

Representational Correspondence as a Basic Principle of Diagram Design

Christopher F. Chabris* and Stephen M. Kosslyn

Department of Psychology, Harvard University,
Cambridge, MA 02138, USA
{cfc, smk}@wjh.harvard.edu

Abstract. The timeworn claim that a picture is worth a thousand words is generally well-supported by empirical evidence, suggesting that diagrams and other information graphics can enhance human cognitive capacities in a wide range of contexts and applications. But not every picture is worth the space it occupies. What qualities make a diagram an effective and efficient conduit of information to the human mind? In this article we argue that the best diagrams depict information the same way that our internal mental representations do. That is, “visual thinking” operates largely on relatively sketchy, cartoon-like representations of the physical world, translating sensory input into efficient codes before storing and manipulating it. Effective diagrams will assist this process by stripping away irrelevant detail while preserving or highlighting essential information about objects and their spatial relations. We discuss several examples that illustrate this “Representational Correspondence Principle,” and we consider its implications for the design of systems that use diagrams to represent abstract, conceptual knowledge, such as social networks, financial markets, or web content hierarchies.

1 Introduction

Diagrams are uniquely powerful tools for communication. Everyone has heard the adage that a picture is worth a thousand words (or ten thousand, according to the Chinese version), and this adage bears repetition because it so often is correct. Diagrams, which we take to include information graphics and all other non-photographic forms of visual communication, are usually good examples of this adage at work—but not simply because they save time or space. Rather, diagrams are useful because they make explicit and accessible *to human users* patterns among facts. Our brains have particular properties that must be respected if a diagram (or anything else) is to communicate effectively, and a good diagram exploits our strengths and does not fall prey to our limitations.

Conceptual maps, such as systems of labeled nodes and links used to represent the relationships between individuals, ideas, or other abstract content, are becoming increasingly common as the computing power and tools needed to design them become more available. Indeed, the power of software to implement increasingly

* Corresponding author.

complex, colorful, and rapidly accessible knowledge visualizations may encourage researchers and information designers to develop representational schemes that will overwhelm, or at least unduly tax, the cognitive powers of the very users they are supposed to be serving. In this chapter we first discuss diagrams as general tools for representing information, and propose a principle of effective diagram design. After explaining the key to this new principle, we illustrate its force with examples drawn from different domains, and then apply this principle to the design of conceptual maps and similar knowledge visualization tools.

1.1 The Power of Diagrams

Tufte (1983) describes an early example of the power of diagrams that is still very relevant today: the solution of the September 1854 cholera epidemic in London. The epidemic showed no signs of abating, and many people puzzled over its causes and possible solutions. Dr. John Snow (along with others) pondered tables of data that



Fig. 1. A portion of Dr. John Snow’s original map of the 1854 London cholera epidemic. Note the clustering of cases near the Broad Street pump




Fig. 2. Visual stimulus used by Bransford & Johnson (1972; redrawn from original)

described the locations of victims, and then hit on the idea of converting these data into a graphical display, as shown in Figure 1. Once he displayed them as symbols on a map, Dr. Snow noticed a pattern in the data: the deaths clustered around the location of a particular water pump, which he inferred was the most likely source of the infection. This sort of diagram can help the user discover spatial patterns that otherwise are extraordinarily difficult to discern, and today this form of graphic display is implemented in many database software packages and is seen widely in all types of publications, from epidemiology journals to the popular press.

In addition to such dramatic anecdotal evidence of the power of diagrams in visual communication, there is a longstanding and overwhelming confluence of quantitative research on the topic. Bransford and Johnson (1972) provided a now-classic illustration, if you will, of the power of diagrams in a study that relied on the picture shown in Figure 2. They showed this picture to participants along with a printed story that included sentences like “If the balloons popped, the sound wouldn’t be able to carry since everything would be too far away from the correct floor,” “Since the whole operation depends on a steady flow of electricity, a break in the middle of the wire would also cause problems,” and “With face to face contact, the least number of things could go wrong.” They also asked other, less fortunate, participants to study the text without ever seeing the picture. All were later asked how well they understood the story and tested to see how much of it they could recall. The results were clear: adding the picture improved comprehension and memory by over 50%, and showing the picture before the text was over 100% better than giving the text alone.

This experimental situation relied on unusually ambiguous text, but we can still learn valuable lessons from it. First, pictures can be useful in helping readers to interpret and remember text. Indeed, Levie and Lentz (1982) surveyed 46 experiments comparing text with pictures to text alone and found that 45—all but one—showed that pictures did in fact improve memory or comprehension. In one case, a group following directions in text illustrated with diagrams did an amazing 323% better than a group following the same directions without the illustrations. Second, to be maximally effective, the diagram should be examined *before* the reader encounters the relevant text, in part because the diagram helps to organize the text and in part because the reader may try to visualize what the text is describing, and the results may not match the diagram. However, adding pictures to prose is not a panacea. Levie and Lentz found that whereas illustrations that were merely “vaguely related” to accompanying text led readers to score 25% better on later tests of understanding and memory than text alone, truly irrelevant illustrations had a minimal effect (5% improvement). Worse yet, pictures serving a purely decorative purpose actually cause readers to perform more poorly than those who received unadorned text.

1.2 Understanding Effective Diagrams

In short, it’s not that all diagrams are inherently good: some are, but some are not. But why? What makes one diagram the key to good comprehension, but another a puzzle in its own right? Research that shows the value of adding illustrations to text does not reveal the design characteristics that make some illustrations more effective than others.

We have argued elsewhere that research in cognitive psychology and neuroscience yields principles of display design that play to the strengths of human perception, memory, comprehension and reasoning while avoiding their weaknesses (Kosslyn, 1985, 1994a; Kosslyn & Chabris, 1992, 1993; Kosslyn, Chabris, & Hamilton, 1990). We could present here a long catalogue of facts about how the mind and brain work that are relevant to diagram design. On the one hand, many of these facts point out specific weaknesses in human information processing, such as our notoriously poor abilities to hold information in short-term memory or to discriminate subtle changes. Clearly, good displays should not require that the user have super-human abilities, and thus designers should eschew displays that are too complex or that don’t have

good contrast among separate portions. On the other hand, relevant features of our information processing systems can be used to overcome its limitations. For example, the Gestalt Laws of organization dictate how perceptual units are formed, and the limits of short-term memory are defined in terms of such units. Symmetrical shapes, for example “[]”, are organized into a single unit, whereas asymmetrical ones, such as “_ l”, are not. Thus, by cleverly organizing a display to minimize the number of perceptual units, a designer can pack a lot of information into it. Similarly, by using contrast to define a “foreground” (such as a line in a graph that specifies data from a country of interest) and “background” (all other lines), a designer can direct the viewer’s attention to the most important features, thereby not requiring him to understand everything else in detail.

