that might occur, suggesting some sort of priming mechanism at work.

Another paradigm that does not match ours quite so obviously but may have more theoretical relevance, is that used by Phillips (1974). Using checkerboard stimuli, he demonstrated that it was easy to detect a difference of one square changed in a given stimulus if the two variants were presented successively with minimal delay and no masking. Complementing this result, he found that the task became very hard indeed if a delay containing a mask intervened between the two checkerboard presentations. The experiments pre-date those by Rensink et al (1997) which also showed (but this time for images in general) that change detection was easy when stimuli were alternated with no delay or mask, but much harder and seemed to require serial visual search when a mask filled delay was inserted.

These studies, and many like them, suggest that we have an ability to detect change very easily under certain circumstances. Suret and McLaren (2001) made use of this idea to perform an experiment that has considerable overlap with ours (though the paradigm used is quite different). They used the change detection paradigm, making participants in one group familiar with the background by using only one such checkerboard background for that purpose. The experiment involved alternating between two stimuli, both of which could be the background (half the trials), or one of them could be the background with one square changed (other half of the trials). The group familiarized with the background were better at detecting the changed square than participants in another group that had random checkerboards (with or without a changed square) used as a background in each trial. Familiarization with the checkerboard background clearly conferred an advantage in detecting changes to that checkerboard in that experiment and suggests that this might also be the case in our rather different target detection experiment.

If we view the target as simply a change to the familiar stimulus, then the prediction that target detection will be superior for the familiar rather than the random background naturally follows from these earlier results. So, this is by no means a counter-intuitive prediction, but it remains to be seen whether or not it is borne out.

One possible (and rather obvious) explanation for the improved change detection when a familiar background is used in Suret and McLaren's experiment is that the change is in some way distinctive or salient in the context of that familiar background. Suret and McLaren (2001) speculate that detection on the basis of novelty may be important here. One of their results was that introduction of a change in both checkerboards in the familiar stimulus condition (so that actually there was no "change" as such from one checkerboard to the other) provoked an unusually high number of false alarms on no change trials. This fits with some sort of novelty-based detection mechanism being triggered by this change to the familiar checkerboard, an idea examined next.

The original Suret and McLaren work was inspired by a theory of representation development that quite naturally provided a mechanism for change detection. McLaren, Kaye and Mackintosh (1989, henceforth MKM) offer a connectionist theory of latent inhibition and perceptual learning that uses error-based modulation of salience as one of its key mechanisms. For present purposes, the theory states that well predicted features (low error) of a stimulus will be relatively low in salience, but novel features that are not predicted by other features present will be relatively high in salience. It's easy to see how this would apply to the paradigm we will use here and to Suret and McLaren's experimental results. A familiar checkerboard background will possess features that are well predicted by the others present in that stimulus, and so those features will be relatively low in salience. A change, however, will be novel, and hence salient, and stand out. Thus, on this analysis, we can expect our target to be relatively more salient than other features present in the whole stimulus when presented on the familiar background, and hence stand out and draw attention to itself. A random background will not confer this advantage because all features will be equally unpredicted whether the target is present or not, and so they will all be equal in salience and serial search will be needed to locate it.

It was this logic that led to the final part of the experimental design we will employ here. We have recently shown that anodal tDCS for 10 minutes to

Fp3 at 1.5 mA has an effect that can be interpreted in terms of the MKM model as preventing the error-based modulation of salience we have just referred to. This claim follows from work both with checkerboards (McLaren, 1997; McLaren and Civile, 2011; Civile et al, 2014; Civile et al, 2016) and with faces (Civile, McLaren and McLaren, 2018; Civile, Obhi and McLaren, 2019). In these experiments the authors show that this tDCS procedure either reduces (faces) or eliminates (checkerboards) the expected inversion effect in a recognition experiment. It does this by reducing performance to upright stimuli, leaving that to inverted stimuli largely unaffected. This is exactly what would be expected if perceptual learning for the upright stimuli had been abolished, and, on the MKM model, this would equate to no longer having error-based modulation of salience enhancing the discriminability of those upright stimuli.

