The effect of visualization on visual search performance

Does visualisation trump vision?

Alasdair D. F. Clarke, Courtney Barr
and Amelia R. Hunt

Abstract Striking results recently demonstrated that visualising search for a target can facilitate visual search for that target on subsequent trials [Reinhart et al., 2015]. This visualisation benefit was even greater than the benefit of actually repeating search for the target. We registered a close replication and generalisation of the original experiment. Our results show clear benefits of repeatedly searching for the same target, but we found no benefit associated with visualisation. The difficulty of the search task and the ability to monitor compliance with instructions to visualise are both possible explanations for the failure to replicate, and both should be carefully considered in future research exploring this interesting phenomenon.

Keywords mental imagery · visual attention, · learning · visual search · perception · open materials

1 Introduction

Mental imagery is the term used to describe the representation of, or process of representing, an experience or activity in the absence of the represented sensory stimuli or behaviour. The processing mechanisms and experience of mental imagery of perceptual events can strongly resemble the experience of perceiving the actual external event. Indeed, many studies suggest mental imagery of perception is simply a weak form of perception (for a recent review,

F. Author
first address
Tel.: +123-45-678910
Fax: +123-45-678910
E-mail: fauthor@example.com

S. Author
second address

see Pearson et al. [2015]). Consistent with this idea, imagining visual information (visualisation) can lead to classical conditioning [Lewis et al., 2013] and perceptual learning [Tartaglia et al., 2009]. This general framework dovetails with promising work suggesting that mental imagery can be a useful tool in clinical settings [e.g. Foa et al., 1980] and in musical and athletic training [e.g. Zatorre et al., 2007, Guillot and Collet, 2008].

An interesting recent study by Reinhart et al. [2015] demonstrated that attentional mechanisms can be effectively trained to select target objects through visualisation. Moreover, they found that visualising search for a particular target in an imagined array can speed reaction time to find the target on a subsequent trial even more than actually searching for the target in an array presented on the screen. In their experiments, participants were first shown a green C oriented in a particular direction, and then shown an array of 12 C's organized in a circle, only one of which was green. Participants had to press a key to indicate whether the one green C in the array was oriented in the same or a different direction as the C they had just been shown. Reaction times become faster when the C of a particular orientation repeated over a run of up to 7 trials. On some runs of trials participants were asked to "visualise search" and were shown the oriented C but not the array for the first two trials of the run. On trial 3 of the run, reaction time to find the target was significantly faster – faster not only compared to the first trial of the run, but also faster when compared to trial 3 of runs during which the search array had actually been shown on trials 1 and 2. The authors also measured EEG during the experiment and found that the N2PC (an early component of stimulus-evoked EEG associated with attentional focus) followed a similar pattern as reaction time, with an enhanced N2PC to the third target in a run following two visualisation trials – again, relative not only to the first trial in the run, but also to the third target following trials in which the search array was physically present rather than just visualised. The results suggest, as the title of the article states, that "visalization trumps vision in training attention."

On the one hand, this result is consistent with previous work reviewed above demonstrating that visualisation has meaningful effects on subsequent perception and behaviour. On the other, this result is unique in demonstrating that the effects of visualisation can surpass those of actual perception – on its face, this is inconsistent with the idea that mental imagery is "weak" perception. One possible reason for the superiority of visualisation over perception, however, is that the perceptual learning over a run of trials when the target and distractors are displayed in fact has components of both interference and facilitation on the training of attention, and visualising removes some of the interference but maintains the facilitation. For example, the participant may may not imagine the distractors very vividly, and this could strengthen the representation of the target orientation, speeding reaction time to match it green C in the array on trial 3. Consistent with this interpretation, when the black distractor C's were removed from the display (Experiment 4), the difference between visualization and practice with actual stimuli was eliminated. This could reflect a floor effect, however, as reaction times in general in this

experiment were much faster than in the experiments with distractors. It is also the case that the target can be rapidly selected on the basis of its color in this experiment, so processing of distractors would be expected to be minimal even when they are presented. If visualization was facilitated relative to search because the interference from distractors was minimised, even more facilitation for visualisation should be observed when search is difficult and the distractors during "real" search need to be processed in more detail.

