

Infomages: Embedding Data into Thematic Images

D. Coelho¹ and K. Mueller¹

¹Computer Science Department, Stony Brook University, NY, United States

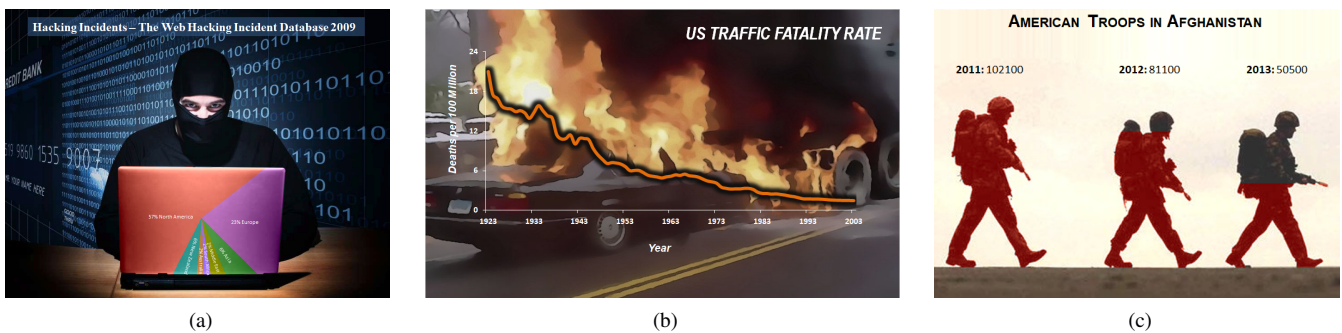


Figure 1: Infomages created with our design tool. (a) A pie chart showing the number of hacking incidents by region embedded into an image of a hacker using a laptop (b) A line chart showing the traffic fatality rate in the US embedded into an image of a burning car crash where the structure of the fire supports the trend of the line and (c) A bar chart showing the number of US troops in Afghanistan over three years embedded into an image of three soldiers walking.

Abstract

Recent studies have indicated that visually embellished charts such as infographics have the ability to engage viewers and positively affect memorability. Fueled by these findings, researchers have proposed a variety of infographic design tools. However, these tools do not cover the entire design space. In this work, we identify a subset of infographics that we call infomages. Infomages are casual visuals of data in which a data chart is embedded into a thematic image such that the content of the image reflects the subject and the designer's interpretation of the data. Creating an effective infomage, however, can require a fair amount of design expertise and is thus out of reach for most people. In order to also afford non-artists with the means to design convincing infomages, we first study the principled design of existing infomages and identify a set of key chart embedding techniques. Informed by these findings we build a design tool that links web-scale image search with a set of interactive image processing tools to empower novice users with the ability to design a wide variety of infomages. As the embedding process might introduce some amount of visual distortion of the data our tool also aids users to gauge the amount of this distortion, if any. We experimentally demonstrate the usability of our tool and conclude with a discussion of infomages and our design tool.

CCS Concepts

• **Human-centered computing** → Information visualization; Visualization systems and tools; • **Applied computing** → Media arts; • **Computing methodologies** → Image manipulation;

1. Introduction

Infographics have recently gained mainstream popularity. They are widely employed in newsrooms, businesses, classrooms, and social media primarily due to their ability to gain a viewer's attention. Unlike traditional visual data representations, infographics contain visual elements that do not represent data. These elements are mostly human recognizable objects that are meant to embellish the data graphic. While it has been a longstanding rule that

visual representations should only show data and not contain any distractors [Tuf83], recent studies have shown that visually embellished charts have a high capacity of piquing interest [BARM*12] [HRC15] [HAS11] and also improving memorability [BMG*10] [BVB*13] [BBK*16].

An inherent shortcoming with these enriched charts, however, is that their creation requires the skill of a designer. While novice users have some level of creativity, they lack the skills to create

sophisticated embellishments. This deprives them of the opportunity to take part in these eye-catching data sharing activities, reducing them to consumers only. However, nowadays many people generate data of some sort, such as fitness data, game scores, etc. or do own personal research on the internet that yields personal perspectives of some data they would like to share. Researchers have worked towards addressing this issue by proposing various design tools [KSL*17] [WZH*18] [XHRC*18] [WSZ*20] [CWW*20]. While these tools are effective at creating certain infographics, they do not cover the entire design space and at times still require some level of design expertise.

In the work subject of this paper, we focus on developing a design tool for an unexplored category of infographics that we call *infomages* or *information images*. Infomages essentially consist of a commonly used data chart (e.g. line, bar, pie etc) that is embedded into a thematic image whose content is related to the subject of the data and support the designer's interpretation of the data. For example, consider figure 1a which has been created with our tool. Here the data – the number of hacking incidents by region – is visualized as a pie chart, embedded into an image of a hacker using a laptop by filling the laptop with the pie chart. By using a topically related image containing human recognizable objects, the graphic can reinforce the subject of the data while at the same improve viewer memorability and engagement [BAW16] [BVB*13] [HRC15]. This makes infomages beneficial in situations where one would like to communicate data to a broad audience. In such situations, the data might be of interest to the audience but they might be too disengaged to look at the data in a standard format. While such charts can be beneficial in terms of their ability to engage an audience, it should also be noted that at times they can potentially distort the data they present. Additionally, their visual contents may distract users and draw their attention away from finer data details. Thus as E. Tufte advocates [Tuf83], when data is of the utmost importance such graphics should be avoided.

In the research leading to our tool we first studied existing infomages we collected from the web. We identified the different types of data charts they employ and the methods they use to embed these charts into images. We found that while some of the embedding techniques can be achieved with existing tools, others are not easily achievable and at times even require a scene to be physically created and photographed. Furthermore, some of the existing tools also require a certain level of design knowledge which is absent in most people. To bridge this gap we began by formulating a set of goals for a tool that would enable novice users to (1) locate suitable thematic images, (2) embed the charts into these images, and (3) create artistically inspired infomages.

In the following, we describe how the design tool that emerged from our activities accomplishes these goals via a set of semi-automated techniques inspired by methodologies employed by artists. These techniques allow users to integrate data charts into images while also informing them of the chart distortion potentially caused by this process. We evaluate the usability of our design tool and report on the feedback we received from a qualitative user study we conducted. We conclude with a discussion about the possible benefits and disadvantages of infomages.

Our main contributions are summarized as follows:

- We survey a set of existing infomage-like graphics to identify design practices involved in the creation of infomages.
- We propose an infomage design tool for novice users that helps them to retrieve contextual images and semi-automatically embed data charts into these images.
- We propose methods to gauge the level of distortion in these infomages as well as methods to minimize the distortion.
- We conduct a user study to evaluate our design tool.

2. Related Work

Traditionally the information visualization community has been aligned with E. Tufte's view that visual representations should only contain data related ink. Tufte opines that visually embellished charts are distracting to users [Tuf90]. In his work [Tuf83], he states that all ink that is not used to represent data should be removed from the chart. He believes that imagery and visual elements in the chart would distract users while they read the data and make the chart difficult to read. Tufte's views have been confirmed by studies such as that by Cleveland et al [CM84] which investigated how users perceive graphs as well as by Kosslyn's study [Kos89] which researched how users read charts and graphs. In his work, Kosslyn states that visual elements such as the background could be used to reinforce the chart's information. But he also advised that a poorly chosen background could interfere with the chart's readability.

While Tufte's view that a clear chart is easily readable and allows the user to completely focus on the data is unopposed, researchers do believe that at times additional visual elements could benefit the chart (see e.g., Moere et al [MTW*12]). Nigel Holmes [Hol84], a well-known designer, advocated that embellished charts will pique a viewer's interest in the chart, and more recently, multiple studies have shown that visual embellishments do have some benefits. Another visualization expert, Stephen Few [FE11], has criticized some of these studies but ultimately states that visual embellishments do engage a user's interest and make the chart more memorable, but they should only be used when they do not lead to distractions or misinterpretations. Thus based on findings in these studies, researchers in the field have been inspired to create tools that aid in the design of appropriately embellished charts. We discuss these studies and tools below.

2.1. The Effect of Visual Embellishments

In recent years, a number of researchers have been studying the impact of visual embellishments on a visual representation's appeal and incentive for engagement. An early study by Cathwon et al. [CM07] found that visually appealing charts were able to better engage viewers. Another study by Viegas et al. [VW07] observed that artistic visualizations have the ability to change the consumer's way of thinking. Similar observations were made by Borgo et al. [BARM*12] and Hullman et al. [HAS11] who found that visual embellishments pique a viewer's interest in charts while also improving their memorability. Recently Harrison et al. [HRC15] studied the first impression of infographics on people. They concluded that colorful infographics that contain limited text and image areas tend to appeal to most people.

