Neural Correlates of 4-D visualizations. A comparison of Balloon Races and 3-D surface graphs

Eric Walden, Gabriel Cogo, David Lucus, Elshan Moradiabadi, Roozmehr Safi

Abstract

In this research we put forth an evolutionary argument to explain how people comprehend graphs. We test this theory by comparing brain activation during a graphic comprehension task using two different graph types: balloon race graphs and 3-D surface graphs. In accordance with our hypotheses we find that comprehension of 3-D surface graphs results in greater activation of the ventral stream and greater accuracy in graphical comprehension than balloon race graphs. We argue that this is because the human visual system is evolutionarily adapted to the comprehension of 3-D surfaces. The implication is that (1) even when representing the same data, the choice of graph matters and (2) choosing graphical representations that match what the brain is evolutionarily designed to process can enhance graphic comprehension.

Keywords: graphs, visualization, fMRI, NeuroIS, 3-D.

Neural Correlates of 4-D visualizations. A comparison of Balloon Races and 3-D surface graphs

INTRODUCTION

Big data offers new opportunities for organizations that can make sense of the massive amounts of data being generated across all business functions (Chen et al. 2012). One of the most important ways organizations can help people to make sense of big data is through visualizations (Baker et al. 2009). Estimates suggest that the human visual system can transfer information at roughly the speed of a 10 Mb Ethernet cable (Koch et al. 2006); making the human visual system an excellent candidate to transfer large amounts of data to a human processor. This has led to proliferation of both visualization research (Tractinsky and Meyer, 1999; Lim and Benbasat, 2000; Kumar and Benbasat, 2004) and new types of data visualization systems (i.e. dashboards). The development of visualization tools and techniques has outpaced research on the fundamental question of what makes a high quality data visualization. This is unfortunate, because the abundance of visualization choices makes it ever more important to understand both the cognitive and the perceptual characteristics of a high quality data visualization.

To better understand this interaction between human visual processing and data visualizations, this research makes use of functional magnetic resonance imagining (fMRI) to observe brain activation while subjects answer questions based on two different types of graphs. The research begins with the premise that the human visualization system has evolved to process visual displays that are encountered in nature and proved crucial for humans' survival. We demonstrate that people are both more accurate and more efficient in the sense that they use fewer neural resources when viewing data representations that take advantage of evolutionarily hard-wired parts of the visual system.

The next section goes into more detail on the literatures in both data visualization comprehension and human visual processing. This section is followed by an exploration of neural theory of data visualization comprehension based on the structure and function of the human visual system, followed by a description of the experiment and presentation of the findings. The paper concludes with a discussion of the implications and limitations of the work, and suggestions for future research.

LITERATURE REVIEW

Most research in data visualizations has been done in the geographical information system (GIS) literature. One of the main cognitive components of comprehending geographical visual display is graphicacy. Graphicacy is the ability to perceive, comprehend, and present information in a graphical form (Balchin 1976). It is analogous with literacy, numeracy, and articulacy, but represents a visual-spatial dimension of human intelligence rather than a narrative or mathematical dimension. Scholars in different fields have investigated whether graphs[1] could offer better comprehension of data than other types of representations (Lambon Ralph et al. 1999; Larkin et al. 1987; Mayer et al. 1994; Walther et al. 2001). Graphics have three major advantages in comparison to text (Aldrich et al. 2000). First, visual information is concise: a graph or picture gives the user a holistic view of the scene instantly. Second, graphics are memorable so that it easy to recall when needed. Third, graphics make relationships obvious; in fact, that is the raison d'etre for graphics.

Despite all the advantages inherent in graphs, scholars argue that the aphorism "A picture is worth a thousand words" may be an overgeneralization (Lambon Ralph et al. 1999; Mayer et

[1] Data visualization is a relatively new term, traditionally the term research has used is graph or graphic. We use these terms interchangeably in this work.

al. 1994; Walther et al. 2001). While pictures are a good way of conveying information, they are not a panacea—text, numbers, and the spoken word also convey information well. Moreover, people do not always interpret the meaning of graphics in the same way or with the same level of proficiency. Furthermore, previous studies have argued that background knowledge has an effect on graphicacy, rather than innate talent (Boardman 1990; Cox et al. 2004; Postigo et al. 2004). Thus, while graphics have advantages, they do not enable perfect communication.

Studies involving graphs comparing between the effectiveness of different types of visualization have been conducted for several decades now (Benbasat et al. 1977). The MIS literature on graphics largely derives from the Minnesota experiments (Dickson et al. 1977), which recognizes the effect of mode of presentation and decision-makers' characteristics on efficiency of decision-making. Furthermore, it identifies a set of dependent and independent variables for research in management information systems (MIS). Following the guidelines of the Minnesota experiments, several studies investigated the effect of different modes of presentations on decision-making quality (accuracy and speed of a decision) in different task settings (Benbasat et al. 1985; Benbasat et al. 1986; Remus 1987). The equivocal results of these studies called for a theory to encompass all these tasks and mode of presentation under the same tenant. Currently the leading theory to understand the interplay between task demand and mode of presentation in MIS is Cognitive Fit Theory (Vessey 1991; Vessey et al. 1991). Cognitive Fit theory posits that if the internal and external environment of the situation match, the problem-solving performance with respect to accuracy and speed will be maximized. Although accuracy and speed may not be the only goals of viewing graphics (Tractinsky et al. 1999), both are often important.

By acting on an unexplored gap in the graph construction literature (Pinker 1990), Kumar & Benbasat (2004) found that three-dimensional graphs performed better than 2D even for simpler tasks, leading to the conclusion that three-dimensional graphs perform better than two-dimensional graphs under all task conditions that involve more than two variables. It is actually a surprising finding that under *all* conditions three-dimensional graphs out-performed two-dimensional graphs. Even though there are several comparative studies between two and three-dimensional graphs (Chen and Yu, 2000; Kelton et al. 2010), most research does not consider four-dimensional graphs. This is regrettable because with more data being collected it becomes increasingly important to represent more dimensions on a single graph. Big data applications will increasingly require better presentations in order to be able to convey all the information in the best way possible.

While much of this research has examined how data is modeled in graphs and how they affect the performance of viewers (Correll et al. 2012), less is known about the cognitive processing of different types of graphs (Scaife et al. 1996). However, some research has begun to look more closely at the cognitive aspects of graphical comprehension (Kelton et al. 2010). Baker and his colleagues (2009) proposed a model based on sense making—an ability to comprehend, assimilate, and make inferences—that they considered to be the main goal of data visualization. Furthermore, Baker and his colleagues propose a model with four characteristic of visual representation which affect sense making. The four characteristics that best facilitate sense making of visualized data are: support for basic human visual perceptual approaches[2], strong gestalt properties, consistence with the viewer's prior knowledge, and support for analogical reasoning. Parsons et al. (2014) formulated 10 properties of interactive visual representations

[2] The four basic visual perpetual approaches are: Association, Differentiation, Ordered perception, and Quantitative perception (Bertin, 1983)

whose values could be adjusted through interaction. Kim and colleagues (Kim et al. 2010; Kim et al. 2011) showed the importance of colors in line graphs, while Chau (2011) showed how techniques used in visualization could enhance Web search results and evaluation of results. Zhu et al. (2010) showed that cognitive fit interacts with working memory capacity and information load. Thus, while research on graphic comprehension his turned increasingly toward understanding the role of the human information processing, there is still a great deal of work to do on the subject.

