Path-Space Differentiable Rendering of Implicit Surfaces: Supplemental Document

Siwei Zhou
siwei.zhou@ist.osaka-u.ac.jp
Osaka University
Japan

Youngha Chang
ychang@tcu.ac.jp
Tokyo City University
Japan

Nobuhiko Mukai
nmukai@tcu.ac.jp
Tokyo City University
Japan

Hiroaki Santo
santo.hiroaki@ist.osaka-u.ac.jp
Osaka University
Japan

Fumio Okura
okura@ist.osaka-u.ac.jp
Osaka University
Japan

Yasuyuki Matsushita
yasumat@ist.osaka-u.ac.jp
Osaka University
Japan

Shuang Zhao
shz@ics.uci.edu
UC Irvine & NVIDIA
USA

1 DERIVING $\partial_\pi J$

Let us derive $\partial_\pi J$ for our setting, i.e., with $\pi = 0$ and $\mathcal{B} = \mathcal{M}$. We derive $\partial_\pi J$ in the direct illumination setting without loss of generality:

$$I_{\text{direct}} = \int_{\mathcal{M}} f_{\text{direct}}(\boldsymbol{x}) \mathrm{d}A(\boldsymbol{x}). \tag{1}$$

We skip the definition of f_{direct} since it is irrelevant to this derivation.

The corresponding material integral is as follows:

$$I_{\text{direct}} = \int_{\mathcal{B}} \hat{f}_{\text{direct}}(\boldsymbol{p}) \mathrm{d}A(\boldsymbol{p}), \tag{2}$$

where $\hat{f}_{\text{direct}} \coloneqq f_{\text{direct}} \cdot J$ and $J(\boldsymbol{p}) = 1$.

Their derivatives are derived by Zhang et al. [2020] and we first prove the boundary components are equal:

$$\int_{\Delta\mathcal{M}} f_{\text{direct}} V_\partial \, \mathrm{d}\ell(\boldsymbol{x}) = \int_{\Delta\mathcal{B}} \hat{f}_{\text{direct}} V_\partial \, \mathrm{d}\ell(\boldsymbol{p}), \tag{3}$$

where $\Delta\mathcal{M}$ and $\Delta\mathcal{B}$ are the discontinuities of $\mathcal{M}$ and $\mathcal{B}$, respectively.

First, $\Delta\mathcal{M} = \Delta\mathcal{B}$ since $\mathcal{M} = \mathcal{B}$ and the only discontinuities of $\mathcal{M}$ are those in the mutual visibility function as is the case with $\mathcal{B}$, because $\mathcal{M}$ is closed and the normal map of $\mathcal{M}$ is continuous. Second, $f_{\text{direct}} = \hat{f}_{\text{direct}}$ since $\mathtt{X}(\cdot, 0)$ is the identity map.

The curve normal velocity V_∂ is defined as

$$V_\partial(\boldsymbol{x}) \coloneqq \partial_\pi \boldsymbol{x} \cdot \boldsymbol{n}_\partial(\boldsymbol{x}), \quad \boldsymbol{x} \in \Delta\mathcal{M}. \tag{4}$$

Since $\boldsymbol{x} = \mathtt{X}(\boldsymbol{p}, \pi)$ for $\boldsymbol{p} \in \Delta\mathcal{B}$,

$$\partial_\pi \boldsymbol{x} = \partial_\pi \boldsymbol{p} + \boldsymbol{v}(\boldsymbol{p}). \tag{5}$$

Since $\boldsymbol{n}_\partial(\boldsymbol{x}) = \boldsymbol{n}_\partial(\boldsymbol{p})$ and $\boldsymbol{v}(\boldsymbol{p}) \cdot \boldsymbol{n}_\partial(\boldsymbol{p}) = 0$ because $\boldsymbol{v}$ points along the surface normal,

$$V_\partial(\boldsymbol{x}) = V_\partial(\boldsymbol{p}). \tag{6}$$

Therefore, the boundary integrals are equal, and that means the interior ones are also equal:

$$\int_{\mathcal{M}} \partial_\pi f_{\text{direct}} + f_{\text{direct}} \kappa V \, \mathrm{d}A(\boldsymbol{x}) = \int_{\mathcal{B}} \partial_\pi \hat{f}_{\text{direct}} \, \mathrm{d}A(\boldsymbol{p}), \tag{7}$$

