

Learning about Shadows from Artists

Elodie Fourquet

Abstract

Renaissance artists discovered methods for imaging realistic depth on a two dimensional surface by re-inventing linear perspective. In solving the problem of depth depiction, they observed how shadows project and volumes flatten in nature. They investigated how controlled illumination projects volumes onto walls, exploring the phenomena long before physical optics, such as the camera, existed. This paper specifically examines artists' constructions for depicting shadows, a 3D double projection problem that artists solved completely within two dimensions. The larger goal is to develop new computational methods for creating 3D perceptions without having to leave the 2D canvas. Those methods have potential application in constructing user interfaces, in 2D image compositing and in simultaneous 2D/3D composition.

This paper develops geometric constructions for casting shadows onto planar surfaces, adapted from artists' methods. Their algebraization for integration into imaging software is demonstrated, and their optically accuracy is shown. Finally, resulting images are included, along with a discussion of limitations.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture; J.4 [Social and Behavioural Sciences]: Psychology;—J.5 [Arts and Humanities]: Fine arts—

1. Introduction

Renaissance painters opened up pictorial space into depth by re-inventing linear perspective. They gave objects 3D relationships in depth while controlling picture plane composition. Achieving this double challenge, they created a large collection of geometric constructions realisable with a straight-edge, in effect developing projective geometry.

Essential for easy depth perception and construction is a strong ground plane, often a tiled floor, which roots objects in depth. An earlier paper [Fou08] showed how the ground plane can be combined with a compositional grid to integrate 2D and 3D composition. This paper examines shadow placement, another technique painters used to strengthen the perception of depth. Research in picture perception shows that shadows enhance understanding of the relative position of objects [CL89, LGTT09]. Furthermore, they are seldom the focus of vision and admit some freedom in depiction, which permits to simulate them convincingly with solely 2D computations.

The constructions painters use to create correct shadow placement perform perspective calculations geometrically. To replicate painters' constructions within computer graph-

ics tools two important specifications are needed: the attachment point of an object, and the light vanishing point. These two points on the picture plane provide enough depth information for solving the 3D double perspective projection, which determines shadow location and shape. Thus, the paper demonstrates that it is unnecessary to work in 3D in order to create accurate shadows. Indeed, the paper introduces 2D matrix transformations that algebraically replicate the geometric constructions of artists.

This research is important because artists creating 2D images prefer to manage composition within 2D environments, such as Photoshop or Illustrator, where for realism they use construction lines to get correct depth, sometimes at great cost [Hus10]. Thus, since compositing and editing images in 2D is ubiquitous throughout the graphic arts, the results presented in this paper provide interesting opportunities for better supporting depth in 2D tools.

The paper begins by describing typical artists' constructions for placing shadows, which are more in the nature of rules of thumb than algorithms. Therefore, those descriptions are complemented with the special considerations needed to encode artists' tacit knowledge and practice within the computer. It then shows how to algebraize the most ba-

sic construction. A full implementation is used to create a variety of images with shadows, illustrating the range of the methods and how they anchor objects in space. The main contribution of this paper is formalizing and implementing artists' geometrical rules for calculating shadows, and demonstrating that the formalization is optically accurate.

2. Previous Work

Despite many well-featured commercial 2D image creation environments, many of which, such as Photoshop or Java2D, support a variety of affine transformations, there is little research on manipulating 3D perspective in 2D. The explanation, most likely, is the development of computer graphics primarily by scientists and engineers. The earliest formalizations of geometry in computer graphics were 3D, modelling using affine transformations with a single perspective transformation for projecting to the 2D image. Image creation applications based on 3D geometry have been technically successful, but, based on poor uptake by visual artists, they seem to offer inadequate support for the creation of artistically interesting images.

More specifically, there is negligible support for realistic shadows in 2D graphics. The single exception is $2\frac{1}{2}$ D user interfaces, which use drop shadows extensively for separating interface components in depth [Wil91]. The techniques for making them are, however, *ad hoc*, and unconnected with explicit depth or illumination.

