Computational Aesthetics in Graphics, Visualization, and Imaging (2010)
O. Deussen and P. Jepp (Editors)

Pannini: A New Projection for Rendering Wide Angle Perspective Images

Thomas K. Sharpless[1], Bruno Postle, and Daniel M. German[2]

[1]tksharpless@gmail.com [2]dmg@uvic.ca, Dept. of Computer Science, University of Victoria.

Abstract

The widely used rectilinear perspective projection cannot render realistic looking flat views with fields of view much wider than 70°. Yet 18th century artists known as 'view painters' depicted wider architectural scenes without visible perspective distortion. We have found no written records of how they did that, however, quantitative analysis of several works suggests that the key is a system for compressing horizontal angles while preserving certain straight lines important for the perspective illusion.

We show that a simple double projection of the sphere to the plane, that we call the Pannini projection*, can render images 150° or more wide with a natural appearance, reminiscent of* vedutismo *perspective. We give the mathematical formulas for realizing it numerically, in a general form that can be adjusted to suit a wide range of subject matter and field widths, and briefly compare it to other proposed alternatives to the rectilinear projection.*

1. Introduction

Modern photographic and computer technology make it easy to acquire wide angle images of the world, even up to fully spherical 360° images, and to synthesize wide images of imaginary worlds. But we still have trouble displaying those images on a flat surface [GBDL*07]. The rectilinear perspective projection, universally used for rendering realistic images at moderate fields of view, is simply not suitable for very wide angle views. Its magnification increases too rapidly with the view angle, with the result that objects near the edges of wide field images appear too large, and are stretched radially. These effects are known as rectilinear perspective distortion.

Modern manuals of perspective drawing suggest that to avoid rectilinear perspective distortion, fields of view should be limited to no more than 60 to 70°. Renaissance artists observed smaller limits – 30 to 40° [Kub86]. Today the "market limit" on the rectilinear field of view appears to be near 90°. Drawings and photographs that wide are published regularly. The widest broadcast television lenses cover 94°, and some rectilinear still camera lenses can take pictures over 100° wide. It is easy to see the perspective distortion in such wide images. The availability of tools to create photographic panoramas (which can have fields of view of up to 360°) increases the need for realistic alternatives to the rectilinear projection.

Certain artists of the Baroque period (1650-1800) produced pictures with wide fields of view, in what looks like correct perspective, without any visible sign of rectilinear perspective distortion. Their style, called *vedutismo* in Italian, view painting in English, is highly recognizable. However there are no written records of how these remarkable perspectives were constructed.

In this paper we describe a simple but effective alternative to the rectilinear perspective projection, derived from an analysis of *vedutismo* perspective, that we call the Pannini projection.

2. Related Work

Digital panoramic photographers now use many alternative projections to render wide views of their work [GdGP07]. Most of those cannot be considered replacements for the rectilinear projection, because they produce images that violate our sense of correct perspective by curving lines we expect to be straight. However several methods have been developed that can, in favorable circumstances, render wide views that resemble rectilinear perspectives.

Zorin and Barr [ZB95] described the first such method in 1995. It uses nonlinear optimization to find a locally varying

DOI: 10.2312/COMPAESTH/COMPAESTH10/009-016

(a) Interior San Pietro (b) Interior Santa Maria Maggiore (c) Interior Pantheon ’32 (d) Interior Pantheon ’34

Figure 1: *Gian Paolo Pannini, (a)* Interior San Pietro *c. 1754 (National Gallery of Art, Washington) horizontal field of view c. 96°; (b)* Interior Santa Maria Maggiore *c. 1753 (Hermitage Museum, St. Petersburg) hfov c. 74°; (c)* Interior of the Pantheon *c. 1732 (private collection) hfov c. 83°; (d)* Interior of the Pantheon*, c. 1747 (National Gallery of Art, Washington) hfov c. 57°*

mixture of two projections, that minimizes a statistical measure of ‘perceptual distortion’. One of the projections preserves straight lines, the other preserves local shape. The optimization balances estimates of local and global distortion. The method is effective at removing perspective distortion from wide angle (90 - 100°) rectilinear images. This work demonstrated that transformations which respect perceptual, as opposed to purely geometrical, rules can produce realistic looking wide angle images.

