Author copy of:
Ricardo Langner, Lonni Besançon, Christopher Collins, Tim Dwyer, Petra Isenberg, Tobias Isenberg, Bongshin Lee, Charles Perin, and Christian Tominski. "Introduction to Mobile Data Visualization". In: Mobile Data Visualization. Chapman and Hall/CRC, 2021. Chap. 1, pp. 1–32. doi: 10.1201/9781003090823-1

CHAPTER 1

An Introduction to Mobile Data Visualization

Ricardo Langner
Technische Universität Dresden, Germany

Lonni Besançon
Linköping University, Sweden

Christopher Collins
Ontario Tech University, Oshawa, Canada

Tim Dwyer
Monash University, Australia

Petra Isenberg
Université Paris-Saclay, CNRS, Inria, LISN, France

Tobias Isenberg
Université Paris-Saclay, CNRS, Inria, LISN, France

Bongshin Lee
Microsoft Research, Redmond, Washington, USA

Charles Perin
University of Victoria, British Columbia, Canada

Christian Tominski
University of Rostock, Germany

CONTENTS

WE may have an intuitive understanding of what is meant by mobile data visualization. Yet, in the context of data visualization, the term *mobile* can be interpreted in several ways. For example, it may describe visual representations shown on devices that are inherently mobile. It may also describe visualizations meant to react to viewers who are mobile relative to the display. Alternatively, it may describe visualizations that are themselves mobile across devices and screens, or in space.

In this chapter, we propose several characteristics that help us to identify and describe the scope of mobile data visualization. We focus on the characteristics that, particularly in their extremes, differentiate mobile data visualization from other forms of data visualization. These characteristics give rise to dimensions of a design space for mobile data visualization, against which instances may be classified and positioned. We discuss a number of examples to illustrate how the design space makes it possible to describe and compare mobile visualizations.

1.1 INTRODUCTION

Gleaning knowledge from data, so-called data analytics, has become a massive industry. Understanding data is no longer the concern of only government and business, but has become a significant component in the life of most people in the developed world. This rise of data analytics has been driven to a large extent by increased computer automation in every aspect of industry and modern life, as well as by the ease of sharing data over the internet. Given the ubiquitous availability of data, people have turned to, amongst other methods, data visualization as a critical tool to understand, experience, explore, and communicate data.

Data and visualizations can now be accessed from almost anywhere, anytime and a vast array of different devices can allow viewers to see and explore data leveraging visualizations. As such, interacting with data visualization is no longer an activity that can only happen on desktop computers, or even laptops or tablets. Smartphones, fitness trackers, and smartwatches—also, more exotically, e-readers, handheld gaming devices, smart glasses, and even augmented-reality headsets—represent the variety of mobile computing devices available today. It is time, therefore, to reconsider data visualization as it relates to this diversifying ecosystem of possibilities.

Mobile data visualization is a nascent research area that aims to take advantage of the new forms of ubiquitous data analysis and communication offered by technological advancements. Hence, it is perhaps not surprising that a shared understanding of its scope is not yet established among researchers and practitioners. Several well-founded perspectives can be adopted. For example, the term mobile data visualization may refer to *visualizations hosted on devices that are mobile* or *visualizations that react to viewers who are mobile relative to the display.* One might also describe it as *visualizations that are themselves mobile across devices and screens.*

The purpose of this chapter is not to identify a connotative definition but to introduce our interpretation of mobile data visualization and lay a conceptual foundation by discussing its scope. The chapter is motivated by our efforts to build a common understanding, language, and discussion basis. This basis is fundamental to building a community of researchers and practitioners around the topic of mobile data visualization.

We started by collecting several core cases, which intuitively would be considered clear examples of mobile data visualization. The set of core examples was then expanded through in-depth discussions of niche examples for which it is more difficult to agree upon whether they are cases of mobile data visualization. To better articulate **why** we considered certain cases to be *core* archetypal examples of mobile data visualization and others to be *edge* cases that in some way stretch the term, we derived a set of characteristics. With the help of these characteristics, the scope of mobile data visualization can be defined, and existing approaches, methods, and techniques can be discussed and categorized more easily. The dimensions that characterize the scope of mobile data visualization will be described next in Section 1.2. Illustrating the central aspects and also the extremes of mobile data visualization, we will discuss core examples and edge cases in Sections 1.3 and 1.4, respectively.

1.2 CHARACTERIZING DIMENSIONS

As modern computing is both portable and ubiquitous, a data visualization can be *mobile* in many different ways. To facilitate discussions on ‘mobile-visness,’ we identify seven descriptive dimensions that are relevant for categorizing existing work:

- Physical data display size
- Data display mobility
- Data source
- Reaction of visualization to display movement
- Intended viewing timespan
- Visualization interaction complexity
- Intended sharing

The dimensions emerged from examining characteristic similarities but also differences of existing mobile data visualization examples. We note that, while we think each dimension has its relevance when discussing mobile data visualization, there are certainly differences in their importance and complexity. In addition, there is some overlap: for example, solutions designed for a rather short *Intended Viewing Timespan* (Section 1.2.5) often also call for a more passive or simple level of *Visualization Interaction Complexity* (Section 1.2.6). Nonetheless, the identified dimensions allow us to reason about existing visualization techniques and devices in terms of mobility. What is more important than thinking about what already exists is to consider the possibility the dimensions' extremes bring, which might give insight into what mobile data visualization might become in the future, with advances in technology and imagination. Let's look at the dimensions in detail.

1.2.1 Physical Data Display Size

Looking at today's mobile devices, it is clear that the *Physical Data Display Size* is an essential aspect. Mobility seems closely related to physical screen size. We distinguish pixel-sized, watch-sized, phone-sized, tablet-sized, monitor-sized, and wall-sized displays as illustrated in Figure 1.1.

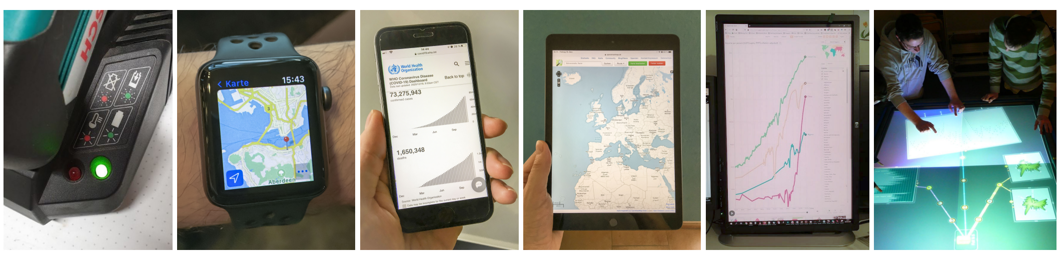

(a) pixel-sized (b) watch-sized (c) phone-sized (d) tablet-sized (e) monitor-sized (f) wall-sized

Figure 1.1: *Physical Data Display Size* dimension, ordered from smallest to largest.

Even though we focus on examples of flat digital data displays because they are most common, we acknowledge that non-flat data displays exist, for example, in the context of data physicalization. Another important thing to note is that the size of a data display with respect to a viewer's field of view depends on the viewer's distance to the display. This is particularly evident for augmented reality (AR) and virtual reality (VR) headsets. They are physically in the phone-sized display range but are worn so close to the wearer's eyes that they practically cover all of the wearer's field of view. Here, we care about the physical size of the display representing the data and not the apparent size for the viewer. We also assume that resolution is not a limiting factor as the trend to high pixel density displays continues across form-factors. Moreover, the following categories show that this dimension also has a considerable influence on how users interact with the device and presented information.

Pixel-sized: Very small data displays, in the range of a few to several millimeters fit into this category. Examples include single LEDs that show battery charging levels or error states of a machine as illustrated in Figure 1.1a. There is typically no interaction with these very small displays.

Watch-sized: Smartwatch displays are typically around 3–4 cm wide or high. Some deviate from a standard rectangular form to a circular geometry (Figure 1.1b), which is an interesting design constraint for visualization. Smartwatches are large enough to convey information to someone being relatively close to the display. Interaction is typically direct on the display using touch, through buttons or a digital crown on the device, and sometimes with speech.

Phone-sized: Smartphone screens as shown in Figure 1.1c are now commonly around 15 cm on the diagonal and have a high resolution of more than 150 pixels per cm. Compared to watch-sized displays, phone-sized displays can convey more information but still require a relatively close proximity of the viewer. Interaction is typically direct through the display via touch or through buttons on the device. AR and VR headsets also fit into this category but, by design, cover a large field of view and require dedicated forms of interaction.

Tablet-sized: Tablet-sized data displays (Figure 1.1d) typically have a book-like form factor. These displays can easily show more than one information panel or view, and thus support more complex visualizations. Interaction is typically direct through the display via touch, pen, or through buttons on the device.

Monitor-sized: Monitor-sized displays can cover a fair portion of the viewer's field of view even when the viewer is positioned further away (Figure 1.1e). Visualizations consisting of multiple views become more practical on these displays. While interaction typically is indirect using a mouse and keyboard, modern monitors can also be touch-enabled.

Wall-sized: Large displays commonly used in conference room, control room, or trading floor belong to this category. Wall-sized displays (Figure 1.1f) are large enough for multiple people to comfortably view them from varying viewing distances. If interaction is available (sometimes it is not), it often is direct via either touch or pen input, or indirect through connected devices.

1.2.2 Data Display Mobility

As already noted, a small display size can afford mobility. Another key aspects of mobile data visualization is the *Data Display Mobility*, which captures the movement of the display(s) containing visual representations of data. Fixed, movable, carryable, wearable, and independently moving displays can be differentiated along this dimension. Corresponding examples are given in Figure 1.2.