Definitively establishing why visualisation is more effective than actually performing the task is of critical importance not only for understanding the nature of visual imagery, but also because it could lead to improved imagery-based training or therapy regimes like those noted above. If it is the case that visualisation is more effective because it isolates the facilitative components of the task, for example, vague or incomplete representations may in fact be superior to less vivid or detailed ones. We therefore plan to first replicate the original study by Reinhart et al. [2015], focusing specifically on the behavioural results. Omitting the recording of EEG allows us to use a slightly simplified stimuli design to test the original results, without the second irrelevant coloured distractor (only included in the original study to allow for comparison of the lateralised N2PC component between the two hemifields). This will provide independent confirmation of this interesting result. We will also analyse the data in more depth using linear-mixed-effect models and explore how the visualisation effect changes over time and is affected by to serial dependencies [Fischer and Whitney, 2014]. Provided we replicate the original findings, we also plan to run a second experiment using a more traditional visual search paradigm using an array of C's of uniform color, in which the observer's task is to determine whether the target (a C of the cued orientation) is present or absent. If the superiority of visualisation over practice at search is due to proactive interference from the distractors when the search array is presented, there should be even larger benefits of visualisation when serial search of all the items is required.

2 Experiment

The aim of this experiment is to directly replicate the behavioural effect of visualisation in visual search [Reinhart et al., 2015]. We have made two minor changes to the original paradigm: (i) we have removed the red distracter element used in the original study (included to measure the N2PC component of the EEG; we will not measure EEG). (ii). We will use runs of length three, four and five, rather than three, five and seven. This was done because 1) the practice effects asymptotes by trial 5 of the runs, and 2) runs of longer length are unimportant for testing the critical result: that reaction times for the first three trials in a run will be faster in the visualise condition than those in the practise condition.
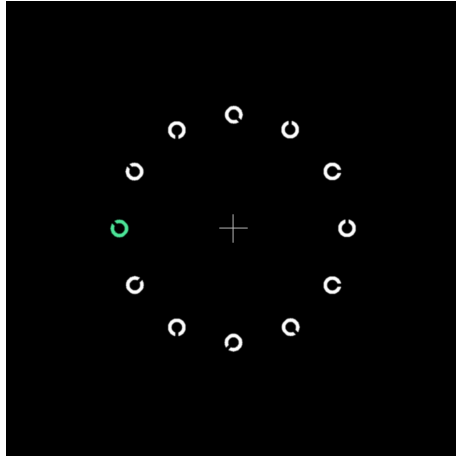
Fig. 1 Example stimulus from Experiment 1. Before the stimulus is shown, the observer is shown a green Landolt C. Their task is to decide if the green C in the stimulus matches the orientation of the cued C.

2.1 Methods

2.1.1 Participants

Thirty participants will be recruited via from amongst the student population at the University of Aberdeen. All participants will have normal or corrected-to-normal vision.

The sample size of $n = 30$ is based on a power analysis for a one-tailed paired t-test, with a power of 0.9 and effect size of $d = 0.55$. This should be sufficient for the replication given the original effect sizes of $0.61 < d < 0.72$ ($n = 18$).

2.1.2 Stimuli

The search stimuli consisted of twelve Landolt Cs arranged in a circle (radius 8°) around a central fixation cross. The Landolt C's had a radius of 0.62° (with thickness = 0.25°, gap width= 0.19°) and had one of eight possible orientations ($\phi = n\frac{\pi}{4}$, $n = 0, \ldots, 7$). One of the C's was coloured green which indicated that it was the target. An example stimulus is shown in Fig. 1.

2.1.3 Procedure

The first block of trials was pre-empted by a set of ten practice trials. The practice trials included the visual search condition only. There were then four blocks of 30 trials (15 for each condition). Each trial began with the presentation of a fixation cross (1200-1600ms) followed by the presentation of a green Landolt C (the cue stimuli) for 100ms, followed by a interval of 1000ms. The

search array was then presented in the visual search condition for a maximum of 2000ms. In the visualise search condition the cue stimuli was followed by instructions to 'visualise search' before the presentation of a further fixation cross. The orientation of the green C, the location of the green C within the search array and whether the green C matched the cue was randomised for each trial.

