Embedded Data Representations

Wesley Willett, Yvonne Jansen, Pierre Dragicevic

► To cite this version:

Wesley Willett, Yvonne Jansen, Pierre Dragicevic. Embedded Data Representations. IEEE Transactions on Visualization and Computer Graphics, 2017, pp.461 - 470. 10.1109/TVCG.2016.2598608 . hal-01377901

HAL Id: hal-01377901

<https://hal.inria.fr/hal-01377901>

Submitted on 29 Nov 2016

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

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

Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike| 4.0 International License

Embedded Data Representations

Wesley Willett, Yvonne Jansen, Pierre Dragicevic

Fig. 1. Examples of embedded data representations¹: (a) dye-based flow visualization on a 1/48 scale airplane model, (b) Yelp's *Monocle* application uses a mobile phone to display business ratings in front of the establishments, (c) concept image of an augmented reality visualization of urban wind flow [42], (d) a visualization of wifi signal strength made with a moving LED rod and long exposure photography, (e) MRI overlay for joint arthrography [13], and (f) a concept image of a drone swarm visualizing crop health [51].

Abstract—We introduce *embedded data representations*, the use of visual and physical representations of data that are deeply integrated with the physical spaces, objects, and entities to which the data refers. Technologies like lightweight wireless displays, mixed reality hardware, and autonomous vehicles are making it increasingly easier to display data in-context. While researchers and artists have already begun to create embedded data representations, the benefits, trade-offs, and even the language necessary to describe and compare these approaches remain unexplored. In this paper, we formalize the notion of *physical data referents* – the real-world entities and spaces to which data corresponds – and examine the relationship between referents and the visual and physical representations of their data. We differentiate *situated representations*, which display data in proximity to data referents, and *embedded representations*, which display data so that it spatially coincides with data referents. Drawing on examples from visualization, ubiquitous computing, and art, we explore the role of spatial indirection, scale, and interaction for embedded representations. We also examine the tradeoffs between non-situated, situated, and embedded data displays, including both visualizations and physicalizations. Based on our observations, we identify a variety of design challenges for embedded data representation, and suggest opportunities for future research and applications.

Index Terms—Information visualization, data physicalization, ambient displays, ubiquitous computing, augmented reality.

1 INTRODUCTION

Data associated with specific locations, people, and objects is becoming increasingly common thanks to the emergence of networked sensing technologies [1]. Yet, despite the fact that much of this data originates in and refers to the physical world, we still tend to explore and analyze it from afar. The vast majority of data analysis and visualization still takes place on desktop computers located far from the objects or locations the data refers to. However, many data-driven tasks – such as debugging a sensor, performing a complex surgery, or assessing the effectiveness of an in-store product display – can be greatly enriched by viewing data relevant to the task in its original context.

Today, a growing array of technologies including lightweight wireless displays, digital fabrication, and autonomous vehicles have begun

to provide new opportunities for integrating data with physical objects and environments. In fact, researchers, engineers, and artists have already started to create a variety of systems that visualize data in the context of its physical referents. However, the benefits, trade-offs, and even the language necessary to describe and compare different approaches remains limited.

In this paper, we present a new conceptual framework that unifies existing research on visualization systems that connect to the physical world. We formalize the notion of *embedded data representations*, the use of visual and physical representations of data that are deeply integrated with the data's *physical referents* – the physical spaces, objects, and entities to which the data refers.

We first discuss the notion of *situated data representation*, which builds on previous work [50] and extends it to include physical (as opposed to screen-based) data representations presented in-context. We then introduce *embedded data representations* – a subcategory of situated data representations that integrate data and physical environments at a deeper level – overlaying or even blending multiple data representations with their physical referents (see examples in Figure 1).

- Wesley Willett is with University of Calgary. E-mail: wj@wjwillett.net.
- Yvonne Jansen is with University of Copenhagen. E-mail: yvja@di.ku.dk.
- Pierre Dragicevic is with Inria. E-mail: pierre.dragicevic@inria.fr.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.

Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx/

¹Photos courtesy: (a) NASA Dryden Flight Research Center photo collection, Photo number ECN-33298-03, 1985; (d) Timo Arnall et al. (yourban.no)

Examples of embedded data representations include *embedded visualizations*, a class of existing systems that use see-through displays, projection, and other augmented reality (AR) technologies to overlay data visualizations with their physical referents. *Embedded physicalizations* go even further by giving data physical form and blending it with physical referents in the real world environment. Although data physicalization is an emerging research topic [25], embedded physicalizations have so far received very little attention. We identify potential tradeoffs associated with these different approaches, and discuss a range of historical, contemporary, and fictional systems that illustrate the broad potential of embedded data representations.

2 DATA VISUALIZATION AND THE PHYSICAL WORLD

Traditional data visualization systems mostly ignore the physical world. Consider the *Dynamic HomeFinder* [52], a classic example of an information visualization system from research. The Dynamic HomeFinder offers interactive scatterplot visualizations to help people find a house to buy. The houses exist physically (in the demo dataset they are all located in the Washington DC area), and in practice, the visualization can be shown on a physical display at almost any location. However, the location and format of the final display is largely considered irrelevant to the design and demonstration of the system.

In fact, the total decoupling of most traditional visualization systems from the real world is generally deliberate. The intent is to design visualization systems that are compatible with a range of concrete datasets and devices, and that can be used in a range of contexts. This view of data visualization has made it possible to visualize almost any dataset anywhere (e.g., through Web apps), and has greatly contributed to bringing data visualization to a large audience. Situated and embedded data representations can be a useful complement to traditional visualization systems. However, designing them requires explicitly accounting for the physical world.

2.1 Physical Referents and Physical Presentations

Figure 2 shows two ways (marked ① and ③) in which a visualization system can relate to the physical world. The top of the diagram shows a *visualization pipeline*, a sequence of operations that turn raw data into a visual representation [4, 6]. The details of the pipeline are left out, only the *raw data* stage has been taken out of the pipeline and shown separately, on the left. Visualization pipelines are computational information processing models: they make up the *logical world* of data visualization systems.

In any concrete instance of a visualization system, this logical world is also linked to the *physical world*, shown at the bottom of the diagram. Visualization systems extend into the physical world in two major ways: through the data's *physical referent* and through the data's *physical presentation*.

The **physical referent** is the physical object or physical space to which the data refers. For example, the referents could be the physical houses described in a real estate dataset or the set of employees described in a company directory. The referents may produce the data themselves (e.g., via sensing instruments or manual entry), but this is not always the case. Often the relationship between the raw data and

its physical referent is purely conceptual. We will discuss the nature of this first relationship (marked ① in Figure 2) in more detail shortly.

The **physical presentation** is the physical instantiation of the visualization produced (©) by the visualization pipeline – i.e., the “*object or apparatus that makes the visualization observable*” [23]. For a traditional visualization system like the Dynamic HomeFinder, this consists of a physical display on which the visualization appears. The physical presentation can also take many other physical forms, such as a mobile device, a large public display, a poster, or a *data physicalization* like a 3D printout or interactive object.

The physical referent and the physical presentation can be related to one another in a variety of different ways in the physical world (see arrow ② in Figure 2). As we will see in the next section, whether or not a data representation can be considered as *situated* entirely depends on the nature of these relationships.

