



Contents lists available at ScienceDirect

Behavioural Processes

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



Visualizing search behavior with adaptive discriminations

Robert G. Cook*, Muhammad A.J. Qadri

Department of Psychology, Tufts University, USA

ARTICLE INFO

Article history:

Received 5 July 2013
Received in revised form
12 December 2013
Accepted 13 December 2013

Keywords:

Adaptive discriminations
Genetic algorithms
Pigeons
Discrimination learning
Visual search

ABSTRACT

We examined different aspects of the visual search behavior of a pigeon using an open-ended, adaptive testing procedure controlled by a genetic algorithm. The animal had to accurately search for and peck a gray target element randomly located from among a variable number of surrounding darker and lighter distractor elements. Display composition was controlled by a genetic algorithm involving the multivariate configuration of different parameters or genes (number of distractors, element size, shape, spacing, target brightness, and distractor brightness). Sessions were composed of *random* displays, testing randomized combinations of these genes, and *selected* displays, representing the varied descendants of displays correctly identified by the pigeon. Testing a larger number of random displays than done previously, it was found that the bird's solution to the search task was highly stable and did not change with extensive experience in the task. The location and shape of this attractor was visualized using multivariate *behavioral surfaces* in which element size and the number of distractors were the most important factors controlling search accuracy and search time. The resulting visualizations of the bird's search behavior are discussed with reference to the potential of using adaptive, open-ended experimental techniques for investigating animal cognition and their implications for Bond and Kamil's innovative development of virtual ecologies using an analogous methodology.

This article is part of a Special Issue entitled: CO3 2013.

© 2014 Published by Elsevier B.V.

1. Introduction

Understanding how animals solve discriminations has been at the heart of the study of learning since its inception (Köhler, 1925; Thorndike, 1898; Tolman and Honzik, 1930). An important theoretical concept for understanding any discrimination revolves around the idea of a *solution space*. For any discrimination, there is a multidimensional space that represents the different potential operations used by an animal to solve and perform a specific task, their likelihoods, and potential order of occurrence. These solutions operate on the features available within a task and consist of those cognitive operations and processes by which stimuli come to differentially control behavior. Because understanding these complex psychological spaces is critical for a complete account of animal behavior in the real world, the development of procedures to explore, map, and analyze them is important. The controlled parametric manipulation of experimental variables has been one time-honored means of doing so. In recent years, however, other computational approaches have offered new

possibilities for investigating these questions. Genetic algorithms represent one powerful type of this approach.

Genetic algorithms (GAs) are a form of optimization procedure that can do an open-ended search of a problem's potential solution(s). GAs have had an increasingly widespread impact across different areas of science and engineering, including biology, economics, bioinformatics, robotics, and machine learning. We think these and associated evolutionary methods have substantial potential for the investigation and identification of the cognitive processes and operations engaged by animals in discriminative settings. Their application and impact in psychology has only been nominal and their utilization for the experimental control and measurement of either human or animal behavior limited.

We think these adaptive procedures hold a number of advantages. First, they are ideal for quickly searching large multivariate, parametric spaces. This allows the analysis of situations that are much closer in their complexity and dynamics to the real world than most experimental settings. Second, they are open-ended, automatic, and make few prior assumptions about the nature of the space or its appropriate solution. Thus, any demand characteristics imposed by the procedures or the shape of the problem space are reduced. Third, by their fundamental organization, GAs are temporally extended in nature allowing the possibility to observe dynamics and changes in cognition over time. Fourth, they are subject-driven. Because the selection is done by the animal's

* Corresponding author at: Department of Psychology, Tufts University, 490 Boston Avenue, Medford, MA 02155, USA. Tel.: +1 617 627 2546; fax: +1 617 627 4132.

E-mail address: Robert.Cook@tufts.edu (R.G. Cook).

behavior, the animal has more direct control of the form of the discrimination. In one sense, GAs causes the procedure to “adapt” to the animal rather than vice-versa. Finally, by reconceptualizing how we can approach the study of animal behavior, it generates new ideas and theories about the best way to understand how animals interact with the environment.

In an innovative and initial application of a GA to the study of animal behavior, Bond and Kamil developed a “virtual ecology” to investigate the evolution and maintenance of polymorphism in prey species (Bond, 2007; Bond and Kamil, 1998, 2002, 2006; Kamil and Bond, 2001, 2002). In these studies, blue jays searched for and detected different kinds of digital “moths” on a computer screen over a series of trials. The visual properties of different “species” were then controlled by a set of underlying *parameters or genes* that controlled their *phenotypic* appearance and structure. Analogous to the real world, if a moth was detected on the display screen, it was “eaten” by the jay and its genes did not contribute to the next generation of possible prey items used to create future trials. Thus, over trials, the featural composition of the moth stimuli gradually transformed because of the differential selection behavior of the blue jays.

Using this GA technique, Bond and Kamil examined the role of predator-generated crypsis and frequency-dependent selection on the polymorphic nature of prey populations. They found that the jays caused the prey populations to become perceptually more cryptic, with a graded relationship between detection time and reproduction. Furthermore, the prey items became phenotypically more diverse when presented on varying backgrounds. The jays demonstrated a tendency to select those prey items that were more similar to recently encountered items. This “overselection” of certain moth phenotypes was beyond what would be expected given their density in the underlying population. This suggests that that sequential selective attention or a search image for different visual features of the moths may have been employed by the blue jays. The conjoint operation of these aspects of blue jays’ visual cognition resulted in the dynamic maintenance of a polymorphism among “moth” stimuli within the virtual ecology. This research elegantly demonstrated how adaptive techniques can be used to understand how cognition functionally influences the evolution and structure of the natural world.

Because of its potential, we have recently been employing a GA to investigate visual cognition in a different bird species, the pigeon. While we have several projects ongoing looking at different phenomena using this technique, our first project of this type investigated how pigeons solve a visual search problem in which the GA was continuously employed to govern the evolution and organization of the trials over sessions (Cook and Qadri, 2013). On each trial, the animal had to accurately locate and peck an intermediate gray *target* element from among a variable number of surrounding darker and/or lighter *distractor* elements of varied spacing. The displays were generated from parametric variables or genes that controlled distractor number, element size, shape, spacing, target brightness, and distractor brightness. The GA resulted in the composition of the visual displays evolving because of the pigeon’s differential accuracy with the large number of diverse trials that were possible from the different combinations of display genes. One important difference from Bond and Kamil’s GA procedure is that we selectively retained the genes from successful target identification displays and eliminated those associated with incorrect responding, whereas in their procedure a successful search by the blue jays resulted in that display’s genes being eliminated from reproducing and the resulting population. The differences produced by such “positive” versus “negative” selection operators are an important area for exploration with GAs in the future.

Cook and Qadri (2013) reported two experiments examining how a pigeon’s selection behavior dynamically altered displays

within the search task. In their experiments, the size of the elements and the number of distractors in the displays were the principal factors determining the pigeon’s search accuracy. The brightness of the distractors, the shape of the elements, and the relative spacing of items on the display made secondary and sometimes variable contributions. This outcome suggested the existence of a set of values within the larger multivariate stimulus space of possible displays that were best given the bird’s solution to the task. One way to conceptualize this solution is as a set of cognitive operations that form a type of stable psychological “attractor”. This attractor is the pattern or solution from within the larger multidimensional solution space that represents the perceptual, cognitive, and decisional processes currently brought to bear on the task by the animal. By examining the different stimuli within the stimulus space that were differentially affected by this attractor, one can infer the nature of the operations comprising the implemented solution.

