

Visualization Methods for Data Analysis and Planning in Medical Applications

Robert Kosara Silvia Miksch

Institute of Software Technology and Interactive Systems,

Vienna University of Technology

Favoritenstraße 9–11/E188

A-1040 Vienna, Austria

{[rkosara](mailto:rkosara@asgaard.tuwien.ac.at), [silvia](mailto:silvia@asgaard.tuwien.ac.at)}@asgaard.tuwien.ac.at

<http://www.asgaard.tuwien.ac.at/>

May 8, 2002

Keywords: User-Computer Interface [L01.700.568.810.910], Decision Making, Computer-Assisted [L01.700.508.100]

Abstract

Time plays an important role in medicine, both the past and the future. The medical history of a patient represents the past, which needs to be understood by the physician to make the right decisions. The past contains two different kinds of information: measured data (such as blood pressure) and incidents (such as seizures). Planning therapies, on the other hand, requires looking into the future to a certain extent.

Visual representations exist for both the past and the future, and they are very useful for getting a better understanding of data or a plan. This paper surveys visualization techniques for both data analysis and planning, and compares them based on a number of criteria.

1 Introduction

Analyzing data and planning therapy steps are two of the central parts of practical work in medicine. Both require an overview of the information: the recorded data is the basis for further planning, and therefore needs to be understood, but on a more abstract level than the single data points; and the plan must be understood so that it can be followed.

Both tasks can be and are supported with graphical depictions of the data – it is much easier to see a trend on a chart than when reading a row of numbers. But for complex information, more sophisticated techniques are necessary to create effective visual representations that ease understanding of the data.

Information Visualization (InfoVis) [22] deals with the effective display of abstract information (in contrast to directly understandable information like MRI (magnetic resonance imaging) or CT (computer tomography) data). Other examples for typical InfoVis data are file systems, data bases (medical and other), and computer networks.

This paper surveys visualization techniques for three different applications: recording measured data (blood pressure, oxygen levels, etc.), recording incidents (seizures, pain attacks, etc.), and planning actions (therapeutic steps). Some techniques can be used for more than one application, and are therefore discussed in the different parts of the paper, with different requirements.

Such surveys have been done already [19], but with different requirements and a different focus.

This paper deals with the three different applications of information visualization for time in three sections. Section 2 discusses methods for visualizing measurements, Section 3 presents techniques for displaying incidents, and Section 4 deals with methods that depict planned actions.

2 Measured Data

Data that is measured either continuously (such as the parameters of intensive care patients) or when that information is needed (when the patient goes to see the doctor) later needs to be analyzed. The focus here is on the quantitative data, not so much on periods with certain symptoms – these are treated in Section 3.

2.1 Requirements

We believe the following features to be needed for visualizing both types of data. Depending on the concrete application, they can differ however.

Intuitiveness. Depending on the intended users of the visualization, it can be necessary that it be as intuitive as possible, so that even a user who is completely unfamiliar with the visualization (but familiar with the application field) can immediately understand it. It is usually desirable to use intuitive visualization methods so that also casual users of the system do not need help.

Focus+Context Time/Data. An important problem in InfoVis is to provide methods that make it possible to see details of very large data (many data points, e.g., after a patient has been monitored for weeks) without losing the context of the data, i.e., where the zoomed area is. Some of the standard methods from visualization, like distortion-oriented methods [14] cannot be used for measured data, because they change the way developments are perceived (see next point).

In the context of this paper, two F+C dimensions are distinguished: F+C can be provided for time, i.e., data from long time ranges can be displayed effectively, and many data dimensions can be visible for the same point in time, so the user might want to choose which ones to see.

Combination of Values Displaying only one value in one graph is easy, but combining several in a way that their interaction can be understood is a different problem.

Seeing Developments. For measured data, it must be easy to see how values change – if a value increases or decreases, or if it stays (roughly) the same. Judging the speed of change is also important, so the visualization must not create a wrong impression by using several different zoom factors in one display, for example.

Finding Patterns. Sometimes, distinctive patterns in the data (e.g., two weeks of fever, then a period without symptoms, and then one week of fever and other symptoms) give a hint as to what the underlying problem is. Such patterns should be easy to find, and to recognize in a visualization.

Discovering Intervals. Determining the cause of a problem can also involve finding a correlation of

symptoms with other events. For this, it is necessary to be able to find the intervals in which problems occur (i.e., every monday, or every 16 hours). This is quite hard to do with tabular-style layouts, such as calendars, especially when the time span between neighboring lines cannot be changed. This requirement is independent of the previous one, because intervals can also be found when no really repeating patterns are present, but there are times where incidents accumulate, and times without (or with few) incidents.

