

# Chapter 1

## Introduction

Computers should also help us warp time, but the challenge here is even greater. Normal experience doesn't allow us to roam freely in the fourth dimension as we do in the first three. So we've always relied on technology to aid our perception of time.

---

[Udell \(2004, p. 32\)](#)

Space and time are two outstanding dimensions because in conjunction they represent four-dimensional space or simply the world we are living in. Basically, every piece of data we measure is related and often only meaningful within the context of space and time. Consider for example the price of a barrel of oil. The data value of \$129 alone is not very useful. Only if assessed in the context of where (space) and when (time) is the oil price valid and only then is it possible to meaningfully interpret the cost of \$129.

Space and time differ fundamentally in terms of how we can navigate and perceive them. Space can in principle be navigated arbitrarily in all three spatial dimensions, and we can go back to where we came from. Humans have senses for perceiving space, in particular the senses of sight, touch, and hearing. Time is different; it does not allow for active navigation. We are constrained to the unidirectional character of constantly proceeding time. We cannot go back to the past and we have to wait patiently for the future to become present. And above all else, humans do not have senses for perceiving time directly. This fact makes it particularly challenging to visualize time – making the invisible visible.

Time is an important data dimension with distinct characteristics. Time is common across many application domains as for example medical records, business, science, biographies, history, planning, or project management. In contrast to other quantitative data dimensions, which are usually “flat”, time has an inherent semantic structure, which increases time's complexity substantially. The hierarchical structure of granularities in time, as for example minutes, hours, days, weeks, and months, is unlike that of most other quantitative dimensions. Specifically, time comprises different forms of divisions (e.g., 60 minutes correspond to one hour, while 24 hours make up one day), and granularities are combined to form calendar systems

(e.g., Gregorian, Julian, business, or academic calendars). Moreover, time contains natural cycles and re-occurrences, as for example seasons, but also social (often irregular) cycles, like holidays or school breaks. Therefore, time-oriented data, i.e., data that are inherently linked to time, need to be treated differently than other kinds of data and require appropriate visual and analytical methods to explore and analyze them.

The human perceptual system is highly sophisticated and specifically suited to spot visual patterns. Visualization strives to exploit these capabilities and to aid in seeing and understanding otherwise abstract and arcane data. Early visual depictions of time-series even date back to the 11th century. Today, a variety of visualization methods exist and visualization is applied widely to present, explore, and analyze data. However, many visualization techniques treat time just as a numeric parameter among other quantitative dimensions and neglect time's special character. In order to create visual representations that succeed in assisting people in reasoning about time and time-oriented data, visualization methods have to account for the special characteristics of time. This is also demanded by [Shneiderman \(1996\)](#) in his well-known task by data type taxonomy, where he identifies temporal data as one of seven basic data types most relevant for information visualization.

Creating good visualization usually requires good data structures. However, commonly only simple sequences of time-value-pairs  $\langle (t_0, v_0), (t_1, v_1), \dots, (t_n, v_n) \rangle$  are the basis for analysis and visualization. Accounting for the special characteristics of time can be beneficial from a data modeling point of view. One can use different calendars that define meaningful systems of granularities for different application domains (e.g., fiscal quarters or academic semesters). Data can be modeled and integrated at different levels of granularity (e.g., months, days, hours, and seconds), enabling for example value aggregation along granularities. Besides this, data might be given for time intervals rather than for time points, as for example in project plans, medical treatments, or working shift schedules. Related to this diversity of aspects is the problem that most of the available methods and tools are strongly focused on special domains or application contexts. [Silva and Catarci \(2000\)](#) conclude:

It is now recognized that the initial approaches, just considering the time as an ordinal dimension in a 2D or 3D visualizations [sic], are inadequate to capture the many characteristics of time-dependent information. More sophisticated and effective proposals have been recently presented. However, none of them aims at providing the user with a complete framework for visually managing time-related information.