Recently, Carroll, Agrawala and Agarwala [CAA09] used numerical optimization to compute general image warping transformations that straighten a set of lines, designated by the user, while minimizing some measures of ‘distortion energy’ including terms similar to Zorin and Barr’s. Not surprisingly, given the very specific nature of the problem and the very general nature of the solution, the optimization procedure is complex and highly tuned. Nevertheless in many cases it was able to transform fish eye and panoramic images into satisfactory perspective views.

Zelnik-Major *et. al.* [ZMPP05] take a more direct approach, combining two or more standard projections to render different parts of a single image. Their “multi-plane” method divides a wide cylindrical image into vertical panels, and renders a rectilinear projection centered on each panel. Conceptually the panels are hinged together at the edges to form a continuous viewing screen. A human user has to position the boundaries and set the panel angles. Good results can be obtained when the subject matter allows the panel boundaries to be well hidden; otherwise they may present corners that do not exist in the original. This method is similar to artists’ techniques for combining multiple points of view in one perspective. Several panorama stitching programs now offer this and other composite projections, collectively know as ‘hybrid’ projections, each of which works well for a limited range of subjects.

3. Analysis of *vedutismo* Perspective

We first encountered the Pannini projection in the spectacular painting by Gian Paolo Pannini shown in figure 1(a). Pannini (1691-1765) was a successful practitioner of *vedutismo* and a professor of perspective at the French Academy of Rome. His students included two other famous painters of wide views, Canaletto (Venice) and Hubert Robert (Paris); and he had a strong influence on the best known purveyor of Roman views, Giambattista Piranesi. Unfortunately there is no record of what Pannini taught, or of how any of the *vedutisti* constructed their perspectives. Instead, we have tried to reverse engineer their methods by studying their works, and discussions of their perspective such as those found in [Wri83, Rap08]. *Vedutismo* perspective exhibits the following characteristics:

- There is almost always a strong central vanishing point.
- Although the field of view looks wide, there is no sign of perspective distortion: everything appears to have its proper width and shape, no matter where in the picture it is located.
- Depth seems compressed: things near the central vanishing point look unexpectedly large and close, but the outer parts of the scene do not seem overly enlarged.

In a flat cylindrical projection the angular scale is constant across the picture, and everything appears at its natural width. However, in *vedutismo* images, the angular magnification increases steadily from center to edge – but less rapidly than in the rectilinear projection. Angles at the edges appear smaller, and those in the center appear larger. There are many projections of which that is true, such as the stereographic. However we knew of no projection with this property that could render both vertical and radial straight lines as straight.

In December, 2008 Bruno Postle realized that a perspective view of an unflattened cylinder has the required properties, and demonstrated that Pannini’s image of San Pietro could be such a view. He envisioned this procedure: paint the scene on a large transparent cylinder, with the view point at its center (this could be done with a camera obscura). Then step back, and draw a rectilinear perspective of the painting, from a projection center located on the surface of the cylin-

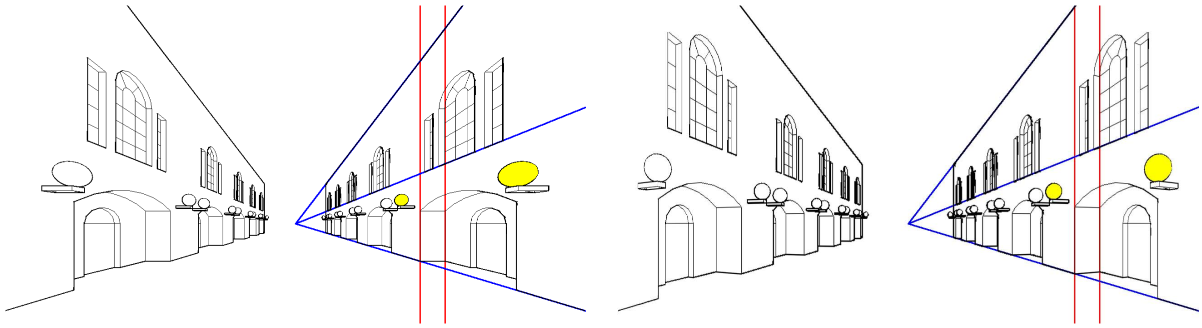

Figure 2: *Comparison of rectilinear and stereographic Pannini projections. Horizontal field of view 120°.*

der, opposite to the view point. In other words, create the cylindrical analog of the spherical stereographic projection, which we refer to as the stereographic Pannini projection. It has several desirable properties:

- Horizontal angular compression compared to the rectilinear projection: horizontal position is proportional to the tangent of half the angle of view, rather than to the tangent of the angle of view.

