Path Space Warp Field For Differentiable Rendering

Spring 2023

I wrote this manuscript while considering a potential path space natural warp field. However, I later realized that the velocity, so crucial in evaluating the boundary integral, is the velocity of the boundary in the material space, not the velocity of the path in the spatial space. Although the proposed method is problematic, the study on path space velocity and divergence has merits on its own and is still worth reading.

1 Introduction

The *divergence theorem* states that

$$\int_\Omega \nabla \cdot \mathbf{F}(x)\ d\Omega(x) = \int_\Gamma \mathbf{n}(x) \cdot \mathbf{F}(x)\ d\Gamma(x) \tag{1}$$

Here $\mathbf{F}$ is some vector field defined over Ω, Γ is the boundary of Ω, and ∇ computes the divergence. In differentiable rendering, one can use the divergence theorem to rewrite the boundary integral as an area integral, thus solving the difficulty of sampling zero-measure boundaries. However, it is previously unclear what $\mathbf{F}$ should be inside the boundary, and people have come up with clever ways of constructing the *warp field* $\mathbf{V}(x) = \nabla \cdot \mathbf{F}(x)$ as a whole.

This warp field approach to differentiable rendering was first introduced by Bangaru et. al. [1]. Wenzel et. al. [3] later adopted it to SDF representation and greatly fastened the algorithm by resolving the need to trace auxiliary rays. However, up to current literature, one needs to artificially construct the warp field $\mathbf{V}(x)$. Although careful construction can give unbiased results, it is not intuitively straightforward and often prone to high variance. In this work, we make a crucial observation of a natural warp field in the path space.

2 Path Space Warp Field

2.1 Theory

In the *path integral formulation* [2], the recursive integrals of the light transport equation are expanded to a single integral over the path space

$$I = \int_\Omega g(\vec{x})\ d\mu(\vec{x}) \tag{2}$$

where Ω denotes the path space, $\vec{x} = (x_0,\ x_1,\ ...,\ x_k,\ z) \in \Omega$ denotes a light path, $g(\vec{x})$ is the path contribution of $\vec{x}$, and μ is some measure over the path space.

Zhao et. al. [4] later proposed the *differential path integral*

$$\frac{\partial I}{\partial p} = \int_\Omega \frac{\partial}{\partial p} g(\vec{x})\ d\mu(\vec{x}) + \int_\Gamma \langle n(\vec{x}), \frac{\partial \vec{x}}{\partial p} \rangle g(\vec{x})\ d\sigma(\vec{x}) \tag{3}$$

where Γ is the boundary path space, $n(\vec{x})$ is the normal of $\vec{x}$ in Γ and σ is a measure over Γ induced by μ. Note here our expression directly follows from the *Reynold's Transport Theorem* and is equivalent to but slightly different from the expression used by Zhao et. al.

We know that the boundary path space is the boundary of the visible path space, i.e. $\Gamma = \partial\Omega$. This allows us to apply the divergence theorem

$$\int_\Gamma \langle n(\vec{x}), \frac{\partial \vec{x}}{\partial p} \rangle g(\vec{x})\ d\sigma(\vec{x}) = \int_\Omega \nabla \cdot (g(\vec{x}) \frac{\partial \vec{x}}{\partial p})\ d\mu(\vec{x}) \tag{4}$$

2.2 2D Example

Consider a triangle Ω of height p and width 1, where p is the single parameter we want to differentiate. We want to integrate the function $f(x, y) = p$ over the triangle.

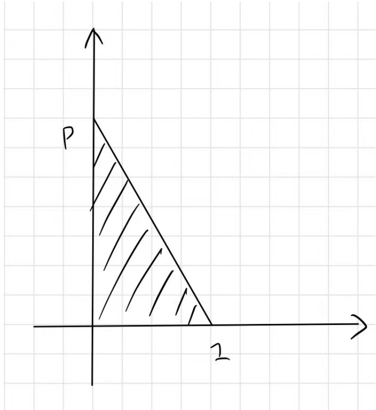

Figure 1: 2D Example

On one hand, we can do a simple calculus exercise and get

$$I = \int_{\Omega} p \; dxdy = \frac{1}{2}p^2 \text{ and } \frac{\partial I}{\partial p} = p$$

On the other hand, we can use the differential path integral and get

$$\begin{aligned}\frac{\partial I}{\partial p} &= \int_{\Omega} 1 \; dxdy + \int_{(1,0)}^{(0,p)} \langle n(s), \frac{\partial s}{\partial p} \rangle p \; ds \\ &= \frac{1}{2}p + \int_{(1,0)}^{(0,p)} \langle n(s), \frac{\partial s}{\partial p} \rangle p \; ds\end{aligned}$$

