Photon surfaces for robust, unbiased volumetric density estimation

Supplemental document

XI DENG*, SHAOJIE JIAO*, BENEDIKT BITTERLI, and WOJCIECH JAROSZ, Dartmouth College, USA

1 DERIVATION OF PHOTON SURFACE NORMALS

In the following, we provide full derivations for the normals of the photon surface estimators in Sec. 5 of the main paper.

Generalized Photon Planes. For two arbitrary distance sampling dimensions $\xi_{a_1} = t_i$ and $\xi_{a_2} = t_j$ where $l \geq i \geq j \geq 1$, the offset vector $\mathbf{g}$ becomes

$$\mathbf{g}(\overline{\xi_a}) = \mathbf{x}_0 - \mathbf{y}_0 \tag{1}$$

$$= \underbrace{\left(\mathbf{x}_l + t_i \omega_i + t_j \omega_j + \sum_{k \neq i, k \neq j} t_k \omega_k \right)}_{\mathbf{x}_0(t_i, t_j)} - \underbrace{(\mathbf{y}_1 + s_1 \psi_1)}_{\mathbf{y}_0(s_1)} \tag{2}$$

$$= \left(\mathbf{x}_l - \mathbf{y}_1 + \sum_{k \neq i, k \neq j} t_k \omega_k \right) + \mathbf{A}\overline{\xi_a}, \tag{3}$$

where $\mathbf{A} = [t_i, t_j, s_1]$. The Jacobian is

$$\mathbf{J}^{\mathbf{g}}_{t_i, t_j, s_1} = |\mathbf{n} \cdot \psi_1|, \quad \text{with} \quad \mathbf{n} = \left(\frac{\partial \mathbf{x}_0^*}{\partial t_i} \times \frac{\partial \mathbf{x}_0^*}{\partial t_j} \right), \tag{4}$$

where obtaining $\mathbf{n}$ requires computing the directions and lengths of the two partial differentials. Since for any distance term t_i we have $\partial \mathbf{x}_0^* / \partial t_i = \omega_i^*$, the scaled normal is

$$\mathbf{n} = \omega_i \times \omega_j. \tag{5}$$

---

*Joint first authors.

Authors' address: Xi Deng, xi.deng.gr@dartmouth.edu; Shaojie Jiao, shaojie.jiao.gr@dartmouth.edu; Benedikt Bitterli, benedikt.bitterli.gr@dartmouth.edu; Wojciech Jarosz, wojciech.k.jarosz@dartmouth.edu, Dartmouth College, 9 Maynard Street, Hanover, NH, 03755, USA.

Photon Spheres. We obtain photon spheres by integrating out one directional sampling dimension. Since a direction is actually a two-dimensional variable, we first decompose it into spherical coordinates. If we choose the direction ω_1, we have $\xi_{a_1} = \cos \theta_1$ and $\xi_{a_2} = \phi_1$ where $d\omega_1 = d\phi_1 d\cos \theta_1$. This gives

$$\mathbf{g}(\overline{\xi_a}) = \mathbf{x}_0 - \mathbf{y}_0 \tag{6}$$

$$= \underbrace{(\mathbf{x}_1 + t_1 \omega_1)}_{\mathbf{x}_0(\cos \theta_1, \phi_1)} - \underbrace{(\mathbf{y}_1 + s_1 \psi_1)}_{\mathbf{y}_0(s_1)} \tag{7}$$

$$= (\mathbf{x}_1 - \mathbf{y}_1) + \mathbf{A}\overline{\xi_a} \tag{8}$$

where $\mathbf{A} = [\cos \theta_1, \phi_1, s_1]$, and

$$\mathbf{J}^{\mathbf{g}}_{\cos \theta_1, \phi_1, s_1} = |\mathbf{n} \cdot \psi_1|, \quad \text{with} \quad \mathbf{n} = \left(\frac{\partial \mathbf{x}_0^*}{\partial \phi_1} \times \frac{\partial \mathbf{x}_0^*}{\partial \cos \theta_1} \right), \tag{9}$$

where $\partial \mathbf{x}_0^* / \partial \phi_1$ and $\partial \mathbf{x}_0^* / \partial \cos \theta_1$ are the tangents at $\mathbf{x}_0^*$ on the rings formed by sweeping ϕ_1 and θ_1. Since these tangents are mutually perpendicular, with lengths:

$$\left| \frac{\partial \mathbf{x}_0^*}{\partial \phi_1} \right| = |t_1^* \sin \theta_1^*|, \quad \text{and} \quad \left| \frac{\partial \mathbf{x}_0^*}{\partial \cos \theta_1} \right| = \left| \frac{t_1^*}{\sin \theta_1^*} \right|. \tag{10}$$

we have

$$\mathbf{n} = \omega_1^* t_1^2. \tag{11}$$

Photon Cones, Disks, Cylinders, and beyond. By decoupling the azimuthal and polar angle terms, we can choose one angle and one distance dimension, or two angles but each coming from a different path vertex. Depending on which of these dimensions we choose for $\overline{\xi_a}$, we get different photon surface estimators.

If we choose ϕ_1 and t_1, $\mathbf{x}_0^*$ will be allowed to move along ω_1^* and the tangent $\omega_1^* \times \omega_2$, which will sweep a *photon cone* with the normal perpendicular to both ω_1^* and ω_2:

$$\mathbf{n} = \underbrace{\omega_1^*}_{\frac{\partial \mathbf{x}_0^*}{\partial t_1}} \times \underbrace{\left(\frac{\omega_1^* \times \omega_2}{|\omega_1^* \times \omega_2|} \right) t_1^* \sin \theta_1}_{\frac{\partial \mathbf{x}_0^*}{\partial \phi_1}} = (\omega_1^* \times (\omega_1^* \times \omega_2)) t_1^*. \tag{12}$$

The simplification in the last step is because $|\omega_1^* \times \omega_2| = |\omega_1 \times \omega_2| = |\sin \theta_1|$.

