# Understanding Visualisation: A Formal Foundation using Category Theory and Semiotics

Paul Vickers, *Member, IET* Joe Faith, and Nick Rossiter.

**Abstract**—This article combines the vocabulary of semiotics and category theory to provide a formal analysis of information visualisation. It shows how familiar processes of information visualisation fit the semiotic frameworks of both Saussure and Peirce, and extends these structures using the tools of category theory to provide a general framework for understanding information visualisation in practice, including: relationships between systems, data collected from those systems, renderings of those data in the form of representations, the reading of those representations to create visualisations, and the use of those visualisations to create knowledge and understanding of the system under inspection. The resulting framework is validated by demonstrating how many familiar concepts used in visualisation arise naturally from it; and used to identify some less intuitive distinctions which are useful in comparing visualisation methods. Finally, some suggestions are made regarding further uses to which this framework might be put, particularly as regards the study of multi- and cross-modal representations.

**Index Terms**—Visualization, category theory, semiotics

## 1 Introduction

Information visualisation is a catch-all term that embraces a wide range of activities that are all concerned with representing, or making visible (that is, perceptible), aspects or features of a given set of data or system. It is classically defined as "the process of forming a mental image of some scene as described" [20, p. 320]. Visualisation embraces a wide variety of activity, from the graphical analysis of scientific data, through the 'infographics' used to communicate in the popular media, to data art. It has recently grown in scale, popular currency, and theoretical discussion due to a combination of factors including the growth in the importance of data mining and processing in industry and science, and the availability of popular and powerful computer visualisation tools, such as Processing.[1]

Information visualisation in practice combines a range of skills and disciplines, including statistics, aesthetics, HCI, and computer science. And, perhaps because of this diversity, there has been relatively little discussion of the theoretical basis of the practice. Purchase et al. [19] remarked that information visualisation "suffers from not being based on a clearly defined underlying theory", and that "formal foundations are at a nascent stage". The danger of neglecting the theoretical foundations is that the discipline will fragment into isolated communities of practice that fail to learn from one another and replicate work unnecessarily. One of the few exceptions to this has been the work of Robert Kosara [15], [28] who has combined an understanding of the major strands within information visualisation with a critical theoretical awareness. In this article we contribute to this process of developing theoretical foundations for information visualisation by employing two existing tools.

The first is semiotics, the study of signs. As devised by Saussure and Peirce, semiotics has developed into a powerful theoretical framework for understanding the relationships between signs, sign systems, the consumers of those signs, and the systems they represent (though this is a point of contention discussed in section 2). Information visualisations are signs par excellence, and thus seem obvious candidates for semiotic analysis. In section 2 we provide an extremely brief introduction to the frameworks used by semiotics to analyse sign systems, and discuss how this applies to information visualisation. The result is to show how information visualisation can be understood using a series of relationships, or mappings, from one domain to another, summarised in a semiotic triad.

The second tool is category theory, the mathematical study of systems of structures and their mappings in terms of their formal relationships. This theory is introduced in section 3, and applied to information visualisation in section 4. The most powerful weapon in category theory is the notion of commutativity, which forces one to try to extend and construct structures in such a way as to reach algebraic closure by considering the consequences and implications of a structural description of a system. By applying the criterion of commutativity to our semiotic triad in sections 4.2 and 4.3 we extend our

- *P. Vickers, J. Faith and N. Rossiter are with the School of Computing, Engineering, and Information Sciences, Northumbria University, Newcastle-upon-Tyne, NE2 1XE, United Kingdom.*
*E-mail: paul.vickers@northumbria.ac.uk*
- *Joe Faith carried out this work whilst at Northumbria University but is now with Google Corp. Inc., Mountain View, CA 94043, USA.*

*Manuscript received ; revised .*

1. See http://www.processing.org

framework in a natural way to cover (or uncover) some aspects which are already familiar to practitioners of information visualisation, and some other aspects which are less obvious. The end result, or closure, is the general description of information visualisation given in section 4.4.

## 2 Semiotics and Information Visualisation

Semiotics is the study of the creation and interpretation of signs. Signs are words, images, sounds, smells, objects, etc. that have no intrinsic meaning and which become signs when we attribute meaning to them [4]. Signs stand for or represent something beyond themselves. Modern semiotics is based upon the work of two principal thinkers, the Swiss linguist Ferdinand de Saussure and the American philosopher and logician Charles Sanders Peirce. In Saussure's semiology the sign is a dyadic relationship between the *signifier* and the *signified*.[2] In Saussure's linguistic system a sign is a link between the signified (a concept) and the signifier (a sound pattern) both of which are psychological constructs having non-material form rather than material substance [4]. For example, /tree/ is a signifier for the concept of the thing we know as a tree. The sign thus formed is a link between the sound pattern and the concept. However, modern applications admit material form for the signifier (e.g., road signs, printed words, etc.). Saussure's scheme explicitly excludes reference to objects existing in the real world. The signified is not directly associated with an object but with a mental concept.

Peirce's semiotics is based upon a triadic relationship comprising:

- The *object*: the thing to be represented (note, this need not have a material form);
- The *representamen*: the form the sign takes (word, image, sound, etc.,) and which represents the object;
- The *interpretant*: the sense we make of the sign.

So, Peirce admits the referent that Saussure brackets. Fig. 4 shows two Peircean triads drawn as 'meaning triangles'.[3] It should be noted that the Saussurean signifier and signified correspond only approximately to Peirce's representamen and interpretant; unlike Saussure's signified, Peirce's interpretant itself becomes a sign vehicle in the mind of the interpreter. Fig. 1(a) shows the basic structure of a Peircean sign and Fig. 1(b) shows the sign formed by the name Agamemnon which represents a specific individual cat with that name.

To relate this to visualisation consider Fig. 2 which shows a semiotic relationship that exists between a set of student marks and an external representation. Using a spreadsheet program we take a data set which has been collected from the real world system of a cohort of students studying a course. These data are then presented to the user via a chart representation. It should be noted that this visual representation is not the data set, but a particular representation of it. So, we have a sign (in the Peircean sense) in which the data set is the referent object, the chart view serves as the representamen of the data set, and the interpretant is the sense we make of the student marks by looking at the chart.

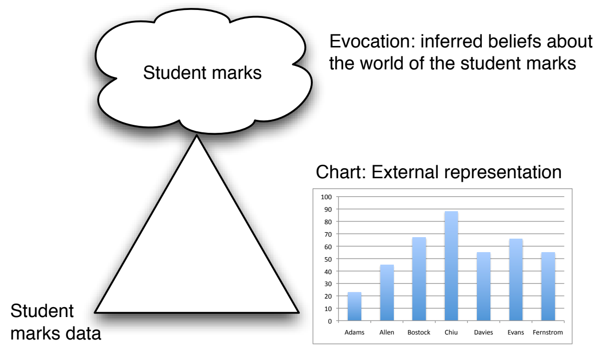


Fig. 2. A chart is an external representation of the underlying data. The interpretant is the concept that is formed in our mind when we view the chart representation.

It is important to note that, contrary to Saussure's original structuralist view, sign systems exist within a social and cultural context which, the post-structuralists would argue, needs to be taken into account. Peirce's semiotics, through the notion of a ground, admits context. It is important because visualisation requires the producer (the addresser in semiotic terminology) and the consumer (the addressee) to share some contextual knowledge in order for successful meaning making to take place. In Peirce's semiotics meaning is mediated such that the "meaning of a sign is not contained within it, but arises in its interpretation" [4]. Hjelmslev [13] recognised that no sign can properly be interpreted without first contextualising it so that in addition to a sign's denotative meaning its context also lends it connotative meaning. For instance, in the example of a student marks system the ground would include knowledge about what constitutes a pass mark in this assessment scheme, where the grade boundaries lie, and so on.

Note that this semiotic framework would exclude some examples of what is popularly regarded as 'information visualisation'. For instance, take Radiohead's "House of Cards" video [21]. The video was shot without any cameras being derived solely from data obtained from 3D images produced by Geometric Informatics for close proximity objects and Velodyne Lidar for landscapes. The data sets used to make the video are available for anyone to download from the project's web site where we are encouraged "to create your own visualizations" [21]. The problem with this concept is

2. Saussure used the term *semiology*. This word has become associated with the European school whilst *semiotics* is more associated with the American school of Peirce et al. For the sake of simplicity we will use semiotics throughout this article.

