Neuropsychologia 113 (2018) 85–94

Contents lists available at ScienceDirect

Neuropsychologia

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

Delineating resetting and updating in visual working memory based on the object-to-representation correspondence

Halely Balaban[a,*], Trafton Drew[b], Roy Luria[a]

[a] Sagol School of Neuroscience and the School of Psychological Science, Tel Aviv University, Tel Aviv 6997801, Israel
[b] Psychology Department, University of Utah, Salt Lake City 84112, UT, United States

ARTICLE INFO

Keywords:
Visual working memory
Resetting
Updating
Correspondence
Contralateral delay activity

ABSTRACT

When an object we represent in visual working memory (VWM) changes, its representation is modified accordingly. VWM can either access and change the existing representation by an updating process, or it can reset, by encoding the object in its novel status as a new representation. Our goal was to show that the determining factor of updating versus resetting is the availability of a stable correspondence between the object and its VWM representation. Here, we demonstrate that updating relies on the object-to-representation mapping to access and modify the appropriate representation, while losing this mapping triggers a resetting process. We compared very similar situations of object separation that either allowed the mapping to hold, or caused it to be lost. When an object that was mapped to one representation separated, VWM reset, manifested by a sharp drop in the contralateral delay activity (CDA) amplitude (an electrophysiological marker of VWM contents; Experiment 1), and a behavioral cost to detect salient changes that co-occurred with the resetting-triggering event (Experiment 2). When each part was mapped to a different representation, the separation resulted in updating, with a gradual rise in CDA amplitude (Experiment 1), and a reduced behavioral cost (Experiment 2). Thus, while updating and resetting resulted in similar final representations (corresponding to the post-change objects), their dynamics were different, depending on the availability of the mapping. Our results reveal the triggering conditions of resetting and updating, establish methods to study these online processes, and highlight the importance of the object-to-representation correspondence in VWM.

1. Introduction

As items around us constantly move and change, how can our representation of the world be both stable and consistent with the dynamic environment? A central mechanism that tackles this challenge is visual working memory (VWM), our online workspace. VWM holds a limited set of visual representation in an active state, allowing higher cognitive functions to access and manipulate these representations. The online status of the encoded representations is one of VWM's defining features, but most of the research in the field does not target this aspect. While many studies examined, for example, what limits information from entering VWM (e.g., Brady et al., 2011; Luck and Vogel, 2013; Ma et al., 2014), relatively little is known about the interactive nature of the representations within the VWM workspace. Thus, an important open question is how VWM keeps track of the objects it represents, modifying the represented information according to changes that occur in the environment (Ecker et al., 2010; Fallon et al., 2018).

Broadly speaking, we can identify two distinct processes that are responsible for modifying active VWM representations following changes in the environment (see Fig. 1). First, VWM can *update an existing representation*. This updating process involves accessing and modifying the relevant representation to match changes in the world. Evidence suggest that updating can occur in a wide range of situations, such as changes in items' locations (Drew and Vogel, 2008; Drew et al., 2011, 2012), features (color and orientation; Blaser et al., 2000), Gestalt grouping cues (Balaban and Luria, 2015, 2016b; Luria and Vogel, 2014), and even the interpretation implied by the context in which the items are presented (Balaban and Luria, 2016a). For example, when two halves of a random polygon were presented separately, and then moved towards each other and met to form a whole polygon that moved as a coherent object, VWM representations were transformed online from independent to chunked, i.e., from two representations to one representation (Balaban and Luria, 2015).

The second way in which VWM can keep track of the changes in the environment is by *'resetting' its workspace*. In a resetting process, VWM discards the original representation and encodes the novel input as a new representation. Logically, not all changes can (or should) be assimilated into the existing representations, so VWM is expected to "start

* Corresponding author.
E-mail address: halelyba@mail.tau.ac.il (H. Balaban).

https://doi.org/10.1016/j.neuropsychologia.2018.03.038
Received 14 November 2017; Received in revised form 13 March 2018; Accepted 28 March 2018
Available online 29 March 2018

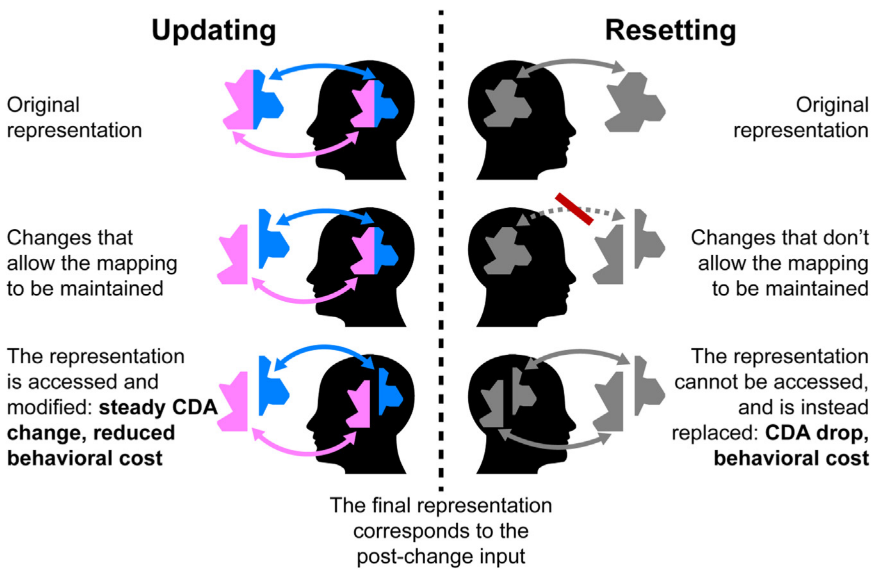

Fig. 1. A schema of the differences between updating and resetting. Changes that allow the unique objects-to-representations mapping to be maintain (e.g., the separation of two parts that were easily distinguishable to begin with, as in the present experiments, see below) result in VWM updating (left): the original representation is accessed (via the mapping) and modified according to the change. Conversely, if the change doesn't allow the mapping to be maintained (e.g., a single polygon separating into two independent polygon-halves), the representation cannot be accessed, and hence a resetting process is triggered (right). In Resetting, the original representation is replaced with a new representation and a new valid mapping. This process has unique neural and behavioral markers.

over" sometimes. A resetting process of some sort must be at the heart of our ability to learn new information and remove outdated schemas. Despite the importance of this process, it has been surprisingly neglected in the context of VWM. Recently, Balaban and Luria (2017) demonstrated that when a polygon moved as a coherent object but then separated into two independently moving halves, VWM did not update the original representations, but instead reset, i.e., removed the previously encoded information and created novel representations. Thus, prior research suggests that when a represented item changes, VWM can either include the change in the item's original representation, or remove the original representation and create a new representation. However, the underlying cause of using one process over the other is still not clear: do different types of changes (e.g., integration versus individuation) trigger updating while others trigger resetting, or does a deeper principle govern the two different processes?

The goal of the current study was to identify the triggering conditions for the updating and resetting processes, to better understand the capability of VWM to cope with the dynamic environment. We argue that the updating process relies on a mapping between each VWM representation and a specific object in the world (Bae and Flombaum, 2013; Kahneman et al., 1992; Levillain and Flombaum, 2012). This one-to-one mapping allows the appropriate representation to be accessed and altered when the object changes. If the unique correspondence is lost, updating is impossible because the representation cannot be accessed, and instead a resetting process is triggered.

Examining previous findings through the prism of the objects-to-representations mapping indeed provides indirect support for this claim. When a whole polygon separates into two independently-moving halves, the initial single mapping cannot support the two objects (none of which correspondence to the original item), resulting in resetting (Balaban and Luria, 2017). Conversely, two halves that meet to form a whole polygon can be updated (Balaban and Luria, 2015), because the original two mappings can still be accessed after the meeting and become grouped together. Furthermore, object-separation was shown to result in updating and not resetting when the parts are easy to individuate, i.e., can initially support two mappings to begin with. For example, updating was found when two overlaid colored squares moved together and then separated, or when two halves of a shape first moved separately and only then met and re-separated (Balaban and Luria, 2017; see also Balaban and Luria, 2016a). Thus, we argue that a change that keeps the object-to-representation mapping valid would trigger an updating process. Conversely, when the change renders the object-to-representation mapping unusable, a resetting process should take place.