To test for the existence of such an attractor in the previous experiments, we tested the bird from three different sets of randomized initial conditions and from one set of controlled initial conditions. These initial conditions represented the beginning seed values for the different genes during the first session of any phase. We found that the bird’s selection behavior resulted in the visual search stimuli repeatedly evolving towards the same point in the performance space regardless of the initial conditions or starting values of the genes. This suggests the pigeon employed the same visual search solution during each iteration of the GA.

The current experiment focused on better understanding the stability and shape of this attractor. This was accomplished by capturing a high quality portrayal of the bird’s discriminative behavior over the entirety of the stimulus space tested as collected over an extended period of testing. In the prior experiments, we concentrated on tracing out the trajectory of the bird’s ongoing selection and resulting alteration of the gene populations. For that purpose, a considerable proportion of the daily trials were derived from variations of previously successfully displays. This resulted in a reduced sampling of the entire stimulus space, which prevented us from identifying the diffuse influences of the attractor. Further, because we reset the discrimination several times, we did not collect observations regarding the attractor’s stability over any extended period of time. Thus, it is unknown if the bird would stay in the vicinity of its initial solution or would or could the bird shift its solution with experience?

In the current test, we altered the GA procedure to examine better the bird’s search behavior across the entire set of stimuli by increasing the proportion of randomly generated trials within a session sixfold. As a consequence, 67% of trials within a session were randomly generated from the entire range of genotypic values rather than just the 10% “sampling rate” used in the previous experiments. This increase in the proportion of randomly generated trials tested the entire stimulus space more thoroughly than was possible before. To measure of the stability of this solution, we tested the pigeon for 200 60-trial sessions. This extended period of testing resulted in 8000 random and 4000 selected trials for analysis. Given this large collection of data, another of our goals for this experiment was to map and characterize the shape of the attractor used by the bird. Better visualizing its structure and organization over the entire stimulus space should provide a better understanding of the bird’s solution and its influence on selection and performance.

2. Method

2.1. Animal

A male White Carneaux pigeon (*Columba livia*) was tested. This bird was familiar with the task and procedures (Cook and Qadri,

2013). The bird was maintained at 85–90% of its free-feeding weight and individually caged in a colony room (12 h LD cycle) with free access to water and grit.

2.2. Apparatus

The pigeon was tested in a flat-black Plexiglas touchscreen chamber. The stimuli were displayed via computer on an LCD color monitor (NEC MultiSync LCD1525X; 1024 × 768 pixels) visible through a 29.0 cm × 21.5 cm viewing window in the middle of the chamber's front panel. Pecks to the display 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 located food hopper (Coulbourn Instruments, Whitehall, PA). A 28-V houselight was centrally located in the chamber ceiling and was always illuminated.

2.3. Procedure

2.3.1. Intermediate brightness search task

The task and its controlling properties were equivalent to those used by Cook and Qadri (2013). Here we provide an overview, and note the major differences in the two procedures. The reader is directed to the first paper for a detailed description of the procedures. The search task required the pigeon to locate and identify a *target* element of intermediate brightness (gray) from a variable number of surrounding *distractor* elements of mixed higher or lower brightness values. The composition of the size, shape, spacing, brightness, and number of distractors, in the displays was controlled by the GA procedure described below.

Each trial was initiated by a single peck to a 2.5 cm circular white ready signal, which was immediately followed by the presentation of the search display. The display was comprised of one *target* element whose brightness was within the middle 20% of the computer's brightness range and a variable number of *distractors* whose brightness values occupied the remaining 80%. The bird's choice of display elements was determined to be the first element that received three pecks. After choosing an element, the entire display was removed. A correct choice resulted in 2.1 s access to mixed grain. There were no consequences associated with an incorrect choice, except lack of reinforcement. The trial was ended after 20 s, if no element had received three total pecks. These “no choice” trials were treated as a non-selected trial and not allowed to reproduce. An inter-trial-interval of 3 s separated each trial.

2.3.2. Gene expression and display properties

Each trial's visual display was computer-generated from a set of parameters or *genes* that determined its phenotypic appearance. The displays were composed of multiple elements that varied in distractor number, element size, element shape, element spacing, target brightness and distractor brightness (see Fig. 1 for a set of randomly configured displays). These various features were

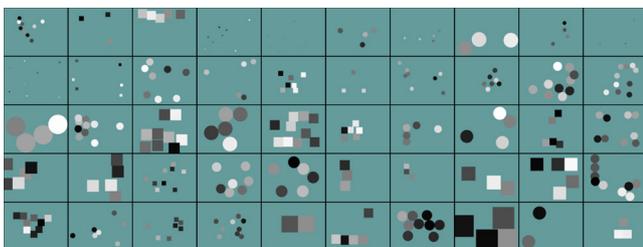


Fig. 1. Representative examples of the large number of possible random displays generated and tested in the current experiment.

controlled by independent genes that were expressed in the following ways (gene values varied from 0 to 1).

The *shape* gene probabilistically controlled whether the visual elements within a trial were either all squares or all circles. For example, a gene value of .75 would give a display a .75 chance of being composed of circles and a .25 chance of being made of squares.

The *distractor number* gene determined the number of distractor elements in a display. Between one and nine distractors could potentially appear in a display.

The *size* gene determined the size of elements in the display. Its expression was determined in conjunction with value of the distractor number and element spacing genes because of the algorithm used to spatially pack the elements in the displays. This algorithm could yield elements from 0.1 to 15.2 cm in size. All elements in each trial were the same size.

The *spacing* gene controlled how much of the overall display could be occupied by the elements. This gene's upper range started with 100% of the display's maximal size (30.5 cm × 23 cm) to 25% of this area (15.25 cm × 11.5 cm) positioned in the center of the screen. The elements were then randomly located with a uniform probability over this area. Depending on the number of elements and their size, their random positions were iteratively adjusted to reduce, but not necessarily eliminate, overlap between elements.

Each element of a display had a *brightness* gene that determined its appearance as an expressed percentage of the brightness range of the display from black to white (0–255 gray scale values or gsv). The expression of target's brightness gene range was restricted to an intermediate level within the middle 20% of the total brightness range, while the distractor elements were restricted to the upper or lower 40%. For analytic tractability, we treat these as a single *target brightness* gene and a pooled composite *distractor brightness* “gene” that measured average brightness relative to the closest bound of the target's range, taking values from 0 (very close to the target's bounds) to .5 (maximally distant) with an expected value of .25 for a randomly generated set of distractors.

2.3.3. Session generation, organization and, genetic variation

Each session's trials were composed from two types of trials. The displays of *random* trials were created from genes that were randomly assigned their value (between 0 and 1). The displays of *selected* trials were randomly drawn from the pool of possible displays (PPD). This PPD was populated with descendent displays from previously correct trials. The genome of a previously correct trial was replicated four independent times with a ±5% random variation around its gene values and added to the PPD. All trials, whether they were randomly generated or selected trials, contributed to the PPD if the pigeon correctly located the target during its presentation. The PPD retained the descendants from the five previous sessions. If the bird was incorrect or failed to respond within 20 s, the trial was considered non-selected and its genome was not transmitted to the PPD.

