Illusory conjunctions reflect the time course of the attentional blink

Juan Botella · Jesús Privado ·
Beatriz Gil-Gómez de Liaño · Manuel Suero

Published online: 9 March 2011
© Psychonomic Society, Inc. 2011

Abstract Illusory conjunctions in the time domain are binding errors for features from stimuli presented sequentially but in the same spatial position. A similar experimental paradigm is employed for the attentional blink (AB), an impairment of performance for the second of two targets when it is presented 200–500 msec after the first target. The analysis of errors along the time course of the AB allows the testing of models of illusory conjunctions. In an experiment, observers identified one (control condition) or two (experimental condition) letters in a specified color, so that illusory conjunctions in each response could be linked to specific positions in the series. Two items in the target colors (red and white, embedded in distractors of different colors) were employed in four conditions defined according to whether both targets were in the same or different colors. Besides the U-shaped function for hits, the errors were analyzed by calculating several response parameters reflecting characteristics such as the average position of the responses or the attentional suppression during the blink. The several error parameters cluster in two time courses, as would be expected from prevailing models of the AB. Furthermore, the results match the predictions from Botella, Barriopedro, and Suero's (*Journal of Experimental Psychology: Human Perception and Performance*, 27, 1452–1467, 2001) model for illusory conjunctions.

Keywords Illusory conjunctions · Attentional blink · RSVP · Binding errors

Illusory conjunctions in the time domain are binding errors produced as migrations of features between stimuli presented sequentially, in the same spatial position, by means of the so-called rapid serial visual presentation (RSVP; see, e.g., Botella, Barriopedro, & Suero, 2001). RSVP has also been the tool for studying the attentional blink (AB), an impairment of performance for the second of two targets (T2) when presented within 200–500 msec from the first target (T1). Explanatory models for both effects, illusory conjunctions and the AB, must be compatible; we will show in this article that, in fact, they are compatible.

Several models have been proposed to account for AB. The core ideas of those models include inhibition (Raymond, Shapiro, & Arnell, 1992), interference (Shapiro, Raymond, & Arnell, 1994), bottlenecks (Chun, 1997b; Chun & Potter, 1995; Dux & Harris, 2007a, 2007b; Jollicœur, 1998; Jollicœur, Dell'Acqua, & Crebolder, 2001), temporary loss of control (Di Lollo, Kawahara, Ghorashi, & Enns, 2005), attentional boost and temporary suppression (bounce) of items by way of an input filter (Olivers & Meeter, 2008), and impaired ability for individuating tokens from simultaneous processing of types (Bowman & Wyble, 2007; Wyble, Bowman, & Nieuwenstein, 2009), among others (see Dux & Marois, 2009, for a review).

Since the term *attentional blink* was first used by Raymond et al. (1992), the theoretical models have evolved to accommodate the increasing database of results from hundreds of experiments intended to elucidate the necessary and sufficient conditions for producing an AB. Special attention has been paid to the role of the T1 + 1 stimulus (Olivers & Meeter, 2008) or the awareness of T1 (Nieuwenstein, Van der Burg, Theeuwes, Wyble, & Potter,

J. Botella · B. Gil-Gómez de Liaño · M. Suero
Autónoma University of Madrid,
Madrid, Spain

J. Privado
Complutense University,
Madrid, Spain

J. Botella (✉)
Facultad de Psicología, Universidad Autónoma de Madrid,
Ciudad Universitaria de Cantoblanco, Campus de Cantoblanco,
c/ Ivan Pavlov, 6,
28049 Madrid, Spain
e-mail: juan.botella@uam.es

2009). The size of the blink has been found to be linked to factors such as the discriminability between the targets and the distractors or the degree of heterogeneity of the distractors.

Considerable efforts have been also made to grasp the very nature of the impairment produced—in the terms introduced above, the differences between the *cognitive scenarios* when T2 arrives during the blink or in the control conditions. Among others, the proposals include impaired perceptual coding, serial consolidation of items in short-term memory, limited capacity of working memory or of the ability to individuate tokens from types, and even involuntary oscillations in the attentional enhancement of the items (Dux & Marois, 2009).

Despite those differences, the view generally accepted for the AB includes at least the following ideas: (1) The AB is an attentional (not a perceptual) effect; (2) the detection of T1 (or perhaps the couple T1 + Filler 1) triggers some processes directed to yield a response to it, and if T2 arrives before those processes have finished, the processing of T2 is different (poorer performance), as compared with what happens in the control conditions; (3) the impairing effect of those processes vanishes gradually and disappears with long T1–T2 lags, so that with long enough lags, the performance in the two-target condition is indistinguishable from performance in the control conditions (without any T1).

The AB has been usually analyzed by means of the hit rate for the T2, with two targets embedded within an RSVP. Specifically, it is characterized by the shape of the function that relates the percentage of hits for T2 (conditionalized to a hit for T1) and the T1–T2 lag. Less attention has been paid to the nature of the errors and how they change along the time course of the AB. The analysis of errors should reach conclusions converging with those from the standard analysis of the performance with T2, given a hit with T1. At least, they should be compatible with the view described above.

However, the few experiments in which the nature of the errors along the time course of the AB has been studied have yielded results that have sometimes been at odd among themselves and with the dominant view for the AB. Our main purpose in this research was to show how the view of the AB that errors yield can be accommodated by the dominant view of the AB. In the next section, we review the experiments in which the illusory conjunctions during the AB have been studied. Then we will highlight the discrepancies between those results and the dominant view of the AB and how they can be experimentally addressed.

Illusory conjunctions during the attentional blink

In many AB experiments, the task for T2 has been detecting a specific stimulus (e.g., an X from a stream of

letters), so that each response could be categorized only as correct or incorrect and there was no room for any analysis of the type of errors. When the task has been one of identifying T2, most times the responses have been registered just as correct or incorrect; again, the recorded responses cannot be analyzed for errors, beyond the analysis of how frequent they are (and this is the complement of the hit rate).

