

A Model of Spatial Directness in Interactive Visualization

Stefan Bruckner, Tobias Isenberg, Timo Ropinski, and Alexander Wiebel

Abstract—We discuss the concept of directness in the context of spatial interaction with visualization. In particular, we propose a model that allows practitioners to analyze and describe the spatial directness of interaction techniques, ultimately to be able to better understand interaction issues that may affect usability. To reach these goals, we distinguish between different types of directness. Each type of directness depends on a particular mapping between different spaces, for which we consider the data space, the visualization space, the output space, the user space, the manipulation space, and the interaction space. In addition to the introduction of the model itself, we also show how to apply it to several real-world interaction scenarios in visualization, and thus discuss the resulting types of spatial directness, without recommending either more direct or more indirect interaction techniques. In particular, we will demonstrate descriptive and evaluative usage of the proposed model, and also briefly discuss its generative usage.

Index Terms—Visualization, direct interaction, human-computer interaction (HCI).



1 INTRODUCTION

THE ability to interactively explore a visual representation is a core aspect of all visualization systems [61, Ch. 1.7]. It facilitates the inspection of data by users and thus increases their understanding of the data “by orders of magnitude.”¹ Some researchers even argue that visualizations cannot exist without interactive exploration abilities [15, Ch. 1]. Yet, the interaction with visualizations has repeatedly [36], [43] and recently [28] been named as still being one of the field’s major research challenges.

Naturally, many different ways of interaction with data visualizations exist (e. g., [96]), each with its own characteristics. In this article we focus on those interaction techniques that provide some sort of spatial input to control the outcome of an interaction. For example, we are concerned with interaction techniques such as view manipulation and zooming, navigation, object and cutting plane manipulation, picking, selection by means of subspace specification, selection by means of brushing, and line/curve placement and manipulation. One of the most frequently used ways to classify these interaction techniques is whether they are *direct* (e. g., [12], [79]) or *indirect* (e. g., [50], [53]). For example, direct manipulation [74], [75] is typically described as a property that is beneficial for interactive systems, in visualization and otherwise. Unfortunately, the terms referring to the *directness*² of the spatial interaction are largely used by intuition and without a clear definition. For example, “direct-touch” interaction with a visualization could be seen either as direct if one explores, for instance, a 2D scatterplot on a touch table. On the other hand, it could also be seen as indirect if one navigates through the (projected)

visualization of a 3D point cloud on a touch-sensitive display, even though in both cases a form of input is being used that is said to be “direct” (e. g., [48], [54], [59], [64], [72], [73], [78], [92]). These examples demonstrate that we need a clear definition and a conceptual model of spatial interaction directness, in particular for the domain of visualization because it deals with multiple spaces and mappings between them.

In the past, researchers in the human-computer interaction (HCI) community have discussed the notion of directness of interaction and manipulation in general. Hutchins et al. [38], in particular, identified *distance* and *engagement* as two phenomena that give rise to the feeling of the degree of directness with respect to a manipulation interface, as well as two forms of distance that are relevant—*semantic* and *articulatory* distance. They also discuss advantages *as well as* disadvantages for *both* direct and indirect manipulation. On the one hand, this discussion from HCI can be applied to our field of visualization. On the other hand, however, it is not specific enough for the multiple pipeline steps (e. g., [15], [35]) and spatial reference frames (e. g., [57]) of data, graphical representation, interaction input, and interaction effect/result.

In this article we thus propose a model that, instead of assigning the generic tag “direct” or “indirect” to a given interaction or to an entire visualization system, discusses the different mappings between the involved spaces and describes the directness of each of these mappings independently. With our model we hope to assist both researchers and practitioners in better understanding the implications of directness and indirectness when building interactive visualization systems, as both can have specific benefits and costs.

2 PRACTICAL EXAMPLES

To better illustrate the problem we are addressing, to explain the different frames of reference, and to support the following introduction of our new model, we begin by describing two practical

- Stefan Bruckner is with the University of Bergen, Norway. E-mail: stefan.bruckner@uib.no.
- Tobias Isenberg is with Inria and Université Paris-Saclay, France. E-mail: tobias.isenberg@inria.fr.
- Timo Ropinski is with Ulm University, Germany. E-mail: timo.ropinski@uni-ulm.de.
- Alexander Wiebel is with Hochschule Worms University of Applied Sciences, Germany. E-mail: wiebel@hs-worms.de.
- All authors contributed equally to the article and are named alphabetically.

Manuscript received 16 May 2017; revised 16 May 2018; accepted 1 June 2018. Recommended for acceptance by M. A. Otaduy. Digital Object Identifier no. 10.1109/TVCG.2018.2848906

1. <https://youtu.be/Sua0xDCf8MA?t=309>, Miriah Meyer, 2011.

2. We use the term *directness* in its generic form to refer to both attributes, i. e., for both *direct* and *indirect* interaction.

interactive visualization scenarios: one that deals with spatial data and a second one based on non-spatial data.

Spatial data scenario. Consider the task of specifying a 3D location in a visualization that is based on direct volume rendering³ (DVR; e. g., [70]). In this scenario, the data (e. g., a medical CT or MRI scan or a physical simulation) is defined on a grid in 3D (data) space. Next, a transfer function defines how the captured or simulated scalar data values are converted into optical properties, in particular opacity and color. In addition, the data volume is arranged with respect to a virtual camera such that the human observer sees a projection of the optical properties, governed by the DVR process. Typically, this projection is shown on a regular monoscopic PC monitor, but alternatives such as stereoscopic rendering (using one projection per eye; e. g., [23], [85]) and true 3D rendering (for which no algorithmic projection and composition are needed; e. g., [32], [63]) exist as well. For all these cases, however, the observers interpret what they see to understand the depiction as a representation of the 3D data, and to explore it accordingly.

Based on this scenario, we have several possibilities to interactively specify a 3D position in the data/visualization. For visualizations displayed on a monoscopic screen of a typical workstation, users can point to a position in the projection for which then the 3D location in the data has to be deduced algorithmically from the specified 2D position. The act of pointing itself can be facilitated by moving a mouse on the table next to the PC which controls a cursor, or by placing a finger on a 2D (touch) screen that shows the projection. Alternatively, the 3D location could also be specified by means of dedicated sliders that are part of the interface, each slider controlling one of the three spatial coordinates. In case of stereoscopic rendering, the 3D position can also be taken from a tracked 3D interaction device, held either in front of (e. g., CAVE [23]) or behind (e. g., HMD [14]) the stereoscopic projection. Alternatively, in the same case, the hands of the interacting person can be tracked in 3D space (e. g., [63]).

Non-spatial data scenario. As a second example illustrating the problem we are addressing, we examine brushing-and-linking. Consider a linked multi-view visualization [5], [11], [94], where—for the purpose of this example—a parallel coordinates view [39] is linked with a 2D scatterplot view to facilitate the analysis of multi-dimensional data. The user shall be able to apply brushing to any of these views in order to update the other, such that only a subset of the data is visualized. The data can be considered to be provided by means of a table, where the columns represent the individual data attributes, and the rows the actual data items. A single data item is then mapped to a polyline that intersects the individual coordinate axes at positions determined by the corresponding attribute values. For a scatterplot, two data attributes are chosen and mapped to a Cartesian coordinate system where individual data items are typically represented as points. While, as in our first example, different output displays would also be possible for this scenario, we will only consider a conventional 2D display screen as the data and the views are inherently 2D. Based on this scenario, the user would have several possibilities to perform the brushing interaction. For both views the user can use the mouse to specify a region, to filter the data. In the case of parallel coordinates, this region would result in a selection of data items as well as data attributes, depending on which parts of the view are inside the region. In the 2D scatterplot, the selection region could either contain a subset

of the data items, or a data range for one of the two depicted data attributes, in case it is performed on an axis. In addition to this region-based filtering, also more focused selection strategies are possible. In both representations, individual data items or an axis can be selected with the mouse.

Common observations. The two examples described above show that there are several spaces—and mappings between them—involved in an interactive visualization. To be able to understand the directness of the underlying interaction, all these mappings need to be considered. For this purpose we may think of these mappings as having a distance. A mapping distance of zero would mean that two spaces are identical, while a large distance can be understood as the spaces to be mapped being considerably different. Hutchins et al.’s [38] model does not cover these distances between the mentioned spaces. Instead, Hutchins et al. use the term *distance* only to refer to the “relationship between the task the user has in mind and the way that task can be accomplished via the interface” [38]—their *semantic* and *articulatory distances* refer to expression of the user’s intention/goal with the help of the provided interface. Ward and Yang [91], in contrast, have previously discussed interaction spaces in the context of the visualization of abstract data including screen space, data value space, data structure space, attribute space, (3D) object space, and visualization structure space. Their discussion, however, only covers manipulations within these spaces (navigation, selection, and distortion), not the relationship between different spaces. We thus need to extend the discussion of spatial directness to the different distances for interactive visualization scenarios which exist, in one form or another, for virtually any type of interactive visualization. As this discussion results in a new model, we first discuss related visualization models and theories.

3 RELATED WORK

The notion of *directness* of interaction or manipulation and, in particular, the term *direct manipulation* in the context of computers has been around for over forty years [75]. Thus, it is not surprising that others [38] have examined this notion, have identified different types of directness, have proposed quantifications of the cognitive distance in direct manipulation interfaces [29], [30], and have discussed the advantages and disadvantages of such interfaces [45]. The goal of our model is to characterize the role of spatial interaction in visualization systems, with a particular focus on arriving at a more formal definition of the notion of directness. Thus, we want to develop a thorough understanding of the factors involved in this notion and how it potentially affects the fidelity of a visualization system, thus adding a discussion of issues with respect to the (mental and physical) mapping between different *spatial* spaces—these issues are orthogonal to models and concepts in the context of Norman’s [62] gulfs of execution and evaluation [38] as well as his seven stages of action [52] or the related user action framework [34]. Our work is inspired by two main lines of investigation discussed below.

Theoretical visualization frameworks. Bertin’s highly influential work [6] was one of the first to systematize different data types and suitable retinal variables. Since then, several conceptual models for visualization have been proposed, such as Silver’s object-oriented model [77], the P-set model of visual exploration by Jankun-Kelly et al. [41], the distributed cognition model by Liu et al. [56], or the predictive data-centered theory by Purchase et al. [66]. Vickers et al. [89] employed category theory and semiotics

3. The term *direct* in “direct volume rendering” relates to how data is transformed into a visualization; it does not refer to the interaction with it.

to facilitate an improved understanding of visualizations in practice. The conceptual framework of visual multiplexing by Chen et al. [18] facilitates the study of different mechanisms for integrating multiple pieces of visual information. Information theory has been shown to be a fruitful framework for conceptualizing the visualization process, as demonstrated by Chen and Jänicke [17], Xu et al. [95], Wang and Shen [90], and, most recently, Chen et al. [16]. Demiralp et al. [25] proposed visual embedding as a constructive model for visualization and demonstrated it using applications such as neural tract coloring and the evaluation of tensor glyphs. Kindlmann and Scheidegger [47] presented an algebraic framework for arguing about the effectiveness of visual representations. Our work builds on these foundations, but specifically addresses how discrepancies between the interaction and visualization mappings may affect the visualization process.

