

News from the field

© Psychonomic Society, Inc. 2011

Object files

How do we maintain object persistence?

Moore, C.M., Stephens, T. & Hein, E. (2010) Features, as well as space and time, guide object persistence. *Psychonomic Bulletin & Review* 17(5):731–736.

Our everyday visual experience is highly dynamic. Objects move around and often disappear and reappear due to occlusion, eye movements, or simply blinking. To maintain a coherent representation of the world we need to maintain a steady representation of objects that will prevail such constant changes. Kahneman, Treisman, and Gibbs (1992) suggested the theoretical construct of object file as a possible solution for these continuing changes. Object file was defined as a temporary episodic representation of the object, in which the current state of the object is linked to its prior states. Various studies have suggested that object files are established and updated solely on the basis of spatiotemporal properties, and that surface features such as color or shape do not play a role in maintaining object persistence. In contrast to these studies, Moore, Stephens and Hein were set to demonstrate that surface features also play a role in the establishment and updating of object files. They employed a similar object-reviewing paradigm to that employed in the past. The initial display included two objects (e.g., two differently colored squares) and two symbols embedded within the objects. Following this initial display, the symbols disappeared and the display went through some changes. In the spatiotemporal condition the squares moved with no color changes. The two squares also moved in the feature-switch and flash conditions, but with the former they switched colors at the last frames of motion, and with the latter they disappeared in the last frames and then reappeared in their original colors. Finally in the feature

condition there was no motion—the squares disappeared immediately after the initial display for the entire ‘motion interval’ and reappeared in the final positions. In the final display, two symbols were presented again within the squares and the observers had to indicate whether both symbols appeared in the initial display.

The results of the spatiotemporal condition replicated previous findings: When the final symbols matched the initial symbols, performance was better when the final symbols reappeared in the same squares as in the initial display (i.e., in the same objects, where object is defined based on its motion history). However, this congruency effect was reversed in the feature-switch condition but not in the flash condition. Thus, when the objects that were defined by spatiotemporal information changed colors the typical congruency effect was disrupted, suggesting that surface features do play a role in the maintenance of object persistence. Moreover, unlike prior studies, a significant congruency effect was also found in the feature condition: performance was better when the final symbols reappeared on squares with the same colors as in the initial display. This suggests that surface features also play a role in the establishment of object files.

Interestingly, when Moore et al. employed letters instead of their original, less familiar symbols, and presented only a single letter in the final display (i.e., following the exact experimental conditions of previous experiments), no congruency effect for the feature condition was found. Hence, when the memory demands posed by the task were lower, the results of previous studies were replicated. Only when memory demands were high, was the role of surface features revealed. This finding suggests that performance in this object-reviewing paradigm, that is commonly used to study the nature of object files, may not truly reflect online maintenance of object persistence. Future study of object files seems to be in need of a new paradigm.—Y.Y.

Film studies

Paying attention to the movies

Attention and the Evolution of Hollywood Film. James E. Cutting, Jordan E. DeLong and Christine E. Nothelfer. *Psychological Science* 2010 21: 432

A search through PsycInfo on the term “cinema” yields 1143 publications, but very few of these overlap with terms like “attention”, “perception”, and/or “psychophysics”. The analysis of film in psychology has generally been left to those of a psychoanalytic or sociological bent. But film has a lot to teach us about perception. For example, the recent revival of the study of change blindness had its roots in the study of transitions in film (Levin & Simons, 1997).

In this tradition, Cutting, DeLong, and Nothelfer (2010) take a “cinematics” approach to film, analyzing local and global structure of shot length in 150 Hollywood films evenly sampled from the decades between 1930 and the present. In a Herculean effort, each film (averaging 114 min long, not counting title and credit sequences) was segmented into a series of shots by a semi-automated process which took 15–36 hours per film, and yielded a d' of 5.5 against a test sample comprising *The Revenge of the Sith* and *Spies Like Us*.

What can shot length tell us about attention? A curmudgeon might point to the fact that mean shot lengths have been decreasing over the last 70 years to argue that films are increasingly catering to people with short attention spans. Cutting et al. have something a little more sophisticated in mind. Analyzing the local shot structure with autocorrelation analyses, they show that shot length autocorrelations have been increasing linearly since the 1930s. That is, the length of a given shot is now a better predictor of the length of the shots that precede and follow it than it was 70 years ago. This suggests that shots are becoming more clustered, so that shots in a given sequence (a car chase, for example) will tend to be of the same length, and this tendency has increased over time. The film in the sample with the most coherent shot-length structure, on this criterion, was *Rocky IV*. In general, action films led the way in this transformation, but the effect can be seen across genres.

