Evolving the stimulus to fit the brain: A genetic algorithm reveals the brain's feature priorities in visual search

Erik Van der Burg University of Sydney, Sydney, Australia
VU University, Amsterdam, the Netherlands

John Cass University of Western Sydney, Sydney, Australia

Jan Theeuwes VU University, Amsterdam, the Netherlands

David Alais University of Sydney, Sydney, Australia

How does the brain find objects in cluttered visual environments? For decades researchers have employed the classic visual search paradigm to answer this question using factorial designs. Although such approaches have yielded important information, they represent only a tiny fraction of the possible parametric space. Here we use a novel approach, by using a genetic algorithm (GA) to discover the way the brain solves visual search in complex environments, free from experimenter bias. Participants searched a series of complex displays, and those supporting fastest search were selected to reproduce (survival of the fittest). Their display properties (genes) were crossed and combined to create a new generation of "evolved" displays. Displays evolved quickly over generations towards a stable, efficiently searched array. Color properties evolved first, followed by orientation. The evolved displays also contained spatial patterns suggesting a coarse-to-fine search strategy. We argue that this behavioral performance-driven GA reveals the way the brain selects information during visual search in complex environments. We anticipate that our approach can be adapted to a variety of sensory and cognitive questions that have proven too intractable for factorial designs.

Introduction

When searching for objects in our visual environment, we face the problem of clutter. While objects presented in isolation are trivially easy to find, the complexity and density of natural images impairs search performance (Bex, Solomon, & Dakin, 2009). Difficult visual search tasks are therefore a common feature of our daily lives, whether finding a friend in a crowd, a word on a page, or a needle in a haystack. For over 30 years, visual search has been studied using factorial designs, whereby certain aspects of the stimulus are systematically manipulated by the researcher (for recent reviews, see Awh, Belepolsky, & Theeuwes, 2012; Wolfe & Horowitz, 2004). Although this approach allows us to infer causal relationships between stimulus attributes and performance, testing the enormous number of stimulus combinations is extremely inefficient. Consequently, most studies examine a small subset of possible feature permutations within a much larger search space, leading to potentially unwarranted theoretical generalizations. To overcome this constraint, we introduce an alternative method well suited to searching large spaces known as a genetic algorithm (GA; Holland, 1975).

A GA is a search and optimization technique that mimics natural selection. In general, a GA begins with a randomly sampled set of potential solutions for a complex problem which then evolves over generations toward a set of more optimal solutions. A fitness function evaluates how optimal each solution is in the set. Within a sample set, solutions that are poor are left to die out, while better solutions are selected (i.e., "survival of the fittest") and allowed to mate and propagate their advantageous traits over subsequent generations using a crossover technique. A small mutation rate is applied to maintain a diverse population and to avoid convergence to local minima (Holland, 1975).

Citation: Van der Burg, E., Cass, J., Theeuwes, J., & Alais, D. (2015). Evolving the stimulus to fit the brain: A genetic algorithm reveals the brain's feature priorities in visual search. *Journal of Vision, 15*(2):8, 1–15, http://www.journalofvision.org/content/15/2/8, doi:10.1167/15.2.8.

doi: 10.1167/15.2.8 Received May 23, 2014; published February 6, 2015 ISSN 1534-7362

To date, GAs have been frequently applied to solve a plethora of complex problems within the fields of mathematics, economics, physics, engineering, chemistry, and even psychology. Although GAs are in general used as an optimization technique to find an optimal solution for a complex problem, in our case we are not interested in finding an optimal search display as such, as there is already an extensive literature about pop-out search (for recent reviews regarding visual search, see Awh et al., 2012; Wolfe & Horowitz, 2004). Instead, we are interested in how we search through more complex displays, and we therefore use the GA to manipulate distractor properties and thus the amount of competition in the search displays. We use correct mean reaction times (RTs) to measure performance (although different dependent variables could be used, e.g., error rate or eye movements) because RTs are the most common measure to identify the amount of competition in a given display (see, e.g., Cass, Van der Burg, & Alais, 2011; Duncan & Humphreys, 1989; Folk, Remington, & Johnston, 1992; Nakayama & Silverman, 1986; Theeuwes, 1992; Treisman & Gelade, 1980; Van der Burg, Olivers, Bronkhorst, & Theeuwes, 2008). Thus, whereas other GA studies have used a fitness function to evaluate the optimality for each solution, the correct mean RT reflects the fitness value directly.

Our interest in manipulating the amount of competition in the display can be achieved by adopting one of two possible selection strategies. One is to select the fastest searched displays from each generation (i.e., the "survival of the fittest" principle), which then evolve over generations towards display sets in which the amount of competition decreases. A second strategy is to select the slowest displays from a generation (i.e., the "survival of the worst" principle), which then evolve towards display sets in which the amount of competition increases. In the present study, we decided to select the best displays from each generation because we expected search to be rather difficult for the initial displays and because this approach would be less time consuming, as RTs would decrease over generation. Finally, and most importantly for our research question, we could determine which properties of the display evolved and thus what factors were driving efficiency gains in visual search. Here we show that GAs provide a quick and efficient way to investigate human search behavior in complex displays and provide informative insights into which are the salient factors in driving competition.

Experiment 1

Observers searched a complex visual array (Figure 1A) as quickly as possible for a horizontal red line (the target) among 72 distractors of various orientations (0°, 10°, and 90°) and colors (red, green, and blue). This combination leads to 8^{72} different display configurations, too many to be tested in a factorial design. In our GA approach (see Figure 1B), 12 unique displays were randomly drawn, with the only constraint being that a distractor could not be the target (i.e., horizontal and red). Observers searched each display six times, and mean correct RTs determined the fitness value for each display. Applying a "survival of the fittest" principle, we selected the three displays yielding fastest search performance as parents to reproduce the next generation. The parents spawned a new generation of 12 displays (children) according to a "crossover and mutate" procedure (Figure 1C). For the crossover, the genotype of each parent was split into two equal segments and new children were created by combining segments from different parents. Note that we decided to select the best three displays, as all possible crossovers created 12 new genotypes (i.e., maintaining the same number of displays over generations). Theoretically, more or fewer parents could be selected to produce the next generation. More research is required to examine what the optimal number is when using an interactive GA, however, in the present experiment 12 children per generation provided sufficient diversity without requiring us to test impractically large stimulus sets. A mutation rate of 4.1% was added by randomly changing three values in the 72 numbers making up the genotype sequence to new random integers in the range 1–8 (i.e., new color/orientation combinations). Mutation introduces diversity to the population and allows the algorithm to avoid local minima. Typically, the mutation rate is chosen in the range 0.5%–5% (see, e.g., Srinivas & Patnaik, 1994). The 12 genotypes were then converted to phenotypes (i.e., new search images) by randomly placing 72 line segments in the display with the color/orientation combinations indicated in the genotype, plus the target. These 12 "evolved" displays formed the next generation and the cycle repeated: Each display was searched six times and the fittest were selected to reproduce. The displays in our experiment evolved through 14 generations, with the 12 displays in the first generation all beginning from a unique location in the color/orientation search space.