Interactive visualization models. The most common view of the visualization process is that of a pipeline, and Upton et al. [86] were the first to propose the still ubiquitous abstraction consisting of the filtering, mapping, and rendering stages. In his seminal work, van Wijk [87], [88] presented a model which incorporates perception and cognition as well as interactive exploration in order to argue about the value of visualization in an economic sense. Green et al. [31] further revised this model to make the interaction between perception, knowledge, and exploration more explicit. In their work, they also argue for the central role of direct interaction, but do not provide a more detailed explanation of what this entails. Keim et al. [44], focusing on the visual analytics process, consider two interacting loops for data mining and visual data exploration. Jansen and Dragicevic [42] proposed an interaction model designed to incorporate physical visualizations. They discuss directness in the context of *instrumental manipulation* as introduced by Beaudouin-Lafon [4], distinguishing between the *degree of compatibility*, the *degree of indirection*, and the *degree of integration* that a manipulation device may exhibit. Cordeil et al. [22], finally, discuss the mapping of physical interactions to the data space. In this paper, we develop a more detailed model of interaction directness and its dimensions and show how it can be used to explain the perceived differences among particular scenarios.

4 A MODEL OF SPATIAL INTERACTION DIRECTNESS

Based on the observations made in Sec. 1 and the examples discussed in Sec. 2, we now derive our model of directness for interactions in visualization. Our approach is inspired by the work of Kindlmann and Scheidegger [47] who devised an algebraic model of visualization design to guide the design of visual encodings and to characterize their shortcomings. While their model focused on the process of visual mapping in terms of faithful representation of the underlying data by analyzing data symmetries and visualization symmetries, we aim to illuminate correspondences between visualization and interaction mechanisms. Hence, while we also use a diagrammatic representation of our model, its semantics and goals differ.

As adumbrated in Sec. 2, our model describes the relationship between several spaces that are involved in an interactive visualization system.

- The *data space* \mathcal{D} is the space the data to be visualized lives in. For a volumetric data set or 3D scattered points this would be \mathbb{R}^3 .
- The *visualization space* \mathcal{V} represents the space a visualization (i. e., the result of the visual mappings applied

to the data) lives in. In many cases, this space will be two-dimensional.

- Analogously, the *interaction space* \mathcal{I} represents the space in which interaction schemes operate. For mouse interactions, for instance, this would also be a two-dimensional space.

As we want to also capture the physical aspects of visualization and interaction and how they relate to their algorithmic counterparts, we consider two additional spaces:

- The *output space* \mathcal{O} , which represents the space the graphical output representations are generated in. This space will most frequently correspond to the positions of the pixels on a desktop monitor, but it can, e. g., also be three-dimensional for true 3D displays [32].
- The *manipulation space* \mathcal{M} represents the space the input modalities operate in. For mouse movements on tables or touch input this would be a planar two-dimensional space, while it would be a three-dimensional space for a tracked 6DOF-device.

The distinction between \mathcal{V} and \mathcal{I} on the one side and \mathcal{O} and \mathcal{M} on the other side is meaningful as the results of the same visualization technique or interaction mechanism may be presented on or controlled by different types of physical devices operating in spaces of different scale and dimensionality.

- Finally, the *user space* \mathcal{U} represents the model generated in the user's mind facilitated by an interactive visualization system. It hence represents how the user ultimately understands the data, and hence will be greatly affected by any transformation that the data undergoes, how different interaction mechanisms affect the visual representation of the data, based on the involved perceptual and cognitive processes.

To understand the relationships between these spaces we can now examine the mappings between them as represented below. These mappings are initially based on relations between single points in the respective spaces. They can be also applied, however, to sets of points in parallel to account, for example, for interactions like selecting regions or sub-spaces or manipulating other objects such as cutting planes or selection brushes that have a non-zero extent or that require the specification of more than a single point.




Fig. 1. Schematic view of our model of interaction directness.

The *visualization mapping* v is the mapping from the data space \mathcal{D} to the visualization space \mathcal{V} , i. e., it describes how the visual encoding transforms data values and their locations into visual representations. v can be thought of as the *entire visualization pipeline*, and must not be confused with the single stage of the same name in some visualization pipeline descriptions. The *interaction mapping* i , on the other hand, describes the mapping from the interaction space \mathcal{I} to \mathcal{D} . The *output mapping* o describes the mapping from the visualization space \mathcal{V} to the output space \mathcal{O} , i. e., the transformations involved in the process of converting a

visualization into physical visual stimuli on an output device, e. g., colored pixels on a monitor. Likewise, the *manipulation mapping* m is the mapping from the manipulation space \mathcal{M} to the interaction space \mathcal{I} , e. g., the conversion of mouse movements to 2D cursor positions. Finally, we can also regard *the user's perception* u_p as a mapping from the output space \mathcal{O} to the user space \mathcal{U} . This mapping is meant to capture all transformations involved in the perceptual and cognitive processes translating from visual stimuli to the user's mental model. For instance, it incorporates the fact that, as mentioned before, a pair of projections of 3D data will be perceived stereoscopically if presented to the user's eyes in a suitable manner (e. g., using shutter glasses). In the other direction, we can regard *the user's action* u_a as a mapping from the user space \mathcal{U} to the manipulation space \mathcal{M} . This mapping also underlines the dependence between visualization and interaction, as users' actions will clearly be influenced by their mental model which is, in turn, affected by the visualization and output mapping. The formation of a mental model from the presented data can thus be thought of in terms of the function composition (to be read from right to left) as $u_p \circ o \circ v$, while a user's interaction with the data based on a conscious action can be written as $i \circ m \circ u_a$. While u_p and u_a are, in principle, specific to each individual user and may exhibit considerable variations across the population,⁴ general knowledge from perception and cognition research allows us to at least consider some aspects of these mappings as constant. In any case, as visualization researchers or practitioners we can affect the design of the technical aspects of a system such that they represent the best match for a given set of perceptual and cognitive constraints. Hence, from a visualization designer's point of view, our aim is typically to select an appropriate v and i as well as—although somewhat less frequently considered— o and m based on a given u_p and u_a to help users generate a mental model which best represents the underlying data. Our discussion thus mostly focuses on the diagram's left and center parts. The main purpose of our model is to provide a means for understanding the correspondences between visualization and interaction and their implications, but does not aim to address the suitability of a particular visualization mapping, which is a separate issue that is addressed in the work of Kindlmann and Scheidegger [47], for instance.

To understand the implications of our proposed model, we can now further examine the direct mappings between the visual and the interactive spaces. First, the *visualization-interaction mapping* vi , which is the mapping between visualization space \mathcal{V} and interaction space \mathcal{I} , characterizes the degree of mismatch between v and i . If $i = v^{-1}$, i. e., the interaction mapping is the inverse of the visualization mapping, then $vi = \text{id}_{\mathcal{V}}$ which corresponds to the highest degree of *visualization-interaction directness*. As an example, consider the visualization of a 2D scalar field as a color-mapped image. The visualization mapping would then include a 2D viewport transformation T which maps grid points to viewport coordinates. A 2D picking interaction on the output image is the inverse transformation T^{-1} . If, on the other hand, we consider the visualization of a 3D scalar field using direct volume rendering, the mapping is generally not invertible as several grid points may be covered by a single pixel. Such a case would exhibit a lower visualization-interaction directness.

For the second major aspect of directness, we consider the

output-manipulation mapping om , i. e., the mapping between the output space \mathcal{O} and the manipulation space \mathcal{M} . If $vi = \text{id}_{\mathcal{V}}$, and the manipulation mapping is the inverse of the output mapping, i. e., $m = o^{-1}$, then the two spaces collapse with $om = \text{id}_{\mathcal{O}}$. This means that the achievable level of om directness is constrained by the vi directness. This constraining happens, for instance, on a 2D touchscreen interface displaying 2D data where the user is, in principle, able to directly manipulate the visual elements on screen. A similar setup, but using a common desktop monitor and a mouse, still exhibits high vi directness, but here \mathcal{O} and \mathcal{M} remain separate, potentially imposing a higher cognitive load on the user to integrate them. In fact, as discussed by López et al. [57], human task performance is affected by the complexity of this mapping: A rotation around a single axis (e. g., for typical vertical desktop monitors where the mouse is manipulated on a horizontal surface in front of the screen), for instance, is much easier to handle than if a second rotation axis is involved (e. g., if the monitor is, in addition, rotated by 90 degrees around the viewing axis).

It can be seen that the separate discussion of v and o as well as m and i does not mean that these mappings are independent. Quite on the contrary, a given choice of output or input modalities will clearly limit the meaningful number of visualization and interaction mappings, and the converse is also true. For instance, an inherently two-dimensional visualization will generally not benefit from a stereoscopic output device just as a two-dimensional interaction mapping is not aided by a 3D spatial input modality. On the other hand, if the data space has a higher dimensionality, stereoscopic visualization and interaction can be beneficial. The work of Alper et al. [1] is a good example for this. Their study showed that for non-planar graphs, which exhibit many edge crossings in conventional 2D embeddings, user performance can be significantly enhanced by using stereoscopic highlighting technique where selected nodes and edges are lifted from the plane in order to resolve ambiguities.

4.1 Thoughts on Measuring Directness

While it is possible to construct simple examples where a high degree of directness is achievable, i. e., the involved visualization and output mappings are indeed invertible, for most visualization scenarios this is not possible. In particular, technical constraints often prevent the realization of invertible mappings, especially if the data dimensionality is higher than 2. For instance, a visualization mapping of $v: \mathbb{R}^3 \rightarrow \mathbb{R}^2$ will generally not be invertible. If, however, we are primarily interested in a 2D manifold embedded in 3D space, an invertible mapping is possible. We can often use additional information to make a particular mapping a *quasi-inverse* of its counterpart, i. e., to closely approximate the inversion by exploiting knowledge about its algorithmic characteristics. For instance, while it is generally not possible to reconstruct a single 3D position from a 2D pixel location of a semi-transparent DVR image because multiple 3D samples are blended into a single pixel, Wiebel et al. [93] identified the object along the ray profile which maximally contributes to the pixel's appearance to derive such a position.

