The Conjunction Fallacy in Perceptual Decision Making

Jared M. Hotaling, Jerome R. Busemeyer, and Richard M. Shiffrin
{jhotalin, jbusemey, shiffrin}@indiana.edu)
Department of Psychological & Brain Sciences, Indiana University
1101 E. Tenth St., Bloomington, IN 47408 USA

Abstract

In cognitive science there is a seeming paradox: on the one hand researchers studying judgment and decision making have repeatedly shown that people employ simple and often less than optimal strategies when integrating information from multiple sources. On the other hand, researchers have had great success using optimal models to account for information integration in fields such as categorization, memory, and perception. This apparent conflict could be due in part to different materials and designs that lead to differences in the nature of processing. Stimuli that require controlled integration may lead to sub-optimal performance. In contrast, the images often used in perceptual domains may lend themselves to automatic processing, resulting in integration that is closer to optimal. The present research explores these ideas within a new perceptual decision making paradigm. We investigate one canonical example of sub-optimal information integration, the conjunction fallacy. Surprisingly this effect was quite reliable across all stimulus presentations. However, we found evidence that manipulations meant to encourage controlled integration of information produced larger conjunction errors, as well as evidence that slower, more deliberate processing yielded performance closer to optimal.

Keywords: conjunction fallacy, information integration, decision making, perception.

Introduction

Information integration is the combination of different sources of information for the purpose of performing some task. To understand speech one must integrate its visual and auditory components; a doctor needs to integrate information from a patient's symptoms, and contextual factors such as family history, diet, and exercise habits. In these examples each piece of information, on its own, provides some predictive or diagnostic value. It is obviously valuable to combine different sources of information in an optimal fashion. The extant literature provides examples in which near-optimal information integration occurs, but just as many examples in which integration is sub-optimal. The tasks in which the different results are found often differ qualitatively, so the present research explores information integration success and failure within a common perceptual decision making paradigm. Within this paradigm we ask which conditions give rise to rational integration, and which conditions hinder it.

Sub-optimal Information Integration in Judgment and Decision Making

On the one hand, there is abundant evidence from many studies of judgment and decision making that information from different sources is often integrated via heuristic strategies. Sometimes simple heuristics produce results that approach optimal decision making (Gigerenzer & Todd, 1999), but in many other cases the results fall well short of optimal combination, relative to standard rational theories of inference (Gilovich, Griffin, & Kahneman, 2002). Findings of deviations from rational behavior have had a great influence on theorizing and within a large subset of judgment researchers it is now common practice to assume sub-optimal performance as a starting point for theories.

For example a dilution effect occurs when weak positive evidence is combined with strong positive evidence to decrease the probability of a hypothesis (Shanteau, 1975). Another example is the "conjunction fallacy" in which the probability of a conjunction is judged greater than one of the individual events in the conjunction. The examples above are only a small sample from the research showing violations of rational principles of information integration in traditional judgment and decision making tasks.

Rational Models of Perception, Categorization, and Memory

On the other hand, there are numerous highly successful applications of optimal or rational models of information integration in fields such as categorization, memory, and perception. In contrast to the judgment literature, the success of such models has moved researchers in these more perceptual fields to begin with an assumption of optimality and only later investigate sub-optimal performance. Examples of optimal information integration are easy to find and the following are a very few of many potential examples. Ashby and colleagues (Ashby & Gott, 1988; Ashby & Maddox, 1990) have found categorization performance to be well described by a decision bound that either nearly optimally integrates information across two dimensions or uses a rule of the same form as the optimal bound. Using an exemplar model of categorization, Nosofsky (1986) suggested that observers tend to distribute attention among dimensions so as to optimize categorization performance. Tenenbaum (1999, 2000) has successfully utilized a Bayesian framework for modeling human concept learning. Anderson (1991) has demonstrated that many common results in both the categorization and memory literatures can be well described by his Bayesian Rational model. Shiffrin & Steyvers (1997) introduced a model built on a Bayesian foundation that predicts phenomena of explicit, implicit, episodic, and generic memory. The Fuzzy Logic Model of Perception (Massaro, 1998) is based on the idea that observers optimally integrate different sources of information and has been applied to a large range of data.

In sensory science, optimal information integration is usually the default assumption and the starting point for research. Of thousands of potential examples, we mention Burgess, Wagner, Jennings, & Barlow (1981), who found the ability of human observers to discriminate visual patterns in noisy backgrounds to be very close to that of an ideal observer. Kersten, Mamassian, & Yuille (2004) provide a review of recent developments in using Bayesian approaches to model people's ability to perceive objects in complex and noisy environments.