Four sessions were successively conducted each day over a span of approximately two hours. Each 60-trial session consisted of 40 *random* trials and 20 *selected* trials randomly chosen from the PPD. Each new session allowed a generation of displays to be assessed and generated by the GA. Within a day, sessions were separated by approximately one minute for these calculations. The current experiment reports the results of 200 sessions collected over the course of three months. The first session's trials were randomly generated (e.g. as in Fig. 1), but after that the GA controlled all trial and session composition. A total of 11,997 trials were completed over the experiment (three trials were missing because of two unfinished sessions). The pigeon recorded 8158 correct target or incorrect distractor responses over all trials, with the remaining

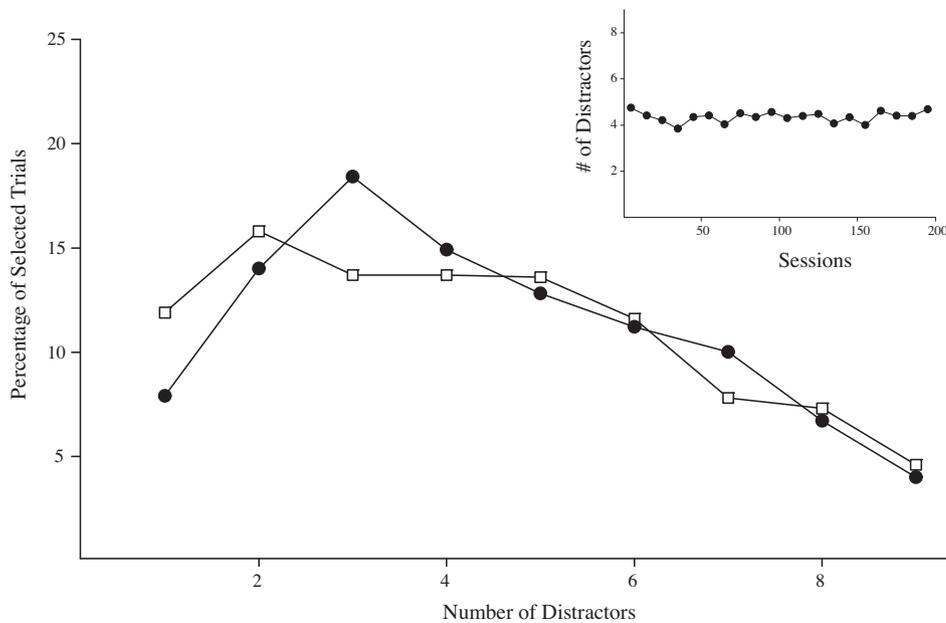


Fig. 2. The main panel shows the percentage of selected trials for the different possible distractor counts in this experiment. The white squares show the population of such displays from the first 50 sessions of the experiment, while the black circles show the same distribution from the last 50 sessions. The inset panel shows the mean phenotypic value of distractor count for selected displays across the entire experiment in 10-session blocks.

3839 resulting in no criterial response of three pecks to either a single target or distractor element.

3. Results

3.1. Overall search accuracy

Because of the bird's prior experience with the task, search accuracy was very good from the beginning of the experiment. Search accuracy was 87% over all sessions for trials with a recorded response to an element. Accuracy on selected trials (92%) was significantly higher than on random trials (84%, $t(18)=8.4$; paired t -test of ten 20-session divisions of the experiment). This accuracy difference between selected trials and random trials represents the conferred advantage of gene values from a previously correct ancestor supporting more successful search behavior. Because documenting such selection shifts in the population of these different genes was a major focus of the earlier paper (Cook and Qadri, 2013), those details are not redundantly recounted here. As a result, the next section only briefly highlights the general pattern of this selection.

Analyzing only the selected trials, the genes that had the most influence on accuracy and resulting selection mirrored those found previously. The genes moving the greatest distance from their non-selected values were related to the number of distractors, element size, and distractor brightness. The values of these genes for selected trials significantly deviated from their mean non-selected values as determined from the averaging of the random trials over the entire experiment (single mean $t_s(9) > 12.1$, again derived from ten 20-session divisions). As reflected in the phenotypes of the selected displays, these genetic values produced displays having fewer distractors, larger elements, and distractors that were generally blacker and whiter than the intermediate gray targets. Element shape and spacing made detectable secondary contributions, but these were not to the same degree as reported in the previous experiments, suggesting that this bird's extensive experience with these factors perhaps reduced their impact. This ordering and relative contribution of these genes to search accuracy for

this experiment was confirmed by conducting multiple regressions identical to those described in Cook and Qadri (2013).

3.2. Attractor stability

The results just described indicate that the phenotypic nature of the selected displays moved towards the same attractor as identified by Cook and Qadri (2013). Perhaps not surprisingly, these changes in gene values happened quickly given the bird's prior experience with the task and its highly likely carryover to the current experiments. The selected gene values quickly reflected their final values within the first ten or so sessions of the start of the experiment. Figs. 2 and 3 show this over the course of the experiment in different ways. The main panel of Fig. 2 shows the distribution of phenotypic values for the number of distractors on selected trials from the first and last quarters of the experiment. Compared to the expected flat distribution of the random trials, the pigeon's higher accuracy with displays having fewer distractors caused this gene's population of selected values to shift towards fewer distractors ($\chi^2(8)=252$). Furthermore, the distribution of these values did not significantly differ between the beginning and end of the experiment, as supported by a two sample Kolmogorov–Smirnov (KS) test ($D = .058, p > .05$). The inset panel of Fig. 2 shows the mean number of distractors in the selected displays over the entire 200 sessions in 10 session blocks. Overall, the mean number of distractors in these displays was stable. Once the bird had selected the displays with lower numbers of distractors, the distribution of these displays did not significantly change with extensive testing. The best fitting trend line to these mean values had a slope of 0.005 as calculated over the entire course of the experiment (divided into 20 10-session bins) and was not significantly non-zero, $F(1,18) < 1$.

This same pattern of selection effects held true for element size. The main panel of Fig. 3 shows the distribution of phenotypic values for element size on selected trials from the first and last quarters of the experiment, while the inset panel again shows the mean value for this gene over the entire experiment. Again, there was no significant change over the experiment after its initial selection at the

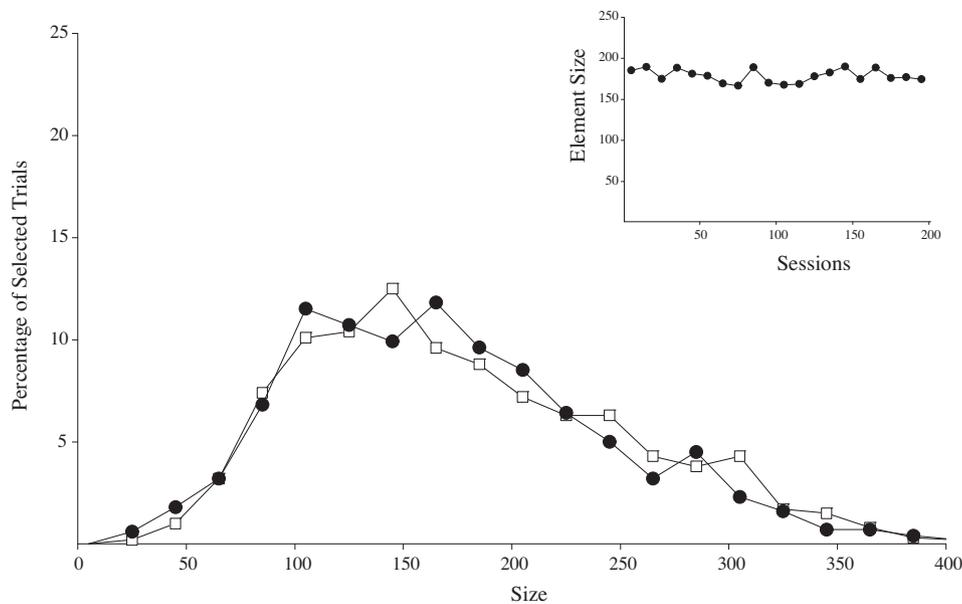


Fig. 3. The main panel shows the percentage of selected trials as a function of element size in 25-pixel bins during this experiment. The white squares show the population of such displays from the first 50 sessions of the experiment, while the black circles show the same distribution from the last 50 sessions. The inset panel shows the mean phenotypic value of element size for selected displays across the entire experiment in 10-session blocks.

beginning of the experiment. The distribution of values for selected trials was not significantly different between the first and last quarter of the experiment (KS test, $D = .055$, $p > .05$), nor did the mean value of this gene change across blocks of sessions (best fitting linear slope = $.02$, $F(1,18) < 1$). A similar examination of the phenotypic values of the other genes also revealed little or no change over these extended observations.

