Chapter 10

The Dance of Meaning

Meaning is what the brain performs in a dance with the external environment. In this dance tokens of meaning are spun off into electronic and social media and tokens of meaning are likewise picked up. New meaning is constructed when patterns already stored within the brain are combined with patterns constructed from external information. Increasingly, new meaning is also constructed by inanimate computers that do at least partial analysis and synthesis of patterns and tokens of meaning, and then present the results using a visual display.

In this book, the dance partners are considered to be individual people interacting with visual displays. In reality the dance is far more intricate and ELABORATE; there is a constant stream of new meaning being developed by people interacting with one another. Visual thinking is but a small part of the dance. Nevertheless, because of the special power of the visual system as a pattern-finding engine, visual thinking has an increasingly important role. This book itself is part of the dance, as is everything that is designed to be accessed visually.

Visual Thinking for Information Design. DOI: https://doi.org/10.1016/B978-0-12-823567-6.00010-0

One of the main themes we have explored in this book is that at every level of description visual thinking can be thought of as active processes operating through the neural machinery of the brain, which, through interconnections and neuron firing patterns, embody executable models controlling our interactions with the world. The purpose of this chapter is to review and summarize what we have covered so far and then discuss some of the broader implications of how the theory of perception applies to design.

REVIEW

The following few pages give a twelve-point summary of the basic machinery and the major processes involved in visual thinking. With the twelfth point we shall segue into observations that go beyond what has been said before.

1. The eye has a small high-resolution area of photoreceptors called the fovea. We see far more detail in the fovea than off to the side and we sample the world by making rapid eye movements from point to point. Eye movements rotate the eyeball so that imagery from different parts of the visual world falls on the fovea to be analyzed by the brain.
2. Our brains construct visual queries to pick up what is important to support what we are doing cognitively at a given instant. Queries trigger rapid eye movements to enable us to pick up information that answers the query.

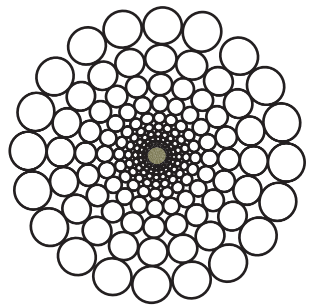

How affluent?

What are their ages?

"Unexpected Returns" by Ilya Repin, taken from A. Yarbus, 1967. Eye movements during perception of complex objects. L.A. Riggs, ed., *Eye Movements and Vision*, Plenum Press. NY. Chapter VII, 171–196.

On the previous page is a picture by the Russian realist painter Ilya Repin titled, "They did not expect him." Further to the right, shown in red, are eye movement traces from one person asked to perform different analytic tasks. When asked about the material well-being of the family, the eye movements fixate on clothes, pictures on the walls, and other furnishings (top). When asked about the ages of the family members, the eye movements fixate on faces almost exclusively (bottom).

3. The first stage of cortical visual processing is a local feature analysis done simultaneously for every part of the visual field. The orientation, size, color, and motion of each part of the image falling on the retina are determined all at once by feature selective neurons. Smaller-scale features are only analyzed in the fovea at the center of vision. All other visual processing is based on the initial division of the visual world into features.

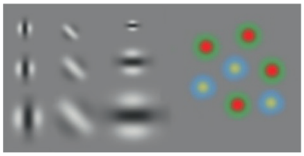

These records were made by the Russian psychologist Alfred Yarbus, who in 1967 used a mirror attached by means of a suction cup to his subjects' eyeballs to reflect light onto photographic paper and thereby record eye movements.

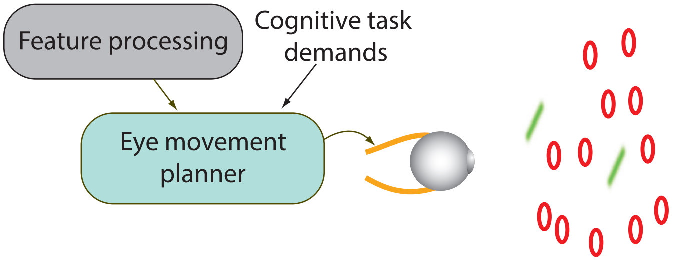

Understanding feature-level processing tells us what will stand out and be easy to find in a visual image. In the image above, the green bars are distinct in terms of several feature dimensions: color, orientation, curvature, and sharpness. We can easily execute eye movements to find the green bars because orientation, color, and curvature are all feature properties that are processed at an early stage, and early-stage properties are the ones that can be used by the brain in directing eye movements.

A simple motion pattern can also be thought of as a feature. Something that moves in a field of mostly static things is especially distinctive and easy to find.

4. There are two major processing pathways called the *where* and *what* pathways.

 The *where* pathway has connections to various regions in the parietal lobe responsible for visually guided actions, such as eye movements, walking, and reaching to grasp objects.