While the *Data Display Mobility* dimension is generally organized from least mobile to most mobile, the spectrum is not quite linear. For example, the difference between a wearable and a carryable display is fuzzy and fluid. A runner strapping a smartphone showing their running data onto their arm makes the display wearable

(a) fixed (b) movable (c) carryable (d) wearable (e) independently moving

Figure 1.2: *Data Display Mobility* types ordered roughly from least to most mobile.

while someone else may put their smartwatch showing weather data into their pocket, thereby, carrying it rather than wearing it.

In this dimension we limit our discussion to visualizations that rest on the displays that show them—or that are themselves the data displays (as in physical representations of data, or *data physicalizations*). We do not consider cases in which visualizations move from one to another display. For example, we disregard the scenario of a large display room in which someone analyzes a visualization on a tablet and then pushes the visualization to a shared large display to be worked on with a group.

Fixed displays: A desktop monitor, a tabletop display, a wall-size display, a large data sculpture, a public display—or any other large, typically stationary, data display—all fit in this category. While displays in this category are not inherently mobile, they still may relate different notions of mobility in some way or the other. The digital tabletop display in Figure 1.2a, for example, is set up to accommodate mobile viewers. A fixed display may also react to mobility in the environment, for example, by showing more or less detail depending on a mobile viewer's distance to the display [19].

Movable displays: Displays that cannot be carried for extended periods, but can be moved with the help of some supporting device belong to the category of movable displays. Examples include the display of computers on wheels (COWs) commonly used in hospital settings (see Figure 1.2b), as well as displays and devices, such as the Microsoft Surface Hub 2 that come with a movable stand.

Carryable displays: A data display is considered carryable if it can be moved without supporting devices. Such displays can be easily carried, for example in a bag or a pocket. As everything from a laptop computer to a phone (see Figure 1.2c) belongs to this category, it represents many of today's consumer mobile devices. Consequently, most of the existing mobile visualization research has been conducted in this category.

Wearable displays: Data displays that can be worn on a person and thus do not need to be actively carried belong to this category. Examples include data jewelry, smartwatches (Figure 1.2d), fitness bands, smartglasses, or augmented clothing. Wearable data displays may have a greater degree of responsiveness to user movement or gaze (for example, becoming immediately available when movement is detected) than carryable data displays. This class of devices is becoming increasingly important for visualization research.

Independently moving displays: Data displays that move autonomously or without direct human propulsion belong to this category. Examples include data displays attached to drones or robots as shown in Figure 1.2e. Humans may be controlling the movement remotely, but to the viewers the data displays would seem to be moving independently. See Section 1.4 for existing examples of independently moving data visualizations.

1.2.3 Data Source

In addition to *Physical Data Display Size* and *Data Display Mobility*, it is also possible to distinguish the source of the data being visualized. Some visualizations show pre-loaded data, while others need a connection to fetch data from cloud storage, for example to display weather and stock data. Yet others visualize live data captured from sensors, for example, step counts, GPS, and WiFi signals. It is also possible to combine different data sources, for example, when showing live location data and cloud-based traffic data on a map. Accordingly, we have divided this dimensions into four categories: pre-loaded, connected, captured, and combination (Figure 1.3).

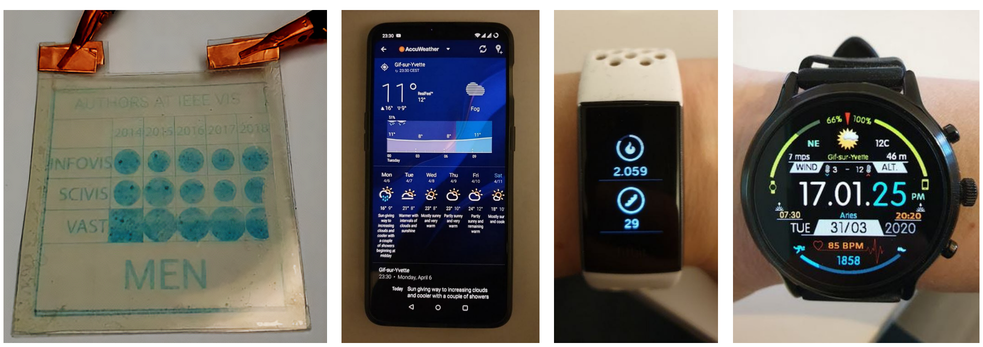

(a) pre-loaded (b) connected (c) captured (d) combination

Figure 1.3: Categories in the *Data Source* dimension.

Pre-loaded data: The data have been previously loaded on the device and thus are static. Pre-loaded data are usually neither time critical nor affected by the environment of the device. The translucent display in Figure 1.3a, for example, can by design only show this and one other dataset depending on the flow of the current through the display.

Connected data: The data arrive dynamically over data connections from online services, such as a WiFi, Bluetooth, or cellular connection. The visualization responds to and may highlight updates to the data. Weather visualizations on smartphones, as in Figure 1.3b, or smartwatches are common examples where data is transmitted from external servers.

Captured data: The data shown in the visualization are generated by the device itself, for example, by capturing them through on-board sensors. Captured data are a form of dynamic data source but without a connection to a server. The mobile device itself needs to mediate the data (i. e., aggregate, filter, or present the data in a consumable form). The fitness band in Figure 1.3c shows captured data as two simple radial progress bars of calories burned and floors climbed.

Combination: It is also common to complement a primary data source with one or more secondary data sources. This is helpful when a mobile device can only partially sense the environment, is not able to process larger amounts of data, or is used in an environment with restricted or limited connectivity. Smartwatches, for example, often show dashboards that combine data captured from the device (battery life) and data from external servers (weather) as shown in Figure 1.3d.

1.2.4 Reaction of Visualization to Display Movement

Another distinctive aspect of mobile data visualization is if and how a visualization changes due to the movement of the display in the environment (not any movement of the viewer). If the visualization changes, the question is if these changes are directly related to the movement or if there is a rather indirect connection. We identify four broad categories for this dimension: no change, indirect change, direct change, and direct + indirect (Figure 1.4). More details on how mobile visualizations may react to movement and other dynamic factors can be found in Chapter 2.

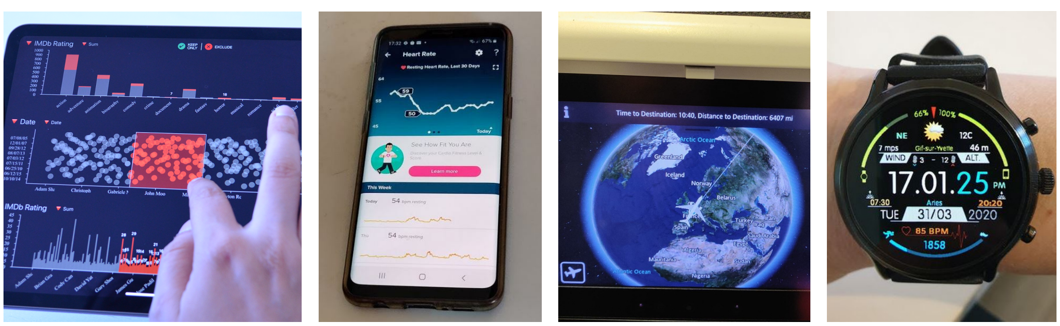

(a) no change (b) indirect change (c) direct change (d) direct + indirect

Figure 1.4: Categories in the *Reaction of Visualization to Display Movement* dimension. *Figure a) reprinted with permission from Ramik Sadana.*

No change: Visualizations in this category are not linked to and hence do not change on movement of the display. The visualized data typically has nothing to do with potential display movements. An example are movie data as shown in Tangere [34] in Figure 1.4a.

Indirect change: Visualizations in this category change due to movement but they visualize data that is only indirectly related to the movement. In other words, they show and update data that is *affected* by the movement, such as heart rate and EEG signals. Figure 1.4b shows a fitness app's heart rate visualization, in which the heart rate increases when the owner is moving more vigorously.

Direct change: Visualizations that show data being related to the movement directly belong to this category. Movement-related data include step counts, velocity, or position and location as for example in the in-seat display in Figure 1.4c, which shows a plane's location relative to the Earth.

Direct + indirect: Visualizations that show both directly and indirectly movement-related data belong to this category. Smartwatch faces as in Figure 1.4d are a common type of data display dashboard that includes both direct (e.g., step counts) and indirect (e.g., heart rate) visualizations reacting to movement.

1.2.5 Intended Viewing Timespan

Another perspective of mobile data visualization opens up when considering the time available for viewing a visual representation. The *Intended Viewing Timespan* can be decoupled from the screen size and is likely highly related to the context of use. For example, a smartwatch visualization intended for being viewed while running the outdoors will need a different design than a tablet visualization to be viewed while sitting in a comfortable armchair. Similarly, a view on a smartphone has different characteristics when intended for glancing while running compared to focused analysis during a meeting. According to how much time someone spends with a mobile data visualization, we consider the categories: sub-second (glance), seconds, minutes, hours or more (Figure 1.5). More details on glanceable mobile visualizations from the first category will be given in Chapter 5.

(a) sub-second (glance) (b) seconds (c) minutes (d) hours or more

Figure 1.5: Categories in the *Intended Viewing Timespan* dimension.

Sub-second (glance): An intended viewing timespan of a few hundred milliseconds often arises in situations when people's attention is directed elsewhere and they can only briefly take their eyes away from their primary task. GPS devices in moving vehicles, such as the car GPS in Figure 1.5a, belong to this category.

Seconds: Contexts with relatively simple information needs such as today's or tomorrow's weather, as shown by the app in Figure 1.5b, require visualizations to be read within a few seconds. Visualizations in this category typically support simple comparisons, timelines, and incorporate familiar chart types.