THEORETICAL DEVELOPMENT

The underlying theoretical basis for this work is evolution (Darwin 1859; Darwin et al. 1858). The basis of understanding cognitive neuroscience is the idea that the brain is built according to instructions contained in an organism's genome. This genome contains the blueprints for structures, such as hands, noses, and even visual cortexes. By virtue of their design, these structures allow an organism to solve problems, such as grasping objects, smelling, or processing visual information (Marr 1982). Within the brain, neurons and glia are arranged into predefined patterns and connections based on the genome of the species. Within *Homo sapiens*, the visual cortex is in the posterior part of the brain (i.e. the back of the head), while the motor cortex is in the superior part of the brain (i.e. top of the head). The visual cortex processes visual data and the motor cortex controls movement. Except in the case of mutation or mutilation, these areas are always in the same part of the brain and always perform the same tasks. Also, the structure of neurons in these specialized regions of the brain differ: the motor cortex always contains much larger neurons than the visual cortex, but the visual cortex contains many more, smaller, neurons.

There are variations in the brain from person to person, just as there is variation on the sizes of noses or hands. Furthermore, just as one can, through practice or inactivity, change the size and strength of muscles, the brain can also be modified by practice or inactivity. However, like the rest of the body, the basic blueprint for the structure of the brain within a species is determined by the genome of the organism, and it is very similar across members of the same species. The genome of a species is modified over time by selection pressures so that it builds structures to solve adaptive problems (Cosmides et al. 1994). An adaptive problem is one that members of a species face repeatedly over many generations, and that may be solved by a structural change resulting from a genetic change.

Hominid feet are a good example of a solution to an adaptive problem. Compared to other apes, hominids have long feet. This is an adaption to moving around on the ground rather than in trees. Long feet enable hominids to balance without falling over. Whereas, the hand-like feet of other apes are extremely awkward for upright walking. Over many generations, hominid species faced the problem of walking or running over open ground. There was an advantage in long feet over short ones. Members of the species with long feet were more likely to reproduce and hence the instructions for long feet were passed down to the next generation and over time the genome of the species changed.

Similar changes in the brain distinguish *Homo sapiens* from other primates. While the size of the occipital cortex (primarily visual areas) is roughly the same relative to body size for all primates including *Homo sapiens*, the temporal cortex (which houses many language regions) is much larger relative to body size for *Homo sapiens* (Rilling 2006). Thus, just as *Homo sapiens* hands are similar to ape hands, but *Homo sapiens* feet are different from primate feet, *Homo sapiens* occipital cortexes are similar to primate occipital cortexes but *Homo sapiens* temporal cortexes are different. Moreover, these differences are due to differences in the genome of *Homo sapiens* relative to other primates. Furthermore, these differences are solutions to adaptive problems of walking for feet and language for the temporal lobes. Both walking long distances and being able to communicate are problems faced repeatedly by *Homo sapiens*.

For the purposes of this paper, the adaptive problem of interest is gathering information from the photons around us. The sun showers the earth in a continuous stream of photons which are absorbed and reflected by, objects around us in ways that convey a great deal of information about the environment to organisms with the appropriate brain structures to process this information. Photons have been showering the earth for billions of years, so making sense of these photons is clearly a problem that has been faced by most species for many generations. The fact that eyes and vision exist suggests there is a genetic solution to building brain structures to make sense of the information contained in photons. We would expect to find specific brain structures that allow us to make sense of the visual information contained in photons. In fact, the human visual cortex is roughly 20% of the cortical area (Wandell et al. 2008). Not only are there areas dedicated to processing visual information, but also there are many such areas.

In general, visual processing in the human brain is divided into the dorsal stream and ventral stream (Goodale et al. 1992). The optic nerve runs from the eyes to the posterior part of the brain where the primary visual cortex resides. This area transforms visual signals from the eye into a brain map of what is being observed. This information is then sent forward and up into the dorsal stream which is the *where* stream that converts the map to locations in space and begins processing information about how to take action on the scene (Goodale et al. 1992; Ungerleider et al. 1994). Thus, the dorsal stream converts the two dimensional images that are registered on the cornea and represented in the primary visual cortex into three-dimensional images, so that an organism can navigate or act on the environment that is being perceived. Figure 1 shows the localization of both the Dorsal Stream and the Ventral Stream in the brain.

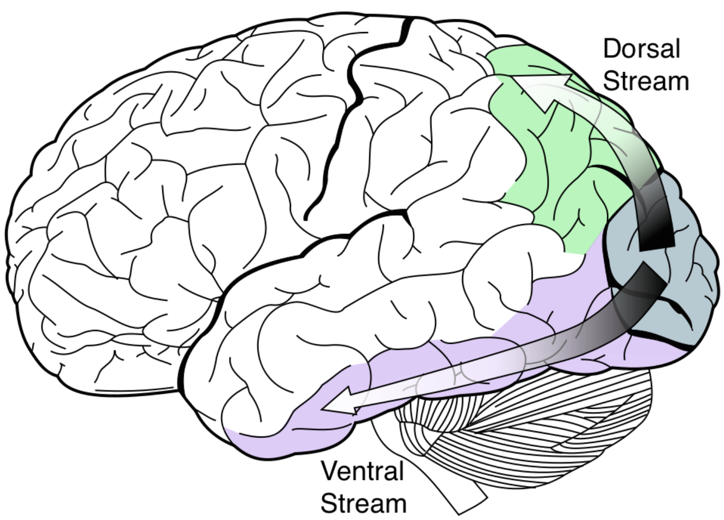

Figure 1: Dorsal and Ventral Streams (Image from Skelket on Wikipedia GNU free license)

While the dorsal stream is implicated in both the *where* and *how* dimensions of visual processing, “the characterization of the ventral stream as a 'What' pathway is relatively uncontroversial (Kravitz et al. 2011 pg. 217).” The ventral stream receives the brain map of the vision from the primary visual cortex and determines *what* is being seen. The ventral stream contains special regions that process faces, body parts, and even tools or houses (Dekker et al. 2011). Whereas the dorsal stream may tell us that there is an object five meters away moving toward us, the ventral stream will tell us whether it is another person or an automobile, and whether it is a person we know or an automobile that we own. Thus, the ventral stream primarily handles the meaning of the scene and the sense making about a scene.

When viewing a graph the goal is generally not to take action on the graph, but to understand the meaning of the graph. Understanding a graph should depend mainly on the *what* of the graph, rather than the *where* of the graph. Clearly, there is important spatial information that must be processed in a graph, but more importantly, one must make sense of the gestalt properties of the graph (Baker et al. 2009). The first, and thus far only, neuroscience paper on graphical comprehension supports this finding that graphical comprehension depends mainly on perceptual processing in the ventral stream (Li et al. 2014).

We have argued that the human brain has been equipped, through the process of natural selection, with structures that allow it to not only process where things are in the environment, but also to make sense of scenes from the environment. Furthermore, it is the structures that make sense of scenes from the environment that mainly allow a person to understand a graphical representation. We would like to build upon this idea, which has been expressed to varying degrees elsewhere (Li et al. 2014; Pinker 1990), to hypothesize about what sorts of graphical representations will be most easy to understand. This theorizing is an original contribution to scientific knowledge.

There are two different ways in which brain structures can be used to solve problems. The first is native processing, which uses the brain to solve an adaptive problem. For example, processing human faces is an adaptive problem. Many generations of *Homo sapiens* faced situations where they had to identify not only whether a face was familiar, but also what was being expressed by the face. There is certainly survival value to this, as angry faces are more likely to be correlated with danger, and familiar faces are associated with additional valuable information about the potential behavior of a person. The processing of faces requires not only integrating the same basic sets of parts (eyes, a nose, lips, etc.) but making distinctions about subtle rearrangements of these same set of parts (Renzi et al. 2013). Within the ventral stream is an area of the brain called the fusiform face area that seems to be tuned to processing faces (Herrington et al. 2011). This area is specifically designed to process faces and that is what we mean by native processing.

A second way brain structures can be used to process information is to appropriate an area that is already adapted to solving problems similar to the current problem. For example, expert chess players seem to use the fusiform face area to recognize chessboards (Bilalic et al. 2011). Chess board recognition, like facial recognition requires making sense of the spatial relationships of a fixed number of component parts. It would seem that expert chess players appropriate the inherent structure and capabilities of the fusiform face area to better understand chessboards. Presumably, this is because there is not evolutionarily evolved structure to recognize chessboards.