If we keep t_1 but choose $\cos \theta_1$ instead of ϕ_1, we get a *photon disc*. The normal will point towards $\omega_2 \times \omega_1^*$, and based on the partial differentials we get

$$\mathbf{n} = \frac{(\omega_1 \times \omega_2) t_1^*}{\sin^2 \theta_1^*}. \tag{13}$$

If we keep ϕ_1 but choose t_2 instead of t_1, we get yet another different photon surface: the *photon cylinder*:

$$\mathbf{n} = (\omega_2 \times (\omega_1^* \times \omega_2)) t_1. \tag{14}$$

The photon cone, cylinder, and disk are only special cases of more general classes of photon surfaces. Naturally, we'd like to generalize all the photon surfaces the same way we generalized photon planes: allowing the sampled terms we include in $\overline{\xi_a}$ to be from any arbitrary position on the photon subpath. In fact, we have more choices, as for any i where the spin happens, our choice ranges from spinning the whole following portion of the photon subpath, to spinning one single segment and leaving anything else unchanged. If we choose two directional terms, then all the importance-sampled directional terms remain properly importance sampled as all the angles between adjacent segments stay unchanged. If, however, we also include distance terms, then some importance-sampled directions might change while some following directions stay the same, which makes some estimators worse than others, and here we exclude this type of surfaces for simplicity. In the end of this section we show the derivation of 2 generalized classes of photon surfaces: the photon toroid and hyperboloid.

The photon toroid corresponds to choosing 2 arbitrary ϕ samples. If we choose ϕ_i and ϕ_j where $l \geq i > j \geq 1$, the normal will be the cross product of tangents $\omega_{i+1} \times \omega_i^*$ and $\omega_{j+1}^* \times \omega_j^*$ whose lengths are both equal to the distance from $\mathbf{x}_0^*$ to the spinning axis. This gives us

$$\mathbf{n} = \frac{(\omega_{i+1} \times \omega_i^*) \times (\omega_{j+1}^* \times \omega_j^*)|(\mathbf{x}_0^* - \mathbf{x}_i) \times \omega_{i+1}^*||(\mathbf{x}_0^* - \mathbf{x}_j^*) \times \omega_{j+1}^*|}{\sin \theta_i \sin \theta_j} \tag{15}$$

Similarly, the photon hyperboloid corresponds to choosing the last distance sample and an arbitrary ϕ sample. If ϕ_i is chosen, the normal would be the cross product of the tangent $\omega_i \times \omega_{i+1}$ with length equal to the spinning radius, and ω_1^* with unit length, which gives us

$$\mathbf{n} = \frac{(\omega_{i+1} \times \omega_i^*) \times \omega_1^*|(\mathbf{x}_0^* - \mathbf{x}_i) \times \omega_{i+1}|}{\sin \theta_i} \tag{16}$$

1.1 Photon Surfaces for Low-Order Scattering

Since photon spheres and cones only require one bounce in medium, they can be directly applied for single scattering from any light source. The scaled normal and Jacobian remain exactly the same, we only need to replace the phase function term $f_\omega(\omega_l^{*r})$ with the emission function $L_e(\mathbf{x}_l, \omega_l^{*r}) \cos \theta_l^{*r}$ on the light. Photon cones require a central axis. This can be an arbitrary direction for point lights, and sensible choice for area lights and spot lights are the surface normal and spotlight direction, respectively, since this allows importance sampling the directional emission.

Photon spheres and cones are also applicable to the first bounce off of surfaces. In this case the BSDF is used instead of the phase function or emission function.

Planar Light Source. Area light sources enable additional estimators since they provide additional surface area sampling dimensions for possible analytic integration. To derive single-scattering photon planes, we introduce an orthonormal coordinate system with origin $\mathbf{o}$ and vectors $\mathbf{u}$, $\mathbf{v}$ spanning the whole plane containing the geometry of the light source. Any point within the plane can be denoted with a (u, v) coordinate. For simplicity, we consider a quad

light with $\mathbf{u}$ and $\mathbf{v}$ aligned to the edges and $\mathbf{o}$ being a corner of the quad. In this case, the position of $\mathbf{x}_l$ would be $(\mathbf{o} + u\mathbf{u} + v\mathbf{v})$. Now we configure $\overline{\xi_a}$ by choosing two dimensions from $\{u, v, t_l\}$ together with a dimension s_1. There are three combinations of $\overline{\xi_a}$, including $\{u, t_l, s_1\}$, $\{v, t_l, s_1\}$ and $\{u, v, s_1\}$, from which we derived three types of single scattering photon plane estimators.

By setting $\overline{\xi_a} = \{u, t_l, s_1\}$, we get the (u, t_l)-plane estimator with offset vector $\mathbf{g}$:

$$\mathbf{g}(\overline{\xi_a}) = \underbrace{(\mathbf{o} + v\mathbf{v} + u\mathbf{u} + t_l\omega_l)}_{\mathbf{x}_0(u, t_l)} - \underbrace{(\mathbf{y}_1 + s_1\psi_1)}_{\mathbf{y}_0(s_1)} \tag{17}$$

$$= (\mathbf{o} + v\mathbf{v} - \mathbf{y}_1) + \mathbf{A}\overline{\xi_a}, \tag{18}$$

where $\mathbf{A} = [\mathbf{u}, \omega_l, -\psi_1]$. Therefore, with this combination of integration domain, the change of variable process gives us this Jacobian:

$$\mathbf{J}_{u, t_l, s_1}^{\mathbf{g}} = |\mathbf{n} \cdot \omega_1| = |(\mathbf{u} \times \omega_l) \cdot \psi_1| \tag{19}$$

Following the same procedure with the other two choices will result in the (v, t_l)- and (u, v)-planes with Jacobians:

$$\mathbf{J}_{v, t_l, s_1}^{\mathbf{g}} = |\mathbf{n} \cdot \psi_1| = |(\mathbf{v} \times \omega_l) \cdot \psi_1| \tag{20}$$

$$\mathbf{J}_{u, v, s_1}^{\mathbf{g}} = |\mathbf{n} \cdot \psi_1| = |(\mathbf{u} \times \mathbf{v}) \cdot \psi_1| \tag{21}$$

Since we can position and rotate our orthogonal coordinate system however we like, the (u, t_1)- and (v, t_1)-planes are the same estimator, but under different rotations of the coordinate system. In fact, we can construct a single-scattering plane estimator for an arbitrary rotation angle $\alpha°$ of the coordinate system, which call an $\alpha°$ (v, t_1)-plane.

