

Embedding Information Visualization Within Visual Representation

Caroline Ziemkiewicz and Robert Kosara

Abstract Despite its often technical nature, visualization is in many ways a form of visual representation. Just how visualization relates to illustration, information graphics, digital art, visual languages, etc., is nonetheless poorly understood. We propose a theory that embeds information visualization within other visual traditions in terms of criteria that are not purely technical: dependence on data, mapping, interactivity, and notationality. In addition to providing the means for a classification, these criteria also foster a different understanding of information visualization. We further adapt our criteria to differentiate within visualization, using mapping, readability and information loss, and notationality as the criteria. Both sets of criteria are demonstrated in a number of case studies. We believe that our novel taxonomies of visualization methods serve as a step towards a more comprehensive theoretical context to understanding the essential purposes, properties, and functions of information visualization.

1 Introduction

Information visualization is a young field which draws from disparate sources in computer science, design, statistics, data mining, psychology, and the visual arts. While interest in visualization has grown greatly in recent years, its appropriate position within this broader space is unclear, as is the nature of the boundaries that set it off from its neighbors. In order for new research and new ideas to be understood within this larger context, a theoretical foundation is needed. General computing theory breaks algorithms down to their essential properties, identifying those which

Caroline Ziemkiewicz
UNC Charlotte, e-mail: caziemki@uncc.edu

Robert Kosara
UNC Charlotte e-mail: rkosara@uncc.edu

are fundamentally similar or different independent of implementation. Similarly, a theory of visualization which identifies the primitives and essential properties of visualization methods could offer new perspectives on current problems and identify paths that lead forward from existing research.

This theory should include a definition of information visualization as distinct from other forms of visual representation. Often, visualization is defined as any translation of data to image, but intuitively its purpose seems substantially different from those of information graphics or data-based artworks. Locating visualization within the larger context of visual representation will contribute to a greater understanding of fundamental research goals.

Also needed are meaningful divisions among specific data visualization methods. Traditional categories within visualization, most notably scientific visualization, information visualization, and visual analytics, are themselves vaguely defined and tend to overlap. The term *data visualization* itself is ambiguous, as it sometimes refers to all of the above and is sometimes used to refer specifically to scientific visualization and not information visualization. For our purposes, *data visualization* will be used to refer to all visualization derived from data, whether that data is abstract or bound to locations. Our primary focus is on information visualization (InfoVis), although our theory has applications outside of that focus. More precise definitions are needed not only to clarify the concepts already in use, but also to show in sharper relief potential research areas which exist beyond those concepts, as suggested by Tory and Möller [24].

With this in mind, we have constructed a map of the visualization field that contains not only the internal structures of the field, but also the outer borders that separate visualization from similar ventures such as infographics, visual languages, and data art. Our theory, which draws from existing visualization models as well as broader theories of aesthetics and visual grammar, takes active readability of information to be the central goal of visualization, distinguishing it from other categories of visual representation. Readability is defined in terms of four criteria. To be readable, a visualization must be a mapping from property data to a corresponding image. This mapping must be *bijective*, so that there is a one-to-one correspondence between elements of the image and pieces of information. We also argue that readability requires nontrivial interactivity, so that the possible information a visualization can convey is not fixed by the designer. Finally, we argue that information visualization in particular must be *syntactically notational* by Nelson Goodman's definition [6], meaning that it is composed of discrete and disjoint visual symbols. We use these criteria to define both information visualization and several closely related classes of visual representation.

We then propose a set of three classifications based on aspects of readability in order to formalize the internal distinctions among visualization methods. These include whether the mapping is linear or nonlinear, whether it is designed to minimize or control information loss, and whether the visualization is discrete or continuous, which we define formally using Goodman's idea of *semantic notationality*. We go on to apply this taxonomy to some key examples from existing visualization research, and show how it can form meaningful groupings of visualization methods

based on type of readability that echo and expand upon earlier taxonomies. We believe that this classification will serve to inspire innovative research directions and illuminate new connections within existing research. We also hope to stimulate interest in a more comprehensive theoretical basis to visualization, leading to greater understanding of the field as a whole.

2 Related Work

Our work builds on taxonomies of visualization methods and also takes ideas from work in visual languages and the automated design of visualizations.

2.1 *Visual Mappings and Languages*

A major goal of our work is to understand the primitives of visual representation and the mappings by which they are combined into meaningful symbol systems. A seminal inspiration for this area of research is Bertin [3], who analyzes diagrams and maps as a semiotic system of marks, retinal encodings, and spatial properties. Other formal theories of visual mapping include Goodman [6], whose definition of a notational system is meant to distinguish formal visual languages, such as those that make up certain kinds of diagrams, from the broader class of representational pictures. These two theories serve as important foundations for our work.

One relevant branch of research is that which strives to create automated or semi-automated visualization systems based on formalized systems of visual mapping—even though we do not want to go nearly as far. This includes Mackinlay’s APT system for automatically generating visualizations based on properties of the data [16], a concept which has recently been incorporated into the commercial software Tableau [17]. Such applications suggest the immediate practicality of creating consistent and precise theories of visual representation.

