

Situated Analytics

Bruce H. Thomas¹, Gregory F. Welch², Pierre Dragicevic³, Niklas Elmqvist⁴,
Pourang Irani⁵, Yvonne Jansen⁶, Dieter Schmalstieg⁷, Aurélien Tabard⁸, Neven
A. M. ElSayed¹, Ross T. Smith¹, and Wesley Willett⁹

¹ University of South Australia
[bruce.thomas,neven.elsayed,ross.smith]@unisa.edu.au

² University of Central Florida
multi-layers and con welch@ucf.edu

³ Inria, France
pierre.dragicevic@inria.fr

⁴ University of Maryland, College Park
elm@umd.edu

⁵ University of Manitoba
irani@cs.umanitoba.ca

⁶ Sorbonne Universités, UPMC Univ Paris 06, CNRS, ISIR
jansen@isir.upmc.fr

⁷ Graz University of Technology
schmalstieg@tugraz.at

⁸ Univ Lyon, Université Lyon 1, CNRS UMR5205
atabard@liris.cnrs.fr

⁹ University of Calgary
wesley.willett@ucalgary.ca

Abstract. This chapter introduces the concept of *situated analytics* that employs data representations organized in relation to germane objects, places, and persons for the purpose of understanding, sensemaking, and decision-making. The components of situated analytics are characterized in greater detail, including the users, tasks, data, representations, interactions, and analytical processes involved. Several case studies of projects and products are presented that exemplify situated analytics in action. Based on these case studies, a set of derived design considerations for building situated analytics applications are presented. Finally, there is an outline of a research agenda of challenges and research questions to explore in the future.

1 Introduction

People are increasingly interested in understanding data directly associated with objects, locations, or persons in their everyday life. For example, imagine hunting for a house by walking through the neighborhood in which you want to live. Your search could be informed by social media posts about the area appearing in virtual signs on the ground, dynamic census data rising above the houses, and historical traffic data rendered on the street in a way that reflects the ebb and flow

throughout a typical day. In an industrial context, imagine a team of professionals collaboratively reorganizing machines and stations on a factory floor. As they walk around together, dynamic data would appear embedded in and around the floor. To demonstrate this concept, the data could be related to physical space, safety, economics (e.g., layout constraints related to past manufacturing data), power, ventilation, ergonomics and worker preferences, past accidents, and legal concerns.

In these examples, the individuals would be interested in discovering, interpreting, and communicating meaningful patterns in data that is directly relevant to, and integrated into, the physical space all around them. For the first time both the technology—such as sensors, wearable displays, natural user interfaces, and augmented reality devices—as well as the data sources—such as dynamically updating social media, ubiquitous sensor information, and large-scale movement data—exist to make this vision a reality. To cope with this massive quantity of data, analytics techniques are required to help the user, much in the same way visual analytics grew out of the visualization research domain. Visual Analytics (VA) has been defined as “*the science of analytical reasoning facilitated by visual interactive interfaces*” [124].

Analogously, the concept of *Situated Analytics* [18] (SA) is the use of data representations organized in relation to relevant objects, places, and persons in the physical world for the purpose of understanding, sensemaking, and decision-making. For example, virtual labels on a physical container provide semantic information for the analyst, as does the proximity of different objects. Situated analytics allows users to access the power of the cloud (data and analysis) seamlessly analyzing virtual data situated in the physical world simultaneously.

Imagine walking into a pharmacy and placing an order for your prescriptions. While you are waiting for your order, you browse the shop for other purchases. Through the use of Internet of Things technology [2] and situated analytics, you will be reminded of items you might wish to replenish. This is performed by highlighting items in your field of view using augmented reality to draw your attention. You notice a natural supplement you have not tried that looks interesting. Situated analytics can provide a number of supporting functions for you. First, it can recognize the supplement and determine if the ingredients are compatible with your prescription and if you are allergic to the item. Second, you can compare different brands of the supplement by placing them next to one another. The system will automatically perform an analysis and compare the products. Third, a visualization summarizing the reviews, keywords, and ratings can be embedded on the store shelf right next to the supplement. Finally, the user can inspect more detailed information about the product by selecting portions of the label to bring up Augmented Reality (AR) [55] information in greater detail.

Situated analytics and immersive analytics are complementary techniques that have emerged at similar times. Situated analytics draws from the domains of Visual Analytics and Augmented Reality to support a new form of in-situ interactive visual analysis. First published as a novel interaction, and a visual-

ization concept for reasoning support which incorporates four primary elements: situated information, abstract information, augmented reality interaction, and analytical interaction [19]. Where immersive analytics [7] covers a broad range of display techniques and technologies [33, 51, 55].

The degree of immersion you might associate with an instance of situated analytics depends on your definition of the word *immersion*. If one considers immersion to reflect deep mental involvement, it is unclear that situated analytics would be any different (from an immersion standpoint) from non-situated analytics. In other words, the situated nature of the analytical task would not necessarily bring any new immersion-related considerations or affects to the task. If one considers immersion to reflect the degree to which the analyst is surrounded by an engrossing total environment—including the real environment, virtual data, and analytics—then situated analytics could be considered inherently immersive, as the environment where the virtual data and analytics are situated surrounds the analyst. If instead one considers immersion to be associated solely with the virtual data and analytics (not including the real environment), then one's sense of the degree of immersion might depend on the degree to which the virtual data and analytics surround the analyst, i.e., the degree of immersion is proportional to the spatial extents of the virtual data and analytics.

In 1997, Slater and Wilbur defined the concept of *immersion* as being more generally related to the characteristics of a *system*, and in particular to the *sensorimotor contingencies* that the system supports [56]. Sensorimotor contingencies refer to the available actions humans employ to perceive things using vision and other sensory modalities, for example moving one's body, head, or eyes to obtain a better visual or aural perspective [48, 49]. This notion of immersion has been widely used in the Virtual Reality and Augmented Reality communities. If one adopts this perspective, then the degree of immersion is related to the sensorimotor contingencies afforded by the system presenting the virtual data and analytics to the user. If, for example, a person is using an AR system with a relatively narrow visual field of view, the user might not be able to see the complete virtual data and analytics without scanning their view left-right or up-down. From a Slater and Wilbur perspective, such a system would be less immersive than a system that presented the virtual data and analytics completely (in the same resolution) without the need for explicit movement on the part of the user. This same concept would apply beyond vision to other senses including sound, smell, and touch (haptic senses). In this view, the degree of immersion is essentially proportional to the sensorimotor contingencies.

1.1 Comparison to Other Fields and Concepts

Situated analytics has a close relationship with a number of fields. Table 1 compares a number of fields from low to high of their *Situatedness* versus their *Analytic Level*. Where *situatedness* is the degree the information and person are connected to the task, location, and/or another person, and the *analytic level* is the quantity of analytic processing of the information. You will notice situated analytics requires high levels of both situatedness and analytics.

Situatedness	Analytic Level Low	Analytic Level High
High	Situation Awareness	Situated Analytics
Low	Information Displays Ambient Displays	Visual Analytics Traditional Analytics

Table 1: Situatedness versus Analytic Level

Below is a set of brief overviews of some of the related fields and concepts:

1. **Augmented/Mixed reality** [9, 55] is a dominant form of presentation of information in situated analytics, and interaction techniques from these fields can inform new interaction techniques for situated analytics.
2. **Wearable and mobile computing** [57] support the user operating in unprepared physical locations and leverages such technologies as interaction, device form factor, and display techniques. This mobile nature is required by many situated analytics applications.
3. **Situated computing** [26] investigates computational devices that detect, interpret and respond to aspects of the user's local environment.
4. **Situated visualization** [51, 64] refers to data representations that are related to and portrayed in their physical environment. Sensemaking is achieved through the combination of the visualization and the relationship of that visualization to the immediate physical environment.
5. **Embedded data representations** [51] focus on the use of visual and physical representations of data that are deeply integrated in the physical spaces, objects, and entities the data refers to. This closer association with physical objects and virtual information is critical for situated analytics.
6. **Contextual computing** overlaps with situated analytics; Chen and Kotz [15] defines context as a set of environmental states and settings that 1) determines an application's behavior or 2) when an application event arises then is of interest to the user.
7. **Ambient displays** [66] employ the user's complete physical environment as an interface to their virtual information space. Situated analytics endeavors to bring the users analytic information space to be in-situ to physical objects and spaces of interest.
8. **Ubiquitous computing**, according to Mark Weiser's [63] vision, consists of embedding numerous computers in the user's physical environment, and as such make the computation device fade into the user's background. The goal of SA is to bring visual analytics to bear on problems away from the user's workstation and into the physical world. This can be performed through either ubiquitous computing or wearable computer or even mobile computing.
9. **Visualization beyond the desktop** [20, 39, 54] is a broad research agenda of which situated analytics and situated visualization is a component. Instead of the ubiquitous mouse and keyboard, such visualization systems focus on touch-based, pen-based, or gestural interaction methods with multiple form factors (smartphones, smartwatches, tablets) and both large and small displays (HMDs, powerwalls, wearable displays, etc.).