Bateman et al. [BMG*10] conducted one of the first studies that

tested the usefulness of chart junk or embellished charts. The study concluded that while embellished charts and plain graphs were equally comprehended by users, the embellished charts were more memorable. Building on Bateman's work, Borkin et al. [BVB*13] investigated which aspects of a chart make it memorable. They observed that properties such as low data-ink ratio, human recognizable objects, and high visual density tend to make charts more memorable. Borkin et al. [BBK*16] then extended their work allowing users to view the charts for a longer period while tracking their eyes to understand what components of the charts attracted the user's attention. Their conclusion reinforced the memorability results of their previous study and showed that the presence of human recognizable objects as well as text (titles, labels, etc.) had a major impact on memorability. Another extension to Bateman's study was that by Li et al. [LM14] who had similar observations regarding comprehension and memorability. Hullman et al. [HAS11] also studied how designers could introduce 'visual difficulties' to help user comprehension and recall as well as explore the trade-offs involved. Borgo et al. [BARM*12] showed the positive impact of visual embellishments on working and long term memory. Haroz et al. [HKF15] showed that pictographs can improve memorability and recall when designed appropriately. These studies and their findings have encouraged the use of embellished charts and served as inspiration for our work. Specifically, we aim to design a tool that assists users in creating embellished charts while minimizing the distortions and distractions the embellishments may cause.

2.2. Tools for Creating Embellished Charts

Just as the studies discussed above served as an inspiration for our tool, they have also inspired other tools. Many of these have used different approaches to create visual representations with a low data-to-ink ratio. Some of these approaches are relatively simple and easy to use while others are more involved and require the expertise of a designer.

The most prevalent embellished visualizations mainly consist of an assortment of data charts and non-data related visual elements placed on a canvas. Tools such as Piktochart [pik] allow users to create such charts. They employ the easy to use drag and drop technique that lets users place visual elements, graphs, and text onto a canvas. They also offer various templates to users that assist with the infographic creation process. Very recently researchers have designed systems to automate the design of such infographics. Chen et al. [WSZ*20] created DataShot which learns data facts and generates multiple fact sheet infographics for a user to choose from. Cui et al. [CZW*20] use natural language understanding techniques to create infographics that communicate simple proportion related statistics. Both of these tools rely on pre-defined templates and layouts to produce a set of infographics that a user can pick from and modify. A more automatic approach was proposed by Chen et al. [CWW*20] who applied machine learning techniques to demonstrate how the style of timeline infographics can be extracted and applied to new data.

Researchers have also suggested more involved techniques to create embellished visualizations, specifically the data-driven transformation and the replication of visual elements. Brosz et al. [BNP*13] used graphic transforms on charts for the purpose of ca-

sual information visualization. More recently Kim et al. [KSL*17] presented a technique which allows users to design graphical objects and bind data to them. The object's dimensions are then accurately deformed based on the data. Bigelow et al. [BDFM17] proposed bridges between existing tools to ease the current process followed by designers. For example, D3 could be used to create a chart, to be then beautified by Illustrator. However, most of these tools target skilled designers. On the other hand, recent tools conceived in research, such as InfoNice [WZH*18] and DataInk [XHRC*18], do not presume this type of expertise. Both of these tools create pictograph style graphs and both replicate, fill, and color graphical objects. InfoNice creates more rigidly aligned graphs while DataInk allows the user to use touch input to replicate the graphical objects in an arbitrary direction. A similar automatic approach was put forth by Setlur et al. [SM14] who represent categorical data values with icons retrieved from the web.

The work closest to ours is the Graphoto framework by Park et al. [PKM18]. Their approach, however, only integrates line charts into photographic images and the images themselves are not chosen to go along with the theme of the data. Further, their approach distorts or invents features in the image to match the shape of the embedded line chart which can lead to somewhat bizarre feature formations.

In our work, we focus on the creation of infomages. Unlike the tools mentioned above, we do not generate and transform graphical objects or icons to represent data. Rather, we use an image's natural structure to encode the data without making structural changes to the image itself. We make use of areas in an image to encode data or we use the orientation of objects in an image to reinforce trends in a data chart. Additionally, we provide users with a guide that informs them of the extent to which a chart is distorted by the merging method.

3. Preliminary Survey of Infomages

Having identified infomages as a subset of infographics and choosing to focus on their creation, we conducted a formative study. The goal of the study was to understand how existing infomages were created. We first collected a set of 224 infomages which were then coded by both authors. Each infomage was coded based on the type of visual encoding used, the chart embedding technique, and if the background image was thematic. Conflicts between the coders were resolved through verbal discussions.

Our collection of infomages was sourced through Google images and Pinterest. These are image search engines that efficiently retrieve and rank relevant images given an appropriate search term. The search term we employed - "photographic infographic" - proved to be the most suitable to retrieve infomage-like graphics. It proved to be more effective and encompassed results returned by other search terms we experimented with such as "Chart Junk", "Image based infographics", "Picture based infographics" and more. The initial set of results returned did not contain more than 30 relevant candidates per source. To collect more candidate graphics, we made use of the "related images" feature of the search engines. For each result in the initial set, we retrieved all related images as well as images related to these related images. This resulted in a set of 1,638 candidate graphics retrieved from the web.




Figure 2: Examples of graphics from our formative study. Peter Ørntoft's [pet] award winning Infographics shown in (a) and (b) use the fill technique to encode data in a pie chart and bar chart, respectively. Oliver Uberti's [CU14] infographic shown in (c) uses a transform to deform an image of smoke to match an overlaid line chart. Images courtesy of Peter Ørntoft and Oliver Uberti.

In addition to infomage-like graphics, this set of retrieved candidates contained other irrelevant items such as regular icon and chart based infographics and other images that did not communicate any data. We removed all irrelevant candidates and duplicates. A collected candidate was considered a valid infomage if it (1) showed some data, (2) used an image (photographed or computer generated but no icons) of a human recognizable object to encode the data or aspect (such as trend) of the data, and (3) was a single panel graphic (i.e. it only showed one data visual). Multi-panel graphics were broken into individual single-panel graphics and then re-validated. It should be noted that while we stipulated that infomages should contain thematic images, we relaxed this condition for the purpose of this study. This allowed us to collect more graphics from which we could study different chart embedding techniques. This process resulted in a set of 224 valid infomage-like graphics. Additionally, we used the reverse look-up functionality of Google Images to identify the creators of these graphics. We found that a majority (188 of 224) of them were created by professional designers or commissioned by newspapers or magazines. Some examples of these valid graphics are shown in figure 2, while the others can be viewed in the supplementary material along with a list of the designers.

Next the coders reviewed each graphic and categorized it based on the type of chart and the method used to incorporate the chart into the scene. We first categorized the graphics into seven chart types described as follows:

Bar (41.1%; 92): Graphics that used multiple similar objects arranged in a vertical or horizontal sequence and whose lengths or areas were used to encode data were considered bar or column charts embedded into an image.

Pie (21.4%; 48): Graphics that radially divided an object or used radial cutouts of different objects in an image to represent data values were considered pie charts embedded into images.

Line (9.4%; 21): Graphics that showed a boundary or component of an object as a line chart or used an image object to match the general trend of a line were considered line charts embedded into images.

Scatterplot (2.2%; 5): Graphics that used multiple image objects to represent points in a scatterplot or used an image object as the canvas for a scatterplot were considered scatterplots embedded into images.

Single Divided Object: (12.5%; 28): These graphics divide a

single image object into multiple parts to represent data values. The division typically occurs along a vertical or horizontal axis of the object and can be based on area or length. These graphics are akin to single bars divided to represent data values.

Multiple Resized Object (11.2%; 25): These graphics are akin to pictograms; they use the size or count of multiple image objects that have been arranged in a scene to represent data values. It should be noted that if these objects are arranged in a horizontal or vertical sequence, they are classified as bar charts unless objects were circular and representative of bubble-like charts.

Others (2.2%; 5): Graphics for which only one chart type was identified were classified as others. For example we found only one contour map based graphic and one treemap based graphic; we thus placed them into the *others* category.

From this categorization we see that the *bar* and *pie* graphics were the most popular followed by the *line*, *single-divided* and *multiple resized objects* type graphics. The *scatterplots* and the *other* graphics were the least popular. After categorizing the graphics based on their chart type, we studied the embedding techniques used to create them. We identified four different embedding techniques one or more of which could be used to achieve the final graphic. These are described as follows:

Fill (29.5%; 66): In the fill technique, designers select an object in the scene, divide it into segments and fill these segments with a color or texture. This technique can be used to partially fill objects such that they represent bars for bar charts, or divide objects radially and fill them to represent pie charts and they can also be used to divide and fill a single object.

