Differentiable Rendering of Neural SDFs through Reparameterization: Supplementary Material

1 PIXEL BOUNDARY SAMPLING

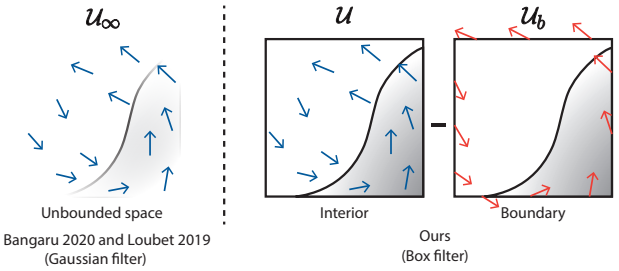

Fig. 1. On the left, existing methods [Bangaru et al. 2020; Loubet et al. 2019] use an unbounded Gaussian filter to avoid the need to handle the boundary of the pixel filter $\mathcal{U}_b$, but this comes at the cost of increased variance in the interior due to the derivative of the Gaussian weights. On the right, our method uses a box filter and explicitly removes the discrepancy in the warp field $\mathcal{V}$ through a boundary integral over $\mathcal{U}_b$.

In Eqn. 3 of the main text, we use a box which implies the pixel domain $\mathcal{U}$ is *bounded*. An implication of this is that we must also consider the boundaries of the pixel filter support (denoted by $\mathcal{U}_b \subset \mathsf{R}^2$) as discontinuities in $\mathcal{U}_{\text{sil}}$. Previously, to avoid this additional complexity, Bangaru et al. [2020] used a Gaussian filter that has infinite support. We have found that this introduces extra variance due to the variation in the pixel filters in the divergence.

We instead keep the box filter as well as exclude the pixel boundary from the area integral $\mathcal{U}_{\text{sil}}$. This means that Eqn. 3 of the main text is no longer valid since the product $(L\mathcal{V})$ does not vanish smoothly at the pixel filter boundary $\mathcal{U}_{\text{b}}$. We must instead rewrite the integral domain as an unbounded space $\mathcal{U}_\infty$. We can further split the unbounded integral into two parts, one inside the pixel filter domain $\mathcal{U}$ and one outside (we omit parentheses here for brevity)

$$I_{\text{sil}} = \int_{\mathcal{U}\setminus\mathcal{U}_{\text{sil}}} \nabla\cdot(L\mathcal{V}) + \int_{(\mathcal{U}_\infty\setminus\mathcal{U})\setminus\mathcal{U}_{\text{b}}} \nabla\cdot(L\mathcal{V})\,. \tag{1}$$

We can then use the divergence theorem on the second area integral to turn it into a boundary integral over $\mathcal{U}_{\text{b}}$

$$I_{\text{sil}} = \int_{\mathcal{U}\setminus\mathcal{U}_{\text{sil}}} \nabla\cdot(L\mathcal{V}) - \oint_{\mathcal{U}_{\text{b}}} L\,(\mathcal{V}\cdot\mathbf{n}_b)\,, \tag{2}$$

where $\mathbf{n}_b$ is the outward pointing normal of the pixel filter boundary, and the negative sign comes from the fact that we consider regions outside of the pixel filter instead of inside. Unlike silhouette boundaries in $\mathcal{U}_{\text{sil}}$, $\mathcal{U}_b$ is easy to sample since it only contains axis-aligned line segments of equal length. Fig. 1 illustrates the difference between using a smooth unbounded filter and using a box filter with pixel boundary sampling.

