

Visual Search Gets Real

August 02, 2010

Moving Visual Research Out of the Lab and Into Real-World Settings

It might not seem like it when you are looking for those missing car keys (for the third time this week), but humans are adept at visual searching. Our attention processes zero in on likely candidates while collecting only the gist of everything else that might be around; and we've also learned to give up when the search really is hopeless. APS 22nd Annual Convention Keynote Speaker Jeremy Wolfe of Harvard Medical School and Brigham and Women's Hospital has devoted his research to understanding these visual search processes. APS Fellow Wolfe has been at the forefront of moving visual research out of the lab and into hospitals, airports, and other real-world settings.

While Wolfe focuses on very modern search problems, he began his address by highlighting how these issues have shown up throughout human history. By Wolfe's calculation, the first mention of visual search is from chapter three of Genesis. In possibly the most famous example of impulse control failure (see coverage of research on self-control elsewhere in this issue), Adam and Eve ate fruit from the forbidden tree, prompting God to look for them. But as Wolfe said, "the experiences of an omniscient searcher, possessed of unlimited parallel processing capabilities, may not be very relevant to human search." Wolfe claimed that visual search showed up again in Shakespeare's 1597 account of the Battle of Shrewsbury in Henry IV. To protect himself during the battle, Henry disguises several of his men in royal robes. His opponent, Douglas, swears that he will kill all the "kings," possibly, said Wolfe, the first known published example of a serial, self-terminating search.

The modern era of visual search research was ushered in by Anne Treisman, who laid the groundwork for the entire field, said Wolfe. Every moment we have our eyes open, we're barraged by visual information. What processes are involved when we look for one specific thing in all this noise? What gets past our processing bottleneck? It turns out there is a gradual continuum of reaction times from easy search tasks, like finding the red object among blue ones, to very hard search tasks, like examining satellite images for possible nuclear weapon facilities. The source of this change seems to be how much guidance you have when searching — that is, how much basic visual information you can use to guide your attention. Apparently there are only a few visual features that we pick up on when searching. We process shape, object orientation, and dimensions, but not intersection type (it's very hard to distinguish an X intersection from a T intersection when looking for an X among Ts), combinations of different colors, and probably not faces or letters. We also use our previous experience to direct our attention. For example, when looking for a stove in a photo of a kitchen we draw on our previous knowledge of where stoves generally are (this has been termed semantic guidance) and where our own stove is at home (episodic guidance) to direct our attention. We go right to objects built into the kitchen counter, not the ceiling.



APS President Linda Bartoshuk presents Keynote Speaker Jeremy Wolfe with a memento of his address.

Once something has made it past the processing bottleneck to our conscious attention, we have to ask ourselves two questions. First, is this my target? If it is, we're done with a successful search. If it's not, another question arises: Is it time to quit searching? If yes, we stop, possibly missing the target if it was, in fact, present. If the answer is no, we find another possible target and start again with the initial question.

Wolfe has been busy taking all of this knowledge from the lab and applying it through a series of ingenious studies involving real-life searchers, such as radiologists looking for cancer and airport luggage screeners looking for terrorists. Thankfully, finding cancer in a mammogram or a gun in someone's luggage are rare events. But this low prevalence affects search processes. Wolfe has found that the more often you don't see something (like never seeing a weapon after months on the job as a luggage screener), the less likely you are to see it, even when it is there. Not only are items more likely to be missed in these low-prevalence searches, but false alarm errors decrease as well (also a potential problem in searches like cancer screenings or security screenings where you would want screeners to lean toward false alarms rather than missing the object entirely). Wolfe asserts that prevalence affects both the "Is it a target?" and "Should I quit?" criteria. Searchers find fewer targets, and they quit earlier because they don't expect to find them.

That work was done in the lab. Is it possible that real experts are not prone to the effects of prevalence? Maybe working continually in a low prevalence condition has conditioned them to avoid those errors? This is where Wolfe's work outside the lab is crucial, but it comes with challenges. Getting the data is difficult and expensive. "You try to get board-certified radiologists to work for \$10 an hour and do enough cases to have some actual results in the low-prevalence condition," quipped Wolfe. To get around this problem, he and his colleagues have conducted studies where they slip their test searches into the general workflow of a group of radiologists looking for cancer. While this seems like it might be simple, inserting tests into today's digital health workflow is also a challenge. "The first problem," said Wolfe "is that one computer wants to bill someone and a second computer wants to alert the state board of health, which is a felony. And don't even ask about IRB issues." But, the researchers soldiered on, successfully conducting a study of radiologists at Massachusetts General Hospital in Boston, and another facility in Wales. They found that experts are not immune to the prevalence effects; they miss more and have lower false alarms just like everyone else.



Perhaps because searching processes are mainly unconscious, they are very difficult to change even when the searcher is made aware of the potential for these errors. When research subjects are told about the effects of looking for low-prevalence objects, they will, as Wolfe put it, “say ‘yes, yes, yes’ and then miss it.” Even Wolfe, who has designed the studies, experiences the same effects when he serves as an observer.

In working with expert searchers, Wolfe noted that “they believe in something like ‘The Force.’” After finding something possibly cancerous in a mammogram, a radiologist might report that she knew that there was something wrong on first glance at the image, even before finding the troublesome spot. How can this be? Are there really psychic radiologists? Maybe this is just a selection bias. Maybe they only remember the times when they had the feeling that there was something wrong and then found that something while forgetting all the other times when they had a bad feeling but found nothing or found something without having that mysterious feeling. Wolfe and his colleagues tested this concept by showing a group of screeners mammogram images for a split second and asking them to make a decision about whether the person should go back for further testing (The radiologists originally balked at this, saying “I would never make a decision that quickly, it’s not responsible.” Wolfe’s response: “I know it’s not responsible, this is science.”) It turns out that the radiologists performed above chance. They were not nearly as accurate as they would be with their normal scrutiny of the images, but they were able to glean some pertinent information from seeing the mammograms for a fraction of a second. With further research on just what it is the radiologists are picking up on, this knowledge could someday lead to a new type of computer-aided detection system.

With all the research being conducted by Wolfe and his colleagues, much has been learned about visual searching. But as Wolfe said, “We have just begun to get real.” Wolfe predicts that researchers will need to broaden their study of search to include different searches, like foraging tasks, and complex search tasks, like looking through satellite imagery, which can take days. Meanwhile, Wolfe’s work will continue to bring us closer to understanding the fundamental psychological processes involved in visual searching and to addressing the very real-world challenges in which those processes play a central role.