Overlay (15.6%; 35): The overlay technique is straightforward. In this case a data chart is overlaid on an image such that the shape or the edges of the image objects resemble the shape of the chart or the data trend in the chart. This technique is often used to overlay bar and line charts such that the orientation of image objects matches the trend of the bars or line.

Replicate and Transform (49.6%; 111): The replicate and transform operation is the most widely used technique when generating almost any type of data graphic. In this technique an object may be replicated and then some transform is applied to the object such that it represents a data value. The transform can be as simple as resizing or re-positioning objects as in the case of scatterplot graphics or multiple resized object graphics. Or they

can be complex transforms such as morphing the shape of an object to match a chart's shape or exploding an object into pieces that represent data values.

Cutout (7.1%; 16): Like the overlay technique, the cutout technique is straightforward. Here a single image is 'cut out' based on the shape of the chart with the contents of the image that overlap with chart objects (bars, pie slices) being kept and the rest of the image being discarded. This essentially makes the chart a stencil where chart objects are filled with image fragments instead of colors or textures.

The results of this categorization is shown in table 1. From the results we can see that the *fill* and *replicate and transform* operations were the most used techniques while the *overlay* and *cutout* techniques were rarely used. It should be noted that in rare cases a graphic could be assigned to one or more embedding techniques. For example, in the case of scatterplots image objects were *replicated and transformed* to represent points while these objects were *overlaid* on a canvas object that provided context to the subject of the image. In addition to categorizing graphics based on the chart type and embedding technique we also studied if the image or image objects were thematically related to the data. Thematic images refer to images that contain objects that are topically or semantically related to the data. To explain with an example, consider figure 2c, here the chart shows the number of people who quit smoking along with a background image of smoke, we consider this background image to be thematic. However, if an image of flowing water (instead of smoke) was used instead, we would not consider it to be thematic. We found that 52.6% (118) of the graphics used thematic or topic-related images while the other 47.4% (106) graphics seemed to use images that were not thematically related but contain easy-to-embed objects.

Table 1: This table shows the number of infomage-like graphics categorized by chart type and embedding technique categories.

Charts	Embedding Type			
	Fill	Overlay	Cutout	R & T
Bar	17	10	8	57
Pie	27	6	1	16
Line	1	12	7	1
Scatterplot	0	2	0	5
Single Divided Object	18	0	0	10
Multiple Resized Object	3	0	0	22
Others	0	5	0	0

4. The Infomages Design Tool

The infomage-like graphics collected in our formative study were primarily created by professional designers. From their designs, we observed that the creation of these data graphics requires a substantial amount of time, effort and skill. This puts the creation of infomages out of reach for most people. Thus we set out to create a design tool for novices to create infomages.

4.1. Design Goals

Based on observation made during our formative study we identified a set of design goals for our design tool listed as follows:

G1. Provide easy-to-use data embedding techniques. Designers who created the infomage-like graphics in our formative study often setup a physical scene to represent data and then photographed it to create the data graphic. They used one or more of the embedding techniques described above to physically place, color or transform objects in the real-world to represent data values. This is a time consuming and laborious process that also requires a fair amount of creativity. Our first goal is to provide a computer-based technique that allows users to easily embed data into images. The design tool should allow the user to embed a visual representation of the data into a selected image using the techniques discussed in the formative study. For the purpose of this work we chose to focus on the *fill* and *overlay* style embedding techniques. We chose to exclude the *replicate and transform* technique as it has been well studied and implemented in tools such as InfoNice [WZH*18], DDG [KSL*17] and DataInk [XHRC*18]. While these tools replicate and transform icons or objects that have been drawn within the tool, they can be extended to work with objects segmented out from images. We also exclude the *cutout* technique as it is straightforward and does not require design expertise.

G2. Inform users of possible distortions. While most infomage-like graphics in our formative study are attractive, at times they distort the data being communicated [SHK15] [SK16] [PRS*15]. This is primarily due to the embedding techniques being incapable of accurately embedding the data. While it is difficult, if not impossible, to remove these distortions, our tool should at the very least inform users of their existence. Thus the next goal of our tool is to provide users with an estimate of the amount of distortion introduced in an infomage data representation. Furthermore, we attempt to minimize this distortion.

G3. Provide techniques that ensure the data encoding stands out. The infomage-like graphics we studied in section 3 are visually dense and contain multiple objects some of which do not encode data. Additionally, these objects may appear over backgrounds with complex textures. In such images it is important to ensure that the objects representing data are clearly visible and stand out from the rest of the scene. We observed that artists used bright and colorful objects to represent data while the scene mostly consisted of a plain or uncluttered background and at times had dark-colored objects. Additionally, we found graphics in which a transparency or blur was applied to the background scene while keeping the data objects in focus. We also found graphics in which the objects representing data did not really stand out. In such cases, we concluded that these objects could have been more pronounced if they were brightly colored or highlighted with a glow. Thus, our third goal was to provide the user with similar techniques to make the data stand out. Our tool should allow users to color or highlight data items in the image while also being able to subdue the effects of the details in the background image.

G4. Ensure accessibility for novices. Most infomage-like graphics we studied were created by designers who were proficient photographers or were experts in design softwares (e.g. Adobe Photoshop, Adobe Illustrator). Our final goal was to ensure that novices




Figure 3: The workflow of our design tool. First, the user loads a data file. Next, the title and attribute labels of the data are extracted and used to retrieve contextual images. The user selects one of these images and moves on to embed a chart into the image. After embedding the chart, the tool computes the amount of distortion introduced to the chart. The user can choose to reduce this distortion with the help of our tool. Once the user is satisfied, he or she can apply some visual enhancements to the graphic. Finally, the resulting infomage is published.

or users with minimal design experience could use our tool to create infomages. Our tool should allow users to use a small number of simple point, drag, and click mouse functions to create infomages. Additionally, a novice user may lack the imagination to think of a thematic scene for the infomage. In such cases, our tool should provide users with suggestions for a background image or that spark the user's imagination to create a new image for the infomage.

4.2. Design Tool Overview

Based on the design goals above we created an infomage design tool for users who lack design expertise. The tool inspires users in the design process by immersing them into a large set of images which for the most part are all contextual to the data. It provides users with tools to tightly integrate the data into structural elements of the images while keeping them informed of the chart distortion introduced by the integration process.

The workflow of our tool is shown in figure 3. To create an infomage the user must first load a data file into the design tool. The data must either be a univariate or bivariate and if the data is categorical it is reduced to the counts of each categorical value. Next, a set of relevant backdrops for the infomage is retrieved using a web-scale text-based image search engine (addressing G4). The images returned are most often, but not always, related to the descriptors. Thus we rely on the user to pick a relevant image or load their own image file which might be inspired by the retrieved images. The user must then select the type of embedding process and the chart (pie, bar, or line) based on which the data will be embedded into the image. Having set the chart type and background image for the infomage, the user can now proceed to design the infomage. The first step is to integrate the chart into the image (addressing G1). Based on the embedding process and chart type, we support this integration task with a set of techniques described in detail in section 4.3. We also inform the user of the distortion caused by the integration process (addressing G2). Once the chart has been integrated into the image, the user is free to apply a diverse set of visual effects (grayscale, cartooning, glow, transparency, etc.) to the chart as well as to the backdrop. The purpose is (1) to give it an illustrative look and feel, and (2) to reduce image detail. The illustrative effects set the data apart from the image and so reduce possible interference of the infomage's image content with its data content. This addresses the aforementioned concerns that embellishments are too dominant and distracting from the data (addressing G3). Once completed, the infomage can be exported as a PNG file to be published. The entire

process is supported by simple mouse interactions thereby minimizing the effort required from the user (addressing G4).

4.3. Merging Data and Images

The crucial step in the creation of infomages is the process of integrating the data chart into the image. As described in our first design goal - G1, we identified two main techniques for this purpose: (1) filling a suitable object in the image with the chart or chart element, or (2) aligning and overlaying the chart on top of the image. We describe each method in detail below.

4.3.1. Fill Technique

A portion of the artist created graphics that we studied, such as those in figures 2a and 2b, use the fill technique to embed data into image objects. Designers follow two main approaches to embed data with this technique. In the first approach a single object in the image is selected, then it is divided into multiple parts and these parts are filled with colors or textures to represent data values. This approach is mostly used to embed pie charts or areas divided along an axis into images. The second approach requires the designer to select multiple objects and then partially fill each object with a color or texture. Here each object and the extent to which it is filled represents a single data value. This approach is often used to embed bar or column charts into images. Similarly, our tool must allow users select image regions and use one of these two approaches to encode data into them.

