

Probabilistic Directed Distance Fields for Ray-Based Shape Representations: Supplemental Material

APPENDIX A

PROOFS OF GEOMETRIC PROPERTIES

For a shape S , we consider visible oriented points (p, v) (i.e., $\xi(p, v) = 1$), such that $q(p, v) = p + d(p, v)v \in S$, unless otherwise specified.

A. Property I: Directed Eikonal Equation

First, note that for any visible (p, v) (with $p \notin S$), there exists an $\epsilon > 0$ such that any $\delta \in (\delta_0 - \epsilon, \delta_0 + \epsilon)$ satisfies $q(p, v) = q(p + \delta v, v) \in S$. In such a case, by definition of the directed distance field (DDF), $d(p + \delta v, v) = d(p, v) - \delta$. Restrict δ to this open interval. The directional derivative along v with respect to position p is then

$$\nabla_p d(p, v)^T v = \lim_{\delta \rightarrow 0} \frac{d(p + \delta v, v) - d(p, v)}{\delta} \quad (23)$$

$$= \lim_{\delta \rightarrow 0} \frac{d(p, v) - \delta - d(p, v)}{\delta} \quad (24)$$

$$= -1, \quad (25)$$

as required.

1) *Gradient Norm Lower Bound:* For any visible (p, v) , since $|\nabla_p d| = \|\nabla_p d\|_2 \|v\|_2 |\cos(\theta(\nabla_p d, v))| = 1$, where $\theta(u_1, u_2)$ denotes the angle between vectors u_1 and u_2 , we also see that $\|\nabla_p d\|_2 = 1/|\cos(\theta(\nabla_p d, v))|$, and thus $\|\nabla_p d\|_2 \geq 1$.

2) *Visibility Gradient:* Consider the same setup as in A-A. The visibility field satisfies a similar property: $\xi(p + \delta v, v) = \xi(p, v)$, as the visibility cannot change when moving along the same view line (away from S). Thus, we get $\nabla_p \xi(p, v)^T v = 0$.

B. Property II: Surface Normals

Consider a coordinate system with origin $q_0 \in S$, with a frame given by $(\hat{i}, \hat{j}, \hat{k})$ where $\hat{k} = n(q_0)$ and $\hat{i}, \hat{j} \in \mathcal{T}_{q_0}(S)$ spans the tangent space at q_0 . Locally, near q_0 , reparameterize S in this coordinate system via

$$S(x, y) = (x, y, f_S(x, y)), \quad (26)$$

where f_S controls the extension of the surface in the normal direction. Notice that

$$\partial_\alpha S|_{q_0} = (\delta_{x\alpha}, \delta_{y\alpha}, \partial_\alpha f_S|_{q_0}), \quad (27)$$

where $\alpha \in \{x, y\}$ and $\delta_{x\alpha}$ is the Kronecker delta, but since $\partial_x S|_{q_0} = \hat{i}$ and $\partial_y S|_{q_0} = \hat{j}$ are in the tangent plane,

$$\partial_\alpha f_S|_{q_0} = 0. \quad (28)$$

Consider any oriented position (p, v) that points to $q(p, v) \in S$ on the surface. Locally, the surface can be reparameterized in terms of (p, v) :

$$q(p, v) = p + d(p, v)v \in S. \quad (29)$$

Yet, using $q = (q_x, q_y, q_z)$, we can write this via Eq. 26 as

$$S(q_x, q_y) = (q_x, q_y, f_S(q_x, q_y)), \quad (30)$$

where q_x is the component along the x direction in local coordinates, which depends on p and v . In other words, the z component of q depends on p via:

$$q_z(p, v) = f_S(q_x(p, v), q_y(p, v)). \quad (31)$$

Let (p_0, v_0) point to q_0 (i.e., $q_0 = q(p_0, v_0)$). Then:

$$\begin{aligned} \partial_{p_i} f_S|_{p_0} &= \underbrace{\partial_{q_x} f_S|_{q_0}}_0 \partial_{p_i} q_x|_{p_0} + \underbrace{\partial_{q_y} f_S|_{q_0}}_0 \partial_{p_i} q_y|_{p_0} \\ &= 0 \quad \forall i \in \{x, y, z\}, \end{aligned}$$

since $\partial_{q_\alpha} f_S|_{q_0} = 0$ using Eq. 28.

Derivatives with respect to position are then given by

$$\partial_{p_z} q_z|_{p_0} = 1 + \partial_{p_z} d(p, v)|_{p_0} v_z = \partial_{p_z} f_S|_{p_0} = 0 \quad (32)$$

$$\partial_{p_\alpha} q_z|_{p_0} = \partial_{p_\alpha} d(p, v)|_{p_0} v_z = \partial_{p_\alpha} f_S|_{p_0} = 0, \quad (33)$$

using Eq. 29 and Eq. 31, with $\alpha \in \{x, y\}$. Thus, using Eq. 32 and Eq. 33,

$$\frac{\partial}{\partial(p_x, p_y, p_z)} d(p, v)|_{p_0} = (0, 0, -1/v_z), \quad (34)$$

in local coordinates. Since the z component here is along the direction of the surface normal, $n(q_0) = n(p_0, v_0)$, this can be rewritten as

$$\nabla_p d(p, v)|_{p_0, v_0} = \frac{-1}{v_z} n(q_0)^T = \frac{-n(q_0)^T}{n(q_0)^T v}. \quad (35)$$

For any C^1 surface S , and point $q_0 \in S$ (that is intersected by visible oriented point (p, v)), we can always construct such a coordinate system, so we can more generally write:

$$\nabla_p d(p, v) = \frac{-n^T}{n^T v}. \quad (36)$$

C. Property III: Gradient Consistency

Consider the same setup (coordinate system) and notation as in the Proof of Property II above (§A-B). Since $v \in \mathbb{S}^2$, it suffices to consider an infinitesimal rotational perturbation of some initial view direction v_0 :

$$dR(t) = [\omega]_\times dt, \quad (37)$$

where $[\omega]_\times$ is a skew-symmetric cross-product matrix of some angular velocity vector ω , so that $\tilde{v} = (I + dR(t))v_0$ and $dv = \tilde{v} - v_0 = [\omega]_\times dt v_0$ is the change in the view (given dt) and a velocity of $u := \partial_t v = [\omega]_\times v_0$. The change in surface position, $q(p, v(t)) = p + d(p, v(t))v(t)$, with respect to t is then

$$\begin{aligned} \partial_t q &= v \partial_t d + d \partial_t v \\ &= (\partial_v d \partial_t v) v + d \partial_t v \\ &= (\partial_v du) v + du. \end{aligned}$$

The z component, in local coordinates, is then

$$\begin{aligned} \partial_t q_z|_{q_0} &= \partial_t f_S(q_x, q_y)|_{q_0} \\ &= \partial_{q_x} f_S|_{q_0} \partial_t q_x + \partial_{q_y} f_S|_{q_0} \partial_t q_y \\ &= 0, \end{aligned}$$

via Eq. 28. Thus, $\partial_t q_z|_{q_0} = (\partial_v du)v_z + du_z = 0$, meaning

$$du^T n = -(\partial_v du)v^T n \quad (38)$$

using $u_z = u^T n$ and $v_z = v^T n$ at q_0 . Recalling that $\partial_p d = -n^T/(n^T v)$ and rearranging, we get

$$\partial_v du = -d \frac{n^T}{v^T n} u = d \partial_p du. \quad (39)$$

Notice that $u = \partial_t v = [\omega]_{\times} v_0$ is orthogonal to v_0 and the arbitrary vector ω . Hence, for any visible oriented point (p, v) away from viewing the boundary of S (i.e., where $n^T v = 0$) and for all $\omega \in \mathbb{R}^3$, we have

$$\nabla_v d[\omega]_{\times} v = d \nabla_v d[\omega]_{\times} v. \quad (40)$$

This constrains the derivatives with respect to the view vector to be closely related to those with respect to position, along any directions not parallel to v .

1) *Alternative Expression:* Note that the inner product with δ_v primarily serves to restrict the directional derivative in valid directions of v . A cleaner expression can be obtained with a projection operator, which removes directional components parallel to v .

First, let us define d to normalize v :

$$d(p, v) := d\left(p, \frac{v}{\|v\|_2}\right), \quad (41)$$

for all $v \in \mathbb{R}^3 \setminus \{0\}$. Then, consider a perturbation along the view direction of size $|\delta| < 1$:

$$d(p_0, v_0 + \delta v_0) = d(p_0, v_0(1 + \delta)) = d(p_0, v_0), \quad (42)$$

with $\|v_0\| = 1$ and using Eq. 41 for the last step. This means the directional derivative along v must satisfy

$$\partial_v d(p, v)v = 0. \quad (43)$$

In the previous section, we showed that $\partial_v du = d \partial_p du$ for all $u \perp v$. Let $\mathcal{P}_v = I - vv^T$ be the orthogonal projection removing components parallel to v . Then we can rewrite the result of the previous section as $\partial_v d \mathcal{P}_v = d \partial_p d \mathcal{P}_v$. But $\partial_v d \mathcal{P}_v = \partial_v d - \partial_v d v v^T = \partial_v d$ by Eq. 43. Thus, we may write

$$\partial_v d = d \partial_p d \mathcal{P}_v, \quad (44)$$

which agrees with Eq. 40.

D. Property V: Local Differential Geometry

Given a visible oriented point (p, v) , we have shown that the surface normal $n(p, v)$ on S (at $q = p + d(p, v)v$) is computable from $\nabla_p d$ (see Supp. A-B). Curvatures, however, require second-order information.

We first construct a local coordinate system via n , by choosing two tangent vectors at q : $t_x, t_y \in \mathcal{T}_q(S)$, where $\|t_x\|_2 = 1$ and $t_x^T t_y = 0$. In practice, this can be done by sampling Gaussian vectors, and extracting an orthogonal tangent basis from them. We can then reparameterize the surface near $q_0 = q(p_0, v_0)$ (with surface normal $n_0 = n(q_0)$) via $S(x, y) = q_0 + xt_x + yt_y + f_S(x, y)n_0$, where x and y effectively control the position on the tangent plane. Alternatively, we can write $S(x, y) = q(p(u), v_0)$, which parameterizes S

about the oriented point (or viewpoint) $p(u)$, where $u = (x, y)$, and $p(u) = p_0 + xt_x + yt_y$. Notice that $p(u)$ is essentially a local movement of p parallel to the tangent plane at $q_0 \in S$ (which is “pointed to” by (p_0, v_0)). Note that p_0, v_0 , and q_0 are fixed; only u and $p(u)$ are varied. Further, notice that the plane defined by $p(u)$ (in the normalized tangent directions t_x and t_y) is parallel to the surface tangent plane at q_0 (and thus orthogonal to n_0), but *not* necessarily perpendicular to v .

Using this local frame on $\mathcal{T}_{q_0}(S)$, the first-order derivatives of the surface are

$$\partial_i S|_{u=0} = \partial_i (p(u) + v_0 d(p(u), v_0))|_{u=0} \quad (45)$$

$$= t_i + \partial_i d(p(u), v_0)|_{u=0} v_0 \quad (46)$$

$$= t_i + (\nabla_p d(p, v_0)t_i)v_0 \quad (47)$$

where $i \in \{x, y\}$ and p_j is the j th component of p . Then the metric tensor (first fundamental form) is given by

$$g_{ij} = \partial_i S|_{u=0}^T \partial_j S|_{u=0} \quad (48)$$

$$= \delta_{ij} + c_i c_j + c_i v_0^T t_j + c_j v_0^T t_i, \quad (49)$$

where $i, j \in \{x, y\}$, $u = (x, y)$, δ_{ij} is the Kronecker delta function, and $c_k = \nabla_p d t_k$. Notice that $\nabla_p d$ is parallel to n (see Property II), and thus orthogonal to t_1 and t_2 . Thus, $c_x = c_y = 0$, and so, in theory, $g_{ij} = \delta_{ij}$, but this may not hold for a learned DDF, if one computes directly with Eq. 49. However, in practice, we use the theoretical value (where $g = I_2$ in local coordinates), as deviations from it are due to error only (assuming correct normals).

Next, the shape tensor (second fundamental form) \mathbb{I}_{ij} can be computed from the second-order derivatives of the surface [95] via:

$$\mathbb{I}_{ij} = n_0^T \partial_{ij} S|_{u=0} \quad (50)$$

$$\begin{aligned} &= n_0^T \partial_{ij} (p(u) + v_0 d(p(u), v_0))|_{u=0} \\ &= n_0^T v_0 \partial_{ij} d(p(u), v_0)|_{u=0} \\ &= n_0^T v_0 \partial_i \sum_k \partial_{p_k} d(p(u), v_0) \underbrace{[\partial_j p(u)]_k}_{[t_j]_k}|_{u=0} \\ &= n_0^T v_0 \sum_k \partial_i \partial_{p_k} d(p(u), v_0) [t_j]_k|_{u=0} \\ &= n_0^T v_0 \sum_k [t_j]_k \sum_\ell \partial_{p_\ell} \partial_{p_k} d(p(u), v_0) [\partial_i p(u)]_\ell|_{u=0} \\ &= n_0^T v_0 \sum_k [t_j]_k \sum_\ell \underbrace{\partial_{p_\ell, p_k} d(p(u), v_0)}_{\mathcal{H}_p[d]_{k\ell}} [t_i]_\ell|_{u=0} \\ &= (t_j^T \mathcal{H}_p[d] t_i) n_0^T v_0, \end{aligned} \quad (51)$$

where $i, j \in \{x, y\}$, $u = (x, y)$, $n_0 = n(q_0)$, $[t_i]_k$ is the k th component of t_i , $\mathcal{H}_p[d]$ is the Hessian of d with respect to p (at $u = 0$), and $p(u) = p_0 + xt_x + yt_y$. Note that our parameterization of the surface using the DDF (i.e., $q(p(u), v_0)$) defines surface deviations in terms of v_0 ; the $n_0^T v_0$ term undoes this effect, rewriting the deviation in terms of n instead.

Curvatures can then be computed via the shape tensor (see, e.g., [95]). Gaussian curvature is given by $\mathcal{C}_K = \det(\mathbb{I})/\det(g)$, while mean curvature is written $\mathcal{C}_H = \text{tr}(\mathbb{I}g^{-1})$. Notice that these quantities can be computed for

any visible (p, v) , using only $d(p, v)$ and derivatives of d with respect to p (e.g., with auto-differentiation). Thus, the curvatures of any visible surface point can be computed using only local field information at that oriented point.

E. Neural Depth Renderers as DDFs

A natural question is how to parallelize a DDF. Consider a function $f : \Lambda_\Pi \times \Lambda_S \rightarrow \mathcal{I}_{d,\xi}^m$, where $\mathcal{I}_{d,\xi} = \mathbb{R}_+ \times [0, 1]$, from continuous camera parameters (i.e., elements of Λ_Π) and some space of shapes Λ_S to an m -pixel depth and visibility image. Then f is implicitly a conditional DDF, as each pixel value can be computed as $(d(p, v), \xi(p, v))$, where p and v are a fixed function of $\Pi \in \Lambda_\Pi$. In other words, the camera intrinsics and extrinsics determine the p and v value for each pixel, with p determined by the centre of projection (camera position) and v determined by the direction corresponding to that pixel, from the centre of projection out into the scene. We may thus regard the depth pixel value as a function of p and v , via the camera parameters. This holds regardless of the architecture of f . Thus, all the properties of DDFs hold in these cases as well: for example, for a depth image I_d , camera position C_p and viewpoint C_v , Property II relates $\partial_{C_p} I_d$ to the surface normals image, while Property III constrains $\partial_{C_v} I_d$. We believe this point of view can improve the framework for differentiable neural rendering of geometric quantities (e.g., [56]–[59]).

APPENDIX B SINGLE-ENTITY DDFs

A. Mesh Data Extraction

We display visualizations of each data type in Fig. 7. We briefly describe how each is computed (note that we only need to obtain (p, v) , after which (ξ, d) can be obtained via ray-triangle intersection):

- Uniform (U): simply sample $p \sim \mathcal{U}[\mathcal{B}]$ and $v \sim \mathcal{U}[\mathbb{S}^2]$.
- At-surface (A): start with $q_0 \sim \mathcal{U}[S]$ (obtained via area-weighted sampling) and $v_0 \sim \mathcal{U}[\mathbb{S}^2]$. Sample p on the line between q_0 and its intersection with \mathcal{B} along v_0 . Set $v = -v_0$. Note that the final output data may not actually intersect q_0 (since one may pass through a surface when sampling p).
- Bounding (B): sample $p \sim \mathcal{U}[\partial\mathcal{B}]$ and $v \sim \mathcal{U}[\mathbb{S}^2]$, but restrict v to point to the interior of \mathcal{B} .
- Surface (S): simply use $p \sim \mathcal{U}[S]$ and then take $v \sim \mathcal{U}[\mathbb{S}^2]$.
- Tangent (T): the procedure is the same as for A-type data, except we enforce v_0 to lie in the tangent plane $\mathcal{T}_{q_0}(S)$.
- Offset (O): take a T-type (p, v) and simply do $p \leftarrow p + \varsigma_O \epsilon_O n_0$, where n_0 is the normal at q_0 , and we set $\epsilon_O = 0.05$ and sample $\varsigma_O \sim \mathcal{U}[\{-1, 1\}]$.

Since we assume \mathcal{B} is an axis aligned box (with maximum length of 2; i.e., at least one dimension is $[-1, 1]$), sampling positions on it, or directions with respect to it, is straightforward. See also §V-A1, which examines the effect of ablating each data type.

B. Single-Entity Fitting

Given a mesh, we first extract data of each type, obtaining (p, v) tuples in the following amounts: 250K (A and U) and 125K (B, T, O, and S). Since rendering outside \mathcal{B} uses query points on $\partial\mathcal{B}$, we bias the sampling procedure for T, A, and O data, such that 10% of the p -values are sampled from $\partial\mathcal{B}$. For each minibatch, we sample 6K (A and U) and 3K (B, T, O, and S) points, across data types. In addition to these, we sample an additional 1K uniformly random oriented points per minibatch, on which we compute only regularization losses (\mathcal{L}_V and \mathcal{L}_{DE}), for which ground truth values are not needed.

We note that not all loss terms are applied to all data types. As discussed in §IV, the transition loss \mathcal{L}_T is only applied to S and T type data (since we do not want spurious field discontinuities, due to the field switching between components, except when necessary). As briefly noted in Property II, the DDF gradient $\nabla_p d$ (and hence field-derived surface normals) are not well-defined when n and v are orthogonal. The gradients are also not well-defined on S-type data (since d is explicitly discontinuous for $p \in S$, and thus $\nabla_p d$ does not exist there). Hence, we do not apply the directed Eikonal regularization \mathcal{L}_{DE} or the normals loss \mathcal{L}_n to S, T, or O data. Similarly, the weight variance loss \mathcal{L}_V (designed to reduce weight field entropy) should not be applied to S, T, or O samples, as the weight field should be transitioning near those samples (and thus have a higher pointwise Bernoulli variance). The remaining losses are applied on all data types.

We then optimize Eq. 2 via Adam [96] (learning rate: 10^{-4} ; $\beta_1 = 0.9$, $\beta_2 = 0.999$), using a reduction-on-plateau learning rate schedule (reduction by 0.9, minimum of 5K iterations between reductions). We run for 100K iterations, using $\gamma_d = 5$, $\gamma_\xi = 1$, $\gamma_n = 10$, $\gamma_V = 1$, $\gamma_{E,d} = 0.05$, $\gamma_{E,\xi} = 0.01$, and $\gamma_T = 0.25$. Note that we double γ_d on A and U data. The field itself is implemented as a SIREN with seven hidden layers, initialized with $\omega_0 = 1$, each of size 512, mapping $(p, v) \in \mathbb{R}^6 \rightarrow (\{d_i\}_{i=1}^K, \{w_i\}_{i=1}^K, \xi) \in \mathbb{R}^{2K+1}$ (note that the set of w_i 's has $K - 1$ degrees of freedom). We use $K = 2$ delta components. An axis-aligned bounding box is used for \mathcal{B} . Since $w_2 = 1 - w_1$, we output only w_1 and pass it through a sigmoid non-linearity. We apply ReLU to all d_j outputs, to enforce positive values, and sigmoid to ξ . We use the implementation from [97] for SIREN. In terms of data, we test on shapes from the Stanford Scanning Repository [98] (data specifically from [99]–[101]).

C. Rendering Efficiency

As noted in the main paper (§V-A4), we tried to use approximately comparable MLP-based architectures. Both models were run on an NVIDIA GeForce RTX 2080 Ti. We used the default settings of DIST [7] for rendering, including no aggressive marching, and used 50 marching steps, the lowest value we found in the released code. To mitigate the effect of the choice of shape and the camera position, we ran DIST on five shapes, each with two azimuths (90 degrees apart) and at two distances from the shape (0.25 and 2.5). At the far distance, DIST renders in 5.34 seconds (standard deviation 0.35); at the close distance, DIST takes 9.78 seconds (standard

deviation 1.58), due to differences in which heuristics can be applied to optimize runtime. In contrast, the DDF renders in 0.006 seconds, with negligible difference across shapes or cameras, since we call the MLP once regardless of whether or where the shape exists along that particular ray.

APPENDIX C UDFS AND MDFs

A. MDF Extraction

We obtain v^* via a fitting procedure, with a forward pass similar to composition (see §V-A). Starting from an already trained (P)DDF, we define a small new SIREN network g_v (five hidden layers, each size 128), which maps position to a set of K_c candidate directions, such that $g_v : \mathbb{R}^3 \rightarrow \mathbb{S}^{2 \times K_c}$. Given a position p and candidates $\{v_i^*\}_{i=1}^{K_c} = g_v(p)$, we can compute $\zeta_{v^*} = \{d(p, v_i^*(p)), \xi(p, v_i^*(p))\}_{i=1}^{K_c}$. To obtain a UDF depth estimate, we use

$$\widehat{\text{UDF}}(p) = \sum_i \omega_{\zeta_{v^*}}^{(i)}(p) d(p, v_i^*(p)), \quad (52)$$

where the weights $\omega_{\zeta_{v^*}}$ are based on those from our explicit composition method:

$$\omega_{\zeta_{v^*}}(p) = \text{Softmax} \left(\left\{ \frac{\eta_T^{-1} \xi(p, v_i^*(p))}{\varepsilon_s + d(p, v_i^*(p))} \right\}_i \right). \quad (53)$$

We use the same η_T and ε_s as for composition. Note that this formulation essentially takes the candidate directions, computes their associated distances, and then (softly) chooses the one that has lowest distance value while still being visible. We can also obtain $v^*(p)$ via these weights, using

$$\tilde{v}^*(p) = \sum_i \omega_{\zeta_{v^*}}^{(i)}(p) v_i^*(p) \quad (54)$$

$$v^*(p) = \frac{\tilde{v}^*(p)}{\|\tilde{v}^*(p)\|_2}, \quad (55)$$

Note that, in the ideal case, $v^*(p) = -\nabla_p \text{UDF}(p)^T$, meaning we could compute it with a backward pass, but we found such an approach was noisier in practice.

