

Goal-Directed Zoom

Allison Woodruff, James Landay, Michael Stonebraker
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA 94720 USA
email: {*woodruff,landay,mike*}@cs.berkeley.edu

ABSTRACT

We introduce a novel zoom method, *goal-directed zoom*. In a goal-directed zoom system, users specify which representation of an object they wish to see. The system automatically zooms to the elevation at which that representation appears at appropriate detail. We have extended a database visualization environment to support end-user construction of visualizations that have goal-directed zoom. We present a sample visualization we have constructed using this environment.

Keywords

Clutter, information density, semantic zoom, visualization

INTRODUCTION

Visualization systems commonly represent objects in a two-dimensional canvas over which the user may pan and zoom. In such systems, zooming changes the user's distance from the canvas, also known as the perceived *elevation*. This spatial metaphor implies a camera-like model in which elevation determines (1) the proportion of the canvas visible in the display and (2) the proportion of display space occupied by the visible objects (*i.e.*, the size of each object).

Because the amount of display space available to an object varies with elevation, a graphical representation of an object that has appropriate visual complexity at one elevation may have inappropriate visual complexity at other elevations. Many zooming systems address this issue by supporting multiple graphical representations of objects. Multiple representations allow a balance between display density, elevation, and representation.

Existing systems reflect this balance between elevation and object representation in one of two ways. In naïve systems, elevation and representation are decoupled; the user chooses a representation and then zooms until the display “looks right.” More sophisticated systems couple elevation and representation using an operation known as *semantic zoom*, *e.g.*, [4]. In semantic zoom systems, an expert user defines a range of elevations in which each representation is valid. When the end user zooms to a given elevation, the

system displays each object using the representation valid at that elevation.

Consider the relationship between zooming and choice of representation in such systems. In naïve systems, elevation and choice of representation are controlled independently by the user. In semantic zoom systems, the elevation determines the choice of representation. An alternative is a system in which the choice of representation determines the elevation. We call the functionality provided by such a system *goal-directed zoom*.

In the next section, we describe the desirable characteristics of goal-directed zoom. We then present our design and implementation of a system to support the end-user construction of goal-directed zoom visualizations. Finally, we present an extended example and conclude.

CHARACTERISTICS OF GOAL-DIRECTED ZOOM

We propose that a goal-directed zoom system ideally possesses at least the following three properties:

- 1a. *Menus*: The system presents the user with a selection of possible object representations. (The user should be able to invoke this selection using some lightweight mechanism.)
- 2a. *Previews*: The selection mechanism graphically depicts the possible representations, *e.g.*, an icon of a house could be used to indicate a representation of a city in which houses are shown.
- 3a. *Automatic zoom*: When the user selects a representation, the system zooms to an elevation “appropriate” to that representation. (This elevation may be determined by ranges specified in a semantic zoom system, or they may be calculated in some other manner. One such method is described below.)

Note that each of these properties can apply either to a set of objects present in a visualization or to individual objects in a visualization. Ideally, these properties apply to an individual object of interest, yielding the following additional guidelines:

- 1b. The menu is appropriate for the specific type of object selected.
- 2b. The graphical options presented to the user are based on the specific object selected.
- 3b. The elevation to which the system zooms depends on the specific object selected.

A number of systems provide some subset of the functionality listed above, but to our knowledge, no system meets all of these criteria. For example, BigBook (<http://www.bigbook.com>), an online directory, provides maps with iconic zoom buttons. However, the buttons zoom to fixed elevations and the icons are abstract scale indicators. As another example, our work has some similarities to Magic Lenses [2] or portal filters [4]. However, while these mechanisms show different representations of objects, they do not zoom automatically to display appropriate detail for a given object. Finally, Pad provides primitives that could be used to support goal-directed zooming, but does not support it directly. The Pad++ web browser, for example, partially supports property 3a [1]. For a more extended discussion of related work, see [5].

A GOAL-DIRECTED ZOOM SYSTEM

We have designed and implemented a system that supports goal-directed zoom. When users click on an object in a visualization, our system presents a menu of the possible graphical representations of that object. When the user selects a representation, the system automatically pans so the selected object appears in the center of the visualization and then zooms to the elevation at which the selected representation has appropriate visual density.