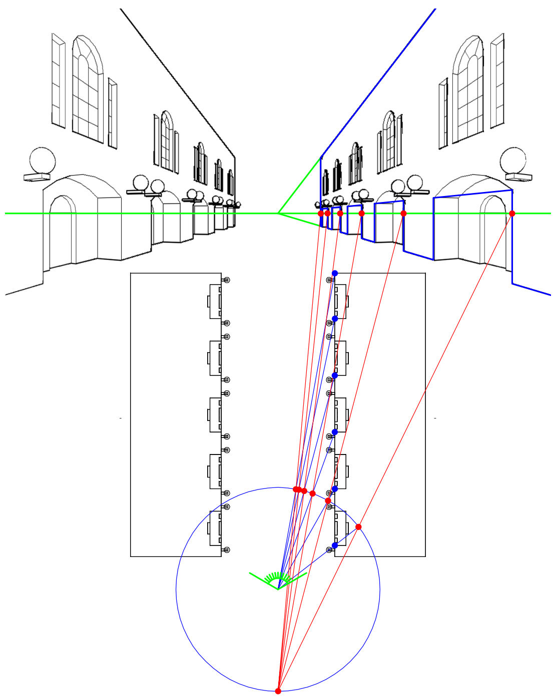

Figure 3: *Horizontal construction for stereographic Pannini projection. The point of view is at the center of the blue circle. Lines of sight (blue) from that point project key points of the plan (black) onto the circle. The red radial lines then project those points from the circle onto the picture plane (green horizontal line).*

- Straight verticals, due to the cylindrical intermediate image.
- All radial straight lines (those passing through the view center) are rendered as straight.
- Easy to draw with ordinary drafting tools.

The straight radial lines property, which is not intuitively obvious, is the reason why this projection so much resembles a rectilinear one when the subject has a strong central perspective. Figure 2 shows two perspective views of an imaginary scene, 120° wide. The left drawing is a rectilinear projection, the right one a stereographic Pannini projection. Note how the rectilinear projection distorts the door arches and lamp globes in the outer parts of the image, while in the Pannini projection those features seem to be the right size and shape. Note, too, how the far end of the buildings seems farther away in the rectilinear view. The vanishing point is at the same place in both views, and all lines that radiate from it are straight.

Being easy to draw is a necessary condition for this projection to have been used by working 18th century artists. In fact a Pannini perspective can be constructed like a rectilinear one, with one extra projection step that determines the compressed horizontal layout, as shown in Figure 3.

The stereographic form just described is only one member of a continuous family of projections, that we collectively call the Pannini projection. All are rectilinear perspective projections of a cylindrical image. The only difference is the distance of the rectilinear projection center from the center line of the cylinder. Varying that distance changes the amount of horizontal compression, without changing the other properties of the projection: vertical and radial lines remain straight, and the projection remains easy to draw. With the projection center at the center of the cylinder, there is no compression; the Pannini projection is identical to the rectilinear projection. Moving it back (while holding the projected field of view constant) gradually magnifies the middle of the image and compresses the edges. This reaches a limit as the distance approaches infinity, when the Pannini pro-

Artist	View	Year	Unlikely	Best	FOV°	RMS%	RMS°	D	N
Pannini	Interior Pantheon	1732	rect,ster,cyl	orth	82.8	2.61	1.98	13787	17
	Interior Pantheon	1734	none	ster	57.0	0.57	0.33	1.30	12
	Sta Maria Maggiore,	1753	rect,ster	orth	73.9	1.58	1.09	715365	21
	Piazza San Pietro	1754	rect,ster,cyl	orth	112.8	1.53	1.46	6056	12
	Interior San Pietro	1756	rect,ster,cyl	orth	96.4	0.50	0.43	12.3	21
Hodge	Wyatt's London Pantheon	1772	rect,orth	cyl	171.5	0.57	0.97	2.08	26

Table 1: *Horizontal projection analysis of* vedutismo *paintings. Key: FOV°= horizontal field of view in degrees; RMS% = root mean square error as % of view width; RMS°= same as angle in degrees, D = compression parameter of fitted Pannini projection, N = number of data points.*

jection becomes an orthogonal (parallel) projection of the cylinder.

