SIGGRAPH ASIA 2018 TOKYO

CONFERENCE 4 – 7 December 2018
EXHIBITION 5 – 7 December 2018
Tokyo International Forum, Japan
SA2018.SIGGRAPH.ORG

Sponsored by

Selective guided sampling with complete light transport paths

Florian Reibold, Johannes Hanika, Alisa Jung, Carsten Dachsbacher

Karlsruhe Institute of Technology, Germany

Karlsruhe Institute of Technology

Motivation

- Path tracing can efficiently handle the majority of rendering problems in practice (Fascione et al. [2017])

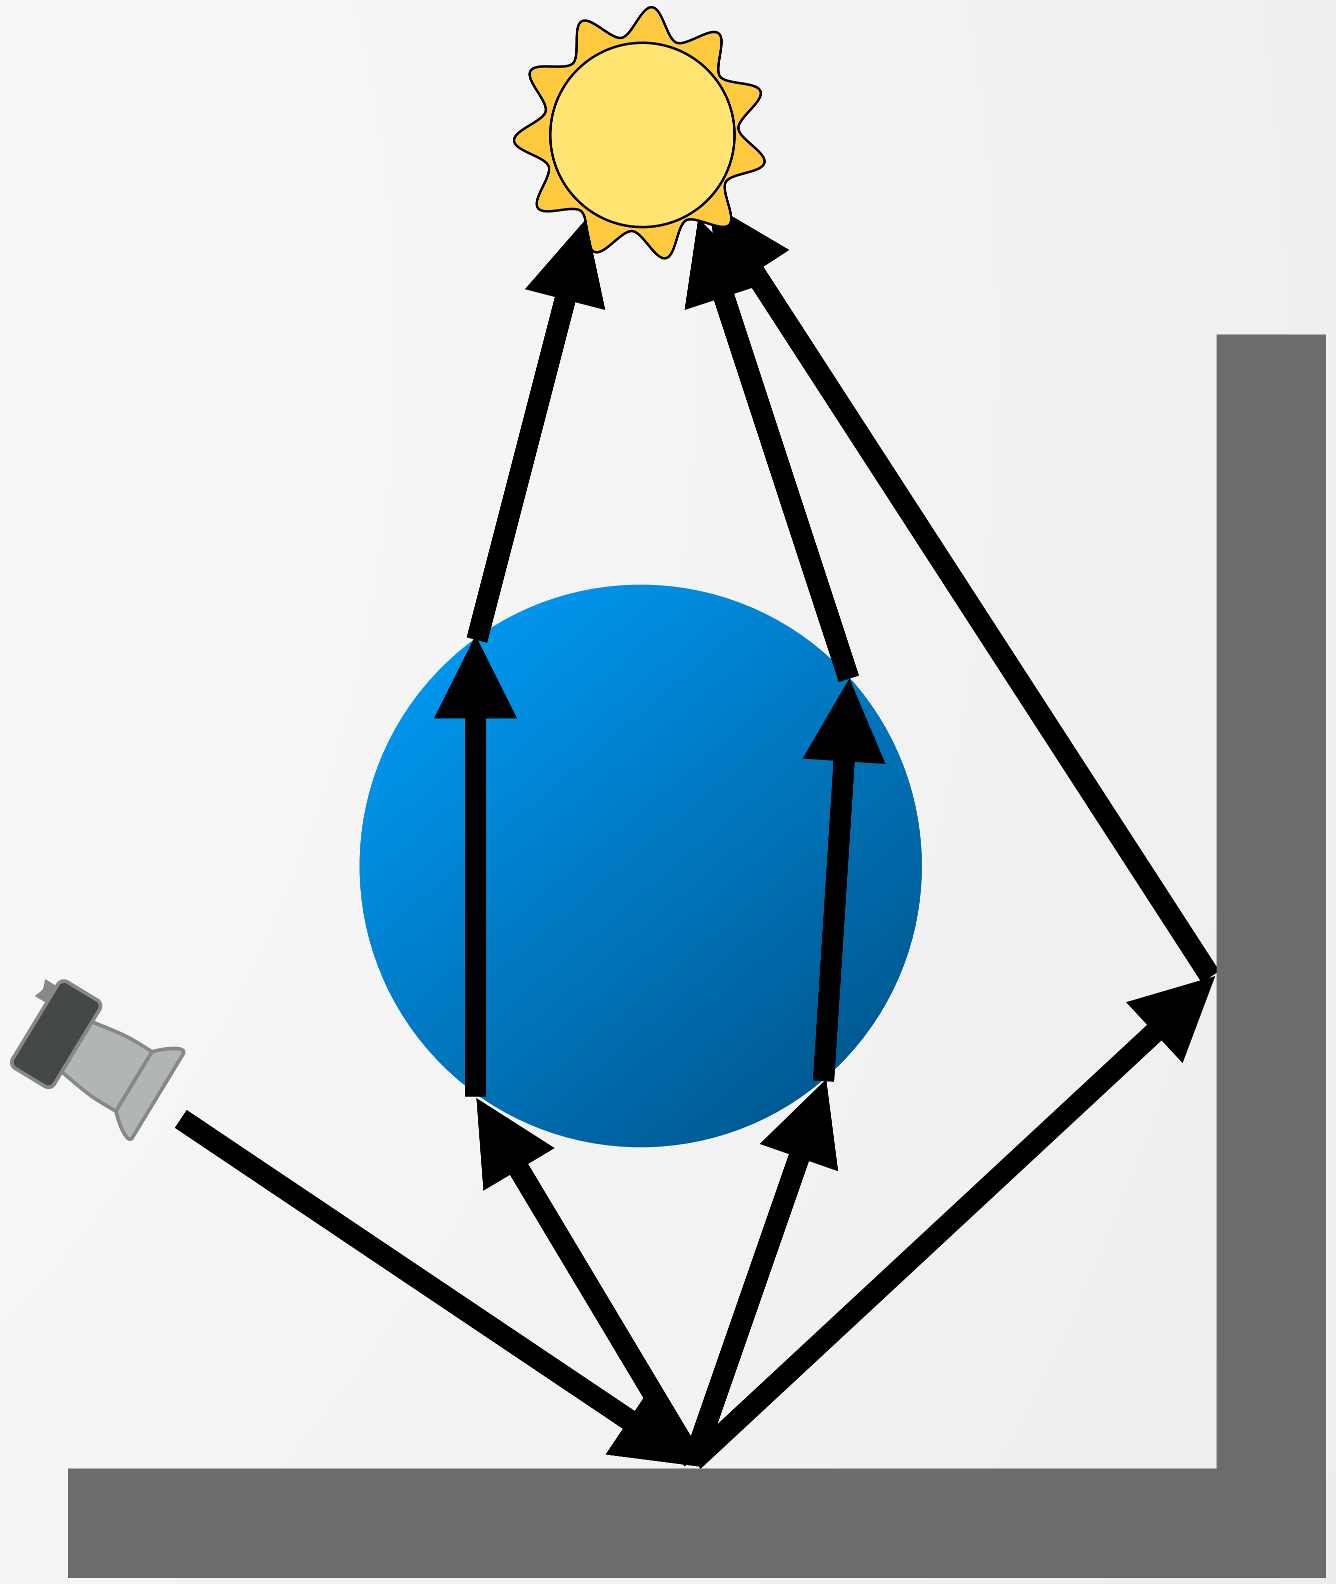

Motivation

- What about the rest?

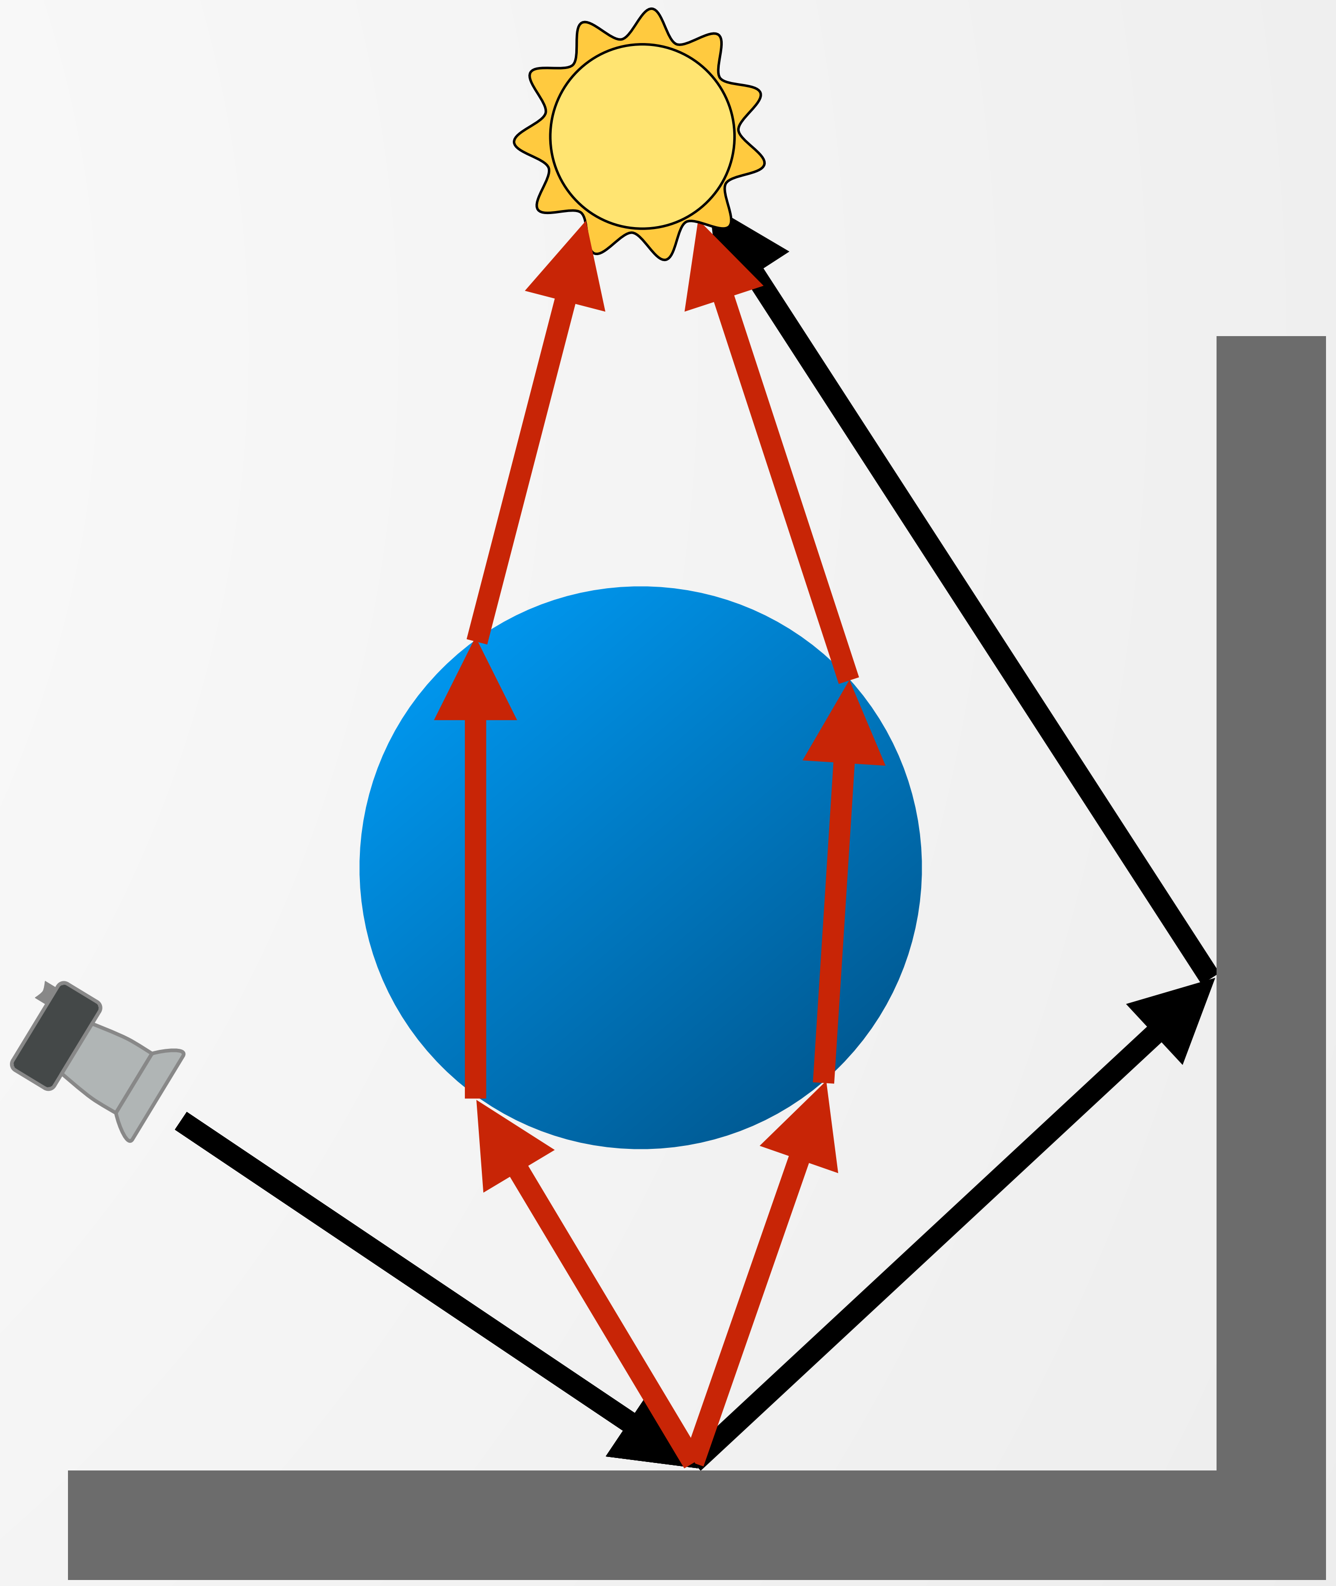

Motivation

- What about the rest?

Motivation

Unproblematic light transport

Difficult light transport outliers

Motivation

Difficult light transport
(**guided sampling**)

Unproblematic
light transport
(**unguided sampling**)

Recent related guiding methods

- Vorba et al. [2014], Müller et al. [2017], …
  - path guiding using incident radiance information
- Product sampling (Herholtz et al. [2016]) and application to participating media is not straight forward

Path Integral

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

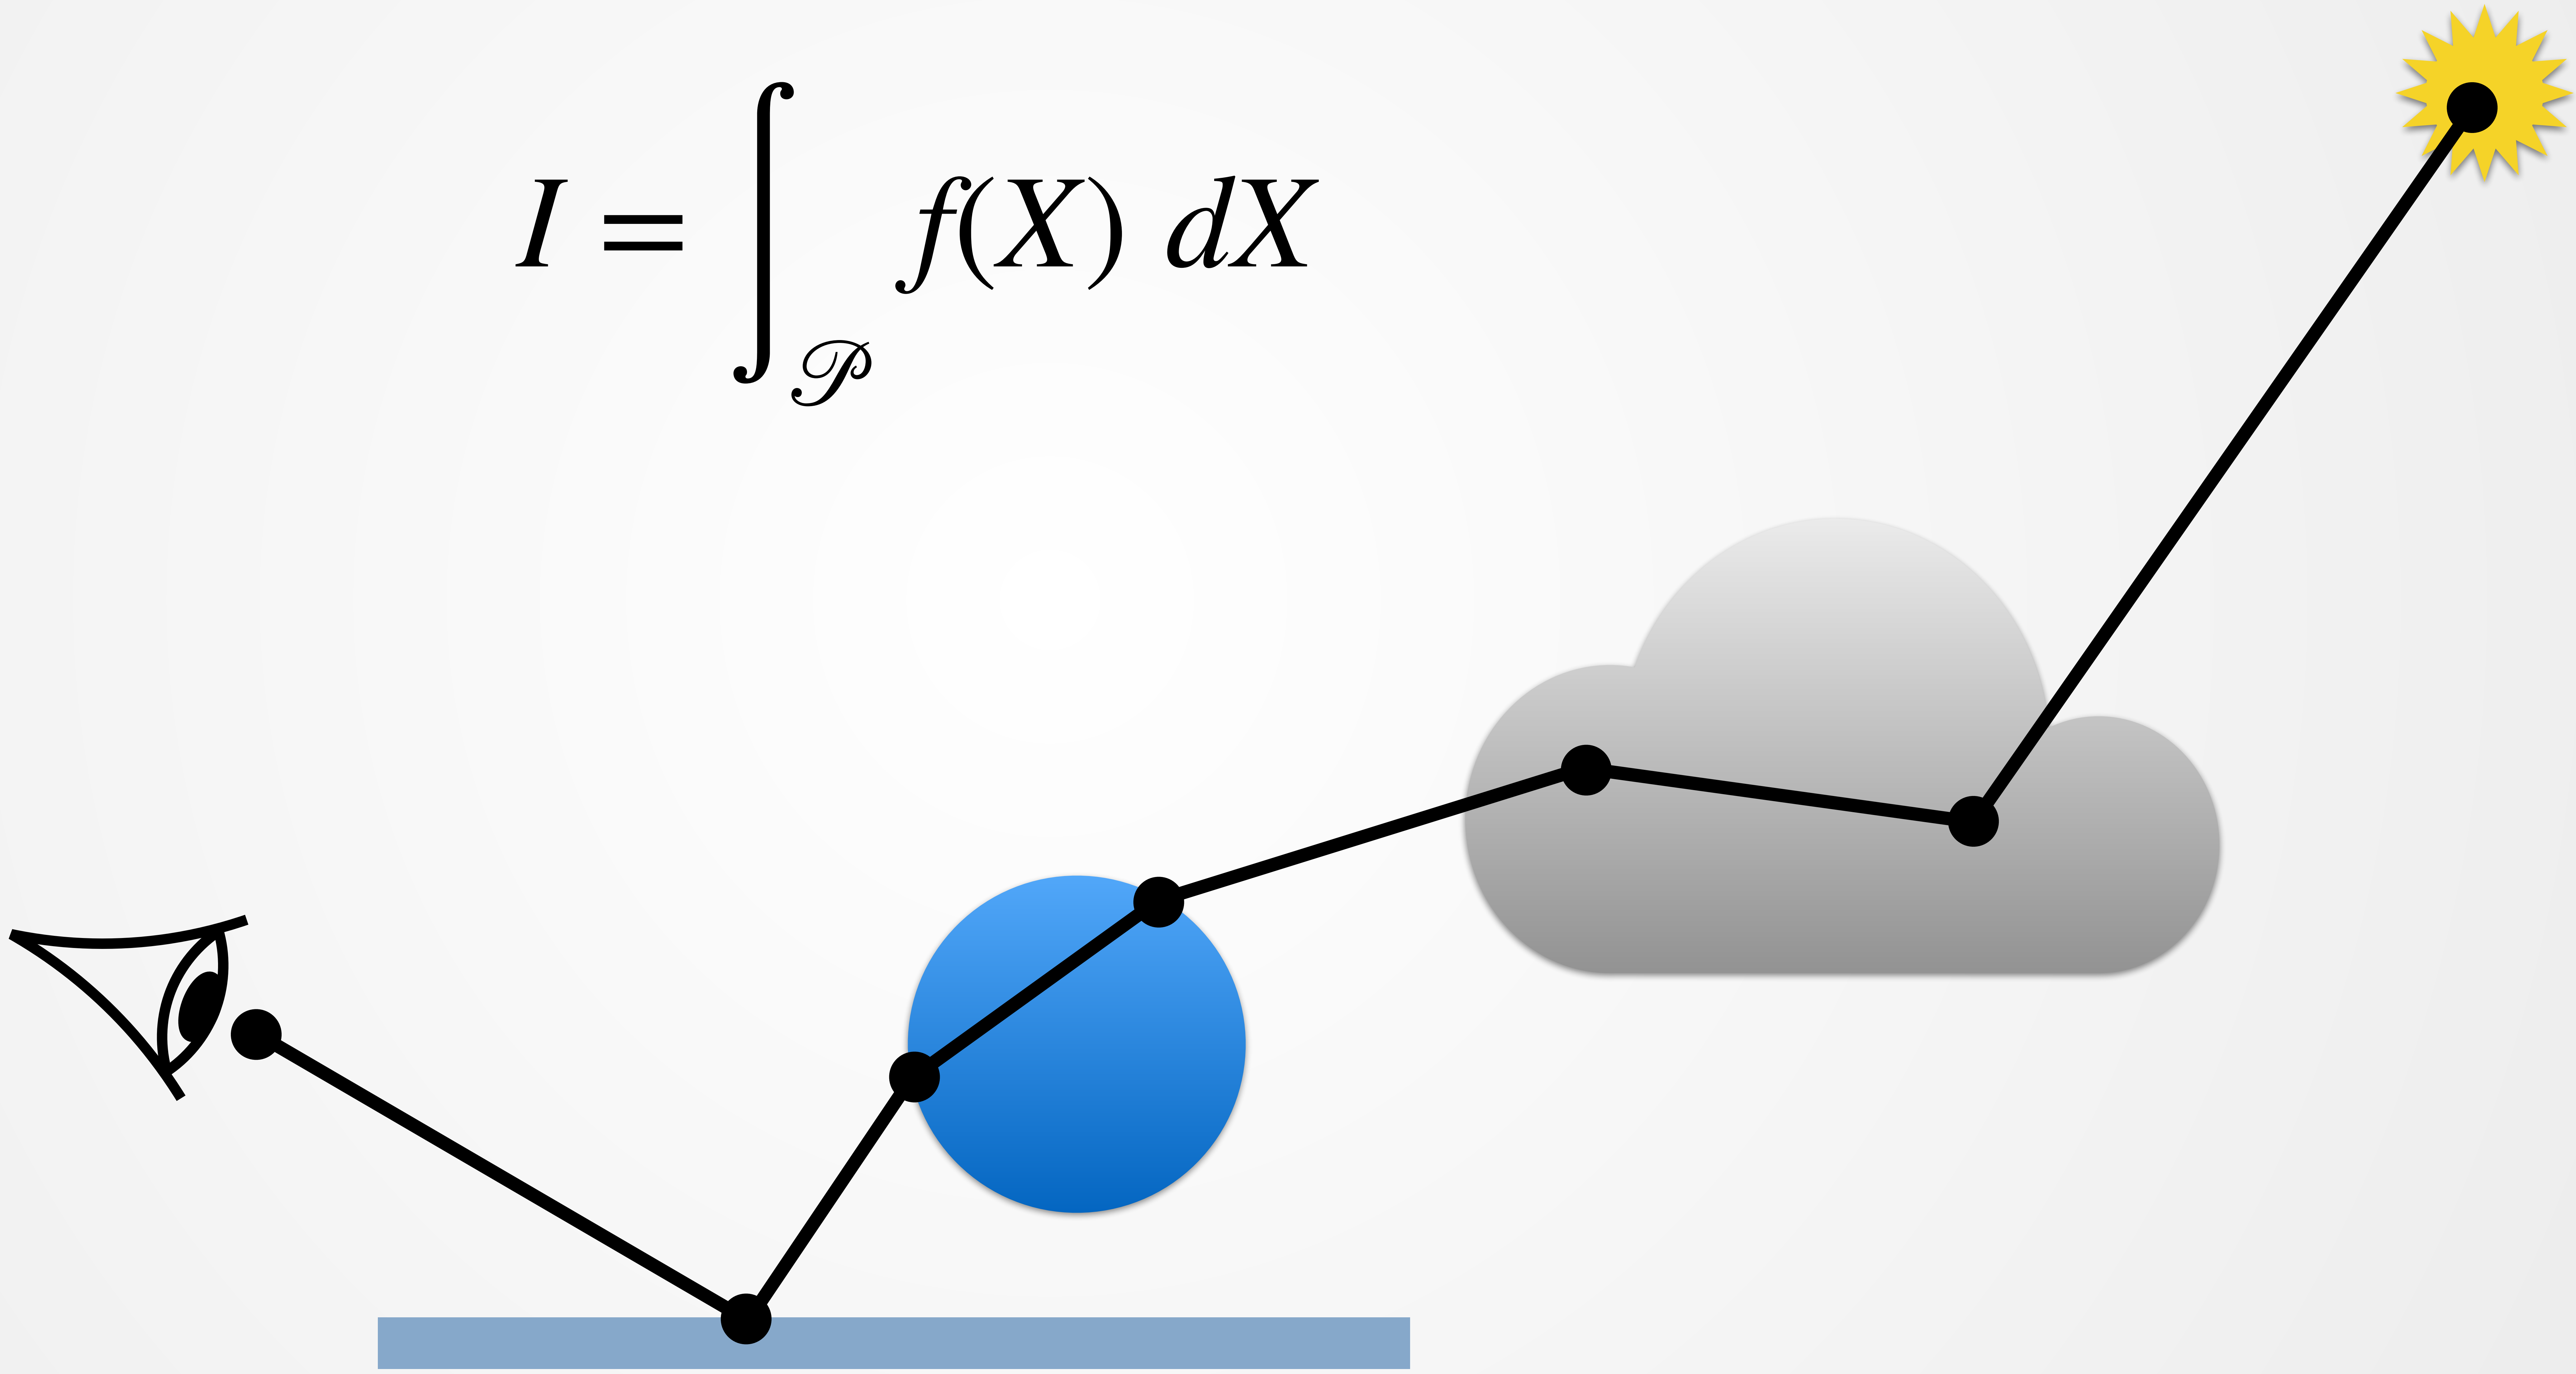

Monte Carlo Integration

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

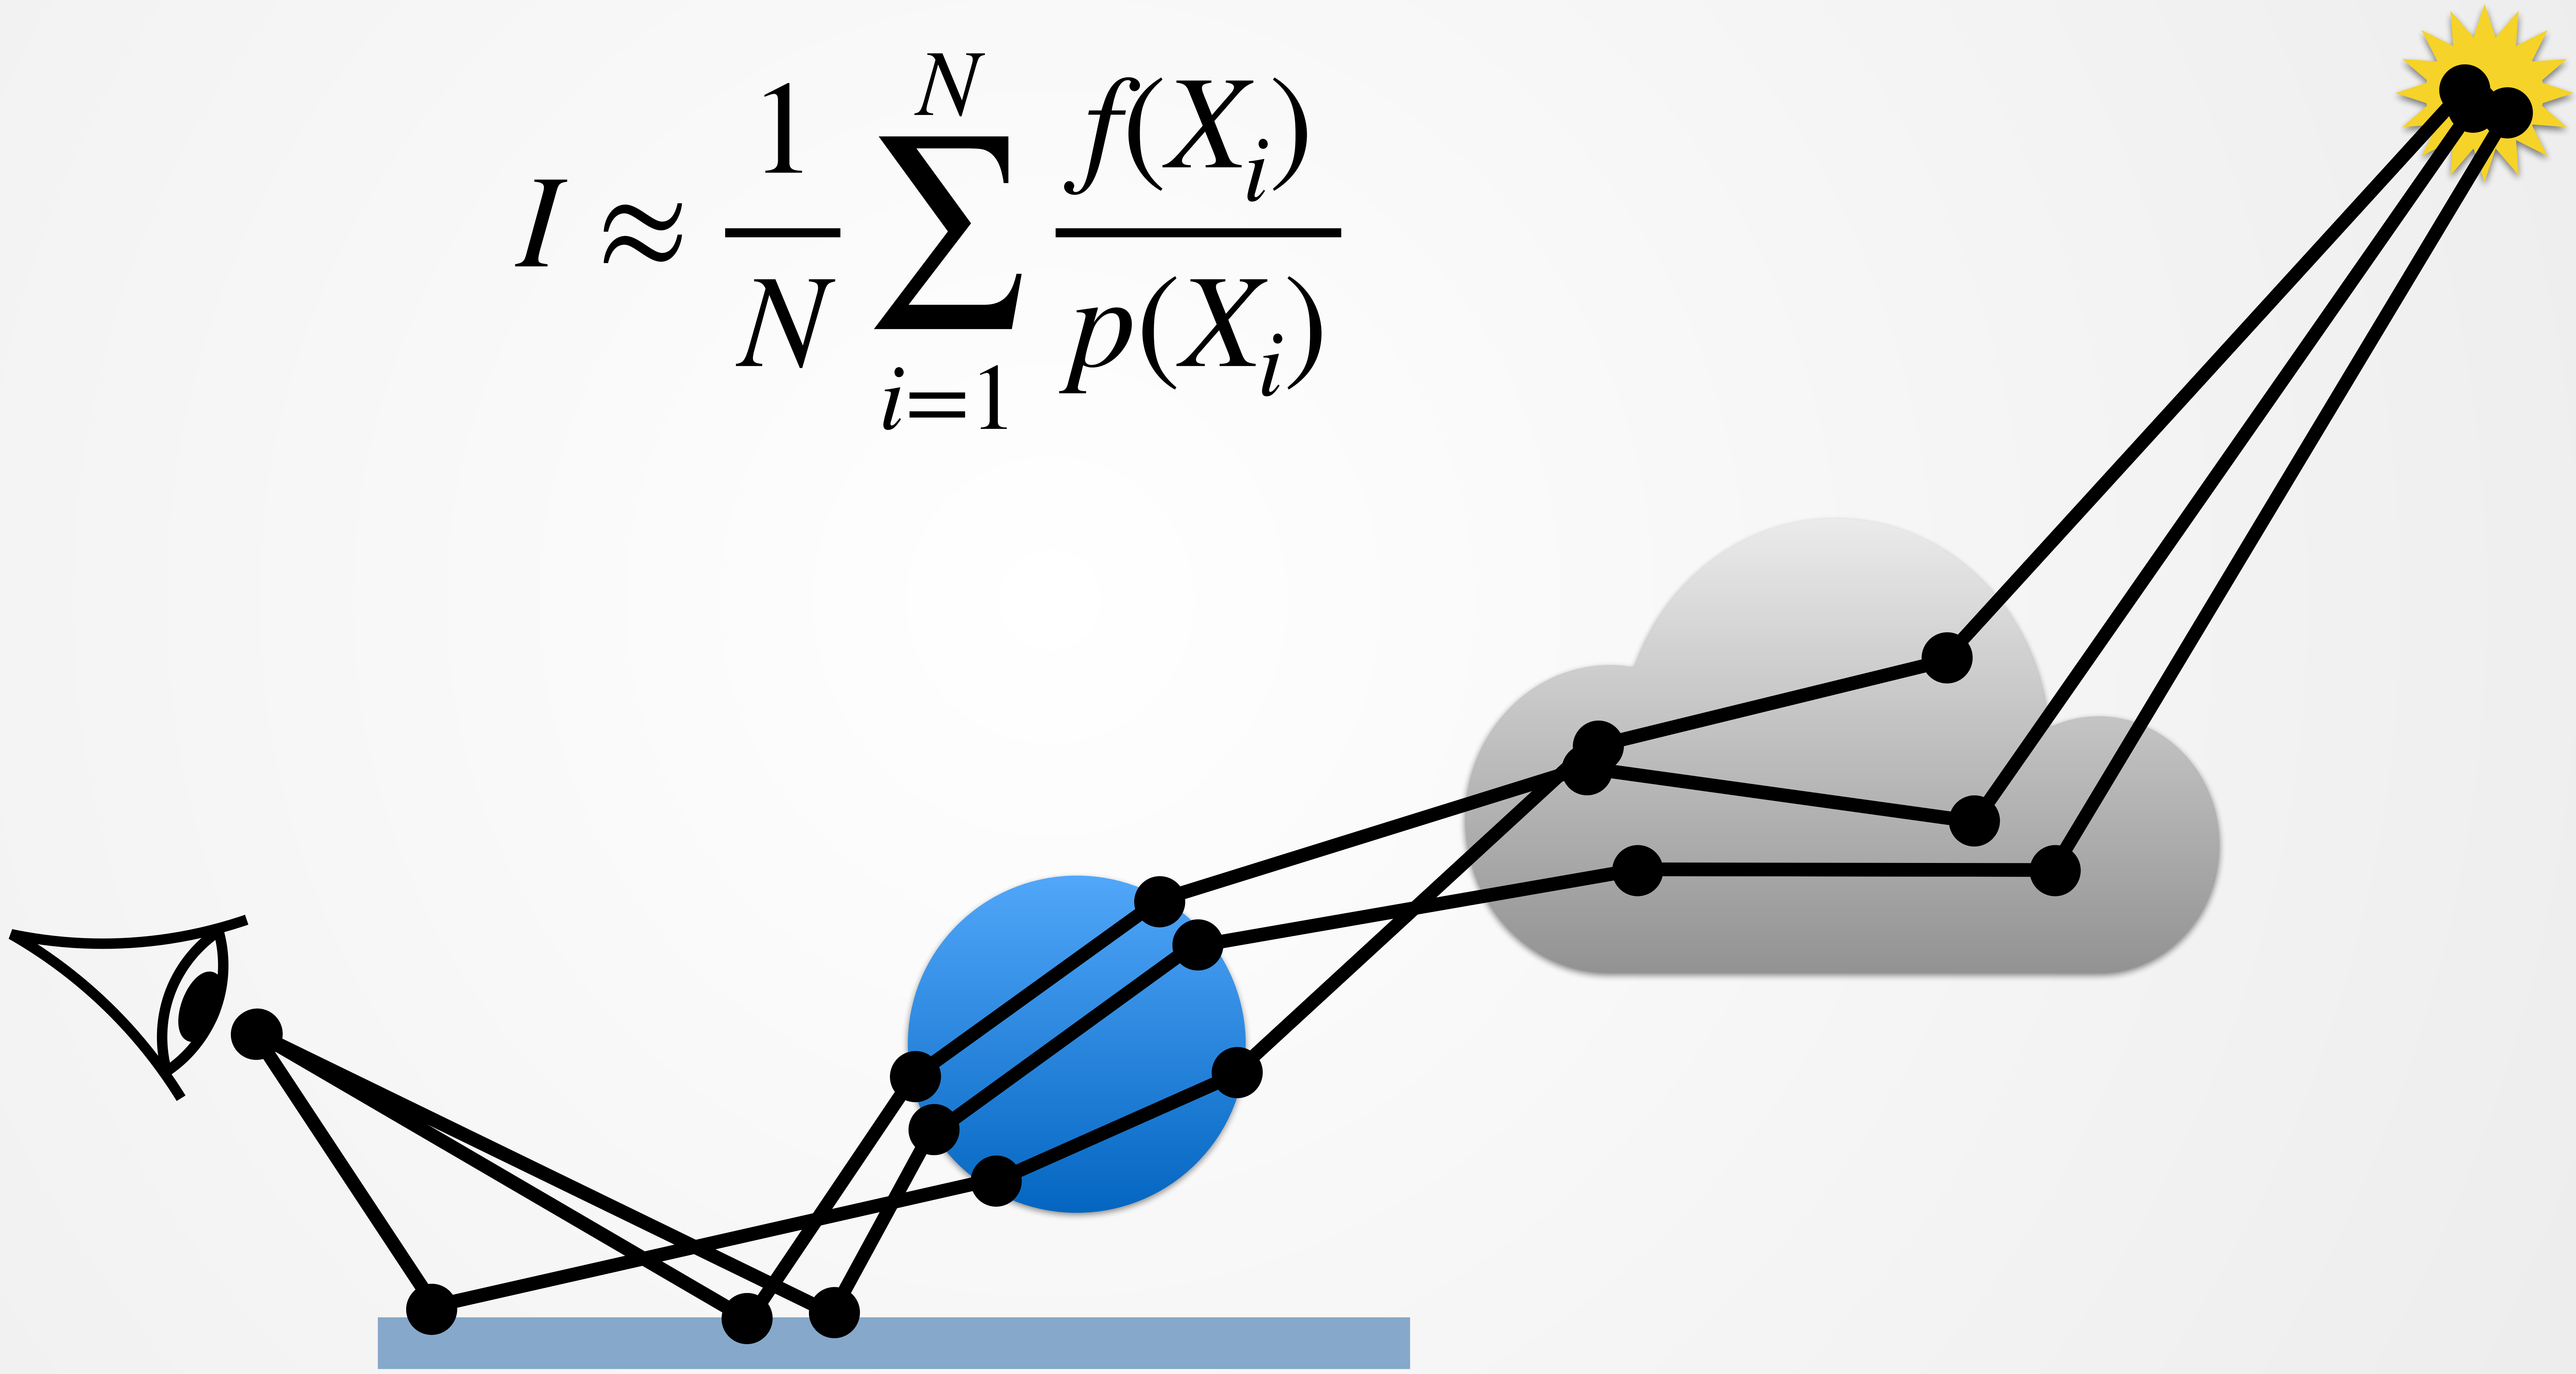

Illustration in 1D

- Importance sampling?