[Silva and Catarci \(2000, p. 9\)](#)

The aim of this book is to present and discuss the multitude of aspects which are relevant from the perspective of visualization. We will characterize the dimension of time as well as time-oriented data, and describe tasks that users seek to accomplish using visualization methods. While time and associated data form a part of *what* is being visualized, user tasks are related to the question *why* something is visualized. *How* these characteristics and tasks influence the visualization design will be explained by several examples. These investigations will lead to a systematic categorization of visualization approaches. Because interaction techniques and

analytical methods also play an important role in the exploration of and reasoning with time-oriented data, these will also be discussed. A large part of this book is devoted to a survey of existing techniques for visualizing time and time-oriented data. This survey presents self-contained descriptions of techniques accompanied by an illustration and corresponding references on a per-page basis.

Before going into detail on visualizing time-oriented data, let us first take a look at the basics and examine general concepts of information visualization.

## 1.1 Introduction to Visualization

Visualization is a widely used term. [Spence \(2007\)](#) refers to a dictionary definition of the term: *visualize* – to form a mental model or mental image of something. Visual representations have a long and venerable history in communicating facts and information. But only about twenty years have passed since visualization became an independent self-contained research field. In 1987 the notion of visualization in scientific computing was introduced by [McCormick et al. \(1987\)](#). They defined the term *visualization* as follows:

Visualization is a method of computing. It transforms the symbolic into the geometric, enabling researchers to observe their simulations and computations. Visualization offers a method for seeing the unseen. It enriches the process of scientific discovery and fosters profound and unexpected insights.

[McCormick et al. \(1987, p. 3\)](#)

The goal of this new field of research has been to integrate the outstanding capabilities of human visual perception and the enormous processing power of computers to support users in analyzing, understanding, and communicating their data, models, and concepts. In order to achieve this goal, three major criteria have to be satisfied (see [Schumann and Müller, 2000](#)):

- expressiveness,
- effectiveness, and
- appropriateness.

*Expressiveness* refers to the requirement of showing exactly the information contained in the data; nothing more and nothing less must be visualized. *Effectiveness* primarily considers the degree to which visualization addresses the cognitive capabilities of the human visual system, but also the task at hand, the application background, and other context-related information, to obtain intuitively recognizable and interpretable visual representations. Finally, *appropriateness* involves a cost-value ratio in order to assess the benefit of the visualization process with respect to achieving a given task. While the value of a visual representation is not so easy to determine (see [Van Wijk, 2006](#)), cost is often related to time efficiency (i.e., the computation time spent) and space efficiency (i.e., the exploited screen space).

Expressiveness, effectiveness, and appropriateness are criteria that any visualization should aim to fulfill. To this end, the visualization process, above all else, has

to account for two aspects: the data and the task at hand. In other words, we have to answer the two questions: “What has to be presented?” and “Why does it have to be presented?”. We will next discuss both questions in more detail.

### What? – Specification of the data

In recent years, different approaches have been developed to characterize data – the central element of visualization. In their overview article, [Wong and Bergeron \(1997\)](#) established the notion of *multidimensional multivariate data* as multivariate data that are given in a multidimensional domain. This definition leads to a distinction between *independent* and *dependent* variables. Independent variables define an  $n$ -dimensional domain. In this domain, the values of  $k$  dependent variables are measured, simulated, or computed; they define a  $k$ -variate dataset. If at least one dimension of the domain is associated with the dimension of time, we call the data *time-oriented data*.

Another useful concept for modeling data along cognitive principles is the *pyramid framework* by [Mennis et al. \(2000\)](#). At the level of data, this framework is based on three perspectives (also see Figure 3.29 on p. 63): *where* (location), *when* (time) and *what* (theme). The perspectives *where* and *when* characterize the data domain, i.e., the independent variables as described above. The perspective *what* describes what has been measured, observed, or computed in the data domain, i.e., the dependent variables as described above. At the level of knowledge, the *what* includes not only simple data values, but also objects and their relationships, where objects and relations may have arbitrary data attributes associated with them.