10. **Ubiquitous analytics** [17] is “... *amplif[ying] human cognition by embedding the analytical process into the physical environment to enable sensemaking of big data anywhere and anytime.*” Several tools and toolkits have been proposed to support this practice [3, 4]. Situated analytics builds on this concept by focusing on place as an index.
11. **Personal visual analytics** (PVA) is concerned with assisting visual analytics with a personal context [25]. The aim of personal visual analytics is to support people with the ability to acquire an awareness of, explore, and learn from data around them and from their personal context. This is a similar goal for situated analytics, yet again without the immersive aspect.

1.2 Visual Analytics and Augmented Reality

Situated Analytics leverages two research domains—VA and AR—to deliver analytical reasoning in the world around the user [21]. VA is a multidisciplinary research domain spanning analytical reasoning techniques with visualization, while AR enhances the physical worldview with a visual overlay of registered contextual information in real-time. Most of situated analytics combines VA techniques with AR techniques for in-situ registration of information onto the physical space.

AR has been shown to be a useful tool for visualization [58]. Kalkofen et al. [34] considered three types of AR visualization techniques: 1) data integration, 2) scene manipulation and 3) context-driven. Data integration techniques enhance the smooth mixing of the virtual information with the physical world [23, 33]. Scene manipulation techniques manipulate the real scene to augment information. A few examples of this smooth mixing are as follows: the relocation of physical objects [35], color corrections [23], and diminished reality techniques [40]. To incorporate some of the user's influence on the visual presentation of information, context-driven techniques have been employed [64].

A major research challenge remains with the limitations of current AR display technologies [37]. AR display technologies and techniques primarily fall into two categories: Visual Augmented Reality (VAR) and Spatial Augmented Reality (SAR). In the case of VAR the virtual content is overlaid into the user's visual field, for example via a head-worn (head-mounted) display (HMD) device, a handheld device such as a mobile phone, and—some day—special AR contact lenses [45], see the example in Figure 1. This is the most common form of AR, and thus often the implied form when someone refers to AR.

In the case of SAR, virtual content is displayed directly on objects in the user's physical space [11, 51, 52]; typically using digital projectors and a mapping technique initially known as *shader lamps* [50, 53]. More recently this concept is referred to as *projection mapping*, see examples in Figure 2. When the virtual content is not associated directly with any physical object, e.g., as would be the case with air flow visualization information or floating virtual labels attached to physical objects, then VAR would be more appropriate. When the virtual content is to appear (or is intended to appear) directly on the surface of a physical object, SAR becomes an option. Some advantages of SAR in such situations




Fig. 1: An example of Visual Augmented Reality (VAR), where the virtual imagery is overlaid onto the user’s visual field. The left image (a) shows a Microsoft HoloLens device. The right image (b) shows a person wearing a Microsoft HoloLens device, and a depiction of what they would see—Minecraft [44] game objects visually overlaid onto the real world table, couch, etc.

include the natural coincidence of vergence and accommodation of the human visual system (a problem with VAR), and not requiring devices to be worn or held—a particular advantage in group settings, where one would like to see other individuals naturally as they discuss the issues at hand.

1.3 Motivation

Why do we need situated analytics? This method of sensemaking has a great potential to have a major impact on people’s use and application of Big Data in their everyday lives. The significance of this method of sensemaking is a new research domain that provides the intersection of many research concepts. Situated analytics can be beneficial for data exploration and information comprehension. There are three ways situated analytics can enhance sensemaking:

1. more understandable information presentation by immediately associating information with the germane physical objects (i.e., place acting as a *spatial index*),
2. more natural method for information exploration interactions by allowing the user to touch and manipulate the germane physical objects (i.e., the use of *natural interaction*), and
3. more comprehensive information analysis providing contextual and overview information (i.e., *contextual synthesis* of data).

Critical to the success of situated analytics for the appropriate application of this new technique for casual and expert users with real world tasks in actual physical settings are techniques that enhance the user’s ability and increase their effectiveness. There are limitations to the current technologies (displays [42] and computer vision [9] for example). Barring these limitations, people increasingly want to base their decisions on data at the location of the decision, such as a purchase.




Fig. 2: Examples of Spatial Augmented Reality (SAR) [11, 51, 52]. The upper pair of images show a one meter square physical model of the Taj Mahal (a) before and (b) after augmentation via the *shader lamps* technique [50, 53]. The lower pair of images similarly show a vase (c) before and (d) after augmentation.

1.4 Structure of the Chapter

This chapter starts by characterizing Situated Analytics in greater detail, including the users, tasks, data, representations, interactions, and analytical processes involved. A set of case studies of projects and products are examined that exemplify best practice situated analytics in action. Extending the interaction technologies of SA are the presented blended situated analytics controls with a set of example applications. Based on these case studies, a set of derived design considerations for building situated analytics applications are presented. Finally, there is an outline of a research agenda of challenges and research questions to explore in the future.

2 Characterization of Situated Visualizations

While *visual analytics* aims at supporting analytic reasoning through the use of visualisations, *situated analytics* aims at supporting analytic reasoning through the use of *situated* visualizations. This section introduces a conceptual framework and a terminology that help characterize and reason about situated visualizations, while temporarily leaving out the analytic aspects. This framework is based largely on Willett et al. [51]¹⁰. The section first starts by explaining what it means for a

¹⁰ While the basic model is the same (Figure 3), the text and definitions have been fully reworked, and an illustration has been added to clarify the notion of embedded visualization (Figure 4). The interaction model (Figure 5) is new.

data visualization to be spatially situated. It then discusses physically-situated vs. perceptually-situated visualizations, embedded visualizations, and temporally-situated visualizations. Finally, a model of interaction is presented.

2.1 Spatially-Situated Visualizations

Since situated analytics involves data visualizations that are integrated in the physical environment, a model of data visualization is required that accounts for the existence of the physical world. The most widely used model of data visualization, i.e., the information visualization reference model (or “visualization pipeline”) [12,16], essentially ignores the physical world. A conceptual model is described that unifies two recently introduced models that capture the physical world around visualizations: the embedded data representation model from Willett et al. [51], and the beyond-desktop pipeline model from Jansen et al. [20].

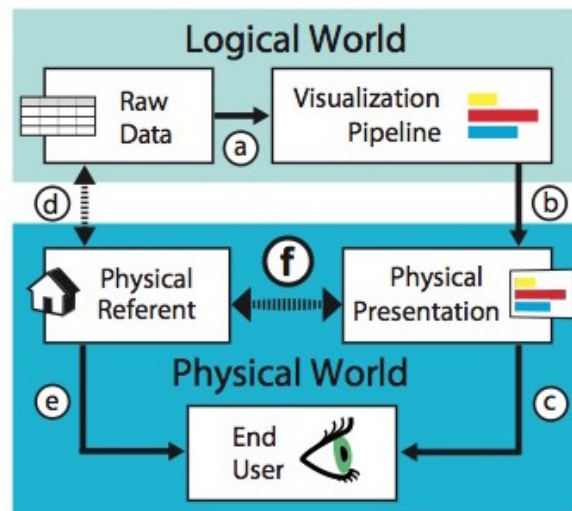


Fig. 3: Conceptual model of a spatially-situated visualization and analytics (adapted from Willett et al. [51]).

The conceptual model is illustrated in Figure 3. It covers both the logical world (top) and the physical world (bottom). Black arrows show information flows, while dashed arrows refer to conceptual relationships. Only information that flows from the data to the end user is shown here—other information flows will be covered in subsection 2.5 on interaction.

The flow of information starts with the raw data, on the top left. The conceptual model assumes that a visualization system turns this raw data into a visual form that humans can understand ($a \rightarrow b$). In information visualization, this process is generally computer-automated and can be characterized by a *visualization*

pipeline. The visualization pipeline applies a sequence of transformations to the raw data until a final image is produced. The different stages of the pipeline have been covered extensively in the past [12, 16, 20], but for the sake of simplicity, Figure 3 shows the entire visualization pipeline as a single block.

Both the raw data and the visualization pipeline exist in a logical world. This logical world relates to the physical world in two major ways [51]: through the data's *physical presentation* (b) and through the data's *physical referent* (d).

A *physical data presentation* (or physical presentation for short) is “*the physical object or apparatus that makes the visualization observable*” [20]. For example, suppose a person is viewing data about a house for sale. The data consists of information such as the house's size, number of bedrooms, price, or energy efficiency. The visualization pipeline (Figure 3, a→b) describes the process by which this data is turned into a visual form (e.g., a bar chart representation, a numerical table or a starplot). For the observer to be able to see it (c), the visual form needs to be brought into existence in the objective world (b). A physical presentation can be an image displayed on a particular computer screen or projected on a particular physical surface, or ink on the surface of a newspaper page. It can also be a physical artifact, e.g., a data physicalization [30]. Virtual presentations as seen in AR systems will be discussed in subsection 2.2.

The second way in which data is connected to the physical world is through *physical referents* (Figure 3-d). A physical referent is a *physical object or physical space to which the data refers* [51]. In the case of our house buyer, the dataset refers to a particular house that exists in the physical world. As the relationship (d) is a conceptual relationship, a dataset can have many possible referents [51]. For example, one can decide that the house dataset refers to the house owner, or to the headquarters of the real estate company that manages it. Both exist in the physical world and could be at a very different location than the house. Finally, the physical referent may or may not be visible to the observer (e).