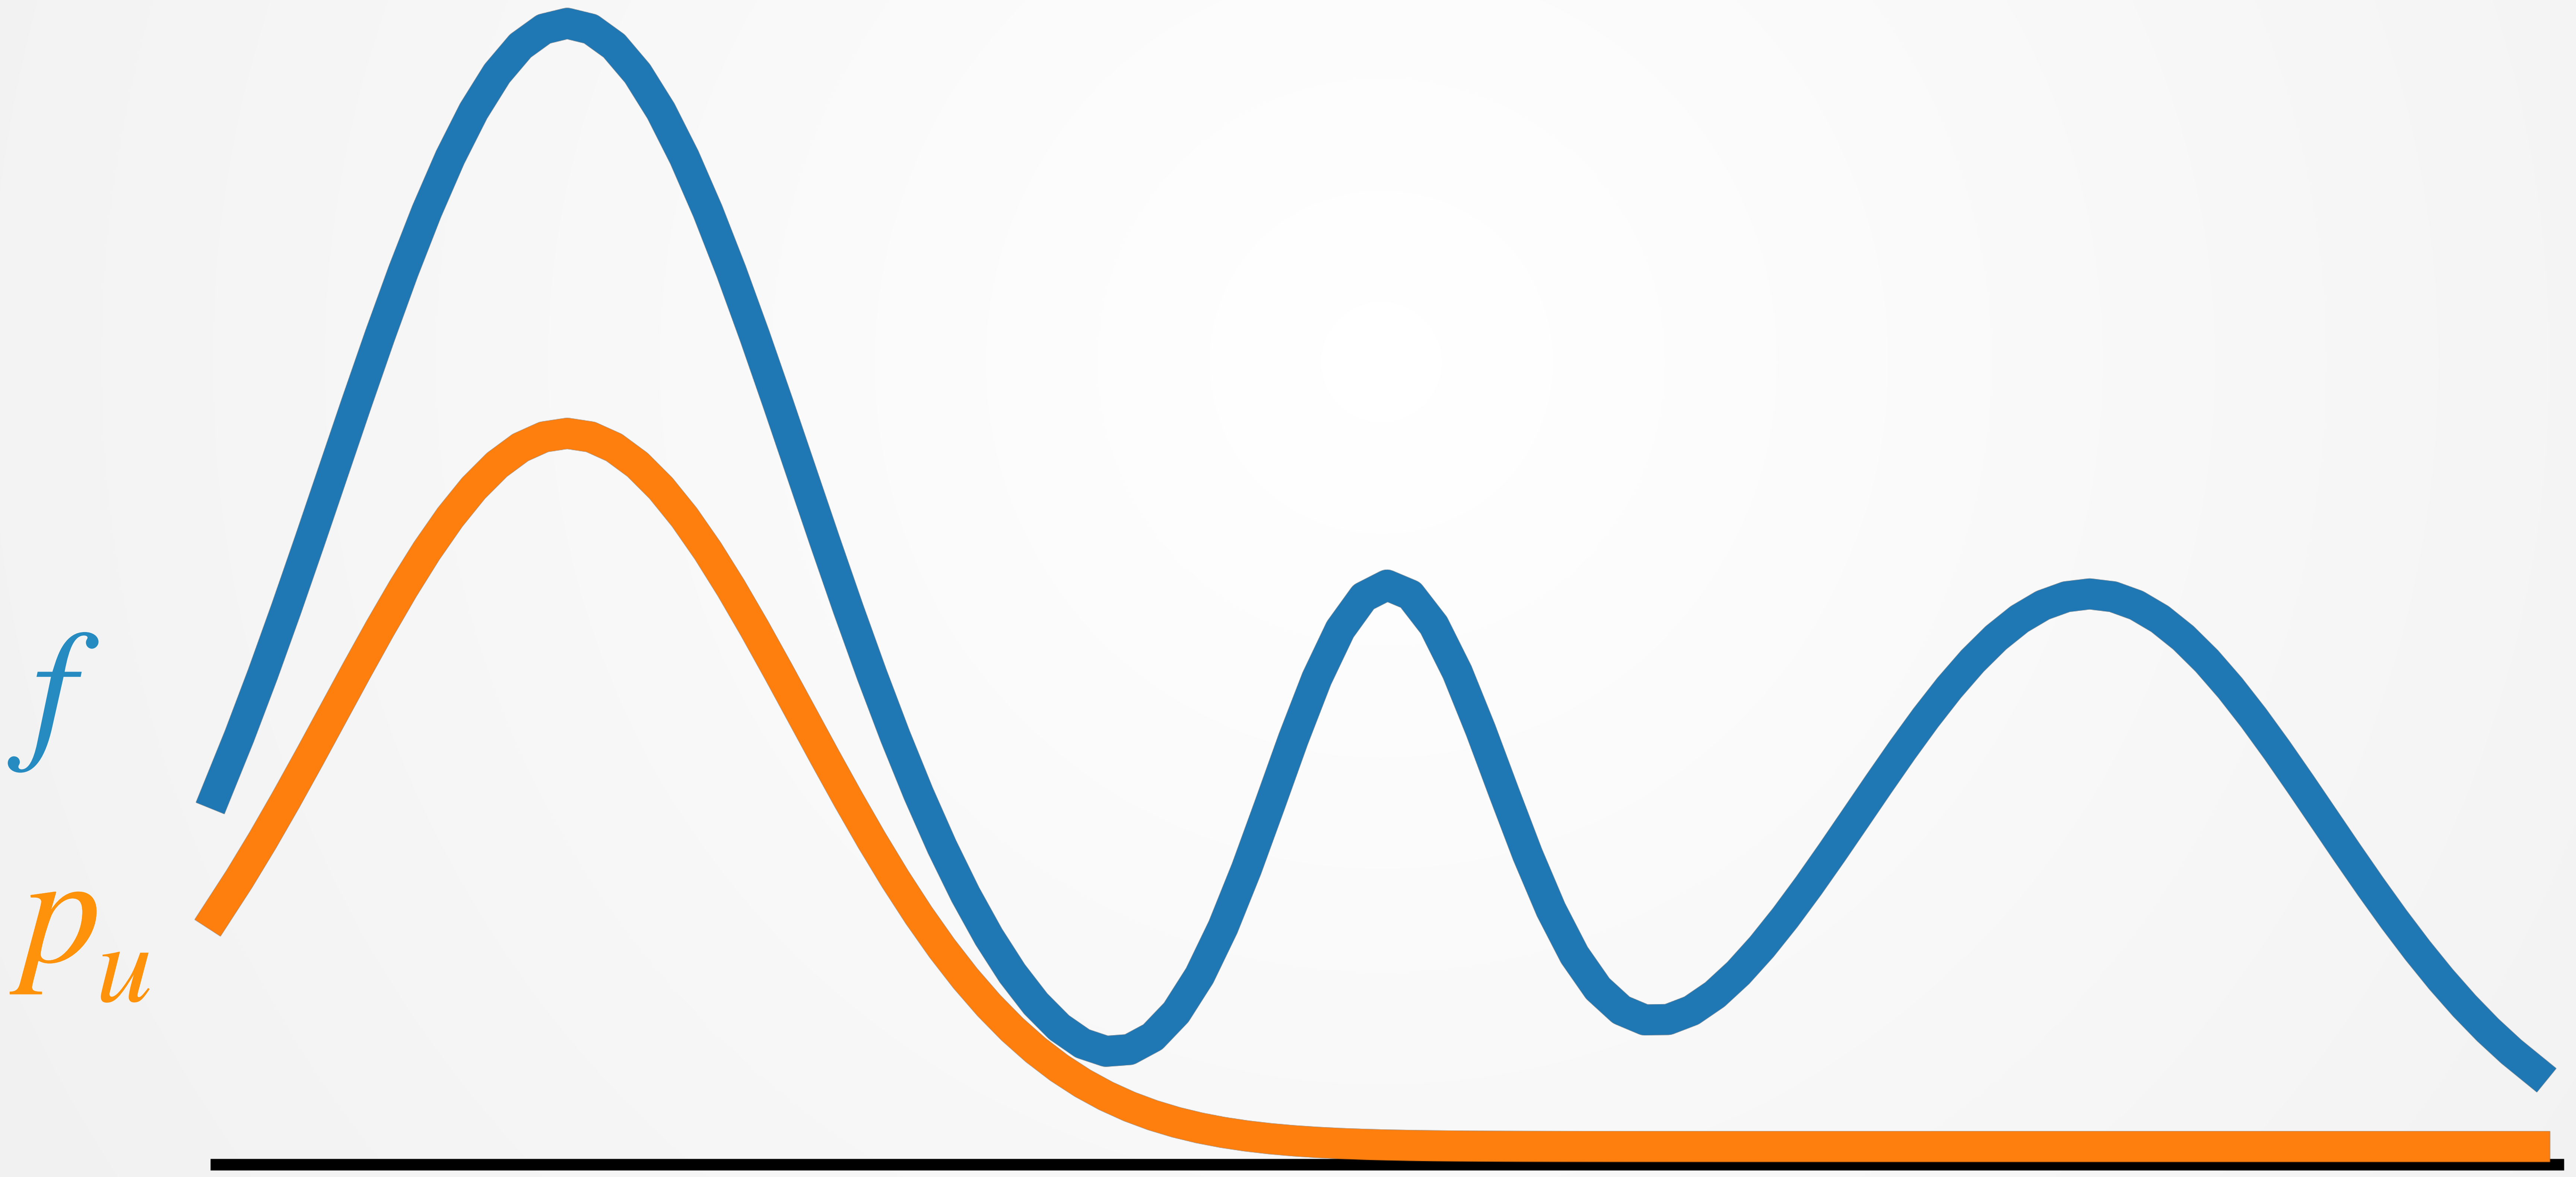

Guided Sampling in 1D

- Create samples $X_i \sim p_u$

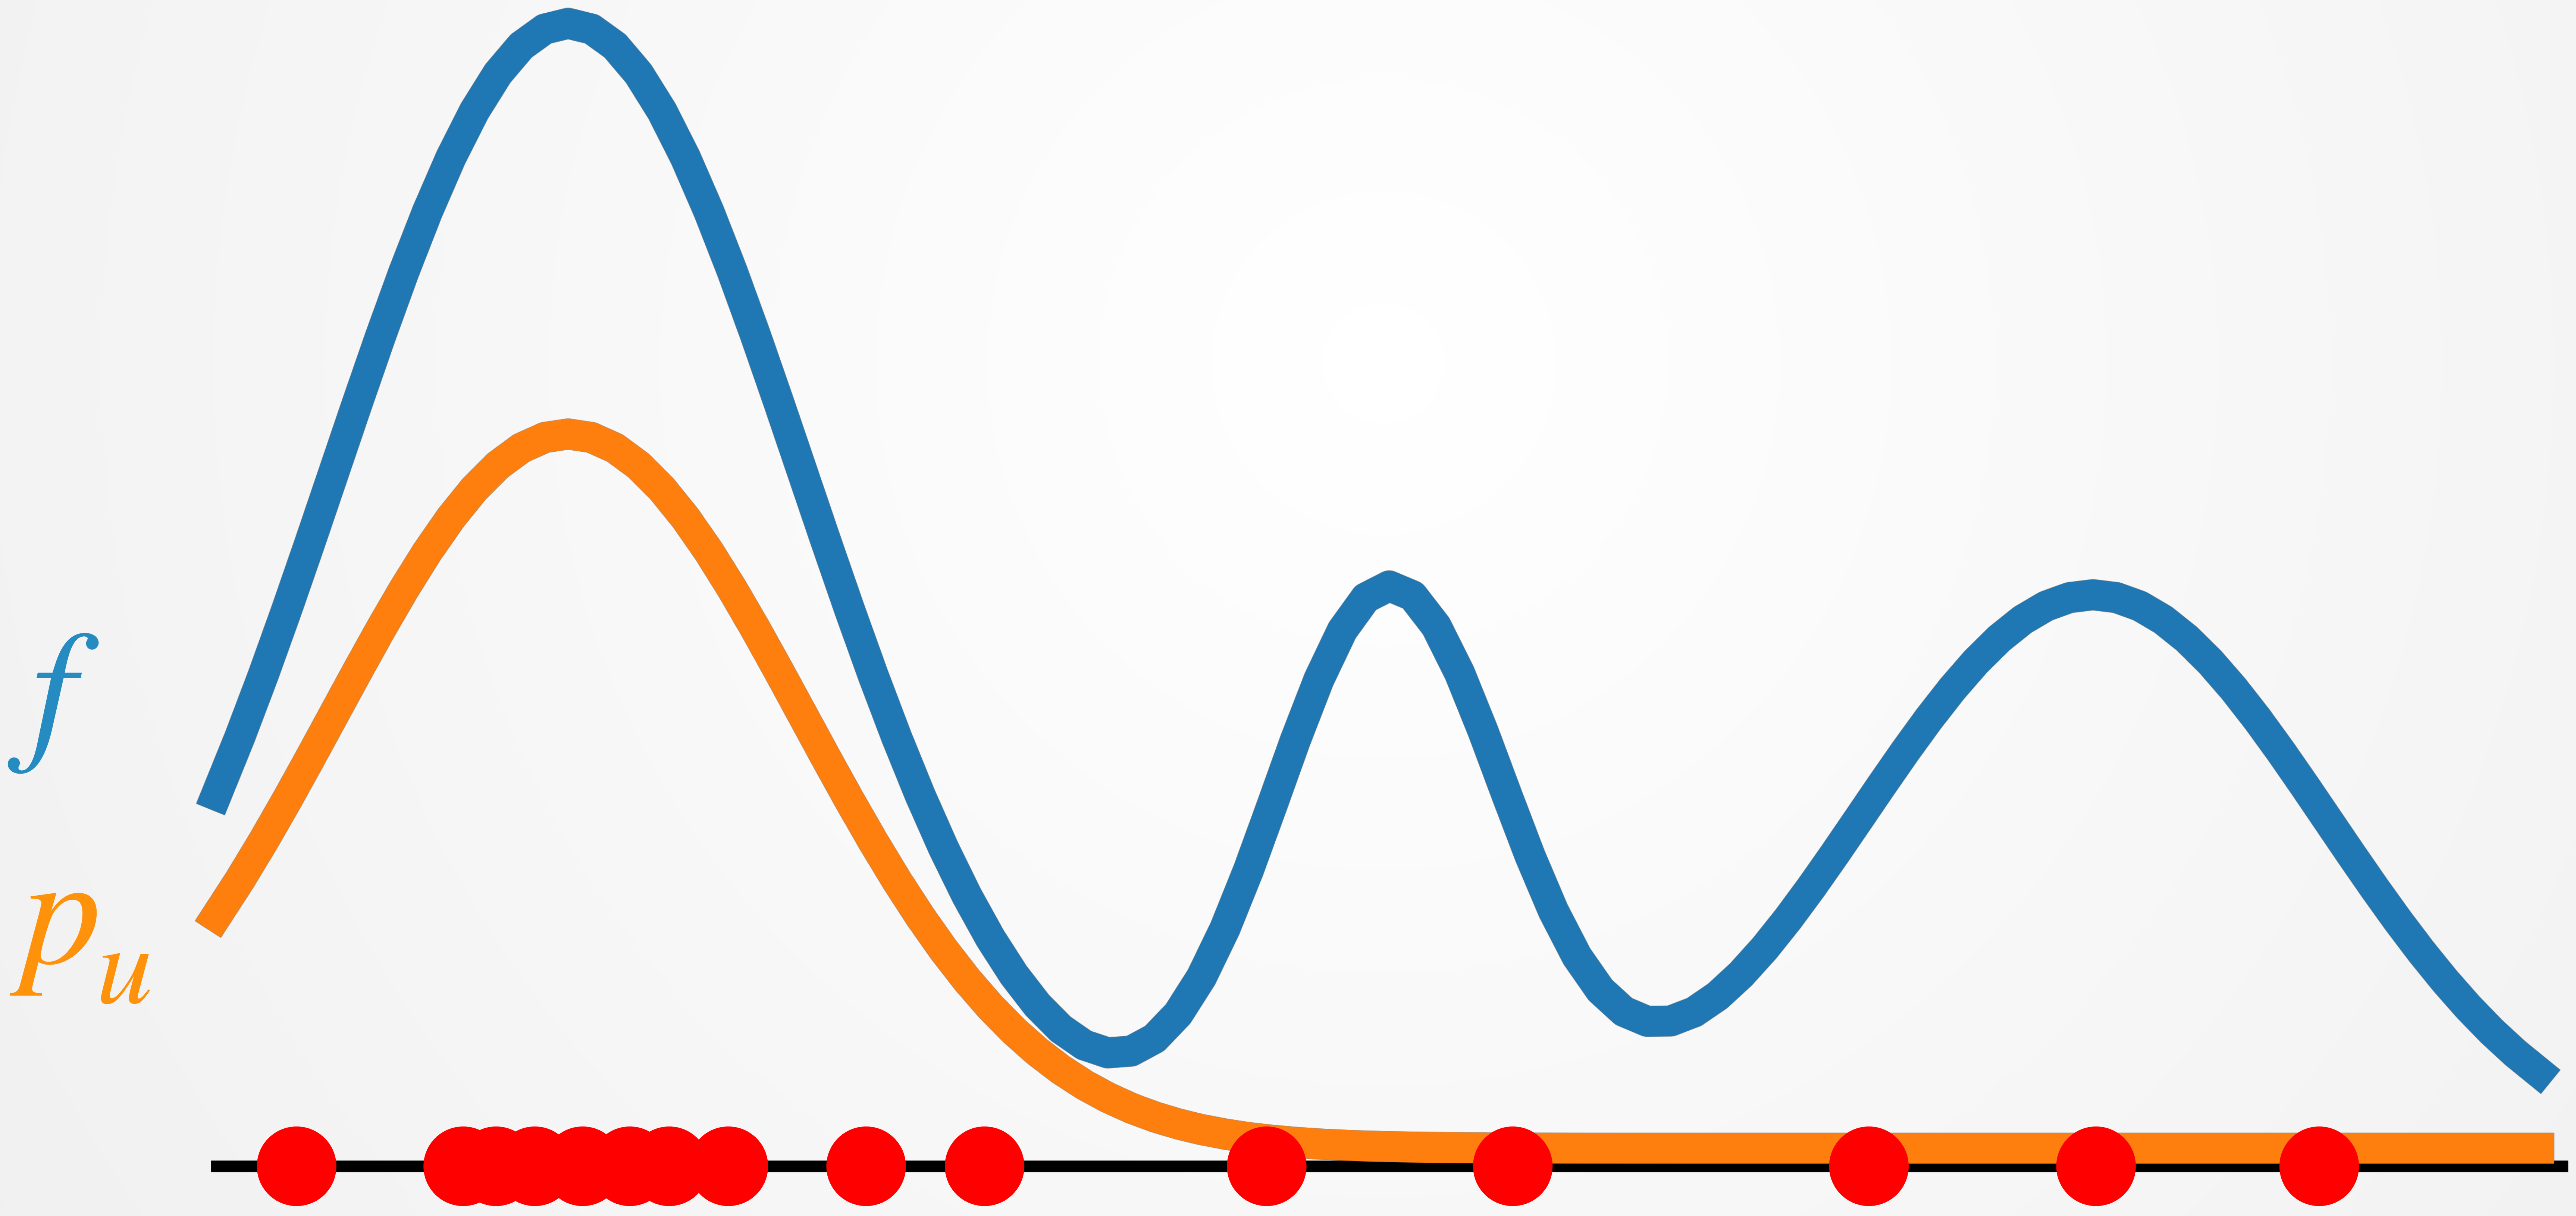

Guided Sampling in 1D

- Keep outliers

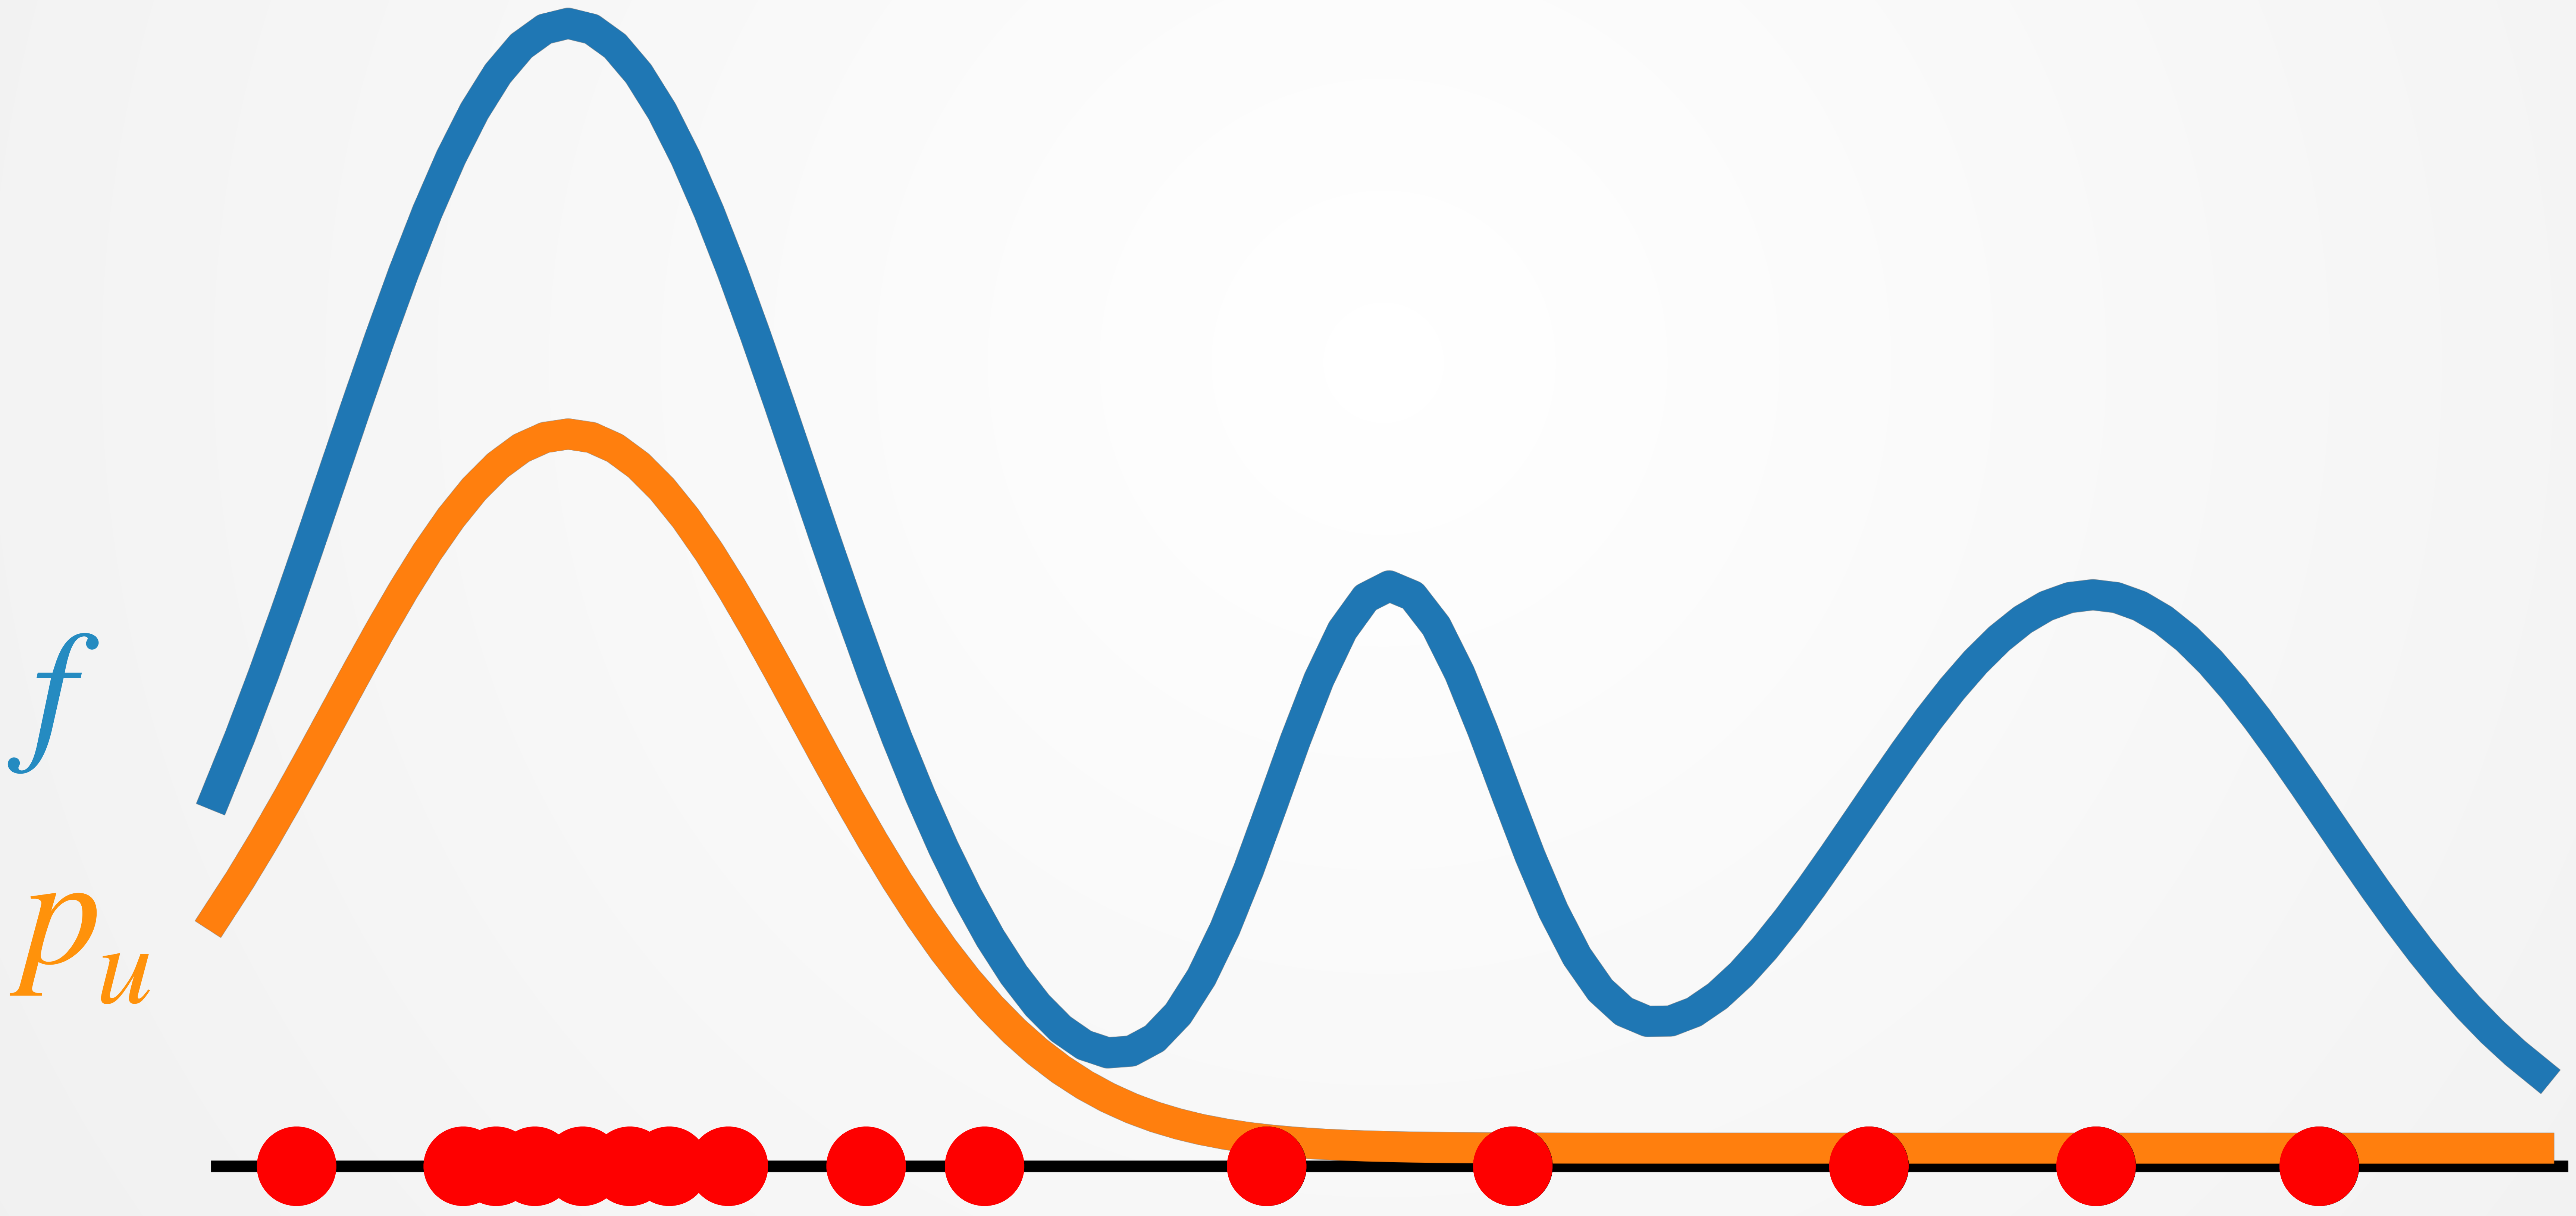

Guided Sampling in 1D

- Place a Gaussian around each outlier

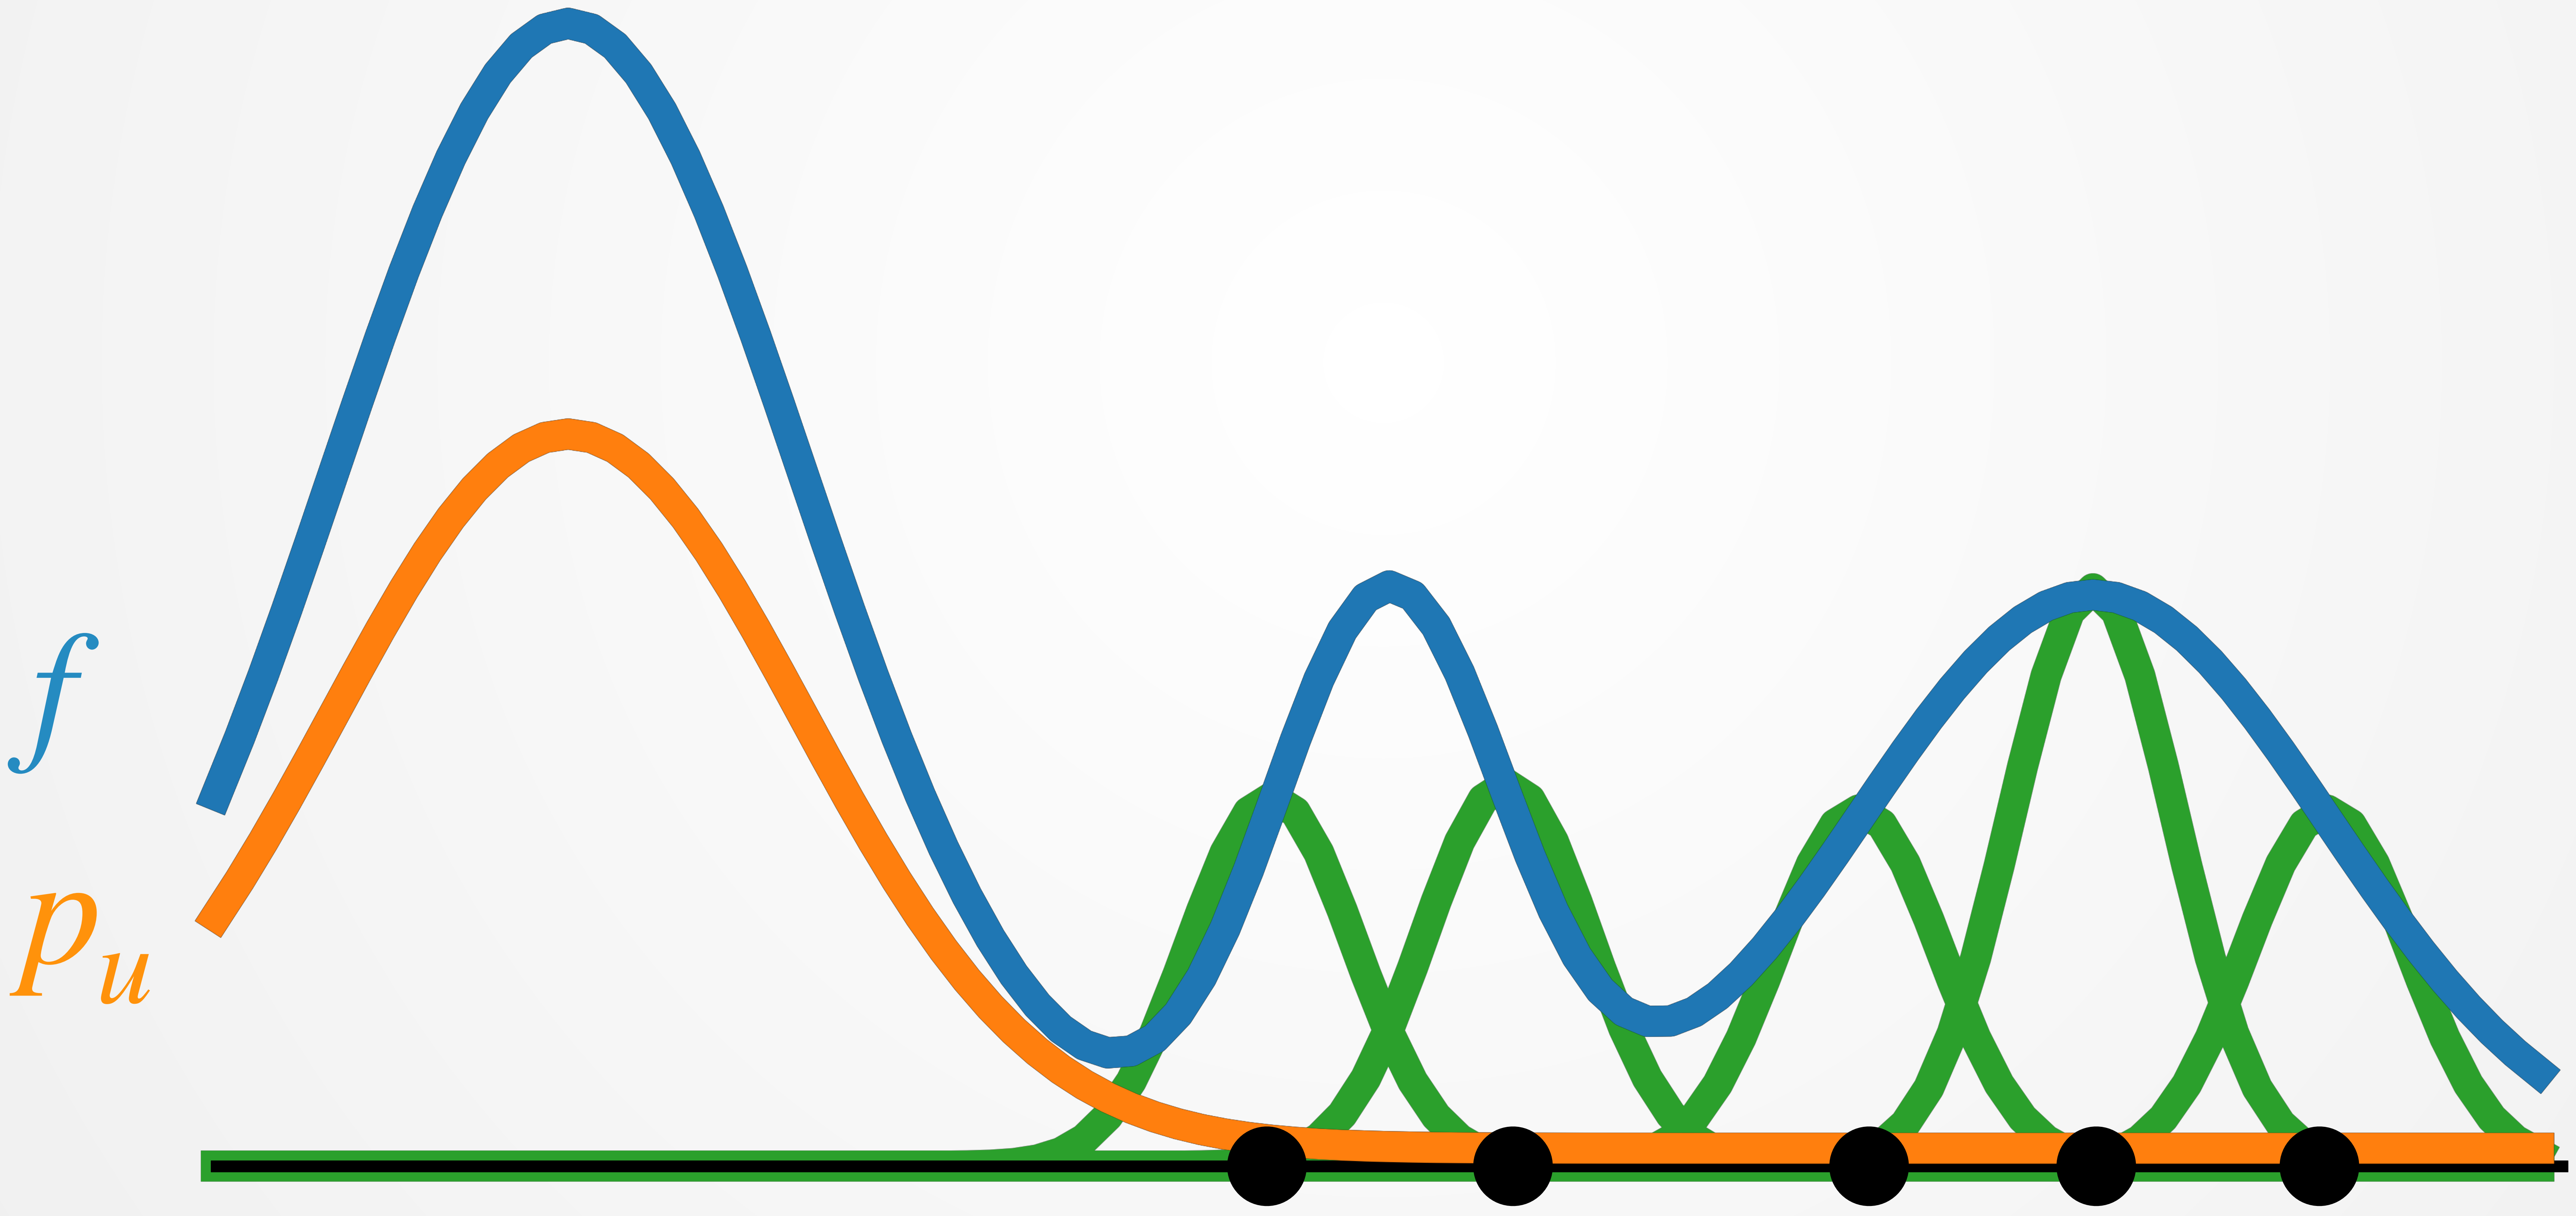

Guided Sampling in 1D

- Define guided PDF as sum over all Gaussians

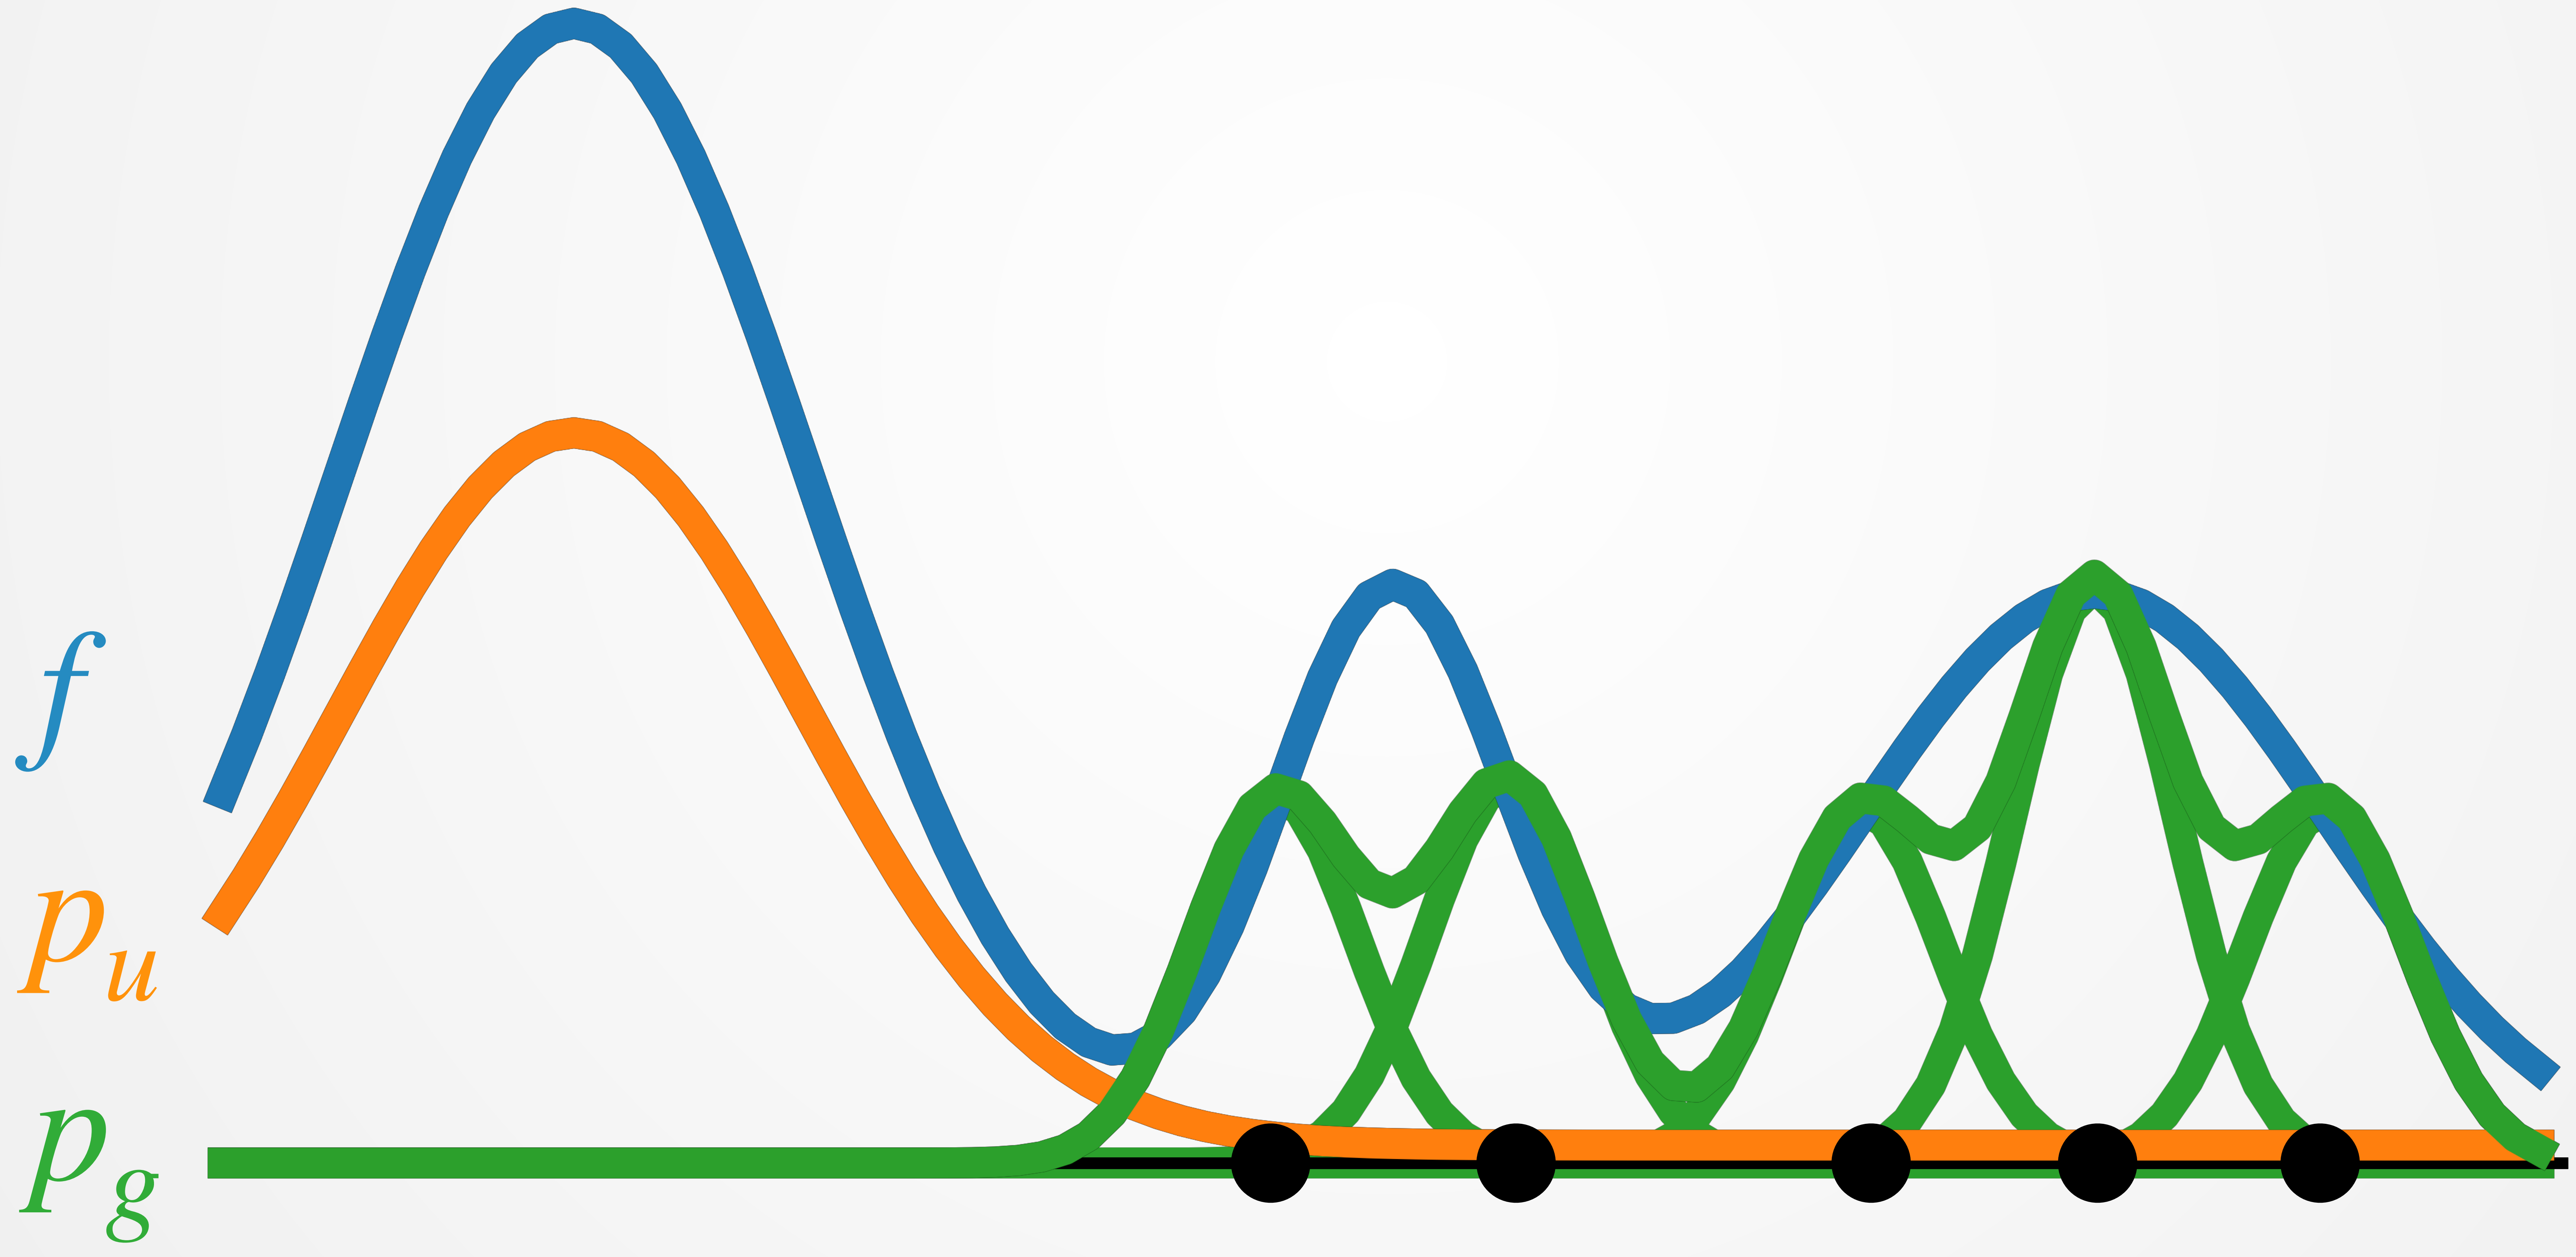

Guided Sampling in 1D

- Iterate by sampling p_u and p_g

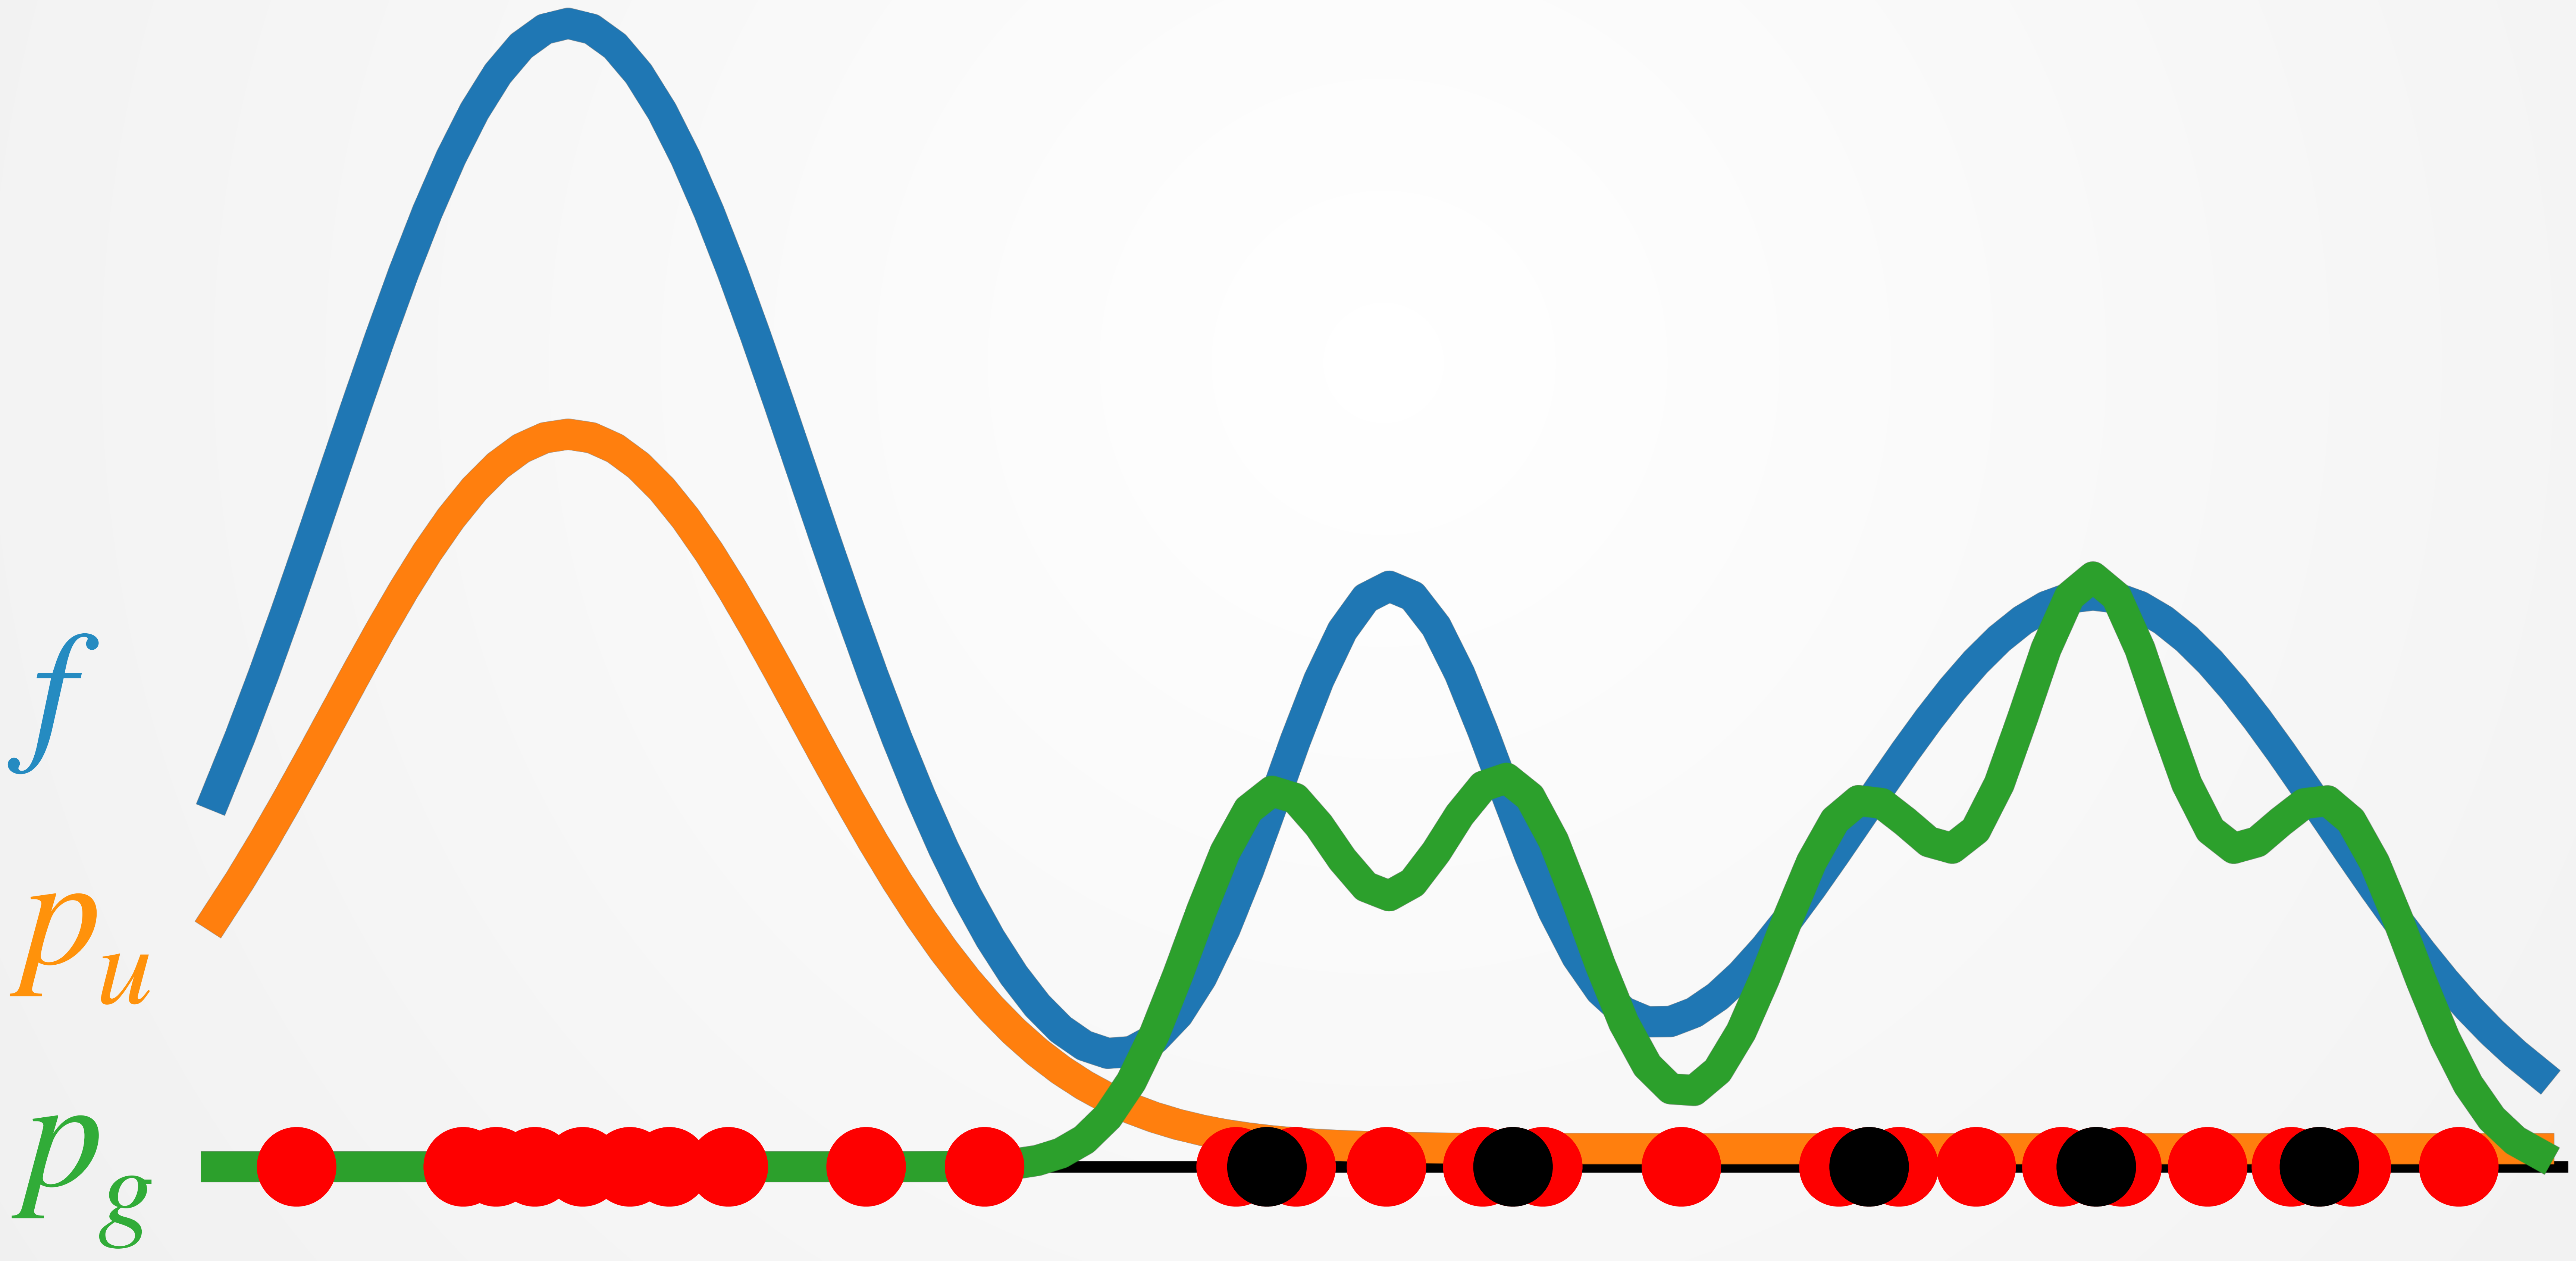

Guided Sampling in 1D

- Keep outliers wrt. p_u and p_g

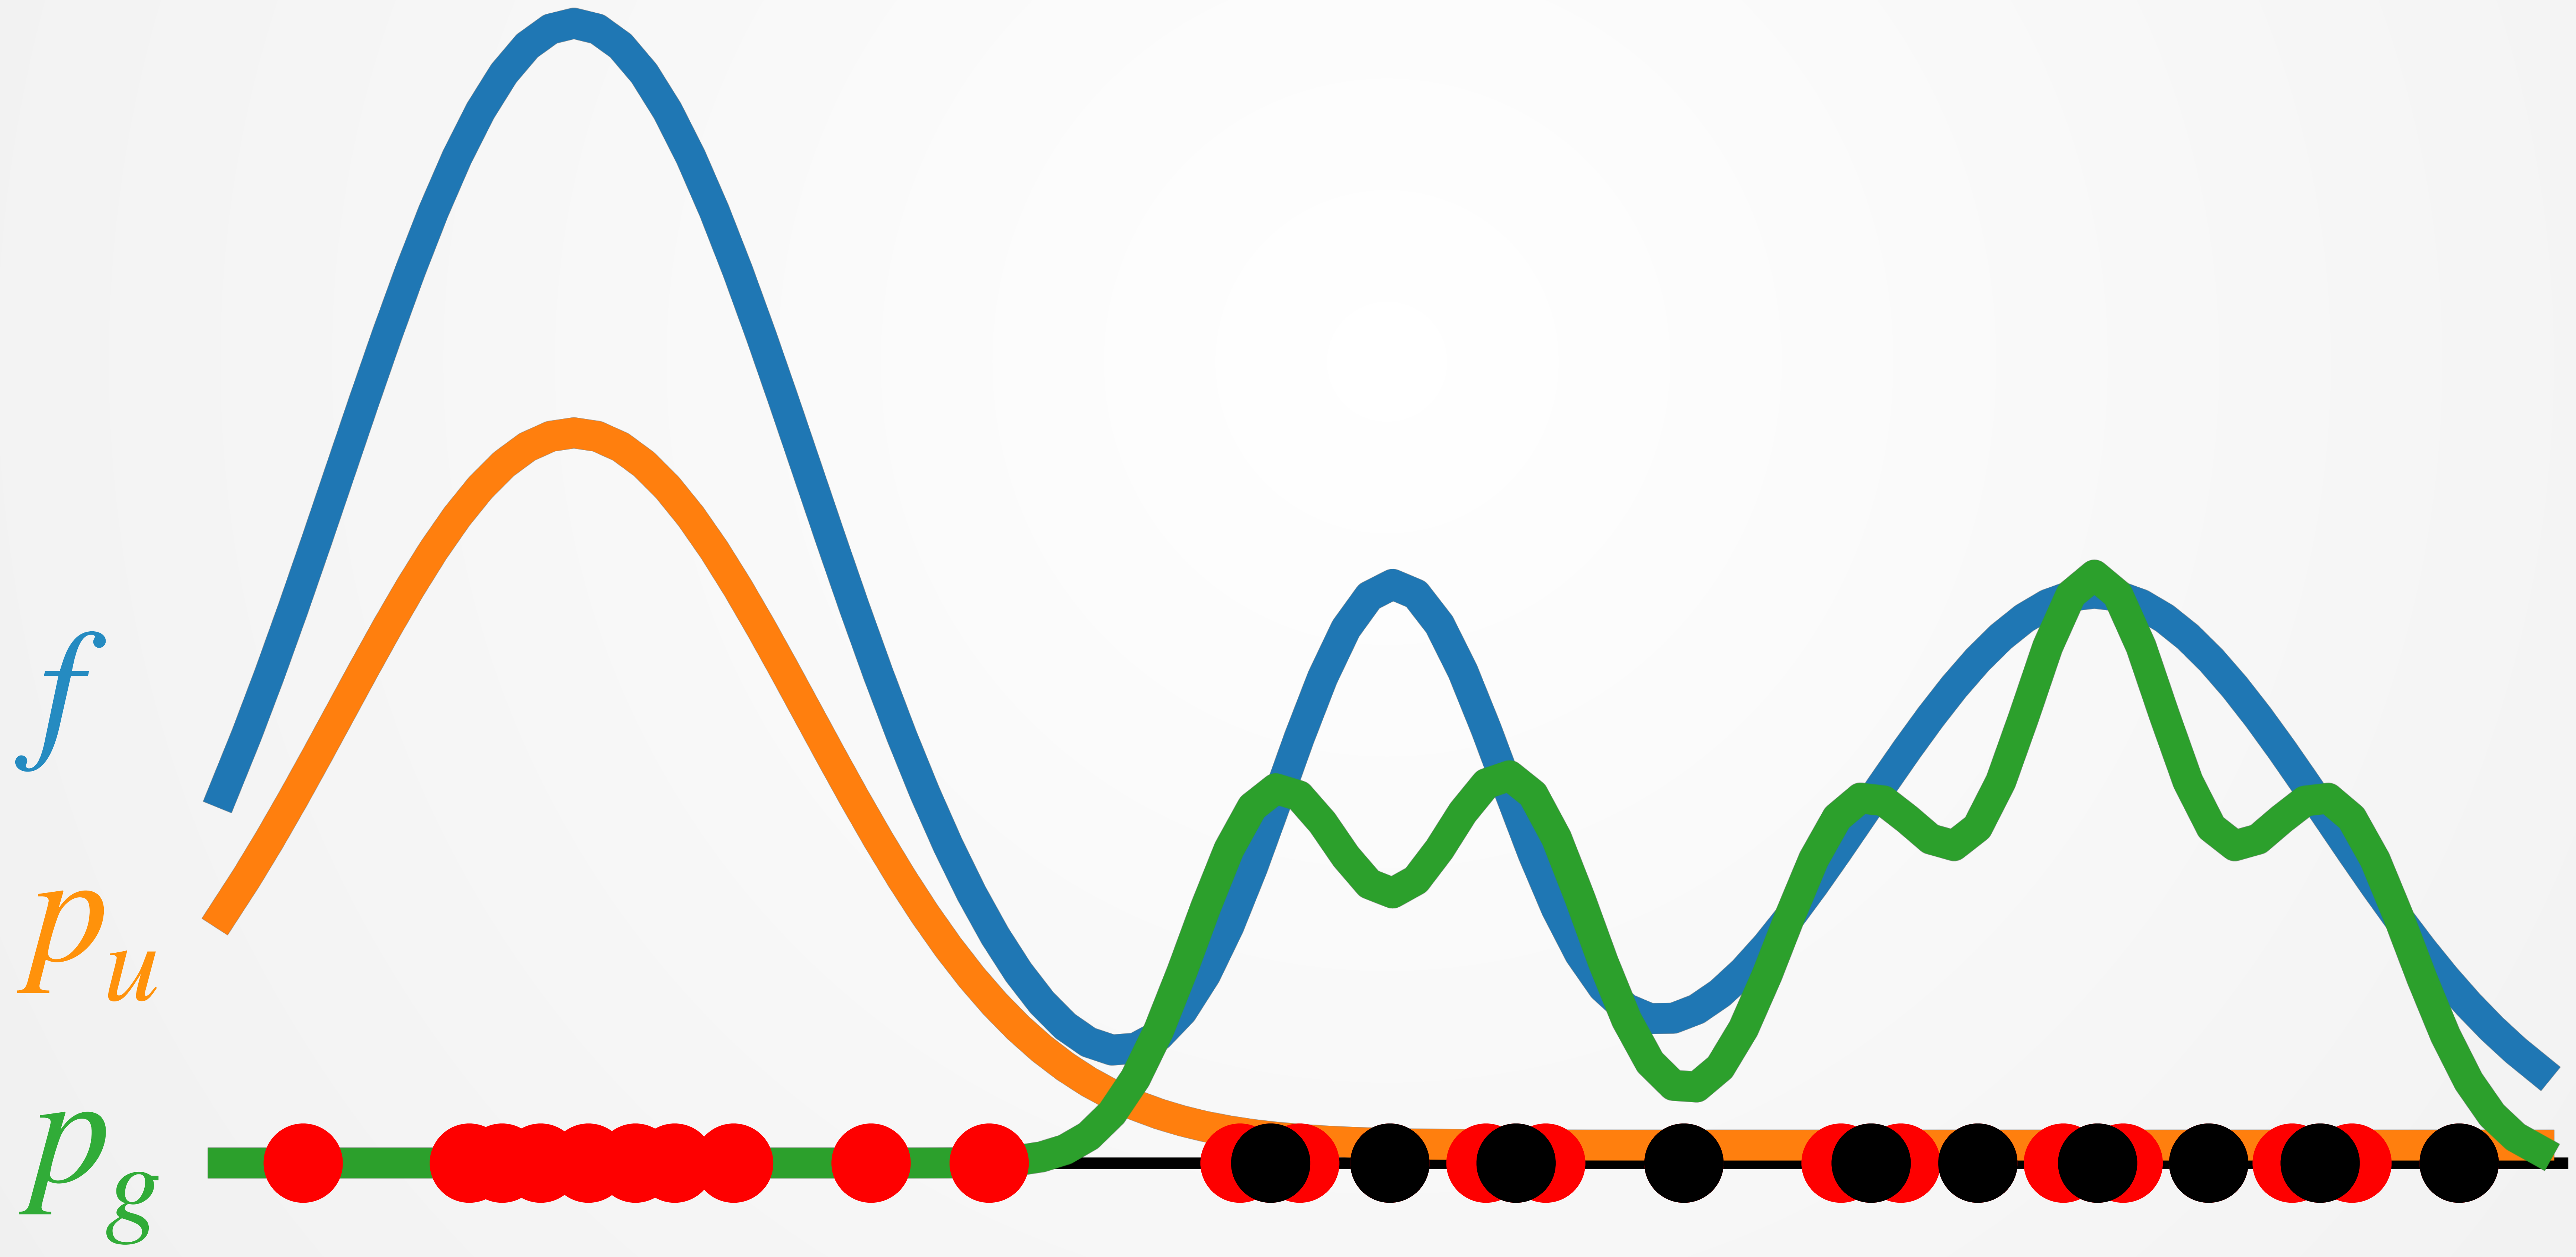

Guided Sampling in 1D

- Update p_g

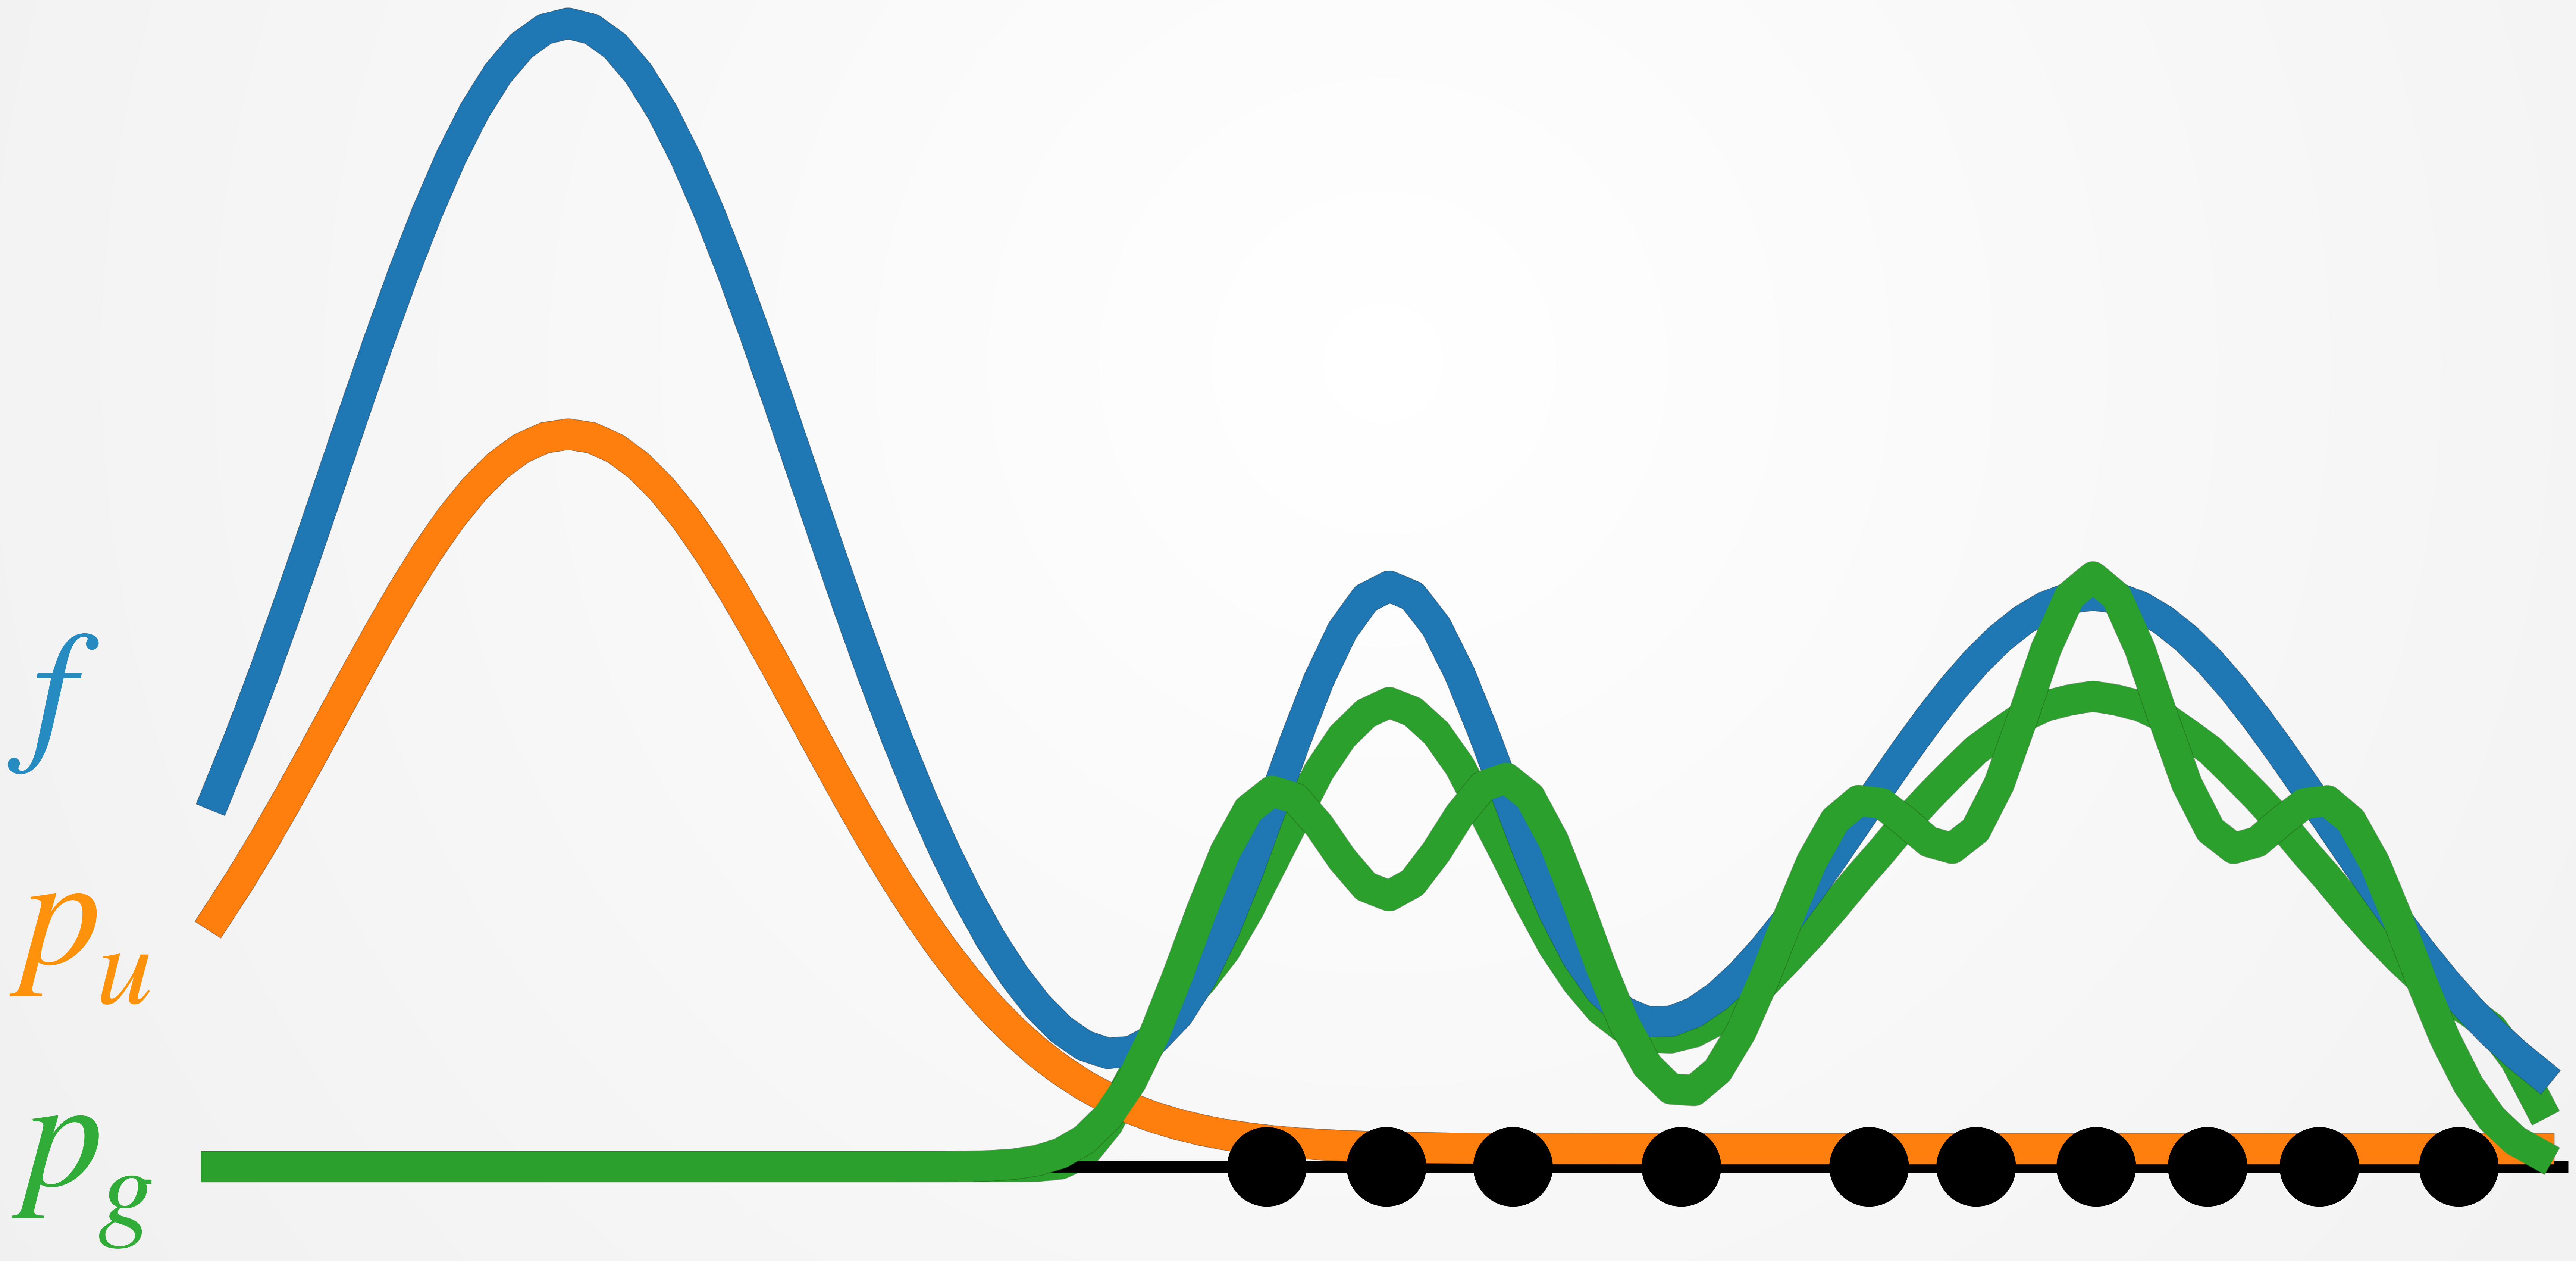

Guiding Overview

- In each iteration, sample paths from p_u and p_g

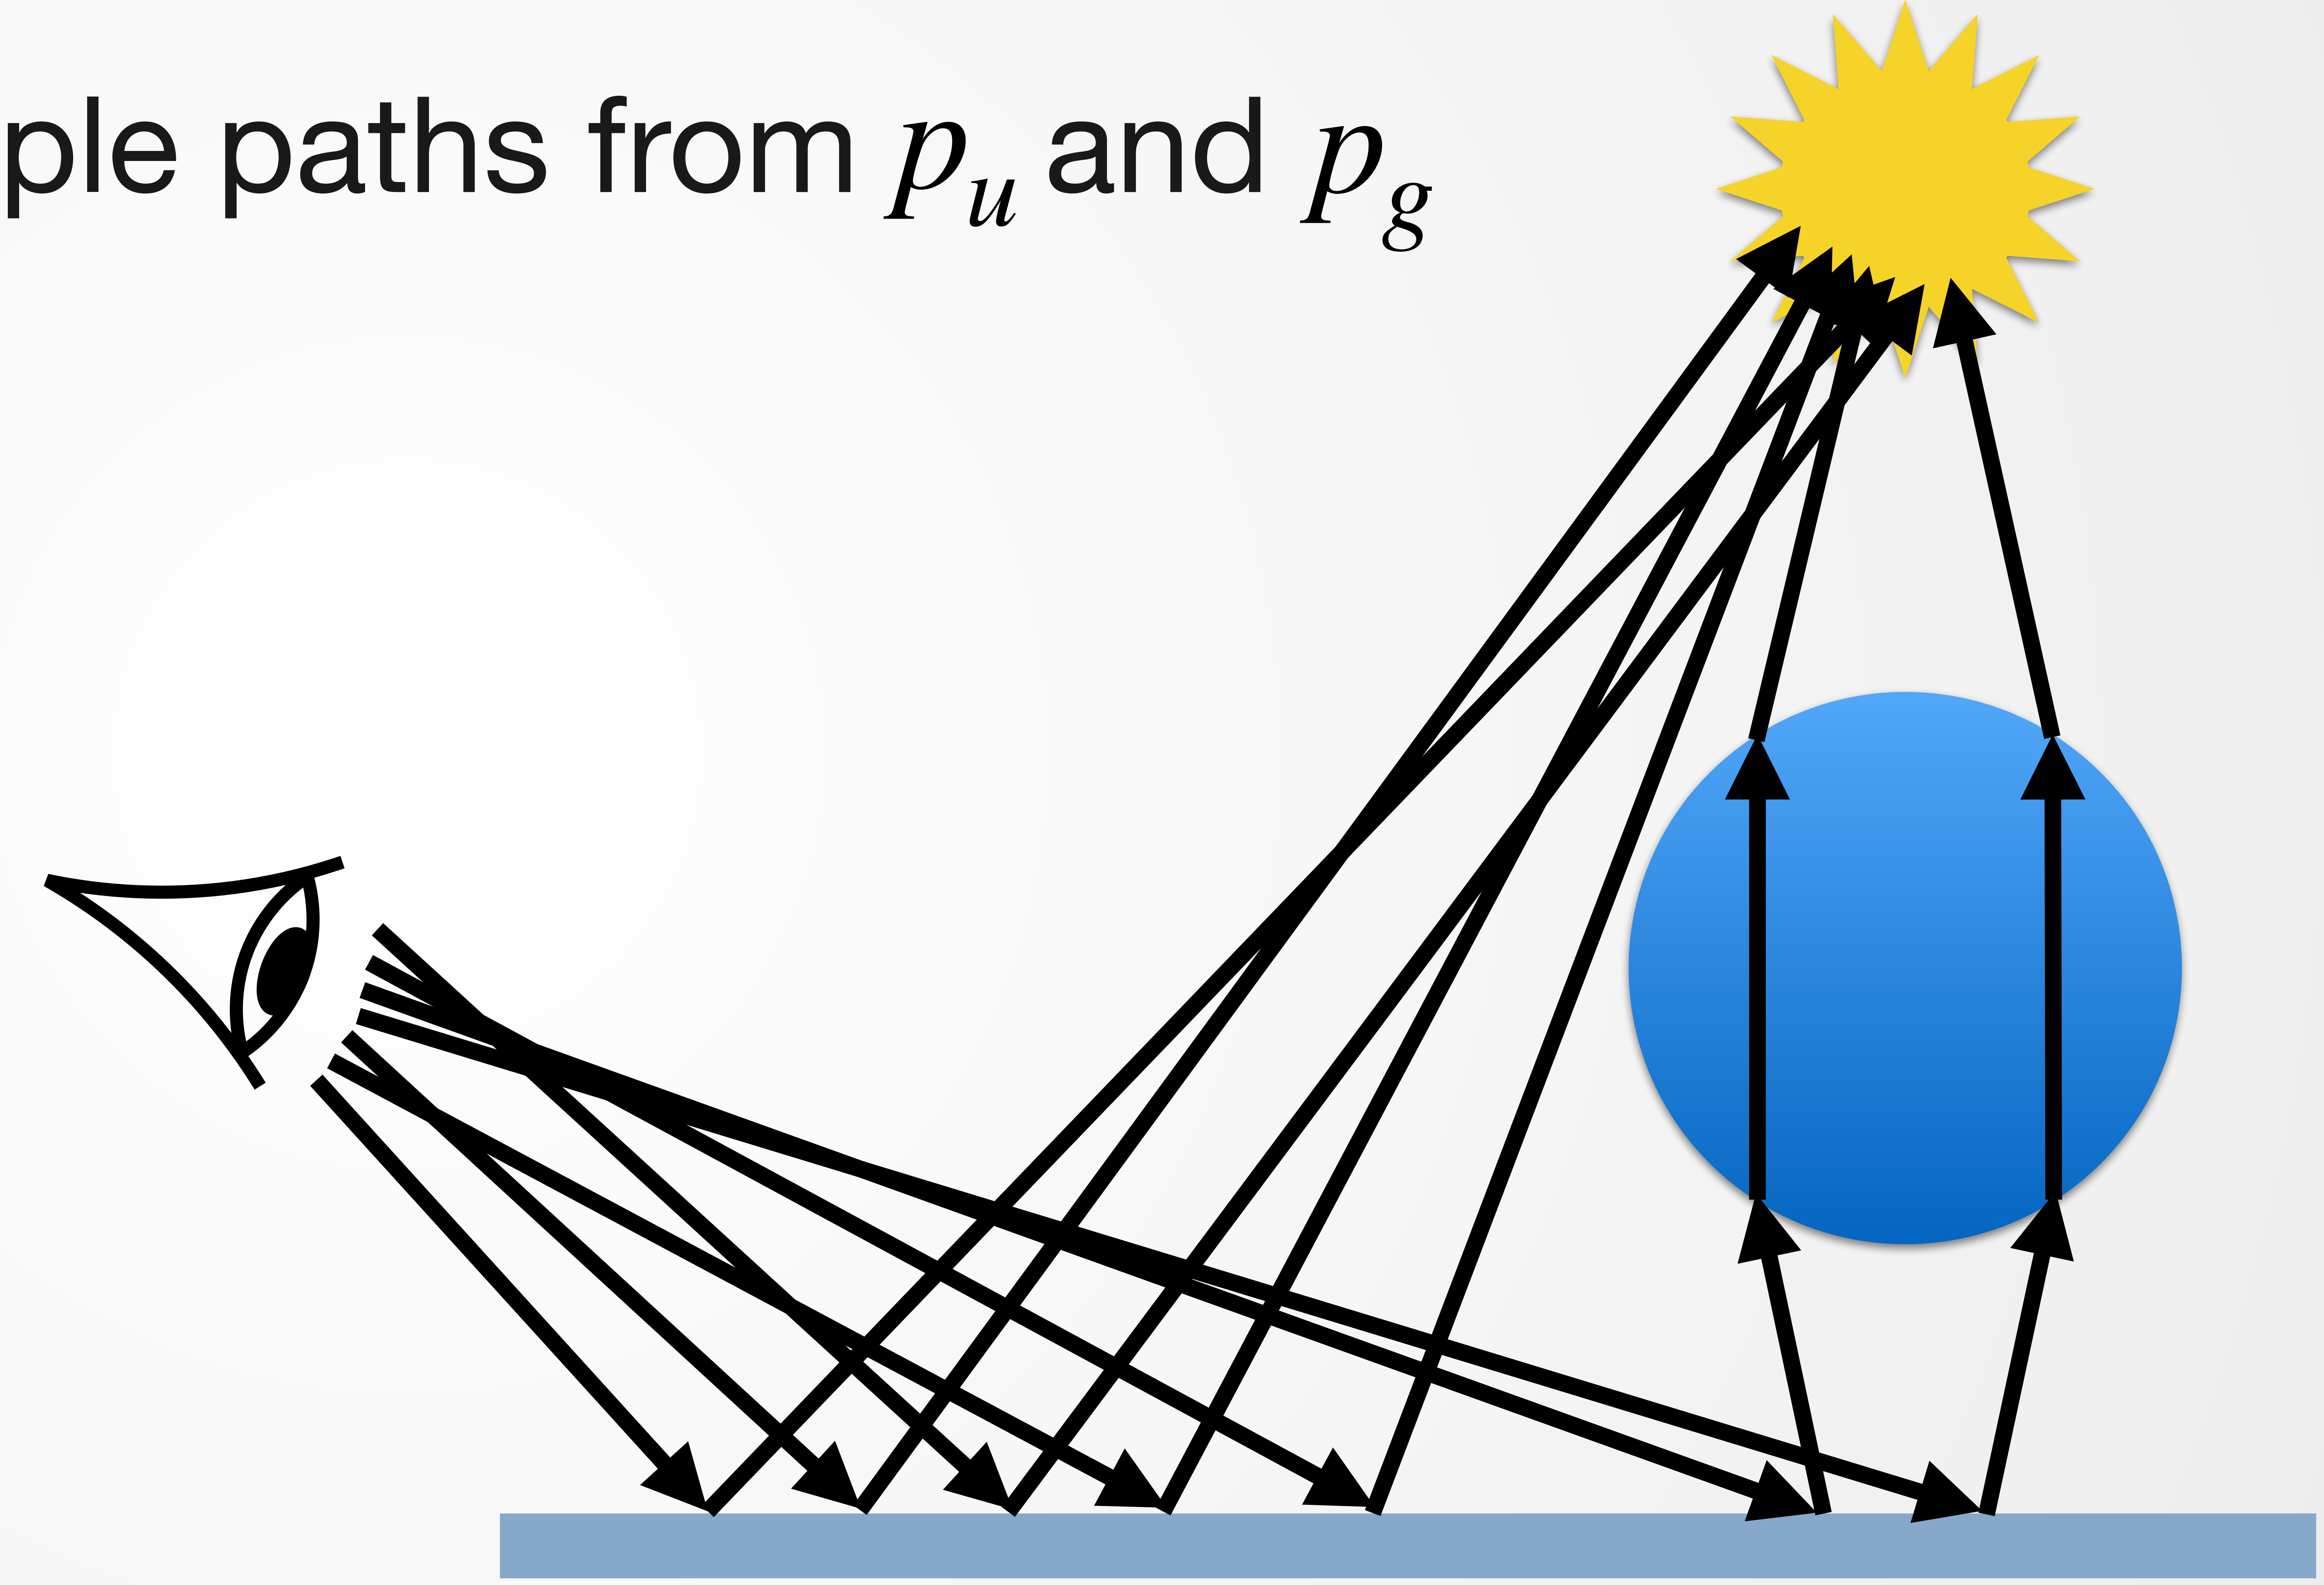

Guiding Overview

- Keep outliers with highest contribution

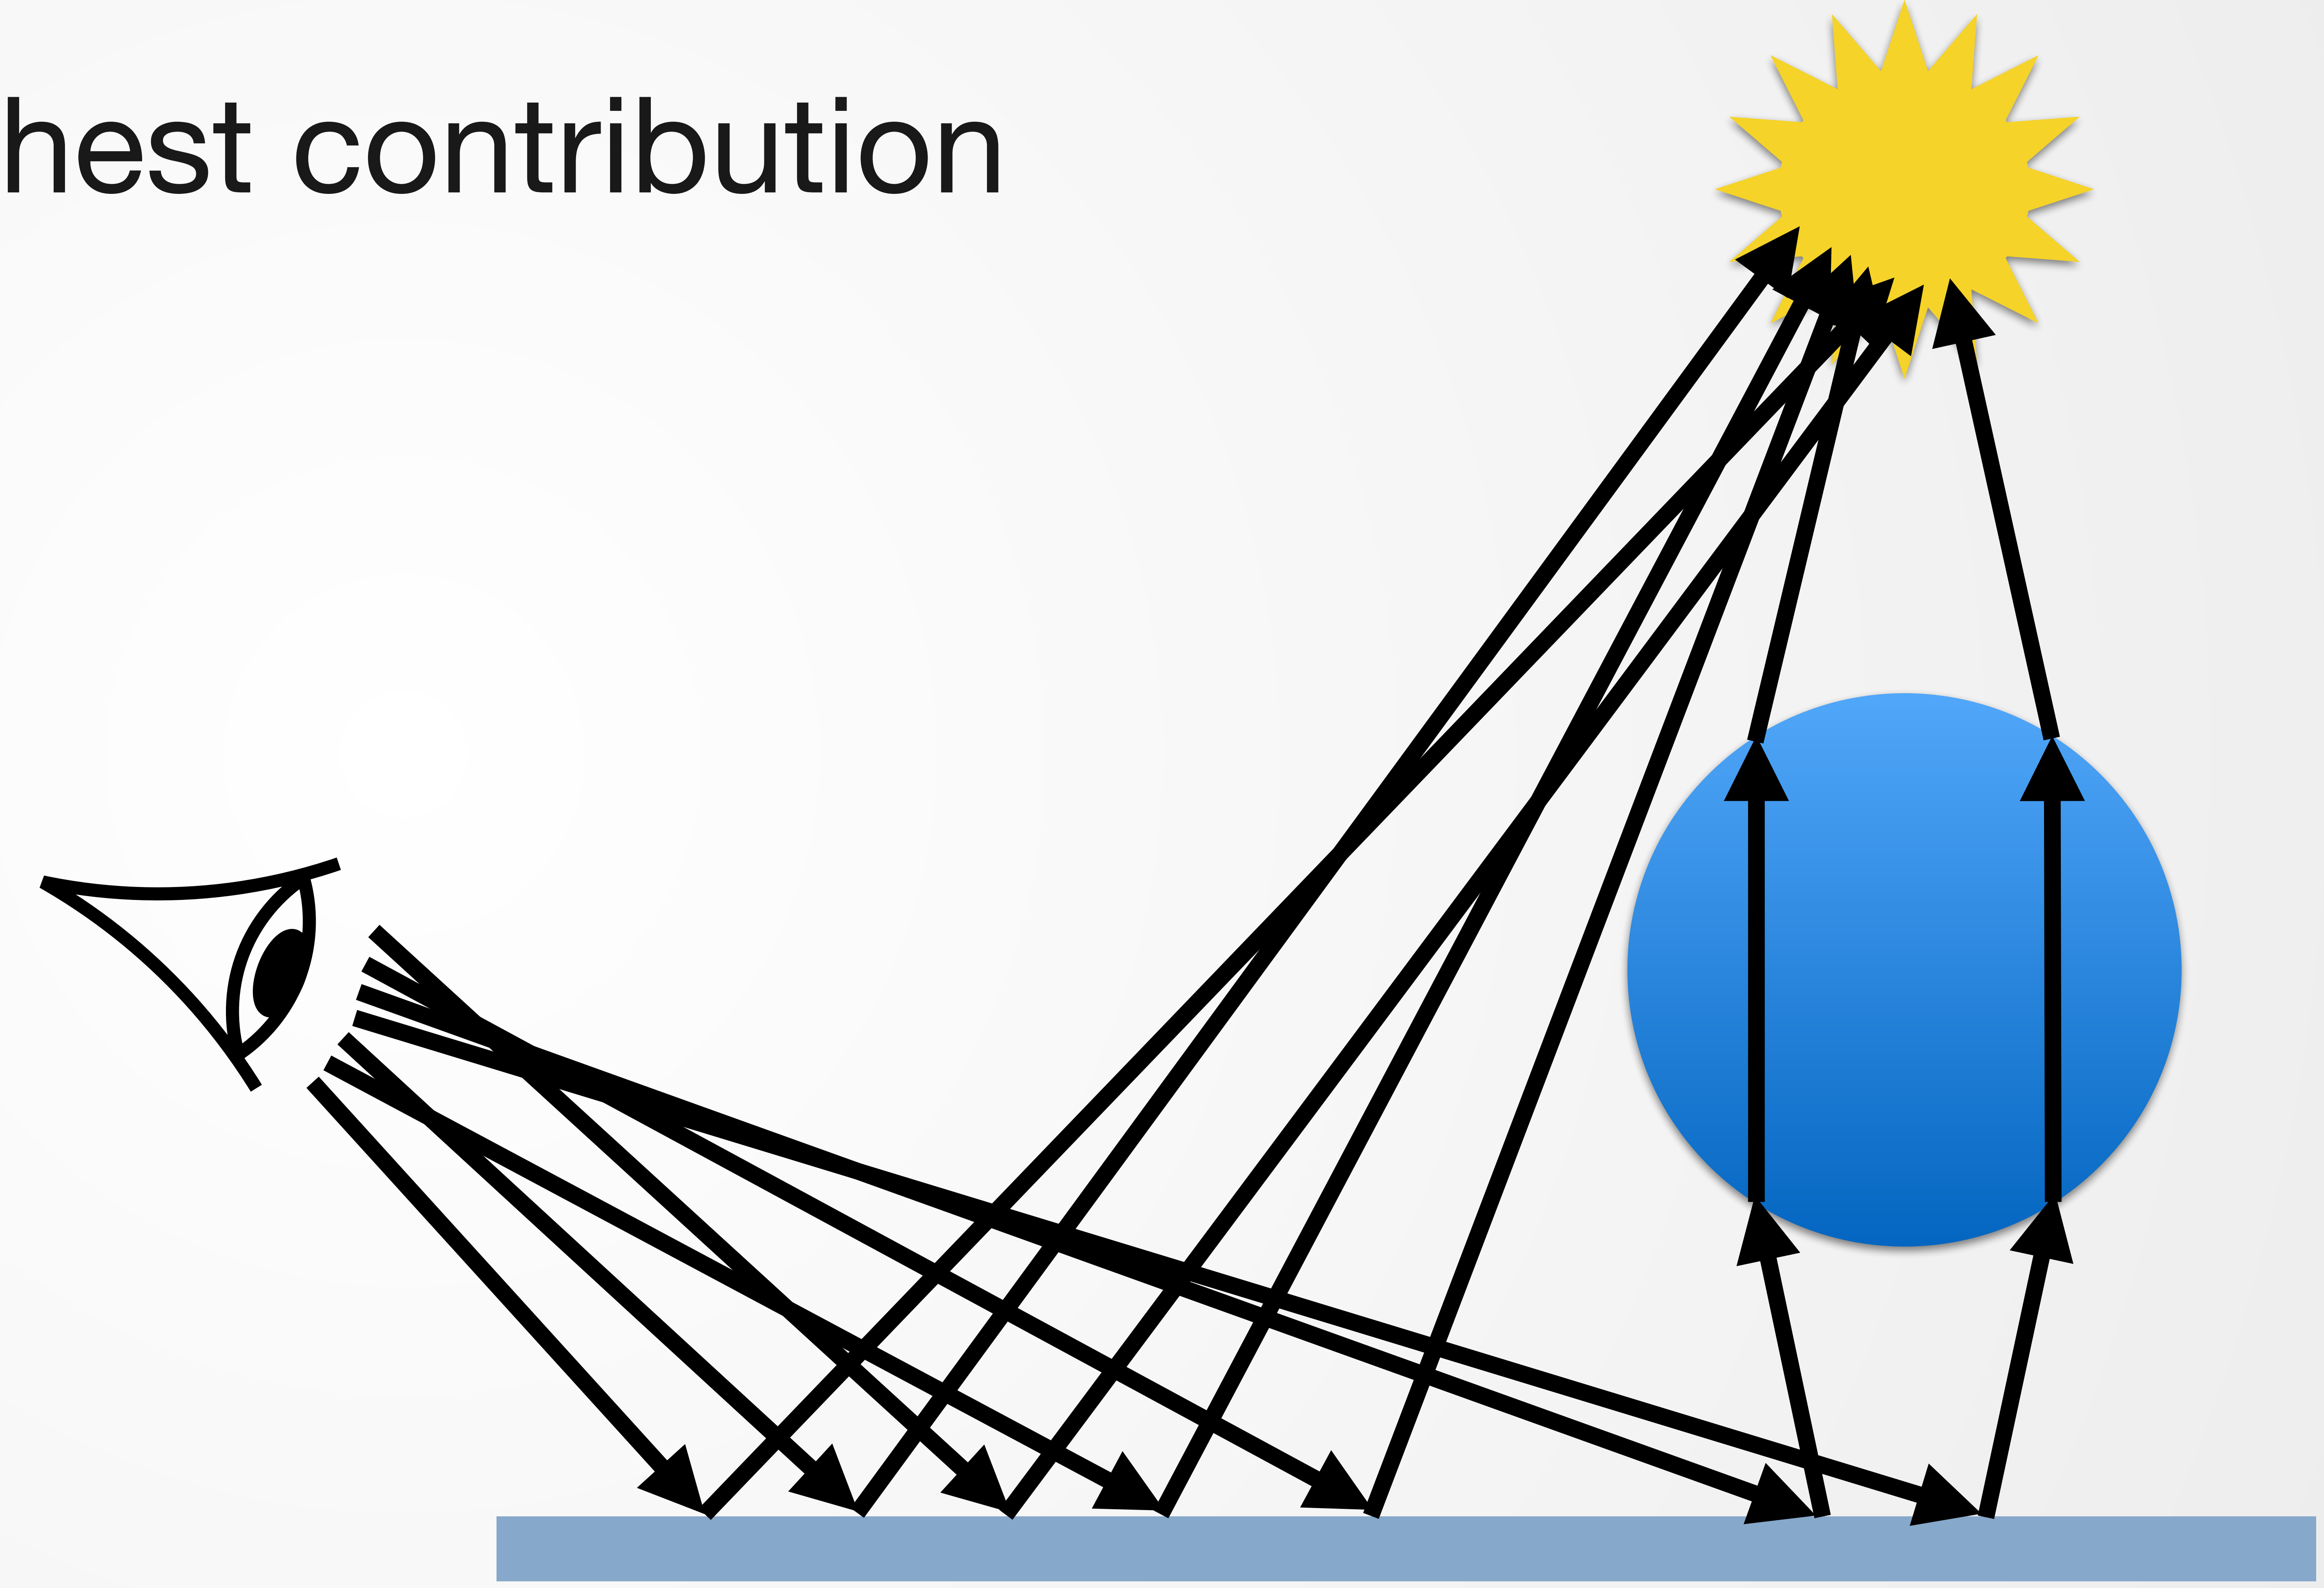

Guiding Overview

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

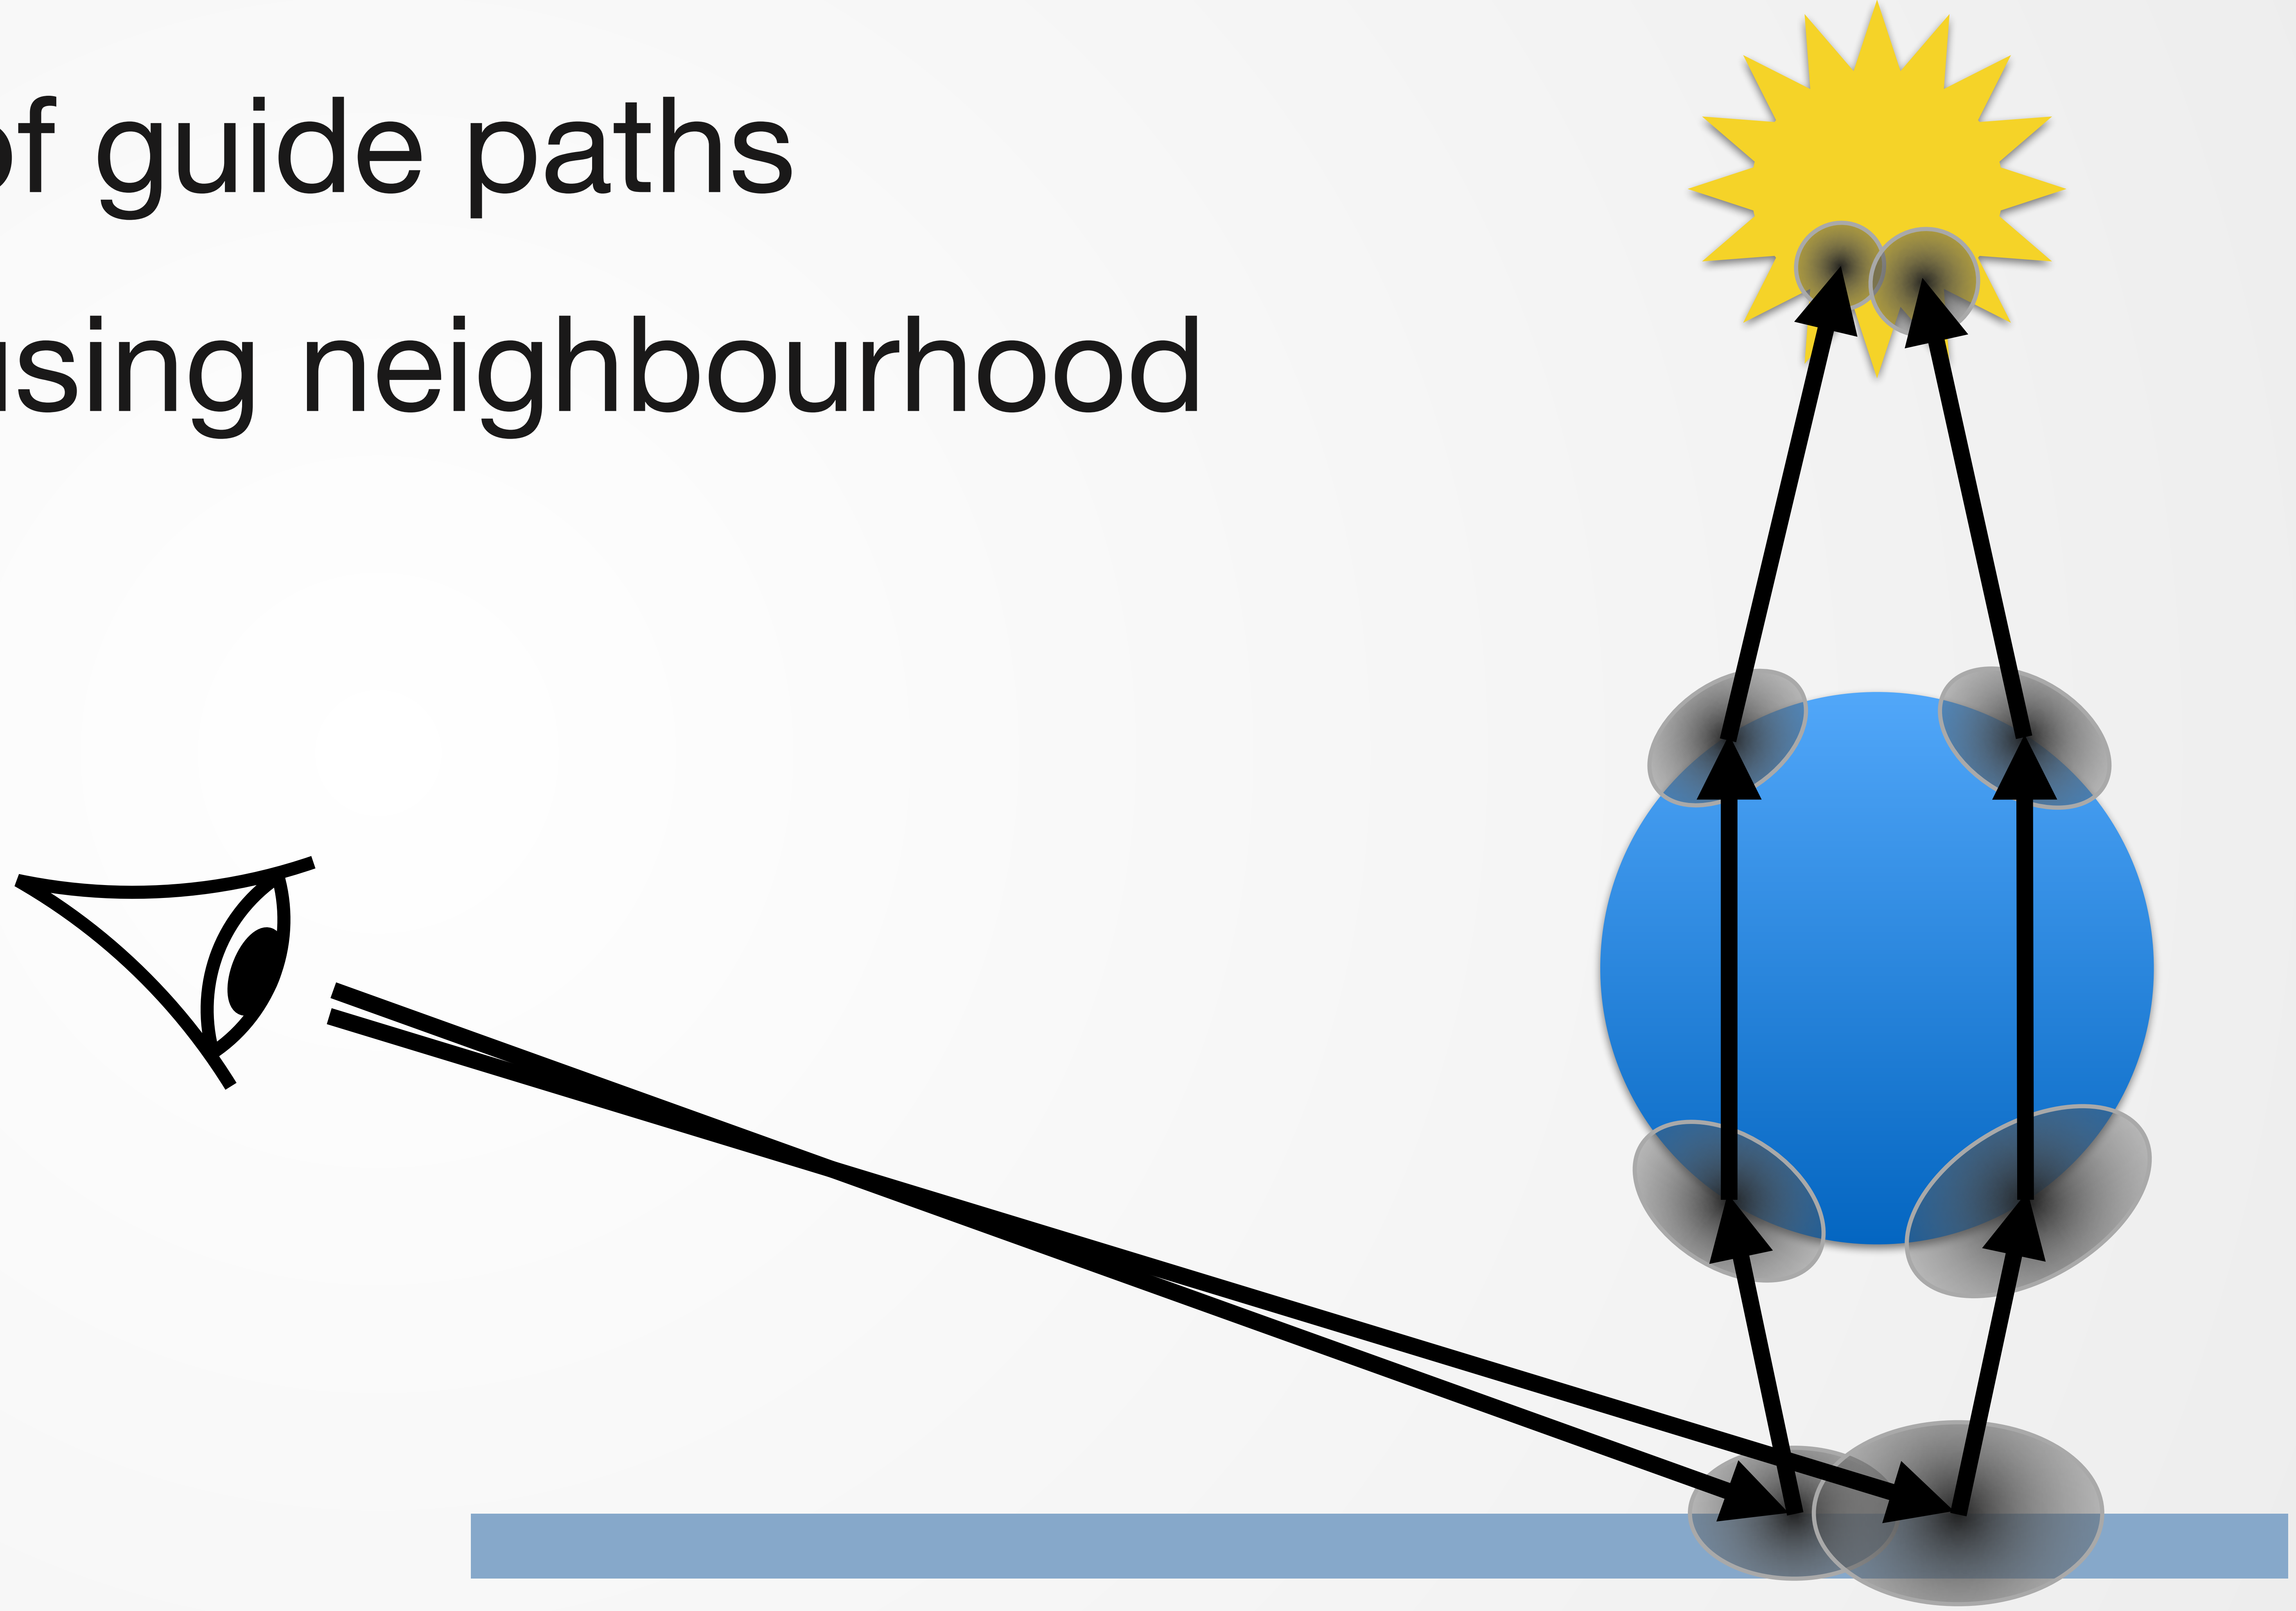

Guiding Overview

- Guided sampling:
  - Choose guide path randomly and
  - Sample Gaussians incrementally