There exist several studies of the effectiveness of manipulating shadows as a method for controlling illumination: for example, users dragging shadow volumes [PF92] or shadows [PTG02] in the image. Direct manipulation of lights in the image seems most effective. Other research adds shadows to existing images, allowing users to add shadows from non-existent sources of illumination [Bar97] or to add shadows cast by non-existent objects [PFWF00, KP09]. Notable among this research are algorithms that stylize shadow mattes calculated from 3D illumination models [DCFR07]. This research follows the practice of artists, who eliminate inessential details of shadows, keeping only details that help the user to perceive the object.

There is a large body of recent research that uses scenes derived from 2D material as the source of new scenes or as additions to other scenes, which is too extensive for review. Occasionally – Shesh et al. [SCRS09] is one such example – shadows are added. However, the methods are explicitly 3D, and operate on partial 3D information reconstructed by estimating the homography that created the scene, using a computer vision technique [HZ04]. The research described here uses homographies, but for a quite different purpose, for image creation rather than scene reconstruction, using image plane illumination key points. These transformations go under a wide variety of names, such as projective transformations, projections [Sti05] or projectivities [Cox74]. We call them collineations, the term that seems in widest use.

3. Artists' Constructions




Figure 1: Three illustrations from chapter XV, *Perspective of Shadows*, of Cole's manual.

Artists draw shadows by placing light sources in the image, then drawing construction lines from them. This section describes constructions that draw optically accurate shadows using only a straight-edge. The constructions are geometrical solutions of double projection equations.

Published early in the twentieth century, Rex Vicat Cole's manual [Col21] collects many illustrations of artists' constructions, including ones for calculating shadow geometry. Figure 1 shows three illustrations of shadow constructions. Notice the vanishing point of each light source, which minimally encodes enough information to define projections from the lights. Cole's descriptions are incomplete: he expects the artist to complement them with tacit knowledge. Implementing the constructions the tacit must be made explicit, one contribution of this paper.

3.1. Placing Lights in 2D

Ground plane geometry is defined by two points. The *scene principal vanishing point*, at eye level in the direction of the viewer's gaze, sets the height of the horizon. The *scene distance point*, on the horizon often on the frame, defines the viewer's distance from the scene [Fou08].

Lights, which are centres of projection for shadows [Bli88], also have *light vanishing points*, which are again centres of projection, directly below the light on the ground plane. Artists position a light in the image, then draw a vertical line terminating on the ground plane at the light vanishing point, defining the light's depth. In contrast with lights near the image plane, distant lights, like the sun, have vanishing points on the horizon.

In two cases lights cannot be located on the canvas straightforwardly. In the first case the light is in front of the viewer beyond the frame, high overhead or far to the side. Such lights produce minimally foreshortened shadows and parallel construction lines from the direction of the light source correctly describe the shadows. In the second case the light is behind the viewer. In this configuration artists use a *pseudo-light*, positioned approximately where the light projects through the viewer's eye onto the image plane. With it they can construct shadows produced by illumination from behind. Indeed, in practice, the artist does not project exactly the light through the eye: the light vanishing point is positioned on the ground plane with the pseudo-light below it,

and the two positions are moved until the shadows enhance the overall composition.




Figure 2: Illumination shading a point on the ground floor. Both the illumination vanishing point and the point attachment point must be known for the construction.

3.2. Shadows on the Ground Plane

Objects in the image need *attachment points*, which do for them what light vanishing points do for lights. They are positioned on the ground plane directly below each object, defining its depth. Lines drawn from a light vanishing point through an object's attachment point show the direction in which the shadow is cast.

The basic shadow construction is illustrated for a point in Figure 2. A *ground ray* is drawn from the light vanishing point through the attachment point. Then, a *light ray*, is drawn from the light source through the shadow-casting point. The point's shadow falls on the ground plane at their intersection. Of course, if the point happens to be on the ground plane, it coincides with both its attachment point and its shadow.