Alternative tasks where the observers can experience illusory conjunctions (e.g., Botella et al., 2001; Botella & Eriksen, 1992; McLean, Broadbent, & Broadbent, 1983) allow for a richer analysis of errors. In those tasks, the response is not only a present-absent one. The observers must identify the target item, so that the origin of an item mistakenly reported can be located along the series, relative to the target. For instance, consider a single-target task in which the observer is asked to report the only red letter in a series of colored letters (Fig. 1a; color is the target-defining dimension, and shape is the to-be-reported dimension). Reporting the S is categorized as a hit, but reporting F or B would be categorized as a +1 or –1 error, respectively. After a number of trials, the distribution of the origin of the reported letters could be as in Fig. 1b. In the experiment reported below, there were two targets, and the typical AB

Fig. 1 **a** Sequence of trial events in a task designed to yield illusory conjunctions. The several tones of gray represent different colors. Observers must identify the only letter in a specified color. **b** Typical frequency distribution of the origin of the responses for an observer, limited to the central set (target ± 2 positions in the series)

was observed. However, we studied not only the hit rate for T2, but also other types of responses. Thus, we studied the distribution of the responses for T2 and how it changed as a function of the lag. Among other goals, we attempted to demonstrate the impact the AB has on the amount and type of binding error. Studying how distributions such as that in Fig. 1b change as a function of the lag between the targets can help us to understand the AB.

A few experiments have employed several variations of such a procedure. In the first one, Chun (1997a) presented a series of 17 letters within white frames; the frames of the two target letters were red and green, respectively. The results showed that besides the well-known U-shaped function for the hit rate, the origin of the errors also changed along the range of lags where typically the AB is produced. Popple and Levi (2007) presented as the two targets two red (or gray) letters embedded within a series of gray (or red) letters. Again, the results showed both the typical AB and shifts of the origin of the reported letters along the AB. Vul, Hanus, and Kanwisher (2008; see also Vul, Nieuwenstein, & Kanwisher, 2008) presented a series with the 26 English letters, but the two designated targets were presented surrounded by a white ring. Once again, besides the expected AB found for the hit rate on T2, the origin of the reports systematically changed along the time course of the AB.

Although, in all of these experiments, important effects of the T1–T2 lag over the distribution of the origin of the responses have been reported, there are a number of conflicting results and inconsistencies between them that deserve further attention. Moreover, some results are not compatible with the general view for the AB depicted above.

Some conflicting results

Even though the conclusion that the origin of the responses (in terms of their positions with regard to T2) changes systematically along the AB is firmly established, there are a number of issues that deserve yet a closer look.

Since some of the inconsistencies in the results we are going to discuss could be due to important differences in the tasks used, we first highlight those methodological differences. For example, the variants of the task more frequently employed do not produce what we consider to be “genuine” illusory conjunctions of percepts. When a frame or a ring is presented around one letter and the observer must identify it, the task can be considered to be one related more to the psychophysics of temporal order. The observers must first detect the ring and then identify and retain T1; then the second ring must be detected, and finally, the identity of T2 must be processed and retained. Attention must shift 3 times between four different pieces of

information. Of course, errors are a kind of illusory conjunction in temporal order judgments, but the task involves too many different elements, and inferences about the source of any observed empirical effect are uncertain.

The most outstanding feature of the illusory conjunctions of single percepts is that the two relevant dimensions incorrectly joined are integral dimensions of the stimulus (Treisman & Gelade, 1980; see also Quinlan, 2003). It is difficult to perceive the shape without processing the color or vice versa. The two features are extracted from the same stimulus, but surprisingly, they migrate between items that are close in the series; this is the reason why we consider them as “genuine” illusory conjunctions.

In the experiments described above, both Chun (1997a) and Vul and colleagues (Vul et al., 2008a, 2008b) employed tasks of synchronicity categorization between composed stimuli. Only Popple and Levi (2007) employed integral dimensions of the stimuli as the two relevant features. Unfortunately, all filler letters were in the same color (gray or red), whereas the targets were in a different color (red or gray). The processing of targets embedded within homogeneous distractors can benefit from a *pop-out* effect. The conclusions from the only experiment with integral stimuli (Popple & Levi, 2007) are different from those from the experiments based on synchrony categorizations of separable elements of composed stimuli. Our first goal in this research was to address a more complete analysis of the results, employing a task more suitable for producing what we have called “genuine” illusory conjunctions. Specifically, we employed the shape and the color as integral dimensions. Furthermore, distractors of heterogeneous colors were used in order to avoid any pop-out effect. This task has been shown to be reliable for producing frequent and predictable illusory conjunctions (Botella & Eriksen, 1991; Botella, García, & Barriopedro, 1992).

With respect to inconsistent results, whereas, after a much more complete analysis, Vul et al. (2008b) concluded that the responses are suppressed, delayed, and diffused (three parameters of the distribution of the responses; see Fig. 1), Popple and Levi (2007) concluded that the AB effect can be understood simply as a redistribution of the origin of the responses around the target. That is, according to Popple and Levi (2007), the identities of the letters presented during the AB are equally available, but the order information is lost. In fact, they analyzed the responses including, progressively, the hits, the hits plus ± 1 responses, ± 2 responses, and ± 3 responses. In the last analysis (responses from positions between -3 and $+3$), the AB almost vanished. This conclusion contradicts that of Vul et al. (2008b), who concluded that during the AB, not only is the temporal order information lost, but also the identification of the items around the target is suppressed to some degree.

The way in which the suppression has been analyzed might be improved. It has been analyzed by computing the frequency of responses from positions far from T2; it has been assumed that those responses reveal a difficulty in identifying the letters during the AB episode. In the present experiment, we included two new features, as compared with previous experiments. The first one was that the observers could choose *don't know* as a response, instead of forcing them to make a guess or rerunning the trial at the end of the block. The second one was that we generated a series of 24 letters but the observers could choose any 1 from 26 letters for the response. In that way, we could assess the quality of the perceptual categorization, since a low quality would produce misidentifications and sometimes yield a report of a nonpresented letter (e.g., Popple & Levi, 2008, have shown that the amblyopic eyes generate more confusions of this type than do normal eyes). In experiments in which the complete set of letters is employed in all the series, errors cannot be disaggregated into order position errors and misidentification errors. In the experiment below, the rate of reports of nonpresented items was analyzed as a function of the T1–T2 lag. The attentional nature of the AB effect must be reflected in the fact that the rate of reporting of nonpresented items does not change as a function of the lag.