Toward Reconciliation

Although information integration is the object of study by researchers in both judgment and decision making domains and in cognitive and sensory domains, these fields often seem to be operating independently of each other. We speculate that one source of such independence is the wide difference in experimental paradigms. The judgment literature focuses mainly on linguistic or quantitative statements of probabilities and is concerned with the ways in which people use information to assess, estimate, and infer what events will occur (Hastie, 2001). The other lines of research finding more optimal performance discussed above typically rely on perceptual stimuli, such as images and sounds, and concentrate on how the information is produced from external stimulation. The first goal of the present research is to bridge this divide through use of a common experimental paradigm.

Recent work by Hotaling, Cohen, Busemeyer & Shiffrin (2010) initiated an investigation of probability judgment errors using a task more similar to the types that supported optimal performance in the past. They investigated the dilution effect, but instead of traditional stimuli they used this perceptual decision making paradigm. They found that sizeable dilution effects continue to occur in this new paradigm, but they also identified perceptual factors that moderate the size of the effect. The experiment presented here continues this line of work but focuses on another phenomenon in which sub-optimal information integration has been observed: the conjunctive fallacy. This effect (e.g., Sides, Osherson, Bonini, & Viale, 2002; Tversky & Kahneman, 1983) is typically illustrated with the famous "Linda the bank teller" story. The laws of probability imply that the probability of a conjunction of two events cannot be larger than the probability of either event separately: e.g., $P(X \cap Y \mid Z) \leq P(X \mid Z)$. Yet a majority of respondents across a variety of studies claimed that statements such as "Linda is a bank teller and is active in the feminist movement" were more likely than "Linda is a bank teller."

Much debate has surrounded the conjunction fallacy since its first demonstration, most of which has focused on its implications for the rationality of human cognition. Rather than focus on the irrationality of the effect, this paper investigates whether or not it is limited to the kinds of the tasks and stimuli typical of traditional judgment and decision making research. Can the conjunction fallacy be found in a more perceptual domain, and if so, what factors affect its prevalence and strength? Although the conjunction fallacy has only been explored using verbal or quantitative stimuli, it easily lends itself to investigation in a perceptual setting. The present research explores the combination of weak and strong evidence from different parts of an image. Consider the example of an oncologist examining a worrisome spot on a brain MRI. The doctor may be interested in two things, whether or not the spot is a tumor, and whether or not it is operable. He may use the color and patterning of the spot to determine the first question, while the area around the spot informs him of how close the object is to other critical brain areas. When deciding on a course of action, the doctor must evaluate the conjunction of these questions. Is this information combined in an optimal fashion, or will the results show sub-optimal combination, as exemplified by the conjunction fallacy? It would be natural to expect the former based on the many studies showing near-optimal combination of perceptual information. Our study compares conditions that are more aligned with normal perception and should foster closer to optimal information integration with conditions that deviate from normal experience and might lead to further from optimal information integration.

Another benefit of using perceptual stimuli is that issues of interpretation and language understanding do not come into play. For example, the conjunction law is much less likely to be violated if participants interpret "Linda is a bank teller" to mean that she is a bank teller and *not* active in the feminist movement (see Sides, et al., 2002 for a short review). In addition, observers often misunderstand how to interpret probabilities; for example more optimal information integration sometimes occurs when data is presented as frequencies rather than probabilities (Gigerenzer & Hoffrage, 1995). Our present tasks require perceptual matching from memory, thereby greatly reducing any confounding influence of language conventions.

Within our present paradigm there is hope that we can ferret out differences in processing that might account for any differences we observe across conditions. It is fairly common to distinguish automatic and controlled processing, both in theory and in empirical research (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Cohen, 2006; Erickson & Kruschke, 1998; Jacoby, 1991; Runeson, Juslin, & Olsson, 2000; Shiffrin & Schneider, 1977). Most often automatic processing is assumed to be fast and independent of conscious manipulation, and controlled processing is assumed to be slow and conscious. In the perceptual literature, automatic processing is usually assumed to be more robust, less prone to large errors, less based on heuristics, and closer to optimal than controlled processing. The judgment and decision literature also use a similar dichotomy but with different implications for optimal processing (e.g., Devine, 1989; Epstein, Lipson, Holstein, & Huh, 1992; Hammond, 2000; Sloman, 1996; Stanovich & West, 2000). Many of these researchers have introduced their own unique terminology to describe the different processes. In this paper we adopt the automatic/controlled language used by Shiffrin & Schneider (1977). This is done primarily for convenience sake, rather than to make a strong claim that information integration is either entirely automatic or completely controlled.

