

Normative Representation of Objects: Evidence for an Ecological Bias in Object Perception and Memory

Talia Konkle (tkonkle@mit.edu)
Aude Oliva (oliva@mit.edu)
Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139

Abstract

An apple can be held in your hand if you are eating it, or be perceived in a fruit bowl from a few meters away. While the apple's real-world size is constant, the visual angle it subtends in your visual field varies based on your distance to the apple. Given the range of visual angles that an object can subtend in the visual field, is there a visual angle that is preferred? The current experiments show evidence supporting the idea of a privileged visual size for object representation, termed the *normative size*. In Experiment 1, observers adjusted the visual size of objects on a monitor, selecting the subjectively 'right' size to see the object. Inter-observer selections were strikingly consistent, and showed a correlation between the size selected on the screen and the actual physical size of the objects in the real world. These size selections were taken as estimates of the normative size for each object. In Experiments 2 and 3, using a size memory task and a change detection task, we found evidence that both long-term and short-term memory for the visual size of objects is biased towards the normative size. Altogether the results support the claims that perception of objects is sensitive to a normative size and that object memory is biased toward this perceptual norm.

Keywords: Object representation; perception; memory; size.

Introduction

"For each object, as for each picture in an art gallery, there is an optimum distance from which it requires to be seen, a direction viewed from which it vouchsafes most of itself: at a shorter or greater distance we have merely a perception blurred through excess or deficiency."

- Merleau-Ponty

We see each object in the world from a variety of viewing distances, and this distance changes as we move or as the objects move. Consequently, objects can appear at almost any size in the visual field depending on how close or far you stand in relation to them. For example, when crossing the road, a car in the distance will initially appear quite small in your visual field, but as the car approaches it will fill increasingly more of the visual field. In general, the physical size of an object and your viewing distance to that object determine the angle it subtends in your visual field.

Given the range of visual angles that an object can subtend in our visual field, is there a visual angle that is preferred? Research on perceptually privileged or canonical object views has traditionally looked at viewpoint (Tarr,

1998; Palmer, Rosch, & Chase, 1981), which has shown that the $\frac{3}{4}$ view is preferred in object identification.

However, the role of visual angle in object representation has largely been ignored, as much research has suggested that object representation is size-invariant (e.g. Biederman & Cooper, 1992). Here, we explore the possibility that each object has a perceptually privileged visual size, which we term the *normative size*.

Our hypothesis that there is a normative size has two motivations. First, resolution limits across the visual field might enforce this: something too big or too small in your visual field will be hard to see. Second, our experience with an object is statistically constrained: a fork is typically interacted with in a particular range of distances which exposes the visual system to a particular distribution of visual angles for forks (Figure 1).

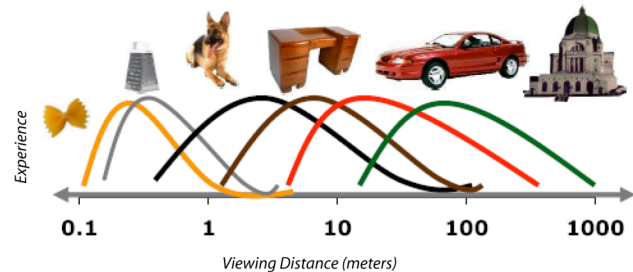


Figure 1: Hypothetical distribution of visual experience with sample real-world objects.

In the current studies, we examined the role of visual angle by varying the presented size on a monitor, as a way to simulate viewing distance. Specifically, observers could simulate moving closer or farther from objects by making them larger or smaller on the screen. In Experiment 1, observers were presented with a picture of an isolated object and adjusted the visual angle of the object until the object was the "right size" on the screen. Despite the subjectivity of the task, the selected size for each object was remarkably consistent across observers.

