

The Influence of Salience and Similarity on Selective Attention

by

Caitlyn Marie McColeman

B.A. (Hons.), York University, 2010

Thesis Submitted in Partial Fulfillment
of the Requirements for the Degree of
Master of Arts

in the

Department of Psychology
Faculty of Arts and Social Science

© **Caitlyn Marie McColeman 2012**

SIMON FRASER UNIVERSITY

Summer 2012

All rights reserved.

However, in accordance with the *Copyright Act of Canada*, this work may be reproduced, without authorization, under the conditions for "Fair Dealing." Therefore, limited reproduction of this work for the purposes of private study, research, criticism, review and news reporting is likely to be in accordance with the law, particularly if cited appropriately.

Approval

Name: Caitlyn Marie McColeman
Degree: Master of Arts (Psychology)
Title of Thesis: *The Influence of Salience and Similarity on Selective Attention*

Examining Committee:

Chair: Ralph Mistlburger, Professor

Mark Blair
Senior Supervisor
Associate Professor

Thomas Spalek
Internal Examiner
Associate Professor
Department of Psychology

Chung-Hye Han
External Examiner
Associate Professor, Department of Linguistics
Simon Fraser University

Date Defended/Approved: August 15, 2012

Partial Copyright Licence



The author, whose copyright is declared on the title page of this work, has granted to Simon Fraser University the right to lend this thesis, project or extended essay to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users.

The author has further granted permission to Simon Fraser University to keep or make a digital copy for use in its circulating collection (currently available to the public at the "Institutional Repository" link of the SFU Library website (www.lib.sfu.ca) at <http://summit/sfu.ca> and, without changing the content, to translate the thesis/project or extended essays, if technically possible, to any medium or format for the purpose of preservation of the digital work.

The author has further agreed that permission for multiple copying of this work for scholarly purposes may be granted by either the author or the Dean of Graduate Studies.

It is understood that copying or publication of this work for financial gain shall not be allowed without the author's written permission.

Permission for public performance, or limited permission for private scholarly use, of any multimedia materials forming part of this work, may have been granted by the author. This information may be found on the separately catalogued multimedia material and in the signed Partial Copyright Licence.

While licensing SFU to permit the above uses, the author retains copyright in the thesis, project or extended essays, including the right to change the work for subsequent purposes, including editing and publishing the work in whole or in part, and licensing other parties, as the author may desire.

The original Partial Copyright Licence attesting to these terms, and signed by this author, may be found in the original bound copy of this work, retained in the Simon Fraser University Archive.

Simon Fraser University Library
Burnaby, British Columbia, Canada

revised Fall 2011

Ethics Statement



The author, whose name appears on the title page of this work, has obtained, for the research described in this work, either:

- a. human research ethics approval from the Simon Fraser University Office of Research Ethics,

or

- b. advance approval of the animal care protocol from the University Animal Care Committee of Simon Fraser University;

or has conducted the research

- c. as a co-investigator, collaborator or research assistant in a research project approved in advance,

or

- d. as a member of a course approved in advance for minimal risk human research, by the Office of Research Ethics.

A copy of the approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for approval and letter of approval are filed with the relevant offices. Inquiries may be directed to those authorities.

Simon Fraser University Library
Burnaby, British Columbia, Canada

update Spring 2010

Abstract

A previously reported effect of similarity on reaction times in a categorization task (Hahn et al., 2010) is tested and extended by using eye trackers as a measure of visual attention. The original effect was not found using a new stimulus set, suggesting that delayed reaction times are not a result of dissimilarity but are due, in part, to the properties of the stimuli. The elements of the stimulus reflecting similarity to the training set, but irrelevant to categorization, are made salient to test this idea. The reaction time effect is not replicated, but a surprising result is found in the eye tracking data: attention to irrelevant information is *less* likely when the irrelevant information is salient. This finding cannot be explained completely by existing models, so a new way of thinking about visual attention is provided in a proposed model of integrated bottom-up and top-down visual attention processes.

Keywords: Visual attention; eye tracking; salience; categorization; similarity

Dedication

To my parents, for their patience, brilliance, guidance and support.

Acknowledgements

The work presented in this document is informed and inspired by a large number of collaborators and friends. Members of the Cognitive Science Laboratory were a great source of inspiration and assistance at every step of this project, and in my experience as a graduate student in general. Of particular note is my supervisor Mark Blair who went above and beyond the call of duty in helping me prepare a project and document that challenged me, and in general for teaching me a great deal about the art of science.

This research project is supported by the National Science and Engineering Research Council (NSERC) Strategic Project Grant, the British Columbia Knowledge Development Fund, and the Canadian Foundation for Innovation (CFI) Leaders Opportunity Fund all awarded to Dr. Mark Blair.

Table of Contents

Approval.....	ii
Partial Copyright Licence	iii
Abstract.....	iv
Dedication.....	v
Acknowledgements.....	vi
Table of Contents.....	vii
List of Figures	ix
1. Introduction	1
1.1. Processes in Categorization	1
1.2. Selective Attention and Eye Tracking	3
1.3. Summary of Experiments.....	4
2. Experiment 1: Baseline Saliency Test of Experiment Features	6
2.1. Participants	6
2.2. Materials	6
2.3. Stimuli	7
2.4. Procedure	7
2.5. Results	8
2.6. Discussion.....	10
3. Experiment 2: The Influence of Similarity in Categorization	
Performance	11
3.1. Participants	12
3.2. Materials	12
3.3. Stimuli	12
3.4. Procedure	12
3.5. Results	14
3.6. Discussion.....	18
4. Experiment 3: Saliency versus Similarity.....	19
4.1. Participants	19
4.2. Materials	20
4.3. Stimuli	20
4.4. Procedure	20
4.5. Results.....	21
4.6. Discussion.....	24
5. General Discussion	25
5.1. Problems with Traditional Accounts of Attention.....	27
5.2. Existing Work Exploring Integrated Bottom Up and Top Down Attention	27
5.2.1. Filtering and Pigeon-holing.....	27
5.2.2. Integrating the Feedforward Sweep	29
5.2.3. Extensions of the Feedforward Sweep.....	30

5.3. An Integrated Account of Visual Attention in a Goal Directed Task.....	31
5.4. The Future of the Suggested Series of Processes	36
6. Summary and Conclusion.....	37
References.....	39
Appendices.....	42
Appendix A. Feature Values	43
Appendix B. Example Stimuli.....	46
Appendix C. Feature Images	47

List of Figures

Figure 1. Experiment Procedure	8
Figure 2. Tests of Equal Salience	9
Figure 3. Experiment 2 Procedure	13
Figure 4. Experiment 2 Accuracy	15
Figure 5. Experiment 2 Reaction Time	16
Figure 6. Experiment 2 Optimization	17
Figure 7. An Example Trial in the High Salience Condition.....	21
Figure 8. Experiment 3 Reaction Time	22
Figure 9. Experiment 3 Optimization.....	23
Figure 10. The proposed processes and mechanisms underlying eye movements	33

1. Introduction

To navigate the environment, we must be able to identify regularities in our surroundings and categorize objects that we come across. Categorization is the process by which we tell the difference between friend and foe, or edible versus inedible by using informative parts of the environment to help make a decision. If there is a large, skittish mammal with hooves, a long face, large eyes, and a long tail we are likely to call it a horse, because these are all properties that are associated with our categorical understanding of horses.

During many visual categorization problems, the information that is most important to making a decision appears in conjunction with useless or distracting features. This is exemplified by airport security staff who are trained to screen luggage using X rays (Brunstein & Gonzalez, 2010). Their task is to identify dangerous items on an X ray by looking for properties that are common to prohibited items, such as the sharp edge of a knife. If such properties exist in a set of luggage, they will be present among common objects like clothing and toiletries that do not warrant additional attention and only distract from the task at hand. For airport luggage screeners and casual observers of the everyday world alike, knowing how to effectively distribute attentional resources is an important part of visual categorization (Meier & Blair, 2011).

1.1. Processes in Categorization

Understanding how humans categorize objects in the environment is important to developing a theory of human cognition, since categorization is a fundamental process in learning (Blair, Watson & Meier, 2009), memory (Grossman, Smith, Koenig, Glosser, DeVita, Moore, & McMillan, 2002) and perception (Ashby & Gott, 1988) alike. One of the long-standing debates about the processes underlying categorization is whether people use rule-based or similarity-based processes to make category decisions. The rule-

based account of categorization suggests that people rely on a set of discrete, formal criteria to sort objects into different groups; while the similarity-based account holds that people sort objects based on the overall likeness of items to members of an existing category. Empirical evidence for either account is important for informing theories and formal models, and understanding if similarity-based processing, rule-based reasoning or a hybrid of the two encompass the processes employed in a categorization task.

Hahn and colleagues (2010) present evidence in favour of the similarity-based processing account through four experiments. In three of their experiments, participants are trained on stimuli belonging to Category A, and then identify members of Category A during a transfer task. All four experiments use the same stimuli: they have six features that together create an image resembling a cartoon robot. The six features are central symbol, body shape, ear shape, hair style, an extension from the top and bottom of the figure, and decal of the robots' apparent 'limbs' (an example is available in Hahn *et al.* 2010).

