

Art and Nonlinear Projection

John Brosz*¹

Sheelagh Carpendale¹

Faramarz Samavati¹

Hao Wang²

Alan Dunning²

¹Department of Computer Science,
University of Calgary

²Alberta College of Art and Design

Abstract

Nonlinear projection is a current area of research in computer graphics. It provides a meeting place between science and art. After describing motivations for use of nonlinear projection, we provide a brief survey of different techniques for nonlinear projection described in computer graphics literature. We then examine several cases of how and when artists choose to use nonlinear projection, as opposed to perspective projection, and how these might be achieved with current computer graphics techniques. Lastly, we provide a synopsis of the Flexible Projection Framework and then describe how it was used in a collaborative project between an artist and a computer scientist.

1 Introduction

One of the most difficult aspects of creating images is determining how to represent three dimensional scenes on a two dimensional medium. This process of moving from three dimensions down to two dimensions is known as projection and it is not easy. The specific type of projection (if any) and our adherence to this type can drastically affect the impression an image makes. Some projections can make images seem very realistic; others can make the images surrealistic. Use of particular projections can create images that make you feel like you are within or close to the image; other projections create images that exaggerate a divide between you, the viewer, and the scene being portrayed [9].

Projection is an interesting meeting point between art and science. The scientific approach to projection is to view it as the mathematical reduction of dimension and organization of the reduced dimension's space. From here formulae, geometries, and spaces are theorized and derived. Artistically, projection is a means of organizing space, relationships between objects, and the observer's presence in an image.

By far the most familiar type of projection is perspective projection. This is the type of projection produced by most cameras and is seen widely in Western art. However, both before and after the discovery and formalization of perspective projection, artists have made use of other types of projection. The most prevalent examples are the many projections that have been developed to represent the earth on a flat surface. A long history of cartographic development has been explored by Snyder [15].

In computer graphics, projection has mostly been limited to projections that are linear in nature; that can be represented as linear transformations. The two common examples of such projections are perspective and orthographic. Recently a number of researchers have begun exploring other possibilities for nonlinear projection.

In this paper we begin with a brief survey of current computer graphics approaches to nonlinear projection. Then we look at some examples of the use of nonlinear projection in art and discuss how these might be reproduced with computer graphics techniques. We then describe our nonlinear projection technique, Flexible Projection [3] and how this technique was used in a collaborative art project.

2 Nonlinear Projection in Computer Graphics

Salomon [12] describes a wide variety of nonlinear projections including a variety of panoramas as well as fisheye, false perspective, telescopic, and microscopic projections. Wyvill and McNaughton [19] describe projections used in ray-tracing as a mapping between image space and rays that define and sample a volume. They demonstrate how this mapping can be defined to produce fisheye and panoramic projections. Similarly Glassner [6] defines ray tracing projections with two NURBS (Non-Uniform Rational B-spline) surfaces; one surface positioning the rays' origin, the other surface providing direction.

Nonlinear ray-tracing allows light to follow curved paths based on vectors fields [7, 18]. These projections have been used to visualize the behavior of dynamical systems [7] and to visualize relativistic effects [18].

Yu and McMillan [21] describe a framework capable of reproducing a wide variety of linear projections; that is, projections with linear equations. Linear projections such as orthographic and perspective have useful features such as mapping straight edges to straight lines in the image and are easily implemented in graphics hardware. Subsequently Yu and McMillan provide a framework [20] for creating nonlinear projections by using many of these linear projections together to create a single image. Image continuity is maintained by careful construction and selection of neighboring projections. A similar approach is that of Trapp and Döllner [16] where planar projections and non-planar projections are combined by placing tiles that cover the image. Each tile corresponds to a planar or non-planar projection that is used to fill in the area of the tile. These approaches can be considered multi-camera projections.

Multi-camera projections combine images produced by several cameras into a single image. The simplest example of this is putting together a collage of images such as that expertly done in David Hockney's *Pearl-blossom Highway No. 2*, 1986 [9]. Examples in computer graphics that make use of this sort of technique include Cubist Style Rendering [5], Automating Joiners [22], and Photographing Long Scenes with Multi-Viewpoint Panoramas [1]. Rademacher and Bishop [11] provide an interesting variation where an image is created by moving a camera along a path. At regular points along this path columns of pixels are captured and placed beside one another to form an image.