First the user should be able to select image regions. To do this programmatically, we made use of a user-guided segmentation algorithm, the well-established Grabcut method [RKB04]. We also considered other newer learning-based and superpixel-based segmentation methods, but in practice, these methods only create a rough mask that still needs to be refined by the user using grabcut. Since the effort required by the user is approximately the same or more in case of these methods, we chose to use grabcut exclusively. Figure 4 illustrates how we use Grabcut to extract and fill areas in an image. For it to work, the user must first draw a box around the desired region. In most cases, this is sufficient for the algorithm to detect the desired object. However, at times part of the object is not extracted or unwanted objects or pieces are extracted. To refine the selection, users can scribble some foreground pixels and some background pixels (figure 4a). Grabcut uses these pixels as examples to separate the foreground from the background creating a refined mask (figure 4b). The user goes on marking pixels until a




Figure 4: Designing an infomage with the fill based technique: (a) The user roughly selects the region into which the chart should be placed; (b) The segmentation algorithm creates a mask for the chart; (c) The mask is divided and filled based on a pie chart and merged with the image; (d) The labels are added and the user puts on any finishing touches (e.g. background filters, font changes etc.)

satisfactory result is reached. Finally, with the region identified, the user can proceed with laying out the chart into the image.

The next step is to allow the user to divide and fill in the segmented regions. As described by Cleveland [CM84] and Skau et al. [SK16] [SHK15], people primarily read data values represented by visual objects based on their length, area, or angles in the case of radial charts. Thus we also allow the user to select one of these encodings based on which a select region will be divided and filled proportional to the data. For the approach of dividing a single selected region, we allow the user to divide the region along a vertical axis, horizontal axis or radially. In the case of division along an axis, we allow the user to select from division based on length or area of the region. But in the case of radial division we allow the user to select from division based on angle, area or arc length. In case of the area-based and arc-based radial division it should be noted that we perform the division on the convex hull of the region rather than the actual region. We choose to use the convex hull to avoid misrepresentations of the area due to irregularities. For example, irregularities such as thin “cracks” along the edge of a region may be ignored by a user when visually estimating its area. The use of a convex hull removes such irregularities. For the approach of filling multiple selected regions, we allow the user to fill the regions along a vertical or horizontal axis. Here we allow the regions to be filled based on their areas or lengths. Using area is straightforward; we partially fill each region’s area based on the data it represents. When using length a bounding rectangle is drawn around all regions. The bottom-most and top-most or left-most and right-most edges (depending on the orientation) among the rectangles are used to represent the minimum and maximum values, respectively. Using the bounding rectangle avoids errors with minor vertical misalignment of the regions.

Once the data is encoded into the region, we color its divisions in CIE LAB color space. We transform the A and B channels to the new color while leaving the L channel unchanged. This retains the texture of the region but changes the color as shown in figure 4c. Users can select a color of their choice or use a color palette from Color Brewer [HB03]. Lastly, the labels are added as shown in figure 4d. The user is allowed to interactively change axis limits or start angles (for radial division) as well as refine label positions.

4.3.2. Overlay Technique

The overlay technique is another method that was used by artists and was widely used to embed line and bar charts in the graphics we studied. This technique is much less involved than the fill technique. We observed that in most cases, artists made an attempt to align the chart with some objects in the image. The alignment was such that either the trend of the data matched the orientation of an image object (in case of line or bar charts) or the chart was just placed over or based on an object contextual to the data. Similarly, our design tool should assist the user in finding the most suitable position in an image for a chart.

To achieve such functionality, we reduced it to the problem of searching for image objects whose outlines or orientations match the trends or outlines of a data chart. As the bar, line, and pie were the most frequently used charts with the embedding technique, we focus on finding objects that match these types of charts. We treat line and bar similarly; the positions of their graphical elements (line joints and bar tops) representing data form a trend, and thus we search for image objects that match this trend. Pie charts are simpler; they are essentially circles and thus we search for circular regions in the image where the pie chart can be placed.

We attempted to automate the shape matching process by analyzing the edge maps. A promising method involved using the ant colony algorithm to evaluate edges (see supplementary material for details), however this method failed very often. While a perfect automatic method could not be identified for the general case, we found that the Hough Transform [DH72] did a reasonable job at estimating an approximate trend or shape. Figure 5 illustrates how we use the Hough transform to overlay a trend-based chart in an image. First we extract Hough lines in the image. An example is shown in figure 5a. Following, we compute the linear regression line of the data values as in figure 5b and find the closest Hough line to it in terms of slope. The chart is then positioned along the matching Hough line. Similarly, we can search an image for circular areas using Hough circles and place the pie chart in the largest circle. It should be noted that finding structures in pre-existing images that match the exact positions of these values is very unlikely if not impossible and thus any method will fail for images that do not contain appropriate structures.




Figure 5: Designing an infomage with the overlay technique: (a) Given the background image the Hough lines are extracted (shown in red); (b) The linear regression line is computed based on the data provided; (c) The chart is placed over the image based on the best matching Hough Line; (d) The user repositions the line chart in the image.

This approach usually only generates a rough positioning for the chart along an image structure. For example see figure 5c. We thus rely on the user to finely reposition the chart to his or her satisfaction. For this purpose we provided an interactive interface that allows users to move and scale charts to and in appropriate positions in the image. We also chose to prevent the user from rotating trend based charts as this may lead to trends looking more dramatic than they really are. Figure 5d shows the final infomage designed by the user using this method. Yet, other users might have preferred the initial configuration of figure 5c, or other translations or scalings of the chart. Our tool can support these personal preferences by simple mouse interactions.

4.4. Estimating Distortion

Merging data and images using the techniques discussed above can distort the data representation at times. An infomage without distortion can only be created if the image contains perfect shapes (circles or rectangles) or trends that match the data. In all other cases there is some amount of distortion. As described in G2 the user must be made aware of such distortions. Hence we provide users with certain measures of distortion that assist them in judging if their creations are too distorted. The measures and how they are computed are explained below.

4.4.1. Lengths, Areas, and Angles

The fill based charts encode data through an object's length, area or in case of radial divisions angle or arc length. Viewers may take into account some or all of these measures to judge the values they represent [CM84] [SK16] [SHK15]. As infomage designers often work with irregularly shaped objects, optimizing for one of these measures causes some level of distortion in the others. Thus we must detect and communicate these distortions to the user. We do so by first calculating the length (or arc length), area, and angles of the segmented region in the infomage. Next, for each of these values we compare their ratios to those in the original data. The differences between the measured values and the original data are represented as percent error. The percent error for each encoding, length (or arc length), area, and angles, and the average error over all measures are then reported to the user.

4.4.2. Trends

Trends are often communicated through line charts or ordinal column charts. These charts encode data through the positions of mul-

iple sequential points. By observing the slope of the line that would connect two adjacent points, a viewer learns of the type of change and its rate between data values. Distortion in such charts is often caused by truncated axes or an inappropriate aspect ratio [PRS*15]. These techniques change the slopes between sequential points and thus give the viewer a wrong impression about the trend.

In infomages, truncated axes may arise when the designer wants to exaggerate or understate a trend to align with an image object when using the overlay technique. The designer can also achieve similar alignment results by changing the charts aspect ratio. To communicate the distortion caused in such cases we use two methods. The first method targets the truncated axes distortion. Computing errors here is straightforward. We measure the distance of a point from the reference axis and then compare these values to the original data values. The differences are represented as percent error. The second method targets the aspect ratio distortion. While there is no fixed rule for choosing an aspect ratio, it has been shown that banking to 45 degrees [HA06] can help users to compare line slopes. Using the banking to 45-degree rule, we report an average orientation error which is the absolute difference between 45 degrees and the average orientation of the line. The percent errors for truncated axes and average orientation error as well as an average of these errors is reported to the user.

4.4.3. Controlling the Distortion Error

All measures of distortion for each chart discussed above are presented to the user individually as well as a combined average (see bottom right of figure 6). In addition to presenting these values, we give users the option to optimize the values for each chart.

Optimizing for the overlay charts (line and bar) is straightforward. First, to avoid truncation errors we set the minimum y-axis value (or location of the x-axis) to 0. Second, we compute the optimal aspect ratio using the 45 degree banking method [HA06]. This is computed using the formula $\alpha = M.R_x/R_y$, where α is the aspect ratio, M is the mean absolute slope of all line segments in the chart, and R_x and R_y are the range of x and y values respectively.