2.2 Composite Systems, Referents and Presentations

Of course, datasets often contain many individual *records* (or *data cases*) that themselves correspond to physical objects or spaces. For example, while a real estate dataset as a whole refers to a set of houses, each individual record in the dataset refers to a single house. As shown in Figure 3, it can often be useful to consider these individual objects and spaces, rather than the set of all of them, as referents. We mark the relationship between records and their individual referents as ①...①.

Standard visualization pipelines often transform each record into a corresponding visual mark. For example, in a scatterplot visualization, we might represent data about each house using a dot. In doing so, these systems collapse information about many referents into a single physical presentation – usually a self-contained visualization displayed on a screen far from the physical referents themselves.

However, visualizing this data using a single physical presentation, while often practical, is not our only option. Just as we can partition a dataset (e.g., all housing data) into a set of individual records that correspond to individual referents (physical houses), we can also partition the physical presentation into smaller sub-presentations. For example, rather than visualizing the record for one house as a dot within a single master visualization, we could instead display it as a separate visualization shown on a separate device. Similarly we could present data from a single house as a stand-alone *physicalization* – a physical object that encodes information about the residence [25]. In Figure 3, we mark the relationship between individual referents and their respective physical presentations as ②...②.

When designing applications, we can present and position individual data elements independently and even distribute them in the physical world. Moreover, we can partition presentations into pieces whose size is appropriate for the current task. This provides greater control over the relationship between each piece of the larger presentation and their respective referents. As we will see, whether or not a data representation can be considered as *embedded* entirely depends on the nature of the relationships ①...①.

Fig. 2. A traditional visualization pipeline extended to the physical world. Raw data and data presentations are both linked to a physical referent.

Fig. 3. A visualization with multiple physical referents and physical presentations that can be composed and interpreted together.

3 SITUATED AND EMBEDDED DATA REPRESENTATIONS

Examples of situated and embedded data representations have been independently explored in a variety of research areas, including information visualization, augmented reality, and ubiquitous computing. To date, however, research has not clearly differentiated between these two approaches, nor has it clearly contrasted them against traditional data visualization. We address these issues by offering a conceptual framework that defines and distinguishes between these distinct classes of data representations.

Fig. 4. Different classes of data representations.

Figure 4 shows the relationships between the classes of data representation we consider. A data representation can be either situated or non-situated, and a situated representation can be either embedded or non-embedded. At the same time, any data representation can be either a regular visualization (to the left), a data physicalization (to the right), or a mix of both (in the center). We start by discussing situated data representations.

3.1 Situated Data Representations

A *situated data representation* is a data representation whose physical presentation is located close to the data’s physical referent(s). Looking back at Figure 2, in a situated data representation the relationship ⑥ between the physical referent and the physical presentation is a relationship of spatial proximity.

Note that the term “situated” has been previously used in various ways, generally in a broader sense. For example in HCI, “situated computing” was used to refer to various concepts such as context-aware computing [20] and contextually-informed system design [3]. Barrett and Irani refer to in-situ data analysis as “*access to situationally appropriate data at an ideal time and place*” [12], while Vande Moere and Hill [35] refer to a situated visualization as “*embedded in a real-world, physical environment*”. Our use of “situated” is more specific but consistent with White [50], who refers to a situated visualization as “*a visualization that is related to its environment*” and “*is based on the relevance of the data to the physical context*”. Also note that we are surrounded by human-made information displays that are situated according to our definition (road signs, price labels, etc.), but we more specifically focus on data representations, which involve the process of encoding and abstracting data into visual variables [36].

3.1.1 Situated Visualizations

A *situated visualization* is a situated data representation for which the presentation is purely visual – and is typically displayed on a screen.

Visualizations displayed on phones, tablets, or laptops can become situated simply by placing the device in a relevant space or near a relevant object (see Figure 5a). Mobile maps with traffic and other data are a very common case, and are often used to access additional data about conditions in the user’s immediate vicinity. Modern automobile instrument panels (Figure 5b) are also examples of permanently situated visualizations, since they display speed, fuel economy, distance, and other information about that specific vehicle.

Researchers working in the areas of ambient and peripheral displays [53, 40] have created and discussed many examples of situated visualizations. For example, Skog’s *Activity Wallpaper* (Figure 5c) visualizes the recent history of acoustic activity in a café [43]. Wouters et al. [54] discuss the design of public situated visualizations that convey

Fig. 5. Examples of situated visualizations: (a) mobile charts of restaurant ratings; (b) an automobile’s instrument panel; and (c) Skog’s Activity Wallpaper [43], all of which show data in the context of their physical referents.

hyperlocal content, that is “*information that is affiliated with a specific geographic area, and with a particular relevance to local community members*”, while VandeMoere et al. [35] discuss the design of urban visualizations that “*reflect on issues that are closely relevant to the social-cultural reality in its vicinity*.”

However, not all public, ambient or peripheral displays are situated. Many ambient systems actually relay remote data, such as the activity of non-collocated collaborators or of significant others [53, 40].

A few situated visualizations use *video see-through* technologies from augmented reality [32]. For example, White [50] discusses a mobile application that can recognize plants and overlays visualizations of related data on the camera image. Similarly, ElSayed [11] proposes a tablet application that can recognize products and overlays a bar chart showing data on the product. Video see-through makes a visualization appear more situated than if the mobile device is used as a separate display of data. We will elaborate on the distinction between perceived and physical situatedness later in the paper.

3.1.2 Situated Physicalizations

Roughly speaking, a data physicalization is a “*physical artifact whose geometry or material properties encode data*” [25], as opposed to traditional visualizations where data is mapped to “*pixels or ink*” [24]. However, there is a large gray area between physicalizations and visualizations, which includes media such as paper that have both virtual and physical qualities, and hybrid representations that combine solid objects with video-projected overlays [25].

Fig. 6. Examples of situated physicalizations: (a) a data-driven jewelry made by meshu.io [22]; (b) mechanically-actuated charts showing local poll results [46]; and (c) a group of living plants showing the usage frequency of a nearby recycling bin [18], all of which show data in the context of its physical referents.

An example of situated data physicalization is meshu.io (Figure 6a), a 3D-printed necklace that encodes a wearer’s past travels [22]. Here the physical referent is the person wearing the necklace. The data physicalization is static, but many situated physicalizations can be updated with new data. For example, someone who tracks her activities using LEGO bricks will update the physicalization manually [23]. In fact, most traditional scientific and measuring instruments are simple situated physicalizations that update in real-time. Another example of dynamically-updated situated physicalization, shown in Figure 6b, is a device that displays and continuously updates results from community polls conducted on the same street [46].

Researchers in related areas such as ubiquitous computing (Ubi-Comp) [48] and ambient/ peripheral displays [53, 40] have created a variety of physical information displays. Many artifacts proposed in these research areas are used to relay remote information, and thus are not situated according to our definition. Even a public or urban display that blends in its physical environment or relates to its environment in semantically meaningful ways [38, 35] is not necessarily a situated data representation. For example, a water fountain where water height maps to the rate of currency exchange [34] is not situated according to our definition. However, examples such as Jeremijenko’s *Dangling String* [48] and Ren et al.’s *Pinwheels* [21], both of which use physical motion to convey the flow of local network traffic, are clear examples of situated data physicalizations.