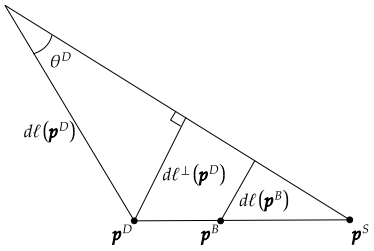

Figure 1: In the tangent plane at $\boldsymbol{p}^{\mathrm{B}}$, we perspectively project $\mathrm{d}\ell(\boldsymbol{p}^{\mathrm{D}})$ onto $\mathrm{d}\ell(\boldsymbol{p}^{\mathrm{B}})$.

where κ is the *total curvature* of $\mathcal{M} = \mathcal{B}$ and

$$V(\boldsymbol{x}) = \boldsymbol{n}_{\mathcal{B}}(\boldsymbol{p}) \cdot \boldsymbol{v}(\boldsymbol{p}), \tag{8}$$

is the "normal velocity" of $\boldsymbol{x} = \mathtt{X}(\boldsymbol{p}, \pi)$ with respect to π at $\pi = 0$. Note κ has the opposite sign in our case than in [Zhang et al. 2020] because we assume our implicit surface has negative sign inside and positive sign outside and thus has its gradient pointing outward aligning with the surface normal.

Since $\mathcal{B} = \mathcal{M}$, the integrands are equal:

$$\partial_\pi f_{\text{direct}} + f_{\text{direct}} \kappa V = \partial_\pi \hat{f}_{\text{direct}}. \tag{9}$$

And since $\hat{f}_{\text{direct}} \coloneqq f_{\text{direct}} \cdot J$ and $J(\boldsymbol{p}) = 1$, by the product rule,

$$\partial_\pi \hat{f}_{\text{direct}} = \partial_\pi f_{\text{direct}} + f_{\text{direct}} \partial_\pi J. \tag{10}$$

Substituting Eq. (10) for the right hand side of Eq. (9),

$$\partial_\pi J = \kappa V. \tag{11}$$

2 DERIVING BOUNDARY SEGMENT

We derive

$$J^B(\boldsymbol{p}^{\mathrm{B}}, \boldsymbol{\omega}^{\mathrm{B}}) \coloneqq \frac{\mathrm{d}\ell_t(\boldsymbol{p}^{\mathrm{S}}) \mathrm{d}\ell_n(\boldsymbol{p}^{\mathrm{S}}) \mathrm{d}\ell(\boldsymbol{p}^{\mathrm{D}})}{\mathrm{d}\theta(\boldsymbol{\omega}^{\mathrm{B}}) \mathrm{d}\ell_n(\boldsymbol{p}^{\mathrm{B}}) \mathrm{d}\ell(\boldsymbol{p}^{\mathrm{B}})} \tag{12}$$

in this section.

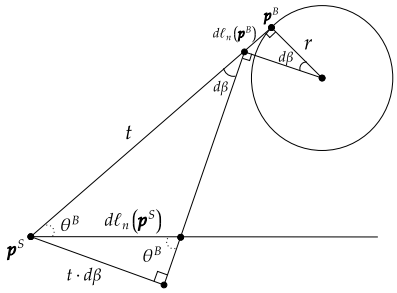

Figure 2: We go from $\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{B})$ to $\mathrm{d}\beta$, to $t \cdot \mathrm{d}\beta$, and finally to $\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{S})$.

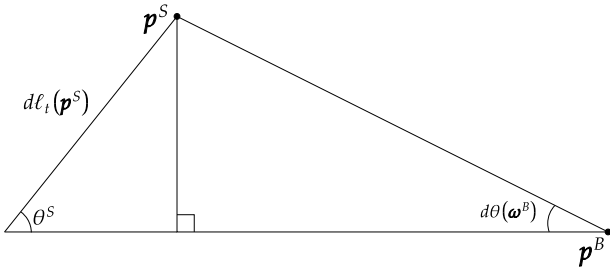

Figure 3: In the tangent plane at $\boldsymbol{p}^\mathrm{B}$, we project $d\theta(\boldsymbol{\omega}^\mathrm{B})$ onto the tangent plane at $\boldsymbol{p}^\mathrm{S}$.