Conversely, optimizing for the fill charts is more complex as most often there will always be some amount of error in the encodings (for example, angle and arc length for a filled pie chart) when optimizing for one of them (in this case, area). To determine a good compromise we utilize the simulated annealing algorithm [KGV83] to minimize the average error. The algorithm is first initiated with the original scaled data values and settings, for

example, in radial division these are the angles to which the data has been scaled. Next, the average distortion error is computed (for example, the average of angle, area, and arc length errors in radial charts) and used as an initial reference error value. We then add a random small amount of change to the original scaled data values. Here the change is up to 5% of the original scaled data values or the maximum change in scaled values when optimizing for one parameter, whichever is greater. For example, if the change between scaled data values when optimizing for angle only and arc length only is greater than 5% of the original scaled value then we use this difference. We ensure that the sum of the changed values is equal to that of the original scaled data values. Next, the chart is generated and the error is recomputed. If the new error is lower, we update the reference error with this value. We keep following this process until the average change between error values is very small (0.01%) or the decay function ($T = T * 0.9$, where T is the annealing temperature initially set to 1.0) reaches a minimum value (in this case $T_{min} = 0.00001$).

The errors computed above are meant to inform the infomage designer that there is some distortion and raise awareness to the amount of error a reader might make when viewing a infomage. While the measures of perceived distortion are somewhat approximate, they can help designers weigh compromises between artistry and possible minor distortion, and also deter them from using a particularly error-prone creation. We ran a preliminary user study (see supplementary material) and found that when the average error is 15% or lower for radial fill charts (10% for bar based charts) or less than 20% for overlay charts, it becomes difficult for viewers to notice the difference between erroneous charts showing the same data. The infomage in figure 6 has an average error of just under 10% which is markedly below this critical threshold. This error is mostly due to the area error of 20%, but the angular error is a perfect zero and the arc length error is only 10%. Since all three measures are taken into account by a viewer when reading the chart [SK16], the influence of the area error will be significantly weakened. We communicate the distortion measure through gauges that are horizontal bars colored from green (low distortion percentage error) to red (high distortion percentage error) with a pin marking the distortion level (shown in figure 6 step 4). We set the transition from green to red at the 20% mark of the bars as our preliminary study shows that distortion below these levels may be acceptable. Note that the level could be higher for fill-style pie layouts and lower for the overlay layouts, however, further testing is required to attain more accurate estimates. This design proved to provide better guidance to novices than displaying a raw percentage error value. For a detailed report on the distortion the user can click the “?” button. This generates a popup that reports the differences in the individual data values with a grouped bar chart, nested donut chart or multi-line chart depending on the embedding technique.

4.5. Post Processing Charts

Most often, it is the image details that can draw a viewer’s attention away from the data, and as described in G3 our design tool should allow designers to address this issue. To reduce the details or draw attention away from them we provide the designer with a wide set of post processing functions to better expose the data. We provide

several image filters such as a grayscale filter, blurring filter, sharpening filter, and a cartooning filter. The grayscale filter converts the image to grayscale while allowing the chart to remain in color. This makes the chart stand out in the resulting infomage (see Figure 7b). The cartooning filter and blurring filter, on the other hand, flatten or lower the details in the background image thus shifting the focus to the chart (see Figure 1b).

Additionally, we provide designers with a variety of well designed color palettes sourced from Color Brewer [HB03]. However, images may have multiple regions that may have different levels of brightness and details. This makes it difficult to use a predefined color palette with the color blending employed in the fill style charts as this causes some colors to look too washed out or too bright in different regions. We address this issue by allowing users to replace colors in a color palette with custom selections. In overlay-style charts this issue cannot be solved as easily since the chart may traverse multiple regions with different brightness levels. We address this issue by allowing users to add a contrast enhancing glow around the chart element thus giving it a “popping” effect as in figure 1b. Additionally, it was shown by Borkin et al. [BBK*16], viewers always look at the chart’s text elements such as the title and axis labels. Thus we also allow designers to control the font style, font size, color, highlight color and transparency.

4.6. Interface and Implementation

Our fourth and final design goal (G4) was to provide users with an easy to use interface to employ and control the techniques discussed above. We accomplish this by implementing our design tool as a Windows based graphical user interface (GUI). Our user interface design takes cues from multiple image editing tools available today. Figure 6 shows an instance of the interface.

Our design tool’s main components are a large area to display the image and interactively reposition elements, and a toolbox to its right that allows users to control the data embedding techniques. On the bottom left is a frame that shows a table with the data and a regular chart visualizing them. The regular chart is provided to enable a comparison with the infomage and to judge whether it is good enough in terms of aesthetics and its amount of distortion. To the right of this frame is another frame containing all the images retrieved by our tool. To start creating an infomage the user can select one of these images or import an image sourced through some other method. In the bottom right corner is the list of distortion measures to help users gauge the degree to which the current infomage distorts the original chart. Finally the narrow toolbar on the top of the interface provides a set of post processing tools by which users can change the font and apply filters to the image.

The implementation was performed in C++ for its performance and vast library support. For all the image processing functionality we use OpenCV [ope]. It includes an implementations of the Hough Transform which we use in the overlay technique as well as in the Grabcut segmentation algorithm and other contour manipulation algorithms for our regions fill based techniques. OpenCV also provides the image filter functions to de-emphasize image details. To suggest images to the user we use the Google Custom Search API [goo]. Through this API, we send requests to Google Images