More recently, researchers have begun to explore the impact of situated data physicalizations for eco-feedback [14] and community engagement [46, 19]. For example Figure 6c shows a situated physicalization that uses computer-controlled lights and real plants placed on top of a recycling bin to display changes in recycling behavior.

3.2 Embedded Data Representations

While the situated representations we have discussed thus far display data via only one physical presentation, *embedded data representations* are made up of multiple physical presentations that each independently display data related to their respective physical referent. Using multiple physical presentations makes it possible to place each presentation closer to its corresponding referent (see Figure 3 (a)–(c)).

Taken individually, each physical presentation encodes data related to its respective referent in close proximity to the referent itself, allowing a viewer to examine both simultaneously. However, the entire embedded data representation can be interpreted globally by considering multiple presentations together. Thus an embedded data representation spatially integrates information more tightly with relevant objects, people, or locations.

3.2.1 Embedded Visualizations

Embedded visualizations use overlays, projection, see-through video, and other virtual presentation techniques to position individual presentations of data close to their corresponding physical referents. Situated visualizations, by contrast, place the entire visualization in a relevant location, but do not necessarily physically align individual data presentations or visual marks with their corresponding referents.

Fig. 7. Examples of embedded visualizations which use projection and video see-through to visualize data about (a) physical inventories [41], (b) faces², and (c) air quality [49] on top of their physical referents.

The area of augmented reality (AR) has demonstrated a number of embedded visualization systems [32, 47]. Figure 7 presents three such examples. Figure 7a shows a prototype application for warehouse management that uses handheld projectors to visualize the contents of boxes in a shelving unit [41]. The visualization is embedded because it displays a physical presentation of each box’s contents directly on the box itself. Similarly, Figure 7b shows a video see-through application that overlays facial expression data on top of each individual person in the scene. Figure 7c shows *SiteLens* [49], a hand-held device that displays air quality measurements in space at the location where they were collected. Figure 1b also shows a related example from Yelp’s *Monocle* application, a mobile phone application that display business ratings in front of the establishments.

²Photo ©Fraunhofer IIS/Kurt Fuchs.

3.2.2 Embedded Physicalizations

Embedded physicalizations achieve a similar effect as embedded visualizations, except that their data presentations consist of physical materials or objects that are associated with the data referents.

Fig. 8. Examples of embedded physicalizations⁴: (a) arrays of physical anemometers; (b) heat-sensitive color changing tiles; and (c) tracking powder on animals, all physically present information using output embedded in the environment.

Scientists, engineers, and artists have long used physical and mechanical methods to observe otherwise-invisible phenomena in-place. Simple examples of these kinds of embedded physicalizations can be composed of discrete physical objects. For example, Figure 8a shows Iñigo Manglano-Ovalle’s *Weather Field No. 1*, an art installation containing a grid of anemometers that allow a viewer to see changes in wind speed and direction in real time across a set of discrete points in space. A related approach used for studying vehicle aerodynamics involves attaching strings or “tufts” to the surface of the vehicle in a grid pattern. When placed in a wind-tunnel, the motion of each string (the physical presentation) communicates the air flow (the data) at a particular location on the vehicle (the referent).

Interestingly, many existing examples of embedded physicalizations use low-level material properties of surfaces or fine-grained particles to make non-visible information discernible. For example, Figure 8b shows a shower with tiles which are coated in a color-changing pigment that visualizes the current temperature at millions of individual points on the wall. Similarly, Figure 8c shows the use of tracking chalk to illustrate the paths and activities of wildlife specimens in exactly the location they occurred. Figure 1a shows a related example, which uses dye to dynamically visualize airflow in the continuous space around a model aircraft.

Presenting information using particles or material properties with extremely small scales makes it possible to very tightly integrate physical presentations and their referents. Because the visual output in these cases is near-continuous, we can conceptually partition the physical presentations and referents into finer resolutions to understand their relationship at an appropriate scale. At one extreme, we could treat each individual grain of pigment in Figure 8b as a separate physical presentation. Each of these tiny presentations is collocated very tightly with its physical referent – the microscopic area of tile it covers. The set of all of these microscopic data presentations, considered together, creates a single highly-embedded wall-sized physicalization.

Individual situated representations at human scales can also serve as building blocks for embedded representations. For example, simple wearable representations such as the data jewelry in Figure 6 or even conference name tags, considered alone, are situated physicalizations. However, a group of people wearing them together (as in Figure 9) become an embedded physicalization. While a single conference name tag presents situated information about just one attendee, the set of tags together allows viewers to compare the roles, affiliations, and seniority of the larger group, while still associating individual pieces of information with each person.

Because computer-driven physical displays are still quite difficult to create, there are currently very few examples of digital embedded

⁴Photos courtesy: (a) Public Art in Public Places Project / K. M. Williamson (publicartinpublicplaces.info); (b) Moving Color (movingcolor.net); (c-left) Carmen Blumlaugh (entm.purdue.edu); (c-right) Lumidust™ (lumidust.co.uk).

Fig. 9. An embedded data representation: a group at a conference where the color of people's badges and the length of ribbons attached to them communicates information on the composition of this group.

physicalizations. However, there are many opportunities to use current and emerging technologies to create them. Embedded shape displays, smart materials, drones, and other forthcoming technologies all provide possible mechanisms for physically rendering data in-place. We discuss a variety of possible future applications of these technologies later in the paper.

4 EXAMPLE SCENARIO

To illustrate how a range of situated and embedded representations can support the same analysis task, we will explore a simple example scenario. Suppose a store manager is interested in optimizing the placement of products in her branch. She has sales data for each product for the past two years, including the locations where each product was displayed in the store.

Non-Situated visualization. By visualizing the data on a computer screen (Figure 10-left), the store manager can easily sort the available data and identify which products sell well and how this varies over time and across different seasons. A screen-based visualization of the sales data can be more or less situated depending on where she uses her computer. For example, if her office is a separate room, she will never be able to see both the sales floor and the visualization at the same time. If she wants to compare the physical placement of various high- or low-selling products, she must do so using abstract descriptions like aisle numbers or overlay the data on a map or 3D simulation of the store [26]. Mapping the data like this can help her to identify locations of interest but still does not allow her to consider the data in the context of the real physical location. For example, if she identifies an area where the sales for all products declined over the last few months, she cannot see whether any unmeasured properties of the physical environment – such as low lighting, a draft, or a smell – could have played a role.

Fig. 10. From left to right: A desktop setting with non-situated visualization. A situated visualization of the same data on a tablet in the store itself. An embedded visualization overlays the data on top of individual products as a heat map. An embedded physicalization displays data by changing properties of the shelves themselves.

Situated visualization. Using a tablet (Figure 10 center-left), she can look at the data for various displays and products from inside the store. The analysis application could even use her current location on the sales floor to filter or highlight data. Doing so allows her to focus on the data in the local context and to compare data from the aisle in which she is currently standing against data from other aisles. However, she still needs to explicitly make the connection between the data points she sees on her tablet and their corresponding referents on the shelves.

Embedded visualization. If the manager wants to see sales information for each product directly alongside the products themselves, she might use an embedded visualization (Figure 10 center-right). Using AR glasses, she could overlay a heatmap of sales information directly onto each of the shelves in the store. Alternatively, if the lights in the store are programmable, she might also use them to display sales or foot traffic data for each aisle. These embedded visualizations could give the manager an even clearer sense of the relationship between the sales data and other environmental factors.