While images consist of points: the points are grouped into representations of objects. The simplest object is a rectangular face, often a wall. Cole shows shadows of walls for different light placements. The shadows are quadrilaterals with vertices defined by the shadow points of the corners. Walls on the ground plane have two corners on the ground, coinciding with two shadow vertices. The shadow points of the top corners are the other vertices of the shadow. If, in defiance of gravity, the wall hovers in mid-air, four shadow points must be found, using four light rays, but only two ground rays. Figure 3 illustrates these constructions. Note that artists do not think of attachment points, but simply draw ground rays by visualizing the object's position. This paper introduces attachment points because they are required for computing shadows, and the user must position them explicitly.




Figure 3: Difference in the shadow placement relative to the position of the attachment point with respect to the object.

3.3. More Complex Shadows

Realistic painters often assemble an image from drawings called cartoons, which are placed within a 2D representation of architecture or nature. The drawn content then requires shadows. Precise shadows are impossible, because only one silhouette is available, but artists create shadows that are very successful all the same. The construction for shadows of cartoons is described below in some detail, so the reader can get a feeling for the challenges of implementing construction lines in imaging software.

3.3.1. Constructions for Panels

Cole provides constructions for planar surfaces such as arches. The most relevant points have their shadows projected, and the artist interpolates between them free-hand. To automate such a construction 2D renderings of objects are mapped onto *panels* [Fou08], which are transparent rectangles. Panels work best when the view of the content is robust, usually a central perspective of a 3D object oriented to a canonical view [FCM07].

Each panel has an attachment point, aligned with its central vertical axis. The attachment point encodes the panel's depth from the image plane and its distance above the ground plane. Almost always, one attachment point suffices because the panel by default is parallel to the image plane.

The shadow geometry of panels is simple. First, a mask is automatically constructed from the panel content. This mask is mapped onto the ground plane by the collineation defined by the panel vertices. (Artists usually transfer a complex shadow into the panel shadow using a grid: 2D texture mapping is an automated solution.) A typical result is shown on the left of Figure 4. The shadow differs slightly from a 2D rendering of full 3D content. However, artists have found such shadows acceptable for centuries, and the shadow conserves all qualities that are important for vision.

The above construction is simple, but can fail. For example, if a vertical edge of a panel is aligned with a light, intersections between light rays and ground rays are undefined. An artist immediately notices the problem and uses equivalent oblique rays to find an intersection that is extrapolated horizontally. Our computational version of the artist's construction always uses oblique rays and provides robust results.




Figure 4: Constructions for ground shadows. On the left, the illumination is a sun (horizon vanishing point). On the right, baseline intersection points from the ground rays are used to only consider the region that shades behind the image plane.

Another case needing a more robust construction occurs when the shadow of a panel extends beyond the picture frame. The simple algorithm – draw the entire shadow and clip to the frame – is impossible in general because shadows become infinitely long as the distance of the light from the ground plane approaches the height of the panel. Thus, the shadow must be cut off at the frame initially. The intersections of ground rays with the frame mark the extend of the visible shadow. Projected back to the panel they are the vertices of the polygon containing the content to be projected.

A more extreme case occurs when the distance of the light above the ground plane is less than the height of the panel. Then, the ground ray/light ray intersections are between the light's vanishing point and the horizon. The shadow then incorrectly lies on the lighted side of the panel and its edges intersect one another. This case is easy to identify because the light vanishing point lies inside the shadow. When this condition is detected the construction described in the previous paragraph gets the shadow correct.

3.3.2. Other Shadows

In addition to completing Cole's constructions it was necessary to create many new constructions for more complex shadow geometries. These include shadows cast on planes parallel to the ground plane, which requires the attachment points and the light vanishing point to be defined on the shadow plane (Figures 5 and 6), shadows cast on vertical surfaces parallel to the image plane, shadows split between two surfaces (Figures 7 and 8), shadows on planes perpendicular to the floor and the image plane (Figure 8), and shadows of rotated panels (Figure 8).