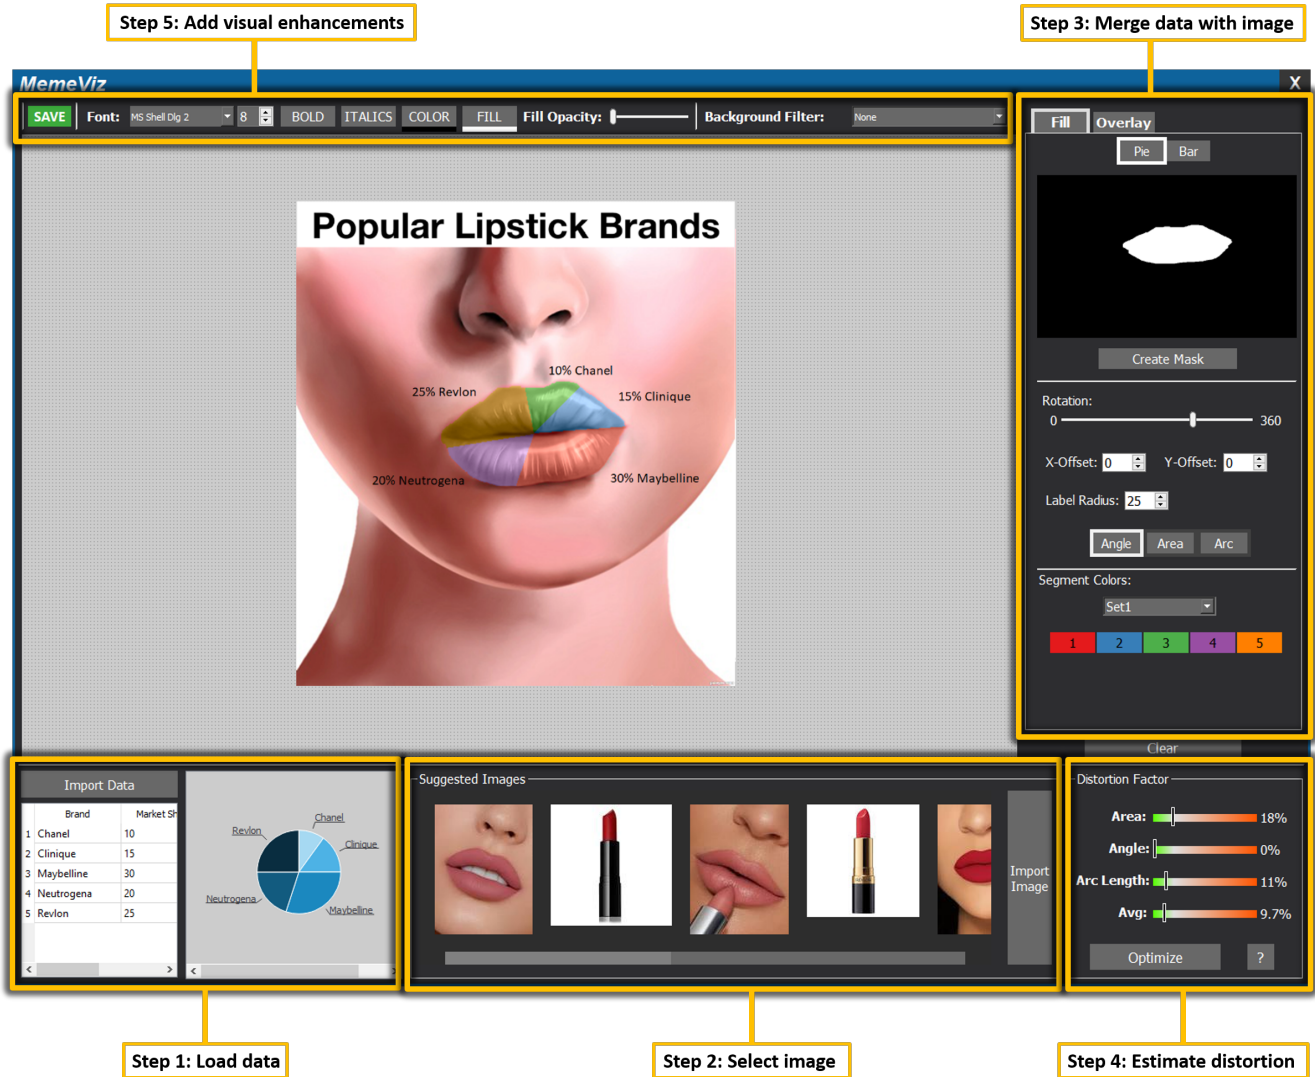


Figure 6: The design tool user interface. Areas in the figure have been marked based on the user’s workflow. At **step 1** the user loads in a data file and previews the data with a basic chart. Next, at **step 2** the user either selects one of the many contextual images retrieved from a web-based search engine or imports a custom image. Once the data is loaded and an image selected, the user moves on to **step 3** which involves employing easy-to-use tools to embed a chart into the image. Next the users moves to **step 4** where he or she uses our distortion estimator to gauge if the levels are acceptable or if it needs to be improved are necessary. Once the user is satisfied, he or she moves on to the last stage, **step 5**, where the user can add further visual enhancements to the text and background image.

with the search string generated from dataset-related information, such as header, axes labels, and other textual information that may come with it. It should be noted that the results of the query depend on the text it is provided, thus the dataset title and attribute labels must be meaningful and should not exclusively contain acronyms or obscure label names such as alphanumeric terms (for example, an attribute indicating patient id labeled as PID). Within a few seconds, the search engine returns a set of results which can be filtered by multiple parameters. One of these is the Creative Commons license which allows us to only present images labeled for reuse in the image gallery. This ensures that the infomages can be legally created and disseminated. The graphical user interface

was implemented with Qt [qt] which has immense functionality and supports all of our interaction tasks as well as charting. The source code for the tool is available at <https://github.com/darius-coelho/Infomages-Design-Tool>.

5. Usability Study

We evaluated the usability of our design tool by ways of a user study in which participants were asked to design infomages with the tool. To the best of our knowledge there is no single tool in existence that can accomplish this task. To create an infomage without our design tool one would need to pair multiple currently avail-

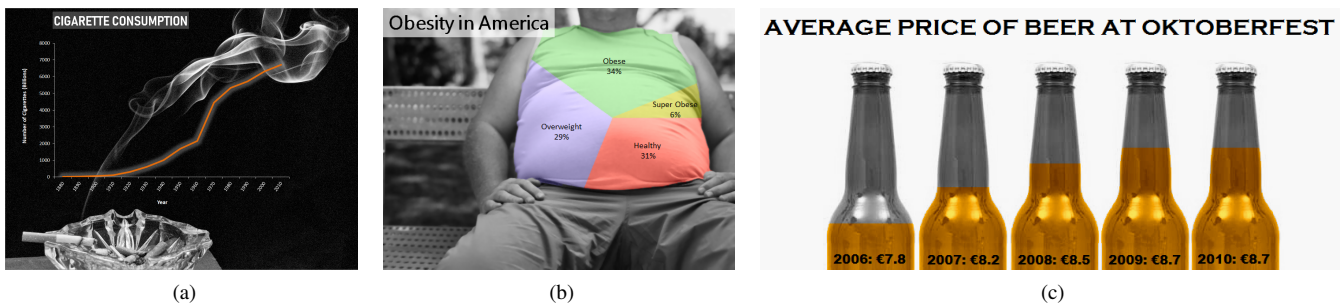


Figure 7: Infomages used for the replication task in the usability study: (a) An overlay graphic using a line chart. (b) A fill style chart with a single object radially divided and (c) a fill style chart with multiple objects partially filled.

able tools such as Adobe Photoshop, Microsoft Excel and an image search engine, which is a setup we deemed too complex to serve within a comparative study. To test our design tool we recruited five participants from our university campus. All participants were graduate students (2 economics majors, 2 computer science majors and a physics major). Four participants were male and one female, all five were in the age range of 25-34. Two of the participants (D1 and D2) had design experience with Adobe Photoshop or a similar image editing tool while the other three participants (N1, N2, and N3) had no experience with any professional design tool. However, all participants had experience creating data charts with Microsoft Excel and PowerPoint but no data visualization expertise.

5.1. Procedure

Our usability study consisted of two main tasks - replication and creation. To begin, the participants were asked to watch a video tutorial of our tool. This was followed by an interactive session in which they could explore the tool and experiment with the Lipstick dataset. This dataset has five data points as shown in figure 6. During the exploration they were free to ask the author about the functionality of any of the buttons in the interface. On average, the participants took 21 minutes to watch the tutorial video and warm up to the tool. Next, the participants proceeded to the two tasks.

We began with the replication task of the study. Here the participants were shown previously created infomages (figure 7) and were asked to re-create them with our design tool. They had to perform three replication tasks - two fill based replications and one overlay based replication. One fill method involved selecting a single object and dividing it radially and the other involved selecting multiple objects and filling them to represent a bar chart. We used only one overlay task with a line chart as the mechanics for this technique were identical to other charts. Once a participant completed the replication task he or she could move on to the creation task. For this task each participant was asked to select any dataset from a set of 17 to create two infomages. No limitations were imposed on the participants for this task; they could choose any chart type and embedding technique they preferred. They were also allowed to use the web to retrieve images if they were not satisfied with the tool's suggestions.

After replicating or creating a chart the participants were al-

lowed to take a break if they wanted to. They were also free to stop the study at any time. During the study we measured the time it took each participant to complete the various tasks. At the end of the study the participants were asked to complete a questionnaire based on the USE questionnaire [Lew95]. The responses were collected on a 5-point Likert scale with ratings from strongly agree to strongly disagree. Additionally, they were also encouraged to give us verbal feedback and suggestions.

5.2. Results

It was encouraging to see our study participants interact with our design tool. All users, even those with no design tool expertise, were able to quickly generate quality infomages with it. In the following, we report the feedback we received from our participants.

5.2.1. Functionality and Ease of Use

Overall, the participants performed the replication task quickly and with minimal help. Assistance was requested mostly for coloring the title of the chart. An average of 4:49 minutes was needed to re-create a single infomage, with the overlay-style (figure 7a) taking 5:31 minutes, the single object fill-style (figure 7b) taking 4:22 minutes, and the multiple object fill-style (figure 7c) taking 4:34 minutes. The overlay style took longer since the graph re-positioning takes more effort than the object filling and line charts had many more configurable parameters (such as line thickness, color, shadow, interpolation type). For creating a new infomage, the participants spent an average of 13:11 minutes to complete the task. All participants created fill-style charts with all five of them (D1, D2, N1, N2, N3) using the radial division of a single object; only one participant (D2) used the multi-object bar-style. Only three participants (D1, N1, N3) chose to use the overlay-style charts; the other two (D2, N2) felt that they were not interesting enough to create.

The participants said the tool was easy to learn (average rating of 4.6/5.0) as well as easy to use (average rating of 4.3/5.0). N1 stated, "This is simple to use, I don't have to do much to create the visualization." Participants also found the design operations were direct and required only very simple mouse interactions. N3 noted, "Now I can place charts in an image with just a few clicks." And experienced photoshop user D2 noted, "This is much easier than

photoshop, I don't need to spend too much time outlining the objects I want to select." The other participants agreed that the design tool was efficient at creating infomages and saved time, rating it at 4.5/5.0 on average for this category. All participants also strongly agreed (5.0/5.0) that our tool was fun to use and that they would recommend it to a friend.

