

# Put a Label On It! Approaches for Constructing and Contextualizing Bar Chart Physicalizations

Kim Sauvé  
kim.sauve@lancaster.ac.uk  
Lancaster University  
United Kingdom

Argenis Ramirez Gomez  
argenis.ramirez-gomez@port.ac.uk  
University of Portsmouth  
United Kingdom

Steven Houben  
s.houben@tue.nl  
Eindhoven University of Technology  
The Netherlands

## ABSTRACT

Physicalizations represent data through their tangible and material properties. In contrast to screen-based visualizations, there is currently very limited understanding of how to label or annotate physicalizations to support people in interpreting the data encoded by the physicalization. Because of its spatiality, contextualization through labeling or annotation is crucial to communicate data across different orientations. In this paper, we study labeling approaches as part of the overall construction process of bar chart physicalizations. We designed a toolkit of physical tokens and paper data labels and asked 16 participants to construct and contextualize their own data physicalizations. We found that (i) the construction and contextualization of physicalizations is a highly intertwined process, (ii) data labels are integrated with physical constructs in the final design, and (iii) these are both influenced by orientation changes. We contribute with an understanding of the role of data labeling in the creation and contextualization of physicalizations.

## CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in visualization**.

## KEYWORDS

Data Physicalization, Physical Visualization, Data Labels, Constructive Visualization

### ACM Reference Format:

Kim Sauvé, Argenis Ramirez Gomez, and Steven Houben. 2022. Put a Label On It! Approaches for Constructing and Contextualizing Bar Chart Physicalizations. In *CHI Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3491102.3501952>

## 1 INTRODUCTION

*Physical data visualizations* or *physicalizations* are tangible and three-dimensional artifacts that represent or encode data in their material and physical form [30]. These physical data representations are useful tools to support collaborative scenarios and the exploration of data, but also open up new ways to interact with, transform, and inspect data [14]. For physicalizations – as with 2D visualizations – the inclusion of data labels, axes values, legends and annotations are in many cases fundamental to contextualizing

the presented data. They provide people with context, frame of references, and visual guides on how to interpret the data. Particularly because of the challenges associated with multi-orientation interpretation of data to groups of people [40], such labels and annotations are instrumental in helping people make sense of the presented data. We refer to this as the *contextualization* of physicalizations, which is the inclusion of contextual elements such as data labels, axes, legends, and annotations to support the extraction of information from physical representations of data.

Despite the obvious importance of providing guiding context to visualizations, related work on physicalizations [13, 14, 30], does not actively consider the labeling of physical data points and structures. Work that does consider labeling of physicalizations in some form (e.g., [38, 45, 47]), often use very different approaches that are not systematic or even consistent with each other. From a conceptual viewpoint, the current definition of physicalization [30] indeed focuses on materiality and does not highlight ‘data labeling’<sup>1</sup> as an explicit part of the physicalization itself. However, a physicalization cannot do without context; the physicality and spatiality of physicalizations explicitly opens up questions such as (i) where to locate different kinds of labels (i.e. title, axes labels, and data values) in relation to the canvas and/or other data points, and (ii) how this is affected by user orientation (e.g. when multiple people are looking at the physicalization from different perspectives). More fundamentally: *why, how, and when should ‘data labels’ be included in the design, construction of, and interaction with physicalizations?*

Research on Constructive Visualization [26] focuses on explicitly understanding the translation process from raw data to physical form. While this approach has provided detailed insights into the construction of data points and structures of the physicalization, they similarly do not actively include data labeling in the authoring of physicalizations. For example Fan et al. [17] provide ready-made braille labels but leave contextualization of data open to participant’s choice, and both Huron et al. [27], and Wun et al. [52] include the annotation of data as a subsequent task to the construction task. As observed by Wun et al. [52], the creation of physical data representations results in an *interrelation principle*: the placement and rearrangement of physical data objects in space – *loading data* – simultaneously influences the *visual mapping* and *presentation mapping* of a visualization. In line with this observation, we propose and argue that the act of ‘data labeling’ should be an active part of this process, further intertwining the construction and contextualization of physical data visualizations.

In this paper, we aim to answer the question: *how does data labeling play a role in the physicalization creation process, visualization*

CHI '22, April 29-May 5, 2022, New Orleans, LA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. 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 Conference on Human Factors in Computing Systems (CHI '22)*, April 29-May 5, 2022, New Orleans, LA, USA, <https://doi.org/10.1145/3491102.3501952>.

<sup>1</sup>Not to be mistaken with the term ‘data labeling’ as used in Machine Learning to describe the annotation of raw data to train a classifier.

*design, and when viewed from different orientations?* As studying data labels in isolation is artificial, our research explores the role of ‘data labeling’ in the entire construction process of barchart-type 3D physicalizations [27]. We follow the approach from Constructive Visualization research, and “*study human behavior independently from the design of specific software tools*” [27] to inform the design of future physical visualizations. To study this process, we designed a toolkit that allows for the creation of data visualizations in 3D space and includes data labeling as ‘building block’ alongside the use of physical 3D tokens. We conducted a study with 16 participants who completed a total of 32 construction tasks. We contribute (i) an understanding of the role of data labeling in the construction and contextualization of physical visualizations, (ii) an overview of how textual and physical constructs coexist in visualization designs, and (iii) reflections on coping strategies for contextualizing bar chart physicalizations across orientations in physical space.

## 2 RELATED WORK

Herein, we discuss (i) existing concepts and knowledge on labeling from 2D Information Visualization (InfoVis), (ii) current use of data labels in existing physicalizations, and (iii) approaches from Constructive Visualization.

### 2.1 Labeling in Information visualization

Information visualization [3, 7, 8, 35, 50] has a long standing tradition, rooted in a history of cartography and later in Computer Graphics, of labeling and annotating visual representations of data. Many of these labeling practices have now been operationalized into toolkits, default visualizations, and best practices [e.g. 19, 23]. As described by the ‘Data Design Standards’ [23]: “*Labels make it easier for users to understand data visualizations by using text to reinforce visual concepts. Labels are traditionally used to label axes and legends, however, they can also be used inside of data visualizations to communicate categorical, sequential, or value attributes*”. In recent years, labeling research has mainly focused on novel forms of graphic algorithms and approaches to handle label placement in complex visualizations [1, 9, 15, 31] including a focus on automation [32], 2D graph layout techniques [20], or best practices for ‘good’ label placement [49]. Nonetheless, as suggested by Brath [5] “*3D InfoVis is here to stay*”, meaning work has also looked at labeling Interactive 3D visualizations [2] or 3D geo-referencing [11].

With the move to a more interactive ‘human-data interaction’ approach, new insights around semantic or interactive versions of information visualization labels have been introduced [48], opening new possibilities for touch-based or even physical data visualizations. A recent concept in the field of information visualization that operationalizes this increased interactivity is the *extended infovis pipeline model* [28]. This model explains the translation from raw data to a visualization that can be rendered in the physical world. It distinguishes between *data transformation*, *visual mapping*, *presentation mapping*, and *rendering*. Especially visual and presentation mapping are of importance to discuss here, as it explains the difference between creating the initial *abstract physical form* and the fully-specified *visual presentation* [28]. According to the infovis pipeline, elements such as axis labels, grid lines, legends and captions are *decoration* operations as part of *presentation mapping*.

However, the precise way in which these grids, legends or captions should be designed in physical 3D space is not specified nor defined.

### 2.2 Labeling in Physicalization

Looking at the use of data labels in existing physicalization research, we observed that state of the art (summary in [14, 30]) pays little attention to data labeling. Therefore, other means are often required to contextualize the data represented, such as prior knowledge, the use of an external device to reveal data, or no means to extract details (i.e. because the intention is purely aesthetic and/or by estimation). Particularly exploratory physicalizations such as *data sculptures* [4, 34, 44, 54] or data installations with complex ecosystems [25, 33, 38, 39] do not provide on-physicalization labeling.

Physicalizations that do use labeling in some form, do so in a myriad of different ways. Examples of interactive systems are work from Veldhuis et al. [47] that presented textual information in a single direction, or Taher et al. [45] that used multiple displays to provide two duplicates for x-axis and y-axis (and only shows categorical/sequential data but no values or legend for values). Examples of static physicalizations are work from Jansen et al. [29] that compared on-screen 3D bar charts with labels floating in space in the reading direction of the viewer, with physical 3D bar charts that represent the same labels sideways in a counterclockwise direction (with the addition of an engraved transparent acrylic back wall to show scale); and Danyluk et al. [10] that used similar physical 3D models but then with alternating reading directions on different sides of the base. Gourlet et al. [22] built a physicalization where the reading direction was aligned in 4 different directions, oriented by each side of the table. Stusak et al.’s [43] work on physicalizations used numeric values on the physical bar-charts, labels for countries on the flat surface, and a transparent background panel with scales. Finally, recently Ren et al. [37] explored physicalizations that were annotated with a basic legend on one side of the visualization.

While these labeling approaches are generally well designed, they are very different and inconsistent with each other, opening up questions around what strategies or approaches can be used for labeling of physicalizations? Furthermore, because of the intrinsic three-dimensionality and physical nature of physicalizations, they can be used, observed, perceived and approached from different directions, making the process of labeling even more challenging. From a conceptual and theoretical perspective, we also observe that labeling is never explicitly included in the definition and scope of physicalization [30], the rendering process [12], or a recent reflection on the research domain of physicalization [14]. Hence, there is currently no principles or standard ways to label in physical space when it comes to reading direction, text orientation, and location in relation to physical data points and the canvas.

