

# State of the Art in Photon Density Estimation

---


Progressive Expectation–Maximization for  
Hierarchical Volumetric Photon Mapping

**Wojciech Jarosz**

(slides courtesy of Wenzel Jakob)

# Motivation

## Volumetric photon mapping



## Issues

- high-frequency illumination requires *many* photons
- time spent on photons that contribute very little
- prone to temporal flickering

# Motivation



Beam radiance estimate : 917K photons



Per-pixel  
render time

# Motivation

Jakob et al. 2011. Proceedings of EGSR.

Beam radiance estimate : 917K photons



Render time: 281 s



Per-pixel  
render time

Our method: 4K Gaussians



Render time: 125 s




Per-pixel  
render time

## Our approach:


- represent radiance using a Gaussian mixture model (**GMM**)
- fit using progressive expectation maximization (**EM**)
- render with multiple levels of detail

# Motivation

Beam radiance estimate : 4M photons



Our method: 16K Gaussians





## Our approach:

- represent radiance using a Gaussian mixture model (**GMM**)
- fit using progressive expectation maximization (**EM**)
- render with multiple levels of detail



# Related work

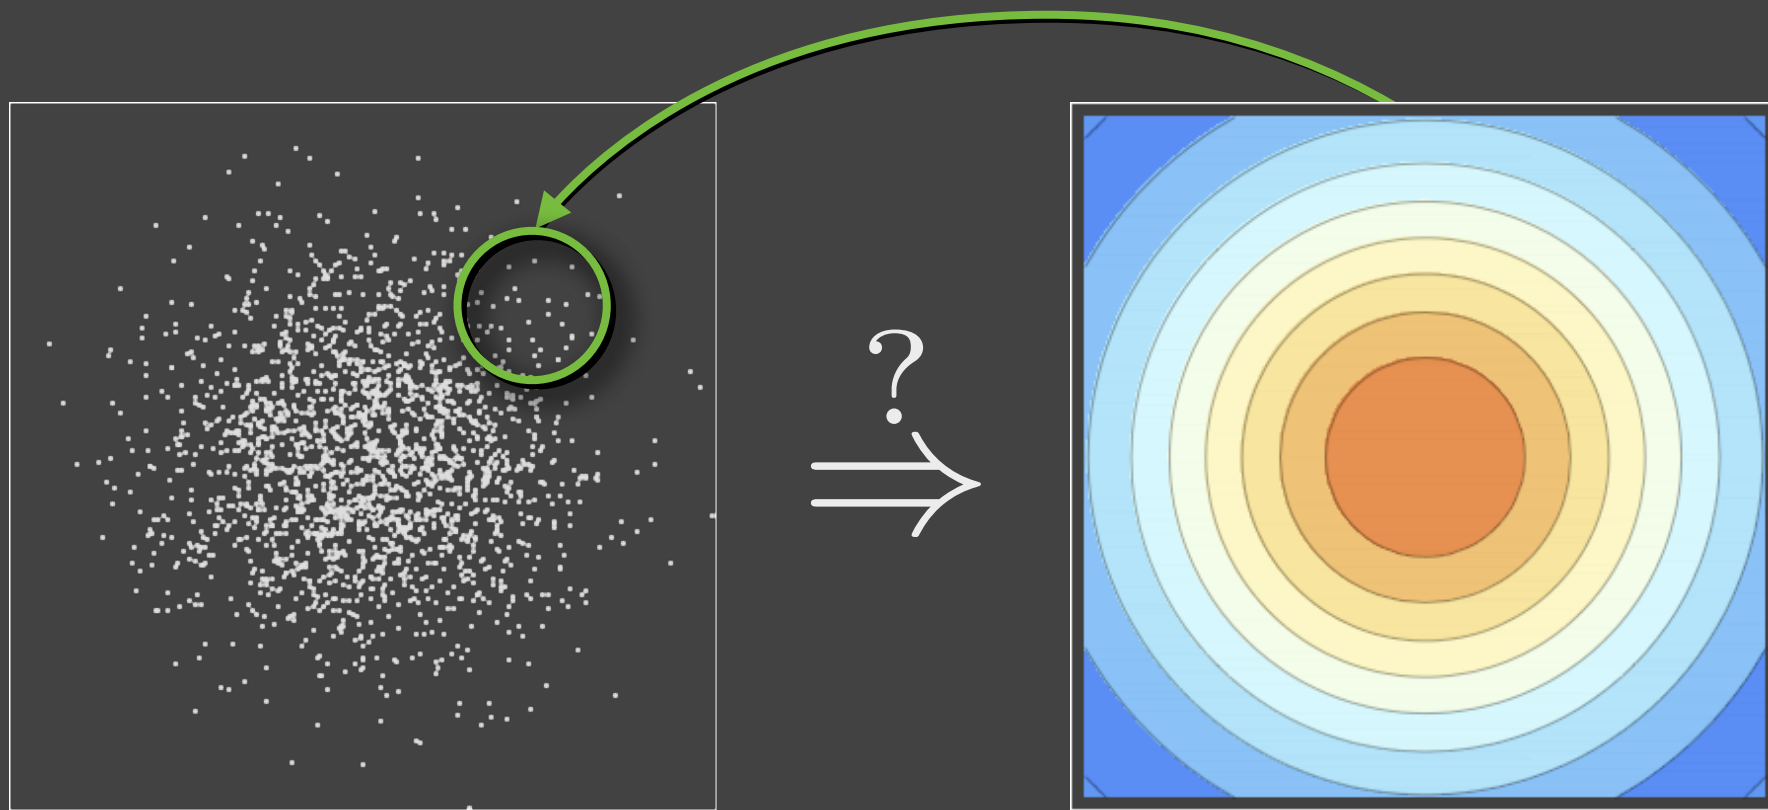
- Hierarchical photon mapping  
[Spencer and Jones 09]
- Photon relaxation  
[Spencer and Jones 09]
- Progressive photon relaxation  
[Spencer and Jones 13]
- Photon parameterisation for  
robust relaxation constraints  
[Spencer and Jones 13]



