Views on Visualization

Jarke J. van Wijk, *Member*, *IEEE*

Abstract—The field of visualization is maturing. Many problems have been solved and new directions are sought. In order to make good choices, an understanding of the purpose and meaning of visualization is needed. In this paper, visualization is considered from multiple points of view. First, a technological viewpoint is adopted, where the value of visualization is measured based on effectiveness and efficiency. An economic model of visualization is presented and benefits and costs are established. Next, consequences and limitations of visualization are discussed (including the use of alternative methods, high initial costs, subjectiveness, and the role of interaction). Example uses of the model for the judgment of existing classes of methods are given to understand why they are or are not used in practice. However, such an economic view is too restrictive. Alternative views on visualization are presented and discussed: visualization as an art, visualization as design and, finally, visualization as a scientific discipline.

Index Terms—Visualization, evaluation, validation, methodology, survey, challenges.

1 Introduction

Modern society is confronted with a data explosion. Acquisition devices such as MRI-scanners, large scale simulations on supercomputers, but also stock trading at stock exchanges produce very large amounts of data. Visualization of data makes it possible for researchers, analysts, engineers, and the lay audience to obtain insight into these data in an efficient and effective way thanks to the unique capabilities of the human visual system, which enables us to detect interesting features and patterns in a short period of time.

Many of us will have written paragraphs like the preceding one, where I attempted to give the standard rationale of our field. In 1987, when the US National Science Foundation's (NSF) influential ViSC report [20] appeared, the expectations were high. Visualization was considered vital and highly promising for the scientific process. In the two decades since the report, much progress has been made. The advances in graphics hardware are astonishing; most laptop computers are graphics superworkstations according to the standards of just a decade ago. Many new methods, techniques, and systems have been developed. Some of them, such as slices, height-surfaces, and isosurfaces are now routinely used in practice.

On the other hand, many of these new methods are not used in real-world situations. Many research results are nowadays considered as incremental by reviewers and our prospective users rarely attend our conferences. So, are we, as researchers in visualization, on the right track?

With this paper, I want to contribute to the discussion on the status and possible directions of our field. Rather than pinpointing specific topics and activities, my aim is to detect overall patterns, to find a way to understand and qualify visualization in general, and to enumerate a number of perspectives on visualization. This is an ambitious and vague plan, although the basic point of departure for this is highly practical.

I have to make decisions on visualization in many roles. As a researcher, these decisions range from in which area to invest time to which particular solution to implement; as a supervisor, guidance to students must be provided; as a reviewer, new results and proposals for new research must be judged and opinions are expected on whether they are worth publishing or funding; as an advisor to a start-up company, novel and profitable directions must be spotted. All these cases imply judgment of the value of visualization in varying senses.

How does one assess the value of visualization? Visualization itself is an ambiguous term. It can refer to the research discipline, to a technology, to a specific technique, or to the visual result. In the following, I use "value of visualization" as shorthand for "the value of one particular visualization method, technique, system, or result," in contrast to the value of visualization in general. One aim, however, is to find a generic approach to assess the value of such particular visualizations.

If visualization is considered as a technology, i.e., as a collection of methods, techniques, and tools developed and applied to satisfy a need, then standard measures apply: Visualization has to be *effective* and *efficient*. In other words, visualization should do what it is supposed to do and has to do this using a minimal amount of resources. One immediate and obvious implication is that we cannot judge visualization on its own, but have to take into account the context in which it is used.

In Section 2, a short overview is given of the background of the topic discussed here. In Section 3, an economic model of visualization is proposed. The basic elements are identified first, the associated costs and gains are added next. Various implications of the model are discussed in Section 4. In Section 5, this model is applied to several cases. However, considering visualization as a technology is just one of the possible views. In Sections 6, 7, 8, and 9, visualization is considered from the viewpoint of innovation, art, design, and science, respectively. Conclusions can be found in Section 10.

• *The author is with the Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, PO Box 513, 5600 MB Eindhoven, The Netherlands. E-mail: vanwijk@win.tue.nl.*

Manuscript received 13 Nov. 2005; revised 30 Jan. 2006; accepted 6 Feb. 2006; published online 10 May 2006.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-0162-1105.

This paper is an extended version of [31]. The first version was primarily focussed on valuation of visualization; here, the emphasis is shifted to multiple perspectives on visualization.

Finally, the views presented here are, on one hand, very general, high-level, and abstract; on the other hand, they are also very personal, in the sense that they concern (subjective) values, valuation of my own work, and personal perspectives. To reflect this, I use the first person in this paper, to emphasize that the opinions given are personal. Most examples that I use come from my own work, often developed collaboratively with coworkers. The main reason for this presentation style is simply that I am most familiar with it, not only with the techniques and results, but also with the context in which it took place.

2 Background

If we use the publication of the ViSC report [20] in 1987 as the year when visualization started, our discipline celebrated its 18th anniversary in 2005. In The Netherlands, this is the age at which a person is considered mature. Many things have changed since 1987. Graphics hardware developments are amazing and, so, is the large number of techniques that have been developed to visualize data in a variety of ways. For new cases, visualization experts can concentrate on interesting and hard problems, instead of having to write straightforward code from scratch.

However, there are signs that there is a need to reconsider visualization. First of all, there seems to be a growing gap between the research community and its prospective users. Few attendants at the IEEE Visualization Conference are prospective users looking for new ways to visualize their data and solve their problems. Second, the community itself is getting both more specialized and critical, judging from my experience as paper cochair for IEEE Visualization 2003 and 2004. In the early nineties, the field lay fallow and it was relatively easy to come up with new ideas. The proceedings in the early nineties show great diversity. Today, the field is getting more specialized, submitted work often consists of incremental results. This could be a signal that our field is getting mature. Advances in the hard sciences are often highly incremental, with breakthroughs sometimes being small changes on basic concepts, which are viewed as incremental by nonexperts. On the other hand, it is not always clear that these incremental contributions have merit and reviewers are getting more and more critical. Third, some major problems have been more or less solved [18]. For volume rendering of medical data, sophisticated industrial packages that satisfy the needs of many users are available.

These trends suggest an urgent need to reconsider the field and to consider new directions. Several researchers have presented [10], [12], [21], [8] overviews of current challenges. Another great overview of the current status of visualization and suggestions for new directions is provided by the position papers [3] contributed by the attendants of the joint NSF-NIH Fall 2004 Workshop on Visualization Research Challenges, organized by Terry Yoo. One contribution [18] is particularly disturbing, for its title, the name and fame of the author, and the vivid description that indeed the field has changed and new directions are needed. Finally, a very impressive overview of the field and recommendations for future directions is provided in *Illuminating the Path*, edited by Jim Thomas and Kristin Cook [24]. Visual analytics is the central topic here, defined as the science of analytical reasoning facilitated by interactive visual interfaces.

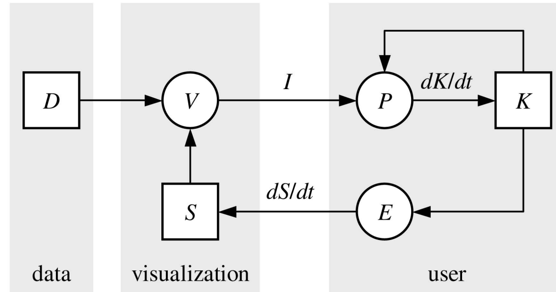

Fig. 1. A simple model of visualization.

