Rearranging the Environment to Facilitate our Goals

Justin Yang

2021-2022

An undergraduate thesis submitted to the department of Cognitive Science at the University of California, San Diego in fulfillment of requirements for graduation with *Distinction in Cognitive Science*.

Abstract

When making a meal, experienced cooks will initially invest time arranging their workspace so that the relevant ingredients are easy to access, saving time and energy in the long run. This act of organizing—rearranging the objects in our environment to facilitate our goals—allows us to make smarter and more efficient decisions in our day-to-day lives. How do people exert control over their environment to facilitate visual search, and to what extent are people able to estimate the benefits of organization such that efficient task performance is achieved? In this thesis, we conducted a series of behavioral experiments investigating the effect of organization on conjunctive visual search, discovering that people overestimate the benefits of an organized environment relative to the cost of organizing. By examining how performance and organization preference change across experience, we find that people are able to learn from and calibrate their experiences to improve performance overall.

Keywords:

Spatial organization, visual search

Acknowledgements

I would like to thank the Cognitive Science and Psychology community for all their feedback, support and random invaluable pieces of advice you all have imparted on me in my four years at UCSD. To all the members of the Cognitive Tools Lab and Mattar Lab, I am deeply grateful for all the support that you have given me, both intellectually and morally. You all have offered me patience and friendship, and for that, I am thankful. To Marcelo Mattar, Judy Fan, and Felix Binder, as of last week we have officially reached meeting number 40! Your questions, suggestions, and input have been invaluable, and the freedom and trust you all have placed in me in pursuing this project has helped me expand my horizons and grow as a person like never before. Thank you Marta Kutas for giving me a platform to sharpen my scientific reasoning skills and for teaching everyone in the cohort how to concisely communicate the insights we have learned this past year with your guidance. Thank you, Judy Fan. Your mentorship, friendship, and constant encouragement for me to think deeply about my interests and plan for the future throughout the past three (!) years has helped unlock doors and shaped me into who I am today :)

I thank my parents Jie Yang and Suhui Ho, and my brother Michael Yang, for their unbounded support. I don't know how you all have put up with me for 22 years, but you managed to do the impossible and I am grateful for your constant support in anything I do. Finally, I dedicate this thesis and my future work to Meghan Shen. Thank you so much for putting up with my bizarre behavior and absentmindedness throughout this endeavor. You are the reason I can work so hard, and I hope the next years are just as fruitful as we continue on this journey together.

Introduction

When making a meal, experienced cooks will initially invest time arranging their workspace so that the relevant ingredients are easy to access, saving time and energy in the long run. This act of organizing—rearranging the objects in our environment to facilitate our goals—allows us to make smarter and more efficient decisions in our day-to-day lives (Kirsh & Maglio, 1994). How do people engage in these acts of spatial organization, and what problems does it solve?

Towards explaining our ability to quickly and flexibly adapt to new scenarios, a large part of decision-making research emphasizes our ability to learn about and adapt to our environment. While such abilities may be important in explaining goal-directed behavior, much overlooked is the notion that humans also adapt the environment to themselves (Kirsh, 1996).

Our approach builds on a growing literature in psychology, emphasizing the importance of the environment on an agent's behavior. For example, some have investigated how people's spatial habits compete with perceived task effort deciding whether to rearrange their environment (Zhu & Risko, 2016), while others have examined how people adapt their environments to facilitate competing task objectives (Solman & Kingstone, 2016, 2019). While these investigations give insights into how people organize to accomplish varying goals, it remains unclear to what extent people have the ability to accurately and effectively weigh the perceived benefits of an *organized* environment against the cost of *organizing*.

In this thesis, we investigate how individuals organize to facilitate visual search, examining the extent and accuracy by which people weigh the costs and benefits of organizing to exert control over their environments. In particular, we evaluate to what extent people have metacognitive access to what spatial configurations best facilitate task performance in a visual search task, and ask whether people are able to accurately engage in some cost-benefit analysis to successfully improve overall task performance.

Towards answering these questions, we conducted a pair of behavioral experiments which systematically investigated the effect of organization on conjunctive visual search using a combination of online crowdsourcing and model-based statistical analyses. Addressing prior limitations disentangling environment structure and agent familiarization, we used a between-subjects design to directly compare visual search performance between the organizers of the environment and untrained users of the organized environments.

How do individuals organize to facilitate search?

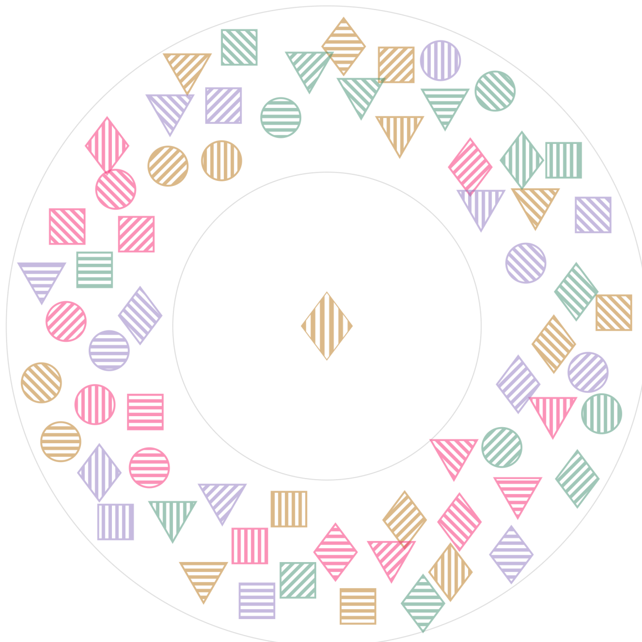

Example display of search trial

The goal of our first study was to explore the extent to which people have the ability to organize their environment to facilitate visual search, and more generally, are able to rationally weigh the costs and benefits of organizing to determine when it is worth it to organize. Towards this end, we manipulated the number of search trials a participant completed in an organized environment in a novel conjunctive visual search task where subjects were given the opportunity to rearrange the objects in a display before searching for targets in that display for the aforementioned number of trials. Insofar as people have metacognitive access to what visual displays are facilitative of visual search, we predicted that there *is* a benefit to organization — namely that participants are faster at finding objects in an organized display than in an unorganized display. Moreover, insofar as people engage in some form of cost-benefit analysis that is sensitive to the expected rewards of organization, we further predicted that participants would organize to a greater extent when searching for more trials in the organized display.

Methods

Organization task

Participants

61 english-speaking adults recruited via the university's undergraduate study pool (SONA) completed the study (43 female, 20.6 years). Each participant received one credit-hour worth of SONA course credit for their participation and were provided information consent as per our institution's IRB. We excluded all data from incomplete experimental sessions, along with any search trials with no responses when conducting reaction time analyses. In total, this yielded 610 instances of organization and 15,276 trials of visual search.