3.1. Quantitative Analysis of Paintings

Table 1 summarizes the results of an analysis of the scaling of horizontal angles in six wide 18th century views. We considered five possible relationships between position and angle of view:

Rectilinear $x/K = \tan(\phi)$
Stereographic $x/K = 2\tan(\phi/2)$
Orthographic $x/K = \sin(\phi)$
Cylindrical $x/K = \phi$
Pannini $x/K = \sin(\phi)(D+1)/(D+\cos(\phi))$

x is the horizontal coordinate of a point in the picture and ϕ is the angle of view to that point, both measured from the center of the picture. K is an angular scale factor, effectively the "focal length" of the projection in the same units as x. The first four functions have one adjustable parameter, K. The Pannini formula has an additional parameter, D, and can match any of first three functions, as well as intermediate forms, according to the value of D.

To evaluate a formula, we need the true angles of view ϕ for a set of points in the picture. We could measure those angles on a plan of the subject, if we knew the true point of view, and the true direction of view. As those are unknown, we must include them, along with K, as parameters to be fitted. The analysis then takes the following form. The raw data are the horizontal coordinates x_i of some points on the picture, and the coordinates (u_i, v_i) of the same points on the floor plan of the building. The unknown parameters are the plan coordinates (U,V) of the point of view; the angle θ of the direction of view on the plan; the angular scale factor K; and, in the case of the Pannini formula, D. In terms of those parameters, the model to be fitted is:

$$\phi_i = \arctan((v_i - V)/(u_i - U)) - \theta$$
$$error_i = x_i - KP(\phi_i)$$

where $P()$ is the projection formula being tested. Note that this model is independent of the units in which the plan coordinates are measured. It is a nonlinear model, so to fit it we used a nonlinear least squares optimizer ("solver" in Microsoft Excel).

After fitting the model for each formula to the data for a given picture, we can identify the one with the smallest residual sum of squared errors as the most likely of our proposed projections. Because it has an extra degree of freedom, the Pannini formula usually achieves the best fit, but for the same reason, its error cannot be directly compared to those of the fixed models. So we restrict our choice for "best fit" to the four fixed formulas. The fitted value of D may help indicate how plausible that choice is: we expect D to be 0 for a rectilinear projection, 1 for stereographic, and > 30 for orthographic. In most cases the error of the best fit is small enough to give us confidence that the perspective was constructed, rather than just 'eyeballed'. The fitted values of D for the Pannini model tend to reinforce this conclusion. We report a projection as "unlikely" if its RMS error is at least 10% greater than the best fit, and the plot of residual errors has a visibly different shape. In many cases several projections could be considered equally likely. However the consensus is clear: these works were created using horizontally compressed projections.

It appears that all but one of the Pannini pictures were constructed with an orthographic Pannini projection, which is the most highly compressed form and "brings the center forward" most strongly. The exception is his 1734 view of the interior of the Pantheon (figure 1(d)). Expert perspectivist Lawrence Wright [Wri83, pp.167-170] accepted that the point of view is where it seems to be, inside the building, right against the back wall. That would make the horizontal field of view 110°. However, according to our analysis, the true point of view is well outside the building, and the true field of view is only 57°. The horizontal scale is probably compressed, but less so than in the other Pannini views. If our analysis is correct, two of those views (the 1734 Pantheon and the Sta. Maria Maggiore) must have been constructed on plans of the buildings, as it is not possible to see the depicted scenes from the fitted points of view.

The Wyatt's Pantheon picture has the widest field of view of any 18th century painting we know of. It is certainly not a rectilinear perspective. Figure 4 shows an "analysis by synthesis" of this picture. We constructed a 3D CAD model from plan and section drawings of the building, then fit a Pannini projection view of the model to the painting by adjusting d only. The fit is remarkably good, much better than a flat cylindrical projection (not shown); and no other common projection even comes close.

