COGNITIVE ETHOLOGY:
GIVING REAL LIFE TO ATTENTION RESEARCH

Alan Kingstone
University of British Columbia

Daniel Smilek
University of Waterloo

Elina Birmingham
University of British Columbia

Dave Cameron
University of British Columbia

&

Walter Bischof
University of Alberta

Address correspondence to:

Alan Kingstone
2136 West Mall
University of British Columbia
Vancouver, British Columbia
Canada
V6N 1Z4

Email: alan.kingstone@ubc.ca
Phone: (604) 822 - 9230

ABSTRACT

Studies of attention, often conducted in artificial laboratory experiments, may have limited validity when performance in the natural world is considered. For instance, for over two decades investigations of "reflexive" and "volitional" attention have tended to be grounded in methodologies that do not capture the demands of attention in everyday life. Recent studies suggest these laboratory investigations have lost touch with real life contexts and accordingly they may generate fundamental misunderstandings regarding the principles of human attention and behaviour. We identify the basic assumptions of laboratory research that has led to this state of affairs, and suggest a new set of assumptions that lead to a new research approach, which we call "cognitive ethology". The implication is that if one is to understand human attention in everyday life then research needs to be grounded in the natural world and not in experimental paradigms.

Patrick Rabbitt has never been one to accept the status quo. None of his students escaped this strong part of his personality, and for many of us, it helped to define who we became. One of Pat's most enduring lessons centered around the idea that the models in psychology are largely "snapshots" of idealized static cognitive sytems. Pat often punctuated this point by noting that experimental psychologists routinely throw away the first set of trials in an experiment as "practice trials". As Pat was fond of saying, these initial trials routinely produce the largest performance changes that will be found in a study. And yet, these initial trials are discarded and left unanalyzed because researchers desire stable systems that can be controlled, manipulated and modeled.

This issue raised by Pat has plagued me over the years because seeded deeply within it was the notion that experimental psychologists were not really getting it right, nor were they going to get it right doing research the way they were doing it. That is, research in the pursuit of stability and control was not going to tell us what we really want to know – how people function in real life where things are highly variable and often outside the domain of experimental control. The present chapter represents an initial attempt by my lab to take on Pat's challenge to get things right and to learn how people function in real life. Our particular area of interest is human attention.

A LABORATORY PARADIGM FOR STUDYING ATTENTION

We begin by closely considering one of the most well known laboratory paradigms for studying attention: The Posner cueing paradigm (Posner, 1978; 1981). In this paradigm, a central fixation dot that is flanked by two boxes is presented at the center of a computer screen. The task is simply to press the spacebar on a computer keyboard as quickly as possible when a visual target object appears inside one of the boxes. This target object is preceded by an attentional cue, which is either a brief peripheral flash surrounding one of the two boxes or a central arrow pointing toward one of the boxes (see Figs. 1a and 1b). The standard and highly robust finding is that a target is detected fastest when it appears in the box that was cued by the peripheral flash or central arrow. On the assumption that the brain processes attended items more quickly than unattended items, it is concluded that target detection time was speeded for a target at the cued location because attention had been committed to the box that was cued.

--- Insert Figure 1 about here ---

Countless studies of this sort have led to the conclusion that there are two categories of attention, exogenous (reflexive) attention and endogenous (volitional) attention, both of which can be manipulated and measured by the Posner paradigm. When one of the two boxes is flashed briefly, as depicted in Figure 1a, attention is considered to be oriented reflexively to the box that brightened. This attention shift is thought to be reflexive because people are faster to detect a target in the cued box even when flashing the box does not predict where the target will occur (e.g., the target appears in the cued box 50%

of the time and in the uncued box 50% of the time). Attention can also be oriented volitionally. When a central arrow points toward one of the two boxes, as depicted in Figure 1b, attention is thought to be oriented volitionally to the box pointed at by the arrow. This attention shift is considered volitional because people are thought to be faster to detect a target in the cued box only when the arrow predicts where the target will occur (e.g., the target appears in the cued box 80% of the time and in the uncued box 20% of the time).

THE POSNER PARADIGM AND REALITY

It is perhaps instructive at this point to compare the displays used in the Posner paradigm, which are shown in Figure 1, with the type of scenes encountered outside the laboratory, examples of which are shown in Figure 2. Even a cursory comparison begs the following type of question: To what extent does the simple, impoverished and highly artificial experimental task of detecting a light within a cued or uncued box have to do with the many complex, rich real-life experiences that people share? In other words, what does the cuing paradigm have in common with everyday situations such as children playing under the supervision of an adult, a father teaching a son how to tie and tie, or a little girl watching another two friends kiss? On the face of it, not very much.