At the global level, Cutting et al. looked for pink noise, defined as a $1/f$ slope in the Fourier spectrum analysis. A $1/f$ noise pattern is thought to reflect natural fluctuations in human attention. By fitting a mixture model with both white and pink noise components to their shot length data, Cutting et al. found that films have been approaching the $1/f$ pattern in recent decades. While films noir such as *Detour* and *Sunset Boulevard* exhibited slopes around $0.09/f$, *Revenge of the Sith* clocked in at $1.14/f$. Again, while action movies are in the vanguard, the effect can be

seen across genres. Interestingly, however, unlike the local, autocorrelation analyses, the effects seen in the global analysis were not monotonic over time. Slopes were actually higher in the 1930s and 1940s than in the subsequent two decades.

Rather than catering to or encouraging short attention spans, per se, Cutting et al. argue that the global trends in shot organization reflect filmmakers gradually adapting their craft to the intrinsic temporal structure of attention. They predict that over the next 50 years, the trend towards a $1/f$ structure in shot durations will become more pronounced.—T.S.H

Bibliography

Levin, D., & Simons, D. (1997). Failure to detect changes to attended objects in motion pictures. *Psychonomic Bulletin & Review*, 4(4), 501–506.

Spatial vision

Finding your center (or centroid)

Juni MZ, Singh M & Maloney LT (2010) Robust visual estimation as source separation, *Journal of Vision* 10(14):2, 1–20, <http://www.journalofvision.org/content/10/14/2>.

How does human vision integrate information across space? A number of experiments have used tasks requiring judgments about the centroids of briefly presented dot clouds to investigate this question. Previous studies have tended to focus on spontaneous, mandatory centroid extraction processes, often emphasizing that such computations may underlie various visual illusions (e.g., Morgan & Glennerster, 1991; Bulatov et al., 2009). A recent study of Juni, Singh & Maloney (2010) points in a new direction. The original motivation for this work was to investigate the possibility that in locating the center of a dot cloud, human vision may use a “robust” estimator—i.e., one that selectively down-weights peripheral dots relative to central ones. (Juni, Singh & Maloney provide a nice review of the important statistical reasons to prefer robust center-estimators to a standard, non-robust center-of-gravity computation.) The first experiment reported by Juni, Singh & Maloney is a straightforward test of this hypothesis: observers judge the location (right vs left of an implicit vertical line) of the centroid of a dot cloud. A variant of the classification image method is used to estimate the impact different dots exert on these judgments as a function of their horizontal distance from the true centroid. In this context, most participants seem to use an ordinary, non-robust center-of-gravity computation (a small number use robust computations, but just as many use “anti-robust” computations which give more weight to peripheral vs central dots

in the cloud). This result prompts two interesting follow-up experiments. Having failed to find evidence of a spontaneously “robust” center-estimator, Juni, Singh & Maloney (2010) proceed to ask what sorts of dot discounting people *can* achieve in computing centroids, given explicit instructions. In Experiment 2, participants view stimuli composed of a mixture of 40 dots uniformly distributed across the entire display and 100 dots drawn from a bivariate Gaussian distribution. The task is to judge the centroid of the 100-dot Gaussian cloud, ignoring the dots from the uniform distribution. In this context, participants do very well at down-weighting peripheral dots. In Experiment 3, participants attempt to judge the center of gravity of a 100-dot Gaussian cloud while striving to ignore a tight, 15-dot “contamination” cluster added somewhere to the display. Again (except, interestingly, when the contamination cluster occurred around 2 standard deviations right or left of the centroid), participants’ judgments were largely immune to the contamination cluster. I find this paper very interesting mainly because of the new terrain it opens up. Clearly, participants have a great deal of latitude in selecting the computations they can use to extract the center of dispersed target. And when one reflects on the many different purposes that might motivate such a judgment, this makes a lot of sense. The challenge we face now is to determine the scope and limits of this strategic flexibility.—C.F.C

Morgan MJ, Glennerster A (1991) Efficiency of locating centres of dot-clusters by human observers, *Vision Research* 31(12), 2075–2083.