Such facts are useful, and can lead to nuts-and-bolts recommendations for how to produce comprehensible displays (e.g., see Kosslyn, 1994a). We expect principles based on such facts to be developed as more researchers study the anatomy of effective displays. However, there is another approach to understanding the “cognitive ergonomics” of effective diagrams, which we develop in the following section.

2 The Representational Correspondence Principle

Rather than focus on particular characteristics of our information-processing systems for vision, memory, comprehension and reasoning, in this article we describe a more general principle of graphic communication. This principle may help designers, software developers, and researchers to take advantage of psychological research in a novel way. We call this principle the *Representational Correspondence Principle*, which states that *effective diagrams depict information the same way that our internal mental representations do*. This principle is rooted in the observation that all visual input is translated into internal codes before it is operated on by reasoning processes. Although these translation steps seem effortless in many everyday situations (for example, we are typically completely unaware of all the brain activity going on in the split-second before we recognize a familiar face), they can require a surprising amount of effort in other situations, such as when we must decode a confusing diagram in order to install a new component in a computer. Here’s our crucial idea: Information will pass through the translation bottleneck faster and less painfully if it starts out in a form that corresponds as closely as possible to the one in which it eventually will be specified.

Depending on one’s point of view, the Representational Correspondence Principle may seem obvious or vacuous. We argue that it is neither: First, if it were obvious, then we wouldn’t expect to find that other principles have been proposed that are inconsistent with it. In fact, such inconsistent principles have been seriously proposed. Tufte (1983, p. 93) offered a salient example when he suggested that “a large share of the ink on a graphic should present data-information, the ink changing as the data change. *Data-ink* is the non-erasable core of a graphic, the non-redundant ink arranged in response to variation in the numbers represented.” He went on to argue that the best graphic maximizes the ratio of “data-ink” to “total ink used to print the graphic,” and that “most important ... is the idea that other principles bearing on graphical design follow from the idea of maximizing the share of data-ink.” Notice

that this principle suggests that instead of showing a complete bar in a bar graph, the designer would be best advised simply to present one side of the bar and a line demarcating its top. Never mind that the resulting bracket would not be symmetrical, and thus would involve two perceptual units instead of the one that is formed by a complete bar; and never mind that thin lines are more difficult to detect and discriminate from the background than are bars. We will revisit Tufte's specific recommendations that flow from his data-ink principle later. For now it should be clear that Tufte's principle, which has had wide influence, says nothing about human internal representations of information or the limitations on how human beings process information; instead, it addresses itself exclusively to the ink on the page (or, nowadays, the pixels on the screen).

Second, could the Representational Correspondence Principle be vacuous? If one believes that it is impossible or currently beyond science's reach to understand how the mind and brain represent information, then one would conclude that there cannot be any way to implement the principle in practice. Or, one might readily concede that science is revealing properties about how information is represented in the human mind and brain, but be skeptical (as Tufte himself was) that it can teach us any lessons about how to design effective diagrams. In our view, both concerns are misguided. Although it is true that there are many unsolved problems in neuroscience (as in every other branch of science), there is more than enough knowledge about how information is represented to give us detailed guidance in how to design effective diagrams. In the following section, we will hint at some of these discoveries and the various methods that have been developed to plumb the workings of the mind and brain.

2.1 Visual Images in the Brain

Although we cannot yet say definitively how visual information is stored and processed in the brain, considerable progress has been made in such research. One line of research hinges on the idea that much of the visual information we store in memory can be recalled in the form of visual mental images, and thus the study of visual mental imagery can reveal the nature of internal visual representations. This hypothesis is supported by many forms of research. For example, for well over 100 years researchers have reported that visual imagery interferes with visual perception—as expected if the same system is used in both cases. For instance, researchers showed that visualizing impairs the ability to see (Perky, 1910; Segal & Fusella, 1970). Later researchers documented that people falsely remember having seen an object when they in fact only visualized it (e.g., Johnson & Raye, 1981). And yet other researchers focused on functional similarities between imagery and perception (for reviews, see Finke & Shepard, 1986; Kosslyn, 1980, 1994b). For example, objects in mental images require more time to imagine rotating greater amounts, as if the images were literally rotating about an axis (Shepard & Cooper, 1982). Similarly, people require more time to scan greater distances across imagined objects, as if the imagined objects are arrayed in space. Moreover, they require more time to “zoom in” greater amounts when “viewing” imagined objects (Kosslyn, 1980).

With the advent of modern brain-scanning technologies, researchers moved beyond purely behavioral studies of imagery and perception to studies of the underlying neural mechanisms (e.g., Kosslyn et al., 1997). The results are remarkable: At least

two-thirds of the same brain areas are used in visual imagery and visual perception—which is a far greater amount than either shares with language or memory processing. Visual imagery isn’t called “seeing with the mind’s eye” only for poetic reasons; it really does use most of the neural machinery used in actual seeing.

Although many brain areas are shared by visual imagery and perception, of particular importance is the fact that imagery usually recruits the first parts of the brain to register visual input from the eyes (Kosslyn et al., 2001; Kosslyn & Thompson, 2003; Thompson & Kosslyn, 2000). During perception, light strikes the retinas and neural signals are sent back into the brain. These signals in turn evoke a pattern of activation on the surface of the cerebral cortex (the thin outer covering of the brain, where most of the neural cell bodies are located). In the first areas to receive such input, the pattern of activation is literally spatial: It preserves the layout of the pattern of activity on the retina itself, which is why these areas are said to be *retinotopically mapped*. There actually are “pictures in the head.”