For the first two experiments, three features (central symbol, ear shape and hair style) are important for determining if the stimulus is a member of Category A, and the remaining three features are irrelevant to determining category membership. The difference between these experiments is only that the second provides feedback during the transfer phase, while the first does not. In the third experiment, only one feature is used to determine the category, in that any stimulus with the 'ear shape' feature appearing as upside down triangles is a member of Category A. The fourth experiment is not a traditional categorization task, in that participants are asked only to identify the presence of three features. In Experiment 4, participants identify the same three features that were relevant to categorization in the first two experiments. During the transfer phase of all four experiments, half of the trials contain similar stimuli, while the others contain dissimilar stimuli. The difference in similarity is defined only on the irrelevant features, such that similar stimuli have more irrelevant features in common with the training set than the dissimilar stimuli. In all four experiments, they find that reaction times to dissimilar stimuli are slower than reaction times to the similar stimuli.

There are critical limitations in interpretation of the results from Hahn and others (2010) due to the design of their study and analyses. The irrelevant features in transfer

differ from the irrelevant features in the training set to create dissimilar transfer stimuli, but this might be problematic because novelty attracts attention (Daffner, Mesulam, Scinto, Cohen, Kennedy, West & Holcomb, 1998; Laurent, 2008) and the delayed reaction time to dissimilar transfer stimuli might reflect increased attention to the salient parts of the environment rather than a global effect of similarity. Another concern is in the measure reported. Due to ceiling effects in accuracy, the authors focus their analyses on the participants' reaction time. Reaction time is an opaque indicator of how participants categorize an object, since it occludes which features are fixated within a trial. Knowing which information participants choose to fixate is important in establishing the attentional processes underlying each category decision.

1.2. Selective Attention and Eye Tracking

Cognitive access to an object is gated by selective attention (Huestegge & Koch, 2011). In the visual modality, selective attention precedes an eye movement to foveate important information if it is not already in the centre of the visual field (Chelazzi, Della Libera, Sani, Santandrea, 2010; Tatler, Hayhoe, Land & Ballard, 2011). Given the tight coupling between selective attention (Itti & Koch, 2001) and eye movements (e.g. Corbetta, Akbudak, Conturo, Snyder, Ollinger, Drury, Linenweber, et al., 1998; Hoffman & Subramaniam, 1995; Liversedge & Findlay, 2000; Sheliga, Riggio & Rizzolatti, 1994; Shepherd, Findlay & Hockey, 1986), an eye tracker is used in the following experiments as an index of attentional allocation. If the earlier effect of similarity on reaction time (Hahn *et al.*, 2010) should replicate, the eye tracking measures can indicate which features of the stimulus are accessed (Rehder & Hoffman, 2005; Blair, Watson & Meier, 2009) to influence participants' slowed responses to dissimilar stimuli. Beyond measuring the goal directed eye movements of participants, an additional advantage of using eye tracking measures is that they can reflect the impact of the physical properties of the environment on eye movements and the corresponding deployment of attention.

For example, unimportant information often has properties that stand out from the environment more than the relevant sources of information do. A highway driver can recognize an example of this in bright, visually compelling billboard advertisements on the side of the road. The driver's task is to safely control the vehicle, but billboards might

stand out from the periphery and capture the attention of any of the passengers in the car, and perhaps even distract attention from the driver's primary task. Categorization problems such as the ones in the following experiments capture goal directed access to an important subset of features. As with the work from Hahn and colleagues (2010) all of the necessary information in the following experiments is available through three relevant features. The three other features provide irrelevant information, acting as distractors. If the irrelevant information exogenously captures attention due its salience relative to the rest of the environment, it is expected that the distractors will negatively impact performance and the efficiency of attention patterns.

1.3. Summary of Experiments

In three experiments, the effect of novel, irrelevant features in a category learning task is explored using measures from eye tracking data, reaction times and accuracy as outcome variables. In Experiment 1, the stimulus set is tested to ensure a good foundation for subsequent experiments. Six features are shown to participants, and response data and eye tracking data are collected to test for biases to one or more of the feature types with the goal of identifying equally salient features.

Using the six equally salient features identified in Experiment 1, Experiment 2 examines the effect of similarity on a learning task by presenting never-before-seen feature values to participants after they complete a training phase. The novel feature values appear on the irrelevant dimensions such that Experiment 2 presents the same abstract feature values used to define similar or dissimilar stimuli in Hahn and colleagues' work (2010). Provided the finding of similarity is robust to the corrections to the methodology, it is expected that reaction times should be slower to stimuli that are dissimilar to the training set. If this first hypothesis is supported by the data, then it will provide further evidence in support of the effect of similarity. If the hypothesis is incorrect, then the effect that is purportedly due to similarity might be due to a confound in the experimental design from Hahn and others' earlier work.

The final experiment uses salient, irrelevant features to explore the impact of novelty in a categorization task. Half of the participants see exactly what participants in

Experiment 2 see: a training set, and then a new set of stimuli at transfer wherein the same types of features take on new values of a pseudo-continuous dimension (Appendix C). The other half of the participants are provided with the same training set, but after training they are presented with entirely new feature types developed specifically to grab attention away from the relevant dimensions (Appendix C). In all three experiments, the relevant dimensions take on the same feature types, and for Category A they take on the same values as in the training set. These experiments are designed to correct for the limitations that were present in Hahn and colleagues' study (2010) by using warps of previously presented features instead of novel feature types, and employing eye tracking as a temporally sensitive measure of attention. Exploring highly salient distractors using this task structure is advantageous in that the distracting information is present on many trials over the learning task. The studies can then investigate both trial-by-trial measures such as reaction time or within-trial measures such as the tendency to look at irrelevant information during each trial. The primary hypothesis in Experiment Three is that there is an alternative explanation of the reaction time effect reported in the earlier study where feature novelty might have captured attention to elicit the delayed reaction time. Novelty is a salient attribute of a feature (Laurent, 2008), and so it is expected that irrelevant feature salience will elicit a slower reaction time to stimuli containing salient irrelevant features, than to stimuli containing less salient irrelevant features (Experiment 3). Data supporting this hypothesis would indicate that the source of Hahn and co-workers' reported effect of similarity on reaction time was actually due in part to feature novelty.

2. Experiment 1: Baseline Saliency Test of Experiment Features

Because this is the first time in which the proposed stimulus set will be used, it is imperative that the images be tested before proceeding to experimental manipulations. Experiment 1 tests for pre-existing biases in the features that may arise from perceptual saliency or the image complexity. If there are biases to some of the feature types then they cannot be used in subsequent experiment, but instead will be replaced with new images and which would be tested using the design described below.

2.1. Participants

Participants are 24 undergraduate students from Simon Fraser University enrolled in a lower level psychology class who receive partial course credit for their involvement in the experiment through the university's Research Participant Pool. Prior to the start of the experiment, participants have corrected, or normal-to-corrected vision and are asked to avoid the use of heavy eye make-up on the day of the experiment.

2.2. Materials

The experiment is run at Simon Fraser University on iMac computers. The computers all run Windows XP, and the experiment is created and presented using code created with ePrime 2.0. Participants' positions are calibrated to a Tobii x120 eye tracker immediately before the experiment. Participants are seated such that their eyes are 60cm from the eye tracker.

2.3. Stimuli

The stimulus set is of original design, created for Experiments 1 through 3, and has not been previously reported. The features are all naturalistic looking stimuli that vary on one pseudo-continuous dimension (Appendix C). The pseudo-continuous change is obtained using an in-house JavaScript function and Photoshop. Six values for each feature are selected (Appendix A), and will be presented randomly between subjects during the experiments. These features define categories based on an absolute conjunction rule. The values of three of the features determine the category membership, and the remaining three are irrelevant to the category.

The stimulus presented to the participant always has six features equally separated in space on a circular background image. Each feature subtends approximately 1.48° of visual angle, and are spaced by 6.2° of visual angle on the background that is 25.3° of visual angle in diameter. The backgrounds have nine possible rotations, all of which are selected at random, trial-by-trial, to increase the naturalistic context of the experiment (e.g. Appendix B). Their random presentation precludes the influence of any systematic variance in possible asymmetries in the stimulus background. The features are randomly assigned to a location, and the relevance of each location is also randomly assigned between subjects. This means that each subject will see the same relevant information in the same location throughout the experiment, but between subjects that relevant information might be anywhere in the six possible feature locations (Appendix B) and may take on any of the six possible feature forms (Appendix C).


2.4. Procedure

Participants are provided with instructions at the start of the experiment to explain how to make a response, the possible feature types and the values of particular features that members of Category A stimuli will have. The instructions follow a story about having to identify samples from far-off planets. Participants will be instructed in a manner similar to the following three proposed experiments, so that any existing feature

biases can be examined in the appropriate context. The experiment itself is a series of categorization trials.

During each trial (Figure 1), the participant sees a central fixation cross, and then the stimulus. The participant makes a self-timed response on a game pad, and then is presented with feedback. The feedback takes the form of red or green letters in the top two corners, for incorrect and correct trials respectively. Because the purpose of this experiment is to test whether the features are appropriate for the studies, the setup for training trials in Experiment 1 matches the training phase in subsequent experiments. For twelve trials, the presented stimuli (e.g. Appendix B) will be members of category A, and the only accepted response will be 'A'.

