Eurographics / IEEE Symposium on Visualization 2011 (EuroVis 2011)
H. Hauser, H. Pfister, and J. J. van Wijk
(Guest Editors)

Volume 30 (2011), Number 3

The Undistort Lens

John Brosz, Sheelagh Carpendale, and Miguel A. Nacenta

Department of Computer Science, University of Calgary, Calgary, Alberta, Canada

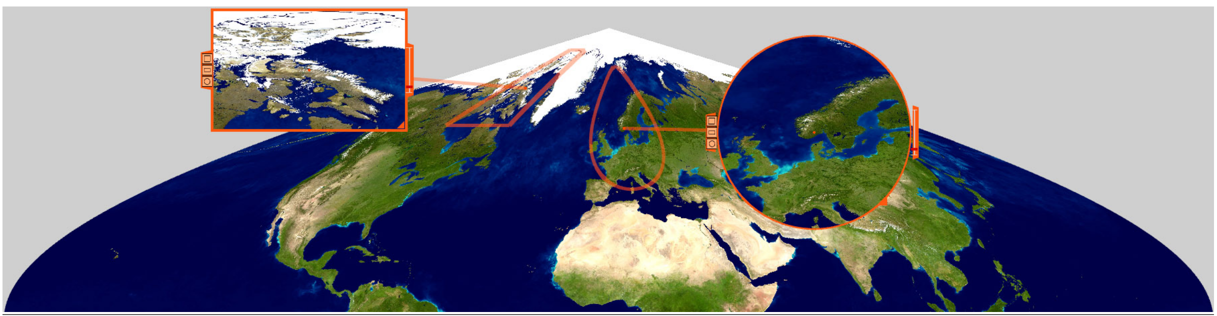

Abstract

Detail-in-context lens techniques can be useful for exploring visualizations of data spaces that are too large or have too much detail to fit in regular displays. For example, by bending the space in the right way we can bring together details from two separate areas for easy comparison while roughly keeping the context that situates each area within the global space. While these techniques can be powerful tools, they also introduce distortions that need to be understood, and often the tools have to be disabled in order to have access to the undistorted data. We introduce the undistort lens, a complement to existing distortion-based techniques that provides a local and separate presentation of the original geometry without affecting any distortion-based lenses currently used in the presentation. The undistort lens is designed to allow interactive access to the underlying undistorted data within the context of the distorted space, and to enable a better understanding of the distortions. The paper describes the implementation of a generic back-mapping mechanism that enables the implementation of undistort lenses for arbitrary distortion based techniques, including those presented in the lens literature. We also provide a series of use-case scenarios that demonstrate the situations in which the technique can complement existing lenses.

Categories and Subject Descriptors (according to ACM CCS): H. [H. Information Systems]: Information Interfaces and Presentation—H.5.2 User Interfaces

1. Introduction

As our understanding of the possibilities and benefits of detail-in-context viewing improves, more variations continue to arise. However, these information exploration techniques are still usually achieved through some type of distortion and while the accompanying distortions may have become more sophisticated, they are also frequently complex. Into this interactive milieu, we introduce the undistort lens. The undistort lens is an interactive technique that provides access to the undistorted data from within a distorted visualization.

Consider the following scenario: we want to use a map (Fig. 1A) to find the best route between two distant small towns that are hundreds of miles from each other. This is a typical detail-in-context problem, since a map scale that allows us to recognize and see the towns (the detail) would be too big to fit on any screen, and a scale that makes the distance between the towns fit within a reasonable space (the context) could make the towns invisible or indistinguishable from other similar objects. There are a number of visualization techniques based on distortion that can help us solve this problem. For example, magnification lenses can be placed on

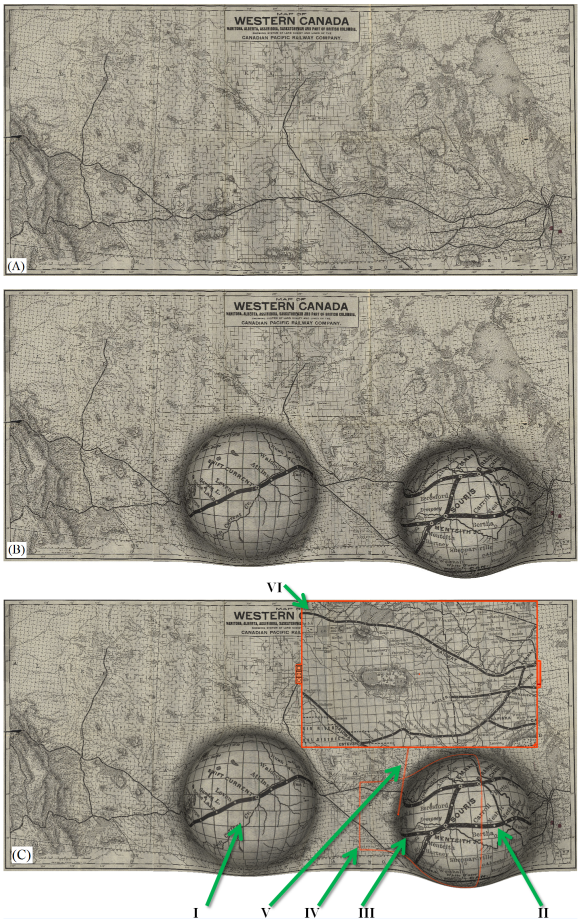

Figure 1: *(A) Finding specific towns in this large map is difficult. (B) Detail-in-context magnification lenses reveal the towns but obscure the railway connections. (C) An undistort lens clarifies the railway connections in the transition region between the magnification lenses. The arrows point out the magnification lenses (I & II), a road obscured by the lens' drop-off (III), the selected region to undistort (IV), the undistort lens connection (V), and displayed undistorted information (VI).*

top of the towns to enlarge the areas where detail is needed while retaining the context (Fig. 1B).