The challenge in distinguishing between the online processes of updating versus resetting is that their end result is similar: whether the change in the item is assimilated into the existing representation (i.e., updating) or causes it to be replaced by another representation (i.e., resetting), the final representation in VWM will correspond to the post-change situation. Therefore, accuracy in a classic VWM task (e.g., change detection) cannot be used to distinguish between the two processes. This is because accuracy will reflect the end result of all preceding processes, and both resetting and updating are expected to result in the same final representation. Here, we rely on two approaches that are specifically suited for studying online VWM processing and allow us to distinguish between resetting and updating: an ERP marker of VWM, and a behavioral task that probes VWM concurrently with the changes to the items. These two methods were recently used to support the view that resetting is a distinct process from the more well-characterized updating process (Balaban and Luria, 2017). We now turn to describe these methods and how they can be employed to test our predictions.

The first method designed to reveal differences between updating and resetting utilizes an electrophysiological marker of VWM contents, namely the contralateral delay activity (CDA; McCollough et al., 2007; Vogel and Machizawa, 2004). This ERP component is a sustained posterior negativity whose amplitude rises as more information is held in VWM, reaching a stable plateau that is tightly correlated with the individual capacity limit (for a review, see Luria et al., 2016). Many studies have validated the CDA as a specific marker of VWM, and have shown that it is unaffected by purely perceptual manipulations, such as the brightness of the items (Ikkai et al., 2010; Luria et al., 2010; Ye et al., 2014) or their proximity (McCollough et al., 2007), and does not simply reflect the number of attended locations (Balaban and Luria, 2016a; Balaban and Luria, 2016b; Ikkai et al., 2010; Luria and Vogel, 2014). The precise temporal resolution of ERP, along with the fact that the CDA reflects an online index of the amount of information in VWM, make the CDA an excellent marker of online processes in VWM (e.g., Balaban and Luria, 2016b; Drew et al., 2012; Peterson et al., 2015; Vogel et al., 2005).

Since the CDA can be measured not only during a retention interval, but also when the items are visible (e.g., Drew and Vogel, 2008; Emrich et al., 2009; Tsubomi et al., 2013), it can be used to examine how VWM representations are transformed following changes to the presented items. Previous research demonstrated that when VWM updated its representations, the CDA amplitude steadily changed until it reflected the changes in the environment. For example, when Gestalt objecthood cues were introduced, online grouping was manifested by a *gradual* decrease in CDA amplitude that reflected chunking (Balaban and Luria, 2015, 2016a, 2016b; Luria and Vogel, 2014). Another example for updating is found in studies that sequentially added to-be-encoded items. This manipulation resulted in a *gradual* increase in CDA amplitude (Drew et al., 2011; Vogel et al., 2005), to account for the novel number of items.

Contrary to the gradual amplitude changes in updating, a resetting process is characterized by a *sharp drop* in the CDA amplitude, about 200 ms after the resetting-triggering event. This sharp drop is followed by a rise in amplitude until the CDA represents the current number of attended objects. The transient decrease in amplitude is in line with the removal of the original representations from VWM, followed by encoding the novel form of the presented items. In a recent study, we demonstrated that when a polygon separated into two halves, instead of the steady increase in CDA amplitude that is expected from an updating process, the CDA amplitude sharply dropped and only then recovered until it reached the amplitude of two items (Balaban and Luria, 2017). This result was replicated with different stimuli and a different number of items. This drop was highly robust, and was present in the individual data sets of every one of the subjects. Furthermore, a drop was found also when one object was abruptly replaced by a different object, showing that the CDA-drop is associated with resetting, and not with object-separation per se. Notably, the CDA can dissociate updating and resetting even though both processes result in the same final representation (e.g., two distinct representations following object-separation).

The second method for studying the online dynamics of resetting and updating relies on changing the represented items and probing VWM *while these changes occur* (unlike in a classic change detection task which probes VWM after the changes took place and a final representation, corresponding to the post-change input, could be established). We argue that resetting is triggered by the loss of the objects-to-representations mapping. Thus, changes that co-occur with the resetting-triggering event should go unnoticed, presumably since the representations cannot be accessed while the object-to-representation mapping is unavailable. Conversely, events that occur during an updating-triggering event should be noticed, because there is a valid mapping that allows accessing and modifying the representations. To demonstrate this, Balaban and Luria (2017) used a novel online change detection task, in which a polygon moved on the screen and could change its shape *during this movement* (for evidence that similar VWM processes operate regardless of the retention interval, see Tsubomi et al., 2013). Thus, the vulnerability of VWM representations during updating and resetting, instead of after the processes were completed, could be tested. On half of the trials, the polygon separated, triggering a resetting process. Critically, the time of the shape-change was manipulated: either at the moment of resetting, when VWM should be vulnerable, or at two baseline times of either 250 ms before separation or 250 ms after separation. The results confirmed that there was a pronounced decrease in hit-rate for detecting the shape change only when it co-occurred with the separation, despite this salient change taking place when the items were visible on the screen. This is somewhat reminiscent of the famous "invisible gorilla" effect (Simons and Chabris, 1999), although note that here what was "invisible" was a change in the very item that subjects were tracking, and not in unattended stimuli (for further findings that when several items are tracked but only some of them separate, the cost is specific to the separated item/s, see Balaban et al., 2017). Presumably, this was because the representations could not be accessed during the resetting process (Balaban and Luria, 2017). This behavioral cost was diminished in situations which allowed VWM to update, meaning the cost can differentiate updating and resetting.

While previous data is in line with our argument that resetting and updating are determined by the correspondence of objects and representations, this claim was never directly tested. The goal of the current study was to directly manipulate the mapping that is formed between the presented stimuli and VWM representations. Specifically, we created two types of separating polygons (see Fig. 1): one in which each half can be mapped to a different representation even prior to their separation, and one in which both parts form one coherent object that is mapped to a single representation. That way, we could compare two situations that are very similar but afford two different types of correspondences between an item's parts and VWM representations. This contrasts with previous research that compared different types of stimuli (i.e., polygons vs. colors), or different kinds of circumstances (e.g., one object that splits vs. two halves that first meet and only then split). Each of the present experiments included two very similar conditions of a polygon separating into two-halves, which critically differed in terms of the mappings they encouraged – either a single mapping to the whole polygon (as in our previous work), or two distinct mappings, one to each half of the polygon. We utilized both the distinct neural signature of resetting versus updating (Experiment 1), and the differential behavioral costs the two processes produce (Experiment 2), to provide converging evidence for the underlying factors causing VWM to use each process. We hypothesized that the polygon's separation would lead to resetting when its two halves were mapped to a single VWM representation, but to updating when each half supported an independent representation. Resetting is expected to result in a sharp drop in CDA amplitude, and a pronounced behavioral cost. Conversely, in an updating process the CDA should gradually rise, and the behavioral cost is expected to decrease.

2. Experiment 1: the neural signature of resetting versus updating

The goal of Experiment 1 was to dissociate the updating and resetting processes, based on the mapping between the items and their VWM representations. We argue that an updating process relies on these mappings to access and modify the appropriate representation. We reasoned that if VWM can maintain two separate mappings (one to each polygon-half) before the separation, the mappings would not be lost because they are still relevant even after the separation. Note that the task is to detect a shape change (that can happen only to a single half), and that the halves maintain their shape throughout the trial. Thus, the movement and the separation are task-irrelevant.

To test this prediction, we created bicolored polygons (see Fig. 2) such that each half of the polygon had a different color. Despite being task-irrelevant, the colors supported an independent representation of each half, because they allowed a clear distinction of each part. Using the CDA as a neural marker of VWM, we compared this bicolored polygon condition to a unicolored black polygon. Assuming that one moving black polygon encouraged VWM to represent one item using one mapping (Balaban and Luria, 2015), we expected to find the CDA-drop when a black polygon separated, replicating Balaban and Luria (2017), indicating a resetting process. The reason is that the one mapping did not correspond to any of the polygon-halves, a situation that we argue triggers resetting. Conversely, in the bicolored polygon condition, each half could be represented as an independent object to begin with, creating two mappings even before the separation. Hence, we predicted an updating process in this condition because the polygon's halves could be accessed throughout the movement. This updating process should cause a steady rise in CDA amplitude, without a drop (see Balaban and Luria, 2016a). Another goal of the present experiment was to directly examine the gradation of the CDA in updating, which should be manifesting in a shallower slope compared to the steeper rise of resetting (due to the lack of drop in this condition).