 The *what* pathway is responsible for identifying objects through a series of stages in which increasingly complex patterns are processed, each stage building on the previous one.

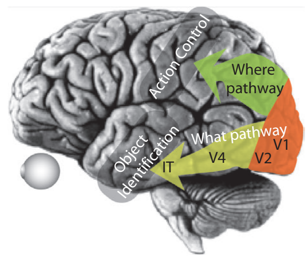

Between the *low-level* feature analysis and *high-level* object recognition is an intermediate pattern-finding stage. This divides visual space into regions bounded by a contour and containing similar textures, colors, or moving features. In V4 more complex compound shapes are identified from patterns of features. In the inferotemporal cortex, neurons respond to specific meaningful patterns such as faces, hands, letters of the alphabet, and automobiles.

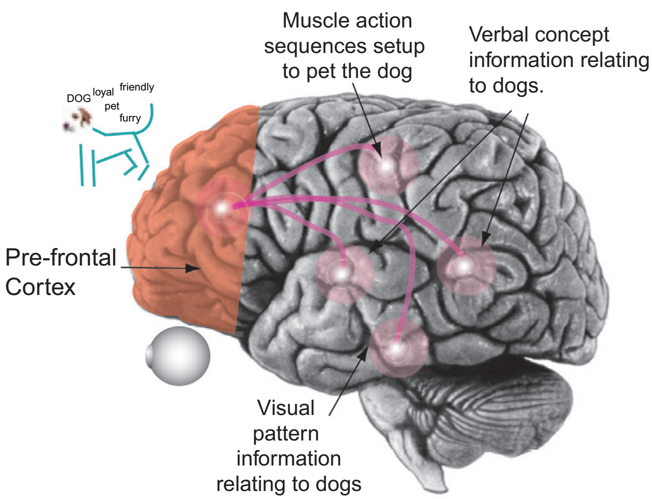

5. As we think visually, various kinds of information are combined in temporary nexuses of meaning. This information can include visual pattern information, language-based concepts, and action patterns. Nexuses are short lived and they make up the contents of the working memories.

From one to three meaningful nexuses can be formed every time we fixate on a part of a scene. Some of these may be held from one eye fixation to the next, depending on their relevance to the thought process. This three-object visual working memory limit is a basic bottleneck in visual thinking.

6. One way that visual displays support cognition is by providing aids to memory. Small images, symbols, and patterns can provide *proxies* for concepts. When these proxies are fixated, the corresponding concepts become activated in the brain. This kind of visually triggered activation can often be much faster than retrieval of the same concept from internal long-term memory without such aids. When an external concept proxy is available, access to it is made by means of eye movements which typically take approximately one-tenth of a second. Once the proxy is fixated, a corresponding concept is activated within less than two-tenths of a second. It is possible to place upwards of thirty concept proxies in the form of images, symbols, or patterns on a screen providing a very quickly accessible concept buffer. Compare this to the fact that we can hold only approximately three concept chunks in visual or verbal working memory at a time. There is a major limitation to this use of external proxies—it only works when there are learned associations between the visual symbols, images, or patterns and particular concepts.

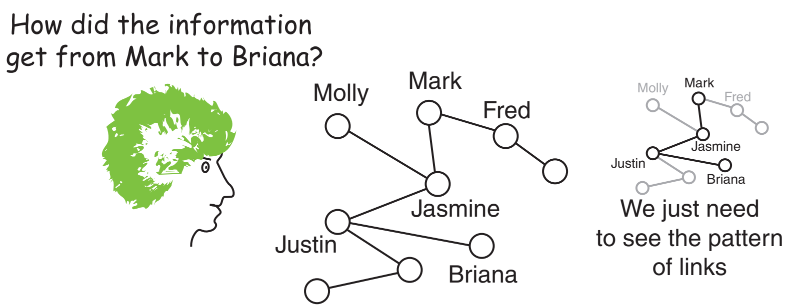

7. Another way in which visual displays support cognition is through pattern finding. Visual queries lead to pattern searches and seeing a pattern, such as a connection between two objects, often providing the solution to a cognitive problem.
8. Language processing is done through specialized centers in the left temporal lobe. Language understanding and production systems are specialized for a kind of informal logic exemplified by the “ifs,” “ands,” and “buts” of everyday speech. This is very different and complementary to the visual logic of pattern and spatial arrangements.