Figure 1. Experiment Procedure



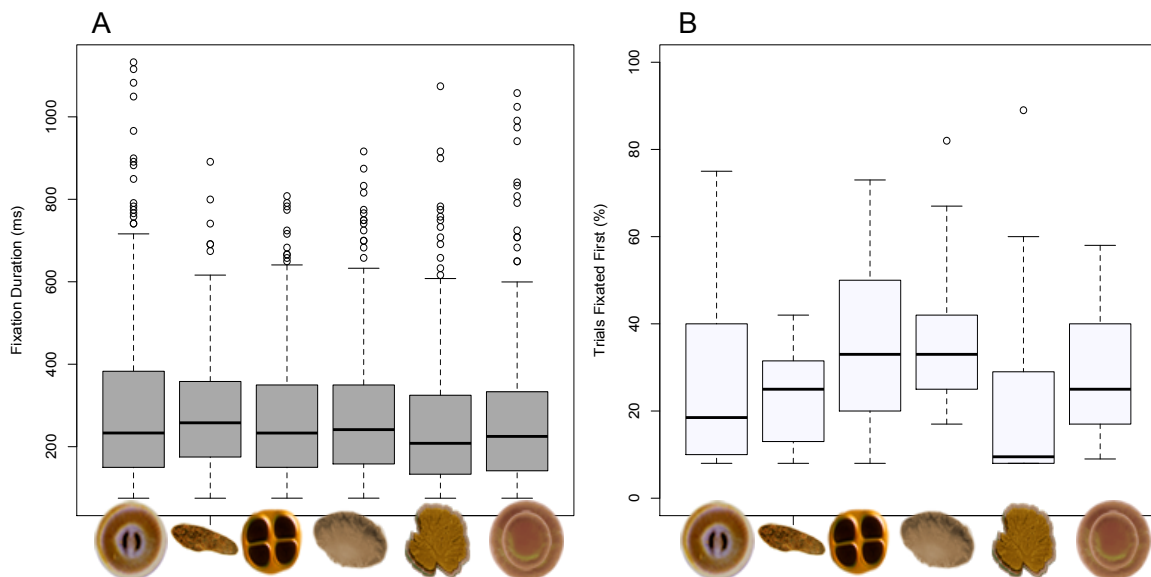
Note. An example of a trial. Each trial starts with a fixation cross, followed by a stimulus and response options. Participants make a response and then receive feedback with stimulus on screen.

2.5. Results

There are two main tests for biases to any one or more of the features. The first test to check for bias to one or more features is the fixation durations to each of the

feature types. If the features are fixated for approximately equal durations during the training phase of the experiment, then it is expected that the features are equally salient. An ANOVA was conducted to determine if one or more of the features was fixated for a different length of time than the others during the training phase of the experiments. Because the distributions of scores violated the assumption of normality, the data included in the test are a log transformation from the original observed values. The test shows that the features are fixated for approximately the same duration, $F(5,1976)=2.51$, $p=0.03$, $\eta^2=0.006$. Although the omnibus hypothesis is rejected, the effect size is very small and the pairwise comparisons between features fail to reach significance after Bonferroni correction, $p_s>0.002$. The first test of salience suggests that no one feature draws more gaze than the others. This indicates that the features are all equally salient.

Figure 2. Tests of Equal Salience



Note. Salience tests for the six low salience features. Salience is tested by average fixation duration (A) and by the probability of each feature being the first feature fixated in a trial (B).

The second test for feature bias examines the first fixation in each trial. If a feature stands out from the environment, and is more salient than the rest of the features it is expected that it should be the first feature fixated more often than the other features. This measure is simply the proportion of the trials in which each of the six features was fixated first. The participants' probability of fixating each of the six features is shown in Figure 2. As indicated by the figure, no one feature was more likely or less likely to be

fixated first, $F(5,83)=0.21$, $p=0.96$. This measure provides additional evidence that the six features are equally salient.

2.6. Discussion

Because the ANOVAs and the subsequent comparisons show that no one feature is more salient than the rest, all features are deemed equally salient at baseline. They will not introduce any biases in the following experiments due to salience or some configuration unique perceptual qualities. To further reduce the opportunity for any systematic biases, the feature images will be assigned task relevance randomly between participants. For example, the eye feature will be relevant to determining the category for some participants, and irrelevant to category membership for the other participants. The same features are distributed across different locations, and those locations are assigned different values for the tasks in the following experiments. Any analyses conducted over a number of subjects will not be influenced by a systematic bias as a function of the baseline salience of these task features. With the six stimulus features selected, the experiment extending earlier work from Hahn and colleagues (2010) can proceed.

3. Experiment 2: The Influence of Similarity in Categorization Performance

The purpose of this experiment is to test the effect of similarity reported by Hahn and colleagues (2010) where reaction times to stimuli that are different from the training set tend to be higher. A within subject manipulation of similarity is introduced in this experiment, following the experimental design of Hahn and colleagues. Similarity is defined on the features that are irrelevant to category membership, such that similar stimuli have features that are identical to the irrelevant features that were shown in the training set. Dissimilar stimuli have features that are different from the irrelevant features shown during training. Experiment 2 expands on previous work by Hahn and colleagues by using eye tracker to provide an index of attention deployed to features while participants make their category decisions.

This experiment differs from the work reported by Hahn and her collaborators (2010) in some notable ways. This study defines similarity on the irrelevant dimensions like Hahn *et al.*, but it is different in that it designs dissimilar stimuli in a way that does not invite extra attention as a function of some perceptual cueing or obvious novelty. The dissimilar features are warped versions of the same images used during an earlier training phase. There should be no perceptual orienting to these distracting features unless the effect of similarity holds true under difficult perceptual discrimination. Given the claim that their reaction time effect is a robust one elicited by dissimilar stimuli, it is expected that this experiment should act as a replication of Hahn and co-workers' findings. That is, it is expected that reaction times to dissimilar stimuli will be higher than reaction times to similar stimuli.

3.1. Participants

Participants are 79 undergraduate students from Simon Fraser University enrolled in introductory psychology classes. All participants are volunteers who receive partial course credit for their time. Participants have corrected, or normal-to-corrected vision and are asked to avoid the use of heavy eye make-up on the day of the experiment. Eight participants are excluded because they exhibited random responding behaviour. They were identified as having abnormally fast reaction times (one standard deviation from the mean) and poor performance (accuracy near chance) which indicate that they had not followed instructions. As in Experiment 1, no demographic information is collected because it is not expected to influence the outcome variables of interest.

3.2. Materials

The materials are the same as were used in Experiment 1.

3.3. Stimuli

The stimuli are the same ones used in Experiment 1 during the training phase. Additional values for each of the different feature types are created using in house JavaScript functions and Photoshop (Appendix C).

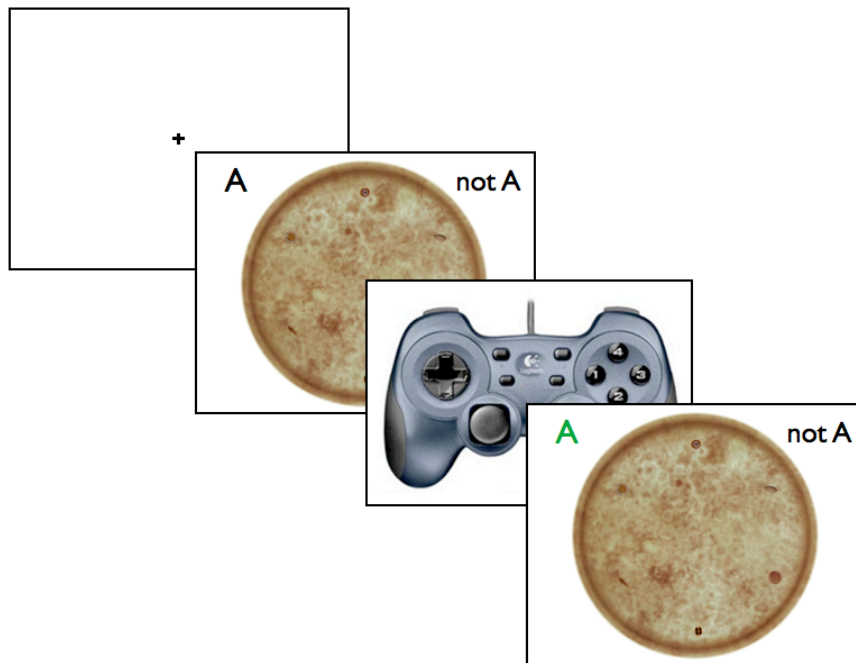
3.4. Procedure

The procedure is similar to Experiment 1 for the training phase (Figure 1). Experiment 1 is repeated for the training phase in Experiment 2. Prior to training, participants are provided with the feature values and the relevant feature types that belong to Category A. In Experiment 2, there is a supervised transfer phase, where participants will see a central fixation cross, then a stimulus, make their response, and then will see feedback to tell them if they are right or wrong and providing the correct category label when the participant is incorrect. In this phase, participants see both

Category A stimuli and non-Category A stimuli, inter-spliced in equal proportions across 192 trials.

For half of the category A trials, the stimuli are very similar to the training set: the irrelevant features take on the same value. For the other half of Category A trials, the stimuli will be very different from the training set. Two of the irrelevant dimensions take on novel values since one irrelevant dimension is the same all through training. Similarly, for the non-A trials, the stimuli are either similar to the training set, or different with the inclusion of novel features. Because the Category A is defined by the conjunction of three features, even one feature value that is inconsistent with the rule definition means that the stimulus is not Category A.

Figure 3. Experiment 2 Procedure



Note. An example of a trial during training, supervised and unsupervised transfer phases. Each trial starts with a fixation cross, followed by a stimulus and response options. Participants make a response and then receive feedback with stimulus on screen.

The supervised transfer phase is followed by an unsupervised transfer phase. The purpose of this is to reduce the information available in the learning environment and to test the participants' knowledge. The unsupervised transfer phase follows the same procedure as the supervised transfer phase. Each of the 96 trials starts with a


fixation cross, then the presentation of a stimulus, then the participant makes a response and finally the participants' response is highlighted in lieu of feedback while the stimulus is re-presented.