The results from the experiments above raise some additional intriguing questions. Thus, in Vul et al.'s (2008b) experiment, the delay of the responses to T2 did not recover the values shown by the responses to T1, even when the lag between them was as long as 10 positions (more than 800 msec). On the contrary, in the experiment by Chun (1997a), the delay reached that level, whereas the data from Popple and Levi (2007) were not conclusive about this specific point. Another of our goals was to check in our experiment whether the delay at long lags was similar to that for T1; if not, the discrepancy could be taken as an effect of the type of task employed ("genuine" illusory conjunctions vs. synchronicity tasks).

Another intriguing result has to do with the several time courses found within each experiment. In Chun's (1997a) experiment, the results were quite conclusive as to the fact that the AB and the average delay of the reported items followed different time courses (the peak was at lag 2 for the first one but was between lags 3 and 4 for the second one). Popple and Levi (2007) did not make such comparisons. However, Vul et al. (2008b) made much richer analyses and found cues of processes that were influenced by the AB that followed three different time courses (well represented by the three statistics of the distribution of the responses reflected in the delay, the suppression, and the diffusion). That was surprising, because what would be expected from virtually all theoretical models of the AB is that it must be possible to group the processes involved in

two different time courses. If the delay at long lags actually reaches the level shown by the responses to T1, the processes could be grouped around only two time courses, as would be expected from the theoretical models.

The two-stage model of the formation of illusory conjunctions (Botella et al., 2001) allows making specific predictions about the time course of the response parameters. According to that model, when the target-defining feature is detected, an attempt to focus attention on the target is made. With the typical stimulus onset asynchronies employed in these tasks, this attempt succeeds on some percentage of the trials, and then the target is completely and accurately identified (a hit is recorded). However, on the other percentage of trials, the system does not finish quickly enough the process of focalization, and a second attempt to generate a response takes place. This second attempt is based on partial, degraded information; essentially, it is based on isolated floating features. The system makes a sophisticated guess related to the strength of the activation of each feature at the critical moment (see Botella et al., 2001). In terms of this model, the AB reduces the probability of a successful focalization of attention on T2 and makes it more probable that the response is based in the second attempt. Furthermore, the critical moment that determines the degree of activation on which the guessing process is based is delayed. Since the hit rate and the average delay of the intrusions are the primary indices of those attempts, the prediction is that the lag at which the delay peaks will be longer than the lag at which the hit rate peaks (the AB itself).

The following experiment was intended to clarify the puzzling results described above. We present some new results based on the combinations of the patterns of both hits and errors, and their time courses, in a task that produces "genuine" illusory conjunctions. We expected that when the results were analyzed within the framework of Botella et al.'s (2001) model, they would be compatible with the general view of the AB, reached with the standard analyses based on the hit rates.

Experiment

According to the discussion above, the main goal of the experiment was to show how the evidence found when the errors are analyzed is compatible with the general view of the AB, on the basis of the analysis of the hit rate. Specifically, our predictions were that (1) the average delay of the errors would peak later than the hits; (2) other response parameters would cluster with those two, so that only two time courses would be identified; (3) the response parameters would recover the "baseline" (control condition level) with the long lags; and (4) a proper analysis of the

reports of nonpresented items would reflect the attentional nature of the effect.

In the pilots for the experiment, we tried several colors for the targets and found that performance with the red and white targets was statistically different: White targets produced reliably worse performance. We thought that this was an opportunity for additional comparisons, defining four experimental conditions: T1 red and T2 white, T1 white and T2 red, both targets red, and both targets white. The rationale was that if the AB was produced because, when T2 arrived, the cognitive system was still under the influence of the episode triggered by T1 (or T1 + Filler 1), the observed effects in the several response parameters would reflect how loading and lasting the episode associated to T1 was. As a consequence, if white targets were more difficult to process, when T1 was a white target, the response parameters for T2 should reflect poorer performance and a longer effect, as compared with when T1 was red. In the same vein, for the same T1 color, the level of performance with a white T2 should be worse than that with a red T2. The differences between the four experimental conditions should be reflected not only in the hit rate, but also in the other response parameters based on errors. In short, we expected that the predictions formulated above would show in different degrees in the four experimental conditions, according to their differential difficulty.

Method

Participants Forty-four observers participated in the experiment, distributed in four groups. These were undergraduate students of the Autonomous University of Madrid, who received €15 for participating in two sessions of about 45 min. All the observers reported normal or corrected-to-normal visual acuity and normal color vision. They were not aware of the purpose of the experiment. Informed consent was obtained from all of them.

Materials and design Each trial consisted of a series of 24 different letters randomly taken from the pool of 26 letters. As is explained below, the other 2 letters were employed to estimate the probability of reporting a nonpresented letter. The letters were presented in four different colors—green (RGB: 173,255,47), black (RGB: 0,0,0), yellow (RGB: 255,255,0), or blue (RGB: 0,0,255),—also taken at random, with the only restriction that two consecutive letters never shared the same color. The letter designated as T1 was equiprobably that in position 8, 9, or 10. The letter designated as T2 was that in positions T1 + 1 to T1 + 9. The colors of the designated targets were changed to red (RGB: 255,0,0) and/or white (RGB: 255,255,255), depending of the experimental condition. The experimental conditions

(RR, WW, RW, WR) were defined according to the colors of the targets (both red, both white, T1 red and T2 white, or T1 white and T2 red). The background remained gray throughout the experiment.

Procedure Each observer participated in two sessions. In the first one, they were instructed to report only the letter corresponding to T2 (control condition), whereas in the second session, they were instructed to report both T1 and T2 (experimental condition). The control condition for the RW and WR groups consisted of reporting the white (RW) or the red (WR) letter, but the series were identical to those used in their experimental condition. However, in the RR and WW conditions, it was not possible to employ the same materials as in the corresponding experimental condition. The reason is that the participants would be asked to report the second of two letters in the target color, and that would not be a real one-target control condition. That is why we also employed the series of the WR and RW groups in the control conditions for the RR and WW groups.