Text orientation and readability of labels is also a concern for work on virtual reality [6, 42]. While a full review of this work is beyond the scope of this paper, previous work has combined physicalization or visualization with VR environments. For example, Ren and Hornecker [37] explored the differences between physicalization and VR simulation and use basic text labels next to the bar-chart in both approaches. Ulusoy et al. [46] explored VR-models of bar-chart physicalizations that were annotated with labels and presented on different scale (i.e. hand-size versus room-size) in



virtual space. Finally, Danyluk et al. [10] compared physical and VR visualizations, again leveraging data annotations and labels around 3D bar-charts.

Lastly, outside the context of physicalization, work has explored how to position and orientate text, illustrating that there are different ways in which text can be represented in 2D and 3D space. These studies discuss for instance text orientation [24], horizontal versus vertical reading [36], left-to-right versus top-to-down reading [21], and the influence of 3D rotations on reading speed [51]. These findings from HCI studies agree with literature from the vision community that also demonstrates the impact orientation has on reading speed [53].

## 2.3 Constructive Visualization

Within the research space of Physicalization, ‘Constructive Visualization’ work explored how to author and construct physical data presentations [27, 52]. This work is concerned with describing and exploring the methods, strategies and tools that help people transform data into physical representations. However, currently these models and approaches for constructive visualization do not include data labeling as an active component in the construction process (visualization mapping), but rather treat annotation of data (presentation mapping) as a secondary process after the construction of the physical form factor. Both Huron et al. [27] and Wun et al. [52] included annotation as a subsequent task to the construction task, while Fan et al. [17] left it up to the participant to use pre-made braille labels in their visualization.

Wun et al. [52] observed that the construction of physicalizations results in an *interrelation principle* [52], as moving physical elements influences multiple parameters of the visualization pipeline at once. For example, when *loading* data (placing data objects in the canvas), one simultaneously has to consider the *visual mapping* (where to place the object in relation to other data objects), and *presentation mapping* (object placement within the canvas). We suspect that because of this *interrelation principle*, the labeling of physicalizations will similarly be intertwined into the overall process.

Hence, for validity we do not want to and/or cannot investigate the labeling of data in isolation. Our methodology is, thus, based on constructive visualization work, with the difference that we treat labeling as an active component in the authoring process.

## 3 STUDY RATIONALE

The focus of this study is to build a better understanding of the role of *labels* in physicalizations. With **data labels** or **data labeling** we refer to annotations that, like visualizations on a screen, highlight axes, data points, legends, and other visual structures that support people in reading and interpreting data effectively. While prior work has considered the labeling of physicalizations in various forms, these have almost always been post-hoc activities from a necessity to counter some of the open challenges or common problems in physicalizations. Therefore, there are no real insights or principled approaches into how, if, and when to label physicalizations. While we can borrow initial insights from screen-based visualizations [2, 15, 23, 35], many of these do not translate directly to the context of physicalizations. Because of their physicality, people have very different strategies to perceiving, using, and

interacting with physicalizations. This implies that more systematic research into labeling strategies and practices is needed to explore how physicalizations can be labeled effectively – taking into account their specific challenges around spatiality, user orientation, and perception.

As labeling is difficult and artificial to understand in isolation, we specifically examine labeling as part of the general construction process of physicalizations. Because the *interrelation principle* [52] suggests that constructing physicalizations is a highly intertwined process that combines various aspects of the *extended infovis pipeline model* [28], we argue that it cannot be understood or studied in isolation. By studying and documenting the strategies that people take for labeling of data, axes, clusters and entire physicalizations, we can learn more about the role of labeling in the overall construction process, but also about how non-experts view physical structures and data points in relation to a given dataset. While studying the labeling of existing physicalizations might help build some insights into how data labeling works, we argue that this would also be a post-hoc activity that reduces labeling to a second class aspect of physicalization – where we suggest it should be a fundamental and inherent part of the overall physicalization design. As such, our study methodology studies labeling in combination with other constructive visualization processes [27, 52].

## 4 METHODOLOGY

The goal of this study is to investigate the role of data labeling (i) during the creation process of a physical visualization, (ii) within the resulting visualization design, and (iii) when viewing the visualization from different orientations and perspectives. Our study is designed to document and highlight strategies and approaches towards constructing a 3D bar chart physicalization and annotate them with contextual labels using a custom design toolkit and methodology. We designed a task that required participants to build two physical visualizations given a toolkit including a set of physical colored blocks and textual labels. During the task, the researcher presented the participant with one dataset at a time and prompted them to build the resulting data using the toolkit. After the creation process, participants were asked to reflect on their visualization design during an integrative process, as the canvas was rotated in increments of 90 degrees. On each rotation, participants were required to observe their visualization from the new viewing angle, and (if desired) make changes to their labeling.

### 4.1 Apparatus


We created a custom-made toolkit including plastic building blocks and paper labels inside a storage box. The toolkit follows [27, 52] in providing a set of custom tools aiming to avoid the artificial constraints introduced by existing systems, as they are often limited by the technologies used. We discuss each of the components below:

**4.1.1 Building blocks.** The design of the building blocks is inspired by interlocking maths learning cubes such as Snap Cubes<sup>2</sup> and Edx Education Linking Cubes<sup>3</sup>. Each block has three different types of faces: 1 square stud, 3 square holes, and 2 regular faces (Figure 1).

<sup>2</sup><https://www.learningresources.co.uk/snap-cubes-set-of-100>

<sup>3</sup><https://edxeducation.com/portfolio-item/2cm-linking-cubes-1000pcs-12012/>

The goal of this design is to allow for enough freedom of the creation in 3D space – the set of different faces per block supports attachment in multiple directions – but also keep them simple and consistent in appearance. Each block is 2x2x2cm in size and is made of 3D printed plastic. The storage box contained 25 blocks of each out of 5 colors (red, orange, yellow, green, blue), totally 125 blocks.



**Figure 1: A 3D rendering of the block design, 3D-printed plastic blocks, and the study setup.**

**4.1.2 Data labels.** The set of paper data labels included: a title label, a label for each categorical (i.e., seasons, countries) and sequential attribute (i.e., years), and a label for each single value attribute. We purposely provided a minimal set of data labels with no duplicates and no inclusion of axis labels (i.e., ‘X’, ‘Y’, ‘Country’, ‘Season’, ‘Year’) to reduce the possibility for redundancies. Lastly, we provided participants with some sticky tack to allow for freedom in placing labels, i.e., sideways on blocks or other midair placements.

**4.1.3 Canvas.** We designed a building area made of a white plastic 40x40cm canvas with square holes at every 2cm so that the building blocks could be snapped in. Figure 1 shows the experimental setup for all tasks. The participant was seated in front of the white square canvas with the toolkit on their left.

## 4.2 Datasets

We used two datasets of similar structure and complexity as Huron et al. [17] (included in supplementary material). The first dataset<sup>4</sup> represented CO<sub>2</sub> emissions in tons per person for five different countries, across three years. The second dataset<sup>5</sup> represented rainfall in the United Kingdom in millimeters for four seasons, across four years. All values are rounded derivatives from the raw data. The datasets were selected so that they are understandable, interpretable and transferable for non-expert participants.

## 4.3 Participants

We recruited 16 participants (8 identified as male, 7 as female, and 1 as non-binary), of which 5 were 18-24 years, 4 were 25-34 years, 6 were 35-44 years, and 1 was between 45-54 years old. Participation was voluntary and without compensation. There were no particular requirements for participation other than that participants were (corrected to) fully sighted and physically able to construct a visualization with objects. Of all participants, 12 were familiar with the concept of data visualization, 13 were experienced in reading data visualizations, and 10 experienced in creating data visualizations.

<sup>4</sup><https://www.gapminder.org/data/>

<sup>5</sup><https://www.metoffice.gov.uk/research/climate/maps-and-data/uk-and-regional-series>

## 4.4 Procedure

At the start of the study, we introduced participants to the study, asked them to sign a consent form, and collected their demographics. We explained the goal of the study: to understand how people construct and label physical visualizations using an exemplar toolkit. We gave participants a set of general instructions and in total asked them to visualize two datasets using the toolkit. Participants were asked to think out loud during the creation process. If participants indicated to have finished but forgot to contextualize their physical constructs they were prompted by the researcher, for example about the topic “*how would someone else know what your visualization is about?*” or the created encoding “*how would they know what one block represents?*”. When finished with the first task (T1), the researcher would ask them to take two pictures of the end result and explain their visualization design. Afterward, the participant was asked to rotate the canvas either 90 degrees clockwise or counterclockwise, and indicate if they would like to make any changes to the labeling of the visualization and if so, they were requested to perform these changes, and take two pictures (from different angles) to capture the current state of the visualization. We repeated this process twice so eventually the participant had seen all 4 orientations of the square canvas. This whole process was repeated during task 2 (T2) with a second dataset. The mapping between the two datasets and two directions was counterbalanced across participants using a balanced Latin square (yielding 4 participant groups). The whole experiment lasted between approximately 60 to 90 minutes, depending on the participants’ performance.

## 4.5 Data Collection

During the study, we collected three different types of data:

**4.5.1 Video.** With participants’ consent we took video and audio recordings of their interactions using two GoPro’s: from a top-down viewing angle and a view from the side. We used these videos to capture participants’ actions during the creation process.

**4.5.2 Pictures.** After each task, and after the changes made upon each rotation of the canvas we asked participants to take two pictures from different viewing angles to capture the current state of the visualization. The first picture was a representation of their viewing angle while seated, and the second picture from any angle they preferred to view their visualization most comfortably and/or effectively. We used these pictures to extract (i) the properties of their visualization design, and (ii) any changes to the labeling across different orientations.

