Seminar on the Spatial Resolution of Attention
John Palmer and Cathleen Moore
Spring 2006
20 June 2006

This is a summary of a seminar that considered the spatial resolution of visual attention. This document contains the syllabus, weekly question sets, and a summary of the discussion for each week.

Seminar on the Spatial Resolution of Attention
Course Syllabus
John Palmer and Cathleen Moore
20 June 2006 (updated after seminar)

Synopsis

This seminar will analyze the literature on the spatial resolution of visual attention. In brief, the spatial resolution of attention is the minimum spatial separation needed to attend to one location and not a nearby location. We will begin with background readings on visual attention and then address several different paradigms that attempt to measure the spatial resolution of attention. The readings include work from visual psychophysics, cognitive psychology and visual neuroscience.

General Information

Psychology Graduate Seminar 555C, 2 Credits, CR/NC
Weekly meetings: time and place to be determined.
Organizational meeting at 2:30 on Monday 27 March in Art 004
Entry Codes and other information available from John Palmer
Readings at: http://faculty.washington.edu/jpalmer/files/ARSeminar/

E-mail: jpalmer@u.washington.edu
Phone: 206-543-0706
Address: Psychology, Box 351525, U of W

Course Structure and Responsibilities

John Palmer will jointly lead the discussion with one discussant selected for the week. To help guide discussion, study and discussion questions will be distributed a week ahead. The study questions have answers that can be found in the reading while the discussion questions are open ended. Seminar participants will be expected to read the material in detail before the class meeting. Be sure to understand the study questions and think about the discussion questions. Everyone is expected to be a discussant once or twice. Discussants are expected to meet with John Palmer a few days before class and afterwards to draft a summary of the class discussion. This summary will be distributed to all participants by e-mail before the next meeting.

Schedule

Week	Topic
1	Organizational
2	Background on attention Dobkins & Bosworth, 2001; Hafter, et al., 1998; Treue & Maunsell, 1999
3	Background on gain control models of attention Reynolds, et al., 2000; Huang & Dobkins, 2005
4	Cueing paradigms Sagi & Julesz, 1986; Exp. 4 in Cutzu & Tsotsos, 2003; Hopf, et al., 2006; Niebergall, Tzvetanov & Treue (SFN abstract, 2005)
5	Two-target separation paradigms Bahcall & Kowler, 1999; Sagi & Julesz, 1985, spatial vision); Muller & Kleinschmidt, 2004
6	Tracking paradigms Intriligator & Cavanagh, 2001; Moore et al (2006, manuscript)
7	VSS break
8	Background on spatial resolution, localization and channel theory Levi & Tripathy, 1996; Graham, 1992
9	Example theory Baldassi & Verghese, 2005; Palmer & Moore (2006, manuscript)
10	Finale

References

Bahcall, D. O., & Kowler, E. (1999). Attentional interference at small spatial separations. *Vision Res, 39*(1), 71-86.

Baldassi, S., & Verghese, P. (2005). Attention to locations and features: Different top-down modulation of detector weights. *J Vis, 5*(6), 556-570.

Bonnel, A. M., & Hafter, E. R. (1998). Divided attention between simultaneous auditory and visual signals. *Percept Psychophys, 60*(2), 179-190.

Cutzu, F. & Tsotsos, J. K. (2003). The selective tuning model of attention: Psychophysical evidence for a suppressive annulus around an attended item. *Vision Res, 43*, 205-219.

Dobkins, K. R., & Bosworth, R. G. (2001). Effects of set-size and selective spatial attention on motion processing. *Vision Res, 41*(12), 1501-1517.

Graham, N. (1992). Breaking the visual stimulus into parts. *Current Directions in Psychological Science, 1*, 55-61.

Hafter, E. R., Bonnel, A. M., Gallun, E., & Cohen, E. (1998). A role for memory in divided attention between two independent stimuli. In A. R. Palmer, A. Rees, A. Q. Summerfield & R. Meddis (Eds.), *Psychophysical and phsyiological advances in hearing* (pp. 228-238). London: Whurr Publishing.

Hopf, J. M., Boehler, C. N., Luck, S. J., Tsotsos, J. K., Heinze, H. J., & Schoenfeld, M. A. (2006). Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proc Natl Acad Sci U S A, 103*(4), 1053-1058.

Huang, L., & Dobkins, K. R. (2005). Attentional effects on contrast discrimination in humans: Evidence for both contrast gain and response gain. *Vision Res, 45*(9), 1201-1212.

Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognit Psychol, 43*(3), 171-216.

Levi, D. M., & Tripathy, S. P. (1996). Localization of a peripheral patch: The role of blur and spatial frequency. *Vision Res, 36*(23), 3785-3803.

Muller, N. G., & Kleinschmidt, A. (2004). The attentional 'spotlight's' penumbra: Center-surround modulation in striate cortex. *Neuroreport, 15*(6), 977-980.

Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of v4 neurons. *Neuron, 26*(3), 703-714.

Sagi, D., & Julesz, B. (1985). Fast noninertial shifts of attention. *Spat Vis, 1*(2), 141-149.

Sagi, D., & Julesz, B. (1986). Enhanced detection in the aperture of focal attention during simple discrimination tasks. *Nature, 321*(6071), 693-695.

Treue, S., & Maunsell, J. H. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *J Neurosci, 19*(17), 7591-7602.

Seminar on the Spatial Resolution of Attention
Questions for Week 1: Background on Attention
John Palmer and Cathleen Moore
2 April 2006

The goal of this session is to understand key elements of attention research:

a. Introduce four behavioral paradigms and one physiological paradigm.

b. Discuss instances of perfect and imperfect selection.

c. Contrast theories of unlimited and limited capacity.

d. Consider hypotheses specifying the locus of selection and capacity limits as sensory, memory or decision.

Readings

Dobkins & Bosworth, 2001; Hafter, et al., 1998; Treue & Maunsell, 1999
Further background can be found in the first 5 chapters of Pashler's "The Psychology of Attention". See especially pages 13-28, 101-124, 176-191, and 217-219.

Study Questions

1. What is measured in each of the four behavioral paradigms:
    a. set-size effects with multiple stimuli,
    b. cue effects with multiple stimuli,
    c. cue effects with a single stimulus, and
    d. dual task vs single task performance?

2. How are attentional effects distinguished from nonattentional sensory effects? (Hint: consider the cueing experiments)

3. How good is selection in D&B's experiments? (Hint: compare cueing and set-size effects)

4. What assumptions are made for theories of unlimited versus limited capacity? How do the details of these assumptions differ in the set-size vs dual-task paradigms?

5. Why do D&B argue that their set-size effects are due to decision processes rather than capacity limits on sensory processes?

6. On what basis does Hafter et al. argue for a memory rather than sensory basis of their dual task effects? (Hint: last experiment)

7. How are attentional effects measured by T&M? Are they larger with two stimuli in the receptive field?

8. How do T&M argue that selective attention causes cell responses to increase for the relevant stimulus and decrease for the irrelevant stimulus?

9. How are the effects found by T&M accounted for by either selective attention to space or by selective attention to a feature (direction of motion)?

Discussion Questions

10. How can one establish that an effect is attentional?

11. Are the ideas in these papers well captured by two attentional phenomena: one for selection and another for capacity?

12. What are the pros and cons of using D&B's method of quantifying attention with threshold ratios?

13. In general, to what extent are effects of selective and divided attention due to sensory, memory and/or decision process? Would a complete theory have to allow such effects for all three kinds of processing?

14. What behavioral experiment is most like the T&M physiology experiment?

15. Suppose you wanted to do physiology, what are the pros and cons of using each of the discussed behavioral paradigms?

16. How would the behavioral and physiological measures have to be modified to allow a direct comparison of the magnitude of attention effects?

Seminar on the Spatial Resolution of Attention
Summary of Discussion for Week 1: Background on Attention
Iris Zemach, John Palmer and Cathleen Moore
11 April 2006

Papers

Dobkins and Bosworth (2001)
Hafter et al. (1998)
Treue and Maunsell (1999)

Discussion Summary

We began with the comment that the Hafter et al. study was less clear than the others. John Palmer mentioned that it was essentially a talk that appeared as a book chapter and that more background can be found in the earlier Bonnel papers (e.g. Bonnel & Hafter, 1998, P&P). The motivation to use the Hafter et al. paper was to introduce the dual task methods that will be used quite a bit in upcoming papers on the spatial resolution of attention. In particular, the Hafter and Bonnel line of work is much more concerned about the role of memory and decision in dual tasks than most other dual-task studies.

To begin, we discussed the benefits of considering a dual task from two modalities. When you use a dual task from the same modality such as two vision tasks, several things other than attentional capacity can limit performance. For example, gaze, location on the retina, etc. Whereas using a vision-audition dual task minimizes those effects so one can be more certain that we are measuring effects that are common to most task combinations. Hafter et al. also chose tasks where quite a bit of theory from SDT to rely on when comparing detection and identification.

The first experiment replicated a comparison between detection and identification that had been explored in their prior studies. Here, they pursued the possibility that transients can aid detection but not identification. Thus the difference observed for detection and identification may have to do with differences in the set of stimuli used rather than the kind of task in general. The use of a transient cue in detection may explain why detection thresholds have been lower than identification when SDT theory predicts that identification thresholds will be lower than detection.

To eliminate the transients, they used a task that always has transients so that the presence of a transient does not distinguish the signal. As a result, detection thresholds increased in this version of the task. The change had no effect on identification.

Hafter et al. tried several types of stimuli without differential transient cues. They settled on the following sequence: a brief presentation of a standard, an empty gap, and then either the standard again or a comparison stimulus. These experiments allow the strategy that the subject can memorize the standard. Under these conditions, both detection and identification showed dual-task deficits.

In contrast to this condition, they added a condition where the standard was varied from trial to trial. Now using a long-term memory standard was impossible. Under these "roving" conditions, performance was worse and there was no dual-task deficit.

Hafter et al. argue that under ideal conditions, there is no capacity limit for perception. The dual-task deficit they find may instead be due to interference between the two tasks' memory or decision processes. The take home message is that to examine the effect of attention on perception, one must make sure the role of memory is minimized.

Cathleen Moore (Penn State) pointed out by email that the use the terms "sensory trace mode" and "context coding mode" were more specifically suggestive of mechanism than the often used sensory vs. memory capacity limits. She wants to know the specific mechanisms that mediate these two modes.

Discussion then moved to Dobkins & Bosworth. They distinguish between limited- and unlimited-capacity models of attention in a search paradigm rather than a dual-task paradigm. They also interpret their results as consistent with unlimited capacity for perception. The set-size effects they observe are attributed to the many-to-one decision process that is necessarily a part of visual search tasks. Specifically, the introduction of additional stimuli introduce errors in the task as a whole even though the perception of individual stimuli is unchanged.

Roozbeh Kiani questioned how reasonable are the assumptions in the various models. The models assume both independence properties and specific noise properties. In particular, they assume normal noise with equal variance for all stimuli. If the noise is not normally distributed or has unequal variance it would change the magnitude of the predicted set-size effects. Some of this is quantified in Palmer, Verghese and Pavel (2000, Vision Research) and in Graham, Kramer and Yager (1987, Math Psych).

