Visual Reasoning With Graphs

Yusuf Pisan
Qualitative Reasoning Group, The Institute for the Learning Sciences
Northwestern University, 1890 Maple Avenue, Evanston, IL 60201, USA
e-mail: y-pisan@nwu.edu

Abstract

Understanding diagrams is an important part of human cognition. Computer programs need to understand and reason using diagrams to communicate effectively with people. This paper explains how line graphs can be interpreted in a domain independent manner. We present a computer program called SKETCHY that reasons about physical phenomena visually by using line graphs. SKETCHY can interpret graphs to recover functional relationships, answer comparative analysis questions and generate qualitative descriptions using geometric models.

1 Introduction

People use diagrams to solve problems, to give explanations, to summarize information and to represent spatial relations. Computer programs that can represent, interpret and reason with diagrams will have a great impact on education, cognitive science and artificial intelligence. By making spatial relationships explicit, diagrams can reduce the amount of search and inference required[12]. Diagrams serve both as devices to aid in visualization of the situation[14] and as short-term fast access memory devices for holding information[7]. Although we do not yet have computer programs that can do general diagrammatic reasoning, diagrams have been successfully integrated with computer programs in a number of areas: to explain complex mechanical and dynamic systems[5,2,6,8], to solve geometry and physics problems[7,13], to constrain the search space in problem solving[9] and to understand the role of visual reasoning in problem solving[15].

Diagram understanding requires being able to identify objects, determine the relevant features for a particular problem and map the graphical features to the domain. A graph is a specialized form of diagrammatic representation that does not involve object recognition. Different graph formats emphasize different relationships between variables[10]. For instance, pie graphs are used to show relative percentages, bar graphs and step graphs to show relative amounts, scatter plots to show trends in data and line graphs to show continuous changes. In this paper, we only consider line graphs.

We are interested in interpreting line graphs because of their extensive use in thermodynamics, physics, economics and other fields. For example, an intelligent tutoring system that reasons with line graphs can use the graph as a shared medium for communication. The system can explain concepts by constructing graphical explanations. The user can also represent her understanding of the concepts graphically and the computer program can check the correctness of her understanding. The user and the computer can both make modifications to the line graph, therefore, providing an additional communication channel between the user and the program. Similar advantages would occur for engineering analysis and knowledge acquisition.

We present a computer program called SKETCHY that reasons about physical phenomena visually by using graphs. SKETCHY provides an interactive drawing environment, answers questions about graphs and interprets user modifications to the graph. SKETCHY has generated interpretations for all the graphs in a college level thermodynamics textbook[17], as well as a number of thermodynamic graphs from [18] and economics graphs from [1]. SKETCHY's interpretations are similar to graph descriptions found in textbooks. To our knowledge, SKETCHY is the first computer program that interprets graphs in a domain independent manner.

The next section (Section 2) gives examples from SKETCHY. Section 3 discusses the theory underlying SKETCHY. Section 4 explains the algorithms and design choices made in building SKETCHY. Finally, Section 5 discusses possible extensions to SKETCHY.

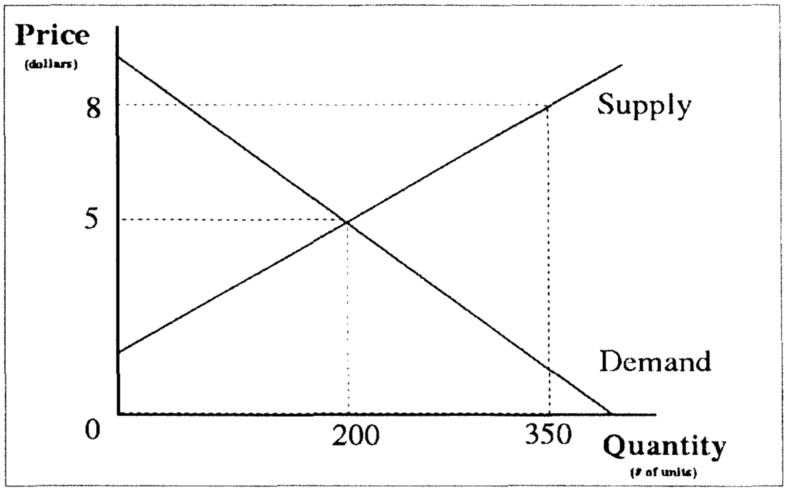

Figure 1: Supply and Demand graph for cassette tapes (Ekeleund & Tollison, p. 84)

2 Examples from SKETCHY

The input to SKETCHY is a graph drawn and labeled by the user. The supply-demand graph, shown in figure 1, is a typical graph found in many introductory economics textbooks. Supply-demand graphs are used to show three important relationships: how price effects the supply for the product, how price effects the demand, and the market price of the product.