While it is wonderful that the brain has enough flexibility to make use of naturally selected brain structures for novel tasks, there are some issues with the second type of processing. The obvious one is the fact that one needs expertise to do this. The brain does not automatically do it; it takes a lot of practice. Novices do not use the fusiform face area to recognize chessboards (Bilalic et al. 2011). Another problem is that even with training, the naturally selected area is better at performing its selected function than in performing the new one. In other words, even chess experts are better with faces than with chessboards. Thus, native processing is easier and more effective than appropriative processing.

With respect to graph comprehension our theory suggests that graphics that are more native will be better both in terms of ease and effectiveness. By native we mean graphs that are similar to adaptive environmental problems. A human brain will naturally contain structures designed to solve evolutionary problems (i.e. adaptive structures) (Cosmides et al. 1994). If a graph can represent information in the same way the environment represents the problem, then graph viewers can apply these adapted structures to solve the problem. The graphs that are more native are those that represent information in a manner similar to the environment.

When comparing sets of two dimension graphs to three-dimensional graphs Kumar and Benbasat (2004) found a "…surprising total dominance of 3D graphs (p 278)". Indeed, it was this surprising finding that motivated this paper. However, in light of an evolutionary theory of graph comprehension, this finding is not surprising. Nature is composed entirely of three-dimensional surfaces. All organisms since the beginning of life have faced the problem of three-dimensional surfaces. All sighted animals face the problem of processing visual information about three-dimensional surfaces. Primates in particular evolved many adaptations for an arboreal environment, which is much more three-dimensional than the environments most land animals occupy (Changizi et al. 2008). There is clear survival value for primates to have brain structures that natively solve problems of three-dimensional visual processing. While it is true that in the modern world we often deal with two-dimensional surfaces (computers and books for example), such surfaces were virtually absent in the environment in which our brains evolved. Thus, it makes sense that in human brains three-dimensional representations should be more natively processed than two-dimensional representations.

In this work we compare three-dimensional surface graphs to balloon race graphs, which have been popularized by Hans Rosling (Rosling 2006). We will offer specifics below. For our theoretical purposes the important issue is that we believe that there are brain structures evolved to process three-dimensional surfaces natively, but there are no such structures evolved to process balloon race graphs. This is because three-dimensional surfaces occur in the evolutionary environment and being able process them has survival value, thus they are the solutions to an adaptive problem. However, there is no naturally occurring analog to balloon race graphs in the evolutionary environment, so there will be no brain structures that deal directly with understanding balloon race representations. This means that balloon race graphs will have to be processed using appropriative processing. Thus we have two behavioral hypotheses.

Accuracy hypothesis*: Subjects will be more accurate at processing three-dimensional surface graphs than balloon race graphs.*

Speed hypothesis*: Subjects will be faster at processing three-dimensional surface graphs than balloon race graphs.*

Furthermore, we theorize that native processing of the meaning of a three-dimensional graphical representation will occur in the ventral stream. The visual content of a three-dimensional surface graph will be processed natively in the ventral stream. However, the visual content of a balloon race graph will be processed appropriatively using other structures. We do not know what these structures will be, but we can hypothesize that for three-dimensional surface graphs there will be more processing in the ventral stream. Our brain data hypothesis is:

Native Processing Hypothesis*: Three-dimensional surface graphs will produce greater activation in the ventral stream than balloon race graphs.*

It is important to note that we are not claiming that the ventral stream is an area of the brain designed to interpret graphs. To the contrary we are claiming that this is not the part of the brain that is designed to interpret graphs. Instead, the ventral stream is designed to interpret natural scenes. However, if a certain graph mimics a natural scene, then the ventral stream can interpret it more effectively. Thus, we do not expect a correlation between ventral stream activation and accuracy, we only expect a correlation between graph type and accuracy. An analogy will help to clarify this difference.

Imagine that you have a toolbox, which contains a variety of tools, including a hammer, but does not contain a screwdriver. Further imagine that you are asked to drive some nails into wood. You would pick the hammer to perform this task, and you would do a good job. You may not drive all of the nails perfectly, but even for nails that you failed to drive perfectly, you would use the hammer a similar amount. Now imagine that you were asked to drive screws into wood. Unfortunately, your toolbox contains no screwdriver. Maybe you could try the hammer, at least to start the screw. Then you could use some pliers, or perhaps a knife. You might be able to get some screws in, but you would not perform as well at driving screws into wood as you would at driving nails into wood. Use of the hammer would not be correlated with success, after controlling for task, because hammers are not tools for driving everything into wood, they are tools for driving nails into wood. If there were another task that were similar to driving nails into wood, you would also be able to perform well because you have a tool for that task.

If we change the toolbox to the brain, the hammer to the ventral stream, and a task similar to driving nails to interpreting a three-dimensional surface graph, we can make the same argument. It is a fit between the tool and the task that increases success, not simply the use of the tool. The ventral stream is, we hypothesize, good at interpreting three-dimensional surfaces, and interpreting three-dimensional surface graphs is similar enough to that task to allow the ventral stream to perform well. We do not expect a correlation between the accuracy and speed of graph reading and activation in the ventral stream (and indeed we do not find one as shown in the appendix). We do expect correlation between activation in the ventral stream and the similarity of a graph to a natural phenomenon. We also expect a correlation between having a tool to solve a problem and the accuracy and speed of the problem solution.

We would also like to note that just as hammers may be useful for solving other problems, such as pulling nails, cracking nuts, or slaying giants, the ventral stream does interpret more than three-dimensional surfaces. Face recognition, object recognition and scene recognition occur in the ventral stream, and there may be other visual tasks that are done there. Some of these visual tasks may be similar enough to other types of graphs to allow them to be processed effectively in the visual stream.

Procedures

The subjects for this study were 20 students in a US university[3], who received course credit for participating in a research subject pool. There were 11 males, and 17 right-handers. The average age was 22.05. Subjects were part of a subject pool and received course credit for participating. IRB approval was obtained prior to the experiment. Upon arrival the participants were greeted and asked to fill out a brief demographic questionnaire along with a consent form. After completion of the questionnaire and consent form, the participants were given an explanation of the task to be done when inside the fMRI machine. During this explanation, the participant was allowed to use a laptop that simulated the controls and interface to be encountered during the experiment to remove confusion about the task.

Once inside the fMRI machine, the participant received a scout scan which captures a low resolution outline of the head to allow the technician to aim the scanner during subsequent

[3] The best estimate of how many subjects a neuroimaging study needs based on power analysis is twelve. “12 subjects were required to achieve 80% power at the single voxel level for typical activations (Desmond et al. 2002)”. “[A]t least 12 subjects are required to detect the effect of stimulation on the fMRI signal change in M1 using a ROI-based analysis (Zandbelt et al. 2008).” “[F]or a threshold of p<. 05 with .80 power for a Level 2 analysis, N=12 was needed for the 1st level analysis (Dimoka 2012).”

scans, followed by a magnitude and phase scan to allow for B0 correction, and finally, a high resolution T1 scan. These scans took about 10 minutes. After these scans, the participant began the experiment by answering questions related to balloon graphs and three-dimensional surface graphs, while being scanned. The participant had 35 minutes to answer the questions presented and would be stopped if the participant exceeded the 35 minute timeline. 76% of participants were able to answer all questions in the stated time period.