The Jacobian takes an identical form, simply using the rotated $\mathbf{u}(\alpha)$ in place of $\mathbf{u}$:

$$\mathbf{J}_{u_\alpha, t_l, s_1}^{\mathbf{g}} = |\mathbf{n} \cdot \psi_1| = |(\mathbf{u}(\alpha) \times \omega_l) \cdot \psi_1| \tag{22}$$

Non-planar Light Source. Although it is hard to explicitly parameterize an arbitrarily shaped light surface, extending the (u, v)-plane estimator to such lights only requires the existence of an arbitrary locally area preserving parametrization $\mathbf{x}_l(u, v)$, while the final Jacobian will not depend on u or v. By taking $\overline{\xi_a} = (u, v, s_1)$, we reform the offset vector $\mathbf{g}$ as

$$\mathbf{g}(\overline{\xi_a}) = \underbrace{\mathbf{x}_l(u, v) + \sum_{i=l}^{1} t_i\omega_i}_{\mathbf{x}_0(u, v)} - \underbrace{(\mathbf{y}_1 + s_1\psi_1)}_{\mathbf{y}_0(s_1)} \tag{23}$$

where photon surface $\mathbf{x}_0(u, v)$ has exactly the same shape as the light surface $\mathbf{x}_l(u, v)$. Concluding from Eq. (23), applying a delta function on $\mathbf{g}$ is equivalent to finding the intersections between the query ray $\mathbf{y}_0(s_1)$ and photon surface $\mathbf{x}_0(u, v)$. Thus for uv-surfaces,

$$\mathbf{J}_{u, v, s_1}^{\mathbf{g}} = |\mathbf{n} \cdot \psi_1| \tag{24}$$

where $\mathbf{n}$ is simply the surface normal at the intersecting point.

It is theoretically possible to also extend the (u, t_l)- and (v, t_l)-plane estimators to arbitrary curved light sources. The shape of the resulting photon surfaces would, however, depend on the isocontours of the u, v parameterization. This would be challenging for arbitrary mesh lights, but might be practical for some simple analytic area light shapes.

1.2 Continuous MIS

For a planar area light source, using (u, t)- or (v, t)-planes can introduce singularities when the dot product of the plane normal and the camera ray approaches zero. Since there are infinitely many possible orthogonal parameterizations of the planar light source (through rotation), we combine all possible rotations of (u, t)-planes to compensate for the singularities of any individual orientation. In this section, we derive the Jacobian and MIS weight for these rotated planes. We start with the natural coordinate system of the light source, in which the unit vectors $\mathbf{u}$ and $\mathbf{v}$ are the two edges of the light, and $\mathbf{o}$ marks a corner of the light source. To obtain a new coordinate system, we rotate the edges by an angle α around the center of light source, yielding new axis $\mathbf{u}_\alpha, \mathbf{v}_\alpha$.

The parametrization using the rotated edges along $\mathbf{u}_\alpha, \mathbf{v}_\alpha$ results in a new (u_α, t)-plane (see below).
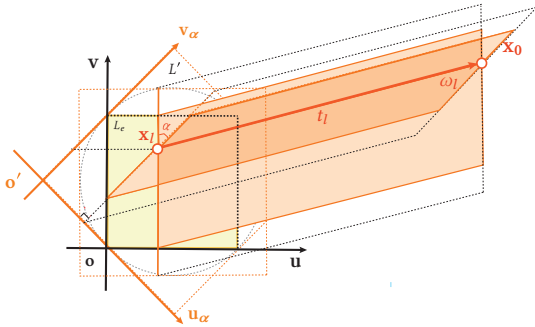
To evaluate the contribution of such a rotated parametrization, we treat it as a new square light source with edges $\mathbf{u}_\alpha, \mathbf{v}_\alpha$ that exactly envelops the original light source. The emission of this light source is zero outside the bounds of the original light, and otherwise matches the original emission.

The estimator score of a (u_α, t)-plane is then

$$\frac{f(\overline{\mathbf{z}})}{p(\overline{\mathbf{z}}, \alpha)} = \frac{f(\overline{\xi}_n)}{p(\overline{\xi}_n)} I_a(u_\alpha, t_l, s_1) \tag{25}$$

$$= \frac{f(\overline{\xi}_n) f_\omega^{1,1} f_t(t_l^*) L_e((u_\alpha^*, v(\alpha)), \omega_l)}{p(\overline{\xi}_n) \mathbf{J}_{u_\alpha, t_l, s_1}^{\mathbf{g}}} \tag{26}$$

Where $p(\overline{\mathbf{z}}_j, \alpha)$ is the probability of generating path $\overline{\mathbf{z}}_j$ with $u(\alpha)t$-plane. Similar to the main paper, we use the estimator contribution as a proxy for the path PDF. Now imagine using m different estimators, corresponding to m randomly rotated parameterizations of the original light source, and combine these m strategies, then by balance heuristic, the pixel intensity is

$$I = \frac{1}{N} \sum_j^N \frac{\dfrac{p(\overline{\mathbf{z}}_j, \alpha)}{f(\overline{\mathbf{z}}_j)}}{\dfrac{1}{m} \sum_{i=1}^m \dfrac{p(\overline{\mathbf{z}}_j, \alpha_i)}{f(\overline{\mathbf{z}}_j)}} \frac{f(\overline{\mathbf{z}}_j)}{p(\overline{\mathbf{z}}_j, \alpha)} \tag{27}$$

$$= \frac{1}{N} \sum_j^N \frac{f(\overline{\mathbf{z}}_j)}{\frac{1}{m} \sum_{i=1}^m p(\overline{\xi}_n) \mathbf{J}_{u_{\alpha_i}, t_l, s_1}^{\mathbf{g}}} \tag{28}$$

Therefore, for a light path $\overline{\mathbf{z}}$, the path contribution evaluated by MIS among randomly-oriented (u_α, t)-planes is

$$\langle I \rangle_{MIS} = \frac{f(\overline{\mathbf{z}})}{p(\overline{\xi}_n) \frac{1}{m} \sum_{i=1}^m \mathbf{J}_{u_{\alpha_i}, t_l, s_1}^{\mathbf{g}}} \tag{29}$$

In the equation above, we regard the term $\frac{1}{m} \sum_{i=1}^m \mathbf{J}_{u_{\alpha_i}, t_l, s_1}^{\mathbf{g}}$ as the Jacobian term of the new estimator produced by MISing. We now take the limit of this expression as the number of strategies m goes toward infinity. This corresponds to performing MIS between the uncountably infinite possible parametrizations of the light source. Recall the properties of the Riemann sum

$$\lim_{m \to \infty} \frac{b-a}{m} \sum_{i=1}^m f(x) = \int_a^b f(x) \mathrm{d}x \tag{30}$$

$$\lim_{m \to \infty} \frac{1}{m} \sum_{i=1}^m f(x) = \frac{1}{b-a} \int_a^b f(x) \mathrm{d}x \tag{31}$$