2.2 *Visualization Taxonomies*

In addition to understanding visual mappings as they relate to information visualization, our work builds upon previous efforts to construct a taxonomy of visualization methods. Shneiderman [22] introduces a 2-axis classification of methods based on type of task and type of data. A similar taxonomy is proposed by Card and Mackinlay [4], who combine a more formal definition of data types with a system of visual properties drawn from Bertin [3].

Chi [5] takes a different approach by focusing on the transformation from data to visual representation, and goes on to differentiate visualization techniques by

the actions taken at each transformation step. Tory and Möller [24] build on previous work that considers the degree of spatial constraint and continuity of data as important axes, but take a more usability-inspired perspective in viewing these as attributes of the designer’s internal model of the data rather than inherent properties of the data themselves.

Our work differs from these previous taxonomies in our focus on visual mapping as independent from data and task type. In this sense it is most similar to Chi, although our interest lies in defining properties of the mapping as a whole, while Chi focuses on dividing it into stages. We extend Kosara’s [12] definition of visualization primarily through the fact that it is a representation of data that is readable, i.e., that provides a mapping from the visual representation back to the data. Our work is also unique in that it attempts to place visualization within the larger realm of visual representation. We believe that this effort is vital in guiding progress within the field, by illuminating both the limitations and unexplored directions of data visualization.

3 Visualization in Visual Representation

Our theory covers two major subjects: first, how to distinguish visualization from other forms of visual representation, and second, how the concepts from this definition may be used to differentiate visualization methods from one another. Building on previous work [12], we have developed a definition of information visualization as that class of visual representation which:

- is derived from data,
- represents a bijective mapping from information to image,
- provides for nontrivial interactivity, and
- is a syntactically notational symbol system.

We elaborate on each of these items and show how they contribute to the overall definition of readability.

3.1 *Data-Driven*

A data visualization must consist of a mapping function from some existing data source to an image. We define this mapping as a two-step transformation: in the first step, data is processed to produce information in the form of an analytical abstraction, and in the second step, the information is transformed to visual elements. This view of the visual mapping follows that in Chi [5], although for our purposes we treat the visual mapping as a single step rather than dividing it into a visualization transformation and a visual mapping transformation.

The requirement that a data visualization be based on data is perhaps an obvious one, although the definition of *data* may require more elaboration. For example,

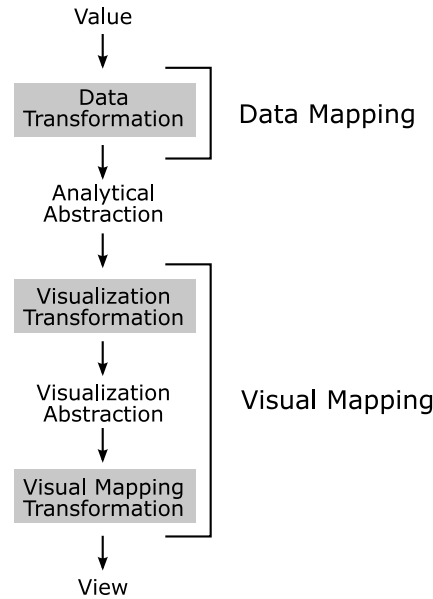


Fig. 1 Our model of the visual mapping process, which is a simplified version of the data state reference model [5].

a mathematical visualization that shows a fractal expanding over time may be a visual representation of information, but not of data in the sense that visualization researchers generally use the term. Generally, the term *data* in our sense refers to a collection of properties about objects or concepts which exists independently of the visualization. This can come in the form of tabular data containing attributes of some collection of subjects, scientific data which represent a sampling of the properties of an object or phenomenon, text, or multimedia data which contain properties of images or sound.

It may be useful, then, to make a distinction between data visualization, which represents properties of information, and function visualization, which represents the visual analog of a process. In the latter case, the information being visualized is treated as a definition of the mapping itself, while we are concerned with cases in which the mapping is independent of specific data. For our purposes, we use the term *visualization* to mean data visualization, and consider function visualizations such as mathematical visualizations, algorithm visualizations, visual programming languages, and visual proofs to be separate ventures.

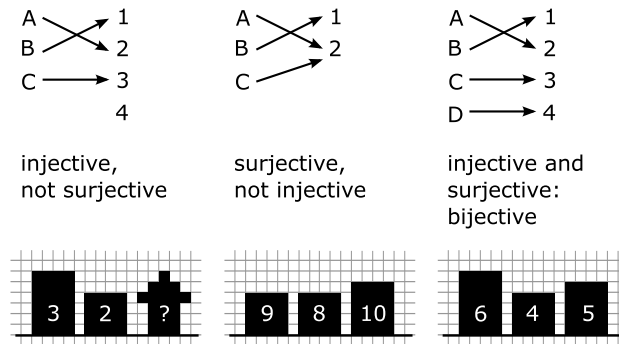


Fig. 2 Three types of mapping (Section 3.2) and resulting information loss (Section 5.2, Figure 4), using the example of a bar chart.