2.2 Methods

This section presents methods that have been proposed for visualizing time-series data. How well they fulfill the requirements is discussed in Section 2.3.

2.2.1 A Simple Chart

Charts (Figure 1a) are drawn all the time for many different applications. And while simple and ubiquitous, charts give a very good idea of the development of values. In this discussion, they only serve as a reference, however, for the features of more elaborate visualizations to become visible.

2.2.2 Graphical Patient Record Summary

The graphical patient record summary [17] (Figure 1b) is a complete overview of all known patient parameters on one sheet of paper. The single charts for each value are drawn in such a way that the values of the current week are drawn with the highest resolution, with values of the last year and the year(s) before getting much less room. This is a similar idea to distortion-oriented focus+context techniques in visualization, like fisheye views [9] or hyperbolic space [13]. Because there are only three magnification levels, and they are clearly separated, they do not interfere with the understanding of developments.

All charts are normalized so that normal, elevated and critically elevated (or reduced) values can be spotted immediately. Thumbnail sized x-ray and other images are also included to give a first impression. A list of recent notes in short keywords is also included to provide a more complete picture.

This summary does not replace a more detailed patient record, of course, but it serves as an overview,

so that the physician can then go into the details where they are needed.

2.2.3 VIE-VISU

VIE-VISU [10] (Figure 1c) uses what some authors call *metaphor graphics* [6], and what is generally known as a *glyph* in information visualization [5]: an object whose features change with data, and which has some kind of relation to the values it represents (e.g., restrictions of the values relative to each other can be seen). These relations can be exploited to make the glyph easier to understand and to remember which parts represent which values. In this case, the glyph consists of three parts (top, middle, and bottom) that represent circulation, respiration, and fluid balance, respectively, and that are subdivided further to show the single values.

2.2.4 Spirals

For discovering intervals, spirals [3, 23] (Figure 1d) are best suited. They make it possible to display data close to each other that is close in the temporal dimension, and close in terms of an interval. When coupled with a quick way of changing the displayed interval, it is very easy to see when data points line up, and therefore an interval was found.

2.3 Discussion

Depending on the concrete requirements, the different methods have their strengths and weaknesses. The simple chart is easy to draw even by hand and can be understood by practically everybody; and at the same time, it literally gives a much better picture of data than a list of numbers. It does not work very well for longer periods of time, however, and does not offer any support when possibly periodic data is to be investigated.

The graphical patient status summary was designed for longer periods of time, with a focus on the more recent past (i.e., the last week), and therefore includes a certain F+C functionality. It was also made for printing many similar diagrams onto one sheet of paper.

The glyphs in VIE-VISU are shown in lines, where each line represents a certain time interval. Similar