Whether the physical referent and the physical presentation are simultaneously observable largely depends on whether they share the same space, i.e., on the physical distance that separates them (Figure 3-f). For example, our house buyer can choose to visualize the house data on a laptop in her own house, in which case the house of interest will likely be visually inaccessible. If in contrast, the user stands in front of the house of interest and visualizes the data on her smartphone, or if the data is visualized on a sign placed on the house, the distance (f) will be small enough that both the data and its referent can be examined (c,e). In such a case, the visualization is referred to as *spatially situated*.

A visualization is *spatially situated* if its physical presentation is close to the data's physical referent.

The term “close” is left vague on purpose, as spatial situatedness lies on a continuum: the visualization shown on a sign placed on the house is spatially more situated than the visualization shown on a bystander's smartphone, which is spatially more situated than the visualization shown on a desktop computer.

Although the definition presented here is far from capturing the full richness of the term “situated” (see, e.g., [6, 27]), it has the merit of clarifying what is situated with respect to what. It also clarifies that spatial situatedness cannot be a property of the data, since data is a purely logical entity. Similarly, when a visualization is referred to as being spatially situated or non-situated, this is really referring to the *physical presentation* of a visualization system, not to the visualization as a representation. For example, two different smartphones can display the same bar chart about car consumption data, with one smartphone being far from the car and the other one being inside. While the two smartphones show the *same* visualization, one is situated and the other one is not. For similar reasons, it would be meaningless to ask whether a bar chart is a more situated visualization than a scatterplot, at least within the present framework of spatial situatedness. Other forms of situatedness will be discussed later in the chapter.

Finally, the conceptual model makes it clear that situated visualizations do not need to assume a particular technology. A situated visualization can be created with rudimentary means, e.g., by printing a visualization of a house’s data and bringing the printout to the house. Conversely, an AR visualization system can be non-situated, e.g., when two users interact with a 3D visualization that shows data about a physical entity located far away. It is clear, however, that new and emerging technologies make it possible to create elaborate forms of situated visualizations.

2.2 Physically- vs. Perceptually-Situated Visualizations

The *physical distance* separating a physical presentation and its physical referent may not necessarily match the *perceived distance* between them [51]. One reason is that distances are perceived in a relative manner. Thus, a one-meter separation may appear large if both the physical presentation and the physical referent are small and the observer is standing close to them (e.g., visualizing data about a rare stone), while the same distance could be negligible in the opposite case (e.g., visualizing data about a distant mountain).

Discrepancies between physical and perceived distances are very common in AR setups [55]. For example, consider a person wearing an HMD who stands in front of a house and sees a data visualization overlaid on the house. The physical data presentation is literally the array of pixels on the surface of the physical display worn by the observer, and could be dozens of meters away from the data’s physical referent (the house). However, the AR system could be designed in such a way that the visualization appears to physically coincide with the house.

The previous definition of spatial situatedness can be either left ambiguous on purpose or refined to distinguish between physical and perceptual distance:

A visualization is *physically situated in space* if its physical presentation is physically close to the data’s physical referent.

A visualization is *perceptually situated in space* if its (physical or virtual) presentation appears to be close to the data’s physical referent.

As stated by the last definition, perceptual situatedness can refer to virtual presentations. For example, if a visualization is rendered next to a house using an HMD, the visualization seen by the user is virtual rather than physical. Since the physical/virtual distinction is not without conceptual difficulties [88], it is often easier to consider the *percept* elicited by the physical presentation rather than the presentation itself [20]. Thus, an alternative definition is as follows:

A visualization is *perceptually situated in space* if its percept appears to be close to the percept of the data's physical referent.

This last definition works for all setups, irrespective of the display technology.

2.3 Embedded Visualizations




Fig. 4: Examples of embedded and non-embedded visualizations (all situated).

Embedded visualizations are situated visualizations that are deeply integrated within their physical environment [51]. Figure 4 shows examples of embedded and non-embedded situated visualizations. If data about a house is shown on a single visualization placed next to (or inside) the house as in Figure 4-(a), the visualization is simply situated. If, however, different sub-elements of the visualization align with different sub-elements of the physical house (b), the visualization becomes embedded. For example, energy consumption data could be displayed within each room of the house, or next to every power socket. Thermal isolation data could even be visualized as “heat maps” on the walls themselves using AR displays or simply thermochromic paint. Willett et al. [51] discuss several such examples of highly-embedded data representations, while Hanrahan [24] and Offenhuber [46, 47] specifically discuss physical implementations.

A visualization is *spatially embedded* if each of its physical sub-presentations is close to its corresponding physical sub-referent.

Embedded visualizations assume multiple sub-presentations that are aligned with their corresponding sub-referents. Thus, if energy consumption is displayed near each power socket, as proposed before, but a house only has a single power socket, the visualization ceases to be embedded and becomes a regular situated visualization. Conversely, it is possible to create an embedded visualization simply by duplicating situated visualizations. For example, if the setup in Figure 4-(a) is duplicated across an entire neighborhood, the entire set of physical presentations becomes an embedded visualization, as shown in Figure 4-(c). The physical referent becomes the set of all houses for sale, and each house becomes a sub-referent. If, however, the same data is shown on a single visualization placed somewhere in the neighborhood (e.g., as a map of all houses for sale), as shown in Figure 4-(d), then the visualization would be situated but not embedded.

2.4 Temporally-Situated Visualizations

The same way data can be thought of as referring to a concrete region in space (i.e., the region occupied by the physical referent), data can be thought of as referring to a concrete region in time. For example, an energy consumption display can show data for the present day, from the day before, for several consecutive days (e.g., as a time series), or can even show forecast data about the future. This region in time can be referred to as the data's *temporal referent*. It is then possible to compare the temporal referent with the moment in time a visualization is observed, derive a measure of temporal distance, and characterize a visualization's temporal situatedness [51]:

A visualization is *temporally situated* if the data's temporal referent is close to the moment in time the physical presentation is observed.

An example of a visualization that is highly situated both spatially and temporally is a car's speedometer, because it is located within the data's physical referent (the car) and shows real-time data. A car's mileage display is similarly highly situated spatially but less so temporally, since it shows data about the present but also about a large segment of the past.

2.5 Interaction

Figure 5 shows the different ways a user can interact with a spatially-situated visualization. On any interactive visualization system (situated or not), an analyst can generally interact with the visualization by altering its pipeline (a). Operations such as filtering data, changing the visual representation, highlighting data points, or zooming, are all modifications to the visualization pipeline and have been extensively discussed by Jansen and Dragicevic [20]. Usually, these interactions




Fig. 5: Interaction with a spatially-situated visualization. Red arrows represent possible flows of information from the user to the system. The flows (a) and (b) are supported by any interactive visualization system (situated or not), while the flow (c) is specific to situated visualizations.

are implemented through *instruments*, i.e. combinations of software and hardware elements that interpret users' actions into changes in the pipeline [5, 20].

A second way of interacting with a visualization system is by directly altering its physical presentation (b). While desktop visualization systems generally offer limited interaction possibilities at this level, physical visualizations can let users filter, compare or reorder data by rearranging physical elements [20, 30]. In addition, users can alter their percept of a physical presentation by moving it or by moving around [20, 28]. Thus, a rich set of interactions can take place in the real world outside the visualization pipeline. Some of these physical interactions can be sensed and reflected back to the pipeline (right black arrow on Figure 5).

A situated visualization system offers a third mode of interaction through the physical data referent (c). Not only does a situated visualization make the physical referent observable, but it also generally makes it reachable and manipulable [51]. Thus, an analyst can use insights gained from the visualization to take immediate action, such as fixing a thermal leak in a room or removing cancerous cells. In contrast to visualizations that are not situated in space or in time, physical action can immediately follow analytical reasoning and decision-making.

In case the physical referent is the data source and the system implements real-time sensing (see Figure 3-d), analysis and action can be intertwined. For example, a thermal leak visualization could dynamically update itself as thermal leaks get fixed, and assuming a sufficiently advanced technology, a 3D body scan visualization could update itself as cancerous tissue gets removed. In these examples, the end user interacts with a visualization *by modifying the data itself*,




Fig. 6: A water tunnel visualization where a wing model can be rotated to examine the impact of orientation on aerodynamism [32]. Here the physical referent is the wing and the physical presentation is the set of water bubbles.

a mode of interaction that classical visualization systems generally do not support. Although in these examples the ultimate task is to take action on the physical referent, this mode of interaction is compatible with purely epistemic tasks. For example, an airplane designer could use a physical or virtual wind tunnel visualization on a malleable or articulated model of an airplane, and physically manipulate the model to explore how different shapes or orientations impact aerodynamism (Figure 6).

2.6 Levels of Situatedness

In situated analytics, *situatedness* is a multidimensional property, with each dimension lying on a continuum. As previously discussed, a key element to consider is the *spatial situatedness* of the data visualization employed, i.e., to what extent its physical presentation is close to (or appears to be close to) the data's physical referent. When this distance is sufficiently low for the visualization to qualify as spatially situated, a finer way of assessing its situatedness is by considering its level of *spatial embedding*. As discussed in Section 2.3, spatial embedding captures to what extent the geometry of the visualization's physical presentation *aligns* with the geometry of the physical referent. Since spatial situatedness is a necessary condition for being spatially embedded, spatial embedding can be seen as a stronger form of spatial situatedness. For example, in Figure 4, the setups a) and d) are spatially situated but not spatially embedded, while the setups b) and c) are spatially embedded. The level of spatial embedding lies on a continuum and depends on several factors such as the number of physical sub-presentations and their distance to their physical sub-referent. For example, a system that overlays a continuous heat map on a physical surface to display its temperature at every single point (e.g. using AR techniques or thermochromic paint) [47, 51] is more deeply embedded than a system that uses an array of thermometers.