**4.5.3 Participant Observations.** During the task, the researcher made notes of participant comments while thinking out loud. After each task we asked participants to (i) elaborate on the dataset using their visualization, (ii) explain the visualization they created, and (iii) if there was anything they struggled with while creating it. This was to understand participants’ creation process and the properties of their visualization design. After both tasks, we asked them about their overall experience with the toolkit.

## 4.6 Method of Analysis

To be able to extract information on (i) the construction and contextualization process and (ii) the properties of the final visualization designs we developed coding schemes for the videos and pictures:

**4.6.1 Analysis of the Creation Process.** We analyzed the videos, using a qualitative and iterative approach, inspired by the approaches of Wun et al. [52] and Huron et al. [27]. We used the ethogram as created by Wun et al. [52] as a reference, but refined it to meet our apparatus (3D blocks instead of 2D tiles and the inclusion of labels) and study aim (role of labeling in the creation of physicalizations).

The first pass involved two researchers performing open coding to identify the behaviors of interest. Once the coding scheme was established, there was primarily one coder, with random checks to verify researcher agreement.

In total, we coded 13 types of actions across 3 activity categories (Table 1). Additionally, we captured when which out of 4 label types (title, sequential, categorical, and value) was interacted with.

Activity category	Action	Description
Data activities	Read	Read the data table.
	Verify	Verification of visualization, i.e. compare with data table and/or count blocks.
	Correct error	Correct an error.
Block activities	Collect	Collect (and count) blocks in hand, canvas or workspace.
	Organize	Organize (constructs of) blocks spatially in the canvas, without placing.
	Build in hand	Build block constructs in hand.
	Build in canvas	Build block constructs in the canvas, without placing.
	Place in canvas	Place block constructs in the canvas.
	Rearrange	Rearrange and place block(s) in the canvas.
	Placeholder	Place placeholder block(s) in canvas for labeling purposes.
Label activities	Order	Order labels in the workspace.
	Label	Place labels in canvas.
	Relabel	Rearrange label(s) in canvas.

**Table 1: Ethogram of activity categories and actions identified in the video data.**

**4.6.2 Analysis of Visualization Design.** We analyzed the pictures taken by participants after the completion of the physicalization creation process to identify (i) the visualization type; (ii) composition; (iii) color association; (iv) axis mapping; (v) data labeling position; and (vi) labels' reading direction. These codes emerged during an iterative process of analysis of the resulting physicalizations and aim to describe how the blocks and labels were mapped and distributed on the canvas to visualize the provided dataset.

*Visualization type* describes the distribution of blocks and the use of the multi-direction stacking affordance of the toolkit in the canvas. *3D visualizations* utilize multiple levels of stacked blocks to distribute data values using height (z-axis) within the 3D space. On the other hand, *planar visualizations* were constructed using a single level of blocks, thus distributing them only in the 2D space (flat surface, x and y-axis). For instance, blocks organized in towers

(stacked) are described as *3D*, whereas visualizations that do not stack more than one block in the canvas are *planar*.

*Composition* refers to visualization archetypes based on the distribution, dispersion, organization, and/or positioning (location) of blocks and groups of blocks within the canvas space. Composition archetypes emerged from the analysis of all the resulting physicalizations, grouping them by look-alike block distributions as new archetypes appeared. For instance, blocks organized equidistantly and dispersed across the canvas belong to a different archetype than those not organized equidistantly; or those clustered in one corner of the canvas.

*Color association* describes how participants use color affordance of the toolkit. Generally, the color of blocks could be used to map *sequential* or *categorical* attributes from the dataset into the canvas space. In contrast, the number of stacked/grouped blocks is used to represent values.

*Axis mapping* refers to the use of the canvas space to map sequential and categorical attributes into the x and y-axis (from the viewer's point). For instance, a physicalization that utilizes the horizontal direction (x-axis) to spread year values (sequence), whereas the canvas depth (y-axis) is used to map seasonal values (categories).

*Data labeling position* provides information about the *location* of each of the 4 label types: title, sequential, categorical, and value labels. For instance, whether a label is located on the canvas, next to a block, on top of a value block or a placeholder block, or onto one of its faces (in the z-axis).

*Labels' reading direction* registers the orientation of each label type from the participant's point of view. This describes if the label can be read from their perspective (in a default direction), it is upside down, or rotated on an approximately 90 degrees angle; and whether all labels follow a consistent direction pattern or are in mixed directions.

**4.6.3 Analysis of Influence of Orientation.** We analyzed the changes participants made to the physicalization's labels after each shift in orientation (three instances) using the pictures they took at the end of each iteration. We followed the analysis of visualization design and registered the changes in *data labeling position* and *reading direction* for each of the four types of labels (title, sequential, categorical, and values). In addition, we compiled a list of actions as descriptors of the changes in position or tweaks and their occurrence per participant. For instance, a title label moved from the back of the canvas to the front, or value labels moved from the canvas to the top of towers of blocks were described as a "relocation". Similarly, changes in orientation of labels or placeholders to preserve their reading direction were described as "rotations". Finally, we refined the list of actions as new ones emerged and organized the resulting dictionary in clusters when appropriate, e.g., grouping actions of low occurrence.

## 5 FINDINGS

To answer our research question we structured the findings in three sections. The first section presents an overview of the construction and contextualization process when creating a physical visualization. The second section elaborates on the relation between the physical and textual properties of the visualization design. The last




Figure 2: Activity categories over time for each participant for tasks 1 and 2: data (■), block (■), and label activities (■).

section shows the influence of orientation on the changes to the textual properties of the physicalization.

Overall, we found that (i) the creation of physicalizations is an intertwined process of label and block activities and unique per participant, (ii) the final visualization design is an integration of data labels and physical constructs, and (iii) the relations between these labels and constructs are influenced by orientation changes.

## 5.1 Construction & Contextualization Process

Herein, we discuss the role of labeling during the physicalization creation process. We first discuss the actions observed in general, after which we go into further detail on behavioral patterns observed within the label and block activities, and across activity categories.

**5.1.1 Overall creation process.** Across all 32 tasks (16 participants  $\times$  2 tasks), participants spent on average 13 minutes to complete the task ( $\sigma = 4.5$  minutes). 9 participants performed task 2 (T2) faster than task 1 (T1) on average by 4 minutes, whereas 7 participants performed T2 slower than T1, on average by 3.5 minutes.

Looking at the occurrence of activities over time, we observed that the construction and contextualization of physicalizations is an intertwined process, as illustrated in Figure 2. This means that labeling happens throughout the creation process rather than at

the end. Across all participants and tasks, on average 53.5% of their time was spent on any type of block activities, 22.7% on any type of labeling activities, and 23.7% on any type of data activities.

Data activities such as looking at the data table generally happened throughout the process, as can be seen from the short time periods throughout the task (Figure 2). Block activities appear in longer periods of time clustered together. Lastly, label activities vary from short time periods throughout to clusters of longer time periods spread across the task, for example at the very beginning of a task to plan out the visualization design or at the end to complete the block constructs. Figure 4 provides a further detailing of the activities observed and the average time spent on each.

Following the overall process observations, we zoom in on the behavioral patterns within and between the different activity categories. For example, some participants built all constructs first (block activities), and then labeled the whole visualization (label activities), whereas others applied a more parallel process in which block and label activities alternated and/or intertwined. For an overview of the timelines per participant per task please refer to the supplementary material.

**5.1.2 Label activity patterns and label types over time.** For each task, we extracted when which out of 4 label types was handled, and analyzed the relation between ordering, labeling, and relabeling.

**Ordering.** For 16 tasks (50%) we observed the ordering of labels at the beginning of the creation process (before any block activities). For example, P14-T2 in Figure 3 and as illustrated in Figure 4 by ‘Ordering labels’. In contrast, we found that for 7 tasks (21.9%) ordering happened either along the creation process – such as P1-T2 in Figure 3 – or at the end (after block activities took place) – see P10-T2 in Figure 3. Lastly, the 9 remaining tasks (28.1%) did not involve any ordering of labels at all.

**Labeling.** Looking at the use of each label type over time we observed different strategies:

- **Title labels:** For more than half of the tasks the title label was placed at the very end ( $f = 19$ ; 59.4%), whereas for 13 tasks (40.6%), the title label was placed at the beginning or first half of the task.


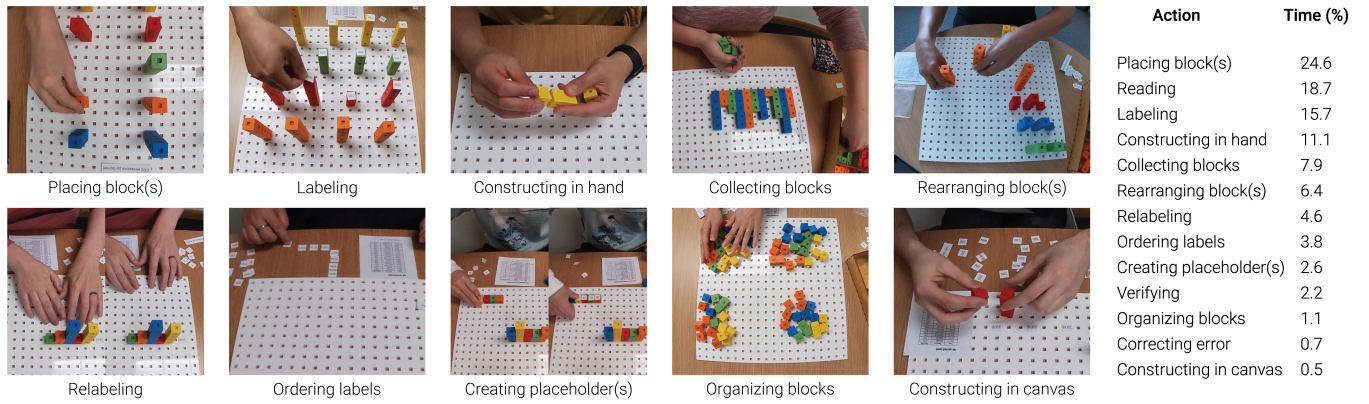


Figure 3: Different approaches to ordering labels (■): at the very beginning of the task (P14-T2), after block activity (■) took place (P1-T2), or at the very end of the task (P10-T2).