Methods

Participants

Six VU University (Amsterdam) students (two women, mean age: 21.5 years, range: 20–23 years) participated in Experiment 1. Participants were naive as to the purpose of the experiment and were paid 8 euros/h. Informed consent was obtained from each participant after the nature of the study was explained

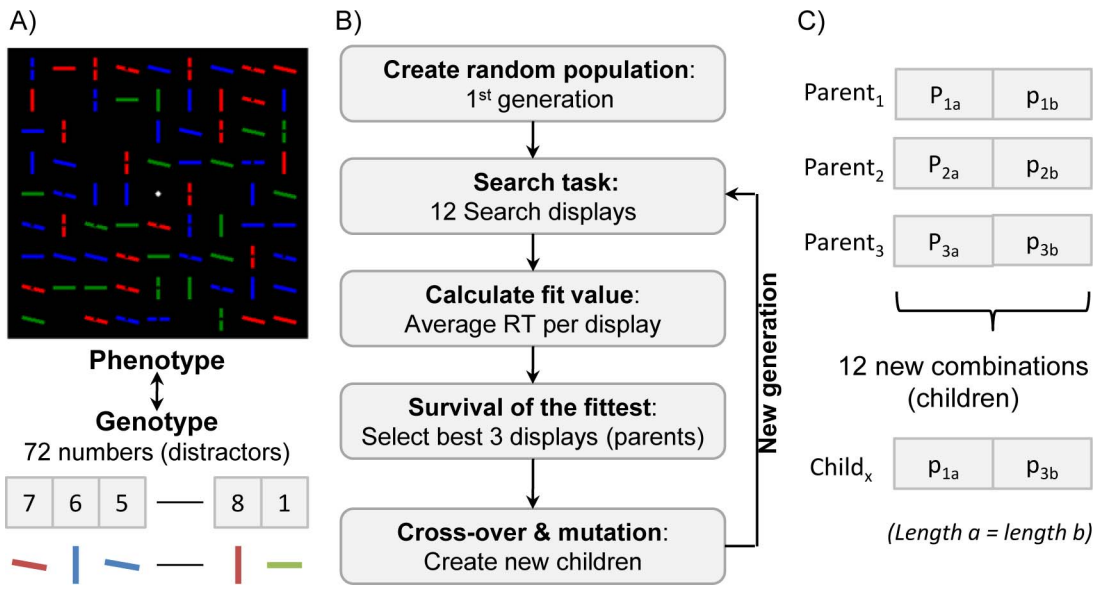

Figure 1. (A) The search display consisted of 73 lines (one target, 72 distractors) presented on a grid. Participants searched for the single horizontal red line and indicated the presence or absence of a gap in this line. Each display (phenotype) is represented by its own gene (genotype): a string of 72 numbers in the range 1–8 representing each distractor's color/orientation combination (3 orientations × 3 colors − 1 combination reserved for the target). (B) The genetic algorithm used in this study began with 12 randomly determined displays forming the first generation. Observers searched a random selection of displays, and those combinations of features (gene sequences) producing the three fastest search times were then selected to be "parents" to reproduce a new generation of "children" according to the "crossover and mutate" procedure shown in panel C. (C) Each parent's genotype was divided into equal segments (a and b), and new children were created by crossing segments from different parents (here, $Child_x$ combines Parents 1 and 3: $C_x = P_{1a} + P_{3b}$). A mutation component was included during intergenerational crossover by changing three randomly selected numbers in the genotype sequence into new random integers in the range 1–8. With three parents, there were 12 crossover combinations leading to 12 evolved children (representing the next generation).

to them. The research was approved by the Ethics Committee of the VU University. The experiments were conducted according to the principles laid down in the Helsinki Declaration.

Apparatus and stimuli

Experiments were run in a dimly lit cubicle. Participants were seated approximately 80 cm from the monitor. E-Prime software was used to program and run the stimuli. The search displays consisted of 73 red (28.5 cd/m^2), green (65.4 cd/m^2), and blue (7.7 cd/m^2) line segments (length: 0.5°; width: 0.08°) of various orientations (0°, 10°, and 90°). Participants were required to search for the target line (a horizontal [0°] red line) among the distractor lines (all other color/orientation combinations) and to respond to the presence or absence of a small gap in the target line (0.08°) by making a speeded response (Z and M keys, respectively). Half of the distractor lines, randomly selected on each trial, also had a gap. All lines were presented on an invisible 9 × 9 grid (8.2° × 8.2°) centered on a white (0.1°; 76.7 cd/m^2) fixation dot (see Figure 1A). The background color was black (<0.5 cd/m^2) and kept constant during the experiment.

Design and procedure

Participants performed the search task in 14 blocks of trials, with each block representing a generation in the evolution of the stimuli. Within each block, 12 different configurations of the search display (i.e., a population) were tested. Each display was tested six times (72 trials per block), with the locations of the line segments randomly assigned on every trial. For the first generation, 12 search displays were generated by assigning a position and a color/orientation combination in a completely randomized manner for each of the 72 distractor lines. The only constraint was that none of the distractors could possess the target-specific combination: horizontal and red. Each search display (phenotype) was represented by its own gene (genotype), a string of 72 numbers defining the characteristics of each distractor. Each number took a value ranging from 1 to 8 (3 colors × 3 orientations − 1 target combination) representing a given distractor's feature combination (e.g., 1 indicates a green vertical distractor, 2 a blue vertical distractor, etc.).

In the first generation, participants performed the search task six times on each of the 12 different displays in the starting population and a fit value was calculated. The fit value was defined as the mean

reaction time (RT) on correctly searched trials and was used to quantify search performance and thus to select which displays were the fittest and should be selected to reproduce (the "survival of the fittest" principle). We chose the three fittest displays (i.e., those with the lowest RTs) as parents for the next generation. The parents spawned a new generation of 12 displays (children) according to the "crossover and mutate" procedure illustrated in Figure 1C. The crossover part of this procedure splits the genotype of the three parents into two equal segments (e.g., Parent 1: P_{1a} and P_{1b}; Parent 2: P_{2a} and P_{2b}; Parent 3: P_{3a} and P_{3b}) and combines them in all possible combinations to create twelve new genotypes (e.g., $\text{Child}_x = P_{1a} + P_{3b}$). The mutation involves changing three randomly selected numbers in the genotype sequence into a new random integer in the range 1–8. Once the 12 new genotypes are created, each is converted to its corresponding phenotype (i.e., the 72 line segments indicated by the color/orientation combinations in the genotype, as well as the target, are randomly placed in the display). These 12 "evolved" search displays represent the next generation of stimuli and the cycle is repeated: Again each display is searched six times, and the fitness value of each is used to select three further displays to reproduce. In our experiment, the displays evolved through 14 generations. Participants performed one practice block of 72 trials before commencing the experiment.

Results

Reaction time

Figure 2 shows the evolution of overall correct mean RT (i.e., the fitness value) over 14 generations averaged across observers. Overall mean error rate was low (1.4%) and not further examined. For each participant, a two-parameter power function[1] was fitted to the correct mean RT to quantify the efficiency of the genetic algorithm:

$$f(x) = ax^b \quad (1)$$

Here, variable x is an integer representing the number of generations. Parameter a represents the time required to search for the target through the random first generation, and parameter b represents the rate at which RTs decrease or increase over generations. Note that a negative b reflects a decrement in performance over generation, whereas a positive value indicates performance increment.

Parameter a was 3100 ms on average. As is clear from Figure 2, participants showed a strong tendency to respond faster over successive generations. This was statistically confirmed by a two-tailed one-sample t test, as the overall mean b (−0.31) was significantly smaller than zero, $t(5) = 10.2$, $p < 0.0005$. This is consistent

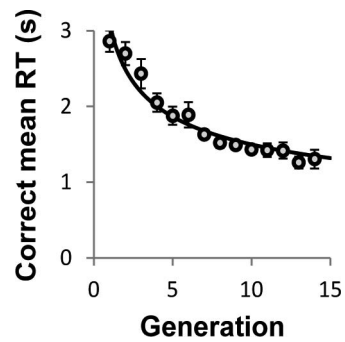

Figure 2. Correct mean response time (RT; i.e., fitness value) as a function of the number of generations shows a decline in search times as the displays evolve. The continuous line is the best fitting power function to the mean RT over generations. Error bars show ±1 standard error of the group mean ($n = 6$).

with the notion that the GA guides stimulus evolution in a way that greatly facilitates visual search. It might be countered that the overall search times are relatively long (>1 s) and thus not indicative of parallel or efficient search; however, the use of a difficult, slow task is deliberate, as our primary interest is not in finding the fastest optimized display but rather in revealing the processing strategies the brain employs to help facilitate a difficult visual search task. In any case, the search display is unlikely to ever achieve an optimal state for search performance, because we apply a small random mutation to every display. Note that small levels of mutation are routinely applied to genetic algorithms to avoid converging upon local minima in the search space (Holland, 1975).