A more advanced multi-camera approach, instead of combining images, is to combine objects in normalized device coordinates after projection. Agrawala et al. [2] do exactly this, allowing each object to be individually projected by a different perspective camera. These projections are then combined in normalized coordinates where depth testing and clipping operations occur. Singh [14] and Coleman and Singh [4] expand upon this, allowing objects to be projected by several cameras at once. Objects projected by more than one camera are only projected once, with a camera that interpolates between camera settings.

The Flexible Projection Framework [3], discussed in more detail in Section 4, operates by using parametrically modelled surfaces to define a viewing volume. We refer to Flexible Projection as a framework because it provides common techniques and tools for creating a wide variety of linear and nonlinear projections. This framework is able to reproduce all of the previously mentioned projections that incorporate a single viewing volume (excluding most multi-camera techniques). It also allows curved paths through the volume that allow for reproduction of the effects seen in nonlinear ray tracing.

Deformation is the manipulation, warping, or bending of the geometry of virtual objects. While often both nonlinear projection and deformation can be used to achieve the same result, the major difference is that nonlinear projection works with the camera; the geometry of the scene remains undisturbed and lighting features such as shadows remain true to the geometry. An interesting work halfway between nonlinear projection and deformation is that of Rademacher [10] where deformation is made dependent upon viewing

direction.

3 Nonlinear Projection in Art

Nonlinear projection has been used widely in art, beginning with cave art and continuing to the present day. In this section we investigate a few examples of nonlinear projection from the fine arts. We have chosen these examples as they provide insight into why and how artists choose to use nonlinear projection in their own creations.

Nonlinear Projection from Optics. In *Secret Knowledge* [9] David Hockney examines a variety of evidence that leads him to conclude that many artists used concave mirrors and later lenses to construct images from the early fifteenth century and onward. While these primitive lenses would allow artists to paint from a moving image, a key limitation was the small size of the projected images. In general these images could be no larger than 30 cm across.

As a consequence, if an artist wished to create a larger image, the artists would either have to try to repaint the image larger without aid of optics, or project the scene a piece at a time and then blend the pieces together into a unified whole. Images constructed in this fashion no longer strictly adhere to perspective projection. Rather each object appears to be directly viewed (i.e., projected as if positioned in the center of the field-of-view of a perspective projection). Hockney [9] points out that these composite projections present a compelling closer view of the world because all parts of the scene are viewed straight on.

It seems quite clear that this type of nonlinear projection performed by artists is the inspiration for the multi-camera projection systems discussed in Section 2. As this technique was originally created to bypass the technical restrictions of early optics, one wonders if this truly is a good technique for composing nonlinear projections. On the other hand, these techniques tie their approach to perspective projection, are based upon a clear, physical metaphor (that of constructing a collage) and, lastly, they allow the user to concentrate on a specific sections of the image, one at a time.

Van Gogh. Heelan [8] proposed that many of Vincent Van Gogh's paintings from 1888 onwards portray the subjective perception of constant curvature hyperbolic geometry; one such painting is that shown in Figure 1. In these paintings the projection is divided into three zones: near, intermediate, and far. In the near zone of Figure 1 vertical and horizontal surfaces such as the bed's footboard and the near chair appear to protrude in a convex bulge. In the intermediate zone, containing the bed and the nightstand, objects appear much as they do in a perspective image, while in the far zone, vertical and horizontal surfaces appear to be concave and depth differences (e.g., foreshortening) become less noticeable, as seen in the far chair and the back wall. Another difference is that the horizon will appear at a finite distance from the observer, rather than at infinity as in a perspective projection [8]. Heelan maintains that Van Gogh's paintings provide a visual embodiment of the subjective spatial perception of binocular vision.

To examine how this sort of projection could be handled with graphics techniques, the key point is that the projection changes as a function of the depth of the image, between the foreground, middle, and background. This is well suited to Flexible Projection where the viewing volume can be altered to produce these changes in projection at the desired depth in the scene. Additionally continuity of the viewing volume will ensure smooth changes between the differing areas seen in the painting.