Many issues are mentioned several times, including handling of complex and large data sets, uncertainty, validation, and a better understanding of the visualization process itself. Very high on the lists is integration with the tasks of the user. This holds especially for visual analytics [24], where the focus is shifted from visualization to the overall sense-making process.

In this paper, no attempt is made to summarize or to provide an overview of these challenges, but the field is considered from a greater distance. First, an attempt is made to find a model or procedure to judge, in general, if a method is worthwhile or not. In the following sections, a first step toward such a model is presented. Much of it is evident and obvious, but some statements made are more surprising and sometimes contrary to main stream thinking. To stimulate the debate, I have taken the liberty of presenting these more extreme positions also, hoping that some readers will not be offended. After these critical reflections, the point of view is shifted and from these different perspectives other criteria emerge.

3 Model

In this section, I propose a generic model for visualization. First, the major ingredients are identified; then, costs and gains are associated. The model is abstract and coarse, but it can be used to identify some relevant aspects, patterns and trends.

3.1 Visualization and Its Context

Fig. 1 shows the basic model. Boxes denote containers, circles denote processes that transform inputs into outputs. The aim here is not to position different visualization methods, for which a taxonomy would be a more suitable approach, but rather to describe the context in which visualization operates. No distinction is made, for instance, between scientific visualization and information visualization. At this level, they are considered to have much in common.

The central process in the model is visualization V. Data D is transformed according to a specification S into a time varying image $I(t)$. All these variables and functions should be considered in the broadest sense. The type of data D to be visualized might vary from a single bit to a time-varying 3D tensor field; the specification S includes a specification of the hardware used, the algorithms to be applied (in the form of a selection of a predefined method or in the form of code), and the specific parameters to be used; the image I will often be an image in the usual sense, but it can also be an animation, or auditory or haptic feedback. In other words, this broad definition encompasses both a humble LED on an electronic device that visualizes whether the device is on or off, as well as a large virtual reality set-up to visualize physical and chemical processes in the atmosphere. The image I is perceived by a user, with an increase in knowledge K as a result. The amount of knowledge gained depends on the image I, the current knowledge K of the user, and the particular properties of the perceptual and cognitive abilities P of the user. Concerning the influence of K, a physician will be able to extract more information from a medical image than a lay-person. But, also, when much knowledge is already available, the additional knowledge shown in an image can be low. A map showing the provinces of the Netherlands provides more new information to a person from the US than to a Dutch national. Also, the additional value of an image of time-step 321 is probably small when time-step 320 has just been studied before. Concerning the influence of P, a simple but important example is that a colorblind person will be less effective in extracting knowledge from a colorful image than a person with normal visual capabilities. But, also, some people are much better than others in spotting special patterns, structures, and configurations.

The current knowledge $K(t)$ is the sum of the initial knowledge K_0 and all knowledge gained from the images viewed so far.

An important aspect is interactive exploration, here, represented by $E(K)$. Starting from an initial specification S_0, the user may continuously change the specification of the visualization, based on his current knowledge, in order to explore the data further.

3.2 Economic Model

To assess whether a visualization method is worthwhile, we must consider its value. I propose using profitability in an economic sense as a measure for this. I simplify this by assuming that there is a homogeneous user community, consisting of n users which use a certain visualization V to visualize a data set m times each, where each session takes k exploratory steps and time T. This is a crude simplification of course. In the real world, the user community will often be highly varied, with different K_0s and also with different aims. The costs associated with using V come at four different levels:

- $C_i(S_0)$: *Initial development costs*. The visualization method has to be developed and implemented, possibly new hardware has to be acquired.
- $C_u(S_0)$: *Initial costs per user*. The user has to spend time on selection and acquisition of V, understanding how to use it, and tailoring it to his particular needs.
- $C_s(S_0)$: *Initial costs per session*. Data have to be converted and an initial specification of the visualization has to be made.
- C_e: *Perception and exploration costs*. The user has to spend time to view and understand the visualization, as well as to modify and tune the specification, thereby exploring the data set.

The total costs are now given by

$$C = C_i + nC_u + nmC_s + nmkC_e.$$

The return on these investments consists of the value $W(\Delta K)$ of the acquired knowledge $\Delta K = K(T) - K(0)$ per session, multiplied by the total number of sessions:

$$G = nmW(\Delta K)$$

and, hence, for the total profit $F = G - C$, we find

$$F = nm(W(\Delta K) - C_s - kC_e) - C_i - nC_u.$$

This gives us a recipe to decide on the value of a visualization method. Positive are high values for n, m, $W(\Delta K)$, and low values for C_s, C_e, C_i, C_u, and k. Or, in other words, a great visualization method is used by many people, who use it routinely to obtain highly valuable knowledge, without having to spend time and money on hardware, software, and effort. Indeed, quite obvious.

4 Implications

Quantification of the elements of the model is hard. In this section, we discuss this in more detail, as well as a number of other issues implied by this model.

4.1 Valuable Knowledge

Insight is the traditional aim of visualization. The term itself is great and suggests a high-level contribution to the advance of science. Users are enabled to see things they were not aware of, and this insight helps them to define new questions, hypotheses, and models of their data. However, from an operational point of view, the term insight does not help us much further in assessing the value of visualization. One problem is that we cannot directly observe or measure how much insight is acquired and, also, it is difficult to assess what the value of that insight is. In the model, we use the term knowledge, but this suffers from the same limitations. Also, there is a strange paradox in the basic paradigm of visualization. We do not know what information is contained in the data, hence we use visuals to get insight. But, if we do not know which specific aspects or features should be visible, we cannot assess if we are successful or not.

Nevertheless, we should try to measure or estimate $W(\Delta K)$ if we want to assess the value of visualization, especially because it is the only term in the model for F with a positive sign. An operational approach is to consider the use of visualization as an element in problem solving. Users have a problem, they must decide which action to take, and, to make that decision, they need information. The visualization should enable them to extract the relevant information from the data.

Decisions are typically about actions to be taken or not. Such decisions include, for instance: "Should a stock be bought or sold?", "Should this patient receive an operation?", or "Which people in an organization are candidates for promotion?" Hence, I recommend to my students to search for and enumerate possible actions of users after using their prospective tools. If such actions cannot be found or defined, the value of visualization is in doubt. Just claiming that a visualization gives insight is not enough if we want to offer additional value.

If we know to which actions the visualization should lead, the next steps are assessing whether the knowledge derived from the visualization does indeed support the decision and, also, to assess the economic value of this decision. This is not easy, but one can try, for instance, to estimate how much time is saved or try to quantify the consequences of a wrong decision.

4.2 Alternative Methods

Efficiency is relative, an aspect that is not captured explicitly in the model. One could predict a high value for F for a new method; however, if other methods are available to obtain the same knowledge with lower costs, then it is very likely that the value of n is overestimated. Or, stated simply, if a better solution already exists, nobody will use the newer one. The model is too simple here. The effective value of n itself is not a parameter, but a function of, among others, the perceived benefit by potential users. Also, alternative methods can have higher benefits against the same costs. In economic terms, opportunity costs should also be considered here, and not only accounting costs.

Developers of new visualization methods should be aware of alternative solutions and carefully study their advantages and limitations. New methods are not better by definition. Especially, when existing methods are heavily used in practice, they have been proven to have value. It is often hard to beat straightforward solutions; for instance, in many cases, just using a line graph is the best way to show a time-varying signal.