This raises the question of whether it is meaningfully possible to quantify the directness of a particular mapping. So far, we have seen that the maximum directness both in terms of the visualization-interaction mapping and in terms of the output-manipulation mapping is achieved if the involved spaces collapse, and we have shown that in some cases it may be possible to increase the degree of directness by using additional information. The ways in which *indirectness* can be increased are countless.

4. For instance, some individuals are incapable of stereoscopic vision or suffer from color blindness; likewise, not all types of manipulation devices or mechanisms may be equally accessible to all users due to physical limitations.

Let us now assume an importance function $I(x) \in [0, 1]$ which quantifies degree of interest in each data element x , for a given interaction technique and intent as well as visual mapping. A measure for the *indirectness* of vi (and, analogously, om) can then be defined as

$$\Delta_{v,i} = \int I(x) |x - i(vi(v(x)))| dx. \quad (1)$$

A visualization mapping v maps a data location x into visualization space (e. g., a 2D plane), while the interaction mapping i maps a location in interaction space (e. g., also a 2D plane) back into the data space. Ideally, the distance between the original data space location and a location that has undergone both mappings $|x - i(vi(v(x)))|$ is zero. Hence, by accumulating and weighting the differences, we obtain a measure that characterizes the behavior of these mappings over all relevant data space locations. Our visualization-interaction (or output-manipulation) mapping is maximally direct for all relevant data elements if $\Delta_{v,i}$ is zero. For non-zero but sufficiently small values, we can consider i to be a quasi-inverse of v , while large values correspond to a lower degree of directness. This discussion assumes a visualization-centric point of view where we aim to quantify the degree of indirectness of an interaction mapping for a given visualization mapping. It is equally correct, however, to say that v is a quasi-inverse of i —in particular in the case of o and m the choice of reference may be less clear in some cases. As the value of Δ , due to the countless ways indirectness can be achieved, is unbounded in general, we can see that there is no distinct boundary between “direct” and “indirect,” but rather a gradual transition towards increasingly indirect interaction. The importance function $I(x)$ represents the fact that, even for mappings that are clearly not invertible in a mathematical sense, a suitable quasi-inverse may be found if additional knowledge is incorporated. While we chose to formulate Eqn. 1 in a continuous manner, in practice we almost always will have to use a numerical approximation and hence will draw samples from the data space and sum up the weighted distances. Similarly, when dealing with discrete data, the integral in the equation can be interpreted in the sense of discrete calculus, and becomes a sum. We discuss details on a possible implementation of the indirectness measure we proposed in Eqn. 1 and corresponding experimental results in Sec. 6.3.

4.2 Composite Mappings

While our discussion has so far only focused on simple setups which only feature a single type of mapping, most real-world systems incorporate multiple distinct mappings. For instance, a typical multi-view visualization application as in Fig. 2 will feature several windows displaying the same data using multiple visualization techniques and a set of corresponding interaction mechanisms. A volume dataset, for example, may be presented in a 3D direct volume rendering view, a 2D slice view, as well as a histogram view. A visualization system for abstract data also typically contains multiple linked views such as scatterplots, parallel coordinates, bar- or line graphs, map views, etc. Each of these views may exhibit a different degree of directness. Similarly, often multiple manipulation and/or output mappings are combined, e. g., keyboard and mouse with head-tracking and a wand device.

The goal of employing multiple distinct visual mappings is often to represent different aspects of the data (e. g., spatial distribution of data values in a 3D rendering vs. frequency information in a histogram) and specific interaction tasks will




Fig. 2. Standard interactive volume rendering can be captured by the proposed model. It contains two types of visualization spaces \mathcal{V} , which are represented by the 3D rendering and the histogram display shown in the transfer function editor.

be localized to the individual views. As the principle mappings for each view are distinct, a multi-view system will exhibit a set of distinct degrees of directness for vi and, potentially, om . Composite mappings can thus facilitate increased degrees of directness for particular tasks. For instance, precise localization tasks for 3D data are often performed in slice views where a high level of vi directness is achievable. On the other hand, it has been shown that the required visual attention switching may have detrimental effects on performance [67] and this trade-off needs to be taken into account in the overall design of a visualization system. Within the framework of our model, we can regard multiple mappings from/to the same data as being associated with different importance functions (see Eqn. 1) related to the respective task.

5 MODEL INSTANCE DERIVATION

To ease application of the model proposed in the previous section to existing interactive visualization scenarios, we describe the evolution of particular instances of the model by going back to the initial examples of this paper from Sec. 2. With these derivations we also demonstrate the descriptive usage of our model, i. e., its application to existing visualization systems.

Spatial data scenario. In this example a user interactively explored a volumetric data set using two major visual representations (Fig. 2). First, and most obviously, the system provides the 3D volume-rendered representation as seen on the left. Second, the transfer function editor shown on the right also incorporates a representation of the data by means of the histogram displayed in blue. In the context of these two visual representations, the user can employ several spatial interaction techniques. Volume rendering offers camera navigation by means of a virtual trackball, while the transfer function widget supports navigation of the transfer function space and enables selections/highlighting in the value domain. While the volume rendering can be considered the main visualization, the transfer function editor with the histogram rather serves as an interaction widget. Therefore, we focus our model instance derivation on the volume rendering view and only point out that a model instance can be derived similarly for the histogram view.

We start our derivation by taking into account the visualization mapping v . Direct volume rendering applies a comparably simple global transformation to the data, since no aggregation or clustering as well as deformation of the Cartesian grid is applied, and thus the neighborhood relation of the actual data points is conserved. Hence, v is mainly parameterized by the camera parameters, which are defined by the model, view, and projection matrices as well as the used compositing function. Accordingly, we can express v

as an affine transformation which can be represented as a 4×4 matrix together with the subsequent blending of semi-transparent structures. While often only a single depth value is preserved, we acknowledge the research that has been conducted regarding volumetric picking [49], [93] and thus can consider i to be a *quasi-inverse* of v , as already discussed in Sec. 4. This consideration takes into account that a single 3D position that matches the content depicted in the image can often be approximated. Therefore, a relatively high degree of vi directness is achieved in this example.

The subsequent output mapping o is given by the presentation of the image on a desktop monitor. As o leads to a 2D image which represents a projection of 3D space, the user can examine it to gain insights about the 3D data, i. e., applying u_p , the mapping from the output space \mathcal{O} to the user space \mathcal{U} . The manipulation mapping m , however, is defined by a standard mouse as the primary input device, which is operated in a different physical space than the presented image. Accordingly, as discussed in Sec. 4, this desktop-based interaction setup renders the manipulation mapping m invertible, but not direct. Based on the cursor movements on the screen, the mouse movement can be derived, but it is transformed w. r. t. the mouse and screen coordinate systems. The fact that manipulation occurs in a different physical space than the presentation of visual results means that this example exhibits a lower degree of om directness. The user's action u_a has to occur in a different physical space than the one where its perceived effects will be seen.

An interesting aspect of this particular setup is the degree of mental integration between these two distinct spaces that occurs as part of the u_p and u_a mappings. It can often be seen that trained users can perform remarkably well at this, while novice users frequently show initial problems. This implies that there may be different classes of o and m pairs for which the degree of perceived directness may increase with sufficient training and experience. We leave further exploration of this aspect as an interesting direction for future work and point to hints for this discussion in the work by Hutchins et al. [38].

Abstract data scenario. The second example described in Sec. 2 employed the combination of parallel coordinates with a 2D scatterplot, to explore multi-variate data initially stored in a table.

Again we start our analysis by considering the visualization mapping v . For both cases, we can observe a transformation of the data—from the table to the visual representation—which can be described with a rather simple mapping. In the parallel coordinates case, we can directly interpret the columns of the table as parallel coordinate axes. In the 2D scatterplot a similar mapping can be observed, but one of the two selected axis needs to be transformed, such that they are orthogonal to each other. When analyzing the possibility to select data points in the two visual representations, it can be observed that in the parallel coordinates plot no unambiguous data selection may be possible, depending on the clutter of the shown data entries. However, in the past techniques have been developed to reduce the clutter in parallel coordinate plots (e. g., [58]). While stereoscopic techniques might help to overcome these selection ambiguities, involving stereoscopy would result in another scenario. Hence, in a monoscopic scenario, the unambiguous selection is a problem, which cannot be overcome. In principle an unambiguous selection is also impossible for 2D scatterplots; i. e., when a high amount of overdraw occurs. Nevertheless, here the overdraw of two points would mean that they have the same values—or that the screen resolution could not resolve their differences—and thus could not be distinguished by a spatial selection anyway. For the scatterplot representation we




Fig. 3. Classification of methods according to the invertibility of the mappings using the coloring of Fig. 1 shown on the left (see also Sec. 6.1). Top to bottom: v and i are non-inverse, quasi-inverse or inverse. Left to right: \mathcal{O} and \mathcal{M} are different or identical.

thus consider i to be an inverse of v , and achieve a high degree of vi directness in this example, while we have a less degree of vi directness when it comes to parallel coordinates. Depending on the output device, \mathcal{O} and \mathcal{M} are identical—for the stereoscopic case—or not. The same observation thus holds for om .

6 APPLICATION EXAMPLES AND CLASSIFICATION

Based on the practical derivation of a specific single instance of our model in the previous section, we can now further benefit from the descriptive power of the model and apply it to several examples of interactive visualization scenarios found in literature. In this section we present how these scenarios can be expressed through our model, mainly to demonstrate two qualities of the proposed model. First, we demonstrate that the proposed model is general enough to capture a wide variety of scenarios. Second, we show that, despite its generality, the model can still be applied within a practical context. To reach both goals, in the first subsection we propose a way to group application examples based on our model. While this is done based on a specific example, i. e., data selection, we discuss further real-world scenarios with respect to this grouping in the second subsection. The third subsection is dedicated to the description of an exemplary implementation of the quantitative indirectness measure $\Delta_{v,i}$ and a discussion of experiments with this implementation.

6.1 Scenario Classification

To classify existing scenarios by means of our proposed model we take into account the visualization-interaction mapping vi and the output-manipulation mapping om . As stated in Sec. 4, if the interaction mapping i is the inverse of the visualization mapping v , the two spaces \mathcal{V} and \mathcal{I} collapse and we have the most direct form of vi which is the identity mapping. In many cases, however, the mappings v and i are not inverse and the mapping vi is thus less direct. As also discussed above, we consider quasi-inverse mappings which facilitate an inversion by taking into account contextual information. We thus obtain three representative cases for vi , while—with respect to om —we only distinguish two cases: the spaces \mathcal{O} and \mathcal{M} either collapse (i. e., the om mapping is the identity) or those spaces do not collapse. We have not taken quasi-invertible mappings into account for om because this part of our model directly relates to the user's manipulation movement: it can either coincide with the output space, or it does not. In total we thus obtain $2 \times 3 = 6$ classes based on the directness of vi and om . The model graph from Fig. 1 lends itself to visually represent the different classes. In Fig. 3 we provide such a representation where shorter distances correspond to higher directness of mappings. To further illustrate the classes, in the following discussion we provide example scenarios for each class by considering data selection as a commonly performed interaction in visualization.

