

Supplementary Document on Revectorization-Based Soft Shadow Mapping

M. C. F. Macedo, A. L. Apolinário Jr., and K. A. Agüero

Department of Computer Science, Federal University of Bahia (UFBA), Bahia, Brazil
marciocfmacedo@gmail.com, antonio.apolinario@ufba.br, karl@dcc.ufba.br

1. Shadow Map Traversal for Shadow Revectorization

In this section, we present details of the shadow map traversal algorithm to estimate the relative position \mathbf{r} of a fragment located in a shadow silhouette. The pseudocode shown in Algorithm 1 allows one to implement both functions **estimativeRelativePosition** (Lines 1-10 of Algorithm 1) and **computeDists** (a summary of Lines 2-7 of Algorithm 1) discussed in the main paper.

Given the coordinate \mathbf{c} of a camera-view fragment in its corresponding shadow map texel (Equation 2 in the main paper, Line 2 of Algorithm 1) and the minimum horizontal and vertical distances between neighbour texels in the shadow map (Line 3 of Algorithm 1), we can traverse the shadow map and compute four signed values to measure the distance between a fragment and a shadow silhouette for left, right, bottom, and top directions (Lines 4-7 of Algorithm 1). Then, the relative position \mathbf{r} can be estimated by the normalization of those distance values (Lines 8-9 of Algorithm 1).

The pseudocode that shows how signed distance values can be computed from shadow map traversals is listed in the procedure **computeDistance**, Lines 11-31 of Algorithm 1, and is described as follows. For each shadow map sample neighbour of a given fragment (the symbol \odot in Line 16 of Algorithm 1 represents the element-wise multiplication of vectors), RBSM computes the shadow test of that sample (Line 17 of Algorithm 1) and evaluates whether the shadow test result of the shadow map sample is different from the one estimated for the given fragment projected in the shadow map space (Line 18 of Algorithm 1). If that is the case, the traversal must be ended for the current direction (Line 20 of Algorithm 1), since the sample is not in the same lit/umbra side of the given fragment on the shadow silhouette. Moreover, as visible in the Figure 1 of the main paper, the fragments with the smallest normalized relative distance are the ones localized near the transition between umbra and lit fragments. To take this fact into consideration, we save that information in the algorithm (Line 19 of Algorithm 1). Even if the shadow map sample has the same visibility condition of the given fragment (Line 21 of Algorithm 1), we perform the 4-connected neighbourhood shadow test comparison (Line 22 of Algorithm 1) to check whether the shadow map sample is still located in a shadow silhouette (Lines 23-25 of Algorithm 1) and the shadow map traversal must continue (Line 27 of Algorithm 1). When one of those two boundary conditions is

Algorithm 1 Shadow Map Traversal for Shadow Revectorization

```

1: procedure ESTIMATERELATIVEPOSITION( $\tilde{\mathbf{p}}$ ,  $\mathbf{S}$ ,  $V_{SM}$ )
2:    $\mathbf{c} \leftarrow$  COMPUTECOORDINATES( $\tilde{\mathbf{p}}$ ,  $\mathbf{S}$ );
3:    $\mathbf{step} \leftarrow$  DETERMINESHADOWMAPSTEPsize( $\mathbf{S}$ );
4:    $d_l \leftarrow$  COMPUTEDISTANCE( $\tilde{\mathbf{p}}$ ,  $V_{SM}$ ,  $\mathbf{S}$ ,  $(-1, 0)$ ,  $\mathbf{step}$ ,  $\mathbf{c}_x$ );
5:    $d_r \leftarrow$  COMPUTEDISTANCE( $\tilde{\mathbf{p}}$ ,  $V_{SM}$ ,  $\mathbf{S}$ ,  $(1, 0)$ ,  $\mathbf{step}$ ,  $1 - \mathbf{c}_x$ );
6:    $d_b \leftarrow$  COMPUTEDISTANCE( $\tilde{\mathbf{p}}$ ,  $V_{SM}$ ,  $\mathbf{S}$ ,  $(0, -1)$ ,  $\mathbf{step}$ ,  $\mathbf{c}_y$ );
7:    $d_t \leftarrow$  COMPUTEDISTANCE( $\tilde{\mathbf{p}}$ ,  $V_{SM}$ ,  $\mathbf{S}$ ,  $(0, 1)$ ,  $\mathbf{step}$ ,  $1 - \mathbf{c}_y$ );
8:    $\mathbf{r} \leftarrow \left\{ \frac{\max(d_l, d_r)}{|d_l| + |d_r|}, \frac{\max(d_t, d_b)}{|d_t| + |d_b|} \right\}$ ;
9:   return  $\mathbf{r}$ ;
10: end procedure

11: procedure COMPUTEDISTANCE( $\tilde{\mathbf{p}}$ ,  $V_{SM}^{\tilde{\mathbf{p}}}$ ,  $\mathbf{S}$ ,  $\mathbf{dir}$ ,  $\mathbf{step}$ ,  $\Delta$ )
12:    $d \leftarrow 0$ ;
13:    $\tilde{\mathbf{s}} \leftarrow \tilde{\mathbf{p}}$ ;
14:   has $V_{SM}$ changed  $\leftarrow 0$ ;
15:   while traversing  $\mathbf{S}$  do
16:      $\tilde{\mathbf{s}} \leftarrow \tilde{\mathbf{s}} + \mathbf{dir} \odot \mathbf{step}$ ;
17:      $V_{SM}^{\tilde{\mathbf{s}}} \leftarrow$  COMPUTESHADOWTEST( $\tilde{\mathbf{s}}$ ,  $\mathbf{S}$ );
18:     if  $V_{SM}^{\tilde{\mathbf{p}}} \neq V_{SM}^{\tilde{\mathbf{s}}}$  then
19:       has $V_{SM}$ changed  $\leftarrow 1$ ;
20:       break;
21:     else
22:        $\mathbf{D} \leftarrow$  COMPUTESILHOUETTEDIRECTION( $V_{SM}^{\tilde{\mathbf{s}}}$ );
23:       if  $\mathbf{D} = \{0, 0, 0, 0\}$  then
24:         break;
25:       end if
26:     end if
27:      $d \leftarrow d + 1$ ;
28:   end while
29:    $d \leftarrow d + \Delta$ ;
30:   return  $-d \times (1 - \text{has}V_{SM}\text{changed}) + d \times \text{has}V_{SM}\text{changed}$ ;
31: end procedure

