Eurographics2019

Analysis of Sample Correlations for Monte Carlo Rendering

Gurprit Singh, Cengiz Oztireli, Abdalla G. Ahmed, David Coeurjolly

Kartic Subr, Oliver Deussen, Victor Ostromoukhov, Ravi Ramamoorthi, Wojciech Jarosz

Good morning everyone, thank you for being here. This is a joint survey work done with many people from the rendering and sampling community [CLICK]

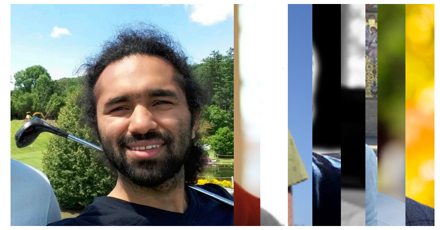

Good morning everyone, thank you for being here. This is a joint survey work done with many people from the rendering and sampling community [CLICK]

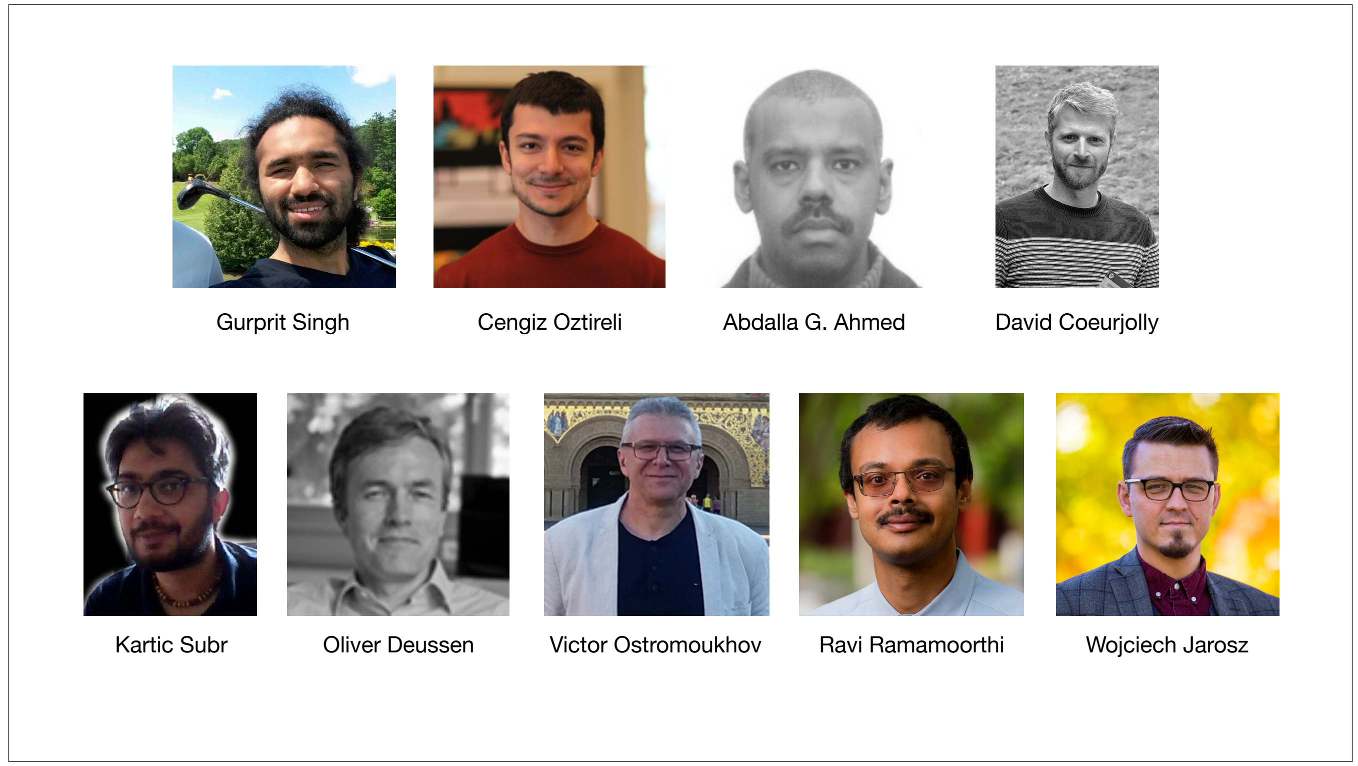

Good morning everyone, thank you for being here. This is a joint survey work done with many people from the rendering and sampling community [CLICK]

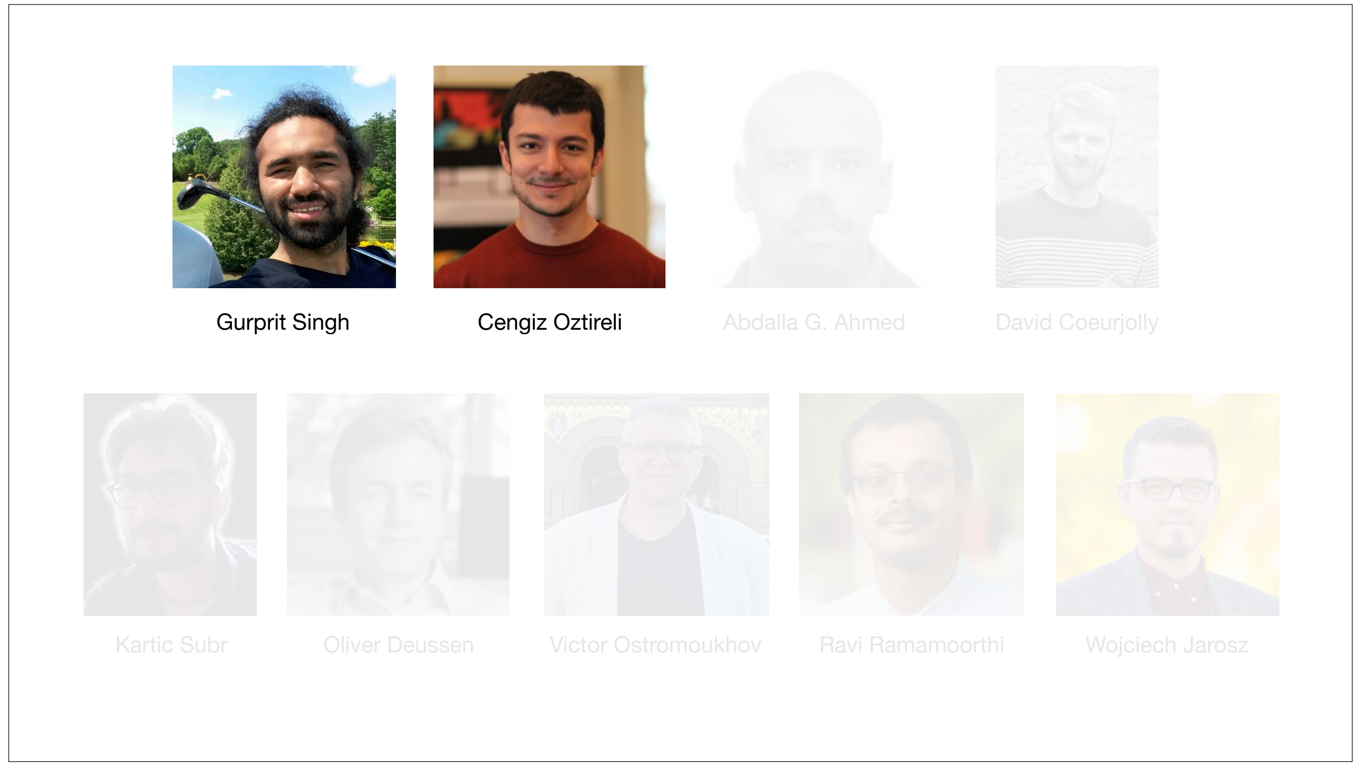

Good morning everyone, thank you for being here. This is a joint survey work done with many people from the rendering and sampling community [CLICK]

Rendering = Geometry + Radiometry

Geometry / Projection

for pin-hole model is known since 400BC

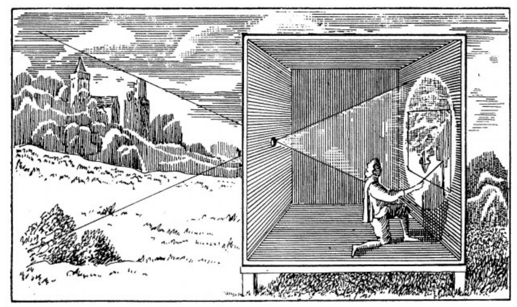

The idea of projecting real world on a 2D surface has a long history, where a pin-hole model allows projection of real world onto a screen (or a wall).