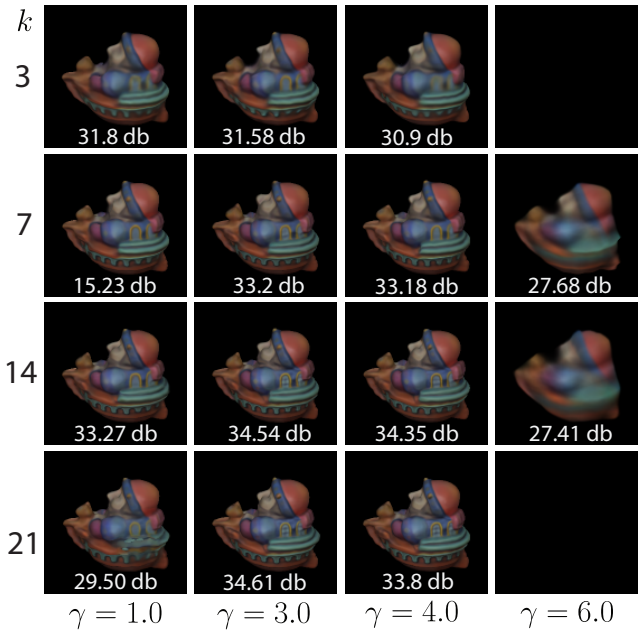

Fig. 2. Ablation Study. We find that $\gamma = 3$ is ideal with both low and and high γ values affecting the reconstruction significantly, with $\gamma \geq 6$ often not converging. For a sphere tracer step count of 22, we find that for subset size $k < 7$, the quality drops quickly but for $k \geq 14$, we see diminishing returns.

2 ABLATION STUDY OF k AND γ

Our pipeline has multiple parameters such as the subset size k and the weight exponent γ. These parameters can affect the variance of the warp field and, as a result, affect the reconstruction quality or time to converge. In Fig. **??**, we look at 16 different combinations of γ and k, and their corresponding reconstructions (this is an extension of Sec. 4.3 in the main text).

Note that there are other parameters such as step count for the sphere tracer N and the adjustment ϵ (used to avoid infinite weights) that can affect the gradient quality. We leave it to future work to further explore these parameters to identify the optimal combination.

3 CORRECTNESS SKETCH OF $\mathcal{V}^q$

To show that $\mathcal{V}^q$ is valid we need to show that it is (i) Continuous and (ii) Boundary consistent. Here, we show that our weights are correct for an ideal C_1 continuous SDF and for an ideal sphere tracer $\mathcal{T}(u)$. Here, $\mathcal{T}(u)$ denotes the infinite series of points generated by the sphere tracer. Note that, in general, none of these points will actually satisfy $f(\mathbf{x}) = 0$ since an ideal sphere tracer never reaches the surface of an ideal SDF. Instead we will deal in *limits*. That is, $\lim_{n\to\infty} f(\mathbf{x}_n) = 0,\ \mathbf{x}_n \in \mathcal{T}(\mathbf{u})$

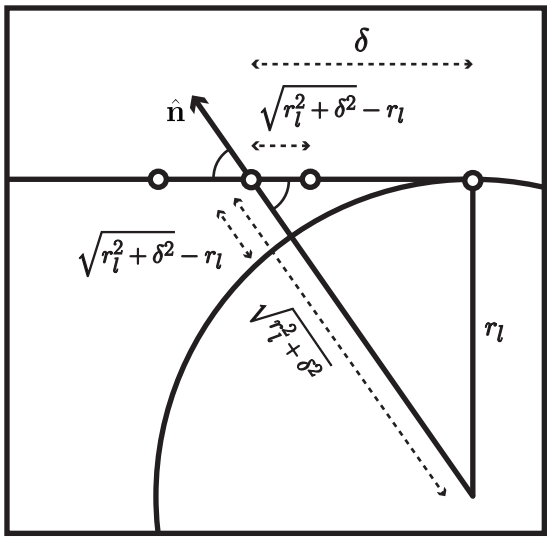

Fig. 3. Circle SDF Bound

Assumption 1. *(General Position Assumption). For a given u, there exist no two points $\mathbf{x}_i, \mathbf{x}_j \in \mathcal{T}(u), i \neq j$ such that both $f(\mathbf{x}_i) = f(\mathbf{x}_j)$ and $\partial_{\mathbf{x}} f(\mathbf{x}_i) = \partial_{\mathbf{x}} f(\mathbf{x}_j)$.*

Lemma 3.1. *(Spherical Lower Bound). There is an ϵ-neighbourhood around every 3D silhouette point $\mathbf{x}_{sil}$, such that the SDF $f(\mathbf{x})$ can be lower bounded by the SDF of a sphere with some fixed radius r_l*