Task

We created 30 data sets. Both a balloon graph and a three-dimensional surface graph were created for each data set. For each of these datasets, two questions were created, which represented an extraction question (answer existing on the graph) and an integration question (answer existing off the graph). Thus, there were 60 trials presented to subjects. To account for ordering effects or issues with "easier" questions for certain graphs, the questions (extraction or integration) were randomly asked to for each graph. Each subject received one extraction and one integration question from each data set and one balloon race and one three-dimensional surface graph from each data set. The questions were balanced so half of the three-dimensional graphs were extraction-based and half of the balloon race graphs were extraction-based. However, the actually pairings were random across subjects. All questions consisted of presenting the participant with 3 dimensions of information and asking for a value for the 4th dimension of data (see Figure 2). All answer options were presented as ranges of values, with four choices. The ranges were all the same size and each of the choices was correct 25% of the time. After answering a question, the participant would be returned to a blank screen with a random jitter, uniformly distributed, ranging between 0 and 2.5 seconds (which is one TR). After completing all the questions or reaching the 35-minute mark, the participants were thanked for their help and released.

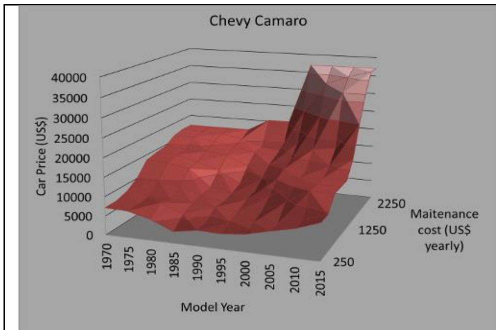

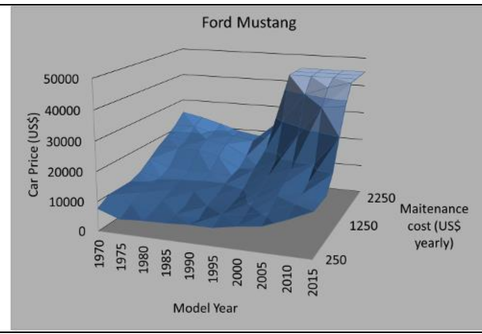

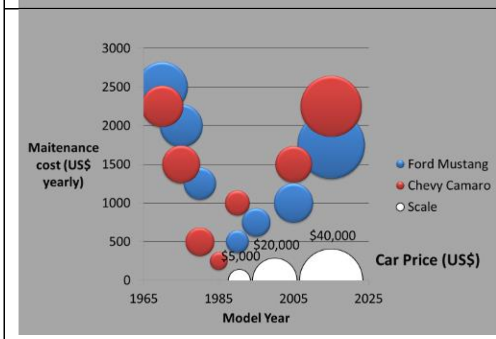

Sample stimuli. Subjects either saw both of the 3-D surface graphs, or just one balloon race graph. Each set of graphs came with two questions. The Z-axis (height) is matched with the size of the balloon.

Sample Question 1: What would be the price of a 2020 Ford Mustang, considering paying $2500 of yearly maintenance?

Sample Question 2: A 1975 Chevy Camaro costs $1500 a year to maintain, what is it`s price?

Figure 2: Sample stimuli

The questions were self-paced (i.e. a subject could take as long as they needed to answer the question, and the screen advanced after the subject answered) for several reasons. First, in pretesting we found significant variation in how long it took people to answer a question. Thus, we would either need a very long interval for the answers or have a large number of unanswered questions, both of which are problematic. Second, reading graphs is not typically a time constrained activity in the natural environment, so adding a time pressure component could result in confounding effects. People might begin guessing, which is not what we wanted to study or simply give up, which is also undesirable. Ultimately, we want to study how people solve graphs, not how people solve graphs under extreme time pressure. Third, self-paced graphs allowed us to explore the behavioral consequences of different types of graphs—specifically, does the time taken to solve a graph depend on the type of graph. Forth, it allowed us to collect more data. Not only do we have answer accuracy and speed for each question presented, but we are not wasting time having people look at the screen after they made a decision. Finally, it reduces unwanted noise. If a subject chooses and answer, then has to wait before moving on, their brain might be doing something other than solving graphs. If we average together the time they are solving graphs with the time they are doing something else, then we are adding noise to our variable of interest. Thus, for this experiment self-paced answers made more sense than fixed time answers often used with simpler stimuli. Figure 3 represents one set of the task performed by the subjects.

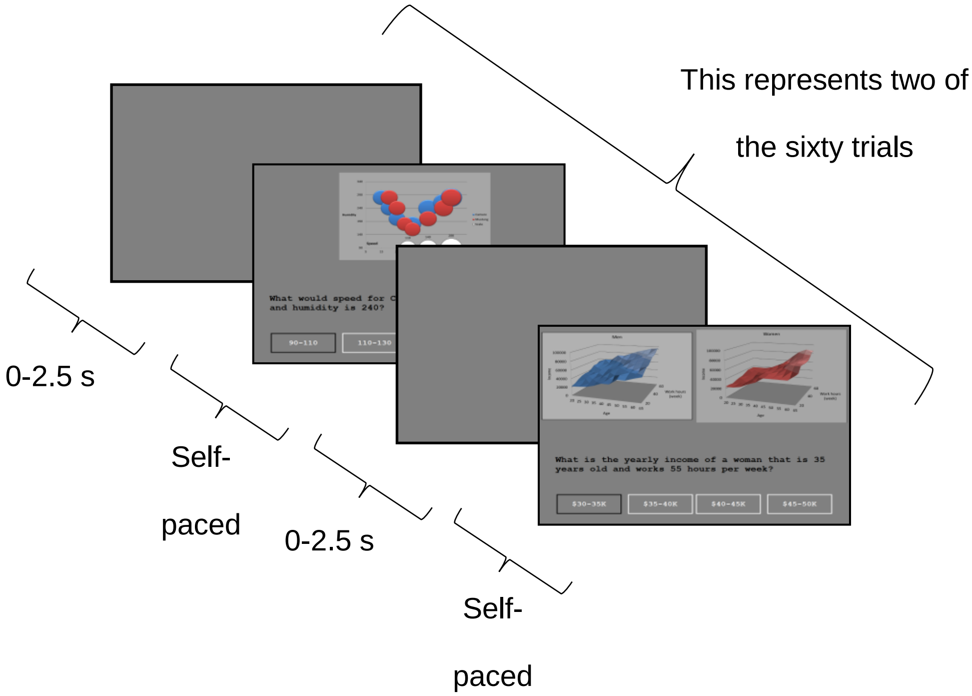

Figure 3: Task

Behavioral Results

To test the Accuracy Hypothesis, we used logistic regression with random effect to model the probability of a person being correct, as dependent variable, with graph type, question type, and graph type question type interaction as regressors. Since multiple responses from the same subject cannot be independent from each other, we added the subject into this model as a random effect. The regression equation for this logistic regression with random effect is shown in equation (1):

$$logit\ (Accuracy) = ln\left(\frac{\pi}{1-\pi}\right) = \beta_0 + \beta_1 \times \text{Question Type}_{ij} + \beta_2 \times \text{Graph Type}_{ij} + \beta_3 \times \text{Question Type}_{ij} \times \text{Graph Type}_{ij} + \xi_j + e_{ij} \quad (1)$$

Where π is the probability of a correct answer by jth participant, for a given question type and graph type, while e_{ij} and ξ_j are a trial error and a subject specific error.

Table 1 presents the results for response accuracy. Three-dimensional surface graphs showed a significantly positive effect on the probability of being correct (p-value < 0.001). However, neither question type nor the interaction term were significant (p-value 0.3544 and 0.2482, respectively). The results hold with a general linear model instead of a logistic model, and with a logistic regression model without considering subjects as random effects. Subjects answered three-dimensional graph questions correctly 48.4% of the time and balloon race graphs 34.6% of the time as shown in Figure 4. The results support the accuracy hypothesis.