Embedded physicalization. As a final alternative, the store owner could choose to embed the sales data into the store environment as an embedded physicalization (Figure 10-right). For example, she might use integrated displays to change the color of shelves themselves, or use mechanically actuated shelving to physically emphasize over- or under-performing products. These physically-embedded representations, while currently more technically challenging to implement, could make it easier to share and discuss the sales data with other employees, or even use it to direct customers to particular products.

5 DATA-REFERENT RELATIONSHIPS

The relationships between data records and physical referents (Ⓐ in Figures 2 and 3) are conceptual links, and many types of connections are possible. In fact, depending on the task, the same data case could reasonably refer to a number of different referents. For example, store sales data could be connected to a product, but also to a shelf, the customer who made the purchase, or something else entirely. Likewise, a given referent may have many different pieces of associated data. A physical product might be connected to sales data, but also information about its manufacturer, ingredients, etc. To understand the various forms these relationships can take, it can be useful to consider them from both a referent-centric perspective and from a data-centric one.

5.1 Referent-Centric Perspective

When taking a referent-centric perspective, a visualization designer first considers which physical referents might be relevant to a particular task. Next, the designer considers which pieces of data about these referents might be interesting or relevant, and how exploring that data could support the task in-place. For example, if we are interested in enriching a specific *location* such as a city with information relevant to people who seek to buy a house, then we could choose to show variations in air quality, noise levels, cost of living, crime, or the age distribution or education level of residents. Similarly, if we focus on augmenting a particular *object* like a consumer product, then we might want to show the history of how it has moved through time and space, its sales value in different places, or customer ratings. Likewise, a *person* who wants to share data about herself might want to display her social media activity, her current mood and willingness to engage in conversations, or data about her social status or her engagement in a specific community.

Physical referents may need to be considered in combination, as they constantly interact with each other. For example, *locations* are frequented by various *people* every day, while *people* move through the world using different *objects* as means of transportation. They also interact with other *people*, and can create, own, use, change, and destroy *objects*.

5.2 Data-Centric Perspective

Conversely, a designer who takes a data-centric approach may already have some data, and seeks to design a visualization that helps other people analyze it and learn from it. However, to create a situated or embedded representation, the designer needs to consider where the data originates from and how it relates to the physical world – and uncovering these relationships can help design better systems. In a data-centric approach, a key question is: what are the physical referents associated with this data, and what additional insights could those referents provide when interpreting the data?

There is rarely a single, natural physical referent for a data case or dataset but rather a large space of possible referents. At the most basic level, there are the sensors with which the data was collected: for example a thermometer, an accelerometer, an activity tracker, or a camera. Examining these physical referents can provide information (such as their configuration or condition) that can help interpret the data, and having direct access to them can make it possible to collect more data. For manually-recorded observations (e.g., sports statistics, measurements of an animal specimen, or observations of solar transits), the “sensor” is the person who observed and recorded the data, and that person too can provide useful contextual information.

However, sensors are rarely the objects of interest themselves. For example, in most cases, heart rate data is seen as “about” a person, rather than the monitor that collected it. Similarly, a photo is usually associated with the specific subject, place, and time it captures—not the camera that took it. Thus, most physical referents are connected to data not from an information processing standpoint, but through a “semantic relationship” [50]. There can be many sorts of semantic relationships between physical referents and data.

Consider the case of a single data point—the reading on a car’s odometer. There are a number of different possible real-world objects or locations with which to associate this number, including the automobile itself, the roads over which the mileage was accumulated, parts of the automobile (e.g. the engine or the tires), or the driver of the vehicle. The choice of which physical referent(s) the data should be associated with depends on the task at hand and the current analytic or observational goals of the viewers.

Looking back at the scenario from Section 4, store sales are typically recorded at the cash register and data about customer foot traffic may be recorded by cameras and proximity sensors situated throughout the store. However, a manager attempting to use sales data to optimize the placement of products is likely to be concerned with the products themselves, the locations they occupy, or the customers who bought them. Displaying sales data in close proximity to these physical referents can support reasoning about outside factors that might correlate with sales. For example, this juxtaposition could help the manager assess the impact of many external factors like light, air circulation, or foot traffic not captured in the sales data. Alternatively, for a technician attempting to debug the store’s camera system, the data source (the camera itself) may be a more relevant physical referent.

The strength of relationship between data and referent can vary across applications. For example, White [50] describes a system that can recognize plant specimens in the field and overlay visualizations on top of them. Since the data is drawn from a species database, the data and the referent are linked through a class/instance relationship that is weaker than if the data was measured directly from the plant.

5.3 Connecting Data and Referents

Links between data and referents can be stored explicitly, as part of the dataset or as metadata. For example, most digital cameras store images with metadata about the camera and its settings, geolocalized sensors record the place where the data was captured, and datasets about people often store people’s identity. This extra information can be displayed or visualized together with the dataset, in order to help people link data to referents and cross-reference information.

Systematically maintaining links between data and referents can help viewers interpret data in non-situated cases [8]. However, because it is hard to record everything about the physical referents themselves,

viewing them in-place can provide richer and more complete information. Furthermore situated representations can clarify the link between data and referents. For example, Figure 6a shows a data representation worn by the person who generated the data. The representation by itself does not encode to whom the data refers, but viewers can nonetheless infer that information based on physical proximity as well as social norms and conventions.

Nevertheless, explicit linking can help designers or viewers verify that a situated representation is placed next to the correct referent, and can help create and maintain embedded data representations. For example, in Figure 9, the names on people’s badges helps people collect their own badge and prevents them from mixing them up.

6 REFERENT-PRESENTATION RELATIONSHIPS

While designers must consider the conceptual relationships between referents and data records when creating embedded data representations, viewers tacitly experience these connections via the physical relationships between referents and data presentations. These strength

and interpretability of these relationships (Ⓔ in the Figures 2 and 3) depends on the level of *indirection* between each physical data presentation and its referent.

In the simplest case, situated and embedded data representations can minimize indirection by reducing the physical distance between the referent and presentation. However, the notion of distance can be interpreted in many ways. Here we refine the notion of distance by instead considering *levels of indirection*, a concept borrowed from Beaudouin-Lafon’s instrumental interaction framework [2]. We distinguish between *spatial* and *temporal* indirection, and consider cases in which *perceived indirection* may matter more than actual distances.

6.1 Spatial and Temporal Indirection

Spatial indirection refers to the distance between a physical presentation and its corresponding physical referent. For example, displaying data from an air quality sensor on a screen at an entirely different location from where it was collected leads to high spatial indirection. In contrast, if we visualize the same data using a small screen on the sensor itself, the degree of spatial indirection will be very low.

Analogously, *temporal indirection* refers to the temporal distance between the *moment in time* a physical presentation is shown and the original time it refers to. For example, temporal indirection is low when air quality readings are displayed live on the sensing device. However, temporal indirection is higher when displaying historical data or showing future forecasts or predictions.