Additionally, we explored the possibility that, if perception is sensitive to a normative size, then memory for that object might also be biased toward a normative size. In Experiment 2, observers were presented with pictures of objects subtending a range of visual angles. Afterwards, participants adjusted the visual angle of the object to match the presented size. The remembered objects showed systematic errors biased towards the normative sizes from

Experiment 1. Experiment 3 measured observer’s sensitivity to detect when an object changed in visual size. The object could either change to be slightly larger or slightly smaller. Results show that when an object changed size in a direction towards the normative size, this change was more difficult to detect than when the object changed size away from the normative size. Experiments 2 and 3 both reveal predictable biases in memory for an object’s presented size, biased toward the normative size for that object.

Experiment 1: Existence of a statistically constrained perceptual norm

The first experiment examined if there is a perceptually privileged size to view an object. Specifically, if you are viewing an object on the screen, is there a visual size at which you prefer to view it? The aim was to determine how consistent people are when they choose an intuitively “best size” to view a real-world object on a monitor.

Based on visual acuity constraints we might expect observers to set all objects to the same size. One prediction, motivated from a perceptual stance, is that observers will show a foveal bias, resizing each object to fill the fovea or parafovea, so that each object will have the same selected normative size (e.g. 3 degrees). Another prediction, motivated from a conceptual stance, is that observers will select a size that is correlated with the real-world size of the object. So, objects that are larger in the world might be preferably viewed as larger on the screen.

For clarity, the term “visual size” will refer to the visual angle of the object on the screen, and the term “real-world size” will refer to the physical size of the object in the world (e.g. airplanes have a large real-world size, paperclips have a small real-world size, but both these objects can be any visual size on the screen).

Method

Object Norming Experiment A: Seven observers with normal or corrected vision (18-35 years old) were presented with 40 norming trials, each consisting of one object shown on a blank background. Observers could freely increase and decrease the visual angle of the object by pressing the up and down arrow keys. The following instructions were given: “For each object, select the best size to see it. Intuitively, when the object is at the smallest extreme, this is too small. When the object is at the largest extreme, this is too big. Use the keys to adjust the object’s size and then press enter when the size of the object is not too big or too small, but just right. There is no right answer, so select the size that is best for you.” Observers were seated approximately 57 cm from a 20 inch monitor, and the range of visual angles each object could subtend was 0 to 30 degrees. The objects were all color photographs of real world objects and were selected from a commercial database (Hemera Photo-Objects, Vol. I & II) to span the range of real-world sizes from small objects (e.g. paper clip) to large objects (e.g. statue). The order of objects was randomized across observers.

Object Norming Experiment B: We repeated the norming task using a sliding-mouse method of adjustment instead of the key press method. Six observers with normal or corrected vision (18-35 years old) participated. 100 objects were used in this experiment, including the 40 objects from Object Norming Experiment A. To adjust the visual angle of the object, participants moved the mouse up and down, and clicked to select the intuitively right size. All other procedures and instructions were the same.

Size Sorting Experiment: The purpose of this experiment was to obtain estimates on the real-world size of objects. This allows us to examine if there is a relationship between the visual angle of the normative size and the real-world size. Six observers gave ground truth rankings on the real-world size of 100 objects using a hierarchical sorting method. Thumbnails of the 100 objects were put on a 30” monitor and participants iteratively divided the images into two groups by dragging and dropping the thumbnails, until there were 8 groups of objects, ranked by real-world size. The number of objects per rank was not constrained to be equal.



Figure 2: 100 objects were sorted into 8 groups by their real-world size. Example objects are shown for each group.