Rendering = Geometry + Radiometry

Geometry / Projection
for pin-hole model is known since 400BC

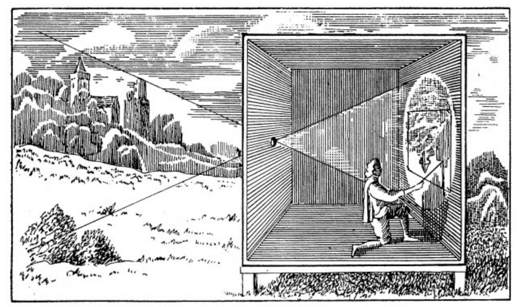

Radiometrically accurate simulation
is importance of realism

However, adding radiometric entities to a geometry is equally important to simulate realism

Rendering = Geometry + Radiometry

Geometry / Projection

for pin-hole model is known since 400BC

Radiometrically accurate simulation

is importance of realism

Many rendering algorithms are developed using this pin-hole camera.

Radiometric fidelity improves photorealism

Papas et al. [2013]

To show the relevance of photometric accuracy, here is one example where one of the object is real and the other one is fabricated.

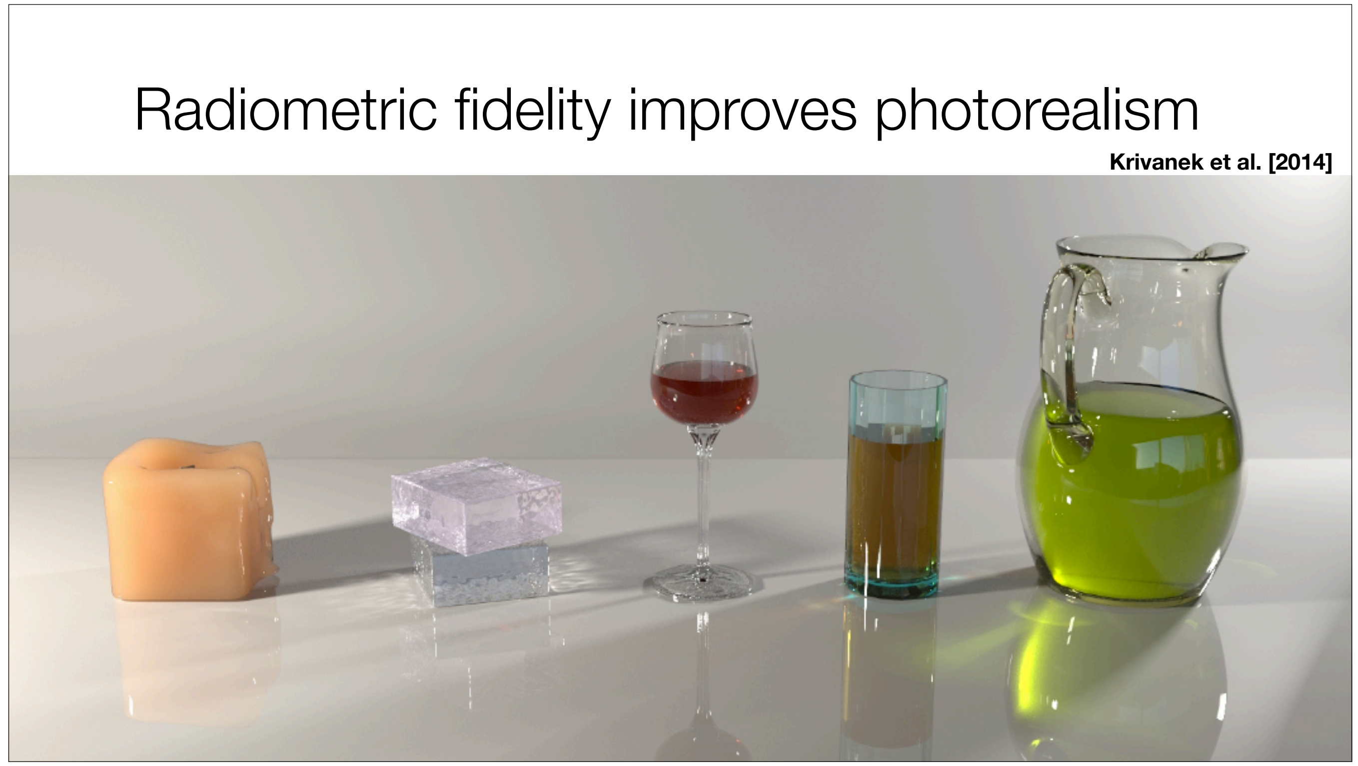

The fidelity of the virtual scene becomes unquestionable.

Reconstruction: Estimate image samples

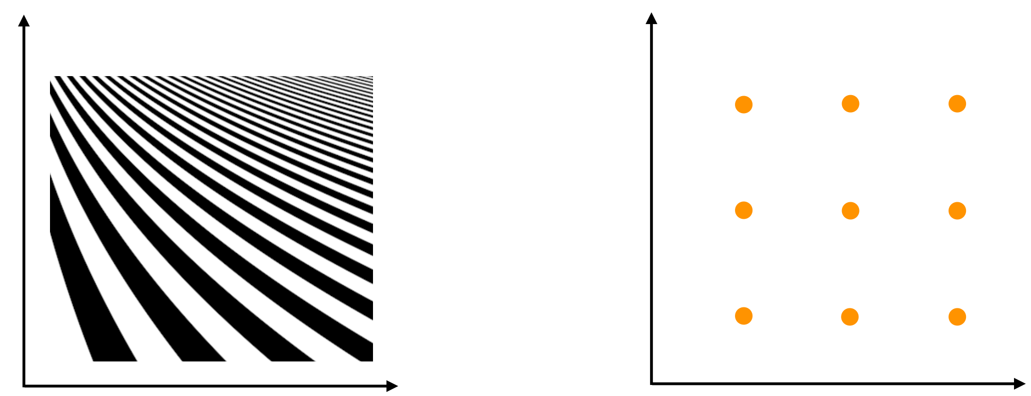

We start from the very basic. Let's start by looking at reconstruction. On one side, we are looking at a simple function with black and white stripes and on the right side we are looking at the samples (or pixel centers) of a image where we want to reproduce this function.