3. In 1923 Ogden and Richards [18] gave this visual interpretation of Peirce's triad since when it has become the conventional representation (see also Eco [7] and Sowa [23]).

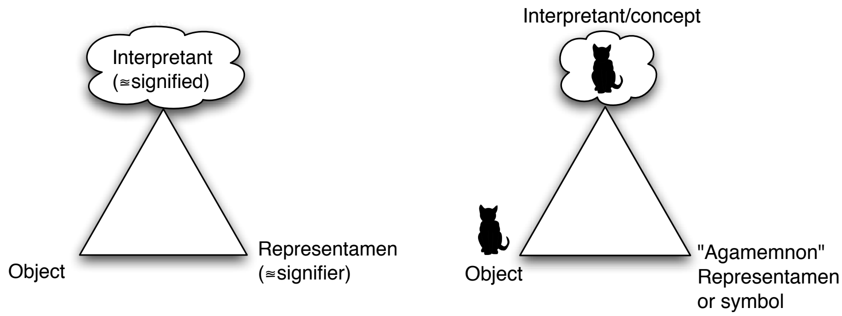


(a) A Peircean semiotic triad (after Sowa [23]). Approximations to Saussure's semiotic terminology given in parentheses.

(b) Example: A real individual cat is the object. It is signified by the symbol "Agamemnon" which brings into our mind the concept of Agamemnon the cat.

Fig. 1. Two semiotic 'meaning triangles'

that the video, whilst unarguably data-driven, ought not be considered a visualisation as it provides no insight into the data, it is pure spectacle [1]. As Card, Mackinlay, and Shneiderman put it: "The purpose of visualization is insight, not pictures" [3].[4] Visualisation, then, is a process that begins with the real world, or more narrowly, a system in the real world about which we are interested. The system could be a mechanical system or it could be something like a cohort of students on a degree course. From the system we gather data which are then mapped via some transformation rules to an external representation (a graph or chart, an interactive 3D model, statistical box plots, etc.) and this representation is then 'read' by the person who wants to gain insight into the system.

This reading of the representation evokes concepts and ideas in the mind and inferences are drawn leading to understanding of the system and, as we shall see later, knowledge of the truth as it pertains to that system. This process is encapsulated in Fig. 3.

### 2.1 Data

At the heart of the visualisation process lies the real world with its events, entities, concepts, etc. Contrary to what some current advertisers would have us believe, the real world is not data and clouds of data do not follow us around in any ontological sense.[5] Rather, the data that are used in computer systems, visualisations, and so on have been generated and collected by *instruments* with a specific *purpose* in mind. For example, census data are collected using a specially designed instrument (the census questionnaire) and the data gathered allow governments to plan spending priorities and the future provision of services. Engineers collect performance data from machines; banks create transaction data to enable them to keep bank accounts accurate and up-to-date. The main point here is that data have no *a priori* existence, they have to be created by collection instruments and stored as a data set in an appropriate storage medium (account ledgers, computer hard drives, optical discs, etc.).[6] This means that data *always* have a context (a *ground*), a scenario in which they were intended for use. These data are abstractions of the real world from which they were collected.

## 3 Category Theory

In this section we introduce the main concepts of category theory and briefly examine one initial attempt to apply it to information visualisation. The Stanford Encyclopedia of Philosophy [17] describes the potential of the theory:

> Category theory has come to occupy a central position in contemporary mathematics and theoretical computer science, and is also applied to mathematical physics. Roughly, it is a general mathematical theory of structures and of systems of structures. As category theory is still evolving, its functions are correspondingly developing, expanding and multiplying. At minimum, it is a powerful language, or conceptual framework, allowing us to see the universal components of a family of structures of a given kind, and how structures of different kinds are interrelated. Category theory is both an interesting object of philosophical study, and a potentially powerful formal tool for philosophical investigations of concepts such as space, system, and even truth.

Category Theory (CT) is a branch of mathematics developed to analyse systems of structures, and mappings between those structures, in their most general form. This level of generality means that category theory is

4. This was itself a recasting of R. W. Hamming's 1962 insistence that "the purpose of computation is insight, not numbers" [12].

5. See Mitchell Whitelaw's blog at http://teemingvoid.blogspot.com/2010/05/this-is-data-arguing-with-data-baby.html?showComment=1282871974120. Here is the 'This is data' example he mentions (the link is broken on his blog) — http://www.youtube.com/watch?v=8Ctv-5n9b6U.

6. By this we mean they have no *a priori* material existence. Whether or not they exist in some ideal Platonic sense is outside the scope of this discussion.

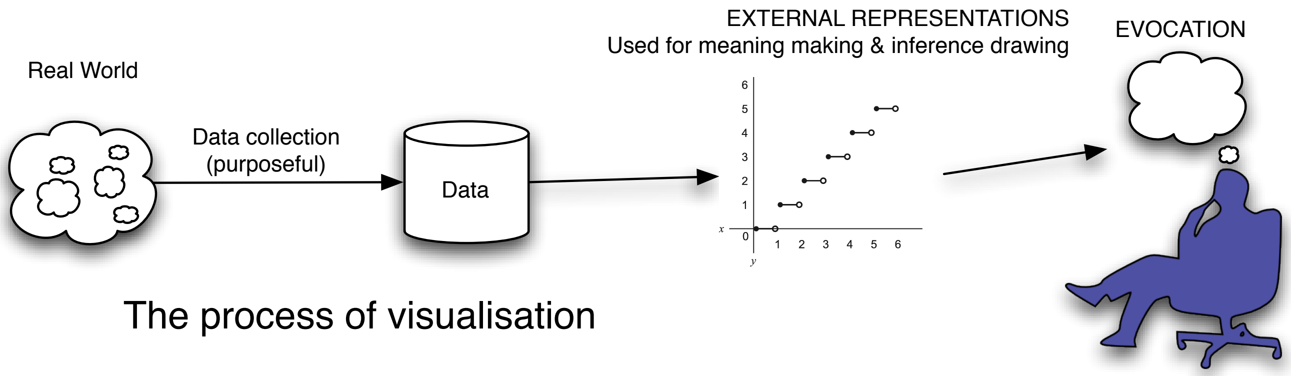


Fig. 3. Understanding the real world through visualisation: Data are collected with purpose from the real world and mappings are used to produce external representations. These representations are then used for meaning making and drawing inferences about the data. The visualisation process encompasses cognition in the mind of the observer.

capable of demonstrating similarities between disparate fields of mathematical enquiry — from set theory to theoretical computer science — in such a way that allows insights from one to be translated to another. It can thus act as a Grand Unified Theory in mathematics. The corresponding weakness, and a common complaint about category theory, is that it can appear as just a theoretical or descriptive superstructure atop the actual domain of study, in which the insight comes from understanding of the domain base rather than the theoretical superstructure. However, as described below, through building upon the underlying concept of an arrow or morphism any construction can be developed in principle to any level of detail. In this article we argue that the 'cash value' of applying CT concepts to the process of information visualisation is that:

- By considering intuitive concepts (such as 'visualisation') from a formal point of view it can enforce conceptual hygiene by exposing unclear definitions.
- By understanding the universal properties of the resulting structure (rather than our day-to-day experience of these kinds of phenomena as practitioners) then it can force us to step back and consider aspects or possibilities that might otherwise be neglected.

An introduction to the use of category theory in practice can be found in Mac Lane's text for the working mathematician [16]. Category theory is built from just two classes of entity: objects and morphisms (we follow the convention that Objects are Capitalised and *morphisms* are *italicised*). Almost anything can be considered as an Object: physical objects, abstract objects, or entire systems. Indeed, a large part of the power of category theory comes from its recursive ability to treat ever more complicated systems as building blocks in the next level of abstraction. All that is required for an entity to qualify as an Object is that it can be individuated, that is, we have some method for determining whether two objects are identical. This method is represented as a morphism from the object to itself, known as the identity morphism.