Stimuli

We algorithmically generated 16 displays in the familiarization phase and 10 displays for the main task. The displays were composed of 64 distinct objects, whose positions are randomly assigned to a lattice point on an ablated sunflower configuration such that the shapes were evenly distributed along an annulus. We used an annulus for our visual search display in order to account for the eccentricity effect — where search times for objects near the center a display are lower than search times for objects far from the center of the display in conjunctive search (Carrasco, et. al, 1995). Each object was composed of three features: color, shape, and texture. For each feature, there are four possible options. In order to approximate perceptual uniformity, the four colors were equidistantly drawn from a circle in CIE $L^*a^*b^*$ color space, centered at L = 54, a = 21.5, b = 11.5, with a radius of 49, at 0, 90, 180, and 270 degrees (Suchow et. al, 2013). The four shapes selected were circle, equilateral triangle, square, and diamond (3:4 side ratio). Finally, there were four possible textures on each object: left-diagonal hatches, right-diagonal hatches, vertical hatches, and horizontal hatches. Thus, every object in the display was unique, so that the $4^3 = 64$ possible feature combinations were all present in each display.

Design

Each experimental session consisted of ten blocks, where each block consisted of an *organization* phase and a *search* phase. In eight of the blocks (*short* blocks), the search phase consisted of 16 search trials over the organized environment. In the other two blocks (*long* blocks), the search phase consisted of 64 search trials. In total, this amounted to 256 search trials, evenly split between short and long blocks. The order of the blocks were counterbalanced, such that participants either completed the long blocks first, followed by the short blocks or the short blocks first, followed by the long blocks. The 10 randomly generated displays were randomly shuffled to each block, so that each starting display was associated with 61 organized displays varying in block index and block length.

Procedure

When beginning the experimental session, each participant was informed that they were going to play a matching game where they match an object to its unique copy on a display. The participant's goal is to match an object to its copy as quickly and accurately as possible; they could also gain points by completing a "search trial" within a certain time window in one try (i.e., that no points would be gained if they selected multiple objects). After each search trial, the participants were told how long they took, along with feedback whose positivity varied depending on how quickly they completed the trial (see supplemental). In order to emphasize search speeds, there was an additional 10 second time limit to match each object to its pair in the display. To familiarize the participant with this search task, they first completed 16 familiarization trials. Then, the participants were given an option to organize their visual display before completing a set of either 16 or 64 search trials in the organized display. They were told to take "*as much or as little time as you think will be helpful when organizing the display*" and that their goal was to minimize the overall time taken in a block of trials, which includes both the time spent organizing in addition to the time spent completing the search trials in the organized display. In total, the participants completed 10 total blocks, where after each block they were given a minute-long break and their organized displays were reset — two blocks containing 64 search trials and eight blocks containing 16 search trials. At the end of the experimental session, participants were prompted with a brief survey where they were asked to provide the following information about their gender, age, input device (e.g., mouse vs. trackpad), effort, study difficulty, strategies used, and any other comments (see supplemental).

Measuring the degree of organization

In one aspect, we aim to measure the *degree* to which a participant chooses to organize their environment. In this thesis we do so by asking two related, but separate questions: how *long* is a participant willing to spend organizing their environment, and how *different* is the organized environment from the starting display? Addressing the first question, we record the time it takes for the user to end the organization phase (time spent organizing). The amount of *change* between the initial and organized environments is defined as the sum of distances between every object's initial and final location (distance traveled).

Measuring and classifying the types of organization

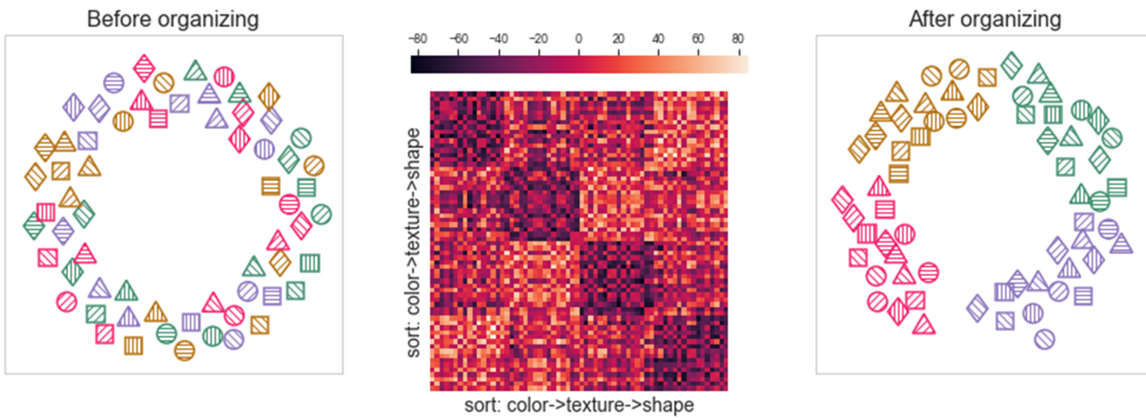

We represent organizing in terms of the change in distance of every pair of objects in the display.

We also aim to understand *what* people are doing when organizing and classify *kinds* of organization in our task insofar as different blocks of organizing can be reduced to certain kinds of actions and objectives. To do so, we must first build a principled, quantitative representation of an act of organizing, which includes relevant information (e.g. the fact that red objects got closer together should be represented) while excluding irrelevant information (e.g., the fact that two red objects are on the left- vs. right-side of the visual display is irrelevant to the task). Because we are more interested in the actual effects of organization and the organized environment, we also do not find it necessary to incorporate within-trial information about *when* in an organization phase a subject moved one object. We only care about the resulting change in the display and that display's effect on subsequent visual search.

With these factors in mind, we decided to represent the action of organization as the difference between two *Euclidean Distance Matrices* (EDMs, Dokmanic et. al, 2015). Specifically, the action of organizing is represented as a 64x64 matrix, where each element in the matrix stores the change in distance between a pair of objects in the display before and after organizing (Org figure; part A). So a negative value in this matrix represents a reduction in the distance between two objects. Being a collection of changes in distances, specific information like the *(x,y)* change in position of objects in the display are ignored.