3.4. Summary

This section describes by example a formalization of artists' constructions for shadows as algorithms, which successfully complete them in accord with artists' intuitions, and provides methods for creating many new constructions. Everything described was implemented using only line intersec-




Figure 5: Panel projection on an elevated ground. Only the panel part above the height defined by the volume is considered. Ground rays and light vanishing point on the elevated surface need to be used.




Figure 6: Constructions for shadows on an elevated ground. The left image shows how the projected shadow quad is evaluated. The right image shows how the elevated light vanishing point is evaluated from the volume perspective.

tion calculations, to prove that all these algorithms are implementable by artists working on the image plane, and without measurement. It shows that artists can work exclusively in 2D while creating the projection of an optically correct 3D volume.

The implementation is important as ground truth against




Figure 7: Pseudo-illumination creating a ground shadow that continues onto a face that is parallel to the image plane, and is visible to the viewer. Light rays are used to map the relevant regions.




Figure 8: Rotated panels creating shadows on the ground plane and on a lateral face of a volume. Parallel lines on the rotated panel converge on the horizon.

which the algebraic 2D formalism introduced in the next section can be tested. It reformulates the basic construction in terms of 2D projective geometry, which is easier to implement, and which is well-adapted to current graphics hardware.

4. Mathematical Theory

2D collineations, represented as homogeneous matrices, perform the geometric computations described in the previous section more flexibly and efficiently than simulated Euclidean geometry. The collineations demonstrate the generality of artists' constructions, ease the implementation, and integrate artists' constructions with existing 2D graphics software. This section presents the 2D matrix corresponding to the geometric construction of shadows on the ground. Then, it is shown that the 2D construction is optically correct by solving the equivalent 3D problem. To show the technique the mapping from a panel to its shadow on the ground plane is derived in detail, but collineations that perform other shadow projections are omitted owing to restricted space.

4.1. Collineations for Ground Shadows

This section calculates homogeneous matrix representations of collineations for the geometric shadows of Section 3 in the artists' coordinate system: origin at the centre of the baseline, x increasing to the right, y upward. The scene vanishing point is centered on the horizon, at $(0, \mathbf{E}_h)$. The scene distance point is at $(\pm \mathbf{E}_d, \mathbf{E}_h)$. The light source is labeled \mathbf{I} , its vanishing point \mathbf{J} , shadow casting points \mathbf{P} , and their attachment points \mathbf{Q} . The notation is illustrated in Figure 2.

The 2D formalism is based on a transformation that maps a point \mathbf{P} on a panel with attachment point \mathbf{Q} , to \mathbf{S} , its shadow on the ground plane. The shadow lies at the intersection of two lines: $\mathbf{L}_1(t) = \mathbf{P} + t(\mathbf{I} - \mathbf{P})$ and $\mathbf{L}_2(s) = \mathbf{Q} + s(\mathbf{J} - \mathbf{Q})$. By definition, $\mathbf{Q}_x = \mathbf{P}_x$ and $\mathbf{J}_x = \mathbf{I}_x$. Thus, $s = t$, and

$t = (\mathbf{Q}_y - \mathbf{P}_y) / (-\mathbf{P}_y + \mathbf{I}_y - \mathbf{J}_y - \mathbf{Q}_y)$, resulting in

$$\begin{aligned} \mathbf{S}_x &= \frac{\mathbf{P}_x \mathbf{I}_y - \mathbf{P}_x \mathbf{J}_y - \mathbf{P}_y \mathbf{I}_x + \mathbf{I}_x \mathbf{Q}_y}{-\mathbf{P}_y + \mathbf{I}_y - \mathbf{J}_y + \mathbf{Q}_y} \quad \text{and} \\ \mathbf{S}_y &= \frac{-\mathbf{P}_y \mathbf{J}_y + \mathbf{I}_y \mathbf{Q}_y}{-\mathbf{P}_y + \mathbf{I}_y - \mathbf{J}_y + \mathbf{Q}_y}, \end{aligned}$$

a Möbius transformation with real coordinates. Thus, the following matrix represents the transformation that maps the coordinates of \mathbf{P} to the coordinates of its shadow \mathbf{S} .