Table 1: Logistic random effects model for response accuracy

Variable	**Coefficient**	**Standard Error**	**t-value**	**P-value**
Intercept	-0.34745	0.09364	-3.711	<0.001***
Question Type	0.1142	0.1234	0.926	0.3544
Graph Type	0.5842	0.1235	4.731	<0.001***
Graph Type X Question Type	0.2850	0.2468	1.155	.2482

Signif. Codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 '' 1 Note: The dependent variable is participants' time spent on each question

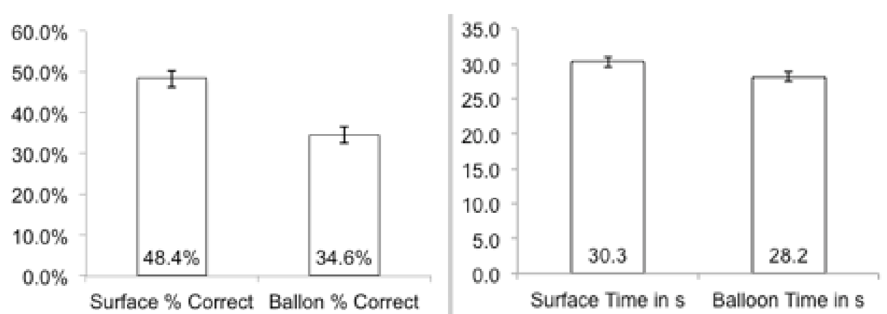

Figure 4: Comparison of accuracy and speed of answering questions from the two types of graphs.

To test the speed hypothesis, we ran a random effects model with response time as the dependent variable and graph type, question type, graph type question type interaction, as regressors, and the subject as a random effect. Type of graph was significant (p-value = 0.0279), but in the wrong direction (surface graphs took 3.705 seconds longer to solve). Type of question did not show a statistically significant effect (p-value= 0.1440). The interaction term was significant (p-value= 0.0253) and negative, indicating that integration questions displayed as surface graphs took less time to solve than would be predicted by each regressor alone. This indicates that subjects were only slower at answering three-dimensional surface graphs of the extraction type. They were actually faster at three-dimensional integration questions than at balloon race graph integration questions. Thus, the speed hypothesis was supported for integration questions but not for extraction questions. We discuss this below. Table 2 presents the results.

Table 2: Random effects model for response speed

Variable	**Coefficient**	**Standard Error**	**t-value**	**P-value**
Intercept	30.1625	2.0423	14.7688	<0.001***
Question Type	1.2134	0.8305	1.461	0.1440
Graph Type	1.8238	0.8298	2.198	<0.0279*
Graph Type X Question Type	-3.7118	1.6601	-2.236	0.0253*

Signif. Codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 '' 1 Note: The dependent variable is participants' time spent on each question

Behaviorally, subjects were more accurate answering three-dimensional graph questions than balloon race questions. Speed showed a mixed effect, subjects were fastest at answering integration questions in three-dimensional surface graphs (29.49 seconds), but slowest at answering extraction questions based on three-dimensional surface graphs (31.99 seconds). Integration and extraction questions based on balloon race graphs took on average 31.38 and 30.16 seconds respectively. This is an unexpected result, and rather than offer an ad hoc explanation we simply note that the speed hypothesis was supported for integration questions, but not for extraction questions.

Imaging Results

Please see the appendix for the technical details of the fMRI acquisition and analysis, as well as for robustness testing using permutation testing and alternative model specifications. In addition to a basic model contrasting graph types, the results were also robust to models controlling for the correctness of answers and the visual complexity of the display. All of this is reported in the appendix as requested by the review team.

The contrast examining where three-dimensional surface graphs showed greater activation than balloon race graphs is shown below in Figure 5 and **Table 3**. It clearly shows increased activation in the ventral stream for three-dimensional surface graphs. The activation runs from the intracalcarine cortex, through the lingual gyrus and down the fusiform cortex. There are no areas of greater activation in the dorsal stream or in the frontal regions. Activation is concentrated very strongly in the dorsal ventral stream. This provides support for the Native Processing Hypothesis.

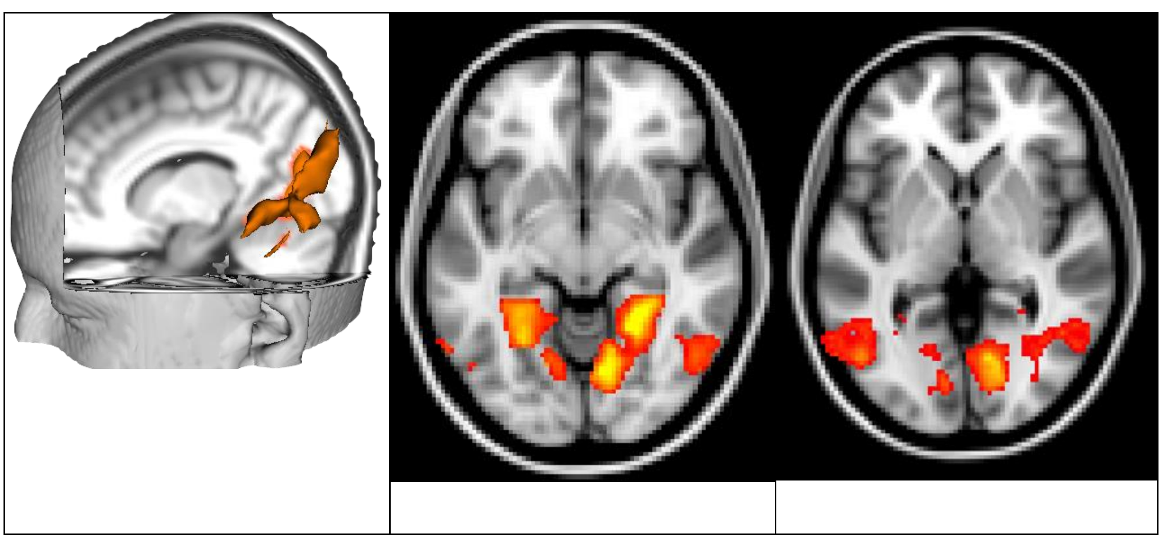

Z = -10

Z = 0

Z score 2.3  5.0 Left is on the left

Figure 5: Areas of greater activation for 3D surface graphs than for Balloon Race Graphs

Table 3: Details of clusters for 3-D surface graphs > Balloon Race Graphs

Cluster	**# of Voxels**	**P-value Cluster**	**Largest Z-score**	**MNI152 Coordinates of Largest Z**	**Region**
4	3985	2.82E-14	6.09	24, -46, -8	Right Temporal Occipital Fusiform Cortex
3	1427	1.43E-06	4.85	-26, -58, -8	Left Temporal Occipital Fusiform Cortex
2	679	1.65E-03	4.96	-50, -70, 2	Left Inferior Lateral Occipital Cortex
1	671	1.80E-03	3.98	-34, -84, 32	Left Superior Lateral Occipital Cortex

The contrast exploring where balloon race graphs produced greater activation than three-dimensional surface graphs is depicted in Table **4**. It tells an interesting story. First, there is activation in the primary visual cortex suggesting that balloon race graphs require more primitive visual processing. Second, there is a great deal of activation in the dorsal stream suggesting that special characteristics (i.e. where) are more important. There is also a great deal of activation in the frontal regions suggesting that people are doing a more conscious processing of the balloon race graphs. In addition, there is a wide dispersion of activation across the entire brain, this is common in compensatory settings where subject need to engage additional regions to solve problems (Bokde et al. 2008; Desmond et al. 2003; Staffen et al. 2002). Finally, we note that there are no areas of greater processing in the ventral stream. Figure 6 shows the difference in activation for the Balloon Race graphs in contrast to the three-dimensional surface graphs. Table 4 shows the details of each of these clusters of activation.