**Figure 4: Illustrations of the observed Block and Label actions, and the percentage of the average time spent on each action that appeared during the creation process.**

- *Sequential & Categorical labels:* We observed the placement of sequential labels was performed (i) throughout the task whilst building sequential block constructs ( $f = 12$ ; 37.5%), or (ii) at the beginning or first half of the task ( $f = 12$ ; 37.5%). For the remaining tasks, this happened at the end of the process ( $f = 8$ ; 25%). For categorical labels, we observed that they are placed either at the beginning ( $f = 9$ ; 28.1%) or first half of the task ( $f = 8$ ; 25%); during the final half ( $f = 3$ ; 9.4%) or at the end of the task ( $f = 9$ ; 28.1%); or spread out during the task ( $f = 3$ ; 9.4%). When we cross-referenced the placing of sequential and categorical labels, we observed some participants placed both of them at the beginning of the task to plan the visualization ( $f = 9$ ; 28.1%); whereas others preferred to place both at the end ( $f = 6$ ; 18.8%). Moreover, some participants chose to place categorical labels in the beginning ( $f = 6$ ; 18.8%) or the end of the task ( $f = 5$ ; 15.6%) whilst sequential labeling was spread across the task, placing them either before or after a sequential construct was created.
- *Value labels:* For the majority of tasks, the labeling of values happened at the end of the task ( $f = 25$ ; 78.1%), after the physical constructs were created. Of these tasks, 7 spent a longer period of time on placing all value labels, 5 spent a shorter period of time on creating a single key, and 2 involved the placement of value labels at first after which a key is created as well (P4, P9). For 2 tasks (6.3%) a longer period of time is spent on value labeling at the beginning or first half of the task. For instance, P14 spent time placing labels to plan out their visualization design, whereas P3 did the same to create a ‘legend tower’ (Figure 6). Lastly, for 5 tasks (15.6%) the value labeling happened throughout the task.


**Relabeling.** We observed that relabeling generally occurred for categorical and sequential labels rather than for value and title labels. To give an example, P8 placed categorical labels on the first bar charts they build, but as they got occluded by the subsequent constructs, they updated the categorical labeling after all physical constructs were finished. In contrast, P14-T1 relabeled each value label as they built physical constructs, after they had placed all labels at the beginning of the creation process to plan their visualization.

**5.1.3 Block activity patterns.** For each task we extracted which block action(s) involved the largest percentage of time and whether

or not they occurred in a chain of actions. To give an example, Figure 5 shows that for P6-T2 the most occurring chain of actions is *collect*, *build in hand*, and *place in canvas*. Overall, we observed four general strategies:

- *Collect – build in hand – place in canvas* ( $f = 10$ ; 31.3%).
- *Place in canvas* ( $f = 9$ ; 28.1%).
- *Collect – place in canvas* ( $f = 8$ ; 25%).
- *Build in hand – place in canvas* ( $f = 5$ ; 15.6%).

The occurrence of these different strategies to build and place constructs can be explained by the affordances of the apparatus. The physical blocks allow for the construction and ‘clicking’ together in multiple ways (in contrast to stackable tiles).




**Figure 5: P6-T2 illustrates the block activity pattern *collect* (■) – *build in hand* (■) – *place in canvas* (■).**

**Organization.** We observed different strategies in the organization of blocks. For instance, P5-T1 organized multiple block constructs on the canvas before placing them (see Figure 6). Moreover, P13-T2 first repeats the collection and organization of blocks within the canvas (Figure 4; ‘Organizing blocks’), after which they start placing all of them.

**Rearrangement.** For 4 tasks (12.5%), we observed that a longer period of creation time was dedicated to the rearrangement of one or more blocks after their placement, for instance halfway through and/or at the end of the task.


**5.1.4 Patterns across block and label activities.** If we look at the relation between block and label activities, generally, we observed that for 6 tasks (18.8%) all block activities were performed first, after which label activities were done (for example Figure 7; P16-T1). For the remaining 26 tasks (81.3%) we observed an alternating and/or intertwined process of block and label activities; meaning




**Figure 6:** P3-T1 showing their ‘legend tower’, P5-T1 organizing block constructs on the canvas before placing them, P7-T1 simplifying construction through rotation of the canvas, and P15-T1 using a label to assist in reading the data table.

that participants were alternating between longer periods of time spent on label or block activities (Figure 7; P4-T2) or spent shorter periods of time on label and block activities subsequently, resulting in a more intertwined process (Figure 7; P14-T1).

Other examples show longer time periods of isolated label or block activities at first, that become shorter and more intertwined over time (Figure 7; P1-T1), or vice versa, planning out the visualization using an intertwined process, after which isolated block and label activities are performed (Figure 7; P5-T2).



**Figure 7:** Different patterns across block (■) and label (■) activities: performed subsequently (P16-T1), in alternation (P4-T2), labeling after which block and label activities are intertwined (P14-T1), from alternation towards intertwined (P1-T1), and from intertwined towards alternation (P5-T2).



**Figure 8:** Examples of intertwined patterns across block and label activities: P3-T2 developed a pattern of collect (■), build in hand (■), place (■), and label (■). P8-T1 used placeholders (■) while labeling each sequence label.

An example of a fully intertwined process of block and label activities is P3-T2 (Figure 8). They mentioned that they first used the sequential and categorical labels to plan out the canvas, and placed each value label as they build constructs for each data point.

Lastly, looking at the placement of placeholder blocks meant for labeling, we observed that this often occurs in parallel or in close proximity of label activities (Figure 8; P8-T1).

**5.1.5 Patterns across data and label activities.** Looking at the relation between data and label activities, we observed that when block activities occur before label activities, this can influence the need for data activities, as physical constructs can be used as reference and/or means of verification.

For the 6 tasks that participants first performed all block activities and then labeling, we found that they did not look at the data table while labeling, as they could use their physical constructs as reference for extracting values. Similarly, we observed this for time periods throughout the alternating and/or intertwined processes, and especially at the end of a task when placing value labels. The placement of value labels at the end of a task was regularly accompanied by verification before, during, or after the labeling.

**5.1.6 Other activities.** We observed that participants sometimes used creative methods to support the creation process. For example, P7 rotated the canvas repeatedly to bring the area of interest closer to them and simplify construction (Figure 6; P7-T1), whereas other participants used the storage box or other attributes to cover up parts of the paper data table to guide reading (Figure 6; P15-T1).

Regarding the use of the different block faces, we observed that participants either cared much or not at all about the direction of the open and closed block faces. Participants that paid close attention to the order of block faces tended to build slower and/or more carefully as precision was required. Lastly, P4 and P8 regularly clicked the wrong block faces together and had to correct themselves. They are the only two participants that showed some minor struggles when constructing the blocks in 3D space, due to their affordance of being attachable in multiple directions. Participants identified different advantages for the open and closed faces: they mentioned that closed faces could create more “neat” or “peaceful” visualizations, whereas the open faces could simplify comparison through counting. P16 mentioned the potential of the block faces (open and closed) to encode further information/detail, i.e., meaning (“to communicate a food item with or without sugar”).





**Figure 9: Overview of all visualization designs created by participants. An enlargement is available in supplementary material.**

## 5.2 Visualization Design

In this section, we elaborate on the visualization type and composition, color association, axis mapping, and use of data labels as part of the final visualization designs created by the participants.

**5.2.1 Visualization type and composition.** Overall, we observed 5 different visualization archetypes across all 32 tasks. Figure 9 shows an overview of the visualization designs created by the participants and their corresponding archetypes, including:

- **Grid:** Equidistant blocks dispersed across the canvas ( $f = 11$ ; 34.4%), for example, P2-T1.
- **Line:** Blocks placed subsequently in a single direction ( $f = 8$ ; 25%), for example, P10-T1.
- **Clusters:** Blocks systematically organized in multiple graphs ( $f = 6$ ; 18.8%), for example, P1-T1.
- **Collection:** Blocks randomly organized in multiple graphs ( $f = 4$ ; 12.5%) for example, P6-T1.
- **Compact:** Blocks ‘clumped together’ with no dispersion across the canvas ( $f = 3$ ; 9.4%), for example, P4-T1.

Out of all 32 physicalizations created, 27 physicalizations used the physical 3D space to visualize data in an upward direction (height). Only 5 physicalizations were created within the plane, by 4 different participants (Figure 9; indicated by ‘planar’). 4 of these physicalizations were of the *line* archetype, either horizontal or vertical within the canvas, whereas outlier P13-T2 created a *collection* of waffle charts in the canvas (Figure 9; P13-T2).

We observed that for 6 physicalizations diagonal spacing was introduced into the x and/or the y-axis (Figure 9; indicated by ‘diagonal’). P10-T2 created a complete diagonal line visualization,

P6-T2 created a collection of diagonal graphs, P9-T2 and P14-T1 created a grid with a diagonal offset in the x-axis, and P16 created a line of diagonal graphs (T1) and diagonally spaced clusters (T2).

Lastly, P12-T2 created a special case of a *collection*, as the spatiality in the canvas was used to represent a geographical map of the countries, to create a more “impactful” visualization to represent carbon emissions (Figure 9; P12-T2).

