Reinventing the redundant target paradigm to distinguish serial and parallel processing of written words

Alex L. White*[1], John Palmer[2], Genevieve Sanders[1], Jannat Hossain[1] & Zelda B. Zabinsky[3]

[1] Department of Neuroscience & Behavior, Barnard College, Columbia University, 76 Claremont Ave, New York, NY 10031

[2] Department of Psychology, University of Washington, Seattle, WA 98195

[3] Department of Industrial & Systems Engineering, University of Washington, Seattle WA 98195

* Corresponding author: Alex L. White, alwhite@barnard.edu

JP: Alternatively put your address under "corresponding author".

Draft as of: 3/29/24

ABSTRACT

The visual system can encode many stimuli across the visual field simultaneously, but the number of objects that can be fully identified in parallel is limited. At the extreme, some objects might have to be identified serially. One useful tool for distinguishing parallel from serial processing is the redundant target paradigm, which compares responses to displays containing one target to displays containing two targets. Many parallel models predict a positive redundant target effect: faster responses to two targets. Here we revisit the redundant target paradigm by developing and testing predictions for a standard self-terminating serial model that includes errors. Specifically, it predicts *slower* responses to two-target displays than one-target displays. To test that prediction, we conducted three experiments that each measured performance for three different judgments of written words: color detection (detecting colored letters), lexical decision (detecting real words among pseudowords), and semantic categorization (detecting nouns that refer to living things). In all the experiments, only the color detection task yielded a positive redundant target effect, while the lexical and semantic tasks yielded zero or negative effects. These results are consistent with low-level features (color) for two stimuli being processed in parallel, while the meanings of two words are processed serially. Altogether, this study informs models of reading and suggests opportunities for general theories of response time that include errors.

The study of perception has long been animated by the question of whether multiple stimuli can be processed in parallel, or whether strict processing capacity limits require serial processing. Here we re-examine that question in the context of visual word recognition: can two words be recognized simultaneously? This is an important question because competing models of natural reading disagree as to whether multiple words are processed in parallel during each gaze fixation (Engbert et al., 2005; Reichle et al., 2006; Reilly & Radach, 2006; Snell & Grainger, 2019b). To investigate whether it is possible to recognize two words simultaneously, we use the *redundant target paradigm.* This experimental paradigm differs from natural reading, assessing instead how well participants can process two words at once when they are encouraged to try. We compare task performance to the predictions of serial and parallel processing. This complements other approaches that use dual-tasks or measure spatial attention effects (Johnson et al., 2022; White et al., 2018, 2020; White, Palmer, et al., 2019).

Fundamentals of redundant target effects

The redundant target paradigm grew out of a larger visual search literature to investigate whether observers can process multiple stimuli presented simultaneously at different visual field positions (van der Heijden, 1975). The observer's task is to view a display and report the presence or absence of stimuli that belong to a target category. Non-target stimuli are termed "distractors." On some trials, one target is presented. On other trials, multiple targets are presented simultaneously – that is, the display contains redundant targets. The *redundant target effect* is a speeding of correct response times on trials with multiple targets compared to trials with a single target. Such an effect, also termed a "redundancy gain," can be taken as evidence that the targets were processed in parallel.

Studies that have used the redundant target paradigm can be divided into two broad categories~~: first, those that~~ seek to distinguish between a parallel model and a serial

Studies in the first category seek

model (van der Heijden, 1975). These studies compare response times between displays that consist of two (or more) targets, versus trials that contain one target and no other stimuli. Studies in the second category seek to distinguish between two flavors of parallel models: those with separate activations caused by each stimulus, versus those with interactive "coactivations" (C. W. Eriksen et al., 1989; Miller, 1982; Mordkoff & Yantis, 1991). To do so, response times are compared between displays that contain two targets and displays that contain one target and one distractor. These "mixed" trials are not useful for testing the serial model, which is the focus of the present study.

Therefore, we focus on the classic redundant target paradigm that compares displays with a single target presented alone to displays with two targets. As shown in **Figure 1**, contrasting predictions for correct response time arise from a standard unlimited-capacity parallel model and a standard serial model. They are called "standard" models because of strong assumptions about the independence of the processes for each stimulus. Both standard models assume that search is *self-terminating*: the observer responds as soon as they detect a target. The parallel model assumes that when two targets are present, they are independently processed in separate channels that race to produce the response. The completion time of each process is variable across trials. On two-target trials, the response time is determined by the faster of the two processes, so the observer is faster on average than when only one target is present (Raab, 1962). Thus, the parallel model predicts a positive redundant target effect: a speeding of responses.