The second term

$$\begin{aligned}\int_{(1,0)}^{(0,p)} \langle n(s), \frac{\partial s}{\partial p} \rangle p \; ds &= \int_{0}^{1} \langle n(s), \frac{\partial s}{\partial p} \rangle p \frac{ds}{dx} \; dx \\ &= \int_{0}^{1} \langle (\frac{p}{\sqrt{1+p^2}}, \frac{1}{\sqrt{1+p^2}}), (0, 1-x) \rangle p \sqrt{1+p^2} \; dx \\ &= \frac{1}{2}p\end{aligned}$$

This attests to the correctness of evaluating boundary path integral.

However, we can also apply the divergence theorem

$$\begin{aligned}\int_{(1,0)}^{(0,p)} \langle n(s), \frac{\partial s}{\partial p} \rangle p \; ds &= \int_{\Omega} \nabla \cdot (\frac{\partial (x,y)}{\partial p} p) \; dxdy \\ &= \int_{\Omega} \nabla \cdot (0, y) \; dxdy \\ &= \int_{\Omega} 1 \; dxdy \\ &= \frac{1}{2}p\end{aligned}$$

And we get the same result! Here we assume that the triangle stretches linearly with p.

3 Evaluating The Divergence

The above expression asks us to compute three things: (1) the path contribution $g(\vec{x})$, (2) the velocity of the path $v(\vec{x}) = \frac{\partial \vec{x}}{\partial p}$, and (3) the divergence inside path space. The path contribution can be evaluated during forward rendering and is easy to obtain.

3.1 Computing The Velocity

Recall that the path space of length k is the union of the scene manifold for k times

$$\Omega_k = \underbrace{\mathcal{M} \cup \mathcal{M} \cup \ldots \cup \mathcal{M}}_{k \text{ times}} \tag{5}$$

And a path $\vec{x} \in \Omega_k$ can be expressed as $\vec{x} = (x_1, x_2, \ldots, x_k)$, where $x_i \in \mathcal{M}$. Note here x_i refers to an abstract point on the manifold. The velocity is then

$$\frac{\partial \vec{x}}{\partial p} = (\frac{\partial x_1}{\partial p}, \frac{\partial x_2}{\partial p}, \ldots, \frac{\partial x_k}{\partial p}) \tag{6}$$

Thus it suffices to know how to compute the derivative of a point x_i on the manifold $\mathcal{M}$, and this is something that we are very familiar with.

An evolving manifold can be characterized as a homotopy

$$H : \mathcal{M} \times \mathbb{R} \to \mathbb{R}^3$$

where $\mathcal{M}$ is the initial manifold (material space) when the parameter $p = p_0$. $H(\cdot, p)$ describes the manifold when the parameter is p, and $H(x, \cdot)$ describes the trajectory of a point $x \in \mathcal{M}$ as the parameter changes. The velocity of x on this evolving manifold is then $\frac{\partial}{\partial p} H(x, p_0)$. Eventually, we want the velocity of x on $\mathcal{M}$, so we project the velocity of the evolving manifold back onto the manifold. That is,

$$\frac{\partial x}{\partial p} = \frac{\partial H(x, p_0)}{\partial p} - \langle \frac{\partial H(x, 0)}{\partial p}, n(x) \rangle n(x) \tag{7}$$

In practice, say $x = (a, b, c)$ in $\mathbb{R}^3$, then $\frac{\partial}{\partial p} H(x, p_0) = (\frac{\partial a}{\partial p}, \frac{\partial b}{\partial p}, \frac{\partial c}{\partial p})$. The normal is simply the surface normal. In other words, we directly compute the derivative and project it back onto the image plane. We introduce the concept of the evolving manifold here because p might also change $\mathcal{M}$. Note that when $\mathcal{M}$ is independent of p, the derivative must be a tangent vector over the surface and no projection is needed. Finally, the path velocity we discuss here is similar to the scalar normal velocity in [4], albeit we rephrase the language from continuum mechanics to topology.