From the visualization point of view, all aspects need to be taken into account: The aspect *where* to represent the spatial frame of reference and to associate data values to locations, the aspect *when* to show the characteristics of the temporal frame of reference and to associate data values to the time domain, and the aspect *what* to represent individual values or abstractions of a multivariate dataset. As our interest is in time and time-oriented data, this book places special emphasis on the aspect *when*. We will specify the key properties of time and associated data in Chapter 3 and discuss the specific implications for visualization in Chapter 4.

### Why? – Specification of the task

Similar to specifying the data, one also needs to know why the data are visualized and what tasks the user seeks to accomplish with the help of the visualization. On a very abstract level, the following three basic goals can be distinguished (see [Ward et al., 2010](#)):

- explorative analysis,
- confirmative analysis, and
- presentation of analysis results.

*Explorative analysis* can be seen as undirected search. In this case, no a priori hypotheses about the data are given. The goal is to get insight into the data, to begin extracting relevant information, and to come up with hypotheses. In a phase of *confirmative analysis*, visualization is used to prove or disprove hypotheses, which can originate from data exploration or from models associated with the data. In this sense, confirmative analysis is a form of directed search. When facts about the data have eventually been ascertained, it is the goal of the *presentation* step to communicate and disseminate analysis results.

These three basic visualization goals call for quite different visual representations. This becomes clear when taking a look at two established visualization concepts: filtering and accentuation. The aim of filtering is to visualize only relevant data and to omit less relevant information, and the goal of accentuation is to highlight important information. During explorative analysis, both concepts help users to focus on selected parts or aspects of the data. But filtering and accentuation must be applied carefully, because it is not usually known which data are relevant or important. Omitting or highlighting information indiscriminately can lead to misinterpretation of the visual representation and to incorrect findings. During confirmative analysis, filtering can be applied more easily as the data which is relevant, that is, the data that contribute to the hypotheses to be evaluated are usually known. Accentuation and de-accentuation are common means to enhance expressiveness and effectiveness, and to fine-tune visual presentations in order to communicate results and insight yielded by an exploratory or confirmative analysis process.

Although the presentation of results is very important, this book is more about visual analysis and interactive exploration of time-oriented data. Therefore, we will take a closer look at common analysis and exploration tasks. As [Bertin \(1983\)](#) describes, human visual perception has the ability to focus (1) on a particular element of an image, (2) on groups of elements, or (3) on an image as a whole. Based on these capabilities, three fundamental categories of interpretation aims have been introduced by [Robertson \(1991\)](#): *point*, *local*, and *global*. They indicate which values are of interest: (1) values at a given point of the domain, (2) values in a local region, or (3) all values of the whole domain. These basic tasks can be subdivided into more specific, concrete tasks, which are usually given as a list of verbal descriptions. [Wehrend and Lewis \(1990\)](#) define several such low-level tasks: identify or locate data values, distinguish regions with different values or cluster similar data, relate, compare, rank, or associate data, and find correlations and distributions. The task by data type taxonomy by [Shneiderman \(1996\)](#) lists seven high-level tasks that also include the notion of interaction with the data in addition to purely visual tasks:

- Overview: gain an overview of the entire dataset
- Zoom: zoom in on data of interest
- Filter: filter out uninteresting information
- Details-on-demand: select data of interest and get details when needed
- Relate: view relationships among data items
- History: keep a history of actions to support undo and redo
- Extract: allow extraction of data and of query parameters

Yi et al. (2007) further refine the aspect of interaction in information visualization and derive a number of categories of interaction tasks. These categories are organized around the user's intentions to interactively adjust visual representations to the tasks and data at hand. Consequently, a *show me* prefaces six categories:

- show me something else (explore)
- show me a different arrangement (reconfigure)
- show me a different representation (encode)
- show me more or less detail (abstract/elaborate)
- show me something conditionally (filter)
- show me related items (connect)

The *show me* tasks allow for switching between different subsets of the analyzed data (explore), different arrangements of visual primitives (reconfigure), and different visual representations (encode). They also address the navigation of different levels of detail (abstract/elaborate), the definition of data of interest (filter), and the exploration of relationships (connect).

