

BodyLenses – Embodied Magic Lenses and Personal Territories for Wall Displays

Ulrike Kister, Patrick Reipschläger, Fabrice Matulic, Raimund Dachsel

Interactive Media Lab Dresden
Technische Universität Dresden, Germany
{firstname}.{lastname}@tu-dresden.de



Figure 1. *BodyLenses* are body-controlled magic lenses that support parallel work in multi-user scenarios through territoriality (a), can be of various shapes, e.g. body-centric shapes being modified by arm gestures (b), and provide personal tools and clipboards for local interaction on the display (c).

ABSTRACT

Magic lenses are popular tools to provide locally altered views of visual data. In this paper, we introduce the concept of *BodyLenses*, special kinds of magic lenses for wall displays that are mainly controlled by body interactions. After motivating the rationale for body-centric lenses, we present a comprehensive design space of *BodyLenses*, where we analyse fundamental aspects such as appearance, function, interaction and use in multi-user contexts. Within that space, we investigated and implemented a number of design alternatives and propose solutions for lens positioning, dynamic shape modification, distance-based parameter mappings and the use of *BodyLenses* as portable tool belts. We demonstrate the practicality of our novel concepts with four realised application scenarios. With this work, we hope to lay the foundation for future research and systems based on body-driven lenses.

ACM Classification Keywords

H.5.2. Information Interfaces and Presentation: User Interfaces – Interaction styles

Author Keywords

Body-centric interaction; embodied interaction; territoriality; magic lenses; proxemic interaction; co-located collaboration.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
ITS'15, November 15 - 18, 2015, Funchal, Portugal.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3899-8/15/11...\$15.00
DOI: <http://dx.doi.org/10.1145/2817721.2817726>

INTRODUCTION

Magic lenses [4] are a convenient focus and context technique that present another view onto the data within a local area of interest. They have been used for very diverse tasks in visualisation [36] but are also applicable to many other domains [4].

Large displays are increasingly available and support both the visualisation of large information spaces as well as multi-user scenarios. However, while most magic lens applications in research rely on mouse and keyboard interaction, this is obviously not convenient for display walls and multi-user scenarios. They rather call for more natural user interface techniques. In addition to multi-touch, researchers have analysed mid-air gestures, body movement and body-controlled manipulations [17, 31]. It has been shown that physical navigation, i.e. the actual physical movement of the user in front of a display, improves user performance when navigating information spaces [3]. For collaboration, Tang et al. [34] found that territories on tabletops [29] are transient, a result that Jakobsen and Hornbæk confirmed for interactive high resolution wall displays [18]. Therefore, we believe that a new version of flexible, dynamic territories should be considered.

By creating interactive body-enhanced magic lenses we combine these diverse research topics: We apply the knowledge of body-centric interactions [17, 31], proxemic interactions [14] and the advantages of magic lenses [4, 36] to create *BodyLenses*, flexible work territories with various functions and tools. *BodyLenses* are body-controlled magic lenses that appear when users move in front of a display wall and which follow the users' movements, adapting their properties to the user's motion (see Figure 1b). They can additionally be adjusted through other interaction modalities, e.g. using touch on the display when in close proximity (see Figure 1c).

It is possible to apply our proposed *BodyLenses* in various application scenarios, e.g. for the analysis of bio-medical images, for information visualisation tools, for graph exploration and mind-map creation. In these cases, lens functions can be helpful, e.g. to reduce edge congestion in graphs [35], to apply colour filters to images [12], or to identify and show users' annotations [28]. However, *BodyLenses* are more than magic lenses as users can seamlessly navigate the information space and focus on the data while the tool automatically follows. Moreover, the user can navigate not only along the display, but also towards and back from the wall display. This can be used to transition between different parametrisations of the lens function, e.g. to switch between different levels of contrast when looking for possible tumours in an MRI image. Further, in multi-user scenarios, users are made aware where and what collaborators are currently working on, thanks to their *BodyLenses* (see Figure 1a). Also, the personal properties of the lens allow access to personalised tools or items from individual collections. In collaborative mind-map or brainstorming activities, for example, users can add previously prepared sketches from their lenses.

In this paper, we contribute the concept of *BodyLenses* unifying aspects of magic lenses [4, 36] with territoriality [29], personal tools and proxemics [14]. We introduce the characteristics of *BodyLenses* and define the concept in the next section. This is followed by a thorough analysis of the design space based on prior work as well as our own contributions. We suggest several novel concepts derived from the design space that have not been covered in the literature and present our own practical investigations of the field: We examine different possible *BodyLens* shapes and manipulations of the lens as well as the modification of its parameters using the distance between users and the wall display. We further extend the concept of body-controlled magic lenses by including a tool belt for all kinds of manipulations within the lens and personal clipboard areas. We implemented these novel concepts in four prototype applications that incorporate the discussed scenarios. Finally, we explore various challenges and research question that arise when working with *BodyLenses*.

BODYLENS CONCEPT

We first describe the basics of the *BodyLens* concept by discussing its background in magic lenses and body-controlled interaction. Then, we introduce and define the term *BodyLens* and characterise these body-enhanced lenses.

Background

The concept of magic lenses, as introduced for user interfaces by Bier et al. [4], describes an arbitrarily shaped area where filters can be applied to present another view onto the data of that region of interest. These filters can be any selection or alteration of the data or its view, e.g. a distortion of all or parts of the selection (e.g. [1, 5]) or a revelation or reduction of information (e.g. [22, 35]). This principle has been used for various application contexts, tasks and data types. They have recently been summarised for information visualisation in a comprehensive survey by Tominski et al. [36]. The interactive quality of the lens that allows selecting different parts of the data set, while typically leaving the context visualisation




Figure 2. Conceptual depiction of the *BodyLens* interaction space. In a seamless way, users can move in front of the display, perform mid-air gestures and interact directly on the screen to control the lenses.

unchanged, makes it a very convenient tool for exploration for very diverse application cases.