Creative Perspective. Watson's book, *How to use Creative Perspective* [17] provides interesting insight into techniques used by illustrators to create perspective projections as well as to modify them for various



Figure 1: Vincent van Gogh, Dutch, 1853-1890, *The Bedroom*, 1889, Oil on canvas, 29 x 36 5/8 in. (73.6 x 92.3 cm), Helen Birch Bartlett Memorial Collection, 1926.417, The Art Institute of Chicago. Photography © The Art Institute of Chicago.

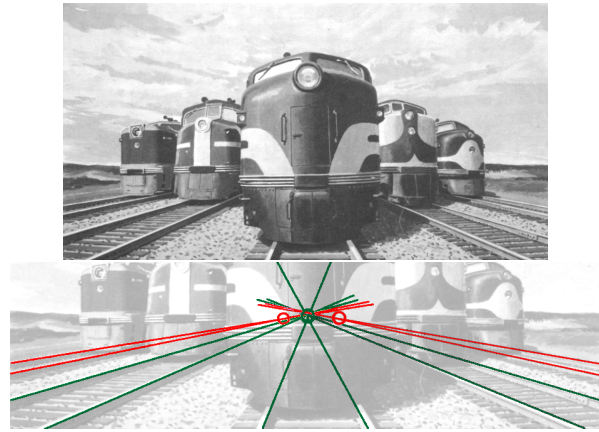


Figure 2: Top: a painting for an advertisement for the Association of American Railroads. Image from [17]. Bottom: an image where we have extended lines from each pair of rails to examine the vanishing points. In a perspective projection, assuming that the rails are parallel, the rails should all converge to the same point; our drawn lines show that this is not the case.

purposes. He calls such modified projections *creative perspective* and notes that an artist who uses “perspective creatively [can] bend it to his uses rather than be limited always to strict conformity” [17]. These manipulations are used: to allow the viewer to see structures that should be hidden, to change the scene (usually by manipulating vanishing points) and improve its overall composition, to allow the viewpoint to shift to more closely re-create human visual experience, and, lastly, to make the illustrations more dramatic and interesting. This last point, while somewhat vague, is important. Manipulating the projection allows illustrators to present objects in precisely the way they feel will have the desired impact.

We will examine two examples from this text. The first is the illustration of trains shown in the top of Figure 2. According to Watson [17] the artist has purposefully altered the vanishing point of the rails in the projection to compensate for eye movement across the picture; consequently this change makes the scene seem more natural.

The second example, shown in Figure 3, has a diagram showing how objects within the illustration have been moved and reoriented. In the diagram we can see that the illustrator has slightly altered the orientation of the wagon in relation to the horses, as well as the orientation of the horses to one another in order to present the most dramatic views of each [17]. With these subtle changes the illustrator has maintained realism while improving the produced image.

These two examples are difficult to describe as being best handled by any single computer graphics technique. The first example can be achieved with a multi-camera approach but care needs to be taken to achieve a smooth transition between changing viewpoints. Flexible Projection could also be applied as its parametric nature makes continuity easily achieved; however, adjusting the viewing volume to achieve the exact change in vanishing points requires finesse. The second example is interesting in that it could be achieved through a multicamera projection (one camera for the horses, the other for the engine), through Flexible Projection (where the viewing volume is shaped to produce a bend between the horses and engine), or through deformation of the models and then using perspective projection.



Figure 3: Left: illustration by Fred Freeman. Middle: sketch by Watson presents the illustration as if wagon and horses are travelling in the same direction in perspective. Right: top down diagram of the horses' position relative to the wagon as interpreted from the left image, assuming it was created with a perspective projection. This diagram indicates how Freeman has positioned and adjusted the size of the horses from the expected orientation of parallel to the engine. All images from [17].

4 Flexible Projection

This Section provides a synopsis of the Flexible Projection Framework as described in [3]. We have included this to assist in understanding of how Flexible Projection was used in the project described in Section 5.

Flexible Projection is unusual as it approaches projection, a rendering technique, with modelling tools. The primary concept is to define the viewing volume (the 3D volume that will appear in the 2D image) as a parametric volume $Q(u, v, t)$. This volume is defined so that the parameters u, v , and t correspond to the width, height, and depth positions that will result in the projected image. This volume can be stretched and reshaped to affect the projection and consequently the image.