Bulatov A, Bertulis A, Bulatova N, Loginovich Y (2009) Centroid extraction and illusions of extent with different contextual flanks, *Acta Neurobiologiae Experimentalis* 69 (4), 504–525.

Juni MZ, Singh M & Maloney LT (2010) Robust visual estimation as source separation, *Journal of Vision* 10(14):2, 1–20, <http://www.journalofvision.org/content/10/14/2>.

Face perception

Faces and bodies in motion

Pilz, Vuong, Bulthoff & Thornton (2011). Walk this way: Approaching bodies can influence the processing of faces. *Cognition*, 118, 17.

As we move about our environments, faces are typically seen in the context of the whole person, including head, neck, body, and characteristic movement. Yet, research on face perception has most often focused on how we recognize and process disembodied static faces. Pilz et al. investigated whether viewing faces in the context of moving bodies influences the perception of faces and found

that faces were identified more quickly following presentation atop a body walking toward or approaching the viewer. Across experiments, participants were presented with the same identical animated figure with head models that had different faces. The figures were presented as either approaching or receding from the viewer, or as static. The results showed that viewers responded faster in a sequential matching task when face targets followed animated sequences of a figure approaching the viewer than when following a receding or static figure. The benefits for viewing a face atop an approaching body lasted over time, influencing a delayed visual search task in which participants searched for a static target face among varying numbers of distractors. Better performance after exposure to approaching figures occurred despite the use of identical moving bodies, which eliminated individual body movement cues as a source of information in the task. Face processing was facilitated in situations in which familiar movement and contexts were provided. The authors argue that the particular sensitivity of the visual system to biologically relevant movement, such as found in approaching bodies, uniquely influences the processing of facial identity.—L.C.N.

Perceptual organization

Objects Distort Space:

Vickery & Chun. (2010). Object-based warping: An illusory distortion of space within objects. *PSYCHOL SCI*, 21, 1759.

Perceptual organization has been actively studied since the early 20th century when the Gestalt psychologists presented some of the basic phenomena. Modern perspectives on perceptual organization, or grouping, highlight the functional importance of the processes that produce organization. These grouping processes determine which visual features belong together by virtue of being on the same object, and these grouped features provide the input to high-level visual processes responsible for object recognition and attentional allocation.

Given the important role of perceptual organization in mediating between features and objects, one might expect organizational processes to provide a spatially accurate clustering of features. However, as with many visual processes, perceptual organization does not appear to form a veridical representation of visual space. Vickery and Chun report a new illusion, termed object-based warping, in which a perceptual group (i.e., a perceptual object) alters spatial perception.

In a basic demonstration of object-based warping, Vickery and Chun present a figure with a pair of dots inside a

rectangle and another pair outside the rectangle. Although the dots inside the rectangle are separated by the same amount of space as the dots outside the rectangle, most viewers see the dots inside the rectangle as being farther apart than those outside the rectangle. This informal observation is bolstered by more detailed measures of spatial separation, in which participants were shown a pair of ‘standard’ dots and a pair of adjustable dots. Participants used a mouse to vary the distance between the adjustable dots to match the separation of the standard dots. To assess object-based warping, the standard dots appeared inside a rectangle or on a uniform background; additional conditions placed the standard dots inside an occluded object, inside a Kanizsa-style illusory object, or inside either the same object or different objects.

Across all of the various conditions, Vickery and Chun found that strong perceptual structure produced a distortion of the perceived distance between the dots. For example, when the standard dots appeared inside a rectangle, participants overestimated their separation by approximately 17% of the actual distance. These separation estimates were significantly smaller when the dots appeared against a blank background or within a weakly organized perceptual group. Vickery and Chun discussed several possible explanations for their object-based warping effect. They presented results that ruled out a ‘contrast’ account, in which dots inside an object appeared more distant because they occupied more of the object’s space than dots outside the object occupied of the larger background. When dots were placed inside a smaller rectangle, the dots occupied even more of the object’s space, yet this manipulation did not affect the amount of object warping. Vickery and Chun conclude with two possible accounts that await further investigation, an attentional spread difference within an object that affects perceived space, or an exaggerated cortical representation for objects compared to backgrounds. Irrespective of the exact mechanism of object-based warping, the phenomena may prove useful as a measure of the strength of perceptual organization.—S.P.V.

Global scene processing

Does this desert make me look cold?

