UNBIASED WARPED-AREA SAMPLING FOR DIFFERENTIABLE RENDERING

UCSD CSE 272

Advanced Image Synthesis

Tzu-Mao Li

with slides from Sai Bangaru

CHALLENGES: EDGE SAMPLING IS HARD!

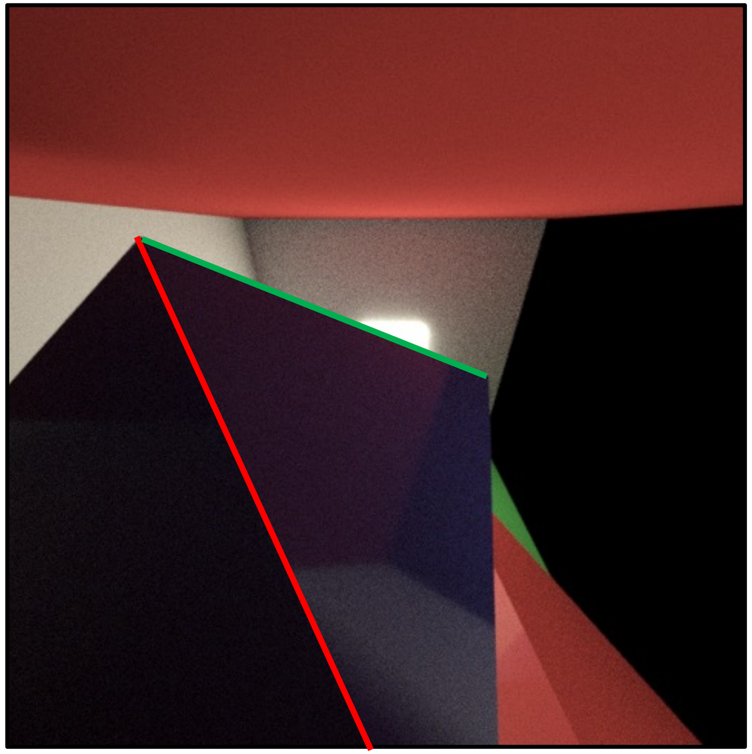

Silhouette classification

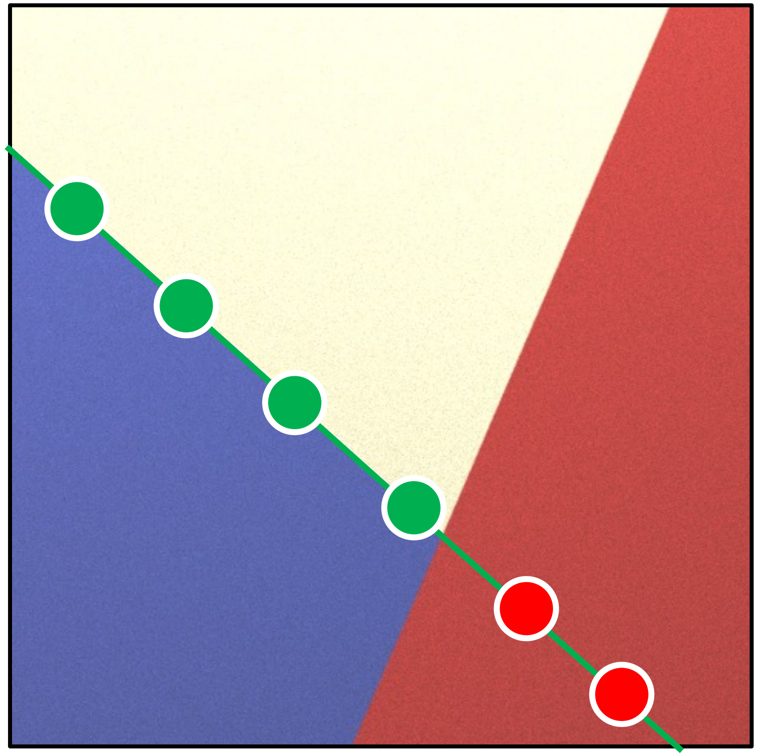

Occlusion

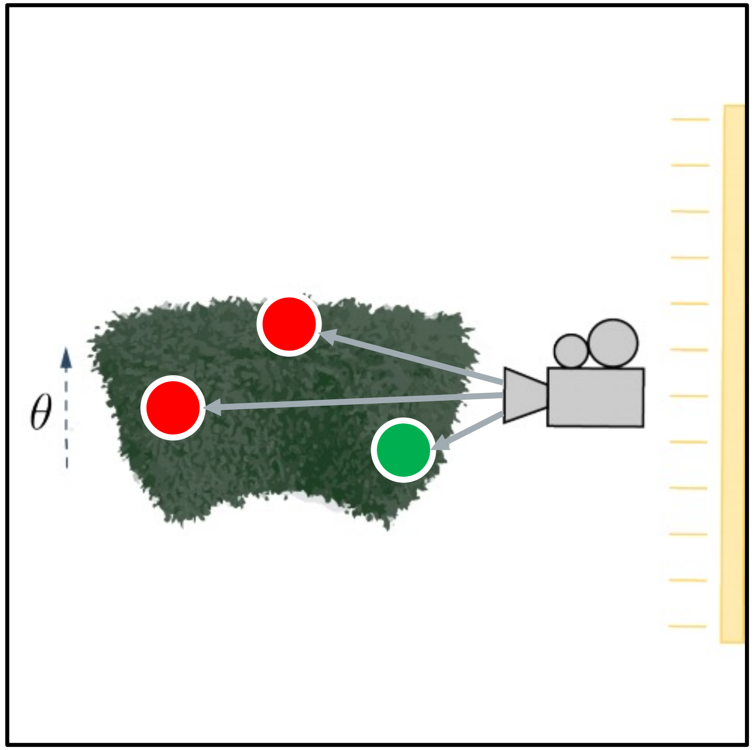

Depth complexity

EDGE SAMPLING HAS TROUBLE WITH SPECULAR REFLECTIONS

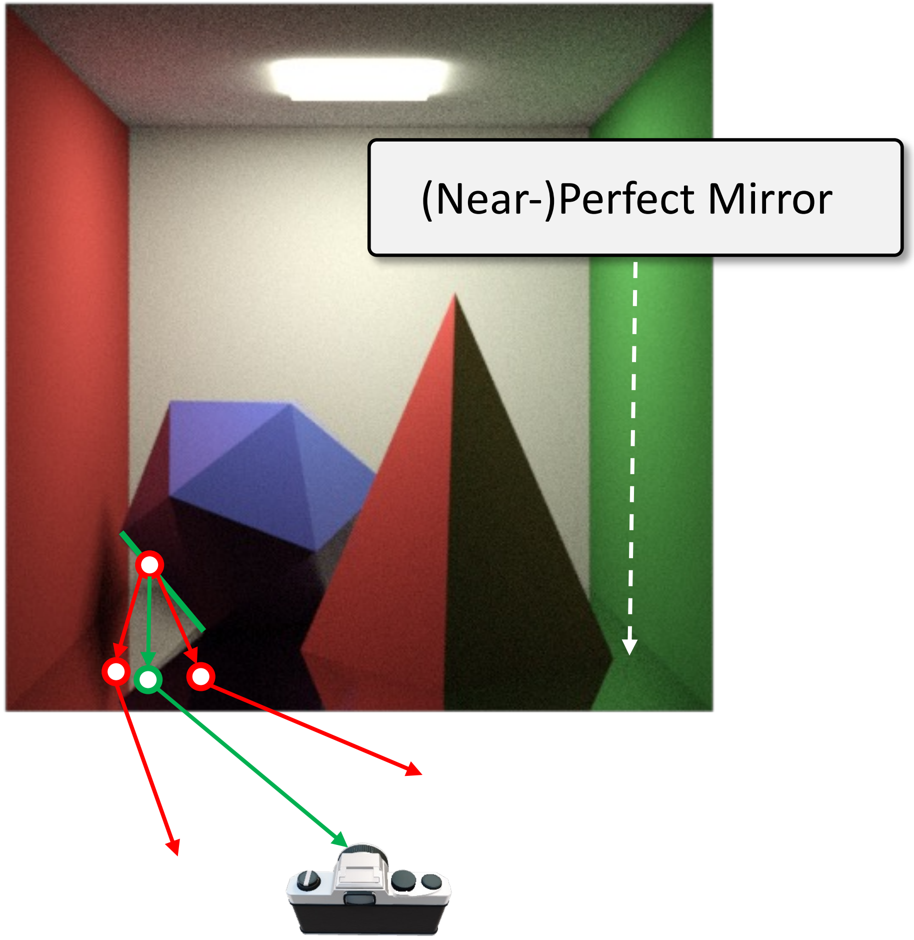

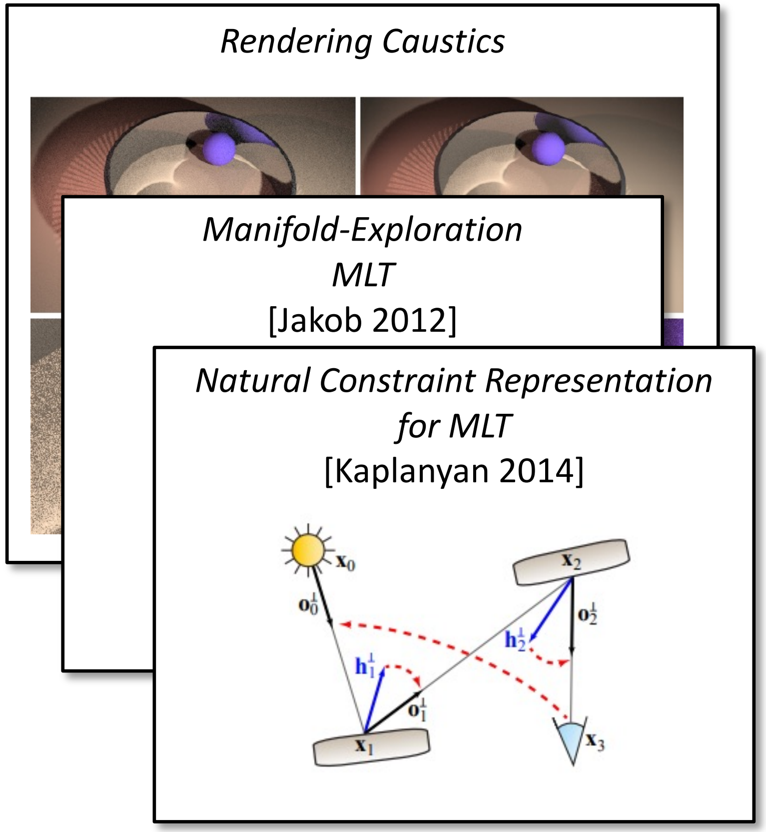

SILHOUETTE EXTRACTION IS DIFFICULT FOR IMPLICIT REPRESENTATIONS

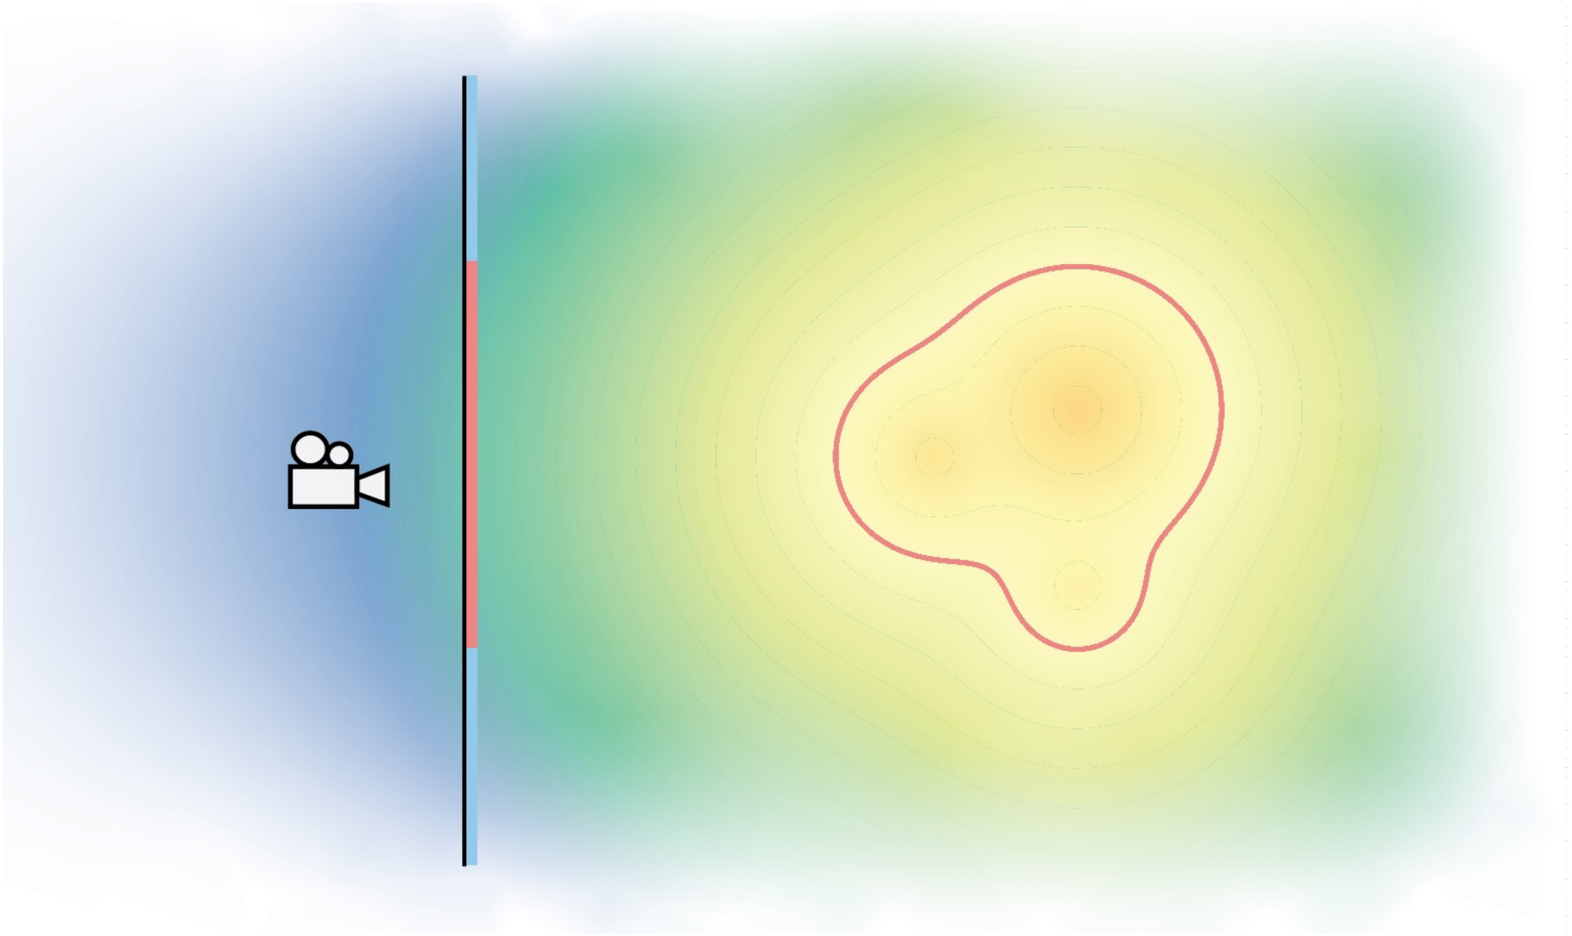

CAN WE DESIGN AN UNBIASED AREA SAMPLING METHOD?

Reparameterizing Discontinuous Integrands for Differentiable Rendering
[Loubet 2019]

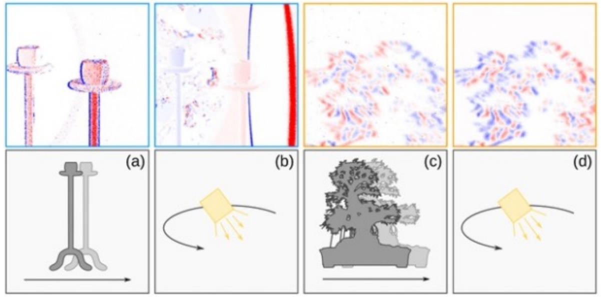

Transform samples with **θ**. Avoids discontinuities.

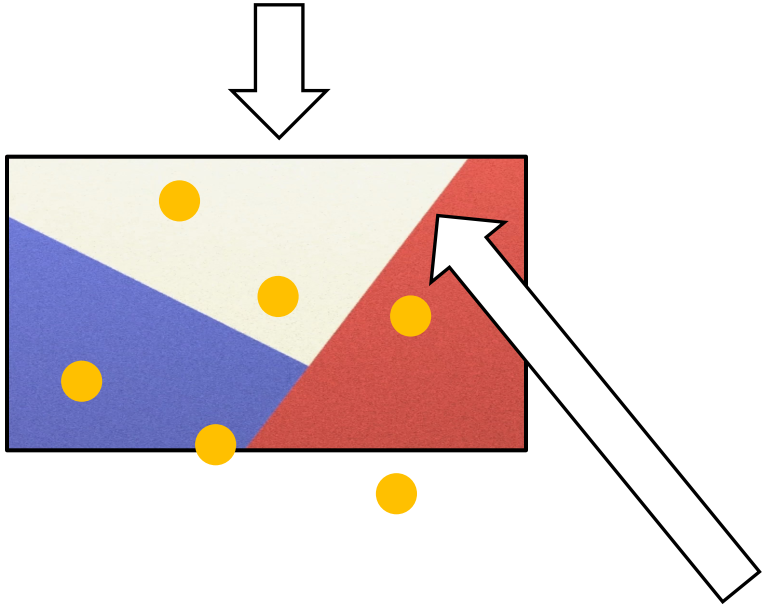

Heuristic Approximation!
May not work for all samples.

OUR APPROACH

THE REYNOLDS TRANSPORT THEOREM

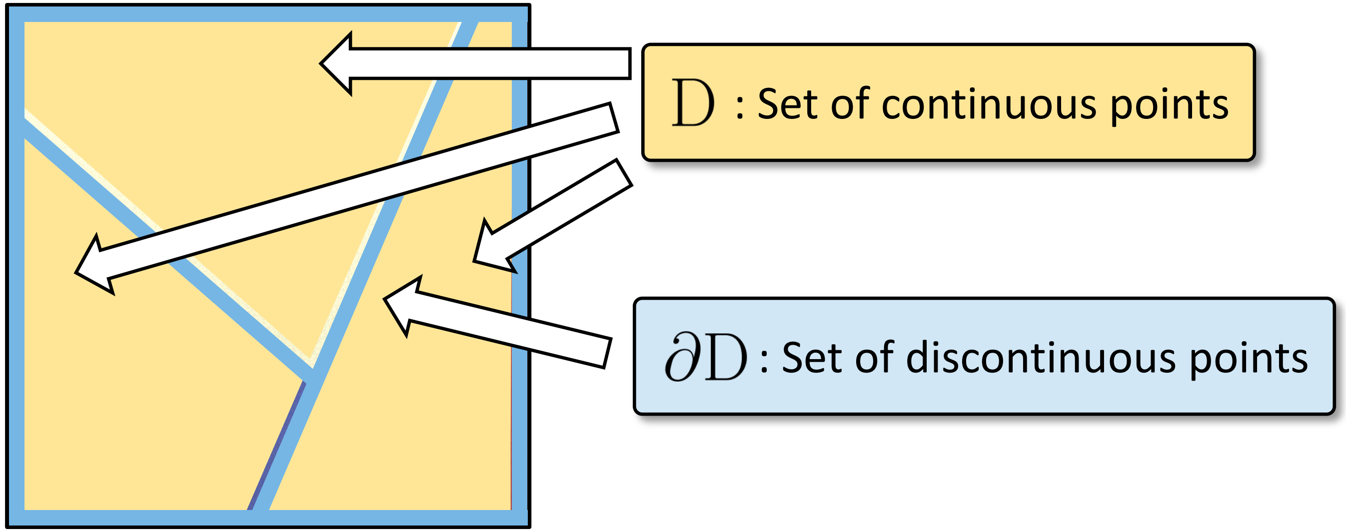

$$\partial_\theta \int_{\mathrm{D}} f = \int_{\mathrm{D}} \partial_\theta f + \int_{\partial\mathrm{D}} f \vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$$

Interior term

Edge term

CONVERTING EDGE SAMPLES TO AREA SAMPLES

$\int_{\partial D} f\vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$ is estimated through edge samples

Goal: Rewrite $\int_{\partial D} f\vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$ into area integral $\int_{D} g$

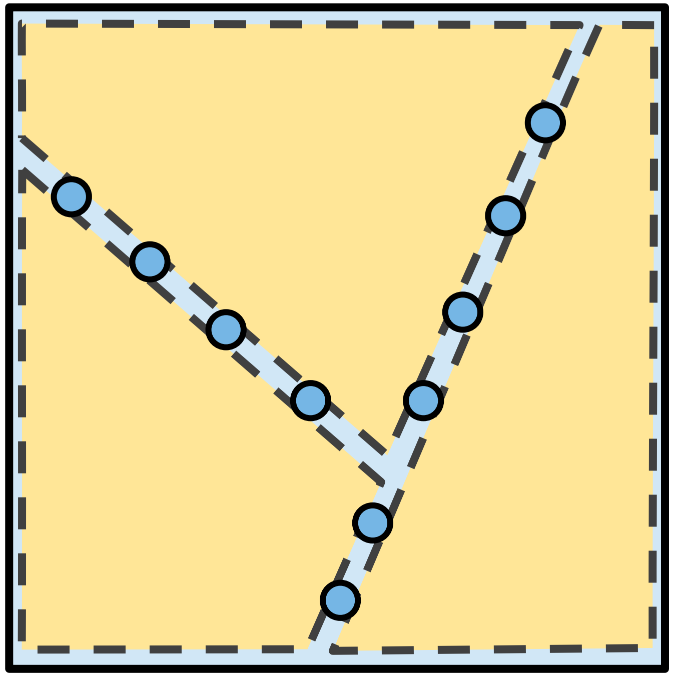

THE DIVERGENCE THEOREM

[Gauss 1813]