To train g_v to obtain good candidates, we use the following loss:

$$\mathcal{L}_{v^*} = \frac{1}{K_c} \sum_i d(p, v_i^*(p)) - \xi(p, v_i^*(p)) \quad (56)$$

$$+ \frac{2\tau_n}{K_c^2 - K_c} \sum_i \sum_{j \neq i} v_i^*(p)^T v_j^*(p) \quad (57)$$

$$+ \frac{\tau_d}{K_c} \sum_i [v_i^*(p)^T n(p, v_i^*(p)) + 1]^2 \quad (58)$$

where the first sum encourages obtaining the direction with minimum depth that is visible, the second prevents collapse of the candidates to a single direction (e.g., local minima), and the third encodes alignment with the local surface normals (similar to the directed Eikonal regularizer; see Supp. C-B). Note that the first sum (Eq. 56) includes a d term, which linearly penalizes longer distances, and a $-\xi$ term, which pushes the visibility (i.e., probability of surface existence in the direction $v_i^*(p)$, from position p) to be high. We use $K_c = 5$, $\tau_n = 5 \times 10^{-3}$, and $\tau_d = 0.1$. Optimization was run for 10^4

iterations using Adam (LR: 10^{-4} ; $\beta_1 = 0.9$, $\beta_2 = 0.999$). Each update used 4096 points, uniformly drawn from the bounding volume ($p \sim \mathcal{U}[\mathcal{B}]$).

B. Surface Normals of the MDF v^*

Let v^* be defined as in §V-A3 and choose p such that $v^*(p)$ is not multivalued (i.e., neither on S nor on the medial surface of S). Recall that, by definition,

$$\text{UDF}(p) = d(p, v^*(p)) \quad (59)$$

and

$$n(p, v) = n(q(p, v)) = \frac{\nabla_p d(p, v)^T}{\|\nabla_p d(p, v)\|} \quad (60)$$

where $n(p, v)^T v < 0$ (see Property II). Then, by Eq. 59,

$$\begin{aligned} \nabla_p \text{UDF}(p) &= \nabla_p d(p, v^*(p)) \\ &= \|\nabla_p d(p, v^*(p))\| \frac{\nabla_p d(p, v^*(p))}{\|\nabla_p d(p, v^*(p))\|} \\ &= \|\nabla_p \text{UDF}(p)\| n(p, v^*(p))^T \\ &= n(p, v^*(p))^T, \end{aligned}$$

via the S/UDF Eikonal equation in the last step.

Finally, by the directed Eikonal property (I), we have:

$$\nabla_p \text{UDF}(p) v^*(p) = \nabla_p d(p, v^*(p)) v^*(p) = -1, \quad (61)$$

meaning $\nabla_p \text{UDF}(p) = -v^*(p)^T$. Indeed, recall that the gradient of a UDF has (1) unit norm (i.e., satisfies the Eikonal equation) and (2) points away from the closest point on S , which is precisely the definition of $-v^*(p)$.

APPENDIX D SINGLE-IMAGE 3D RECONSTRUCTION

A. Model Architecture and Training Details

Architecture and Optimization. Our CPDDF is a modulated SIREN [74] with layer sizes (512, 512, 512, 256, 256, 256, 256). Note that we use a softplus activation instead of ReLU, when multiplying the modulator to the intermediate features. The encoders (one ResNet-18 for the camera, and another for inferring the latent conditioning shape vector) take 128×128 RGBA images as input. However, when rendering the visibility mask for \mathcal{L}_M , we output 64×64 images. We set $\gamma_{R,S} = 1$, $\gamma_{R,\Pi} = 5$, and $\gamma_{R,M} = 10$. We also used small weight decays on the camera and shape predictors (10^{-6} and 10^{-3} , respectively). Training ran for 100K iterations with AdamW [102] and a batch size of 32. We used $\dim(z_s) = 512$ and a SIREN initializer of $\omega_0 = 1$.

Shape-fitting Loss. We use the same settings as the single-shape fitting experiments, with slight modifications: $\gamma_d = 5$, $\gamma_\xi = 10$, $\gamma_n = 10$, $\gamma_V = 1$, $\gamma_{E,d} = 0.1$, $\gamma_{E,\xi} = 0.25$, and $\gamma_T = 0.1$. An additional loss on the variance of ξ is also applied (to reduce the visibility entropy at each point, which leads to fuzzy renders):

$$\mathcal{L}_{V,\xi} = \gamma_{V,\xi} \xi(p, v) [1 - \xi(p, v)], \quad (62)$$

where we set $\gamma_{V,\xi} = 0.25$. This encourages less blurriness in the shape output due to uncertainty in the visibility field. For




Figure 20: Single-image 3D reconstruction visualizations on held-out data. Per inset, columns represent (i) the input RGB image, (ii) the visibility $\hat{\xi}$, (iii) the depth \hat{d} , (iv) the normals \hat{n} , (v) the sampled point cloud (PC) from the DDF, and (vi) a sample from the ground-truth PC. Quantities (ii-v) are all differentially computed directly from the CPDDF and $\hat{\Pi}$, per point or pixel (i.e., no post-processing needed). PC colours denote 3D coordinates. A high-error example is in the lower-right of each category. See also Fig. 12 and Supp. §D-D.

minibatches, we sample 1.2K (A and U) and 300 (S, B, T, and O) oriented points per shape, with each minibatch containing 32 shapes.

Camera Loss. The camera fitting loss \mathcal{L}_{Π} utilizes a camera representation only involving extrinsic position (the camera is assumed to be looking at the origin). In particular, we use the azimuth, elevation, and radius representation of the camera position. Before computing the L_2 loss, we z-score normalize each element, based on the training set statistics. We also restrict the predicted camera $\hat{\Pi}$ to be within the range of the parameters observed in the dataset.

Rendering Scale Factor. We note that an additional scale factor is needed for rendering DDFs for ShapeNet. Since ShapeNet shapes are normalized with an instance-dependent measure (the bounding box diagonal), one needs to know the scale to reproduce the output image. This is an issue as our CPDDF always outputs a shape in $[-1, 1]^3$, with training data normalized to have longest axis-aligned bounding box length equal to two. Further, it creates an ambiguity (with respect to the output image) with the camera position (radius from the shape). At train time, before rendering to compute \mathcal{L}_M , we use the ground-truth scale factor. At test time, we estimate it by sampling a point cloud and measuring the diagonal.

Data Extraction. We use the ShapeNet-v1 [77] data and splits from Pixel2Mesh [70], [71], with the renders by Choy et al. [72]. For DDF training, per data type, we sample 20K (A and U) and 10K (S, B, T, and O) training samples per shape, using the watertight form of the meshes (via [103]), decimated to 10K triangles. These samples are only used for training, not evaluation, and are in a canonical aligned pose. We set the maximum offset size for O-type data as $\epsilon_O = 0.02$; remaining parameters are the same as those used for single shapes (see §B).

Explicit Sampling Details. We also remark that our explicit sampling algorithm slightly oversamples points initially (when requesting a point set of size n_p , we actually sample $(1+\epsilon_p)n_p$ via $p \sim \mathcal{U}[\mathcal{B}]$). The final point cloud, however, is sorted by visibility (i.e., by $\xi(p, \hat{v}^*(p))$) and only the top n_p points are returned. We used $\epsilon_p = 0.1$ in all experiments. This is to prevent outputting non-visible points.

PC-SIREN Baseline Details. Recall that the PC-SIREN is our architecture-matched baseline, with identical encoders and a nearly identical decoder architecture to the DDF one. Here, the decoder is a mapping $f_b : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, which is trained to compute $f_b(p) \in S$ from $p \sim [-1, 1]^3$, but uses an identical set of SIREN hidden layers as the DDF. A set of sampled random

		DDF					PC-SIREN				P2M	3DR
		Π_g -L	Π_g -S	$\hat{\Pi}_{\nabla}$ -S	$\hat{\Pi}$ -L	$\hat{\Pi}$ -S	Π_g -L	Π_g -S	$\hat{\Pi}$ -L	$\hat{\Pi}$ -S		
Chairs	$D_C \downarrow$	0.459	0.512	0.823	0.855	0.919	0.431	0.465	0.876	0.915	0.610	1.432
	$F_\tau \uparrow$	55.47	48.40	42.28	47.51	41.08	62.25	56.36	50.56	45.57	54.38	40.22
	$F_{2\tau} \uparrow$	72.82	67.75	60.39	63.81	58.98	77.38	74.56	65.43	62.76	70.42	55.20
Planes	$D_C \downarrow$	0.210	0.239	0.673	0.793	0.836	0.215	0.227	0.829	0.844	0.477	0.895
	$F_\tau \uparrow$	80.46	76.62	63.32	63.75	60.54	81.49	80.11	63.64	62.34	71.12	41.46
	$F_{2\tau} \uparrow$	90.05	88.55	76.16	74.96	73.47	89.71	89.18	74.76	74.18	81.38	63.23
Cars	$D_C \downarrow$	0.231	0.288	0.390	0.541	0.606	0.371	0.400	0.737	0.768	0.268	0.845
	$F_\tau \uparrow$	70.91	59.93	54.16	62.69	52.47	64.57	57.82	56.01	50.04	67.86	37.80
	$F_{2\tau} \uparrow$	86.57	79.66	74.68	79.71	72.78	78.72	76.00	71.22	68.52	84.15	54.84
Avg	$D_C \downarrow$	0.300	0.346	0.629	0.730	0.787	0.339	0.364	0.814	0.842	0.452	1.057
	$F_\tau \uparrow$	68.95	61.65	53.25	57.98	51.36	69.44	64.76	56.74	52.65	64.45	39.83
	$F_{2\tau} \uparrow$	83.15	78.65	70.41	72.83	68.41	81.94	79.91	70.47	68.49	78.65	57.76

Table V: Single-image 3D reconstruction results. Rows: ShapeNet categories and performance metrics. Columns: L/S refer to sampling 5000/2466 points for evaluation (2466 being the output size of P2M), $\Pi_g/\hat{\Pi}$ denote using the true versus predicted camera for evaluation (the former case removing camera prediction error), and $\hat{\Pi}_{\nabla}$ test-time camera correction from the predicted position using gradient descent. Metrics: D_C is the Chamfer distance ($\times 1000$), F_τ is the F-score ($\times 100$) at threshold $\tau = 10^{-4}$. PC-SIREN is our matched-architecture baseline; Pixel2Mesh (P2M) [70], [71] and 3D-R2N2 (3DR) [72] are baselines using different shape modalities (numbers from [71]). Note that scenarios using Π_g (effectively evaluating shapes in canonical object coordinates) are not directly comparable to P2M or 3DR. Overall, DDF-derived PCs (1) perform similarly to directly learning to output a PC and (2) underperform P2M overall, but outperform it in terms of shape quality when camera prediction error is excluded.

points can thus be mapped into a point cloud of arbitrary size. The baseline has 25,645,574 parameters, while the DDF-based model has 25,647,367 parameters. The camera loss \mathcal{L}_{Π} is unchanged and the mask-matching loss \mathcal{L}_M is not used. The shape-fitting loss \mathcal{L}_S is replaced with a standard Chamfer loss D_C [104], computed with 1024 points per shape with a batch-size of 32. The remaining aspects of training remain the same.

B. Table of Detailed Results

We show complete results across all categories in Table V.

C. Ablation with $N_H = 1$

Recall that N_H is the number of times to “cycle” the points (projecting them towards the surface via the DDF) when sampling an explicit point cloud shape from a DDF (see §V-B). We show results with $N_H = 1$ in Table VI. In most cases, it is slightly worse than using $N_H = 3$, by 1-3 F -score units; occasionally, however, it is marginally better: on Planes- $\hat{\Pi}$, it has slightly lower D_C , though this does not translate to better F -score.

D. Additional Visualizations

Some additional visualizations are shown in Fig. 20. For highly novel inputs, we also observe that, sometimes, the

		DDF			
		Π_g -L	Π_g -S	$\hat{\Pi}$ -L	$\hat{\Pi}$ -S
Chairs	$D_C \downarrow$	0.477	0.532	0.861	0.928
	$F_\tau \uparrow$	54.37	47.26	46.81	40.35
	$F_{2\tau} \uparrow$	71.62	66.33	63.09	58.09
Planes	$D_C \downarrow$	0.201	0.231	0.748	0.799
	$F_\tau \uparrow$	80.69	76.71	63.86	60.48
	$F_{2\tau} \uparrow$	90.23	88.49	75.30	73.55
Cars	$D_C \downarrow$	0.235	0.309	0.545	0.628
	$F_\tau \uparrow$	68.21	57.36	59.84	49.98
	$F_{2\tau} \uparrow$	83.99	76.41	76.87	69.32

Table VI: Single-image 3D reconstruction results with “single-hop” point sampling ($N_H = 1$).

network does not adapt well to the shape (e.g., see the chairs example in the second row and second column). While much error is due to the incorrectly predicted camera, the DDF outputs can also be a bit blurrier, especially when it is uncertain about the shape. This can occur on thin structures, which are hard to localize (e.g., the chair legs in either row one and column one, or row three column two), or atypical inputs (e.g., row three, column two of the cars examples), where the network does not have enough examples to obtain high quality geometry. One can also see some non-uniform densities from our point cloud sampling algorithm (e.g., concentrations of points on the chair legs in column two, or on the wheels of several examples of cars).




Figure 21: Two-stage unpaired generative modelling architecture. *Upper inset:* VAE formulation mapping a point cloud P and associated normals N to latent shape vector z_s via PointNet encoder E , and decoding into depth values via conditioning a PDDF. *Lower inset:* latent shape z_s and camera Π are randomly sampled and used to render a surface normals image I_n , via the derivatives of the learned conditional PDDF (see Property II). The normals map I_n is then concatenated (\oplus) with a sampled latent texture z_T , and used to compute the final RGB image $I_T = f_T(I_n, z_T)$. Coloured circles indicate random variables sampled from a particular distribution; dashed lines indicate computation with multiple forward passes (per point or pixel).

APPENDIX E GENERATIVE MODELLING

See Fig. 21 for a diagram of our two-stage 3D-to-2D modality translation architecture. See also Fig. 22 and 23 for additional sample visualizations (as in Fig. 13).

A. Shape VAE

The first stage learns a VAE [84], [105] on 3D shape, in order to (i) obtain a latent shape variable z_s that is approximately Gaussian distributed, and (ii) learn a CPDDF using z_s that can encode a shape in a form that is easily and efficiently rendered, yet still encodes higher order shape information. We learn a PointNet encoder [86] E that maps a point cloud with normals (P, N) to a latent shape vector $z_s = E(P, N)$. The conditional PDDF (CPDDF) then acts as the decoder, computing depth values as $d(p, v|z_s)$ for a given input. To implement conditional field computations, we use the modulated SIREN approach [74]. This can be trained by the following β -VAE loss [106]:

$$\mathcal{L}_{\text{VAE}} = \mathcal{L}_S + \beta \mathcal{L}_{\text{KL}}, \quad (63)$$

where \mathcal{L}_{KL} is the standard KL divergence loss between a Gaussian prior and the approximate VAE posterior, and \mathcal{L}_S acts analogously to the reconstruction likelihood. We use ShapeNet cars [77] to fit the Shape VAE. Our goal is to show that we can rapidly train a model that can sample renders of surface normal geometry (from a unified 3D model), which is either difficult (or computationally costly) for most implicit approaches, or it uses a model directly on 2D surface normal images (missing a unified 3D shape model). These images could be useful for other downstream models, which may expect 2D input, but still propagate gradients to an underlying 3D object model.

Data. Data is extracted from 1200 randomly chosen shapes from ShapeNet-v1 cars [77], sampling 60k (A and U) and 30K

(S, B, T, and O) oriented points. We downsample shapes to 10K triangles before extraction. For minibatches, we sample 1.2K (A and U) and 300 (S, B, T, and O) oriented points per shape, with each minibatch containing 32 shapes. We also sample point clouds P (with normals N) of size 1024 to send to the PointNet at each minibatch.

Architectures. The PointNet encoder follows the standard classification architecture [86], with four Conv1D-BatchNorm-ReLU blocks (sizes: 64, 64, 128, 1024), followed by max pooling and a multilayer perceptron with Linear-BatchNorm-ReLU blocks (two hidden layers of size 512; dropout probability 0.1). No point transformers are used. The final output is of size $2\dim(z_s)$. We then compute the approximate variational posterior $q(z_s|P, N) = \mathcal{N}(z_s|\mu(P, N), \Sigma(P, N))$ with two networks for mean μ and diagonal log-variance matrix Σ , each structured as Linear-ReLU-Linear (in which all layers are of dimensionality $\dim(z_s)$; note that each posterior parameter network takes half of the vector output from the PointNet encoder as input).

For the decoder, we use eight layers (512, 512, 512, 512, 256, 256, 256, 256) (with the modulated SIREN [74]). We set $\dim(z_s) = 400$ and $\dim(z_T) = 64$.

Training. We run for 100K iterations, using $\beta = 0.05$, $\gamma_d = 5$, $\gamma_\xi = 10$, $\gamma_n = 10$, $\gamma_V = 1$, $\gamma_{E,d} = 0.1$, $\gamma_{E,\xi} = 0.1$, $\gamma_{V,\xi} = 0.1$, and $\gamma_T = 0.1$. Adam is used for optimization (learning rate 10^{-4} ; $\beta_1 = 0.9$, $\beta_2 = 0.999$).

B. Image GAN

After training a Shape VAE, the CPDDF decoder can be used to render a surface normals image I_n (see §III-C and Fig. 8). To perform generation, we first sample latent shape $z_s \sim \mathcal{N}(0, I)$ and camera Π , which includes extrinsics (position and orientation) and focal length, following with normals map rendering to get I_n . We then use a convolutional network f_T to obtain the RGB image I_T , based on the residual image-to-image translation architecture from CycleGAN [87]. This is done by sampling $z_T \sim \mathcal{N}(0, I)$ and computing $I_T = f_T(I_n, z_T)$, where one concatenates z_T to each pixel of I_n before processing. Notice that z_s , z_T , and Π are independent, while the final texture (appearance) depends directly on z_T , and indirectly on z_s and Π , through I_n .

For training, a non-saturating GAN loss [85] with a zero-centered gradient penalty for the discriminator [107], [108] is used (as in [66], [109]). To ensure that f_T preserves information regarding 3D shape and latent texture, we use two consistency losses:

$$\mathcal{L}_{C,S} = \text{MSE}(I_n, \hat{I}_n) \quad \& \quad \mathcal{L}_{C,T} = \text{MSE}(z_T, \hat{z}_T), \quad (64)$$

where MSE is the mean squared error, $\hat{I}_n = g_T(I_T)$ (with g_T having identical architecture to f_T), and $\hat{z}_T = h_T(I_T)$ (with h_T implemented as a ResNet-20 [73], [110]). The first loss, $\mathcal{L}_{C,S}$, encourages the fake RGB image I_T to retain the 3D shape information from I_n through the translation process (i.e., implicitly forcing the image resemble the input normals) while the second loss, $\mathcal{L}_{C,T}$, does the same for the latent




Figure 22: Additional example samples from the ShapeVAE and translational image GAN.



Figure 23: Example interpolations from the ShapeVAE and translational image GAN.

texture (implicitly strengthening textural consistency across viewpoints). The final loss for the GAN image generator is

$$\mathcal{L}_{\text{GAN}} = \mathcal{L}_A + \gamma_{C,S} \mathcal{L}_{C,S} + \gamma_{C,T} \mathcal{L}_{C,T}, \quad (65)$$

where \mathcal{L}_A is the adversarial loss for the generator and the last two terms enforce consistency (see also [75], [81], [82]). We fit to the dataset of ShapeNet renders from Choy et al. [72]. Note that information on the correspondence with the 3D models is not used (i.e., the images and shapes are treated independently).

Generation process. Recall that our generation process can be written $I_T = G(z_s, z_T, \Pi) = f_T(I_n(z_s, \Pi) \oplus z_T)$, where $I_n(z_s, \Pi)$ denotes the CPDDF normals render (see §III-C) and \oplus refers to concatenating z_T to every pixel of I_n . For latent sampling, $z_s, z_T \sim \mathcal{N}(0, I)$, while the camera Π is sampled from the upper hemisphere above the object, oriented toward

the origin at a fixed distance and with a fixed focal length. The image size was set to 64×64 .

Networks. The translation network f_T exactly follows the architecture from CycleGAN [87] consisting of residual blocks (two downsampling layers, then six resolution-retaining layers, followed by two upsampling layers), using the code from [111]. The normals consistency network, g_T , has identical architecture to f_T , while the texture consistency network, h_T , is a ResNet-20 [73], [110]. We utilize the convolutional discriminator implementation from Mimicry [112], based on DCGAN [113].

Training. Our image GAN is trained in the standard alternating manner, using two critic training steps for every generator step. The non-saturating loss [85] was used, along with a zero-centered gradient penalty [107], [108] (with a weight of 10 during critic training). We used the following loss weights: $\gamma_{C,S} = 1$ and $\gamma_{C,T} = 1$. For optimization, we use Adam (learning rate 10^{-4} ; $\beta_1 = 0.0$, $\beta_2 = 0.9$) for 100K iterations (with the same reduce-on-plateau scheduler as in Supp. B).

2D GAN Comparison. As mentioned in the paper, we trained a convolutional GAN with the same loss and critic architecture using the Mimicry library [112]. We evaluated both our model and the 2D GAN with Frechet Inception Distance (FID), using torch-fidelity [114] with 50K samples.