There are also non-spatial forms of situatedness. For example, the level of *temporal situatedness* is a non-spatial form of situatedness, i.e., to what extent the data's temporal referent matches the time of observation. The interactions supported (especially when the physical referent can be directly manipulated for pragmatic or for epistemic purposes, as discussed in Section 2.5) also likely participate in the observer's subjective impression of situatedness. Section 4 will cover more specific examples of interaction styles involving the manipulation of physical referents for epistemic purposes. For now we will go through simpler and more classical examples of situated analytics systems.

3 Examples of Situated Analytics

This section presents examples of situated analytics used in real world applications. Key characteristics include situated virtual data and associated analytics.

3.1 Pollution Monitoring

NoxDroid is a sensor system aimed at monitoring air quality in cities. As shown in Figure 7, NoxDroid is a small mobile sensor device built and mounted on bicycles by volunteers. The sensor provides low fidelity real-time feedback on air quality as people ride their bicycles (green: Nox level are well below the official limit; yellow: Nox level are just below the limit; red: Nox level is around or above the limit). The sensor connects to an Android application to upload its data, share it with others, and more advanced functionalities. This enables cyclists to analyze pollution level and navigate in the sensor history through their mobile-phones, in situation, and chose cycling routes accordingly.




Fig. 7: NoxDroid situated analytics. Left: sensor with embedded pollution indicators. Right: smartphone application with contextual measures.

3.2 Personal Protective Equipment (PPE) Donning/Doffing

During 2014-2015, Prof. Greg Welch directed a small group of graduate computer science students at the University of Central Florida in the development of a system to allow volunteers around the world to help check the integrity of personal protective equipment (PPE) being worn by healthcare providers caring for patients with deadly viruses such as Ebola, before the providers come into contact with any patients. See Figure 8(a) for an example of a provider in their PPE. The system, called *SterilEyes*, consists of a smartphone app and a back-end server system that allows the provider to quickly capture video imagery of themselves in their PPE, that is instantly made available to certified volunteer observers from around the world who can instantly check the provider's PPE and vote on the quality of the protection over the entire body.



Fig. 8: Left (a): Example of healthcare provider wearing personal protective equipment (PPE) such as would be used when caring for patients during an outbreak of a deadly contagious disease such as Ebola. Right (b): Example observer (crowd) analytics situated on the appropriate body parts.

The smartphone app would be used in two different circumstances and corresponding modes: the provider mode and the observer mode. After donning the PPE, and before entering the potentially contagious space around the patient, the provider or a colleague would capture video of the PPE on the provider—in particular in locations known to be problematic such as the neck, wrist, and ankle connections/seals. Each observer would be notified and if possible/willing would respond by selecting each video—associated (situated) with a part of the body—and rating it. Progressively, as the observers around the world respond, the back-end system would calculate and update the displayed confidences associated with each critical body location. See Figure 8(b) for an early example of a visualization presented to the provider as the observer votes evolve.

3.3 Future Farming

In a recent TechCrunch article [36], contributor Jeff Kavanaugh discussed a futuristic (but not far off) vision for farming where sensors collect data about plant and soil health for example, machine learning or approaches continually perform some analytical analyses, and Mixed (Augmented) Reality is used to allow a farmer to see and interact with the data in place to “help both farmers and gardeners to monitor and manage crop health.” Kavanaugh also described an Infosys open-source digital farming project called Plant.IO*. Kavanaugh describes the vision where PVC pipes provide a frame for devices such as sensors and plant growth lights, a remote server continually analyzes and predicts plant health, and “AR-capable glasses” like the HoloLens could be used to both visualize crop analytics in place and affect plant health via AR actuator interfaces that control fertilizer, water flow, growth lights and more. See also <http://plantio.de> (website in German).

There are additional elaborations on these ideas. For example, on the heels of the TechCrunch article, Rob LaPointe of SDI presented some related ideas [38]. LaPointe pointed out how sensors for monitoring crops, weather stations, satellite information, etc. can be cross-referenced to specific crops, and analyzed by AI algorithms that are informed by the latest agricultural publications, with the results being “wirelessly transmitted to a set of AR Goggles” that provide the farmer with information about water, light, and fertilizer needs for each plant in (or region of) the field.



Fig. 9: Blended user interface.

4 Blended Situated Analytics Controls

This section explores a particular style of situated analytics interaction methodology, *Blended Situated Analytics Controls* [20]. Existing solutions for AR interaction techniques provide users with a limited number of predefined interaction perspectives for the presented data and the input controls are either static for all objects or have a limited number of controls that can be associated. Working

with abstract information in AR requires more methods of interaction than the traditional approach, allowing the user to manipulate the data freely and explore relationships between data in two different spaces: physical and virtual.

The blended user interface is a promising SA tool, which fuses the controls into the physical referents (physical objects), and derives the controls' appearance from the physical context, affording dynamic widget appearance and layout techniques [18, 21] (see Figure 9). The appearance of the controls is dynamic depending on their placement and function on the physical object. The novelty of the techniques is their context-aware dynamic blending of physical/virtual user interface controls allowing seamless transition between the physical and information spaces. The blended user interface has three main components: blended controls, blended views, and the blended model.




Fig.10: Situated analytics blended controls. (a) Users can view the attached information, (b) interact with physical objects to explore more information, (c) and view/compare the information associated with multiple physical objects.

4.1 Overview of Blended Controls

The blended controls allow users to 1) view in a meaningfully fashioned abstract data with their relationships and 2) apply operations such as select, zoom, search, filter, and analyze. Figure 10 demonstrates the blended controls within the context of a shopping task, enabling user exploration and interaction with information in novel ways. In Figure 10-a, a user explores a product's overall information, presented to them as a virtual annotation overlaid on the top of the physical box. Information of interest is highlighted—for example, if the user is searching for Australian-made products, the Australian logo can be highlighted. Figure 10-b shows a user interacting with the physical referent to explore more information (Details-on-Demand). The user can explore information, such as touching the ingredients listing printed on the product's box, and the SA system

will display additional detailed visual analytical information as an AR overlay. This information representation is based on the product's ingredients (for example, the percentage of the user's daily recommended intake (RDI) for a nutritional category such as sugar or fat that the product contains). SA also allows a user to analyze and compare information between products (seen in Figure 10-c). As an example, when a user selects two products and places them side-by-side, the SA system presents a comparison of the two products to the user.



Fig. 11: Blended zoom.

Blended interactions differ in that the user is not restricted to the planar screen of a tablet device, keyboard, or an indirect pointing device. Instead, the user is guided by the form factor of the underlying physical referent such as a product box, physical lever or button, or tangible artifact. The user is therefore not disconnected from the physical referent as in the mid-air interactions. In this scenario, the haptic feedback is given through the product while an HMD provides the overlay of visual information. Figure 11 depicts the user employing the pinch gesture, zooming on a supermarket product to provide an easier to read portion of the label. Using context-aware blended controls alters the assigned UI control. For instance, a blended selection can be altered based on the ratio between the width and the height of the tracked image. Figure 12-a shows a calculation of the percentage of the juice in the cup, computing the calorie content by using the one-dimensional slider. In the box scenario, the type of slider for the same operation has been changed; a two-dimensional slider has been used due to the physical shape of the box (Figure 12-b). This shape adaptation feature reduces the complexity of UI design, enabling the storing of the UI properties and constraints, which will automatically blend the relevant UI components to the physical referents.

Blended interaction can also provide intuitive interaction with the physical space, such as proximity and collision. Proxemics is an interaction based on the user's view, by calculating the distance between the user's view and the tactile physical referents. By moving the referent nearer and further to the user's view, the amount of data presented changes. Where holding a physical referent closer the user's view, this closer view reflects an interest in the object. Figures 13-a and 13-b show an implementation of proximity exploration, between overview and




Fig.12: Blender sliders changing their appearance based on the physical object's shape.

detailed representations. The overview shows the book's ranking, and the detailed view shows the book's table of content. Through an AR display, the user can see the virtual overview annotation (see Figure 13-a), and by bringing the book nearer to the user's view, this will invoke the detailed view. The registered table of content is color-coded based on the user's query and preferences, as the green shows highly related book sections moving towards red for the least preferable ones. The white text means that the section title is not about the user's entered query (see Figure 13-b). Another example is shown in Figures 13-c and 13-d demonstrating depth-level adaptation to override the small FOV challenge of the optical see-through devices. The technique arranges the data into multi-layers controlled by the distance between the user and the physical referents. Proxemic interactions invoke the visualization based on the cue's ratio to the physical referent's size. When the physical object is near to the user, its size will increase, which will invoke more visual cues.

