



Published in final edited form as:

J Exp Psychol Anim Behav Process. 2013 October ; 39(4): 357–376. doi:10.1037/a0034074.

The Adaptive Analysis of Visual Cognition using Genetic Algorithms

Robert G. Cook and Muhammad A. J. Qadri

Tufts University

Abstract

Two experiments used a novel, open-ended, and adaptive test procedure to examine visual cognition in animals. Using a genetic algorithm, a pigeon was tested repeatedly from a variety of different initial conditions for its solution to an intermediate brightness search task. On each trial, the animal had to accurately locate and peck a *target* element of intermediate brightness from among a variable number of surrounding darker and lighter *distractor* elements. Displays were generated from six parametric variables, or genes (distractor number, element size, shape, spacing, target brightness, distractor brightness). Display composition changed over time, or evolved, as a function of the bird's differential accuracy within the population of values for each gene. Testing three randomized initial conditions and one set of controlled initial conditions, element size and number of distractors were identified as the most important factors controlling search accuracy, with distractor brightness, element shape, and spacing making secondary contributions. The resulting changes in this multidimensional stimulus space suggested the existence of a set of conditions that the bird repeatedly converged upon regardless of initial conditions. This psychological "attractor" represents the cumulative action of the cognitive operations used by the pigeon in solving and performing this search task. The results are discussed regarding their implications for visual cognition in pigeons and the usefulness of adaptive, subject-driven experimentation for investigating human and animal cognition more generally.

The experimental investigation of behavior and cognition has been of intense interest for over a century. The preponderance of these investigations have used a parametric factorial approach that involves manipulating one or more independent variables over a selected range of values to observe the effects on participants' behavior (Fisher, 1935). Observations are typically followed by additional experiments conducted in the same way to further investigate the question at hand. When combined with the testing of multiple hypotheses (e.g., Kuhn, 1996; Platt, 1964), such factorial experiments have revealed much about the nature of cognition in human and non-human animals.

Despite its success, this sequential factorial approach to behavioral science has several issues and disadvantages. First, this conventional method is best suited for the exploration of psychological spaces of low dimensionalities, due to the small number of variables and values that can be managed at any one time in such designs. This limitation has recurrently raised concerns about the ecological validity of such controlled experimental approaches

(e.g. Neisser, 1976). Further, its success is highly dependent on the experimenters' initial selection of the variables and values to investigate. While this gives experimenters tight control over the situation and its replication across subjects and experiments, miscalculations can have a number of negative empirical and theoretical consequences. For example, if inappropriate or unrelated variables to a phenomenon are investigated, null results are likely. Contrastingly, if variables are not manipulated over a sufficient range, aspects of a phenomenon may be mischaracterized or missed entirely. Such problems are further compounded when testing animals, as the selection and implementation of the variables is frequently biased by our anthropocentric perspective and intuitions about what animals might or should do. Additionally, because the experimenter directly interacts with the situation, this approach is more susceptible to different forms of experimenter effects and biases (Barber & Silver, 1968; Nickerson, 1998; Orne, 1962). Finally, this approach is minimally interactive for the participants. Much of the richness of behavior is removed by allowing only indirect interaction via measured behavior (most often a binary choice or reaction time in cognitive studies).

Less frequently employed in the study of cognition have been experimental approaches that use more interactive, open-ended, and dynamic methods. Free response production tasks are the most open-ended of such approaches, where the participants have free reign in responding to the experimenter's instructions and constructed situation. The results of such tasks are typically revealing and informative (e.g., Goodenough, 1963). Their open-ended nature and complexity, however, often makes them difficult to quantify, analyze, and systematize. They are also difficult to conduct with animals, although research with controlled reproduction does mimic some of its features (Basile & Hampton, 2011).

A controlled manifestation of the same logic is the classic psychophysical methods of adjustment (Fechner, Howes, & Boring, 1966). In these, a participant interactively adjusts a condition until it is equivalent (or not) to an established standard or reference. Such interactive procedures can rapidly produce results about the equivalence and discriminability of stimuli. Variations of this adjustment-based approach have even been used with animals. In these cases, animals are cycled through variable conditions and compared against a reference value that is adjusted as a function of the animal's responses. This type of adjustment approach has been used with animals to examine temporal discounting and self-control, for example, where the length of a delay is dynamically adjusted relative to a reference to establish the equivalence and relation of time and reinforcement value (e.g., Mazur, 1987, 1996; Richards, Mitchell, De Wit, & Seiden, 1997). A few animals have been taught to directly adjust stimuli in the same manner as humans (e.g. D. S. Blough & Schrier, 1963).

Here, we test a new adaptive approach to the experimental measurement of cognition. It involves a continuous and dynamic adjustment method capable of generating open-ended, subject-driven resolutions to experimental tasks. The procedure relies on advances in evolutionary computing to implement a genetic algorithm (GA) as a means to continuously explore and evaluate a subject's experimental responses to a discriminative context. Genetic algorithms are search heuristics with evolutionary properties that can search among and evaluate a wide variety of candidate solutions to a problem independent of the shape of the

problem space. GAs have been having an increasingly widespread impact across many areas of science and engineering, including biology, economics, bioinformatics, robotics, and machine learning. Researchers are also beginning to use them to explore and solve a number of questions related to the optimization of experimental designs (Cavagnaro, Myung, Pitt, & Kujala, 2010; Gallo et al., 2006; Müller, Berry, Grieve, Smith, & Krams, 2007).

Their impact in psychology has been minimal, however. To our knowledge, they have not been much utilized for control and measurement in experimental settings. Bond and Kamil (1998, 2002; Bond & Kamil, 2006) have used a genetic algorithm to examine the effects of predatory selection in a “virtual ecology”. In their studies, blue jays searched for digital “moths” on a computer screen. Analogous to the real world, if a moth’s stimulus features were detected and then “eaten” by the jay, that moth did not transmit its stimulus features/genes to the future generations of trials. Using this technique, Bond and Kamil were able to examine predator-generated, frequency-dependent crypsis and polymorphism in their prey populations of moths as a result of the blue jays differential selection pressures (Bond, 2007; Kamil & Bond, 2002). Their primary focus, however, has been on the impact of selection on the structure of prey populations rather than the nature of cognition in the blue jays. Our idea was to turn these powerful techniques to investigate cognition more specifically, especially for designs evaluating stimulus spaces of high dimensionality.

For any discrimination, there is a solution space that represents all the potential resolutions to a task and their relative probabilities of occurrence. Because understanding such complex psychological spaces are critical to a complete account of behavior in the real world, the development of procedures to efficiently explore, map and analyze them is an acute and important need. The piecemeal nature of parametric experiments is not particularly well suited for the investigation of this kind of complex space. GAs offer a possible alternative method for this type of analysis.

Here we report results where a GA was used to continuously control and evaluate the discriminative behavior of a pigeon in a complex multidimensional stimulus space. By having the GA dynamically react, adjust, and track large number of parameters in parallel, we can capture the best traditions of simpler adaptive psychophysical techniques with the power of modern computing technology for an unbiased and open-ended investigation of behavior within an environment containing higher and more realistic degrees of dimensionality and complexity than heretofore possible.

In the current experiment, for the first time, we evaluated how one animal processed this kind of open-ended visual discrimination, and how the dynamic interaction between subject and algorithm altered the discrimination over time. Because GAs use concepts derived from evolution and genetics, this metaphor is quite useful for describing our experimental procedures. The composition of the stimulus display for each trial was determined by a *genetic code* or *genome*. Separate *genes* encoded information about the shape, size, brightness, number and spacing of the display’s visual elements. If the bird successfully discriminated a trial’s display by locating, identifying, and pecking at the designated target, that display’s genome was allowed to *reproduce with variation*, and its descendants included in the *population* of trials were used to compose future sessions. If the bird was unsuccessful

with a display, then its genetic material was not reproduced. Accordingly, over successive sessions, the composition of the displays evolved as driven by the pigeon's selection behavior. The resulting changes in the factors can be tracked to isolate those display variables (genes) critical to the animal's discrimination and their relative importance. The objectives of this research were to examine the evolution of these discriminative properties and their implications for the underlying visual and cognitive mechanisms used by birds in solving these types of search problems, and to investigate how adaptive techniques can be employed for addressing experimental questions in animals.

Experiment 1

Experiment 1 tracked the ability of a pigeon to perform a GA-controlled visual search task. In the task, the pigeon had to locate and identify a variably located, shaped, and sized *target* element of intermediate brightness from a variable number of surrounding darker and lighter *distractor* elements. Locating, identifying and pecking the target resulted in food reward, while incorrect choices to the distractors did not. After the randomly-generated first session, the GA determined the composition of all trials based on the pigeon's selection behavior.

Figure 1 diagrams the procedures used to generate the trials for each session. Different independent genes controlled the brightness, shape, size, spacing, and number of elements in each display. If the bird successfully located the target in a display, the genetic "descendants" of this *selected* trial were reproduced with variation for probabilistic inclusion in future trials. Unsuccessful trials, in which the pigeon did not find the target or failed to respond, were not reproduced. All sessions tested a new *generation* of trials created from the descendant of *selected* displays and *new* or *random* displays that were stochastically created from the display genes. Shown in Figure 2 is a sample of the variety of displays that could be randomly generated from this genome.

This particular visual search task was used in the initial investigation for several reasons. First, it would likely be easy for the bird to learn. Second, its display properties could be varied over a wide range of dimensions and values to explore the pigeon's selection behavior. Third, a solid foundation of previous research involving visual search processes already existed for pigeons using conventional factorial designs (D. S. Blough, 1982, 1986, 1990, 2001; P. M. Blough, 1989; Cook, 1992, 2000; Cook & Hagmann, 2012; Katz & Cook, 2003; Spetch & Friedman, 2006; Spetch & Mondloch, 1993). These provided reliable reference points for understanding our results in this new approach.

All discriminations possess a multidimensional space of potential solutions. We were interested in identifying and isolating the pigeon's potential solution(s) to this visual search task based on how the populations of the displays' parameters or genes changed with selection over time. Before the experiment, one could imagine a range of plausible ways for the pigeon to solve the task (Katz & Cook, 2003). Given their general local bias (Cavoto & Cook, 2001; Lea, Goto, Osthaus, & Ryan, 2006), for instance, the pigeon might solve the task using only small elements with large numbers of homogeneous distractors surrounding it, so that the target's oddity readily emerges (Cook, 1992; Zentall, Hogan, Edwards, & Hearst, 1980). Alternatively, the pigeon might prefer a task with elements that have large

areas to compare or having few distractors in order to reduce search load (D. S. Blough, 1977; D. S. Blough & Blough, 1997). Alternatively, perhaps both solutions could co-evolve and a new and unexpected solution could emerge. The advantage of the open-ended adaptive procedure is that we should be able to determine the structure of the animal's solution from the impact of its selection behavior, rather than from our own closed selection of the variables to manipulate. If in the resulting solution there were such a set of conditions that the bird repeatedly converged upon regardless of initial conditions, it would strongly suggest the existence of a *psychological attractor* within the solution space of this task. Attractors are sets of conditions over which dynamical systems tend to evolve towards regardless of initial conditions. While this terminology is not widely used in psychology, it is a useful framework for thinking about the organization of complex dynamical systems, such as pigeons and the cognitive operations they bring to bear on the solution of a discrimination.

Experiment 1 consisted of three cycles of testing from randomly generated initial conditions. After the initial acquisition phase, the GA conditions were *reset* twice such that there were only random differences in the populations of values for each display gene. This resetting allowed us to determine whether the attractor located during the terminal portion of the acquisition phase was reliable and consistent (cf. Wright & Cumming, 1971).

Method

Animals—A male White Carneaux pigeon (*Columba livia*) was tested. This bird had participated in two unrelated visual choice procedures. The bird was maintained at 85–90% of its free-feeding weight and individually caged in a colony room (12 hour LD cycle) with free access to water and grit.

Apparatus—Testing was conducted in a flat-black Plexiglas chamber. Stimuli were presented by computer on an LCD color monitor (NEC MultiSync LCD1525X), which was visible through a 29.0 × 21.5 cm viewing window in the middle of the front panel of the chamber. Pecks to the monitor (1024 × 768 pixels) were detected by an infrared touch screen (Carroll Touch Systems). The viewing window's bottom edge was 18.5 cm above the chamber floor. Mixed grain was delivered through a centrally placed food hopper (Coulbourn Instruments, Whitehall, PA). A 28-volt houselight was centrally located in the chamber ceiling and was always illuminated.

Procedure

Intermediate brightness search task: In this task, the pigeon had to locate and identify a variable *target* element of intermediate brightness (gray) from a variable number of surrounding *distractor* elements of higher or lower brightness. Each trial in a session began with a peck to a centrally located 2.5 cm circular white ready signal. This was replaced by a search display determined and controlled by a GA procedure (detailed below). Within this display, there was a *target* element of variable intermediate brightness, defined by the middle 20% of the computer's brightness range (without gamma correction). This gray target was presented in a context of surrounding *distractor* elements that randomly varied in brightness from above (60–100%) or below (0–40%) this range. A correct response was

defined as locating and pecking at the target element three times, turning off the display and resulting in 2.1-s access to mixed grain. An incorrect response was defined as three accumulated pecks to a distractor element, turning off the display and terminating the trial. Because our random trial generation procedure raised the possibility of displays being tested that might be difficult for the bird to respond to (e.g., very small elements), if no element had received three pecks within 20-s of display onset, the display was turned off, the trial terminated, and its outcome scored as no choice. Such *no choice* trials were treated as a non-selected trial and not reproduced because they likely represented poor display conditions for the bird. A 3-s inter-trial-interval (ITI) separated each trial.

Gene expression and display properties: Each trial's display was composed from multiple elements that varied in distractor number, element size, element shape, relative spacing, target brightness and distractor brightness. Their generation was derived from a set of independent parametric variables or *genes* that determined the display's appearance.

The *distractor number* gene controlled the number of distractor elements in the display. Each display had at least one distractor. This gene varied over a range of 0 to 1 (expressed as 1 to 10). On randomly generated trials, it was subsequently discovered that counts of 1 and 10 occurred half as often as planned due to Visual Basic's rounding operations (adjusted for in-computing chance expectations for this factor).

The element *shape* gene determined whether all of the elements of a display were squares or circles (mixed displays were not allowed). The shape gene quantitatively varied over range of 0 to 1, and used to probabilistically determine the expressed shape. For example, a gene value of .75 would give a display a 75% chance of being composed of circles, while a value of .2 had an 80% chance of being composed of squares.

The element *size* gene determined size of elements. It varied over a range 0 to 1, with larger values leading to larger elements. The phenotypic expression of this gene was monotonic, but not directly linear, with its value. Its expression was determined in combination with distractor number and element spacing genes because of the spatial packing algorithm used to generate the displays and resulted in elements that varied from 0.1 to 15.2 cm in diameter or width

The display *spacing* gene controlled the percentage of the display that these elements could occupy. This gene varied over a range 0 to 1, with larger values leading to denser displays. This gene's range started with 100% percent of the display's maximal size (30.5×23 cm) to 25% of this value (15.25×11.5 cm) as measured from its center. Elements were randomly located within the display as derived from a uniform probability distribution. Depending on the distractor number and the element size gene, a grid-based spatial packing algorithm was iteratively adjusted so that the expressed size and positions of the elements reduced their overlap to ensure that approximately 90% of each element was visible (roughly 65% of displays had an overlap between at least two elements). We did examine whether the presence of this minor overlap had any effect on choice accuracy, but found no differences related to this factor.