3.2 Bijective Mapping

Although a visualization must be the result of a mapping function, not all visual mappings necessarily result in visualizations. We take the concept of injective, surjective, and bijective functions from mathematics (Figure 2) as an analogy in distinguishing visual mappings that form visualizations from those which do not.

A function is defined as a mapping from a domain to a codomain. Those mappings in which each element in the (source) domain is mapped to a unique element in the (target) codomain are injective. In a non-injective mapping, several elements in the domain may map to one element in the codomain. Because of this, a non-injective mapping is irreversible, i.e., there does not exist a function which perfectly inverts the mapping. Mappings in which each element in the codomain is mapped to by some element in the domain are surjective. This means that there are no possible values in the codomain which do not correspond to a value in the domain. Those mappings which are both injective and surjective are bijective, so that every element in the domain maps to exactly one element in the codomain, and every element in the codomain can be mapped back to a single element in the domain.

We use these classifications as a metaphor in defining properties of the visual mapping step. In this metaphor, the domain represents the information extracted from the data in the processing step and the codomain is the set of image elements in the visual representation. An injective visual mapping is therefore one in which every piece of information maps to a single visual element, although there can appear visual elements that do not correspond to information. Likewise, a surjective visual mapping consists only of visual elements that relate to information, although each unique visual element can correspond to one or more pieces of information. And a bijective visual mapping is one in which every piece of information corresponds to exactly one unique visual element. In Section 3.4 we discuss the definition of visual elements more thoroughly.

We find that both injectivity and surjectivity are necessary conditions for a visual representation of information to be readable. In a non-injective mapping from

information to image, more than one piece of information can be mapped to a single unique visual element. Since this can lead to ambiguity and information loss (Section 5.2), such a mapping would not be readable. Non-surjectivity, on the other hand, implies that there can be significant elements in the visual representation that are not derived from the data. These elements may be accidental artifacts of the visual mapping, or intentional aesthetic elements included for decorative rather than informational purposes (what Tufte calls *chartjunk* [25]). Since such elements can distract or mislead a user, we require that the visual mapping be surjective as well. An exception should be made, however, for visual elements that provide information about the mapping itself, such as legends, scales, or labels, provided they do not interfere with readability.

3.3 Nontrivial Interactivity

We argue that interactivity is an essential aspect of active readability. By *active readability* we mean that a user should be able to actively seek information, rather than merely accessing existing information determined by the designer. This would exclude static infographics, diagrams, illustrations, and maps from the class of visualization. The ability for the user to control her view of the data within the context of the given visual mapping has long been considered one of the foundations of visualization, as evidenced by Shneiderman’s oft-cited Visual Information-Seeking Mantra [22] and the emphasis on focus-and-context strategies in both information and scientific visualization.

	InfoVis	SciVis	Infographics	Function Vis	Vis Art	Iconography	Art
Data-Driven	Yes	Yes	Yes	No	Yes	No	No
Type of Mapping	Bijjective	Sur- or Bijjective	In- or Bijjective	None	Any	None	None
Interactive	Yes	Yes	No	Maybe	Maybe	No	Maybe
Syntactically Notational	Yes	No	Yes	Maybe	No	Yes	No

Table 1 Embedding information visualization within visual representation, depending on whether or not a representation is data-driven, its mapping type, interactivity, and syntactic notationality.

Additionally, we distinguish between trivial and nontrivial interaction, which changes the parameters of the visual mapping itself. For example, semantic zooming, moving from node to node on a hyperbolic tree, and multidimensional scaling would be nontrivial interactions. While a visual representation may include any number of interaction methods, both trivial and nontrivial, the presence of at least some nontrivial interactivity is an especially strong argument for a given representation being a visualization.

3.4 Syntactically Notational

As discussed in Sack [20], the usability goals of visualization have been described as opposed to ideas of the “sublime” as a criterion for aesthetic works. Although Sack proposes several other aesthetic criteria by which to judge artistic visualizations, we distinguish between data-based representations with the goal of readability and those which are to be read as more pictorial depictions. Therefore, for the last item in our definition of visualization, we have drawn from Nelson Goodman’s definition of notational systems. A notational system is a visual mapping in which discrete and disjoint concepts map to discrete and disjoint symbols, making the system a visual grammar of sorts, rather than an image which is to be understood in a more holistic, pictorial fashion. For example, musical notation, in which each mark maps to a unique musical note, is a notational system, while photography is not. Goodman’s formal criteria for a notational system, as summarized by Lee [14], are as follows:

1. *The system must consist of symbols (utterances, inscriptions, marks) that form equivalence classes (characters), which can be exchanged without syntactical effect. [...] The characters have to be disjoint, so that no mark qualifies as an instance of more than one character. [...]*
2. *Characters have to be “finitely differentiable” (or “articulate”) in the sense that their disjointness is feasibly testable; this rules out, in particular, “dense” systems in which any two (ordered) characters have another between them.*
3. *Notational systems must be unambiguous, so that the extension (i.e., what is referred to, which Goodman calls the “compliance-class”) of an inscription is invariant with respect to time, context, and so on.*
4. *The compliance-classes of all characters must be disjoint. (Also, the system will ideally be nonredundant.)*
5. *Compliance-classes must also be finitely differentiable. Thus, for example, any system that is “semantically dense,” in that its compliants form an ordering such that any two have another between them, is excluded.*

We do not require that any given information visualization be a complete notational system, but that it be at least syntactically notational, meaning that it fulfills the first two criteria in Goodman’s definition. The discrete characters to which Goodman refers can be thought of as glyphs [2], where the term “glyph” is being used more generally than usual, to refer to the primitive symbols of any visualization. That is, a glyph is the smallest portion of a visualization whose visual attributes are fully meaningful in terms of mapping data.