9. The term "attention" is used to refer to the focus of brain activity as a particular moment. This focus is multifaceted and can be externally or internally controlled. At the early stage of visual processing, attention biases the patterns which will be constructed from raw imagery. Different aspects of the visual images are enhanced or suppressed according to the top-down demands of the cognitive thread. An art critic's visual cortex at one moment may be tuned to the fine texture of brush strokes, at another to the large-scale composition of a painting, and at yet another to the color contrasts. Attention is also the very essence of eye-movement control because looking is a prerequisite for attending. When we wish to attend to something, we point our eyeballs at it so that the next little bit of the world we are most interested in falls on the fovea. The planning of eye movements is therefore the planning of the focus of attention, and the sequence of eye fixations is intimately tied to the thread of visual thinking.

Visual working memory is a process that is pure attention. It is a momentary binding together of a few nexuses of visual features and patterns that seem most relevant to the cognitive thread. These constructions are short lived, most lasting only a tenth of a second, or less than the duration of a single fixation. A few persist through a series of successive fixations, being reconnected to the appropriate parts of the visual image after each one. Some visual working memory constructions are purely imaginary and not based on external imagery. When an artist is contemplating a stroke of a pen, or an engineer is contemplating the alteration of a design, mentally imaged potential marks on the paper can be combined with information from the existing sketch. The combination of imaginary with the real is what makes visual thinking such a marvel and is a key part of the internal-external dance of cognition.

There is a basic cycle of attention. Between one and three times a second our brains query the visual environment, activate an eye movement to pick up more information, process it, and re-query. The information picked up becomes the content of visual working memory. This cycle provides the basic low-level temporal structure of the cognitive thread. At a higher level the cognitive thread can have a wide variety of forms. Sometimes the cognitive thread is largely governed by external information coming through the senses. In the case of a movie, the cinematographer, directors, and actors can, to a large extent, control the cognitive thread of the audience. This is not to deny each person in an audience their own thoughts and opinions about what they are seeing, but they will mostly look at the same images, and even the same parts of those images, in the same order. The objects constructed in a hundred people's visual and verbal working memories will be similar and occur in the same order.

Most of the time our visual cognitive thread is occupied with mundane things, like finding the path in front of us, picking up an object, or watching someone's face as they talk so as to integrate information from their facial expression with what they are saying. These skills are highly learned and so most of the time we are not aware of doing them; nevertheless, they do take up our attentional capacity.

Sometimes the cognitive thread is up for grabs. Someone walking, bored, along a flat unobstructed sidewalk only requires a small amount of visual attention to keep to the path and the rest of his or her attentional capacity is a kind of free resource. Such a person is likely to give a second glance to anything that is even slightly novel, especially if it is moving.

There is also the visual thinking process that occurs when someone reasons with a graphic design as an external aid. A map reader is carrying out a very goal-directed query loop that involves both the rough logic of the language-processing system together with the pattern-finding capability of the visual system.

To add to the diversity of attention, the cognitive thread shifts back and forth between visual-processing and language-processing modalities. Sometimes it is primarily visual, or visual and motor, as when we are doing something like tracing out a line with a pen. Sometimes it is driven by the language modules of the brain. We cannot do more than one visual task at a time, and we cannot do more than one verbal task at a time. But we can carry out a verbal task, such as talking on a cell phone, and another visual motor task, like folding laundry, at the same time, if one of them is a highly learned skill. The separation is not perfect and people driving while talking on phones are undoubtedly more dangerous, but these are examples of real multitasking using visual and verbal channels to semi-independently carry out two cognitive tasks at once. Nevertheless, it is when visual and language modalities are combined that the brain is most effective. A well-designed presentation, for example, will use words and graphics, each to convey different kinds of information and the two kinds of information will be linked using pointing, or simple proximity in space and time.

The natural way of linking spoken words and images is through deixis (pointing). People point at objects just prior to, or during, related verbal statements, enabling the audience to connect the visual and verbal information into a visual working memory nexus.

10. One basic skill of designers can be thought of as a form of constructive seeing. Designers can mentally add simple patterns to a sketch to test possible design changes before making any changes to the sketch.
11. Long-term episodic memories are executable predictive mental models, rather than fixed repositories like books or CD ROMs. Their primary purpose is action, not reminiscing. The pathways that are activated when a cognitive task is carried out become stronger if that task is successfully completed. These pathways exist in all parts of the brain and on many levels; they are responsible for feature detection, pattern detection, eye movement control, and the sequencing of the cognitive thread. Activated long-term memories are partially reconstructions of prior sequences of neural activity in particular pathways. Certain external or internal information can trigger these sequences. In the case of pattern recognition, the sequences are triggered by visual information sweeping up the "what" pathway.

Some visual skills, such as seeing closed shapes bound by contours as "objects" or understanding the emotional expressions of fellow humans, are basic in the sense that they are to some extent innate and common to all humans, although such skills are refined with practice and experience.

Understanding this scene requires the ability to perceive social interactions. (Photo: Josh Eckels)

Most importantly, long-term memories enable us to plan and act through their execution. We cannot do anything without having some ideas of the consequences of our actions.