Collision is an interaction based on the spatial relationship of multiple objects to provide information pertinent to the objects' combination. Collision can be to aggregate the virtual data associated with the physical referents. Figure 14 depicts a collision-based example, allowing the user to calculate the total calorie intake of two products by aligning the physical objects side-by-side. As the user calculates the total star point of two products, by putting them side-by-side. When the user holds the chips by the crackers box, it shows a low star point value, a low nutrition outcome. Then the user checked the juice with the cracker by aligning the juice by the crackers box, which showed a better star nutrition value than the chips combination.

4.2 The Blended Views

The blended views hold the GUI elements and are responsible for generating the blended widgets. The uniqueness of the blended views is attaching the widgets and the visuals based on the physical context to leverage their meaning. The




Fig. 13: Proxemic interaction.

semantic fusion of the UI elements to the physical world allows physical referents to be part of the interactive information process, working in concert with the controls to achieve the blending aim. Figure 15 shows a blended menu that changes its appearance and items based on the physical context. The menu can be dragged and relocated to any place on the physical referents, with dynamic size, shape, and color of the menu based on the physical context. These menus use pre-stored regions' meta-values to restrict the location of the menu on the physical box.



Fig. 14: Collision interaction to combine products' nutritional value.

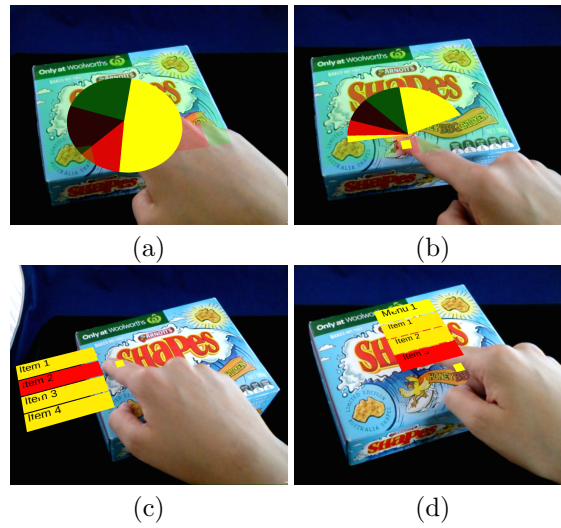


Fig. 15: Blended menu.

4.3 The Blended Model

The blended model allows a two-way, real-time association between the physical and the virtual information, enabling contextual and situation awareness for the interactive information process. Figure 16 depicts the user experience during a series of interaction states in the blended space for picking a meal. The user moves between the states based on the predefined parameters, defining the invoking trigger for each state, permissions, and parameters associated with the mapped contextual feature. In Figure 16-(a), the user moves the AR display with a camera to scan products on the shelf and select the product. The selected product is highlighted by a green frame. In Figure 16-(b), the user takes one of the products off the shelf, as they are interested in more detailed information about this particular product. This user's interaction will invoke a detailed view of the product; the user is holding, enabling region selection. The user selects the product's logo, then tilts the box to select the flavor region. Finally in Figure 16-(c), the user starts to interact with two fingers on the box surface, the interaction control changes to a magnifying pinch zoom.

5 Design Considerations for Situated Analytics Systems

This section discusses design considerations for SA. As usual when designing interactive systems, no single perfect solution that would fit all intended use cases exists. The design of SA systems must account for the physical environment, data, and viewers with which it will be used, as well as the presentation and sensing technology used to implement it. However, because only a small number of SA systems have been developed so far, few guidelines and best practices




Fig. 16: Interaction states in the blended space.

for situated tools currently exist. This section goes on to discuss a variety of practical design decisions that, based on the authors' experience, have important trade-offs and repercussions for situated tools. The discussions are structured around the components of situated analytics systems as introduced earlier in this chapter (see Section 2 Characterization of Situated Visualizations). The section starts by discussing components of the physical world—the physical referents, the physical presentation, and the users—and then discusses design considerations for components in the logical world—the data and the visualization pipeline.

5.1 Physical World

Situated analysis tools are characterized by their relationship to the physical world. As a result, these tools must account for the environments and physical referents to which the data are related. Moreover, the visual feedback from SA tools themselves must ultimately emanate from some source (typically an object, projector, or display) in the physical world. Therefore, the design of situated tools must take into account the physical characteristics of the referents and environments with which the systems will be used, as well as the physical limitations of the presentation technologies, such as with Blended Situated Analytics Controls.

Physical Referents and Environments. Essentially any object, person, or environment can potentially become a referent, given a dataset that somehow relates to them. However, some referents and environments lend themselves more readily to situated analysis than others.

Size and visibility. In some cases, the physical and visual characteristics of the referents themselves may dictate whether or not situated analysis is practical.

For example, referents that are very large or very small may be difficult to examine or compare and, as a result, may not provide useful context. Similarly, situated analysis may be challenging in cases where environmental constraints like distance or occlusions make it difficult or impossible to examine important physical referents simultaneously.

Identifiability and dynamicity. Similarly, the viability of a situated analysis system may be dictated by how easy it is to distinguish, track, and connect physical referents with data about them. In order to display the appropriate data in context, it must be possible for either the system or the viewer to identify corresponding referents and environments. As a result, environments with many similar referents or high dynamicity (such as quickly moving swarms or crowds) may pose serious implementation challenges.

Safety and security. In other situations, the risks to the physical safety of viewers and those around them need to be considered. Situated analysis may not be viable in locations where visual augmentations might distract viewers' attention away from the environment and disrupt critical tasks like driving, flying, or operating machinery.

Physical Presentation. When considering the design of a situated system, one of the most important decisions is whether to display information virtually—using projection or overlays that are visually superimposed on top of the environment—or physically—via physical output mechanisms that are situated in the environment itself. This distinction is not strictly binary. For example, physical screens placed in an environment may be concrete objects, but provide largely virtual content. As a result, the choice to situate data displays via primarily physical or primarily virtual means will likely have a considerable impact on the scalability of the system, as well as the kind of observations and interactions it supports.

Virtual output. At one extreme are overlays produced by HMDs. While these kinds of hardware can make it possible for wearers to superimpose data on top of environments and objects, the relationship between the presentation and referent is purely visual and largely individualized. As a result, the physical presentation provides no tactile or physical feedback and has no direct physical relationship in the environment. Computer vision and other tracking techniques can be used to align virtual overlays and controls with specific physical referents (as with Blended Situated Analytics Controls [20]). However, correctly aligning presentations and referents can be challenging in complex environments, and providing appropriate depth cues and haptic feedback may be difficult.

This degree of independence may be useful in situations where an analytic system needs to display very large numbers of data points, or where data must be displayed in areas with no corporeal physical referents. Virtual displays can also deal with environments that contain dynamic referents whose form or identity may change over a short period of time. Tasks like displaying data about tens of thousands of parts in a manufacturing plant or visualizing air quality data in the center of an open space may benefit from these kinds of virtual presentation mechanisms. In fact, virtual presentations may be the only viable options if

referents are physically very large or small, or if they are fragile, distant, or otherwise inaccessible.

Physical output. Alternatively, more concrete physical presentations may be beneficial when analyses are centered around a smaller number of stable referents. In these cases, presenting data via physical output in the real world may support a more direct coupling between the information and the referents it relates to. Displaying visual output via displays which are physically attached to their referents provides an immediate coupling between data and context, and can make it easier for viewers to examine the two simultaneously. Tight physical connections may also provide more physical affordances for interaction, making it easier for viewers to interact with the analytic tools using the referents themselves as input controls.

Physical output may also be helpful in collaborative situations where multiple viewers need to examine and interact with the analytics tools simultaneously. Because physical presentations preserve real-world visibility and interaction cues, they may make it easier for individuals to understand what their collaborators can see and manipulate, and what they cannot. However, physical outputs like these are also more difficult to secure, especially when displays are not dynamic or the identities of viewers cannot be determined. In these situations, systems may need to rely on restricting physical access to the space or encoding data displays to reduce their intelligibility.

More generally, physically attaching presentations of data to referents may be a practical technical solution—especially in situations where the set of referents is human-scale, small in number, and relatively static (small numbers of people, animals, tools, objects, rooms, etc.). In these cases, physically associating data presentations with their referents may eliminate or reduce the need to track the referents' locations in order to display data at the right place and time.

Embedded vs. Situated Output. When designing situated analysis tools, developers may also need to consider the degree to which the system is connected to individual physical referents. On one hand, systems may be only lightly situated—presenting data in a relevant environment, but ignoring the specific orientations and positions of related people, objects, and spaces. Alternatively, systems can be more deeply embedded, placing presentations of data on or near their referents. Determining what kind of embedding is appropriate may depend not only on the available presentation technologies but also on the complexity of the environment and viewers' likely tasks.

Embedded. Choosing to embed presentations of data directly alongside their corresponding referents may present a number of benefits. For example, embedding output on or near relevant objects, people, and environments can make it easier for viewers to understand the relationship between data and physical referents, and take action based on it. Doing so also increases the likelihood that viewers can correctly identify which referent the data corresponds to. Similarly, embedding makes it easier for viewers to perceptually integrate

information from the dataset with relevant contextual information from the physical environment [41].