```

achieved, the shadow map traversal is stopped for the current direction and the computed distance is added to the coordinate \mathbf{c} previously computed (Line 29 of Algorithm 1), oriented towards the shadow silhouette (Line 30 of Algorithm 1).

2. Additional Results

In this section, we provide additional visual and performance comparisons between distinct soft shadow techniques for two new scenarios: Rungholt (Figure 1-(a)) and Vokselia (Figure 1-(b)). Both are large-scale scenarios composed of simple, cube-like shaped objects that cast shadows on each other and on the scene.

2.1. Rendering Quality

In this section, we provide an additional visual quality evaluation for both Rungholt and Vokselia scenarios.

In the Rungholt scenario (Figures 2, 3, and 4), all real-time soft shadow techniques estimate the penumbra size incorrectly (Figures 2-(a, b, c, d, e)). As we have discussed in the paper, that happens because these real-time techniques approximate the area light source by a single point light source, causing the inaccurate penumbra size computation. Moreover, while the MSSM technique provides the most accurate anti-aliasing for this scenario (Figure 3-(c)), the small light leaking artifacts generated by the technique (see the regions pointed by red arrows in Figure 2-(c)) diminish the perceived realism of the shadow rendering (Figure 4-(c)).

For the Vokselia scenario (Figures 5, 6, and 7), MSSM and SS-ABSS suffer from high numerical and perceptual errors due to the presence of aliasing and light leaking artifacts that are easily perceived by the user (see the aliased red pixels in Figure 7-(b, c)). The aliasing artifacts are reduced by SSRBSSM, as mainly shown in Figure 5-(e). Similarly to SSABSS, PCSS suffers from aliasing artifacts (see the aliased red pixels in Figure 7-(a)), that are accurately reduced by RBSSM (Figures 6-(d) and 7-(d)). Nevertheless, Figure 7 shows that, for the background part of the scenario, none of the techniques are able to faithfully reproduce the quality attained by the ground-truth.

As can be seen in Figures 2, 3, 4, 5, 6, and 7, MSSM suffers from the generation of light leaking artifacts along the scene. SS-ABSS is slower and more inaccurate than PCSS. RBSSM produces shadows with the lowest numerical and perceptual errors among the techniques tested in the paper, due to the effective reduction of aliasing and light leaking artifacts. Finally, SSRBSSM is able to produce soft shadows about 2 to 4 times faster than RBSSM, while decreasing the quality of the soft shadow rendering.

Finally, in Figure 11, we show that the quality of the different soft shadow mapping techniques evaluated in the paper can be improved by increasing the resolution of the shadow map.

2.2. Rendering Performance

In Figures 8, 9, and 10, we show that the performance behaviour of the soft shadow techniques under varying shadow map and viewport resolutions is similar for the four scenarios evaluated in this document and in the main paper.




Figure 1: Two additional scenarios evaluated in this supplementary document. Accurate shadows were computed using the average of 1024 area light source samples.

