

Neural Sequence Transformation

Sabyasachi Mukherjee¹ Sayan Mukherjee² Binh-Son Hua^{3,4} Nobuyuki Umetani¹ Daniel Meister¹




Background

Physically Based Rendering



Background

Monte Carlo Integration




We want to find

$$\int_a^b f(x) dx$$

Background


Monte Carlo Integration



1. Draw n samples randomly in (a, b)

Background

Monte Carlo Integration



2. Calculate

$$\hat{I} = \frac{1}{n} \sum_{i=1}^n f(x_i)$$

$$\rightarrow \int_a^b f(x) dx \text{ as } n \rightarrow \infty$$

Disadvantages of Monte Carlo Integration

- Converges at a slow rate of $O(\sqrt{n})$
- We propose to improve convergence using [Sequence Transformation](#)

Sequence Transformation

- A mapping \mathcal{T} from a sequence (s_n) to another sequence (t_n)

$$\begin{array}{c} \left(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \frac{1}{6}, \dots \right) \\ \downarrow \mathcal{T} := t_n = s_n^2 \\ \left(1, \frac{1}{4}, \frac{1}{9}, \frac{1}{16}, \frac{1}{25}, \frac{1}{36}, \dots \right) \end{array}$$

Background

Sequence Transformation

Example: Aitken's Δ^2 Process

$$t_n = s_n - \frac{(s_{n+1} - s_n)^2}{s_{n+2} - 2s_{n+1} + s_n}$$

Background

$$I = \int_0^1 f(x) dx = \int_0^1 e^{-x^2} dx \approx 0.746824$$

$$\hat{I}_1 = 0.991071$$

$$\hat{I}_2 = 0.990075$$

0.874699

... 0.757849

... 0.767839

... ...

$$\hat{I}_{16} = 0.746318$$

Motivation

$$I = \int_0^1 f(x)dx = \int_0^1 e^{-x^2} dx \approx 0.746824$$

0.991071 , 0.990075 , 0.874699 , 0.791553 , 0.767839 , ... , 0.746318

- Monte Carlo integration → a sequence of terms converging to I
- Apply sequence transformation methods to Monte Carlo integration

Related Work:

$a_n g_n$ transformation [BZ91]

$$S_n = \frac{1}{n} \sum_{i=1}^n f(x_i)$$

- Transformation:

$$T_{n+1} = S_n - \frac{S_{n+1} - S_n}{g_{n+1} - g_n} g_n$$

- g_n is arbitrary function of n
- However, we show (T_n) has slightly *higher variance* than (S_n) regardless of choice of g_n

Related Work:

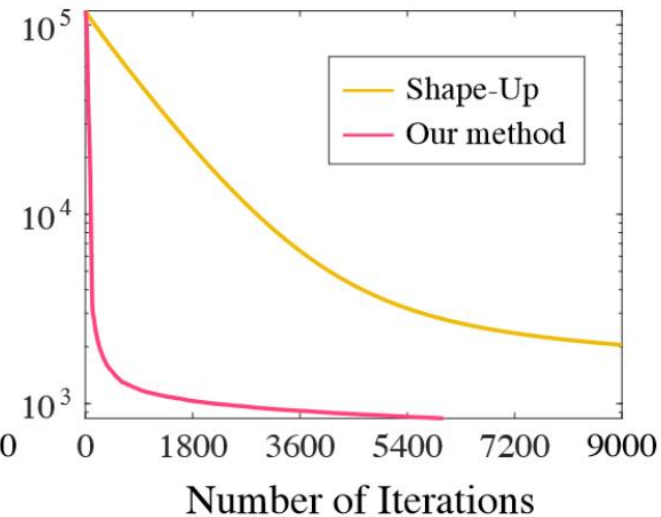

Using Anderson Acceleration [PDZ*18]

- Showed improvements in geometry processing and physics simulation using Anderson Acceleration [And65]
- However, applicable to fixed point methods
- Monte Carlo integration is **difficult to formulate as one**



