

# GUIDING IN PATH SPACE

# Johannes Hanika, Weta Digital



# Selective guided sampling with complete light transport paths

Florian Reibold, Johannes Hanika, Alisa Jung, Carsten Dachsbacher

Karlsruhe Institute of Technology, Germany




# Motivation

- we have heard earlier: need to consider all decisions along a path together to work really well
  - radiance, BSDF, distance, RR, ..
- NEE works well and is advancing (guided, too!)
  - no need to learn this twice!
- 1) all aspects: work with full paths?
- 2) working well except in which area?



# Motivation


- Path tracing can efficiently handle the majority of rendering problems in practice (see path tracing course Wednesday)





# Motivation

- What about the rest?





# Motivation

- What about the rest?





# Motivation



↑  
Unproblematic  
light transport

Difficult light transport  
outliers





# Motivation



↑  
Unproblematic  
light transport  
(**unguided sampling**)

Difficult light transport  
(**guided sampling**)






# Related methods

- we've seen many approaches today:
  - guided NEE (by Alex and Jaroslav)
  - 2D marginalised guiding caches (by Jirka and Thomas)
  - extensions to products/BSDF (by Sebastian)
- main difference: low dimensional!



# Example: marginal vs. high dimensional

- indirect lighting




- conclusion: can augment marginal caches, but why discard data in the first place?



# Path Integral


$$I = \int_{\mathcal{P}} f(X) dX$$





# Monte Carlo Integration


$$I \approx \frac{1}{N} \sum_{i=1}^N \frac{f(X_i)}{p(X_i)}$$





# Illustration in 1D


- Importance sampling?





# Guided Sampling in 1D


- Create samples  $X_i \sim p_u$





# Guided Sampling in 1D


- Keep outliers





# Guided Sampling in 1D


- Place a Gaussian around each outlier





# Guided Sampling in 1D


- Define guided PDF as sum over all Gaussians





# Guided Sampling in 1D


- Iterate by sampling  $p_u$  and  $p_g$





# Guided Sampling in 1D


- Keep outliers wrt.  $p_u$  and  $p_g$






# Guided Sampling in 1D

- Update  $p_g$



# Overview


- In each iteration, sample paths from  $p_u$  and  $p_g$






# Overview

- Keep outliers with highest contribution



# Overview


- Add paths to the set of guide paths
- Compute Gaussians using neighbourhood information






# Overview

- Guided sampling:
  - Choose guide path randomly and
  - Sample Gaussians incrementally
- Guided and unguided sampling combined with multiple importance sampling



# Guiding Behaviour


Reference






# Guiding Behaviour

Path tracing




Guided path tracing





# Guiding Behaviour


Reference






# Guiding Behaviour

Path tracing



Guided path tracing




# Guiding PDF


- how to derive a good sampling density around a guide path?



# Path Correlation



Specular/glossy




Rough/diffuse

# Gaussians

- Compute Gaussians for sampling using nearest neighbours

— Guide path  
— Nearest neighbours






# Gaussians

- Compute Gaussians for sampling using nearest neighbours


— Guide path  
— Nearest neighbours  
— New path



# Gaussians

- Sample 3D Gaussian at next vertex?


— Guide path  
— Nearest neighbours  
— New path





# Gaussians


- Sample 3D Gaussian at next vertex?



# Gaussians

- Compute 6D covariance matrix for path segments

- Guide path
- Nearest neighbours
- New path






# Gaussians


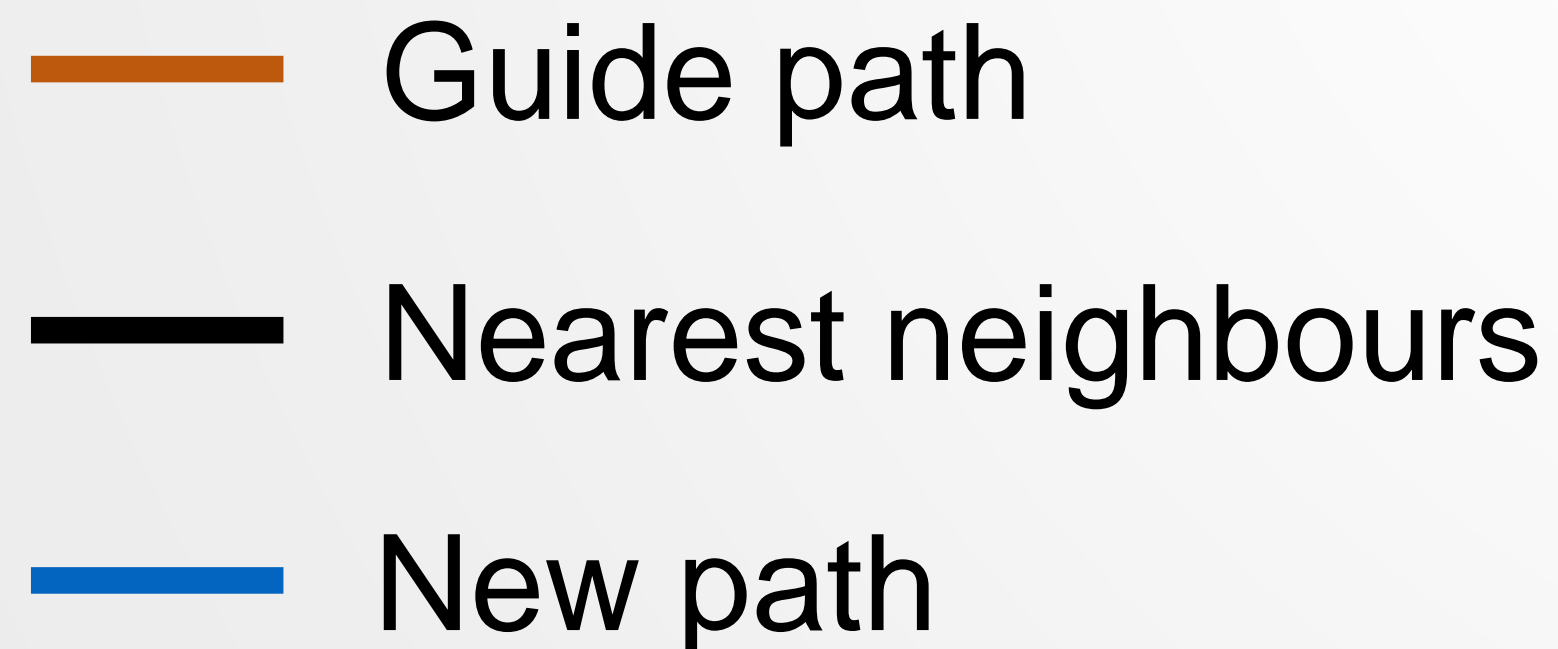
- Compute 6D covariance matrix for path segments
- And conditional Gaussian using  $x$