Then we can compute the denominator of Eq. (29) by

$$\lim_{m \to \infty} \frac{1}{m} \sum_{i=1}^m \mathbf{J}_{u_{\alpha_i}, t_l, s_1}^{\mathbf{g}} = \frac{1}{\pi} \int_0^\pi \mathbf{J}_{u_\alpha, t_l, s_1}^{\mathbf{g}} \mathrm{d}\alpha \tag{32}$$

where $\mathbf{u}_\alpha = \sin \alpha \mathbf{u} + \cos \alpha \mathbf{v}$. Replacing $\mathbf{u}$ with $\mathbf{u}_\alpha$ in Eq. (19) and insert into Eq. (32) we obtain

$$\frac{1}{\pi} \int_0^\pi |((\sin \alpha \mathbf{u} + \cos \alpha \mathbf{v}) \times \omega_1) \cdot \psi_1)| \mathrm{d}\alpha \tag{33}$$

$$= \frac{1}{\pi} \int_0^\pi |k_1 \sin \alpha + k_2 \cos \alpha| \mathrm{d}\alpha \tag{34}$$

$$= \frac{2}{\pi} \sqrt{k_1^2 + k_2^2} \tag{35}$$

where $k_1 = (\mathbf{u} \times \omega_l) \cdot \psi_1, k_2 = (\mathbf{v} \times \omega_l) \cdot \psi_1$. This results in the continuously MIS'd estimator

$$\langle I \rangle_{MIS} = \frac{f(\overline{\mathbf{z}})}{p(\overline{\xi}_n) \frac{2}{\pi} \sqrt{((\mathbf{u} \times \omega_l) \cdot \psi_1)^2 + ((\mathbf{v} \times \omega_l) \cdot \psi_1)^2}} \tag{36}$$

2 ADDITIONAL RESULTS

Comparison with path tracing. In Fig. 4, we compare MISed combinations of some of our estimators with path tracing for single scattering. Our estimators generally lead to lower variance. It is worth pointing out that the noise characteristics of our estimators differ significantly from path tracing: Due to the reuse of the photon paths for many pixels, the noise in our method is spatially coherent and appears smoother, whereas the noise appears at much higher frequencies for path tracing.

Individual comparisons. We additionally compare all of our single scattering estimators (spheres, cones, (u, t)-planes, (u, v)-planes and (v, t)-planes) and various MISed combinations to photon beams and path tracing in Fig. 2 and Fig. 3 for differently sized light sources to better evaluate invidual vs. combined performance.

Estimator overhead. In Table 1, we measure render times of different estimators at equal photon count. This outlines the performance overhead of our new estimators individually, as well as their MIS weighted combinations. The cost of each estimator mostly depends on its average size in screen-space, and hence how many times it is evaluated by a pixel. On the other hand, MIS does not contribute significantly to the estimator cost, which is encouraging.

3-plane MIS. In Fig. 1, we show a full-sized version of Fig. 3 in the main paper, comparing averaging and MIS for different plane estimators.
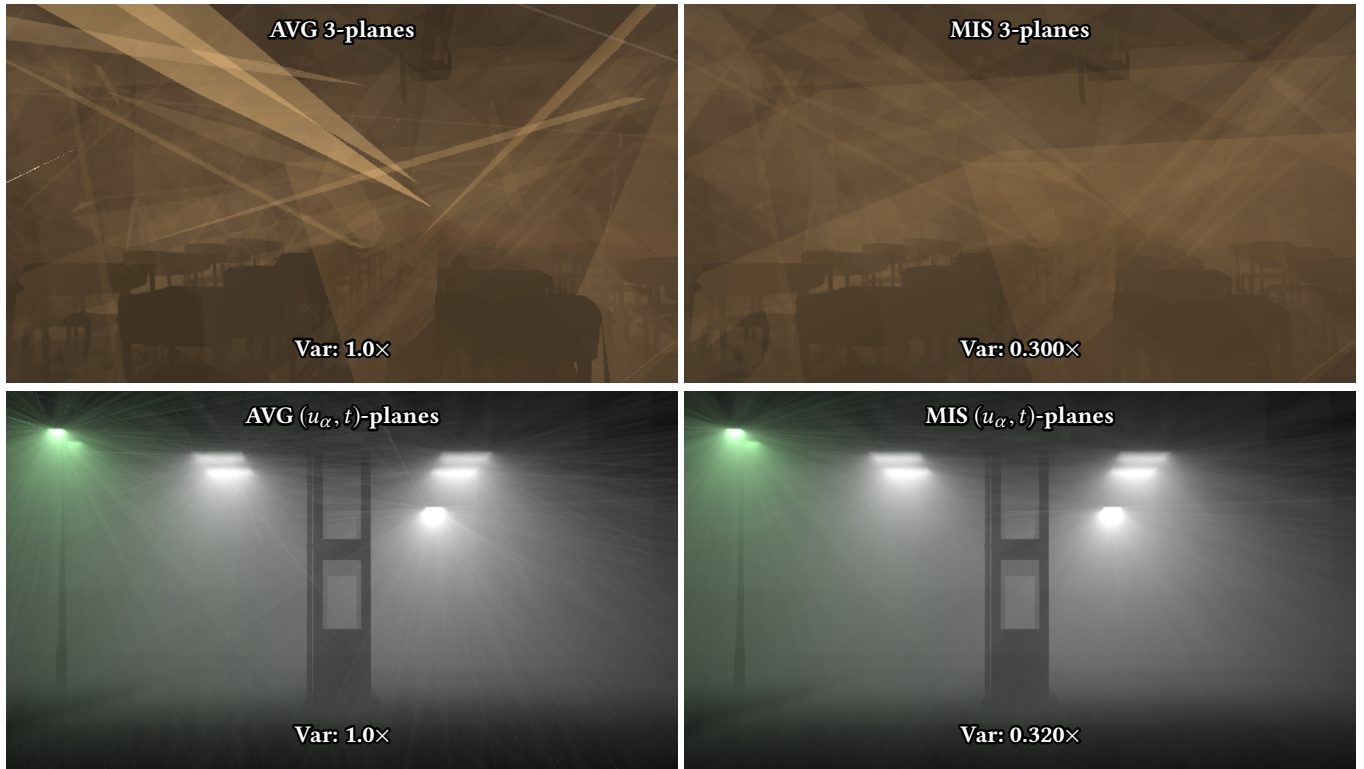

Fig. 1. Full size version of Fig. 3 in the main paper.