Results

In object norming Experiment A, the selected visual angle of the height dimension for each object ranged on average from 4 degrees (e.g. peanut, thumbtack) to 14 degrees (e.g. Arc de Triumph, crane). We refer to the average selected size for each object as the *normative size*. Inter-rater reliability was calculated based on an adaptation of the Spearman-Brown formula (Rosenthal & Rosnow, 1987). Interestingly, despite the subjectivity of the task, inter-rater reliability for each object across observers was remarkably consistent, ($R=.9$, $p<0.05$). This correlation indicates consistency in which objects observers set the smallest and which objects they set the largest. However, two observers appeared to use more of the total range, making several very large or very small settings. For the purposes of Experiments 2 and 3, we set the normative size of each object equal to the average size selected, excluding these two of the seven subjects. This was done solely to select a normative size for each object that would be representative of most observers in the absolute value of the normative

sizes. In fact, without the two observers included the inter-reliability measure was slightly lower ($R=0.88, p<0.05$).

In the object norming Experiment B with 100 objects, inter subject reliability was again very high ($R=.7, p<0.05$). This study was repeated to confirm the reliability of the normative size selections with a new method of response, and to obtain a normative size for a larger number of objects. For the 40 objects that were tested in both experiments, the average selected sizes were within 3 degrees for all items, with no significant bias to be either smaller or larger.

With the normative sizes for 100 objects, we examined if there was a relationship between the visual angle of the normative size and the object's real-world size. To do so, we used the rankings from the Size Sorting Experiment. Across all ranks there was a minimum of 7 and a maximum of 23 objects. An object's size-rank was taken to be the mode rank across observers. A rank of 1 is the smallest real-world size and a rank of 8 is the largest (see Figure 2).

Using the selected visual angles from the norming experiment with all 100 objects, we averaged the normative size for the objects in each rank group, and plotted it against the rank size. As shown in Figure 3, the normative size of the object was highly correlated with the real-world size rank of the object ($r^2=0.96, p<0.05$).

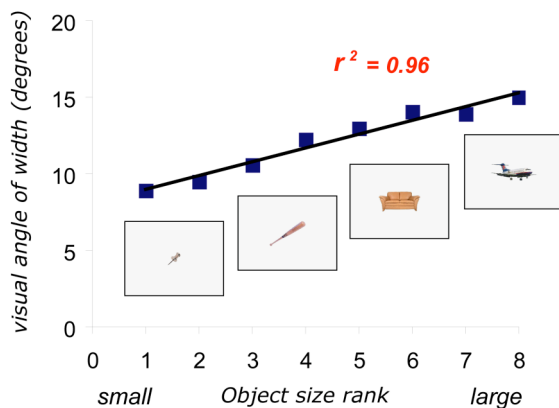


Figure 3. Correlation of visual angle of the normative size and real-world size ranking.

Discussion

The results show striking reliability for normative sizes across observers and across reporting methods, despite the subjectivity of the task to select the “intuitively right size” to see the object. This provides initial evidence for the existence of a preferred or *normative size* for each object. Critically, these normative sizes across observers cannot be explained solely by visual acuity constraints, which predict that all objects would be set to fill the fovea or parafovea. Moreover, the normative size was significantly correlated with real-world size (i.e. the larger the real-world size of the object, the larger the selected visual angle). This suggests that observers’ conceptual representation of the object’s real-world size influenced the intuitively right size to see the object on the screen. This demonstrates that the normative

size is influenced by knowledge or experience of these objects in the world. Interestingly, the normative sizes were relatively small on the screen (4 – 14 visual degrees); thus participants were not using the whole range of the monitor, even for the largest objects. This is suggestive that some acuity factors of the visual field may be also playing a role. Thus, the normative size may reflect a balance between conceptual and perceptual factors.

Experiment 1 gives a *perceptual* norm for each object, corresponding to the average size that observers selected as the best size to see it. If this size is indeed privileged in perception, then it might also be privileged in memory. Experiment 2 tests this hypothesis directly, and predicts that long-term memory for an object’s visual size will be biased towards the normative size found in Experiment 1.