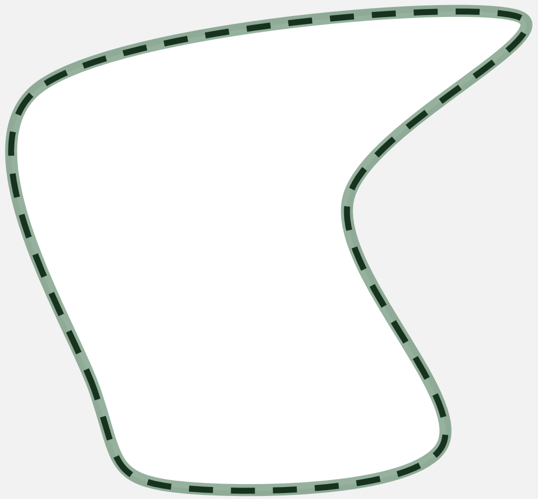

$$\int_{\partial D} \vec{f} \cdot \vec{n}$$

$$\Leftrightarrow$$

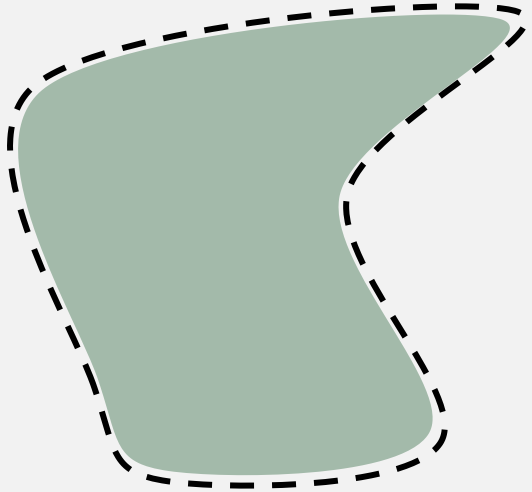

$$\int_{D} \nabla \cdot \vec{f}$$

APPLYING THE DIVERGENCE THEOREM TO THE EDGE INTEGRAL

Goal: Rewrite $\int_{\partial D} f \vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$ into area integral $\int_{D} g$

Solution: Rewrite $\int_{\partial D} f \vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$ into $\int_{D} \nabla \cdot (\vec{\mathcal{V}}_\theta f)$

$\int_{D} \nabla \cdot (\vec{\mathcal{V}}_\theta f)$ can be estimated through area samples

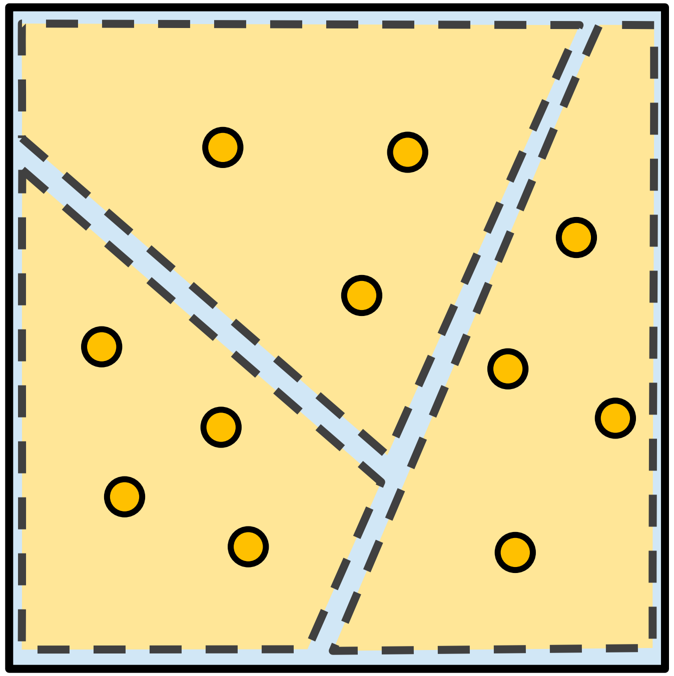

QUICK RECAP

- Used *Reynolds transport theorem* to find the boundary integral $\int_{\partial \mathrm{D}} f \vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$
- Rewrote $\int_{\partial \mathrm{D}} f \vec{\mathbf{v}} \cdot \vec{\mathbf{n}}$ to $\int_{\mathrm{D}} \nabla \cdot (\vec{\mathcal{V}}_{\theta} f)$ using the *divergence theorem*.
- Have to define the *vector field* $\vec{\mathcal{V}}_{\theta}$ over domain D

A 2D EXAMPLE SCENE

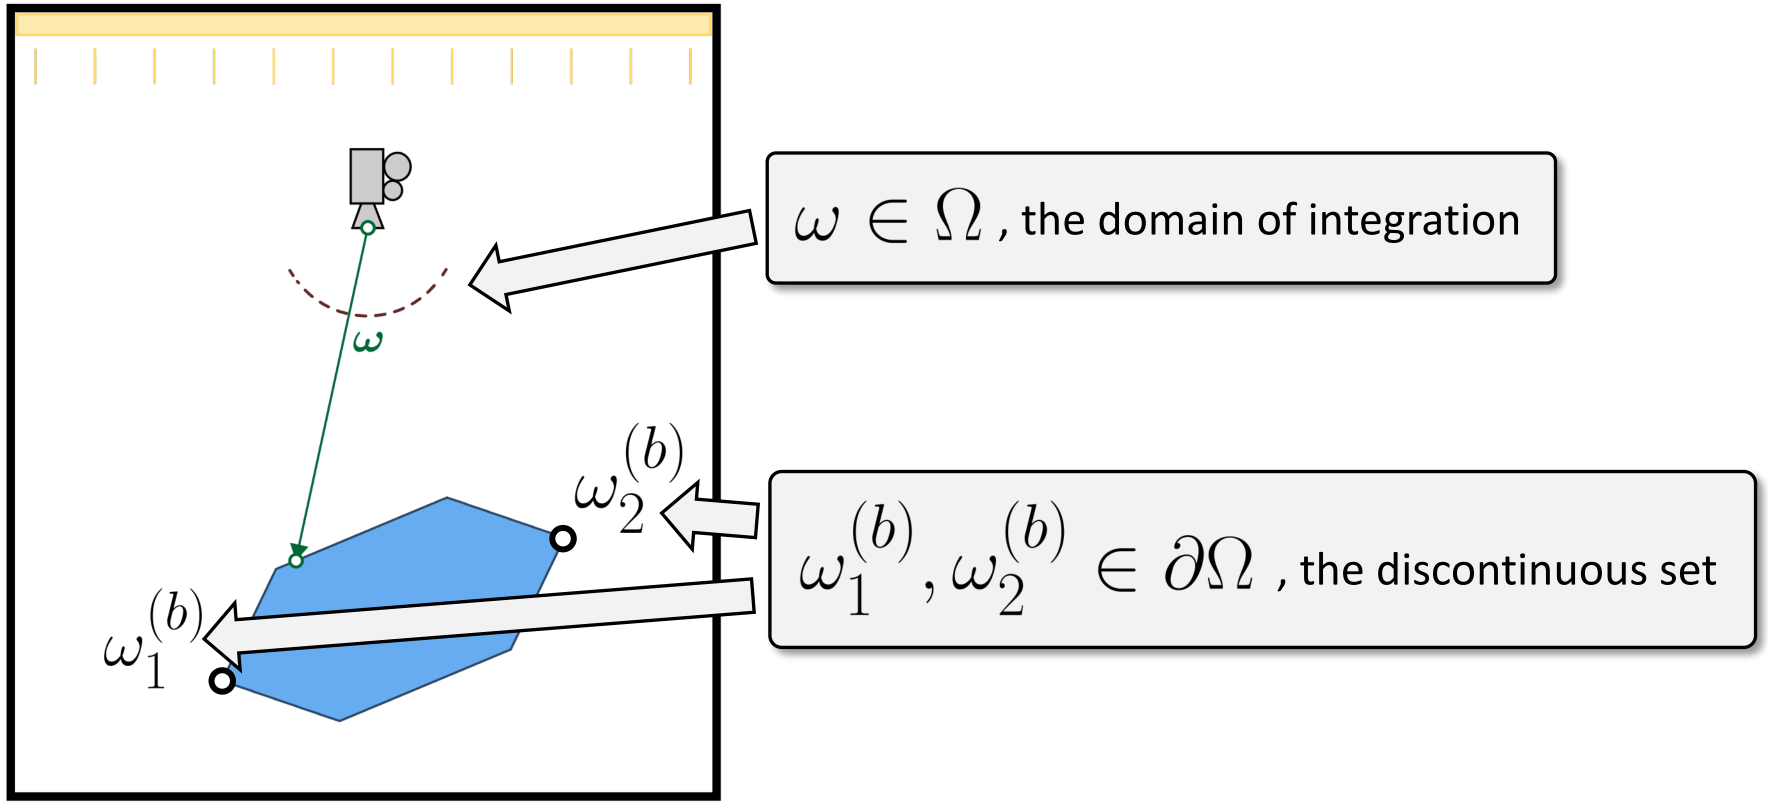

VELOCITY $\vec{\mathbf{V}}$: THE BOUNDARY DERIVATIVE

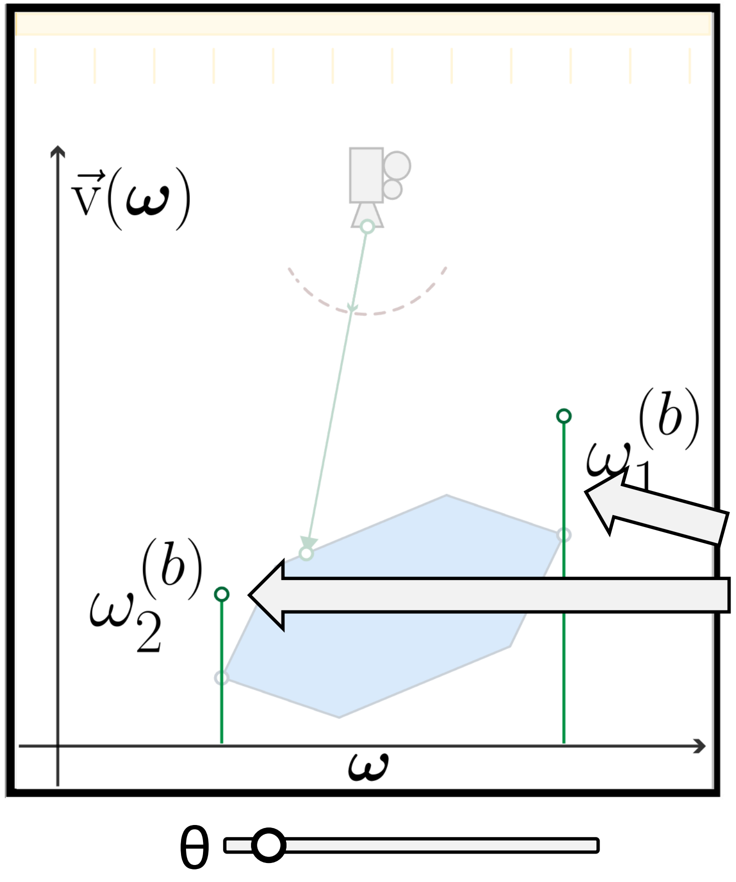

$\partial_\theta \omega_i^{(b)}$: Derivative of boundary position w.r.t **θ**

WARP FIELD $\mathcal{V}_\theta$: EXTENSION OF $\vec{\mathbf{V}}$ TO ALL POINTS

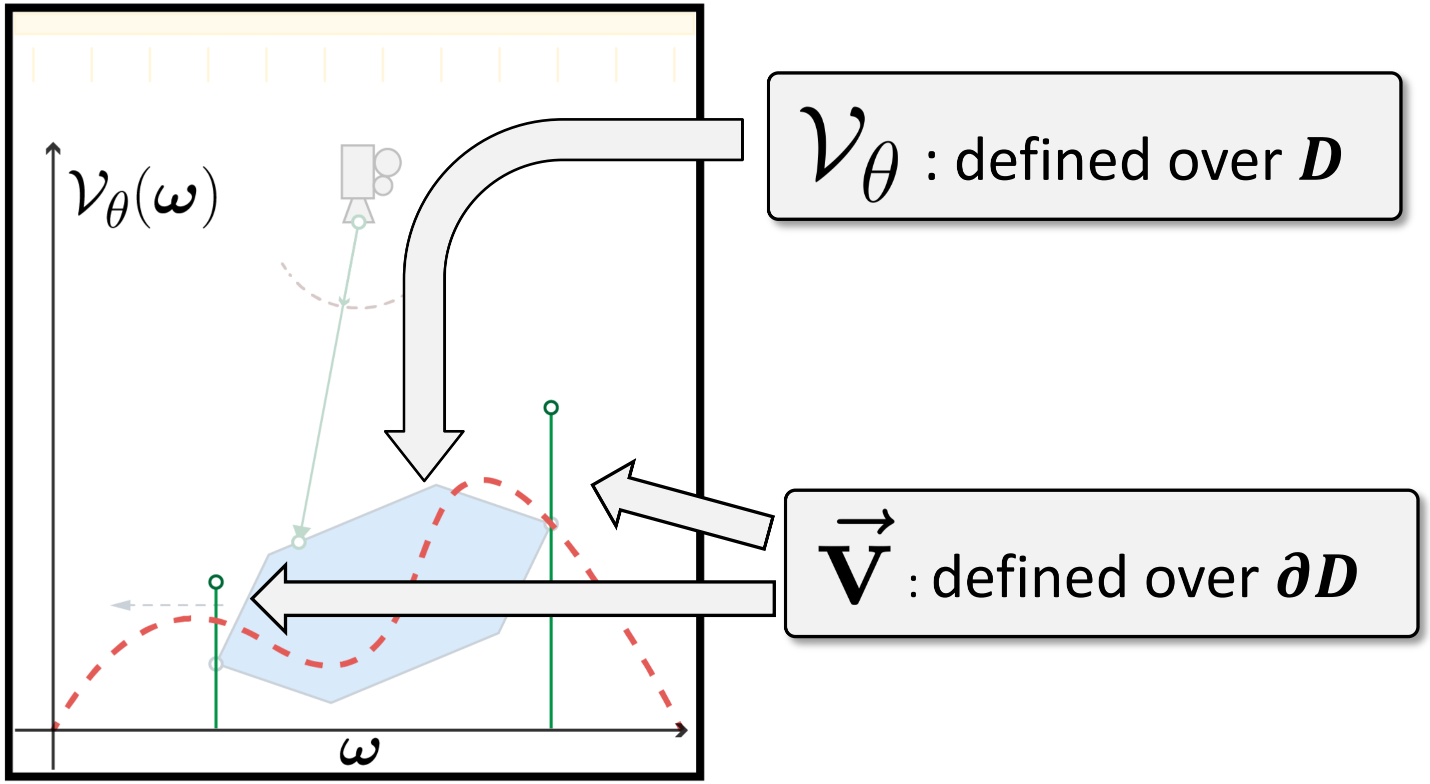

VALIDITY OF $\vec{\mathcal{V}}_\theta$

Rule 1: Continuous

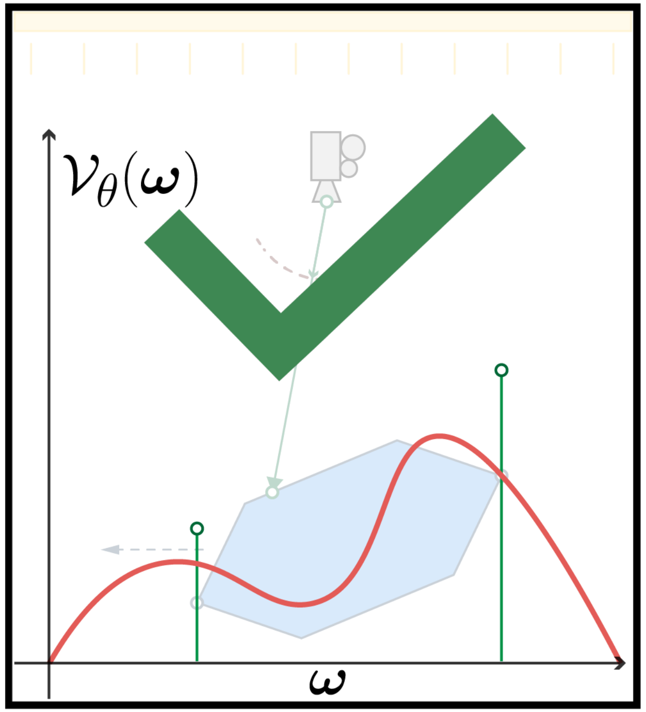

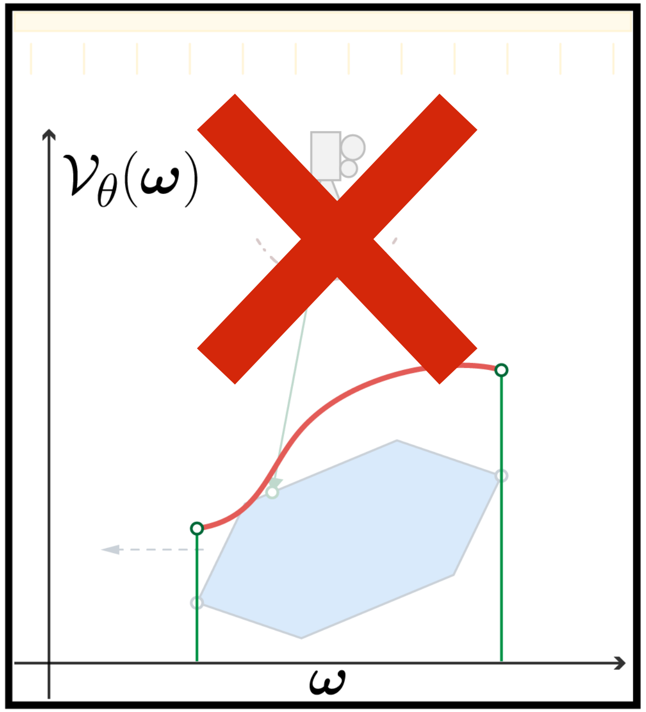

VALIDITY OF $\vec{\mathcal{V}}_\theta$

Rule 2: Boundary Consistent

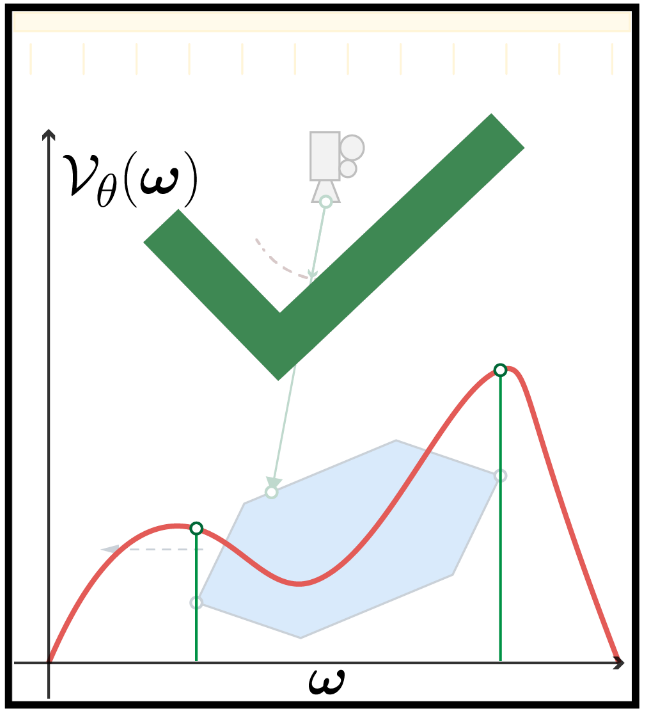

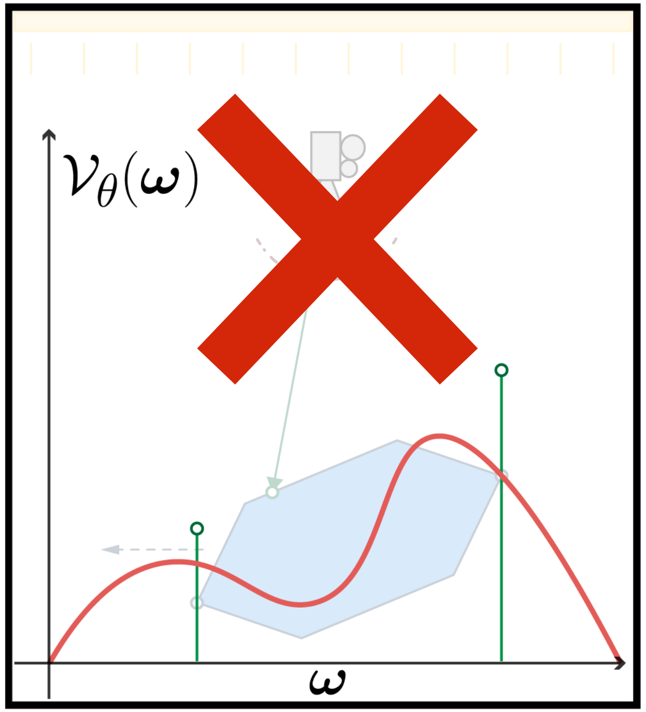

INTERPOLATION WITHOUT KNOWLEDGE OF BOUNDARIES

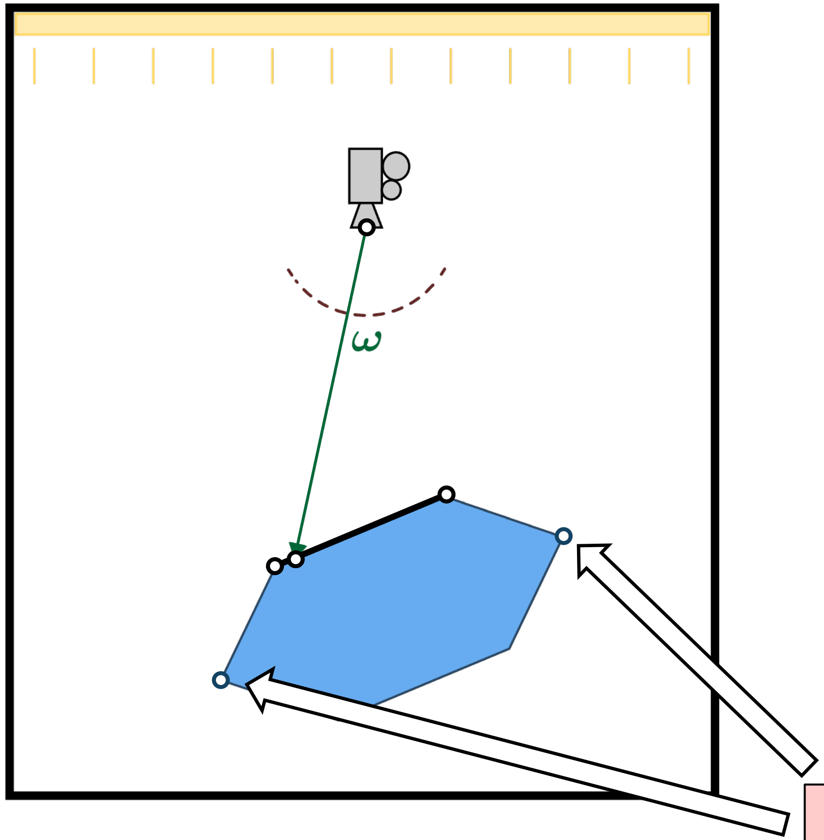

Available quantities

- Origin point
- Ray
- Intersection
- Primitive

No access to discontinuity points

CONSTRUCTING $\vec{\mathcal{V}}_\theta$

Attempt 1 ➡ Find $\partial_\theta \boldsymbol{\omega}$ through *implicit derivative*