In addition to the *show me* categories, Yi et al. (2007) introduce three further interaction tasks:

- mark something as interesting (select)
- let me go to where I have already been (undo/redo)
- let me adjust the interface (change configuration)

*Mark something as interesting (select)* subsumes all kinds of selection tasks, including picking out individual data values as well as selecting entire subsets of the data. Supporting users in going back to interesting data or views (undo/redo) is essential during interactive data exploration. Adaptability (change configuration) is relevant when a system is applied by a wide range of users for a variety of tasks and data types.

As we have seen, the purpose of visualization, that is, the task to be accomplished with visualization, can be defined in different ways. The above mentioned visualization and interaction tasks serve as a basic guideline to assist visualization designers in developing representations that effectively support users in conducting visual data exploration and analysis. In Chapter 4 we will come back to this issue and refine tasks with regard to the analysis of time-oriented data. The aspect of interaction will be taken up in Chapter 5.

## How? – The visualization pipeline


In order to generate effective visual representations, raw data have to be transformed into image data in a data-dependent and task-specific manner. Conceptually, raw data have to be mapped to geometry and corresponding visual attributes like color, position, size, or shape, also called *visual variables* (see Bertin, 1983; Mackinlay, 1986). Thanks to the capabilities of our visual system, the perception of visual stimuli is mostly spontaneous. As indicated earlier, Bertin (1983) distinguishes three

levels of cognition that can be addressed when encoding information to visual variables. On the first level, elementary information is directly mapped to visual variables. This means that every piece of elementary information is associated with exactly one specific value of a visual variable. The second level involves abstractions of elementary information, rather than individual data values. By mapping the abstractions to visual variables, general characteristics of the data can be communicated. The third level combines the two previous levels and adds representations of further analysis steps and metadata to convey the information contained in a dataset in its entirety.

To facilitate generation of visual output at all three levels, a flexible mapping strategy is required. Such a strategy has been manifested as the so-called *visualization pipeline*, first introduced by [Haber and McNabb \(1990\)](#). The visualization pipeline consists of the three steps (see Figure 1.1(a)):

1. filtering,
2. mapping, and
3. rendering.


The filtering step prepares the raw input data for processing through the remaining steps of the pipeline. This is done with respect to the given analysis task and includes not only selection of relevant data but also operations for data enrichment or data reduction, interpolation, data cleansing, grouping, dimension reduction, and others. Literally, the mapping step maps the prepared data to appropriate visual variables. This is the most crucial step as it largely influences the expressiveness and effectiveness of the resulting visual representation. Finally, the rendering step generates actual images from the previously computed geometry and visual attributes. This general pipeline model is the basis for many visualization systems.



**Fig. 1.1:** The visualization pipeline.


The basic pipeline model has been refined by [dos Santos and Brodlie \(2004\)](#) in order to better address the requirements of higher dimensional visualization problems. The original filtering has been split up into two separate steps: data analysis and filtering (see [Figure 1.1\(b\)](#)). The data analysis carries out automatic computations like interpolations, clustering, or pattern recognition. The filtering step then extracts only those pieces of data that are of interest and need to be presented. In the case of large high-dimensional datasets, the filtering step is highly relevant because displaying all information will most likely lead to complex and overloaded visual representations that are hard to interpret. Because interests may vary among users, tasks, and data, the filtering step has to support the interactive refinement of filter conditions. Further input like the specific analysis task or hypothesis as well as application specific details can be used to steer the data extraction process.