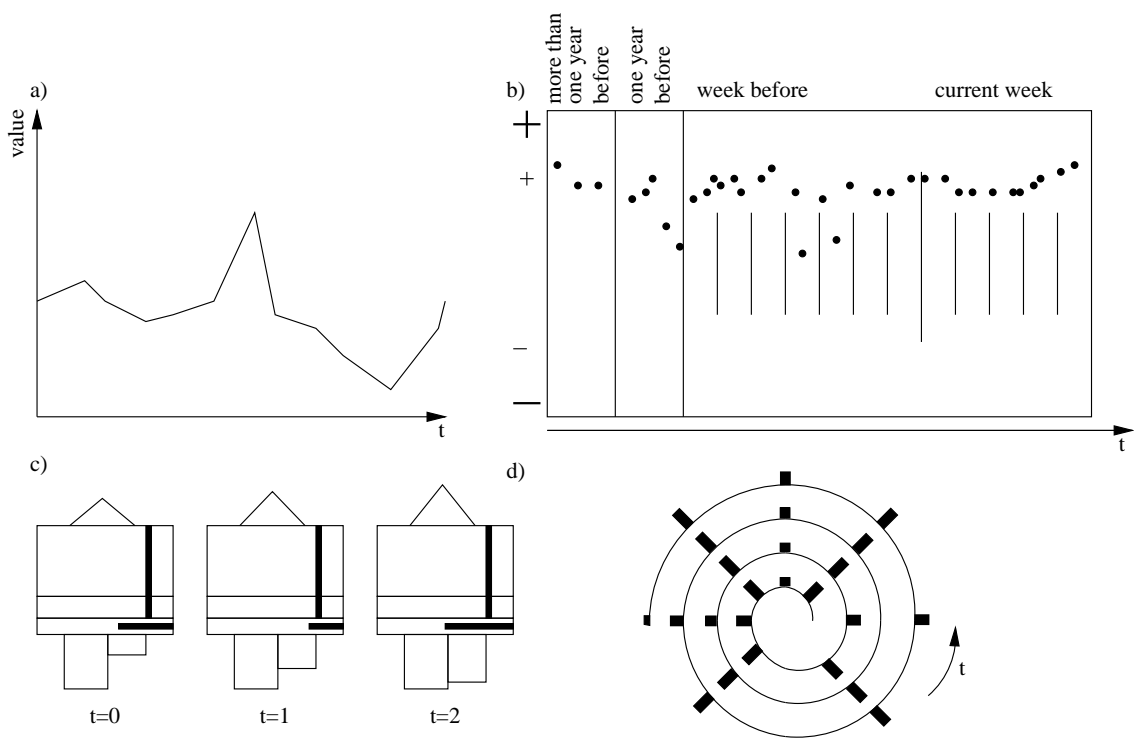


Figure 1: Methods for recording data over time. **a)** a simple chart; **b)** the graphical patient summary [17], Section 2.2.2; **c)** VIE-VISU [10], Section 2.2.3; **d)** Spirals [3, 23], Section 2.2.4.

Method	Intuitiveness	Focus+Context Time	Focus+Context Data	Combining Values	Seeing	Developments	Finding	Patterns	Discovering	Intervals
Chart	•					•	•			
Graphical Summary	•	•		•	•					
VIE-VISU		•		•	•					
Spirals	•	•								•

Table 1: Comparison of Visualization Methods for Measured Data.

values lead to similar shapes, which makes it possible to see developments when many glyphs are drawn next to each other. It is principally also conceivable that a large number of such glyphs is used to find intervals – but that would require a display that can contain an arbitrary number of glyphs per line, so that similar shapes lining up in columns can be seen. But in its current design, VIE-VISU does not allow the user to do this.

Combining several data channels in a chart can be done, but there are limits to how many values can be drawn over each other before the diagram becomes unreadable. Combining different data in a spiral display does not seem to be possible, however. VIE-VISU contains more than one value by design, but its metaphorical power can only be exploited when visualizing ventilation data.

3 Incidents and Symptoms

Apart from quantitative data, incidents (like seizures, pain attacks, etc.) and periods with certain symptoms (pain, fever) are also important information for finding causes and planning therapies. Visualizing information about incidents is less common than data visualization (with its ubiquitous chart), but by no means less interesting.

3.1 Requirements

The requirements for the methods discussed in this section are identical to the ones for recording data (Section 2.1), except for “Seeing Developments”, which only plays a role for quantitative data.

3.2 Methods

Methods for visualizing incidents are less common, so some of the presented techniques here are slight abuses of the original ideas.

3.2.1 Time Lines and LifeLines

A very intuitive way of drawing incidents is by simply drawing a line for the time span the incident took (Figure 2a). This idea is called a time line [20], and a further development of the idea is the LifeLine [16]. Lines in different distances from the time scale stand for different types of incidents, and groups of lines belong to so-called facets that can be collapsed, so that less information is visible, but has more screen space for it.

One of the extensions introduced in [16] are facets. Facets are vertical segments that group similar events. They can be opened and closed to provide the user with information on different aspects of the same structures without cluttering the display with too much information. They are usually used to show different aspects (views) of the same information.

3.2.2 Graphical Patient Summary

The graphical patient summary (Section 2.2.2) also records incidents, and therefore also needs to be taken into consideration here.

3.2.3 Shinkansen Timetables

The timetables for the Japanese Shinkansen train [21] depict many different trains and stations in one diagram. A horizontal time axis and a vertical axis enumerating all stations lying along a particular track form a grid of time-location data points. For a single train, all such points it passes through are connected

with a line (Figure 2b). This way, one can not only see which direction a train goes, but also, how fast it is (compared to others) and if it stops at a station (if it does, there is a short horizontal segment at the vertical position of the station).

3.2.4 Circles

Concentric circles [8] (Figure 2c) are another useful time visualization method, which draws concentric circles for different kinds of information (location, symptoms), and draws many such diagrams for different time periods.

Concentric circles presume that a certain interval exists: in the proposed uses, a circle usually depicts one day or one week, so that repeated structures at certain fixed intervals can be found, but not arbitrary ones.

3.2.5 Time Cube

Another extension of time lines is the display of spatiotemporal information in a cube [11] (Figure 2d). In this pseudo-3D diagram, two spatial and one temporal dimensions are drawn so that an event can be located in time as well as in space.