Situated. In contrast, situated views may often be less technically difficult to implement—as they only need to be presented in the appropriate environment at the right time, but do not need to be aligned with the individual referents in any particular fashion. In fact, simply placing an existing data display or analysis tool into an appropriate environment (for example, on a phone, tablet, or head-mounted display) can be enough to produce a situated analysis experience. Using a purely situated approach also ensures that the presentation of data is not limited by the physical positions of the referents. As a result, this strategy can make it easier to guarantee that viewers will be able to see and access the data, regardless of the environmental configuration.

Users. Important design decisions need to be made when considering who will be the end user(s) of the system. Will it be a single person or does the system need to be able to accommodate multiple users? If it is a multi-user system, will these users collaborate or work in parallel? Should all users be able to access the same data or are there restrictions on who can access what data? While some of these questions have already been mentioned above when considerations were discussed for choosing physical presentations, the discussion now focuses on considerations taking the users' point of view.

Privacy and collaboration. Situated analytics can be performed on many different types of data. Some of these data can be private or confidential and may need to remain hidden from other people sharing the same environment—a problematic shared with any “sensing system” [7]. Privacy can be assured implicitly by using an HMD instead of publicly visible displays. Yet such HMDs make collaboration more difficult as they require that each collaborator have their own device. Furthermore, content and changes need to be dynamically coordinated across all of the individual displays. As a result, this type of setup can make it more difficult for multiple users to determine whether they are seeing the same data. Thus, shared situated displays may be preferable for applications that involve collocated collaboration.

Access-controlled environments, for example in a corporate setting, present a special case, particularly if the SA system has access to information on who is allowed to view what data and who is currently in a room, for example, through tracked badges [61]. In such cases, a SA system could ensure confidentiality by adapting the visual output so that it shows only content accessible to all current viewers. In turn—with HMDs—data with different levels of detail can be shown to different people according to their access rights without requiring them to leave the physical environment before certain data can be viewed.

5.2 Logical World

When creating situated analysis tools, designers and developers must also consider the logical world—including the data and the visualization pipeline. While these

underlying constructs are shared with other non-situated visualization systems, situated analysis tools introduce a number of new complexities.

Data. Data in an SA system are viewed in spatial proximity to the physical referents. Beyond spatial proximity, temporal proximity of the data can be relevant as well. For example, White and Feiner report a small field study where their users would have preferred to access real-time data about air pollution in a city instead of “stale”, previously made measurements [64]. In cases where the intended purpose of the SA systems is to (possibly) take action in the physical environment, for example, to explore how such modifications affect sensor readings, live data become crucially important.

Tracking physical referents. In order to display data at the appropriate time and place, situated tools may require considerably more information about the environment in which they are used than traditional analysis tools. Desktop data analysis packages can render the same visual output on a wide variety of different commodity hardware regardless of their surroundings. In contrast, situated tools will typically display different information (often in different configurations) depending on where they are used. As a result, situated tools need mechanisms for uniquely identifying and tracking physical referents around them, and for associating referents with related data. This means that SA tools may often need more elaborate data models that can represent referents as logical objects within the system, as well as mechanisms for authoring and updating relationships between referents and data. Unless the visual output of an analysis system is physically connected to the referents, systems must also be able to track and process referents’ position, movement, and visibility, and use this data to update the situated presentations in real time.

Visualization pipeline. The visualization pipeline transforms the data into a visual representation that can be displayed using physical presentations as was discussed earlier. The physical presentation and the physical referents need to be taken into account when designing the visualization pipeline. Particularly, the geometry of the physical referents and physical presentations are important to consider when choosing appropriate visual encodings to ensure that the SA system informs but not hinders users’ actions.

Visual encodings. Most of previous work studying the perception of visual encodings focused on two-dimensional encodings which are best adapted to the presentation of 2D screens. With situated analytics, data are shown in physical, three-dimensional environments thus it may be beneficial to consider 3D encodings for such systems especially in the case of head-mounted displays or physical output. White and Feiner [64] gauged preferences from their field study users who expressed a preference for representations that facilitated making the link between the location of sensor readings and their value. For example, they preferred the use of spheres over cylinders to represent sensor readings associated to a particular position in 3D space.

A first study on size perception of physical 3D marks found that the size perception of 3D bars and spheres is not—as previously assumed—systematically biased if an appropriate transfer function is chosen [31]. However, if the physical referents feature flat 2D surfaces, then a 2D visual encoding matching these surfaces would be preferable. Yet prior work on the perception of two-dimensional visual marks displayed on large wall displays suggests that these are not accurately perceived when viewed at a non-perpendicular angle [8].

5.3 Summary

Situated analytics is an emerging theme and as of yet few empirical studies can give validated advice on how to best design such systems. This discussion refrains from making prescriptive design recommendations and instead focused on laying out the different aspects a designer of a situated analytics system should consider. At the same time the need for more empirical studies illustrates rich research opportunities within situated analytics. The open challenges are described in the next section.

6 Challenges and Research Agenda

Analytics moving into “the real world” raises challenges at multiple levels: technical ones, methodological ones, and conceptual ones. As has been seen in the previous sections, new typologies of analytical tasks are also emerging to account for a more “casual” approach to analytics. People conducting analysis in short situated bursts instead of long focused sessions in front of a computer. Supporting new analytical tasks should lead us to rethink how we design analytical tools, i.e., how do we prototype situated tools, and how we evaluate them. Besides new design methods, conceptual tools and technical frameworks will be needed to support development of these tools.

The envisioned pervasiveness of SA will require attention to new domains. Understanding expert tasks will not be enough, and designers will have to consider pleasure, engagement, or social acceptability. Finally, we should not forget to ask, what are the benefits and limitations of Situated Analytics, i.e., when is it worthwhile to offer such analytics, and when would people be better off with traditional analytics tools? There are clear trade-offs in terms of attention and information overload, privacy risks, and ethical concerns in case badly situated analytics could reinforce prejudices by only displaying a partial view of the situated data.

6.1 Visual Display

The emergence of novel display form factors and capabilities bears a direct impact on possibilities latent in Situated Analytics. From using HMDs to pico-projectors that facilitate group-based sense-making, necessary consideration needs to be placed on where to project information, how to embed content in the environment.

There is a difference between personal and group-based displays. In terms of personal displays, attention is needed to issues involving color blending (for superimposing virtual content on physical objects) [37], on issues of environmental saliency to suitably embed content [33], as well as on placement strategies for effective interaction with content. In terms of group-based displays, attention is needed on identifying how best to position the display to suit group work [13], on how best to offer shared displays with HMDs [10], and approaches for facilitating group interaction. With current advances, many of the above areas can be further explored in more depth. For example, brighter displays affect how content can be fused into the environment and advances in steerable pico-projection can facilitate more fluid approaches to SA.

6.2 Interactions Techniques

One primary challenge that needs addressing is the design of novel interactive tools for SA. Unlike traditional desktop environments, where devices and tools have been entrenched in the fundamentals of areas such as visual analytics, interaction interfaces and devices for SA is uncharted territory. As described above, one approach could include the use of Blended Situated Analytics Controls that fuse the user's interaction onto the physical objects in the environment. Such an approach relies on no more than the sensors HMDs are already equipped with, and thus offers an attractive solution for SA interactivity. However, when the embeddings are loosely connected to the physical objects, such as an entire environment, novel approaches are necessary. Ens and Irani [22] offer a preliminary discussion into the types of interactive devices possible for situated analytics. These include finger-worn sensors [60], digital pens [1], on-body interactivity [62] and the use of physical objects, that can be tracked by the displays worn by users [22]. Such forms of devices have shown little application to SA, and therefore re-examining these from the standpoint of specific usage scenarios (see Section 3) in SA is an important first step. Furthermore, such forms of interactivity have not been explored in the context of mobility or applications involving Augmented and/or Virtual Reality. Therefore pressing questions include: "can such devices be appropriated for tasks in SA?", "what environment properties affect the use of such forms of interactivity?", "how can such interfaces be made more efficient and optimized in the context of SA?". These and other questions can formulate the basis of a new research agenda in Situated Analytics.

6.3 Rethinking the Design Cycle

Designing visual analytics tools has traditionally centered on answering the specialized needs of experts. Situated analytics shifts the focus from experts to a much broader user population. Moreover, the context for which to design is also much more ill-defined. If visual analytics is typically conducted on a desktop computer with mouse and keyboard, exceptionally on a large wall display, situated analytics is, by definition, associated with any possible context. This radically changes the way analytical tools should be designed. New methods are

needed to better account for the situated aspect of visualisation: consider space, consider the unexpected, consider social acceptance, etc.

With ill defined tasks, the design process cannot rely as heavily on the collection of requirements, and the specification of needs. Designers will have to explore opportunities for design in a much more iterative way, sketching possible applications, testing them out and figuring whether they fit the needs of people. Such an iterative process requires new tools that enable quick sketch solutions to explore a design space.

Beyond sketches, developing fully functional situated analytics tools is still particularly complex, and costly. New frameworks guiding development, offering ready to use building blocks, could speed up development significantly.

Finally evaluation methods whether it is for early sketches or fully fledged applications must be refined. Because of the situated nature of tools, evaluations methods will have to incorporate some forms of field work to assess situated analytics applications on open-ended activities. This differs widely from well-defined tasks typically supported by visual analytics, for which methods to measure time, errors or insights have been developed. [Reference the evaluation chapter](#)

7 Future Work

Situated analytics is a new and emerging research field. Investigations are required into new display technologies, application domains, forms of data, interaction methodologies; just to name a few. Two particular research directions of interest are moving beyond spatial situatedness and tackling the ethical challenges this new research field presents.