A defense often heard for the lesser performance of new methods compared to existing ones is that the users have not had enough time to get accustomed to them. In some cases, this might hold, but an equally viable hypothesis is that an existing method is simply better. For instance, just showing a set of objects in a list enables linear scanning, whereas scanning a fancy 2D or 3D display where the objects are distributed over space is much harder [23].

Alternative methods are not limited to visualization methods. For instance, when an automatic method exists to extract the relevant information, visualization is useless. Visualization is not "good" by definition, developers of new methods have to make clear why the information sought cannot be extracted automatically. One reason could be that such automated methods are not guaranteed to give the target result. In this case, integration of automated methods, for instance, from statistics or data-mining, and visualization is a great idea, see, for instance, the work underway and led by Jim Thomas in the visual analytics arena [24].

Fig. 2 shows an example where we used standard methods in a new combination [34]. For the analysis of a time-series of one year, daily patterns were clustered, i.e.,

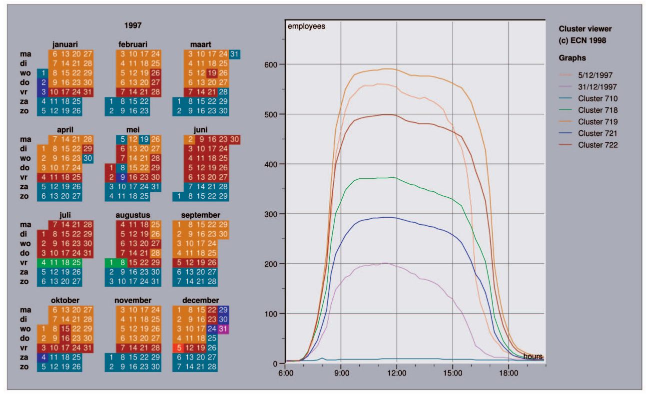

Fig. 2. Visualization of daily patterns [34], an example of the combined use of conventional statistical and graphical methods.

finding similar daily patterns was automated. The results are shown using two conventional representations: Average daily patterns of clusters are shown as graphs and the days per cluster are shown on a calendar. The approach is straightforward and conventional and very effective.

4.3 High Initial Costs

One important reason that new visualization techniques are not used in practice is the high initial cost per user $C_u(S_0)$ involved. Let us consider potential customers for visualization, for instance, researchers doing complex simulations. First, they have to realize that maybe visualization can help them to understand their data. This is not obvious. They already use some methods to extract information from their results in a condensed form. For instance, in molecular dynamic simulations, one typical aim is to derive large scale quantities (temperatures, porosity, etc.) via simulation from the properties on a small scale (size of ions, fields, etc.). Such large scale quantities can be calculated fairly easily from the raw data. Mathematicians working in computational fluid dynamics are often not interested in particular flow patterns, but rather in the convergence of numerical methods and conservation of quantities, which again can be calculated easily and summarized in a few numbers.

The easiest way to visualize data is to use postprocessing capabilities that are integrated into the software used. For instance, commercial packages for computational fluid dynamics or finite element simulation offer such visualization postprocessing capabilities. From a visualization point of view, the techniques offered are far from state of the art: Usually, just options like isosurfaces, color mapping, slicing, streamlines, and arrow plots are provided. But, if these meet the demands of our users, then this is highly cost-effective.

Suppose that this option is not available or is inadequate. The next step is to find alternatives. Our researchers have to get acquainted with possible solutions. Unfortunately, there are no books that present and compare novel visualization techniques (like volume rendering or topology-based flow visualization) at an introductory level. So, they have to study research papers, or search and get in contact with experts in the field.

Following steps are also costly. Maybe they can get a research prototype to work with or else they have to (or let somebody) implement the novel techniques. Often, additional software has to be developed to convert their data to a suitable format.

This all takes much time and effort, while it is unclear whether the new method will indeed solve their problem. Hence, a rational decision is to abstain from this.

There are, of course, ways to share the initial costs with others. A group of researchers can take advantage of an initial investment by one of them. Also, providers of simulation software can be asked to integrate new methods. Visualization does not seem to have a high priority here, however. For an impression of what providers think to be important for their customers, we can have a look at Web sites of companies like MSC or Fluent, and observe that features like advanced simulation capabilities and tight integration are promoted much more than visualization, which is just mentioned in passing labeled as postprocessing.

4.4 Visualization Is Subjective

In the ideal case, one would hope that extraction of knowledge from data is an objective process, in the sense that the outcome does not depend on who performs it, and that the analysis can be repeated afterward by others, with the same outcome. Statistics aims at this, a typical pattern is the use of statistical tests to validate hypotheses on the data. Such tests make assumptions on the data (such as a normal distribution) and have free parameters (like the confidence level), but, furthermore, they do meet the criteria for objectiveness.

Unfortunately, visualization often does not meet this aim. Consider

$$\frac{dK}{dt} = P(V(D, S, t), K).$$

This simply means that the increase in knowledge using visualization not only depends on the data itself, but also on the specification (for instance, which hardware has been used, which algorithm has been used, and which parameters), the perceptual skills of the observer, and the a priori knowledge of the observer. Hence, the statement that visualization shows that a certain phenomenon occurs is doubtful and subjective.

An even harder case is the statement that a certain phenomenon does not occur. I have often spent hours visualizing data, searching for patterns and structure. Sometimes some result could be produced using a particular setting of the parameters, in other cases, I failed to do so. When a visualization does not show clear patterns, it is hard to decide if this is a limitation of the visualization method or an indication that the setting of the parameters was wrong or that the data simply does not contain significant patterns.

This does not mean that visualization is useless. If there are no better alternatives to inspect complex data, visualization has to be used. Another line of defense is that visualization should not be used to verify the final truth, but rather to inspire new hypotheses, to be checked afterward. Part of the subjectiveness can be eliminated by simply showing the visualization to the audience so that they can view and judge it themselves. However, this does not take away the subjectiveness inherent in S as a secondhand viewer; we do not know how sensitive the ultimate visualization is to changes in scales and/or selections of the data.

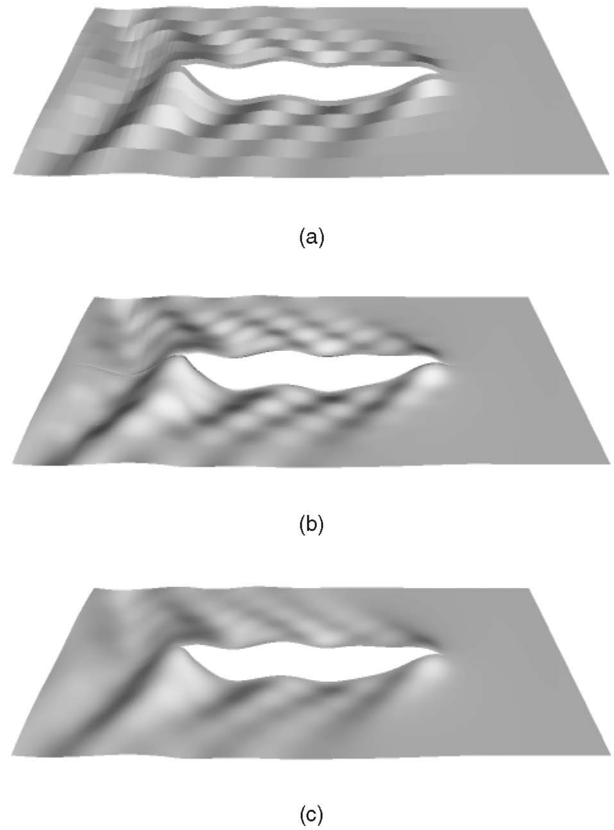

Fig. 3. Wave surface, from top to bottom: (a) bilinear interpolation, (b) cubic interpolation, and (c) cubic approximation. Incorrect interpolation leads to artifacts.