Class 1— v, i : inverse, om : identity Both v and om are the identity if the spaces \mathcal{V} and \mathcal{S} as well as \mathcal{O} and \mathcal{M} collapse, respectively. This is the case, for example, for abstract data (e. g., data values in tables or relationship information) which is encoded as a 2D representation, shown on a 2D screen, on which the input is physically provided, which again is interpreted as manipulations on the 2D representation. For instance, interactions such as selection within a 2D scatterplot visualization on a 2D touchscreen⁵ fall into this category (e. g., [71]). A related case for 3D data representations exists as well, when all spaces are truly three-dimensional, for instance when making a selection in a 3D dataset that is displayed in a stereoscopic CAVE⁶ while using a 3D-tracked input device (e. g., [8], [51], [60]). In both examples not only do \mathcal{V} and \mathcal{S} as well as \mathcal{O} and \mathcal{M} collapse, respectively, but all four spaces are virtually identical. Interactive visualization scenarios that are similar to these examples thus lead to users perceiving a high degree of interaction directness.

Class 2— v, i : quasi-inverse, om : identity When the v and i mappings are not inverse such as for the case of displaying a 2D projection of a 3D visual representation, contextual information can be used to relate interaction input from the 2D input space to the 3D space of the visual representation; i thus becomes quasi-inverse of v . An example is a context-sensitive selection within a 3D point cloud dataset displayed on a 2D touchscreen (e. g., [98], [99])—contextual information is used to derive a 3D selection volume from the 2D input. Another example is the selection within a semi-transparent 3D dataset displayed on a powerwall using a tracked input device (e. g., [51]).

Class 3— v, i : non-inverse, om : identity In particular when the data to be visualized has no particular spatial component, the mapping from data space to visualization space often offers a higher degree of flexibility, as no predetermined spatial arrangement exists for visualization and interaction. It must be established by the designer. When combining such visualizations with more direct display and manipulation hardware, we obtain a specific group, where v and i are non-inverse, and om is the identity. Common example applications would, for instance, be the selection of data points through a histogram on a 2D touchscreen, or the manipulation of data points through parallel coordinates in a stereoscopic CAVE with a 3D-tracked input device.

Class 4— v, i : inverse, om : not identity Assuming v and i as inverse and om as not identity, we can naturally group all typical desktop-based scientific visualization applications which do not display the data semi-transparently. Accordingly, representative examples for this group would be the selection within a 2D dataset in a standard desktop environment, or the selection within a projection of an opaque 3D surface in a standard desktop environment—in both cases using mouse, remote pointer, or certain forms of tangible or even gaze-based input.

Class 5— v, i : quasi-inverse, om : not identity One example of this class is picking within a semi-transparent 3D dataset displayed on a standard PC screen with the input provided by the mouse and where contextual information is used to derive the perceptually most likely picking depth (e. g., [93]). The difference to the second class above with quasi-inverse v

and i is that, here, the spaces \mathcal{O} and \mathcal{M} are different: the projection of the visual data representation is on the screen (\mathcal{O}), while the input with the mouse takes place on the desktop (\mathcal{M}).

Class 6— v, i : non-inverse, om : not identity The last of the six groups is potentially the biggest, since it captures all combinations where no particular consideration was given to directness. We consider all scenarios where v and i are non-inverse and where om is not the identity. Due to the fact that the latter is true for all common desktop systems, this group spans a large number of scenarios. However, the constraint with respect to v and i is that they are non-inverse, so mappings having a lower degree of directness fall within this group. With respect to our data selection example, we consider the selection of data points through a histogram in a standard desktop environment by means of mouse input as a representative. For instance, systems that incorporate complex analysis functionality initiated by menus or toolbars would likely fall into this class.

The attentive reader has observed that the given example scenarios are roughly ordered with decreasing directness. Accordingly, more *true-to-the-data* representations together with appropriate displays appear at the beginning, while more abstract representations in standard desktop environments appear towards the end of the list. Furthermore, it can be observed that an invertible om is, in most cases, realized by means of virtual reality technology. Finally, it can be seen that the type of data also has an effect on the invertibility and thus the overall directness. When, for instance, one is performing a selection within a projection of an opaque data set, a 2D screen is sufficient for an invertible selection interaction, while this is not the case for semi-transparent representations where multiple data points coincide in screen space.

6.2 Example Visualization Scenario Discussion

To better illustrate the implications of our model for practical interactive data visualization scenarios, we now discuss several common real-world environments with respect to the proposed model. In particular, we discuss VR-based environments for the exploration of 3D data, PC environments that rely on 2D projections of 3D visualizations, tactile interaction with the same data representations, and interaction with 2D visualizations of 2D spatial or abstract data.

6.2.1 3D-Tracked Interaction in VR with Volumetric Data

While we have derived a model instance for desktop-based volume interaction already in Sec. 5, we now focus on VR-based interaction. Within immersive environments, interactions such as specifying 3D locations, selecting and manipulating 3D objects, and moving within the 3D space are essential. Many researchers and practitioners intuitively perceive a 3D-tracked interaction with a 3D representation of volumetric data in a virtual environment (e. g., [51] or Fig. 4) to be direct and, in particular, more direct than the interaction with the same representation on a desktop screen. Applying our model supports this intuition. In a VR environment, the presentation of 3D data and the tracked interaction specifying a certain position in space by simply moving the tracked object to the desired position happen in the same space. In terms of our model this situation means that om is the identity mapping, and essentially that all the spaces \mathcal{D} , \mathcal{V} , \mathcal{O} , \mathcal{M} , and \mathcal{S} collapse, at least from a conceptual perspective that abstracts from a more technical point of view (Class 1). This allows users to directly relate physical input, the virtual effects of this input, the changes to the data (space), and

5. In this specific example a small vertical offset exists between the displayed 2D image (\mathcal{O}) and the input sensing (\mathcal{M}) but this offset is typically ignored by the user—it is smaller than the error introduced by the “fat fingers” and the perceived horizontal offset in input sensing [37].

6. Here, \mathcal{O} technically comprises two 2D images, but these images correctly displayed together form a 3D view from the perspective of the observer.




Fig. 4. Visualization of volumetric data in CAVE environment controlled using head tracking and a tracked wand device [33]. Image courtesy of and © Claudia Hänel and Bernd Hentschel, used with permission.

the resulting changes to the visual representation with each other, with the space \mathcal{U} also being closely related to \mathcal{D} .

Different metaphors for selecting positions or objects in virtual environments exist (e. g., [2], [9]) and not all of them exhibit the same degree of directness. A common example is pointing with a wand-like device towards an object and determining the selected object by ray intersection. Due to the distance between the location of the input device in \mathcal{M} to its effect in \mathcal{I} (Class 4) it is less direct than the interaction described before because now vi is not necessarily the identity mapping.

6.2.2 Exploration of Volumetric Data on a 2D Touch Screen

The previous scenario is a common environment for the exploration of 3D data. With the advance of new input sensing technology over the past decade or two, however, additional spatial exploration settings have become available such as touch-sensitive displays, and several approaches for 3D exploration have been developed and investigated [40]. In comparison to, in particular, desktop/workstation-based settings and environments, this tactile form of spatial interaction is often cited to be more direct (e. g., [48])—not only for the interaction with 3D data (e. g., [46], [48]).

The left part of our model's instance for this case resembles that of desktop-based interaction discussed in Sec. 5. The mappings v and i are inverse, quasi-inverse, or non-inverse depending on the specific interaction technique. For instance, a translation parallel to the view plane for an orthographically projected data representation would be invertible (Class 1; e. g., [19], [20]). Most interactions, however, require at least some contextual information to make the mappings quasi-inverse (Class 2). Translations parallel to the view plane in perspectively projected representations, for instance, need to infer a specific plane in 3D space to which the 2D translation input is being mapped (e. g., [21], [80], [83], [100]), while ARCBALL-like rotations [76] need to infer a center of rotation (e. g., [21], [27], [80], [97], [100]; see Fig. 5(a)). Similarly, a spatial selection based on a drawn 2D lasso can use contextual information to infer the specific 3D selection volume likely meant by the user (e. g., [98], [99]; see Fig. 5(b)).

The right part of our model, however, differs for touch-based input from the desktop scenarios. In contrast to the previous case of using input devices that operate in a different space from that of the displayed image (Classes 4–6), touch-based input is captured in the same space where the 2D image is displayed: \mathcal{O} and \mathcal{M} are identical/collapse (Classes 1–3). This means that the om mapping




Fig. 5. Examples of direct-touch interaction with projected 3D data: (a) navigation [100] and (b) spatial selection [98]. Both images © IEEE (2010 and 2012, respectively), used with permission.




Fig. 6. Illustration of the direct-touch interaction with 3D data showing the mental mapping necessary for the mappings v and i (and, respectively, u_p and u_a): The 2D display space in the illustration corresponds to the spaces \mathcal{V} , \mathcal{O} , \mathcal{M} , and \mathcal{I} simultaneously, while the 3D data shown behind the screen corresponds to the spaces \mathcal{D} and \mathcal{U} . Image © 2013 Tobias Isenberg, used with permission.

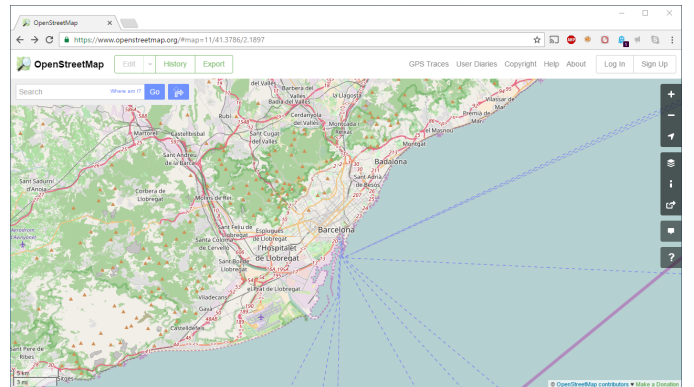


Fig. 7. Interface to explore geospatial 2D visualizations—typically displayed on a 2D screen and controlled with the mouse.

is more direct than that for desktop environments, requiring less mental processing for the person working with the interface.