So far, lenses have been mainly controlled in desktop environments using mouse interaction (e.g. [1, 35]), but they have also been implemented for touch-enabled tabletops (e.g. [19, 21]) or made tangible and spatially-aware [32]. In this paper, we use the *body* as the main input device to control the lens, which is what characterises our *BodyLens* concept. Similar to magic lenses, body-centric interaction has been extensively studied in the literature. Embodied interaction for wall displays is the focus of design space analyses conducted by Shoemaker et al. [31] and Müller et al. [24], but those research efforts do not address the specific case of lenses. Using the body's shadow as a magic lens has been proposed very shortly by Shoemaker et al. in their work on shadow reaching [30]. However, we call this idea into being, considerably extend it and generalise the concept. We combine principles of territoriality [29] and magic lenses [4, 36] to form a new concept using both (mid-air) movement of individual body parts as well as proxemic dimensions [14], such as distance and movement. Additional interaction modalities, e.g. touch and pen, are used to manipulate the lens' parameters. While these individual components are known, we unify them and, with *BodyLenses*, provide a coherent local tool for the flexible exploration and manipulation of information spaces, thereby also creating adaptable personal territories for users.

Definition and Characteristics

We define the term *BodyLens* as any type of personal, body-controlled magic lens used on a vertical display. Thus, the *magic lens* can apply any function to a user-centred region of interest, including selection, zooming, focus+context, advanced visual modifications, or semantic changes to the data. "*Body-controlled*" refers to a) any type of body movement with regard to the fixed display as well as b) gestures of any body part – most prominently arms and hands, but also head, legs and feet – and any combinations thereof. Furthermore, provided the display is an interactive surface, c) direct interaction *on* the display is also possible, e.g. by means of multi-touch, pen or tangible input. While horizontal movements along the wall display typically influence the position of the lens, distal movements, body gestures and direct interactions may influence any other parameter associated with the lens.

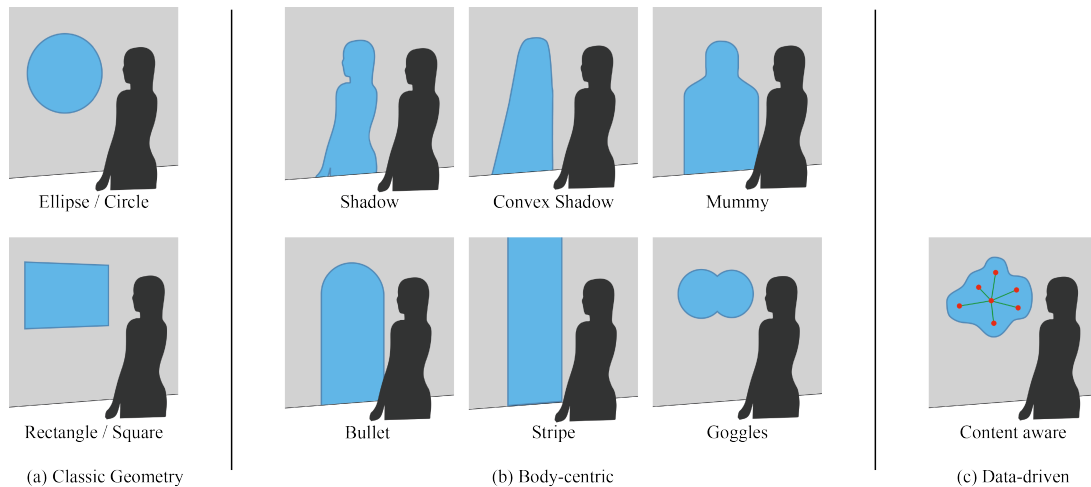


Figure 3. Three categories of *BodyLens* shapes with typical and promising examples.

We believe that the combination of the human body and a magic lens is more than a simple mapping of novel means of interaction to well-known functions. A *BodyLens* is indeed not just a filter, but also a personal tool and a work territory. *BodyLenses* support effortless, *implicit* navigation within an information space as they automatically move with the user. They can also be adjusted *explicitly* using either direct manipulation, e.g. touch on the display, or mid-air gestures. In multi-user contexts, the lens function alters a local view into the data space, not disturbing other users. Additionally, the lens presents a personal territory that supports mutual awareness and fosters collaboration. Through personalisation we also support individual tools along the lens, private annotations and ownership of associated elements.

BODYLENS – DESIGN DIMENSIONS

Both magic lenses and body-centric interaction have been extensively studied in the literature. The design space we propose here extends that prior work and establishes a general frame of reference for body-centric interactive lenses in a variety of contexts. We divide our exploration of the space into four main dimensions: *appearance*, *function*, *interaction* and *multi-user contexts*. They are described with their properties in the following sections using examples from prior work when relevant aspects are covered.

Appearance

The appearance of a *BodyLens* refers to its characteristic visual features, specifically position on the display, shape and rendering effects. The most common representations of magic lenses are detailed in [36].

Position

The choice of the position of the lens on the display with respect to the user’s body, i.e. the anchor point of the lens, is an important design aspect. Perhaps the most obvious materialisation of a *BodyLens* is a circle or a rectangle roughly centred at eye-level of the user (Figure 3a). If such a lens is mapped to the user’s physical movement in front of the screen, it can only move horizontally, since the user’s head remains at the same height. Therefore, additional means have to be devised to also allow the lens to move vertically and to enable 2D

positioning in the whole screen space. This problem can be avoided at the cost of additional space monopolised by the lens, if the latter is made to span the entire height of the display, e.g. by using vertical stripes (Figure 1a). Such kind of lenses might be particularly appropriate in multi-user scenarios, where each person working on the display occupies a dedicated personal territory delimited by such a band.

Shape

The shape of magic lenses is typically round or rectangular, but they can also have arbitrary silhouettes either explicitly determined by the user (e.g. with a sketch) [5, 20] or by the underlying data [16, 26]. Those types of shapes are of course also applicable to *BodyLenses*, where the contours can be dynamically modified using body gestures (see *Interaction*).