To draw the graph shown in figure 1 using SKETCHY, the user starts with an empty graph consisting only of two axes. The user labels the axes as price and quantity and chooses scales for the axes. The supply and demand lines are drawn by choosing two points on the screen. SKETCHY requires the user to label the lines, so they can be referred in question and answers as well as in the graph interpretation. SKETCHY does not have any other knowledge about the labels price, quantity, supply and demand. The labels are only used to indicate the graph object that is being referred to. SKETCHY produces the following interpretation for the supply-demand graph.

```
For edge SUPPLY:
   QUANTITY and PRICE are directly
proportional.
For edge DEMAND:
   QUANTITY and PRICE are inversely
proportional.
```

Because the intersection point is not labeled, SKETCHY does not include that information in the graph description. This information can be obtained by asking questions to SKETCHY.

Q: At what point is SUPPLY equal to DEMAND?
```
The SUPPLY equals DEMAND at the point
when QUANTITY equals 200 and PRICE
equals 5.
```
Q: What is the PRICE for the SUPPLY line when the QUANTITY is 350?
```
For SUPPLY line, when QUANTITY is 350
the PRICE is 8.
```

Because supply-demand graph is relatively simple, SKETCHY's description is concise. A more complicated graph taken from a thermodynamics textbook is shown in figure 2. Since all substances exhibit the same qualitative behavior shown in the pressure-volume graph, understanding this graph is essential for solving many thermodynamics problems. The curved temperature lines (31°C, 40°C and 50°C) are drawn by choosing control points for the curve. Lines with discontinuities, such as 10°C and 20°C, are drawn by combining curved and straight lines. Lines 10°C to 50°C are grouped together and identified as a temperature contour by the user since their labels indicate a third dimension to the graph. The critical point is drawn by choosing a single point and labeling it. SKETCHY infers that the critical point is on line 31°C from its location. The pressure-volume graph also has three regions— liquid, liquid-and-vapour and vapour— corresponding to the state(s) the substance is in. The regions are labeled by points at their boundaries. The complete graph interpretation SKETCHY generates is:

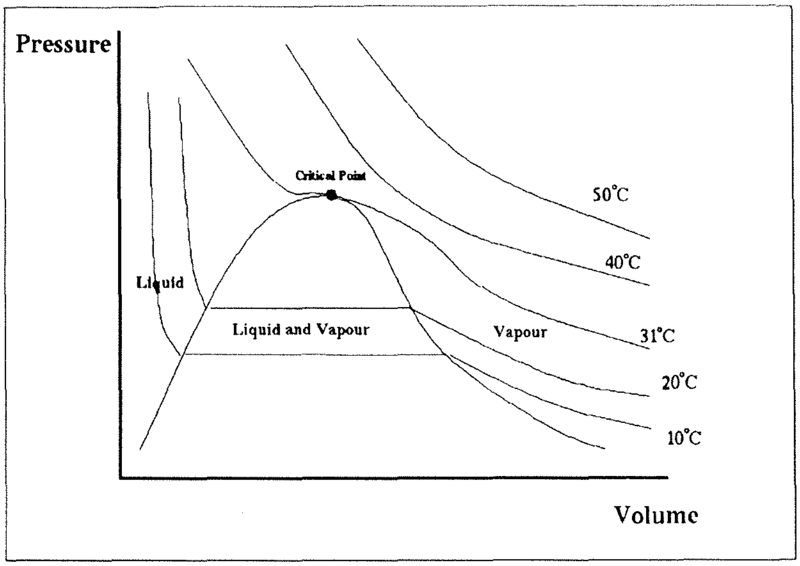

Figure 2: Compression of carbon dioxide (Whalley, p. 22)

```
For line 50-C:
   VOLUME and PRESSURE are inversely
proportional.
For line 40-C:
   VOLUME and PRESSURE are inversely
proportional.
For line 31-C:
   VOLUME and PRESSURE are inversely
proportional.
For line 20-C:
  The slope of 20-C has discontinuities
        ;associating discontinuities with regions
  Inside region LIQUID:
       VOLUME INCREASE and PRESSURE
DECREASE.
  Inside region LIQUID-AND-VAPOUR:
       VOLUME INCREASE and PRESSURE
CONSTANT.
  Inside region VAPOUR:
       VOLUME INCREASE and PRESSURE
DECREASE.
For line 10-C:
       The slope of 10-C has
discontinuities.
  Inside region LIQUID:
       VOLUME INCREASE and PRESSURE
DECREASE.
  Inside region LIQUID-AND-VAPOUR
       VOLUME INCREASE and PRESSURE
CONSTANT.
  Inside region VAPOUR:
       VOLUME INCREASE and PRESSURE
DECREASE.
CRITICAL-POINT is on lines (31-C)
CRITICAL-POINT is on regions (LIQUID
LIQUID-AND-VAPOUR VAPOUR)
For TEMPERATURE contour:
  As TEMPERATURE increases
      the slopes of TEMPERATURE lines
become more LINEAR.
        ;basis for Boyle's Law
  For a constant PRESSURE:
     As VOLUME increases TEMPERATURE
INCREASE.
  VOLUME and TEMPERATURE are directly
proportional.
  For a constant VOLUME:
    As PRESSURE increases TEMPERATURE
INCREASE.
      PRESSURE and TEMPERATURE are
directly proportional.
```

The examples described above are interpretations of static graphs, but graphs do not have to be static. Graph designers typically superimpose graphs or show sequence of graphs to describe changes in the situation. SKETCHY demonstrates that this natural form of comparative analysis[16] can be done via visual processes on a graph.

Analyzing engineering cycles is an important task in thermodynamics. The basic cycle for a steam power plant is the Rankine cycle, shown in figure 3. A common modification to the Rankine cycle is superheating the steam in the boiler to increase the efficiency of the cycle. The net work of the cycle before modification is represented by area 1-2-3-4-1 and after modification by area 1-2-3'-4'-1.

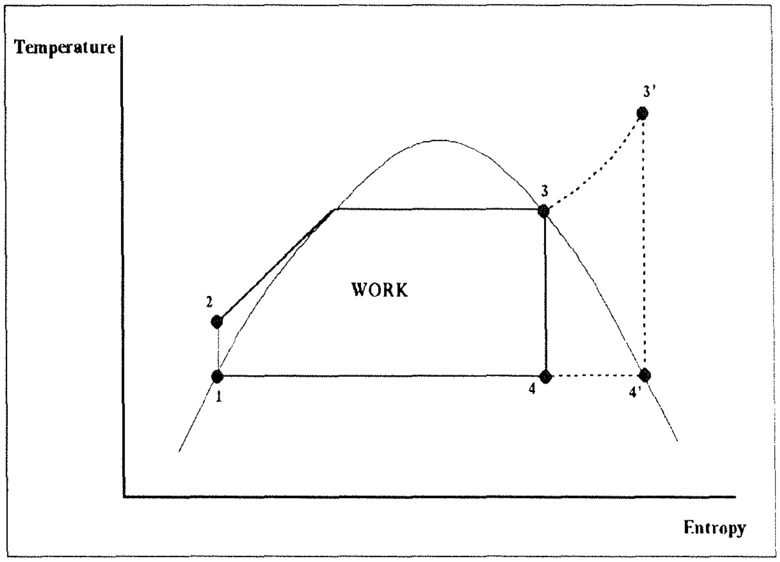

Figure 3: Effect of superheating on Rankine cycle

SKETCHY produces the following interpretation when 3 is moved to 3' and 4 to 4':

```
As a result of moving 3, 4:
For point 3:
  The ENTROPY of 3 INCREASE.
  The TEMPERATURE of 3 INCREASE.
For point 4:
  The ENTROPY of 4 INCREASE.
  The TEMPERATURE of 4 CONSTANT.
For region WORK:
  The area covered by WORK INCREASE.
```

3 The theory underlying SKETCHY

SKETCHY demonstrates that visual reasoning via graphs is a domain independent process. Most graphs are made up of three basic elements: points, lines and regions. Although the meanings of these elements could change from one domain to another, the interactions among graph elements does not change. Just as our perceptual mechanism can easily compute the graph properties, SKETCHY derives them using geometric models to perform visual reasoning.

Qualitative results provide valuable insights to understanding of a domain. Graphs represent qualitative information effectively by presenting it in a spatial format that our perceptual mechanism can process easily. SKETCHY qualitatively describes line slopes and curvatures. A straight line represents a numerical proportionality, and a vertical or a horizontal line represents that one of the properties is constant while the other is increasing. Smooth curves are described in terms of qualitative proportionalities[3,11]. Lines with discontinuities are described by dividing them up into qualitatively distinct regions. If the change in line slope corresponds to an intersection with another object, a relation between the intersection and the change of slope can be inferred (e.g. the existence of phase transition points).