3.5. Results

The first set of analyses is conducted on the accuracy data. The data from the supervised and unsupervised transfer levels are tested separately since the supervised transfer phase data is from a part of the experiment where participants received feedback, while no feedback was provided at unsupervised transfer level.

The accuracy during supervised transfer was subjected to a two-way ANOVA with Block (1-6) and Similarity (high, low) as factors. The block data result in a violation of the sphericity assumption for Block, $W(2)=0.106$, $p<0.001$ and for the interaction between Similarity and Block, $W(2)=0.442$, $p<0.001$ for the ANOVA. Due to this violation, Huynh-Feldt correction is applied to the degrees of freedom of the tests involving these factors. There was no significant interaction between Stimulus Similarity and Block, $F(1,522.29)=0.332$, $p=0.894$. There was a main effect of Block, $F(1,522.29)=12.465$, $p<0.001$, $\eta^2=0.083$; but no main effect of Similarity, $F(1,5)=1.174$, $p=0.279$. The accuracy from the unsupervised transfer phase was also subjected to a two-way ANOVA with Block (1-3) and Similarity (high, low) as factors. There was no significant interaction between Stimulus Similarity and Block, $F(1,2)=0.098$, $p=0.91$. There was a main effect of Block, $F(1,2)=4.915$, $p=0.008$; but no main effect of Similarity, $F(1,2)=0.098$, $p=0.907$.

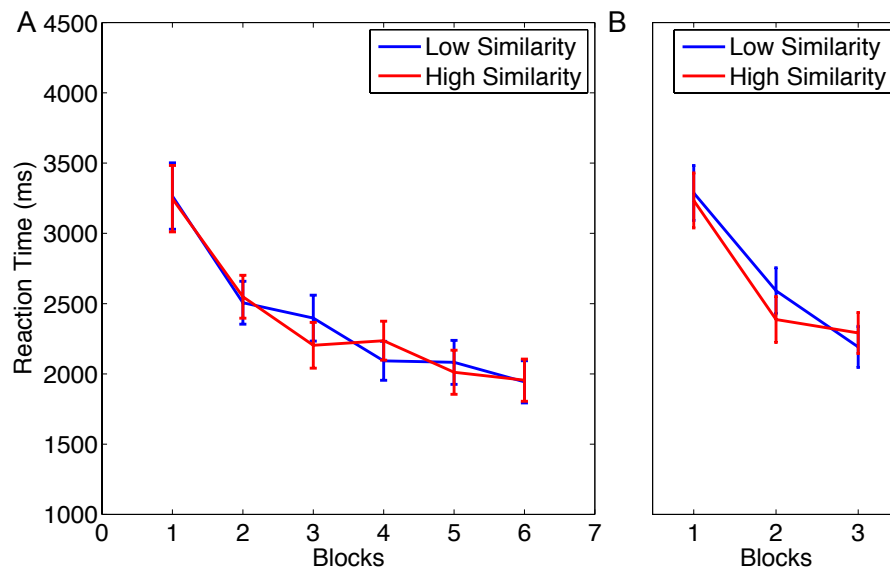
Figure 4. Experiment 2 Accuracy



Note. Accuracy during supervised transfer (A) and unsupervised transfer (B) for trials that are either similar or dissimilar from the training set. Each block represents the average accuracy across 32 trials. Error bars reflect standard error of the mean. Similarity to the training set does not influence accuracy.

This study was designed as an attempted replication of Hahn and colleagues (2010). For this reason the analyses on reaction time are based on the same trials that Hahn and coworkers did. They only included reaction times to correct trials identifying Category A, so the analyses are conducted only on these trials for the present experiment. The reaction times from the supervised transfer phase of the experiment were subjected to a two-way ANOVA with Block (1-6) and Similarity (high, low) as factors. The data result in a violation of the sphericity assumption for Block, $W(2)=0.106$, $p<0.001$ and for the interaction between Similarity and Block, $W(2)=0.442$, $p<0.001$. Due to this violation, Huynh-Feldt correction is applied to the degrees of freedom of the ANOVA tests involving these factors. There was no significant interaction between Block and Similarity, $F(3.911,269.885)=0.987$, $p=0.414$. There was a main effect of Block, $F(2.529,174.481)=19.281$, $p<0.001$, $\eta^2=0.218$. There was no main effect of Similarity, $F(1,69)=0.088$, $p=0.768$. Participants' responses tend to be the same speed regardless of the similarity of the stimulus to the training set.

Figure 5. Experiment 2 Reaction Time




Note. Reaction time during supervised transfer (A) and unsupervised transfer (B) for trials that are similar or dissimilar from the training set. The reaction times are calculated only for trials in which participants correctly identify a member of Category A. Each block represents the average reaction time across 32 trials. Error bars reflect standard error of the mean. There is no difference in reaction time as a function of similarity to the training set.

Reaction times from the unsupervised transfer phase of the experiment were also subjected to a two-way ANOVA with Block (1-6) and Similarity (high, low) as factors. As with the training phase, there was a violation in the assumption of sphericity for Block $W(2)=0.556$, $p<0.001$, and so the Huynh-Feldt correction is reported. The assumption of sphericity holds for the interaction between Block and Similarity, $W(2)=4.628$, $p=0.099$. There was no significant interaction between Block and Similarity, $F(2,134)=1.488$, $p=0.230$. There was a main effect of Block, $F(1,405,94.123)=24.922$, $p<0.001$, $\eta^2=0.315$. There was no main effect of Similarity, $F(1,67)=0.584$, $p=0.447$. The outcome of these two ANOVA tests show that this experiment fails to replicate the findings by Hahn and colleagues (2010).

The eye tracking data is used to elucidate the cognitive processes that might have been the basis of the reaction time difference originally reported by Hahn and colleagues (2010). Although this reaction time effect failed to replicate, the eye tracking data was extracted according to the original analysis plan. The data are consolidated into an optimization score that allows for comparison between trials and between

conditions (Blair, Watson & Meier, 2009). Optimization ranges from -1 to 1 and it reflects the relative amount of time spent on relevant versus the irrelevant features. Optimization is calculated by subtracting the time spent fixating irrelevant dimensions from the time spent fixating relevant dimension. The resulting difference is then normalized by dividing by the total time spent fixating any type of feature.

Figure 6. Experiment 2 Optimization



Note. Optimization scores during supervised transfer (A) and unsupervised transfer (B) for trials that are similar or dissimilar from the training set. Optimization is calculated only for trials in which participants correctly identify a member of Category A. Each block represents the average optimization across 32 trials. Error bars reflect standard error of the mean. There is no difference in optimization as a function of similarity to the training set.

The optimization scores from the supervised transfer phase of the experiment were subjected to a two-way ANOVA with Block (1-6) and Similarity (high, low) as factors. The data result in a violation of the sphericity assumption for Block, $W(2)=0.38$, $p<0.001$ and for the interaction between Similarity and Block, $W(2)=0.491$, $p<0.001$. Due to this violation, Huynh-Feldt correction is applied to the degrees of freedom of the ANOVA tests involving these factors. There was no significant interaction between Block and Similarity, $F(4.178,254.888)=1.636$, $p=0.163$. There was no main effect of Block, $F(2.077,126.708)=2.503$, $p=0.084$. Finally, there was no main effect of Similarity, $F(1,61)=0.508$, $p=0.479$. Optimization scores from the unsupervised transfer phase of the experiment were also subjected to a two-way ANOVA with Block (1-6) and Similarity

(high, low) as factors. As with the training phase, there was a violation in the assumption of sphericity for Block $W(2)=0.651$, $p<0.001$, and so the Huynh-Feldt correction is reported. The assumption of sphericity holds for the interaction between Block and Similarity. There was no significant interaction between Block and Similarity, $F(2,126)=1.804$, $p=0.169$. There was no main effect of Block, $F(1.509,95.096)=2.080$, $p=0.142$. There was no main effect of Similarity, $F(1,63)=0.014$, $p=0.907$. The results from Experiment 2 show no effect of similarity.

3.6. Discussion

There were no main effects of the similarity on accuracy (Figure 4), reaction time (Figure 5) or optimization (Figure 6). Experiment 2 failed to replicate the results presented by Hahn and colleagues (2010). Stimuli were either very similar or less similar from the training set. For both Hahn and colleagues' work and the present study, the similarity is defined independently from category membership. Both a member and a non-member of Category A can be either similar or dissimilar from the training set. The design of the present study differs from Hahn and colleagues' work only in the appearance of the stimuli. The stimuli in the present study are created so that the dissimilar trials are not defined through obvious novelty, but by a subtle shift in existing types of features. The features that inform dissimilarity vary on a pseudo-continuous dimension from the features that were present in the training set. This was to correct for a possible confound in Hahn and colleagues' work in that the new features during the trials might have drawn attention above and beyond the effect of similarity itself.

Failure to replicate the reaction time finding might suggest that the finding reported by Hahn and co-workers was confounded with stimulus properties such as novelty on dissimilar stimuli. Considering the motivation for the change in the methodology between the present study and Hahn and colleagues' work (2010), this possibility is particularly concerning. If the reaction time effect is a function of novelty and not of global dissimilarity, then the results that were reported do not actually speak to similarity, and the reaction time effect is a result of a confound in the design.

4. Experiment 3: Salience versus Similarity

The purpose of this study is to further examine the findings from Hahn *et al.* (2010) and to build on the results from Experiment 2. As with the previous experiments, similarity is defined independently from the features that inform category membership. Similarity is determined by matching a transfer stimulus set closely to an earlier training set on features that are irrelevant to category membership. As the match increases, similarity is higher; and as the irrelevant features become more different from the irrelevant features in training, similarity decreases. In Hahn and colleagues' study, dissimilar stimuli had irrelevant features that were not present in the training task. This experiment tests if their reported effect of similarity on reaction time might be due to the novelty of the irrelevant features. The experiment assigns one group of participants to a set of less salient irrelevant features at transfer; and another group of participants to a set of highly salient irrelevant features at transfer.