3.2 Example 1

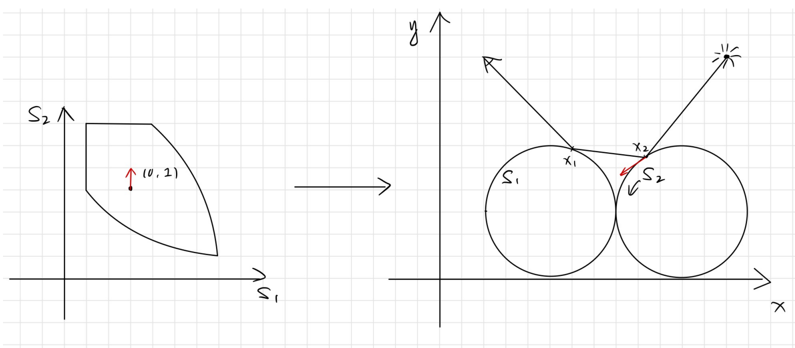

Figure 2: Example of computing path velocity

Consider the 2D scene with 2 spheres $S_1,\ S_2$ touching each other horizontally. We can parameterize each sphere by its radian. The path space of length 2 is then 2-dimensional. Now given a path from the camera to a point $x_1 \in S_1$ to a point $x_2 \in S_2$ to the point light, we can compute analytically its position in the path space. If we assign a velocity of $(0, 1)$ in the path space, this would be rotating S_2 counter-clockwise in $\mathbb{R}^2$ (the image in the wrong direction). x_1 remains static, so the velocity of x_1 is 0; x_2 rotates and has a velocity (v_x, v_y) tangential to S_2. Already we see $(0, 0, v_x, v_y)$ is the velocity of the path. We can further compute the local coordinate of $(0, 0,), (v_x, v_y)$ at x_1, x_2, which is simply the dot product of the velocity vector with the tangent vector. This tells us how to map $(0, 0, v_x, v_y)$ back to the path space and get exactly $(0, 1)$.

3.3 Example 2

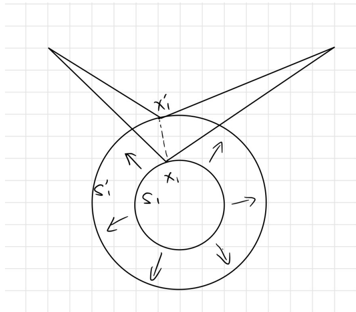

Figure 3: Another Example of computing path velocity

If you are not yet convinced by the above example, consider the scene where a sphere increases its radius with time. A path that goes from the camera to $x_1 \in S_1$ to the point light would have x_1 move along the normal direction at the surface. By our theory, the path velocity should be 0, since the velocity of x_1 in $\mathbb{R}^2$ is orthogonal to the tangent plane. And this is indeed the case: if we map a later point x_1' back to the initial manifold (material space), we see that the point stays at x_1 and does not move at all!

3.4 Computing The Divergence

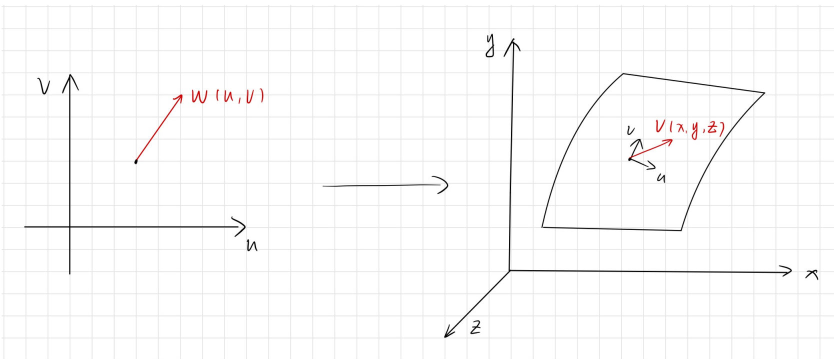

Figure 4: Illustration of local parametrization of manifold

Let us again first simply to the path space of length 1. For any path that goes from the camera to the intersection point x_1 and then to the point light, we can identify the path with the point x_1 on the manifold. That is, $\Omega_1 \cong \mathcal{M}$. From above we now have a vector field $f(x_1)v(x_1)$ spread over the manifold, and we want to compute the divergence of this vector field.