2.1. Materials and methods

Data and code are available in the Open Science Framework: https://osf.io/m6ryg/.

2.1.1. Participants

Participants were Tel Aviv University students who received partial course credit or 40 NIS (approximately $10) per hour for participation.

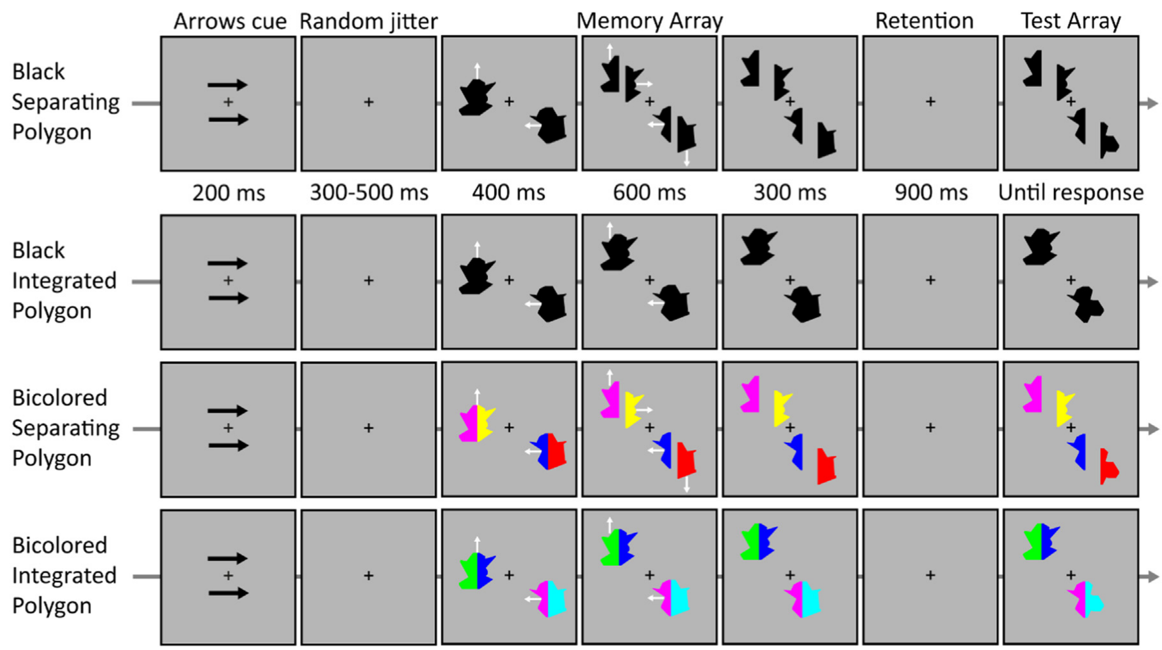

Fig. 2. The trial sequence in the different conditions of Experiment 1. Subjects performed a shape change detection task (color was irrelevant) with polygons. Items moved for 1000 ms (which was task-irrelevant) and remained stationary for an additional 300 ms. The black arrows indicated the relevant side for the upcoming trial. White arrows indicate movement direction and were not visible on the screen. At test, a single polygon-half could change its shape (here, only change trials are presented). Black and bicolored polygons were presented in different blocks, and regardless of the color, on half of the trials the polygon separated into two independently moving halves.

All subjects had normal or corrected-to-normal visual acuity and normal color-vision. The experiment included 12 subjects (8 females, mean age 24.9). All subjects' rejection rate (see below) was below our predefined criterion of 25%, and hence all of them were included in the final analysis.

2.1.2. Stimuli and procedure

We used a bilateral shape change detection task (Fig. 2). Stimuli were presented on a 23-in. LED monitor with a 120 Hz refresh rate, using a resolution of 1920 × 1080.

Each trial started with a 750 ms fixation display, with a black cross (0.4° × 0.4° of visual angle, from a viewing distance of approximately 60 cm) in the center of a grey screen. Then, two white arrows (1.9° × 0.4°) appeared for 200 ms, pointing to the to-be-attended side for the upcoming trial (left or right, with an equal probability). After a fixation display that lasted either 300, 400, or 500 ms (randomly determined with an equal probability), the memory array of moving random polygons was presented for 1300 ms. Their trajectories were straight, and the movement was restricted to a single side of the screen, i.e., the items never crossed the center of the screen. The allowed directions (randomly determined with an equal probability) were up, down, or horizontally towards the fixation. The two sides always included the same type of stimuli. The items appeared at random locations inside an imaginary 4.5° × 3.5° rectangle, which was placed 1.65° next to the center of the screen (one in each side of the screen). Following a 900 ms retention interval including only the fixation cross, the items reappeared, and subjects had to indicate whether a single polygon-half in the attended side changed to a new one (by pressing the "Z" or "/" buttons on a standard computer keyboard for "same" or "different", respectively), which had a 50% probability.

Stimuli were random polygons (1.6° × 1.6°), constructed by presenting one of 6 left-side and one of 6 right-side polygon-halves next to each other (the right side of each left-side half and the left side of each right-side half subtended the full length of the shape, so each left-side half could combine with each right-side half, for a total of 36 possible polygons). Stimuli were randomly drawn without replacement at the beginning of the trial, separately for the left and right sides of fixation.

All trials started with 400 ms of joint movement, with the two polygon-halves moving as a coherent polygon. Then, there were two main movement conditions, with an equal probability. In the Integrated Polygon condition, the two halves continued to move together for another 600 ms (however, note that eventually a change could only happen to one of the halves). In the Separation condition, the two halves separated after 400 ms, such that each half moved independently for another 600 ms. Items moved at a constant rate of about 1.5° per second. Both of these trial types included another 300 ms of stationary display at the end of the movement, for a total of 1300 ms. To ensure subjects paid attention to the initial 400 ms of movement, 10% of the trials were catch trials, in which the memory array disappeared after the initial 400 ms. These trials were not further analyzed.

In separate blocks, the polygon-halves were either all black, or in a different color for each half (randomly selected without replacement from a set of 6 highly discriminable colors: yellow, green, cyan, blue, magenta, and red). The color was always task-irrelevant, and it never changed between the memory and test arrays. Thus, there were 4 conditions: Integrated Black Polygon, Separating Black Polygon, Integrated Bicolored polygon, and Separating Bicolored polygon.

The experiment started with 12 practice trials, followed by 7 blocks of one type (i.e., black or bicolored items), and then 7 blocks of the second type, each with 60 trials, for a total of 840 experimental trials. The order of the blocks was counter-balanced between subjects.

2.1.3. EEG recording and analysis

EEG was recorded inside a shielded Faraday cage, with a Biosemi ActiveTwo system (Biosemi B.V., The Netherlands), from 32 scalp electrodes at a subset of locations from the extended 10–20 system (mostly parietal and occipital sites in which the CDA is most pronounced: Fp1, Fp2, AF3, AF4, F3, F4, F7, F8, Fz, FCz, C3, C4, Cz, T7, T8, P1, P2, P3, P4, P5, P6, P7, P8, Pz, PO3, PO4, PO7, PO8, POz, O1, O2, and Oz), as well as from two electrodes placed on the mastoids. EOG

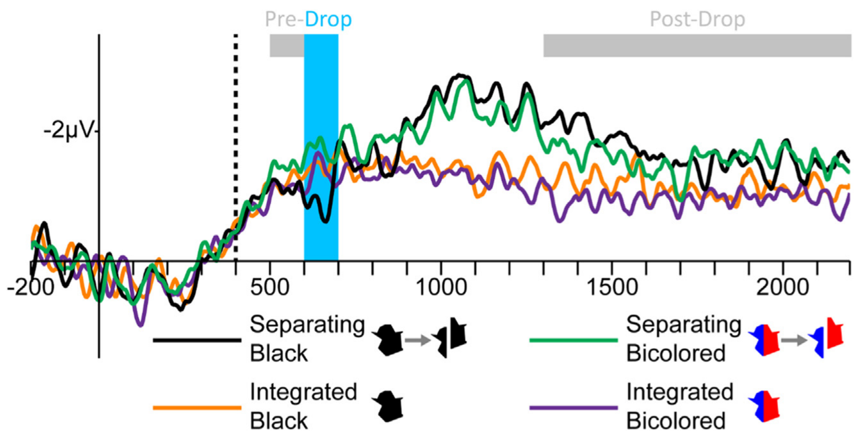

Fig. 3. The CDA results of Experiment 1 (Negative voltage is plotted upwards). Averaged amplitude over the P7/8, PO3/4, and PO7/8 electrode pairs. The black dashed line depicts the time of separation, and the colored rectangles depict the analyzed time windows: Drop (200–300 ms after separation) in light blue, and Pre- and Post-Drop (100–200 ms after separation, and during the retention interval, respectively). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

was recorded from two electrodes placed 1 cm laterally to the external canthi, and from an electrode beneath the left eye. Data was digitized at 256 Hz.