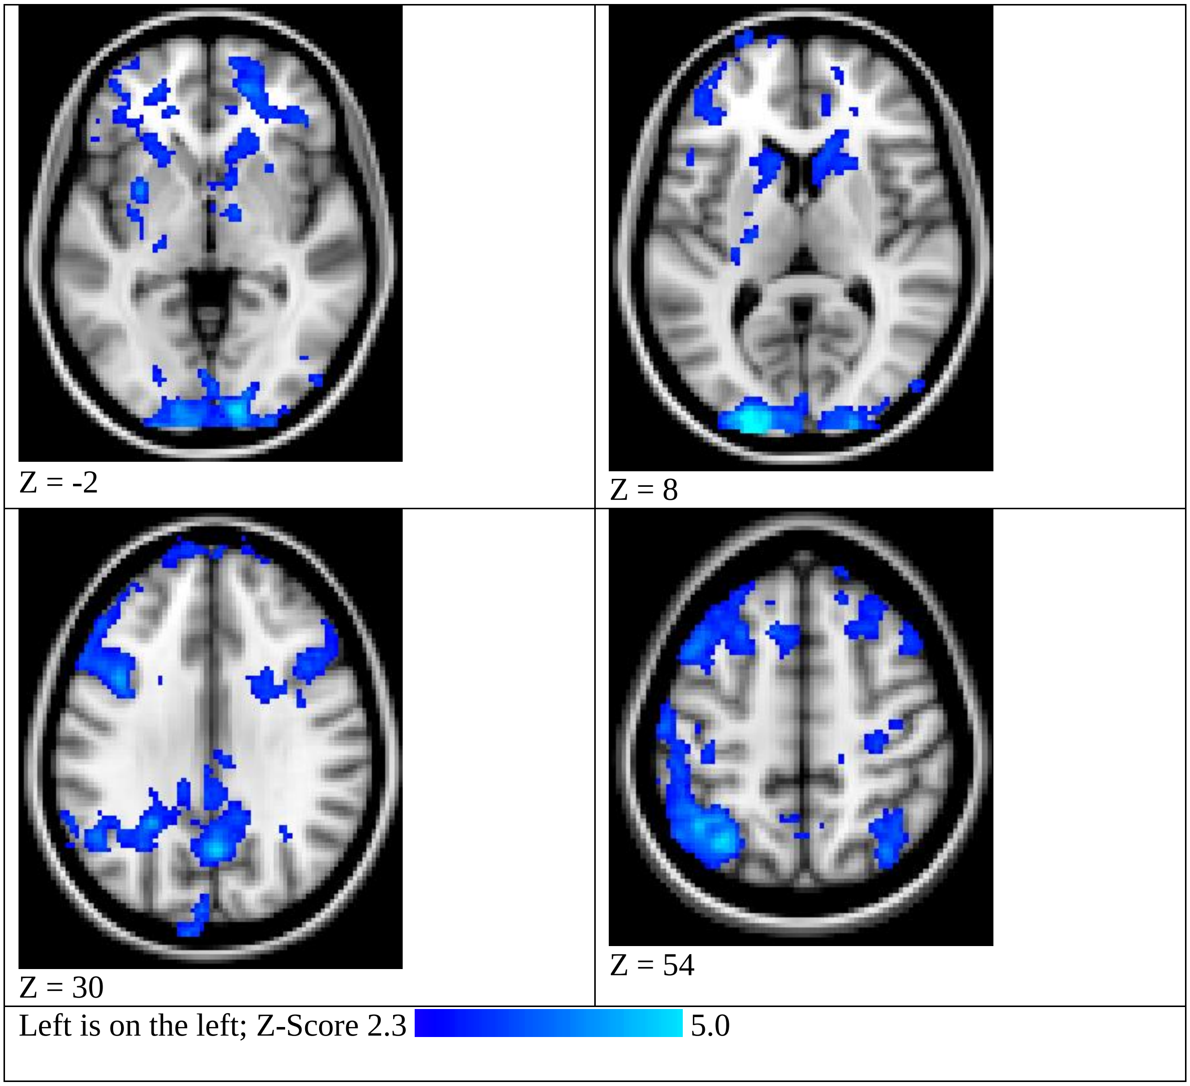

Figure 6: Areas of greater activation for Balloon Race graphs than for 3D surface graphs.

Table 4: Details of clusters for Balloon Race Graphs > 3-D Surface Graphs

Cluster	# of Voxels	P-value Cluster	Largest Z-score	MNI152 Coordinates of Largest Z	Region

5	8309	6.41E-23	4.72	-32, -64, 52	Left Superior parietal lobule
4	6071	2.33E-18	4.13	-40, 10, 30	Left Middle Frontal Gyrus
3	4321	2.19E-14	5.46	-20, -102, 8	Left Occipital Pole
2	2091	2.15E-08	3.65	24, 10, 32	Right Middle Precentral Gyrus
1	1550	1.19E-06	3.74	18, 48, -2	Right Anterior Cingulate

When comparing the areas of greater activation for three-dimensional surface graphs to the greater areas of activation for balloon race graphs, we find that there are 5,825 voxels of greater activation for three-dimensional surface graphs, but 21,136 voxels of greater activation for balloon race graphs. This suggests that subjects were "thinking" a great deal more about the balloon race graphs. Unfortunately, they were not using the ventral stream, but rather using a widely dispersed network of areas in both the frontal and far occipital regions, as well as the dorsal stream. This suggests that they were paying closer attention to the individual details of the graphs including the relative positions of different graphical elements, then using executive function and logical reasoning to determine the answers to the questions. However, these are more general cognitive areas that are not specialized for understanding visual information. It is more akin to a person trying to navigate a dark room full of furniture via touch than a lit room via sight. Certainly, it can be done, but touch is much less effective for navigation of a room full of furniture as countless shins can attest.

LIMITATIONS

This work has several limitations related to the use of fMRI methods. FMRI requires repeated measures within a single individual, and it requires these measures in an uncomfortable environment. For graph reading, each observation can be fairly time consuming. This limits the number of observations that can be collected, and because of this we were only able to look at two different types of graphs. It would be very valuable for future researchers to look at more than just two types of graphs. It would be worthwhile to directly compare many types of graphs in terms of speed and accuracy outside of the scanner, and then to identify candidates for research in a scanner where only few comparisons are possible.

Related to this, we are limited in only considering one type of natural graphical representation. There could certainly be more. The ventral stream is known to process faces, so something like the Wong Baker pain scale (Wong et al. 2001) could be natural, and maybe there are others. It would certainly be worthwhile to establish some measure of naturalness and apply it to graphs to establish what sorts of graphs are more natural. Such inquiry could lead to new types of graphs.

Another issue we face that arises from evolutionary studies in general is controlling for experience. Determining the relative contributions of nature and nurture to human behavior is a long and ongoing question across many sciences. It would be ideal to have a group of graphs that were completely novel to the viewers. While ideal, this is probably infeasible as it would likely require an aboriginal population who did not study graphs as part of their regular education. Transporting sizable number of such people to large city and putting them in an MRI machine would be traumatic for those individuals. An alternative would be to develop entirely new types of natural and unnatural graphs, but that might be infeasible too. Nonetheless, it would be useful for future research to look at the effect of training on graphical comprehension and how and if the neural circuits used for comprehension changed over time.

IMPLICATIONS FOR PRACTICE

Data visualization has become an important method for representing information in user interfaces. From graphics in Excel, to proprietary packages like Periscope or open source solutions like D3.js graphics are increasingly being added to information systems. While there is a great deal of interest in capturing data and representing it graphically, there is little knowledge about how human brains make use of these graphs. By making inroads into the exploration of how human use graphs, this paper offers several implications for IS practice.

First, we have found that three-dimensional surface graphs facilitate comprehension more than balloon race graphs. Balloon race graphs are based on Gapminder's Trendalyzer software, which was acquired by Google in 2007 for an undisclosed amount. Thus, at least one company spent money to specifically acquire balloon race graph software. It would certainly be worth it for Google to have known before acquiring the company how much balloon race graphs facilitate comprehension relative to three-dimensional surface graphs, which they could have developed for free. More generally, it is worthwhile for designers of interfaces to know what effect their chosen graphs will have on the users of the interface.