Experiment 2: Bias of Long-Term Memory towards the Normative Size

Previous work has shown a systematic bias in memory for objects and scenes called boundary extension: people tend to remember a picture of an object or scene as farther away than it was originally viewed (Intraub & Richardson 1989). We hypothesize that memory for an object’s size is biased toward the normative size, and that boundary extension is only one possible result of this memory bias. On this view, when an object is presented larger than the norm, memory for that object will be biased toward its smaller normative size. Thus, at test, the object will be remembered as smaller than it actually was presented, which is consistent with the classic boundary extension effect. However, the normative view also predicts that, if the object is presented smaller than the norm, memory for that object will be biased toward its larger normative size. Thus, at test, the object will be remembered as larger than it actually was. This effect in the opposite direction from boundary extension has not been observed, and would not be predicted given standard interpretations of boundary extension (Gottesman & Intraub, 2003). However, it could be accounted for naturally within the normative representation framework.

Experiment 2 used a classical boundary extension paradigm in which a stream of objects were presented and observers were told that they were going to be tested on their memory for these objects, without explicit instruction about the kind of memory test. Afterwards, subjects had to report the visual size of the object that was presented. In the current experiment, objects were presented either “too big” or “too small” relative to each object’s normative size from Experiment 1.

Method

Twenty categorically unique objects from the set of 40 used in Experiment 1 were selected, uniformly across the eight real-world size groups. For each of the 24 participants, ten objects were randomly selected for the too-big condition and ten were randomly selected for the too-small condition. The experiment consisted of two phases. During the learning phase, each object was presented on the screen for

5 seconds with a 1 second inter-stimulus-interval. Participants were informed that after a learning phase of 20 objects, they would be “tested on their memory of the objects.” As with the classic boundary extension paradigm, observers were *not* informed that the memory test was specifically going to be for the object’s size on the screen. Following the learning phase, participants were presented with each of the 20 objects, one at a time in randomized order, and used the up and down arrows on the keyboard to resize the object to match the size they saw during the learning phase.

Object ‘step-sizes’ were linear steps in visual angle of the height dimension corresponding to approximately 1 visual degree in height increase per step. In the learning phase, each object was presented 5 steps larger or 5 steps smaller than its normative size. Key presses advanced the object size one step. For each object there were 40 possible step sizes. During the memory test, the object was initially presented jittered around the middle step position.

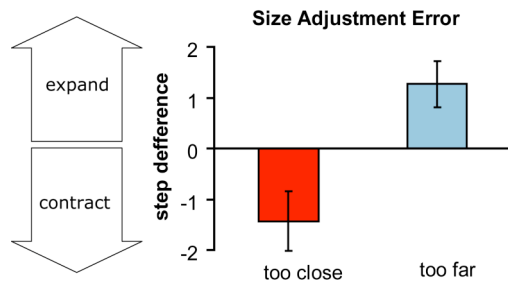


Figure 4. Deviation between presented visual size and selected visual size. Error bars are 1 S.E.M.

Results

Memory performance was quantified by calculating the number of steps between the object size selected during the testing phase and the object size presented during the learning phase. *Negative* numbers indicate that the object was reported as *smaller* than at learning (object contraction); *positive* numbers indicate that the object was reported as *larger* than at learning (object expansion).

Participants showed significant contraction (remembering a smaller object) for the objects presented larger than the normative size (too-big condition: $t(23)=2.45, p<0.05$). This is consistent with the known boundary-extension effect. Critically, participants also showed significant expansion (remembering the object bigger) for objects presented smaller than the normative size (too-small condition: $t(23)=2.80, p<0.05$, Figure 4). None of the 24 subjects showed memory errors with the opposite trend. In an item analysis, 15 of the 20 objects presented “too big” during learning showed significant compression and 18/20 objects presented “too small” showed significant expansion.

Discussion

Long-term memory for objects showed a systematic bias toward the normative size, both for objects presented too big

and objects presented too small. The data show that when an object is seen larger than its normative size, it is remembered as smaller (closer to the norm), whereas when an object is seen smaller than its normative size, it is remembered as bigger (again, closer to the norm).