Most of the information SKETCHY extracts from the graph can be easily represented using constructs of qualitative physics[3,11]. Directly proportional lines can be represented via qualitative proportionalities, points on the graph can be represented via correspondences. As a result, the output of SKETCHY can feed easily into qualitative reasoners.

Even when a verbal explanation is sufficient, textbooks use graphs to emphasize the verbal explanations. In these cases, although the graphs do not contain additional information, they create an additional representation in the form of an image. The image provides additional cues which makes remembering the information easier.

Our perceptual mechanisms are good at comparing the size and orientation of objects. Graphs exploit this ability to help perform comparative analysis. The modification made to increase the work output of the Rankine cycle (see Fig. 3), is an essential concept in understanding power plants. The graphical demonstration of the modification makes a comparative argument which is a lot more lucid than a numerical argument could have been. Graphical representations take advantage of peoples ability to follow qualitative and comparative arguments by using the graph as a shared medium for communication.

Similar lines in graphs are grouped as contours by our perceptual mechanism. SKETCHY treats contour lines, identified by the user, as a single element. Any changes of slope (or curvature) among contour lines is detected in the same way changes in line slopes are detected. Because contours are generally used to represent a third dimension, SKETCHY derives the relationship between the contour and each of the variables represented on the axes.

Graphs do not have to be drawn to scale when the representing only qualitative information. A common graph convention used in the absence of scales is to assume that moving from left to right on the horizontal axis and bottom to top on the vertical axis implies an increase in the variable represented by that axis. SKETCHY uses this assumption in making qualitative interpretations when numerical scales are not specified.

To determine what people notice in graph, we have examined graph descriptions in various textbooks and identified graph properties that the authors consider important. We claim that any computer program will need to detect all of these properties to effectively understand and reason using graphs. The common graph properties are:

1. Relative orientation of points, lines and regions
2. Intersection point of lines
3. Slopes of lines
4. Changes in line slopes
5. Minimum, maximum and inflection points of lines
6. Relative size of regions

SKETCHY implements these operations and based on our sample of 65 graphs (from [17],[18] and [1]), we believe SKETCHY provides strong support that these operators are sufficient and necessary for general graph interpretation.

4 Algorithms and design choices

The graph components of SKETCHY are labels, axes, points, lines, contours and regions. The labels on the graphs provide the vocabulary for SKETCHY's interpretations. The distinction between spatial relations represented in a diagram and the interpretation of the diagram has been made in earlier work in spatial reasoning[4]. Except when there is a single component of a particular type in the graph, unlabeled components can not be referred to either by the user or by SKETCHY. The labels on the axes are required to make meaningful interpretations since axis labels represent the physical properties which provide the framework for the graph. When the axes have scales, SKETCHY includes numerical information in addition to qualitative interpretation of the graph.

Points represent discrete state information. Points can be connected to lines (whether they are on the line or not), thus move the lines when they are moved, preserving the orientation relationship between them. They also serve as boundary markers for regions and modify the shape of the region when moved. Moving a point changes the values of physical properties represented. SKETCHY interprets the modification by comparing the property values for the old and the new location.

Internally, lines are represented by an ordered list of segments (for convenience, the user can enter equations which are converted to line segments). Since segments can be of any length, using segments provides arbitrary precision for line curvatures without needing to manipulate complex equations. Like points, lines also serve as boundary markers for regions. Moving or extending the line changes the shape and the size of the region.

When the shape of the region changes, SKETCHY compares the old area to the new area in the interpretation. If the change in the region's shape results in including (or excluding) an object, this information is also included in the region. Any changes in the line slope while traversing a region is included in the graph interpretation as significant. In the pressure-volume graph, noticing that the pressure stays constant inside the liquid-vapour region and decreases outside the region is an essential observation for understanding phase changes.

SKETCHY compares the slopes of the contour lines to find any trends of changes. Lines with discontinuities, such as temperature lines 10°C and 20°C in pressure-volume graph, are approximated by curves to be able to make qualitative statements about all the contour lines.

A common intersection point of contour lines which indicates a convergence to a specific value is detected and included in SKETCHY's graph interpretation. When all contour lines exhibit similar qualitative changes in their slopes, the contour is described in terms of these qualitative changes.