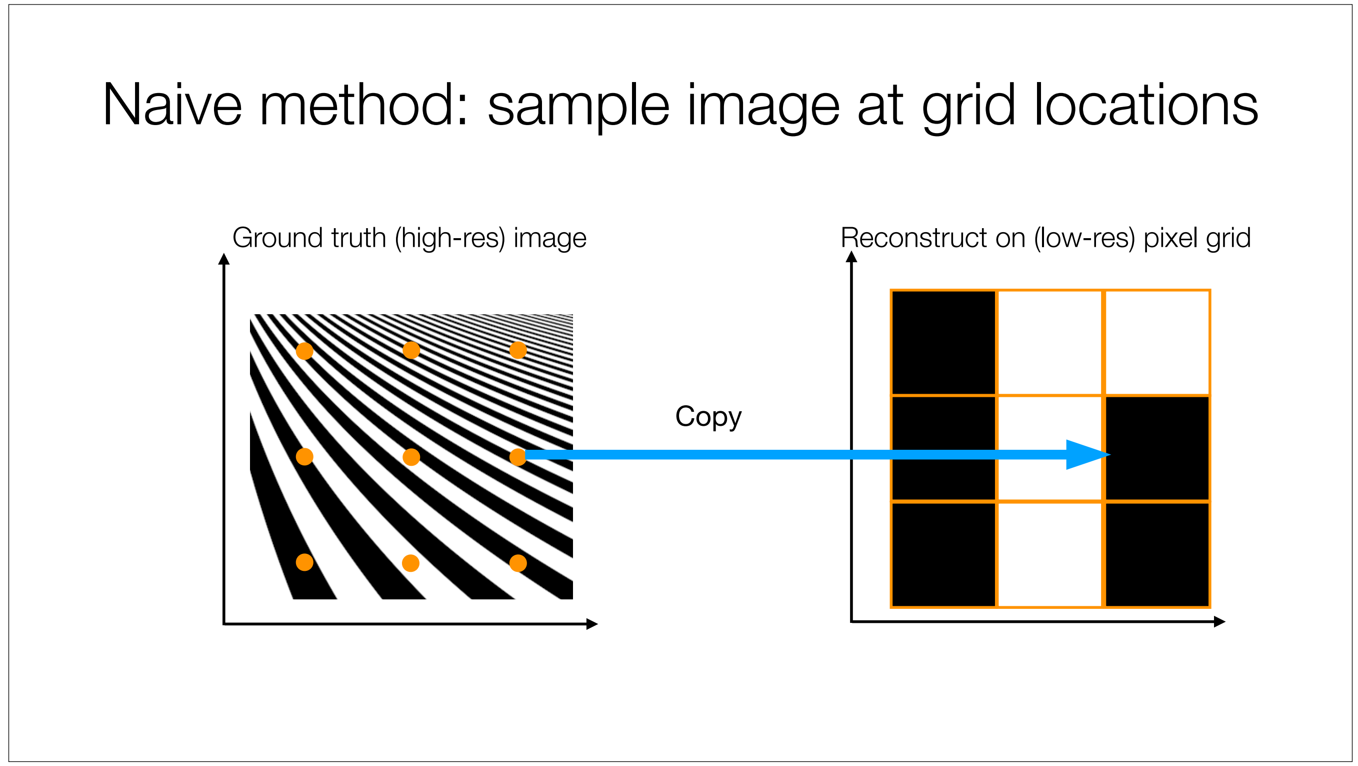

Naive approach goes by simply "copying" the values of the underlying function values to the pixels. This is bad as it gives...

Naive method: sample image at grid locations

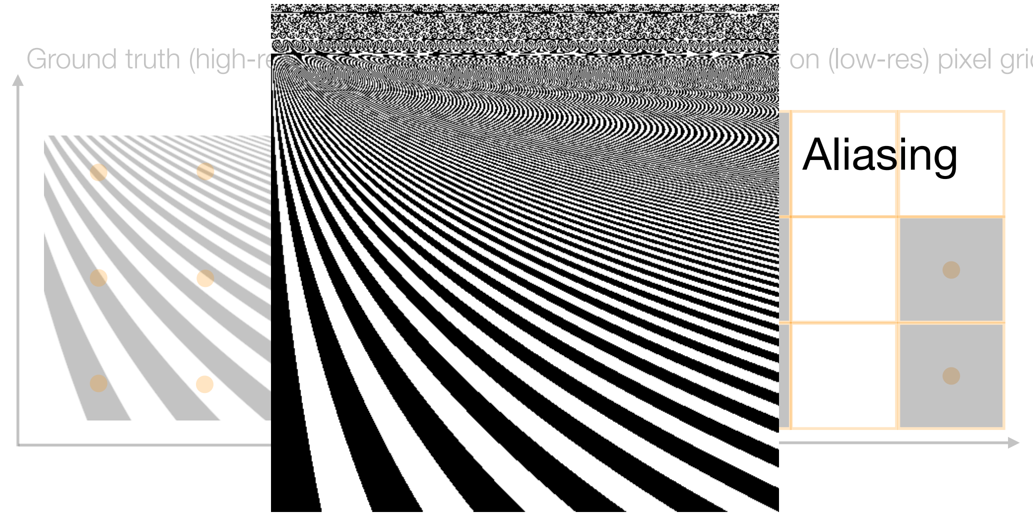

...structured noise,, also known as aliasing.

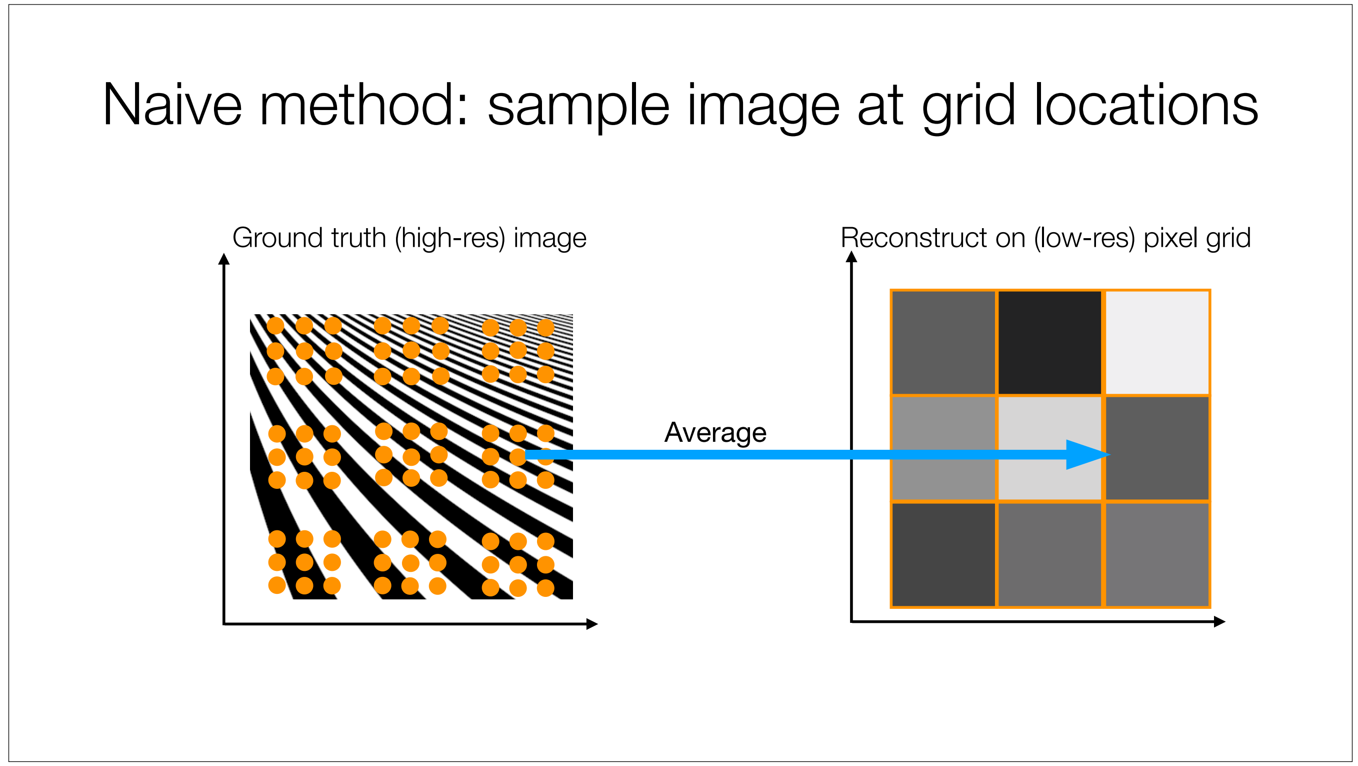

An easy way to get rid of this aliasing artifacts is to perform supersampling which involves generating multiple grid samples per pixel, evaluating the function values for each sample and average their values.

Antialiasing using general reconstruction filters

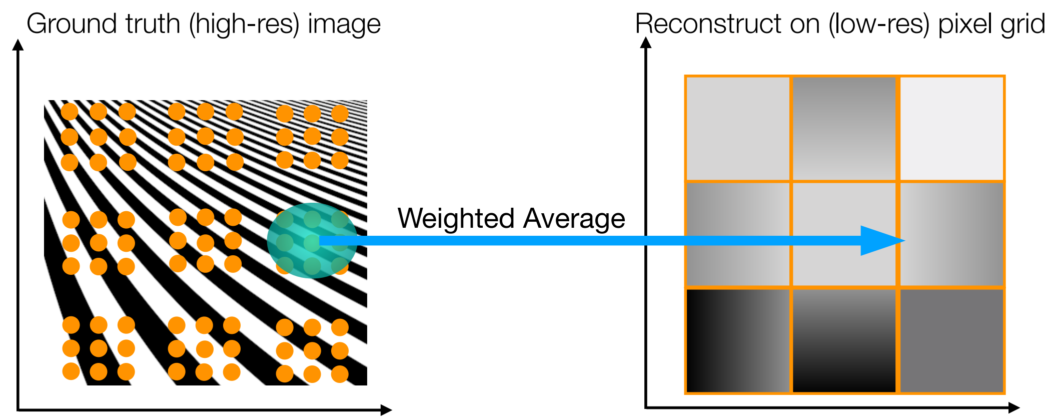

This average could be done using a reconstruction kernel which assigns some weights to each sample.