Minutes: Longer analysis times in the range of minutes are expected in visualizations that have more complex information needs. For example, the star map in Figure 1.5c requires viewers to orient the phone at the part of the sky they are interested in and to compare the sky with the representation on their phone. More generally, common tasks required in these contexts are comparing many items, navigating a large dataset, or analyzing multiple attributes.

Hours or more: In-depth analysis of complex data, which require extensive interaction and highly specialized visualization techniques may require viewing timespans of hours or more. For example, dedicated analysis environments such as Tableau Mobile (Figure 1.5d) or tools for situated awareness in law enforcement [30] allow for in-depth data exploration.

1.2.6 Visualization Interaction Complexity

In relation to the impression that mobile data visualization might imply simpler or minimal interaction, this dimension describes the complexity of the main interaction as an interplay between a person and the visualization. Compared to traditional data visualization, mobile data visualization needs to support usage scenarios where people are on-the-go or are engaged in other activities. Analyzing a data visualization may not always be a person's primary task, but rather an auxiliary step to enhance or support other mobile tasks. We categorize four levels of visualization interaction complexity ranging from passive interaction to highly interactive as shown in Figure 1.6. For an in-depth discussion of interaction for mobile data visualization, we refer to Chapter 3.

Passive interaction: The visualizations in this category allow for minimal or no interaction. The focus is on providing pure consumption of information as for the viewer in Figure 1.6a. This category is also related to glanceable visualizations that are characterized in Chapter 5.

Simple view specification: Here, visualizations allow for discrete view switching, which involves changing between the display of different types of data, often also involving a change of representation. The person interacting with a visualization in Figure 1.6b is 'swiping' to switch between different representations.

View specification & manipulation: At this stage, visualizations support standard interactions, such as selection, details-on-demand, navigation (for example, with pan and zoom), and so on. The person interacting with a heatmap in Figure 1.6c just selected a single cell to call a tooltip with more information.

(a) passive interaction (b) simple view specification (c) view specification & manipulation (d) highly interactive

Figure 1.6: Four categories in the *Visualization Interaction Complexity* dimension.

Highly interactive: Visualizations in this category include a broad spectrum of sophisticated interactions. Beyond the types of interactions mentioned above, they may support the full sense-making workflow, including visualization authoring, keeping track of the analysis process, as well as annotation and externalization of insight. For example, the analyst in Figure 1.6d is performing a series of interactions to systematically review information for a group of cells, while bookmarking cells of interest and importance.

1.2.7 Intended Sharing

Mobile data visualizations offer the opportunity for a variety of sharing scenarios, from highly personal use to collaborative use of shared displays. Our focus is on synchronous sharing in which visualizations are viewed at the same time through a shared device or a set of connected devices. Here, we discuss to which extent a visualization is meant to be shared with others: personal use, a few people, larger groups, and the general public (Figure 1.7).

While the intended sharing is affected by the display form-factor (it is more challenging to share a view on a smartwatch than a tablet), these two dimensions are separable. The mobility of the visualization may affect the scenarios in which sharing can take place, for example, from more fixed location-dependent groups at home to dynamically formed opportunistic groupings in a public setting. The sharing might also depend on the personal nature of the data being displayed.

Personal use: The visualization is viewed in a private context, for example, on a small display that cannot be shared with others easily. These visualizations often contain personal data which may have strict privacy considerations or may be of utility only to a single individual. For example, period calendar visualizations or health visualizations such as the migraine visualization in Figure 1.7a are sensitive and meant primarily for a single person.

(a) personal use (b) a few people (c) larger groups (d) general public

Figure 1.7: Categories in the *Intended Sharing* dimension. *Figure d) © Ron Levit, reprinted with permission [24].*

A few people: Visualizations can also be for sharing with co-located people in a small group setting, for example, when viewing a map together on a smartphone. In Figure 1.7b, we see a visualization of a visitor's trip progress that can be looked at together with the family.

Larger groups: Visualizations may be shared with larger groups to support teamwork and decision making. For example, the large display in Figure 1.7c can be wheeled into a meeting room to support the coordination of larger group activities by means of a visualization of project plans.

General public: The visualization is displayed in a public setting or can be viewed by many people simultaneously on a personal device. Figure 1.7d shows an example of a public opinion visualization [24]. It can be driven through town and collect data by people passing by.

Overall, we have now presented seven dimensions each with four to six categories according to which mobile data visualizations can be organized. Along the described dimensions, we can mark ranges that are typical for mobile data visualization. We will find small, mobile displays with visualizations that show data that are fetched from the cloud or captured via sensors. Mobile data visualizations will have interactions of moderate complexity and also react to the display movement. A mobile visualization is typically in use only for shorter time spans and also touches upon the aspect of sharing it with others. These seven dimensions also allow us to reason about extremes. For example, a mobile data visualization is not necessary small as can be seen from the truck display in Figure 1.7d. A mobile data visualization might be used only for the fraction of a second on one extreme, for example, when viewers only glance at them as for the navigation display in Figure 1.5a. On the other end of the spectrum, a mobile data visualization might also be used for hours, for example, in the context of supporting law enforcement as in Figure 1.5d.

These few examples already illustrate that mobile data visualization covers a considerable range of designs. We will expand on concrete examples and corresponding categorizations along the characterizing dimensions in the next section.

1.3 TYPICAL EXAMPLES OF MOBILE DATA VISUALIZATION

As mentioned earlier, a crisp definition of mobile data visualization is difficult to articulate, as computing devices converge and propagate to more and more different types of activity. The academic literature and web collections such as MobileVis [32] and Mobile Infovis [35] showcase various mobile data visualizations developed by researchers, practitioners, and technology companies. In this section, we discuss selected examples and relate them to the dimensions presented in the previous section.

We will first look at typical examples of visualizations on smartwatches, phones, and tablets, where most of us would immediately agree they are clearly and intuitively categorized as mobile data visualizations. On the other hand, the dimensions' categories are sometimes blurry and the extremes lead us to examples that point us in Section 1.4 to more creatively think about how mobility affects visualization.

1.3.1 Early Mobile Data Visualizations for PDAs

Research on mobile data visualization began in the nineties [13, 20]. The term mobile computing emerged to describe people's interactions with computing devices that are wirelessly connected, able to exchange information, and portable. With the wide availability of personal digital assistants (PDAs), the topic of mobile data visualization also began to gain momentum [10].

The driving slogan for the works in this early research period was to provide access to information anytime anywhere. The key research challenges were primarily focused on technical issues due to limited computing power and memory. As can be seen in the example on an early PDA in Figure 1.8a, the display capabilities were limited. The display resolution was low (240×320 pixels in our example) and early PDAs could display only a few colors—if color was available at all. Therefore, the visualization had to be designed carefully to avoid wasting precious pixels or relying on color-related visual channels. At the same time, the implementation had to be efficient: any overhead by run-time libraries or interpreted languages had to be avoided to keep the visualization reactive and to reduce battery drain.

Another key difference to today's mobile data visualizations is the much lower network bandwidth, causing data to reach the mobile device at a snail's pace. Therefore, the bulk of the information to be visualized had to be stored on the mobile device. Only small pieces of information could be transmitted over the network on demand.

A classic example of a mobile data visualization from this time period is the DateLens [3] (Figure 1.8b). DateLens is a focus+context interface that grants users access to their calendar at any time on carryable handheld PDAs. While PDAs had a form factor that is similar to that of today's smartphones (*Physical Data Display Size*), DateLens had to cope with a screen resolution of only 240×320 pixels.

Although some PDAs were equipped with GPS sensors, DateLens did not react to the device owner's mobility (*Data Display Mobility*). The primary input modalities of PDAs were pen-based interaction and dedicated buttons. In terms of *Visualization Interaction Complexity*, DateLens supported a rich set of interactions that allowed users to navigate in time and adjust the focus+context display of the calendar.

The schedule visualized in DateLens resided on the PDA, yet, the schedule was

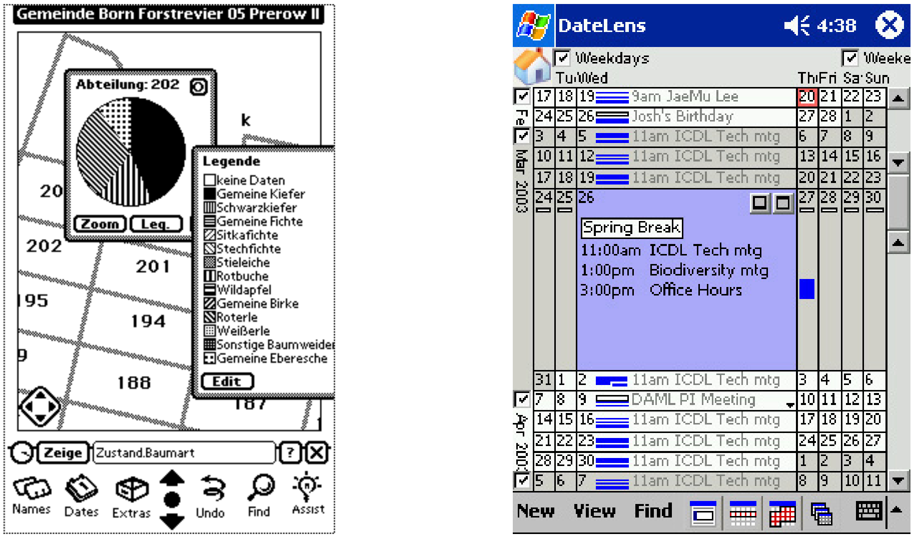

(a) Monochrome visualization of forestry-related data on a PDA. *Reprinted from Kirste and Rauschenbach [20] with permission from Elsevier.*

(b) DateLens visualizing a person's calendar [3]. *Taken from Windsor Interfaces, Inc. [43] with permission from Ben Bederson.*

Figure 1.8: Early mobile data visualizations on PDAs.

not static. It was possible to create new appointments or change or delete existing ones. Naturally, DateLens had an *Intended Viewing Timespan* between seconds (for example, when glancing at the calendar for quick confirmation that there are no upcoming events today) and minutes (for example, when searching the calendar for specific instances of events). Sharing the calendar display with a few colleagues was certainly possible (*Intended Sharing*), but not mentioned in the paper.