- Guide path
- Nearest neighbours
- New path



# Gaussians

- Compute 6D covariance matrix for path segments
- And conditional Gaussian using  $x$






# Gaussians

- Compute 6D covariance matrix for path segments
- And conditional Gaussian using  $x$

- Guide path
- Nearest neighbours
- New path



# Sampling next vertex

Next vertex in volume

Next vertex on surface

Sampling  
of  
Gaussian



Sampling  
of  
BSDF






# Sampling next vertex

Next vertex in volume

Sampling  
of  
Gaussian



Next vertex on surface



Sampling  
of  
BSDF





# Sampling next vertex

Next vertex in volume

Next vertex on surface

Sampling  
of  
Gaussian



Sampling  
of  
BSDF







# Sampling next vertex


Next vertex in volume

Next vertex on surface

Sampling  
of  
Gaussian




Sampling  
of  
BSDF




# Sampling next vertex

Next vertex in volume



Sampling  
of  
Gaussian



Next vertex on surface



Sampling  
of  
BSDF






# Guided PDF

- Many guide paths could sample the same path  $X$ 
  - We have to sum up all individual probability densities
- For fast evaluation, we truncate Gaussians ( $\approx 3\sigma$ )
- Acceleration structure for fast pruning

— Guide paths  $X_l$   
— Path  $X$




# What is a good guide path?

- add incrementally, but how to pick?



# Selecting Guide Paths



- Outliers  $\neq$  samples with high contribution  $\mathcal{C} = f/p$
- Outliers classification: Density based outlier rejection (DBOR, Zirr et al. [2018])






# Selecting Guide Paths


- DBOR: DeCoro et al. [2010] / Zirr et al. [2018]
- Framebuffer cascade with histogram
- Samples split according to throughput
- bad sampling? count nearest neighbours




$$C \in [0, 8^1]$$



$$C \in [8^1, 8^2]$$



$$C \in [8^2, 8^3]$$









$$C \in [8^3, \infty)$$



# Pool DBOR Cascade

path tracing



guided path tracing



# Selecting Guide Paths

- collect candidate outliers for every iteration/thread
- accept a fixed number of new guide paths every iteration
- pick the largest contributions (biggest effect on rMSE)



# Constructing a full estimator

- combine guided and unguided sampling

# Multiple importance sampling

- Balance heuristic:

$$p(X) = u \cdot p_u(X) + (1 - u) \cdot p_g(X)$$

$$\langle I(X) \rangle = \frac{f(X)}{p(X)}$$

- $u$ : ratio of sampling unguided paths
- what if this combination still doesn't cover full path space efficiently?



# Remaining outliers

- Outliers contribute fully to the image
- We remove outliers with DBOR to get clean images







Outliers  
removed





# Pool

## Reference








56



# Pool - 30min

## Path tracing

PT




57



# Pool - 30min




## Guided path tracing



104k guide paths ( $\approx 310\text{MB}$ )

PT


Guided PT






# Pool - 30min

## Guided path tracing + DBOR





104k guide paths ( $\approx 310\text{MB}$ )

PT



Guided PT




Guided PT+DBOR




# Pool - 30min


## Path tracing + DBOR



PT



guided PT



PT + DBOR



guided PT+DBOR





# Dragon

## Reference







# Dragon - 10h

## Path tracing


PT






# Dragon - 10h

## Guided path tracing






69k guide paths ( $\approx 207\text{MB}$ )

PT




Guided PT






# Dragon - 10h

## Guided path tracing + DBOR




69k guide paths ( $\approx 207\text{MB}$ )






# Dragon - 10h


## Path tracing + DBOR



PT



Guided PT



PT + DBOR




Guided PT+DBOR






# Limitations & Future Work

- When every path is an outlier, no path is an outlier ☹️
- Impossible to cover all of path space with guide paths




path tracing



7921spp, RMSE 0.75

guided path tracing



4405spp, RMSE 1.30



# Dense scattering media?



- Works only in principle. Sebastian's method is better here!





# Limitations & Future Work


- Temporal stability is challenging






# Limitations & Future Work

- Temporal stability is challenging
- Improvement: Resample guide paths from previous frame



**independent guide path cache**




**resampled guide path cache**



# Conclusion

- Data driven path sampling with local exploration behaviour
- Path construction using information of multiple existing paths
- Other Monte Carlo samplers possible as the unguided sampler
- Similarities to Sequential Monte Carlo
- Guide paths could be hand picked (artist) or from Markov Chain without detailed balance







THANK YOU!