Morphisms are mappings between objects. They are represented diagrammatically by, and often described as, arrows. (Indeed, in some notations objects are represented by their identity morphisms, meaning that the arrow can be used to represent both objects and transformations.) All that is required for a mapping to qualify as a morphism is that there is a unique target for each domain object: i.e., one object at the base of the arrow and one at the head. Ontologically, a morphism is understood as a generalisation of a mathematical function, in other words as an association between its source object and the target. Morphisms may represent physical, causal, or temporal processes, or purely formal relationships.

Objects and morphisms can be combined into diagrams, the simplest of which, and the most basic tool in category theory, is a triangle, as in Fig. 4.

$$\begin{tikzcd} \mathrm{A} \arrow[r, "f"] \arrow[dr] & \mathrm{B} \arrow[d, "g"] \\ & \mathrm{C} \end{tikzcd}$$

(a) The basic triangle

$$\begin{tikzcd} \mathrm{A} \arrow[r, "f"] \arrow[dr, "g \circ f"'] & \mathrm{B} \arrow[d, "g"] \\ & \mathrm{C} \end{tikzcd}$$

(b) A triangle that commutes

Fig. 4. Category theory triangles

Given objects A, B, and C and morphisms $f : \mathrm{A} \to \mathrm{B}$ and $g : \mathrm{B} \to \mathrm{C}$, the first question a category theorist would ask, looking at Fig. 4(a), is whether there is another morphism, $h : \mathrm{A} \to \mathrm{C}$, that completes the triangle. This morphism is described as the composition of $f$ and $g$, or $g \circ f$. Note the ordering: compositions are read right to left as the morphism $g$ is applied to the result of $f$. If there is such a morphism $h$ then the triangle is said to commute, the objects and morphisms together form a category and we can write the equation

$h = g \circ f$.[7]

Consider an informal example, that of the category of familial relations. Suppose A, B, and C are persons, $f$ is the mapping of 'motherhood', and g is 'sisterhood'; then $g \circ f$ corresponds to 'aunthood', the diagram commutes, and we have a category. But now suppose that $g$ is 'friendship'. We do not have a well defined mapping for $g \circ f$ in this case (other than just 'the relationship I have with a friend of my mother', a circular definition in which the definiendum invokes the definiens). Hence we cannot form a commuting triangle or create a category. In order to form a category that encompasses both kinship and friendship then we must enrich our vocabulary of morphisms to describe those relationships. Societies that are based on communal kin-groups rather than atomic families will tend to develop richer vocabularies to capture these composite relationships; hence their set of morphisms will tend to form well-defined categories. This is an example of how category theory can be used in practice: no one would suggest that kin-groups form vocabularies because of their category-theoretic properties, but the category-theoretic perspective suggests a set of questions that could be asked of that vocabulary, and a conceptual framework for analysing it.

Objects in a category must have an identity morphism. This means that for an object A there is an identity morphism $1_{\mathrm{A}} : \mathrm{A} \to \mathrm{A}$ (also written as $id_{\mathrm{A}}$) such that for every $f : \mathrm{A} \to \mathrm{B}$ we can state $1_{\mathrm{B}} \circ f = f = f \circ 1_{\mathrm{A}}$. We can represent this diagramatically thus:

$$\begin{tikzcd} \mathrm{A} \arrow[loop left, "1_{\mathrm{A}}"] \arrow[r, "f"] & B \arrow[loop right, "1_{\mathrm{B}}"] \end{tikzcd}$$

The final requirement for a category to be valid is that where three or more morphisms (e.g., $f$, $g$, $h$) are composed together, the overall operation is associative. That is, the order in which the evaluation is made is immaterial: $(f \circ g) \circ h = f \circ (g \circ h)$. This is trivially true in the case of familial relations.

### 3.1 Types of Morphism

Morphisms are of several types, viz monomorphic, epimorphic, endomorphic, isomorphic, and automorphic. A morphism can be monomorphic (monic) or epimorphic (epic) depending on whether it is left- or right-cancellative, that is, how it exposes differences in morphisms with which it is composed. For example, consider the category **F** of female relations with objects $X$. The relation $r : \mathrm{X} \longrightarrow \mathrm{X}$ where X is a female and $r$ is motherhood is not monic (that is *into*, 1:1, injective in sets, see Fig. 5(a)) as there are potentially many paths from each female to the female mother, one for each child, so the left component is therefore not cancellative.A relation that would be monic is *first-born* as this is 1:1 and the left component is cancellative. The relation *motherhood* may though be epic (that is *onto*, surjective in sets, see Fig. 5(b)) if all females in the category **F** are mothers, that is, the right component is cancellative. The tentative results here illustrate how sensitive the properties of monomorphism and epimorphism are to the definitions employed for the underlying types.

In sets a morphism that is both injective and surjective is said to be bijective (see Fig. 5(c)). Bijection may also be used as a term in category theory in the same way as in set theory but a term more often used is isomorphism, which is a stronger concept as it additionally involves identities. For example, for the relation $s : \mathrm{C} \longrightarrow \mathrm{D}$ (C is an object containing females, $s$ is marriage, D is male), where there exists an inverse $s^{-1} : \mathrm{D} \longrightarrow \mathrm{C}$ then if both $s^{-1} \circ s = 1_C$ and $s \circ s^{-1} = 1_D$, $s$ is said to be isomorphic to $s^{-1}$. A morphism that is both monic and epic is said to be isomorphic if both identities hold. Note that these relations will only be isomorphic if C and D respectively are the collections of females and males actually married at any given point in time and the laws governing marriage permit only one partner each.[8]

The mapping $r$ described earlier is an endomorphism in that it maps an object X to itself in a recursive manner. If an inverse $r^{-1}$ can be defined such that $r$ is isomorphic to $r^{-1}$ then the relationship between $r$ and $r^{-1}$ is said to be an automorphism.

Categories themselves can be considered as objects in higher-level categories. For example, consider two sets of individuals, each comprising a mother, daughter and aunt. There is a natural morphism between the individuals: from daughter to daughter, from mother to mother, and from aunt to aunt. But this morphism can also be defined between the morphisms: from the motherhood relationship in the first set to that in the second, etc. A morphism between categories that preserves structure in this way is known as a functor, and most practical applications of category theory come from studying the structure preserving properties of such mappings. We could, for example, study possible functors between languages, including how they describe familial relationships, revealing which languages are structurally similar and in what ways.

### 3.2 'Structuralist' Algebraic Semiotics

Both Peirce and Saussure understand signs in terms of relationships and mappings between signs and sign systems; thus they seem natural candidates for the category theory treatment. The first effort to apply category theory to semiotics was that of Goguen and Harrel [11], in which they attempted a formalist treatment of the semiotics concerned with information visualisation and user interface design, an attempt that they described as 'algebraic semiotics'.

The objects in Goguen's algebraic semiotics are sign systems. The actual definition of a sign system in

7. A characteristic activity of category theorists, and the basis of many of the resulting proofs, is known as 'diagram chasing', that is, tracing the arrows of morphisms around such diagrams, checking for commutativity.

8. $s^{-1}$, the right inverse of $s$, is sometimes called a section of $s$. $s$, the left inverse of $s^{-1}$, is sometimes called a retraction of $s^{-1}$.