(Incorrect)

$$\mathbf{y} = \textsc{Intersect}(\boldsymbol{\omega}, \theta) \Longrightarrow \partial_\theta \boldsymbol{\omega} = \frac{\partial_{\boldsymbol{\omega}} \mathbf{y}}{\partial_\theta \mathbf{y}}$$

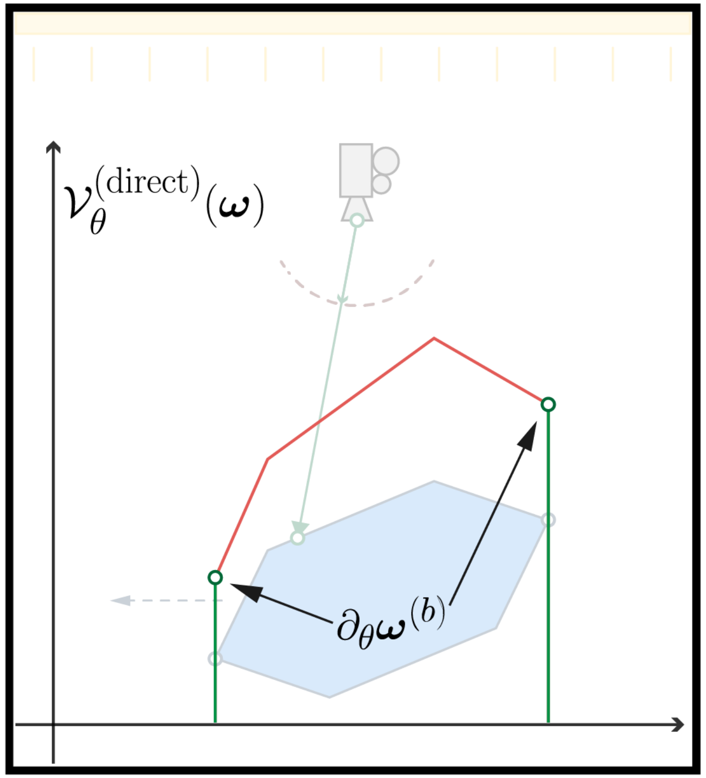

At all points (not just boundaries)

+ Boundary consistent

- Not continuous

CONSTRUCTING $\vec{\mathcal{V}}_\theta$

Attempt 2 ➞ Filter *Attempt 1* with a Gaussian filter

(Incorrect)

$$\int_{\Omega'} k(\boldsymbol{\omega}, \boldsymbol{\omega}') \frac{\partial_{\boldsymbol{\omega}} \mathbf{y}}{\partial_\theta \mathbf{y}}$$

k(.,.) = Gaussian filter

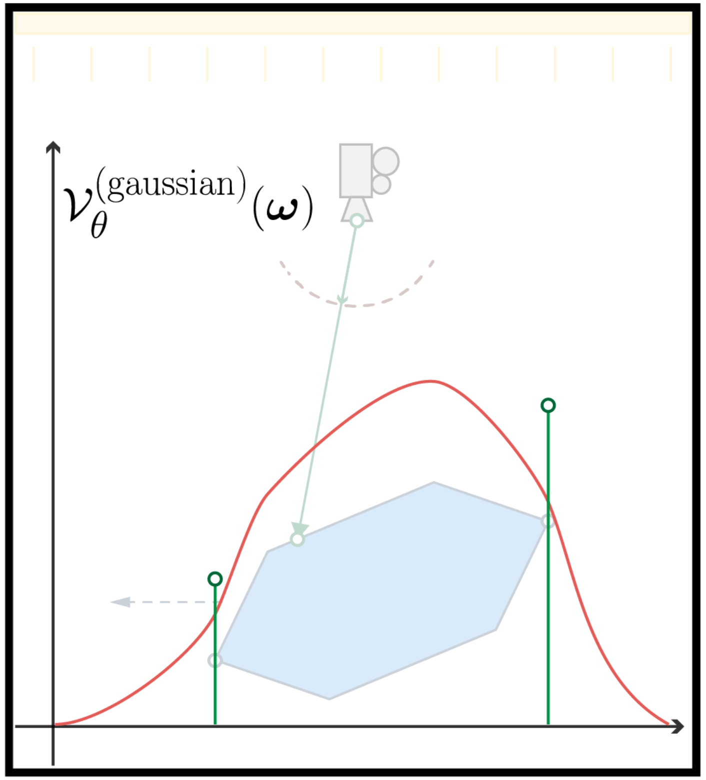

+ Continuous

- Not boundary consistent

BOUNDARY-AWARE WEIGHTING

<u>Goal</u>: Find weights $k(\boldsymbol{\omega}, \boldsymbol{\omega}')$ s.t. $\vec{\mathcal{V}}_\theta = \frac{\partial_\omega \mathbf{y}}{\partial_\theta \mathbf{y}}$ at boundaries.

Ideal weighting function

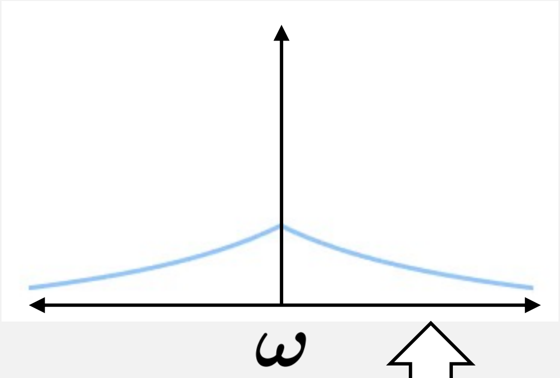

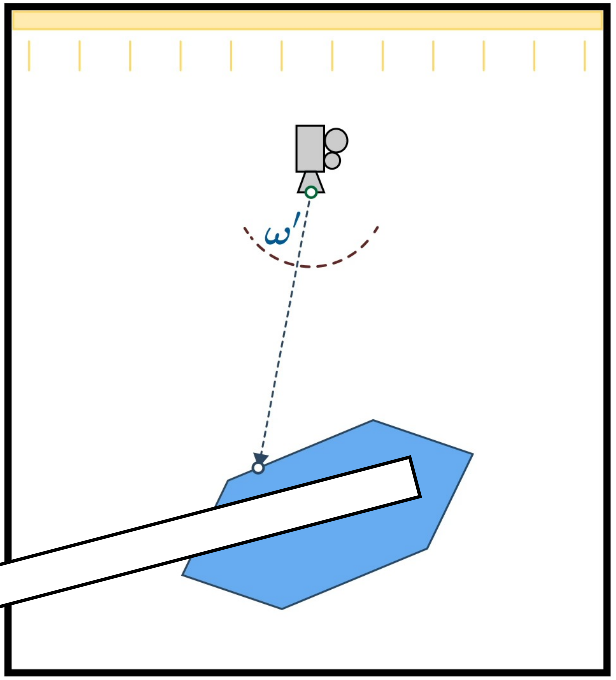

Approach Dirac delta near boundaries

BOUNDARY-AWARE WEIGHTING

Implicit Boundary through *geometric normals*

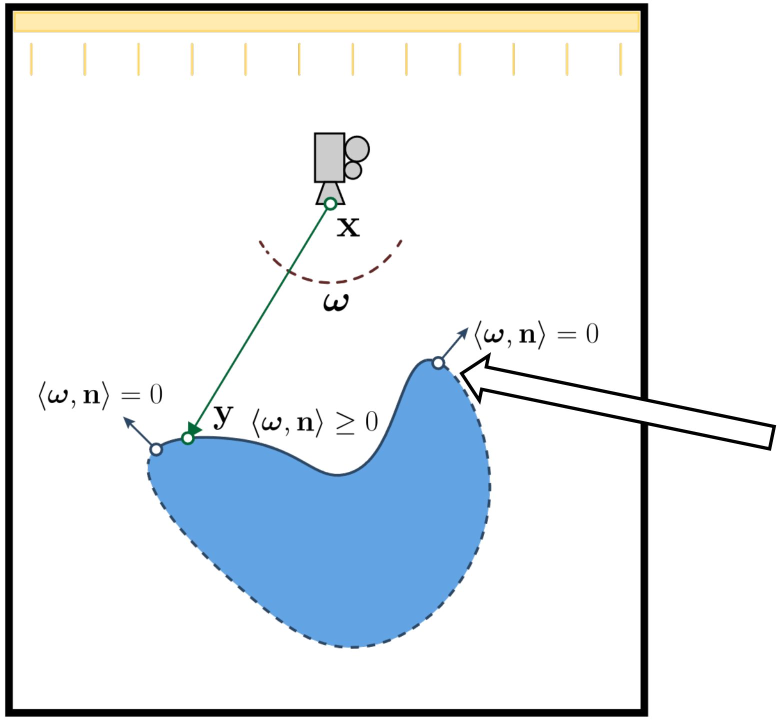

$$\langle \boldsymbol{\omega}, \mathbf{n} \rangle = 0$$

at boundaries

y sampling)

boundary
g)

CONSTRUCTING $\vec{\mathcal{V}}_\theta$

Our Approach ⟶ Filter *Attempt 1* with harmonic weights

$$k(\boldsymbol{\omega}, \boldsymbol{\omega}') = \frac{1}{\mathcal{D}(\boldsymbol{\omega}, \boldsymbol{\omega}') + \mathcal{B}(\boldsymbol{\omega}')}$$

Distance function

Boundary test

+ Boundary consistent

+ Continuous

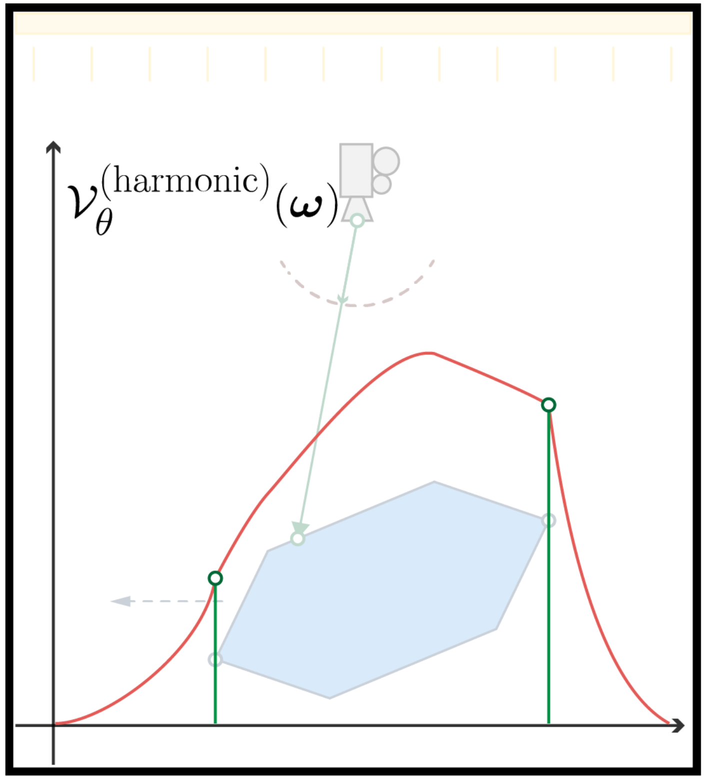

COMPUTING $\vec{V}_\theta$

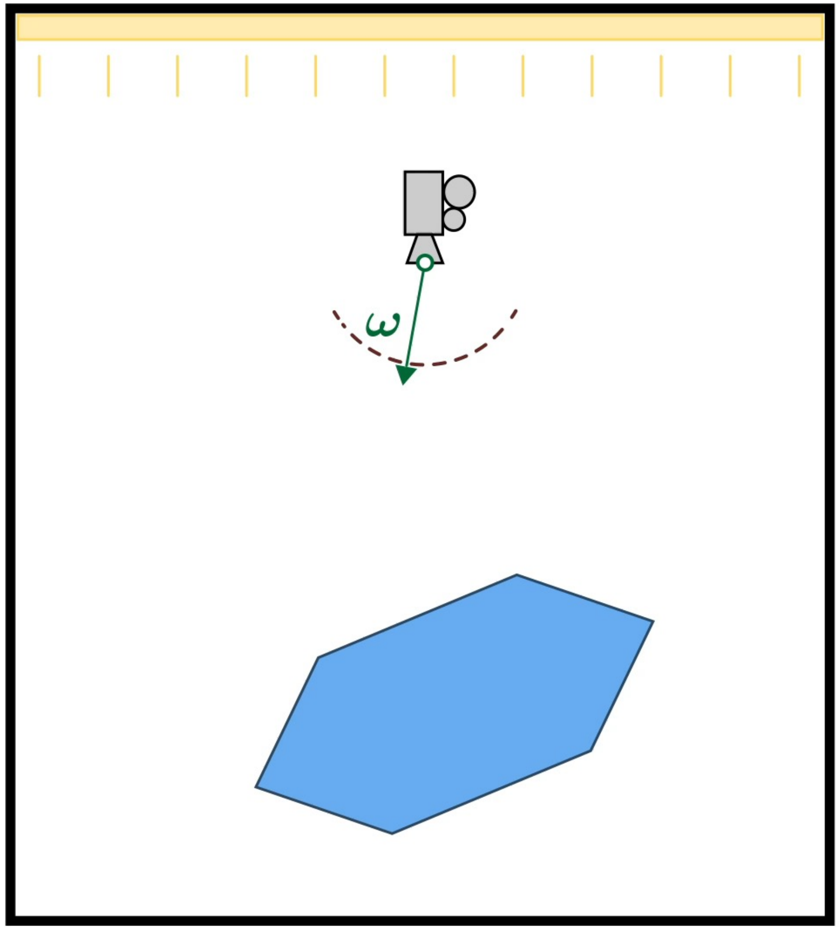

1. Sample **path** using path tracer *(N paths)*

For each bounce:

2. Sample **auxiliary** rays *(N' rays)*

3. Compute boundary term ***B()*** locally

4. Compute weight **k(.,.)** and $\partial_\theta \boldsymbol{\omega}$

5. Find weighted mean

QUICK RECAP

- Used *Reynolds transport theorem* to find the boundary integral $\int_{\partial D} f\vec{\mathbf{v}}\cdot\vec{\mathbf{n}}$
- Rewrote $\int_{\partial D} f\vec{\mathbf{v}}\cdot\vec{\mathbf{n}}$ to $\int_{D} \nabla\cdot(\vec{\mathcal{V}}_\theta f)$ using the *divergence theorem*.
- Estimate **consistent** and **continuous** $\vec{\mathcal{V}}_\theta$ over domain D using auxiliary rays

MORE INTUITION: WARP-AREA SAMPLING CAN BE SEEN AS A CHANGE OF VARIABLE

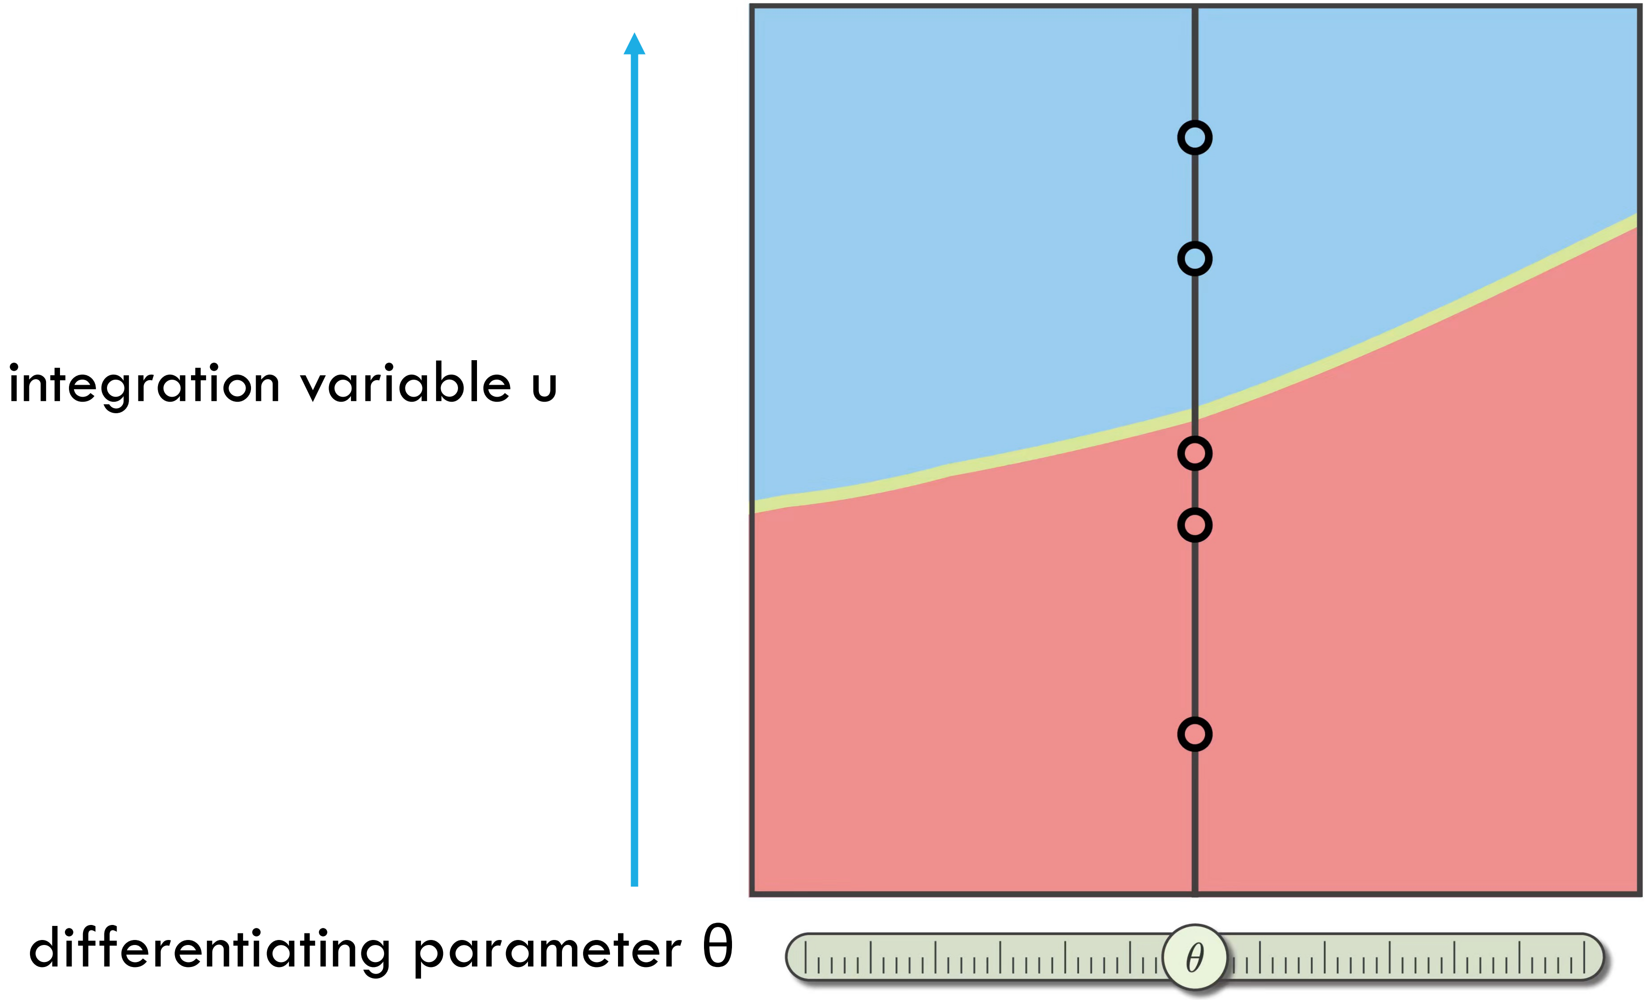

TRANSFORM SAMPLES

$$u = T(u'; \theta)$$

MORE INTUITION: WARP-AREA SAMPLING CAN BE SEEN AS A CHANGE OF VARIABLE

$$\frac{\partial}{\partial\theta}\int_D f \mathrm{d}u = \int_D f_\theta + \nabla \cdot \left(\vec{\mathcal{V}}_\theta f\right) \mathrm{d}u$$

MORE INTUITION: WARP-AREA SAMPLING CAN BE SEEN AS A CHANGE OF VARIABLE

$$\frac{\partial}{\partial\theta}\int_D f\,\mathrm{d}u = \int_D f_\theta + \nabla\cdot\left(\vec{\mathcal{V}}_\theta f\right)\mathrm{d}u$$

$$= \int_D \frac{\partial}{\partial\theta}\left(f\left(T(u';\theta)\right)J_T\right)\mathrm{d}u'$$