The standard serial model, in contrast, assumes that one stimulus is processed at a time (Townsend & Nozawa, 1995; van der Heijden, 1975). If the first target of two simultaneously presented targets is correctly identified, then the response time is on average the same as when only one target is present—the redundant target has no effect on performance. Previous descriptions of this serial model stop there; but as we explain in our new theory section below, a serial model that incorporates errors predicts a *slowing* of correct response times if the first target to be processed is misidentified, and search continues to process the second target correctly.

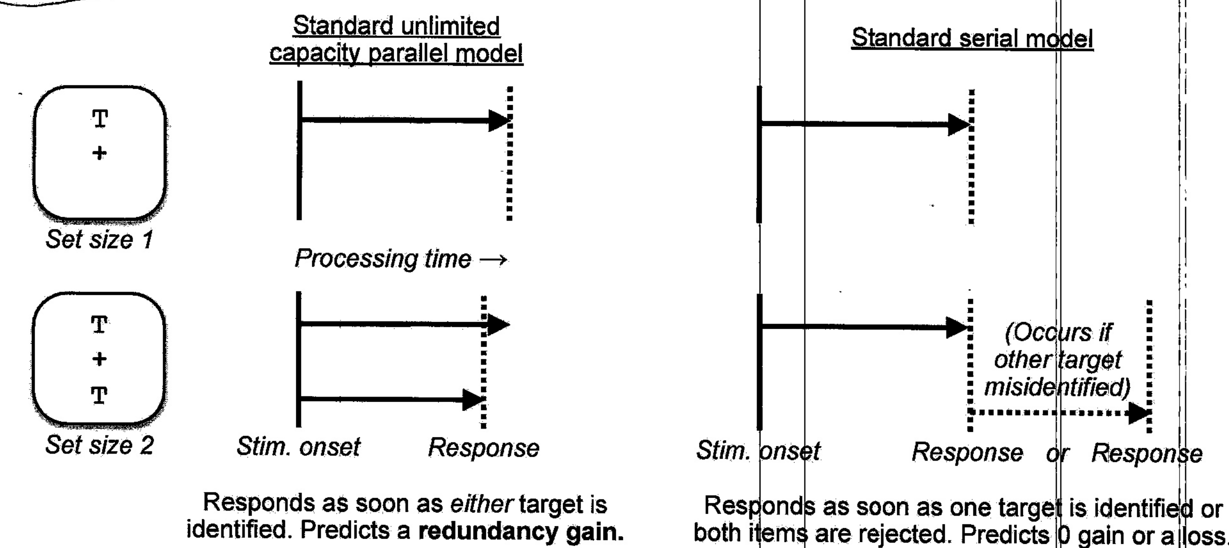

Figure 1: Diagram of a standard parallel and a standard serial model processing displays containing one target (set size 1) or two targets (set size 2). The parallel model predicts faster responses for set size 2 because the response is triggered by whichever process happens to finish sooner. The serial model predicts either zero effect of set size, or a loss (slower responses to set size 2).

Redundant target effects have been used to reject the standard serial model for processing simple visual features, such as detecting lights, discriminating colors, orientations, and motion directions (Corballis, 2002; Donkin et al., 2014; Egeth et al., 1989; Ridgway et al., 2008; Schwarz, 2006; Thornton & Gilden, 2001). Redundant target effects have also been found with auditory stimuli (e.g. (Schröter et al., 2007) and with bimodal stimuli (e.g. (Gondan et al., 2010; Hershenson, 1962). Face recognition has also been studied with redundant target effects (Fitousi, 2021).

Letters are an interesting case, being the building block of words. Several studies have used tasks that require the participant to distinguish one target letter from other letters. When the task uses a "go-no/go" design—to press a button when a target is detected and otherwise make no response—there are positive redundant target effects (e.g. Grice & Reed, 1992; Mordkoff & Yantis, 1991; van der Heijden et al., 1983). That is facilitation at a sub-lexical level (Abrams & Greenwald, 2000)[1]. Mullin & Egeth (1989)'s third experiment included a condition with two real word targets that were different from each other (also in a lexical decision task). In that experiment, the redundant target effect was significantly *negative,* meaning that response times were slowed by the addition of a second target.