These highly stable patterns in the selected properties of the displays indicate that the pigeon was using and maintaining the same solution over the entire testing period. This suggests the attractor reflected in the bird's solution was highly stable and was not fluctuating to any great degree during this experiment.

3.3. Shape of the attractor

The other goal of the experiment was to evaluate and visualize the shape and extent of this stable attractor. Our method of doing this was to create a complete map documenting the effects of the bird's search solution over the entire stimulus space controlled by the GA. The current study afforded this opportunity because the inclusion of a large number of random trials evenly distributed testing of all combinations of gene values over the stimulus space. Because accuracy was the determinant of discriminative responding and its consequences, the pigeon's behavior using this measure was mapped first. Fig. 4 displays separate three-dimensional contour maps showing the pigeon's accuracy on random trials (top panel) and on selected trials (bottom panel). The illustrated stimulus space is defined by the phenotypic values of the number of distractors and element size genes. This space was selected because both the current and prior results identified these parameters as the most influential on the bird's search performance. The maps were generated from log-transformed averages of accuracy binned into equal divisions along each parameter using a built-in radial basis function gridding algorithm (Surfer 9, Golden Software Inc.). The contour lines depict the interpolated accuracy levels across the space. A region in the upper area of the surface contains no value because such displays were not possible given the design of the computer algorithm generating the stimulus displays. These *behavioral surfaces* are among the most

detailed maps of an animal's discrimination behavior over a complex, multivariate stimulus space (Blough, 1969; Herbranson et al., 1999).

The upper map of accuracy on random trials isolates a profile of the bird's solution without the direct effects of selection. The lower map of accuracy on selected trials shows accuracy when such prior selection has directly operated on the displays. Search accuracy is generally elevated on such trials compared to the random trials over the entire surface. Given that the two surfaces depict the same stimulus space as defined by the values on the two most influential genes, this increase demonstrates the benefits derived from the selected values of remaining genes not included in the space (e.g., distractor brightness).

Accuracy across both surfaces was highly similar in its basic shape: generally smooth and monotonic along its two defining parameters. There was a marked increase in accuracy from left to right that reflects the large and significant benefit of element size across a large portion of this space. The poor performance recorded for displays having the smallest element sizes was primarily a function of the pigeon's inability to direct responses to individual elements in these displays. For elements below 1.5 cm in size, the bird had a difficult time responding to the elements, with 80% of such displays not recording a criterial response to an element. On these trials, the bird often pecked the display, but had trouble recording sufficient numbers of pecks to be judged as a specific element response; hence these displays were treated as "incorrect" and not retained by the GA. Likely a combination of perceptual/cognitive (too small), mechanical (touchscreen sensitivity and resolution), and response (pecking directed at the edges of elements) factors contributed to this specific difficulty. As the size of the elements increased, the bird's accuracy rapidly increased regardless of the number of distractors in the display. Displays made of elements larger than 3 cm supported good and stable performance as the accuracy surface flattens out at this point. With elements of this size or larger, there was a very large region or plateau where the bird is reasonably accurate at the task, although the influence of distractor number became greater with increasingly larger elements. The effect of the number of distractors is harder to see in this map. The bird was consistently better with

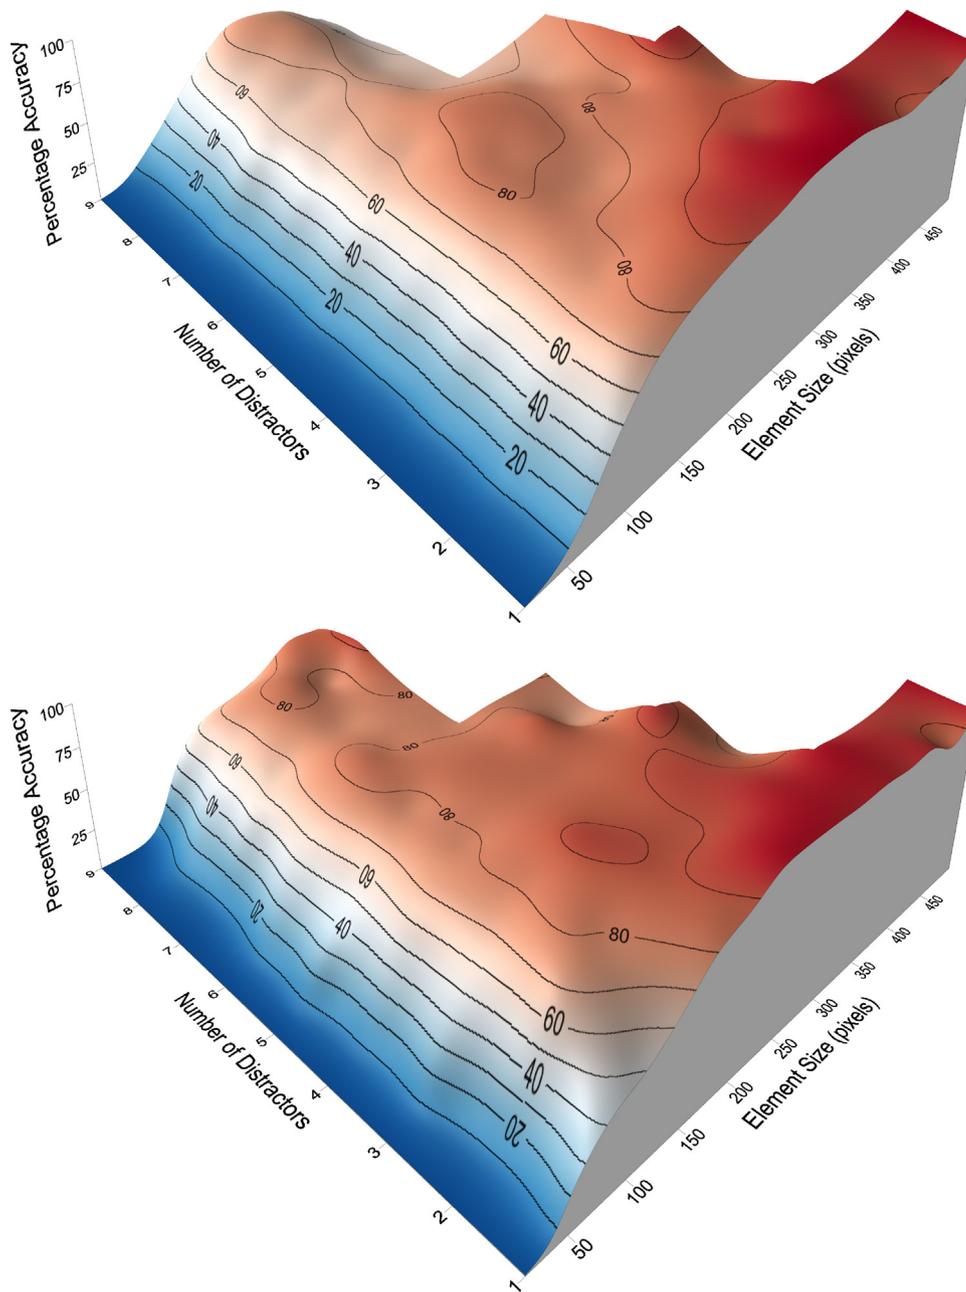


Fig. 4. Three-dimensional contour surfaces of estimated mean search accuracy as a function of a multivariate stimulus space, depicted along the values of number of distractors and element size. The upper panel shows accuracy on *random* trials. The lower panel shows accuracy on *selected* trials. Red areas indicate regions of higher accuracy and blue regions areas of lower accuracy.

fewer distractors (i.e., the strong selection effects for this gene), but was also able to perform reasonably well even when displays have numerous distractors.