The next question is to examine how the GA guides stimulus evolution: *Which* display properties that evolve during evolution afford faster search performance? One possibility is that the target and distractors become dissimilar over generations so that featural selection via competition declines (Duncan & Humphreys, 1989). According to this view, decreasing the number of red distractors or horizontal distractors would make the red horizontal target more unique, progressively approaching a color or orientation pop-out (Foster & Ward, 1991; Treisman & Souther, 1985).

Evolution of distractor combinations

Figure 3A shows the mean number of each distractor combination in the display plotted as a function of generation. The continuous lines represent best fits.

To analyze display properties further, power functions (Equation 1) were fitted for each of the eight distractor combinations, another three analyzing color with orientation collapsed, and three more analyzing orientation with color collapsed. This was done for each subject to examine whether display properties increased or decreased over successive generations.

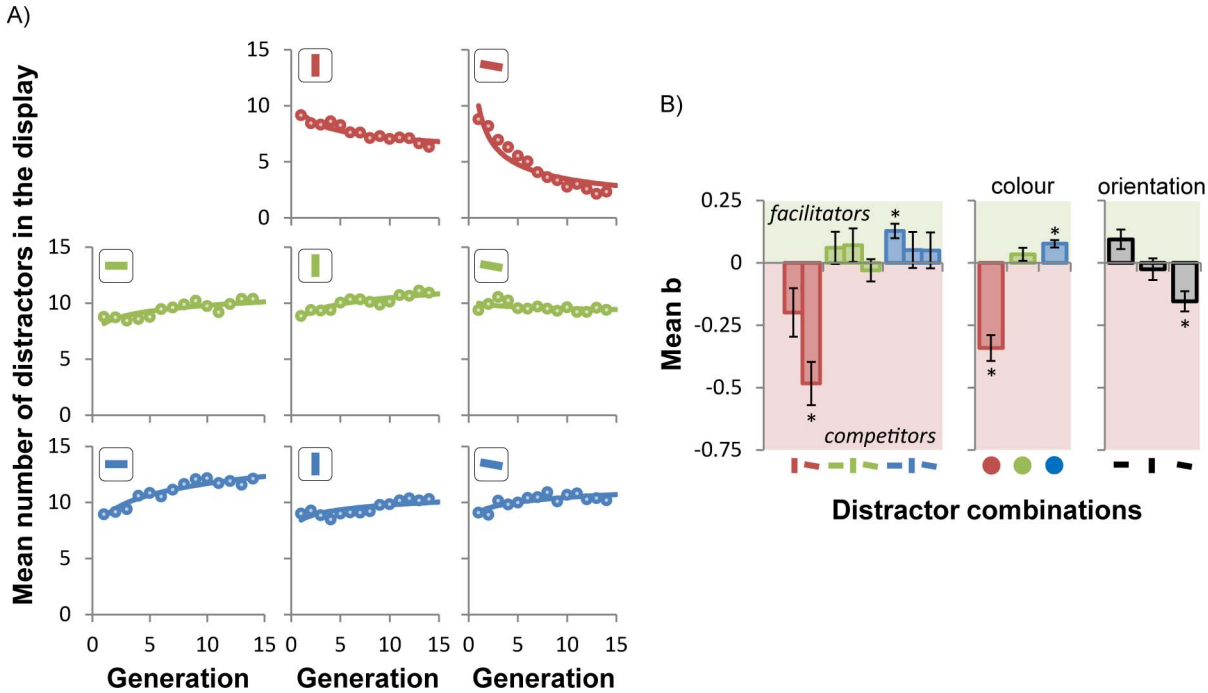

Figure 3. Evolution of distractors. (A) Mean number of each distractor type in the display as a function of generation. Continuous lines represent best fitting power functions (see Equation 1). (B) The group mean value of parameter *b* as a function of each distractor combination and as a function of each color and orientation. A negative value for *b* signifies that these distractors were detrimental for search (red shaded area) and declined over generations, whereas a positive *b* indicates that these distractors facilitated visual search (green shaded area) and increased in number. Asterisks (*) indicate that those distractors were significantly different from zero (see Table 1 for a summary of all statistical results).

Parameter *a* reflects the mean number of initial elements in the display, and *b* once again reflects the increase or decrease in the number of elements in the display over successive generations. The mean *b* for each distractor combination is shown in Figure 3B. Two-tailed *t* tests were conducted on parameter *b* for each display property. Table 1 summarizes these results.

Although an evolution towards an orientation pop-out stimulus is a possibility (i.e., one in which the target is the sole horizontal element), the data do not support it. Horizontal distractors (leftmost panels in Figure 3A) actually become more numerous over generations, $p = 0.06$ (data collapsed over color). Figure 3B shows that the number of blue, $p = 0.001$, and green, not significant, distractors also increased over generations. The only distractors to become less numerous were the red 10° distractors, $p = 0.003$, those close to the target's orientation and sharing its color. Importantly, no changes in the number of red 90° distractors, not significant, were observed, implying that the orientation—not the color—of the red 10° distractor was the primary source of competition with the target.

Although the GA was unconstrained, our results are consistent with earlier studies that used a classic factorial design. For instance, Kaptein, Theeuwes, and

Distractor combination				
Color	Orientation	*a*	*b*	*t*(5)
Red	90°	9.7	−0.20	2.0
Red	10°	10.0	−0.48	5.6**
Green	0°	8.3	0.06	0.9
Green	90°	8.7	0.07	1.0
Green	10°	10.0	−0.03	0.7
Blue	0°	8.7	0.13	4.4**
Blue	90°	8.5	0.05	0.7
Blue	10°	9.2	0.05	0.7
Red		9.7	−0.34	6.6**
Green		9.0	0.03	1.3
Blue		8.7	0.08	5.2**
	0°	8.4	0.09	2.4
	90°	8.9	−0.03	0.6
	10°	9.6	−0.15	3.8*

Table 1. Summary of curve-fitting results from Experiment 1. Table 1 illustrates the mean value for parameters *a* and *b*: one for every distractor combination, one for each color (orientation collapsed), and one for each orientation (color collapsed). T-values are shown for each test of whether the number of elements in the display significantly increased (i.e., a positive *b*) or decreased (i.e., a negative *b*). * $p < 0.05$, ** $p < 0.005$.

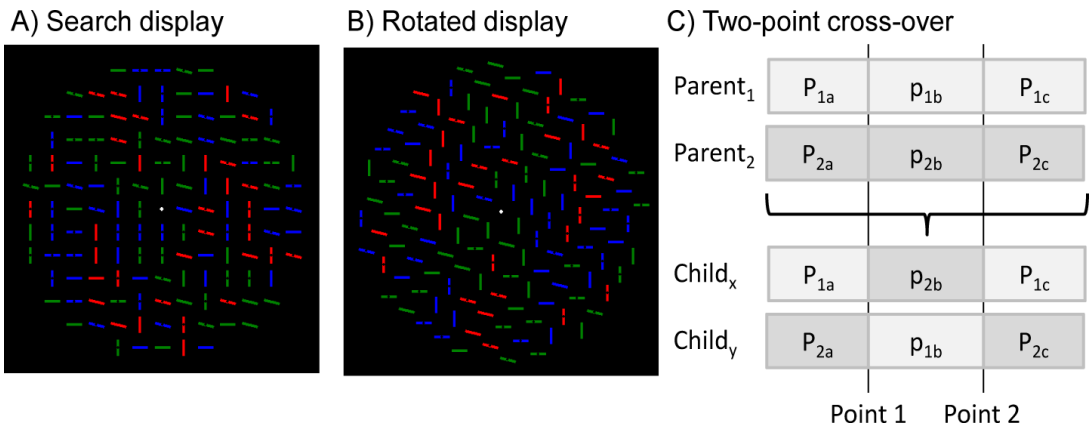

Figure 4. (A) The display in Experiment 2 consisted of a circular array (6° radius) of 136 line elements, and participants searched for the horizontal red line (always 3.9° from the center) and indicated the presence or absence of a gap in this target line. (B) Same display as in panel A, but rotated around fixation by 216° to change the element locations. Note that intertrial rotation changed the position of objects in viewer-centered coordinates but not object-centered coordinates, i.e., retinal positions were affected but their relative locations were not. Note also that this rotation had no affect on the local orientation of individual elements (e.g., vertical elements remained vertical regardless of intertrial rotation). (C) For the crossover procedure, two points along the genotype were randomly selected, splitting the sequence into three segments (a, b, and c). New children were created by exchanging the middle segment between each parent. With four parents selected to reproduce each generation, there were 12 unique combinations and therefore 12 evolved children.