Offline signal processing was performed using the EEGLAB Toolbox (Delorme and Makeig, 2004), ERPLAB Toolbox (Lopez-Calderon and Luck, 2014), and custom Matlab (The Mathworks, Inc.) scripts. All electrodes were referenced offline to an average of the mastoids. The continuous data was segmented into 2400 ms epochs, from 200 ms before memory array onset to the end of the retention interval (2200 ms from memory array onset). Artifact detection was performed using a sliding window peak-to-peak analysis, with a threshold of 80 μV for the EOG electrodes, and 100 μV for the analyzed electrodes (P7, P8, PO3, PO4, PO7, and PO8). These procedures resulted in a mean rejection rate of 8.8%. Additionally, to ensure the results were not driven by eye-movement, we examined the horizontal EOG. The mean HEOG deflection was roughly 4μV (corresponding to about 0.25° of eye-movement; Hillyard and Galambos, 1970), and, critically, did not differ systematically between conditions. These findings, along with an HEOG plot, are summarized in the Supplementary material. Only trials with a correct response were included in the analysis. For illustration purposes, the epoched data was low-pass filtered using a noncausal Butterworth filter (12 dB/oct) with a half-amplitude cutoff point at 30 Hz. This was done before plotting, and all statistical analyses were performed on the unfiltered data.

Epoched data were averaged separately for each condition. The CDA difference wave was calculated by subtracting the average activity at electrodes ipsilatersal to the attended side from the average activity at electrodes contralateral to the attended side. Our dependent measure was mean amplitude at the average of P7/8, PO3/4, and PO7/8 (there were no differences in the qualitative pattern observed in the critical analyses in each electrode-pair separately), at three time-windows (which were the same as in our previous study; Balaban and Luria, 2017): ‘Pre-Drop’, 100–200 ms after separation (note that the CDA usually takes about 200 ms to develop after the items are first presented, e.g., Vogel and Machizawa, 2004), ‘Drop’, 200–300 ms after separation, and ‘Post-Drop’, during the entire range of the retention interval (which translates into 500–600, 600–700, and 1300–2200 ms from memory array onset, respectively).

2.1.4. Statistical analysis

We analyzed mean CDA amplitude as a dependent measure using a three-way analysis of variance (ANOVA), with Movement Type (Integrated Polygon vs. Separation), Block Type (Black vs. Bicolored), and Time (Pre-Drop, Drop, and Post-Drop) as within-subject variables. We then broke down this analysis into 3 simple two-way ANOVAs, by Time, examining Movement Type and Block Type within each time-window. An additional ANOVA was conducted on behavioral accuracy. Finally, we performed planned comparisons (contrasts) between the different conditions, namely between the Integrated and Separated conditions within each Block Type, and between the Black and Bicolored items within each Movement Type. We also report effect sizes: partial η^2 for ANOVAs, and Cohen's d for pairwise comparisons.

An additional analysis was aimed at showing that updating involves a steadier rise in CDA amplitude than resetting. For each subject, we calculated the mean amplitude in each of 9 consecutive time-windows of 50 ms, from 650 to 1100 ms after trial onset. For each of the Separation conditions, we examined the linear rise in amplitude, by calculating the Pearson correlation between amplitude and time (defined as the mid-point of the time-window) across the 9 time-windows. Finally, we compared the slopes of the best-fitting regression line between the Black and Bicolored items, to establish the steepness of the CDA rise in each condition.

2.2. Results

2.2.1. EEG results

As predicted, we observed a drop in CDA amplitude after a black polygon separated, but not after a bicolored polygon separated (Fig. 3). There was a significant interaction of Movement Type, Block Type, and Time, $F(2, 22) = 6.88$, $p < 0.005$, partial $\eta^2 = 0.38$, suggesting that the different effect of the separation on the black and bicolored polygons depended on the time-window. To further elucidate the pattern of results, we examined the effects within each time-window separately.

Critically, in the “Drop” time-window (200–300 ms after separation), we found a significant interaction of Movement Type and Block Type, $F(1, 11) = 17.25$, $p < 0.005$, partial $\eta^2 = 0.61$. There was no significant effect for Movement Type or Block Type, $F < 1$, $p = 0.58$, partial $\eta^2 = 0.03$, and $F(1, 11) = 2.79$, $p = 0.12$, partial $\eta^2 = 0.20$, respectively. There was a temporary significant decrease in CDA amplitude for the separating black polygon (−0.96 μV, SD: 1.22) compared with the black integrated polygon (−1.51 μV, SD: 1.1), $F(1, 11) = 10.99$, $p < 0.01$, Cohen's $d = 0.47$, replicating our previous findings (Balaban and Luria, 2017), in line with the loss of VWM contents associated with a resetting process. In contrast, there was a slight *increase* (which was not statistically significant) in amplitude for the separating bicolored polygon (−1.69 μV, SD: 1.16) compared with the bicolored integrated polygon (−1.35 μV, SD: 0.9), $F(1, 11) = 1.88$, $p = 0.2$, Cohen's $d = 0.33$, indicating VWM did not reset when the halves of the polygon could be easily individuated, based on the different color of each half. These results suggest that when the polygon could be represented as two independent objects even prior to separation, the separation did not trigger a resetting process, presumably because VWM held two mappings that could still be accessed after separation. Conversely, when the polygon was represented as a single object because of its uniform color, a single mapping could not support the two independent halves and a resetting process was triggered.

Conversely, In the Pre-Drop time-window (100–200 ms after separation), there was no significant interaction of Movement Type and

Block Type, $F(1, 11) = 3.43$, $p = 0.09$, partial $\eta^2 = 0.23$. There was no significant effect of Movement Type or Block Type, $F < 1$, $p = 0.60$, partial $\eta^2 = 0.03$, and $F(1, 11) = 1.10$, $p = 0.32$, partial $\eta^2 = 0.09$, respectively. The lack of an interaction in the Pre-Drop time-window demonstrates that resetting only resulted in a transient decrease in amplitude, instead of a generally lower amplitude. Unlike in the Drop time-window, there was no significant difference between the Black Separating Polygon (−1.12 µV, SD: 1.3) and the Black Integrated Polygon (−1.25 µV, SD: 1.2), $F < 1$, $p = 0.52$, Cohen's $d = 0.11$, suggesting that at this time-window the polygon was still represented as a single unit in VWM, and indicating that the CDA drop is specific to the Drop time-window.

The lack of a drop in CDA amplitude when the separation allowed the mapping to hold suggests that no resetting process is triggered, meaning that VWM could update. However, this conclusion is based on a null result, which is problematic. We therefore contrasted updating and resetting in a complementary way, to support the claim that updating involves a steadier rise in CDA amplitude than resetting (e.g., Drew et al., 2011; Luria and Vogel, 2014). For that aim, we calculated the CDA mean amplitude in 9 consecutive time-windows of 50 ms each, from 650 to 1100 ms from trial onset. This time window subtends the full time-course of the rise in CDA amplitude, starting from the lowest point of the drop in the resetting condition. We then examined the slope of the linear trend for each separation condition (Fig. 4). The slope should be shallower for updating compared with resetting, because in updating the CDA amplitude is expected to gradually rise while resetting results in a sharp drop followed by a quick rise in amplitude. Indeed, we found that the slope was significantly shallower for the Bicolored Separating Polygons (i.e., updating) than for the Black Separating Polygons (i.e., resetting), $t(11) = 4.61$, $p < 0.001$, Cohen's $d = 0.99$, supporting the claim that updating involves a steadier rise in CDA amplitude without a drop (notably, similar analyses with window sizes of 25 ms and 100 ms led to the same result).