4.5 Negative Knowledge

In the previous section, we considered subjective aspects of visualization. There is another problem: Visualizations can be wrong and misleading. That is, in the terminology introduced here, negative knowledge ($|\Delta K| < 0$) can be produced. Tufte has introduced the *lie-factor* [28], which he defined as the ratio of the size of an effect shown in the graphic to the size of the effect in the data.

Here, it suffices to give an example of my own experience. A long time ago, I visualized the waves produced by ships for a maritime research institute. The data were the result of simulations. Fig. 3a shows the result of bilinear interpolation of the data. I found these results unclear, hence, I decided to use an interpolating spline, thereby smoothing the surface while remaining faithful to the data. Fig. 3b shows clearly that two sets of waves are generated: the standard waves as well as a set of waves orthogonal to this. I proudly presented this discovery to the researcher, who immediately replied that this was physically totally impossible. A much better visualization is shown in Fig. 3c, where an approximating spline is used. The artifacts in the middle image are the result of aliasing. The data orthogonal to the ship are sampled close to the Nyquist frequency; interpolation gives rise to aliases, which corresponding waves have in this 2D case a different direction than the original wave. An interpolating spline smoothes away the high frequencies, but the first aliases survive and give rise to wrong interpretations. I learned from this that interpolation is not by definition better than approximation and also that the judgment of an expert, with a high K_0, is vital for proper interpretation and validation. I never published this and, also, articles on limitations and pitfalls of visualization are scarce. For the advancement of the field, more such reports would be highly beneficial.

4.6 Interaction

Interaction is generally considered as "good." One could advocate the opposite: Interaction should be used carefully and sparingly for two reasons. First of all, as mentioned before, allowing the user to modify S freely will lead to subjectiveness. It is tempting to tune the mapping so that the desired result comes out strongly, but this can be misleading. Also, high customization can make it hard to compare different visualizations. Second, interaction is costly and leads to a high C_e. Rerendering the image after a change of the mapping or the point of view taken often requires a few seconds, viewing it again also. If many options are available to modify the visualization, trying all of them out can take hours. A developer of a new method, therefore, should think carefully about good defaults or automatic ways to set the visualization parameters so that as much knowledge is transferred as possible.

Obviously, in many cases, interaction strongly enhances the understanding of the data. The most important case is simply when the amount of data to be shown does not fit on the screen or is too large to be understood from a single image. In this case, navigation and selection of the data has to be supported to enable the user to interactively explore the data.

Another case for interaction is during development of new methods. I encourage my students to make every aspect of their new methods customizable via user interface widgets so that the total solution space can be explored. However, for the final versions of their prototypes, I recommend offering suitable presets under a few buttons so that a good visualization can be obtained with little effort.

5 Examples

In this section, a number of (classes of) techniques are considered and the cost model is used to explain their adoption in practice.

5.1 Texture-Based Flow Visualization

The use of texture to visualize fluid flow was introduced in the early nineties. The idea is that dense textures enable viewers to judge the direction of flow at all locations of the plane, whereas the standard arrows and streamlines only give discrete and hard to interpret samples. The topic has been studied heavily in the visualization community, a recent nonexhaustive overview [16] has 90 references. The progress made in this decade is substantial. The early Spot Noise technique [29] was an interesting first attempt; in 1993, Cabral and Leedom introduced Line Integral Convolution (LIC), which gave high quality renderings of 2D fluid flow [6]. Many other variations and additions have been presented since then, for instance, to handle flow on surfaces and in volumes and also to boost the performance, using software or hardware acceleration [16]. Nowadays, high quality 2D texture images of flow fields can easily be generated on standard hardware at 50 or more frames per second (Fig. 4, [30]). This seems a success story, but, on the other hand, these methods are not integrated in commercial software, users of Computational Fluid Dynamics (CFD) are typically completely unaware of their existence, let alone that they routinely use them to solve their problems. Here, I use texture-based flow visualization because I am most familiar with it, but, for other classes of methods, such as topology-based flow visualization and feature-based flow visualization, similar patterns seem to apply.

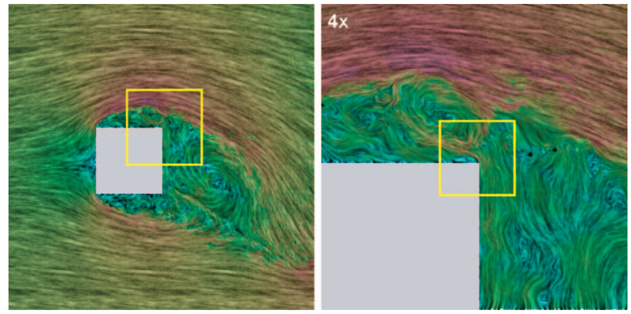

Fig. 4. Image-based flow visualization [30]. Nice visual effects, but rarely used in practice.

How can we explain this? Let us consider the parameters of the cost model. The number of users n is not too great. CFD is vital for some areas, but there are few cases where CFD is routinely used for screening, compared to, for instance, medical applications. The frequency of use m is also not very high. Often, CFD-users spend much time on defining the model, simulations can also take a long time. By then, they are very familiar with their models (high K_0). For the analysis of the results, many alternative options are available, including composite quantities (such as lift of an airfoil) and straightforward cross-sections and arrow plots, with low costs. The use of texture-based visualization incurs at least a high value for C_u (see Section 4.3). The additional ΔK that texture-based visualization offers is unclear. Laidlaw et al. [15] have compared different vector visualization methods. LIC turned out to yield better results for critical point detection, but worse results for other aspects, such as estimation of the angle of the flow. Also, standard LIC does not give the sign of the direction of the flow. Hence, we can have doubts about the value of ΔK. And, finally, it is not clear what the real value is of this ΔK, in the sense that better visualization leads to better decisions. At least, so far, there does not seem to be such a strong need for better visualization methods in the CFD community that they have attempted to integrate these methods into their packages.

5.2 Cushion Treemaps

Also, in the early 1990s, Johnson and Shneiderman introduced the concept of a treemap [11] to visualize large hierarchical data sets. The base algorithm is straightforward: A rectangle is recursively subdivided according to the hierarchical data in such a way that the size of each

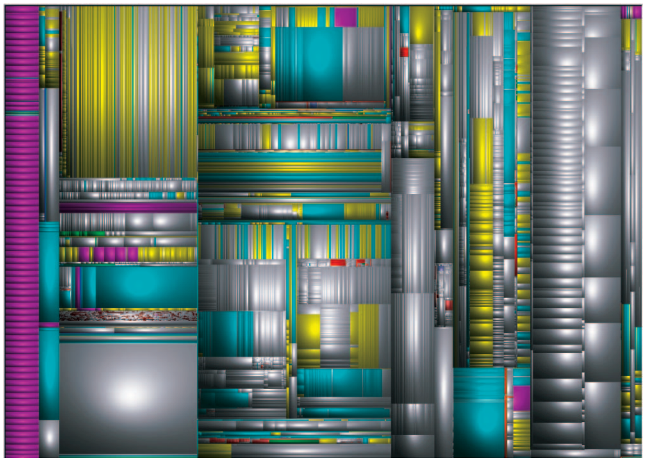

Fig. 5. Visualization hard disk using SequoiaView [1], [32], [33], an example of an application that has found an audience.

rectangle corresponds to the size of each leaf element. In the late 1990s, we proposed using hierarchical cushions to show the underlying hierarchical structure more clearly [32]. We packaged this technique in SequoiaView in 2000 [1], a tool for the visualization of the contents of a hard disk (Fig. 5), and made this publicly available as freeware. Since then, SequoiaView has been downloaded over 500,000 times from our site. Also, it has been distributed three times via CD with the German computer magazine $C't$. This is an example of how visualization has reached an audience.