Step one. We first derive $\frac{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{D})}{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{B})}$, illustrated in Figure 1. We perspectively project $\mathrm{d}\ell(\boldsymbol{p}^\mathrm{B})$ onto $\mathrm{d}\ell^\perp(\boldsymbol{p}^\mathrm{D})$. Their ratio is the same as the ratio of their depths, i.e.,

$$\frac{\mathrm{d}\ell^\perp(\boldsymbol{p}^\mathrm{D})}{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{B})} = \frac{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{D}\right|}{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}\right|}. \tag{13}$$

Letting the angle made by $\mathrm{d}\ell(\boldsymbol{p}^\mathrm{D})$ and $\boldsymbol{p}^\mathrm{D}\boldsymbol{p}^\mathrm{S}$ be θ^D, by the Pythagorean theorem,

$$\frac{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{D})}{\mathrm{d}\ell^\perp(\boldsymbol{p}^\mathrm{D})} = \frac{1}{\sin\theta^D}. \tag{14}$$

Combined together,

$$\frac{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{D})}{\mathrm{d}\ell(\boldsymbol{p}^\mathrm{B})} = \frac{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{D}\right|}{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}\right|\sin\theta^D}. \tag{15}$$

Step two. Now we will derive $\frac{\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{S})}{\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{B})}$. We focus on the slice plane spanned by the vector $\boldsymbol{p}^\mathrm{B}\boldsymbol{p}^\mathrm{S}$ and the surface normal vector at $\boldsymbol{p}^\mathrm{B}$ as in Figure 2. Let r be the radius of the osculating circle at $\boldsymbol{p}^\mathrm{B}$, t the distance between $\boldsymbol{p}^\mathrm{S}$ and $\boldsymbol{p}^\mathrm{B}$, and θ^B the angle made by $\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}$ and $\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{S})$. For brevity, let's denote $\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{B})$ and $\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{S})$ as $\mathrm{d}\ell_n^B$ and $\mathrm{d}\ell_n^S$ respectively. Then

$$\frac{\mathrm{d}\ell_n^B}{\mathrm{d}\beta} = r, \tag{16}$$

and

$$\frac{\mathrm{d}\ell_n^S}{\mathrm{d}\beta} = \frac{t}{\sin\theta^B}. \tag{17}$$

Therefore

$$\frac{\mathrm{d}\ell_n^S}{\mathrm{d}\ell_n^B} = \frac{t}{r\cdot\sin\theta^B}. \tag{18}$$

The radius of the osculating circle at $\boldsymbol{p}^\mathrm{B}$ is the reciprocal of the normal curvature κ along $\overrightarrow{\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}}$. Substituting κ for $1/r$ and $|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}|$ for t,

$$\frac{\mathrm{d}\ell_n^S}{\mathrm{d}\ell_n^B} = \frac{\kappa\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}\right|}{\sin\theta^B}, \tag{19}$$

Step three. Lastly, as illustrated in Figure 3, we perspectively project differential angle $d\theta(\boldsymbol{\omega}^\mathrm{B})$ onto the tangent plane at $\boldsymbol{p}^\mathrm{S}$, and the ratio again is the ratio of the depths modulated by a sine factor as in the first step, i.e.,

$$\frac{\mathrm{d}\ell_t(\boldsymbol{p}^\mathrm{S})}{d\theta(\boldsymbol{\omega}^\mathrm{B})} = \frac{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}\right|}{\sin\theta^S}, \tag{20}$$

where θ^S is the angle made by the curve tangent at $\boldsymbol{p}^\mathrm{S}$ and $\overrightarrow{\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}}$. Putting everything together,

$$J^B(\boldsymbol{p}^\mathrm{B}, \boldsymbol{\omega}^\mathrm{B}) \tag{21}$$

$$:= \frac{\mathrm{d}\ell_t(\boldsymbol{p}^\mathrm{D})\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{S})\mathrm{d}\ell_t(\boldsymbol{p}^\mathrm{S})}{\mathrm{d}\ell_t(\boldsymbol{p}^\mathrm{B})\mathrm{d}\ell_n(\boldsymbol{p}^\mathrm{B})d\theta(\boldsymbol{\omega}^\mathrm{B})} \tag{22}$$

$$= \frac{\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{D}\right|\kappa\left|\boldsymbol{p}^\mathrm{S}\boldsymbol{p}^\mathrm{B}\right|}{\sin\theta^D\sin\theta^B\sin\theta^S}. \tag{23}$$

3 DIFFERENTIABLE PIXEL FILTER WEIGHT

Vicini et al. [2022] discovered that making the pixel filter normalization weight differentiable reduces noise in the interior integral. Yu et al. [2022] claim the path-space differentiable rendering formulation cannot benefit from differentiable pixel normalization weight.