A more specific question is what models can be rejected. Dobkins and Bosworth calculate the prediction of a fixed capacity model that predicts much larger set-size effects than are observed. Such a model probably cannot be saved by variations in the noise assumptions. It has been rejected by several studies of set-size effects.

Another kind of model that can be rejected is the combination of some degree of limited capacity and a "high-threshold" assumption that eliminates any effect of decision on set-size effects. This model can be rejected on the basis of ROC experiments and the like (Palmer, Verghese and Pavel, 2000, Vision Research). Some of the newest studies try to estimate a parameter that describes the degree of limited capacity that is consistent with the set-size effects (e.g. McLean, 1999, dissertation). What cannot be rejected are models with either unlimited capacity or a small degree of limited capacity combined with a reasonable decision model.

Dobkins and Bosworth also investigated selective attention using a cue paradigm in the context of visual search. Such experiments allow one to evaluate the quality of selection which is at the heart of this seminar's topic. In general, Dobkins and Bosworth find effects with cues as well as with set-size manipulations.

Quantitatively, Dobkins and Bosworth find better performance in a cued multiple display than in a cued single display. If this discrepancy was in the opposite direction, one might think it was due to imperfect selection. In this direction, it poses a different kind of problem for the theory. If selection is perfect and the cue helps you to completely ignore the distracters then the ratio between thresholds in single and multiple displays should be 1. However, in the experiment it was less than 1. Somehow the presence of irrelevant stimuli must be facilitating performance. Perhaps this is related to induced motion effects by distant stimuli.

John Palmer mentioned that sometimes a subject can use unanticipated strategies to help them do the task. For example, in a set-size experiment using size judgments and a regular array of stimuli, a smart subject was visualizing multiple stimuli around a semi-circle as an arc and her performance improved with increased set-size because it improved the detection of deviations in the visualized arc. A solution was to randomize the positions of the targets so the subject wouldn't be able to see a systematic pattern across the stimuli. A similar thing might be happening in the Dobkins and Bosworth display where the performance was better with surrounding distracters.

What other paradigms get at the capacity issue in perception? The dual task and set-size experiments are two examples of experiments that address the capacity issue. Other examples are the successive vs. simultaneous paradigm and the redundant target paradigm.

The Dobkins and Bosworth experiment includes a manipulation of duration. Modeling this effect has its own complications. The effect of stimulus duration levels off at some point (the critical duration of Bloch's law). This duration effect has its own story that needs to be considered before modeling these effects in detail.

Elisabeth Hein (Penn State) pointed out by email that predictions similar to the Dobkins and Bosworth time-to-orient hypothesis can be made by a signal enhancement hypothesis. The ability to differentiate these hypotheses appears to depend on the details of the duration effect. She cited papers by Reinitz (1990, P&P) and Smith and Wolfgang (2004; JEP:HPP). After the seminar, John Palmer adds another paper by Smith, Ratcliff and Wolfgang (2004, Vision Research). Differentiating these hypotheses requires specifying more about how duration affects performance.

Discussion next moved to the physiological paper of Treue and Maunsell. The key result in this paper was to demonstrate a cueing effect on the firing rate of MT and MST neurons. In particular, there was a nearly 100% modulation of MST neurons when the cue disambiguated two stimuli that were both in the receptive field of the cell.

One possible worry is that the cue is close in time and space to the display. In fact the cue is also the target, it could cause a problem if the neuron is responding to the cue which is in the receptive field during the experiment. The authors discuss this in the paper an argue against any sensory effect of the cue.

Another concern about the paper is that the attention indexes used in the first and second experiments are difficult to compare. In fact, the authors also compute multiple indexes for the second experiment to provide a better comparison between the results of the two experiments. A more principled choice of "indexes" is needed. For example, might the threshold ratios used in Dobkins and Bosworth be better grounded in theory?

The discussion then moved to larger issues in attention. We began by discussing how to distinguish attentional and sensory effects. The preferred way used in several of these papers is to keep the stimuli the same and change only instructions and cues. Then the sensory input is exactly the same so any effects can not be due to difference in the sensory input.

Roozbeh Kiani asked what is the definition of attention. John Palmer gave an answer for what are attentional phenomena. They are phenomena that are influenced by voluntary control. This definition ignores phenomena that are due to exogenous cues and stimulus-driven attention more generally.

How can this definition be expanded to include stimulus-driven attention without including phenomena of adaptation and perceptual learning? John Palmer cited studies by Steve Yantis that showed that under certain condition you can ignore an exogenous cue that is supposed to drag away your attention. Therefore, it might exogenous cues may still be under a degree of voluntary control.

We talked briefly about how to do behavioral and physiological experiment that can be analyzed similarly so we can compare the magnitude of the attentional effects in neurons and in behavior. John pointed a problem with the Treue and Maunsell study in that the speed increment task using in this experiment is probably mediated by population coding and is not directly represented in individual neurons. The comparison may be a bit easier for contrast tasks where one can argue that the firing rate does represent contrast (see the Reynolds paper for next week).

Seminar on the Spatial Resolution of Attention
Questions for Week 2: Example Models of Attention
John Palmer and Cathleen Moore
5 April 2006

The goal of this session is to understand example mechanisms of the attentional effect on perception. Specifically, we consider multiplicative models such as contrast gain or response gain. Such a mechanism must be specified to define an appropriate measure for a spatial tuning function of attention.

Readings

Reynolds, Pasternak, & Desimone, 2000; and Huang & Dobkins, 2005

Study Questions

1. What is a contrast response function and how does modulating contrast gain and response gain differ?

2. What is the task in Reynolds et al. and how does this "filtering" task differ from those discussed last week? (Hint: in the filtering task one must ignore "foils" at the irrelevant location.)

3. What are the details of the behavioral task in Reynolds et al.? How do they differ from a standard contrast discrimination task such as used at the heart of H&D? Might these differences undermine the interpretation of results?

4. Can we estimate three effects of attention from Reynolds et al.? These would be a large effect for contrast gain; a small effect for the spontaneous firing rate; and a small or absent effect for response gain.

5. According to H&D, why do contrast gain models predict that attention improves behavioral performance for low and moderate contrast and worsens it for high contrast?

6. What is the deltaContrastThreshold–vs–Contrast (ΔCvsC) function that is measured in H&D to summarize contrast discrimination? (See Legge, 1981 for a good example – on web site).

7. What are the traditional accounts for the effects of the base contrast in contrast discrimination? (Hint: spatial uncertainty or transducer nonlinearity for the "dip", contrast adaptation for the rise with larger contrasts.)

8. What results do H&D obtain for their ΔCvsC function?

9. What results do H&D obtain for their various attentional manipulations? Does the dependence on the central task complicate the interpretation of these results?

10. What are H&D's conclusions regarding contrast and response gain?

Discussion Questions

11. Given a contrast gain model, what is an easy way to measure the change in gain? (Hint: Detection threshold ratios equal gain ratios). How can this fact be used to measure the spatial tuning function for attention.

12. Why do contrast response functions asymptote at different levels for preferred vs. nonprefered stimuli? (Hint, look at Geisler & Albrecht, 1995 – on web site).

13. What are the details of the behavioral model outlined by H&D? (Hint: look at Ross & Speed, 1991 – on web site).

14. How can models based on the contrast response function be reconciled with signal detection theory that assumes performance is limited by the noise in the system? Specifically, what implicit noise assumptions are made the by the simple theories described in H&D?

15. Is there any possibility of nonperceptual contributions to the dual-task effects in H&D?

16. H&D attribute differences in the literature to the use of different "central tasks" in the dual-task paradigm. What tasks and what controls are needed to sort out the role of the "central task"?

17. What human behavioral study is most like the Reynolds et al. study? How would these studies need to be modified to allow a comparison of the magnitude of effects on the physiology and on behavior?

18. How can one adapt the filtering task used by Reynolds et al. to measure contrast discrimination functions as done by H&D? This would avoid the use of dual tasks. What are the pros and cons of this alternative approach?

Seminar on the Spatial Resolution of Attention
Summary for Week 2: Background on gain models of attention
Roozbeh Kiani, John Palmer and Cathleen Moore
18 April 2006

Papers

Reynolds, et al., 2000
Huang and Dobkins, 2005

Overview

Our discussions went back and forth between the shortcomings of the tasks and models used in the two papers and ways to improve them. At the end of the session, we tried to develop a unified framework to understand the results of different attention papers we have studied so far.

Discussion of the models

Why model attention as a gain change? Although there are alternative hypotheses for attentional mechanisms, many of them are more-or-less equivalent to a multiplicative mechanism. This is why gain-type mechanisms are so popular in modeling the effects of attention. Nevertheless, there are also alternatives that are not multiplicative. One of them is an "all or none" model. In this model there is a probability for each stimulus not to be processed but when it is processed it is processed completely.

In Huang and Dobkins' paper, the presence of the two gain effects is inferred by the lack of a good fit to the data by either of the gain models alone. The equation that they use to fit their data is quite complex especially because of the exponents in the equation (Ross and Speed 1991). Unfortunately, their data does not have enough power to distinguish the change in different parameters of the model. Because of the compression that can be achieved by the exponents, the equation fits the data well even when they force C_{50} (the parameter representing the contrast gain) to be the same for the full-attention and poor-attention conditions. As a result, the reliability of the reported changes in C_{50} and R_{max} (the parameter representing the response gain) can be questioned.

Another question about their modeling is why the contrast gain model should not cause a vertical shift in addition to a horizontal shift? John Palmer had tried to derive the prediction of these gain models and was concerned that the reported predictions of the contrast gain model were incorrect.

There is a strong assumption in inferring the TvC curves from the contrast response functions. It is assumed that detection of a change in contrast relies on a fixed change in the response. This assumption follows the work of Speed and Ross. However, the variance of neural responses increases with the mean response. As a result, one might expect that the required increase in the response for reaching the threshold be larger when the pedestal gets bigger. There have been some modeling attempts for inferring the TvC curve from the contrast response function with the assumption that response variance changes as a function of pedestal contrast (see Geisler and Albrecht, 1995). The predicted TvC curve does not show much qualitative difference compared to when we assume response variance remains the same across different pedestal contrasts. However, the quantitative differences may be important for fitting the data to explore the effects of attention. The simple contrast response function that was used by Huang and Dobkins is just a first step toward more realistic modeling.

A typical TvC curve has a dip: the threshold for pedestal contrasts that are slightly larger than zero is smaller than the threshold for a zero contrast pedestal. Because the latter is the detection threshold, a typical TvC curve indicates that the discrimination threshold can be sometimes lower than the detection threshold. One early hypotheses that predicted the dip ascribed it to spatial uncertainty effects (Pelli, 1985). Later studies, however, argue it is due to processing nonlinearities. The other interesting feature of the TvC curve is the rise of threshold for high pedestal contrasts. This may be due to contrast adaptation (see Bowen, 1997, Vision Research). Our understanding in this seminar is hampered by our limited knowledge of the theory behind these TvC curves.

Alternative gain models may soon be appearing in the literature. Reynolds and Desimone in a recent review suggested a combination of response gain and contrast normalization that may result in effects similar to a contrast gain model.

