Visualization Taxonomy and Display.

Name:	Paul Gokin
Course:	HF 730, Spring 2005
Instructor:	William Gribbons, Ph.D.
Date:	May 15, 2005

Three ways to classify visualizations.

After completing the visualization matrix, I started looking for ways to organize the long list of visualizations in a way that would be meaningful to someone trying to select the most effective visualization to support a particular user task. I have identified three major ways in which all visualizations can be classified: by primary purpose, task area, and the individual perceptual and cognitive sub-tasks supported by the visualization.

Way #1. Classifying visualizations by primary purpose: exploration vs. communication.

John Tukey has identified two purposes a graph can serve: prospecting and transfer. After reviewing the visualizations we've studied this semester, I noticed that not only graphs, but other visualizations as well seem to serve *one* of these two purposes better than the *other*:

1. **Exploration**: allowing the viewer to explore a set of ideas, data, or outcomes to see relationships and patterns in order to arrive at conclusions, decisions, or discoveries. Some characteristics of these visualizations include high data-ink ratio and little or no reduction of detail to ensure data integrity, accurate decisions, etc.
2. **Communication**: communicating ideas, discoveries, and conclusions to the viewer, most often with the intent to teach or persuade the viewer. Some of the characteristics of visualizations that fall into this category include high visual impact, extensive precomputation (to display the data in the way that conveys the message most clearly), familiarity and use of metaphors (to reduce learning required to benefit from the visualization), extensive use of imagery (to enhance memorability of the message delivered by the visualization), and so on.

While many visualizations can be effectively accomplish both purposes, either "out of the box" or altered to better serve each purpose, I do believe it is useful to think of all visualizations in this way, not just graphs. Having said that, most visualizations appear to reside somewhere between communicating well and poorly, or enabling exploration well and not. This makes classifying them this way a challenge, but I believe I have the answer to this riddle in the visualization created for this paper.

Way #2. Classifying visualizations by general task areas.

This course (HF730) is organized around four general task areas: creative thinking (which includes creative problem solving), decision making, instruction, and data exploration (which includes both static and dynamic visualizations). These areas are perhaps the most mutually exclusive groups into which the visualizations we've covered this semester can be placed. They serve as the primary criterion in the visualization winnowing tool proposed for this paper.

Way #3. Classifying visualizations by sub-tasks: *The sub-task taxonomy.*

The rationale for using tasks as the basis for classifying visualizations is based on the simple fact that visualizations don't exist for their own sake—they exist to facilitate tasks. While visualizations are already broken up into broad task areas (way #2), I wondered whether a finer-grained classification—using finer, but still universal, tasks—could be developed. First, I needed to come up with a list of universal tasks. I used three sources: Bloom's educational objectives taxonomy, the visualizations themselves and the tasks they claim to facilitate, and the general task areas (creating thinking, decision making, etc.) we've covered this semester. Extracting tasks from the latter two sources helped fill in the gaps in Bloom's taxonomy. Now that I had the list, I eliminated tasks that were redundant and ordered them.

I've decided to order the tasks roughly in the order of increasing cognitive resources involved. For example, detection involves the sensory systems almost exclusively. Shifting of attention, estimation of stimulus magnitude, juxtaposition, comparison involve both the visual and cognitive systems. Discovering relationships and structure involve higher order cognitive capabilities. This ordering means that, for the most part, in order for any task to be carried out, all or most of the tasks listed **above it** must be completed as well (see the illustration on the next page).