However, the participants only gave an average rating of 3.8/5.0 when asked whether the tool's functions performed as expected. D2 stated, "The object selection tool is easy to use but not as accurate as photoshop." But he also agreed that "Using photoshop would be more tedious as I would have divide the area myself. Maybe I could use photoshop to refine the infomage." Some users (D1, D2, N2) also complained about the label positioning being unexpected at times. While the participants did not discuss the overlay method as much, N1 and N3 did "like that the tool suggests a position for the line chart." But D1 did note that "it is not accurate" all the time as he previewed multiple backgrounds. Additionally, participants (D1, D2, N2) "wished there were better color palettes."

5.2.2. Distortion Measures

When asked about the distortion measures, all participants stated that they "understand how it was computed and what it means." Most participants (D1, D2, N1, N2) explicitly stated that they understood that the goal was to minimize the error and "get it as close to zero as possible." However, if they were not able to lower an error value, they were "not sure what a good value is." D1 wanted to know "if there a way the tool can tell me that an error value is acceptable for most viewers." This is a difficult task due to the dynamic nature of infomages and is still being explored by the research community.

5.2.3. Image Suggestion

The participants appreciated that our tool suggests images related to the topic of the data. N3 stated, "I like the image suggestion, I dont need to look for one myself." However, most participants (D1, D2, N1, N2) stated, "picking an image is difficult." Some participants (N1, N2) said that this was because they were "not sure if this image would look better than others." But our design oriented participants (D1, D2) said, "I can definitely find better images on my own" and they could "imagine some scenes." This was illustrated by D2 who was inspired by an image suggested when using the lipstick dataset. He saw an image of a single lipstick (second image in the image suggestion panel in figure 6) and used a web browser to find an image that had multiple lipsticks lined up so that he could fill them with the bar style method. All participants believed that they could find better images if they had more time.

6. Discussion

6.1. Impact of Infomages

Researchers in the visualization field have published multiple recent studies on the benefits and shortcomings of embellished visualizations. However, none of these studies feature a large amount of infomages. Nonetheless we can expect certain findings in these studies to apply to infomages as well. Multiple studies have shown that the colorfulness and presence of human recognizable objects

make embellished visualizations engaging [BARM*12] [HRC15] [HKF15] and memorable [BMG*10] [BVB*13] [BBK*16]. Both of these properties apply to infomages as well; they tend to be colorful and often the image chosen contains human recognizable objects. These properties are the main reason for using infomages. However, the process of embedding a chart can cause distortions to the data, irrespective of the designer's intentions. Such distortions have been studied by Skau et al. who show that embellishments that change the area and shape of bar [SHK15] and pie charts [SK16] can have a negative impact on a viewer's judgement of data values. These distortions primarily affect the fill-style graphics. Additionally, distortion techniques that affect regular charts [PRS*15] can affect overlay-style graphics as well. Thus designers should use infomages only when their engagement and memorability value outweighs the need to show detailed and accurate information.

We also discussed infomages with a five member focus group that consisted of a business school professor (male), a graphic designer (female), a visualization expert (male), and two marine science researchers (both female). The focus group was shown examples of infomages created with our tool and those shown in figure 2. All group members stated that they very much liked the infomages created with our tool. However as expected, they also remarked that the images were not as refined as those created by designers. The group members confirmed our design premise that the "backdrops were contextual" in infomages created with our tool. However the remark was made that "better color palettes could have been used." All five participants echoed our belief that infomages in general can be engaging and memorable. They stated that using infomages in blogs, introduction to news articles or presentations, social media and advertising could reinforce a message and "make it stick." The business professor when examining our infomages noticed that some of them distort data and stated that he would avoid using them in "serious scenarios" where the data being communicated needs to be accurately presented. The graphic designer commented that infomages succeed in "pulling people in, using two languages, images and data." But she also remarked that "This can be tricky because each person can be affected by an image in different ways. They can send mixed messages when paired with the data visualization." In our work we did not focus on this aspect and it must be studied in the future. Existing studies in advertising, have shown that images have the power to bring out different emotions in people [LGBH93] and thus influence them in a variety of ways [MBL*91].

6.2. Infomage Design tool

Our design tool adds to the growing set of chart embellishment tools that is accessible to both experts and novices. Our work focuses on the design of infomages which to the best of our knowledge is a type of infographic that has not been studied before. The tool implements two main design techniques - the fill and overlay techniques. While these techniques have primarily been used to create pie, line and bar chart infomages by both designers and our design tool, they can easily be used to embed other charts into images. For example the fill based technique can be used to embed other area based charts such as treemaps, voronoi charts, funnel charts, and more. We would only need to change the

way in which we divide the area of selected objects. However, it should be noted that these charts are considerably more complex and using them in infomages may cause more difficulties for the viewer. Furthermore, we believe that the techniques in our design tool could be combined with other methods such as the replication technique [KSL*17] [XHRC*18] [WZH*18] and morphing technique [BNP*13] [CLC*15] to create more sophisticated infomages. For example, the fill and replication techniques could be combined to fill or pack areas with replicated and resized objects.

Our usability study showed that users were mostly positive about the tool; they stated that it was easy and enjoyable to use. However, the users did comment about ways to improve the tool. First, they stated that they would like to have improved algorithms which would reduce or eliminate the effort required to manually refine a segmentation result or chart alignment. While automatic approaches do exist, the state-of-the-art is still not accurate enough to operate without human intervention. Next, the users appreciated that our tool suggests images for their data, however some users were unsure of which image would be the most attractive. Additionally, users stated that they could find better images on the web, however, our tool does not recommend copyright protected images which tend to be much more attractive. We agree that it is difficult to pick an image and as part of future efforts we plan to develop an image ranking suite based on image aesthetics and appropriateness with respect to the chosen merging technique. Finally, as pointed out by our focus group, users still need assistance when selecting appropriate chart colors. Although we use Color Brewer [HB03] palettes, the background colors of the image may interfere with them. Thus in future efforts it is important to develop better color selection methods.

Finally, we aimed to address the issue of distortion. Our design tool estimates the amount of distortion being introduced by the embedding techniques. It was also able to bring these distortions to a designer's attention. While our distortion estimation methods are based on previous studies [CM84] [SK16] [SHK15] [HA06], these studies did not test infomages. Thus further testing is required to understand how viewers are affected by such distortions and if some distortion measures are more important in infomages. For example, viewers may ignore arc length when viewing pie style infomages but may consider it in regular pie charts or in infomages with objects closely resembling circles. We also provided a technique to automatically optimize the distortion, but users were still not sure of how much distortion was acceptable as indicated by our usability study. Due to the complex nature of images, it is challenging to define an acceptable level of distortion for the general case. We find that the initial study we conducted to evaluate the effect of distortion on readability (see supplementary material) was quite promising. We plan to expand it to better understand the specific thresholds to be used as acceptable levels of distortion.

7. Conclusion

We have presented infomages as a new type of data graphics in which an image is augmented with a data chart such that the structure of the image supports the designer's intended meaning or interpretation of the data. Through our discussion we show that infomages are suited for more casual purposes (social media and blogs)

and in situations where one needs to capture the viewers attention (presentations to a general audience). However, infomages are not suited for targeted analytics tasks as they can potentially distract and mislead viewers.

To support the proliferation of infomages we created an interactive tool for designing infomages. It enables general users who lack the professional skills of a designer to create fairly decent infomages. It also keeps the user informed of the distortions being introduced in the graphic. Merging the chart into the image and adding illustrative enhancements is straightforward with our tool, requiring just a few simple mouse interactions. Users stated that the tool was easy and enjoyable to use. However, they did point out areas for improvement such as more accurate segmentation and alignment techniques and better image recommendations. They also stated that having a predefined threshold for the amount of distortion would be beneficial.

As a future effort we intend to improve our design tool by creating a more sophisticated image retrieval module that includes structure analysis, implementing better matching algorithms, and adding more visual customization to the chart. We would also like to derive more perception-based bounds for the various distortion errors. Finally, while we have discussed the expected effects of infomages, deeper studies need to be conducted to confirm and test the extent of these effects.

8. Acknowledgments

This research was partially supported by NSF grant IIS 1527200 and by the MSIT (Ministry of Science and ICT), Korea, under the SUNY Korea's ICT Consilience Creative program (IITP-2020-2011-1-00783) supervised by the IITP (Institute for Information communications Technology Planning & Evaluation).

