Report
[COV877]

Reparameterization Discontinuous Integrands for Differentiable Rendering

Instructor:

Rahul Narain

narain@cse.iitd.ac.in

By:

Rajeev Gupta

2019CS51018

Introduction:

The paper "Reparameterizing Discontinuous Integrands for Differentiable Rendering" presents an innovative method aimed at addressing a significant challenge in the field of differentiable rendering - handling discontinuous integrands. In this report, we'll delve into the method introduced in the paper, highlighting its distinctive features in comparison to previous approaches. Additionally, we'll explore potential limitations of this method.

Method Introduced:

The central method introduced in the paper revolves around the idea of reparameterization using kernels, with a primary focus on spherical convolutions.

1. **Sampling the Monte Carlo Samples:** In the traditional rendering process, Monte Carlo integration is used to estimate complex integrals involving various parameters. However, when dealing with non-differentiable functions, computing gradients directly with respect to these parameters becomes problematic. Attempts where:
    - Calculating visibility of triangles analytically.
    - Blurring the discontinuity.

    Satisfying solution so far:

    - Samples around discontinuities (edge samples).
    - Require sampling Silhouette edges in the integrands.
        - Problematic when we have complex geometry or non-trivial visibility.
2. **Leveraging the Integral:** To overcome this limitation, the paper proposes a novel approach. Instead of calculating gradients with respect to the parameters of interest directly, the integral itself is utilized as a "bridge" between the parameters and the Monte Carlo samples. This is a critical insight as the integral, while non-differentiable with respect to a parameter, remains unchanged regardless of how it's computed.

3. **Moving Discontinuities:** By calculating the Monte Carlo samples with respect to the integral, the discontinuities in the integrand are effectively made to move in relation to the Monte Carlo samples. This dynamic adjustment prevents visible discontinuities during the rendering process. In essence, the method reparametrize the problem so that the discontinuities are no longer stationary with respect to the Monte Carlo samples.

    $$I = \int_{\mathcal{X}} f(x, \Theta)\, \mathrm{d}x, \quad \text{Or} \quad I \approx E = \frac{1}{N} \sum_{i=1}^{N} \frac{f(x_i, \theta)}{p(x_i)},$$

    But when sampling techniques depend on a scene parameter θ:

$$I \approx E = \frac{1}{N} \sum \frac{f(x_i(\theta), \theta)}{p(x_i(\theta), \theta)}.$$

Differentiating Monte Carlo estimates and a probability density function p(x) :

$$\frac{\partial I}{\partial \theta} \approx \frac{\partial E}{\partial \theta} = \frac{1}{N} \sum \frac{\partial}{\partial \theta} \frac{f(x_i(\theta), \theta)}{p(x_i(\theta), \theta)}.$$

Here, it is often the case that the original integrand f (x, θ) is non-differentiable in θ due to visibility changes.

In the proposed method we can overcome this problem using a change of variables that removes the discontinuity of the integrand in θ. Consider it as a transformation T : Y → X, the re-parameterized integral becomes.

$$\int_{\mathcal{X}} f(x, \theta)\, \mathrm{d}x = \int_{\mathcal{Y}} f(T(y, \theta), \theta)\, |\det J_T|\, \mathrm{d}y \quad = \int_{\mathcal{Y}} \mathbb{1}_{y>0}\, k(y + \theta)\, \mathrm{d}y.$$

where:

$$f(x) = \mathbb{1}_{x>\theta}\, k(x), \;\; \text{with} \int_{\mathcal{X}} k(x)\, \mathrm{d}x = 1.$$

we obtain:

$$\frac{\partial I}{\partial \theta} \approx \frac{1}{N} \sum \frac{\partial}{\partial \theta} \frac{\mathbb{1}_{y_i>0}\, k(y_i + \theta)}{p(y_i)}.$$

4. **Differentiable Sampling:** Additionally, the paper emphasizes the importance of making the sampling procedures differentiable. This is achieved through the use of kernels and fixed parameters, ensuring that gradients can be accurately computed even when dealing with non-differentiable integrands.

Unlike conventional rendering techniques, which struggle to compute gradients accurately when integrands contain sudden discontinuities, this approach employs kernels to smooth out these irregularities.

Differences from Previous Methods:

1. **Avoidance of Edge Sampling**: The paper eliminates the need for explicit sampling of silhouette edges, which is a common approach in previous methods for handling visibility changes at discontinuities. By doing so, the authors assert that their technique simplifies the rendering process and avoids the complexity associated with edge sampling.
2. **Spherical Convolutions:** One of the key distinctions of this method is the utilization of spherical convolutions. By convolving the integrand with a kernel on the unit sphere, it becomes possible to handle the non-differentiable behavior of functions defined on the sphere. This is in stark contrast to previous methods that often employed ad-hoc techniques for handling such discontinuities.
3. **Differentiable Sampling:** The paper emphasizes the importance of making sampling procedures differentiable. While prior approaches focused on non-differentiable sampling methods, this method carefully constructs the sampling process using kernels and fixed parameters. This ensures that gradients can be computed accurately, even in the presence of non-differentiable integrands.

Limitations of the Method:

While the paper presents a promising approach to address non-differentiability, it's essential to consider potential limitations:

1. **Kernel Selection:** The choice of an appropriate kernel can significantly impact the effectiveness of the method. Selecting the right kernel may require domain-specific knowledge and could be challenging for novice users.
2. **Robustness to Scene Complexity and parameter sensitivity:** The paper mainly focuses on relatively controlled scenarios. It remains to be seen how well the method performs in highly complex scenes with multiple interacting objects, complex lighting, and intricate materials. Also for some scenes, the parameters may have a stronger impact on the integrand's discontinuities than others.
3. **Generalization:** The proposed method is tailored for specific integrals with known characteristics. Generalizing it to a broader range of rendering scenarios may be a non-trivial task.

Conclusion:

In conclusion, the paper introduces an intriguing approach to address the challenge of non-differentiable integrands in differentiable rendering. By incorporating spherical convolutions and emphasizing differentiable sampling, it sets itself apart from previous methods. However, like any novel technique, it comes with its own set of potential limitations, which should be carefully considered when applying it in practical rendering scenarios.

1. **Reparameterization Technique:** The authors introduce a reparameterization technique that transforms non-differentiable integrals into differentiable forms. This approach

involves careful changes of variables that eliminate the dependence of discontinuities on scene parameters, enabling differentiation under the integral sign.

2. **Differentiable Rendering:** A novel differentiable rendering algorithm is proposed, which applies the reparameterization technique to handle integrals with previously problematic discontinuities. The method does not require explicit sampling of discontinuities, making it efficient and applicable to various scenes.
3. **Variance Reduction:** The paper presents a variance reduction technique for gradient estimation, enhancing the robustness of the differentiable rendering approach, even in complex scenes with high geometric complexity.
4. **Implementation:** An open-source implementation of the proposed method is provided as part of the Mitsuba renderer, making it accessible to the computer graphics community.