Naive method: sample image at grid locations

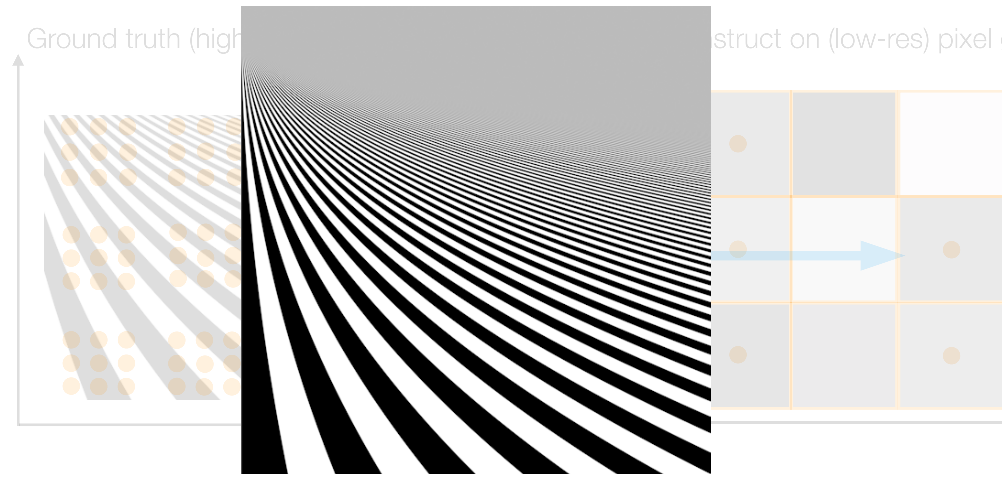

The result looks better to the human eye with more smooth transition from low frequency texture to the high frequency texture.

Rendering: reconstructing integrals

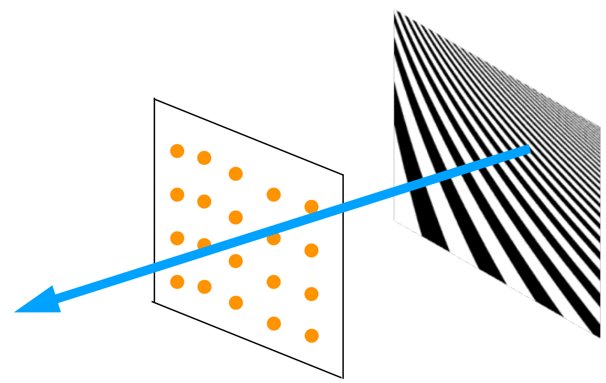

In rendering context, we can look at this as a ray shot from a sample within a pixel or image plane (or vice versa) and hitting this texture function. The function value is then stored in a pixel for a given sample. However,...

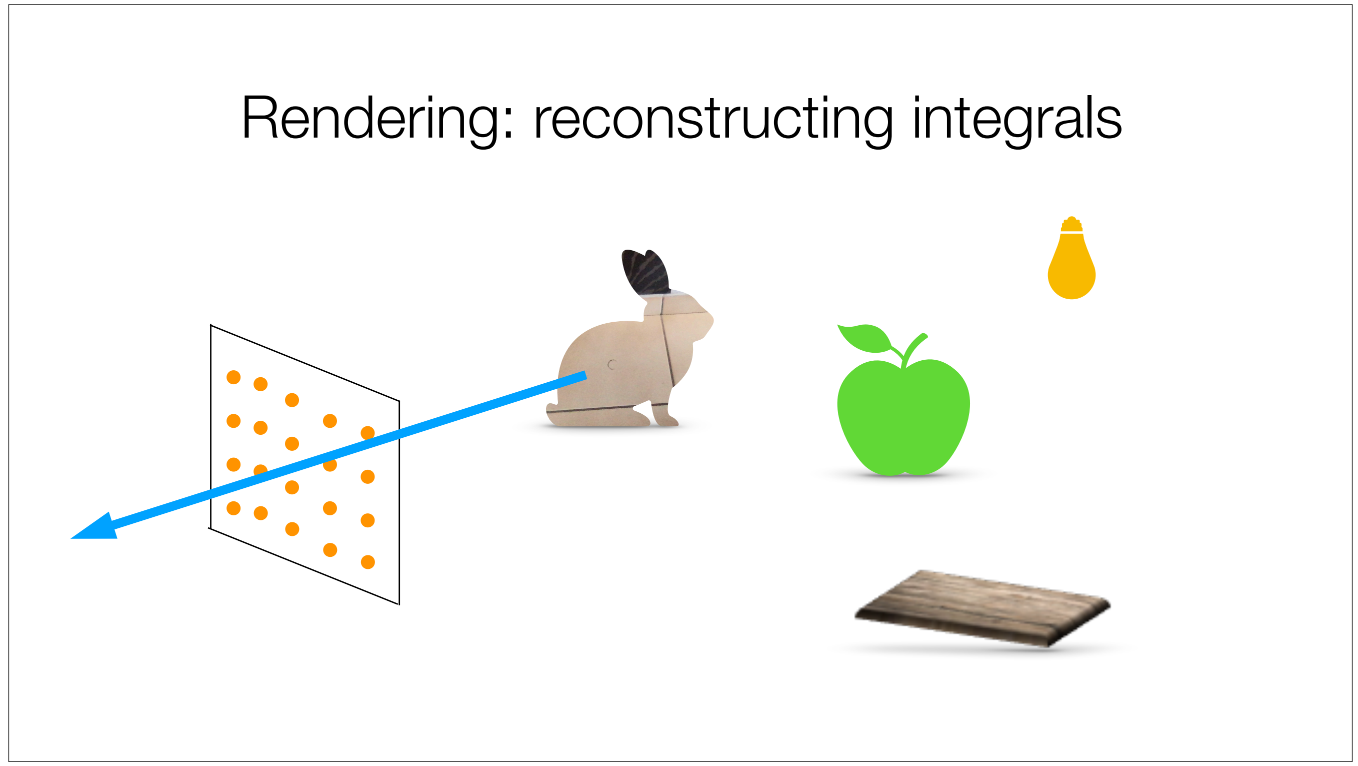

...in rendering we have more complex setup with 3D objects and multiple light sources. In this simple illustration, we assume a single light source.