We give a mathematical formulation for this method and demonstrate that path-space differentiable rendering can also benefit from pixel normalization weight as well as reparameterization can.

Let the value of a pixel be formulated as follows:

$$I = \frac{\int_\mathcal{P} w(x)L(x)\mathrm{d}x}{\int_\mathcal{P} w(x)\mathrm{d}x}, \tag{24}$$

where $\mathcal{P}$ is the support of the filter and L is the incoming radiance.

To make the filter weight w differentiable, we have to project $\mathcal{P}$ onto some other domain we denote $\mathcal{X}$. With reparameterization, the domain is the unit sphere $\mathbb{S}^2$. With path-space differentiable rendering, the domain is the reference surface $\mathcal{B}$. Let the projection be T. We change the integration domain by the projection T, i.e.,

$$I = \frac{\int_\mathcal{X} w(\mathrm{T}(x))L(\mathrm{T}(x))J_\mathrm{T}\mathrm{d}x}{\int_\mathcal{X} w(\mathrm{T}(x))J_\mathrm{T}\mathrm{d}x}. \tag{25}$$

Let the numerator be F and the denominator be W. Differentiating I gives

$$\partial_\pi I = \frac{\partial_\pi F}{W} - \frac{F\partial_\pi W}{W^2}. \tag{26}$$

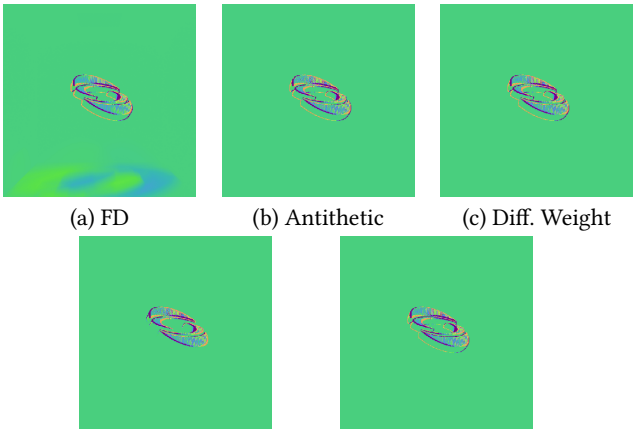

(a) FD (b) Antithetic (c) Diff. Weight

(d) Antithetic (Interior) (e) Diff. Weight (Interior)

Figure 4: We also confirmed that differentiable pixel normalization weight reduces variance in the interior integral and the primary boundary integral. (b) and (c) show differentiable pixel normalization weight achieves less variance with half as many samples as antithetic sampling. (d) and (e) show that differentiable pixel normalization weight makes the interior integral to include part of the primary boundary integral, which reduce the primary boundary integral's value and thus its variance.

The first term on the right hand side is what we would get if we did not let w differentiable. So it follows there is less variance with the second term than without. Note the expectation of the second term is zero since W is a constant and its derivative $\partial_\pi W$ is thus zero.

$\partial_\pi W$ introduces another primary boundary integral. This primary boundary integral needs to be directly sampled and estimated by path-space differentiable rendering methods, but is automatically handled by reparameterization methods. And this is why Yu et al. [2022] find differentiable pixel normalization weight does not apply to path-space differentiable rendering.

REFERENCES

Delio Vicini, Sébastien Speierer, and Wenzel Jakob. 2022. Differentiable Signed Distance Function Rendering. *ACM Trans. Graph.* 41, 4 (2022), 125:1–125:18.

Zihan Yu, Cheng Zhang, Derek Nowrouzezahrai, Zhao Dong, and Shuang Zhao. 2022. Efficient Differentiation of Pixel Reconstruction Filters for Path-Space Differentiable Rendering. *ACM Trans. Graph.* 41, 6 (2022), 191:1–191:16.

Cheng Zhang, Bailey Miller, Kai Yan, Ioannis Gkioulekas, and Shuang Zhao. 2020. Path-space differentiable rendering. *ACM Trans. Graph.* 39, 4 (2020), 143:1–143:19.