There are several properties of the surfaces worth noting. One is that the same elevated region of highest accuracy occurs in the upper right of both surfaces. These higher levels of search accuracy in this portion of the stimulus space reflect the displays representing the configuration of parameters best suited to the bird's ongoing solution to finding and identifying targets in this task. This region consists of the displays having a limited number of larger target and distractor elements. This similarity of convergence across maps indicates that the implemented solution by the bird was identical for both random and selected trials. This location represents the psychological attractor for this task.

Besides this convergence, a second feature to note is the bird's generalized capacity to search accurately over a reasonably large and widespread area of the stimulus space, especially given that chance responding decreases as the number of elements increases. Despite there being a "point" in the space producing, in some sense an "optimal" display condition, towards the upper right-hand side, the bird was 80% accurate or better on selected trials over a fairly extensive combination of display values. Thus, while there was a display configuration that was best based on the bird's solution to the task, this solution was simultaneously general and flexible enough to be effective over a substantial portion of the stimulus space. Being able to see and characterize the size and shape of an animal's solution and its applicable range is an important advantage available when using open-ended procedures. The capacity to compare the generality, shape, and spatial extent of different behaviors

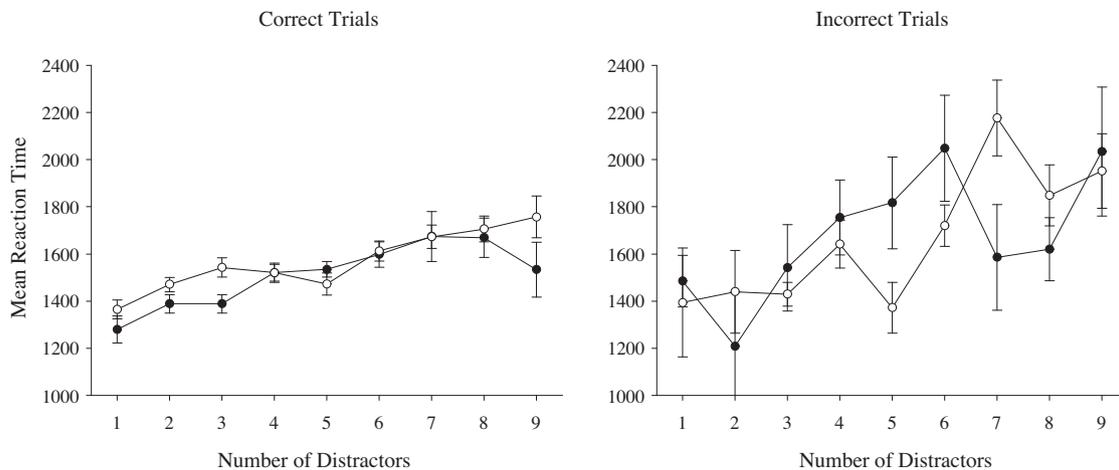


Fig. 5. Mean search reaction times as a function of number of distractors in the display, choice accuracy, and type of trial (filled = selected, open = random). Error bars indicate standard errors as determined from 10-session blocks.

over complex stimulus spaces is one of the important applications of this visualization approach to mapping discriminative behavior.

3.4. Temporal properties

While appropriately concentrating on search accuracy, because it was critical in determining whether the pigeon was rewarded or not, we thought it would be valuable to examine the temporal properties of its search behavior. Search time was based on the time to first peck the display (with second peck times replacing first responses of <200 ms; about 7% of responses) as capped at 5000 ms (<1% of trials). Overall, we found search accuracy and search time to be highly correlated, with accurate responding also being generally rapid.

Fig. 5 shows mean search time as a function of the number of distractors in the displays and choice accuracy. Overall, search times linearly increased as the number of distractors correspondingly increased, with search times being faster on correct trials than incorrect trials. A repeated measures (RM) ANOVA (number of distractors \times accuracy, using 20-session blocks as the repeated factor) revealed a significant linear effect of number of distractors, $F(1,9)=39.1$, and a marginally significant main effect of accuracy, $F(1,9)=4.65$, $p=.059$. The latter is likely attributable to the greater variability observed on the incorrect trials resulting in a somewhat less powerful comparison. Similar effects have been reported in other investigations of pigeon visual search. There also seemed to be a slight search time advantage for selected trials when compared to random trials, at least for displays with fewer distractors, but statistical evidence of this trend could not be found.

Fig. 6 is the surface map for search time over the same space as used earlier for accuracy. Because of the greater variability in our recorded temporal measurements, this surface combines search times from both correct and incorrect outcomes and for selected and random trials in order to produce a sufficiently stable map. The main feature of this surface is that it more or less directly complements the one based on accuracy, with the fastest search times occurring in those regions that supported the highest accuracy. This inverse linear relationship was confirmed by correlating the 75 bins used to generate the RT surface with the spatially corresponding data from the accuracy surface in a linear regression ($r=-.62$, $F(1,73)=44$). The similarity of the general shapes of these surfaces indicates that the influences of the attractor can be simultaneously reflected in different measures of performance.

3.5. Searches and surfaces

Finally, we thought it would be interesting and valuable to see if the results of the present experiment could be combined with those collected earlier by Cook and Qadri (2013) using a more tracking-oriented approach. The present experiment tested large numbers of random trials in order to examine performance over the entire stimulus space, resulting in the surfaces presented above. Cook and Qadri (2013), on the other hand, tested greater numbers of selected trials in their procedures and in their second experiment used a hill-climbing procedure to isolate the trajectory of the genes towards the attractor from a controlled distant location in the stimulus space.

Fig. 7 shows one such integration of these different approaches. In doing so, it is important to remember that the frequency distribution of displays produced by the computational algorithm used to construct the displays, while genetically uniform, was not phenotypically uniform. As mentioned previously, the display packing algorithm for the elements used a combination of three gene values to determine a display's final element size and distribution. As a result, very large element sizes, for example, occurred less frequently by chance than displays with slightly smaller elements. To adjust for this procedural issue, Fig. 7 shows the phenotypic frequency distribution of displays from selected trials in which we subtracted out the expected distribution of random displays over the experiment. The result best isolates the actual effect of selection on the different gene values independent of their relative frequency. As before, the higher red regions reflect displays that were repeatedly selected by the bird's accurate choice behavior, while the lower blue values reflect areas of poorer performance in which displays were not retained. Note that the peaks and troughs of this gene frequency display are similar, but not identical, to the early maps shown for accuracy (see Fig. 4). Overall, the bird's choice behavior still retained displays with fewer distractors and larger elements. Because displays with several distractors and smaller elements occurred more frequently and the bird was reasonably good at them, they were frequently retained by the GA, and thus could not be displaced by displays with even fewer distractors and somewhat larger elements which supported the highest accuracy. The gene frequency values in Fig. 7 thus represent an interaction between genotypic frequency and phenotypic selection. It nicely provides a common map, however, for integrating the results of the different experiments.