--- Insert Figure 2 about here ---

Even a cursory look at the naturalistic scenes shown in Figure 2 yields many hypotheses regarding attention that are unlikely to ever be generated from laboratory studies of attention. For instance, inspection of the scenes suggests that people's direction of gaze might be a critical cue, perhaps even a reflexive one, for orienting attention in the real world. In 1998, Chris Friesen and I tested this idea by modifying the Posner paradigm in two important ways. First, arrows pointing to the left and right were replaced by a schematic face that looked left or right. Second, the predictive value of the central cue was eliminated; that is, eye direction did not predict where a target item would appear (see Fig. 1c; also Friesen and Kingstone, 1998). Note that because the eyes were centrally located and spatially nonpredictive, the traditional line of thinking predicted that they would not lead to shifts of attention. The result was that, contrary to this prediction, eye gaze did trigger shifts of attention, with all types of responses (target detection, target localization and target identification) being enhanced almost immediately for targets that appeared at the gazed-at location. This rapid onset of an attention effect, and the fact that it occurred in response to a spatially nonpredictive stimulus, demanded the conclusion that the attentional shift was reflexive (Cheal & Lyon, 1991; Müller & Rabbitt, 1989; Jonides, 1981).
This conclusion, that centrally nonpredictive eye-gaze can direct trigger reflexive shifts of attention, led us to reconsider the fundamental notion that arrows only direct attention when they are spatially predictive, i.e., the shifts are volitional in nature. Since a classic study by Jonides (1981, Experiment 2), which failed to find positive evidence that nonpredictive central arrows will trigger a shift in attention, researchers have assumed that arrows do not produce a shift in attention unless they predict where an item will appear. We considered the possibility, for the first time, that central arrows might be much like eyes and

that they might orient attention even when they are not predictive. We tested this possibility by replacing spatially nonpredictive eyes with spatially nonpredictive arrows (Fig. 1c; Ristic, Friesen, & Kingstone, 2002). The results were unequivocal. People attend to where arrows point even when they know that the arrows do not predict where a target will appear. In other words, like eyes, arrows produce a reflexive shift in attention to the cued location. This result has been confirmed by several investigators (Tipples, 2002; Friesen, Ristic & Kingstone, 2004; Hommel et al. 2001) establishing that attention can be directed reflexively to locations indicated by nonpredictive arrows. The failure of Jonides (1981) to obtain this significant outcome may be attributed to several factors, including his arrow cue being hard to discriminate (it was flashed for only 25 ms) and his study lacking sufficient experimental power (fewer than 10 participants were tested and RTs at the cued location were based on less than 20 trials).

We believe that the recent experiments on eye gaze and nonpredictive central arrows have important implications for laboratory studies of attention. Specifically, the findings indicate that the Posner cuing paradigm is fundamentally flawed. Central cues such as eyes and arrows, which were assumed to tap volitional attention, actually do not reflect volitional orienting. Thus, what was taken as a fundamental truth turned out to be completely in error. How could this mistake happen? How could this error be overlooked for over 20 years of research? We believe that the answer rests with the fact that laboratory research in the field of attention is based on two crucial assumptions – both of which are problematic.

ASSUMPTIONS UNDERLYING LABORATORY RESEARCH

Studies of attention in the laboratory are grounded on two basic assumptions. One is that human attention is subserved by processes that are stable across different situations. For example, the processes that are studied in the lab are assumed to be the same as the processes that are expressed in the real world. Second, one can maximize analytical power of a process by minimizing all variability in a situation save for the factor being manipulated. Note that the first assumption enables one to lay claim to processes in the real world without ever leaving the lab. And the second assumption demands that the laboratory situation be as controlled as possible. Ironically, the combination of these two assumptions has the effect of driving researchers further away from real life situations into highly contrived and artificial laboratory environments, all the while seducing the investigators into the belief that they are getting closer and closer toward a true understanding of how attention operates in real life.

While the assumptions of process-stability and situational-control are commonly held and readily applied in studies of attention, adopting them comes with a high degree of risk. The assumption of stability, for example, eliminates any need or obligation by the scientist to confirm that the factors being manipulated and measured in the lab actually express themselves in the real world. The field does of course check routinely that the effects being measured are stable within the lab environment, by demanding that results in the lab be replicable. Unfortunately, a result that is stable within a controlled laboratory environment does not necessarily entail that it is stable outside the lab. Indeed, there are many examples within the field of human attention indicating that even the most minor changes within a laboratory situation will compromise the replicability of an effect (e.g., Soto-Faraco, Morein-Zamir & Kingstone, in press; Kingstone et al., 1993). Indeed, the fragility of laboratory findings whereby small changes in an experimental situation results in radically different experimental outcomes, should perhaps not be particularly surprising. There is now a growing body of literature indicating that "process stability" is tied intimately to the specifics of a situation, with brain re-configurations occurring continuously in response to even subtle environmental changes. Neisser (1976) referred to the dynamic re-configurations as "schemata", Mosell (1996) referred to them as "task-set reconfigurations", and Di Lollo et al. (2001) have referred to them as "configurable input filters". Thus, while the assumption of process stability is remarkably convenient because it allows one to think that the work in the lab has some relevance to real world situations, the fact is that there is a very real risk, indeed a high likelihood, that the process being studied in the lab will not exist outside the lab.

Associated with the above is the fantastic risk of wasting many years and dollars conducting research that is replicable but which has little, if anything, to do with cognitive processes beyond the lab in real life environments. This follows because by simply assuming a process is stable across situations – specifically from the lab to real life –means that one never has to check whether the results in the lab are at all relevant to real life performance. This seems to be the reason why for more than 20 years researchers failed to notice that the Posner cuing paradigm, while producing results that were highly replicable, did not measure what people believed they were measuring. It was only after considering the results derived in the Posner paradigm against real life situations (e.g., how eye-gaze might influence attention) that this error was detected (Kingstone et al, 2003).

The problem gets even worse, however, because it is often the case that attempts to test the first assumption of stability against real life situations are immediately met with apparently insurmountable obstacles. These obstacles arise from the second assumption of experimental control. The first obstacle is that processes in the lab often become defined by the experimental controls that were used to examine them. In addition, if this obstacle were somehow overcome and the controls that define a phenomeon were reproduced in the real world, a researcher is immediately posed with the challenge of making the case that the data collected in a real life situation is in fact a manifestation of the process that was measured previously in the lab. This is a daunting, and perhaps ultimately impossible, obstacle to surmount. This is because data that are attributed to a particular process in the lab can always be re-attributed to other factors that were left free to vary in a real life situation.