There are two main hypotheses in this experiment. The first is that irrelevant information will elicit a reaction time effect similar to Hahn and colleagues' report only if the irrelevant dimensions are highly salient. If this hypothesis is supported by the data, it is likely that their reported effect of similarity on reaction time is due to irrelevant dimension salience rather than by similarity. The eye tracking data serves as a more detailed measurement to explore the within trial information access that might precede slower reaction times to dissimilar stimuli. The second hypothesis is that there will be improved efficiency reflected in optimization. It is expected that optimization will be worse when distracting information is salient.

4.1. Participants

There are 74 volunteer participants in this experiment, recruited and screened for in the same manner as the above two experiments.

4.2. Materials

The materials are the same that were used in Experiments 1 and 2.

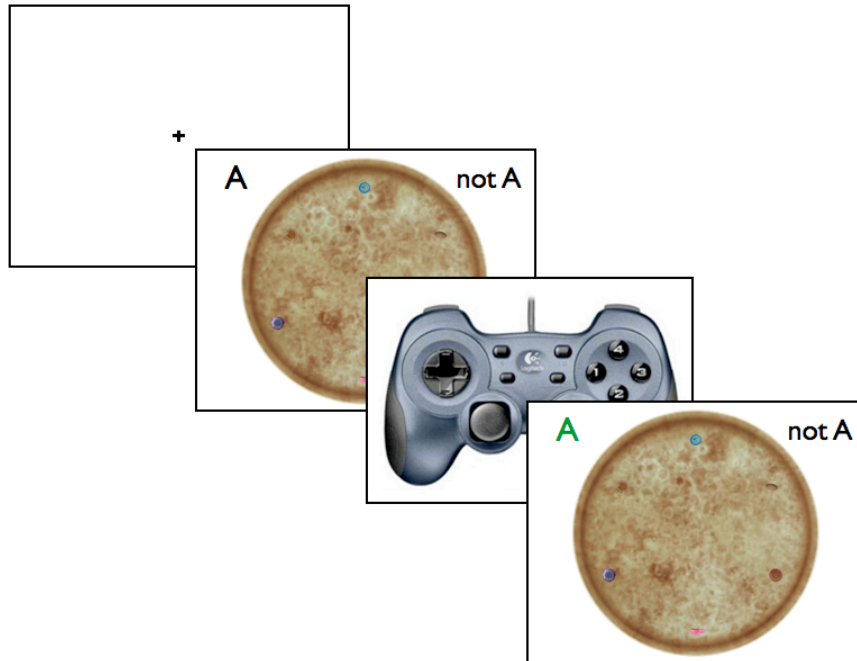
4.3. Stimuli

The stimuli are the same ones used in Experiment 1 and 2 during the training phase. For half of the participants, the transfer stimuli are the same for Experiment 2. However, for the second half of participants, there are three features that are used for the irrelevant dimensions that were not used in Experiment 2 (Appendix C). These are salient in comparison to the training set stimuli, and are changed using an in-house JavaScript and Photoshop. They are labeled here as salient irrelevant features and they are changed by rotation, width, or texture from the original training features.

4.4. Procedure

The procedure is the same as in Experiment 2, except that half of the participants will be presented with the highly salient irrelevant features rather than the less salient transfer features (Figure 7). The second half of the participants undergo the same experimental procedure as was outlined in Experiment 2. The experiment includes a training phase, a supervised transfer phase and an unsupervised transfer phase. In Experiment 3, participants are not provided with the explicit rule before training.

Figure 7. An Example Trial in the High Salience Condition




Note. An example of a trial during the supervised transfer and unsupervised transfer phases for a participant assigned to the high salience condition. Each trial starts with a fixation cross, followed by a stimulus and response options. Participants make a response and then receive feedback with stimulus on screen. The brightly coloured features are the distractors and the less salient features are the ones that are important for the category decision.

4.5. Results

As with Experiment 2, the analyses on reaction times and optimization are based on the correct identifications of Category A. The reaction times from the supervised transfer level of Experiment 3 are subjected to a two way ANOVA with Block (1-6) and Distractor Salience (high, low) as factors. There is a violation of the assumption of sphericity, and so degrees of freedom are reported with Huynh-Feldt correction on the within subjects factor. There is an interaction between Block and Distractor Salience, $F(3.519, 221.679)=2.901, p=0.028, \eta^2=0.044$. There is a main effect of Block, $F(3.519, 221.678)=37.863, p<0.001, \eta^2=0.375$. No main effect of Distractor Salience was found, $F(1, 63)=377.373, p=0.608$. Participants are equally likely to respond correctly regardless of the salience of the distractors.

Figure 8. Experiment 3 Reaction Time



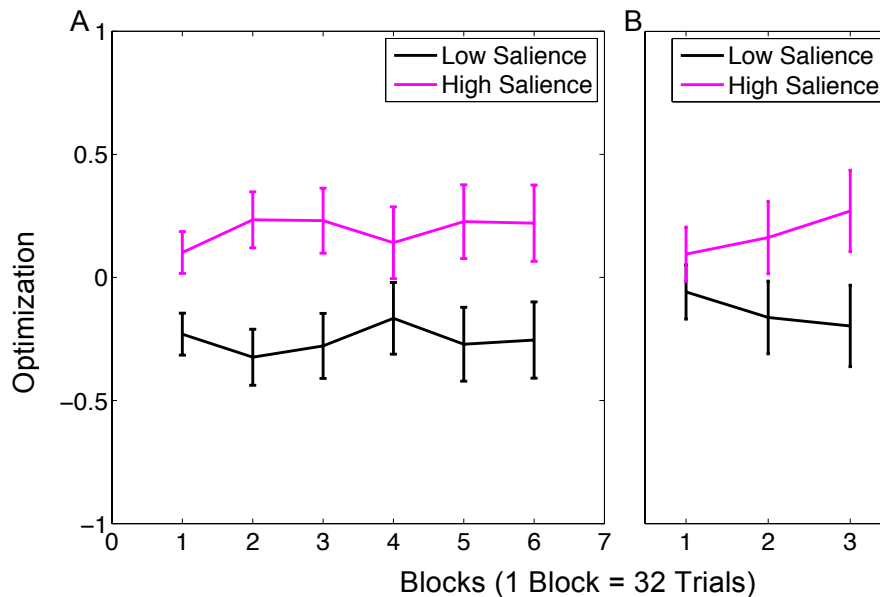
Note. Reaction times during supervised transfer (A) and unsupervised transfer (B) for participants with highly salient distractors or less salient distractors. The reaction time is calculated only for trials in which participants correctly identify a member of Category A. Each block represents the average reaction time across 32 trials. Error bars reflect standard error of the mean.

Data from the unsupervised transfer phase are also analysed. The reaction times from the transfer phase are subjected to a two way ANOVA with Block (1-3) and Distractor Salience (high, low) as factors. Due to a violation of the assumption of sphericity, $W(2)=0.781$, $p=0.025$, a Huynh-Feldt correction is applied to the degrees of freedom on Block, the within-subject factor. There is no interaction between Block and Distractor Salience, $F(1.776, 55.07)=1.996$, $p=0.150$. There is a main effect of Block, $F(1.776,55.07)=22.274$, $p<0.001$, $\eta^2=0.418$ but no main effect of Distractor Salience, $F(1,31)=1.147$, $p=0.293$. Even with the highly salient irrelevant dimensions, the reaction time effect reported by Hahn and colleagues (2010) could not be replicated.

The optimization scores from the supervised transfer phase of the experiment were subjected to a two-way ANOVA with Block (1-6) and Distractor Salience (high, low) as factors. The data result in a violation of the sphericity assumption for Block, $W(2)=0.38$, $p<0.001$. Due to this violation, Huynh-Feldt correction is applied to the degrees of freedom of the ANOVA tests involving this factor. There was no significant

interaction between Block and Distractor Saliency, $F(2.648,161.527)=1.594$, $p=0.198$. There was no main effect of Block, $F(2.648,161.527)=0.269$, $p=0.823$. However, there was a main effect of Distractor Saliency, $F(1,61)=6.846$, $p=0.011$, $\eta^2=0.101$.

Figure 9. Experiment 3 Optimization.



Note. Optimization scores during supervised transfer (A) and unsupervised transfer (B) for participants with highly salient distractors or less salient distractors. Optimization is calculated only for trials in which participants correctly identify a member of Category A. Each block represents the average optimization across 32 trials. Error bars reflect standard error of the mean. Participants with the highly salient distractors optimize more effectively than participants with less salient distractors.

Optimization scores from the unsupervised transfer phase of the experiment were also subjected to a two-way ANOVA with Block (1-6) and Similarity (high, low) as factors. As with the training phase, there was a violation in the assumption of sphericity for Block $W(2)=0.754$, $p=0.014$, and so the Huynh-Feldt correction to the degrees of freedom is reported. There was no significant interaction between Block and Distractor Saliency, $F(1.734,53.744)=2.944$, $p=0.068$. There was no main effect of Block, $F(1.734,53.744)=0.163$, $p<0.820$. During the unsupervised test phase there was no main effect of Distractor Saliency, $F(1,31)=2.600$, $p=0.117$. In contrast to the original hypothesis, the optimization is improved in the case where there are salient distractors.