Beyond the classic examples, other forms more specific to body-centric lenses are conceivable. A *BodyLens* can be considered as a kind of augmented shadow, with the lens contour more or less matching that of a virtual shadow cast by the body on the display (e.g. [30]). “True shadow” *BodyLenses*, i.e. lenses, whose outline exactly mirrors the contour of the user’s body, presumably find more relevance in artistic installations or games such as [6, 9]. *BodyLenses* as virtual user embodiments might also be useful for collaborative telepresence [40] and to represent personal spaces [11]. For co-located computing tasks, however, in particular information visualisation applications, the projected shape of the human body may not constitute an adequate view space for magnified or filtered data. Thus, for those contexts, body-based lens shapes are likely to take more organic or geometric forms, e.g. vertical stripes or bands, whose width roughly corresponds to that of the body (as in [8]), to vertically elongated ellipses or other oblong shapes.

In Figure 3, we propose a number of typical *BodyLens* shapes that designers may choose to use in their applications, ranging from conventional formats such as ellipses and rectangles (a) to more body-centric silhouettes such as convex shadows, mummies, stripes etc. (b), without forgetting data-driven, content-aware shapes (c). We elaborate on practical aspects of some of those proposed shapes further below.

Rendering Effects

Various rendering effects can be applied to the inner area as well as to the borders of lenses. Borders, for instance, can be fully rendered (thus mimicking real lenses) or they can be implicit, i.e. the lens area is determined and visualised only by the filtered content within or outside its bounds. Smooth transitions between focus and context can also be introduced, similar to soft vs hard shadows and Sigma Lenses [25], which use time and translucence to smooth transitions between the filtered content and the surrounding non-modified context. Similarly, the area within the lens can exhibit various degrees of opacity and filling patterns, e.g. a translucent image of the user, depending on whether the lens should strictly display processed content or also fulfil other purposes (such as providing visual feedback for awareness).

3D BodyLenses

While this paper mainly focuses on *BodyLenses* in 2D spaces, the concept can readily be extended to 3D. In virtual 3D worlds and augmented reality scenarios, the user's body can conceivably act as a 3D *BodyLens* that selectively filters volumetric information in the target space. In such contexts, the standard shapes of 2D lenses, i.e., ellipses and rectangles, respectively become ovoids and cuboids that "cut through" chunks of volume data (3D magic lenses [37]). Similar to vertical stripes that stretch from the bottom to the top of the display, 3D stripes or cuboids mapped to the user's physical position can extend completely across one or more dimensions (that is, not only height, but also depth) to cover infinite 3D slices, and hence reduce the need for physical movement or additional interactions to reach distant areas within the 3D space. Closer to embodied representations still, one can also conceive of a 3D avatar-lens of the user (the 3D version of a shadow) that performs truly body-based 3D-filtering operations through the volumetric information space.

Function

The function of the *BodyLens* is the operation that is performed on the selected data. Most lenses are used for visualisation purposes [36], the most prominent example being magnification (with a variety of scaling styles) emulating a physical lens. Zooming lenses allow finer details of objects within the magnified region to be observed, but they can also allow more precise interaction with the content within that zone [1, 19, 27]. Such kinds of scenarios are well suited for *BodyLenses*, since the body naturally provides an interaction area within which the user can perform fine-grained (multi-)touch and/or pen operations on the screen. In case the *BodyLens* is entirely controlled by body (torso) or head movements, the user's hands and fingers can fully devote themselves to the main task, as they are not required for (re-)positioning the lens. An immediate extension of that concept for lenses, whose source data and visualisation view are decoupled, is reaching assistance, that is local lenses used as proxies to interact with physically distant or unreachable elements [13].

Further popular uses of visualisation lenses are to provide different views of the underlying data, i.e. filters [33], and to isolate desired items inside the lens (e.g. node links in

graphs [21]). Most lenses modify the elements inside them, but they can also be used as samplers to modify the surrounding content [12]. Again, *BodyLenses* can be directly utilised in those cases, e.g. as "magic shadows" that offer alternative visualisations of data within their body-based bounds [30] or to uncover portions of a hidden image overlapping with the user's body passing in front of the display [8, 9].

Lenses often provide tools to modify their functional parameters (e.g. zoom factor) or to manipulate their associated content. In this sense, interactive lenses also function as convenient tool palettes. The individual tools can be placed on the lens itself as semi-transparent overlays (as in the initial toolglass interface [4] and in [11, 31]) or attached to its edges [21]. Thus, a *BodyLens* can have a tool belt that hugs its contour, so that the tools remain constantly at the user's fingertips as they follow him/her around when s/he moves to different areas of the workspace (Figure 6).

Interaction

Interaction here deals with manipulation of *BodyLenses* themselves and their generic properties and not with the underlying data, which is task-specific. The design space of the former is however obviously influenced by the gestural vocabulary of the latter, as the two have to form a coherent and conflict-free whole in the final application. In principle, we can differentiate between fine-grained, finger-based interaction *on* the screen (such as multitouch or pen input) and coarse input through whole-body movements and arm or hand gestures *in front* of the display, with possible transitions between both types of input.

Position control

As introduced above, the standard mapping of body motion to a lens is to couple the user's physical position in front of the display to the location of the lens on the screen. In such a configuration, the *BodyLens* follows the user like a shadow as s/he moves in front of the screen. There are several ways in which that mapping can be realised. Depending on the shape of the lens and the application requirements, a simple unidimensional mapping along the x-axis (Figure 2) suffices [8, 9, 38]. If the lens needs to be steerable over the entire 2D coordinate space of the screen, movement along the y-axis also needs to be supported. One way to achieve this that we propose is to map the user's distance from the display to vertical movements (Figure 4a). Other approaches include using gaze direction [2], head tilt or arm gestures (see below), depending on the availability of those interactions for lens operation. Further, the moving element might not be the lens, but the content underneath, e.g. a "cockpit view" with a fixed lens where the background changes according to user movement. In further applications, both the location of the lens and of the data on the display might be simultaneously modifiable, for instance, when the lens view is decoupled from its source data.