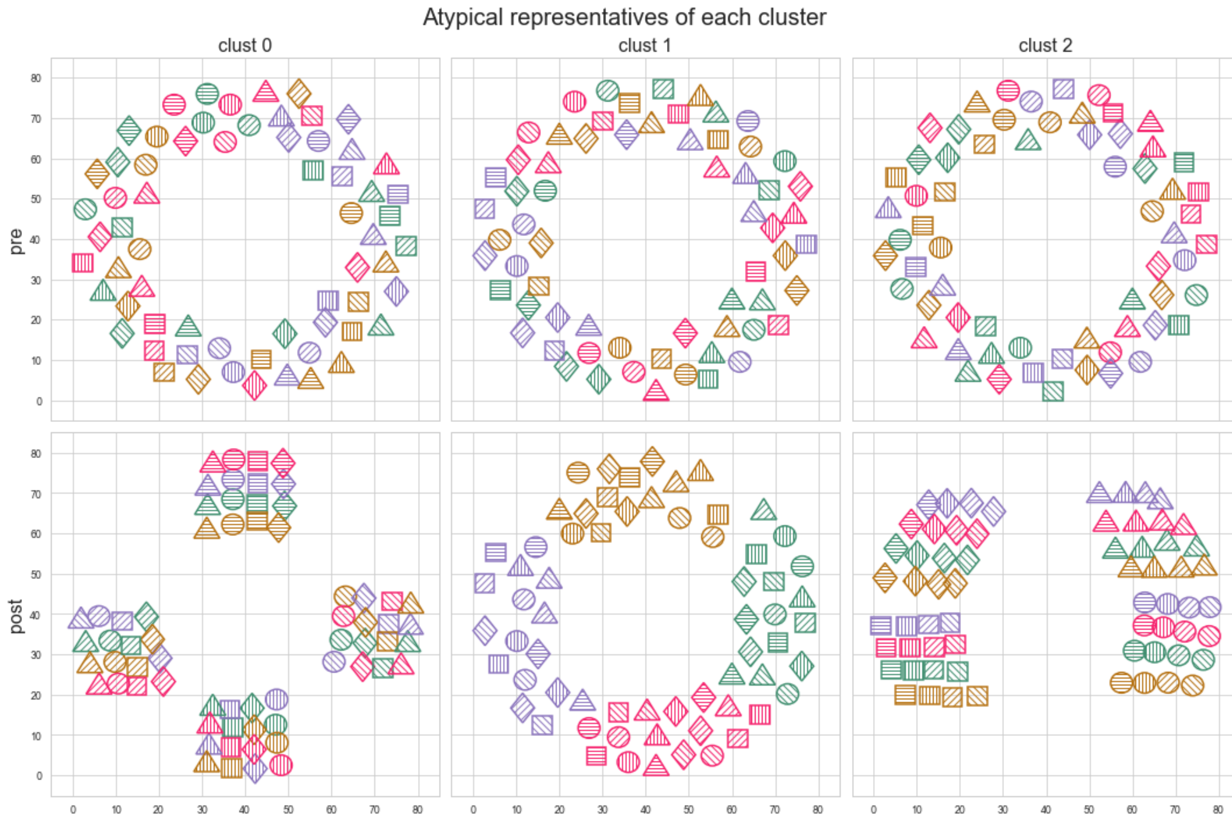

Atypical representatives of each unsupervised cluster

To gain insight into how people organized, we apply dimensionality reduction (PCA) and *k*-means clustering over the matrix representations to obtain three clusters corresponding to three types of actions: primary clustering by color, by shape, and by texture (for an interactive visualization, see [here]()). With this insight, here we use four measures to describe the extent to which people organize by clustering objects along their feature dimensions. To capture how much participants clustered along color, shape, and texture respectively, we associated each act of organization with its *reduction in intra-cluster variance* along each of the three features. So for example, the reduction in intra-cluster variance for *color* is defined by the difference in sum of within-color sum of square distances in the organized and starting configurations. So a larger value represents tighter clusters for each color. The fourth measure of clustering is agnostic to the particular feature being clustered over, and aims to capture further clustering of secondary features within each cluster: the reduction in *hierarchical* intra-cluster variance.

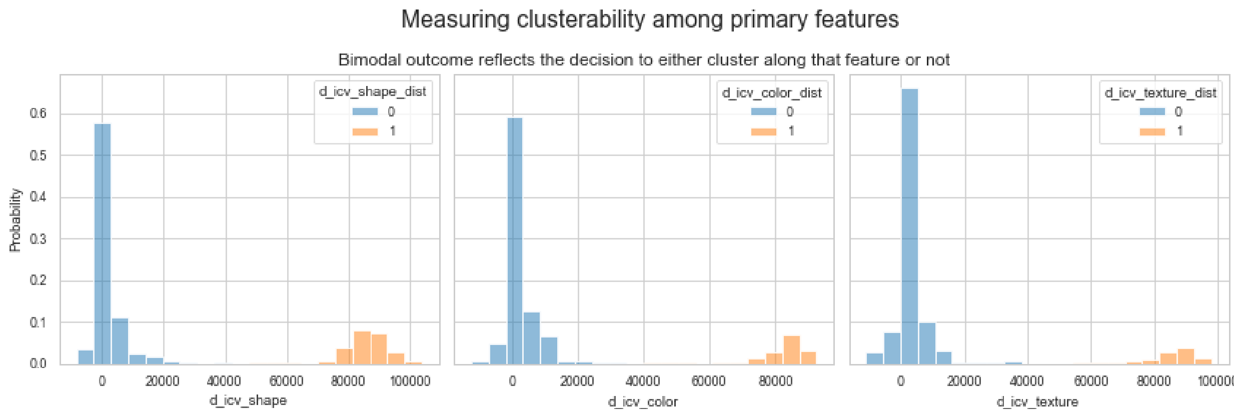

Intra-cluster variance measures are highly bimodal, suggesting that organization in this task is more of a binary condition along each feature.

We found, by visually examining the data, observing the strongly bimodal distributions of our intra-cluster variance measures, and noting the high clusterability of organization actions in PCA-space, that intermediate levels of organization along one feature (e.g., organizing two of the four possible shapes) was rare. Either participants fully organized along a feature dimension or did not organize at all. Moreover, when subjects did organize, they did so by clustering along one of three primary features; either clustering by shape, color, or texture. Because of this, we find it useful to also use a *categorical* notion of organization, which we call "organization type" for the rest of this thesis. Organizational actions which did not reduce hierarchical intra-cluster variance substantially (organizational actions in the first minor mode of the distribution) were classified as unorganized. The rest were classified by either organizing by shape, color, or texture depending on which dimension reduced the largest reduction in intra-cluster variance. We note that within each cluster, it is possible that further strategies of organizing were implemented (e.g., Gestalt similarity) however in this thesis we decide not to investigate further (Koffka, 1935).

Results

Organizing reduces trial-level search times

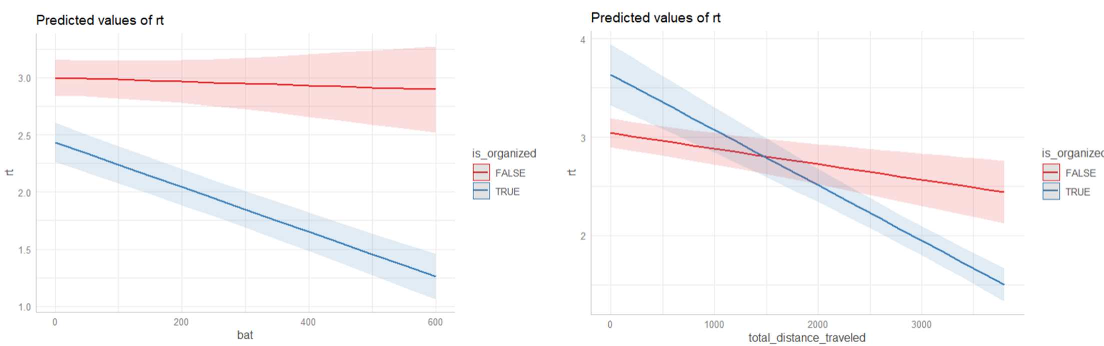

We found that as the time spent organizing (bat - block arrange time) and total distance traveled increased, trial-level search times decreased

To investigate how people leverage the benefits of organization relative to its costs, we must first determine whether organization *is* beneficial in the first place. Towards this end, we first analyzed how the amount of time spent organizing and the amount of change in an organized environment relate to how effectively subjects are able to search in that environment.

To do so, we fit two linear mixed effects models predicting search trial reaction times from (a) time spent organizing and (b) the distance traveled, with random intercepts for each participant and starting display, and nested random intercepts for each search trial within each block. We compared these to a null model using the same random effects structure, finding that models including time spent organizing (χ^2=635.94, p<2.2e-16) and distance traveled (χ^2=911.6, p<2.2e-16) outperformed the null model. So we have found positive evidence that both notions of organization influenced the speed of visual search. To control for the fact that some subjects did not find it worthwhile to organize, as well as the study's counterbalanced design and the eccentricity effect, we considered augmented models trained only on trials where subjects organized, with two covariates: counterbalanced condition and distance from the center of the screen. We found that these fuller models outperformed the model containing only organization time (χ^2=22.2, p=1.47e-5) and distance traveled (χ^2=20.5, p=3.45e-5) by a large margin, suggesting that these variables account for additional variation in search times.