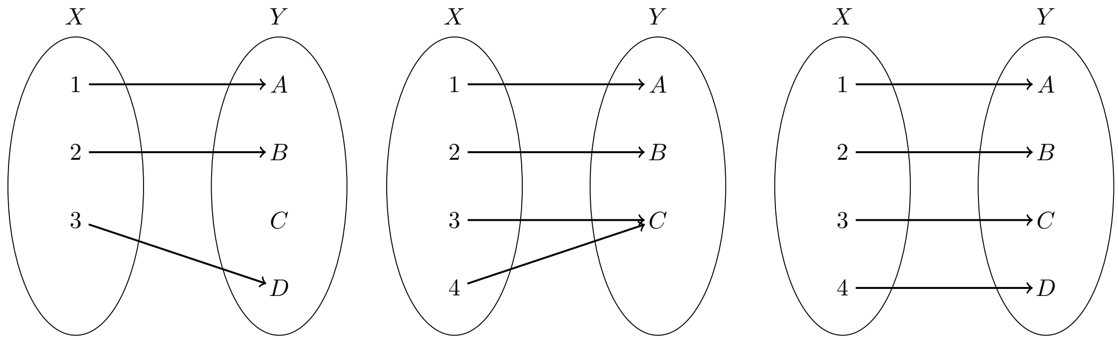


(a) Injection (into, or 1:1) (b) Surjection (onto) (c) Bijection (1:1 onto)

Fig. 5. Types of mapping. In an injection each element of the co-domain is mapped to by *at most* one element of the domain. In a surjection each element of the co-domain is mapped to by *at least* one element in the domain. In a bijection each element of the co-domain is a mapping of *only one* element in the domain

Goguen's sense is complex, involving much theoretical apparatus taken from mathematical algebra but, for our purposes, the key elements are signs, constructors, and axioms. The signs form the vocabulary, or set of all possible signs. The constructors provide a systematic way of generating those signs. And the axioms constrain those signs. For example, consider Goguen's example of a simple time of day sign system which just shows the number of minutes past midnight. The signs are the set of natural numbers generated using two constructors: the constant 0 (representing midnight) and a successor operation, $s$, where for a time $t$, $s(t)$ is the next minute. The only axiom, which constrains the set of generated signs to a 24 hour day, is that $s(1439) = 0$. Goguen thus defines signs in terms of generalised structures, from which the validity of any particular sign is derived. In the vocabulary of first-order logic, the sign system is a theory of which any particular sign is a model.

Having defined sign systems, Goguen then considers mappings between them, which he describes as semiotic morphisms. For example, if we had a very slow, regular, sand glass containing 1440 grains of sand, in which precisely one grain of sand fell every minute, and which is turned when the last grain falls, then we can define a mapping to this from the time of day sign system defined above, in which the elements of the former (constant 0, the $s$ constructor, the set of numbers, the $s(1439) = 0$ axiom) are mapped respectively to elements of the latter (empty lower glass, the falling of a grain of sand, the possible piles of sand, turning the glass).

Goguen's is a strongly structuralist theory, in two senses. The first is that signs are defined as such in virtue of their membership of, and role within, a sign system. It is the structure of the sign system that defines its constituents as signs. The second is that the only relationships considered — the semiotic morphisms — are between sign systems rather than between sign systems and either external or mental states. It all happens within the 'third order' [6]. In particular, there is no distinction in this framework between a set of data and a visualisation of that data. Data is just another sign system.

### 3.3 'Post Structuralist' Algebraic Semiotics

We find Goguen's structuralism problematic. For example, how can this framework be used to discuss the quality of a visualisation? Goguen attempts an answer to this 'narrowly' by characterising how well his semiotic morphisms preserve the structure of a sign system.We argue that the solution is a richer characterisation of the visualisation process, to expand the use of category theory to explicitly incorporate other elements of the visualisation process, including the context of the visualisation, and how the visualisation is used in practice. The end result of this process can be seen in the commutative diagram in Fig. 13 but, as is usual with category theory, we proceed by constructing the diagram in stages, checking for commutativity at each step.

## 4 Category Theory Applied to Information Visualisation

Peircean semiotics is based on a triadic relationship between object, representamen, and interpretant. We can draw our semiotic triad as the commutative diagram in Fig. 6 (where O, R, and I stand for object, representamen, and interpretant respectively):

$$\begin{tikzcd} \mathrm{O} \arrow[r, "f"] \arrow[dr, "g \circ f"'] & \mathrm{R} \arrow[d, "g"] \\ & \mathrm{I} \end{tikzcd}$$

Fig. 6. The Peircean semiotic triad as a commutative diagram

We say that the object in visualisation is the data that has been collected from a given system, the representamen is the representation, and the interpretant is the mental state evoked by the representation in the mind of the interpreter. So, we will call the objects of our commutative diagram Data, Representation, and Evocation respectively. The morphisms between them we define as follows:

- The transformation from Data to Representation is a morphism called $render$ because the data are rendered in a given way so as to represent the Data whilst maintaining structure and content (this is a partial ordering).[9]
- The morphism between Representation and Evocation is called $read$ because the interpreter reads the Representation.
- To the composition $read \circ render$ we assign the name $understanding$ for a proper Representation will lead to the reader understanding some aspect of the Data.

Thus, Fig. 7 shows us the Peircean semiotic triad presented as a commutative diagram which forms the core of the visualisation process.

$$\begin{tikzcd}
\text{Data} \arrow[r, "render"] \arrow[dr, "understanding"'] & \text{Representation} \arrow[d, "read"] \\
& \text{Evocation}
\end{tikzcd}$$

Fig. 7. Visualisation as a semiotic triad

However, in the particular case of information visualisation the object can be further decomposed into a set of data and a system which that data measures in some way. So, the starting point for the visualisation process is not data but the system from which the data have been gathered. Furthermore, the interpretation of the representation leads not only to understanding the data but also to beliefs and inferences about the system, the truth of which can be tested. In other words, if we combine Figs 3 and 7 we get Fig. 8. So, a prototypical visualisation process consists of the following entities and processes, which we will describe using the category theoretic terms of objects and morphisms:

- **System**: a real world system, object, or phenomenon, such as a class of students.
- **Data**: a set of data that describes some aspect of that System, produced by a *measure*, such as test scores for those students.
- **Representation**: some visual, aural, haptic, or literal artefact of that Data, produced by a process of *render*ing, such as a bar chart of their performance.
- **Evocation**: what that Representation evokes in the mind of the user, produced by the user's reading of that Representation, such as the teacher's viewing of that bar chart.
- And that Evocation is thus an understanding of the original data ($understanding = read \circ render$).

Fig. 8 shows this expansion, though it should be noted that this diagram is incomplete and is to be read only as a stepping stone on the way to Fig. 13.

9. It is partial because not all structure and content is necessarily carried over.

$$\begin{tikzcd}
\text{System} \arrow[r, "measure"] & \text{Data} \arrow[r, "render"] \arrow[dr, "understanding"'] & \text{Representation} \arrow[d, "read"] \\
& & \text{Evocation}
\end{tikzcd}$$

Fig. 8. System, Data, Representation, Evocation and the morphisms between them

### 4.1 The Visualisation Process is a Category

In order for these objects and morphisms to form a category certain conditions must apply.

1) **Object Identity**: each of {System, Data, Representation, Visualisation} must have an identity operation defined, and the kernel of this identity function will, in turn, identify an equivalence class of objects that are considered identical under this mapping.[10] In particular this requires that we are able to decide unequivocally if two instances of the same System, Data, Representation, and Visualisation are identical. This simple requirement forces a great deal of conceptual hygiene. Unless we are able to answer the following questions then we are vulnerable to the accusation that we do not have a well-defined visualisation process.
   a) $1_{\text{System}}$: The problem of identifying systems is a common and urgent one in most empirical science, especially biology and medicine. How else can we talk about replicating results unless we are doing the same things to the same systems? Unlike a data set which captures a snapshot of a system and is static, a System may experience change over time. For example, in a System of a class of students the individual students age and mature, change their clothes and their hairstyles, and some may even drop out of the course. How then do we know when two System objects have the same denotation? Whitehead talked in terms of *continuants* and *occurrents*. A continuant is an object that persists over time whilst an occurrent is an event or a process, something that does not persist. Whitehead [27] argued that whether we consider something to be a continuant that retains its identity over time or an occurrent in a constant state of flux depends, as Sowa puts it "more on the viewer than on the thing itself" [23]. Therefore, agreed criteria are needed for establishing System identity which will allow a System to experience change over time whilst still being considered to be the same System.

10. In category theory it is strictly not possible to show that things are the same or identical. The strongest statement possible is that two sets are naturally isomorphic (unique up to natural isomorphism), that is, indistinguishable.

b) $1_{\text{Data}}$: When do we say that two sets of data, such as test scores, are the same? Do the absolute scores matter or is it the same set of scores when expressed as a percentage? What degree of precision is required? If we had a very large set of scores (for example if we are choosing to understand the changes in national exam performance) then identity might be defined in cases where there is no statistically significant difference between samples or aggregate distributions.

c) $1_{\text{Representation}}$: When are two representations the same? Is a printed version the same representation as an on-screen version? Do the rendering resolution or particularities of the hardware matter?

d) $1_{\text{Evocation}}$: What does it mean to say that two users form the same mental picture of the data? The obvious problems in determining internal psychological states mean that this issue is usually operationalised in terms of the ability to answer questions about the data ('which student performed best?' 'Has average performance declined or improved'), where the same answer implies the same understanding.