We believe that these two obstacles actually reflect a more fundamental problem that arises from the assumption of control when it is applied to nonlinear systems like human cognition and attention. While experimental control can be effective at revealing basic characteristics of simple linear systems, general systems theory has established that experimental control is unlikely to be effective at revealing important characteristics of complex, non-linear systems such as the human attention system. This is because certain characteristics of complex systems are only revealed, or emerge, when several variables vary together in highly specific ways (see Ward, 2002; Weinberg, 1975). This is precisely what is not allowed to occur in highly controlled laboratory situations.

Considered together, it appears that there are both practical and principled reasons to question whether selecting tasks based on the assumptions of stability and control is likely to inform us about cognitive processes as they are expressed in real life situations. Moreover, it is sobering to contemplate the fact that these two assumptions can lead researchers to unwittingly and unthinkingly commit their lifetime to studying processes that may be expressed only within artificial lab situations.

COGNITIVE ETHOLOGY: MOVING TOWARDS A DIFFERENT WAY OF STUDYING ATTENTION

In light of these considerations our most recent work has begun to move away from studying human attention in standard artificial laboratory paradigms. While we are still conducting laboratory-based investigations we have sought to conduct studies that remove as many constraints on the behaviour of our participants as possible. Thus rather than making causal claims about fundamental processes, our goal initially is simply to <u>observe</u> and <u>describe</u> behavior as it occurs. As such, our new approach, what we call <u>cognitive ethology</u>, is characterized by giving up the assumptions of stability and control for the scientific and objective study of human behavior under natural conditions. Our laboratory tasks, though not identical to real world tasks, have increased substantially in their ecological validity (see Koch, 19xx; Neisser, 19xx; Hutchins, 1995). Our hope is that by using this new cognitive ethological approach we might uncover new aspects of attention that have eluded researchers wedded to standard laboratory-based paradigms.

In our first attempts to move in this new direction we sought to study a task that individuals appear to engage in effortlessly every day -- that of inferring where other people are attending. Our specific aim was simply to get a feeling for the types of cues that people use when they infer the attentional states of others. In our study two groups of participants viewed photographs that depicted one or more individuals involved in everyday behaviors, such as painting or playing basketball. An example of a picture used in the study is shown in Figure 3a.

--- Insert Figure 3 about here ---

One group of participants (the "inference" group) was required to view the picture and then answer the following question: Where are the people in the picture directing their attention, and how do you know? Another group (the "describe" group) was required to view the pictures and simply describe each picture. Our goal was not to make inferences about underlying cognitive mechanisms but to simply observe and describe behavior as it occurred. We observed and described what information participants used to infer the attentional states of the individuals in each picture using a combination of objective (third-person) and subjective (first-person) measures. The objective measure involved monitoring participants' eye movements and examining the patterns of eye fixations that emerged as they viewed each picture. We reasoned that participants would fixate the information that they considered to be important for either making inferences about attentional allocation in the picture or for describing the pictures. In addition to observing the third-person eye fixation data we also measured the first-person reports that participants provided after they viewed each image. Of particular interest were the subjective reports of those participants who inferred where the individuals in the pictures were allocating their attention and who also reported what information they used to make such inferences. We thought that participants' subjective reports might provide a valuable insight into the cues that they were use when making inferences about the attentional states of others.

As our approach did not involve controlling behavior, we were challenged with how to effectively analyse the complex patterns of behavior that emerged as participants engaged in our open-ended investigation. To examine the eye-movement data we divided the scenes into regions of interest (e.g. the eyes, heads, and bodies of people in the picture). Once a scene was divided into regions, we evaluated both the static fixation patterns (e.g., fixation frequency, fixation duration) for each region, as is done in most studies of scene viewing (e.g., Buswell, 1935; Antes, 1974; Henderson & Hollingworth, 1999), as well as the transition-patterns between various regions (see Liu, 1998). The regions were comparable across the two groups of participants, which then allowed us to directly compare how the patterns of eye fixations among these regions differed as a function of whether participants were inferring attentional allocation in the picture or merely describing the picture.

Figures 3b and 3c provides an illustration of the type of performance patterns we obtained. Figure 3b shows the distribution of fixations (coloured dots) that occurred in the describe condition and Figure 3c shows the fixation pattern for the inference condition. The white arrows indicate transitions between different regions that occurred more oftenthan would be expected by chance.

The fixation plots in Figures 3b and 3c suggest that individuals were primarily fixating the faces of the individuals in the picture, the drawings of the artists, and the tripod held by the model. Thus, the fixation data suggest that people use information from faces and the actions that the people in the picture are performing both to describe the picture and to make inferences about where the people in the picture are directing their attention. A comparison of the fixation data for the inference group (Figure 3c) to the fixation data for the describe group (Figure 3b) revealed some particularly interesting findings. First, observers who inferred the attentional states of people in the picture fixated on the eyes of these people more frequently than observers who described the picture. That is, relative to the total number of fixations, observers in the inference group fixated the eyes of the model and artists more often than observers in the describe group. Second, this effect was specific to the eyes. When other body regions (such as the head, torso, arms, and legs) were analyzed, the inference and describe groups showed no difference in fixation frequencies. Thus, there appears to be a specific increase in the use of eye gaze information when observers are inferring the attentional states of others.