Finally, the pattern of results in the Post-Drop time-window (during the retention interval) also differed from the Drop time-window. CDA amplitude towards the end of the trial was higher for both of the Separated Polygon conditions than for both of the Integrated Polygon conditions, resulting in an effect of Movement Type, $F(1, 11) = 56.45$, $p < 0.00002$, partial $\eta^2 = 0.84$, no significant effect of Block Type, $F(1, 11) = 1.24$, $p = 0.30$, partial $\eta^2 = 0.10$, and no significant interaction, $F < 1$, $p = 0.38$, partial $\eta^2 = 0.07$. This suggests both the Black and Bicolored Separated Polygons were ultimately represented as individuated halves in VWM. Amplitude was higher for the Separated Polygons (Black: −1.64 µV, SD: 0.97; Bicolored: −1.57 µV, SD: 0.83) than for the Integrated Polygons (Black: −1.2 µV, SD: 0.92; Bicolored: −0.96 µV, SD: 0.72), $F(1, 11) = 13.42$, $p < 0.005$, Cohen's $d = 0.47$ for the Black Polygons, and $F(1, 11) = 27.6$, $p < 0.0005$, Cohen's $d = 0.80$ for the Bicolored Polygons. The Post-Drop amplitude did not significantly differ between the Black and Bicolored Separating Polygons, $F < 1$, $p = 0.67$, Cohen's $d = 0.08$, and also did not significantly differ between the Black and Bicolored Integrated Polygons, $F(1, 11) = 1.8$, $p = 0.21$, Cohen's $d = 0.30$. Thus, VWM could recover and represent the correct number of items (i.e., more representations following the individuation signal of the separation), not only when updating but also after resetting, as expected. This demonstrates the advantage of the CDA as an online marker of these processes in VWM over simple behavioral measures: amplitude of this component reveals the entire time course of VWM representations and not only their final outcome.

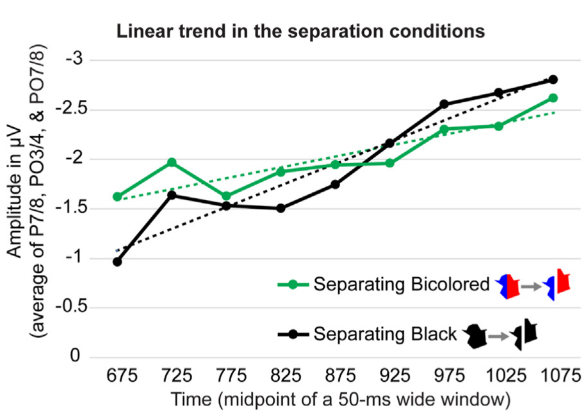

Fig. 4. Mean amplitude in each of 9 consecutive time-windows of 50 ms in the two separation conditions. The slope of the linear trend (shown in dashed lines) was shallower for the Bicolored polygon compared with the Black polygon.

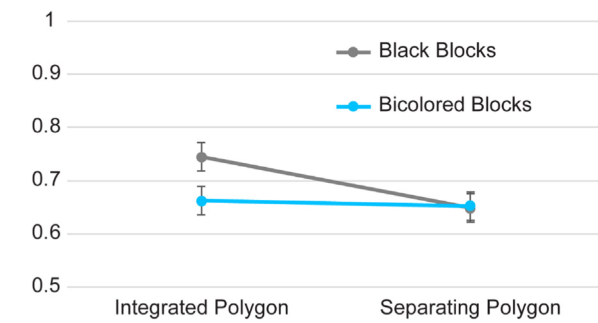

Fig. 5. The accuracy results of Experiment 1, by Movement Type (Separated vs. Integrated) and Block Type (Black Polygons vs. Colored Polygons). Error bars denote standard deviation.

2.2.2. *Behavioral results*

Accuracy in the change detection task resulted in a significant effect of Stimuli Type (Black Polygons versus Colored polygons), $F(1, 11) = 24.25$, $p < 0.0005$, partial $\eta^2 = 0.69$, a significant effect of Movement Condition (Separation versus Integrated), $F(1, 11) = 23.52$, $p < 0.001$, partial $\eta^2 = 0.68$, and a significant interaction, $F(1, 11) = 32.36$, $p < 0.0002$, partial $\eta^2 = 0.75$ (Fig. 5). This pattern of results was driven by a higher accuracy for the Black Integrated Polygon (0.74, SD: 0.06) compared with both the Black Separating Polygon (0.65, SD: 0.06), $F(1, 11) = 49.13$, $p < 0.00005$, Cohen's $d = 1.58$, replicating our previous findings (Balaban and Luria, 2017), and the Bicolored Integrated Polygon (0.66, SD: 0.05), $F(1, 11) = 35.68$, $p < 0.0001$, Cohen's $d = 1.44$. Conversely, the Bicolored Integrated Polygon did not significantly differ from the Bicolored Separating Polygon (0.65, SD: 0.05), $F < 1$, $p = 0.46$, Cohen's $d = 0.20$, and the Black and Bicolored Separating Polygon also did not significantly differ, $F < 1$, $p = 0.57$, Cohen's $d = 0.08$.

We interpret the accuracy results pattern by noting that the Integrated Black Polygon is the only condition that included a single object and hence a single object-to-representation mapping throughout the trial, which resulted in a better performance. The other conditions involve two mappings: in the Bicolored conditions, the two halves could be individuated starting from the trial's onset, and in the Black Separating Polygon condition, the two halves were individuated following their separation. Thus, the higher accuracy for the Black Integrated Polygon is in line with previous findings of object-based advantages in VWM (e.g., Balaban and Luria, 2016b; Delvenne and Bruyer, 2006; Gao et al., 2016).

2.3. *Summary*

The results of Experiment 1 support the notion that the factor determining which online process will take place following a change in the environment is the stability of the object-to-representation correspondence. Using the CDA as a neural marker of VWM, we delineated when resetting versus updating are triggered. We showed that the updating process depends on a stable correspondence between the representations and the items, allowing access to each representation. If an active mapping can be maintained despite of the items' separation (here, due to the different color of each polygon-half), VWM can update, resulting in a gradual change in the CDA amplitude, with a shallow rise. If this essential mapping is no longer relevant (e.g., when a unicolored polygon separated into two halves), a resetting process is initiated in VWM, marked by the sharp drop in CDA amplitude. Notably, the situations involved in resetting and updating in the present study were very similar (i.e., in both cases a polygon separated into two halves), and the end result of both processes was the same, namely VWM representations that correspond to the correct number of objects. Nevertheless, the CDA could differentiate between updating and resetting based on the distinct dynamics they have, making the CDA an excellent tool for the study of online processes in VWM.

3. Experiment 2: behavioral performance during resetting versus updating

The goal of Experiment 2 was to further dissociate resetting from updating, in terms of their differential reliance on a stable object-to-representation mapping. Specifically, we argue that resetting is triggered when this correspondence is lost, and hence the representations cannot be accessed via the initial mapping and updated in response to changes. Specifically, our CDA results suggest that VWM should be vulnerable during resetting, and subjects should be more likely to miss salient changes in the environment that occur during the resetting process. Conversely, in situations that allow the correspondence to hold, and therefore result in updating, VWM should more easily cope with simultaneous changes in the environment. To test this, we used an online change detection task, in which subjects observed moving polygons while monitoring for a shape-change in one of the polygon's halves (during the movement, see Fig. 6). Recently, we found a cost to detect salient changes when they occurred during a resetting process (Balaban and Luria, 2017). When a uniform black polygon separated, causing the single-object mapping to be lost and triggering a resetting process, we expected to replicate the finding of a cost in performance. Critically, we compared this to a situation where a separate mapping could be associated with each half of the polygon. We encouraged a separate mapping by drawing a thin red frame (2 pixels wide) around each half. We hypothesized that this very similar situation would result in a smaller cost, since the original mapping can be maintained and no resetting process is needed.

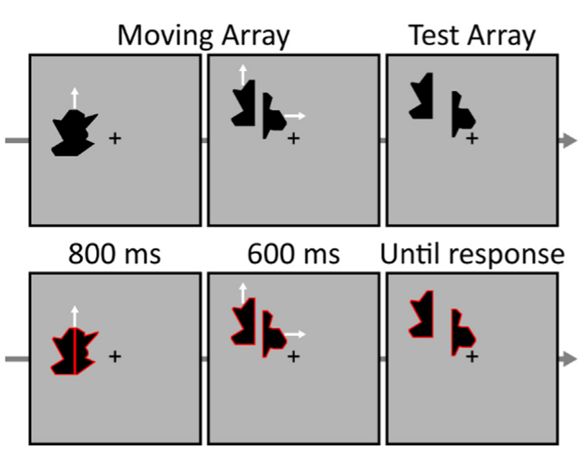

Fig. 6. The trial sequence in the different conditions of Experiment 2. Subjects performed an online shape change detection task with polygons. Items moved for 1400 ms, and during this movement, a single polygon-half could change its shape (only change trials are shown). For half of the subjects, each polygon-half had a thin red frame around it (which was task-irrelevant), and for the other half there were no frames. Regardless of the group, on half of the trials the polygon separated into independent halves after 800 ms.