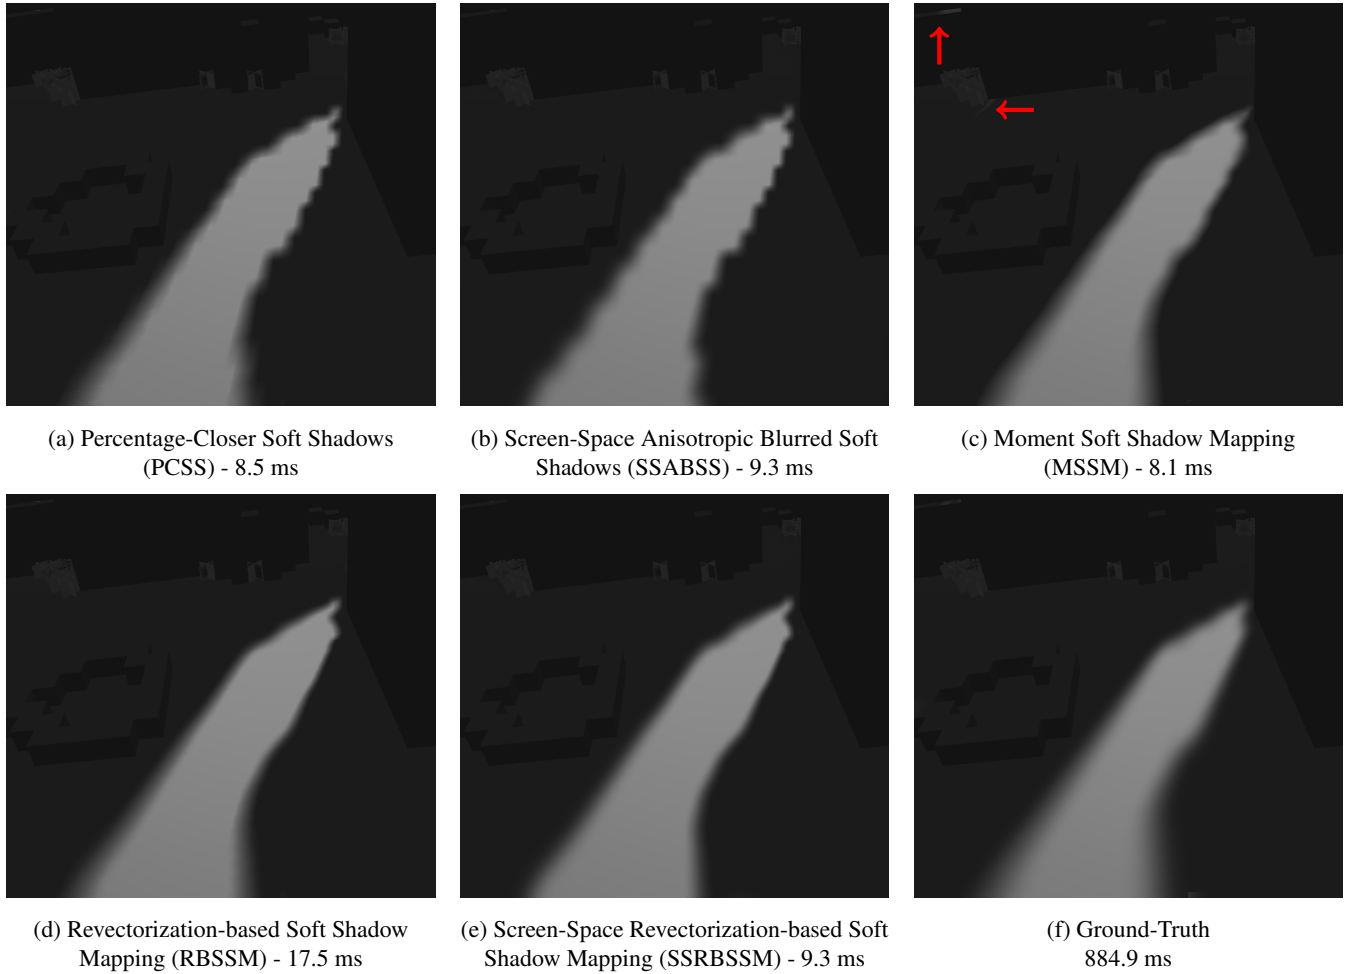


Figure 2: An equal low shadow map resolution comparison between soft shadow techniques. Light leaking artifacts are pointed by red arrows in (c). Images were generated for the Rungholt model using a 1024^2 shadow map resolution. The ground-truth image was computed using the average of 1024 area light source samples.

