Cognition 174 (2018) 19–27

Contents lists available at ScienceDirect

Cognition

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

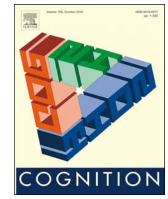

Adding statistical regularity results in a global slowdown in visual search

Anna Vaskevich[a,*], Roy Luria[a,b]

[a] School of Psychological Sciences, Tel Aviv University, Tel Aviv 6997801, Israel
[b] Sagol School of Neuroscience, Tel Aviv University, Tel Aviv 6997801, Israel

ARTICLE INFO

Keywords:
Statistical learning
Implicit learning
Practice
Visual search
Predictive value

ABSTRACT

Current statistical learning theories predict that embedding implicit regularities within a task should further improve online performance, beyond general practice. We challenged this assumption by contrasting performance in a visual search task containing either a consistent-mapping (regularity) condition, a random-mapping condition, or both conditions, mixed. Surprisingly, performance in a random visual search, without any regularity, was better than performance in a mixed design search that contained a beneficial regularity. This result was replicated using different stimuli and different regularities, suggesting that mixing consistent and random conditions leads to an overall slowing down of performance. Relying on the predictive-processing framework, we suggest that this global detrimental effect depends on the validity of the regularity: when its predictive value is low, as it is in the case of a mixed design, reliance on *all* prior information is reduced, resulting in a general slowdown. Our results suggest that our cognitive system does not maximize speed, but rather continues to gather and implement statistical information at the expense of a possible slowdown in performance.

1. Introduction

Since the pioneering work of Reber (1967), a large body of evidence has accumulated regarding our ability to pick up regularities from the environment. This evidence comes from primarily two fields- statistical learning and implicit learning. Although the two fields rely on different learning paradigms, the general objective is similar: to explore our ability to extract and use regularities from the environment (for a review see, Perruchet & Pacton, 2006). Most of the studies conducted so far have focused on the following questions: What types of regularities can be acquired (Cohen, Ivry, & Keele, 1990; Fiser & Aslin, 2002; Pothos, 2007), under which conditions regularities are identified (Turk-Browne, Jungé, & Scholl, 2005), and how implicit is this learning process (Bertels, Franco, & Destrebecqz, 2012). The results suggest that both visual and auditory regularities can be acquired (Frost, Armstrong, Siegelman, & Christiansen, 2015), and that the extraction and use of these regularities can occur incidentally and implicitly, so that the observer is unaware of the learning (Bertels et al., 2012; Buchner & Wippich, 1998).

One important aspect of implicit learning of regularities is how it affects performance. Here, it is important to separate effects regularity may have on performance in the task in which it is embedded (i.e., ongoing performance) from effects regularity may have on performance in subsequent tasks. When it comes to the impact regularity has on subsequent tasks, both facilitating and interfering effects have been demonstrated. For instance, Otsuka and Saiki (2016) showed that objects that were previously encountered in structured sequences were remembered better than objects from random sequences, while distractor items that were inserted into random sequences were remembered better than those inserted into structured sequences. Regarding ongoing performance, interference effects originating from an additional regularity that was not beneficial to the task, were found in a task that required summary statistics (Zhao, Ngo, McKendrick, & Turk-Browne, 2011; Zhao & Yu, 2016). This type of interference may be caused by a competition between two statistical operations – statistical learning and summary statistics (Zhao et al., 2011).

In the present work, we focused on the impact regularity may have on online performance from a different perspective: we tested the impact of a single, potentially beneficial regularity on general practice effects. We define practice effects as faster Reaction Times (RTs) and/or higher accuracies that are a result of repeatedly performing the task, without the presence of any regularity. For instance, in sequence learning tasks (Cohen et al., 1990), or visual search tasks (Clark, Appelbaum, van den Berg, Mitroff, & Woldorff, 2015), participants become faster as the task progresses without the presence of regularity. Current statistical learning theories suggest that introducing regularity that is relevant to the task should result in even better performance as the regularity contains additional beneficial information (Goujon, Didierjean, & Thorpe, 2015; Perruchet & Pacton, 2006).

While this assumption seems to be obvious, we argue that the

* Corresponding author.
E-mail address: annav@mail.tau.ac.il (A. Vaskevich).

https://doi.org/10.1016/j.cognition.2018.01.010
Received 5 July 2017; Received in revised form 22 January 2018; Accepted 23 January 2018
Available online 03 February 2018
0010-0277/

relationship between general practice and learning with regularity was never properly assessed. This is because current statistical learning paradigms that assess online performance lack the necessary baseline conditions reflecting only practice effects (Chun & Jiang, 1998). Such conditions are necessary in order to evaluate separately performance in a task with and without regularity. To overcome these limitations, we rely on a task in which practice performance can be assessed both in the presence and absence of regularity – visual search (Wolfe, 1998). We now turn to describe this paradigm and how performance was measured in the present work.

In a visual search task, participants search for a predefined target stimulus among distractors and are asked to respond as fast as possible (Wolfe, 1998). Typically, participants become faster to find the target as the task progresses (Sireteanu & Rettenbach, 1995). A recent ERP study suggested that this general task improvement (i.e., practice) is a result of modulations of several processing stages that involve early attentional processes, target discrimination processes and response selection (Clark et al., 2015). Together, these results show that while executing the task, there is an ongoing updating of task-set parameters that leads to a gradual improvement in performance. We consider the gradual improvement in performance in such tasks to result from practice alone because it is achieved under conditions that do not involve any regularity.