Since we can choose both ϵ and r_l, we can reduce the former and increase the latter until this lemma is satisfied. The only way this scheme fails is if the curvature of the surface is 0 at $\mathbf{x}_{\text{sil}}$. That cannot be the case, because then the surface would be a plane parallel to the ray direction, which means all points along the ray contradict Assumption 1

Lemma 3.2. *(Weight Lower Bound). For a quadrature point along the silhouette ray that is distance δ away from the silhouette point, the weights can be lower bounded.*

$$\lim_{\mathbf{u}\to\mathcal{U}_{sil}} w^q(\mathbf{x}(\mathbf{u}, t_{sil}-\delta)) \geq \left(\sqrt{r_l^2+\delta^2}-r_l+\frac{\delta}{\sqrt{r_l^2+\delta^2}}\right)^{-\gamma}\cdot\left(\sqrt{r_l^2+\delta^2}-r_l\right) \tag{3}$$

Fig. 3 illustrates the slice of the sphere SDF that contains the center of the sphere and the ray direction. For a point δ away from the silhouette point, the sphere SDF at $f(\mathbf{x})$ is $\sqrt{r_l^2+\delta^2}-r_l$. Using the property of similar triangles, the dot product of the normal with the direction is $\delta/\sqrt{r_l^2+\delta^2}$.

Lemma 3.3. *(Unboundedness of the Lower Bound). For $\gamma > 2$, the lower bound in Lemma 3.2 is unbounded in the limit $\delta \to 0$*

$$\lim_{\delta\to 0}\left(\sqrt{r_l^2+\delta^2}-r_l+\lambda_d\cdot\frac{\delta}{\sqrt{r_l^2+\delta^2}}\right)^{-\gamma}\cdot\left(\sqrt{r_l^2+\delta^2}-r_l\right)=\infty \tag{4}$$

To see this, notice that the limit above can be written as

$$\lim_{\delta\to 0}\frac{\sqrt{r_l^2+\delta^2}-r_l}{\left(\sqrt{r_l^2+\delta^2}-r_l+\lambda_d\cdot\frac{\delta}{\sqrt{r_l^2+\delta^2}}\right)^{\gamma}}. \tag{5}$$

Taking the Taylor expansion at $\delta = 0$, for the numerator we have:

$$\sqrt{r_l^2+\delta^2}-r_l=\frac{\delta^2}{2r_l}+O(\delta^4). \tag{6}$$

For the denominator we have:

$$\left(\sqrt{r_l^2+\delta^2}-r_l+\lambda_d\cdot\frac{\delta}{\sqrt{r_l^2+\delta^2}}\right)^{\gamma}=\frac{\delta^{\gamma}}{r_l}+O(\delta^{\gamma+1}) \tag{7}$$

Substituting, we have that the ratio is asymptotically equivalent to:

$$\frac{\frac{\delta^2}{2r_l}+O(\delta^4)}{\frac{\delta^{\gamma}}{r_l}+O(\delta^{\gamma+1})}=\frac{\frac{\delta^{2-\gamma}}{2r_l}+O(\delta^{4-\gamma})}{\frac{1}{r_l}+O(\delta)}, \tag{8}$$

which diverges as long as $\gamma > 2$.

Lemma 3.4. *(Kronecker Delta Behaviour). For a ray exactly at the silhouette, the limiting point of the sphere tracer is assigned all the weight, given $\gamma > 2$ and $\lambda_d > 0$*

$$\lim_{\mathbf{u}\to\mathcal{U}_{sil}}\lim_{\mathbf{n}\to\infty}\frac{w^{(q)}(\mathbf{x}_n(\mathbf{u};t))}{\sum_{\mathbf{x}_i\in\mathcal{T}(u)} w^{(q)}(\mathbf{x}(\mathbf{u};t'))\cdot \mathrm{d}t'}=1 \tag{9}$$

We can show this through contradiction. Since the number of sphere tracer points are countably infinite, let's consider some point $\mathbf{x}_i \in \mathcal{T}(u)$ that is not the limiting point. From 3.4 and 3.3, since $\mathcal{T}(u)$ is an infinite series, we can necessarily find a point $\mathbf{x}_j$, $j > i$ such that $\frac{w^{(q)}(\mathbf{x}_i)}{w^{(q)}(\mathbf{x}_j)} < p$ for *any* $p > 0$ Therefore, in the limit of $n \to \infty$, the normalized weight of $\mathbf{x}_i$ is 0.