A line represents a person or object that moves back and forth between different locations or modes denoted by cylinders. This diagram is used to show typical behavior patterns of tourists in ski resorts, for example. When incident types are used on the base plane instead of locations, this diagram can of course also show events or periods of symptoms.

3.2.6 Spirals

Spirals (Section 2.2.4) can also be used for incidents, and therefore need to be considered here.

3.3 Discussion

LifeLines are a very intuitive way of thinking about incidents or periods, and allow the user to see patterns.

Shinkansen Timetables appear to be useful in cases with many different incidents to be tracked, because

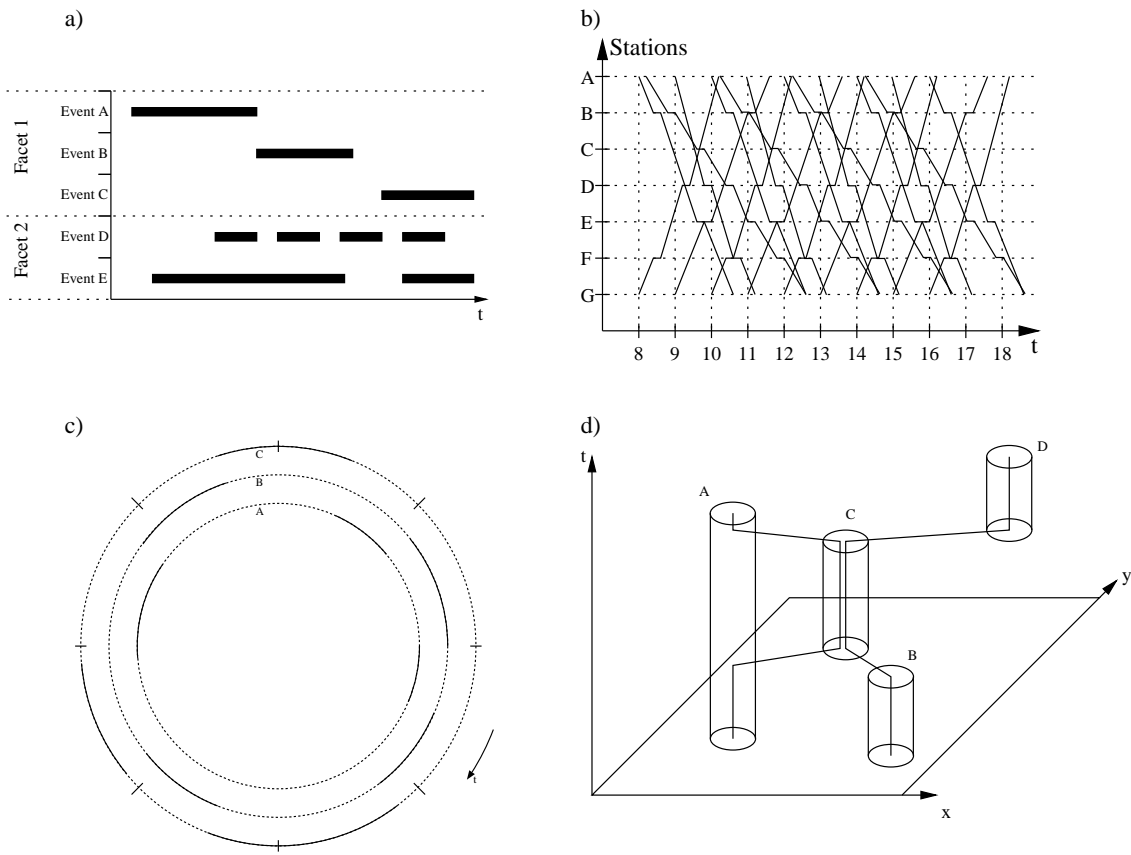


Figure 2: Methods for recording incidents over time. **a)** LifeLines [16], Section 3.2.1; **b)** Shinkansen Timetables [21], Section 3.2.3; **c)** Concentric Circles [8], Section 3.2.4; **d)** Time Cube [11], Section 3.2.5.

Method	Intuitiveness	Focus+Context Time	Focus+Context Data	Seeing	Developments	Finding	Patterns	Discovering	Intervals
LifeLines	•		•	•		•			
Graphical Summary	•	•							
Shinkansen							•		
Circles	•								(•)
Time Cube				•		•			
Spirals	•	•						•	

Table 2: Comparison of Visualization Methods for Incidents

they do not need much space, and therefore can accommodate a larger number of objects in the same screen space.