Each element of a display had a *brightness* gene that determined its appearance expressed as a percentage of the total brightness range of the display (0–255 gray scale values). The target's brightness was restricted to an intermediate level within the middle 20% of the total range (gray scale values 102–153). The brightness of the distractor elements were restricted to the upper 40% (>153 gray scale level) or lower 40% (<102 gray scale level) of the total range. The brightness of each distractor was determined by one gene from a set of independent genes. This biased the displays towards a mixture of different distractor brightness values to hamper the bird reducing the discrimination to an easy brightness discrimination by driving all the distractors towards one of their maximal values. To simplify analysis and discussion, we created a representative “composite” gene for distractor brightness by averaging the value of all distractors in a display. Because the distractors were typically combinations of values that were darker and lighter than the target's value, we transformed these values to a common scale by measuring their brightness relative to the closest bound of the target's range. As a result, this averaged distractor metric could vary from a value of 0 (the target's upper or lower bound) to .4 (maximally distant), with an expected value of .2 for a randomly generated set of distractors.

Session generation, organization and, genetic variation: Each session was composed of selected and random trials. *New* or *random* trials were generated by independently assigning a random value to each gene, resulting in a constant influx of highly variable trials. *Selected* trials were randomly drawn from a pool of trials derived from the pigeon's past selection behavior. Whenever the bird responded correctly, the genes of that trial were replicated four independent times with a $\pm 5\%$ random variation in each of their values. These altered “descendants” were then added to the pool of possible displays (PPD) for potential inclusion in future sessions. If the bird was incorrect or failed to respond, the trial was considered non-selected and its genome was not transmitted forward to the PPD. The PPD retained up to 250 total descendants of the five previous sessions. In the rare occasion where the PPD was emptied before all selected trials for a new session were drawn, random trials were substituted in their place.

Four 50-trial sessions were successively conducted each day over a span of approximately two hours. Each session consisted of 45 *selected* trials randomly chosen from the PPD and five randomly generated *new random* trials. Each new session allowed a generation of trials to be assessed and generated by the GA. Sessions were separated by approximately one minute for these calculations. The first seven days employed two 100-trial sessions per day that doubled this same organization of trials.

Acquisition and resetting: No preliminary training was needed because of the bird's prior experience. All trials of the first session were randomly generated, but following this, the GA determined all trial and session composition. The acquisition phase (also called cycle 1) lasted 80 sessions. After completing cycle 1, the setup was restarted twice with new randomized initial conditions. For these resets (cycles 2 & 3), the PPD was emptied and all 50 trials of the first session were randomly generated. After this, the GA again determined all trial and session composition. Cycles 2 and 3 each lasted 108 sessions.

Results

Because of the novelty of the approach, the next sections use several different approaches to describe the bird's behavior and the selection in the discrimination. The first part recounts the accuracy and selection across the three cycles. The second describes the valuation of the different genes or parameters controlling the displays as determined by the GA. The last section examines the bird's performance as a function of each gene's phenotypic expression and its effect on the underlying population of gene values. These different sections document that a highly similar set of terminal display conditions repeatedly emerged as a result of the bird's selection behavior, despite random initial conditions that did not necessitate these converging evolutions. The recurring convergence of the GA indicates the existence of a set of display conditions that best supported the bird's search and identification behavior and the relative importance of the different display factors controlling this outcome.

Choice Accuracy—The dynamics of the accuracy of the bird's search behavior over the three cycles are outlined in this first section. Because reproduction in the GA was controlled by correct choices (selected), incorrect choices, and no-choice responses (both treated as non-selected and "incorrect"), we used all three trial outcomes to compute accuracy. Overall, the pigeon had little difficulty learning this search task. Shown in Figure 3 is the pigeon's percentage of correct target identifications during acquisition (cycle 1) and the two resets to randomized initial conditions (cycles 2 & 3) in 8-session blocks. Accuracy is shown separately for *selected* (black circles) and *random* (white circles) trials.

Three separate estimates of chance-level responding are included in Figure 3 to compare to the bird's performance. For all three estimates, chance was determined by computing the expected value of accuracy based on randomly selecting among the varying number of elements in each display. The theoretical expectation for chance guessing based on the average number of elements in a display is 20.4% (dashed line). A direct empirical estimate can be computed based on the number of distractors that were actually randomly generated from just the random trials (white squares). Figure 3 confirms that this empirical value hovered at its theoretical expectation. Because the pigeon was more accurate when there were fewer distractors in the displays, and these occurred with increasing frequency because of differential selection, the estimated value for guessing based on selected trials is also included in Figure 3 (black squares). This chance expectation is computed from the actual number of distractors appearing in selected displays as adjusted by the GA. Except for a brief period at the beginning when the bird was learning the discrimination, accuracy on both selected and random trials was significantly above chance during the entire experiment regardless of the chance estimate employed (based on a combination of paired and single-mean t-tests comparing accuracy in the different eight-session blocks of each cycle against these estimates of chance; an alpha of $p < .05$ was used to judge the significance of all statistical tests in this paper).

Given the mechanisms of GAs, accuracy on *selected* trials was expectedly high across each cycle. With the reintroduction of the more difficult randomized trials, it briefly dropped at the beginning of the resets for cycles 2 and 3, but quickly recovered. Accuracy on *random*

trials was significantly lower than for *selected* trials across the three cycles (paired t-tests across blocks, cycle 1- $t(9)=7.4$; cycle 2- $t(13)=18.4$; cycle 3- $t(13)=14.1$). Nonetheless, random trial accuracy significantly improved over the three cycles, acquisition/cycle 1=36%, cycle 2=47%, cycle 3=55%, which is indicative of the bird's increasing capacity to identify the target independent of the visual parameters. The greater variability in accuracy on random trials was caused by their smaller N and randomized generation, leading to unpredictable mixtures of "easy" and "hard" trials in any session.

Effects of Selection on Gene Composition—The comparatively high accuracy on selected trials in comparison to the random trials reflects the result of the pigeon's selection behavior, as the varied descendants of the accurate trials were differentially retained. The impact of this differential accuracy can be seen in the resulting distribution of the different genes controlling the composition of the selected displays. Shown in Figure 4 is a summary of the movement of the genes over the three cycles of the experiment. The different panels show the mean values of each of the six genes as computed from the selected trials for each cycle. Included in each panel is the expected value for each gene in the absence of any selection by bird (dashed lines, distractor count was derived by simulation; details below).

These data were subjected to a stepwise binary logistic regression analysis using a likelihood ratio-based forward method to identify the relative weighting of the different genes in determining accuracy. Standardized values of each of the six genes were regressed onto accuracy for all trials conducted across the three cycles ($n=5286$). This regression analysis identified that five of the six genes had significant contributions to the final logistic model for accuracy. The most critical factors included first in the predictive model were the genes for distractor count ($\text{Exp}(\beta) = .5$; for ease of comparison, this converts to a common odds ratio of 2.0), and element size ($\text{Exp}(\beta) = 1.58$). The greater weights reflect their consistent directional selection by the pigeon over all three cycles and reflect the bird's strong selection of trials composed of fewer distractors and larger elements. Secondary to these factors were the genes for distractor brightness ($\text{Exp}(\beta) = 1.43$) and shape ($\text{Exp}(\beta) = .7$; odds=1.42), which next entered the model, as the bird was somewhat more accurate when more of the distractors were farther away from the target's gray value (i.e. whiter or blacker) and the elements were square. Their inclusion in the model, however, explained less of the variation than the first two factors. The least weighted factor was display spacing ($\text{Exp}(\beta) = .77$; odds=1.29). As can be seen in Figure 4, one of the reasons the latter genes had less impact was that they had smaller and more inconsistent directional trends across the three cycles. Target brightness was excluded as a factor in the final model as its variation did not significantly improve the fit.

The top panels of Figure 4 show the mean values of the three genes with the greatest impact: distractor count, element size, and distractor brightness. All three of these genes showed consistent directional trends in their selection across cycles, reflecting their significance to the bird's accurate target identification behavior. The genes controlling distractor count and element size both exhibited large and rapid changes within just a few sessions of the reset at the beginning of each cycle, as the pigeon consistently selected displays composed of fewer distractor items and larger element sizes early on, or similarly did not select trials with many very small shapes. The composite distractor brightness gene shown in the third panel also

moved consistently in the same direction in each cycle, but more slowly compared to the first two genes.

The distractor number gene moved quickly because the pigeon consistently performed better with displays having increasingly fewer distractors. It is worth noting that the mean value for distractor count could move, unlike for any other gene, for reasons independent of the bird's selection. Even if just guessing, displays with one distractor have a greater chance of being selected than displays with ten distractors. To model this type of gene movement in this case, we used Monte Carlo simulations (see last section) of the GA procedure that randomly selected among the available elements on each trial to obtain the chance expectations for this gene that are shown in Figure 4.

Distractor Number—Shown in Figure 5 is a summary of behavior as a function of distractor number and its impact of the distractor number gene from the combined last 20 sessions of each cycle (60 total sessions). The bars in the left panel show the bird's mean accuracy as a function of this gene's phenotypic expression, as measured by the number of distractors in the displays (the symbols depict results from the different cycles). Accuracy significantly declined with increasing numbers of distractors in the display, $F(9,18)=11.4$, (repeated measures ANOVA of accuracy using cycle as the repeated factor). Regardless of the number of distractors, however, accuracy was significantly greater than expected by chance (the "C" symbol in each bar) based on single-mean t-tests using accuracy from the last 20 sessions of each cycle, all $t_s(19)>2.43$. A linear regression on mean accuracy averaged across all three cycles revealed a significant linear effect of distractor number, $r(9)=.93$; adjusted $r^2=.87$. The best fitting line from this regression is included in Figure 5 (solid black line).

The greater accuracy with fewer distractors caused the composition of the selected trials to shift over the three cycles. This effect can be seen in the frequency distribution in the right panel of Figure 5. As computed from the last 20 sessions of each cycle, it shows the population distribution of this gene's value for selected trials as a function of distractor number for each cycle. The filled gray background shows the expected distribution given no selection, as determined empirically from distractor counts of all random trials created during the experiment ($n=1765$). The final distributions for each cycle were significantly different from the distribution prior to selection, $\chi^2_s(9)>1863$, and heavily skewed towards displays with only few distractors. On average, by the end of each cycle, 83% of the selected trials for a typical session had three or fewer distractors, with the majority of trials having only one or two (68%). Hence, during each cycle, the bird's selection behavior resulted in the rapid evolution of displays with increasingly fewer distractors.

Element Size—Element size was the other gene exhibiting strong differential selection over testing. Figure 6 shows a summary of the impact of element size, again derived from averaging results from the last 20 sessions of each cycle. The left panel shows the bird's accuracy as a function of this gene's phenotypic expression (measured in pixels of diameter/width). Accuracy significantly increased with larger elements, $F(16,32)=18.4$, (repeated measures ANOVA of accuracy using cycle as the repeated factor; several bins from each end of the size dimension were excluded from this analysis because of their small Ns).

Displays with elements of 200 pixels or larger in size supported similar levels of high accuracy. Accuracy declined between element sizes of 100 to 200 pixels. Elements smaller than 70 pixels, supported the poorest performance. The latter result was largely due to the pigeon failing to respond to displays having elements 2 cm or less in size (72% of the non-responses on random trials).

These differences in the pigeon's accuracy caused the selection and retention of displays with larger elements. The right panel of Figure 6 shows the mean percentage of selected displays as a function of the expressed element size (grouped in 20-pixel bins). Because the display packing algorithm used a combination of element size, distractor count, and spacing gene values to determine element size, the expected distribution of size values, given no selection, was not directly calculable. Consequently, this distribution was again derived empirically from the random trials created during the experiment (the filled gray in Figure 6). Across the three cycles, the pigeon repeatedly selected displays with larger elements as shown by the significant rightward shift of the final size distributions for each cycle, $\chi^2_{s(24)} > 143$. Across the three cycles the modal value of element size over the distribution of selected displays converged towards comparable values (acquisition=190 pixels; cycle 2=190; cycle 3=170) as did the median values (acquisition=254 pixels; cycle 2=191; cycle 3=250 as computed from finer grain 5-pixel bins). Thus, 5-cm elements were consistently converged upon as the most frequently occurring element size in the displays.

Element Brightness—The composite average brightness of the distractors was the third parameter with consistent directional effects each cycle. Overall, the mean brightness of the distractors in the displays gradually shifted away from the target's intermediate range. Figure 7 shows a summary of this selection behavior and its impact on distractor brightness over the last 20 sessions of each cycle. The left panel shows the bird's accuracy as a function of this gene's expression, which is measured by the average difference from the closest target boundary to all distractors in a display. As the average brightness of the distractors moved away from the target's range, accuracy significantly and linearly increased, $F(7,25)=4.3$ (repeated measures ANOVA of accuracy using cycle as the repeated factor with linear polynomial contrasts). A linear regression on accuracy as a function of distractor brightness averaged over the three cycles confirmed this pattern, $r(9)=.98$; adjusted $R^2=.953$, and the best fitting line from this regression is included in the left panel of Figure 7 (solid black line).

We examined this effect further to see if the polarity of the difference contributed to performance (whether distractors were lighter or darker than the target). Using displays with smaller numbers of distractors to simplify the analysis, we could find no evidence that polarity affected the bird's behavior, as the subject performed equally well with generally lighter or darker distractors overall. We also examined whether having distractor homogeneity mattered (all darker or all lighter than the target vs. mixed displays). Using displays with two and three distractors where there were sufficient numbers of each type to compare, the percentage correct with homogeneous displays (acquisition/cycle 1=88%; cycle 2=98%; cycle 3=94%) was slightly, but inconsistently, better than with heterogeneous displays (acquisition=89%; cycle 2=96%; cycle 3=91%).

The right panel of Figure 7 shows the relative frequency of selected displays as a function of the distractor brightness gene for the last 20 sessions of each cycle. Again, differences in the pigeon's accuracy caused the selection and retention of displays with increasingly contrasting distractors. The expected distribution given no selection (filled gray area) was again derived empirically from all random trials. The distribution of selected values differed significantly from this expected distribution for each cycle, $\chi^2_{s(15)} > 274$.

As indicated by the logistic analysis, the three genes in the bottom panels of Figure 4 were less influential and their selection less consistent regarding their influence on performance. Compared to those in the top row, they moved little or inconsistently across the three cycles. The bottom left panel shows that average target brightness remained essentially constant over the cycles. The next two bottom panels show the changes in mean gene value for display spacing and element shape. The mean value of these genes shifted within cycles, but inconsistently across cycles.

Consider first the gene for target brightness. The logistic analysis indicated that this gene made no significant contribution to differences in accuracy over the experiment. The mean brightness of target varied little within its bounded range, although it may have gotten slightly darker with time. Given that the brightness of the distractors fell equally to either side of the target, both within and between trials, holding to the middle of this range is not unexpected.

Element spacing—While the logistic analysis identified the display spacing gene as a significant component towards accounting for differences in accuracy across the experiment, it was the least weighted factor in the model, indicating its smaller impact on explaining variation in search accuracy. The changes in this gene across the cycles are shown in the middle bottom panel of Figure 4. The most noticeable feature is that it did not move in the same way each cycle. Given perhaps a selection bias towards sparser displays early in the first two cycles, the mean value of this gene did not markedly vary from its unselected value over the middle portions of each cycle, suggesting any effects of this factor emerged only after other factors more influential to responding had settled to their best values. Over the latter sessions of the three cycles, the changes in this gene diverged more clearly from its unselected value, but in different directions. In the first cycle, it remained close to the expected value for non-selection. By the end of second cycle, a significantly greater proportion of denser displays had been selected. By the end of the third cycle, a significantly greater number of sparser displays had been selected.