4.6. Discussion

There was no main effect of salience on either reaction time or accuracy. The reaction time effect reported by Hahn and co-workers (2010) was not elicited by the salient distractors as was expected if the initial report was confounded with stimulus novelty. For this reason, it is believed that the reaction time reported by Hahn and colleagues is due to neither similarity (Experiment 2) or to novelty on the irrelevant dimensions.

There was a main effect of distractor salience on optimization. The effect is opposite what would be expected given earlier research on attentional deployment (e.g. van Rullen & Koch, 2003) in that the participants assigned to a condition with salient distractors were more efficient with their eye movements than participants with less salient distractors.

5. General Discussion

There are two general contributions of this research. The first is to show that the similarity effect on reaction time reported by Hahn and colleagues (2010) cannot be replicated using the stimulus set in Experiments 2 (Figure 5) and 3 (Figure 8). It might be that the delayed reaction time that is interpreted as being due to dissimilarity, is actually elicited by novel features that create dissimilar stimuli. These novel features in Hahn and colleagues' study might have drawn attention above and beyond what would be invoked by dissimilarity itself, but if this is the case then the reaction times in the salient condition of Experiment 3 should have been slower than the less salient condition. Another possibility is that similarity effects are only elicited by a certain spatial configuration of features. Some features occupy more space in the stimulus than the others, and this might have encouraged participants to look at the irrelevant features of the dissimilar stimuli and subsequently slow their response times. If the similarity effect is a function of feature configuration, then the report from Hahn and her collaborators should be extended to say that the similarity effects are for the different task they presented to participants using their stimuli, but not for different stimulus sets. A possible next step in this series of experiments is to explore the influence of having a single relevant feature separated in space.

In the measures reported throughout this study, participants show worse performance in beginning the unsupervised transfer portion of this experiment. This was an unexpected result in that performance should have held constant throughout the transfer task since participants were highly trained in selecting members of Category A. The performance deficit might be a function of the uncertainty that is introduced in the environment after feedback is removed. It might also be due to the small break between experiment phases where the participant is told that they would not receive feedback for the rest of the experiment.

The second main contribution of this work is the investigation of salient distractors in a categorization task. Attention in a visual categorization task is necessarily affected by stimulus properties such as salience since the information needed to perform the task enters the system through the visual modality before spatial attention is deployed. There are two main factors that contribute to the allocation of spatial attention: stimulus driven and goal directed attention (Du & Abrams, 2008). Stimulus driven attention is deployed to salient parts of an object or scene, whereas goal directed selection is aligned with the important or meaningful parts of the environment. Experiment 3 identifies a surprising ability for people to use salient distractors to increase the efficiency of their eye movements (Figure 9). This is surprising since it is expected that attention is drawn to the salient image through the orienting response (Laurent, 2008). Basic accounts of the deployment of attention would predict that salient distractors would be more distracting than less salient distractors which is consistent with an expansive literature spanning visual search (Li, 2002; Poise, Spalek & Di Lollo, 2008), scene perception (Elazary & Itti, 2008; Parkhurst, Law & Niebur, 2001; Tatler, Hayhoe, Land & Ballard, 2011), cue utilization in categorization (Kruskhe & Johansen, 1999; Nosofsky, Palmeri, & McKinley, 1994), and oddball paradigms (Suwazono, Machado & Knight, 2000).

Given that participants in Experiment 2 share the same stimuli as the participants in low salience condition of Experiment 3, it was expected that they would have exhibited similar optimization patterns. However, in Experiment 3, the low salience group optimizes worse than the participants in Experiment 2. This might reflect the difference in the instruction set between the two experiments: Experiment 2 starts with the provision of the explicit rule, while Experiment 3 does not show the category rule. The difference in optimization between the low salience group in Experiment 3 and the participants in Experiment 2 might be due to the difference in the instruction set, which is an interesting avenue for future research. Of particular interest for the current study, however, is the difference between the high and low salience groups in Experiment 3 wherein optimization is better when the distractors are salient. Although the low salience group is less effective in their optimization than the Experiment Two, it is still an effective baseline from which to investigate the effect of highly salient distractors. The finding

(Figure 9) indicates a co-ordination of top-down and bottom-up attention underlying eye movements.

5.1. Problems with Traditional Accounts of Attention

A purely bottom-up account of eye movements cannot account for the observed results. If bottom-up attention were solely responsible for eye movements, then the results would be essentially the inverse of what was reported: it would be expected that the salient distractors would draw the majority of eye movements. The observed results show that this is not the case, since eye movements to distractors are less common in the high salience group.

The opposite idea is that only top-down attention driving eye movements, and it is also problematic. If top-down attention exclusively drives eye movements (Brockmole & Henderson, 2008), then there should have been no difference between the two distractor conditions. This is because the relevant dimensions have the same qualities in both conditions, and the relevant dimensions are the only ones that should elicit top down attention. Because there are distinct eye movement patterns for each of the two groups, where the difference lies only in the distracting features, a pure account of top-down attention can be ruled out.

5.2. Existing Work Exploring Integrated Bottom Up and Top Down Attention

There are a number of theories that consider the deployment of attention as an integrated phenomenon of both top-down and bottom-up processes, and they are summarized through this section as potential descriptions of the attentional processes underlying the findings from Experiment 3 (Figure 9).

5.2.1. *Filtering and Pigeon-holing*

Work by Bundesen (1990) unifies low-level perceptual attention selection with higher level category biases. Filtering increases the probability of selecting an element

in the environment but does not increase the probability of assigning that element to any one category. Because of its separation from higher level decision processing, it is deemed to be a lower-level perceptual mechanism in this model. Bundesen also includes a pigeon holing mechanism that biases responses to a certain category. Filtering and pigeon holing work in together to produce selective attention: first the filter increases the likelihood to select a subset of elements in the environment, and then one element is selected through a winner-take-all rule. After selection, a perceptual category judgment is made on the selected element with a bias to one of the potential categories from the pigeon holing mechanism. For example, if a participant is presented with a display of red and green letters and symbols, and told that their job is to list the all of the red letters present in the display, they will first filter the display for potential candidates (the red items) by setting the pertinence values for red letters higher than the others, and set a category bias for each of the 26 candidate letters from the filtered set to more efficiently identify letters from the symbols.

The findings from Experiment 3 are consistent with the unified theory presented by Bundesen (1990). In Experiment 3, the filtering process would eliminate the salient distractors as contenders for attentional selection more easily than the less salient distractors. This might be because the brightness of the distractors is a single property of the image that can be more easily eliminated as a candidate for attention. If the difference exists only at the filtering stage, then the bias to categorization at the higher, performance level would be stable across conditions. This is consistent with the reaction times (Figure 5) and accuracy reports, in that there was no difference between the groups in either measure.

Practically speaking, this theory is advantageous because it is flexible, easily extensible, and for the purpose of this research it can be used to account for the results that would be problematic from a purely top-down or a purely bottom-up account of attentional deployment. The processes outlined by Bundesen are good descriptions in a computational framework, and a next step is to consider how plausible the filtering and the pigeon-holing mechanisms are at the neural level. It remains unclear how the information enters the visual system in order to be filtered back out with the filter mechanism in Bundesen's model. An appropriate extension of Bundesen's ideas would be to implement a way for low level visual information to build up a more complex

representation of the visual environment as the information travels to point where it can be subjected to the filtering and pigeon holing mechanisms, as is the case with the feedforward sweep. It can be mapped onto neurophysiology and has been supported by a number of empirical findings. Its role in integrated accounts of attention deployment is discussed in the following papers.

5.2.2. *Integrating the Feedforward Sweep*

Lamme and Roelfsema (2000) suggest that the feedforward sweep is not as simple as information being brought up through the visual system through simple neuronal activation. They suggest that the tuning of a neuron can change throughout a response. The neurons in some parts of the inferotemporal cortex are highly responsive to faces, meaning that they are tuned to face perception and are more likely to fire when there is a face in the visual environment (Lamme & Roelfsema, 2000). A number of these face-sensitive cells can fire differently over a response following a short delay after their initial activation. If they fire early in face perception it is in response to the presence of a face in the environment, but at a second pass they can be tuned for identity or facial expression. For this to be possible, the information must be coming back to the same area for further processing. The lateral information transfer occurs through recurrent connections. Provided that they are present throughout the visual system's hierarchy, recurrent connections are a potential source of closely integrated top-down and bottom-up processing. It might be that as information is passed up through the hierarchy, that some information is recurrently brought back into the network and processed by the same area twice before being sent through the vision hierarchy with the rest of the information.

The temporal sensitivity of the optimization measure used for Experiment 3 is too coarse to make any direct claims about recurrent connections, but the idea is brought up to present a biologically plausible account of how information in the visual field and goal directed information might work simultaneously to inform selective attention. The question remains, though, if the feedforward sweep is sufficient for the deployment of early goal-directed attention, or if some top down influence must still be considered. Van Rullen and Koch explore this issue (2003).

5.2.3. Extensions of the Feedforward Sweep

Masking experiments remove information that was presented to the visual field and replace it with a visual cover such as a checkerboard. The purpose of the mask is to make it so the original stimulus is no longer available to the observer. The advantage of the masking design is that experimenters can control how much information is introduced to the participant in both static terms (e.g. what was in an image that was presented) and temporal terms (e.g. apply the mask after 30 ms). In van Rullen and Koch's study (2003), participants were presented with a natural scene (26ms) and then a mask. The participants are instructed to report whether or not an animal was present in the natural scene.