We can use Goodman’s definition of notational characters, then, as a definition of glyphs as visualization primitives. The claim that glyphs must be disjoint implies that each possible glyph must qualify as a unique character regardless of its placement in the image, and no mark in the image qualifies as more than one unique glyph. The claim that they are finitely differentiable implies that the space of possible glyphs must be discrete. Glyphs in our formulation need not be the multidimensional icons to which the term is most often applied; they could just as well be bars, points, lines, areas, pixels, or voxels. As long as they are discrete pieces of visual information that can be differentiated from other elements of the visual field, they

qualify. This is even true for the extreme case of pixel-based visualization [11], in which each single pixel can be considered a glyph.

4 Applying the Definition

The common definition of visualization as the presentation of information in visual form can be argued to include a large variety of visual representations with wildly differing goals. A partial list of these representations, and their classification according to our system, can be seen in Table 1. According to our definition, information visualization contains those visual representations which are based upon property data, represent a bijective mapping of information to image, include non-trivial interactivity, and are syntactically notational, distinguishing it from several closely related categories. While some of these distinctions are self-explanatory, others warrant further discussion.

4.1 *Scientific Visualization*

Data visualization as a whole can be seen as comprising all visual representations that fit the first three criteria: being data-driven, having a bijective mapping, and being nontrivially interactive. Through the final requirement of syntactic notationality, our theory makes a distinction between information visualization and most kinds of scientific visualization as the terms are usually used. Many instances of flow and especially volume visualization cannot be broken into disjoint glyphs. In all these cases, the continuous nature of the visual representation makes the resulting image pictorial rather than notational. Understanding a scientific visualization is more like taking in a picture than like reading a visual language.

4.2 *Information Graphics*

Information graphics, or infographics, are often mentioned in connection with information visualization, and there are cases where it is unclear which category a given visual representation falls under. The term historically refers to visual presentation of data in the context of a newspaper or magazine, however. It may be thought of as a useful category covering those visual representations whose intent is communication of fixed information rather than interaction with fluid information: in other words, readability, but not active readability. Therefore, the main distinction between infographics and information visualization is the lack of interactivity. In addition, many infographics as seen in magazines and newspapers contain decorative elements which detract from the bijectivity of the representation.

4.3 *Function Visualization*

Another category closely related to data visualization contains what might be thought of as function visualizations. Visual analogues of mathematical functions, algorithms, or logical proofs are not based on property data, and as such are not readable in our sense. A visualization of a fractal, for example, is a representation of the equation on which it is based, but it is not possible to retrieve that function by looking at the image. Since we see visualization as the mapping from properties of data to properties of image, visualization of processes represents an interesting but separate field of visual representation. Visual programming languages may fall under this category as well, since they are closely related in purpose to algorithm visualization and visual proofs.

4.4 *Artistic Visualization*

Artistic visualizations, such as *Artifacts of the Presence Era* [27] or *The Dumpster* [15], provide a final borderline case. While these vary widely, in general they treat data in a visualization-like way with the goal of achieving an aesthetic rather than a readable effect. *Artifacts* can perhaps be broken into glyphs in the form of the “strata” that build on top of one another to show the images and sound waves captured in a museum over time. These strata are not, however, finitely differentiable; since the glyphs are based directly on the image and sound wave, and the precise nature of the images and sounds captured can be nearly infinitely varied, the range of possible glyphs is syntactically dense, and therefore non-notational. In contrast, if properties of the image and sound (such as illumination, number of people visible, or volume) were mapped to visual properties, the disjointness of the glyphs could be feasibly tested.

Our theory also suggests an additional answer to the questions raised by Skog et al. [23]. This ambient visualization of bus schedules inspired by the art of Piet Mondrian was not automatically perceived by passersby as being a visualization at all, but rather was often assumed to be just a piece of digital art. Since only a few of the many visual elements on display in the image corresponded directly to data, the mapping was not surjective. This may be a characteristic of many decorative or artistic visualizations where visual elements are chosen for aesthetic reasons (including creating a sublime experience, Section 3.4) rather than those of usability. As Sack points out [20], the two goals can be seen as diametrically opposing each other. The fact that many passersby saw the work as aesthetic rather than readable may be a result of its non-bijectivity and non-notationality, in addition to its resemblance to existing artworks.




Fig. 3 A scatterplot is a linear mapping, because position has an absolute meaning. A scatterplot superimposed over another scatterplot (of the same scale) is no different from a scatterplot of both datasets combined, and multiplying all the values in a dataset by a constant factor yields a uniform stretch of the visual representation.