Hotaling et al. (2010) investigated the effects of automatic and controlled processing on information integration in perceptual decision making. They found that the dilution effect was stronger in conditions that promoted controlled combination of evidence, while stimulus conditions aimed at fostering automatic integration yielded performance closer to optimal. However, it would be premature to conclude that one mode of processing is simply *better* than another. Rather, some modes are probably better suited to certain kinds of information and tasks. There is reason to believe that automatic processing may actually increase the prevalence of conjunction errors in perceptual decision making. Tversky & Kahneman invoked the *representativeness* heuristic to explain the conjunction fallacy. They claimed that participants saw Linda as more representative of feminist bank tellers than of bank tellers in general. Rather than view "Linda is a bank teller and is active in the feminist movement" as a conjunction of two events participants based their responses on a global similarity judgment. In the experiment reported below we use conditions that manipulate face images in ways designed to bias processing toward or away from automatic processing. Most would think it intuitively obvious that the features of a normal face image would be automatically combined through vision. This automaticity may impede an observer's ability to separately evaluate parts of the image in order to judge a conjunction statement. The top and bottom halves of the image may be automatically treated as a whole, just as *feminist* and *bank teller* were combined into a single concept.

We chose to use face images for several reasons: in their normal form they should promote automatic processing (face detection and identification is an over-learned visual tasks). Face images can be altered from their normal form (without significantly changing the information content) in a way that should disrupt automatic processing and thereby promote intentional processing (e.g., turning a face upside down, or separating a face into halves that are not spatially aligned can disrupt automatic face processing, Young, Hellawell, & Hay, 1987). The parts of faces can be adjusted independently to provide different evidence strengths, say by altering contrast, or by manipulating similarity to choice targets. To summarize: For well-learned perceptual stimuli, such as faces in their normal form, automatic processes ought to dominate. However, for faces split in a fashion never experienced, or for faces turned upside-down, information integration probably will rely more on controlled processes. Such processing differences should affect the degree to which information is integrated in accordance with the conjunction law, but it is difficult to predict the direction of these effects a priori.

Method

Participants

Nine students at Indiana University were paid to participate in this study. Seven were female and two were male. All participants had normal or corrected vision.

Stimuli

All of the stimuli used in the experiment were derived from two "target faces" (Targets A and B) selected from the FERET database (Phillips, et al, 2000). These faces were warped so that their major facial features aligned and were then cropped to remove the hair and head outline. Once this is done, a morph is essentially a linear combination of the grayscale values of the two faces at each pixel. The cropped areas of the 256 × 384 pixels images were filled with a sinusoidal grating. These two target faces were used to construct the experimental stimuli, separately for each participant, during the calibration phase of the experiment.

On Day 1 participants completed two blocks of trials during which a staircase algorithm was used to find the half face morphs that produced roughly 75% accuracy. There were 19 levels of morph, spanning the range from Target A to Target B. The participant saw a test face appear on the computer screen for two seconds. This image was then masked and the two target faces appeared on the screen. On each trial participants chose the target face that most closely resembled the test. Auditory feedback was given, indicating a correct or incorrect answer. The first block involved upright faces, while the second used inverted faces. Each block consisted of 72 half face trials. Faces favoring Target A and Target B were initialized to 94.44% A and 5.56% A morphs, respectively. 48 whole face trials were interspersed as fillers to prevent participants from developing novel strategies tailored only to half faces. Filler trials did not affect calibration. Half face morphs were adjusted with a 2-up-1-down algorithm, whereby the morph proportion was stepped closer to 50-50 after every two correct responses, and was moved further from 50-50 after each incorrect response. The result was that for each target and each orientation, top and bottom morphs were found that produced an intermediate level of accuracy (roughly 75%). These morphs were used to create medium (M) half faces, while weak (W) and strong (S) morphs were derived by extrapolating from these values. W halves were defined by the morph coefficient halfway between the M morph and 50-50 morph. S halves used the coefficient two thirds of the distance between the M morph and a 100% morph.