2) **Morphisms are Maps**: the target object of each morphism is determined by the source object. In the case of *representation* this requires that a single set of data (to within $1_{\text{Data}}$) generates a single Representation. In particular, where the representational tool can be manipulated, steered, or interacted with then it is the tool, bound to that Data, that is considered to be the Representation rather than any particular state or view that it produces.
3) **Commutativity of Morphisms**: Given a triangle of objects and morphisms, such as that formed by (Data, Representation, Visualisation, *rendering*, *reading*, and *understanding*) the first question posed by category theory is whether the diagram commutes, i.e., $read \circ render = understanding$. That is, whether the result of *read*ing the Representation produced by *render*ing the Data is an *understanding* of the Data. If the Data is *render*ed and then *read*, but as a result the reader does not understand the data then the process of visualisation has failed.
4) **Morphism Associativity**: where three or more morphisms are composed together, the order of evaluation is immaterial. This is a trivial constraint in this context.

Visualisation processes for which conditions 1–4 are satisfied can be considered as valid categories, satisfying the axioms of category theory.

### 4.2 Category-Theoretic Properties of the *representation* Morphism

Once we have constructed a category corresponding to our visualisation process, then we can start to use the concepts of CT to consider its properties. First we will consider what it means, in visualisation terms, for the morphism *representation* to have each of the following category-theoretic properties. In each case we take a standard definition from category theory and apply it to the Visualisation Process category (or rather, that subset of it shown in Fig. 8). In each case we find that the formally defined property yields a property or issue that is important when considering visualisation.

#### 4.2.1 Monomorphism Corresponds to Sensitivity

A *representation* morphism is monic iff for all $measure_1, measure_2 : \text{System} \to \text{Data} : representation \circ measure_1 = representation \circ measure_2 \implies measure_1 = measure_2$.

In other words, if the same System is measured in two different ways then the resulting Representations will necessarily be different. For example, suppose we are to measure the performance of our students using two different tests ($measure_1$ and $measure_2$), and represent them in two different ways: a simple textual description (for example "John was the best student"), and a bar chart showing the relative performance of each. The differences in the test would not make any difference to the textual description but it would to the bar chart (assuming that identity morphisms on the Representations and Data are well-defined). We would normally describe this in terms of the sensitivity of the visualisation. Sensitivity is normally assumed a desirable property of a representation — and this assumption may often be valid — but the point here is to show that this important property corresponds to a category-theoretic property of the Visualisation Process.

#### 4.2.2 Epimorphism Corresponds to Non-Redundancy

A *representation* morphism is epic iff for all $read_1, read_2 : \text{Representation} \to \text{Evocation} : read_1 \circ representation = read_2 \circ representation \implies read_1 = read_2$.

In other words, if two individuals *read* the same Representation in different ways then they will reach different *understandings* of the Data. Although this may seem tautological, there are important cases when it is not true. Consider a set of data in which three attributes are measured for each sample (for example, student performance on three different tests). This data could be represented using a conventional scatter plot in which the $x$-coordinate corresponds to test 1, $y$ to test 2, and *both the size and shade of each point correspond to test 3* (see Fig. 9). One individual may notice the position and size of each point, whereas another may notice the shade. They would draw identical conclusions about the data, but they have read the representation in different ways. The representation in this case is redundant in the sense that the same information is represented in two different ways; hence there is more than one way of gaining an understanding of that information from a single representation. Redundancy is usually considered to be an undesirable property of a Representation, but, again, that is not the issue here.

| height | weight | strength |
|---|---|---|
| 34 | 18 | 56 |
| 26 | 21 | 77 |
| ... | ... | ... |

(a) The data set

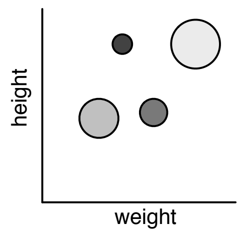


(b) A redundant representation in which the *strength* dimension is mapped to both shade and size

Fig. 9. Redundancy in representations

#### 4.2.3 Endomorphism Corresponds to Literalness

A *representation* morphism is endomorphic iff $Data = Representation$. Now, there are issues about the metaphysical status of logical Data compared to that of a physical representation, and hence whether such an identity is ontologically valid; but if we bracket those issues then the concept that the category theoretic property of endomorphism encapsulates is that of *literalness*. Contrast this with the notion that visualisations are always metaphorical, as Cox argues [5]. Cox suggests that there is a direct relationship between visualisation and the mapping process (cognitive and creative) in metaphor theory. She says:

> Linguistic and visual metaphors are defined as mappings from one domain of information (the source) into another domain (the target). Likewise, data-viz maps numbers into pictures, resulting in visaphors, digital visual metaphors [5].

According to Cox then, visualisation may be explained in terms of a mapping process in which some features of the source domain are mapped onto certain features of the target domain (n.b., these mappings do not produce a one-to-one correspondence: some characteristics may be mapped, others not). However, some Representations of a data set consist of simply presenting the data set, for example, in a spreadsheet or printed table. These seem to be valid representations, but they lack the metaphorical nature posited by Cox.

#### 4.2.4 Isomorphism Corresponds to Non-ambiguity

A *representation* morphism is an isomorphism iff there is an inverse morphism (which we will call *decode*) such that $decode \circ representation = 1_{Data}$ and $representation \circ decode = 1_{Representation}$ Thus the *decode* morphism allows us to recover the original Data from a Representation generated from it (to within the degree of accuracy determined by $1_{Data}$). Representations that are not isomorphic are thus ambiguous in the sense that two different data sets may generate the same representation. In some situations this is a valuable property, for example in representations that aggregate or filter large or complex data sets into simpler forms. In other situations it is less desirable.[11]

### 4.3 The Intension of the Visualisation Process Category

Fig. 8 represents the process of visualisation as a category; however it is important to note that this represents a single concrete instance of this process. It could represent, for example, the production and consumption of a single particular scatter plot from a single particular set of data. We now need to generalise this notion to describe, for example, the properties of scatter plots in general. In the case of the Data we have the familiar notion of a Schema, which refers to the structure of data as opposed to its particular values. For example, a relational database of tables and attributes is defined using a database schema, which is then filled with values during its lifetime of use. The Schema is the intension (the constant conditions that capture the set of all possible values for that data set) while the Data is the extension (the variable set of actual values).[12] We can also generalise the notion of a particular Representation to that of a Layout, that is, the way in which Data belonging to a Schema are represented. For example, the two scatter plots in Fig. 11(a) are different Representations, but share a Layout. The notion of Layout is familiar from data visualisation tools (such as the charting feature in spreadsheet programs) that allow the user to choose which of many options are used to graphically represent a selection of data.

We then have a set of rules that allow us to map from a data Schema to a representational Layout, which is a generalisation of a particular Representation. Fig. 10 shows the Data, Representation, and Evocation objects from our visualisation category at the extensional level and Schema and Layout as the generalisations, or intensions, of the Data and Representation objects respectively.

#### 4.3.1 Layouts Require Schemata

Not all Data has a corresponding Schema. Data that does not is generally known as unstructured data. Examples

11. A representation morphism for which the first condition holds is a *section*. A representation for which the second holds is a *retraction*. There are categories and morphisms for which only one of these conditions holds, but it does not seem possible in the case of a visualisation process category.