Because of its computational involvement with element size and distractor number in the packing of the display, the examination of the phenotypic variation of element spacing is perhaps more valuable than examining its mean genetic value. We computed the average pixel distance from the center of every element to every other element to measure the phenotypic density of the display. This measure revealed that during cycles 1 and 2, the average distance between the elements grew smaller over each cycle. Compared to the average pixel distance among elements for randomized new displays (352 pixels), this value was marginally reduced for selected trials over the last 20 sessions of the first two cycles, cycle 1= 327 pixels; cycle 2=302 pixels. This was not the case for the third cycle, where

average spacing grew slightly larger (371 pixels). When considered in the context of the gene's movement, these patterns suggest that during the first two cycles the bird was selecting displays that consisted of larger elements that were packed slightly more closely together. By the third cycle, however, while the bird was still selecting larger elements, the necessity of packing them more closely together appeared to be diminishing, causing the distribution of this gene population to move towards sparser displays. One possible reason for this is that with experience, the pigeon was more tolerant of a wider set of display conditions.

Element shape—Finally, the right bottom panel of Figure 4 shows the movement in the gene controlling element shape. This gene also moved in contrasting directions during the three cycles. During acquisition, an increasing proportion of the displays involved circular elements, while for the two subsequent cycles, an increasing proportion of the displays involved square elements. As derived from last 20 sessions of each cycle, there were differences in accuracy depending on the elements' shape, acquisition/cycle 1: circle=85%, square=79%; cycle 2: circle=71%, square=91%; cycle 3: circle=79%, square=86%. The better performance over the last two cycles of the experiment with square elements may indicate that this is the better of the two shapes for the present discrimination, perhaps because it permitted a better visual comparison of the target/distractor contrasts due to its geometrically straight edges and greater area. The greater control by the circle early in the first cycle may reflect a type of “founder effect” where it dominated the better square because of chance frequency, or greater success early in training.

Performance Trajectory & Monte Carlo Simulation—Lastly, Figure 8 shows a plot of the trajectory of the pigeon's solution for each cycle through a stimulus space defined by the two genes with the most influence on search performance. It shows that the pigeon consistently evolved the discrimination to the same location in this space, a location characterized by displays having few distractors and large elements. To capture this area, Figure 9 contains five representative examples of the selected displays from the last five sessions of each cycle. The displays shown are those that minimized the total difference from the median of all six gene values for the trials tested within a specific session and hence give a good representation of the underlying population of values at that time. Collectively, they illustrate the pigeon's general selection of displays containing three or fewer distractors that were generally lighter and/or darker than the target, and approximately 5 cm in size.

Included for comparison in Figure 8 are the results of a Monte Carlo simulation that modeled the evolution of the discrimination given random guessing (i.e., no selection done by the pigeon). We conducted 3500 runs of 108 sessions using randomized initial conditions for each gene. The figure plots the regions where 95% of the final location of these unselected trajectories resided. The brightness of an area indicates the relative density of these terminal locations in the solution space (black=highest density). As noted before, the movement along distractor count reflects the inherent selection bias for fewer distractors even when just guessing, but otherwise the trajectories of the other genes all engaged in random walks of the stimulus space. The bird's solution in each cycle consistently resided

outside this region, providing further evidence of the bird's active and consistent selection among the different displays tested during the experiment.

Discussion

Experiment 1 revealed that a pigeon can successfully learn and perform with a constantly evolving dynamic and variable discrimination governed by a genetic algorithm. Across three sets of different randomized initial conditions, the pigeon's solution to the target search and identification problem repeatedly caused the trials to converge upon similar display conditions for the most critical parameters. Overall, the pigeon's selected conditions consisted of displays characterized by a gray square target with one to three surrounding blacker or whiter distractors, most frequently being about 5-cm in size. This suggests that these display conditions represent the area within the space of all possible solution that is used by the bird to solve the task. This region can be thought of as a "psychological attractor," representing the best conditions for the operation of the cognitive processes deployed in solving and performing the discrimination. Regardless of the task's initial conditions, the combined dynamical system of pigeon and GA evolved towards this same area each time. This is reflected in the similar shifts across cycles of the means and population distributions of the genes controlling display organization, most importantly distractor count, element size, and distractor brightness. Their importance to performance is indicated by their consistent directional selection, corresponding shifts in underlying population values, and the degree to which they influenced accuracy as indicated by the regression analysis.

The pigeon's consistent selection resulting in the reduction of distractor number was foreseeable based on prior studies. Studies of pigeon visual search have consistently found that choice accuracy and RT benefit from decreases in the number of distractors (D. S. Blough, 1977, 1979, 1989). That pigeons are better discriminators as distractors become increasingly unlike the target would also be consistent with previous results (D. S. Blough, 1988, 1990). The rapid emergence of larger sizes and their considerable impact here as determined by the GA was not as easily anticipated as the effect of number of distractors, as there is much less systematic information about the effects of size on discrimination. That pigeons might do better with larger elements would not be entirely unexpected (Cook, Goto, & Brooks, 2005; Peissig, Kirkpatrick, Young, Wasserman, & Biederman, 2006), although given their strong local bias or the size of their of natural food items, one could have argued that smaller sizes might have been selected (Cavoto & Cook, 2001; Lea et al., 2006). A factor to note is that bird was quite good with displays larger than the ones that eventually dominated. Because of the combination of the spatial packing algorithm and the selected gene values, such large sizes occurred more infrequently, however, which may have underrepresented their potential contribution in the bird's final solution.

The fact that the three remaining genes/parameters changed little (target brightness) or inconsistently across cycles (shape and spacing) clearly suggests their role as secondary factors. While these genes had measurable impacts on performance (denser being slightly better than sparse displays; squares being slightly better than circles), the revealing new

product of the GA is the indication that these factors made less of a contribution to the pigeon's visual processing than those outlined above.

It is significant that the final evolved form of the pigeon's solution corresponds in many respects to what we would have predicted based on prior visual cognition research with these animals. It is essential to remember that these correspondences needn't have occurred at all. The open-ended nature of the GA procedure permitted an infinite number of singular or multiple regions to potentially emerge as terminal "best" conditions to each evolution. Thus, these correspondences validate that the use of the GA procedure itself did not alter the pigeon's basic approach to this kind of visual search task, at least as compared to simpler factorial designs. Therefore we can be reasonably confident that when a GA procedure is used to examine the effects of variables that are not as well researched, the results will reflect the cognitive processes underlying the discrimination, rather than being a product of the GA procedure itself or the dynamic quality of the composition of the displays.

The higher experimental dimensionality possible with the current GA approach revealed several new facets of our parameters that would have taken longer to discern with a series of smaller factorial designs. It identified the strength of a large set of different display parameters relative to each other quickly, and it suggested that element size is perhaps more critical to pigeons than previously suspected. The results indicate that there is substantial flexibility in the pigeon's cognitive processes for dealing with a wider and more diverse set of displays than normally tested in factorial setups.

As judged in isolation from other factors, for example, the pigeon was slightly better when only one distractor was on the display, and was slightly better when the elements were larger. One could easily see those as outcomes of factorial designs varying just distractor number or just element size, and concluding that one distractor of the largest size would be best and would come to dominate selection. Such displays certainly occurred, but they never came to dominate as compared to many to the other forms of the displays that also survived, although nothing in the procedure prevented that from occurring. There was not so great an advantage to having only a single distractor or very large elements that instances of a few or intermediated-sized distractors also could not thrive. Thus, there was not sufficient selection pressure to drive the displays to such extreme values, and the pigeon's solution to the task was far more flexible and useful in processing display variability.

One possible reason for these effects is that several factors were interacting with each other in controlling the bird's behavior. In the present case, the pigeon was perhaps best suited by displays with a small number of distractors of intermediate size because it allowed for a better computation of the brightness differences between the target and distractors, possibly by permitting multiple comparisons to be compared. By themselves, larger elements and one distractor may indeed be best in simpler controlled situations, but in a context with combinatorial complexity and variability similar to the real world, it turns out that less extreme values might be the best or equally good solutions. Such complex interactions will be easier to identify and locate in the current GA approach than in isolated factorial designs. While it is true that with all other things being equal, for example, the pigeon was better when the elements were square, circular elements could be discriminated just fine (and

therefore retained) when occurring frequently in the context of displays composed of the more critical factors (such as fewer distractors and intermediate sized elements).

One factor to consider was whether we tested the bird long enough, which might have changed these apparent bounds, even with small differences in accuracy. Working out the exact and proper ratio between length of testing, precision, and optimality of the measured solution is an area that still needs to be evaluated for our GA approaches, as the appropriate procedure will depend on the variability and shape of the performance space. What can be said here is that over the several months that we tested here, all of the genes appeared to have slowed down in their approach to their terminal values. This indicates that we had likely reached a reasonably stable equilibrium point for this discrimination and at worst, were close to the final attractor of its solution.

Another experience-related factor to examine concerns the possible effects of repeatedly resetting the initial conditions of the discrimination during the experiment. While such retesting is a classic psychophysical tactic, the bird's succession of experiences with varying degrees of success from harder to easier may have long-term sequential effects. Such effects may be revealed by looking at the consistency of the directional movement of genes across repeated cycles. The effects of distractor count, size, and brightness were easily identified by their consistent effects. While a firm conclusion is not possible at the moment, the bird's greater selection in the second and third cycles for square elements might represent such an experience-related change related to long-term exposure to the task relative to the bird's more "naïve" approach at the beginning of experiment.

Experiment 2

Experiment 2 investigated the control exerted by the presumed psychological attractor identified in Experiment 1. For the purposes of this experiment, we envisioned the same stimulus space as identified in Experiment 1, using the factors of element size and distractor count to track the bird's solution to the task. Specifically, we tested the strength and repeatability of this attractor by selecting initial conditions that were maximally distant from its assumed location in the space. The initial conditions in Experiment 1 were functionally in the middle of the stimulus space because of the randomization procedure used for generating random trials across the three cycles. For Experiment 2, we purposely seeded the initial condition for random trials to be in the corner opposite and farthest away from the location of the hypothetical attractor isolated in Experiment 1. We modified the GA to produce new random displays characterized by large numbers of sparsely packed, small distractors. If the attractor identified in Experiment 1 exists and is powerful enough, the composition of the selected displays should move away from this seeded location and evolve towards the types of displays that dominated the first experiment. If other viable solutions or potential intervening local maxima closer to this starting point existed, however, we should see this trajectory move towards a different and potentially stable point in the stimulus space. Under the right conditions, for instance, it has been found in a texture-like discrimination that pigeons can benefit from having many distractors in the display rather than fewer (Zentall et al., 1980). Consequently, it is important to determine if there are multiple stable solutions to the task's stimulus space.

We also thought it was essential to examine the impact of extended experience on selection behavior. Could we, for instance, persuade the animal to remain in a particular location of the space by overexposing it to constituent trials? The hypothesis was that by repeatedly testing trials from the same area of the stimulus space, the greater familiarity/practice/associations/attention resulting from this experience might cause the bird to remain in that area. To accomplish this, the modified GA composed random trials from a limited range of values located near the initial conditions used to seed the experiment. This restricted values caused random trials to be consistently more alike themselves and quite distinct from the solution that evolved in the first experiment. If familiarity-based selection could be induced, the average value of the selected trials should remain stable in the same location over many sessions.

Experiment 2 had two phases. In Phase 1, the bird was tested with a set of experimenter-determined initial conditions and a modified GA that restricted random trials to a limited range of gene values around that seed. As will be described, no effect of familiarity-based selection could be established by this restriction. As a result, in Phase 2, we modified the random trial generation procedure so that random trials were composed from values located around the moving average of the pigeon's selected trials. This modification transformed the procedure into a type of hill-climbing algorithm with which we tracked the trajectory of the bird's selection behavior. This trajectory allowed us to again to determine the location of any psychological attractor governing the bird's selection behavior and its similarity to that identified previously.

Methods

Animals & Apparatus—Same subject and apparatus were used as in Experiment 1.

Procedure—The search task remained exactly the same for the pigeon; locate, identify, and peck at the gray target element amidst the variable lighter and darker distractors. The experiment consisted of two phases. In phase 1, the GA was modified compared to Experiment 1 for making random trials. Instead of the completely randomized technique used there, this experimental phase employed a set of experimenter-determined initial conditions. *Random* trials were now generated from within a distribution of gene values that restricted the distractor count, element size, and display spacing genes to a small range of values. The distractor number and spacing genes were restricted to the highest 10% of their range, and the size gene was restricted to the 10% – 20% values of its range. These restrictions resulted in seed values that consistently produced displays characterized by large numbers of sparsely packed distractors about 1-cm in size. Such displays were located far from the evolved solution in the stimulus space observed in the first experiment, while creating elements large enough for the pigeon to accurately target. All other genes controlling display appearance otherwise varied at random in the same manner as in Experiment 1.

Phase 1 lasted 108 sessions. Sessions consisted of 50 trials, with four sessions successively conducted each day. For each session, 45 selected trials and 5 random trials were tested. The selected trials were generated from the PPD of the bird's prior selection behavior, as

reproduced with a randomized $\pm 5\%$ variation in their genetic values. After 80 sessions, the variability involved with reproducing selected trials was changed to $\pm 10\%$ and five additional random trials replaced five selected trials in each 50-trial session.

Phase 2 simply removed the limitation that random trials come from with the restricted seed conditions enforced during Phase 1. Instead, random trials for each session were derived from the mean gene values ($\pm 5\%$) of the trials in the PPD for the three restricted genes, while the unrestricted genes reproduced over the same range as used in Phase 1. Phase 2 lasted 164 sessions.

Results

Choice Accuracy—Despite the procedure for restricting the proportion of new random trials to a distant region of the task’s stimulus space, the pigeon’s evolution of the discrimination again converged on values similar to those of Experiment 1. Figure 10 shows the pigeon’s mean accuracy across the experiment for both random and selected trials in 8-session blocks. Because the pigeon was already good at the task, the increase in accuracy on the selected trials reflects the bird’s selection of progressively “better” displays. Accuracy on random trials during the two different phases reflects the procedures operating at those times. During Phase 1, the bird had considerable difficulty with the restricted high count/small size/sparse placed random trials. Accuracy with these trials was low and improved little with experience. This is not surprising given that the restricted structure of these displays was antithetical to the bird’s selections in Experiment 1. This changed with the advent of the mean-based hill-climbing algorithm during Phase 2. At this point, because the random trials were directly influenced by the selected trials in the PPD, the accuracy on the random trials began to strongly track performance on the selected trials, since their sources were functionally equivalent.

Effects of Selection on Gene Composition—The data from Experiment 2 were subjected to the same binary logistic regression analysis used in Experiment 1. Standardized values of each of the six genes were regressed onto selection accuracy for all trials from both phases ($n=13,604$). This analysis again identified that five of the six genes made significant contributions to the final logistic model for accuracy. The two most highly weighted factors were again the genes for distractor count ($\text{Exp}(\beta) = .62$; odds=1.61) and element size ($\text{Exp}(\beta) = 1.46$) as the bird again strongly selected displays with larger elements and fewer distractors (detailed below). Next were the genes for distractor brightness ($\text{Exp}(\beta) = 1.21$) and shape ($\text{Exp}(\beta) = .81$; odds=1.23), as the bird tended to be more accurate when more of the distractors were farther away from the target’s gray value and the display’s elements were square. The contribution of display spacing ($\text{Exp}(\beta) = .83$; odds=1.20) was slightly greater than in Experiment 1, with the bird showing a tendency to select denser displays this time around. Target brightness was again excluded in the final model, as its addition and variation did not significantly improve the fit.