Table 1. We compare the overhead of each of our estimators by rendering scenes at equal photon count and measuring the rendering time (in seconds). The computational expense of an estimator depends both on the complexity of intersecting/evaluating it, as well as the number of pixels it overlaps with on average. The latter depends on the scene setup (phase function, occlusion, placement of light sources), but generally dominates the computational cost for estimators such as photon spheres that tend to have a large screen-space extent. On the other hand, the overhead of MIS is not significant.

| Scene	(t_1, t_2)-plane	(t_1, t_3)-plane	(t_2, t_3)-plane	3-planes	Cone	Cylinder	Sphere	Cone Sphere	0D Plane Cone Cylinder	3-planes Cone Cylinder
HorseRoom	97.9	173.5	172.9	166.4	76.3	95.4	58.6	76.9	108.0	154.3
LivingRoom aniso	83.3	112.9	97.4	116.6	152.8	141.2	434.0	377.7	154.3	161.1
LivingRoom iso	121.5	184.1	182.7	188.6	238.5	269.7	381.7	374.1	252.8	260.6
Dining room	88.8	181.4	153.9	149.1	146.4	185.4	346.6	247.6	147.7	169.2
Bathroom	28.0	175.8	177.8	131.1	34.9	46.6	57.1	49.7	38.7	101.6
Kitchen	154.9	228.3	229.7	237.3	340.5	390.1	660.6	613.4	357.6	346.5
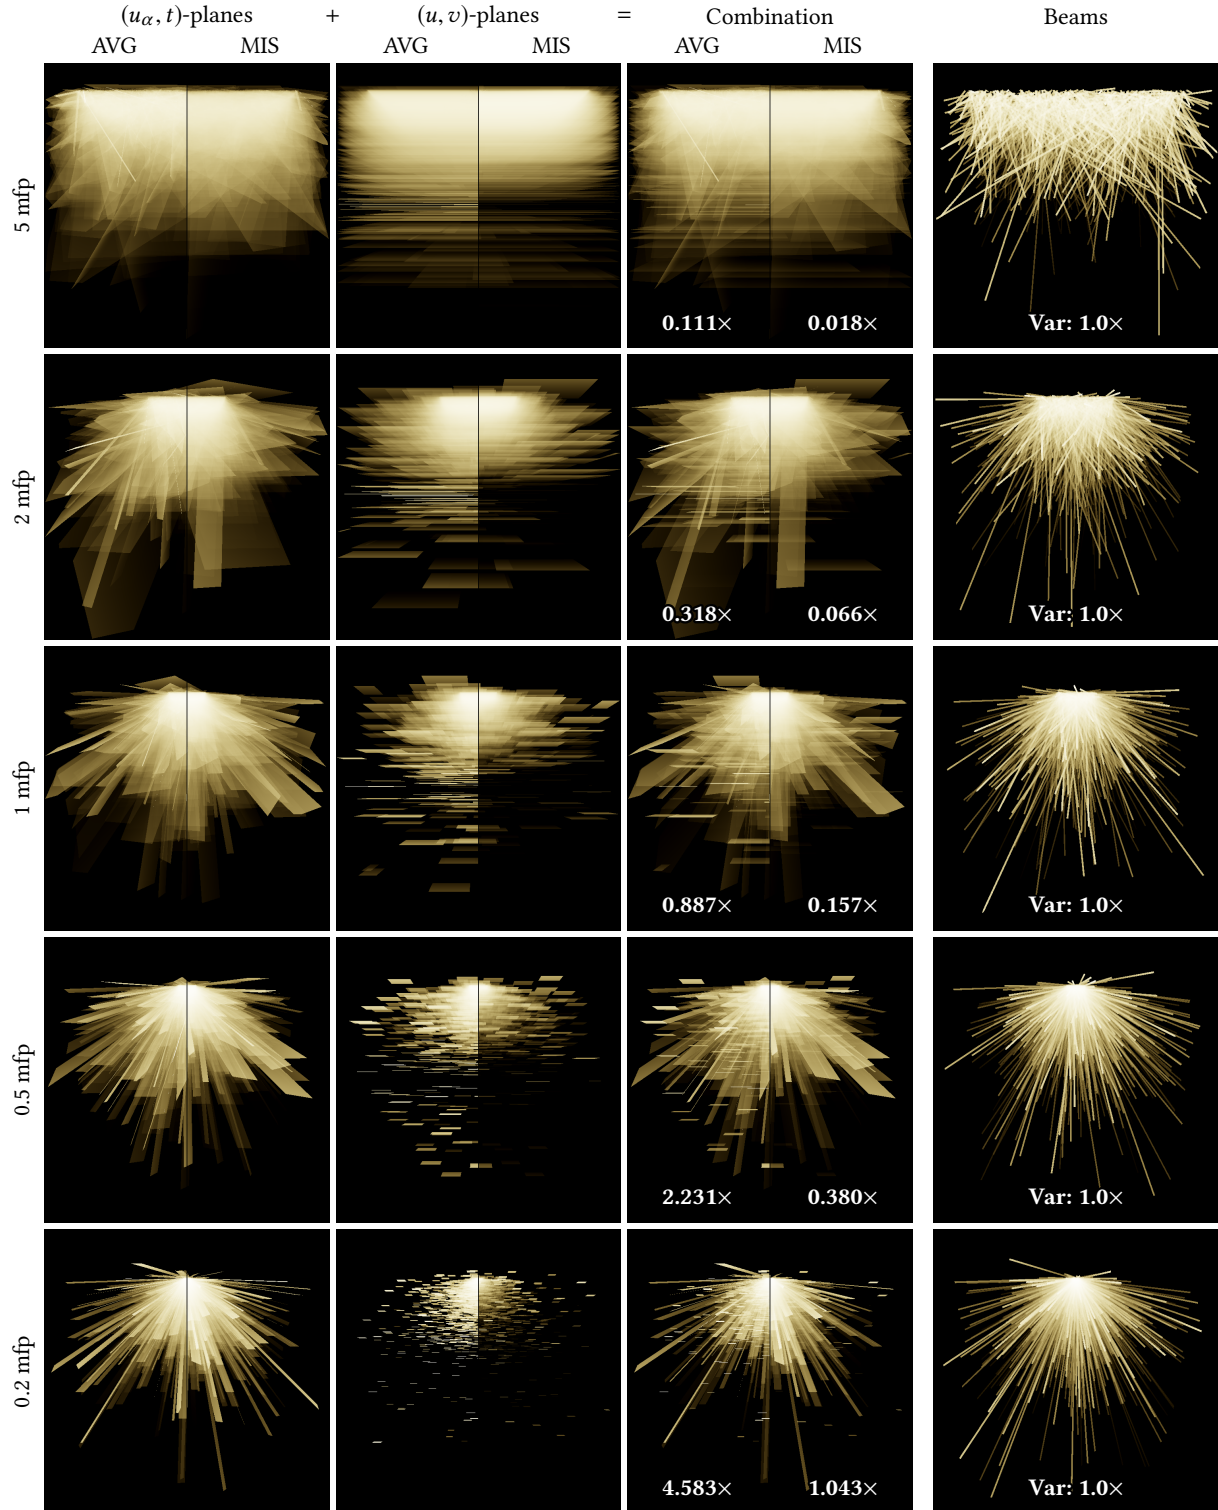
Fig. 2. An expanded version of Fig. 6 from the main paper examining the effect of more light source sizes (rows) on the relative performance of continuously rotated (u_α, t)-planes (first column) and (u, v)-planes (second column). For each estimator we show its contribution individually (left split, AVG), as well as weighted by MIS (right split). The third column shows the weighted and unweighted combination of these two estimators. The (u, v)-planes perform well when viewing a large light source, but for smaller lights their weight diminishes compared to continuously rotated (u_α, t)-planes, which start to resemble the photon beams baseline, though without bias.
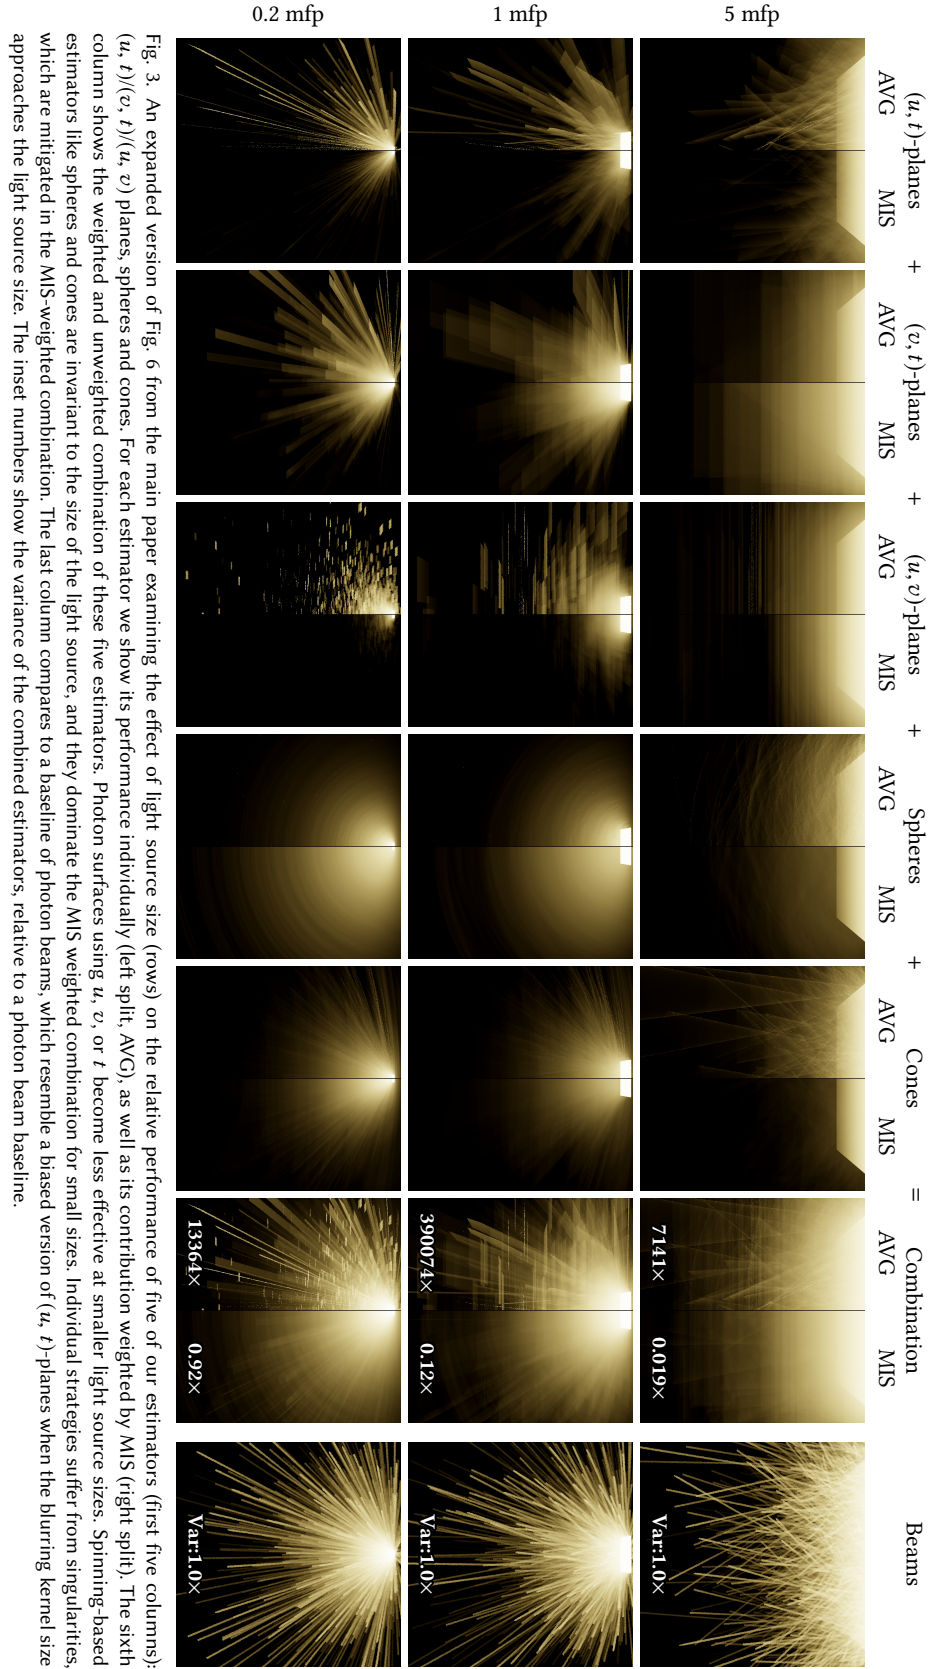
Fig. 3. An expanded version of Fig. 6 from the main paper examining the effect of light source size (rows) on the relative performance of five of our estimators (first five columns): $(u, t)/(v, t)/(u, v)$ planes, spheres and cones. For each estimator we show its performance individually (left split, AVG), as well as its contribution weighted by MIS (right split). The sixth column shows the weighted and unweighted combination of these five estimators. Photon surfaces using u, v, or t become less effective at smaller light source sizes. Spinning-based estimators like spheres and cones are invariant to the size of the light source, and they dominate the MIS weighted combination for small sizes. Individual strategies suffer from singularities, which are mitigated in the MIS-weighted combination. The last column compares to a baseline of photon beams, which resemble a biased version of (u, t)-planes when the blurring kernel size approaches the light source size. The inset numbers show the variance of the combined estimators, relative to a photon beam baseline.