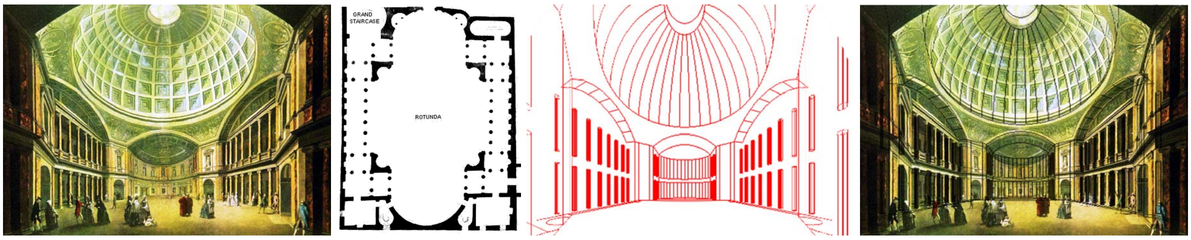

Figure 4: *Wyatt's London Pantheon. l. - r.: Painting, probably by Hodge after a drawing by Wyatt; floor plan; 3D CAD model, viewed in Pannini projection, $d = 2.5$, hfov 150°; projected model overlayed on the painting.*

4. Mathematics of the Pannini Projection

The Pannini projection is a family of partial mappings between the surface of the sphere and the plane. The sphere surface holds a true image of a 3-dimensional scene, generated by a linear projection on the center of the sphere. The plane holds a synthetic perspective view of part of the scene.

We will use the following terminology and geometrical framework.

- The *panosphere* holds the world image. Its radius is 1. Its center is the origin of Cartesian world coordinates (x, y, z) and the center of projection for the world image.
- The *view* is a plane tangent to the panosphere at $(0, 0, -1)$, which is the origin of 2D view coordinates (h, v). We call that point the *view center*, and the z-axis the *view axis*. The direction of view is toward negative z.
- We use 2D equirectangular coordinates for points in the world image: ϕ is the azimuth angle, measured in the plane $y = 0$ from the negative z-axis, θ is the altitude angle above $y = 0$.
- The 2D coordinates (h, v) and (ϕ, θ) are linearly related to pixel positions; however in this discussion the unit for all coordinates is the radius of the panosphere. Thus all angles are in radians, and all other values are conformable to the trigonometric functions.

4.1. Basic Coordinate Mappings

The basic Pannini projection is a rectilinear projection of a 3-dimensional cylindrical image, which is a linear projection of the panosphere onto a tangent cylinder. The cylinder axis coincides with the y axis. The center of the rectilinear projection is on the view axis at distance d from the cylinder axis. In our reference frame d is the z coordinate of that point.

The parameter d, which can be any non-negative number, determines the specific form of the projection. When $d = 0$ the view is rectilinear. $d = 1$ gives the cylindrical stereographic projection, and $d \to \infty$ gives the cylindrical orthographic projection.

The Cartesian coordinates of a point on the cylinder are

$$x = \sin(\phi),\ y = \tan(\theta),\ z = -\cos(\phi)$$

The distance from projection center to view plane is $d + 1$, and the distance from projection center to the parallel plane containing the cylinder point is $d + \cos(\phi)$. Their ratio,

$$S = \frac{d+1}{d+\cos(\phi)} \tag{1}$$

is the rectilinear projection scale factor for the point. Thus the mapping from sphere to plane is

$$h = S\sin(\phi) \tag{2}$$

$$v = S\tan(\theta) \tag{3}$$

The inverse horizontal mapping involves a quadratic that results from (2) and $\sin^2(\phi) + \cos^2(\phi) = 1$. The best plan is to solve it for $\cos(\phi)$, which is independent of the sign of h, then compute S and evaluate the analytic inverses of (2) and (3) with the atan2 function, to avoid the inaccuracies of arcsin() and arccos() for arguments near 1. Letting

$$k = h^2/(d+1)^2$$

the quadratic discriminant reduces to

$$\Delta = k^2 d^2 - (k+1)(kd^2 - 1)$$

There is no solution if $\Delta < 0$, otherwise

$$\cos(\phi) = \frac{-kd + \sqrt{\Delta}}{k+1} \tag{4}$$

$$S = \frac{d+1}{d+\cos(\phi)} \tag{5}$$

$$\phi = atan2(h, S\cos(\phi)) \tag{6}$$

$$\theta = atan2(v, S) \tag{7}$$

The maximum horizontal field of view varies with d. For $d \leq 1$, the practical limit is image width, because the projection is parallel to the view plane at the theoretical limit. At $d = 0$ the theoretical limit is 180° and at $d = 1$ it is 360°. For $d > 1$ the maximum field of view shrinks again, approaching 180° as $d \to \infty$, and the corresponding image width is finite. The theoretical limit (in radians) is

$$F = 2\arccos\left(-\begin{cases} 1/d, & d > 1 \\ d, & \text{otherwise} \end{cases}\right) \tag{8}$$