By examining the full model's coefficients, we found that both greater amounts of organization time (b=-1.57e-3, t=-11.5. p<2e-16), as well as larger changes to the display (b=-3.92e-4, t=-9.82. p<2e-16) lead to faster search times (Fig.). This confirms that organizing the display facilitates visual search.

People are sensitive to the expected rewards of organization

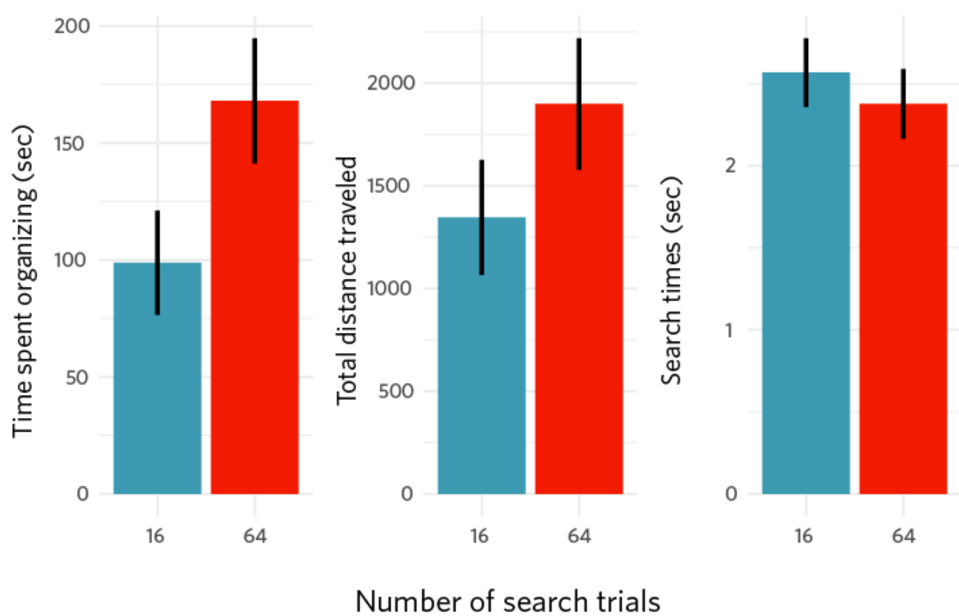

Varying across expected rewards, we see that participants spent more time organizing and modified the display to a greater extent when the expected rewards are larger. We also see that this results in a reduction in search times.

If people do weigh the costs and benefits of organizing when deciding how much to organize, then we should expect people to be comfortable organizing more when the expected rewards of organization are higher. To evaluate this possibility, we analyze how the time spent organizing and the distance traveled related to how many search trials the participants needed to complete after organizing. When participants needed to search for 64 trials, we expected them to organize more than when searching for 16 trials, since they had a greater opportunity to reap the benefits of an organized environment.

To test this hypothesis. We first fit a linear mixed effects model predicting the time spent organizing from the number of search trials in that display, with random intercepts for each participant and initial configuration. We found that this model outperformed a null model with the same random effect structure (χ^2=45.1, *p*=1.84e-11), suggesting that people are taking the expected rewards of organization into account when deciding how long to organize. Controlling for counterbalancing and task calibration (block repetition number), we found that the fuller model outperformed our minimal model using only the number of search trials to predict organization times (χ^2=52.0, *p*=2.99e-11). Examining the model coefficients, we found that indeed, when participants expected to search for 64 trials after organizing, they chose to organize longer than when they searched for only 16 trials (*b*=89.6 *seconds*, *t*=5.64. *p*=2.70e-8). Moreover, we also found that as people completed more organization blocks, they organized less and less with time (*b*=-11.3 *seconds*, *t*=-6.24. *p*=8.75e-10).

Repeating this same procedure with the total distance traveled between the unorganized and organized displays, we found a similar effect. Fitting a linear mixed effects model predicting the total distance traveled from the number of search trials in that display, with random intercepts for each participant and initial configuration was better able to explain variation in distance traveled relative to a null model (χ^2=26.5, *p*=2.69e-7). This suggests that in addition to the amount of time spent organizing, people took the expected rewards of organization into account when deciding how much to change the unorganized display as well. Controlling for the same covariates, we again found that the full model outperformed the minimal model using only the number of search trials to predict distance traveled

(χ^2=56.3, p=3.61e-12). Examining the full model showed that that participants exerted greater amounts of change on the display when the expected rewards of organization was higher (b=789, t=4.87. p=2.22e-6). We also found that subjects changed their displays less and less as they gained more experience with the task (b=-124, t=-6.56. p=1.25e-10).

Greater expected rewards of organization leads to more efficient search

Combined with our earlier result that people are able to modulate their search efficiency by organizing their environment, we then inferred that when participants expected to benefit more from organizing their environment, they *did*. Confirming this hypothesis, we fit another linear mixed effects model predicting trial-level search times from expected rewards, using random intercepts for each participant and starting display, with nested random intercepts for each search trial within each block. We found that such a model outperformed a null model with the same random effects structure, indicating that the expectation of searching for more trials after organizing leads to *some* change in search efficiency (χ^2=31.3, p=2.26e-8). Including a covariate accounting for the eccentricity effect (click distance from the center of the screen) also reliably accounted for variation in search times (χ^2=52.9, p=3.49e-13), so we use the full model to estimate effects of expected reward on reaction times. Examining the full model, we confirm that indeed, when participants expected to search for more trials after organizing, they also produced displays more facilitative of visual search (b=-0.191 *seconds*, t=-5.64. p=3.22e-8).

People calibrate their organization preferences with more task experience

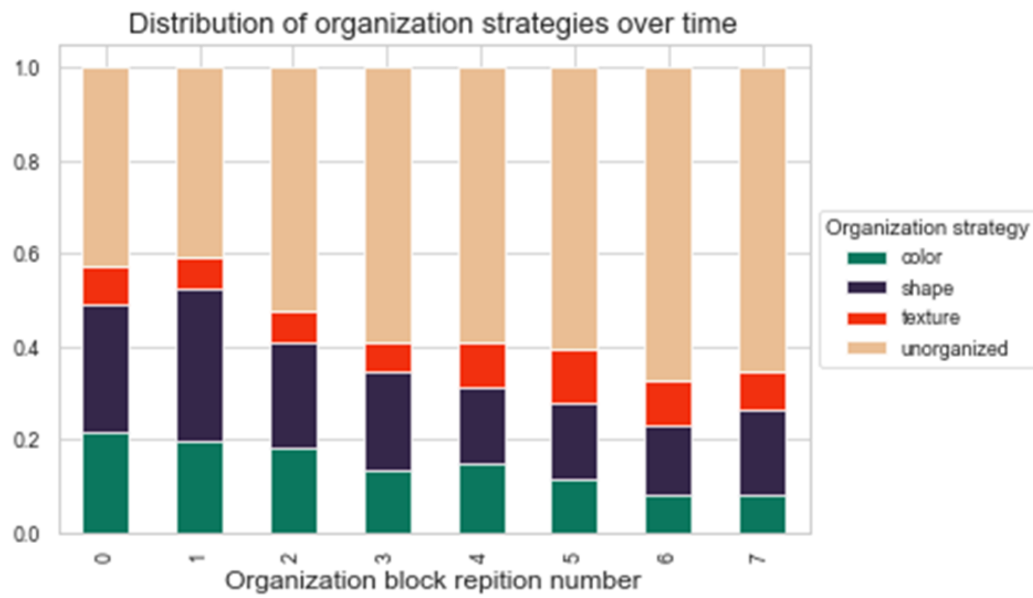

Participants organized less and less as they gained experience organizing their displays.