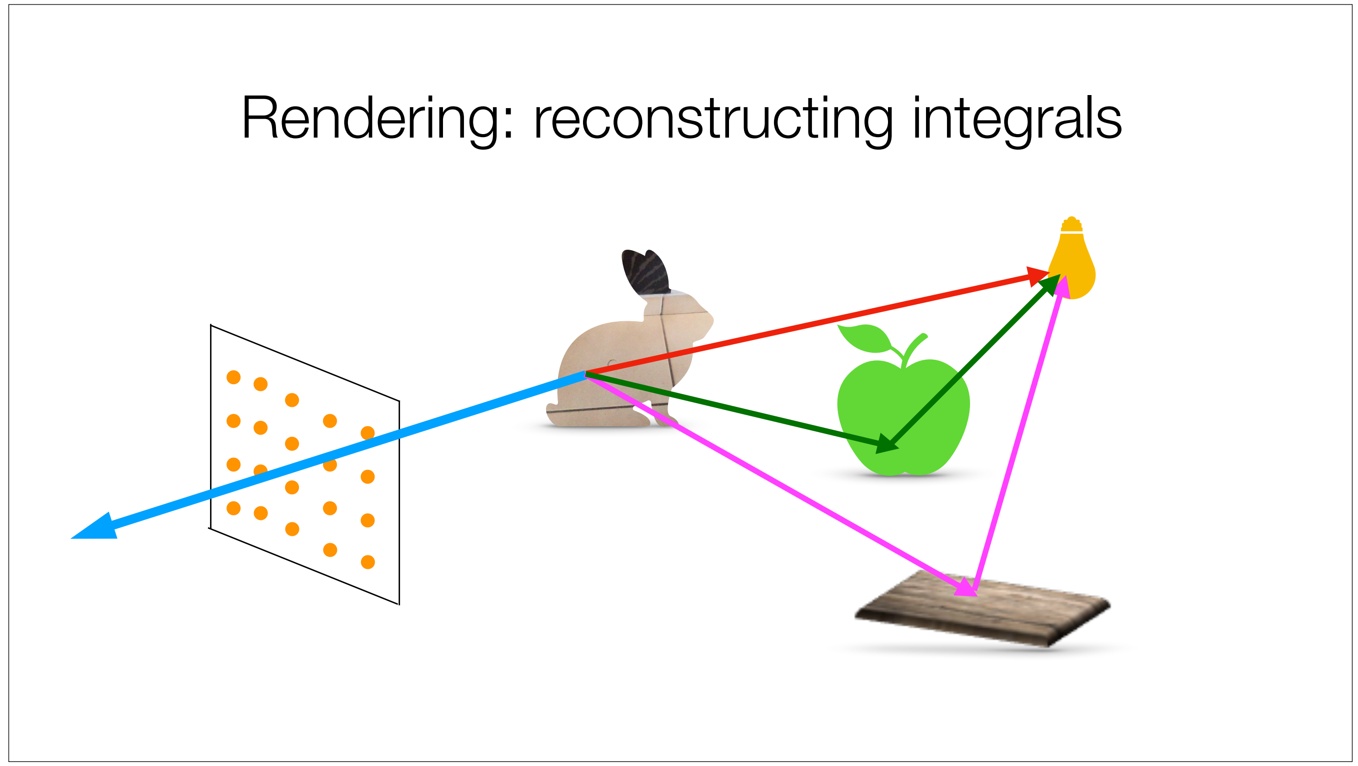

Here, the radiance or color value projected back to the image plane has been already reflected from multiple objects [CLICK], generating multiple paths.

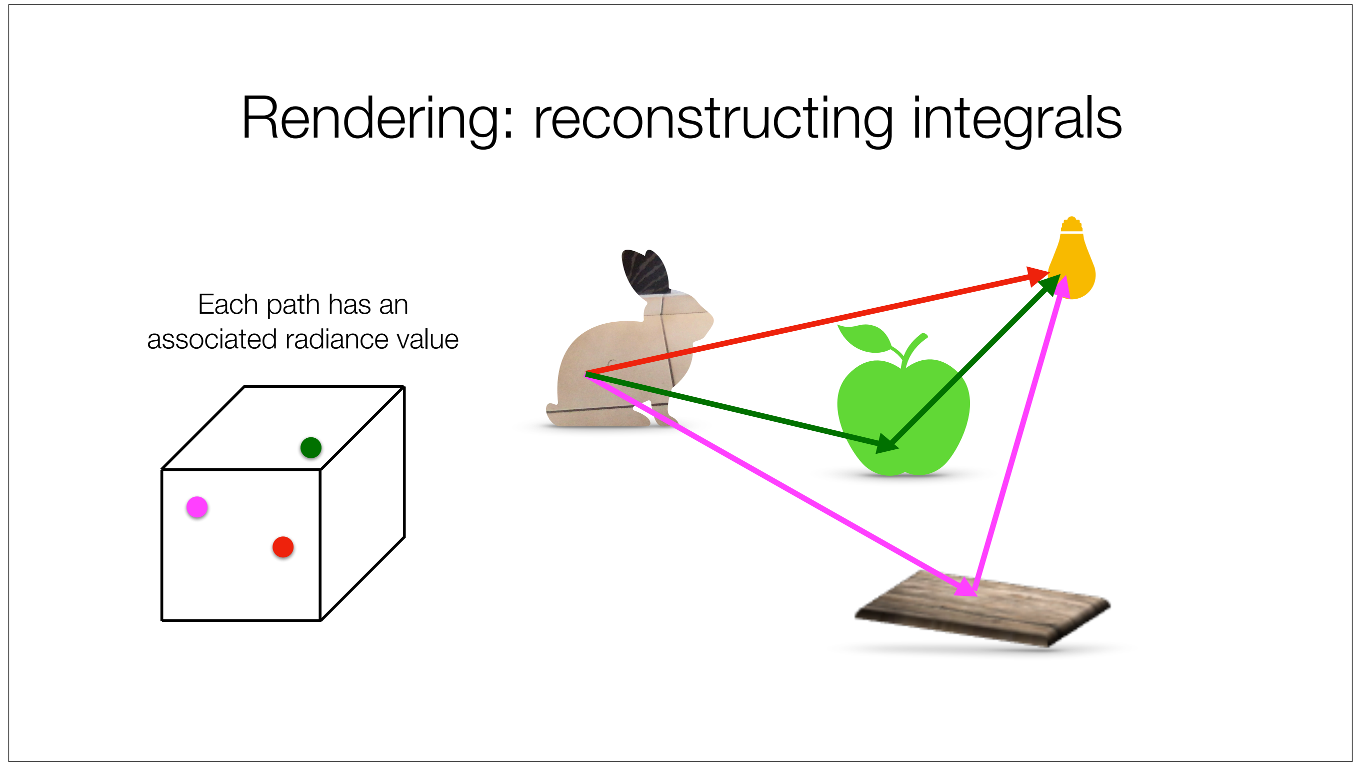

Each path has an associated radiance value.

Global Illumination: Participating media

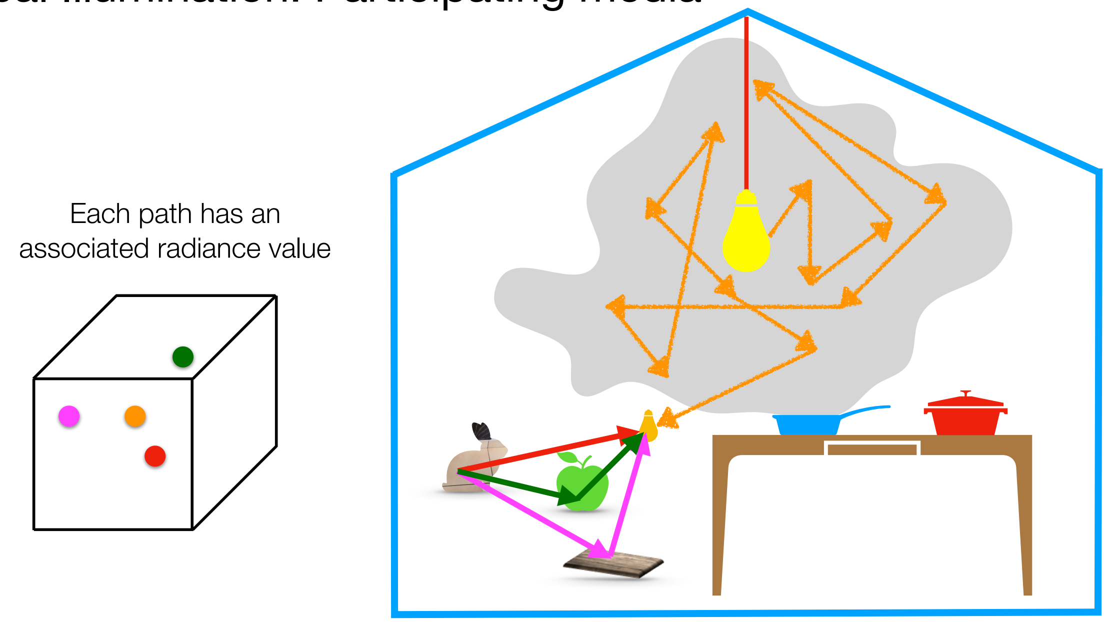

In a more complex setting, with participating media like smoke, the paths can be quite long.