(a) Grand Central Terminal, hfov 220°

(b) London Eye, hfov 160°

Figure 5: *(a) Cristian Marchi,* Grand Central Terminal, New York *2009. stereographic Pannini projection. (b) Alexandre Duret-Lutz,* London Eye Jubilee Gardens *2009 (detail), general Pannini projection (H 100, T 22, B 22).*

4.2. Vertical Compression

The basic Pannini projection renders transverse horizontal lines as curves, with maximum curvature at image center. Although for some subjects this will pass unnoticed, for many others it creates an un-natural appearance. It is likely that the *vedutisti* dealt with this simply by drawing straight lines in places where curves would be disturbing, for example the transverse lines of the floor tiles in Pannini's church interiors are all straight. What we call the general Pannini projection makes it possible to emulate this practice to a degree. Either of two vertical compression functions can be applied separately in the upper and lower halves of the image, to reduce the curvature of transverse lines. One function, that we call "hard" vertical compression, can exactly straighten those lines, but is limited to fields of view less than 180°; the other, called "soft" compression, works on wider fields of view, but cannot eliminate all curvature.

A $v \leftrightarrow \theta$ mapping that straightens transverse horizontal lines is given by

$$v = S\tan(\theta)/\cos(\phi) \tag{9}$$

$$\theta = \arctan(v\cos(\phi)) \tag{10}$$

where S is given by (1). This mapping is degenerate when ϕ is an odd multiple of 90°. The general Pannini projection applies "hard" compression by computing v as a weighted average of (9) and (3). The weight of (9) can vary from 0 (no compression) to 1 (full straightening).

The "soft" compression scales v by a factor that depends on $\cos(\phi)$ and d as well as the weighting parameter. We do not specify it here because, unlike the other formulas of the Pannini projection, it has no firm theoretical basis in geometry, and may be subject to experimental improvement (for details of the current implementation please refer to the supplemental materials website).

Straightening horizontal lines necessarily displaces and bends radial lines. Fortunately the resulting curvature is strongest at extreme angles of view, and is hardly noticeable on fields of view less than 135°, or on larger fields when the vertical compression is small. As a result, vertical compression does not usually diminish the perspective illusion, and indeed often improves it.

Figure 6: London Eye*, rectilinear projection, hfov 160°.*

5. Applications of the Pannini Projection

The basic Pannini projection (without vertical compression) has been in use since December 2008. The first implementations were a script for the MathMap image processing language, and a freely available panorama viewer called *Panini* (`http://sourceforge.net/projects/pvqt/`). In *Panini* the compression is adjustable interactively, along with other viewing parameters, and the view shown on the screen can be saved to a file at moderate resolution. It has proved popular with panoramic photographers for preparing views for printing, and thumbnail views that give a good impression of the panorama. *Panini* is also used to convert photos taken with fish eye lenses to perspective form. In April 2009 the stereographic Pannini projection was available in three panorama stitching programs (both open source and commercial): Hugin, PTGui and PTAssembler, and a web panorama viewer, KRPano.

Figure 5(a) is an example of the stereographic Pannini projection. It covers 220° horizontally. Such an extreme field of view needs a deep central perspective. The angled walls visible at the sides are actually the ends of a transverse balcony that stands well behind the point of view; the Pannini projection makes even these appear straight.