In addition to our findings, Kumar and Benbasat (2004) also found three-dimensional surface graphs to dominate two-dimensional graphs. Thus, the simple finding is to use three-dimensional surface graphs when possible, or at least in preferences to balloon race or line graphs.

More generally, our prescription is to use graphical interfaces that are processed natively by the human brain. The only type we are currently confident about is three-dimensional surface graphs, but there are potentially others. The important thing is that not all graphs are comprehended equally. Therefore, organizations that use and organizations that design graphics need to understand which graphs are natively processed by human brains to encourage comprehension.

Of course, comprehension does not need to be the goal of including a graph in an interface. Persuasion, comfort or even entertainment could also be reasons to include graphs. Thus, before including a graph in an interface the system designer should ask what is the purpose of the graph. If the purpose of the graph happens to be comprehension, then graphs that are processed natively in general and specifically three-dimensional surface graphs are good choices.

DISCUSSION

In this research, we hypothesized that people should perform better on graphical tasks when they have neural structures that are natively able to handle the graphical representation. To test this, we compared three-dimensional surface graphs to balloon race graphs, and argued that because the world is made up of three-dimensional surfaces the brain should contain structures specifically adapted to processing three-dimensional surfaces. What we found is that subjects are much more accurate (though not consistently faster) at answering questions represented in three-dimensional graphs than in balloon race graphs.

More importantly, we found brain activation that supported our theory. When answering questions about three-dimensional graphs subjects used a very concentrated set of structures in the ventral stream, which is where the brain integrates the various components of a visual image to understand what is being viewed. Conversely, balloon race graphs gave rise to a diffuse set of brain areas associated with compensatory effort. This diffuse set of brain areas comprised nearly four times as much brain volume as the focused set of regions used for three-dimensional surface graphs.

Overall, these results suggest that when subjects viewed three-dimensional surface graphs, they made more accurate decisions with less activation because they used the best-optimized set of brains areas. They used regions of the brain that are known to be important for understanding what is being viewed. We believe they use these brain regions because evolution has endowed *Homo sapiens* (and most sighted animals) with very powerful machinery for processing the meaning and implications of three-dimensional surfaces.

Our findings should not be misinterpreted as saying three-dimensional surface graphs are the best sorts of graphs. Rather what we claim is that graphs, which can be processed natively, are the best sorts of graphs. Natively means that there are brain structures, developed over time by evolutionary pressures in the natural environment, to address the sort of visualizing produced by the graphic. Certainly, there were few evolutionary pressures related to balloon race type visualizations, but many related to three-dimensional surfaces. However, there are likely to be other types of visualizations that are also adaptive. One example might be graphs based on facial recognition, like the Wong-Baker pain scale. The important point is that there are graphs that represent data in a way the brain is evolved to interpret, and graphs that represent data in a way the brain is ill equipped to interpret, because there was never any evolutionary pressure to develop an arrangement of neurons that would quickly and efficiently process that sort of image.

To be sure, brains are plastic and can recruit regions that are specifically designed for task X to help with task Y. With training and practice a brain's ability to do the processing needed for task Y will increase. However, the whole point of a graphical representation is to present data in a way that is intuitively obvious.

Our theory is slightly different than cognitive fit theory (Vessey 1991) that is often applied in the study of graphical comprehension (Dennis et al. 1998; Huang et al. 2006; Umanath et al. 1994). Cognitive fit argues for a correspondence between task and information presentation. Evolutionary fit argues for a correspondence between information presentation and evolutionarily adaptive brain structures. These theories are not in conflict;, they simply examine different links in the chain from task to information presentation to comprehension. Our theory predicts that the important dimension of information presentation is how well it matches adaptive information in the evolutionary environment. Of course, many tasks in the modern world have no evolutionary counterparts, but to the degree that one can restructure a task to take advantage evolutionarily endowed brain structures, it makes sense to do so. Cosmides (1985) research on cheater detection is a good example of such design. In this problem people are asked to make logical deductions about a statement in the form of If P then Q. When presented in abstract terms, most people miss modus tollens (i.e. not Q implies not P). However, when presented in terms of cheater detection (i.e. I do work, you pay me) then people notice modus tollens (if you don't pay me I don't do work).

In addition to a comparison of graphs, we propose a neural correlate of good graphs. Good graphs are those that are processed in the ventral stream. The ventral stream is the part of the brain built by evolution to solve the adaptive problem of understanding what is being seen. The ventral stream contains multiple structures that process different parts of the meta-task of understanding what is being seen, so there are likely multiple different ways of understanding what is being seen. Nonetheless, the ventral stream as a whole is the primary area for structures that understand visual inputs, and it is the ventral stream that forwards the meaning of visual input to the rest of the brain for further processing. Thus, graphs that stimulate the ventral stream more will tend to be those, which are being processed natively, and hence that are best understood.

REFERENCES

Aldrich, F., and Sheppard, L. 2000. "Graphicacy; The fourth 'R'?," *Primary Science Review*), pp 8-11.

Baker, J., Jones, D., and Burkman, J. 2009. "Using visual representations of data to enhance sensemaking in data exploration tasks," *Journal of the Association for Information Systems* (10:7), p 2.

Balchin, W. 1976. "Graphicacy," *The American Cartographer* (3:1), pp 33-38.

Benbasat, I., and Dexter, A. S. 1985. "An Experimental Evaluation of Graphical and Color-Enhanced Information Presentation," *Management Science* (31:11), pp 1348-1364.

Benbasat, I., and Dexter, A. S. 1986. "An Investigation of the Effectiveness of Color and Graphical Information Presentation under Varying Time Constraints," *MIS Quarterly* (10:1), pp 59-83.

Benbasat, I., and Schroeder, R. G. 1977. "An Experimental Investigation of Some MIS Design Variables," *MIS Quarterly* (1:1), pp 37-49.

Boardman, D. 1990. "Graphicacy Revisited: mapping abilities and gender differences," *Educational Review* (42:1) 1990/01/01, pp 57-64.

Bokde, A. L. W., Lopez-Bayo, P., Born, C., Dong, W., Meindl, T., Leinsinger, G., Teipel, S. J., Faltraco, F., Reiser, M., and Möller, H.-J. 2008. "Functional abnormalities of the visual processing system in subjects with mild cognitive impairment: an fMRI study," *Psychiatry Research: Neuroimaging* (163:3), pp 248-259.

Changizi, M. A., and Shimojo, S. 2008. "“X-ray vision” and the evolution of forward-facing eyes," *Journal of theoretical biology* (254:4), pp 756-767.

Chau, M. 2011. "Visualizing web search results using glyphs: Design and evaluation of a flower metaphor," *ACM Transactions on Management Information Systems (TMIS)* (2:1), p 2.

Chen, H., Chiang, R. H., and Storey, V. C. 2012. "Business Intelligence and Analytics: From Big Data to Big Impact," *MIS quarterly* (36:4), pp 1165-1188.

Correll, M., Albers, D., Franconeri, S., and Gleicher, M. Year. "Comparing averages in time series data," Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM2012, pp. 1095-1104.

Cosmides, and Tooby 1994. "Beyond intuition and instinct blindness: Toward an evolutionarily rigorous cognitive science," *Cognition* (50:1) **Need Date, pp 41-77.

Cosmides, L. 1985. *Deduction or Darwinian Algorithms?*, Citeseer.

Cox, R., Romero, P., du Boulay, B., and Lutz, R. 2004. "A Cognitive Processing Perspective on Student Programmers’ ‘Graphicacy’," in *Diagrammatic Representation and Inference,* A. Blackwell, K. Marriott and A. Shimojima (eds.), Springer Berlin Heidelberg, pp. 344-346.

Darwin, C. 1859. "On the origins of species by means of natural selection," *London: Murray*).

Darwin, C., and Wallace, A. 1858. "On the tendency of species to form varieties; and on the perpetuation of varieties and species by natural means of selection," *Journal of the proceedings of the Linnean Society of London. Zoology* (3:9), pp 45-62.