Having defined S, M, W top and bottom morphs for both targets and orientations, all test stimuli were created by combining these half faces. 72 test stimuli for the integration phase were constructed as follows. The W, M, and S top half faces for Target A were crossed with the W, M, and S bottom halves for Target A. The same was done for the Target B half faces. Each of these faces was presented in a normal configuration, i.e., directly above or below the other half face, or horizontally split by 60 pixels.

All stimuli were presented in both the upright and inverted orientations. Sample stimuli are given in Figure 1.

Figure 1: Example test stimuli. From left to right: upright together, upright split, inverted together, and inverted split.

Procedure

Participants completed twelve sessions of the following experiment on all but Day 1 and Day 6 of a 14 day study. Participants were told that they would see a series of faces, each of which belonged to one of two families. All family members were descended from one of two patriarchs (the target faces). Participants were told that on each trial they would see a test face followed by one target patriarch. Their task was to decide if the test face was a member of the target's family based on how much it resembled the patriarch. They were warned that on some trials they would be asked about only one half of the test face, while on others they would be asked about both halves. They were also told that test faces would appear in both together and split configurations. Participants were shown several example trials meant to familiarize them with the task and stimuli. It was explained that for trials asking about only one half they should ignore the other half face. After viewing the example trials participants underwent a training phase during which they were given auditory feedback. Once they had responded to ten consecutive trials correctly they began the integration phase of the experiment. Each individual completed one block of trials during each session. Upright and inverted faces appeared in odd and even numbered blocks, respectively. Each block contained 288 trials. Each test face, for the appropriate orientation, appeared once for each combination of target face and question type per block, with the exceptions of W/W, M/M, and S/S stimuli, which appeared twice. Test stimuli were fully crossed with the two levels of target face (Smith and Jones), and with the three levels of question type (top, bottom, and conjunction).

Trials began with a test face appearing in a random position in the middle of the screen. After two seconds the face was masked with one of two scrambled feature sets from the target faces. After 250ms the mask disappeared and a target face appeared. Participants were asked one of three questions: "Is this the family of the TOP half you just saw?", "Is this the family of the BOTTOM half you just saw?", or "Is this the family of the TOP half AND the BOTTOM half you just saw?". After giving a yes/no response they were asked, "What is the likelihood that you are correct?" on a scale of 50% to 100%. A fixed number of points were awarded for each correct choice and the observer with the highest final score received a $20 bonus.

Results

Choice data were analyzed as follows. First, yes/no responses were recoded as correct/incorrect, collapsing across target face and test face family. Since we anticipated that processing time might affect performance, the data were divided into fast and slow trials. This was done by finding the median response time (RT) for each individual and coding slower trials as *slow*, and faster trials as *fast*. We calculated deviation scores for each conjunction trial. This involved comparing the response an observer gave to the conjunction question, when presented with a given test face, to her response to the *top* and *bottom* questions, given the same test face. Mean P(correct) was calculated for each combination of top and bottom strength for each individual, separately for each combination of orientation, split, and question type. Next, for each conjunction trial the weaker of the two corresponding half face means was subtracted from the response. For example, an upright, together, M/S, conjunction trial was compared to the minimum of the upright, together, M/S, top and bottom means. These conjunction scores were then averaged within conditions for each participant. Mean scores above 0 indicate that a participant was more likely to say that both halves of a face came from a given family than they were to say that one of the halves can from the family, i.e., the conjunction fallacy.

Mean deviation scores, averaged across individuals, are presented in Figure 2. For display purposes (but not for analysis), the data were collapsed across similar strength conditions (e.g., W/S and S/W). Note that all fast scores are well above 0, indicating strong and reliable conjunction errors. The case is different for slow conditions, where most scores are very close to 0, except for W/S.

To simultaneously evaluate the effects of our various manipulations we conducted a 3 (test top strength) x 3 (test bottom strength) x 2 (orientation) x 2 (split) x 2 (speed) repeated measures analysis of variance (ANOVA) on conjunction scores. A positive intercept coefficient was obtained for ANOVA model, signifying that, in general, conjunctions were rated as more likely than one of the conjuncts. Split, but not orientation, had an effect on the optimality of information integration consistent with Hotaling et al. (2010). However, in the present experiment the split effect was only marginally significant, $F(1,8) = 3.51$, $MSE = .06$, $p < .1$. Conjunction scores were slightly higher for split faces than for together faces. A highly significant effect of speed was found, $F(1,8) = 33.2$, $MSE = 1.61$, $p < .001$. When participants responded quickly they were much more likely to commit conjunction errors compared to when they responded more slowly. Finally, there was a marginally significant top strength x bottom strength interaction, $F(4,32) = 2.39$, $MSE = .04$, $p < .1$. Previous investigations of the conjunction fallacy have found the largest effects in cases where probabilities of the conjuncts were most dissimilar. Figure 2 suggests that this was the case for slow trials, with the highest scores coming in the W/S condition and smallest appearing in conditions where top and bottom halves were roughly equal in strength.