MORE INTUITION: WARP-AREA SAMPLING CAN BE SEEN AS A CHANGE OF VARIABLE

$$\frac{\partial}{\partial\theta}\int_D f \mathrm{d}u = \int_D f_\theta + \nabla \cdot \left(\vec{\mathcal{V}}_\theta f\right) \mathrm{d}u$$

$$= \int_D \frac{\partial}{\partial\theta}\left(f\left(T(u';\theta)\right) J_T\right) \mathrm{d}u'$$

$$T = u' + (\theta - \theta_0)\,\mathcal{V}_\theta$$

RESULTS

VARIANCE COMPARISON WITH EDGE-SAMPLING

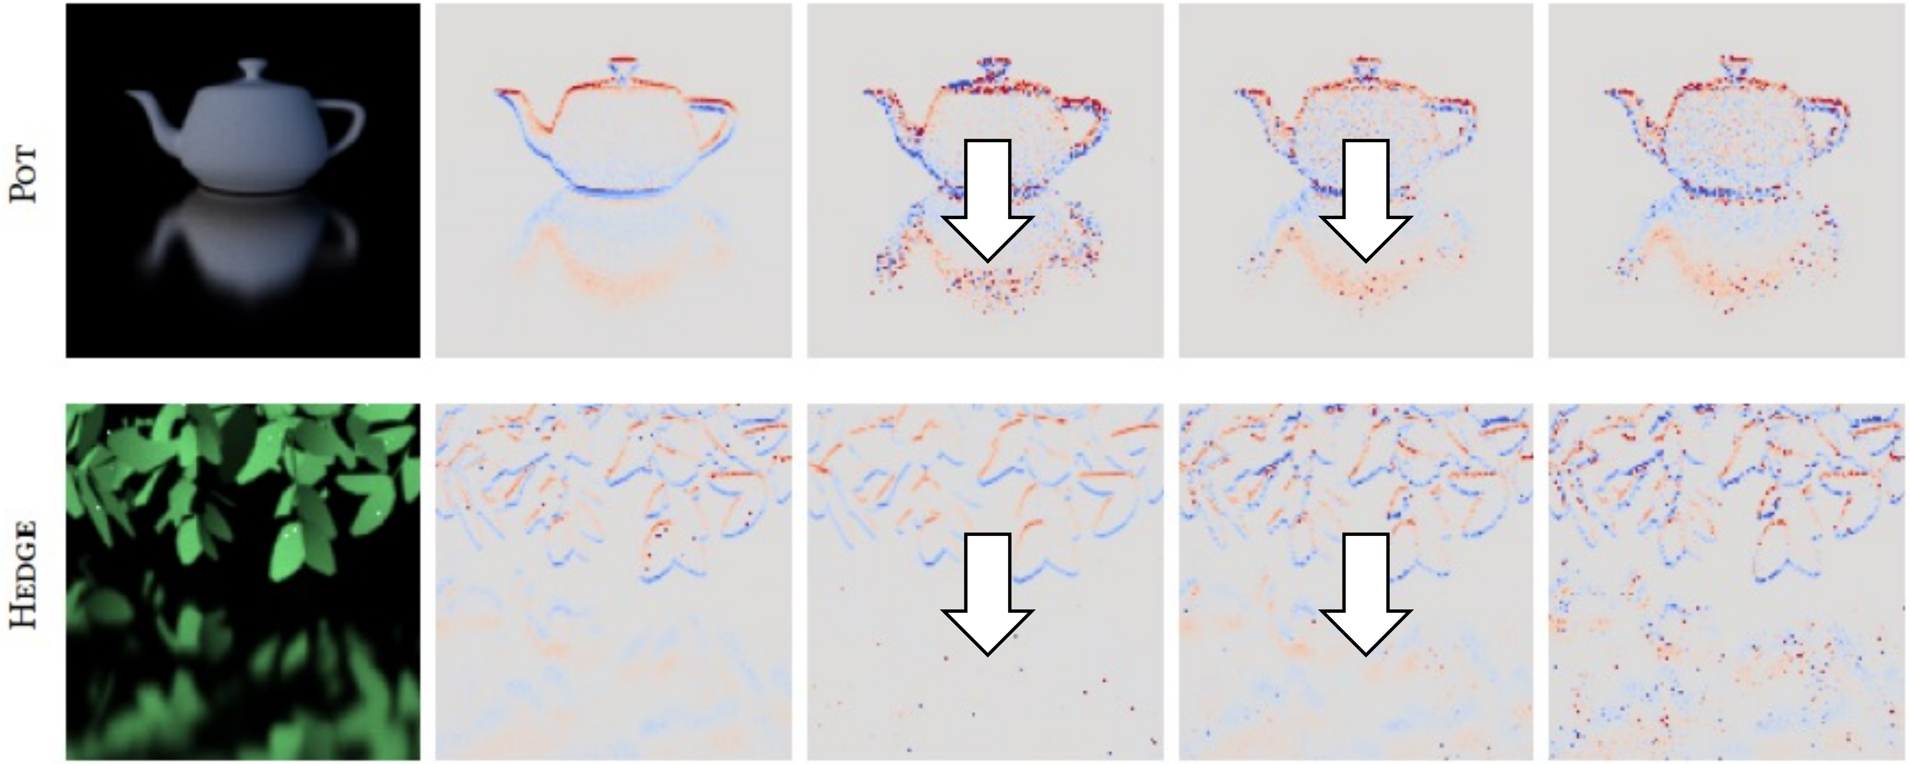

Image I | Reference Derivative | Li et al. 2018 | Ours *without* Russian roulette | Ours *with* Russian roulette

BIAS COMPARISON WITH REPARAMETERIZATION

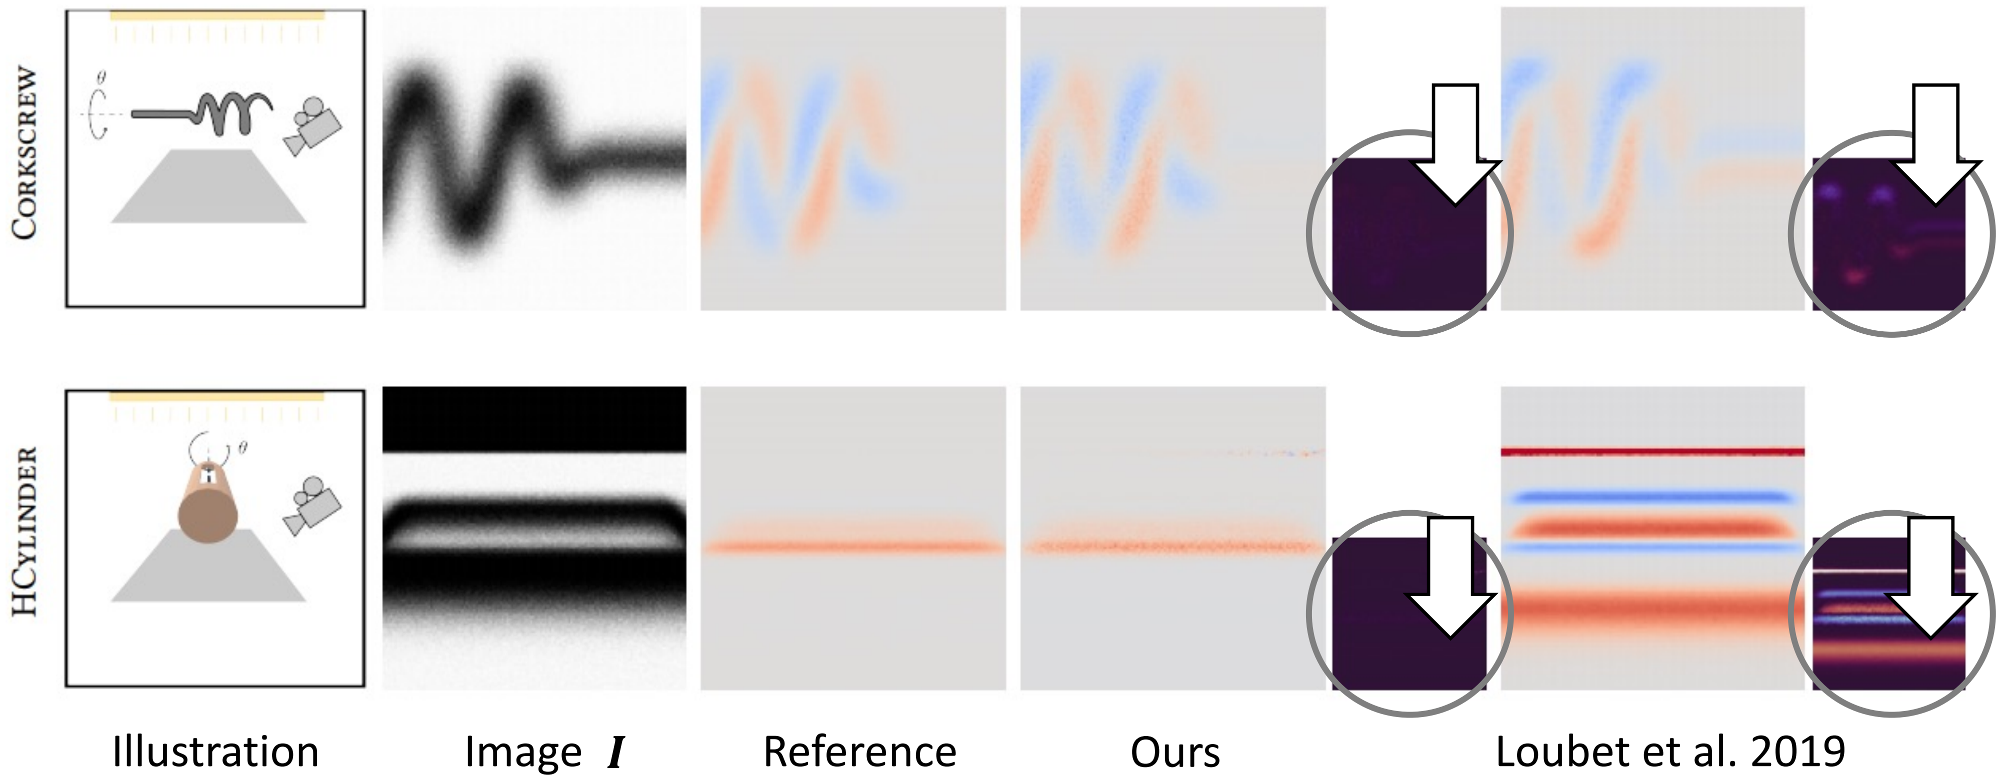

Rotating cylindrical objects present a complicated scenario for area-sampling

BIAS COMPARISON WITH REPARAMETERIZATION

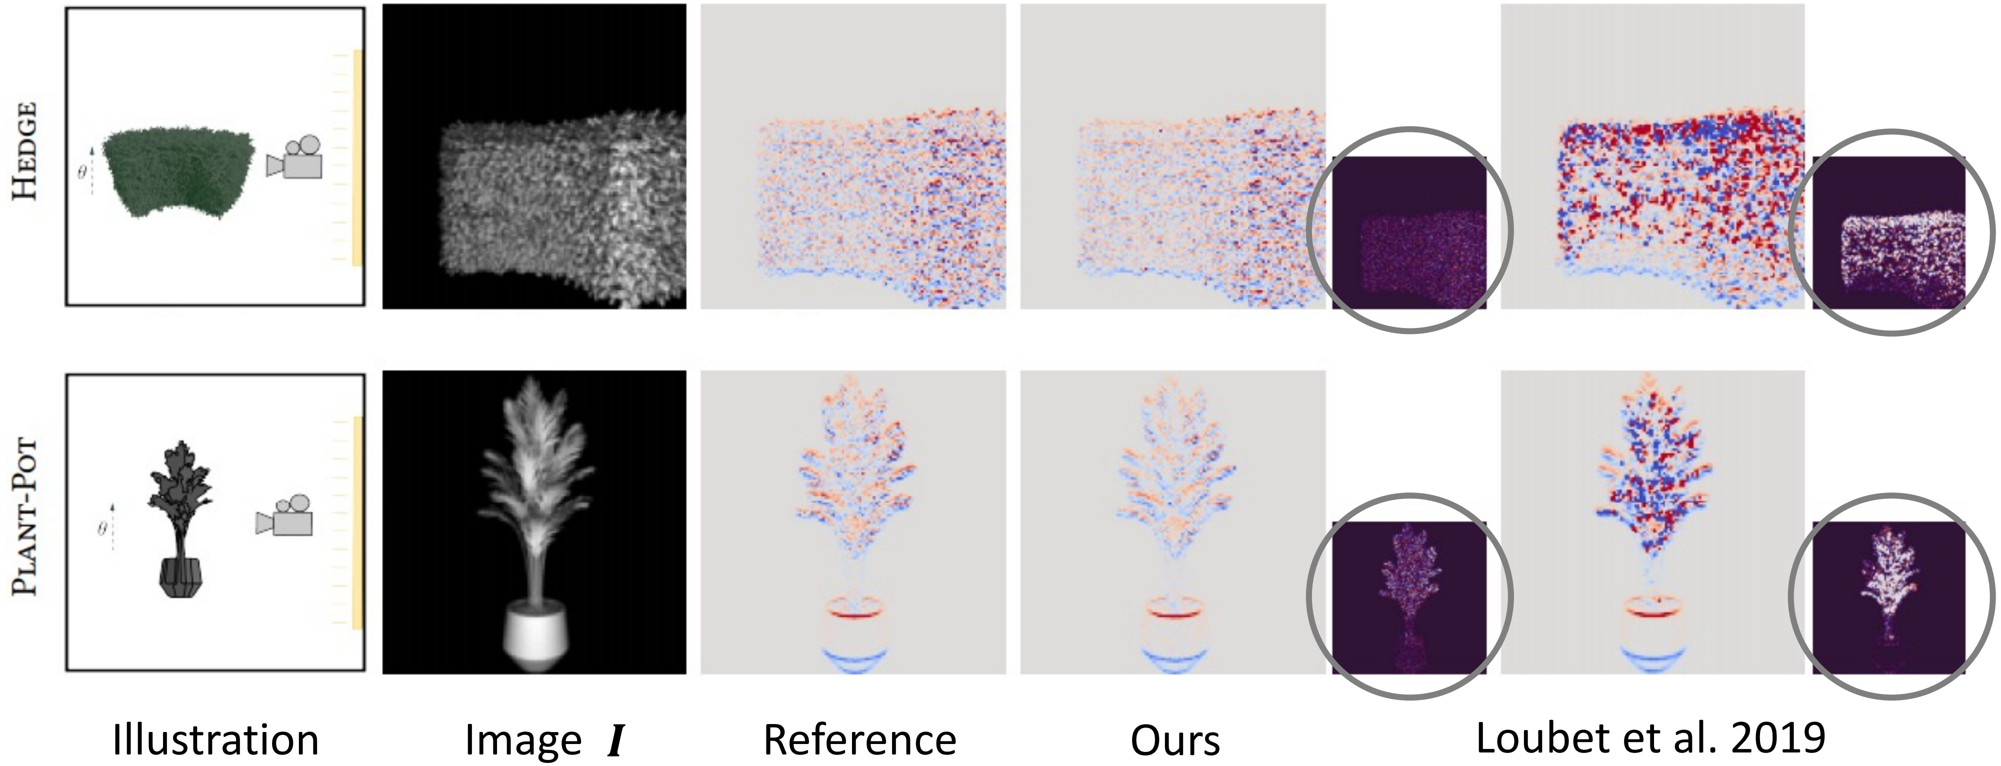

Extremely complex geometry like foliage can cause heuristic to fail

POSE ESTIMATION CAN FAIL WITH BIASED GRADIENTS

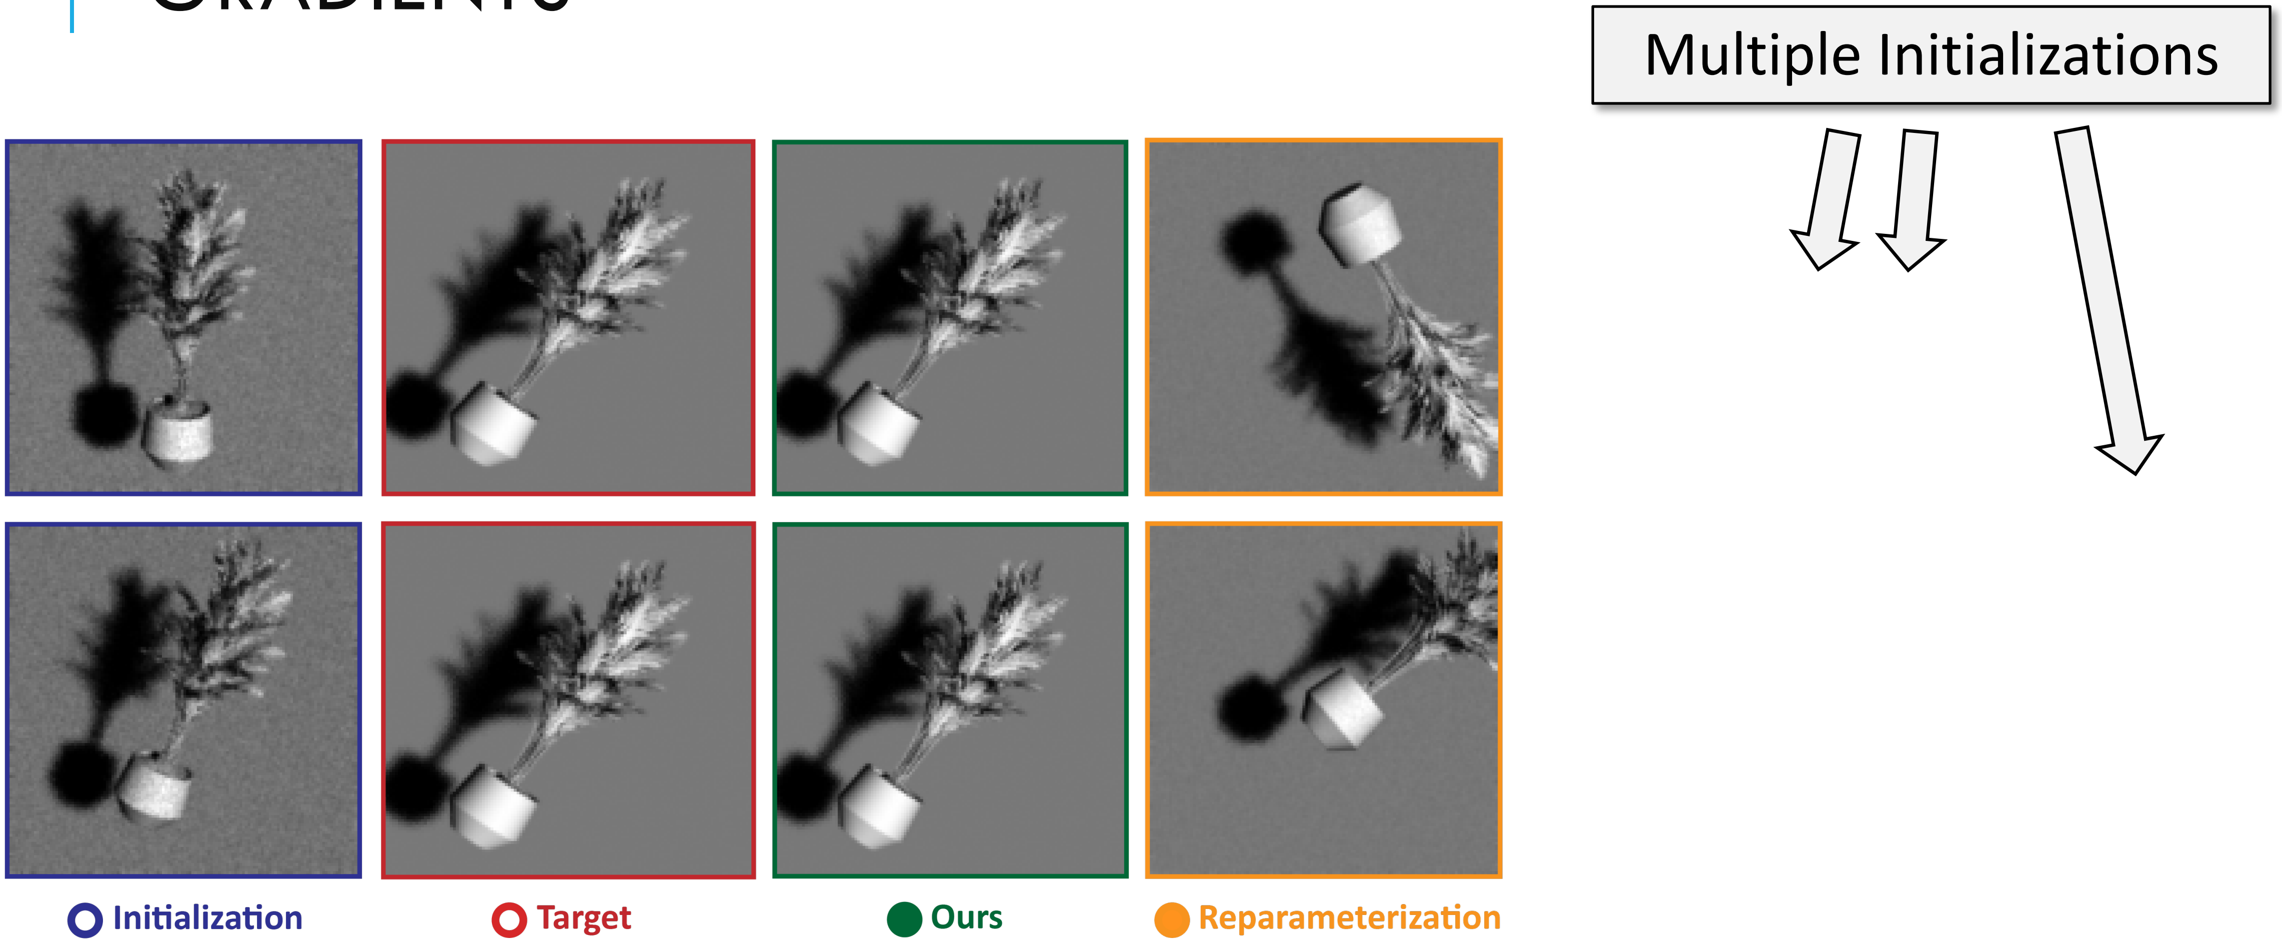

WARPED-AREA SAMPLING CAN BE USED FOR SIGNED DISTANCE FIELDS RENDERING

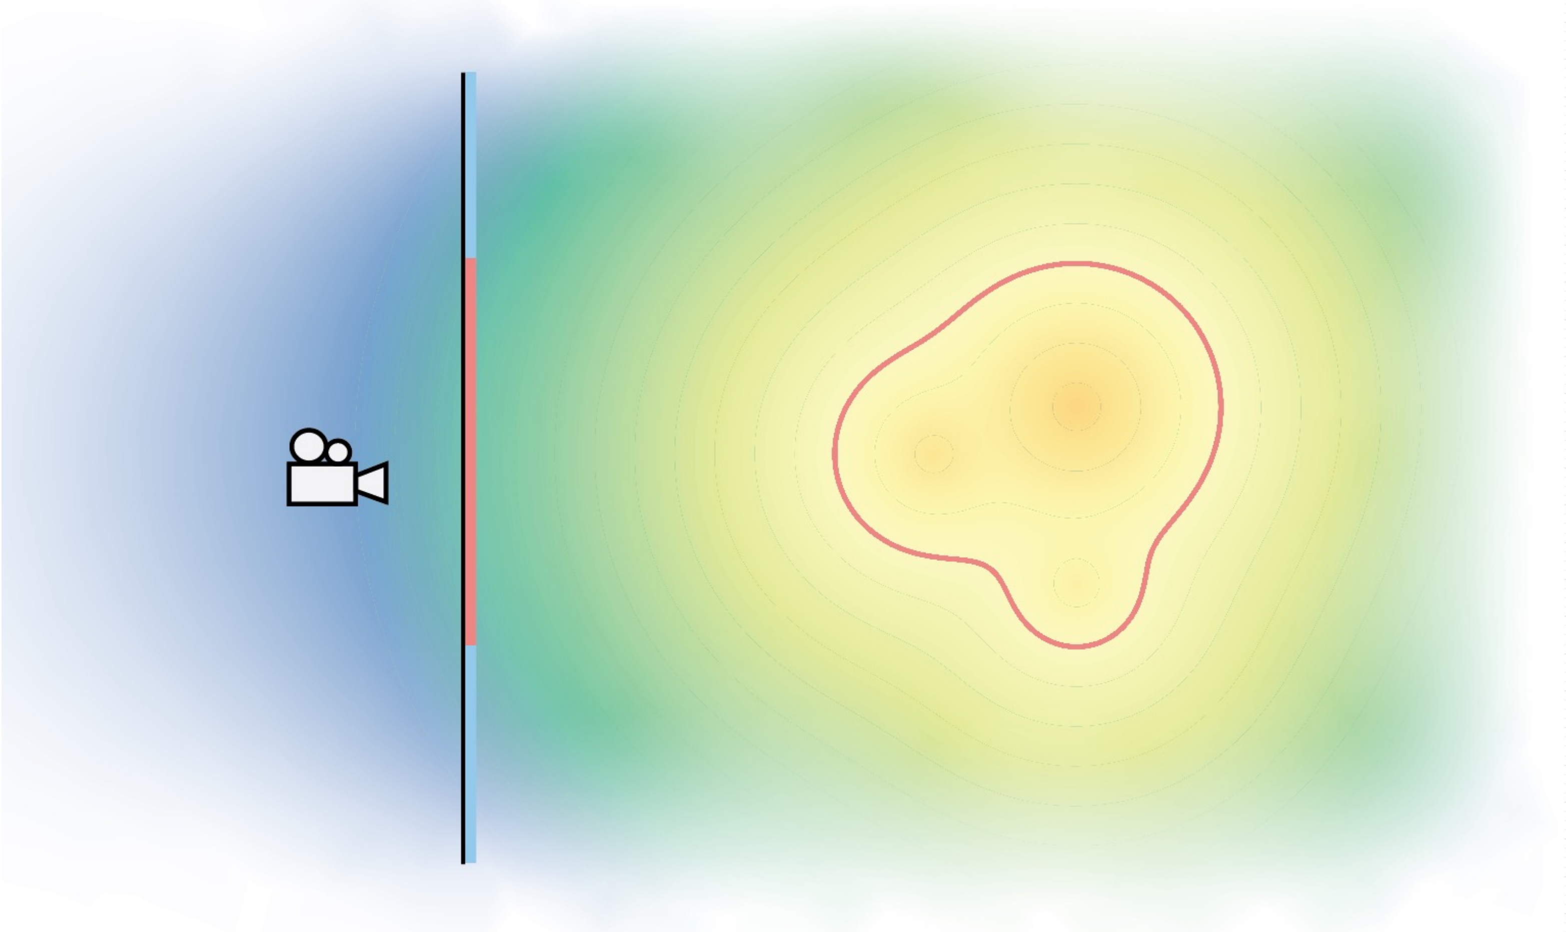

Differentiable Signed Distance Function Rendering

DELIO VICINI, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
SÉBASTIEN SPEIERER, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
WENZEL JAKOB, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland

Differentiable Rendering of Neural SDFs through Reparameterization

Sai Praveen Bangaru, MIT CSAIL, USA, sbangaru@mit.edu
Michaël Gharbi, Adobe Research, USA, mgharbi@adobe.com
Tzu-Mao Li, UC San Diego, USA, tzli@ucsd.edu
Fujun Luan, Adobe Research, USA, fluan@adobe.com
Kalyan Sunkavalli, Adobe Research, USA, sunkaval@adobe.com
Miloš Hašan, Adobe Research, USA, mihasan@adobe.com
Sai Bi, Adobe Research, USA, sbi@adobe.com
Zexiang Xu, Adobe Research, USA, zexu@adobe.com
Gilbert Bernstein, MIT CSAIL & UC Berkeley, USA, gilbo@berkeley.edu
Frédo Durand, MIT CSAIL, USA, fredo@mit.edu

Computing A Consistent $\mathcal{V}(u)$ For An Arbitrary SDF

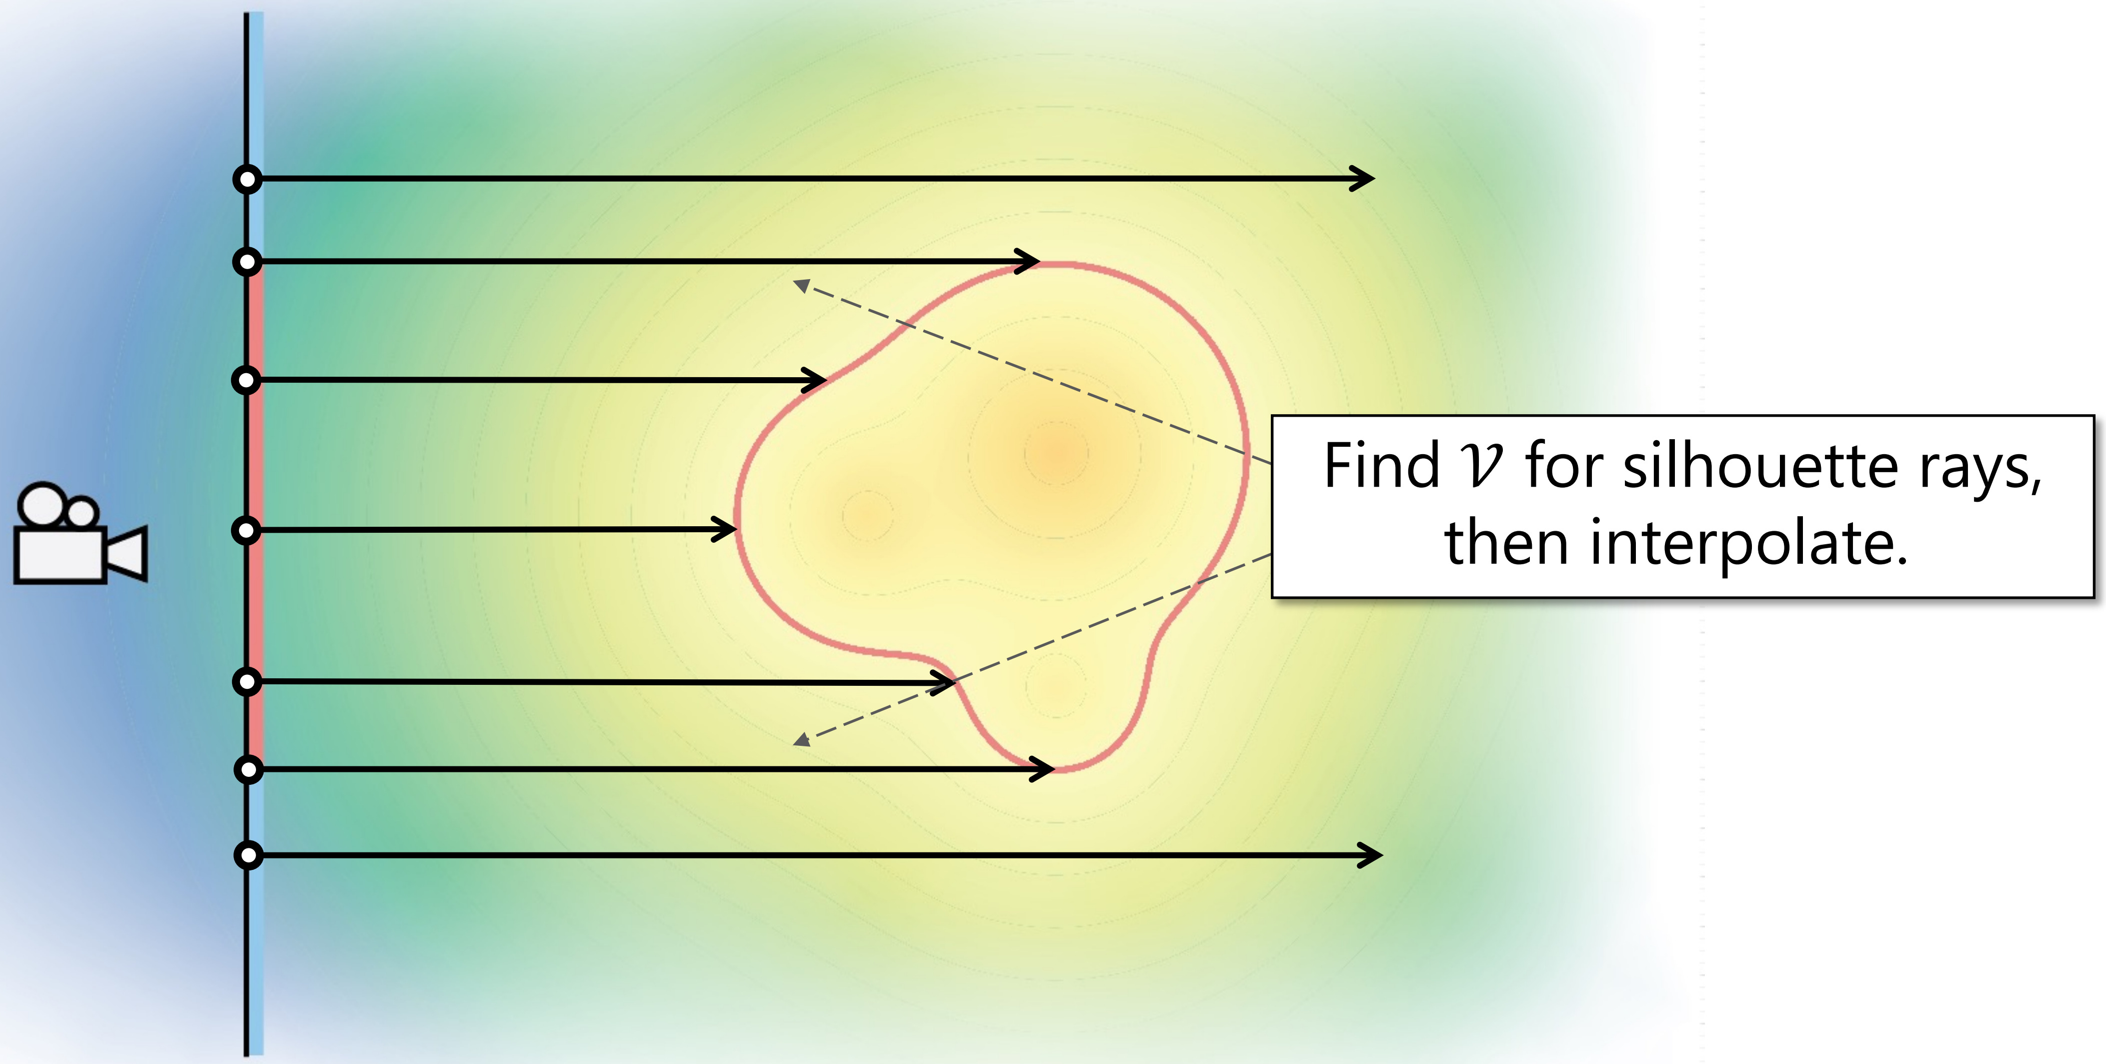

Computing $\mathcal{V}$ By Differentiating The Silhouette Position u^*

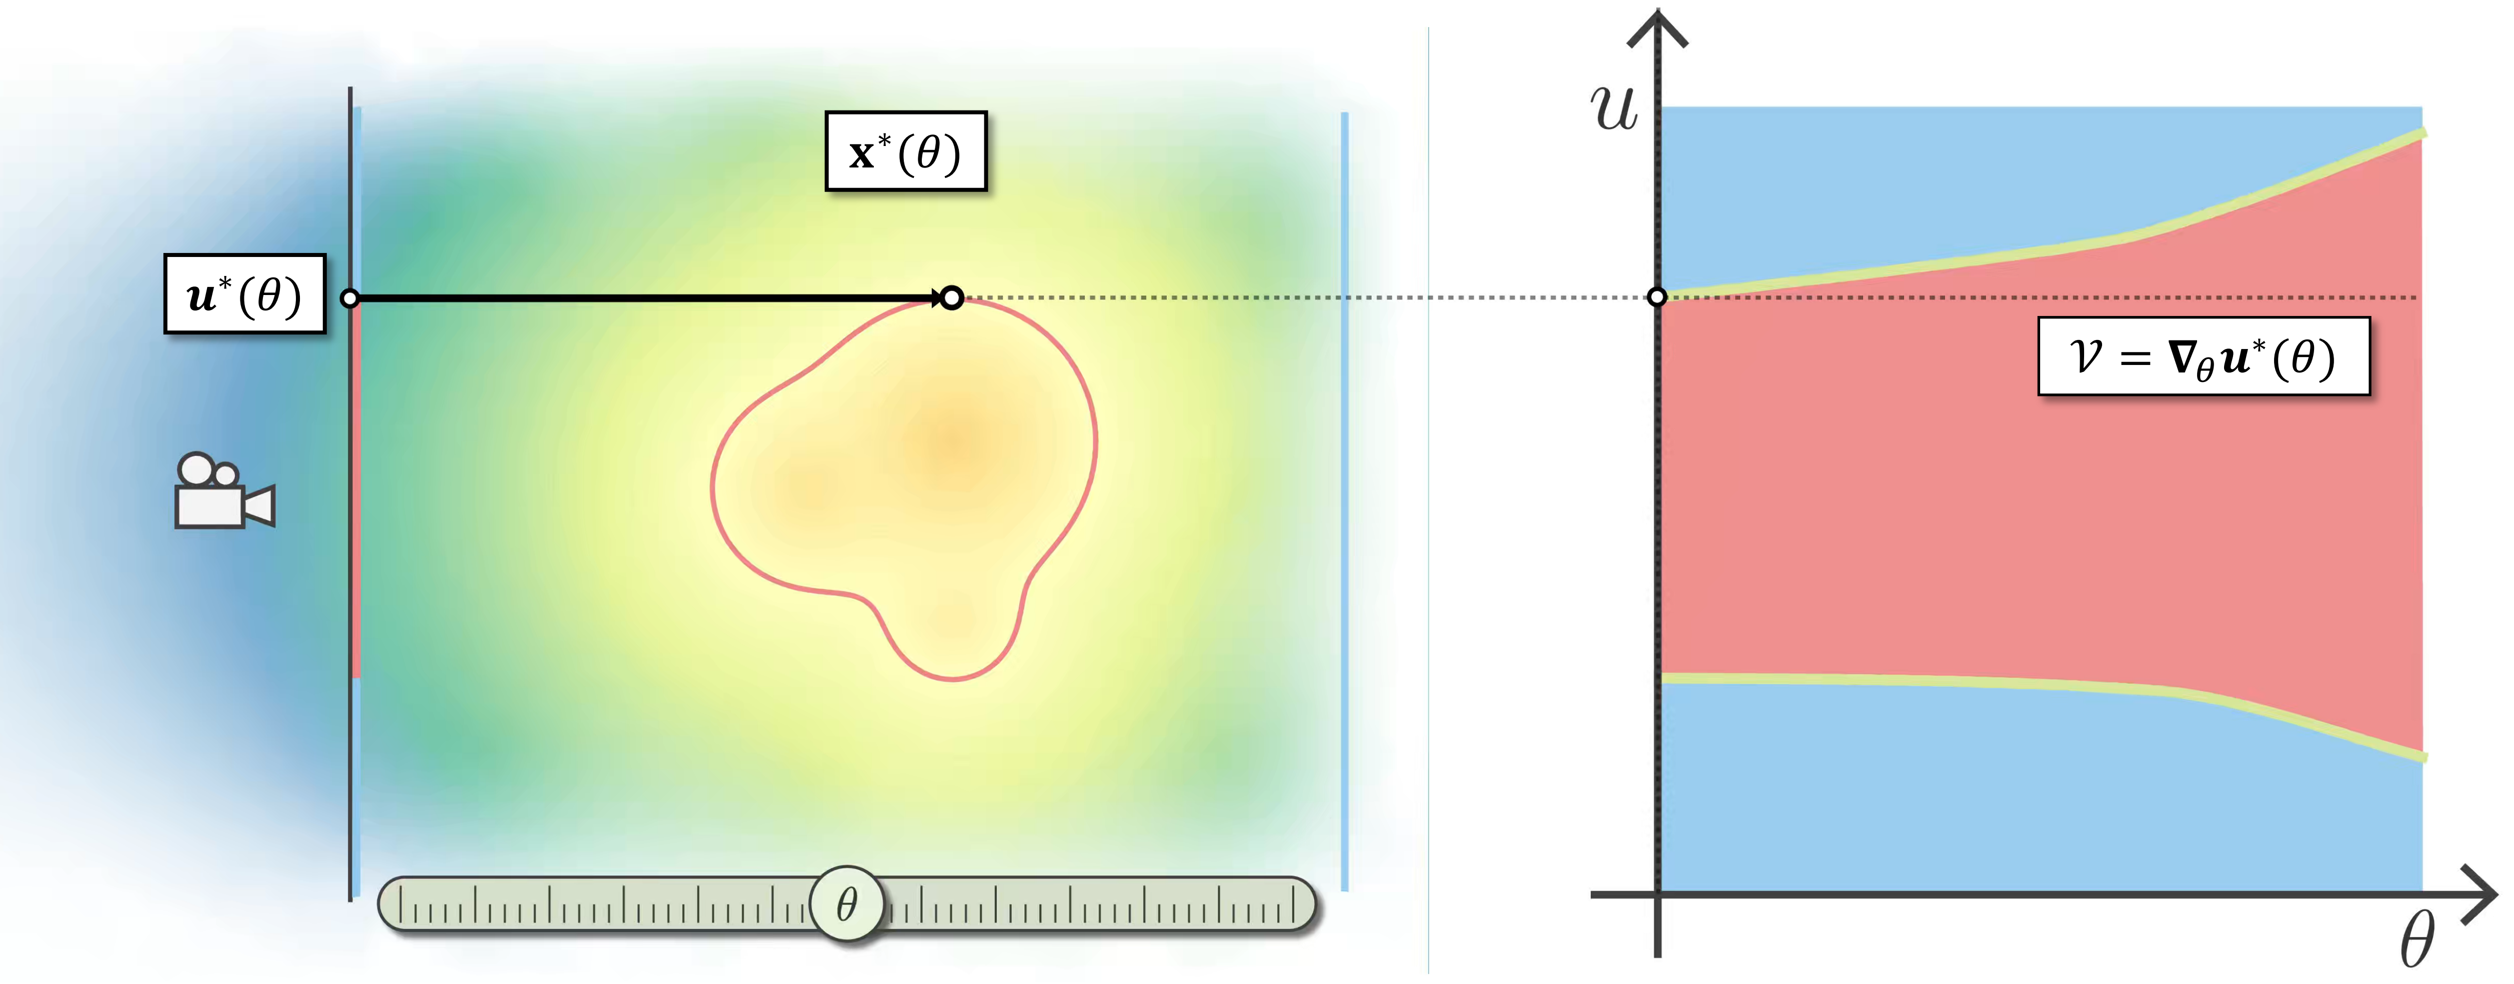

Computing $\mathcal{V}$: Implicit Fn. Theorem + Chain Rule

1. **Compute $\nabla_\theta \mathbf{x}^*(\theta)$ using implicit fn. theorem:**
   Derivative of any point in SDF can be computed by differentiating SDF function **f**

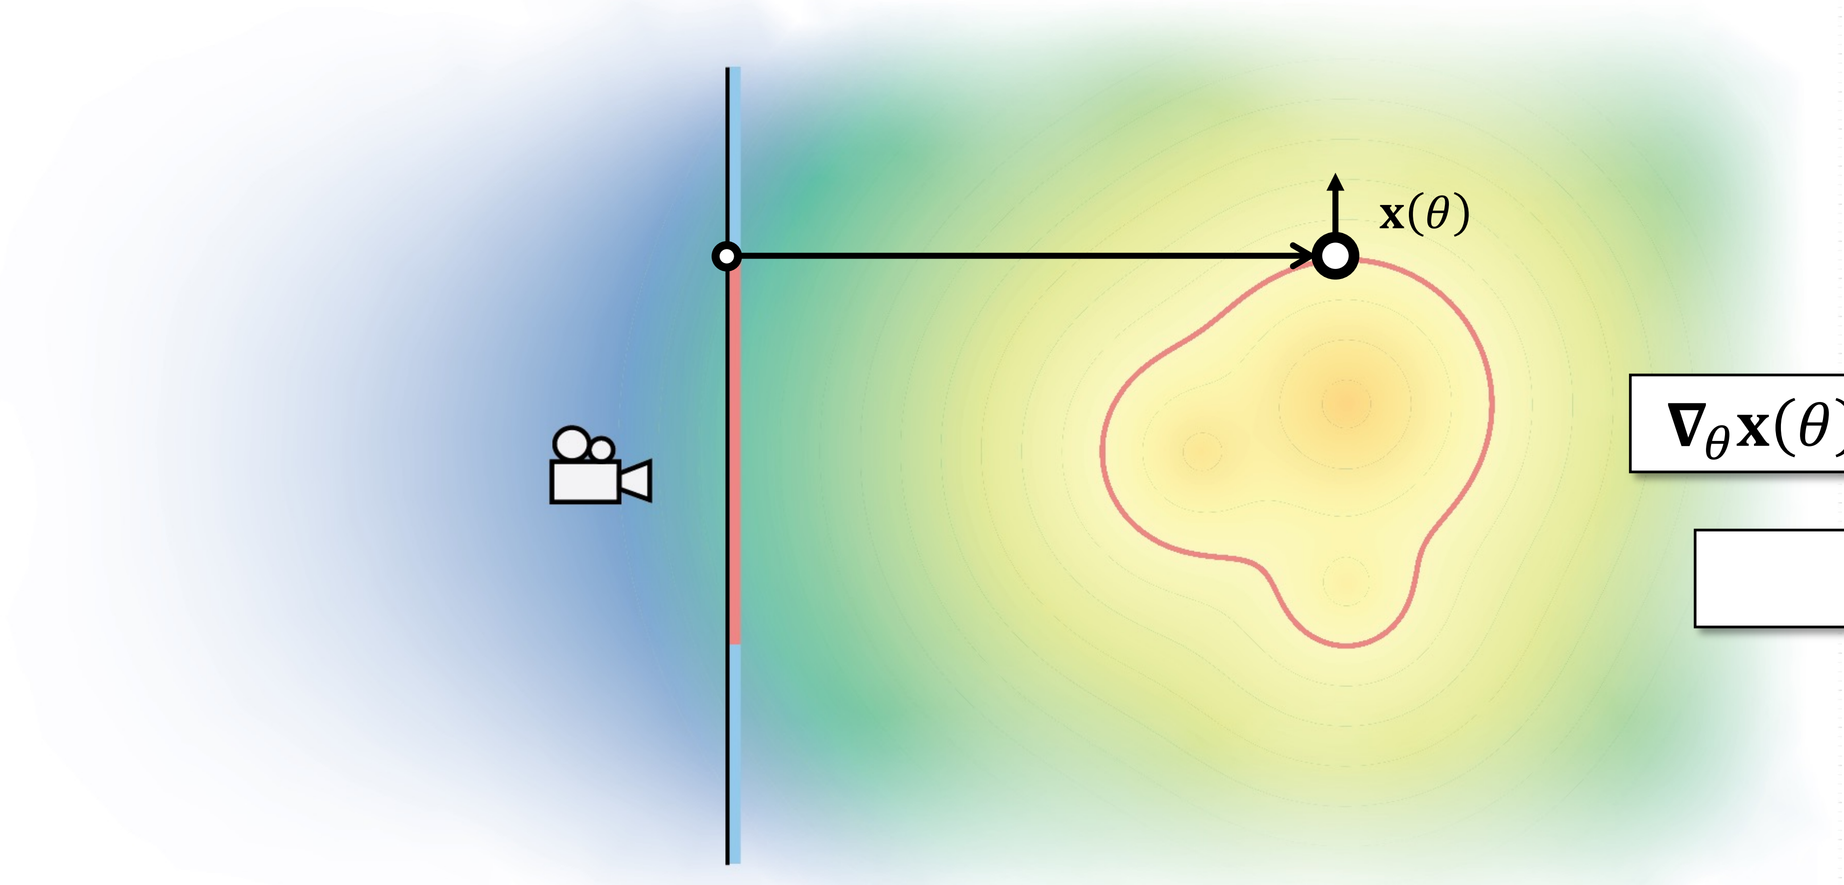

$$\nabla_\theta \mathbf{x}(\theta) = \nabla_{\boldsymbol{\theta}} \boldsymbol{f}(\mathbf{x}; \boldsymbol{\theta}) \cdot \hat{\boldsymbol{n}}$$

$$= \boldsymbol{G}(\mathbf{x})$$

Computing $\mathcal{V}$: Implicit Fn. Theorem + Chain Rule

1. **Compute $\nabla_\theta \mathbf{x}^*(\theta)$ using Implicit Fn. Theorem:**
2. **Propagate $G(\mathbf{x}^*)$ to sample space through chain rule ($u \to \mathbf{x}$):**