Van der Heijden (1995) showed that when searching for conjunction targets, observers restrict their search to color-defined subsets, not to the orientation-defined subsets. In our displays, when searching for a red horizontal line among red, green, and blue lines of various orientations, our results suggest that observers limited their search to red items, ignoring other colors. This explains why the number of red 10° items decreased over generations (since they are highly similar to the target and compete with it) and why red 90° items remained stable (as they provide the pop-out contrast for the horizontal target). Searching through the red subset would be an efficient strategy, as it reduces the search space by two-thirds.

Experiment 2

In Experiment 1, the target and distractor locations were randomly determined in each display, making it unclear whether search performance was affected by local factors around the visual target or by larger scale perceptual grouping effects (Humphreys, Quinlan, & Riddoch, 1989; Treisman, 1982). For instance, it is possible that search performance improved because the red elements evolved at locations more distant from the target in order to spatially isolate it from distractors. Experiment 2a addressed this by fixing the relative positions of target and distractors. To prevent observers learning the actual retinotopic target location instead of the relative position of target and distractors (which would be trivially easy), displays were rotated on each trial by a random angle (0°–360°) to maintain task difficulty. Except for keeping the relative positions of target and distractor constant, the GA and task were identical to those used in Experiment 1.

For the first generation, 12 displays were generated by allocating random color/orientation combinations to 135 distractor locations, creating a genotype sequence of 135 numbers in the range 1–8. Unlike in Experiment 1, each position in the genotype sequence represented a unique display location. Observers ($n = 27$) searched the displays, and the four best displays (fastest mean RTs after six trials) were selected to reproduce. The next generation was created by a two-point crossover which maintained location-specific information (see Figure 4C). Two points on the genotype were randomly selected, and everything between them was swapped between the parents to create two new child genotypes. Each child's genotype was mutated at five points by randomly reassigning the color/orientation combination.

We also conducted a control experiment to examine whether the search benefits accruing over generations were due to a practice effect. In Experiment 2b, 14 new participants performed the same task, except that we did not apply the GA. Instead, we used 12 new random displays each generation. After a practice block (72 trials), participants did four experimental blocks (equivalent to four generations).

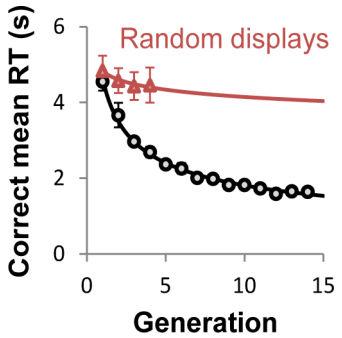

Figure 5. Effect of the genetic algorithm on search performance in Experiment 2. Mean correct response time measured as a function of generation: Black circles signify displays subject to the genetic algorithm; red triangles, displays in which each generation is a randomly generated combination of distractor features not subject to the genetic algorithm (control condition). Error bars show ±1 standard error of the group mean ($n = 27$). Continuous lines represent best fitting power functions.

Method

Participants

27 VU University students (15 women, mean age: 21.6 years, range: 18–26 years) participated in Experiment 2a. In Experiment 2b, 14 VU University students (eight women, mean age: 23.9 years, range: 22–27 years) participated in the control experiment. In both experiments, participants were naive as to the purpose of the experiment and were paid 8 euros/h. None of the students participated in both Experiments 2a and 2b.

Experiment 2a was similar to Experiment 1, except for the following details. Again participants were asked to search for the horizontal red line and to respond to the presence or absence of a small gap in the target line. Rather than randomizing the location of each display element as in Experiment 1, the present experiment used displays with fixed relative locations between distractors and target but whose retinotopic locations were varied randomly in each trial by rotating the entire display by a randomly chosen angle (0°–360°). Rotating the display maintained a high

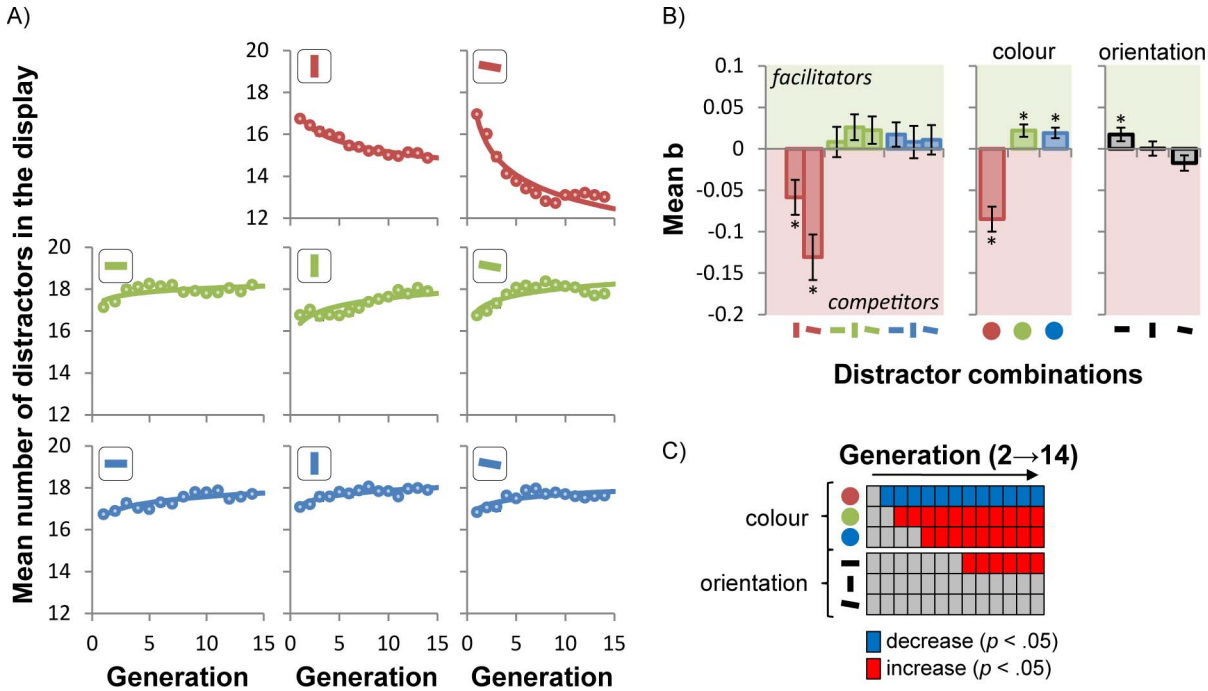

Figure 6. Evolution of distractor features in Experiment 2. (A) Mean number of items of each distractor combination in the display as a function of generation. (B) The group mean value of parameter *b* as a function of each distractor combination and as a function of each color and orientation. A negative value for *b* signifies that these distractors were detrimental for search (red shaded area), whereas a positive *b* indicates that these distractors facilitated visual search (green shaded area). Asterisks (*) indicate distractors that were significantly different from zero (see Table 2 for a summary of all statistical results). Error bars show ±1 standard error of the group mean ($n = 27$). Continuous lines represent best fitting power functions. (C) Time course of the evolution process in Experiment 2a. Reliable differences between each generation compared to the initial (random) generation: Blue shading signifies that the number of distractors possessing a particular feature decreased significantly; red indicates a significant increase, $p < 0.05$. The top three rows show effects of each color (collapsed across orientation), and the bottom three rows show effects of orientation (collapsed across color).