In our previous analysis, we found that there was a negative relationship between both the amount of time spent organizing and the total distance traveled and the amount of experience one had with the task. Why is this the case? It is possible that participants do not have a perfectly accurate representation of how beneficial it is to organize and how much time it takes to organize, and therefore must learn to better estimate both the costs and benefits of organization. For example, insofar as participants try to minimize the total time spent both searching and organizing in a block, if participants overestimate the benefits of organization and improve over time, we should expect that they organize less and less over time. This is hinted at in our previous models where both time spent organizing and distance traveled reduced as participants gained experience, and we describe this effect in detail by analyzing how the type of organization exhibited changes across trials.

We fit a logistic mixed effects regression predicting whether or not the participant organized from the within-condition (i.e., 16- vs 64- search trials) block number, number of search trials, their interaction, and counterbalancing condition, with random intercepts for each participant and initial configuration. We found that there was indeed an effect of block number compared to a null model (χ^2=54.1, p=4.93e-11). Looking at the model coefficients, we found that as participants gained experience, they organized less and less (b=-0.336, t=-5.11, p=3.25e-7). This suggests that participants had inaccurate initial representations of how beneficial it was to organize, relative to the costs. In particular, insofar as participants are calibrating their expectations to be more realistic over time, that the benefits of organization does not outweigh its costs.

Did people effectively organize to achieve their goals?

In the last study, we discovered that people are engaging in some form of cost-benefit analysis that is sensitive to the expected rewards of organization when deciding when and to what extent to organize. Moreover, we hypothesized that because people organize less as they gain more experience, that their estimates about the benefits of organizing are calibrated over time, in particular suggesting that *people have inaccurate representations over the benefits of organization*. So it remains unclear: to what extent does organization benefit one's capacity to search over that environment? Insofar as people have metacognitive access to what spatial configurations are facilitative towards conjunctive visual search, and insofar as our estimates of *how much reward* we gain from organizing in that way are accurate, we should expect people to be able to leverage organization to consistently reduce the overall time spent in a block, such that the time-cost of organization does *not* outweigh the search time benefits of searching in the organized environment. However, insofar as our estimates of the costs and benefits of organization are skewed, we may expect subjects to either *over-organize* the environment if they discount the costs of organizing and inflate expected benefits or organization, or *under-organize* the environment if they inflate the costs of organizing and discount the benefits of searching in an organized environment.

Moreover, it is reasonable to ask: how does an organized environment support visual search? Is it the case that people create perceptually salient structure or groupings that are broadly useful to anyone who accesses the environment? Or is it also possible that people organize to reflect the idiosyncratic ways in which they reason about conjunctive visual search, such that organization is a mechanism to extend an individual's mind into the world and subsequently offload computation, offering great benefit to the organizer but little benefit to others? By analogy, when an executive chef methodically lays out their spices, meats, vegetables, and cookware in preparation for a large volume of customers, does this help them rapidly find ingredients because the placement of these items just "makes sense to them"? Or is it because it accommodates their more general visual search abilities, such that another chef unfamiliar with the intentions of the executive chef is also able to efficiently act in their organized kitchen?

To dissociate the effects of intentional design and contextual cueing from the physical effects of having an organized environment on our ability to perform conjunctive search, we recruited a separate sample of participants to search over both the unorganized displays and a subset of the organized environments produced in the previous study. By doing so, we are able to (1) estimate the time saved from organizing relative to a between-subjects baseline, and also (2) quantify the extent to which the benefits of organization is restricted to the organizer (who has a mental representation of the organized environment) over an outside observer (who relies on the perceptual feature of the environment).

Methods

Search Task

Participants

362 english-speaking adults recruited via the university's undergraduate study pool (SONA) completed the study (266 female, 20.5 years). Each participant received one credit-hour worth of SONA course credit for their participation and were provided information consent as per our institution's IRB. We excluded all data from incomplete experimental sessions, along with any search trials with no responses when conducting reaction time analyses. In total, this yielded 87,605 trials of visual search, so that each display we collected search time responses for had a median of 352 valid search trial responses.

Stimuli

We used the same 16 algorithmically generated displays in the familiarization trials as in Study 1, and each participant also completed one search trial for each of the 10 initial, unorganized displays. Then, we subsetted 120 organized displays from each condition (note that there were 122 organized displays where organizers searched for 64 trials afterwards) to obtain between-subjects search times for this second study. Aside from being organized by another set of subjects, the displays were the same as in Study 1.

Design

Each experimental session used the same 10 unorganized and 240 organized displays, where the presentation order was randomly shuffled for each participant. Importantly, each display was only presented one time to each subject, unlike in Study 1 where participants were shown the organized display either 16 or 64 times in a row (see discussion for implications and future work).

Procedure

Each participant completed 16 familiarization trials (reusing the same displays as in Study 1) and 250 search trials (10 unorganized, 240 organized). The 250 search trials were broken into five blocks of 50 trials each, where on completion of each block they were given a minute-long rest period. They were given the same instructions as in Study 1, except that there was no organization phase so there was no concept of minimizing time spent in a "block" of trials. Instead, they were prompted to complete each search trial as quickly and accurately as possible. At the end of the experimental session, participants were prompted with the same exit survey as in Study 1, where they were asked to provide the following information about their gender, age, input device (e.g., mouse vs. trackpad), effort, study difficulty, strategies used, and any other comments (for more details, see the supplemental materials).

Evaluating the empirical utility of organization

To evaluate how much time the participants saved when organizing, we leveraged the between-subjects trial-level search times of the unorganized configurations. Specifically, we used the mean search times as baseline estimates of how organizers *would have* performed if they did organize. Thus, we define the total amount of time saved within a block as:

organization time saved =
\# *search trials* ∗ ((*outsider search time*) - (*organizer search time*)) – *organization time*

So if participants *saved* more time than they look at organizing, the time saved should be positive. However, if participants took more time *organizing* than they saved by searching in that organized environment, the time saved should be negative. Using this measure, we determine the extent to which participants are able to accurately engage in cost-benefit analysis to determine how much organization is worthwhile.

Results

Participants do not effectively leverage organization to improve task performance

We found that overall, participants did not save time in this task. Fitting a null mixed effects model predicting block time saved, with random intercepts for each participant and starting configuration, we found that on average participants took an extra 104 seconds when given the opportunity to organize (*b=-104 seconds,* $t=-9.08$, $p=2.33e\text{-}9$). To understand how the amount of organizing affected the block time saved, we then Included the total distance traveled when organizing as a predictor, finding that this model outperformed the null model by a large margin ($\chi^2=144.23$, $p<2.2e\text{-}16$). Inspecting the model parameters, we found that greater amounts of organization corresponded to *less* time saved in each block (b=-0.061, t=-14.1, p<2e-16). In other words, moving a single object halfway across the display reduced the time saved by a predicted -3.03 seconds.

