EUROGRAPHICS 2015/ B. Bickel and T. Ritschel *Short Paper*

Content-Aware Projection for Tiny Planets

M. Brown

University of Bath

Abstract

Tiny Planets visualise the world looking down at the ground, with physically unrealisable projections that curve the ground plane to look like small worlds. Whilst certain geometries, such as Stereographic, are known to give good Tiny Planet visualisations, the best projection to use depends on the image content. In this work we define a family of Tiny Planet projections that includes several commonly used projection types, but allows for data-dependent adaptation to best present the image content to the viewer. We show how to select optimal content-aware projections from this set, minimising distortions from conformality whilst closing gaps and emphasising salient areas in the scene.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—Display Algorithms

1. Introduction

Presenting very wide-angle imagery to the viewer is a challenging problem involving trade-offs between various types of distortion to coerce the view-sphere onto a 2D plane. Though humans perceive a wide field of view (approximately 135×200 degrees), standard projection choices such as rectilinear (straight-line preserving) or equidistant (preserving angles from the central point) look unnatural and unappealing over this field of view. Given the increased popularity of wide-angle cameras and the ability to stitch images into very wide angle views [BL07, Sze06], there is an increasing need to effectively visualise large field-of-view images.

In searching for suitable projections a natural place to begin is with the canon of cartographic projections designed to transform the earths surface to a map, as this is analogous to mapping the view sphere to a plane. German et al. [GdGP07] describe a range of such projections, and additionally propose a “hybrid” approach that combines 2 or more of these projections along lines where the mappings align (e.g, “Architectural Cylindrical”, merging Miller and Lambert Equal Area projections either side of the horizon).

An early content-aware approach was suggested by Zorin and Barr [ZB95], who find image projections that trade-off between minimum curvature (deformation of straight lines) and direct viewing (local stretch). These mappings are, however, cumbersome to specify, and the technique is limited in field of view.

Zelnik Manor et al. [ZMPP] address this problem in the context of wide angle panoramas by allowing the user to define multiple regions for a multi-plane perspective rendering of the scene. This works well with scenes in which there are clear transitions between planar surfaces (e.g., indoors), however, in many cases the optimal choice of projections is difficult for the user to define. [GKB] extended this idea to automatically find 2 planar regions per image, typically a ground plane, and the plane at infinity.

A more general solution to wide-angle image visualisation was proposed by Carroll et al., who formulate the problem as an optimisation over a set of spatially varying projections. Their optimisation aims to maintain salient structures and respect user specified constraints, whilst minimising deviations from a conformal mapping [CAA09, CAA10]. An alternative approach was proposed by Kopf et al [KLD*09], who allow the user to manipulate the attitude of scene planes and find a smooth deformation of a cylindrical projection surface that satisfies these constraints. Our approach is most similar to the former in that we use image-based constraints and user-input and optimise for a mapping that minimises distortions from conformality.

Other authors have addressed specific content-based rendering issues, for example, photographing entire street

© The Eurographics Association 2015.

DOI: 10.2312/egsh.20151008

delivered by EUROGRAPHICS DIGITAL LIBRARY www.eg.org diglib.eg.org

Figure 1: *Content-Aware projections can be used to close gaps (left), and emphasise salient regions in the scene (right)*

scenes with optimal projections to "see" down each street [AAC*06, KCSC10], or discovering optimal projections in a video sequence [WLH*12]. In this work, we consider the specific problem of Tiny Planet renderings, where the optical axis is aligned with gravity and the ground plane mapped to a circle (see Section 2). Tiny planets are normally captured using multiple stitched images, as generally the full view-sphere is required, for this we use an algorithm based on [BL07].

The novel contributions of this work are: 1) a geometric characterisation of Tiny-Planets that includes the popular stereographic projection, but allows for many other variants, 2) a technique to optimise content-aware projections from this group, satisfying various objectives: minimum distortion from conformality, gap closing, emphasis of salient regions, user constraints, smoothness of mapping.

2. Tiny Planet Projections

Tiny Planets are formed when an ultra wide-angle virtual camera looks directly at the ground, causing images to appear like tiny worlds. To give an undistorted spherical planet, the optical axis must be aligned with gravity, and projection must be radially symmetric about this axis. Such projections are radially-symmetric azimuthal, since azimuthal angles (directions on the ground plane) are preserved, and deformation on the altitude angle must be rotationally symmetric (i.e., independent of azimuth). Mapping from altitude (ϕ) and azimuth (θ) to the image plane (u,v) is thus given by:

$$\begin{bmatrix} u \\ v \end{bmatrix} = sr(\phi) \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix} \tag{1}$$

where $r(\phi)$ is a monotonic function of ϕ and s is an arbitrary scale factor. An important radially-symmetric azimuthal projection is stereographic $(r(\phi) = \sin\phi/(1+\cos\phi))$, which is the conformal azimuthal projection, but there are several other well known projection types of this form, such as downward-gnomonic $(r(\phi) = \tan\phi)$ and equidistant $(r(\phi) = \phi)$. Our task will be to select a good distortion function $r(\phi)$ that satisfies various criteria to present a pleasing image to the viewer.