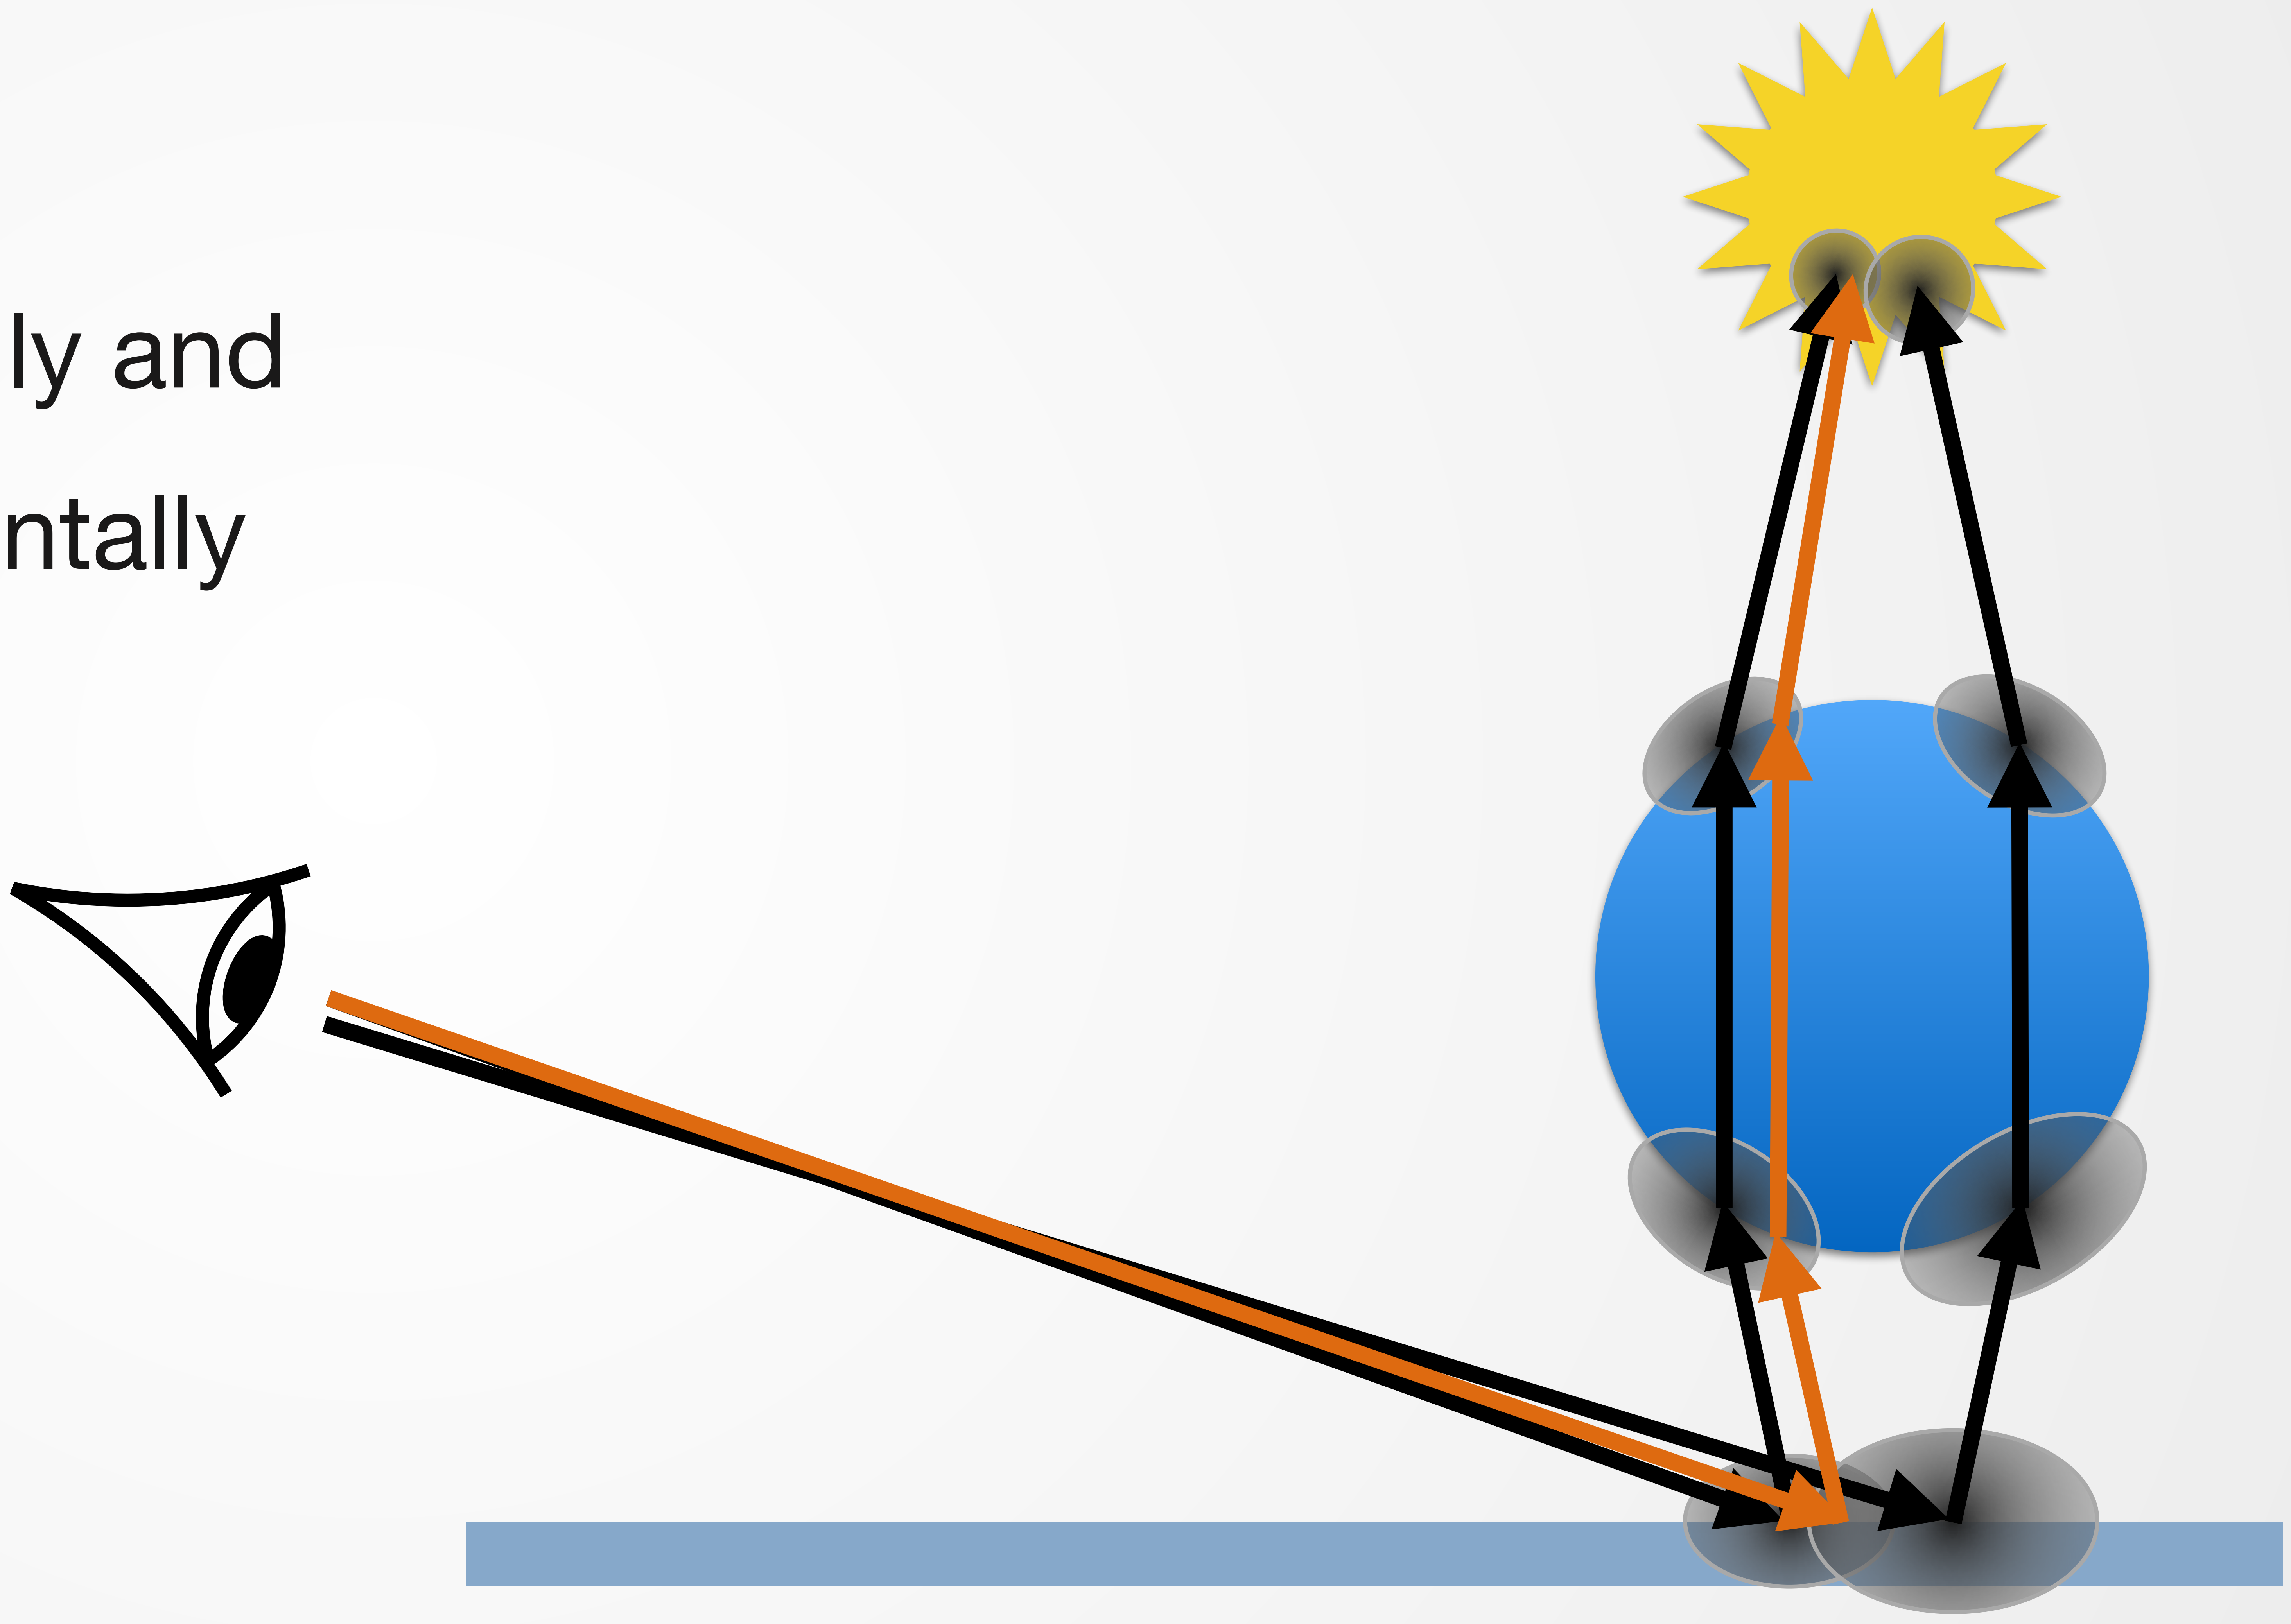

- Guided and unguided sampling combined with multiple importance sampling

Guiding Behaviour

Reference

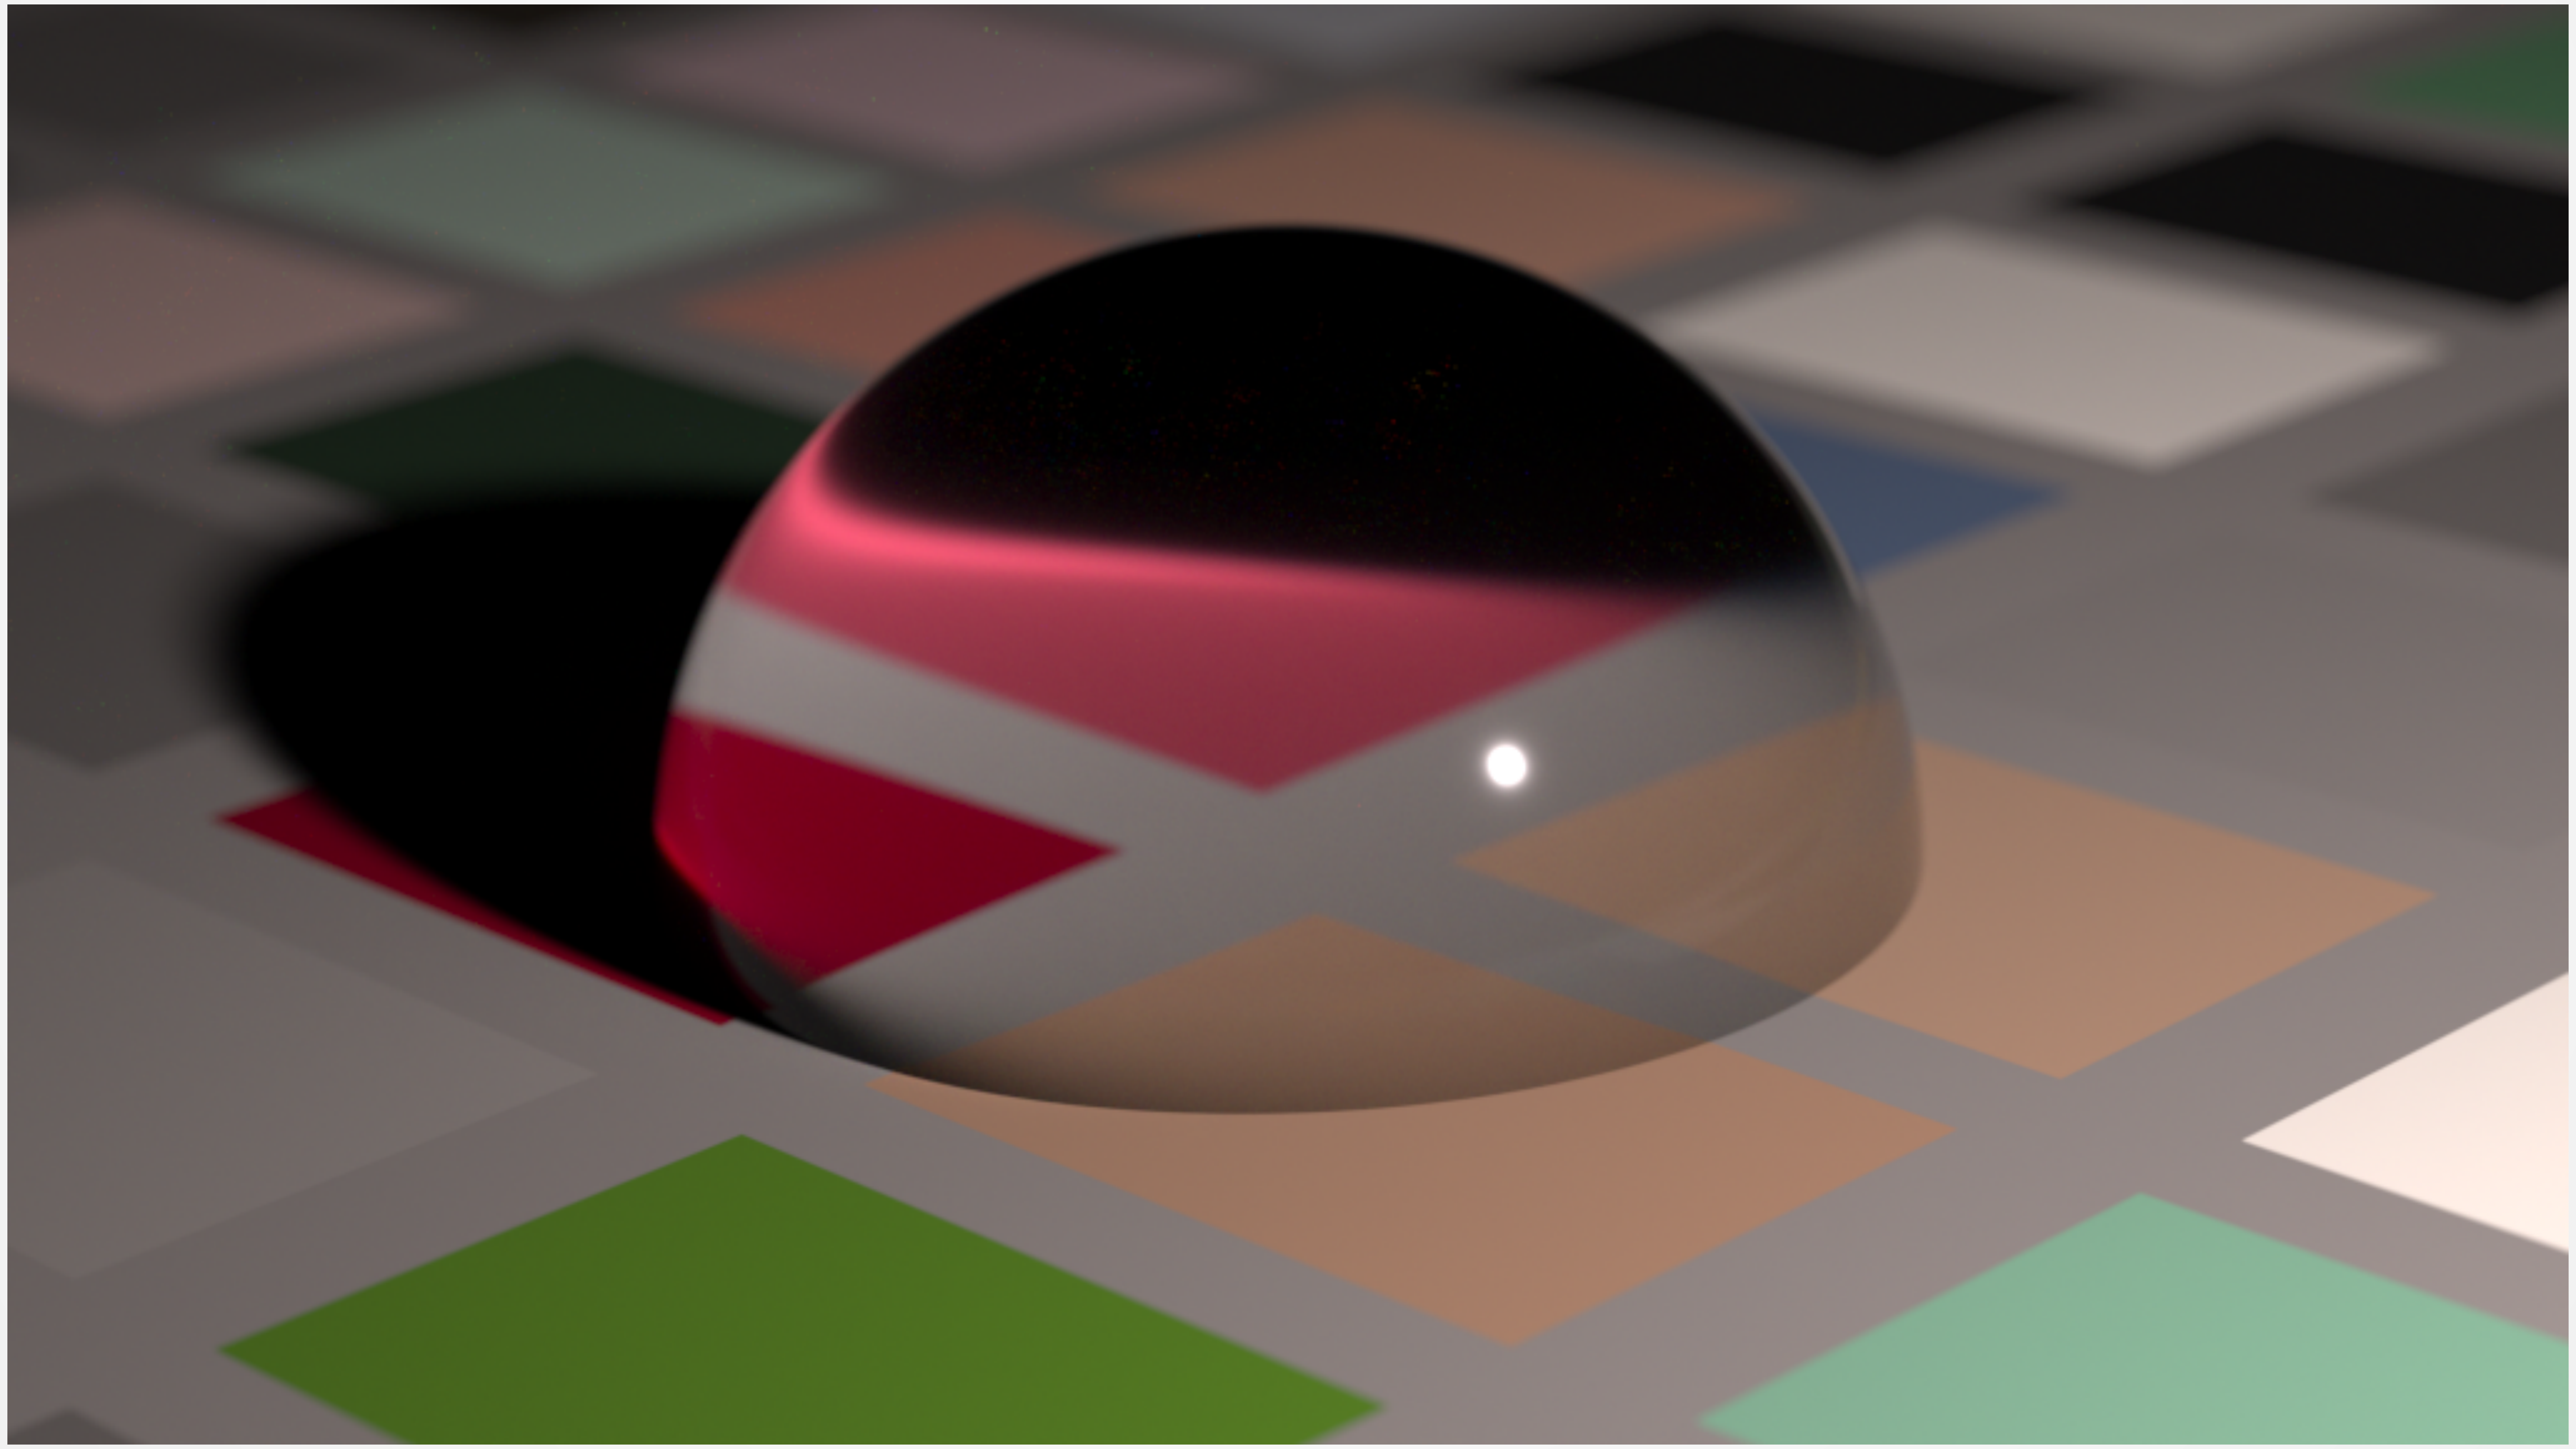

Guiding Behaviour

Path tracing

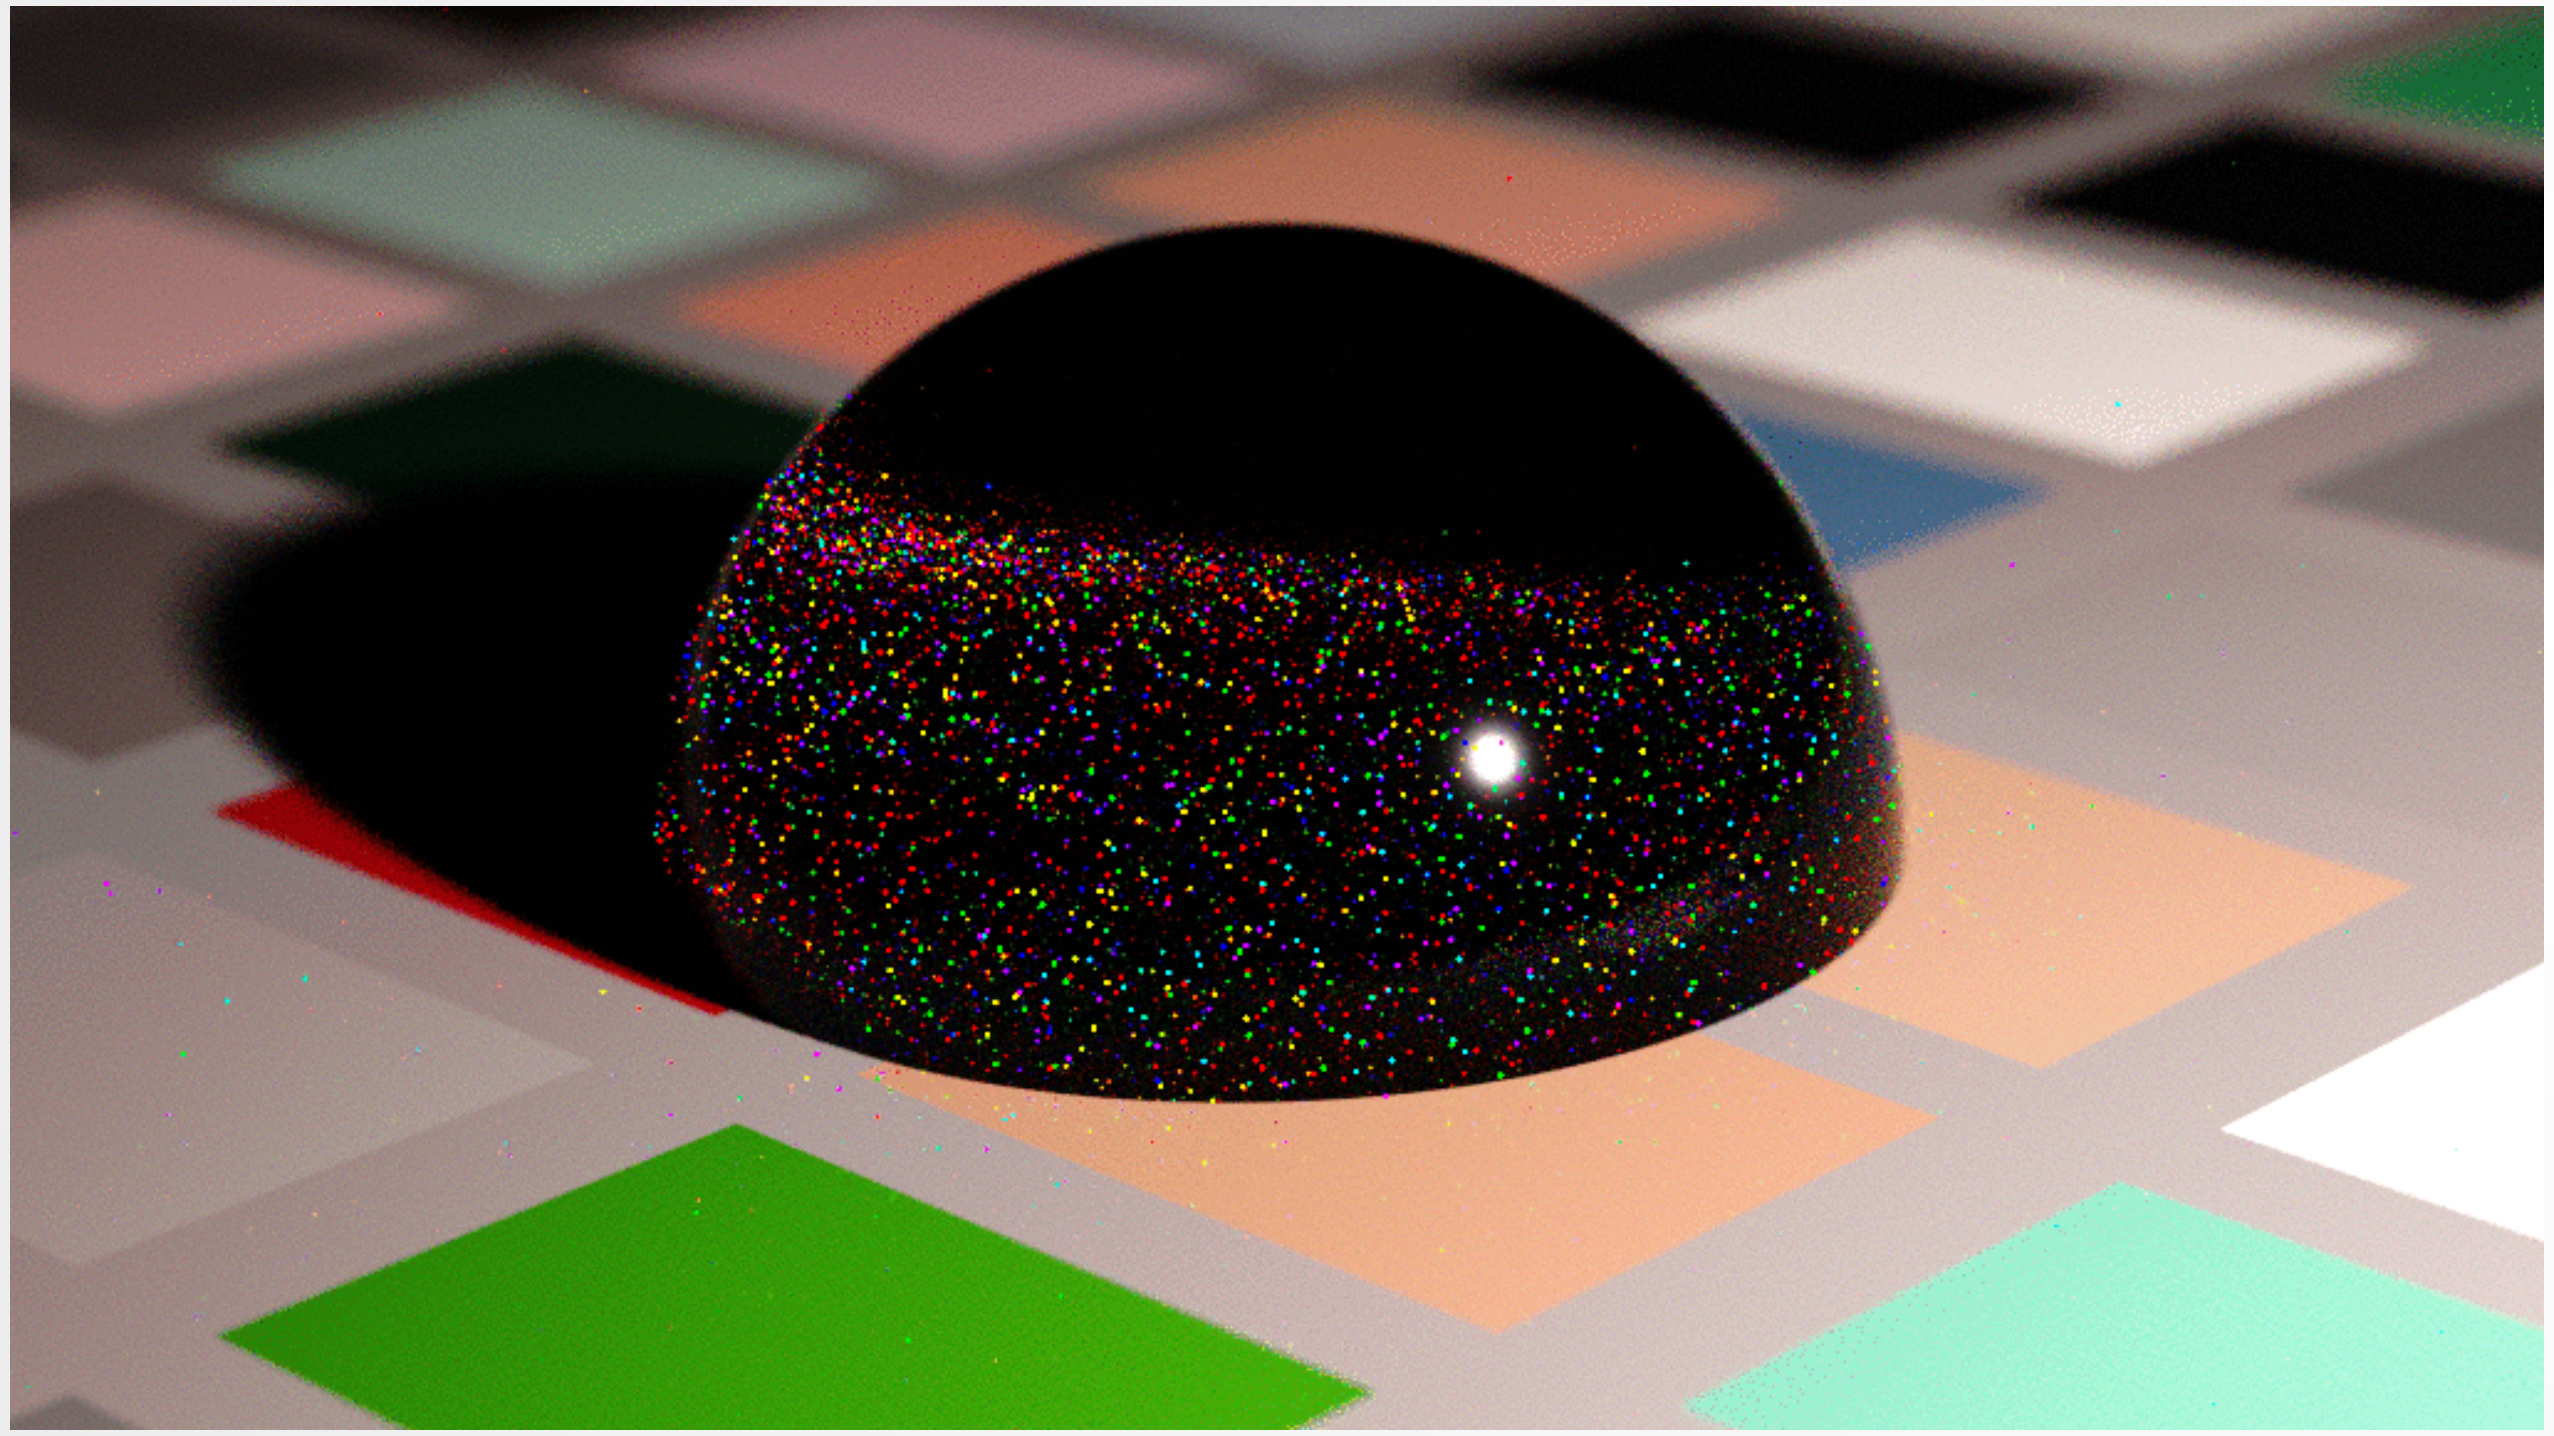

Guided path tracing

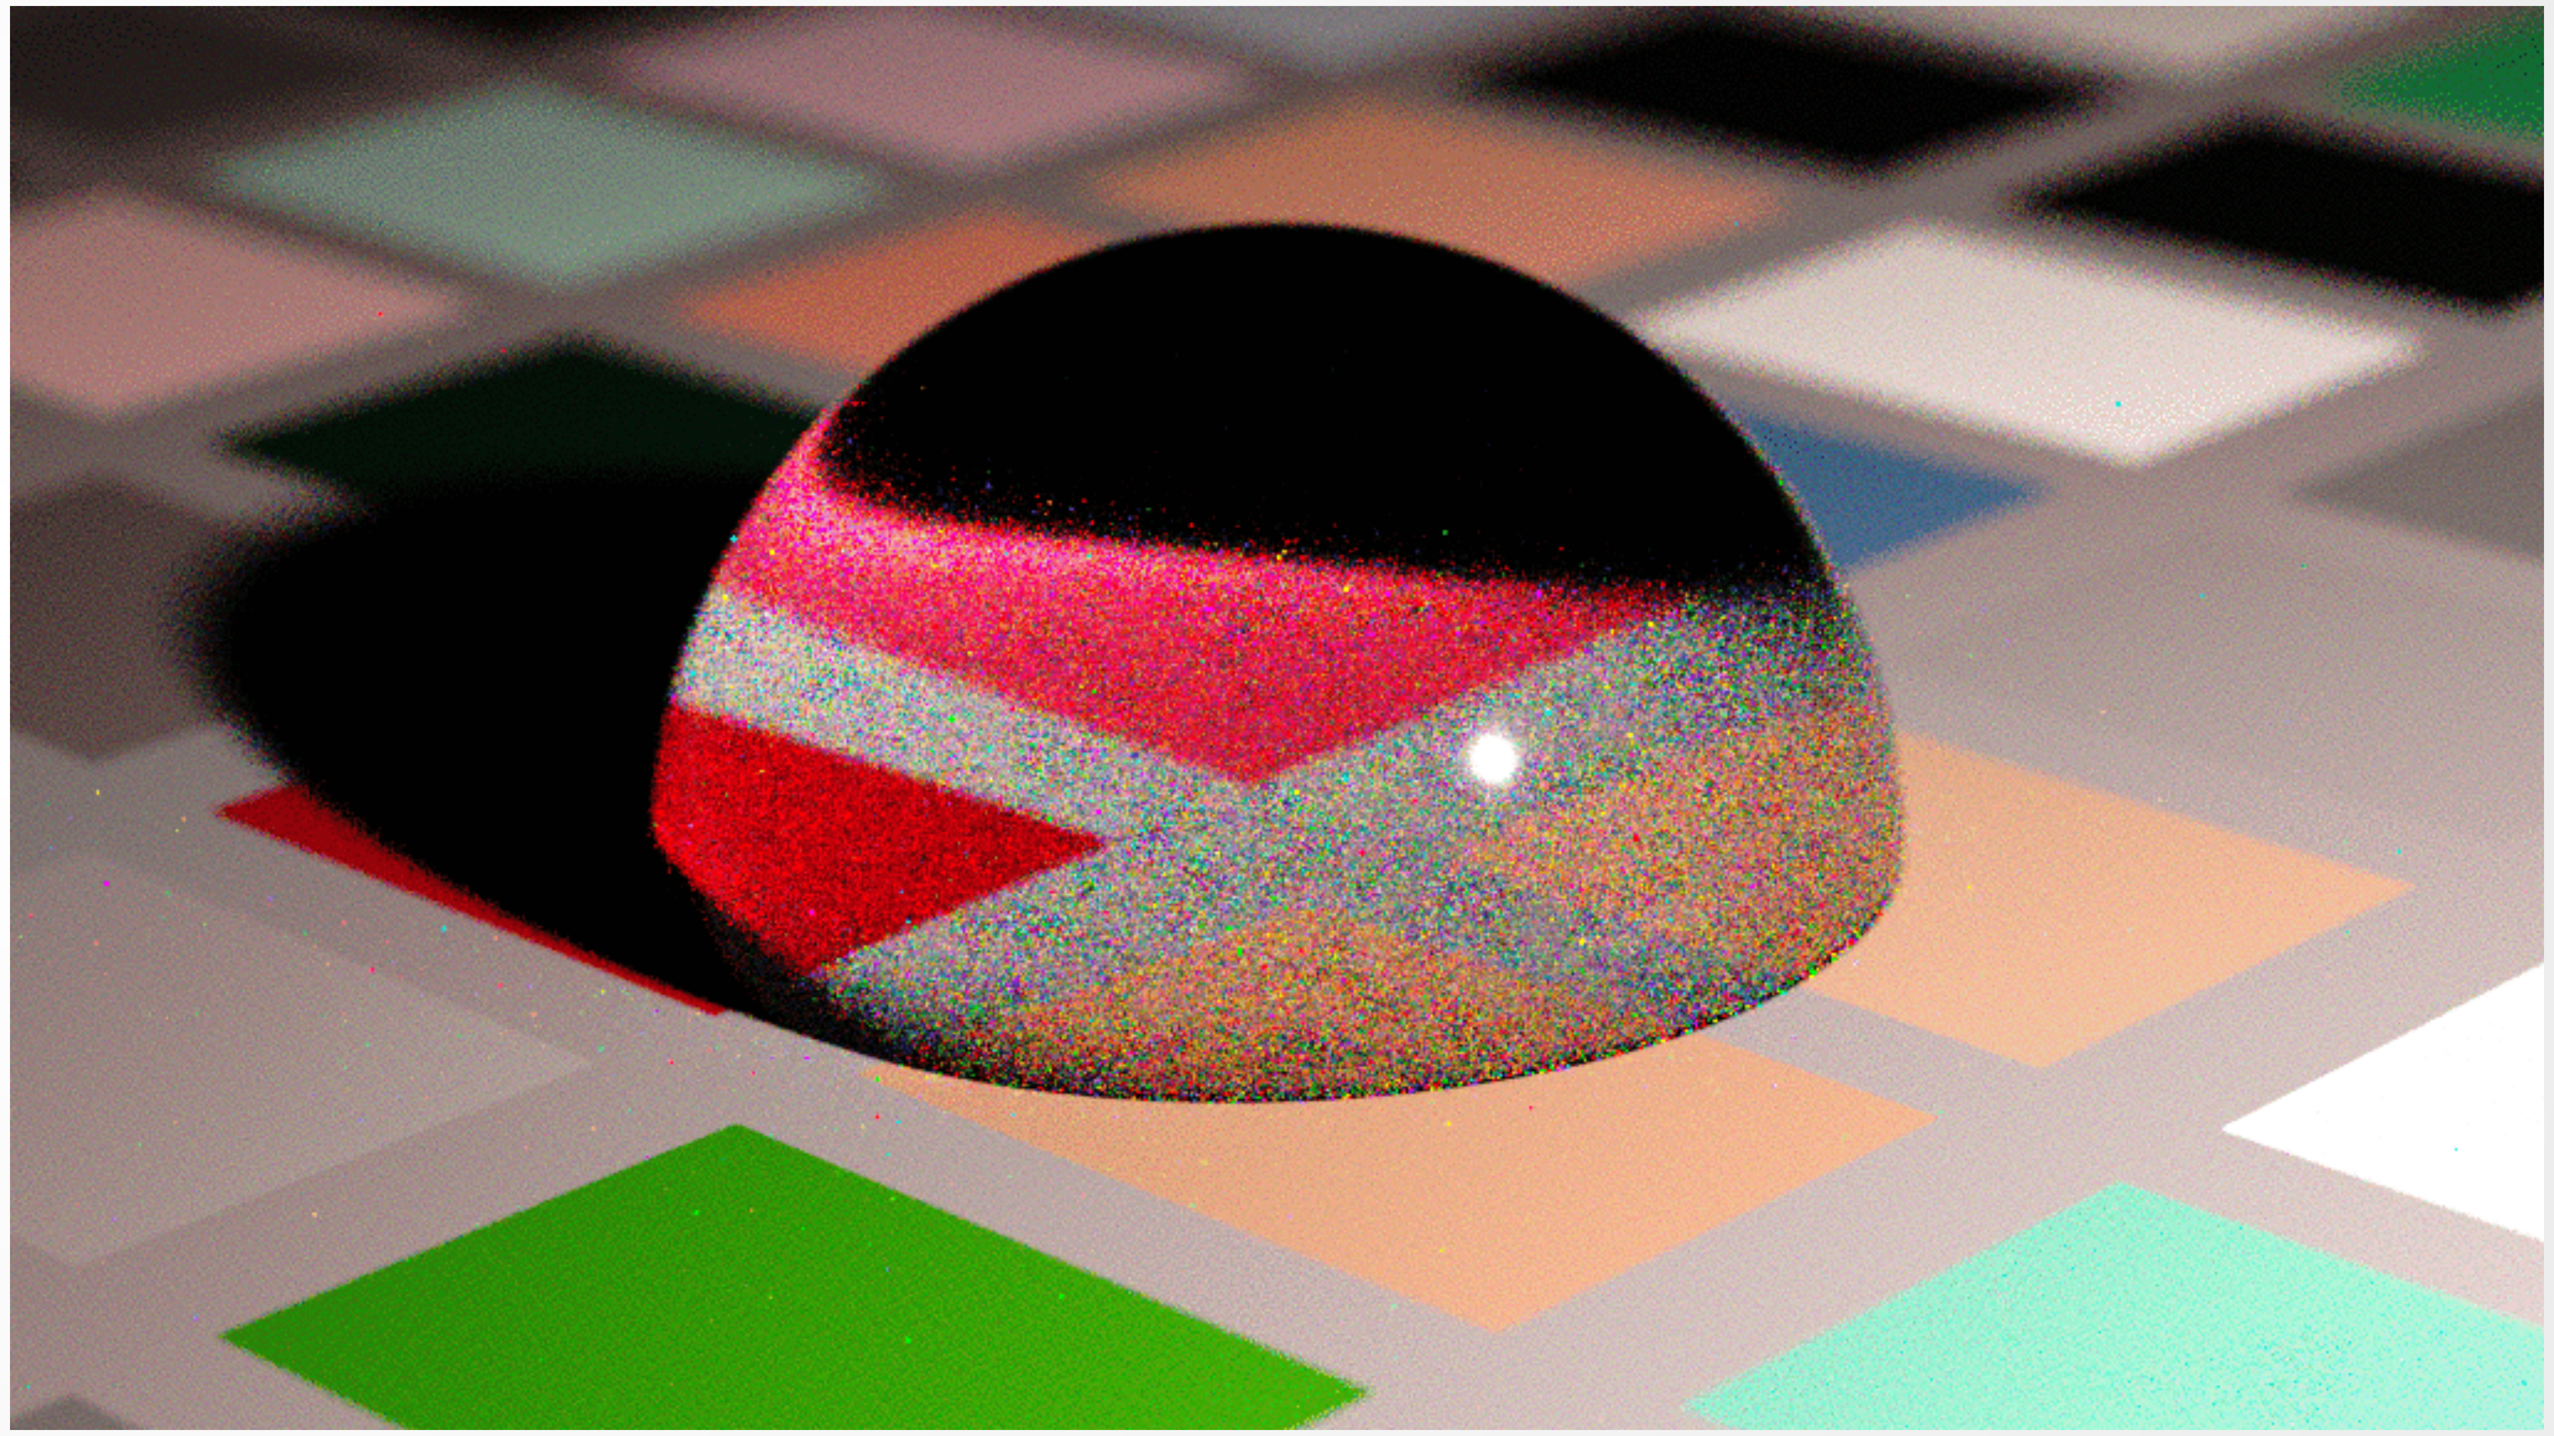

Guiding Behaviour

Reference

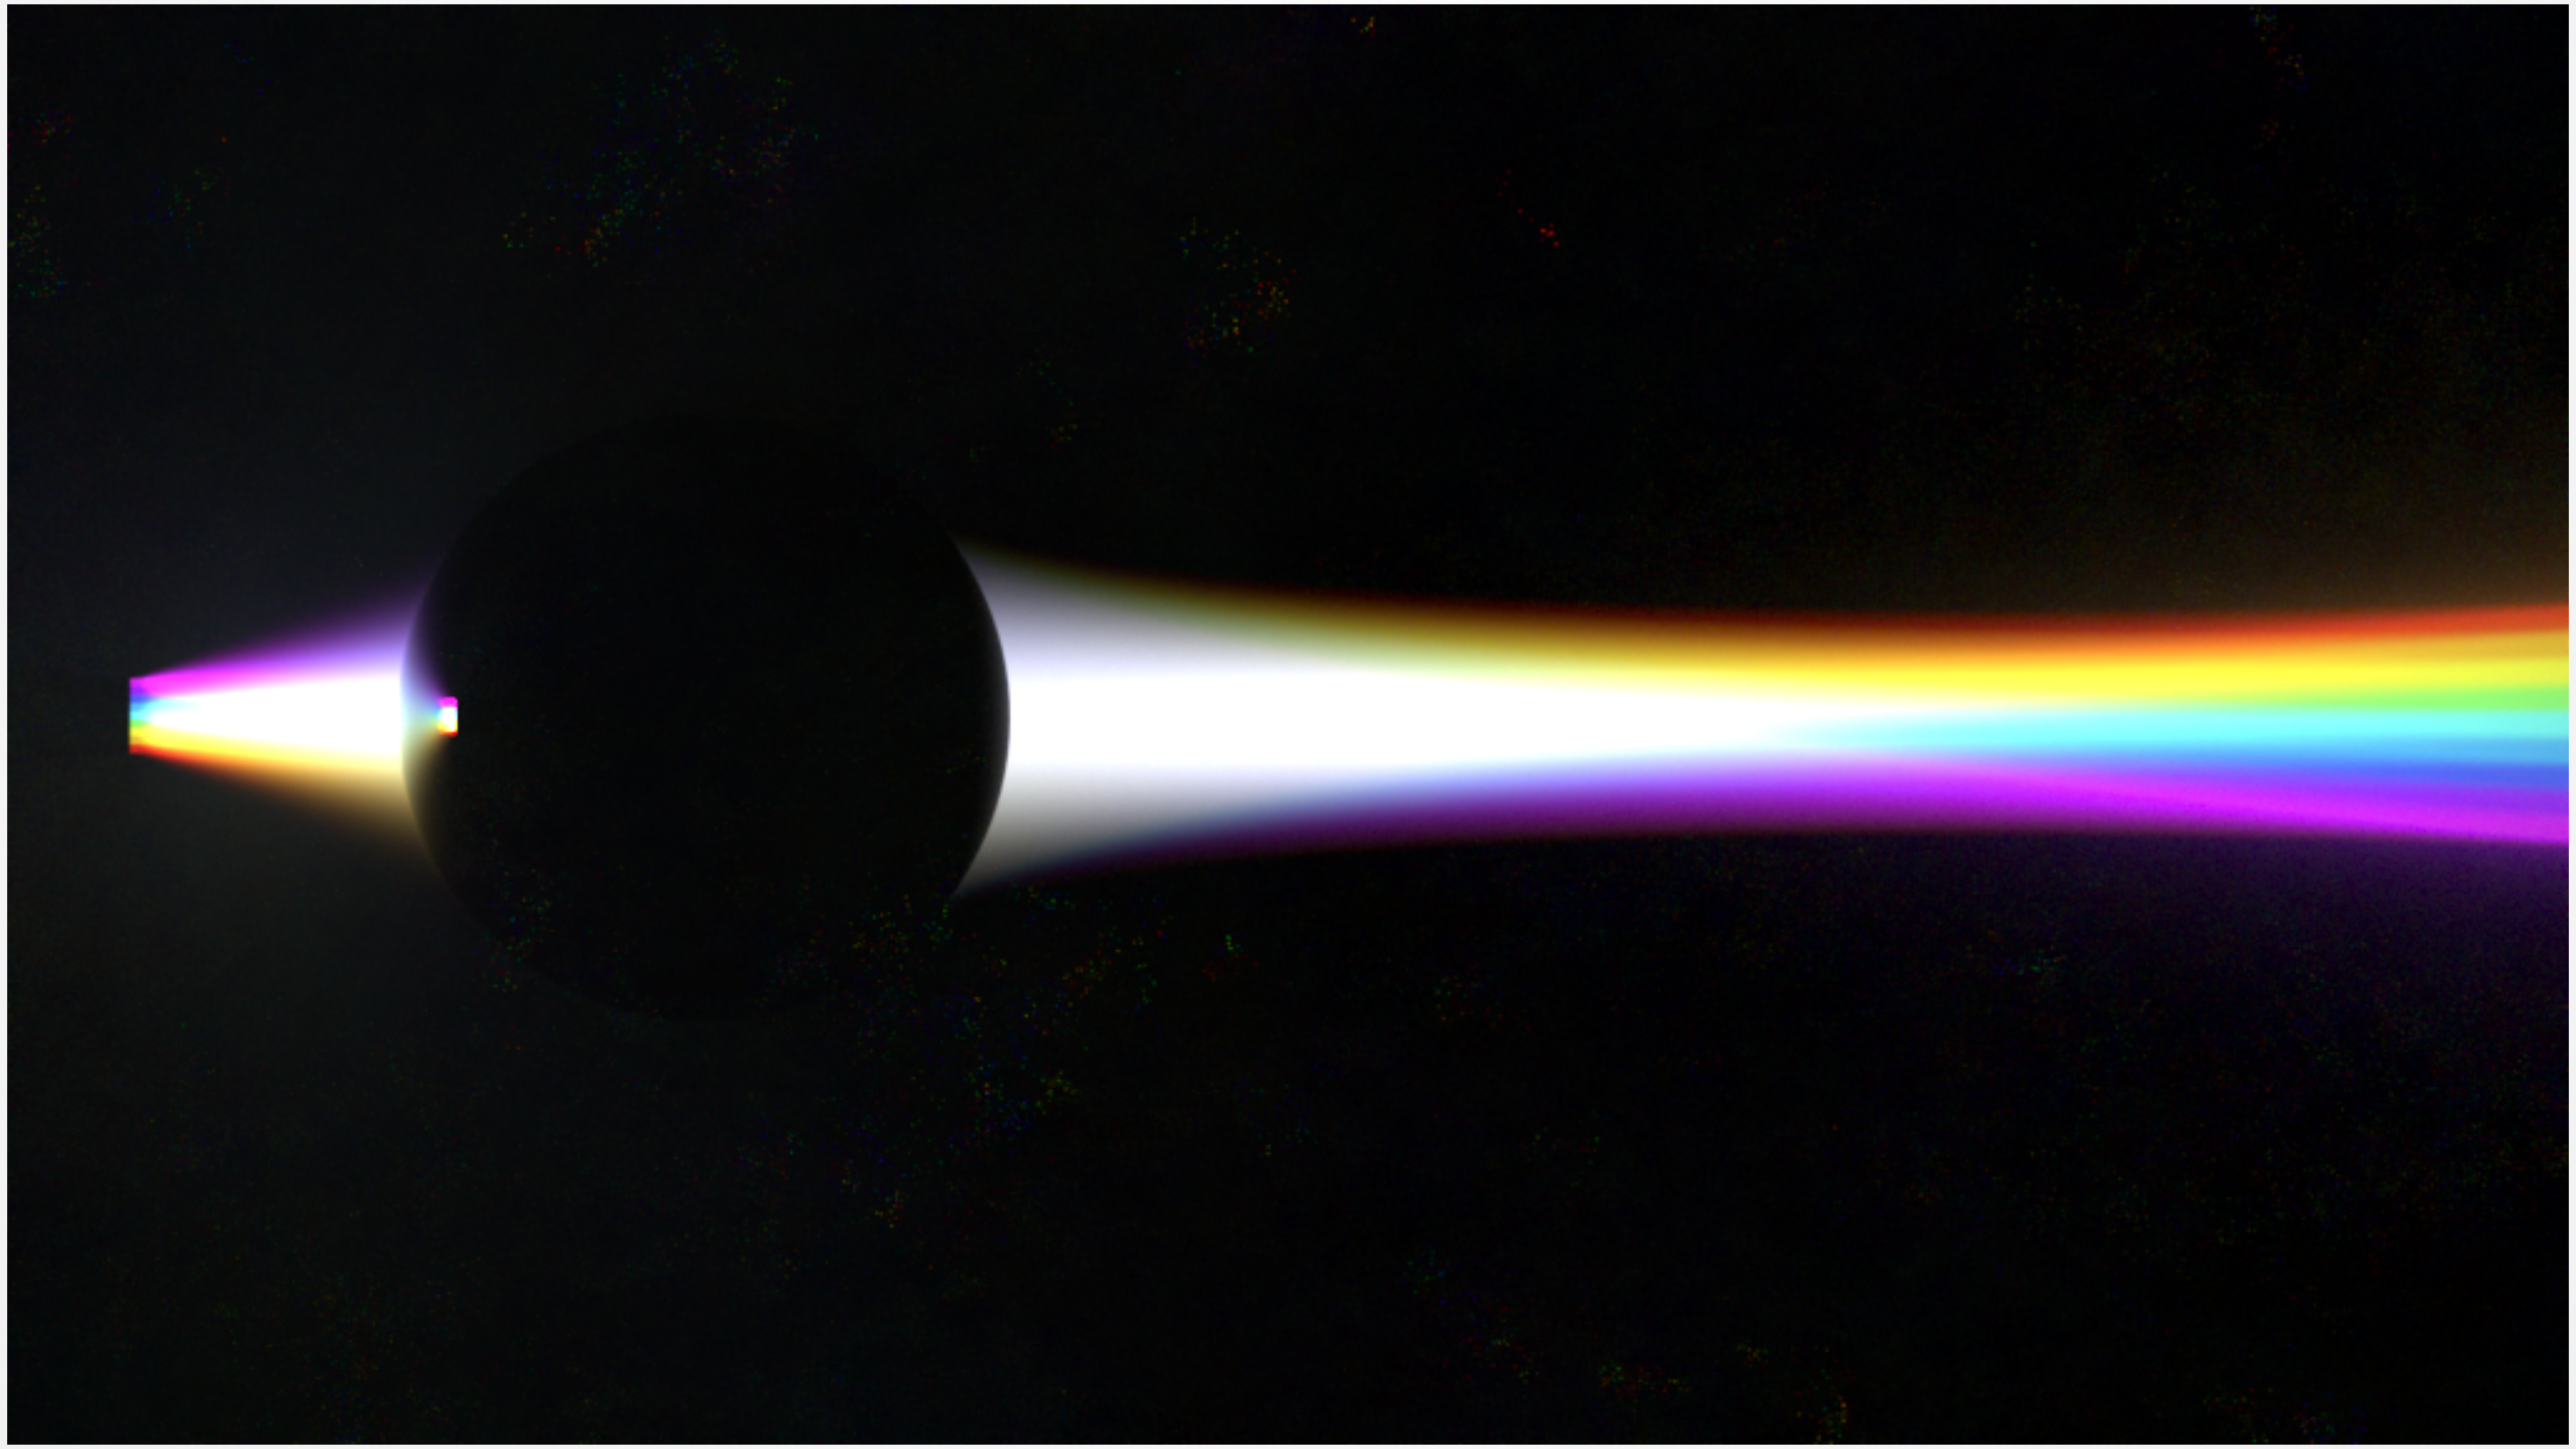

Guiding Behaviour

Path tracing

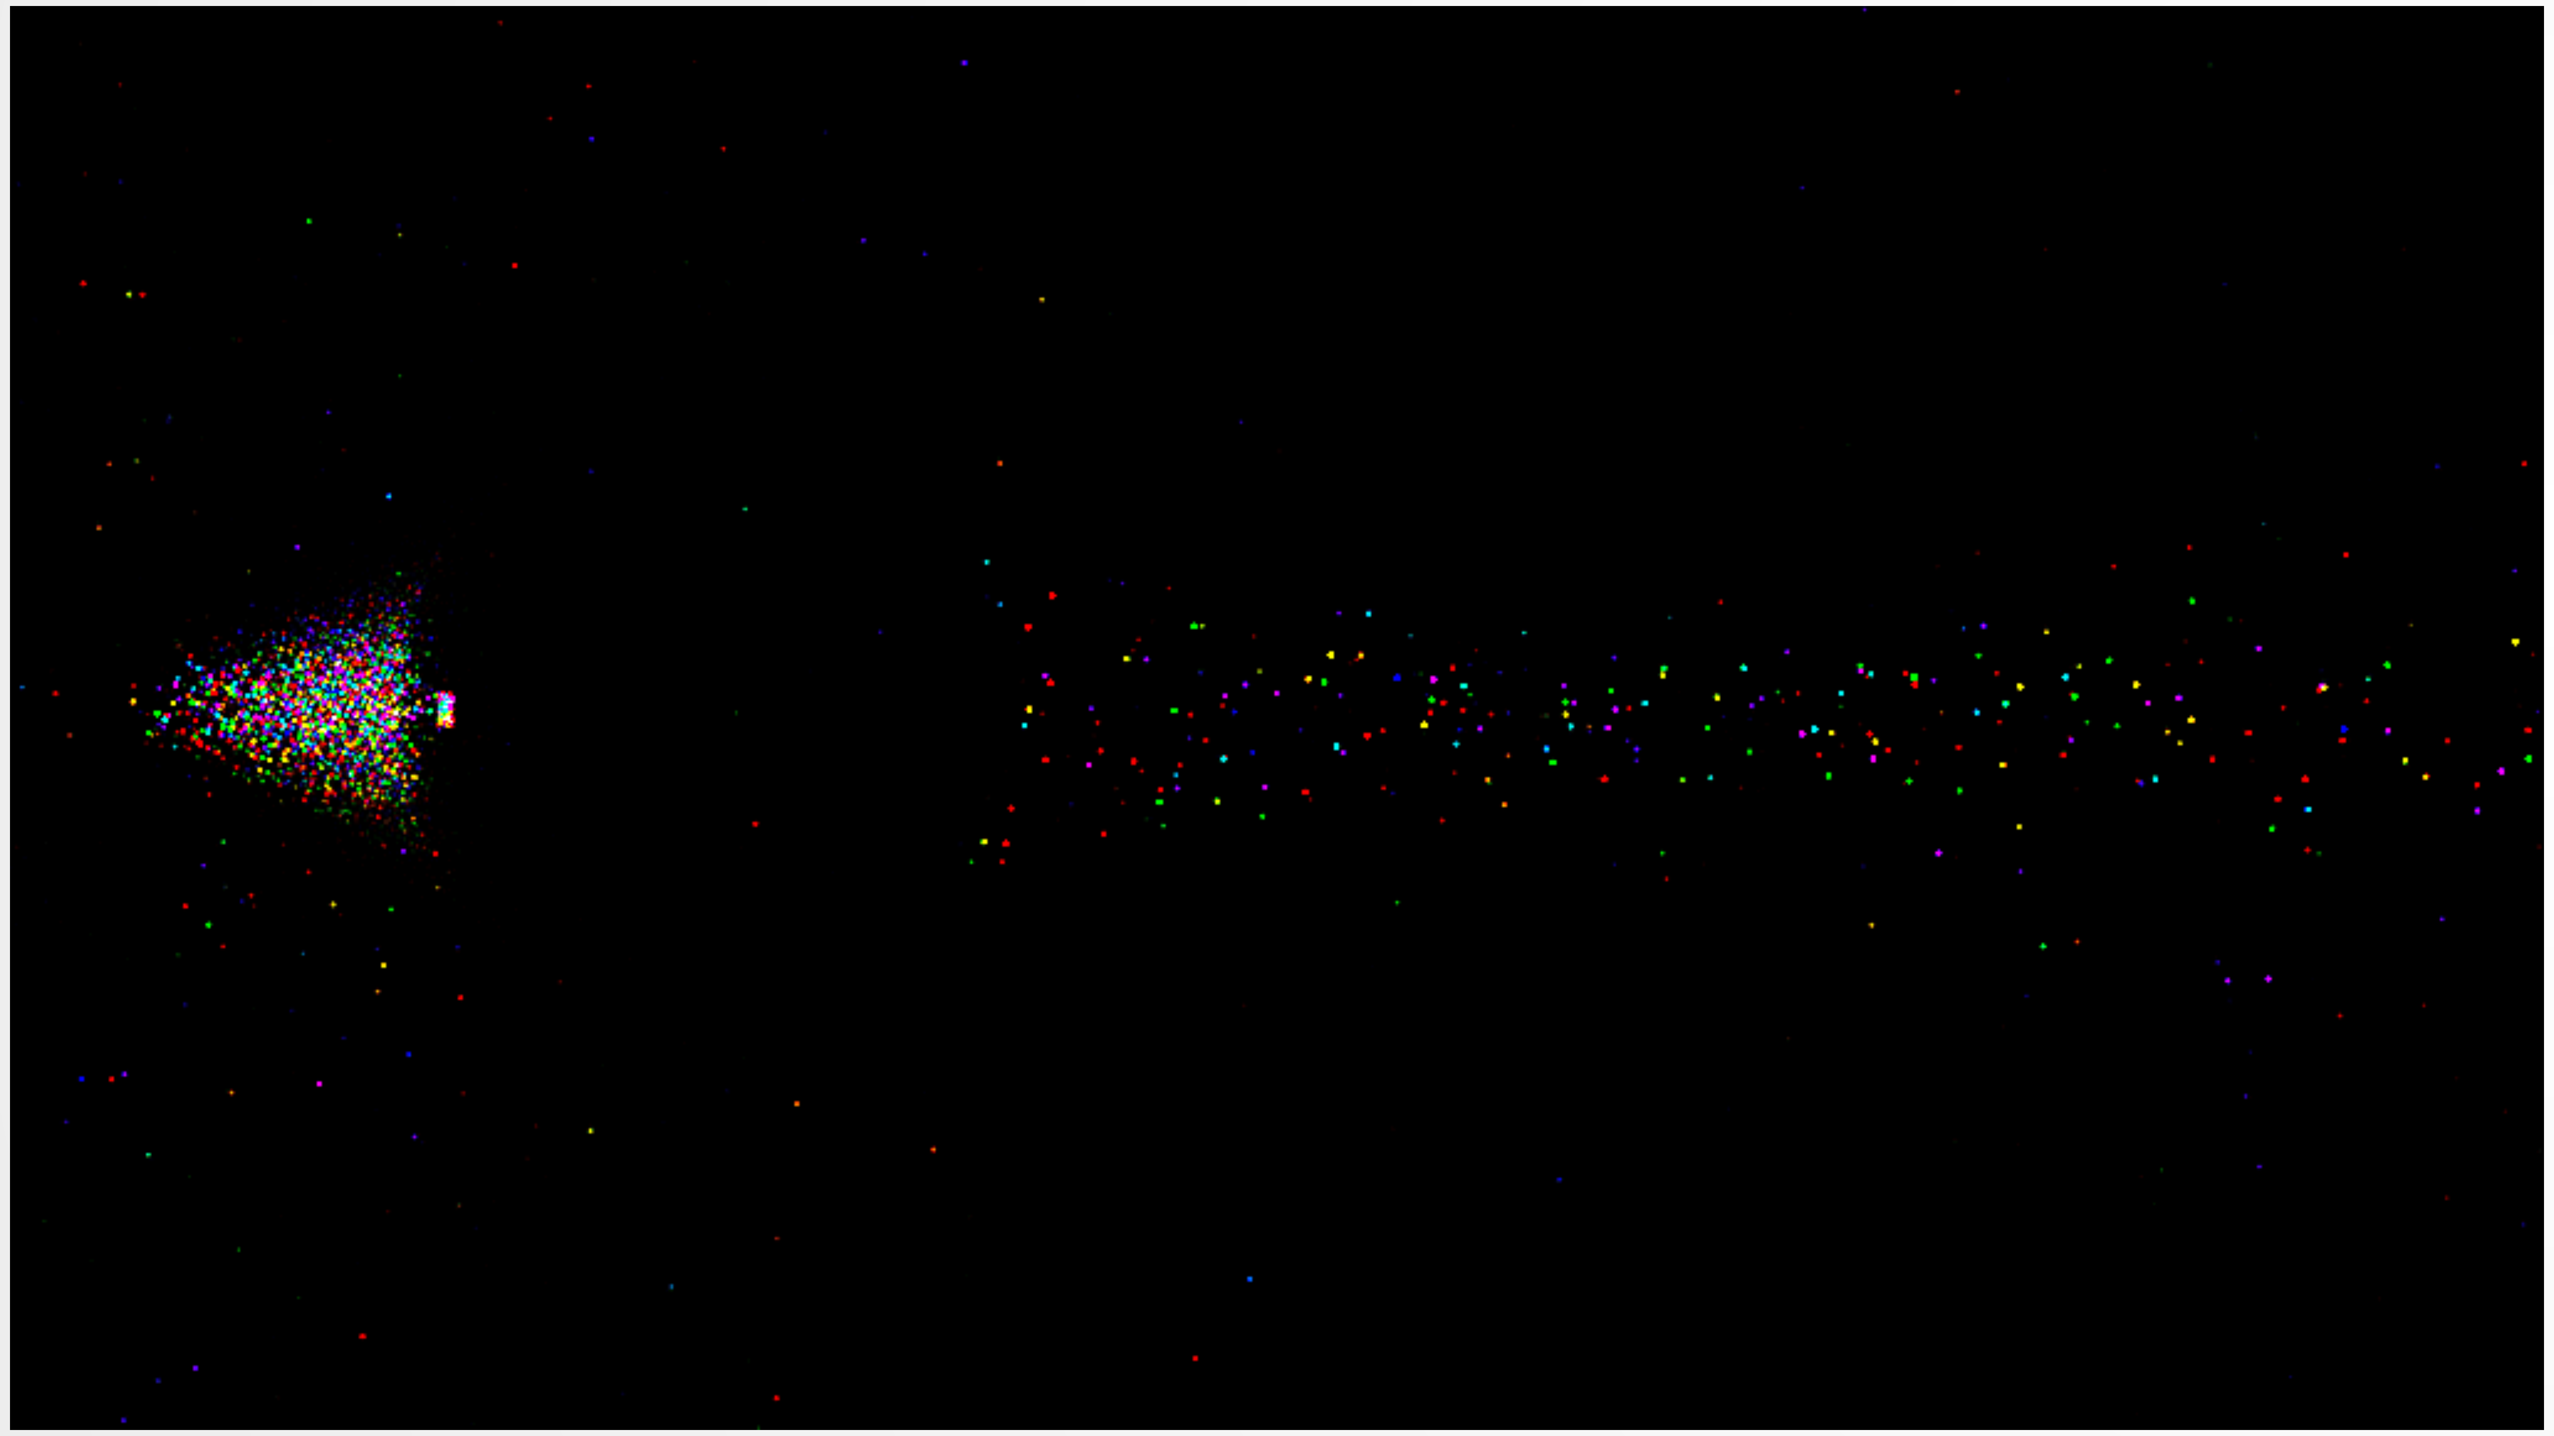

Guided path tracing

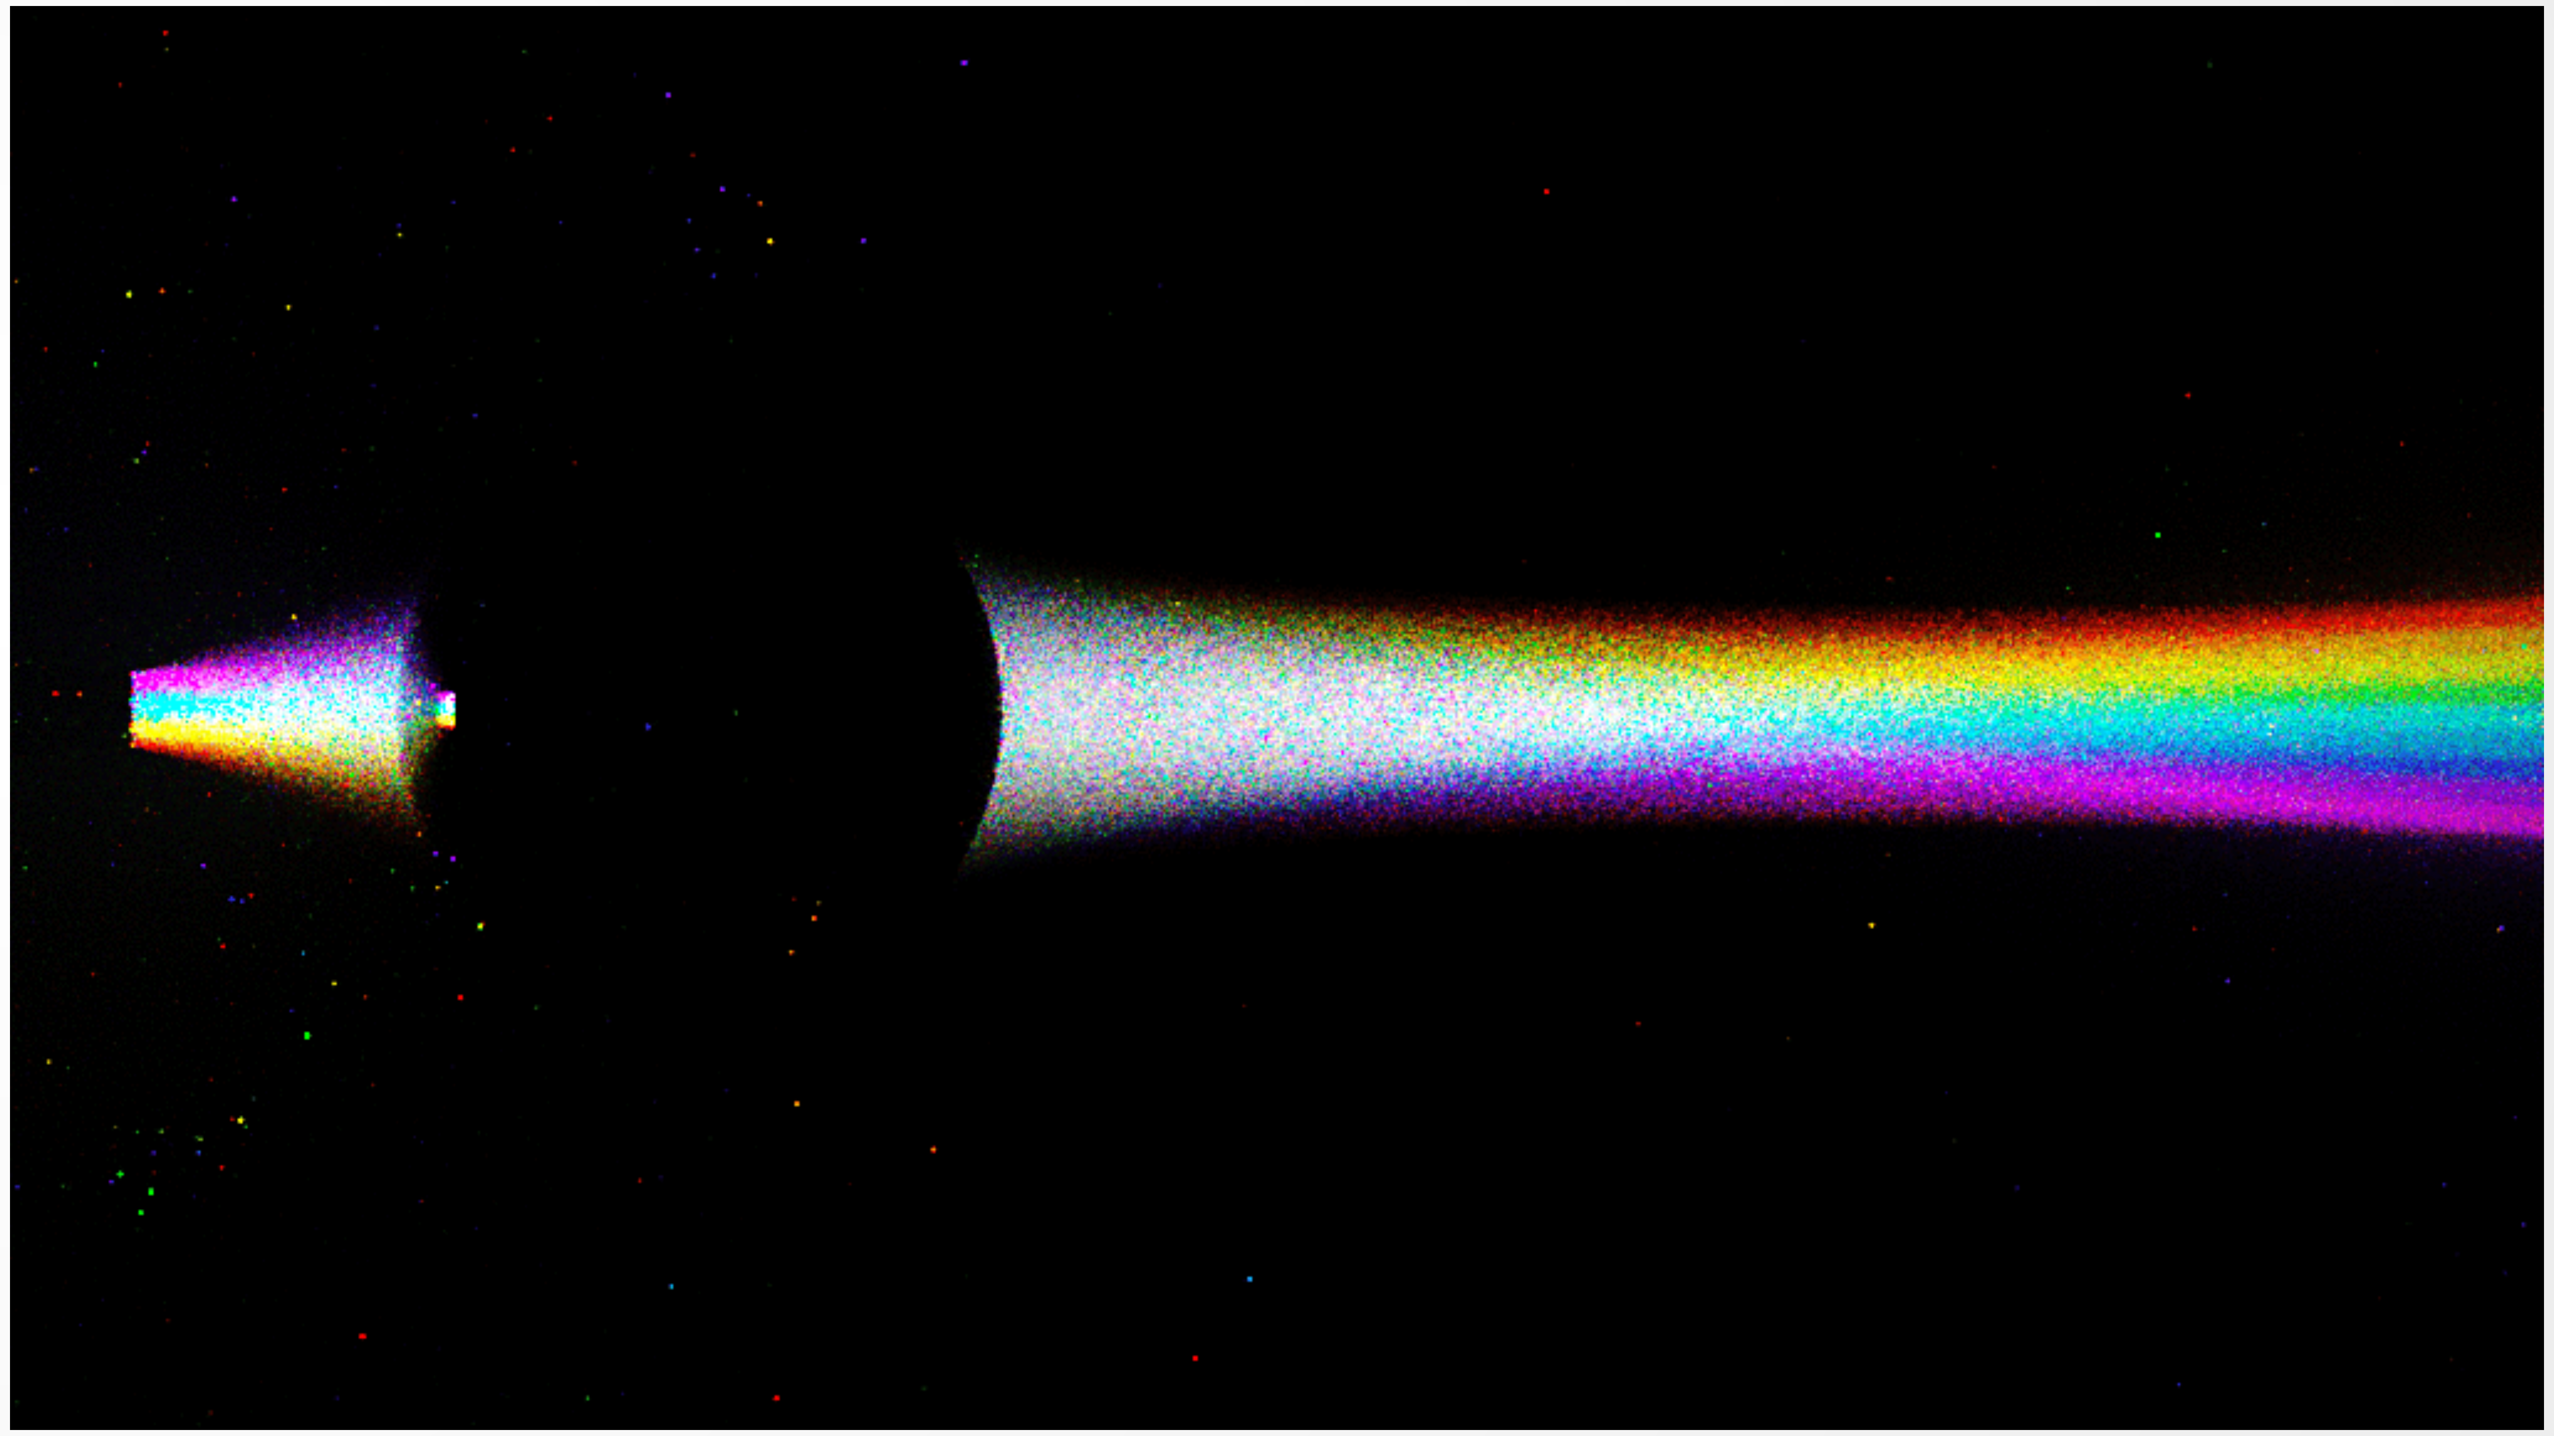

Path Correlation

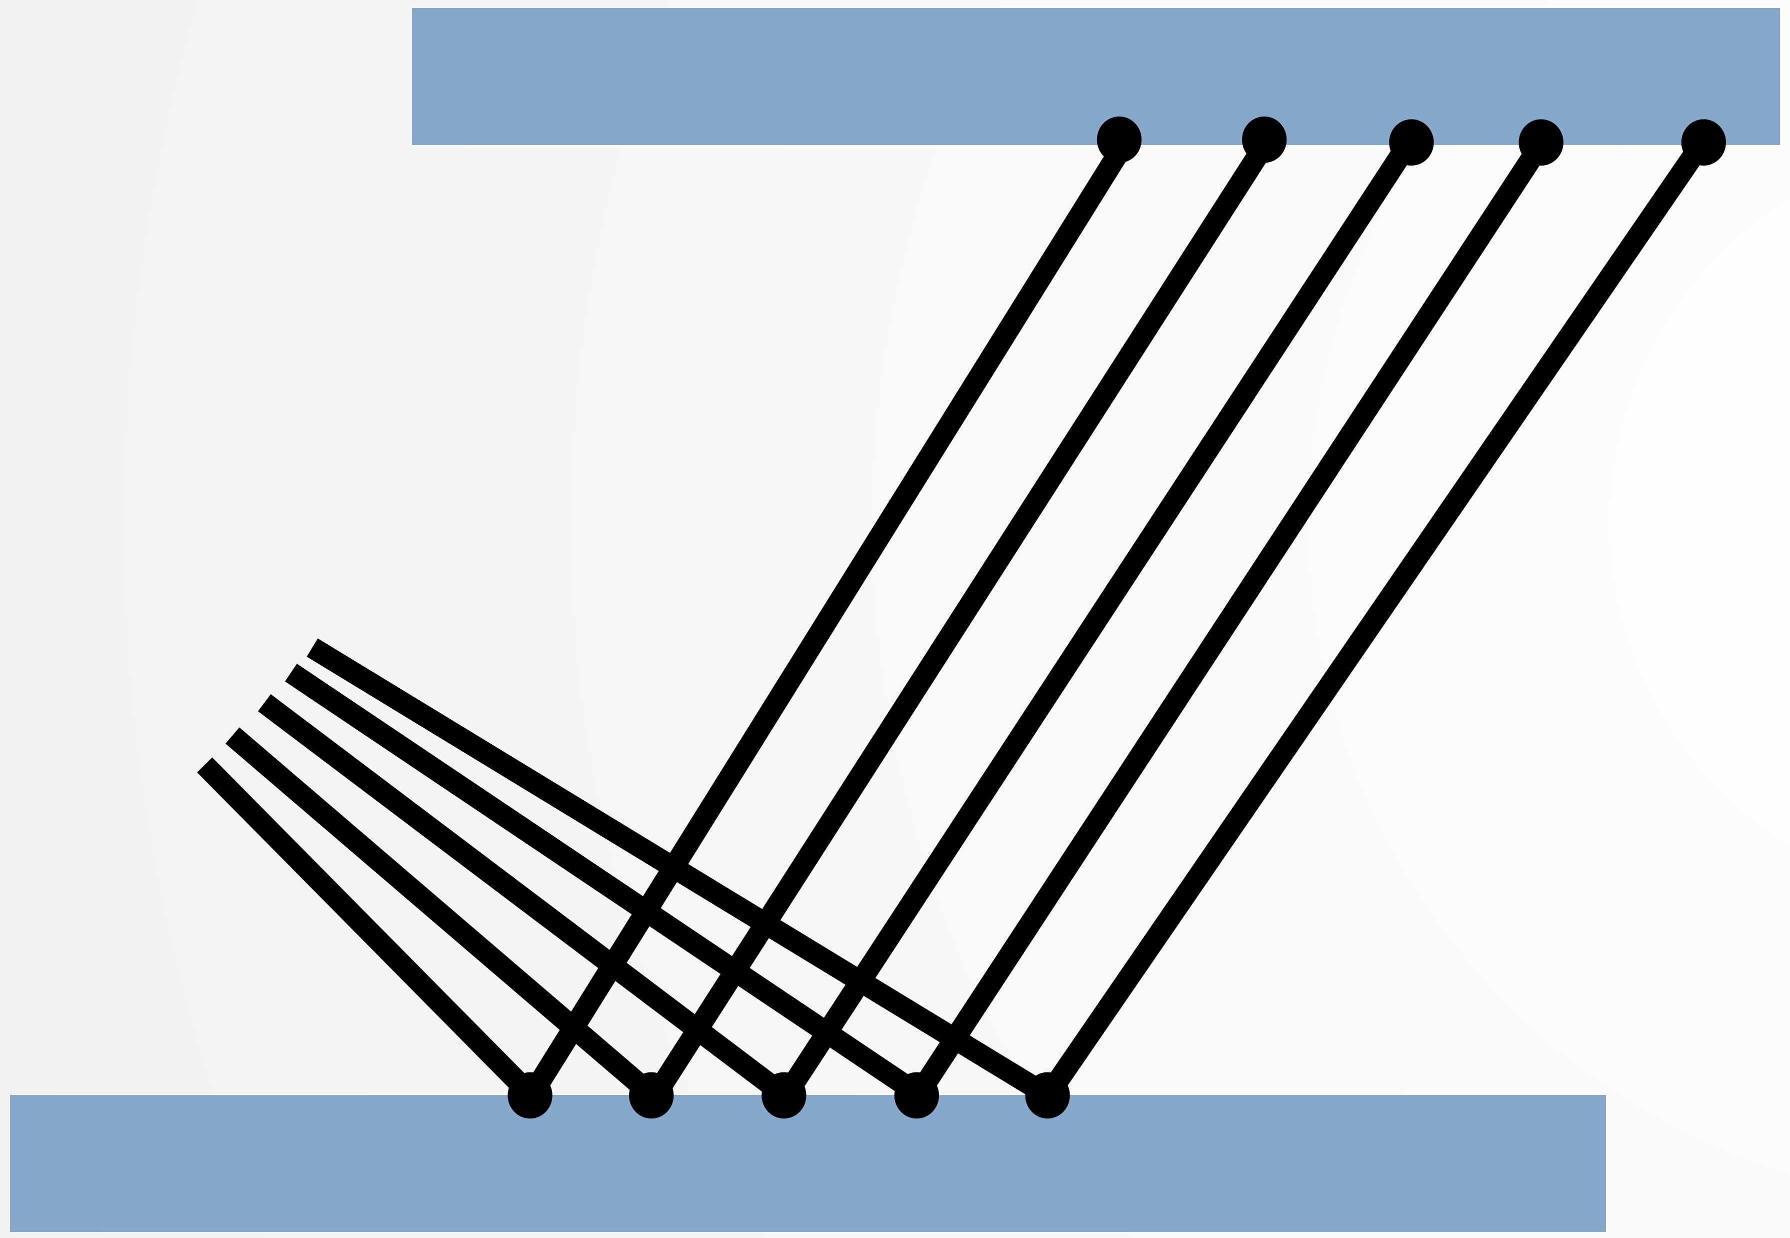

Specular/glossy

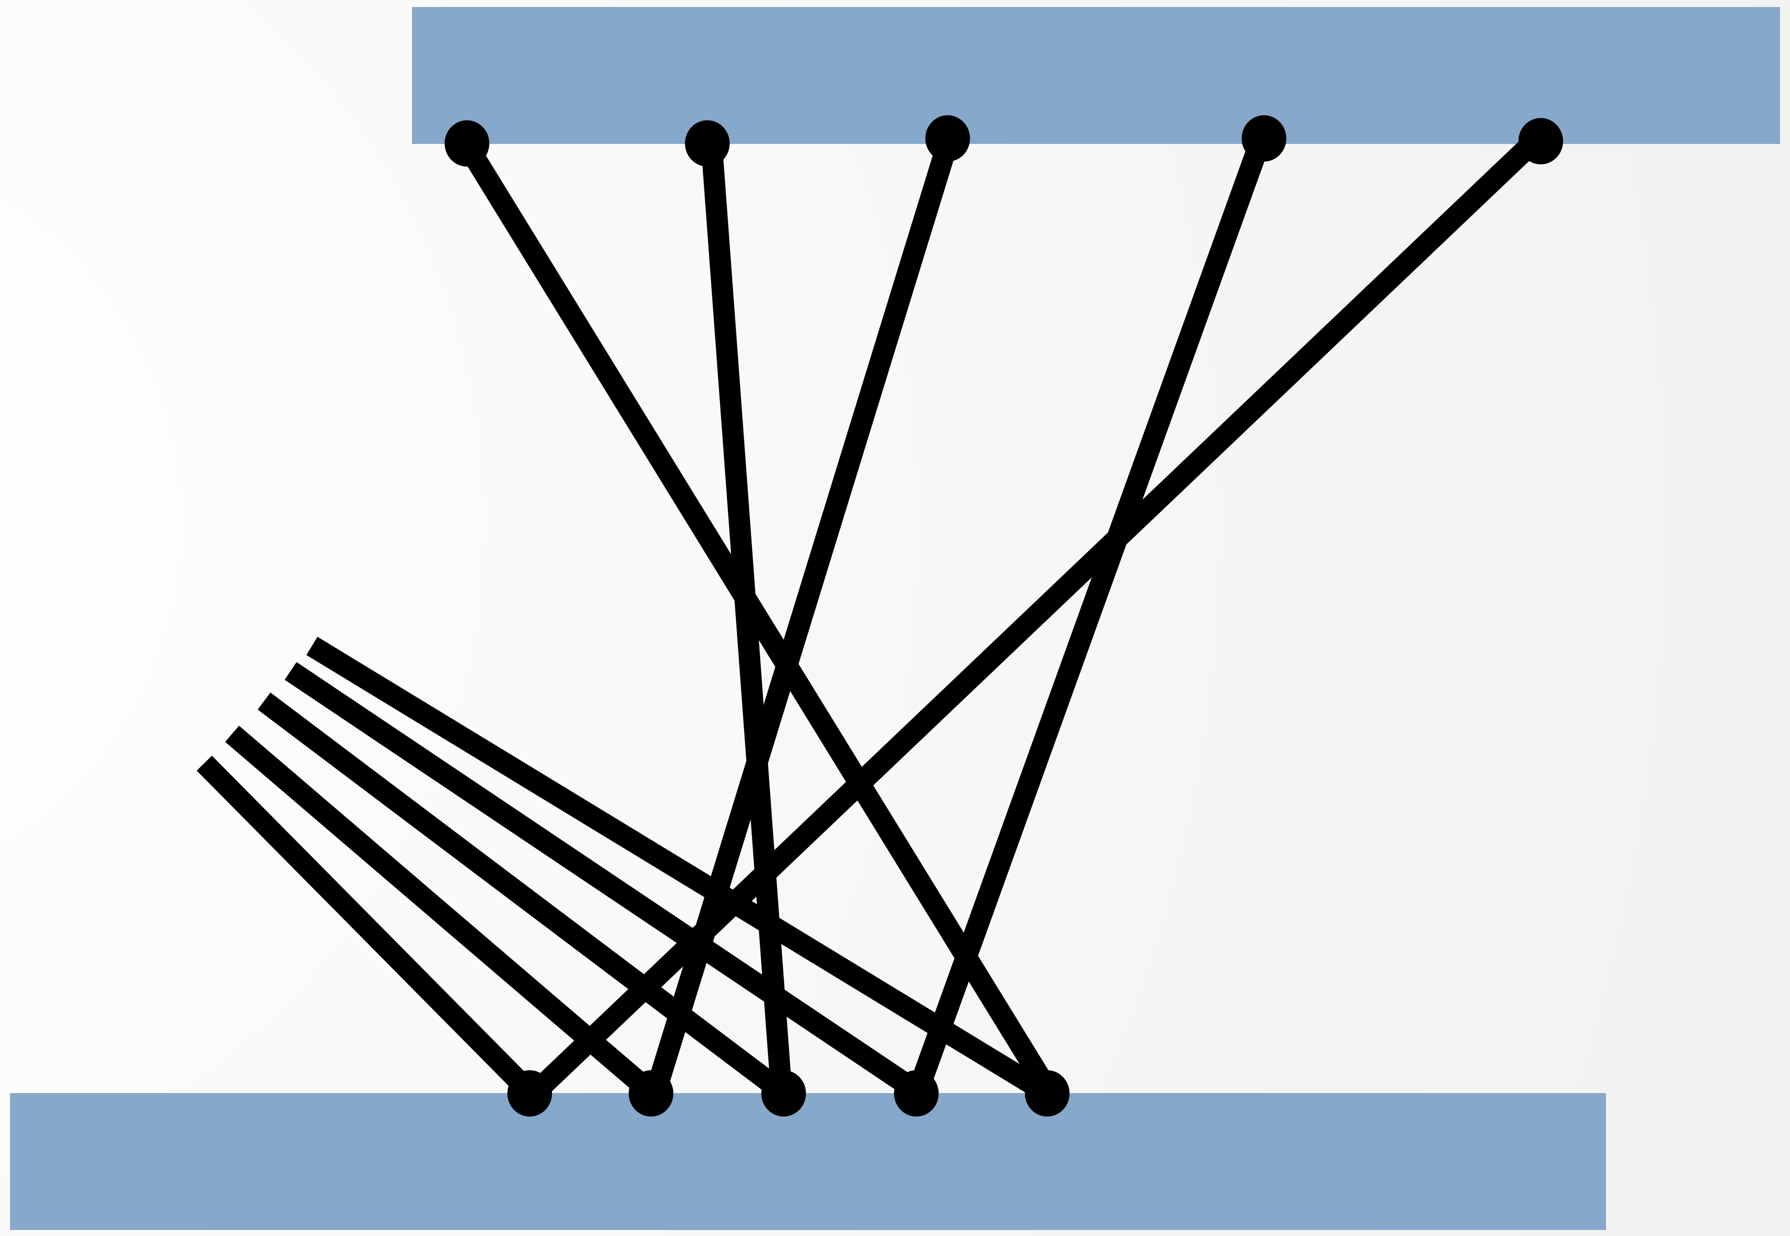

Rough/diffuse

Gaussians

- Compute Gaussians for sampling using nearest neighbours

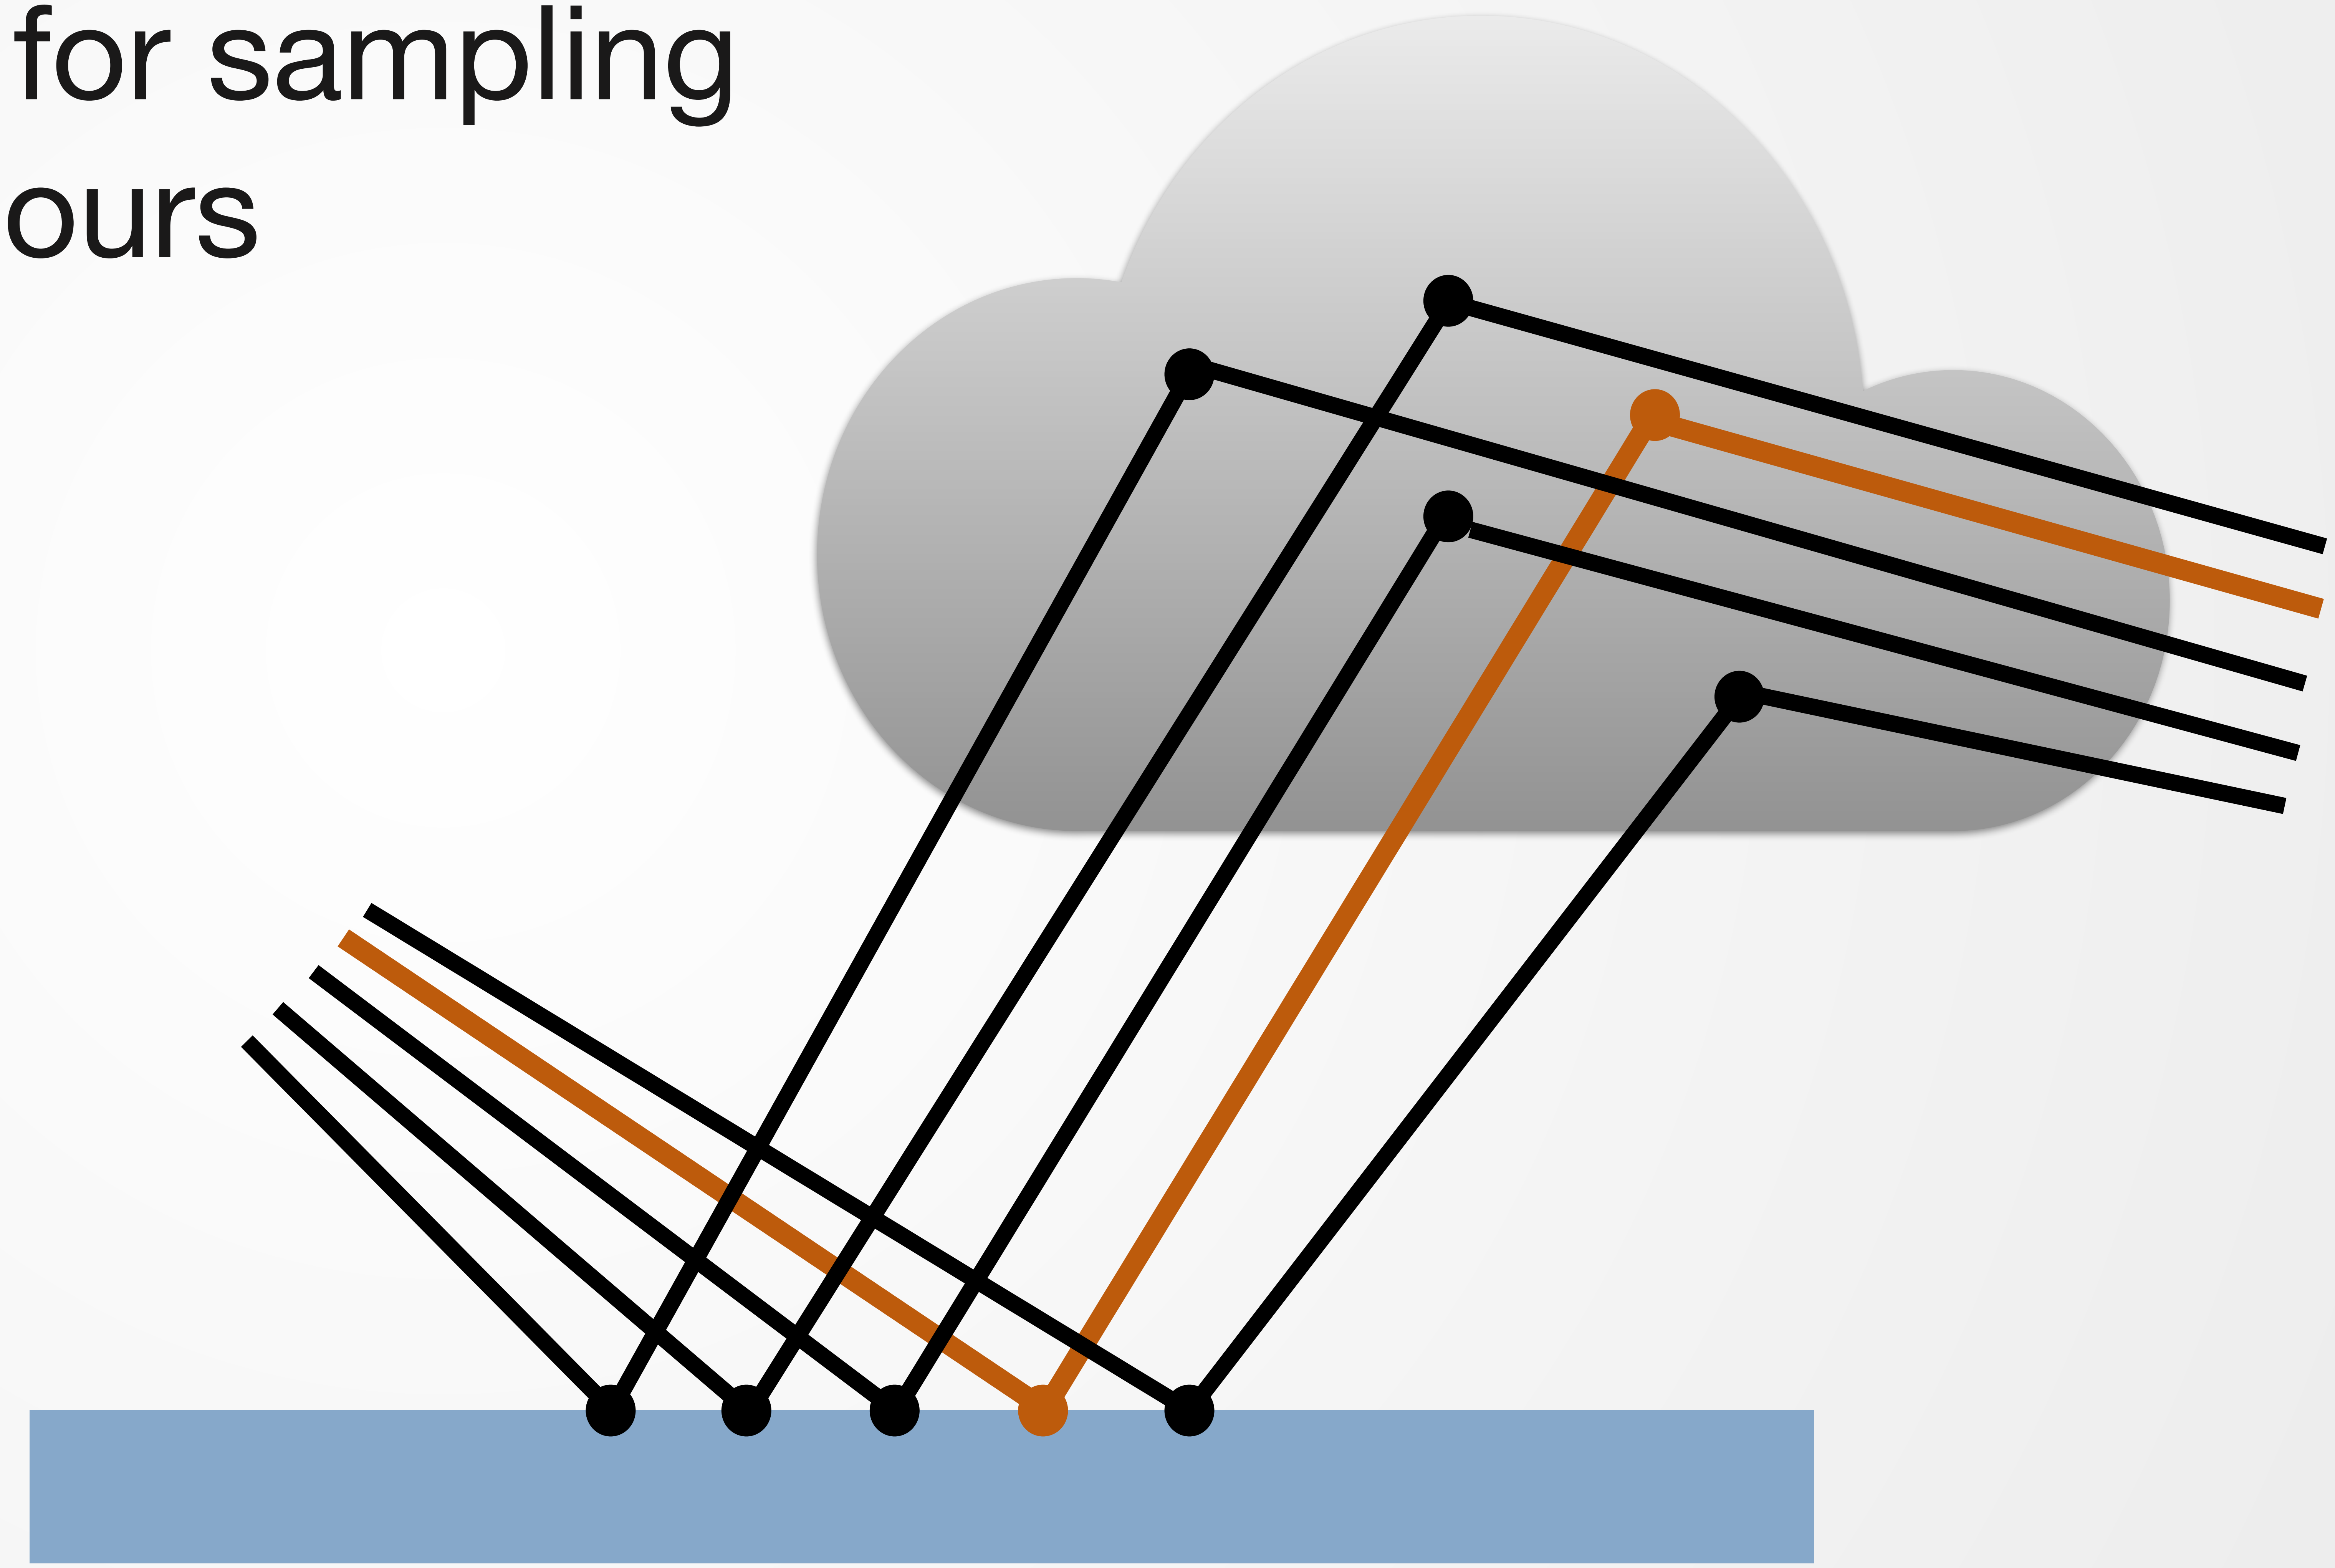

Guide path

Nearest neighbours

Gaussians

- Compute Gaussians for sampling using nearest neighbours

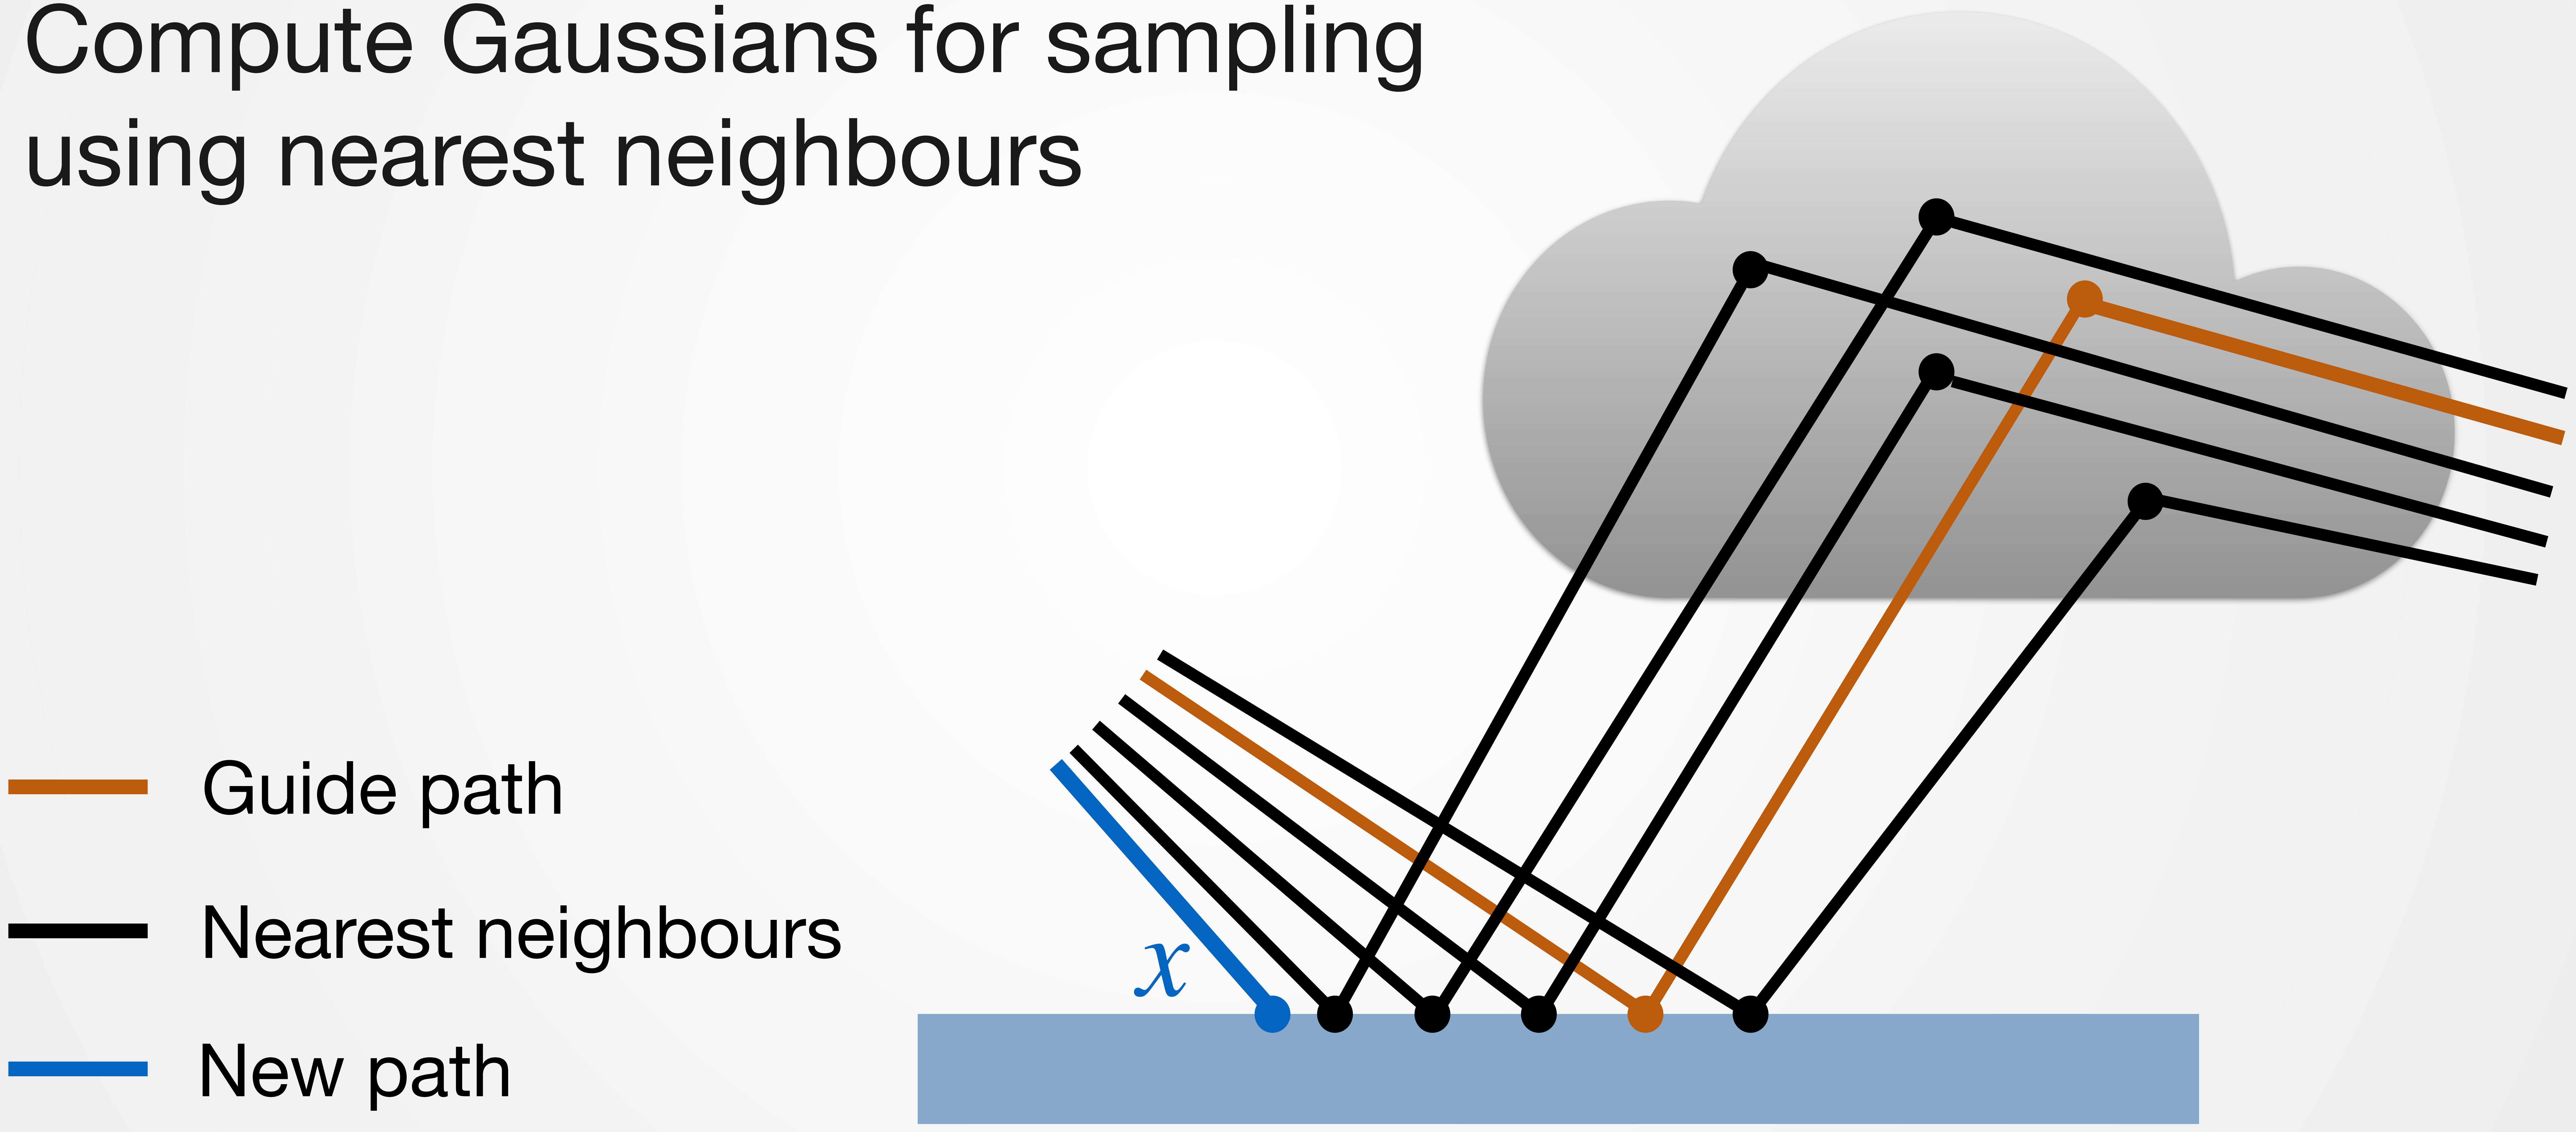

Gaussians

- Sample 3D Gaussian at next vertex?

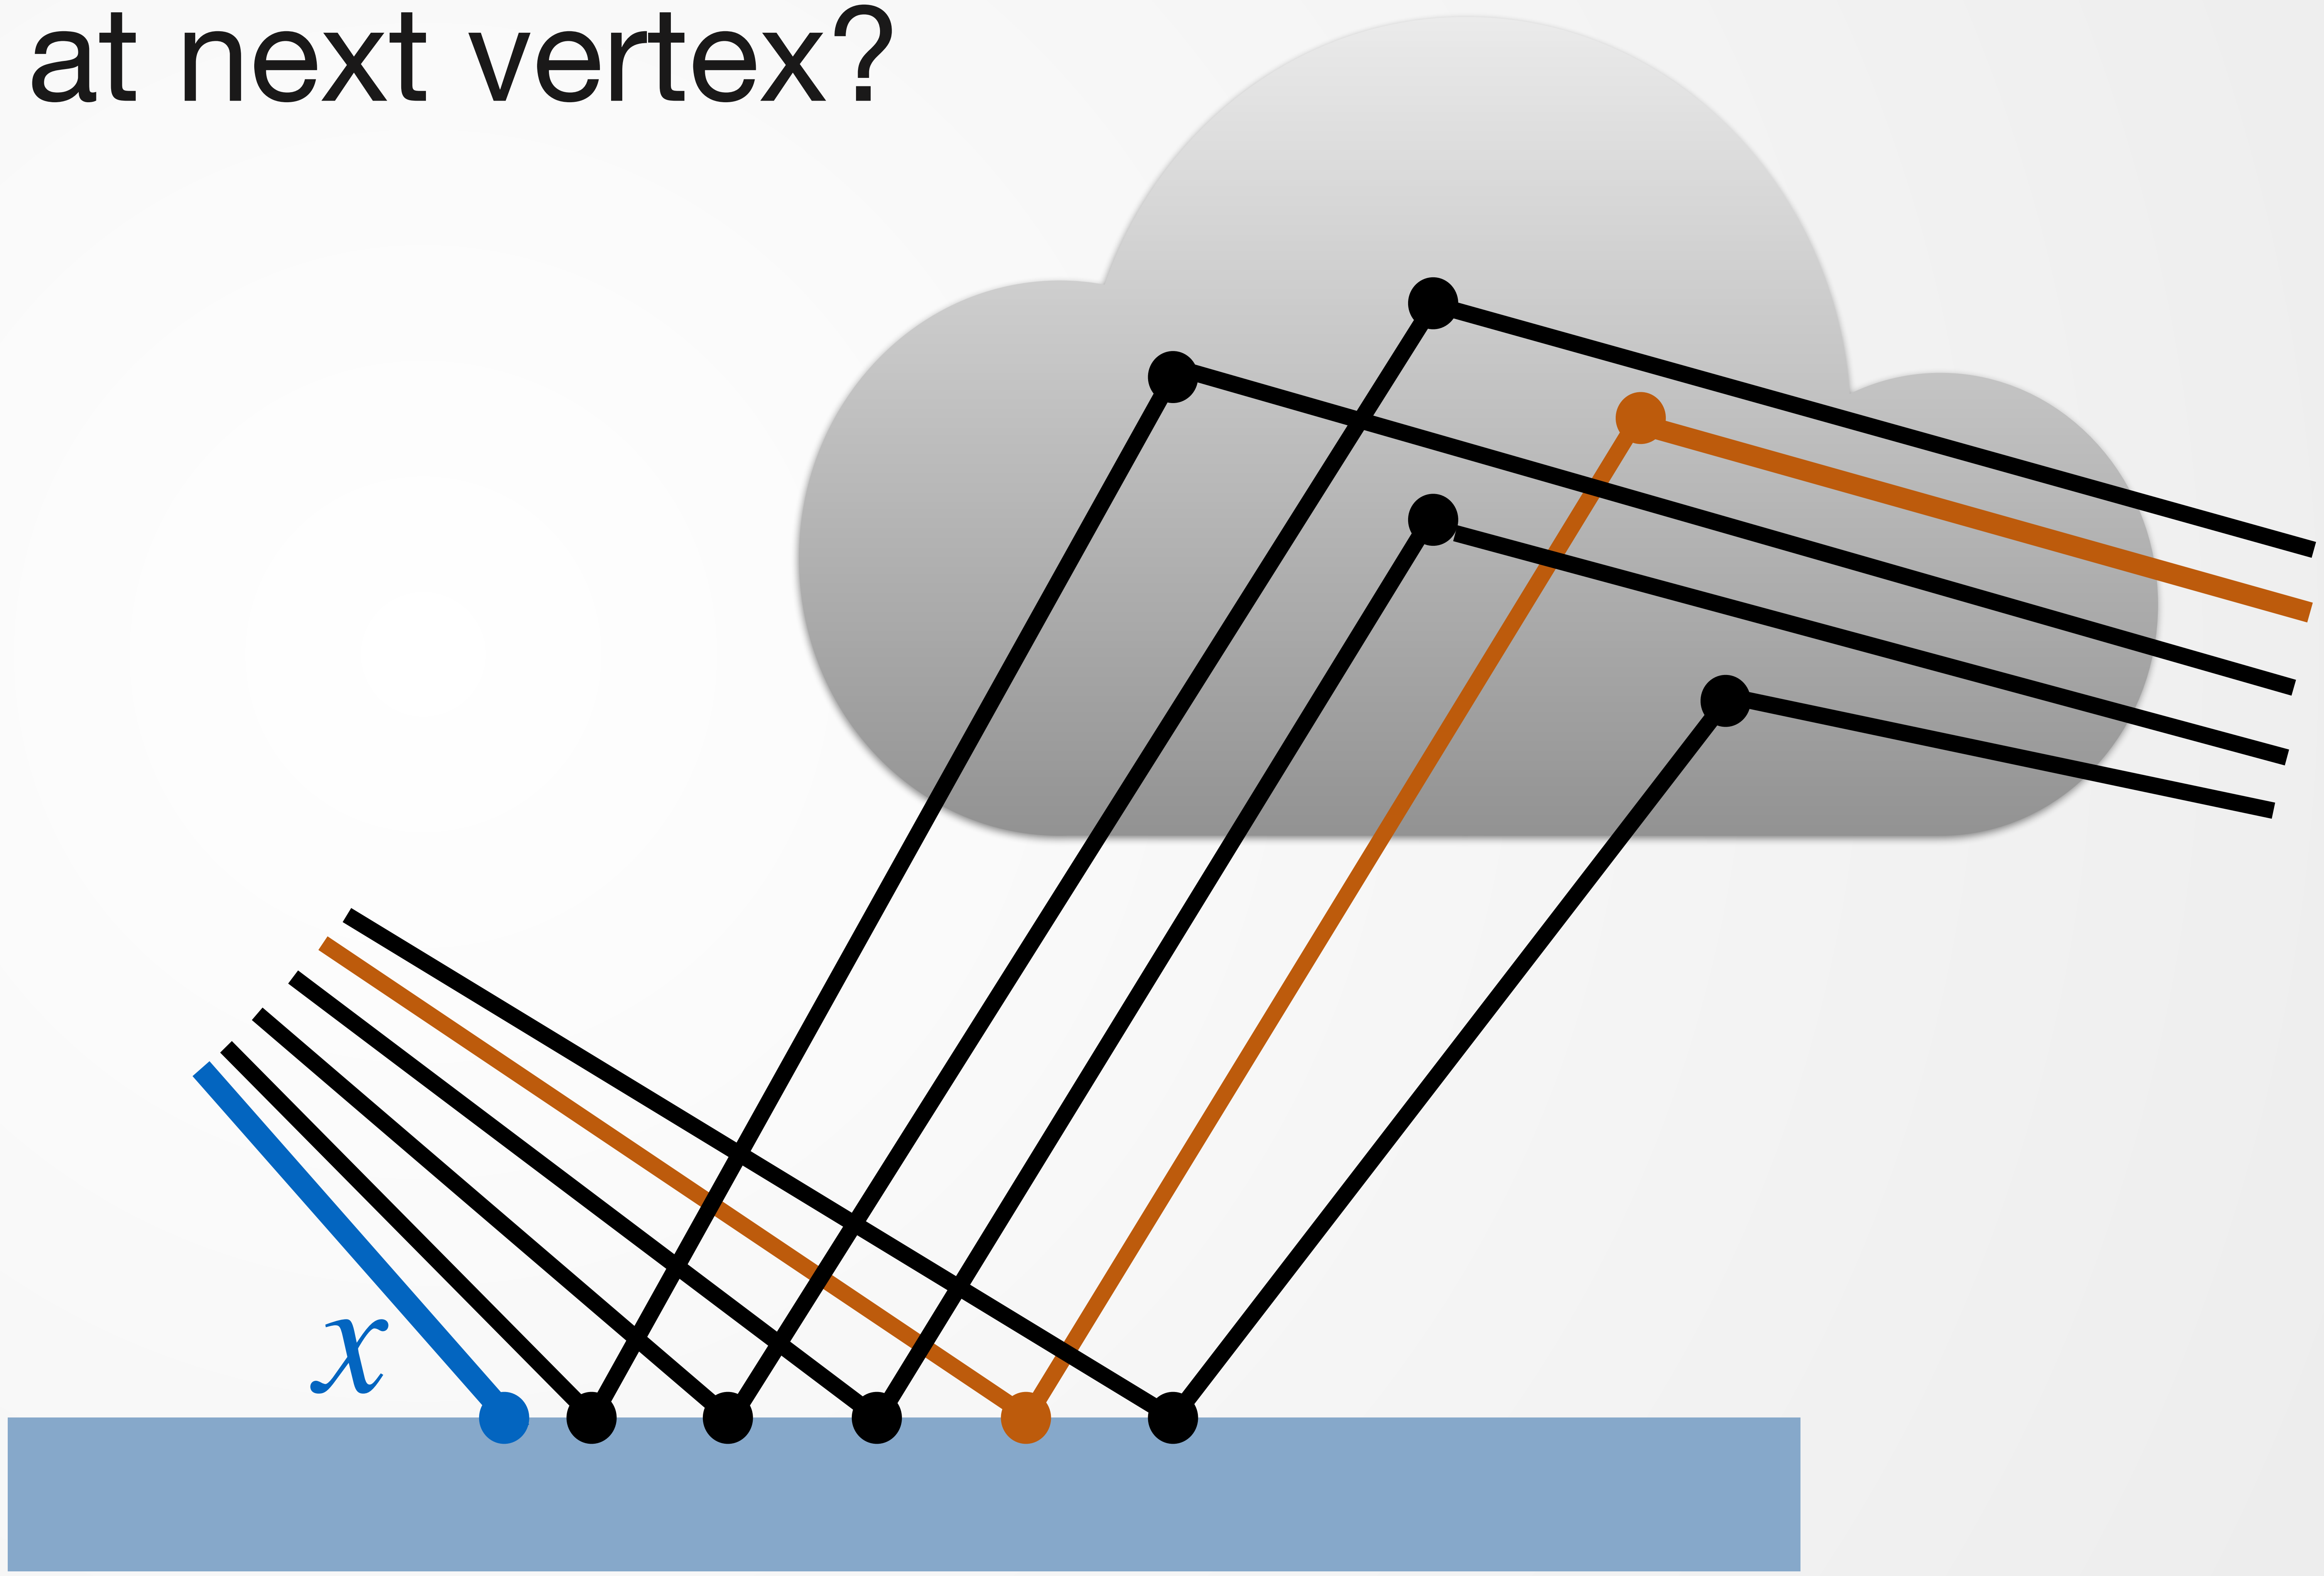

Guide path

Nearest neighbours

New path

Gaussians

- Sample 3D Gaussian at next vertex?