We know that given a vector field V in $\mathbb{R}^3$, the divergence of the vector field is

$$\nabla \cdot V(x, y, z) = \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y} + \frac{\partial V_z}{\partial z} \tag{8}$$

where V_x denotes the x-component of $V(x, y, z)$. However, it is not immediately clear how to compute the divergence of a vector over a surface in a time-efficient way. Since divergence is a local property, a natural idea is to compute a local parametrization of the manifold and maps the vector field to the flat $\mathbb{R}^2$. Suppose $W(u,v)$ is the vector field mapped to the uv-coordinate, then the divergence

$$\nabla \cdot W(u,v) = \frac{\partial W_u}{\partial u} + \frac{\partial W_v}{\partial v} \tag{9}$$

From differential geometry we know this uv-coordinate actually describes the tangent plane, so $W(u,v)$ under the uv-coordinate describes the same tangent vector as $V(x,y,z)$ in $\mathbb{R}^3$. We can also observe that $\frac{\partial W_u}{\partial u}$ takes the directional derivative of $W_u(u,v)$ in the direction of the u-axis. Thus we can perform the same computation in $\mathbb{R}^3$ by taking the dot product between $V(x,y,z)$ and the u vector to get the W_u and then taking the directional derivative of the dot product in the direction of the u vector. In other words,

$$\nabla \cdot V = \frac{\partial}{\partial u}\langle V, u\rangle + \frac{\partial}{\partial v}\langle V, v\rangle \tag{10}$$

It is known that divergence does not depend on the choice of coordinates, so the final value does not depend on our choice of the u, v vector.

We now derive how to generalize to path space of higher length. A crucial observation is the divergence is associative. We know $\mathbb{R}^3 = \mathbb{R}\times\mathbb{R}\times\mathbb{R}$, so we can decompose $V(x_0, y_0, z_0)$ into three vectors $V_x(x_0)$, $V_y(y_0)$, $V_z(z_0)$. $V_x(x_0)$ lies in the subspace $\mathbb{R} \times \{y_0\} \times \{z_0\}$. Then we have

$$\begin{aligned}
\nabla_{\mathbb{R}^3} \cdot V(x_0, y_0, z_0) &= \frac{\partial V_x}{\partial x} + \frac{\partial V_y}{\partial y} + \frac{\partial V_z}{\partial z} && (11)\\
&= \nabla_{\mathbb{R}\times\{y_0\}\times\{z_0\}} \cdot V_x(x_0) + \nabla_{\{x_0\}\times\mathbb{R}\times\{z_0\}} \cdot V_y(y_0) + \nabla_{\{x_0\}\times\{y_0\}\times\mathbb{R}} \cdot V_z(z_0) && (12)
\end{aligned}$$

Recall that Ω_k is the self-union of $\mathcal{M}$ for k times. Thus given a path $\vec{x} = (x_1, x_2, \ldots, x_k)$, we can apply the associative law and get

$$\nabla_{\Omega_k} \cdot v(\vec{x}) = \sum_i \nabla_{\mathcal{M}} \cdot v(x_i) \tag{13}$$

And this brings us back to the simple case of Ω_1 we discussed above!

3.5 Example 3

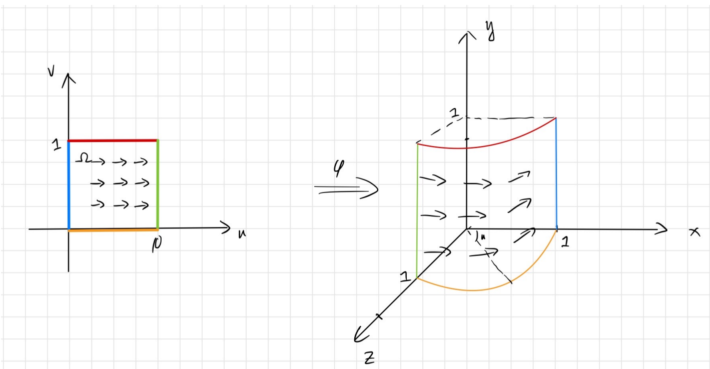

Figure 5: Example of computing surface divergence

Consider a cylinder $\mathcal{S}^2 \times \mathbb{R}$ and a parametrization

$$\varphi(u,v) = (\cos u, v, \sin u)$$

Suppose we want to compute the divergence of a uniform vector field $(1,0)$ in the uv-coordinate, but all we have is a vector field V over the cylinder. Then given a point (x_0, y_0, z_0) on the cylinder, the first we can do is to compute an orthonormal basis of the tangent plane at (x_0, y_0, z_0). Say that we pick u to align with $V(x_0, y_0, z_0)$ and v to be parallel to the y-axis. Then

$$\langle V, u\rangle = 1 \text{ and } \langle V, v\rangle = 0$$

Now if we slide the orthonormal basis along the cylinder, then for all points of the same height $\langle V, u\rangle = 1$ and $\langle V, v\rangle = 1$. Thus

$$\frac{\partial}{\partial u}\langle V, u\rangle = \frac{\partial}{\partial v}\langle V, v\rangle = 0$$

And this matches the divergence we would compute in the uv-coordinate!