#V: 58752000

#F: 58736640



Our Method


Background:

Analysis of Monte Carlo convergence

- There is **no universal sequence transformation** method for all sequences [DGB82]
- Determine *type of convergence* before applying sequence transformation methods

Type of Convergence [Wen89]

$$\rho = \lim_{n \rightarrow \infty} \left| \frac{s_{n+1} - s}{s_n - s} \right|$$



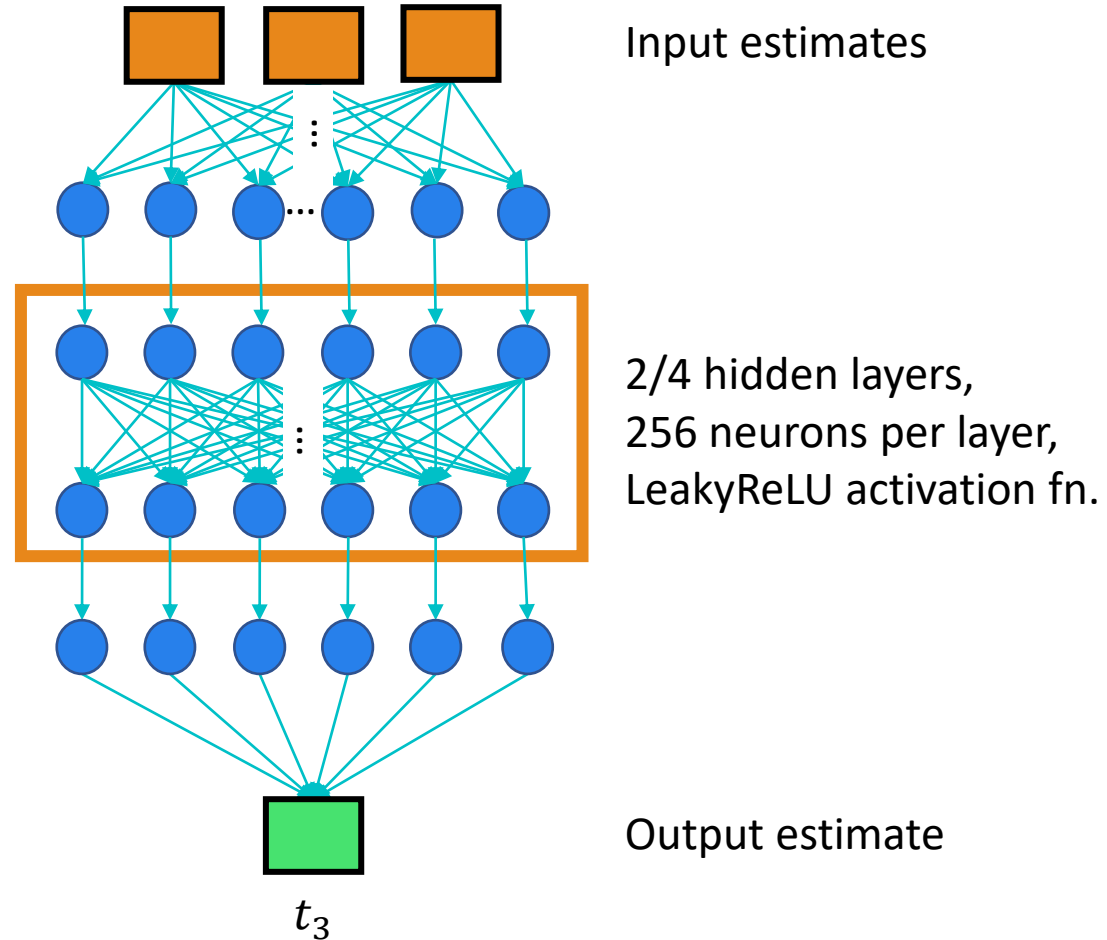

Contribution 1: Analysis of Monte Carlo convergence

- *Type of convergence* is defined for deterministic sequences only
- We show that **Monte Carlo estimates converge like a logarithmic sequence**