APPENDIX F DDF-BASED PATH TRACING DETAILS

A. Path Tracing with Intrinsic Appearance DDFs

We depict the path tracing process in Fig. 24. We describe the path tracing process, material models, and lighting specifiers of the Intrinsic Appearance DDF (IADDF) in §F-B, §F-C, and §F-D. We also describe pseudocode of the algorithms in Figs. 25 and 26. Additional implementation details are in §F-E and §F-G. Finally, we present additional discussion in §F-F.

B. The Intrinsic Appearance DDF (IADDF)

Recall the classical notion of intrinsic images, which separated a 2D image into depth, “orientation” (or surface nor-




Figure 24: Illustration of our simple DDF-based path tracing algorithm. To simulate the pixel value obtained in a given direction, we use recursively query the DDF (i.e., the mapping q). Upon encountering a light source, the light value is propagated back to the eye through the sequence of rays. The entire process is differentiable, including the normals computation.

mals), reflectance, and illumination [88]⁶. In this section, we show how to define a scene description covering these pieces, by combining a DDF with simple INRs for appearance. In particular, taking advantage of the light transport interpretation of DDFs, we factorize rendered image colour into fields for intrinsic appearance and lighting. This combined formulation can be thought of as an “intrinsic image field”, because it acts as a forward model for the set of all intrinsic images, whether of geometry or appearance (i.e., it can continuously render such intrinsic images, from any viewpoint). We therefore denote the combined model an *Intrinsic Appearance DDF* (IADDF), as it incorporates a model of lighting-aware intrinsic appearance along with the geometric information held by the DDF. In the context of vision, maintaining an IADDF model of the scene automatically provides colour constancy as a side-effect, simply by replacing image appearance with intrinsic appearance. The IADDF consists of the following fields:

- The geometry field, \mathcal{G} , modelled by a DDF, consisting of
 - The distance field $d : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{R}_+$.
 - The visibility field $\xi : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow [0, 1]$.
- The material appearance model, \mathcal{M} . We consider a simple approach consisting of a continuous Bidirectional Reflectance Distribution Function (BRDF) and a directional sampling function, which provides guidance on bounce directions (e.g., for the sake of computational efficiency when rendering or learning). \mathcal{M} includes⁷
 - The BRDF $f_B : \mathbb{R}^3 \times \mathbb{S}^2 \times \mathbb{S}^2 \rightarrow \mathbb{R}_+^{|\mathcal{C}|}$.
 - The importance sampler⁸ $\Psi : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{P}[\mathbb{S}^2]$, where $\mathbb{P}[\mathbb{S}^2]$ is the set of probability distributions on \mathbb{S}^2 .
- The lighting model, \mathcal{L} . To simplify our setting, we consider two fields: one for entities locally emitting light

and another for measuring the contribution of the external environment. Thus, in our case, \mathcal{L} consists of

- The emission field, which handles objects directly emitting light (e.g., glowing), $M_L : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathcal{C}$.
- The environment light, which represents light “incoming” to the scene from some faraway source, $E_L : \mathbb{S}^2 \rightarrow \mathcal{C}$.

In our simplified scenario, the IADDF is simply the tuple $(\mathcal{G}, \mathcal{M}, \mathcal{L}) = (d, \xi, f_B, \Psi, M_L, E_L)$.

C. IADDF Path Tracing

Given an IADDF, we have disentangled models for the scene geometry, lighting, and intrinsic appearance. However, unlike other neural or differentiable rendering methods, the presence of an illumination model now necessitates some form of simulated light transport. In particular, we need to solve the classical rendering equation [115], [116], which can be formalized in the recursive integral

$$\mathcal{L}(q, \omega_o) = \mathcal{L}_M(q, \omega_o) + \int_{\Omega[n(q)]} f_B(q, \omega_i, \omega_o) \mathcal{L}_i(q, \omega_i) d\omega_i^\perp, \quad (66)$$

where q is a surface point, $n(q)$ is the surface normal at q , $\Omega[n(q)]$ is the hemi-sphere about n , $\mathcal{L}(q, \omega_o)$ is the outgoing radiance from q in direction ω_o , $\mathcal{L}_M(q, \omega_o)$ is the emitted radiance (from q in direction ω_o), $f_B(q, \omega_i, \omega_o)$ is the BRDF (which computes how irradiance from ω_i is converted to radiance along ω_o), $\mathcal{L}_i(q, \omega_i)$ is the incoming irradiance to q from direction ω_i , and $d\omega_i^\perp$ is the projected solid angle measure (where $d\omega_i^\perp = |\omega_i \cdot n(q)| d\omega_i$ under Lambert’s cosine law, foreshortening the area element from oblique viewpoints) [117], [118]. In our simple setting, the recursion and connection to DDFs becomes more clear by rewriting the incoming radiance to

$$\mathcal{L}_i(q, \omega_i) = \begin{cases} \mathcal{L}(q - d(q, -\omega_i)\omega_i, -\omega_i) & \text{if } \xi(q, -\omega_i) > 0.5 \\ \mathcal{L}_E(q, -\omega_i) & \text{otherwise} \end{cases}, \quad (67)$$

where (d, ξ) are the DDF and $\mathcal{L}_E(q, \omega)$ is the light from the environment falling on q from direction ω . (e.g., [119]). In other words, the incoming radiance at a point is either (i) the outgoing radiance reflected off of another surface (i.e., inter-reflections) or (ii) the external light from the environment (when another surface is not visible).⁹

Clearly, for a given pixel with an eye ray through it, corresponding to an oriented point (p, v) , we want to obtain the light coming back to the eye: $\mathcal{L}(q(p, v), -v)$. We use a straightforward form of solution, detailed in Fig. 25: an eye ray tracer or backward path tracer¹⁰, which is a form

⁶Given that we consider the material reflectance of an entity to be its intrinsic appearance, we could also reasonably call the illumination-derived image component, coming from the external environment, its extrinsic appearance.

⁷We use \mathcal{C} to denote the radiance emitted from a position (commonly relegated to, e.g., a simple RGB colour value).

⁸We remark that Ψ depends on the local surface normal n through the spatial position q at which it is evaluated. (We make this indirect dependence explicit in Figs. 25 and 26).

⁹We remark that there is an abuse of notation and terminology. In particular, \mathcal{L} and \mathcal{L}_M should be radiant *densities*. Under the pinhole camera assumption, for a given direction and position on a pixel sensor, the irradiance is proportional to radiance from a single direction. One then integrates over the pixel area to get total irradiance (i.e., flux or power), which is integrated over time to measure radiant energy at that pixel (for a given exposure).

¹⁰We mean “backward” here in the sense of rays starting from the camera and bouncing to a light source (i.e., the reverse of what physical light is thought to do) [120], but note that this nomenclature is not uniformly followed (e.g., [121]).

of MC rendering. Benefits of such MC path tracing methods include excellent generality (in terms of the scene descriptions and lighting effects they can handle), relative simplicity of implementation, and lack of statistical bias (i.e., zero error in expectation), while downsides include the slow convergence (depending on the algorithm and scene) and image noise [122].

We briefly describe the stochastic path tracing algorithm (see Figs. 25 and 26 for details). For a given pixel and camera (eye), we cast a ray into the scene, finding the scene point q and normal n via the DDF. Based on Ψ , we sample new outgoing directions from the hemisphere about q , and use the DDF to find the next surface in that direction. These “bounces” are the segments of a light path, being built backwards from the eye to the light source. This continues until the ray misses any local geometry (i.e., $\xi = 0$), returning a value from E_L in that direction. The final sampled pixel irradiance corresponds to this value (attenuated by the material f_B through the bounces), plus any emissions (from M_L) encountered along the way.

D. Basic Intrinsic Appearance Models

In our simplified setting, we merely need to specify the BRDF f_B and bounce sampler Ψ (see §F-B). We consider three basic cases, but note that any differentiable field (ideally satisfying certain priors, particularly those of physical realism) f_B and sampling procedure (with accompanying density function) Ψ could be used, which would be necessary for modelling arbitrarily complex observed materials. Note that we treat $\mathcal{C} = \mathbb{R}^3$; instead of including a continuous dependence on wavelength, which would be more realistic, we instead output three discrete colour values corresponding to the RGB channels in the end.

The first case is Lambertian (or diffuse) material:

$$f_{B,\text{Lam}}(q, \omega_i, \omega_o) = \frac{\rho_a}{\pi} \quad (68)$$

$$\Psi_{\text{Lam}}(q, \omega_o | n) = \mathbb{U}[\Omega[n(q)]], \quad (69)$$

where \mathbb{U} is the uniform distribution and $\rho_a \in [0, 1]^3$ is the per-channel albedo. We write $\Psi(q, \omega_o | n)$ to make the dependence on $n(q)$ more clear, but note¹² that n depends on q , through d .

The second case is a mirror material (e.g., see [118], [123]):

$$f_{B,\text{Mir}}(q, \omega_i, \omega_o) = \frac{\delta(\omega_r(\omega_i, n) - \omega_o) \rho_m}{|\omega_r(\omega_i, n) \cdot n|} \quad (70)$$

$$\Psi_{\text{Mir}}(q, \omega_o | n) = \mathbb{P}_\delta[\omega_r(-\omega_o, n)], \quad (71)$$

where δ is the Dirac Delta function (distribution) on \mathbb{S}^2 , \mathbb{P}_δ is the Dirac probability measure, ρ_m is the albedo (Fresnel reflectivity), and $\omega_r(u, v)$ is the outgoing reflection direction of an incoming direction u about n . Recall that, when we sample $v \sim \Psi_{\text{Mir}}(q, \omega_o | n)$ (see Fig. 25), we are interested in sampling directions that will contribute to outgoing light back

¹¹While p may initially be the camera position (i.e., C_p in Π), recall from §III-C that rendering proceeds by first projecting p along the camera ray to the domain of the PDDF, denoted \mathcal{B} , and then utilizing that boundary point as p instead.

¹²Additionally, since the DDF is not view consistent, n is not necessarily either.

Algorithm 1: PathTracePixelDDF(u, Π)

Input: Pixel u and Camera Π

Models: DDF (d, ξ), Appearance (f_B, Ψ), and Illumination (M_L, E_L)

Hyper-Parameters: Maximum bounces n , Number of MC Samples m

Output: The pixel colour c at u

```

1  $p, v \leftarrow \text{GetOrientedPoint}(u, \Pi)$  // Camera ray
   for  $u$ 
2 if  $\xi(p, v) < 0.5$  then // Check for ray
   intersection with scene
3   return  $E_L(v)$  // Return environment
   light if no intersection
4  $q \leftarrow p + d(p, v)v$  // Scene point for the
   camera ray
5  $n \leftarrow \text{Unit}(\nabla_p d(p, v))$  // Scene point
   surface normals
6  $\omega_o \leftarrow -v$  // Direction back to the eye
7  $c_E \leftarrow M_L(q, \omega_o)$  // Colour emitted from  $q$ 
   towards  $\omega_o$ 
8 for  $k \leftarrow 1$  to  $m$  do // Iterate over MC
   samples
9    $v_n^{(k)} \sim \Psi(q, \omega_o | n)$  // Sample light
   direction from  $q$  about  $n$ 
10   $\omega_i^{(k)} \leftarrow -v_n^{(k)}$  // New incoming
   direction, into  $q$ 
11   $B \leftarrow f_B(q, \omega_i^{(k)}, \omega_o)$  // Query the BRDF
   field
12   $c_k \leftarrow B \odot \text{Trace}(q, v_n^{(k)}, \ell = 1)$  // Obtain
   colour by “bouncing” the ray
13   $w_k \leftarrow \Psi(v_n^{(k)} | q, \omega_o)^{-1} |n \cdot \omega_i^{(k)}|$  // Weight
   in MC averaging
14  $c = c_E + \frac{1}{m} \sum_{i=1}^m w_i c_i$  // Final pixel colour
15 return  $c$ 

```

Figure 25: Algorithm used for DDF rendering via backward path tracing. We devise a simple approach based on stochastic backward ray tracing, where we start rays from the eye and follow them into the scene until a light source is encountered.¹¹ For efficiency, we only use multiple Monte Carlo (MC) samples for the first bounce, since the contribution afterwards often falls off quickly. Notice that (i) we can compute the colour of each pixel *independently*, using a series of differentiable field (i.e., DDF, BRDF field, and lighting field) calls per pixel ray (in a parallelizable manner similar to NERFs), (ii) substituting the DDF as the geometric representation does not preclude utilizing the many advances in path tracing methods, and (iii) this framework naturally enables handling reflections, shadows, and glowing objects. See Fig. 26 for the definition of the recursive Trace function and Fig. 24 for a visual depiction.

Algorithm 2: Trace(p, v, ℓ)

Input: Initial Scene Point p , Direction v , Bounce Level ℓ
Models: DDF (d, ξ), Appearance (f_B, Ψ), and Illumination (M_L, E_L)
Hyper-Parameters: Maximum bounces n
Output: Colour c at p from direction v

```

1 if  $\xi(p, v) < 0.5$  then           // Check for ray
  intersection with scene
2   return  $E_L(v)$                  // Return environment
  light if no intersection
3 if  $\ell \geq n$  then                 // Check current bounce
  level
4   return  $\vec{0}$  // Stop bouncing after level
   $n$ 
5  $q \leftarrow p + d(p, v)v$  // Scene point along ray
  from  $(p, v)$ 
6  $n \leftarrow \text{Unit}(\nabla_p d(p, v))$  // Scene point
  surface normals
7  $\omega_o \leftarrow -v$  // Direction back to  $p$  from  $q$ 
8  $c_E \leftarrow M_L(q, \omega_o)$  // Colour emitted from  $q$ 
  towards  $\omega_o$ 
9  $v_n \sim \Psi(q, \omega_o | n)$  // Sample light
  direction from  $q$  about  $n$ 
10  $\omega_i \leftarrow -v_n$  // New incoming direction
  (i.e., into  $q$ )
11  $B \leftarrow f_B(q, \omega_i, \omega_o)$  // Query the BRDF field
12  $c \leftarrow B \odot \text{Trace}(q, v_n, \ell + 1)$  // Recursive
  colour from  $q$  (via  $\omega_i$ ) to  $p$ 
13 return  $c_E + \Psi(v_n | q, \omega_o)^{-1} |n \cdot \omega_i| c$ 
  // Single-sample MC irradiance at  $p$ 

```

Figure 26: DDF-based recursive path tracing process, used by the algorithm in Fig. 25. Each recursive call corresponds to a single bounce of the light path (and thus one call to the DDF). Notice that only a single forward and backward pass is needed from the DDF per bounce. See also Fig. 24 for a visual depiction.

to the eye (ω_o); thus, for a mirror, we should of course sample the reflection of ω_o . Notice that the BRDF contains a cosine term that cancels out the geometric attenuation from Lambert’s cosine law; this is needed to ensure energy conservation in the rendering integral (i.e., we do not need to correct for the differential contribution of directions based on foreshortening of the area element, because no scattering occurs about the reflection direction). We remark that modelling reflections is an important and difficult topic in CV; modern Analysis-by-Synthesis methods based on NERFs cannot easily model them without significant modification (e.g., [124], [125]), despite applications in terms of editing (including, e.g., removing reflections from glass or making an object reflective). Moving to a path tracing framework, such as the IADDF or an extension that can model transmission, immediately enables disentangling or generating reflections.

The third case is a simple glossy material, based on the classical Phong model [126]. This formulation, while non-physical, simply mixes a Lambertian and specular reflection term, as has been previously applied in computer graphics (e.g., see [118], [127]):

$$f_{B,G}(q, \omega_i, \omega_o) = \eta_L \frac{\rho_a}{\pi} + \frac{(1 - \eta_L) \rho_m}{N(\alpha, \omega_i, n)} (\omega_o \cdot \omega_r(\omega_i, n))^\alpha \quad (72)$$

$$\Psi_G(q, \omega_o) = \eta_L \mathbb{U}[\Omega[n(q)]] + (1 - \eta_L) P_{\text{Proj}}(q, \omega_o), \quad (73)$$

where $\eta_L \in (0, 1)$ balances diffuseness and specularity, $N(\alpha, \omega_i, n) = \pi^{\frac{1}{1+\alpha}} |\omega_i \cdot n|^{\frac{\alpha}{1+\alpha}}$, $P_{\text{Proj}}(q, \omega_o) = \mathcal{N}_{\text{Proj}}(\omega_r(-\omega_o, n(q)), \Sigma_p(\alpha))$, $\alpha \in [0, \infty)$ is the specular exponent (controlling “glossiness”), and $\mathcal{N}_{\text{Proj}}$ is a distribution centered about the reflection direction, based on a spherical projection of a Gaussian. In the second line, we use the sum to denote a convex combination (mixture) of the two measures. We relegate the details of the Gaussian model to §F-G2, but remark that its sampling procedure can be back-propagated through¹³, via the reparameterization trick. The second term effectively “interpolates” between a mirror ($\alpha \rightarrow \infty$) and a Lambertian material ($\alpha \rightarrow 0$). Notice that, in order to obtain visually reasonable specular highlights, the DDF has to provide fairly accurate surface normals (i.e., the derivatives of the mapping need to be correct), just as with the mirror case.

E. Path Tracing Post-Processing

We remark that our algorithm records the light flowing through a scene at a pixel, effectively measuring irradiance at that position. It is thus akin to a raw image (in “linear RGB”), which is often not perceptually agreeable. In real cameras, of course, the standard camera image processing pipeline transforms such images to standard RGB, through a complex series of specialized algorithms (e.g., [128], [129]). In our case, we apply a very simplified differentiable post-processing for visualization, applying denoising (Gaussian filtering), min-max normalization, and then a gamma transform via $\text{rgb} = \text{rgb}_{\text{linear}}^{1/2.2}$. In order to fit to real images, one would likely need a learnable mapping to mimic the camera pipeline (to produce the images to which we would be fitting).

F. Additional Discussion

In conclusion, we have devised a simple approach to integrating DDFs into a path tracing framework, via the IADDF. The interpretation of recursive DDF applications as covering the space of light paths lends itself well to this combination. The *internal structure modelling* capability of PDDFs, which other directed fields often ignore, is fundamental to our light transport algorithm; without it, we could not model the intra-scene paths between surfaces.

Similar to NERFs, one can differentially render each pixel and MC sample independently (i.e., in parallel), with a fixed

¹³Notice that such samples depend on the DDF through n and material parameters, like α , can be back-propagated to via pathwise derivatives.

number of queries per ray. More exactly, we incur $\mathcal{O}(nm)$ DDF calls per pixel, for n bounces and m MC samples. (see Fig. 25). While this is distinctly more costly than the *single* call used for *geometric* neural rendering of DDFs (and then using a simple INR, such a texture field [130]), it enables a principled connection between DDFs and appearance modelling. Moreover, rendering via light transport in this manner enables (a) decomposition of illumination, reflectance, and shape; (b) automatic handling of many natural phenomena, including surface inter-reflections, glowing objects, multiple light sources, complex material reflectances, and mirrors; and (c) differentiability through the rendering process (i.e., potentially enabling gradient-based learnability). When editing IADDF scene descriptions, one can easily relight scenes or alter intrinsic material appearance of objects (even changing surfaces to mirrors or enabling them to emit light). The compositionality of DDFs suggests inserting new objects should be relatively straightforward as well. Non-local changes in appearance due to such edits, including shadows or inter-reflections, will be naturally accounted for by the path tracing algorithm. This is in contrast to, e.g., standard NeRF editing procedures, for which such non-local physical effects are not easily modelled or corrected.

The most important challenge is in deriving IADDFs from image data. The Analysis-by-Synthesis multiview image setup, as in NeRFs, is likely to be amenable to inverse rendering with IADDFs, but powerful priors will be needed to constrain the inherent ambiguities between shape (shading), illumination, and reflectance, as well as neural architectures able to learn higher dimensional (here, 5D) scene descriptions and learning methods for view consistency. For instance, priors on smoothness and material-type parsimony have already proven effective in decomposing lighting from reflectance (e.g., [44], [131]); constraining the space of BRDFs will also necessitate learning which materials are likely to appear in the world. Similarly, more sophisticated learned models of lighting and intrinsic material appearance will be needed, compared to the simple ones used here. In addition, the use of a DDF does not preclude most techniques for improving MC path tracing, (e.g., bidirectional tracing [120], importance sampling [132], [133], radiosity techniques [119]). Research on such methods has been a major topic within computer graphics for decades (and increasingly enmeshed with differentiable and/or learning-based techniques); combining such advances with the IADDF should hopefully lead to a disentangled neural scene representation that (i) can be efficiently rendered, (ii) interoperates with other ML algorithms, (iii) models a myriad of physical phenomena, (iv) can be learned in a weakly supervised Analysis-by-Synthesis manner, and (v) is easily edited in terms of shape, material, and lighting.

G. Technical Details

1) *Hyper-Parameters and Implementation Details*: The fitting process was the same as in §V-A, except we use double the number of extracted samples (as well as 500K W-type samples; see §F-A). In all cases, we used number of bounces

$n = 3$, MC samples¹⁴ $m = 256$, and post-processing Gaussian blur $\sigma = 1$ (images are 512^2), except for the glossy material with $\alpha = 3$, which used $m = 512$, and $\alpha = 1$, which used $m = 1024$ and $\sigma = 3$. Fig. 14, row one, used $\alpha = 1$ and $\eta_L = 0.75$.

2) *Projected Gaussian Importance Sampler*: Recall the simple glossy material defined in Eq. 72, specifically the importance sampler Ψ :

$$\Psi_G(q, \omega_o | n) = \eta_L \mathbb{U}[\Omega[n(q)]] + (1 - \eta_L) \mathcal{N}_{\text{Proj}}(\omega_r(-\omega_o, n), \Sigma_p(\alpha)). \quad (74)$$

Two operations need to be implemented: density evaluation and sampling.

We use a basic simplification to compute the inverse density, by assuming we know the origin of a sample:

$$\Psi_G(v | q, \omega_o, n)^{-1} = \frac{\mathbb{1}[v \sim \mathbb{U}[\Omega[n(q)]]]}{\eta_L} P_{\mathbb{U}}^{-1}(v) + \frac{\mathbb{1}[v \sim \mathcal{N}_{\text{Proj}}]}{1 - \eta_L} P_{\mathcal{N}_{\text{Proj}}}^{-1}(v | \omega_r, \Sigma_p(\alpha)), \quad (75)$$

where $P_{\mathbb{U}}$ and $P_{\mathcal{N}_{\text{Proj}}}$ are density functions, $\omega_r = \omega_r(-\omega_o, n)$, and $\mathbb{1}[\cdot]$ is the indicator function. We discuss the actual density used by $P_{\mathcal{N}_{\text{Proj}}}$ later.