2.2 Planned Analyses

We plan to carry out the analysis in several different ways. Firstly, we will repeat the analysis from Reinhart et al. [2015] in which the median[1] reaction time for trial 1,2 and 3 within a (normal) run is compared to the median reaction time for trial 3 in a visualise-run. A one-tailed paired t-test will be used for this comparison. We will include a Bayesian t-test[2] which will enable comparison of the expected difference between visualisation and trials 1-3 in the run to the null hypothesis (no difference). We will also directly compare/combine our results with those of Reinhart et al. [2015], who have already provided us with the summary data from their Experiment 1.

Additionally, we will also analyse our data in more detail using a linear mixed-effect model (`lme4` [Bates et al., 2015, R Core Team, 2015]), following the guidelines on model design given by Barr et al. [2013] This allows us to include trial-to-trial variation in the analysis rather than using aggregate statistics. As the distribution of reaction times is expected to be skewed, we use log reaction times in the analysis. We will treat *trial number* (within run) as a numerical factor, and will investigate non-linear regression if required (e.g. if there is a clear asymptote). However, given the results presented by Reinhart et al. [2015], a linear model should suffice. We will follow the model simplification procedure put forward by [Crawley, 2012, chapter 9]. p-values will be obtained via the `Anova` function from the `car` package [Fox, 2011].

Finally, we will also compare the visualisation effect to the effect of serial dependency. More specifically, we will investigate how reaction time varies on trial three depending on what preceded it (present-present, absent-absent, present-absent, absent-present, visualise-visualise). This will be modelled using another linear mixed-effect model with a five-level factor describing the previous two trials, and a two-level factor for whether trial three is a target absent or present trial[3].

---

[1] This is a mistake made in our pre-registered submission. Reinhart et al. [2015] used mean reaction time rather than median, and we follow suit.

[2] `http://www.lifesci.sussex.ac.uk/home/Zoltan_Dienes/inference/bayes_normalposterior.swf`

[3] Due to the failure to replicate the original findings, we did not carry out this part of the planned analysis
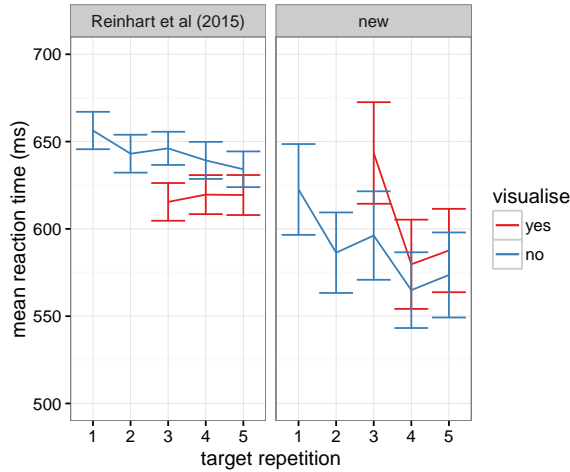
Fig. 2 Results over all participants. The mean reaction time was first computed for each participant, and then the mean of the means and associated standard error were plotted. Note: in order to facilitate comparison with Reinhart et al. [2015], the error bars indicate +/-1 standard error, rather than 95% confidence intervals.

Table 1 *t*-test (*p*-value) results. All *t* tests are paired and one-tailed (is the second value less than the first).

		difference	*t*-value	*p*-value
Stimulus 1	Stimulus 2	36 ms	5.78	< 0.001
Stimulus 1	Stimulus 3	26 ms	4.25	< 0.001
Stimulus 1	Visualise 3	-21 ms	-1.83	*n.s.*
Stimulus 2	Visualise 3	-57 ms	-4.80	*n.s.*
Stimulus 3	Visualise 3	-47 ms	-3.50	*n.s.*

2.3 Results and Discussion

All participants were accurate at the task $(88\% - 99\%)$ and incorrect trials were removed from further analysis. The reaction times from our study are presented together with those of Reinhart et al. [2015] in Figure 2 and the *t*-test based analysis is given in Table 1. The first two lines of this table compare performance across successive trials of search for the same target and demonstrate that we have a repetition benefit of around 30 ms on each trial where the target repeats from the last trial. The remaining lines of the table repeat the t-tests used by Reinhart et al. [2015] to compare the effect of visualising search to actually performing search. Participants were slower on trial 3 after visualization than compared to all three trials within a run of actual search. Because we planned to do a one-tailed test, and the direction of these effects is the opposite of what was predicted, they are deemed non-significant.