Contribution 2: A data-driven approach to learn sequence transformation

- A **simple MLP architecture** is proposed
 - Pass in sequence values in a **sliding window fashion**
 - Output of the network is a single value per sliding window



Our Method




Our Method



Our Method



Our Method




Contribution 3: Loss function

- We propose a novel loss function tailored to Monte Carlo integration
- Output sequence requirements:
 - Must have **lower error** than input sequence
 - Should **converge at a faster rate** than input sequence

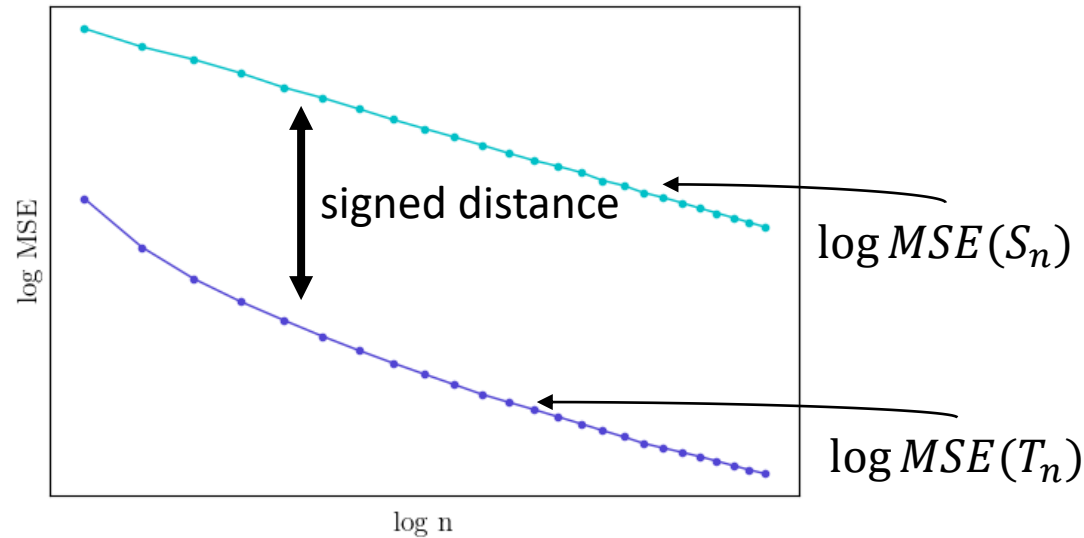
Our Method

Convergence Graphs

- Plot $\log(\text{MSE})$ vs. $\log(\text{Sample Count})$



Loss Function Requirement 1: Lower Error



- Minimize the total **signed distance** between the output and input sequences:

$$\log(\text{MSE}(T_n)) - \log(\text{MSE}(S_n))$$

Loss Function Requirement 2: Faster Convergence


Faster convergence :=

Better “slope” of $\log MSE(T_n)$


$$\log MSE(T_n) \approx a \log n + b$$

↑
slope

↑
intercept



Loss Function Requirement 2: Better Slope



Minimize $a \log n + b$, n small \Rightarrow
 b dominates


Want optimizer to make a more
negative

Proposal: minimize $a + \frac{b}{\log n}$ instead

\Rightarrow Minimize $\frac{\log \text{MSE}(T_n)}{\log n}$

(Since $\log \text{MSE}(T_n) \approx a \log n + b$)

Loss Function Requirement 2: Better Slope



Proposed loss function:

$$\mathcal{L}_n = \frac{\log \text{MSE}(T_n)}{\log n} - \log \text{MSE}(S_n)$$

Our Method


A Quick Recap

- Monte Carlo integration has **slow convergence** and **can be viewed as a sequence**
- We consider a **data-driven approach** to sequence transformation to improve it
- We **design a neural network to learn sequence transformation**
- Now we **apply our method** to 1D integrals and images