5 Differentiation Within Information Visualization

Having established the boundaries of the field of visualization, we will now examine how aspects and implications of that definition may be used as the foundation for a novel visualization taxonomy. This taxonomy is composed of three axes:

- whether the mapping is linear or nonlinear,
- the treatment of information loss in the mapping,
- and whether the system is semantically notational.

We believe that these axes – while not immediately obvious – encapsulate many of the intuitive dimensions on which researchers tend to sort visualization methods.

5.1 Linear and Nonlinear Mappings

Another characteristic of the visual mapping function that illuminates fundamental differences among visualization methods is again based on an analogy from mathematical functions. In this case, we explore what it means for a visual mapping to be linear or nonlinear.

In mathematics, a linear transform is one which has the following properties [19]:

- Additivity, so that $f(x + y) = f(x) + f(y)$
- Homogeneity, so that $f(ax) = af(x)$

We extend these properties to visual mappings as follows. A visual mapping is additive if a visualization of one dataset, combined visually with the same visualization of another dataset, is equivalent to the visualization of the union of the two datasets. Likewise, a visual mapping is homogenous if a constant transformation over the data would yield a similarly constant transformation over the image.

Because this is only possible when a given data point's mapped location and appearance remain constant when other portions of the dataset change, this implies that a linear visual mapping is one in which each visual element of a glyph has absolute meaning. In nonlinear mappings, in which a glyph's location and appearance change with changes in the dataset, information in the visualization must arise to some degree from visual relationships among glyphs.

5.2 *Mapping and Information Loss*

The mapping of data to its visual representation is a central aspect of visualization, and thus also a criterion for differentiating within visualization. In the definition section, we discussed the requirement that a visualization implement a bijective mapping of data to visual parameters (Section 3.2). But what exactly does this mean for a practical visualization that uses a real-world output system with a limited resolution both spatially and in terms of color? Visualization leads to a loss of information in most cases, but exactly how much is lost, and when do we accept this?

We distinguish between two types of information loss: intentional and unintended.

Intentional Information Loss

Information loss is intended in many cases where data is aggregated so as to get a better overview over the structure of the data rather than its exact details (e.g., histograms [13], pivot graphs [29], etc.). Information is also intentionally lost in the case of non-numerical data like text documents. Many non-linear mappings (see above) like the multi-dimensional scaling used in IN-SPIRE [30] reduce the data to a tiny fraction of the original to be able to visualize them at all.

Unintended Information Loss

Unintended information loss comes from the limits of the output device and/or human perception. All current displays use pixels of finite size as their basic elements, and have a limited number of colors or shades of gray they can show. Human perception is certainly also limited in the number of points it can discern, and more than that in the number of colors that can be differentiated [28].

There are many different kinds of unintended information loss, which can result from projections of higher dimensional data, overplotting (which can have different reasons), etc. For the purposes of our model, we only investigate the loss of information caused by the quantization of values when they are displayed on an output device.

A linear transformation of data to visual properties on screen (Figure 4) can be defined like the following. The n th data dimension $D_n = \{d_i\}$ of an unstructured data set contains a largest value d_{\max} and a smallest value d_{\min} . The data is mapped onto screen values s , which can be pixel positions or colors. For simplicity, we only consider a single visual attribute here, but this discussion can be easily generalized. The smallest value for s is s_{\min} , the largest s_{\max} (Figure 2) – these are always whole numbers. The square bracket indicates rounding to the nearest integer.

$$s(d) = \left\lceil \frac{d - d_{\min}}{d_{\max} - d_{\min}} (s_{\max} - s_{\min}) \right\rceil + s_{\min}$$

Especially in the case of color, a more general quantization function [7] may be introduced that further reduces the number of possible values.

$$s(d) = q \left(\frac{d - d_{\min}}{d_{\max} - d_{\min}} (s_{\max} - s_{\min}) \right) + s_{\min}$$

There is clearly a loss of information from the quantization that results from the rounding that needs to be performed when the value is drawn onto the screen; in the case of anti-aliasing, the rounding will be to a fraction of an integer, but the resolution of the display is still finite.

The resulting limited resolution limits the user's ability to read values (see below). We can quantify the resolution in data space as follows. All values within one bin of size Δd end up looking exactly the same to the user.

$$\Delta d = \left\lfloor \frac{d_{\max} - d_{\min}}{s_{\max} - s_{\min}} \right\rfloor$$

In case a quantization function q is used, and this function is a uniform quantizer, $\Delta d = \Delta_q$ (Δ_q being the bin size of the quantizer).

Actual information loss comes from the difference between data values and possible visual mappings. We define additional variables: $r_D = d_{\max} - d_{\min}$ is the range of values in the data dimension; u_D is the number of unique (distinct) values in data dimension D ; u_S is the number of distinct values on screen, which for the linear case means $u_S = s_{\max} - s_{\min} + 1$, and in case of using a quantization function $u_S = N_q(D_n)$, the number of distinct values the quantizer can produce for the given data dimension D_n – in the case of a uniform quantizer, $N_q(D_n) = \lceil \frac{r_D}{\Delta_q} \rceil$.