3.6 Example 4

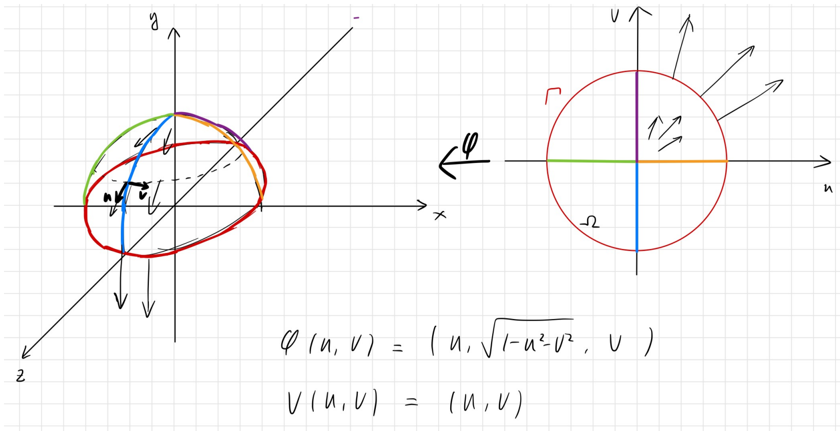

Figure 6: Another example of computing the divergence

Consider another example where we have a hemisphere in $\mathbb{R}^3$ and a parametrization by projecting points on the hemisphere onto the xz-plane.

$$\varphi(u, v) = (u, \sqrt{1 - u^2 - v^2}, v)$$

Now we assign a vector field $V(u, v) = (u, v)$ in the uv-coordinate. We can first compute the boundary integral

$$\begin{aligned}
\int_\Gamma \langle V(s), n(s)\rangle \; d\Gamma(s) &= \int_0^{2\pi} \langle V(\cos\theta, \sin\theta), n(\cos\theta, \sin\theta)\rangle \; d\theta \\
&= \int_0^{2\pi} 1 \; d\theta \\
&= 2\pi
\end{aligned}$$

We can compute the divergence

$$\nabla \cdot V(u, v) = \frac{\partial u}{\partial u} + \frac{\partial v}{\partial v} = 2$$

and apply the divergence theorem in the uv-coordinate

$$\int_\Omega \nabla \cdot V(x) \; d\Omega(x) = \int_\Omega 2 \; d\Omega(x) = 2\pi$$

Finally, we can map the vector field V onto the hemisphere and compute in $\mathbb{R}^3$. The vectors point downward, and their magnitude increases as they approach the xz-plane. Given an arbitrary point x, we can have $v(x)$ as the tangent vector of the latitude and $u(x)$ as the tangent vector of the longitude. For all points on the same latitude, note $\langle V(x), v(x)\rangle = 0$ and thus $\frac{\partial}{\partial v}\langle V(x), v(x)\rangle = 0$. For all points on the same longitude, note their magnitude and thus $\langle V(x), u(x)\rangle$ increases linearly from 0 at $(0, 1, 0)$ to 1 at the xz-plane. Thus $\frac{\partial}{\partial u}\langle V(x), u(x)\rangle = 1$.

Integrate over the hemisphere and we get

$$\int_{\mathcal{H}^2} \left(\frac{\partial}{\partial v}\langle V(x), v(x)\rangle + \frac{\partial}{\partial u}\langle V(x), u(x)\rangle\right) \; dA(x) = \int_{\mathcal{H}^2} 1 \; dA(x) = 2\pi$$

And we see once again we obtain the same result in all three methods!

3.7 Summary

Let us state a brief summary of how to compute the divergence—

1. Given a path $\vec{x} = (x_1, x_2, \ldots, x_k)$, compute derivatives $\frac{\partial x_i}{\partial p}$, where p denotes the parameter.
2. Project $\frac{\partial x_i}{\partial p}$ onto the tangent plane at x_i to get $v(x_i)$, and the path velocity is $v(\vec{x}) = (v(x_1), v(x_2), \ldots, v(x_k))$
3. Compute an orthonormal basis $u,\ v$ at x_i
4. Compute $\nabla \cdot f(\vec{x})v(x_i) = \frac{\partial}{\partial u}\langle f(\vec{x})v(x_i), u\rangle + \frac{\partial}{\partial v}\langle f(\vec{x})v(x_i), v\rangle$
5. Return $\nabla \cdot f(\vec{x})v(\vec{x}) = \sum_i \nabla \cdot f(\vec{x})v(x_i)$