Dekker, T., Mareschal, D., Sereno, M. I., and Johnson, M. H. 2011. "Dorsal and ventral stream activation and object recognition performance in school-age children," *NeuroImage* (57:3) 8/1/, pp 659-670.

Dennis, A. R., and Carte, T. A. 1998. "Using geographical information systems for decision making: Extending cognitive fit theory to map-based presentations," *Information Systems Research* (9:2), pp 194-203.

Desmond, J. E., Chen, S. A., DeRosa, E., Pryor, M. R., Pfefferbaum, A., and Sullivan, E. V. 2003. "Increased frontocerebellar activation in alcoholics during verbal working memory: an fMRI study," *Neuroimage* (19:4), pp 1510-1520.

Desmond, J. E., and Glover, G. H. 2002. "Estimating sample size in functional MRI (fMRI) neuroimaging studies: Statistical power analyses," *Journal of Neuroscience Methods* (118:2) Aug 30, pp 115-128.

Dimoka, A. 2012. "How to Conduct a Functional Magnetic Resonance (Fmri) Study in Social Science Research," *Mis Quarterly* (36:3) Sep, pp 811-840.

Goodale, M. A., and Milner, A. D. 1992. "Separate visual pathways for perception and action," *Trends in neurosciences* (15:1), pp 20-25.

Herrington, J. D., Taylor, J. M., Grupe, D. W., Curby, K. M., and Schultz, R. T. 2011. "Bidirectional communication between amygdala and fusiform gyrus during facial recognition," *Neuroimage* (56:4) Jun 15, pp 2348-2355.

Huang, Z., Chen, H., Guo, F., Xu, J. J., Wu, S., and Chen, W.-H. 2006. "Expertise visualization: An implementation and study based on cognitive fit theory," *Decision Support Systems* (42:3), pp 1539-1557.

Kelton, A. S., Pennington, R. R., and Tuttle, B. M. 2010. "The effects of information presentation format on judgment and decision making: A review of the information systems research," *Journal of Information Systems* (24:2), pp 79-105.

Kim, G., and Lui, C. 2010. "The impacts of blink on graphical information processing," *International Journal of Information Processing and Management* (1), pp 124-132.

Kim, G., and Lui, S. M. 2011. "The effect of color combination on visual attention and usability of multiple line graphs," *Journal of Communications and Information Sciences* (1), pp 11-21.

Koch, K., McLean, J., Segev, R., Freed, M. A., Berry, M. J., Balasubramanian, V., and Sterling, P. 2006. "How much the eye tells the brain," *Current Biology* (16:14), pp 1428-1434.

Kravitz, D. J., Saleem, K. S., Baker, C. I., and Mishkin, M. 2011. "A new neural framework for visuospatial processing," *Nature Reviews Neuroscience* (12:4), pp 217-230.

Lambon Ralph, M. A., Graham, K. S., Patterson, K., and Hodges, J. R. 1999. "Is a Picture Worth a Thousand Words? Evidence from Concept Definitions by Patients with Semantic Dementia," *Brain and Language* (70:3) 12//, pp 309-335.

Larkin, J. H., and Simon, H. A. 1987. "Why a Diagram is (Sometimes) Worth Ten Thousand Words," *Cognitive Science* (11:1), pp 65-100.

Li, M., Lu, S., Wang, J., Ma, L., Zhang, M., and Zhong, N. 2014. "Ventral Stream Plays an Important Role in Statistical Graph Comprehension: An fMRI Study," in *Brain Informatics and Health*, Springer, pp. 12-20.

Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, W.H. Freeman and Company: San Francisco, CA.

Mayer, R. E., and Sims, V. K. 1994. "For whom is a picture worth a thousand words? Extensions of a dual-coding theory of multimedia learning," *Journal of Educational Psychology* (86:3), pp 389-401.

Parsons, P., and Sedig, K. 2014. "Adjustable properties of visual representations: Improving the quality of human-information interaction," *Journal of the Association for Information Science and Technology* (65:3), pp 455-482.

Pinker, S. 1990. "A theory of graph comprehension," *Artificial intelligence and the future of testing*), pp 73-126.

Postigo, Y., and Pozo, J. I. 2004. "On the Road to Graphicacy: The learning of graphical representation systems," *Educational Psychology* (24:5) 2004/09/01, pp 623-644.

Remus, W. 1987. "A Study of Graphical and Tabular Displays and Their Interaction with Environmental Complexity," *Management Science* (33:9) 1987/09/01, pp 1200-1204.

Renzi, C., Schiavi, S., Carbon, C. C., Vecchi, T., Silvanto, J., and Cattaneo, Z. 2013. "Processing of featural and configural aspects of faces is lateralized in dorsolateral prefrontal cortex: A TMS study," *Neuroimage* (74) Jul 1, pp 45-51.

Rilling, J. K. 2006. "Human and nonhuman primate brains: are they allometrically scaled versions of the same design?," *Evolutionary Anthropology: Issues, News, and Reviews* (15:2), pp 65-77.

Rosling, H. 2006. "Debunking third world myths with the best stats you've ever seen," *online]. Ted Talks-Ideas Worth Spreading. Retrieved July* (25), p 2011.

Scaife, M., and Rogers, Y. 1996. "External cognition: how do graphical representations work?," *International journal of human-computer studies* (45:2), pp 185-213.

Staffen, W., Mair, A., Zauner, H., Unterrainer, J., Niederhofer, H., Kutzelnigg, A., Ritter, S., Golaszewski, S., Iglseder, B., and Ladurner, G. 2002. "Cognitive function and fMRI in patients with multiple sclerosis: evidence for compensatory cortical activation during an attention task," *Brain* (125:6), pp 1275-1282.

Tractinsky, N., and Meyer, J. 1999. "Chartjunk or Goldgraph? Effects of Presentation Objectives and Content Desirability on Information Presentation," *MIS Quarterly* (23:3), pp 397-420.

Tversky, B., Morrison, J. B., and Betrancourt, M. 2002. "Animation: can it facilitate?," *International journal of human-computer studies* (57:4), pp 247-262.

Umanath, N. S., and Vessey, I. 1994. "Multiattribute Data Presentation and Human Judgment: A Cognitive Fit Perspective*," *Decision Sciences* (25:5-6), pp 795-824.

Ungerleider, L. G., and Haxby, J. V. 1994. "'What'and 'where'in the human brain," *Current opinion in neurobiology* (4:2), pp 157-165.

Vessey, I. 1991. "Cognitive Fit: A Theory-Based Analysis of the Graphs Versus Tables Literature*," *Decision Sciences* (22:2), pp 219-240.

Vessey, I., and Galletta, D. 1991. "Cognitive Fit: An Empirical Study of Information Acquisition," *Information Systems Research* (2:1) 1991/03/01, pp 63-84.

Walther, J. B., Slovacek, C. L., and Tidwell, L. C. 2001. "Is a Picture Worth a Thousand Words?: Photographic Images in Long-Term and Short-Term Computer-Mediated Communication," *Communication Research* (28:1) February 1, 2001, pp 105-134.

Wandell, B., Dumoulin, S. O., and Brewer, A. A. 2008. "Visual cortex in humans," *Encyclopedia of neuroscience* (10), pp 251-257.

Wong, D. L., and Baker, C. M. 2001. "Smiling face as anchor for pain intensity scales," *Pain* (89:2-3), pp 295-297.

Zandbelt, B. B., Gladwin, T. E., Raemaekers, M., van Buuren, M., Neggers, S. F., Kahn, R. S., Ramsey, N. F., and Vink, M. 2008. "Within-subject variation in BOLD-fMRI signal changes across repeated measurements: quantification and implications for sample size," *Neuroimage* (42:1) Aug 1, pp 196-206.

Zhu, B., and Watts, S. A. 2010. "Visualization of network concepts: The impact of working memory capacity differences," *Information Systems Research* (21:2), pp 327-344.