For obtaining a single sample from this combined material (i.e., $v \sim \Psi_G(q, \omega_o | n)$), we first choose which mixture mode (specular or Lambertian) to pick: $b \sim \mathcal{B}(\eta_L)$, where $\mathcal{B}(\eta_L)$ is the Bernoulli distribution. If $b = 0$, we sample uniformly, from $\mathbb{U}[\Omega[n(q)]]$. Otherwise, if $b = 1$, we take a Gaussian sample $\tilde{\zeta} \sim \mathcal{N}(\omega_r, \Sigma_p(\alpha))$ and then normalize it via $v = \tilde{\zeta}/\|\tilde{\zeta}\|_2$.

To compute the density $P_{\mathcal{N}_{\text{Proj}}}$, we use an approximation. We assume that we know $\zeta = \tilde{\zeta} - \omega_r$, and then compute $P_{\mathcal{N}_{\text{Proj}}}(v | \omega_r, \Sigma_p(\alpha)) \approx C(P_{\mathcal{N}}(\zeta | 0, \Sigma_p(\alpha)))$. We use a simple clamping function, C , to keep the density numerically stable.

The key parameter that controls the glossiness of the material is the Phong specular exponent α , which is related to the “spread” of the bounce directions about the reflection direction. This affects the density through the noise covariance $\Sigma_p(\alpha) = f_\alpha \exp(-\kappa_0 \alpha) I_3$, where $f_\alpha = 3.5 \times 10^{-\alpha}$ and $\kappa_0 = 1$. We remark that this ad hoc formulation means that one should not interpret α the same way as in other Phong models.

In terms of differentiability, the reparameterization trick can be used to optimize α , since we are using a diagonal Gaussian (as in most VAEs). However, notice that back-propagation to η_L through the sampling process is non-trivial; since it controls which mixture component (b) to sample, it describes a fundamentally discrete random indicator function (i.e., Bernoulli variable). Techniques for handling discrete random variables in stochastic computation graphics (e.g., continuous relaxations [134], [135] or score functions [136]) would be necessary for obtaining derivatives.

¹⁴Note: the mirror material only requires one sample (though, for instance, non-mirror materials encountered *after* a mirror could lead to requiring more).




Figure 27: Examples of path-traced DDFs, as in Fig. 14. Top row: Lambertian material. Second row: glossy material with $\alpha = 3$ and $\eta_L = 0.25$. Notice the “sheen” added by the mirror term, which injects the glossiness. Third row: no external lighting is used. Instead, a simple “glow” (constant $M_L > 0$) is emitted by the geometry; hence the areas of the shape that would normally be shadowed due to self-occlusion are actually highlighted.

APPENDIX G

THEORETICAL DETAILS: VIEW CONSISTENCY OF DDFs

This section is an extended and more complete version of the theoretical analysis in §VII of the main paper.

A. Preliminaries

In this section, our intention is to explore the conditions under which a DDF, or one of its constituent fields, represents a shape. In particular, we show an equivalence between certain local properties of such fields and the existence of a 3D shape to which the field has been perfectly “fitted”. We begin by defining the basic setting, including the domain in which our fields operate, the oriented points that characterize our ray-based geometric representation, and the precise definition of a shape.

Definition G.1: Domain \mathcal{B} , \mathcal{B}_ε , Γ , and Γ_ε

We consider fields with a geometrically simple domain, $\mathcal{B} \subset \mathbb{R}^3$. In particular, we assume that \mathcal{B} is convex, compact, and has a smooth boundary, $\partial\mathcal{B}$, with outward-facing surface normals. Finally, for $\varepsilon > 0$, we denote \mathcal{B}_ε as the ε -domain of \mathcal{B} , defined via

$$\mathcal{B}_\varepsilon = \left\{ p \in \mathcal{B} \mid \min_{b \in \partial\mathcal{B}} \|p - b\| \geq \varepsilon \right\}, \quad (76)$$

which will be useful for ensuring certain boundary conditions on our fields of interest.

Often, we will be interested in the 5D Cartesian product space, $\Gamma = \mathcal{B} \times \mathbb{S}^2$. Similarly, we can define $\Gamma_\varepsilon = \mathcal{B}_\varepsilon \times \mathbb{S}^2$.

The role of \mathcal{B}_ε is to enable us to define a notion of “away from the boundary” (i.e., $\mathcal{B} \setminus \mathcal{B}_\varepsilon$ is a thin ε -width shell around the edges of the domain). Notice that we restrict our attention to a small part of \mathbb{R}^3 , where our shape of interest will reside,

but the domain of our fields will be the 5D space Γ , comprising oriented points or, equivalently, rays.

Definition G.2: Oriented Points and Rays

We generally denote an *oriented point* via $\tau = (p, v) \in \Gamma$ (or in Γ_ε when specified). Any oriented point induces (or is equivalent to) a 3D ray: $r_\tau(t) = p + tv$, $t \geq 0$. In a slight abuse of notation, we may refer to the oriented point and ray forms of such 5D elements interchangeably.

Fundamentally, our interest is in representing 3D shapes, which are defined simply as follows.

Definition G.3: Shapes

We define a *shape* to be a compact set $S \subset \mathcal{B}$. Often, we will be interested in shapes $S \subset \mathcal{B}_\varepsilon$, for $\varepsilon > 0$.

Note that using \mathcal{B}_ε does not strongly constrain the shape: for example, one can simply use the bounding sphere of the given point set, and then inflate it enough to satisfy the ε condition.

Next, for the sake of notational simplicity, we define an “along-ray” form of any function (or field) that operates on oriented points (i.e., in 5D), by considering its restriction to a given ray.

Definition G.4: Along-Ray Functions

We define the “along-ray” form of a function $g : \Gamma \rightarrow \mathcal{X}$, which maps into a set \mathcal{X} , to be:

$$f_g(s \mid \tau) := g(p + sv, v) = g(r_\tau(s), v), \quad (77)$$

where $s \geq 0$, $\tau = (p, v) \in \Gamma$, and r_τ is the ray form of τ .

Finally, we introduce notation for the intersections of rays and point sets.

Definition G.5: Intersecting Rays and Point Sets

Consider a point set $S \subseteq \mathcal{B}$ and an oriented point $\tau \in \Gamma$. *Intersected Points.* Let $S_\tau \subseteq S$ denote the set of points in S intersected by r_τ (i.e., $q \in S_\tau$ iff $\exists t \geq 0$ such that $r_\tau(t) = q$). For notational clarity, we may write $[S]_\tau$ as well.

Intersecting Rays. Let $\mathcal{I}_S \subseteq \Gamma$ be the set of rays that intersect at least one member of S (i.e., $\tau \in \mathcal{I}_S$ if $\exists q \in S$ such that $r_\tau(s) = q$ for some $s \geq 0$). For $q \in \mathcal{B}$, we denote $\mathcal{I}_q := \mathcal{I}_{\{q\}}$.

B. View Consistency for Simple DDFs

Our first consideration will be a simplified form of DDF, which does not utilize a visibility field.

Definition G.6: NNBC Field

An *NNBC field* on Γ is a non-negative, bounded scalar field that is piece-wise continuously differentiable along rays, written $d : \Gamma \rightarrow \mathbb{R}_{\geq 0}$. The latter condition means the along-ray form is piecewise C^1 (i.e., it satisfies $f_d(s \mid p, v) = d(p + sv, v) \in C^1$).

In general, such fields need not have an obvious connection to any shapes (in the sense of the definition above). In this section, we treat such fields as *putative* representations for shapes – the goal is to understand the exact conditions under which such fields are “equivalent to” (or “represent”) some shape.

Constraint Definition G.1: Positive Boundary Condition for Simple DDFs (BC_d)

An NNBC field d satisfies the *Positive Boundary Condition* BC_d iff

$$\inf_{v \in \mathbb{S}^2} d(p, v) > 0 \quad \forall p \in \mathcal{B} \setminus \mathcal{B}_\varepsilon. \quad (78)$$

Recall that $\mathcal{B} \setminus \mathcal{B}_\varepsilon$ is the “outer shell” of the bounding domain. In other words, BC_d demands that d cannot have any zeroes close to $\partial\mathcal{B}$.

Constraint Definition G.2: Directed Eikonal Condition for Simple DDFs (DE_d)

An NNBC field d satisfies the *Directed Eikonal Constraint* DE_d if $\partial_s f_d(s \mid \tau) = -1$, except at along-ray zeroes of d (i.e., at any (q, v) such that $\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0$).

The second constraint is on the derivative of d : it says that, along any ray, d must decrease linearly, at unit rate, unless $d = 0$ (a value at which it may potentially stay).

Constraint Definition G.3: Isotropic Opaqueness Condition for Simple DDFs (IO_d)

An NNBC field d satisfies the *Isotropic Opaqueness Constraint* at a point $q \in \mathcal{B}$ if

$$\left(\lim_{\epsilon \downarrow 0} \inf_{\substack{s \in (0, \epsilon) \\ v \in \mathbb{S}^2}} d(q - sv, v) = 0 \right) \implies \left(\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0 \quad \forall v \in \mathbb{S}^2 \right). \quad (79)$$

We say d is isotropically opaque, denoted IO_d, if it satisfies the Isotropic Opaqueness Constraint everywhere in \mathcal{B} .

Notice that this is stronger than (i.e., it implies) the following constraint, which refers only to local properties in the neighbourhood of some oriented point (q, v) :

$$\left(\exists v \in \mathbb{S}^2 \text{ s.t. } \lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0 \right) \implies \left(\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0 \quad \forall v \in \mathbb{S}^2 \right). \quad (80)$$

Namely, if the field approaches zero at some position q from any direction v , then it must approach zero at that point from all directions. Note that, for the left-hand side of Eq. 79, a single particular direction v along which d converges to zero need not exist (see Supp. §H-D for an example). The weaker, second form of IO_d means that convergence to a zero from one direction must be accompanied by a zero from all directions.

Together, these field constraints enable us to define a *simple DDF*, which we will show is capable of representing shapes with theoretical consistency guarantees.

Definition G.7: Simple DDF

We define a simple DDF to be an NNBC field such that IO_d, BC_d, and DE_d are satisfied.

The remainder of this section will be devoted to showing that the three field constraints, which define a simple DDF, are sufficient to ensure a form of “view consistency”.

Before that, however, we define some notational constructs around Simple DDFs. First, we define a simple function based on d , which outputs the *position* that our depth field takes us to, along a given ray.

Definition G.8: Positional Field

Any simple DDF has an associated positional field, or q -mapping, defined by: $q(p, v) := p + d(p, v)v$.

Next, because DDFs can be discontinuous at zeroes, for convenience of notation, whenever a simple DDF d approaches zero along a ray, we denote this as a zero of the field (regardless of the nature of any discontinuities at such points).

Definition G.9: Left-Continuity Convention

For any $q \in \mathcal{B}_\varepsilon$, we write:

$$d(q, v) = 0 \iff \lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0. \quad (81)$$

In words, left continuity for d defines a slightly looser notion of a zero along rays (here, we say “left” with respect to the direction v). We also apply this convention to q , so that if $d(p, v) = 0$, we write $q(p, v) = p$. Finally, we define a special set of points that is associated to any simple DDF.

Definition G.10: Positional Zero Set for Simple DDFs

We denote the set of *positional zeroes* of a Simple DDF as

$$Q_d = \left\{ p \in \mathcal{B} \mid \lim_{\epsilon \downarrow 0} \inf_{\substack{s \in (0, \epsilon) \\ v \in \mathbb{S}^2}} d(p - sv, v) = 0 \right\}. \quad (82)$$

Note that the condition for being in Q_d is slightly weaker than that of being a directional zero, in the sense of Eq. 81. I.e., there may be $q \in Q_d$ such that there is no v for which $d(q, v) = 0$. However, any q that satisfies $d(q, v) = 0$ for some v is necessarily in Q_d .

Our first lemma concerns these positional zeroes, namely that it is a closed set bounded in \mathcal{B}_ε .

Lemma G.1: Properties of Simple DDF Positional Zeroes

Let d be a simple DDF. Then $Q_d \subseteq \mathcal{B}_\varepsilon$ and Q_d is a closed set.

Proof. See Supp. §H-A2. ■

Notice this shows that Q_d is a shape (i.e., compact point set). We will use these properties later, connecting this point set to the shape represented by the field (as we will later show). The first step along this line is to define what it means for a simple DDF to “exactly” represent a shape.

Definition G.11: Shape-Induced Simple DDFs

Let S be a shape such that $S \subseteq \mathcal{B}_\varepsilon$. We define a set of “induced fields” associated to S , which each act as a form of *representation* for the shape. Recall that S_τ denotes the points in S intersected by r_τ . Then S induces a field $d : \Gamma \rightarrow \mathbb{R}_{\geq 0}$ iff

- (i) for any $\tau = (p, v) \in \Gamma$ such that $S_\tau \neq \emptyset$, $d(\tau) = \min_{q \in S_\tau} \|q - p\|$ and
- (ii) $\exists \varepsilon_E > 0$ such that for any $\tau \in \Gamma$ with $S_\tau = \emptyset$, we have that $d(\tau) > \varepsilon_E$, $q(\tau) \notin \mathcal{B}$, and f_d is C^1 along r_τ with $\partial_s f_d(s \mid \tau) = -1$.

Intuitively, a point set S induces a field d if (i) the q -mapping always takes p to the closest point in S along v , unless (ii) no intersection exists, in which case q must predict a point *outside* the valid domain. Notice, of course, that a large class of d can therefore satisfy this for a given S , differing

only in where outside \mathcal{B} they output.¹⁵ We next show that any induced field is a Simple DDF.

Theorem G.1: Shape-Induced Fields Are Simple DDFs

Consider a field d that is induced by some shape S . Then d is a simple DDF. Further, the positional zeroes of d in \mathcal{B} satisfy $S = Q_d$.

Proof. See Supp. §H-A3. ■

Recall that the definition of a shape-induced field need not be unique (as two induced fields may differ on rays that do not intersect S) and therefore allows a given S to induce a *set* of d fields, which differ only in terms of their out-of-domain behaviour. Hence, we need a measure of “equivalence” between Simple DDFs. A given shape S can then induce an equivalence class of Simple DDFs. To define this notion of equivalence, we first define the set of empty rays.

Definition G.12: Empty Rays

Let d be a Simple DDF. Then we let $\mathcal{E}[d] \subseteq \Gamma$ be the set of *empty rays*, defined by $\mathcal{E}[d] = \{\tau \in \Gamma \mid q(\tau) \notin \mathcal{B}\}$, where q is the q -mapping of d . Notice that an empty ray $\tau \in \mathcal{E}[d]$ must have $f_d(s \mid \tau) > 0$ as long as $r_\tau(s) \in \mathcal{B}$. (I.e., we cannot have a zero of d on an empty ray).

For a given d , the empty rays are those that d maps outside of the domain. To check equality, we can ignore the depth *values* of empty rays, as long as *which* rays are empty is the same.

Definition G.13: Equivalence of Simple DDFs

Consider two Simple DDFs, d_1 and d_2 . Then d_1 and d_2 are *equivalent* iff (i) both fields have the same set of empty rays, written \mathcal{E} (i.e., $\mathcal{E} := \mathcal{E}[d_1] = \mathcal{E}[d_2]$) and (ii) they agree on non-empty rays (i.e., $d_1(\tau) = d_2(\tau) \forall \tau \in \Gamma \setminus \mathcal{E}$).

Later we will see a slightly different form of equivalence between *full* DDFs (i.e., including a visibility field), similar to that of Simple DDFs. First, however, we formally define what it means for a simple DDF to be a shape representation.

Definition G.14: Shape Representations as View-Consistent Simple DDFs

A Simple DDF d is View Consistent (VC) iff there exists a shape S such that S induces d .

We can therefore call any such d a *shape representation* for the point set S .

Recall that an NNBC field (see Def. G.6) maps rays (in 5D) to positive scalars, with very little constraints. Most

¹⁵ An alternative form of shape-induction (see Supp. §H-A1) avoids defining notions of equivalence by including the boundary $\partial\mathcal{B}$ with the shape S . This enables a simpler form of equality, but the approach above is more analogous to what we will do in the case of the full DDF.

such fields will not be view-consistent, in the sense that they will not represent a shape (point set), in a self-consistent manner (e.g., violating the basic inequality¹⁶ from §III-B). The definition above formalizes what it means for a field to be view-consistent.

The next theorem is the main result of this section. It shows that the local¹⁷ field constraints (Constraints G.1, G.2, and G.3) are sufficient to ensure a Simple DDF represents a shape (i.e., in a view-consistent manner).

Theorem G.2: Every Simple DDF is a Shape Representation

Any simple DDF d is view consistent (i.e., there exists a shape $S \subseteq \mathcal{B}_\varepsilon$ such that d is equivalent to any field induced by S). Further, d is a shape representation for the set formed by its positional zeroes, Q_d (i.e., Q_d induces d).

Proof. See Supp. §H-A4. ■

The two theorems above (G.1 and G.2) show that being a simple DDF (i.e., non-negativity, no zeroes close to the boundary, isotropic zeroes, and unit along-ray derivatives) is equivalent to view consistency (i.e., to being induced by some shape). In other words, the constraints on the implicit field of being a simple DDF is synonymous with being a shape representation. Later, these results will be useful when showing that *full* DDFs can be constrained in a similar manner, in order to guarantee view consistency.

C. View Consistency for Visibility Fields

We next look at the second constituent field in a DDF: the visibility field, which demarcates whether or not a shape is present along a given ray. Similar to the previous section, our interest is in determining sufficient constraints on a given visibility field such that one can be assured it is “view-consistent”, meaning compatible with some 3D shape.

As for simple DDFs, we begin by defining a relatively unconstrained field, for which we shall find constraints relating to view consistency.

Definition G.15: BOZ Field

Let $\xi : \Gamma \rightarrow \{0, 1\}$ be a binary-valued field on Γ . We restrict ξ such that its zero set is open; i.e., $\forall \tau \in \Gamma$, if $\xi(\tau) = 0$, then $\exists \varepsilon > 0$ such that $\xi(\tilde{\tau}) = 0 \forall \tilde{\tau} \in B_\varepsilon(\tau) \cap \Gamma$, where $B_\varepsilon(\tau)$ is the open ball centered at τ of radius ε .

We call any such binary field, with an open zero set, a BOZ field.

Similar to NNBC fields (Def. G.6), a BOZ field does not have an obvious connection to 3D shapes or their multiview silhouettes – one can imagine BOZ fields that do not represent

¹⁶Recall that if $q(p_1, v_1)$ is a predicted surface point, then $d(p_2, v_2) \leq \|p_2 - q(p_1, v_1)\|$, for any ray $r(p_2, v_2)$ that intersects $q(p_1, v_1)$.

¹⁷Local in the sense that each constraint can be defined about a single point, particularly IO_d (a point that is zero along one ray must be zero on all of them) and DE_d (the along-ray derivative at any point is -1).

a consistent 3D point set from all viewpoints. Our goal is to understand when (or under what conditions) a BOZ field acts as a continuous representation of the silhouettes of some coherent shape (i.e., when it assigns every ray a binary indication as to whether a shape point exists along that ray or not). It turns out a very similar situation to that of Simple DDFs arises. Thus, similarly, we begin by defining field constraints, which will eventually connect BOZ fields to coherent silhouettes of shapes.

Constraint Definition G.4: Non-Visible Boundary Condition for Visibility Fields (BC_ξ)

A BOZ field ξ satisfies the *Non-Visible Boundary Condition*, denoted BC_ξ , if $\xi(p, v) = 0 \forall (p, v) \in \mathcal{O}[\mathcal{B} \setminus \mathcal{B}_\varepsilon, \mathcal{B}_\varepsilon]$, where $\mathcal{O}[V, U] = \{\tau = (p, v) \in \Gamma \mid p \in V, r_\tau(s) \notin U \forall s \geq 0\}$ is the set of “outward rays” from a given set.

Recall that $\mathcal{B} \setminus \mathcal{B}_\varepsilon$ is the “outer shell” of the bounding domain and, essentially, $\mathcal{O}[V, U]$ contains oriented points (or, equivalently, rays) that *start* in V and do *not* intersect U . In other words, BC_ξ demands that any rays that (i) start at a point close to the boundary and (ii) do not intersect the inner part of the domain (i.e., \mathcal{B}_ε) must have a value of zero. Intuitively, if we imagine ξ to be representing an object in \mathcal{B}_ε , this condition prevents the object from being visible on rays that start outside of \mathcal{B}_ε and “look away” from \mathcal{B}_ε .

Constraint Definition G.5: Directed Eikonal Constraint for Visibility Fields (DE_ξ)

A BOZ field ξ satisfies the *Directed Eikonal Constraint* DE_ξ if ξ is always non-increasing along a ray (i.e., $f_\xi(s_1|\tau) \leq f_\xi(s_2|\tau) \forall \tau \in \Gamma, s_1 > s_2$).

Intuitively, along a ray, this constraint means that the visibility can only ever go from “seeing” to “not seeing” an object (i.e., along a ray, one should not suddenly see a new point in the shape become visible).

Constraint Definition G.6: Isotropic Opacity for Visibility Fields (IO_ξ)

A BOZ field d satisfies the *Isotropic Opacity Constraint* at a point $q \in \mathcal{B}$ if the following implication holds:

$$\begin{aligned} \exists v \text{ s.t. } & \left(\lim_{\epsilon \downarrow 0} \sup_{s \in (0, \epsilon)} \xi(q - sv, v) = 1 \right) \wedge \\ & \left(\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} \xi(q + sv, v) = 0 \right) \\ \implies & (\xi(\tau) = 1 \forall \tau \in \mathcal{I}_q), \end{aligned} \quad (83)$$

where \mathcal{I}_q is the set of rays in \mathcal{B} (i.e., r_τ , for oriented points $\tau \in \mathcal{B}$) that intersect q . An *isotropically opaque field* (IO_ξ) satisfies the Isotropic Opacity Constraint everywhere.