Visual thinking is based on a hierarchy of skills. Sophisticated cognitive skills build on simpler ones. We cannot begin to play chess until we can identify the pieces, and we will not become an expert until we have learned patterns involving whole configurations relating to strategic advantage or danger. As we get skilled at a particular task, like chopping onions, the operation eventually becomes semiautomatic. This frees up our higher-level control processes to deal with higher-level problems, such as how to deal with an extra person coming to dinner. The process whereby cognitive activities become automated is absolutely critical in the development of expertise because of the fundamental limitations of visual and verbal working memory capacities. If a set of muscle movements involved in drawing a circle on paper becomes automated, then the designer has free capacity to deal with arrangements of circles.

Patterns of neural activation are not static configurations, but sequences of firing. At the highest levels, involving the prefrontal cortex at the front of the brain and the hippocampus in the middle, these sequences can represent action plans. This, too, is hierarchical. Complex tasks, like cooking a meal, involve high-level plans that have the end goal of getting food on the table, together with mid-level plans, like peeling and mashing potatoes, as well as with low-level plans that are semiautomatic, such as reaching for and grasping a potato.

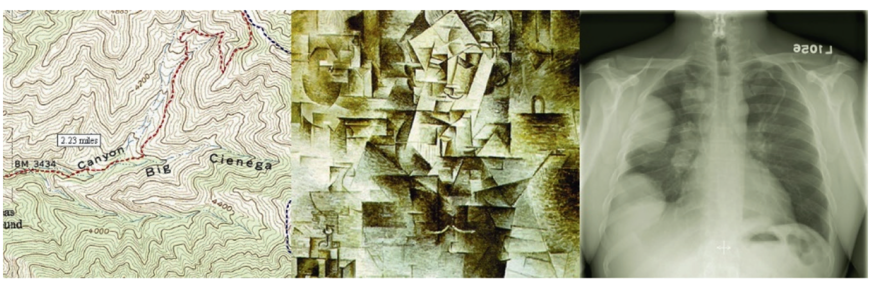

Other perceptual tasks, such as reading a contour map, understanding a cubist painting, and interpreting an X-ray image, require specialized pattern recognition skills. These higher-level skills are much easier to acquire if they build on more basic skills, which means that an artist cannot be too radical and still expect to be widely understood.

12. Effective presentations based on visualization consist of the transmission of predictive mental models from the presenter to the audience. To be successful they must build on existing mental models.

> There are many kinds of visualization that can help with this, from simply presenting data, to using cartoons and animation showing how things work, to photographs that provide rich information about the complexity of some part of the world.

The sum of the cognitive processing that occurs in problem solving is moving inexorably from being mostly in the head, as it was millennia ago before writing and paper, to being a collaborative process that occurs partly in the heads of individuals and partly in computer-based cognitive tools. Computer-based cognitive tools are developing with great speed in human society, far faster than the human brain can evolve. Any routine cognitive task that can be precisely described can be programmed and executed on a computer, or on millions of computers. This is like the automation of a skill that occurs in the brain of an individual, except that the computer is much faster and less flexible.

A cognitive tool can be a map or a movie poster, but increasingly cognitive tools are interactive and computer based. This means that every visual object shown on the screen can be informative in its own right, as well as be a link, through a touch or mouse click, to more information. We may also be able to manipulate that information object with our computer mouse, to literally organize our ideas.

IMPLICATIONS

The active vision model has four broad implications for design.

1. Designs should support the pattern-finding capability of the brain. Information structures should be transformed into easily identified patterns.
2. Designs optimize the cognitive process as a nested set of activities.
3. Designs should take the economics of cognition into account, considering the cost of learning new tools and ways of seeing.
4. Designs should take the mental models of the consumers into account.

The following sections elaborate these principles with examples.

DESIGN TO SUPPORT PATTERN FINDING

Properly exploiting the brain's ability to rapidly and flexibly discover visual patterns can provide a huge payoff in design. The following example, from my own work, illustrates this. Over the past few years I have been fortunate to be involved in a project to discover the underwater behavior of humpback whales. The data was captured by a tag attached to a whale with suction cups. When the tag came off it floated to the surface and was retrieved. Each tag provided several hours of data on the position and orientation of the whale as it foraged for food at various depths in the ocean. This gave us an unprecedented opportunity to see how humpback whales behave when they are out of sight underwater.

Our first attempt to provide an analysis tool was to create a program that allowed ethologists to *replay* a moving three-dimensional model of the whale at any desired rate.◆ They were initially thrilled because for the first time they could see the whale's underwater behavior. Some fascinating and previously unknown behavioral patterns were identified by looking at these replays. Nevertheless, although the visualization tool did its job, analysis was extremely time-consuming. It took many hours of watching to interpret an hour's worth of data.