Importantly, the normative size was taken from Experiment 1, which showed a range of visual angles. So, the large real-world size objects (airplane, crane), when presented in the too-small condition are in fact comparable to the small real-world size objects (button, tack) when presented in the too-big condition. Despite the roughly equivalent presentation size on the monitor for stimuli in these conditions, the large and small real-world size objects show memory errors in *opposite* and *predictable* directions, based on the normative size found from Experiment 1.

The normative hypothesis poses an alternate explanation for the current interpretation of the boundary extension phenomenon. In the boundary extension effect, close-up scenes are remembered as farther away than they were actually perceived (Intraub & Richardson, 1989). While classically this effect is thought to be a phenomenon about the visual information at the edges of the scene, more recent evidence suggests this effect is driven by the central object (Bertamini et al, 2005). Our results suggest that boundary extension actually reflects a memory bias towards the normative size of the central object.

Experiment 3: Bias of Short-Term Memory towards the Normative Size

In Experiment 2, observers were not explicitly informed that they would be tested on memory for the objects visual-size. Further, observers were required to remember 20 objects before being tested. This shows that *long-term* memory for object size is biased toward the normative size, in the *absence of explicit encoding* of the size of the object. The aim of Experiment 3 was to generalize the effects of the normative size to a situation where memory is tested immediately after the presentation, and where observers explicitly know they are being test on memory for size. A change detection paradigm is suitable to evaluate short-term memory of a single event (Luck & Vogel, 1997).

Suppose an object is initially presented at a larger visual angle than its normative size. We hypothesize that memory for this object will be shifted slightly smaller, toward the direction of the normative size. Thus, we predicted that if the object is presented again at a slightly smaller visual angle than the first, this change should be more difficult to detect than if the object is presented again at a slightly larger visual angle. Similarly, suppose your first view of an object is smaller than its norm; again, the normative theory predicts that your memory of that object’s visual size will be larger. Thus, if the object reappears slightly larger, this change should be more difficult to detect than if the object appears slightly smaller.

Put succinctly, we hypothesized that a change in size *toward the normative size* should be more difficult to detect than a change of visual-size *away from the normative size*. If

there is no systematic bias in the memory of the first stimulus, then there should be no difference in detection if the object changes toward or away from the normative size.

Method

Twelve observers participated in the change detection task. On each trial, an object was presented for one second, masked for 200ms, followed by a blank screen for one second, and then re-appeared at the same or a different size. The object remained on the screen until participants pressed a key indicating whether the size of the object was the same or different. Forty objects were used in the experiment. Each object was repeated in 12 trials: on six trials the object was presented as too-big relative to its norm, and on six trials the object was presented as too-small. In two of the six trials, the object changed toward the norm, in two of the six trials, the object changed away from the norm, and in the remaining two trials, the object was presented at the same size. The first image of the object was presented at 5 steps smaller or larger than its normative size from Experiment 1. The second image could change by 3 steps toward or away from the norm. Figure 5a shows an example object and the size of the changes.

Results

A measure of sensitivity (d' -prime) was calculated for each type of change (change toward the norm, change away from the norm), by taking the z -score of the percent hits minus the z -score of the percent false alarms. An ANOVA was computed on d' -prime with change type (toward vs. away) and starting size (too-big vs. too-small) as factors. There was a significant effect of change type ($F(1,11)=34.94, p<0.05$). Paired t -tests show the change toward the norm condition was more difficult than the change away from the norm condition ($t(11)=2.80, p<0.05$, Figure 5b).

Additionally, there was a significant effect of starting size ($F(1,11)=15.31, p<0.05$). Paired t -tests reveal that the too-big condition was more difficult overall than the too-small condition ($t(11)=8.50, p<0.05$). There was no significant interaction between change type and starting size ($F(1,11)<1$). When the starting size was too small, a change toward the norm was significantly more difficult than a change away from the norm ($t(11)=2.20, p<0.05$); however, when the starting size was too big, the post-hoc paired t -test did not reach significant ($t(11)=1.36, p=0.2$).