We previously found a smaller cost when the halves first moved separately and only then joined to form a polygon that subsequently separated (Balaban and Luria, 2017). However, in this situation the separation could be ignored because all trials in which the halves joined also involved a subsequent separation. Therefore, it is possible that neither updating nor resetting were triggered. Hence, it remains possible that when VWM has to update, a cost similar to resetting will be found. Here, we compared two separation situations which were the same in terms of their predictability (overcoming the limitations from our previous study), but differed in terms of the units of mapping (i.e., the whole polygon or the separate halves), thus allowing us to directly examine the behavioral consequences of a separation that leads to either updating or resetting.

3.1. Materials and methods

3.1.1. Participants

Participants were Tel Aviv University students who received partial course credit or 40 NIS (approximately $10) per hour for participation. All subjects had normal or corrected-to-normal visual acuity and normal color-vision. Since this was a between-subjects design (see below), the experiment included 60 participants (46 females, mean age 24.7), 30 in each group.

3.1.2. Stimuli and procedure

We used an "online" shape change detection task, in which the changes occurred while the items were visible on the screen (Fig. 6). Each trial started with a fixation display, with a black cross ($0.4^\circ \times 0.4^\circ$) in the center of a grey screen, for 800 ms, followed by the presentation of a black polygon ($1.6^\circ \times 1.6^\circ$), constructed from two halves, as in Experiment 1 (see above). Stimuli were randomly selected for each trial. The polygon moved on the screen for 1400 ms, covering 2°, in straight trajectories, either up, down, left, or right. During the polygon's movement, one of its halves could change to another half (a right-side half could only change to a new right-side half, and a left-side half only to a new left-side half), with a 50% probability. When the polygon stopped, subjects indicated whether they saw a change in shape during the movement (by pressing the "Z" or "/" buttons on a standard computer keyboard for "same" or "different", respectively).

Subjects were randomly assigned to one of two groups, each completing the same experiment, but with different stimuli. For the 'No-Frames' group, stimuli were the same black polygon-halves used in Experiment 1. For the 'Frames' group, stimuli were the same, except that a thin red line was drawn around each polygon-half. When the two halves moved together as one polygon, this created the appearance of a thin (2 pixels, 0.05°) red line in the middle of the polygon, marking the different halves. We used a between-subjects design because seeing the polygon-halves surrounded by frames could have caused VWM to represent them as individuated objects in the No-Frames blocks as well (this was less of a concern in Experiment 1, since in an EEG experiment subjects cannot move their eyes to look directly at the shapes).

All trials started with 800 ms of joint movement, with the two polygon-halves moving as a coherent polygon. Then, the halves could either continue to move together, or separate and move independently for the remaining 600 ms (with an equal probability). If the polygon separated, one of the halves continued in the same trajectory, and the other half turned to a different direction, 90° from the original direction.

Regardless of the movement condition (Integrated Polygon vs. Separation), changes could occur either 550, 800, or 1050 ms after the polygon's onset, which translates to 250 ms before separation, during separation, or 250 ms after separation (if it occurred).

Twelve practice trials were followed by 8 experimental blocks of 60 trials each, for a total of 480 trials.

3.1.3. Statistical analysis

We analyzed hit rate as a dependent measure, with a three-way ANOVA, with Movement Type (Integrated Polygon vs. Separation) and Time of Change (−250, 0, or 250 ms relative to separation) as a within-subject variable and Stimuli Type (No Frames vs. Frames) as a between-subjects variable. We followed these tests by pairwise comparisons of the different conditions at time 0 (i.e., the time of separation, and the onset of the resetting process).

3.2. Results

Our dependent measure was hit rate, meaning only trials that included a shape-change were analyzed. Notably, determining the time of change is only possible for change trials. Accuracy involves both change and no-change trials, and therefore lower accuracy for the separation condition might be driven by a high false alarm rate, indicating that the separation itself tended to be perceived as a change. Conversely, our focus here was not on how people perceive the separation, but rather on how their perception of the shape-change is affected by the separation. For a detailed analysis and discussion of false alarms (FA) in the online change detection paradigm, see Balaban et al., 2017. As in our previous work (Balaban et al., 2017), FA rate was low (7% and 5% in the Frames and No-Frames groups, respectively), suggesting that participants did not often guess during this task.

As hypothesized, the behavioral cost to detect salient changes that coincide with a polygon's separation was reduced by adding frames around each half (Fig. 7), leading to a significant interaction of Stimuli Type (Frames vs. No-Frames), Movement Type (Integrated vs. Separation) and Time of Change (−250, 0, or 250 ms relative to the separation), $F(2, 116) = 5.2$, $p < 0.01$, partial $\eta^2 = 0.08$.

Hit rate for the Separating Polygon at time 0 was higher when each polygon-half had a frame around it (i.e., in the Frames group) than in the No-Frames group, $F(1, 58) = 5.09$, $p < 0.03$, Cohen's $d = 0.59$. For the No-Frames group, hit rate at time 0 was lower in the Separating Polygon condition than in the Integrated Polygon condition, $F(1, 58) = 52.42$, $p < 0.000001$, Cohen's $d = 1.38$, replicating our previous results (Balaban and Luria, 2017). There was a cost in performance also for the Frames group, $F(1, 58) = 13.02$, $p < 0.001$, Cohen's $d = 1.01$. However, critically this cost (i.e., the difference between the Separating Polygon and Integrated Polygon conditions) was smaller than the cost observed for the No-Frames group, $t(58) = 2.57$, $p < 0.02$, Cohen's $d = 0.68$, corroborating our prediction. These results were specific to time 0, while at −250 or 250 ms there was no significant cost in performance for the Separating Polygon compared with the Integrated Polygon condition (hit rate was *higher* for the Separating Polygon at time −250, $F(1, 58) = 6.65$, $p < 0.02$, Cohen's $d = 0.40$, which was due to a higher hit rate for the Separating Polygon in the No-Frames group, $F(1, 58) = 6.00$, $p < 0.02$, Cohen's $d = 0.60$, and at time 250 the Separating Polygon and Integrated Polygon did not significantly differ, $F(1, 58) = 1.33$, $p = 0.25$, Cohen's $d = 0.18$), and no significant difference in the costs between the groups, $t < 1$, $p = 0.38$, Cohen's $d = 0.23$, and $t(58) = 1.07$, $p = 0.29$, Cohen's $d = 0.29$ for time −250 and 250, respectively.

The frames indeed reduced the behavioral cost to detect changes that coincide with the polygon's separation, but we still found some cost in the Frames group. A possible reason could be that the common-fate grouping cue (i.e., the integrated movement phase) was strong enough to override the individuation cue of the frames. This would have caused the two halves to be perceived as a single object prior to separation, and hence the separation might have triggered a resetting process on some of the trials or for some of the subjects, resulting in a cost. This cost was still smaller when compared to the No-Frames group, in which no contrasting individuation cue was present. Notably, this hypothesis does not go against our claim that the resetting process and its behavioral cost are determined by the object-to-representation mapping, but simply that the mapping itself is determined by an interplay of several factors and cues.

The key finding of Experiment 2 is that the behavioral cost of resetting was significantly reduced simply by adding thin red frames around each polygon-half. Presumably, this was because the frames allowed each half to be represented as a distinct object in VWM (perhaps not completely, see above), along with an independent mapping between the polygon-half and its representation. Thus, a subtle manipulation of a 2-pixels frame allowed the original pre-separation mapping to hold more strongly throughout the separation, leading VWM to update instead of resetting. This made it possible to still access the original representations via their mapping, diminishing the cost in performance.