Unfortunately, the same distortion that helps us fit detail and context in a constrained space also makes certain tasks difficult; in the case of these magnification lenses, the drop-off areas deform the map in ways that make it difficult to decide if a road or railway will connect the towns or not. This can be circumvented by temporarily disabling the magnification lenses, but this is likely to confuse us even more, since disabling a lens will make the town impossible to locate and will force us to recognize the distorted shapes of the roads that we were following from their undistorted shapes. In other words, we were relying on the distorted representation as our context, and removing the context will make our task difficult.

The undistort lens is a simple and elegant way to address this problem. Instead of temporarily disabling distortion, forcing the traveler to commit to memory the two representations and establishing relationship between them on the go, the undistort lens provides an undistorted representation of the interest area that can be compared concurrently with the distorted context. Figure 1C represents the use of the undistort lens in the town route scenario. The magnification lenses (Elements I and II of Fig. 1C) make the towns visible, but it is difficult to trace a route through their drop-off areas (e.g., III). The undistort lens displays an area of the undistorted map through a rectangular offset window (VI), that is connected through a line with the area being undistorted (V). The undistort lens also provides an outline of the undistorted area within the distorted space (IV) which can help us relate the elements visualized in the undistorted area to the distorted representation.

In this paper we present the undistort lens: a technique that can be used interactively within any complex multiscale viewing environment, regardless to the complexity of the current distortion, to provide access to the original undistorted visualization. The next section outlines the related work. This is followed by the details of the inverse mapping algorithm and the implementation. In Section 5 we illustrate the versatility of the undistort lens through several examples of its use with a variety of different distortion viewing techniques.

2. Related Work

Since the undistort lens is a technique that both belongs to and complements the existing lens and detail-in-context viewing techniques, we provide a brief overview of these techniques to set the context for this research.

Early in the information presentation and exploration research, two eloquent arguments were made about the importance of context in information work. One is Spence and Apperley's [SA94] discussion of the importance of spatial memory in information search, which suggested that there are benefits to displaying full context even if it must be compressed. The other is Furnas' [Fur86] explanation of how it is common practice for people to set information they are interested in the context that they deem most important. From this Furnas suggested use of a degree-of-interest (DOI) function to create what he termed Fisheye views that set selected regions in an appropriate context. The research community's response to these ideas has been many and varied, giving rise to all kinds of viewing and interaction possibilities most of which create some type of focus and context view and make use of some form of spatial distortion.

The following are some examples of multi-scale focus and context presentations. Bifocal display [SA94] and the rubber sheet stretching metaphor [SSTR93] provide methods for distorting 2D space. Graphical Fisheye views [Fur86] apply Furnas' fisheye concepts to simple vector graphics. Table Lens [RC94] applies distortion viewing concepts to tabular data. PerspectiveWall [MRC91] and DocumentLens [RM93] use 3D perspective to create focus and context views of predominately 1D data and of large documents respectively. Accordion Drawing [MGT*03, SHM06] extends the rubber sheet concept to offer guaranteed visibility. Carpendale et al. [CCF95], Keahey et al. [KR96] and Shoemaker and Gutwin [SG07] offer constrained lenses that provide independent local magnification lenses and limit the spread of distortion into the rest of the context. Carpendale and Montagenese [CM01] introduced Elastic Presentation Framework (EPF) giving an algorithmic basis for unifying these approaches, making it possible to include any number of them in a given interface. DTLens [FS05] introduced local magnification lenses to multi-touch tables. More recently Sigma Lenses [PA08] combined the concept of Magic Lenses [BSP*93] with EPF to create a greater variety of integrated visual lens types. Melange [EHRF08], also looked at the EPF basis but considered the interaction metaphor of making regions smaller or of pushing regions away from view. This, similarly to folding in EPF, made it possible to juxtapose different regions of information that might initially be quite far apart. The undistort lens relates most closely to Magic Lenses [BSP*93], offering its own brand of magic to invert existing distortions.

In parallel with the development of the new focus and context viewing techniques there has been an ongoing discussion about the effect of distortion on comprehension [CCF97a, Fur06, MELS95]. In additional studies [Gut02, CG03, SG05] have shown that distortion viewing has both positive and negative effects. This has contributed to the discussion about whether the use of distortion outweighs the advantages of the focus and context presentations. Within this context, we explore the use of a layered interaction that integrates the ability to invert distortions, providing an undistort lens that can operate interactively in conjunction with all existing distortions.

3. The Undistort Lens

The primary concept of the undistort lens is to provide interactive access to the original undistorted data no matter what the state of the current distortion. The problem that motivates the undistort lens is that, although distortion-based visualizations can be disabled to reveal the undistorted data, we wish to examine a chosen section of the data free of distortions without necessitating the removal of a complex viewing configuration. As noted in Section 1 the added functionality of undistort can enhance the interactive exploration possibilities when combined with magnification lenses. The undistort lens provides access to original data within a multi-scale viewing environment.

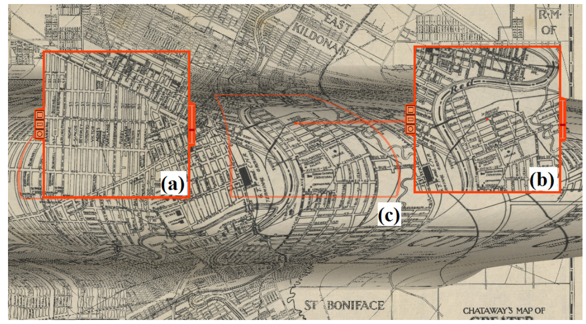

Figure 2: *The base of the undistort lens can either be directly under the distorted area (a) or connected (b) with a line between the region of interest and the undistort lens' display. The region shown within the undistort lens is indicated by an outline (c).*

Figure 3: *Undistort lenses can have a rectangular, square, or circular shape.*

The visual design of the undistort lens corresponds to its function of providing undistorted data within its distorted context. Similarly to other lenses, it has a base or region of connection to the current visual presentation and a focus or region that presents the undistorted visuals that directly correspond to the visuals indicated by the base. The focus or the undistorted region ((a) in Figure 2) presents the undistorted view of the underlying data that is within the base. Figure 2 shows how the undistort lens can display an undistorted region directly over the distorted area (a) or offset and connected through a line to the region of interest (b).