While this volume can be defined in any manner whatsoever; we recommend using familiar surfaces to aid in the definition. For example, consider perspective projection. This projection uses a frustum as its viewing volume and is usually defined by an eye position, near surface, and far surface as shown. Such a frustum is shown in Figure 4 with the axes for the u, v, t parameters marked. Due to the familiarity of this construction, it makes sense to define Flexible Projections by defining a near and a far surface. We can then create our volume by interpolating between these two surfaces:

$$Q(u, v, t) = (1-t)S_0(u, v) + tS_1(u, v) \quad u_{min} \leq u \leq u_{max}, v_{min} \leq v \leq v_{max}, 0 \leq t \leq 1$$

where S_0 and S_1 are the near and far surfaces respectively. We do not include an eye position as many projections do not feature a single eye position.

Interesting variations on projection can be obtained by altering the interpolation between surfaces. For instance, by adding surface(s) between the near and far surfaces we can use Bézier curves to interpolate between the surfaces, consequently obtaining greater control over the volume, as is shown in Figure 5. This results in a definition of our volume as:

$$Q(u, v, t) = \sum_{i=0}^{n-1} B_{i,n-1}(t)S_i(u, v)$$

where $B_{i,d}(t)$ is the Bernstein polynomial $\binom{d}{i} t^i (1-t)^{d-i}$ and S_i is the i th surface (making S_{n-1} the far surface). Certainly other curve schemes could prove useful for this interpolation, however the ability of




Figure 4: A perspective projection's viewing volume. The volume is bounded by the near and far planes (green) and the sides of the frustum (yellow).




Figure 5: Flexible projection using nonlinear interpolation. Left: the projection volume. The large green curved surface is the near surface, the red rectangle in front of the car is used to curve the volume, and the green rectangle behind the car is the far surface. Right: the projected result.

Bézier curves to interpolate the near and far surfaces makes this curve scheme a useful place to start. Figure 5 shows how a curved volume can be created to show the front, back, and side of a car at the same time.

Projection in the volume occurs based on surface parameterization. We can imagine projection occurring by following light paths through the volume. The paths the light follows are iso-curves within the volume, known as *projectors*, that are defined by fixing u, v values $Q(u, v, t) = Q_{u,v}(t)$. For surfaces that flatten easily, image creation is a matter of scaling parameter values by the desired pixel dimensions of the image. For surface parameterizations that do not map directly to an image (i.e., the case of a hemisphere parameterized by polar coordinates) an additional step of reparameterization is necessary. This reparameterization step can additionally be used to apply distortions to the final image (e.g., zooming into particular areas, etc).

In implementation, projection of these viewing volumes occurs in one of two ways. The first and simplest is through ray tracing. Upon determining the parameter values u_0, v_0 that correspond to a particular pixel, the corresponding ray is the projector $Q_{u_0, v_0}(t)$. If the volume has not been linearly interpolated, the projector is not a linear curve and thus not suitable for the usual ray-tracing algorithm. There are two options to proceed, the first as Gröller [7] suggests, is to break the curve into linear pieces that approximate the curve and proceed with ray-tracing. The option suggested by Brosz et al. [3] is to limit our volumes to quadratic curves and compute object-quadratic intersections.

The other approach to implementation is to derive projection equations by finding the inverse of Q :

$$Q^{-1}(x, y, z) = (u, v, t).$$

While such an inverse is not always possible to derive analytically, there are several cases where analytical solutions are possible. Once a projection equation has been determined, projection can then be performed by vertex or geometry shaders in graphics hardware to achieve realtime rendering through rasterization. One last note is that rasterization makes use of linear interpolation, causing inaccuracies when triangles are deformed to curved shapes. These inaccuracies can be minimized with the use of high resolution models.

5 Using Flexible Projection in an Art Installation

As a part of a collaborative course between the Alberta College of Art and Design (ACAD) and the University of Calgary Hao Wang and John Brosz set out to complete a project combining artistic and scientific purposes. The project's goal became an art installation, *Perspectives*, that aimed to explore the 2008 confrontation between Tibetans and the Chinese authority. In particular, this project portrayed news images from distorted viewpoints representing the biases that exist in the media as well as between the involved parties. The produced images do not just represent pictorial organization but are instead intended to reflect different constructions of this issue and the world in general.