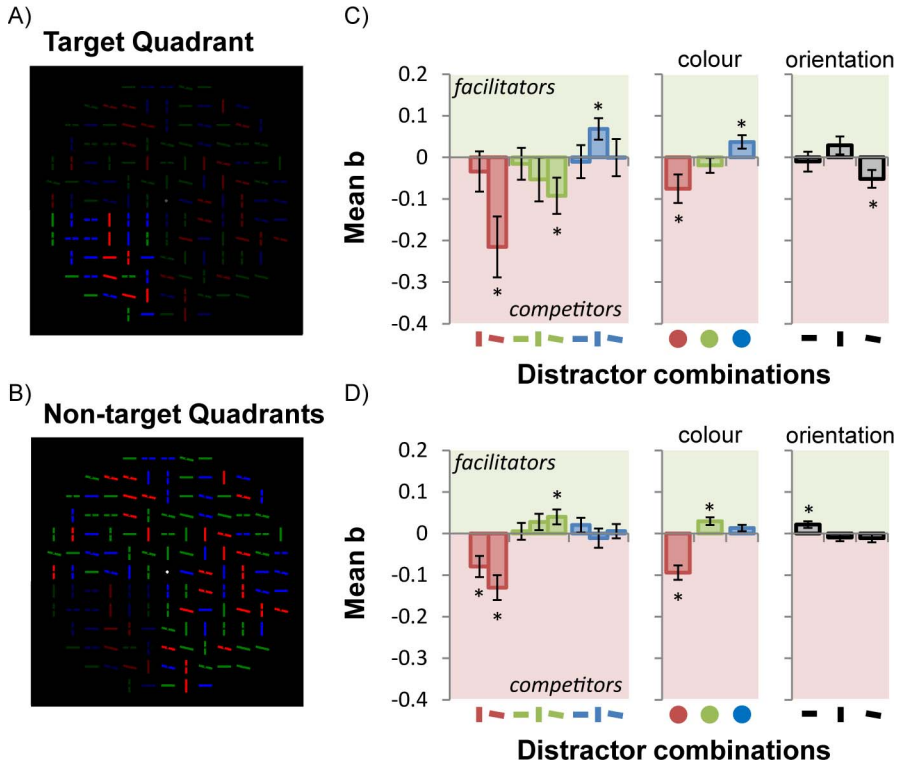

Figure 7. (A) An illustration of the target quadrant. A total of 27 distractors (plus the target) fell within the target quadrant (lines located on the vertical and horizontal were not included). (B) The remaining lines formed the nontarget quadrants and numbered 108. (C–D) Mean parameter *b* as a function of each distractor combination and as a function of each color and orientation, for the target quadrant (panel C) and the nontarget quadrants (panel D). A negative value for *b* signifies that these distractors were detrimental for search (red shaded area), whereas a positive *b* indicates that these distractors facilitated visual search (green shaded area). Asterisks (*) indicate that those distractors were significantly different from zero, $p < 0.05$ (see Table 3 for a summary of all statistical results). Error bars show ± 1 standard error of the group mean ($n = 27$). Continuous lines represent best fits.

degree of task difficulty (without rotation, finding the target would be trivial after trial one) while also maintaining the fixed target and distractor locations needed to examine local and global effects. Rotation only affected the retinotopic location of each element, not the identity (e.g., a horizontal distractor remained horizontal after rotation), and a circular array was used to remove rotation cues. Examples of two identical but rotated displays are shown in Figure 4A and B. The display contained 136 elements and had a radius of 6°, and the target was always presented on an imaginary circle centered on the display with a radius of 3.9°.

In Experiment 2b, the participants performed the same task as in Experiment 2a except that we did not apply the GA to evolve the stimuli from one generation to the next. Instead, as a control, we used 12 new random displays for each generation. After a practice block, participants did four experimental blocks (equivalent to four generations). The displays were equivalent to the displays in Experiment 2a.

Analysis

In Experiment 2a, to facilitate visualization of the stimulus as it evolved, each generation's image was smoothed using a two-dimensional Gaussian filter ($SD = 1.8°$) as in Figure 8. The size of the red cluster around the target was estimated by finding the circular region in the display with maximum mean proportion of red, with location and radius free to vary. This produced a circle centered over the red cluster with a radius of 1.37° (centered 4.5° from fixation) and enclosing four elements.

Results

Reaction time

Figure 5 shows the change in correct mean RT over generations. The black circles signify displays subject to the genetic algorithm (Experiment 2a); red triangles, displays in which each generation is a randomly generated combination of distractor features not subject to the genetic algorithm (Experiment 2b).

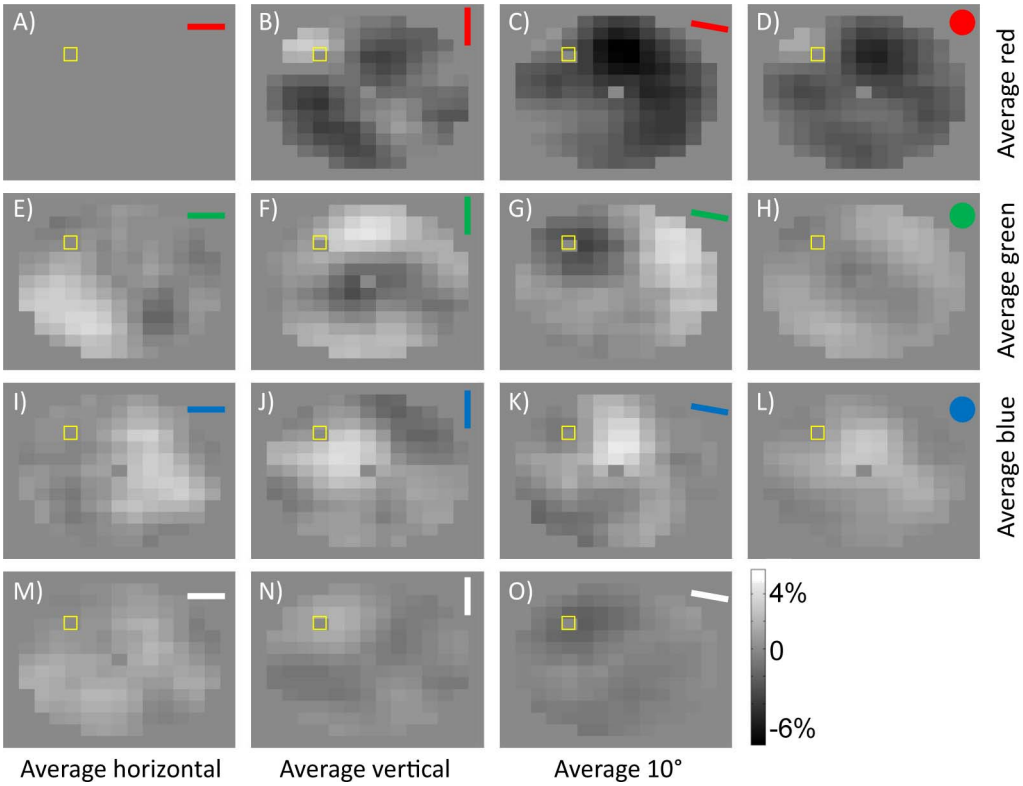

Figure 8. Evolution of distractor combinations. An image of the final evolved stimulus (corresponding to Generation 14; see also Supplemental Movie S1), showing how the number of elements changes over the evolution process compared to the first (random) generation. White indicates an increase in the number of elements for Generation 14 compared to the first (random) generation, while black indicates a decrease. The yellow box indicates the location of the target (i.e., the horizontal red line) in stimulus-centered coordinates. The three right panels represent increments and decrements in the number of red, green, and blue distractors (top, middle, and bottom, respectively) collapsed over distractor orientation. The three bottom panels represent increments and decrements in the number of horizontal, vertical, and 10° distractors collapsed over distractor color.

Overall mean error rates were low (Experiment 2a: 2.4%; Experiment 2b: 2.9%) and not further analyzed.