Each trial began with a cross for fixation, appearing for 500 ms at the center of the screen. One second after the fixation went off, the stream of letters appeared successively, without interstimulus blanks, at the same location for 83 ms each. After the last letter disappeared, the observers had to make one (control condition) or two (experimental condition) responses, pressing the letter(s) identified as the target(s), in the order in which they were perceived. They could also make a *don't know* response, pressing the space bar on the keyboard. No feedback was given for accuracy. The experimental trials were preceded by a practice block of 20 trials. The control condition was always administered in the first session, and the experimental condition in the second one. The experimental condition consisted of 60 trials with each of the nine T1–T2 lags, totaling 540 experimental trials, and 20 practice trials (randomly from the several lags). The control condition consisted of the same series, although the observers had only to report T2.

Apparatus The program for running the experiment was written by the authors in VisualBasic. The experiment was run in individual cabins equipped with PC-compatible computers and monitors with a resolution of 1,280 × 600 pixels; the participants sat at a 40-cm viewing distance from the screen, from which the letters of the series had a visual angle of 0.46° in height and 0.29° in width.

Results

This section is organized into four parts. In the first one, the differential performance for the two target colors is highlighted, since that difference becomes a tool for subsequent analyses. Then the performance with T1 and

its relationship with task switching between the targets are analyzed. In the third part, the AB itself is analyzed in the traditional way; a mathematic model is also fitted in order to make more accurate comparisons between the experimental conditions. In the last part, the several types of errors are analyzed as a function of the experimental condition and the T1–T2 lag. When indicated, the Bonferroni post hoc comparison procedure has been applied.

Differential performance for red and white targets

Differential performance for different target colors has been previously reported (e.g., Popple & Levi, 2007, with the report better performance for red than for gray targets). As was expected from the pilots, in the present experiment, the performance was better with red targets than with white targets. This conclusion can be statistically supported by analyzing the results of the control conditions. A 4×9 (group $\times$ lag) ANOVA on the hits rates in the control condition revealed a main effect of group, $F(3, 37) = 24.608$, $MSE = 2.241$, $p < .001$, but not of lag, $F(8, 296) = 1.191$, $MSE = 0.005$, $p = .304$, and no interaction, $F(24, 296) = 0.907$, $MSE = 0.004$, $p = .593$. Post hoc comparisons confirmed that performance in the control conditions was better with the red target letters (WR and RR groups) than with the white target letters (RW and WW groups) (83.0% vs. 56.4%, on average; see Fig. 2).

This effect was probably due to a higher contrast between the specific red used and the other colors employed for the distractors than between the white and those same filler colors. As was anticipated above, this differential difficulty is useful for later analyses, since it becomes predictive for ordering the conditions according to several response parameters.

Performance for T1

A main issue related to performance for T1 deserves being highlighted: Performance with T1 is impaired if there is a task switch between the targets. For both red and white T1s, performance with T1 is better when T2 is in the same color rather than in a different color, even for the longest lags (see Fig. 2). This is not a typical task-switch effect (see Kiesel et al., 2010; Koch, Gade, Schuch, & Philipp, 2010; Monsell, 2003), since it does not decrease as the T1–T2 delay increases. The switch in the target colors between T1 and T2 has an influence on T1 because the response to T1 is not made until the end of the trial. Our interpretation of this effect is that in the WR and RW conditions, the observer must perform an additional operation after detecting T1 that is not necessary in the WW and RR conditions: a change in the cognitive set when waiting for a target with a color different from that of T1. The cost of such an operation is

that the identity of T1 is sometimes lost, and this is reflected in an extra impairment of performance with T1. However, the lag between the targets does not matter, since it has an influence on the probability of losing T1 from short-term memory and the switch in the cognitive set is done as soon as T1 is detected, no matter how long the arrival of T2 takes. On the contrary, that operation is not done in the nonswitch conditions WW and RR (and of course, in the control conditions), so this source of loss is absent.

The drop of performance for lag1 in the WW and RR conditions (see Fig. 2) is due to a higher rate of permutation errors—that is, trials on which the identities of T1 and T2 are correctly reported, but in the wrong order (Juola, Botella, & Palacios, 2004).

Attentional blink

The four groups showed the typical U-shaped AB. Fig. 3 shows the hit rates for T2, conditionalized to a hit on T1, graphed jointly with the hit rate in the control condition for each group. In all of them, the AB shows first the so-called lag 1 sparing, then the deepest zone of the effect in position 2 (conditions RW, RR, and WR) or positions 2 and 3 (condition WW), and finally performance growth, gradually, until it reaches the level of the control condition in a temporal window between lags 5 and 6. A 4×9 (group $\times$ lag) ANOVA on the hit rates for T2, conditionalized to a hit on T1, reveals main effects of group, $F(3, 37) = 17.308$, $MSE = 1.804$, $p < .001$, and lag, $F(8, 296) = 136.2$, $MSE = 1.749$, $p < .001$, but, more important, also of the interaction, $F(24, 296) = 4.431$, $MSE = 0.057$, $p < .001$. The significant main effect of lag, jointly with the shape of the averages shown in the figure, are by themselves the AB effect. The interaction is due to the fact that the time course of the AB varies between the conditions. When the conditions that share the same T2 color target are compared post hoc, the conditions with the more difficult color target in T1 (white) show lower performance in the *recovery* segment of the functions. Specifically, the performance in the WW condition is significantly worse than that in the RW condition for lags 4 and 5 ($p < .01$; no significant differences for the other lags). In the same vein, performance in the WR condition is significantly worse than that in the RR condition for lags 4, 5, 6, and 7 ($p < .01$; no significant differences for the other lags). It must be highlighted that there are no significant differences between the experimental conditions in the lag 2 condition. This could be due to a floor effect, although it is clear that the deepest effect is observed at lag 2 for all the experimental conditions.