In January, 2010 we implemented the general Pannini projection in the open source Hugin (`http://hugin.sourceforge.net`) and Panotools (`http://panotools.sourceforge.net`). In February 2010 it was added to Helmut Dersch's fast GPU-based panorama stitcher, PTStitcherNG. In the PanoTools implementation, the user controls

(a) Louvre, Paris, hfov 180° (b) The City of London, hfov 259°

Figure 7: *Very wide views from panoramas by Alexandre Duret-Lutz. (a) H 100, T -10, B -20 and (b) H 100, T 0, B -6.*

the projection with three parameters: H – horizontal compression, T – top compression, and B – bottom compression. H is non-linearly scaled so that H = 0 gives the rectilinear projection, $H = 100$ the stereographic Pannini projection, and $H = 150$ the orthographic Pannini projection. The vertical compression parameters T and B range from -100 to 100. Negative values select "soft" compression, positive ones "hard" compression. Using Hugin's fast previewer, the user can experiment with different values of the three parameters of the projection: horizontal compression (H) and top and bottom compressions (T, B). This is done using sliders, and the feedback is immediate (the previewer renders a lower resolution version of the image, but detailed enough to evaluate the impact of the parameters). It usually takes seconds to find good values for these parameters.

Figure 5(b) is an example of the kind of result that can be obtained routinely with the general Pannini projection. The field of view is 160°x 90°; the horizontal mapping is the default (stereographic) and a mild "hard" vertical compression has been applied. Figure 6 shows the same field in rectilinear projection. The contrast is dramatic. The extreme rectilinear perspective distortion vanishes completely in the Pannini view. The middle of the picture appears much closer, so that the wheel and the park in front of it assume their proper role as the focus of attention. The perspective of the Pannini view looks normal; indeed, without having seen the rectilinear view, one could easily mistake it for an ordinary wide angle photo. That kind of natural look is a hallmark of *vedutismo*. The two views have one defect in common: the great wheel is visibly stretched upward and to the right. Figure 7 shows two more examples of ultra-wide Pannini projections. The Louvre view has a hfov of 180°, and the one of The City of London 259°. Yet both retain a natural look across the entire image.

The Pannini projection is also useful for correcting rectilinear perspective distortion in images with smaller fields of view, such as normal wide angle photographs, and for "defishing" fish eye photos. The fact that it is continuously adjustable from rectilinear to more compressed forms makes it easy to find the correction that best suits a given image.

6. Comparison with Carroll's Method

Due to lack of space, we present only a brief comparison of the Pannini projection to the method of Carroll *et. al.*. The supplemental materials website (see last page) gives more thorough comparison against that and other methods.

Like the Pannini projection, Carrol's method transforms wide angle images into quasi-perspective views. A specific transformation is computed for each image, based on the user marking curved lines that should be shown as straight. The examples in the paper required from 5 to 28 marked lines each. In most cases very plausible perspective views were obtained, however the method did fail completely on a few example images.

Figure 8 presents two of Carroll's examples. The first column is the original image, with Carroll's control lines superimposed, the second is Carroll's result. The third column shows the basic stereographic Pannini projection, and the fourth the general Pannini projection with parameters adjusted to suit the contents of the photograph. For the Pannini projections, Hugin's focal length, pitch and roll parameters were first adjusted to get an undistorted rectilinear view with the vertical direction correctly aligned.

In the first row of figure 8, Carroll's result shows several vertical lines pointing in different directions. The lamp post is at a different angle than the vertical walls of the buildings, and leans inwards, and the tall building at the back seems to stand at an odd angle. In contrast, the Pannini projection renders all vertical lines as straight and parallel. These differences can be ascribed to the fact that the Pannini projection is derived from an intermediate spherical projection with the vertical direction correctly aligned, while Carroll's method essentially ignores lens focal length and the geometry of the depicted space. The general Pannini removes the curvature of the lines in the road with the help of a "soft" bottom compression.

The image on the second row is a challenge for the Pannini projection, whose basic form renders the horizontal lines of the cabinet as curves. With some "soft" vertical compression, the general Pannini removes most of the curvature,