12. In philosophy a distinction is drawn between a term's intrinsic meaning or its *intension* and its denotation, or *extension*. Frege gave the example of the *morning star* and *evening star*. The two terms have different meanings (intensions) but both have the same denotation (extension), the planet Venus [23].

$$
\begin{tikzcd}
\mathit{Intension} & \text{Schema} \arrow[r, "rules"] & \text{Layout} & \\
\mathit{Extension} & \text{Data} \arrow[u, "generalisation_S"] \arrow[r, "render"'] & \text{Representation} \arrow[u, "generalisation_L"'] \arrow[r, "read"'] & \text{Evocation}
\end{tikzcd}
$$

Fig. 10. Introducing the intensional level of the category

include natural language text which can certainly be Represented using, for example, pictorial illustrations. However there can be no rules governing the Layout of this Representation: there may be a set of Representations of a similar style — for example several illustrations in a single book — but each is individually inspired by the text it is designed to Represent. Unless the Data being represented can be generalised into a Schema, then there can be no corresponding generalisation of the Representation into a Layout. This criterion also helps to distinguish between visualisations that are data- or data-driven art.

#### 4.3.2 *Non-surjectivity of Layout Corresponds to Chart Junk*

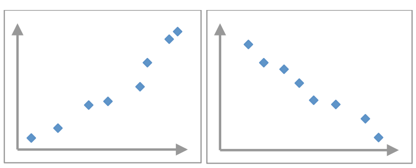

(a) Plain charts

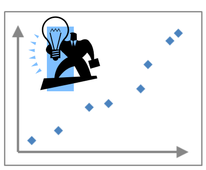

(b) Arbitrary chart junk

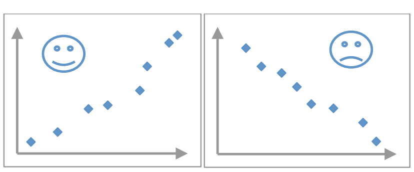

(c) Schema-derived chart junk

Fig. 11. Chart Junk

In the scatter plots in Fig. 11(a) every element of the Layout was derived from some element of the Schema. If we write these objects as sets, then we can say that the morphism between them is a surjection:

$$
\begin{aligned}
\text{Schema} &= \{Instances, Attributes\} \\
\text{Layout} &= \{Points, Axes\} \\
rules(Instances) &= Points \\
rules(Attributes) &= Axes
\end{aligned}
$$

However, this is not always the case. For example, suppose we decorate one of our scatter plots with a figure as in Fig. 11(b). What is the Layout in this case? If the Layout is a *generalisation* of this particular Representation then it will include the decoration. But this decoration is not derived using a *rule* from any part of the Schema — it is an arbitrary addition (perhaps inspired from some property of the System not captured in the Data). It is 'chart junk'.

The following aspects of this definition of chart junk should be noted.

1) This use of the term is not intended to be pejorative: chart junk can be useful in aiding understanding of the system [2]. The purpose of this category-theoretic definition is to highlight the difference between decorative elements in Layouts that communicate Data (those that are derived from a Schema), and those that do not.
2) It is a much narrower use of the term than Tufte's original definition [24] which defined chart junk as all unnecessary, redundant, or non-data ink. Consider the example in Fig. 11(c) in which scatter plots are decorated with faces indicating the movement of the data. The decoration is redundant chart junk in Tufte's sense, but not in ours since it communicates something about the data. There is a rule for deriving this element of the Layout from the Schema which thus ensures that Fig. 10 commutes.
3) There is a difference between redundancy at the extensional level of the Data and Representation introduced in Fig. 10 and arbitrary chart junk at the intensional level of the Schema and Layout. A Representation may include some redundancy even though there is no chart junk. In the example of the scatter plot in which a single Data attribute is represented using two retinal attributes such as in Fig. 9(b) — or, indeed, the example in Fig. 11(c) of a scatter plot decorated with a face representing the polarity of the correlation — we have redundancy in the Representation, but the Layout *rules* are not surjective: every element of the Layout is the product of some aspect of the Schema. Redundancy at the level of Data and Representation may not be arbitrary at the level of Schema and Layout. Conversely, chart junk is not (necessarily) redundant, it is arbitrary.

### 4.4 Closure of the Information Visualisation Category

Combining our partial diagrams (Figs. 8 and 10) we get the diagram in Fig. 12.

From a category-theoretic point of view, that is, by considering its formal structural properties, this diagram is incomplete; it is not closed in the mathematical sense. Seeking closure suggests the following questions:

$$\begin{tikzcd}
& \text{Schema} \arrow[r, "rules"] & \text{Layout} \\
\text{System} \arrow[r, "measure"] & \text{Data} \arrow[u, "generalisation_S"] \arrow[r, "render"] \arrow[dr, "understanding"'] & \text{Representation} \arrow[u, "generalisation_L"'] \arrow[d, "read"] \\
& & \text{Evocation}
\end{tikzcd}$$

Fig. 12. Intermediate, mathematically un-closed visualisation category

1) Is there an equivalent of Evocation at the intensional level? If so, does this form a category with Schema and Layout, so that we can form a functor from the extension of the visualisation process to its intension?
2) Are there initial and terminal objects? That is, can the diagram be completed such that all arrows originate from a single object and terminate at a single object?[13]

We may now answer these questions.

#### 4.4.1 Questions are the Generalisation of Evocation

*Generalisation* is the process of splitting the extension into an unsaturated (or incomplete) and a saturated (complete) part (in Frege's sense of the terms [9]). It is the former that constitutes the intension. Data, for example, may be associated with a Schema, such as a table, and the values that can fill the empty spaces. Representations can be split into a Layout, such as a set of axes, and the markers that are placed in the space that those axes define. In the case of the contents of the mental states evoked by *reading* a Representation, the equivalent is to split the proposition describing that mental state into a property for which there may be some object of which it can be truthfully predicated. To put it simply: if *reading* a Representation, such as a bar chart of student exam results, evokes the thought that 'Alan got the best mark' (or $best\_mark(Alan)$), then the generalisation of this is the question 'Who got the best mark?' (or $best\_mark(\_)$). In other words, Questions are system predicates at the intensional level whilst Evocation involves the extension of those predicates for specific cases. Evaluation of these extensional predicates allows truth statements about the System to be tested.

Completing the *generalisation* functor, from extension to intension, thus provides a salutory reminder of the importance of the underlying Question in visualisation; "From Killer Questions Come Powerful Visualisations" as Johnstone puts it [14]. In software development terms, answering a question is the *requirement* of the Representation. This is even true in the case of *exploratory* data analysis (the terminology is due to Tukey [25]), where the purpose of the data analysis process is not to answer a specific prior Question or hypothesis, but to discover hypotheses worth subsequent testing using conventional *confirmatory* data analysis. Exploratory data analysis is what happens when you don't know what question you're trying to answer. Confirmatory data analysis starts with Data and a Question and then seeks an answer using a Representation. Exploratory data analysis starts with Data and a Representation, and then seeks a Question worth asking.

#### 4.4.2 Knowledge is the Terminal Object of Visualisation

Data is not the start of the visualisation process, and nor is the ability to answer Questions about that Data the end. What we are looking for is Knowledge of the System. Visualisation starts with a System that we measure in various ways in order to generate Data. That Data will always be partial (in both senses of the word) but it's all we've got. Similarly the only Knowledge we can gain is what we can deduce from the evidence presented in the Representation, and the Questions define what Knowledge we can gain from the visualisation process. That is, Knowledge operationalises the Questions by allowing the abstract nature of the Questions to be practically measured or assessed. It is by answering questions that one gains knowledge. This relationship is also analytic: Knowledge is not well-defined unless it is capable of answering questions, and those questions are epistemologically prior to the knowledge of the answers. (To see this, one can imagine a question to which there is no answer; but not an answer for which there is no question.)

The ultimate purpose of the visualisation process is to gain Knowledge of the original System. When this succeeds (when the diagram commutes) then the result is a *truth* relationship between the Knowledge and the System. When this process breaks down and we fail to deduce correct conclusions then the diagram does not commute.[14]

#### 4.4.3 The Completed Category

The full diagram describing the visualisation process in general is shown in Fig. 13 and a version showing a specific example is shown in Fig. 14 (the positions of the Layout and Representation objects have been changed but otherwise the commutative diagram in Fig. 14 is the same as that in Fig. 13). In this example we start with students in a Maths class (System). We want to know how well they are performing, overall and individually, (Knowledge) so we determine that we will have to gather data about their performance in a test (Schema). This Schema, as well as being a generalisation of the Data is also a time-invariant $descriptive$ abstraction of

13. Most well-behaved categories in mathematics have initial and terminal objects.

14. Note, that although this is a strongly realist and representational use of terms such as truth and knowledge it does not necessarily imply a commitment to objectivism about the status of that knowledge. See, for example, Faith [8].