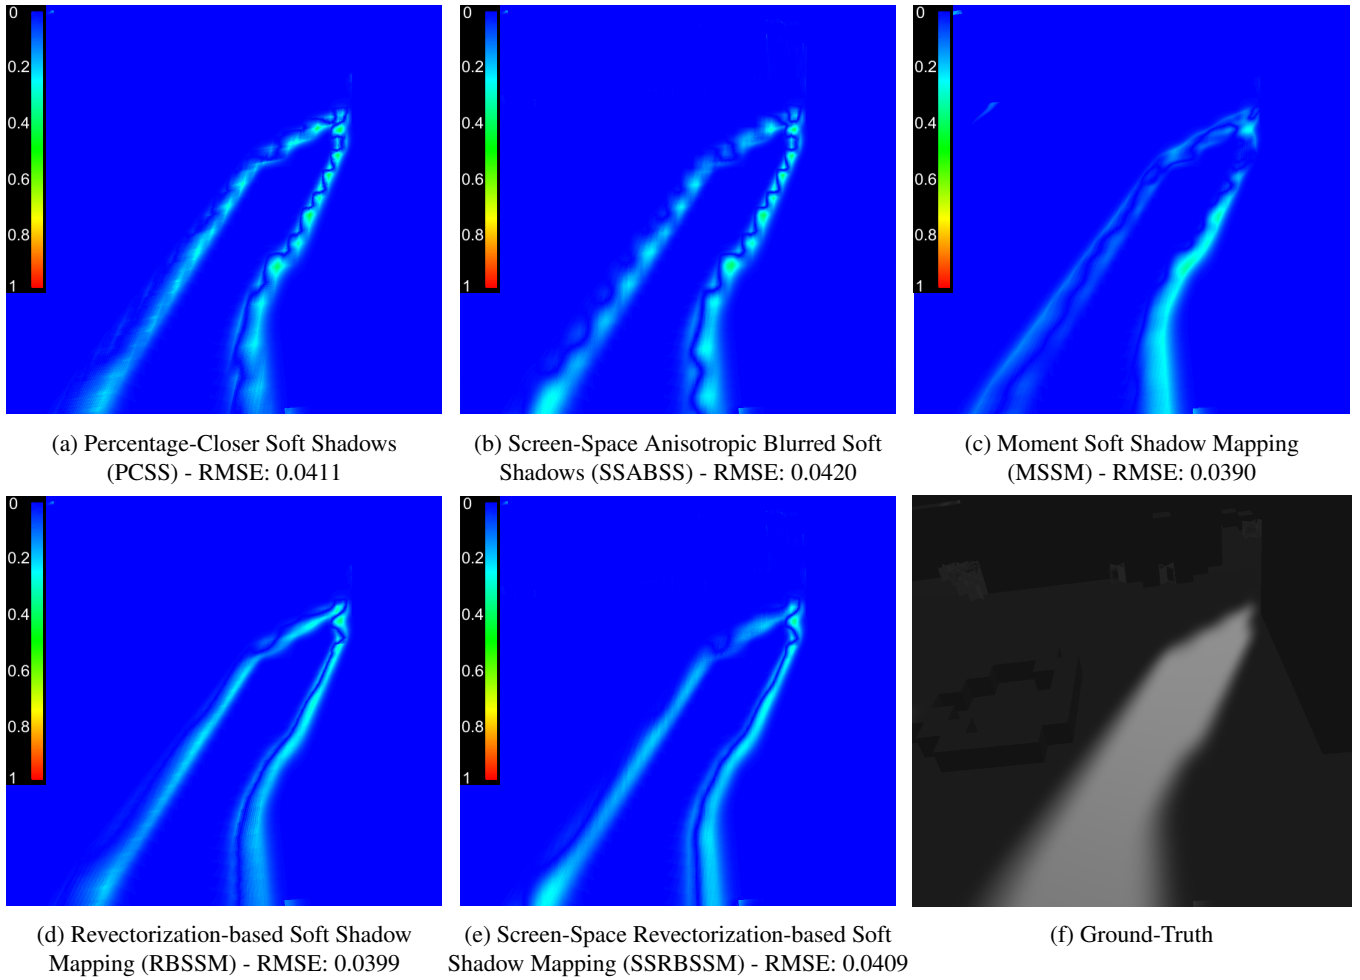


Figure 3: False color visualizations show the difference between the shadows produced by different soft shadow techniques (a, b, c, d, e) and the ground-truth one (f). Images were generated for the Rungholt model using a 1024^2 shadow map resolution and evaluated using the root mean squared error (RMSE). The ground-truth image was computed using the average of 1024 area light source samples.