The time cube is an interesting idea for tracking the change in incidents and finding interesting patterns.

Circles are quite intuitive when used for “natural” periods such as days, weeks, or months. Depending on the implementation, circles could also be used for dynamically changing the interval, but would have a discontinuity where the values “jump” from one circle to the next. So they are normally used not to analyze intervals, but to present them as results.

4 Planning the Future

The most complex time visualization is needed for planning. This is due to the fact that planning, by definition, involves uncertainties and needs to be robust enough to also work when unanticipated things happen.

4.1 Requirements

Because this study was done in the scope of the Asgaard Project, being able to depict the language constructs of Asbru (the language used for specifying therapy plans) was the main requirement. In terms of time, the key part here is the Time Annotation (TA), which consists of seven parts: the reference point, which all other times are defined relative to; the earliest starting shift (ESS), which is the earliest point in time the action can start; the latest starting shift (LSS), which is the latest possible starting point; the earliest finishing shift (EFS), which is the earliest time the action can end; the latest finishing shift (LFS), which is the latest possible end; the minimum duration (MinDu); and the maximum duration (MaxDu). MinDu and MaxDu are not necessarily equal to the difference between EFS and LSS, or ESS and LFS, respectively, but are bounded by these time ranges.

These features of the Asbru Time Annotation are reflected in the following requirements:

Allen's Relations. Any visual representation of time must of course be able to visually represent all possible relationships between intervals. There are 13 such relations [1]: A before B (A ends before B starts), A meets B (B starts at the same instant that A ends), A overlaps B (B starts before A ends), A starts B (A and B start at the same time), A contains B (A is shorter than B and starts after B starts), A finishes B (B ends at the same time A ends). For each of these six relations, there also is a reverse one plus the commutative relation A equals B.

Temporal Uncertainty. The representation must be able to deal with temporal uncertainty as described in the previous section.

Undefined/Unknown Parts. If parts of a time annotation are unknown (denoted in Asbru as $_$), they must be depicted in a way to make this easily recognizable without at the same time being too dominant.

Resolution. In addition to the temporal uncertainty of actions (see previous section), it should also be possible to see to what precision a point in time has been specified relative to the scale it is currently being viewed on.

Hierarchical Decomposition. It must also be possible to communicate the fact that plans are made up of sub-plans. This hierarchical decomposition is important not only to structure a plan and to make

parts reusable, but also to be able to select the amount of information visible by showing more or fewer levels at the same time.

Different Dimensions. Different kinds of information about the same object should be visible at the same time. One of the key ideas for this are facets.

4.2 Methods

4.2.1 Time Lines, LifeLines

LifeLines (Section 3.2.1) can also be used for planning, and indeed are in different ways. One is a visualization of a vaccination plan [2], another example are Gantt charts (see below).

4.2.2 Gantt and PERT Charts

A popular type of diagram in project management are PERT charts (Project Evaluation Review Technique — usually used for CPM (Critical Path Method)). While PERT charts show dependencies of tasks (and hence their “topology”), they do not show the temporal extents of tasks directly. (Figure 3a).

Another common diagram in this context are Gantt charts (Figure 3b), which are heavily based on time lines, but have the additional ability of displaying hierarchies. Gantt charts are used in many project management packages, where it is useful and necessary to be able to choose between seeing a tasks parts or just the higher levels of a project. Gantt charts usually do not have the ability of showing different facets of the same task (i.e., not like LifeLines).

4.2.3 Sets of Possible Occurrences (SOPOs)

An interesting way of looking at time are sets of possible occurrences (SOPOs) [18]. This diagram uses two time axes that represent the begin and end times of an interval, respectively (Figure 3c). Any point in this diagram represents a whole interval, specified by its start (x coordinate) and end time (y coordinate). The area a SOPO covers (black in the figure) contains all intervals that fit the specification given by means of an earliest start, latest start, earliest end, latest end, and minimum and maximum durations.

4.2.4 AsbruView's Time Annotation Glyph

This diagram is a part of AsbruView [12], which is a user interface for the plan representation language Asbru. For temporal information, AsbruView uses glyphs (Section 2.2.3). The one described here is an extension of LifeLines and was specifically designed to represent Asbru's time annotation.

It is based partly on a simple metaphor (Figure 3d). Four vertical "pillars" along the time axis represent the earliest and latest starting and ending times. On top of these vertical bars lies a bar that is as long as the maximum duration. On top of the MaxDu bar, supported by two diamonds (which lie above the LSS and EFS supports), lies another bar representing the minimum duration. Undefined parts of a time annotation are displayed in gray, and in addition the diamonds supporting the MinDu bar become "rolls".