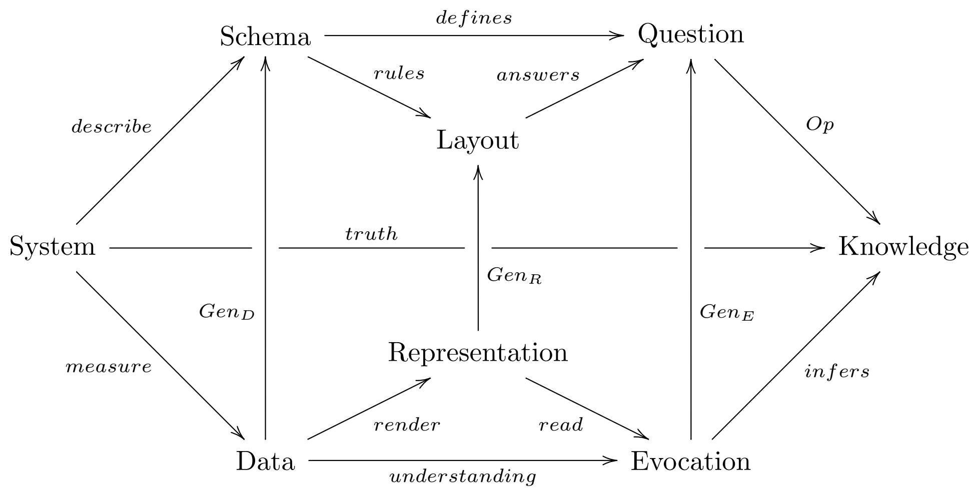


Fig. 13. The visualisation process as a category

the System.[15] Having gathered the results (Data) we then choose a bar chart (Layout) that will be capable of showing the distribution of their overall, and individual marks (Questions). The Schema acts as a frame and determines, or $defines$ what Questions can be answered by the Layout. For example, by seeing we have a table with two columns labelled 'Grade' and 'Name' we can tell that the Data denoted by this Schema will enable us to answer questions like "What was the average mark?", but not questions like "Who was the tallest?".

The Layout, in turn, is used to $answer$ the Questions, but only some Layouts derived from the Schema are capable of $answer$ing the Questions. For example, a bar chart will enable us to spot the best student; but a pie chart would not.

Showing the data (Representation) in this way reveals that the cohort achieved an average mark of more than 70, that Alan got the highest mark, etc. (Evocation). So now we have Knowledge about our System: we know who performed best of all, that nobody failed, the the average mark was at a satisfactory level (this requires prior context), etc.

# 5 Conclusions

In this article we have shown how the process of visualisation can be studied formally using the mathematical language of category theory. By identifying the principal objects implicated in visualisation and the morphisms that map between the structures in these objects we are able to define in a precise and unambiguous manner some of the common terms and concepts used in visualisation. Prior to this many of these concepts were undertsood intuitively or heuristically, but now we have formal theoretical descriptions that can be used to make principled judgements about the visualisation process in general and about individual specific instances in particular. The category shown in Fig. 13 defines a well-formed visualisation process; if all the assumptions underlying the morphisms and objects are met then the diagram will commute and we can be confident that the knowledge gained is reliable and reflective of the system. This category theory description gives us reasons for why certain attempts at visualisation are not as successful as their designers had hoped.

15. Formally, a Schema has "the structure of a continuant which does not specify time or timelike relationships" [23, p. 73].

## 5.1 Where Do We Go From Here?

This article has provided a foundation for discussing visualisation formally and theoretically. Using these theoretical foundations it is now possible to explore other properties of the visualisation process. The category shown in Fig. 13 provides a high-level view of the visualisation process, but there is more detail that can be explored using category theory. For example, the relationship between the Schema and the Layout merits further investigation. More details on a multi-level categorical architecture for the intension-extension relationship are given elsewhere [22].

### 5.1.1 *Layout is a Cartesian Product Corresponding to Individuation of Attributes and Instances*

As an example of how the ideas presented here might be applied, consider the following. One of the most common types of Schema for Data is the familiar tabular form comprising a set of instances (rows) for which we have values of a set of attributes (columns). In category-theoretic terms this Schema is a Cartesian product of Instances and Attributes:

$$\text{Schema} = Instances \times Attributes$$

In such cases we can also ask whether the Layout can be considered as a product in its own terms. For example, the Layout of a scatter plot could be considered as a product of the points ($L_I$, representing the instances) and axes ($L_A$, representing the attributes), with projection morphisms $\Pi_{LI}$ and $\Pi_{LA}$ which respectively pick out

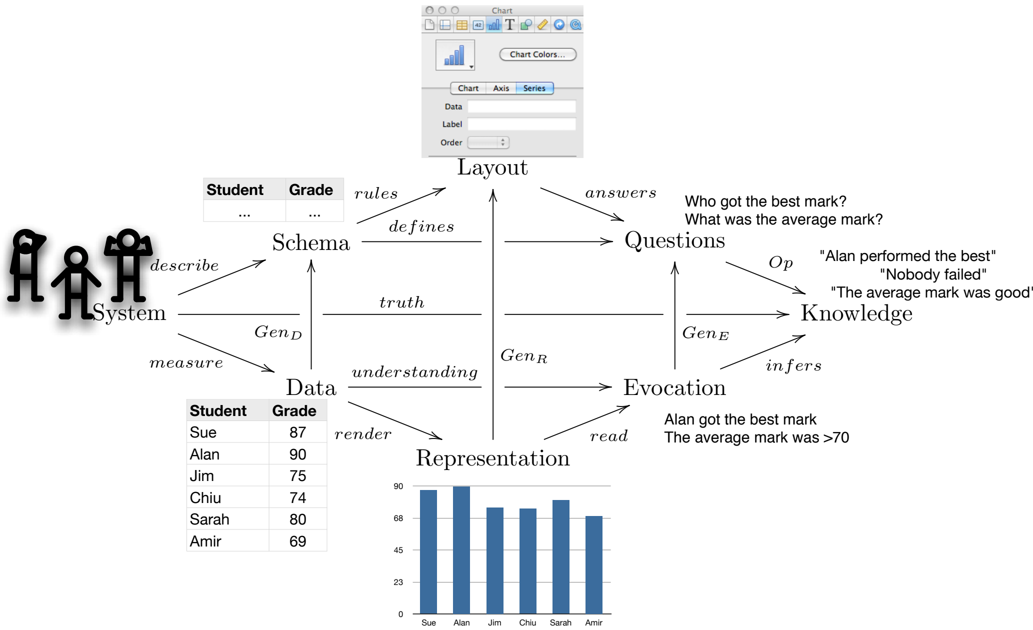


Fig. 14. Example of the visualisation category in practice

the points and axes from a given Layout. In category theoretic terms, this is a valid product if for every Schema for which we can define *rules* for mapping onto the points and axes the diagram in Fig. 15 commutes.

$$\begin{tikzcd} & L_I & \\ \text{Schema} = I \times A \arrow[ur, "rules_I"] \arrow[dr, "rules_A"'] \arrow[rr, dotted, "rules"] & & L_I \times L_A \text{ Layout} \arrow[ul, "\Pi_{LI}"'] \arrow[dl, "\Pi_{LA}"] \\ & L_A & \end{tikzcd}$$

Fig. 15. Individuation of Attributes and Instances in the Layout

How could this fail to be the case? Consider the following example. Suppose we have a Data set in which two instances share the same attribute values, such as two students with identical scores. In this case the points representing those two instances would coincide and the set of points picked out by $\Pi_{LI}$ from the Layout would not be the same as those constructed from the Schema: one would be missing. In order to enforce the condition that the Layout is a product then some way of representing coincident points is required.[16]

There is also a symmetric problem of individuating attributes. This is not apparent in most Layouts, such as scatter plots, which treat Instances and Attributes asymmetrically. In particular, they assume a fixed number of attributes (normally one or two) but don't fix the number of Instances.