In order to better characterize the AB functions and make more accurate comparisons, we empirically adjusted a set of mathematical models, inspired by some previous

Fig. 2 Identification performance for the first target (T1) in the four two-target experimental conditions and the two control one-target conditions, plotted across the T1–T2 lags. The data points for the control conditions are averaged for the two groups that shared the same control condition (T2R and T2W)

models (e.g., Olivers & Meeter, 2008). In the model that fits the data better (see the Appendix), the probability of a correct report of T2, given a correct response for T1, is defined as

$$f(\text{lag}) = d \cdot (1 - (a \cdot (\text{lag} - b)^2) \cdot e^{-c \cdot \text{lag}}). \quad (1)$$

In this model, the parameter d is estimated as the average performance in the one-target condition. As lag increases, performance approaches the asymptotic value set by this parameter. The model allows one to characterize the functions according to several special values. Table A1 in the Appendix includes the parameters and several main

values for the four two-target conditions, together with the R^2 fit index. The functions are plotted in Fig. 4.

The main results are as follows. First, the lag at which the minimum is reached ($\text{lag}_{\min}$) is shorter for an easy (red) T1 than for a difficult (white) T1, when the conditions with the same T2 color are compared. Furthermore, comparing the conditions that share the same T1 color, the minimum is again reached earlier for an easy (red) T2 than for a difficult (white) T2.

Second, the AB itself is deeper [smaller $f(\text{lag}_{\min})$] for a difficult (white) T2 than for an easy (red) T2. As above, comparing the conditions that share the same T2 color, the

Fig. 3 a Identification performance for the first and second targets (T1 and T2) for the four groups, plotted across the T1–T2 lags. The results for T2 performance are based on trials on which T1 was correctly identified, T2/T1. **b** Combined figure for the T2/T1 results of the four experimental conditions

Fig. 4 Functions fitted for the four experimental conditions according to Equation 1. They relate the performance in identifying T2, given a hit in identifying T1, as a function of the T1–T2 lag

minimum is again deeper with a difficult (white) T1 than with an easy (red) T1.

Third, the lag at which performance recovers a criterion level of 95% of the performance in the control condition (column lag95 in Table A1) is shorter for an easy (red) T1 than for a difficult (white) T1. Furthermore, comparing the conditions that share the same T1 color, the criterion is again reached earlier for an easy (red) T2 than for a difficult (white) T2.

In short, the ABs produced with easy and difficult targets (for both T1 and T2) shows predictable differences in their time courses. As compared with easy targets, the difficult targets show a deeper AB that reaches the minimum performance later and recovers the performance of the control condition later. In the next section, we check whether several indexes based on the errors also reflect this consistent effect of the differential difficulty of the target-defining colors.

Analysis of errors in T2, given a hit in T1

Errors in T2 were analyzed by calculating a set of response parameters that reflect several facets of information processing during the task. But a caution must be noted before the analysis of errors. Suppose that the total frequency of responses coming from positions neighboring T2 (e.g., from positions between ± 2) is calculated for each T1–T2 lag. If this value is calculated conditionalized to a hit for T1, there is a bias for lags +1 and +2. For lag 2 the item from position –2 has been already reported correctly for T1, and the same happens for the item from position –1 in the lag 1 condition. Then, the probability of reporting an item from positions ± 2 is reduced in those two conditions, since it can come only from four positions, instead of five. If this is not taken into account, the comparison between the lag conditions will be biased against lag 1 and lag 2 conditions. Nevertheless, in the analysis of most of the response parameters, we do not need to include the

conditions of lags +1 and +2 to reach a conclusion. In those cases, the statistical analyses will be restricted to lag 3 or longer.

Delay We calculated for each observer the average position of the letters reported as T2, conditionalized to a hit on T1, as a function of lag. Taking only lag 3 and higher lags, a 4×7 (group $\times$ lag) ANOVA on the average of the positions of the responses to T2 (within $T2 \pm 2$), conditionalized to a hit on T1, reveals main effects of group, $F(3, 37) = 9.804$, $MSE = 1.039$, $p < .001$, and of lag, $F(6, 222) = 9.718$, $MSE = 0.392$, $p < .001$, and also of the interaction, $F(18, 222) = 1.673$, $MSE = 0.067$, $p < .05$. The shape of the functions in Fig. 5 mirrors that of the AB, although shifted to the right. In fact, the highest delay is observed at lag 4 in the four experimental conditions; as was expected, the peaks of the delay are produced later than the peaks of the AB (lag 2 for three conditions [RR, RW, and WR] and between lag 2 and lag 3 for one condition [WW]). Furthermore, the delay of the responses is generally higher for the condition with a larger AB (WW group) and lower for the condition with a smaller AB (RR group). In short, the temporal course of the delay is different from that of the AB, although they are closely related (the delay is larger as the AB deepens).

The values for the largest lag (lag 9) are close to those for their corresponding control conditions. This result is different from the conclusion of Vul et al. (2008a; although they did not compare the results with a control condition, since they did not include such a condition; instead, the comparison was made against the performance with T1).

Suppression Suppression was operationalized in two ways. First, we calculated for each observer the percentage of responses to T2 that were letters from the central set ($T2 \pm 2$), conditionalized to a hit on T1, as a function of lag. Again, we took only lag 3 and later positions, for the same reason as that explained above. A 4×7 (group $\times$ lag) ANOVA on the percentage of $T2 \pm 2$ responses to T2, conditionalized to a hit on T1, reveals main effects of group, $F(3, 37) = 14.274$,

Fig. 5 Delay of the responses to T2, plotted across the T1–T2 lags, for the four experimental conditions. It is based on trials on which T1 was correctly identified. The delay is calculated as the average position of the responses (APR) within $T2 \pm 2$. Given that there are many correct responses (delay = 0), the typical average delay usually is between –0.5 and +1 (see Botella et al., 2001; Botella & Eriksen, 1991; Botella et al., 1992). The values for C are the averages in the control conditions

$MSE = 0.562$, $p < .001$, of lag, $F(6, 222) = 104.08$, $MSE = 0.907$, $p < .001$, and the interaction, $F(18, 222) = 4.158$, $MSE = 0.036$, $p < .001$. The shape of the functions in Fig. 6a mirrors that of the AB. The higher suppression is produced around lag 2, that of the maximum AB; furthermore, suppression is higher for the condition with a larger AB (WW) and lower for the condition with a smaller AB (RR). The time course of the suppression is the same as that of the AB, although different from the time course of the delay.