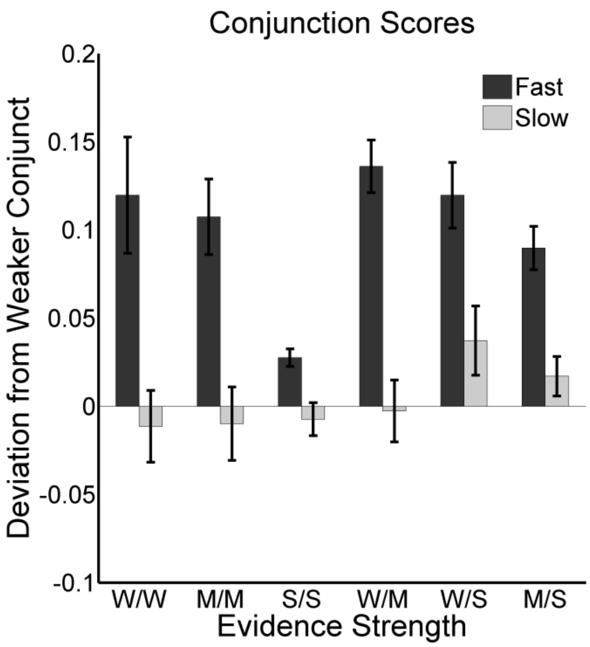

Figure 2: Mean conjunction scores across participants for the W/W, M/M, S/S, W/M, W/S, and M/S fast and slow conditions. Error bars are between-subject standard errors.

Discussion

The aim of the research described above was twofold. First, we were interested in whether or not observers would produce errors analogous to the conjunction fallacy in a perceptual decision making paradigm. In judgment and decision making, these errors are typically taken as evidence that humans process information in a way that differs fundamentally from normative models. On the other hand, the vast majority of research in perceptual domains predicted near-optimal performance in accordance with the conjunction law. Our results suggest otherwise. Participants often indicated that it was more likely that both halves of a test face came from a given family than it was that one of the halves came from the family. These violations of the conjunction law were found across all levels of evidence strength, orientation, and split. Most striking is the fact that participants committed the fallacy in the upright together condition. There is significant evidence that humans are inherently expert face processors and this condition was designed to closely mimic natural face perception. The prevalence of conjunction errors in even this situation calls into question any claims that perception can be well described by a "rational" model in all situations. This result gives hope for bridging the gap between research in judgment and decision making and perceptual domains.

The second goal of our experiment was to address the inconsistency in empirical results across these different fields by investigating what stimulus and processing factors affected the optimality of information integration. We hypothesized that processing may proceed through different channels, and that the automaticity of integration may underlie divergent empirical results.

Consistent with Hotaling et al. (2010) we found a marginally significant effect of split. This effect is somewhat puzzling because splitting a face apart was hypothesized to promote controlled processing, but it had a negative effect on the optimality of integration. Since this effect was small it is probably best not to overinterpret it. Individual difference may exist, with some observers being more sensitive to the manipulation than others.

We flipped faces upside-down in some conditions with the expectation that this would inhibit automatic processing of information from the top and bottom halves. Consistent with the findings of Hotaling et al. (2010) this manipulation did not significantly affect performance.

The speed at which participants responded did have a highly significant effect. Using each individual's median RT as a baseline, trials with a faster RT were much more likely to produce large conjunction errors compared to slower trials. This is consistent with our hypothesis that rapid automatic processing of the image features may result in a more holistic judgment based on global similarity between the test and target stimuli. In contrast, when participants took longer to respond they appeared better able to combine the probabilities of top and bottom halves in accordance with the conjunction law. It would appear that additional processing in needed to separate out this information from the initial holistic representation. It is worth pointing out that this pattern is opposite to what is typically found in psychological experiments, where errors produce larger RTs than correct responses. In the present experiment a fast RT did not indicate that the problem was easy, but rather that more automatic, perhaps more holistic, processing was used to combine the evidence from the top and bottom halves.