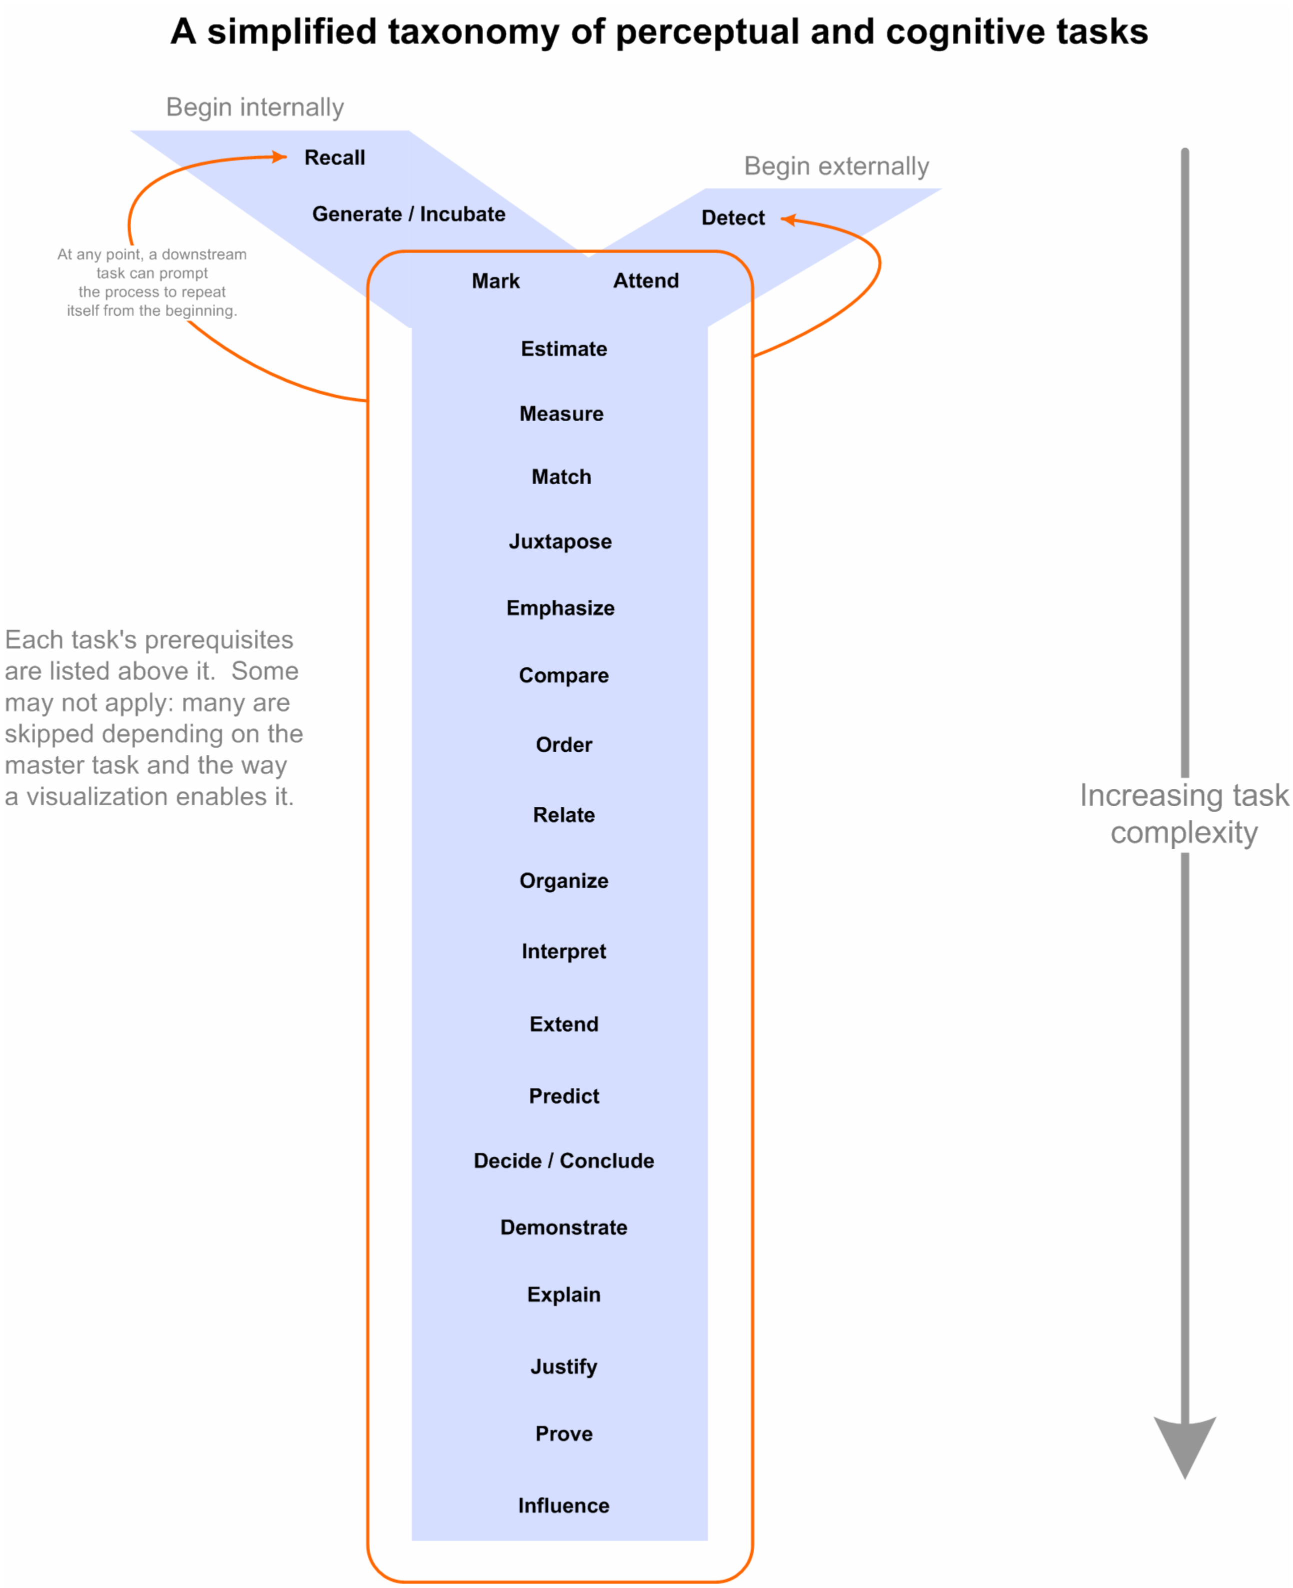
A simplified taxonomy of perceptual and cognitive tasks
Begin internally
Recall
Generate / Incubate
Begin externally
Detect
At any point, a downstream task can prompt the process to repeat itself from the beginning.
Mark
Attend
Estimate
Measure
Match
Juxtapose
Emphasize
Compare
Order
Relate
Organize
Interpret
Extend
Predict
Decide / Conclude
Demonstrate
Explain
Justify
Prove
Influence
Each task's prerequisites are listed above it. Some may not apply: many are skipped depending on the master task and the way a visualization enables it.
Increasing task complexity