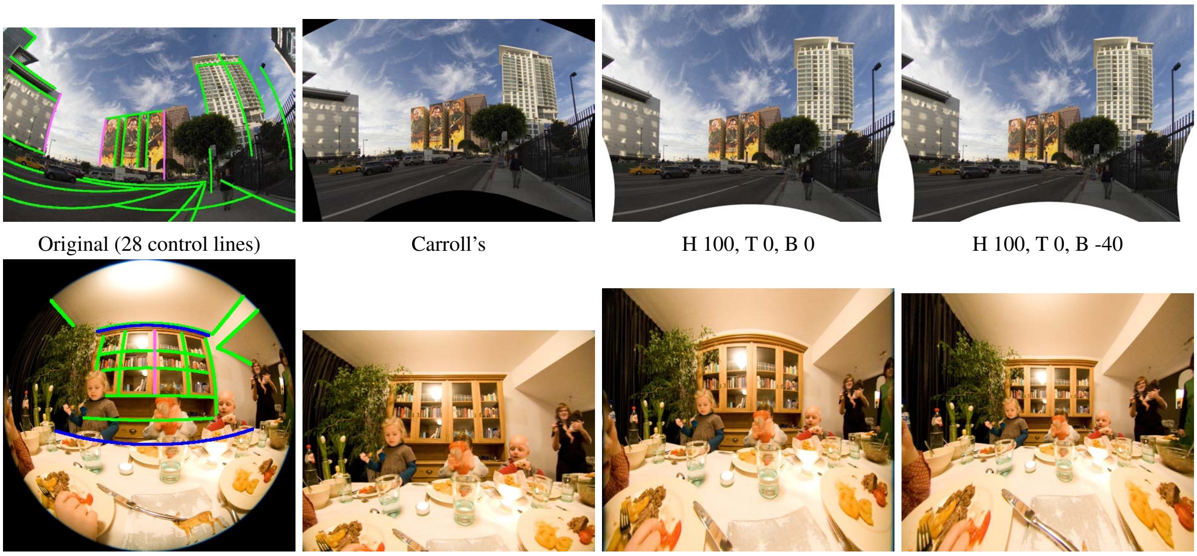

Figure 8: *Comparison of Carroll et al.'s results to the Pannini projection.*

at the cost of displacing and curving the ceiling lines. The field of view here is almost 180°, which precludes using "hard" compression to exactly straighten the cabinet.

7. Conclusions

We discovered the Pannini projection as a result of efforts to "reverse engineer" *vedutismo* perspective, and have been able to develop it into a generally useful method for rendering very wide photographic views. The resulting perspectives strongly resemble *vedutismo*, even when the field of view is far larger than any the 18th century artists drew.

We certainly have not proved that the *vedutisti* actually used this projection, but we have provided evidence that they could have. What really matters, however, is that their work inspired the Pannini projection, which adds a useful new tool to our repertoire of methods for making wide field images that "look right".

Acknowledgements

We are grateful to Cristian Marchi and to Alexandre Duret-Lutz for permission to publish their photographs (mr. Duret-Lutz made them available to us under the Creative Commons license Attribution-Noncommercial-Share Alike 2.0), and to Georges Lagarde and Luca Vascon for helpful discussions on perspective.

Supplemental Materials Website

We invite the reader to visit `http://vedutismo.net/Pannini/`where we have placed more examples of uses of the Pannini Projection and a comprehensive comparison with Carrol et al. method.

References

[CAA09] Carroll R., Agrawal M., Agarwala A.: Optimizing content-preserving projections for wide-angle images. In *SIGGRAPH '09* (2009), pp. 1–9.

[GBDL*07] German D. M., Burchill L., Duret-Lutz A., Pérez-Duarte S., Pérez-Duarte E., Sommers J.: Flattening the Viewable Sphere (artistic contribution). In *Comp. Aesthetics in Graphics, Visualization, and Imaging 2007 (CAe 2007)* (2007), pp. 23–28.

[GdGP07] German D. M., d'Angelo P., Gross M., Postle B.: New methods to project panoramas for practical and aesthetic purposes. In *Computational Aesthetics in Graphics, Visualization, and Imaging 2007 (CAe 2007)* (June 2007), pp. 13–22.

[Kub86] Kubovy M.: *The Psychology of Perspective and Renaissance Art*. Cambridge University Press, 1986.

[Rap08] Rapp J. B.: A geometrical analysis of multiple viewpoint perspective in the work of Giovanni Battista Piranesi: an application of geometric restitution of perspective. *The Journal of Architecture 13*, 6 (2008), 701–736.

[Wri83] Wright L.: *Perspective in Perspective*. Routledge & Kegan Paul, 1983.

[ZB95] Zorin D., Barr A. H.: Correction of geometric perceptual distortions in pictures. In *SIGGRAPH '95* (1995), pp. 257–264.

[ZMPP05] Zelnik-Manor L., Peters G., Perona P.: Squaring the Circle in Panoramas. In *Tenth IEEE International Conference on Computer Vision* (2005).