DateLens is just one classic example of a mobile data visualization. Visual representations of data can nowadays be found on a wide variety of devices with different form factors. Next, we describe three examples using the most common form factors with high device mobility: smartwatch, smartphone, and tablet.

1.3.2 Mobile Data Visualizations for Smartwatches

When thinking about mobile data visualization, one might picture a small, simple, glanceable visualization on a smartwatch (Figure 1.9). While watch-sized screens are not new, they have gained an increased popularity in recent years with, for example, the Apple watch or Fitbit activity trackers. Recently, we have seen some research efforts in understanding how data and representations are currently displayed on smartwatch faces [18] and how people perceive small-scale visualizations on a smartwatch [7, 15, 27].

Far from the pixelized black & white devices they used to be, smartwatches are becoming powerful computing devices. Despite their small screen size, (high-end) smartwatches can display colorful images at a high resolution (often more than 150

Figure 1.9: Mobile data visualizations on a Garmin Forerunner 245 watch.

pixels per cm). They are equipped with a state-of-the-art GPS and several on-device sensors, such as accelerometer, gyroscope, and magnetometer. They can be connected to other devices to retrieve and display a wider range of information, including emails, calendar events, and phone call history. This plethora of data being available on smartwatches is a great opportunity for visualization. Smartwatches offer various faces and widgets for people to customize what data to show and how to show them.

Visualizations on GPS watches for runners, as shown in Figure 1.9, are canonical examples of mobile data visualization. The visualizations are on a wearable (*Data Display Mobility*), watch-sized device (*Physical Data Display Size*), and change both directly and indirectly due to movement (*Reaction of Visualization to Display Movement*). The visualized data are typically captured from the device through GPS and sensors, but can be combined with data downloaded from a server (*Data Source*). The visualization design is optimized for on-the-go use while the person wearing the device is running. That said, glancing at the visualization is only a subordinate activity (*Intended Viewing Timespan*). Also the *Visualization Interaction Complexity* is low, mostly involving only passive interaction and simple view specifications. The reason behind this simplicity is obvious: Fiddling with the visualization for too long would increase the risk of tripping or causing other dangerous situations. The visualizations are designed for personal use (*Intended Sharing*). In short, visualizations on GPS watches for runners are a good exemplar of a small, personal visualization, providing glanceable information with limited interaction, for on-the-go use.

1.3.3 Handheld Mobile Data Visualizations

Many existing mobile data visualizations are designed for carryable rather than wearable devices. Examples of existing smartphone visualizations can be reviewed through the web collections, *MobileVis* [32] and *Mobile InfoVis* [35]. Nevertheless, examples of handheld mobile data visualization also include early applications for PDAs (see Section 1.3.1) as well as those for other smartphone-sized tools such as handheld GPS. Just like visualizations for smartwatches, those for handhelds are also designed mostly for personal use.

Smartphones, however, are not as specialized and can be described as a universal personal device with broad functionalities. The successful combination of mobility, portability, and display size as well as the subsequent extended interaction capabilities are reasons why nowadays people are accessing and consuming a majority of digital information with smartphones. The use of smartphones in particular is often linked to actual personal activities and usage scenarios that include a broad range of mobility types: standing at a bus station, sitting on a bench, or lying in bed.

Applications in the context of handheld mobile data visualizations allow people to visually inspect, for example, performance, such as in sports, health measures, such as sleep quality and blood sugar level, product information, such as its components and ingredients, route and navigation information, and departure and transfer times in public transportation. One particularly illustrative example for handhelds is Goddemeyer and Baur's Subspotting app [14]. It visualizes the available mobile phone reception along tracks of the New York City Subway. The motivation behind this application is that the expected constant network connectivity of smartphones makes times or places with poor network coverage a challenge for many use cases. The Subspotting app, therefore, shows the previously measured and recorded network coverage for and along a specific route (Figure 1.10), allowing train riders to better understand and decide "where to send the next text or make the next call" [14].

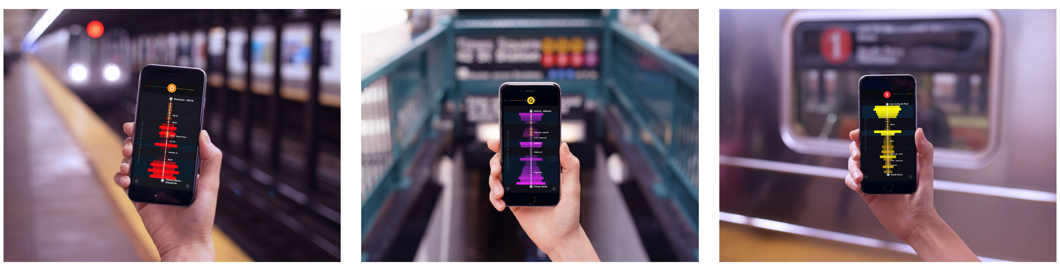

Figure 1.10: The Subspotting app by Goddemeyer and Baur [14] visualizes available mobile phone reception along the lines and stations of the New York City Subway. *Images © 2016 OFFC NY, used with permission from Dominikus Baur.*

With regard to our dimensions of mobile data visualization, this example shows that a visualization on a smartphone (*Physical Data Display Size*) can change due to the movement of the device, both directly and indirectly (*Reaction of Visualization to Display Movement*). The varying strength of a WiFi signal can be displayed during the journey as an indirect indication of position, and the actual physical position on the route can be marked. The *Visualization Interaction Complexity* is at a level of view specification and manipulation: Users can switch between discrete views, and each view also allows for navigation via zooming and panning. The *Data Source* is mainly static and pre-loaded into the application, although dynamic on-board sensor data such as strength of a WiFi signal could be integrated. The visualizations are designed in such a way that the data can be viewed briefly for a few seconds or for several minutes (*Intended Viewing Timespan*). The application and its underlying data are available for the public but the visualizations shown on a device are not designed for sharing (*Intended Sharing*).

1.3.4 Mobile Data Visualizations on Tablet Devices

Tablet devices are like a hybrid of a smartphone and a laptop, with the smartphone's enhanced mobility and touch input and the laptop's larger display. Both companies and developers of data visualization products—for example Tableau, Microsoft Power BI, and Datawrapper—are well aware of the importance of mobile solutions for their customers and, therefore, provide mobile versions of their products as well. Many research examples also are specifically designed for tablet devices, such as Tangere [33, 34] in Figure 1.11a and InChorus [39] in Figure 1.11b, and clearly belong to this core group of mobile data visualizations. At the same time, many of them are expected to be used when the device is in a stationary setting, unlike those designed for smartwatches or smartphones. Visualization research on tablet devices so far has been centered on designing and developing (mostly touch) interactions with existing visualizations rather than developing novel visual representations that might be more appropriate for tablet devices, especially while they are carried around. More details on novel interactions for visualizations on tablet devices can be found in Chapter 3.

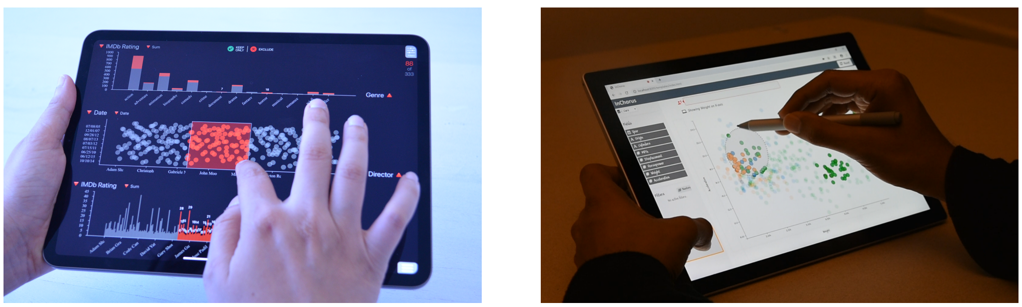

(a) Tangere [34]   (b) InChorus [39]

Figure 1.11: Example visualizations specifically designed for tablet devices. *Left image © Ramik Sadana, used with permission.*

Tablet devices and smartphones share a few characteristics (especially from the technical specification point of view): both are carryable (*Data Display Mobility*), can be equipped with a similar set of sensors, can have similar network connectivity (WiFi, Bluetooth, cellular connection), and can have similar pixel resolutions. In addition, in the extreme case, the size of the largest smartphones is close to that of the smallest tablets (*Physical Data Display Size*). However, unlike visualizations for smartphones, those for tablet devices are typically designed for scenarios that require a set of sophisticated interactions (*Visualization Interaction Complexity*) lasting for more than several minutes (*Intended Viewing Timespan*). It is possible to visualize the data captured by a sensor and to react to display movement (*Reaction of Visualization to Display Movement*), for example, a map visualization with the current position overlaid. Visualizations on tablet devices typically allow people to preload (tabular) data to explore from files (*Data Source*). In terms of *Intended Sharing*, they are more for a personal use but since tablet devices are a little bigger than smartphones they can also be shared with a few people, if needed. We note that visualizations—including their visual representation and interaction design—for tablets are not readily transferable to smartphones, which calls for more research on responsive visualization as further discussed in Chapter 2.