Figure 11 shows a summary of the change in the means of the different genes over sessions. Each panel shows the mean gene values of the selected trials for Phase 1 (white circles) and Phase 2 (black circles). Also included for comparison are the mean values for the new

random trials for the three initially restricted genes (distractor count, element size, and spacing) used to create these displays during Phase 1 (white squares), and also the result of Monte Carlo simulations for those genes. Like in Experiment 1, both the distractor count and element size genes changed over time as the bird selected displays with few intermediate-sized distractors. Further, compared to the mixed results observed in Experiment 1, the spacing gene changed in an apparently more directional manner from the selection of denser displays over sessions. All three genes changed significantly between the sessions at the beginning and end of the experiment, $t_s(22) > 59.1$. This figure shows that during the restricted conditions of Phase 1 (i.e., the constrained random trials in phase 1) designed to potentially induce familiarity-based experience effects by overexposing a limited range of displays, had little impact. All three gene values significantly diverged from the “fixed” location of the random trials, and did so from the beginning of the experiment. This occurred despite the repeated sampling of restricted random trials and the small reproductive variation allowed among the selected trials, designed to possibly “encourage” the bird to stay put. This shows that in determining the bird’s accuracy display, familiarity was less critical than display discriminability.

While the restricted genes moved in the same directions over both phases of the experiment, the three unrestricted genes related to element shape and brightness moved in seemingly different directions across the phases. These contrasting movements suggest that the bird perhaps compensated for the imposed organization of the restricted trials in Phase 1 by relying on different features of the displays than in Phase 2 or Experiment 1. During Phase 1, the movement in the mean values of these three unrestricted genes indicates the pigeon was better with displays that had darker targets, distractors that were farther away from the target bound, and square elements. These features can be seen by the significant deviation of these three genes from their random expectation over the last 20 sessions of Phase 1 for the selected trials, distractor brightness $t(19) > 26.9$; target brightness $t(19) > 7.8$; shape $t(19) > 47.4$. Thus, in a situation where the bird was forced to discriminate displays with large numbers of distractors and small elements, it was most successful with (and thus selected for) brightness differences with darker targets and whiter distractors as presented in square elements. We think this suggests that the bird was strongly emphasizing the target/distractor brightness differences in the displays.

In Phase 2, with the institution of the hill-climbing algorithm and the removal of the restrictions on the random trials, the selection pressure from this “brightness-based solution” apparently disappeared. This resulted in the three unrestricted genes moving in new directions inconsistent with such a solution. This change can be seen in more circular elements reentering the mix of selected trials, the return of target brightness back to the middle of its range, and a reduced selection pressure for distractor brightness. Over the last 20 sessions of Phase 2, two of these genes, target brightness and shape, moved to locations that were not different from their chance expectations, single mean $t_s(19) < 1$. The distractor brightness gene was still significantly different from its chance expectation given no selection, $t(19) = 3.1$, but was significantly reduced compared to these values recorded at the end Phase 1, $t(38) = 13.1$ (two-sample t-test compared with last 20 sessions of Phase 1). These movements back towards a less selected state suggest that selectively enhancing the

contrast between the target and the distractor was not a relatively important ingredient in the bird's solution and selection after the restricted range for the random trials was removed with Phase 2.

Figure 12 provides a summary of mean accuracy and the expressed phenotypic populations of the genes identified by the logistic analysis. In each case, as evidenced by differential selection, the bird's accuracy varied as a function of distractor count, element size, distractor brightness, and element spacing (see the left hand panels of Figure 12). The bird was also slightly better overall with square elements (72%) than circular ones (68%). The results for distractor brightness are broken down by distractor count to clarify their presentation. All of these effects of display properties on accuracy were comparable to those observed at the end of Experiment 1.

The right hand panels of Figure 12 show the effect of this differential accuracy on the selection and distribution of the genes compared between the first and last 32 sessions of the experiment. Again, for each gene, there are significant population shifts in the distribution of their values as the bird once again evolved the discrimination. χ^2 -tests confirmed these significant population shifts between the beginning and end of the experiment. The shape gene (not shown) also showed changes, with 68% of the early sessions being composed of squares moving toward a more 50%–50% split by the last 32 sessions (square=48%). The simple summary is that by the end of the experiment, the pigeon had once again strongly selected displays predominately characterized a combination of one to three discriminable circular or square distractors per display, separated by an average of about 2-cm from each other from elements of approximately 4–5 cm in size.

Figure 13 displays the trajectory of the pigeons' selection behavior in the two-dimensional stimulus space of distractor count and element size used to summarize Experiment 1. The white (Phase 1) and black (Phase 2) symbols show the pigeon's path during Experiment 2. The examples of the displays included in this figure are representative of the types of displays near their locations. Also included are the results of 3500 Monte Carlo simulations of the evolution of the discrimination given no differential selection. The figure plots the regions where 95% of the terminal location of these unselected random walks eventually resided for each phase, which were the upper left region for Phase 1 and the bottom region for Phase 2 (the initial condition for the simulations of Phase 2 used the state of the pigeon's value at that point in the experiment). The brightness of the areas indicates the relative density of these terminal locations at that point in the stimulus space (black=highest density).

The bird's trajectory through the stimulus space was not random. Some indication of the relative importance of the different variables can be seen in the figure. Given the trajectory across sessions, the pigeon's selection behavior seemed to first increase the size of the elements, as reflected in the stronger initial rightward movement of the trajectory. This may partly reflect the apparently early attempt to implement a "brightness solution" to the task, or just the importance of this basic feature. Once the elements become large enough, the trajectory takes a downward turn as the bird began to reduce the number of distractors in the display. Overall, even with the remote location of the initial conditions from Experiment 1's

attractor and the forced and repeated sampling from that location over Phase 1, the pigeon's selection behavior over sessions resulted in a steady movement towards the same location in the stimulus space as in Experiment 1.

Discussion

Experiment 2's results confirmed the existence of a robust psychological attractor in this pigeon's solution space of this discrimination. Despite initial conditions picked to be maximally distant from the hypothesized attractor, and a sampling and reproduction procedure that attempted to bias the pigeon away from its solution, the trajectory of the pigeon's evolving discrimination showed a steady return to display configurations equivalent to those observed during the three cycles of Experiment 1. This psychological attractor, from the converging trajectories across the independent resets, was characterized by displays with a gray target element, surrounded by one to three distractors, separated by an average of about 2 cm from each other, and typically 4–5 cm in size. The parallel weighting and ordering of the different genes in each experiment's best fitting logistic model indicates that the bird was consistently selecting the same factors and values each time. This did not have to happen. The task's solution was open-ended from a procedural perspective. If the bird were equally good with any display, there would have been only random movement in the trajectory through the stimulus space in each experiment. If the bird had been better with smaller elements or when having many distractors to compare, for example, those displays would have emerged and dominated the selected trials. They did not. Thus, from the pigeon's perspective, the task was not open-ended, but one characterized by a "sweet spot" in the discrimination as identified by the pigeon's repeated convergence towards the same parameter values.

While the remote distance of the initial conditions in Experiment 2 did not change the solution that emerged, restricting random trials to that location did have an unexpected impact in the first phase. The original purpose of the restriction was to see if overexposure to a small range of displays might encourage the bird to stay in that region of the stimulus space. That didn't appear to happen, as the bird slowly moved away from this origin. In doing so, however, the resulting displays seemed consistent with a brightness-based solution to the task. This "brightness" solution can be seen in the greater directional extent the evolution followed on target/distractor contrasts during Phase 1. Given the greater predominance of displays with many small elements, this kind of solution is reasonable. With the subsequent removal of the count and size restrictions on the structure of the random trials, this solution seemed to disappear, as the same visual aspects of the discrimination that controlled performance in Experiment 1 strongly emerged after that point. The disappearance of this strategy is confounded, however, with the introduction of the "hill-climbing" methodology in Phase 2. This eliminated the large difference between random trials and selected trials, in which case the powerful contributions of element size and distractor count may have overwhelmed the effect of any previous or ongoing target/distractor contrast effect. Interesting questions to explore will be how the parameter distributions and visual distinctiveness of different features affects selection and the development of different solutions.

Other studies using maintained discriminations have reported similar shifts in the nature of processing during the course of testing (cf. Krechevsky, 1932). Blough (1993), for example, was able to isolate “drifts” of attention to different features as a function of correlations among the RTs during a visual search task testing pigeons. Other experiments have detected shifts in attention to different dimensions depending upon reinforcement outcomes (D.S. Blough, 1969; Brown, Cook, Lamb, & Riley, 1984; Cook, Riley, & Brown, 1992; Lamb & Riley, 1981; Leith & Maki, 1975). This ability to detect and track the presence and absence of processing, selection, and attention to different aspects of a complex discrimination is one of the potentially valuable attributes of using GA procedures to examine the dynamics of cognition.

General Discussion

These experiments represent a novel use of open-ended adaptive techniques for examining cognitive processing in non-human animals. Using continuous and simultaneous variation in multiple independent parameters and their differential selection by the animal, the approach interactively explored a pigeon’s processing of a multidimensional search task. This discrimination evolved as a direct function of the animal’s differential selection and from different initial conditions repeatedly converging towards the same psychological attractor in the task’s stimulus space. To our knowledge, this is the first time this type of psychological attractor has been identified in a non-human species. The dynamic shaping of the task exhibited by the animal in these experiments is a highly useful means for understanding the underlying cognitive mechanisms involved.

In dynamical systems, attractors are considered to be a stable set of variables toward which such systems converge upon over time. The pigeon’s recurring trajectories towards the same region of the stimulus space (Figures 9 & 13) used here suggest that such a term is applicable to describing the bird’s behavior and reflects its cognitive processing of the task. Neither the GA nor the procedural organization of experiments determined this outcome, as the pigeon was free to drive the task in whatever direction its choice behavior selected and potentially “preferred”. There were a large number of different and viable solutions available that the pigeon could have adopted to increase its rate of reinforcement. It did not do so.

Several factors were hierarchically involved in the pigeon’s consistent solution to the task. The three parameters that were most prominent for the bird were the number of distractors, the size of the elements, and to a lesser degree the contrast between the target and distractors. The spacing and shape of the elements, while influential, were secondary to these other parameters.

The factors of distractor number and element size were reliably and strongly selected, and principally controlled the bird’s trajectory through the task’s stimulus space. This attractor was characterized by elements of approximately 4 to 5 cm in size with a target surrounded by one to three distractors. In Experiment 2, it looks like element size was adjusted prior to reducing distractor count (i.e., the initial rightward movement in the trajectory). Regardless, these two variables primarily defined the psychological attractor that represented the bird’s

solution to the discrimination over the four independent evolutions. This repeated “unsupervised” solution reflects the pigeon’s fundamental approach to the discrimination and the importance of these variables in this solution.

An effect of target/distractor contrast was also present. When the distractors were generally whiter and blacker, the bird was better at identifying the gray targets. Our intentional dampening of the impact of this variable by forcing mixtures of distractors of different values may cause us to underestimate its relative importance in the solution hierarchy. If the bird had been allowed to select the directional brightness of the entire set of distractors in a trial, the bird may have tended towards consistently high contrast displays.

The secondary nature of element spacing and shape is reflected in their consistently being less weighted factors in the logistic analyses and the inconsistent nature of their gene movement in different cycles. While it is clear that with all other things being equal, the pigeon was slightly better when the elements were square and generally preferred the elements to be spaced closer together, these factors did not dominate the selection of the displays. Spaced circular elements, for example, could be discriminated well, especially when occurring in a display with a few large distractors. Correspondingly, densely packed, square elements could not save a display having many small distractors. Thus, the underlying population of these secondary genes could not easily eliminate their “weaker” values precisely because of the greater importance and compensating influences of displays with better values from the more dominant controlling genes. Although we did not explicitly include it in our algorithm’s genetic structure (and thus it could not be selected for), another possibly influential secondary variable was display homogeneity, as the bird showed a small advantage when homogenous sets of distractors occurred by chance. The capacity to prioritize the importance of such different display factors is one of the advantages of the adaptive open-ended method, especially in evaluating complex discriminations with higher dimensionalities. The unbiased, interactive, open-ended, and subject-driven approach of our procedure has the significant advantage of allowing us to precisely pinpoint such differences, and identify the most influential parameters and their respective ranges. A rich avenue to explore with such procedures would be how restrictions on different variables, such as their mean value, relative stability, location, and variability, impact the expressed solution(s) of any task.

Although we are using a quite novel approach to examining discriminations, our experiments obviously employed another classic psychological approach – the detailed psychophysical examination of a single subject. Because we were more interested in exploring and validating the adaptive approach, the psychophysical strategy of extensively studying a few subjects in great detail seemed a conservative place to start. Our intentional selection of a problem that was already reasonably well-explored using factorial methods gives us considerable faith that our procedure indeed captured aspects of behaviors that have been established to be more important in such settings (e.g., the effect of distractor number). It will be essential, however, to replicate these findings with additional animals and species to examine how much individual and group variability there is in the task, whether the attractor identified here is characteristic of pigeons in general or reflects this bird’s unique solution, and how such solutions vary across different animals and species.

These issues raise one of the possible benefits of using open-ended GA techniques. It may allow for the easy isolation and study of the similarity and differences among individuals, tasks, and species. Individual variability is an acknowledged part of any experiment. For the most part, these are regulated to the error terms in our factorial designs. Because of their subject-driven nature, adaptive techniques provide a way to identify rapidly the manner and degree to which individuals vary. Within the confines of factorial experiments, subjects can only vary by being simply better or worse along some aspect of the task (e.g., fast versus slow learners). With the richer dimensionalities of the adaptive procedure, this variation can be profiled by the nature and weighting of parameters selected. With such information, we can easily identify the common, characteristic, and idiographic aspects of each subject's psychological profile. By being able to reset or relocate the initial conditions, it would also be possible to judge the relative stability of these for each subject. Expanding on this same logic, adaptive techniques would be equally valuable for profiling and comparing different species. Again, being able to identify the variable(s) most vital to each species and its relative weighting in the same problem would be useful in determining the relative similarity and differences among species. The open-ended nature of GA procedures has the potential for rapidly discovering and isolating such properties.

One continuing challenge in the study of comparative cognition is whether any task covers the right range of variables and values for the population under study. Various methodologies of accomplishing this have been proposed (e.g. control by equation; control by systematic variation, Bitterman, 1975). Adaptive techniques may function to advance this issue by identifying those variable(s) most critical to any species' success. As a result, the "best" conditions could then be used in subsequent testing, or the resulting profiles compared with greater confidence that the proper range of conditions for testing each species had been available for inclusion and selected by each species. Thus, adaptive experimentation offers a potentially transformative new tool for investigating complex discriminations across different species. By the same rationale, GAs have the potential of revealing how different clinical or developmental populations differ, as well.

Besides its application to examining cognitive processes, the GA procedure itself is worthy of investigation. As with any new procedure, a number of questions remain to be examined about the best methods for using GAs to measure cognition and behavior, and about its relations to traditional factorial experiments. It will be valuable, for instance, to examine the influence of other dependent variables. Besides the discrete measure of accuracy used here, examining analog measures such as peck rate or reaction time, either alone or in conjunction with accuracy, may reveal other aspects of behavior and the relative weighting of the different variables that might differentially influence them. Adding time pressure by selecting trials based on RT would emphasize those display parameters, for example, most connected to processing speed.