References

- [BARM*12] BORGIO R., ABDUL-RAHMAN A., MOHAMED F., GRANT P., REPPA I., FLORIDI L., CHEN M.: An empirical study on using visual embellishments in visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec 2012), 2759–2768. 1, 2, 3, 12
- [BAW16] BYRNE L., ANGUS D., WILES J.: Acquired codes of meaning in data visualization and infographics: Beyond perceptual primitives. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan 2016), 509–518. 2
- [BBK*16] BORKIN M., BYLINSKII Z., KIM N., BAINBRIDGE C., YEH C., BORKIN D., PFISTER H., OLIVA A.: Beyond memorability: Visualization recognition and recall. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (Jan 2016), 519–528. 1, 3, 9, 12
- [BDFM17] BIGELOW A., DRUCKER S., FISHER D., MEYER M.: Iterating between tools to create and edit visualizations. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 481–490. 3
- [BMG*10] BATEMAN S., MANDRYK R., GUTWIN C., GENEST A., MCDINE D., BROOKS C.: Useful junk?: The effects of visual embellishment on comprehension and memorability of charts. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2010), pp. 2573–2582. 1, 2, 12
- [BNP*13] BROSZ J., NACENTA M., PUSCH R., CARPENDALE S., HURTER C.: Transmogrification: Causal manipulation of visualizations. In *ACM Symp. on User Interface Software and Technology (UIST)* (2013), pp. 97–106. 3, 13

- [BVB*13] BORKIN M., VO A., BYLINSKII Z., ISOLA P., SUNKAVALLI S., OLIVA A., PFISTER H.: What makes a visualization memorable? *IEEE Trans. on Visualization and Computer Graphics* 19, 12 (2013), 2306–15. 1, 2, 3, 12
- [CLC*15] CHI M., LIN S., CHEN S., LIN C., LEE T.: Morphable word clouds for time-varying text data visualization. *IEEE Transactions on Visualization and Computer Graphics* 21, 12 (Dec 2015), 1415–1426. 13
- [CM84] CLEVELAND W., MCGILL R.: Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association* 79, 387 (1984), 531–554. 2, 7, 8, 13
- [CM07] CAWTHON N., MOERE A.: The effect of aesthetic on the usability of data visualization. In *Information Visualization, 2007. IV '07. 11th International Conference* (July 2007), pp. 637–648. 2
- [CU14] CHESHIRE J., UBERTI O.: *London: The Information Capital*. Particular Books, 2014. 4
- [CWW*20] CHEN Z., WANG Y., WANG Q., WANG Y., QU H.: Towards automated infographic design: Deep learning-based auto-extraction of extensible timeline. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 917–926. doi:10.1109/TVCG.2019.2934810. 2, 3
- [CZW*20] CUI W., ZHANG X., WANG Y., HUANG H., CHEN B., FANG L., ZHANG H., LOU J., ZHANG D.: Text-to-viz: Automatic generation of infographics from proportion-related natural language statements. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 906–916. doi:10.1109/TVCG.2019.2934785. 3
- [DH72] DUDA R., HART P.: Use of the hough transformation to detect lines and curves in pictures. *Commun. ACM* 15, 1 (Jan. 1972), 11–15. 7
- [FE11] FEW S., EDGE P.: The chartjunk debate. *Visual Business Intelligence Newsletter*, no. June (2011), 1–11. 2
- [goo] Google Custom Search API. <https://developers.google.com/custom-search/>. Accessed: 2019-06-30. 9
- [HA06] HEER J., AGRAWALA M.: Multi-scale banking to 45 degrees. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (Sept 2006), 701–708. 8, 13
- [HAS11] HULLMAN J., ADAR E., SHAH P.: Benefitting infovis with visual difficulties. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2213–2222. 1, 2, 3
- [HB03] HARROWER M., BREWER C.: Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal* 40, 1 (2003), 27–37. 7, 9, 13
- [HKF15] HAROZ S., KOSARA R., FRANCONERI S.: Isotype visualization: Working memory, performance, and engagement with pictographs. In *ACM Conf. on Human Factors in Computing Systems (CHI)* (2015), pp. 1191–1200. 3, 12
- [Hol84] HOLMES N.: *Designer's guide to creating charts & diagrams*. Watson-Guptill, 1984. 2
- [HRC15] HARRISON L., REINECKE K., CHANG R.: Infographic aesthetics: Designing for the first impression. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2015), pp. 1187–1190. 1, 2, 12
- [KGV83] KIRKPATRICK S., GELATT C. D., VECCHI M. P.: Optimization by simulated annealing. *SCIENCE* 220, 4598 (1983), 671–680. 8
- [Kos89] KOSSLYN S.: Understanding charts and graphs. *Applied Cognitive Psychology* 3, 3 (1989), 185–225. 2
- [KSL*17] KIM N., SCHWEICKART E., LIU Z., DONTCHEVA M., LI W., POPOVIC J., PFISTER H.: Data-driven guides: Supporting expressive design for information graphics. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 491–500. 2, 3, 5, 13
- [Lew95] LEWIS J.: Ibm computer usability satisfaction questionnaires: Psychometric evaluation and instructions for use. *Int. J. Hum.-Comput. Interact.* 7, 1 (Jan. 1995), 57–78. 11
- [LGBH93] LANG P., GREENWALD M., BRADLEY M., HAMM A.: Looking at pictures: Affective, facial, visceral, and behavioral reactions. *Psychophysiology* 30, 3 (1993), 261–273. 12
- [LM14] LI H., MOACDIEH N.: Is “chart junk” useful? an extended examination of visual embellishment. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 58, 1 (2014), 1516–1520. 3
- [MBL*91] MINIARD P., BHATLA S., LORD K., DICKSON P., UNNAVA H.: Picture-based Persuasion Processes and the Moderating Role of Involvement. *Journal of Consumer Research* 18, 1 (06 1991), 92–107. arXiv:<http://oup.prod.sis.lan/jcr/article-pdf/18/1/92/5133182/18-1-92.pdf>. 12
- [MTW*12] MOERE A. V., TOMITSCH M., WIMMER C., CHRISTOPH B., GRECHENIG T.: Evaluating the effect of style in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (Dec 2012), 2739–2748. 2
- [ope] OpenCV. <https://opencv.org>. Accessed: 2019-06-30. 9
- [pet] Peter Ørntoft. <http://www.peterorntoft.com/>. Accessed: 2019-06-30. 4
- [pik] Piktochart. <https://piktochart.com/>. Accessed: 2019-06-30. 3
- [PKM18] PARK J. H., KAUFMAN A., MUELLER K.: Graphoto: Aesthetically pleasing charts for casual information visualization. *IEEE computer graphics and applications* 38, 6 (2018), 67–82. 3
- [PRS*15] PANDEY A., RALL K., SATTERTHWAITE M., NOV O., BERTINI E.: How deceptive are deceptive visualizations?: An empirical analysis of common distortion techniques. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2015), pp. 1469–1478. 5, 8, 12
- [qt] Qt. <http://www.qt.io/>. Accessed: 2019-06-30. 10
- [RKB04] ROTHER C., KOLMOGOROV V., BLAKE A.: “grabcut”: Interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.* 23, 3 (Aug. 2004), 309–314. 6
- [SHK15] SKAU D., HARRISON L., KOSARA R.: An evaluation of the impact of visual embellishments in bar charts. In *Eurovis* (2015), pp. 221–230. 5, 7, 8, 12, 13
- [SK16] SKAU D., KOSARA R.: Arcs, angles, or areas: Individual data encodings in pie and donut charts. *Computer Graphics Forum* 35, 3 (2016), 121–130. doi:10.1111/cgf.12888. 5, 7, 8, 9, 12, 13
- [SM14] SETLUR V., MACKINLAY J.: Automatic generation of semantic icon encodings for visualizations. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2014), pp. 541–550. 3
- [Tuf83] TUFTE E.: *The Visual Display of Quantitative Information*. 1983. 1, 2
- [Tuf90] TUFTE E.: *Envisioning Information*. Graphics Press, 1990. 2
- [VW07] VIÉGAS F., WATTENBERG M.: Artistic data visualization: Beyond visual analytics. In *Online Communities and Social Computing* (Berlin, Heidelberg, 2007), Springer Berlin Heidelberg, pp. 182–191. 2
- [WSZ*20] WANG Y., SUN Z., ZHANG H., CUI W., XU K., MA X., ZHANG D.: Datashot: Automatic generation of fact sheets from tabular data. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (Jan 2020), 895–905. doi:10.1109/TVCG.2019.2934398. 2, 3
- [WZH*18] WANG Y., ZHANG H., HUANG H., CHEN X., YIN Q., HOU Z., ZHANG D., LUO Q., QU H.: Infonice: Easy creation of information graphics. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2018), pp. 335:1–335:12. 2, 3, 5, 13
- [XHRC*18] XIA H., HENRY RICHE N., CHEVALIER F., DE ARAUJO B., WIGDOR D.: Dataink: Direct and creative data-oriented drawing. In *ACM Conference on Human Factors in Computing Systems (CHI)* (2018), pp. 223:1–223:13. 2, 3, 5, 13