The results from each of the 30 participants are shown separately in Figure 3. As the reaction times give a skewed distribution, we calculated the mean and 95% confidence intervals after log-transforming the data. We fur-

Table 2 Linear mixed effects model fit. Reaction time was log transformed. All effects were statistically significant under a Type II Wald χ^2 test.

Effect	β	std. error	t-value	p-value
(intercept)	6.38	0.04	163.64	-
trial number	-0.02	0.003	-8.11	< 0.001
visualise	0.17	0.04	3.92	0.007
interaction	-0.03	0.01	-3.39	< 0.001

ther analysed these data with a linear mixed effects model. Reaction times were log-transformed, trial number was treated as a continuous variable, and we allowed for a maximal (crossed) random effects structure. All effects were statistically significant ($p < 0.05$) and are shown in Table 2. The negative effect of trial number confirms that we have a practise effect: as with Reinhart et al. [2015], reaction times decrease over the course of a run of trials. However, as visualisation has a positive effect, we conclude that participants are slower in the visualisation condition. The interaction is also statistically significant, with reaction times decreasing faster for trials after the visualisation condition. Upon inspection of Figure 3, it appears to be primarily driven by a few participants with relatively slow reaction times on Trial 3 in the visualisation condition, followed by an improvement for Trials 4 and 5. This pattern is consistent with task-switching costs (e.g. Rogers and Monsell, 1996).

Finally, we use Bayesian methods [Dienes, 2008] to compute the posterior distribution of effect, taking the previously published findings from Reinhart et al. [2015] as the prior. As can be seen in Figure 4, when we update our beliefs based on the new data presented here, we end up with a posterior distribution centred around zero and conclude that there is no evidence of an effect in either direction.

3 General Discussion

We find no evidence that "visualizing trumps vision in training attention." Indeed, mean reaction times to find a target are actually slower when performing search following visualisation compared to an equivalent amount of practice actually searching for that target. We offer two explanations for the discrepancy between our findings and the original results.