The full argument justifying these claims is lengthy, covers many years of research and we cannot give it in detail here. But the outline just offered serves to make clear the reasons for the final manipulation in this experiment. We will run two groups on our paradigm, one undergoing anodal tDCS as described, the other in a sham stimulation condition. The experiment will be run double blind, and we are interested in seeing what impact tDCS has on the predicted advantage for target detection on the familiar checkerboard. It may be that a loss of relative salience for the target on a familiar background will make it harder to detect and so reduce that advantage, but this is not the only possibility as we shall see.

The Experiment

Participants

Fifty-three undergraduates from the University of Exeter (39 female, 14 male, age range = 19-29) took part in the experiment and were given either course credit or cash for their participation. The participants were all right-handed and met the safety screening criteria for tDCS participants (approved by the Research Ethics Committee at the University of Exeter). They were allocated to either sham or anodal tDCS groups using a double-blind procedure for the neurostimulation (see Civile et al., 2018 for details). Five outlier participants were excluded before data analysis, leaving 48. The participants were also allocated to one of eight participant groups in order to counterbalance the stimuli and target locations used (see further details below).

Stimuli

The study used sets of 16 x 16 cell (each cell 16 x 16 pixels) checkerboard stimuli, containing roughly half black and half white squares, presented on a grey computer screen background on an iMac computer using Superlab 5 software. Two prototype checkerboards (A and B) were generated and, from these, exemplar checkerboards were created by changing 48 cells at random. These exemplars were used for the categorisation task at the beginning of the study. For each participant, one of the prototype checkerboards, A or B (counterbalanced across participants), was used in the target detection part of the study as the “familiar” checkerboard together with a further set of “random” checkerboards. These random checkerboards comprised of completely randomised black and white cells within the 16 x 16 cell grid (so they were unrelated to the exemplars in the categorisation part of the study). The dimensions of all the checkerboards were 256 by 256 pixels, presented at a resolution of 1680 by 1050 pixels. The participants viewed the computer screen from approximately 70 cm.



Figure 1a. Target stimulus




Figure 1b.
Example background checkerboard with target stimulus present in lower middle region. The target was presented on 50% trials at a different random location each time.

Figure 1: Diagram of the two trial types in the experiment.

The target pattern (Fig. 1a) had the dimensions 80 x 80 pixels and was a 5 x 5 cell pattern designed to be symmetrical (both horizontally and vertically) and to have similar numbers of black (13) and white (12) cells which would line up exactly with the cells in the checkerboards (Fig. 1b). It was unlikely to occur by chance in any checkerboard ($p=4.29 \times 10^{-6}$).

The position of the target pattern within each checkerboard was determined by calculating the co-ordinates (defined within the Superlab 5 programme as the number of horizontal pixels and vertical pixels

measured from the centre of each checkerboard) which would allow the target to line up exactly with the cells in each checkerboard. These co-ordinates were divided into 64 “odd” positions and 64 “even” positions for the target. Participants were counterbalanced across these target position types.

tDCS

The stimulation was delivered by a battery driven constant current stimulator (neuroConn) using a pair of surface sponge electrodes (35 cm²) applied to the scalp at the target areas for stimulation. We adopted a bilateral bipolar-non-balanced montage with one of the electrodes (anode) placed over the target stimulation area (Fp3) and the other (cathode) on the forehead over the reference area (right eyebrow). The study was conducted using a double-blind procedure reliant on the neuroConn study mode. In the anodal condition, a direct current stimulation of 1.5mA was delivered for 10 mins (5s fade-in and 5s fade out) starting as soon as the behavioral task began and continuing throughout the study. In the sham group, participants experienced the same 5s fade-in and 5s fade-out, but with the stimulation intensity of 1.5 mA delivered for just 30s, following which a small current pulse was delivered every 550ms for the remainder of the 10 minutes.