As was explained above, we also took the percentage of *don't know* responses as a second index of suppression. We calculated for each observer the percentage of such responses to T2, conditionalized to a hit on T1, as a function of lag. A 4×9 (group $\times$ lag) ANOVA on those percentages reveals main effects of group, $F(3, 37) = 6.086$, $MSE = 0.551$, $p < .01$, of lag, $F(8, 296) = 33.944$, $MSE = 0.444$, $p < .001$, and of the interaction, $F(24, 296) = 2.951$, $MSE = 0.039$, $p < .01$. The shape of the functions in Fig. 6b mirrors again that of the AB. The higher suppression is produced for lag 2, which is the lag with the maximum AB for three conditions (RR, RW, and WR), and is the same for lag 2 and lag 3 for one condition (WW). Furthermore, suppression is higher for the condition with a larger AB (WW) and lower for the condition with a smaller AB (RR). The time course of the suppression is the same as that of the AB, although different from the time course of the delay.

Diffusion Following Vul et al. (2008a), we calculated for each observer the variance of the positions of the letters reported for T2, conditionalized to a hit on T1, as a function

Fig. 6 Suppression of the responses to T2, plotted across the T1–T2 lags, for the four experimental conditions. It is based on trials on which T1 was correctly identified. The suppression is assessed in two ways: a as the total percentage of the responses originated within the $T2 \pm 2$ interval and b as the percentage of *don't know* responses

of lag. As above, we took only lag 3 and later positions. A 4×7 (group $\times$ lag) ANOVA on the variances of $T2 \pm 2$ responses to T2, conditionalized to a hit on T1, reveals main effects of group, $F(3, 36) = 3.792$, $MSE = 0.456$, $p < .02$, of lag, $F(6, 216) = 26.344$, $MSE = 2.021$, $p < .001$, and of the interaction, $F(18, 216) = 2.411$, $MSE = 0.185$, $p < .01$. As is shown in Fig. 7a, the data are uncertain as to when the maximum diffusion is produced, but it is around lag 2. However, what seems clear is that the maximum diffusion is not produced in lag 4, as in the delay. What seems more plausible is that the time course of this parameter is the same as that of the AB and the suppression parameters.

Confusions Confusions were measured as the percentage of responses reporting a letter not presented in the series, conditionalized to a hit on T1, as a function of lag. A 4×9 (groups $\times$ lags) ANOVA on those percentages reveals no significant main effects of group, $F(3, 37) = 0.310$, $MSE = 0.001$, $p = .818$, or of lag, $F(8, 296) = 1.487$, $MSE = 0.001$, $p = .161$, and no interaction, $F(24, 296) = 1.468$, $MSE = 0.001$, $p = .076$. Apparently, confusions are not significantly involved in the AB effect (Fig. 7b).

Discussion

Taken as a whole, the results of the present experiment converge with the general view of the AB when the

Fig. 7 Diffusion and confusions in the responses to T2, plotted across the T1–T2 lags, for the four experimental conditions. They are based on trials on which T1 was correctly identified. a Diffusion is calculated as the variance of the positions of the responses within $T2 \pm 2$. b Confusions are calculated as the percentage of reports of items not presented in the series

frequencies of illusory conjunctions and its distribution are analyzed, taking the model of Botella et al. (2001) as a framework.

The hit rates show the typical U-shaped function of the AB in the conditions created from the combinations of two targets with an easy (red) and/or a difficult (white) color. Furthermore, the depth and the time course of the AB functions are predictable from those levels of difficulty.

More important for our goals, the set of response parameters employed to analyze the responses produced during the AB (including hits, several types of errors, and misses) show that they can be characterized by two different time courses. The first one has its peak at about lag 2 and includes the AB (T2/T1) per se, the two parameters of *suppression* (the *don't know* responses and the availability of the items presented during the blink), and, probably, the *diffusion* (variability of the positions of the items reported). The second group includes only the *delay* (average position of the items reported). The *confusions* parameter does not show any effect at all.

Roughly speaking, the first group of response parameters and its time course reflect what does not happen during the AB (or happens less), while the delay and its separate time course reflect what actually happens. On one side, during the AB, it does not happen (or happens less) that the observer identifies T2 at the level of consciousness for a voluntary report. Fewer hits and more suppression are found as observable consequences of that impairment in processing T2. There is less successful processing and reporting of T2, as compared with the control conditions. However, on the other side, it is also shown that another attempt to give a response, probably based on partial and degraded information (isolated features, or types), is triggered after the first attempt fails. This second attempt happens more frequently when T2 is presented during the blink, as compared with the control conditions. The average delay returns to a size similar to that for the control condition when the T1–T2 lag is so long that any AB effect has vanished. The idea of two sequential attempts to generate a response is at the core of the model of illusory conjunctions of Botella et al. (2001).

Although the two time courses are different, they are not independent. The WW condition shows the largest impairment, as reflected in the parameters of the first group, whereas the RR condition shows the smallest impairment (longer and deeper blink and larger suppression for the WW condition; see Figs. 3 and 6). As was expected, the same happens to the parameter with a different time course (the largest delay is also observed for the WW condition, whereas the smallest is again observed for the RR condition; see Fig. 5). Apparently, the delay and its time course are time dependent on the time course of the first

group of processes. The delay reflects processes that are triggered after the processes reflected in the AB.

At the longest lag (lag 9), the average delay of the responses recovers the values reached in the control condition. This result contradicts that of Vul et al. (2008a), who found a persistent delay effect with even longer lags. This difference is probably due to the fact that we employed a task that produced “genuine” illusory conjunctions, instead of a cuing task. In tasks in which one stimulus works as a cue for reporting another stimulus, the migrations are allowed in time but also involve shifts of attention in the space domain (Reeves & Sperling, 1986).

Our results are not a by-product of the specific task employed, since we found many results familiar from AB experiments. Besides the standard AB depicted in Fig. 3, we found that the duration and size of the AB is a function of the processing load from both targets. Generally speaking, the task is more difficult for white targets than for red targets (as revealed by comparing the levels of performance in the control, one-target conditions; see Fig. 2). As a consequence, for a given T2 color (white or red), performance is worse when T1 is white than when it is red (the AB is longer and/or deeper). Simultaneously, for a given T1 color (white or red), performance is worse when T2 is white than when T2 is red (again, the AB is longer and/or deeper).