In our implementation, the shape, both base and focus, of the undistort lens can be switched between rectangle, square, or circle (Figure 3) through small buttons attached to left of the lens. Note that an undistort lens can easily be created in any arbitrary shape; we have chosen these three simple shapes because they are commonly-used in lens-related literature. The undistort lens provides an outline of its base within the distorted space of the current visualization. This indicates the area the undistorted lens is presenting, supporting a better understanding of the relationship between the base and the focus. Also, it is possible to independently adjust the size of the undistort lens area and the level of magnification within the undistort lens. This scale adjustment is accessible through a slider on the right of the undistort lens. Supporting changes of scale within the undistort lens allows presentation of the full range from the original visualization with no non-linear distortions, to a highly magnified small area. This freedom allows it to be also used in a similar manner to either an overview or a radar view [GRG96]. Finally, the area of interest within the current visualization is selected by choosing the center of interest and then dragging to establish either the radius of the circle or the position of a corner of the square or rectangle from the base of the undistort lens.

4. Implementation

The main challenge for the implementation of the undistort lens is the back-mapping algorithm. Although an undistorted version of the data space is available for most distortion-based transformations, the undistort lens must be able to refer specific points in the arbitrarily distorted space to their undistorted counterparts. It is necessary that this back-mapping works for any arbitrary distortion-based technique, including lenses covered in survey papers [LA94, CM01, CKB09] as well as recent variations such as sigma lenses [PA08] and melange [EHRF08]. To further illustrate undistort's flexibility, we also introduce new distortion factors such as lens rotation and lens wedges, combining these with lens folding [CCF97b].

Distortion techniques describe mappings from the original data to the distorted presentation. Essentially, these mappings are functions that transform locations on the original data space (x,y) into locations in the distorted space (x^*,y^*). Distortion techniques use a variety of functions to achieve these deformations of the original data into distorted viewing presentations. However, in order to relate a specific point on the distorted presentation to the point in the original data, we need the inverse function that maps (x^*,y^*) to (x,y). This inverse mapping is not trivial because a separate formula would have to be independently derived for each distortion technique. Such a derivation is not possible in many cases, since analytic inverse formulas do not exist for all distortion-based transformations [LA94, SSTR93].

To address this challenge we closely approximate this inverse mapping through a rendering technique very similar to that devised by Cowperthwaite et al. [CCF99], which was in turn based on Hanrahan and Haeberli's technique [HH90] for directly painting textures on projected 3D objects.

The inverse mapping technique works in two steps. First we encode the position of each pixel of the undistorted space into the colour channels of an equally-sized image. This means that by looking at the colour values of the pixel at (x,y) we can recover the exact value of x and y. Second, we distort this auxiliary image through the same process

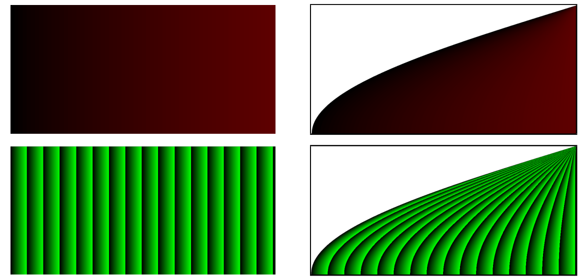

Figure 4: *Left: width coordinates encoded into red (top) and green (bottom) channels of a lookup texture. While normally combined into a single image, these channels have been presented as separate images to make changes in the red channel easily visible; blue and alpha channels are not shown. The black bands in the green channel mark where 8-bit overflow causes an increment in the red channel. Right: width lookup texture channels after deformation by top-left quarter of a Sinusoidal map projection [Sny97].*

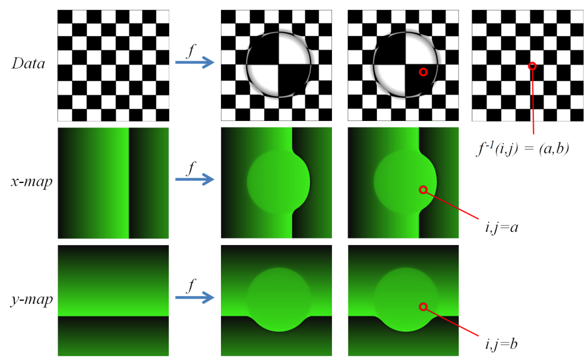

Figure 5: *Lookup textures (x-map and y-map, left column) encode x and y coordinates in RGB image values. Data is distorted by a function f, that is also applied to the lookup textures (middle column). To find the original coordinates of a point (i,j) on the distorted presentation (red circle, right column) we examine the channel values of pixel (i,j) in the x-map and y-map to determine the undisorted x and y respectively.*

that transformed the original data into a distorted presentation. This transformation deforms the auxiliary image, but the colour of the distorted pixels preserves the information of their original location in the undistorted space. If we want to find out the original location of any position in the distorted space, we can lookup the colour value of its coordinate in the distorted auxiliary image. Decoding this colour value provides the original coordinate from the undistorted space.