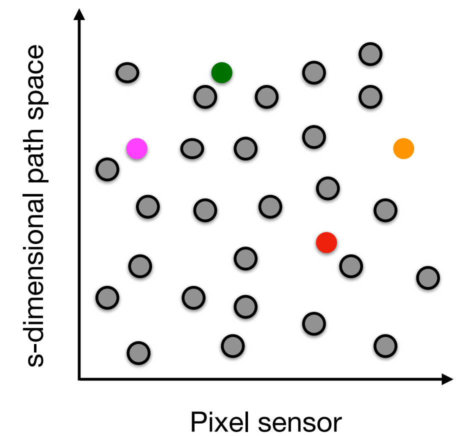
s-dimensional path space
Pixel sensor

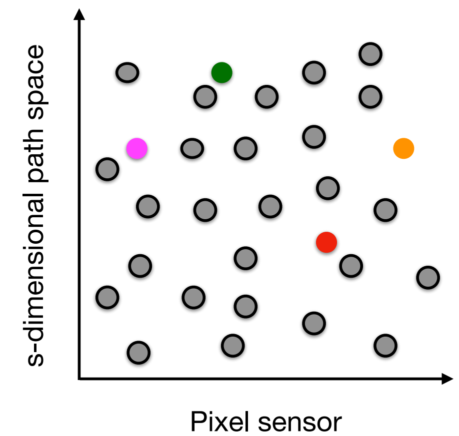

We can look at contribution of all these paths on this flatland illustration where the vertical axis represents the radiance value of s-dimensional paths and the horizontal direction represents the pixels of our image or a sensor.

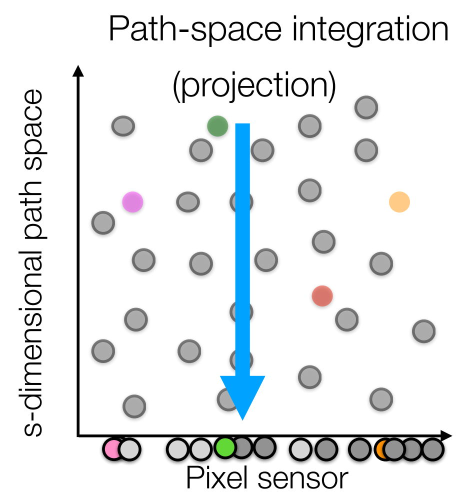

To get a value at a pixel, we project these path values [CLICK] to the corresponding pixels.

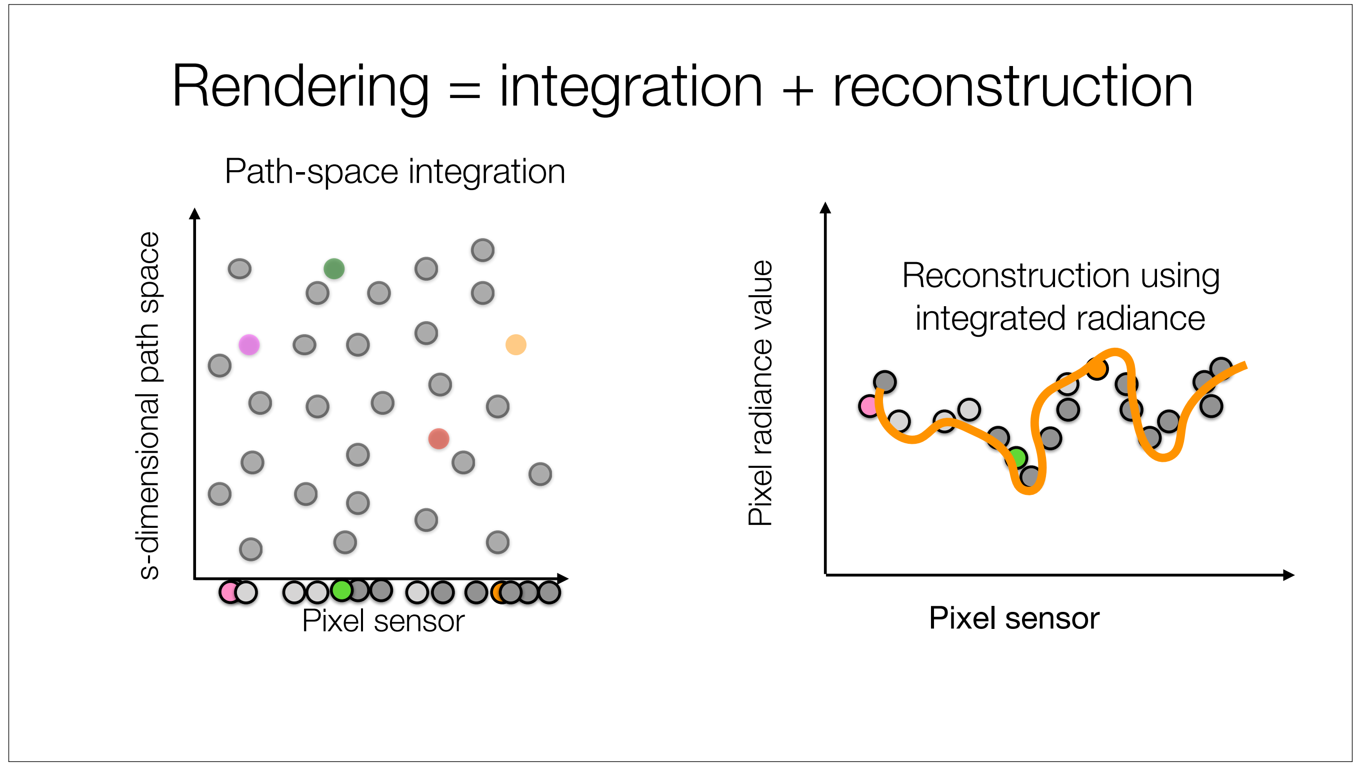

These pixel sensor values are then reconstructed to get a more smooth appearance on the image plane of the scene.

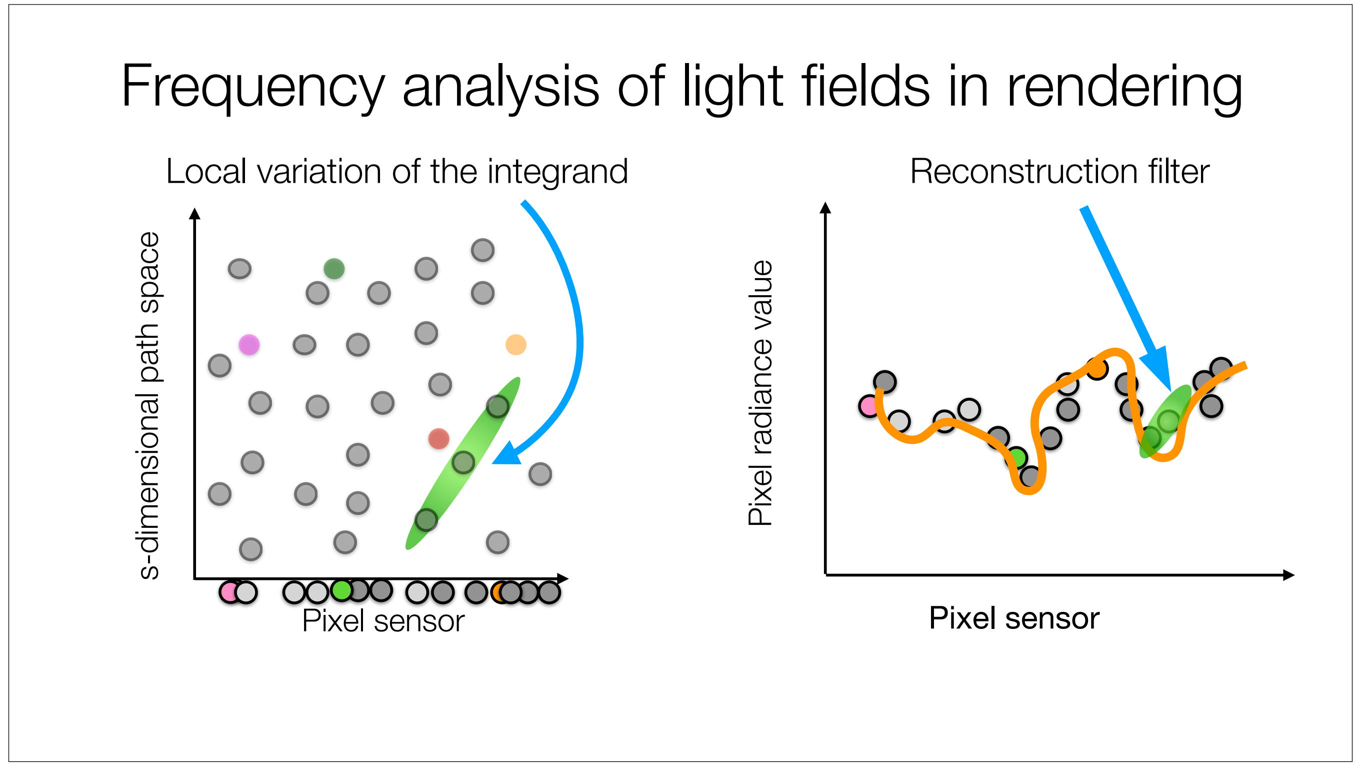

There is a quite a lot of work has been done over more than a decade in the Frequency domain where Fourier tools are used to better understand the local variation [CLICK] of the integrands in the path space or ray space to design or orient filters for better reconstruction.