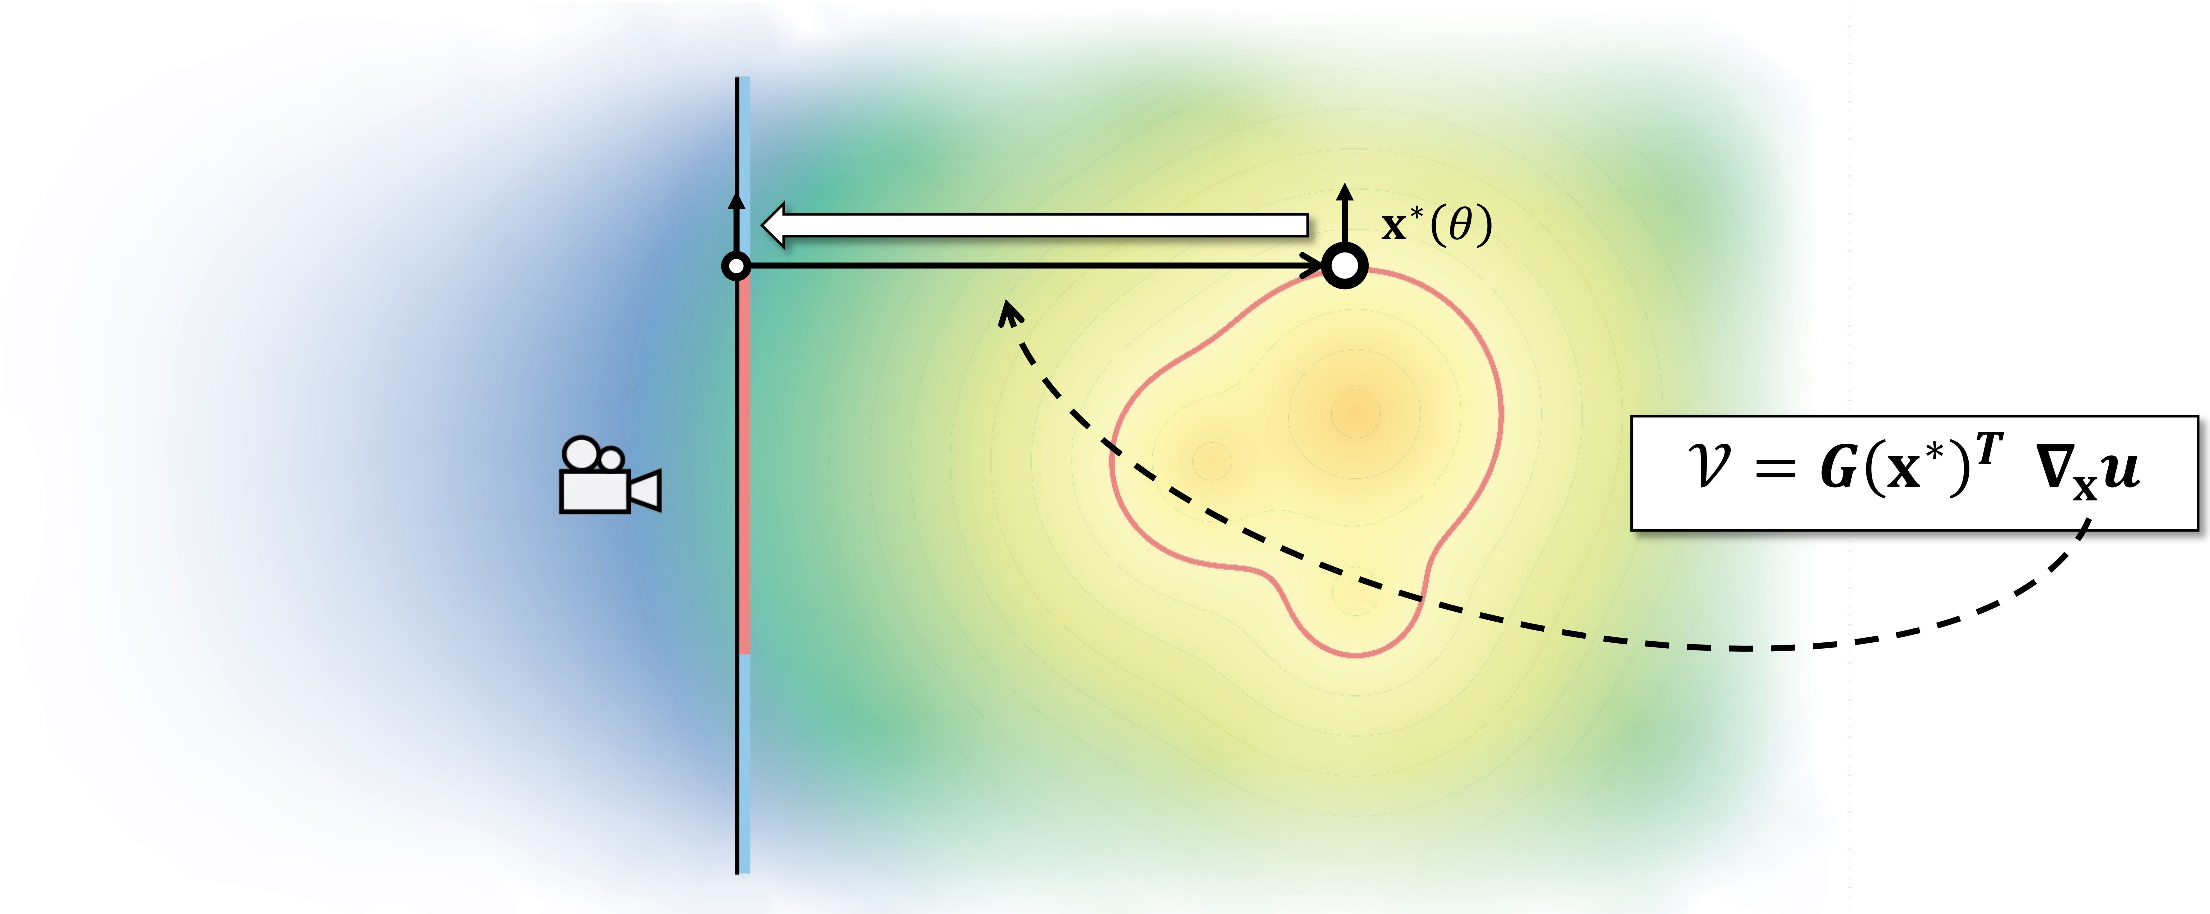

Computing $\mathcal{V}(u)$ For An Arbitrary Ray

Ray-SDF Intersection: Sphere Tracing

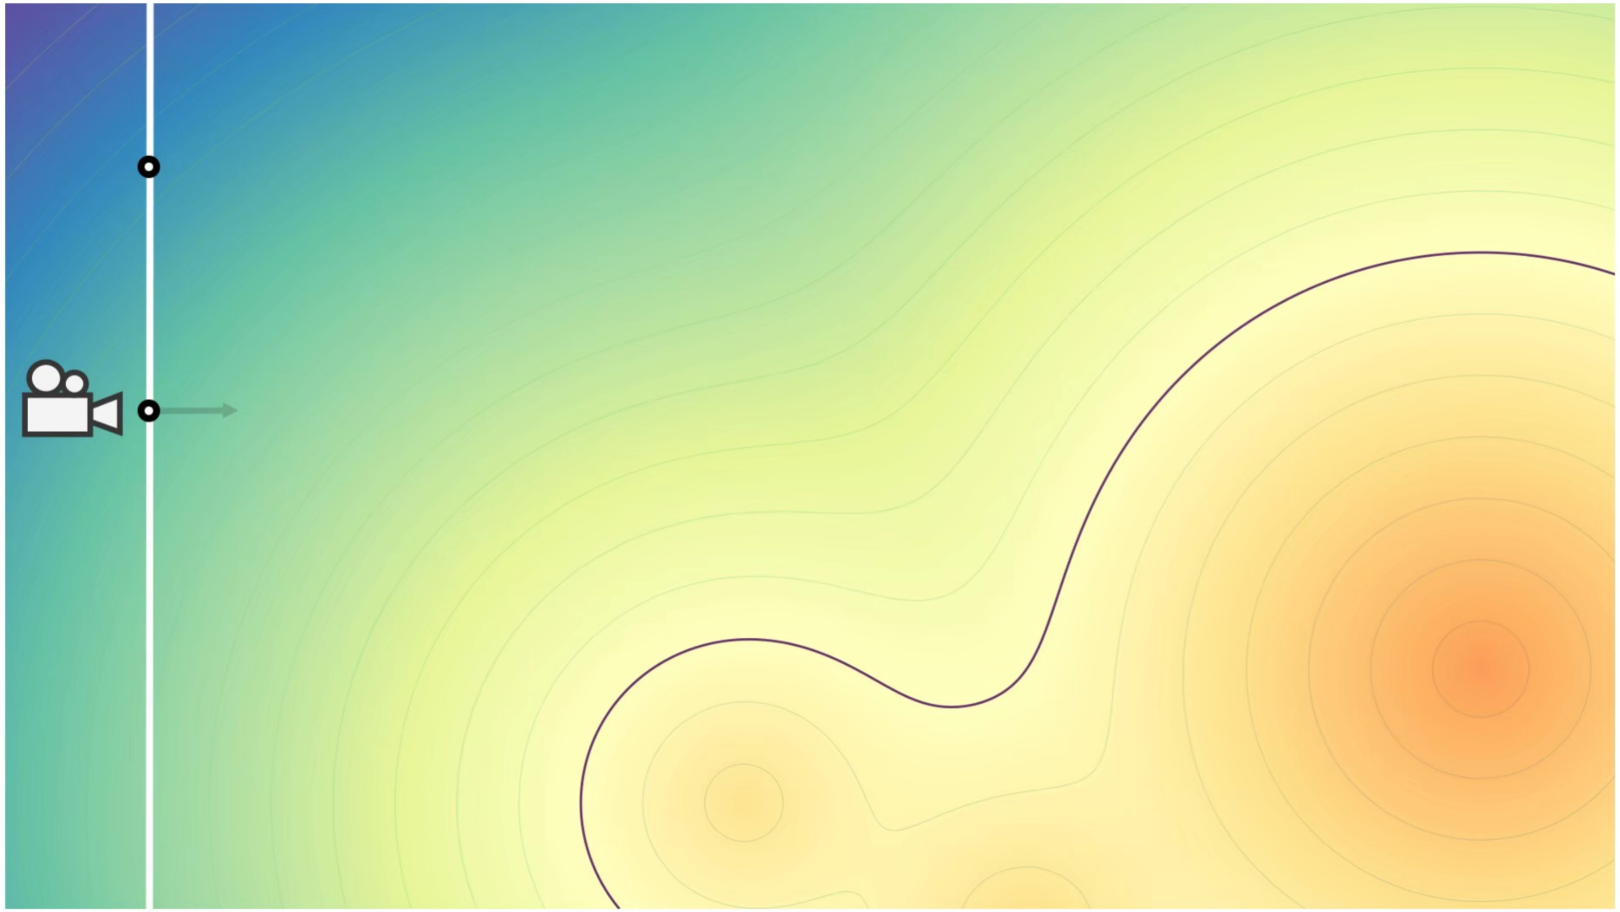

about
ette rays?

Can Compute $\mathcal{V}(u)$ using the Geometry Derivative $G(\mathbf{x})$ of **any** Sphere Tracer point

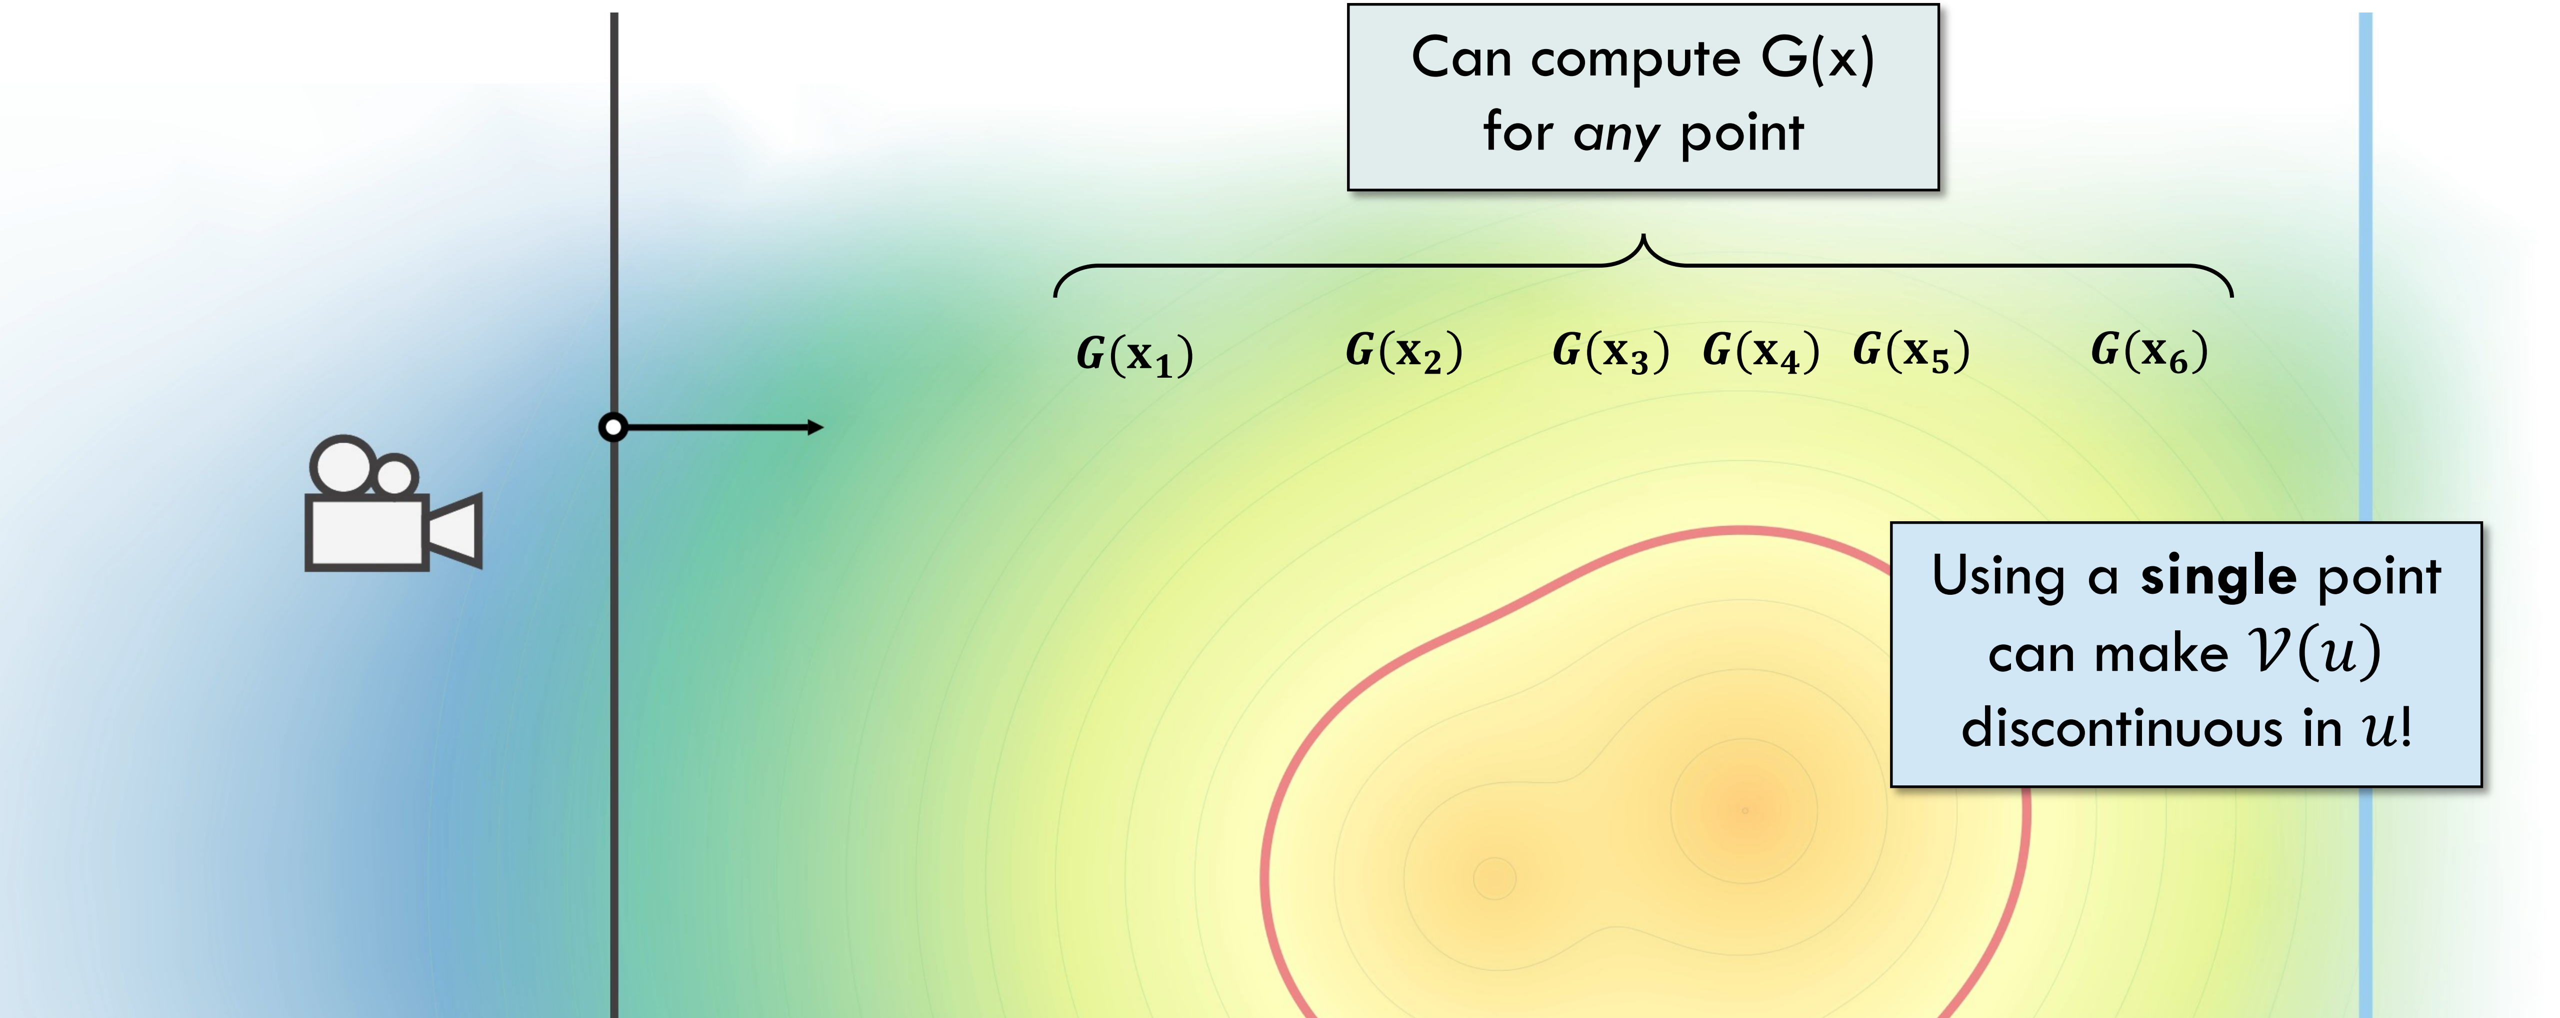

Computing $\mathcal{V}(u)$ as Weighted Mean of $G(\mathbf{x})$ over Sphere Tracer Points

Solution: Compute silhouette weights w(x)

