



HAL
open science

Suggested Interactivity: Seeking Perceived Affordances for Information Visualization

Jeremy Boy, Louis Eveillard, Françoise Detienne, Jean-Daniel Fekete

► **To cite this version:**

Jeremy Boy, Louis Eveillard, Françoise Detienne, Jean-Daniel Fekete. Suggested Interactivity: Seeking Perceived Affordances for Information Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2016, 22 (1), pp.639-648. 10.1109/TVCG.2015.2467201 . hal-01188973

HAL Id: hal-01188973

<https://inria.hal.science/hal-01188973>

Submitted on 31 Aug 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Suggested Interactivity: Seeking Perceived Affordances for Information Visualization

Jeremy Boy, Louis Eveillard, Françoise Detienne, and Jean-Daniel Fekete

Abstract—In this article, we investigate methods for suggesting the interactivity of online visualizations embedded with text. We first assess the need for such methods by conducting three initial experiments on Amazon’s Mechanical Turk. We then present a design space for Suggested Interactivity (*i. e.*, visual cues used as perceived affordances—SI), based on a survey of 382 HTML5 and visualization websites. Finally, we assess the effectiveness of three SI cues we designed for suggesting the interactivity of bar charts embedded with text. Our results show that only one cue (SI₃) was successful in inciting participants to interact with the visualizations, and we hypothesize this is because this particular cue provided feedback.

Index Terms—Suggested interactivity, perceived affordances, information visualization for the people, online visualization.

1 INTRODUCTION

Information visualization (infovis) designers often tend to consider visualizations as individual artifacts that users willingly come to view and interact with. However, many visualizations end up embedded in webpages with other media like text, and it is unclear whether web users actually have a natural propensity to interact with these. While the importance of interaction is well recognized within the infovis community [16, 31], and is often emphasized in the design of successful visualizations (*e. g.*, [54]), until recently most mainstream data graphics have been static (*e. g.*, *infographics* in news papers)—and many still are, even on the web. Thus, in a context like that of a data-journalism article, where text and visualization can co-exist, it seems optimistic to assume that people will know they can or should interact with a visualization to find information.

Most interactive features on the web use standard widgets (essentially buttons), which usually rely on metaphors of physical objects to suggest how they operate; they borrow *affordances* from their real world counterparts. These are not ‘real’ affordances in a ‘Gibsonian’ sense [24, 42], since they do not support the physical actions of pointing, clicking, and (possibly) dragging with a mouse device; but they do suggest that an interaction is possible. Buttons, for example, designed with embossments and drop shadows (illustrating their mechanic origin) suggest that ‘pressing’ them is possible.

While effective, these analogies fall short when it comes to more abstract or symbolic interactive features, which mainly depend on design conventions. Hyperlinks, for example, use by default a specific visual variable, *i. e.*, color hue, and an additional visual mark, *i. e.*, an underline, to suggest that they can be clicked. This is ‘heavy’ design, as it requires two visual attributes to highlight a single difference with other textual elements, *i. e.*, interactivity; and the fact that a user knows that such highlighted text can be clicked on is purely conventional.

Interactive visualizations however, have neither convention nor real world counterparts that can help suggest their interactivity—a pie chart does not afford eating! Thus, assuming that interacting with visualizations embedded with text is not obvious to everyday internet users, we ask: “*how can we attract these users’ attention to a visualization and suggest its interactivity through design?*”

- Jeremy Boy is with INRIA, Telecom ParisTech, and EnsadLab, e-mail: myjyby@gmail.com.
- Louis Eveillard is with EnsadLab, e-mail: louis.eveillard@gmail.com.
- Françoise Detienne is with Telecom ParisTech, e-mail: francoise.detienne@telecom-paristech.fr.
- Jean-Daniel Fekete is with INRIA, e-mail: jean-daniel.fekete@inria.fr.

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

To address our assumption, we first conducted three experiments on Amazon’s Mechanical Turk (AMT) that confirm that a majority of people do not interact with visualizations embedded with text, even if these are more efficient for performing given tasks. To address our main research question, we then surveyed 382 HTML5 and visualization websites to see how interaction designers make use of different visual cues to suggest the interactivity of abstract graphical objects and areas—where several interactive objects may be aggregated into a whole (*e. g.*, the bars in a bar chart). Based on this, we developed a design space for Suggested Interactivity (SI). Finally, we conducted a follow-up study to evaluate the effectiveness of three SI cues applied to bar charts, which we believe are most representative of the diversity of our design space. As such, our main contributions are as follows:

- an assessment of the need for SI in cases where visualizations are embedded with text;
- a design space for SI;
- an evaluation of three different SI cues applied to bar charts, which we created using specific design consideration derived from our design space; which led us to
- initial recommendations for the design of SI cues for infovis.

This article is organized in the following way: it begins with a background section that introduces the idea of passive interaction with media and software; reviews the concepts of affordances; describes several new graphic standards for interface design, and discusses how these consider suggesting the interactivity of specific features; and presents previous work on the use of motion and icons in interface design to attract users’ attention and convey meaning. Section 3 verifies our initial assumption. Section 4 describes our survey of existing SI cues, and introduces our design space; it then provides a set of design considerations, which we use for creating three SI cues that we apply to bar charts. Finally, Section 5 presents the design and evaluation of these cues, and discusses the implications of our early results for future designs.

2 RELATED WORK

Finding ways to suggest the interactivity of graphical objects is an old issue for interface design. However, with the development of new graphic standards, especially on the web where designers and artists can easily create any new kind of interactive visuals, the concern is becoming of importance once again. In addition, while Segel & Heer have pointed out the use of “*markers of interactivity*” in certain narrative visualizations [48], this issue has hardly been considered by the infovis community, which seems unfortunate as visualizations are now spreading out to new users who cannot all be formally trained.

2.1 Passive Interaction

As mentioned in the introduction, information visualizations on the web often end up embedded with other media like text. Common examples of this can be found in data journalism articles like those of the Guardian [28]. Most of the time, these other media are static and do not suggest any interaction.

Passive interaction defines the changing or enhancement of a user's mental model while interacting with a system, without modifying the system's model [25, 49]. In simpler terms, passive interaction occurs when reading a text, looking at an image or a visualization, or more generally, when receiving, decoding, and interpreting a new piece of information, without having to manipulate the medium—it requires little or no input device manipulation.

Passive interaction is a main component of what Victor names *information software* [53]. According to his definition, such software “serves the human urge to learn;” it initiates cognitive processes like learning, comparing, and decision making. He opposes this to *manipulation software*, which “serves the human urge to create;” it helps a user construct and manipulate an external model: that of a system. Most visualizations on the web are intended to show viewers something new: they serve the urge to learn. However, they are not necessarily static, and users can manipulate them to display data in different ways. Thus, while Victor speaks of software, we posit the same goes for online media. People have certain expectations from media, and if the main medium requires no interaction (*e. g.*, text), it seems unlikely that they will seek to determine whether other components of a webpage are ‘manipulable.’

2.2 Affordances and Perceived Affordances

To suggest that an action is possible, interaction and interface designers often call upon certain graphical attributes, which are generically named *perceived affordances*.

The term *affordance* was coined by the perceptual psychologist J.J. Gibson to define certain properties of the world that induce action in an organism [24]. Norman introduced the term to the field of Design to define the specific attributes of physical artifacts that help people understand how they ought to be manipulated [43]. However, due to inappropriate use of the term, he later distinguished *real affordances*, *i. e.*, the actual physical properties of an artifact that call for action; from *perceived affordances*, *i. e.*, the perception and/or understanding a person has of the actions that can be performed with that artifact.

Hartson makes a further distinction between *cognitive*, *physical*, *sensory* and *functional affordances* [30]. He emphasizes the importance of dealing with all four types, and suggests a list of questions to consider when designing interactions. Based on this, Vermeulen *et al.* have proposed the following complete definition for perceived affordances: “*Perceived affordances are cognitive affordances that are understandable through well-defined sensory affordances (e.g. a door's handle) and reveal a physical affordance (an action possibility), which is coupled to a functional affordance (the action's purpose). Perceived affordances occur before the user's action and invite them to an appropriate action*” [52]. Tang *et al.* [50] have also extended Hartson's model with *perceived affective affordances* and *perceived control affordances*. The first describes attributes of the artifact that can trigger or stimulate an emotional reaction in the user; and the latter reflects attributes of the artifact that give the user a certain level of perceived control over the interaction.

In parallel, Vermeulen *et al.* have highlighted the importance of providing feedforward in interaction design [52]; it “*tells users what the result of their action will be.*” However, while important, perceived affordances and feedforward are not easily designed, especially as new interface design standards lean towards more abstraction. This is typically the case for visualizations, as they are abstract representations by nature. In addition, perceived affordances are generally linked to distinct graphical objects that can be manipulated. Visualizations are interactive ‘areas’ composed of multiple objects (visual marks, axes, *etc.*), which may all be manipulable. As such, it seems more appropriate to find ways to suggest the interactivity of the ‘visualization area’ than to design perceived affordances for each individual object.

2.3 New Interface Design Standards

Recent developments of ‘platform-centric’ user-interface guidelines have approached the issue of perceived affordances in greater detail than before. Historically, Apple was one of the first companies to provide such guidelines—the Human Interface Guidelines (HIG)—in order to help third parties create consistent experiences when developing for their platforms. These guidelines explained which perceived affordances to use, and when to use them. For example, iOS 6's HIG emphasized the use of embossed, skeuomorphic cues to suggest that buttons are actionable. However, iOS 7 and 8's HIG propose a radical reorientation: they suggest designers “*embrace borderless buttons,*” and mention that the use of “*context, color and a call-to-action title [is enough] to indicate interactivity*” [6].

Similarly, Microsoft's *Visual Studio 6.0's* instructions for “*creating a user interface*” [40] mentioned that “*a user interface also makes use of affordances. For instance, the three-dimensional effects used on command buttons make them look like they are meant to be pushed. If you were to design a command button with a flat border, you would lose this affordance and it wouldn't be clear to the user that it is a command button.*” However, ten years later Microsoft reverted its course and presented the *Metro Design* standard which states that “*while lack of affordance and discoverability may sound discouraging, note that drop-down menus and context menus—other mechanisms for initiating actions—suffer similar problems*” [38].

More recently, Google published its own set of guidelines for *Material Design*, which also describes a variety of ways to suggest interaction [27]. Depending on the importance of the element and its associated action, it can be made into a colored “*floating action button,*” a “*raised button*” or a “*flat button*”—the number of cues is proportional to the importance of its function. A fundamental aspect of *Material Design* is the use of motion to convey interactivity: “*Perceiving an object's tangible form helps us understand how to manipulate it. Observing an object's motion tells us whether it is light or heavy, flexible or rigid, small or large. Motion in the world of material design is not only beautiful, it builds meaning about the spatial relationships, functionality, and intention of the system*” [55]. Rather than copying the textures and shadows of physical objects, applying realistic motion to graphical objects helps with the perception of affordance.

2.4 Attracting Attention and Conveying Meaning

Motion is known to have major psychophysical benefits for capturing attention. As humans, our visual system is extremely sensitive to fragments of natural motion [41]; our peripheral vision is also highly receptive to movement, and operates at a much lower resolution than our fovea [29]. We are able to view and interpret several motions simultaneously [46], and we can perform complex grouping tasks of moving elements. Research in vision science has shown that we perceive causality and animacy in motion, primarily through perceptual, and possibly through modular, processes [47]. This means we attribute ‘life,’ intentionality, and behavior to moving objects without impeding higher-level cognitive processes or judgements. However, it is important to note that the graphical representation of motion (*i. e.*, animation) is not always easy to perceive accurately or to conceptualize when it conveys more abstract meaning [51].

Pictographic symbols are common in interface design for conveying meaning, and more broadly in machine display design in general [20, 35]. Such pictograms are commonly referred to as *icons*. Although there was originally some debate about their effectiveness [8, 26, 36], today icons are a major part of any interface, whatever the system or device, and have undoubtedly greatly contributed to the success of personal computers. Huang *et al.* claim that “*icons offer the perception of affordance, which can facilitate human-machine interaction in terms of ecological perception*” [33].

Lodding [35] has proposed a taxonomy of icons for user interfaces, which is composed of three dimensions: representational, abstract, and arbitrary. *Representational icons* are “*typical*” images that serve as examples for general classes of objects. *Abstract icons* represent concepts, and use illustrations of real objects to refer to abstract ideas.

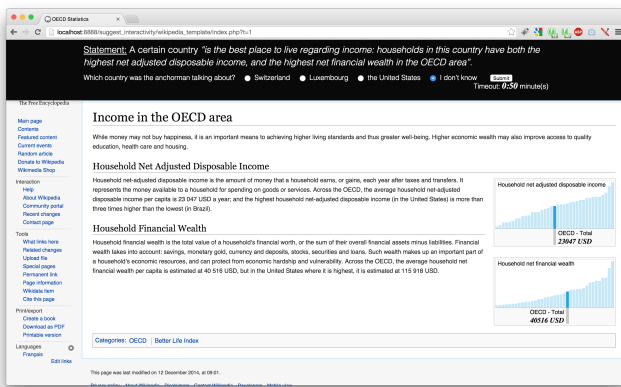


Fig. 1: Screen capture of an AMT experiment

Finally, *arbitrary icons* do not relate to any object in the world; they are “invented” and assigned a conventional meaning.

The combination of motion (or animation) and icons can help clarify meaning, explain the purpose of a given tool, demonstrate its capabilities, and even convey its method of use [8, 10, 29]. There are essentially two kinds of animated icons: icons that incorporate animated graphics and kineticons [29]. This distinction is based on the fact that icons are generally composed of a bounding box which contains a pictogram. If the pictogram is animated, then it is an icon that incorporates animated graphics. If the bounding box is animated, then it is a kineticon.

Baecker *et al.* propose some high-level considerations for the design of the first kind of animated icons [8]; they identify ten basic ways in which they can be useful, and illustrate these with relevant questions.

Harrison *et al.* describe “six popular sources” of inspiration for the design of kineticons [29]. Five of these use references to real-world motions: *biological motion*, *gestures*, *organic motion*, *mechanical motion* and *physics and natural effects*. These effectively use our ability to perceive causality and animacy in the way things move. The sixth is *cartoon conventions*, which are described as commonly accepted caricatures or exaggerations of real-world motions.

All these general HCI design guidelines should certainly inspire the design of perceived affordances for online visualizations. However, general user interfaces are much more mature and familiar to the general public than visualizations, so the history and evolution of their perceived affordances should certainly be taken into account as well.

3 TESTING INTERACTION PROPENSITY

To verify whether online users are naturally inclined to interact with visualizations embedded with text, and to assess the need for SI, we conducted three initial experiments on AMT. For each, we used a series of seven simulated Wikipedia articles we created, which included both visualizations (bar charts) and text. We used data and text from the OECD’s Better Life Index website [44], and grouped similar-topical indicators into specific articles (as is done on the OECD’s website); topics were Housing, Income, Education, Environment, Health, Safety, and Work-Life balance. We reduced the original text for each article to limit the amount of contextual information it provided, in order to create a better information-balance between text and visualizations. We created specific sections in the articles for each indicator (either two or three, depending on the topic), and we displayed the bar charts next to the corresponding paragraphs; this design follows the traditional Wikipedia layout, with text to the left, and images (in this case charts) to the right (see Figure 1).

We chose to simulate Wikipedia articles for ecological validity, since Wikipedia will undoubtedly soon provide tools for building interactive visualizations—the markup already supports the creation of static charts (*e. g.*, [9]). While it is arguable that this choice may bias

users’ propensity to interact with visualizations (since there are few in contemporary Wikipedia articles), it is a realistic setting. In the same line, online news articles like those of the Guardian, which heavily rely on traditional media like text, now integrate more and more interactive graphics. Thus, we believe Wikipedia is a good and timely environment for testing people’s propensity to interact with visualizations.

To save screen-real-estate, we limited the default labeling of the charts to OECD averages, and highlighted the corresponding bars. As all bar charts in an article presented data about the same OECD countries, and as there were either two or three per article, we implemented a simple brushing-linking technique to highlight identical countries in the different charts simultaneously; this also displayed their labels and precise values. Thus, interaction was necessary to extract specific values from the charts, and this hover interaction was what we expected participants to discover. We assumed that if such a simple interaction was difficult to discover, more complex ones would be even more so.

We created a simple fact-checking task for each article, as fact-checking is a common activity on the web; this was also done to encourage people to go straight to the point, *i. e.*, to find the specific datapoint we asked them to target, without needing to seek for contextual information. Each article corresponded to a trial, resulting in seven trials per participant, the order of which we simply randomized to prevent carry-over effects.

Our first experiment was conducted to assess whether people are inclined to interact with charts to carry out fact-checking tasks. Our second experiment was conducted to make sure charts are more efficient than text (as we expected) for performing the tasks we created, and that if given the chance to discover their interactivity, participants would be more inclined to use them. Finally, our third experiment was conducted to make sure our choice of simulating Wikipedia articles did not bias participants’ propensity to interact with charts.

3.1 Experiment 1

3.1.1 Design

Tasks We designed the fact-checking tasks specifically to make the visualizations more efficient than the text for retrieving the necessary information (provided that participants interacted with them). For each article (or trial), we set a multiple-choice extraction task with 3 possible answers and an “I don’t know” option, which required participants to consider all the indicators presented at once; a typical example for an article that presents three indicators would be: “Which country has the highest rate of X, the lowest rate of Y, and an average rate of Z.” Each possible answer was hand-picked to complicate reliance on potential background knowledge. For example, the possible answers for the article on Income, where the task was to find the country in which households have both the highest income and financial wealth, were Luxembourg, Switzerland, and the United States (see Figure 1), *i. e.*, three countries where one might expect that income is high, but does not necessarily know in which it is the highest. To find the correct answer, participants simply needed to brush over one of the visualizations until they found a country that met the first requirement, and then check its performance across the other indicators. We also purposefully chose combinations of min/max tasks to make the questions as highly-congruent as possible [13]. However, to balance out the study design, and to make it possible to perform the tasks using the text alone, we added specific mentions of the possible answers in the paragraphs corresponding to each indicator. This would require reading through all paragraphs and memorizing how well each possible country performs across all indicators.

Procedure Upon accepting the Human Intelligence Task (HIT) on AMT, participants were directed to an external page for the study. On this page, they were first asked to complete a pre-study to make sure that they had the necessary English skills to participate in our experiment, and that they were willing to comply with instructions—the pre-study was an intermediate English reading comprehension test taken from [21]. Participants who failed the pre-study were not allowed to continue. Those who succeeded were then asked to fill out a short, anonymous demographic survey; were given a simulated task scenario

to consider in order to situate the fact checking activity; and were administered the study. Before each trial, participants were instructed the task and were invited to click on a “Display Wikipedia page” button to display the article. This triggered a countdown, which we set to force participants to be quick. Finally, upon completion, participants were asked to fill out another short survey about the study.

Hypotheses We had two simple hypotheses for this experiment:

- **H1.1:** a majority of participants will not know that the charts are interactive, and therefore they will not use them to complete the trials; and
- **H1.2:** a majority of participants who ‘discover’ the interactivity of the charts will use them throughout all subsequent trials, as they are in principle more efficient.

Participants We recruited 70 participants on AMT who were required to have a 98% acceptance rate and a total of 1000 or more HITs approved. We removed the work of 2 participants from the collected data, as these had taken the HIT on a mobile device—such devices do not support brush interactions, *i. e.*, for hover interactions—but we paid them nonetheless. This resulted in a subset of 68 participants, all of which were native english speakers.

Coding We traced participants’ low-level activity on the external page. First, we counted the number of brush (hover) interactions each participant performed. However, since such indicators are often noisy, we also counted what we refer to as *decisive brushes*, *i. e.*, brush interactions over bars related to the answers participants gave that lasted more than 250 ms (so that participants had time to see their effect on the display). For example, if a participant answered “the United States” for the article on Income (mentioned above), we coded every brush interaction that lasted longer than 250 ms over a bar encoding US data as a decisive brush. Since we did not reset the display to its original state after user interactions (with the OECD average highlighted and labeled), we also counted one decisive brush if the last bar to be highlighted in a series of brush interactions that each lasted less than 250 ms was related to the answer participants gave. Note that in this article, when we claim participants ‘used’ the charts to find answers, we mean they performed at least one decisive brush. Next, we counted both the number of trials in which each participant performed brush interactions and decisive brushes, and for each we coded 1 when such interactions were performed in all subsequent trials to the one in which the interactivity of the charts was discovered, and 0 otherwise. Finally, we coded participants’ answers 1 when correct, 0 when the “I don’t know” option was submitted, and -1 when incorrect.

3.1.2 Results

All the analyses and discussions in this article are based on estimation, *i. e.*, point estimates and effect sizes with confidence intervals (95% CI), with respect to the concerns and recommendations in [5, 18, 19]. Point estimates and 95% CI are based on 10,000 percentile bootstrap replicates of the statistic (in this case percentages in means) applied to the data [15]. The proportion 95% CI are calculated using the nicely documented `VassarStats.net` web-application.

We first inspected participants’ scores. As the questions and tasks were overall quite simple, and as the default answer was “I don’t know” (= 0), we removed all participants whose total score was below or equal to 0 from further analysis, as we considered these to be either random clickers, or people who only provided answers based on *a priori*. This resulted in a subset of 59 participants.

42.4%, 95% CI [30.6%, 55%] of these participants performed at least one brush interaction, and **28.8%** [18.8%, 41.4%] performed at least one decisive brush. Thus, **68%** [48.4%, 84.3%] of participants who performed a brush interaction also performed a decisive brush.

52% [33.5%, 69.9%] of participants who performed brush interactions performed at least one in all seven trials, and **60%** [40.7%, 76.6%] performed at least one brush in every subsequent trial to the one in which they discovered the interactivity of the charts.

Finally, **58.8%** [36%, 78.4%] of participants who performed decisive brushes performed at least one in all seven trials, and **88.2%**

[65.7%, 96.7%] performed at least one in every subsequent trial to the one in which they first performed a decisive brush.

3.1.3 Discussion

Although admittedly we had expected that less participants would discover the interactivity of the charts, our results still confirm **H1.1**. We suspect this higher number may be due to the layout of the Wikipedia articles, as the charts were set to the right hand side of the page, next to the scroll bar. Participants whose displays were too small to show the whole webpage may have hovered over the charts while moving their cursor to scroll. This is interesting though, as it could suggest that the layout itself can be used to help people discover interactive content.

H1.2 is also confirmed, as a majority of participants who discovered the interactivity of the charts continued to brush them throughout all subsequent trials; and this is particularly true for participants who performed decisive brushes. This seems to indicate that participants who discovered the charts’ interactivity and understood how to use them perceived the charts as more efficient than the text.

3.2 Experiment 2

3.2.1 Design

To extend our confirmation of **H1.2**, and to ensure that our task-design did indeed make the visualizations more efficient for extracting the necessary information, we conducted a second experiment on AMT. The design of this experiment was identical to that of the previous, with the exception that in trials 3, 4, and 5 (out of seven, and whatever the article) we removed all the textual information, and made the visualizations much larger, laid them out on the left hand side of the screen, and explicitly mentioned that they were interactive. This was done to force participants to use the charts. Thus, this experiment consisted of two initial trials ([1,2]) in which the charts were embedded with text, followed by three trials ([3-5]) in which there was no text, and completed with two final trials ([6,7]) in which the charts were once again embedded with text. The scenario, tasks, and procedure were kept the same as before.

Hypothesis We had four hypotheses for this experiment:

- **H2.1:** all participants will interact with the charts in trials [3-5];
- **H2.2:** a majority of participants will use the charts in trials [6,7];
- **H2.3:** there will be good evidence that more participants interact with the charts in trials [6,7] than in [1,2]; and
- **H2.4:** as more participants should interact with the charts in trials [6,7] than in [1,2], and as the charts are in principle more efficient, there will be good evidence that participants complete trials [6,7] faster than [1,2].

Participants We recruited 70 different participants on AMT, in order to make sure they would not be biased by the first experiment. However, this time we only retained the work of 47 participants whose total score was higher than 0; all were native english speakers. We then coded the data in the same way as in Experiment 1.

3.2.2 Results

We first inspected the number of participants who interacted with the charts in trials [1,2]. In trial 1, **12.7%**, 95% CI [5.9%, 26.4%] performed at least one brush interaction, and **6.4%** [2.2%, 18.6%] performed at least one decisive brush. In trial 2, **8.5%** [3.4%, 21.3%] performed at least one brush interaction, and **2.1%** [0.4%, 12.7%] performed at least one decisive brush.

We then inspected the number of participants who interacted with the charts in trials [3-5], where no text was displayed. In trial 3, **80.5%** [72.3%, 92.6%] (40/47) performed brush interactions, and **72.3%** [58.2%, 83.1%] performed decisive brushes. In trial 4, **95.7%** [85.7%, 98.8%] performed brush interactions, and **78.7%** [65.1%, 88%] performed decisive brushes. In trial 5, **95.7%** [85.7%, 98.8%] performed brush interactions, and **91.5%** [80%, 96.6%] performed decisive brushes. Out of the seven participants who did not interact with

the charts in trial 3, four discovered the interactivity in trial 4 and continued to interact with the charts in trial 5; one discovered the interactivity in trial 5; one discovered the interactivity in trial 4, but oddly did not continue to interact in trial 5; and one simply never interacted.




Fig. 2: Participants’ progression between trials [3,5].

These last two cases made us wonder whether there may have been a visualization literacy problem [13], since participants not only had to interact with the charts to find the answers, but also had to know how to search for them visually. To check for this, we inspected whether participants showed signs of ‘progress’ through trials 3 to 5. To estimate progress, we isolated participants who performed brush interactions, and calculated the proportions of these who performed decisive brushes in trials 3 and 5; we then calculated the difference between proportions for both trials (Figure 2a). We also compared the mean time participants spent in trials 3 and 5 (Figure 2b), as well as their mean scores (Figure 2c).

After that, we inspected the number of participants who interacted with the charts in trials [6,7]. In trial 6, **82.9%** [69.9%, 91.1%] performed brush interactions, and **70.2%** [56%, 81.3%] performed decisive brushes. In trial 7, **89.4%** [77.4%, 95.4%] performed brush interactions, and **72.3%** [58.4%, 83.1%] performed decisive brushes.

Finally, we aggregated the results for trials [1,2] and for trials [6,7], in order to compare what participants did before and after they were ‘forced’ to use the charts; we inspected the number of unique users who interacted with the charts (Figure 3a), and the mean time they spent completing the trials (Figure 3b).




Fig. 3: Comparisons of aggregated results for trials [1,2] and [6,7].

3.2.3 Discussion

Our results do not support **H2.1**. An important majority of participants did interact with the charts in trials 3, 4 and 5, but it seems they needed to progressively get ‘used’ to them. Figure 2a shows some evidence of an increase in decisive hovers between trials 3 and 5, as the lower boundary of the 95% CI is only slightly below 0. This suggests that participants needed the three trials to elaborate strategies for finding the answers in the charts. Similarly, Figure 2b shows good evidence of a reduction of time spent completing the trials, which suggests that participants progressively became more efficient in finding the answers in the charts. As fewer participants performed hover interactions in trial 3, we hypothesize that they may have initially preferred to avoid the charts, not necessarily because of lack of propensity to interact, but because of lack of strategies for finding the answers in charts *i. e.*, because of low visualization literacy. If a person lacks visualization literacy, the cost of interacting with a chart could be perceived as greater than the benefit, since the benefit is unknown.

H2.2, **H2.3**, and **H2.4** however, are all confirmed. This suggests that the charts are indeed more efficient, and that given the chance to discover their interactivity (and to elaborate effective search strategies), most participants will keep using them instead of going back to

the text. Thus, we consider that participants’ lack of propensity to interact with the charts is not due to an efficiency problem; and while perceived efficiency may be an initial concern (due to low visualization literacy), it is rapidly overcome.

3.3 Experiment 3

Finally, to ensure our results were not biased by the Wikipedia template, we conducted a third experiment which replicated Experiment 1, but for which we removed all Wikipedia styling attributes from the articles. This was done to check that our results were not confounded by expectations participants may have had from Wikipedia articles. We kept the same overall layout, scenario, tasks, and procedure.

Hypothesis We had one simple hypothesis for this study:

- **H3**: results will be consistent with Experiment 1, meaning that the Wikipedia styling did not bias participants’ behavior.

Participants Once again, we recruited 70 different participants on AMT, in order to be able to establish comparison with the results of Experiment 1 in a between subjects design. We retained the work of 51 participants whose total score was higher than 0; all were native english speakers. We then coded the data in the same way as before.

3.3.1 Results

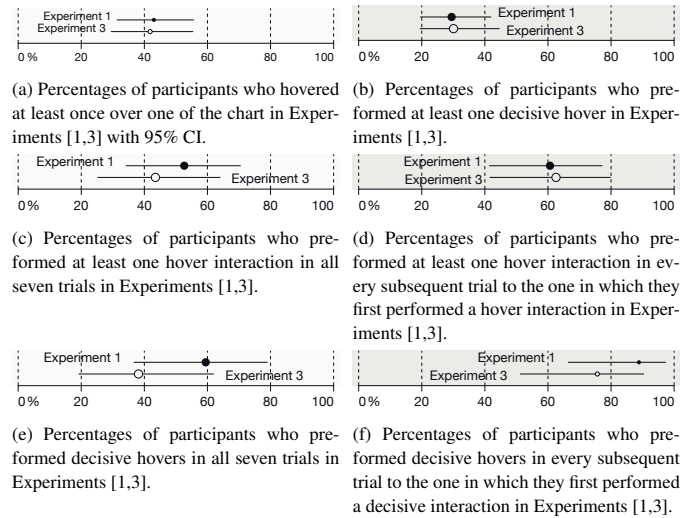


Fig. 4: Between subjects comparisons of results in Experiments [1,3].

Our analysis was exactly the same as for Experiment 1; all results are shown in Figure 4, and are compared with those of Experiment 1 in a between subjects fashion.

3.3.2 Discussion

Although the 95% CI in Figure 4 are quite wide—which seems normal as the estimation is based on only one bit of information, *i. e.*, percentages—results show no real evidence of a difference between Experiments 1 and 3. On the contrary, they show a high similarity (with a slight exception in Figure 4e). This confirms **H3**, and suggests that the Wikipedia template did not bias participants’ propensity to interact with the charts.

3.4 Initial Experiments’ Discussion

Overall, our results show that a majority of people lack initial propensity to interact with charts when these are embedded with text, whatever the styling of the webpage. This may in part be due to visualization literacy problems, but if people discover the interactivity, they are likely to rapidly learn how to perform the necessary visual queries for finding answers in charts. They will then continue to interact with them, as the charts are indeed more efficient for fact-checking tasks

like the ones we created. This indicates that people can be motivated to interact with visualizations if they are shown the possibility, and highlights the need for suggested interactivity (at least when visualizations are embedded with text).

4 SUGGESTED INTERACTIVITY

In light of the work presented in Section 2, we propose the following definition for Suggested Interactivity (SI): *Suggested Interactivity is a set of methods for indicating that a graphical area can be interacted with by subtly directing a user’s attention so as not to impede too heavily on this person’s focus or on the rest of the interface design.* SI cues are then specific graphical elements or attributes that are used for suggesting interactivity.

While the concepts of *perceived affordance* and *suggested interactivity* are similar, we make the distinction based on the fact that the perception of affordance is generally related to design attributes of a unique/distinct interactive graphical object (e.g., a widget), whereas SI is related to visual cues that do not necessarily pertain to an individual object: SI cues can be icons or text labels (i.e., external objects) placed on top or next to an interactive area (e.g., a visualization), which can be composed of several interactive graphical objects.

With these definitions in mind, we conducted a survey of a variety of highly interactive websites to identify how designers create and make use of SI cues for abstract interactive features. We also surveyed several standard widgets, since, as we have mentioned above, new graphic standards tend to move away from the traditional embossments and drop shadows design. From this survey, we extracted a set of important dimensions for the design of SI cues, and constructed the design space presented in Figure 5. Note that while we focus on SI cues for visualizations here, this design space can be used to describe and generate SI cues for any kind of ‘abstract’ user-interface.

4.1 Survey and Design Space

We surveyed 230 HTML5 websites listed in [7, 11, 32], 150 data-journalism visualization websites listed in [37], and 2 of the Gapminder visualizations available for download [23]. This resulted in a total survey of 382 websites. We recorded all the different techniques used to suggest the interactivity of the webpages, or of specific graphical objects and areas within these (other than standard textual hyperlinks). Some websites did not include any SI cues, and many included similar ones. Overall, we identified 45 distinct cues, from which we extracted the following five main dimensions:

- Attractor:** the object that attracts attention to the interactive area;
- Animation:** the state of the attractor over time—note that in some cases the attractor is not animated;
- Trigger:** the event that initiates the animation;
- Visual attributes:** the specific visual variable(s) and/or mark(s) the animation is applied to; and
- Persistence:** the ongoing display or not of the cue once the interaction has been performed.

The **attractor** can either be the *object of interest*¹, i.e., the interactive graphical object or area itself, or an *external object* (e.g., an overlaid icon or text label). Its **animation**, when it exists, can be *staged*, i.e., a predefined on/off *blink* or *interpolation*—which is either *unique* (one-shot) or *looped*; or *dynamic*, i.e., dependent on specific ‘page-level’ user input (e.g., mousemove or mousewheel). The animation can be **triggered** by a *system-event* (e.g., pageload), or by a *user-event* (e.g., mousemove, mouseover, click, or mousewheel), and can be applied to various **visual attributes** of the attractor, i.e., to *visual variables* and/or to *extra visual marks* (similarly to hyperlinks, which use both). The **persistence** of the cue then determines whether it remains displayed after the intended interaction has been performed—in some cases it is removed immediately afterwards, as it can be considered that the user has discovered the interaction and will remember it

¹Note that if the attractor is the object of interest and that this object is a unique graphical object, the SI cue can be considered a perceivable affordance.

throughout the rest of the session. Note however, that persistent SI cues can also be temporarily hidden while the user is interacting with the interactive graphical object or area, or while this object is in focus. For example, the “play” button displayed on top of a video temporarily disappears while the user is watching the content.

To illustrate these dimensions, consider a standard hyperlink (even though we did not record these in our survey). The attractor is the object of interest, i.e., the clickable text, to which no animation is applied. A visual variable and an extra visual mark are used, i.e., color and underline. Finally, the cue is persistent, as it remains visible after a user clicks on the link, and later comes back to the webpage.

In addition to these main dimensions, we also coded the **intended interaction**, i.e., the interaction the user is invited to perform; as well as the presence or not of **feedforward**, i.e., a hint to the outcome of the interaction. Finally, although it is not directly accounted for in our design space, we identified the distinction between icons that incorporate animated graphics and kineticons described in Section 2.4. In some cases, we even found combinations of both. Generally, an attractor to which a staged or dynamic animation is applied on an extra visual mark is an icon that incorporates animated graphics; and an external object attractor to which a staged or dynamic animation is applied on a visual variable is a kineticon.

Our final design space is presented in Figure 5. Due to space limitations, we only reference one website per entry, but we provide a count (in brackets) of the number of websites that use each SI cue. We only include one row for the attractor dimension, as its levels are mutually exclusive: black cells encode objects of interest, and white cells external objects. Similarly, the persistence and feedforward dimensions are binary, so black cells indicate ‘true.’ We stress that the visual attributes we coded are only the ones to which the animation is applied. For example, an attractor may have a textual component, but if this component is not directly subject to the animation, it is not accounted for in our design space. Finally, we did not include *mousemove* in the intended interaction dimension, as we did not find, and could not think of any graphical object or area that simply relies on a mousemove to be interacted with.

In the following subsection, we present several specific cases we encountered in our survey, and discuss how these fit into our design space (when applicable).

4.1.1 Discussion

A majority of the SI cues we found (27/45) are applied to the object of interest; and in most cases (33/45), the type of animation is determined by what triggers it: staged animations are triggered by system-events (26/38) and dynamic animations are triggered by user-events (7/7). However, staged animations (12/38) can also be triggered by user-events. This occurs when SI cues are subject to sequenced interactions, i.e., predefined linear series of interactions the user is invited to perform. Each interaction triggers the display of a new SI cue for the subsequent interaction. We highlighted these cases using red and blue cells. Sequences can focus on different interactive graphical objects or areas (blue cells) or on a same graphical object (red cells).

Sequenced interactions with different graphical objects or areas
To illustrate this, suppose a user is required to click on a first graphical object g_{01} before clicking on a second graphical object g_{02} . On page-load (a system-event trigger), a staged animation SI cue is applied to g_{01} , but no cue is applied to g_{02} . When the user identifies the SI cue and clicks on g_{01} (i.e., a user-event trigger), then a new staged animation SI cue is applied to g_{02} . An example of this is SIcue9: the user has to click on a ‘play’ button (g_{01}) to reveal an SI cue applied to a slider (g_{02}). We found this case hard to code as a user-event triggered SI cue, since the initial required interaction (performed on g_{01}) is unrelated to the second interactive graphical object (g_{02}). Thus, for simplicity we consider that any user-event that is not performed on the ‘whole page’ level (e.g., mousemove or mousewheel—yellow cells) and that reveals an SI cue for an interactive graphical object (g_{02}) other than the one the user is already interacting with (g_{01}) triggers a new system-event (page update), which in turn triggers the animation of the SI cue. This way, we consider all staged animations to be triggered

			SIcuel (1)	SIcuel2 (1)	SIcuel3 (14)	SIcuel4 (1)	SIcuel5 (1)	SIcuel6 (5)	SIcuel7 (1)	SIcuel8 (1)	SIcuel9 (2)	SIcuel10 (1)	SIcuel11 (1)	SIcuel12 (1)	SIcuel13 (1)	SIcuel14 (2)	SIcuel15 (2)	SIcuel16 (1)	SIcuel17 (6)	SIcuel18 (1)	SIcuel19 (1)	SIcuel20 (2)	SIcuel21 (3)	SIcuel22 (1)	SIcuel23 (1)	SIcuel24 (1)	SIcuel25 (1)	SIcuel26 (2)	SIcuel27 (2)	SIcuel28 (1)	SIcuel29 (1)	SIcuel30 (1)	SIcuel31 (1)	SIcuel32 (1)	SIcuel33 (4)	SIcuel34 (2)	SIcuel35 (1)	SIcuel36 (2)	SIcuel37 (2)	SIcuel38 (1)	SIcuel39 (1)	SIcuel40 (2)	SIcuel41 (3)	SIcuel42 (6)	SIcuel43 (1)	SIcuel44 (1)	SIcuel45 (3)
Attractor	Object of interest		[Grid of black cells representing attractor occurrences]																																												
Animation	None		[Grid of black cells]																																												
	Staged	Blink	Unique	[Grid of black cells]																																											
		Looped	[Grid of black cells]																																												
	Dynamic	Interpolated	Unique	[Grid of black cells]																																											
			Looped	[Grid of black cells]																																											
	System		[Grid of black cells]																																												
Trigger (event)	User	Payload/update	[Grid of black and gray cells]																																												
		Mousemove	[Grid of yellow and black cells]																																												
		Mouseover	[Grid of red and black cells]																																												
		Click	[Grid of blue and black cells]																																												
		Drag	[Grid of black cells]																																												
	Variable	Mousewheel	[Grid of yellow and black cells]																																												
Visual attribute	Extra mark	Non-textual	[Grid of black cells]																																												
		Textual	[Grid of black cells]																																												
		Persistence		Mouseover	[Grid of black cells]																																										
Intended interaction	Click		[Grid of black cells]																																												
	Drag		[Grid of black cells]																																												
	Mousewheel		[Grid of black cells]																																												
Feedforward			[Grid of black cells]																																												

Fig. 5: A design space of SI, based on our survey of 382 HTML5 and visualization websites. Due to space limitations, the table has been transposed, so entries are columns and dimensions are rows. The count of occurrences of each SI cue is shown in brackets, and the cue names refer to minified URLs (e.g., SIcuel1 can be retried at <http://tiny.cc/SIcuel1>). Note that several URLs direct to the same websites, as these include multiple distinct SI cues. Finally, an interactive version of the design space with animated GIFs of each cue is available at [1].

by system-events, which is why we have projected those that do not occur on pageload onto the system-event dimension in gray. As such, only dynamic animations are truly triggered by a user-event (which can occur anywhere on the page—yellow cells, with the exception of SIcuel45, which we discuss in the Mouse cursor cases below).

Missing step cases Sequenced interaction can also occur with a same graphical object. With regard to Buxton’s three-state model of graphical input [14], we had expected that the SI cues for such sequences would follow a specific order, *i. e.*, hover first, then click, then drag (e.g., SIcuel18 then SIcuel16). Standard hyperlinks do this, as they have an initial appearance that suggests a first interaction is possible, *i. e.*, mouseover, and a second appearance when a user hovers them that suggests another interaction is possible, *i. e.*, click. However, in several cases (e.g., SIcuel4, SIcuel6) the SI cues skip the initial steps, for example inviting users directly to perform a drag interaction. While in many cases these steps are implicit, we stress that for abstract interactive features, each step should be carefully considered, especially if the SI cue is not very ‘expressive,’ *i. e.*, if it does not explicitly indicate what interaction is expected. For example, if a cue does not convey the idea of dragging and is not sequenced when a user hovers or clicks on the feature, s/he may not move on to the next step, and will not understand the purpose of the interactive graphical object. Similarly, in some cases of sequenced interactions with different graphical objects, we found no SI cue for the first object the user must interact with in order to trigger the SI cue for the second object (e.g., SIcuel8, SIcuel9).

Mouse cursor cases In some cases, the mouse cursor itself is used as an SI cue (e.g., SIcuel45); the design of the cursor is modified to

indicate a specific interaction the user can or should perform (e.g., click-and-drag, with a hand icon and arrows indicating the direction in which the content can be dragged). These cases were particularly hard to code, as they can either be considered as inanimated (*i. e.*, no animation) attractors, or dynamic animation attractors. Here, we decided to code them as the latter, since obviously the mouse cursor is affected by the user-event mousemove, but it shows an exception since the cue is visible (or triggered) on pageload (a system-event). This kind of cue is often used when a ‘whole page’ interaction is dependent on clicking and dragging, like a swipe on a touch device.

Zoom widgets Another ambiguous case was found in interactive maps with zoom-in-and-out widgets (e.g., SIcuel43). These can either be considered as standard widgets that can be directly manipulated (by pressing on a + or – button, or by using a slider), in which case they are not really SI cues; or as cues for suggesting a mousewheel interaction, since zooming commonly relies on using the mousewheel. In our design space, we consider these widgets as the latter, and have coded them as external object attractors that invite users to scroll.

Combinations Sometimes, both an object of interest attractor and an external object attractor are used simultaneously. The intended interaction is often a mouseover, which highlights a specific region of the display and shows a tooltip (e.g., [17]). Such combinations are very effective for providing feedforward. However, for the purpose of our design space, we have separated out the individual SI cues used in these combinations (e.g., SIcuel19).

Misleading cases We encountered two cases in which SI cues were misleading. The first was in an NY Times graphic (SIcuel5) that

presents an interactive 3D map. The SI cue is a staged animation triggered on pageload that scales the map (the object of interest) into view, suggesting the possibility to zoom in and out. However, no mouse-wheel interaction is implemented for this purpose; only a click-and-drag is possible, but it only rotates the map in 3D space. The second was in another NYTimes graphic (SIcue7) that presents what seems to be a slider. The SI cue is a staged animation triggered on pageload that moves the slider thumb to a specific location, suggesting the possibility to click-and-drag it. However, the interactive area is actually a series of buttons which only allow for clicking; this SI cue provides a false idea of a continuous scale.

Edge cases Finally, we encountered two edge cases, which required careful consideration for integration in our design space. The first was buttons, which like the zoom widgets discussed above, can either be considered as standard widgets, or as external object attractors. We consider that when the intended interaction can be performed in a region beneath or around a button (like in the case of a “play” button on top of a video where a user can click anywhere on the video frame to play the content), then the button is an external object attractor (e.g., SIcue42). However, if the button requires clicking on directly to trigger something, then it is a widget, and we do not included it in our design space. The second edge case was found in [39], and is difficult to classify as an SI cue, as it relates to the layout of the interactive elements on the page. The interactive features on this webpage are quite large and centered horizontally below the page-fold so that users, while scrolling with the mousewheel, accidentally end up hovering the features; this then triggers another SI cue inviting users to click (SIcue11), creating a sequenced interaction with the same feature. Interestingly, this example can be related to the hypothesis we formulated at the end of Experiment 1 (see Section 3.1.3), in which we supposed that the layout of features could be used to help discover interactions. However, we consider this **layout** dimension outside of our present scope, as it only relates to very specific edge cases.

4.2 Design Considerations for Visualizations

From this design space, we derived several considerations to operationalize the creation of SI cues for visualizations. While such low-level deconstructions are useful for describing existing designs, we find they are often too complex when it comes to creating new ones. As we are interested in suggesting the interactivity of charts embedded with text, we propose it is possible to use either the visualization itself as an attractor, *i. e.*, the object of interest, or an overlaid icon, *i. e.*, an external object. In this subsection, we discuss our considerations, and introduce several metaphors which we believe may assist designers.

4.2.1 Visualizations as attractors

Setting the visualization as the attractor limits the possible visual attributes and animations that can be used. Indeed, visualizations already depend on visual marks and variables to encode data, so those used for the SI cue should not overlap or interfere. In the case of an inanimated visualization (*i. e.*, an object of interest attractor with no animation), simply playing with *free* visual attributes [34] should be avoided, as this could be considered more a stylistic choice than an invitation to interact. For example, using a red hue instead of a blue hue for a bar chart is unlikely to be more effective for suggesting interactivity. Thus, using the visualization as the attractor requires applying a staged or a dynamic animation.

For staged animations, we propose the metaphor of *organic motion*, which consists in small repetitive animations, simulating the motion of organic processes to which we are “*intimately familiar*” [29]. Organic motions range from a heart beat to a timelapse of a blossoming flower. For dynamic animations, we propose the metaphor of *attractive motion*, which can consist of orienting, squeezing, or stretching a visualization, depending on how far the mouse cursor is from it. A typical example is SIcue24. Attractive motions range from a cat’s head following the trajectory of a moth, to the orientation of a sunflower according to the sun, or of a metallic object attracted to a magnet. These different motions can be applied to any visual attribute of the visualization, and can be persistent or not—although we do recommend

stopping the animation when the user is interacting with the visualization, as this can be distracting and annoying.

4.2.2 Icons as attractors

External object attractors are generally icons, which may or may not be accompanied by text. These can be animated or not.

Inspired by Baecker *et al.*’s considerations [8], we identify three kinds of icons: focal icons, identifier icons, and demonstrator icons. The first two generally use no animation, while the third uses staged animations. A *focal icon* is an icon displayed on top of a multimedia artifact like a video. When the artifact is ‘out of focus,’ *i. e.*, when the user is not interacting with it, the icon is shown. When the artifact is ‘in focus,’ the icon is removed. A typical example is the “play” button displayed on top of a video (e.g., SIcue42). An *identifier icon* is usually an icon displayed next to an interactive feature with a textual label indicating what should be done (e.g., “navigate years” in [22]—SIcue40). However, identifier icons can also be used to replace the mouse cursor, in which case they are dynamically animated (as discussed in the Mouse cursor cases in Section 4.1.1). Finally, *demonstrator icons* are generally icons that incorporate animated graphics or kineticons (see Section 2.4); they show the user what to do in tutorial-like fashion (e.g., SIcue31).

5 TESTING THREE SI CUES APPLIED TO BAR CHARTS

To make an initial assessment of the effectiveness of SI cues applied to bar charts embedded with text, we generated a series of examples (available at [12]) and we tested the three we believe to be most representative of the diversity of our design space in a follow-up of Experiment 1. The first cue (SI₁) we tested uses the object of interest (*i. e.*, the visualization) as the attractor; the second (SI₂) uses an external object attractor; and the third (SI₃) uses a combination of both to provide feedforward (as is done in [17]—see paragraph on Combinations). In this section, we first describe the design of these cues and present their evaluation. We then provide some initial recommendations for design. Note that while it should prove interesting to test the full spectrum of variations that can be generated from our design space to see which are most effective, here we simply intend to assess whether SI cues actually have an effect on users’ propensity to interact with charts.

5.1 Three SI Cue Designs

SI₁ uses the visualization as the attractor and applies an organic motion to it (see [2]). We simulated a heart beat that first slowly stretches out, then bounces back into its original state. This staged animation is looped, triggered on pageload, and applied to the width of the bar chart. The cue is not persistent, *i. e.*, it is removed as soon as the chart is hovered over.

SI₂ uses a *focal icon* as the attractor (see [3]). This respects the considerations mentioned in Section 4.2.2, and shows an open hand to suggest manipulation. The cue is persistent, as it is displayed again when the visualization is ‘out of focus.’

Finally, SI₃ uses both the visualization and a *demonstrator icon* as attractors (see [4]). For the visualization, we sequentially highlighted different bars using a looped blink animation. For the demonstrator icon, we mimicked a black pointer cursor to which we applied a looped staged animation to its horizontal position, simulating a brushing interaction. We also added an extra textual mark, *i. e.*, the label for the highlighted bar, to provide feedforward. This way, users have a sense of what they will find when performing the interaction. The cue is not persistent.

5.2 Experiment 4

5.2.1 Design

To assess the effectiveness of these SI cues, we conducted a follow-up study on AMT. We reproduced Experiment 1 three times, respectively applying SI₁, SI₂, and SI₃ to the bar charts. We then used the results of Experiment 1 as a baseline for comparison.

Three groups were tested, each with one of the SI cues, in a between subjects design. Group 1 (G₁) was assigned SI₁, Group 2 (G₂) was

assigned SI_2 , and Group 3 (G_3) was assigned SI_3 . The scenario, tasks, and procedure were kept exactly the same as in Experiment 1.

Hypotheses We had the same simple hypothesis for each group:

- **H4:** participants perform more brush interactions and decisive brushes when an SI cue is applied to the charts.

Participants For each group, we recruited 40 different participants, making sure they had not participated in our initial studies. We retained the work of 33 participants in G_1 whose total score was higher than 0; of 35 in G_2 ; and of 40 in G_3 (this was the only group in which all scores were above 0). All participants were native english speakers. We then coded the data in the same way as in our initial experiments.

5.2.2 Results

For each group, we calculated the difference between proportions of participants who performed brush interactions in this experiment and in Experiment 1; we did the same for decisive brushes. Results are shown in Figures 6, 7, and 8 (for each group, respectively).




Fig. 6: Results for G_1 .




Fig. 7: Results for G_2 .




Fig. 8: Results for G_3 .

5.2.3 Discussion

Our results do not support **H4** for G_1 and G_2 . The fact that the 95% CI are well below 0 shows no evidence that adding SI_1 and SI_2 to the charts enticed more users to interact, or for that matter to use the charts for finding the answers (Figures 6 and 7). However, **H4** is confirmed for G_3 : SI_3 successfully incited more participants to interact, and to perform decisive brushes (Figure 8). We hypothesize that the success of SI_3 is due to the fact that it provided feedforward. As shown in Experiment 2, people may need a short amount of time to learn how to use the charts, and we believe the feedforward helped them identify the benefit of interacting with the charts. However, it should be noted that this was ‘heavy’ design, as it required a combination two SI cues: one applied to the object of interest, and one applied to an external object attractor.

5.3 Initial Recommendations for Design

Overall, it seems that providing SI is necessary, especially when visualizations are designed for online audiences who may not be accustomed to the interactivity or different interaction techniques information visualizations may provide. While we have focused here on visualizations embedded with text, we strongly believe the same applies to other ‘independent’ online visualization applications. More people may expect these to be interactive, but there are no real conventions that can help them identify what can be done, *i. e.*, *how* to

interact with the display; and interactions with visualizations are generally more advanced than those required for other web-based media. Note that in our experiments, participants were only expected to discover a hover interaction, which can be considered as the simplest kind of interactions—as it can be performed ‘involuntarily’—and even this was problematic.

Concerning design, while our results are only preliminary, we have found that our more subtle cues (*e. g.*, SI_1 and SI_2) were unhelpful, so similarly to hyperlinks, we believe a somewhat ‘heavy’ approach is necessary (*e. g.*, SI_3). External object attractor SI cues can be combined with object of interest SI cues to provide feedforward, which can show a user what s/he can or should do with a visualization. A simple way of implementing such cues could be to create situated animated GIFs on the webpage (as done in [45]). However, we stress that this is not an immutable guideline: more work is needed on the evaluation of SI cues applied to visualizations. This could reveal that other/more subtle techniques and/or adjustments may be just as effective. For example, our initial experiments have led us to consider that for simple hover interactions, the position of the visualization in the webpage might have an effect. This should be properly tested.

Finally, it is still unclear how much interaction can or should be suggested. As a broader guideline, we encourage designers to aim for simple yet effective interaction techniques (*e. g.*, to implement complex infovis interaction techniques using only the array of standard interactions people usually perform on the web), as the current levels of *interaction literacy* and/or *propensity* seem generally low.

6 CONCLUSION & FUTURE WORK

In this article, we have shown that most people lack initial propensity to interact with visualization when these are embedded with text. To address this issue, we have introduced the concept of Suggested Interactivity, and have presented a survey of 382 HTML5 and visualization websites, in which we identified 45 distinct cues used to suggest the interactivity of abstract features. From this survey, we extracted a set of important dimensions for the design of SI cues, and have constructed a design space. We have then evaluated the benefit of using three representative SI cues for visualizations embedded with text, and have shown that an SI cue that provides feedforward can successfully entice more users to interact with charts. Our results also suggest that while certain people may lack initial visualization literacy, this problem can be rapidly overcome when questions and charts are highly-congruent.

We see four main avenues for future work on SI. The first concerns extending our evaluation of SI cues to establish whether other/more subtle cues can be as effective as SI_3 . The second concerns testing the cues we designed with different tasks. Those used in this article were very specific and focused, *i. e.*, fact-checking tasks. We intend to continue exploring the effectiveness of our SI cues for more open-ended tasks. The third concerns using SI cues for more complex visualizations, which are not embedded in text. While most visualizations still rely on widgets to perform dynamic queries, many now also propose direct manipulation techniques which are applied to the visual representation itself. These interactions need to be suggested to users, and we believe SI can be an effective means for doing so. Finally, the fourth concerns extending our design space by evaluating the ‘expressiveness’ of cues, *i. e.*, how well they communicate the intended interaction to a user. This will require a more qualitative approach.

Ultimately, we consider the need for SI may simply be a ‘transition’ phase: animated icons were necessary for a time in graphical user-interfaces, but have now mostly disappeared as users have become accustomed with such designs. Nevertheless, we believe it is important to guide users through this transition phase to accelerate the adoption of information visualization for the people.

ACKNOWLEDGMENTS

The authors thank Ferdinand Dervieux and Jérémie Lasnier for their initial help with the survey; they also thank Google for funding this work with a Research grant on Visualization for the People.

REFERENCES

- [1] http://peopleviz.gforge.inria.fr/trunk/SI_cues/.
- [2] http://deputeviz.fr/other_apps/SI/wikipedia_template/example.php?t=1&s=1.
- [3] http://deputeviz.fr/other_apps/SI/wikipedia_template/example.php?t=1&s=2.
- [4] http://deputeviz.fr/other_apps/SI/wikipedia_template/example.php?t=1&s=3.
- [5] American Psychological Association. *The Publication manual of the American psychological association (6th ed.)*. Washington, DC, 2010.
- [6] <https://developer.apple.com/library/ios/documentation/UserExperience/Conceptual/MobileHIG/>.
- [7] <http://www.awwwards.com/>.
- [8] R. Baecker, I. Small, and R. Mander. Bringing icons to life. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '91*, pages 1–6, New York, NY, USA, 1991. ACM.
- [9] <http://en.wikipedia.org/wiki/Help:Barchart>.
- [10] L. Bartram, C. Ware, and T. Calvert. Moticons: detection, distraction and task. *Int. J. Hum.-Comput. Stud.*, 58(5):515–545, May 2003.
- [11] <http://bestwebgallery.com/category/html5/>.
- [12] J. Boy. Examples of suggested interactivity. http://peopleviz.gforge.inria.fr/trunk/SI_examples/.
- [13] J. Boy, R. A. Rensink, E. Bertini, and J.-D. Fekete. A Principled Way of Assessing Visualization Literacy. *IEEE Transactions on Visualization and Computer Graphics*, pages –, Nov. 2014.
- [14] W. Buxton. A three-state model of graphical input. In *Proceedings of the IFIP TC13 Third International Conference on Human-Computer Interaction, INTERACT '90*, pages 449–456, Amsterdam, The Netherlands, The Netherlands, 1990. North-Holland Publishing Co.
- [15] A. Canty and B. Ripley. *Bootstrap Functions*, 2014.
- [16] S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1999.
- [17] http://codecanyon.net/item/pixel-map/full_screen_preview/1243869.
- [18] G. Cumming. The new statistics: Why and how. *Psychological science*, 25(1):7–29, 2014.
- [19] P. Dragicevic, F. Chevalier, and S. Huot. Running an HCI Experiment in Multiple Parallel Universes. In *CHI '14 Extended Abstracts*, pages 607–618. ACM, 2014.
- [20] R. S. Easterby. The perception of symbols for machine displays. *Ergonomics*, 13(1):149–158, 1970. PMID: 5416867.
- [21] http://englishteststore.net/index.php?option=com_content&view=article&id=241&Itemid=285.
- [22] <http://www.evoenergy.co.uk/uk-energy-guide/>.
- [23] <http://www.gapminder.org/downloads/>.
- [24] J. Gibson. *The Ecological Approach to Visual Perception*. Resources for ecological psychology. Erlbaum, 1986.
- [25] S. Gilroy, J. Porteous, F. Charles, and M. Cavazza. Exploring passive user interaction for adaptive narratives. In *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces, IUI '12*, pages 119–128, New York, NY, USA, 2012. ACM.
- [26] D. Gittins. Icon-based human-computer interaction. *Int. J. Man-Mach. Stud.*, 24(6):519–543, June 1986.
- [27] <http://www.google.com/design/spec/material-design/introduction.html>.
- [28] <http://www.theguardian.com/data>.
- [29] C. Harrison, G. Hsieh, K. D. Willis, J. Forlizzi, and S. E. Hudson. Kinetics: using iconographic motion in graphical user interface design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, pages 1999–2008, New York, NY, USA, 2011. ACM.
- [30] H. R. Hartson. Cognitive, physical, sensory, and functional affordances in interaction design. *Behaviour and Information Technology*, 22:315–338, 2003.
- [31] J. Heer and B. Shneiderman. Interactive dynamics for visual analysis. *Commun. ACM*, 55(4):45–54, Apr. 2012.
- [32] <http://html5gallery.com/>.
- [33] S. M. Huang, K. K. Shieh, and C. F. Chi. Factors affecting the design of computer icons. *International Journal of Industrial Ergonomics*, pages 211–218, Apr. 2002.
- [34] Y. Jansen and P. Dragicevic. An Interaction Model for Visualizations Beyond The Desktop. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2396–2405, Nov. 2013.
- [35] K. Lodding. Iconic interfacing. *Computer Graphics and Applications, IEEE*, 3(2):11–20, March.
- [36] S. Manes. Pushing picture-perfect programs: Smash that icon! *PC Magazine*, page 64, June 1985.
- [37] <http://collection.marijerooze.nl/#!/>.
- [38] <http://blog.jerrynixon.com/2012/03/windows-8-does-metro-even-make-sense-to.html>.
- [39] <http://www.michelbergerbooze.com/>.
- [40] [https://msdn.microsoft.com/en-us/library/aa733613\(v=vs.60\).aspx](https://msdn.microsoft.com/en-us/library/aa733613(v=vs.60).aspx).
- [41] P. Neri, M. C. Morrone, and D. C. Burr. Seeing biological motion. *Nature*, 395(6705):894–896, Oct. 1998.
- [42] D. A. Norman. Affordances and design. http://www.jnd.org/dn.mss/affordances_and.html.
- [43] D. A. Norman. *The design of everyday things*. Basic Books, New York, 2002. Previously published under : The psychology of everyday things, 1988.
- [44] OECD. Oecd better life index. <http://www.oecdbetterlifeindex.org/>.
- [45] C. Perin, R. Vuillemot, and J.-D. Fekete. À Table! Improving Temporal Navigation in Soccer Ranking Tables. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'14)*, Apr. 2014.
- [46] Z. Pylyshyn, J. Burkell, B. Fisher, C. Sears, W. Schmidt, and L. Trick. Multiple parallel access in visual attention. *Can J Exp Psychol*, 48(2):260–83, 1994.
- [47] B. J. Scholl and P. D. Tremoulet. Perceptual causality and animacy. *Trends Cogn Sci*, 4(8):299–309, Aug. 2000.
- [48] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1139–1148, Nov. 2010.
- [49] R. Spence. *Information Visualization: Design for Interaction (2nd Edition)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2007.
- [50] J. Tang, X. Zhao, and P. Zhang. Perceived affordances of web advertisements: Implications for information artifacts design. In *Proceedings of the Fifth China Summer Workshop on Information Management, CSWIM*, 2011.
- [51] B. Tversky, J. B. M. Y, and M. Betrancourt. Animation: Can it facilitate. *International Journal of Human-Computer Studies*, 57:247–262, 2002.
- [52] J. Vermeulen, K. Luyten, E. van den Hoven, and K. Coninx. Crossing the bridge over norman’s gulf of execution: Revealing feedforward’s true identity. In *The ACM SIGCHI Conference on Human Factors in Computing Systems, CHI '13*, 2013.
- [53] B. Victor. Magic ink: Information software and the graphical interface. <http://worrydream.com/#!/MagicInk>, 2006.
- [54] M. Wattenberg. Baby names, visualization, and social data analysis. In *Information Visualization, 2005. INFOVIS 2005. IEEE Symposium on*, pages 1–7, 23–25 2005.
- [55] <http://www.webdesignerdepot.com/2014/10/what-web-designers-can-learn-from-googles-material-design/>.