Examination of the transitions suggests that while there are some significant transitions from people's faces to the objects that the people are attending, this certainly did not occur for all of the faces. It seems that significant transitions mirroring gaze paths of the people in the picture only occurred in situations where direction of gaze was ambiguous, such as the direction of gaze of the model. A comparison of the significant transitions across the describe and inference conditions is also revealing. One striking observation that can be made by comparing the pictures in Figure 3 is just how similar the fixation areas and the significant transitions are across the two conditions. We think these findings are interesting when compared with the findings of Yarbus (1967) who reported that scan patterns differ substantially when participants are given different viewing instructions. The similarity in scan paths between the describe and inference conditions in our study suggests the tentative hypothesis that when participants view pictures with the purpose of describing them,they might actually be inferring the attentional states of the individuals in the picture.

Importantly, this is precisely what was indicated by the subjective reports. Moreover, the subjective report data provided an interesting insight into the cues that people use to infer attention that would not be expected either from an analysis of the eye-movement data alone or from previous studies of attention. Notably, when participants were asked how they knew where people in the picture were directing their attention, participants reported using cues such as the body orientation of the people in the picture and what appeared to be the focus of the scene, e.g., the model in the painting scene. In addition to these observations, it is worth noting that many of the subjective reports echoed what was suggested by the eye movement data. For instance, participants reported that direction of gaze of the people in the scene was an important indicator of where people appeared to be attending.

Although the findings that we are reporting here merely represent our first, preliminary steps in a new research direction, they have revealed already several interesting points worthy of further investigation, e.g., are people always inferring attentional states of others when first perceiving scenes? Importantly we expect that performance differences will emerge not only with changes in the task that participants are asked to perform, as demonstrated here, but also with the type of scene that is shown to participants and with the type of participant that is viewing a scene. In keeping with our new approach, the challenge is not to eliminate this variance by imposing experimental control, but rather to observe, describe, and finally understand it.

THE ULTIMATE GOAL: STUDYING ATTENTION IN THE REAL WORLD

Although it is difficult to leave the comfortable confines of a controlled laboratory setting, our ultimate goal is to study attention outside the laboratory as individuals perform everyday tasks in their natural environments. Rather than being constrained by our laboratory settings and their associated paradigms, we are moving in the direction of exploring how people behave as they function within the environment one ultimately intends to understand -- the real world as it exists outside the laboratory. Of course, such an approach involves studying what people actually do in everyday life, how they appear to be doing it, what they think that they are doing, how they feel about what they are doing, and so forth, and looking for ways of understanding these performance outcomes.

We believe that a central component to studying attention in the real world is to become comfortable with giving up the assumptions of process stability and situational control, as we have begun to do in the study described briefly above. The goal of this approach is to simply observe and describe the complex patterns of behaviour that people produce as they engage in complex everyday tasks, in order to ultimately gain a true understanding of human attention and performance.

A Nature publication by Land and Lee (1994) provides another excellent illustration of this approach. These investigators were interested in understanding where people look when they are steering a car around a corner. This simple issue had obvious implications for human attention and action, as well as for matters as diverse as human performance modeling, vehicle engineering, and road design. To study this issue, Land and Lee monitored eye, head, and steering wheel position, as well as car speed, as drivers navigated a particularly tortuous section of road. Their study revealed that drivers rely on a "tangent point" on the inside of each curve, seeking out this point 1-2 seconds before each bend and returning to it reliably. This finding was new to the field, and interestingly, the drivers themselves were unaware of the fact that they were searching for and applying a tangent point when navigating curves.

The study conducted by Land and Lee illustrates elegantly that meaningful research can be conducted without falling into the standard experimental assumptions of stability and control. By stating that one is interested in understanding how an individual performs a particular task in the real world, like driving around a corner, one is implicitly acknowledging that this task may be unlike any other task that one performs within the specific domain of driving; and indeed, it may have no counterpart in any other cognitive domain. Thus it is neither claiming to be a model task for other situations, nor is it assuming that there is a model task that can speak to this particular driving situation. In other words, a real world research approach rejects the assumption of process stability. In doing so it assumes that processes may be contextualized to the situation within which they occur. As noted previously, there is a wealth of evidence, ranging from Neisser's "schemata" to Di Lollo's "configurable input filters", indicating that this is precisely the case.

The Land and Lee study is also important because by choosing to measure performance as it occurred naturally in the environment, they were rejecting the standard a priori assumption that any variance that is not manipulated experimentally is something to be controlled and minimalized. This alternative way to deal with variance, to let it occur naturally and measure it, requires the assumption that variance may reveal key characteristics of cognitive processing. In other words, it requires the assumption that variance is part of the cognitive signal and as such it must be understood. Interestingly, this assumption dovetails with the most basic tenet of general systems theory, that complex systems are only revealed when several variables are permitted to co-occur.
In sum, the real world approach rejects the standard of assumptions of stability and control, and in their place we find a commitment to understanding the situation, and the variance within that situation. According to the real world approach, the initial job of the researcher is simply to observe and measure what people do in the situation of interest. Of course, this approach is unlikely to be of much value in artificial laboratory situations where human behaviour is typically highly constrained. For example, in a typical attention experiment people are only allowed to move one finger to press one key, with all other movements being banned including even minor movements of the eyes. Yet observation of real world behaviour is very different. As Koch puts it "description is no lowly or easy task; it is in fact the very basis -- indeed, the flesh -- of non-spurious knowledge" (Koch, 1999, p. 27). In other words, description of attentional behaviour in the real world is intrinsically valuable because it is grounded in what people really do and as such the questions related to it can be both anchored and meaningful rather than abstract and trivial.