A second set of word recognition studies have used semantic categorization tasks. In Mullin and Egeth (1989)'s first and fourth experiments, one or two words were presented above and below fixation. The task was a go/no-go response to the presence of a word belonging to a given semantic category (e.g., 'animals'). Notably, each of the four categories contained only 5 words. In one experiment, when two targets were presented they were identical words, and in the other experiment the two targets were different words. There was no significant redundant target effect in either experiment, consistent with the serial model.

The second relevant study using semantic categorization was by Shepherdson & Miller (2014). We focus on their Experiment 3. Stimuli were presented to the left and right of fixation. On each trial, the subject had to make a yes/no response to report the presence of a word belonging to a target semantic category. Most importantly, the "single-target" trials also contained a "filler" stimulus that was a pseudo-word. They found an advantage for the redundant target condition compared to this modified baseline. Interpreting this experiment is difficult. The critical question is whether a serial, self-terminating model predicts no effect between these conditions. Such a prediction holds only if the target is always processed first by the serial process and the filler pseudoword is never processed first, which seems unlikely. Thus, we conclude this experiment should be considered along with experiments that used a mixed-trial baseline to test co-activation models and not as a test of the serial, self-terminating model.

where should this footnote go??

[1] Nonetheless, presenting copies of the same word across the field provides a redundancy gain that might help patients with macular degeneration to read, with rapid serial visual presentation (Snell et al., 2022)

In summary, the effects of redundant targets in word recognition tasks require further investigation. The presence of a redundant target effect might depend on the subject's task (lexical decision or semantic categorization), the mode of response (go/no-go or forced-choice), and on whether the experiment includes trials in which a target is paired with a distractor. In the new experiments reported below, we investigate all those factors, and compare lexical and semantic tasks to a font color task. In doing so we also test new models of parallel and serial processing that consider both response time and accuracy.

Response time and accuracy in redundant target effects

Most of the redundant target studies reviewed above focused on only correct response times and use tasks in which accuracy is near ceiling. Other studies have focused on accuracy, for instance in the context of spatial summation (e.g., Robson & Graham, 1981; Verghese & Stone, 1995). One important study compared redundant target effects on accuracy and response time (Mordkoff & Egeth, 1993). This work has shown that typical parallel models predict positive redundant target effects for accuracy as well as response time.

As shown in ~~our~~ new theory section below, serial processing of the individual stimuli can lead the opposite effect on response time. *This hinges on the possibility of errors:* if the first target to be processed is misidentified as a distractor, then search continues, and the second target may be correctly identified. Those correct responses increase the mean response time for two-target displays compared to correct responses to single target displays. Thus, our new theory ~~below~~ explicitly considers the accuracy of each stimulus recognition process when predicting response times.

New Theory

The new theory builds on previous models of pure response time by adding the possibility of errors (misclassifying targets as distractors or vice versa). The **Appendix** contains full mathematical descriptions of three classes of models. Here we describe them in more intuitive terms and emphasize the qualitative redundant target effects that they each predict.

Consider the standard self-terminating serial model developed for response time (Townsend & Nozawa, 1995). Like others of its type, it assumes discrete component processes for each stimulus. In addition, it allows time for residual processes before a response is made that do not depend on the stimulus. This model has been called standard because it includes a number of independence properties (see Appendix). We add to this model the possibility of an error and additional independence assumptions concerning the errors and the relationship between error and response time. The four primary results from this new work are described in the following paragraphs, and **Figure 2** illustrates the typical range of predicted redundant targets for each class of model.

Our goal in building this new theory is to compare the *qualitative* predictions of the various models: whether they predict positive, negative, or zero redundant target effects on response time and accuracy. Our goal is not to quantitatively fit models to our data; that is a larger endeavor (never before attempted for redundant target effects) which we leave for the future. For now, it is sufficient to generate qualitative predictions that allow experimental data to rule out some models.

those trials in which the first target was correctly identified and the response was made quickly, and those trials in which the first target was not correctly identified and processing continued to the second target. This new result makes the serial model with errors more distinctive from the parallel models, in terms of response times. The predictions for errors are discussed below.

Quantitatively, this model's predicted redundant target effect is given by Equation 5 in the Appendix, which is equivalent to:

$$\mu_{t,correct} - \mu_{tt,correct} = -\frac{(1-p_t)}{(2-p_t)}\mu_{t,incorrect}$$

The redundant target effect is the difference between the mean correct response time for a single target ($\mu_{t,correct}$) and the mean correct response time for two targets ($\mu_{tt,correct}$). This effect depends on only two factors, the probability correct on single-target trials (p_t) and the mean component processing time for a single target when the participant makes an error (a 'miss' response) ($\mu_{t,incorrect}$).