No evidence that people save more across expected benefits of organization

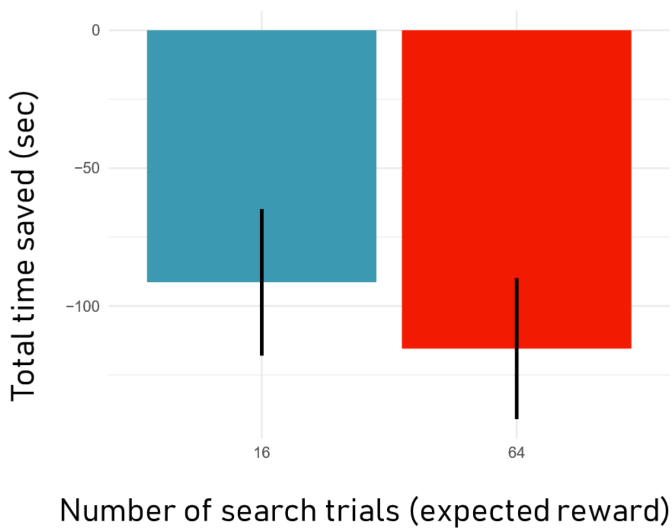

The effect of search trials after organizing did not explain substantial variation in the time saved when organizing.

From the first study, we found that participants organized their environment and successfully facilitated visual search as a result. This, however, came at a large time cost and we found that when given the opportunity to organize, participants failed to properly balance the time saved searching in an organized environment against the time spent organizing. However, it does appear that participants were sensitive to the expected benefits of organization, both organizing for a longer amount of time and also reducing the time taken to find an object in the display. How do subjects integrate these competing factors: is it the case that as the expected benefits of organization grow larger, participants inflate the benefits of organization relative to its costs to a greater extent? Or perhaps it is the case that the time spent organizing does not increase as much relative to the time savings, so that the block time saved would be reduced as the expected benefits increase? To test these possibilities, we fit a linear mixed effects model predicting the time saved in each block from the number of search trials, with random intercepts for each participant and initial display. To our surprise, we found that using the number of search trials after organizing as a predictor explained only slightly more variability in block time saved than the null model (χ^2=3.01, *p*=0.083), such that any effect of search trials on time saved is likely small. Adding two covariates — our counterbalancing measure and its interaction with the number of search trials — explained more variability in the time saved when organizing than just including the search trials (χ^2=6.90, *p*=0.032), so we use these to examine our model coefficients. We found no reliable effect of the number of search trials on the time saved in each block (*b*=9.50 *seconds*, *t*=-0.49, *p*=0.62) — either a result of our specific experimental parameters or suggesting that the additional time spent organizing and the additional time saved from organizing as the expected rewards increase is equal. For this thesis, we make no claims about this, noting that the expected rewards and counterbalancing parameters interact (*b*=-68.3 seconds, *t*=-2.49, *p*=0.0135).

Participants improve with experience

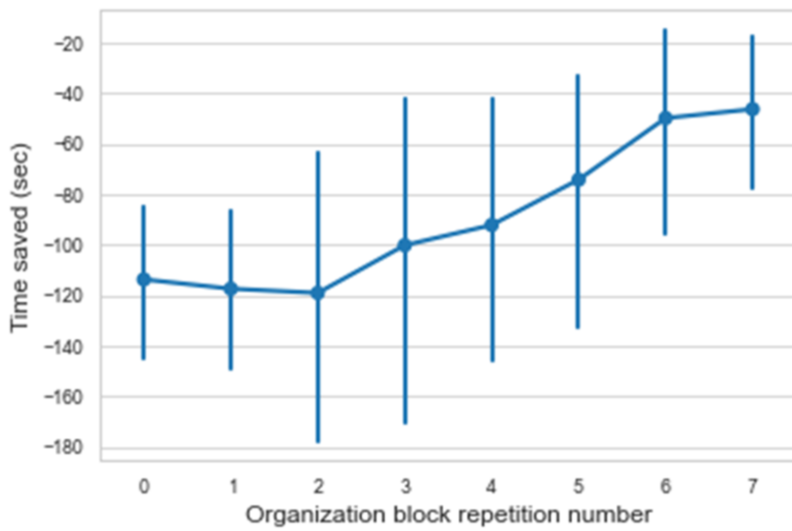

We find that with practice, participants improve on the task

In the organization task, participants were assigned to one of two counterbalanced conditions: either they encountered eight blocks of organization with 16 search trials, followed by two blocks with 64 search trials, or it was reversed. The result of this was that in each condition there was an imbalance in the amount of experience participants had in order to calibrate their expectations about how worthwhile it is to organize. What is the effect of experience as it relates to our counterbalancing, and more generally how does the ability to gradually calibrate our expectations with experience affect our ability to search? We found that as participants gained more experience, they chose to organize less and less. So we predicted that because there is no cost (nor reward) for not organizing, people gradually improved their performance over time.

To test this, we again fit a linear mixed effects regression, this time using within-condition block numbers to predict the amount of time saved in each block. We found that our model, containing predictors for within-condition block number, the number of search trials (condition), and their interaction was better fit to the data than a null model containing the same random intercepts for each participant and initial configuration (χ^2=14.5, *p*=0.002). We found a reliable effect of within-condition block number (*b*=12.9 *seconds*, *t*=3.155, *p*=0.002) showing overall that task performance improved over time. Additionally, we also found an interaction between within-condition block number and number of search trials search trials (*b*=-38.0 *seconds*, *t*=-2.017, *p*=0.045) suggesting that only when completing the eight blocks with 16 search trials did performance improve; when completing the two blocks with 64 search trials, it may be the case that participants chose not to organize less thinking that it was still beneficial to do so. Caution should be advised with this interpretation however, since our task design only allowed for two 64-search trial blocks per participant (i.e., only they were restricted to one instance of calibration, as opposed to the seven instances when there were 16 search trials per block).

Discussion

In this thesis, we investigated how people organized their environments in order to minimize the overall time spent on a conjunctive visual search task. We asked whether people were able to rearrange these displays in order to facilitate trial-by-trial visual search, and found that both spending more time organizing and modifying the display to a greater extent were reliable predictors reducing the amount of time searching in the display. We then asked whether people were sensitive to the expected benefits of organization when deciding how much to organize, and found that participants organized to a greater extent when they expected to search for 64 objects in that display rather than 16 objects, and that doing so resulted in more efficient visual search. Probing how people adapted their strategies to the task as they gained more experience, we learned that people organized less as they gained more experience, suggesting that we initially overestimate the benefits of organization relative to the costs, but can learn to calibrate these expectations with experience. We validated this hypothesis in a second study by collecting search times of both unorganized and organized displays from a separate sample of participants, learning that people failed to accurately determine the costs and benefits of organization, spending more time organizing than they saved searching in that organized environment. This changes as a function of experience, however, as we found that participants learned to more effectively leverage organizing the environment over time. Finally, despite spending more time organizing, which made the displays more facilitative to visual search, the overall inefficiency of organizing did not change substantially despite changing their behaviors to accommodate varying expected rewards.