The economic model helps to explain this result. First, the number of (potential) users is very large, in principle equal to the number of PC users. Typically, such a tool is used several times per year, which is not very high, but not negligible. Alternative solutions for this problem are scarce (SpaceMonger, also using treemaps, is an example) and getting an overview of a hard disk is hard using Windows Explorer.

Information can be derived fairly easily from the visualization. It is easy to spot large files, large directories, and large collections of files. Furthermore, this information is directly valuable to the user: The tool can help (and many users have confirmed this) to delay buying a new hard disk. The action is clear here: removal of files. We offer an option to start up Windows Explorer from SequoiaView to remove files manually. The initial costs per user are low: The tool itself is freeware, it only has to be downloaded and installed. The costs per use case are minimal as well. By default, the tool starts to collect data from the last folder specified and an image is shown automatically. Exploration is easy: Extra information per file can be obtained by hovering the pointer over the rectangles.

In summary, F is high in this case. We would like to think that this is a result of our visualization method, however, the main reasons are probably that our tool meets a real need of real users and that the costs, in all respects, are minimal.

5.3 Presentation versus Exploration

Next, we consider a more general case. The main use cases for visualization are exploration (where users do not know what is in the data) and presentation (where some result has to be communicated to others). It is hard to quantify this, but my impression is that many researchers in visualization consider exploration as the major raison d'etre for visualization, whereas presentation is considered as something additional and not too serious. However, from my own experience, presentation is at least as important as exploration. Many users find videos and images attractive for presenting their work at conferences; the popularity of visualization tools and demos often rises sharply just before open days. For years, I had a pleasant and fruitful cooperation with Flomerics Ltd. in the United Kingdom. This company develops CFD-based tools for, among others, thermal assessment for the electronics industry. My major contact there was the marketing manager, who could use visualization to show the benefits of the CFD tools to managers.

In a broader sense, we can view visualization everywhere. Commercial television uses visualization to show the chemical miracles of new cosmetics, the ingenuity of vacuum-cleaners, and why a new fitness device does not harm your back. Obviously, such visualizations are not the result of visualizing data, but rather the result of the fantasy of advertisement agencies. Selling stuff is not only the realm of business, but also of science itself. Once I heard someone state: *The purpose of visualization is funding, not insight*. We can explain the value of visualization for presentation with the cost model. If we consider the viewers of such visualizations as the users, we see that n is high; K_0 is low (the viewers know little about the topic, so much can be gained); the action to be taken is clear (buy a product, fund research) and has direct economic consequences; the costs for the viewers are low (they just have to view the visualization), although they can be high for the presenter. And, furthermore, for presentation purposes, there are almost no alternative or competing techniques. Pure facts (product X saves Y percent of time) can be convincing, but, to make plausible why and also to show that this is all scientifically sound, visualization is the way to go.

6 Visualization as Innovation

In the preceding sections, a number of questions were raised and various disturbing statements were made. There are many objections that can be made and, in the following sections, some of them are given. One important distinction is which viewpoint to use. Besides a purely technological and economical viewpoint, we can also consider our field as innovation, art, design, or as science. Associated with these are a number of questions, but also routes for future work.

In the cost model, visualization is considered as a technology, to be measured for utility. In this context, research in visualization should lead to new solutions that are useful in practice. Not all the work done is successful in this respect, but we can find a number of reasons to explain this. In this section, I consider our work from an innovation point of view.

First of all, innovation is a merciless process, where only a few new solutions survive. A rule of thumb in product development is that a thousand ideas lead to a hundred prototypes, which lead to ten products, out of which just one is successful. The visualization research community operates at the start of this pipeline, hence it should come as no surprise that not everything finds its way into industry. We can see it as a mission to develop inspiring new ideas, which are a primary fuel in the innovation process.

However, creativity consists of two parts: creation of new ideas as well as selection of the best ones. The first task is fulfilled properly by the visualization community, the second is not. The number of careful validations of visualization methods is still low, although this seems to be improving in the last years.

Second, innovation is a long chain. Eick states [8] that an emerging discipline progresses through four stages: first as a craft practiced by skilled artisans, followed by researchers formulating scientific principles and theories, which are refined by engineers to determine production rules, and, finally, the technology becomes widely available. According to Eick, we have just entered the second phase. Also, developing new methods is quite different from turning these into products and marketing them. There is a gap between our prospective users and the research community. Neither have the proper stimuli to bridge this gap: Individual researchers are too busy increasing the number of publications they are judged on and, for the end-users, implementing new methods is far too costly. The gap can be filled in different ways. One way is via commercial companies (spin-off companies or companies that integrate visualization in their simulation packages), an alternative is via open source and academic development and maintenance, funded by government agencies. VMD [2] is an example of the latter category. As a corollary, if we think that visualization is useful and that this gap causes the lack of adoption, we should aim at increasing funding for more practical activities. Or, we should start companies.

Third, one could state that this is all a matter of time.

It takes time before new ideas penetrate, before new users become aware of new methods, before initiatives are taken to integrate new methods into existing systems. This might be true in some cases, however, it is also too easy to use this as an excuse. It could be used for any method, hence it does not help us to distinguish between good and bad ones.

Fourth, the focus in the model is on large numbers of users and use cases. One can also consider cases where the number of users is small, but where the value of the result is very large. In the books of Tufte, some great cases are presented, such as Snow's discovery of the cause of a cholera epidemic in 1854 [26]. Are there recent cases for new visualization methods? Cases that enabled the researcher to obtain a major scientific insight, to save many lives, or to solve a crucial technological problem? One would like to read more case studies in this spirit that show that visualization is worthwhile and can make a difference.

Finally, one defense is that maybe we are not doing bad at all. A large amount of technology is available and parts of it are routinely used. Also, many disciplines (for instance, mathematics) often do not care about practical usability at all; for some computer science fields that do claim to have practical relevance, it is also sometimes hard to see the adoption in practice. Why should we bother? This notion is explored further in the next section.

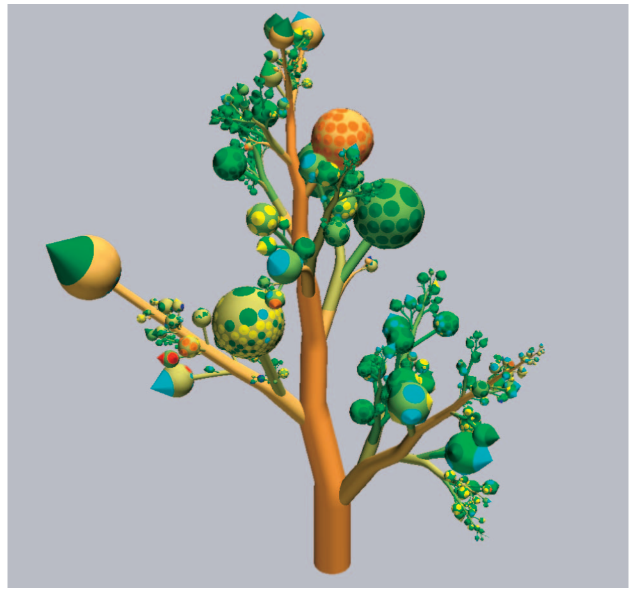

Fig. 6. Botanic visualization contents of a hard disk [13], [33]. Useful or just a nice picture?