Crucially, the serial model always predicts a negative redundant target effect. An illustration of this prediction is in **Figure 2.** For this illustration, we assume that for single-target trials, accuracy p_t is 0.95 (5% errors) and the mean correct response time is 800 ms. The upper end of the range is predicted with the assumption that the mean component processing time for an error is equal to the mean component processing time for a correct response. The lower end of the range is predicted with the assumption that the mean component processing time for an error is twice that for a correct response. See the Appendix for more detail.

Predictions of Our Unlimited-Capacity, Parallel Model

Our second result concerns the standard, self-terminating, unlimited-capacity parallel model with errors. The corresponding model without errors predicts a positive redundant target effect. We show that this generalizes to models with errors. Errors can reduce the size of the effect, but it always remains positive. Thus, there remains a sharp contrast in the predictions for this parallel model and the standard self-terminating serial model.

The predictions of this parallel model are given by equation 11 in the Appendix:

$$\mu_{t,correct} - \mu_{tt,correct} = \left(\frac{1}{2-p_t}\right) E\left[\boldsymbol{D}_{t,correct}\right] - \left(\frac{p_t}{2-p_t}\right) E\left[\min\{\boldsymbol{D}_{t_1,correct}, \boldsymbol{D}_{t_2,correct}\}\right]$$

Here, p_t is the probability correct on single-target trials, E[$\mathbf{D}_{t,correct}$] is the mean correct component processing time for a single target, and E[min{$\mathbf{D}_{t1,correct}$, $\mathbf{D}_{t2,correct}$}] is the mean of the minimum of the two component processing times when two targets are presented and judged correctly.

The predicted effect is always positive. This is primarily because the model's response to two targets is driven by whichever of the two stimulus processes finishes first, hence the "min" function in the equation above. On average this is faster than the response to a single target. The equation also makes clear why the predicted effect is always positive: it is the difference between two products, and the first is always larger. This must be the case, as is clear when examining each part of the two products separately. First:

$$\left(\frac{1}{2-p_t}\right) \geq \left(\frac{p_t}{2-p_t}\right)$$

because $0 \leq p_t \leq 1$. Second:

$$E\left[\boldsymbol{D}_{t,correct}\right] > E\left[\min\{\boldsymbol{D}_{t_1,correct}, \boldsymbol{D}_{t_2,correct}\}\right]$$

because the mean difference between two identically distributed (non-negative) variables is always less than the mean of one of those variables alone.

An illustration of this prediction is the middle bar in Figure 2. For this illustration, we used predictions of two specific models described in the Appendix. The upper point is for a diffusion model with parameters that generate large redundant target effects. The lower point is for a linear ballistic accumulator model with parameters that generate relatively small redundant target effects. While not strict limits, these model outputs illusrate the range of predictions from the standard, self-terminating, unlimited-capacity parallel model. They are always positive.

Predictions of Our Fixed-Capacity, Parallel Models

The parallel model that can most mimic a serial model is one that has limited capacity. The limited capacity slows processing when there are two stimuli and thus reduces and possibly eliminates the redundant target effect. Unfortunately, this model is so general that it does not make very specific predictions. Here, we consider a special case of the limited-capacity parallel model: the fixed-capacity, parallel model. The idea of fixed capacity is that a set of parallel processors extract the same total amount of information from multiple stimuli as they do from a single stimulus. Thus, splitting a fixed set of 'resources' between multiple stimuli introduces a cost. Most of the prior work with this model has been in the domain of accuracy (Scharff et al., 2011; Shaw, 1980; White et al., 2018).

We investigated two special cases of self-terminating fixed-capacity parallel models in which we assume a particular stochastic process for each stimulus being processed. Predictions of these two special cases define the range of redundant target effects plotted in Figure 2 (rightmost bar). First, with a diffusion process of sensory evidence accumulation (Palmer et al., 2005), the model yields positive redundant target effects on correct ? response times, for all relevant parameter values (as well as a positive effect on accuracy). This prediction defines the upper end of the range of effects predicted by the fixed-capacity parallel model in Figure 2. However, with a linear ballistic accumulator process (Brown & Heathcote, 2008), the fixed-capacity, parallel model can predict a *negative* redundant target effect on correct response time (a slowing), despite a positive effect for accuracy. This is illustrated in Figure 2 by the lower end of the range of predicted effects for the fixed-capacity model. In essence, because of fixed-processing capacity, the addition of a second target slows processing of both stimuli. Thus, among many parallel models that generate positive response time effects of redundant targets, there are models with fixed-capacity limits that yield the opposite result.