In words, if there is *any* direction along which the field flips from one to zero, then the point at which that occurs must be “isotropically opaque” (i.e., from any direction, any ray that hits that point must produce a visibility value of one). Intuitively, IO_ξ says that such “one-to-zero” discontinuities in ξ are special. As we shall see, they act similar to surface points (or zeroes of d , in the case of simple DDFs). In particular, for any point $q \in \mathcal{B}$, if even a single direction v exists such that a one-to-zero discontinuity occurs along v at q , then any ray that contains q must also be visible.

As for simple DDFs, with these three constraints, we can now define an appropriately limited set of fields, called visibility fields, which we will show is intimately tied to shapes.

Definition G.16: Visibility Field

We define a Visibility Field to be a BOZ field such that IO_ξ , BC_ξ , and DE_ξ are satisfied.

Analogously to simple DDFs with left continuity (Def. G.9), we utilize a notational convention for visibility, which slightly broadens when an oriented point is considered visible.

Definition G.17: Upper-Continuity Convention for Visibility Fields

Whenever a visibility field ξ approaches one, we denote the field value as one, regardless of the nature of any discontinuities at such points. In particular, for any $q \in \mathcal{B}$, we write:

$$\xi(q, v) = 1 \iff \lim_{\epsilon \downarrow 0} \sup_{\substack{||\delta|| < \epsilon \\ q + \delta \in \mathcal{B}}} \xi(q + \delta, v) = 1. \quad (84)$$

In other words, we assume “upper continuity” for ξ , along all directions. Notice that this is consistent with the requirement on BOZ fields that the zero set $\xi^{-1}(0)$ is open. Further, its complement, $\xi^{-1}(1) \subseteq \Gamma = \mathcal{B} \times \mathbb{S}^2$, is closed (using the fact that \mathcal{B} is compact).

Definition G.18: Positional Discontinuities of the Visibility Field

Consider a visibility field ξ . Then we define its positional discontinuity set as follows:

$$Q_\xi := \mathcal{C} \left(\left\{ q \in \mathcal{B} \mid \exists v \text{ s.t. } \lim_{\epsilon \downarrow 0} \sup_{s \in (0, \epsilon)} \xi(q - sv, v) = 1 \right. \right. \\ \left. \left. \wedge \lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} \xi(q + sv, v) = 0 \right\} \right), \quad (85)$$

where \mathcal{C} is the closure operator on sets.

Intuitively, given a visibility field ξ , the set Q_ξ simply contains all of the one-to-zero discontinuities in ξ . More specifically, for any point $q \in \mathcal{B}$, if there exists some direction v such that the field ξ flips from one to zero (i.e., from “seeing something” to “seeing nothing”) at q , then $q \in Q_\xi$.

Lemma G.2: Properties of Positional Discontinuities of Visibility Fields

A visibility field ξ satisfies (i) $Q_\xi \subseteq \mathcal{B}_\epsilon$, (ii) Q_ξ is a shape, and (iii) there cannot be discontinuities in ξ from zero-to-one along a ray.

Proof. See Supp. §H-B1. ■

Given these properties, the discontinuities Q_ξ are a good candidate for a potential shape that ξ is implicitly representing. First, like for simple DDFs, we need to define what it means for a shape to *induce* a field (i.e., “exactly” represent it).

Definition G.19: Shape-Induced Binary Fields

Let $S \subseteq \mathcal{B}_\epsilon$ be a shape. Recall that $\mathcal{I}_S \subseteq \Gamma$ is the set of rays that intersect S . Then we say a BOZ field ξ is *induced* by S if the following holds:

$$\xi(\tau) = 1 \iff \tau \in \mathcal{I}_S. \quad (86)$$

Further, given a fixed shape S , its associated induced field ξ is unique.

Notice that the induction relation is an “iff”, meaning $\xi(\tau) = 0$ implies that *no* $q \in S$ can be present along the ray r_τ . In addition, importantly, while the S -induced field ξ is unique, the inducing shape for a given ξ is *not* necessarily unique. For instance, imagine a sphere with another shape (say, another sphere) inside it – such a shape will induce the same visibility field as using the outermost sphere alone for the inducement. Given this definition, our first result in this section shows that any shape-induced field satisfies the constraints of a visibility field.

Theorem G.3: Shape-Induced Binary Fields Are Visibility Fields

Let $S \subseteq \mathcal{B}_\epsilon$ be a shape, with an induced field ξ . Then ξ is a visibility field (i.e., IO_ξ , BC_ξ , and DE_ξ hold). Also, the positional discontinuities of ξ satisfy $Q_\xi \subseteq S$.

Proof. See Supp. §H-B2. ■

So far we have shown that any shape can be used to generate or induce a visibility field (in the sense of Def. G.16). We can use this to define a notion of shape representation for visibility fields.

Definition G.20: Shape Indicators as View-Consistent Visibility Fields

A visibility field ξ is View Consistent (VC) iff there exists a shape S such that S induces ξ . We can therefore call any such ξ a *shape indicator* for the point set S .

Note we call it a shape indicator instead of a representation because we cannot (necessarily) reconstruct S from ξ (i.e., it would be an incomplete representation). Instead, we call it an indicator because it “indicates” whether or not the shape (or some shape) exists along the given ray (i.e., acts as an indicator

function for each ray). We next show that the local constraints of a visibility field are sufficient for it to be a representation of some shape.

Theorem G.4: Every Visibility Field is a Shape Indicator

Let ξ be a visibility field. Then there exists a point set S such that ξ is induced by S . Further, we can take the positional discontinuities as an inducer; i.e., $S := Q_\xi$ induces ξ .

Proof. See Supp. §H-B3. ■

So far, we have shown a close duality between visibility fields and shape indicators (i.e., between binary fields with specific local constraints on the field, namely IO_ξ , DE_ξ , and BC_ξ , and binary fields constructed from a given point set). However, it is clear that many inducing point sets can induce the same field. Unlike the case with simple DDFs, this is not merely a question of behaviour outside the domain (e.g., recall the example of a sphere within a sphere inducing the same field as the outer sphere alone). We can therefore ask for a “minimal example” over an equivalence class of inducing point sets (i.e., given many shapes S that all induce the same ξ , which is the simplest?).

This leads us to the following corollary, which shows that the simplest (i.e., smallest) shape (i.e., closed point set) that induces a given visibility field is the set of positional discontinuities.

Corollary G.1: Minimal Characterization of Visibility Fields

Consider the set of shapes $\mathfrak{S}[\xi] = \{S \mid S \subseteq \mathcal{B} \text{ induces } \xi\}$ that all induce the same field ξ . I.e., $\mathfrak{S}[\xi]$ is an equivalence class of inducing shapes. Then, $Q_\xi \in \mathfrak{S}[\xi]$ is the smallest closed point set among all such inducers.

Proof. See Supp. §H-B4. ■

This shows that Q_ξ is special among shapes that induce ξ , in that it forms a subset of any other shape that also induces ξ . Next, we provide some additional intuition concerning Q_ξ , which is derived from the visibility field itself. Instead, given any shape S , we can describe which subset of points actually forms $Q_\xi \subseteq S$, without referring to any ξ directly.

Definition G.21: Directly Lit Points

Consider a shape S contained in \mathcal{B}_ε . Imagine a set of uniform lights placed at each $p \in \partial\mathcal{B}$ (or, equivalently, imagine each such p casting rays in all possible directions. We denote the set of points in S that are “directly lit” by such lights (or intersected *first* by such cast rays) as $\mathfrak{D}_\ell(S)$. Mathematically, the directly lit points (DLPs) can be written

$$\mathfrak{D}_\ell(S) = \mathcal{C} \left(\left\{ q \in S \mid \exists \tau = (p, v) \in \Gamma_\partial \text{ s.t.} \right. \right. \\ \left. \left. q = \arg \min_{\tilde{q} \in S_\tau} \|p - \tilde{q}\| \right\} \right), \quad (87)$$

where $\Gamma_\partial = \partial\mathcal{B} \times \mathbb{S}^2$ and S_τ is the set of points in S that are intersected by the ray r_τ .

Intuitively, any point in S that can be hit by a light ray from the boundary, directly and without occlusion, is a DLP. One could, equivalently, call $q \in \mathfrak{D}_\ell(S)$ an *observable* point from the boundary. Similar to Q_ξ , we include a closure operator, to ensure that $\mathfrak{D}_\ell(S)$ contains its limit points and boundary. Without this, the directly lit set is not necessarily closed.

We next show that directly lit points are an alternative characterization of the positional discontinuities. Intuitively, it relates the local field property of a one-to-zero discontinuity to the “observable” behaviour of the shape from the boundary $\partial\mathcal{B}$. In other words, DLPs are defined via the point set S , *not* its induced field ξ . However, we next show that $\mathfrak{D}_\ell(S)$ is a special subset of S , in that it forms a minimal inducer for the equivalence class $\mathfrak{S}[\xi]$, and it is therefore equal to Q_ξ .

Corollary G.2: Directly Lit Points Minimally Characterize Visibility Fields

Given any shape S in the equivalence class $\mathfrak{S}[\xi]$ of inducers for the field ξ , the set of DLPs $\mathfrak{D}_\ell(S)$ is a *minimal* inducer of ξ and furthermore $\mathfrak{D}_\ell(S) = Q_\xi$.

Proof. See Supp. §H-B5. ■

In words, to get the minimal inducing point set from a shape S , one need not construct the induced field and then find its discontinuities. Instead, one can compute the DLPs directly from S , by “looking inwards” from the boundary and finding the observable (i.e., directly lit) points.

Summary. We have shown that (i) we can devise a notion of “inducement”, under which any 3D shape gives rise to an associated visibility field, ξ , which acts as its binary indicator function along rays; (ii) the local constraints that define a visibility field ξ (i.e., IO_ξ , DE_ξ , BC_ξ) are sufficient to guarantee that *some* point set must exist that induces ξ (meaning, every visibility field is the indicator of some shape); (iii) the positional discontinuities Q_ξ play a special role, as the *minimal* closed subset of S that can induce a given field ξ (in some sense, Q_ξ represents the best reconstruction of the initial shape that one can obtain from ξ); and (iv) an alternative but equivalent characterization of the minimal closed inducing subset of S can be given by the set of directly lit points (DLPs).

1) *Shape-from-Silhouette, Visual Hulls, and Visibility Fields*: The connection to DLPs (or boundary-observable points) allows us to interpret ξ more intuitively as the continuous field of silhouette images. In this setting, each pixel is associated to a ray that starts, say, from a camera outside the domain, and assigns a binary pixel value that decides whether or not the given shape is intersected by that ray. As already noted, given ξ , one can compute Q_ξ , which recovers as much information about the shape (or, rather, the equivalence class of shapes) that could have induced it as is possible. This is therefore closely related to prior work on shape-from-silhouette and visual hulls, which has seen significant theoretical investigation in prior works.

In particular, estimating 3D shape from visibility images alone (i.e., silhouettes) has been studied in computer vision for decades [2] (see also [92]), complemented by research into how humans overcome the inherent ambiguity of the problem (e.g., [137], [138]). Given that an infinite set of shapes can give rise to an identical silhouette, additional constraints are needed for shape inference, either in the form of purely geometric considerations [91], priors [139], or simply more viewpoints. In general, the basic idea is to use back-projection from the silhouettes into 3D space, with each such constraint forming a 3D *visual cone*, in which the shape must lie. The volumetric intersection of such cones generates the *visual hull* of the object, a point set that contains the shape of interest with increasing precision as the number of views increases. Significant theoretical work has characterized the geometry of visual hulls, particularly in the shape-from-silhouette context (e.g., [90], [93], [140]–[144]). Numerous techniques have been applied to visual hull computation (e.g., [145]–[148]), as well as applications thereof (e.g., [149]–[151]). Generalizations of the shape-from-silhouette problem consider errors in the camera poses, inconsistencies or noise in the silhouettes themselves, or issues related to non-convexity of shapes (e.g., [147], [152]–[156]). Space carving [94] generalizes shape-from-silhouette geometric constraints to utilize a notion of photo-consistency, producing a “photo-hull” that is more constrained than the visual hull; this relates more closely to recent AbS approaches to multiview 3D scene reconstruction, such as NeRFs.

In contrast, our work focuses on potentially under-constrained 5D visibility fields, which can be rendered into silhouettes from any continuously varying viewpoint, and the conditions under which such fields properly correspond to some underlying shape (i.e., when ξ is a shape representation, guaranteeing view consistency). The most common shape-from-silhouette situation involves so-called *external* visual hulls, which involve only viewpoints (i.e., cameras) outside the convex hull of the shape [93]. In the context of this thesis, a visibility field ξ considers the presence of a shape along any possible ray. This enables ξ to represent aspects of shape that are missed by external visual hulls, such as concavities, and is thus more similar to *internal* visual hulls (e.g., see [93], [142]), which assigns the “viewing region” to be anywhere outside the shape (not just its convex hull).

Nevertheless, questions still arise regarding visibility fields and their representational capacity. One aspect of shapes “missed” by visibility fields is *internal structure*: shape points

that are completely surrounded by other shape points cannot be preserved by the field. In other words, given a shape S with internal structure, any S -induced field will not be able to recover such structure.¹⁸ One can see this by noting that the directly lit points of ξ (equivalently, the set of one-to-zero discontinuities) will not include such internal structure. These points, however, fully encode the field (as they induce it). Combining a visibility field with a distance field, as we will do in the next section, enables a complete shape representation (i.e., including internal structure).

D. View Consistency for Directed Distance Fields

In this section, we investigate the *combination* of a visibility field and a slightly modified form of simple DDF, which define the full DDF discussed in the other sections of this chapter. Similar to the prior two sections, we define view consistency through a notion of “shape representation”; i.e., whether or not some shape (point set) exists that induces the DDF. In this case, the visibility field will be used to control where (in Γ) the distance field needs to be strongly constrained. This is different from the simple DDF case, in that it alleviates constraints on d . Given a simple DDF, one can imagine constructing a visibility field based on the presence or absence of a zero of the d -field along the ray (i.e., whether q maps a ray outside of \mathcal{B}). However, this section considers the opposite: rather than deriving ξ from d , we specify ξ first, and use it to derive where and how d must be constrained (reducing the demands on d). As a result of this, we need additional conditions to ensure that the visibility field and depth field are compatible with one each other.

Definition G.22: Visible and Non-Visible Ray Sets

Let ξ be a visibility field. We define the sets of visible and non-visible rays, respectively, as

$$\mathcal{R}_V[\xi] = \{\tau \in \Gamma \mid \xi(\tau) = 1\} \quad (88)$$

$$\mathcal{R}_N[\xi] = \{\tau \in \Gamma \mid \xi(\tau) = 0\}, \quad (89)$$

written via upper-continuity (Def. G.17). Further, $\mathcal{R}_V[\xi]$ is closed, while $\mathcal{R}_N[\xi]$ is open.

The closed or open nature of these sets is due to the BOZ field definition, which requires $\xi^{-1}(0)$ to be open, and so its complement $\xi^{-1}(1) \subseteq \Gamma = \mathcal{B} \times \mathbb{S}^2$ is closed, since Γ is compact.

Define $\mathcal{A}_{d,\xi}(q)$ as the *visible* infimum of d at q :

$$\mathcal{A}_{d,\xi}(q) = \lim_{\epsilon \downarrow 0} \inf_{\substack{s \in (0, \epsilon) \\ v \in \mathbb{S}^2 \mid \xi(q, v) = 1}} d(q - sv, v).$$

¹⁸This is why there is an equivalence class of shapes that induce a given visibility field (i.e., $\mathfrak{S}[\xi]$).

Definition G.23: ξ -Coherent Simple DDFs

Given a visibility field ξ , an NNBC field d is a ξ -coherent simple DDF iff it satisfies

- 1) (IO $_{d,\xi}$) *Isotropic opaqueness on visible rays.*

$$\begin{aligned} \forall (q, v) \in \mathcal{R}_V[\xi] : \\ (\mathcal{A}_{d,\xi}(q) = 0) \implies (d(q, v) = 0). \end{aligned} \quad (90)$$

- 2) (BC $_{d,\xi}$) *Positive boundary condition on visible rays.*

$$((p, v) \in \mathcal{R}_V[\xi] \wedge p \in \mathcal{B} \setminus \mathcal{B}_\varepsilon) \implies d(p, v) > 0. \quad (91)$$

- 3) (DE $_{d,\xi}$) *Directed Eikonal constraint on visible rays.*

For any $\tau \in \mathcal{R}_V[\xi]$ such that $d(\tau) > 0$, we have $\partial_s f_d(s \mid \tau)|_{s=0} = -1$.

In words, a ξ -coherent depth field d satisfies the isotropic opaqueness, boundary, and directed Eikonal constraints whenever $\xi = 1$ (i.e., it is a simple DDF on $\mathcal{R}_V[\xi]$). Outside of the visible rays $\mathcal{R}_V[\xi]$, d is essentially unconstrained.

Definition G.24: Locally Visible Depth Zeroes

Let ξ be a visibility field and d be ξ -coherent. Then the *locally visible zeroes* are given by

$$Q_{d,\xi} = \mathcal{C}(\{\tilde{q} \in Q_d \mid \exists v \in \mathbb{S}^2 \text{ s.t. } \xi(\tilde{q}, v) = 1\}), \quad (92)$$

where \mathcal{C} is closure and we use upper-continuity for ξ (Def. G.17).

This set includes any zero of d (i.e., $q \in Q_d$; see Def. G.10) that is “locally visible” along some direction (meaning, $\lim_{\epsilon \downarrow 0} \sup_{\|\delta\| < \epsilon} \xi(q + \delta, v) = 1$). The locally visible zeroes ($Q_{d,\xi}$) will play an analogous role to the zeroes Q_d for simple DDFs and the discontinuities Q_ξ for visibility fields.

Definition G.25: Directed Distance Field (DDF)

A BOZ field ξ and an NNBC field d form a *Directed Distance Field* (ξ, d) iff:

- 1) ξ is a visibility field.
- 2) d is a ξ -coherent simple DDF.
- 3) Locally visible depth zeroes are isotropically opaque with respect to ξ :
 $q \in Q_{d,\xi} \implies \tau \in \mathcal{R}_V[\xi] \forall \tau \in \mathcal{I}_q$, where $\mathcal{I}_q \subset \Gamma$ is the set of rays that intersect q .
- 4) Every visible ray must hit a locally visible depth zero:
 $\tau \in \mathcal{R}_V[\xi] \implies \mathcal{B}_\tau \cap Q_{d,\xi} \neq \emptyset$, where \mathcal{B}_τ is the set of points in \mathcal{B} intersected by r_τ .

The definition of a DDF can be intuited as combining a visibility field with a simple DDF that is constrained only along visible rays, along with two additional restrictions that force the two to be consistent. The consistency constraints hinge on special zeroes of d that are “locally visible” (i.e., at least one ray exists, along which $q \in Q_d$ is visible at an infinitesimal distance). Such zeroes are (i) always visible and

(ii) every visible ray must intersect at least one. The next lemma clarifies the relation between the fundamental point-sets associated with these various fields: the zeroes of d , the visibility discontinuities, and the locally visible depth zeroes.

Lemma G.3: Fundamental Point Sets

Let (ξ, d) be a DDF. Then $Q_\xi \subseteq Q_{d,\xi} \subseteq Q_d$.

Proof. See Supp. §H-C1. ■

Note that locally visible zeroes (i.e., points in $Q_{d,\xi}$) require all intersecting rays to be visible rays, but they do *not* require such points to be one-to-zero visibility discontinuities (i.e., Q_ξ). However, such discontinuities *do* have to be zeroes of the depth field (i.e., in Q_d).

Next, we connect the locally visible depth zeroes to a closely related set of points, the directly visible zeroes, where the relation is more non-local. Later, this relation will be useful in determining the relationship between shape-induced field tuples (soon to be defined, in Def. G.27) and DDFs.

Definition G.26: Directly Visible Depth Zeroes

Let (ξ, d) form a DDF. Then the set of *directly visible depth zeroes* is given by

$$\mathcal{D}_{d,\xi} = \mathcal{C}(\{\tilde{q} \in \mathcal{B} \mid \exists (p, v) \in \mathcal{R}_V[\xi] \text{ s.t. } q(p, v) = \tilde{q}\}), \quad (93)$$

where $q(p, v)$ is the q -mapping of d .

In words, $\mathcal{D}_{d,\xi} \subseteq \mathcal{B}$ records positions where at least one ray τ exists such that (i) $\xi(\tau) = 1$ and (ii) \tilde{q} is the *first* zero (i.e., member of Q_d) along r_τ . Thus, equivalently, using the properties of ξ -coherence, we may write

$$\begin{aligned} \mathcal{D}_{d,\xi} = \text{Closure}(\{\tilde{q} \in \mathcal{B} \mid \exists \tau \in \mathcal{R}_V[\xi] \text{ s.t.} \\ \tilde{q} = p + \tilde{s}(\tau \mid Q_d)v\}), \end{aligned} \quad (94)$$

where $\tilde{s}(\tau \mid S) = \min_{q \in S_\tau} \|p - q\|$, $\tau = (p, v)$, and S_τ is the set of points in $S \subseteq \mathcal{B}$ that are intersected by r_τ . This latter form helps illuminate the connection to shape induction (e.g., recall Def. G.11 for simple DDFs). First, however, we show that not only are the locally and directly visible depth zeroes equivalent, but that the outputs of d and q (the positional field of d) are effectively controlled by them.

Lemma G.4: DDFs Map to Locally Visible Depth Zeroes

Let (ξ, d) be a DDF. Then $Q_{d,\xi} = \mathcal{D}_{d,\xi}$. Further:

$$\begin{aligned} \forall \tau = (p, v) \in \mathcal{R}_V[\xi] : d(\tau) = \min_{\tilde{q} \in [\mathcal{D}_{d,\xi}]_\tau} \|\tilde{q} - p\| \text{ and} \\ q(\tau) \in Q_{d,\xi}, \end{aligned} \quad (95)$$

where $[\mathcal{D}_{d,\xi}]_\tau$ is the set of points in $\mathcal{D}_{d,\xi}$ that intersect r_τ and q is the q -mapping of d .

Proof. See Supp. §H-C2. ■

This lemma will be used later, to link DDFs to shape representations. Following the previous two sections, on simple DDFs and visibility fields, we next define an approach to “field

induction”: given a shape S , how can we construct a DDF, consisting of a “field pair” (ξ, d) , that appropriately represents S ? We then show that inducing a field in this manner, versus having a DDF with constraints on local field behaviour (i.e., Def. G.25) are equivalent.