Although the first version of the iPad came with a display of 1024×768 pixel resolution in 2010, the resolution has continuously and significantly improved over the last decade: the resolution of the latest iPad devices is as good as (if not better than) laptop displays and desktop monitors. One thing to note is that, as evidenced by the 2-in-1 laptops, the line between tablet devices and laptops has blurred even though tablet devices tend to imply the absence of a mouse and physical keyboard. Somewhat reflecting this hardware trend, earlier visualization research on tablet devices concentrated on showing only one visualization that fit the entire display (for example, TouchWave [2], TouchViz [12], or the initial version of Tangere [33]), while some of the more recent research started to include multiple views (for example, the later version of Tangere [34], SmartCues [40], or InChorus [39]).

1.4 EDGE CASES OF MOBILITY-RELATED VISUALIZATIONS

Apart from the previous examples for smartwatches, phones, and tables, which in a certain way describe the stereotype of mobile data visualization, there are also cases that extend and even stretch the boundaries of what mobile data visualization seems to be in different directions. Likewise, some of these examples also illustrate what mobile data visualization might be in the future.

Sometimes mobility is not in the device itself. For example, an interactive floor [41] is a fixed room-sized device, making it hardly a candidate for mobile data visualization. However, it is linked to the notion of mobility in that it reacts to people's movements. Such an interactive floor can visualize pre-loaded, connected and captured data, support multiple intended viewing timespans, and allow sharing with a relatively large number of people.

In the following, we will see four more exceptional examples of visualizations in mobile contexts: self-propelled visualizations, visualization in the context of micro-mobility, visualization in hybrid virtual environments, and large movable visualizations.

1.4.1 Self-propelled Visualizations

Devices that can bring themselves to a person to show them a data visualization—rather than the person approaching or simply carrying the device—may seem fanciful, but prototypes of exactly this idea are being tested in research labs. Advances in autonomous robotics are likely to make such devices more and more practical. The appeal of self-propelled billboards to advertisers is likely to create a ready market for such devices, as has arguably been the case for public display walls and projection technologies. Data visualization might be able to take advantage of the same technology.

Yamada et al. [45] developed a self-propelled display device, called iSphere. As shown in Figure 1.12, the iSphere device is a flying drone, surrounded by an array of rotating LED strips to create spherical persistence of a vision display with an effective resolution of 144×136 pixels, at 24 frames per second and with 32 bit color. Despite

Figure 1.12: A more advanced version of the iSphere [45] with an effective display resolution of 760 × 320 pixels. *Images © 2020 NTT Docomo, used with permission from Wataru Yamada.*

its short battery life (only a few minutes long) and the relatively pixellated image quality, the authors were already able to demonstrate some compelling and potentially important applications—other than advertising. Perhaps the most compelling is the idea of such display drones guiding survivors to safety, or perhaps first responders to survivors, in a disaster situation. The display can show arrows and text, such as "Follow," but it is easy to imagine scenarios involving the display of more complex information to first responders. It is ideal in this scenario, because it does not require the public to have any specifically preconfigured device, it can *actively* get their attention, and its physical presence can be both reassuring and commanding.

The following analysis of the iSphere device according to our seven dimensions of mobile data visualization assumes such a disaster response scenario. The device is independently moving in terms of the *Data Display Mobility* dimension. The *Physical Data Display Size* (88 cm in diameter) is monitor-sized. The *Reaction of Visualization to Display Movement* is direct: for example, displaying information pertinent to the current location, to face the user. *Visualization Interaction Complexity* is passive: while not likely to be interactive (a user should avoid attempting to touch the highspeed spinning LEDs), voice interaction might be a possibility (if a mounted microphone can pickup speech over the motor noise). *Data Source* is a combination of connected and captured, for example, on-board carbon monoxide and thermostat sensor information could be displayed to firefighters. The *Intended Viewing Timespan* is seconds: given the battery life, but also the motor noise, viewing timespan is likely brief. Finally, *Intended Sharing* is general public: this device is intended to be viewed by many people simultaneously in public settings.

1.4.2 Micro-mobility for Visualization

In addition to the dimensions mentioned above, the scope or form of mobility is another relevant and characterizing factor of mobile data visualizations. Most people typically relate mobility of a device to the idea of carrying and using it while they move in the world. However, as with most other physical artifacts around us, we can move mobile devices in a much more local or limited scope. Marquardt et al. [25] write about such micro-mobility, as "the fine-grained orientation and repositioning of objects so that they may be fully viewed, partially viewed, or concealed from other persons." There are several examples, such as Conductor [16], Thaddeus [44], Is Two Enough?! [28], VisTiles [23], or The Role of an Overview Device [9], that all make general use of the mobility of mobile devices but in mostly stationary settings.

VisTiles [23], as an example, is based on the idea of enabling co-located collaborative work with information visualizations by using the combination and spatial arrangement of multiple mobile devices (Figure 1.13). Essentially this allows the use of coordinated & multiple views [31, 42] that are displayed and linked across devices. By repositioning and orientating devices, or by even setting up specific side-by-side device arrangements, users can adapt the interface on the table according to requirements of actual situations.

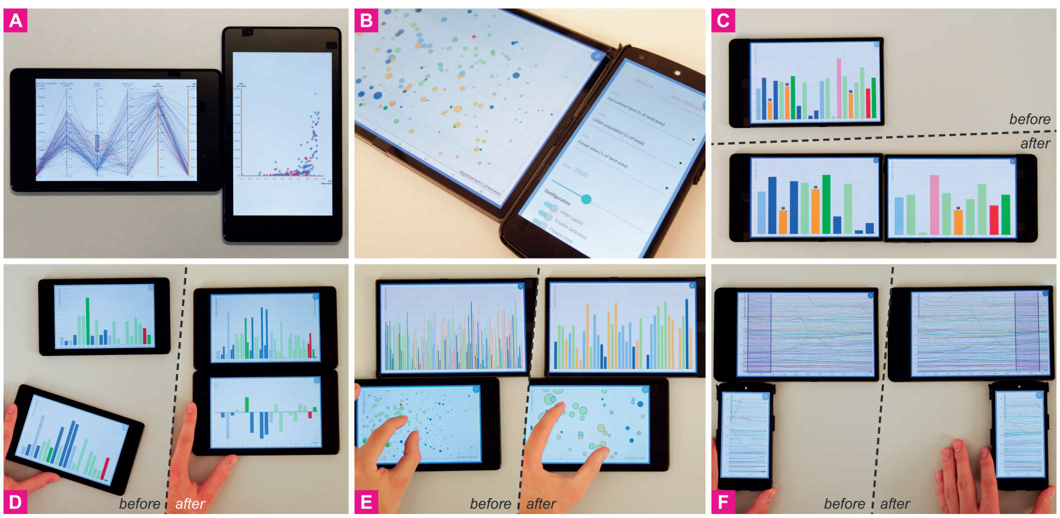

Figure 1.13: VisTiles [23] allows to interact with coordinated & multiple views that are distributed across multiple mobile devices. *© 2018 IEEE. Reprinted, with permission, from Langner et al. [23].*

With regard to *Data Display Mobility* and *Physical Data Display Size*, this system works with carryable and phone-sized or tablet-sized devices. *Data Source* is preloaded because devices download and then visualize data. Interestingly, visualization views respond when a side-by-side device arrangement is recognized, which is why the *Reaction of Visualization to Display Movement* is indirect. In contrast to many other core mobile data visualizations, VisTiles's *Visualization Interaction Complexity* is highly interactive and its design considers an *Intended Viewing Timespan* of hours or more. The *Intended Sharing* is personal and—more importantly—for a few people, as colleagues might use such a system collaboratively at a meeting table.

1.4.3 Mobile Data Visualizations in Multi-display Environments

Combining multiple output devices to form enhanced and augmented viewing spaces has long since been an interesting prospect for visualization research. So called multi-display environments, display ecologies, or hybrid virtual environments offer plenty of display space for visualization and various ways for interactively exploring data [11, 29]. Still more visualization scenarios unfold when considering mobile devices in addition to the mostly stationary devices in such environments.

The appeal of mobile devices in multi-display environments is threefold. First, people nowadays use smartwatches and smartphones regularly, so they are quite proficient in operating these devices [5]. Second, bringing mobile devices to multi-display environments makes it possible to utilize the devices' mobility as well as their output and input capabilities to augment the environment, which in turn can make certain tasks easier to accomplish. Finally, a combination of displays potentially solves problems of sharing information on mobile devices and drastically increases their originally limited display space.

Typically, the devices used in multi-display environment are tablets or smartphones; nevertheless, Horak et al.'s David Meets Goliath [17] show that smartwatches could also be integrated. Mobile devices can be combined with a variety of additional displays and devices. Song et al. [38], Langner and Dachselt [22], Besançon et al. [4, 6], Sollich et al. [37], and Kister et al. [21] envisioned combining spatially-aware tablets or smartphones with large vertical screens (Figure 1.14) for different visualization tasks. Badam et al. developed Munin [1], a framework to seamlessly transition between mobile devices and desktop environments. Instead of large vertical screens, Sereno et al. [36] combined tablets or smartphones with head-mounted displays, while Miguel et al. [26] investigated using such devices in CAVEs.