Traditional visual analysis tools are largely screen-based and usually focus on past data. Thus, most are effectively non-situated, and have little unique relationship to either the current place or the current time. Perhaps the biggest exception are tools for real-time event monitoring and streaming data analysis [37]. Interestingly, however, most traditional scientific instruments (including thermometers, barometers, seismometers, etc.) also have very low spatial and temporal indirection. For these instruments, situatedness is a byproduct of the physical mechanisms used to locally and dynamically turn a non-visual physical quantity into visual form. Over the past century, there has been a widespread transition from these kinds of mechanical instruments to digital sensing tools which record and/or transmit readings rather than displaying them in-place. In fact, for much of the past century, the technologies available for processing and displaying sensed data have often made it difficult to dynamically visualize or examine sensed data with a low spatial or temporal indirection.

Simulated and forecast data represent a special case. Predicted data (such as air quality forecasts) generally map to specific areas in space and specific times in the future, so both spatial and temporal indirection are well-defined. However, viewers may also need to consider the level of uncertainty associated with the prediction. Other simulations (such as air flow simulations for physical models) may have no clear

relationship to the current timeline, and thus no clear degree of temporal indirection. Abstract mathematical simulations and models (such as simulations of dynamic systems in chaos theory) may also have no clear physical referent, and thus no clear degree of spatial indirection.

6.2 Perception and Indirection

Viewers' perception of spatial indirection is related to the real world differences in position between the physical presentation of data and a referent. However, viewers' interpretations can vary based on a number of different factors. For example, the perception of spatial indirection can be biased by the scale of the referent. A distance of 30cm between the presentation and referent may seem very large when examining medical imagery embedded on a patient's body, as in Figure 1e [13]. Yet the same 30cm distance might be barely noticeable on a city-scale visualization like Figure 1c. Temporal indirection varies similarly—a 30-minute lag in measurements might seem very large when monitoring stock market transactions, but is negligible when analyzing geopolitical data.

In some cases the display technology can also lead to systematic disconnects between the *perceived spatial indirection* and the real spatial indirection. For example, a head-mounted display that overlays a visualization on top of the environment can create a subjective experience of low spatial indirection. Presentations of data can appear to be directly on top of their referents, even though the physical display (worn on the viewer's head) is actually far from them [47]. Augmented reality applications like these often try to elicit the appearance of low spatial indirection, but for many setups, this experience is incomplete or imperfect. For example, Yelp's Monocle app (Figure 1b) renders visual presentations of restaurant data in front of the corresponding storefronts. On screen, this creates an embedded visualization that at first appears to have relatively low spatial indirection. However, because viewers remain aware that the smartphone display is far from the actual storefronts, they simultaneously experience both low and high spatial indirection.

Other perceptual cues like depth, shading, and motion can also contribute to the perception of spatial indirection. In the Monocle app, the fact that labels remain aligned with businesses even when the phone is moved helps reinforce the impression of low spatial indirection. However, the app does not adjust the scale or overlap of marks based on their distance from the viewer, making it more difficult to associate presentations with their referents when the physical storefronts are far away or obscured from view. White et al.'s SiteLens (Figure 7c), on the other hand, uses scale, shading, and occlusion between marks to more accurately map them to the appropriate physical space. However, neither tool accurately handles occlusion of virtual marks by real-world objects. Because many factors participate in perceived indirection [47, 22], building perceptually realistic situated or embedded representations is difficult without minimizing the real-world spatial indirection between them.

In contrast to spatial indirection, it is hard for designers of data representations to systematically manipulate *perceived temporal indirection*. Temporal indirection is generally not perceived directly but instead inferred from explanations and higher-level cognitive cues (e.g., labels indicating days or years). However, for practical purposes, temporal indirection can usually be considered low if the state of the physical referent when the data is displayed closely resembles its state when the data was collected. For example, a visualization that reflects a persistent or durable state of a physical referent (such as the geological composition of a mountain) can be seen as having low temporal indirection, even if data was collected long ago. In other cases, designs can use the cyclical nature of time to decrease indirection by displaying data at the same time of day, day of the week, or season in which it was originally collected.

6.3 Visibility and Reachability

The level of spatial indirection between the presentation and referent affects what a viewer can see (*visibility*) and what they can act upon (*reachability*). When the level of spatial indirection is high, a viewer is typically limited to examining either the data representation or the

physical referent. For example, in Section 4, in the non-situated case the store owner cannot simultaneously explore her data visualizations and her store. On the other hand, exploring the two together becomes increasingly easy in the situated and embedded cases when the level of spatial indirection is lower.

Analogously, if the level of spatial indirection is low, viewers can generally reach both the physical data presentation and the physical referents, and physically act on them. However, it is possible for a presentation and referents to be simultaneously visible but not reachable. For example, with AR technology, a viewer may be able to examine astronomical data overlaid on the sky, or to view city population statistics from the window of an airplane. Here the perceived spatial indirection remains low, despite the fact that it is impossible to reach or manipulate the referents directly.

Situated and embedded representations often appear to be the most useful at human-accessible scales, where the physical size and distribution of the referents maximizes visibility and reachability. While larger- and smaller-scale representations may still be useful, big differences in scale can make it increasingly difficult for viewers to assess, explore, and interact with data presentations and/or their referents.

6.4 Facsimiles

A common approach for displaying data in-context involves using facsimiles as stand-ins for the original physical referent(s). A facsimile is generally a scale model that either stands for an actual physical referent (e.g. a solid terrain model of a particular mountain), or is linked through a class/instance relationship to many physical referents. Facsimiles may be useful when a physical referent is too large (like the airplane in Figure 1a), too small (a molecule [16]), too distant (a rover on the surface of Mars), too dangerous (a nuclear reactor) or too fragile (a rare archaeological artifact or a famous painting). There are also many examples of virtual facsimiles in use, such as 3D models in scientific visualization applications, maps backgrounds in geovisualization, and video images in AR video see-through applications such as Figure 1c [42].

In situations like these, presenting data on a facsimile with a different size, location, or material properties can help reduce apparent spatial indirection and increase reachability (at least with respect to the replica). In some cases, a facsimile can be considered a physical referent in and of itself, especially if it replicates the important features of the original at sufficiently high-fidelity. However, using a facsimile tends to reduce a viewer's ability to manipulate or make changes to the original referent, or to see all of its details. However, telepresence and teleoperation techniques [31] may help circumvent these limitations, allowing viewers to either view or manipulate referents that would otherwise be too far away or at too large or small of a scale.

7 ENABLING TECHNOLOGIES & FUTURE SCENARIOS

Current examples of situated and embedded data representations almost exclusively rely on using head-mounted and handheld displays or projectors to overlay information onto existing surfaces. Forthcoming high-fidelity mixed reality displays, such as Microsoft's HoloLens system [5] promise to make it even easier to overlay data on specific objects and features in the surrounding environment. On top of this, a number of promising new technologies have the potential to make embedded data physicalizations more practical.

7.1 Enabling Technologies

Location tracking and spatial mapping techniques are increasingly making it easy to identify and track objects and locations. Technologies like low-power beacons, RFID, and small integrated sensors [29] can already provide precise indoor location information for people, devices, and other objects. Moreover, visual Simultaneous Localization and Mapping (SLAM) systems [9] like Google's Project Tango⁵ are now able to build robust, updatable 3D models of new locations and

⁵get.google.com/tango

objects on-the-fly. Together these technologies will help designers accurately situate presentations close to their referents, even in dynamically changing environments.