Overall, these findings show that we reformat the way we interact with our world in a nuanced manner, dependent on both our prior experiences as well as the specific demands of the task at hand. We find a bias towards over-organization in a single-session conjunctive visual search task, perhaps suggesting a bias towards maintaining environments that are efficient in a broad set of goals, despite the fact that sometimes environments are only used for single tasks. While the present study focused on a narrow range of activities, future work addresses questions about generalization, looking at how well organized environments designed to meet the demands of one task transfer to other tasks. For example, how well environments designed for conjunctive search facilitate efficient search and vice versa.

Other directions for future work aim to address the fact that environments are organized with a designer's *intentions*, and serve as a way to offload cognition onto the world. Current exploratory analyses suggest that in order to reap the benefits of an organized environment, users of that environment must first have a functional mental representation of that environment. Comparing trial-level search times between the organizers of the environment and outside observers, we note that organizers strongly outperform outside observers seeing that display for the first time. Examining how people communicate the use of a display they organized allows us to gain detailed insights into how what is being represented in the mind is intimately linked to what is constructed in the world.

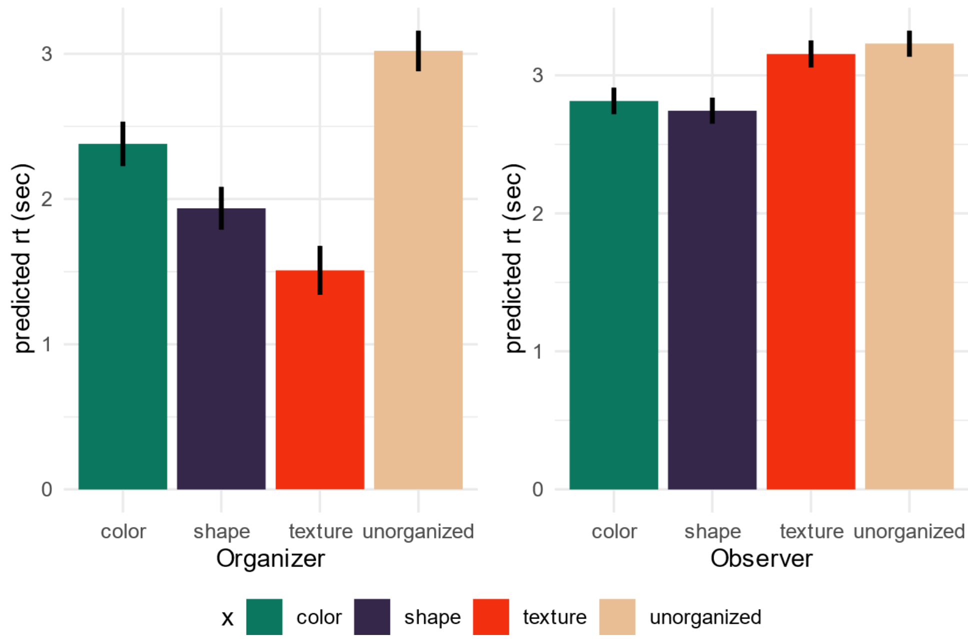

Organizers, particularly those who clustered by texture, outperformed outside observers seeing the organized display for the first time.

More broadly, the findings presented in this thesis point to an understudied mechanism by which individuals are able to overcome the limitations of time and computation. By redefining our relationship with the space around us, endowing it with rich cognitive meaning, an organized environment has the potential to extend their cognitive capacities contributing to the richness of human cognition.

Data and code availability

The anonymized data and code is available in the following GitHub repository: https://github.com/cogtoolslab/environment_structuring. We include all code used in web experiments, data analysis and stimulus generation in hopes that it may be of use to others in the future.

References

Kirsh, D. and Maglio, P. (1994), On Distinguishing Epistemic from Pragmatic Action. *Cognitive Science*, 18: 513-549. https://doi.org/10.1207/s15516709cog1804_1

Kirsh D. Adapting the Environment Instead of Oneself. *Adaptive Behavior*. 1996;4(3-4):415-452. doi:10.1177/105971239600400307

Zhu MJH, Risko EF. Spatial habit competes with effort to determine human spatial organization. *Quarterly Journal of Experimental Psychology*. 2016;69(7):1255-1264. doi:10.1080/17470218.2016.1158301

Solman, G.J.F. and Kingstone, A. (2017), Arranging Objects in Space: Measuring Task-Relevant Organizational Behaviors During Goal Pursuit. *Cogn Sci*, 41: 1042-1070. https://doi.org/10.1111/cogs.12391

Solman GJF, Kingstone A (2019) Spatial organization to facilitate action. *PLOS ONE* 14(5): e0216342. https://doi.org/10.1371/journal.pone.0216342

Carrasco, M., Evert, D.L., Chang, I. et al. The eccentricity effect: Target eccentricity affects performance on conjunction searches. *Perception & Psychophysics* 57, 1241–1261 (1995). https://doi.org/10.3758/BF03208380

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of vision*, 13(10), 9-9.

I. Dokmanic, R. Parhizkar, J. Ranieri and M. Vetterli, "Euclidean Distance Matrices: Essential theory, algorithms, and applications," in *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 12-30, Nov. 2015, doi: 10.1109/MSP.2015.2398954.

Koffka, K. (2013). *Principles of Gestalt psychology*. Routledge.

Supplemental information

Study design

Instructions

Here are the instructions given to every participant:

Search instructions

Hello! In this study, you will be playing a matching game where you are trying to match a unique object to its copy. We expect the average game to last under 60 minutes, including the time it takes to read these instructions.

By completing this study, you are participating in a study being performed by cognitive scientists at the University of California, San Diego, USA. If you have questions about this research, please contact the **Cognitive Tools Lab** at cogtoolslab.requester@gmail.com. You must be at least 18 years old to participate. There are neither specific benefits nor anticipated risks associated with participation in this study. Your participation in this research is voluntary. You may decline to answer any or all of the following questions. You may decline further participation, at any time, without adverse consequences. Your anonymity is assured; the researchers who have requested your participation will not reveal any personal information about you.

In this study, we are interested in your ability to match a specific object to its copy among a set of other objects as quickly and accurately as possible.

Specifically, you will be shown a display containing a bunch of unique objects like this, and asked to match the object in the middle to the one in the surrounding display:

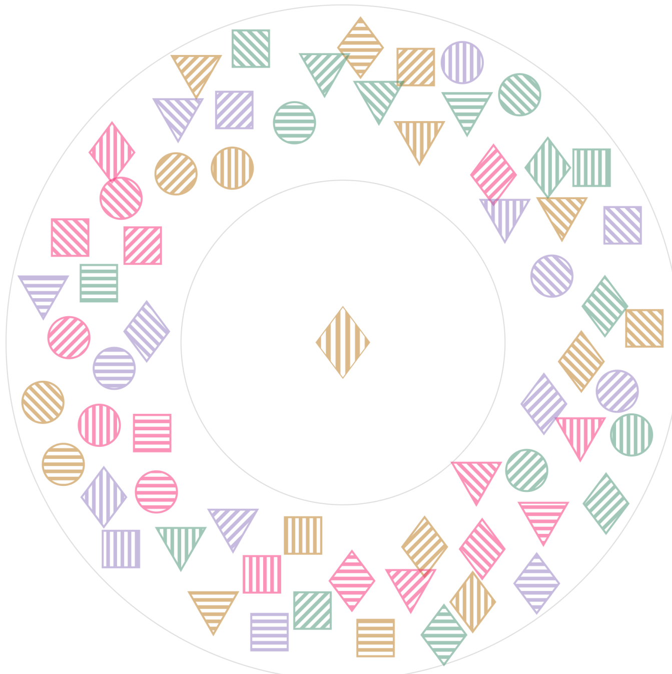

For example, you might be asked to find and click on the golden diamond with vertical stripes as quickly as possible.