Shape control

While position is likely the most obvious parameter of a *BodyLens* that the user's body can control, embodied interactions may also govern other attributes of the lens, including and especially its shape. In Figure 3, we presented a

number of examples of typical *BodyLens* shapes. Some of those shapes, such as the shadows, continuously change by nature. The other more geometric forms can also be modified by appropriate body gestures. For example, one of the interactions that we propose for stripe-shaped lenses is to allow their width to be increased by stretching the arms sideways, thereby expanding the area covered by their function. For more organic shapes, the arms can be used to punch bulges in the lens contour to stretch it in desired directions, similar to sculpting gestures used in *BodyAvatar* to grow virtual limbs [39]. To ensure that the lens size does not constantly change with any arm movement, an explicit hand pose can be required for shape-modifying interactions. Our contribution includes implementations of several dynamic shape-modifying techniques, which we describe further below.

Distance mapping and proxemics

The ability of users to move freely in the physical space in front of the wall display to control a *BodyLens* is at the heart of its interaction paradigm. We have seen that such control can be based on a simple unidimensional mapping along the x-axis to allow the lens to follow the user like a shadow, or that it can also involve user movement towards and away from the display, i.e. along the z-axis (see Figure 2), to allow it to move vertically on the screen. Depth movement can also be used to modify other lens parameters, especially that of the filter function via which the target data is viewed. While not explicitly making use of magic lenses, there are examples in the literature, in which user-to-display distance has been used to modify zoom [15], detail [10] and abstraction [22] levels of visual data. Depending on the function and the parameter to modulate, the distance-to-variable mapping can be continuous or discrete. While the former associates a variable with continuous values to the distance between the user and the screen (e.g. zoom factor), the latter often implies the partition of the physical space in front of the display into several zones corresponding to different discrete states of the variable (e.g. layer type). We discuss novel distance-based functional mappings for *BodyLenses* in more detail in the next section.

Distance-based *BodyLens* interactions can also be designed to mirror interpersonal engagement, based on familiar spatial relationships governing social interactions, a principle known as proxemics [14]. Thus, following the observation that approaching a person or a device signifies an increasing interest in interacting with that entity, *BodyLenses* can be made to progressively react to users nearing the display following a gradual engagement design pattern [23]. Perhaps the most straightforward realisation of that pattern is a progressive fade-in of the lens as the user comes closer to the screen. When the lens is completely materialised, further distance-based interactions controlling a functional parameter can be performed, as described above. If the display detects contact input as well, a further “at-the-wall” stage might exist, where users can directly operate lenses and content on the screen using touch, pen or tangible interactions.

Multi-User Contexts

Another important dimension to consider is the number of lenses and users interacting simultaneously. We can differ-

entiate between single-user and multi-user scenarios which include parallel individual and collaborative work. For instance, with multiple active lenses, users have the possibility to overlap them to combine their effects. Examples using classic lenses include operation composition in collaborative visual queries [20] and, in the case of personal lenses with user-specific IDs, sharing of private spaces [28, 31]. Adapted to *BodyLenses*, those applications could rely on multi-user embodied interactions to collaboratively create compound lens functions and grouped spaces. Thus, *BodyLenses* that completely or partially merge through concerted user actions can form common embodied territories with shared properties and elements, e.g. a coalesced convex shadow lens within which users work together on a common task. Thanks to the very dynamic nature of *BodyLenses*, ad hoc groups with temporarily joint territories can be quickly formed and split as users move towards and away from each other. This scenario can be viewed as an application of F-formations [23] and more generally of proxemics [14].

Multiple non-overlapping lenses (and therefore presumably also *BodyLenses*) have notable benefits in InfoVis situations where visual datasets need to be viewed side by side and compared, with the help of dynamically varying filter functions. An example of such parallel body-based data view control (albeit not through lenses per se) is presented in [10]. Extended to flexible and versatile *BodyLenses*, this paradigm allows groups of people to collaboratively examine multiple arbitrary facets of a common dataset through vignettes that can be freely moved, scaled and adapted using natural proxemic and body-driven interactions (Figure 4).

A further advantage of *BodyLenses* in multi-user contexts is their ability to provide mutual awareness. Much like digital shadows, which are known to possess that property [30], *BodyLenses* that follow the user continuously indicate his/her current locus of attention (i.e. “I am currently working on the content in that lens”). Those moving lenses are noticed by the other collaborators, who can then adapt their behaviour accordingly. As with shadows and visual feedback reflecting users’ motions and postures, we imagine that the awareness benefits of *BodyLenses* particularly manifest themselves in co-located situations, in which users interact in a contactless manner with the display. This is likely true to an even greater extent in remote collaboration conditions, where the lack of physical co-presence of other participants needs to be further compensated (similar to what virtual embodiments achieve for remote collaborative whiteboard applications [40]). In such settings, the lens would therefore function both as a tool to interact with the data as well as a mechanism to maintain mutual awareness between workers (co-located or remote).

Finally, we mention a particular category of multi-user scenarios in which the user’s body plays a role not only as input to control a *BodyLens*, but also as a physical obstacle in front of the display. Those cases occur particularly in situations, when privacy needs to be safeguarded in (semi-)public environments. Specifically, the body blocks access to other people in the vicinity, while the lens creates a confined view of the private content that can only be visualised by the owner [7].

NOVEL CONCEPTS AND PRACTICAL INVESTIGATIONS

Based on our design space already including several suggestions that advance prior work, we implemented a number of application prototypes. They serve as examples for the feasibility of our concepts and demonstrate the potential of *BodyLenses*. Several HCI experts regularly provided feedback during our iterative development process, leading to a number of novel concepts which, to the best of our knowledge, have not been reported before. Those concepts are introduced below along with a more in-depth exploration of the dimensions of *BodyLenses* and solutions for some of the practical challenges that we faced.

We address adjustment of dynamic shapes to the user's position and pose as well as the mapping of their distance to the wall to various properties and parameters of the lens, as we see these aspects as the main differentiators between magic lenses and *BodyLenses*. We also explore the role of the lens as a personal territory and tool belt with menus and clipboards. We end this section by discussing open research questions.