The scientific aspect of this project was the use of Flexible Projection to provide a variety of distorted outlooks in realtime. Flexible Projection was particularly useful in creating projections for this purpose in two ways. The first was that these projections could be adjusted to affect different depths of the scene differently; as a result, the background of the scene could feature a different kind of distortion than the foreground. Another aspect was that Flexible Projections can be animated by moving the projection surfaces over time. This added an extra impression of change and distortion. Together, this allowed creation of projections that provided very different outlooks into the 3d virtual world.

As our goal was to produce realtime renderings, we needed to derive projections equations (i.e., $Q^{-1}(x, y, z) = (u, v, t)$) allowing us to take an arbitrary point $p = (x, y, z)$ and then project it into device coordinates. As mentioned, for arbitrary volumes, finding Q^{-1} is difficult or impossible. To solve this, we limited ourselves to projection with viewing volumes where t (depth) can be easily calculated given p . With t determined in the viewing volume $Q(u, v, t)$ we can extract a parametric surface $Q_t(u, v)$. The next step is to find the specific values of (u, v) such that $Q_t(u, v) = p$. This leads to solve for two variables given three equations (one for each coordinate: x , y , and z). This can be solved for many parametric surfaces including the two that we made use of: parametrized spheres and bilinear patches.

To ensure we could calculate t , we relied upon two scenarios. The first is to define the viewing volume with surfaces that are pieces of spheres with a common center. Then t can be calculated by comparing the distance of p from this center, to the spheres' radii. The other scenario is to use surfaces that are planar and have a common surface normal. Then by projecting p onto this normal, and calculating distances between the surfaces, t can again be determined.

Sixteen different projections were developed. The only projection making use of parameterized spheres was a hemispherical projection (also known as fisheye) that was defined by a small hemisphere as the near surface and a larger one as the far surface. The other projections were defined by bilinear patches that, in turn, were defined by their corner points. The simplest projection recreated a perspective projection using two rectangular surfaces as shown in Figure 6(a).

The remaining fourteen projections were essentially chosen to be variations on perspective projection. This was a deliberate decision, made because it allowed the resulting projections to be different enough to be noticeable and eye-catching, but retain enough familiar characteristics to be intelligible. The modifications we made included: swapping the near and far surfaces to create an inverse perspective projection, moving the far surface relative to the near surface to create off-axis perspective projection, making the tops of the near and far surfaces narrower than the bottoms to expand the top of the image, and decreasing the height of the near surface as shown in Figure 6(b) to cause stretching of objects close to the near surface. More complex modifications were created by adding an intermediate surface between the near and far surfaces to allow for nonlinear projectors. In one projection, shown in Figure 6(c), this intermediate surface has been placed below the near and far surfaces causing projectors to travel downward before curving upward to reach