Information loss occurs in all cases where $u_D > u_S$, i.e., there are more distinct values than can be mapped into the visual domain. For $u_D \leq u_S$, there are two different cases (Table 2). If $r_D > u_S$, information loss occurs if the data is to be mapped onto a continuous scale. In the case of the data domain being categorical, however (which is rather likely for all $u_D \leq r_S$ cases), the data can be mapped without loss. The best case in terms of information loss is if $r_D \leq u_S$, because no information loss will occur: there are enough unique screen variables to which the data can be mapped.

Since a true bijective mapping is only occasionally possible, visualization designers must make a choice between trying to minimize the amount of information

Information Loss	$r_D \leq u_S$	$r_D > u_S$
$u_D \leq u_S$	No	Yes, unless categorical
$u_D > u_S$	Yes	Yes

Table 2 Information loss in the linear case, depending on: number of distinct screen values u_S , range of data values r_D , and the number of unique data values u_D .

Information Loss	Type of Mapping	Continuity	Examples
Loss Is Unintended	Linear	Discrete	Scatterplot, Parallel Coordinates [10], Bar Charts
		Continuous	Timelines [18]
	Nonlinear	Discrete	TreeMap [21], Cartograms, Node-Link Diagrams
		Continuous	Information Flocking Boids [26]
Loss Is Intentional	Linear	Discrete	Histograms [13], PivotGraph [29]
		Continuous	Brush Strokes [9]
	Nonlinear	Discrete	IN-SPIRE [30]
		Continuous	ThemeRiver [8], TextPool [1]

Table 3 Differentiation within information visualization along the axes of information loss, mapping, and continuity.

loss as much as possible and controlling information loss through data processing techniques such as clustering, summary, or filtering. Whether a visualization contains intentional information loss or only unintended information loss is the second axis in our taxonomy.

5.3 Semantic Notationality and Continuity

Following Tory and Möller, we see a discrete or continuous data model as an additional important variable in distinguishing among visualization methods. We further formalize this variable using the second part of Goodman’s definition of notationality (Section 3.4). As the first two items referred to syntactic notationality, the last three define a system with semantic notationality. A semantically notational system is one in which the meaning of a symbol is unambiguous and does not overlap with the meanings of other symbols. In terms of visualization, this means that any given glyph must maintain a distinct meaning regardless of context, and each glyph refers to a unique and disjoint portion of the data.

In visual mappings that are based on continuous data models, semantic notationality does not apply. A continuous data model implies the assumption that data points are not disjoint, but rather can be thought of as samplings along a continuum, so that the portions of the continuum referred to by any two consecutive sample points cannot be clearly divided. Furthermore, a continuous model implies that con-

text is significant in understanding the meaning of a data point. For example, in a visualization showing change over time along a timeline, each glyph’s meaning is dependent upon the change it represents from the glyphs that come before and after. Semantic notationality, therefore, can be used as a definition of the intuitive sense of continuity that some visualizations represent.




Fig. 4 An illustration of information loss in a linear mapping for two cases a) $u_D > u_S$ and b) $u_D < u_S$; with $r_D > u_S$ in both cases.

To elaborate further upon the notion of continuity, it is also useful to consider whether the data are treated as sampled or fully represented. In medical volume data, for example, the data are generally discrete images, but represent a sampling from a continuous structure; therefore, the design model usually treats the data as continuous and sampled, implying the need for interpolation. Data in the form of a continuous distribution function, on the other hand, could be treated as continuous and fully represented. This distinction applies similarly to discrete data. A dataset of attributes of different cars, for example, is usually treated as both discrete and fully represented. Census data, on the other hand, are a common example of discrete data which are often treated as sampled, since not all of the people represented are actually polled. Another example is stock market data, which can represent a sampling over time of the aggregation of many individual actions (i.e., buying and selling at specific prices).

Since a fully represented and continuous data model requires a continuous symbol space, a visual mapping that employs such a model is not syntactically notational and therefore lies outside of information visualization. Of the other three cases, only the discrete and fully represented data model is semantically as well as syntactically notational. Both discrete and sampled and continuous and sampled models are not semantically notational, and are therefore relatively more pictorial than other visual mappings.

A visualization method that employs a discrete and fully represented data model, and is therefore a complete notational system, is a more literal mapping of the data than one that does not. In other words, the metaphor being used is more direct, since it maps more closely to a translation from one language (that of the data, be it numerical or otherwise) to another (i.e., the symbol system of the visualization). As a visualization moves closer to Goodman’s idea of the pictorial, the mapping from data to image becomes less a translation than an abstraction.

6 Taxonomy of Visualization Methods

The three major axes described above map the design space of valid visualization methods according to our definition. In Table 3, we show a number of information visualization techniques sorted according to our three axes. This produces an informal spectrum of methods from the most direct at the top of the list to the most abstract at the bottom. We have divided this spectrum into four loosely defined groups as an initial taxonomy of visualization methods by type of readability. Note that these are not rigidly defined boundaries between subgroups, and that some methods may draw from several categories.