Another future topic to examine concerns the method of variation and selection used to control the algorithm. We used two methods of variation with different strengths. The first involved randomized trials that sampled the entire parameter space. This method proved faster in reaching the attractor, and it repeatedly converged on the same solution. The second method used a hill-climbing procedure in which random trials were generated from within a

limited range of the bird's current solution. The second method produced slower movement towards the attractor, but allowed for a most detailed mapping of its trajectory. Further, this more restricted method of variation did have the benefit of allowing us to confine the bird to a region of the solution space for an extended period of time. How the extent and type of stimulus variation influences acquisition and learning of discrimination is not well explored, but is likely a significant factor (e.g., Kendrick, Wright, & Cook, 1990; Wright, Cook, Rivera, & Sands, 1988). A related issue concerns the best ratio of selected trials to random trials and their relative "life span" in the PPD. Further, should the PPD reflect trials selected for each subject or represent the group selection of all subjects combined? Finally, it will be important to investigate how past decisions affect current performance (path dependence). The set of alternatives for any given circumstance is limited by past choices and how past selection and founder effects influence the resulting solutions of task will need to be understood.

The twin themes of variation and selection have a long history in biological and psychological thought (Darwin, 1859; Thorndike, 1898). Based on these fundamental principles, the adaptive experimental approach tested here offers a promising and exciting new tool for investigating the behavioral and cognitive structure of virtually any complex discrimination in human and non-human animals. While this approach clearly is within the tradition of adjustment methodologies (Treutwein, 1995), it differs in several significant respects. These new procedures allow us to test parameter spaces with much richer and more natural dimensionalities than is possible in traditional factorial designs. Recent simulations, for example, have suggested that such continuous adaptive sampling techniques are more efficient than factorial approaches for parametric estimation (Cavagnaro, Pitt, & Myung, 2011). By being open-ended and subject-driven the outcomes are conceptually similar to a production task for animals, as their selection behavior shapes the final form of the task. Thus, the current GA approach provides a rich and dynamic new way to explore animal behavior and its underlying cognitive processes. Its techniques capture the best traditions of simpler psychophysical techniques and combine it with the power of modern digital technology and evolutionary computing to reveal new aspects of cognition and behavior.

Acknowledgments

This research was supported by a grant from the National Eye Institute. The authors thank Gabe Rothman for his helpful comments on an earlier draft of the manuscript. Robert.Cook@tufts.edu. Home Page: www.pigeon.psy.tufts.edu

References

- Barber TX, Silver MJ. Fact, fiction, and the experimenter bias effect. *Psychological Bulletin*. 1968; 70:1.
- Basile BM, Hampton RR. Monkeys recall and reproduce simple shapes from memory. *Current Biology*. 2011; 21:774–778. [PubMed: 21530257]
- Bitterman M. The comparative analysis of learning. *Science*. 1975
- Blough DS. Attention shifts in a maintained discrimination. *Science*. 1969; 166:125–126. [PubMed: 5809588]
- Blough DS. Visual search in the pigeon: Hunt and peck method. *Science*. 1977; 196:1013–1014. [PubMed: 860129]

6. Blough DS. Effects of the number and form of stimuli on visual search in the pigeon. *Journal of Experimental Psychology: Animal Behavior Processes*. 1979; 5:211–223. [PubMed: 528887]
 7. Blough DS. Pigeon perception of letters of the alphabet. *Science*. 1982; 218:397–398. [PubMed: 7123242]
 8. Blough DS. Odd-item search by pigeons: Method, instrumentation, and uses. *Behavior Research Methods, Instruments, & Computers*. 1986; 18:413–419.
 9. Blough DS. Quantitative relations between visual search speed and target-distractor similarity. *Perception & Psychophysics*. 1988; 43:57–71. [PubMed: 3340500]
- Blough DS. Odd-item search in pigeons: Display size and transfer effects. *Journal of Experimental Psychology: Animal Behavior Processes*. 1989; 15:14–22. [PubMed: 2926332]
- Blough, DS. Form similarity and categorization in pigeon visual search. In: Commons, ML.; Herrnstein, R.J.; Kosslyn, SM.; Mumford, DB., editors. *Behavioral approaches to pattern recognition and concept formation Quantitative analyses of behavior*. Vol. 8. Hillsdale, NJ, England: Lawrence Erlbaum Associates, Inc.; 1990. p. 129-143.
- Blough DS. Reaction time drifts identify objects of attention in pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes*. 1993; 19:107–120. [PubMed: 8505592]
- Blough DS. Some contributions of signal detection theory to the analysis of stimulus control in animals. *Behavioural Processes*. 2001; 54:127–136. [PubMed: 11369465]
- Blough DS, Blough PM. Form perception and attention in pigeons. *Animal Learning & Behavior*. 1997; 25:1–20.
- Blough DS, Schrier AM. Scotopic spectral sensitivity in the monkey. *Science*. 1963; 139:493–494. [PubMed: 13971824]
- Blough PM. Attentional priming and visual search in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 1989; 15:358–365. [PubMed: 2794871]
- Bond AB. The evolution of color polymorphism: crypticity, searching images, and apostatic selection. *Annual Review of Ecology, Evolution, and Systematics*. 2007; 38:489–514.
- Bond AB, Kamil AC. Apostatic selection by blue jays produces balanced polymorphism in virtual prey. *Nature*. 1998; 395:594–596.
- Bond AB, Kamil AC. Visual predators select for crypticity and polymorphism in virtual prey. *Nature*. 2002; 415:609–613. [PubMed: 11832937]
- Bond AB, Kamil AC. Spatial heterogeneity, predator cognition, and the evolution of color polymorphism in virtual prey. *Proceedings of the National Academy of Sciences of the United States of America*. 2006; 103:3214–3219. [PubMed: 16481615]
- Brown MF, Cook RG, Lamb MR, Riley DA. The relation between response and attentional shifts in pigeon compound matching-to-sample performance. *Animal Learning & Behavior*. 1984; 12:41–49.
- Cavagnaro DR, Myung JI, Pitt MA, Kujala JV. Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science. *Neural computation*. 2010; 22:887–905. [PubMed: 20028226]
- Cavagnaro DR, Pitt MA, Myung JI. Model discrimination through adaptive experimentation. *Psychonomic Bulletin & Review*. 2011; 18:204–210. [PubMed: 21327352]
- Cavoto KK, Cook RG. Cognitive precedence for local information in hierarchical stimulus processing by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 2001; 27:3–16. [PubMed: 11199512]
- Cook RG. Acquisition and transfer of visual texture discriminations by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 1992; 18:341–353.
- Cook RG. The comparative psychology of avian visual cognition. *Current Directions in Psychological Science*. 2000; 9:83–89.
- Cook RG, Goto K, Brooks DI. Avian detection and identification of perceptual organization in random noise. *Behavioral Processes*. 2005; 69:79–95.
- Cook, RG.; Hagmann, CE. Grouping and early visual processing in avian vision. In: Lazareva, OF.; Shimizu, T.; Wasserman, EA., editors. *How Animals See the World: Behavior, Biology, and Evolution of Vision*. London: Oxford University Press; 2012.

- Cook RG, Riley DA, Brown MF. Spatial and configural factors in compound stimulus processing by pigeons. *Animal Learning & Behavior*. 1992; 20:41–55.
- Darwin, C. *The origin of species*. London: Murray; 1859.
- Fechner, GT.; Howes, DH.; Boring, EG. *Elements of psychophysics*. Adler, HE., translator. New York: Holt, Rinehart & Winston; 1966.
- Fisher, RA. *The design of experiments*. Oxford, England: Oliver & Boyd.; 1935.
- Gallo P, Chuang-Stein C, Dragalin V, Gaydos B, Krams M, Pinheiro J. Adaptive designs in clinical drug development—An executive summary of the PhRMA working group. *Journal of Biopharmaceutical Statistics*. 2006; 16:275–283. [PubMed: 16724485]
- Goodenough, FL. *Goodenough-Harris drawing test*. Harcourt: Brace & World New York; 1963.
- Kamil, AC.; Bond, AB. Cognition as an independent variable: Virtual ecology. In: Bekoff, M.; Allen, C.; Burghardt, GM., editors. *The cognitive animal: Empirical and theoretical perspectives on animal cognition*. Cambridge, MA: MIT Press; 2002. p. 143-149.
- Katz, JS.; Cook, RG. The multiplicity of visual search strategies in pigeons. In: Soraci, S., Jr; Murata-Soraci, K., editors. *Visual information processing*. Westport, CT: Praeger Publishers/Greenwood Publishing Group, Inc.; 2003. p. 201-222.
- Kendrick DF, Wright AA, Cook RG. On the role of memory in concept learning by pigeons. *Psychological Record*. 1990; 40:359–371.
- Krechevsky I. "Hypotheses" in rats. *Psychological Review*. 1932; 39:516–532.
- Kuhn, TS. *The structure of scientific revolutions*. University of Chicago press; 1996.
- Lamb MR, Riley DA. Effects of element arrangement on the processing of compound stimuli in pigeons (*Columba livia*). *Journal of Experimental Psychology: Animal Behavior Processes*. 1981; 7:45–58.
- Lea SEG, Goto K, Osthaus B, Ryan CME. The logic of the stimulus. *Animal Cognition*. 2006; 9:247–256. [PubMed: 16909234]
- Leith CR, Maki WS. Attention shifts during matching-to-sample performance in pigeons. *Animal Learning & Behavior*. 1975; 3:85–89.
- Mazur, JE. An adjusting procedure for studying delayed reinforcement. In: Commons, JEMML.; Nevin, JA.; Rachline, H., editors. *Quantitative analysis of behavior*. Vol. 5. Hillsdale, NJ: Erlbaum; 1987. p. 55-73.
- Mazur JE. Choice with certain and uncertain reinforcers in an adjusting-delay procedure. *Journal of the Experimental Analysis of Behavior*. 1996; 66:63. [PubMed: 8755700]
- Müller P, Berry DA, Grieve AP, Smith M, Krams M. Simulation-based sequential Bayesian design. *Journal of statistical planning and inference*. 2007; 137:3140–3150.
- Neisser, U. *Cognition and reality: Principles and implications of cognitive psychology*. WH Freeman/ Times Books/Henry Holt & Co.; 1976.
- Nickerson RS. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology; Review of General Psychology*. 1998; 2:175.
- Orne MT. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist; American Psychologist*. 1962; 17:776.
- Peissig JJ, Kirkpatrick K, Young ME, Wasserman EE, Biederman I. Effects of varying stimulus size on object recognition in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*. 2006; 32:419. [PubMed: 17044744]
- Platt JR. Strong inference. *Science*. 1964; 146:347–353. [PubMed: 17739513]
- Richards JB, Mitchell SH, De Wit H, Seiden LS. Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the Experimental Analysis of Behavior*. 1997; 67:353. [PubMed: 9163939]
- Spetch ML, Friedman A. Comparative cognition of object recognition. *Comparative Cognition & Behavior Reviews*. 2006; 1:12–35.
- Spetch ML, Mondloch MV. Control of pigeons spatial search by graphic landmarks in a touch-screen task. *Journal of Experimental Psychology: Animal Behavior Processes*. 1993; 19:353–372.

- Thorndike EL. Animal intelligence: An experimental study of the associative processes in animals. *Psychological Monographs: General and Applied*. 1898; 2
- Treutwein B. Adaptive psychophysical procedures. *Vision Research*. 1995; 35:2503–2522. [PubMed: 8594817]
- Wright AA, Cook RG, Rivera JJ, Sands SF. Concept learning by pigeons: Matching-to-sample with trial-unique video picture stimuli. *Animal Learning & Behavior*. 1988; 16:436–444.
- Wright AA, Cumming WW. Color-naming functions for the pigeon. *Journal of the Experimental Analysis of Behavior*. 1971; 15:7–17. [PubMed: 16811491]
- Zentall TR, Hogan DE, Edwards CA, Hearst E. Oddity learning in the pigeon as a function of the number of incorrect alternatives. *J Exp Psychol Anim Behav Process*. 1980; 6:278–299. [PubMed: 7391753]

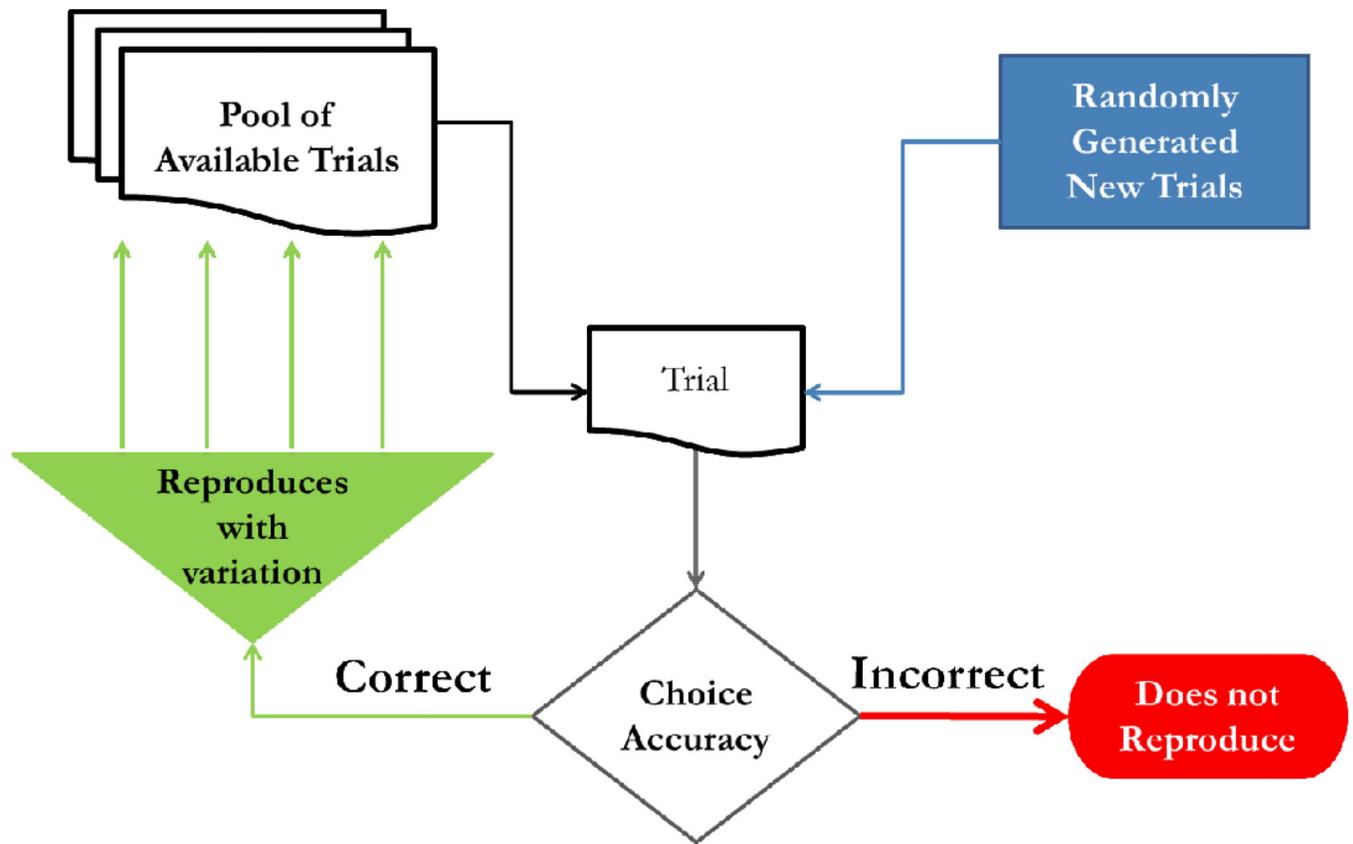


Figure 1.