It is possible to understand a few simple constraints based on this metaphor. One is that the minimum duration can never be shorter than the difference between LSS and EFS – if it was, the MinDu bar would fall down between its supports. If LSS or EFS are undefined, the corresponding diamonds become rolls, which means that the implicitly defined EFS moves if the MinDu changes, for example. On the right side of Figure 3d, the depiction of different temporal scales can be seen. If the current scale is coarser than the precision to which the time points were defined, a circle appears (similar to the notion of an open interval in mathematics, where the end point of the interval "from 1 to, but not including, 2" cannot be depicted directly). On a smaller scale, a zigzag line appears that covers the area of "additional uncertainty" due to the lower resolution of the actual point in time.

4.2.5 Temporal Objects

Temporal objects [7] are similar in many respects to AsbruView (but were developed independently). It does, however, lack a way of displaying undefined parts of a time annotation. A way of displaying the granularity to which the time points were specified is missing as well.

But it is probably easier to understand as an extension of time lines than the AsbruView glyph (Figure 3e).

4.2.6 Paint Strips

Another extension to LifeLines are paint strips [4], which consist of a part that is already painted (the minimum duration, latest start and earliest end) and a part that is not yet painted, but that paint rollers can add when they move over the diagram (Figure 3f). These define the maximum duration and the earliest start or latest end, depending on which end of the painted strip they are attached to. Thus it is not possible to constrain the durations independently of the differences between start and end points.

4.3 Discussion

Even if only a part of the functionality of the time annotation is used in each temporal pattern or action specification, none of its capabilities can be ignored. Lifelines clearly do not satisfy the requirements for displaying time annotations. They cannot represent uncertainty or even undefined values. As Rit points out [18], this is simply due to the fact that they are one-dimensional, so any information exceeding one-dimensional time results in an ambiguous diagram. For recorded patient information, however, they are very useful. This mostly includes information about events, rather than recorded device readings (but a combination of both is easily imaginable). SOPOs were designed for the easy graphical propagation of temporal constraints, not for making a complex notion of time easy to understand. Specifically, parallel plans and hierarchical decomposition are very hard to depict and to work with (if plans from several levels are drawn into the same diagram, their relationship is not immediately visible; parallel plans cover the same area in the diagram). A notion of undefined parts is missing in the original design.

We extended SOPOs so that they met almost all the requirements described in Section 4.1 [15]. In a usability study we performed with these extended SOPOs, one of the results was that this type of diagram is hard to understand for people outside of computer science (and even many within).

AsbruView's time annotation glyph was also specifically designed to meet the requirements, so it does satisfy all of them. The simple metaphor it is based on has proven to be very helpful, without it limiting the representative power of the glyph (which can happen if a glyph is too tightly coupled to a metaphor: then, the constraints of the metaphor become constraints of the glyph. This is of course very undesirable). Other methods like Gantt and PERT charts are quite common in many areas, including medicine. Gantt charts