For each participant, we fitted power functions to the reaction time data over generations to examine whether the GA was effective or not. Overall, parameter a was 4650 ms in Experiment 2a and 4814 ms in Experiment 2b. This difference was not significant, $t(39) = 0.36$, $p = 0.757$ (independent-samples t test), indicating that there was no reliable difference regarding RTs to the first, random generation. Overall, parameter b was significantly smaller in Experiment 2a (–0.40) than in Experiment 2b (–0.07), $t(39) = 8.6$, $p < 0.0001$ (independent-samples t test), indicating that the application of a GA had a major impact on search performance. Moreover, when displays were subject to the GA, participants responded faster over successive generations, as indicated by parameter b being significantly smaller than zero, $t(26) = 18.4$, $p < 0.0001$ (one-sample two-tailed t test). In contrast, when displays were randomly determined, parameter b was not significantly different from zero, $t(13) = 2.1$, $p = 0.056$, indicating that performance did not improve across generations.

Evolution of distractor combinations

Figure 6A illustrates how the number of each distractor combination evolved over generations (collapsed over location).

As in Experiment 1, power functions (Equation 1) were fitted for each of the eight distractor combinations, another three analyzing color with orientation collapsed, and three more analyzing orientation with color collapsed. This was done for each subject to examine whether display properties increased or decreased over successive generations. Parameter a reflects the mean number of initial elements in the display, and b once again reflects the increase or decrease in the number of elements in the display over successive generations. The mean b for each distractor combination is shown in Figure 6B.Two-tailed t tests were conducted on parameter b for each display property. Table 2 summarizes these results.

Figure 6A and B shows that while the green and blue distractors increased in number over generations, $p = 0.007$ and 0.005, respectively, the red distractors decreased, $p < 0.0001$, indicating that red distractors

Distractor combination				
Color	Orientation	*a*	*b*	*t*(5)
Red	90°	17.0	−0.06	2.8*
Red	10°	16.9	−0.13	4.7**
Green	0°	17.5	0.01	0.5
Green	90°	16.4	0.03	1.7
Green	10°	17.0	0.02	1.4
Blue	0°	16.7	0.02	1.2
Blue	90°	17.2	0.01	0.4
Blue	10°	17.0	0.01	0.6
Red		16.9	−0.08	5.6**
Green		16.9	0.02	3.0*
Blue		16.9	0.02	3.1*
	0°	17.1	0.02	2.1*
	90°	16.8	0.00	0.0
	10°	16.9	−0.02	1.8

Table 2. Summary of curve-fitting results from Experiment 2. Table 2 illustrates the mean value for parameters *a* and *b* for all distractor combinations. T-values are shown for each test of whether the number of elements in the display significantly increased (i.e., a positive *b*) or decreased (i.e., a negative *b*). * $p < 0.05$, ** $p < 0.005$.

were a source of competition in searching for the target. Consistent with Experiment 1, this decrease was strongest for red 10° distractors, $p < 0.0001$, and somewhat less strong for red vertical distractors, $p = 0.01$. Furthermore, collapsing across color, the number of horizontal elements increased as a function of generation, $p < 0.05$. Thus, with larger set sizes and rotated displays, the data generally replicate those of Experiment 1 in showing an unexpected increase in horizontals even though the target was also horizontal.

Time course of the evolution process

The current results provide a unique insight into how the brain solves visual search over time: It reveals the features that are prioritized and the sequence in which selection occurs. Figure 6C shows the time course of evolution for Experiment 2a for each feature, with reliable differences in the number of distractors between each generation and the initial random generation indicated by color shading. Significant increases are shown in red and decreases in blue, with each shaded square indicating that the same change occurred over two sequential generations at $p < 0.05$ (for a similar analysis where multiple sequential comparisons are required, see also Van der Burg, Talsma, Olivers, Hickey, & Theeuwes, 2011). The figure clearly shows that observers select color over orientation, as color change was the first to evolve. Red, green, and blue became reliable at Generations 3, 4, and 6, respectively, while the only significant change in orientation occurred later, with horizontal elements increasing from Generation 9.

Spatial effects

To investigate whether spatial factors around the target affected search, we compared the density of distractor features within the target quadrant with the three nontarget quadrants (see Figure 7A and B, respectively). Once again, but now separately for the target quadrant and nontarget quadrants, power functions (Equation 1) were fitted for each of the eight distractor combinations, another three analyzing color with orientation collapsed, and three more analyzing orientation with color collapsed. Mean parameter *b* was subjected to one-sample *t* tests to examine whether the number of distractors increased or decreased significantly (see Table 3 for an overview of mean parameters *a* and *b* as well as results from statistical tests). Figure 7C and D provides a summary of parameter *b* to illustrate the competitors and facilitators in the displays.

The target quadrant contained 27 distractors as well as the target, while nontarget quadrants contained 108 distractors (to standardize the number of elements, the nontarget quadrants were multiplied by 0.25). These analyses yielded some global effects and some location-specific effects. Globally, red distractors of both kinds (90° and 10°; Figure 6A, upper panels) became less numerous across generations, $ps < 0.05$, indicating that both kinds of red distractors competed with the target. Surprisingly, however, in the target quadrant only 10° red distractors declined in number, $p < 0.01$, while red verticals remained constant, $p = 0.489$. Furthermore, whereas blue verticals in the target quadrant increased in number over generations, $p < 0.05$, green distractors increased as well, but only in nontarget quadrants, $p < 0.005$. As in Experiment 1, the total number of horizontal lines (independent of color) increased over generations, but here we show that it occurs in nontarget quadrants, $p < 0.05$, and not target quadrants, not significant. Table 3 summarizes all statistical analyses.

Supplementary Movie S1 illustrates how each distractor combination in Experiment 2 evolves at each (target-centered) stimulus location (see also Figure 8 for an image of the final generation). This movie shows changes in the number of distractors relative to the first (random) generation, with white indicating an increase and black a decrease in element numbers. The small yellow box in each frame indicates the target location in object-centered coordinates. Each generation's image was smoothed using a Gaussian filter to facilitate visualization. The last frame of Supplementary Movie S1 is shown in Figure 9, with each panel showing a spatial map of significant correlations (Bonferroni

Distractor combination		Target quadrant			Nontarget quadrants		
Color	Orientation	*a*	*b*	*t*(26)	*a*	*b*	*t*(26)
Red	90°	3.4	−0.03	0.7	3.4	−0.08	3.1**
Red	10°	3.4	−0.21	2.9**	3.4	−0.13	4.3**
Green	0°	3.5	−0.02	0.4	3.5	0.01	0.3
Green	90°	3.5	−0.05	1.0	3.3	0.03	1.4
Green	10°	3.4	−0.09	2.1*	3.4	0.04	2.3*
Blue	0°	3.4	−0.01	0.3	3.3	0.02	1.2
Blue	90°	3.6	0.07	2.7*	3.4	−0.01	0.5
Blue	10°	3.3	0.00	0.0	3.4	0.01	0.3
Red		3.3	−0.08	2.2*	3.4	−0.09	5.4**
Green		3.4	−0.02	1.1	3.4	0.03	3.1**
Blue		3.4	0.04	2.3*	3.4	0.01	1.7
	0°	3.5	−0.01	0.4	3.4	0.02	2.7*
	90°	3.4	0.03	1.3	3.4	−0.01	1.0
	10°	3.3	−0.05	2.4*	3.4	−0.01	1.1

Table 3. Overall spatial results from Experiment 2a. Table 3 illustrates the mean value for parameters *a* and *b* for all distractor combinations, comparing target and nontarget quadrants of the display. T-values are shown for each test of whether the number of elements in the display significantly increased (i.e., a positive *b*) or decreased (i.e., a positive *b*). * $p < 0.05$, ** $p < 0.005$.