Diagram of the genetic algorithm used in these experiments. Sessions were composed of mixture of trials generated from those selected based on the bird's prior correct choices (top left path) and trials generated at random (top right path). Incorrect choices caused those displays to not reproduce. These were not permanently eliminated as the same display could return in future random trials.

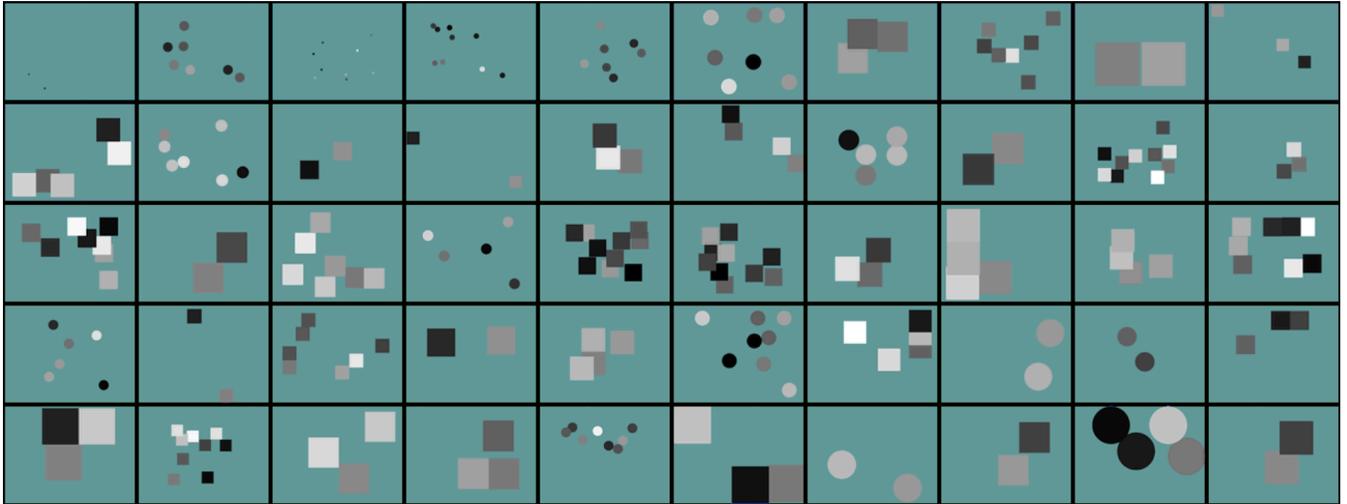


Figure 2.
Representative sample of random new displays typical of the initial conditions used at the beginning each cycle.

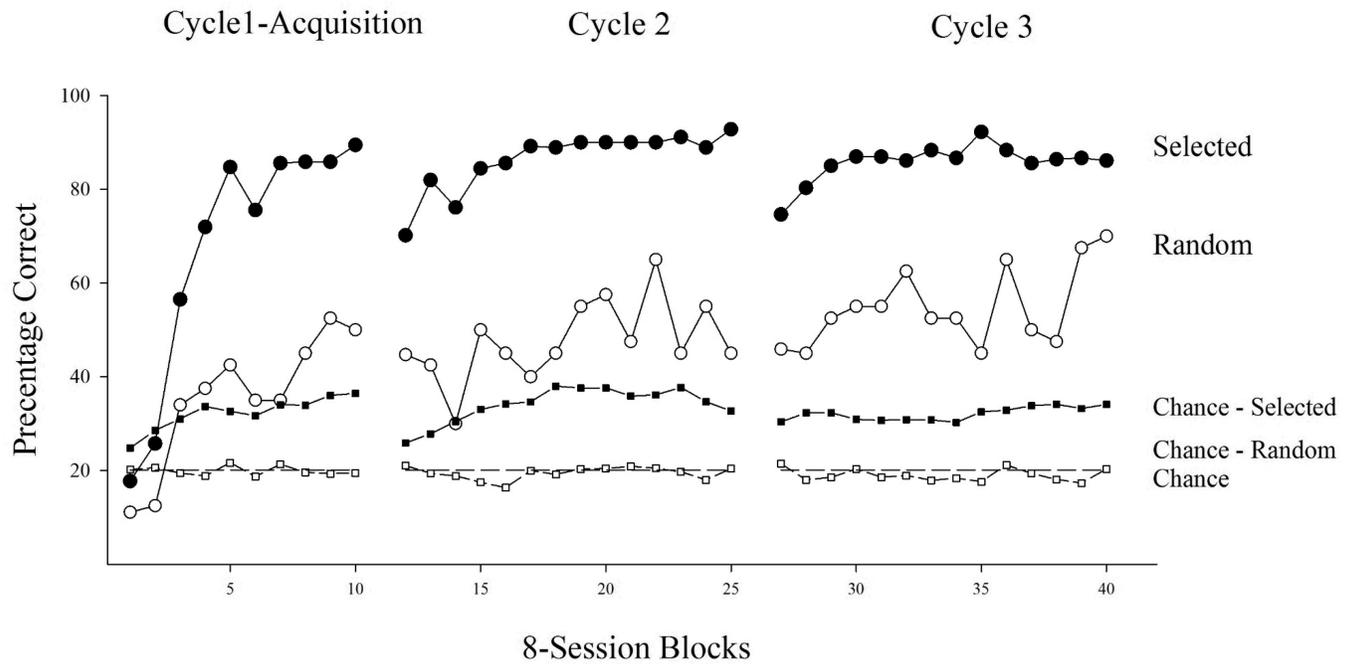
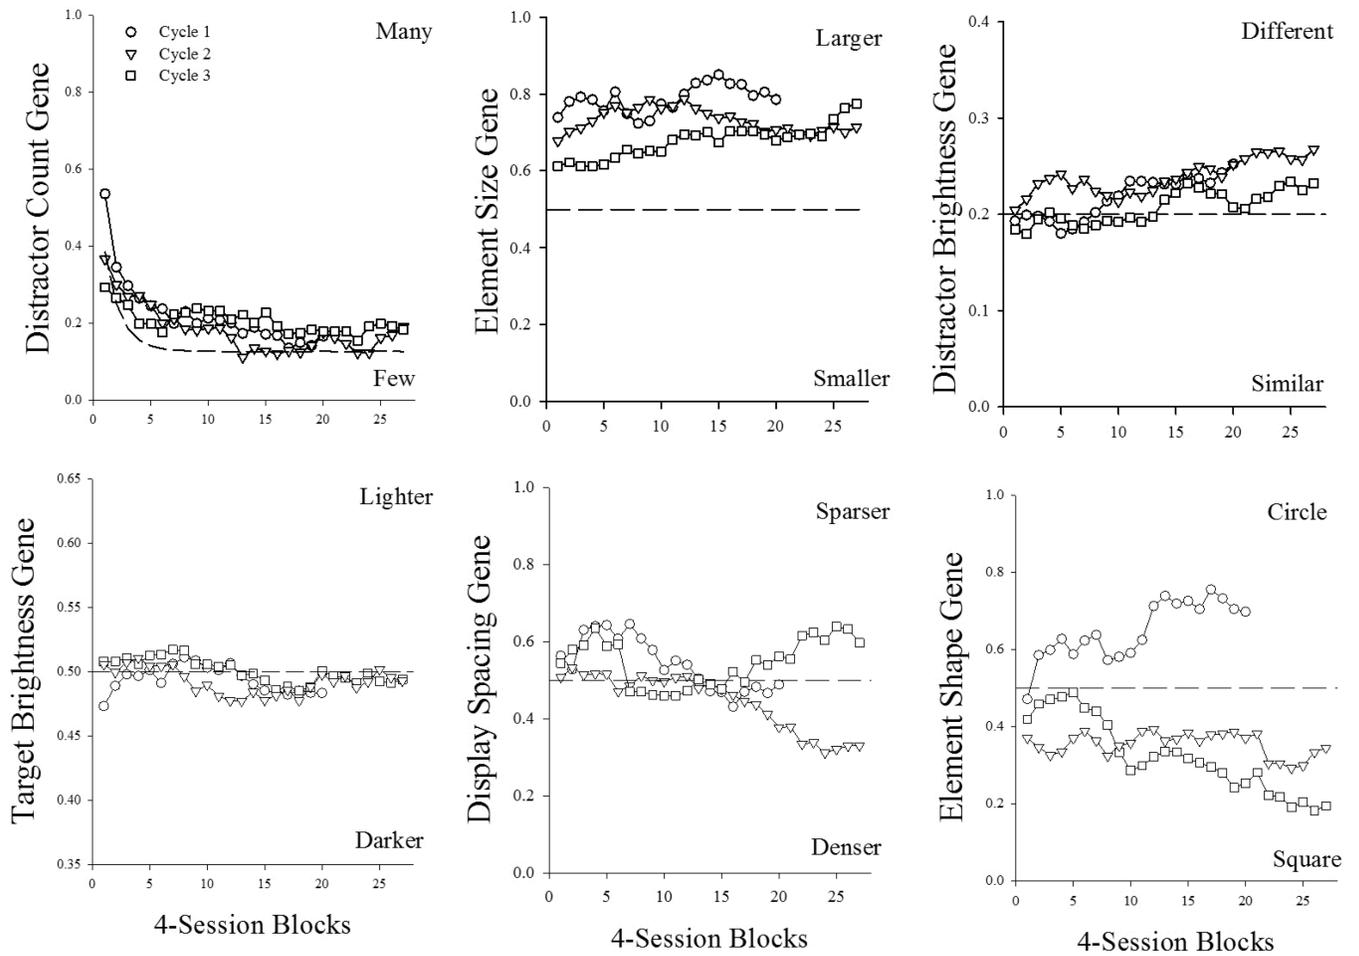


Figure 3.

Choice accuracy for selected (black circles) and new (open circles) trials over the successive cycles of Experiment 1. The bottom three lines show the theoretical (dotted line) and empirical (open squares) chance expectation for random trials and empirical chance expectation for selected trials (black squares) based on the number of elements appearing the actual displays of a specific session.

**Figure 4.**

The mean value for the six genes used to generate selected trials during the three cycles of Experiment 1 (circle = cycle 1/acquisition, triangle = cycle 2; square = cycle 3). The dashed lines show the expected value of each gene given no differential selection.

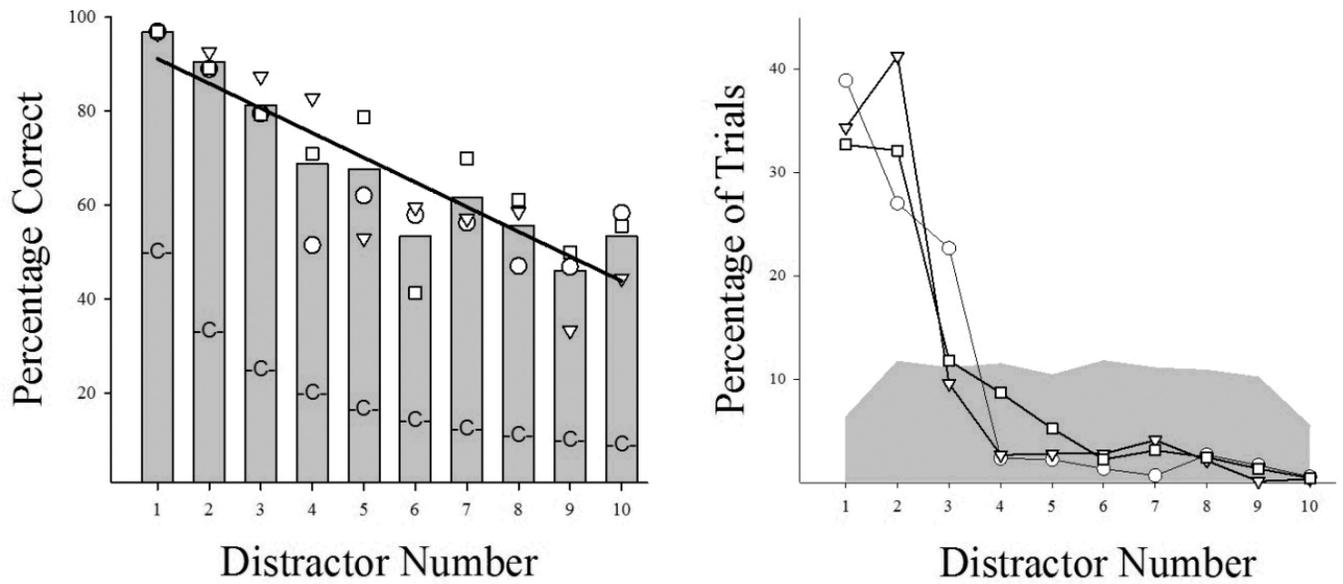


Figure 5.

The left panel shows mean accuracy as a function of distractor count over the last 20 sessions of each cycle of Experiment 1. The bars show averaged accuracy across all cycles (the basis for the included best-fitting regression line), while the symbols show accuracy within each cycle (circle = cycle 1, triangle = cycle 2; square = cycle 3; c = chance). The right panel shows the population distributions of distractor number for selected trials for each cycle over the same 20 sessions. The gray region shows the empirical population distribution of distractor count from random trials.

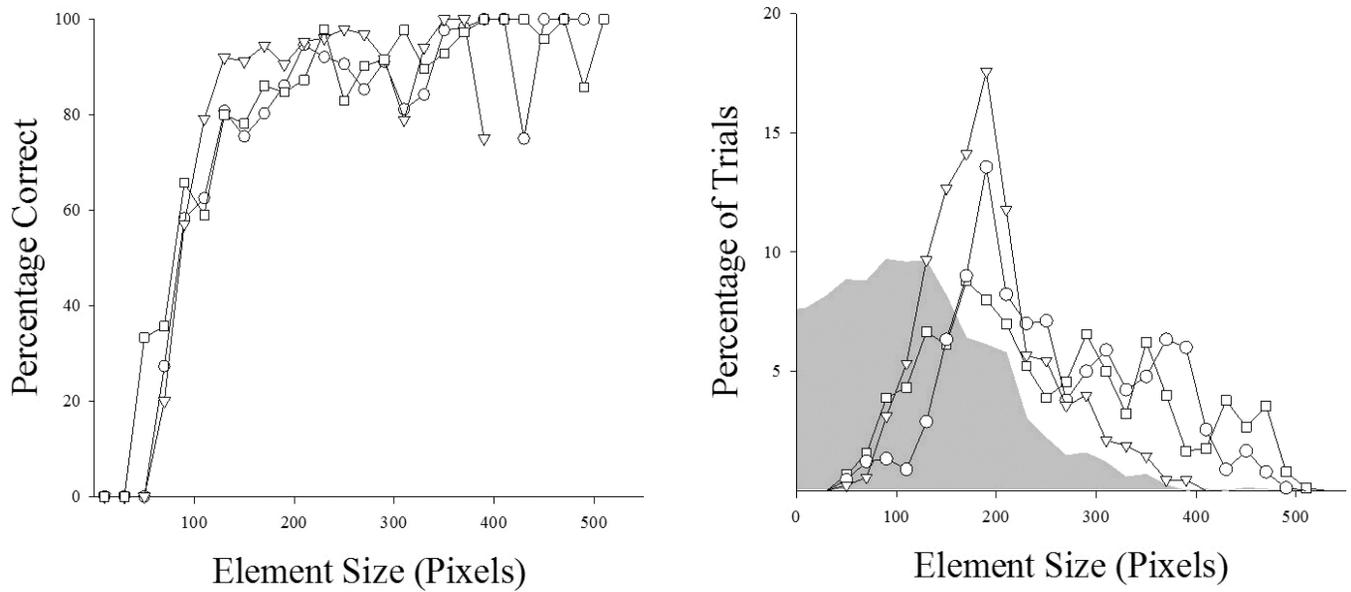


Figure 6.