Since we want to be able to apply undistort to very large images, we need enough colour bits to encode any coordinate. To achieve this, we use two auxiliary images, one for the horizontal coordinate and one for the vertical. We call these images the horizontal and vertical lookup textures. Because each image has 4 channels (Red, Green, Blue, Alpha) of 8 bits each, we can encode coordinates of a space of up to $(2^{32})^2$ pixels. Figure 4 displays the red and green channels of a horizontal lookup texture before (left) and after (right) distortion; the blue and alpha channels are not shown since they are not used due to the image being less than 2^{16} pixels wide. Figure 5 summarizes the back-mapping process with horizontal and vertical lookup textures.

Backmapping is used in our interface to calculate the center of the undistorted content shown by the undistort lens, which is the backmapped location of the point of interest (V in Figure 1). To draw the outline of the area shown by the undistort lens in distorted space we take the original (undistorted) coordinates of points displayed along the inside edge of the undistort lens, apply the distortion f to these points, and draw them within the distorted presentation.

5. Undistort in Action

We introduced this paper with a scenario that used undistort with two EPF Gaussian lenses, showing how undistort could interactively offer scaled views of the lenses' distorted regions clarifying the positions of the roads and railway lines. In this section we illustrate how undistort can work with other distortions and show four additional uses: to explore and compare alternate distortions arising from map projections; to ease touch-up editing; when re-aligning magnification lenses for better juxtapositions; and in collaborative scenarios.

To illustrate undistort in use with different types of distortions, Figure 6 shows undistort in use with Melange type lenses. Melange uses rectilinear distortions to create radically folded views where regions of the map that were previously far apart can now be seen adjacently, albeit with the distortion between them providing an indication of the amount of information that exists between the two focal regions. The Melange lens employs shading to make the distinctions between focal, and thus scaled only, regions and regions of distortion. Note how undistort also removes the shading giving a clear view of the selected distorted region.

Figures 7 and 8 show undistort with sigma-type lenses. Sigma Lenses expand distortion lenses to include varied use of transparency. Here undistort is applied across both non-linear distortion and transparency. Figure 9 introduces a wedge lens. A wedge lens can combine the advantages of a Manhattan Lens with that of a blended lens. In this figure the lens has been chosen to fit the angle between the fork in the railways, employing the Manhattan's abrupt drop-off to the bottom and to the right and using a Gaussian blend on the left. The undistort lens is oblivious of these changes in type of distortion used and presents the area from the top of the focal region down the steep side to include part of the context.

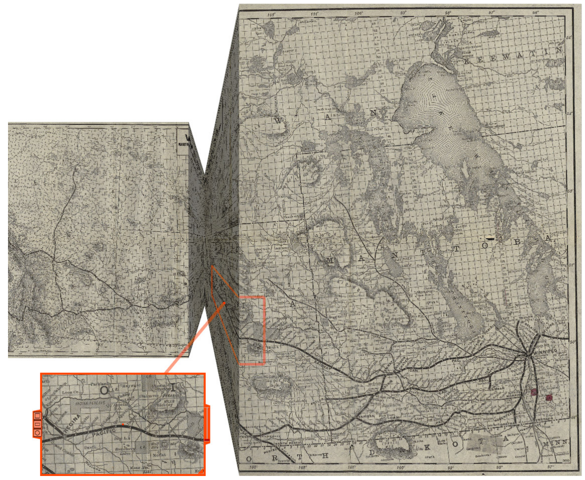

Figure 6: *An undistort lens used with a Melange-like space fold.*

Figure 7: *A sigma-type lens is positioned over the north sea with an undistort lens centered on Olso.*

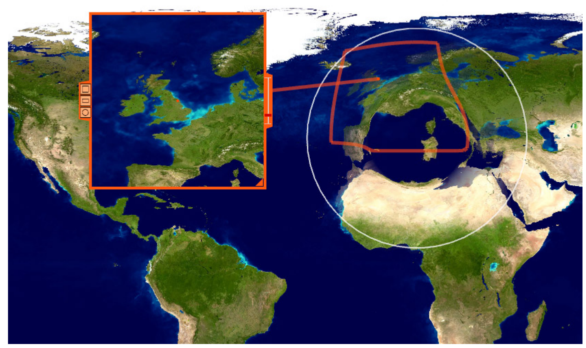

Figure 8: *An undistort lens is used to clarify the details on Great Britain which was currently less clear on the side of a sigma-type lens.*

High magnification exploration is also an important scenario, especially as more and more extremely high resolution images are made available. For example, this high resolution exploration involves examining the details in digitally preserved artworks. Here the high resolution detail is necessary to make out features of interest such as fine detail, brush strokes, paint composition, tapestry weave, etc.

In exploring extremely high resolution imagery with a detail and context lens, even with lenses designed for this pur-

Figure 9: *A wedge lens positioned to fit the angle formed by a fork in the railway. An undistort lens presents a region on either side of a sharp, Manhattan-style drop-off.*

Figure 10: *Exploration of high magnification imagery (*The Arnofini Portrait *by Jan van Eyck, oil on oak, 1434) with an EPF detail and context lens and an undistort lens. The detail and context lens features* $10X$ *magnification and a power function dropoff with a parameter value of* 2 *as per [CLP04]. One undistort lens is used to clarify the position of the high magnification area in the original scene, the other is used to make text legible.*

pose such as those described by Carpendale et al. [CLP04], it still can be difficult to maintain useful context. For instance in Figure 10 if we consider the $10X$ magnification detail+context lens on its own, the context is greatly distorted making comparisons between the magnified reflection in the mirror to the rest of the scene difficult. The addition of the circular undistort lens and the independent adjustment of its magnification level provides a simple means of comparison without interfering with the focus area. An additional rectangular undistort lens is used to make text in the context area legible at a middle-level magnification. While use of a multiplicative lens on top of another lens [CLP04] can also provide this middle-level magnification, adjustments to the lower lens are propagated to the upper lens consequently discouraging exploration with the middle magnification as any such changes causes the loss of the current focus. The undistort lens in contrast acts independently allowing any its parameters (magnification, shape, position, point-of-interest) to be changed without causing disruption to the high magnification focus of the detail and context lens.