3. Conformality Constraints

Conformality is a highly desirable property of an image projection, resulting in no local stretching or aspect ratio changes, and this sort of geometry is very pleasing to the eye despite potentially large changes in scale over the image. However, the only true conformal mapping that is a radial symmetric azimuthal projection is Stereographic. In a similar manner to [CAA09], we work from differential definition for conformality, and construct an energy function that penalises local deviations from conformality.

Consider the mapping of an elemental patch of the view sphere at (ϕ,θ) to polar render coordinates (r,θ). It is straightforward to show that for no local distortion (stretch)

$$\frac{dr}{d\phi} = \frac{r}{\sin\phi} \tag{2}$$

This is the conformality condition for radially-symmetric azimuthal projections (Tiny Planets). Following [CAA09] we form conformality constraints by discretising this equation

$$\frac{r_{i+1} - r_{i-1}}{2} \approx \frac{r_i}{\sin\phi} \tag{3}$$

where the index i steps over equal increments in elevation angle ϕ.

4. Content Aware Projection

Conformal mappings are attractive in that they cause no local stretching in the projection, so Stereographic projection is in general a good choice for Tiny Planet renderings. However, in many cases other projections that are close to conformal are more appropriate. For example, image data is commonly missing at the ground pole because of a tripod used in the capture process, causing black holes in the centre of the rendering. Also, large blank areas in the ground plane or sky may be uninteresting to look at, and we would like to compress these to occupy a smaller portion of the rendered image. We thus propose an objective function that allows some violation of conformality in areas of the image that are uninteresting, penalising a weighted sum-squared deviation from

© The Eurographics Association 2015.

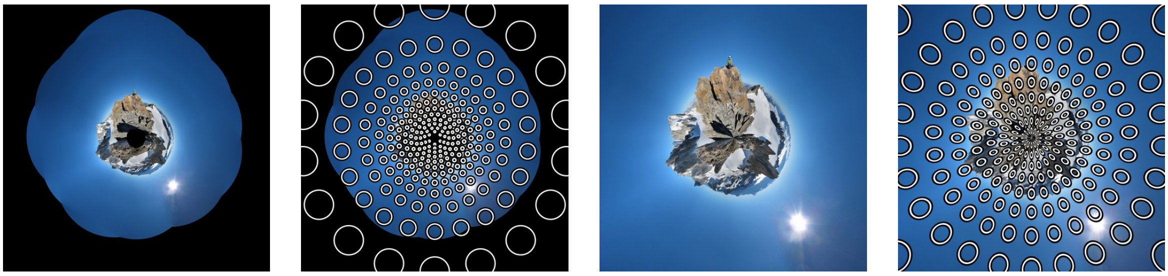

Figure 2: *In Stereographic projection (left images), circles tangent to the view sphere are mapped to circles in the render view. Our Content Aware projections (right images) remain close to conformality whilst manipulating the projection based on the content (e.g., closing gaps). Tangent circles now map to ellipses with small eccentricity.*

conformal coordinate mapping. We define the "conformality energy" E_c as:

$$E_c = \sum_i w_i \left(\frac{r_{i+1} - r_{i-1}}{2} - \frac{r_i}{\sin\phi} \right)^2 \tag{4}$$

where w_i is a spatially varying weight. For example,

$$w_i = \frac{1}{N} \sum_\theta \sum_{r \in \{r_{i-1}, r_{i+1}\}} (I(r,\theta) - \bar{I})^2 \tag{5}$$

weights the conformality penalty at radius r by the variance in the image in a ring at that radius. We have also used saliency weightings based on user input, see Section 5.

4.1. Smoothness and Constraints

The formulation above can lead to very sharp changes in the radial distortion function $r(\phi)$ where the weights w_i are small. We mitigate this by a smoothing term that favours solutions with smooth changes in the radial distortion function (by minimising the curvature $\int r''(\phi)^2 d\phi$):

$$E_s = w_s \sum_i (r_{i+1} - 2r_i + r_{i-1})^2 \tag{6}$$

This has a global weighting relative to the conformality term (w_s). We also support hard constraints on the mapping in the form of linear constraints on r_i. This can be used to close gaps at the poles, e.g., $r(\phi) = 0, \phi < \phi_{min}$ or for more general user-specified constraints of the form $\mathbf{Cr} = \mathbf{d}$ (see Section 5). The final objective function is

$$\mathbf{r}^* = \min_{\mathbf{r}} \; E_c + E_s, \quad s.t. \; \mathbf{Cr} = \mathbf{d} \tag{7}$$

This gives a constrained least-squares problem for the optimal distortion function $\mathbf{r}^*$ which is solved in closed-form by projecting to the feasible space of the constraints via SVD.

5. Results

We have tested our content-aware projection technique using a large database of stitched images. Simple variance based weighting was found to work well for gap closing and simple cases such as compressing sky (see Figure 1, left images, and Figure 2). However, in more complex cases, user specified constraints were helpful. We experimented with two forms of user constraints:

User Defined Saliency User strokes are applied in equidistant coordinates to specify weightings over the entire view sphere. This allows high-level specification of areas where conformality should not be violated (Figure 3).