This last result introduces an asymmetry in using the redundant target paradigm to test the serial model and fixed-capacity, parallel models. All our models are assumed to be standard, self-terminating models. A positive redundant target effect rejects the serial model, but a negative redundant target effect does not reject all possible fixed-capacity, parallel models. Thus, the redundant target paradigm is a good test for rejecting both the serial model and the unlimited-capacity parallel model but not the fixed-capacity parallel model. Nevertheless, it is relevant to distinguishing serial and parallel models. Indeed, many experiments with simple feature tasks have used this test to rule out the standard serial model (e.g., van der Heijden, 1975).

Predictions About Errors

Our last result concerns the usefulness of models that incorporate errors. All of the models described above predict a positive redundant target effect for errors (that is, fewer errors on trials with 2 targets than on trials with 1 target). This is not a surprise for typical parallel models that have been investigated in summation experiments of accuracy alone (Graham et al., 1978). What is new is that this result also occurs for our serial model, even though that model predicts slower response times for two targets. The reason is that when two targets are present and processed sequentially, there are two chances to correctly detect target presence, so accuracy increases compared to when only 1 target is present – even though doing so takes more time on average. While this result for errors does not distinguish between the serial and parallel models, it introduces a result that is specific to errors and that is not accounted for by pure response time models.

In each experiment, we also measured performance in three different tasks (color detection, lexical decision, and semantic categorization). The *color task* required participants to judge a low-level visual feature of the words, and served as a control condition for which we expected positive redundant target effects. The *lexical decision task* requires the subject to distinguish real English word targets from pseudoword distractors. The *semantic categorization task* requires categorizing words either as targets that belong to a category of "living things" and distractors that belong to a category of "non-living" things. The semantic and lexical tasks might tap into different levels of linguistic processing, and have both been used in prior redundant target studies (Egeth et al., 1989; Mullin & Egeth, 1989). In sum, within each of our three experiments, we carry out a side-by-side comparison of redundant target effects that arise in three tasks using the same stimuli. The tasks differ in which they require low-level color feature detection, lexical access, or semantic categorization.

Experiment 1: Go/No-Go procedure with correlated stimuli

Methods

Participants: Participants were recruited using Prolific (www.prolific.co, accessed August 2021-May 2023). Participants gave informed consent in accordance with the Declaration of Helsinki and Barnard College's Institutional Review Board. All participants indicated being fluent speakers who learned English as their first language, with no literacy difficulties, and normal or corrected-to-normal vision. For each task, we aimed to recruit an independent sample of 28 participants, half male and half female. That sample size was chosen on the basis of a power analysis of an independent pilot data set, seeking at least 95% power to detect a redundant target effect on response time of 15 ms.

Experiment	Task	N recruited	N excluded	N included	N Female	Mean age [min max]
1: Go/No-Go, correlated	Color	30	2	28	13	32 [19 50]
	Lexical	28	0	28	12	32 [20 47]
	Semantic	28	0	28	14	30 [20 47]
2: Choice, correlated	Color	29	0	29	16	27 [18 50]
	Lexical	28	0	28	14	29 [19 48]
	Semantic	28	0	28	18	27 [19 50]
3: Choice, uncorrelated	Color	28	0	28	22	20 [18 24]
	Lexical	29	1	28	15	31 [20 48]
	Semantic	29	1	28	14	31 [20 48]

Table 1: number of subjects in each experiment and task, as well as the ages in years of the included subjects.

Table 1 indicates the number of subjects and exclusions for all experiments in the study. Across all experiments in this study, our criteria for exclusion were: overall proportion correct <0.6, or proportion correct <0.5 in more than 1 block of 60 trials. In Experiment 1, two subjects in the color task were excluded for the latter reason: one of them had 3 blocks with accuracy <0.5, and the other had 4 such blocks (out of 10 blocks). This was a risk of the go/no-go task conducted over the web browser: if the participant gets distracted mid-block, the experiment carries on without them.