At the top (beginning) of the task continuum are the two ways of entering it, based on the two general sources of initial stimulation of the mind: internal, through recall, and external, through detection by the senses. This separation became clear when, after initially having created the taxonomy for data exploration tasks alone, I was augmenting it to incorporate creative thinking, where tasks are initiated through recall as much as they are through detection. Another thing to notice is that any task listed in the main task continuum may prompt recollection or detection, which leads the mind to the appropriate starting point at the top. One example where this type of recursion occurs by design is during mind mapping when, after putting down several ideas and juxtaposing them we recall another associated idea.

Using the task taxonomy to organize visualizations.

This is all very nice and neat, but how does this taxonomy (or a two-headed task complexity continuum if you prefer) help us organize visualizations and help someone trying to select the one that would support the users' task(s) most effectively?

Let's consider the main purpose of a visualization—converting serially-executed cognitive tasks into concurrently-executed visual tasks. Now consider that each task in the taxonomy imposes a cognitive load that can be reduced by offloading this task to the visual system. Therefore, each visualization we wish to classify can be rated in terms of how well it enables each of these tasks to be completed *visually*. So the continuum can be used to select the master task, that task's "upstream" pre-requisites, and evaluate a visualization based on how well it enables the pre-requisites. If it passes muster on all of them, then it is likely to be a good choice.

One example is a task where the user must *predict* the value of a variable based on historical data (*predict* one of the higher-level tasks in the taxonomy). One of the best ways to do this visually is by extrapolating a trend using a regression line. Visualizations that have this capability—and allow the user to take precise readings at different points on the line (in other words *measure* and *compare*—both lower level tasks)—will be successful in supporting this task. On the other hand, a visualization that doesn't support taking measurements—even if it has a lovely regression line—will fail.

Requirements for the visualization selection tool.

In creating an application that would work for the uninitiated analyst, we must, in turn think about that person's goals and abilities. For example, the solution must encourage divergent thinking, be an effective decision aid, take advantage of the user's existing knowledge, and of course, not subject the analyst to undue cognitive load. The requirements of this assignment also state that the system must serve an instructional role by informing the designer how each visualization exploits the capabilities and compensates for the weaknesses of the human mind. The VizMaster© 5000™ is a dynamic visualization tool designed to meet these requirements.

Putting it all together: The VizMaster© 5000™ (see next page for screen illustration)

Regarding the analyst's prior knowledge, the interface assumes very little. The main assumption is that the user of the system is familiar with the general task area: creative, decision making, etc. Once this selection is made, the analyst is presented with a two-dimensional plot of all visualizations that fall into this general area. The visualizations are plotted along the exploration and communication axes. This approach solves the problem of most graphics falling somewhere in the middle of the scale in both respects (this is the riddle from bottom of page 2). A sortable, tabular list is also presented below. Since the number of visualizations can be quite high, there are multiple ways in which the analyst can filter the list:

- resize and drag the scroll bars to filterzoom (perform dynamic, zoom-linked filtering);
- the analyst can also check off individual sub-tasks in the top right area to narrow the list down to just those visualizations that support the checked sub-tasks effectively.

Each visualization's collapsible table entry includes a brief summary of the visualization's strengths and weaknesses in enabling the subtasks the analyst has selected. Because the summary is customized to the tasks the analyst is interested in enabling, the summary is more relevant to the analysis's choice, reducing workload associated with sifting through long lists of irrelevant information. Checking the compare box and clicking the compare button brings up a screen (not shown) where the selected visualizations can be compared side by side in detail. This is meant to happen at the later stages in the process of winnowing the visualization database; once a group of highly suitable candidates has been identified via the application's filtering tools.

As I was finishing up this assignment I realized that just as some visualizations are better at supporting exploration while others communication, one visualization may support a subtask better than another. The system could be expanded to take this into account and calculate the total utility of each visualization in the comparison view to help the analyst make the choice.