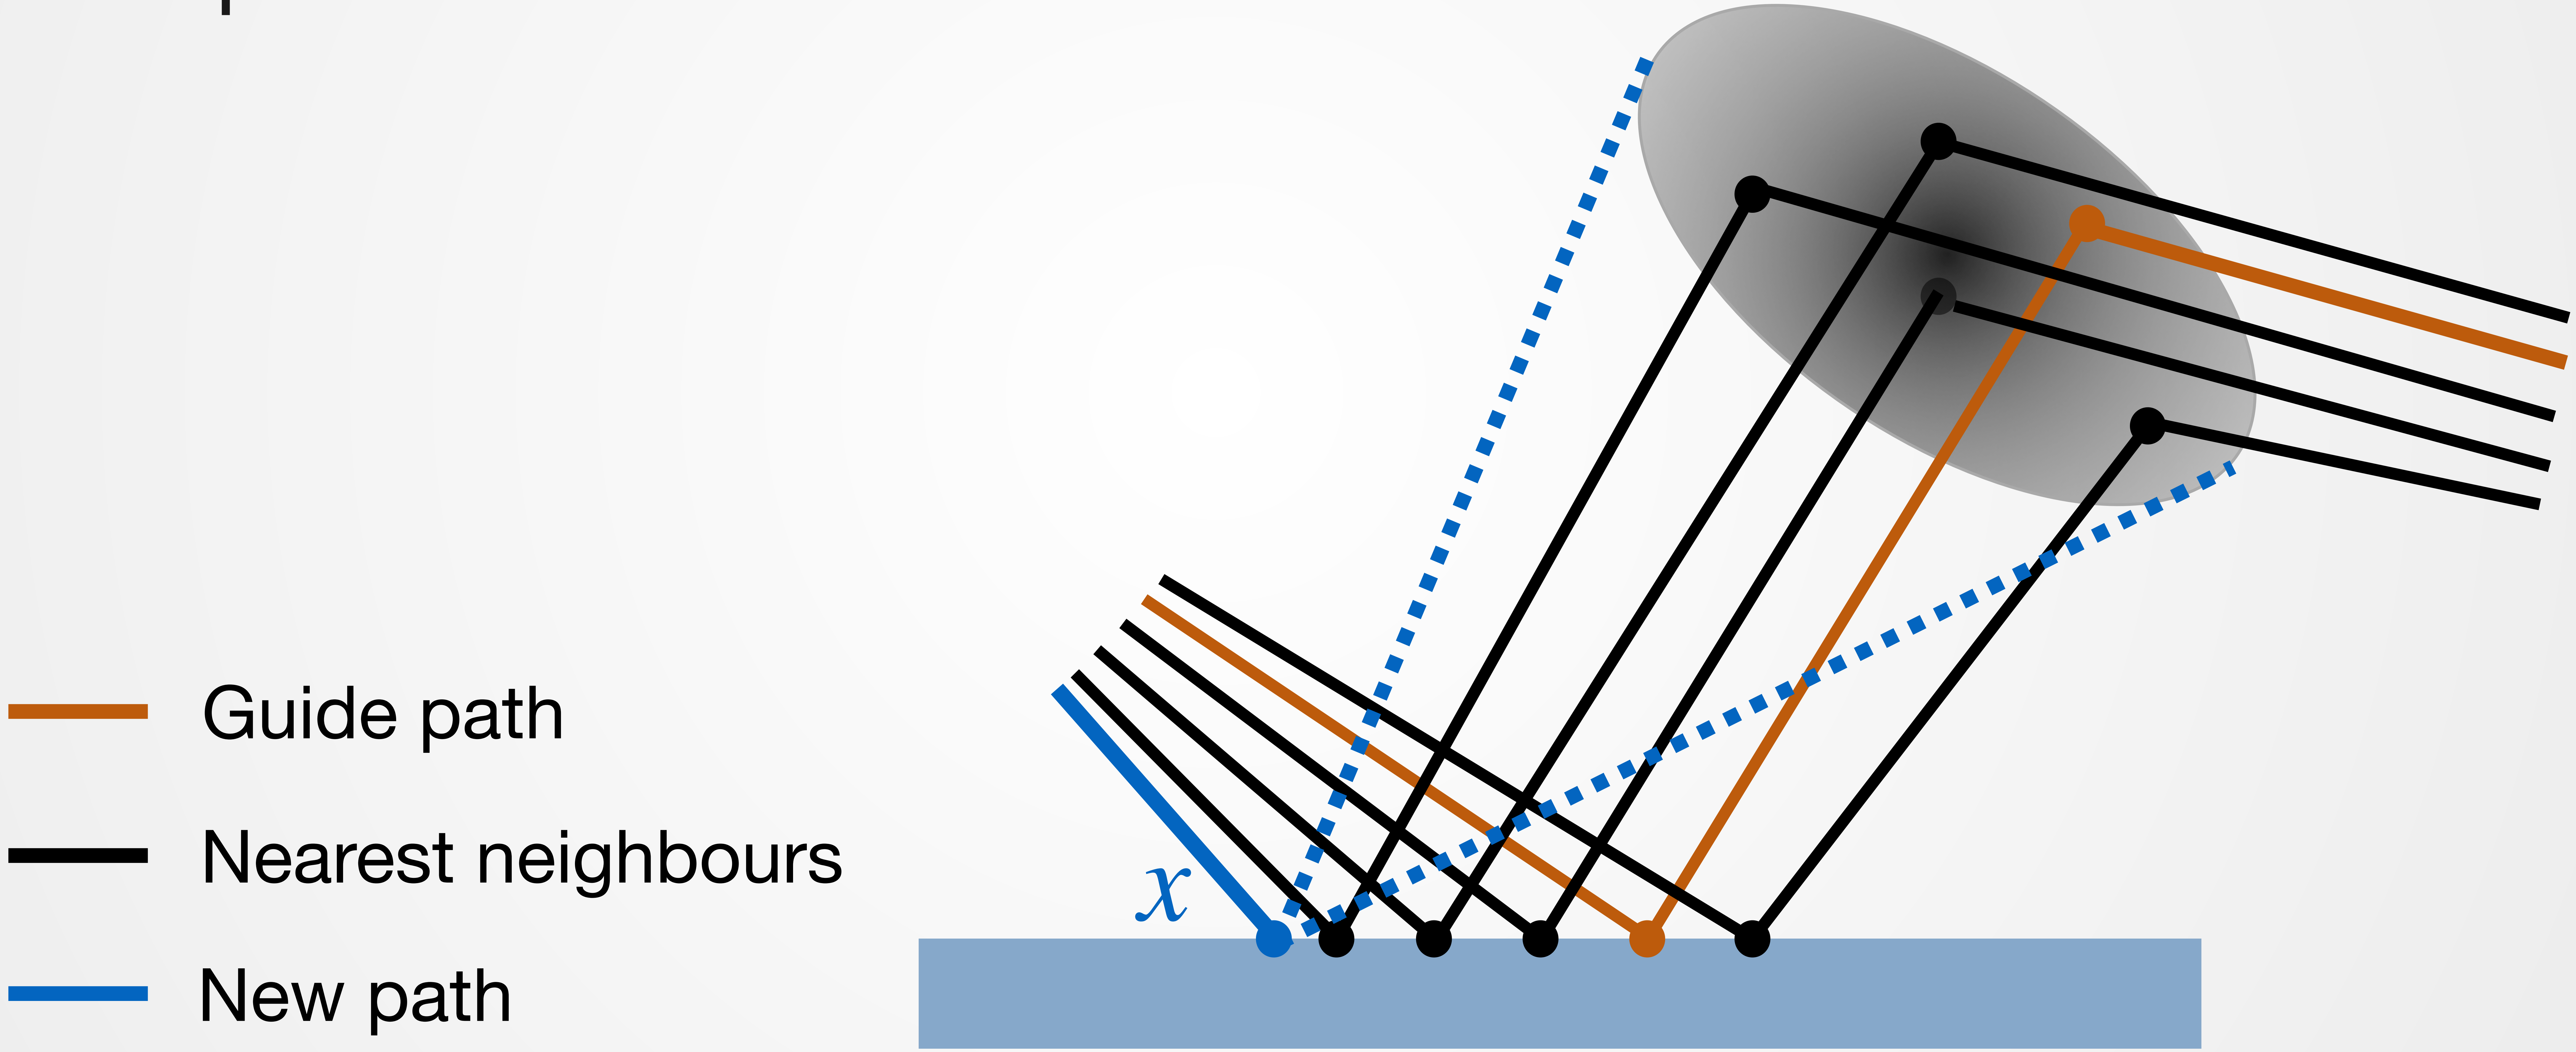

Gaussians

- Compute 6D covariance matrix for path segments

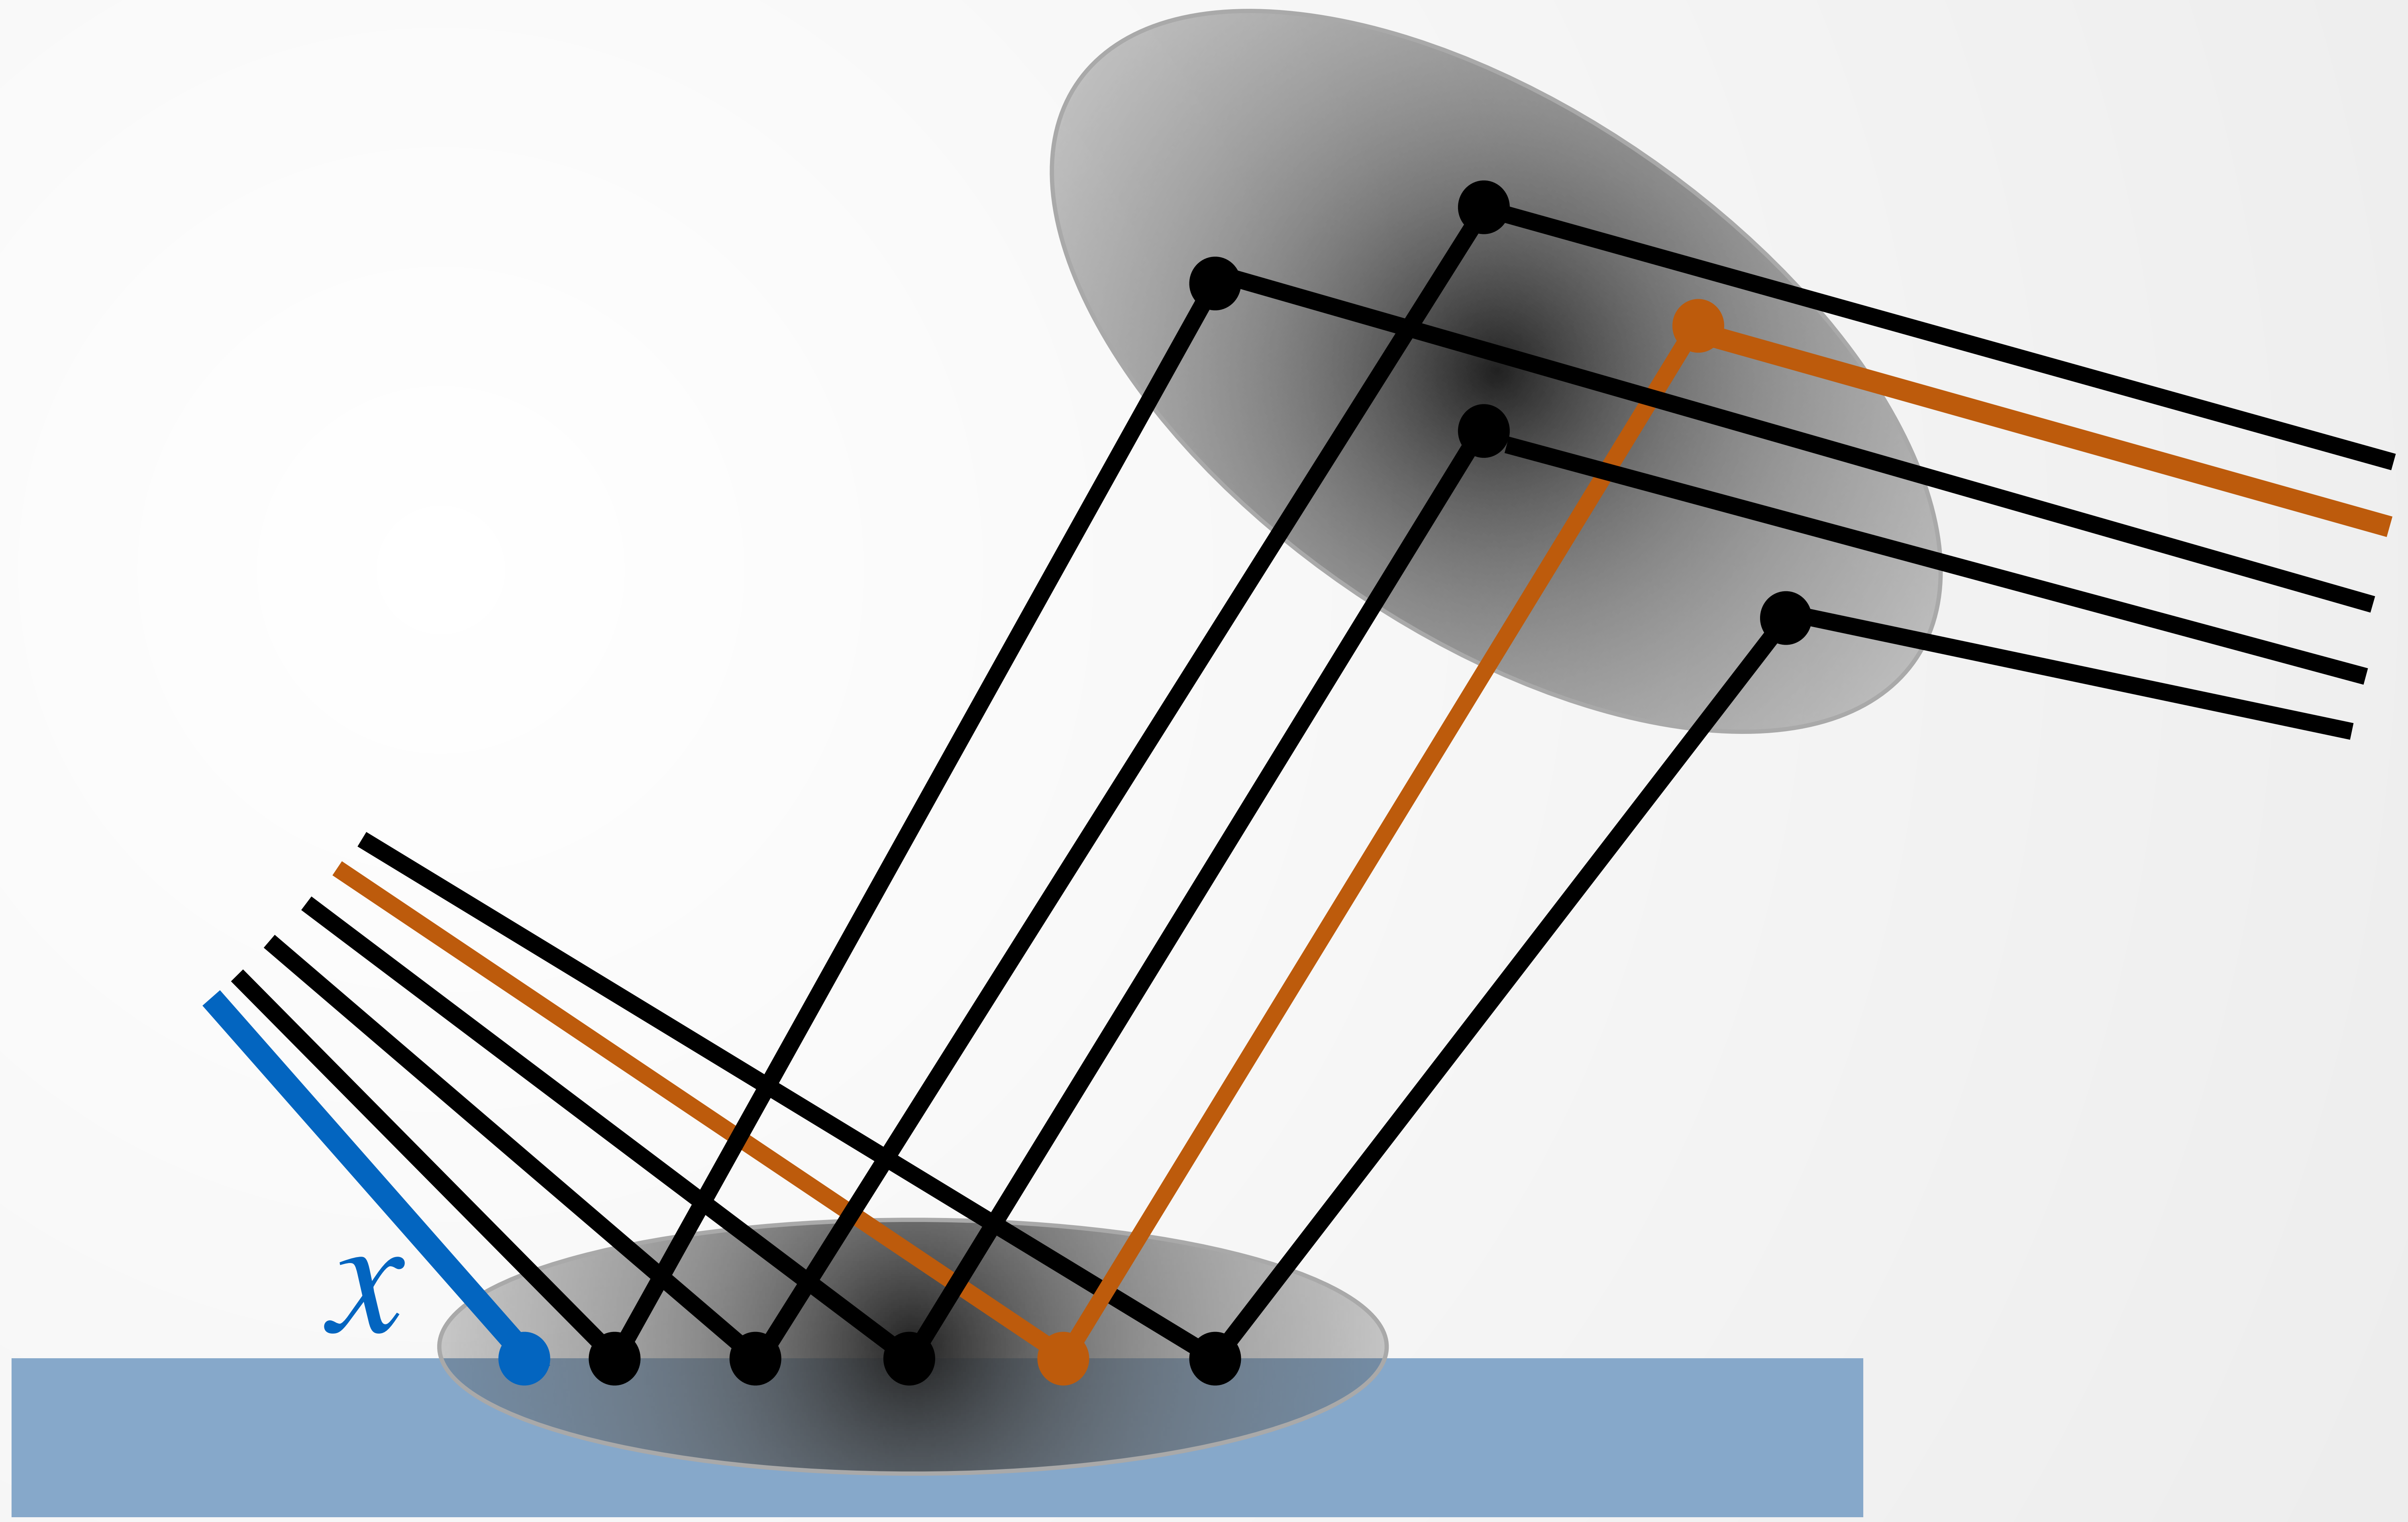

Guide path

Nearest neighbours

New path

Gaussians

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

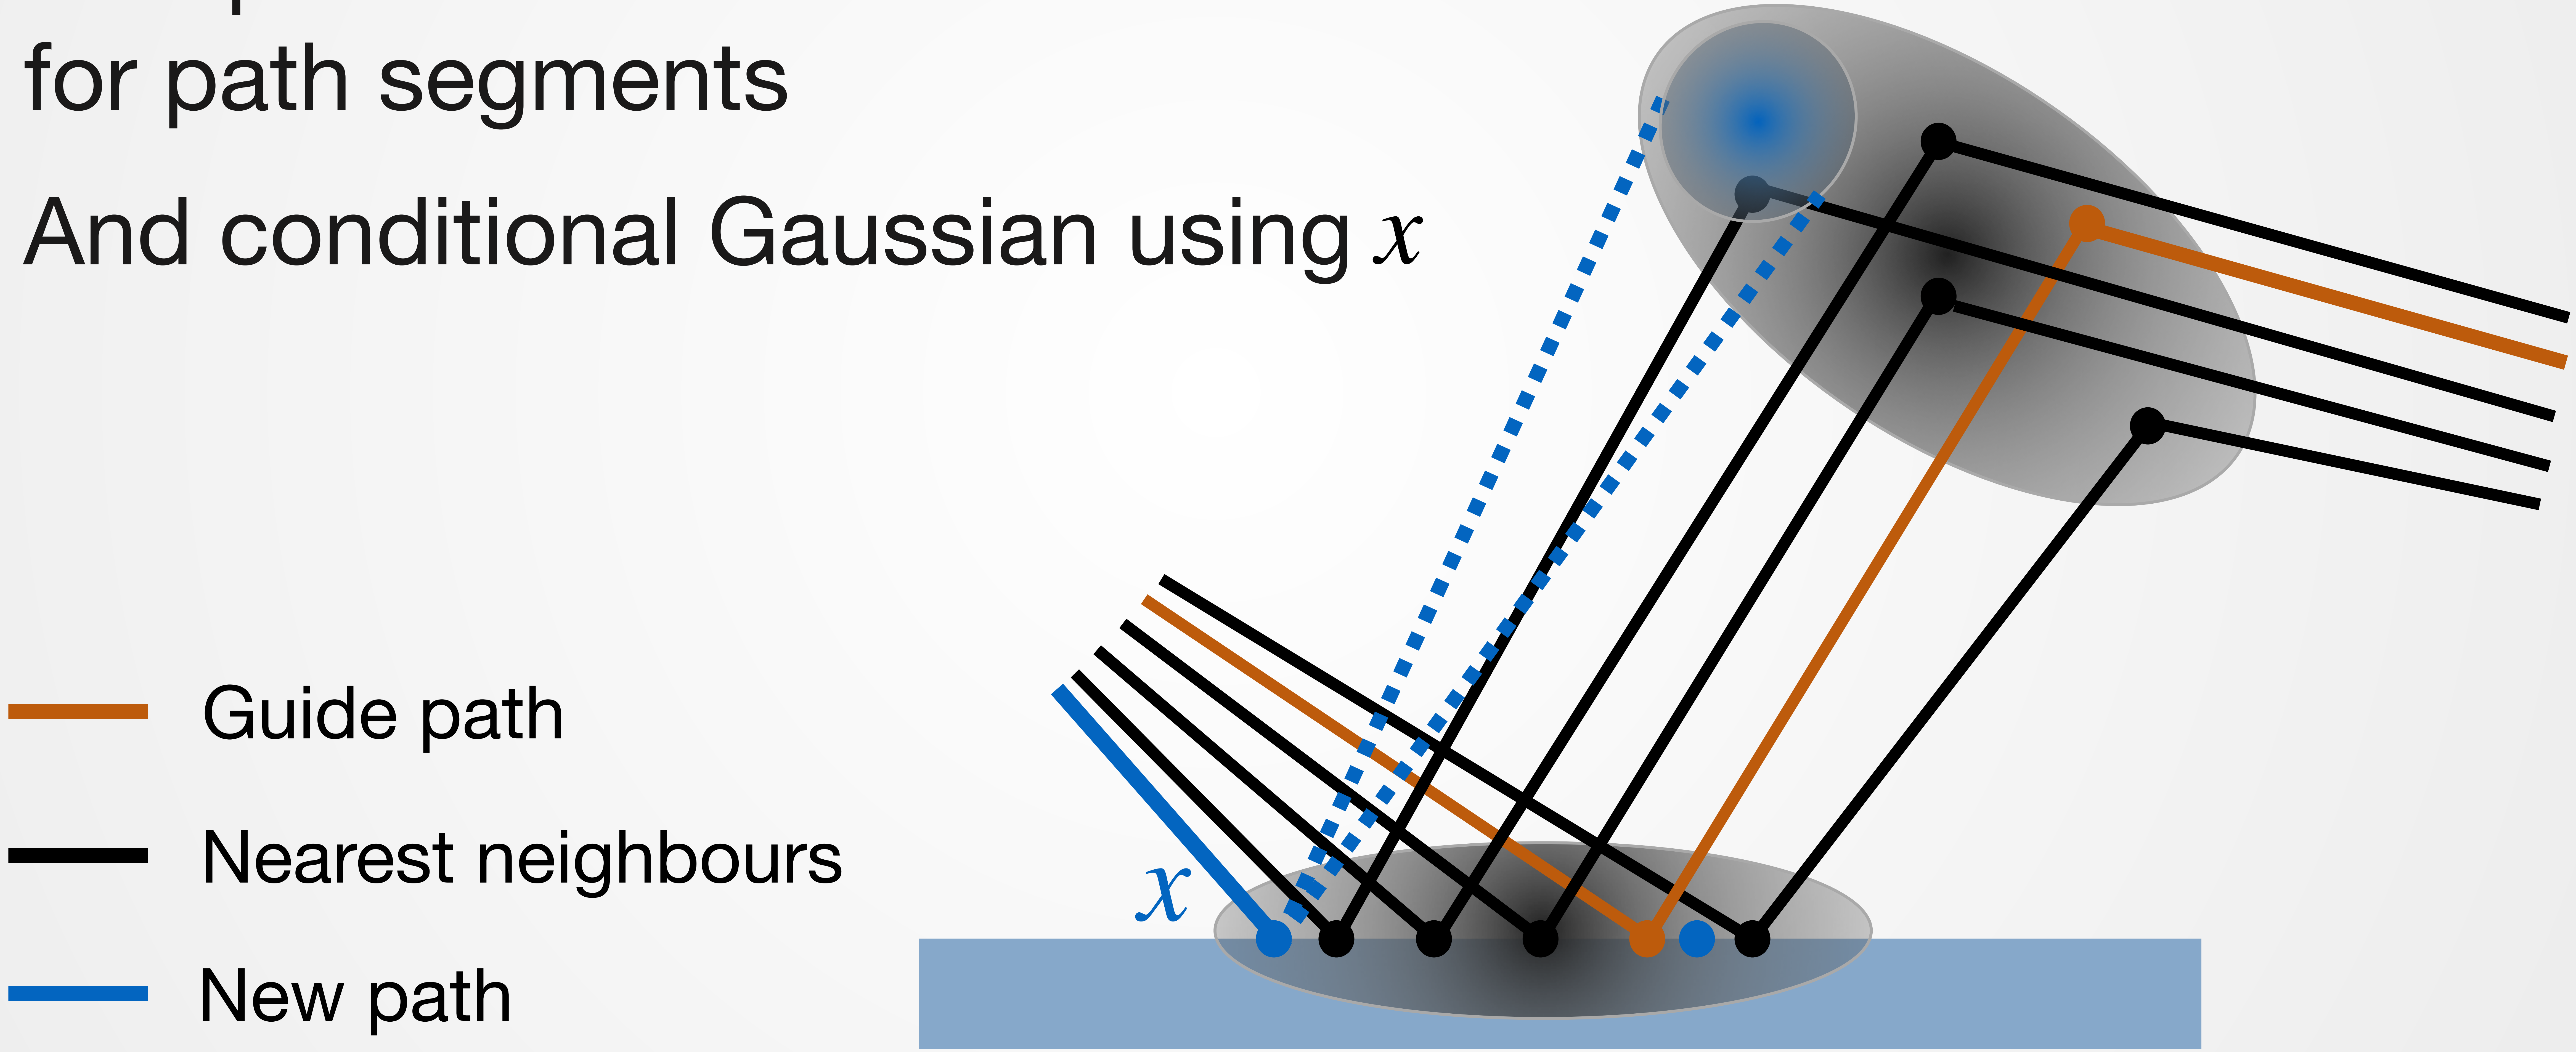

Gaussians

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

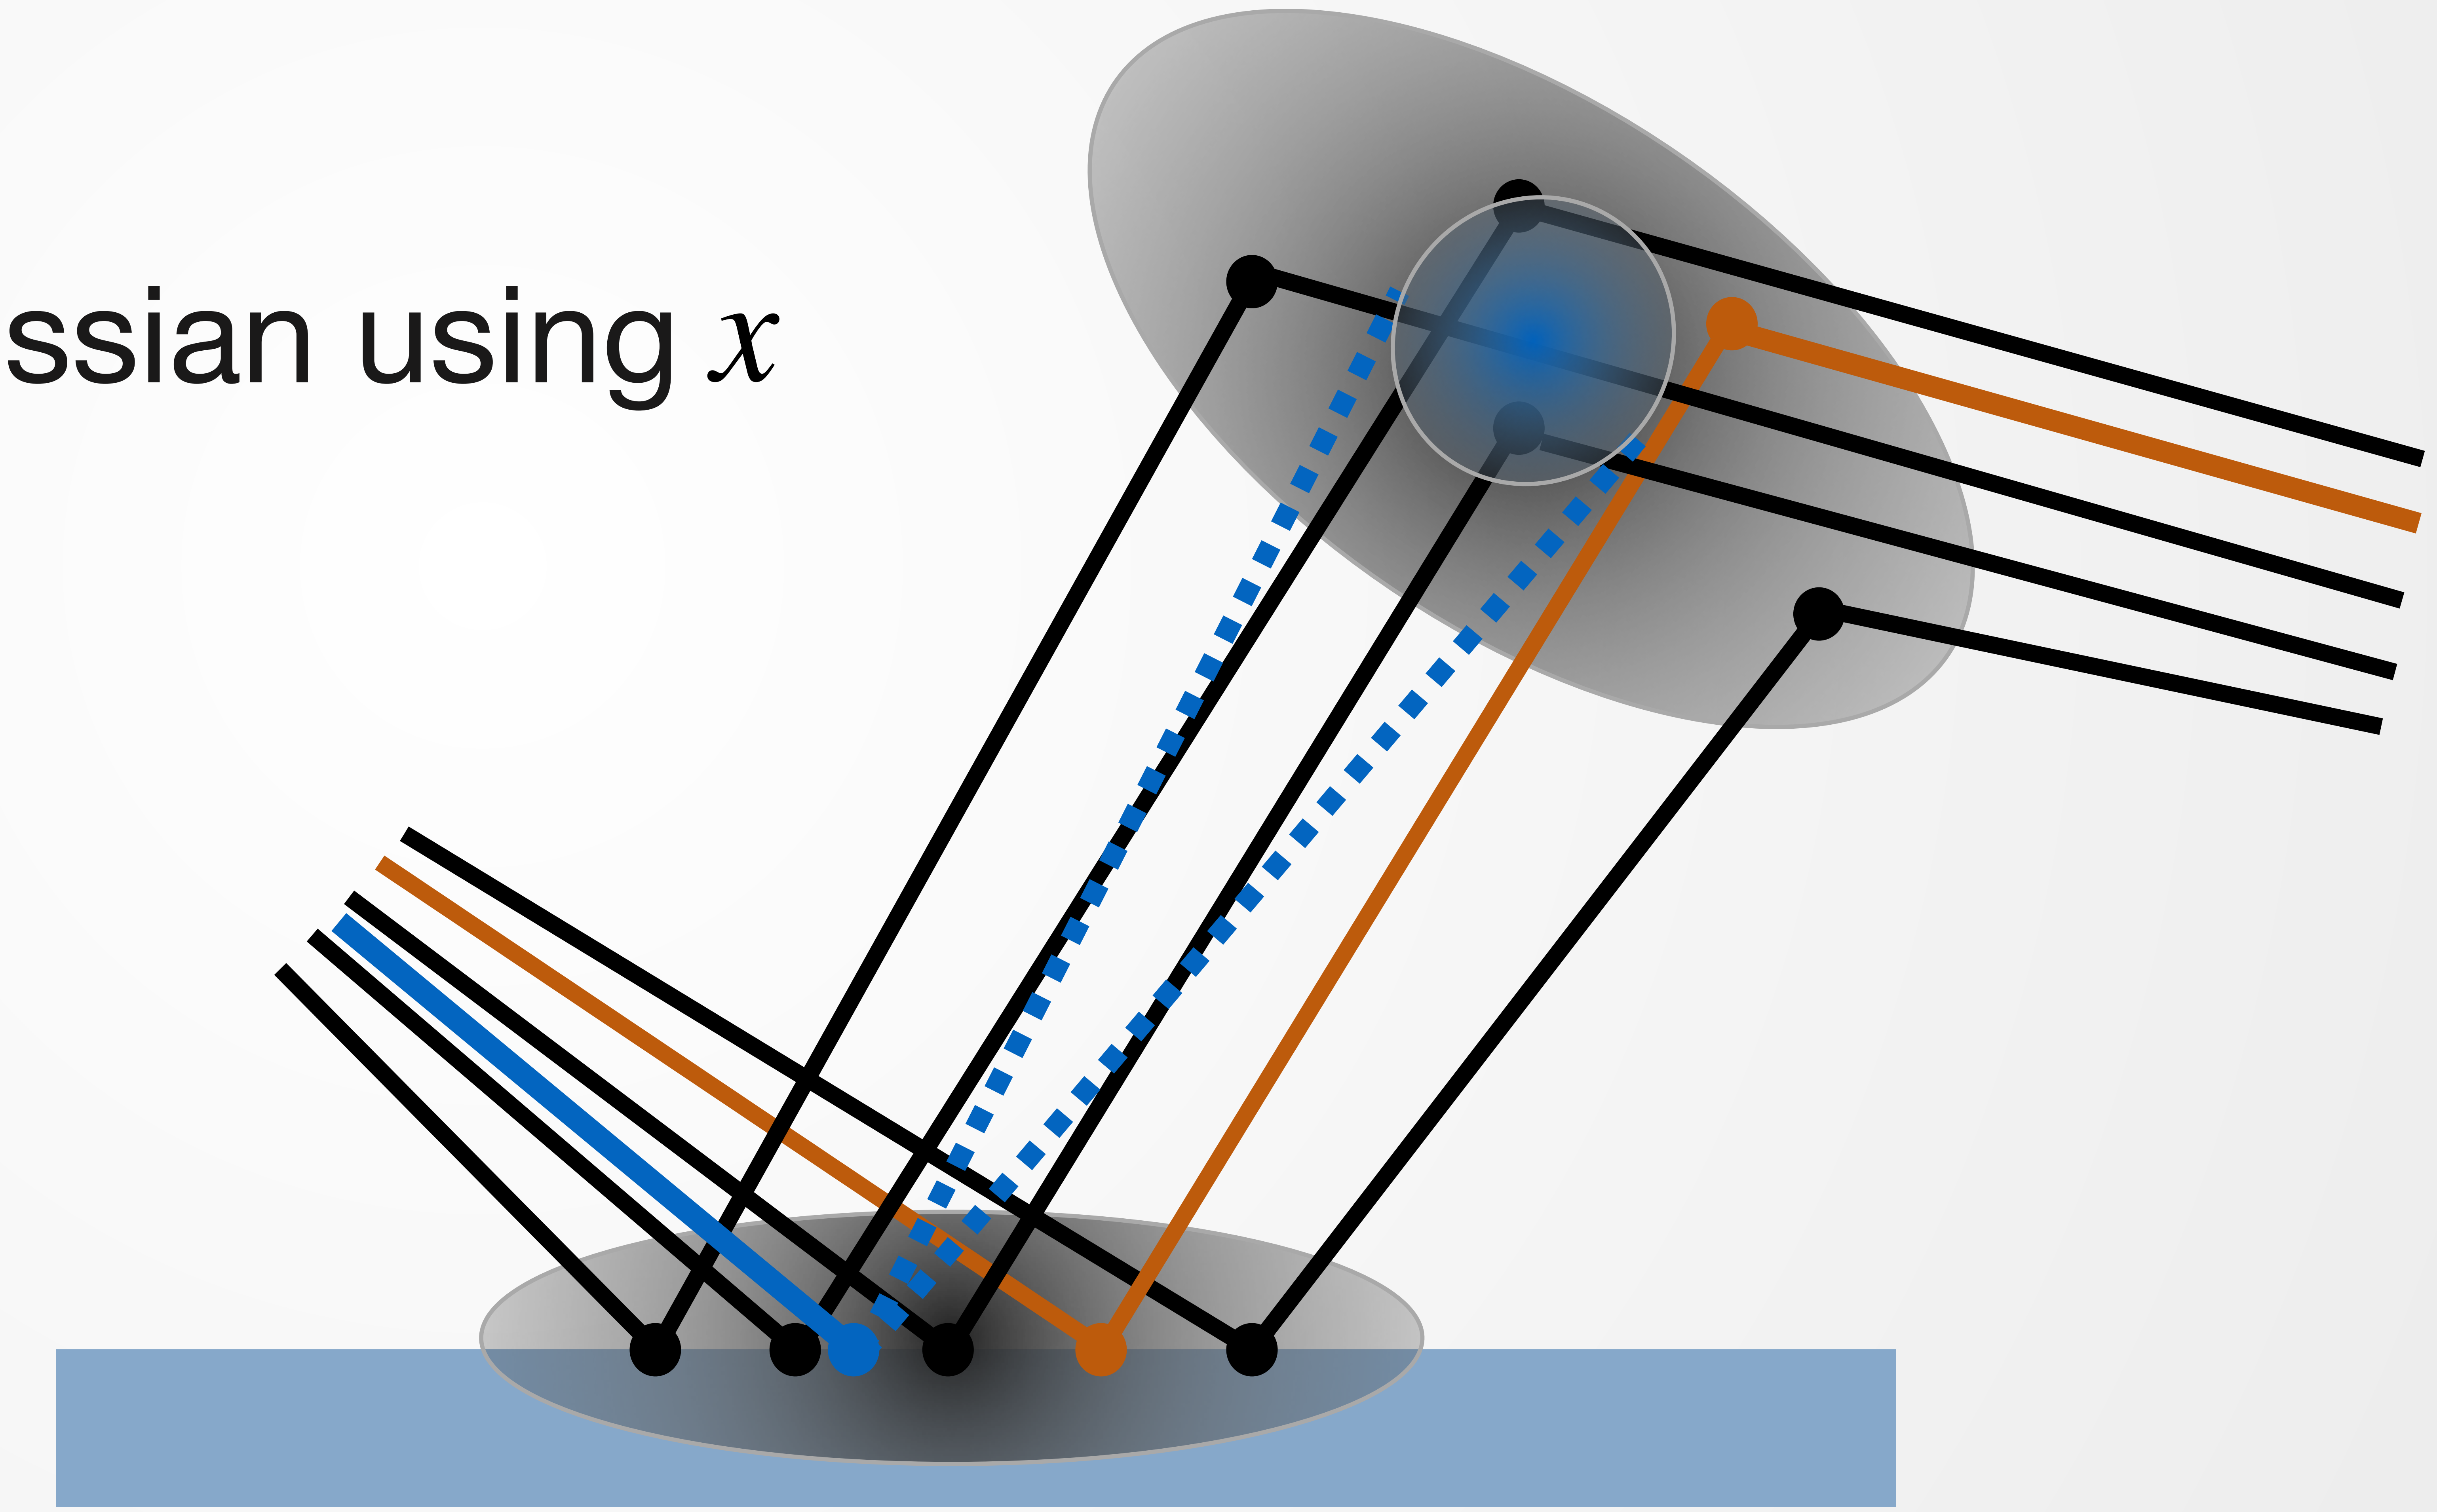

Guide path

Nearest neighbours

New path

Gaussians

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

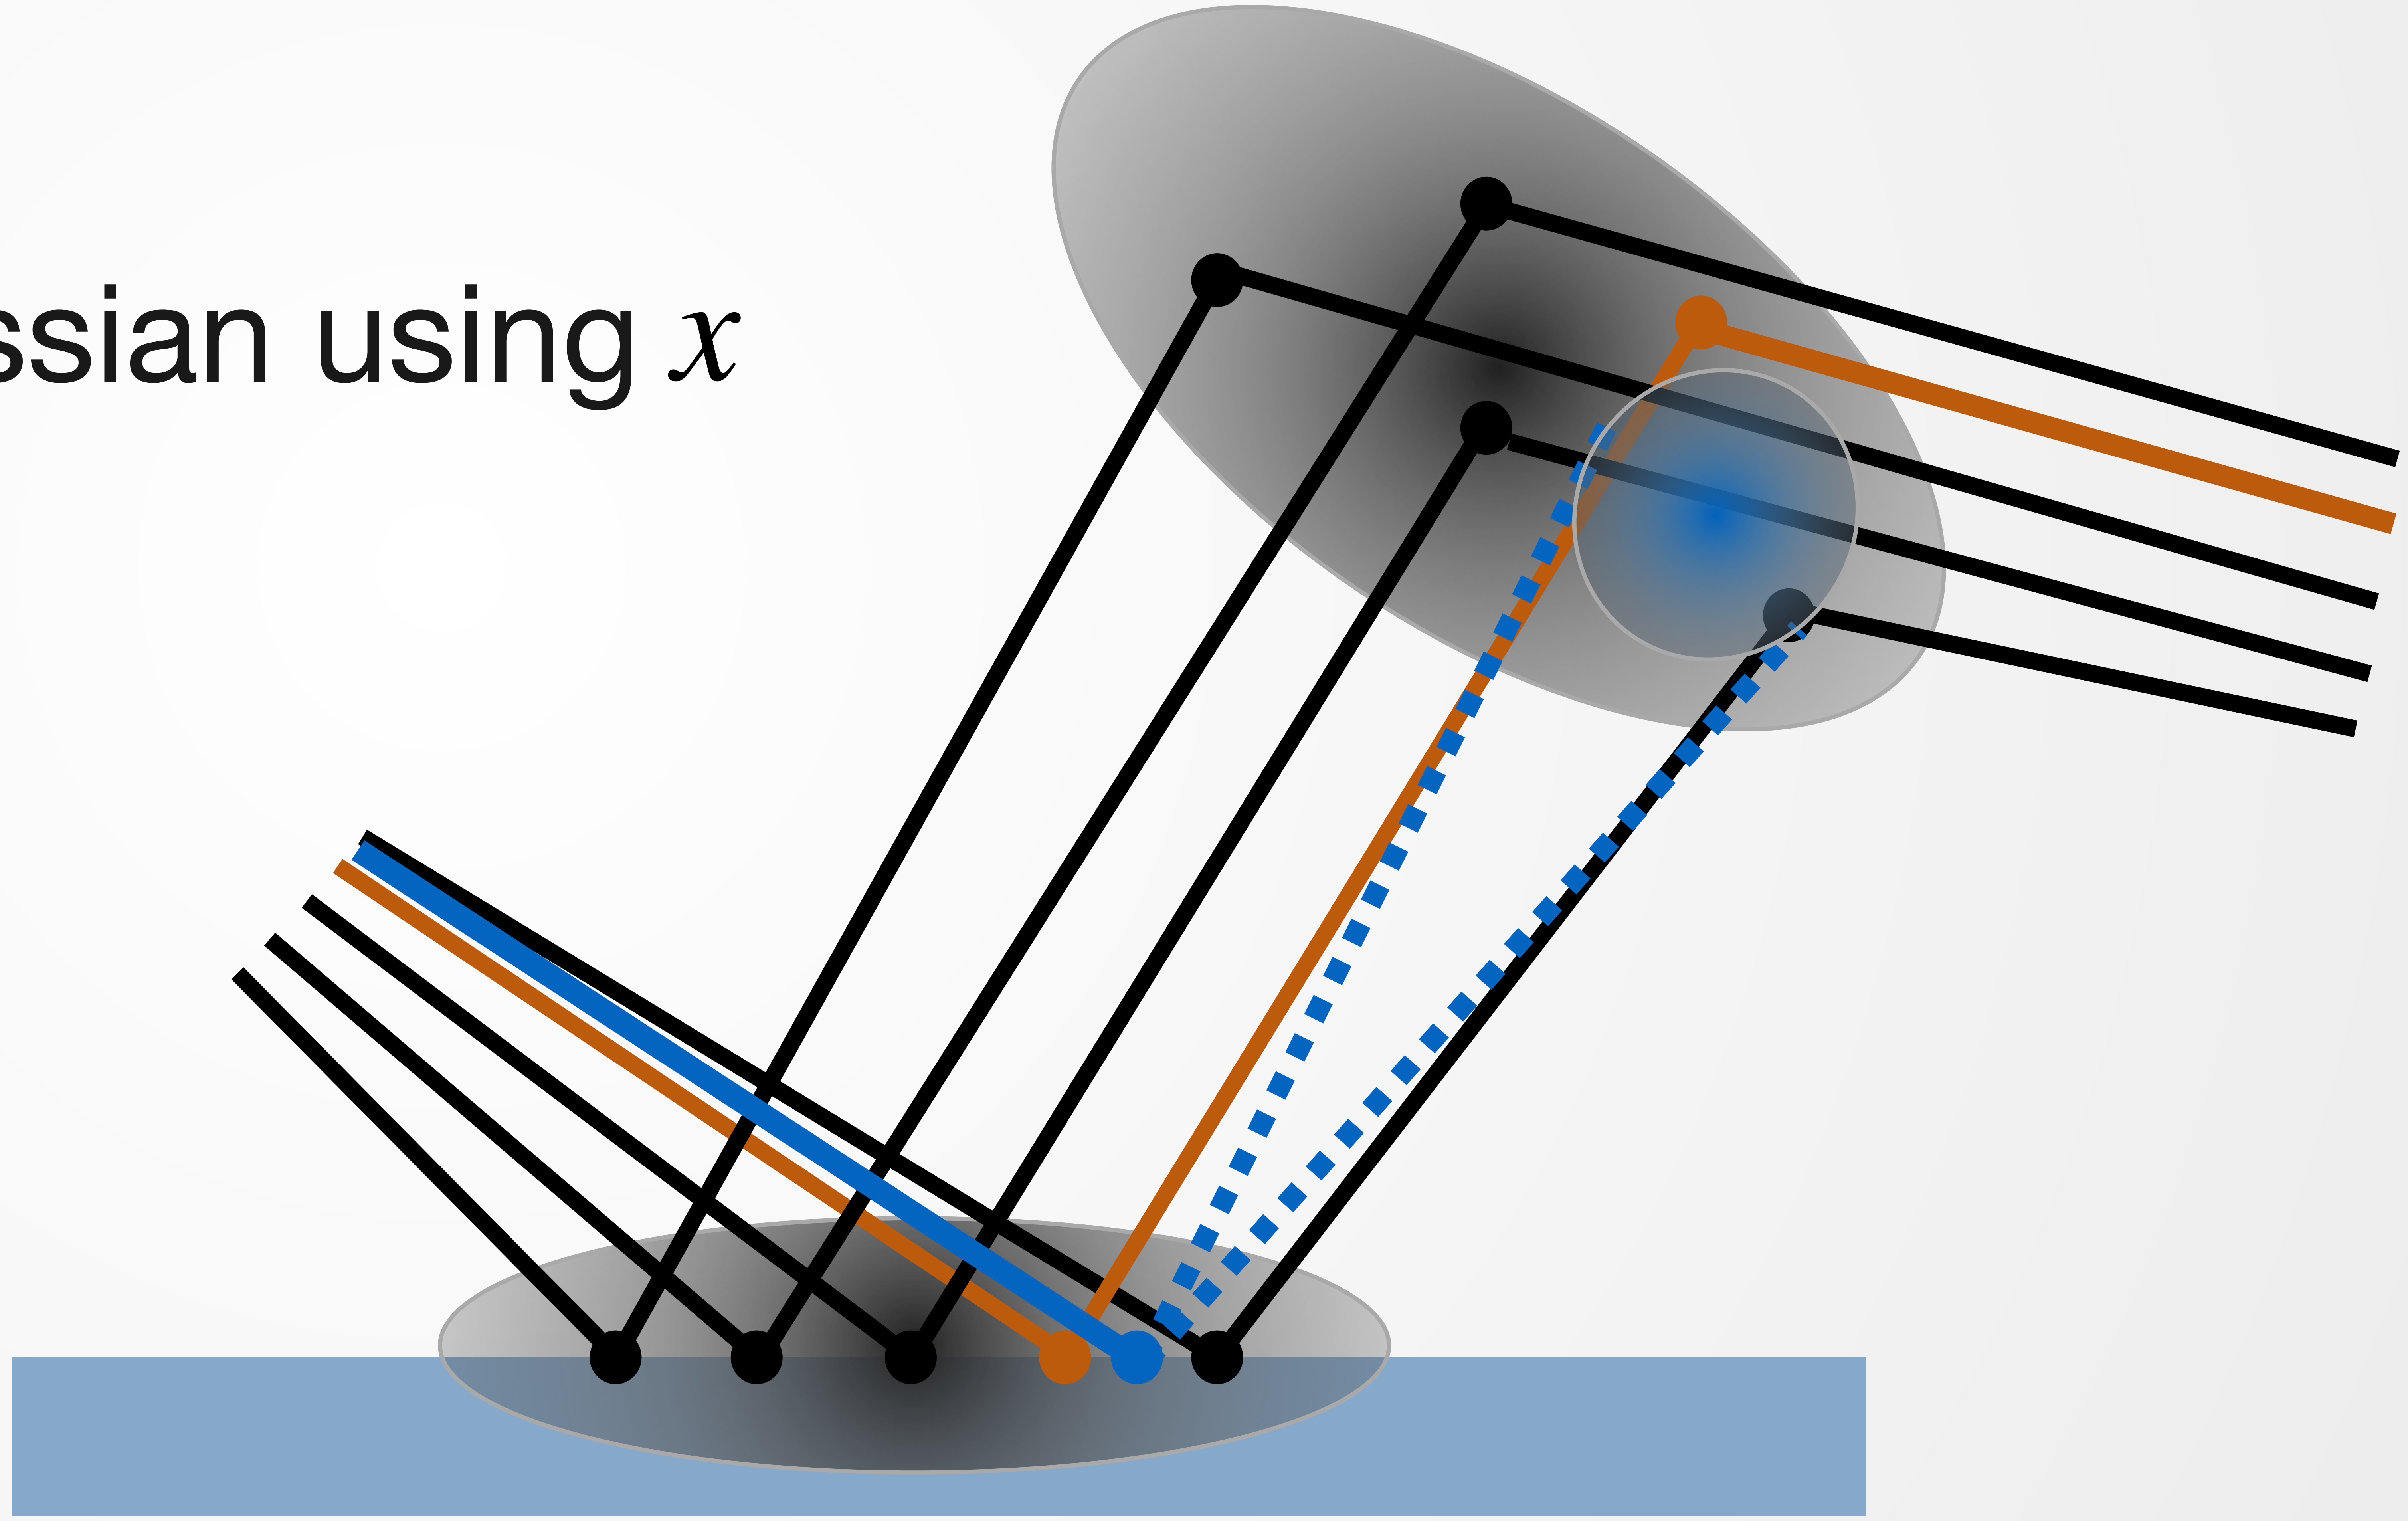

Guide path

Nearest neighbours

New path

Sampling next vertex

Next vertex in volume | Next vertex on surface

Sampling of Gaussian

Sampling of BSDF

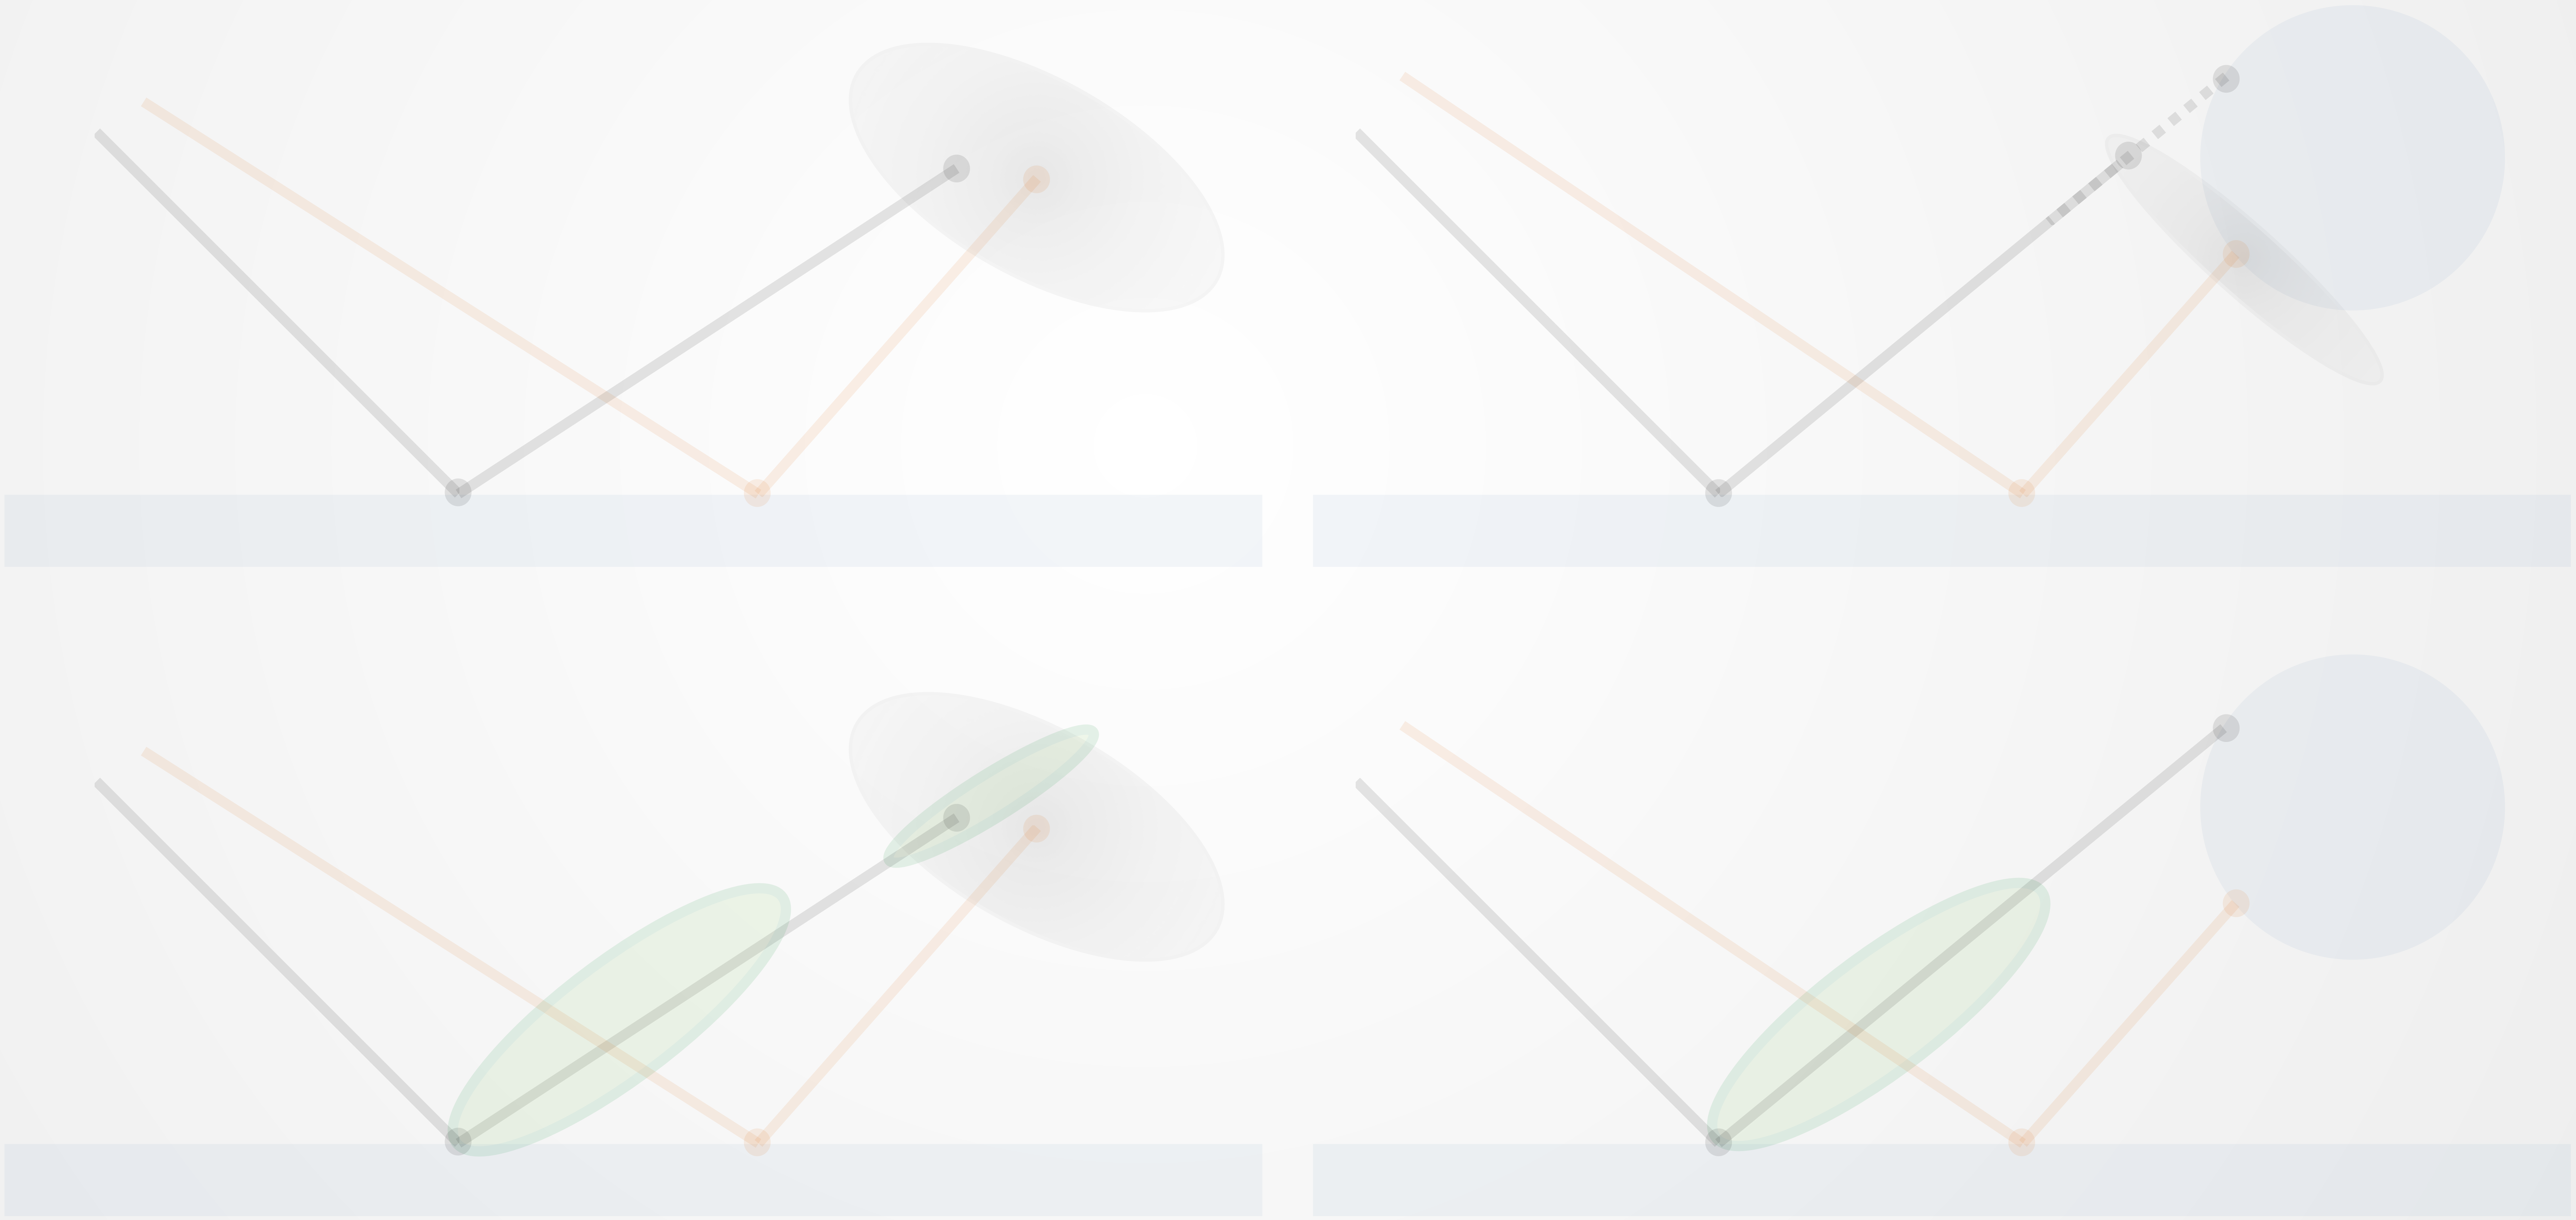

Sampling next vertex

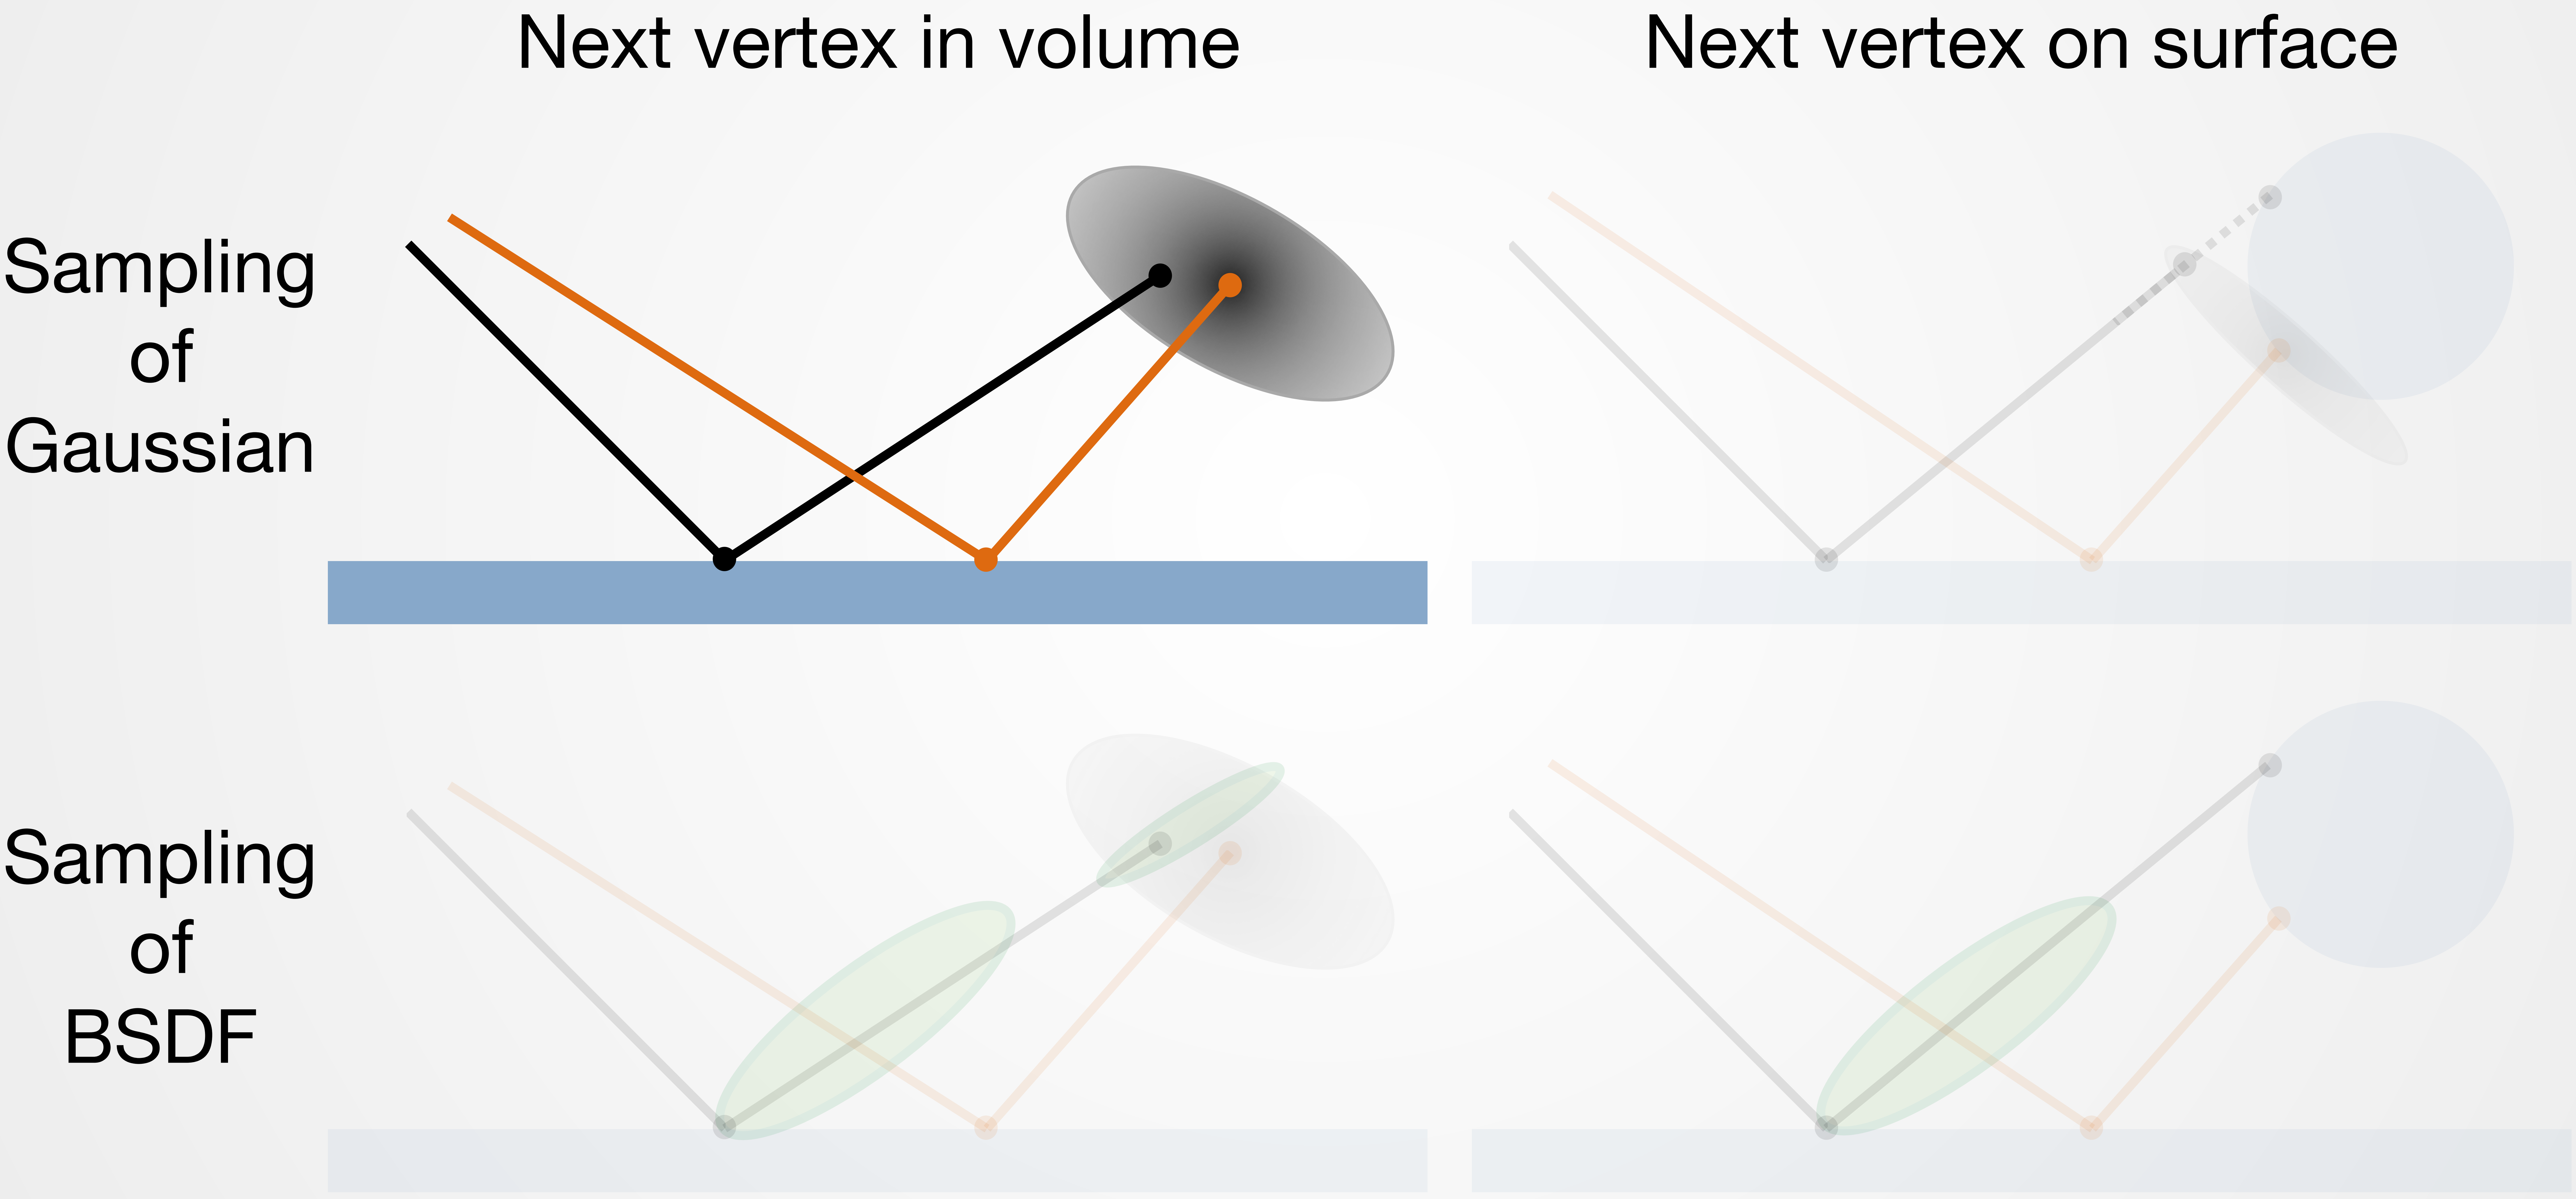

Sampling next vertex

Next vertex in volume

Next vertex on surface

Sampling of Gaussian

Sampling of BSDF

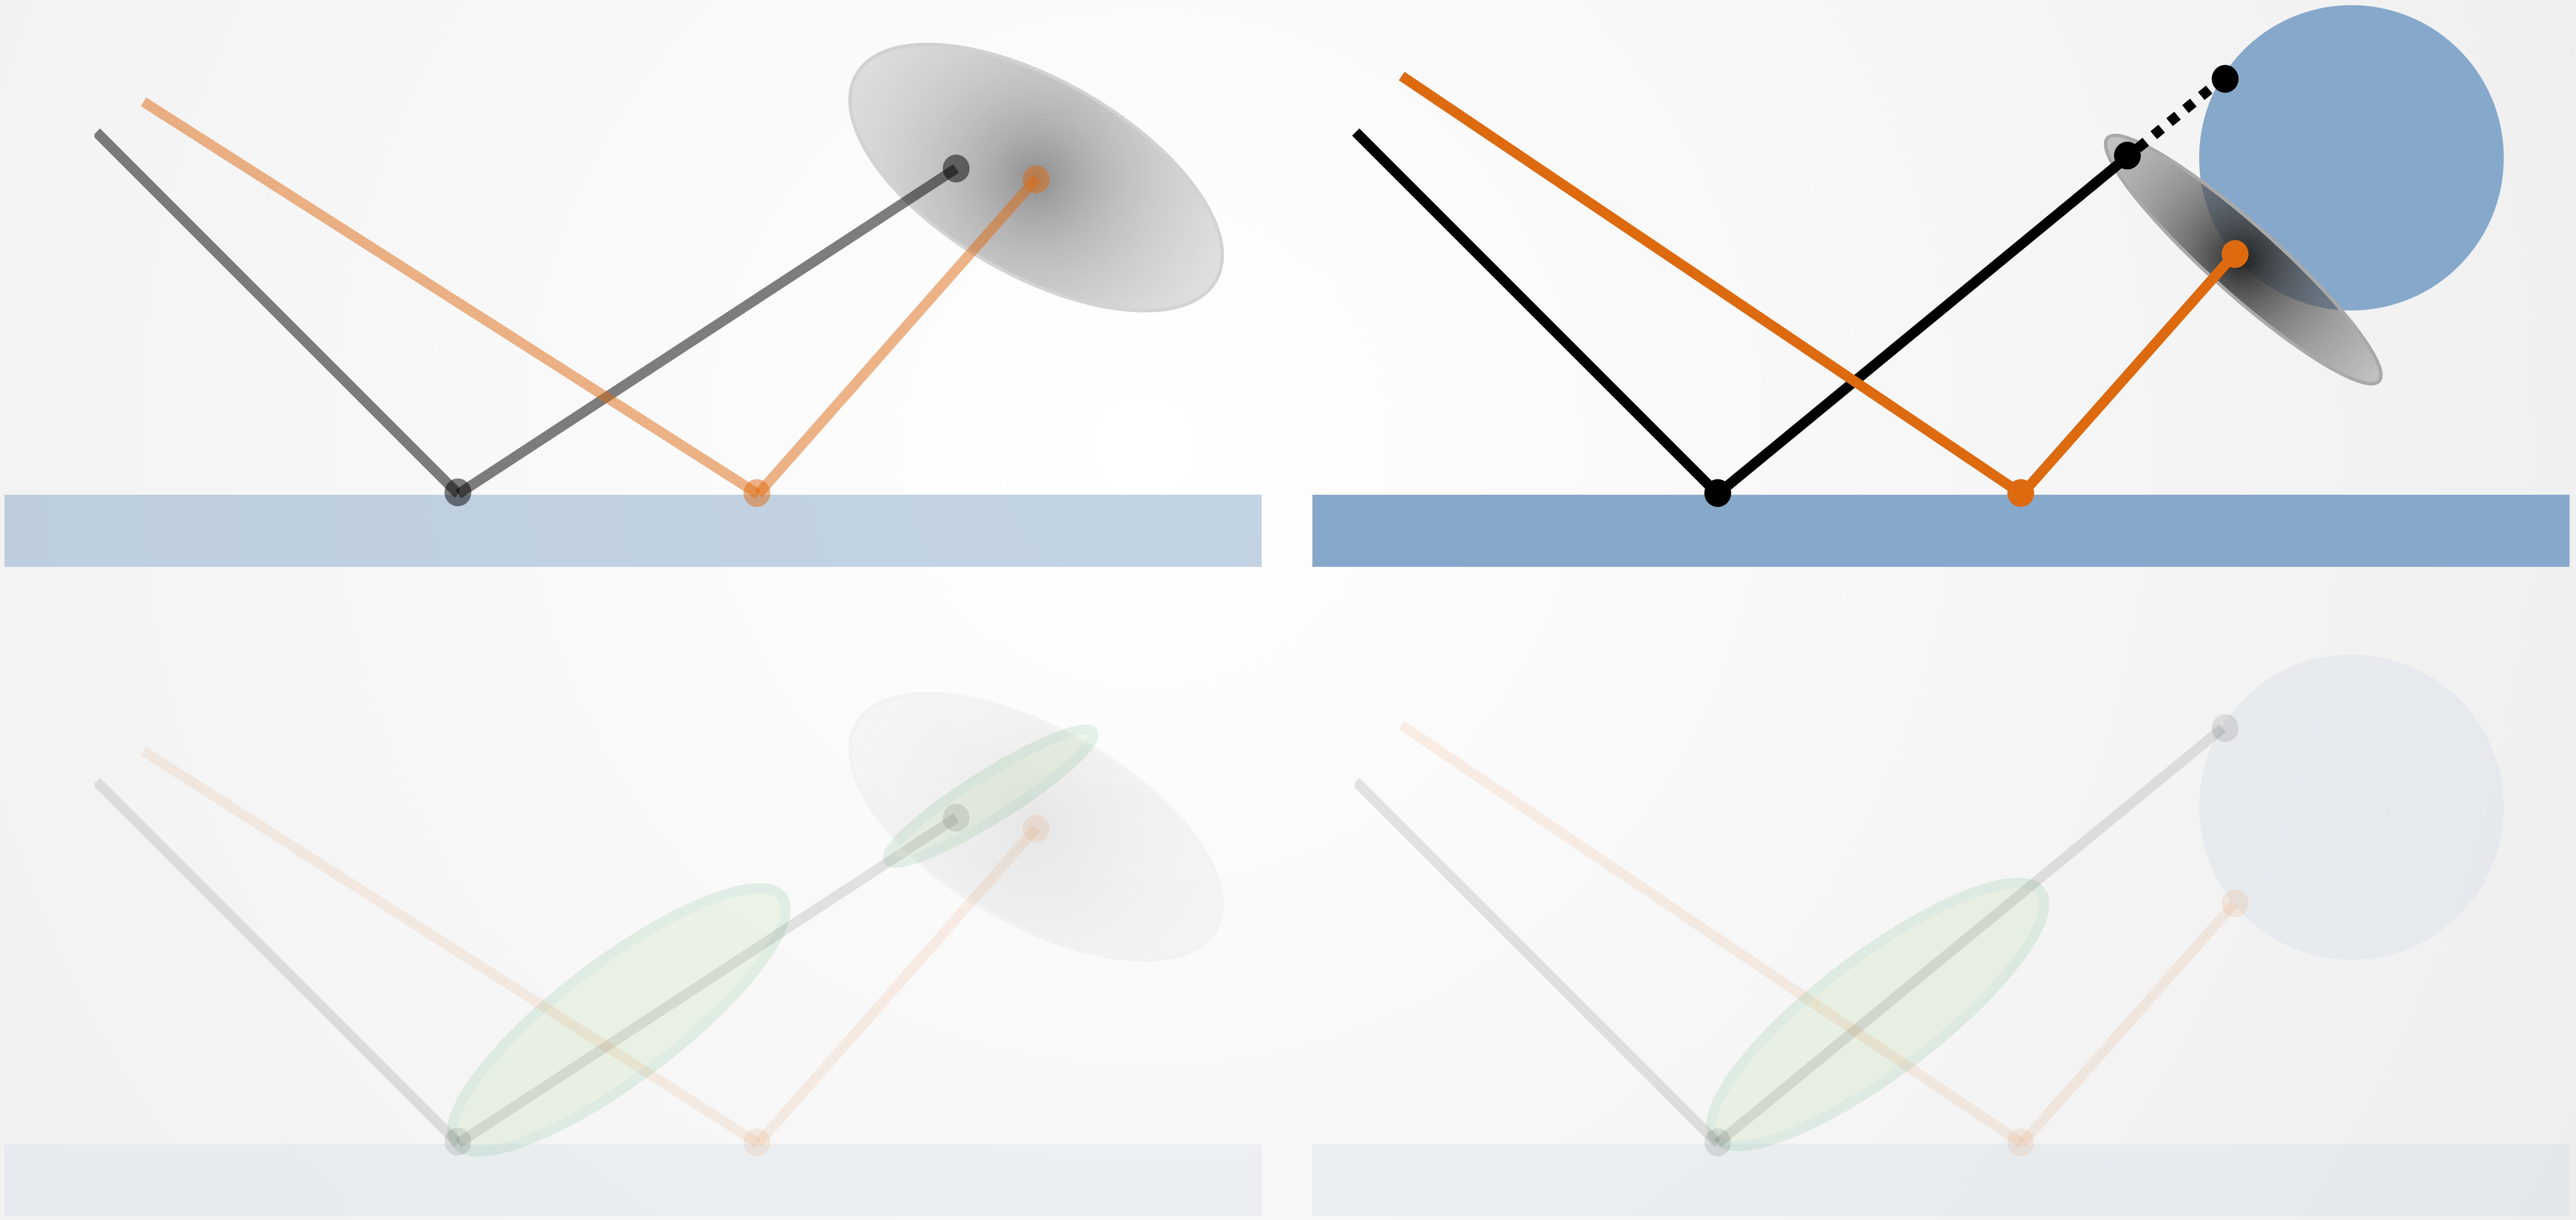

Sampling next vertex

Next vertex in volume | Next vertex on surface

Sampling of Gaussian

Sampling of BSDF

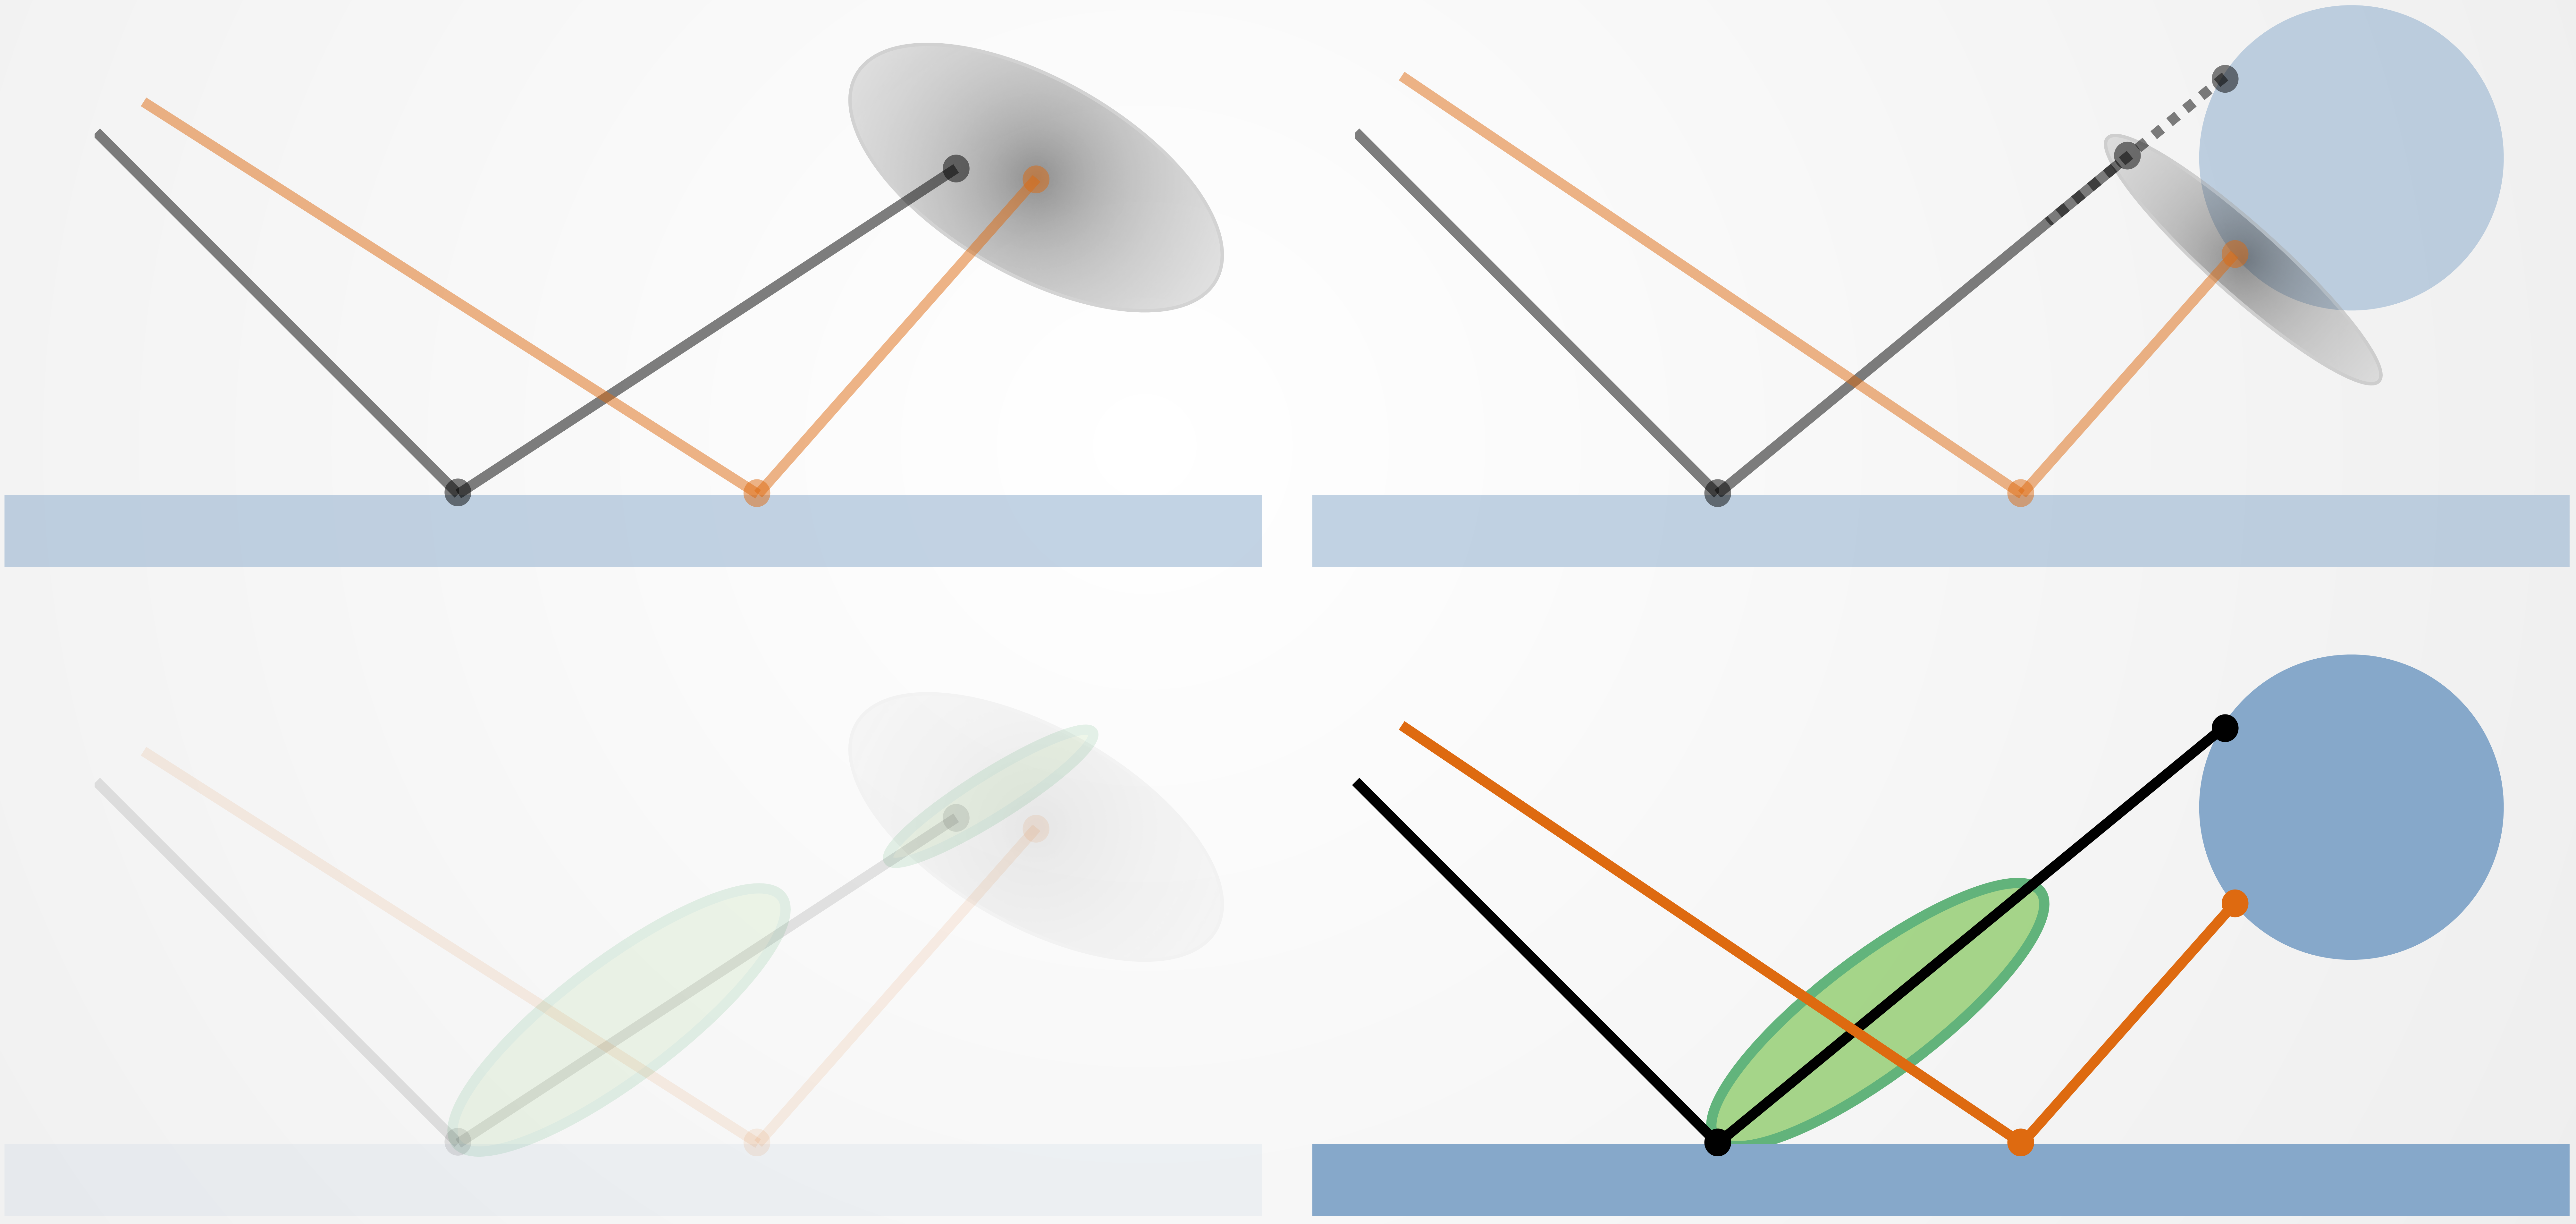

Sampling next vertex

Next vertex in volume

Next vertex on surface

Sampling of Gaussian

Sampling of BSDF

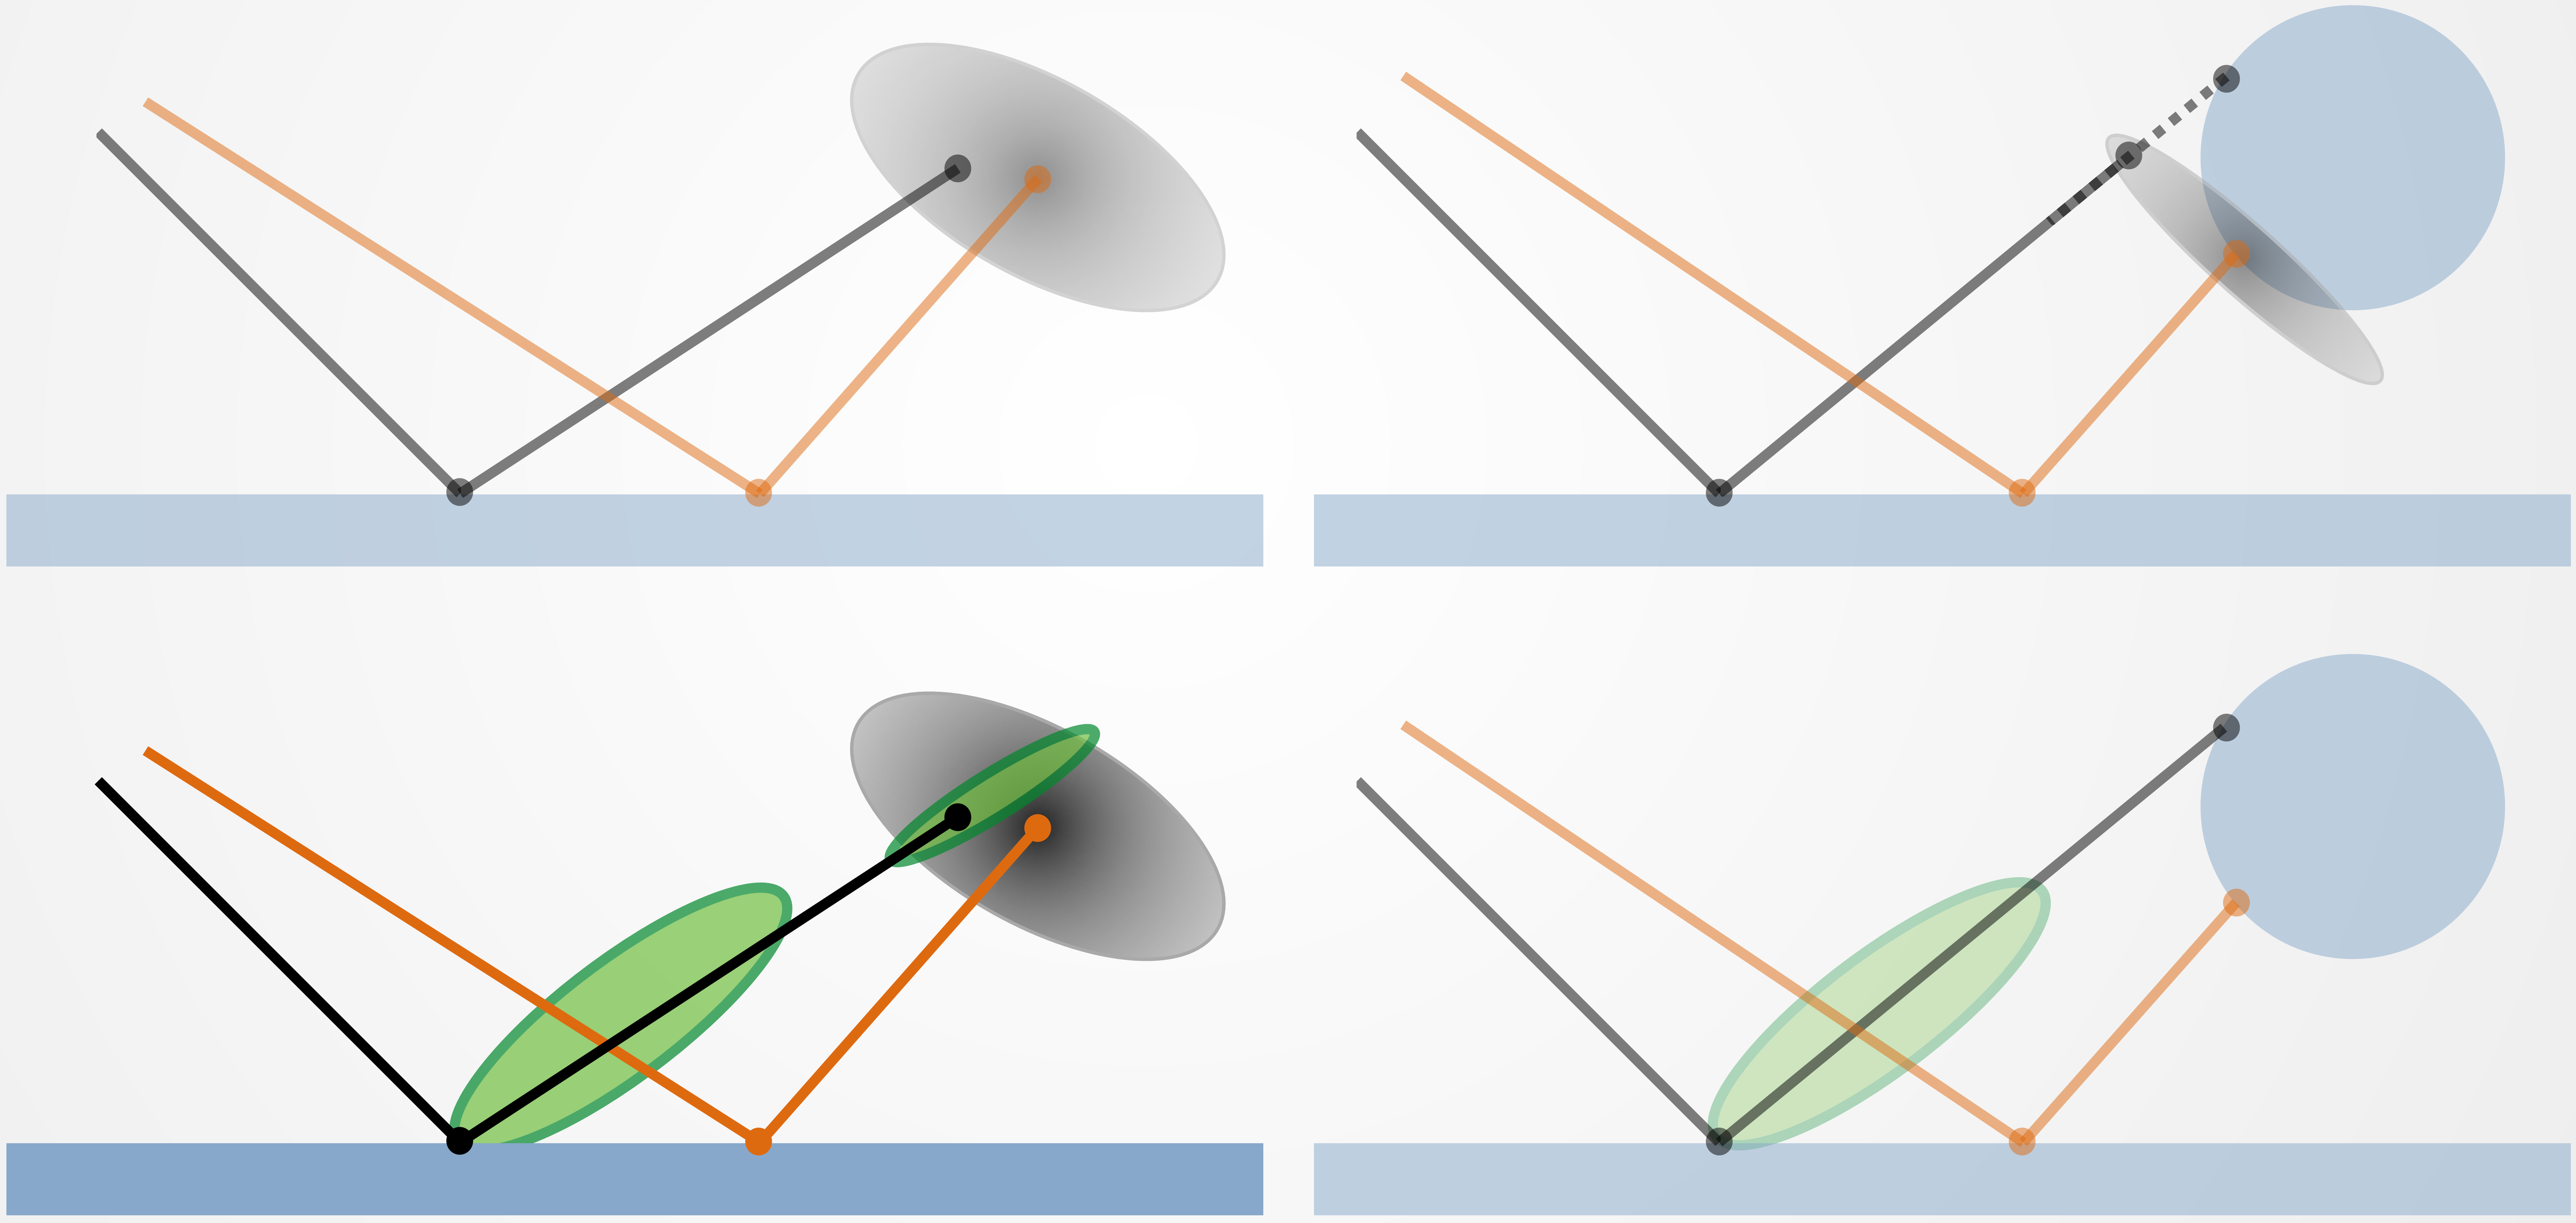

Guided PDF

- 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

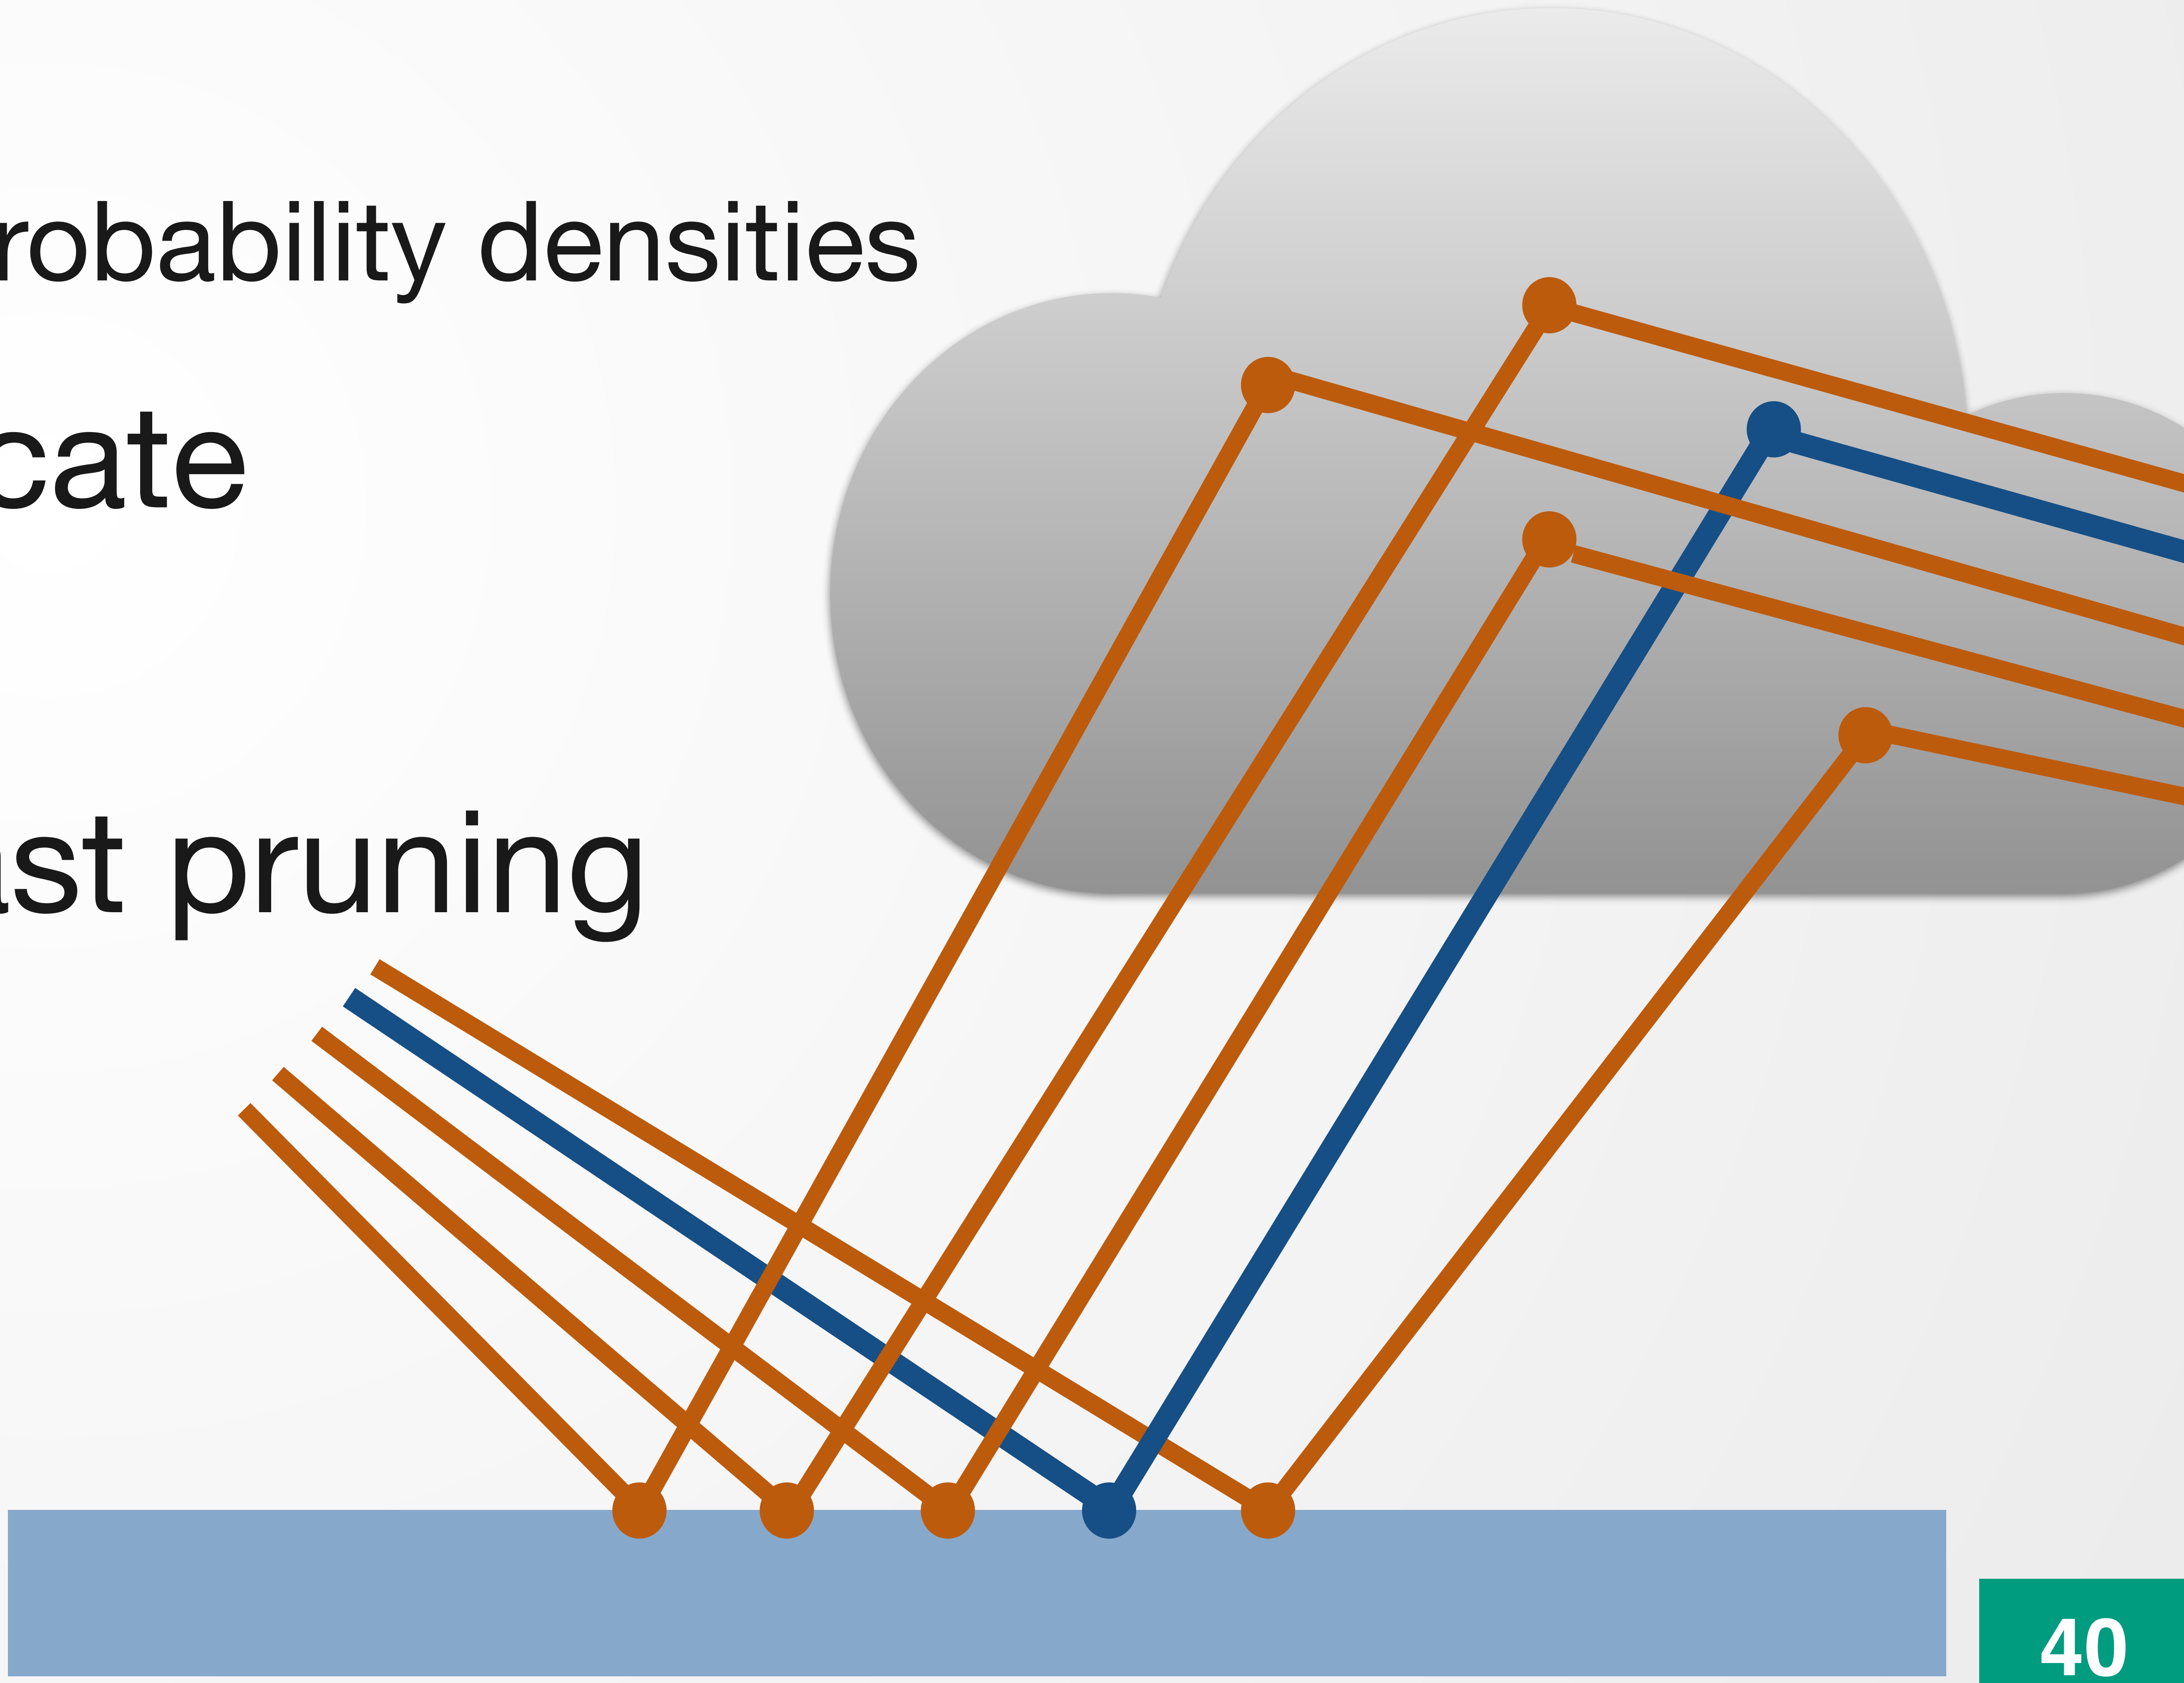

Guide paths X_l

Path X

Selecting Guide Paths

- Outliers $\neq$ samples with high contribution
- Outliers classification: Density based outlier rejection (DBOR, Zirr et al. [2018])