First, the most salient difference between our paradigm and the one presented by Reinhart et al. [2015] is the additional task-irrelevant colour singleton, included in the original design to rule out physical-stimulus confounds in the lateralised ERP component of interest (the N2PC). This additional coloured item in the array seemed like an unimportant detail when we planned our (purely behavioural) replication, so we did not include it. However, this small modification may have made our search task slightly easier than theirs. Reaction times in our version of the experiment were indeed slightly faster than in the original, although it is important to note that it is not the case that participants who responded relatively slowly in our experiment were more
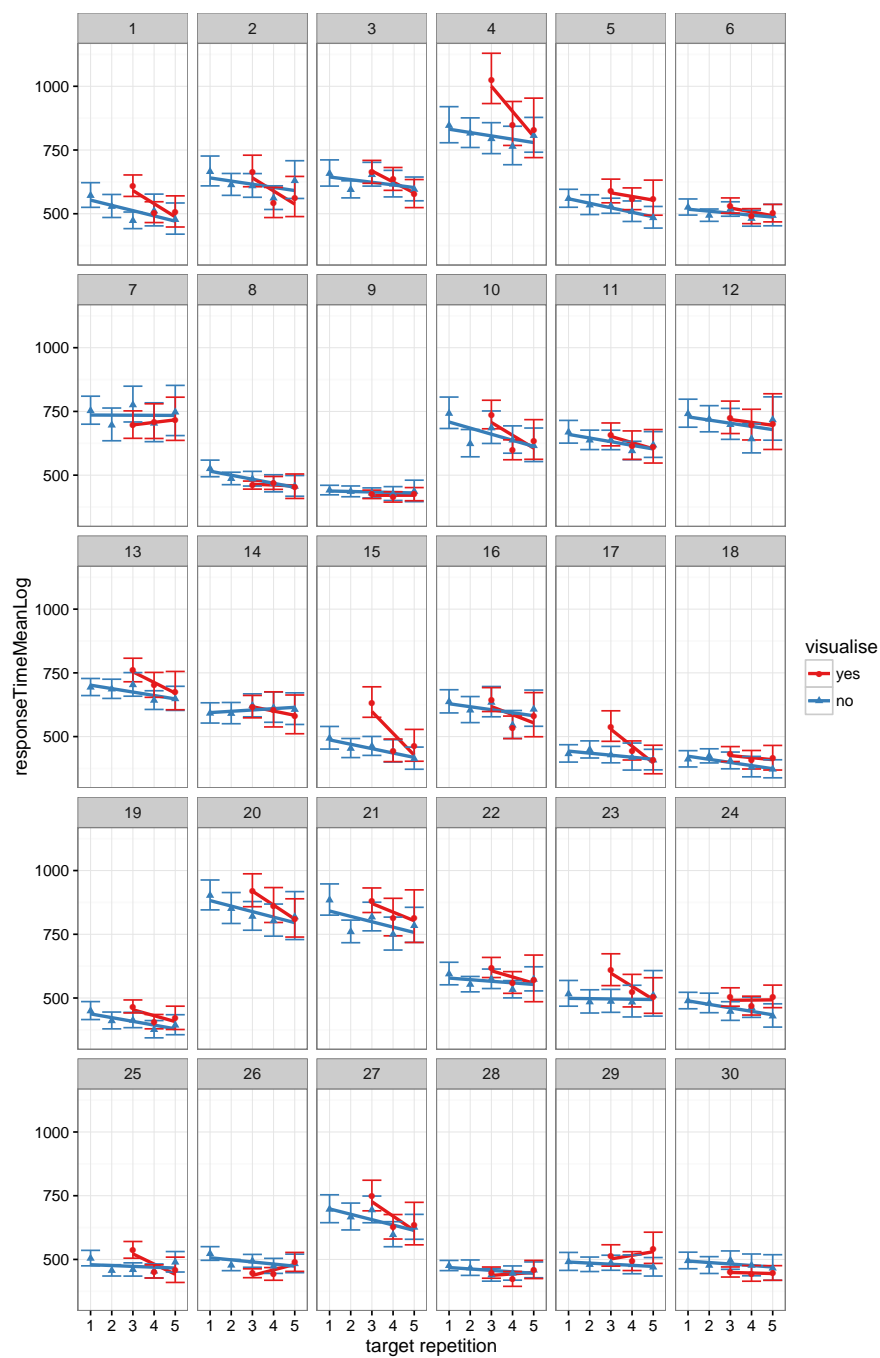
Fig. 3 Results from each participant. Means and 95% confidence intervals are computed after a log transform. The results have been transformed (using the exponential function) back into the original units for ease of reading the plots and comparing to the original data.
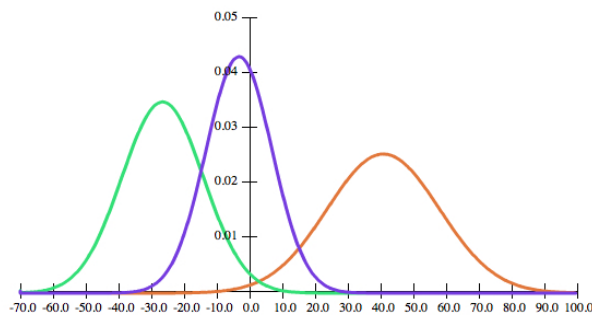
Fig. 4 The purple graph shows the posterior distribution of the difference in mean reaction times between Trial 1 in the stimulus condition and Trial 3 in the visualise condition. The red graph shows the prior, taken from Reinhart et al. [2015], while the green line shows the likelihood of our data.

likely to show an effect in the original direction (quite the opposite in fact, see Figure 3). Obviating the need to search, as they did in their Experiment 4, eliminated the benefit of visualisation. It may be the case that removing one salient distractor, as we did in our study, has a similar impact. If so, this is an important detail to note for any researcher interested in trying to study the effect of visualisation on search: the search may need to be of a minimum level of difficulty for the effect of visualisation to be observed.