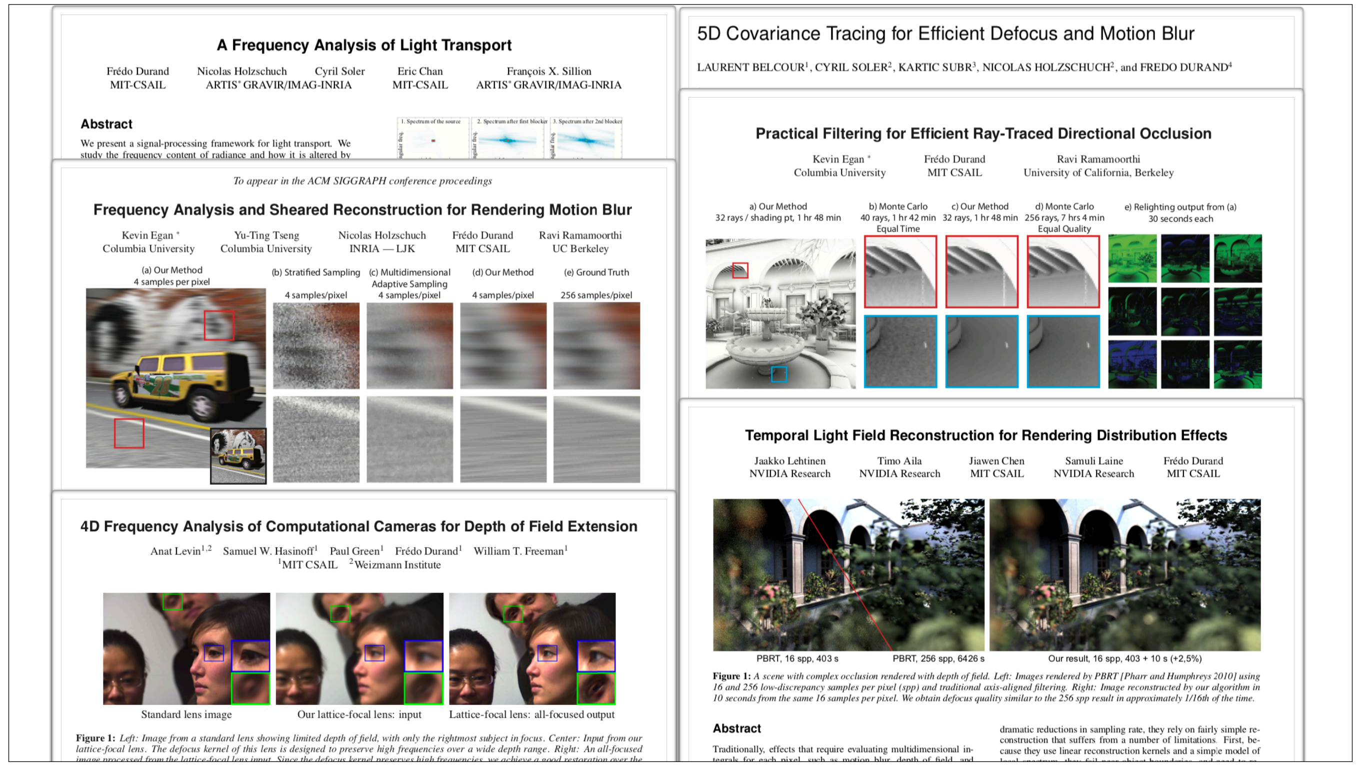

You can go over some of these papers to get an idea about these methods.

This STAR: Analyze sample correlations for MC sampling

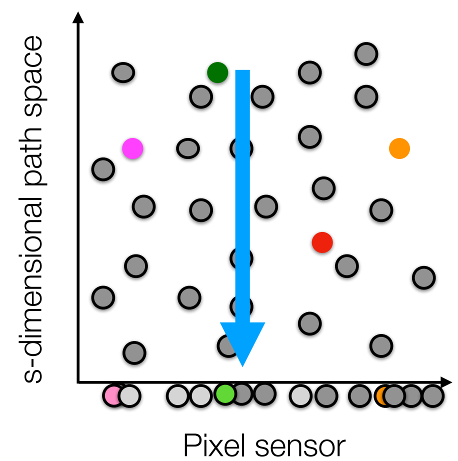

Assessing MSE, bias, variance and convergence of Monte Carlo estimators using spatial and spectral tools

In this presentation, we are actually interested in the projection or the integration aspect of these path space samples. We would like to asses the error in the form of variance and bias due to different sampling strategies used during Monte Carlo estimation techniques (that will be introduced in the next part of the presentation by Cengiz) using spatial and spectral tools.

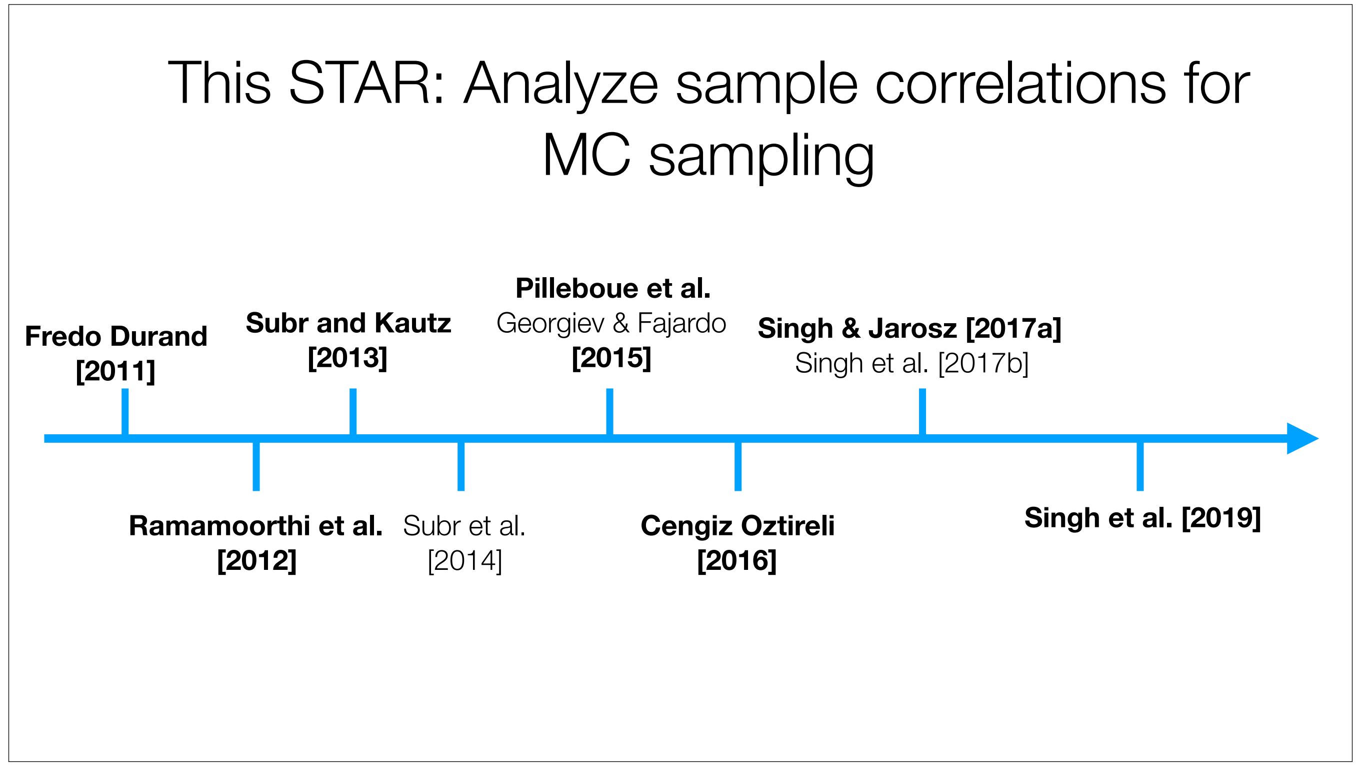

We will survey the works done in this past decade, starting from [CLICK] Fredo Durand's tech report in 2011 to this year's paper that makes the first attempt in analyzing importance samples using Fourier tools.