Just as there were practical problems for the assumptions of stability and control when they are applied to understanding real world phenomena, one finds that there are also practical problems for the assumptions of situation and variance. The key problem is quite simply that it is very hard to do research at the real world level. It is hard for several reasons. First, it is difficult because there are no "off the shelf" model-tasks to use when one conducts research in this way. Hence, one cannot, for instance, simply manipulate the Posner cuing paradigm (Hunt & Kingstone, 2003) or the visual search paradigm (Fecteau, Enns & Kingstone, 2001; Eastwood, Smilek, & Merikle, 2001; Smilek, Eastwood & Merikle, 2000) or the attentional blink paradigm (Giesbrecht, Bischof & Kingstone, 2003), and believe that one is gaining new insights into how people allocate their attention in everyday life. Instead, one has to spend a good deal of time simply observing and describing what people appear to be doing with their attention. And because one cannot control what people do, one finds that there is a tremendous amount of variation in the behaviour that people produce, not only between people but for the same person at different times. It is also very difficult because there is very little data in the literature on how people actually allocate their attention in the real world rather than in artificial laboratory environments. This means that what questions and approaches are most interesting and likely to bear fruit are largely unknown. It also means that there may be little or no previous work performed on how to go about analyzing the data one collects, and therefore, one is often having to create new tools to understand the data that have been collected.

We would argue, however, that these problems can be viewed as exciting opportunities for researchers interested in discovering how attention is really allocated in the world. All the data one collects, all the questions that one explores and answers, provide a foundation for future investigations and a benchmark against which all other studies will need to be measured.

SUMMARY AND CONCLUDING COMMENTS

The goal of attention research is, and has always been, to understand how human attention operates as people behave in their everyday lives. As we have noted, however, in practice, the current laboratory-based research approach does not seem to be reaching this goal, and in fact, it may be moving the field further away from it. Our discussion of the underlying assumptions of the standard laboratory-based research approach – the assumptions of stability and control – brings to light both principled and practical reasons for why the laboratory-based approach is unlikely to reach its goal of understanding how attention is allocated in everyday life.

We have articulated what we believe to be an alternative approach to the standard lab-based methodology. This alternative approach replaces the assumptions of stability and control with the assumptions related to understanding individuals as they behave unconstrained in their natural environments and the variance that emerges with it. These latter assumptions are polar-opposites to the traditional assumptions of stability and control.

The present chapter also suggests principled and practical ways to focus attention research and to identify research questions that are meaningful. It is often the case that research questions are paradigm-driven or perhaps driven by a set of informal introspections of the researcher. In contrast to this usual way of generating research questions, we suggest that research might benefit by being grounded in systematic observations of real-world behavior, and that research concepts might best be grounded in what people really do objectively and what they perceive themselves to be doing. Indeed, William James clearly identified lay understanding as the starting point for attention research in his much quoted statement "everyone knows what attention is". Ironically, although many researchers use this quote when attempting to define attention, the field has rarely committed its research effort to determining what people think attention is and how they allocate it in everyday life. To be clear, we are not suggesting that attention research simply be more applied. Though it might focus on applied problems, such as driving cars and flying planes, attention research ultimately needs to study every aspect of attention as individuals define it and as it is used in real life settings. Pure research on the nature of real cognitive practices is needed.

In this chapter we have emphasized an approach that we call cognitive ethology. By observing and measuring what people really do in their natural environments one can begin to define the problem space in a manner that becomes grounded in real life situations and not contrived laboratory environments. This mean that subsequent laboratory investigations can be protected from pursuing behaviour and questions that are merely paradigm-driven and paradigm-specific. It also means that one is free to ask questions that are meaningful and relevant to real life rather than being constrained and limited to artificial laboratory environments. Our hope is that our cognitive ethological research approach will lead to pure research that will ultimately expand and deepen our understanding of human cognition.

REFERENCES

Antes, J.R.(1974). The time course of picture viewing. *Journal of Experimental Psychology*, *103*, 62-70.

Buswell, G.T. (1935). How people look at pictures. University of Chicago Press, Chicago.

Duda, R., P. E. Hart & D. G. Stork (2000) Pattern Classification. 2nd Edition, New York: Wiley.

Eastwood, J.D., Smilek, D., & Merikle, P.M. (2001). Differential attentional guidance by unattended faces expressing positive and negative emotion. *Perception & Psychophysics, 63(6)*, 1004-1013

Fecteau, J. H., Enns, J. T., & Kingstone, A. (2001). Competition-induced visual field differences in search. *Psychological Science*, *11*, 386-393.

Friesen, C.K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin & Review, 5(3*), 490-495.

Friesen, C. K., Ristic, J., & Kingstone, A. (2004). Attentional effects of counterpredictive gaze and arrow cues. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 319-329.

Giesbrecht, B., Bischof, W. F. & Kingstone, A. (2003). Visual masking during the attentional blink: Tests of the object substitution hypothesis. *Journal of Experimental Psychology: Human Perception and Performance, 29*, 238-258.

Henderson, J.M. & Hollingworth, A. (1999). High-level scene perception. *Annual Review of Psychology*, *50,* 243-271.