6.1 *Direct Translation*

The methods on the unintended loss, linear, and discrete end of the spectrum are the most direct translations from data to image. These methods perform minimal processing of the data, so that the visual mapping is executed directly on data points, rather than on intermediate pieces of information. In parallel coordinates [10], for example, each data point is projected to a unique glyph (in this case, a polyline); therefore, any information loss that occurs due to overplotting is unintentional, not inherent to the mapping. Likewise, this mapping is linear since the arrangement of each polyline glyph has absolute meaning, and it is discrete since there is no meaningful interpolation between data points. This is a direct translation because if there were no unintended information loss, one could theoretically reverse the mapping to produce the original data. Readability of data in this case is straightforward, except as it is confounded by information loss.

6.2 *Relative Meaning*

As mappings employ more intentional information loss, nonlinear mapping, and continuous data models, the visualization becomes more abstracted in different ways. In the case of TreeMaps [21], there is no intentional data loss, and the data model is discrete; however, the mapping is nonlinear since the position of a glyph (that is, a rectangle representing a node in the hierarchy) only has meaning relative to the position of other glyphs. Adding data to a hierarchy would result in an entirely different TreeMap layout. This makes for a slightly more indirect mapping than in the linear case. The readability of such a mapping is less trivial than in the case of more direct methods, because data is mapped not only to visual properties, but to relationships among visual properties. The user must look at more than one glyph to extract information about a single data point.

An unusual case which shows the abstraction in continuous methods is the Information Flocking Boids [26] system. In this visualization of time-varying stock

market data, glyphs are treated as agents which attempt to gather near similar glyphs without colliding with them. This is meant to produce a flocking behavior among clusters of similar stocks over time. Since the position of glyphs is so highly dependent on other glyphs, this is clearly nonlinear, and the importance of context to meaning indicates a continuous data model. While this method still attempts to show every data point, it represents a high level of abstraction. Since data properties are being mapped to physical attraction of glyphs rather than to simple image properties, readability is not a matter of reversing the mapping to data points, but rather of seeing overall patterns.

6.3 Aggregation

In mappings where the loss is intentional, steps are taken in the data processing stage to pare down the amount of information in order to reduce the amount of unintended loss. The simplest of these are histogram methods, which aggregate data into bins based on categorical dimensions or ranges of numerical dimensions, with size representing the number of members in each category or range. PivotGraphs [29] combine the idea of node-link diagrams with histogram representations, resulting in aggregated nodes and links, rather than nodes and links which refer to one data point or one relationship each. Properties of the dataset are readable, while individual properties of data points may be lost.

6.4 Abstraction

The most abstract mappings are found at the nonlinear and continuous end of the intentional-loss methods. This includes methods such as ThemeRiver [8], which employs a continuous model of highly filtered and aggregated document data in a nonlinear mapping. These methods sacrifice readability of any given data point for an understanding of the dataset as a whole, meaning they are moving closer to being pictorial rather than purely notational symbol systems.

While many useful methods have already been established in the more direct subgroups of our taxonomy, examples in the more abstracted subgroups are relatively rare. These areas may suggest directions for novel future research that seeks innovative mappings of data to image.

7 Conclusion

We present a theory in which visual representations are classified by properties of their mapping from data to image, rather than by properties of the image or proper-

ties of the data. This theory embeds visualization as a field within the larger visual culture, and serves as a basis for a taxonomy of information visualization methods. This work contributes the beginnings of a more solid theoretical foundation for the visualization field. Further work can build on this foundation by describing these mappings in more detail; for example, a framework of visual representation as visual metaphor has been fruitful in exploring cognition in information visualization [31]. Additionally, this theory provides a new perspective on visualization as a whole as well as on the similarities and differences among visualization methods. It can be used to guide design decisions and to inspire debate that will lead to greater understanding of the visualization field.

The application of our theory to existing forms of visual representation demonstrates its viability as an organizational tool. In particular, using the distinction between notational and pictorial systems, we have shown how the classic categories of scientific visualization and information visualization can be distinguished from one another without relying on properties of the data, which can be problematic [24]. This definition also allows us to formalize to some degree the intuitively apparent differences among data visualization, artistic visualization, and infographics, in addition to the less frequently considered category of function visualizations.

In extending our theory to an internal taxonomy of information visualization, we also build upon the work of Tory and Möller in constructing a new perspective on visualization which is less dependent upon inherent properties of the data. Specifically, their axis of discrete vs. continuous models of the data is included in our taxonomy, while we use the notion of semantic notationality to formalize and expand upon the original concept. Our axis of linear vs. nonlinear mappings also addresses some of the intent of Tory and Möller's axis of spatially given vs. spatially chosen data models. Both axes end up describing visualizations with varying levels of spatial constraint assumed in a given visualization, though the axis of linearity does so independent of data type.

Future work in this research includes a more thorough examination of the taxonomy of information visualization and a more formal characterization of the groupings it produces. This work may also serve as the foundation for a semi-automated tool for designing glyphs and visual representations, in order to test its practical applicability. Such an effort would entail a greater exploration of the definition of glyphs taken from Goodman's syntactic notationality.