Here are examples of the kinds of objects you might be asked to pair. Note that there are four possible colors, textures, and shapes:

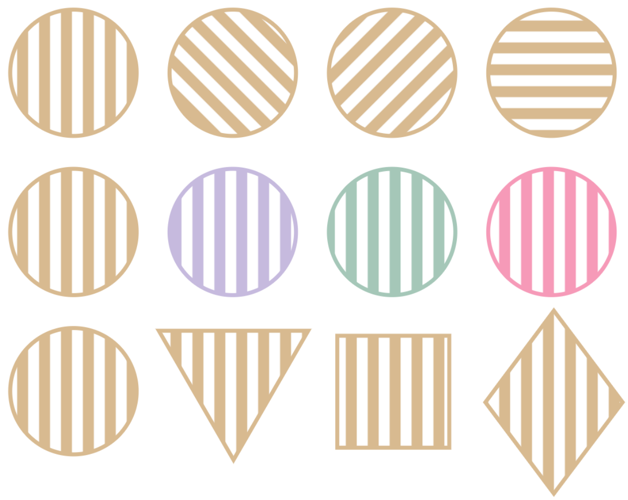

Try to match the object you are shown as fast as you can! If you match the object fast enough, you will receive one point. But please only click on the object you are asked to match, since clicking on any wrong object will result in gaining zero points that round.

If you cannot find the object you need, the trial will automatically time out after **10 seconds**.

Here is an example of what playing one round, or "search trial," looks like: https://github.com/cogtoolslab/environment_structuring/blob/main/experiments/s3_tan_optimization/assets/3_search_trial.gif

Each search trial will automatically begin once you finish the previous trial. To get used to this procedure, you will first complete 16 trials. After you complete those trials, you will be asked to rearrange the objects in a way that helps you find and select the right shape as quickly as possible, before completing a set number of search trials in that organized display.

In total, this study should take around 60 minutes. Once you are finished, the study will be automatically submitted for approval. If you encounter a problem or error, please send us an email at **cogtoolslab.requester@gmail.com** and we will make sure you're compensated for your time. Thank you for contributing to our research! Let's begin!

Organization instructions

Good job! For the rest of this study, you will be completing ten **groups** of trials. Each group of trials has two phases: an organization and a search phase. In the organization phase, you have the opportunity to organize the display in a way that helps you complete the search phase as quickly as possible. In the search phase, you will complete a set of search trials just like the ones you completed using the organized display.

Take as much or as little time as you think will be helpful when organizing the display. Your goal is to minimize the overall time taken in a **group** of trials. Keep in mind that once the search phase of each group is completed, your organized display will be reset and you will have to organize from scratch again in the next group.

Here is an example of what the organization phase looks like:
https://github.com/cogtoolslab/environment_structuring/blob/main/experiments/s3_tan_optimization/assets/4_arrangement_trial.gif

Note that in some groups of trials you will complete 16 search trials, while in others you will complete 64 trials. You can take as long or as little time organizing the display as you think is helpful, but be sure to remember that your display will reset after you finish each group!
For the first 2/8 groups, you will be completing 64/16 search trials after organizing. For the last 8/2 groups, you will complete 16/64 search trials. After each group, you will be given a one-minute break!

Thank you again for taking the time to participate in our research. When you are ready, let's begin!

Search trial feedback

After each search trial, participants were given feedback dependent on how well they did in the previous trial. The possible responses are as follows:

Participant clicks on more than one object

You took "search_time" seconds to respond.

Search time < 1.2 seconds:

Great job! 😆

You took "search_time" seconds to respond.

Search time < 1.6 seconds:

Good job! 😄

You took "search_time" seconds to respond.

Search time >= 1.6 seconds:

Nice job. 🙂

You took "search_time" seconds to respond.

Survey questions

At the end of the experimental session, participants were asked the following questions in the post-experiment survey:

- What is your gender?
    - Options:
        - Male
        - Female
        - Non-binary
        - Prefer Not to Say

- How many years old are you?
  - Placeholder: 18
- Which of the following devices did you use to complete this study?
  - Options:
    - Mouse
    - Trackpad
    - Touch Screen
    - Stylus
    - Other
- How difficult did you find this study?
  - Likert scale from 1 (Very Easy) to 5 (Very Hard)
- How much effort did you put into the game? Your response will not effect your final compensation.
  - Likert scale from 1 (Low Effort) to 5 (High Effort)
- How did you decide to play the game? (Please describe any strategies you used)
  - Placeholder:

    When deciding how long to organize, I...

    When searching for the objects, I...
- Thank you for participating in our study! Do you have any other comments or feedback to share with us about your experience?
  - Placeholder: I had a lot of fun!
- If you encountered any technical difficulties, please briefly describe the issue.
  - Placeholder: I did not encounter any technical difficulties.

Individual differences in organization strategies

We found that there were four types of displays across all participants: displays that organized by color, shape, texture, and displays that were not organized at all. Across individual participants, we can plot the proportion of blocks where they used a particular organization strategy, finding that participants largely stuck to one search strategy and stopped organizing after a certain number of search trials, as shown below:

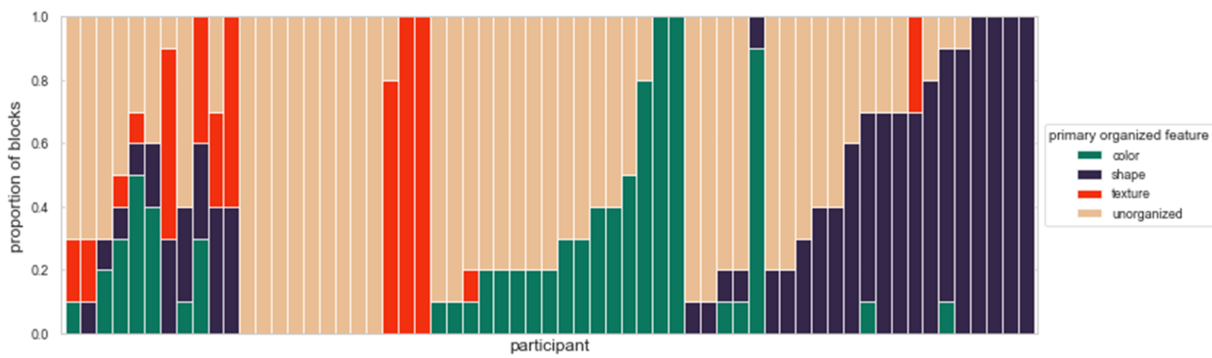

One subset of the participants — those who persevered and organized in every trial — is particularly interesting. We find that they become highly efficient in their organization strategy, and are able to outperform most other participants after learning how to efficiently organize.

Organization preference explains search times across expected rewards

What explains the fact that when participants expect to search over 64 search trials, they are able to organize their environment to better facilitate search? Exploratory analysis suggests that this is explained by different decisions over which organization strategy to use in a block:

As we can see, when subjects search over 16 trials, they largely choose to not organize, while when searching over 64 trials, they choose to organize more. So if we calculate the trial-level search time across conditions, faceted across organization strategy, we should see no difference in search times. Indeed, this is what we see: