

Exploring Interactions with Physically Dynamic Bar Charts

Faisal Taher¹, John Hardy¹, Abhijit Karnik¹, Christian Weichel¹, Yvonne Jansen²,
Kasper Hornbæk², Jason Alexander¹

School of Computing and Communications, Lancaster University, United Kingdom

{f.taher, j.hardy, a.karnik, c.weichel, j.alexander}@lancaster.ac.uk

Department of Computer Science, University of Copenhagen, Denmark

{yvja, kash}@di.ku.dk

ABSTRACT

Visualizations such as bar charts help users reason about data, but are mostly screen-based, rarely physical, and almost never physical and dynamic. This paper investigates the role of physically dynamic bar charts and evaluates new interactions for exploring and working with datasets rendered in dynamic physical form. To facilitate our exploration we constructed a 10×10 interactive bar chart and designed interactions that supported fundamental visualisation tasks, specifically: annotation, navigation, filtering, comparison, organization, and sorting. The interactions were evaluated in a user study with 17 participants. We identify the preferred methods of working with the data for each task (e.g. directly tapping rows to hide bars), highlight the strengths and limitations of working with physical data, and discuss the challenges of integrating the proposed interactions together into a larger data exploration system. In general, physical interactions were intuitive, informative, and enjoyable, paving the way for new explorations in physical data visualizations.

Author Keywords

Shape-changing interfaces; physical visualizations; tangible user interfaces; information visualization; shape displays.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Effective visualizations help users “*use vision to think*” [3] and aim to improve reasoning about data. They increase the effectiveness of our ability to process information by transforming it into visual structures that leverage our biological ability to detect patterns and trends. While today's visualizations are optimised for 2D screens and computers, the information visualization community has begun to

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version of Record was published in CHI '15 which can be found at:

<http://dx.doi.org/10.1145/2702123.2702604>

discuss how to design 3D visualizations for physical devices and non-traditional inputs [15, 17, 18, 23].

As part of this effort, researchers have investigated the benefits of 3D physical charts in comparison with 3D on-screen visualizations [16] and found that the rich qualities of physical objects can play an important role in the data inspection process. Huron et al. [13] showed how people construct, manipulate, and update visualizations based on tangible tokens. Other examples include data sculptures [1] and tactile cartographic maps [27]. The main drawback of current physical visualizations is that they are inert, being either fabricated (i.e. laser cut [15], 3D printed [3], constructed from passive building blocks [13]) and thus disconnected from the data-source once constructed. Shape-changing interfaces (i.e. Relief [22], inFORM [6]) have the potential to alleviate these drawbacks. However, the community lacks an understanding of how data can be interacted with to achieve effective and intuitive data exploration with physically dynamic displays.




Figure 1 – EMERGE using actuating physical rods and RGB LEDs to display international export data.

This paper presents a first exploration of user interactions with data using our custom built dynamic bar chart: EMERGE (Figure 1). EMERGE contains self-actuating rods capable of RGB colour output and touch detection for pushing and pulling of the data itself, in addition to traditional touch detection on the surrounding surface. As a first step towards guidelines for designing physically dynamic data we carried out a user study with 17 participants and investigated fundamental interaction tasks that are common in data exploration scenarios: annotation, filtering, organisation, and navigation. Within these we

explore directly manipulating data points (e.g. by pulling to select data points), as well as using gestures on projected axis labels (e.g. moving rows of data by dragging labels). These explorations form the baseline for more complex and diverse interaction techniques.

The key contributions of the paper are (a) the identification and design of 14 baseline interaction techniques designed to be used with physically dynamic bar charts, (b) a user study that evaluates these interactions, determining ones that are most preferred and effective, and (c) a discussion of important design considerations and challenges for physically dynamic visualizations.

RELATED WORK

Physical data visualizations, or ‘Physicalizations’ [17], are informed by two strands of research: visualization [15, 25] and tangible and shape-changing interfaces [22]. Below we review these strands to argue that few papers have used dynamic, physical interfaces to visualize data, and that no study we are aware of has investigated the possibilities and usefulness of interactions for such interfaces. We also outline the research questions for this paper.

Static Physical Visualizations

Physical visualizations aim to extend the benefits of visualizations by tapping into active perception skills and integrating sensory information in addition to the visual sense. Jansen and Dragicovic [16] curate a *list of physical visualizations*¹ that includes examples more than six millennia old. These are mostly used to show quantity in physical form, for instance by mapping a number to physical height. Panels and Roberts [26] reviewed approaches that focus on haptic data visualization, such as using audio, texture/friction, and enclosures to show quantity.

Previous research has examined the efficiency of physical visualisation [15] and approaches for designing physical visualizations [35]. Jansen et al. [15] found that hand-held 3D printed physical visualizations improved users’ efficiency at information retrieval tasks, with physical touch and visual realism key being advantages. Stusak and Aslan examined various physical visualization prototypes and found that these representations can support analytical tasks through mature design, emphasising the importance of stability and affordances [35].

Physical visualizations can also increase the accessibility of data to blind or low-vision users [3, 7, 19, 24]. Examples include tactile pin arrays to show graphics [39], VizTouch [3] that allows blind users to 3D print visualizations of line graphs, and tactile maps [27] that show cartographic data using physical properties (physical height corresponds to elevation of the terrain).

Other physical visualizations have more artistic aims. The term ‘data sculptures’ [1] describes visual pleasing artefacts that communicate information, such as jewellery that shows internet connection rates. Khot et al. [30] describe Sweatatoms, a 3D modeling and printing system that can turn activity patterns in sports into 3D objects that support reflection and aesthetic pleasure. Stusak et al’s ‘activity sculptures’ artistically visualised running activity for discussion and reflection [36].

Physical visualizations, such as those described above, have a range of potential benefits over their purely visual counterparts [16, 25, 38]. First, as physical objects they can be manipulated more directly than through a mouse or even through touch screens [16]. Second, the interplay of vision and touch can facilitate cognition [15]. Finally, the physical modality opens a range of new interaction possibilities compared to on-screen visualizations.

Dynamic Physical Visualizations

While static, physical visualizations are useful and attractive, they must be fabricated before use. Modifications of their physical form are often limited once created and many of the computational and interactive benefits of screen-based visualizations are lost. Work into tangible user interfaces (TUIs) and shape-changing interfaces has attempted to address these disadvantages by showing non-static data.

Shape displays typically have a physical equivalent of pixels, either binary (on/off) or continuous (being able to show a range of values). Typical examples include Sublimate [21], Relief [22, 23], Lumen [28], Feelex [14], Taxel [18], inFORM [6], Tilt Displays [1], and Physical Charts². These physical pixels are often implemented via motorized pins that can extrude from a surface; pneumatics [10] and shape-memory alloys [5] can serve a similar purpose. These extrusions can be mapped to data values. Resolution varies from a few (<10) pixels to inFORM’s 900 motorized pins. Although these papers mainly demonstrate systems, they do show tasks relating to visualization, including showing bar chart data (Physical Charts), mathematical functions [6] and wind tunnel flow [21].

The interactions for shape displays often use well-known input, such as direct touch, and pulling and pushing [6, 23]. Relief [23] supports free-hand gestures to interact with 3D models: the user performs hand gestures to translate, scale, and rotate models. In addition to direct touch, inFORM [6] supports remote gestural interaction and the use of on-surface objects to control interactions such as menu selections. While free-hand gestures are useful in many scenarios, they do not capitalise on the rich haptic dimension touted as a key benefit of physical visualizations,

¹ <http://dataphys.org/list> (last accessed 07/01/2015)

² <http://research.microsoft.com/en-us/um/cambridge/projects/physicalcharts/> (accessed 23/12/2014)

nor do they support the delicate manipulation of data observed in by Jansen et al. [16].

Open Questions for Physically Dynamic Bar Charts

The above literature instigates several research questions regarding physically dynamic bar charts. In particular, the following two questions are crucial for interaction design. First, it is an open question how physically dynamic bar charts should support tasks in visualization, such as comparing specific values [2] or gaining an overview of a dataset [33]. Previous work has suggested multiple task models for use in visual data analysis [2, 11], yet these are directed towards 2D data visualizations.

Second, it is unclear which interactions with physically dynamic bar charts are useful and usable. In contrast to previous work on static physical visualizations and data sculptures, support for user interaction is as crucial as the ability of the interface to actuate itself. Further, the key aim of these artefacts is to help users think about data and the focus is *on* the visualizations, rather than on the control. This poses questions such as: are the useful interactions with static physical visualizations (e.g., [16]) transferrable to dynamic physical visualizations? Thus, identifying useful interactions with physical dynamic graphs is still very much an open question, beyond the straightforward uses of touch, pushing/pulling, and mid-air gestures.

IMPLEMENTATION

To investigate data analysis tasks and interactions, we developed a dynamic physical bar chart: EMERGE. Our design follows that of the inFORM system [6], differing through details such as the inclusion of LEDs, hardware control architecture, and software model.




Figure 2 – The setup of the EMERGE system.

Hardware

The hardware consists of a 10 × 10 array of actuated plastic rods that are individually linked to 100 motorized potentiometer sliders. The sliders provide 100mm travel length. The linkage consists of push-rods supported by

hollow tubes such that the sliders at the bottom are able to control the plastic rods at the top.

The sliders are stacked in two layers to minimize the size footprint (0.55×0.44×1.20m). Each rod is illuminated by a dedicated RGB LED (WS2812B) attached to rod guides, which are designed to keep them vertically stable. Additionally, a Microsoft Kinect® and projector were mounted above the system to project information (i.e. axes, labels, and controls) and detect touch interaction on the surface surrounding the rods.

Software

The EMERGE software stack is shown in Figure 3. The overall software architecture separates three logic layers: application, model, and firmware. The separation between application and model uses the model-view-controller pattern to allow multiple applications to connect simultaneously. This was useful in the debugging and development of applications.

The application layer consists of multiple user study applications and a 3D viewer/simulator. The study applications are written in HTML5 and hosted by the Ubi Displays toolkit [9] to map the projection and detect touch on the surrounding surfaces. The EMERGE API (model) is written in C# and used to manage the state of the actuators and LEDs. It also handles calibration, animation, and touch detection, which was based on the changes in potentiometer readings (thereby enabling push and pull interactions). The model communicates with the EMERGE Driver—an ABI (Application Binary Interface) that interfaces with the control firmware for sliders and LEDs.




Figure 3 - The EMERGE software stack. Application layer: orange; Model layer: blue; Firmware: green; hardware: red.

The control firmware consists of 2 Arduino Mega2560 (sliders) and a single Arduino Uno (LEDs). These send and receive data frames from the ABI via USB. Connected to the Mega2650s, are 17 bespoke motor driver boards. Each driver board carries 3 ATTiny84 microcontrollers that control the position of 6 sliders at ≈8Mhz using a PID controller. When selected on an SPI bus, these update the PID set point and report byte position values back to the ATmega2560s.

APPROACH

To reason about effective methods for exploring data with dynamic physical visualizations, we first needed to evaluate

interactions that are specific to this type of display. We present a series of novel methods for interacting with physical data and test whether they are useful and usable. Our investigations are based on Heer and Shneiderman’s taxonomy [11], and explore different types of commonly used visualization tasks such as annotation, organization, filtering, and navigation. A user study was structured to elicit formative feedback about the usefulness and usability of specific features, capturing initial reactions and experiences. We specifically avoided early experimentation as others have warned against this approach [8, 12], especially given the immaturity of this area.

Task Structure

We focus on ‘physicalizing’ data using bar charts as they map naturally to our hardware setup and have been used for previous studies on physical visualizations due to their low requirements for visualization literacy [15]. We base our exploration on Heer and Shneiderman’s taxonomy for visual data analysis [11]. This taxonomy structures the required interactive dynamics for visual analysis tasks into three high-level categories: *Data and View Specification*, *View Manipulation*, and *Process and Provenance*. These break down into sub-categories, which we utilize to derive tasks that are relevant and applicable in a physical space.

Table 1 provides an overview of the task-sets that we adopted from Heer and Shneiderman’s taxonomy (expressed in parentheses) of visual data analysis tasks [11]. Each task-set consists of multiple tasks/interactions.

Table 1 – Task-sets and interaction techniques explored during the user study.

Task	Overview	Interaction Techniques
Annotation (<i>Process & provenance</i>)	Selecting and marking individual data points.	Point, pull, press.
Filtering (<i>Data view & specification</i>)	Hiding and refining data for enhanced perception and comparison.	Swipe away, manual press, assisted press, press shortcut, and press to compare.
Organization (<i>View manipulation</i>)	Data arrangement by moving rows and columns.	Drag and drop with immediate transition and hide-all with transition, press with instant transition and hide-all with transition.
Navigation (<i>View manipulation</i>)	Controlling the view of large data sets.	Scroll, directional arrows, directional press, and paging.

Tasks such as *Coordinate* from *View Manipulation*, and *Share* and *Guide* from *Process and Provenance* are not applicable to our setup, and are therefore excluded. *Coordinate* typically concerns multiple views of data, whereas we experiment with one view. *Sharing* involves actions such as exporting data, which is software-based, and *Guiding* leans towards helping users through workflows. Complex task explorations such as *Sort* and *Derive* from *View Manipulation* are, for now, excluded as we first need to understand basic interactions rather than, for instance,

exposing patterns and forming data models. Further, *Record* typically involves showing historical data (e.g. through undo or redo features) and forms part of a larger ‘data explorer’ system, rather than a fundamental interaction.

Participants

The user study was carried out with 17 participants (6 female) with a mean age of 27 years. None of the participants had previous experience interacting with shape-changing displays. Two participants had previously seen demonstrations of shape-changing technology. Each session lasted approximately 40 minutes and participants were compensated £8 for their time.

Materials

We utilized the EMERGE system (as described in the previous section), which enables users to interact with an actuating interface of LED-lit rods (representing the bars of bar charts). Information such as labels were projected around the bars of the shape-display (e.g. see Figure 4).

Procedure

Participants were welcomed individually and introduced to the EMERGE system. A short demographic questionnaire was initially provided. Participants were then asked to carry out a number of task-sets pertaining to annotation, organisation, filtering, and navigation. We explored a variety of interactions for each task-set, some of which were adapted from Jansen et al. [15] (e.g. physically selecting data points of interest). The data sets were downloaded from Jansen et al’s [15] materials page³, and included country indicator data. The data sets consisted of HIV prevalence, GDP percentage in exports, annual electricity consumption, and a UK rainfall dataset downloaded from the MET office website⁴ (as we required a larger data set for the Navigation task). The data shown to users was encoded as follows: each row is discerned by a unique colour, and the height of a rod represents a numerical value (i.e. the y-axis value). The study procedure included the following steps:

- To elicit initial thoughts and perceptions, and before showing any interactions for a task-set, participants were asked how they could achieve the task. This followed an informal guessability approach [40] to gather insights without being influenced by demonstrations. Participants were shown a sample data set with interactivity disabled.
- Interaction techniques were then demonstrated to participants. They were also asked to carry them out to build familiarity. A task was then verbally given to participants by the experimenter. We were interested in participant behaviours and feedback whilst carrying out

³ <http://www.aviz.fr/phys> (last accessed 7/01/2015)

⁴ <http://www.metoffice.gov.uk/climate/uk/datasets/> (last accessed 22/09/2014)

the interactions rather than in how accurately they were able to carry out the task.

- The ordering of task-sets and interaction techniques were counterbalanced in order to reduce the influence on subsequent tasks.
- Immediately after participants carried out a task, they filled out a 5-point Likert scale questionnaire, and asked to comment on what they found useful and problematic.
- At the end of a session, a short discussion was carried out in order to receive any additional insights.

Each session was recorded using audio and video to document participant feedback and interaction behaviours with the EMERGE system.

EXPLORING PHYSICAL BAR CHART INTERACTIONS

We used EMERGE as a test-bed for exploring user interactions with physically dynamic bar charts. Below, each of the task-sets (see Table 1) is described along with the results from the user study.

Annotation

Annotation (see Figure 4) allows users to select and mark data-points for later reference [11]. This is useful when users wish to document and return to subsets of data, or to communicate interesting observations. During the study, we asked participants to mark single data points using three techniques: point, pull, and press. The data set used for this task was the electricity consumption in 10 different countries (rows) between 1971 and 1998 (columns).

- *Point*: Users position one finger on the row label and a second finger on the column label projected beside the physical columns. All data points are dimmed except the intersection point, which is marked by a unique colour (e.g. white). Participants were asked to select the year and country with the highest electricity consumption in the complete data set.
- *Pull*: Users pull the data point that they wish to select. The selected data point becomes emphasized by dimming the unselected data points. Participants were asked to select the year in which Qatar had the highest electricity consumption.
- *Press*: Similar to *Pull*, however, users press the data point (similar to a tap) to select it. Participants were asked to select the year in which Iceland had the lowest electricity consumption.

Results and Feedback

Participants preferred the *press* technique for annotating data points (5-point *Likert M*: 4.29, *SD*: 0.99, see Figure 5). This also matched all 16/17 participants' initial perceptions of how they would select an individual data point. One participant stated they would prefer to pull a column in order to "pick out" the value.




Figure 4 - The annotation task. Here the user is using the *Point* technique to select a data point.




Figure 5 - Likert scale ratings for helpfulness of interaction techniques. Range = 1: Strongly Disagree, 5: Strongly Agree.

"This was my favourite technique as it was very simple to use. I like how only that particular point is highlighted with colour to make it easier to see. It feels more like I am actively selecting the data point I want to see, more than when I just touched the screen."

The visual feedback (i.e. dimming the unselected data points) also allowed participants to distinguish the selected data point from the others (as commented by 5 participants). Although users will be physically aware of selecting a data point, the visual confirmation (as feedback) is necessary to confirm the selection.

Filtering

Data filtering allows users to control the items displayed and enables focus by eliminating irrelevant data [33]. Filtering is particularly useful in a setup of physical bars. Zacks et al. [41] found that for height estimates of bar charts printed on paper, distortions introduced by neighbouring elements were most problematic. Therefore, clearing the display of unneeded bars can aid users to correctly read and compare values. Although Zacks et al.'s [41] observations were based on printed bar charts, we expect similar issues with physical bars. The data set used

for this task was the export percentage of goods and services in 10 different countries (rows) between 1999 and 2008 (columns). During the study we explored filtering irrelevant rows using the techniques below (see Figure 6):

- *Swipe away*: Swipe the projected row label off the edge of the surface to hide the row. Participants were asked to hide eight selected rows that were deemed irrelevant and to state a general comparison overview between the remaining two countries.
- *Manual press*: Directly and manually push down all irrelevant data points. Similar to *swipe away*, participants were asked to press down the data points of eight different rows and compare the GDP of the remaining two.
- *Assisted press*: Similar to manual push, except the system detects downward pressure and hides the pressed data point. Participants were asked to hide the row with the highest exports percentage.
- *Press to compare*: This technique explores pressing down on any data point on two rows, which then hides all other rows. Participants were asked to press and compare the GDP of two selected countries.
- *Press shortcut*: The data point at the beginning of each row acts as a control point. Once this data point is pressed, the remainder of the data in the row becomes hidden. Participants were asked to hide eight different rows and compare the GDP of the remaining two.




Figure 6 - Filtering task showing the *press to compare* technique.

Results and Feedback

All 17 participants initially stated that they would prefer to only tap the row labels that they are interested in observing and comparing, which would subsequently hide the other irrelevant rows. The study showed that the *press shortcut* technique was preferred by most participants (*Likert M*: 4, *SD*: 0.63). This was regarded to be especially useful when smaller numbers of rows are required to be hidden.

“I found this one the best way, as it was easy, you weren't relying on touch screen and you had to click the column closest to the name so it was good for accuracy.”

We also found that using the *press to compare* technique was preferred for comparing smaller data sets such as two rows (*Likert M*: 3.94, *SD*: 1.06). This builds on the premise that it is more efficient to select a small set of rows that are of interest, rather than hiding multiple irrelevant rows.

“I wanted this for the previous comparison task and it did not disappoint. It also helped cut down on the sheer amount of repetitive selection, since it makes more sense to select the 20% of objects I want rather than the 80% I don't.”

A noteworthy outcome of studying the filtering techniques is that, unlike participants' initial perceptions of how to filter data (i.e. tapping the projected labels), both highly rated interaction techniques involved physical interaction with the data points. Furthermore, it is clear that a tradeoff is required between hiding small amounts of data versus a larger amount of data.

Organization

This task-set explores user preferences for organizing data sets and points in the physical space (see Figure 7). Although Heer and Shneiderman's [11] taxonomy describe *Organize* in terms of multiple views, we adapt this concept and look at organization within a single view. For instance, users might want to bring some data closer to them (e.g. from the last row to the first row). The data set used for the organization task was prevalence of HIV in 10 countries (rows) between 1920 and 2012 (columns). The following interaction techniques were explored:

- *Drag and Drop*: Analogous to drag-and-drop on touch devices, users touch-down on a row label and drag it to a new position. Participants were asked to compare HIV prevalence between two countries by organizing the rows and placing them next to each other.
- *Press to swap*: Users simultaneously select a target row by tapping on a column, and select the destination by tapping on a column in a different row. Similar to *drag and drop*, participants were asked to compare two countries side-by-side.




Figure 7 – The organize task, showing the *drag and drop* technique.

In addition, we use the organisation task techniques to explore how users would like to see physical data transitions (i.e. visual feedback of the data points 'moving' to another location). Participants were asked to compare different rows (to the ones above) for techniques involving

the transitions. Both interaction techniques had the following transition options:

- *Instant transition:* New data simply replaces old data (bars immediately adjust to correct new heights).
- *Hide-all and transition:* All irrelevant rows hide at the beginning of a swap action; the selected and target rows swap values, then the other rows re-emerge.

Using the nominal dataset, participants were asked to compare various combinations of data by moving the comparison rows/points beside each other.

Results and Feedback

All 17 participants initially stated that they would prefer to drag a row by touch the projected labels to reorder them. The study confirmed their initial perceptions and it was clear that dragging the label was the preferred method (*Likert M* 3.88, *SD*: 1.22).

“This interaction has a sequential logic and this seems to facilitate this task ... It is also easier to look at the labels as opposed to interacting with the columns directly in rearranging data.”

The *hide-all and transition* technique scored low in both organize techniques (*drag and drop*, and *press to swap*). Participants preferred to have faster feedback rather than wait for a more drawn-out transition.

“I don't think the added animation was very helpful - prefer instant reaction. Also the travel of the row isn't the important bit.”

One participant also stated that in a professional setting, accidentally selecting the incorrect rows, and having to wait for the transition to finish, would be embarrassing.

Navigation

Navigation typically involves exploring data sets in more detail (e.g. geographic maps that present an overview, but also further details on data subsets). One of the limitations of developing shape-changing display is the high cost to achieve higher resolution due to the one-actuator-per-column architecture (e.g. as faced by inFORM [6]). Therefore, these displays require mechanisms to allow users to navigate large data sets. The rainfall data set was used for this particular reason during our study, showing 10 different regions within the UK (rows) between 1920 and 2012 (columns). This allowed participants to look at trends over time (i.e. 92 years) by navigating through the data. The following techniques were investigated (see Figure 8):

- *Scrollbars:* scrollbars, identical to those found on the desktop were projected on the x-axis. Tapping within the scrollbar trough moved the dataset to that location. All of the data points were also shown (to show continuity) until the selected position was reached. Participants were asked to identify any patterns and relationships that can be observed between the entire

dataset (i.e. 1920 to 2012) between regions in the North and South of the UK.

- *Directional Arrows:* Projected left and right arrows were shown on the EMERGE surface. Analogous to the desktop environment, tapping on the arrows moved the all data set rows by a single column. Participants were asked to compare rainfall between two regions during a specific year.
- *Directional Press:* All the data points on the right-half and left-half act as navigation mechanisms; when pressed, the data set shifts on column to the left or right. Participants were asked to compare rainfall between two regions during a specific year.
- *Paging:* Similar to *Direction Press*, all the data points on the right-half and left-half act as navigation mechanisms for shifting the data set by one page, that is, 10 columns to the left or right. The term paging is used because of the 10×10 grid of the EMERGE display. Participants were asked to compare rainfall between two regions during a specific year.




Figure 8 - The navigation task, showing the scroll technique.

Results and Feedback

The initial user perceptions stage showed that 15/17 participants prefer to navigate a large data set by using a scrollbar display combined with a swipe gesture. Intriguingly, two participants stated that they would prefer to physical shift the columns to the left or right to view the next or previous data points. The study showed that the scrollbar technique was preferred (*Likert M*: 4, *SD*: 1).

“It was very simple and you could scroll a lot quicker if you had to jump say, 40 years. It took a little while for the [scrollbar] to catch up with where you wanted it to be, but I found this interesting as you got to see the bar chart changing over time.”

Furthermore, by showing the actual values whilst the data is scrolling was useful for allowing participants to look at trends over time.

“The [transition] that followed the projected scrollbar was most useful as it also showed progress towards the intended date. The tangible interface indicated certain patterns

between the years which were interesting to observe (such as seasonality)."

Although the ratings were lower for the single-column navigation tasks, participants stated that this would be more useful for more fine-grained control of the data. It is therefore necessary to support continuous data navigation, as well as to allow more fine-grained control.

"Simpler for fine tuned detail. A combination of this and scroll bar would simulate the standard web scroll bar, this would cut down the learning curve."

DISCUSSION

The user study demonstrated a number of insights from user feedback and observed behaviours that indicate how bar charts can be effectively combined with shape-changing technology. We discuss the following themes: gestural vs. physical interaction, combining interaction modalities, effect of preconceptions, user reactions, technological challenges, limitations, and future directions.

Gestural vs. Physical Interaction

The interaction techniques explored during the user study comprised of directly touching the data points (or the plastic rods) as well as using gestures such as swiping to manipulate data. There was no clear difference in user preferences (i.e., the Likert scale ratings) for directly touching the rods and using gestures; each provides strengths and weaknesses based on the context of use. For instance, it was far more effective for participants to directly press and annotate a data point rather than to use gestures. In contrast, larger gestures such as row organization benefits from dragging their corresponding projected labels. The positive feedback received from physical interactions removes uncertainty about whether users might feel that they are interfering with the data points. This provides designers with more freedom to integrate different types of interactions.

Combining Interaction Modalities

We developed a number of insights on combining interactions for physically dynamic bar charts. For instance, we found that switching between overview and more fine-grained interaction modalities is highly useful. The *scrolling* technique in the navigation task revealed that all participants noticed a clear distinction over time between the northern and southern regions of the UK (i.e. more rainfall in the north). However, it is important to be able to switch to a mode that allows scrolling through one data point at a time (e.g. if users want to compare specific years). One participant commented that they would like to expand a data point showing rainfall by year and region into the months that make up the yearly value. Similarly, the filtering task necessitated the ability to combine the *press to compare* technique with *press shortcut* to hide a row. One participant's comment particularly highlights this:

"I did like this and felt it was easier to be sure I had selected the right rows for what I wanted to see. I would

like this method alongside being able to keep only the rows I wanted, depending on how many I wanted to keep/hide."

Effect of Preconceptions

Each participant was initially asked, before being shown any techniques for the tasks, how they would achieve, for instance, selecting an individual data point or organizing rows by moving the furthest one closer to them. All responses consisted of using swipe, drag or tap on a projected label, similar to the interactions on touch-screen devices. The prevalence of touch-screen devices have provided users with preconceptions of how to carry out interactions such as scrolling, selecting, etc. As a result, these interactions have become familiar and intuitive. An interesting outcome from the study concerns the filtering task, where participants initially suggested that actions such as using swipe or tapping on a row label would be preferred. However, after exploring direct touch interactions with the data points, the opinions shifted to the more physical side. Whilst designers should capitalize on the familiarity of touch-screen gestures and incorporate them into shape-changing displays, it is also important to realize that physical interactions can expand the interaction space in innovative ways.

User Reactions

As shape-changing displays are uncommon, we were interested in how users physically approach and react to this type of technology. During the study, we observed that nearly all participants were surprised or startled by the actuation (e.g. by moving back, moving their hands out the way), as well as from the noise generated by the motors. Three out of seventeen participants also moved around the system to carry out their interactions. For instance, one participant moved to the side of the display to select press down on two rows to compare them (filtering task). Another participant bent down to align themselves and select the highest value in a row (annotation task). We believe that movement around the system was likely reduced by the fixed text alignment of the labels, and therefore, to encourage users to move around the display, text orientation could be adjusted based on user position.

Limitations and Technological Challenges

For this first study, we chose to focus on fundamental low-level tasks and a limited number of interaction techniques. Further studies will be necessary to explore the possible interaction space for physically dynamic graphs (including different variations of actions and gestures, and combinations of tasks and techniques). The EMERGE system uses a grid of 10×10 data points and therefore only a limited amount of data could be presented to participants. Thus, the scalability of the interactions studied requires further investigation. A higher resolution might, for instance, afford different types of interactions (pagination might suffice rather than scrolling during navigation). Finally, we excluded vertical axis data (essentially the y-axis) and it is therefore difficult to anticipate how this might change user interactions and behaviours.

There were numerous technological challenges during the development of the EMERGE system, such as choosing an appropriate actuation speed, selecting an appropriate spacing for the plastic rods, as well as determining the size of the entire setup. The user study showed, for instance, that almost all participants were sometimes hesitant to interact with the system either due to the speed of actuation (e.g., when hiding data points, the rods actuated downwards quite rapidly) or due to the noise generated by the actuators. The spacing between rows and columns was also fixed and predetermined by technological constraints. The large actuator size also forced us to increase the height of the system to reduce angles and enable smooth actuation.

Generalizability and Future Directions

The findings from the user study have identified baseline interactions that users prefer, and for which data manipulation tasks they prefer them. For instance, smaller interactions such as annotation afforded physical interaction, whereas larger motions such as row organization afforded touch-screen style swipe interactions. We believe that understanding such interactions with fundamental data manipulation tasks can aid researchers to incorporate them in the development of similar systems.

There are numerous avenues for further work with physical interactions, data analysis tasks, and system functionality. We intend to carry out deeper explorations into the physical aspect, such as data manipulation with external objects, multi-finger input, and pressing over time. More complex task explorations are required (e.g. from Heer and Shneiderman's taxonomy [11]) such as undo/redo, different types of filtering (e.g. thresholding), and combining interactions (annotation with filtering). Furthermore, we require controlled studies with performance metrics (e.g. task completion times, accuracy) to measure how well users can analyze data by using physical visualizations.

CONCLUSION

The key objective of this paper was to uncover means through which physically dynamic bar charts can support data analysis-based interaction techniques. The user study provides initial insight into physical data exploration by testing 14 interaction techniques that formed part of four task-sets: annotation, filtering, organization, and navigation. We report on several ways of combining interaction modalities, utilizing the physicality of the system, incorporating familiar interactions (i.e. from touch-screens), and discussing the challenges that arise from this type of technology. By exploring these fundamental interactions we hope to lay the groundwork for future investigation into physically dynamic data visualizations.

ACKNOWLEDGEMENTS

This work forms part of GHOST, a project funded by the European Commission's 7th Framework Programme, FET-Open scheme (grant #309191). Thank-you to Matthias Schittenhelm, Leanne Bates, Pauline Anthonysamy, and John Vidler for their help at various stages in this project.

REFERENCES

1. Alexander, J., Lucero, A., & Subramanian, S. (2012). Tilt displays: designing display surfaces with multi-axis tilting and actuation. In *Proc. MobileHCI*, pp. 161-170.
2. Brehmer, M. and Munzner, T. (2013). A multi-level typology of abstract visualization tasks. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12, pp. 2376–2385.
3. Brown, C. and Hurst, A. (2012). VizTouch: automatically generated tactile visualizations of coordinate spaces. In *Proc. TEI*, pp. 131–138.
4. Card, S.K., Mackinlay, J.D., and Shneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Pub.
5. Coelho, M., Ishii, H., & Maes, P. (2008). Surfex: a programmable surface for the design of tangible interfaces. In *CHI'08 EA*, pp. 3429-3434.
6. Follmer, S., Leithinger, D., and Ishii, A.O.A.H.H. (2013). inFORM: dynamic physical affordances and constraints through shape and object actuation. In *Proc. UIST*, pp. 417–426.
7. Fritz, J. P., & Barner, K. E. (1999). Design of a haptic data visualization system for people with visual impairments. *IEEE Transactions on Rehabilitation Engineering*, 7(3), pp. 372–384.
8. Greenberg, S. and Buxton, B. Usability evaluation considered harmful (some of the time). (2008). In *Proc. CHI*, pp. 111–120.
9. Hardy, J., Ellis, C., Alexander, J., & Davies, N. (2013). Ubi Displays: A Toolkit for the Rapid Creation of Interactive Projected Displays. In *The International Symposium on Pervasive Displays*.
10. Harrison, C., & Hudson, S. E. (2009). Providing dynamically changeable physical buttons on a visual display. In *Proc. CHI*, pp. 299-308.
11. Heer, J. and Shneiderman, B. (2012). Interactive Dynamics for Visual Analysis. *Commun. ACM* 55, 4, pp. 45–54.
12. Hornbæk, K. Some Whys and Hows of Experiments in Human–Computer Interaction. (2013). *Foundations and Trends in HCI* 5, 4, pp. 299–373.
13. Huron, S., Jansen, Y., and Carpendale, S. Constructing Visual Representations: Investigating the Use of Tangible Tokens. (2014). *IEEE Transactions on Visualization and Computer Graphics* 20, 12, 1.
14. Iwata, H., Yano, H., Nakaizumi, F., and Kawamura, R. (2001). Project FEELEX: adding haptic surface to graphics. In *Proc. SIGGRAPH*, pp. 469–476.
15. Jansen, Y., Dragicevic, P., and Fekete, J.-D. (2013). Evaluating the Efficiency of Physical Visualizations. In *Proc. CHI*. pp. 2593-2602.