<u>Stimuli</u>: We created and presented stimuli with PsychoPy 3 (Peirce et al., 2019), run through the web browser using Pavlovia (https://pavlovia.org/). Each stimulus size and position were defined as a fraction of the height of the participant's screen; thus, the dimensions in degrees of visual angle likely varied across participants. Participants were asked to sit with their head roughly 1 arm's length from their screen. A central black fixation cross, 4.5% of screen height in width, was present throughout each trial except during feedback. The stimuli consisted of letter strings, of length between 4 and 6 letters. The word lists are described below and provided fully in **the supplemental data repository (<<link>>).** They were drawn in Courier font, with the height of an "o" or "x"

occupying 3.8% of the screen height. That is roughly 0.7 degrees visual angle on a typical laptop. The specific words (or pseudowords) and font color varied across tasks, as described below.

Trial sequence: An example trial is illustrated in **Figure 3A**. Each trial began with just the fixation mark present for 750 ms. Then either 1 or 2 words were presented for 183 ms. There were two possible word positions, centered horizontally and either just above or just below the fixation mark. The distance from the fixation mark to the center of each word was 10% of the screen height (roughly 1.8 degrees visual angle on a typical laptop at arm's length). About two letter o's would fit stacked vertically in the empty space between the screen center and the words.

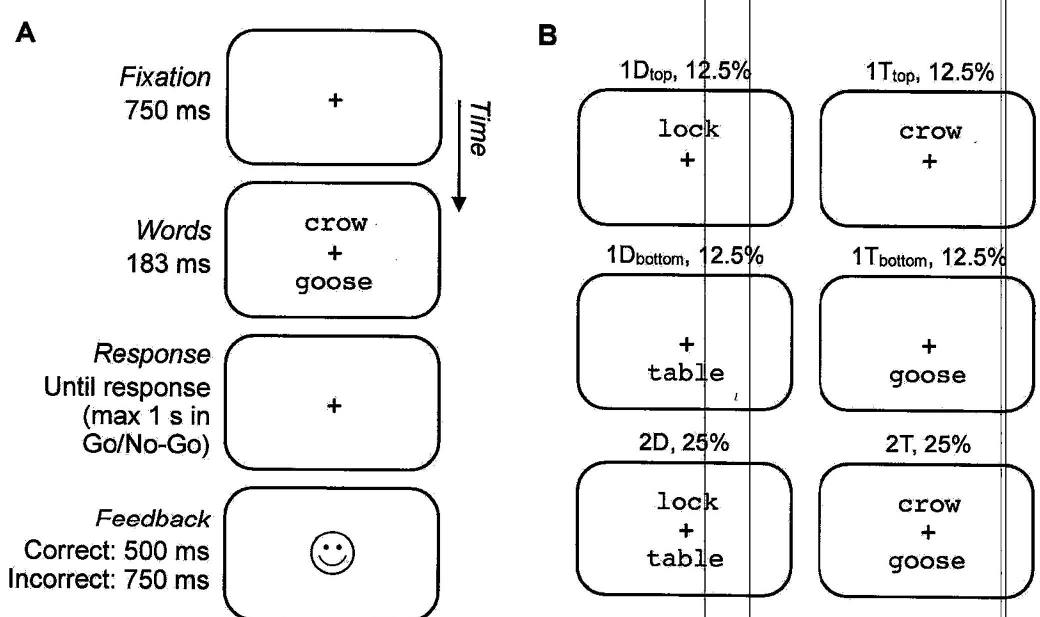

Figure 3: Stimuli and Design. (A) Example trial sequence with two color targets. **(B)** Examples of the trial types in Experiments 1 and 2. The text above each panel indicates the percentage of trials that were of that condition. "D" = distractor, "T" = target. In this illustration, the colored words all refer to living things, thus could be color or semantic targets. But in the actual color task, the semantic category and letter colors were uncorrelated. In the semantic and lexical tasks, the letters were all dark gray.

The trials were evenly distributed between these 4 conditions: 1 target, 1 distractor, 2 targets, and 2 distractors. **Figure 3B** illustrates examples of each trial type, and Table 2

This took me awhile to understand — because it was not clear that lock and table were considered distractors because they are not living things. I think the emphasis on color confused the primary issue of a target being living things.

lists the proportion of trials assigned to each combination of stimuli at the top and bottom locations. Each location could have no word ('None'), a distractor word, or a target word. When just 1 word was present, it was equally likely to be in the top or bottom location. Unlike in Experiment 3, there were never any mixed pairs of 1 target and 1 distractor. Thus, in this experiment, the words in each display were "correlated," meaning that when there were 2 words present, they were either both distractors or both targets.