Hommel, B., Pratt, J., Colzato, L. & Godign, R. (2001). Symbolic control of visual attention *Psychological Science*, *12*, 360-365

Hunt, A. & Kingstone, A. (2003). Inhibition of return: Dissociating attentional and oculomotor components. *Journal of Experimental Psychology: Human Perception and Performance*, *29*, 348-354.

Hutchins, E. (1995). Cognition in the wild. MIT Press.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In J. B. Long and A. D. Baddeley (Eds.), *Attention and Performance IX* (pp. 187-203). Hillsdale, NJ: Erlbaum.

Kingstone, A., Smilek, D., & Ristic, J. (2003). Attention, Researchers! It Is Time to Take a Look at the Real World. *Current Directions in Psychological Science, 12(5)*, 176-180.

Koch, S. (1999). *Psychology in Human Context: Essays in Dissidence and Reconstruction*. D. Finkelman and F. Kessel (Eds). University of Chicago Press.

Land, M. F.; Lee, D. N. (1994). Where we look when we steer. *Nature, 369(6483*), 742-744.

Liu, Andrew (1998). What the driver's eye tells the car's brain. *In:* Eye guidance in reading and scene perception. Underwood, Geoffrey; Oxford, England: Elsevier Science Ltd, 1998. pp. 431-452. [Chapter]

Neisser, U. (1976). *Cognition and Reality.* San Francisco: Freeman.

Posner, M.I. (1978). *Chronometric explorations of mind*. Lawrence Erlbaum Associates, Inc., Hillsdale, N.J.

Ristic, J., Friesen, C.K., & Kingstone, A. (2002). Are eyes special? It depends on how you look at it. P*sychonomic Bulletin & Review, 9(3),* 507-513.

Smilek, D., Eastwood, J. D., & Merikle, P. M. (2000). Does unattended information facilitate change detection? *Journal of Experimental Psychology: Human Perception and Performance, 26*(2), 480-487.

Tipples, J. (2002). Eye gaze is not unique: Automatic orienting in response to uninformative arrows. *Psychonomic Bulletin & Review, 9*, 314-318.

Treisman, A.M. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12(1),* 97-136.

Ward, L. M. (2002). *Dynamical Cognitive Science*. Cambridge: MIT Press.

Weinberg, G. M. (1975). *An Introduction to General Systems Thinking*. New York: Wiley.

Yarbus, A. L (1967). *Eye movements and vision* (B. Haigh,Trans.). New York: Plenum Press. (Original work published 1965)

ACKNOWLEDGEMENTS

This work was funded by the Natural Sciences and Engineering Research Council of Canda, and the Michael Smith Foundation for Health Research. Thank you also to Patrick Rabbitt for the many wonderful years spent in his lab, and to Peter McLeod for his insightful comments on an earlier version of this chapter.

FIGURE CAPTIONS

FIGURE 1. Variations of the Posner paradigm. Each panel presents three stages of a typical trial (start, cue onset, and target onset); in these examples, the target (a small black square) appears at the cued box. In one variation of the paradigm (a), at the start of each trial, a central fixation dot is flanked by two squares (boxes). The left or right box is cued by a brief flash (illustrated by the thick black line), and then a target (the black square) is presented. The task is to press a key as quickly as possible when the target appears. The target appears in the cued (flashed) box 50% of the time and in the uncued (not flashed) box 50% of the time. Thus, the cue does not predict where the target will appear. In another variation (b), the left or right box is cued by an arrow pointing toward it, and the target appears in the cued box 80% of the time and in the uncued box 20% of the time. Thus, the cue predicts where the target will appear. In a third variation (c) the left or right box is cued by eyes looking toward it, and the target appears in the cued box 50% of the time and in the uncued box 50% of the time. Thus, the cue does not predict where the target will appear.

FIGURE 2. Prototypical examples of everyday scenes that occur in life.

FIGURE 3. Aggregate eye movement data from one of the scenes we presented to participants. Figure 3a depicts the original scene presented to participants. Figure 3b shows data from participants who were asked to "describe the picture". Figure 3c show data from participants who were asked: "Where are people in the picture directing their attention, and how do you know?" The arrows depict transitions between regions that occurred more often than would be expected by chance.