Discussion of the tasks

The tasks used by Reynolds et al. and Treue and Maunsell were filtering tasks. Such tasks have not been extensively explored by psychophysicists. For example, Reynolds' study used foil stimuli that the monkey must ignore and instead make its decision based on only the relevant stimulus. Foils usually make the task more difficult. In the Reynolds' paper, monkeys appear to be very conservative. They adopt very high thresholds for responding and tend to ignore targets at the cued location even at contrasts that are well above the typical detection threshold. John Palmer is running similar experiments and believes that subjects are affected by the occasional appearance of the target at the non-cued locations. Apparently, it takes effort to ignore these distracters. Perhaps a bit like the Stroop effect.

The task used by Huang and Dobkins has a twist that make the interpretation of the results more difficult. The two gratings in the contrast discrimination task were surrounded with high contrast frames that were presented simultaneously with the gratings. The purpose of these frames was to make subjects certain about the presentation of a target when the gratings had low contrast and were hard to see. the frames were also meant to reduce the advantage of high contrast gratings in grabbing the subject's attention. However, such high contrast frames can change the visual response to the contrast of the grating stimuli through mechanisms such as contrast adaptation. Although the frames were present in both full- and poor-attention conditions, they complicate the quantitative modeling.

Improvement of tasks for future studies

It would be interesting if in electrophysiology experiments one could use a task similar to what is used by H&D. There are least two obstacles in using H&D's task. First, training an animal for a dual task is difficult. Second, simultaneous presentation of the two relevant stimuli in the inside and outside of the receptive field makes the interpretation of the data somewhat more difficult. In this task, both stimuli are relevant because they must be compared with one another. Perhaps it would be better to use a sequential presentation?

It would be interesting if we could avoid dual tasks. Perhaps we can modify Reynolds' filtering task to contrast discrimination. Following his design, use cues to mark one location as relevant. Require the discrimination of a contrast increment at the cued location. Increments at the irrelevant location must be ignored. There are several choices of how to present the standard and comparison stimuli: simultaneous as in H&D, sequentially as in 2IFC, or a method of single stimuli. Now one can measure the effect of the contrast discrimination stimulus in the RF for both relevant and irrelevant conditions as Reynolds did for contrast detection. Does the modulation vary with the change from the standard or just the current contrast?

Is a unified framework for attention possible?

In the last meeting, we reviewed papers that suggested performance of subjects in the cueing paradigm is not limited by attentional resources for the visual processing of targets. In those experiments, set-size effect can be accurately described by errors at a decisional stage. If the capacity of attention is not limited, why should the performance of the subject decline in a dual task similar to Huang and Dobkins? Moreover, why should the type of the central task, RSVP or pop-out, affect the performance of the subject by different amounts?

One possible explanation is based on the resources that different tasks require after the visual processing stage of the stimuli. It is possible that both the stimuli of in the central RSVP or pop-out tasks and the peripheral contrast discrimination task be fully processed visually but the central and peripheral tasks compete for resources in decision or memory encoding.

Another possibility is that subjects modify the size of their attended region depends on the task. In tasks that study the set-size effect (e.g. Dobkins and Bothworth), the subjects may expand their attended region to include all of the stimuli. But in a dual task (e.g. Huang and Dobkins), they may shrink the attended region around the central task. However, this hypothesis does not predict that subjects shrink their attended region around the central task if there is an unlimited capacity of attention for the visual processing of stimuli. It seems that the first hypothesis which emphasizes the decision and memory resources needed by tasks is a bit more successful in reconciling the various psychophysical studies.

Seminar on the Spatial Resolution of Attention
Questions for Week 3: Spatial Extent of Cueing Effects
John Palmer and Cathleen Moore
12 April 2006

Goals

Examine attempts to measure the spatial extent of visual selective attention using a cueing paradigm. Primary questions: What is the spatial profile of the attended region: a single-peaked spatial profile or an opposing center-surround spatial profile? What is the spatial extent of the attended region?

Readings

Sagi & Julesz (1986)

Cutzu & Tsotsos (2003)

Hopf et al. (2006)

Niebergall, Tzvetanov & Treue (2005), NEW ABSTRACT on web site

Study Questions

1. Is the way S&J plotted their data essentially equivalent to plotting P(hit)?

2. How did S&J quantify the spatial extent of their effects? They suggest that the diameter of the attended region is about a 0.75 fraction of its eccentricity.

3. How did S&J test if their effects were attentional?

4. The results of S&J and C&T, Exp 4 appear to contradict one another. One shows a monotonic fall off of performance with distance from the cued location and the other shows a nonmonotonic effect. What are the possible differences in the two studies that may account for the different results?

5. In Hopf et al., the observed ERMF is interpreted as falling off with target-probe distance with an additional drop specific to the shortest distance. What sensory factors are correlated with target-probe distance?

6. How did Hopf et al. try and separate sensory and attentional factors in Exp 2?

Discussion Questions

7. S&J tested if their cueing effect was attentional by comparing conditions where subjects were instructed to judge a peripheral stimulus or a central stimulus. They found cueing effects near the peripheral stimulus only when the peripheral stimulus was being judged. Is this a convincing control?

8. Is the dual-task design important in S&J? Why not just cue a location for detection like C&T? Pros and Cons?

9. Given our discussion of dual-task effects, are the results of S&J likely to be general to other dual-task combinations?

10. Might the pooled-false-alarm design used in S&J and C&T lead to inaccurate interpretations? For example, suppose one reduces the criterion for a "Yes" response at the cued location and nowhere else. How could this strategy be detected?

11. Can we improve S&J by controlling response bias using a localization version of the detection task? For example, one could ask subjects to respond whether the flash was on the left side or the right side of the display. A possible further improvement would be to estimate a detection threshold at each spatial location. What are the pros and cons of such changes? (See abstract from Niebergall, Tzvetanov & Treue, 2005.)

12. Is there any way to modify the cueing paradigm to obtain larger effects? How about adding foils as done by Reynolds, et al. (2000)?

13. For Hopf et al., what assumptions are necessary to interpret the variation of probe ERMF with target-probe distance as a pure measure of attention?

14. What would have to be added to Hopf et al. to get a behavioral measure relevant to selective attention? (Hint: measure the detectability of the probe)

15. What needs to be done to relate the magnitude of a probe ERMF effect to the magnitude of a behavioral effect of the probe (e.g. probe detection)?

16. The stimuli in Hopf et al. seem "crowded" relative to the other studies. How might this change the results?

17. What can be done to determine if these cueing effects are determined by the objects visible in the scene rather than by their spatial separation?

Seminar on the Spatial Resolution of Attention
Summary for Week 3: Spatial Extent of Cueing Effects
Heather Knapp, John Palmer and Cathleen Moore
25 April 2006

Readings

Sagi & Julesz (1986)
Cutzu & Tsotsos (2003)
Hopf et al. (2006)
Niebergall, Tzvetanov, & Treue (2005 abstract)

Discussion of Sagi and Julesz

Prompted by Michael Lee, we began with a clarification of the stimuli, decision task, and response used in the Sagi and Julesz study. In this experiment, a participant simultaneously viewed two differently-oriented bars, one at fixation and one at 4 degrees eccentricity. S/he is instructed to attend to only one of the targets (central or peripheral, blocked). On 66% of trials, following a short SOA a small probe dot appeared at one of 12 positions adjacent to one of the bars for 10 ms. The entire display was then masked with a random dot pattern to which are added two carats at the (former) locations of the bars. The timing of the carat onset was manipulated to control task difficulty. This was chosen instead of contrast or size manipulations. As the SOA between the target and mask is reduced, the task becomes harder. The participant has two tasks: (a) discriminate whether the orientation of the bar that appeared at the instructed location was vertical or horizontal; (b) report the presence or absence of the test flash. S/he did so via button presses. The primary dependent variable is the probability of detecting of the probe dot. The primary manipulation is the distance of the probe dot from the attended bar. S&J anticipate that dot detection will be greatest in the region of the attended bar, and will decrease with distance from the attended location. This is the result they got.

Roozbeh Kiani inquired as to whether the lower left panel of Figure 2 is strange. This panel depicts probe detection (p[detect] - p[FA]) as a function of the distance between the peripheral bar and the probe dot. This panel differs from the one above it in that it depicts trials on which participants attended to the central bar, rather than the peripheral bar. This manipulation was used by S&J to demonstrate that the effect seen in the upper panels is attentional rather than sensory. In both upper and lower panels, detection is plotted as a function of the distance between the peripheral target and the probe dot. In the upper panels, attention is fixated on the peripheral target; in the lower panels; attention is fixated on the central target. Differences between the two conditions regarding the spatial proximity of the peripheral target to the probe dot must be driven by attention directed to the location of the peripheral target and not the physical presence of the peripheral target. This is because the stimuli were identical in the two conditions.

The lack of a systematic relation between probe detection and the probe's distance from the (unattended) peripheral target is both predicted and critical to the theory. On the other hand, this is based on the assumption that spatial attention can be controlled, and not unwittingly captured by the presence of either the peripheral target or the probe dot. If any such exogenous cuing is present, it can't be seen in the lower panel of Figure 2 because performance is at floor. One way to fix this is to make the flash probe more detectable. In sum, we wonder whether there will be an effect of stimulus driven attention in an improved version of the central control of S&J.

Roozbeh Kiani asked why the test flash isocontours depicted in Figure 4 are not symmetrical around the peripheral target. John Palmer pointed out that these were created from the data such as that depicted in Figure 2a (left), although the exact procedure is unclear. Symmetry is present along one axis and not the other. Scott Murray pointed out that similar maps exist for crowding, and these also contain asymmetries (Toet & Levi, 1992, Vision Research). The discussion then segued into what would happen if the central target was removed. John Palmer cited a sentence in the paper that argued that the isocontour is more symmetric when the central target is removed. Perhaps this is because the central target serves, to some extent, as an irrelevant distractor that skews the shape of the enhancement area. When it is removed, one has a simpler depiction of the shape of the enhancement area around the peripheral target.

Discussion of Cutzu & Tsotsos

At this point, we moved on to the Cutzu & Tsotsos paper. We focused on Experiment 4 because the focus of this session of the seminar was on the cueing paradigm and not separation effects between two targets. That will be taken up next week. In this experiment, the participants receive a cue that consisted of a light gray disc for 180 ms. The cue was followed by a 100 ms–duration test image in which 11 more discs appeared on the screen in a circular pattern, so that 12 discs were present. Each was filled with a T or an L, oriented in various directions. In half the trials, all the letters in the test display were the same. In the other half, one disc contained a different letter from the others. The test image was followed by a random letter mask. The participant's task was to indicate whether a target letter (a T or an L) was present in the test display. As in the S&J paper, the effect of attention was measured as a function of the distance between the cue and the target. Unlike S&J, these authors find that performance does not decrease monotonically with distance. Instead, it is strikingly bad when the target is immediately adjacent to the cue and then improves gradually with distance. This result is consistent with Tsotsos' selective tuning model. This model predicts an excitatory-center, inhibitory-surround pattern of attentional modulation of physiological activity.