7.1 Moving Beyond Spatial Situatedness

Situating analytics in the physical space, i.e., spatially close to the objects of interest, is the predominant strategy used for situated analytics. Section 2 discussed alternative dimensions, such as using time as the frame of reference. However, these physical dimensions only account for a limited part of users' situation. In the field of ubiquitous computing, situatedness has often been modeled alongside several dimensions describing what is generally referred to as *Context*. The notion of context helps to account for broader phenomena and people's activity.

For instance, while picking up a citybike at a station, the relevant information may not be the availability of bicycles in the pick-up spot since the user can already see whether there are free bicycles in its surrounding, but the number of free spots at the destination so that the user can adjust the travel goal, to the closest station with free parking spots. To develop a better understanding of people's intention, context should include information about the people involved in the activity, the set of devices available, the active applications, as well as sensor information such as light, noise, or temperature.

With such an understanding of context, it becomes possible to bring analytics relevant to users' activity rather than their location. For instance, a museum guide and visitors could benefit from SA about an art piece they stand in front of while touring the museum. But as the guide goes through the museum after hours with colleagues, to revise the guiding plan, they might get more relevant analytics about the time visitors spent in front of the piece, what other pieces they were interested in, etc. At the same location, situated analytics could take various form depending on the activity people are involved in.

7.2 Ethical Challenges

Widespread use of SA raises a number of ethical concerns that we should be aware of and ideally consider in the early phase of any project. Situating analytics whether it is in space, time or activities requires having rich datasets that have such detailed properties. Situated visualizations of the books borrowed in a library could be highly valuable to patrons, but could also lead to "leaks" revealing sensitive books rented by individuals. Similarly, SA of health related information could improve patients' understanding of treatments, their adherence, and their overall experience of illness, but capturing sensitive information about health and "projecting it into the world" should be done with extreme care to potential side effects.

Another concern of SA relates to the reinforcement of prejudices. Data collection is never exhaustive, and datasets offer a partial view of reality, however faithful we try to be. Selectively displaying data, is a way to introduce some prejudice by over-emphasizing some elements. For instance, designers working on urban situated analytics could decide to display information about criminality, but only display crimes against people or property, and not white-collar crimes which is harder to locate physically, leading to an emphasis on crime from a given population, while ignoring another.

8 Conclusion

As discussed in Chapter ?? and described above, there are two ways in which analytical activities can be immersive: either perceptually or cognitively. Perceptually speaking, a situated visualization can be thought of as more immersive than a non-situated visualization because the user is exposed to extra perceptual (visual or otherwise) information from the *physical* world. Naturally, a user is *always* situated in a physical environment (e.g., a desktop computer user can be situated in an office space), and this environment can be extremely rich in perceptual information (e.g., a messy desktop, a loud office). However, in non-situated systems this information is irrelevant to the analytic task — it is either filtered out if the user is focused, or disruptive if the user is not. In contrast, for the user of a situated visualization, a larger portion of the physical environment is task-relevant, and therefore the user can be considered as perceptually more immersed in the task. This is all the more true if both the physical referent and

the visualization occupy a large area (e.g., a tourist who walks in a city and explores city data using an AR display). This perceptual immersion and the relatively lower proportion of task-irrelevant stimuli can reduce the opportunities for distraction, and in turn increase the likelihood of being *cognitively* immersed in the analytic task.

This chapter defines a new method of Immersive Analytics referred to as Situated Analytics. The concept is characterized in greater detail, including the users, tasks, data, representations, interactions, and analytical processes involved. A set of case studies is examined in detail to elicit the best practices for situated analytics in action. Blended situated analytics controls are detailed as a particular method of developing situated analytics user interactions. A set of derived design considerations for building situated analytics applications is described. A research agenda of challenges and research questions to be explored in the future are presented.

Acknowledgments

[Acknowledgments go here.]

References

1. Aliakseyeu, D., Irani, P., Lucero, A., Subramanian, S.: Multi-flick: an evaluation of flick-based scrolling techniques for pen interfaces. In: Proceedings of the SIGCHI conference on Human factors in computing systems. pp. 1689–1698. ACM (2008)
2. Atzori, L., Iera, A., Morabito, G.: The Internet of Things: A survey. *Computer Networks* 54(15), 2787–2805 (2010)
3. Badam, S.K., Elmqvist, N.: PolyChrome: A cross-device framework for collaborative web visualization. In: Proceedings of the ACM Conference on Interactive Tabletops and Surfaces. pp. 109–118. ACM(2014), <http://dl.acm.org/citation.cfm?id=2669485>
4. Badam, S.K., Fisher, E.R., Elmqvist, N.: Munin: A peer-to-peer middleware for ubiquitous analytics and visualization spaces. *IEEE Transactions on Visualization and Computer Graphics* 21(2), 215–228(2015), <http://dx.doi.org/10.1109/TVCG.2014.2337337>; <http://doi.ieeecomputersociety.org/10.1109/TVCG.2014.2337337>
5. Beaudouin-Lafon, M.: Instrumental interaction: an interaction model for designing post-wimp user interfaces. In: Proceedings of the ACM Conference on Human Factors in Computing Systems. pp. 446–453. ACM (2000)
6. Beaudouin-Lafon, M., Mackay, W.E.: Research directions in situated computing. In: Extended Abstracts on Human Factors in Computing Systems. pp. 369–369. ACM (2000)
7. Bellotti, V., Back, M., Edwards, W.K., Grinter, R.E., Henderson, A., Lopes, C.: Making sense of sensing systems: Five questions for designers and researchers. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. pp. 415–422. CHI '02, ACM, New York, NY, USA(2002), <http://doi.acm.org/10.1145/503376.503450> doi: 10.1145/503376.503450
8. Bezerianos, A., Isenberg, P.: Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE Transactions on Visualization and Computer Graphics* 18(12), 2516–2525 (2012)

9. Billinghamurst, M., Clark, A., Lee, G.: A survey of Augmented Reality. *Foundations and Trends in Human-Computer Interaction* 8(2-3), 73–272 (2015)
10. Billinghamurst, M., Kato, H.: Collaborative augmented reality. *Communications of the ACM* 45(7), 64–70 (2002)
11. Bimber, O., Raskar, R.: *Spatial Augmented Reality: Merging Real and Virtual Worlds*. A. K. Peters, Ltd., Natick, MA, USA (2005)
12. Card, S.K., Mackinlay, J.D., Shneiderman, B.: *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann (1999)
13. Cauchard, J.R., Löchtefeld, M., Irani, P., Schoening, J., Krüger, A., Fraser, M., Subramanian, S.: Visual separation in mobile multi-display environments. In: *Proceedings of the 24th annual ACM symposium on User interface software and technology*. pp. 451–460. ACM (2011)
14. Chandler, T., Cordeil, M., Czauderna, T., Dwyer, T., Glowacki, J., Goncu, C., Klapperstueck, M., Klein, K., Marriott, K., Schreiber, F., Wilson, E.: Immersive analytics. In: *Proceedings of the IEEE Symposium on Big Data Visual Analytics*. pp. 73–80. IEEE (2015)
15. Chen, G., Kotz, D., et al.: A survey of context-aware mobile computing research. Tech. Rep. TR2000-381, Department of Computer Science, Dartmouth College (2000)
16. Chi, E.H.h., Riedl, J.T.: An operator interaction framework for visualization systems. In: *Proceedings of the IEEE Symposium on Information Visualization*. pp. 63–70. IEEE (1998)
17. Elmqvist, N., Irani, P.: Ubiquitous analytics: Interacting with big data anywhere, anytime. *IEEE Computer* 46(4), 86–89 (2013)
18. Elsayed, N., Thomas, B., Marriott, K., Piantadosi, J., Smith, R.: Situated analytics. In: *Proceedings of the IEEE Symposium on Big Data Visual Analytics*. pp. 1–8. IEEE (2015)
19. Elsayed, N., Thomas, B., Smith, R., Marriott, K., Piantadosi, J.: Using augmented reality to support situated analytics. In: *Proceedings of the IEEE Conference on Virtual Reality*. IEEE (2015)
20. Elsayed, N.A.M., Smith, R.T., Marriott, K., Thomas, B.H.: Blended UI controls for situated analytics. In: *Proceedings of the IEEE International Symposium on Big Data Visual Analytics*. pp. 1–8. IEEE (2016)
21. Elsayed, N.A., Thomas, B.H., Marriott, K., Piantadosi, J., Smith, R.T.: Situated analytics: Demonstrating immersive analytical tools with Augmented Reality. *Journal of Visual Languages & Computing* 36, 13–23 (2016)
22. Ens, B.M., Finnegan, R., Irani, P.P.: The personal cockpit: a spatial interface for effective task switching on head-worn displays. In: *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. pp. 3171–3180. ACM (2014)
23. Gruber, L., Richter-Trummer, T., Schmalstieg, D.: Real-time photometric registration from arbitrary geometry. In: *Proceedings of the IEEE International Symposium on Mixed and Augmented Reality*. pp. 119–128. IEEE (2012)
24. Hanrahan, P.: Self-illustrating phenomena. In: *Visualization, 2004*. IEEE. p. xix. IEEE (2004)
25. Huang, D., Tory, M., Aseniero, B.A., Bartram, L., Bateman, S., Carpendale, S., Tang, A., Woodbury, R.: Personal visualization and personal visual analytics. *IEEE Transactions on Visualization and Computer Graphics* 21(3), 420–433 (2015)
26. Hull, R., Neaves, P., Bedford-Roberts, J.: Towards situated computing. In: *Proceedings of the International Symposium on Wearable Computers*. pp. 146–153. IEEE (1997)