Procedure

The study consisted of a categorisation phase followed by a target detection phase. Neurostimulation commenced at the start of this categorisation phase and continued for 10 minutes. In the categorisation phase, the participants were instructed to sort the 128 exemplar checkerboards (64 exemplars from each prototype A and B) into two categories by pressing one of two keys (1 or 2 on the numerical keypad) in response to each checkerboard appearing one at a time on the screen. Participants were encouraged to scan the whole of each checkerboard before categorising it. The checkerboards were presented in random order for up to 4s or until the participant had responded. The computer gave immediate feedback as to whether their response was correct or not.

In the target detection phase, the participants were shown the target pattern (Fig 1a) and then instructed to look for this target on the checkerboards that they would subsequently see. They were told that the target could be anywhere on the checkerboard but that it would always be the same size and in the same orientation. They were instructed to press the “x” key if they thought the target was present or the “>.” key if they thought the target was absent. The participants

were then shown 256 checkerboards one at a time in random order, 128 with the target present (64 familiar and 64 random checkerboards) and 128 with the target absent (64 familiar and 64 random checkerboards). There were 8 blocks, each with 32 trials. Each checkerboard was presented on screen for 4000ms, preceded by a fixation cross presented in the centre of the screen for 500ms. The participants were given immediate feedback according to whether their response was correct or not.

Data were collected on the accuracy (our primary measure) but also on the latency of responses for the target detection phase. The accuracy data were then used to perform a signal detection analysis appropriate for this type of target detection task, giving us d' (sensitivity) and C (criterion) measures of performance. The d' measure gives us an index of how easy or hard the task is, with higher values indicative of better performance, whilst the criterion, C , is independent of the sensitivity measure and tells us something about how people are making the decision on target presence / absence.

Results

The main prediction was confirmed. Performance in the target detection task was superior with the familiar checkerboard background in Sham and Anodal tDCS groups. In both groups the difference was highly significant ($t(23) = 3.98$ and 2.94 respectively), despite performance on the task being close to ceiling (raw means above 90% in all cases). Figure 2 shows these results graphically and reports the d' for target detection separately by group and by checkerboard background. The difference in performance between the two groups is not significant for either condition (familiar or random background) or overall, and nor is there a reliable Group x Condition interaction, $F(1,46) = 1.05$, $p = ns$.




Figure 2: Graph showing d' for target detection in both groups by familiar or random checkerboard background.

The next graph (Figure 3) shows an equivalent analysis for Criterion, C , using the same signal detection approach to the data. This analysis also confirms a difference in performance involving familiar or random backgrounds in the Sham group. Here, the criterion is significantly greater for the familiar condition than in the random condition, whereas there is no significant difference (and a numerical reversal of this effect) in the Anodal tDCS group. This time there is a significant interaction between Group and Condition, $F(1,46) = 5.22$, $p < .05$, indicating that the effect of background is different in the two groups. Comparison of the criterion for Sham performance on familiar backgrounds to that for Anodal on the same backgrounds reveals a trend ($p = .064$) and no other effects are significant. There were no significant effects on latency to respond, with numerically slower times overall for the Anodal (1070ms) group compared to Sham (967ms).




Figure 3: Graph showing C (criterion) for target detection in both groups split by familiar or random checkerboard background.

General Discussion

On the face of it, we have a very straightforward result here. If we just consider the results for the sham group first, then they bear out the predictions made in the introduction. People are better at the task (larger d') when the background is the familiar checkerboard than a random checkerboard. They are good at the task in any case, but there is a highly reliable difference in favor of the familiar background condition. This in itself is a novel finding with this paradigm, albeit one that can to some extent be extrapolated from our earlier research.