By allowing only very brief access to the information in the visual environment, van Rullen & Koch (2003) eliminate the influence of any recurrent or feedback loops that occur after about 20ms for participants who were presented with a mask. There were two conditions in the experiment, one of which is a masked condition and the other in which there is no mask after the brief presentation of the stimulus. Importantly, there is no difference between the masked and unmasked conditions in terms of when participants are able to discriminate between targets and distractors.

This work suggests that large recurrent loops and feedback loops are not necessary for visual perception (van Rullen & Koch, 2003). A feedforward sweep can account for a great deal of perception, but the research conducted by van Rullen & Koch cannot rule out the importance of small, short latency recurrent loops. The interconnected parts of the visual perception hierarchy seem to allow for a reverberation of some parts of the visual environment for about 150ms. This reverberation of information can offset the influence of the mask prior to conscious perception, which helps to explain the relative ineffectiveness of a mask in natural scene viewing. The sufficiency of the feedforward sweep for selective attention championed by van Rullen and Koch (2003) is consistent with bottom-up account of selective attention. This is appropriate for the short time course that is of interest to them through their experiments, but top-down processing will affect the responses of participants at some point in the task. An understanding of the interaction between the two is important to understanding the demands of perception in the real world.


Theeuwes (2010) and Hochstein and Ahissar (2002) agree that the initial perception of an environment is largely bottom-up, stimulus driven processing. This idea acknowledges that top-down processing is a part of the visual selection process. By this account, the stimulus information cascades through the visual system in a feedforward manner. Only after the information has approached the top of the hierarchy will the top-down processes begin to affect visual selection. This is consistent with the order of proposed processes in that eye movements begin after selective attention identifies a saccade target. For this reason, the optimization measure and other possible eye movement data are too coarse to speak to the order of the bottom-up and top-down influences. Even so, any idea of integrated bottom-up and top-down processing is important to consider in determining how the salient distractors may help the efficiency of eye movements in the categorization task from Experiment 3, since the information from the distractors must be processed deeply enough to cause a difference in eye movements between participants with salient distractors and participants with less salient distractors.

5.3. An Integrated Account of Visual Attention in a Goal Directed Task

With the considerations from earlier research exploring the interaction between bottom-up and top-down attention in the deployment of selective attention in mind, a hybrid account wherein both bottom-up and top-down attention work together to deploy selective attention and subsequent eye movements is proposed. In short, this account suggests how participants in the salient distractor condition of Experiment 3 were able to use the salience of the distractors extracted from bottom-up attention to more quickly rule them out as potentially informative parts of the environment and ultimately deploy their eye movements more efficiently. Separating the levels of attention out into smaller sub-processes and mechanisms can help in understanding the elements of attention that yield this effect. The proposed series of processes are a novel combination of attentional mechanisms that are suggested by earlier theories. In recombining and enriching existing attentional processes from both bottom-up visual attention accounts and higher level computational accounts of attention, this idea is at odds with any purely top-down or bottom-up account of selective attention.

The visual world is perceived through a set of neural structures that are specialized for understanding information in the visual modality. The information first lands on the retina, which projects information through the optic nerve into specialized parts of the thalamus, through to visual cortical areas and up to cortical areas associated with higher cognition, which is a pattern known as the feedforward sweep (Theeuwes, 2010). As information moves through up this hierarchy, the information quantity is reduced and the quality is increased: abundant, raw visual information is traded for higher level abstractions and representations. As the bulk of the information approaches the top of the feedforward hierarchy, the visual system has developed a rich representation of the environment that can be used to inform the next steps of visual perception.

Figure 10. The proposed processes and mechanisms underlying eye movements



Note. A schematic of the proposed series of processes. The visual system takes retinotopic property maps as input, and aggregates the information on the property maps into a saliency map. Early selection (A) can tune the saliency map to a subset of features. The global expectations and biases of the system are then integrated with the tuned saliency map to create a priority map. The priority map is the basis for late selection (B) which brings the information to a higher point in the visual system where filtering occurs and a saccade target is determined. Once the target is identified, a saccade is deployed if the target is not already foveated (C) to bring it into the centre of the visual field for detailed processing (D).

Information is carried through the feedforward sweep to create property maps¹ which are topographical representations of the information in one visual channel (e.g. orientation, colour, luminosity in Itti, Koch & Neiber, 1998) determined by the activation of specially tuned neurons or neural ensembles. The property maps are aggregated to create a saliency map that provides a topography upon which higher processes will act

¹ These are historically called 'feature maps' (Koch, Itti & Neiber, 1998), but to reduce confusion with category features the term 'visual property' is used

to determine where coarse spatial attention is deployed. At the point in the feedforward sweep where the salience map is created, the raw visual information is processed to the extent that it can be without influence from higher level attention. Early selection occurs at this stage (Figure 10A). Spatial attention is tuned to salient parts of the image, but there are no saccades or oculomotor activity.

After the creation of a salience map, the global expectations are brought into the system. They can take many forms: a bias to a specific point in the visual environment, or the expectation for certain images to co-occur, for example. In Experiment 2 and 3, the global expectations include the location of the relevant information, since that is constant trial-to-trial. Global expectations re-shape the information from the salience map, transforming it from a bottom-up representation of the environment into a priority map (Fecteau & Munoz, 2006), which combines bottom-up and top-down information into a richer topographical representation of the environment. The priority map provides the information from which the relative pertinence (Bundesen, 1990) of the different parts of the environment can be extracted (Figure 10B).

The parts of the represented space with relatively low pertinence values are filtered out of contention (Bundesen, 1990); those with high pertinence values are held in visual working memory as potential targets for the deployment of a saccade, or for encoding to memory in tasks where saccades are not possible. The part of the space assigned the highest pertinence is then fixated (Figure 10C), if it is not already in a foveated position and saccades are allowed in the task, though a winner-take-all rule. The fixation brings information of the highest resolution possible into the visual system (Figure 10D). This information is taken into the system for detailed processing such as fine-grained perceptual discrimination or encoding into long term memory. At this point in the system, the action and the type of detailed processing is largely determined by task demands.

These processes are a plausible account of the observed findings of increased efficiency of eye movements in the presence highly salient distractors, which will be illustrated by an example trial. At the start of the trial, a participant fixates the centre of the screen and then presented with the stimulus. The initial visual information activates the retinal receptive fields and it begins its journey through the feedforward sweep. The

information from the retinal activation is differentiated into visual property channels through neuronal tuning (e.g. colour and brightness) to create the visual property maps. If the participant is in the low salience condition, and the distractors are similar in luminosity and colour to the relevant features, the property maps will have six moderately active locations: one for each feature. However, if the participant is in the high salience condition, the distractors differ in luminosity, and also differ from one another on the colour property maps. The salience map is an aggregate of the property maps in the system, and so for the low salience condition it will reflect six mildly salient points in space. In the high salience condition, it will aggregate three very salient points in space (the distractors) and the three non-salient areas (the relevant features). In this experiment, the advantage of early selection is to identify the spatial location of features from background noise, in a sense 'sharpening' the presence of six features regardless of the salience of the distractors. Since the experiment begins with a training session, participants are very aware of the location of the three relevant features.

At the stage where the salience map is transformed to a priority map, the locations of relevant information are biased and will be weighted more heavily than other locations. The priority map is then subjected to filtering, where the visual properties of interest to the observer create a subset of possible targets for late selective attention. In the case of the low salience group, there is obvious visual property channel to filter out of the system since the colour and luminosity of the six features are all very similar. In the high salience group, luminosity can be used to filter out the distractors since they all have that property in common. The winner-takes-all rule is applied to the filtered set of features in contention, and late selective attention is deployed to a saccade target, following by an eye movement. The low salience group has a set of six features in contention, and so there is a greater chance of accessing uninformative information. The high salience group will have only three features in contention, all of which are relevant, and will have a better chance of making an eye movement to an informative part of the environment.

5.4. The Future of the Suggested Series of Processes

The proposed integrated account of visual attention provides a possible explanation of the series of the processes underlying the results uncovered in Experiment 3 (Figure 9). Additional work to parse out the time course of the various processes can set the foundation for generalization to different tasks, such as visual search and masking; and to make the case for the appropriateness of different measures, such as electroencephalography and traditional psychometric analysis given experimenters' research questions.

Formalizing the series of processes underlying the results reported here can bring these ideas into a richer empirical realm. Falsifiable predictions derived from particular points in the series of processes may enlighten nuances that are skipped over in long standing debates surrounding visual attention. The predictions can be derived from the formalized version of the suggested integration of processes. For instance, ability of a formal version of the proposed model to fit the results reported by researchers arguing for the relative influence of bottom up (Theeuwes, 2010) and top-down attention (Ansorge, Kiss, Worschech, & Eimer, 2011) at different points in stimulus perception would help provide some basis of understanding the different results supporting two sides of a debate as an integrated body of knowledge.

6. Summary and Conclusion

There are two primary contributions of this study. The first is to report a failed replication. In Hahn and co-authors' (2010) work, there was a robust effect of similarity. The similar stimuli were the ones closely matching a training set, whereas the dissimilar stimuli differ from the training set on dimensions that are irrelevant in determining the categories to which the stimuli belong. Their design is used in the current study, but with a different set of stimuli comprised of equally sized, equally spaced features. Experiment 2 was a direct attempt at replicating the findings while controlling for the irrelevant dimensions' salience and size. The reaction time effect reported by Hahn and colleagues was not elicited by dissimilarity in Experiment 2. Given that a primary difference between Experiment 2 in the current study and the experiments from Hahn and colleagues was the novelty of the irrelevant features for the dissimilar stimuli, Experiment 3 implements novel irrelevant features in an attempt to elicit the reaction time effect reported in the earlier research. As with Experiment 2, Experiment 3 failed to yield an increased reaction time to dissimilar stimuli. However, there was an interesting finding in the eye movement data from the participants assigned to the novel, salient distractor condition in Experiment 3 in that participants with salient distractors showed improved efficiency in their eye movements.