We adopted the automatic/controlled dichotomy for the sake of convenience and to motivate experimental manipulations. Our results suggest that this outlook is overly simple. Some factors thought to inhibit automatic processing had no effect (inversion) while others increased conjunction errors (splitting). Response times, which intuitively reflect the automaticity of a task, indicate that slower (and presumably more controlled) processing actually produces performance more in line with rational norms.

The experiment presented here is just one example of a new research program that allows researchers to address issues of information integration arising from both perceptual and judgment and decision making domains. The prevalence of conjunction errors across all stimulus conditions suggests that the effect is robust and so there is still more work to be done to fully bridge the span between optimal integration in perception and sub-optimal performance in judgment. This paradigm can also be extended beyond the conjunction fallacy in future investigations to other paradoxical phenomena found in judgment research. Such work would continue down the path laid out here toward reconciling the qualitatively different results characteristic of information integration across various cognitive domains.

Acknowledgments

Special thanks to Andrew Cohen for help in developing our experimental paradigm. J.M.H. was supported by an NSF GRF. J.R.B. was supported by NIMH Cognition and Perception MH068346. R.M.S was supported by NIMH Grant 1 R01 MH63993.

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98*(3), 409-429.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105*(3), 442-481.

Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14*(1), 33-53.

Ashby, F. G., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 16*(3), 598-612.

Burgess, A. E., Wagner, R. F., Jennings, R. J., & Barlow, H. B. (1981). Efficiency of human visual signal discrimination. *Science, 214,* 93-94.

Cohen, A. J. (2006). Contribution of invariants, heuristics, and exemplars to visual perception of relative mass. *Journal of Experimental Psycology: Human Perception and Performance, 32*, 574-598.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology, 56*, 5-18.

Epstein, S., Lipson, A., Holstein, C., & Huh, E. (1992). Irrational reactions to negative outcomes: Evidence for two conceptual systems. *Journal of Personality & Social Psychology, 62*, 328-339.

Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General, 127*(2), 107-140.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve bayesian reasoning without instruction: frequency formats. *Psychological Review, 102*, 684-704.

Gigerenzer, G., & Todd, P. M. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Gilovich, T., Griffin, D., & Kahneman, D. (Eds.). (2002). *Heuristics and biases: The psychology of intuitive judgment*. Oxford: Oxford University Press.

Hammond, K. R. (2000). *Judgments under stress*. New York, NY: Oxford University Press, Inc.

Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology, 52*, 653-683.

Hotaling, J. M., Cohen, A. L., Busemeyer, J. R., & Shiffrin, R. M. (2010). Information integration in perceptual decision-making. *Manuscript submitted for publication. Indiana University*.

Jacoby, L. L. (1991). A process dissociation framework: separatig automatic from intentional uses of memory. *Journal of Memory and Language, 30*, 513-541.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as bayesian inference. *Annual Review of Psychology, 55*, 271-304.

Massaro, D. W. (1998). *Perceiving talking faces: from speech perception to a behavioral principle*. Cambridge, MA: MIT Press.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115*(1), 39-57.

Runeson, S., Juslin, P., & Olsson, H. (2000). Visual perception of dynamic properties: Cue heuristics versus direct-perceptual competence. *Psychological Review, 107*(3), 525-555.

Shanteau, J. (1975). Averaging versus multiplying combination rules of inference judgment. *Acta Psychologica, 39*, 83-89.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*(2), 127-190.

Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM--retrieving effectively from memory. *Psychonomic Bulletin and Review, 4*(2), 145-166.

Sides, A., Osherson, D., Bonini, N., & Viale, R. (2002). On the reality of the conjunction fallacy. *Memory and Cognition, 30*(2), 191-198.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3-22.

Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral & Brain Sciences, 23*, 645-726.

Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. In M. S. Kerns, S. A. Solla & D. A. Cohn (Eds.), *Advances in Neural Information Processing Systems* (Vol. 11, pp. 59-68). Cambridge, MA: MIT Press.

Tenenbaum, J. B. (2000). Rules and similarity in concept learning. In S. A. Solla, T. K. Leen & K. R. Muller (Eds.), *Advances in neural information processing systems* (Vol. 12, pp. 59-65). Cambridge, MA: MIT Press.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90*(4), 293-315.

Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, *16*, 747-759.