A second important difference is that we did not measure EEG. While the measurement of EEG itself does not influence behaviour, perhaps participants wearing an EEG cap believed that their compliance with instructions to "visualise" could actually be verified by the experimenter, and hence were more likely to follow that instruction. Effects along these lines have been found in the eye-tracking literature [Nasiopoulos et al., 2015] in which participants modify their viewing behaviour when they know their eyes are being tracked. Instructions to visualise did have an impact on reaction time in our study (albeit opposite to what was found previously) so there was some evidence that our participants were attempting to comply with instructions. However, they may not have been as consistent or effortful in their imagery compared to a group of participants wearing a device that they have been told measures their brain activity. If this is the case, it would be important for researchers to be aware that a putative ability to monitor compliance with visualisation instructions is a key requirement for finding the effect. For the record, the participants in our study were instructed as follows: "When prompted to visualise search you are required to generate a mental image of the search array in your mind and imagine searching through this array for the green cue C you have been shown in the trial." A precise description of how researchers instruct and/or verify visual imagery in their participants should be made available in future studies on this topic so that the importance of this factor can be determined.

By failing to replicate the original effect, it is our hope that this report can provide important guidance to future researchers who may wish to build on the

findings of Reinhart et al. [2015]. To be useful in clinical or training settings, as suggested by the authors, the imagery effect must be robust and the boundary conditions need to be clearly and consistently documented. Moreover, for the potential theoretical implications of their results to be fully developed, we require a thorough understanding of the conditions under which it does, and does not, occur.

References

Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3):255–278, 2013.

Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015. doi: 10.18637/jss.v067.i01.

Michael J Crawley. *The R book*. John Wiley & Sons, 2012.

Zoltan Dienes. *Understanding psychology as a science: An introduction to scientific and statistical inference*. Palgrave Macmillan, 2008.

Jason Fischer and David Whitney. Serial dependence in visual perception. *Nature neuroscience*, 17(5):738–743, 2014.

Edna B Foa, Gail Steketee, Ralph M Turner, and Steven C Fischer. Effects of imaginal exposure to feared disasters in obsessive-compulsive checkers. *Behaviour Research and Therapy*, 18(5):449–455, 1980.

John Fox. *An R and S-Plus companion to applied regression*. Sage, 2011.

Aymeric Guillot and Christian Collet. Construction of the motor imagery integrative model in sport: a review and theoretical investigation of motor imagery use. *International Review of Sport and Exercise Psychology*, 1(1): 31–44, 2008.

David E Lewis, Matthew J OReilly, Sieu K Khuu, and Joel Pearson. Conditioning the minds eye associative learning with voluntary mental imagery. *Clinical Psychological Science*, page 2167702613484716, 2013.

Eleni Nasiopoulos, Evan F. Risko, Tom Foulsham, and Alan Kingstone. Wearable computing: Will it make people prosocial? *British Journal of Psychology*, 106(2):209–216, 2015. ISSN 2044-8295. doi: 10.1111/bjop.12080. URL `http://dx.doi.org/10.1111/bjop.12080`.

Joel Pearson, Thomas Naselaris, Emily A. Holmes, and Stephen M. Kosslyn. Mental imagery: Functional mechanisms and clinical applications. *Trends in Cognitive Sciences*, 19(10):590 – 602, 2015. ISSN 1364-6613. doi: http://dx.doi.org/10.1016/j.tics.2015.08.003.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2015. URL `https://www.R-project.org`.

Robert MG Reinhart, Laura J McClenahan, and Geoffrey F Woodman. Visualizing trumps vision in training attention. *Psychological science*, pages 1114–1122, 2015.

Elisa M Tartaglia, Laura Bamert, Fred W Mast, and Michael H Herzog. Human perceptual learning by mental imagery. *Current Biology*, 19(24):2081–2085, 2009.

Robert J Zatorre, Joyce L Chen, and Virginia B Penhune. When the brain plays music: auditory–motor interactions in music perception and production. *Nature Reviews Neuroscience*, 8(7):547–558, 2007.