The second main contribution of this research is that salient distractors improve the efficiency of attentional deployment in a categorization task. This finding is not predicted by traditional accounts of attention, in that neither a purely top-down or a purely bottom-up model can explain these results. This finding will be extended and explored in future research to elucidate the relative influence of top-down and bottom-up attentional processes and their relative influence on goal directed task performance. Namely, the relative influence of early stimulus driven attention and coarser, goal directed expectations can be teased apart to gain an understanding of the importance of each type of attention in everyday perceptual tasks.

Beyond basic perception, this research illustrates the importance of considering how information comes into the cognitive architecture to inform category decisions. The value of context, prior knowledge, physical stimulus properties, the number of features and the configurations of the features with respect to one another might all influence how participants perceive the stimuli that they are to categorize. Understanding the influence of these factors in categorization is advantageous in extending the interpretation of laboratory results to real world category choices, where sources of distracting information often outnumber the features determining the category membership of an object. Despite the massive amount of task irrelevant information in the real world, we perform a remarkable number of category choices at an impressive level of performance. This research serves as a preliminary step in exploring how our category decisions are affected by the visual properties of the environment and the object in question.

References

- Ansorge, U., Kiss, M., Worschech, F., & Eimer, M. (2011). The initial stage of visual selection is controlled by top-down task set: new ERP evidence. *Attention, perception & psychophysics*, *73*(1), 113-22.
- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of experimental psychology. Learning, memory, and cognition*, *14*(1), 33-53.
- Becker, S. I., & Horstmann, G. (2011). Novelty and saliency in attentional capture by unannounced motion singletons. *Acta psychologica*, *136*(3), 290-9.
- Blair, M. R., Watson, M. R., & Meier, K. M. (2009). Errors, efficiency, and the interplay between attention and category learning. *Cognition*, *112*(2), 330-6.
- Blair, M. R., Watson, M. R., Walshe, R. C., & Maj, F. (2009). Extremely selective attention: eye tracking studies of the dynamic allocation of attention to stimulus features in categorization. *Journal of experimental psychology. Learning, memory, and cognition*, *35*(5), 1196-206.
- Brockmole, J. R., & Henderson, J. M. (2008). Prioritizing new objects for eye fixation in real-world scenes : Effects of object – scene consistency. *Visual cognition*, *16*(2), 37-41.
- Bundesen, C. (1990). A theory of visual attention. *Psychological review*, *97*(4), 523-47.
- Chelazzi, L., Della Libera, C., Sani, I., & Santandrea, E. (2011). Neural basis of visual selective attention. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(4), 392-407.
- Corbetta, M., Akbudak, E., Conturo, T. E., Snyder, A. Z., Ollinger, J. M., Drury, H. A., Linenweber, M. R., et al. (1998). A Common Network of Functional Areas for Attention and Eye Movements. *Neuron*, *21*, 761-773.
- Daffner, K. R., Mesulam, M. M., Scinto, L. F., Cohen, L. G., Kennedy, B. P., West, W. C., & Holcomb, P. J. (1998). Regulation of attention to novel stimuli by frontal lobes: an event-related potential study. *Neuroreport*, *9*(5), 787-91.
- Du, F., & Abrams, R. A. (2008). Synergy of stimulus-driven salience and goal-directed prioritization: Evidence from the spatial blink. *Attention, Perception, & Psychophysics*, *70*(8), 1489-1503.
- Elazary, L., & Itti, L. (2008). Interesting objects are visually salient. *Journal of Vision*, *8*(3), 1-15.
- Fecteau, J. H., & Munoz, D. P. (2006). Saliency, relevance, and firing: a priority map for target selection. *Trends in Cognitive Science*, *10*(8), 382-90.

- Grossman, M., Smith, E. E., Koenig, P., Glosser, G., DeVita, C., Moore, P., & McMillan, C. (2002). The Neural Basis for Categorization in Semantic Memory. *NeuroImage*, *17*(3), 1549-1561.
- Hahn, U., Prat-Sala, M., Pothos, E. M., & Brumby, D. P. (2010). Exemplar similarity and rule application. *Cognition*, *114*(1), 1-18.
- Hochstein, S., & Ahissar, M. (2002). View from the Top : Hierarchies and Reverse Hierarchies in the Visual System. *Neuron*, *36*(3), 791-804.
- Hoffman, J. E., & Subramaniam, B. (1995). The role of visual attention in saccadic eye movements. *Perception and Psychophysics*, *57*(6), 787-795.
- Huestegge, L., & Koch, I. (2011). Eye movements as a gatekeeper for memorization: evidence for the persistence of attentional sets in visual memory search. *Psychological Research*, *18*, 544-49.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *20*(11), 1254-1259.
- Kruschke, J. K., & Johansen, M. K. (1999). A model of probabilistic category learning. *Journal of experimental psychology. Learning, memory, and cognition*, *25*(5), 1083-119.
- Laurent, P. A. (2008). The emergence of saliency and novelty responses from Reinforcement Learning principles. *Neural Networks : The Official Journal of the International Neural Network Society*, *21*(10), 1493-9.
- Li, Z. (2002). A saliency map in visual cortex. *Trends in Cognitive Sciences*, *6*(1), 9–16.
- Liversedge, S. P., & Findlay, J. M. (2000). Saccadic eye movements and cognition. *Trends in Cognitive Science*, *4*(1), 6-14.
- Meier, K.M., & Blair, M.B. (2011). Beyond probability gain: Information access strategies in category learning. *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological review*, *101*(1), 53-79.
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, *42*, 107-123.
- Poiese, P., Spalek, T. M., & Di Lollo, V. (2008). Attentional capture by a salient distractor in visual search: the effect of target-distractor similarity. *Canadian journal of experimental*, *62*(4), 233-6.
- Rehder, B., & Hoffman, A. B. (2005). Eye tracking and selective attention in category learning. *Cognitive psychology*, *51*(1), 1-41.
- Sheliga, M., Riggio, L., & Rizzolatti, G. (1994). Orienting of attention and eye movements. *Experimental Brain Research*, *98*, 507-522.
- Shepherd, M., Findlay, J. M., & Hockey, R. (1986). The relationship between eye movements and spatial attention. *The Quarterly Journal of Experimental Psychology*, *38*, 37-41.

- Suwazono, S., Machado, L., & Knight, R. (2000). Predictive value of novel stimuli modifies visual event-related potentials and behavior. *Clinical Neurophysiology*, 111, 29-39.
- Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision : Reinterpreting salience. *Journal of Vision*, 11(5), 1-23.
- Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta psychologica*, 135(2), 77-99.
- VanRullen, R., & Koch, C. (2003). Visual selective behavior can be triggered by a feed-forward process. *Journal of cognitive neuroscience*, 15

Appendices

Appendix A.

Feature Values

Feature	Value
1	111122223333
2	444444444444
3	111111111111
4	333333333333
5	123412341234
6	555555555555

Feature values for the training stimuli in Experiment 1-3. Each column represents the feature values for one trial for entire stimulus. Stimuli selected randomly without replacement from the 12 possible columns. For this, and all subsequent feature value tables, the bold features are relevant to category membership.

Category A – High Similarity

Feature	Value
1	111111112222222233333333
2	4444444444444444444444
3	234562345623456234562345
4	3333333333333333333333
5	123412341234123412341234
6	5555555555555555555555

Feature values for members of Category A that are more similar to the training set in Experiments 2 and 3. Each column represents the feature values for one trial entire stimulus. Stimuli are selected randomly without replacement from the 24 possible columns. Each stimulus is sampled twice during supervised transfer and once during unsupervised transfer.

Category A – Low Similarity

Feature	Value
1	444444445555555566666666
2	4444444444444444444444
3	223344562234556623445566
4	3333333333333333333333
5	123412341234123412341234
6	5555555555555555555555

Feature values for members of Category A that are less similar to the training set in Experiments 2 and 3. Each column represents the feature values for one trial for the entire stimulus. Stimuli are selected randomly without replacement from the 24 possible columns. Each stimulus is sampled twice during supervised transfer and once during unsupervised transfer.

Non A Stimuli – High Similarity

Feature	Value
1	111122223333111122223333
2	444444441122444411224444
3	1111111111111111111111
4	333311223333112233333333
5	123412341234123412341234
6	223355555555555555552233

Feature values for non A stimuli that are more similar to the training set in Experiments 2 and 3. Each column represents the feature values for one trial for the entire stimulus. Stimuli are selected randomly without replacement from the 24 possible columns. Each stimulus is sampled twice during supervised transfer and once during unsupervised transfer.


Non A Stimuli – Low Similarity

Feature	Value
1	445566445566445566445566
2	444412441244214444444421
3	1111111111111111111111
4	331233123333333321332133
5	5656565656565656565656
6	2355555552355325532555


Feature values for non A stimuli that are less similar to the training set in Experiments 2 and 3. Each column represents the feature values for one trial for the entire stimulus. Stimuli are selected randomly without replacement from the 24 possible columns. Each stimulus is sampled twice during supervised transfer and once during unsupervised transfer.

Appendix B.

Example Stimuli



An example stimulus with low salience distractors.




An example stimulus with highly salient distractors.

Appendix C.

Feature Images



All possible values for the relevant features, and the less salient distractors. The three relevant features are selected randomly from one of the six feature types for each subject (Experiment 1-3).



All possible values for the salient distractors (Experiment 3)