Greene, M. R., & Oliva, A. (2010). High-level aftereffects to global scene properties. *Journal of Experimental Psychology: Human Perception and Performance*, 36(6), 1430–1442.

Once upon a time, we thought that visual aftereffects showed us the simple building blocks of visual perception. If you were exposed to lines tilted to the left of vertical, vertical lines looked like they were tilted to the right

because you had adapted orientation channels tuned to left tilt. Your perception of vertical was based on some balance of channels tuned to the left and to the right and exposure to the left-tilted lines had skewed that balance to the right. With methods of this sort, we mapped out the tuning of processes for color, orientation, size/spatial frequency, motion, and so on. Kohler and Wallach had changed the shape of objects with “Figural Aftereffects” in the 1940s but no one was going to spend time adapting to cows to produce a cow aftereffect. Aftereffects happened with simple stimuli.

Things became less simple when it became clear that, in fact, you could adapt to more complex stimuli like faces. For instance, if you adapt to a clearly male face, a neutral face will appear more female. Now, in a new paper, Greene and Oliva have shown that you can adapt to global scene properties. Global scene properties “represent the structure and function of a scene”. Examples include “mean depth” (Is everything near to the viewer or far away?), “navigability” (Could you walk through this scene?), and “temperature” (How hot is it here?). Pick a random scene and you will see that you could rate it on these sorts of dimensions. Greene and Oliva have vast number of scenes rated. In order to adapt to one property, they took 100 scenes from one end of that dimension; for example, 100 natural scenes from the natural end of the natural/urban continuum. Observers looked at these, one after the other, for 5 min. In order to make sure that they were paying attention, observers had to push a response key whenever the same picture repeated. After this initial adaptation, Os saw a single scene for just 100 msec and were asked to label it, in this example, as “natural” or “urban”. Then they would get 10 sec more adaptation, another test stimulus, and so on. The test stimuli were drawn from the middle of the natural/urban continuum. The results showed that perception of those scenes had been displaced toward the end opposite from the adaption. Thus, after adapting to natural, a relatively neutral stimulus was more likely to be labeled “urban” than before adaptation.

So, global scene properties adapt. Does that mean that like orientation or color, they are basic building blocks of more complex percepts? Greene and Oliva provide some evidence that this is so. Consider forests and fields. Forests tend to be more “closed”; fields, more open. If you adapt to openness, not only will middling scenes be labeled as “closed” more frequently, but images with trees and open spaces, lying between clear “forests” and clear “fields” labels, will tend to be labeled as “forest” more frequently. “Closed” seems to be a building block of the basic level category “forest”.

Over 30 years ago, Colin Blakemore called the aftereffect “the psychologist’s microelectrode”. Greene and Oliva are showing us how this electrode can record from global,

scene-wide properties just as it can record from local processing of basic visual features.—J.M.W

Olfaction

Nasal Mucus affects the Sense of Smell

Nagashima, A., & Touhara, K. (2010). Enzymatic conversion of odorants in nasal mucus affects olfactory glomerular activation patterns and odor perception. *The Journal of Neuroscience*, 30(48), 16391–16398.

The next time you blow your nose, consider that the contents of your Kleenex may shed new light on our ability to recognize and distinguish odors. According to a new study conducted by Ayumi Nagashima and Kazushige Touhara at the University of Tokyo, the enzymes contained within nasal mucus appear to change the molecular structure of certain odorants. More importantly, these enzymatic changes appear to occur rapidly enough to influence the activation patterns in the olfactory bulb thought to reflect the neural correlates of olfactory percepts. These conclusions were supported by a variety of experimental findings obtained from laboratory mice. First, the researchers demonstrated that the molecular structure of two test compounds—benzaldehyde and acetyl isoeugenol—were altered when these compounds were incubated with nasal mucus that had been isolated from the mice as well as when these compounds were inhaled by intact mice. In contrast, no such change occurred when the test compounds were mixed with “boiled mucus.” Next, the researchers attempted to determine which of the many possible candidate enzymes contained in the nasal mucus was responsible for the conversion by introducing various enzymatic inhibitors. The results were clearest for the acetyl isoeugenol conversion which suggested that the conversion process was mediated by carboxyl esterase rather than carbonic anhydrase. The researchers then used this selectively to further examine how patterns of activation in the glomeruli of the mice changed as function of whether the carboxyl esterase inhibitor was introduced or not. The results suggested that the pattern of glomerulus activation that was elicited by acetyl isoeugenol changed when the carboxyl esterase inhibitor was introduced whereas this same pattern of glomerulus activation remained unchanged when the carbonic anhydrase inhibitor was introduced. Because patterns of glomerulus activation are believed to reflect the neural correlates of olfaction percepts, these findings suggest that the effects of nasal mucus are not confined to the nasal epithelium, but may actually modulate olfactory-based behavior. This final hypothesis was supported by a study in which mice were trained to discriminate one of two odors with a sugar