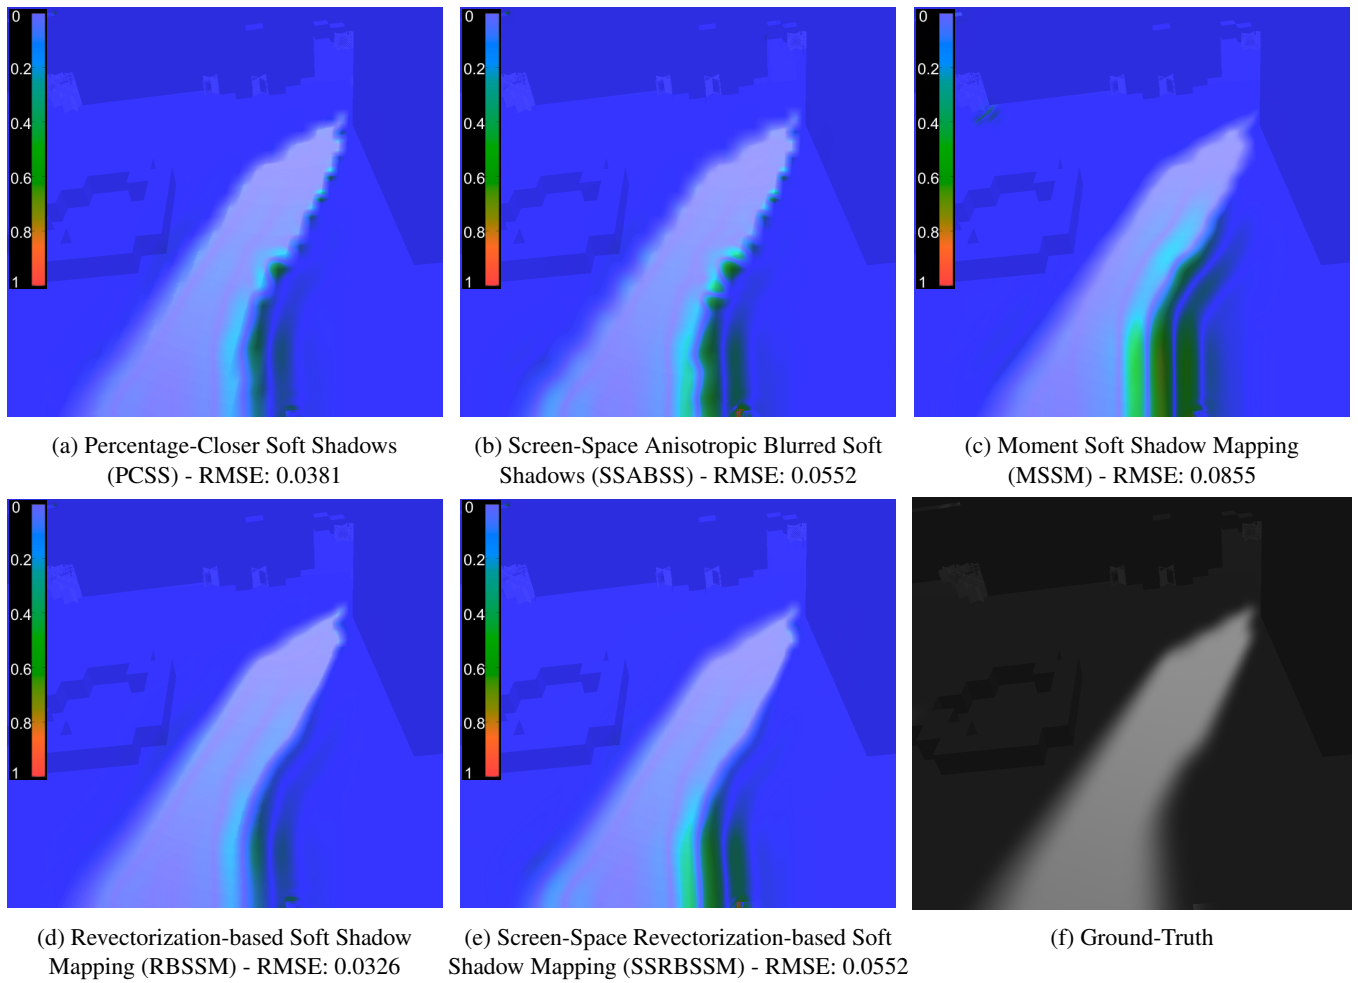


Figure 4: HDR-VDP-2 metric shows the perceptual difference between the shadows produced by different soft shadow techniques (a, b, c, d, e) and the ground-truth one (f). Images were generated for the Rungholt model using a 1024^2 shadow map resolution. In the sub-captions, we show the root mean squared error (RMSE) values measured from the perceptual metric. The ground-truth image was computed using the average of 1024 area light source samples.