Definition G.27: Induced Field Pair

Consider a shape $S \subseteq \mathcal{B}_\varepsilon$. By Def. G.19, S induces a unique visibility field, ξ . We also define an induced distance field d to be any NNBC field that satisfies

$$d(\tau) = \min_{q \in S_\tau} \|q - p\| \quad \forall \tau \in \mathcal{R}_V[\xi], \quad (96)$$

where S_τ is the set of points in S that are intersected by τ .

Then, given a shape S , we define any such pair (ξ, d) , to be an *induced field pair*.

Notice that this construct for d is similar to inducement of simple DDFs; the main difference is that $d(\tau)$ is unconstrained when $\xi(\tau) = 0$ (i.e., d no longer needs to predict distances outside the domain). Further, the induced visibility $\xi(\tau)$ is one iff τ intersects S . Hence, for a visible ray $\tau \in \mathcal{R}_V[\xi]$, d is constrained to always predict the distance to the closest $q \in S_\tau$.

Thus, similar to Theorems G.1 and G.3, when a shape S induces a field pair F , we want to show that F follows the field requirements of DDFs (Def. G.25).

Theorem G.5: Induced Field Pairs are Directed Distance Fields

Let S be a shape and $F = (\xi, d)$ be an S -induced field pair. Then F is a Directed Distance Field, which satisfies $Q_{d,\xi} = S$.

Proof. See Supp. §H-C3. ■

Now that we know that any field pair generated from a shape is a DDF, we next want to work towards showing the converse, that any DDF is a shape representation of some shape. First, similar to simple DDFs, the presence of the visibility field means we need to define a notion of equivalence between fields (where values on non-visible rays can be ignored).

Definition G.28: Equivalence of Directed Distance Fields

Consider two DDFs, $F_1 = (\xi_1, d_1)$ and $F_2 = (\xi_2, d_2)$. Then F_1 and F_2 are equivalent iff (i) $\xi_1(\tau) = \xi_2(\tau) \quad \forall \tau \in \Gamma$ and (ii) $d_1(\tau) = d_2(\tau) \quad \forall \tau \in \mathcal{R}_V[\xi_1]$.

This definition of field equality enables us to establish when a DDF is a shape representation.

Definition G.29: View Consistent Directed Distance Fields as Shape Representations

A DDF F is view consistent (VC) iff it is equivalent to a DDF that has been induced by a shape S . I.e., F is VC iff there exists a shape S such that F is an S -induced field pair. In such a case, we say that F is a *shape representation* for S .

Finally, analogously to Theorems G.2 and G.4, we can assert that the field constraints of Def. G.25 are sufficient to guarantee that a DDF represents some shape.

Theorem G.6: Every Directed Distance Field is a Shape Representation

Let $F = (\xi, d)$ be a DDF. Then F must be view consistent (i.e., there exists a shape S such that F is an S -induced field pair). Further, $Q_{d,\xi}$ is a shape that induces F .

Proof. See Supp. §H-C4. ■

This last theorem finally links *full* DDFs to shape representations, by combining view-consistent visibility fields with depth fields that are constrained (in the same manner as simple DDFs) along visible rays (i.e., in $\mathcal{R}_V[\xi]$) and *unconstrained* along non-visible rays (i.e., in $\mathcal{R}_N[\xi]$). In other words, full DDFs are simple DDFs, but with looser requirements on d when $\xi = 0$ (controlled instead by a visibility field, ξ). The concept of shape representation can therefore be equivalently derived in two ways: (i) by starting from a point set and inducing a field pair from it, or (ii) by starting from a visibility field and a distance field that are “compatible” with one another, and satisfying the basic requirements on the fields (BC, IO, and DE).

The latter case, in particular, is important, as it suggests that one need only ensure certain local constraints, in order to obtain a view consistent 5D field. Specifically, for the boundary conditions (BC), one should demand that, near the boundary of the domain, d does not predict surface points (i.e., zeroes) and the visibility field returns zero whenever it “looks outward” from the domain (i.e., the shape cannot be visible outside the domain). For the directed Eikonal (DE) constraints, one requires $\partial_s f_d = -1$ and f_ξ non-increasing, along visible and all rays, respectively. Isotropic opaqueness (IO) asks that (visible) zeroes of d are zero from all directions and, similarly, that one-to-zero “flips” in ξ are visible from all directions. Finally, two demands are placed to ensure consistency between the visibility and depth fields, relying on the set of locally visible zeroes (i.e., depth zeroes that are visible at infinitesimally close points along some ray): (a) any ray that intersects such a zero must be visible, and (b) any visible ray must hit such a zero. Together, the previous theorems ensure that a field satisfying these conditions must be a shape representation (and vice-versa).

This is reminiscent of solving the Eikonal equation to obtain a distance field, which is loosely followed by some modern distance-based isosurface ISF training methods. Notice, for instance, that BC and IO act like boundary conditions, while

DE forms the constraint that needs to be solved across space. We expect that improvements in techniques for fitting higher-dimensional INRs with differential constraints (in this case, equivalent to solving a 5D partial differential equation) will make neural implementation of (P)DDFs more viable.

In summary, these theorems give requirements for generic learned fields (i.e., implemented as neural networks (ξ_θ, d_θ) with weights θ) that ensure view-consistency if satisfied. By careful design of the field architectures and losses, to ensure fitting to the data and satisfaction of these conditions, we should converge to the desired shape, given enough network capacity and sufficient supervision.

APPENDIX H

PROOFS FOR VIEW CONSISTENCY OF DDFs

This section provides the proofs for the Lemmas, Theorems, and Corollaries of §VII in the main paper and Supp. §G. It also includes an alternative definition of field induction for simple DDFs (§H-A1) and an example of why the positional zeroes need to be defined as converging over all directions, rather than just a single one (§H-D). Note that we denote a shape S to be a compact point set in the domain \mathcal{B}_ε .

A. Proofs for Simple DDFs

1) Alternative Formulation of Field Induction:

Definition H.1: Field Induction via Boundary-Subsumed Shapes

Let S be a shape (closed point set) such that $S \subseteq \mathcal{B}_\varepsilon$. Define $S^\partial = S \cup \partial\mathcal{B}$. Let S_τ^∂ be the set of points in S^∂ that are intersected by r_τ (i.e., for any $q \in S_\tau^\partial$, there exists $t \geq 0$ such that $r_\tau(t) = q$). Then S^∂ induces a field d in the boundary-subsumed sense iff $d(\tau) = \min_{q \in S_\tau^\partial} \|q - p\|$, for any $\tau \in \Gamma$.

Note that S is closed by definition, while S_τ^∂ must be closed and non-empty (because we include the boundary in the definition of S_τ^∂ , there must be at least one intersection); hence, the minimum must exist. In this case, the induced field is unique and equality may be used instead of equivalence, for simple DDFs. Finally, notice that saying S induces d in the boundary-subsumed sense is equivalent to saying $S^\partial = S \cap \partial\mathcal{B}$ induces d (in the non-boundary-subsumed sense). However, in the boundary-subsumed case, S has a *unique* induced field, for a fixed \mathcal{B} . On one hand, this simplifies the notion of equivalence between simple DDFs. However, in the main thesis, we utilize the more complex notion of equivalence, as it (i) keeps the shape separate from the domain boundary and (ii) is more analogous to the scenario we encounter for full DDFs.

2) Properties of Positional Zeroes:

Lemma G.1: Properties of Simple DDF Positional Zeroes

Let d be a simple DDF. Then $Q_d \subseteq \mathcal{B}_\varepsilon$ and Q_d is a closed set.

Proof. By BC_d , $d(p, v) > 0$ for all $p \in \mathcal{B} \setminus \mathcal{B}_\varepsilon$, $v \in \mathbb{S}^2$, so there cannot be any zeros of the field outside of \mathcal{B}_ε .

Next we show that the set is closed. Since d is a simple DDF, it satisfies IO_d , BC_d , and DE_d .

Suppose that Q_d is not closed. Then there must exist a p s.t. there is a sequence $(p_k)_{k=1}^\infty$ that converges to p , where every $p_k \in Q_d$, but $p \notin Q_d$ (i.e., p is the limit of points p_k , all of which are in Q_d). Since any $p_k \in Q_d$, by IO_d , we have that $d(p_k, v) = 0 \forall v \in \mathbb{S}^2$.

Let $\varepsilon > 0$. Choose p_k s.t. $\|p_k - p\| < \varepsilon$ (such a p_k must exist as the sequence is converging to p) and set $v_k := \text{Unit}(p_k - p) \in \mathbb{S}^2$ to be the direction vector pointing from p back to p_k . By IO_d , $d(p_k, v_k) = 0$ and so, by DE_d , we must have that $d(p, v_k) \leq \|p_k - p\| < \varepsilon$. But since ε is arbitrary, we must have that

$$\lim_{\varepsilon \downarrow 0} \inf_{s \in (0, \varepsilon)} d(p - sv, v) = 0, \quad (97)$$

which implies $p \in Q_d$.

This is a contradiction; hence, Q_d must be closed. ■

3) Shape-Induced Fields Are Simple DDFs:

Theorem G.1: Shape-Induced Fields Are Simple DDFs

Consider a field d that is induced by some shape S . Then d is a simple DDF. Further, the positional zeroes of d in \mathcal{B} satisfy $S = Q_d$.

Proof. We need to show that d , induced by S , satisfies the three constraints of a simple DDF.

- Let d be induced by S . Let $\tau = (p, v) \in \Gamma$ be such that $S_\tau \neq \emptyset$, and $\tau_e = (p_e, v_e) \in \Gamma$ be such that $S_{\tau_e} = \emptyset$.
- Since $S \subseteq \mathcal{B}_\varepsilon$ is compact, $d(\tau) = \min_{q \in S_\tau} \|q - p\| \geq 0$ is both non-negative and bounded. By definition, $d(\tau_e) > 0$ is also non-negative and bounded. Finally, for any τ (with $S_\tau \neq \emptyset$), $d(\tau) = \min_{q \in S_\tau} \|q - p\|$ is clearly C^1 along the ray, with potential discontinuities at zeroes of d ; for any τ_e , d is defined to be C^1 along the ray as well.
- (IO_d) First, note that there can be no zeros of d except on rays that intersect S , since a ray with no intersections (e.g., τ_e) is defined to have positive d ; further, those zeroes must be *at* the intersection with S , since, by definition, at all other points, d measures the distance to the closest of those points along the ray. So we need only check points in S .

So, consider any $q_0 \in S$, and choose an arbitrary $v \in \mathbb{S}^2$. Consider the oriented point $\tau(s) := (q_0 - sv, v)$, so that $S_{\tau(s)}$ is the set of points in S intersected by the ray $r_{\tau(s)}$, which is a function of s . Looking at d , as we move along v towards q_0 , shows us that

$$\lim_{\varepsilon \downarrow 0} \inf_{s \in (0, \varepsilon)} d(q_0 - sv, v) \quad (98)$$

$$= \lim_{\varepsilon \downarrow 0} \inf_{s \in (0, \varepsilon)} \min_{q \in S_{\tau(s)}} \|q - (q_0 - sv)\| \quad (99)$$

$$\leq \lim_{\varepsilon \downarrow 0} \inf_{s \in (0, \varepsilon)} \|q_0 - q_0 + sv\| = 0. \quad (100)$$

Since $d \geq 0$ as well, we must have that $\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q_0 - sv, v) = 0$. Since this holds for arbitrary v , IO_d also holds.

- (BC_d) By definition, S is a subset of \mathcal{B}_ϵ . Hence, the induced field d cannot have any zeroes in $\mathcal{B} \setminus \mathcal{B}_\epsilon$. Since d is non-negative, we must have $d(p, v) > 0$ for any $v \in \mathbb{S}^2$, $p \in \mathcal{B} \setminus \mathcal{B}_\epsilon$. Thus, BC_d holds.
- (DE_d) Consider any intersecting oriented point, $\tau = (p, v) \in \Gamma$, s.t. $d(p, v) = \min_{q \in S_\tau} \|q - p\| = \|q_0 - p\| > 0$ (i.e., $d(\tau)v + p = q_0 \in S$). First, note that $f_d(s | \tau) = \|q_0 - (p + sv)\| = \|d(\tau)v + p - (p + sv)\| = \|[d(\tau) - s]v\| = d(\tau) - s$ for $s \in [0, d(\tau)]$. Then, clearly, $\partial_s f_d(s | \tau) = \partial_s (d(\tau) - s) = -1$, as required. For a non-intersecting ray, by construction, we have $d(p, v) > 0$ and $\partial_s f_d(s | p, v) = -1$. Hence, d satisfies DE as well.
- Finally, as noted above (see the section on IO_d), we know that the zeros of an induced field inside \mathcal{B} can only occur at intersections with S . Hence, any $q \in \mathcal{B}$ s.t. $q \in Q_d$ must also satisfy $q \in S$. Since $S \subseteq \mathcal{B}_\epsilon$, this also restricts Q_d to points in \mathcal{B}_ϵ (recall, by the lemma above, that $Q_d \subseteq \mathcal{B}_\epsilon$ for any simple DDF). There may be additional zeros of d outside the domain \mathcal{B} , which we can ignore. Hence, we have $S = Q_d$.
- Altogether, we have shown that any closed point set in \mathcal{B}_ϵ can induce or generate a simple DDF. Further, the zeros of the induced field (in the domain) coincide with the S itself.

4) Every Simple DDF is a Shape Representation:

Theorem G.2: Every Simple DDF is a Shape Representation

Any simple DDF d is view consistent (i.e., there exists a shape $S \subseteq \mathcal{B}_\epsilon$ such that d is equivalent to any field induced by S). Further, d is a shape representation for the set formed by its positional zeroes, Q_d (i.e., Q_d induces d).

Proof. Let d be a simple DDF on \mathcal{B} . We show that its positional zeroes, Q_d , in \mathcal{B}_ϵ form the inducing shape. Recall that any inducing shape must be in \mathcal{B}_ϵ and compact (and note that \mathcal{B}_ϵ itself is compact).

- **Q_d is a Shape**
By Lemma G.1, $Q_d \subseteq \mathcal{B}_\epsilon$ and Q_d is closed; hence, Q_d is a compact point set (i.e., shape). Thus, we can simply set $S = Q_d$.
Now, we also already know that any S -induced field, say d_s , must be a simple DDF. We simply need to show that d_s (which was induced by Q_d) is equivalent to d .
- **Empty rays**
Suppose $\mathcal{E}[d_s] \neq \mathcal{E}[d]$. Then either (a) $\exists \tau_a \in \mathcal{E}[d_s] \setminus \mathcal{E}[d]$ or (b) $\exists \tau_b \in \mathcal{E}[d] \setminus \mathcal{E}[d_s]$.
Case (a): $\tau_a \notin \mathcal{E}[d] \implies \exists q = p + d(\tau_a)v \in \mathcal{B}_\epsilon$. (Note that $q \notin \mathcal{B} \setminus \mathcal{B}_\epsilon$ due to BC_d.) But by DE_d, $d(q, v) = 0 \implies q \in Q_d$ and thus $q \in S$, contradicting $\tau_a \in \mathcal{E}[d_s]$. (I.e., τ_a cannot be in $\mathcal{E}[d_s]$ if it is not in $\mathcal{E}[d]$.)

Case (b): $\tau_b \notin \mathcal{E}[d_s] \implies \exists q = p + d(\tau_b)v \in \mathcal{B}_\epsilon$. In particular, τ_b must intersect some $q \in S$ along r_{τ_b} (where $S \subseteq \mathcal{B}_\epsilon$). But this implies $q = q(\tau_b) \in \mathcal{B}_\epsilon$, which contradicts $\tau_b \in \mathcal{E}[d]$. (I.e., τ_b cannot be in $\mathcal{E}[d]$ if it is not in $\mathcal{E}[d_s]$.)

Thus, both cases lead to a contradiction. Hence, the empty sets coincide: $\mathcal{E}[d_s] = \mathcal{E}[d]$.

• Intersecting rays

Consider an intersecting ray $\tau = (p, v) \in \Gamma$ (i.e., non-empty ray). For the shape-induced field, we have $d_s(p, v) = \min_{q \in S_\tau} \|q - p\| = \|q_s - p\|$, where q_s is the closest intersection point along r_τ . We therefore know $q_s \in S = Q_d$ and there cannot be any other $q \in Q_d$ (i.e., zeros of d) between p and q_s along τ . Thus, by DE_d, along τ , d must decrease at a unit rate, until it reaches the zero at q_s (or, equivalently, f_d must increase at a unit rate, going from q_s to p). Hence, $d(p, v) = \|q_s - p\| = d_s(p, v)$, as required.

This shows that d and d_s are equivalent. In other words, given a Simple DDF d , we are guaranteed that d is a view-consistent shape representation of its zero-set, Q_d . ■

B. Proofs for Visibility Fields

1) Properties of Positional Discontinuities of Visibility Fields:

Lemma G.2: Properties of Positional Discontinuities of Visibility Fields

A visibility field ξ satisfies (i) $Q_\xi \subseteq \mathcal{B}_\epsilon$, (ii) Q_ξ is a shape, and (iii) there cannot be discontinuities in ξ from zero-to-one along a ray.

Proof. Note that the closure operator in Def. G.18 ensures the set Q_ξ is closed by definition. Also, by DE_ξ, there cannot ever be discontinuities from zero to one along a ray. Further, since ξ satisfies BC_ξ, note that $Q_\xi \cap (\mathcal{B} \setminus \mathcal{B}_\epsilon) = \emptyset$, because (by IO_ξ) any ray τ that intersects a $q \in Q_\xi$ (i.e., satisfies $\tau \in \mathcal{I}_q$) must have $\xi(\tau) = 1$ including outward-facing rays, which would violate BC_ξ if they were present. Hence, $Q_\xi \subseteq \mathcal{B}_\epsilon$. This shows that Q_ξ is closed and bounded, and hence is a shape. ■

2) Shape-Induced Binary Fields Are Visibility Fields:

Theorem G.3: Shape-Induced Binary Fields Are Visibility Fields

Let $S \subseteq \mathcal{B}_\epsilon$ be a shape, with an induced field ξ . Then ξ is a visibility field (i.e., IO_ξ, BC_ξ, and DE_ξ hold). Also, the positional discontinuities of ξ satisfy $Q_\xi \subseteq S$.

Proof. We need to show that a shape-induced binary field satisfies the constraints of a visibility field.

- (BC_ξ) Since $S \subseteq \mathcal{B}_\epsilon$, any outward-pointing ray $\tau \in \mathcal{B} \setminus \mathcal{B}_\epsilon$ in the ϵ -boundary cannot intersect a $q \in S$. Hence, such rays satisfy $\xi(\tau) = 0$ as required.

- (DE $_{\xi}$) Consider a ray $\tau_E \notin \mathcal{I}_S$ that does not intersect S . Then $f_{\xi}(s \mid \tau_E) = 0 \ \forall \ s \geq 0$, meaning ξ is non-increasing along τ_E (since no point in S can appear along r_{τ_E}).
Next, consider a ray $\tau = (p, v) \in \mathcal{I}_S$ that does intersect $q \in S$. We must have $f_{\xi}(s \mid \tau) = 1$ until at least $s_q = \|p - q\|$. After s_q , we either must have $f_{\xi}(s \mid \tau) = 0$ iff there are no more points in S along r_{τ} or $f_{\xi}(s \mid \tau) = 1$ iff additional intersection points exist. In the former case, by the same reason as for τ_E above, $f_{\xi}(s' \mid \tau) = 0$ for all $s \geq s_q$, because of the assumption that no further points exist.
- ($Q_{\xi} \subseteq S$) Consider a point q with a one-to-zero discontinuity (i.e., at q , there exists a v such that $\lim_{\epsilon \downarrow 0} \sup_{s \in (0, \epsilon)} \xi(q - sv, v) = 1 \wedge \lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} \xi(q + sv, v) = 0$). This means that there exists a ray that, when approaching q , has a visibility of one, but has a visibility of zero upon exiting it. By the upper continuity convention, we have $\xi(q, v) = 1$ but $\xi(q + sv, v) = 0$ for any $s > 0$ (recalling also that ξ cannot increase along a ray, and hence it must stay at zero thereafter). Recall that $\xi(\tau) = 1$ iff τ intersects $q_s \in S$, but also $\xi(\tau_E) = 0$ iff τ_E does *not* intersect S . This implies we must have $q \in S$. Thus, any positional discontinuity of a shape-induced field must be an element of the shape itself.
- (IO $_{\xi}$) By the definition of Q_{ξ} , IO $_{\xi}$ can be written as $q \in Q_{\xi} \implies \xi(\tau) = 1 \ \forall \ \tau \in \mathcal{I}_q$, where \mathcal{I}_q is the set of rays that intersect q . Now, suppose that $q \in Q_{\xi}$. But, since $Q_{\xi} \subseteq S$, then $q \in S$. Hence, by the definition of inducement, any ray τ containing q must satisfy $\xi(\tau) = 1$.
- Together, these imply that any shape-induced field is a visibility field.

3) Every Visibility Field is a Shape Indicator:

Theorem G.4: Every Visibility Field is a Shape Indicator

Let ξ be a visibility field. Then there exists a point set S such that ξ is induced by S . Further, we can take the positional discontinuities as an inducer; i.e., $S := Q_{\xi}$ induces ξ .

Proof. Our strategy is to show that Q_{ξ} induces ξ directly.

- We have already shown that Q_{ξ} is a shape in Lemma G.2. Next, we simply need to show that Q_{ξ} induces ξ . By definition, ξ is a BOZ field that satisfies IO $_{\xi}$, BC $_{\xi}$, and DE $_{\xi}$. For any $\tau \in \Gamma$, we need to show that $\xi(\tau) = 1 \iff \tau \in \mathcal{I}_{Q_{\xi}}$.
1) Suppose $\xi(\tau) = 1$, where $(p, v) = \tau$. Consider $f_{\xi}(s' \mid \tau)$, which is the visibility along the ray r_{τ} . By BC $_{\xi}$, there must exist an s such that $f(s \mid \tau) = 0$ (since any outward ray in $\mathcal{B} \setminus \mathcal{B}_{\epsilon}$ must have a visibility of zero). This means there must exist $s_q \in [0, s]$ such that $q = p + s_q v$ is a positional discontinuity (i.e., $q \in Q_{\xi}$). By DE $_{\xi}$, there can

only be one such one-to-zero discontinuity. But this means that r_{τ} intersects $q \in Q_{\xi}$, so $\tau \in \mathcal{I}_{Q_{\xi}}$.