reward. Under normal circumstances, the mice spent more time sniffing the rewarded odors, then the unrewarded odors. However, the mice lost their ability to discriminate the two odors when a potent inhibitor was introduced into their nasal cavity. Altogether, these findings are important because they suggest that nasal mucus can alter our sense of smell by changing the odors that enter our noses.—B.S.G.

Lyrics perception

Seeing to avoid *mondegreens*

Jesse, A. & Massaro, D. W. (2010). Seeing a singer helps comprehension of the song’s lyrics. *PB&R*, 17, 323.

It is well known in the field of speech processing that viewing a speaker’s face cannot be ignored by a perceiver. Indeed, there are benefits to expect from viewing, even when the speech signals are not degraded. There is a particular context where speech could be very difficult to decode: lyrics in songs. Not only could comprehending lyrics be difficult, but something different sounding like speech could be heard. Thus, in the Scottish ballad “The Bonnie Earl O’Murray”, “And laid him on the green” could sound like “And Lady *Mondegreen*”. Therefore, the song context sometimes leaves plenty of room for improving the comprehension of words. So far, research revealed that the benefit obtained from visual information with songs’ lyrics was as strong as what is observed with regular speech. In their study, Jesse and Massaro used visual singing instead of visual speech as was the case in previous studies on this topic, for testing the potential benefit of visual information on the comprehension of lyrics. The authors demonstrated that seeing and hearing the singer leads to substantial comprehension benefit (about 35% recognition improvement) when compared to trials where only seeing or only hearing the singer was possible. The magnitude of this benefit compares to what is observed in speech perception. Therefore, spoken and sung languages seem to be a domain-general phenomenon. Interestingly, in the same *PB&R* issue, other authors (Thompson et al., p. 317) report that seeing the face of a singer influences the perception of music. More specifically, the singer’s facial expressions carry information about another aspect auditory processing, namely the pitch relations.—S.G.

Illusions

Wherefore art thou r?

Bosten & Mollon (2010). Is there a general trait of susceptibility to simultaneous contrast? *Vision Research*, 50, 1656–1664.

The tilt illusion (Gibson, *J. Exp. Psychol.* 20:553; Over et al., *J. Exp. Psychol.* 96:25) and the Chubb illusion (Chubb et al., *Proc. Nat. Acad. Sci.* 86:9631; Lotto & Purves, *J. Cog. Neurosci.* 13:547) are but two examples; the exaggeration of feature contrast is so widespread in sensory systems that you might expect there to be a general trait for the susceptibility to repulsion. Call it *r*. Just as Spearman (*Amer. J. of Psychol.* 15:201) went hunting for general intelligence (*g*), Boston & Mollon have gone hunting for *r*. Alas, the *r* factor seems no more real than the *g* factor (or the *g* spot, for that matter; Burri et al., *J. Sexual Med.* 7:1842).

What Boston & Mollon did find were large individual differences in various types of perceptual contrast enhancement (e.g. luminance, chromaticity, tilt, etc.). Within observers, repeated measurements of any one illusion were highly correlated. However, with one exception (yellow/

blue contrast and red/green) there were no significant between-observer correlations.

The exciting implication of Boston & Mollon's result is that correlations between repulsion susceptibility and other psychiatric (e.g. Dakin et al., *Current Biol.* 15:R822) and physiological (e.g. Schwartzkopf et al., *Nature Neurosci.* doi:10.1038/nn.2706) factors are likely due to mechanisms more specific than previously thought. Some of these mechanisms, like the lateral inhibition thought responsible for the tilt illusion, have obvious physiological correlates. Others, like that responsible for the Chubb illusion, do not. Too often failures of correlation crop up in "control" experiments, designed to show that a study's "main result" is not artifactual. That's why Boston & Mollon's results should please sceptics. Despite extensive measurements, an elegant experimental design, and the clear desire to find *r*, it just isn't there.—J.A.S.