Nonetheless, both are 2D spaces so a user still has to mentally integrate the dimensional distance between these 2D spaces and the 3D data space (Fig. 6). This mental integration involves a mapping of the input from \mathcal{I} to \mathcal{D} , which is done algorithmically but really happens in the u_p mapping from \mathcal{O} to \mathcal{U} . It is aided by the fact that \mathcal{V} , \mathcal{O} , \mathcal{M} , and \mathcal{I} are all 2D spaces which are collapsed,⁷ so the input in \mathcal{M} intuitively relates to a position in \mathcal{O} and its effect is algorithmically mapped to \mathcal{D} and, simultaneously, mentally to \mathcal{U} .

The discussion so far relates to stationary display/input device combinations on which spatial input is being provided. Touch-sensitive displays, however, can also be used in the form of mobile,

7. We (here) treat \mathcal{V} and \mathcal{I} as being collapsed for all interaction designs that lead to inverse or quasi-inverse mappings v and i .

tangible displays where the location and orientation of the display itself can be a form of input, while the display shows the projected 3D view depending on this input. For example, the tangible display can show the view of a virtual camera that is attached to the tangible device (e. g., [13], [57]) or the viewer (e. g., [7]), or it can function as a slicing plane through the data (e. g., [81]), in both cases with respect to a (typically) invisible 3D dataset that exists in the same physical space as the tangible device. Yet, its existence is only established in the u_p mapping from \mathcal{O} to \mathcal{U} . If the tangible device location and orientation, for instance, is used as a handle to manipulate the imagined data location and orientation (e. g., [7]) then the manipulations of the device in a three-dimensional \mathcal{M} affect both the mapping v (projected view on the tangible device) as well as o (the physical location of the display with respect to \mathcal{U} and \mathcal{M}). The tangible device then becomes a direct or indirect handle for manipulations of the (projected) view on the data: *Direct* for the virtual camera [13] because the location/orientation of the tangible device in \mathcal{M} matches that of the virtual camera in \mathcal{I} , \mathcal{D} , and \mathcal{U} ; *indirect* for the tangible data manipulations [7] because the manipulations of the device in \mathcal{M} are only used as offsets to manipulations of the location of the data with respect to \mathcal{M} or \mathcal{U} ; i.e. the location of data remains constant w.r.t. the display of the tangible device. Any additional (tactile) input on the tangible, however, is still governed by the same 3D-2D mapping considerations as described before.

6.2.3 Interaction with 2D Representations of 2D Spatial Data

The issue with the mental mapping from 3D space to a 2D data representation does not present itself in the case when data is explored that is inherently two-dimensional. For example, geospatial visualizations (e. g., OpenStreetMap or similar; Fig. 7) are often thought of as 2D data, so the spaces \mathcal{D} , \mathcal{V} , and \mathcal{I} collapse, resulting in the spatial interaction with such data (e. g., translations, rotations, scaling) being perceived to be quite direct. Depending on the specific input device that is being used, \mathcal{O} and \mathcal{M} may also collapse into the same space (for touch-based or pen input, Class 1) or remain separated (e. g., mouse-based or gaze-based input, Class 4), with the implications for the perception of directness for the interacting person already described in Sec. 6.2.2.

6.2.4 Interaction with 2D Representations of Abstract Data

The exploration of multi-dimensional abstract data may initially seem similar to 3D data as discussed in Sections 5 and 6.2.2: The mapping from \mathcal{D} to \mathcal{V} typically involves a projection to a lower-dimensional space, and the mapping from \mathcal{I} to \mathcal{D} , consequently, requires at least some contextual information to make it quasi-inverse (Class 2 or Class 5), otherwise it remains non-inverse (Class 3 or Class 6). While this reasoning is valid from a technical perspective, the case of 2D representations of abstract data is different from that of 3D spatial data because, for abstract data, no inherent mapping and thus no inherent interpretation of the resulting representation exists: The arrangement of visual elements to represent data in 2D space affects the mental model that viewers construct of the representation (e. g., [61, Ch. 7])—people use such representations to form an “external memory” [61, Sec. 1.4] to “amplify cognition” [15]. Some personality traits may even have an influence on a person’s mental model as shown, e. g., by Ziemkiewicz et al. [101].

In these considerations the user space \mathcal{U} is quite different from the data space \mathcal{D} . Consequently, manipulations in \mathcal{I} are no longer necessarily interpreted as manipulations of \mathcal{D} (also since a mental




Fig. 8. Interface to explore abstract data representations (e. g., scatterplots) using a touch display [24]. Image courtesy of and © 2015 Visualization Lab @ Georgia Tech, used with permission.

model of such a high-dimensional space is difficult to impossible to construct) but as manipulations of elements of the mental model, with the latter being heavily influenced by the mapping that leads to \mathcal{V} . This situation results in a quite direct mapping if elements in \mathcal{V} are directly affected due to the manipulations in \mathcal{I} . Interestingly, such manipulations can also occasionally be interpreted as manipulations of a temporally existing 3D data (and thus user) space (e. g., [26]), then requiring the same considerations for directness as we discussed for the exploration of 3D data on 2D surfaces.⁸ Similarly, the use of a specific input paradigm also has an effect on the perception of directness such as the difference between tactile (e. g., Fig. 8) and mouse-based control, with the same considerations as discussed earlier in the section.

Certain forms of tangible interaction are also an interesting case in this context because they often carry the data view on the manipulated device. Spindler et al.’s [82] Tangible Views, for example, can be used to select subsets of abstract data visualizations through translations of a tangible device, while the device itself carries the visualization as it changes during interaction. Depending on how the interaction with the Tangible View is interpreted, however, this interaction can be classified as Class 1—if the device is seen as a window that extends one’s hand and which one moves in the same spatial context as the visualization—or as Class 4—if the motion of the physical representation of the view is interpreted as largely being distant from the finger/hand that is manipulating it.

6.3 A Practical Application of the Quantitative Indirectness Measure

While we have already discussed how, in principle, our indirectness measure $\Delta_{v,i}$ (see Eqn. 1) can be applied to quantify the degree of directness for a particular visualization setup, we now demonstrate the feasibility of this process using a concrete scenario. Our goal here is not to derive new insights; on the contrary, we want to demonstrate that our measure is capable of reflecting intuitive design choices that have been documented in the literature. For this purpose, we choose the comparatively simple and easy to understand scenario of DVR picking and outline two examples below.

Our implementation⁹ evaluates the integral in Eqn. 1 by summing up (and normalizing) the argument values for a large

8. Of course, there is also work on mapping abstract data to a 3D representation [10]; for such visualizations the considerations in Sec. 6.2.2 apply as well.

9. Available in OpenWalnut at <http://source.openwalnut.org/> as of April 2017.




Fig. 9. DVR with transfer functions producing (a) an isosurface and (b) a foggy appearance by low but non-zero opacity values for most data values and high values for the rest.

number (8,000 in our experiments) of samples uniformly distributed within a dataset. We obtain the value for a sample x by first using the DVR representation, i. e., projecting the sample position onto the bounding box of the dataset (*visualization mapping* v), and then applying a picking technique (*interaction mapping* i) which selects a 3D position along the viewing direction according to a clicked position in screen space and a specific criterion. Here vi bridges the gap between the position on the bounding box and the one in screen space. We then weight the distance between the resulting position $i(vi(v(x)))$ and x using a pre-defined *importance function* (details in the examples below) evaluated at x . To obtain view-independent values, we compute projections along different viewing directions (26 in our experiments). Furthermore, to make $\Delta_{v,i}$ comparable across experiments, we scaled the datasets to fit into the unit cube $[0, 1] \times [0, 1] \times [0, 1]$.

Example 1. For this example we computed $\Delta_{v,i}$ for a DVR with an isosurface transfer function for a low isovalue such as shown in Fig. 2 or Fig. 9(a). As importance function $I(x)$ we chose the opacity at the original sample position x because important features are commonly assigned higher opacities to make them visually salient. With the described configuration we compare the results for three different picking criteria [92], namely maximum intensity, maximum opacity, and first-hit (first encountered non-zero opacity).

As expected, the first-hit ($\Delta_{v,i} = 0.000823$) and the maximum opacity ($\Delta_{v,i} = 0.0006895$) criteria result in low indirectness. Locations with high opacity are correctly selected by these criteria, thus contributing only small distances to $\Delta_{v,i}$. Locations with zero opacity result in large distances because they are not selected by the picking criteria (locations with high opacities are selected instead). However, these large distances are weighted by the importance function which is zero. These distances thus do not increase $\Delta_{v,i}$.

The maximum intensity criterion ($\Delta_{v,i} = 0.002407$) yields much larger indirectness because it may produce large distances for locations considered important (large importance function values). As the considered transfer function maps low intensities to high opacities, it assigns low intensity locations a high importance. The picking criterion, however, selects a location with high intensity which may have a large distance to the initial location x .

Example 2. Our second example uses a transfer function producing foggy DVR images (Fig. 9(b)). This transfer function maps most values to small but non-zero opacity. In addition, it maps few values (close to the isovalue) to high opacities.

For the same picking criteria and importance function as above, we expect the first-hit criterion to result in the highest indirectness. Due to the overall non-zero opacity, it results in positions on the dataset boundary for all samples x and thus in large distances to most of these samples. Our experiments confirm this: for maximum

intensity we get $\Delta_{v,i} = 0.001707$, for maximum opacity we get $\Delta_{v,i} = 0.0007630$, and for first-hit we get $\Delta_{v,i} = 0.002079$. The maximum intensity criterion also appears to be relatively indirect again because high intensities have been mapped to low opacities by the transfer function. Although there are no locations with zero opacity and thus zero importance in this example, the maximum opacity criterion benefits from the very low opacity and importance around the high opacity regions which eliminate the influence of the regions where the criterion selects the “wrong” locations.

Other examples. Other importance functions might also be of interest for DVR such as those that use original data values, i. e., intensity, or simply a uniform importance of 1. We also tested such examples but found these results less instructive and thus do not discuss them here. Such functions result in considerably different quantitative results, which shows how strongly $\Delta_{v,i}$ depends on the chosen function that models what is important to users. Thus $\Delta_{v,i}$ has its limitations in cases in which the importance of the samples to the user is difficult to quantify. Nevertheless, the two examples discussed above show that the indirectness measure $\Delta_{v,i}$ can be used to compare different interaction techniques in a setting where an importance function can be given and evaluated.