Furthermore, we found that a task switch between T1 and T2 impairs performance on T1. As is shown in Fig. 2, this happens for both colors. In our experiment, the two targets were reported after the series had finished, so that, probably during the switch, the identity of T1 was sometimes lost. When both targets were in the same color, there was no task switching, and this was not a source of additional errors with T1.

Illusory conjunctions and the AB

As was sketched in the introduction, Botella et al. (2001; Botella, Narváez, Suero, & Juola, 2007) proposed a two-stage model of the formation of illusory conjunctions produced as migrations in the temporal axis. In the first stage, the detection of the target-defining feature triggers an episode of focalization of attention on the target. At presentation rates of 8–12 items per second, this first attempt frequently succeeds. However, when it fails, a second attempt is made (the second stage of the model) on the basis of partial information (essentially, isolated features; Quinlan, 2003; Treisman & Gelade, 1980).

From the perspective of that model, what happens during the blink is that selective attention to T2 is, on average, delayed, an idea previously discussed by several authors (Chun, 1997a; Vul et al., 2008a). The consequences are that fewer trials yield a reportable response on the first attempt,

more trials are finished with a response from the second attempt, and more trials do not yield any reportable response, increasing the rate of *don't know* responses. The processing of T1 and the episode triggered by the +1 distractor suppress several items, which do not enter into short-term visual memory; the suppression includes both T2 (when presented in positions close to T1) and the distractors in the positions immediately following T1. As a consequence, the hit rate is reduced (the AB itself), but other items are also suppressed. Of course, the temporal window of that episode varies as a function of the processing load for T1. Although the episode is triggered by the unexpected entering of a distractor into visual short-term memory (the +1 distractor), once started, its duration is a function of the duration of the processing of T1.

The items that survive the episode are not completely processed; specifically, what survive are often the isolated features that compose them, because attention has not had the opportunity to convert them from types to tokens (Kanwisher, 1991). The second stage of Botella et al.'s (2001) model consists of a *sophisticated guessing* from the relative activation of the representation of the isolated features. Those relative activations are converted into candidates translating them into discrete states (Vul, Hanus, & Kanwisher, 2009). This attempt begins as soon as the first attempt fails, and failures are more frequent during the blink. Its later beginning is reflected in the fact that the maximum delay is reached around lag 4, later than the maximum AB.

One of our results also converges with the generally assumed idea that the AB is actually an attentional (not a perceptual) effect. The frequency of items reported that were not presented in the series was practically zero and independent of the lag. If the quality of a perceptual analysis is impaired, it is expected that, sometimes, the items will be miscategorized (Poppel & Levi, 2008). As a consequence, the frequency of nonpresented but reported items should be higher for lags with a deeper AB effect (lags 2 and 3). Our results show that observers prefer to give a *don't know* response to a T2 presented during the blink than to make what they probably consider a pure guess (pure guesses would be reflected in a higher probability of reporting a nonpresented item).

The picture that emerges from our results matches the general view depicted in the introduction: The AB is triggered by the detection of T1 (and probably filler 1) and disappears gradually as the T1–T2 lag increases. Our results are compatible with several alternative explanations of the mechanisms involved. Of course, the relationship between the processing load for T1 and the length and size of the episode demonstrates that what interferes during the AB episode are the processes associated with T1 (Raymond et al., 1992; Shapiro et al., 1994). The rate of suppression indicates that the effect of the interference is that the items

following T1 are not properly entered (or consolidated) into short-term visual memory (Chun & Potter, 1995). However, our results are also compatible with the idea that the episode is triggered by the +1 item and that it consists in oscillations of the attentional enhancement of the items (Olivers & Meeter, 2008). Furthermore, the effect of a reduced attentional impairment is that the response is based on partial information (feature types), instead of reportable individuated tokens (Bowman & Wyble, 2007; Vul et al. 2008a). Apparently, all those models have elements will prove valuable for a more general, comprehensive view of the temporal dynamics of attention.

However, what is new in the experiments that have analyzed errors made during the AB is that they help us to get a more comprehensive view of this effect. As viewed from a model such as that of Botella et al. (2001), it can be concluded that the AB deeply impairs the first attempt to yield a response (via focal attention). More responses are produced as a second attempt, based on partial, degraded information and prone to generate illusory conjunctions. Those responses are not pure guesses. They mostly involve features that have, indeed, been presented in the series, in positions close to the target, and, on average, have come from positions later than when the target is presented outside the AB window.

Author Notes We thank Julio Lillo for his assistance in managing the colors and Jim Juola for his help with the manuscript. Parts of the results were presented at the 50th Annual Meeting of the Psychonomic Society, Boston, MA.

Notes 1 One of the reviewers highlighted the risks of making such a strong statement, since, in fact, we do not know the average delay for lags 1 and 2. Although this is true, it is also true that in all the experiments and conditions run in our laboratory, we have always found one-mode distributions. Even assuming that we are making some inference (we cannot check it directly), what we are saying is very probably true. In three of the conditions, the function show higher values for lag 4 than for lag 3, and in the other condition, there is no difference between them. Assuming one-peak functions, we can state safely that the maximum delay is produced later than lag 2. We have made similar inferences for other response indexes that cannot be calculated for lags 1 and 2.

Appendix

What follows is a mathematical model that expresses performance for T2/T1 as a function of the T1–T2 lag. Instead of being built from a conceptual model, it is an empirical development intended to allow more accurate comparisons between experimental conditions. The final model expresses performance at any lag as some reduction as compared with a control condition, without any effect from T1. Of course, for long lags, T2/T1 performance asymptotically approaches that level of performance. The

asymptotical performance is estimated from the average performance for several lags in the control, one-target, condition; it is the amount d in the model. The final model is

$$f(\text{lag}) = d \cdot (1 - (a \cdot (\text{lag} - b)^2) \cdot e^{-c \cdot \text{lag}}).$$

In this expression, $f(\text{lag})$ is the probability of a correct response to T2, given a hit for T1; a and c are scale parameters; and b is a parameter of position. Table A1 shows the estimates for the parameters, the R^2 fit index, and three special values for each of the 4 two-target conditions (the four fitted functions are depicted in Fig. 4). The special

values of Table A1 are the minimum value reached [$f(\text{lag}_{\text{Min}})$; this reflects the maximum depth of the AB], the lag at which the minimum is reached (lag_{Min}), and the lag at which the value corresponding to 95% of the performance in the control condition is reached (lag_{95}). The values in Table A1 are obtained from the functions fitted for each experimental condition: $f(\text{lag}_{\text{Min}})$ is the minimum of the function, lag_{Min} is the value of the lag at which the function reaches the minimum, and lag_{95} is the value of the lag at which the function reaches the amount equal to the product $.95 \cdot d$. The last two values reflect the time course of the AB. Larger values reflect a longer AB.