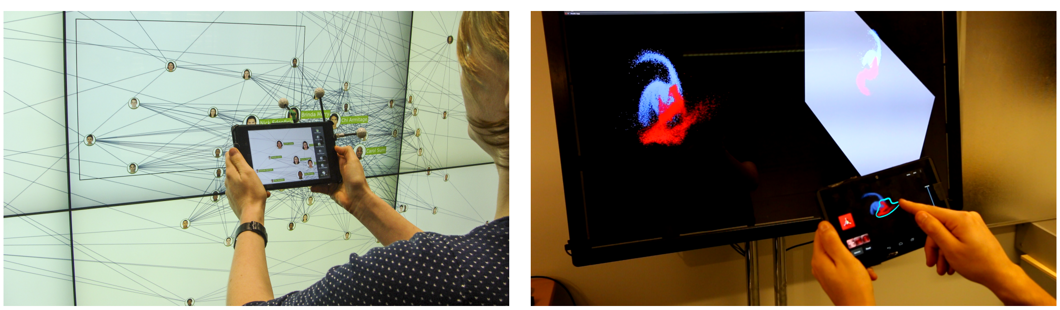

Figure 1.14: Combining a tablet with a larger vertical screen: *(left)* the GraSp system [21], *image © Konstantin Klamka, used with permission;* (right) *Tangible Brush [6], used with permission.*

In terms of our dimensions of mobile data visualization, the use of mobile devices in multi-display environments spans multiple categories. The *Data Display Mobility* is generally in the range of carryable and wearable, as such devices are flexible and afford direct manipulation for interaction—although other output devices in the environment are usually fixed (table or wall) displays. The *Physical Data Display Size* is, again, likely to be small for mobile devices and large for fixed devices. The combination of different displays often allows users to switch between differently sized visualizations or to augment visualizations with additional contextual information. In terms of *Reaction of Visualization to Display Movement*, the primary purpose is to support direct interaction using spatially-tracked devices. Given the focus of many of these hybrid approaches in supporting highly interactive scenarios, the *Visualization Interaction Complexity* is generally quite high. *Data Source* may be pre-loaded or connected. As many of the examples described above are intended for sophisticated visual analytics, the *Intended Viewing Timespan* is likely in the minutes if not hours—although the weight of handheld mobile devices may be a limiting factor. From the point of view of *Intended Sharing*, there is a span from personal use via the handheld and wearable to shared with colleagues or the public.

1.4.4 Large Movable Displays

Visualizations shown on large movable displays certainly touch upon aspects of mobility, but they are not immediately associated with mobile data visualization. As a thought-provoking exercise, the authors of this chapter sketched the idea of a truck loaded with a large visualization cruising the city to inform the public about certain data. The *visualization truck* is clearly a exceptional case of mobile data visualization. And, it turned out that the idea is not new at all.

In Figure 1.15 on the left you see visualizations carried by horse-drawn wagons as part of a New York City parade in 1913 [8]. The small print explanation reads: "Many very large charts, curves and other statistical displays were mounted on wagons in such a manner that interpretation was possible from either side of the street. The Health Department, in particular, made excellent use of graphic methods, showing in most convincing manner how the death rate is being reduced by modern methods of sanitation and nursing."

Figure 1.15: The left image shows statistical displays on wagons during a 1913 parade in New York [8]. *Image is in the public domain.* The right image shows a truck for collecting participant opinion in the form of a stacked bar chart. © *Ron Levit, used with permission [24].*

A modern example of a visualization truck is shown to the right in Figure 1.15. Ron Levit designed this truck as a 'traveling' data visualization game in which participants could leave their opinion creating a unit-based opinion visualization [24]. While the creation of the visualization would be in-place, the truck could later start moving and show the results of people's opinions.

Large movable displays, here in the form of visualization trucks, are clear borderline cases for mobile data visualization. They are displays that move independently of the viewer (*Data Display Mobility*) and are quite large, even wall-sized (*Physical Data Display Size*). One can envision truck-attached visualizations that are connected to the internet or a GPS device and change due to the truck's movement but in our two examples above, there is no change (*Reaction of Visualization to Display Movement*) and the data is pre-loaded (*Data Source*). The *Intended Viewing Timespan* depends largely on the speed of the truck or wagon but is likely in the minutes and sharing is intended for the general public (*Intended Sharing*). As such, the *visualization truck* is quite different from the core cases we discussed above—but it is in the realm of mobile data visualization.

1.5 SUMMARY AND REFLECTIONS

In this chapter, we addressed the question of what is mobile data visualization. We identified seven dimensions for characterizing mobile data visualization and presented a wide range of mobile data visualizations, ranging from common but typical examples, such as visualizations on smartphones and smartwatches, to more exotic examples, such as drone-mounted visualizations, hinting at a pervasive role for mobile data visualization in a future society. All of our examples have in common that they visualize data in some way, and that they are mobile, but in a variety of senses of mobility. Our set of dimensions is useful as a classification system to describe existing approaches in the context of mobility and visualization. A compact overview of our examples in relation to the characterizing dimensions is provided in Table 1.1.

As computing itself becomes more and more diverse, with devices evolving and shrinking further into clothing and other wearable and flexible forms, and as more and more computing is done in the cloud instead of on local devices, the possibilities for mobile data visualization will continue to grow and diversify. Chapter 9 reflects this with a particular view on ubiquitous visualization.

We struggled with the definition of mobility as it is a moving target. In some sense, everything is mobile depending on your frame of reference. A more serious question that we had to contend with in this chapter was what should be the largest frame of reference that we would still consider mobile? Is a visualization truck as just mentioned mobile? Or, consider the flight path visualization on the seatbacks on an airplane as in Figure 1.16. Is it mobile? The answer in both cases is certainly, yes, they are mobile: the seatback display with respect to the Earth and the truck with respect to the observing bystander. Yet, both examples contradict with our intuition of a mobile data visualization being portable and handheld.

In our considerations of what to cover in this chapter, we have taken a common sense approach: If someone looks at a visualization as intended, would they plausibly

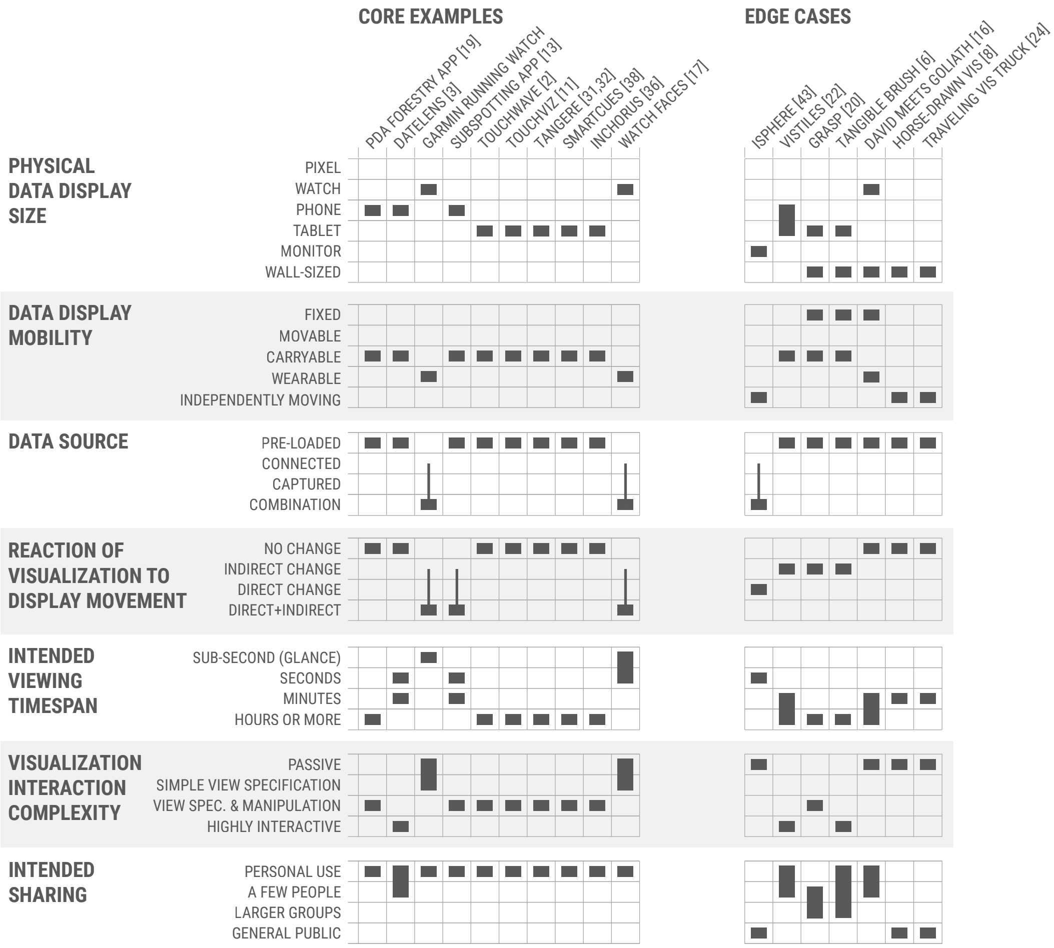

Table 1.1: Overview of core examples and edge cases and their relation to the proposed dimensions of mobile data visualizations.

think it is a mobile visualization? If so, we considered it as part of our investigation. We consciously decided to avoid making solid and arbitrary demarcations of what is and is not mobile data visualization by defining our seven dimensions. The degree of inclusion in the class of mobile data visualizations is then a spectrum with fuzzy borders across these dimensions. It is not about whether something is strictly in or out, but certainly some examples are more clearly mobile data visualizations than others. Our intention is not to find a narrow definition which limits what is considered mobile data visualization, but to present our interpretation and to help suggest in what ways mobile data visualization might grow in the future.

There is still a reductive decision implicit in the choice of only seven dimensions. Other dimensions of description are certainly possible; another dimension we might consider is *utility*. For example, at some point wearable visualizations become more decorative jewelry than useful data displays. An edge case even more exotic than those

Figure 1.16: The flight path visualization in the entertainment system on a passenger airplane. They are mobile with the plane with respect to the planet, but the observer is usually (approximately) static with respect to the visualization.

considered in Section 1.4 is the DNA ring.[1] This ring is a wearable physicalization of one's DNA, but it is likely more a conversation starter than something that really informs about the structure of the underlying data. However, it is still mobile (it moves around) and it is still a representation of the data.