7 DISCUSSION

In the past years, we have seen an increased emphasis on theoretical contributions in visualization, and several important advances have been made in establishing a sound scientific foundation for the field. The work we presented in this paper was motivated by a gap in the literature regarding the notion of spatial interaction directness in visualization and our aim was to contribute to an improved understanding of its nature. The proposed model can be used in three principal ways: (1) *descriptive*—for describing existing interactive visualization solutions, (2) *evaluative*—to help assess design alternatives, and (3) *generative*—to support the creation of novel ideas and solutions.

Descriptive usage. As demonstrated in Sec. 5, Sec. 6.1 and Sec. 6.2, the proposed model, its structural abstraction, and its terminology can be employed to abstractly describe interactive visualization setups and systems with respect to their interaction capabilities. By clarifying the notion of directness and its dimensions, the model will facilitate improved communication with respect to this previously only vaguely defined term. In particular, the distinction between vi and om directness can help researchers to reduce contradictory descriptions and can assist them in highlighting meaningful differences between individual system designs. Indeed, in our analysis of the literature, we came across several examples where vi and om directness were confounded, which can obscure subtle but important differences and thus has a detrimental effect when comparing different setups. We described several visualization scenarios in terms of our model and showed how they correspond to examples in the literature. We are therefore confident that the introduced framework is descriptive for a wide range of instances and that it can contribute to establishing a more precise terminology for the important aspect of interaction in visualization.

Evaluative usage. While the first parts of Sec. 6 focused on the descriptive use of our model, we demonstrated the evaluative usage in Sec. 6.3. By introducing a measure for the degree of indirectness exhibited by a pair of v, i or o, m mappings (Eqn. 1), our framework facilitates the evaluation of different design alternatives. In particular, we believe that the presented model can form the basis

for establishing a quantitative link between mapping indirectness and various measures of performance. Existing studies have already explored the influence of directness on task performance [69] but only made a binary distinction between *direct* and *indirect*, which makes it difficult to explore this relationship in a more detailed manner. We also want to emphasize that *we do not argue that an increased level of directness necessarily leads to better performance in general*. On the contrary, it is our hypothesis that the correspondence between directness and different performance measures is highly task-dependent. We envision future studies which will relate performance and *vi* as well as *om* directness for existing task taxonomies. Given such data, it will be highly interesting to see the predictive power of our measure.

Generative usage. Given the fact that at least our *vi* indirectness measure can be numerically approximated (for *om* directness, this may be more challenging), we foresee the possibility of directly computing estimates for an existing or proposed interaction technique based on a set of given importance functions which model user tasks. A modular visualization framework such as MeVisLab¹⁰ or Inviwo [84], for instance, could automatically suggest interaction schemes for a chosen visualization setup based on their directness. While the realization of these goals certainly requires substantial further work, in this paper we aim to lay the groundwork for such developments. Further exploration of the quantification of interaction directness may also provide us with additional means for steering and/or optimizing interaction methods in visualization. Feature detection methods, for instance, may be used to establish an importance function as used in Eqn. 1 to automatically adapt the directness of an interaction technique with respect to these features. Lindow et al. [55] already presented work along a similar line of thought—they linearize the effect of parameter changes with respect to the visual perception of image changes. It is further interesting to note that, at least to a certain degree, previous work on visual encoding has also already incorporated the fact that particular types of visual mappings can be more easily inverted than others. A part of the appeal of unfolded representations of 3D structures [65, Ch. 14], for instance, is that they facilitate higher *vi* directness in a 2D interaction space.

An aspect not specifically covered by our model that can increase the *perceived* directness for interaction in virtual environments is the sense of touch, i. e., (force) feedback when reaching out for a virtual object. While such aspects are, in principle, captured by u_p and u_a , and can be considered to be part of the *i* and *m* mappings, a more explicit representation of the sense of touch or even other senses may be a viable extension. Touch-based interfaces, for example, benefit from their inherent support of somesthetic feedback [68] while the interaction is taking place, leading to a higher degree of similarity of the interaction to the manipulation of a tangible object. Consequently, the manipulation may feel more direct to some users than when using other forms of input for data exploration. In this sense, a “holodeck”-like virtual environment which completely and perfectly immerses the user in a physical representation of the data could be seen as an example for maximal directness.

8 CONCLUSION

We presented a model of directness in spatial interaction for visualization. Where previous discussions of such interaction

scenarios have mostly used the term *direct* by intuition, our proposed model captures different aspects and degrees of the directness of interaction in a more formal way. Based on a number of mappings connecting different intermediate spaces between data and users, the model enables the analysis of the relations between different steps of the visualization technique and their corresponding parts of the interaction technique in a comprehensive and more detailed way than it has been possible before. In our model, the directness of certain steps of the examined techniques is captured in the invertibility of the mappings and the identity of spaces. We have shown that our proposed model

- is general enough to cover a wide variety of scenarios,
- can be used to describe, discuss, and evaluate interactive visualization solutions and designs, and
- captures characteristics of techniques that allow us more detailed judgments about the directness of spatial interactions with visualizations than have been possible before.

We also envision a generative usage of the model that would allow systems to automatically suggest appropriate interaction techniques for given visualizations or vice versa. Finally, we hope that with this paper we can start a discussion on the concept of interaction directness in visualization and beyond, to ultimately arrive at an understanding of how to best design interaction techniques depending on the type of data, visualization approaches, and interactive environments.

We would like to emphasize again that our model thus complements previous discussions of interaction directness (e. g., [34], [38], [52], [62]) and direct manipulation (e. g., [45], [74], [75]). Our model refers to *explicit spatial input* and does not cover potentially spatial interaction paradigms such as proxemic interaction [3]. Here, the spatial input does not necessarily refer to data elements but, instead, leads to implicit system reactions or only “afford[s] opportunities for explicit interactions” [3].

ACKNOWLEDGMENTS

The authors would like to thank Filip Sadlo for an insightful discussion at EuroVis 2016. This work was partially funded by the Deutsche Forschungsgemeinschaft (DFG) under grant RO 3408/3-1 (Inviwo) as well as the Federal Minister for Economic Affairs and Energy (BMWi) under grant ZF4483101ED7 (VRReconstruct). We would like to acknowledge the financial support of the Carl Zeiss foundation for the stereoscopic multi-touch powerwall installed in Ulm. The research presented in this paper was further supported by the MetaVis project (#250133) funded by the Research Council of Norway.

REFERENCES

- [1] B. Alper, T. Höllerer, J. Kuchera-Morin, and A. Forbes. Stereoscopic highlighting: 2D graph visualization on stereo displays. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2325–2333, Dec. 2011. doi: 10.1109/TVCG.2011.234
- [2] F. Argelaguet and C. Andujar. A survey of 3D object selection techniques for virtual environments. *Computers & Graphics*, 37(3):121–136, May 2013. doi: 10.1016/j.cag.2012.12.003
- [3] T. Ballendat, N. Marquardt, and S. Greenberg. Proxemic interaction: Designing for a proximity and orientation-aware environment. In *Proc. ITS*, pp. 121–130. ACM, New York, 2010. doi: 10.1145/1936652.1936676
- [4] M. Beaudouin-Lafon. Instrumental interaction: An interaction model for designing post-WIMP user interfaces. In *Proc. CHI*, pp. 446–453. ACM, New York, 2000. doi: 10.1145/332040.332473
- [5] R. A. Becker and W. S. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, May 1987. doi: 10.1080/00401706.1987.10488204