$$\begin{pmatrix} \mathbf{S}_x \\ \mathbf{S}_y \\ 1 \end{pmatrix} \sim \begin{pmatrix} \mathbf{I}_y - \mathbf{J}_y & -\mathbf{I}_x & \mathbf{I}_x \mathbf{Q}_y \\ 0 & -\mathbf{J}_y & \mathbf{I}_y \mathbf{Q}_y \\ 0 & -1 & \mathbf{I}_y - \mathbf{J}_y + \mathbf{Q}_y \end{pmatrix} \begin{pmatrix} \mathbf{P}_x \\ \mathbf{P}_y \\ 1 \end{pmatrix},$$

where \sim means 'is the same point as'. The light position (\mathbf{I} and \mathbf{J}) are parameters of the transformation, as is the panel attachment point (\mathbf{Q}). Therefore, the matrix does not vary from point to point, only from panel to panel. The mapping depends only on points in the image plane: no 3D calculations are performed.

Collineations implementing the other shadow constructions mentioned in Section 3, are derived using similar techniques, together with the collineation defining the scene perspective, as encoded by the tiled floor construction. Calculating 3D properties using only 2D collineations is possible because the attachment points and light vanishing points encode just enough 3D information, and encode it so it is accessible in 2D. This important observation is a contribution of the paper.

4.2. 3D Equivalence

The previous subsection derived a homogeneous matrix representing the collineation for geometric constructions of ground shadows. This one proves that the constructions are optically accurate.

There is a direct simple proof. Examining Figure 9, we visualize in 3D the scene it portrays. A line, from the light source through a shadow-casting point, intersects the ground plane at the shadow. A second line, from the light vanishing point through the attachment point, lies on the ground plane and meets the first line where it intersects the ground plane. An intersection is guaranteed because the two lines are coplanar. Projection to the image plane maps preserves intersections, so that the shadow on the image plane lies at the intersection of the projected eye and ground rays.

Such reasoning, plus the image before their eyes, convinces artists that the construction is correct. However, algebraic calculation is the medium of computer graphics, and leads more directly to algorithmic expression. Therefore, a second proof calculates the projection geometry in 3D and derives the corresponding 2D collineation. The presentation is in the style of Blinn [Bli88].




Figure 9: The shadow casting geometry is shown in an orthogonally-projected side view. \mathbf{i} is the light source; \mathbf{j} its vanishing point; \mathbf{p} is the shadow-casting point; \mathbf{q} its attachment point. The shadow falls on the ground at \mathbf{s} . These points are projected to the image plane through the viewer's eye \mathbf{E} . Their images have the same label, but uppercase.

The 3D geometry is shown in Figure 9. We show that the projected shadow point, \mathbf{S} , can be calculated using only the projected points, \mathbf{I} , \mathbf{J} , \mathbf{P} and \mathbf{Q} . To do so we calculate the 3D points from projected points, and substitute them into the projected location of the shadow. In 3D the shadow falls at $\mathbf{s} \sim \Sigma \mathbf{p}$, where

$$\Sigma = \begin{pmatrix} -\mathbf{i}_y & \mathbf{i}_x & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & \mathbf{i}_z & -\mathbf{i}_y & 0 \\ 0 & 1 & 0 & -\mathbf{i}_y \end{pmatrix}.$$