Hard Constraints We also experimented with hard constraints that specify the mapping of an altitude angle in the input images to the output. Here the user drags two circles in the stereographic view to specify the mapping of altitude angles. These constraints are incorporated in the $\mathbf{Cr} = \mathbf{d}$ term in Equation 7 (Figure 4).

Overall the second form was found most effective, with normally just one or two circle mappings needing to be specified to generate good results. The resulting radial distortion functions ($r(\phi)$) are similar to those for Stereographic projection (see Figure 4, right), with small deviations corresponding to the user constraints.

6. Conclusions

We have presented a scheme to generate content-specific Tiny Planet projections. Projections close to Stereographic were found to give the best results, with hard user constraints on the mapping of radii giving an effective form of control.

References

[AAC*06] AGARWALA A., AGRAWALA M., COHEN M., SALESIN D., SZELISKI R.: Photographing long scenes with multi-viewpoint panoramas. In *ACM Transactions on Graphics (TOG)* (2006), vol. 25, ACM, pp. 853–861. 2

[BL07] BROWN M., LOWE D. G.: Automatic panoramic image

© The Eurographics Association 2015.

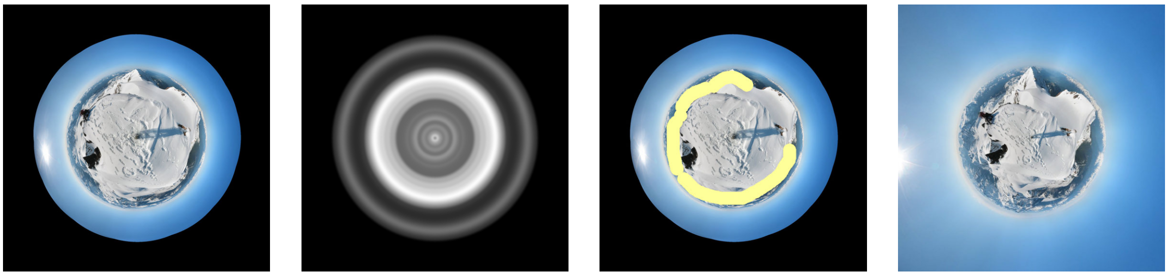

Figure 3: *We modulate the conformality penalty based on image based saliency measures such as standard deviation (centre-left) or user input (centre-right, yellow stroke). In this example, both the variance measure and user input give low weight to the sky and snow, allowing the projection to apply stretch/squash these regions (right).*

Figure 4: *The user specifies constraints mapping the red circles to the green circles to expand/contract areas of interest (results centre-right). Note that the new projections remain close to stereographic (right).*

stitching using invariant features. *International Journal of Computer Vision 74*, 1 (2007), 59–73. 1, 2

[CAA09] CARROLL R., AGRAWALA M., AGARWALA A.: Optimizing content-preserving projections for wide-angle images. In *ACM Transactions on Graphics (TOG)* (2009), vol. 28, ACM, p. 43. 1, 2

[CAA10] CARROLL R., AGARWALA A., AGRAWALA M.: Image warps for artistic perspective manipulation. *ACM Transactions on Graphics (TOG) 29*, 4 (2010), 127. 1

[GdGP07] GERMAN D. M., D'ANGELO P., GROSS M., POSTLE B.: New methods to project panoramas for practical and aesthetic purposes. *Proceedings of Computational Aesthetics* (2007), 15–22. 1

[GKB] GAO J., KIM S. J., BROWN. M. S.: Constructing image panoramas using dual-homography warping. In *International Conference on Computer Vision and Pattern Recognition (CVPR 2011)*, pp. 49–56. 1

[KCSC10] KOPF J., CHEN B., SZELISKI R., COHEN M.: Street slide: browsing street level imagery. *ACM Transactions on Graphics (TOG) 29*, 4 (2010), 96. 2

[KLD*09] KOPF J., LISCHINSKI D., DEUSSEN O., COHEN-OR D., COHEN M.: Locally adapted projections to reduce panorama distortions. In *Computer Graphics Forum* (2009), vol. 28, Wiley Online Library, pp. 1083–1089. 1

[Sze06] SZELISKI R.: Image alignment and stitching: A tutorial. *Foundations and Trends® in Computer Graphics and Vision 2*, 1 (2006), 1–104. 1

[WLH*12] WEI J., LI C.-F., HU S.-M., MARTIN R. R., TAI C.-L.: Fisheye video correction. *Visualization and Computer Graphics, IEEE Transactions on 18*, 10 (2012), 1771–1783. 2

[ZB95] ZORIN D., BARR A. H.: Correction of geometric perceptual distortions in pictures. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques* (1995), ACM, pp. 257–264. 1

[ZMPP] ZELNIK-MANOR L., PETERS G., PERONA P.: Squaring the circle in panoramas. In *Tenth IEEE International Conference on Computer Vision (ICCV 2005)*, vol. 2, pp. 1292–1299. 1

© The Eurographics Association 2015.