The left panel shows choice accuracy as a function of element size over the last 20 sessions of each cycle of Experiment 1 (circle = acquisition/cycle 1, triangle = cycle 2; square = cycle 3). The right panel shows the population distribution of element size for selected trials for each cycle over these last same 20 sessions. The gray region shows the empirical population distribution of element sizes from random trials.

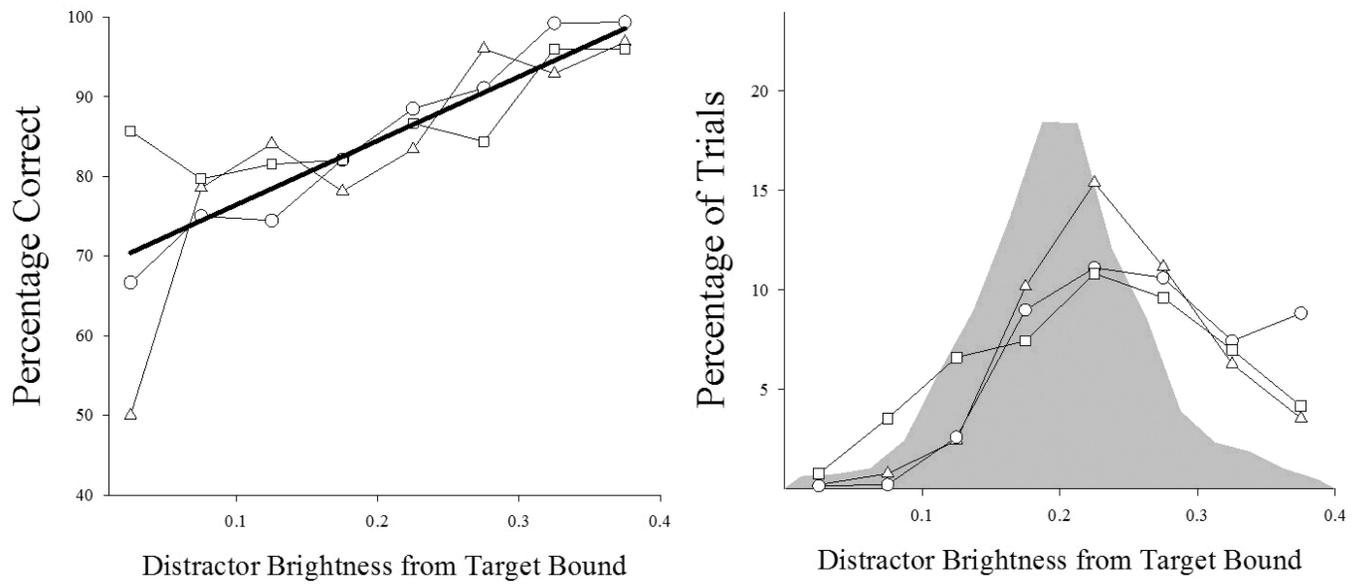


Figure 7.

The left panel shows choice accuracy as a function of average distractor brightness over the last 20 sessions of each cycle of Experiment 1 (circle = acquisition/cycle 1, triangle = cycle 2; square = cycle 3). The right panel shows the population distribution of distractor brightness for selected trials for each cycle over these last same 20 sessions. The gray region shows the empirical population distribution of these values from random trials.

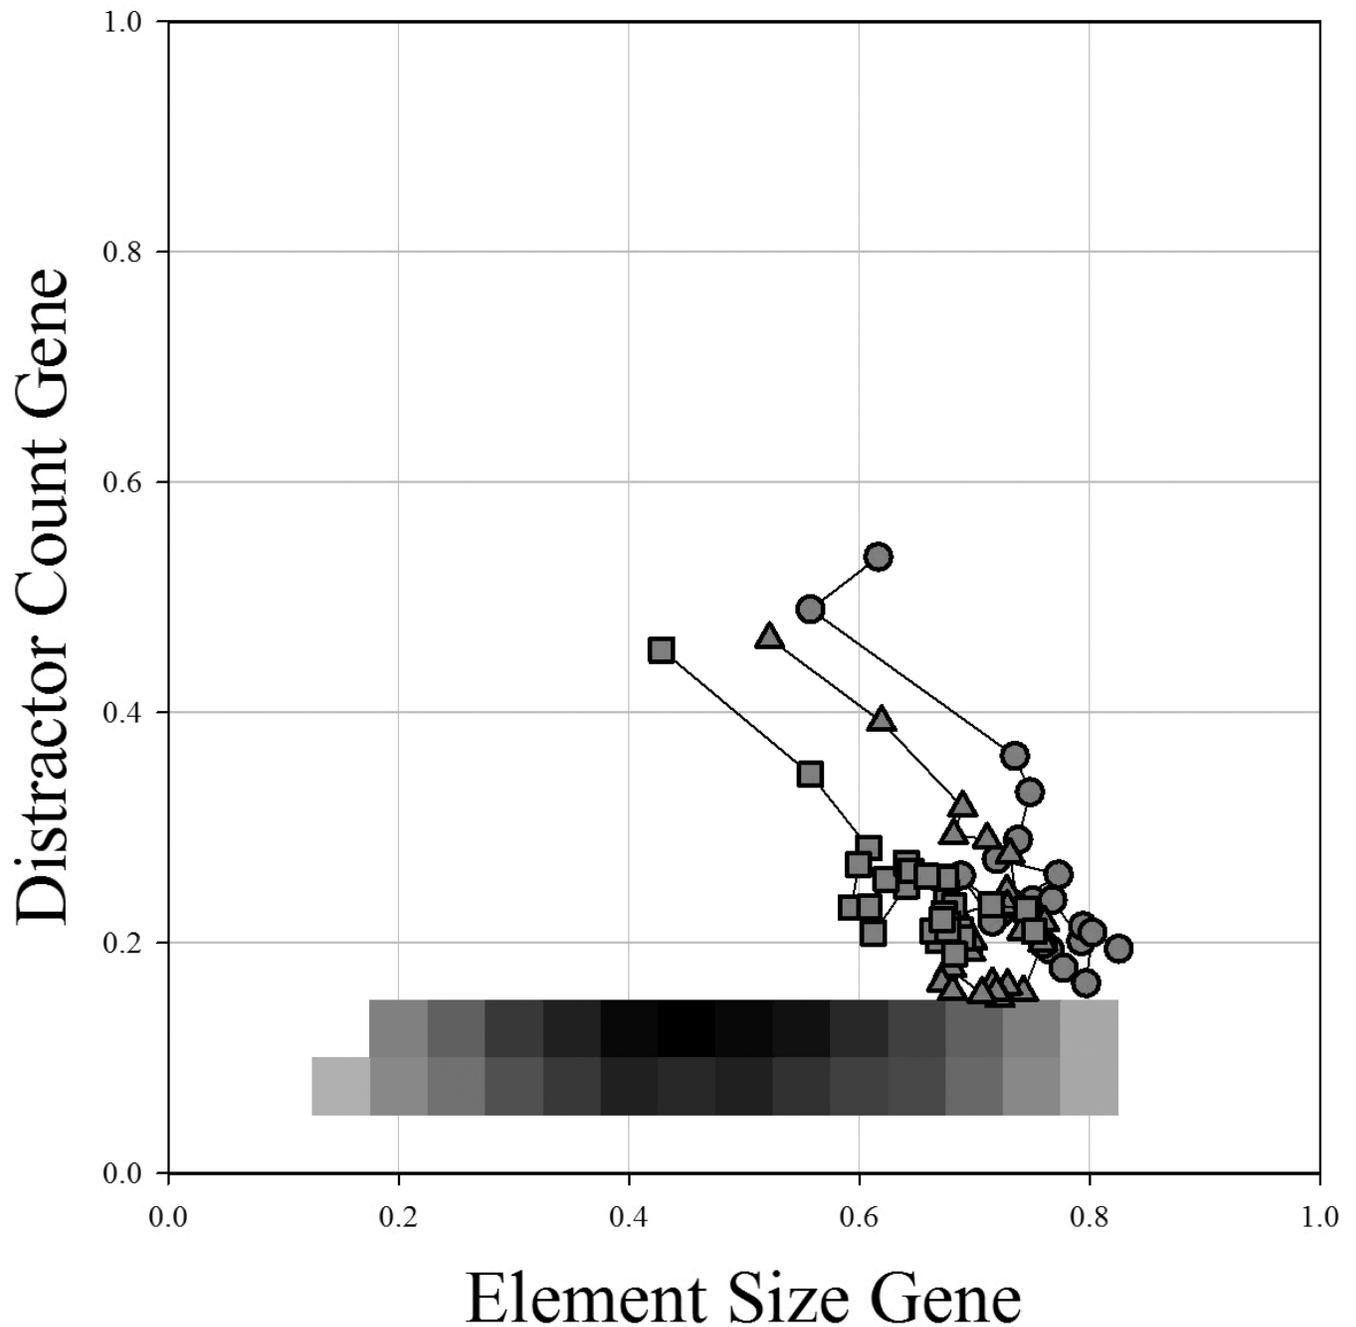


Figure 8.

A solution space for Experiment 1 as defined by the element size and distractor number genes. The three lines show the bird's trajectory from each cycle derived from the mean values of these genes in four-session blocks (circle = cycle 1, triangle = cycle 2; square = cycle_3). The gray region shows the 95% probability area of the terminal solutions derived from 3500 Monte Carlo simulations involving no differential selection. The brightness of area indicates the relative density of these random walk simulations at different points in the space. (black = highest density).

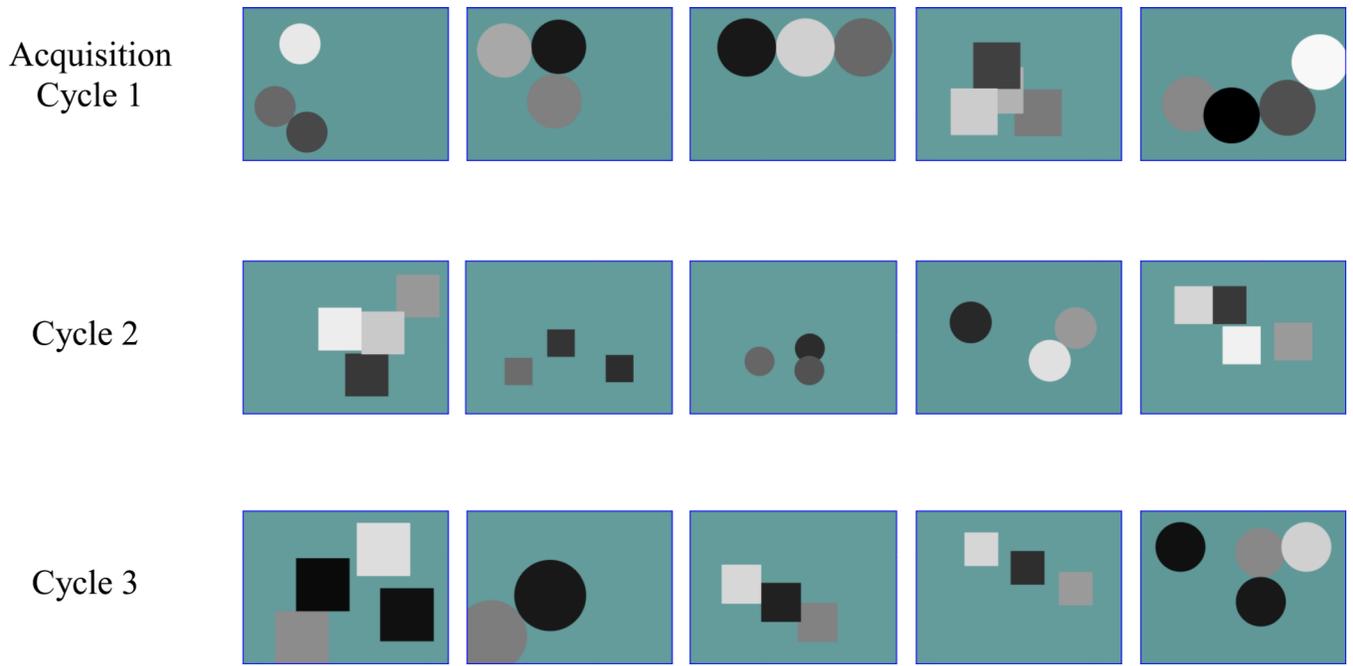


Figure 9. Median display type from each of the last five sessions of the three cycles of Experiment 1. The method for selecting these representative displays is in the text.

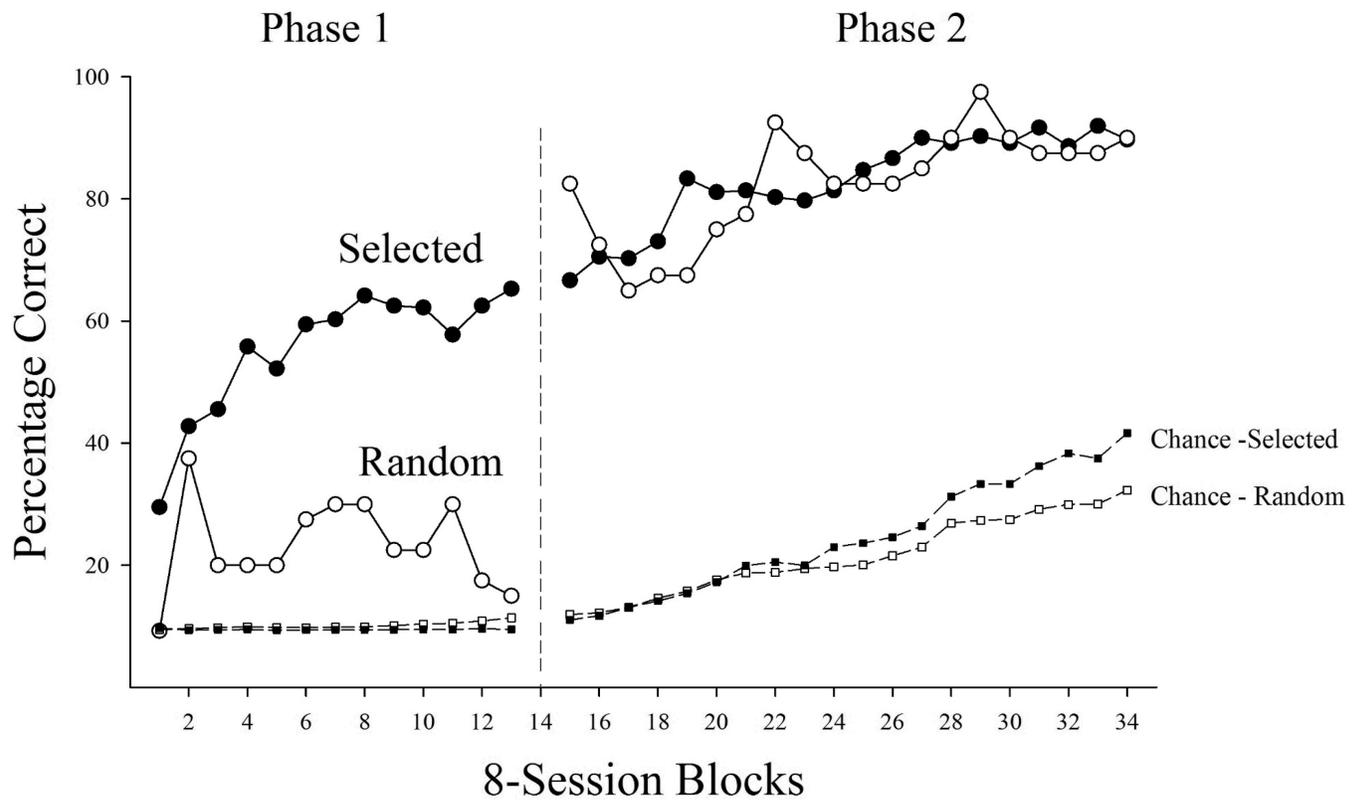


Figure 10.