This sort of code has several important properties. For example, it uses space on the cortex to represent space in the world. As such, it makes explicit and accessible *shape* and the *spatial relations* among shapes and parts of shapes. In addition, the code is tailor-made to function well within the brain’s processing system. This is a crucial point, so we need to emphasize it: The properties of any representation can only be understood in the context of the systems in which it is embedded. If there were no processes that could interpret shapes (e.g., ^ is readily seen as pointed and U as rounded), the representations would have no impact on the rest of the system—and for all intents and purposes would not exist. Thus, when we evaluate properties of a representation, we need to consider them in the context of the types of processes that operate on them. Specifically, some aspects of representations will be easily operated upon by the extant processes, whereas others will not be. For example, it would be easy to *draw* a rounded point, but the brain might find it difficult to *verbally label* such a shape.

The fact that the initial input to the visual system is picture-like is convenient for those who want to use diagrams to convey information, but if we want our diagrams to correspond as closely as possible to the representations used by the brain itself we need to know more about those representations. The patterns of activation in the first visual areas are just the beginning. These representations are converted to a series of other representations as processing continues. We focus on one fact about these conversions: The brain is often in danger of being overwhelmed by too much information, and thus a crucial aspect of processing involves stripping down representations to their core, preserving some aspects and discarding others.¹ An effective diagram should not only map neatly into the representations used early in processing, but also facilitate the processing such representations evoke. In the next section we unpack this idea.

¹ Although the abstraction of relevant information, or “gist,” is a critical component of human perception, there are in fact some individuals who are able to suppress this tendency, or who even have difficulty *not* being overwhelmed by visual detail. Autistic savants with artistic ability are able to draw surprisingly detailed, naturalistic pictures, unlike all normal children their age (e.g., Snyder and Thomas, 1997), who tend to draw schematically, showing the general shapes of critical parts and their spatial relations, but not the visual details.

2.2 Chess Diagrams

An elegant illustration of representational correspondence is found every week in hundreds of newspapers: chess diagrams. As shown in Figure 3, the configuration of pieces on a three-dimensional chess board can be represented by an array of symbols in a two-dimensional diagram. Indeed, this type of diagram is used throughout the world, in virtually every country and culture where the Western form of chess is played. In particular, all chess publications aimed at expert players use this type of diagram, and have for over a century. Internet chess services, computer chess software, and video projection systems at public chess events use the same format. In fact, even though they could easily have realistic three-dimensional displays, chess professionals prefer to use this format to study the game on their computers. And this diagrammatic representation has even yielded an international convention for a notation to communicate the moves of chess games: the symbols used for the different pieces replace the initials of the piece names in the local language, making it possible for literature produced in one country to be read and understood in many others regardless of language differences.

How does this format demonstrate representational correspondence? To answer this question, we must discover how chess experts internally represent the locations and identities of the pieces on the board when they are playing chess. Again, we can appeal to visual mental imagery as a way to study the nature of the internal representations used to this end. The easiest way to begin this process is simply to ask players what they “see in their mind’s eye” when they visualize a chess position or “think ahead.” Is it a veridical, three-dimensional image of a particular chess board and set of pieces (perhaps the ones currently being played on, or the ones the player most often uses)? Taine (1875), based on the report of one amateur player, believed that it was, and characterized the type of imagery used in chess as an “internal mirror” that reflected the precise state of the thing(s) being imagined.




Fig. 3. A photograph of a chess position on a standard board and set, viewed from a player’s visual perspective (*left*); a standard symbolic diagram of the same chess position (*right*)

But further examination of the chess literature suggests a different conclusion. Reuben Fine (1965), a psychoanalyst and former world championship contender, said that “the visualization that takes place must emphasize the chess essentials and must eliminate accidental factors, such as size or different colors [of the board and pieces].” Commenting on the qualities of the chess player’s internal image, the chess master Jacques Mieses (1940) wrote that “it is not a planimetric or stereometric picture that appears before the mind’s eye of the chess player ... His task lies rather in mentally picturing the constantly changing formations of the pieces ... This process is indeed more closely allied to the department of ‘topographic sense’.” Alfred Binet (1894), the father of modern intelligence testing, surveyed many experts who engaged in simultaneous blindfold chess, a popular public “stunt” in the late 19th century in which the performer played several games at once, entirely without sight of a board. One of Binet’s participants, a young master named Goetz, reported that he was “aware only of the significance of a piece and its course ... to the inner eye, a bishop is not a uniquely shaped piece, but, rather, an oblique force” (Binet, 1893).

One of us (CFC) recently asked American grandmaster Sergey Kudrin to reconstruct chess positions with a standard set and board after five seconds of viewing them as printed chess diagrams. Afterwards, he asked this player how he remembered where the pieces were located. The response: “I visually remembered the diagram, but [the pieces] were on this diagram, which for me is almost the same as the board, but in my mind they stayed as the pieces of this diagram.” Did he have any problems in translating between his memory of the diagram and the three-dimensional board and pieces? “It seemed natural, although maybe I lost it somewhere because I didn’t get this [particular] position exactly. Although at the moment when I stopped looking at [the diagram], I was sure I would remember everything, but I didn’t remember, and I knew I was doing something wrong,” he replied. That is, Kudrin formed some sort of visual image of the chess position with the same characteristics as the two-dimensional diagram he studied, and when the stimulus was removed he had a clear image. But as he began to construct the position on the (three-dimensional) board, the image faded, and he could not complete the task as accurately as he initially expected.

A common thread running through these descriptions, and many others in the literature, is that all players denied using a representation of the three-dimensional qualities of the board or pieces, and instead emphasized that the spatial relationships among pieces are more important than their particular shapes, colors, and so on. Chess diagrams discard all the superfluous detail of the chess board and pieces (and the player’s particular perspective), and in the process make more salient the identities of the pieces (high-contrast black and white symbols, instead of lower-contrast dark- and light-brown woods) and their spatial relationships (each piece’s location is clearly visible, instead of occluding nearby pieces as on a real board).²

Further introspective accounts of chess masters suggest uses to which diagrams are put in visual thinking. Binet’s subject Goetz drew a diagram showing what he “saw” when thinking ahead about a particular chess position, as illustrated in Figure 4; note

² Indeed, the diagram seems so optimal for visual thinking about chess that one might ask why (beyond for historical reasons) three-dimensional sets are still used. They are superior to diagrams for actually grasping and moving the pieces with the hands, and for face-to-face interaction. But chess competition is conducted increasingly via the Internet, and there players find it faster to use a diagram depiction and move pieces by clicking with a mouse.

that some squares are highlighted and movements between squares are also shown. According to Steiner (1972, p. 66), “the great chess player does not see squares and pieces as discrete units, or even abstract counters. He internalizes a very special sense of ‘fields of force,’ of regions characterized, differentiated by the fact that certain events can or cannot take place in them. What matters, thus, is not the particular square, or even piece, but a cluster of potential actions, a space of and for evolving events.” The chess diagram, of course, does not represent these “fields of force” explicitly; they are the phenomena that occur during visual “thinking ahead” by chess masters, whereas diagrams are meant to represent snapshots of a game in progress. What is critical is that none of these aspects of chess thinking revealed by these reports depend on a photograph-like representation of the chess board; indeed, they would apparently be impaired if they had to perform in such a cumbersome, overly detailed arena instead of the simplified two-dimensional diagram.