7 The Art of Visualization

One could claim that visualization has value in its own right and for its own purposes. One part of this is in the results: Some of the images we produce have a clear aesthetic value. But, the art of visualization can also be found in the ideas, methods, and techniques developed. We can consider ourselves as a group of puzzle solvers and the challenge is to develop new, simple, and elegant solutions, which provide us all with intellectual and aesthetic satisfaction.

This is not a line of defense that can help us to convince our prospective users and sponsors. Nevertheless, I do want to mention it because it can give a powerful thrust (and, obviously, also because the results of this will possibly find applications in the real world). In the early 1990s, I worked hard on using texture for visualization—not to satisfy users, but simply because the puzzle was tough, challenging, and hard to crack. The work of our student Ernst Kleiberg on botanically inspired tree visualization (Fig. 6, [13]) was not driven by user requests, but just an experiment to find out if it could be done at all. At the Information Visualization Symposium in 2004, we got two messages back. Alfred Kobsa found the usability limited, compared to other methods [14]; on the other hand, Stuart Card showed this image in his keynote speech as an example of a nice visualization. Is this a good visualization or not?

Finally, in my own work, I found aesthetic criteria for new methods to be effective guides. Sometimes, each link of the chain from idea, mathematical model, algorithm, implementation, to visual result is clean, simple, elegant, symmetric, etc. It is amazing how much effort is required to reach this. Developing great ideas is simple, rejection of bad ideas takes all the time.

8 Design and Visualization

Visualization methods, techniques, and systems have to be developed. I was trained (a long time ago) as an industrial designer and, in my opinion, our field can take advantage of insights from design in general and visual design in particular.

I studied industrial design engineering at the Delft University of Technology. The aim of this educational program was not to produce industrial designers per se, but rather managers who can overview the product development process and can communicate with all specialists involved. Hence, the curriculum consists of a wide variety of topics, including human factors, mechanical engineering, business science, and graphic design. Also, product design itself fills a large part. I was not talented enough to become a industrial designer myself, but, nevertheless, I have learned much that can be applied in many fields, including visualization.

Besides research, I have developed, alone or in a team, many systems for pre and postprocessing technical applications, especially when I was employed at The Netherlands Energy Research Foundation. I find developing applications for real users with real problems highly rewarding and often just as challenging and difficult as doing research. For some time, I considered such tasks as completely different from research. The primary aims in practical visualization are a useful tool and a satisfied user; the focus in research is on novelty and on publications. However, looking back, I realize that practical jobs have led to new research questions, whereas, now and then, I could insert findings from research into practical systems.

In my present position, among computer scientists, I am often surprised about the view taken on design. The rigid waterfall model is often still dominant and colleagues and students often do not seem to be aware of basic methodological issues and approaches that are commonplace in product design [22]. Hence, I take the opportunity here to communicate and discuss a few of these in short.

The basic design cycle has the following ingredients. First, set up requirements; second, invent a number of solutions; and, finally, match the solutions against the requirements and select the best one. This seems quite obvious; however, I often encounter reports and theses where these basic steps cannot be identified. Just documenting the final result is not enough to understand why it is a solution to a problem. Let us consider each step in more detail.

In the ideal case, requirements should be complete, discriminative, objective, and measurable. This is difficult to achieve in practice, but, nevertheless, should be attempted. Thinking about measurable requirements (rather than "Tool should give insight") leads to a better understanding of the problem and task at hand. If a fuzzy requirement cannot be turned into something measurable: Keep it on the list as a reminder of something important. Sometimes, people consider writing down requirements just as a pure writing exercise. A more positive attitude is to consider it as a creative act in itself. Defining requirements comes down to creating a solution space. Care should be taken that the requirements are not too constraining or too open ended. Requirements vary in importance, weights and priorities should be assigned to them. Requirements are often conflicting. For instance, a tool should be easy to use and also have many features. This is important; it helps us to identify the real challenges and problems.

Generating multiple solutions is important, but often not that difficult. Most important here is the awareness that the first solution thought of is not necessarily the best. Again, trivial and existing solutions should not be forgotten, just like wild and exciting ideas. Simple methods can be used to increase the number of solutions, for instance, combine solutions, use metaphors, identify aspects, and enumerate solutions. Take a large sheet of paper and do not stop until it is full, scribbled with ideas and partial solutions. Accept that others can have good ideas as well. Postpone evaluation. Do not hesitate to write down solutions that are perfect in one respect, but fail on others.

In the last step of the basic design cycle, the best solution is sought. One often-used approach is to put all requirements and solutions in a matrix, put a score in each cell, and calculate a (weighted) score. This gives a first insight into the value of the potential solutions. Often, the outcome is not what is wanted; for instance, one boring solution survives, while a solution that has much more appeal fails for one requirement. Note that the basic design cycle is already iterative. Feel free to add, remove, or reformulate requirements. If "solution should be appealing" is important for the intended audience (or for motivation of the development team), write that down, thereby making it explicit and open for discussion. Also, careful inspection of the table can lead to new solutions. Often, combinations of the good parts of different solutions fulfill the requirements best.

The design process itself should be top-down and cyclic. The basic design cycle is gone through first on a high level. In the first intake conversation, discuss the problem at hand, get a feeling of the scale of the data, the questions asked, and find out which solutions have been tried already and why they do not satisfy the needs. Next, develop sketchy, global, and rough solutions which require many details to be filled in. Each of these details leads to a new design cycle, which, in turn, again leads to new solutions and in even finer details to be worked out. Also, refinements can lead to new insights such that earlier decisions have to be changed. This is natural and it is better to recognize this situation as early as possible and to take adequate action than to drag a dead horse. Furthermore, it is important that, in each overall iteration, all relevant aspects are addressed and are given the right amount of attention and resources. *Premature optimization is the root of all evil* [5], not only in software development, but also often in design in general.

Visualization can be considered as a human-computer interface design task. Methodologies from the HCI field should be adopted. Continuous involvement of end-users should be sought. Mock-ups and prototypes should be used as early as possible. Often, the hard question is what the users really need and want. The prospective users often cannot make this explicit (or they can only state it in terms of something they know), but they can comment on something that they can see and touch themselves. Enjoy criticism on your brilliant ideas; this guides you in the right direction.

8.1 Visual Design

The previous statements on design are very general and broad and, in my opinion, they apply to most, if not all, design problems. Even broader, they can be applied to all decision processes, ranging from buying a new car to definition of a strategy for a department. Also, they apply to visual design.

In my observation, our community has a somewhat uneasy relationship with visual design. On one hand, there is an overall recognition that visual design is very important for our field. When we design a visualization or an interface, we have to decide on shapes, colors, labels, etc. We invite great graphic designers, like Edward Tufte or Nigel Holmes, to our conferences, hear their inspiring talks breathlessly, and, on the work floor, we decide that visual design is the expertise of gods from a different planet and continue to produce visual junk. Maybe I am exaggerating, but, for instance, a quick scan over the posters presented at Visualization 2005 show that often very interesting work is presented in an amateurish and messy way. David Laidlaw has stated the importance of visual design and visual design education on many panels and I strongly support him in this.