In Σ the row of zeros occurs because the shadow is cast on the floor, $y = 0$. The shadow is projected to the image plane, using the projection centre $\mathbf{E} = (\mathbf{E}_x, \mathbf{E}_h, \mathbf{E}_d)$, with the matrix

$$\Pi = \begin{pmatrix} -\mathbf{E}_d & 0 & \mathbf{E}_x & 0 \\ 0 & -\mathbf{E}_d & \mathbf{E}_h & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -\mathbf{E}_d \end{pmatrix}$$

so that $\mathbf{S} \sim \Pi \Sigma \mathbf{p}$. We now calculate \mathbf{p} and \mathbf{i} . Indeed, \mathbf{p} is derived from the following matrix Γ ,

$$\mathbf{p} \sim \Gamma \begin{pmatrix} \mathbf{P}_x \\ \mathbf{P}_y \\ \mathbf{Q}_y \\ 1 \end{pmatrix} = \begin{pmatrix} -\mathbf{E}_h & 0 & \mathbf{E}_x & 0 \\ 0 & -\mathbf{E}_h & \mathbf{E}_h & 0 \\ 0 & 0 & \mathbf{E}_d & 0 \\ 0 & 0 & 1 & -\mathbf{E}_h \end{pmatrix} \begin{pmatrix} \mathbf{P}_x \\ \mathbf{P}_y \\ \mathbf{Q}_y \\ 1 \end{pmatrix}.$$

In the product $\Pi \Sigma \Gamma$ we substitute $\mathbf{i} = \Gamma(\mathbf{I}_x, \mathbf{I}_y, \mathbf{J}_y, 1)^T$ and place the eye \mathbf{E} in the world frame at $\mathbf{E}_x = 0$. (With this projection a pseudo-light appears on the image plane by projecting it through the eye of the viewer.) Thus, we obtain

$$\mathbf{S} \sim \begin{pmatrix} \mathbf{I}_y - \mathbf{J}_y & -\mathbf{I}_x & \mathbf{I}_x & 0 \\ 0 & -\mathbf{J}_y & \mathbf{I}_y & 0 \\ 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & \mathbf{I}_y - \mathbf{J}_y \end{pmatrix} \begin{pmatrix} \mathbf{P}_x \\ \mathbf{P}_y \\ \mathbf{Q}_y \\ 1 \end{pmatrix},$$

which can be simplified to the transformation obtained in Section 4.1, proving that the artists' construction is optically correct.

5. Implementation and Results

The techniques of this paper have been implemented as a prototype application, with an interaction model that would be useful for artists roughing out volumes in a 2D realistic image, or for stage designers planning the use of space on a stage as seen by the audience. The goal of the prototype, however, is not to provide a practical tool, but to provide a test-bed for the formalism.

5.1. Implementation

The prototype is implemented in Java, using the Java2D API. Both geometric construction methods and 2D collineations are represented as 2D homogeneous matrices. This section highlights non-standard aspects of the implementation.

5.1.1. Panels and Shadows

Drawn on each panel is a 2D image, its content, which is created off-line. The image is often the central axis projection of a 3D object, but may be any 2D content. The content is opaque, the rest of the panel transparent. Shadows use the content as a mask. Classes implementing non-affine matrices and filters map the shadow colour through the mask onto the shadow plane. Shadows that coincide are combined using illumination attributes. The combined shadow can be given a smoke-like outline or made semi-transparent.

The prototype also includes rectangular blocks, which are commonly used to rough out volumes. Blocks are tile-aligned and can be stacked. Panels placed on block faces are scaled so that panel corners coincide with face corners.

5.1.2. Occlusion by Clipping

A painter's algorithm combined with 2D clipping and compositing provides correct occlusion. It uses geometry-guaranteed patterns in block/block and block/panel occlusion among unit blocks. Scene elements are drawn row by row, from the horizon to the baseline. Within a row they are drawn from the frame toward the mid-line of the image. Within a tile, blocks, which may be stacked, are drawn from base to horizon, with any part of the block above the horizon drawn from its top down to the horizon. Except for infrequent floating point round-off artifacts, visible in Figure 8, rendering is flawless, including shadows and panels on block faces.

5.1.3. The Interface

The composition interface is drag-and-drop for panels and for points that define the geometry, like vanishing points, attachment points, and light sources. Each panel can be scaled or translated by dragging. Note that only the y-coordinate of panel attachment points is editable, the x-coordinate being determined by panel location. From a selected set of tiles, blocks are extruded upward using the scroll wheel. The prototype runs at interactive frame rates on a mediocre Linux box, 3 GHz Pentium 4 with 1 Gbyte of memory.