Acknowledgments

This work was performed with partial support from the National Visualization and Analytics Center (NVACTM), a U.S. Department of Homeland Security Program, under the auspices of the SouthEast Regional Visualization and Analytics Center. NVAC is operated by the Pacific Northwest National Laboratory (PNNL), a U.S. Department of Energy Office of Science laboratory.

References

1. C. Albrecht-Buehler, B. Watson, and D. A. Shamma. Textpool: Visualizing live text streams. In *Proceedings of the IEEE Symposium on Information Visualization*, page 215.1, Washington, D.C., USA, 2004. IEEE Computer Society.
2. E. Anderson. A semigraphical method for the analysis of complex problems. *Proceedings of the National Academy of Sciences*, 43(923–927), 1957.
3. J. Bertin. *Semiology of Graphics*. University of Wisconsin Press, Madison, WI, USA, 1983.
4. S. K. Card and J. Mackinlay. The structure of the information visualization design space. In *Proceedings of the IEEE Symposium on Information Visualization*, page 92, Washington, DC, USA, 1997. IEEE Computer Society.
5. E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *Proceedings of the IEEE Symposium on Information Visualization*, page 69. IEEE Computer Society, 2000.
6. N. Goodman. *Languages of Art*. Hackett Publishing Company, 2nd edition, June 1976.
7. R. M. Gray and D. L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(6):2325–2385, October 1998.
8. S. Havre, B. Hetzler, and L. Nowell. Themeriver: Visualizing theme changes over time. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 115–123, Salt Lake City, UT, USA, 2000. IEEE Computer Society.
9. C. G. Healey, L. Tateosian, J. T. Enns, and M. Remple. Perceptually based brush strokes for nonphotorealistic visualization. *ACM Transactions on Graphics*, 23(1):64–96, January 2004.
10. A. Inselberg and B. Dimsdale. Parallel coordinates: A tool for visualizing multi-dimensional geometry. In *Proceedings of the First IEEE Conference on Visualization*, pages 361–378, San Francisco, CA, USA, 1990. IEEE Computer Society.
11. D. A. Keim. Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):59–78, 2000.
12. R. Kosara. Visualization criticism — the missing link between information visualization and art. In *Proceedings of the 11th International Conference on Information Visualisation (IV)*, pages 631–636. IEEE CS Press, 2007.
13. R. Kosara, F. Bendix, and H. Hauser. TimeHistograms for large, time-dependent data for large, time-dependent data. In *Joint Eurographics - IEEE TVCG Symposium on Visualization (VisSym 2004)*, pages 45–54, 2004.
14. J. Lee. Goodman’s aesthetics and the languages of computing. In P. A. Fishwick, editor, *Aesthetic Computing*, chapter 2, pages 29–42. MIT Press, April 2006.
15. G. Levin, K. Nigam, and J. Feinberg. The dumpster. <http://artport.whitney.org/commissions/thedumpster/>, 2006.
16. J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, 1986.
17. J. D. Mackinlay, P. Hanrahan, and C. Stolte. Show Me: Automatic presentation for visual analysis. *IEEE TVCG*, 13(6):1137–1144, 2007.
18. C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Lifelines: Visualizing personal histories. In *Proceedings Human Factors in Computing Systems (CHI)*, pages 221–227. ACM Press, 1996.
19. T. Rowland. Linear function. From MathWorld – A Wolfram Web Resource, created by Eric W. Weisstein. <http://mathworld.wolfram.com/LinearFunction.html>.
20. W. Sack. Aesthetics of information visualization. In C. Paul, V. Vesna, and M. Lovejoy, editors, *Context Providers*. University of Minnesota Press, 2007 (forthcoming).
21. B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Transactions on Graphics*, 11:92–99, 1992.
22. B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualization. In *Proceedings of the IEEE Symposium on Visual Languages*, pages 336–343. IEEE Computer Society, 1996.

23. T. Skog, S. Ljungblad, and L. E. Holmquist. Between aesthetics and utility: Designing ambient information visualizations. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 233–240. IEEE Computer Society, 2003.
24. M. Tory and T. Möller. Rethinking visualization: A high-level taxonomy. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 151–158. IEEE Computer Society, 2004.
25. E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2nd edition, 2001.
26. A. Vande Moere. Time-varying data visualization using information flocking boids. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 97–104. IEEE Computer Society, 2004.
27. F. B. Viegas, E. Perry, E. Howe, and J. Donath. Artifacts of the presence era: Using information visualization to create an evocative souvenir. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 105–111. IEEE Computer Society, 2004.
28. C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.
29. M. Wattenberg. Visual exploration of multivariate graphs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 822–819. ACM Press, 2006.
30. J. A. Wise. The ecological approach to text visualization. *Journal of the American Society for Information Science*, 50:1224–1233, 1999.
31. C. Ziemkiewicz and R. Kosara. The shaping of information by visual metaphors. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1269–1276, 2008. Honorable Mention in Best Paper.