◆This work is described in C. Ware, R. Arsenault, M. Plumlee, and D. Wiley, 2006. Visualizing the underwater behavior of humpback whales. *IEEE Computer Graphics & Applications*. July/August issue. 14–18.

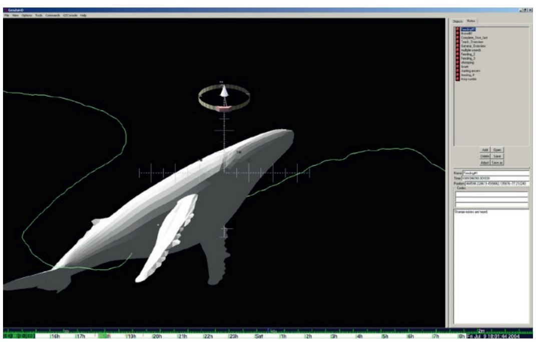

GeoZu4D software allows for the underwater behavior of a humpback whale to be played back along with any sounds recorded from a tag attached to the whale by means of suction cups.

Using this replay tool, the visual thinking process can be roughly summarized as follows. The whale movements were replayed by an ethologist looking for *stereotyped temporal sequences of movements*. When something promising was observed, the analyst had to remember it in order to identify similar behaviors occurring elsewhere in a record. The ethologist continued to review the replay looking for novel characteristic movement patterns.

Although this method worked, identifying a pattern might take three or four times as long as it took to gather the data in the first place. An 8-hour record from the tag could take days of observation to interpret. If we consider the problem in cognitive terms the reason becomes clear. We can remember at most only a half dozen temporal patterns in an hour of video, and these may not be the important or stereotyped ones. Further, we are not nearly as good at identifying and remembering motion patterns as we are at remembering spatial patterns. Also, every time an ethologist formed the hypothesis that some behavior might be stereotypical, it was necessary to review the tracks of all the other whales again looking for instances of the behavior that might have been missed.

Our second attempt at an analysis tool was far more successful because it took much greater advantage of the brain's pattern-finding ability. We transformed the track of the whale into a 3D ribbon. We added a sawtooth pattern above and below the ribbon that was derived from calculated accelerations indicating the fluke strokes of a whale.

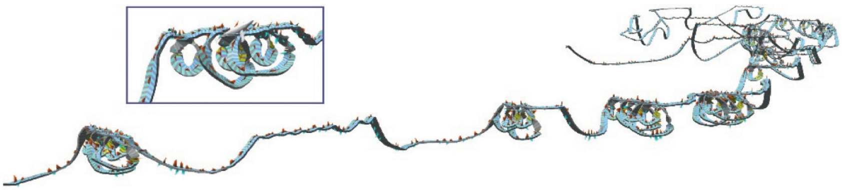

Transforming the whale track into a ribbon allowed for much more rapid identification of behavior patterns. This new tool enabled a much different visual thinking process. The ethologist could quickly zoom in on a region where feeding behaviors were seen and stop to view a static image. Behavior patterns could be identified by visually scanning the image looking for repeated graphic patterns. Because eye movements are so fast and static patterns can be picked out efficiently, this method enabled analysts to compare several patterns *per second*. This process was hundreds of times faster than the replay method. Naturally, a complete analysis still took a great deal of work, but now we could scan the data for familiar or new patterns within minutes of extracting the data files from the tag. The new visual thinking process was hundreds of times faster than the old one based on simple replay.

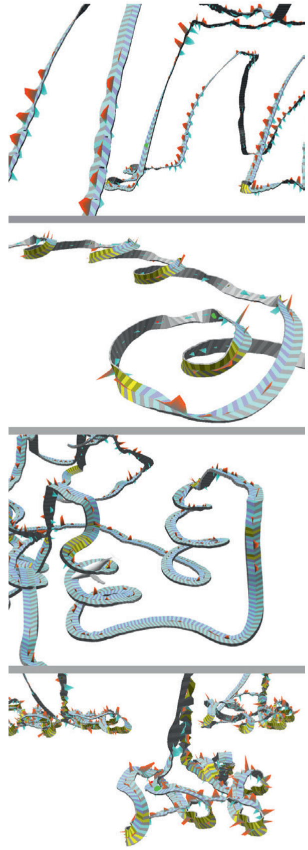

Ribbons make the underwater behavior of a humpback whale accessible through pattern perception.

The whale behavior study clearly demonstrates the general principle that patterns can be identified and compared very rapidly if we can turn information into the right kind of spatial display. The design challenge is to transform data into a form where the important patterns are easy to interpret.

OPTIMIZING THE COGNITIVE PROCESS

Good design optimizes the visual thinking process. The choice of patterns and symbols is important so that visual queries can be efficiently processed by the intended viewer. This means choosing words and patterns each to their best advantage. When designing the visual interface to an interactive computer program we must also decide how the visual information should change in response to every mouse click.