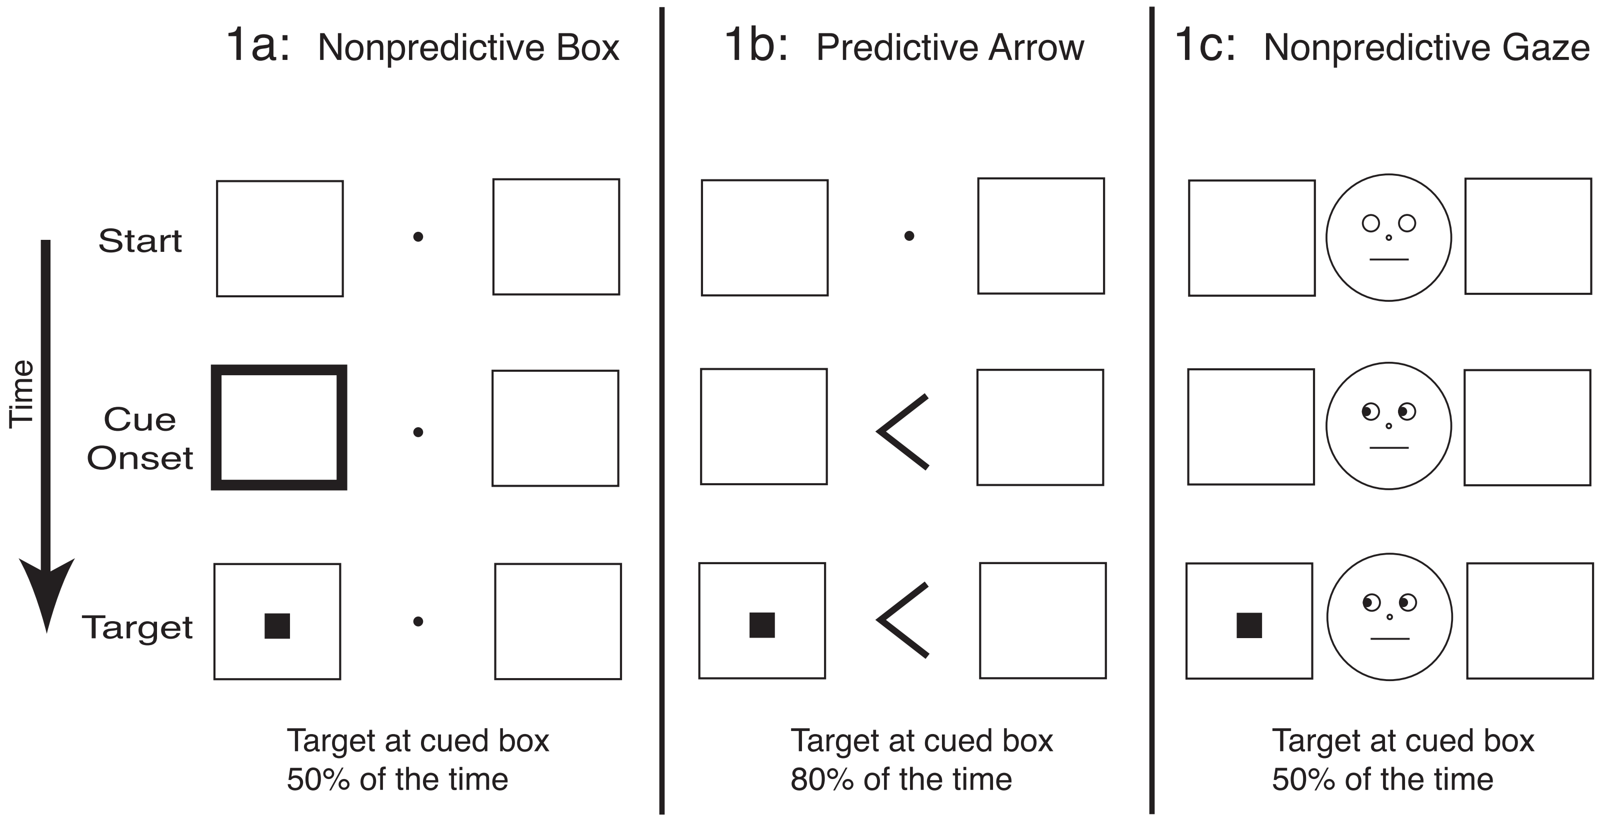
1a: Nonpredictive Box
1b: Predictive Arrow
1c: Nonpredictive Gaze
Time
Start
Cue Onset
Target
Target at cued box 50% of the time
Target at cued box 80% of the time
Target at cued box 50% of the time