Despite many possibilities associated with mobile data visualization, there are also many challenges, which are reflected in the chapters of this book. For example, as we have seen, display characteristics may vary considerably, and thus data visualization researchers can no longer limit their research to the assumption of a relatively large and flat screen. Chapters 2 and 4 will pick up this aspect by discussing responsive visualization design and 3D mobile data visualization, respectively. Similarly, mobile data visualization calls for new ways of interacting with visual representations of data. As will be explained in Chapter 3, interaction for mobile data visualization can be based on a variety of input modalities available on mobile devices, which is an opportunity and a challenge at the same time. Chapter 8 presents an ideation method that might help in creating new mobile data visualization experiences. So, for those interested in the field of mobile data visualization, the possibilities are wide open and exciting.

ACKNOWLEDGMENTS

We thank Harald Reiterer, who was part of our discussion group on "What is mobile data visualization?" and contributed to the preparation of this chapter. We are also very grateful for the helpful comments and constructive critiques given by Tanja Blascheck, Matthew Brehmer, Sheelagh Carpendale, Eun Kyoung Choe, Raimund Dachselt, Alark Joshi, and John Stasko during the writing of this book chapter.

References

[1] Badam, S. K., Fisher, E., and Elmqvist, N. "Munin: A Peer-To-Peer Middleware for Ubiquitous Analytics and Visualization Spaces". In: *Transactions on*

[1] http://dataphys.org/list/jewellery-shaped-by-dna-profile/

Visualization and Computer Graphics (TVCG) 21.2 (Feb. 2015). **Open Access version:** https://www.researchgate.net/publication/273161089, pp. 215–228. DOI: 10.1109/TVCG.2014.2337337 (cited on page 21).

[2] Baur, D., Lee, B., and Carpendale, S. "TouchWave: Kinetic Multi-Touch Manipulation for Hierarchical Stacked Graphs". In: *Proceedings of the Conference on Interactive Tabletops and Surfaces (ITS)*. **Open Access version:** https://innovis.cpsc.ucalgary.ca/innovis/uploads/Publications/Publications/Baur2012ITS.pdf. New York, NY, USA: ACM, 2012, pp. 255–264. DOI: 10.1145/2396636.2396675 (cited on page 18).

[3] Bederson, B. B., Clamage, A., Czerwinski, M., and Robertson, G. G. "DateLens: A Fisheye Calendar Interface for PDAs". In: *Transactions on Computer-Human Interaction (TOCHI)* 11.1 (Mar. 2004). **Open Access version:** https://www.microsoft.com/en-us/research/wp-content/uploads/2004/03/tochidatelens.pdf, pp. 90–119. DOI: 10.1145/972648.972652 (cited on pages 13, 14).

[4] Besançon, L., Issartel, P., Ammi, M., and Isenberg, T. "Hybrid Tactile/Tangible Interaction for 3D Data Exploration". In: *Transactions on Visualization and Computer Graphics (TVCG)* 23.1 (Jan. 2017). **Open Access version:** https://hal.inria.fr/hal-01372922, pp. 881–890. DOI: 10.1109/TVCG.2016.2599217 (cited on page 21).

[5] Besançon, L., Issartel, P., Ammi, M., and Isenberg, T. "Mouse, Tactile, and Tangible Input for 3D Manipulation". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** https://hal.inria.fr/hal-01436206. New York: ACM, May 2017, pp. 4727–4740. DOI: 10.1145/3025453.3025863 (cited on page 21).

[6] Besançon, L., Sereno, M., Yu, L., Ammi, M., and Isenberg, T. "Hybrid Touch/Tangible Spatial 3D Data Selection". In: *Computer Graphics Forum* 38.3 (June 2019). **Open Access version:** https://hal.inria.fr/hal-02079308, pp. 553–567. DOI: 10.1111/cgf.13710 (cited on page 21).

[7] Blascheck, T., Besançon, L., Bezerianos, A., Lee, B., and Isenberg, P. "Glanceable Visualization: Studies of Data Comparison Performance on Smartwatches". In: *Transactions on Visualization and Computer Graphics (TVCG)* 25.1 (Jan. 2018). **Open Access version:** https://hal.inria.fr/hal-01851306, pp. 630–640. DOI: 10.1109/TVCG.2018.2865142 (cited on page 14).

[8] Brinton, W. C. *Graphic Methods for Presenting Facts*. **Open Access version:** https://archive.org/details/graphicmethodsfo00brinrich. The Engineering Magazine Company, 1914 (cited on page 22).

[9] Brudy, F., Budiman, J. K., Houben, S., and Marquardt, N. "Investigating the Role of an Overview Device in Multi-Device Collaboration". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. CHI '18. **Open Access version:** https://eprints.lancs.ac.uk/id/eprint/89679. Montreal QC, Canada: ACM, 2018, 300:1–300:13. DOI: 10.1145/3173574.

3173874. URL: http://doi.acm.org/10.1145/3173574.3173874 (cited on page 20).

[10] Chittaro, L. "Visualizing Information on Mobile Devices". In: *Computer* 39.3 (Mar. 2006). **Open Access version:** http://hcilab.uniud.it/images/stories/publications/2006-03/VisualizingInformationMobile_IEEECOMPUTER.pdf, pp. 40–45. DOI: 10.1109/MC.2006.109 (cited on page 13).

[11] Chung, H., North, C., Self, J. Z., Chu, S. L., and Quek, F. K. H. "VisPorter: Facilitating Information Sharing for Collaborative Sensemaking on Multiple Displays". In: *Personal and Ubiquitous Computing* 18.5 (2014). **Open Access version:** http://infovis.cs.vt.edu/sites/default/files/visporter_puc.pdf, pp. 1169–1186. DOI: 10.1007/s00779-013-0727-2 (cited on page 21).

[12] Drucker, S. M., Fisher, D., Sadana, R., Herron, J., and schraefel m. c., m. c. "TouchViz: A Case Study Comparing Two Interfaces for Data Analytics on Tablets". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.362.7241. New York: ACM, 2013, pp. 2301–2310. DOI: 10.1145/2470654.2481318 (cited on page 18).

[13] Encarnação, J. L., Frühauf, M., and Kirste, T. "Mobile Visualization: Challenges and Solution Concepts". In: *Proceedings of the Conference on Computer Applications in Production and Engineering (CAPE)*. **Open Access version:** https://www.researchgate.net/publication/2785242. Boston: Springer, 1995, pp. 725–737. DOI: 10.1007/978-0-387-34879-7_75 (cited on page 13).

[14] Goddemeyer, D. and Baur, D. *Subspotting*. Web site. 2016. URL: http://subspotting.nyc/ (cited on page 16).

[15] Goodman, S., Kirchner, S., Guttman, R., Jain, D., Froehlich, J., and Findlater, L. "Evaluating Smartwatch-Based Sound Feedback for Deaf and Hard-Of-Hearing Users Across Contexts". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** https://par.nsf.gov/biblio/10172669. Honolulu, HI, USA: ACM, 2020, 279:1–279:13. DOI: 10.1145/3313831.3376406 (cited on page 14).

[16] Hamilton, P. and Wigdor, D. J. "Conductor: Enabling and Understanding Cross-Device Interaction". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** http://www.dgp.toronto.edu/~dwigdor/?acm=http://dl.acm.org/authorize?N96405. Toronto, Ontario, Canada: ACM, 2014, pp. 2773–2782. DOI: 10.1145/2556288.2557170 (cited on page 20).

[17] Horak, T., Badam, S. K., Elmqvist, N., and Dachselt, R. "When David Meets Goliath: Combining Smartwatches With a Large Vertical Display for Visual Data Exploration". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** http://users.umiacs.umd.edu/~elm/projects/david-goliath/david-goliath.pdf. Montreal

QC, Canada: ACM, 2018, 19:1–19:13. DOI: 10.1145/3173574.3173593. URL: http://doi.acm.org/10.1145/3173574.3173593 (cited on page 21).

[18] Islam, A., Bezerianos, A., Lee, B., Blascheck, T., and Isenberg, P. "Visualizing Information on Watch Faces: A Survey With Smartwatch Users". In: *Short Paper Proceedings of the Conference on Visualization (VIS)*. **Open Access version:** https://hal.inria.fr/hal-03005319. Los Alamitos: IEEE, Oct. 2020. DOI: 10.1109/VIS47514.2020.00038 (cited on page 14).

[19] Jakobsen, M. R., Haile, Y. S., Knudsen, S., and Hornbæk, K. "Information Visualization and Proxemics: Design Opportunities and Empirical Findings". In: *Transactions on Visualization and Computer Graphics (TVCG)* 19.12 (Dec. 2013). **Open Access version:** http://www.kasperhornbaek.dk/papers/InfoViz2013_ProxemicVisualization.pdf, pp. 2386–2395. DOI: 10.1109/TVCG.2013.166 (cited on page 6).

[20] Kirste, T. and Rauschenbach, U. "A Presentation Model for Mobile Information Visualization". In: *Computers & Graphics* 20.5 (Sept. 1996). **Open Access version:** http://www.rauschenbach.net/Publications/docs/cag96.pdf, pp. 669–681. DOI: 10.1016/S0097-8493(96)00041-6 (cited on pages 13, 14).

[21] Kister, U., Klamka, K., Tominski, C., and Dachselt, R. "GraSp: Combining Spatially-Aware Mobile Devices and a Display Wall for Graph Visualization and Interaction". In: *Computer Graphics Forum* 36.3 (June 2017). **Open Access version:** https://mt.inf.tu-dresden.de/cnt/uploads/Kister_GraSp_EuroVis17.pdf, pp. 503–514. DOI: 10.1111/cgf.13206 (cited on page 21).