Thus, chess diagrams appear to capture what is important about a board configuration, and strip away the irrelevant details. Moreover, because the diagram has a 1:1 correspondence with the actual board, it captures all of the possible relations among pieces. This representation is easily internalized, given what we know about how visual information is stored.

Note that we are not claiming that modern chess diagrams and notation systems were developed to conform to the Representational Correspondence Principle. They use a convention that evolved over centuries of practice and publishing. We suspect




Fig. 4. A drawing made by Binet’s chess-master subject Goetz of his internal mental imagery while thinking ahead about a chess position (Binet, 1893)

that they are the product of a kind of Darwinian selection, winning out over other competing schemes, and that their continued success as a representational scheme (and their proliferation in an age when computers can depict three-dimensional scenes and complicated images in video games with lifelike detail) reflects their (inadvertent) adherence to the principle.³ If so, we are led to suggest that novel diagrammatic schemes in other fields may be developed by carefully debriefing expert practitioners, trying to discover how they convert relevant stimuli into internal representations. Of course, not all aspects of representations are accessible to introspection, but many of the functional properties of visual images are in fact evident to introspection (Kosslyn, 2001), as in the case of chess.

2.3 Caricatures of Faces

We noted earlier that most human beings can effortlessly recognize hundreds of different individual faces. Face recognition is completed in less than one second, usually with no conscious thought, and—unlike the situation in chess—we have no introspective access to the procedures we use to do it or the way that different faces are represented in our memories. Faces themselves are not diagrams, but they can be represented by diagrams—assuming that caricatures of faces are diagrams. Caricatures are drawings that exaggerate distinctive features (as in political cartoons of George W. Bush that emphasize his ears) and de-emphasize nondistinctive features (such as Mr. Bush’s chin). What can these illustrations teach us about diagram design in general?

An elegant line of research suggests that faces, despite our ability to recognize them rapidly and seemingly without effort, are not internally represented in a veridical, photograph-like format. Rhodes, Brennan, and Carey (1987) took a set of faces that were familiar to their participants, and used a computer program to generate a caricature for each face. The program generated these caricatures by starting with photographs, extracting key points and lines, and then comparing these features to “average” faces. By stretching the features farther from the average, or moving them closer, caricatures and “anti-caricatures” of the faces could be created (see Figure 5 for an example). Participants were then asked to identify the individuals depicted by veridical, caricature, and anti-caricature drawings as quickly and accurately as possible. The findings were straightforward: Viewers identified the caricatures fastest—even faster than they identified veridical depictions—and the anti-caricatures most slowly. When these researchers later showed the participants a range of caricatures and anti-caricatures and asked them to select the best likeness for each person, the average choice was a caricature that exaggerated the distinctive aspects of the face by about 16%.

³ The Darwinian selection hypothesis is not the only possibility, of course. Simple tradition, ease of production, or cost could also explain why a particular sort of display has persisted. However, we doubt that a truly ineffective display would be retained for long if better alternatives are developed. Consider the case of the dial clock: Many predicted its demise when digital watches and displays became common. But the dial conveys information explicitly that needs to be computed from a digital display: Namely the proportion of the hour that has passed and that remains. Dial clocks do something well that isn’t done well by the current alternative, and hence are unlikely to be supplanted by them.




Fig. 5. Caricatures of John F. Kennedy in 0% (*left*), 50% (*center*), and 100% (*right*) exaggeration away from the average face. The exaggerated versions are easier to identify (Rhodes, Brennan, & Carey, 1987; reprinted by permission of Gillian Rhodes)

Accordingly, although we may not be aware of it, our internal representation of a human face appears to magnify what distinguishes it from other faces and minimize what it has in common with them. This makes sense: if the task is to recognize a particular individual's face, the facts that he has a mouth, two eyes, a nose, and so on will not help us much. When presented with a caricature, the unconscious translation process that converts a face into the internal code does not have to work as hard, and we are able to recognize the face faster. Consistent with this view, a subsequent study by Rhodes and McLean (1990) found that people who are expert in identifying birds benefited similarly when they identified caricatures of particular birds compared to when they identified veridical drawings of them.

The lesson for diagram design is clear: Caricatures make use of the Representational Correspondence Principle by matching a stimulus more closely to our internal representation of the represented object, and thereby facilitate our encoding and using the representation. Moreover, our internal representation is not simply a mental photograph; rather, these representations emphasize the most useful aspects of the stimulus and de-emphasize those aspects that will not help us in the most common tasks involving those stimuli.⁴ In fact, caricatures appear to resemble our internal representations more closely than veridical drawings of the same objects.

⁴ How much of what we feel we are seeing at any given time is actually being represented and saved in memory? Recent research on "change blindness" suggests that surprisingly little visual detail, even visual detail we think we must be storing, persists beyond the point when we stop looking at an object or scene. For example, Simons and Levin (1998) conducted a study in which an experimenter approached an unwitting pedestrian and asked for directions to a nearby building on a college campus. While the pedestrian was giving directions to the experimenter, two other experimenters carrying a door passed between the them. As the door passed, the first experimenter switched places with one of the experimenters who had been carrying the door, so that once the door was gone, the pedestrian was now talking to a different person from the one who initially asked for directions. Approximately half of the pedestrians approached in this study did not notice the change at all, and similar results have been obtained many times since in other studies (e.g., Levin, Simons, Angelone, & Chabris, 2002). In general, despite our introspective belief that we perceive the full appearance of someone we are talking to, this line of research suggests that only the most critical information, such as sex, age, and ethnicity, is guaranteed to be stored. Diagram designers should always keep in mind the inherent paucity of detail in human perception and memory.