Heather Knapp commented that the way the paper is written, it is unclear whether the task in Experiment 4 is target detection or oddball detection. For example, in Section 2.2 the task is described as reporting the existence of a target (target detection), whereas in Section 2.6.1 the task is described as the detection of the odd letter, either a T among Ls or an L among Ts (Target identity not blocked). John Palmer pointed out that the participants seem to have been instructed to indicate the presence of an odd item. Heather Knapp said that these tasks might actually engender different strategies on the part of participants, as a fixed target detection (i.e., "report whether a T is present") requires the mapping of an image onto a mental 'template' whereas oddball detection can proceed by a filtering operation, in which one item is perceived to be different from the others. This may not be a true identification of the letter. Instead, a distinguishing set of features may be used to perform the discrimination. This latter strategy is what Pashler sought to avoid by employing a highest digit identification task instead of detection in his simultaneous/successive alphanumeric character identification experiment (Pashler & Badgio, 1987, Attention & Performance). Some think that paper has the best evidence of parallel, unlimited capacity processing of characters.

Iris Zemach asked whether it was possible that the test duration of 100ms allows for an attentional shift to the opposite side of the visual field. She pointed out that participants can use this as a strategy, which would obscure effects of cue-target distance. For instance, it might make the detection of items on the other half of the visual field better, and thus items on the same side worse. If this was mixed with a monotonic decrease in performance without a shift, then the combined result might be the observed non-monotonic pattern. This is a possible explanation of the results. John Palmer agreed, but said that there have traditionally been two views of attention shifts: Some people believe that attentional shifts can be very quick, on the order of 10-20 ms. However, studies by Sperling, Duncan, and others have revealed that the shortest time in which attentional shifts can take place is on the order of 200 ms or more, similar to making an eye movement (see Moore, et al., 1996, Psychonomic Bulletin and Review and the papers it references). But, this may not undermine the proposed explanation. Perhaps the participant guesses a location in response to the cue that is either at the cued location or on the opposite side. That would result in a similar mixture with a relatively longer time for shifting attention.

At this point, John Palmer briefly interjected that there is a possibly odd aspect in the design of this experiment. The authors say they counterbalanced by separation between cue and target (i.e., 0, 1, 2…6 intervening letters). Did they mean to say counterbalance separation rather than target position? For example, there is one opportunity for a cue-target distance separation value of 0 (the cue and the target appear at the same location), but there are two opportunities for a cue-target distance separation value of 1 (the target appearing adjacent to the cue on opposite sides). Balancing separations results in unequal probabilities for target location. If so, this may underlie the result.

Returning to Iris' question, he reiterated that implicit in the question is whether this process is serial or parallel. In the paper, the authors note that letter identification tasks such as these are often taken to be serial processes. For simple stimulus discriminations, there is considerable evidence that with four widely separated stimuli, there are little if any capacity limitations. Results are much more mixed for letter identification. And for word identification, limited capacity is the rule. John Palmer himself believes that rotated T-L discrimination may have unlimited capacity with up to 4 items (Palmer, 1994, Vision Research, next to last experiment). But he laughed that he is likely the only person who does so for the rotated T-L task. But, this point may be irrelevant because in fact C&T used larger set sizes and more crowded displays.

A related issue is the likely complications from crowding in C&T's displays. Iris Zemach asks whether crowding has to do more with character size, size of the display, or number of items in the display. We didn't know the answer.

Contrast of Sagi & Julesz versus Cutzu & Tsotsos

Next we turned to a discussion of ways in which the S&J and C&T experiments were different from one another. It was hoped that by winnowing in on these differences, we can figure out which differences were most critical.

Parameter	**Sagi & Julesz**	**Cutzu & Tsotsos**
Paradigm	dual task	single task
Task	detection	discrimination (odd T-L item)
Critical feature	contrast	shape
Display	empty	filled
Cue type	blocked	exogenous
Cue uncertainty	highly uncertain	one of 12 known positions
Cue relevance	100% relevant	irrelevant
Visual field	quadrant	whole circle
Capacity limitations	less	more

The conversation about these differences quickly turned to one of capacity limits. Roozbeh Kiani inquired as to whether the results of these two experiments are compatible with unlimited-capacity models of attention. John Palmer said that he doubts there are any data in these papers that explicitly addresses this point. The S&J paper does not vary the number of items in the display and did not have a single-task control. However, next week we will discuss capacity limitations more, with respect to dual-task experiments that do have some single-task controls (Bacall & Kowler, 1999) and have clear indications of limited capacity.

Roozbeh Kiani raised the hypothesis that perceptual processes at the spatial locus of attention have unlimited capacity but perceptual processes outside of the locus of attention have limited capacity. If so, there be a monotonic decrease of detection once outside of that region, rather than a Mexican hat function? John Palmer responded with a further complication. Selection and capacity are two different things: In theory one might select something yet still have unlimited capacity outside of the selected region. For example, one might weight more highly the information from some areas and weight less the information of other areas. In short, selection effects can occur without capacity limits. He argued these papers are more about selection than capacity.

Another aspect of selection is whether it occurs before the relevant perceptual process or after. John presented a confused version of the following observation which is hopefully clarified here. Sometimes early selection may be preferred over late selection because it can eliminate the interference between conflicting items. For example, suppose one was to discriminate Ts and Ls and thus the key perceptual process is the identification of T vs. L. Suppose a T was presented at the relevant location, and a L was presented at a nearby irrelevant location. If processing occurs only at the relevant location, then no response competition can result. In contrast, if selection is after perceptual identification, then the two percepts demand different responses and plausibly selection is then more difficult due to response competition. Perhaps some of the differences we are observing in the reviewed studies are due to whether the attentional selection was early vs. late and whether the particular task allowed response competition.

To make things a bit more specific, John Palmer briefly described Tsotsos' selective tuning model. In it, control of attention is determined by both stimulus-driven and top-down processes. Imagine multiple layers of units in a network that has stimulus inputs a the "bottom" and attentional priority inputs at the "top". First, assume the stimulus activates bottom-layer units in the neural network and then activation propagates in a feed-forward manner toward the output layer. Units in higher levels tend to have larger receptive fields. This results in an pyramid of activated units. Second, assume the attentional control signal activate a single relevant unit at the top that indicates a particular spatial location. These top units have relatively fine spatial resolution. The control signal propagates down the network as feedback relative to the stimulus information. Third, these two kinds of signals combine in a salience value using a winner-take-all process that selects the most active unit in each particular layer. This system will evolve iteratively over time and result in a recursive pruning of a "pass zone" of attended stimulus information. This zone is surrounded by a zone of pruned or "inhibited" units.

Roozbeh Kiani wonders whether this pattern of active-center and inhibited-surround is an unavoidable consequence of the model. In other words, if in order to attend to a place, do we always have to NOT attend to nearby places. John Palmer suggests that the model can be constructed differently. As an example he points out that it would be possible to attend to one place and NOT ATTEND to everything else. If the model were constructed this way, he thinks it might predict a monotonic decrease of perceptual enhancement. Moreover, one might get a different effect with the presence of foils. In the presence of foils a participant's goal is to inhibit everywhere but the attended area. However, in the C&T paper there are multiple filled spaces that contain relevant content, so these cannot be conceived of as foils. Maybe if there were real foils the other result would have been obtained?

Given that all the experiments besides S&J have "placeholders", Heather Knapp wonders whether it is the place-holding itself that matters, or the informational content of the placeholders. In other words, are attentional effects due to the presence of objects in the visual field, or due to the relevance of those objects to the task? John Palmer believes that the primary effect comes from changing the perceptual organization of the display. This is probably related to crowding. Cathleen Moore has argued that placeholders seem to increase uncertainty about location under some conditions. An increase in the density of placeholders is associated with an increase in errors if localizing attention (see the Intriligator and Cavanagh (2001) in two weeks).

Iris Zemach notes that the Treue and colleagues abstract discusses a cleaned-up version of the Sagi and Julesz paper in which the task is detection but the display is empty. They also get a monotonic function. She wonders whether there are any discrimination tasks (like the C&T experiment) with empty displays. This would be a way to distinguish whether the differential result hinges on task or display factors (see the Sagi and Julesz, 1995, Spatial Vision paper next week).

More generally, one can ask how the Niebergall, Tzvetanov, & Treue (2005) abstract relates to the table given above. They report monotonic results consistent with S&J. But like C&T, they use a single task, exogenous cueing, irrelevant cues, cue 12 possible distances (locations?), and cue the whole circle of possible displays. On the face of it, this rules out several entries from our table of differences. Those that remain and seem most plausible are:

a. detection versus discrimination tasks,
b. contrast versus shape,
c. empty versus filled space, and

d. unlimited- versus limited-capacity perceptual tasks.

Discussion of Hopf, et al. paper

Lastly, we had a brief discussion of the Hopf et al. paper. In these experiments, participants viewed a filled display of 9 “C” shapes, and discriminated the orientation of the single red C. On half of the trials, a white probe appeared 250 ms after the search frame onset. The probe was a high contrast white circle around a particular one of the C stimuli. Event Related Magnetic Field potentials (ERMF) were measured in response to the probe. The strength of the ERMF potentials were found to vary with target-probe distance in a non-monotonic fashion. Specifically, the location next to the cued location resulted in a lower response than either the cued location or two away from the cued location. This qualitative pattern appears consistent with the center-surround profile.

Roozbeh Kiani was troubled by the lack of consistency in the functions depicted in this paper and the Cutzu & Tsotsos paper. John Palmer agreed. But he argued that one might expect such variability in the behavioral results as well. Thus it is desirable to obtain behavioral data from the conditions of the Hopf et al. experiment. Unfortunately this is inherently difficult: In order to get a neural signal, the probe must be supra-threshold and such a strong signal will always be detected. Thus it is impossible to get probe detection behavioral data below ceiling. Scott Murray pointed out that the same problem holds for fMRI.

This was followed by a short general discussion of the ERMF procedure, including wondering what is actually being measured (what are influx and efflux exactly), why certain electrodes are used, why this particular time window was used. We don't know.

Related email and phone discussions

Cathleen Moore agreed that empty vs. filled space may have an important role in determining these effects. But not everything is consistent with this hypothesis. In particular, consider closely the Bacall and Kowler paper next week. Should not the neighbor errors be the primary source of error? (See also Strasburger's (2005, J. of Vision) paper on the role of attention in crowding.) Also, should not better cueing the location of the relevant information reduce errors? Neither of these conjectures are fully satisfied.

Cathleen Moore and John Palmer agreed that we should be on the lookout for differences between experiments with foils and without. It may be that "inhibitory" effects are more common with foils or foil-like characteristics.

Cathleen Moore and John Palmer agreed that something like "two-target interference" (see next week's questions), may be going on in some of these studies. This seems particularly plausible for the dual-task experiments in next week's readings. This two-target interference is a complication distinct from the usual capacity limits that is found in search experiments with multiple targets and in the attentional blink paradigm which focuses on interference between two targets at different temporal lags.