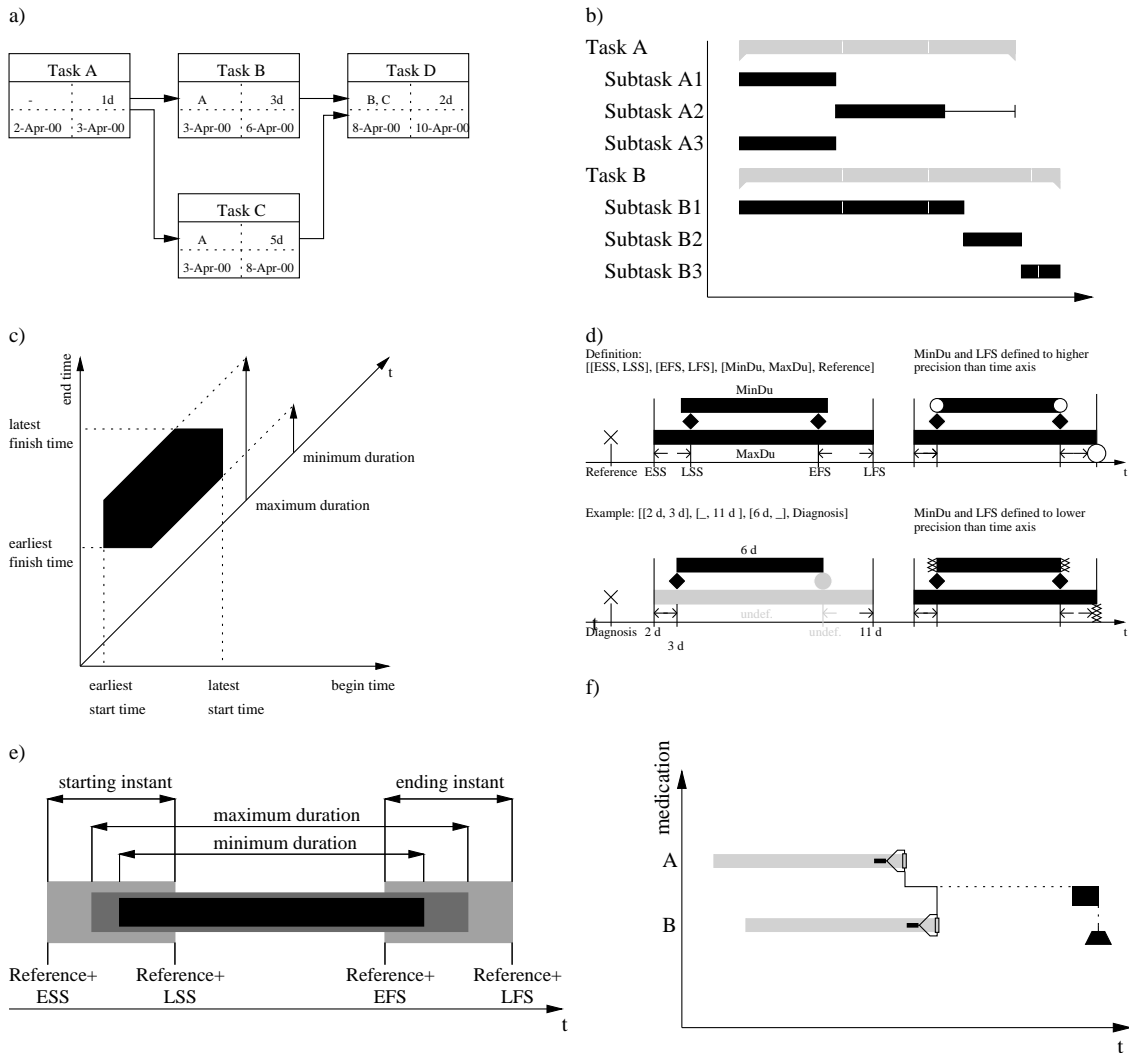


Figure 3: Visualization techniques for planning. **a)** PERT chart, Section 4.2.2; **b)** Gantt chart, Section 4.2.2; **c)** SOPOs [18], Section 4.2.3; **d)** AsbruView [12], Section 4.2.4; **e)** Temporal Objects [7], Section 4.2.5; **f)** Paint Strips [4], Section 4.2.6.

Method	Allen's Relations	Temporal	Uncertainty	Undefined Parts	Resolution	Hierarchical	Decomposition	Different	Dimensions
TimeLines	•								•
Gantt	•						•		
PERT							•		
SOPOs	•	•							
Extended SOPOs	•	•		•	•		•		
Temporal Objects	•	•					•		•
Paint Strips	•	•					•		•
Time Annotation Glyphs	•	•		•	•		•		•

Table 3: Comparison of Time Visualization Methods

are similar to LifeLines, but can display hierarchies, and it is easy to imagine how to add facets (similar to LifeLines) to a Gantt chart display.

PERT charts, on the other hand, do not show any precise temporal information, only the order in which actions should occur. This can be useful, of course, especially in early stages of planning. But the approach is quite limited, especially when dealing with such complex things as treatment protocols. This discussion is summarized in Table 3.

5 Conclusions and Future Work

Visualization supports complex tasks such as the analysis of patient records and treatment planning. More sophisticated techniques provide the user with information that otherwise would not be easily accessible.

A few open problems remain. One is the problem of how to represent cyclical as well as other types of events when planning. This is made complicated through the fact that cyclical events might be specified to be performed until a certain state is reached. In such a case, the temporal extent of the action is not known

(not even possible intervals), and a way of depicting this would have to be found.

The combination of data and plans after execution is also an interesting topic – this is easier than displaying plans in the future (no uncertainties), but should be visually similar to the other plan views, so that the user does not have to learn too many different visualizations.

For recorded data, a method would be useful that fulfills all the requirements but still does not create a display that is too overloaded with information.

6 Acknowledgements

The Asgaard Project is supported by “Fonds zur Förderung der wissenschaftlichen Forschung” (Austrian Science Fund), grant P12797-INF.