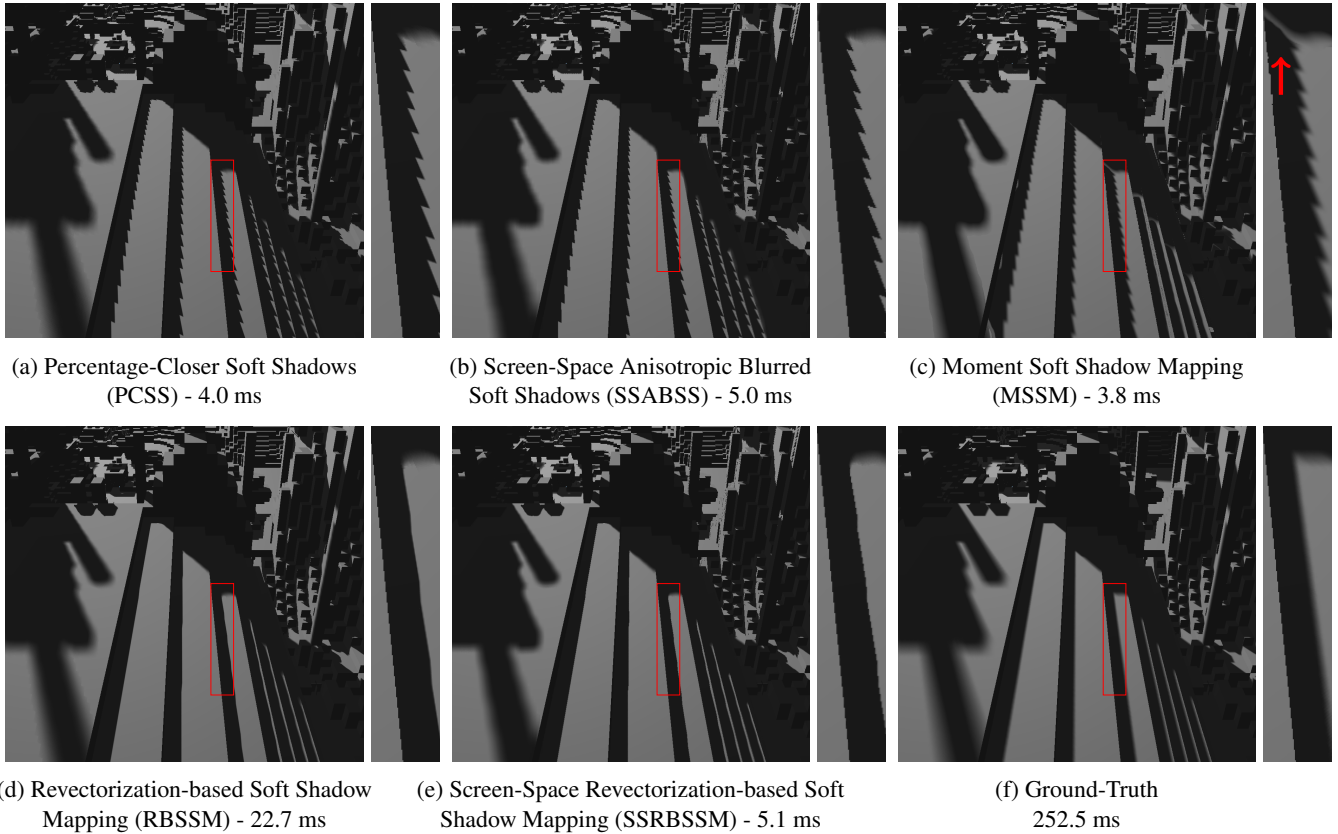


Figure 5: An equal low shadow map resolution comparison between soft shadow techniques. Light leaking artifact is pointed by the red arrow in (c). Images were generated for the Vokselia model using a 1024^2 shadow map resolution. The ground-truth image was computed using the average of 1024 area light source samples.