Lightweight wireless displays that can be attached to objects, surfaces, and people could be a valuable building block for creating embedded visualizations that do not necessarily require tracking. Current-generation electronic shelf labels designed for store shelves⁶ already provide wireless connectivity and dynamic displays that could be repurposed to visualize data in a variety of other contexts.

Mobile digital fabrication is rapidly decreasing the cost of creating data physicalizations. Mobile 3D printers and other fabrication tools may soon make it possible to easily create new physicalizations in the field. For example, experts examining a set of archeological artifacts might create physicalizations of data for each one - allowing them to compare age, excavation data, and other information about each artifact alongside the physical referent.

Drones and autonomous vehicles have the potential to visually embed data at large scale and in difficult locations. For example, small drones could dynamically position themselves in a physical environment, displaying information like air quality for which the data referent is often an empty space. Swarms of self-organizing drones [28] could also use light, motion, or projection to visualize data at environmental scales (as in Figure 1f).

Smart materials can be controlled to manipulate material properties such as temperature, rigidity, or reflectivity [33, 25]. They have the potential to integrate embedded data representation capabilities directly into physical referents. For example, a clear water bottle could turn opaque if it comes in contact with contaminated water, or smart fabric could become darker with increasing air pollution.

Finally, **high-resolution surface coatings** like those contemplated by Holman et al. [17] should eventually make it even easier to display visualizations on their referents. Coating walls, garments, tools, and other objects with high-resolution displays would allow those items to visualize data both about themselves and other nearby referents. Automultiscopic surface coverings [30] capable of projecting different images in each direction have even greater potential. These displays could produce situated 3D visualizations with fixed spatial relationships to the referent—helping bridge the gap between surface-based visualizations and physicalizations. Alternatively, programmable matter could give a solid form to such 3D visualizations [25].

7.2 Possible Applications

We imagine that embedded data representations can support a variety of specialized and everyday tasks. Some possible applications include:

Search and Rescue. Embedded visualization tools could be particularly useful for tasks like search and rescue that take place in large outdoor environments. For example, in the aftermath of a natural disaster, first responders could benefit from augmented reality overlays highlighting areas that have already been searched or which are particularly dangerous.

Autonomous vehicles could also help expedite searches, and provide visual feedback within the environment. For example searchers might deploy a set of flying drones in a disaster area to identify survivors and assess their status. The drones could then illuminate their immediate surroundings, highlighting the victim's location and using color to visualize their level of criticality. Each individual data presentation (a drone) could guide rescue personnel to a survivor and help them prepare to treat the victim's specific needs. Seen together (for example from the air), the complete set of illuminated drones could also serve as an aggregate visualization, highlighting important areas and providing an overview of the environment.

Fitness and Medicine. Displays embedded in garments or attached to the skin could visualize medical or fitness data in-context. For example, physical therapists, coaches, or athletes could use small adhesive displays to visualize the activation or usage of individual muscle groups.

Similarly, doctors could monitor or visually compare blood flow, temperature, and other information across the body.

Food. Plates, cutlery, or serving ware with embedded displays could visualize nutrition information about a meal in real time. Using sensors on the plate or data about the ingredients or recipes, these displays could highlight possible allergens or give feedback on portion sizes for individual items. They could also visualize information about the provenance of the recipe and ingredients or the process of preparing the meal.

8 BENEFITS, TRADE-OFFS AND CHALLENGES

For some tasks, situated and embedded data representations introduce a number of potential benefits when compared to traditional non-situated and screen-based tools. However, these approaches also have trade-offs that affect where and when we can expect them to be useful. We summarize these trade-offs in the form of comparisons between the different classes illustrated in Figure 4, and discuss challenges with embedded data representations that remain to be addressed.

8.1 Situated vs. Non-Situated

Design flexibility. In practice, most current visualization tools are almost entirely space-indifferent and are designed to run on general-purpose, commodity hardware at any location. As a result, visualization developers can design for the affordances of standard computing devices, without needing to consider the physical referents of the data they represent. However, in cases where contextual information on physical referents is relevant to analyze the data, non-situated systems need to explicitly capture, store and visualize this information.

Data richness. With non-situated representations, analysts are constrained by the information conveyed by the dataset alone, and lack the opportunity to integrate new information gleaned from the physical referent or its surroundings. Sensing systems inevitably record only a small amount of the possible information about a physical referent. Thus, when viewing that data later, either on-screen or in a different location, an analyst is limited to considering the information that has been captured, rather than the surrounding context. By displaying data in the context of the original physical referents, situated representations can allow viewers to examine and extract additional information not present in the dataset itself.

Ambient and Casual Use. Situated visualizations can also act as ambient or peripheral information displays, and allow people to remain aware of contextual data or be notified of anomalies in data while they are engaged in other activities [40, 53].

8.2 Embedded vs. Situated

Compared to situated representations, embedded data representations fundamentally change the way analysts can perceive and interact with data representations.

Physical interactions. Using embedded representations, viewers who navigate the physical environment (either through head motions or locomotion) can simultaneously navigate the data. Analysts can not only change their perspective on the data, but can also alter the physical referents and their surroundings. If representations are driven by live data that is actively being collected and refreshed, viewers can even collect more data by changing the environment. For example, an aerodynamics engineer who visualizes air flow on an articulated or clay model can manipulate the model to observe the results.

⁶e.g., www.m2comm-semi.com/electronic-shelf-labels

Visual attention. While situated representations also show data in context of the physical referents, they may require analysts to make explicit connections between the physical presentation and the physical referents, or to split their visual attention between the two. Examples include a person who uses a map visualization on a mobile phone to locate interesting restaurants in her vicinity, or the scenario in Section 4. With embedded representations, data presentations and their referents can be perceived in unison.

Visibility and Physical Accessibility. In most embedded representations, if a physical referent is not visible or not accessible (e.g., behind the viewer or behind a wall), its data will also be invisible. Thus, situated representations may be preferable to embedded ones when visibility and accessibility of the entire dataset is crucial, but physical scale or environmental conditions would make it difficult to show using an embedded representation.

Encoding flexibility. As with scientific visualizations [36], the spatial representation of an embedded visualization is to a large extent predetermined. Thus a situated representation is preferable if the geometry of the referent is unimportant, or when spatial visual variables could be better used to encode other data attributes (as with the data jewelry in Figure 6a).

8.3 Embedded Visualization vs. Physicalization

Choosing whether to implement embedded data representations using virtual or physical techniques also entails a number of trade-offs.

Implementation Difficulty. In cases where the number of physical referents is small, embedded physicalizations might be more practical than embedded visualizations because they may be more mobile, and may require less infrastructure. Most AR systems require building and maintaining an accurate spatial model of the physical world [32, 15] while autonomous objects may not. However, using virtual embeddings may be more practical in cases where the number of physical presentations is very large, changes considerably over time, or where physical presentations occupy a very large or distant space. For example, using virtual AR overlays to embed information about every house in a city may be much less cumbersome than placing a physical screen, light, or other display at each individual residence.

Visibility and Object Presence. When embedding a representation in a real-world environment, managing occlusions and ensuring the visibility of the visual marks can be a challenge. Overlaid visualizations can show data for physical referents otherwise occluded (e.g., Figure 1b). At the same time, such overlays can visually interfere with or even hide the physical referent. Also, in contemporary AR systems, the level of realism or “object presence” [44, 45] for virtual objects tends to be weak, thus people may not strongly experience the embedded visualizations as being connected to the environment.