References

- [1] James Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
- [2] C. A. Brandt, S. J. Frawley, S. M. Powsner, R. N. Shiffman, and P. L. Miller. Visualizing the logic of a clinical guideline: A case study in childhood immunization. *Methods of Information in Medicine*, 36:179–83, 1997.
- [3] John V. Carlis and Joseph A. Konstan. Interactive visualization of serial periodic data. In *Proceedings of the 11th annual ACM symposium on User interface software and technology (UIST'98)*, pages 29–38. ACM Press, 1998.
- [4] Luca Chittaro and Carlo Combi. Visual definition of temporal clinical abstractions: A user interface based on novel metaphors. In Silvana Quaglini, Pedro Barahona, and Steen Andreassen, editors, 8th *Conference on Artificial Intelligence in Medicine in Europe (AIME 2001)*, number 2101 in Lecture Notes in Artificial Intelligence, pages 227–230. Springer, 2001.

- [5] Mei C. Chuah and Stephen G. Eick. Glyphs for software visualization. In *5th International Workshop on Program Comprehension (IWPC '97) Proceedings*, pages 183–191. IEEE Computer Society Press, Dearborn, Michigan, May 1997.
- [6] William G. Cole and James G. Stewart. Metaphor graphics to support integrated decision making with respiratory data. *International Journal of Clinical Monitoring and Computing*, 10:91–100, 1993.
- [7] Carlo Combi, Luisa Portoni, and Francesco Pinciroli. Visualizing temporal clinical data on the www. In Werner Horn, Yuval Shahr, Gerger Lindberg, Steen Andreassen, and Jeremy Wyatt, editors, *Proceedings of the Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making (AIMDM '99)*, pages 301–311. Springer, jun 1999.
- [8] Chaouki Daassi, Marlon Dumas, Marie-Christine Fauvet, Laurence Nigay, and Pierre-Claude Scholl. Visual exploration of temporal object databases. In *Proceedings of 16^{ièmes} Journées Bases de Données Avancées (BDA, French Conference on Databases)*, 2000.
- [9] George W. Furnas. Generalized fisheye views. In *Proceedings of the ACM Conference on Human Factors in Computer Systems*, SIGCHI Bulletin, pages 16–23. ACM Press, 1986.
- [10] Werner Horn, Christian Popow, and Lukas Unterasinger. Support for fast comprehension of icu data: Visualization using metaphor graphics. *Methods of Information in Medicine*, 40:421–424, 2001.
- [11] Manu Konchady, Ray D’Amore, and Gary Valley. A web based visualization for documents. In *Proceedings of the 1999 workshop on New paradigms in information visualization and manipulation*, pages 13–19. ACM Press, 1998.
- [12] Robert Kosara and Silvia Miksch. Metaphors of movement: a visualization and user interface for time-oriented, skeletal plans. *Artificial Intelligence in Medicine*, 22(2):111–131, May 2001.
- [13] John Lamping, Ramana Rao, and Peter Pirolli. A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings Conference on Human Factors in Computing Systems (CHI'95)*. ACM, 1995.

- [14] Y. K. Leung and M. D. Apperley. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction*, 1(2):126–160, June 1994.
- [15] Peter Messner. Time shapes - a visualization for temporal uncertainty in planning. Master’s thesis, Vienna University of Technology, 2000.
- [16] Catherine Plaisant, Brett Milash, Anne Rose, Seth Widoff, and Ben Shneiderman. LifeLines: Visualizing personal histories. In *Proceedings of ACM CHI 96 Conference on Human Factors in Computing Systems*, volume 1 of *PAPERS: Interactive Information Retrieval*, pages 221–227, 1996.
- [17] Seth M. Powsner and Edward R. Tufte. Graphical summary of patient status. *The Lancet*, 344:386–389, August 6, 1994.
- [18] Jean-François Rit. Propagating temporal constraints for scheduling. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 383–388. Morgan Kaufman Publishers, Inc., August 11-15 1986.
- [19] Sônia Fernandes Silva and Tiziana Catarci. Visualization of linear time-oriented data: A survey. In *First International Conference on Web Information Systems Engineering (WISE’00)*. IEEE Computer Society Press, 2000.
- [20] Edward R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, CT, 1983.
- [21] Edward R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, CT, 1990.
- [22] Colin Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, 2000.
- [23] Marc Weber, Marc Alexa, and Wolfgang Müller. Visualizing time-series on spirals. In *IEEE Symposium on Information Visualization 2001 (InfoVis 2001)*, pages 7–13. IEEE Computer Society Press, 2001.