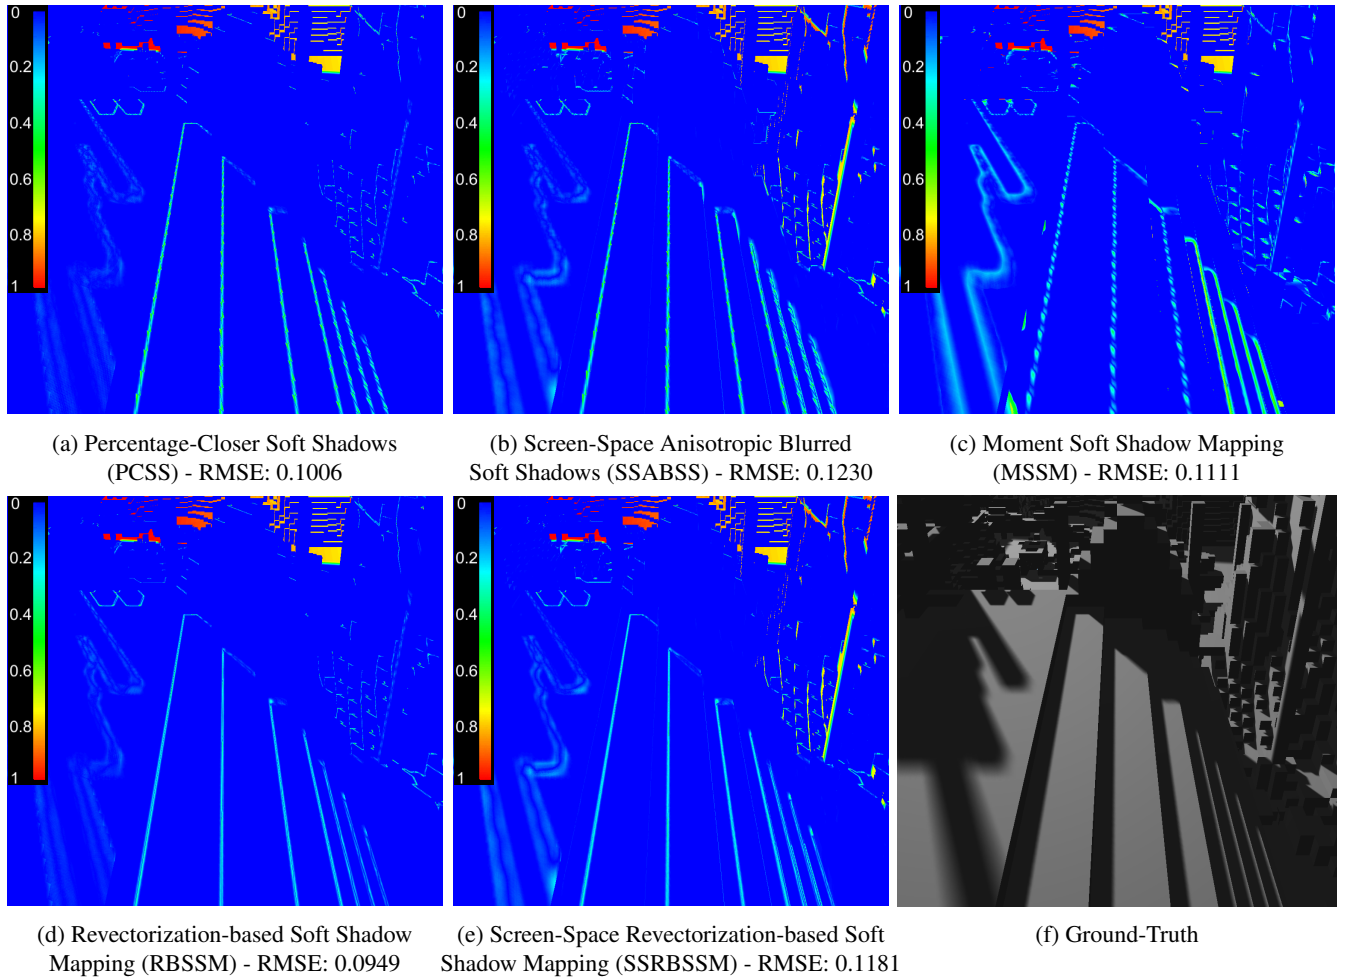


Figure 6: False color visualizations show the difference between the shadows produced by different soft shadow techniques (a, b, c, d, e) and the ground-truth one (f). Images were generated for the Vokselia model using a 1024^2 shadow map resolution and evaluated using the root mean squared error (RMSE). The ground-truth image was computed using the average of 1024 area light source samples.

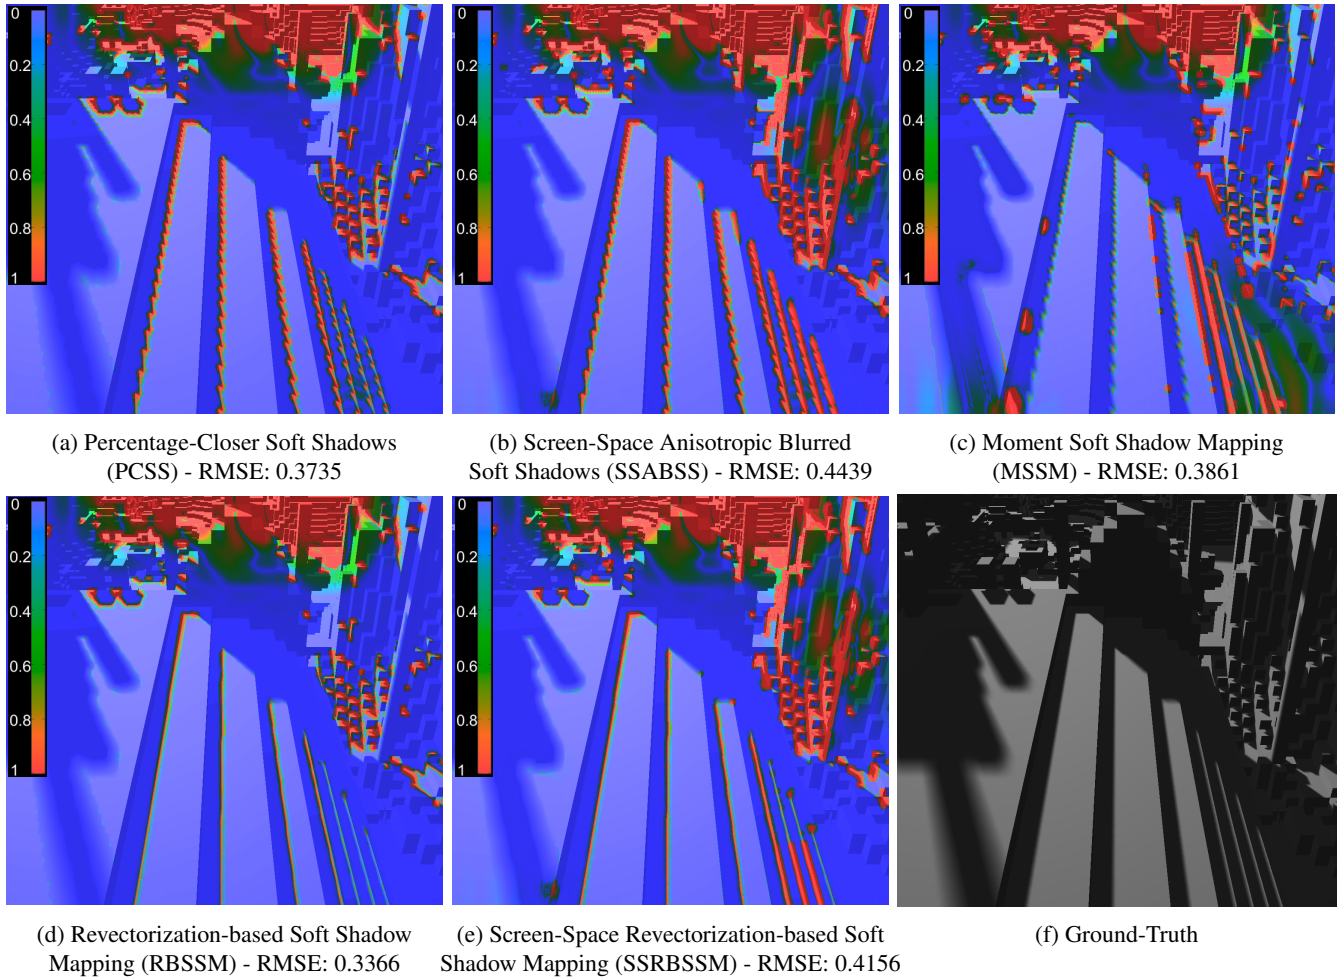


Figure 7: HDR-VDP-2 metric shows the perceptual difference between the shadows produced by different soft shadow techniques (a, b, c, d, e) and the ground-truth one (f). Images were generated for the Vokselia model using a 1024^2 shadow map resolution. In the sub-captions, we show the root mean squared error (RMSE) values measured from the perceptual metric. The ground-truth image was computed using the average of 1024 area light source samples.