**5.2.2 Color association.** For 28 tasks (87.5%) the color of blocks was associated with categorical attributes. Hence, participants used color to differentiate between countries or seasons. In the other 4 tasks (12.5%) color was used to differentiate between years (sequential attributes). Looking at the exact colors that were allocated to categories of the datasets, we observed more consistency in color association with seasons than with countries. Participants explained different approaches to the color mapping, which were either (i) as a utility to separate data ( $f = 12$ ; 37.5%), or (ii) to create a conceptual mapping to familiar concepts ( $f = 20$ ; 62.5%).

For the 16 tasks that involved the dataset on UK rainfall, the most common color allocations were green for spring ( $f = 13$ ), blue for winter ( $f = 12$ ), orange for autumn ( $f = 12$ ), and yellow ( $f = 8$ ) or red ( $f = 5$ ) for summer. For 13 tasks participants consciously allocated color to seasons, based on associations between color and temperature (i.e. blue for a cold winter temperature), or seasonal landscape (i.e. yellow for a “dry climate” during summer). As an outlier, P14-T2 strategically kept the color red aside to highlight extremums in rainfall per year (see Figure 9; P14-T2).

For the 16 tasks that involved the dataset on CO<sub>2</sub> emissions, the most common color associations were orange for Netherlands ( $f =$

7), red for Spain ( $f = 6$ ), red ( $f = 4$ ) or blue ( $f = 4$ ) for the United Kingdom, yellow ( $f = 5$ ) or green ( $f = 4$ ) for Belgium, and blue ( $f = 5$ ) or green ( $f = 4$ ) for Norway. For 6 tasks participants tried to allocate color to countries, based on the colors of their flag (red for Spain) or other colors of national importance (orange for the Netherlands), followed by a process of elimination.


Overall, participants' strategy in the use of color association is dependent on the topic of the dataset. Accordingly, results suggest the adoption of a pragmatic approach to relate colors to familiar concepts first (e.g., color hue with the temperature of seasons, or flags), followed by the association or allocation of the remaining color resources by process of elimination.

**5.2.3 Axis mapping.** The most common mapping of axes we observed was that both sequential and categorical attributes were represented from left to right ( $f = 7$ ; 21.9%). For instance, we observed 6 *line* and 1 *line (diagonal)* archetype displaying this pattern. Besides that, we observed equal occurrences of physicalizations that represented (i) sequential data from left to right, and categorical data from either front to back or back to front, and (ii) categorical data from left to right and sequential data from front to back or vice versa ( $f = 4$ ; 12.5% for each occurrence). Lastly, we observed for 4 physicalizations that one data attribute was represented from left to right, while the other attribute was represented through spatiality. For example, for 3 physicalizations of the *collection* archetype, categorical data was represented from left to right and sequential data was represented using dispersed positioning in the plane.

Although participants generally followed the structure of the data table while constructing their physicalization, the only times a randomization of categorical data took place was for the emissions dataset, for which participants randomized the order of countries, consciously or not ( $f = 6$ ; 18.8%). This also happened a single time for the rainfall dataset, which was adapted by choosing a different season as the starting point for each year.

Hence, we conclude that generally for two data attributes (table top to bottom), if one attribute is represented left to right, the other is represented either back to front or vice versa, with no particular preference for categorical or sequential data in either axis.

**5.2.4 Data labeling position and reading direction.** Overall, for the majority of tasks ( $f = 28$ ; 87.5%) participants placed all labels in their default reading direction (left to right, labels legible from the viewing point). However, we observed different approaches in the positioning and orientation for each label type:



**Figure 10: Different approaches to reference labels:** P4-T1 placed labels for each distinct value on top of the physical bar charts, whereas P5-T2 placed labels on three different sides of the physical constructs to anticipate viewing from multiple orientations.

**Title:** For the majority of tasks the title label was placed on the canvas ( $f = 28$ ; 87.5%) and for 4 (12.5%) on the side or top of placeholder blocks (Figure 9; P1 & P8). Looking at the relative location of the title, for 14 tasks it was placed in the front of the canvas (of which 6 in the center), for 11 tasks in the back (of which 7 in the center), and for 7 tasks in the middle area (of which 3 on the left). Lastly, we observed that 2 participants placed title labels in counterclockwise reading direction ( $f = 3$ ; 9.4%, Figure 9; P1-T1, P11-T1/2) or clockwise direction ( $f = 1$ ; 3.1%, Figure 9; P1-T2).

**Sequential attribute:** For the majority of tasks ( $f = 28$ ; 87.5%) the sequence labels were placed on the canvas alongside the physicalization. For 2 tasks they were placed as a key in the back center of the canvas, either with (Figure 9; P11-T2) or without placeholder blocks to communicate the color mapping (Figure 9; P12-T2). P8-T1 placed the sequence labels on yellow placeholder blocks alongside the physicalization and P10-T2 placed them against the physical data points of the physicalization. Lastly, we observed that 1 participant placed sequential labels in a counterclockwise reading direction ( $f = 2$ ; 6.3%, see Figure 9; P1-T1/2).

**Categorical attribute:** For 18 tasks (56.3%) the category labels were placed on the canvas alongside the physicalization. For 11 tasks they were placed as a key, either on the canvas alongside placeholder blocks ( $f = 6$ ; 18.8%, for example, Figure 9 P2-T2), or on top of the placeholder blocks ( $f = 5$ ; 15.6%, for example, Figure 9 P6-T1). Looking at the relative location of the category key within the canvas, the majority ( $f = 7$ ; 21.9%) was placed in the front of the canvas (of which 4 on the right). For 3 tasks the category labels were placed or attached against data points of the physicalization (Figure 9; P8-T1, P10-T1, and P13-T1). P13 mentioned that for each country bar chart, they placed the country label on the bar with the highest value for visibility. Lastly, we observed that 1 participant placed categorical labels in counterclockwise reading direction ( $f = 2$ ; 6.3%, Figure 9; P1-T1/2).

**Data values:** For 15 tasks (46.9%) all value labels were used to indicate each individual data point, either by placing them on top of each bar chart ( $f = 13$ ; for example Figure 9; P1-T1), or on the canvas in front of each bar chart ( $f = 2$ ; for example Figure 9; P12-T1). For 11 tasks (34.4%) a single value label was used to create a key, either by placing it on the canvas by itself (Figure 9; P10-T1), alongside a placeholder block (Figure 9; P5-T1), or on the top (Figure 9; P2-T2) or the side (Figure 9; P8-T1) of the placeholder. Lastly, there were 6 tasks in which multiple value labels were used to create reference points for data extraction. For example, P3-T1 created a 'legend tower' for sideways height comparison with the bar charts (Figure 6; P3-T1). Likewise, P4 included reference labels for each distinct value on top of the bar charts, as well as included a key at the right side of the physicalizations. However, they explained that when viewed from above, the reference labels allowed for value estimation of bars of similar height (Figure 10; P4-T1). P9-T1 placed reference labels on the canvas in front of the first row of data points (Figure 9; P9-T1), and P13-T1 placed them against the first row of data points (Figure 9; P13-T1). Lastly, in addition to a key, P5-T2 provided reference labels on 3 sides of the bar charts to anticipate for viewing from different orientations (Figure 10; P5-T2). Moreover, we observed that 2 participants placed data value labels in mixed reading directions ( $f = 1$ ; 3.1%, see Figure 9 P1-T1).



In sum, participants placed title labels in a central location on the canvas. Similarly, sequence labels were placed on the canvas, but then alongside one of the sides of the physicalization. In contrast, category labels were placed on the canvas alongside the physicalization, as well as a key separate from the physicalization to encode color mapping. Lastly, for almost half of the tasks all value labels were used to indicate each individual data point, whereas, for a third, a single value was used to create a key.

### 5.3 Influence of Orientation

Herein, we discuss the role of labeling when viewing physicalizations from different orientations. Participants were asked to rotate the canvas with 90 degrees increments and assess their labels (whether they wanted to change the labels to read them effectively and comfortably). We elaborate on the challenges encountered with the physical constructs within the canvas and the coping strategies participants adopted when manipulating labels to more effectively convey the information presented in their physicalizations.


**5.3.1 Challenges of orientation.** The rotation of the canvas introduces viewing perspective challenges that affect the digestion of the presented labels. Taking as the starting point the most common physicalization construction, we will unfold the potential issues encountered during the iterative change of orientation.

We take as a reference a 3D grid of data points with value labels on top of the bars; categorical and sequence labels placed on top of the canvas alongside the bars' rows/columns, and the title label located on the canvas at the front (all labels legible from the viewing point). For instance, after a 90 degrees rotation, all labels are read sideways and categorical/sequential labels are hidden behind stacked blocks. After a second 90 degrees change, labels are displayed upside down and the title label is pushed to the far end of the canvas. As such, each orientation change introduces (i) a change in viewing position affecting label legibility and salience, and (ii) a change in characters/numbers reading direction. These factors introduce the following challenges:

- **Reading Direction** occurs when text is not displayed in the default/legible orientation (characters displayed upwards for ease of reading), but is rotated clockwise, counterclockwise, or is presented upside down, thus introducing higher cognitive demand.
- **Occlusion** occurs when labels are hidden behind block constructs, making viewing from all directions more difficult.
- **Proximity and Organization** occurs when labels are relocated, increasing their distance from the viewing point, and therefore affecting the salience of information and the users' predefined mental model of the physicalization.
- **Ordering and Direction** occurs when the order of labels alters their meaning, hindering the digestion of the information displayed. For example, a sequence of year labels that loses chronological order upon multiple orientation changes.