Figure 10: Shadow effect from a rotating panel. Bird shadow with fuzzy edges.




Figure 11: Shadows of volumes and of a vase.

5.2. Results

All colour figures in the paper were created using the prototype. They illustrate important aspects of the shadow algorithms. In some cases construction lines, dashed and in colour, were manually added to guide the eye.

The left pair of images in Figure 10 compare shadows cast on the ground by a panel parallel to the image plane and by the same panel slightly rotated. Owing to the shadow, the woman is perceived as rooted firmly on the ground, even though her feet are awkwardly drawn. The gap between elbow and body is more visible in the shadow than on the panel. The rotated panel/woman produces an equally satisfactory shadow. Shadows can be surprising. The right image shows the shadow of a dove that is more birdlike than the bird itself.

Figure 11 shows shadows of volume faces. The upper image shows the shadow from an off-canvas directional light, its direction determined by the oblique line. The lower left image is the same scene lit with a back light on the left. In those images the shadows are semi-transparent: idempotent combination ensures that opacity does not add. The lower right image is a vase and its shadow, the vase derived from a painting by Vanderlyn. DeCoro et al. used it to illustrate a method for artistic simplification of 3D shadows [DCFR07]. Projecting from a panel, as artists generally do, gives a similar simplified result, but as the default case.

Figure 12 shows more complex images created using the prototype. The leftmost image shows shadows cast by a pseudo-light, pointing toward the light vanishing point. The centre image has eight panels and two volumes, with shadows that create strong visual relationships between panels and volumes. Finally, the rightmost image is lit from the centre of the picture, with shadows that diverge in all directions.

6. Discussion

For centuries artists and designers have manipulated shadows in 2D using rules of thumb acquired by experience. These rules can be completed and formalized. Essential for formalization are attachment points of objects and vanishing points of lights, which encode in the image just enough depth information to determine shadow geometry.

Lights can be manipulated in 2D by dragging them or their vanishing points. Recently, Kerr and Pellacini [KP09] showed that novice users best manipulate lighting directly, by moving light sources in 3D, or indirectly, by dragging shadows and highlights in 2D, with the latter preferred. Most likely this result shows that the familiarity of 2D overcomes the disadvantage of indirection. If this interpretation is correct interacting directly with lights in 2D is even better, a vindication of learning from artists.

Also worth considering is shadow perception. Working directly in 2D limits the quality of shadows: is the artists' practice an unavoidable compromise of image quality? Or are full 3D shadows overkill, computing perceptually unimportant details? Vision scientists [CL89] suggest that only a few aspects of shadows are important for vision. For example, shadows attached to an object root it on the ground plane, while detached ones place it in the air. Figure 10 shows the difference. Consistent shadow orientations are important [LGTT09]. Average orientation indicates the direction of the light; variation of orientations the distance to the light. Both are preserved by 2D shadows. Other important features, shadow edges, the relative lightness of shadowed and non-shadowed regions [CL89], are also correct. Finally, shadows give information about the geometry of the surface on which they are cast. When 2D shadows extend over several surfaces, as illustrated in Figure 12, they clearly strengthen perception of surfaces.

In summary, while the shadows presented in this paper have limitations, they capture the visually salient aspects of shadows. Presumably, other aspects are less important because shadows are rarely the focus of attention.

7. Conclusion

The goals of this research were three: to capture in algorithms the tacit knowledge with which artists complete their 2D constructions of shadows; to show the algorithms to be equivalent to 2D collineations; and to demonstrate the constructions to be optically correct. These objectives have been




Figure 12: Compositions demonstrating interactions between volumes, panels and shadows.

achieved, providing essential support for a host of applications. The most general are novel tools that will allow an artist to compose a realistic image in 2D with good control over the position in depth of its elements. For example, improving composition by moving lights to control shadows is advantageous, especially when done within the 2D image.