In general, we propose that diagrams drawn to exaggerate or highlight critical distinctions will be more effective than veridical drawings. This recommendation will apply especially when diagrams must be compared or otherwise differentiated from one another. Note that in this case, the important qualities of an internal representation were only understood through measuring speed of performance in cognitive testing (identification), not through introspection. Representational properties can be revealed by a range of different techniques.

3 Consequences of Violating the Principle

What about traditional information graphics? Does the Representational Correspondence Principle apply to situations in which the constituents of the diagram are purely symbolic representations of quantities rather than depictions? If this principle is in fact general, as we suggest, then violating it should render any sort of diagram difficult to encode and understand. Let's consider some examples.

3.1 Keeping the Bar in Bar Graphs

Tufte (1983, p. 96) attacks the traditional bar graph as wasting ink, and thus violating his data-ink efficiency principle. In particular, he notes that:

“The labeled, shaded bar of the bar chart ... unambiguously locates the altitude [the quantity represented by the bar] in six separate ways (any five of the six can be erased and the sixth will still indicate the height): as the (1) height of the left line, (2) height of the shading, (3) height of right line, (4) position of top of horizontal line, (5) position (not content) of number at bar's top, and (6) the number itself.”

He then redesigns a traditional bar graph from a scientific journal by erasing many of the lines that form the bars, leaving essentially a single vertical line for each bar, and connects the baselines of pairs of adjacent bars, a process that “improves the graphic immensely” (p. 101). Figure 6 shows an example of this process.

Is it true that turning bars into lines or points will make communication more effective? Tufte is correct, of course, that much of the ink used to draw a bar is “redundant” in a mathematical sense. But the visual system does not represent the separate elements of a simple object like a bar. In fact, individual neurons respond to bars of different lengths and orientations, which suggests that the bar itself is one of the fundamental stimuli for the brain to process. Deconstructing a bar into a set of lines or points, as Tufte recommends, *converts it from a single object into multiple objects*, which will actually increase the load on processing and memory. This is because, as recent research in psychology and neuroscience has shown, objects (not pixels or ink) are a fundamental unit of representation (for a review, see Scholl, 2001). For example, viewers habitually register spatial relationships between objects, and can pay attention to an individual object more readily than to its parts. Representational correspondence therefore suggests that if what must be depicted is a single quantity, the diagram component chosen should be an object—which will be encoded into memory with less effort than will isolated parts. In fact, Tufte's redesign

rules have created four new objects (connected sets of lines that resemble hooks), which are not easily decomposed by the visual system into their constituent “bars.”




Fig. 6. A bar graph (of fictional data) drawn in the traditional way (*top*) and after applying Tufte’s suggested transformations to increase the “data-ink ratio”

In fact, when Gillan and Richman (1994) actually tested Tufte’s idea, they did not find that people could interpret simple bar graphs faster or more accurately after they were pared down; even when only two bars were present, participants tended to require more time to use the pared-down bar graphs than the standard ones. Gillan and Richman note that “ink can be helpful or detrimental to a graph reader depending on the function and location of the ink, the user’s task, the type of graph, and the physical relations among graphical elements.” (p 637). And they continue, “Thus, these data call into question Tufte’s general rule that graph designers should maximize the data-ink ratio by eliminating non-data-ink and redundant ink.” (p. 638).

A bar is the simplest visual depiction whose properties capture all that we need to know about quantity *in a single object*. Note that a circle at a specified height is not as good as a bar: the viewer must register not only the circle, but also the distance between the circle and the x-axis to estimate the height; in this case, the viewer would need to compare two objects, rather than encode the length of a single object. Also note that making bars appear three-dimensional actually detracts from processing them in bar graphs (Fischer, 2000). Although a 3-D bar is, like a 2-D bar, a single object, it is harder to place these bars properly within the axes, and difficult to avoid making the lines that depict the third dimension seem relevant. Three-dimensional bars also require the reader to ignore truly irrelevant information (such as apparent distance), which taxes processing. Note, however, that the added information is “irrelevant” not in Tufte’s sense of quantitative aesthetics, but in the sense of what will optimize our cognitive performance.

3.2 Face Displays and Train Schedules

As a general rule, diagrams that violate the Representational Correspondence Principle are likely to be ineffective or obfuscating. In this section, we note two more examples, also coincidentally endorsed by Tufte as ink-efficient displays. First, as illustrated in Figure 7, the so-called Chernoff faces (Chernoff, 1973; Flury and Riedwyl, 1981) use different facial features to represent different variables: the expression of the eyes, the size of the mouth, the width of the face, the location of the ears, and so on, all may specify the values of different variables, enabling a compact representation of multidimensional information. Unfortunately, these face displays are almost impossible to use for extracting and comparing this information. Why? Human faces are processed primarily as single objects, not as collections of individual features. It is notoriously difficult to recognize individual facial features, and comparing features to one another is difficult when they must be isolated from their facial context. Given all that is known about visual processing of faces (for recent reviews, see Haxby et al., 2002; Rakover, 2002), it is hard to imagine a worse way to communicate multiple variables.⁵

Second, the 1880 Paris–Lyon train schedule designed by Marey (detail shown in Figure 8) is a grid with time along the x-axis, and location (Paris, Laroche, Dijon, etc.) along the y-axis. Each train is represented by a line that starts at the top when it leaves Paris, or starts at the bottom when it leaves Lyon. The line progresses across and down (or up) the grid until it reaches the final destination. When a line intersects a

⁵ For example, Tanaka and Farah (1993) conducted an experiment in which participants studied labeled pictures of faces and houses (e.g., “this is Bob’s face” or “this is Bob’s house”), and later were shown individual features from these types of pictures (e.g., a pair of eyes, or a door) and asked to decide whether the feature was part of a named face or house (e.g., “are these Bob’s eyes?”). On some trials, the part was shown in isolation, but on others it was shown in the context of the whole studied face or house. Showing house parts in context did not improve accuracy compared to showing them alone, but showing face parts in context helped the participants to decide whether they had seen them before; this finding suggests that individual facial features are not well-processed in isolation from the other facial features or the overall shape and context provided by a face. Perhaps designers who wish to pursue the concept behind the Chernoff faces should consider using houses instead—or just stick with a series of old-fashioned bar graphs.