**5.3.2 Changes to data labeling as a coping strategy.** In our study, participants were invited to modify (as they wished) the display of labels after each viewing orientation iteration. Herein, we elaborate on the changes participants made to the data labeling across the orientation conditions. In total, there were 96 conditions (16 participants  $\times$  2 tasks  $\times$  3 orientations). We did not find any significant

differences between the orientation conditions (clockwise or counterclockwise). Overall, we observed 4 different types of changes made to the data labels (Table 2) listed in order of most occurrence:



**Table 2: Changes made to each label type – title (■), sequence (■), category (■), and value (■) – across the 3 orientations.**

**Rotation** in a (counter)clockwise direction to set the reading direction back to the original default after the orientation change ( $f = 61-72; 63.5-75\%$ ). Although some participants changed the reading direction of all label types (P2, P3, P4), others prioritized changes to the orientation of categorical and sequential labels over title and value labels, specifically when these were upside down after the second orientation change (P6, P13). However, some participants reported not caring about reading direction at all (P1, P5).

**Relocation** of labels within the canvas to avoid occlusion, increase proximity, or preserve organization ( $f = 26-41; 27.1-42.7\%$ ). Generally, participants preferred to relocate title labels over the other types across orientation changes. This could be caused by a desire to maintain the original presentation of the title (P4, P10) or to place the label in a position that is salient and avoids occlusion.

Introduction of an **Offset** in relation to the physicalization to compensate for the occlusion of labels ( $f = 5-6; 5.2-6.25\%$ ). Offset strategies occur when modifying sequential and categorical labels (found alongside the block constructs) as they might get occluded after each change of orientation.

Other outlier changes ( $f = 1-2; 1-2\%$ ), such as the **Re-purposing** of blocks to use them as a key to two different data attributes (e.g., P5-T1 reused the block representing the scale to create a legend for categorical attributes). Moreover, we observed the **Addition** of new blocks to create a category legend and avoid occlusion (P5-T1), or of unused value labels to add detail (P13-T1).

**5.3.3 Changes to physical constructs as a coping strategy.** In addition to the changes to the data labeling during the different orientation iterations, we observed participants' strategies to try and anticipate orientation challenges during the creation process. These strategies emerged from the accumulation of participants' out-loud rationalization of "improvements" across tasks as a response to the changes of orientation experienced and/or anticipated.

The following strategies are a reflection of isolated instances of behaviors observed during the study to provide further evidence of coping mechanisms adopted at the creation level that we aim to be illustrative as much as they could be guiding for future work.

**Space Dispersion and Organization:** 7 participants played with the use of the canvas space (e.g., distancing blocks, centering the physicalization). This affected the organization of data blocks to facilitate the digestion of information and avoid occlusion. For instance, P1-T2 and P2-T2 reported increasing the space between bar charts (**dispersion**), whereas P4-T2 mentioned placing their data blocks in

the middle (*centering*) of the canvas to make it “look good” and have space around them. Similarly, P6-T2 indicated they decided to add space between bar charts so the visualization looked less “messy”, but they were concerned the use of space could convey meaning (e.g. separate different categories) when they aimed to solely improve readability. On the other hand, P9-T2 reported deciding to spread out bar charts so they do not visually block each other, whereas P13-T2 pushed groups of bars as far away possible so they would not “distract” each other. Moreover, P16-T2 described organizing their bar charts so the smallest values (e.g. countries with lower CO<sub>2</sub> emissions) were placed on the outskirts of the canvas, whereas the highest values (e.g. countries with higher CO<sub>2</sub> emissions) were placed at the center so they would not be occluding.

*Introducing Diagonal Offset:* 3 participants experimented with the addition of a diagonal offset between data values. For instance, P10-T2 increased the separation in both the x and y-axis to create a “diagonal” line rather than mapping values on a single axis. Moreover, P9-T2 introduced a diagonal offset to display their grid as a rhomboid rather than a square, whereas P6-T2 introduced a diagonal offset for each bar chart in a collection archetype.

*Addition of Key Placeholder:* 5 participants introduced the use of blocks as key placeholders or legends. This was aimed to avoid the occlusion of labels behind blocks as legends were pulled away from the location of the physicalization structure. For instance, P8-T2 discussed their addition of a key aimed to facilitate looking at it from any possible angle. Additionally, P5-T2 mentioned placing a legend centered within the canvas to anticipate “hidden” labels after a 90 degrees turn, whereas P13-T2 wanted to use the free space available in the middle of the canvas to place all the information necessary to read their visualization (a legend for categories’ color mapping and sequential labels to indicate organization).

*Experimenting with Archetypes:* 3 participants experimented with the use of the canvas space, thus changing the composition of their physicalization and creating a different archetype (e.g., moving from a 3D visualization to a planar one). For instance, P13-T2 mentioned “making it flat” and avoid building different stories to facilitate understanding the data from every angle (and tackle occlusion). P2-T2 discussed the trade-off of the use of planar visualization as it introduces directionality (i.e., once rotated 90 degrees it looks “sideways”), which P15-T2 felt was limiting even though a planar visualization could remove occlusion problems.

*Highlighting:* 1 participant (P14-T2) decided to *highlight* the extremums of the data values with different colored blocks to improve the visualization of minimum and maximum values at a glance without necessitating to estimate height differences in the 3D space.

## 6 DISCUSSION

We investigated the role of data labeling in the physicalization creation process, the visualization design, and the resilience of data labels across orientations. Our findings show that (i) label activities are alternated and/or intertwined with block activities during the creation process, (ii) labels are integrated with physical constructs in the final visualization design, and (iii) this relation between data labels and physical constructs is influenced by orientation changes. Overall, our results suggest that the use of data labels is fundamental to consider for future physicalization designs.

### 6.1 Towards A Principled Use of Data Labels in Physicalization Design

Although physicalizations embody data in their material and physical form [30], they still benefit from the inclusion of contextualizing elements (i.e. data labels, axes, legend, and annotations) to support the extraction of information from the physical representation. However, despite the evident importance of providing context to visualizations, most related work on physicalization is not labeled at all [e.g. 25, 39, 54]. Physicalizations that do use contextual elements are often inconsistent or specific to that individual design [e.g. 22, 29, 45]. As the current definition of physicalization [29] suggests, the focus is on *physicality* and not on ‘data labeling’ or other contextual elements of *the physicalization in use*. Moreover, the physical and spatial nature of physicalizations introduces additional challenges, as it remains unclear where to locate different kinds of labels and how they accommodate multi-user scenarios. Hence, there is currently no principled way of contextualizing physical representations of data.

The field of Information Visualization has established ways to discuss and implement the contextualization of digital data representations [19, 23]. However, it remains unclear how this translates to the field of physicalization. Implementations of 2D visualizations in the field of InfoVis are more homogeneous than 3D representations of data. Hence, some variance will always exist in the data labeling of 3D physical constructs. Nonetheless, it would be useful to aim for the development of a collection of ‘best practices’, guidelines, or at least illustrative work to, as a research field, become more strategic at contextualizing physicalization design.

It is apparent that our specific apparatus aids in the creation of physicalizations of the ‘bar chart aesthetic’. However, it still allowed participants to create a variety of visualization archetypes going beyond the traditional use of bar charts. We observed that across these different archetypes, the use of data labels was consistent: the majority of data labels were placed in default reading direction and were paired and/or integrated with physical constructs (i.e. value labels on top of data points). Moreover, labeling was used in combination with other visualization components such as color encoding and axis mapping. As such, future work could investigate whether similar use of data labels, and similar integration of data labels with physical constructs will occur for a variety of physicalizations.

### 6.2 Utility of Labeling in the Physicalization Creation Process

Constructive Visualization work [17] previously explored how the use of physical tokens can support the authoring of physical data representations. However, these approaches focus on the construction of visual mappings, and thus far did not actively include the use of data labels in the authoring process. Instead, the labeling or annotation of data is treated as a subsequent process to the construction process [27, 52], or their use is left up to participant preference [17]. As a result, it remains unclear what role data labeling can and/or should have in the creation process of physicalizations.

As the act of data labeling is part of a larger process of construction and contextualization, we decided to study it in the context of a constructive visualization process. Hence, we designed a toolkit that follows state of the art methodology [27, 52], with the inclusion

of both physical tokens as well as textual labels to investigate the use of data labels during the creation process. Our findings show that this allowed participants to alternate and/or intertwine label and block activities during the creation process. This illustrates the utility of active inclusion of data labels for physicalization creation. Moreover, we observed that the use of data labels can serve different purposes: to plan the visualization before including physical constructs, to guide the creation of subsequent physical constructs, and to verify constructs afterward. Thus, the use of data labels allows verifying physical constructs ad hoc, in particular when the label and block activities are heavily intertwined. Hence, the inclusion of data labels could provide people with more agency within the creation process of physicalizations.

The *extended infovis pipeline model* [46] describes the contextualization of physicalizations as ‘*decoration*’ operations as part of *presentation mapping*. However, we observed that labeling activities can occur across different pipeline operations, such as the *loading* of data by ordering data labels in the workspace, or as part of *visual mapping* as they are organized as elements in the canvas alongside block constructs, before the final *presentation mapping* takes place.

To explain this, we take interest in the *interrelation principle*. Wun et al. [52] described this principle as the intertwined nature of operations due to the physical nature of the authoring tool. However, as they did not actively include data labels in the toolkits discussed, no reflections are provided on how labeling fits within this principle. We argue that similar interrelated processes occur for data labels as for physical tokens. To give an example, ordering data in the workspace outside the canvas is *loading data* [46]. However, the moment data labels are introduced in the canvas, relations are created between the data label and (i) other data labels, (ii) other block constructs, and (iii) relative position within the canvas. As such, data labels could be considered as building blocks in themselves, not just complementary to physical constructs.