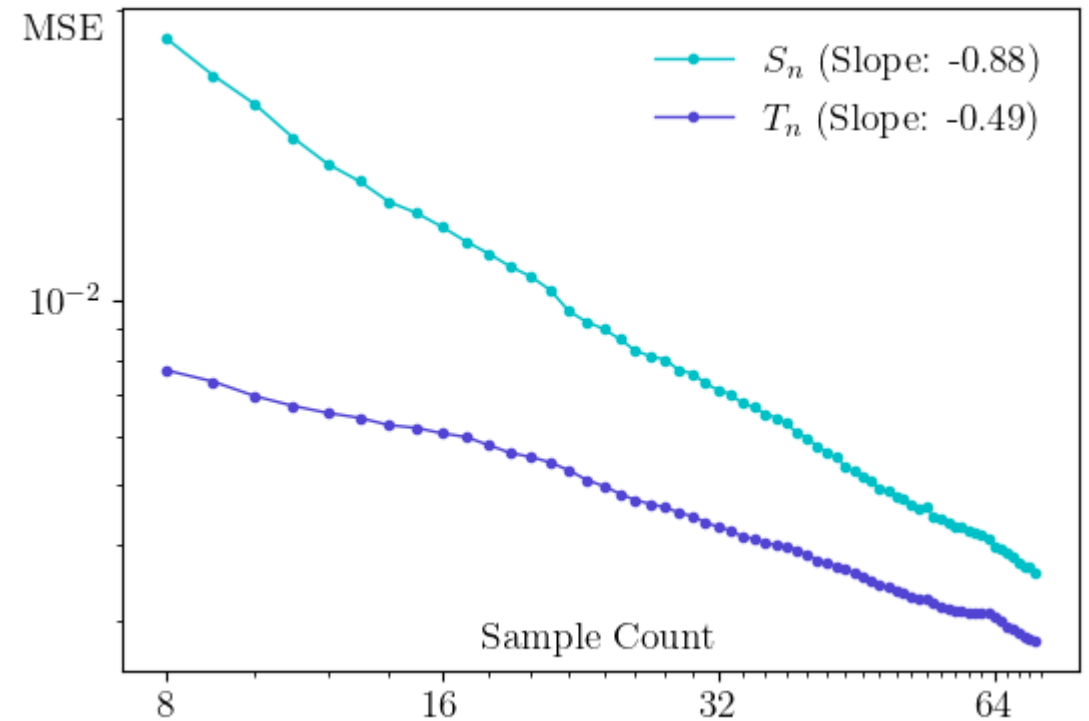
Results

Results

1D Result: Step Function




Shape of the Step function




Convergence graph of Step function

Results


Results: Diffuse Teapot Scene



Input, 8 samples per pixel



Output, 8 samples per pixel



Reference, 8192 samples per pixel

Results

Results: Diffuse Teapot Scene



Input, 8 samples per pixel




Output, 8 samples per pixel



Reference, 8192 samples per pixel


Results: Diffuse Teapot Scene




Convergence graph of Diffuse Teapot scene

Results


Results: Glossy Teapot Scene



Input, 8 samples per pixel




Output, 8 samples per pixel



Reference, 8192 samples per pixel

Results

Results: Glossy Teapot Scene



Input, 8 samples per pixel

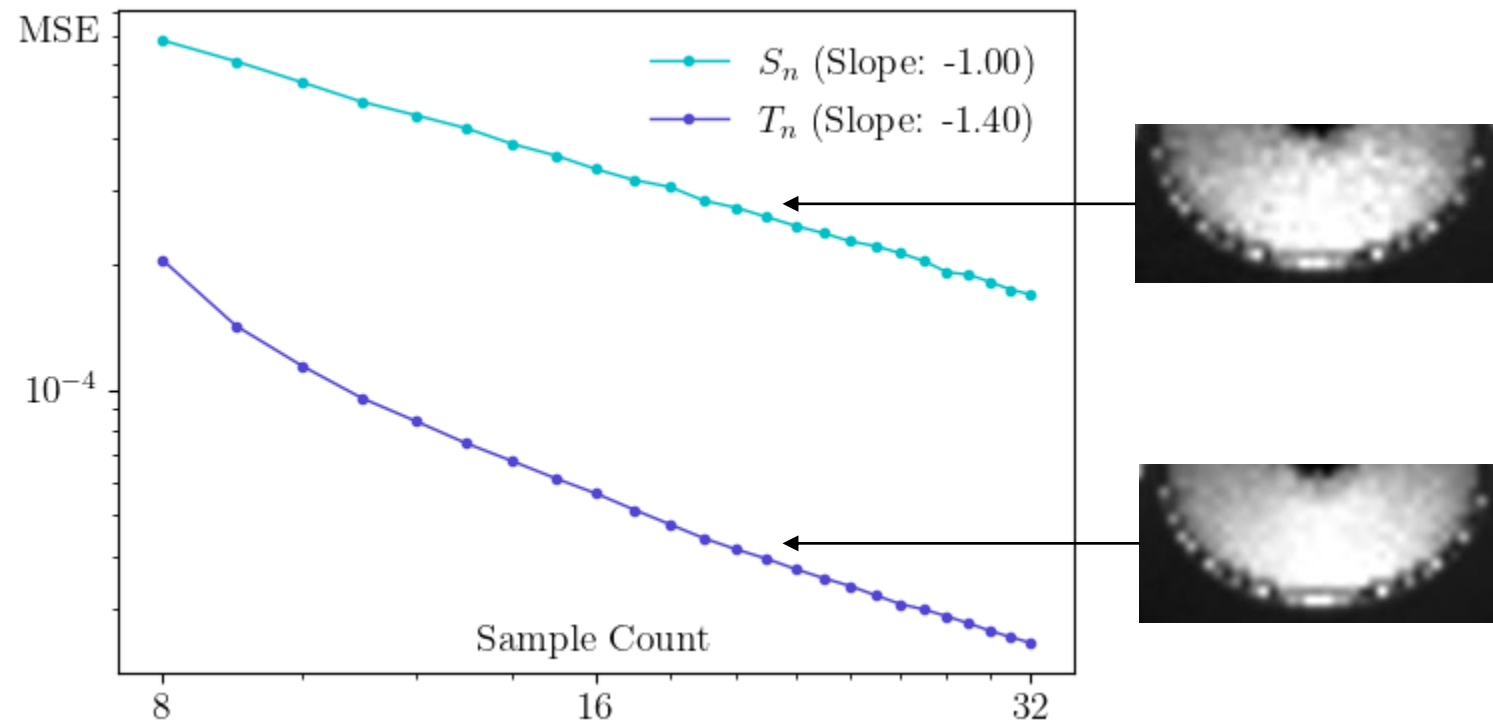


Output, 8 samples per pixel



Reference, 8192 samples per pixel


Results: Glossy Teapot Scene




Convergence graph of Glossy Teapot scene

Results


Results with Denoiser: Cornell Box Scene



Input, 8 samples per pixel




Output, 8 samples per pixel




Denoised output, 8 samples per pixel

Results


Results with Denoiser: Cornell Box Scene



Input, 8 samples per pixel




Output, 8 samples per pixel



Denoised output, 8 samples per pixel

Results with Denoiser: Cornell Box Scene



Convergence graph of Cornell Box scene

Summary and Future Work

- Monte Carlo estimates converge close to logarithmic rate.
- Proposed a data-driven neural network approach to learn a sequence transformation that can improve Monte Carlo integration.
- Proposed a custom loss function tailored to Monte Carlo integration.
- Obtained improvements for both 1D integrals and rendered images.

- Make our method real-time
- Explore possibilities of application along with analytical sequence transformations

Neural Sequence Transformation

Sabyasachi Mukherjee¹ Sayan Mukherjee² Binh-Son Hua^{3,4} Nobuyuki Umetani¹ Daniel Meister¹