Choice accuracy for selected (black circles) and new (open circles) trials over the different phases (dotted dividing line) of Experiment 2. The bottom two lines show the empirical chance expectation for new (open squares) trials and selected trials (black squares) during Experiment 2. The data for one eight session block at the junction between the two phases are not shown.

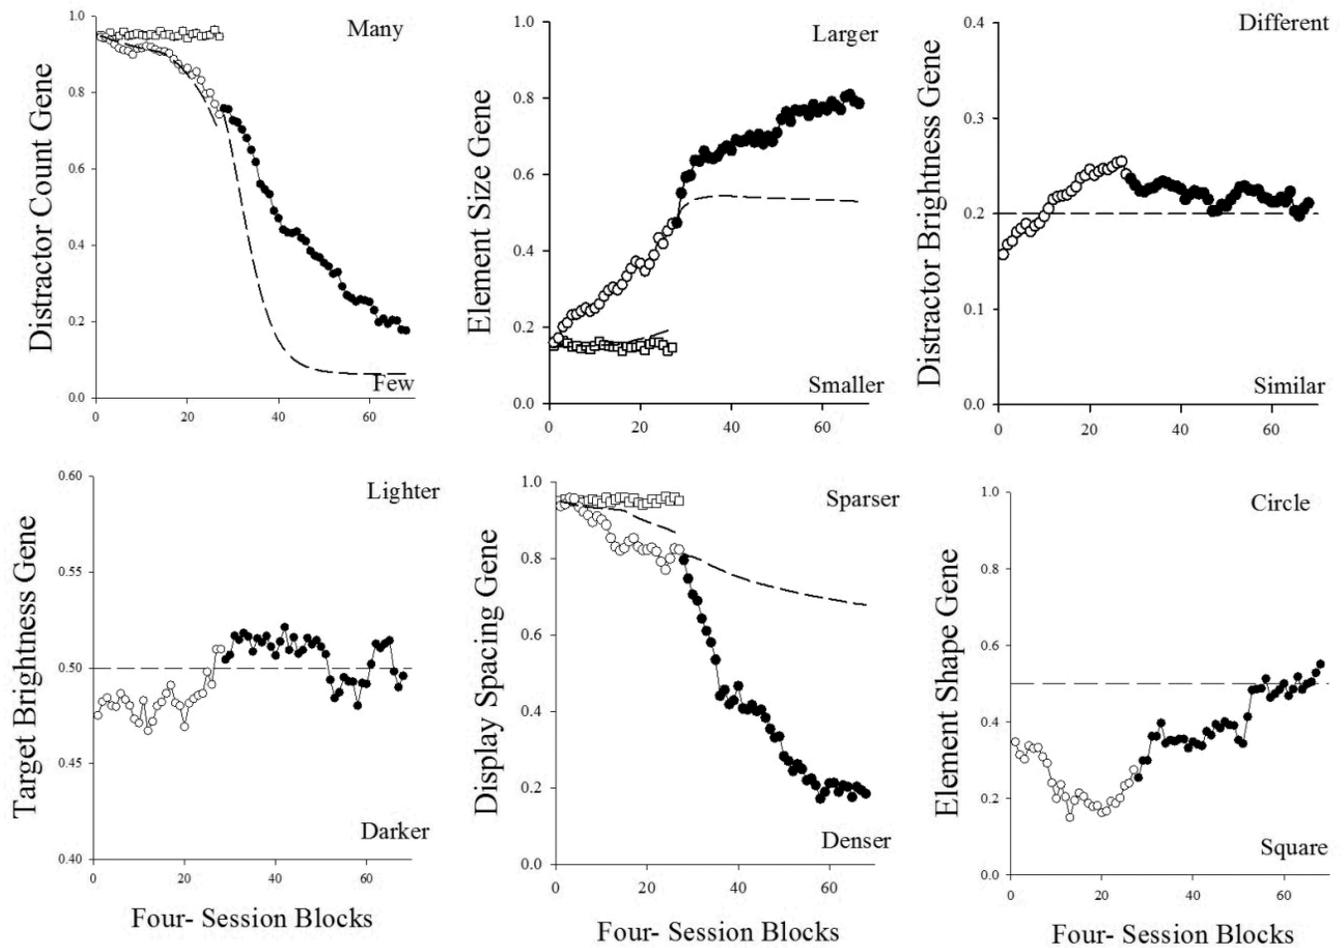


Figure 11.

The mean value for six genes used to generate displays across Experiment 2 (open circles = phase 1 selected trials; open squares phase 1 random trials, black circles=phase 2 selected). These values are derived from just *selected* trials of each session. The dotted lines show the expected gene values given no differential selection.

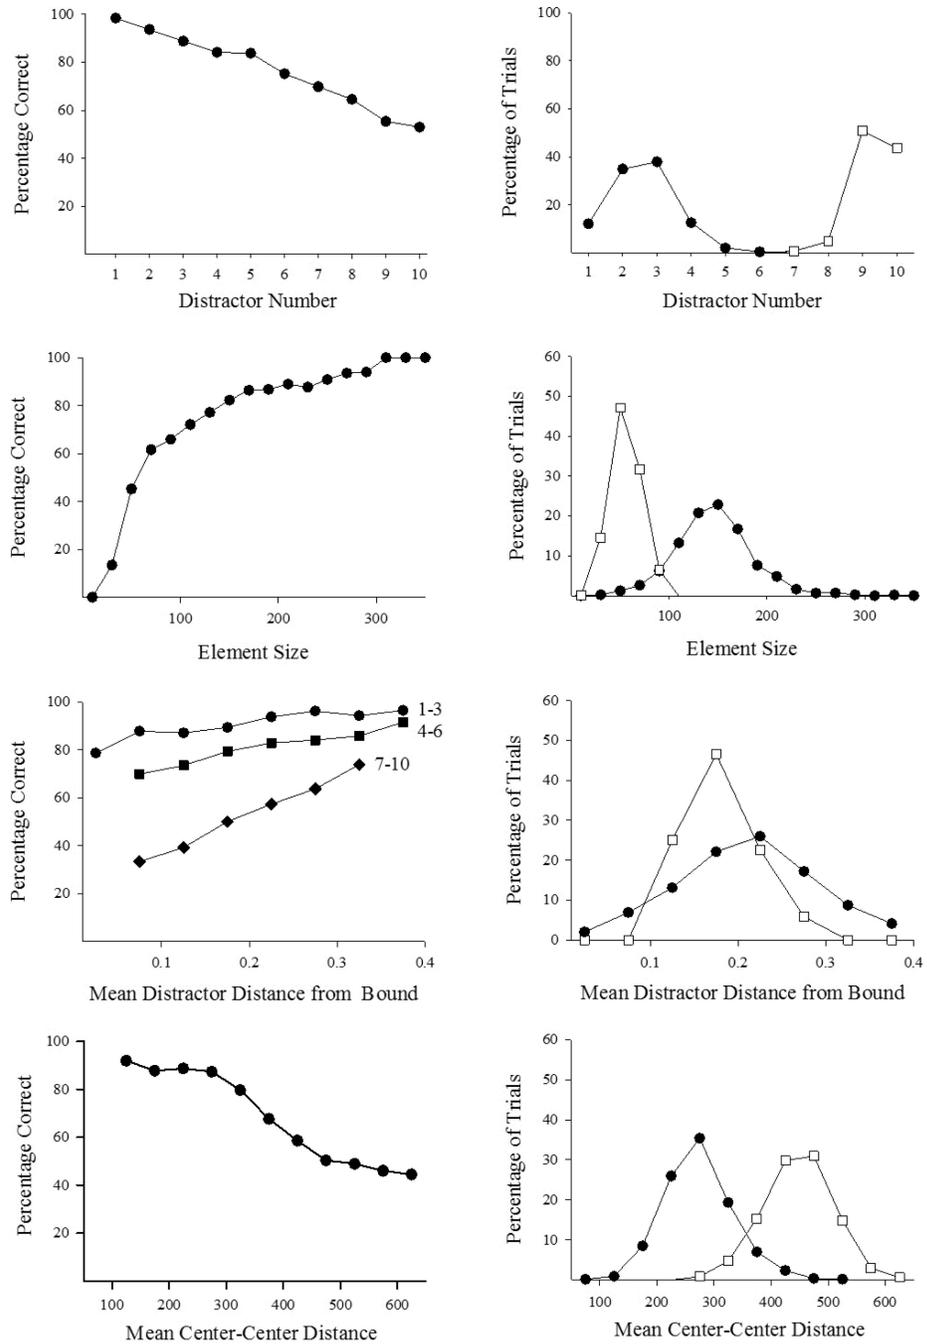


Figure 12.

Each row of panels shows the accuracy (left panels) and population distributions (right panels) for the distractor count, element size, distractor brightness and display spacing genes in Experiment 2. Accuracy is derived from choices made over the entire experiment. The right panels show the shift in population distributions from the first 32 (open squares) the last 32 sessions (black circles) of the experiment for each gene.

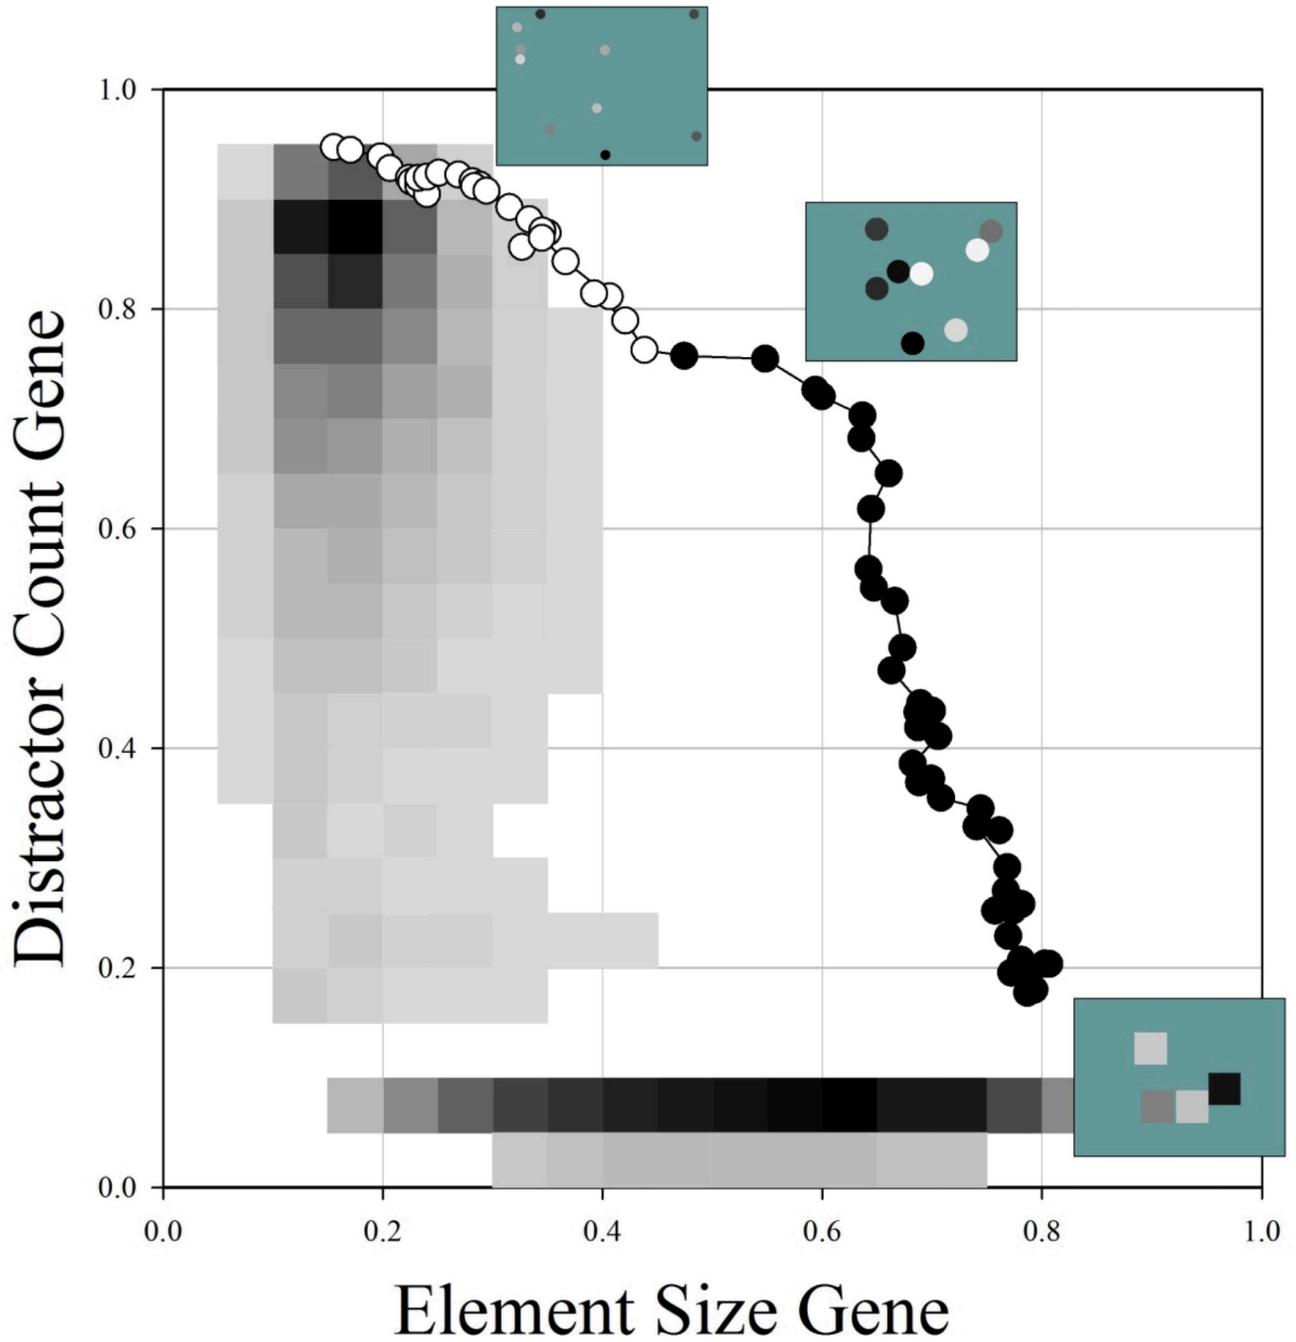


Figure 13. The solution space for Experiment 2. The trajectory shows the mean values of the distractor count and element size genes in four-session blocks over Phase 1 (white circles) and Phase 2 (black circles) of the experiment. The gray region shows the 95% probability area of the terminal solutions derived from 3500 Monte Carlo simulations involving no differential selection. The brightness of area indicates the relative density of these random walk simulations at different points in the space. (black=highest density).