**Feature detection & preservation challenging**

# Density estimation

# Density estimation



Given photons  
 $x_1, x_2, \dots$

approximately determine  
their density  $f$

## Nonparametric:

- Count the number of photons within a small region

## Parametric:


- Find suitable parameters for a **known** distribution



# Gaussian mixture models

- Photon density modeled as a weighted sum of Gaussians:

$$f(\mathbf{x} \mid \Theta) = \sum_{i=1}^k w_i g(\mathbf{x} \mid \Theta_i)$$




# Gaussian mixture models

- Photon density modeled as a weighted sum of Gaussians:


$$f(\mathbf{x} \mid \Theta) = \sum_{i=1}^k w_i g(\mathbf{x} \mid \Theta_i)$$

Unknown parameters  $\Theta$ :



1. Weights



2. Means



3. Covariance matrices



# Maximum likelihood estimation

Approach: find the “most likely” parameters, i.e.

$$\Theta^* := \underset{\Theta}{\operatorname{argmax}} \prod_{i=1}^n f(x_i | \Theta)$$

Diagram illustrating the Maximum Likelihood Estimation (MLE) approach:

- $\Theta^*$  is labeled "Estimated parameters" (indicated by a green arrow pointing up).
- $\Theta$  is the parameter space.
- $n$  is the sample size.
- $f(x_i | \Theta)$  is the likelihood function, labeled "Mixture model" (indicated by a purple arrow pointing down to the function).
- $x_i$  is the observed data, labeled "Photon locations" (indicated by a pink arrow pointing up).


➔ Expectation maximization

# Expectation maximization

- Two components:


**E-Step:** establish soft assignment between photons and Gaussians

**M-Step:** maximize the expected likelihood



- Finds a locally optimal solution  
→ good starting guess needed!
- Slow and scales poorly —  $\mathcal{O}(n^2)$   
(where  $n$ : photon count)