Seminar on the Spatial Resolution of Attention
Questions for Week 4: Separation effects in dual tasks
John Palmer and Cathleen Moore
19 April 2006

Goals

Examine measures of the spatial extent of visual selective attention using separation effects in dual-task paradigms. Primary questions: Is there an effect of separation on performance and does it increase or decrease with separation? What is the spatial extent of such effects?

Readings

Bahcall & Kowler, 1999
NEW: Sagi & Julesz, 1985, Spatial Vision, Fig 4, added to web site readings
Muller & Kleinschmidt, 2004

Study Questions

1. How do the identification tasks measured by B&K differ from the detection and discrimination tasks discussed earlier in the seminar?

2. In B&K, how are neighbor errors analyzed? Why do they interpret errors as being due to both mislocalization and misidentification?

3. How do B&K argue against nonattentional accounts?

4. In S&J Figure 4, there are no effects of separation except for small separations where the effect may be due to masking or crowding. This conflicts directly with the results of B&K. What differences in the studies might account for the difference in results? (Hints: intervening distractor items, task, etc)

5. How did M&K modify the typical cueing experiment to facilitate imaging?

6. M&K report only behavioral data for the cued condition and not their other controls. Should they be criticized for this omission?

Discussion Questions

7. Are there non-perceptual accounts for the separation effect in B&K? How about a perceptual account that is unique to making multiple letter identifications compared to a detection task?

8. How might one quantify the size of the separation effect in B&K?

9. B&K use a dual task but say little about the deficit of a dual task relative to a single task control. Single task data is presented in Figures 7 and 8. Even with large separations, there appears to be a large deficit for the dual task. This appears to contrast sharply with the unlimited capacity results of several of the other studies we have read. This might even be a "serial" task. Implications?

10. How might these dual-task effects be related to the "two-target effect" in visual search (e.g. Duncan, 1980, see Pashler's book p. 124-126, 229-230) or the interaction between targets in the attentional blink (Shapiro, et al., 1994; see Pashler's book p. 237-239)? Accounts of these effects typically depend on memory encoding and other processes distinct from capacity limits on perception.

11. B&K and S&J have very different, largely implicit interpretations of the processes that underlie their tasks. One assumes parallel processing and the other assumes serial processing. Implications? Tests?

12. What are the possible resolutions of the apparent conflict between the results of B&K and S&J?

13. M&K measure change in blood flow for ROIs that are retinatopically nearby, far or at the cued location. The reported data are for time after the cue but before the stimulus. What are the potential differences in interpreting these effects of the cue on blood flow compared to the effect of the relevant stimulus on performance?

14. M&K obtain results consistent with a center-surround spatial profile for V1 and a monotonic profile for V2, VP and V4. What does this diversity of results suggest about mechanism? How do their ideas about receptive field size relate to the differences between brain regions?

15. What minimal behavioral conditions would be helpful to add to the M&K experiment?

16. What would M&K have to add to their study to allow a quantitative comparison between behavior and imaging data?

Seminar on the Spatial Resolution of Attention
Summary for Week 4: Separation effects in dual tasks
Michael Lee, John Palmer and Cathleen Moore
2 May 2006

Readings

Bahcall & Kowler, 1999
Sagi & Julesz, 1985
Müller & Klinschmidt, 2004

Discussion of Bahcall & Kowler

John Palmer described the complexity of the task used by Bahcall and Kowler compared to previous tasks we have reviewed. The subject had to identify the two targets. Specifically, the subject to make two responses out of 26 possible choices from a stimulus which involved 24 letters. The target stimuli were from the same set as the distracters in the cueing and masking frames.

What is a dual task? John Palmer defined a dual task as requiring multiple responses, regardless of whether those responses were the same task. As an example of a possible single task that had two responses, a detection task with a confidence interval would not be considered a dual task. Scott Murray asked whether an attentional blink would be a dual task. John said he considered it a dual task.

Iris Zemach asked about the role of the color cues in the task. The color cue signaled the subject as to which two of the 24 targets was relevant to the task. Responding to any other stimulus would lead to errors. Thus, unlike some of the previous cueing studies, the distractors are also foils in the sense of being related to incongruent responses.

We then had a discussion about cueing methods. Heather Knapp brought up the possibility of trying the experiment without the cues or distractors with just two targets in an empty field to see what possible effect these elements have on attention. During cueing, critical frame, and masking, there are 22 irrelevant stimuli on the screen at a time. These may be potential distractors that increase the load of processing. This brought up the comparison between empty and full fields. Heather also brought up the possibility of cueing the subject with dots instead of letters.

We discussed Section 3.3 which describes a condition with cues in the critical frame only. In this condition, all the letters were green in the pretrial and masking frames. Only during the critical frame was the relevant letter distinguished by a red color cue. Cueing was less effective in this condition, but subjects could still perform the task.

We then moved on the results of the paper. John Palmer thought it was convincing that all the graphs showed increases in performance with increased distance. However, he wanted Bahcall and Kowler to quantify the effect. For example, John would have liked to see their psychometric function of duration used to estimate a threshold duration. Then the threshold duration can be considered as a function of separation. There is enough data reported to attempt such an analysis.

We next turned to a discussion of the neighbor errors. These errors may tell us something about the separation effect that is absent from the other papers we have looked at. With a neighbor error, instead of correctly identifying the target letter, subjects reported the letter that was directly adjacent to the target. This was an interesting finding, because the subject should not have been attending to the green letters, however information about these letters was somehow being confused with the targets and that resulted in errors. Heather Knapp commented that is ironic that neighbors are interrupting the identification process, as the ability to identify these targets should be poor. John Palmer said that the selective attention for this task is imperfect as seen by the neighbor errors, and the error is related by the amount of separation between the targets. The likelihood of neighbors causing an error depended on the neighbors proximity to a target. Heather offered a theory that when the targets are close together, the subject has difficulty in separating the identity of the target and distracter. However, when the targets are far apart, is it easier to separate the target from the distracters.

A possible way to reduce the number of neighbor errors would be to use a shifting attention strategy. When the targets are close together, one directs attention to that region. But, when they are far apart, then one can direct attention to the two relevant regions in turn. This might minimize neighbor errors when the two targets are far apart. However, one might also expect more "reversal errors" when the targets are directly adjacent to each other. Such reversal errors were very infrequent. Also, the point was made that on trials where there was a one distractor in between the two targets, the subject made an equivalent number of errors choosing the distractor in between the targets as choosing a distractor outside of the targets. Again, a difference is expected from the attention shift hypothesis.

Neighbor errors are not possible with empty space, and so it is possible that the neighbors themselves are creating the center-surround spatial profile seen in this and other experiments.

We briefly mentioned the possibility of serial processing occurring in this task.

Heather Knapp asked why there was a difference in the hemifield performance. This suggests some anatomical differences in stimulus location. We did not know what to conclude from these differences, Scott Murray mentioned that it would be more understandable if there were right-left hemi-field differences, but these were up-down differences.

Heather Knapp asked how closely related this attention effect is to the true center-surround model found in the visual system. John Palmer thought these are simply metaphors. This is similar to other attention metaphors of a spotlight, gradient, or zoom lens. In contrast, Tsotsos' network model actually predicts the center-surround profile from basic principles.

One unresolved detail: we could not determine whether the subjects were responding by keypress or vocalization.

Discussion of Sagi & Julesz

We went over the task used by Sagi and Julesz. The stimuli were T & L as opposed to the 26 possible letters of Bahcall and Kowler. Two letters were displayed at a time and the subject had to make a same-or-different response. The separation between the targets was varied and a mask which contained neither target letter was displayed after target presentation. The time between the target presentation and display of the mask, known as SOA (stimulus onset asynchrony), was kept at a level to allow a performance of 80-90%. They found in Figure 4 that besides a decrement at 2 degrees, the performance did not improve with increase in separation distance like the Bahcall and Kowler paper. Michael Lee asked if the experiments were truly comparable because they used different measurements to describe the separation distance. While Sagi and Julesz uses visual degrees, Bahcall and Kowler used circular degrees. John reassured us that the stimuli separation distances were comparable because of a similar design and eccentricity.

Bahcall and Kowler did not vary the location of the placeholders in their task, and thus controlled the masking effects due to location.

Sagi and Julesz conducted an additional experiment looking at the effects of one or two targets in Figure 5. In this experiment, the subject had to identify a single target vs. discriminate between two targets. The results show that performance for two targets is worse than performance on a one target identification task. It is difficult to compare these results in Figure 5 because the response requires different attentional needs. The effect could be explained by the differences in the tasks and not an effect of one vs. two stimuli. We hypothesized that a reasonable comparison could be done by showing two stimuli, where one of them is irrelevant, and measure performance at identifying the relevant target.

We then began to discuss the differences of serial and parallel processing. In serial processing, if there were two targets to identify, attention would be given to one target, then shifted to the other target. In parallel processing, both targets would be processed simultaneously. A classical paradigm to differentiate between parallel and serial processing is the simultaneous-sequential paradigm. In it, the stimuli are presented simultaneously or in sequence. Parallel processing predicts no effect of the kind of presentation, whereas serial processing predicts an improvement with sequential presentation.

Sagi and Julesz used a variation on the simultaneous-sequential paradigm. The amount of time between target onset was varied but the total SOA was constant. If the delay between target presentation is increased, the parallel model to predicts decreased performance because of the reduced processing time. The serial model predicts no effect because the stimuli can be processed in turn without loss. Increasing the delay between targets did not show a decrease in performance up to a delay of more than 40 ms. This result is consistent with serial processing in this discrimination task.

We entertained the possibility of a ceiling effect that could be masking any improvement. Sagi and Julesz's percent correct reached 90%, while in Bahcall and Kowler percent correct was around 40%.

Next, we created a list of the major differences in experimental design.

	Sagi & Julesz	Bahcall & Kowler
Task	Same-Different	Identification
Display	Empty	Filled
Paradigm	Single	Dual
Presence of Foils	No Foils	Foils
Capacity Limits	Less	More
Performance Limits	Ceiling performance	Not ceiling
Targets	T-L	All letters

After completing the list, we discussed which differences had the greatest influence in the results. From the previous week, the dual and single task experiments gave results that were respectively different from this weeks, so it doesn't appear that this parameter is important in the difference in this week's results. Several of us suspect that the empty vs. filled space is playing the biggest role in determining the differential results.

Discussion of Müller & Klinschmidt

We turned to discussing the Müller and Klinschmidt paper. We found it odd that they always cued the upper left square, and sometimes masked 2 or 4 cues.

With this design, the particular cues are confounded with particular locations. Heather Knapp asked whether it was relevant to have a cue in this experiment if the subject knows where to direct attention on every trial. The effect of the cue itself may become irrelevant because blocking has already provided the "cue".

John Palmer brought up the fact that while they conducted control conditions, they did not discuss or show the data from these conditions. He did not like that omission. On the substance of the paper, Scott Murray has experience with fMRI, and stated that the paper seemed solid. He said that a drop in the BOLD response in fMRI can be due to blood flow steal which is an increase in blood flow in one area causing the loss of blood flow in a neighboring area. However, in this paper the response is retinotopically precise to V1. Another possibility to consider arises because the reported measure is a contrast between two conditions, so it may be possible that in areas where they suggest reduced activity, there is simply increased baseline activity in the control condition. In this paper, the baseline is the time before the cue. However, the center-surround profile was not present in other areas like V2. It is common that attentional effects enlarge as one progresses along visual processing. Several of us suggested an alternative baseline measure that used the condition where they cue all 4 squares.