Collaboration. In both situated and embedded scenarios, physicalizations may be support collocated collaboration better than virtual tools. In particular, physical marks could provide concrete, persistent instantiations of data that support shared pointing, manipulation, and reasoning among collaborators. In contrast, visualizations rendered using video see-through or head-mounted displays appear different to each collaborator and may be more difficult to coordinate and manipulate as a group. Video projection does not suffer from this problem but has other limitations. For example, projected content can be easily occluded, projection is difficult on shiny and dark surfaces, and it is hard to project content outside the surfaces themselves [15].

Privacy. In cases where the embedded data is of a sensitive or personal nature, virtual embeddings may be preferable, since they offer more opportunities to hide, obscure, or anonymize sensitive information.

8.4 Hybrids and Transitioning Between Representations

Non-situated, situated, and embedded data representations all have benefits and drawbacks, with trade-offs depending on a variety of factors such as the type of data, referents, and tasks. In many cases, situating or embedding data representations may not make sense. For example, in some situations, data may be abstract with no natural physical referents, or there may not be obvious physical referents whose scale, location, number, or accessibility make them useful for a given task.

For other data, situating or embedding representations may be useful in some circumstances but not in others. In these situations, it appears useful to consider the feasibility of hybrid solutions that would let people transition between non-situated, situated, and embedded modes. A non-situated representation can temporarily become situated if brought in proximity with its physical referents – for example, a person who wants to buy a camera could visualize data on all recent cameras on her tablet device, and bring the device to the store. However, such a temporarily situated representation would be agnostic to its current situated status and, for example, not be able to take into account that half of the cameras shown are not available in that store. The increasing use of location-based services on mobile devices (e.g., Figure 1b) suggests that many non-situated visualization systems could be repurposed as situated visualizations simply by adding filtering or search functionality based on a viewer’s current location.

The step from situated to embedded representations is more profound, as it requires shifting from a single, independent presentation (e.g., on a mobile device) to a collection of presentations integrated with the physical environment. This step requires rethinking not only how users experience and interact with the representation, but also the visualization pipeline for creating the representation (cf. Figures 2 and 3). While transitioning from a situated to an embedded mode is theoretically feasible, it depends heavily on the availability of systems that can accurately render and control multiple visual or physical presentations in arbitrary environments. Emerging technologies, like those mentioned in Section 7.1, will likely facilitate this in the near future.

8.5 Open Research Questions

Our conceptual framework is descriptive rather than prescriptive, and does not consider how to best design situated and embedded data representations. More work is needed to understand best practices and to combine the lessons learned from HCI and information visualization with principles from design and architecture [34, 38, 39, 7]. We also currently have little empirical data on the benefits of situated and embedded data representations, with the exception of a few isolated evaluations of situated AR systems [10, 49] and observational studies on urban visualizations [27, 7, 46]. These studies suggest that situated and embedded data representations can be beneficial, but many open questions remain about how situated and embedded visualizations affect viewers’ perception of data. When can viewers gain richer insights by analyzing data in the context of their physical referents? Do embedded visualizations lead to perceptual distortions? If so, do these distortions differ when using visualizations or physicalizations? How can we make sure that situated and embedded data representations in public spaces are recognizable and readable [38]?

9 CONCLUSION

Emerging technologies like wireless displays, accessible augmented reality hardware, and drones are rapidly creating new opportunities to surface data in the physical world. Although a variety of art pieces, applications, and visualization systems have begun to explore deeper integrations of data with the physical world, the language for characterizing and comparing these approaches remains nascent. As a first step, our work (i) formalizes the distinction between non-situated, situated, and embedded data representations and (ii) provides an initial discussion of the relationships between data, physical referents, and data presentations that define these systems.

For now, designing embedded data representations, particularly physical ones, is both technically and conceptually challenging—and only a few examples exist. However, tools for presenting and exploring data in-context in our everyday environment have the power to change how we approach almost any task, personal or professional. By outlining the challenges, trade-offs, and potential of these approaches, we hope to set the stage for a diverse range of embedded visualization and physicalization systems to come.

ACKNOWLEDGMENTS

We thank Lora Oehlberg, Jean-Daniel Fekete, Frédéric Vernier, and Anthi Dimara for thoughtful discussion, feedback, and proofreading.