An additional and novel finding is the significant difference in criterion between the two conditions. In the Sham condition, the criterion is higher for the familiar checkerboard background than in the random

case. This could be a genuine case of a shift in criterion based on recognition of the checkerboard background. But, given that this is a within-subjects design and that participants don't know what type of background they will be given until the trial actually happens, and they then typically respond in less than a second, it seems unlikely that they set their criterion dynamically from trial to trial. Hence, we suggest that this finding is more likely to reflect a shift in the distributions for the familiar and random conditions, rather than a shift in the absolute location of the criterion used by participants to make their decision. Figure 4 illustrates this idea. If the criterion's absolute position on the dimension remains fixed, but the distributions that apply to the target and non-target themselves shift on that dimension, then this will manifest in a signal detection analysis as a shift in criterion even though the criterion for decision has, in some sense, remained fixed. The key here is that the use of, say, a familiar background has an impact on the target and non-target distributions so that they shift in such a fashion as to increase their separation (hence the increase in d') and also moves their crossover point (giving the increase in C) on the dimension relative to the random background case.




Figure 4: Diagram showing hypothetical distributions for the target on a familiar background (FT), the noise distribution for a familiar background (FN), and target and noise distributions for the random background (RT and RN) in the Sham group. The bold vertical line is an example placement of the criterion used to decide if the target is present (values greater than C) or absent. No scale is given for the x-axis because the zero point depends on which set of distributions is being considered.

In Figure 4 the solid lines represent the target (blue) and noise (red) distributions for the familiar background condition. We've shown the distributions as well separated, indicative of a high d' and hence an easy task, but otherwise have made no attempt to provide a fit to the data here – they are just for

illustration. The solid black vertical line is an example of how a criterion could be set that would allow a decision on whether a target was present or absent to be made. If, on a given trial with the familiar background (solid lines), the value on which a decision is based (whatever it is – this doesn't have to be specified by the theory and can be multiply determined) is greater than C , then a “yes” response is given. Clearly if it is a target present trial, and hence the target distribution is in play, this is more likely (and a hit will be the result) than if no target is present and the noise distribution is in play such that a yes response would constitute a false alarm. The actual value of C is calculated relative to the crossover point of the two distributions which is set to 0. On this basis, C would be positive and around 0.2 for the familiar background condition if it were to correspond to our data.

Now consider what would happen if we were to leave the criterion in the same place but shifted the target and noise distributions because now we were dealing with a random background (dashed lines). One possibility is shown in the diagram, which reduces the mean difference between the two distributions (corresponding to a lower d' as observed in our data) in an asymmetrical fashion. Thus, the noise distribution shifts more than the target one, and this also corresponds to our data as it is the false alarm rate that increases most for the random backgrounds. As a result, the crossover point for the two distributions (which is used to calculate the C for these data) itself shifts, and C will now be computed as being smaller, even though in reality the absolute value of the criterion used by the participant has not changed.

This provides a plausible account of why both d' and C are different for the two conditions in the Sham group. It remains to be explained exactly why the distributions for the random checkerboards should be shifted in something like this fashion relative to the distributions for the familiar checkerboards, but we will return to that issue in due course after we have finished illustrating how shifts in distribution can produce changes in C . The next step in our analysis of the results of this experiment is to consider performance in the Anodal tDCS group. The most striking thing about it is that it is not so very different from that in Sham. There is no significant difference in d' , and no significant difference in terms of reaction times either. The only significant effect that we can point to concerns C , and it relates to the difference in calculated C for the Familiar and Random conditions. This difference is significant in the Sham group (with

a larger value in the familiar condition) and reverses for the Anodal group (not significantly) to give a significant interaction between Group and Condition. How are we to interpret this?