Although Huron et al. [27] provide a conceptual flow diagram of common construction behaviors, this does not include the act of annotation as it happened as a secondary task after construction. Arguably, the act of appropriating data labels within the canvas and in relation to physical constructs can be described through those diagram elements as well (i.e. organize, arrange, merge, align) and should be considered alongside physical tokens in the process. Hence, it might be necessary to expand existing conceptual models and/or introduce new models as data labeling is an interrelated process within itself, and in relation to construction activities.

### 6.3 Data Label Resilience across Orientations

Prior work has demonstrated the influence of orientation, introducing ambiguity when extracting information from physical representations of data [40], and discusses the different types of occlusion that can occur due to user orientation. In line with this work, we observed challenges for effective use of data labels due to orientation changes: the correction of reading direction, prevention of occlusion, and maintenance of proximity and organization.

We argue that the introduction of data labels can mitigate the challenges introduced by physical 3D space, such as directionality, occlusion, and user multiplicity. Whereas the use of duplicate data labels might seem a straightforward solution, the necessity for

duplicates would ‘clutter’ the visualization. To simplify cognitive digestion, we argue for the use of *reactive* and *resilient data labels*.

Reactive data labels can accommodate the point of view of the user, and solve occlusions created through physical constructs. To acquire this, two parallel processes would need to happen: (i) data labels follow the point of view of the user to maintain reading direction and proximity (*user-label relation*), but are also reactive to (ii) the physical composition or layout of the physicalization, to prevent occlusion and maintain effective offsets (*label-layout relation*). If this is done successfully, it results in a *user-label-layout* relation that supports effective extraction of information from physical data representations for any orientation. Our results on coping strategies through a change in data labels (and to some level physical constructs) are illustrative for ways in which future physicalization designs could counteract orientation influences (such as the rotation, relocation, and offset of labels). Depending on the system implementation, these strategies can be informative for the design of reactive data labels and/or adaptable physical constructs:

For static physicalizations [e.g. 29, 43] data label resilience needs to be high, as the physical construct is rigid and cannot adapt to viewing angle and/or perspective changes. Hence, accommodation for orientation influences is fully dependent on data label design and adaptability. To give an example, data labels follow the viewer orientation to adapt reading direction, and if a physical construct gets occluded in a particular orientation, the label can ‘float’ above or aside the construct to notify the viewer of its existence.

For dynamic and interactive physicalizations [e.g. 16, 18, 45] data label resilience can interplay with the specific actuation technologies implemented. Hence, data label design and/or physical construct actuation counteract orientation influences in parallel. For instance, if a physical construct gets occluded, actuation can ‘move’ it aside to maintain the line of sight and the data label follows.

Moreover, on top of the interplay of data labels and actuation, interaction could also play a role. For instance, users could indicate ad hoc what information they require and manipulate the data labels and/or physical constructs accordingly. Our observation of isolated instances of strategies to cope with orientation through the change of physical constructs resonates with prior work on reconfiguration strategies [41]. Herein, they found that proximity change was generally the most used strategy to rearrange physical constructs, which relates to the organization and dispersion of physical constructs we observed.

Lastly, the introduction of multiple users and/or a collaborative context creates new challenges for data labeling as well. For effective information extraction by collaborators, there is a necessity for either maintaining a shared view versus the introduction of individual viewports. For example, a shared view could be accomplished through top-down projection or display integration in each physical data points, whereas individual viewports could be accomplished through an AR overlay or VR environment.

### 6.4 Opportunities for Future Work

In our study, we focused on a subset of physicalizations – 3D bar charts – that are well-established in the field (i.e. [18, 45]). Hence, we can not make conclusive statements on the labeling of other types of physicalizations or even other implementations of 3D bar charts.

Additionally, other label designs (i.e. curved, embossed, transparent, 3D), different ways of attaching labels and construction strategies, and/or more participants' agency in designing their own labels could generate diverse outcomes. Hence, future work is needed to expand on our initial findings for these particular conditions, to further investigate the role of data labeling in the creation, design, and mitigation of physicalizations with orientation challenges. Lastly, in the present study, we did not record further demographics (i.e. occupation, cultural background, native language) that could have been of influence on the observed labeling behaviors.

First, future work could further compare the different strategies for labeling we observed in the creation process. It could be valuable to compare the design outcomes of post-hoc, pre-hoc, and interrelated labeling activities. Moreover, we observed that data labeling can serve different purposes (i.e. to plan, guide, or verify a physicalization), hence, it could be further investigated what other purposes labeling can have beyond the creation process, such as self-reflection or as part of the presentation to others.

Second, although our apparatus allowed for the creation of different visualization archetypes, further investigation would be necessary to explore the data labeling of physicalizations beyond the bar chart aesthetic. Subsequently, our study is illustrative of coping strategies through a change in data labels (and to some level physical constructs), but is not an exhaustive list of how to contextualize physicalizations in general. Hence, future work could investigate the labeling of other types of physicalizations, and expand on coping strategies for challenges due to physical space.

Third, there are some biases introduced by the characteristics of our apparatus: the structure of the data table could influence participants order of creating constructs, and the use of an actual dataset introduces recognition bias for the ones familiar with the specific topics. Moreover, the current dataset was two-dimensional (1 sequential and 1 categorical attribute), hence, we cannot postulate results for other datasets that are more or less complex, i.e. a more complex dataset with multi-dimensional data, requiring creation in multiple axes. Hence, future work would need to investigate how our findings translate for other datasets and toolkits.

Lastly, as our focus was on the use of data labels for contextualization, the methodology was designed to allow for data label alterations but not for changes to physical constructs. Hence, future work is needed to develop further understanding of the interplay between label resilience and adaptability of physical constructs.

## 7 CONCLUSION

In this paper, we investigated the role of labeling in the creation process, final physicalization design, and when viewed from different orientations. We designed a custom toolkit including physical tokens and textual labels, and asked 16 participants to complete a total of 32 construction tasks. Our findings show that (i) the creation of physicalizations is an intertwined process of labeling and construction activities, (ii) resulting in an integrated visualization design of data labels and physical constructs, and (iii) these integrated labels and constructs are influenced by orientation changes. Hence, we argue for further development of contextualization methods for future physicalizations, and propose the introduction of *reactive data labels* to counteract challenges of orientation.

## REFERENCES

- [1] Pankaj K Agarwal, Marc Van Kreveld, and Subhash Suri. 1998. Label placement by maximum independent set in rectangles. *Computational Geometry* 11, 3-4 (1998), 209–218.
- [2] Kamran Ali, Knut Hartmann, and Thomas Strothotte. 2005. Label layout for interactive 3D illustrations. (2005).
- [3] Jacques Bertin. 1983. *Semiology of graphics; diagrams networks maps*. Technical Report.
- [4] Rahul Bhargava and Catherine D'Ignazio. 2017. Data sculptures as a playful and low-tech introduction to working with data. (2017).
- [5] Richard Brath. 2014. 3D InfoVis is here to stay: Deal with it. In *2014 IEEE VIS International Workshop on 3DVis (3DVis)*. IEEE, 25–31. <https://doi.org/10.1109/3DVis.2014.7160096>
- [6] Grigore Burdea and Philippe Coiffet. 2003. Virtual reality technology.
- [7] Chaomei Chen. 2013. *Information visualisation and virtual environments*. Springer Science & Business Media.
- [8] Ed Huai-hsin Chi. 2000. A taxonomy of visualization techniques using the data state reference model. In *IEEE Symposium on Information Visualization 2000. INFOVIS 2000. Proceedings*. IEEE, 69–75.
- [9] Jon Christensen, Joe Marks, and Stuart Shieber. 1995. An empirical study of algorithms for point-feature label placement. *ACM Transactions on Graphics (TOG)* 14, 3 (1995), 203–232.
- [10] Kurtis Danyluk, Teoman Tomo Ulusoy, Wei Wei, and Wesley Willett. 2020. Touch and Beyond: Comparing Physical and Virtual Reality Visualizations. *IEEE Transactions on Visualization and Computer Graphics* (2020). <https://doi.org/10.1109/TVCG.2020.3023336>
- [11] A De Boer. 2007. Label placement in 3D georeferenced and oriented digital photographs using GIS technology. (2007).
- [12] Hessam Djavaherpour, Faramarz Samavati, Ali Mahdavi-Amiri, Fatemeh Yazdankhsh, Samuel Huron, Richard Levy, Yvonne Jansen, and Lora Oehlberg. 2021. Data to Physicalization: A Survey of the Physical Rendering Process. *Computer Graphics Forum* (2021). <https://arxiv.org/abs/2102.11175>
- [13] Pierre Dragicevic and Yvonne Jansen. 2012. List of Physical Visualizations. [www.dataphys.org/list](http://www.dataphys.org/list).
- [14] Pierre Dragicevic, Yvonne Jansen, and Andrew Vande Moere. 2021. Data Physicalization. In *Springer Handbook of Human Computer Interaction*, Jean Vanderdonck (Ed.). Springer. <https://hal.inria.fr/hal-02113248>
- [15] Shawn Edmondson, Jon Christensen, Joe Marks, and Stuart Shieber. 1996. A general cartographic labelling algorithm. *Cartographica: The International Journal for Geographic Information and Geovisualization* 33, 4 (1996), 13–24.
- [16] Aluna Everitt, Faisal Taher, and Jason Alexander. 2016. ShapeCanvas: An Exploration of Shape-Changing Content Generation by Members of the Public. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). ACM, New York, NY, USA, 2778–2782. <https://doi.org/10.1145/2858036.2858316>
- [17] Danyang Fan, Alexa Fay Siu, Sile O'Modhrain, and Sean Follmer. 2020. Constructive Visualization to Inform the Design and Exploration of Tactile Data Representations. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility (Virtual Event, Greece) (ASSETS '20)*. ACM, New York, NY, USA, Article 60, 4 pages. <https://doi.org/10.1145/3373625.3418027>
- [18] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogue, and Hiroshi Ishii. 2013. InFORM: Dynamic Physical Affordances and Constraints through Shape and Object Actuation. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (UIST '13). ACM, New York, NY, USA, 417–426. <https://doi.org/10.1145/2501988.2502032>
- [19] Office for National Statistics (ONS). 2021. Chart Titles and Text. <https://style.ons.gov.uk/category/data-visualisation/titles-and-text/>.
- [20] Helen Gibson, Joe Faith, and Paul Vickers. 2013. A survey of two-dimensional graph layout techniques for information visualisation. *Information visualization* 12, 3-4 (2013), 324–357.
- [21] Silke M Göbel. 2015. Up or down? Reading direction influences vertical counting direction in the horizontal plane—a cross-cultural comparison. *Frontiers in psychology* 6 (2015), 228.
- [22] Pauline Gourlet and Thierry Dassé. 2017. Cairn: A Tangible Apparatus for Situated Data Collection, Visualization and Analysis. In *Proceedings of the 2017 Conference on Designing Interactive Systems* (Edinburgh, United Kingdom) (DIS '17). ACM, New York, NY, USA, 247–258. <https://doi.org/10.1145/3064663.3064794>
- [23] US Government. 2021. Data Visualisation Standards. <https://xdgov.github.io/data-design-standards/components/labels>.
- [24] Tovi Grossman, Daniel Wigdor, and Ravin Balakrishnan. 2007. Exploring and Reducing the Effects of Orientation on Text Readability in Volumetric Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 483–492. <https://doi.org/10.1145/1240624.1240702>
- [25] Steven Houben, Connie Golsteijn, Sarah Gallacher, Rose Johnson, Saskia Bakker, Nicolai Marquardt, Licia Capra, and Yvonne Rogers. 2016. Physikit: Data Engagement Through Physical Ambient Visualizations in the Home. In *Proceedings*