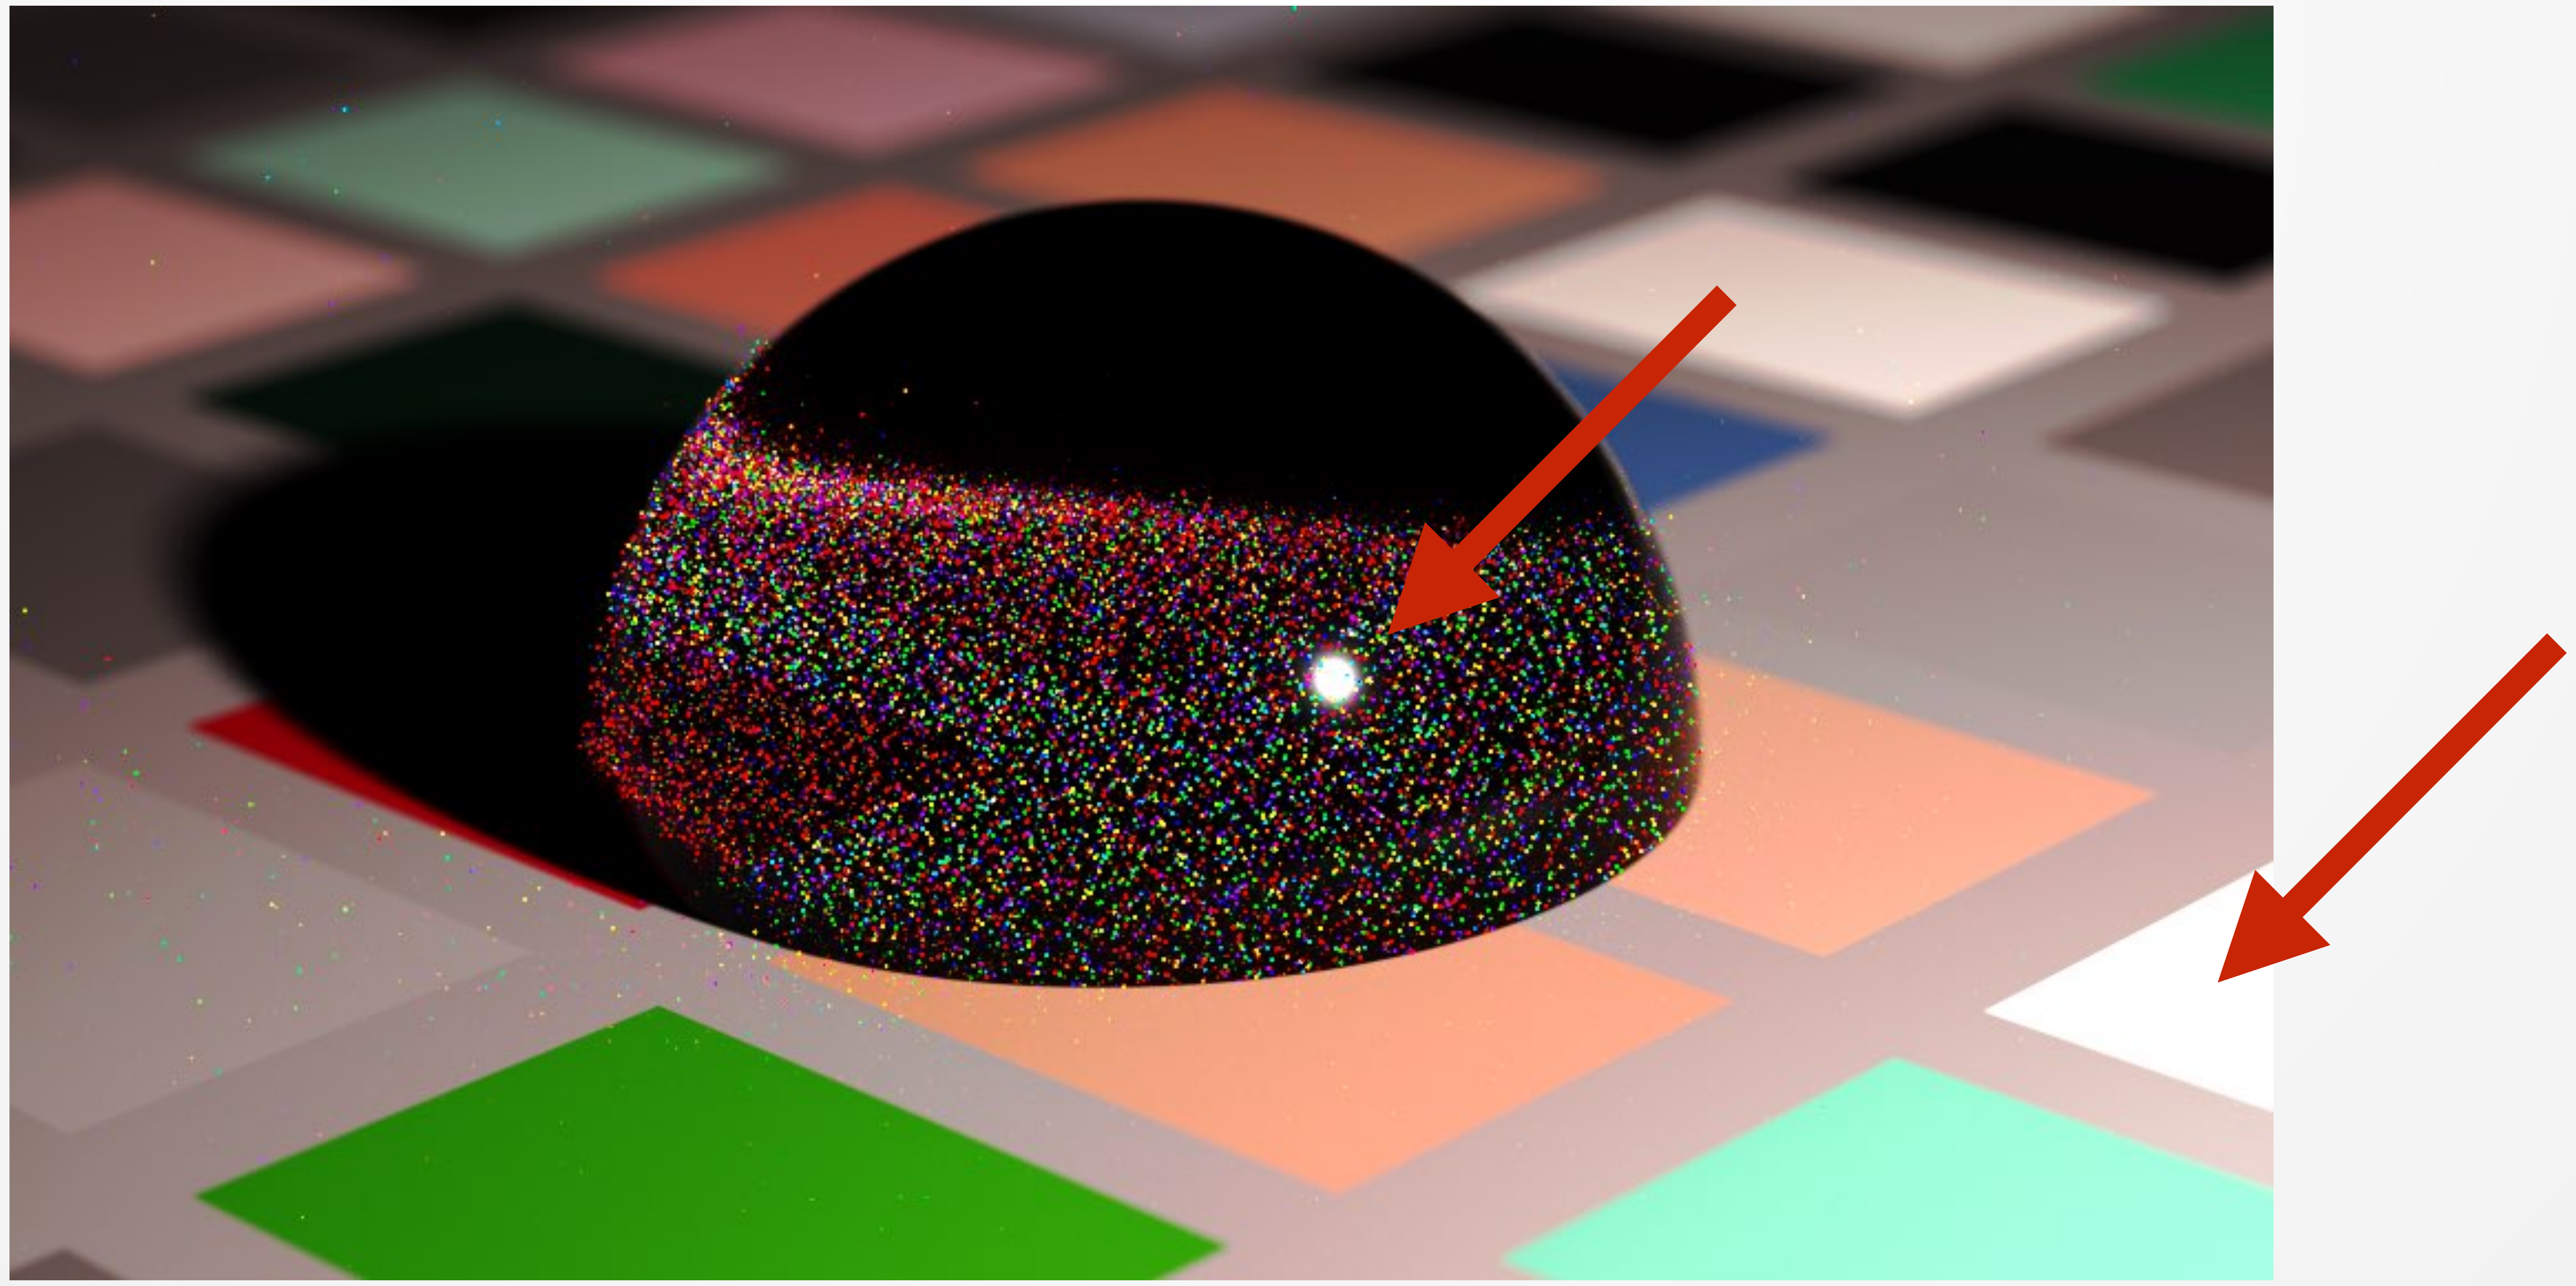

Remaining outliers

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

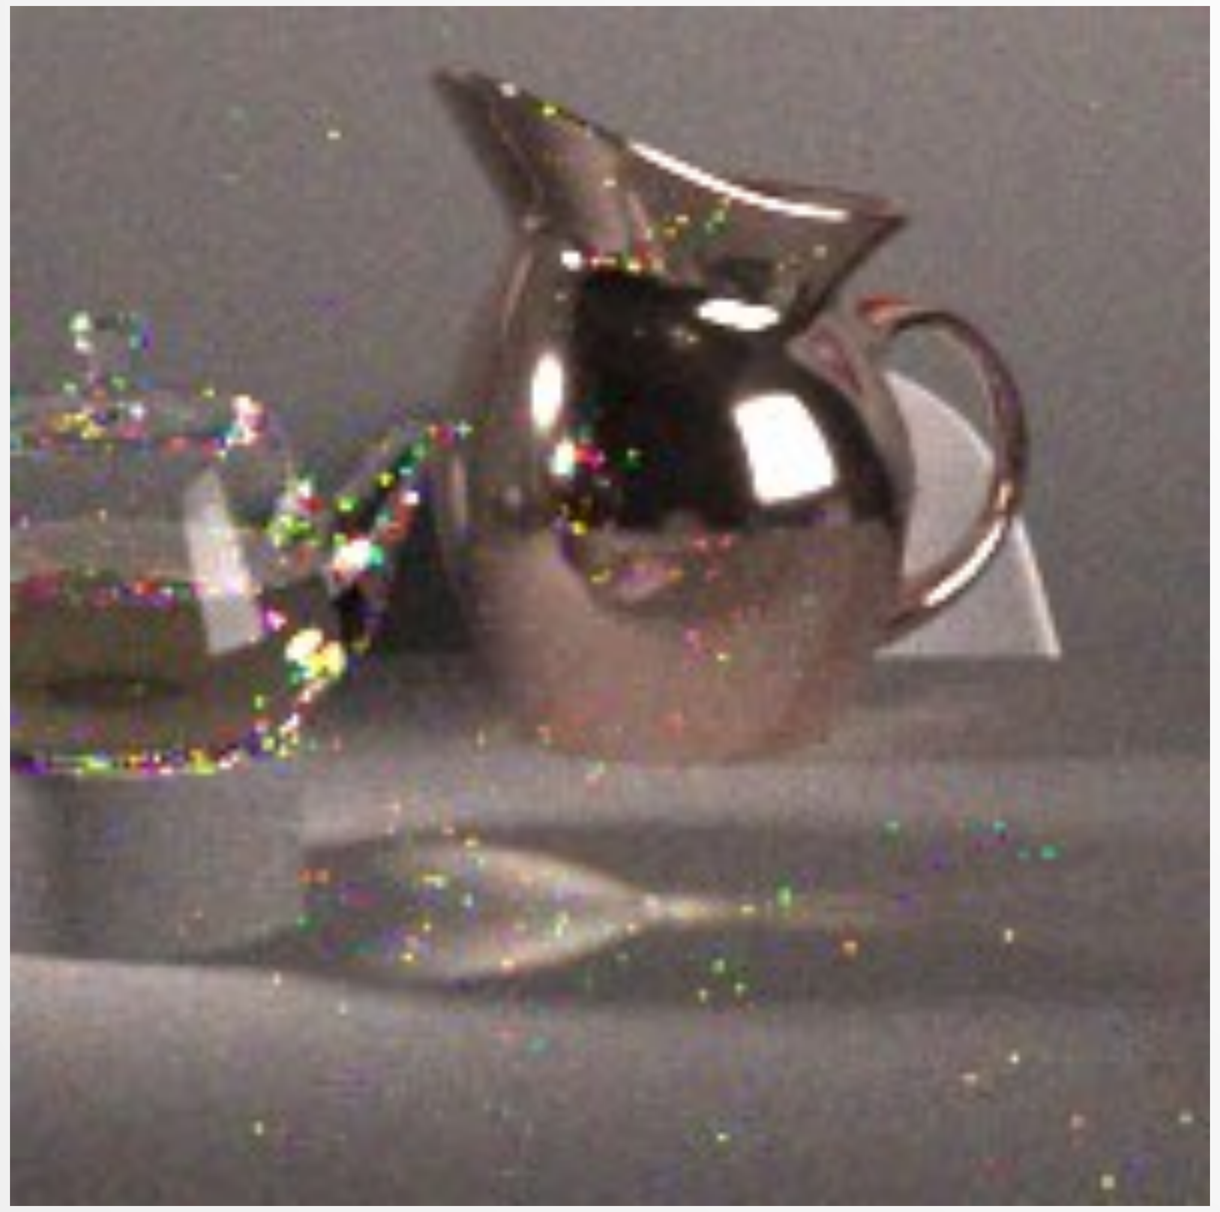

Outliers
removed

Pool

Reference

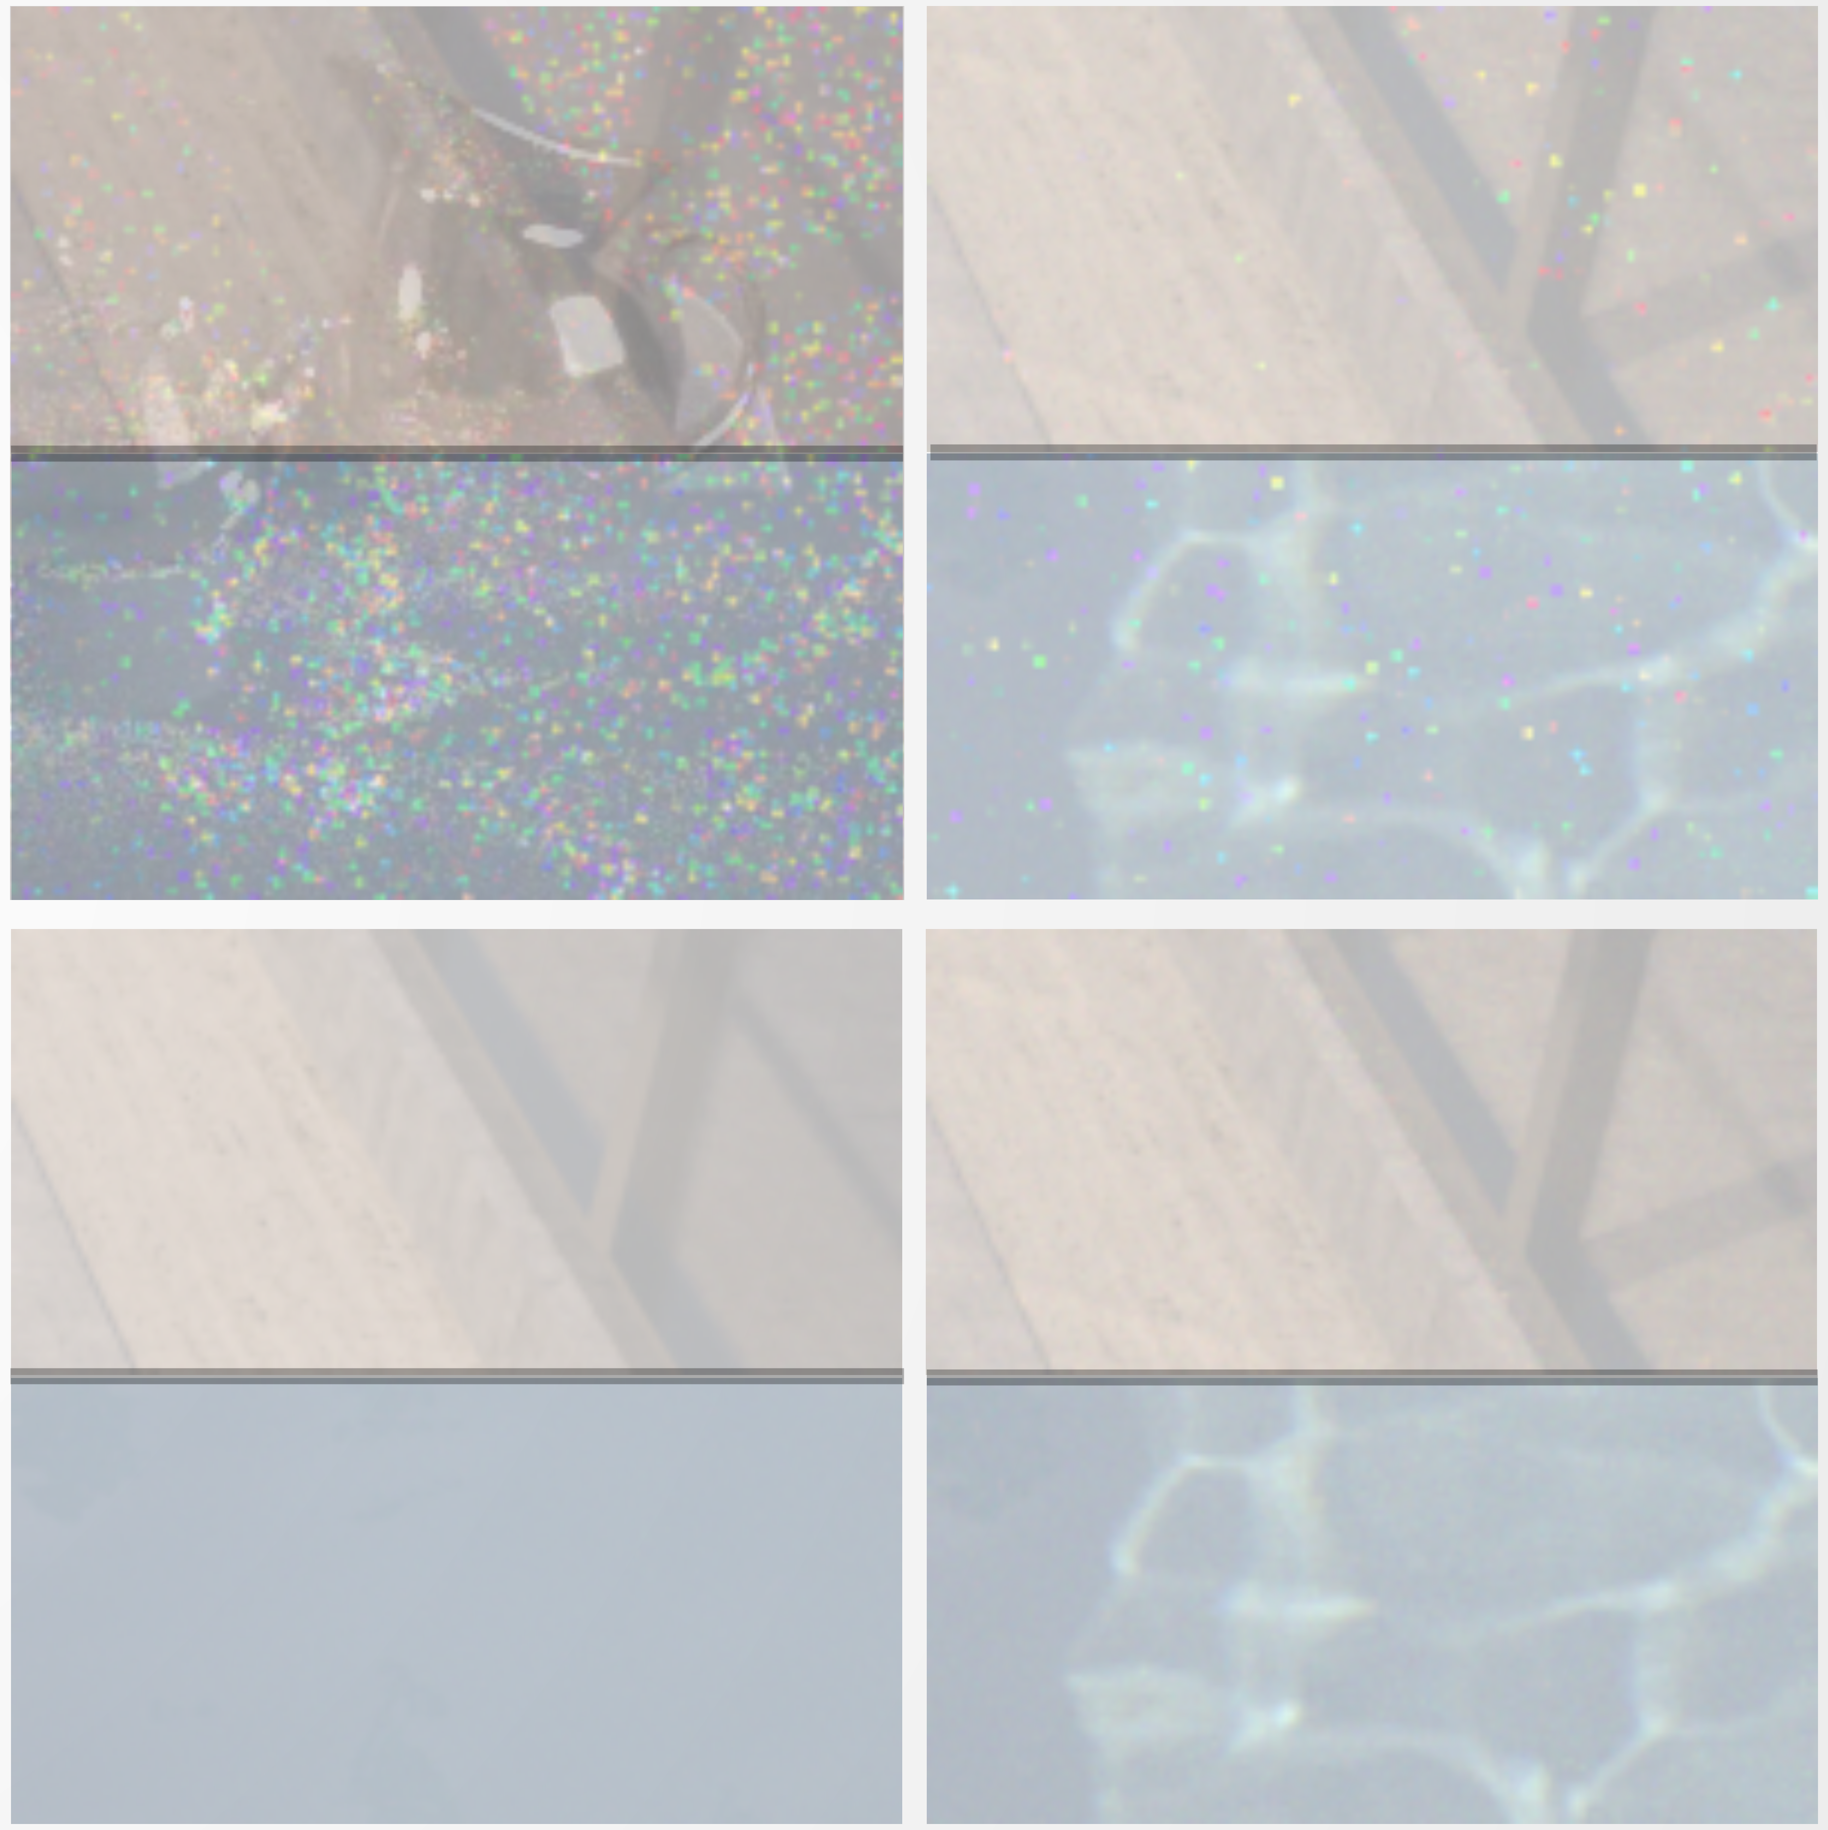

Pool - 30min

Path tracing

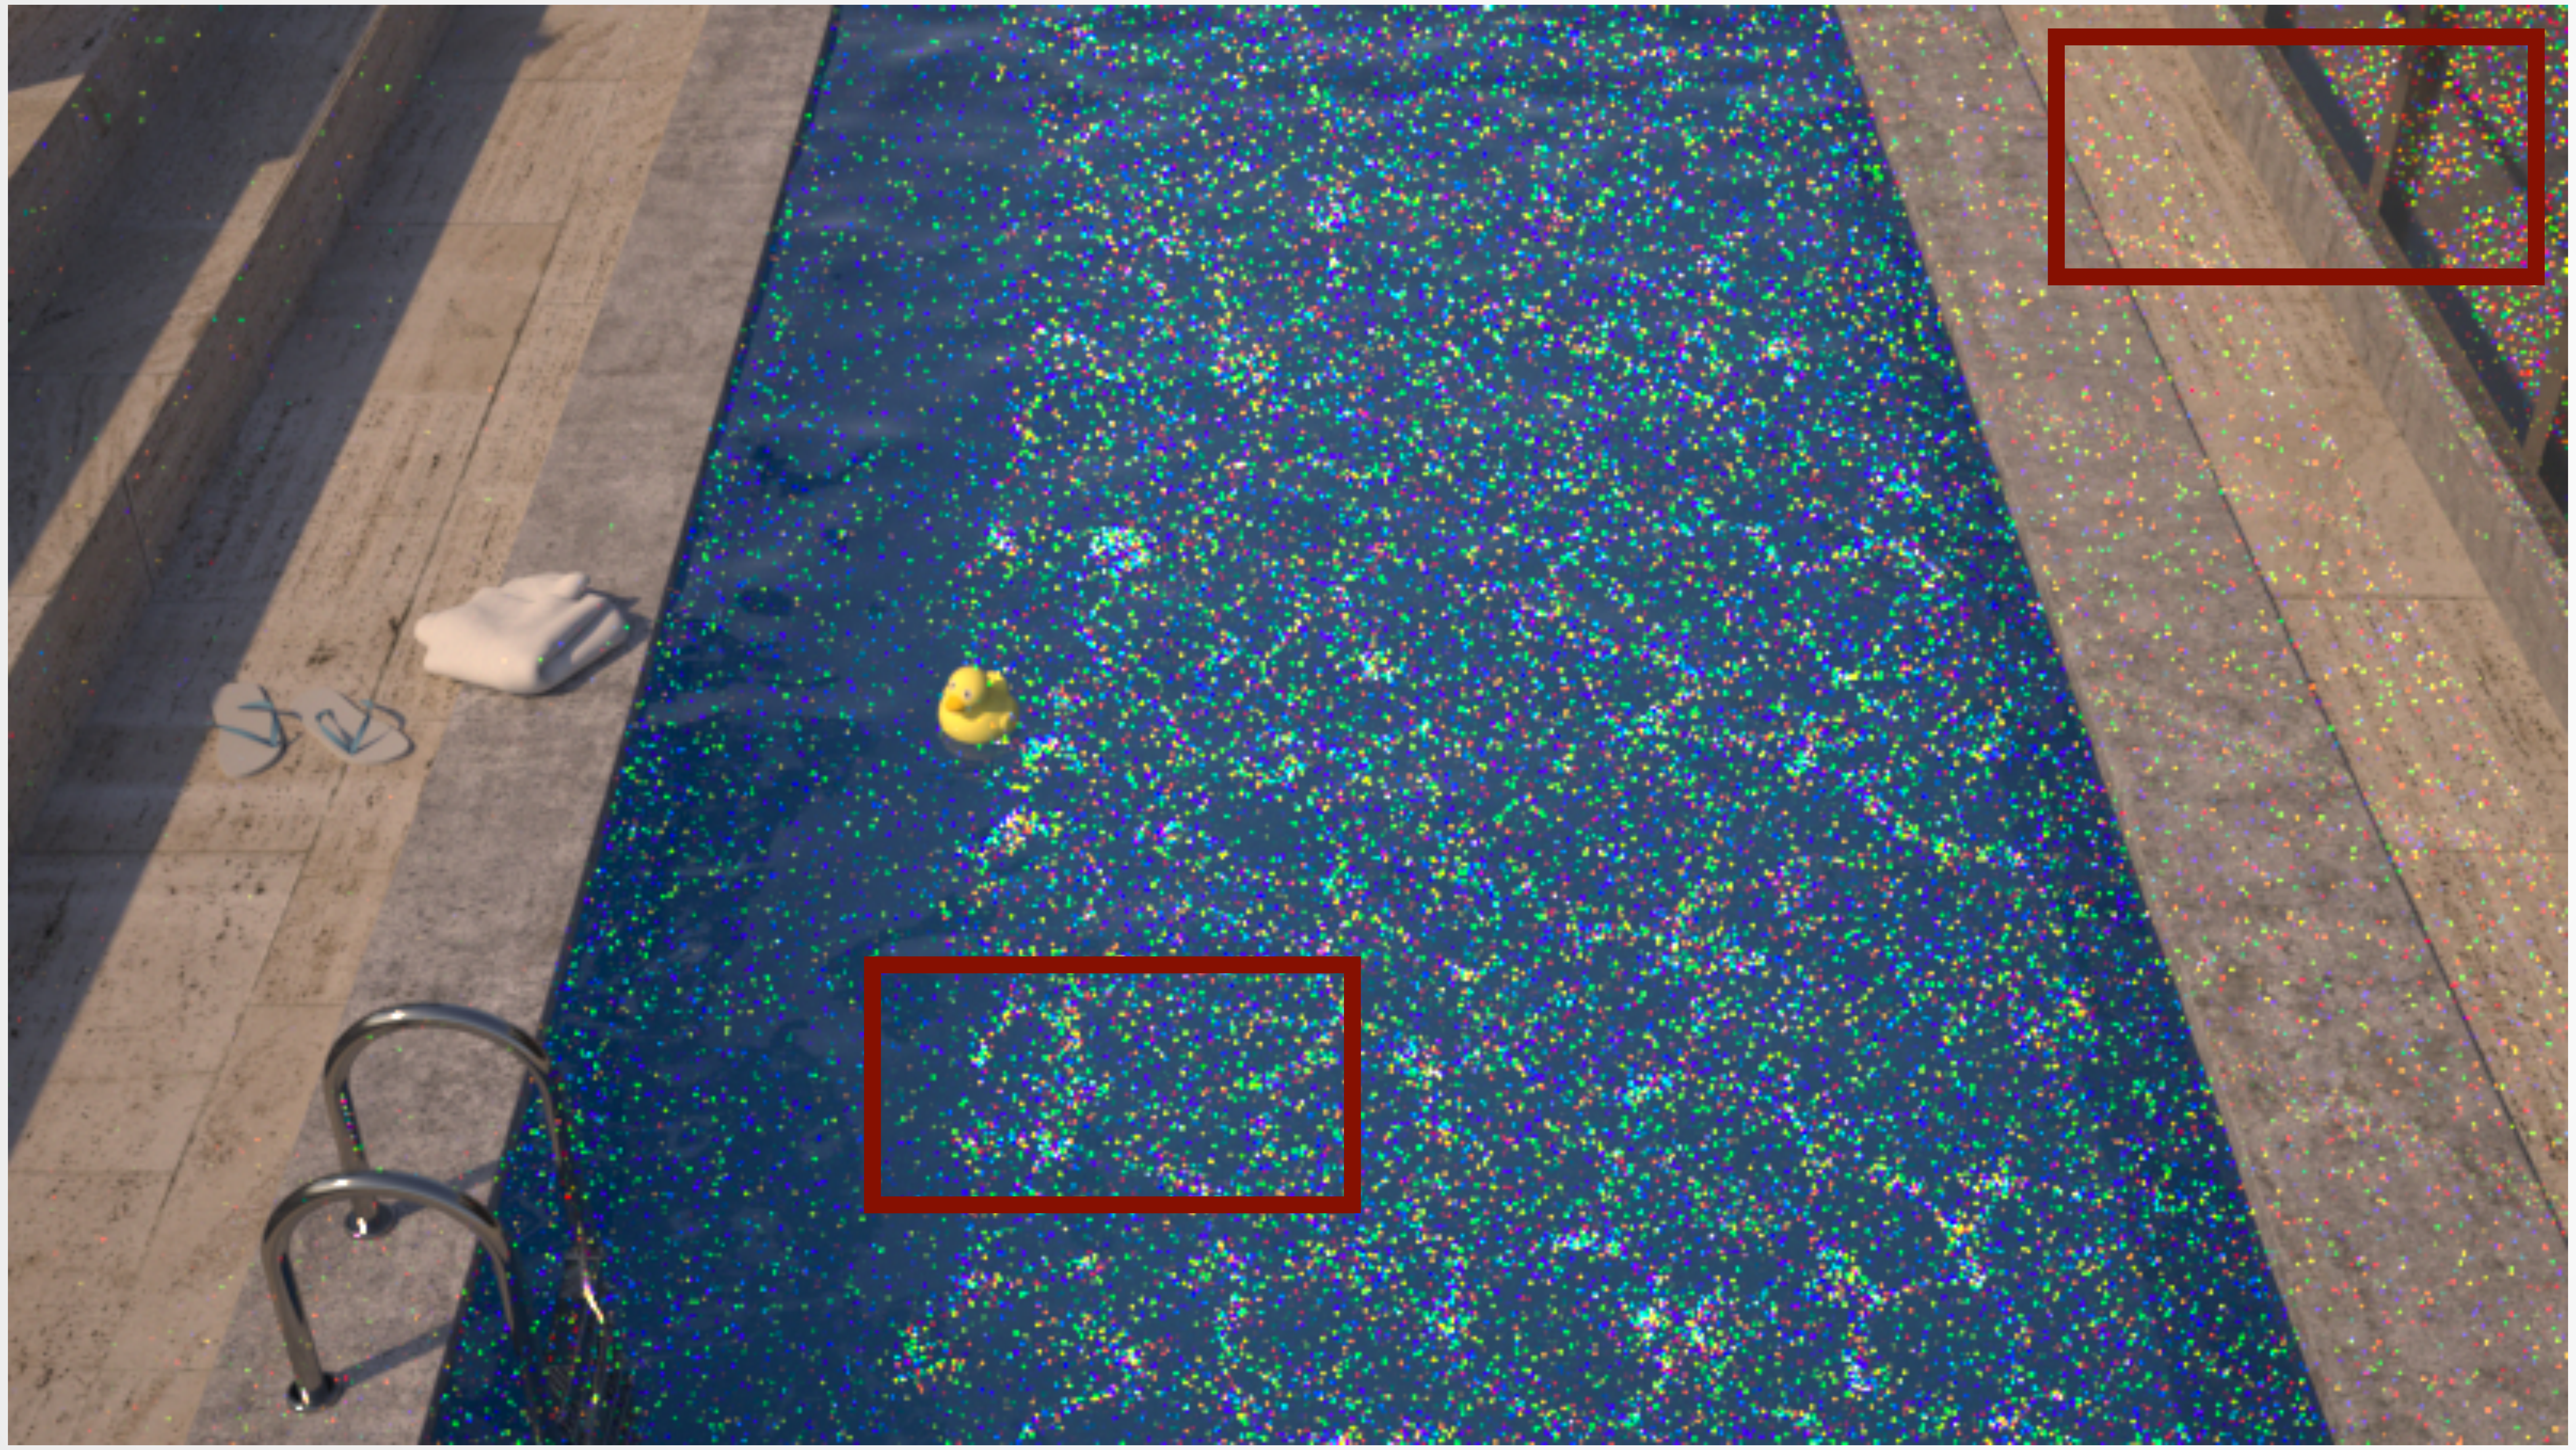

PT

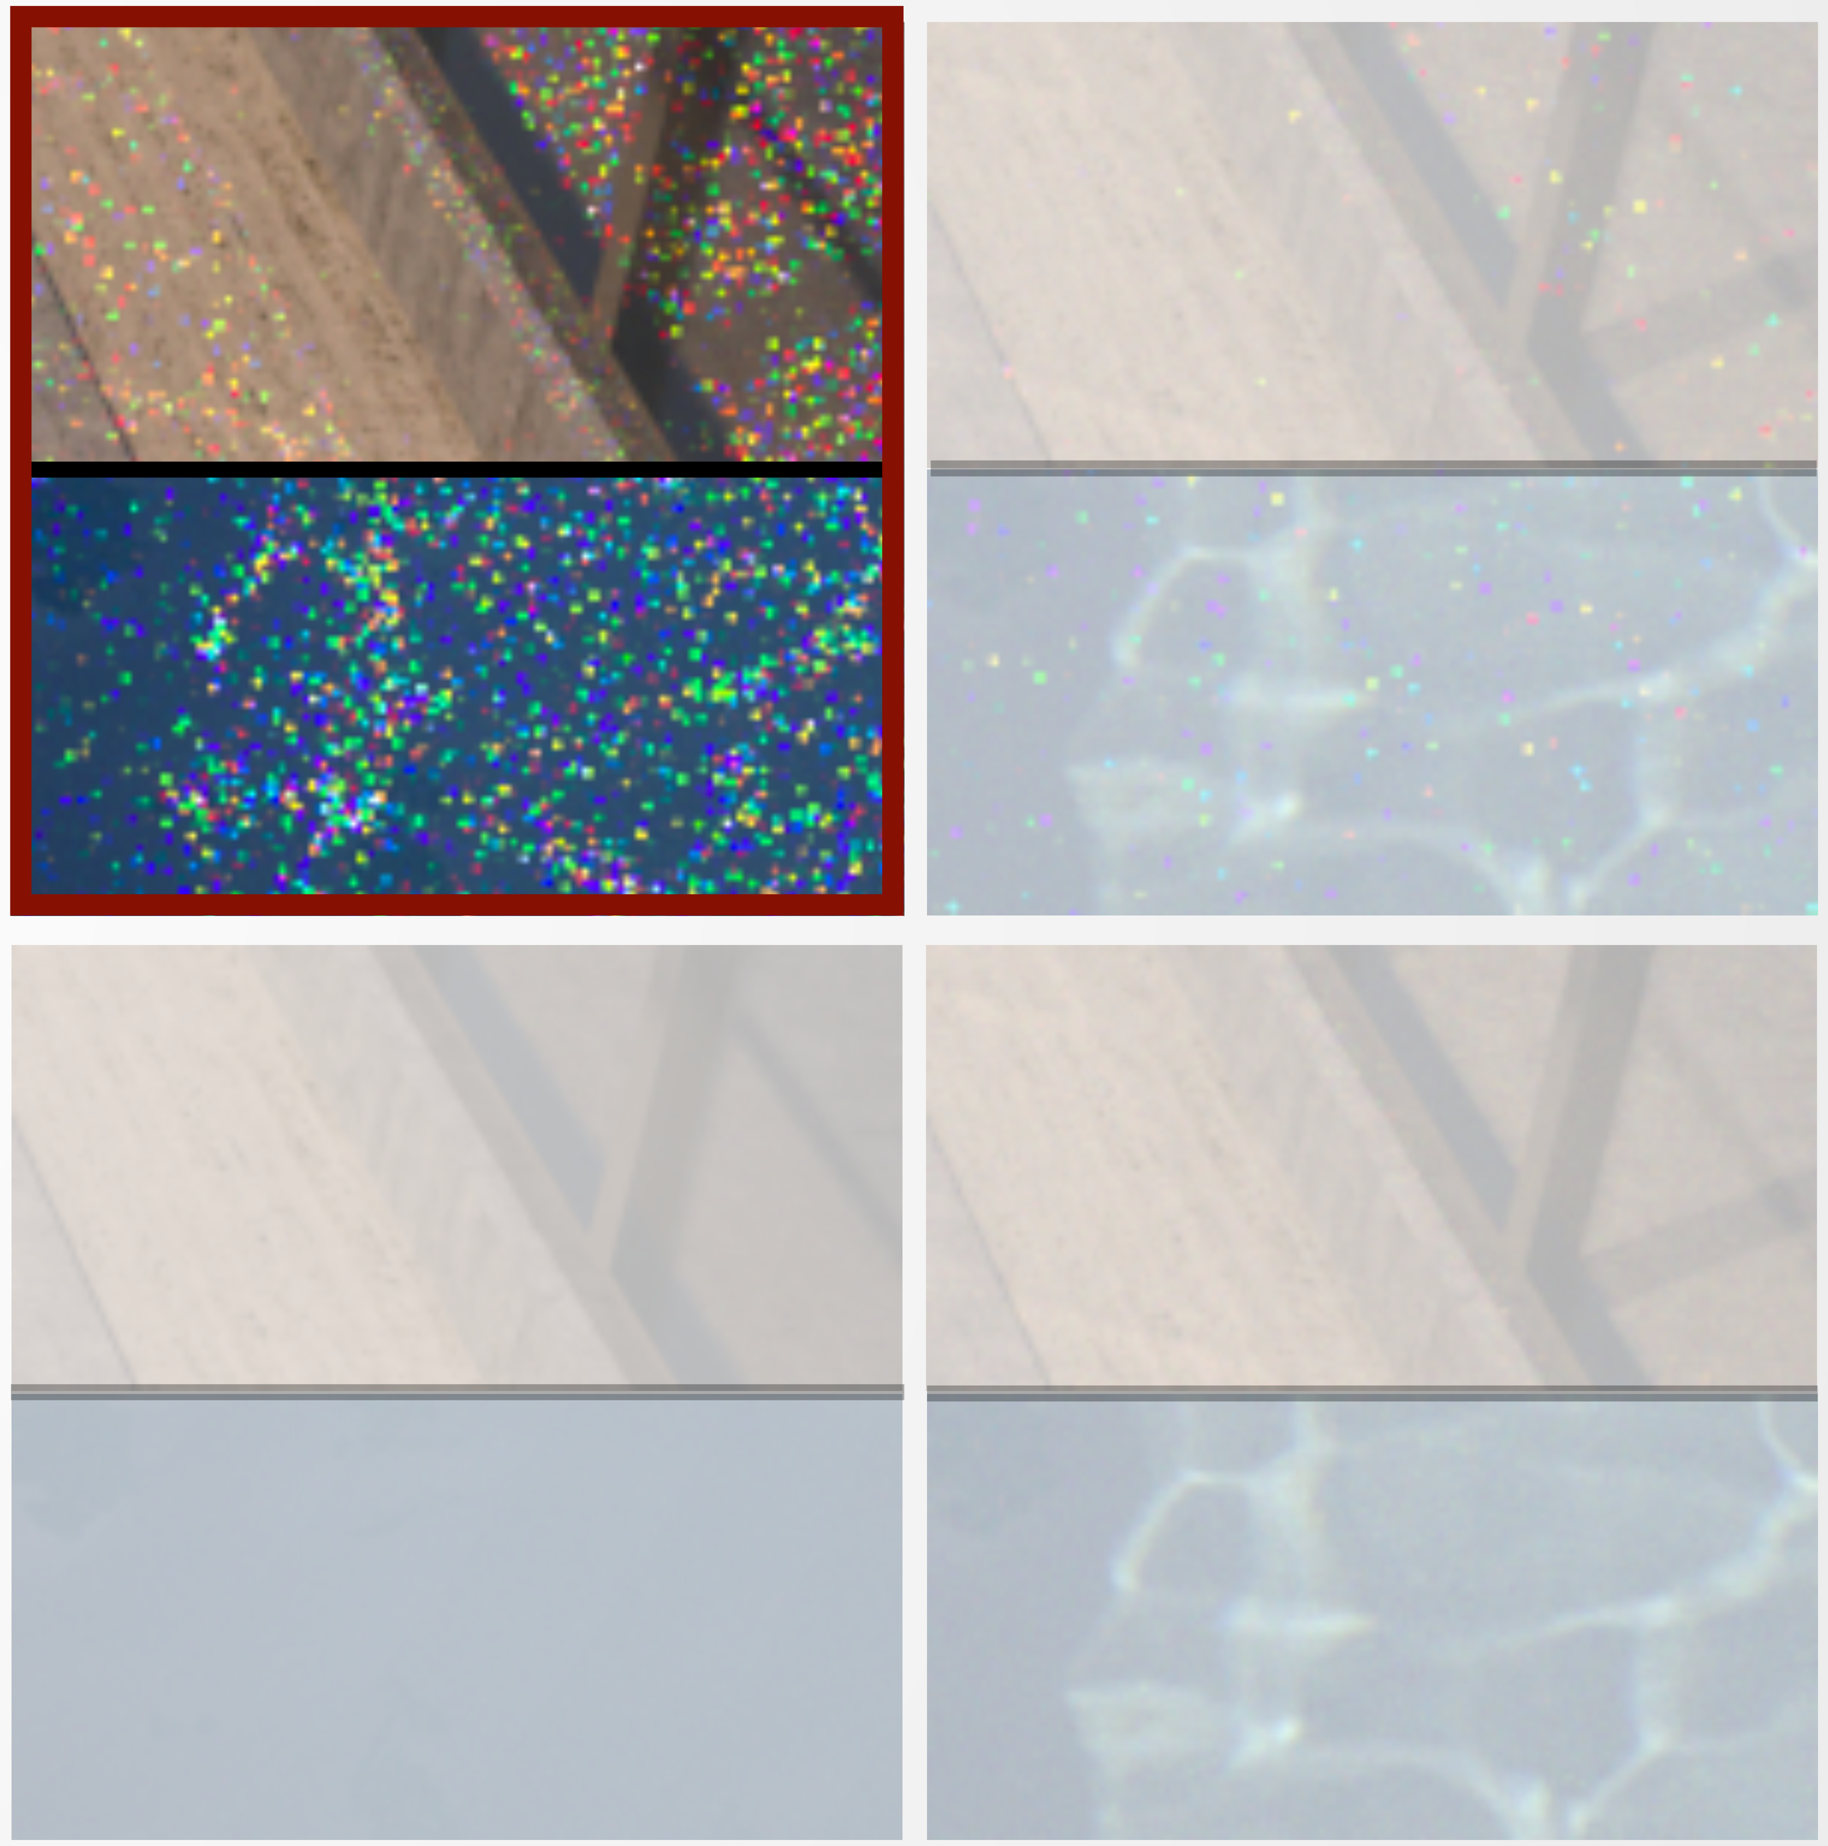

Pool - 30min

Guided path tracing

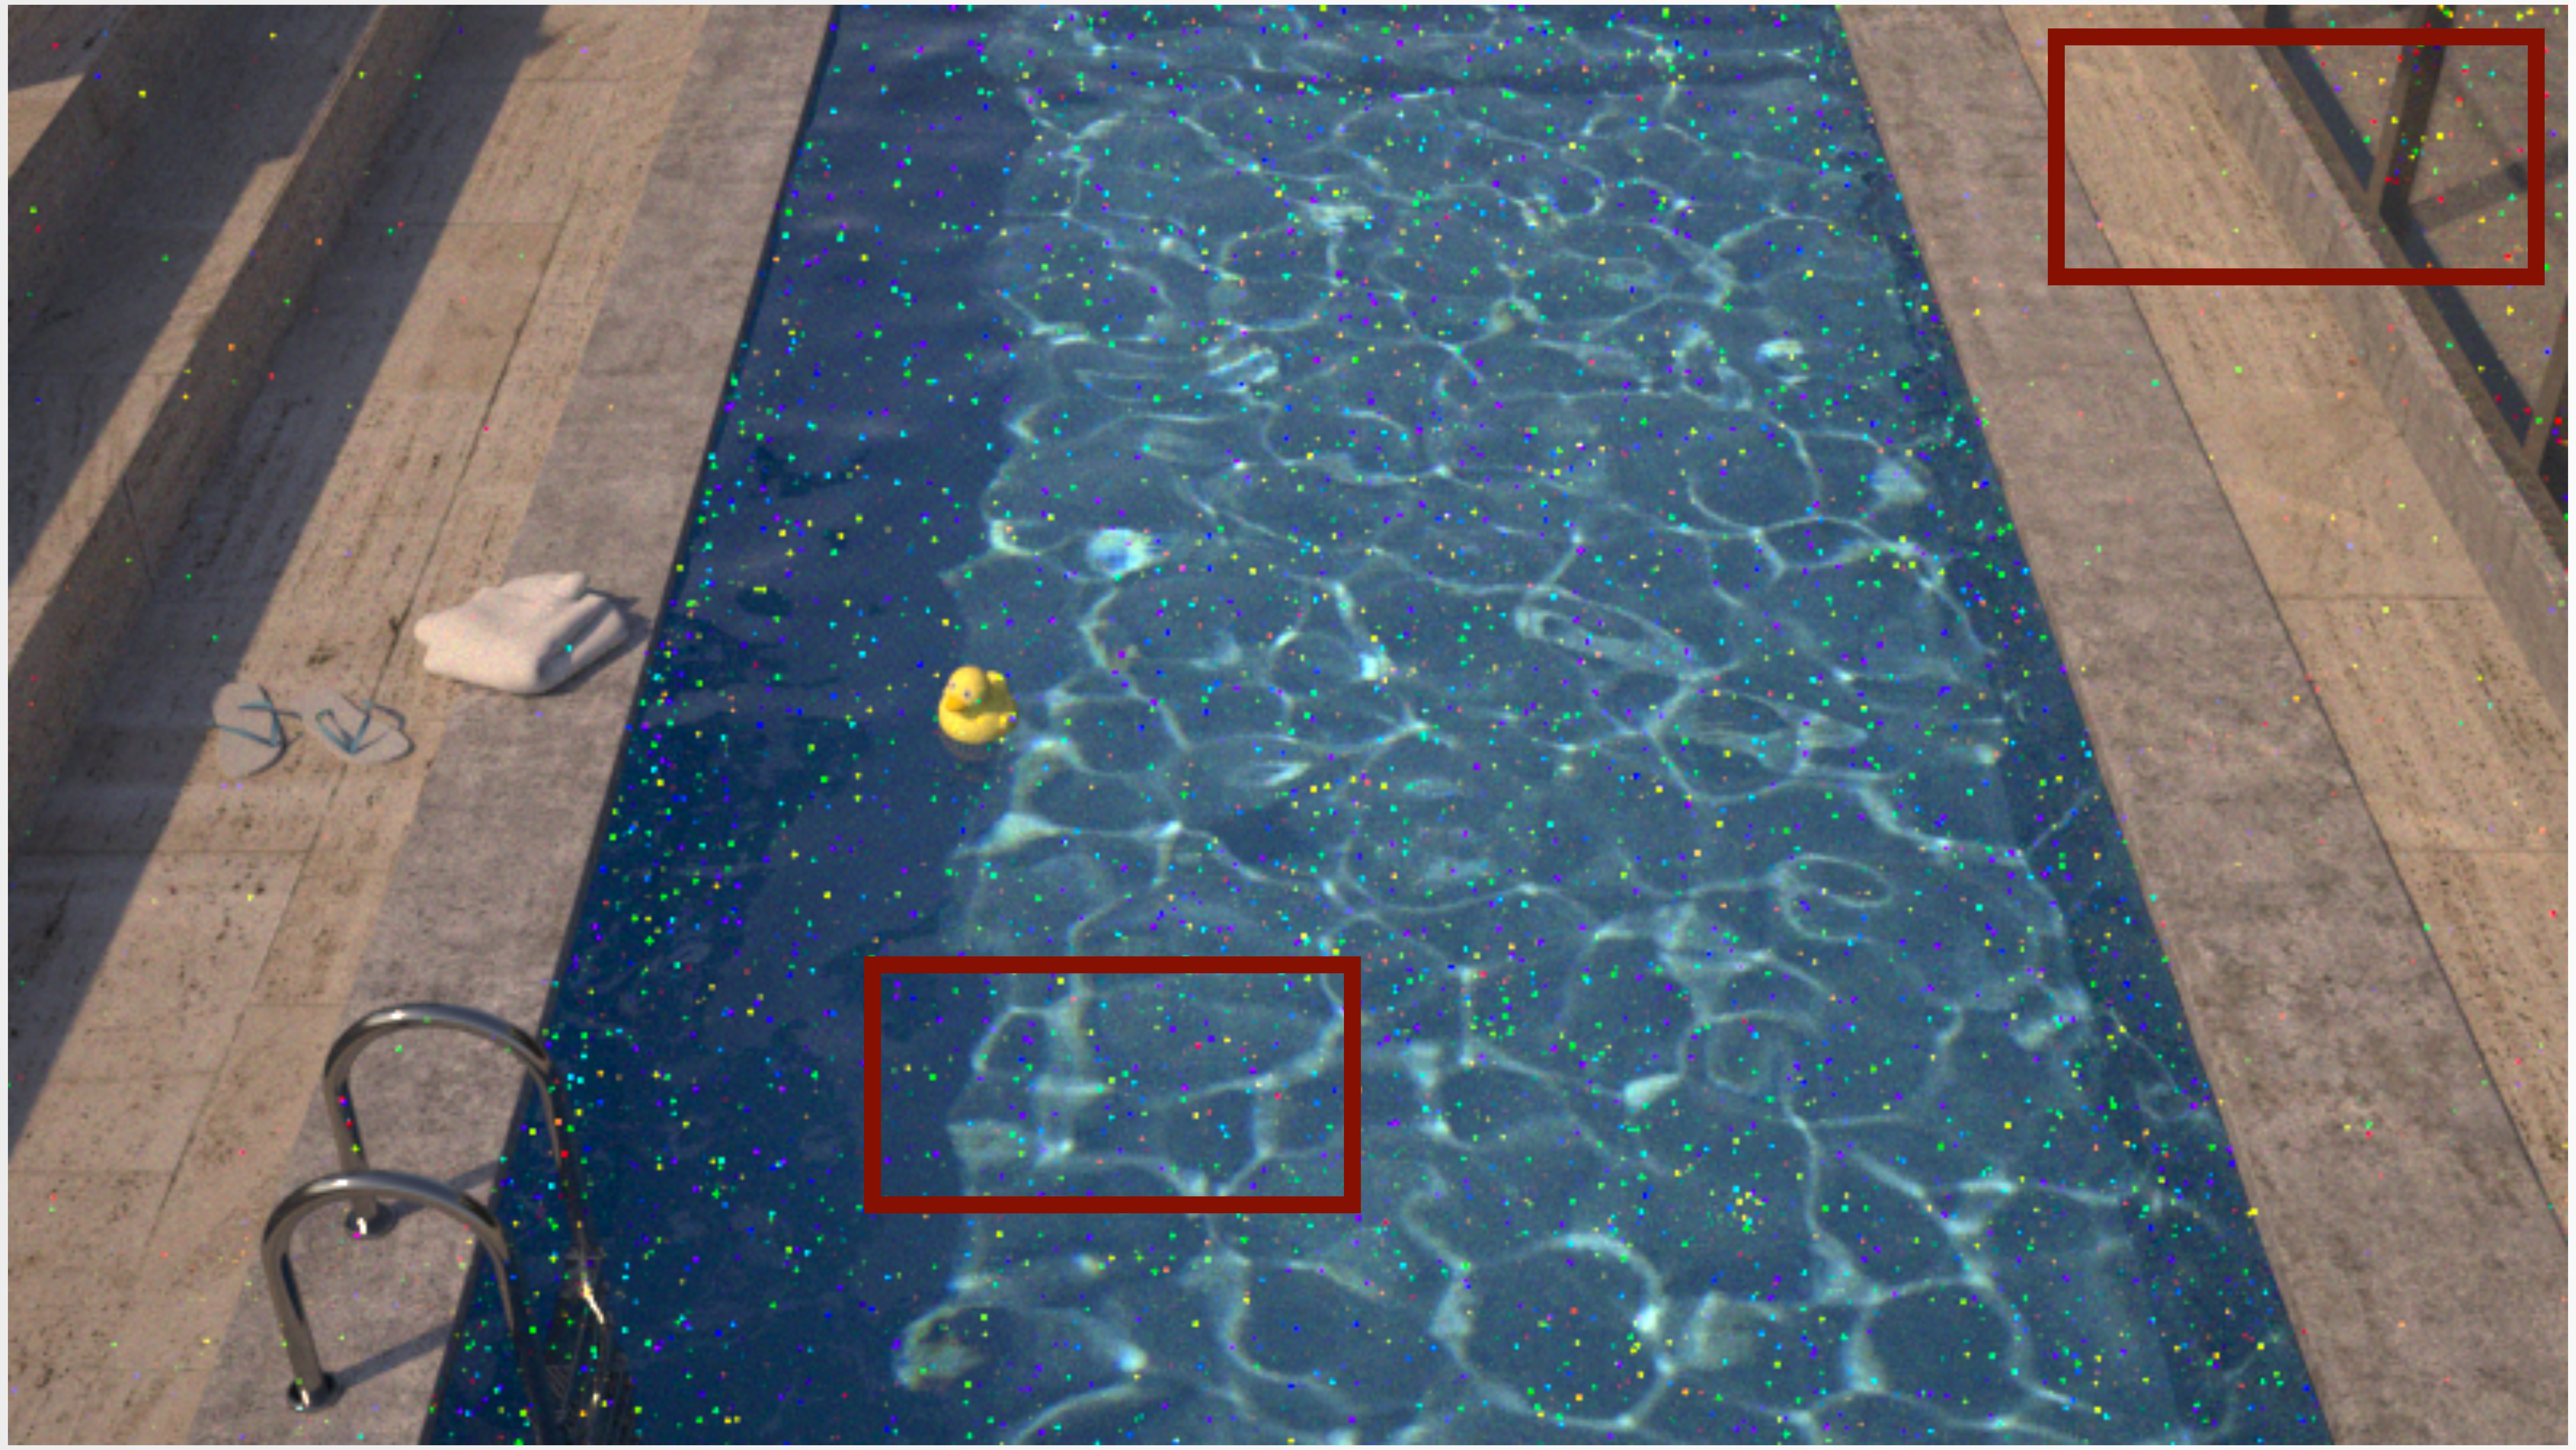

PT | **Guided PT**

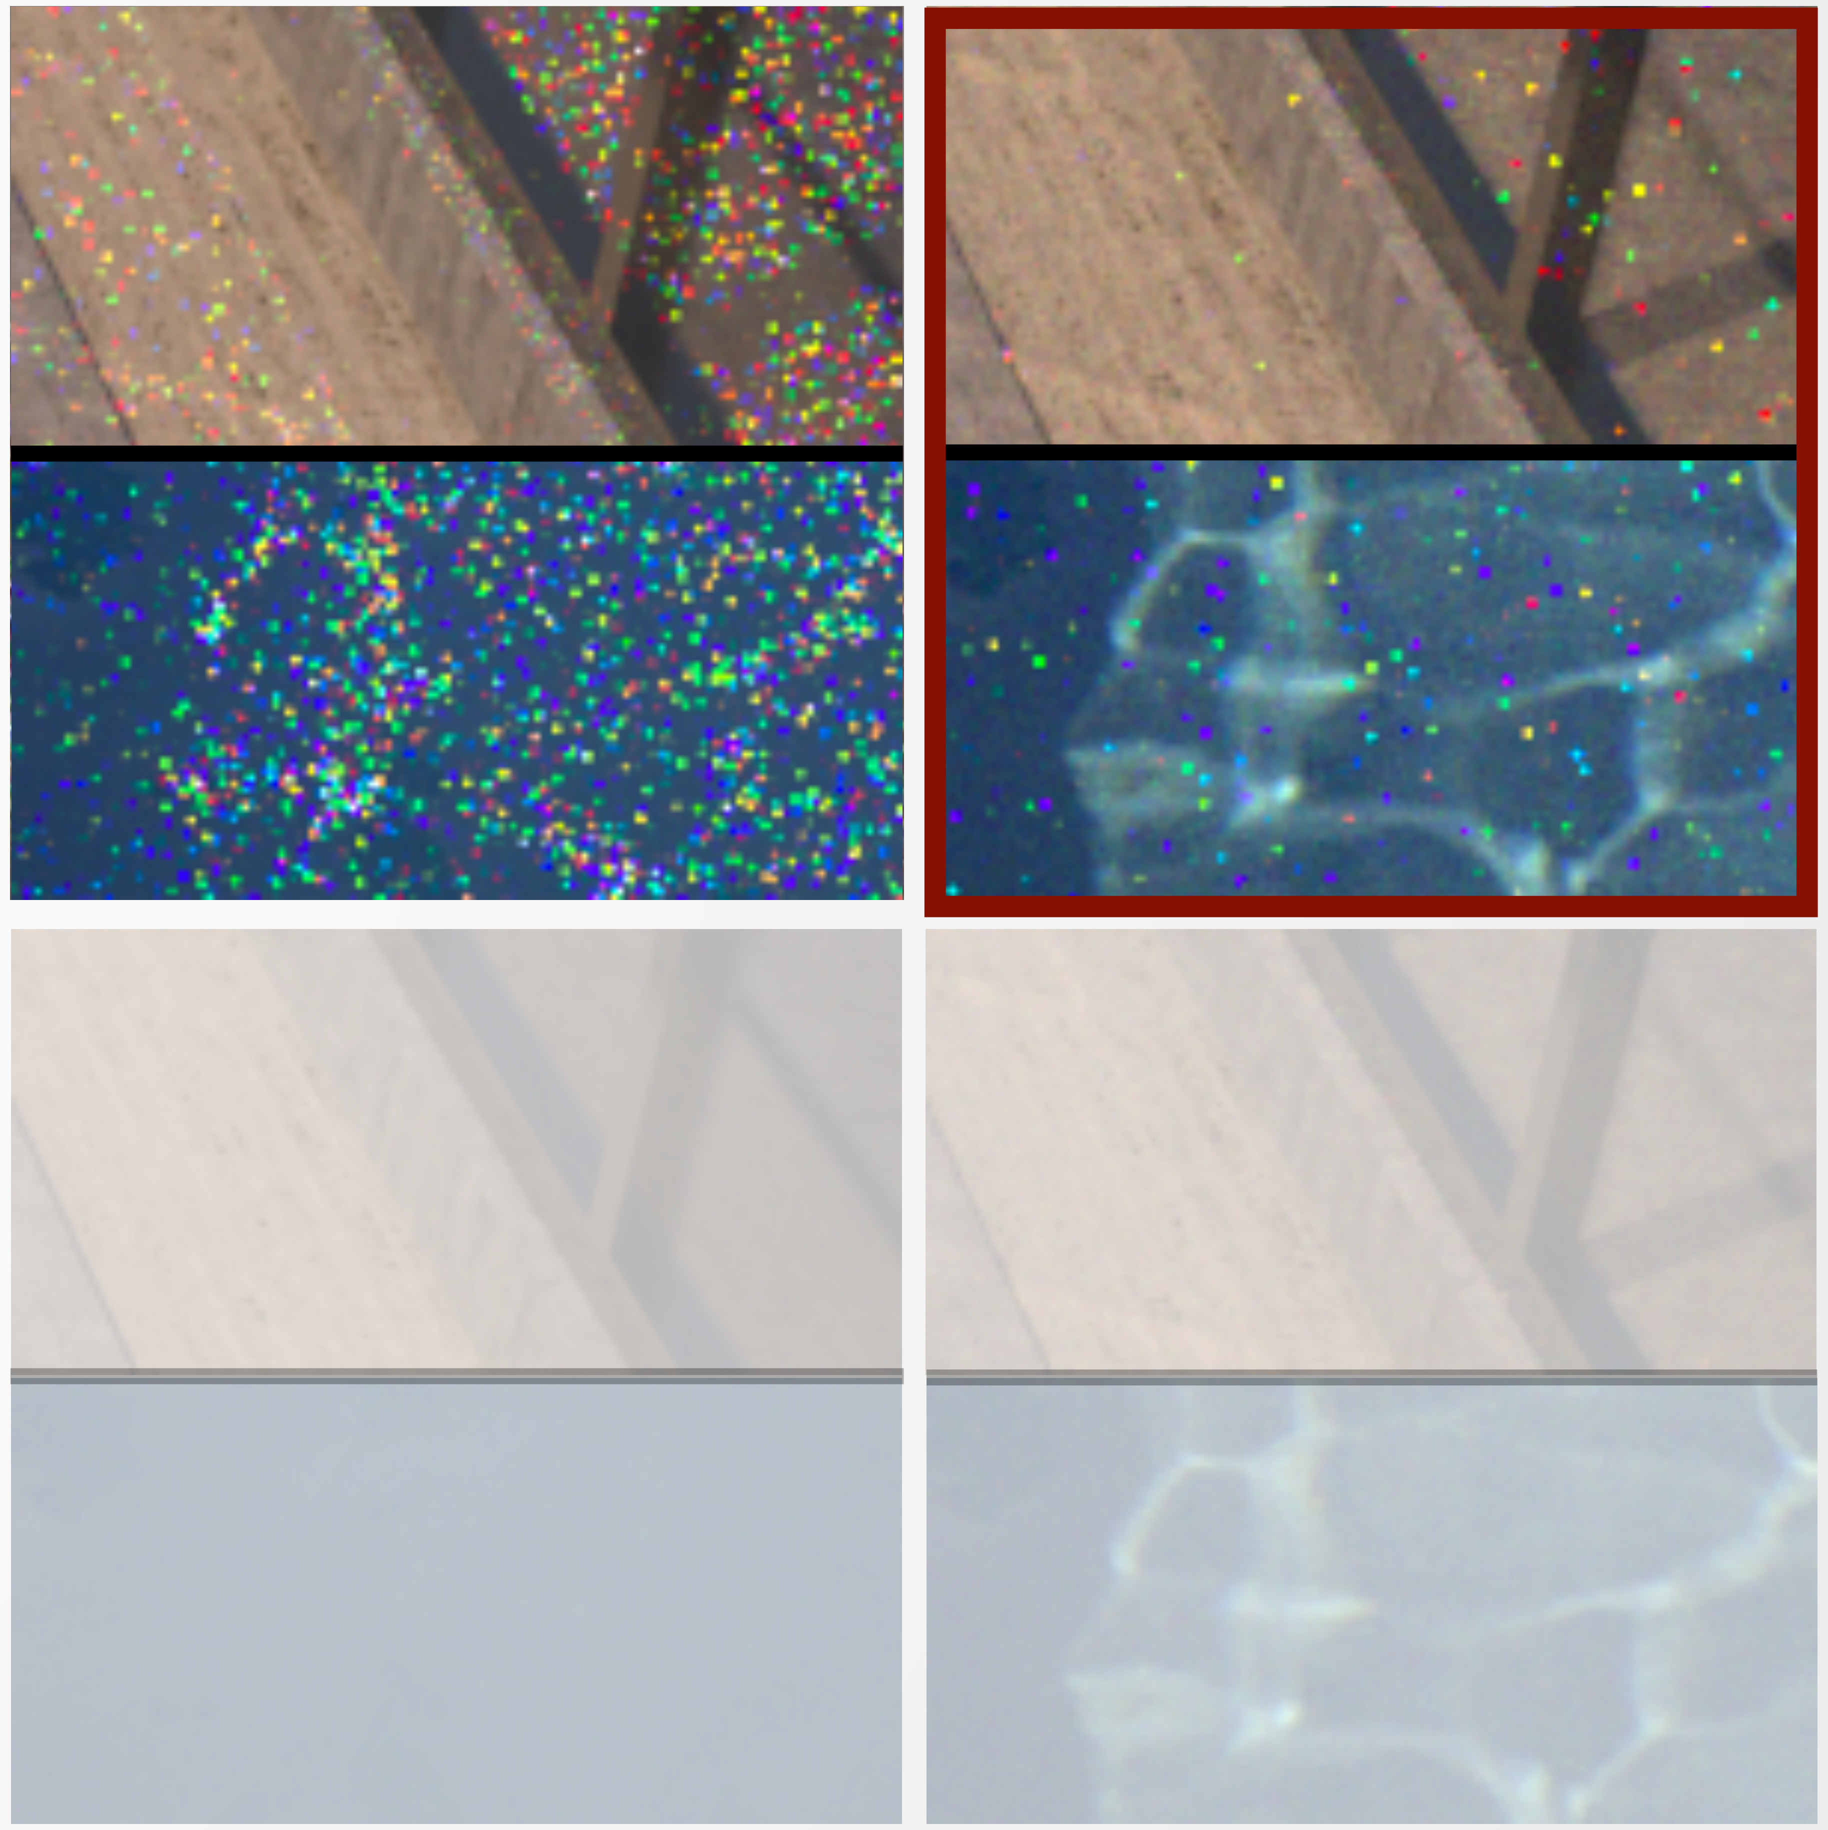

104k guide paths (≈310MB)

Pool - 30min

104k guide paths (≈310MB)