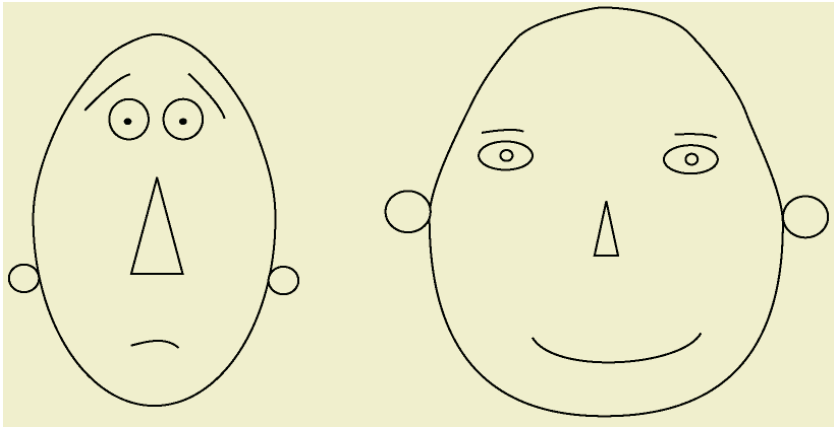


Fig. 7. Two sample “Chernoff faces”

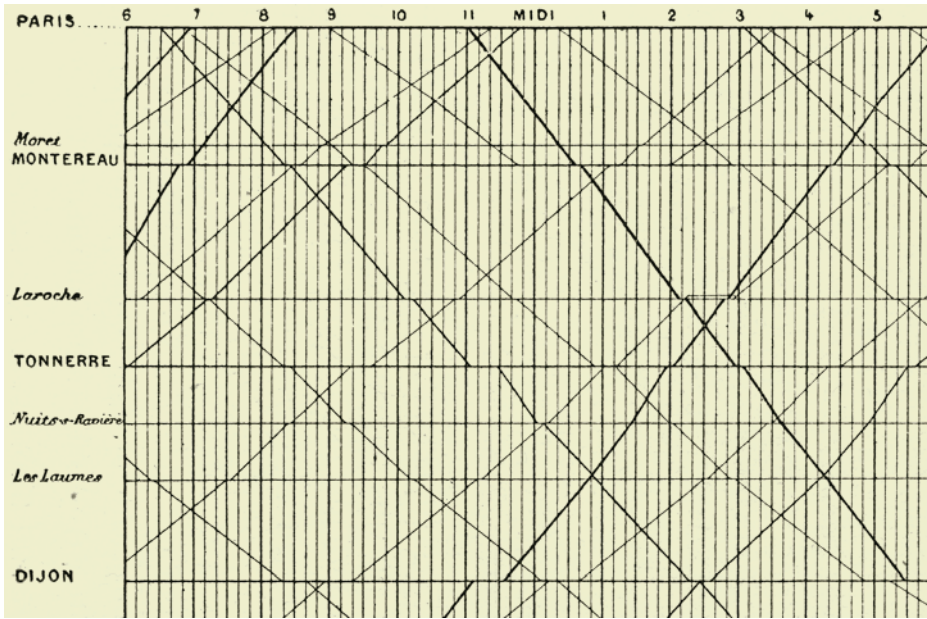


Fig. 8. A detail of Marey’s original 1880 Paris–Lyon train schedule diagram

horizontal line representing a city, that indicates a stop. When viewers look at this diagram, however, all they can see is a mass of oriented lines. Orientation is one of the most salient properties of lines, and lines that differ by only a small amount in orientation are readily discriminated (indeed, line orientation is another one of the properties encoded very early in the brain’s visual system). In this display, orientation conveys information about the train’s *speed*, but the display was constructed to convey information about routes and schedules, so having speed “pop out” of the

display is counterintuitive. Travelers are less interested in a train's speed than when it leaves and arrives, but to determine these facts they must compare multiple objects and scan across the chart to find the corresponding point on the axes. It's no wonder that these displays are rarely seen today.

3.3 Conceptual Networks

Our principle also applies to another sort of display, which is gaining increasing currency: conceptual network diagrams. An early use of such diagrams was the semantic network (e.g., Collins & Loftus, 1975). This now familiar type of diagram shows words or concepts as nodes, with links among them representing how strongly the nodes are associated in human memory. Of course, the "closeness" of two concepts in semantic space cannot be observed directly in the mind or brain, but it can be inferred from the results of different kinds of experiments. These experiments can be as simple as word association tasks, in which many people are asked to say the first word they think of when they hear or see a target word; in this case, the frequency of specific word responses determines the strength of their links to the target word. For example, what do you think of when you see "fork"? "Knife" would be a frequent response, which therefore is inferred to be close to "knife" in semantic space. Various methods can be used to converge on a single representation of a particular semantic network for an individual or a group.

An intriguing use of network diagrams developed more recently is social network analysis. For example, terrorist networks can be analyzed by depicting individual members as nodes, with links indicating the strength of the relationships among the individuals. In such cases, the network diagram can increase our knowledge of the underlying facts used to create it because it uses spatial codes to depict properties and patterns in the data that our limited powers of attention and memory impede us from recognizing in masses of individual facts. For example, Fellman et al. (2003) describe a network analysis of the interactions among the hijackers involved in the September 11, 2001 terrorist attacks. The information used to create the network was extracted entirely from public sources, but the network depiction highlights the central roles of the pilots of the four teams in the planning that preceded the attacks. The network can also be a tool for testing hypotheses by manipulating the data that underlies it; in this case, redrawing the network *without* using the data on the most recent contacts among the hijackers before September 11 suggests that one of the non-pilots was a central hub of interaction, and may thus have served a more critical role in planning than he did in operations. This observation in turn suggests possible new avenues of investigation for intelligence and law enforcement, and might stimulate new thinking about the organization of terrorist networks in general. (See Sageman, 2004, for analysis of terrorist networks on a group, rather than individual, level.)