Discussion

The results of Experiment 3 indicate that a change toward the normative size was significantly harder to detect than a change away from the normative size (Figure 5b). This suggests that short-term memory for the initially presented object is also biased toward the normative size. This is in some sense a stronger demonstration of this effect, because unlike the surprise memory task (Experiment 2), here the observers were explicitly aware that the visual size of the object was the dimension of interest for the task, and were

probed immediately afterwards, with only 1 second delay and no intervening stimuli except for a brief mask.

The too-big condition, in which the stimuli were presented larger than the normative size, was significantly more difficult than the too-small condition. Why might this be the case? Here, the comparison images were 3 steps on either side of the too-big or too-small condition. However, a change of 3 degrees on a small starting size of, for example 5 degrees visual angle is quite different than that same change of 3 degrees on a large starting size of 20 degrees visual angle, according to Weber's law. Thus, it is not surprising that the too-big condition was more difficult than the too-small condition (see Figure 5a). If we assume that this difficulty reflects more noise in all the size estimates, then the lack of significance in the too-big paired t -test could be due to insufficient power.

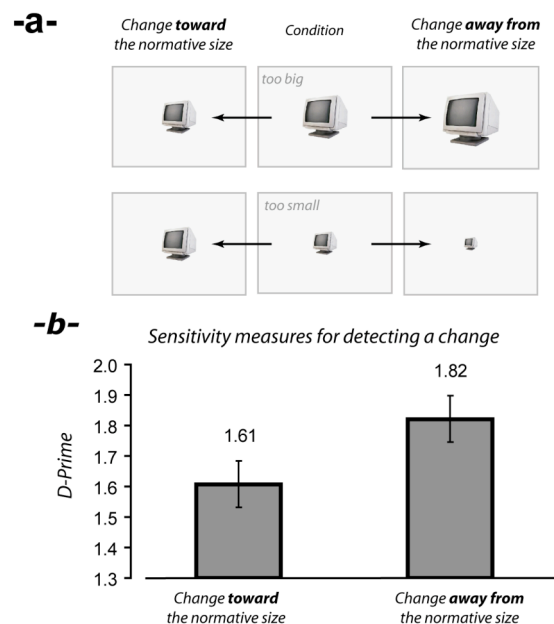


Figure 5. (a) Visual-angle changes are shown for a sample object. The first stimulus is presented 5 steps larger or smaller than its normative size (middle column). The second stimulus changed towards the normative size (left column) or away from the normative size (right column). (b) D' -prime measures are plotted for the change-toward condition and the change-away condition.

General Discussion

Experiment 1 showed that for each object, there is a particular visual-angle, termed the normative size, which is *perceptually* privileged across observers. Experiments 2 and 3 demonstrated that *memory* for an object's visual size is biased in a systematic direction that is predicted by the normative size found in Experiment 1. Experiment 2 showed that long-term incidental memory for object size is biased toward the normative size: objects that were presented larger than their normative size tended to be remembered as smaller, and objects that were presented smaller than their normative size tended to be remembered

as larger. Experiment 3 showed that short-term explicit memory for object size is also biased toward the norm: a change in object size toward the normative size was harder to detect than a change away from the norm. Taken together these experiments provide support for the following claims: (1) Perception of objects is implicitly sensitive to a normative size. (2) The normative size is influenced by the real-world size of objects. (3) Memory for object size is biased towards the normative size.

What are the possible reasons why there is a normative size in perception and memory? One account is a perceptual account: we move around in the world to place the visual information of interest in the proper place in our visual field to do some task. So, if we need to judge the ripeness of an apple on a branch, we may move in close to bring the color and shape into the foveal representation of our visual field. If we need to judge how low the apple is hanging on the branch, we may move farther back. Thus one hypothesis is that our visual-motor perception is implicitly sensitive to this visual size in order to optimally extract information from the world.