In an effort to formally model the visualization process, [Chi \(2000\)](#) built upon the classic pipeline model and derived the *data state reference model*. This model reflects the stepwise transformation of abstract data into image data through several stages by using operators. While transformation operators transform data from one level of abstraction to another, within stage operators process the data only within the same level of abstraction (see [Figure 1.2](#)). This model broadens the capabilities of the visualization process and allows the generation of visual output at all of Bertin's levels. Different operator configurations lead to different views on the data, and thus, to comprehensive insight into the analyzed data. It is obvious that the selection and configuration of appropriate operators to steer the visualization process is a complex problem that depends mainly on the given visualization goal, which in turn is determined by the characteristics of the data and the task at hand.



**Fig. 1.2:** The data state reference model (adapted from [Chi, 2000](#)).

The previous paragraphs may suggest that the image or view eventually generated by a visualization pipeline is an end product. But that is not true. In fact, the user controls the visualization pipeline and interacts with the visualization process in various ways. Views and images are created and adjusted until the user deems them suitable. Therefore, [Card et al. \(1999\)](#) integrate the user in their *information visualization reference model* (see [Figure 1.3](#)).



**Fig. 1.3:** The information visualization reference model (adapted from [Card et al., 1999](#)).


Having introduced the very basics of interactive visualization, we now move on to an application example. The goal is to illustrate a concrete visual representation and to demonstrate possible benefits for data exploration and analysis.

## 1.2 Application Example

Our particular example is in the domain of medicine. A considerable share of physicians' daily work time is devoted to searching and gathering patient-related information to form a basis for adequate medical treatment and decision-making. The amount of information is enormous and disorganized, and physicians might be overwhelmed by the information provided to them. Often, datasets comprise multiple variables of different data types that are sampled irregularly and independently from each other, as for example quantitative parameters (e.g., blood pressure or body temperature) and qualitative parameters (e.g., events like a heart attack) as well as instantaneous data (e.g., blood sugar measurement at a certain point in time) and interval data (e.g., insulin therapy from January to May 2010). Moreover, the data commonly originate from heterogeneous sources like electronic lab systems, hospital information systems, or patient data sheets that are not well integrated. Exploring such heterogeneous time-oriented datasets to get an overview of the history or the current health status of an individual patient or a group of patients is a challenging task.

Interactive visualization is an approach to representing a coherent view of such medical data and to catering for easy data exploration. In our particular example, an active discourse of the physician via interaction with the visual representation is of major importance since most static representations cannot satisfy task-dependent information needs seamlessly. In addition to presenting information intuitively, aiding clinicians in gaining new medical insights about patients' current health status, state changes, trends, or patterns over time is an important aspect.

*VisuExplore* is an interactive visualization tool for exploring a heterogeneous set of medical parameters over time (see [Rind et al., 2010](#) and  $\leftrightarrow$  p. 231). *VisuExplore* uses multiple views along a common horizontal time axis to convey the different medical parameters involved. It is based on several well-known visualization meth-



**Fig. 1.4:** Visualization of heterogeneous medical parameters of a diabetes patient.

ods, including line plots (↔ p. 153), bar graphs (↔ p. 154), event charts, and timelines (↔ p. 166), that are combined and integrated.

Figure 1.4 shows data of a diabetes patient over a period of two years and three months between November 2006 and March 2009. Beneath a panel that shows patient master data, eight visualization views are visible.