Bar graphs and chess diagrams are obvious examples of representational correspondence; what about semantic and social network diagrams? The network diagram uses a spatial code to depict inherently non-spatial information. The fact that "cat" is more closely related to "dog" than it is to "train" is not a spatial fact in the same sense that the number of squares between two chess pieces is spatial; the fact that Mohammed Atta met with one of his conspirators more often than he met with another one is not spatial in the same sense that the difference between today's

temperature and yesterday's temperature is spatial. However, the network diagram does represent information in a cartoon-like code that excludes irrelevant information, and it highlights important distinctions, for example by making obvious the difference between "hubs" and "loners" in social networks. A network diagram of the domestic terrorists involved in the Oklahoma City bombing of 1995 would look much different from the diagram of Al Qaeda, instantly showing that the two events were carried out by much different organizations.

Because the network diagram does not depict spatial information, it cannot match the power of the bar graph, caricature, or chess diagram to convey large quantities of information instantly through the human perceptual systems. Network diagrams are better suited to long-term exploration of complex data sets than to short-term communication of patterns and facts. Nonetheless, designers of knowledge visualization systems should keep in mind the general principle of representational correspondence, as well as the specific strengths and weaknesses of human cognitive architecture, if they want to maximize the satisfaction (and repeat business) of their customers.

Unfortunately, some of the most intriguing visualization systems do not fully succeed in these respects. The "Map of the Market" (updated continuously at www.smartmoney.com/maps) is a case in point. A large rectangular display is broken up into smaller rectangles, each of which represents a specific industry sector. Within each sector, smaller rectangles correspond to individual companies. The size of a rectangle indicates relative market value, and its color indicates daily changes (green for increases, red for decreases). One problem with this display is that the human visual system does not excel at comparing the areas of different objects, especially when they are not presented on a common baseline. General differences are apparent (Microsoft is a very large rectangle, Unisys is a very small one), but a more fine-grained comparison is difficult. Moreover, the spatial layout of the sectors and stocks does not seem to correspond to any "map" observers will have encountered previously. In this case a spatial code is being used when there is no underlying spatial or numerical dimension in the information being depicted.

Contrast this with distorted geographic maps, in which countries or regions are scaled up or down in size to match their share of some quantity, such as economic output, foreign debt, or oil consumption. In such maps two different spatial codes are in competition, but the familiar one (geography) can be used as an index to find the other, unfamiliar one. This is not the case in the Map of the Market, or in similar systems for representing the organization of information available on the web, such as ET-Map (ai3.eller.arizona.edu/ent/entertain1/), in which there is no natural index to help a user search for relevant information. Representational correspondence is not being exploited in these displays because the depiction uses a spatial code that does not match any internal representation that the human mind is likely to use or be familiar with.

Even worse, however, is the situation where a spatial code normally used for another purpose is appropriated for representing an incompatible form of knowledge. For example, NewsMaps (mappa.mundi.net/maps/maps_015/) use the metaphor of a topographic map to organize the news stories from a single time period, resulting in a sort of "representational mismatch" that may do more harm than good relative to a simple grouping of headlines into hierarchical categories.

4 Conclusions

The Representational Correspondence Principle states that to be effective, diagrams should depict information in the same way that our internal mental representations do. These internal representations are not veridical photographs, but instead are sketchy and cartoon-like (in the sense that distinguishing characteristics are emphasized). These internal representations do not include irrelevant detail, but “irrelevant” is defined relative to the task they will be used to accomplish and relative to how information is easily encoded, stored, and used in human cognition. This principle applies to diagrams that are intended to be manipulated by human mental processes, such as chess diagrams, bar graphs, and train schedules, all of which must be studied to extract relevant information.

We have two general concluding observations about the Representational Correspondence Principle. First, it does not trump all other considerations in design; indeed, no single principle could or should. In particular, diagram creators should not overlook longstanding conventions and the individual viewer’s practice and experience with specific types of diagrams. For instance, differences in hue are often used to convey variation on a quantitative dimension; images of brain activation, as one example, use different hues to indicate different amounts of activity. However, hue is a so-called “nomothetic” dimension: Variations in hue do not naturally line up with variations along a single quantitative dimension; indeed, psychologically, hue is arrayed as a circle (the famous “color wheel”), not a single continuum. Nevertheless, many users have mastered the conventions (white indicates the largest amount, yellow next, followed by red, and so on). That said, variations in the other two aspects of color—saturation and intensity—do vary psychologically along a continuum. Even for experts, we expect that if these other two variables are manipulated to line up with the information conveyed (e.g., by using brighter colors to indicate the largest amount), and thereby respect representational correspondence, the display will be even more effective than one that follows an accepted but suboptimal convention of representation. Even a two-dimensional chess diagram only demonstrates representational correspondence for viewers who understand the game and are familiar with the symbolic conventions involved; for beginning players or those who have never seen such a diagram before, it may look like nothing more than an overwhelming jumble of odd figurines.

Second, this principle is not set in stone. The implications of the principle will evolve as researchers learn more about how external representations are converted to internal representations, and how internal representations are used in mental processing. Much of this research will focus on basic science, and those researchers will not consider possible applications. But this does not mean that questions that arise from considering diagram design cannot themselves feed into this research. From our perspective, there should be a rich exchange between researchers who study mental processing and those who design diagrams and the systems that create them. We expect such interactions to become especially productive when truly interactive computerized displays become common. According to our principle, such displays will be most effective when they mimic the corresponding mental processes, allowing them to become, in effect, prosthetic devices for the human mind. If well designed,

displays can seem like extensions of ourselves, as easy and natural to work with as mental images.

Finally, we have commented on the design of conceptual network diagrams and other knowledge visualization systems in the context of the representational correspondence principle. Our aim was not to criticize or discourage the originators of novel information displays; every diagrammatic convention that is now in widespread use, facilitating human communication and understanding, had to be invented by someone at some point. We do believe, however, that technology has made it easier to invent a new form of display than to determine how effective a display is. The most successful display conventions illustrate the principle of representational correspondence, and although conceptual networks and other new visualization forms have great potential to help us make sense of more and varied types of information, they will succeed or fail in large measure based on how well they adhere to representational correspondence and respect the limitations of human information processing capabilities.