Extraordinarily powerful thinking tools can be made when a visual interface is added to a computer program. These are often highly specialized—for example, for stock market trading or engineering design. Their power comes from the fact that computer programs are cognitive processes that have been standardized and translated into machine-executable form. They offload cognitive tasks to machines, just as mechanical devices, like road construction equipment, offload muscle work to machines. Once offloaded, standardized cognitive tasks can be done blindingly fast and with little or no attention. Computer programs do not usually directly substitute for visual thinking, although they are getting better at that too; instead they take over tasks carried out by humans using the language-processing parts of the brain, such as sophisticated numerical calculations. A visual interface is sometimes the most effective way for a user to get the high volume of computer-digested information that results. For example, systems used by business executives condense large amounts of information about sales, manufacturing, and transportation into a graphical form that can be quickly interpreted for planning and day-to-day decision-making.

It is useful to think of the human and the computer together as a single cognitive entity, with the computer functioning as a kind of *cognitive co-processor* to the human brain.◆ Low-bandwidth information is transmitted from the human to the computer via the mouse and keyboard, while high-bandwidth information is transmitted back from the computer to the human for flexible pattern discovery via the graphic interface. Each part of the system is doing what it does best. The computer can preprocess vast amounts of information. The human can do rapid pattern analysis and flexible decision-making.

◆The term "cognitive co-processor" comes from a paper from Stuart Card's famous user interface research lab at Xerox Palo Alto Research Center (PARC) although I am using it in a somewhat different sense here. G. Robertson, S. Card, and J. Mackinlay, 1989. The cognitive co-processor for interactive user interfaces. *Proceedings of the ACM UIST Conference*. 10–18.

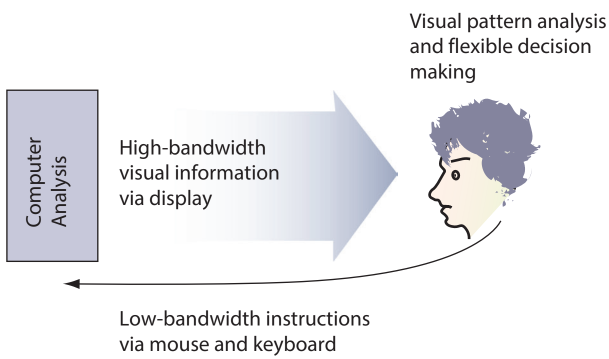

Peter Pirolli and Stuart Card developed a theory of information access to help with the design of interfaces to cognitive tools. They began with the *foraging theory* developed by ethologists to account for animal behavior in the wild. Most wild animals spend the bulk of their time in a highly optimized search for food. To survive they must balance the energy expended in finding and consuming food—just eating and digesting has a high cost for grazing animals—with the energy and nutrients obtained. They found that people forage for information on the Internet much as animals forage for food; they are constantly making decisions about what *information scent* (another term from Stuart Card's influential laboratory) to follow, and they try to minimize how much work they must do, offloading tasks onto the computer wherever possible.◆

◆P. Pirolli and S. Card, 1999. Information foraging. *Psychological Review*. 106: 643–675.

The ideal cognitive loop involving a computer is to have it give you exactly the information you need when you need it. This means having only the most relevant information on screen at a given instant. It also means minimizing the cost of getting more information that is related to something already discovered. This is sometimes called *drilling down*.

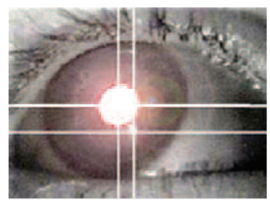

It might be thought that an eye tracker would provide the ideal method for drilling down, since eye movements are the natural way of getting objects into visual working memory. Eye tracking technology can determine the point of gaze within about one centimeter for an object at arm's length. If the computer could have information about what we are looking at on a display, it might summon up related information without being explicitly asked.

There are two problems with this idea. First, tracking eye movements cheaply and reliably has proven to be technically very difficult. Eye trackers require careful and repeated calibration for each user. Second, when we make eye movements we do not fixate exactly; we usually pick up information in an area of about one centimeter around the fovea at normal computer screen viewing distances and this area can contain several informative objects. This means the computer can only "know" that we might be interested in any of several things. Showing information related to all of them would be more of a hindrance than a help.