Application to other problems is equally important. For example, one can build a consistent lighting model in a $2\frac{1}{2}$ user interface without working in 3D. Visualization researchers have said that 2D design of 3D visualizations will be equally beneficial. The formalism might improve the notoriously difficult interfaces of 3D modelling tools. These few examples demonstrate the many benefits of exploiting the hard-won knowledge of practicing artists.

Acknowledgements – This work would not have been possible without the support, dedication and patience of Bill Cowan and Stephen Mann. Bill Cowan also improved the English of this manuscript.

References

- [Bar97] BARZEL R.: Lighting controls for computer cinematography. *J. Graph. Tools* 2, 1 (1997), 1–20.
- [Bli88] BLINN J.: Me and my (fake) shadow. *IEEE Comput. Graph. Appl.* 8, 1 (1988), 82–86.
- [CL89] CAVANAGH P., LECLERC Y. G.: Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance.* 15, 1 (February 1989), 3–27.
- [Col21] COLE R. V.: *Perspective: The Practice & Theory of Perspective as Applied to Pictures, with a Section Dealing with Its Application to Architecture.* Seeley, Service & Co, 1921.
- [Cox74] COXETER H. S. M.: *Projective Geometry.* Springer-Verlag, New York, 1974.
- [DCFR07] DECORO C., COLE F., FINKELSTEIN A., RUSINKIEWICZ S.: Stylized shadows. In *NPAR '07: Proceedings of the 5th international symposium on Non-photorealistic animation and rendering* (2007), pp. 77–83.
- [FCM07] FOURQUET E., COWAN W., MANN S.: On the empirical limits of billboard rotation. In *APGV '07: Proceedings of the 4th symposium on Applied perception in graphics and visualization* (New York, NY, USA, 2007), ACM, pp. 49–56.
- [Fou08] FOURQUET E.: Composition in Perspectives. In *Computational Aesthetics in Graphics, Visualization, and Imaging* (2008), pp. 9–16.
- [Hus10] HUSLEY K.: Adobe Illustrator & Photoshop Tutorials & Lessons. http://www.khulsey.com/demo_1howto.html, accessed on April 2010, 2010.
- [HZ04] HARTLEY R. I., ZISSERMAN A.: *Multiple View Geometry in Computer Vision*, second ed. Cambridge University Press, ISBN: 0521540518, 2004.
- [KP09] KERR W. B., PELLACINI F.: Toward evaluating lighting design interface paradigms for novice users. *ACM Trans. Graph.* 28, 3 (2009), 1–9.
- [LGTT09] LOVELL P. G., GILCHRIST I. D., TOLHURST D. J., TROSCIANKO T.: Search for gross illumination discrepancies in images of natural objects. *Journal of Vision* 9, 1 (1 2009), 1–14.
- [PF92] POULIN P., FOURNIER A.: Lights from highlights and shadows. In *SI3D '92: Proceedings of the 1992 symposium on Interactive 3D graphics* (1992), pp. 31–38.
- [PFWF00] PETROVIĆ L., FUJITO B., WILLIAMS L., FINKELSTEIN A.: Shadows for cel animation. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques* (2000), pp. 511–516.
- [PTG02] PELLACINI F., TOLE P., GREENBERG D. P.: A user interface for interactive cinematic shadow design. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (2002), pp. 563–566.
- [SCRS09] SHESH A., CRIMINISI A., ROTHER C., SMYTH G.: 3d-aware image editing for out of bounds photography. In *GI '09: Proceedings of Graphics Interface 2009* (Toronto, Ont., Canada, Canada, 2009), Canadian Information Processing Society, pp. 47–54.
- [Sti05] STILLWELL J.: *The Four Pillars of Geometry.* Springer, 2005.
- [Wil91] WILLIAMS L.: Shading in two dimensions. In *GI '91: Proceedings of Graphics Interface 1991* (1991), pp. 143–151.