Before tackling this question we first need to address the failure to find a significant difference in d' in this experiment. Our logic in running this experiment was that tDCS would have an impact on the familiar background task that was different to than on the random background version. This would predict that the change in salience modulation would affect the familiar background (features with low error) differently to the random backgrounds (features with high error), and we stand by this prediction from the MKM theory. But, of course, these effects should apply to both target and non-target trials, and so a change in d' is unlikely on this basis. A shift in location of both distributions that preserves d' , however, would be predicted by this analysis, which is why we are pursuing that possibility here. Note that our explanation for target detection being better on the familiar background was based on the target “standing out” on that background by virtue of being novel. But it can still do this in the tDCS condition, the difference in novelty has not disappeared, but its expression in terms of salience may well have altered. But that doesn’t change the fact that the target will be the “new” set of features in that location on a familiar background, and just another set of features on a random background, assuming a location specific coding. Because of this any effect on target detection sensitivity is moot. Now we turn to a possible explanation for the effects of tDCS on C .

It must be admitted at the outset that it is quite possible for there to be a real difference in criterion setting contingent on neurostimulation. We have no prior evidence that points to this, but certainly cannot rule it out as a possibility. The reversal of the numerical difference, however, is intriguing if we view it from a distributional perspective in the way that we have done for the sham data. Figure 5 illustrates this idea. Having set some value of C for the Familiar condition, the implication now is that the reduction in d' is achieved by shifting the two distributions in the opposite manner to that employed for the Sham group. This brings the two distributions closer together, but moves their crossover point to further away from C , thus increasing its calculated value.

It would seem that the interpretation of the effects of tDCS that this naturally implies is for it to make “yes” decisions to target trials on random checkerboard backgrounds harder. The question is why should this

be so? Our earlier analysis of why performance in detecting the target on familiar backgrounds might well be superior to that on random checkerboard backgrounds seemed to imply that losing salience modulation would change this effect. We now need to give this more careful consideration. The idea was that the target would be relatively salient on a less salient background in the Familiar condition. This would then naturally facilitate finding and identifying the target. We see no reason to alter this explanation, it does, after all, explain the basic phenomenon very well. But note that it is the relative difference in salience that is the basis for locating and identifying the target. This difference is not available in the Random condition, hence the advantage for the familiar background.




Figure 5: Diagram showing hypothetical distributions for the target on a familiar background (FT), the noise distribution for a familiar background (FN), and target and noise distributions for the random background (RT and RN) in the Anodal tDCS group. The bold vertical line is an example placement of the criterion used to decide if the target is present (values greater than C) or absent.

But when we apply anodal tDCS, according to the MKM model we lose the greater salience of the target relative to the background in the Familiar condition, but only to replace it with lower relative salience for the target compared to the background. The salience difference is still there, and still absent in the Random condition. It may be this that leads to the roughly equal levels of performance, and the continued superiority of the Familiar condition over Random. It would seem that all that is needed is some difference to provide a signal to guide search to confer this advantage. There is convergent evidence for this idea in the literature. For example, Lubow, Rifkin and Alek (1976) were able to show that enhanced learning resulted from a contrast in familiarity between context and discrete stimuli in both people and rats. If the stimuli were familiar, but the context novel, or the stimuli were

novel and the context familiar, then learning was potentiated. If we extrapolate that this was due to the participants' attention being more easily drawn to the target stimuli for learning, then this is essentially what we are saying here.

If we take the view that it is the difference in salience between target and background that is crucial in our experiment, then this also offers a possible explanation of the effect on both d' and C . The difference in salience in the Familiar condition explains the increase in d' . The direction of that difference potentially explains why the criterion is larger for the Familiar condition than the Random condition in the Sham group (target salience higher than background in the Familiar condition) but smaller for the Familiar condition than for the Random condition in the Anodal group (target salience lower than background in the Familiar condition). Hence, if we postulate that this difference in salience is what modulates both the separation between signal and noise distributions, and whether those distributions are shifted to the right or left of the Random condition distributions, then we have an explanation that covers the facts.