[22] Langner, R. and Dachselt, R. "Towards Visual Data Exploration at Wall-Sized Displays by Combining Physical Navigation With Spatially-Aware Devices". In: *Poster Proceedings of the Conference on Visualization (VIS)*. **Open Access version:** https://imld.de/cnt/uploads/Langner-2018_PhysNav-SpatialMobiles_InfoVis2018-Poster.pdf. 2018 (cited on page 21).

[23] Langner, R., Horak, T., and Dachselt, R. "VisTiles: Coordinating and Combining Co-Located Mobile Devices for Visual Data Exploration". In: *Transactions on Visualization and Computer Graphics (TVCG)* 24.1 (Jan. 2018). **Open Access version:** https://imld.de/cnt/uploads/Langner_VisTiles_InfoVis17.pdf, pp. 626–636. DOI: 10.1109/TVCG.2017.2744019 (cited on page 20).

[24] Levit, R. *Traveling Datavis Game*. Web site: https://www.ronilevit.com/traveling-dataviz-game. 2018 (cited on pages 12, 22, 23).

[25] Marquardt, N., Hinckley, K., and Greenberg, S. "Cross-Device Interaction via Micro-Mobility and F-Formations". In: *Proceedings of the Conference on User Interface, Software, and Technology (UIST)*. UIST '12. **Open Access version:** http://debaleena.com/courses/cs522/GroupTogether-UIST-2012.pdf. Cambridge, Massachusetts, USA: ACM, 2012, pp. 13–22. DOI: 10.1145/2380116.2380121 (cited on page 20).

[26] Miguel, M. M., Ogawa, T., Kiyokawa, K., and Takemura, H. "A PDA-based See-Through Interface Within an Immersive Environment". In: *Proceedings of the Conference on Artificial Reality and Telexistence (ICAT)*. **Open Access version:** https://www.researchgate.net/publication/4304401. Los Alamitos: IEEE, 2007, pp. 113–118. DOI: 10.1109/ICAT.2007.41 (cited on page 21).

[27] Neshati, A., Sakamoto, Y., Leboe-McGowan, L. C., Leboe-McGowan, J., Serrano, M., and Irani, P. "G-Sparks: Glanceable Sparklines on Smartwatches". In: *Proceedings of the Graphics Interface Conference (GI)*. **Open Access version:** https://doi.org/10.20380/GI2019.23. Kingston, Canada: Canadian Human-Computer Communications Society, 2019. DOI: 10.20380/GI2019.23 (cited on page 14).

[28] Plank, T., Jetter, H.-C., Rädle, R., Klokmose, C. N., Luger, T., and Reiterer, H. "Is Two Enough?! Studying Benefits, Barriers, and Biases of Multi-Tablet Use for Collaborative Visualization". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** https://kops.uni-konstanz.de/handle/123456789/42030. Denver, Colorado, USA: ACM, 2017, pp. 4548–4560. DOI: 10.1145/3025453.3025537 (cited on page 20).

[29] Radloff, A., Tominski, C., Nocke, T., and Schumann, H. "Supporting Presentation and Discussion of Visualization Results in Smart Meeting Rooms". In: *The Visual Computer* 31.9 (2015). **Open Access version:** https://www.researchgate.net/publication/264868629, pp. 1271–1286. DOI: 10.1007/s00371-014-1010-x (cited on page 21).

[30] Razip, A. M. M., Malik, A., Potrawski, M., Maciejewski, R., Jang, Y., Elmqvist, N., and Ebert, D. S. "A Mobile Visual Analytics Approach for Law Enforcement Situation Awareness". In: *Proceedings of the Pacific Visualization Symposium (PacificVis)*. **Open Access version:** https://engineering.purdue.edu/~elm/projects/ivalet/ivalet.pdf. Los Alamitos: IEEE, 2014, pp. 169–176. DOI: 10.1109/PacificVis.2014.54 (cited on page 10).

[31] Roberts, J. C. "State of the Art: Coordinated & Multiple Views in Exploratory Visualization". In: *Proceedings of the Conference on Coordinated and Multiple Views in Exploratory Visualization (CMV)*. **Open Access version:** https://kar.kent.ac.uk/14569/. Los Alamitos, CA, USA: IEEE, 2007, pp. 61–71. DOI: 10.1109/CMV.2007.20 (cited on page 20).

[32] Ros, I. *MobileVis: Examples of Data Visualization Usage on Mobile Devices*. Website. Accessed December 2019. 2014. URL: http://mobilev.is/ (cited on pages 13, 15).

[33] Sadana, R. and Stasko, J. "Designing and Implementing an Interactive Scatterplot Visualization for a Tablet Computer". In: *Proceedings of the Conference on Advanced Visual Interfaces (AVI)*. **Open Access version:** http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.473.6882. New

York, NY, USA: ACM, 2014, pp. 265–272. DOI: 10.1145/2598153.2598163 (cited on pages 17, 18).

[34] Sadana, R. and Stasko, J. "Designing Multiple Coordinated Visualizations for Tablets". In: *Computer Graphics Forum* 35.3 (June 2016). **Open Access version:** https://www.cc.gatech.edu/~stasko/papers/eurovis16-mcv.pdf, pp. 261–270. DOI: 10.1111/cgf.12902 (cited on pages 8, 17, 18).

[35] Sadowski, S. *Mobile Infovis and Dataviz Pattern: Best of Data and Information Visualisations for Mobile Devices*. Web site: https://mobileinfovis.com/. 2018. URL: https://mobileinfovis.com/ (cited on pages 13, 15).

[36] Sereno, M., Besançon, L., and Isenberg, T. "Supporting Volumetric Data Visualization and Analysis by Combining Augmented Reality Visuals With Multi-Touch Input". In: *Poster Proceedings of the European Conference on Visualization (EuroVis)*. **Open Access version:** https://hal.inria.fr/hal-02123904. Eurographics, 2019, pp. 21–23 (cited on page 21).

[37] Sollich, H., von Zadow, U., Pietzsch, T., Tomancak, P., and Dachselt, R. "Exploring Time-Dependent Scientific Data Using Spatially Aware Mobiles and Large Displays". In: *Proceedings of the Conference on Interactive Surfaces and Spaces (ISS)*. **Open Access version:** https://imld.de/cnt/uploads/biovis.pdf. Niagara Falls, Ontario, Canada: ACM, 2016, pp. 349–354. DOI: 10.1145/2992154.2996779 (cited on page 21).

[38] Song, P., Goh, W. B., Fu, C.-W., Meng, Q., and Heng, P.-A. "WYSIWYF: Exploring and Annotating Volume Data With a Tangible Handheld Device". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** https://www.researchgate.net/publication/221518364. New York: ACM, 2011, pp. 1333–1342. DOI: 10.1145/1978942.1979140 (cited on page 21).

[39] Srinivasan, A., Lee, B., Riche, N. H., Drucker, S. M., and Hinckley, K. "InChorus: Designing Consistent Multimodal Interactions for Data Visualization on Tablet Devices". In: *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*. **Open Access version:** https://arxiv.org/abs/2001.06423. New York: ACM, 2020, 653:1–653:13. DOI: 10.1145/3313831.3376782 (cited on pages 17, 18).

[40] Subramonyam, H. and Adar, E. "SmartCues: A Multitouch Query Approach for Details-On-Demand Through Dynamically Computed Overlays". In: *Transactions on Visualization and Computer Graphics (TVCG)* 25.1 (Jan. 2018). **Open Access version:** http://haridecoded.com/images/papers/smartcues.pdf, pp. 597–607. DOI: 10.1109/TVCG.2018.2865231 (cited on page 18).

[41] Vermeulen, J., Luyten, K., Coninx, K., Marquardt, N., and Bird, J. "Proxemic Flow: Dynamic Peripheral Floor Visualizations for Revealing and Mediating Large Surface Interactions". In: *Proceedings of the Conference on Human-Computer Interaction (INTERACT)*. **Open Access version:** https://hal.archives-ouvertes.fr/hal-01610802/. Cham, Switzerland: Springer, 2015, pp. 264–281. DOI: 10.1007/978-3-319-22723-8_22 (cited on page 18).

[42] Wang Baldonado, M. Q., Woodruff, A., and Kuchinsky, A. “Guidelines for Using Multiple Views in Information Visualization”. In: *Proceedings of the Conference on Advanced Visual Interfaces (AVI)*. **Open Access version:** http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.34.7883. Palermo, Italy: ACM, 2000, pp. 110–119. DOI: 10.1145/345513.345271. URL: http://doi.acm.org/10.1145/345513.345271 (cited on page 20).

[43] Windsor Interfaces, Inc. *DateLens Calendar*. Online. 2004. URL: www.windsorinterfaces.com/datelens.shtml (cited on page 14).

[44] Woźniak, P., Lischke, L., Schmidt, B., Zhao, S., and Fjeld, M. “Thaddeus: A Dual Device Interaction Space for Exploring Information Visualisation”. In: *Proceedings of the Nordic Conference on Human-Computer Interaction (NordiCHI)*. Helsinki, Finland: ACM, 2014, pp. 41–50. DOI: 10.1145/2639189.2639237. URL: http://doi.acm.org/10.1145/2639189.2639237 (cited on page 20).

[45] Yamada, W., Yamada, K., Manabe, H., and Ikeda, D. “ISphere: Self-Luminous Spherical Drone Display”. In: *Proceedings of the Conference on User Interface, Software, and Technology (UIST)*. UIST ’17. Québec City, QC, Canada: ACM, 2017, pp. 635–643. DOI: 10.1145/3126594.3126631. URL: https://doi.org/10.1145/3126594.3126631 (cited on pages 18, 19).