References

- Binet, A. (1893). Les grandes mémoires: Résumé d'une enquête sur les joueurs d'échecs. *Revue des Deux Mondes*, 117, 826–859. (Translated by M.L. Simmel and S.B. Barron as "Mnemonic virtuosity: A study of chess players," *Genetic Psychology Monographs*, 74, 127–162, 1966.)
- Binet, A. (1894). *Psychologie des grands calculateurs et joueurs d'échecs*. [Psychology of great calculators and chess players.] Paris: Hachette.
- Bransford, J.D., & Johnson, M.K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning & Verbal Behavior*, 11, 717–726.
- Chabris, C.F., & Kosslyn, S.M. (1995). Illustrated editorial is value-added text. *Folio*, February, 28–29.
- Chernoff, H. (1973). The use of faces to represent points in k-dimensional space graphically. *Journal of the American Statistical Association*, 68, 361–368.
- Collins, A.M. & Loftus, E.F. (1975). A spreading activation theory of semantic processing. *Psychological Review*, 82, 407–428.
- Fellman, P.V., Sawyer, D., & Wright, R. (2003). Modeling terrorist networks: Complex systems and first principles of counter-intelligence. Presented at "NATO and Central Asia: Enlargement, Civil – Military Relations, and Security," Kazach American University/North Atlantic Treaty Organization (NATO), May 14–16.
- Fine, R. (1965). The psychology of blindfold chess: An introspective account. *Acta Psychologica*, 24, 352–370.
- Finke, R. A., & Shepard, R. N. (1986). Visual functions of mental imagery. In K. R. Boff, L. Kaufman, & J. P. Thomas (Ed.), *Handbook of perception and human performance* (pp. 37–55). New York: Wiley-Interscience.
- Fischer, M.H. (2000). Do irrelevant depth cues affect the comprehension of bar graphs? *Applied Cognitive Psychology*, 14, 151–162.
- Flury, B., & Riedwyl, H. (1981). Graphical representation of multivariate data by means of asymmetrical faces. *Journal of the American Statistical Association*, 76, 757–765.
- Gillan, D.J., & Richman, E.H. (1994). Miminalism and the syntax of graphs. *Human Factors*, 36, 619–644.

- Haxby, J.V., Hoffman, E.A., & Gobbini, M.I. (2002). Human neural systems for face recognition and social communication. *Biological Psychiatry*, *51*, 59–67.
- Johnson, M.K., & Raye, C.L. (1981). Reality monitoring. *Psychological Review*, *88*, 67–85.
- Kosslyn, S.M. (2001). Visual consciousness. In P. Grossenbacher (Ed.), *Finding consciousness in the brain* (pp. 79–103). Amsterdam: John Benjamins.
- Kosslyn, S.M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S.M. (1985). Graphics and human information processing: a review of five books. *Journal of the American Statistical Association*, *80*, 499–512.
- Kosslyn, S.M. (1994a). *Elements of graph design*. New York: Freeman.
- Kosslyn, S.M. (1994b). *Image and brain*. Cambridge, MA: MIT Press.
- Kosslyn, S.M., & Chabris, C.F. (1993). The mind is not a camera, the brain is not a VCR: Some psychological guidelines for designing charts and graphs. *Aldus Magazine*, September/October, 35–38.
- Kosslyn, S.M., & Chabris, C.F. (1992). Minding information graphics. *Folio*, February, 69–71.
- Kosslyn, S.M., Chabris, C.F., & Hamilton, S.E. (1990). Designing for the mind: Five psychological principles of articulate graphics. *Multimedia Review*, *1*, 23–29.
- Kosslyn, S.M., Ganis, G., & Thompson, W. L. (2001). Neural foundations of imagery. *Nature Reviews Neuroscience*, *2*, 635–642.
- Kosslyn, S.M., & Thompson, W.L. (2003). When is early visual cortex activated during visual mental imagery? *Psychological Bulletin*, *129*, 723–746.
- Kosslyn, S.M., Thompson, W.L., & Alpert, N.M. (1997). Neural systems shared by visual imagery and visual perception: A positron emission tomography study. *Neuroimage*, *6*, 320–334.
- Levie, W.H., & Lentz, R. (1982). Effects of text illustrations: A review of research. *ECTJ*, *30*, 195–232.
- Levin, D.T., Simons, D.J., Angelone, B.L., & Chabris, C.F. (2002). Memory for centrally attended changing objects in an incidental real-world change detection paradigm. *British Journal of Psychology*, *93*, 289–302.
- Miseses, J. (1940). Psychology and the art of chess. *Chess*, April, 154–156.
- Perky, C.W. (1910) An experimental study of imagination. *American Journal of Psychology*, *21*, 422–452.
- Rakover, S.S. (2002). Featural vs. configurational information in faces: A conceptual and empirical analysis. *British Journal of Psychology*, *93*, 1–30.
- Rhodes, G., Brennan, S., & Carey, S. (1987). Identification and ratings of caricatures: Implications for mental representations of faces. *Cognitive Psychology*, *19*, 473–497.
- Rhodes, G., & McLean, I.G. (1990). Distinctiveness and expertise effects with homogeneous stimuli: Towards a model of configural coding. *Perception*, *19*, 773–794.
- Sageman, M. (2004). *Understanding terror networks*. Philadelphia: University of Pennsylvania Press.
- Scholl, B.J. (2001). Objects and attention: The state of the art. *Cognition*, *80*, 1–46.
- Segal, S.J., & Fusella, V. (1970). Influence of imaged pictures and sounds on detection of visual and auditory signals. *Journal of Experimental Psychology*, *83*, 458–464.
- Shepard, R.N., & Cooper, L.R. (1982). *Mental images and their transformations*. Cambridge, MA: MIT Press.
- Simons, D.J., & Levin, D.T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review*, *5*, 644–649.
- Snyder, A.W., & Thomas, M. (1997). Autistic artists give clues to cognition. *Perception*, *26*, 93–96.
- Steiner, G. (1974). *Fields of force: Fischer & Spassky at Reykjavik*. New York: Viking.

- Taine, H-A. (1875). *On intelligence (Vol. I)*. New York: Henry Holt. (“Translated from the French by T.D. Haye and revised with additions by the author.” First published as *De l’intelligence*, Hachette, Paris, 1870. Reprinted by University Publications of America, Washington, 1977.)
- Tanaka, J.W., & Farah, M.J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, 46, 225–245.
- Thompson, W.L. & Kosslyn, S.M. (2000). Neural systems activated during visual mental imagery: A review and meta-analyses. In: A. W. Toga and J. C. Mazziotta (Eds.), *Brain mapping II: The systems* (pp. 535–560). San Diego: Academic Press.
- Tufte, E.R. (1983). *The visual display of quantitative information*. Cheshire, CT: Graphics Press.