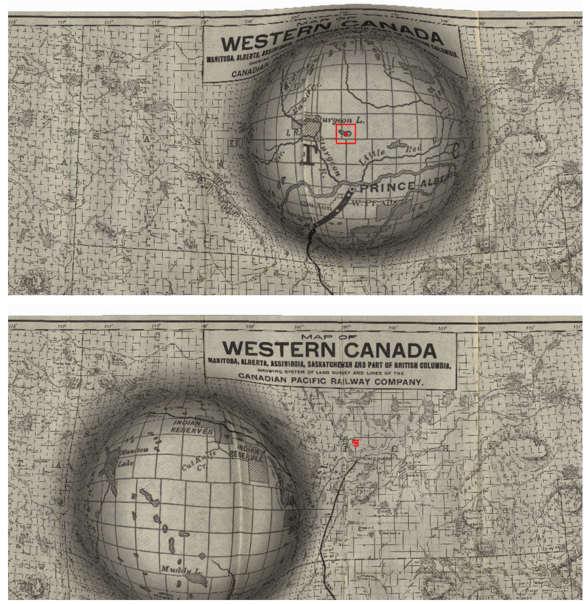

Figure 11: *A simple example of annotation where a point is marked (red) in an area magnified by a context and detail lens (top). Then lens is then moved but the red marker retains its position with respect to the underlying image despite the change in the distortion (bottom).*

5.1. High Resolution Image Editing

Since undistort can directly access the underlying visualization, it can also be used for editing. Figure 11 shows a simple red-marker annotation to denote the location of a point. Figures 12 and 13 illustrate complex touch-up editing as described below.

A task that can benefit from the undistort lens is that of digital touch-up. Removing or adjusting small details in pictures is common among professional and amateur photographers. Due to the high-resolution of current cameras, digital touch-up usually requires switching between two levels of magnification: a highly magnified level where individual pixels can be clearly discerned; and the magnification level at which the image is likely to be viewed as a whole. Changing and picking out individual pixels requires high magnification, but low magnification is required to ensure

Figure 12: *A photograph is edited to remove a cloud. Edits performed at high magnification should be examined at normal magnification to ensure that they blend in. Left: original. Middle: poorly blended edit. Right: well blended edit.*

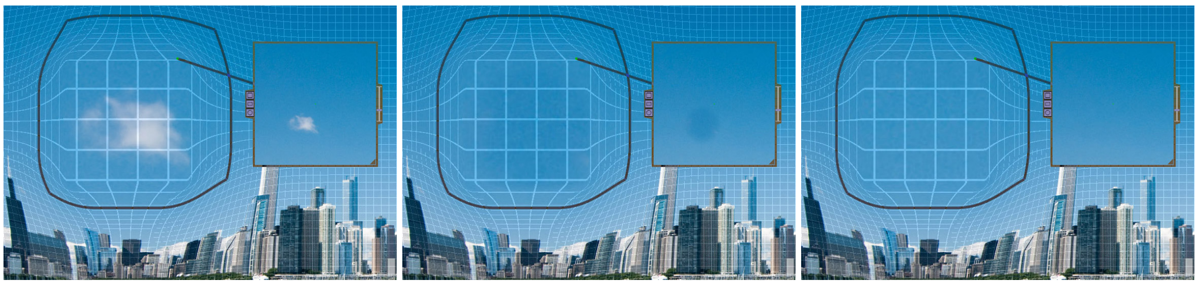

Figure 13: *Mock-up of editing with a* 4*X magnification detail and context lens with assistance from an undistort lens. Left: original image; middle: poorly blended local editing; right: well blended local editing. A grid has been super-imposed to reveal the detail and context lens as shading causes greater disruption to colour values necessary for photo editing.*

that the changes appear seamless over the macro context of the changed area. Figure 12 shows an image being touched up to remove a small cloud from a large city skyline.

Rather than switching between magnifications, this task can be accomplished with a combination of a detail-in-context lens and an undistort lens. As shown in Figure 13, the detail and context lens provides an area of high magnification where local editing operations can take place. The undistort lens is used to check the global context of the local area without tying up the screen space necessary to present another copy of the image. Additionally the use of undistort's backmap (colour-based lookup) for annotation allows pixel operations to be performed on the original image when specified in image space.

5.2. Map Projections

Undistort lenses can be used to examine map projections. All projections of a spherical surface (e.g., the earth) onto a plane introduce distortions. To address these distortions a wide variety of cartographic projections exist with different trade-offs between the preservation of area, scale, angle, and shape [Sal06]. In this scenario, rather than undoing distortion, undistort lenses support exploration, understanding, and comparison between different cartographic projections as shown in Fig. 14. The undistort lens' ability to present the same area and the outline of area's shape on the original projection allows direct comparison between different projections and can lead to a better understanding of the biases of these projections without losing global context, and without having to resort to other visual artifacts such as grid lines.

We can also use the undistort lens to undo the effects of displaying a world map on a sphere. Figure 15 depicts a representation of the earth where Japan appears skewed and thin due to the current orientation of the globe (facing the Indian Ocean). In this case, the undistort lens shows a cylindrical rendition of the area.