4. Discussion

The goal of the present study was to examine the triggering factors of online resetting versus updating in VWM. These two processes allow VWM representations to remain in line with attended information in an ever-changing world, in two different ways – either by incorporating changes into existing representations (updating), or by discarding the original representations and starting anew (resetting). We provide evidence that resetting, as indicated by its neural and behavioral markers, is triggered by a loss of object-to-representation correspondence, while updating can occur as long as this unique mapping holds. We compared very similar situations of object-separation, manipulating the initial mapping between items and their VWM representations. Specifically, we presented either uniform black polygons, or polygons in which each half could be individuated even prior to separation, using either two different colors (Experiment 1) or thin red frames around each half (Experiment 2). When a black uniform polygon separated, a resetting process was triggered: there was a sharp drop in the amplitude of the CDA (Experiment 1), indicating a loss of VWM-contents.

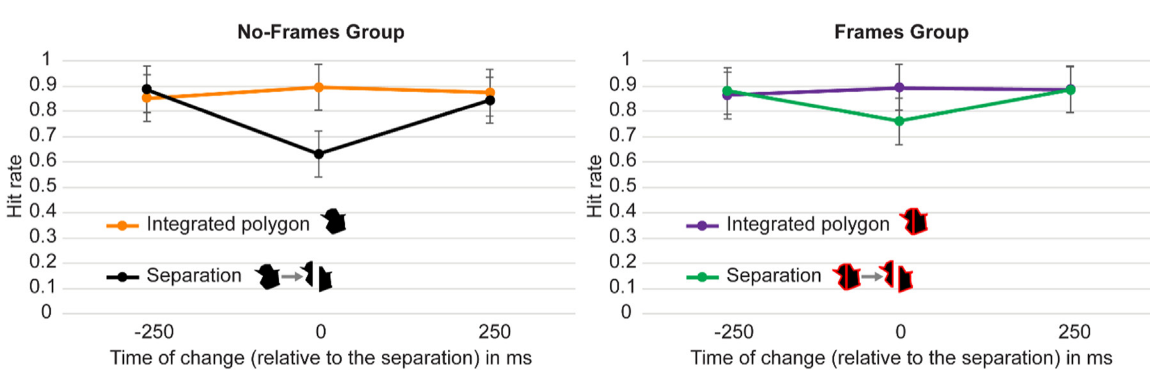

Fig. 7. Hit rate in Experiment 2, by condition (Separation vs. Integrated Polygon). Left panel presents the No Frames group, and the right panel presents the Frames group. Error bars denote standard deviation.

Additionally, changes coinciding with the separation were often missed (Experiment 2), suggesting that the contents of VWM could not be accessed, presumably because the mapping was lost, and the corresponding representation could not be accessed. Conversely, when each polygon-half could be easily identified before the separation, the CDA-drop was eliminated, and instead we found evidence for a steady change in CDA amplitude (Experiment 1), characteristically to an updating process (e.g., Vogel et al., 2005). Moreover, the behavioral cost of detecting changes that co-occur with the separation was reduced (Experiment 2), indicating that the contents of VWM could be more readily accessed via the still-valid mapping.

We argue that these neural and behavioral results are due to the ability to map each polygon-half to an independent representation in VWM even during their joint movement (because of the different colors or frames), unlike the single mapping to the uniform black polygon. Thus, the separation did not destroy the original mapping, making the original representations still accessible, and allowing VWM to update them following separation, instead of triggering a resetting process. Taken together, our results suggest that when an object we represent in VWM changes (e.g., separates into parts), the factor determining which online process will be triggered is the availability of a stable correspondence between the objects and their representations. Updating relies on a previously established mapping to access the representations and transform them. Conversely, resetting is triggered when the mapping in lost, causing VWM to start anew, re-encoding the items in their novel status. Importantly, our behavioral results demonstrate the costs of losing the correspondence: salient changes that occur while the representations are inaccessible are frequently missed.

Previous studies have shown that VWM can update its representations in response to a wide range of changes. For example, changes in item's location (e.g., Drew and Vogel, 2008) or in its features (Blaser et al., 2000) can be incorporated into the object's existing representation, and introducing Gestalt grouping cues can lead the representations of separate objects to be chunked (e.g., Luria and Vogel, 2014). The present results show that in some situations, the *removal* of grouping cues can also result in updating, namely when two parts of an object are separated and start to move independently (see also Balaban and Luria, 2016a; Balaban and Luria, 2017), demonstrating updating in a new situation. Critically, however, we found that updating, i.e., incorporating this individuation cue into the original representations, can occur only if each part supported a distinct representation in VWM before the separation (e.g., based on distinct colors). Only if each part is mapped to a unique VWM representation, does the separation signal VWM to access the original representations and un-chunk them. Thus, the present results go beyond separation per se, putting a strong constraint on the online updating process, by delineating when VWM is unable to update. Namely, updating relies on a correspondence between an object and its VWM representation. Whenever a change in the object causes its mapping to a representation to be lost, updating is no longer possible, and VWM must reset. Hence, it is not the type of change an object undergoes that limits updating (i.e., object separation can be successfully updated under certain conditions), but the stability of the correspondence prior to the change. Therefore, an informative way to understand the updating process itself, and the ways it enables VWM representations to stay in line with the environment, is to study the correspondence mechanism (Pylyshyn, 2000).

The theoretical notion of correspondence relates to the conceptualization of the object file theory (Kahneman et al., 1992). In this framework, findings of a behavioral benefit for consecutive presentations of stimuli sharing all feature were interpreted as stemming from a pointer-like system that connects objects in the environment to temporary representations that include their different features (see also Bae and Flombaum, 2013; Levillain and Flombaum, 2012). The central factor driving this correspondence was assumed to be location. This is echoed by our findings that object-separation can destroy the correspondence, leading to a resetting process. However, the present finding that very similar situations of separation lead either to resetting or to updating strongly suggest that the correspondence process relies on more than spatiotemporal continuity, integrating factors such as the items' history and their features (e.g., distinct color) to individuate objects (cf, Balaban and Luria, 2016a; Balaban and Luria, 2016b; Luria and Vogel, 2014).

This demonstrates the importance of the present findings in establishing the stable neural and behavioral patterns associated with the destruction of the object-to-representation correspondence (i.e., with the VWM resetting process), both of which are highly specific. The drop in CDA amplitude is found around 200 ms after the onset of the resetting process (e.g., after object-separation or object-replacement), in parietal-occipital electrodes. The behavioral cost to detect salient changes appears to be distinct from the perceptual or attentional dynamics of the separation per se, and is found for changes coinciding with resetting but not for changes occurring 250 ms afterwards. We have now replicated both these behavioral and neural patterns in several experiments (see also Balaban and Luria, 2017; Balaban et al., 2017). The specificity and replicability of the CDA-drop and behavioral cost allow them to serve as markers of the resetting process. While different theories previously postulated a process for removing irrelevant information from working memory (Hasher et al., 1999; Oberauer et al., 2012), our novel markers allow directly studying how the resetting process operates. Many aspects of the resetting process, such as its temporal characteristics and subjects' ability to control it in a top-down manner, are yet unknown, and can be the target of future studies.

Furthermore, the present demonstration that resetting is specifically triggered due to a loss of mapping between items and their representations suggests that the resetting markers can be used to study the factors underlying the correspondence process itself. For example, the fact that object-separation leads to resetting suggests that a coherent object might be the unit of correspondence between the continuous visual input and our VWM representations, in line with other object-based dynamics in VWM (e.g., Gao et al., 2016; Luria and Vogel, 2011; Zhang and Luck, 2008). Generally, if the generating cause of resetting is a lost correspondence, resetting can reveal the necessary factors of the pointer system (Pylyshyn, 2000): changes that trigger resetting can be assumed to be critical for the mapping process. Thus, by establishing tools for studying the critical components for the correspondence system, the current approach has promise for understanding one of the fundamental building blocks of continuous visual perception.

5. Conclusion

VWM representations must constantly change to keep track of our dynamic environment. The present results shed light on two online processes that achieve this, namely updating, i.e., modifying an existing representation to fit the novel input, and resetting, i.e., encoding the novel input as a new representation, and presumably discarding the original one. By comparing very similar situations that differ only in the objects-to-representations mapping they afford, we identified the triggering conditions of updating versus resetting. Once the process has been completed, updating and resetting ultimately result in a similar state of VWM representation, presenting a challenge to study these important constructs. The present studies establish reliable electrophysiological and behavioral methods to differentiate updating and resetting, namely the CDA drop versus its gradual rise, and the behavioral cost specifically tied to moments in time when an object representation changes. Finally, the results corroborate the importance of a pointer-like system that is at the heart of the mapping between items in the environment and their VWM representations, determining which online process follows a change.