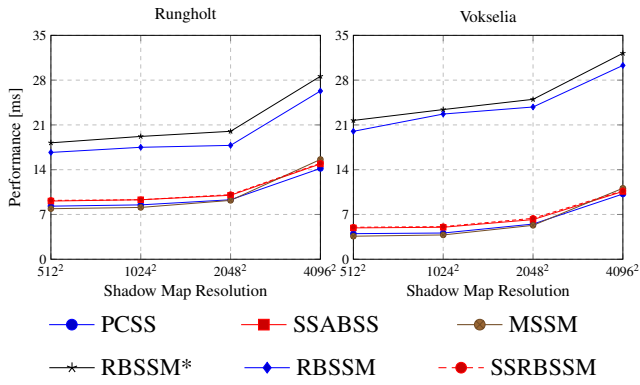


Figure 8: Time usage (in milliseconds) for several soft shadow techniques and scenes shown in Figures 2 and 5 at an output 1280×720 resolution. Measurements include varying shadow map resolution. RBSSM* refers to the unoptimized implementation of RBSSM.

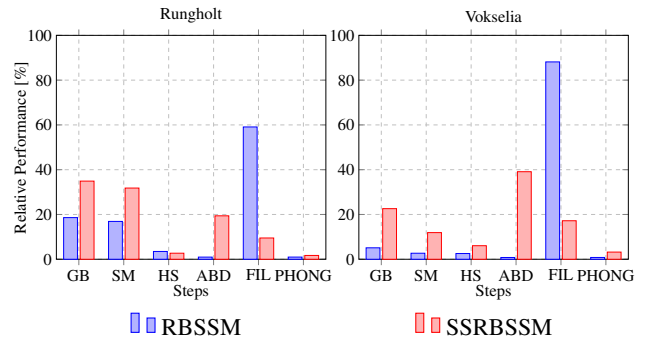


Figure 10: Time usage (in %) for the main steps of both RBSSM and SSRBSSM techniques in the scenes shown in Figures 2 and 5. Steps: GB - G-Buffer rendering. SM - Shadow map rendering. HS - HSM building (for RBSSM) or filtered Hard Shadow rendering (for SSRBSSM). ABD - Average blocker depth computation. FIL - Soft shadow filtering. PHONG - Phong shading. Steps not covered in this figure (e.g., penumbra size estimation) required less than 1% of the total frame time obtained by our experimental setup.

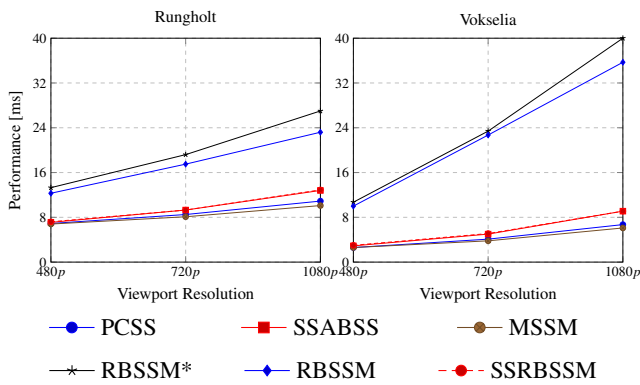


Figure 9: Time usage (in milliseconds) for several soft shadow techniques and scenes shown in Figures 2 and 5 at an 1024^2 shadow map resolution. Measurements include varying output resolution. RBSSM* refers to the unoptimized implementation of RBSSM.













Percentage-Closer Soft Shadows (PCSS) - 256^2 Percentage-Closer Soft Shadows (PCSS) - 1024^2 Percentage-Closer Soft Shadows (PCSS) - 4096^2 Screen-Space Anisotropic Blurred Soft Shadows (SSABSS) - 256^2 Screen-Space Anisotropic Blurred Soft Shadows (SSABSS) - 1024^2 Screen-Space Anisotropic Blurred Soft Shadows (SSABSS) - 4096^2 Moment Soft Shadow Mapping (MSSM) - 256^2 Moment Soft Shadow Mapping (MSSM) - 1024^2 Moment Soft Shadow Mapping (MSSM) - 4096^2 Revectorization-based Soft Shadow Mapping (RBSSM) - 256^2 Revectorization-based Soft Shadow Mapping (RBSSM) - 1024^2 Revectorization-based Soft Shadow Mapping (RBSSM) - 4096^2

Figure 11: Shadows produced by distinct soft shadow techniques under different shadow map resolutions. Images were generated for the San Miguel model.