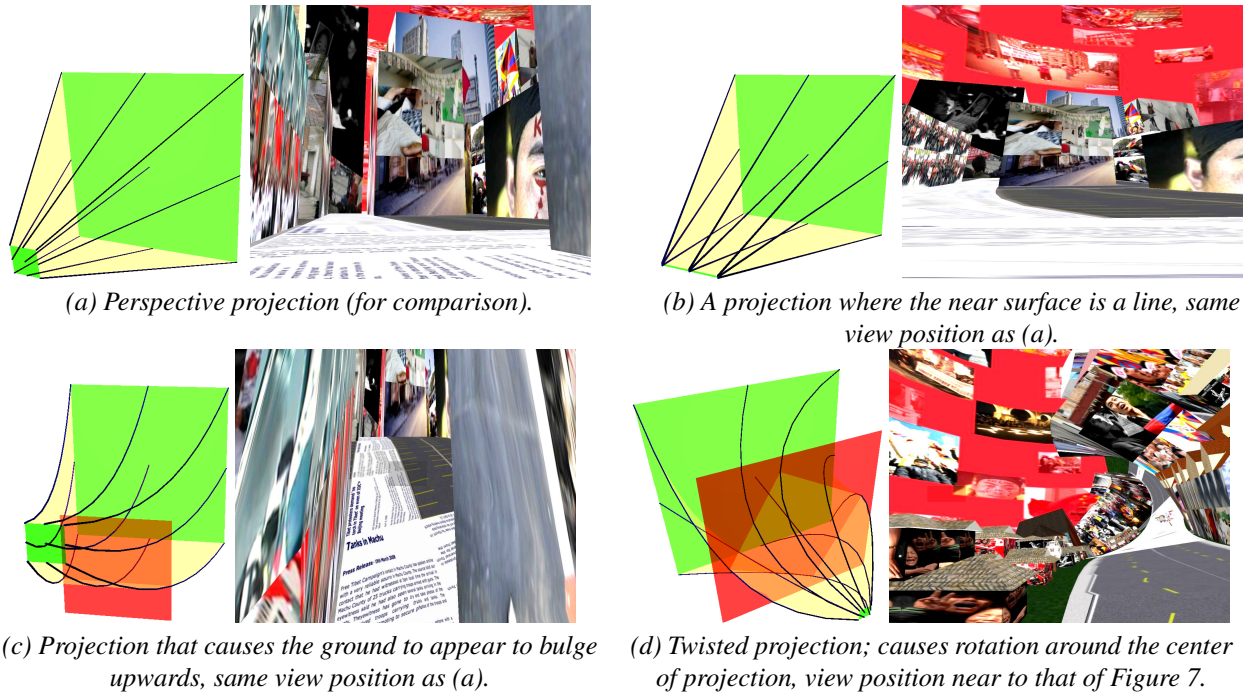


Figure 6: Projection diagrams (left) and resulting images (right). In the diagrams the near and far surfaces are green, the viewing volume is yellow, and projectors are shown with lines. The red middle surfaces in (c) and (d) are used to curve the volume.

the far surface. This produces an upward bulge in the midground of the produced images. A more exciting result is caused by introducing this intermediate surface and then rotating the far surface by ninety degrees. This, as shown in Figure 6(d), results in a twisting through the depth of the image.

Two animated projections were created. The first simply increased and decreased the size of the far plane over time. The second, shown in Figure 7, translated an intermediate surface around the center of projection.

The project itself made use of a game-like 3d environment with limited interaction. The environment, shown in Figure 8, was textured with news stories and images related to recent events in Tibet and related protests in Canada. Two modular ambient display screens [13] were used, each presenting a different virtual camera with a different projection of the environment. Each camera is then translated and rotated along a cyclical path through the environment. The interactive element of the projection was a button that randomly changed the projections used on both screens. Realtime rendering was at a rate of approximately 25 frames per seconds on two display devices of 1280 by 1024 resolution with an Intel Core2 6600 CPU, an NVIDIA 7800 video card, and 2 GB of system RAM. For the installation, the two displays with different projections were placed side by side at eye level. The installation setup is shown in Figure 9 and was presented at the ACAD graduation show in May 2008.

6 Conclusions and Future Work

In this work we have provided motivation for use of nonlinear projection. We have described specific examples where artists have used nonlinear projection. Additionally we have provided a short survey of current nonlinear projection techniques developed for computer graphics and shown how these technique could be




Figure 7: Keyframes of an animated projection. The middle rectangle that controls the curving of the volume is moved in circles causing the middle ground of the produced images to bend in a circular motion.



Figure 8: The virtual environment used in the project.



Figure 9: The Perspectives installation.

used to reproduce the artist created examples. We also have described a collaborative art and science project, describing how it made use of Flexible Projection to create a wide variety of visual effects.

In Section 5 we attempt to ensure that our projections are visually understandable by only introducing variations on perspective projection. This ensures that visual cues such as foreshortening and vanishing points are present (to some degree) in the resulting images. In particular we noted that the inverse perspective projection that causes a reversal of foreshortening, causing near objects to appear smaller than far objects, caused viewers of the installation difficulties. Further experimentation into restrictions or guidelines that allow freedom, while ensuring intelligibility of the resulting image, is a definite area for future study. Also of pertinence is that viewing of the animation (as opposed to a single image) as well as a known environment (such as city sidewalks) assists viewers in interpreting the images.

Another interesting area would be to further explore the use of animated projections; that is, projection where the projection surfaces move over time. One could imagine a camera rotation where near objects came into view first, followed by the appearance of objects further away. This could be achieved by rotating the far surface at a slower rate than the near surface. We were not able to achieve this in Flexible Projection due to the current limitation on calculating Q^{-1} . This is also an area of exploration for multi-camera projections.