Notice also that, if $\xi(\tau) = 0$, along the ray r_{τ} , we can never have a discontinuity from zero to one, by DE $_{\xi}$. Hence, $\xi(\tau) = 0$ implies r_{τ} cannot intersect any point in Q_{τ} as well.

- 2) Suppose $\tau = (p, v) \in \mathcal{I}_{Q_{\xi}}$. This means that there exists $q \in Q_{\xi}$ along the ray r_{τ} (i.e., there exists s such that $q = p + sv \in Q_{\xi}$). By the definition of Q_{ξ} , at q , there must exist at least one v such that that a one-to-zero discontinuity in ξ occurs along v at q . Then, the implication of IO $_{\xi}$ is satisfied, so we must have that $\xi(\tau) = 1$ (since $\tau \in \mathcal{I}_q$).

Together, these show that $\xi(\tau) = 1 \iff \tau \in \mathcal{I}_{Q_{\xi}}$, meaning Q_{ξ} induces ξ , as required. \blacksquare

4) Minimal Characterization of Visibility Fields:

Corollary G.1: Minimal Characterization of Visibility Fields

Consider the set of shapes $\mathfrak{S}[\xi] = \{S \mid S \subseteq \mathcal{B} \text{ induces } \xi\}$ that all induce the same field ξ . I.e., $\mathfrak{S}[\xi]$ is an equivalence class of inducing shapes. Then, $Q_{\xi} \in \mathfrak{S}[\xi]$ is the smallest closed point set among all such inducers.

Proof. We have already shown that Q_{ξ} induces ξ (i.e., $Q_{\xi} \in \mathfrak{S}[\xi]$) in Theorem G.4. However, we have also shown that *any* S -induced field ξ has $Q_{\xi} \subseteq S$ in Theorem G.3. This means Q_{ξ} is at *most* as large as the smallest member of $\mathfrak{S}[\xi]$. Hence, it is the smallest such inducing point set. \blacksquare

5) Directly Lit Points Minimally Characterize Visibility Fields:

Corollary G.2: Directly Lit Points Minimally Characterize Visibility Fields

Given any shape S in the equivalence class $\mathfrak{S}[\xi]$ of inducers for the field ξ , the set of DLPs $\mathcal{D}_{\ell}(S)$ is a *minimal* inducer of ξ and furthermore $\mathcal{D}_{\ell}(S) = Q_{\xi}$.

Proof. Given Cor. G.1, we simply need to show that $Q_{\xi} = \mathcal{D}_{\ell}(S)$.

- Let S be a shape and ξ its induced field. By definition, both Q_{ξ} and $\mathcal{D}_{\ell}(S)$ are closed sets.
(i) Consider $q \in Q_{\xi}$ (and since $Q_{\xi} \subseteq S$, we also have $q \in S$). This means there exists a v such that $\xi(p, v) = 1$ (by upper-continuity) but $\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} \xi(q + sv, v) = 0$. Since by this latter zero value, by DE $_{\xi}$, there can be no $\tilde{q} \in Q_{\xi}$ between q and the boundary along v . There also cannot be any $s > 0$ such that $q + sv \in S$ (else ξ would be one). Let $b \in \partial \mathcal{B}$ be the intersection point of $r_{(q, v)}$ with the boundary. This means that the first intersection with S , from the ray $r_{(b, -v)}$, is q . Hence q is directly lit; i.e., $q \in Q_{\xi} \implies q \in \mathcal{D}_{\ell}(S)$.
(ii) Consider $q \in \mathcal{D}_{\ell}(S)$, with $b \in \partial \mathcal{B}$ being a boundary point from which q can be observed. Let $v := \text{Unit}(q - b)$

point from b to q . We make the same argument as in (i), but effectively in reverse. By the definition of a DLP, $q \in S$ and there can be no other $q' \in S$ between q and b . Hence, $\xi(q, -v) = 1$ (notated via upper-continuity), but $\lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} \xi(q - sv, -v) = 0$. Thus, $q \in Q_\xi$ is a positional discontinuity; i.e., $q \in \mathcal{D}_\ell(S) \implies q \in Q_\xi$. Together, these show that $\mathcal{D}_\ell(S) = Q_\xi$. Since Q_ξ is a minimal inducer, so too is $\mathcal{D}_\ell(S)$. \blacksquare

C. Proofs for Directed Distance Fields (DDFs)

1) Fundamental Point Sets:

Lemma G.3: Fundamental Point Sets

Let (ξ, d) be a DDF. Then $Q_\xi \subseteq Q_{d,\xi} \subseteq Q_d$.

Proof. The second relation is true by definition, so we focus on the first one. Consider a location with a visibility discontinuity, $q \in Q_\xi$. By the definition of Q_ξ (and the upper continuity convention), we may write $\xi(q, v) = 1$ for some $v \in \mathbb{S}^2$. Consider the visible ray $(q, v) = \tau \in \mathcal{R}_V[\xi]$, which starts at q and goes along v (so that $\xi(\tau) = 1$). By property 4 of the DDF (Def. G.25), there must be a $\tilde{q} \in Q_{d,\xi}$ along r_τ (i.e., $\exists s \geq 0$ s.t. $r_\tau(s) = \tilde{q}$). But, since $q \in Q_\xi$, we must have that $\xi(q + sv, v) = 0$ for any $s > 0$. That is, there is *no* locally visible zero along $r_\tau(s)$, for any $s > 0$. Hence, we must have a zero at $s = 0$, meaning $q \in Q_{d,\xi}$; i.e., any $q \in Q_\xi$ also satisfies $q \in Q_{d,\xi}$, as required. \blacksquare

2) DDFs Map to Locally Visible Depth Zeroes:

Lemma G.4: DDFs Map to Locally Visible Depth Zeroes

Let (ξ, d) be a DDF. Then $Q_{d,\xi} = \mathcal{D}_{d,\xi}$. Further:

$$\forall \tau = (p, v) \in \mathcal{R}_V[\xi] : d(\tau) = \min_{\tilde{q} \in [\mathcal{D}_{d,\xi}]_\tau} \|\tilde{q} - p\| \text{ and } q(\tau) \in Q_{d,\xi}, \quad (95)$$

where $[\mathcal{D}_{d,\xi}]_\tau$ is the set of points in $\mathcal{D}_{d,\xi}$ that intersect r_τ and q is the q -mapping of d .

Proof. First, we show that locally visible and directly visible depth zeroes are equivalent; i.e., $Q_{d,\xi} = \mathcal{D}_{d,\xi}$. The rest of the lemma follows straightforwardly.

- Consider $q \in Q_{d,\xi}$. Then, $d(q, v) = 0$ and $\xi(q, v) = 1$. This means $\tau = (q, v) \in \mathcal{R}_V[\xi]$ and $q(q, v) = q + 0v = q$, which implies $q \in \mathcal{D}_{d,\xi}$. Consider $q \in \mathcal{D}_{d,\xi}$. Then there must exist $\tau = (p, v) \in \mathcal{R}_V[\xi]$ such that $q = p + d(p, v)v$ (note that this relation ensures $d(\tau) = \|q - p\|$). Along τ , by $\text{DE}_{d,\xi}$, we have that $f_d(s \mid \tau)$ must be linearly decreasing until $s = \|q - p\|$, at which point $d(p + sv, v) = d(q, v) = 0$. We therefore only need to show that $\xi(q, v) = 1$. We know that $\xi(p, v) = 1$. However, we have already shown that $Q_\xi \subseteq Q_{d,\xi}$ in Lemma G.3. Thus, we can only have a one-to-zero discontinuity at a zero of d . (We cannot change from zero to one because ξ satisfies DE_ξ). Thus, there are only two options: (i) there is a discontinuity in ξ

at q , but this exactly means $q \in Q_{d,\xi}$ (and $\xi(q, v) = 1$ by upper continuity), or (ii) there is no discontinuity between p and q , in which case $\xi = 1$ throughout the ray (i.e., $f_\xi(s \mid (p, v)) = 1 \forall s \in [0, d(p, v)]$). In either case, $\xi(q, v) = 1$ as required.

Hence, $Q_{d,\xi} = \mathcal{D}_{d,\xi}$.

- Consider a visible ray $\tau = (p, v) \in \mathcal{R}_V[\xi]$. We want to show that d computes the distance to the closest member of $Q_{d,\xi}$ (equivalently, $\mathcal{D}_{d,\xi}$). Since (ξ, d) is a DDF and $\tau \in \mathcal{R}_V[\xi]$, there must be a closest zero, $q \in Q_{d,\xi}$, along τ (by requirement four of Def. G.25). By the same logic as the previous part of the proof, we must have $\xi = 1$ along the ray until at least q (i.e., a one-to-zero discontinuity cannot occur until q , at the closest). By definition, $d(q, v) = 0$. Thus, along the ray (where $\xi = 1$), we must have that $\text{DE}_{d,\xi}$ holds, which ensures that $d(p, v) = \|q - p\|$ at p . Thus, d computes the distance to the closest member of $Q_{d,\xi} = \mathcal{D}_{d,\xi}$, as required.
- Since $d(p, v) = \min_{q \in [Q_{d,\xi}]_\tau} \|p - q\|$, where $\tau = (p, v) \in \mathcal{R}_V[\xi]$, we have that $q(\tau) = p + d(p, v)v = p + (\min_{q \in [Q_{d,\xi}]_\tau} \|p - q\|)v = p + \|q^* - p\|v = q^* \in Q_{d,\xi}$, where $q^* = \arg \min_{q \in [Q_{d,\xi}]_\tau} \|p - q\|$ and using the fact that q^* lies on r_τ . \blacksquare

3) Induced Field Pairs are Directed Distance Fields:

Theorem G.5: Induced Field Pairs are Directed Distance Fields

Let S be a shape and $F = (\xi, d)$ be an S -induced field pair. Then F is a Directed Distance Field, which satisfies $Q_{d,\xi} = S$.

Proof. We need to show that F satisfies the properties of a DDF.

- ξ is necessarily a visibility field, because, as shown in Theorem G.3, any shape-induced binary field is a visibility field.
- The induced d is defined in the same way as a simple DDF, except that it is unconstrained on non-visible rays. On visible rays, the same properties therefore apply (recalling that we have previously shown that shape-induced depth fields are simple DDFs; see Theorem G.1). However, restricting IO_d , BC_d , and DE_d to $\mathcal{R}_V[\xi]$ is exactly the definition of ξ -coherence, as required.
- Next, we show that $Q_{d,\xi} = S$.
 - Let $q \in Q_{d,\xi}$ be a locally visible depth zero. We know $\xi(q, v) = 1$ and $d(q, v) = 0$, for some $v \in \mathbb{S}^2$. But for a visible ray $(q, v) = \tau \in \mathcal{R}_V[\xi]$, the only way $f_d(s \mid \tau)$ can approach zero is if $q \in S$, since d is induced and thus constrained along τ . Hence, $q \in Q_{d,\xi}$ implies $q \in S$.
 - Consider $q \in S$. By the definition of induced field pair, $\xi(q, v) = 1$ and $d(q, v) = 0$, for any $v \in \mathbb{S}^2$. This immediately implies $q \in Q_{d,\xi}$.
 Together, these show that $Q_{d,\xi} = S$ for induced field pairs.

- 4) Next, let $q \in Q_{d,\xi}$ and consider \mathcal{I}_q , the set of rays that intersect q . We need to show that every $\tau \in \mathcal{I}_q$ is visible (i.e., $\tau \in \mathcal{R}_V[\xi]$). But $q \in S$ as well, so (since ξ is induced by S), we must have that $\xi(p, v) = 1$ for any $(p, v) \in \mathcal{I}_q$.
- 5) Finally, consider $\tau \in \mathcal{R}_V[\xi]$. But ξ is induced by S , so there must exist some $q \in S$ such that $\tau \in \mathcal{I}_q$. But since $S = Q_{d,\xi}$, $q \in Q_{d,\xi}$ as well. This implies $\mathcal{B}_\tau \cap Q_{d,\xi} \neq \emptyset$, as required.

■

4) *Every Directed Distance Field is a Shape Representation:*

Theorem G.6: Every Directed Distance Field is a Shape Representation

Let $F = (\xi, d)$ be a DDF. Then F must be view consistent (i.e., there exists a shape S such that F is an S -induced field pair). Further, $Q_{d,\xi}$ is a shape that induces F .

Proof. Our strategy is to directly show that $Q_{d,\xi}$ is a shape that induces F .

- We first note that $Q_{d,\xi}$ is a shape. Since $Q_{d,\xi} \subseteq Q_d$ and $Q_d \subseteq \mathcal{B}_\varepsilon$ is bounded, so too is $Q_{d,\xi}$. The closure operator ensures that $Q_{d,\xi}$ is a closed set by definition. Hence, $Q_{d,\xi}$ is a compact point set and therefore a shape.
- We first show that ξ is induced by $Q_{d,\xi}$. We need to show $\xi(\tau) = 1 \iff \tau \in \mathcal{I}_{Q_{d,\xi}}$, where $\mathcal{I}_{Q_{d,\xi}} \subseteq \Gamma$ is the set of rays starting within \mathcal{B} that intersect $Q_{d,\xi}$.
Let $\tau \in \mathcal{R}_V[\xi]$. Since (ξ, d) is a DDF, we have that r_τ must intersect $Q_{d,\xi}$, by requirement four of Def. G.25. In the other direction, suppose $\tau \in \mathcal{I}_{Q_{d,\xi}}$. But by requirement three (of Def. G.25), every locally visible depth zero is isotropically opaque with respect to ξ . Hence, $\xi(\tau) = 1$, as required.
Hence ξ is induced by $Q_{d,\xi}$, and uniquely so. This is the first element of the field pair.
- Next, we need to show that d is induced by $Q_{d,\xi}$ on visible rays; i.e., $d(\tau) = \min_{q \in [Q_{d,\xi}]_\tau} \|q - p\|$ for any $\tau \in \mathcal{R}_V[\xi]$. However, using the directly visible depth zeroes, we have already shown this in Lemma G.4. Since $\mathcal{D}_{d,\xi} = Q_{d,\xi}$, we have that $d(\tau) = \min_{q \in [\mathcal{D}_{d,\xi}]_\tau} \|q - p\| = \min_{q \in [Q_{d,\xi}]_\tau} \|q - p\|$. This means that any d induced by $Q_{d,\xi}$ matches d on visible rays; hence, any induced pair (ξ, \tilde{d}) is equivalent to (ξ, d) .
- Together, these show that any DDF actually forms a shape-induced field pair, via the locally visible zeros $Q_{d,\xi}$. In other words, any DDF is a representation of the shape $Q_{d,\xi}$.

■

D. Convergence Example

We demonstrate the need for defining positional zeroes (Q_d) with

$$\lim_{\epsilon \downarrow 0} \inf_{\substack{s \in (0, \epsilon) \\ v \in \mathbb{S}^2}} d(q - sv, v) = 0, \quad (101)$$

rather than the simpler form

$$\exists v \text{ s.t. } \lim_{\epsilon \downarrow 0} \inf_{s \in (0, \epsilon)} d(q - sv, v) = 0. \quad (102)$$

The latter scenario requires specifying a single fixed direction along which convergence occurs. In the former case, however, such a direction need not necessarily exist (e.g., via a sequence where d converges to zero, but there is no fixed direction along which the converge occurs). This is useful to ensure that the positional zeros Q_d always form a closed set.

Consider the following 2D point set S :

$$p_k = \rho^k R(\theta_k) [1, 0]^T \quad (103)$$

where $k \in \{1, 2, \dots\}$, $\rho < 1$, $R(\theta)$ is a 2D rotation matrix with angle θ , and $\theta_k = \pi/(2k)$. Note that p_k converges to $q = [0, 0]^T$ (but $q \notin \{p_k\}_k$), while θ_k converges to 0. Let $v_k = R(\theta_k) [1, 0]$. As k increases, v_k converges to $u = [1, 0]^T$.

Consider the simple DDF d , which has been induced by $S = \{p_k\}_k$.¹⁹ Is $q \in Q_d$? Note that $p_k \in Q_d \forall k$, so if not, then Q_d will not be a compact set (shape). Further, note that d necessarily satisfies IO_d , in addition to BC_d and DE_d . The former means that if a point is not a zero from one direction, then it cannot be a zero from any direction.

First, consider the definition using Eq. 101. In this case, $q \in Q_d$, because (similar to the proof in §H-A2 above) for any $\epsilon > 0$, we can always choose a p_k s.t. $\|p_k - q\| < \epsilon$ and let $\tilde{v} := \text{Unit}(p_k - q)$. This converges to zero, as ϵ shrinks, in the limit. Hence the equation is satisfied. However, note that \tilde{v} changes continuously; there is not one single direction that is used.

Next, consider the definition using Eq. 102. But notice that $\tau = (q, u)$ is an empty ray, since $q \notin \{p_k\}_k$. Thus, $d(q, u) > 0$ (by IO_d), meaning $q \notin Q_d$. This situation is what we wish to avoid, since it permits the set of positional zeroes to not be closed.

REFERENCES

- [1] L. G. Roberts, “Machine perception of three-dimensional solids,” Ph.D. dissertation, Massachusetts Institute of Technology, 1963. 1
- [2] B. G. Baumgart, “Geometric modeling for computer vision,” Ph.D. dissertation, Stanford University, 1974. 1, 19
- [3] A. Yuille and D. Kersten, “Vision as Bayesian inference: Analysis by synthesis?” *Trends in Cognitive Sciences*, 2006. 1
- [4] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “NeRF: Representing scenes as neural radiance fields for view synthesis,” in *ECCV*, 2020. 1, 3, 4, 10
- [5] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger, “Occupancy networks: Learning 3D reconstruction in function space,” in *CVPR*, 2019. 1, 3
- [6] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, “DeepSDF: Learning continuous signed distance functions for shape representation,” in *CVPR*, 2019. 1, 3, 7
- [7] S. Liu, Y. Zhang, S. Peng, B. Shi, M. Pollefeys, and Z. Cui, “DIST: Rendering deep implicit signed distance function with differentiable sphere tracing,” in *CVPR*, 2020. 1, 3, 4, 7
- [8] M. Niemeyer, L. Mescheder, M. Oechsle, and A. Geiger, “Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision,” in *CVPR*, 2020. 1, 3

¹⁹Note that we define shapes to be compact sets but here S is not closed. However, that is not an issue for this example because we merely want to show that *some* simple DDF can have an open set as its positional zeroes. I.e., we do not need to obtain it by inducement.