Clearly further research will be needed to verify this account or falsify it. We are already working on a conceptual replication of this result as well as a straightforward replication of the novel behavioral effect of background familiarity on C and hope to have the answers soon. Another question of great interest is whether a sufficiently detailed simulation of this task using the MKM model with appropriate representational assumptions and incorporating suitable decision mechanisms would generate this pattern of results. The difficulty here lies in specifying what is "appropriate" and "suitable". Nevertheless, we are optimistic that a simulation effort would be informative and help us understand whether or not the MKM model as applied to our recent work concerning the effects of tDCS on perceptual learning will also provide a good account of these results.

References

- Civile, C., McLaren, R., & McLaren, I.P.L. (2018). How we can change your mind: Anodal tDCS to Fp3 alters human stimulus representation and learning. *Neuropsychologia*, 119, 241-246.
- Civile, C., Obhi, S.S., & McLaren, I.P.L. (2019). The role of experience-based perceptual learning in the Face Inversion Effect. *Vision Research*, 157, 84-88.
- Civile, C., Verbruggen, F., McLaren, R., Zhao, D., Ku, Y., & McLaren, I.P.L. (2016). Switching off perceptual learning: Anodal transcranial direct current stimulation (tDCS) at Fp3 eliminates perceptual learning in humans. *Journal of Experimental Psychology: Animal Learning and Cognition*, 42, 290-296.
- Civile, C., Zhao, D., Ku, Y., Elchlepp, H., Lavric, A., & McLaren, I.P.L. (2014). Perceptual learning and inversion effect: Recognition of prototype-defined familiar checkerboards. *Journal of Experimental Psychology: Animal Learning and Cognition*, 40, 144-161.
- Lubow, R.E., Rifkin, B. and Alek, M. (1976). The Context Effect: The Relationship Between Stimulus Preexposure and Environmental Preexposure Determines Subsequent Learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 2, 38-47.
- McLaren, I.P.L. (1997). Categorization and perceptual learning: An analogue of the face inversion effect. *The Quarterly Journal of Experimental Psychology* 50, 257-273.
- McLaren, I.P.L., & Civile, C. (2011). Perceptual learning for a familiar category under inversion: An analogue of face inversion? In L. Carlson, C. Hoelscher, & T.F. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, (pp. 3320-25). Austin, TX: Cognitive Science Society.
- McLaren, I. P. L., Kaye, H. & Mackintosh, N. J. 1989. An associative theory of the representation of stimuli: applications to perceptual learning and latent inhibition. In: *Parallel Distributed Processing. Implications for Psychology and Neurobiology* (Ed. by R. G. M. Morris), pp. 102-130. Oxford: Clarendon Press.
- McLaren, I.P.L. & Mackintosh, N.J. (2000). An elemental model of associative learning: Latent inhibition and perceptual learning. *Animal Learning and Behavior*, 38, 211-246.
- Phillips, W.A. (1974) On the distinction between sensory storage and short-term visual memory. *Perception and Psychophysics*, 16, 283-290
- Pietrewicz, A. T., & Kamil, A. C. (1979). Search image formation in the blue jay (*Cyanicittacristala*). *Science*, 204, 1332-1333.
- Plaisted, K. C., & Mackintosh, N. J. (1995). Visual search for cryptic stimuli in pigeons: Implications for the search image and search rate hypotheses. *Animal Behaviour*, 50, 1219-1232.
- Rensink, R.A., O'Regan, J.K. & Clark J.J. (1997) To See Or Not To See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, 8, 368-373
- Suret, M. and McLaren, I.P.L. (2001). An Analogue of the Phillips Effect. In Johanna D. Moore and Keith Stenning (Eds.). *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*. Mahawah, New Jersey: LEA 1006-1011.
- Tinbergen, L. (1960). The natural control of insects in pinewoods. Factors influencing the intensity of predation by songbirds. *Archives Neerlandaises de Zoologie*, 13, 265-343.