After the words disappeared, the participant was free to respond. In these go/no-go tasks, the participant was instructed to press the spacebar as soon as they detected a target, and to do nothing except wait for the trial to end if they saw no targets. Up to 1 second was allowed for a response. After the response interval elapsed or was ended by a keypress, feedback was given: the fixation cross was replaced with a smiley face for 500 ms if the response was correct, or a neutral face for 750 ms if the response was incorrect. Then, the fixation cross reappeared and another trial began (except when it came time for a break between blocks, see below).

		Bottom word		
		None	*Distractor*	*Target*
Top word	*None*	*N/A*	*0.125*	*0.125*
	Distractor	*0.125*	0.25	0
	Target	*0.125*	0	0.25

Table 2: The probability of stimulus pairings at the top and bottom locations in Experiments 1 and 2. The word at each location was either absent, a distractor, or a target. The green shading highlights conditions when 2 words were present. In this design, the two words were perfectly correlated, meaning that they were either both targets or both distractors.

Task	Effect mean	Effect SEM	95% CI	t	p	BF
	Correct response time (ms)					
Color	31.98	2.49	[27 36]	12.67	6.35×10^{-12}	1.09×10^{10}
Lexical	-0.11	3.53	[-6 8]	0.03	0.98	0.20
Semantic	-10.67	3.53	[-17 -6]	3.60	0.002	27.69
	Errors (percent)					
Color	7.95	2.10	[4.54 13.23]	3.72	0.001	35.778
Lexical	2.12	0.80	[0.62 4.02]	2.59	0.017	3.22
Semantic	1.90	0.86	[0.49 4.12]	2.19	0.038	1.54

Table 3: Statistics on redundant target effects in Experiment 1, describing the mean *improvement* in response time or error rate, contrasting 1-target displays vs. 2-target displays. The degrees of freedom for the t-tests was 27. For each measure (response time or accuracy), p-values are corrected for false discovery rate across all 9 comparisons including all 3 experiments in the study. BF = Bayes Factor.

Discussion

The redundant target effects in ~~this~~ the first experiment suggest that the colors of the letters within two words can be processed in parallel, leading to speeding of response times. However, the meanings of the two words are not necessarily processed in parallel. This is because the presence of a second word target in the lexical decision task yielded 0 improvement in response time, and the semantic categorization task yielded a significant slowing of response time.

Our new theory shows that such a negative response time effect is in fact consistent with the standard serial model, even when accompanied by an increase in accuracy (see Appendix). It can be explained by participants occasionally mis-categorizing the first target they process as a distractor, then going on to correctly process the other target, with a slower response time compared to correct trials with single targets. It is potentially interesting that this negative redundant target effect was significant in the semantic task but not the lexical task, but we lack strong statistical evidence that those two results were significantly different from each other.

are uncorrelated. This should motivate the participant to process both stimuli as well as they can. Uncorrelated stimuli like this are also common in the wider visual search literature.

Methods

<u>Participants</u>: Participants in the lexical and semantic tasks were recruited and compensated in the same way as in Experiment 1, via Prolific. Participants in the color task were recruited from the Barnard College Introductory Psychology subject pool, and participated in exchange for course credit. See Table 1 for counts. One participant was excluded for overall accuracy being <0.6, and another because they pressed the same key on every trial (but overall proportion correct was >0.6 because targets were present on 75% of trials).

<u>Stimuli and procedure:</u> All details were the same as in Experiment 2, except as noted here. The primary difference is that 30% of trials contained mixed pairs of 1 target and 1 distractor. **Table 2** lists the proportions of trials assigned to each type, which were chosen to ensure that the categories (target vs distractor) of the upper and lower stimuli were independent (uncorrelated). Specifically, on two-word trials, the conditional probability of one stimulus being a target given that the other was a target was 0.5. In contrast, this conditional probability was 1.0 in Experiment 1 and 2. Another difference in Experiment 3 was that across the experiment, the probability of a target being present on any given trial was 0.75, rather than 0.5. That was true both among trials with set size 1 and trials with set size 2. Also, six words were added to the stimulus set for the color and semantic tasks (see Appendix II).

<u>Analysis:</u> We analyzed these data in the same way as Experiment 1 and 2, focusing on the comparison between trials with two targets and trials with a single target presented alone, which provide the best test of our self-terminating models of parallel or serial processing. The mean percentages of trials excluded for sluggish response times

distinguish our self-terminating serial model from the parallel models. Even the serial model predicts faster responses to two targets than to a single target paired with a non-target, because if the non-target is processed first, search must continue. As shown in Figure 9, mean responses to mixed-pair trials were in fact slower (and less accurate) than responses to trials with two targets. Thus, we have focused on comparing two-target trials (set size 2) to single-target trials (set size 1). The result of that contrast was dramatically different across tasks.