A document browser is placed on top that shows icons for medical documents, like for example diagnostic findings or x-ray images. In our example case, the document browser contains progress notes, as at the very beginning of treatment the physicians suspected renal failure. Next, a line plot with semantic zoom (see p. 112) is present which shows blood glucose values. Colored areas below the line provide qualitative information about normal (green), elevated (yellow), and high (red) value ranges which makes this semantic information easy to read. Below that, another line plot with semantic zoom functionality shows HbA1c (an indicator of a patient’s blood glucose condition over the previous several weeks). In this case, more vertical space is devoted to the chart, thus allowing more exact readings of the values. Still, semantic information is added as color annotation of the y-axis, using small ticks to indicate when the variable’s value crosses qualitative range boundaries (e.g., from critically high to elevated, as shown in the screenshot via a horizontal line that is colored red and yellow). Below the blood sugar values, there are two timeline charts showing the insulin therapy and oral anti-diabetic drugs. Insulin is categorized into rapid-acting insulin (ALT), intermediate-acting insulin (VZI), and a mixture of these (Misch). Details about brand name or dosage in free text are shown as labels that are located below the respective timeline. Oral anti-diabetic drugs are shown via an event chart below. There are also free text details about oral diabetes medication. The sixth view is a bar graph with adjacent bars for systolic and diastolic blood

pressure. The bottom two views are line plots related to the body mass index (BMI) and blood lipids with two lines showing triglyceride and cholesterol values.

This arrangement has been chosen because it places views of medical tests directly above views of the related medical interventions. The height of some views has been reduced to fit on a single screen. This is possible because all information that is relevant for the physician's current task can still be recognized in this state.


The shown diabetes case is a 44-year-old patient with initially very high blood sugar values. From the interactive visual representations, several facts about the patient can be inferred as illustrated by the following insights that were gained by a physician using the VisuExplore system. The initially high blood sugar values were examined in detail via tooltips and showed exact values of 428 mg/dl glucose and 14.8% HbA1c. In addition, it can be seen in the bottom panel that blood lipid values are also high (256 mg/dl cholesterol, 276 mg/dl triglyceride). At the same time, the body mass index shown above is rather low (20.1). From the progress notes in the document browser it can be seen that the physician had the suspicion of a nephropathy. But these elevated values are also signs of latent autoimmune diabetes of adults, a special form of type 1 diabetes. After one month, blood sugar has improved (168 mg/dl glucose) and blood lipids have normalized. The patient switched to insulin therapy in a combination of rapid-acting insulin (ALT) and intermediate-acting insulin (VZI). Since April 2007, the insulin dosage has remained stable and concomitant medication is no longer needed. The patient's overall condition has improved through blood sugar management. Furthermore, the physician involved in the case study wondered about the very high HbA1c value of 11.9% in November 2006 and why diabetes treatment had only started four months later.

VisuExplore's interactive features allow physicians to get an overview of multiple medical parameters and focus on parts of the data. Physicians can add visualizations with one or more additional variables. They may resize and rearrange visualizations. Further, it is possible to navigate and zoom across the time dimension by dragging the mouse, by using dedicated buttons, or by selecting predefined views (e.g., last year). Moreover, the software allows the selecting and highlighting of data elements. Other time-based visualization and interaction techniques can extend the system to support special purposes. For example, a document browser shows medical documents (e.g., discharge letters or treatment reports) as document icons (e.g., PDF, Word) that physicians can click on if they want to open the document. VisuExplore integrates with the hospital information systems and accesses the medical data stored there.

This example demonstrated that visual representations are capable of providing a coherent view of otherwise heterogeneous and possibly distributed data. The integrative character also supports interactive exploration and task-specific focusing on relevant information.

### 1.3 Book Outline

With the basics of visualization and an application example, we have set the stage for the next chapters. Before going into detail about the contemporary visualization of time and time-oriented data, some inspiring and thought-provoking historical depictions and images from the arts are given attention in Chapter 2. The characteristics of time and data for modern interactive visualization on computers are the focus of Chapter 3. The actual visualization process, that is the transformation of abstract data to visual representations, will be discussed in Chapter 4, taking into account the key question words *what*, *why*, and *how* to visualize. In Chapters 5 and 6, we go beyond pure visualization methods and discuss cornerstones of interaction and analytical methods to support exploration and visual analysis. A major part of this book is devoted to a survey of existing information visualization techniques for time and time-oriented data in Chapter 7. Throughout the book we use the  $\leftrightarrow$  symbol followed by a page number to refer the reader to a particular technique in the survey. A final summary along with a discussion of open challenges can be found in Chapter 8. Figure 1.5 provides a visual overview of the contents of the book.