The effect of embedding regularities in a visual search task is investigated by mixing a consistent mapping condition with a random mapping condition (Chun, 2000; Chun & Jiang, 1998). In the consistent mapping condition, the target is embedded in an invariant configuration that is repeated across the experiment, while in the random mapping condition the target appears in a novel or unrepeated configuration. Participants are faster to find the target in the consistent mapping condition than in the random mapping condition, an effect termed "contextual cueing" (Chun & Jiang, 1998). Because this effect occurs without instructions, without an intention to learn, and without evidence of conscious memory, it is thought to result from implicit learning (Chun & Jiang, 2003; but see also Vadillo, Konstantinidis, & Shanks, 2016). To date, the contextual cueing effect has been replicated numerous times (for a recent review see Goujon et al., 2015). A widely accepted interpretation of this effect is that the repeating context in the consistent mapping condition is learned implicitly and serves as a cue that guides attention to the target (Chun, 2000; Harris & Remington, 2017; Peterson & Kramer, 2001; but see also, Kunar, Flusberg, Horowitz, & Wolfe, 2007; Schankin & Schubö, 2009).

Given this interpretation of the contextual cueing effect, it is reasonable to assume that the implicit learning of regularities operates in accordance with practice effects. However, we argue that such a conclusion is premature, because within this paradigm the consistent and random conditions are mixed. As such, performance in the random condition cannot be regarded as a pure baseline that reflects only practice effects. Instead, performance in this condition is driven by general practice in a task that also contains regularity (i.e., affected by the consistent mapping condition). Similarly, the consistent mapping condition reflects a situation in which the regularity is diluted by random trials. In order to separate and compare practice performance in a task with and without regularity, which is the aim of the present work, it is essential that performance in the mixed design be contrasted with additional baseline conditions that reflect general improvement alone (i.e., no regularity in the task), and improvement in a task with regularity that is valid on all trials. Employing the above described baseline conditions should allow us to determine whether the presence of regularity drives performance beyond practice.

Importantly, adding these baseline conditions will also allow us to evaluate the influence of regularity on online performance from a predictive value perspective. We define predictive value as the extent to which incoming information is consistent with the system's expectations. In the decision-making domain, it is well established that multiple sources of information, such as the history of items and the current visual input are reconciled according to their predictive values and influence the decision on a given trial (Behrens, Woolrich, Walton, & Rushworth, 2007). The impact of validity has also been demonstrated in statistical learning (Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014): items that were first encountered in a specific context, which was then changed, were more likely to be forgotten than items that appeared in an unchanged context. Presumably, when a previously experienced context is reencountered, a prediction about which item should appear in that context is automatically generated. If this information proves to be invalid, and the expected item does not appear, the representation of the item in memory may become vulnerable (Kim et al., 2014).

The above described study shows how validity of statistical information may affect a consequent memory task. However, according to the predictive value framework, the reliability of information changes during the acquisition task itself, so that the values of the incoming information is constantly reassessed and updated (Behrens et al., 2007). The recently proposed predictive-processing framework, argues that similar processes operate during perception (Clark, 2013; Lupyan, 2015; Lupyan & Clark, 2015). Within this framework the cognitive system is viewed as a probabilistic-prediction system that is continuously estimating and re-estimating its own sensory uncertainty, assigning differential weights to the systems' expectations (i.e., previous experience) versus the current inputs. In other words, the influence of what the system "knows" changes according to the reliability of the incoming information. This adaptive process is described as 'variable precision weighting': a mechanism for tuning the extent to which input is modulated by top-down predictions (A. Clark, 2013; Lupyan, 2015; Lupyan & Clark, 2015). Thus, according to the predictive-processing framework, a task that contains information with low predictive value should lead to a general slowdown.

Relying on the predictive-processing framework, and previous results from statistical learning (Kim et al., 2014), we argue that when regularity is present in the task, the predictive value of the regularity may be crucial, because it determines the extent to which all prior information is taken into consideration. When the regularity applies to all trials in the task the predictive value is high. Counter intuitively, when the task contains no regularity, the predictive value is also high (i.e., no regularity is expected). When the regularity applies to half of the trials, its predictive value is relatively low. Thus, mixing consistent mapping with random trials should result in interference and in slower responses because the incoming information is valid only on 50% of the trials.

2. Experiment 1

In the present study, three groups of participants completed a visual search task that was either random (i.e., without any regularity), completely structured, in which case the regularity was valid on every trial (i.e., consistent mapping), or a visual search with consistent and random mapping conditions mixed.

The key aspect of performance that was assessed is the end-of-session performance. This measurement represents the best performance (i.e., fastest RTs) that is achieved in a given session, and is ideal for our current purpose because it reflects both practice effect and the size of the contextual cueing effect. Several previous studies have successfully used end-of-session performance to estimate the size of the contextual cueing effect across conditions (Chun & Jiang, 1998; Kunar, Flusberg, & Wolfe, 2006; Kunar, Flusberg, & Wolfe, 2008) and experiments (Chun & Jiang, 1998).

Current statistical learning theories predict that the best performance would be observed in the consistent mapping group, when the regularity is present on every trial. Performance in the mixed design group should be worse than in the consistent mapping group because the regularity is present on only half of the trials. Lastly, the worst performance is expected to appear in the random group, as it contains no beneficial regularity. Alternatively, from the perspective of the