10. <http://www.mevisslab.de/>

- [6] J. Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983.
- [7] L. Besançon, P. Issartel, M. Ammi, and T. Isenberg. Hybrid tactile/tangible interaction for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):881–890, Jan. 2017. doi: 10.1109/TVCG.2016.2599217
- [8] D. A. Bowman and L. F. Hodges. Formalizing the design, evaluation, and application of interaction techniques for immersive virtual environments. *Journal of Visual Languages & Computing*, 10(1):37–53, Feb. 1999. doi: 10.1006/jvlc.1998.0111
- [9] D. A. Bowman, E. Kruijff, J. J. LaViola, Jr., and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley, Boston, 2005.
- [10] R. Brath. 3D InfoVis is here to stay: Deal with it. In *Proc. 3DVis*, pp. 25–31. IEEE Computer Society, Los Alamitos, 2014. doi: 10.1109/3DVis.2014.7160096
- [11] A. Buja, J. A. McDonald, J. Michalak, and W. Stuetzle. Interactive data visualization using focusing and linking. In *Proc. Visualization*, pp. 156–163. IEEE Computer Society, Los Alamitos, 1991. doi: 10.1109/VISUAL.1991.175794
- [12] K. Bürger, J. Krüger, and R. Westermann. Direct volume editing. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1388–1395, Nov./Dec. 2008. doi: 10.1109/TVCG.2008.120
- [13] W. Büschel, P. Reipschläger, R. Langner, and R. Dachsel. Investigating the use of spatial interaction for 3D data visualization on mobile devices. In *Proc. ISS*, pp. 62–71. ACM, New York, 2017. doi: 10.1145/3132272.3134125
- [14] O. Cakmakci and J. Rolland. Head-worn displays: A review. *Journal of Display Technology*, 2(3):199–216, Sept. 2006. doi: 10.1109/JDT.2006.879846
- [15] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann, San Francisco, 1999.
- [16] M. Chen and A. Golan. What may visualization processes optimize? *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2619–2632, Dec. 2016. doi: 10.1109/TVCG.2015.2513410
- [17] M. Chen and H. Jänicke. An information-theoretic framework for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1206–1215, Nov. 2010. doi: 10.1109/TVCG.2010.132
- [18] M. Chen, S. Walton, K. Berger, J. Thiyagalingam, B. Duffy, H. Fang, C. Holloway, and A. E. Trefethen. Visual multiplexing. *Computer Graphics Forum*, 33(3):241–250, June 2014. doi: 10.1111/cgf.12380
- [19] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, A. G. Erdman, and D. F. Keefe. Interactive Slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626, Oct. 2012. doi: 10.1109/TVCG.2011.283
- [20] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, and D. F. Keefe. Slice WIM: A multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality. In *Proc. 3D*, pp. 191–198. ACM, New York, 2011. doi: 10.1145/1944745.1944777
- [21] A. Cohé, F. Dècle, and M. Hachet. tBox: A 3D transformation widget designed for touch-screens. In *Proc. CHI*, pp. 3005–3008. ACM, New York, 2011. doi: 10.1145/1978942.1979387
- [22] M. Cordeil, B. Bach, Y. Li, E. Wilson, and T. Dwyer. A design space for spatio-data coordination: Tangible interaction devices for immersive information visualisation. In *Proc. PacificVis*, pp. 46–50. IEEE Computer Society, Los Alamitos, 2017. doi: 10.1109/PACIFICVIS.2017.8031578
- [23] C. Cruz-Neira, D. J. Sandin, T. A. DeFanti, R. V. Kenyon, and J. C. Hart. The CAVE: Audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6):64–72, June 1992. doi: 10.1145/129888.129892
- [24] A. Dai, R. Sadana, C. D. Stolper, and J. Stasko. Hands-on, large display visual data exploration. In *Posters of VIS*, 2015.
- [25] Ç. Demiralp, C. E. Scheidegger, G. L. Kindlmann, D. H. Laidlaw, and J. Heer. Visual embedding: A model for visualization. *IEEE Computer Graphics and Applications*, 34(1):10–15, Jan. 2014. doi: 10.1109/MCG.2014.18
- [26] N. Elmqvist, P. Dragicevic, and J.-D. Fekete. Rolling the dice: Multidimensional visual exploration using scatterplot matrix navigation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1539–1148, Nov./Dec. 2008. doi: 10.1109/TVCG.2008.153
- [27] C.-W. Fu, W.-B. Goh, and J. A. Ng. Multi-touch techniques for exploring large-scale 3D astrophysical simulations. In *Proc. CHI*, pp. 2213–2222. ACM, New York, 2010. doi: 10.1145/1753326.1753661
- [28] I. Fujishiro, B.-Y. Chen, W. Chen, S.-H. Hong, T. Itoh, K. Koyamada, K. Ono, and J. Nonaka. Top computational visualization R&D problems 2015: Panel. In *Proc. SIGGRAPH Asia Symposium on Visualization in High Performance Computing*, pp. 20:1–20:13. ACM, New York, 2015. doi: 10.1145/2818517.2818545
- [29] C. A. Gough, R. Green, and M. Billinghurst. Better realising direct manipulation. In *Proc. New Zealand Image and Vision Computing Workshop*, pp. 455–460. Industrial Research Ltd., Christchurch, N.Z., 1995. doi: 10092/2539
- [30] C. A. D. Gough, R. Green, and M. Billinghurst. Accounting for user familiarity in user interfaces. In *Proc. CHINZ*, pp. 137–138. ACM, New York, 2006. doi: 10.1145/1152760.1152778
- [31] T. M. Green, W. Ribarsky, and B. Fisher. Visual analytics for complex concepts using a human cognition model. In *Proc. VAST*, pp. 91–98. IEEE Computer Society, Los Alamitos, 2008. doi: 10.1109/VAST.2008.4677361
- [32] T. Grossman, D. Wigdor, and R. Balakrishnan. Multi-finger gestural interaction with 3D volumetric displays. In *Proc. UIST*, pp. 61–70. ACM, New York, 2004. doi: 10.1145/1029632.1029644
- [33] C. Hänel, B. Weyers, B. Hentschel, and T. W. Kuhlen. Visual quality adjustment for volume rendering in a head-tracked virtual environment. *IEEE Transactions on Visualization and Computer Graphics*, 22(4):1472–1481, Apr. 2016. doi: 10.1109/TVCG.2016.2518338
- [34] R. Hartson. Cognitive, physical, sensory, and functional affordances in interaction design. *Behaviour & Information Technology*, 22(5):315–338, 2003. doi: 10.1080/01449290310001592587
- [35] H. Hauser and H. Schumann. Visualization pipeline. In L. Liu and M. T. Özsu, eds., *Encyclopedia of Database Systems*, pp. 3414–3416. Springer, Cham, Switzerland, 2009. doi: 10.1007/978-0-387-39940-9_1133
- [36] B. Hibbard. Top ten visualization problems. *ACM SIGGRAPH Computer Graphics*, 33(2):21–22, May 1999. doi: 10.1145/326460.326485
- [37] C. Holz and P. Baudisch. The generalized perceived input point model and how to double touch accuracy by extracting fingerprints. In *Proc. CHI*, pp. 581–590. ACM, New York, 2010. doi: 10.1145/1753326.1753413
- [38] E. L. Hutchins, J. D. Hollan, and D. A. Norman. Direct manipulation interfaces. *Human-Computer Interaction*, 1(4):311–338, 1985. doi: 10.1207/s15327051hci0104_2
- [39] A. Inselberg. The plane with parallel coordinates. *The Visual Computer*, 1(2):69–91, Aug. 1985. doi: 10.1007/BF01898350
- [40] T. Isenberg. Interactive exploration of three-dimensional scientific visualizations on large display surfaces. In C. Anslow, P. Campos, and J. Jorge, eds., *Collaboration Meets Interactive Spaces*, chap. 6, pp. 97–123. Springer, Berlin/Heidelberg, 2016. doi: 10.1007/978-3-319-45853-3_6
- [41] T. Jankun-Kelly, K.-L. Ma, and M. Gertz. A model and framework for visualization exploration. *IEEE Transactions on Visualization and Computer Graphics*, 13(2):357–369, Mar. 2007. doi: 10.1109/TVCG.2007.28
- [42] Y. Jansen and P. Dragicevic. An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405, Dec. 2013. doi: 10.1109/TVCG.2013.134
- [43] C. R. Johnson. Top scientific visualization research problems. *IEEE Computer Graphics and Applications*, 24(4):13–17, July/Aug. 2004. doi: 10.1109/MCG.2004.20
- [44] D. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. Visual analytics: Definition, process, and challenges. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, eds., *Information Visualization: Human-Centered Issues and Perspectives*, pp. 154–175. Springer, Berlin/Heidelberg, 2008. doi: 10.1007/978-3-540-70956-5_7
- [45] D. Kieras, D. Meyer, and J. Ballas. Towards demystification of direct manipulation: Cognitive modeling charts the gulf of execution. In *Proc. CHI*, pp. 128–135. ACM, New York, 2001. doi: 10.1145/365024.365069
- [46] K. Kin, M. Agrawala, and T. DeRose. Determining the benefits of direct-touch, bimanual, and multifinger input on a multitouch workstation. In *Proc. Graphics Interface*, pp. 119–124. CIPS, Toronto, 2009. doi: 10.20380/GI2009.16
- [47] G. Kindlmann and C. Scheidegger. An algebraic process for visualization design. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2181–2190, Dec. 2014. doi: 10.1109/TVCG.2014.2346325
- [48] S. Knoedel and M. Hachet. Multi-touch RST in 2D and 3D spaces: Studying the impact of directness on user performance. In *Proc. 3DUI*, pp. 75–78. IEEE Computer Society, Los Alamitos, 2011. doi: 10.1109/3DUI.2011.5759220
- [49] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Contextual picking of volumetric structures. In *Proc. PacificVis*, pp. 185–192. IEEE Computer Society, Los Alamitos, 2009. doi: 10.1109/PACIFICVIS.2009.4906855
- [50] R. Kosara. Indirect multi-touch interaction for brushing in parallel coordinates. In *Proc. SPIE*, vol. 7868, pp. 786809:1–786809:7. SPIE, Bellingham, WA, USA, 2011. doi: 10.1117/12.872645
- [51] A. Kratz, M. Hadwiger, A. Fuhrmann, R. Splechtna, and K. Bühler. GPU-based high-quality volume rendering for virtual environments. In *International Workshop on Augmented Environments for Medical Imaging and Computer Aided Surgery (AMI-ARCS)*, 2006.