Applications and Implementation

We implemented four applications derived from the scenarios presented in the introduction, which we then use for our practical investigations. As magic lenses have been proved valuable tools for information visualisation, we implemented 1) a graph explorer with various lens functions derived from this background, e.g. *Local Edge*, *Bring Neighbours* and *Fish-eye* [35]. For multi-user brainstorming, we further developed 2) a mind mapping tool that allows creation and manipulation of items using personal lens territories. Further, 3) an image analysis tool allowed us to investigate the application of body-enhanced lenses with different colour filters and optical zoom lenses. Finally, we also implemented 4) a small artistic application that explores the casual exploration of proxemic dimensions and gestural interaction to create a dynamic piece of art. Our prototypes show the variety of possible usage scenarios for *BodyLenses* and feature various shapes from our proposed design space. We will further elaborate the important aspects and challenges in the next sections.

Similar to previous work [22, 23, 24], we divide the area in front of the wall into multiple zones (Figure 2). In our particular setting, we label the area closer than 0.5 m to the wall as "touch-interaction zone" (A). The area from 0.5 m to roughly 2.5 m is the main zone for body-centric interaction ("body-interaction zone", B). Beyond 2.5 m the user is considered outside of the interaction zone (C) and no *BodyLens* is created or the existing lens fades out.

The prototypes were developed in Python using libavg¹ framework for visualisation. Our setup consists of an interactive display wall of 4.86 m in width and 2.06 m in height (frame at 2.3 m), with a resolution of 7680 × 3240 pixels. The movements and body of each individual user is tracked using a Kinect One² fixed to the ceiling at the back, i.e. behind the users facing the wall.

¹<http://www.libavg.de>

²<http://www.kinect.com>

Dynamic Shapes

In our design space, we identified several different shapes for the lenses (Figure 3), divided into three categories. We implemented a selection of them for our prototypes which enables us to investigate multiple levels of embodiment (for some results, see Figure 4). Specifically, we created the elliptical and rectangular shapes (classic geometries), content aware shapes (data-driven), as well as the stripe, bullet, convex shadow and mummy shapes (body-centric). These lenses were used to various degrees in the prototypes. All shapes were tested with the graph explorer while the artistic application was designed for the stripe shape.

Body-driven lens movement

When users step in front of the wall, a new *BodyLens* is created in correlation to the user's position, which follows whenever users move. As previously discussed, our proposed default mapping positions the lens to encompass the area in front of the user at eye level. The lens movement is closely coupled to the central body motion of the user (i.e. the torso) on all axes. For geometric and content-aware shapes, the lens is centred at the user's shoulder position, while the body-centric shapes present a widened shadow in front of the user. When in the "touch-interaction zone", users can drag the lens on the screen, which creates an offset relative to their current position. This offset is maintained when users move, thus allowing them to easily reach top or bottom areas of the display across its entire length. When the distance to the wall is too large for touch interaction, one of the solutions we propose to control the vertical position of the lens is to use up-and-down hand movements. When users are within the "touch-interaction zone", the lens is frozen and its position no longer coupled with their movement. This enables working on details of the data and avoids jitter. Users may still drag the lens while it is frozen.

Body-driven shape adjustments

Depending on the type of shape that is used and the current goal of the user, there are several ways in which the shape of the lens may react to the user's movements. The shape of the lens can be adjusted **implicitly** by using the user's current pose. If, for example, the user raises a hand to access some data, the shape may bulge outward at that location, mimicking the user's movement. The lens may also react in a more subtle way as the relation between body parts changes during movement. However, in most scenarios the user focuses on the underlying data, not on the lens itself. Therefore, the closely coupled shapes might not be suitable. We strongly reduced lens shape movements in our prototypes by not considering the very active body parts, e.g. hands or fingers, but the more body-centric movements (the torso). This further supports the active movement of hands for touch interaction.

Users can also **explicitly** change and adapt the lens shape to include or exclude data. For this, we propose a grabbing gesture, which consists of raising the arms and closing the hands in mid-air. To provide feedback, the shape becomes highlighted and can then be modified as the shape's border is "attached" to the user's fist. For the stripe shape, only horizontal movement is possible, which widens or shrinks the stripe.




Figure 4. User’s *distance* to the display changes a) vertical position of the lens, b) size of the lens, c) function parameter (here: zoom), or d) visualisation parameters (here: selected time steps).

All implemented classic geometric shapes allow adjustment in both horizontal and vertical directions. Embodied shapes can similarly be altered by bulging the contour when grabbing their border and moving the hands. Clutching is possible when the user re-opens, moves and re-forms the fist.

One of the main challenges while implementing this functionality was the placement of the Kinect to guarantee robust tracking of all users. As the recognition is best when the Kinect is placed in front of the user, integrating it directly into the wall (as in [9]) would give the best results. This, however, was not possible in our setup. Placing the Kinect above or below the wall results in unreliable recognition because of the steep angle. Thus, we decided to place the device behind the users, who face the wall. This results in reliable tracking, but makes the detection of certain gestures such as closed fists more difficult.

Mappings of User-to-Wall Distance and Movement

While movement along the wall is tightly coupled with the horizontal position of the lens, the distance can be mapped more freely and in interesting ways. In the following, we propose a number of promising mappings, that we investigated and implemented within our application scenarios.

Mapping distance to lens features

From the design space, we derived four categories of lens control based on distance-mapping: mapping to a) lens position, b) lens properties, c) lens function parameters, and d) visualisation parameters. The suitability of those mappings depends on the application context and task as well as user preferences.

The user’s distance to the display can be used to change the **vertical position** of the lens, moving it up and down by stepping forward or backward (see Figure 4a). This way, all areas of the display can contain data and can still be explored by the user through the lens. However, the mapping should always ensure that the lens is at eye-level when moving at arm-distance to the wall, so that manipulations using touch are still possible. In our implementation, the user can move the lens up and down when far from the wall. When stepping closer to the wall, the *BodyLens* returns to eye-level.