We discussed the result of not seeing the center-surround profile in regions besides V1. The authors suggest that the experiment can be scaled up to seek the center-surround profile over larger regions for of the higher visual processing centers. The receptive fields between the different regions generally increasing in size as one progresses through visual processing. V1 has a much smaller receptive field then V4. Thus, one might expect the center-surround structure to also scale. In our discussion, we did not buy this argument entirely. It seemed that the relevant scale is defined by the stimulus and the problem of selecting among those stimulus is potentially just as relevant to V4 as V1.

Seminar on the Spatial Resolution of Attention
Questions for Week 5: The attentional walk task
John Palmer and Cathleen Moore
26 April 2006

Goals

Examine measurements of the spatial extent of visual attention using the attentional walk paradigm introduced in Experiment 2 of Intriligator & Cavanagh (2001). In addition, Intriligator & Cavanagh (2001) review related work in crowding, flanker and counting tasks.

Readings

Intriligator & Cavanagh (2001): intro, disc, and Exp. 2; (SKIP Exp. 1)

NEW: Moore, et al. manuscript (2006, now in readings folder at web site)

Study Questions

1. In I&C's Figure 1 demonstration, what different kinds of processing are implied by the resolution and individualization tasks?

2. How do I&C describe individualization?

3. What do you think of the definition of attentional resolution on p. 176 of I&C?

4. I&C present representative results of the attentional walk task in Figures 10 and 11. How good is the analogy between Fig 10 and typical psychometric functions? How about the analogy between typical thresholds and the thresholds estimated in Fig 11?

5. I&C estimate the doubling eccentricity of their attentional walk task as about 0.3°. This implies a larger eccentricity effect than typical resolution task (doubling eccentricity~2°). In contrast, this doubling eccentricity is similar to that reported for crowding (0.2–0.4°, Toet & Levi, 1992, added to web site) and localization (~0.6°, Levi & Tripathy, 1996, in the readings). What is their model of eccentricity effects? How is the "doubling eccentricity" defined?

6. How do the attentional walk experiments of Moore et al. generalize those in I&C? How do the Moore et al. results challenge a purely spatial attention hypotheses?

Discussion Questions

7. Regarding I&C's review of cueing studies, do you agree with their criticism that this paradigm does not constrain subjects to use their finest attentional resolution? Can this issue be solved by the use of foils? Can it be addressed other ways?

8. Regarding I&C's review of crowding studies, they are enthusiastic about the potential connection to attentional resolution. Why should crowding studies reveal something about attention? How close is the comparison of the size of attentional resolution using the attentional walk task and the crowding measures?

9. I&C have brief reviews of the flanker and counting paradigms. How can these paradigms be developed to provide converging evidence for a particular attentional resolution?

10. I&C make much of comparisons between attentional resolution and spatial resolution tasks. Is this a fair comparison? What is the theoretical basis of each of these effects? Is spatial resolution the right comparison tasks for these measures of the spatial extent of attentional selection?

11. What mechanisms account for the errors made in the attentional walk task? On page 208, I&C describe two possibilities.

12. Moore et al. replicate the attentional walk results of I&C. In addition, they add a 0-step condition which is similar to a cueing paradigm. How can this 0-step condition be exploited to compare among the various paradigms of interest?

13. Moore et al. find attention resolution estimates for homogenous arrays are less than for heterogeneous arrays. What are the implications of this effect for purely spatial theories of attention?

Preview for the next two weeks

14. What is the relation between the spatial extent of attentional selection and spatial localization?

15. What is the relation between the spatial extent of attentional selection and the spatial "channel width" of the spatial-frequency based channel theories of spatial vision?

Seminar on the Spatial Resolution of Attention
Summary for Week 5: The attentional walk task
Iris Zemach, John Palmer and Cathleen Moore
17 May 2006

Readings

Intriligator & Cavanagh (2001; intro, disc, and Exp. 2; SKIP Exp. 1)

Moore, et al. manuscript (2006; now in readings folder at web site)

Discussion of Intriligator & Cavanagh

We started the seminar with Intriligator and Cavanagh's paper. We talked about the difficulty of distinguishing spatial interactions from attentional limitations. This point is largely ignored by Intriligator & Cavanagh. Roozbeh Kiani thinks that such "contamination" is likely to play a large role in their measurements. Some tools are necessary for distinguishing these contributions. For example, can the density of the relevant stimuli be varied without changing the displays? One could have alternating red, green, blue, and yellow stimuli and either walk over every stimulus or just the red stimuli. Manipulating the relevant density within identical displays provides more control over spatial interactions. If individualizing the red stimuli is key, then one should be much better with the lower relevant density. Alternatively, if the spatial interactions among nearby stimuli is the key (crowding?), then one might have the same performance regardless of the relevant density.

We also talked about visual resolution as opposed to attentional resolution and how they change as a function of eccentricity. In the periphery, visual resolution appears to be much higher than attentional resolution. On the other hand, localization thresholds and crowding measures increase in the periphery in a fashion more like the measures of attentional resolution. The systematicity of these eccentricity measurements is a positive side the Intriligator and Cavanagh paper.

Discussion of Moore, et al.

We discussed the 0-step condition added in Moore et al. John Palmer argued this was a nice simplification of the walk task. However, it needs to be explored further. For starters, higher densities need to be measured to establish a "threshold" density for the 0-step task. Then one could determine the eccentricity effects for this tasks and compare to the other tasks of interest. One can also ask if the heterogeneous-vs-homogenous manipulation affects the 0-step condition when measured with more of errors (off the ceiling?).

General Discussion

In response to Roozbeh Kiani's concerns that the walk task was "complicated", John Palmer suggested the following variation on the task. Present an initial display similar to the Moore et al. experiment. This would have a grid of placeholders with one cued stimulus. Then remove the cue and grid. After a short interval of time, present a single probe stimulus to the right or left of the cued location. The subjects task is to judge the location of the probe relative to the cued element of the grid. The heart of this task is that if the grid was removed from the initial display, the task reduces to a traditional spatial localization task. In this context, the question becomes whether adding the grid to the initial display reduces performance? Given the results of Moore et al., one would expect such a result. If so, then one could measure the spatial localization threshold as a function of the grid density. The advantage of this method is that it uses a traditional localization task rather than the cued identification task of these papers. The disadvantage is that it introduces a temporal interval and thus an explicit need for memory.

Someone raised the question of how closely does saccade localization correspond to the visual localization tasks such as used in the walk and 0-step task. There is one relevant reference in Intriligator & Cavanagh. Elisabeth Fine also had a abstract on this topic at VSS this year. More generally, this question suggests a connection to the saccade literature on how nearby distractors affect saccade accuracy. One can ask whether such effects due to spatial interactions or attention. I don't know if this question has been pursued.

John Palmer raised the more general issue of how can one define attentional resolution. Clearly each of the papers we have read provide an operational definition of the term. But what about a more theoretical definition that can be applied across tasks? To pick on Moore et al., they give a definition that follows closely one used in Intriligator & Cavanagh:

> "[Attentional resolution is] the smallest inter-item distance that allows one to successfully shift attention among individual items."

This statement can be probably be augmented with a bit of theory to specify the meaning of "successfully", "shift attention" and "individual items" in terms of constructs within a particular theory. This augmented definition is probably fine for the walk and the 0-step task. But, how can it be extended to the connection between a wider range of paradigms? Such paradigms include traditional spatial localization, crowding tasks, etc. What one needs is an overarching theory of how to model a variety of tasks. One approach is to combine the channel theory of spatial vision, the signal detection theory treatment of task differences, and to specify how selective attention relates spatial channels to task-specific decision rules. A tall order.

Seminar on the Spatial Resolution of Attention
Questions for Week 6: Spatial localization
John Palmer and Cathleen Moore
3 May 2006

Goals

Examine measures of spatial localization and its dependence on stimulus characteristics. Consider the relation between spatial localization and attentional resolution.

Readings

Levi & Tripathy (1996)

Graham (1992), background only

Study Questions

1. What is the difference between relative spatial localization and "absolute" spatial localization?

2. What is their model of eccentricity effects and what is the doubling eccentricity parameter? (Hint: Equation 4)

3. What is their model of the effect of blur and what is the intrinsic blur parameter? Also, what is the asymptotic threshold? (Hint: Equation 3)

4. What is a local sign model of localization?

Discussion Questions

5. Are measures of spatial localization relevant to attentional resolution?

6. What are the alternative interpretations of the intrinsic blur estimate? how is it related to the spatial extent of the analyzers mediating these judgments?

7. Resolution thresholds have a doubling eccentricity of about 2° while localization as measured in this paper has a doubling eccentricity of about 0.5°. Thus, if they were equal at 1° eccentric, the localization threshold would be a fraction of the resolution threshold in the fovea and nearly twice the size of the resolution threshold at 8° eccentric. What does this imply about visual mechanisms? What might it imply about attentional resolution?

8. In Figure 12, the authors compare the spatial scale of "positional pooling" (intrinsic blur) and various anatomical standards. They suggest that the pooling measure corresponds in size to either V2 receptive fields or V1 receptive fields augmented with the silent surround reached by lateral interactions. Might these regions be the anatomical locus of the relevant analyzers? Might they also determine attentional resolution?

9. Consider the following argument. Localization of single stimuli can be made more precise than the spatial extent of the relevant analyzers because outputs of an array of analyzers centered on different locations can be compared to estimate the location of a single stimulus (an analogous model for spatial frequency discrimination and perception is Davis, et al., 1995, on web site). Such simple comparisons across such analyzers may not work for experiments with multiple stimuli or filled visual fields. Might such multiple stimulus displays result in the spatial extent of the analyzers becoming the relevant resolution of the system?

10. Consider another approach. Suppose one measures spatial localization for stimuli at detection threshold. Such stimuli may affect only the analyzers that are most sensitive to stimuli at a particular position. Thus, localization of such threshold stimuli may tell one more about the spatial extent of the relevant analyzers. An somewhat analogous application of this approach can be found in the "label lines" paper of Watson and Robson (1981, Vision Research). The complication is that the analyzers for position may have more overlap than for spatial frequency.

Seminar on the Spatial Resolution of Attention
Summary for Week 6: Spatial localization
Roozbeh Kiani, John Palmer and Cathleen Moore
24 May 2006

Readings

Levi & Tripathy, 1996

Discussion

There are two key measurements in the paper:

a) the effect of eccentricity on the localization threshold, and

b) the effect of stimulus blur on the localization threshold.

The paper demonstrates that subjects' accuracy for localization falls very quickly as a function of eccentricity. The fall-off rate (E_2 in equation 4 of the paper) or doubling eccentricity is 0.6 degrees. This is the eccentricity at which the foveal asymptotic threshold doubles. A value of 0.6 degrees is quite low and indicates very steep fall-off in localization accuracy increases in eccentricity. It is much lower that the fall-off rate for two-point discrimination or grating resolution. For these tasks, the doubling eccentricity is typically about 2 degrees. In fact, a doubling eccentricity of 0.6 is very close to the doubling eccentricity that Intriligator and Cavanagh (2001) have reported for their attentional walk task. The steep fall-off has been interpreted by Cavanagh as evidence that attentional resolution declines faster with eccentricity than would be expected by limits on spatial resolution. However, the results of Levi and Tripathy provide an alternative and interesting explanation.