Path Tracing   Ours: MIS$((u, t), (v, t), (u, v))$ Path Tracing   Ours

Var: 1.0×   Var: 0.559×

Var: 1.0×   Var: 0.198×
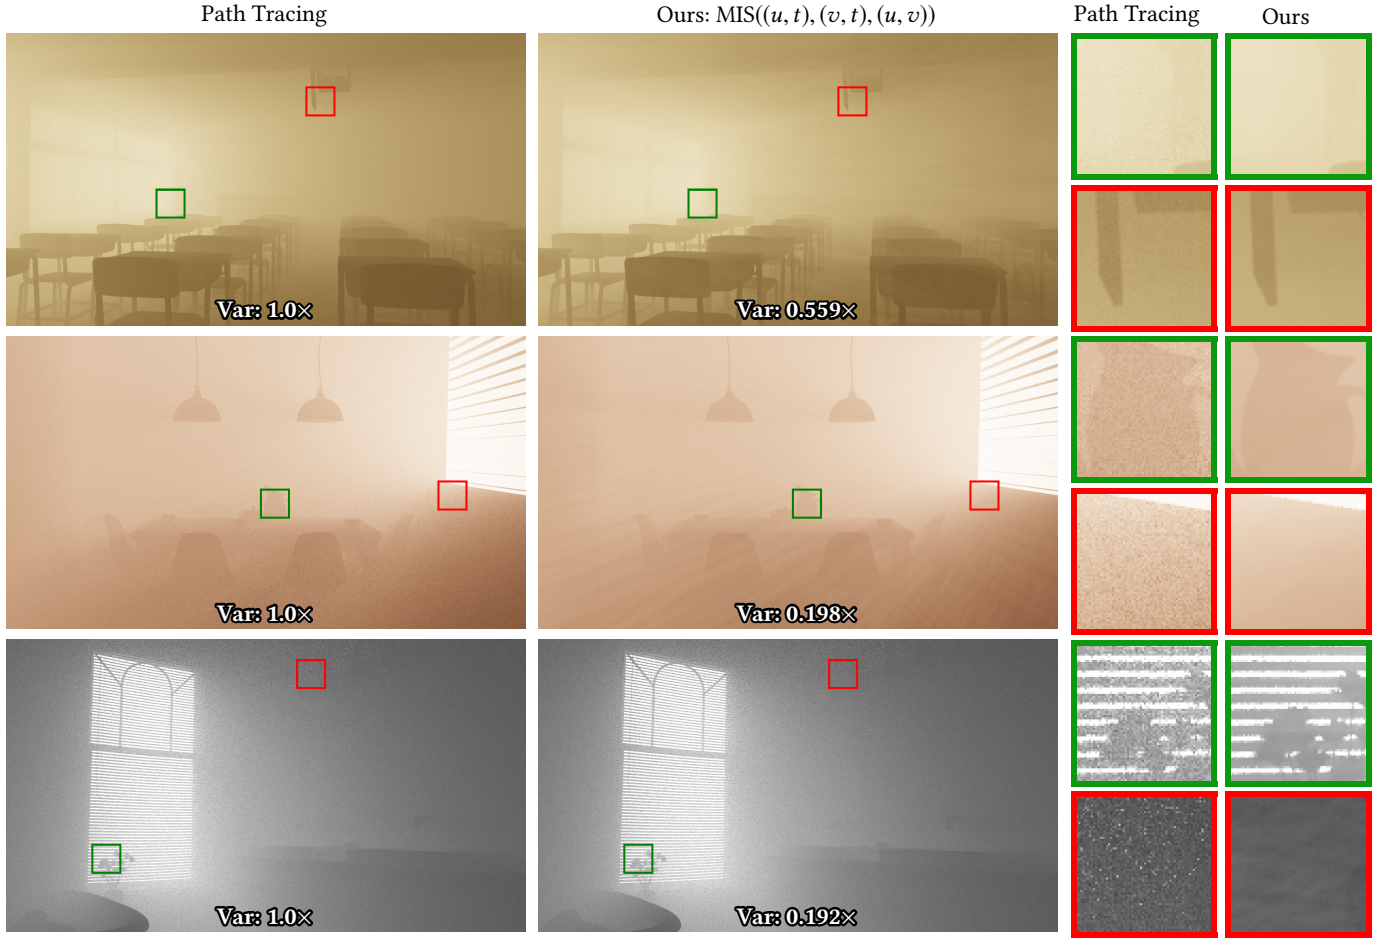
Var: 1.0×   Var: 0.192×