There are many misconceptions on visual design. Some of these are:

Design is a matter of talent. Indeed, some people are extremely gifted and produce breathtaking designs. But, on the other hand, designing is a human activity like many others, such as writing, walking, and programming. Everybody has a ceiling, few become great experts, but, by practicing, many can perform reasonably well. For visual design, the law of diminishing returns works in an advantageous way. Because only few in our field have received training in this, there is much that can be learned in a reasonably short time.

Design is a matter of creativity. Generation of alternatives is often not that difficult. Suppose, for instance, we have to decide on a background color or the position of a label: The alternatives can easily be generated. It is much harder to decide which choice is optimal. Real creativity implies thinking out of the box and rejecting some convention that everybody has taken for granted. But, in many cases, just following conventions is good enough for an effective design.

Design is a matter of inspiration. Often, design is just a matter of hard work. Good ideas often do not come out of the blue, they arrive when the problem at hand has been studied carefully and when much effort has been spent on the generation and evaluation of alternatives.

Design is about aesthetics. Visual design is often purely functional. The main concern is communication: Is the message conveyed to the audience in the way intended? Readability and the attraction of the proper amount of attention to the right elements are the issues at stake. When these issues are addressed in a proper way, often an attractive design will result.

Design is a matter of taste. This is maybe the most important misconception: Designers are supposed to create something just because they like it that way and that's where the discussion stops. However, visual design is much more rational than that. Professionals make conscious decisions all the time and they can rationalize why they have done something. Training helps here. When professional designers review a design, they will, for each aspect, ask the reason why a decision was taken and they will not accept "no answer" or just that this is a matter of taste. I have gone through this experience many times. It is humiliating at first, but it does help to raise awareness and make conscious decisions.

Obviously, only a few of us have the real talent to become experts and, for critical visualization applications, the input of professional visual designers cannot be replaced. Nevertheless, there is much that can be gained for us as amateurs and developing insight and skill does help to improve visualizations and interfaces. Reading about design helps; a great introduction, with many visual examples, is Liddel's recent text [17]. Also, analyzing designs is helpful, not only the highlights (such as Beck's 1933 London Tube map), but also day-to-day material, such as magazines, brochures, and books. One can try to see the structure and decisions taken. And, above all, practicing is important, preferably guided by an expert.

Finally, as a comparison, let's consider writing papers. This skill is important for any academic. The aim is not to produce great fiction, but to communicate results clearly. It is difficult and requires hard work and much practice. Everybody will remember when he or she received his very first draft for a paper back from his supervisor, with strong comments in red all over the pages. Often, lengthy discussions and many iterations are needed before the structure and contents are optimal. Now, my statement is that, for visual design, almost exactly the same patterns apply, in contrast to what many in the field believe.

9 Science

Apart from considering visualization as a technology, innovation, an art for its own sake, or as design, we could consider visualization research as a scientific discipline. If there is something like a science of visualization, with what should it be concerned? Loosely defined, a scientific discipline should aim at a coherent set of theories, laws, and models that describe a range of phenomena, have predictive power, are grounded in observations, and that can be falsified.

If we look at the field now, many algorithms and techniques have been developed, but there are few generic concepts and theories. One reason for the lack of fundamental theories is that visualization is intrinsically complex, has many aspects, and can be approached from different perspectives. In terms of the model proposed, visualization can be observed from the point of view of the data D to be visualized, the various solutions proposed (S and V), from the ΔK aimed at, i.e., the task, purpose or discipline for which it is applied, the images I themselves, or from aspects such as perception P or exploration E. Also, developing good visualization solutions is intrinsically a design problem and closed form solutions for the optimalization problem "Given D find V such that ΔK is optimal" cannot be expected.

Nevertheless, we could and should aim at more generic insights, at several levels. First of all, a *descriptive* approach can be pursued further. Methods are analyzed and categorized, leading to taxonomies that show how they relate to and differ from each other. Such taxonomies span the current solution space and can lead to insight as to where new opportunities can be found. Some examples of good overviews are [37], [9], [16], a great example of a taxonomy is given in [4], where a variety of different marching cube style algorithms are brought under one umbrella using computational group theory.

Even if it were only because the field is still developing and overviews are quickly outdated, more work in this area should be encouraged. Taxonomies need not be confined to methods, also taxonomies on different kinds of data and especially on different types of knowledge that are relevant for end users are useful. Duke et al. [7] give strong arguments for more efforts in this area. Development of an ontology of visualization (as a next level above taxonomy and terminology) gives a more rigorous foundation of the field and is beneficial for collaboration, composition, preservation, and education.

Second, *evaluation* and *validation* are important. Assessment of the effectiveness and efficiency of different methods and techniques is vital from a technological point of view (which method to use), but also as a base for more generic statements on visualization. A science of visualization should be empirical, in the sense that concrete measurements of the phenomena studied are done, which, in our case, concern people making and watching images that depict data. Tory and Möller [25] give a good overview of the current status of the use of human factors research in visualization and identify areas for future research.

Third, in line with the previous, we should ultimately aim at *generic results* (models, laws) that enable us to understand what goes on and to predict why certain approaches do or do not work. In the end, explanations should be based on properties of the environment of visualization, especially the end users. The value of visualization is ultimately determined by their perceptual abilities, their knowledge on the data shown, the value they assign to various insights, and the costs they are willing to spend.

A great example of such a generic result is Ware's statement that a 2 1/2 D attitude toward the design of visualization should be adopted [35], with the strong argument that our human visual system is also 2 1/2 D. Ware's book on Information Visualization [36] is a rich source of insights on perception and how these can be used to improve visualization, Tufte gives many useful guidelines and recommendations in his books [28], [26], [27]. However, many of these are not quantitative and, also, do not explain how to handle conflicting requirements. One operational and practical criterium for guidelines is that they should allow for automated implementation such that the user gets a good, if not optimal view on the data without costs. The early work of Mackinlay [19] on automated generation of visualizations is great in this respect, and it is surprising that the state of the art in this area does not seem to have advanced much further since then.

Finally, *methodological* issues have to be studied further. This concerns questions such as how to design visualizations and how to measure and evaluate the effectiveness of various solutions. And, also, how to assess the value of visualization in general.

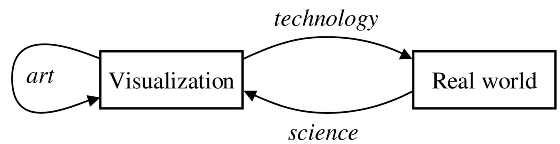

Fig. 7. Views on visualization.

10 Conclusion

In the preceding sections, I have considered visualization from multiple perspectives. None of these is superior. One view is to consider visualization purely from a technological point of view, aiming for effectiveness and efficiency. This requires that costs and benefits be assessed. The simple model proposed enables us to get insight into various aspects of visualization and also to understand why certain classes of methods are successful and others are not. Another view is to consider visualization as an art, i.e., something that is interesting enough for its own sake. The use of insights from design can help us to improve visualizations. Finally, a view on visualization as an empiric science was discussed.

Obviously, these different views, schematically depicted in Fig. 7, are strongly related and results from one view can stimulate work according to the other views. Finally, each view that is adopted does imply playing a different game, and, if we want to win, we should play those games according to their own rules: aim for provable effectiveness and efficiency, aim for elegance and beauty, and aim at generic laws with predictive power.