Superimposed on this surface are the tracks within this stimulus space depicting the evolution of several different *restricted* types of

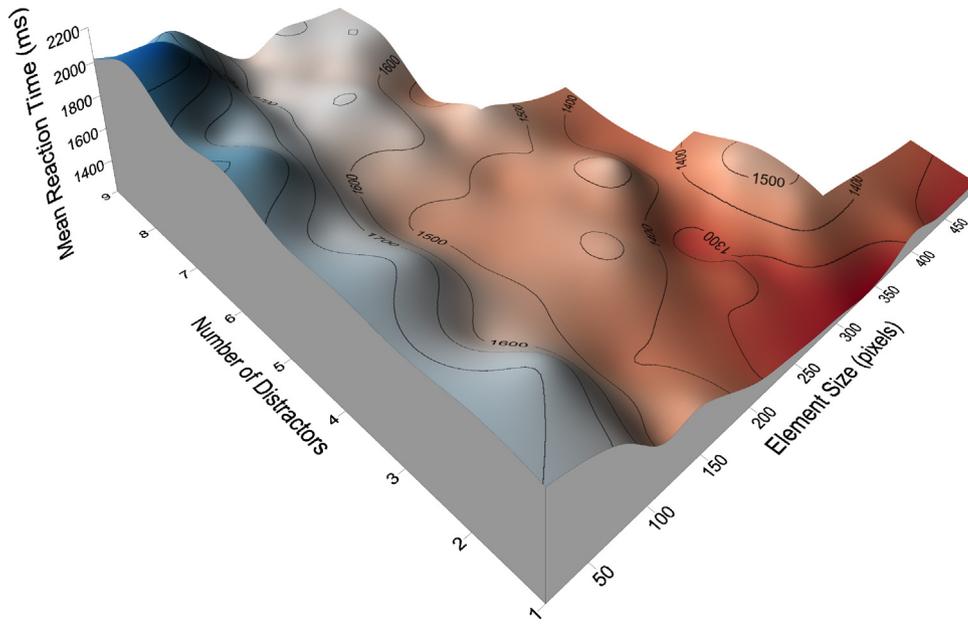


Fig. 6. Three-dimensional contour surface of search time for all trials on which an element received the critical number of pecks as a function of a multivariate stimulus space defined by number of distractors and element size. Red areas indicate regions of faster search times and blue areas indicate regions of slower search times.

displays, or “species”. The depicted tracks include the one reported from Experiment 2 of Cook and Qadri (2013), but also several new additional tracks collected using the same hill-climbing procedure, but started with other different initial conditions. In this procedure, we constrained the production of new random trials to stay within the immediate vicinity of the previously selected displays. This allows for the tracking of the genes as displays evolve over a number of generations in a manner similar to traditional hill-climbing. The new tracks were collected after Cook and Qadri’s (2013) Experiment 2, but before the results reported here. These represent three other display “species” that were initially seeded to start in other

places of the stimulus space relative to the hypothesized attractor region.

Nonetheless, the resulting trajectories for each of the different “species” conform remarkably well to the shape of the selection surface collected from the present experiment. Both species that started with small elements quickly evolved towards having larger elements. In the case of the species most distant from the attractor (circle symbols in the left corner), after increasing in element size, the number of distractors in the displays began to be reduced causing the trajectory of the species to turn right as it evolves towards the putative attractor. The species that were deficient in element

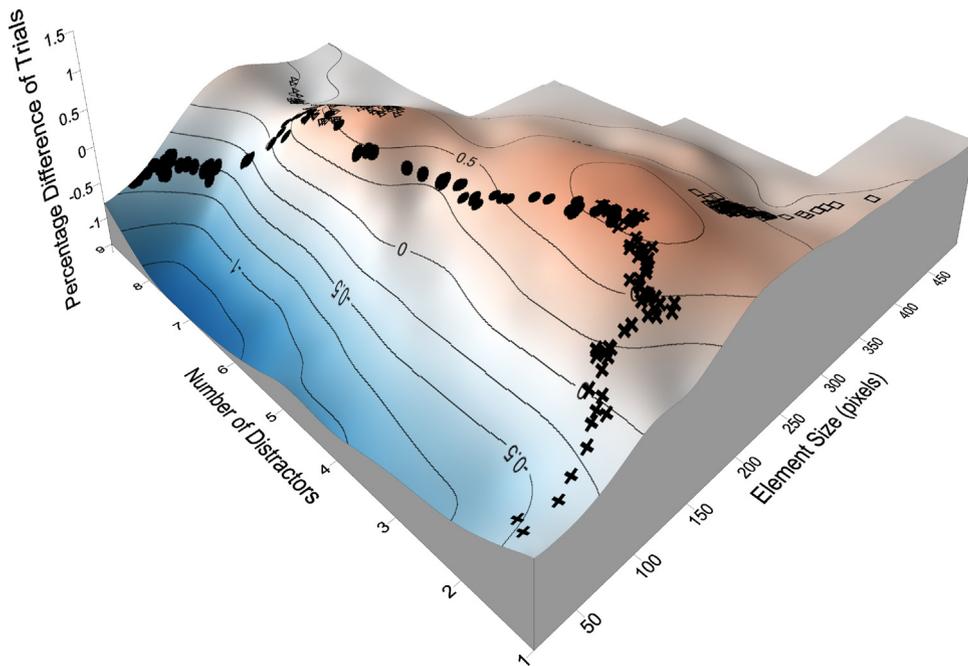


Fig. 7. Three-dimensional contour surface showing relative frequencies of selected displays adjusted by subtracting out the obtained distribution generated from random during the current experiment. Red regions show the areas with the greater relative gene frequency as retained in the selected trials of the experiment. Superimposed on this selection surface are trajectories collected using a hill-climbing procedure (see text for details) from four different sets of displays or “species” that were seeded at different initial conditions in this stimulus space. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of the article.)

size (cross and square symbols) quickly moved towards the plateau and then expanded in its spatial extent upon reaching it, as displays with more distractors successfully survived. The display that started in the top left corner (triangle symbols) moved considerably less distance during its period of testing. It may have become stuck in the possibly peaked region over the middle of the space, akin to hitting a local maximum in the surface.

Some previous research has found that displays with large numbers of distractors can generate beneficial effects of visual oddity (Zentall et al., 1980). If so, the displays in this vicinity may have benefited from this kind of oddity, indicating a potential type of secondary solution to the task. How to identify and deal with such local maxima within a complex behavioral surface is a key procedural and theoretical issue requiring more examination. Currently, it is unclear if the secondary swells and troughs within these surfaces represent noise within the process of collecting data of this type or represent stable psychological effects that need to be considered in thinking about the bird's overall solution(s) to this task. Nevertheless the convergence of the adjusted frequency or selection surface from the current experiment and the trajectories of evolving species of prior displays collected using a more tracking-oriented procedure is compelling.

4. General discussion

In the current experiment, displays representing every possible combination of the different genes were extensively tested to examine the stability and shape of the bird's solution to the search task. Overall, we found the bird's solution to the task was highly stable. The results consistently showed strong and dependable patterns of selection resulting in the evolution of the displays and their underlying population of gene values towards a single attractor or solution. Despite the GA's inherent sensitivity for detecting such changes, there was no significant variation in the selected values of the genes or the resulting displays over the months of extended testing. The subsequent visualizations of this selection behavior provides one of the most detailed maps yet of how a bird's choice accuracy performance varies over a complex, multivariate stimulus space (see also Blough, 1969; Herbranson et al., 1999; Smith et al., 2011).