REFERENCES

- [1] L. Atzori, A. Iera, and G. Morabito. The internet of things: A survey. *Computer networks*, 54(15):2787–2805, 2010.
- [2] M. Beaudouin-Lafon. Instrumental interaction: an interaction model for designing post-wimp user interfaces. In *Proc. CHI'00*, pages 446–453. ACM, 2000.
- [3] M. Beaudouin-Lafon and W. E. Mackay. Research directions in situated computing. In *CHI-EA'00*, pages 369–369. ACM, 2000.
- [4] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.
- [5] H. Chen, A. S. Lee, M. Swift, and J. C. Tang. 3d collaboration method over hololens and skype end points. In *Proc. of the 3rd International Workshop on Immersive Media Experiences*, pages 27–30. ACM, 2015.
- [6] E. Chi and J. T. Riedl. An operator interaction framework for visualization systems. In *Proc. of the IEEE Symposium on Information Visualization*, pages 63–70. IEEE, 1998.
- [7] S. Claes and A. Vande Moere. Street infographics: Raising awareness of local issues through a situated urban visualization. In *Proc. of PerDis'13*, pages 133–138. ACM, 2013.
- [8] S. B. Davidson and J. Freire. Provenance and scientific workflows: challenges and opportunities. In *Proc. of SIGMOD'08*, pages 1345–1350. ACM, 2008.
- [9] A. J. Davison and D. W. Murray. Simultaneous localization and map-building using active vision. *TPAMI*, 24(7):865–880, 2002.
- [10] N. ElSayed, B. Thomas, K. Marriott, J. Piantadosi, and R. Smith. Situated analytics. In *2015 Big Data Visual Analytics (BDVA)*, pages 1–8. IEEE.
- [11] N. A. M. ElSayed, B. H. Thomas, R. T. Smith, K. Marriott, and J. Piantadosi. Using augmented reality to support situated analytics. In *2015 IEEE Virtual Reality (VR)*, pages 175–176. IEEE.
- [12] B. Ens and I. Pourang. Spatial analytic interfaces: Spatial user interfaces for In-Situ visual analytics. *IEEE CG&A*, in press, 2016.
- [13] G. Fischer, A. Deguet, C. Csoma, R. Taylor, L. Fayad, J. Carrino, J. Zinreich, and G. Fichtinger. Mri image overlay: application to arthrography needle insertion. *Computer Aided Surgery*, 12(1):2–14, 2007.
- [14] J. Froehlich, L. Findlater, and J. Landay. The design of eco-feedback technology. In *Proc. of CHI'10*, pages 1999–2008. ACM, 2010.
- [15] R. Gervais. *Interaction and introspection with tangible augmented objects*. Theses, Université de Bordeaux, Dec. 2015.
- [16] A. Gillet, M. Sanner, D. Stoffler, and A. Olson. Tangible interfaces for structural molecular biology. *Structure*, 13(3):483–491, 2005.
- [17] D. Holman and R. Vertegaal. Organic user interfaces: Designing computers in any way, shape, or form. *CACM*, 51(6):48–55, June 2008.
- [18] D. Holstius, J. Kembel, A. Hurst, P.-H. Wan, and J. Forlizzi. Infotropism: living and robotic plants as interactive displays. In *Proc. of DIS'04*, pages 215–221. ACM, 2004.
- [19] S. Houben, C. Golsteijn, S. Gallacher, R. Johnson, S. Bakker, N. Marquardt, L. Capra, and Y. Rogers. Physikit: Data engagement through physical ambient visualizations in the home. In *Proc. of CHI'16*, pages 1608–1619. ACM, 2016.
- [20] R. Hull, P. Neaves, and J. Bedford-Roberts. Towards situated computing. In *First International Symposium on Wearable Computers, 1997. Digest of Papers.*, pages 146–153, Oct. 1997.
- [21] H. Ishii, S. Ren, and P. Frei. Pinwheels: visualizing information flow in an architectural space. In *CHI-EA'01*, pages 111–112. ACM, 2001.
- [22] Y. Jansen. *Physical and Tangible Information Visualization*. PhD thesis, Université Paris Sud-Paris XI, 2014.
- [23] Y. Jansen and P. Dragicevic. An interaction model for visualizations beyond the desktop. *TVCG*, 19(12):2396–2405, 2013.
- [24] Y. Jansen, P. Dragicevic, and J.-D. Fekete. Evaluating the efficiency of physical visualizations. In *Proc. of CHI'13*, pages 2593–2602. ACM, 2013.
- [25] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and challenges for data physicalization. In *Proc. of CHI'15*, pages 3227–3236. ACM, 2015.
- [26] G. Kahl and C. Burckert. Architecture to enable dual reality for smart environments. In *International Conference on Intelligent Environments*, pages 42–49, June 2012.
- [27] L. Koeman, V. Kalnikaitė, and Y. Rogers. Everyone is talking about it!: A distributed approach to urban voting technology and visualisations. In *Proc. of CHI'15*, pages 3127–3136, 2015.
- [28] A. Kushleyev, D. Mellinger, C. Powers, and V. Kumar. Towards a swarm of agile micro quadrotors. *Autonomous Robots*, 35(4):287–300, 2013.
- [29] M. Le Goc, P. Dragicevic, S. Huron, J. Boy, and J.-D. Fekete. Smarttokens: Embedding motion and grip sensing in small tangible objects. In *Proc. of UIST'15*, pages 357–362. ACM, 2015.
- [30] B. Masia, G. Wetzstein, P. Didyk, and D. Gutierrez. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics*, 37(8):1012–1038, 2013.
- [31] P. Milgram and H. Colquhoun. A taxonomy of real and virtual world display integration. *Mixed reality: Merging real and virtual worlds*, 1:1–26, 1999.
- [32] P. Milgram, H. Takemura, A. Utsumi, and F. Kishino. Augmented reality: a class of displays on the reality-virtuality continuum. In *Photonics for Industrial Applications*, pages 282–292, 1995.
- [33] A. Minuto, D. Vyas, W. Poelman, and A. Nijholt. Smart material interfaces: A vision. In *Intelligent technologies for interactive entertainment*, pages 57–62. Springer, 2011.
- [34] A. V. Moere. Beyond the tyranny of the pixel: Exploring the physicality of information visualization. In *Proc. of IV'08*, pages 469–474, 2008.
- [35] A. V. Moere and D. Hill. Designing for the situated and public visualization of urban data. *Journal of Urban Technology*, 19(2):25–46, 2012.
- [36] T. Munzner. *Visualization Analysis and Design*. CRC Press, 2014.
- [37] A. A. Norton, M. A. Rubin, and L. Wilkinson. Streaming graphics. *Statistical Computing and Graphics Newsletter*, 12(1):11–14, 2001.
- [38] D. Offenhuber. The invisible Display-Design strategies for ambient media in the urban context. In *International Workshop on Ambient Information Systems, Colocated with Ubicomp*, 2008.
- [39] D. Offenhuber and S. Seiting. Over the rainbow: Information design for low-resolution urban displays. In *Proc. of MAB'14*, pages 40–47, New York, NY, USA, 2014. ACM.
- [40] Z. Pousman and J. Stasko. A taxonomy of ambient information systems: four patterns of design. In *Proc. of AVI'06*, pages 67–74. ACM, 2006.
- [41] R. Raskar, P. Beardsley, J. Van Baar, Y. Wang, P. Dietz, J. Lee, D. Leigh, and T. Willwacher. Rfig lamps: interacting with a self-describing world via photosensing wireless tags and projectors. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 406–415. ACM, 2004.
- [42] S. Ritterbusch, S. Ronnäs, I. Waltschläger, P. Gerstner, and V. Heuveline. Augmented reality visualization of numerical simulations in urban environments. *International Journal of Advances in Systems and Measurements*, 6(1):26–39, 2013.
- [43] T. Skog. Activity wallpaper: ambient visualization of activity information. In *Proc. of DIS'04*, pages 325–328. ACM, 2004.
- [44] B. Stevens and J. Jerrams-Smith. The sense of object-presence with projection-augmented models. In *Haptic Human-Computer Interaction*, LNCS, pages 194–198. Springer, 2001.
- [45] B. Stevens, J. Jerrams-Smith, D. Heathcote, and D. Callear. Putting the virtual into reality: Assessing Object-Presence with Projection-Augmented models. *Presence: Teleoperators and Virtual Environments*, 11(1):79–92, 2002.
- [46] A. Taylor, T. Regan, D. Sweeney, V. Vlachokyriakos, L. Grainger, J. Lingel, and S. Lindley. Data-in-place: Thinking through the relations between data and community. In *Proc. of CHI'15*. ACM, April 2015.
- [47] D. Van Krevelen and R. Poelman. A survey of augmented reality technologies, applications and limitations. *International Journal of Virtual Reality*, 9(2):1, 2010.
- [48] M. Weiser and J. S. Brown. Designing calm technology. *PowerGrid Journal*, 1(1):75–85, 1996.
- [49] S. White and S. Feiner. Sitelens: situated visualization techniques for urban site visits. In *Proc. of CHI'09*, pages 1117–1120. ACM, 2009.
- [50] S. M. White. *Interaction and presentation techniques for situated visualization*. Columbia University, 2009.
- [51] W. Willett. Cetonia - a dynamic swarm at your fingertips. IEEE VIS Death of the Desktop Workshop. <http://www.wjwillett.net/content/cetonia/>, November 2014. [Online].
- [52] C. Williamson and B. Shneiderman. The dynamic homefinder: Evaluating dynamic queries in a real-estate information exploration system. In *Proc. of SIGIR'92*, pages 338–346. ACM, 1992.
- [53] C. Wisneski, H. Ishii, A. Dahley, M. Gorbet, S. Brave, B. Ullmer, and P. Yarin. Ambient displays: Turning architectural space into an interface between people and digital information. In *Cooperative Buildings: Integrating Information, Organization, and Architecture*, LNCS, pages 22–32. Springer, 1998.
- [54] N. Wouters, S. Claes, and A. V. Moere. Investigating the role of situated public displays and hyperlocal content on place-making. *Interaction Design and Architecture (s)*, (25):60–72, 2015.