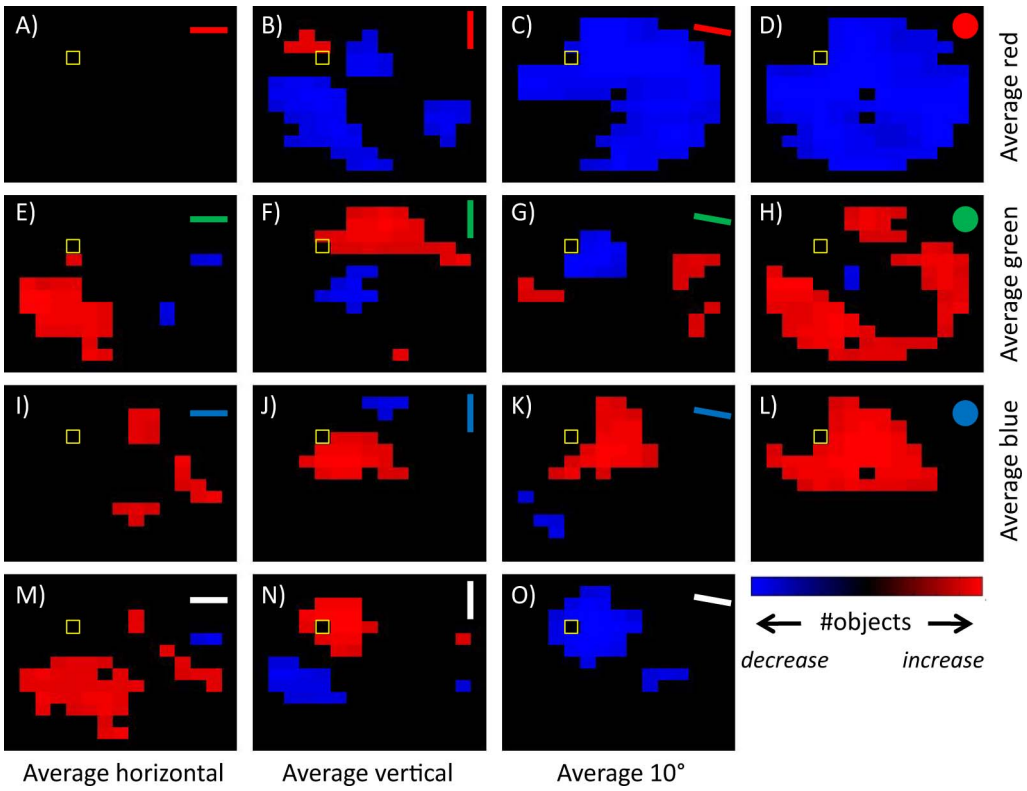

Figure 9. Correlations between the prevalence of a distractor type and search times, spatially mapped with the yellow box indicating target location. Blue and red indicate significant decreases and increases, respectively, in distractor prevalence (Bonferroni corrected), *ps* < 0.0004. Panels D, H, and L show correlations for red, green, and blue, respectively, collapsed over orientation; panels M, N, and O show the correlations for horizontal, vertical, and 10° distractors, collapsed over color. Figure 9 is the final frame from Supplementary Movie S1.

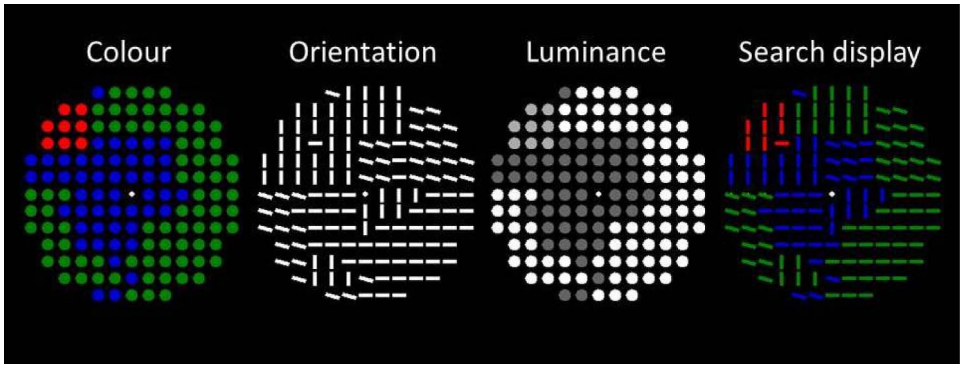

Figure 10. The color, orientation, and luminance dominance maps show the most dominant feature for each location after 14 generations. The far right image is the combination of all three maps. The target is located in the top-left corner of each featural set. For the purposes of computing the dominance maps, the global intertrial rotation was subtracted from each image so that the target is always in the upper left quadrant.

corrected), $ps < 0.0004$, between the number of each distractor combination (rows = color; columns = orientation) and RTs at each stimulus location. Blue pixels show that decreasing numbers of a given distractor at a particular location correlated significantly with decreases in RT (i.e., better search), whereas red pixels show the opposite.

We created dominance maps to gain a better understanding of how observers searched through our displays. Supplementary Movie S2 illustrates the evolution of the most dominant color, orientation, and luminance (included because each color differed in luminance: red = 28.5 cd/m^2, green = 65.4 cd/m^2, blue = 7.7 cd/m^2) for each location and generation. Figure 9 shows the last generation of Supplementary Movie S2. The dominance maps reveal, counterintuitively, that while red distractors decreased in number across generations, it was not to isolate the (red) target from competing red distractors and thus produce a color pop-out. Instead, a cluster of red verticals evolved adjacent to the target (Supplementary Movies S1 and S2; Figures 8 and 9). We estimated the red cluster's radius in the last generation at 1.37°, centered 4.50° from fixation, an area encompassing four elements. The number of red elements in this cluster increased significantly over generations (mean $b = 0.16$, $p < 0.05$). We propose that this red cluster may facilitate efficient search in two ways. First, as the target was peripheral to initial fixation, and early visual receptive fields increase in size with eccentricity, a red cluster would drive a "red" response better than a small isolated target. This would signal the target's approximate location. Subsequent fixations in the direction of this cluster would therefore predictably permit a quick orientation-based target search. Second, the red cluster of four elements around the target might facilitate top-down search for information signifying target featural identity ("red," in this case). Visual search studies indicate that approximately four objects can be attended to at once (Van der Burg, Awh, & Olivers, 2013) and that search is probably conducted in subsets of about four objects (Anderson, Vogel, & Awh, 2013). Our experiments cannot differentiate these possibilities, although future studies manipulating search display density could examine this. For example, increased display density should not affect the cluster size if it is constrained by the receptive field size. By contrast, if the number of elements is the determining factor, then the retinotopic area associated with the cluster should shrink if the display becomes more dense, so that the total number of its constituent elements is preserved.

Dominance maps also reveal that vertically oriented distractors increased in number around the target across generations. By the final generation, nearly the entire target quadrant contained vertical elements (see orientation map, Figure 10). This agrees with findings showing that large differences in target and distractor orientation promote target detection by reducing the deleterious effects of crowding (Whitney & Levi, 2011), and with visual search results showing that it is easier to detect a horizontal line among verticals (Wolfe & Horowitz, 2004).

Another interesting observation is that a large region of horizontal elements (mainly green) evolved in the quadrant opposite to that containing the target (see orientation map, Figure 10). The large number of horizontals is counterintuitive, given that they share the target's orientation, which should impair search. While this clustering into a single large region is curious, it may explain why so many horizontals were retained. Orientation-selective neurons in primary visual cortex are strongly suppressed by iso-oriented stimuli surrounding their receptive fields (surround suppression; DeAngelis, Freeman, & Ohzawa, 1994; Shushruth et al., 2013), an effect that is stronger for horizontal than oblique and vertical stimuli (Mannion, McDonald, & Clifford, 2010). The evolution of uniform horizontal fields away from the target, therefore, would help

suppress neural response at target-irrelevant regions and thereby boost target salience.