Acknowledgements

We would like to thank the anonymous reviewers for their insightful and helpful comments as well as Ruth Hart-Budd for her editing assistance. This research was supported in part by the National Science and Engineering Research Council of Canada and the Informatics Circle of Research Excellence.

References

- [1] Aseem Agarwala, Maneesh Agrawala, Michael Cohen, David Salesin, and Richard Szeliski. Photographing long scenes with multi-viewpoint panoramas. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, pages 853–861. ACM Press, 2006.
- [2] Maneesh Agrawala, Denis Zorin, and Tamara Munzner. Artistic multiprojection rendering. In *Proc. of the Eurographics Workshop on Rendering Techniques 2000*, pages 125–136. Springer-Verlag, 2000.
- [3] John Brosz, Faramarz Samavati, Sheelagh Carpendale, and Mario Costa Sousa. Single camera flexible projection. In *Proc. of the 5th int. symp. on non-photorealistic animation and rendering*, pages 33–42. ACM, 2007.
- [4] Patrick Coleman and Karan Singh. Ryan: rendering your animation nonlinearly projected. In *Proc. of the 3rd int. symp. on non-photorealistic animation and rendering*, pages 129–156. ACM, 2004.
- [5] John P. Collomosse and Peter M. Hall. Cubist style rendering from photographs. *IEEE Transactions on Visualization and Computer Graphics*, 09(4):443–453, 2003.
- [6] Andrew S. Glassner. Cubism and cameras: Free-form optics for computer graphics. Technical Report MSR-TR-2000-05, Microsoft, January 2000.
- [7] Eduard Gröller. Nonlinear ray tracing: Visualizing strange worlds. *The Visual Computer*, 11(5):263–274, May 1995.
- [8] Patrick A. Heelan. Towards a new analysis of the pictorial space of vincent van gogh. *The Art Bulletin*, 54(4):478–492, December 1972.
- [9] David Hockney. *Secret Knowledge: Rediscovering the Lost Techniques of the Old Masters*. Viking Studio, 2 edition, 2006.
- [10] Paul Rademacher. View-dependent geometry. In *SIGGRAPH '99: Proc. of the 26th annual conference on computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1999.
- [11] P. Radmacher and G. Bishop. Multiple-center-of-projection images. In *SIGGRAPH '98: Proc. of the 25th annual conference on computer graphics and interactive techniques*, pages 199–206. ACM, 1998.
- [12] David Salomon. *Transformations and Projections in Computer Graphics*. Springer Verlag, 2006.
- [13] Ryan Schmidt, Eric Penner, and Sheelagh Carpendale. Reconfigurable displays. In *The ACM Conf. on Ubiquitous Computing, Workshop: Ubiquitous Display Environments*. ACM Press, 2004.
- [14] Karan Singh. A fresh perspective. In *Graphics Interface*, pages 17–24, May 2002.
- [15] John P. Snyder. *Flattening the Earth*. University of Chicago Press, 2 edition, 1997.
- [16] Matthias Trapp and Jürgen Döllner. Generalization of single-center projections using projection tile screens. In *Advances in Computer Graphics and Computer Vision - VISGRAPP 2008*, Communications in Computer and Information Science (CCIS). Springer, 2008.
- [17] Ernest W. Watson. *How to Use Creative Perspective*. Van Nostrand Reinhold Company, 1955.
- [18] Daniel Weiskopf. Four-dimensional non-linear ray tracing as a visualization tool for gravitational physics. *VIS*, 00:12, 2000.
- [19] G. Wyvill and C. McNaughton. Optical models. In *Proc. of the 8th international conference of the Comp. Graphics Society on CG International: comp. graphics around the world*, pages 83–93. Springer-Verlag, 1990.
- [20] J. Yu and L. McMillan. A framework for multiperspective rendering. In *15th Eurographics Symposium on Rendering (EGSR04)*, pages 61–68, 2004.
- [21] J. Yu and L. McMillan. General linear cameras. In *Computer Vision - ECCV 2004*, volume 2, pages 14–27. Springer, 2004.
- [22] Lihi Zelnik-Manor and Pietro Perona. Automating joiners. In *Proc. of the 5th int. symp. on non-photorealistic animation and rendering*, pages 121–131. ACM, 2007.