Different **lens properties** can be controlled by changing the distance to the wall. We define a lens property as any parameter of the lens that is independent from its currently selected function or content. One example is the mapping of

the distance to the *size* of the lens, i.e. the width and height of a geometric shape, the width of a stripe, or a factor to enlarge a body’s mapped shadow. However, different mappings and factors to manipulate size are possible: 1) When stepping back from the wall, the lens becomes larger, thus increasing the spatial coverage of the lens function. This type of mapping coincides with the notion of considering the lens as a shadow with the light source behind the user [31]. 2) Alternatively, we can consider the lens as a tool for the data in front of the user. Here, users standing in front occlude, and thereby occupy, more data than users in the back. Hence, when stepping back, the lens area becomes smaller (Figure 4b). This also coincides with the fact that by stepping back we assume that the user’s intention is to get an overview of the data, focusing more on the context, not the lens and its focus area. As with the mapping type, however, its direction may be very dependent on task, context and user preferences.

The lens function itself may also have certain parameters, which can be adjusted, e.g. the zoom level of a lens (Figure 4c). It is possible to map these **parameters of the lens function** to the distance from the wall. A useful application for this are for example lenses for graph visualisations as mentioned above, or a semantic zooming lens: By stepping closer to the wall, nodes expand and the user can see more details. In our small artistic installation, we use lenses to paint and manipulate colour stripes on the screen. Here, the distance changes the colour of the user’s stripe. Based on our implementation, we suggest that when the lens is used not only for exploring but also to manipulate content, it is essential to choose a distance at which the user can simply look at the result without changing it. For the art installation, we additionally use a dwell time before any stripes are painted onto the display.

Finally, general **visualisation parameters** of the data can be influenced by the user’s distance to the wall. One example is the mapping of the third dimension of the information space to the distance, so that movement from and towards the display actually maps to depth movement in the *3D virtual space*. This way, we propose the local exploration of arbitrary slices of volumetric or spatial data by using *BodyLenses*. Another important dimension in today’s big datasets is the special property of *time*. Data such as the measurements of traffic behaviour or the development of a tumour are often analysed with respect to time. Comparing different time steps with each other and slicing through these time-dependent data sets

are essential tasks during analysis. In our image analysis prototype, we map the user's distance from the wall to the time dimension. To the best of our knowledge, we are the first to propose this mapping of body movement to time. It allows the user to individually select a desired time frame, which can often be changed only globally. This further supports parallel work on the whole data set (Figure 4d).

Selecting the direction of the mapping

In our observations, we found that even for the simple case of zooming, some users prefer having the detailed view close to the wall and the less detailed view from afar, supporting the assumption that stepping back equals focusing on the context, while others preferred seeing the details from afar, helping them to better compare them to the context data. Hence, it is not only the application designer's selection that should be considered, but also the personal embodiment character of the lens and therefore users' individual preferences. This can be specifically supported when including user identification, e.g. by using the Kinect, or other biometric approaches, such as the hand contour [28].

Comparing different slices of the visualisation

To compare different slices of the data, a user will generally move back and forth within the visualisation dimension. For example, when time is mapped to the distance to the wall, the user can compare different time cuts by moving towards the display. To see multiple slices of the data at the same time, users can work cooperatively, with each user selecting a value of the dataset to be compared, as implemented in our prototypes. Conceptually, users can also freeze their lens and create a second one to compare two different slices of the visualisation.

Compensating visual shrinking at larger distances

Previously, we discussed using the distance to the wall to change the function, e.g. to create other views onto the data, and imply the user's intentions, e.g. to see details when stepping closer. However, when stepping back, e.g. slicing through time, elements seem smaller from further away. To be able to still properly compare these layers, it is important that users see the same amount of detail and content that they saw before from close by. Hence, we contribute the concept of automatically enlarging the *BodyLens* and the elements within to compensate for the user's movement. Thus the user *perceives* the content to be always the same size (see Figure 5) regardless of the distance, which is used to control another parameter like time.

Reachability of all areas of the wall

A problem that arises when working with body-controlled lenses on and in front of a display wall is the reachability of certain areas of the wall. This especially concerns the upper areas, which might not be reachable because of the user's size, and the lower areas, as kneeling down might be uncomfortable or inappropriate. For example, the top of our wall display is at roughly 2.3 m, which can be unreachable even for people of average height. This reachability problem arises for any shape that does not cover the complete height of the wall, i.e. all but the stripe shape in our selection. We implemented four alternatives in our prototypes, that we found promising




Figure 5. The lens is compensating the distance of the user by enlarging its content so that it appears to be the same size.

in addressing this problem: It is possible to 1) use the hand motion to move the lens up and down. The user can also 2) resize the lens to reach the upper or lower part of the display. Alternatively, it is possible to 3) set an offset by dragging the lens using touch, which is maintained upon moving in front of the display. Finally, the height can be influenced by 4) the user's distance from the wall. We believe the whole subject of reachability for wall displays is an interesting question for future research.

BodyLens as a Personal Tool Belt and Clipboard

The *BodyLens* is not only a tool in itself, it can also provide additional tools depending on the current task and context. For network manipulation, tools for node or edge creation might be necessary. Similarly, our mind-map application requires tools for creating and manipulating items (see Figure 1c). We suggest **personal tool belts** as a means to arrange and access these menus and tools attached to a lens. Thereby the lens becomes a user's personal territory for work on the mind-map. In accordance to the *BodyLens* concept, tool belts will move implicitly with the user. In previous work menu icons or applications were placed within a user's shadow [11, 31]. As we apply this concept to *BodyLenses*, the content within that shadow is very important. Hence, in our implementation, we found placing menus and tools at the border to be beneficial as it enables the content of the lens to be entirely visible (Figure 6a). The selection of required tools is dependent on the role, current task or general preference of the user. If the system supports user identification, content created within a lens can be associated with a specific user. This means not only the tools but also the content can be personalised and highlighted per user, which we also investigated in our prototype. This offers further possibilities, such as users locking their content to prevent changes from other users or removing all their personal content from the display if they decide to no longer participate in the collaboration.