Fig. 4. We perform an equal-time comparison between our MIS weighted (u, t), (v, t) and (u, v) planes (middle column) with path tracing (left column). We show zoomed in comparisons (right columns) as well as variance numbers. We only compare transport that can be handled by our estimators.
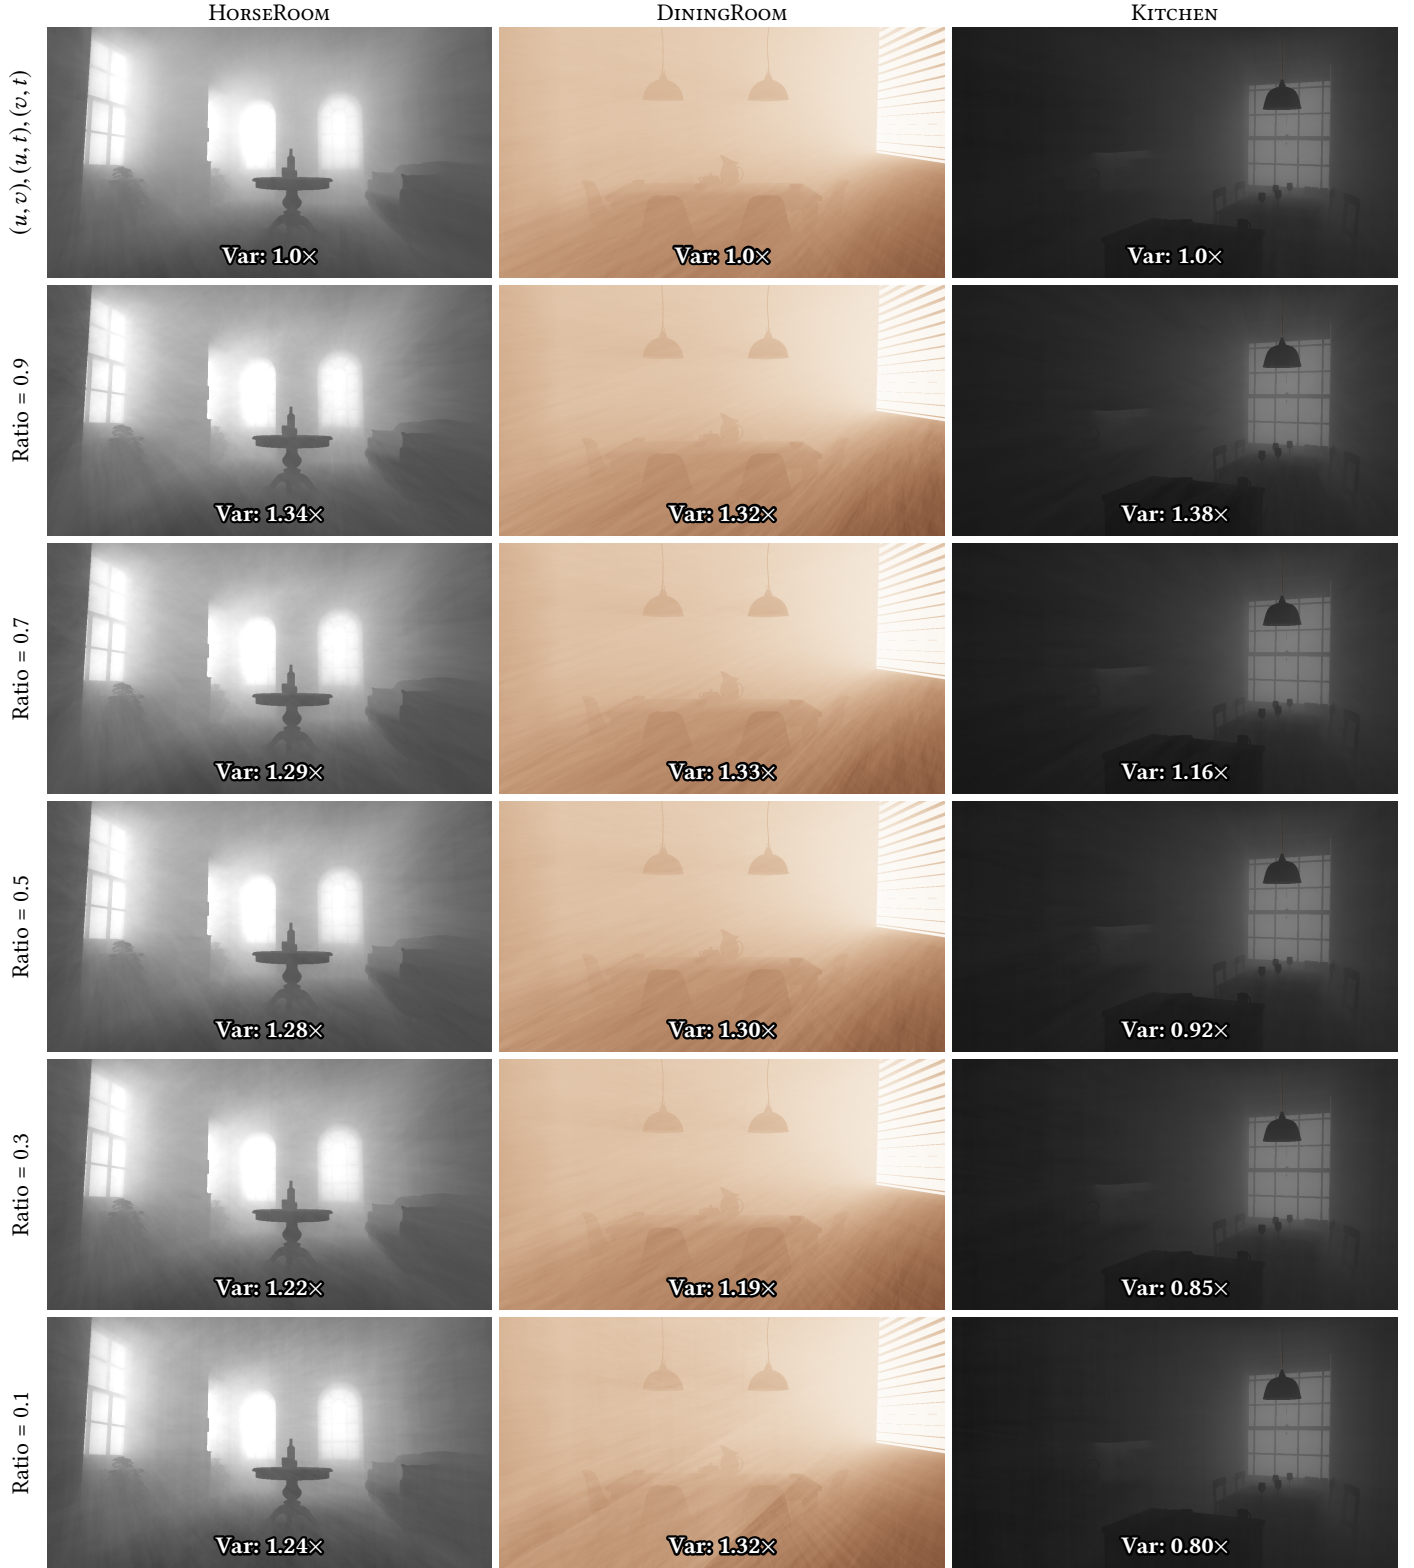
Fig. 5. We perform equal-time comparison between MIS weighted $(u, t),(v, t),(u, v)$ planes (top row) and MIS weighted $(u_\alpha, t), (u, v)$ planes (remaining rows), with the ratio $(u_\alpha, t) : (u, v) + (u_\alpha, t)$ changing from 0.1 to 0.9.