Pool - 30min

Dragon

Reference

KIT
Karlsruhe Institute of Technology

Dragon - 10h

Path tracing

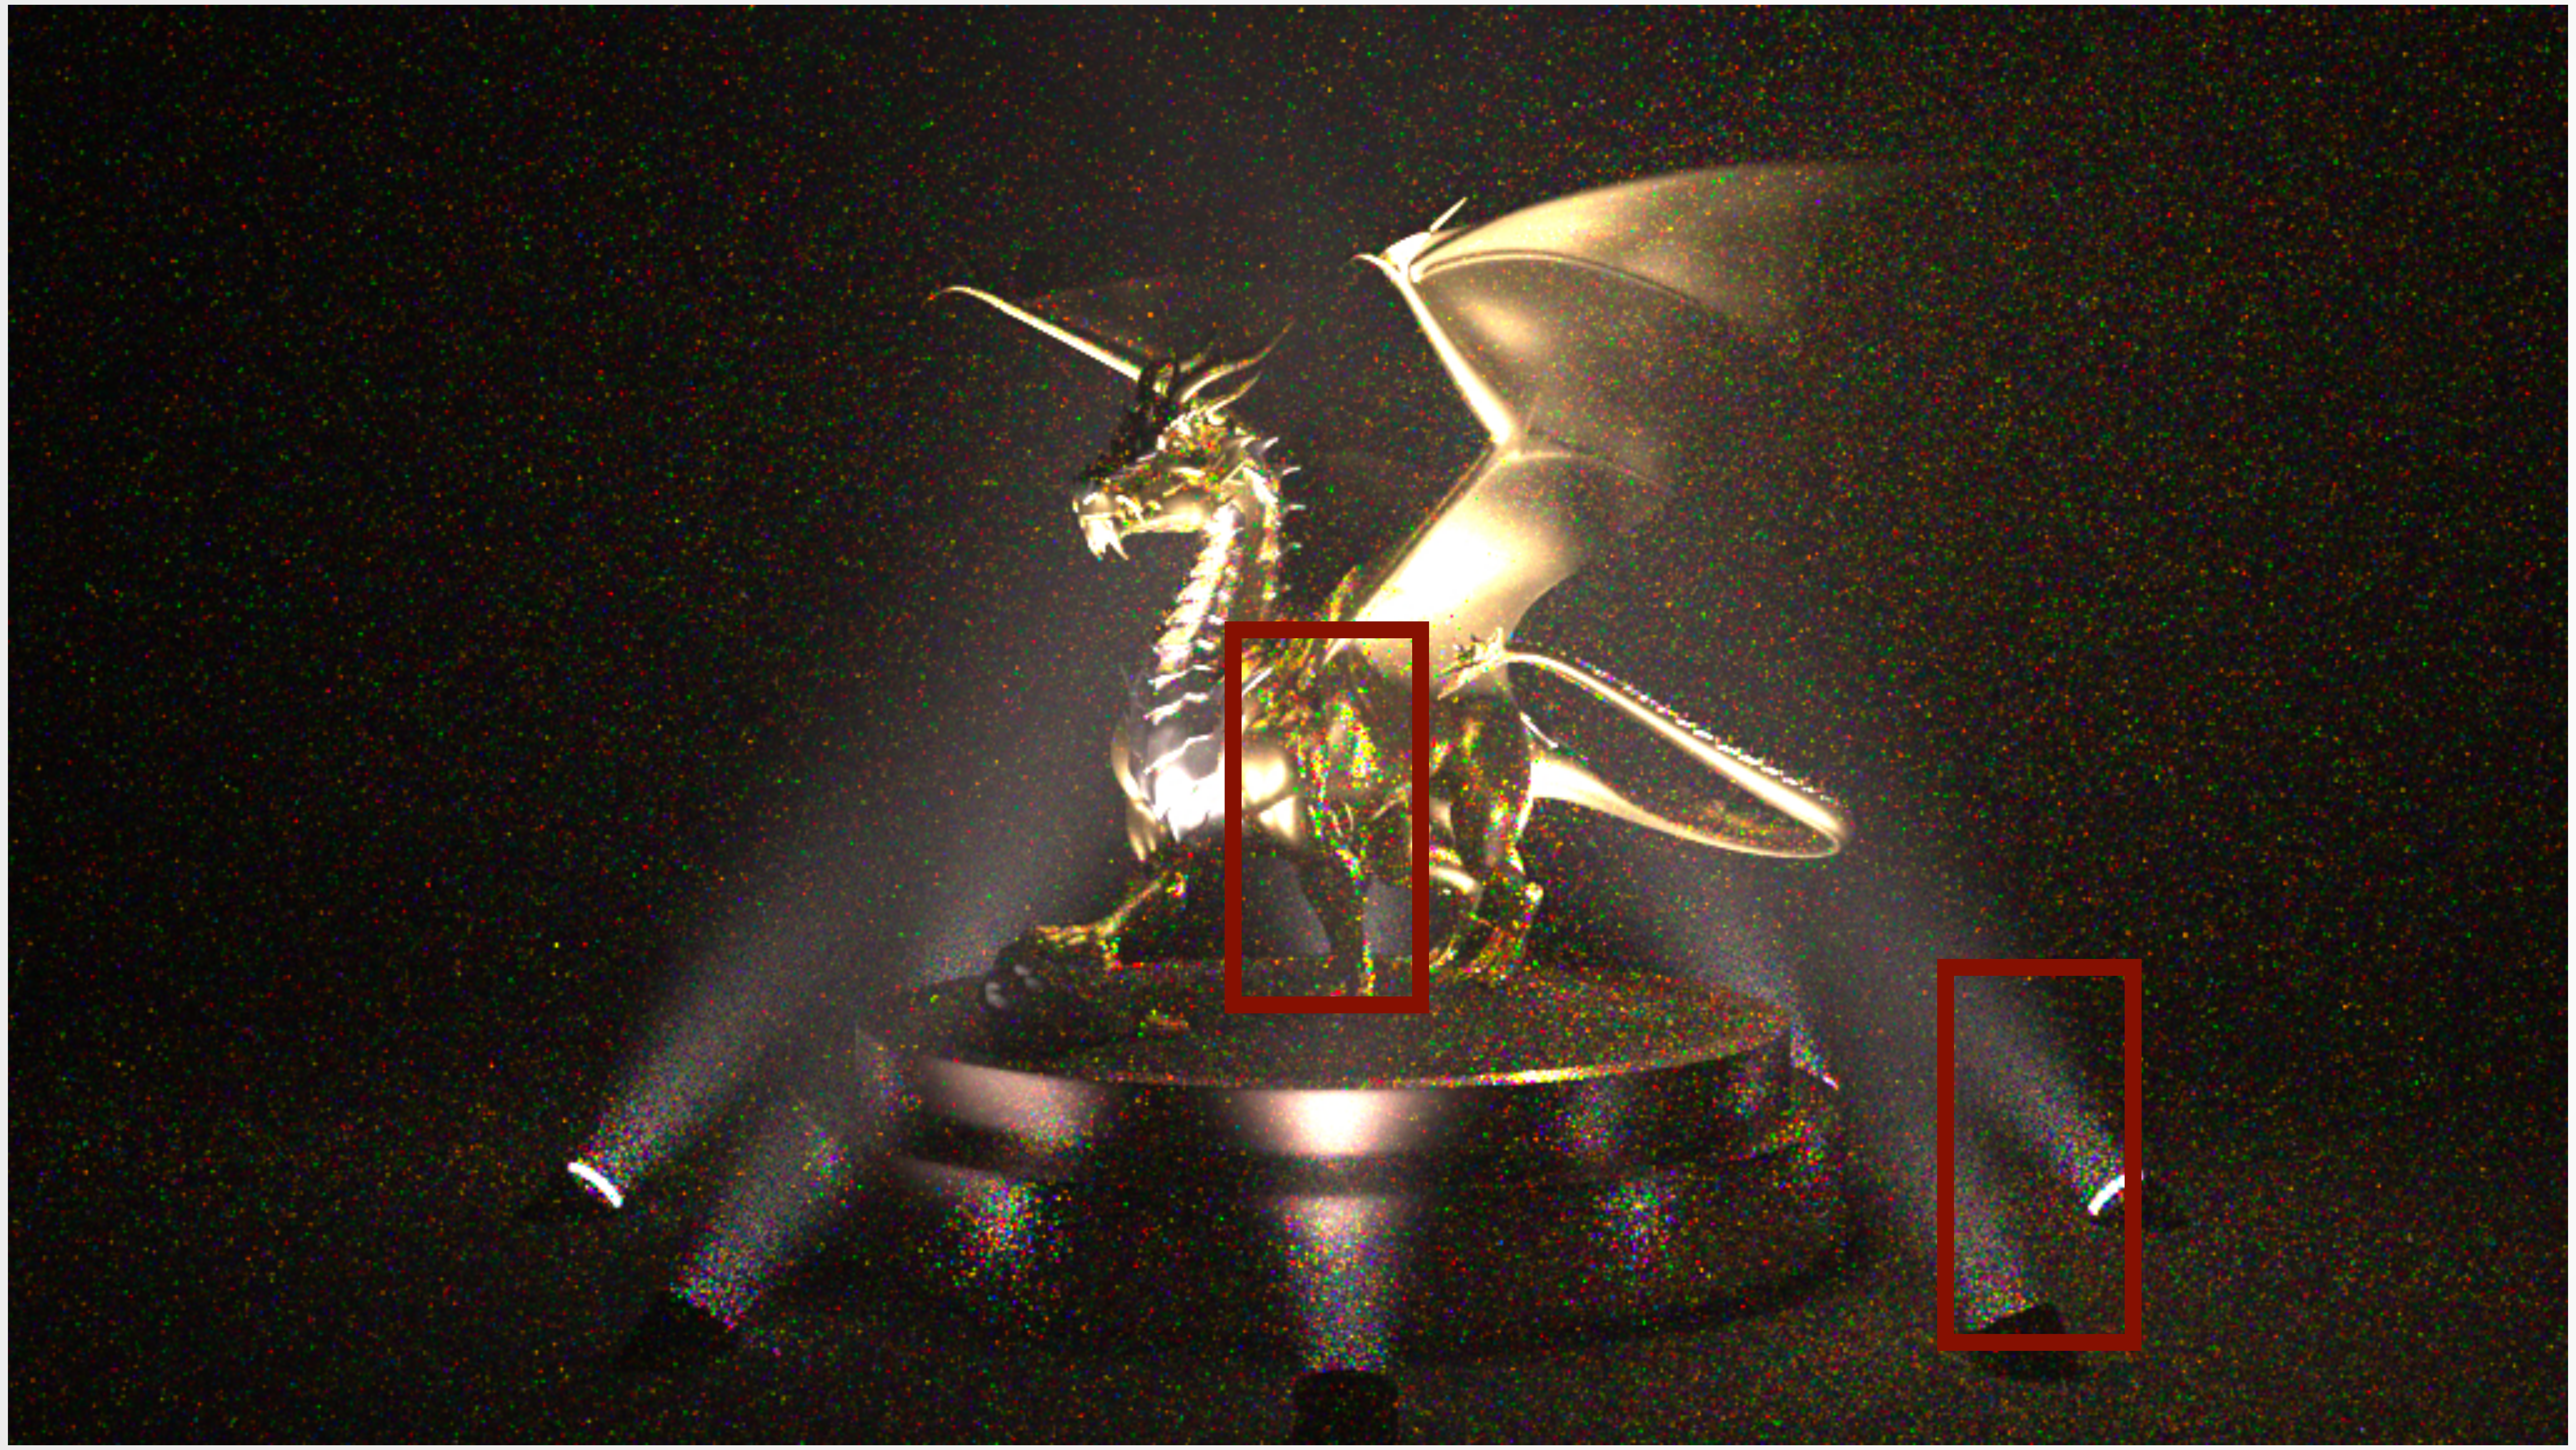

PT

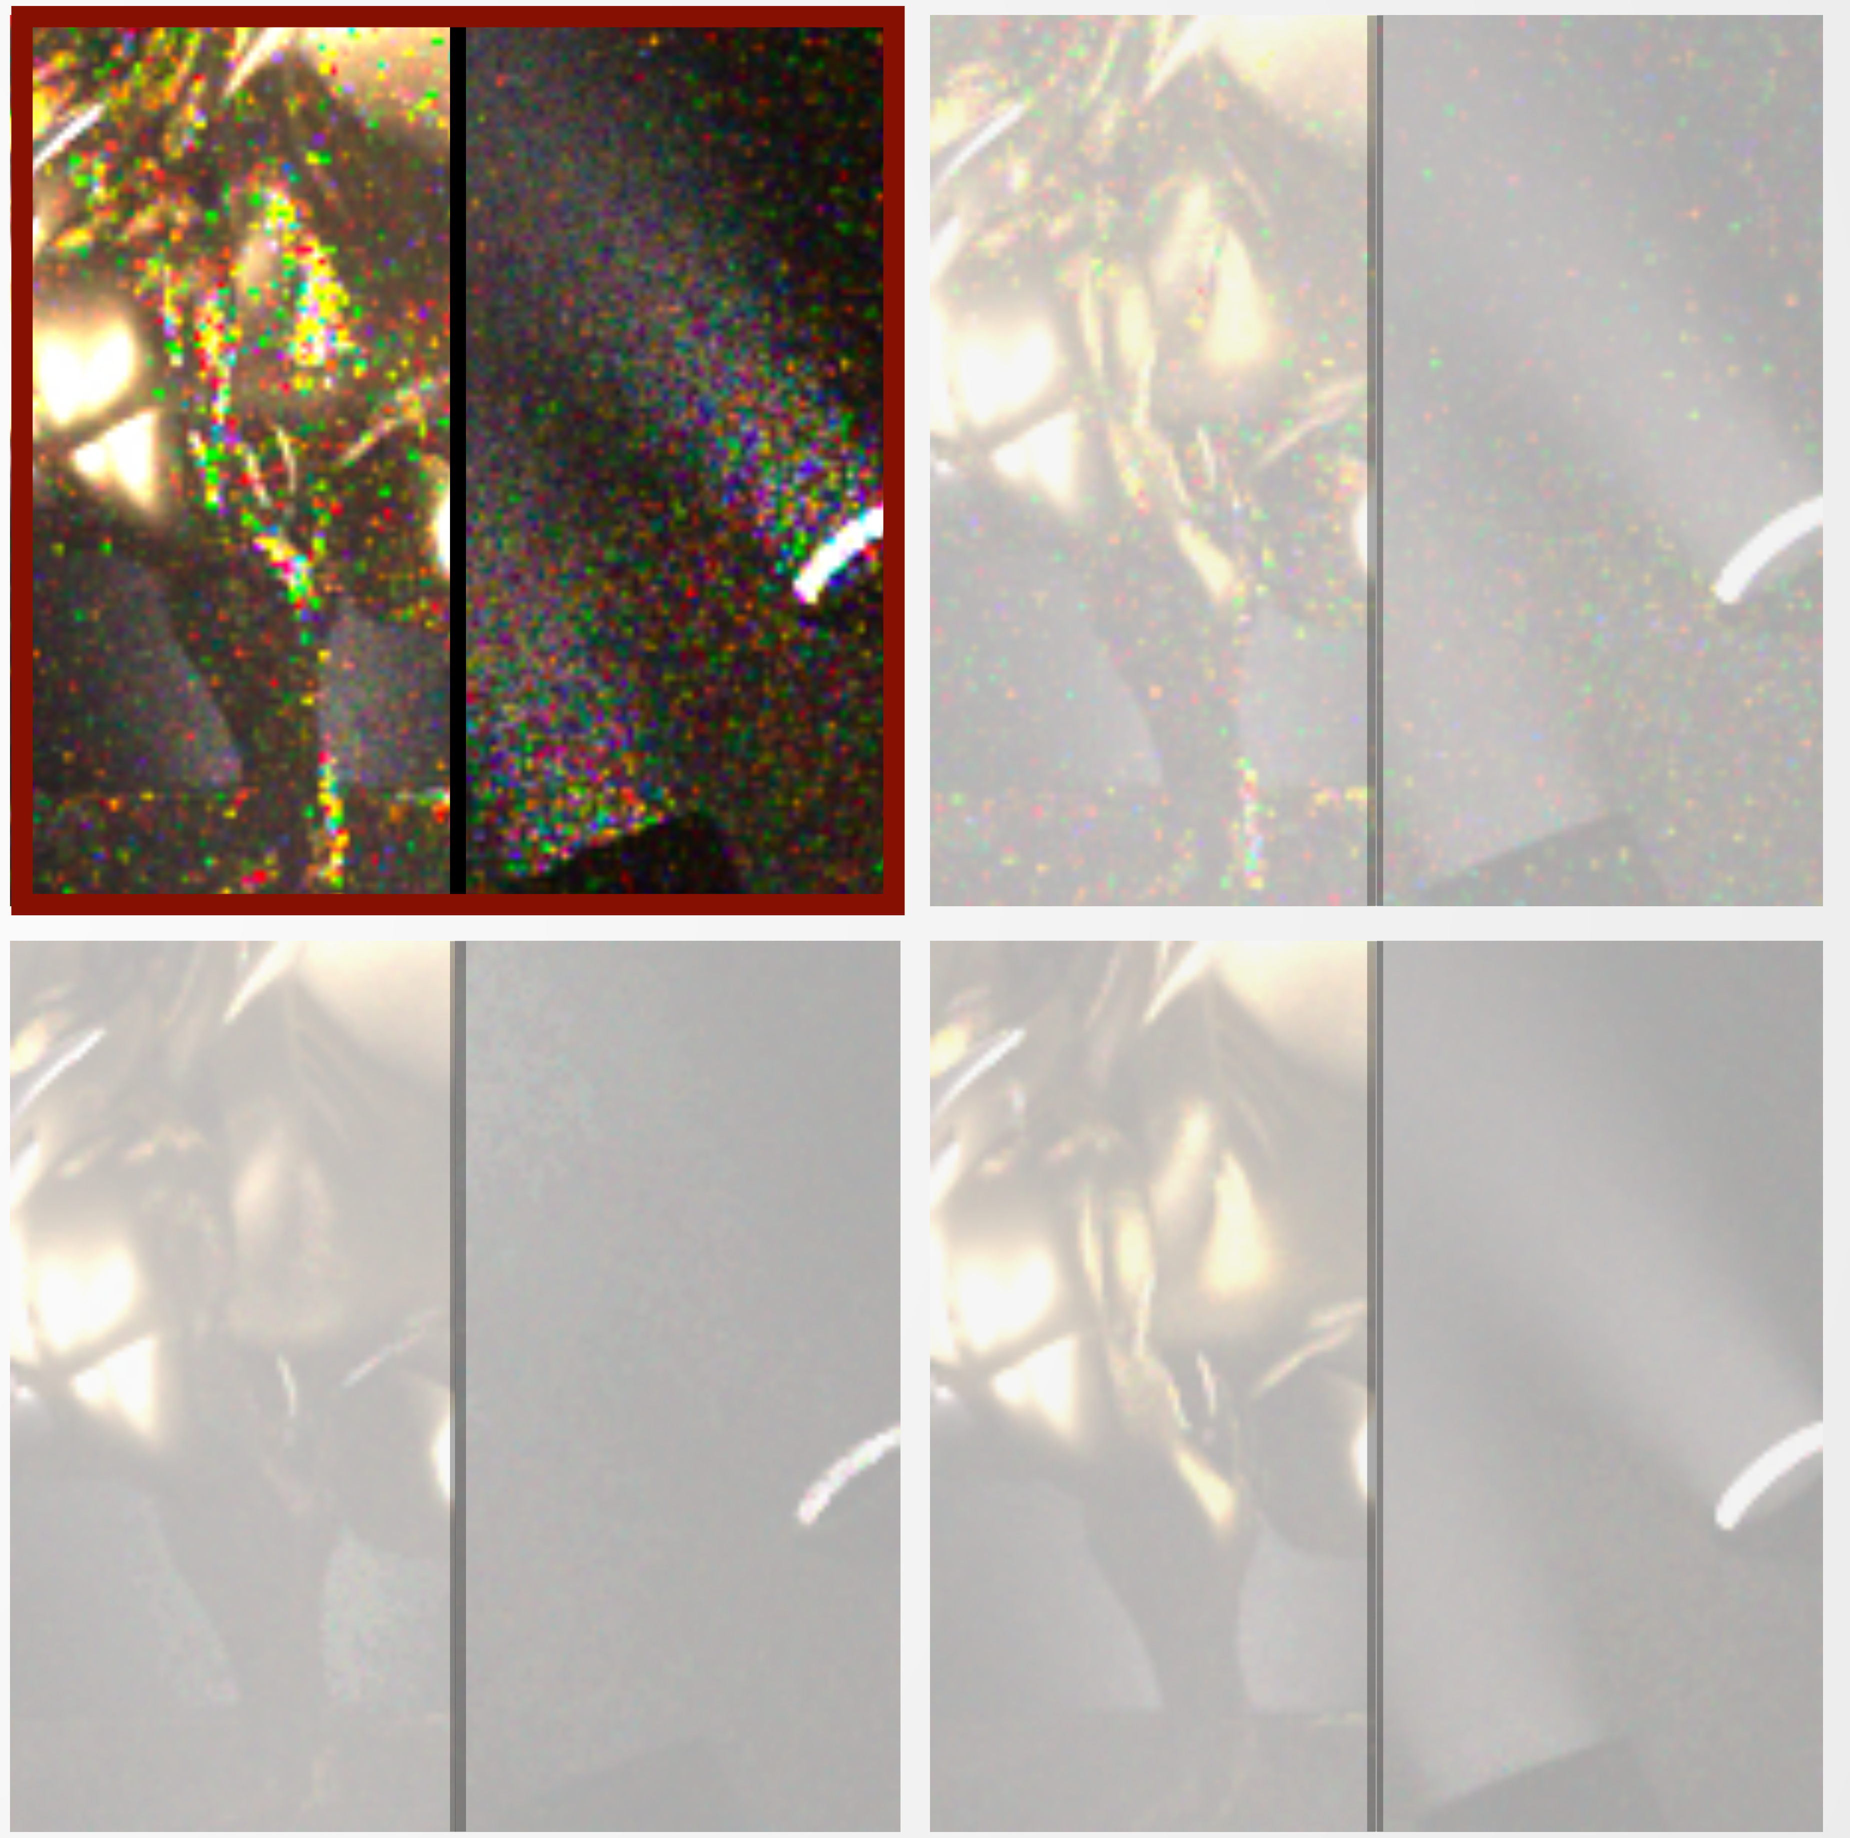

Dragon - 10h

Guided path tracing

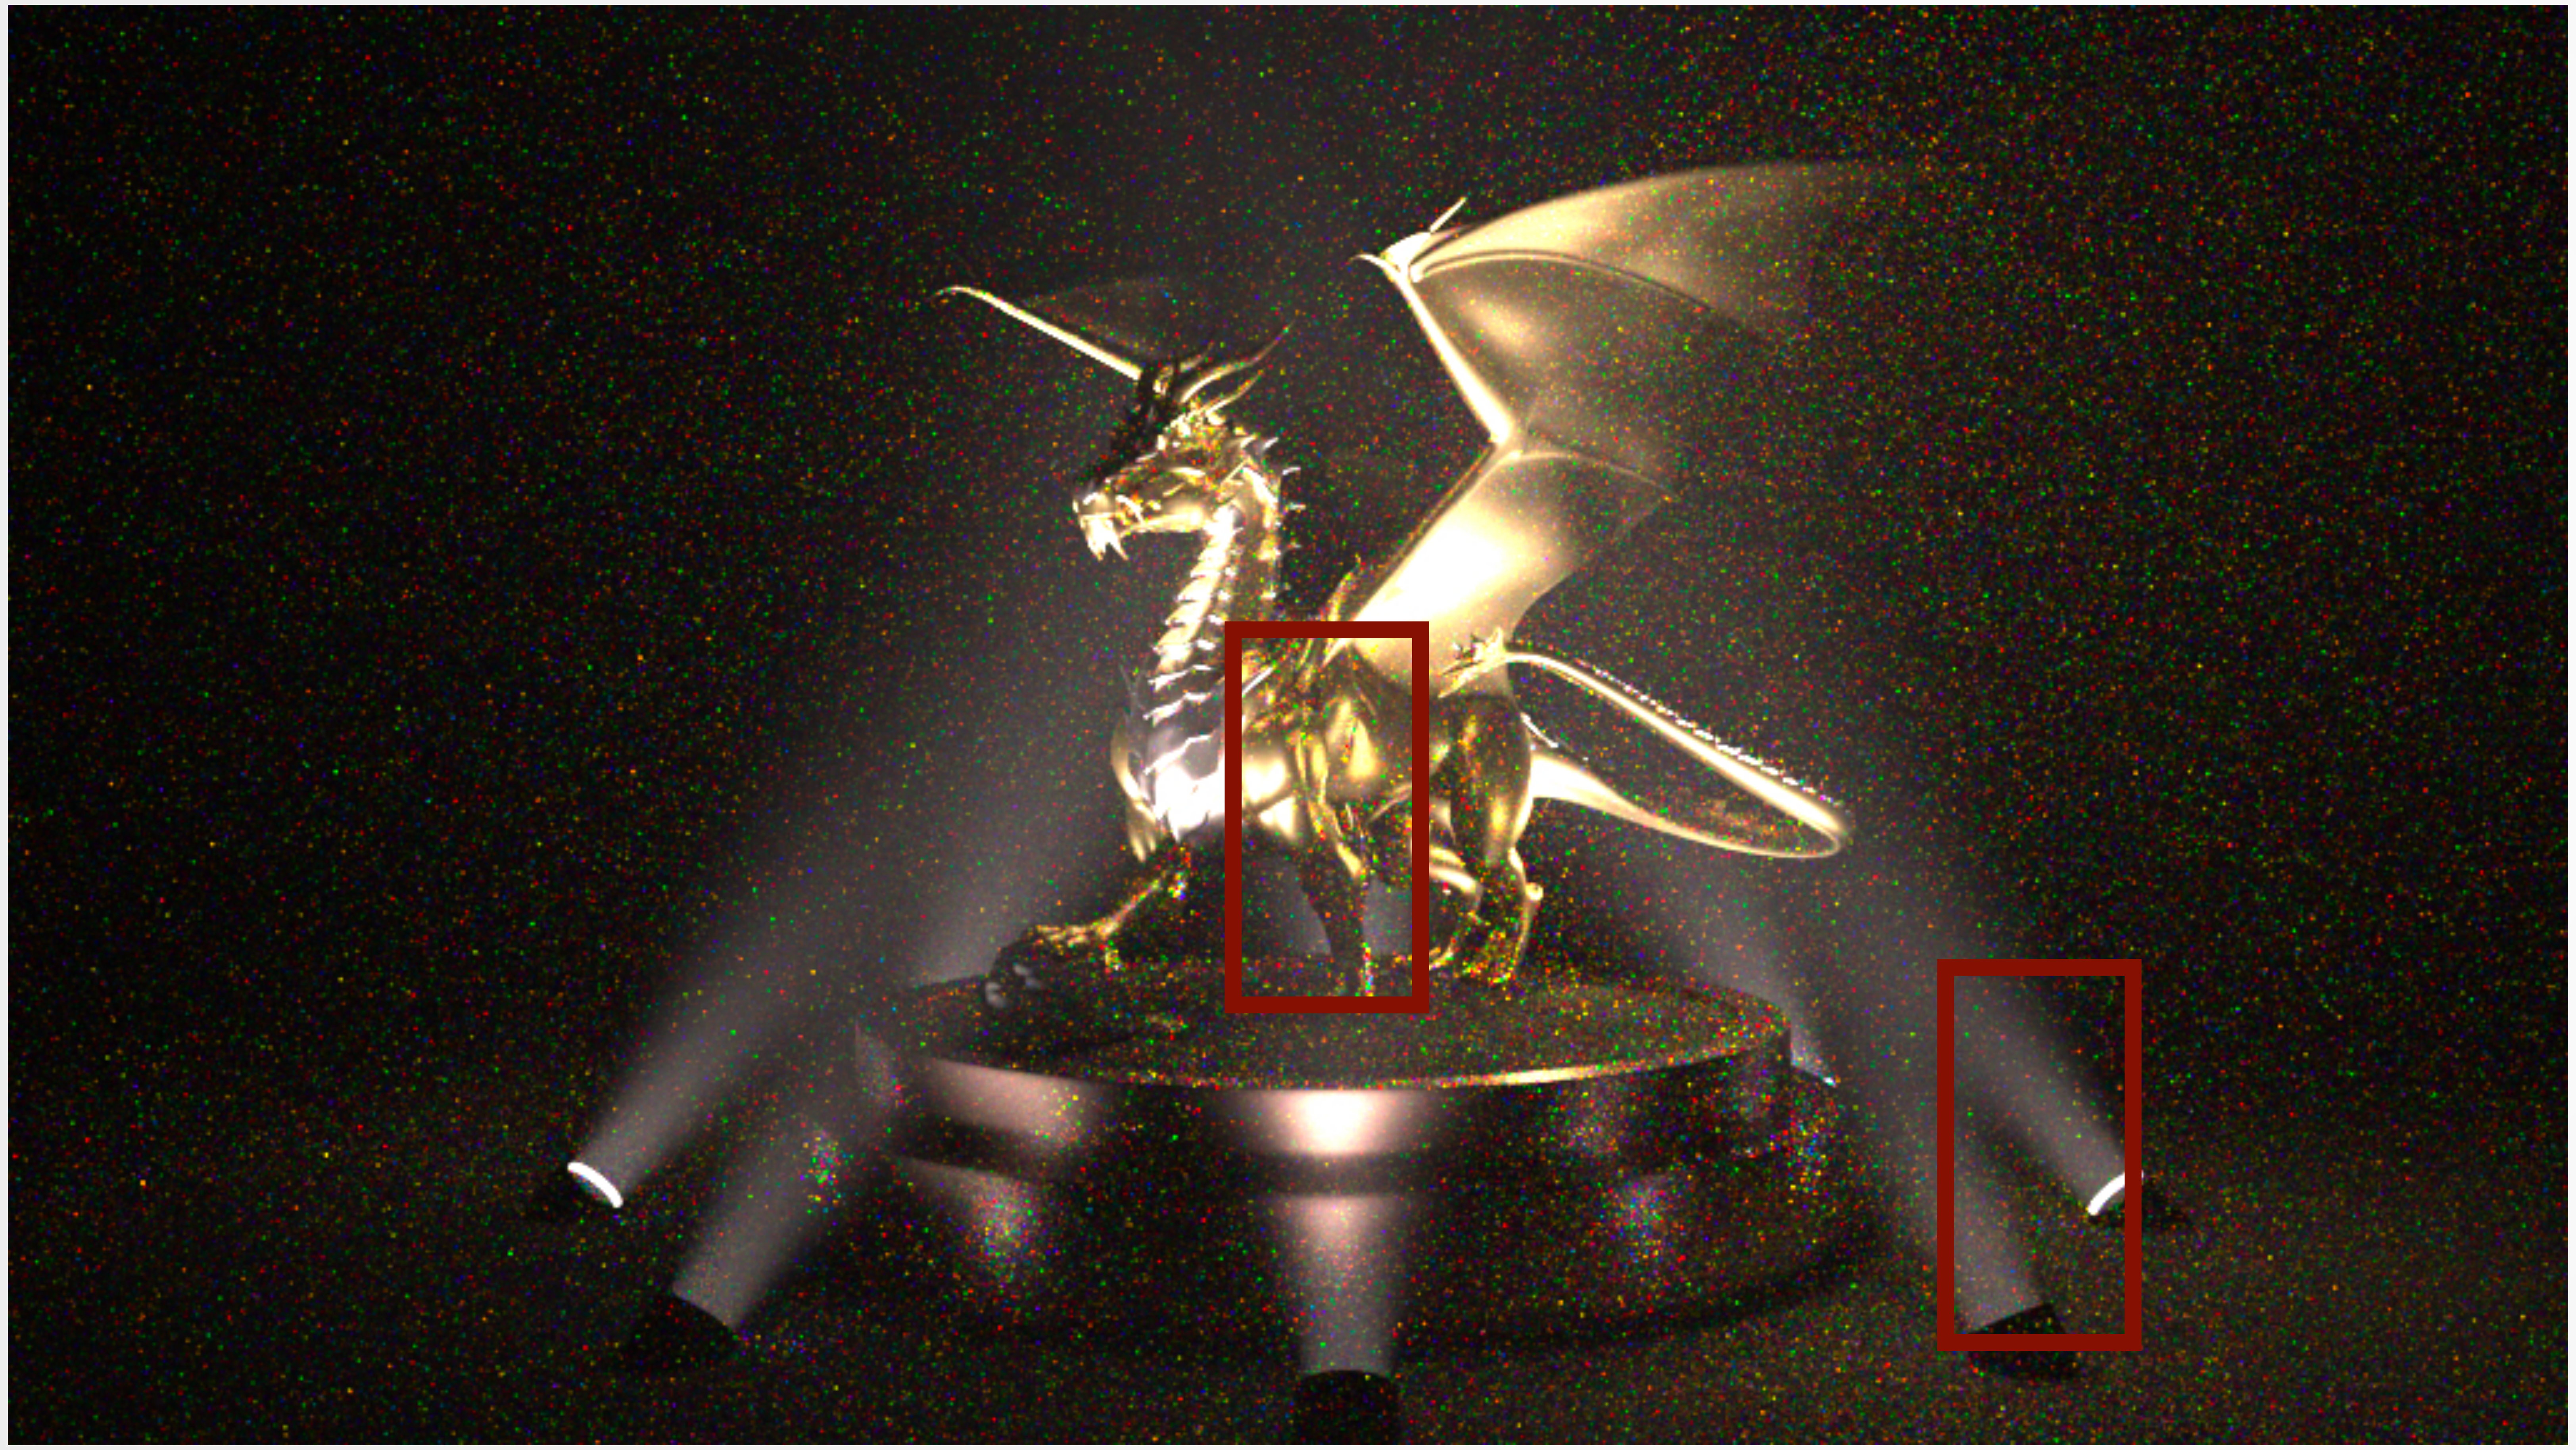

69k guide paths (≈207MB)

PT	Guided PT

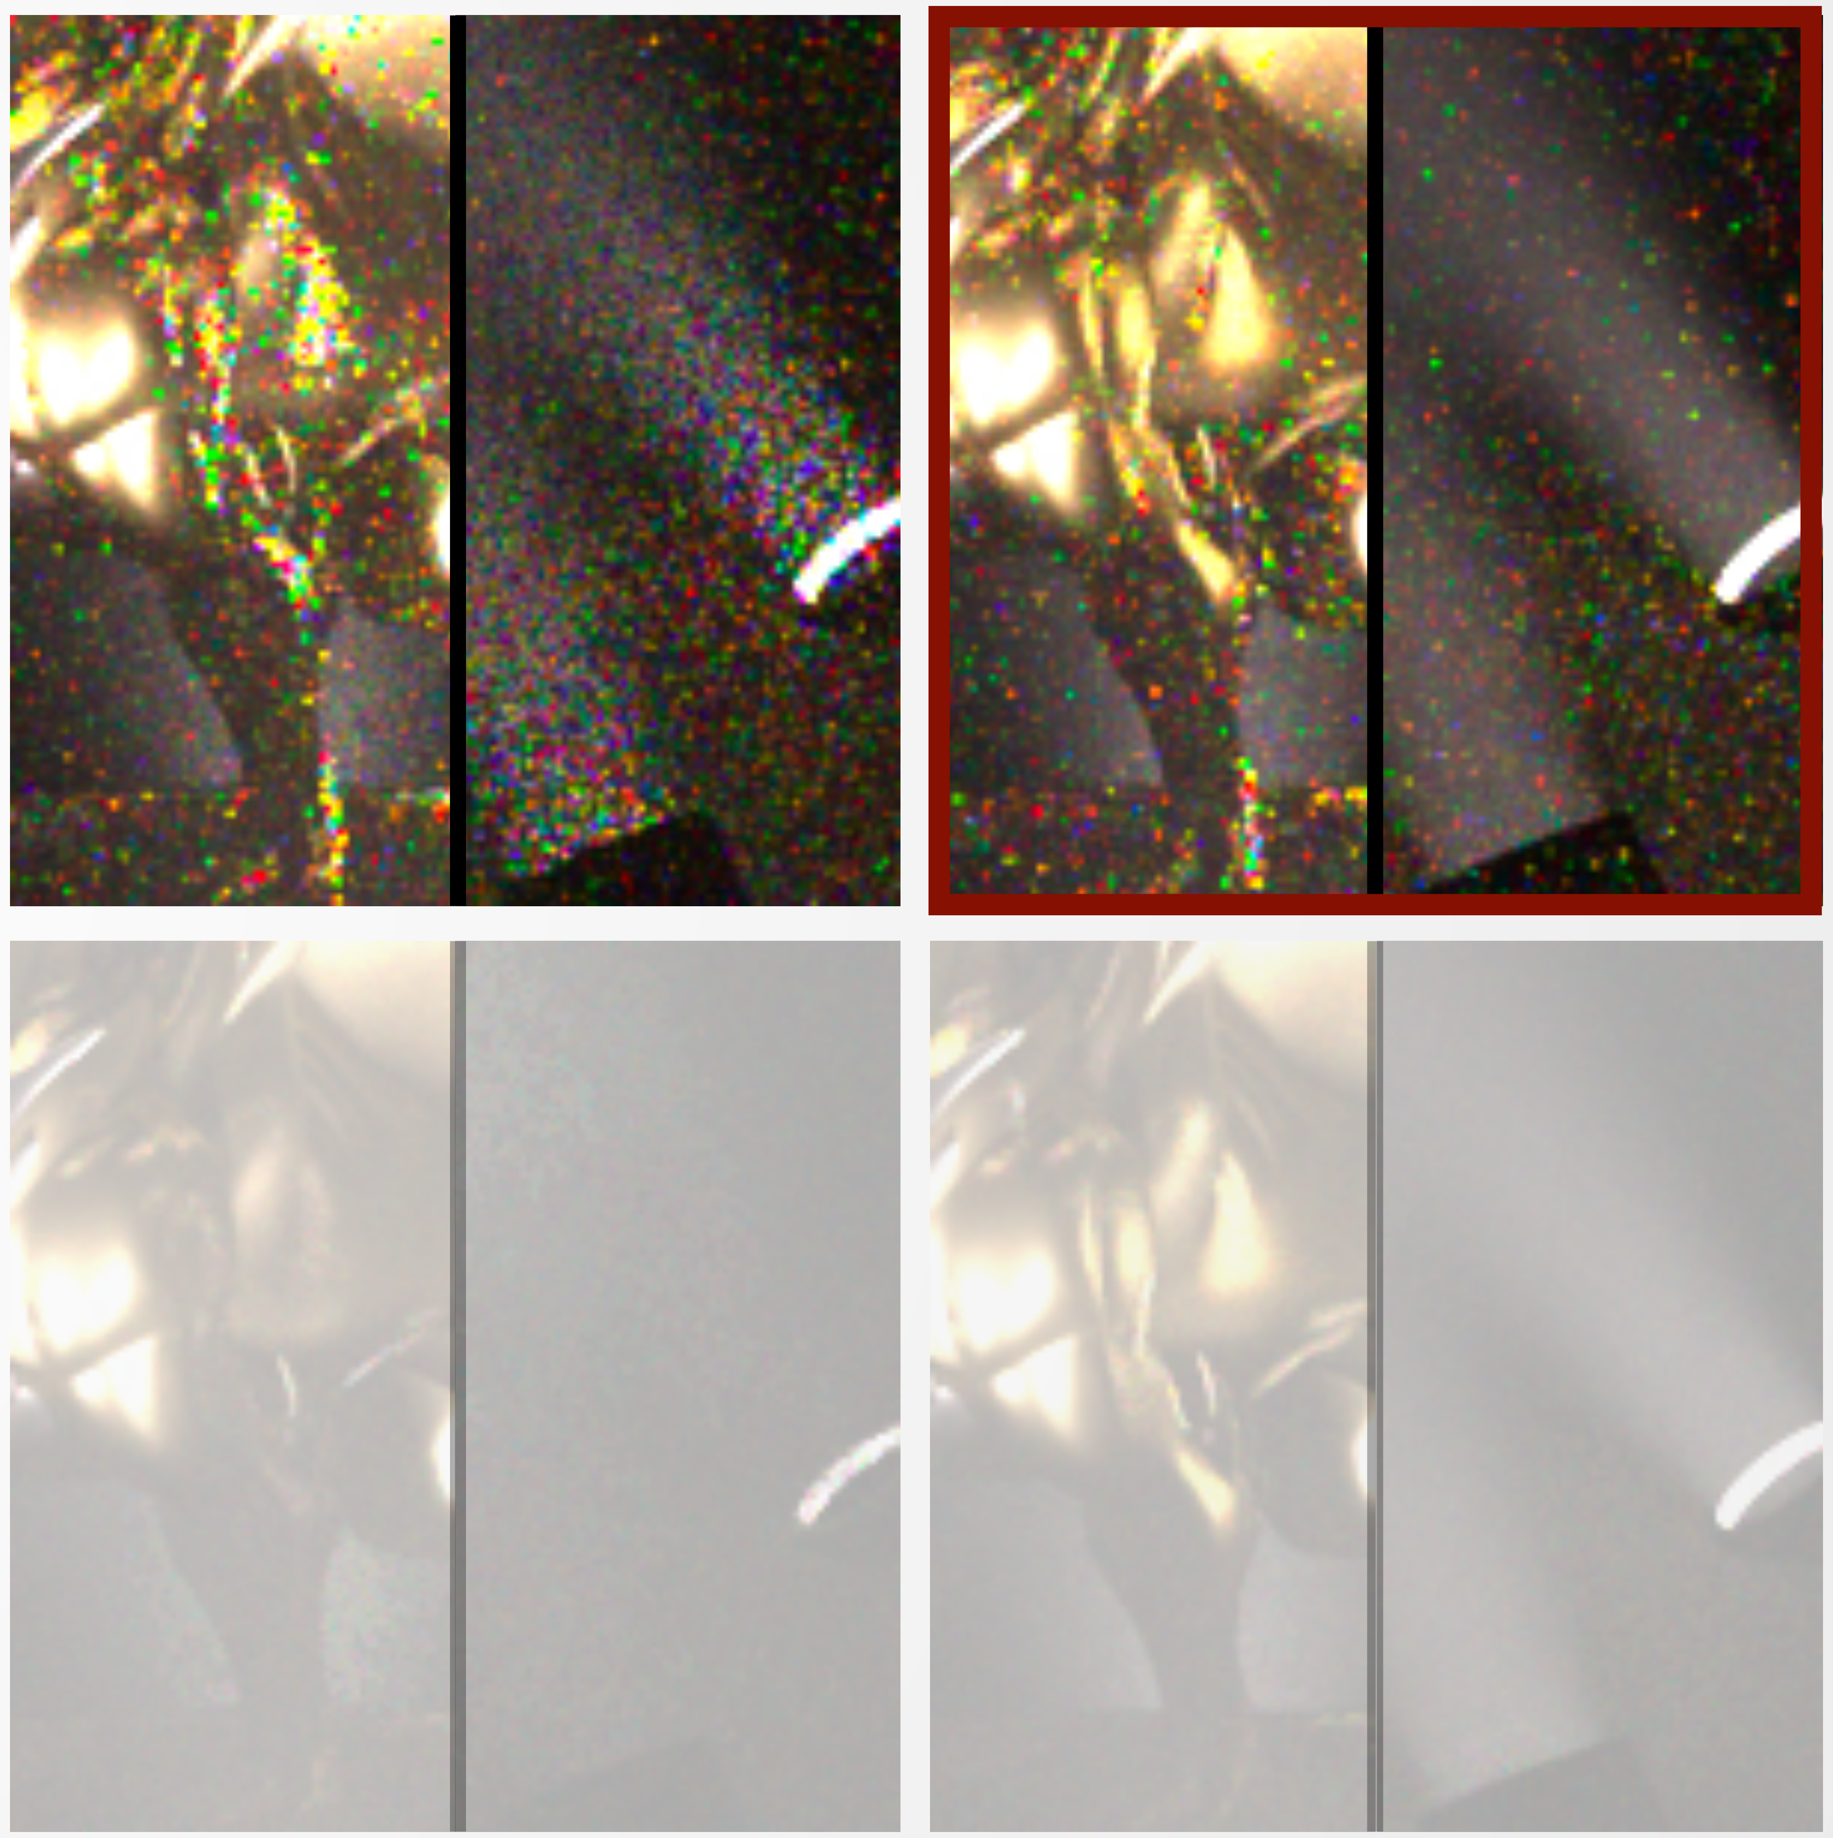

Dragon - 10h

Guided path tracing + DBOR

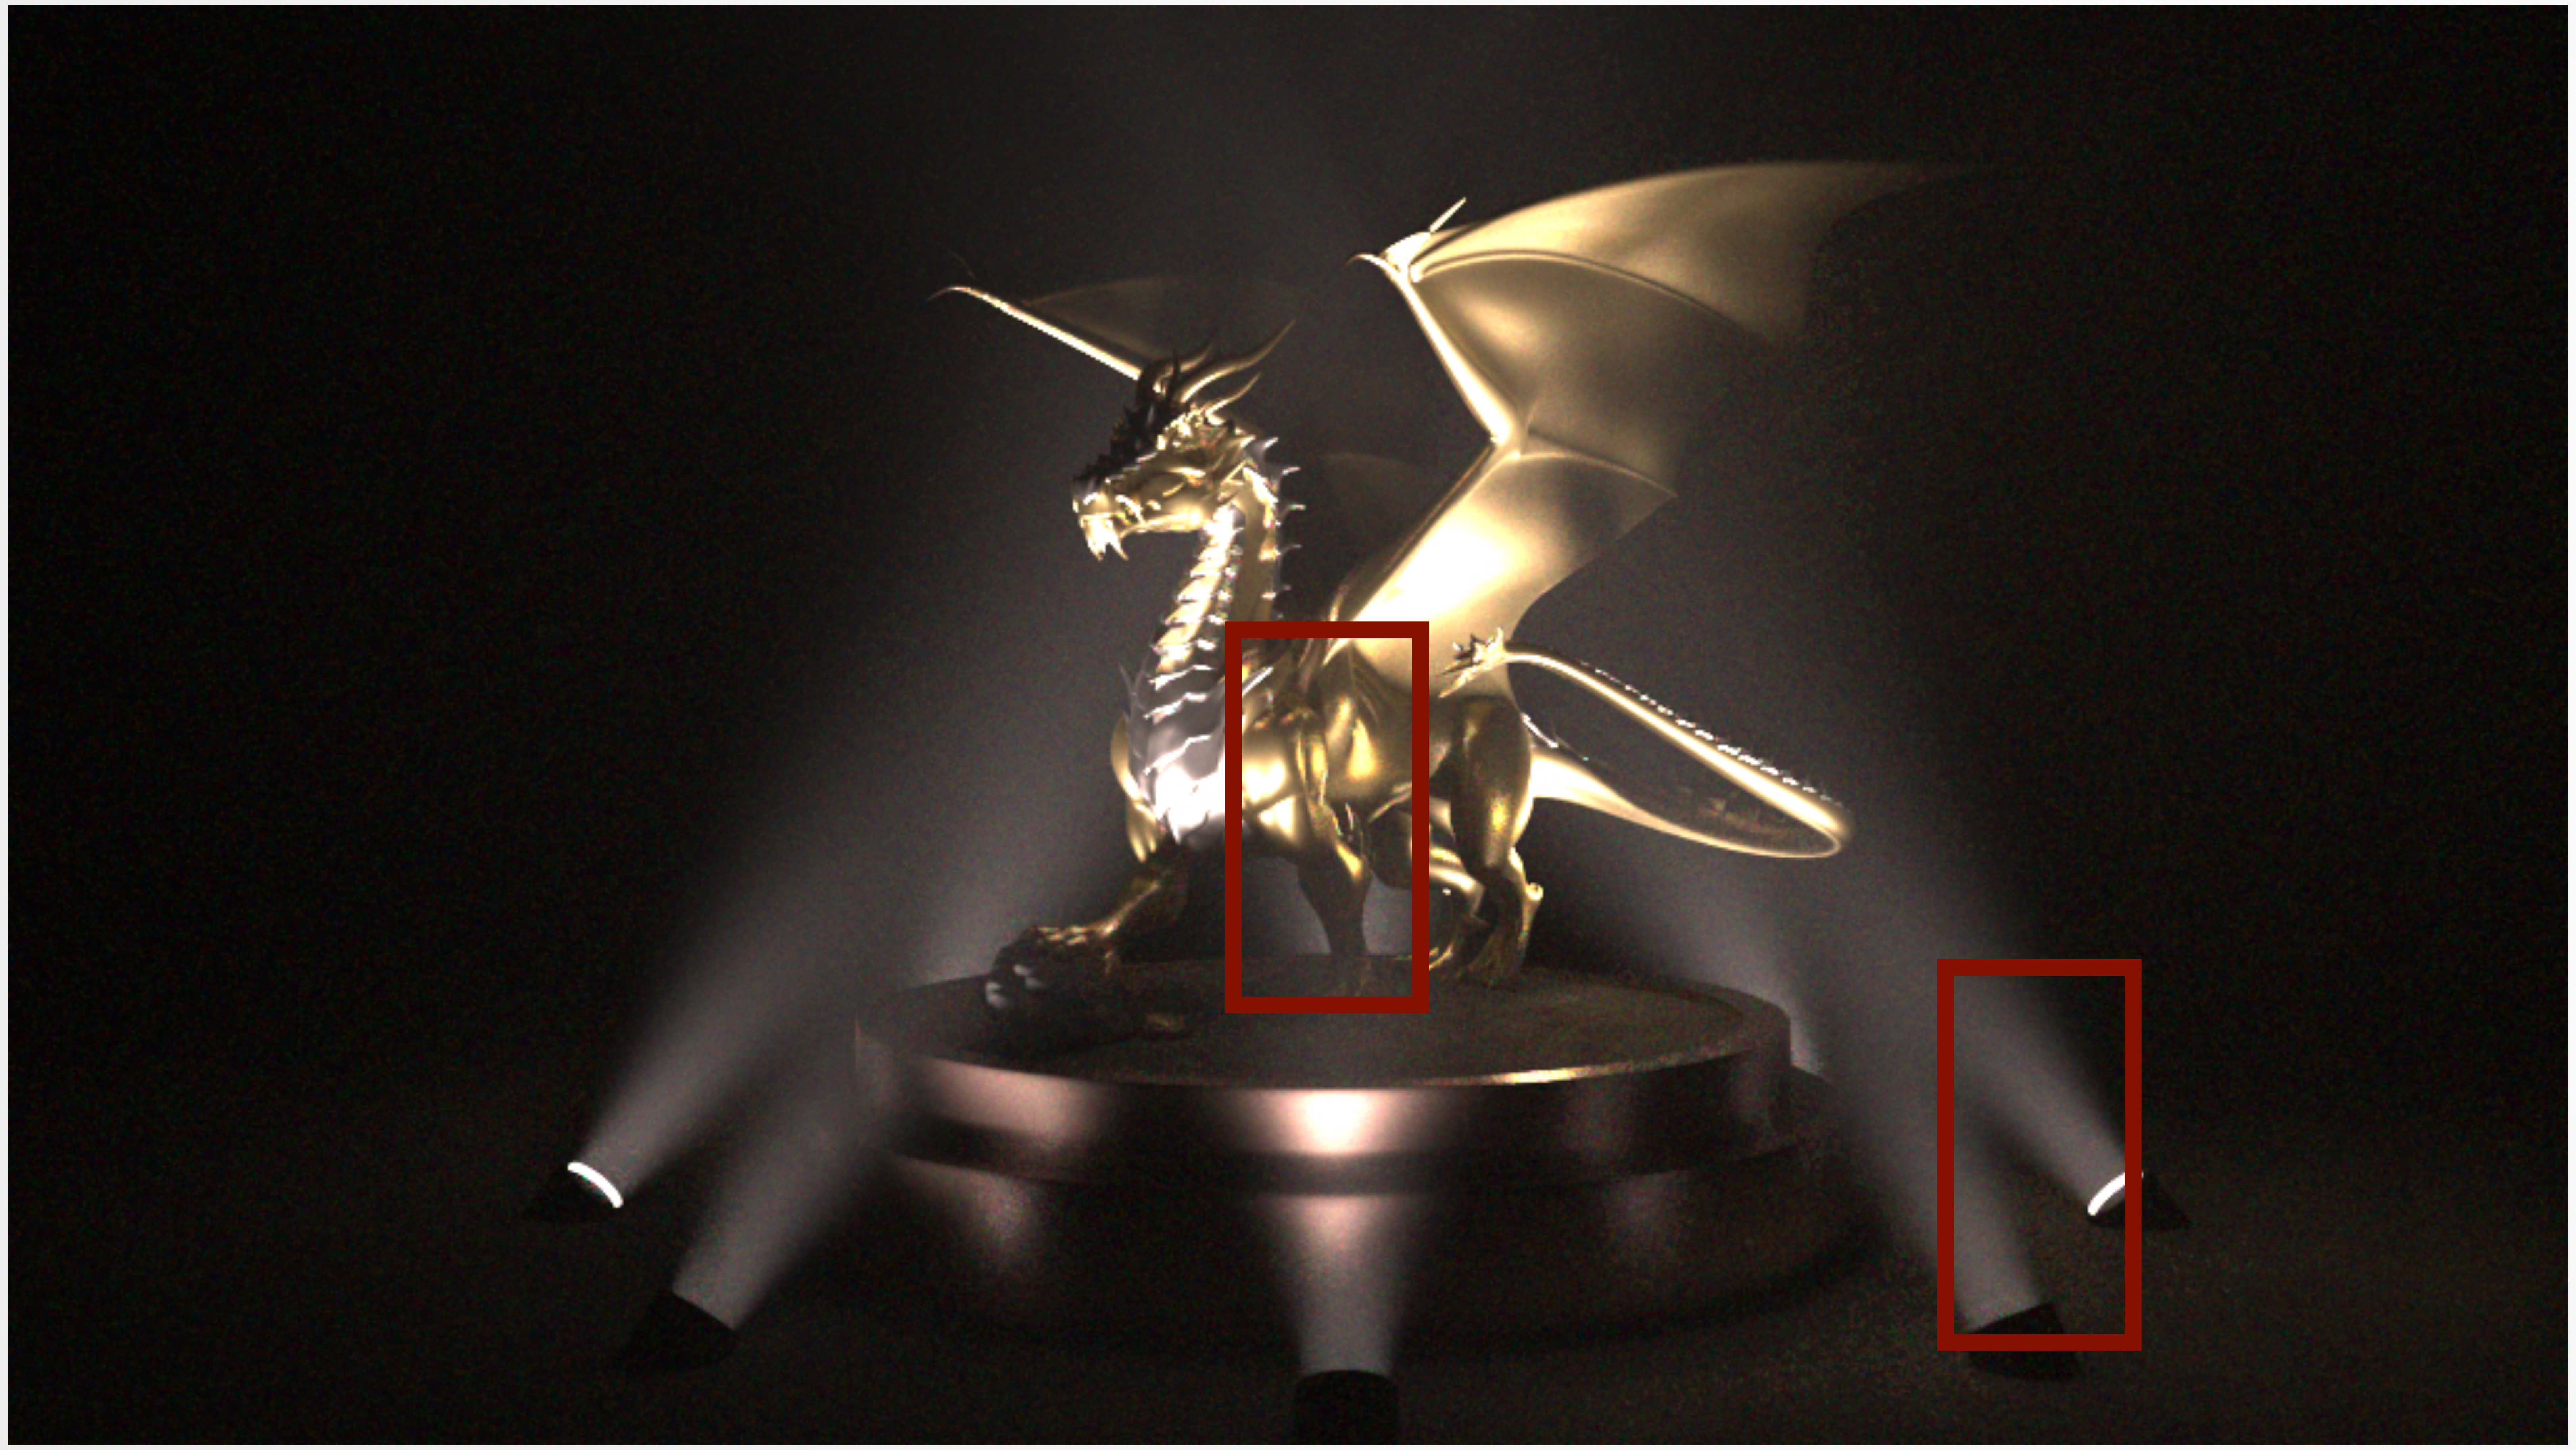

69k guide paths (≈207MB)