This is true for every point $x_i \in \mathcal{T}(u)$ that is not the limiting point itself. Thus, our weights become a discrete version of delta (i.e. the *Kronecker delta*) on the limiting point.

Since the limiting point of $T(u_{\text{sil}})$ is $\mathbf{x}_{\text{sil}}$, it follows from the form of our quadrature weights that, $\lim_{\mathbf{u}\to\mathcal{U}_{\text{sil}}} \mathcal{V}^{\text{q}}(\mathbf{u}) = G(\mathbf{x}(u, t_{\text{sil}}); \theta)^T \partial_{\mathbf{x}} u$. That is, $\mathcal{V}^q$ is boundary consistent.

4 CORRECTNESS SKETCH OF TOP-K WEIGHTS $\bar{w}_k$

Since $T_k(u)$ always contains the k points with the largest weights, boundary consistency follows from the correctness of $\mathcal{V}^{(\text{q})}$. However, *continuity* is non-trivial since the discrete set of points in $\mathcal{T}_k(u)$ can change as u changes. We also only need to consider continuity at non-silhouette points since the resulting warp field is never evaluated exactly at silhouette points.

Lemma 4.1. *Top-K Weight Continuity The weights of the set $\mathcal{T}_k(u)$ are continuous for all $u \notin \mathcal{U}_{sil}$*

We analyze the weights of the points in the set $\mathcal{T}_k(u)$ under two separate cases

(1) Case 1: The indices of points in $\mathcal{T}_k(u)$ change in the infinitesimal neighbourhood around u.
First, note that because of Assumption 1, no two points in $\mathcal{T}(u)$ (and $\mathcal{T}_k(u)$ by extension) can have the same weight. Thus, in an infinitesimally-small neighbourhood, we can assume that there is only one $\mathbf{x}_i \in \mathcal{T}_k(u)$ that is replaced with a new point $\mathbf{x}_j \in \mathcal{T}(u)$, $\mathbf{x}_j \notin \mathcal{T}_k(u)$, as we perturb u. In this neighbourhood, we can assert that:

$$w^{\mathrm{q}}(\mathbf{x}_i) = w^{\mathrm{q}}(\mathbf{x}_j)$$

We can also note that

$$w^{\mathrm{q}}(\mathbf{x}_i) = \min_{\mathbf{x_m} \in \mathcal{T}_k(u)} w^{\mathrm{q}}(\mathbf{x}_m)$$

because by definition of the top-k subset, only the smallest weight is swapped out of the set. However, remember that, from our definition of top-k weights, because we shift every weight by the smallest weight, the smallest weight in the subset is zero, i.e.:

$$w^{\mathrm{k}}(\mathbf{x}_i) = w^{\mathrm{k}}(\mathbf{x}_j) = 0$$

Therefore, because both the swapped points $\mathbf{x}_i$ and $\mathbf{x}_j$ have a weight of 0, the weights of $\mathcal{T}_k(u)$ are continuous in the neighbourhood of u.

(2) Case 2: The order of points in $\mathcal{T}_k(u)$ remain constant in the infinitesimal neighbourhood around u.
Since the points are at the same position in the original series $\mathcal{T}(u)$ and their weights are continuous, it follows that the weights of points in the subseries $\mathcal{T}_k(u)$ are also continuous.

REFERENCES

Sai Bangaru, Tzu-Mao Li, and Frédo Durand. 2020. Unbiased Warped-Area Sampling for Differentiable Rendering. *ACM Trans. Graph.* 39, 6 (2020), 245:1–245:18.

Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing discontinuous integrands for differentiable rendering. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 38, 6 (2019), 228.