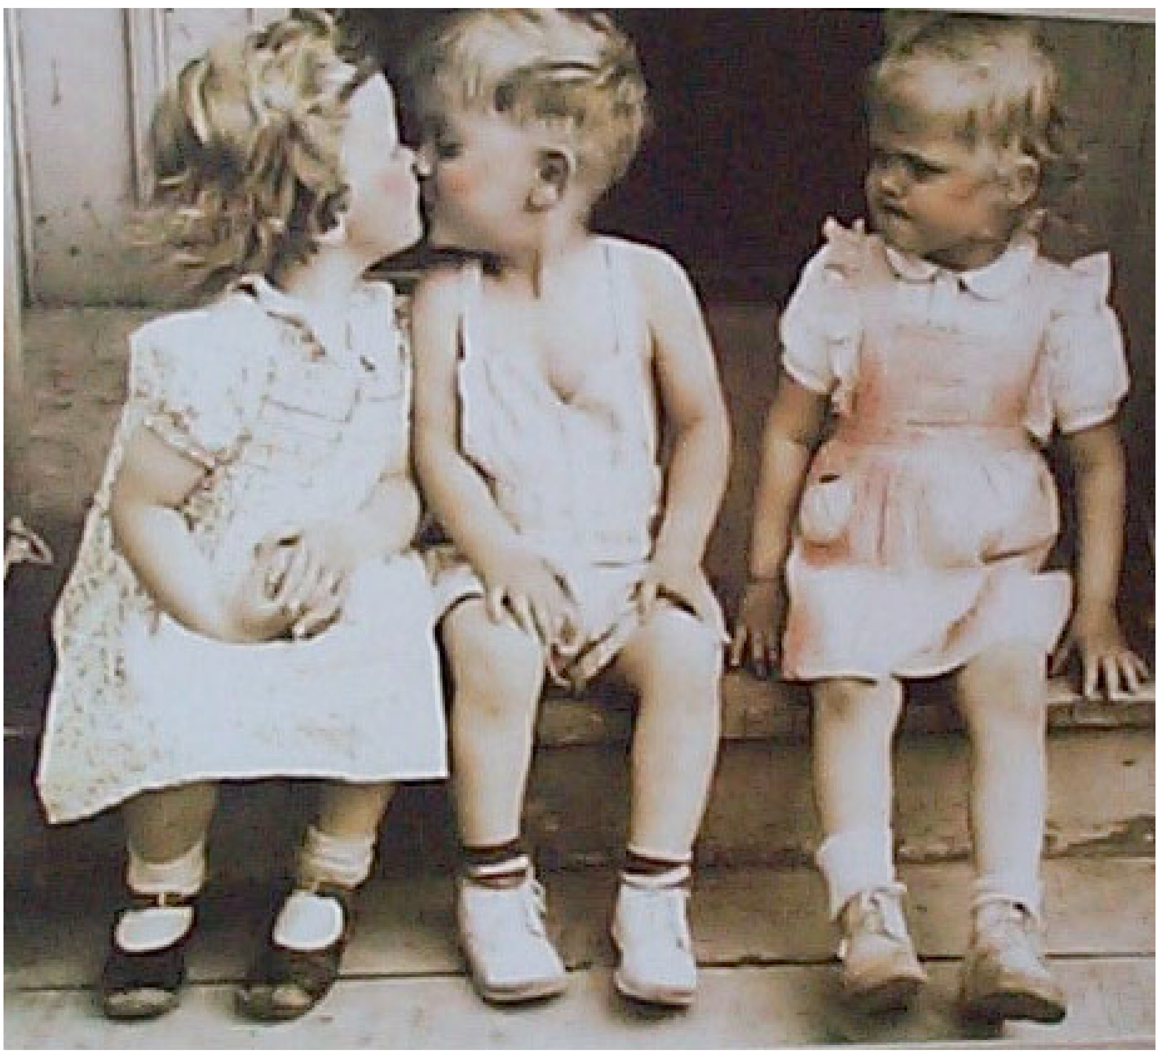

3a

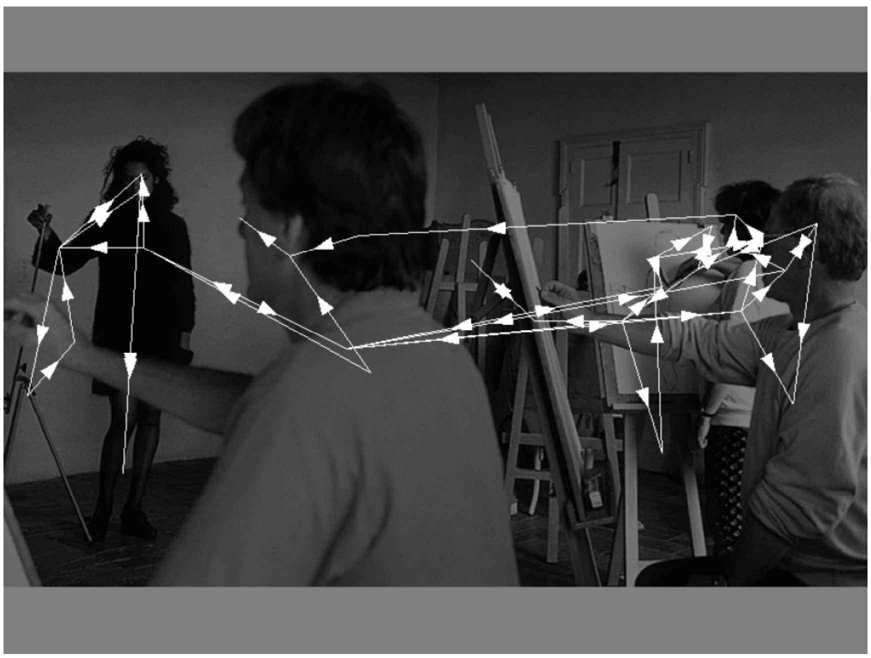

3b

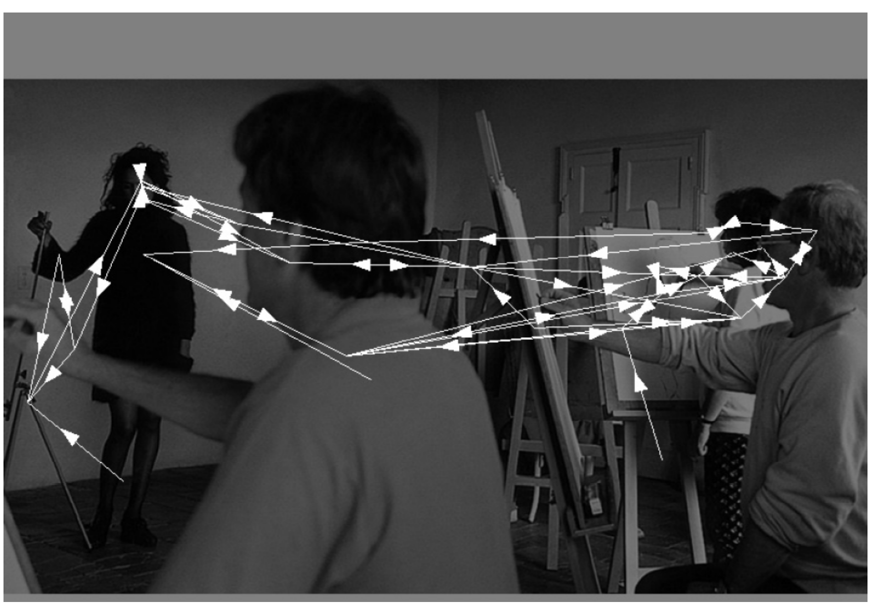

3c