**Fig. 1.5:** Visual overview of the contents of the book.

Please refer to the companion website of the book for updates and additional resources including links to related material, visualization prototypes, and technique descriptions: <http://www.timeviz.net>.

### References

- Bertin, J. (1983). *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press, Madison, WI, USA. translated by William J. Berg.
- Card, S., Mackinlay, J., and Shneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, San Francisco, CA, USA.

- Chi, E. H. (2000). A Taxonomy of Visualization Techniques Using the Data State Reference Model. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis)*, pages 69–76, Washington, DC, USA. IEEE Computer Society.
- dos Santos, S. and Brodlié, K. (2004). Gaining Understanding of Multivariate and Multidimensional Data through Visualization. *Computers & Graphics*, 28:311–325.
- Haber, R. B. and McNabb, D. A. (1990). Visualization Idioms: A Conceptual Model for Scientific Visualization Systems. In *Visualization in Scientific Computing*, pages 74–93. IEEE Computer Society, Los Alamitos, CA, USA.
- Mackinlay, J. (1986). Automating the Design of Graphical Presentations of Relational Information. *ACM Transactions on Graphics*, 5(2):110–141.
- McCormick, B. H., DeFanti, T. A., and Brown, M. D. (1987). Visualization in Scientific Computing. *Computer Graphics*, 21(6).
- Mennis, J. L., Peuquet, D., and Qian, L. (2000). A Conceptual Framework for Incorporating Cognitive Principles into Geographical Database Representation. *International Journal of Geographical Information Science*, 14(6):501–520.
- Rind, A., Miksch, S., Aigner, W., Turic, T., and Pohl, M. (2010). VisuExplore: Gaining New Medical Insights from Visual Exploration. In Hayes, G. R. and Tan, D. S., editors, *Proceedings of the 1st International Workshop on Interactive Systems in Healthcare (WISH@CHI2010)*, pages 149–152, New York, NY, USA. ACM Press.
- Robertson, P. K. (1991). A Methodology for Choosing Data Representations. *IEEE Computer Graphics and Applications*, 11(3):56–67.
- Schumann, H. and Müller, W. (2000). *Visualisierung – Grundlagen und allgemeine Methoden*. Springer, Berlin, Germany.
- Shneiderman, B. (1996). The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 336–343, Los Alamitos, CA, USA. IEEE Computer Society.
- Silva, S. F. and Catarci, T. (2000). Visualization of Linear Time-Oriented Data: A Survey. In *Proceedings of the International Conference on Web Information Systems Engineering (WISE)*, pages 310–319, Los Alamitos, CA, USA. IEEE Computer Society.
- Spence, R. (2007). *Information Visualization: Design for Interaction*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2nd edition.
- Udell, J. (2004). Space, Time, and Data. *InfoWorld*, 26(26):32.
- Van Wijk, J. J. (2006). Views on Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):421–433.
- Ward, M., Grinstein, G., and Keim, D. (2010). *Interactive Data Visualization: Foundations, Techniques, and Applications*. A K Peters Ltd, Natick, MA, USA.
- Wehrend, S. and Lewis, C. (1990). A Problem-Oriented Classification of Visualization Techniques. In *Proceedings of IEEE Visualization (Vis)*, pages 139–143, Los Alamitos, CA, USA. IEEE Computer Society.
- Wong, P. C. and Bergeron, R. (1997). 30 Years of Multidimensional Multivariate Visualization. In Nielson, G. M., Hagen, H., and Müller, H., editors, *Scientific Visualization*, pages 40–62. IEEE Computer Society, Los Alamitos, CA, USA.
- Yi, J. S., ah Kang, Y., Stasko, J., and Jacko, J. (2007). Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231.