Funding

This work was supported by the Binational Science Foundation [grant number 2015301 awarded to Roy Luria and Trafton Drew], the Israel Science Foundation [grand number 0610718231 awarded to Roy Luria], and the Azrieli Fellowship [awarded to Halely Balaban].

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.neuropsychologia.2018.03.038.

References

Bae, G.Y., Flombaum, J.I., 2013. Two items remembered as precisely as one: how integral features can improve visual working memory. Psychol. Sci. 24 (10), 2038–2047. http://dx.doi.org/10.1177/0956797613484938.
Balaban, H., Drew, T., Luria, R., 2017. Visual working memory can selectively reset a subset of its representations. Psychon. Bull. Rev. 1–7.
Balaban, H., Luria, R., 2015. The number of objects determines visual working memory capacity allocation for complex items. NeuroImage 119, 54–62.
Balaban, H., Luria, R., 2016a. Object representations in visual working memory change according to the task context. Cortex 81, 1–13.
Balaban, H., Luria, R., 2016b. Integration of distinct objects in visual working memory depends on strong objecthood cues even for different-dimension conjunctions. Cereb. Cortex 26 (5), 2093–2104. http://dx.doi.org/10.1093/cercor/bhv038.
Balaban, H., Luria, R., 2017. Neural and behavioral evidence for an online resetting process in visual working memory. J. Neurosci. 37 (5), 1225–1239. http://dx.doi.org/10.1523/JNEUROSCI.2789-16.2016.
Blaser, E., Pylyshyn, Z.W., Holcombe, A.O., 2000. Tracking an object through feature space. Nature 408 (6809), 196–199.
Brady, T.F., Konkle, T., Alvarez, G.A., 2011. A review of visual memory capacity: beyond individual items and toward structured representations. J. Vis. 11 (5), 4. http://dx.doi.org/10.1167/11.5.4.
Delorme, A., Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Methods 134 (1), 9–21. http://dx.doi.org/10.1016/j.jneumeth.2003.10.009.
Delvenne, J.F., Bruyer, R., 2006. A configural effect in visual short-term memory for features from different parts of an object. Q. J. Exp. Psychol. 59 (9), 1567–1580 (KX7458M2U0652586)(pii).
Drew, T., Horowitz, T.S., Wolfe, J.M., Vogel, E.K., 2011. Delineating the neural signatures of tracking spatial position and working memory during attentive tracking. J. Neurosci. 31 (2), 659–668. http://dx.doi.org/10.1523/JNEUROSCI.1339-10.2011.
Drew, T., Horowitz, T.S., Wolfe, J.M., Vogel, E.K., 2012. Neural measures of dynamic changes in attentive tracking load. J. Cogn. Neurosci. 24 (2), 440–450. http://dx.doi.org/10.1162/jocn_a_00107.
Drew, T., Vogel, E.K., 2008. Neural measures of individual differences in selecting and tracking multiple moving objects. J. Neurosci. 28 (16), 4183–4191. http://dx.doi.org/10.1523/JNEUROSCI.0556-08.2008.
Ecker, U.K., Lewandowsky, S., Oberauer, K., Chee, A.E., 2010. The components of working memory updating: an experimental decomposition and individual differences. J. Exp. Psychol.: Learn. Mem. Cogn. 36 (1), 170.
Emrich, S.M., Al-Aidroos, N., Pratt, J., Ferber, S., 2009. Visual search elicits the electrophysiological marker of visual working memory. PloS One 4 (11), e8042.
Fallon, S.J., Mattiesing, R.M., Dolfen, N., Manohar, S., Husain, M., 2018. Ignoring versus updating in working memory reveal differential roles of attention and feature binding. Cortex.
Gao, Z., Gao, Q., Tang, N., Shui, R., Shen, M., 2016. Organization principles in visual working memory: evidence from sequential stimulus display. Cognition 146, 277–288.
Hasher, L., Zacks, R.T., May, C.P., 1999. Inhibitory control, circadian arousal, and age.
Hillyard, S.A., Galambos, R., 1970. Eye movement artifact in the CNV. Electroencephalogr. Clin. Neurophysiol. 28 (2), 173–182.
Ikkai, A., McCollough, A.W., Vogel, E.K., 2010. Contralateral delay activity provides a neural measure of the number of representations in visual working memory. J. Neurophysiol. 103 (4), 1963–1968. http://dx.doi.org/10.1152/jn.00978.2009.
Kahneman, D., Treisman, A., Gibbs, B.J., 1992. The reviewing of object files: object-specific integration of information. Cogn. Psychol. 24 (2), 175–219.
Levillain, F., Flombaum, J.I., 2012. Correspondence problems cause repositioning costs in visual working memory. Vis. Cogn. 20 (6), 669–695.
Lopez-Calderon, J., Luck, S.J., 2014. ERPLAB: an open-source toolbox for the analysis of event-related potentials. Front. Human. Neurosci. 8, 213. http://dx.doi.org/10.3389/fnhum.2014.00213.
Luck, S.J., Vogel, E.K., 2013. Visual working memory capacity: from psychophysics and neurobiology to individual differences. Trends Cogn. Sci. 17 (8), 391–400. http://dx.doi.org/10.1016/j.tics.2013.06.006.
Luria, R., Balaban, H., Awh, E., Vogel, E.K., 2016. The contralateral delay activity as a neural measure of visual working memory. Neurosci. Biobehav. Rev. 62, 100–108.
Luria, R., Sessa, P., Gotler, A., Jolicoeur, P., Dell'Acqua, R., 2010. Visual short-term memory capacity for simple and complex objects. J. Cogn. Neurosci. 22 (3), 496–512. http://dx.doi.org/10.1162/jocn.2009.21214.
Luria, R., Vogel, E.K., 2011. Shape and color conjunction stimuli are represented as bound objects in visual `working memory. Neuropsychologia 49 (6), 1632–1639. http://dx.doi.org/10.1016/j.neuropsychologia.2010.11.031.
Luria, R., Vogel, E.K., 2014. Come together, right now: dynamic overwriting of an object's history through common fate. J. Cogn. Neurosci. 26 (8), 1819–1828. http://dx.doi.org/10.1162/jocn_a_00584.
Ma, W.J., Husain, M., Bays, P.M., 2014. Changing concepts of working memory. Nat. Neurosci. 17 (3), 347–356 (10.1038/nn.3655)(doi).
McCollough, A.W., Machizawa, M.G., Vogel, E.K., 2007. Electrophysiological measures of maintaining representations in visual working memory. Cortex; J. Devoted Study Nerv. Syst. Behav. 43 (1), 77–94.
Oberauer, K., Lewandowsky, S., Farrell, S., Jarrold, C., Greaves, M., 2012. Modeling working memory: an interference model of complex span. Psychon. Bull. Rev. 19 (5), 779–819.
Peterson, D.J., Gözenman, F., Arciniega, H., Berryhill, M.E., 2015. Contralateral delay activity tracks the influence of gestalt grouping principles on active visual working memory representations. Atten., Percept., Psychophys. 1–14.
Pylyshyn, Z.W., 2000. Situating vision in the world. Trends Cogn. Sci. 4 (5), 197–207.
Simons, D.J., Chabris, C.F., 1999. Gorillas in our midst: sustained inattentional blindness for dynamic events. Perception 28 (9), 1059–1074.
Tsubomi, H., Fukuda, K., Watanabe, K., Vogel, E.K., 2013. Neural limits to representing objects still within view. J. Neurosci. 33 (19), 8257–8263. http://dx.doi.org/10.1523/JNEUROSCI.5348-12.2013.
Vogel, E.K., Machizawa, M.G., 2004. Neural activity predicts individual differences in visual working memory capacity. Nature 428 (6984), 748–751. http://dx.doi.org/10.1038/nature02447.
Vogel, E.K., McCollough, A.W., Machizawa, M.G., 2005. Neural measures reveal individual differences in controlling access to working memory. Nature 438 (7067), 500–503 (nature04171)(pii).
Ye, C., Zhang, L., Liu, T., Li, H., Liu, Q., 2014. Visual working memory capacity for color is independent of representation resolution. PloS One 9 (3), e91681.
Zhang, W., Luck, S.J., 2008. Discrete fixed-resolution representations in visual working memory. Nature 453 (7192), 233–235. http://dx.doi.org/10.1038/nature06860.