In order to attend to something we must first localize it. Any limitation of the localization process will affect our measure of the spatial resolution of attention. Levi and Tripathy provide evidence for the limited accuracy of localization. The fact that the fall-off rate measured by them is very close to the reported value for the attentional walk task suggest that the inability of subjects in precise localization of targets is the main limitation in the attentional walk task and probably in other crowding experiments. In other words, the resolution limit measured by Intriligator and Cavanagh is due to limits in localization and *not* additional limitations of the spatial resolution of attention.

Levi and Tripathy report a similar eccentricity effect for the intrinsic blur. They assume that for the localization of targets there is an intrinsic source of noise which is uncorrelated with the noise associated with the stimulus (blur). They take the standard deviation of the Gaussian envelope of the Gabor stimuli as the external noise (external blur) and then ask how the intrinsic noise (intrinsic blur) should change to cause localization threshold changes similar to those observed in human subjects' data (equation 3 of the paper). The find a fall-off rate of 0.6 degrees for the intrinsic blur. Similar eccentricity effect on localization threshold and intrinsic blur suggests a channel theory for localization.

John Palmer found the simple additive model used for describing the blur effect a nice analogy from noise models. But he felt that it needed to be developed in more detail in the context of spatial blur. Specifically, how would additivity of intrinsic and external blur be implemented in a typical spatial vision model?

In these models, the spatial channels can indicate a region in space where the visual signal is pooled for localization. The size of these channels is one of the factors that determine the precision of localization. An important question arises about the neural correlate of such channels. In other words, how realistic is the assumption of a channel for localization? One would expect to be able to find neurons that pool information over the same spatial extents that these channels suggest. Levi and Tripathy suggest V2 for this purpose. However, there is a lack of any specific evidence to connect these channels to neurons in a particular cortical area.

One potential source of concern about the results in this paper is that the perimeter of the monitor was not blacked out in the task. It may allow subjects to improve their performance by using the edge of the monitor as a landmark for localization of the target. It would be nice if a more advanced setup which did not use monitors could be used for the experiment. However, it does not seem to be a very important concern because the reported thresholds in the paper remained the same when subjects' distance from the monitor changed and hence changed the proximity of the edge of the monitor.

Another concern about the measurements in the paper is the possibility that the intrinsic blur and external blur not be totally uncorrelated. However, in the absence of a direct measure for the intrinsic blur, equation 3 in the paper is probably the best we can do.

Yet another concern, if the channels suggested by the paper exist, it should be possible to measure their width by other methods. Masking experiments which will be considered next week may provide such an opportunity.

Finally, Roozbeh asked whether localization itself required attention. John Palmer argued from his visual search experiments that relative localization tasks show signs of sharply limited capacity and may even require one-at-a-time serial processing.

Seminar on the Spatial Resolution of Attention
Questions for Week 7: Masking and Foil experiments
John Palmer and Cathleen Moore
22 May 2006

Goals

Examine two additional paradigms that may be related to attentional resolution: cueing effects on masking and selective filtering of foils. Both paradigms are intended to measure spatial tuning functions that may be relevant to attentional resolution.

Readings

Baldassi and Verghese (2005); Focus on location cueing of the location tuning function (Figure 5, left and center columns).

NEW: Palmer and Moore (2006, manuscript distributed by email, not on web)

DROPPED: Buffalo, et al. (2005)

Study Questions

1. What effect of spatial uncertainty is measured by B&V without masking? How sensitive are these measurements?

2. How do B&V measure a location tuning function using masking (Figure 5)? How is it affected by location cues?

3. What is B&V's reweighting model (p. 566)?

---

4. How do P&M measure the effect of selective attention on the perception of foils?

5. What is the magnitude of the spatial selectivity effects measured in P&M?

6. What are the differences in predictions between the multiplicative and all-or-none models described by P&M?

7. What is the estimated critical separation in P&M?

Discussion Questions

8. What is the expected effect of spatial uncertainty under the conditions of B&V's experiment? (Hint: See the derivation for 2IFC in the appendix of Palmer, Ames and Lindsey (1993, added to the web site "additional readings"). A complication specific to detection thresholds is that their psychometric functions are "steeper" than for most other judgments and this reduces the magnitude of the effect.)

9. What theories predict no effect on the width of the location tuning function with cueing?

10. Is there a connection between the location tuning functions measured by B&V and attentional resolution?

---

11. What are the advantages and disadvantages of measuring selective attention using foils instead of targets?

12. What theories of selective attention can be classified as either multiplicative or all-or-none in the sense defined by P&M? What theories do not fit into either of these categories?

13. How can one compare the critical separation observed by P&M to analogous measures in other studies?

14. Is the critical separation measured by P&M related to attentional resolution?

Seminar on the Spatial Resolution of Attention
Summary for Week 7: Mask and Foil Experiments
Heather Knapp, John Palmer and Cathleen Moore
29 May 2006

Readings

Baldassi and Verghese (2005, focus on the location tuning function, Fig. 5).

Palmer and Moore (2006, manuscript distributed by email, not on web)

The Masking Experiments of Baldassi and Verghese

We began with a discussion of Baldassi and Verghese (2005). In this paper, the authors utilize a masking paradigm to explore how location and orientation cues selectively affect the shape of location and orientation tuning functions. They conducted two sets of experiments with three conditions each: unmasked (uncued, location cue, orientation cue), and masked (neutral cue, location cue, orientation cue). In all experiments, subjects were instructed to indicate which of two time intervals contained a test patch. Contrast was also manipulated. Target interval and position was counterbalanced.

The unmasked experiments were conducted to obtain baseline detection thresholds in cued and uncued conditions. Contrast thresholds were found to be relatively similar across uncued, location cued, and orientation cued conditions (Figure 1). The authors conclude that there is no effect of cueing, and remark that they are surprised by this, as they expected a 20% improvement on threshold. John Palmer thinks that 20% is actually too high a prediction for these conditions. This is because contrast detection has an especially steep psychometric function. One would expect a correspondingly smaller shift in threshold, say 10% instead of 20%. Given that the size of the error bars on the observed data was nearly 20%, it appears that they did not obtain enough trials to detect the expected small effect.

Masked experiments were conducted with neutral cues, location cues, and orientation cues. The change in the shape of the tuning function between neutral cue and location cue conditions provide information as to how location cues affect the basic sensitivity of the filter. This change in sensitivity as a function of attention being directed to a location near the target was our primary interest (Figure 5). Neutral cue trials indicate that contrast threshold changes dramatically with the distance of the masks (two masks, one horizontal and one vertical, moved in tandem) to the test patch. Thresholds were highest when the mask and test occurred at the same location, and were lowest when the mask was the greatest distance from the test. Adding location cues does not change the shape of the tuning function, but lowered thresholds by a small amount across the entire range. Thus, effects of location cueing were found to be wide and non-specific. John Palmer pointed out that these results are consistent with physiology papers that find similar effects for orientation cues in V4. Baldassi & Verghese (2005) is one of the few papers to attack this behaviorally.

Scott Murray pointed out that he was troubled by conditions in which the mask and target overlap spatially. With overlap, it is possible that the entire stimulus becomes a different one. The mask and target are no longer two separate objects to be distinguished, but meld to form a single novel object. Thus, the judgment may be fundamentally different.

John then discussed these results in terms of the authors' reweighing model, in which cueing changing the weights of detectors (i.e., neurons) in specific ways. Location cues seem to reweigh the gain of all detectors equally (by a factor of 1.4), rather than increasing the gains of some detectors and not affecting others (as with orientation). John points out that the reweighing concept is consistent with biased competition models of attention.

The Foil Experiments of Palmer and Moore

The seminar then moved to a discussion of Palmer and Moore (in preparation). In this paper, the authors' goal is to explore the spatial resolution of attention by finding the critical separation for attentional resolution. This is operationalized in a spatial filtering task as the minimal separation between a relevant and irrelevant stimulus that allows a subject to respond to the relevant stimulus perfectly and, conversely, to ignore the irrelevant stimulus completely.

The authors use a two-interval-forced-choice (2IFC) paradigm, in which the subject must indicate whether the target stimulus appeared in the first or second temporal interval. Each trial contains a target and an otherwise identical *foil* at an irrelevant location. Both target and foil occur on each trial, but in the same interval (congruent foil) or in different intervals (incongruent foil). The target stimulus is a light at a pre-cued location in a circular array. The foil is a light at one of 2 possible locations on either side of the target. Contrast of the target and foil were varied such that a given target contrast value (10%) occurs with a range of foil contrast values (10, 14, 20, 28, and 100%), and a given foil contrast value (10%) occurs with a range of target contrast values (10, 14, 20, 28, and 100%). Psychometric functions were measured for both relevant and irrelevant stimuli. Finally, this entire design was repeated for different separations between the target and foil.

Two possible mechanisms of selective attention are contrasted: a multiplicative model (e.g. contrast gain), and an all-or-none model (e.g. attention switching). A multiplicative model describes a selection process that attenuates less relevant representations by a multiplicative factor. Its specific predictions are for the psychometric function for irrelevant stimuli, in which a horizontal shift occurs with each separation increment, corresponding to a decrease in gain. The obvious parameter to be measured in multiplicative models is detection threshold. An all-or-none model, on the other hand, describes a selection process in which stimuli are judged to be relevant or not relevant. Irrelevant stimuli do not affect performance. However, the probability that a foil is judged to be irrelevant increases with its distance from the target. Psychometric functions for irrelevant stimuli are predicted to be scaled vertically. The parameter to be judged in all-or-none models is therefore the asymptote of the function.

John Palmer added a new point that is not in the paper. These mathematical models can both be formulated as either early or late selection theories. For the multiplicative model, modulating contrast gain is a possible early selection model while a multiplicative weighting of evidence at decision is a possible late selection model. For the all-or-none model, a spatial attention switching process is a possible early selection model while an all-or-none decision (is this a target or foil) based fully formed percepts is a possible late selection model.

Heather Knapp pointed out that the terminology "relevant stimulus" and "irrelevant stimulus" gave the impression that there was something different about the stimuli. Perhaps it would be better to refer to "relevant location" and "irrelevant location". If one needs to refer the stimuli specifically, then one could use the labels "target" and "foil". This made me (John Palmer) think about how to fix another sticky point of terminology. How about the following notation for the psychometric function labels: P(Response to Target) and P(Response to Foil).

Scott Murray pointed out that Table 1 in the paper was very hard to follow. Nicolle Perisho suggested dropping the "x" notation from the table and instead use explicit contrast values.