A final unexpected finding was the increasing number of blue elements over generations (Figure 9L; luminance map in Figure 10) and their peculiar spatial distribution, forming a path leading from fixation to target (Figures 9L and 10). Because the blue elements had the lowest luminance, this effectively forms an oriented luminance path pointing to the target from fixation. We predict that the direction of this luminance path may serve to direct saccadic eye movements in the direction the target. Moreover, given the luminance variation associated with the different colors in our displays, the higher luminance red cluster adjoining the target generates a higher luminance contrast when bounded by blue elements than it would if surrounded by green elements or red elements, which may also benefit search performance.

Discussion

Our implementation of a GA reverses the traditional logic of visual search experiments. Rather than the experimenter deciding what information is important for the task, the GA allows the observer's brain to select what visual information, and in what sequence, is most relevant. It can be said, therefore, that the stimulus adapts to suit the brain. This paradigm provides new insights into human visual search performance unconstrained by previous theories. Being highly efficient, the GA permits investigation of complex displays, which would be implausible using conventional approaches. Importantly, the GA was unconstrained and yet still replicated several previous findings while also revealing new and counterintuitive findings not predicted by current models (Duncan & Humphreys, 1989; Treisman & Gelade, 1980; Wolfe, 1994). These included (a) the emergence of a small cluster of red vertical distractors around the red target (Experiment 2), (b) a blue path emanating from fixation to target (Experiment 2), and (c) an increase in horizontal distractors (Experiments 1 and 2). We propose that these outcomes reveal the brain's own search strategies when seeking a target in a complex visual display. Importantly, the GA reveals not only *which* contextual features are prioritized by the brain but also the order in which this occurs. We found that color evolved first, followed by orientation. This suggests that color was the primary feature constraining search performance in our complex displays.

This study demonstrates the validity of a GA approach to studying visual search. The efficiency of GAs means that future studies may be able to investigate far more complex visual search environments that are more representative of natural scenes. While other applications of GAs have been aimed at optimizing stimuli (Verma & McOwan, 2009; Wager & Nichols, 2003) or simulating human behavior (Zhang & Eckstein, 2010), our study is the first to use human performance behavior to generate new displays revealing the brain's own search strategies.

In the present study, we used a "survival of the fittest" principle for the selection of the displays (i.e., the best displays survived). As a result, search improved over generations as the competitors in the display disappeared systematically. Whereas this principle worked very well in our case, it is important to note that different display solutions may evolve in response to different task-specific demands or stimulus constraints. One might also consider varying aspects of the GA as conditions require. For instance, if the initial displays were too easy and supported quick search times from the beginning, a selection criterion of "survival of the worst" could be employed (i.e., selecting the worst displays for reproduction) to devolve the stimuli into less efficient displays. Furthermore, although we evolved the displays based on the participants' mean RTs, it is also possible that other fitness values such as accuracy or eye movements (e.g., precision or velocity) could be used to select the best or worst displays. One advantage of the mean RT is that there is considerable trial-to-trial variation. Such short-term variance would be unlikely when using accuracy, necessitating more trials and reducing the efficiency of the genetic algorithm. Another issue to consider is the particular choice of parameters used in the genetic algorithm (mutation factor, crossover procedure). Therefore, more research is needed that seeks to understand the role of each parameter used in the GA. We anticipate that our approach can be widely adapted to a variety of cognitive and neuroscientific questions considered too intractable using conventional paradigms.

Keywords: visual search, attention, evolution, genetic algorithm

Acknowledgments

This work was supported by Australian Research Council Grant DE130101663 (EVdB) and DP120101474 (DA and JC).

Commercial relationships: none.
Corresponding author: Erik Van der Burg.
Email: vanderburg.erik@gmail.com, erik.vanderburg@sydney.edu.au.
Address: School of Psychology, University of Sydney, Sydney, Australia.

Footnote

[1] Although the power function works satisfactorily to quantify the efficiency of the genetic algorithm, it is important to note that this model is theoretically impossible, as the RTs will never converge to 0 ms. One suggestion that may solve this is to use a three-parameter equation with an intercept. However, in cases in which there is no evolution (i.e., concomitant with a reduction in RTs), the best fit will be a flat line whose value equals that of the intercept. As a result, parameter *a* will be equal or close to zero, and parameter *b* can have noninterpretable or extreme values.

References

Anderson, D. E., Vogel, E. K., & Awh, E. (2013). A common discrete resource for visual working memory and visual search. *Psychological Science, 24*(6), 929–938.

Awh, E., Belepolsky, A., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences, 16*(8), 437–443.

Bex, P. J., Solomon, S. G., & Dakin, S. C. (2009). Contrast sensitivity in natural scenes depends on edge as well as spatial frequency structure. *Journal of Vision, 9*(10):1, 1–19, http://www.journalofvision.org/content/9/10/1, doi:10.1167/9.10.1. [PubMed] [Article]

Cass, J., Van der Burg, E., & Alais, D. (2011). Finding flicker: Critical differences in temporal frequency capture attention. *Frontiers in Perception Science, 2*(320), 1–7.

DeAngelis, G. C., Freeman, R. D., & Ohzawa, I. (1994). Length and width tuning of neurons in the cat's primary visual cortex. *Journal of Neurophysiology, 71*(1), 347–374.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Folk, C., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 1030–1044.

Foster, D. H., & Ward, P. A. (1991). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society B, 243,* 75–81.

Holland, J. H. (1975). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. Oxford, UK: University of Michigan Press.

Humphreys, G. W., Quinlan, P. T., & Riddoch, M. J. (1989). Grouping processes in visual search: Effects with single- and combined-feature targets. *Journal of Experimental Psychology: General, 118,* 258–279.

Kaptein, N. A., Theeuwes, J., & Van der Heijden, A. H. C. (1995). Search for a conjunctively defined target can be selectively limited to a color-defined subset of elements. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1053–1069.

Mannion, D. J., McDonald, J. S., & Clifford, C. W. G. (2010). Orientation anisotropies in human visual cortex. *Journal of Neurophysiology, 103,* 3465–3471.

Nakayama, K., & Silverman, G. H. (1986). Serial and parallel processing of visual feature conjunctions. *Nature, 320,* 264–265.

Shushruth, S., Nurminen, L., Bijanzadeh, M., Ichida, J. M., Vanni, S., & Angelucci, A. (2013). Different orientation tuning of near- and far-surround suppression in macaque primary visual cortex mirrors their tuning in human perception. *Journal of Neuroscience, 33*(1), 106–119.

Srinivas, M., & Patnaik, L. M. (1994). Adaptive probabilities of crossover and mutation in genetic algorithms. *IEEE Transactions on Systems, Man and Cybernetics, 24*(4), 656–667.

Theeuwes, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics, 51,* 599–606.

Treisman, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance, 8,* 194–214.

Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114,* 285–310.

Van der Burg, E., Awh, E., & Olivers, C. N. L. (2013). The capacity of audiovisual integration is limited to one item. *Psychological Science, 24*(3), 345–351.

Van der Burg, E., Olivers, C. N. L., Bronkhorst, A. W., & Theeuwes, J. (2008). Pip and pop: Non-spatial auditory signals improve spatial visual search. *Journal of Experimental Psychology: Human Perception and Performance, 34*(5), 1053–1065.

Van der Burg, E., Talsma, D., Olivers, C. N. L.,

Hickey, C., & Theeuwes, J. (2011). Early multisensory interactions affect the competition among multiple visual objects. *NeuroImage, 55*(3), 1208–1218.

Verma, M., & McOwan, P. W. (2009). Generating customised experimental stimuli for visual search using genetic algorithms shows evidence for a continuum of search efficiency. *Vision Research, 49,* 374–382.

Wager, T. D., & Nichols, T. E. (2003). Optimization of experimental design in fMRI: A general framework using a genetic algorithm. *NeuroImage, 18*(2), 293–309.

Whitney, D., & Levi, D. M. (2011). Visual crowding: A fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences, 15*(4), 160–168.

Wolfe, J. M. (1994). Guided search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review, 1,* 202–238.

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience, 5,* 1–7.

Zhang, S., & Eckstein, M. P. (2010). Evolution and optimality of similar neural mechanisms for perception and action during search. *PLoS Computational Biology, 6*(9), e1000930.