BodyLenses can also function as individual **storage containers** with access to personal data, which can be accessed on the tool belt of the personal lens. Within our mind-map application, images or previously prepared sketches can be added to the mind-map through the lens. To categorise, sort and manage these items, a part of the lens can function as a clipboard to which items can be attached (Figure 6b). When multiple people want to work together and exchange items, individual clipboards can be shared.

Since the shape of a *BodyLens* can be very dynamic, the placement of elements like menus and clipboards is not easy. We therefore suggest that for each shape certain parts of the



Figure 6. a) The lens' tool belt is placed along the border, so that the result of the LocalEdge function [35] is visible; b) The user can add posts from the mind-map to the clipboard.

border need to be defined as paths where the tools may be placed. In our prototype, the circular lens defines a circular path for the clipboard and menu, while the rectangular lens defines straight paths (Figure 6). However, the design has to take into account the probable position of the user's arms. When positioning the menu at the top, the user might occlude the data within the lens when switching functions. It is our goal to conduct further research regarding this placement along dynamic body-centric lenses, since early experiments within our prototypes have revealed this to be a challenging topic.

Open Research Questions

With the presentation of the design space, concepts and initial prototypical implementations of *BodyLenses*, we have laid the foundations of a research framework for embodied magic lenses. From this foundational work a number of research questions emerge that require further investigation. Perhaps most fundamental to the validation of the concepts underlying *BodyLenses* are the application contexts within which those lenses should be applied to efficiently support the tasks at hand. A particular challenge is to design efficient and intuitive lens interactions that work as intended by the user by a) avoiding always-on problems and involuntarily activating functions, b) allowing to explicitly trigger desired lens changes (e.g. by raising the arm or performing mid-air gestures) and most importantly c) that do not interfere with the primary data exploration and manipulation tasks.

There is considerable freedom within the design space and it remains to be investigated what lens shapes, attributes, interaction techniques and embodied tools are most suitable for each considered scenario. Specific solutions have to be investigated for the placement of elements along body-centric lenses, to conveniently reach all areas of the display (particular top and bottom areas) and for the freezing functionality which allows to temporarily fix lenses, their properties and their parameters. In terms of experimental validation, appropriate evaluations will have to confirm the advantages of *BodyLenses* compared to traditional magic lenses. The intricate case of multi-user environments, with issues such as territoriality, awareness, group dynamics, privacy and further social factors within both co-located and remote collaborative work, will also have to be tackled specifically. Related to that, technical solutions have to be developed for reliably tracking occluded users in crowded situations and for user identification to allow true personalisation.

CONCLUSION

We presented *BodyLenses*, a novel concept combining the principles of magic lenses with the advantages of body-centric and gestural interaction. Our body-enhanced lenses support implicit movement together with the user in front of the display. They apply the rich set of existing functions of magic lenses in very diverse application scenarios. Thanks to their embodiment capabilities, *BodyLenses* incorporate transient personal work territories that provide awareness of the user's position and work content.

In this paper, we contributed a thorough analysis of the design space dimensions of *BodyLenses* including related work and our novel concepts. We focused on new concepts for lens positioning, dynamic shape adjustment and possible mappings of the user's distance to various lens properties as well as personalisation aspects. We presented multiple prototypes that strengthen and confirm the potential of *BodyLenses* including a graph explorer, an image analysis tool, a mind-map creator, and an artistic application. Further, we discussed challenges and important issues to address when designing for body-enhanced lenses on interactive wall displays.

ACKNOWLEDGMENTS

This project was in part funded by the DFG (DA 1319/3-1) and the Excellence Initiative of the German Federal and State Governments (Institutional Strategy, measure "support the best").

REFERENCES

1. C. Appert, O. Chapuis, and E. Pietriga. 2010. High-precision Magnification Lenses. In *Proc. CHI '10*. ACM, 273–282.
2. M. Ashmore, A. T. Duchowski, and G. Shoemaker. 2005. Efficient Eye Pointing with a Fisheye Lens. In *Proc. GI '05*. 203–210.
3. R. Ball, C. North, and D. A. Bowman. 2007. Move to Improve: Promoting Physical Navigation to Increase User Performance with Large Displays. In *Proc. CHI '07*. ACM, 191–200.
4. E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. 1993. Toolglass and magic lenses: the see-through interface. In *Proc. SIGGRAPH '93*. ACM, 73–80.
5. J. Brosz, M. A. Nacenta, R. Pusch, S. Carpendale, and C. Hurter. 2013. Transmogrification: Causal Manipulation of Visualizations. In *Proc. UIST '13*. ACM, 97–106.
6. Brother System. 2013. Vitrine Interactive "Les Rives de l'Orne". <https://youtu.be/KLOB-T1mgdY>. (2013).
7. F. Brudy, D. Ledo, S. Greenberg, and A. Butz. 2014. Is Anyone Looking? Mitigating Shoulder Surfing on Public Displays Through Awareness and Protection. In *Proc. PerDis '14*. ACM, 1:1–1:6.
8. BUT Dubai and Firmavision. 2011. Interactive Video Wall. <https://youtu.be/b2kqRv1WdEM>. (2011).
9. B. Chung and O. Li. 2013. Domain Mall Interactive Wall Display in Hong Kong. <https://youtu.be/NRhDpDxTsLA>. (2013).