PT | **Guided PT**

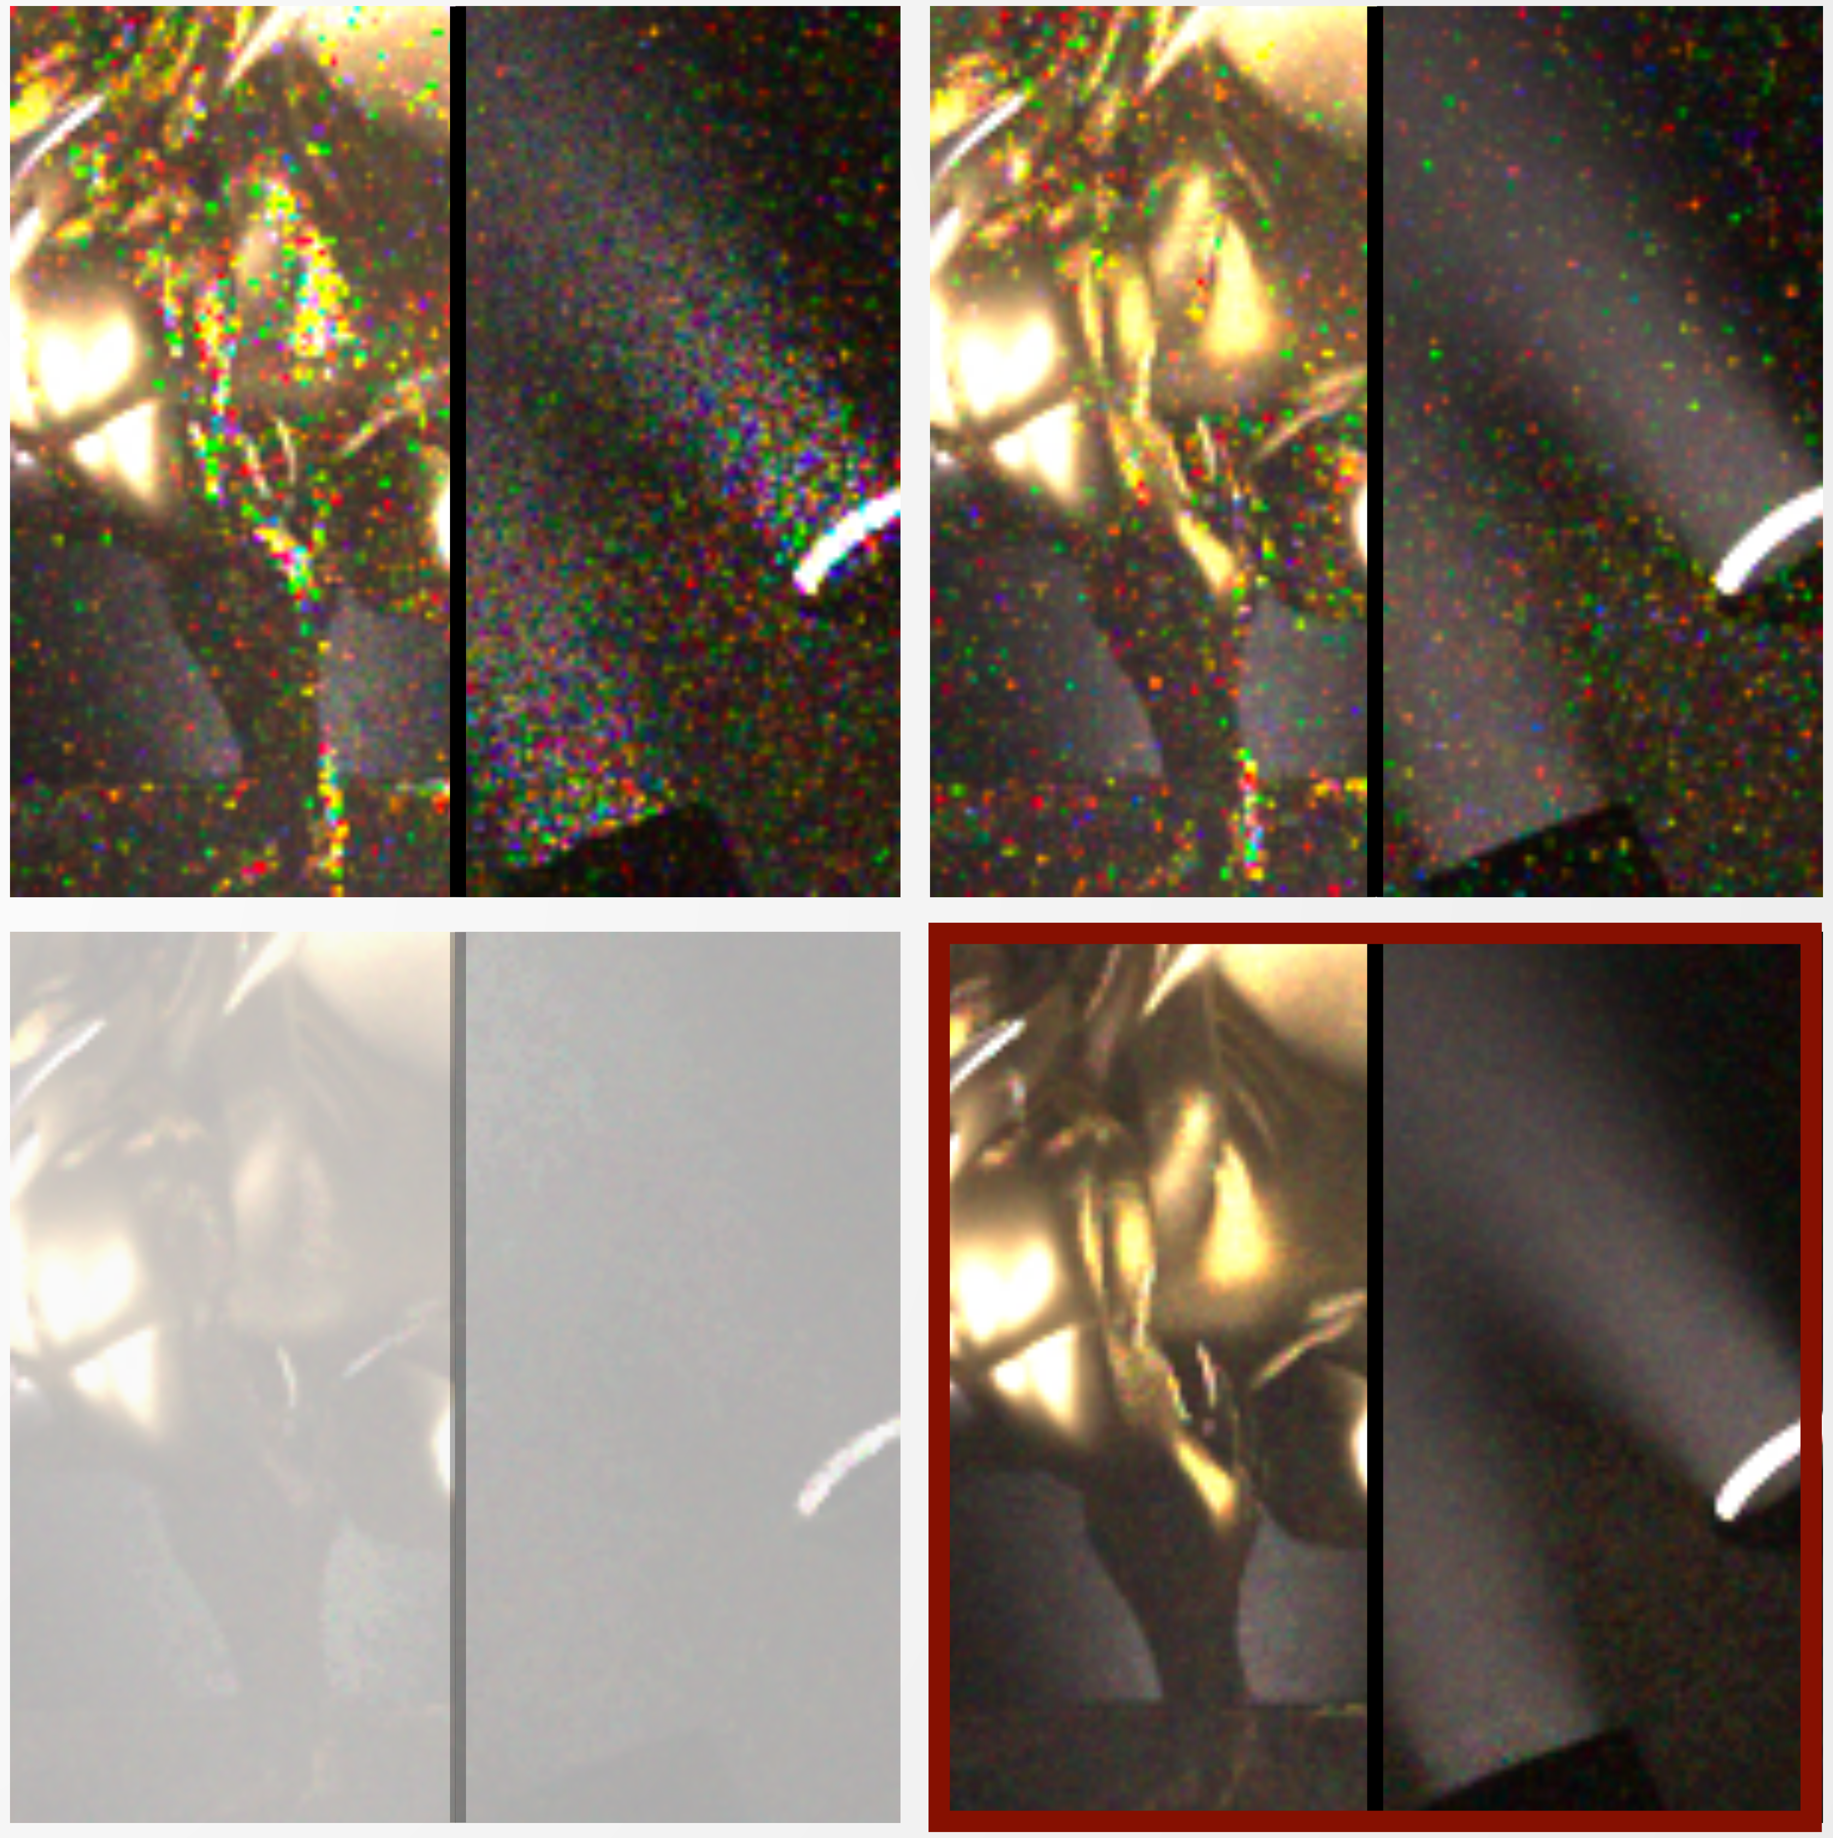

Guided PT+DBOR

Dragon - 10h

Path tracing + DBOR

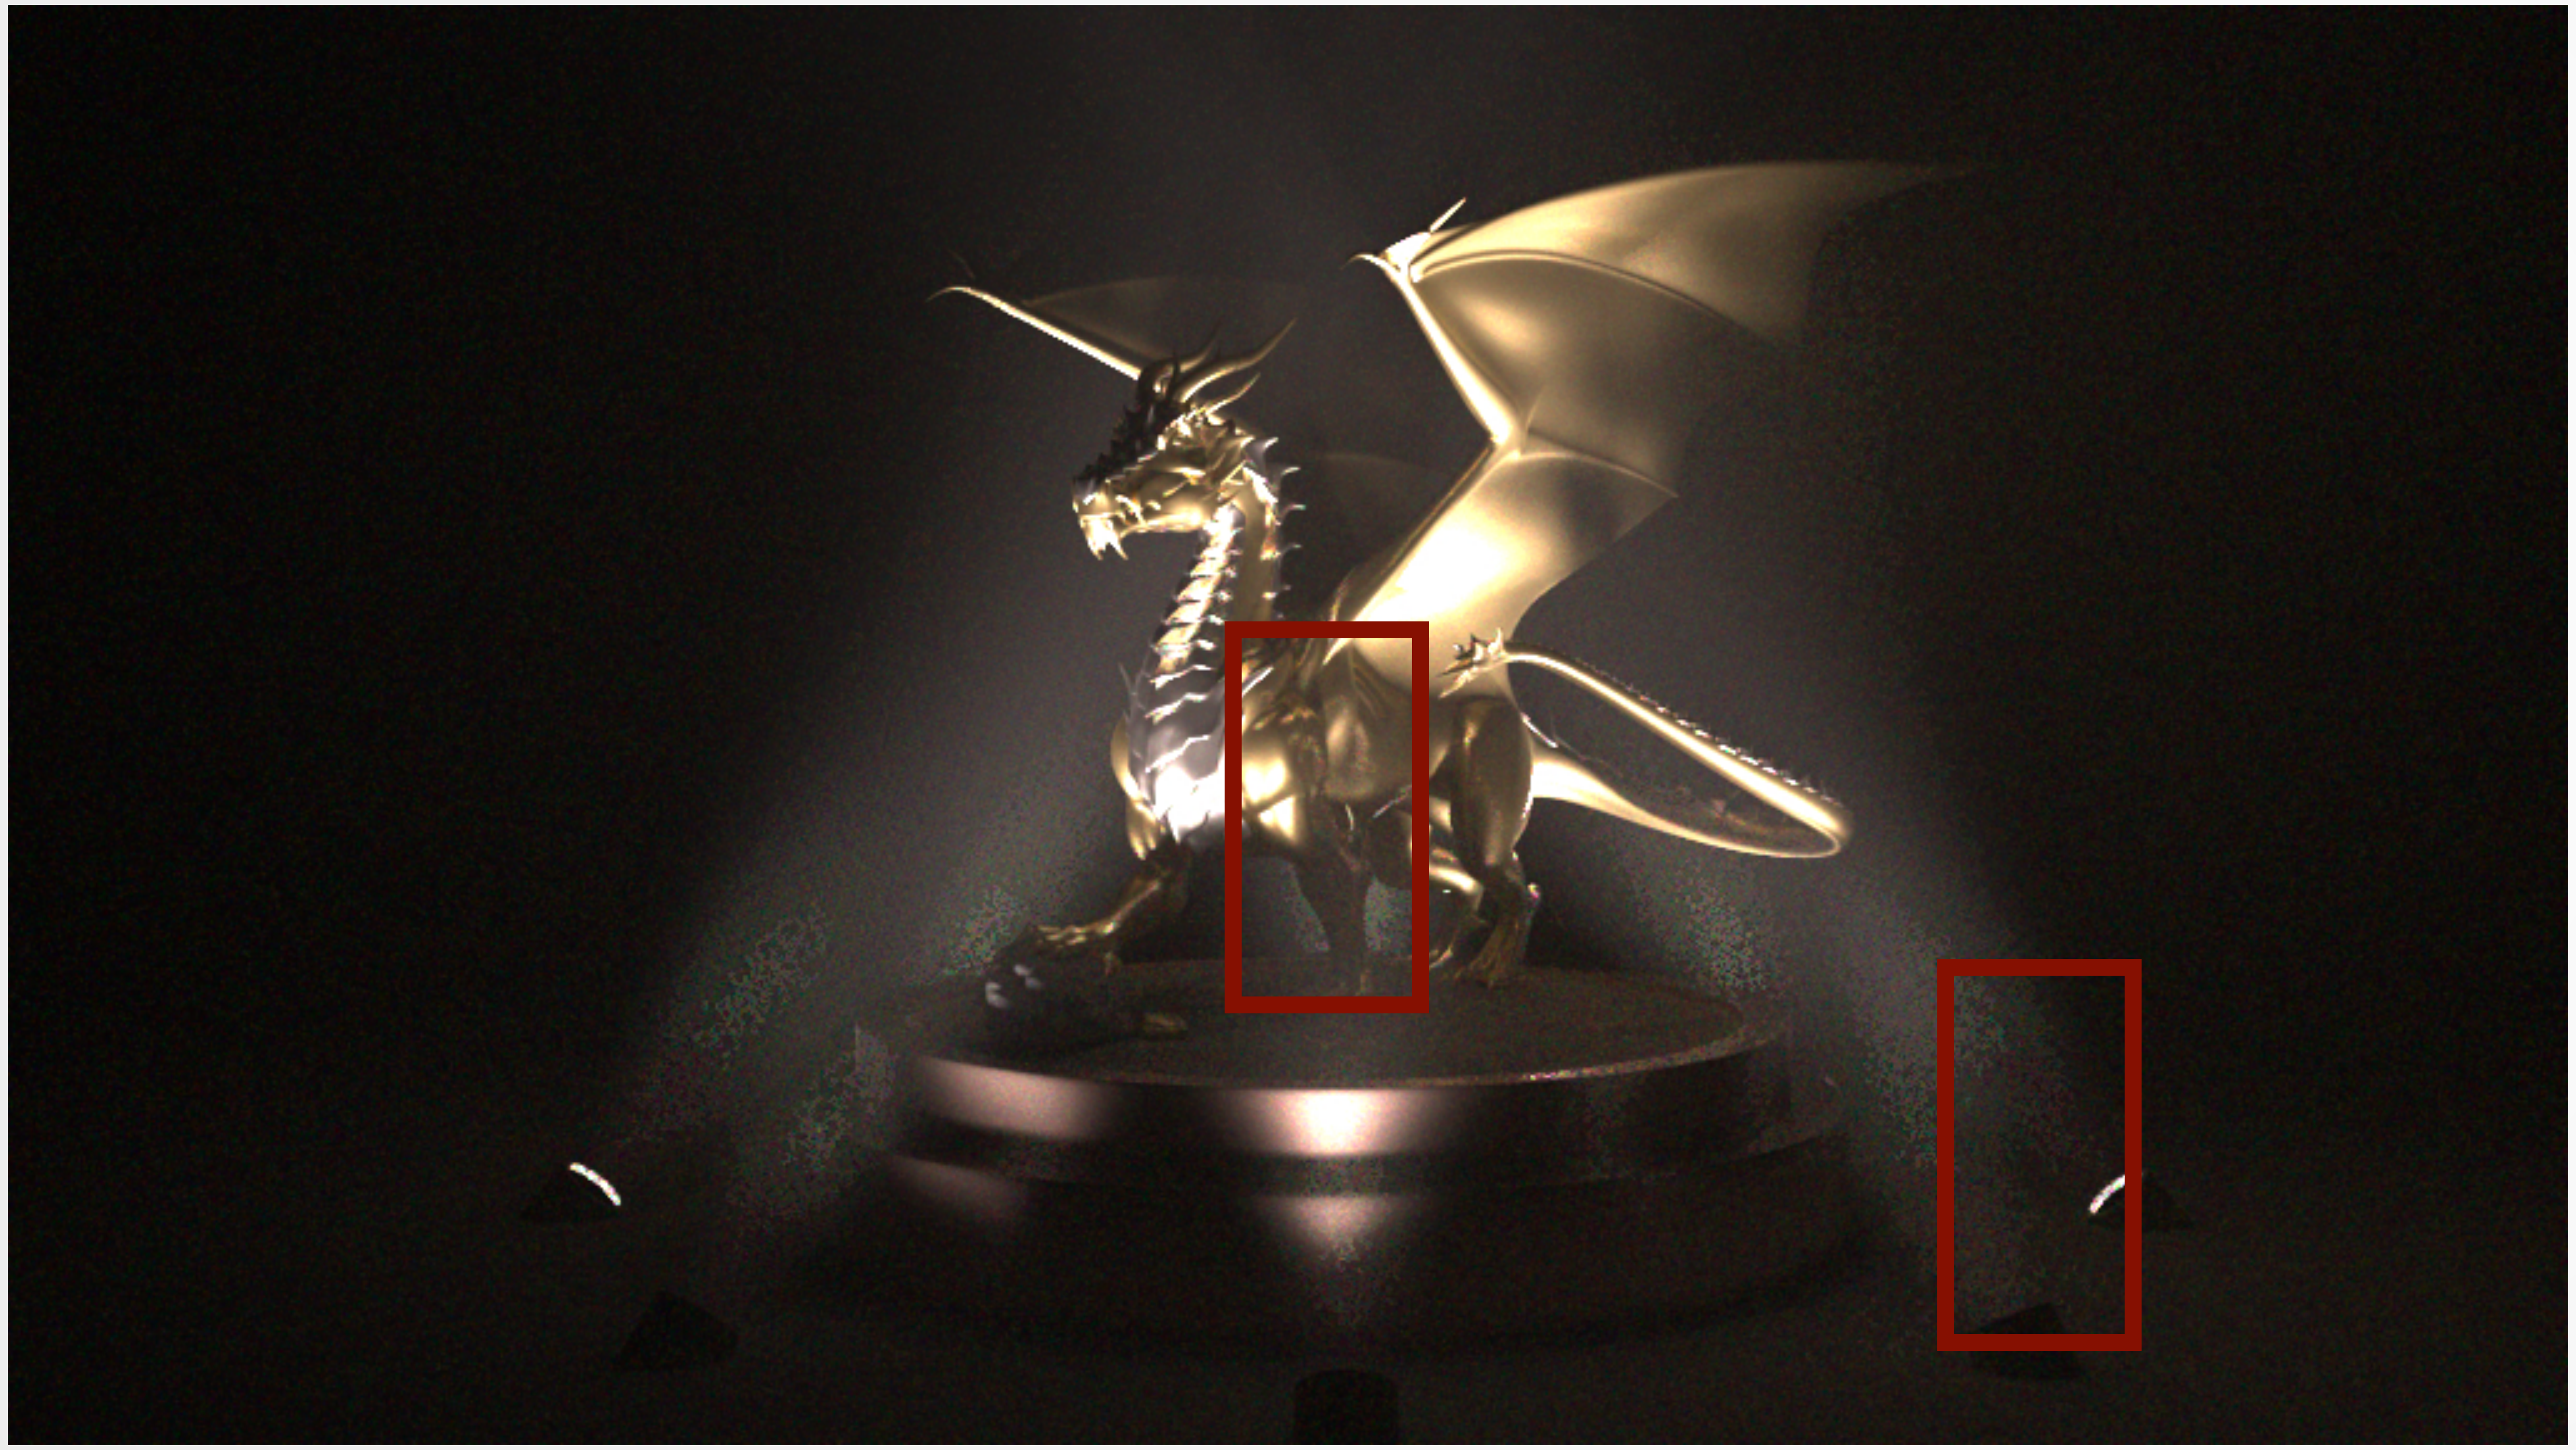

PT

Guided PT

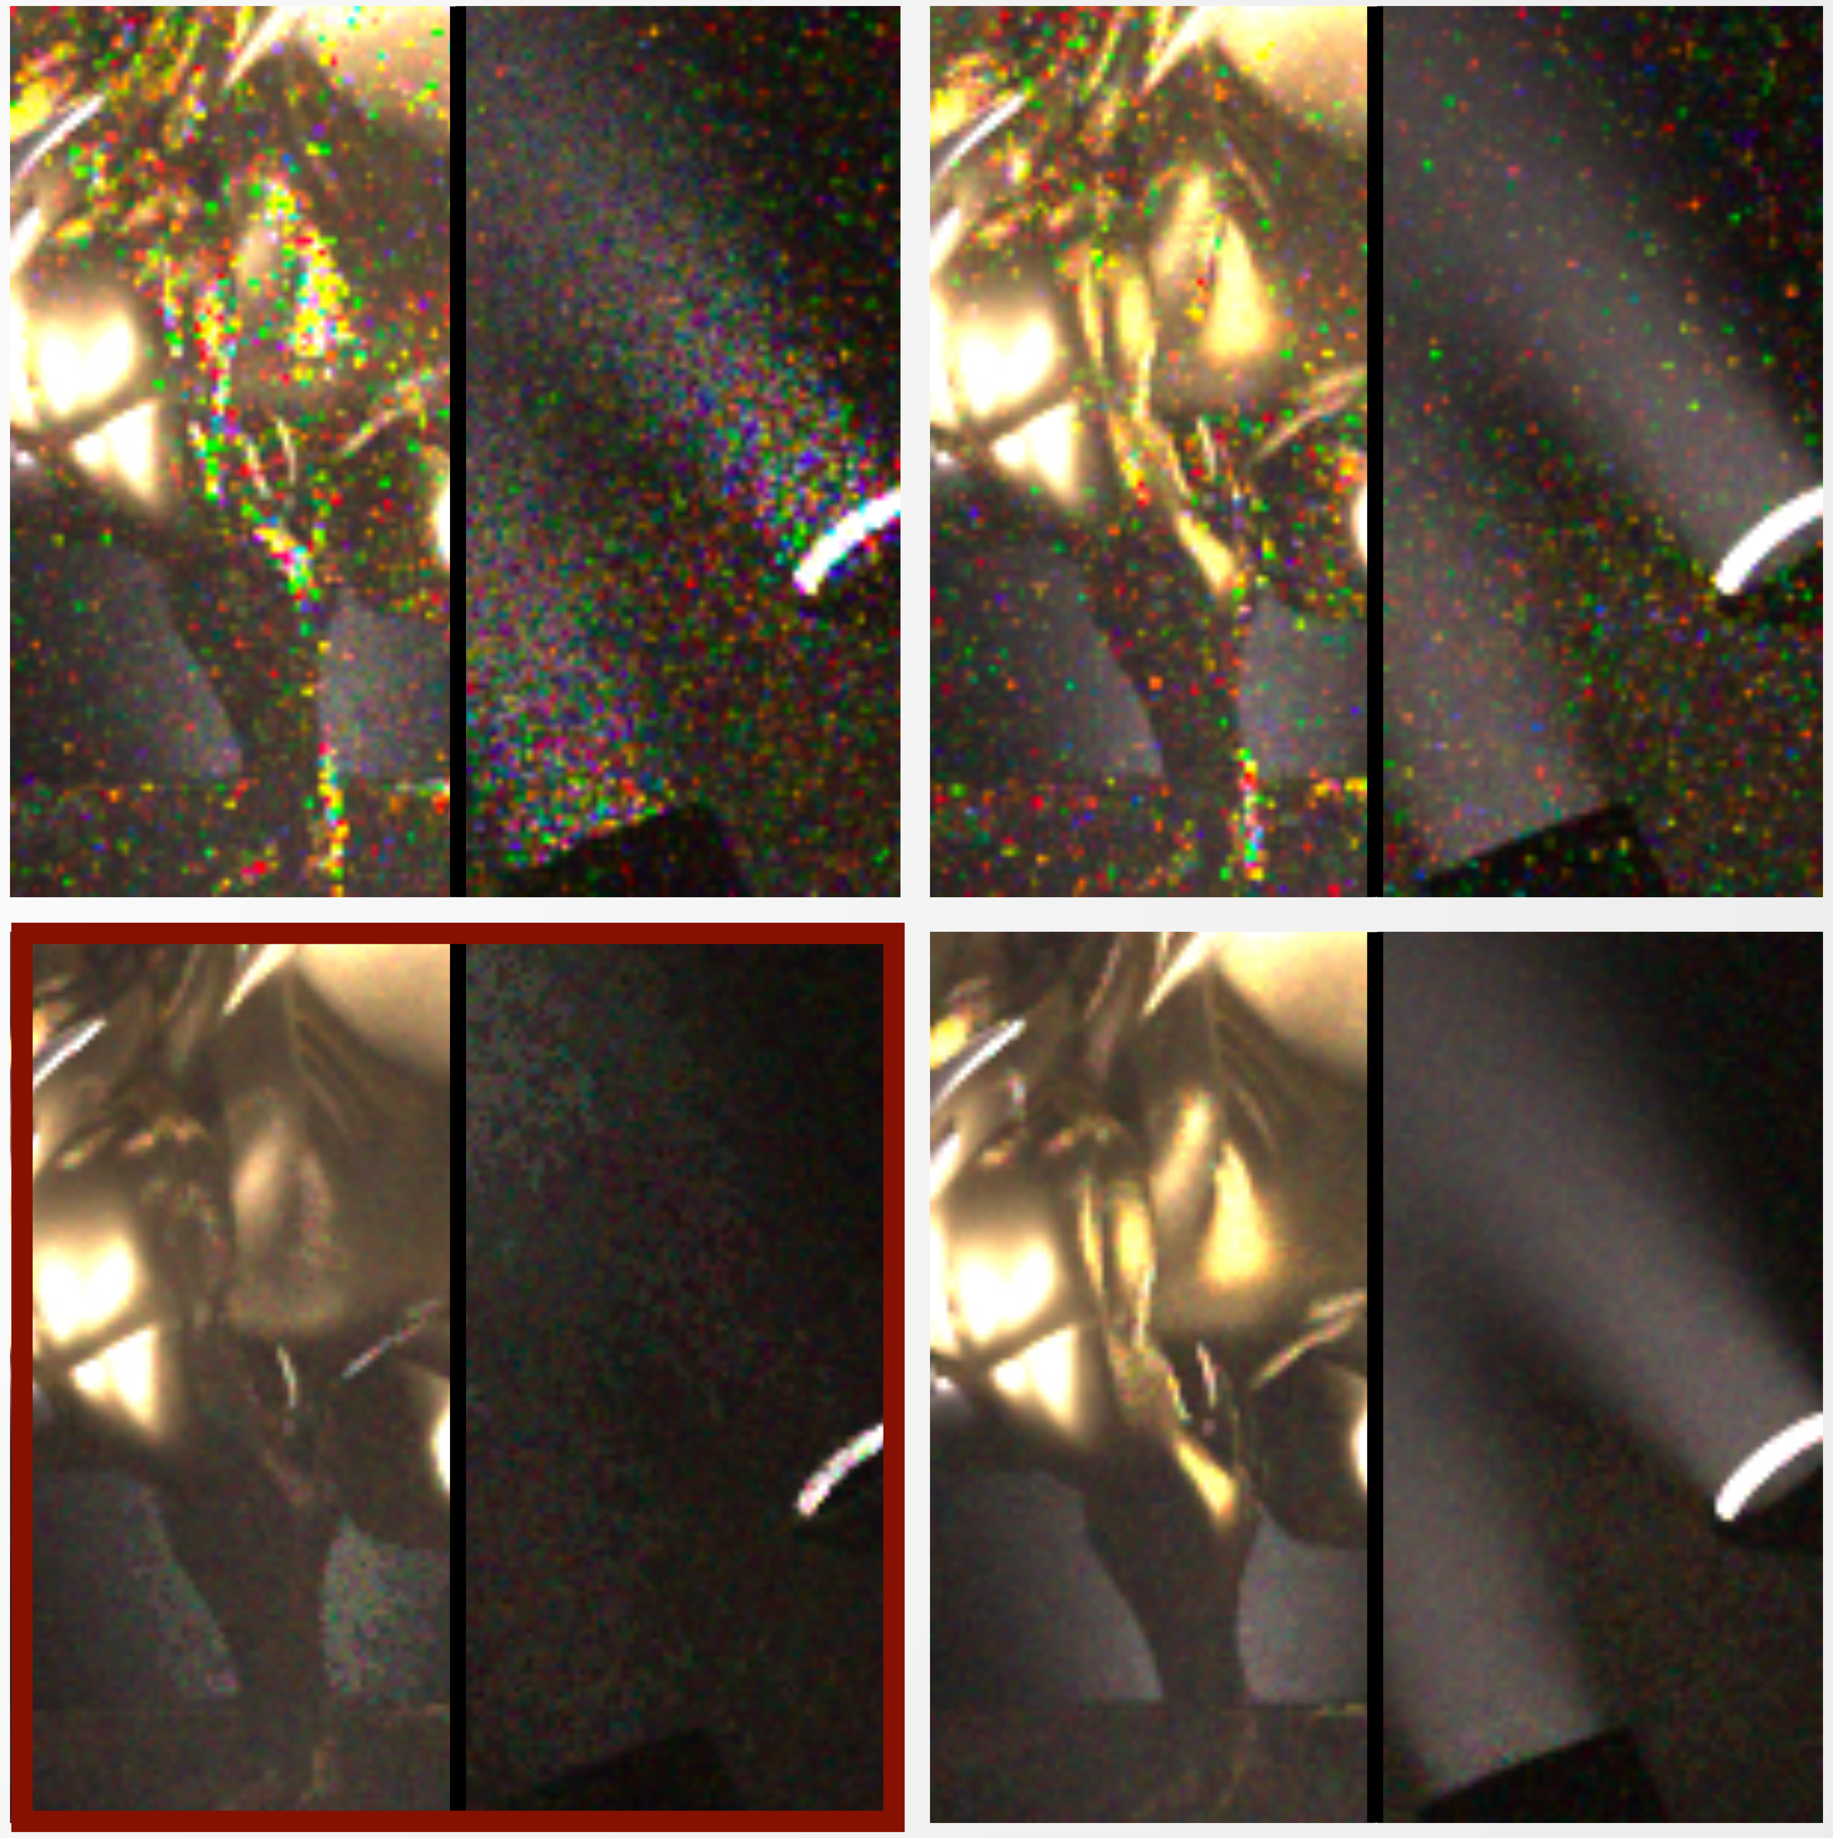

PT + DBOR

Guided PT+DBOR

Results

- More results in the paper/supplemental document

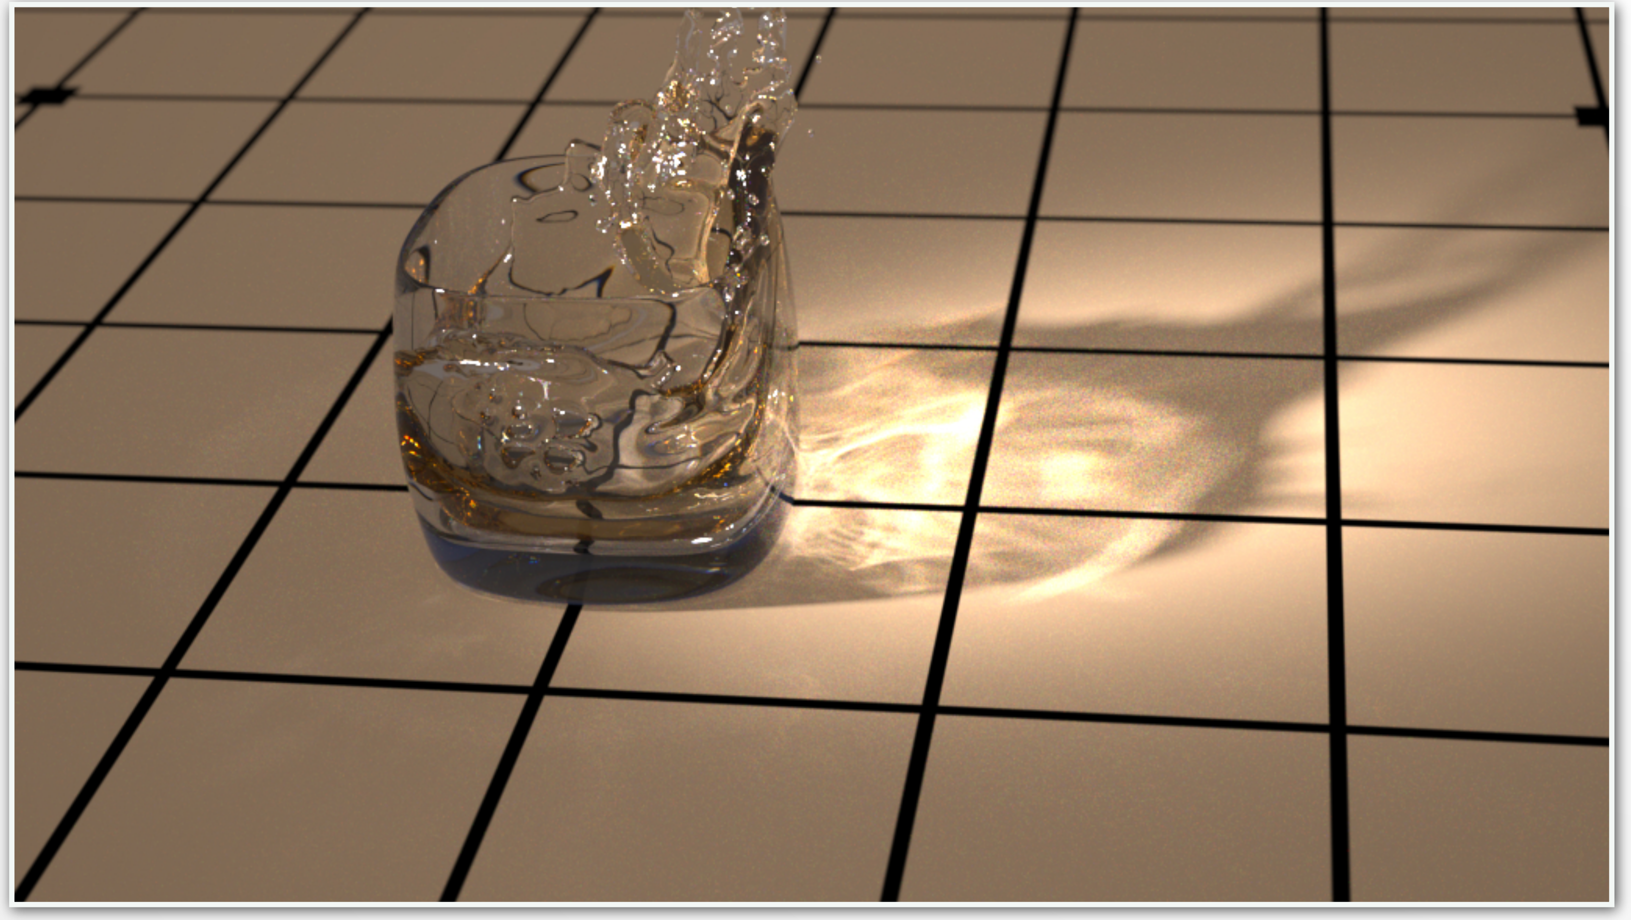

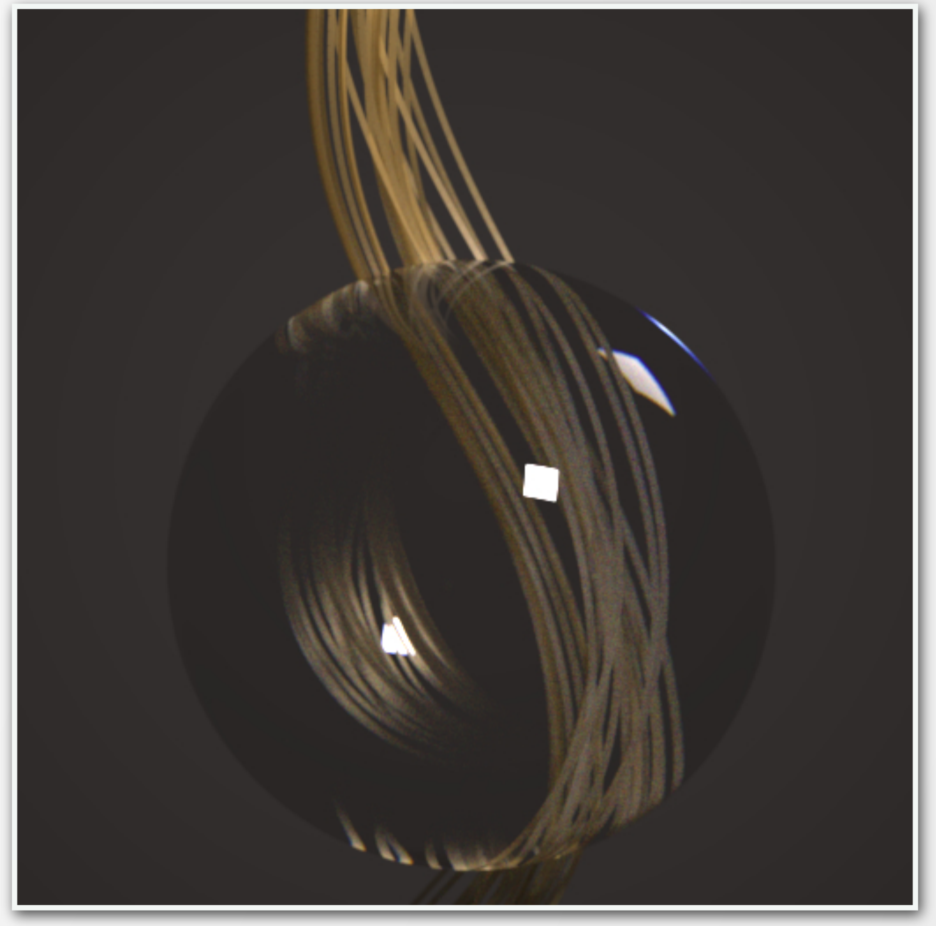

Limitations & Future Work

- 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

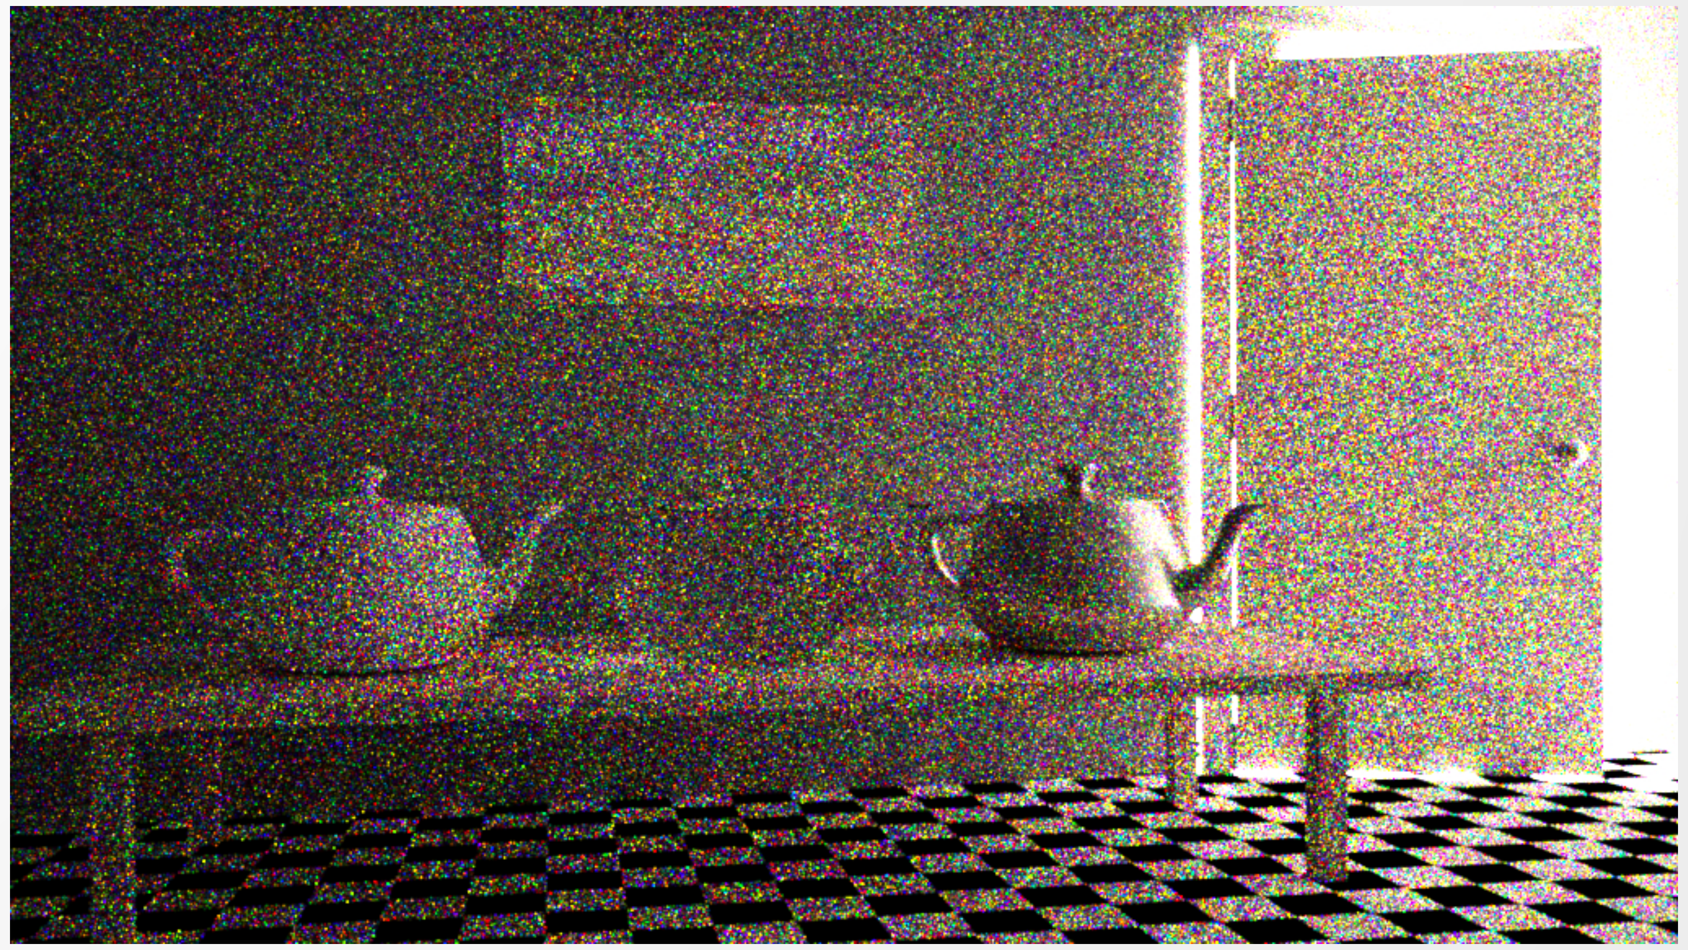

4405spp, RMSE 1.30

Limitations & Future Work

- Temporal stability is challenging

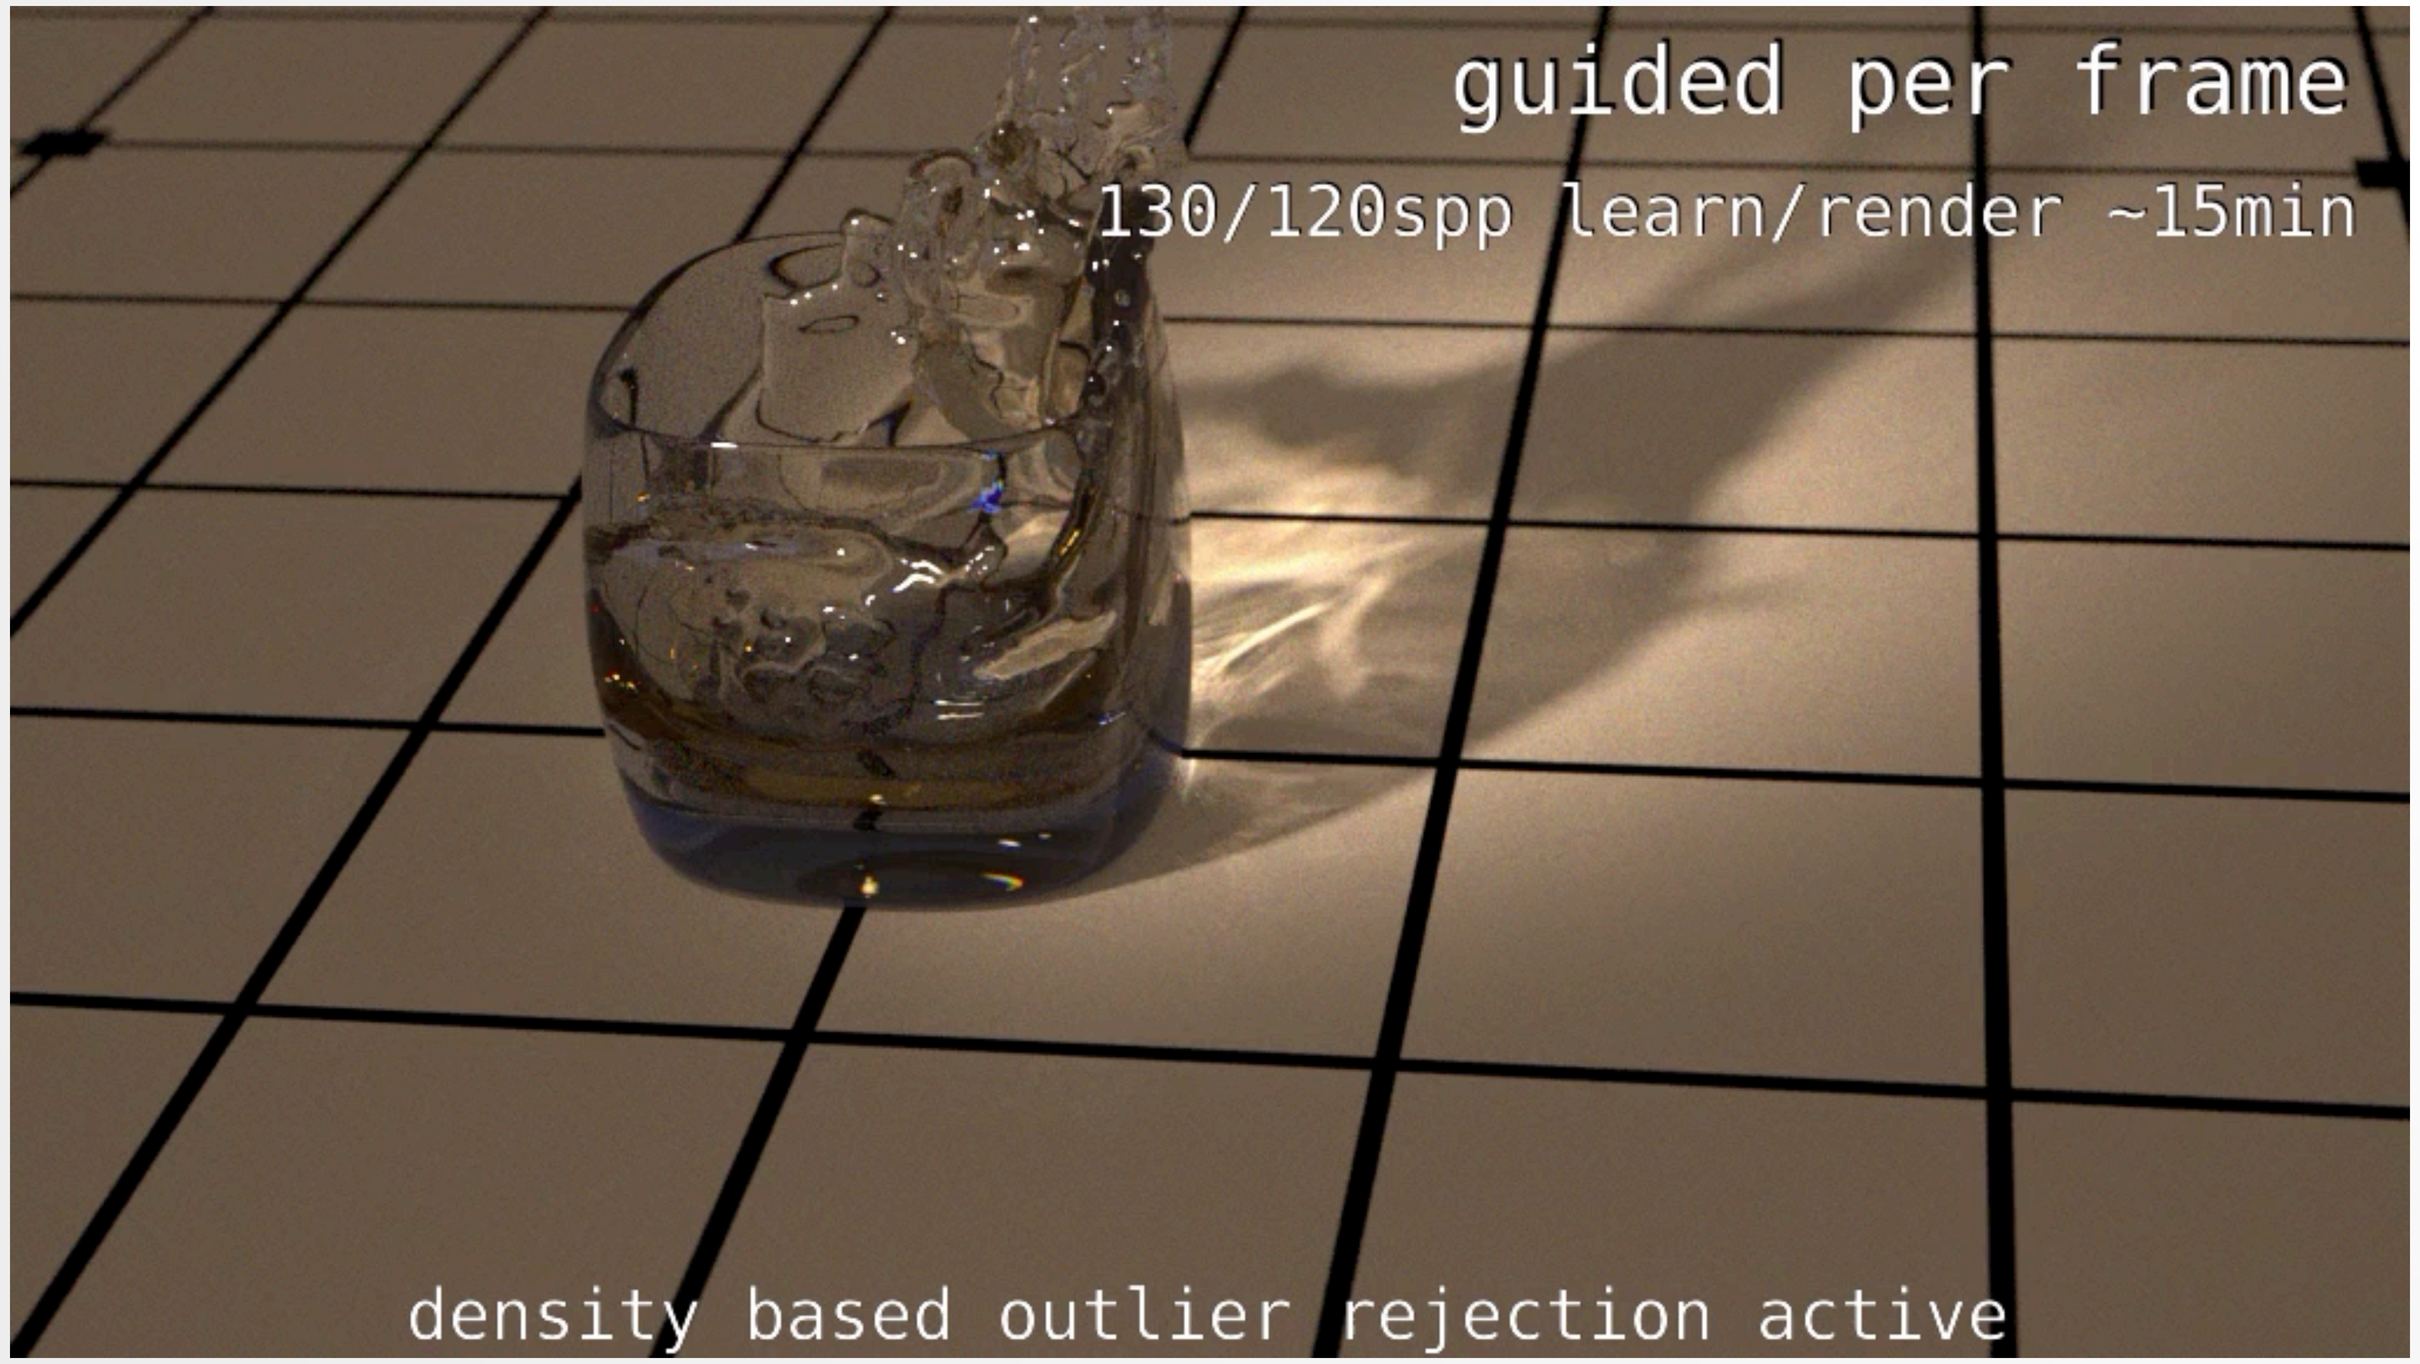

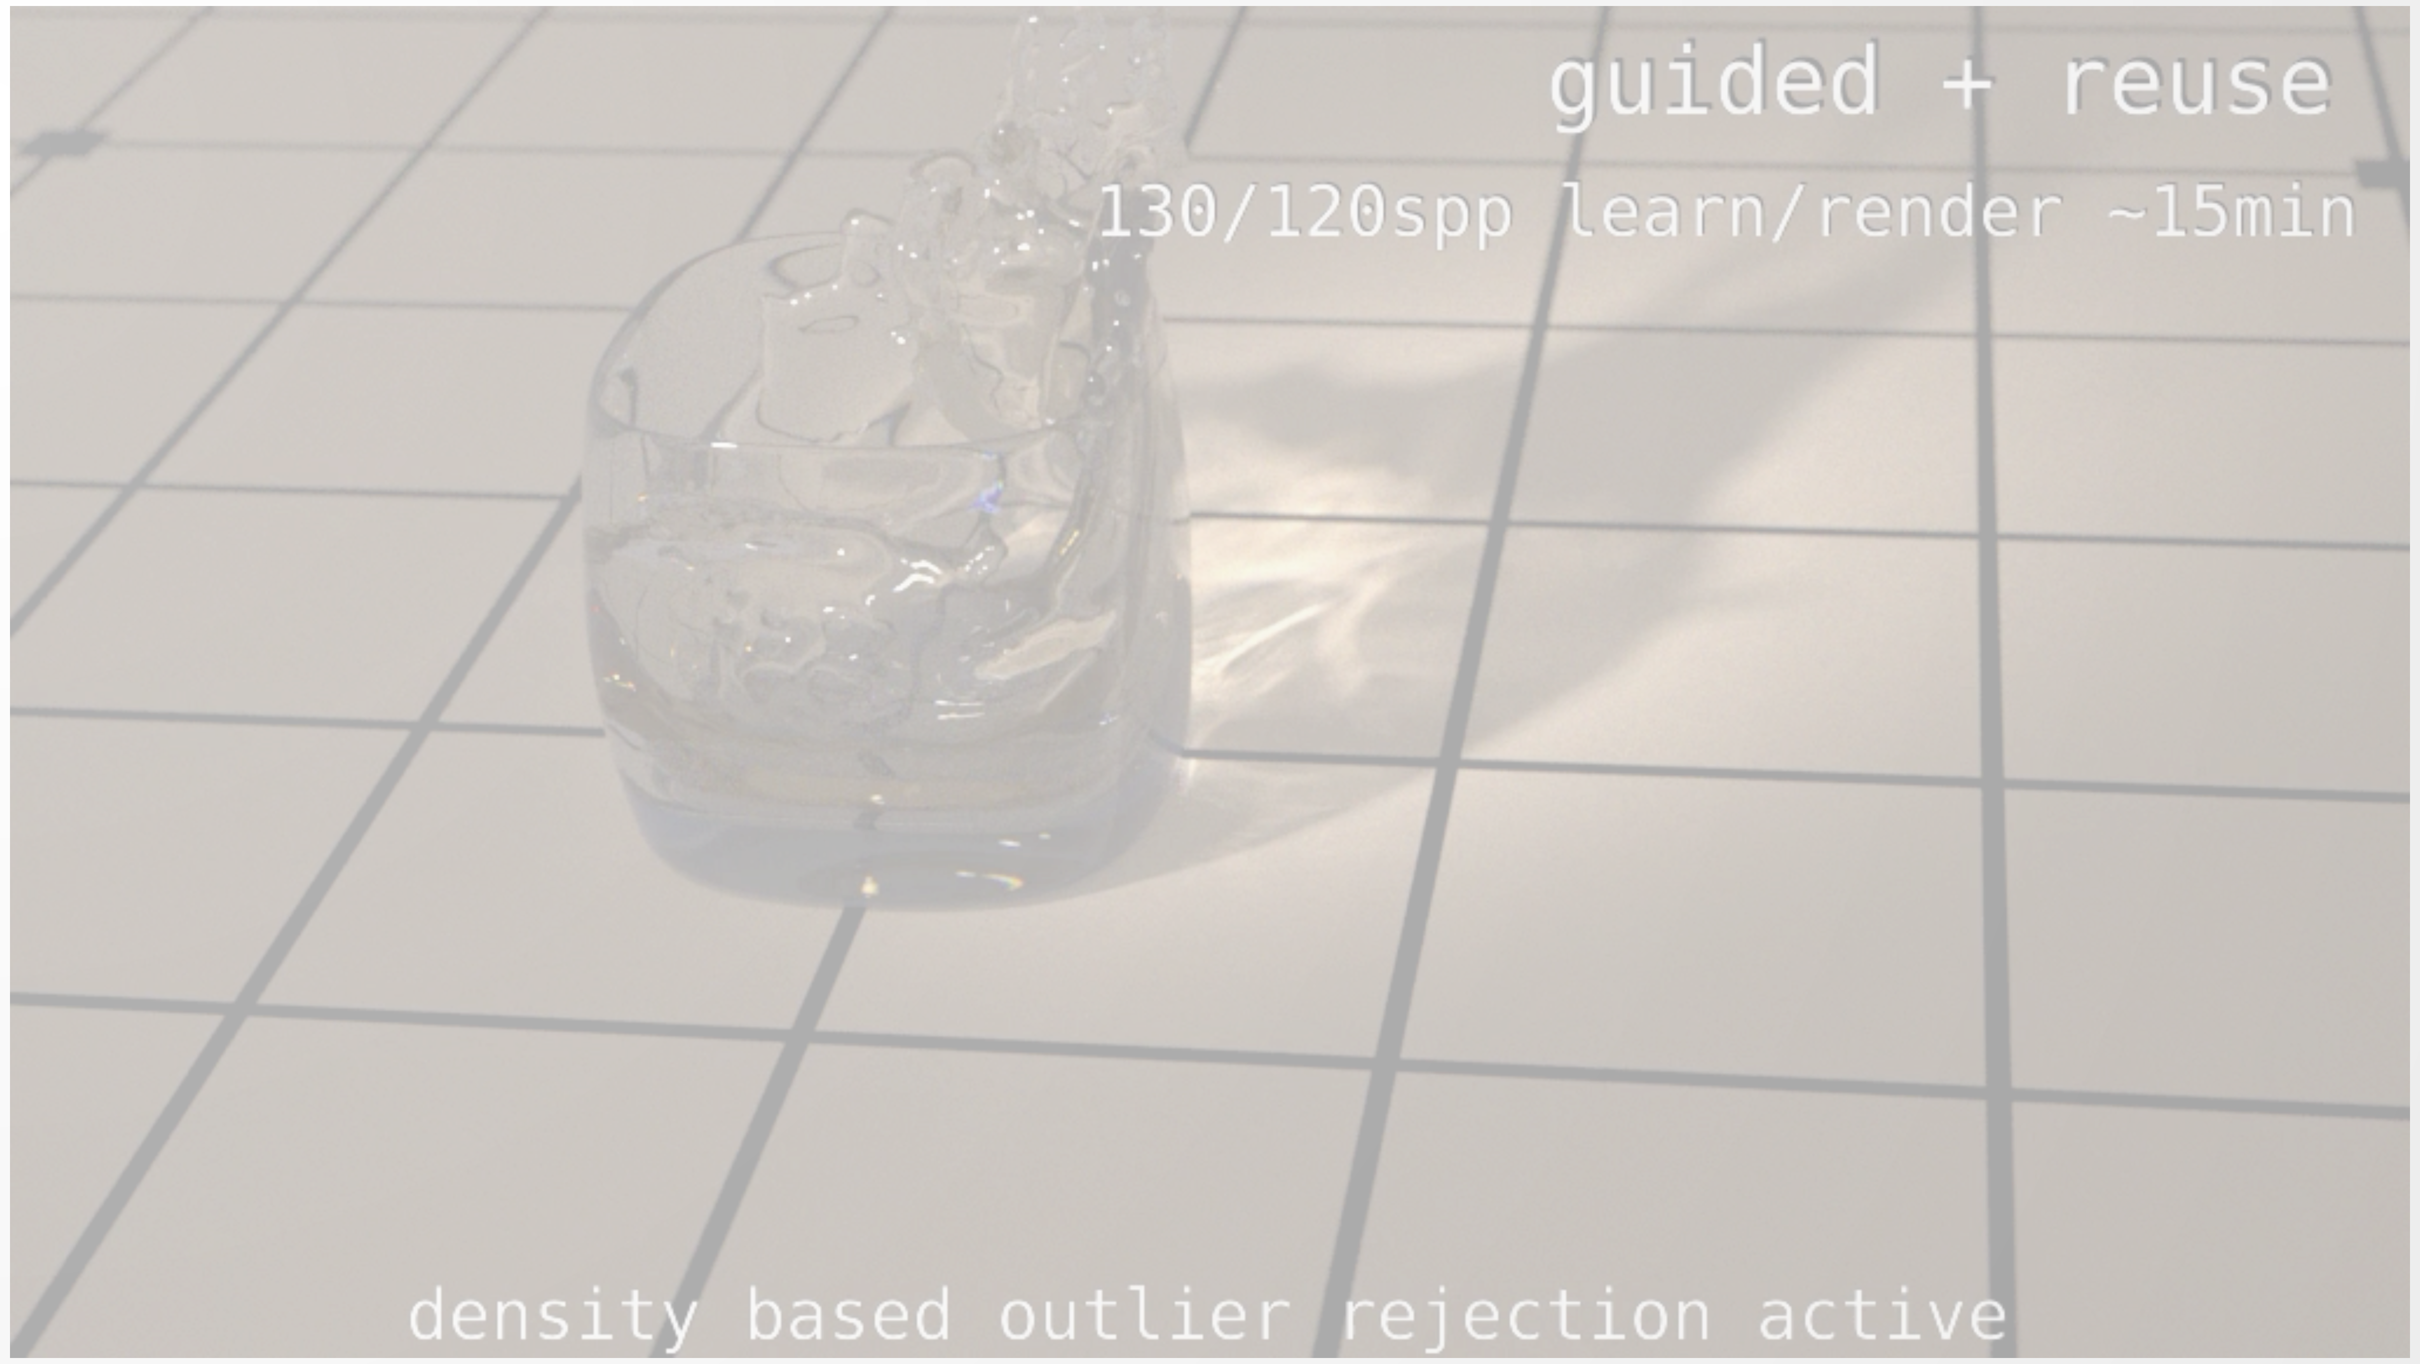

Karlsruhe Institute of Technology

Limitations & Future Work

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

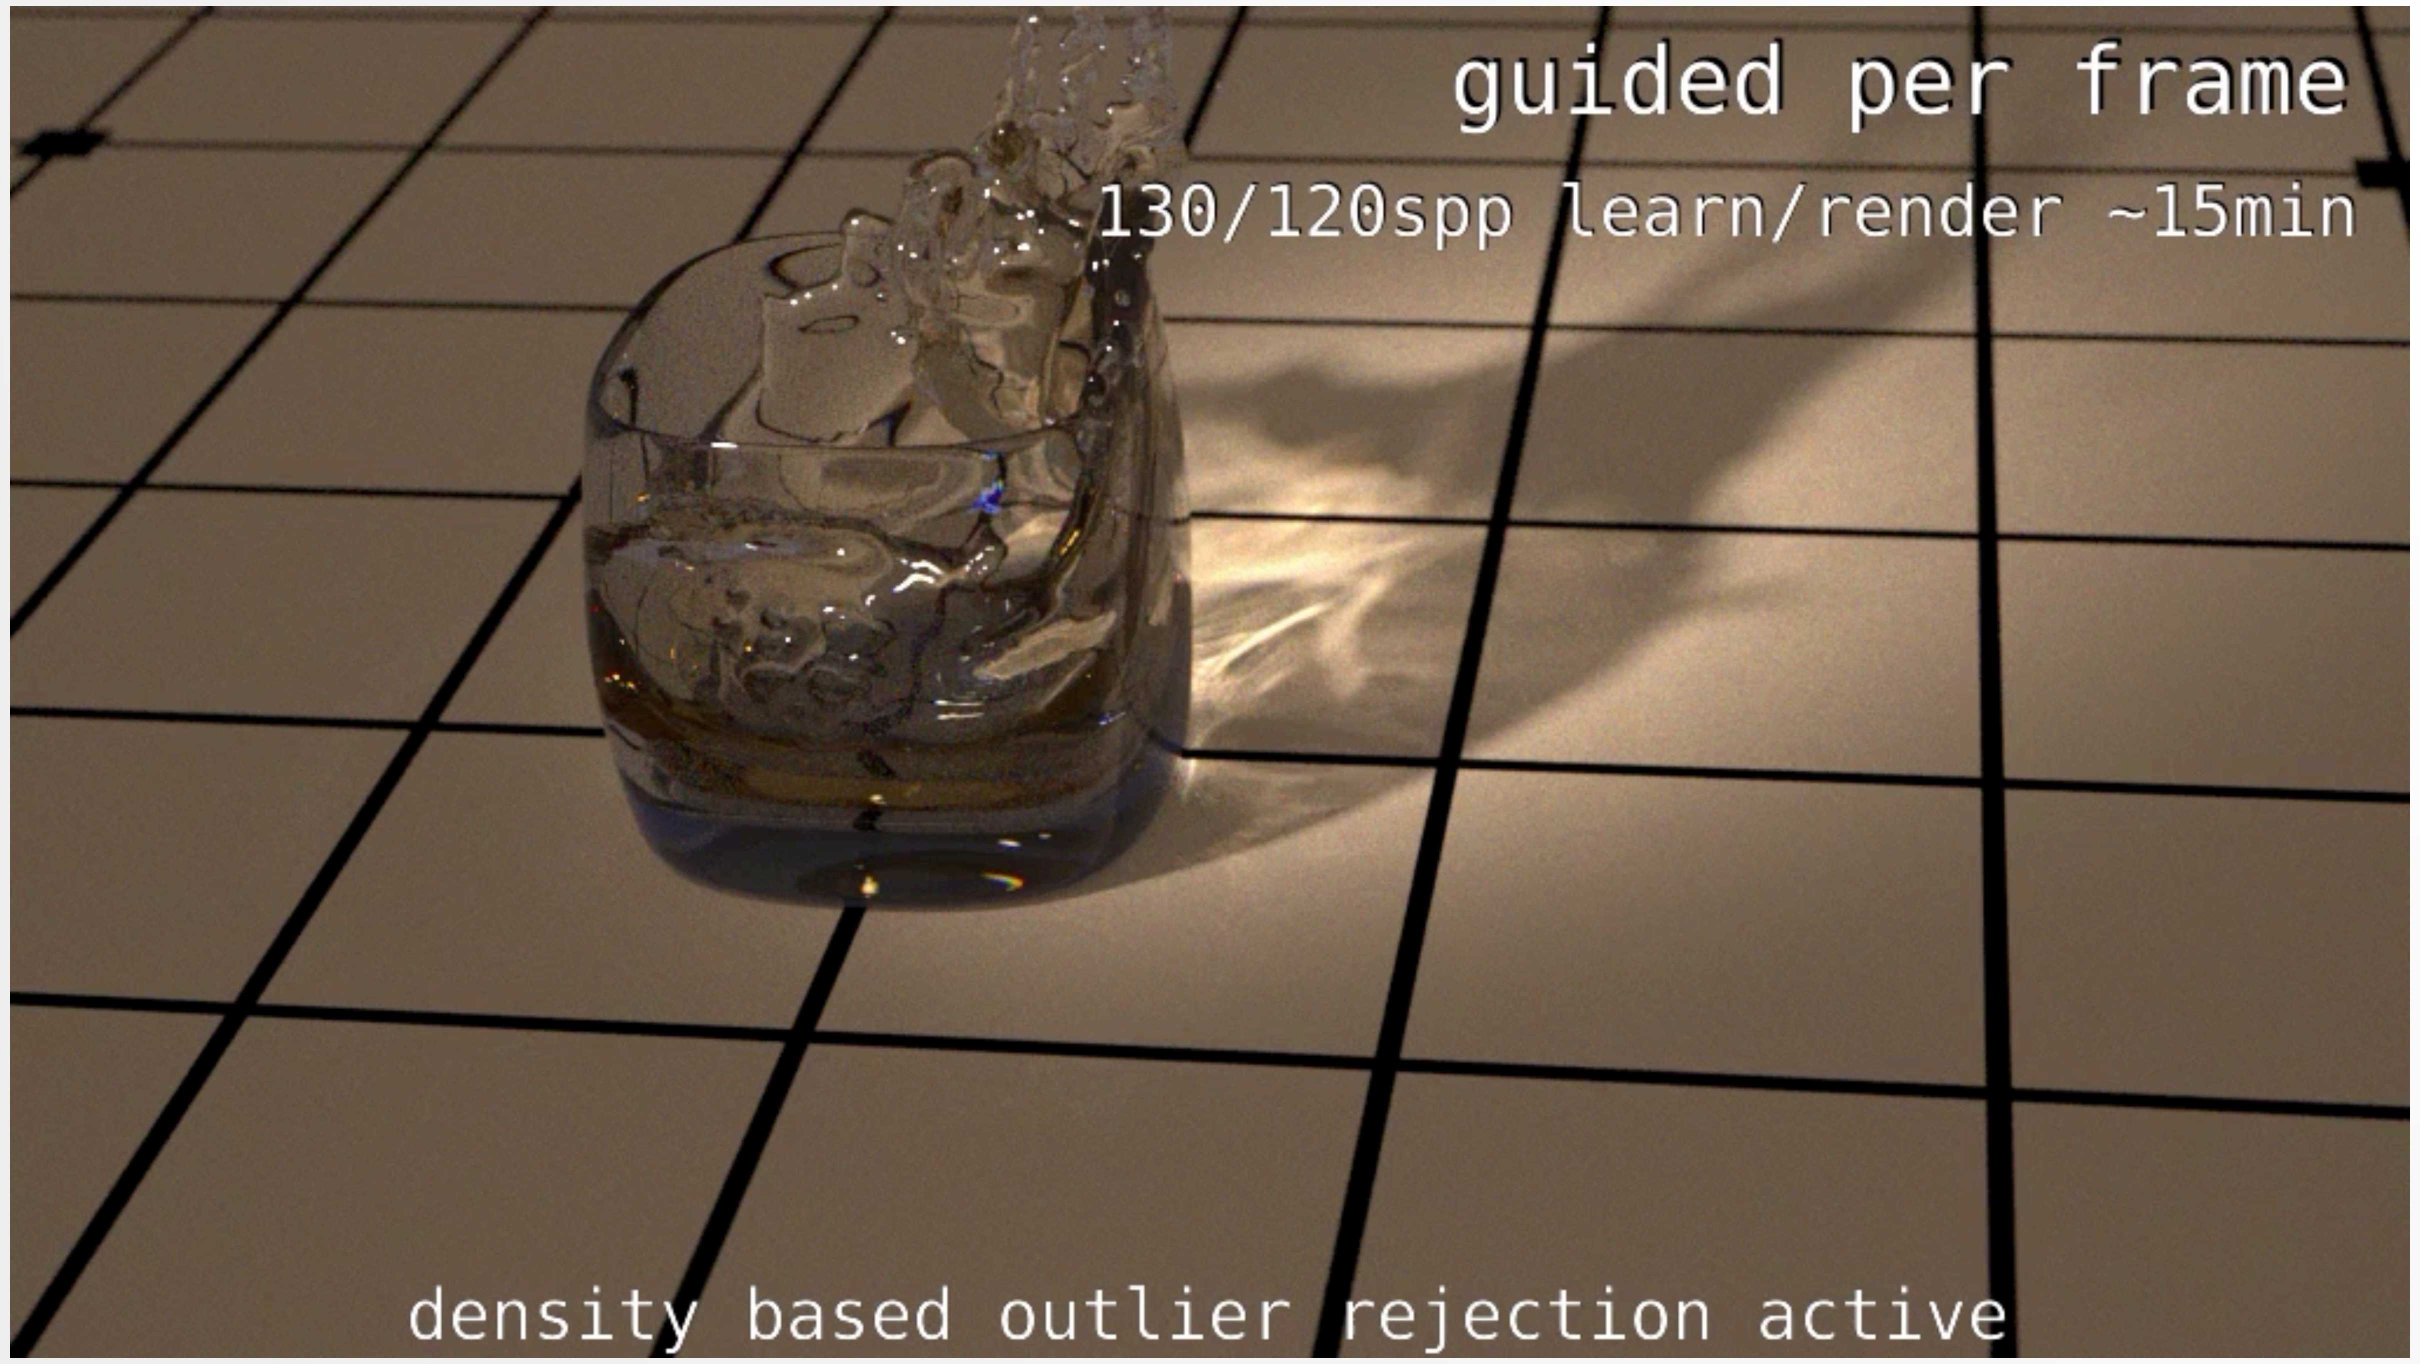

independent guide path cache

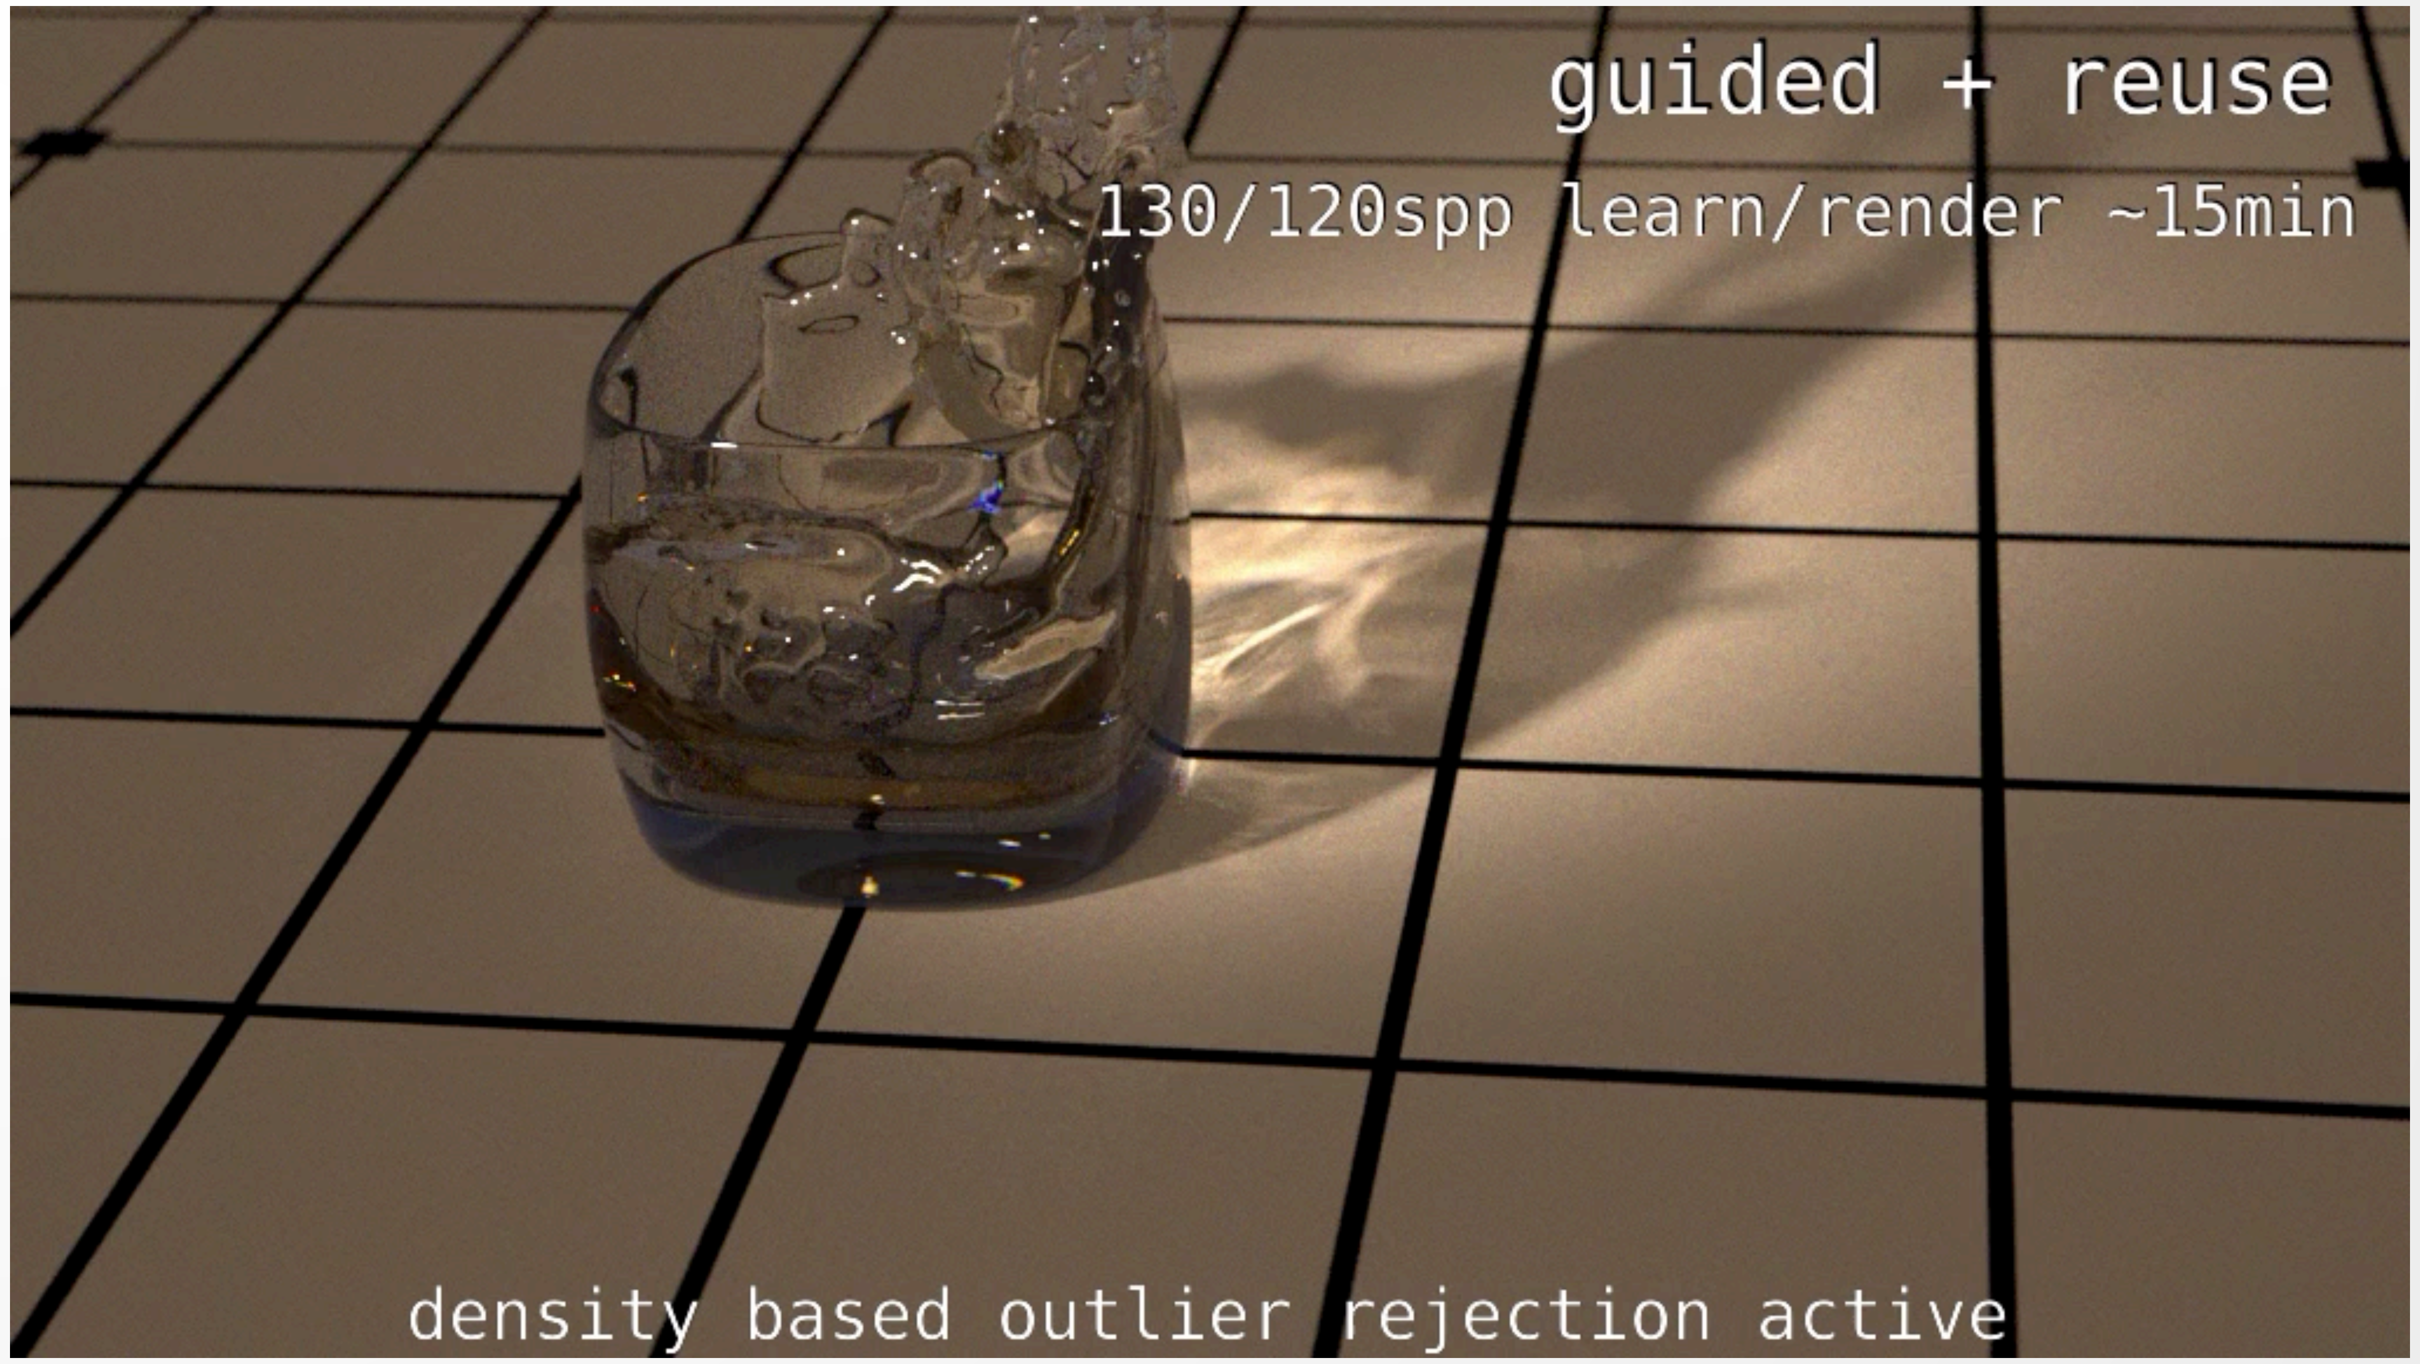

resampled guide path cache

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

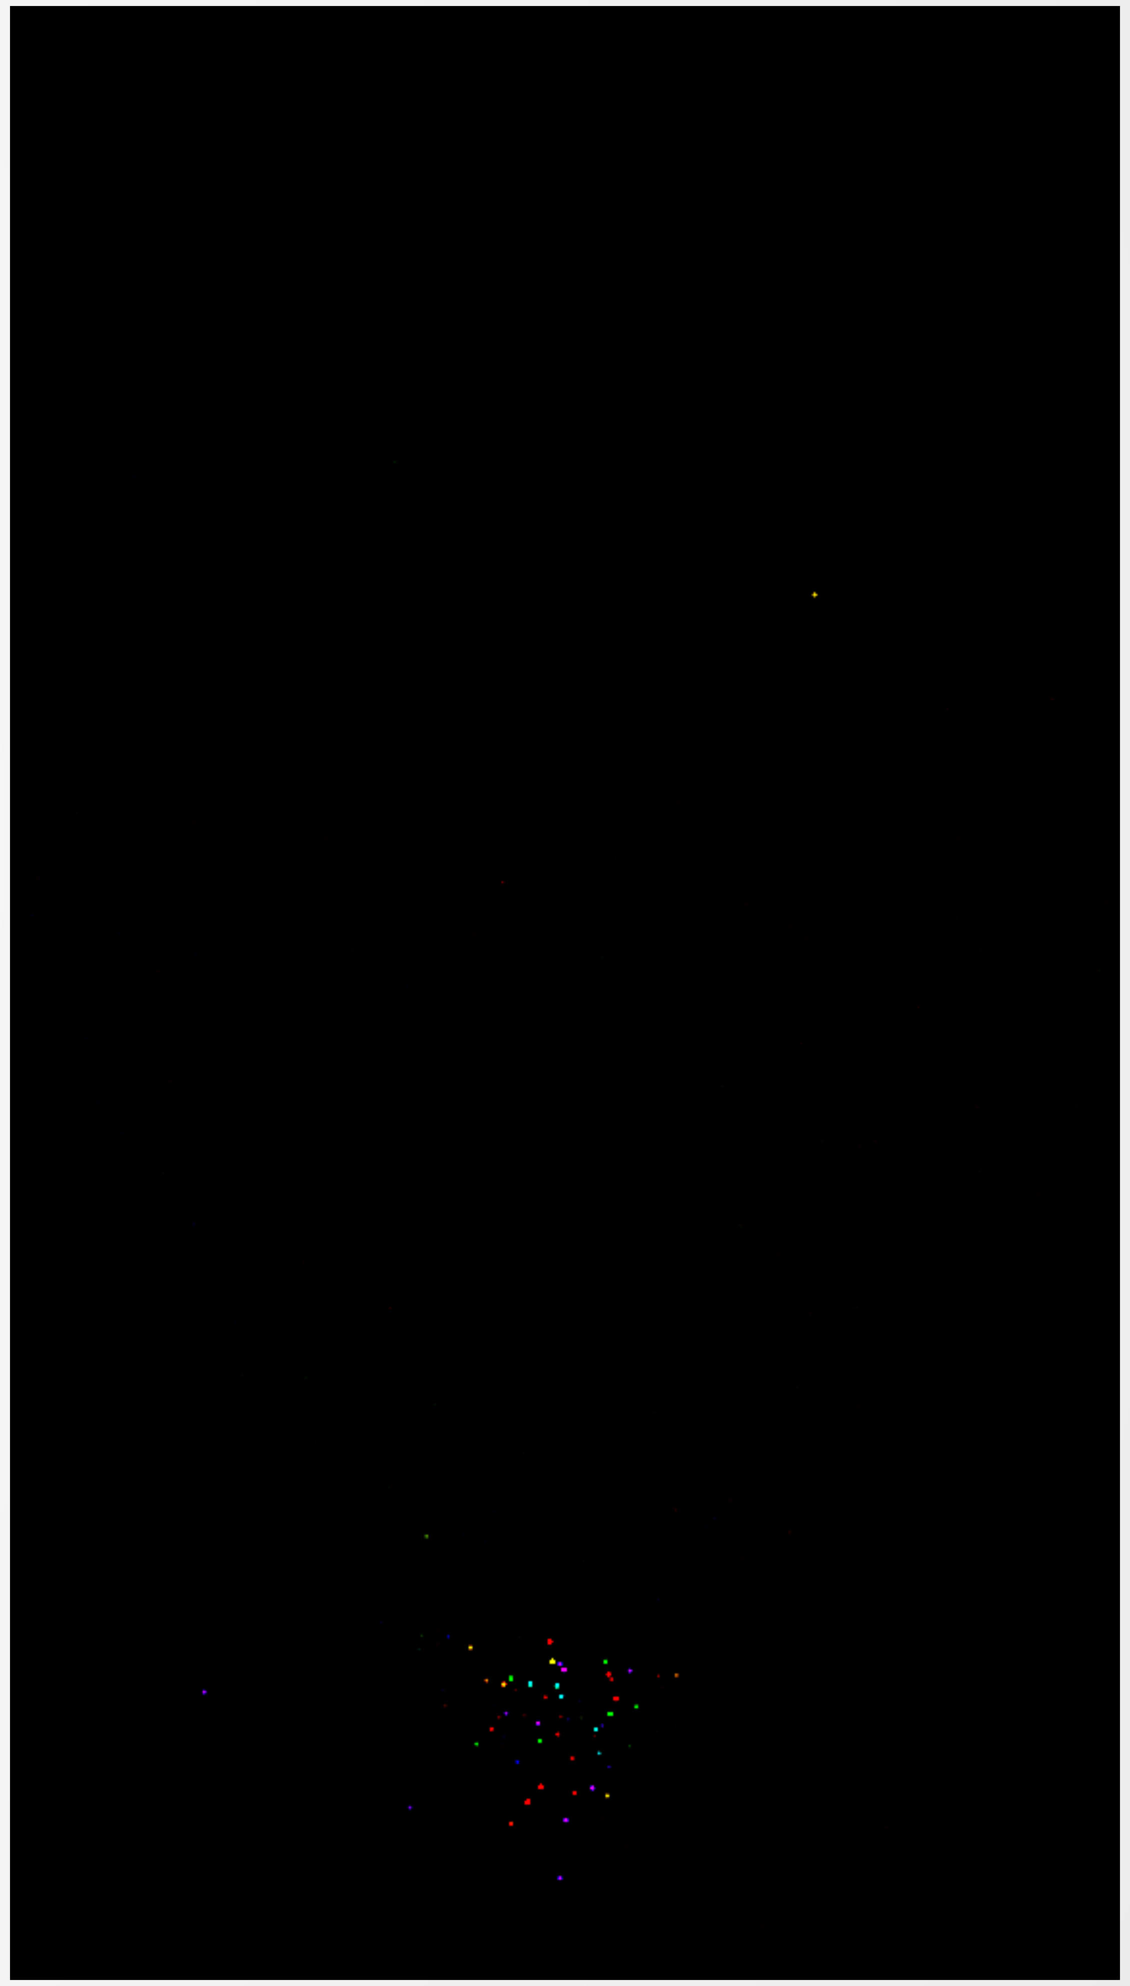

Sponsored by

SIGGRAPH ASIA 2018 TOKYO

CONFERENCE 4 – 7 December 2018
EXHIBITION 5 – 7 December 2018
Tokyo International Forum, Japan
SA2018.SIGGRAPH.ORG

Thank you!

Discarded Ideas

- Constant Kernel
- Repeated 3D search for NN
- Metropolis Hastings to resample guide paths
- Relaxation of guide paths (photon map-style)

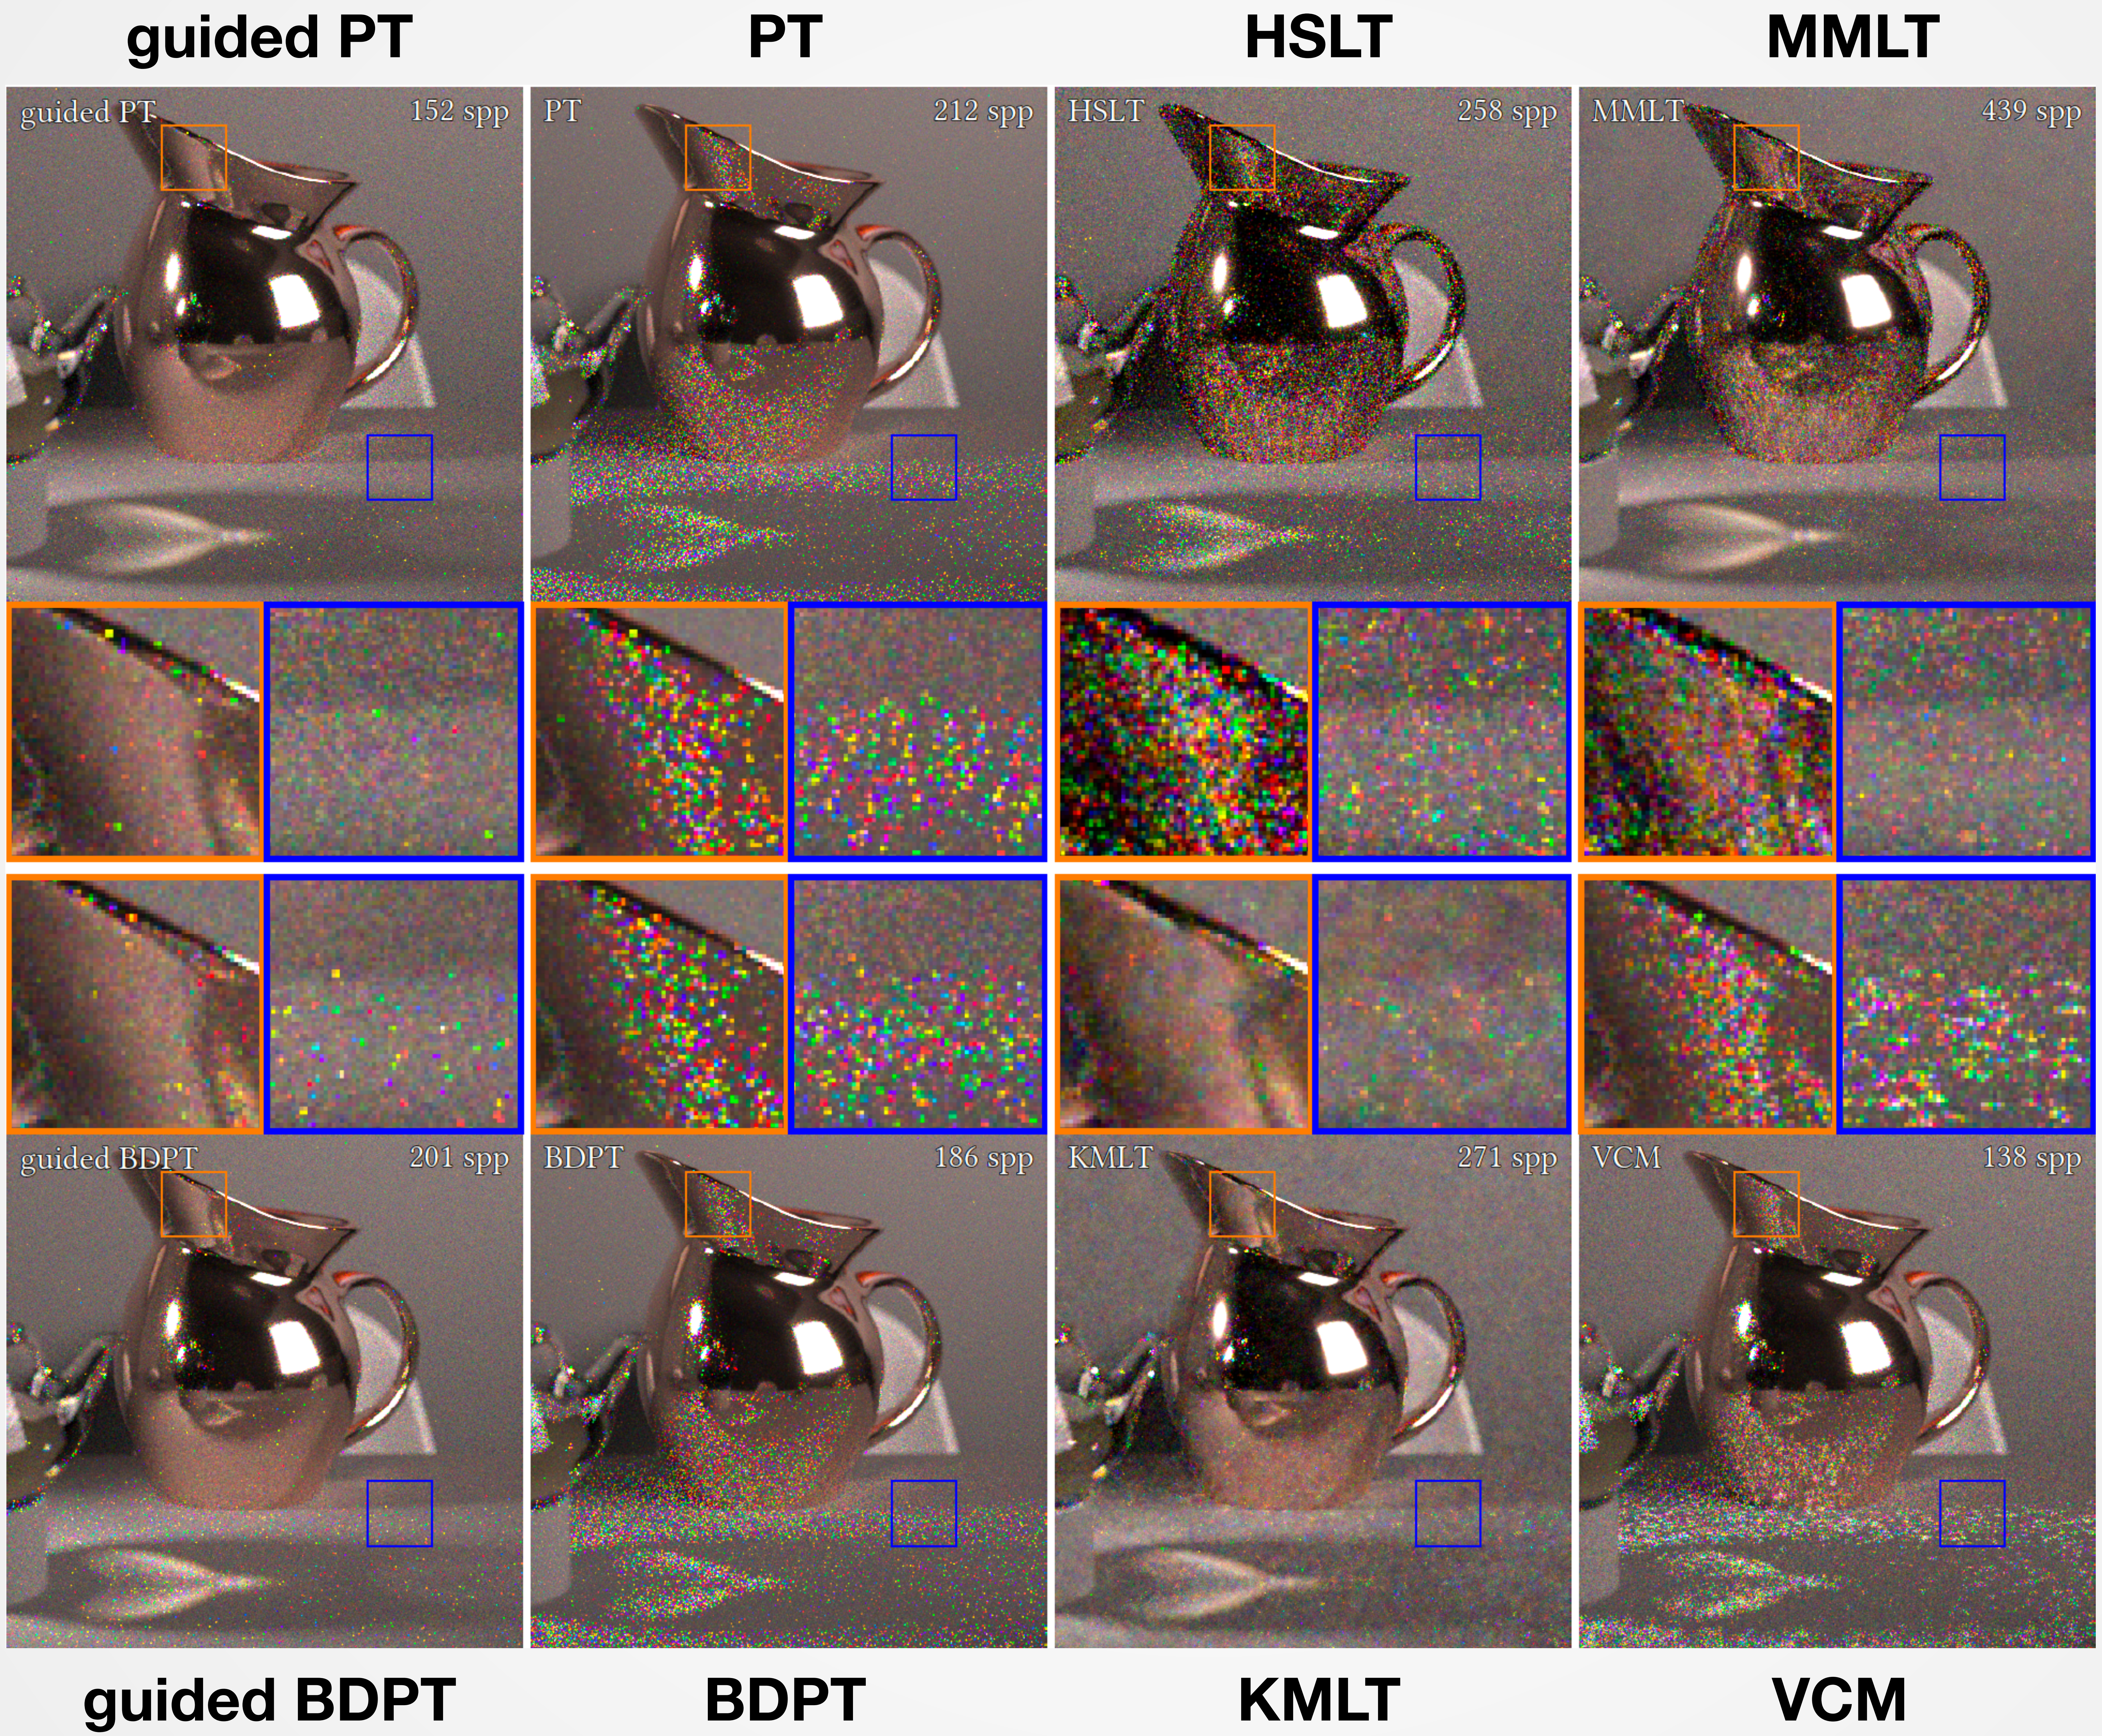
guided PT
PT
HSLT
MMLT
guided PT
152 spp
PT
212 spp
HSLT
258 spp
MMLT
439 spp
guided BDPT
201 spp
BDPT
186 spp
KMLT
271 spp
VCM
138 spp
guided BDPT
BDPT
KMLT
VCM

Recent Guiding Methods

Vorba et al. [2014]