The data from the two subjects in this experiment show that the effect of separation on how the foil affects performance is quite pronounced: By a four-step separation (one-step = 0.6°), the effect on performance is at chance, and an eight-step separation yields below-chance performance. The observed psychometric functions for the response to the foil are vertically scaled, consistent with the all-or-none model. A spatial tuning function constructed from these data suggest a critical separation of about 2 degrees. Of special note is the below-chance performance at 8-steps of separation. Below-chance performance is quite surprising to the authors, who discuss several possible explanations. One plausible mechanism is attentional masking, in which masked information gives less information about the target than no information at all. A way of conceptualizing this is that a 'something' at one location (interval 1) is worse than a ‘nothing’ at that same location (interval 2). Another possible explanation is attentional capture, in which an interval with a foil captures attention from the target, and has a deleterious effect on performance.

John points out that another way of measuring these effects of the foil is via a spatial two-alternative-forced-choice paradigm, rather than the temporal two-interval-forced-choice paradigm. The below-chance effect should go away if it is independent of separation (such as the attentional capture hypothesis). On the other hand, it should remain if it is some kind of spatial inhibition that is dependent on separation.

Roozbeh asked if one could also analyze the spatial effect of the foil on the detection of the target. This is in contrast to the previous analysis which was all about the response to the foil and not to the target. This other effect is presumably the masking effect of the foil on the target. This analysis is indeed possible, but there are some complications. In particular, it is hard to interpret the zero separation case because one needs to have an overlapping target and foil. The control condition is not appropriate for this purpose.

An few final questions were raised. What distinguishes a mask from a foil? For a mask, the mask is typically presented in both intervals. A foil, on the other hand, is only presented in one interval. Cathleen Moore in email discussions favors measuring the effect of masking in a traditional masking paradigm and making an explicit comparison to the effect of the foil in this new spatial filtering paradigm.

Another general question for the future. What are the advantages and disadvantages of using foils?

A final high-priority question inspired by this seminar is to extend the two models to predict "below-chance" performance. With a little care, this appears possible but the interpretations of the possible extensions are not obvious.

Seminar on the Spatial Resolution of Attention
Questions for Week 8: Finale
John Palmer and Cathleen Moore
24 May 2006

Goals

Review issues raised over the course of the seminar. Attempt closure on a few selected issues. Cathleen Moore will be our special guest.

Readings and Assignments

No additional readings. Please pick one discussion question and write an informal answer in 1-3 paragraphs. Email me your answer by the end of the day Tuesday 30 May. Be prepared to discuss them during our final meeting.

Discussion Questions

1. What are the critical attentional concepts for analyzing attentional resolution? For example, consider the ideas of contrast gain and a spatial tuning function. How can these concepts be defined for both behavioral and neural experiments? Other critical concepts?

2. What is the best empirical measure of attentional resolution? Be specific! How about the critical separation as estimated from a spatial filtering task (Palmer & Moore, 2006, manuscript)? Alternatives?

3. What is the best theoretical definition of attentional resolution? Be general! How about the width of the spatial filter mediating a spatial selective attention task? By spatial filter, we mean the weighted combination of detectors as defined in Baldassi and Verghese (2005); by selective attention task, we mean tasks like those described in Palmer and Moore (2006). Alternatives?

4. Over what range of stimulus conditions should one be able to generalize estimates attentional resolution? For example, within reasonable limits, should blur or stimulus size affect attentional resolution? Other variables?

5. Over what range of tasks should one be able to generalize estimates of attentional resolution? For example, should one be able to compare measures of attentional resolution based on localization and identification tasks? Others?

6. What manipulations are expected to affect attentional resolution? For example, Intriligator and Cavanagh (2001) make the case that attentional resolution varies systematically with eccentricity. Moreover, they argue that this variation can be used to relate attentional resolution to other sensory phenomena. What other manipulations go to the heart of attentional resolution? Empty vs. filled space? Limited vs. unlimited capacity tasks?

Seminar on the Spatial Resolution of Attention
Summary for Week 8: Finale
Michael Lee, John Palmer and Cathleen Moore
9 June 2006

1. What are the critical attentional concepts for analyzing attentional resolution? For example, consider the ideas of contrast gain and a spatial tuning function. How can these concepts be defined for both behavioral and neural experiments? Other critical concepts?

We began with a brief discussion about the definition of attention and related concepts. John Palmer brought up the unstated assumption in some of the papers that attention acted upon a representation of the stimulus by an array of spatial-frequency and spatially tuned channels. Attention can act in several ways. One way is to modulate the gain on a channel and another way is to change the profile of the spatial filter for a channel. This gives us a basic language to discuss possible attention models.

Roozbeh Kiani asked how attentional effects can be explained by gain control. The simplest gain change would increase the amount of noise as well as signal, thus not improving the signal-to-noise ratio. Two possibilities were raised in response. One is the possibility of a hierarchical channel theory with multiple layers that converge as one rises in the hierarchy. In such a theory, differential gains on the outputs of a lower layer change the spatial tuning of the upper layer. Thus, gain has it effect by weighting one source of information more than another. If relevant sources are weighted highly and irrelevant ones weighted low, one can improve the overall signal-to-noise ration. A second possibility is that the effect of gain on the signal may not be modeled by a simple amplification. In particular, the Fano factor (ratio of variance to mean) is often observed to remain constant in neural coding. If so, an increase in gain will increase the noise but not as fast as the signal. Thus, there is an improvement in signal-to-noise ratio if gain acts in this fashion. John Palmer mentioned that a recent poster at the 2006 VSS by Reynolds suggested deviations from a constant Fano factor in an attention experiment.

We also talked about a "late selection" style model. Suppose the spatial tuning functions are fixed and there are no early gain modulations. Attention can still have an effect in how these sources of information are combined into a decision. By this hypothesis, the weighting is only at the very end of the processing.

2. What is the best empirical measure of attentional resolution? Be specific! How about the critical separation as estimated from a

spatial filtering task (Palmer & Moore, 2006, manuscript)? Alternatives?

We discussed the best ways to measure attentional resolution. The focus was on three tasks: cueing, walk and foil tasks. Our primary example of a cueing task was Sagi and Julesz who used a dual task and showed relatively small effects of cueing. The alternative walk task requires endogenous cueing to signal a move of attention from one point to another. Performance in this task show large declines if the stimuli are too dense. The third is a cueing task with foils which must be ignored to correctly perform the task. These foil experiments appear to generate large effect when measuring the effect of the foil rather than the effect of the target.

These three task need not measure exactly the same thing. For example, weighting of different relevant sources in the original cueing tasks may have different properties than underlie ignoring a foil in the foil task.

Cathleen Moore was interested in considering these tasks as representing a range of increasing complex sets of cognitive operations. For example, perhaps the cueing paradigm taps a single weighting operation; the foil task taps a somewhat different filtering operation; and the walk task taps several operations involving the initial cue, changes in attention, and retrieval of the attended location.

3. What is the best theoretical definition of attentional resolution? Be general! How about the width of the spatial filter mediating a spatial selective attention task? By spatial filter, we mean the weighted combination of detectors as defined in Baldassi and Verghese (2005); by selective attention task, we mean tasks like those described in Palmer and Moore (2006). Alternatives?

The approach assumed in the question relies heavily on the "standard" spatial–frequency channel theory of early vision. The advantage of this approach is that it is well defined and fairly successful for contrast detection experiments. The disadvantage is that it has not successfully generalized to superthreshold stimuli and tasks. Thus, while this theory is a good starting point, it may not help with understanding attention with conventional perceptual judgments.

Roozbeh Kiani stated that the suggested “working” definition of attentional resolution was too similar to that of localization. John argued that one needed to differentiate the tasks — say cueing and localization — from the underlying mechanisms. Thus, the spatial resolution of both cueing and localization might depend on the underlying channel's spatial tuning function.

We also discussed the role of attention in a localization task.

4. Over what range of stimulus conditions should one be able to generalize estimates attentional resolution? For example, within reasonable limits, should blur or stimulus size affect attentional resolution? Other variables?

This question addresses what stimulus conditions yield common estimates of attentional resolution. Papers we read suggested that these should include stimulus size and blur at least over reasonable ranges.

As with the last question, one approach to this question is to assume the spatial–frequency channel theory of early vision. John Palmer and Cathleen Moore argued that if the finest spatial channel is responsible for attentional resolution, than blur will not affect attentional resolution until it prevents the finest spatial channel from being modulated by the stimuli. Size would similarly not have effects. Indeed, for objects with sharp edges, size should never have an effect because sharp edges always modulate the fine spatial channels regardless of size.

We briefly discussed crowding. Need further study to relate crowding and attentional resolution. They may well share a common basis in the underlying spatial tuning functions. Understanding the role of attention (or its absence) in crowding complicates thinking through these connections.

Roozbeh Kiani asked what attentional phenomena are most used for daily life. John Palmer said that activities such as reading and switching between targets requires the serial application of selective attention. Another natural activity is locomotion such as walking and avoiding obstacles. Under these conditions, one often uses peripheral vision to control gait. Cathleen Moore suggested visual search as a relatively natural task used to scan the environment to find relevant information and ignore irrelevant information. She also talked about the link between attention and eye movements in the context of search.

Roozbeh Kiani asked if there is unlimited capacity for the processes mediating a visual search task, why is there a need to move the eyes at all? John Palmer also wondered about this point. He described an experiment where, on the face of it, eye movements were not necessary because of the minimal effects of peripheral vision, but subjects made them anyway. They may serve other purposes than reducing the effect of peripheral vision. Or they may simply be an overused strategy.

Cathleen Moore talked about the counting tasks investigated by Kowler. Kowler found that counting performance declined for dense stimuli. Curiously, this effect was not reduced when eye movements were allowed.

5. Over what range of tasks should one be able to generalize estimates of attentional resolution? For example, should one be able to compare measures of attentional resolution based on localization and identification tasks? Others?

Michael Lee said that in an identification task, the requirement to access memory may introduce differences from the attentional effects seen in localization tasks.

John Palmer briefly discussed experiments designed to compare identification and localization performance for simple tasks. These are summarized in the chapters in Norma Graham's book on detection-discrimination tasks. The best known is sometimes called the 2-by-2 task and has been studied in detail for situations such as the detection of gratings versus the discrimination of their spatial frequency.

6. What manipulations are expected to affect attentional resolution? For example, Intriligator and Cavanagh (2001) make the case that attentional resolution varies systematically with eccentricity. Moreover, they argue that this variation can be used to relate attentional resolution to other sensory phenomena. What other manipulations go to the heart of attentional resolution? Empty vs. filled space? Limited vs. unlimited capacity tasks?

Roozbeh Kiani commented on eccentricity effects. If localization and identification tasks show the same eccentricity effects, one could subtract the eccentricity effects and cause the attentional effects to be flat. To counter that view, Cathleen Moore suggested that one could also subtract the localization effects and see if the effect was due to attention. In short, a more theoretical analysis is necessary to compare the processes that underlie these tasks. No simple contrast is appropriate without a motivating theory.

John Palmer pursued the ideas of spatial-frequency channel theory for this question. For this theory, the natural stimuli are Gabor patches because they best isolate individual channels. Might one compare attentional resolution for Gabor patches of different size and frequency? Do we predict the same attentional resolution or not? Localization does decline with very low spatial frequency stimuli, so may attentional resolution.