However, another account is one from memory of accumulated visual experiences acting on objects in the real world. We interact with a bow-tie pasta in the range of, say, 1 cm to a 1 meter; we interact with cars on the range of 3 meters to 30 meters. These viewing distance distributions vary based on the size of the object, and lead to corresponding distributions of experienced visual angles. Thus, the normative size might correspond to the mode of our active visual experience with an object, and is influenced by the natural statistics of the world (Gibson, 1979). One broad hypothesis is that the default representation of an object is the statistical mode of visual experience with that object along *any* relevant dimension.

While this work demonstrates a normative concept for an object's visual size, this could be extended to other properties of real world objects, such as elevation, shape, or state. Indeed, research on canonical viewpoints (Tarr et al., 1998; Blanz, Tarr, & Buelthoff, 1999) and color (Tanaka & Presnell, 1999), suggests that there are privileged perceptual views along other spatial and featural dimensions.

Systematic memory errors towards a norm have implications for the nature of object representation. For instance, one account of object memory suggests that when recall the size of a previously presented object, we combine our perceptual estimate with our conceptual *normative* estimate. Further, objects could be represented as a sum of their parts (Biederman, 1987), or as implied here, as deviations from their statistical mode. Some theories of efficient coding (Barlow, 2001) and prototypes (Rosch, 1981), are suggestive and consistent with this theory of object representation.

In the current experiments, visual size on a monitor was used as a proxy for viewing distance in the real world. Future work will be necessary to determine how this normative concept operates in an embodied context, where people walk towards physical objects in a natural setting.

Acknowledgments

The authors wish to thank G. Alvarez, M. Greene, and B. Hidalgo-Sotelo and three reviewers for comments on the manuscript. This research is supported by a NDSEG fellowship awarded to T.K. and a NSF Career Award #0546262 to A.O.

References

- Barlow, H. (2001). Redundancy reduction revisited. *Network: Computation in Neural Systems*, 12, 241-253.
- Bertamini, M., Jones, L.A., Spooner, A., & Hecht, H. (2005). Boundary extension: the role of magnification, object size, context and binocular information. *Journal of Experimental Psychology: Human Perception and Performance*, 31, 1288-1307.
- Biederman, I. & Cooper, E.E. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, 94, 115-147.
- Biederman, I. & Cooper, E.E. (1992). Size invariance in visual object priming. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 121-133.
- Blanz, V., Tarr, M.J., & Buelthoff, H.H. (1999) What object attributes determine canonical views? *Perception*, 28, 575-599.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.
- Gottesman, C.V. & Intraub, H. (2003). Constraints on spatial extrapolation in the mental representation of scenes. View-boundaries versus object boundaries. *Visual Cognition*, 10, 875-893.
- Intraub, H., and Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 179-187.
- Luck, S.J., & Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279:281.
- Palmer, S.E., Rosch, E. & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long and A. Baddeley (Eds.), *Attention and Performance IX*. Erlbaum, Hillsdale, NJ.
- Rosch, E. (1981). Principles of Categorization. In A. Collins & E. E. Smith, *Readings in Cognitive Science, a Perspective from Psychology and Artificial Intelligence*. Morgan Kaufmann Publishers (pp. 312-322).
- Rosenthal, R. & Rosnow, R.L. (1984). *Essentials of Behavioral Research: Methods and Data Analysis* Boston: McGraw-Hill.
- Tanaka, J.W., & Presnell, L.M. (1999). Color diagnosticity in object recognition. *Perception & Psychophysics*, 61, 1140-1153.
- Tarr, M.J., Williams, P., Hayward, W.G., & Gauthier, I. (1998). Three-dimensional object recognition is viewpoint dependent. *Nature Neuroscience*, 1(4), 275-277.