- [9] Y. Jiang, D. Ji, Z. Han, and M. Zwicker, “SDFDiff: Differentiable rendering of signed distance fields for 3D shape optimization,” in *CVPR*, 2020. 1, 3
- [10] D. Vicini, S. Speierer, and W. Jakob, “Differentiable signed distance function rendering,” *ACM TOG*, 2022. 1
- [11] A. Tewari, J. Thies, B. Mildenhall, P. Srinivasan, E. Tretschk, W. Yifan, C. Lassner, V. Sitzmann, R. Martin-Brualla, S. Lombardi, T. Simon, C. Theobalt, M. Nießner, J. T. Barron, G. Wetzstein, M. Zollhöfer, and V. Golyanik, “Advances in neural rendering,” *Computer Graphics Forum*, 2022. 1
- [12] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick, “Zero-1-to-3: Zero-shot one image to 3D object,” *ICCV*, 2023. 1
- [13] A. Bhattad, D. McKee, D. Hoiem, and D. Forsyth, “StyleGAN knows normal, depth, albedo, and more,” *NeurIPS*, 2024. 1
- [14] A. Tewari, T. Yin, G. Cazenavette, S. Rezchikov, J. Tenenbaum, F. Durand, B. Freeman, and V. Sitzmann, “Diffusion with forward models: Solving stochastic inverse problems without direct supervision,” *NeurIPS*, 2024. 1
- [15] N. Müller, Y. Siddiqui, L. Porzi, S. R. Buló, P. Kotschieder, and M. Nießner, “DiffRF: Rendering-guided 3D radiance field diffusion,” in *CVPR*, 2023. 1
- [16] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich, “Score Jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation,” in *CVPR*, 2023. 1
- [17] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, “DreamFusion: Text-to-3D using 2D diffusion,” *arXiv:2209.14988*, 2022. 1
- [18] J. Pan, X. Han, W. Chen, J. Tang, and K. Jia, “Deep mesh reconstruction from single RGB images via topology modification networks,” in *ICCV*, 2019. 1
- [19] J. Tang, X. Han, J. Pan, K. Jia, and X. Tong, “A skeleton-bridged deep learning approach for generating meshes of complex topologies from single RGB images,” in *CVPR*, 2019. 1
- [20] H. Kato, Y. Ushiku, and T. Harada, “Neural 3D mesh renderer,” in *CVPR*, 2018. 1, 9
- [21] S. Liu, T. Li, W. Chen, and H. Li, “Soft rasterizer: A differentiable renderer for image-based 3D reasoning,” in *ICCV*, 2019. 1, 4, 6
- [22] V. Sitzmann, M. Zollhöfer, and G. Wetzstein, “Scene representation networks: Continuous 3D-structure-aware neural scene representations,” in *NeurIPS*, 2019. 1, 3, 4
- [23] T. Aumentado-Armstrong, S. Tsogkas, S. Dickinson, and A. D. Jepson, “Representing 3D shapes with probabilistic directed distance fields,” in *CVPR*, 2022. 2, 3
- [24] A. Rosenfeld and J. L. Pfaltz, “Distance functions on digital pictures,” *Pattern recognition*, vol. 1, no. 1, pp. 33–61, 1968. 3
- [25] J. Chibane, A. Mir, and G. Pons-Moll, “Neural unsigned distance fields for implicit function learning,” *NeurIPS*, 2020. 3, 7
- [26] R. Venkatesh, S. Sharma, A. Ghosh, L. Jeni, and M. Singh, “DUDE: Deep unsigned distance embeddings for hi-fidelity representation of complex 3D surfaces,” *arXiv:2011.02570*, 2020. 3, 7
- [27] R. Cai, G. Yang, H. Averbuch-Elor, Z. Hao, S. Belongie, N. Snavely, and B. Hariharan, “Learning gradient fields for shape generation,” in *ECCV*, 2020. 3
- [28] R. Venkatesh, T. Karmali, S. Sharma, A. Ghosh, R. V. Babu, L. A. Jeni, and M. Singh, “Deep implicit surface point prediction networks,” in *ICCV*, 2021. 3
- [29] S. Liu, S. Saito, W. Chen, and H. Li, “Learning to infer implicit surfaces without 3D supervision,” in *NeurIPS*, 2019. 3
- [30] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler, “Neural geometric level of detail: Real-time rendering with implicit 3D shapes,” in *CVPR*, 2021. 3
- [31] J. Zhang, Y. Yao, and L. Quan, “Learning signed distance field for multi-view surface reconstruction,” in *ICCV*, 2021. 3
- [32] P. Kellnhofer, L. Jebe, A. Jones, R. Spicer, K. Pulli, and G. Wetzstein, “Neural lumigraph rendering,” in *CVPR*, 2021. 3
- [33] M. Oechsle, S. Peng, and A. Geiger, “UNISURF: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction,” *ICCV*, 2021. 3
- [34] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman, “Volume rendering of neural implicit surfaces,” *arXiv:2106.12052*, 2021. 3
- [35] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang, “NeuS: Learning neural implicit surfaces by volume rendering for multi-view reconstruction,” *NeurIPS*, 2021. 3
- [36] J. T. Barron, B. Mildenhall, D. Verbin, P. P. Srinivasan, and P. Hedman, “Zip-NeRF: Anti-aliased grid-based neural radiance fields,” *arXiv:2304.06706*, 2023. 3
- [37] S. J. Garbin, M. Kowalski, M. Johnson, J. Shotton, and J. Valentin, “FastNeRF: High-fidelity neural rendering at 200fps,” *arXiv:2103.10380*, 2021. 3
- [38] P. Hedman, P. P. Srinivasan, B. Mildenhall, J. T. Barron, and P. Debevec, “Baking neural radiance fields for real-time view synthesis,” *ICCV*, 2021. 3
- [39] C. Reiser, S. Peng, Y. Liao, and A. Geiger, “KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs,” *ICCV*, 2021. 3
- [40] A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and A. Kanazawa, “Plenotrees for real-time rendering of neural radiance fields,” *ICCV*, 2021. 3
- [41] T. Aumentado-Armstrong, A. Mirzaei, M. A. Brubaker, J. Kelly, A. Levinstein, K. G. Derpanis, and I. Gilitschenski, “Reconstructive latent-space neural radiance fields for efficient 3D scene representations,” *arXiv:2310.17880*, 2023. 3
- [42] Z. Wan, C. Richardt, A. Božič, C. Li, V. Rengarajan, S. Nam, X. Xiang, T. Li, B. Zhu, R. Ranjan *et al.*, “Learning neural duplex radiance fields for real-time view synthesis,” in *CVPR*, 2023. 3
- [43] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3D gaussian splatting for real-time radiance field rendering,” *ACM TOG*, 2023. 3
- [44] X. Zhang, P. P. Srinivasan, B. Deng, P. Debevec, W. T. Freeman, and J. T. Barron, “NeRFactor: Neural factorization of shape and reflectance under an unknown illumination,” *ACM TOG*, 2021. 3, 12
- [45] V. Sitzmann, S. Rezchikov, B. Freeman, J. Tenenbaum, and F. Durand, “Light field networks: Neural scene representations with single-evaluation rendering,” *NeurIPS*, 2021. 3, 6, 11
- [46] Z. Li, L. Song, C. Liu, J. Yuan, and Y. Xu, “NeuLF: Efficient novel view synthesis with neural 4D light field,” *arXiv:2105.07112*, 2021. 3
- [47] B. Attal, J.-B. Huang, M. Zollhöfer, J. Kopf, and C. Kim, “Learning neural light fields with ray-space embedding,” in *CVPR*, 2022. 3
- [48] E. Zobeidi and N. Atanasov, “A deep signed directional distance function for object shape representation,” *arXiv:2107.11024*, 2021. 3, 6, 11
- [49] T. Houchens, C.-Y. Lu, S. Duggal, R. Fu, and S. Sridhar, “NeuralODF: Learning omnidirectional distance fields for 3D shape representation,” *arXiv:2206.05837*, 2022. 3
- [50] B. Y. Feng, Y. Zhang, D. Tang, R. Du, and A. Varshney, “PRIF: Primary ray-based implicit function,” in *ECCV*, 2022. 3
- [51] X. Yang, G. Lin, Z. Chen, and L. Zhou, “Neural vector fields: Implicit representation by explicit learning,” in *CVPR*, 2023. 3
- [52] N. Kulkarni, J. Johnson, and D. F. Fouhey, “Directed ray distance functions for 3D scene reconstruction,” in *ECCV*, 2022. 3
- [53] J. Xi, Y. Shi, Y. Wang, Y. Guo, and K. Xu, “RayMVSNet: Learning ray-based 1D implicit fields for accurate multi-view stereo,” *CVPR*, 2022. 3
- [54] C. Zhang, Y. Di, R. Zhang, G. Zhai, F. Manhardt, F. Tombari, and X. Ji, “DDF-HO: Hand-held object reconstruction via conditional directed distance field,” *NeurIPS*, 2024. 3
- [55] T. Yenamandra, A. Tewari, N. Yang, F. Bernard, C. Theobalt, and D. Cremers, “FiRe: Fast inverse rendering using directional and signed distance functions,” in *WACV*, 2024. 3
- [56] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee, “Perspective transformer nets: Learning single-view 3D object reconstruction without 3D supervision,” *NeurIPS*, 2016. 4, 3
- [57] J. Wu, Y. Wang, T. Xue, X. Sun, B. Freeman, and J. Tenenbaum, “MarrNet: 3D shape reconstruction via 2.5D sketches,” *NeurIPS*, 2017. 4, 3
- [58] S. Tulsiani, T. Zhou, A. A. Efros, and J. Malik, “Multi-view supervision for single-view reconstruction via differentiable ray consistency,” in *CVPR*, 2017. 4, 9, 3
- [59] T. H. Nguyen-Phuoc, C. Li, S. Balaban, and Y. Yang, “RenderNet: A deep convolutional network for differentiable rendering from 3D shapes,” in *NeurIPS*, 2018. 4, 3
- [60] A. Gropp, L. Yariv, N. Haim, M. Atzmon, and Y. Lipman, “Implicit geometric regularization for learning shapes,” *ICML*, 2020. 5, 16
- [61] C.-H. Lin, C. Wang, and S. Lucey, “SDF-SRN: Learning signed distance 3D object reconstruction from static images,” *NeurIPS*, 2020. 5, 16
- [62] M. Yang, Y. Wen, W. Chen, Y. Chen, and K. Jia, “Deep optimized priors for 3D shape modeling and reconstruction,” in *CVPR*, 2021. 5, 16
- [63] S. P. Bangaru, M. Gharbi, F. Luan, T.-M. Li, K. Sunkavalli, M. Hasan, S. Bi, Z. Xu, G. Bernstein, and F. Durand, “Differentiable rendering of neural sdfs through reparameterization,” in *SIGGRAPH Asia Conference Papers*, 2022. 5, 16
- [64] Z. Yu, S. Peng, M. Niemeyer, T. Sattler, and A. Geiger, “MonoSDF: Exploring monocular geometric cues for neural implicit surface reconstruction,” *NeurIPS*, 2022. 5, 16

- [65] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," *NeurIPS*, 2020. **6, 7**
- [66] E. R. Chan, M. Monteiro, P. Kellnhofer, J. Wu, and G. Wetzstein, "pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis," in *CVPR*, 2021. **6, 7**
- [67] K. Jo, G. Shim, S. Jung, S. Yang, and J. Choo, "CG-NeRF: Conditional generative neural radiance fields," *arXiv:2112.03517*, 2021. **6**
- [68] J. Gao, W. Chen, T. Xiang, A. Jacobson, M. McGuire, and S. Fidler, "Learning deformable tetrahedral meshes for 3D reconstruction," *NeurIPS*, 2020. **6**
- [69] Y.-T. Liu, L. Wang, J. Yang, W. Chen, X. Meng, B. Yang, and L. Gao, "NeUDF: Learning neural unsigned distance fields with volume rendering," in *CVPR*, 2023. **7**
- [70] N. Wang, Y. Zhang, Z. Li, Y. Fu, W. Liu, and Y.-G. Jiang, "Pixel2Mesh: Generating 3D mesh models from single RGB images," *ECCV*, 2018. **8, 5, 6**
- [71] N. Wang, Y. Zhang, Z. Li, Y. Fu, H. Yu, W. Liu, X. Xue, and Y.-G. Jiang, "Pixel2Mesh: 3D mesh model generation via image guided deformation," *IEEE PAMI*, 2020. **8, 5, 6**
- [72] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese, "3D-R2N2: A unified approach for single and multi-view 3D object reconstruction," in *ECCV*, 2016. **8, 5, 6**
- [73] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016. **8, 7**
- [74] I. Mehta, M. Gharbi, C. Barnes, E. Shechtman, R. Ramamoorthi, and M. Chandraker, "Modulated periodic activations for generalizable local functional representations," in *ICCV*, 2021. **8, 4, 7**
- [75] T. Aumentado-Armstrong, A. Levinstein, S. Tsogkas, K. G. Derpanis, and A. D. Jepson, "Cycle-consistent generative rendering for 2D-3D modality translation," *3DV*, 2020. **9, 8**
- [76] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local nash equilibrium," *NeurIPS*, 2017. **9**
- [77] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An Information-Rich 3D Model Repository," *arXiv:1512.03012*, 2015, <https://shapenet.org/>. **8, 5, 7**
- [78] E. Insafutdinov and A. Dosovitskiy, "Unsupervised learning of shape and pose with differentiable point clouds," *NeurIPS*, 2018. **9**
- [79] D. Shin, C. C. Fowlkes, and D. Hoiem, "Pixels, voxels, and views: A study of shape representations for single view 3D object shape prediction," in *CVPR*, 2018. **9**
- [80] M. Tatarchenko, S. R. Richter, R. Ranftl, Z. Li, V. Koltun, and T. Brox, "What do single-view 3D reconstruction networks learn?" *CVPR*, 2019. **9**
- [81] B. Kaya and R. Timofte, "Self-supervised 2D image to 3D shape translation with disentangled representations," in *3DV*, 2020. **9, 8**
- [82] Y. Miyauchi, Y. Sugano, and Y. Matsushita, "Shape-conditioned image generation by learning latent appearance representation from unpaired data," in *ACCV*, 2018. **9, 8**
- [83] J.-Y. Zhu, Z. Zhang, C. Zhang, J. Wu, A. Torralba, J. Tenenbaum, and B. Freeman, "Visual object networks: Image generation with disentangled 3D representations," in *NeurIPS*, 2018. **9**
- [84] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *ICLR*, 2014. **9, 7**
- [85] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, 2014. **9, 7, 8**
- [86] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3D classification and segmentation," *CVPR*, 2017. **9, 7**
- [87] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," *ICCV*, 2017. **9, 7, 8**
- [88] H. Barrow, J. Tenenbaum, A. Hanson, and E. Riseman, "Recovering intrinsic scene characteristics," *Computer Vision Systems*, 1978. **10, 8**
- [89] E. H. Adelson and A. P. Pentland, "The perception of shading and reflectance," *Perception as Bayesian inference*, 1996. **10**
- [90] A. Laurentini, "How many 2D silhouettes does it take to reconstruct a 3D object?" *Computer Vision and Image Understanding*, 1997. **13, 19**
- [91] W. Richards, J. J. Koenderink, and D. D. Hoffman, "Inferring three-dimensional shapes from two-dimensional silhouettes," *Journal of the Optical Society of America A*, 1987. **13, 19**
- [92] K.-M. Cheung, S. Baker, and T. Kanade, "Shape-from-silhouette across time part I: Theory and algorithms," *IJCV*, 2005. **13, 19**
- [93] A. Laurentini, "The visual hull concept for silhouette-based image understanding," *IEEE PAMI*, 1994. **13, 19**
- [94] K. N. Kutulakos and S. M. Seitz, "A theory of shape by space carving," *IJCV*, 2000. **13, 19**
- [95] E. Kreyszig, *Differential Geometry*. University of Toronto Press, 1959, mathematical Expositions, No. 11. Dover edition. **2**
- [96] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *ICLR*, 2015. **3**
- [97] A. Dalmia, "dalmia/siren," <https://doi.org/10.5281/zenodo.3902941>, Jun. 2020, zenodo: 10.5281/zenodo.3902941. <https://github.com/dalmia/siren> (MIT License). **3**
- [98] Stanford Computer Graphics Laboratory, "The Stanford 3D scanning repository," <http://graphics.stanford.edu/data/3Dscanrep/>, accessed: 09/08/21. **3**
- [99] G. Turk and M. Levoy, "Zippered polygon meshes from range images," in *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, 1994, pp. 311–318. **3**
- [100] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 303–312. **3**
- [101] V. Krishnamurthy and M. Levoy, "Fitting smooth surfaces to dense polygon meshes," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 313–324. **3**
- [102] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019. [Online]. Available: <https://openreview.net/forum?id=Bkg6RiCqY7>. **4**
- [103] K. M. Jatavallabhula, E. Smith, J.-F. Lafleche, C. F. Tsang, A. Rozantsev, W. Chen, T. Xiang, R. Lebedev, and S. Fidler, "Kaolin: A PyTorch library for accelerating 3D deep learning research," *arXiv:1911.05063*, 2019, <https://github.com/NVIDIAGameWorks/kaolin> (Apache License). **5**
- [104] T. Groueix, "Pytorch chamfer distance." *Thibault-GROUEIX/ChamferDistancePytorch*, Nov. 2021, <https://github.com/ThibaultGROUEIX/ChamferDistancePytorch> (MIT License). **6**
- [105] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *ICML*, 2014. **7**
- [106] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "Beta-VAE: Learning basic visual concepts with a constrained variational framework," *ICLR*, 2017. **7**
- [107] K. Roth, A. Lucchi, S. Nowozin, and T. Hofmann, "Stabilizing training of generative adversarial networks through regularization," *arXiv:1705.09367*, 2017. **7, 8**
- [108] L. Mescheder, A. Geiger, and S. Nowozin, "Which training methods for GANs do actually converge?" in *ICML*, 2018. **7, 8**
- [109] M. Niemeyer and A. Geiger, "Giraffe: Representing scenes as compositional generative neural feature fields," in *CVPR*, 2021. **7**
- [110] Y. Idelbayev, "Proper ResNet implementation for CIFAR10/CIFAR100 in PyTorch," https://github.com/akamaster/pytorch_resnet_cifar10, accessed: Nov 2021. (BSD-2-clause License). **7, 8**
- [111] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "CycleGAN and Pix2Pix in PyTorch," Nov. 2021, <https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix> (BSD License). **8**
- [112] K. S. Lee and C. Town, "Mimicry: Towards the reproducibility of GAN research," *CVPR Workshop on AI for Content Creation*, 2020, <https://github.com/kwotsin/mimicry> (MIT License). **8**
- [113] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *ICLR*, 2016. **8**
- [114] A. Obukhov, M. Seitzer, P.-W. Wu, S. Zhydenko, J. Kyl, and E. Y.-J. Lin, "High-fidelity performance metrics for generative models in PyTorch," 2020, version: 0.3.0, DOI: 10.5281/zenodo.4957738 (Apache License). [Online]. Available: <https://github.com/toshas/torch-fidelity>. **8**
- [115] J. T. Kajiya, "The rendering equation," in *Proceedings of the 13th annual conference on Computer graphics and interactive techniques*, 1986, pp. 143–150. **9**
- [116] D. S. Immel, M. F. Cohen, and D. P. Greenberg, "A radiosity method for non-diffuse environments," *Acm SIGGRAPH Computer Graphics*, vol. 20, no. 4, pp. 133–142, 1986. **9**
- [117] L. Fascione, J. Hanika, D. Heckenberg, C. Kulla, M. Droske, and J. Schwarzhaupt, "Path tracing in production: part 1: modern path tracing," *ACM SIGGRAPH 2019 Courses*, 2019. **9**
- [118] R. Jain, R. Kasturi, and B. G. Schunck, *Machine vision*. McGraw-Hill New York (ISBN 0-07-032018-7), 1995, vol. 5, chapter 9: Shading. **9, 10, 11**

- [119] S. Hadadan, S. Chen, and M. Zwicker, “Neural radiosity,” *ACM Transactions on Graphics (TOG)*, vol. 40, no. 6, pp. 1–11, 2021. [9](#), [12](#)
- [120] E. Veach and L. Guibas, “Bidirectional estimators for light transport,” in *Photorealistic Rendering Techniques*. Springer, 1995, pp. 145–167. [9](#), [12](#)
- [121] J. Arvo, “Backward ray tracing,” *Developments in ray tracing*, vol. 12, pp. 259–263, 1986, see 1995 addendum. [9](#)
- [122] E. Veach, “Robust Monte Carlo methods for light transport simulation,” Ph.D. dissertation, Stanford University, 1998. [10](#)
- [123] M. Pharr, W. Jakob, and G. Humphreys, *Physically based rendering: From theory to implementation (3rd ed.)*. MIT Press, 2018, (8.2 & 8.4.). [10](#)
- [124] Y.-C. Guo, D. Kang, L. Bao, Y. He, and S.-H. Zhang, “NeRFReN: Neural radiance fields with reflections,” in *CVPR*, 2022. [11](#)
- [125] C. Zhu, R. Wan, and B. Shi, “Neural transmitted radiance fields,” *NeurIPS*, 2022. [11](#)
- [126] B. T. Phong, “Illumination for computer generated pictures,” *Commun. ACM*, vol. 18, no. 6, pp. 311–317, 1975. [Online]. Available: <https://doi.org/10.1145/360825.360839> [11](#)
- [127] R. Montes and C. Urena, “An overview of BRDF models,” *University of Grenada, Technical Report LSI-2012*, 2012. [11](#)
- [128] H. C. Karaimer and M. S. Brown, “A software platform for manipulating the camera imaging pipeline,” in *ECCV*, 2016. [11](#)
- [129] K. Yu, Z. Li, Y. Peng, C. C. Loy, and J. Gu, “ReconfigISP: Reconfigurable camera image processing pipeline,” in *ICCV*, 2021. [11](#)
- [130] M. Oechsle, L. Mescheder, M. Niemeyer, T. Strauss, and A. Geiger, “Texture fields: Learning texture representations in function space,” in *ICCV*, 2019. [12](#)
- [131] J. T. Barron and J. Malik, “Shape, illumination, and reflectance from shading,” *IEEE PAMI*, 2014. [12](#)
- [132] E. Veach and L. J. Guibas, “Metropolis light transport,” in *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, 1997, pp. 65–76. [12](#)
- [133] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, “Neural importance sampling,” *ACM TOG*, vol. 38, no. 5, pp. 1–19, 2019. [12](#)
- [134] E. Jang, S. Gu, and B. Poole, “Categorical reparameterization with Gumbel-softmax,” *arXiv:1611.01144*, 2016. [12](#)
- [135] C. J. Maddison, A. Mnih, and Y. W. Teh, “The concrete distribution: A continuous relaxation of discrete random variables,” *arXiv:1611.00712*, 2016. [12](#)
- [136] W. Grathwohl, D. Choi, Y. Wu, G. Roeder, and D. Duvenaud, “Back-propagation through the void: Optimizing control variates for black-box gradient estimation,” *arXiv:1711.00123*, 2017. [12](#)
- [137] D. Marr, “Visual information processing: The structure and creation of visual representations,” *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, vol. 290, no. 1038, pp. 199–218, 1980. [19](#)
- [138] J. Willats, “Seeing lumps, sticks, and slabs in silhouettes,” *Perception*, 1992. [19](#)
- [139] K. Li, R. Garg, M. Cai, and I. Reid, “Single-view object shape reconstruction using deep shape prior and silhouette,” *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. [19](#)
- [140] A. Laurentini, “How far 3D shapes can be understood from 2D silhouettes,” *IEEE PAMI*, 1995. [19](#)
- [141] —, “Computing the visual hull of solids of revolution,” *Pattern Recognition*, 1999. [19](#)
- [142] S. Petitjean, “A computational geometric approach to visual hulls,” *International Journal of Computational Geometry & Applications*, 1998. [19](#)
- [143] M. Trager, M. Hebert, and J. Ponce, “Consistency of silhouettes and their duals,” in *CVPR*, 2016. [19](#)
- [144] K. N. Kutulakos, “Shape from the light field boundary,” in *CVPR*, 1997. [19](#)
- [145] S. Lazebnik, E. Boyer, and J. Ponce, “On computing exact visual hulls of solids bounded by smooth surfaces,” in *CVPR*, 2001. [19](#)
- [146] S. Lazebnik, Y. Furukawa, and J. Ponce, “Projective visual hulls,” *IJCV*, 2007. [19](#)
- [147] J.-L. Landabaso, M. Pardàs, and J. R. Casas, “Shape from inconsistent silhouette,” *Computer Vision and Image Understanding*, 2008. [19](#)
- [148] L. Yi and D. W. Jacobs, “Efficiently determining silhouette consistency,” in *CVPR*, 2007. [19](#)
- [149] W. Matusik, C. Buehler, R. Raskar, S. J. Gortler, and L. McMillan, “Image-based visual hulls,” in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000. [19](#)
- [150] Y. Furukawa and J. Ponce, “Carved visual hulls for image-based modeling,” *IJCV*, 2009. [19](#)
- [151] J.-S. Franco, M. Lapierre, and E. Boyer, “Visual shapes of silhouette sets,” in *Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT)*, 2006. [19](#)
- [152] A. Bottino and A. Laurentini, “Introducing a new problem: Shape-from-silhouette when the relative positions of the viewpoints is unknown,” *IEEE PAMI*, 2003. [19](#)
- [153] G. Haro, “Shape from silhouette consensus,” *Pattern Recognition*, 2012. [19](#)
- [154] G. K. Cheung, T. Kanade, J.-Y. Bouguet, and M. Holler, “A real time system for robust 3D voxel reconstruction of human motions,” in *CVPR*, 2000. [19](#)
- [155] J. Y. Zheng and F. Kishino, “3D models from contours: further identification of unexposed areas,” in *Proceedings of the International Conference on Pattern Recognition (ICPR)*, 1992. [19](#)
- [156] J. Y. Zheng, “Acquiring 3-d models from sequences of contours,” *IEEE PAMI*, 1994. [19](#)