10. J. Dostal, U. Hinrichs, P. O. Kristensson, and A. Quigley. 2014. SpiderEyes: Designing Attention- and Proximity-aware Collaborative Interfaces for Wall-sized Displays. In *Proc. IUI '14*. ACM, 143–152.
11. I. Elhart, F. Scacchi, E. Niforatos, and M. Langheinrich. 2015. ShadowTouch: A Multi-user Application Selection Interface for Interactive Public Displays. In *Proc. PerDis '15*. ACM, 209–216.
12. N. Elmqvist, P. Dragicevic, and J.-D. Fekete. 2011. Color Lens: Adaptive Color Scale Optimization for Visual Exploration. *IEEE Trans. on Vis. and Comp. Graphics* 17, 6 (June 2011), 795–807.
13. J. Gilliot, G. Casiez, and N. Roussel. 2014. Direct and Indirect Multi-touch Interaction on a Wall Display. In *Proceedings of the 26th Conference on L'Interaction Homme-Machine IHM '14*. ACM, 147–152.
14. S. Greenberg, N. Marquardt, T. Ballendat, R. Diaz-Marino, and M. Wang. 2011. Proxemic Interactions: The New Ubicomp? *interactions* 18, 1 (Jan. 2011), 42–50.
15. C. Harrison and A. K. Dey. 2008. Lean and Zoom: Proximity-aware User Interface and Content Magnification. In *Proc. CHI '08*. ACM, 507–510.
16. A. Ion, Y.-L. B. Chang, M. Haller, M. Hancock, and S. D. Scott. 2013. Canyon: Providing Location Awareness of Multiple Moving Objects in a Detail View on Large Displays. In *Proc. CHI '13*. ACM, 3149–3158.
17. M. R. Jakobsen, Y. S. Haile, S. Knudsen, and K. Hornbaek. 2013. Information Visualization and Proxemics: Design Opportunities and Empirical Findings. *IEEE Trans. on Vis. and Comp. Graphics* 19, 12 (2013), 2386–2395.
18. M. R. Jakobsen and K. Hornbæk. 2014. Up Close and Personal: Collaborative Work on a High-resolution Multitouch Wall Display. *ACM Trans. Comput.-Hum. Interact.* 21, 2 (Feb. 2014), 11:1–11:34.
19. D. P. Käser, M. Agrawala, and M. Pauly. 2011. FingerGlass: efficient multiscale interaction on multitouch screens. In *Proc. CHI '11*. ACM, 1601–1610.
20. K. Kim and N. Elmqvist. 2012. Embodied Lenses for Collaborative Visual Queries on Tabletop Displays. *Information Visualization* 11, 4 (Oct. 2012), 319–338.
21. U. Kister, P. Reipschläger, and R. Dachsel. 2014. Multi-Touch Manipulation of Magic Lenses for Information Visualization. In *Proc. ITS '14*. ACM, 431–434.
22. A. Lehmann, H. Schumann, O. Staadt, and C. Tominski. 2011. Physical Navigation to Support Graph Exploration on a Large High-Resolution Display. In *Advances in Visual Computing*. Lecture Notes in Computer Science, Vol. 6938. Springer Berlin Heidelberg, 496–507.
23. N. Marquardt, T. Ballendat, S. Boring, S. Greenberg, and K. Hinckley. 2012. Gradual Engagement: Facilitating Information Exchange Between Digital Devices As a Function of Proximity. In *Proc. ITS '12*. ACM, 31–40.
24. J. Müller, F. Alt, D. Michelis, and A. Schmidt. 2010. Requirements and Design Space for Interactive Public Displays. In *Proc. MM '10*. ACM, 1285–1294.
25. E. Pietriga and C. Appert. 2008. Sigma lenses: focus-context transitions combining space, time and translucence. In *Proc. CHI '08*. ACM, 1343–1352.
26. C. Pindat, E. Pietriga, O. Chapuis, and C. Puech. 2012. JellyLens: content-aware adaptive lenses. In *Proc. UIST '12*. ACM, 261–270.
27. G. Ramos, A. Cockburn, R. Balakrishnan, and M. Beaudouin-Lafon. 2007. Pointing Lenses: Facilitating Stylus Input Through Visual-and Motor-space Magnification. In *Proc. CHI '07*. ACM, 757–766.
28. D. Schmidt, M. K. Chong, and H. Gellersen. 2010. IdLenses: Dynamic Personal Areas on Shared Surfaces. In *Proc. ITS '10*. ACM, 131–134.
29. S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. 2004. Territoriality in Collaborative Tabletop Workspaces. In *Proc. CSCW '04*. ACM, 294–303.
30. G. Shoemaker, A. Tang, and K. S. Booth. 2007. Shadow Reaching: A New Perspective on Interaction for Large Displays. In *Proc. UIST '07*. ACM, 53–56.
31. G. Shoemaker, T. Tsukitani, Y. Kitamura, and K. S. Booth. 2010. Body-centric Interaction Techniques for Very Large Wall Displays. In *Proc. NordiCHI '10*. ACM, 463–472.
32. M. Spindler, C. Tominski, H. Schumann, and R. Dachsel. 2010. Tangible Views for Information Visualization. In *Proc. ITS '10*. ACM, 157–166.
33. M. C. Stone, K. Fishkin, and E. A. Bier. 1994. The movable filter as a user interface tool. In *Proc. CHI '94*. ACM, 306–312.
34. A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. 2006. Collaborative Coupling over Tabletop Displays. In *Proc. CHI '06*. ACM, 1181–1190.
35. C. Tominski, J. Abello, F. van Ham, and H. Schumann. 2006. Fisheye Tree Views and Lenses for Graph Visualization. In *Proc. IV '06*. IEEE Computer Society, 17–24.
36. C. Tominski, S. Gladisch, U. Kister, R. Dachsel, and H. Schumann. 2014. A Survey on Interactive Lenses in Visualization. In *EuroVis - STARs*. The Eurographics Association.
37. J. Viega, M. J. Conway, G. Williams, and R. Pausch. 1996. 3D Magic Lenses. In *Proc. UIST '96*. ACM, 51–58.
38. D. Vogel and R. Balakrishnan. 2004. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In *Proc. UIST '04*. ACM, 137–146.
39. Y. Zhang, T. Han, Z. Ren, N. Umetani, X. Tong, Y. Liu, T. Shiratori, and X. Cao. 2013. BodyAvatar: Creating Freeform 3D Avatars Using First-person Body Gestures. In *Proc. UIST '13*. ACM, 387–396.
40. J. Zillner, C. Rhemann, S. Izadi, and M. Haller. 2014. 3D-board: A Whole-body Remote Collaborative Whiteboard. In *Proc. UIST '14*. ACM, 471–479.