27. Hull, R., Neaves, P., Bedford-Roberts, J.: Towards situated computing. In: Proceedings of the International Symposium on Wearable Computers. pp. 146–153. IEEE (1997)
28. Isenberg, P., Dragicevic, P., Willett, W., Bezerianos, A., Fekete, J.D.: Hybrid-image visualization for large viewing environments. *IEEE Transactions on Visualization & Computer Graphics* 19(12), 2346–2355 (2013)
29. Jansen, Y., Dragicevic, P.: An interaction model for visualizations beyond the desktop. *IEEE Transactions on Visualization and Computer Graphics* 19(12), 2396–2405 (2013)
30. Jansen, Y., Dragicevic, P., Isenberg, P., Alexander, J., Karnik, A., Kildal, J., Subramanian, S., Hornbæk, K.: Opportunities and challenges for data physicalization. In: Proceedings of the ACM Conference on Human Factors in Computing Systems. pp. 3227–3236. ACM (2015)
31. Jansen, Y., Hornbæk, K.: A Psychophysical Investigation of Size as a Physical Variable. *IEEE Transactions on Visualization and Computer Graphics* 22(1), 479–488 (Jan 2016), <http://dx.doi.org/10.1109/TVCG.2015.2467951> doi: 10.1109/TVCG.2015.2467951
32. Jordan, T.: Water Flow Visualization using Electrolysis Hydrogen Bubbles (youtube video). <https://youtu.be/memvL8NG8jc> (2013), [Online; accessed 17-November-2016]
33. Kalkofen, D., Mendez, E., Schmalstieg, D.: Interactive focus and context visualization for augmented reality. In: Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality. pp. 1–10. IEEE Computer Society (2007)
34. Kalkofen, D., Sandor, C., White, S., Schmalstieg, D.: Visualization techniques for Augmented Reality. In: Handbook on Augmented Reality, pp. 65–98. Springer (2011)
35. Kalkofen, D., Tatzgern, M., Schmalstieg, D.: Explosion diagrams in augmented reality. In: Proceedings of the IEEE Virtual Reality Conference. pp. 71–78. IEEE (2009)
36. Kavanaugh, J.: How mixed reality and machine learning are driving innovation in farming. <https://techcrunch.com/2016/11/17/how-mixed-reality-and-machine-learning-are-driving-innovation-in-farming/> (November 2016)
37. Kruijff, E., Swan II, J.E., Feiner, S.: Perceptual issues in Augmented Reality revisited. In: Proceedings of the ACM/IEEE International Symposium on Mixed and Augmented Reality. vol. 9, pp. 3–12 (2010)
38. LaPointe, R.: How AI and AR apps can change agriculture. <https://softwaredevelopersindia.com/blog/ai-ar-apps-can-change-agriculture/> (November 2016)
39. Lee, B., Isenberg, P., Riche, N.H., Carpendale, S.: Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics* 18(12), 2689–2698(2012), <http://dx.doi.org/10.1109/TVCG.2012.204>; <http://doi.ieeecomputersociety.org/10.1109/TVCG.2012.204>
40. Lepetit, V., Berger, M.O.: An intuitive tool for outlining objects in video sequences: Applications to augmented and diminished reality. In: Proceedings of the International Conference on Mixed Reality (2001)
41. Marner, M.R., Irlitti, A., Thomas, B.H.: Improving procedural task performance with augmented reality annotations. In: Mixed and Augmented Reality (ISMAR), 2013 IEEE International Symposium on. pp. 39–48. IEEE (2013)

42. McGill, M., Boland, D., Murray-Smith, R., Brewster, S.: A dose of reality: overcoming usability challenges in VR head-mounted displays. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 2143–2152. ACM (2015)
43. Milgram, P., Kishino, F.: A taxonomy of Mixed Reality visual displays. *IEICE Transactions on Information and Systems* 77(12), 1321–1329 (1994)
44. Mojang: Minecraft. <https://minecraft.net/en/> (December 19 2016)
45. Nguyen, M.: Augmented Reality: Will 2016 be the year of smart contact lens? <https://www.wearable-technologies.com/2016/02/augmented-reality-will-2016-be-the-year-of-smart-contact-lens/> (December 2016)
46. Offenhuber, D., Bertini, E., Stefaner, M.: Indexical visualization with Dietmar Offenhuber – data stories podcast (2016), <http://datastori.es/80-indexical-visualization-with-dietmar-offenhuber/>
47. Offenhuber, D., Telhan, O.: Indexical visualization—the data-less information display. *Ubiquitous Computing, Complexity and Culture* p. 288 (2015)
48. O'Regan, J.K., Noë, A.: A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences* 24(05), 939–973 (2001)
49. O'Regan, J.K., et al.: What it is like to see: A sensorimotor theory of perceptual experience. *Synthese* 129(1), 79–103 (2001)
50. Raskar, R.: *Projector-Based Three Dimensional Graphics*. Ph.D., University of North Carolina at Chapel Hill (2001)
51. Raskar, R., Welch, G., Chen, W.C.: Table-top spatially-augmented reality: Bringing physical models to life with projected imagery. In: *Proceedings of the IEEE and ACM International Workshop on Augmented Reality*. IEEE Computer Society, Washington, DC, USA (1999)
52. Raskar, R., Welch, G., Fuchs, H.: Spatially augmented reality. In: Behringer, R., Klinker, G., Mizell, D. (eds.) *Augmented Reality: Placing Artificial Objects in Real Scenes*, pp. 63–72. A.K. Peters Ltd., San Francisco, CA, USA (1998)
53. Raskar, R., Welch, G., Low, K.L., Bandyopadhyay, D.: Shader lamps: Animating real objects with image-based illumination. In: Gortler, S.J., Myszkowski, K. (eds.) *Rendering Techniques 2001, Proceedings of the Eurographics Workshop in London, United Kingdom*, pp. 89–102. Springer, New York, University College London (UCL), London, England (2001)
54. Roberts, J.C., Ritsos, P.D., Badam, S.K., Brodbeck, D., Kennedy, J., Elmqvist, N.: Visualization beyond the desktop — the next big thing. *IEEE Computer Graphics and Applications* 34(6), 26–34 (Nov/Dec 2014) doi: 10.1109/MCG.2014.82
55. Schmalstieg, D., Hollerer, T.: *Augmented Reality: Principles and Practice*. Addison-Wesley Professional (2016)
56. Slater, M.: Place illusion and plausibility can lead to realistic behaviour in immersive virtual environments. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1535), 3549–3557 (December 2009)
57. Thomas, B.H.: Have we achieved the ultimate wearable computer? In: *Proceedings of the International Symposium on Wearable Computers*. pp. 104–107. IEEE (2012)
58. Thomas, B.H., Marnier, M., Smith, R.T., Elsayed, N.A.M., Von Itzstein, S., Klein, K., Adcock, M., Eades, P., Irlitti, A., Zucco, J., Simon, T., Baumeister, J., Suthers, T.: Spatial augmented reality—a tool for 3D data visualization. In: *Proceedings of the IEEE International Workshop on 3DVis*. pp. 45–50. IEEE (2014)
59. Thomas, J.J., Cook, K.A.: *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Press (2005)

60. Wang, F., Cao, X., Ren, X., Irani, P.: Detecting and leveraging finger orientation for interaction with direct-touch surfaces. In: Proceedings of the 22nd annual ACM symposium on User interface software and technology. pp. 23–32. ACM (2009)
61. Want, R., Hopper, A., Falcao, V., Gibbons, J.: The active badge location system. *ACM Transactions on Information Systems (TOIS)* 10(1), 91–102 (1992)
62. Weigel, M., Lu, T., Bailly, G., Oulasvirta, A., Majidi, C., Steimle, J.: Iskin: flexible, stretchable and visually customizable on-body touch sensors for mobile computing. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. pp. 2991–3000. ACM (2015)
63. Weiser, M.: Some computer science issues in ubiquitous computing. *Communications of the ACM* 36(7), 75–84 (1993)
64. White, S., Feiner, S.: SiteLens: situated visualization techniques for urban site visits. In: Proceedings of the ACM Conference on Human Factors in Computing Systems. pp. 1117–1120. ACM (2009)
65. Willett, W., Jansen, Y., Dragicevic, P.: Embedded data representations. *IEEE Transactions on Visualization and Computer Graphics* (2017), to appear
66. Wisneski, C., Ishii, H., Dahley, A., Gorbet, M., Brave, S., Ullmer, B., Yarin, P.: Ambient displays: Turning architectural space into an interface between people and digital information. In: International Workshop on Cooperative Buildings. pp. 22–32. Springer (1998)

All links were last followed on October 20, 2016.