# Plain EM



Each photo Accelerate “pull” motion [Verby-Gustafson 06] components




# Accelerated EM



## Stored cell statistics:

- photon count
- mean position
- average outer product

# Progressive EM




## Stored cell statistics:

- photon count
- mean position
- average outer product

## Our modifications:

- better cell refinement

# Progressive EM




## Stored cell statistics:

- photon count
- mean position
- average outer product

## Our modifications:

- better cell refinement
- progressive photons shooting passes

# Progressive EM



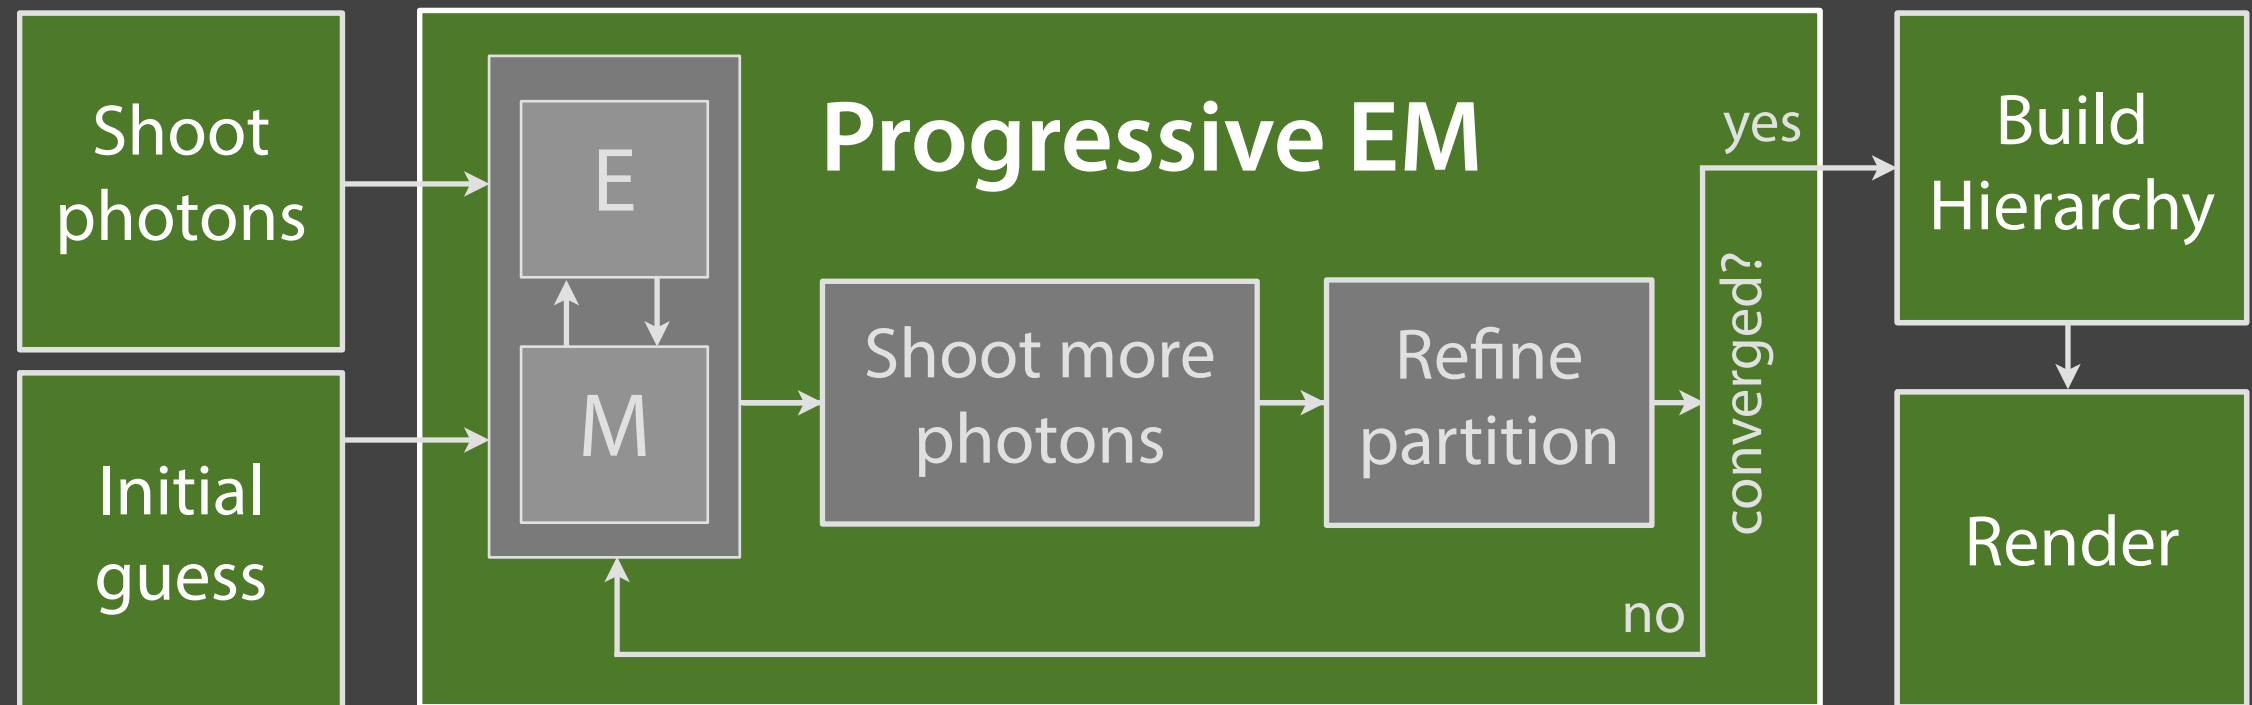
## Stored cell statistics:

- photon count
- mean position
- average outer product

## Our modifications:


- better cell refinement
- progressive photons shooting passes
- reduced complexity  
 $\mathcal{O}(n^2) \rightarrow \mathcal{O}(n \log n)$

# Pipeline overview





# Rendering



$$\text{pixel value} = \sum_{i=1}^k \text{contrib}(i)$$


$$\text{contrib}(i) = \int_a^b g(\mathbf{r}(t) | \bar{\Theta}_i) e^{-\sigma_t t} dt = C_0 \left[ \text{erf} \left( \frac{C_3 + 2C_2 b}{2\sqrt{C_2}} \right) - \text{erf} \left( \frac{C_3 + 2C_2 a}{2\sqrt{C_2}} \right) \right]$$

...

# Level of detail hierarchy


## Agglomerative construction:

- Repeatedly merge nearby Gaussians based on their Kullback-Leibler divergence




# Rendering


Example  
hierarchy:




Criterion 1: bounding box intersected?




Criterion 2: solid angle large enough?




Criterion 3: attenuation low enough?





**BRE:** 1M Photons

23+192 = 215 s






**Our method: 4K Gaussians**  
(fit to 1M photons)

$35+24 = 59 \text{ s}$   
(3.6x)

23








**BRE:** 18M Photons

507+609 = 1116 s


24



**Our method: 64K Gaussians**  
(fit to 18M photons)


868+66 = 934 s  
(1.2x)