We suggest that this set of best conditions can be conceptualized as a form of psychological attractor. This attractor in essence represents the bird's solution to the discrimination task as derived from the much larger candidate space of potential solutions present in the arrangement of stimuli and task parameters. As a function of an animal's internal cognitive states, past experiences, predispositions, the stimuli tested, and reinforcement contingencies in place, one, or potentially more, of these candidate solutions emerges to control behavior in any learned discrimination. This solution is a set of psychological processes that results in differential performance with the displays, which in turn, causes the trials to shift towards a location in the stimulus space that manifests the form of this attractor. Thus, the bird's approach and solution to the task in this open-ended procedure results in systematic changes of performance across the stimulus space and differential selection of the stimuli. The resulting patterns in the stimulus space can be used to make inferences about the structure of the bird's solution (Cook and Qadri, 2013). Our notion of a psychological attractor in some ways shares features with similar ideas developed in computational modeling and neuroscience. In non-symbolic models, for example, artificial neural networks represent memories and categories as a set of steady-state recurrent patterns (Hopfield, 1982; Hu and Zhang, 2010). Such steady states or fixed attractors can then be used as representations of different external stimuli (Corneille et al., 2007; Tanaka et al., 1998). This same pattern of stable recurrent

activity has been proposed to serve networks of cortical neurons as learning and storage mechanisms (Amit, 1992; Hebb, 1949; Yuste et al., 2005).

The current experiment and the previous study both revealed a stable and consistent set of display conditions that generally supported the bird's most accurate search behavior and indicated the attractor's location. These specific stimulus conditions consisted of displays with large elements having a few, preferably lighter or darker, distractors. Both studies have strongly indicated that distractor count and element size are the principal factors controlling selection in this task.

The current experiment, however, suggested a slightly different location for the attractor than the previous one. Using the accuracy surface, here we found that the bird's best search behavior happened with only one distractor and with larger elements than we had estimated before. Previously, we had found that two to three moderate to large elements might be best. We made this estimate by using the frequency shifts in gene populations from the bird's selection in that experiment, since the tracking algorithm employed made that estimate most viable. However, the greater accuracy seen with the long-term, extensive mapping in the current experiment seems the more appropriate location for the attractor.

The difference was created by the many-to-one correspondence between gene frequency and display frequency in our procedure. Because of the interactions between several genes used to spatially pack our displays, displays with several distractors were more frequent than those with one distractor. Because the birds were just about equally good with such displays, the slight accuracy advantage of having one distractor was not sufficient to dominate selection among the relative frequency and reasonable good accuracy with the descendants of displays with only two or three elements. One recommendation for researchers using GAs in the future is to strive to have the contributions of the genes be independent of one another in the construction of the displays. This will ensure equal sampling of the entire stimulus space. Under such conditions gene frequency and accuracy estimates of the attractor's location should converge. When experimental or computational requirements prevent independent gene expression, then we recommend determining the attractor's location and shape in the stimulus space based on accuracy in choice settings.

Conceptualizing solutions in this way offers new avenues for thinking about the animal's behavior. One aspect of this refers to the spatial area or extent of the attractor over the stimulus space. As determined from the accuracy surface, once elements become large enough for the bird to perform the task in our computerized setting, there is a large region or set of stimulus conditions over which the bird is reasonably good to excellent at the search task. This suggests that the bird's strategy to identifying the gray target is highly flexible and capable of solving a wide range of stimulus and contextual conditions. Further, the non-circular shape of the attractor's impact on performance over the space suggests that it is not a singularity, but has a form that asymmetrically impacts performance along different dimensions and their relative weighting. One advantage of the GA is that it strongly encourages measuring behavior over an extensive range of stimulus conditions. Being able to characterize the range of conditions over which a solution can be applied would be useful for isolating differences between solutions, tasks, individuals, and species. For instance, one could test different species with identical problems and stimulus spaces to determine if the same attractors emerge and if their spatial extent and shape are equivalent or if they vary in psychologically revealing and different ways.

Within this context, our results have several implications for Bond and Kamil's examination of prey selection in their virtual ecology setting using a GA procedure. A central difference between their experiments and the one conducted here is the nature of

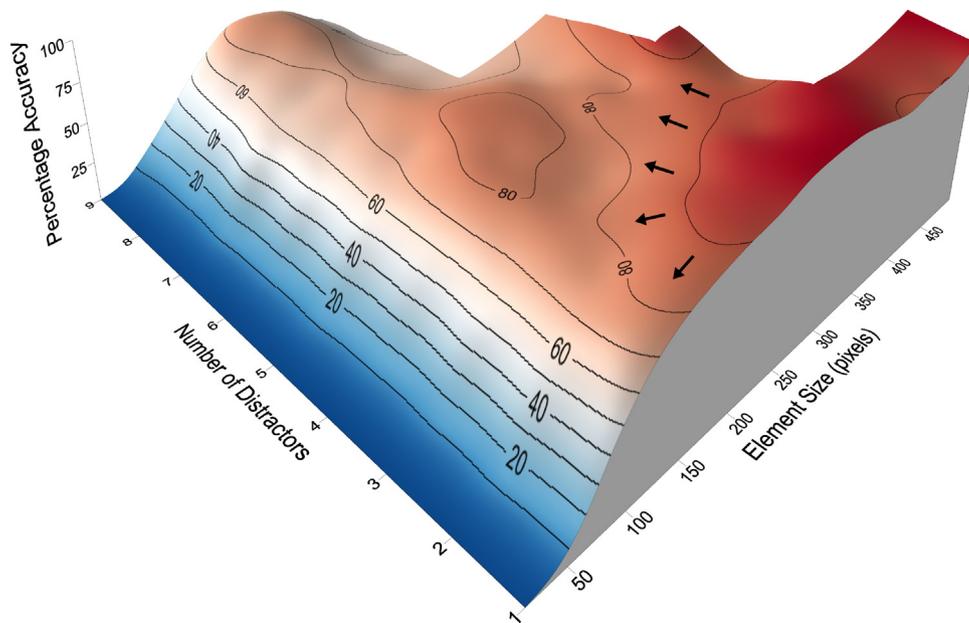


Fig. 8. Hypothetical “negative” selection pressures based on retaining inaccurate displays in the current search discrimination. The three-dimensional contour is the accuracy surface for all trials in this experiment. The arrows are hypothesized directional forces from a GA that uses non-successful reproduction at each generation.

the selection rule employed. In our experiment the pigeon’s accurate “target” selection retained stimuli that were increasingly more discriminable over generations, while the blue jays’ accurate “prey” selection in Bond and Kamil’s experiments retained stimuli that become increasingly less discriminable. The resulting differences in the pool of reproduced descendent displays contrast markedly. Understanding the complementary and contrasting relations between selecting for the “best” and “worst” of conditions is an important factor in GA techniques to be explored.