16. Jansen, Y. and Dragicevic, P. (2013). An Interaction Model for Visualizations Beyond The Desktop. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12, pp. 2396–2405.
17. Jansen, Y., Dragicevic, P., Isenberg, P., Alexander, J., Karnik, A., Kildal, J., Subramanian, S., Hornbæk, K. (2015). Opportunities and Challenges for Data Physicalization. In *Proc. CHI*.
18. Kyung, K.-U., Lim, J.M., Lim, Y.-A., et al. (2011). TAXEL: Initial progress toward self-morphing visio-haptic interface. *IEEE World Haptics*, pp. 37–42.
19. Lederman, S.J. and Campbell, J.I. (1982). Tangible graphs for the blind. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 24, 1, pp. 85–100.
20. Lee, B., Isenberg, P., Riche, N.H., and Carpendale, S. (2012). Beyond Mouse and Keyboard: Expanding Design Considerations for Information Visualization Interactions. *Visualization and Computer Graphics, IEEE Transactions on* 18, 12, pp. 2689–2698.
21. Leithinger, D., Follmer, S., Olwal, A., et al. (2013). Sublimate: state-changing virtual and physical rendering to augment interaction with shape displays. In *Proc. CHI*, pp. 1441–1450.
22. Leithinger, D. and Ishii, H. (2010). Relief: a scalable actuated shape display. In *Proc. TEI* pp. 221–222.
23. Leithinger, D., Lakatos, D., DeVincenzi, A., Blackshaw, M., and Ishii, H. (2011). Direct and gestural interaction with relief: a 2.5 D shape display. *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pp. 541–548.
24. Manshad, M.S., Pontelli, E., and Manshad, S.J. (2012). Trackable interactive multimodal manipulatives: towards a tangible user environment for the blind. In *Proceedings of the 13th international conference on Computers Helping People with Special Needs - Volume Part II*, Springer-Verlag, pp. 664–671.
25. Moere, A.V. (2008). Beyond the tyranny of the pixel: Exploring the physicality of information visualization. *Information Visualisation*, pp. 469–474.
26. Paneels, S., & Roberts, J. C. (2010). Review of Designs for Haptic Data Visualization. *IEEE Transactions on Haptics*, 3(2), pp. 119–137.
27. Perkins, C. (2002). Cartography: progress in tactile mapping. *Progress in Human Geography* 26, 4, pp. 521–530.
28. Poupyrev, I., Nashida, T., Maruyama, S., Rekimoto, J., and Yamaji, Y. (2004). Lumen: interactive visual and shape display for calm computing. *ACM SIGGRAPH Emerging technologies*, 17.
29. Rasmussen, M.K., Pedersen, E.W., Petersen, M.G., and Hornbæk, K. Shape-changing interfaces: a review of the design space and open research questions. In *Proc. CHI*, pp. 735–744.
30. Khot, R. A., Hjorth, L., and Mueller, F. (2014). Understanding Physical Activity through 3D Printed Material Artifacts. In *Proc. CHI*.
31. Roudaut, A., Karnik, A., Löchtefeld, M., and Subramanian, S. (2013). Morphees: Toward High “Shape Resolution” in Self-Actuated Flexible Mobile Devices. In *Proc. CHI*.
32. Shah, P. and Hoeffner, J. (2002). Review of graph comprehension research: Implications for instruction. *Educational Psychology Review* 14, 1, 47–69.
33. Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proc. IEEE Symposium on Visual Languages*, pp. 336–343.
34. Spence, I. (1990). Visual psychophysics of simple graphical elements. *Journal of Experimental Psychology: Human Perception and Performance* 16, 4, pp. 683–692.
35. Stusak, S., & Aslan, A. (2014). Beyond physical bar charts: an exploration of designing physical visualizations. In *CHI'14 EA*, pp. 1381-1386.
36. Stusak, S., Tabard, A., Sauka, F., Khot, R., & Butz, A. (2014). Activity Sculptures: Exploring the Impact of Physical Visualizations on Running Activity. *IEEE Transactions on Visualization and Computer Graphics*, 99, 1.
37. Tversky, B., Morrison, J.B., and Betrancourt, M. (2002). Animation: can it facilitate? *International Journal of Human-Computer Studies* 57, 4, pp. 247–262.
38. Tversky, B. (2001). Spatial schemas in depictions. *Spatial schemas and abstract thought*, pp. 79–111.
39. Wall, S.A. and Brewster, S. (2006). Sensory substitution using tactile pin arrays: Human factors, technology and applications. *Signal Processing* 86, 12, pp. 3674–3695.
40. Wobbrock, J. O., Morris, M. R., & Wilson, A. D. (2009). User-defined gestures for surface computing. In *Proc. CHI*, pp. 1083-1092.
41. Zacks, J., Levy, E., Tversky, B., and Schiano, D.J. (1998). Reading bar graphs: Effects of extraneous depth cues and graphical context. *Journal of Experimental Psychology Applied* 4, pp. 119–138.
42. Zhao, J. and Moere, A.V. (2008). Embodiment in data sculpture: a model of the physical visualization of information. *Proc. DIMEA*, pp. 343–350.