25



**BRE:** 4M Photons

89 + 638 = 727 s

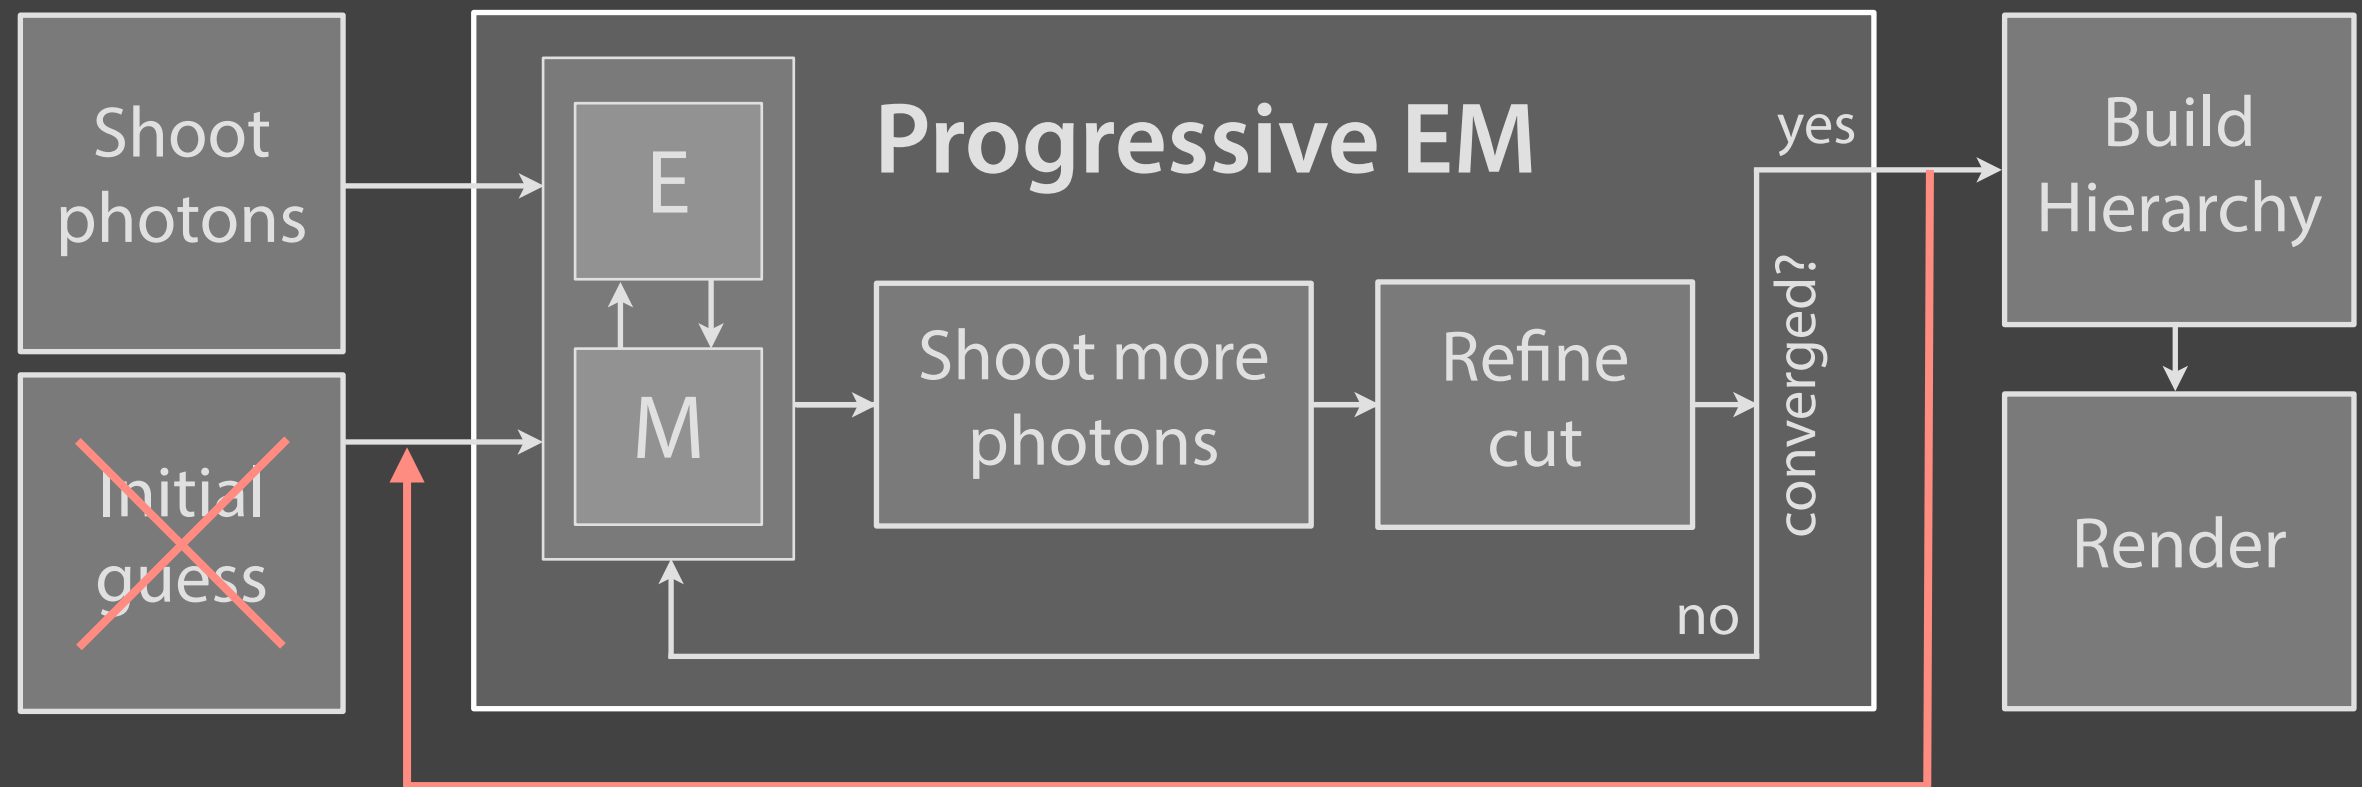


**Our method: 16K Gaussians**

$330 + 127 = 457$  s  
(1.6x)

27

# Temporal Coherence



- Feed the result of the current frame into the next one  
→ Faster fitting, no temporal noise



# Scene 1: **BUMPYSPHERE**

Volume caustics from a rotating light source



## GPU-based rasterizer:

- Anisotropic Gaussian splot shader: 30 lines of GLSL
- Gaussian representation is very compact (4096-term GMM requires only ~240KB of storage)

# Conclusion

- Rendering technique based on parametric density estimation
- Uses a progressive and optimized variant of accelerated EM
- Compact & hierarchical representation of volumetric radiance
- Extensions for temporal coherence and real-time visualization

Questions?