The quickest and most practical method for drilling down is the mouse-over *hover query*. Imagine that by moving the mouse over a part of a diagram all the other on-screen information *relating to the thing the mouse is moving over* becomes highlighted and the relevant text enhanced so that it can be easily read. This is exactly what was done by Tamara Munzner and a team at Stanford University in their experimental *Constellation* system.◆ Their application was a kind of network diagram

◆T. Munzner, F. Guimbretiere, and G. Robertson, 1999. Constellation: A visual tool for linguistic queries from MindNet. Proceedings of IEEE *InfoVis. Conference*, San Francisco. 132–135.

showing the relationships between words. The diagram was far too complex to be shown on a screen in any normal way, but by making it interactive hundreds of data objects could be made usable. Compare this to a typical network diagram, in which only between ten and thirty nodes are represented, and the advantage becomes clear. This kind of mouse-over clickless hover query is the next best thing to moving the eyes around an information space. Clickless hover queries can be made only about once a second, slow compared to three-per-second mouse movements, but this still leads to a very rapid interaction where the computer display seems like part of the thinking process, rather than something to be consulted.

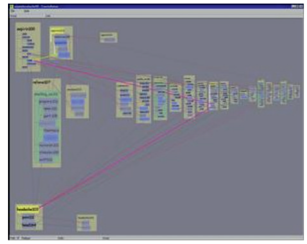

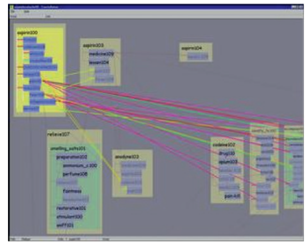

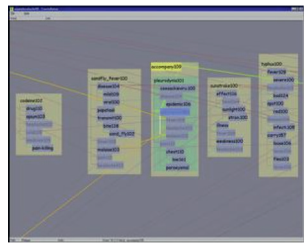

Constellation was a single example of an experimental user interface, but it provides a more general lesson. The ideal computer system that supports visual thinking should be extremely responsive, presenting relevant information just at the moment it is needed. This is not easy to achieve, but something for which to strive.

It should be noted that the most common application of hover queries is in so-called "tooltips." These are the presentation of additional information about a menu item or icon when the mouse is placed over it. Generally these are implemented so that a delay of a second or two occurs between the query and the information display. This is probably appropriate for tooltips because we do not want such information to be constantly popping up, but it would not be appropriate for more tightly coupled human-computer systems. Such a long delay would do serious damage to the efficiency of the cognitive loop that is supported by Constellation.

LEARNING AND THE ECONOMICS OF COGNITION

Many of the visual problems we solve in life seem completely mundane: walking across a room, preparing a salad, looking for a road sign. They are mundane only because we have done them so often that we no longer consider them as requiring thinking at all. When we are born we have little visual skill except the basic minimum needed to identify that something is an "object" and a special propensity to fixate on human faces. Newborns do, however, have the neural architecture which allows the capabilities they can develop.◆ Everything else is a learned skill and these skills make up who we are. Our visual skills vary from the universal, like the ability to reach for and grasp an object, to the very specialized, like determining the sex of chicks, a very difficult visual pattern matching task that can only be done by highly trained experts.

All cognitive activity starts off being difficult and demanding attention. As we develop skills, the neural pathways involved in performing the

◆One of the most important fixed architectural cognitive capacities is the three-item limit in visual working memory. Even this is somewhat mutable. A recent study showed that video game players can enhance their visual working memory capacities from three to four items. But we do not find people with a working capacity of ten items.

See C.S. Green and D. Bavelier, 2003. Action video game modifies visual selective attention. *Nature*. 423: 534–537.

task are strengthened. These neural pathways carry particular patterns of activation, and strengthening them increases the efficiency of sequences of neural firing (the cognitive activities). The cognitive process becomes more and more automatic and demands less and less high-level attention. The beginning typist must use all his cognitive resources just to find the letters on the keyboard; the expert pays no attention to the keys. The use of drawing tools by an engineer or graphic designer is the same.

Sometimes we have a choice between doing something in an old, familiar way, or in a new way that may be better in the long run. An example is using a new computer-based graphic design package, which is very complex, having hundreds of hard-to-find options. It may take months to become proficient. In such cases we do a kind of cognitive cost-benefit analysis, weighing the considerable cost in time and effort and lost productivity against the benefit of future gains in productivity or quality of work. The professional will always go for tools that give the best results even though they may be the hardest to learn because there is a long-term payoff. For the casual user sophisticated tools are often not worth the effort. This kind of decision-making can be thought of as cognitive economics. Its goal is the optimization of cognitive output.

A designer is often faced with a dilemma that can be considered in terms of cognitive economics. How radical should one make the design? Making radically new designs is more interesting for the designer and leads to kudos from other designers. But radical designs, being novel, take more effort on the part of the consumer. The user must learn the new design conventions and how they can be used. It is usually not worth trying to redesign something that is deeply entrenched, such as the set of international road signs because the cognitive costs, distributed over millions of people, are high. In other areas, innovation can have a huge payoff.