5.3. Folded Lens

Sometimes with detail-in-context lenses it is desirable to reposition the foci to align the information in the lens foci with either other foci or other regions of the context. This re-positioning of a magnified focus has been referred to as folding [CCFS95, CM01]. Since a folded lens has moved its focus across the context, it is a distortion-based deformation that makes certain parts of the data disappear under a fold. Figure 16 shows an example of folded lenses where the displacement of the focus areas forces the occlusion of drop-off regions. Folded lenses provide a compromise solution when the focus area needs to be displaced a large distance from its original location (e.g., to compare two distant areas of interest side by side), but still preserve a partial connection with their context. Figure 16 exemplifies this by showing a comparison of the distances between stations in two different areas of a railway line, without losing all the context of how the rail line continues into the rest of the map.

Naturally, the large drop-off region of the folded lens generates substantial distortion but, as Fig. 16 illustrates, we can

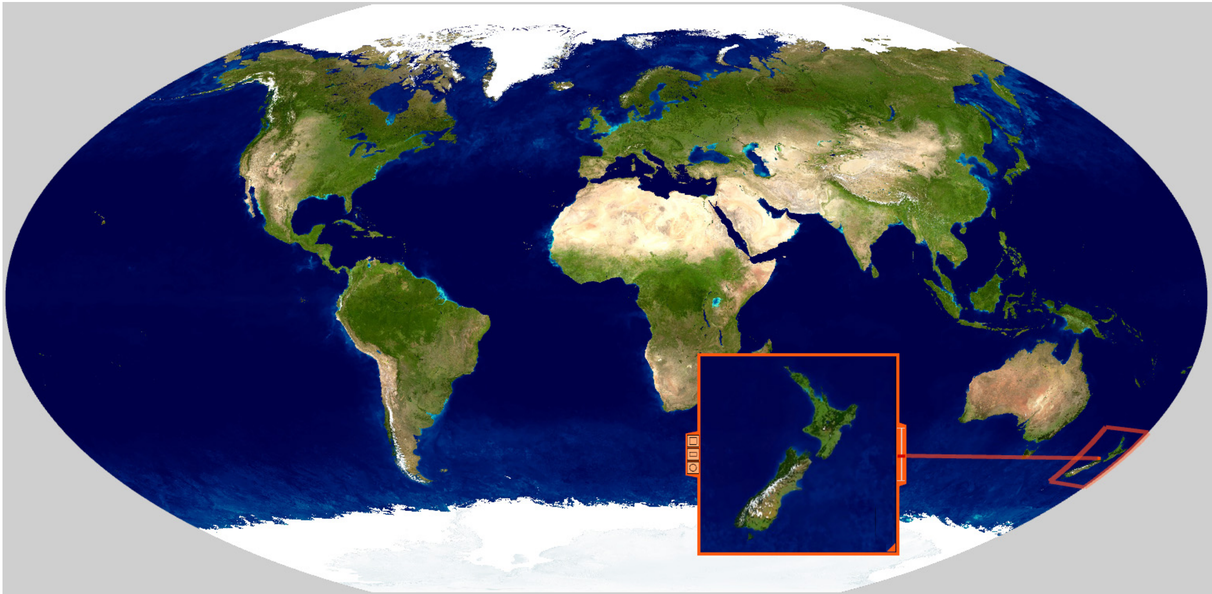

Figure 14: *Exploring the differences between a Winkel Tripel projection and the original image data (a cylindrical projection).*

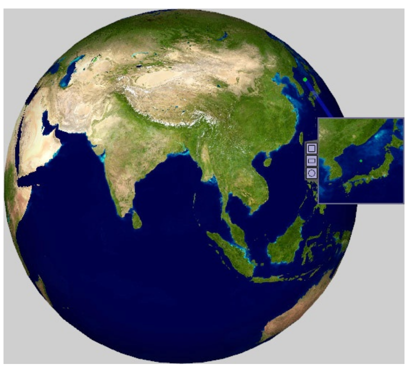

Figure 15: *Use of the undistort lens to present a cylindrical projection next to a 3D globe.*

then apply an undistort lens to explore this area without having to lose the location of the folded lenses. The undistort lens can also be positioned to reveal areas that fall under the fold; moreover, by examining the shape of the undistort base, a better understanding of the distortion and the folded parts can be achieved.

5.4. Collaborative multi-user detail and context

In collaborative environments (e.g., collaborative tabletop information visualization [KCSG04, IC07]) it is often necessary to adapt the content to the different positions and orientations of people. To support this we introduce a distortion lens variation in which the lens focus can be rotated, while maintaining the same information content. This can be useful in showing a lens' content to a collaborator, providing the right viewing orientation for the collaborator. This rotation will introduce a twist within the distorted region of the lens. Figure 17 shows a single rotated lens with the word 'city' now aligned to the right hand side. Lenses can provide rotation and translations that accomplish this while preserving connections to the global context (Fig. 18).

The undistort lens complements these collaborative environments in three main ways. First, undistort lens(es) can be placed to allow one person to view the area's original data without interfering with somebody else's distortion analysis of the area (e.g., Figures 17 and 18). Second, undistort lenses can be rotated themselves to facilitate the analysis of the undistorted data. Third, undistort lenses and their bases can assist in assessing the current orientation relative to the context of rotating lenses.

6. Conclusion

Distortion-based techniques such as detail-in-context magnification lenses, sigma lenses, and Melange can be helpful to address detail-in-context problems where a single uniform presentation of the data space is not sufficient to effectively use the data. In this paper we have introduced the undistort lens, a technique that complements distortion-based techniques to address the problems that they introduce.