General Discussion

The three experiments reported above consistently demonstrate that there is a positive redundant target effect when the targets are defined by color but not when the targets are defined by lexicality or by semantic category. In all three tasks, the stimuli were written words presented singly or in pairs above and/or below fixation. In the color task, the presence of a redundant target (a word written in colored letters) consistently sped responses compared to trials with a single target. In the semantic task, the presence of a redundant target (a word that refers to a living thing) *slowed* responses in each experiment. In the lexical decision task, the redundant target had no effect in two experiments (with correlated stimuli), and a significantly negative effect in the third experiment (with uncorrelated stimuli).

These data are all consistent with the hypothesis that the low-level features—such as color—of multiple stimuli are processed in parallel, but written words are *recognized* serially. The lexical decision and semantic categorization tasks were two ways to assess word recognition: the lexical decision task requires the participant to judge the familiarity of each letter string, and the semantic categorization task further requires further judgment of the word's meaning.

hyphen here are confusing, use commas

again consistent with the standard serial model (Campbell et al., 2024). These data lead to a similar conclusion as the redundant target studies reported in the present article: low-level features are processed in parallel, but linguistic features of words are processed serially. In addition, one neuroimaging study identified a potential neural locus of the serial bottleneck in the left ventral temporal cortex (White, Palmer, et al., 2019).

A related paradigm is called "partially-valid cueing," in which one of two stimulus locations is pre-cued as more likely to be task-relevant. One study using post-masked words found that when participants were asked to judge the semantic category of a word that appeared at the uncued (less attended) location, they performed no better than chance (Johnson et al., 2022). This is again consistent with the standard serial model for word recognition. On the face of it, this result is inconsistent with any parallel model. If processing in parallel, why not acquire some information about the low-probability word? To save the parallel model, one must assume a strategy of processing only one word at a time under some conditions; in other words, the parallel model becomes effectively serial.

Not all studies agree, however (Snell & Grainger, 2019a). Varieties of the "flanker paradigm" (Eriksen & Eriksen, 1974) have demonstrated that judgments of one target word are influenced by the characteristics of nearby, task-irrelevant words (Snell et al., 2017; Snell, Mathôt, et al., 2018; Snell & Grainger, 2018, but see Broadbent & Gathercole, 1990). That is true even when the whole display is flashed for 50 ms and then masked (Snell, 2024). Moreover, there is a "sentence superiority effect": when the words that are flashed along with the target form a sentence, the target is reported more accurately than when the word order is scrambled (Snell & Grainger, 2017; Wen et al., 2019). One interpretation of these results is that the words each display were all processed simultaneously (Snell & Grainger, 2019b). It is important to note that these experiments differ from the redundant target experiments reported above in at least two key ways: (1)

awkward

theory that ignores errors (e.g., Townsend & Nozawa, 1995). This work is also general in not assuming particular stochastic processes or response time distributions.

The second path is to assume a specific stochastic process such as the diffusion process or the linear ballistic accumulator (both described in the Appendix). For example, Blurton, Greenlee & Gondan (2014) built on the diffusion process to model the redundant target effect. The strength of this path is the integrated treatment of response time and accuracy.

Here we sought to expand the general response time models to incorporate errors. The surprising result was that the standard, self-terminating, serial model with errors showed a *negative* redundant target effect. This is not predicted by the corresponding pure response time model. This work complements other recent effects to generalize pure response time models. In particular, Little et al. (2022) extended part of the theory of systems factorial technology to include errors. They examine the prediction of the double-factorial paradigm to distinguish serial and parallel processes. They showed that the previous analysis of exhaustive search models was general to conditions with errors. However, they did not find a similar general result for self-terminating search models. Instead, they examined two special cases and showed that the analysis for pure response time did generalize to those cases. This is important progress, but it remains to be determined if this method of distinguishing serial and parallel models holds for *all* standard, self-terminating models with errors.

In summary, a critical development is the creation of general theories of both response time and accuracy. We have developed such a theory for the redundant target paradigm.

Conclusion

This study makes two primary contributions: first, we developed a new theory of the redundant target effect, which yielded some new results. By accounting for errors as