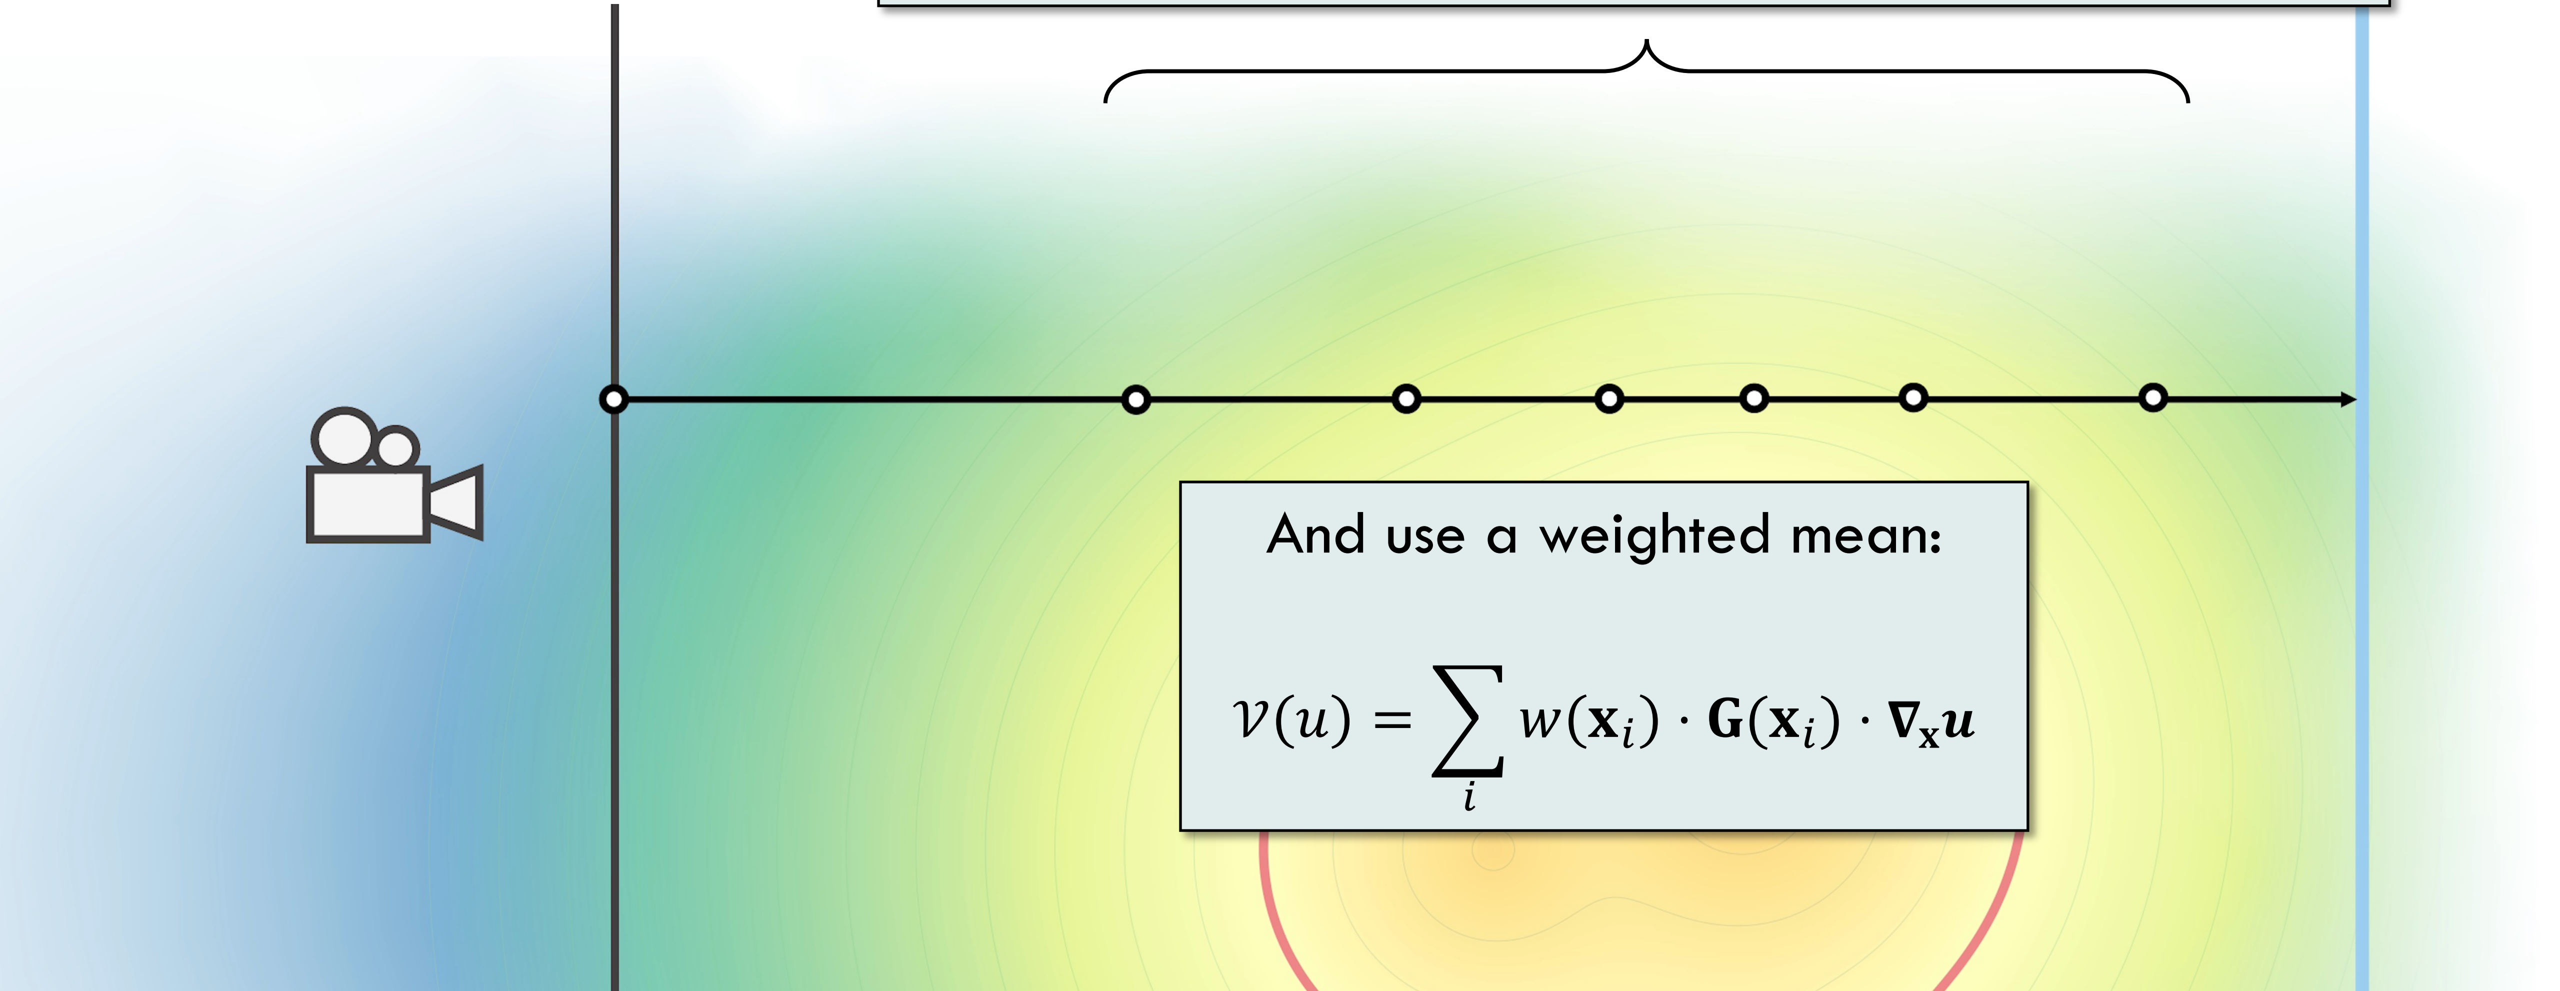

And use a weighted mean:

$$\mathcal{V}(u) = \sum_i w(\mathbf{x}_i) \cdot \mathbf{G}(\mathbf{x}_i) \cdot \nabla_{\mathbf{x}} \boldsymbol{u}$$

Weighted-Mean $\mathcal{V}(u)$ Is Both Consistent And Continuous

✓ Continuous

See **Paper**: Harmonic & Quadrature Weighting

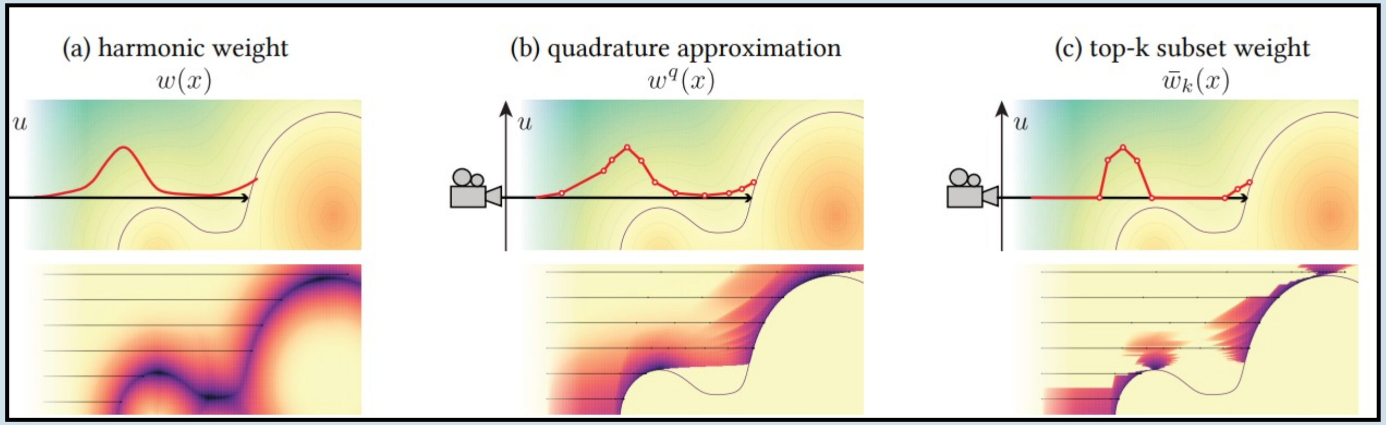

PROJECT PAGE

Scaling Up From Simple 2D To Neural 3D

Simple 2D SDF		Neural 3D SDF
$f(\mathbf{x}, \theta)$	⇨	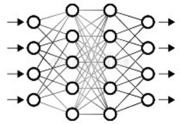
Single Parameter θ	⇨	100,000s of Parameters $\boldsymbol{\theta}$
2D Points $\mathbf{x} : R^2$	⇨	3D Points $\mathbf{x} : R^3$
1D integration coord u	⇨	2D integration coords $\vec{\boldsymbol{u}}$
1D Scalar warp $\mathcal{V}(u)$	⇨	2D Vector warp $\vec{\mathcal{V}}(\vec{\boldsymbol{u}})$
Weight adjustment $1 + \frac{\partial \mathcal{V}}{\partial u}$	⇨	Jacobian determinant $\left\lvert \mathrm{I} + \frac{\partial \vec{\mathcal{V}}}{\partial \vec{\boldsymbol{u}}} \right\rvert$

Putting It All Together: First, Render SDF As Usual

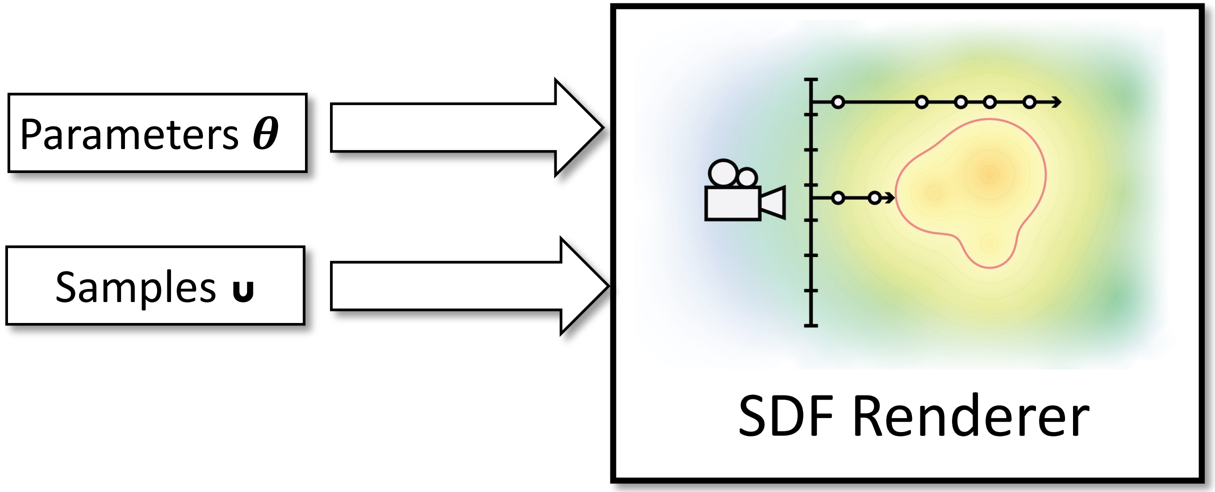

Putting It All Together: Then, Reparameterize Samples

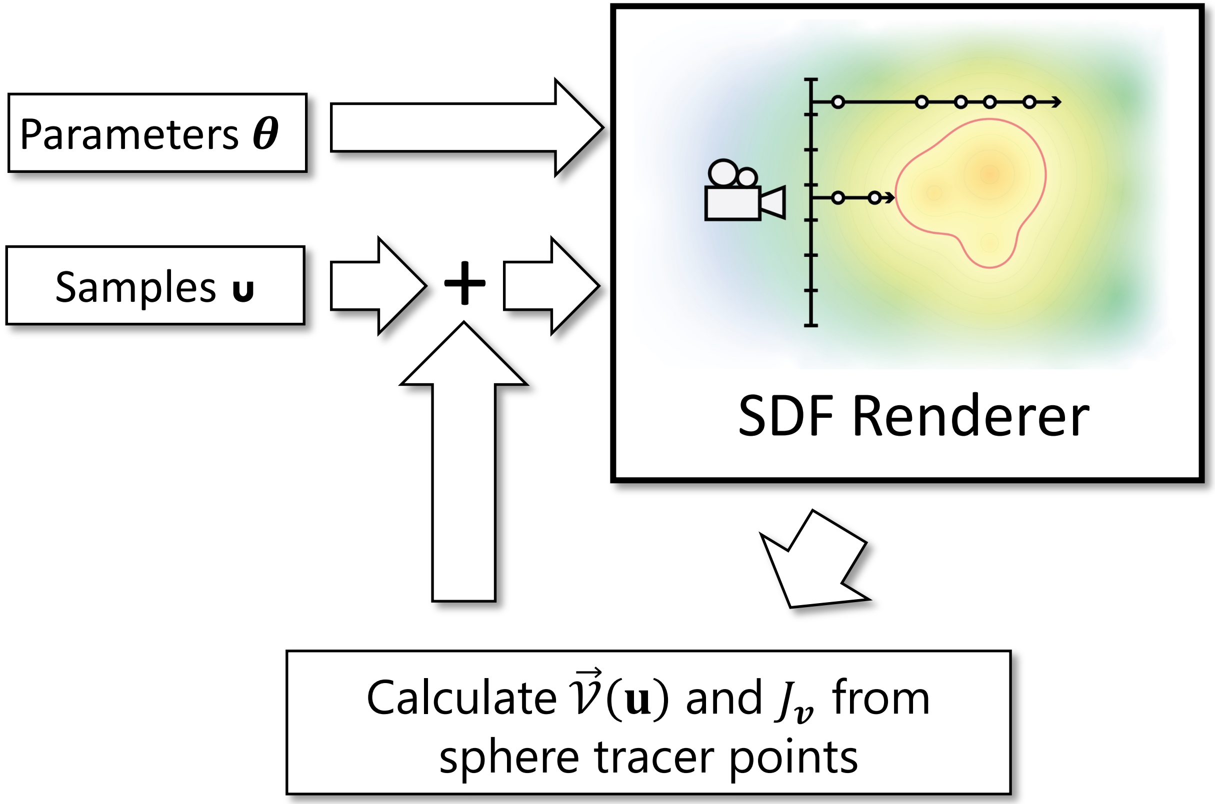

Putting It All Together: Finally, Differentiate With AD

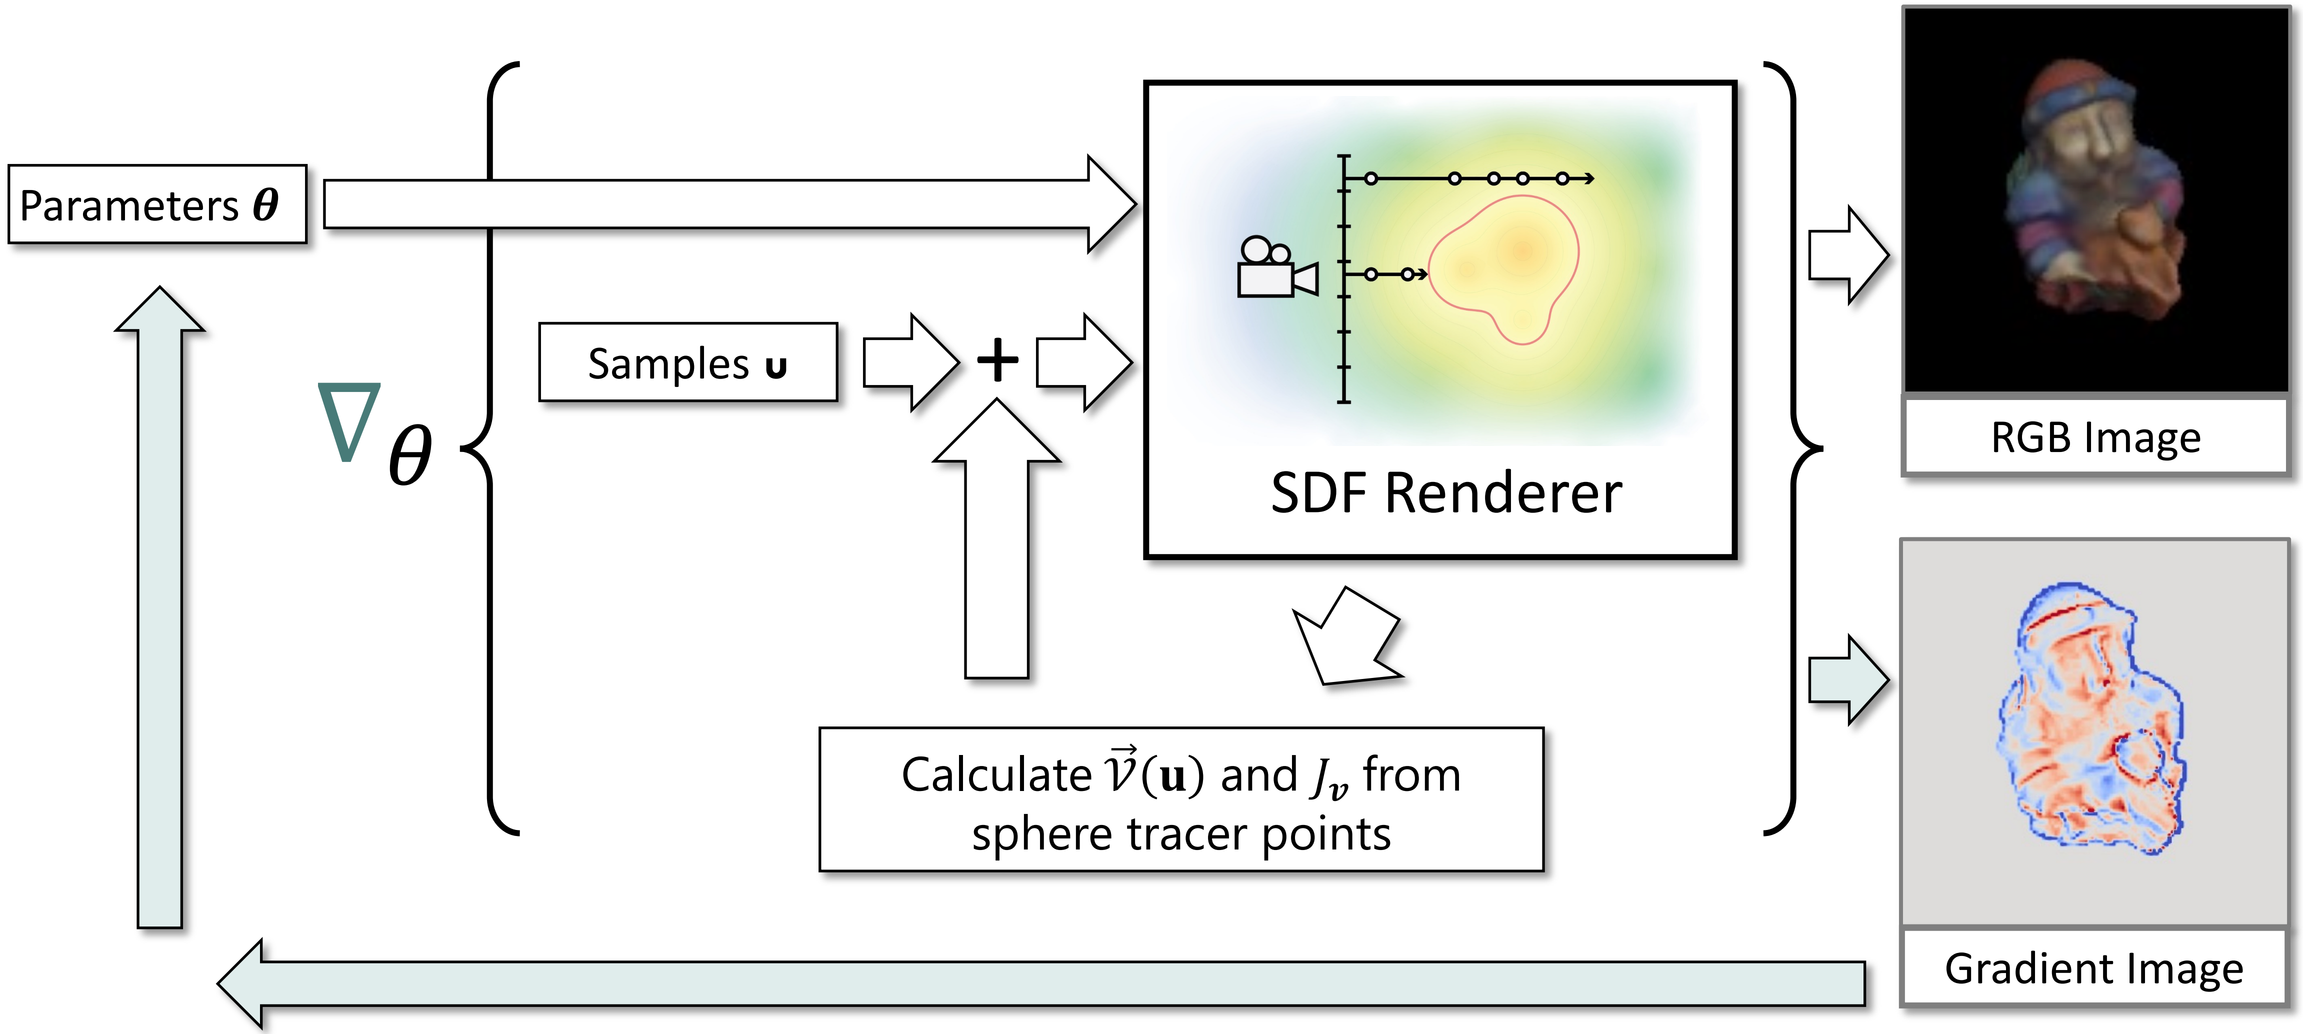

Comparisons Against IDR (Yariv et al. 2020): A Sharp-Surface Model With Segmentation Mask Inputs