The undistort lens provides an auxiliary undistorted visualization of an area of interest that is visually linked to the context of the distorted data space created by a distortion-based technique. The technique enables the exploration of an area in its original undistorted form without having to

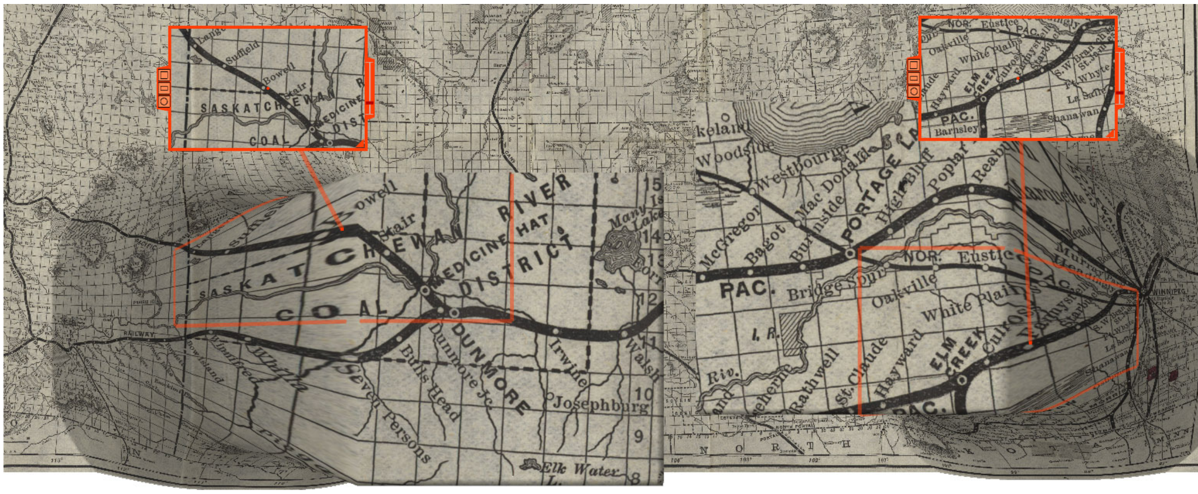

Figure 16: *Two* 4.5 *magnification lenses that have been folded over to compare two separated areas of a map. Two undistort lenses assist in understanding the distortions, reading labels, and revealing obscured areas.*

disable the distortion, and allows comparison of the distorted and undistorted presentations. We demonstrate through a series of examples how the undistort lens can be applied to scenarios where the distorted space is not sufficient, and for situations where a better understanding of a distortion can be useful (e.g., spherical projections).

Finally, we also contribute the description of a backmapping algorithm that enables the implementation of undistort lenses for any arbitrary distortion techniques. We believe that the addition of undistort lens capabilities to existing distortion-based representations can contribute to their usefulness by helping address some of the problems generated by distortion.

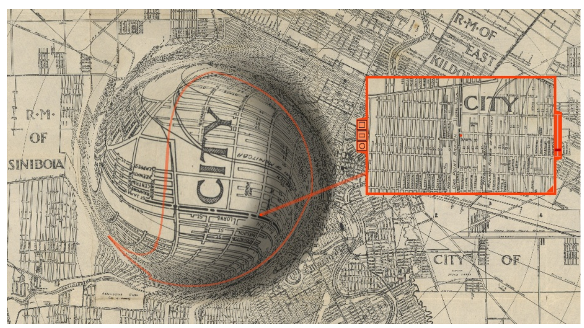

Figure 17: *A detail-in-context lens where the focus can be arbitrarily oriented. The undistort lens can clarify data's original orientation.*

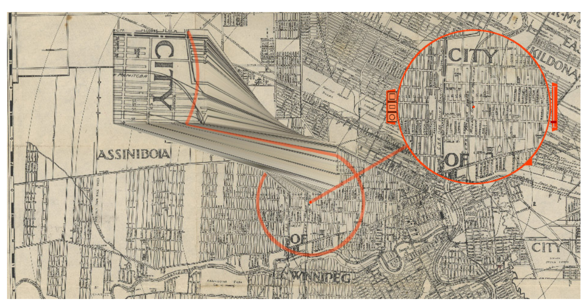

Figure 18: *A detail-in-context lens that automatically orients towards the nearest screen edge based on how the lens has been folded. The undistort lens and its frame make clear the orientation of the focus of the detail and context lens relative to the context.*

Acknowledgements

Three of the images used in this work were provided by outside sources: *The Blue Marble: Land Surface, Ocean Color and Sea Ice* was provided by NASA's Visibile Earth http://visibleearth.nasa.gov/view_rec.php?id=2430. *Chataways's Map of Greater Winnipeg Enlarged and Revised Edition 1917* and *Map of Western Canada Manitoba, Alberta, Assiniboia, Saskatchewan and Part of British Columbia, Showing System of Land Survey and Lines of the Canadian Pacific Railway Company (1893)* used were obtained through a Creative Commons attribution license from http://www.flickr.com/photos/manitobamap.

We would like to thank the reviewers for their insightful comments as well as for support from the Natural Sciences and Engineer Research Council (NSERC), iCORE, and Smart Technologies Inc.

References

[BSP*93] BIER E. A., STONE M. C., PIER K., BUXTON W., DEROSE T. D.: Toolglass and magic lenses: The see-through interface. In *Proceedings of Siggraph 93* (Aug. 1993), ACM, pp. 73–80. 3

[CCF95] CARPENDALE S., COWPERTHWAITE D. J., FRACCHIA F. D.: 3-dimensional pliable surfaces: for the effective presentation of visual information. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (1995), ACM, pp. 217–226. 3

[CCF97a] CARPENDALE S., COWPERTHWAITE D., FRACCHIA F.: Making distortions comprehensible. In *Proc. of IEEE Symp. on Visual Languages: VL* (Sept. 1997), IEEE, pp. 36–45. 3

[CCF97b] CARPENDALE S., COWPERTHWAITE D. J., FRACCHIA F. D.: Extending distortion viewing from 2D to 3D. *IEEE Comp. Graph. Appl. 17*, 4 (July/Aug. 1997), 42–51. 4