Understanding a graph necessitates being able to answer questions about it. By using built in templates, SKETCHY answers questions similar to the ones in the GRE[1] and the SAT[2]. SKETCHY answers questions that involve reading values of the graph (e.g. What is the pressure when the volume is 3 and the temperature is 21?), comparing slopes, areas or locations of objects with respect to each other and describing relations of objects to each other (e.g. Is the critical point on line 20°C?).

It is especially interesting that SKETCHY's graph interpretations are so similar to explanations found in textbooks given its lack of any other understanding of the domain. The interpretations produced using object labels are meaningful and insightful for the domain. The major difference is that SKETCHY's interpretations lack the briefness of experts' explanations because SKETCHY does not differentiate between relevant and irrelevant information.

5 Discussion

Understanding graphs is an important part of human cognition. We have shown how visual reasoning can be used to interpret graphs in a domain independent way using geometric properties. The labels in graphs provide the vocabulary, the points and lines provide the relationships for interpreting graphs. SKETCHY produces output that can be used by qualitative reasoners and other problem solvers. The ability of SKETCHY to interpret graphs suggests that these ideas are robust.

There are a number of avenues we are exploring for extending SKETCHY. One possible extension is to extend SKETCHY's drawing environment and interpretation mechanism to other kinds of graphs, such as scatter plots bar graphs, and pie graphs. By understanding the graph properties exploited by these other graphic representations, we hope to build a model for general graph understanding.

Another avenue we are pursuing is incorporating SKETCHY in an intelligent tutoring system where SKETCHY provides an additional communication medium between the user and the computer. The user and the computer program both can make modifications to the graph and discuss the graph relating the concepts represented with other ideas in the domain. The ideas implemented in SKETCHY can be used to construct graphs that the user will be able to interpret and reach the intended conclusions.

6 Acknowledgments

The research reported was supported by Office of Naval Research. I owe many thanks to my advisor, Ken Forbus, for valuable discussions, comments and encouragement. My gratitude also goes to Meryl McQueen who patiently proofread the original draft and to John Everett, Ron Ferguson and Keith Law who read subsequent drafts.

References

1. R. B. Ekelund and R. D. Tollison. *Economics*. Little, Brown, Boston, 1986.
2. B. Faltings. Qualitative models in conceptual design: a case study. In *1st International Conference on Artificial Intelligence in Design*, 1991.
3. K. D. Forbus. Qualitative process theory. Artificial Intelligence, 24:85-168, 1984.
4. K. D. Forbus. Qualitative reasoning about space and motion. In D. Gentner and A. L. Stevens, editors, *Mental models*, Erlbaum, Hillsdale, NJ, 1983.
5. K. D. Forbus, P. Nielsen, and B. Faltings. Qualitative spatial reasoning: the clock project. *Artificial Intelligence*, 51:417-471, 1991.
6. L. Joskowicz and E. P. Sacks. Computational kinematics. *Artificial Intelligence*, 51:381-416, 1991.
7. G. S. Novak Jr. and W. C. Bulko. Uses of diagrams in solving physics problems. In *AAAI Symposium on Reasoning with Diagramatic Representations*, Stanford, CA, March 1992.
8. H. Kim. *Qualitative reasoning about fluids and mechanics*. PhD thesis, 1993.
9. K. R. Koedinger and J. R. Anderson. Abstract planning and perceptual chunks: elements of expertise in geometry. *Cognitive Science*, 14:511-550, 1990.

[1]Graduate Record Examination
[2]Scholastic Aptitude Test

10. S. M. Kossylyn. *Elements of Graph Design*. Freeman and Company, New York, NY, 1994.
11. B. J. Kuipers. Qualitative simulation. *Artificial Intelligence*, 29:289- 388, 1986.
12. J. Larkin and H. Simon. Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11:65-69, 1987.
13. T. F. McDougal and K. J. Hammond. Representing and using procedural knowledge to build geometry proofs. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, 1993.
14. N. H. Narayanan and B. Chandrasekaran. A computer model of diagrammatic reasoning. In *AAAI Symposium on Reasoning with Diagrammatic Representations*, Stanford, CA, 1992.
15. S. Tessler, Y. Iwasaki, and K. Law. Qualitative structural analysis using diagrammatic reasoning. In *The Seventh International Workshop on Qualitative Reasoning about Physical Systems*, 1993.
16. D. S. Weld. *Theories of Comparative Analysis*. MIT Press, Cambridge, MA, 1990.
17. P. B. Whalley. *Basic Engineering Thermodynamics*. Oxford University Press, New York, 1992.
18. G. J. Van Wylen and R. E. Sonntag. *Fundamentals of Classical Thermodynamics*. Wiley, New York, 3rd edition, 1985.