Table A1 Models fitted from equation 1 for the four experimental conditions (T1/T2): parameter estimates, fit indices and several critical values (see text)

Condition	Parameters				R^2	$f(\text{lag}_{\text{Min}})$	Lag_{Min}	Lag_{95}
	$\hat{a}$	$\hat{b}$	$\hat{c}$	$\hat{d}$				
RW	3.486	0.318	1.218	58.9	0.993	8.306	1.960	6.467
WW	4.215	0.781	1.052	54.4	0.975	5.081	2.682	7.963
WR	2.998	0.464	1.075	80.8	0.993	11.892	2.324	7.415
RR	4.422	0.441	1.349	85.2	0.992	23.379	1.924	5.816

References

- Botella, J., Barriopedro, M. I., & Suero, M. (2001). A model of the formation of illusory conjunctions in the time domain. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 1452–1467.
- Botella, J., & Eriksen, C. W. (1991). Pattern changes in rapid serial visual presentation tasks without strategic shifts. *Bulletin of the Psychonomic Society*, 29, 105–108.
- Botella, J., & Eriksen, C. W. (1992). Filtering versus parallel processing in RSVP tasks. *Perception & Psychophysics*, 51, 334–343.
- Botella, J., García, M. L., & Barriopedro, M. I. (1992). Intrusion patterns in rapid serial visual presentation tasks with two response dimensions. *Perception & Psychophysics*, 52, 547–552.
- Botella, J., Narváez, M., Suero, M., & Juola, J. F. (2007). Parallel processing of stimulus features during RSVP: Evidence from the second response. *Perception & Psychophysics*, 69, 1315–1323.
- Bowman, H., & Wyble, B. P. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, 114, 38–70.
- Chun, M. M. (1997a). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics*, 59, 1191–1199.
- Chun, M. M. (1997b). Types and tokens in visual processing: A double dissociation between the attentional blink and repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 738–755.
- Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 21, 109–127.
- Di Lollo, V., Kawahara, J., Ghorashi, S. M. S., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control. *Psychological Research*, 69, 191–200.
- Dux, P. E., & Harris, I. M. (2007a). On the failure of distractor inhibition in the attentional blink. *Psychonomic Bulletin & Review*, 14, 723–728.
- Dux, P. E., & Harris, I. M. (2007b). Viewpoint costs occur during consolidation: Evidence from the attentional blink. *Cognition*, 101, 47–58.
- Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. *Attention, Perception, & Psychophysics*, 71, 1683–1700.
- Jolicœur, P. (1998). Modulation of the attentional blink by online response selection: Evidence from speeded and unspeeded task1 decisions. *Memory & Cognition*, 26, 1014–1032.
- Jolicœur, P., Dell'Acqua, R., & Crebolder, J. M. (2001). The attentional blink bottleneck. In K. Shapiro (Ed.), *The limits of attention: Temporal constraints in human information processing* (pp. 82–99). Oxford: Oxford University Press.
- Juola, J. F., Botella, J., & Palacios, A. (2004). Task and location switching effects on visual attention. *Perception & Psychophysics*, 66, 1303–1317.
- Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 404–421.
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching—a review. *Psychological Bulletin*, 136, 849–874.
- Koch, I., Gade, M., Schuch, S., & Philipp, A. M. (2010). The role of inhibition in task switching: a review. *Psychonomic Bulletin & Review*, 17, 1–14.

- McLean, J. P., Broadbent, D. E., & Broadbent, M. H. P. (1983). Combining attributes in rapid serial visual presentation tasks. *Quarterly Journal of Experimental Psychology*, 35A, 171–186.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7, 134–140.
- Nieuwenstein, M., Van der Burg, E., Theeuwes, J., Wyble, B., & Potter, M. (2009). Temporal constraints on conscious vision: On the ubiquitous nature of the attentional blink. *Journal of Vision*, 9 (9, Article 18), 1–14.
- Olivers, C. N. L., & Meeter, M. (2008). A boost and bounce theory of temporal attention. *Psychological Review*, 115, 836–863.
- Popple, A. V., & Levi, D. M. (2007). Attentional blink as errors in temporal binding. *Vision Research*, 47, 2973–2981.
- Popple, A. V., & Levi, D. M. (2008). The attentional blink in amblyopia. *Journal of Vision*, 8(13, Article 12), 1–9.
- Quinlan, P. T. (2003). Visual feature integration theory: Past, present, and future. *Psychological Bulletin*, 129, 643–673.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18, 849–860.
- Reeves, A., & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review*, 93, 180–206.
- Shapiro, K. L., Raymond, J. E., & Arnell, K. M. (1994). Attention to visual pattern information produces the attentional blink in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 20, 357–371.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Vul, E., Hanus, D., & Kanwisher, N. (2008). Delay of selective attention during the attentional blink. *Vision Research*, 48, 1902–1909.
- Vul, E., Hanus, D., & Kanwisher, N. (2009). Attention as inference: Selection is probabilistic; responses are all-or-none samples. *Journal of Experimental Psychology: General*, 138, 546–560.
- Vul, E., Nieuwenstein, M., & Kanwisher, N. (2008). Temporal selection is suppressed, delayed, and diffused during the attentional blink. *Psychological Science*, 19, 55–61.
- Wyble, B., Bowman, H., & Nieuwenstein, M. (2009). The attentional blink provides episodic distinctiveness: Sparing at a cost. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 787–807.