Our system, VIDA (Visual Information Density Adjuster), is an extension of the DataSplash database visualization development environment. VIDA allows users to create multiple representations of objects [5].

VIDA applies the cartographic Principle of Constant Information Density [3] to choose which representation is appropriate at a given elevation (and conversely, which elevation is appropriate for a given representation). In its general formulation, the Principle of Constant Information Density states that the amount of visual information per display unit should remain constant as the user pans and zooms. VIDA assumes that the screen is divided into grid cells. As the user pans and zooms, VIDA dynamically chooses graphical representations for the objects in each grid cell such that each grid cell contains a user-specified density (this density can be expressed according to various metrics, e.g., number of objects, number of vertices).

The current implementation supports all properties but 2b. We are currently implementing this functionality as well. Figure 1 shows an example user interaction.

CONCLUSIONS

We have presented a novel zoom method. This method, goal-directed zoom, allows users to directly control the choice of graphical representation of an object. We have implemented goal-directed zoom in the DataSplash database visualization environment. We are pursuing a number of related issues as detailed in [5].

ACKNOWLEDGMENTS

This research was sponsored by NSF under grants IRI-9400773 and IRI-9411334.

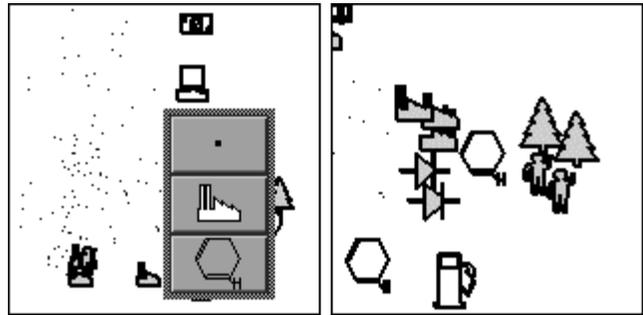


Figure 1: Visualization of Fortune 500 data. The visualization is an interactive scatterplot of selected Fortune 500 companies; on the X and Y axes are %profit growth and number of employees, respectively. Each company has three potential representations: (1) a dot; (2) an icon of the general category of industry to which the company belongs; and (3) an icon of the specific type of industry to which the company belongs.

Suppose the user is browsing the visualization and wants more information about companies with high profit growth and a small number of employees. When the user selects such a company in the visualization, a menu of the possible graphical representations of that company appears above it (Figure 1, left side). In this case, the factory icon in the middle of the menu indicates that the selected company belongs to the heavy industry sector and the benzene ring at the bottom of the menu indicates more specifically that the company is a chemical company.

Now suppose the user selects the most detailed representation (the benzene ring). The system automatically pans and zooms so that the selected company is centered in the screen and zooms until the benzene ring becomes visible (Figure 1, right side). From this vantage point, the user learns that, unexpectedly, high growth companies with few employees are not predominantly high-technology companies. Note that because VIDA chooses representations based on local density, the objects are presented at varying levels of graphical detail. In this case, some of the companies surrounding the benzene ring are represented as a general category, while others are represented as a specific industry. Slightly denser regions are represented as dots.

REFERENCES

1. Bederson, B., *et al.*, "A Zooming web browser," *Proc. SPIE Multimedia Computing and Networking*, 1996, Volume 2667, pp. 260-271.
2. Bier, E., *et al.*, "Toolglass and magic lenses: The see-through interface," *Proc. ACM SIGGRAPH 1993*, Anaheim, Calif., Aug. 1993, pp. 73-80.
3. Frank, A., and Timpf, S., "Multiple representations for cartographic objects in a multi-scale tree - an intelligent graphical zoom," *Computers & Graphics*, Nov.-Dec. 1994, 18(6):823-829.
4. Perlin, K., and Fox, D., "Pad: An alternative approach to the computer interface," *Proc. ACM SIGGRAPH 1993*, Anaheim, Calif., Aug. 1993, pp. 57-64.
5. Woodruff, A., and Stonebraker, M., "End-user interaction with clutter reduction techniques," Tech. Report, Univ. of California, Berkeley, Calif., Jan. 1997.