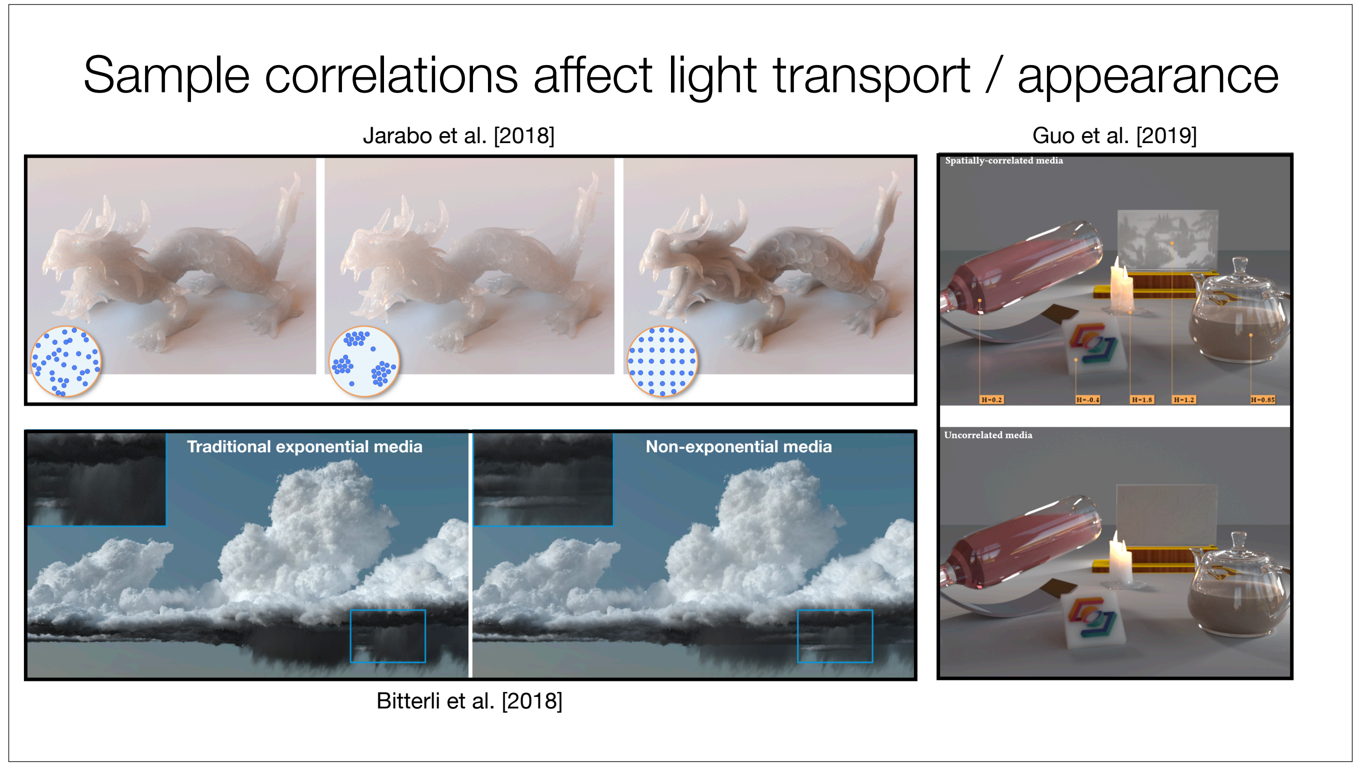

One quick impact of these correlations could be scene on appearance rendering.
When propagating through a participating medium, light is scattered
and absorbed in a very complicated way, and this transmission through a spatially-correlated media has demonstrated deviations from the classical exponential law of the corresponding uncorrelated media. And as you can see, these correlations can affect the overall appearance of the object.
This hints towards a new paradigm where we need to explore other sample correlations that could be useful in tailoring new appearances for artistic purposes.
However, in this talk...

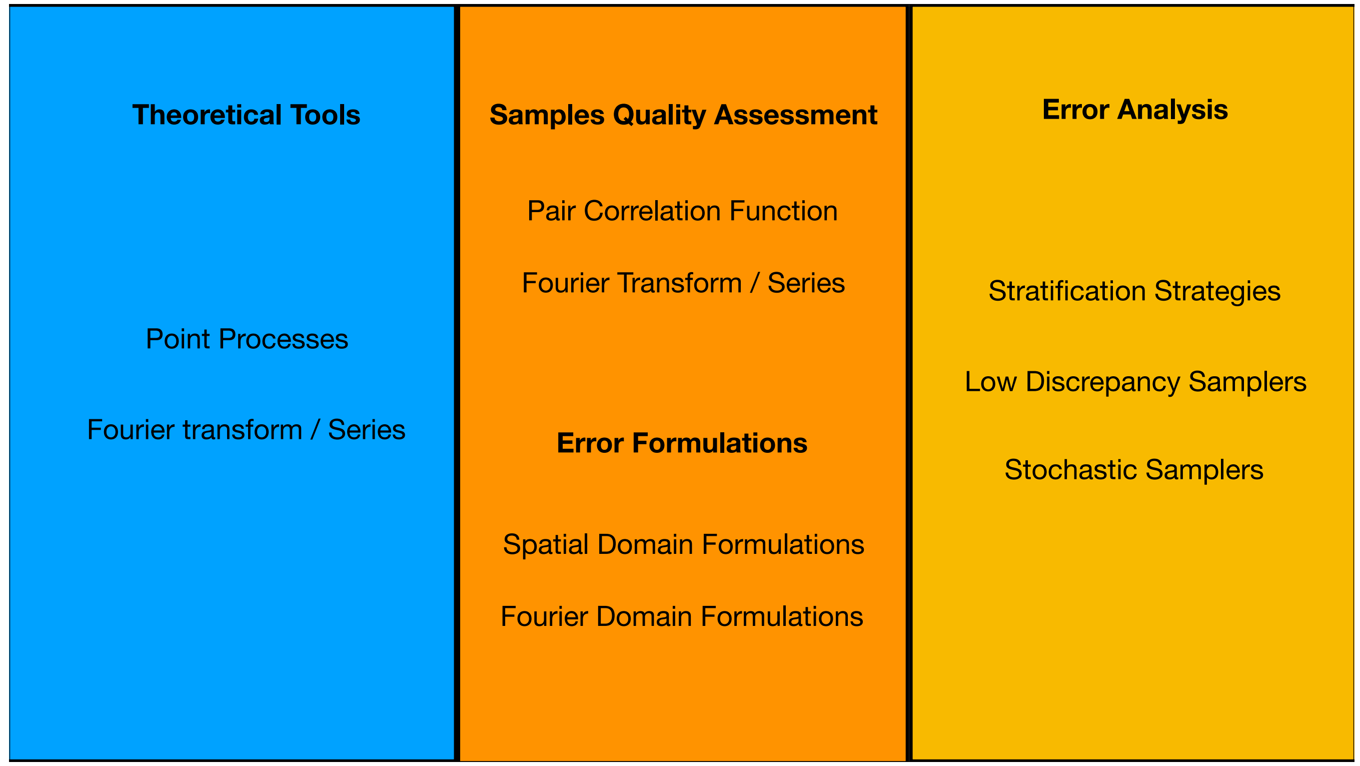

...we will confine ourselves to the following topics.
[CLICK] We first overview the theoretical tools from the stochastic point processes and the Fourier literature.
[CLICK] We then see different ways to assess sample correlations using these spatial and spectral tools. We then present the error formulations developed using these tools.
[CLICK] In the last section, we overview how different sampling strategies affect error during Monte Carlo integration for rendering purposes.
With this, I handover the stage to Cengiz...