- of the 2016 CHI Conference on Human Factors in Computing Systems. ACM, New York, NY, USA, 1608–1619. <https://doi.org/10.1145/2858036.2858059>
- [26] Samuel Huron, Sheelagh Carpendale, Alice Thudt, Anthony Tang, and Michael Mauere. 2014. Constructive Visualization. In *Proceedings of the 2014 Conference on Designing Interactive Systems* (Vancouver, BC, Canada) (DIS '14). ACM, New York, NY, USA, 433–442. <https://doi.org/10.1145/2598510.2598566>
- [27] Samuel Huron, Yvonne Jansen, and Sheelagh Carpendale. 2014. Constructing Visual Representations: Investigating the Use of Tangible Tokens. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2102–2111. <https://doi.org/10.1109/TVCG.2014.2346292>
- [28] Yvonne Jansen and Pierre Dragicevic. 2013. An Interaction Model for Visualizations Beyond The Desktop. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2396–2405. <https://doi.org/10.1109/TVCG.2013.134>
- [29] Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. 2013. Evaluating the Efficiency of Physical Visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Paris, France) (CHI '13). ACM, New York, NY, USA, 2593–2602. <https://doi.org/10.1145/2470654.2481359>
- [30] Yvonne Jansen, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. 2015. Opportunities and Challenges for Data Physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). ACM, New York, NY, USA, 3227–3236. <https://doi.org/10.1145/2702123.2702180>
- [31] Jia Li, Catherine Plaisant, and Ben Shneiderman. 1998. Data object and label placement for information abundant visualizations. In *Proceedings of the 1998 workshop on New paradigms in information visualization and manipulation*. 41–48.
- [32] Simon Lok and Steven Feiner. 2001. A survey of automated layout techniques for information presentations. *Proceedings of SmartGraphics 2001* (2001), 61–68.
- [33] Daphne Menheere, Evianne van Hartingsveldt, Mads Birkebæk, Steven Vos, and Carine Lallemand. 2021. Laina: Dynamic Data Physicalization for Slow Exercising Feedback. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (DIS '21). Association for Computing Machinery, New York, NY, USA, 1015–1030. <https://doi.org/10.1145/3461778.3462041>
- [34] Andrew Vande Moere and Stephanie Patel. 2009. The physical visualization of information: designing data sculptures in an educational context. In *Visual information communication*. Springer, 1–23.
- [35] Tamara Munzner. 2014. *Visualization Analysis and Design*. CRC Press, Boca Raton, FL, USA.
- [36] Kayleigh Porter and Gemma Arblaster. 2020. How Does Vertical Reading Affect Reading Speed? *The British and Irish orthoptic journal* 16, 1 (2020), 38.
- [37] He Ren and Eva Hornecker. 2021. Comparing Understanding and Memorization in Physicalization and VR Visualization. In *Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Salzburg, Austria) (TEI '21). ACM, New York, NY, USA, Article 41, 7 pages. <https://doi.org/10.1145/3430524.3442446>
- [38] Kim Sauvé, Saskia Bakker, and Steven Houben. 2020. Econundrum: Visualizing the Climate Impact of Dietary Choice through a Shared Data Sculpture. In *Proceedings of the 2020 Conference on Designing Interactive Systems* (Eindhoven, Netherlands) (DIS '20). ACM, New York, NY, USA, 1287–1300. <https://doi.org/10.1145/3357236.3395509>
- [39] Kim Sauvé, Saskia Bakker, Nicolai Marquardt, and Steven Houben. 2020. LOOP: Exploring Physicalization of Activity Tracking Data. In *Proceedings of the 2020 Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (NordiCHI '20). ACM, New York, NY, USA, Article 52, 12 pages. <https://doi.org/10.1145/3419249.3420109>
- [40] Kim Sauvé, Dominic Potts, Jason Alexander, and Steven Houben. 2020. A Change of Perspective: How User Orientation Influences the Perception of Physicalizations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '20). ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/3313831.3376312>
- [41] Kim Sauvé, David Verweij, Jason Alexander, and Steven Houben. 2021. Reconfiguration Strategies with Composite Data Physicalizations. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, Article 471, 18 pages. <https://doi.org/10.1145/3411764.3445746>
- [42] William R Sherman and Alan B Craig. 2003. Understanding virtual reality. *San Francisco, CA: Morgan Kaufman* (2003).
- [43] Simon Stusak, Jeannette Schwarz, and Andreas Butz. 2015. Evaluating the Memorability of Physical Visualizations. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 3247–3250. <https://doi.org/10.1145/2702123.2702248>
- [44] Simon Stusak, Aurélien Tabard, Franziska Sauka, Rohit Ashok Khot, and Andreas Butz. 2014. Activity Sculptures: Exploring the Impact of Physical Visualizations on Running Activity. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 2201–2210. <https://doi.org/10.1109/TVCG.2014.2352953>
- [45] Faisal Taher, Yvonne Jansen, Jonathan Woodruff, John Hardy, Kasper Hornbæk, and Jason Alexander. 2017. Investigating the Use of a Dynamic Physical Bar Chart for Data Exploration and Presentation. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (2017), 451–460. <https://doi.org/10.1109/TVCG.2016.2598498>
- [46] Teoman Ulusoy, Kurtis Thorvald Danyluk, and Wesley J Willett. 2018. *Beyond the Physical: Examining Scale and Annotation in Virtual Reality Visualizations*. Technical Report. Department of Computer Science, University of Calgary.
- [47] Annemiek Veldhuis, Rong-Hao Liang, and Tilde Bekker. 2020. CoDa: Collaborative Data Interpretation Through an Interactive Tangible Scatterplot. In *Proceedings of the Fourteenth International Conference on Tangible, Embedded, and Embodied Interaction* (Sydney NSW, Australia) (TEI '20). ACM, New York, NY, USA, 323–336. <https://doi.org/10.1145/3374920.3374934>
- [48] Eliane Zambon Victorelli, Julio Cesar Dos Reis, Heiko Hornung, and Alysson Bolognesi Prado. 2020. Understanding human-data interaction: Literature review and recommendations for design. *International Journal of Human-Computer Studies* 134 (2020), 13–32.
- [49] Frank Wagner, Alexander Wolff, Vikas Kapoor, and Tycho Strijk. 2001. Three rules suffice for good label placement. *Algorithmica* 30, 2 (2001), 334–349.
- [50] Colin Ware. 2019. *Information Visualization: Perception for Design*. Morgan Kaufmann, Cambridge, MA, USA.
- [51] Daniel Wigdor and Ravin Balakrishnan. 2005. Empirical Investigation into the Effect of Orientation on Text Readability in Tabletop Displays. In *ECSCW 2005*, Hans Gellersen, Kjeld Schmidt, Michel Beaudouin-Lafon, and Wendy Mackay (Eds.). Springer, 205–224.
- [52] Tiffany Wun, Jennifer Payne, Samuel Huron, and Sheelagh Carpendale. 2016. Comparing Bar Chart Authoring with Microsoft Excel and Tangible Tiles. *Computer Graphics Forum* 35, 3 (2016), 111–120. <https://doi.org/10.1111/cgf.12887> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12887>
- [53] Deyue Yu, Heejung Park, David Gerold, and Gordon E Legge. 2010. Comparing reading speed for horizontal and vertical English text. *Journal of vision* 10, 2 (2010), 21–21.
- [54] Jack Zhao and Andrew Vande Moere. 2008. Embodiment in Data Sculpture: A Model of the Physical Visualization of Information. In *Proceedings of the 3rd International Conference on Digital Interactive Media in Entertainment and Arts* (Athens, Greece) (DIMEA '08). ACM, New York, NY, USA, 343–350. <https://doi.org/10.1145/1413634.1413696>