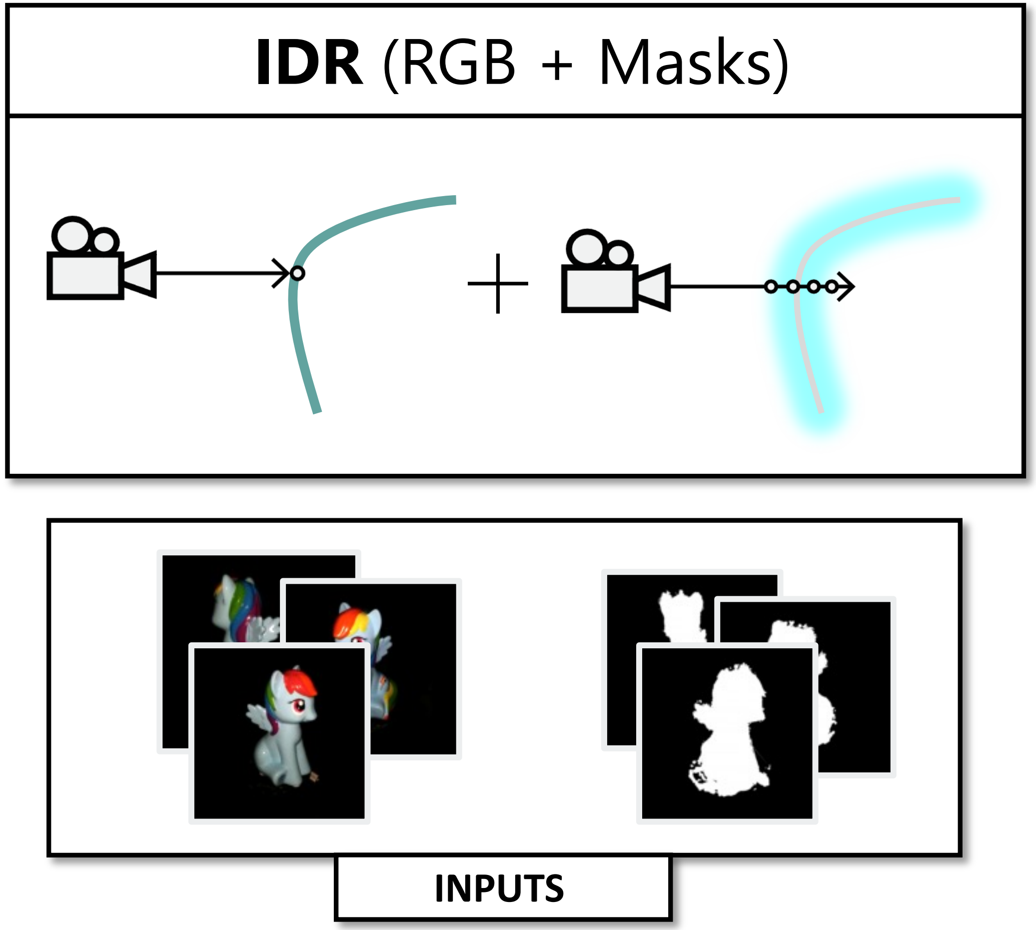

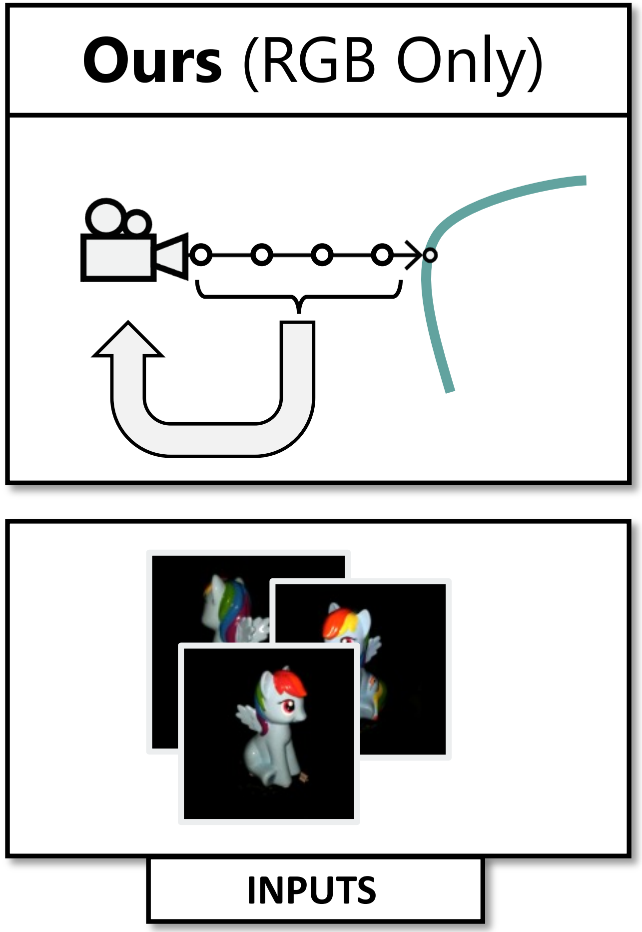

Reconstructions On-Par With IDR *Without* Using Masks

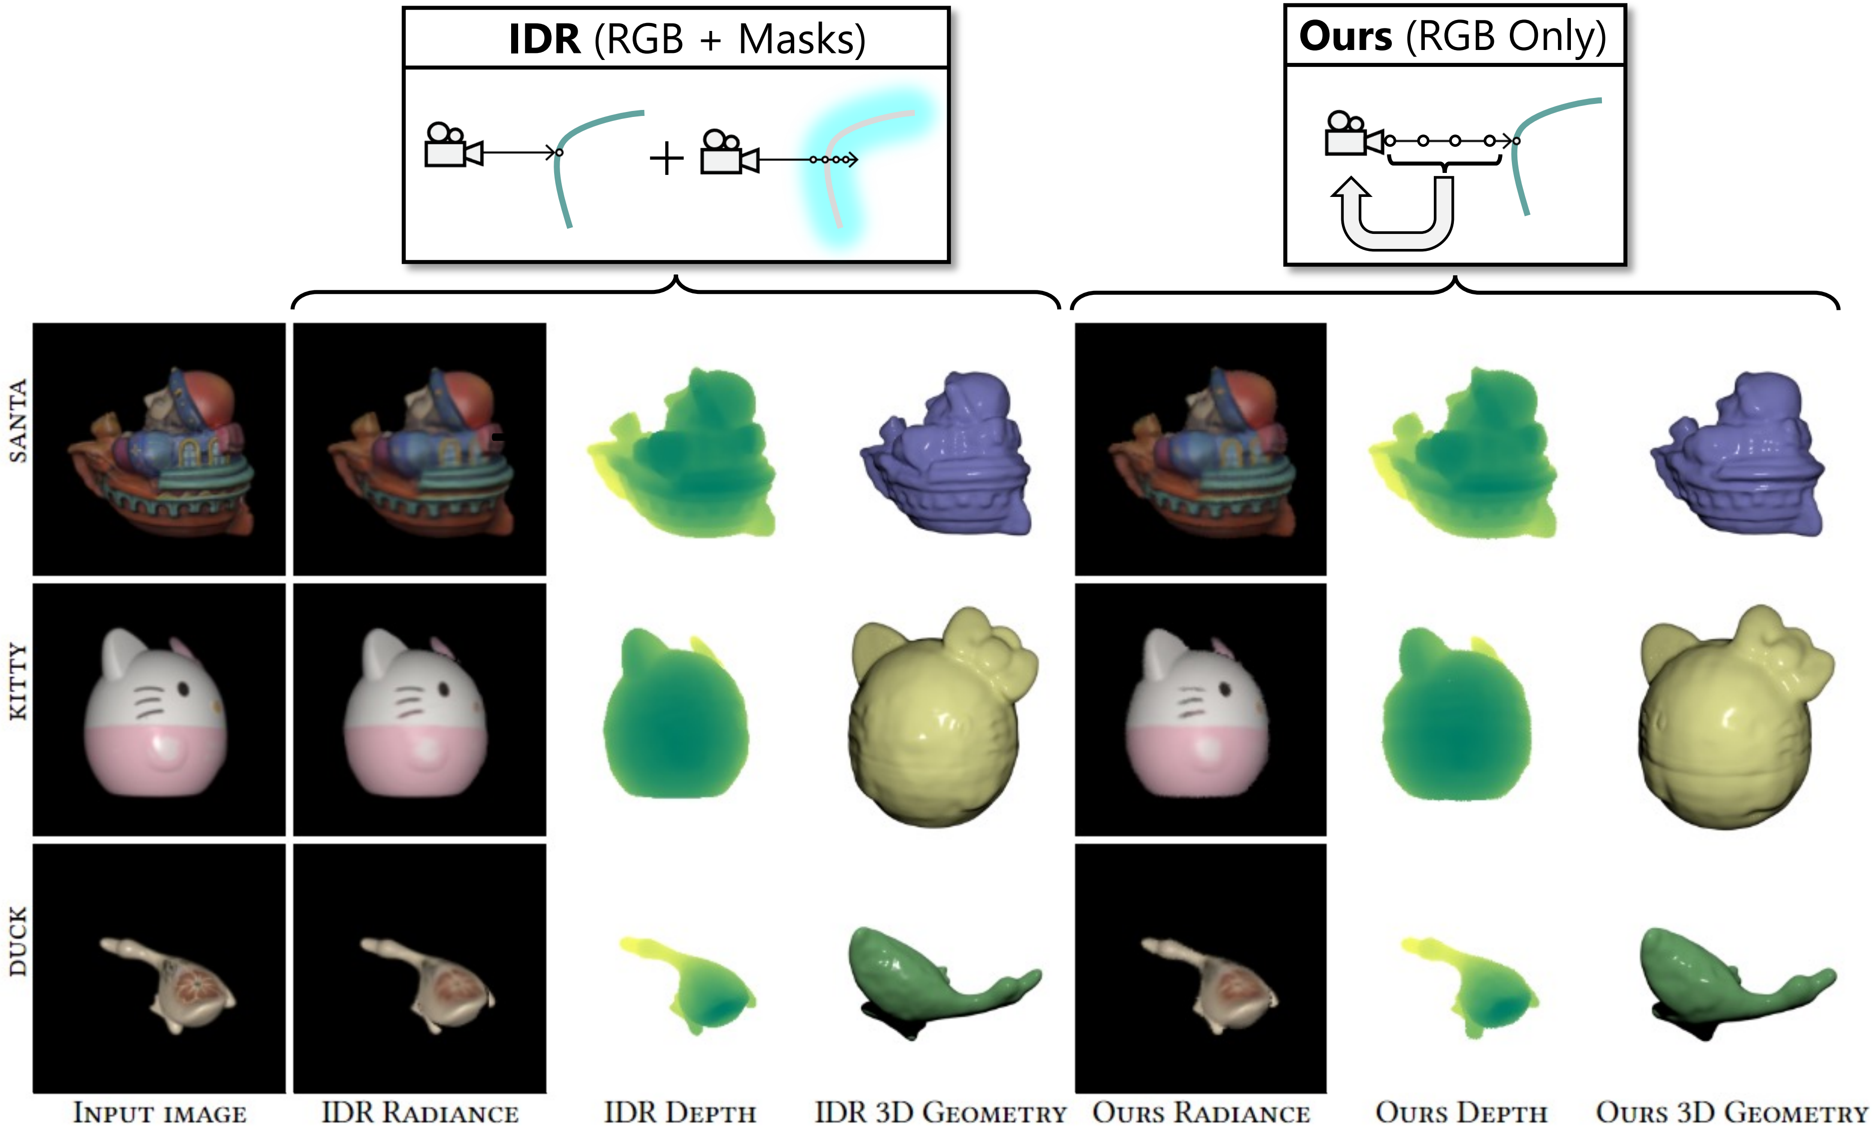

Cleaner Reconstructions Than IDR On Real Data with Poor Segmentation Masks

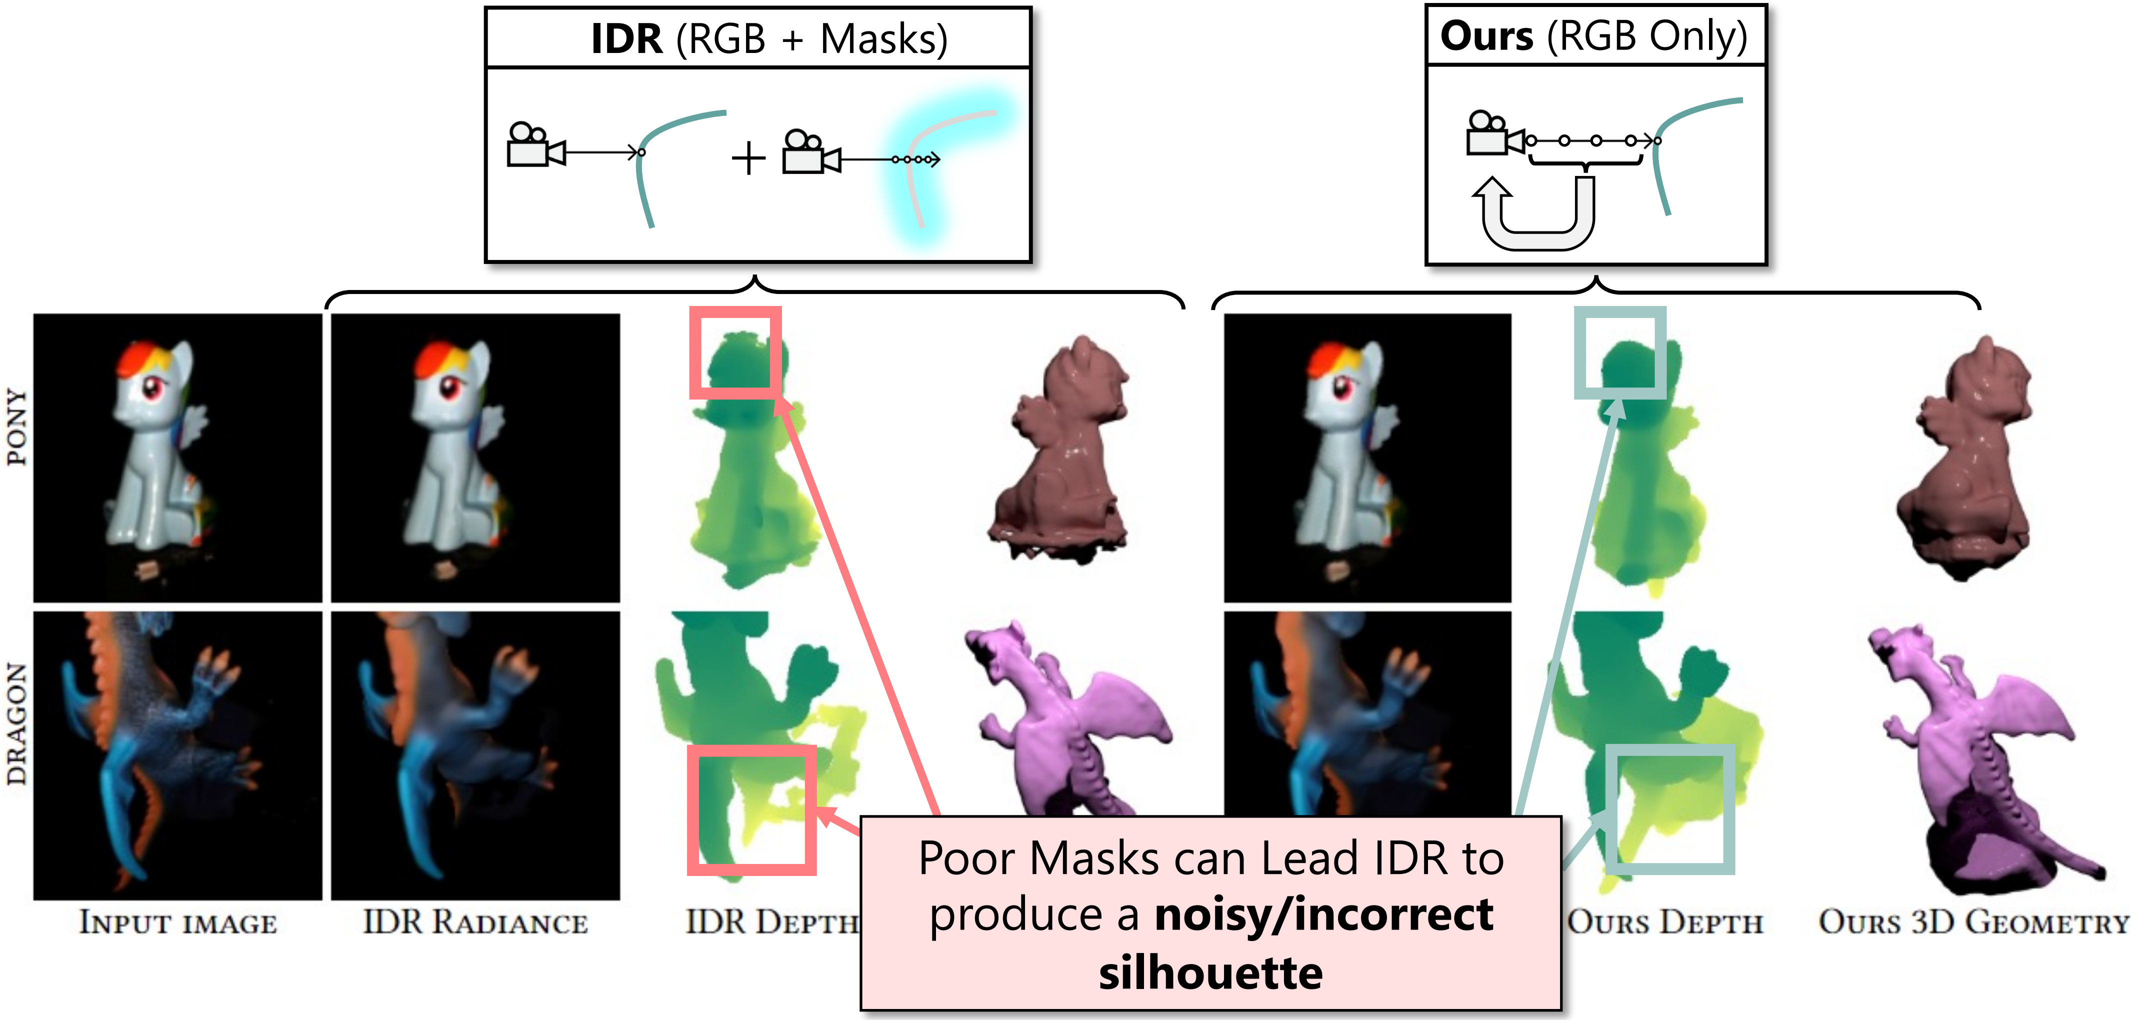

CONCLUSIONS

Problem With **Edge Sampling**	**Warped-Area Sampling**	**Future** Directions
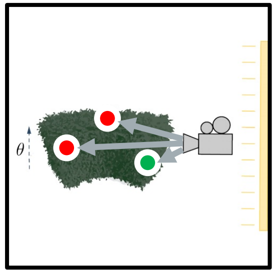 Depth complexity 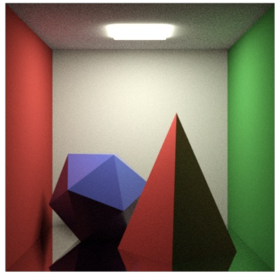 Specularities 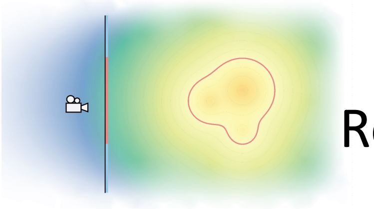 Implicit Representations	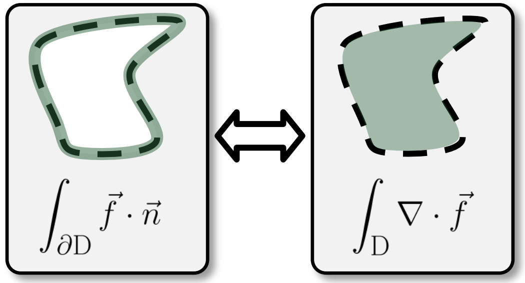 Convert to Area Sampling 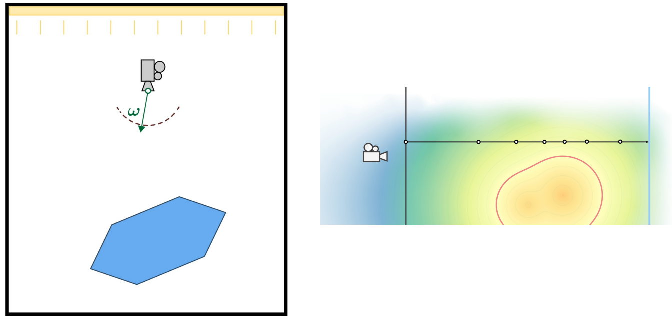 On-the-fly Warp Field Estimation	More SDFs in Physically-based Pipelines 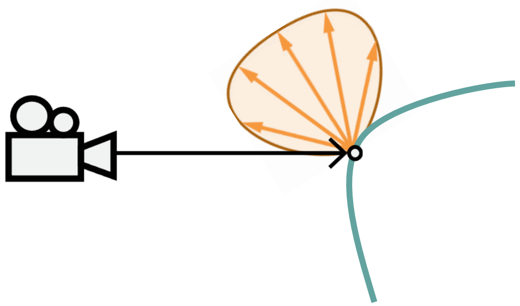 Boundary-Aware Reparameterization For Other Domains