[CCF99] COWPERTHWAITE D. J., CARPENDALE M. S. T., FRACCHIA F. D.: *Editing in elastic presentation spaces*. Tech. Rep. TR 1999-11, Simon Fraser University, Burnaby, BC, Canada, Oct. 1999. 4

[CCFS95] CARPENDALE S., COWPERTHWAITE D. J., FRACCHIA F. D., SHERMER T.: Graph folding: Extending detail and context viewing into a tool for subgraph comparisons. In *Graph Drawing: GD* (1995), Springer-Verlag, pp. 127–139. 7

[CG03] CARL A. S., GUTWIN C.: Finding things in fisheyes: Memorability in distorted spaces. In *Graphics Interface* (2003), pp. 67–75. 3

[CKB09] COCKBURN A., KARLSON A., BEDERSON B. B.: A review of overview+detail, zooming, and focus+context interfaces. *ACM Comput. Surv. 41* (Jan. 2009), 2:1–2:31. 4

[CLP04] CARPENDALE S., LIGHT J., PATTISON E.: Achieving higher magnification in context. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (2004), ACM, pp. 71–80. 6

[CM01] CARPENDALE S., MONTAGNESE C.: A framework for unifying presentation space. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (2001), ACM, pp. 61–70. 3, 4, 7

[EHRF08] ELMQVIST N., HENRY N., RICHE Y., FEKETE J.-D.: Melange: space folding for multi-focus interaction. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2008), ACM, pp. 1333–1342. 3, 4

[FS05] FORLINES C., SHEN C.: Dtlens: multi-user tabletop spatial data exploration. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (2005), ACM, pp. 119–122. 3

[Fur86] FURNAS G.: Generalized fisheye views. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (1986), ACM, pp. 16–23. 2, 3

[Fur06] FURNAS G. W.: A fisheye follow-up: further reflections on focus + context. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2006), ACM, pp. 999–1008. 3

[GRG96] GUTWIN C., ROSEMAN M., GREENBERG S.: A usability study of awareness widgets in a shared workspace groupware system. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work* (1996), CSCW '96, ACM, pp. 258–267. 4

[Gut02] GUTWIN C.: Improving focus targeting in interactive fisheye views. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2002), ACM, pp. 267–274. 3

[HH90] HANRAHAN P., HAEBERLI P.: Direct WYSIWYG painting and texturing on 3d shapes. In *Proceedings of SIGGRAPH 90* (1990), vol. 24, ACM, pp. 215–223. 4

[IC07] ISENBERG P., CARPENDALE S.: Interactive tree comparison for co-located collaborative information visualization. *IEEE Trans. Vis. and Comput. Graphics 13*, 6 (2007), 1232 –1239. 8

[KCSG04] KRUGER R., CARPENDALE S., SCOTT S. D., GREENBERG S.: Roles of orientation in tabletop collaboration: Comprehension, coordination and communication. *Computer Supported Cooperative Work (CSCW) 13* (2004), 501–537. 8

[KR96] KEAHEY T., ROBERTSON E.: Techniques for non-linear magnification transformations. In *Proceedings of the IEEE Symp. on Information Visualization '96,* (Oct. 1996), pp. 38 –45. 3

[LA94] LEUNG Y. K., APPERLEY M. D.: A review and taxonomy of distortion-oriented presentation techniques. *ACM Trans. Comput.-Hum. Interact. 1* (June 1994), 126–160. 4

[MELS95] MISUE K., EADES P., LAI W., SUGIYAMA K.: Layout adjustment and the mental map. *Journal of Visual Languages & Computing 6*, 2 (1995), 183 – 210. 3

[MGT*03] MUNZNER T., GUIMBRETIÈRE F., TASIRAN S., ZHANG L., ZHOU Y.: Treejuxtaposer: scalable tree comparison using focus+context with guaranteed visibility. *ACM Trans. Graph. 22* (July 2003), 453–462. 3

[MRC91] MACKINLAY J. D., ROBERTSON G. G., CARD S. K.: The perspective wall: Detail and context smoothly integrated. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (1991), pp. 173–180. 3

[PA08] PIETRIGA E., APPERT C.: Sigma lenses: focus-context transitions combining space, time and translucence. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2008), ACM, pp. 1343–1352. 3, 4

[RC94] RAO R., CARD S. K.: The table lens: merging graphical and symbolic representations in an interactive focus + context visualization for tabular information. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (1994), ACM, pp. 318–322. 3

[RM93] ROBERTSON G. G., MACKINLAY J. D.: The document lens. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (1993), ACM, pp. 101–108. 3

[SA94] SPENCE R., APPERLEY M.: Data base navigation: an office environment for the professional. *Behaviour and Information Technology 1*, 1 (1994), 43–54. 2, 3

[Sal06] SALOMON D.: *Transformations and Projections in Computer Graphics*. Springer-Verlag, 2006. 7

[SG05] SKOPIK A., GUTWIN C.: Improving revisitation in fisheye views with visit wear. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2005), ACM, pp. 771–780. 3

[SG07] SHOEMAKER G., GUTWIN C.: Supporting multi-point interaction in visual workspaces. In *Proc. of the SIGCHI conf. on Human factors in comp. systems* (2007), ACM, pp. 999–1008. 3

[SHM06] SLACK J., HILDEBRAND K., MUNZNER T.: Prisad: A partitioned rendering infrastructure for scalable accordion drawing (extended version). *Information Visualization 5* (2006), 137–151. 3

[Sny97] SNYDER J. P.: *Flattening the Earth*, 2 ed. University of Chicago Press, 1997. 4

[SSTR93] SARKAR M., SNIBBE S. S., TVERSKY O. J., REISS S. P.: Stretching the rubber sheet: a metaphor for viewing large layouts on small screens. In *Proc. of the ACM Symp. on User Interface Software and Tech: UIST* (1993), ACM, pp. 81–91. 3, 4