- [52] H. Lam. A framework of interaction costs in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1149–1156, Nov./Dec. 2008. doi: 10.1109/TVCG.2008.109
- [53] F. Lauber, A. Follmann, and A. Butz. What you see is what you touch: Visualizing touch screen interaction in the head-up display. In *Proc. DIS*, pp. 171–180. ACM, New York, 2014. doi: 10.1145/2598510.2598521
- [54] V. Levesque, L. Oram, K. MacLean, A. Cockburn, N. D. Marchuk, D. Johnson, J. E. Colgate, and M. A. Peshkin. Enhancing physicality in touch interaction with programmable friction. In *Proc. CHI*, pp. 2481–2490. ACM, New York, 2011. doi: 10.1145/1978942.1979306
- [55] N. Lindow, D. Baum, and H.-C. Hege. Perceptually linear parameter variations. *Computer Graphics Forum*, 31(2):535–544, May 2012. doi: 10.1111/j.1467-8659.2012.03054x
- [56] Z. Liu, N. Nersessian, and J. Stasko. Distributed cognition as a theoretical framework for information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1173–1180, Nov. 2008. doi: 10.1109/TVCG.2008.121
- [57] D. López, L. Oehlberg, C. Doger, and T. Isenberg. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(5):1616–1629, May 2016. doi: 10.1109/TVCG.2015.2440233
- [58] K. T. McDonnell and K. Mueller. Illustrative parallel coordinates. *Computer Graphics Forum*, 27(3):1031–1038, May 2008. doi: 10.1111/j.1467-8659.2008.01239x
- [59] S. Meyer, O. Cohen, and E. Nilsen. Device comparisons for goal-directed drawing tasks. In *Proc. CHI*, pp. 251–252. ACM, New York, 1994. doi: 10.1145/259963.260468
- [60] M. A. Muhanna. Virtual reality and the CAVE: Taxonomy, interaction challenges and research directions. *Journal of King Saud University – Computer and Information Sciences*, 27(3):344–361, July 2015. doi: 10.1016/j.jksuci.2014.03.023
- [61] T. Munzner. *Visualization Analysis and Design*. A K Peters Visualization Series. CRC Press, Boca Raton, FL, USA, 2014. doi: 10.1201/b17511
- [62] D. A. Norman. *The Design of Everyday Things*. Doubleday, New York, 1990.
- [63] Y. Ochiai, K. Kumagai, T. Hoshi, J. Rekimoto, S. Hasegawa, and Y. Hayasaki. Fairy lights in femtoseconds: Aerial and volumetric graphics rendered by focused femtosecond laser combined with computational holographic fields. *ACM Transactions on Graphics*, 35(2):17:1–17:14, Feb. 2016. doi: 10.1145/2850414
- [64] I. Poupyrev and S. Maruyama. Tactile interfaces for small touch screens. In *Proc. UIST*, pp. 217–220. ACM, New York, 2003. doi: 10.1145/964696.964721
- [65] B. Preim and C. P. Botha. *Visual Computing for Medicine: Theory, Algorithms, and Applications*. Morgan Kaufmann Publishers Inc., San Francisco, 2nd ed., 2013. doi: 10.1016/C2011-0-05785-X
- [66] H. C. Purchase, N. Andrienko, T. J. Jankun-Kelly, and M. Ward. Theoretical foundations of information visualization. In A. Kerren, J. T. Stasko, J.-D. Fekete, and C. North, eds., *Information Visualization: Human-Centered Issues and Perspectives*, pp. 46–64. Springer, Berlin/Heidelberg, 2008. doi: 10.1007/978-3-540-70956-5_3
- [67] U. Rashid, M. A. Nacenta, and A. Quigley. Factors influencing visual attention switch in multi-display user interfaces: A survey. In *Proc. PerDis*, pp. 1:1–1:6. ACM, New York, 2012. doi: 10.1145/2307798.2307799
- [68] G. Robles-De-La-Torre. The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia*, 13(3):24–30, July–Sept. 2006. doi: 10.1109/MMUL.2006.69
- [69] E. Rousset, F. Bérard, and M. Ortega. Study of the effect of the directness of the interaction on novice users when drawing straight lines. In *Proc. IHM*, pp. 14:1–14:7. ACM, New York, 2015. doi: 10.1145/2820619.2820633
- [70] P. Sabella. A rendering algorithm for visualizing 3D scalar fields. *ACM SIGGRAPH Computer Graphics*, 22(4):51–58, June 1988. doi: 10.1145/378456.378476
- [71] R. Sadana and J. Stasko. Designing and implementing an interactive scatterplot visualization for a tablet computer. In *Proc. AVI*, pp. 265–272. ACM, New York, 2014. doi: 10.1145/2598153.2598163
- [72] D. Schmidt, F. Block, and H. Gellersen. A comparison of direct and indirect multi-touch input for large surfaces. In *Proc. INTERACT*, pp. 582–594. Springer, Berlin/Heidelberg, 2009. doi: 10.1007/978-3-642-03655-2_65
- [73] A. Sears and B. Shneiderman. High precision touchscreens: Design strategies and comparisons with a mouse. *International Journal of Man-Machine Studies*, 34(4):593–613, Apr. 1991. doi: 10.1016/0020-7373(91)90037-8
- [74] B. Shneiderman. The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology*, 1(3):237–256, 1982. doi: 10.1080/01449298208914450
- [75] B. Shneiderman. Direct manipulation: A step beyond programming languages. *IEEE Computer*, 16(8):57–69, Aug. 1983. doi: 10.1109/MC.1983.1654471
- [76] K. Shoemake. ARCBALL: A user interface for specifying three-dimensional orientation using a mouse. In *Proc. Graphics Interface*, pp. 151–156. Morgan Kaufmann Publishers Inc., San Francisco, 1992. doi: 10.20380/GI1992.18
- [77] D. Silver. Object-oriented visualization. *IEEE Computer Graphics and Applications*, 15(3):54–62, May 1995. doi: 10.1109/38.376613
- [78] A. L. Simeone and H. Gellerseny. Comparing indirect and direct touch in a stereoscopic interaction task. In *Proc. 3DUI*, pp. 105–108. IEEE Computer Society, Los Alamitos, 2015. doi: 10.1109/3DUI.2015.7131733
- [79] D. Smilkov, S. Carter, D. Sculley, F. Viégas, and M. Wattenberg. Direct-manipulation visualization of deep networks. In *KDD 2016 Workshop on Interactive Data Exploration and Analytics (IDEA)*, pp. 115–119, 2016. Also available as arXiv report # 1708.03788.
- [80] P. Song, W. B. Goh, C.-W. Fu, Q. Meng, and P.-A. Heng. WYSIWYF: Exploring and annotating volume data with a tangible handheld device. In *Proc. CHI*, pp. 1333–1342. ACM, New York, 2011. doi: 10.1145/1978942.1979140
- [81] M. Spindler, M. Martsch, and R. Dachsel. Going beyond the surface: Studying multi-layer interaction above the tabletop. In *Proc. CHI*, pp. 1277–1286. ACM, New York, 2012. doi: 10.1145/2207676.2208583
- [82] M. Spindler, C. Tominski, H. Schumann, and R. Dachsel. Tangible views for information visualization. In *Proc. ITS*, pp. 157–166. ACM, New York, 2010. doi: 10.1145/1936652.1936684
- [83] N. Sultanum, S. Somanath, E. Sharlin, and M. C. Sousa. “Point it, split it, peel it, view it”: Techniques for interactive reservoir visualization on tabletops. In *Proc. ITS*, pp. 192–201. ACM, New York, 2011. doi: 10.1145/2076354.2076390
- [84] E. Sundén, P. Steneteg, S. Kottravél, D. Jonsson, R. Englund, M. Falk, and T. Ropinski. Inviwo – An extensible, multi-purpose visualization framework. In *Posters of VIS*, pp. 163–164, 2015. doi: 10.1109/SciVis.2015.7429514
- [85] I. E. Sutherland. A head-mounted three dimensional display. In *Proc. AFIPS (Fall, part I)*, pp. 757–764. ACM, New York, 1968. doi: 10.1145/1476589.1476686
- [86] C. Upson, T. A. Faulhaber, D. Kamins, D. Laidlaw, D. Schlegel, J. Vroom, R. Gurwitz, and A. van Dam. The application visualization system: A computational environment for scientific visualization. *IEEE Computer Graphics and Applications*, 9(4):30–42, July 1989. doi: 10.1109/38.31462
- [87] J. J. van Wijk. The value of visualization. In *Proc. VIS*, pp. 79–86. IEEE Computer Society, Los Alamitos, 2005. doi: 10.1109/VISUAL.2005.1532781
- [88] J. J. van Wijk. Views on visualization. *IEEE Transactions on Visualization and Computer Graphics*, 12(4):421–432, July/Aug. 2006. doi: 10.1109/TVCG.2006.80
- [89] P. Vickers, J. Faith, and N. Rossiter. Understanding visualization: A formal approach using category theory and semiotics. *IEEE Transactions on Visualization and Computer Graphics*, 19(6):1048–1061, June 2013. doi: 10.1109/TVCG.2012.294
- [90] C. Wang and H.-W. Shen. Information theory in scientific visualization. *Entropy*, 13(1):254–273, Jan. 2011. doi: 10.3390/e13010254
- [91] M. Ward and J. Yang. Interaction spaces in data and information visualization. In *Proc. VISSYM*, pp. 137–146. Eurographics Association, Goslar, Germany, 2004. doi: 10.2312/ViSym/ViSym04/137-146
- [92] A. Wiebel, T. Isenberg, S. Bruckner, and T. Ropinski. Tools and techniques for direct volume interaction. In M. Zwicker and C. Soler, eds., *Eurographics Tutorials*. Eurographics Association, Goslar, Germany, 2015. doi: 10.2312/egt.20151047
- [93] A. Wiebel, F. M. Vos, D. Foerster, and H.-C. Hege. WYSIWYP: What you see is what you pick. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2236–2244, Dec. 2012. doi: 10.1109/TVCG.2012.292
- [94] G. Wills. Linked data views. In C.-h. Chen, W. Härdle, and A. Unwin, eds., *Handbook of Data Visualization*, chap. II.9, pp. 217–241. Springer, Berlin/Heidelberg, 2008. doi: 10.1007/978-3-540-33037-0_10
- [95] L. Xu, T. Y. Lee, and H. W. Shen. An information-theoretic framework for flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1216–1224, Nov. 2010. doi: 10.1109/TVCG.2010.131
- [96] J. S. Yi, Y. ah Kang, J. T. Stasko, and J. A. Jacko. Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1224–1231, Nov./Dec. 2007. doi: 10.1109/TVCG.2007.70515

- [97] A. Ynnerman, T. Rydell, D. Antoine, D. Hughes, A. Persson, and P. Ljung. Interactive visualization of 3D scanned mummies at public venues. *Communications of the ACM*, 59(12):72–81, Dec. 2016. doi: 10.1145/2950040
- [98] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. Efficient structure-aware selection techniques for 3D point cloud visualizations with 2DOF input. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2245–2254, Dec. 2012. doi: 10.1109/TVCG.2012.217
- [99] L. Yu, K. Efstathiou, P. Isenberg, and T. Isenberg. CAST: Effective and efficient user interaction for context-aware selection in 3D particle clouds. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):886–895, Jan. 2016. doi: 10.1109/TVCG.2015.2467202
- [100] L. Yu, P. Svetachov, P. Isenberg, M. H. Everts, and T. Isenberg. FI3D: Direct-touch interaction for the exploration of 3D scientific visualization spaces. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1613–1622, Nov./Dec. 2010. doi: 10.1109/TVCG.2010.157
- [101] C. Ziemkiewicz, R. J. Crouser, A. R. Yauilla, S. L. Su, W. Ribarsky, and R. Chang. How locus of control influences compatibility with visualization style. In *Proc. VAST*, pp. 81–90. IEEE Computer Society, Los Alamitos, 2011. doi: 10.1109/VAST.2011.6102445



Alexander Wiebel is professor at Hochschule Worms University of Applied Sciences where he is also co-heading the research group *User Experience & Visualization*. He received his Ph.D. from Universität Leipzig in 2008 for research on flow visualization. As a postdoctoral researcher at Max Planck Institute for Human Cognitive and Brain Sciences, and later at Zuse Institute Berlin (ZIB), he conducted research in interactive visualization of 3D MRI data with a focus on intuitive selection of structures in direct volume renderings. During this time, he lectured at Universität Leipzig and Freie Universität Berlin. Between 2013 and 2015 he was a professor of visual computing at Coburg University.



Stefan Bruckner is professor in visualization at the Department of Informatics of the University of Bergen, Norway. He received his master's degree in Computer Science from the TU Wien, Austria in 2004 and his Ph.D. in 2008 from the same university. He was awarded the habilitation (venia docendi) in Practical Computer Science in 2012. From 2008 to 2013, he was an assistant professor at the Institute of Computer Graphics and Algorithms at TU Wien. His research interests

include interactive visualization techniques for spatial data, particularly in the context of biomedical applications, visual parameter space analysis, illustrative methods, volume visualization, and knowledge-assisted visual interfaces.



Tobias Isenberg is a senior research scientist at Inria, France. He received his doctoral degree from the University of Magdeburg, Germany, in 2004. Previously he held positions as post-doctoral fellow at the University of Calgary, Canada, and as assistant professor at the University of Groningen, the Netherlands. His research interests comprise topics in scientific visualization, illustrative and non-photorealistic rendering, and interactive visualization techniques. He is particularly interested in interactive visualization

environments for 3D spatial data that rely on novel input paradigms such as tactile screens and tangible devices.



Timo Ropinski heads the Visual Computing Group at Ulm University. Before moving to Ulm, he was Professor in Interactive Visualization at Linköping University in Sweden where he was heading the Scientific Visualization Group. Timo Ropinski has received his Ph.D. in computer science in 2004 from the University of Münster, where he has also completed his Habilitation in 2009. His research foci lie in the areas of computer graphics and visualization. He is especially interested in interactive visual problem solving,

multimodal and multivariate visualization, volume rendering as well as volumetric illumination models, all with a special consideration of large data sets.