Just like our pigeon, Bond and Kamil’s jays must have had display conditions that would have best suited their search strategy for finding “moth” targets on any particular background. Because of the differences in the selection rules, our population of selected displays moved progressively towards and reflected our pigeon’s attractor, while Bond and Kamil’s procedure would have continuously pushed their populations of selected displays away from their blue jays’ attractor at any one time. As a consequence, this type of selection resulted in displays becoming more cryptic and more variable as the phenotypes of the moths moved sufficiently far enough from the central area or spatial extent of the attractor to avoid being eaten. This type of selection is illustrated in Fig. 8 for our procedure, in which presumably the retention of inaccurate trials would have been the complement of our attractor, resulting in displays pushing away from that location and likely over a wide area. Assuming there are no strategic differences in processing as items become less or more visible, the resulting phenotypic selections by blue jays presumably traces out a rough outline of the attractor, as the retained virtual moths likely reside just outside the attractor’s location and the spatial extent of its effective region (e.g., Fig. 2 of Bond and Kamil, 2006). Knowing the location and form of the “positive” attractor(s) in this case would be potentially quite revealing, as the features and sets of conditions the jays favored in their positive selection behavior are likely the ones driving the negative selection and retention of the now more distant cryptic moths. If there were asymmetries in the attractor’s shape, however, the resulting polymorphism observed in the virtual moths could change. If some parts of the feature space used by the jays are more important than others, the types of retained moths would be different and move away in different directions depending on the attractor’s extent and shape. Hence, the form of the jays’ attractor could have real

world implications in what type of moths effectively reproduce in nature. As a result, knowing what the jays are searching for or prefer to search for is an important component of the equation within a virtual ecology. Determining the fits between such positive (reproduction from selection) and negative (reproduction from survival) surfaces is crucial for understanding the nature of such selection results. Intuitively they seem like they should be complements of one another, but that may not be the case. If the nature of the positive or negative selection alters the search strategy or solution of the animals (e.g. such as in the relative weighting of different features) then one might not be predictive of the other.

A related and interesting issue concerns the nature and number of potential attractors in any discrimination. In our experiments there appears to be a single and stable attractor that extends over a large spatial extent of the stimulus space. Bond and Kamil’s results indicate, on the other hand, that multiple attractors might exist or can be created within the same task over time (Blough, 1969, 1993; Krechevsky, 1932). In their studies, the blue jays appear to shift their attention to different features of the moth stimulus space based on their encounter rate. These attentional shifts result in different “species” of moths being differentially selected as their populations ebb and flow over successive generations of differential selection (Bond and Kamil, 1998). One way to think of these attentional shifts is that the attractor or current solution to the task moves to a new location in the solution space. Such a shift causes moth species to be selected at locations that were farther away from the previous attractor’s location, but now fall within the range of the new attractor’s parameters. Again, knowing the shape and extent of the current or different attractor(s) would be helpful in understanding why different moths are retained and their characteristics (cf. Blough, 1993). One line of research for the future will be examining the flexibility of attractors, their shape, variability, spatial extent, and the number that can simultaneously exist or be maintained by an animal in any given discrimination. The role of attention in rapidly changing the shape of such discriminative solutions and its interaction with the dynamic properties of GAs would be of considerable interest to measure (cf. Nosofsky, 1986). Taking apart the feedback loop between the animal’s selection of stimuli and the impact of the new stimuli on the animal’s selection behavior will also be of considerable importance.

One challenge for the current research project is to now test more pigeons in order to provide generality to the results reported here. We conservatively started this project by looking at the behavior of a single subject. While this has been informative for our understanding, implementation, and conceptualization of the task, it limits expanding our conclusions to pigeons more broadly. In one sense, what we have found is that a single bird is repeatedly employing the same solution to the current search problem over an extended period of time, initial conditions, and different adaptive procedures. An important next step is to expand the number of birds being tested in such procedures and examine the degree, similarity, and convergence of their solutions on the phenotypic and genotypic structure of the stimulus space tested here. Nonetheless, we hope that the findings presented in this paper show the unrealized potential of using adaptive experimentation for investigating animal cognition generally and the development of theoretical conceptions for visualizing and organizing such dynamic results. Such developments will be highly useful in understanding how the perception and discriminative strategies of predators and mates have influenced the coloration and patterning of different animals.

Author notes

This research was supported by Grant #R01EY022655 from the National Eye Institute. Home page: www.pigeon.psy.tufts.edu.

References

- Amit, D.J., 1992. *Modeling Brain Function: The World of Attractor Neural Networks*. Cambridge University Press, Cambridge, UK.
- Blough, D.S., 1969. Attention shifts in a maintained discrimination. *Science* 166, 125–126.
- Blough, D.S., 1993. Reaction time drifts identify objects of attention in pigeon visual search. *Journal of Experimental Psychology: Animal Behavior Processes* 19, 107–120.
- Bond, A.B., 2007. The evolution of color polymorphism: crypticity, searching images, and apostatic selection. *Annual Review of Ecology, Evolution and Systematics* 38, 489–514.
- Bond, A.B., Kamil, A.C., 1998. Apostatic selection by blue jays produces balanced polymorphism in virtual prey. *Nature* 395, 594–596.
- Bond, A.B., Kamil, A.C., 2002. Visual predators select for crypticity and polymorphism in virtual prey. *Nature* 415, 609–613.
- Bond, A.B., Kamil, A.C., 2006. 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* 103, 3214–3219.
- Cook, R.G., Qadri, M.A., 2013. The adaptive analysis of visual cognition using genetic algorithms. *Journal of Experimental Psychology: Animal Behavior Processes* 39, 20, <http://dx.doi.org/10.1037/a0034074>.
- Corneille, O., Hugenberg, K., Potter, T., 2007. Applying the attractor field model to social cognition: perceptual discrimination is facilitated, but memory is impaired for faces displaying evaluatively congruent expressions. *Journal of Personality and Social Psychology* 93, 335.
- Hebb, D., 1949. *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York.
- Herbranson, W.T., Fremouw, T., Shimp, C.P., 1999. The randomization procedure in the study of categorization of multidimensional stimuli by pigeons. *Journal of Experimental Psychology: Animal Behavior Processes* 25, 113–135.
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America* 79, 2554–2558.
- Hu, X., Zhang, B., 2010. A Gaussian attractor network for memory and recognition with experience-dependent learning. *Neural Computation* 22, 1333–1357.
- Kamil, A.C., Bond, A.B., 2001. The evolution of virtual ecology. In: Dugatkin, L.A. (Ed.), *Model Systems in Behavioral Ecology*. Princeton University Press, Princeton, pp. 288–310.
- Kamil, A.C., Bond, A.B., 2002. Cognition as an independent variable: virtual ecology. *The Cognitive Animal. Empirical and Theoretical Perspectives on Animal Cognition*, 143–149.
- Köhler, W., 1925. *The Mentality of Apes*. Routledge and Kegan Paul, London.
- Krechevsky, I., 1932. Hypotheses in rats. *Psychological Review* 39, 516–532.
- Nosofsky, R.M., 1986. Attention, similarity, and the identification–categorization relationship. *Journal of Experimental Psychology: General* 115, 39–57.
- Smith, J.D., Ashby, F.G., Berg, M.E., Murphy, M.S., Spiering, B., Cook, R.G., Grace, R.C., 2011. Pigeons' categorization may be exclusively nonanalytic. *Psychonomic Bulletin & Review* 18, 414–421, <http://dx.doi.org/10.3758/s13423-010-0047-8>.
- Tanaka, J., Giles, M., Kremen, S., Simon, V., 1998. Mapping attractor fields in face space: the atypicality bias in face recognition. *Cognition* 68, 199–220.
- Thorndike, E.L., 1898. Animal intelligence: an experimental study of the associative processes in animals. *Psychological Monographs: General and Applied* 2.
- Tolman, E.C., Honzik, C.H., 1930. "Insight" in rats. *University of California Publications in Psychology* 4, 215–232.
- Yuste, R., MacLean, J.N., Smith, J., Lansner, A., 2005. The cortex as a central pattern generator. *Nature Reviews Neuroscience* 6, 477–483.
- Zentall, T.R., Hogan, D.E., Edwards, C.A., Hearst, E., 1980. Oddity learning in the pigeon as a function of the number of incorrect alternatives. *Journal of Experimental Psychology: Animal Behavior Processes* 6, 278–299.