Acknowledgments

Many colleagues, including the anonymous reviewers, have given constructive feedback on earlier versions of this paper; their insights have helped me to add new points of view. Also, their positive support shows that there is room for critical reflection, which is a strong signal that our field is healthy and vital. I thank Hannes Pretorius for proofreading and giving valuable advice on improving the readability of this version.

References

[1] www.win.tue.nl/sequoiaview, 2006.
[2] www.ks.uiuc.edu/Research/vmd, 2006.
[3] *Position Papers NIH/NSF Proc. Fall 2004 Workshop Visualization Research Challenges,* 2004, visual.nlm.nih.gov/evc/meetings/vrc2004.
[4] D.C. Banks and S. Linton, "Counting Cases in Marching Cubes: Toward a Generic Algorithm for Producing Substitopes," *Proc. IEEE Conf. Visualization*, pp. 51-58, 2003.
[5] J. Bentley, *Writing Efficient Programs.* Prentice Hall, 1982.
[6] B. Cabral and L.C. Leedom, "Imaging Vector Fields Using Line Integral Convolution," *Computer Graphics,* vol. 27, pp. 263-270, 1993.
[7] D.J. Duke, K.W. Brodlie, D.A. Duce, and I. Herman, "Do You See What I Mean?" *IEEE Computer Graphics and Applications,* vol. 25, no. 3, May/June 2005.
[8] S.G. Eick, "Information Visualization at 10," *IEEE Computer Graphics and Applications,* vol. 25, no. 1, pp. 12-14, Jan./Feb. 2005.

[9] I. Herman, G. Melançon, and M.S. Marshall, "Graph Visualization in Information Visualization: A Survey," *IEEE Trans. Visualization and Computer Graphics,* vol. 6, no. 1, pp. 24-44, Jan./Mar. 2000.

[10] B. Hibbard, "Top Ten Visualization Problems," *PSIGGRAPH Computer Graphics Newsletter,* vol. 33, no. 2, 1999.

[11] B. Johnson and B. Shneiderman, "Treemaps: A Space-Filling Approach to the Visualization of Hierarchical Information Structures," *Proc. IEEE Conf. Visualization '91,* pp. 284-291, 1991.

[12] C. Johnson, "Top Scientific Visualization Research Problems," *IEEE Computer Graphics and Applications,* vol. 24, no. 4, pp. 13-17, July/Aug. 2004.

[13] E. Kleiberg, H. van de Wetering, and J.J. van Wijk, "Botanical Visualization of Huge Hierarchies," *Proc. IEEE Symp. Information Visualization (InfoVis '01),* pp. 87-94, 2001.

[14] A. Kobsa, "User Experiments with Tree Visualization Systems," *Proc. IEEE Symp. Information Visualization (InfoVis '04),* pp. 9-16, 2004.

[15] D.H. Laidlaw, R.M. Kirby, J.S. Davidson, T.S. Miller, M. da Silva, W.H. Warren, and M.J. Tarr, "Quantitative Comparative Evaluation of 2D Vector Field Visualization Methods," *Proc. IEEE Conf. Visualization '01,* pp. 143-150, 2001.

[16] R.S. Laramee, H. Hauser, H. Doleisch, B. Vrolijk, F.H. Post, and D. Weiskopf, "The State of the Art in Flow Visualization: Dense and Texture-Based Techniques," *Proc. Computer Graphics Forum,* vol. 23, no. 2, pp. 203-221, 2004.

[17] W. Lidwell, K. Holden, and J. Butler, *Universal Principles of Design.* Rockport Publishers, 2003.

[18] B. Lorensen, "On the Death of Visualization," *Position Papers NIH/NSF Proc. Fall 2004 Workshop Visualization Research Challenges,* 2004.

[19] J.D. Mackinlay, "Automating the Design of Graphical Presentations of Relational Information," *ACM Trans. Graphics,* vol. 5, no. 2, pp. 110-141, 1986.

[20] *Visualization in Scientific Computing,* B.H. McCormick, T.A. DeFanti, and M.D. Brown, eds., New York: ACM SIGGRAPH, 1987.

[21] T.-M. Rhyne, B. Hibbard, C. Johnson, C. Chen, and S. Eick, "Can We Determine the Top Unresolved Problems of Visualization?," *Proc. IEEE Conf. Visualization '04,* pp. 563-565, 2004.

[22] N.J.M. Roozenburg and J. Eekels, *Product Design: Fundamentals and Methods.* Wiley, 1995.

[23] B. Shneiderman, comment during session at InfoVis, 2004.

[24] *Illuminating the Path: Research and Development Agenda for Visual Analytics,* J.J. Thomas and K.A. Cook, eds., IEEE, 2005.

[25] M. Tory and T. Möller, "Human Factors in Visualization Research," *IEEE Trans. Visualization and Computer Graphics,* vol. 10, no. 1, pp. 72-84, Jan./Feb. 2004.

[26] E.R. Tufte, *Envisioning Information.* Graphics Press, 1990.

[27] E.R. Tufte, *Visual Explanations: Images and Quantities, Evidence and Narrative.* Graphics Press, 1997.

[28] E.R. Tufte, *The Visual Display of Quantitative Information,* second ed. Graphics Press, 2001.

[29] J.J. van Wijk, "Spot Noise: Texture Synthesis for Data Visualization," *Computer Graphics,* vol. 25, pp. 309-318, 1991.

[30] J.J. van Wijk, "Image Based Flow Visualization," *ACM Trans. Graphics,* vol. 21, no. 3, pp. 745-754, 2002.

[31] J.J. van Wijk, "The Value of Visualization," *Proc. IEEE Conf. Visualization '05,* pp. 79-86, 2005.

[32] J.J. van Wijk and H. van de Wetering, "Cushion Treemaps," *Proc. 1999 IEEE Symp. Information Visualization (InfoVis '99),* pp. 73-78, 1999.

[33] J.J. van Wijk, F. van Ham, and H.M.M. van de Wetering, "Rendering Hierarchical Data," *Comm. ACM,* vol. 46, no. 9, pp. 257-263, Sept. 2003.

[34] J.J. van Wijk and E. van Selow, "Cluster and Calendar-Based Visualization of Time Series Data," *Proc. 1999 IEEE Symp. Information Visualization (InfoVis '99),* pp. 4-9, 1999.

[35] C. Ware, "Designing with a 2 1/2D Attitude," *Information Design J.,* vol. 10, no. 3, pp. 255-262, 2001.

[36] C. Ware, *Information Visualization: Perception for Design,* second ed. Morgan Kaufman, 2004.

[37] P.C. Wong and R.D. Bergeron, "30 Years of Multidimensional Multivariate Visualization," *Scientific Visualization—Overviews, Methodologies, and Techniques,* G.M. Nielson, H.H. Hagen, and H. Mueller, eds., IEEE CS Press, 1997.

Jarke J. van Wijk received the MSc degree in industrial design engineering in 1982 and the PhD degree in computer science in 1986, both with honors. He worked at a software company and at The Netherlands Energy Research Foundation ECN before he joined the Technische Universiteit Eindhoven in 1998, where he became a full professor of visualization in 2001. He is a member of the IEEE, ACM SIGGRAPH, and Eurographics. He was paper cochair for IEEE Visualization in 2003 and 2004 and is paper cochair for IEEE InfoVis in 2006. His main research interests are information visualization and flow visualization, both with a focus on the development of new visual representations.