16. This is a common problem when representing large or dense data sets and a range of solutions are suggested in Unwin et al. [26]

There are some exceptions. Dimension reduction visualisation techniques, such as principal component analysis (PCA) and mutlidimensional scaling (MDS) compress data with arbitrary numbers of data dimensions into a two-dimensional plane, which has the effect of obscuring the contribution of each individual attribute. These are not cartesian products since the attributes are not individually represented, hence there is no projection morphism $\Pi_{LA}$. Other techniques, such as Parallel Coordinate Plots, on the other hand have an explicit and separate Layout element for each attribute — though one result of this is that they are not suited to representing very large numbers of attributes.

### 5.1.2 Algebraic Semiotics

Goguen suggested the combination of algebra and semiotics could bring much to user interface design and visualisation. An advantage of formalisation, he argued is that "by forcing one to be explicit, some subtle issues are exposed that usually get glossed over" [10]. Goguen introduced the concept of algebraic semiotics to the study of visualisation, but where he restricted himself to a context-free structuralist view, there would appear to be much to be gained from investigating how semiotic morphisms could be used in a post-structuralist landscape in which context (or ground) is not just admitted but required.

### 5.1.3 The Functors of Multi-Modal Perceptualisation

Another very interesting avenue deserving of further exploration is that of multi-modal representations. If we take visualisation to be the branch of *perceptualisation* that restricts itself to visual representations there arises the question of how the formal treatment presented in this article might be applied to the other modalities, specifically, auditory display and haptic representations. Of particular interest here would be the role category theory and semiotics could play in discussing the relationships between representations across the modalities. For example, much work in the field of auditory display and sonification has been directed at trying to find auditory equivalents of visual representations. This is motivated by several factors, the two principal ones being universal access (making perceptualisation available to the visually impaired, for example) and sensory loading (transferring load to other senses such as hearing and touch as the visual channel becomes increasingly overloaded).

Where this category theoretic approach might be useful is in identifying the functors between the categories of perceptualisation processes. For example, it would enable us to answer questions such as 1) "is this auditory display equivalent to this visual representation?" or 2) "is it possible to produce an equivalent visual representation of this sonification?". In the case of the first question we are, in effect, asking whether two processes in two modalities are isomorphic. If they are, what does this mean in practice? Do we mean the representations are isomorphic or that the interpretants are isomorphic, or both, or neither? How can we know? Is it useful to know? Does knowing about the nature of the mappings in one modality inform what we know about the other modality?

In the case of the second question if the answer is 'no' then we would be able to state why it is not possible. Therefore, further work needs to be done to investigate this aspect as it offers the possibility of being able to reason about external representations across modalities in the perceptualisation field and thus, potentially, brings a very powerful tool kit to bear on perceptualisation design. It is possible that Category Theory offers a framework for deciding such issues.

## References


[1] S. Barrass and P. Vickers, "Sonification design and aesthetics," in *The Sonification Handbook*, T. Hermann, J. Neuhoff, and A. D. Hunt, Eds. Berlin: Logos, 2011, p. in press.

[2] S. Bateman, R. L. Mandryk, C. Gutwin, A. Genest, D. McDine, and C. Brooks, "Useful junk?: The effects of visual embellishment on comprehension and memorability of charts," in *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, ser. CHI '10. New York, NY, USA: ACM, 2010, pp. 2573–2582. [Online]. Available: http://doi.acm.org/10.1145/1753326.1753716

[3] S. K. Card, J. Mackinlay, and B. Shneiderman, "Information visualization," in *Readings in Information Visualization: Using Vision to Think*, K. Card, Stuart, J. Mackinlay, and B. Shneiderman, Eds. Morgan Kaufman, 1999, pp. 1–34.

[4] D. Chandler, *Semiotics: The Basics*, 2nd ed. Routledge, 2007.

[5] D. Cox, "Metaphoric mappings: The art of visualization," in *Aesthetic Computing*, ser. LEONARDO, P. A. Fishwick, Ed. Cambridge, MA: MIT Press, 2006, pp. 89–114.

[6] G. Deleuze, "How do we recognise structuralism?" in *Desert Islands and Other Texts (1953-1974)*, ser. Foreign Agents, M. Taormina, Ed. Semiotext(e), 2002, translated by David Lapoujade.

[7] U. Eco, *A Theory of Semiotics*. Bloomington, IN/London: Indiana University Press/Macmillan, 1976.

[8] J. Faith, "Emergent representations: Dialectical materialism and the philosophy of mind," D.Phil thesis, University of Sussex, 2000.

[9] G. Frege, "Function and concept," in *Translations from the Philosophical Writings of Gottlob Frege*, P. Geach and M. Black, Eds. Blackwell Publishers Ltd, 1960.

[10] J. Goguen, "Semiotic morphisms, representations and blending for interface design," in *AMAST Workshop on Algebraic Methods in Language Processing*, F. Spoto, S. Giuseppi, and A. Nijholt, Eds. Verona, Italy: AMAST Press, 25–27 Aug. 2003, pp. 1–15.

[11] J. A. Goguen and D. F. Harrell, "Information visualization and semiotic morphisms," in *Multidisciplinary Approaches to Visual Representations and Interpretations*, ser. Studies in Multidisciplinarity, G. Malcolm, Ed. Elsevier, 2005, pp. 83–97.

[12] R. W. Hamming, *Numerical Methods for Scientists and Engineers*. McGraw-Hill, 1962.

[13] L. Hjelmslev, *Prolegomena to a Theory of Language*. Madison: University of Wisconsin Press, 1961 (1943).

[14] M. Johnstone, "From killer questions come powerful visualisations," Fell In Love With Data blog at http://fellinlovewithdata.com/reflections/killer-questions, 6 Jun. 2011.

[15] R. Kosara, "EagerEyes blog," http://eagereyes.org.

[16] S. Mac Lane, *Categories for the Working Mathematician*. Springer-Verlag New York, Inc., 1971.

[17] J.-P. Marquis, "Category theory," in *The Stanford Encyclopedia of Philosophy*, spring 2011 ed., E. N. Zalta, Ed., 2011. [Online]. Available: http://plato.stanford.edu/archives/spr2011/entries/category-theory/

[18] C. K. Ogden and I. A. Richards, *The Meaning of Meaning: A Study of the Influence of Language upon Thought and of the Science of Symbolism*, 8th ed. New York: Routledge & Kegan Paul, 1923.

[19] H. Purchase, N. Andrienko, T. Jankun-Kelly, and M. Ward, "Theoretical foundations of information visualization," in *Information Visualization*, ser. Lecture Notes in Computer Science, A. Kerren, J. Stasko, J.-D. Fekete, and C. North, Eds. Springer Berlin / Heidelberg, 2008, vol. 4950, pp. 46–64. [Online]. Available: http://dx.doi.org/10.1007/978-3-540-70956-5_3

[20] A. Quigley, "Aesthetics of large-scale relational information visualization in practice," in *Aesthetic Computing*, ser. LEONARDO, P. A. Fishwick, Ed. Cambridge, MA: MIT Press, 2006, pp. 315–333.

[21] Radiohead. (2007) House of cards. [Online]. Available: http://code.google.com/creative/radiohead/

[22] N. Rossiter, M. Heather, and D. Nelson, "A natural basis for interoperability," in *Enterprise Interoperability: New Challenges and Approaches*, G. Doumeingts, J. Müller, G. Morel, and B. Vallespir, Eds. Springer, 2007, pp. 417–426.

[23] J. F. Sowa, *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks/Cole, 2000.

[24] E. R. Tufte, *The Visual Display of Quantitative Information*. Cheshire, CT: Graphics Press, 1983.

[25] J. W. Tukey, *Exploratory Data Analysis*. Addison-Wesley, 1977.

[26] A. Unwin, M. Theus, and Ho, *Graphics of Large Datasets: Visualizing a Million*. Springer, 2006.

[27] A. N. Whitehead, *Modes of Thought*. Macmillan, 1938.

[28] C. Ziemkiewicz and R. Kosara, "Embedding information visualization within visual representation," *Advances in Information and Intelligent Systems*, vol. 251, no. 307–326, 2009.