The idea of an economics of cognition can be fruitfully applied at many cognitive scales. In addition to helping us to understand how people make decisions about tool use, it can be used to explain the moment-to-moment prioritization of cognitive operations, and it can even be applied at the level of individual neurons. Sophie Deneve of the *Institute des Sciences Cognitives* in Bron, France, has developed the theory that individual neurons can be considered as *Bayesian* operators, "accumulating evidence about the external world or the body, and communicating to other neurons their certainties about these events." It her persuasive view each neuron is a little machine for turning prior experience into future action.◆

There are limits, however, to how far we can take the analogy between economic productivity and cognitive productivity. Economics has money as

◆In 1763 the Reverend Thomas Bayes came up with a statistical method for optimally combining prior evidence with new evidence in predicting events.

For an application of Bayes' theorem to describe neural activity, see S. Deneve, 2005. Bayesian inference in spiking neurons. Published in *Advances in Neural Information Processing Systems*. Vol. 17. MIT Press 1609–1616.

a unifying measure of value. Cognitive processes can be valuable in many different ways and there is potentially no limit on the value of an idea.

DESIGNING FOR MENTAL MODELS

Our brains embody a set of flexibly interconnected mental models of the world and how we operate within it. They are predictive and allow us to take actions knowing the likely consequences. Many of the problems with computer interfaces occur when the user has a mental model of system's operation that does not accurately predict how it behaves. For example, consider the case where a computer appears to be hung up on some task. We do not know if data is taking a long time to arrive, or if the compute program itself has failed. Should we wait or terminate the program? This is a relatively simple problem that can be usually fixed with feedback indicators showing how much data has been loaded. The simple spinning indicators, showing activity, are usually less useful because they can continue to spin even though a program is irretrievably stuck.

Another example where mental models are critical and where visualization is often used is in storm forecasts. Usually weather sites show visualizations of predicted hurricane tracks, wind speeds, and storm surge. However, a study reported in *BAMS* of superstorm Sandy revealed dramatic failures of mental models.◆ People in the affected areas greatly overestimated the danger from high winds, but underestimated the risk from flooding. For one example where action could have been taken, something like 250,000 cars were damaged from flooding and most were subsequently written off, even though in most cases higher ground parking could have been found. In addition, only about 22% of people had an evacuation plan including a place to stay and the traffic conditions associated with late departure were not anticipated.

◆R.J. Meyer, J. Baker, K. Broad, J. Czajkowski, and B. Orlove, 2014. The dynamics of hurricane risk perception: Real-time evidence from the 2012 Atlantic hurricane season. *Bulletin of the American Meteorological Society*. 95(9): 1389–1404.

Visualizations can be extremely effective means of transmitting executable mental models, as we saw in Chapter 8. But just showing the data—or in this case, the forecasts—is often not enough. Other kinds of visualization such as simple cartoons can help develop the mental models needed for people to realize the true risks.

The problem of providing mental models of computer operations is becoming increasingly important with the new abilities of computer pattern recognition that comes with so-called deep learning. In many cases visualization can help; for example, when computer is used to diagnose cancerous cells, a human radiologist can be shown the most heavily weighed examples of cell that lead to the decision.

WHAT'S NEXT?

The science of human perception is continuing to advance and the account of visual thinking given in this book will, to some extent, be obsolete before it hits the bookstores. As research moves forward, our increasing knowledge of the human brain and perception will lead us in new and exciting directions. I believe, however, that the broad picture of how visual thinking works is essentially correct and will stand the test of time. There is a feeling among vision scientists that the secrets of the brain are at last becoming unlocked and that the outlines of lasting theories are beginning to emerge.

Ultimately the science of perception must take design into account because the designed world is changing people's thinking patterns. Real-world cognition increasingly involves computer-based cognitive tools that are designed to support one mode of thinking or another. This cognitive support environment is developing and evolving from year to year in a process that is happening much faster than evolution. Designed tools can change how people think.

The human brain evolved in a world made up of natural objects: plants, rocks, earth, sky, other humans, and animals. None of these were explicitly designed by early humans, although like other animals, humans have always shaped their environments. Most people now live in cities where almost everything has been designed. Our visual thinking skills are shaped by how we interact with this world of designed objects from the moment we first open our eyes. Objects in our modern environment incorporate computer programs, and some of them are explicitly designed as cognitive tools. At work or at play, many of us are in front of computer screens for a significant percentage of our lives. For many people the manipulation of a mouse or video game controller is as skilled and entrenched as any basic life skill. Many modern teenagers are more skilled at navigating three-dimensional virtual worlds with a game controller than they are at running over rough real-world terrain. Visual design must take into account both the relatively fixed capacities of the human brain as well as the evolving skill sets of people who use sophisticated and powerful cognitive tools. Interactive design is becoming ever more important as the loop coupling humans and computers tightens. The visual system will always be the highest bandwidth sense by far, and making full use of its flexible pattern-finding capabilities can provide great benefits.