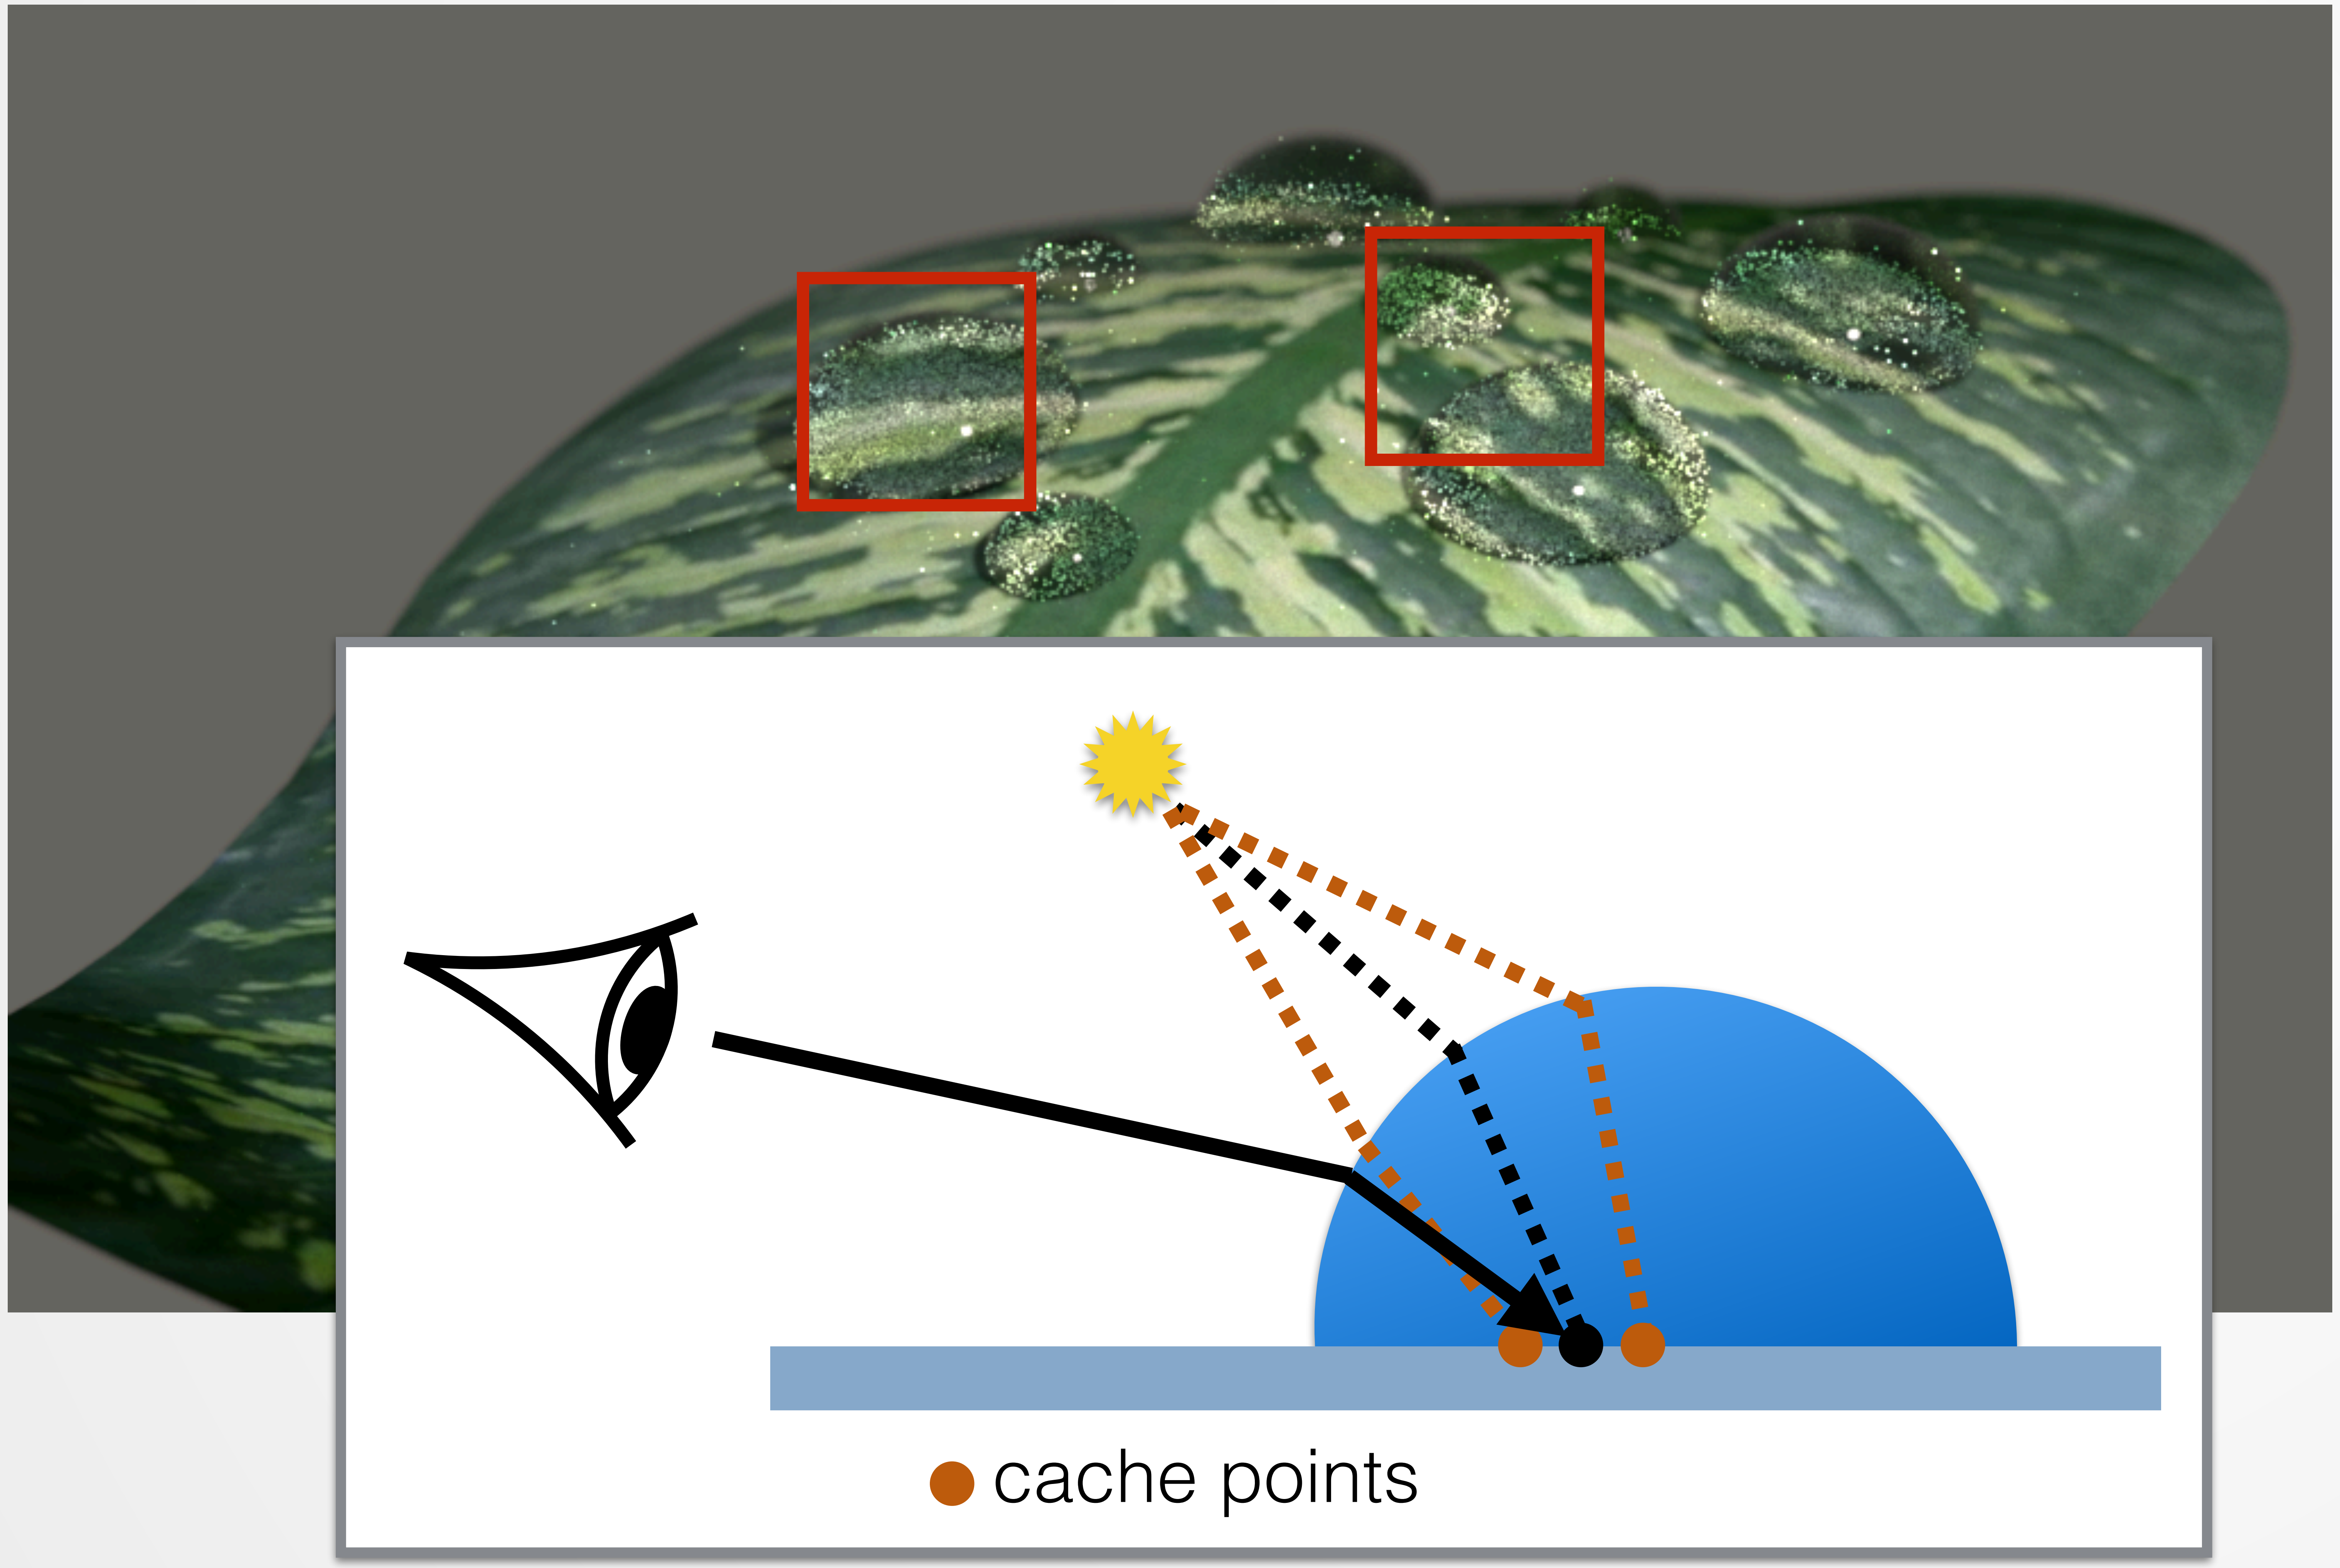

Vorba et al. [2014] | ours

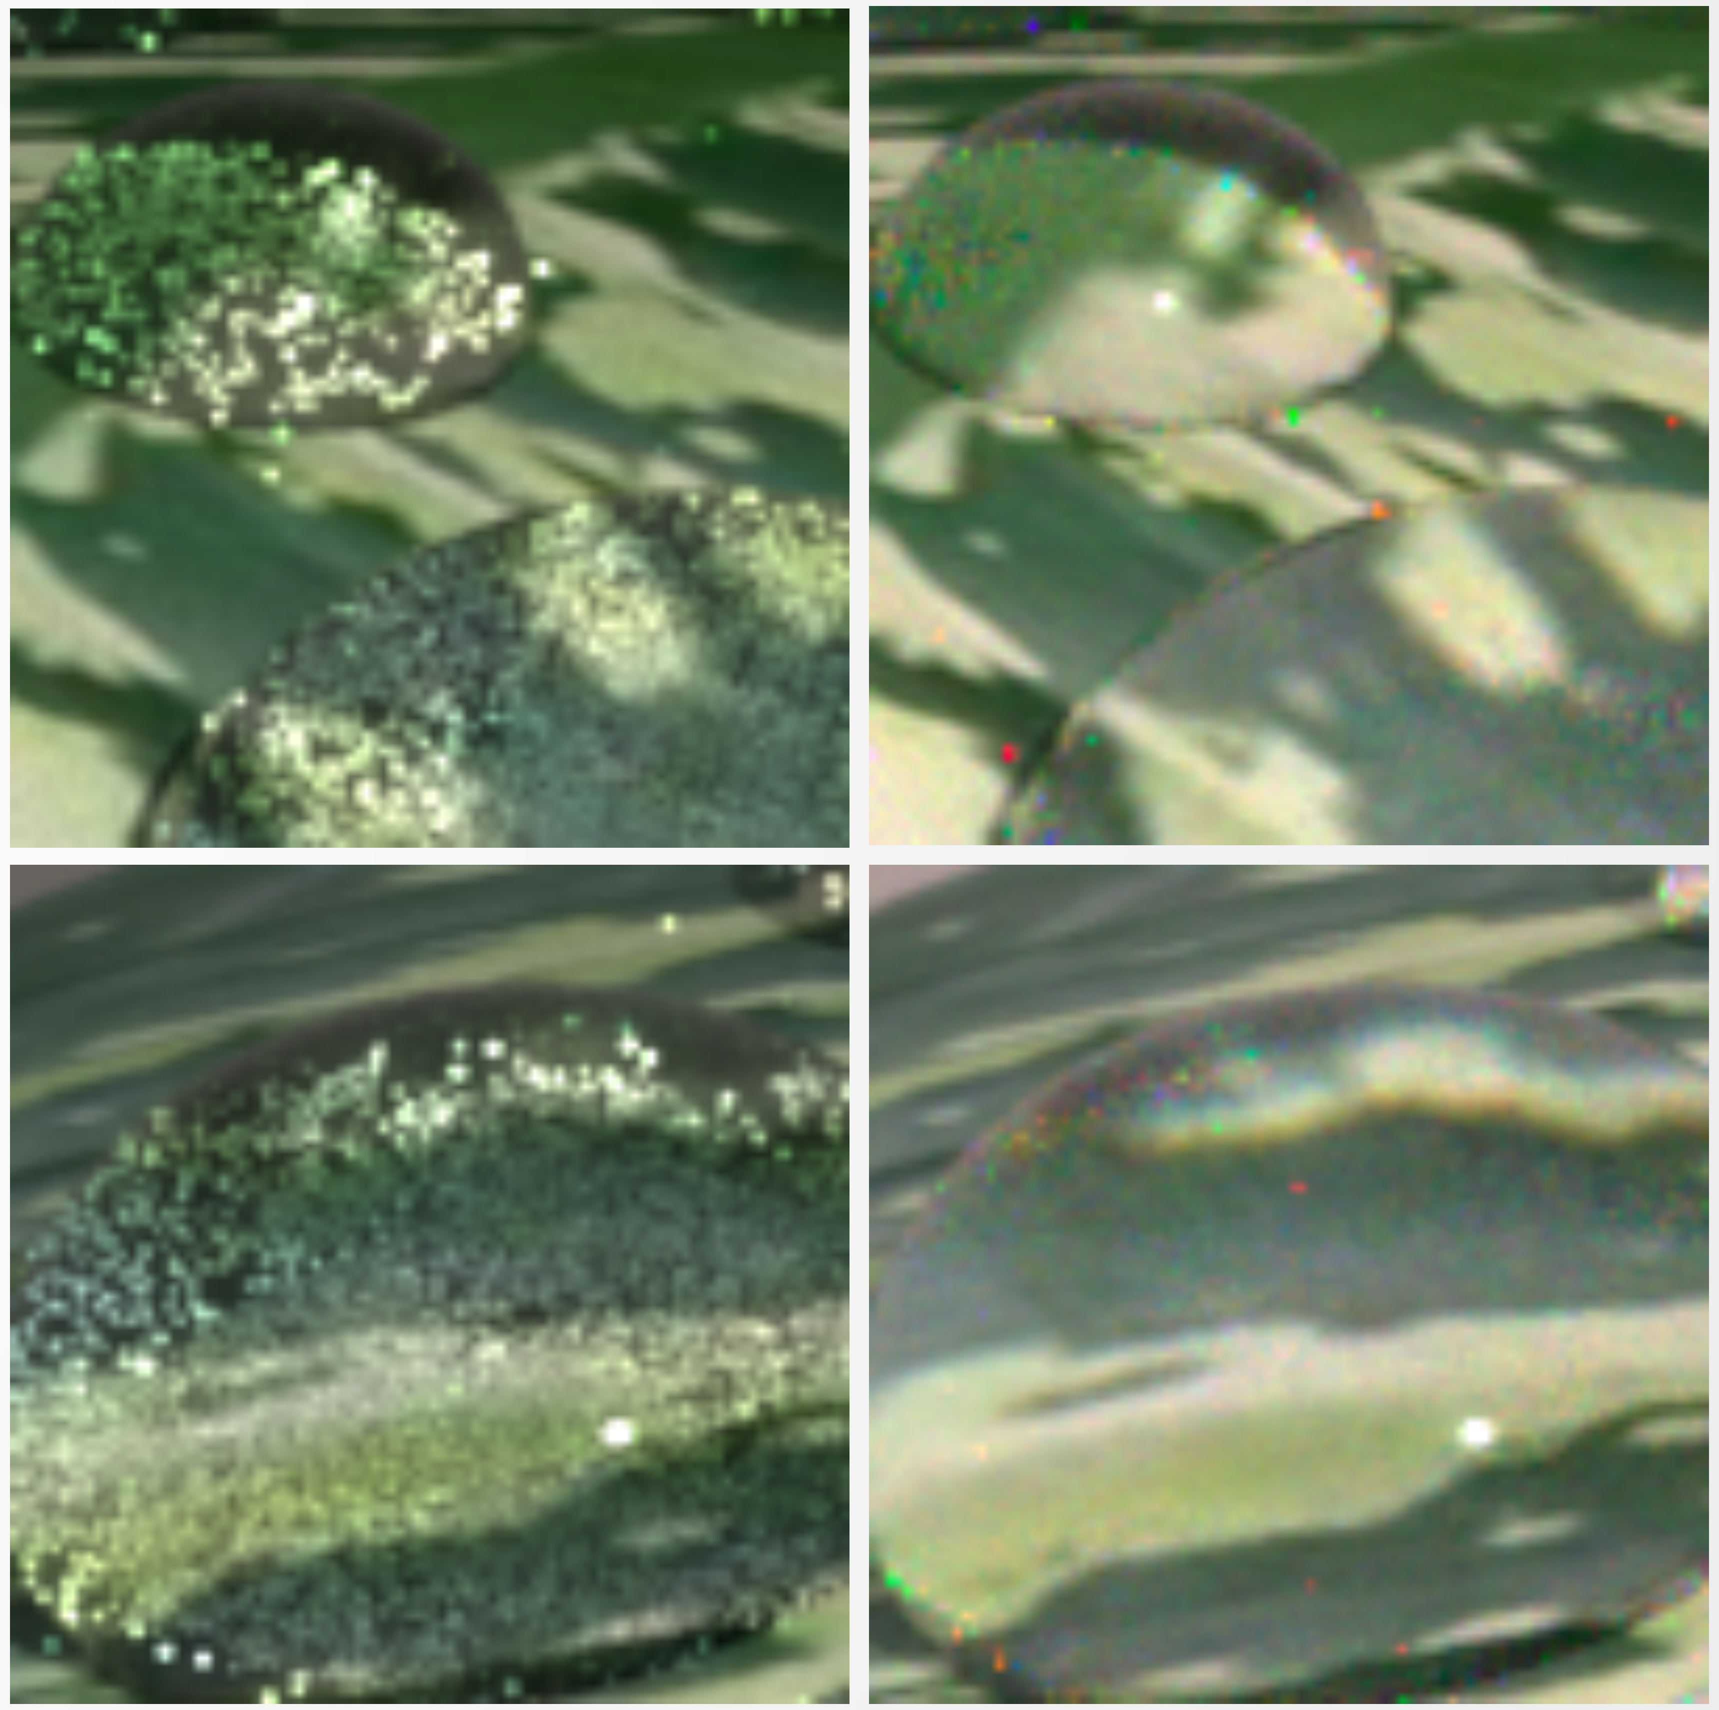

Recent Guiding Methods

Müller et al. [2017] 256spp

ours 260spp

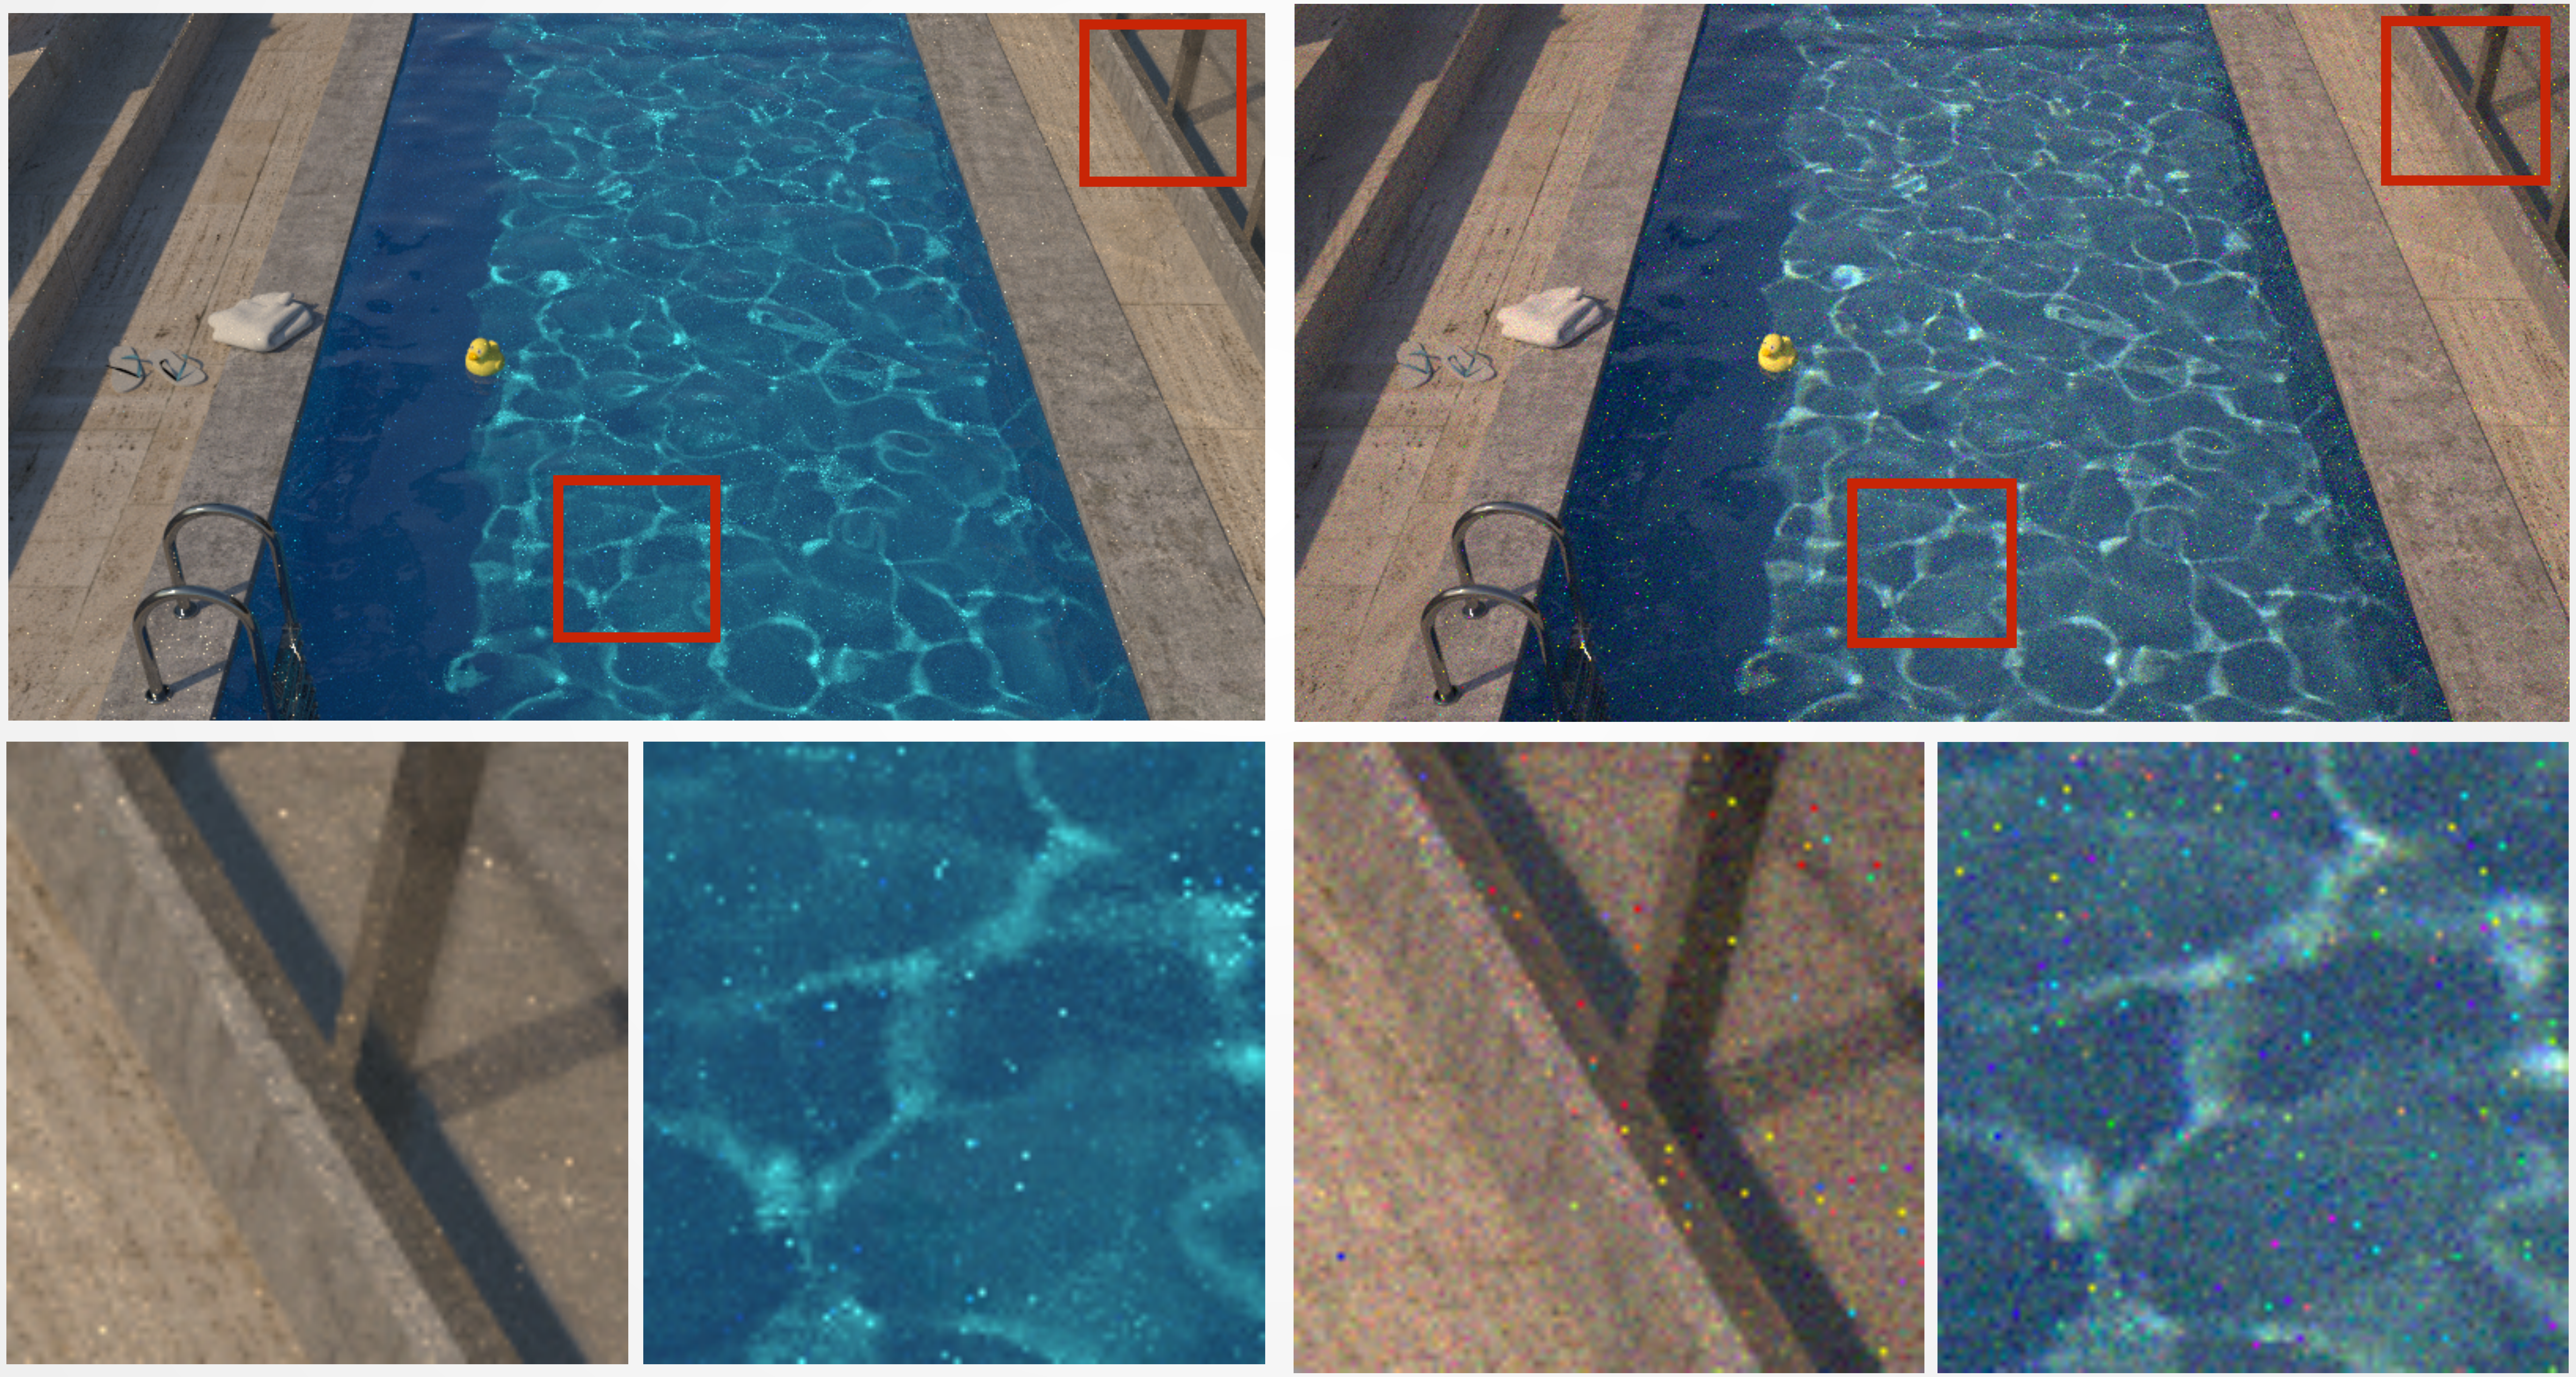

Illustration in 1D

- Uniform sampling

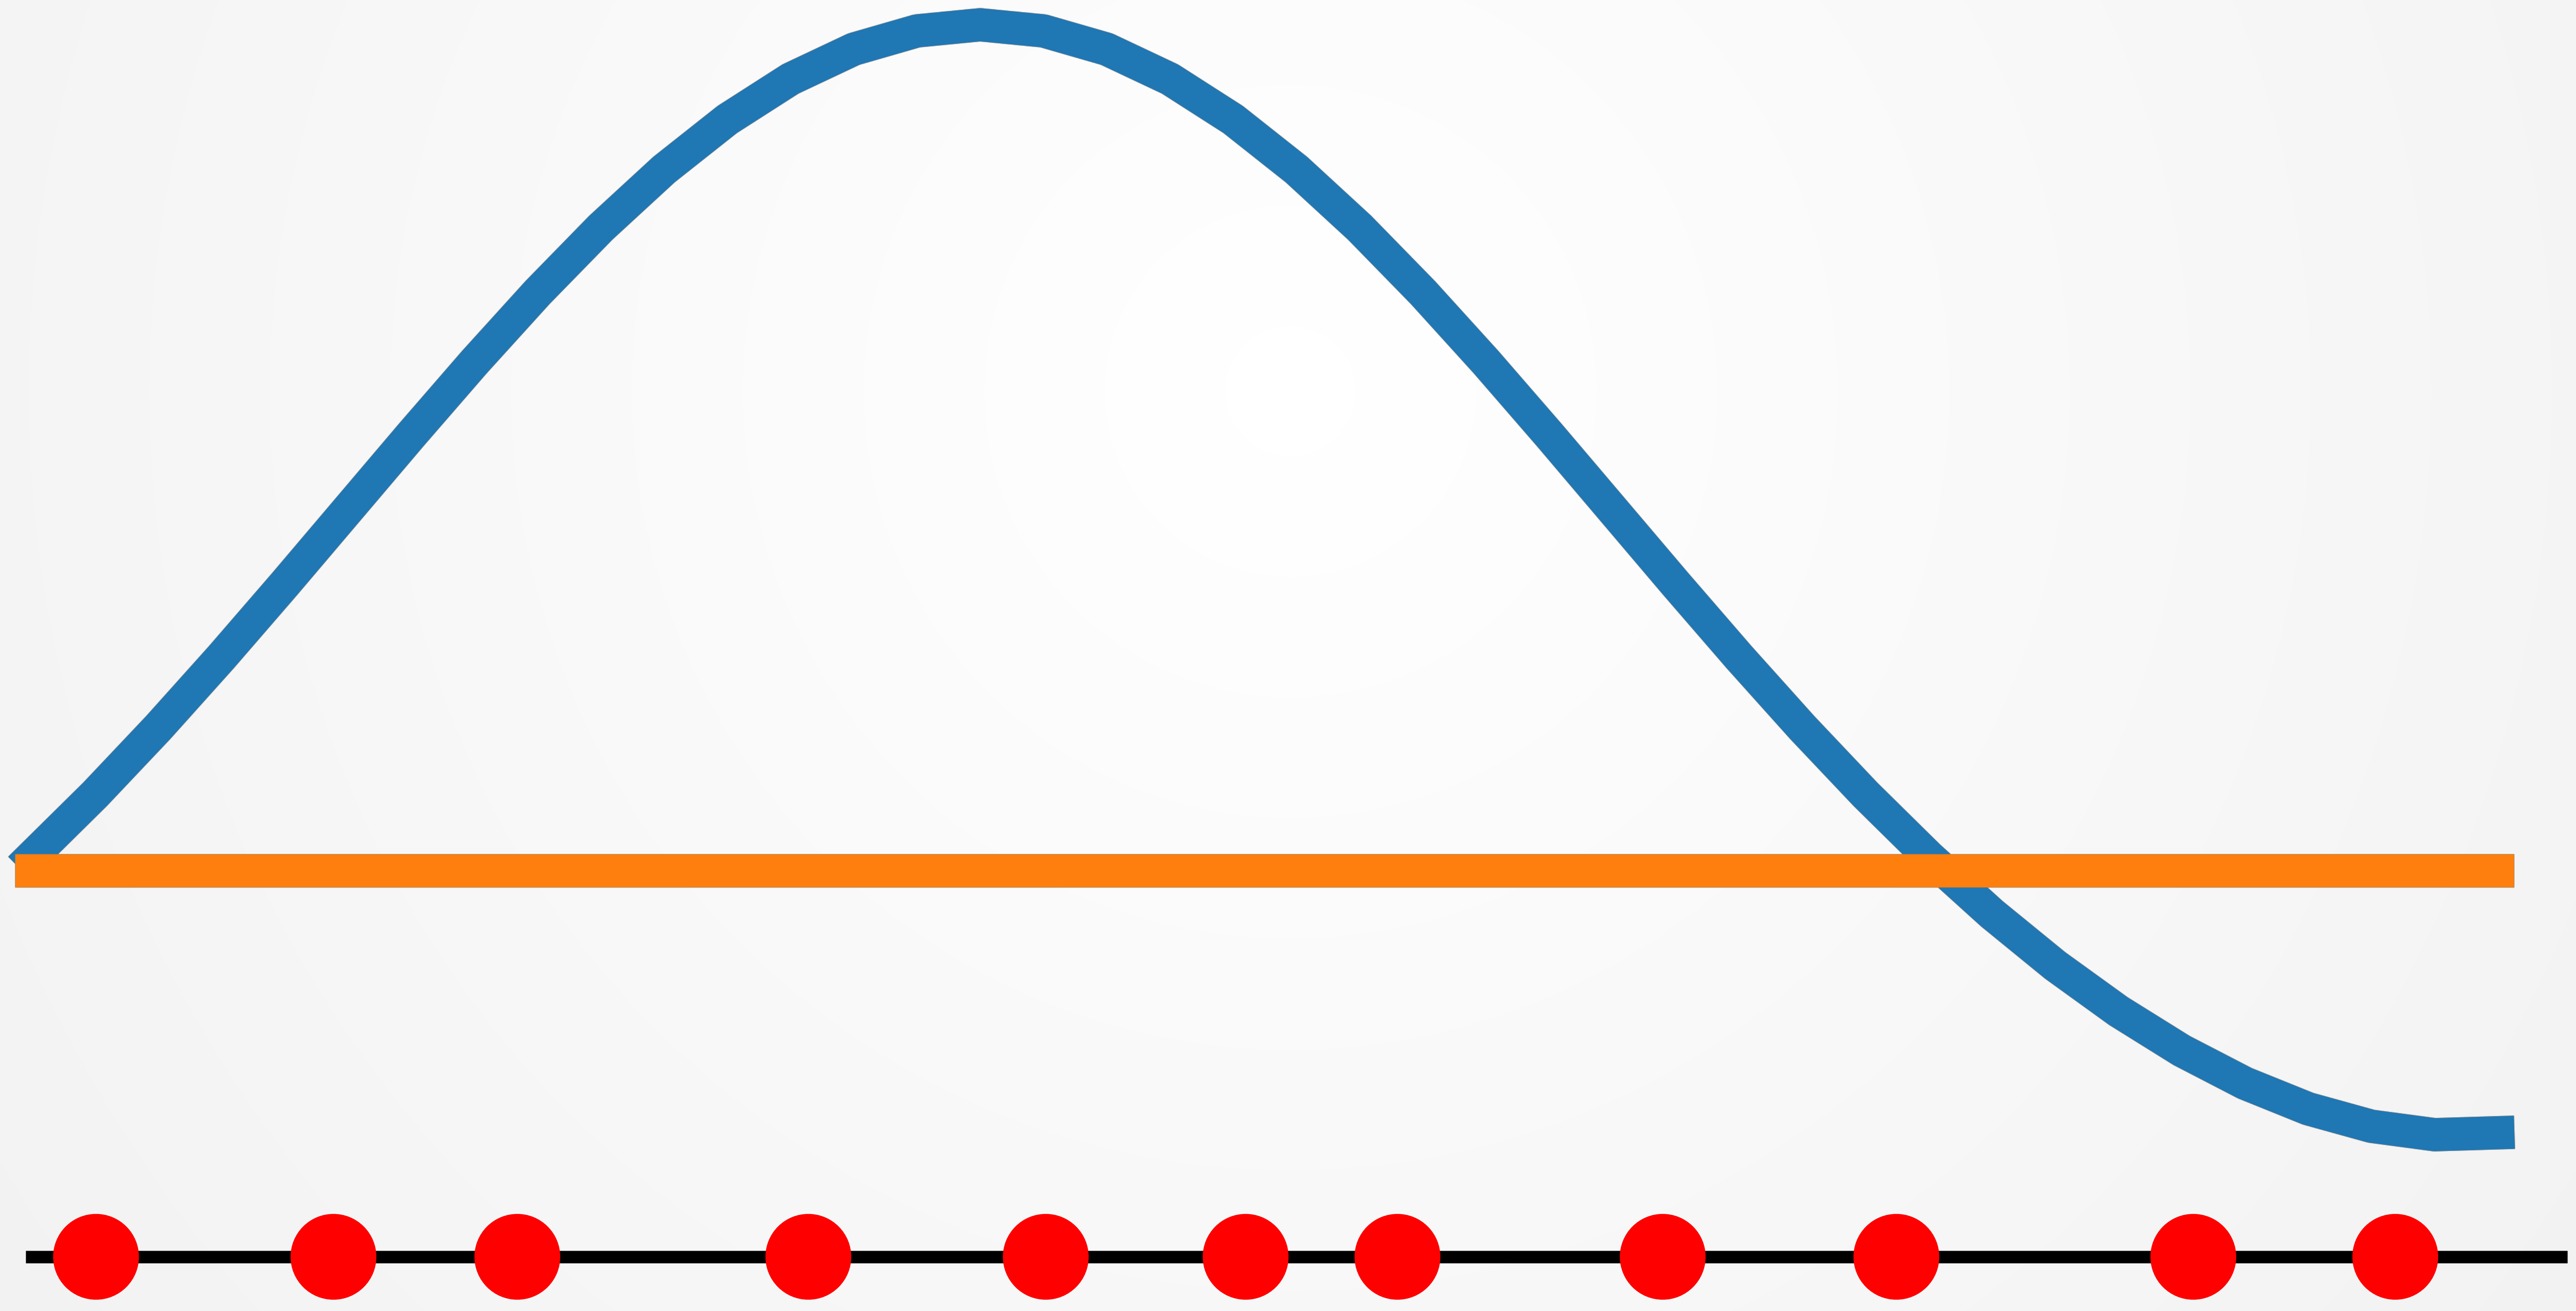

Illustration in 1D

- Importance sampling

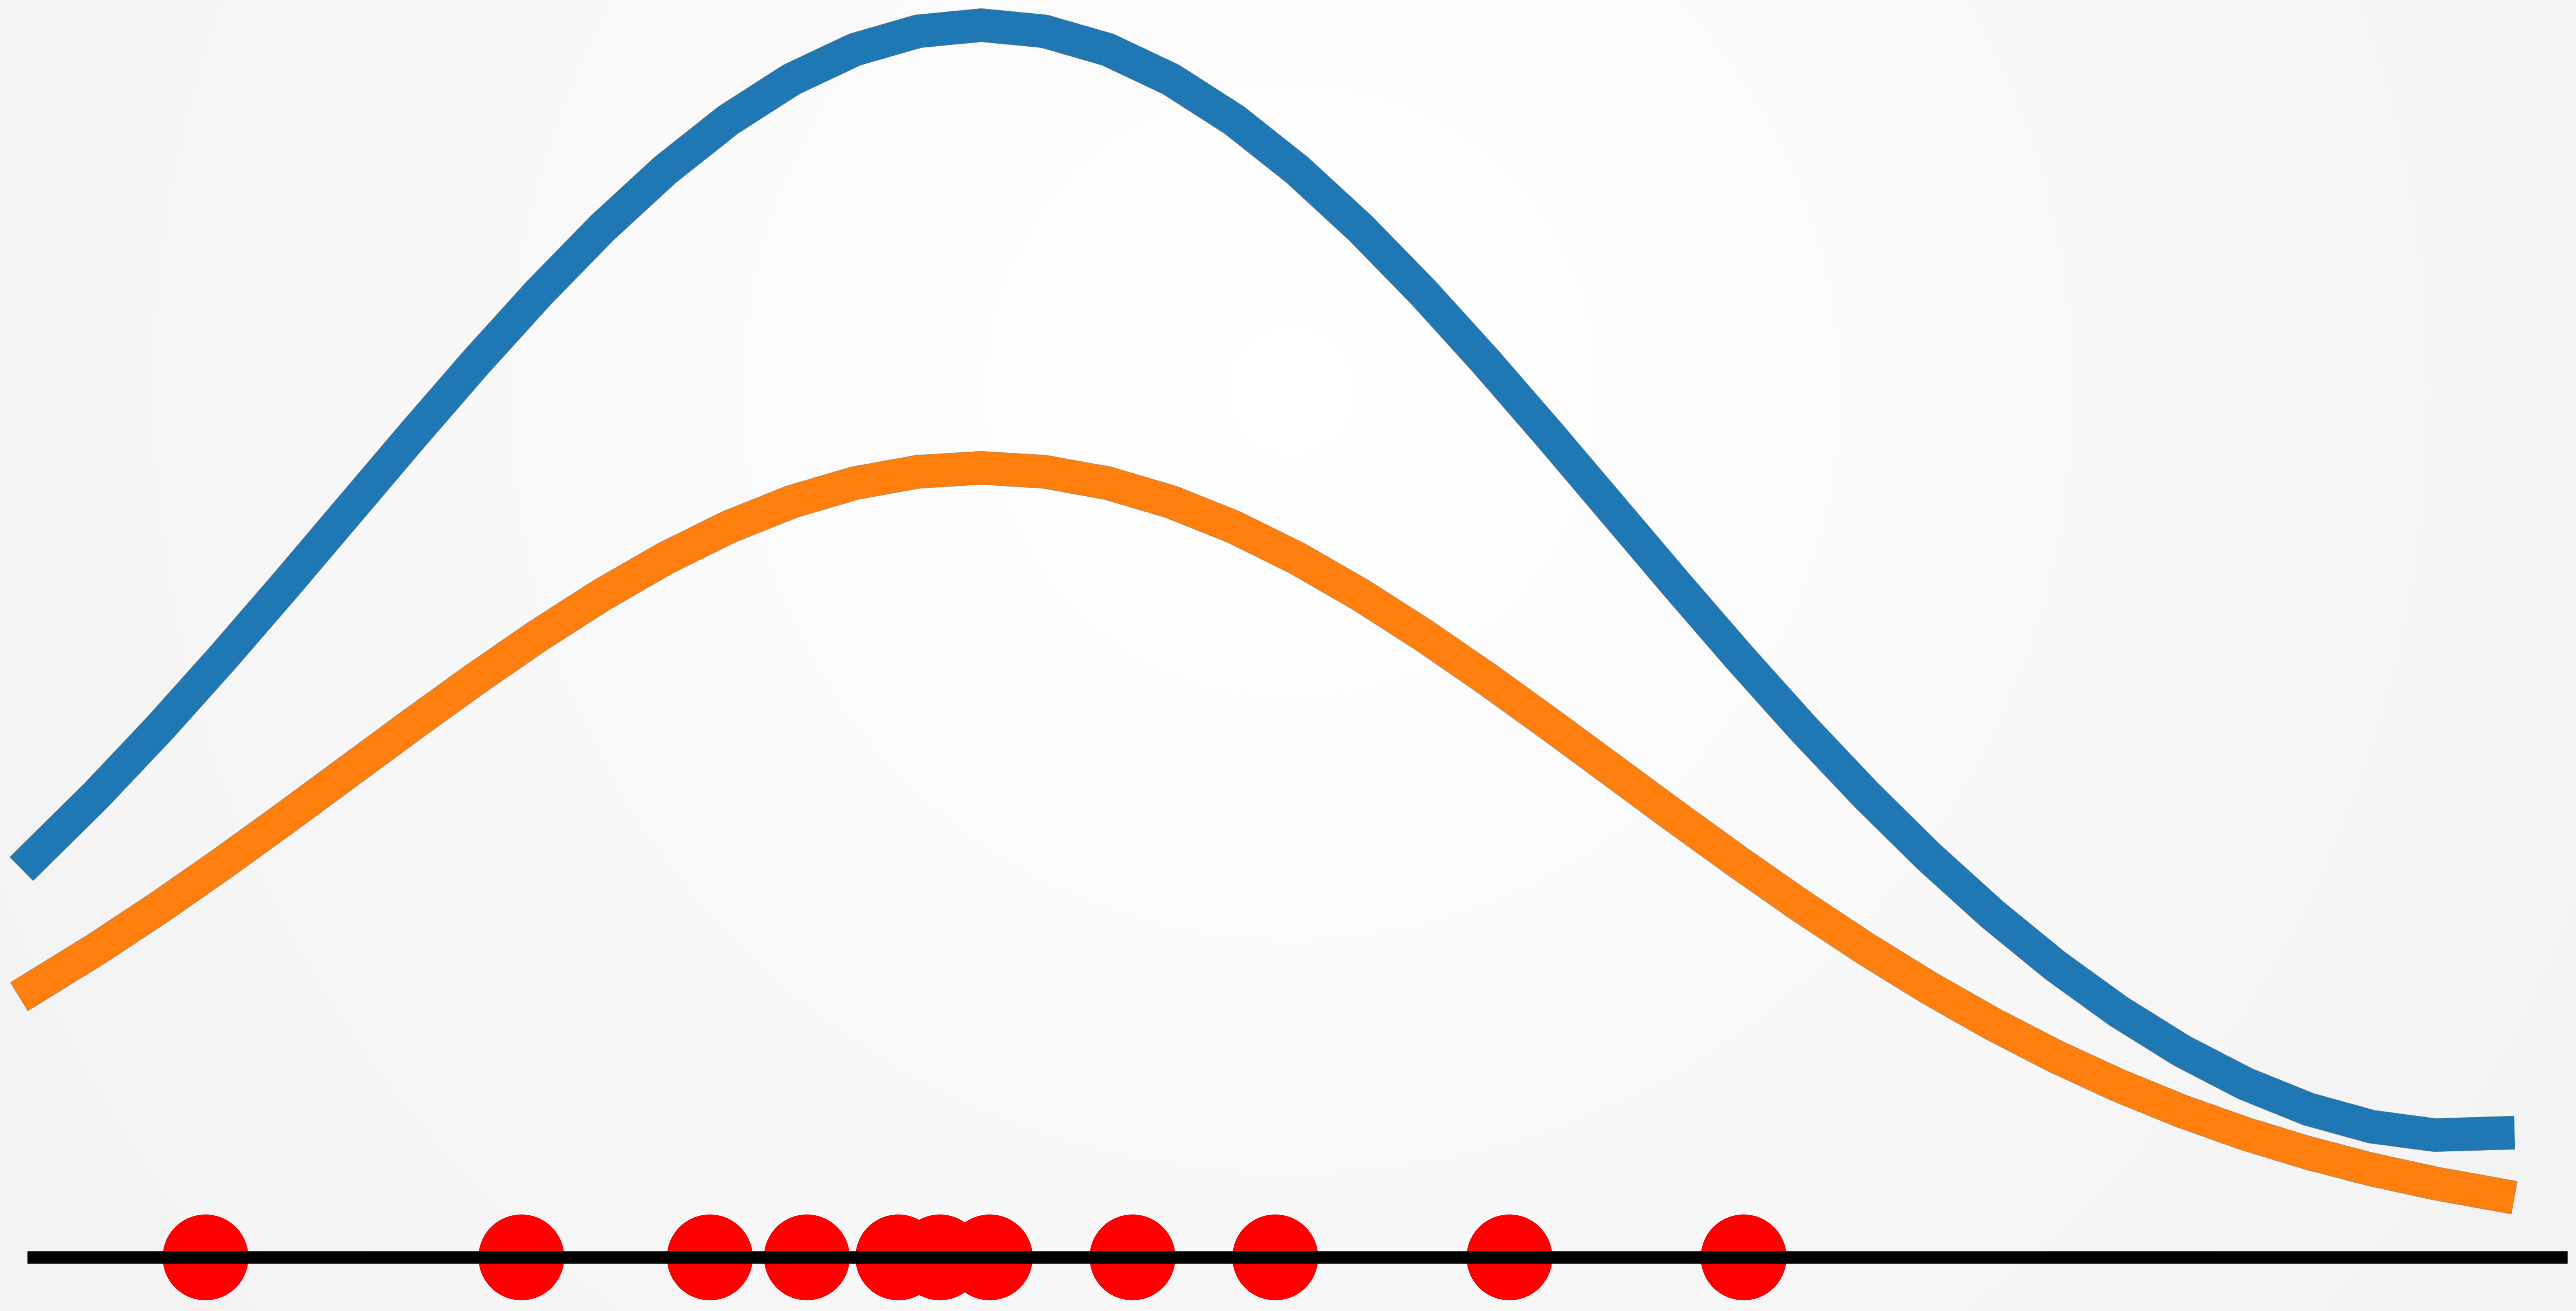

Guiding Behaviour

Guided path tracing - iteration 1

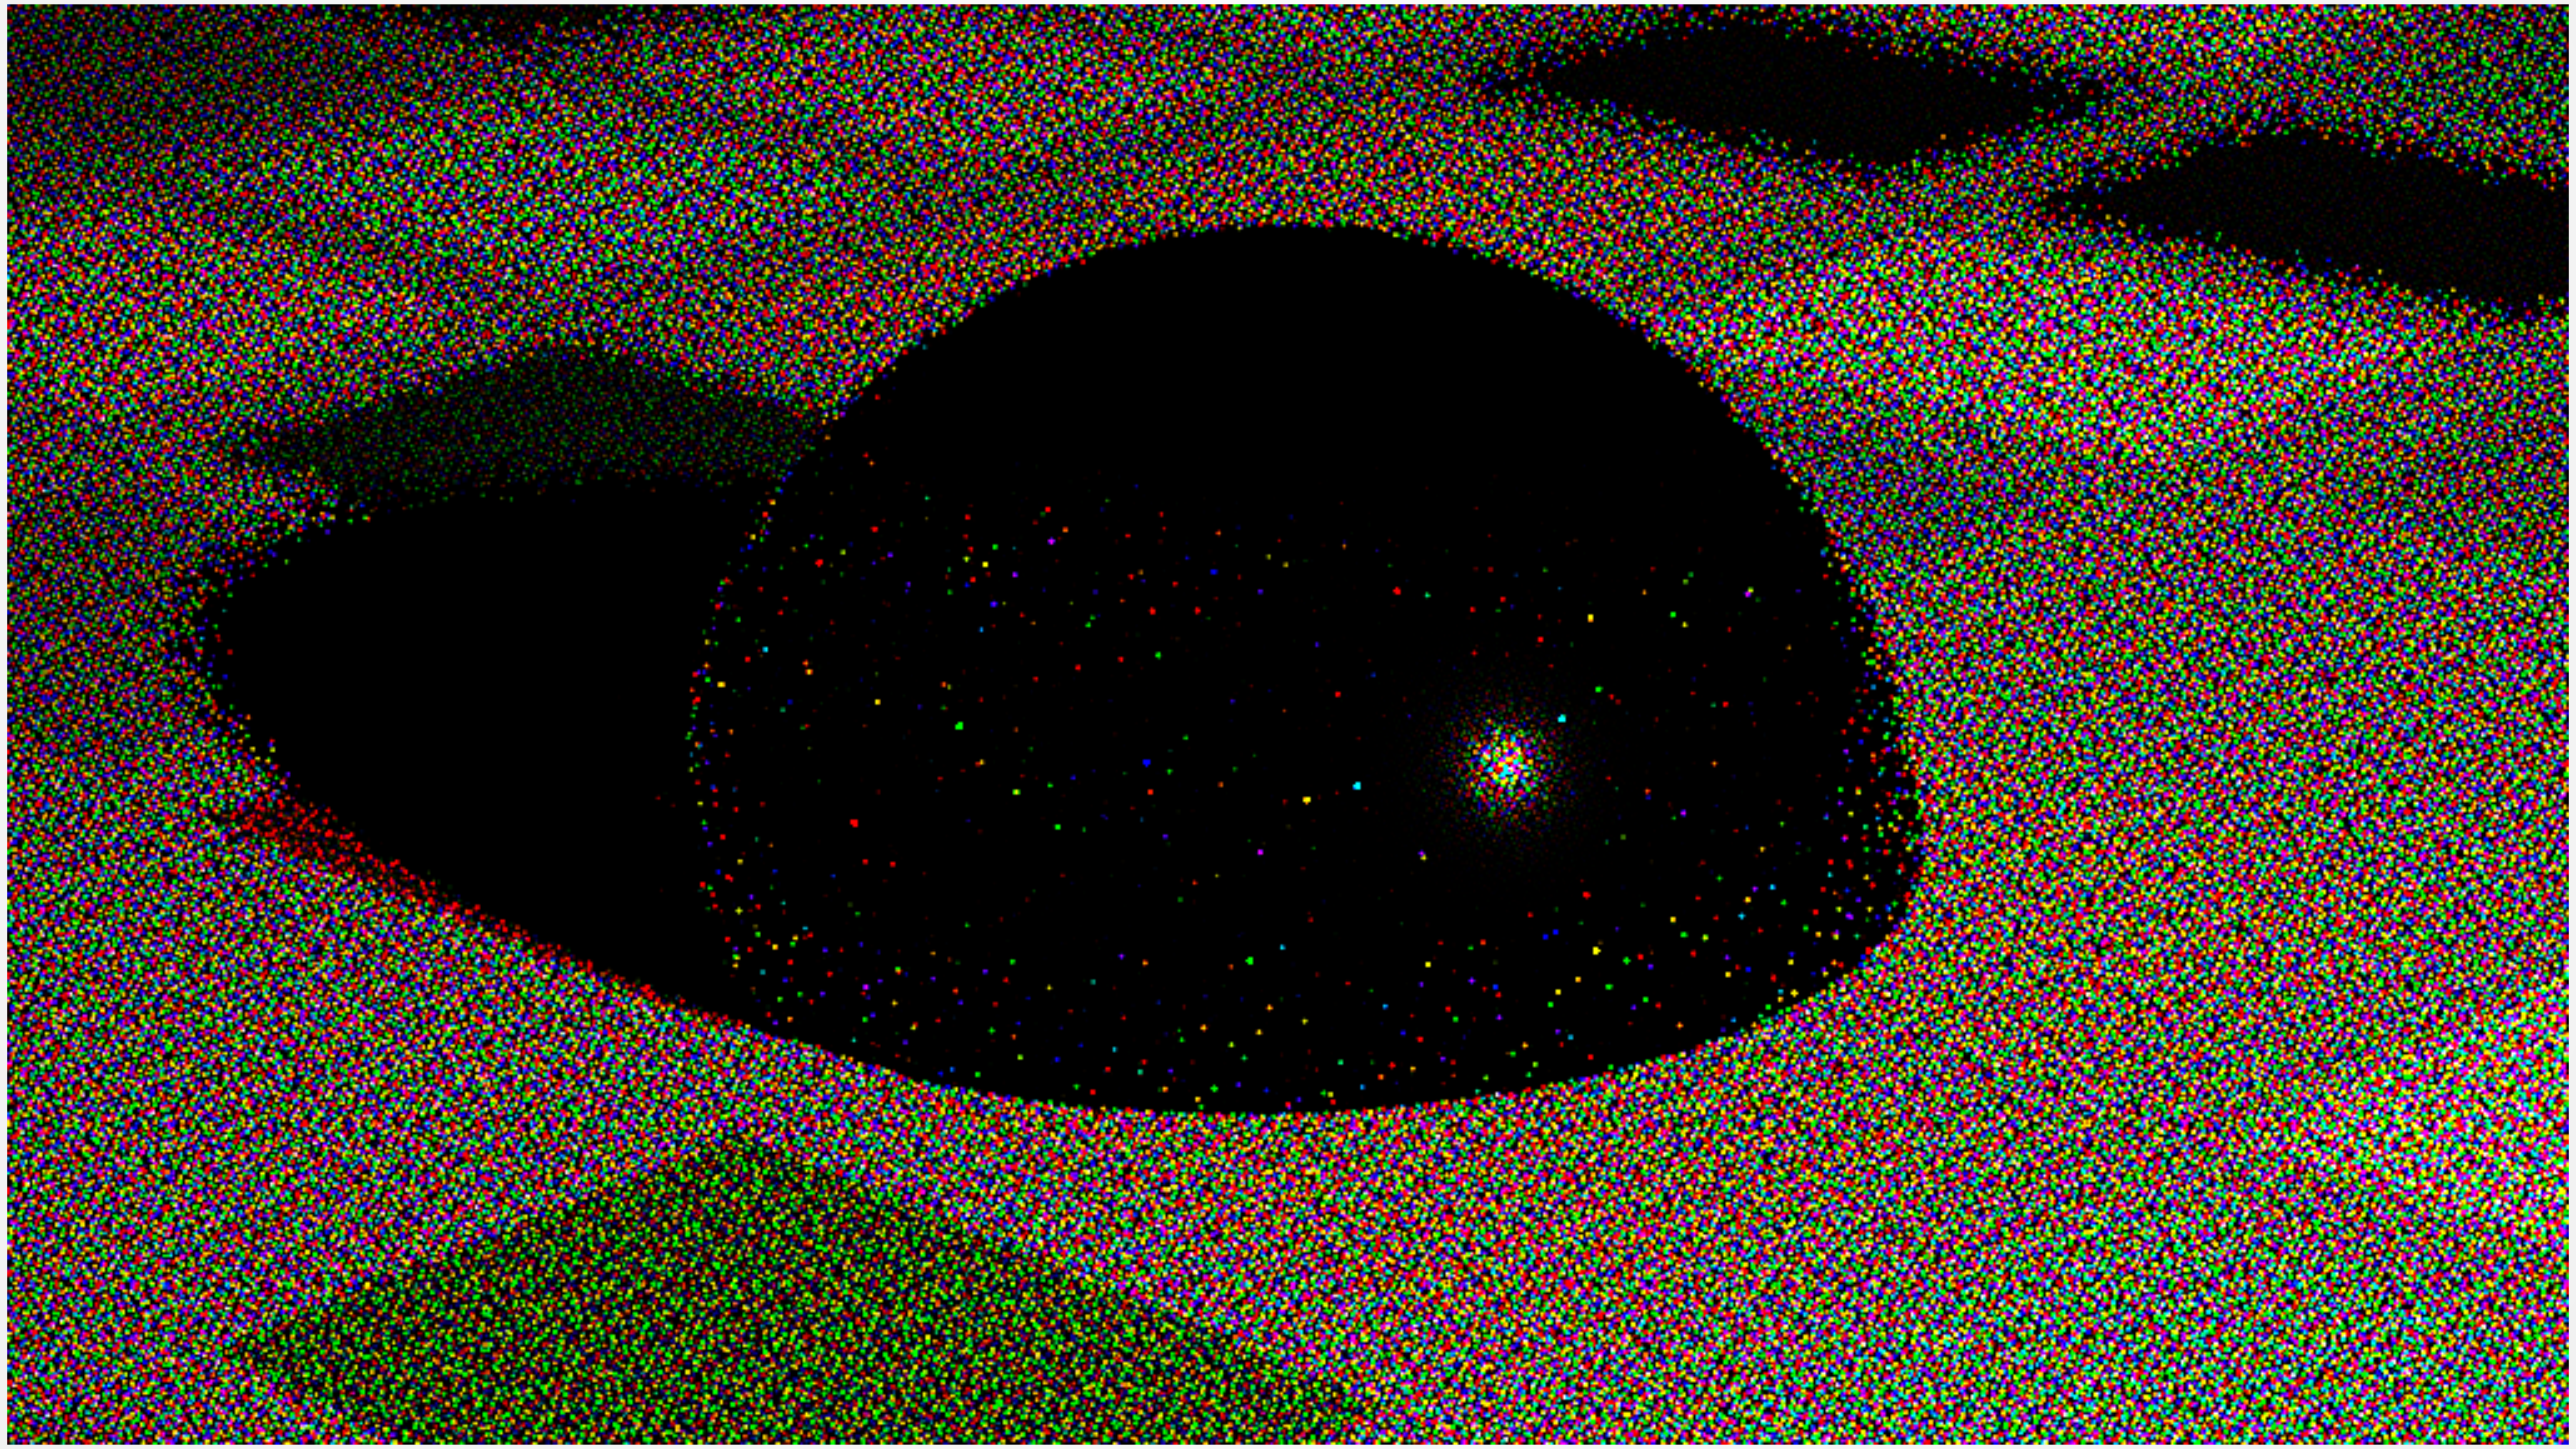

Guiding Behaviour

Guided path tracing - iteration 2

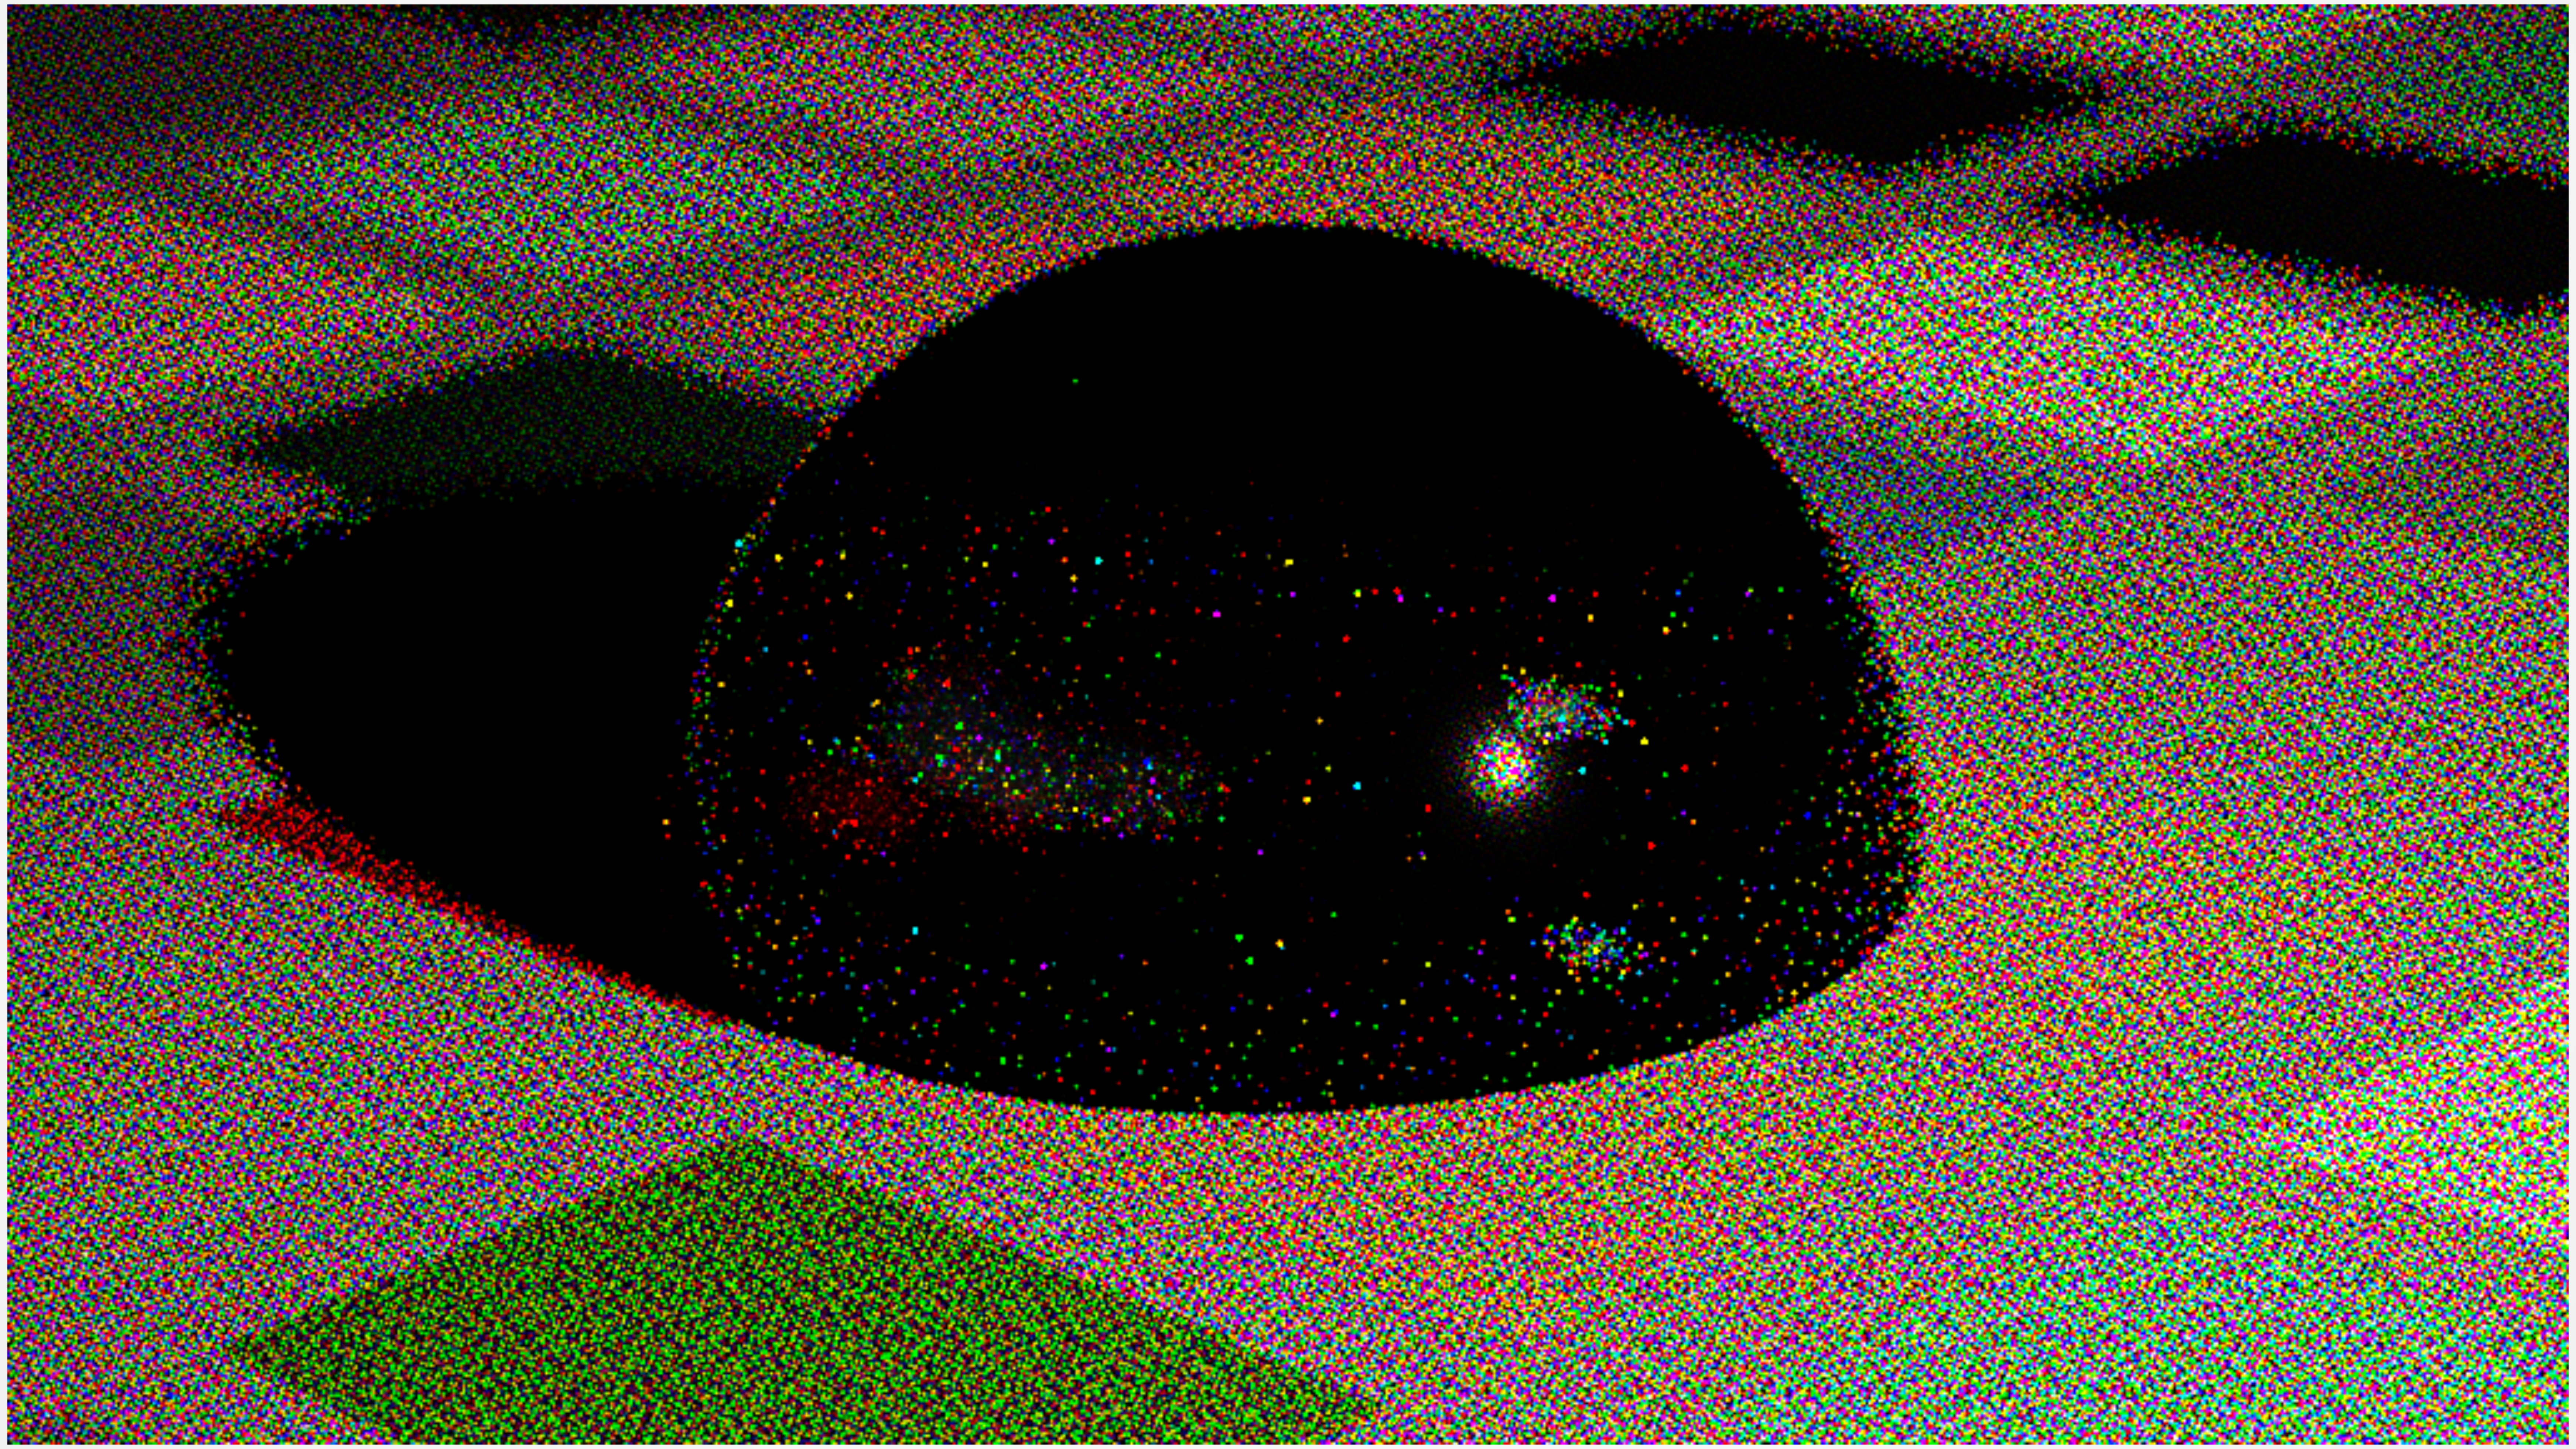

Guiding Behaviour

Guided path tracing - iteration 3

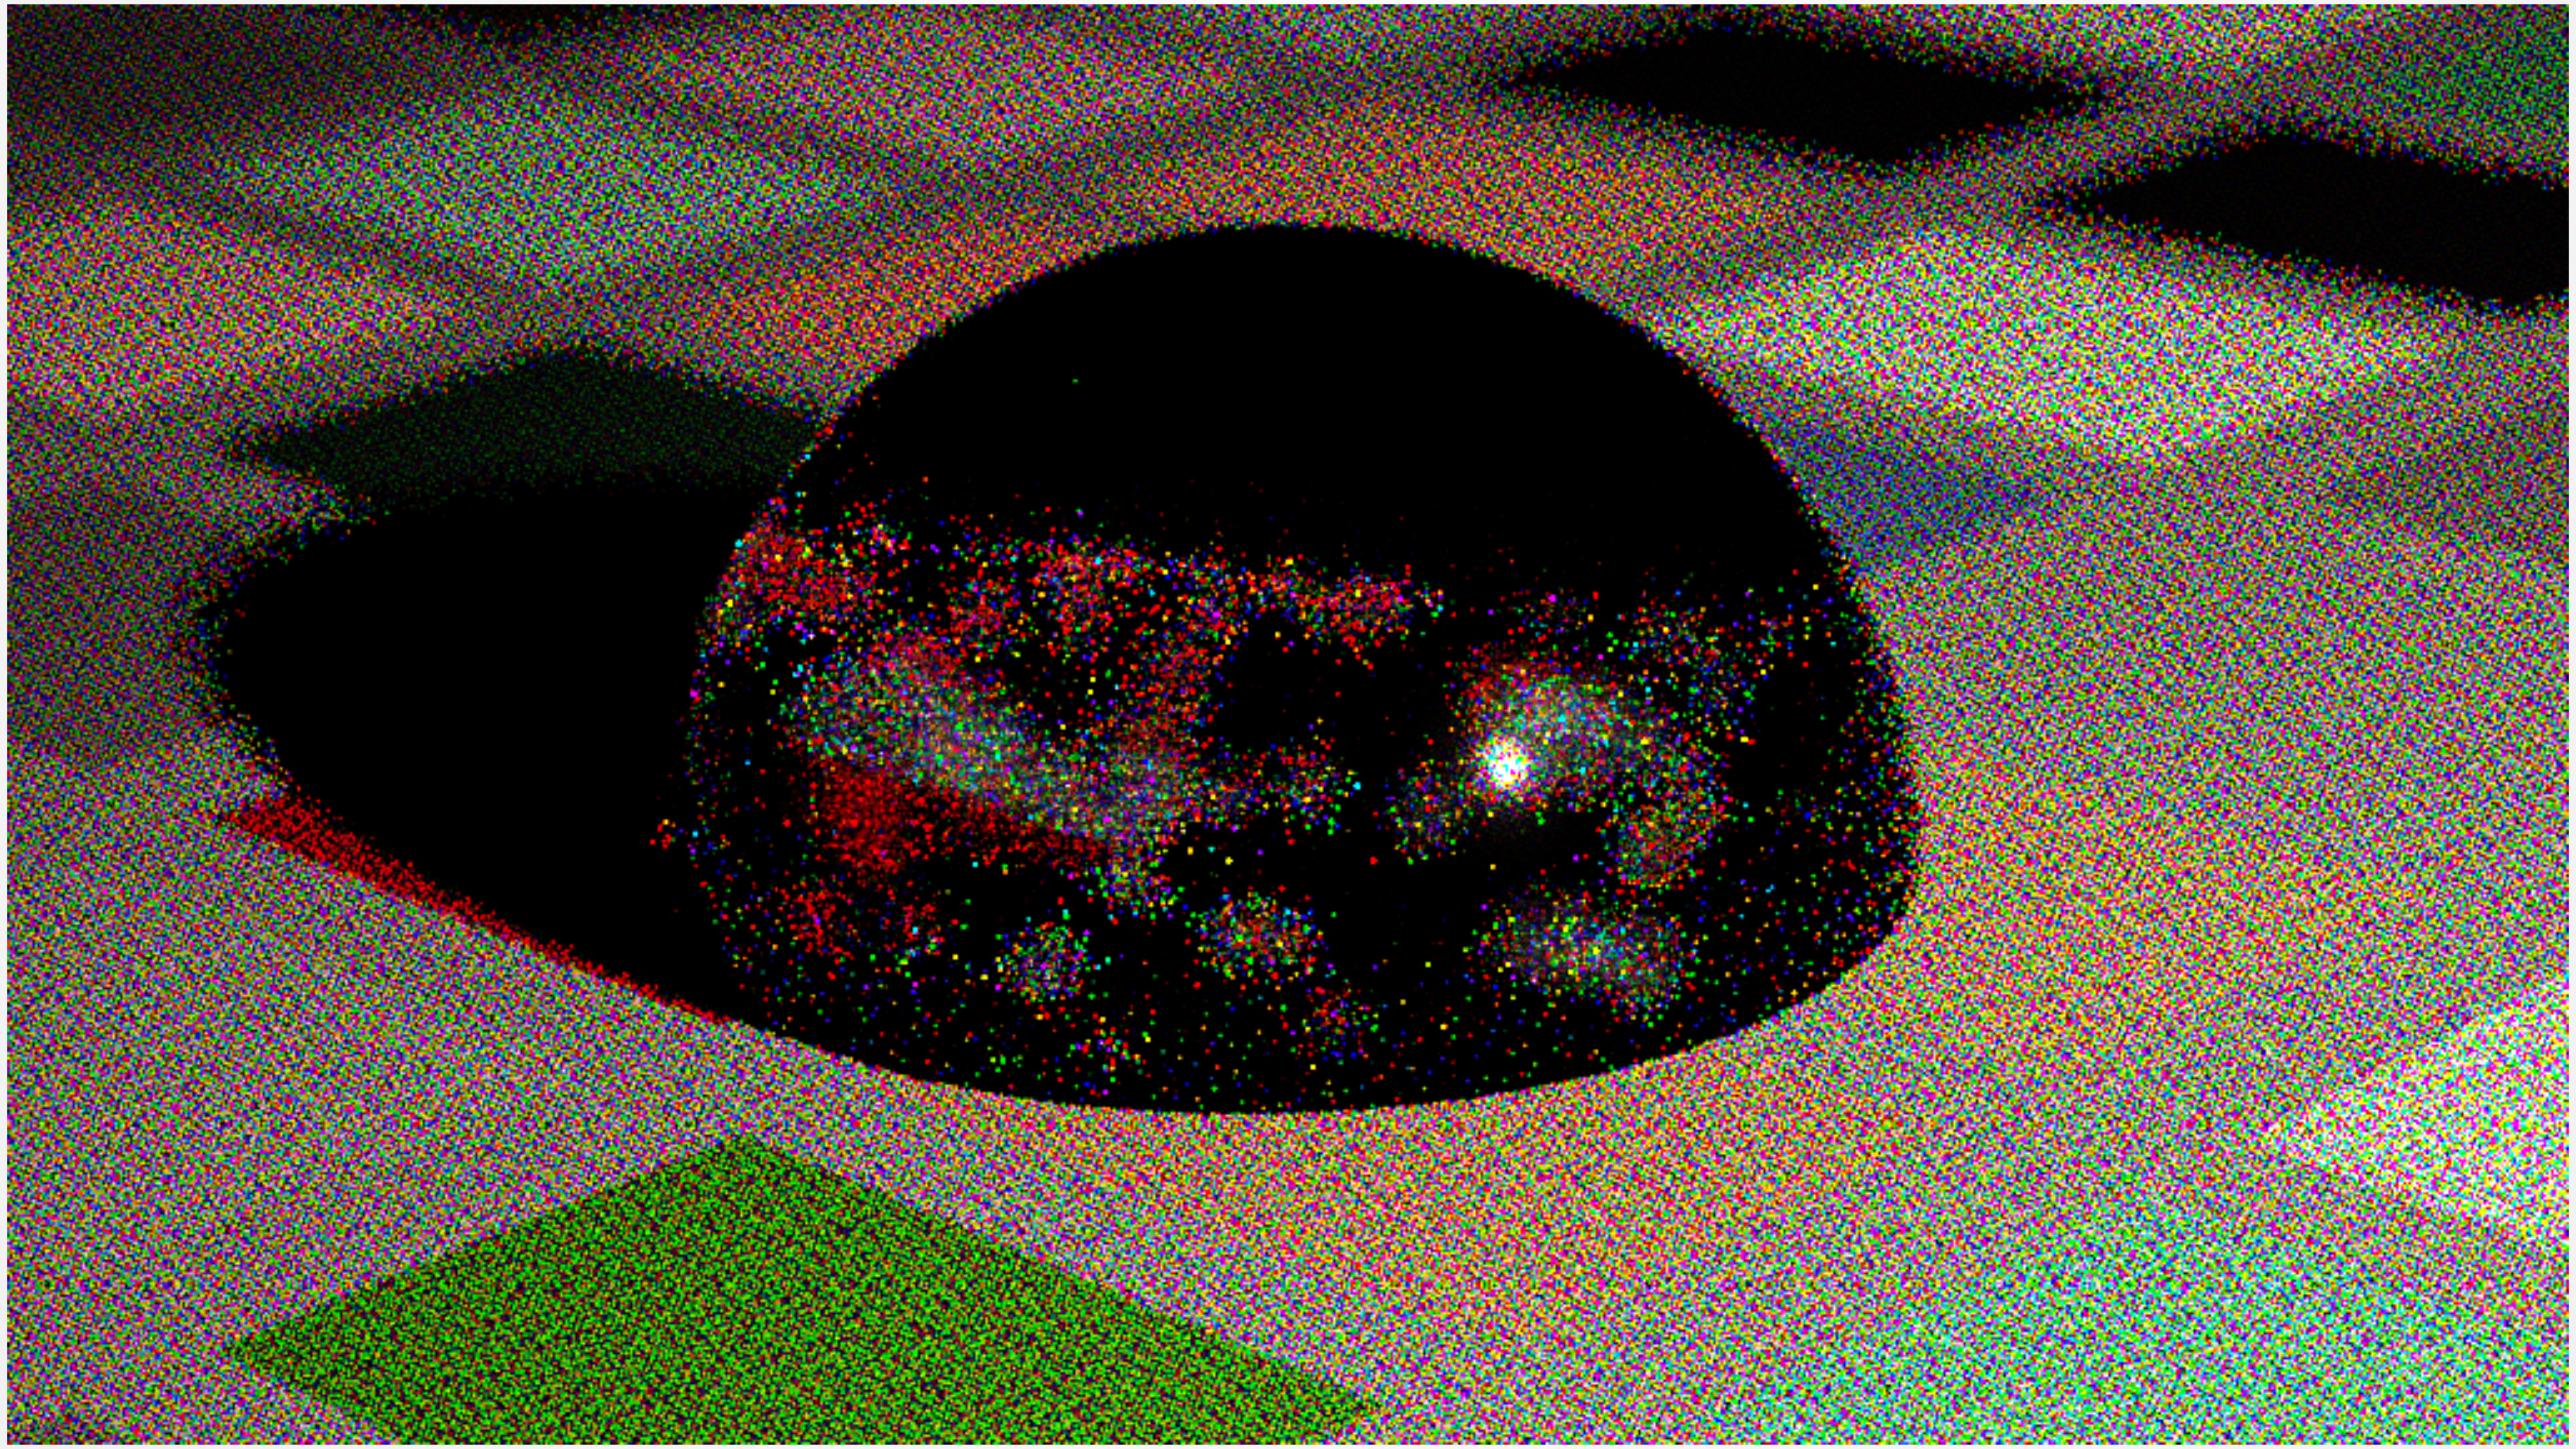

Guiding Behaviour

Guided path tracing - iteration 4

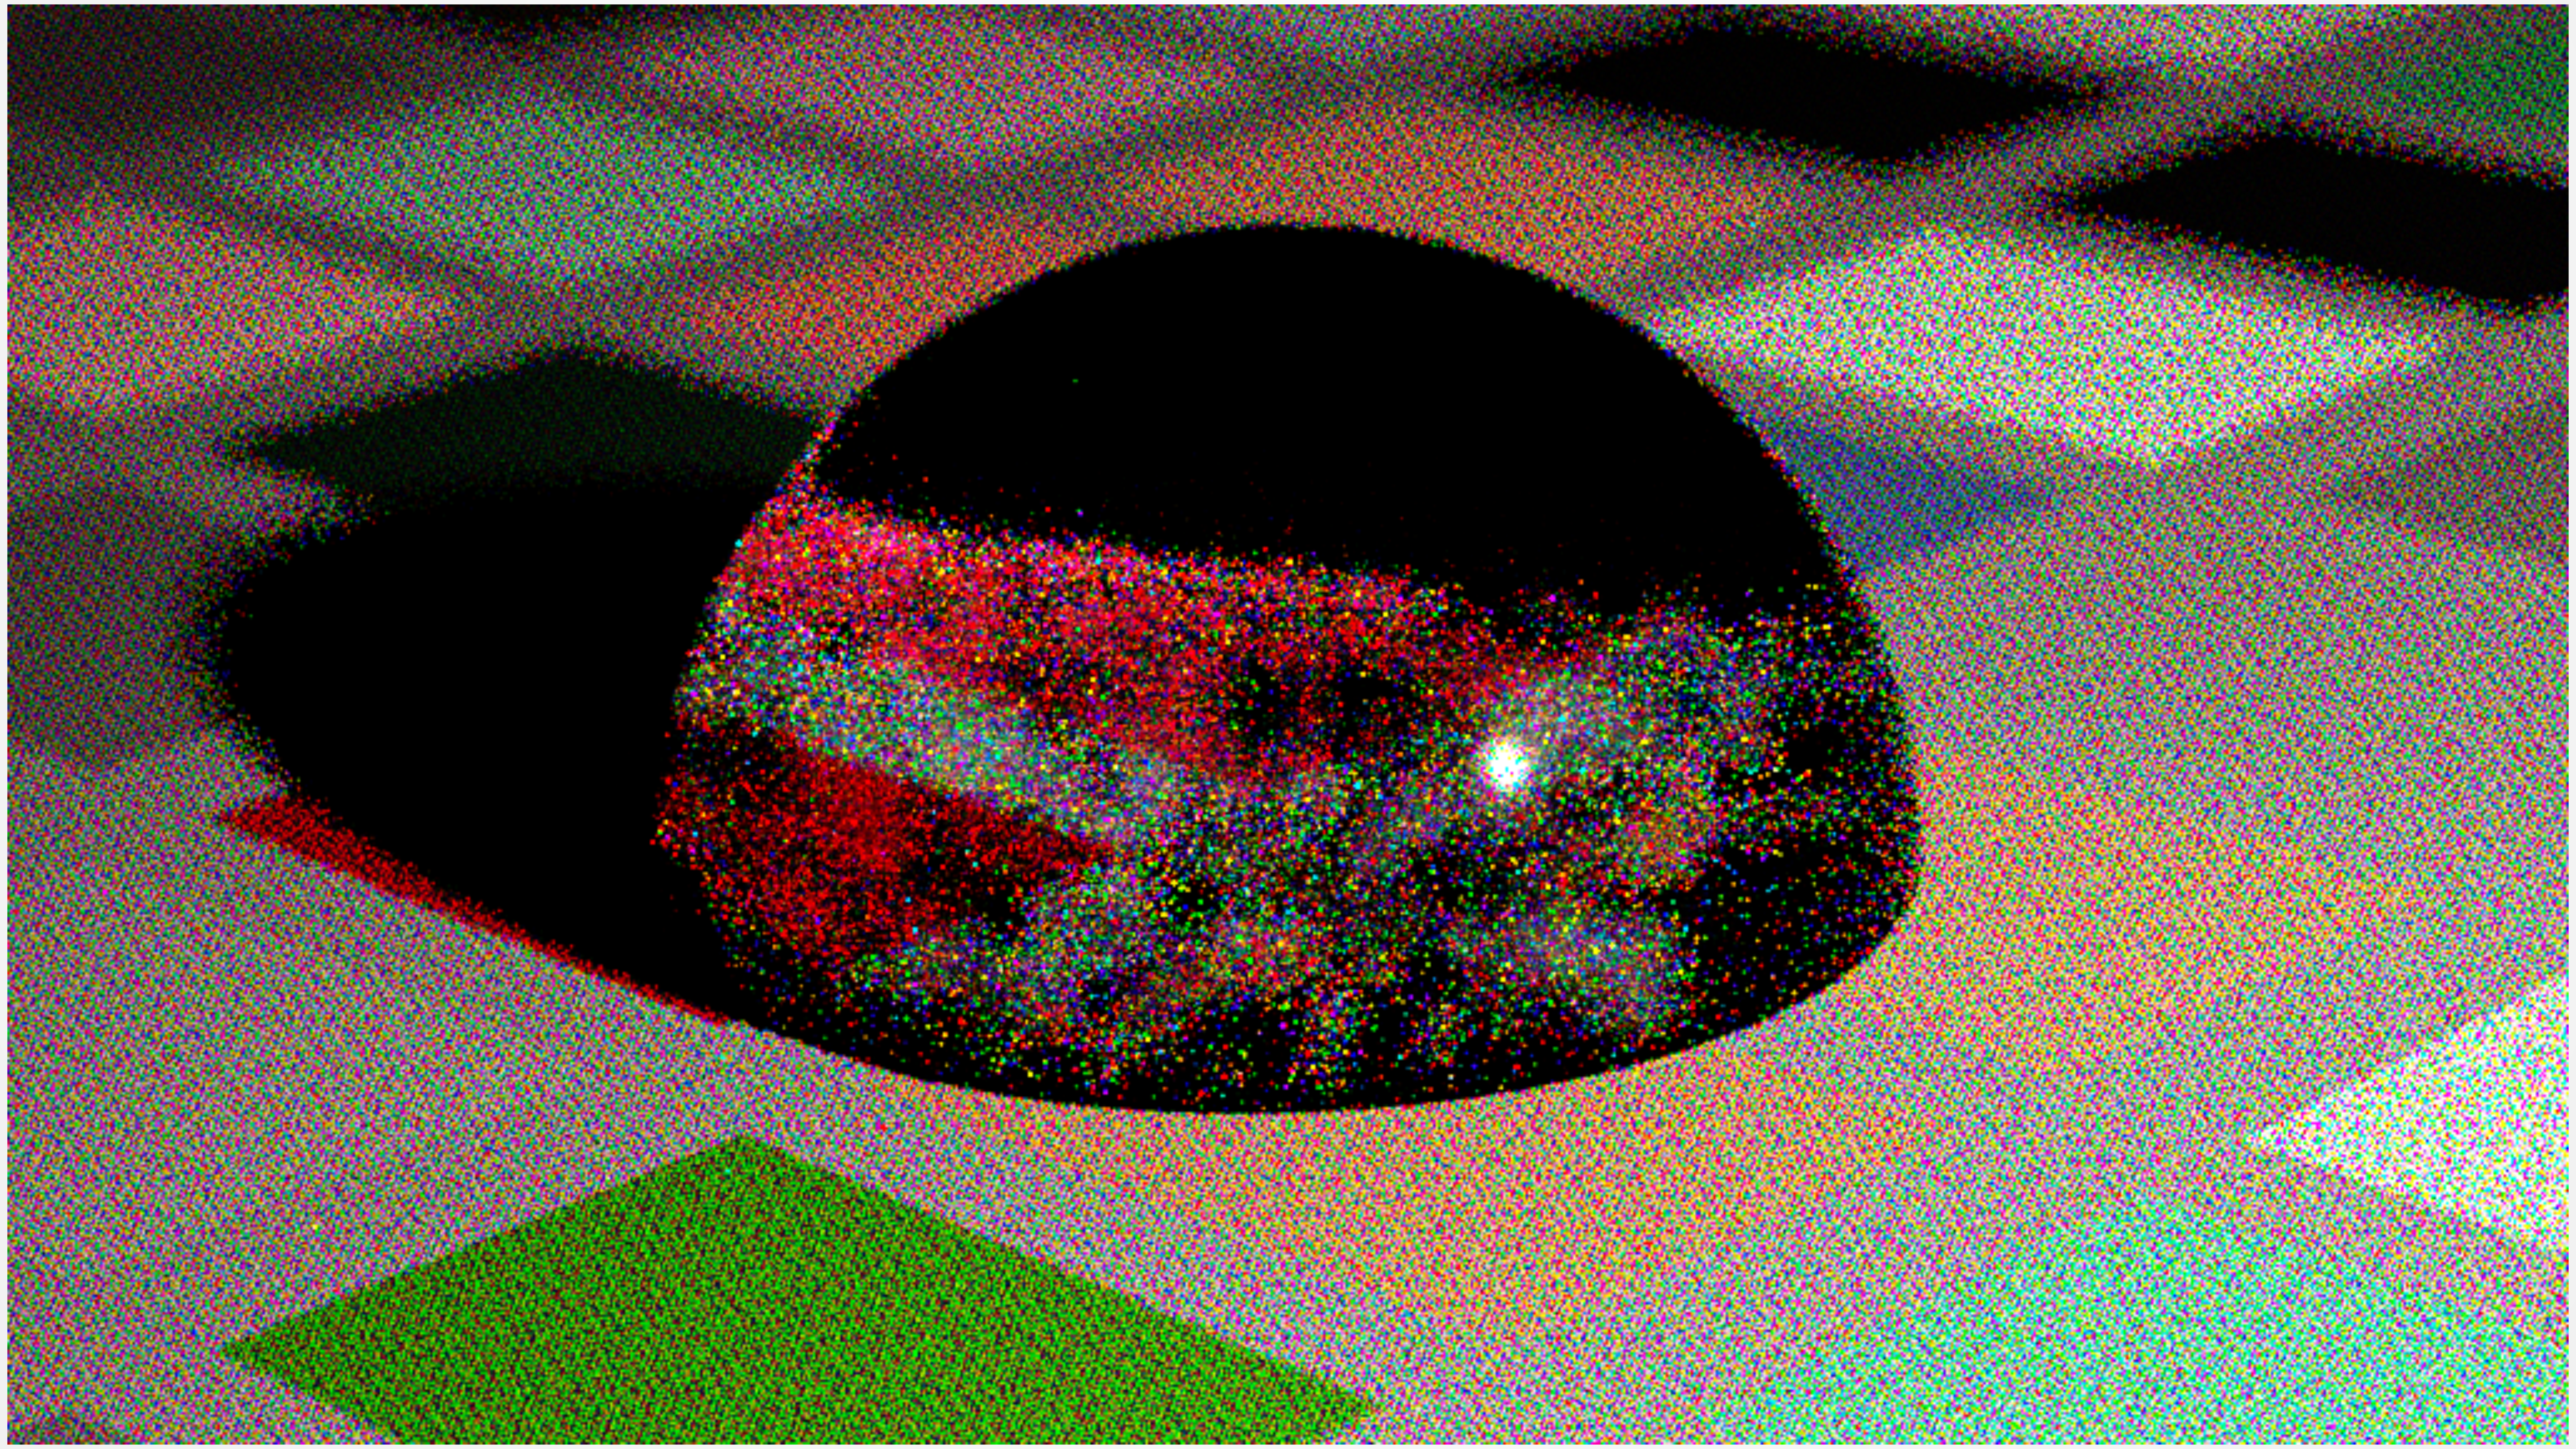

Guiding Behaviour

Guided path tracing - iteration 8

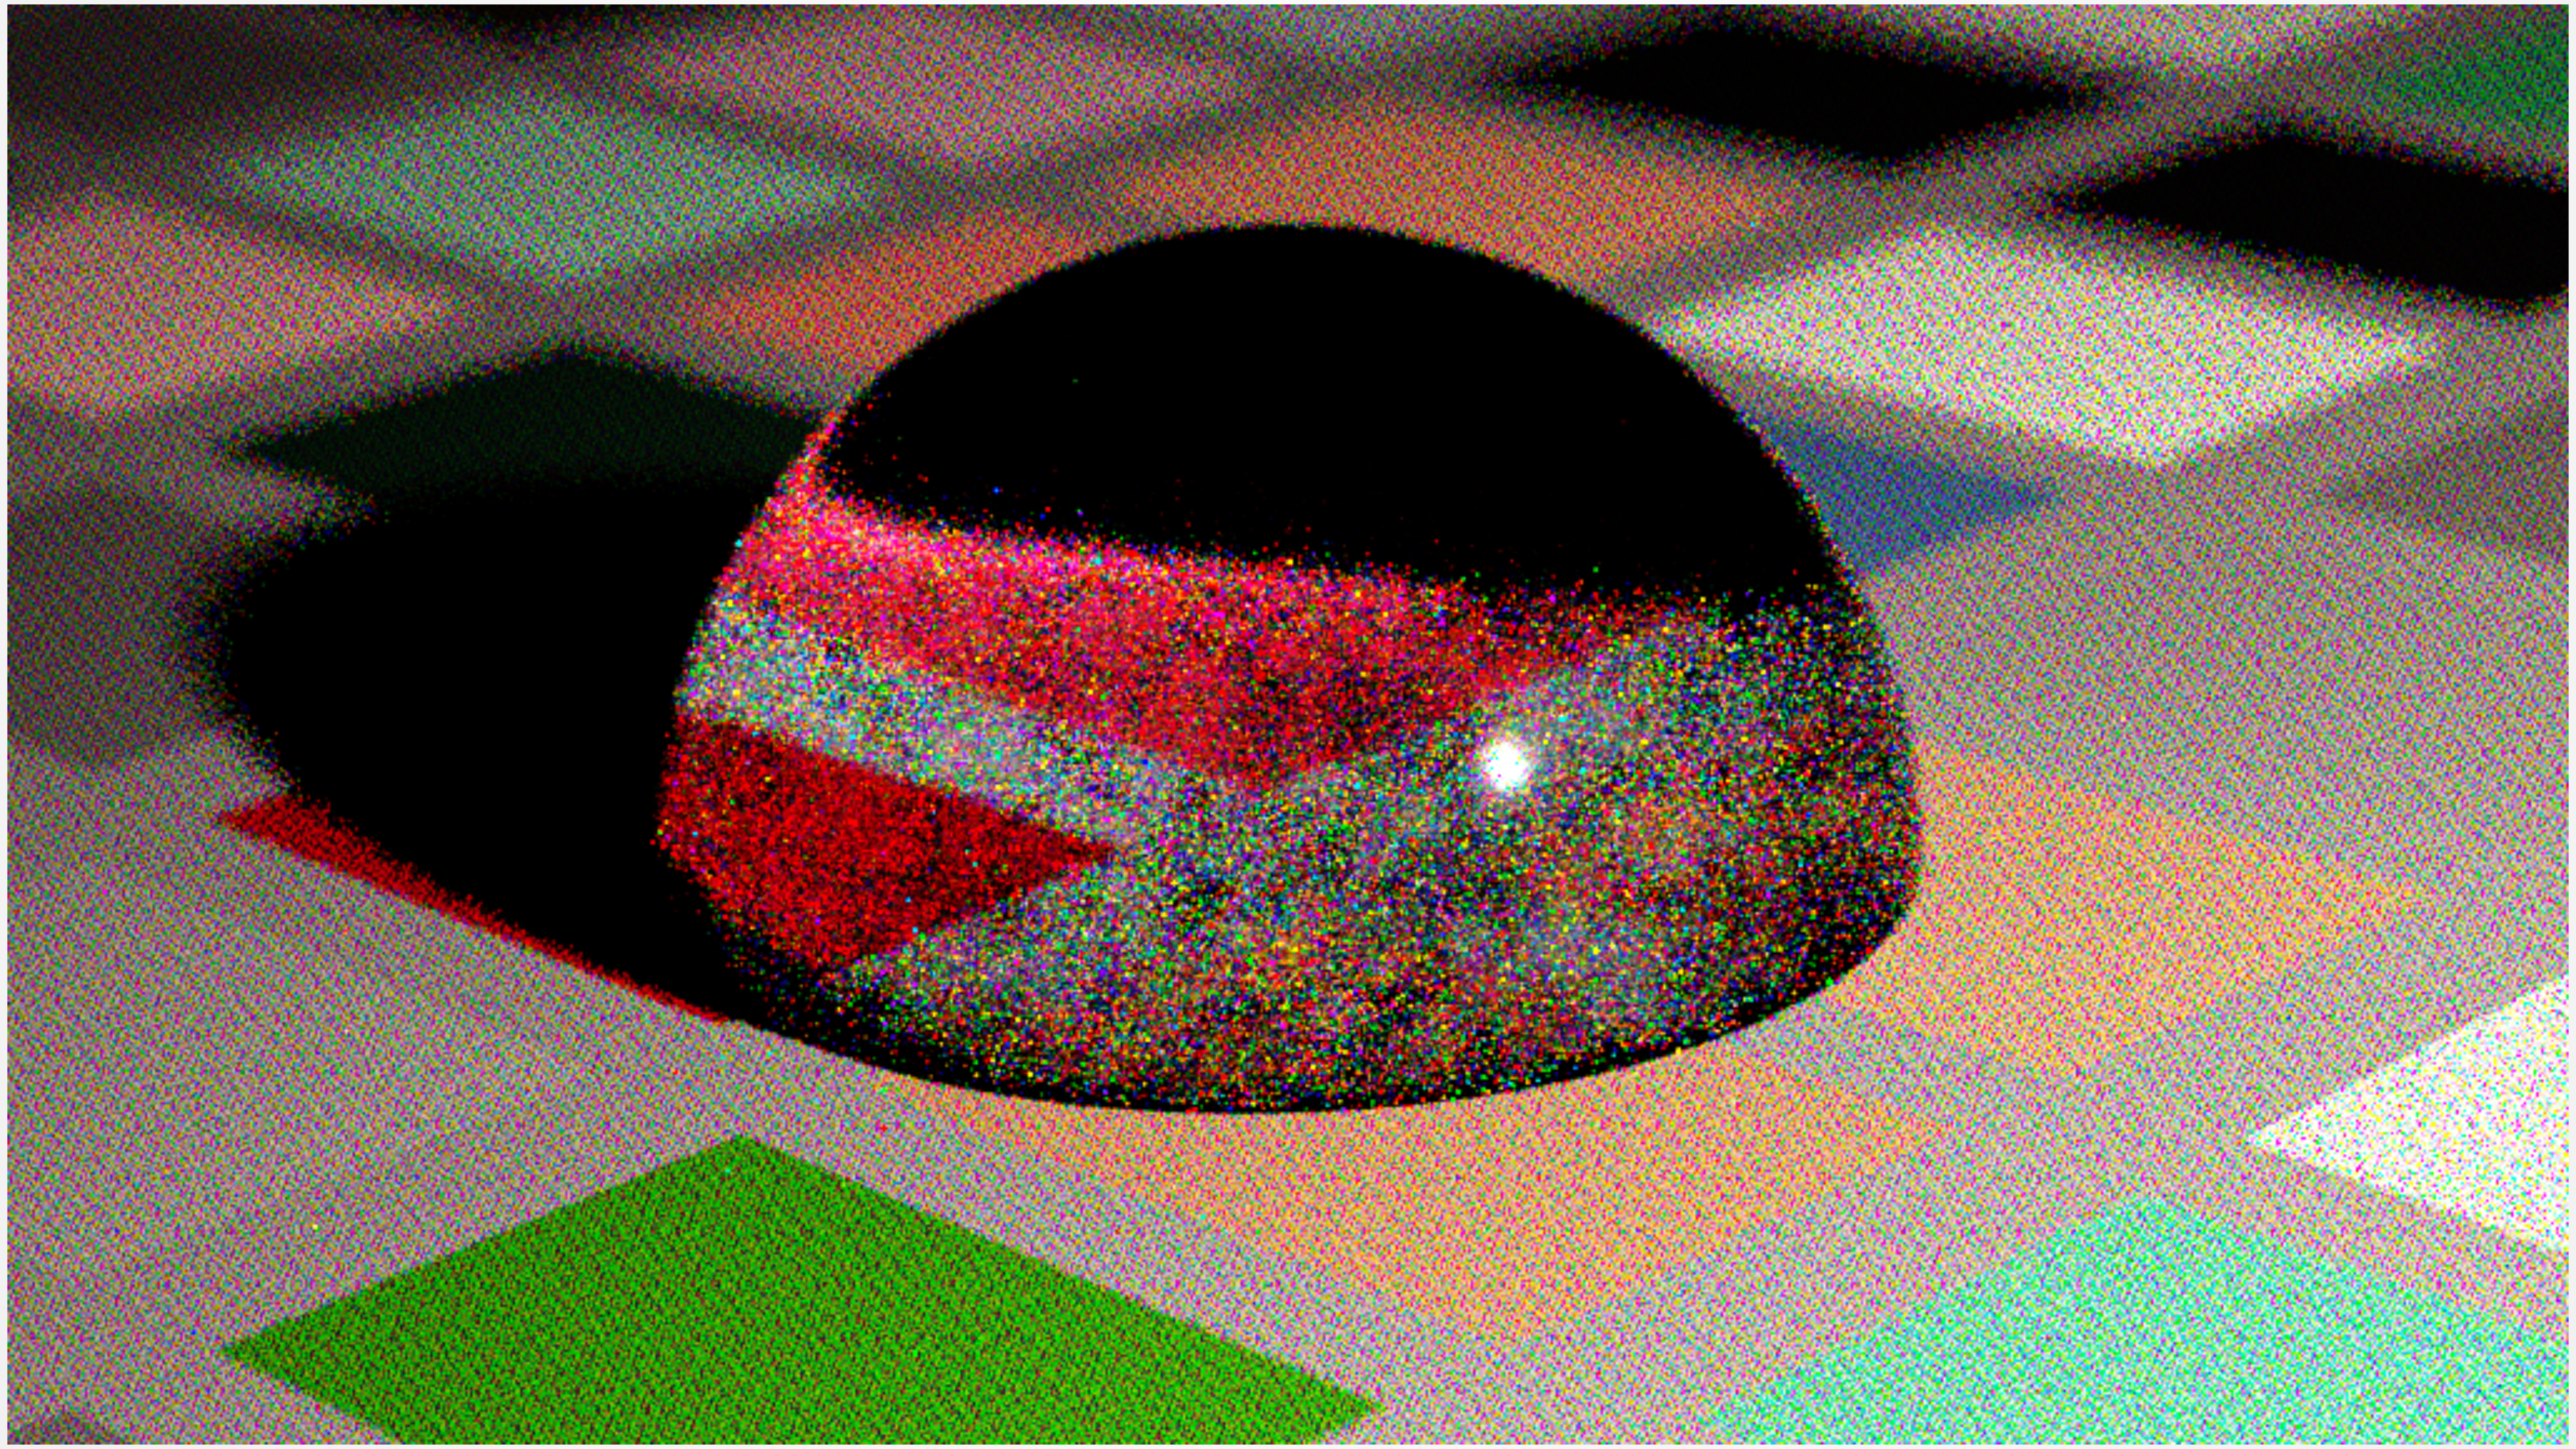

Guiding Behaviour

Guided path tracing - iteration 16

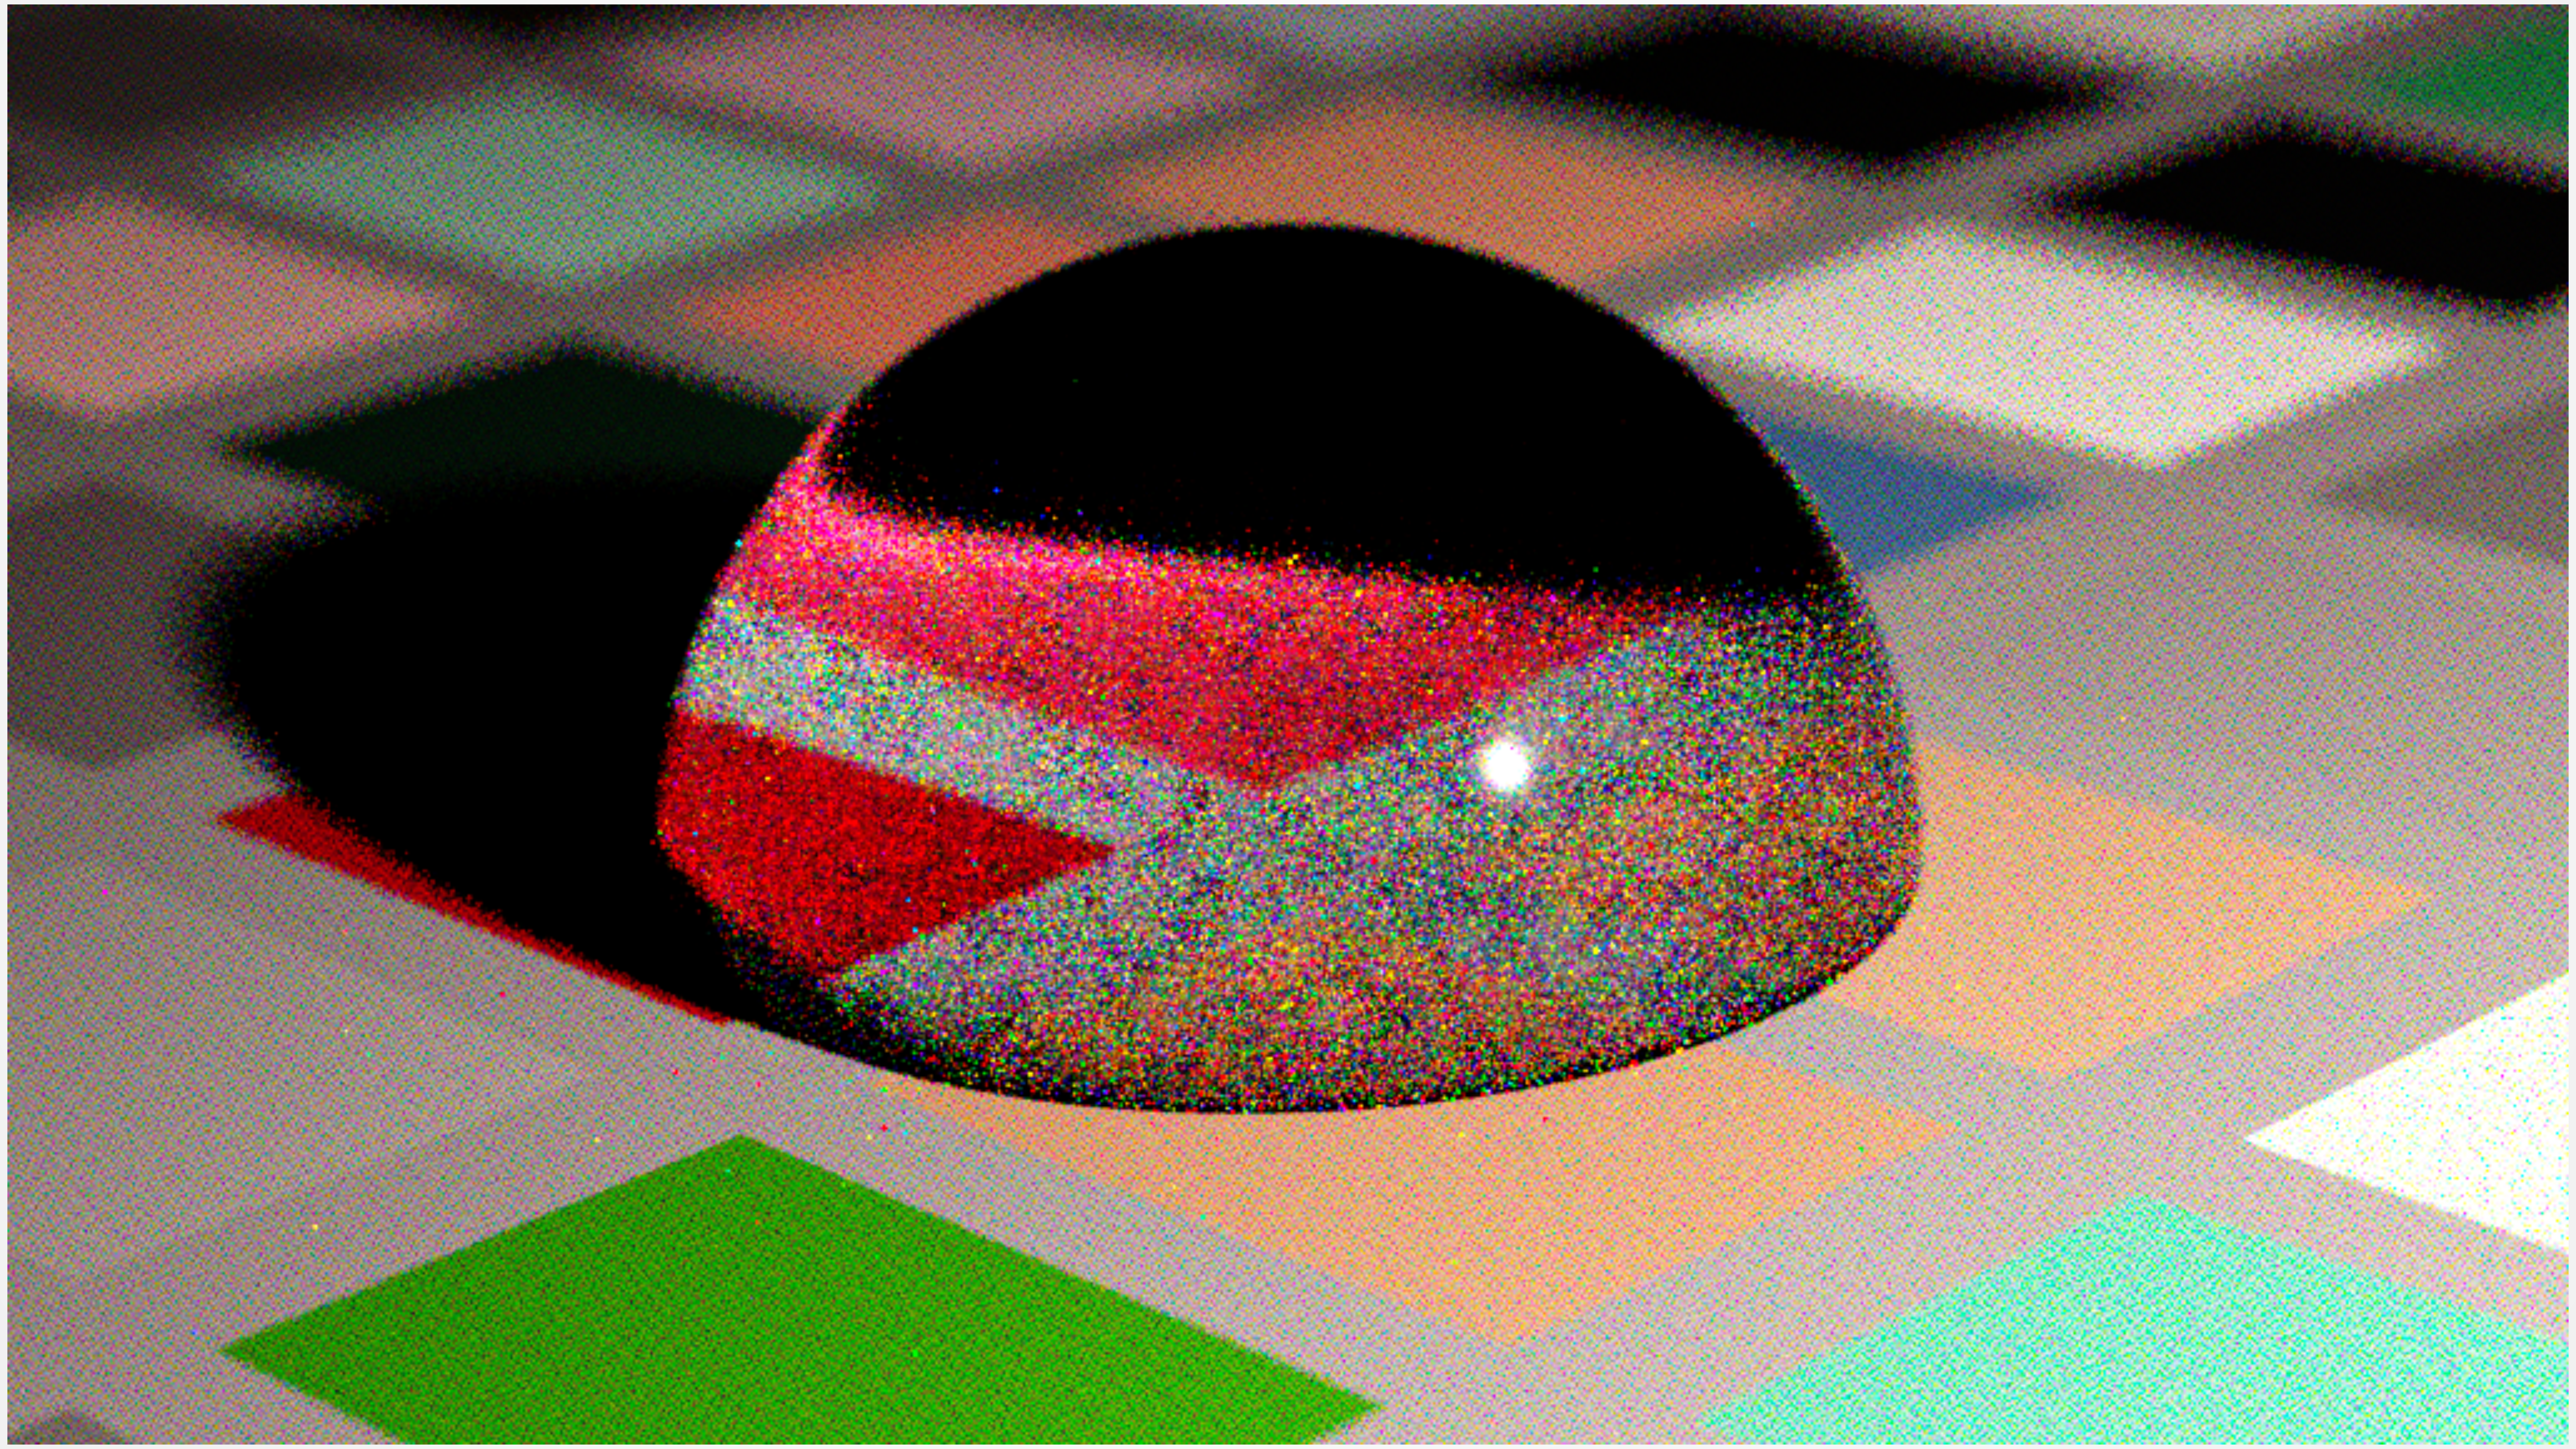

Recent related guiding methods

SIGGRAPH ASIA 2018 TOKYO

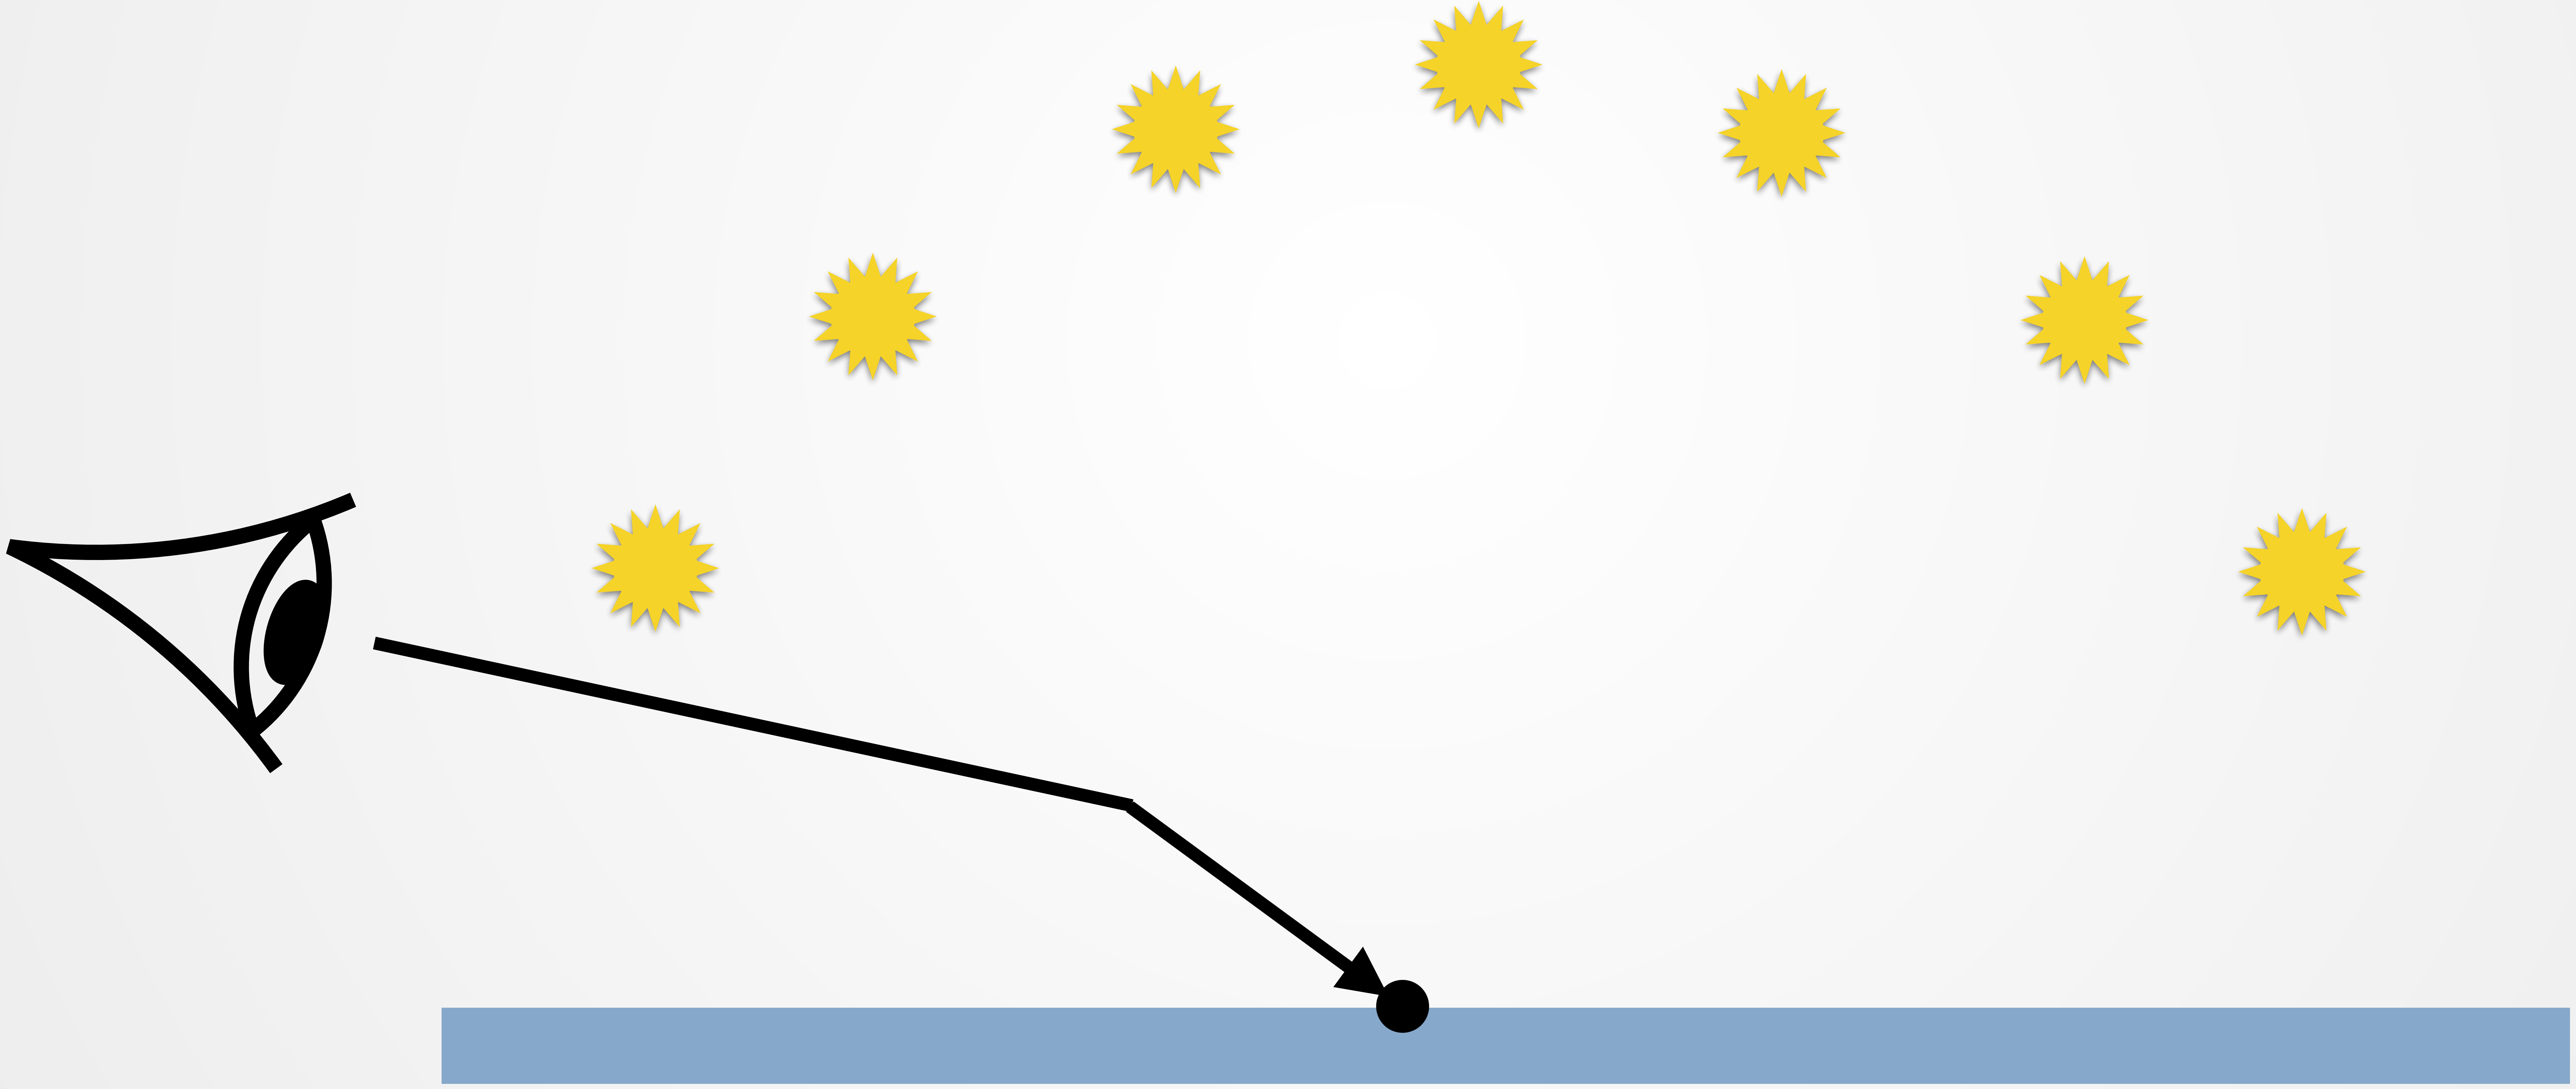

Recent related guiding methods

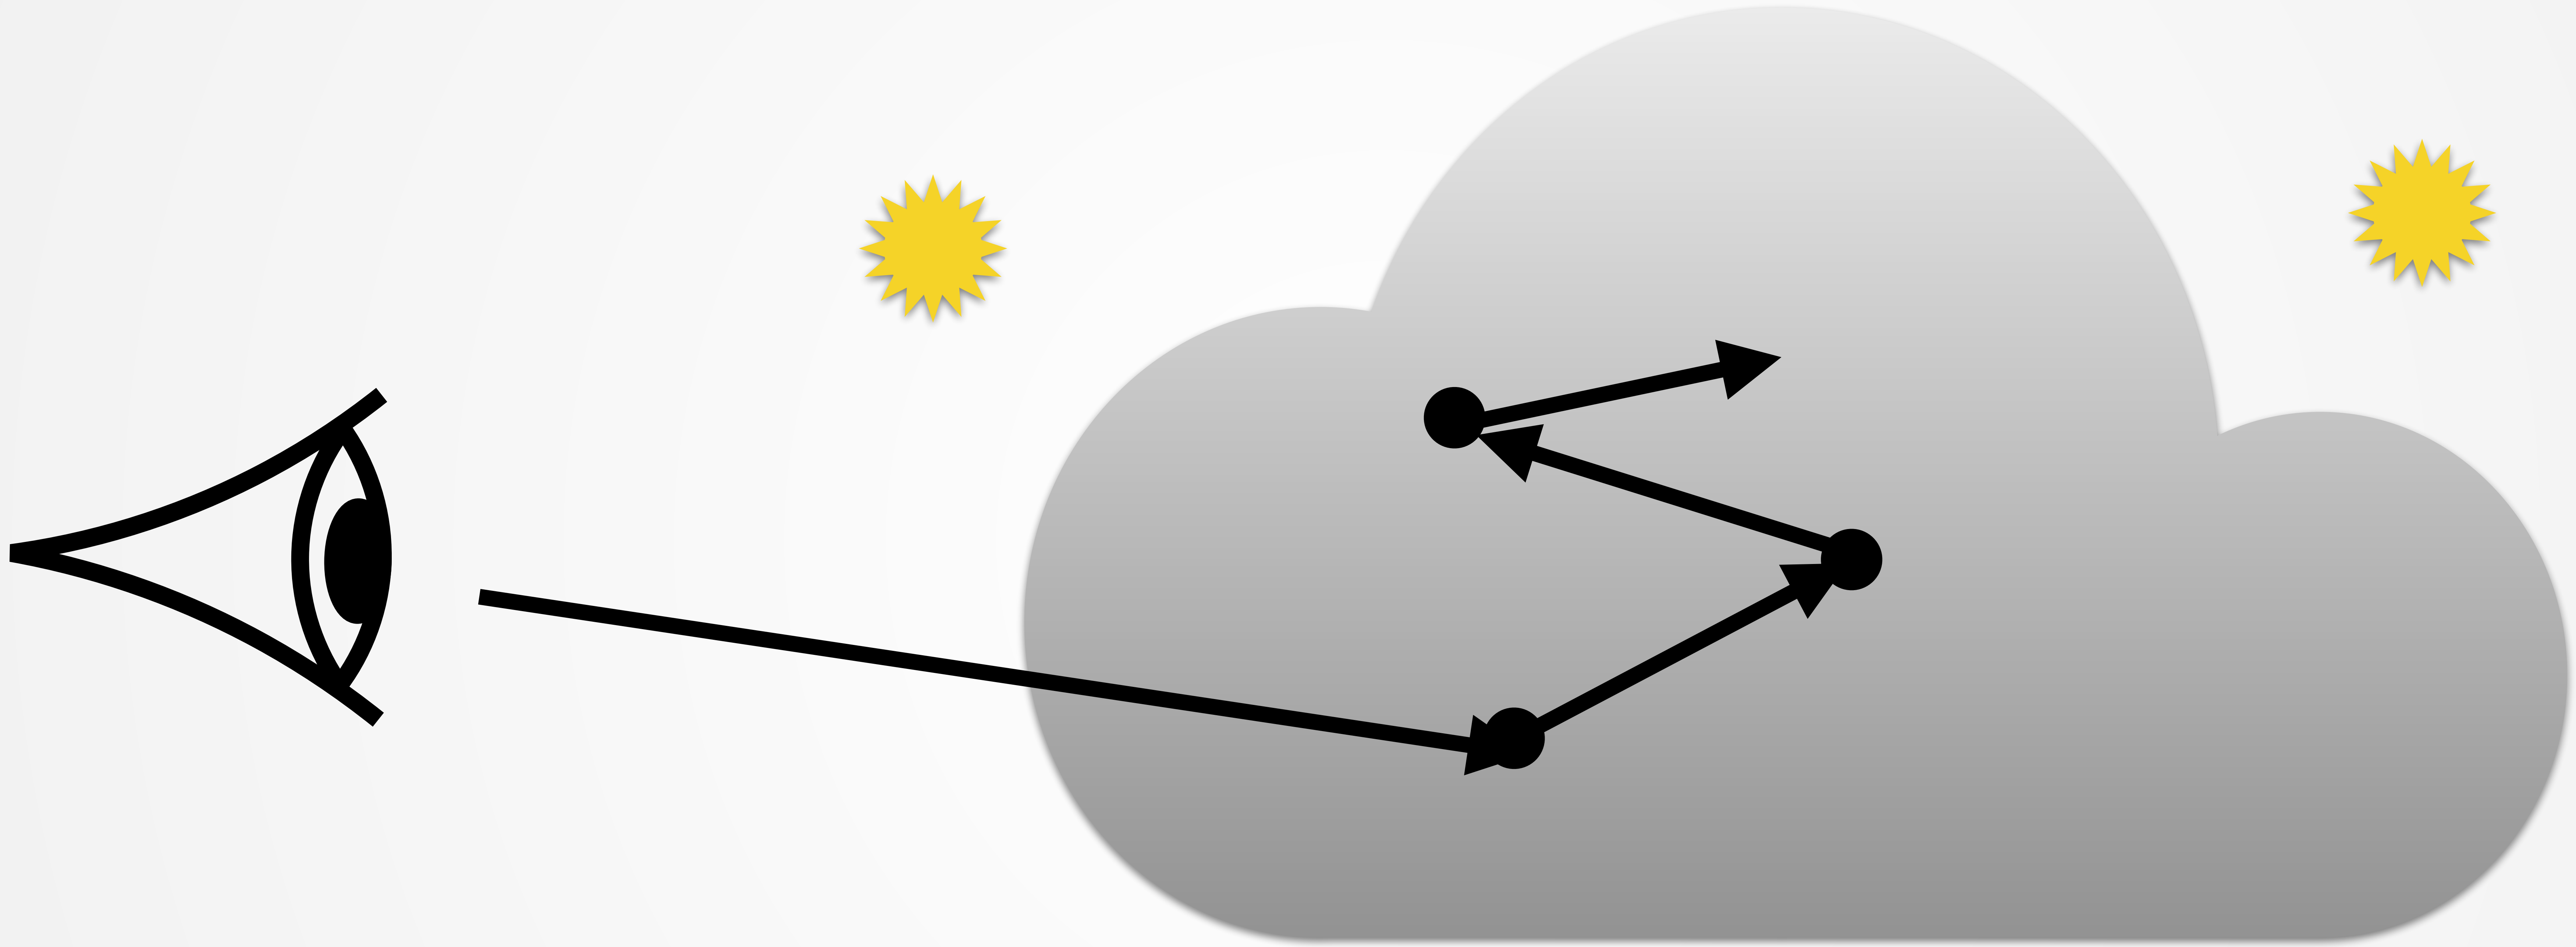

Conditional Gaussians

$$\Sigma = \begin{pmatrix} \Sigma_{11} \Sigma_{12} \\ \Sigma_{21} \Sigma_{22} \end{pmatrix}$$

$$\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}$$

$$\mu_x = \mu \,|\, x$$

$$\Sigma_x = \Sigma \,|\, x$$

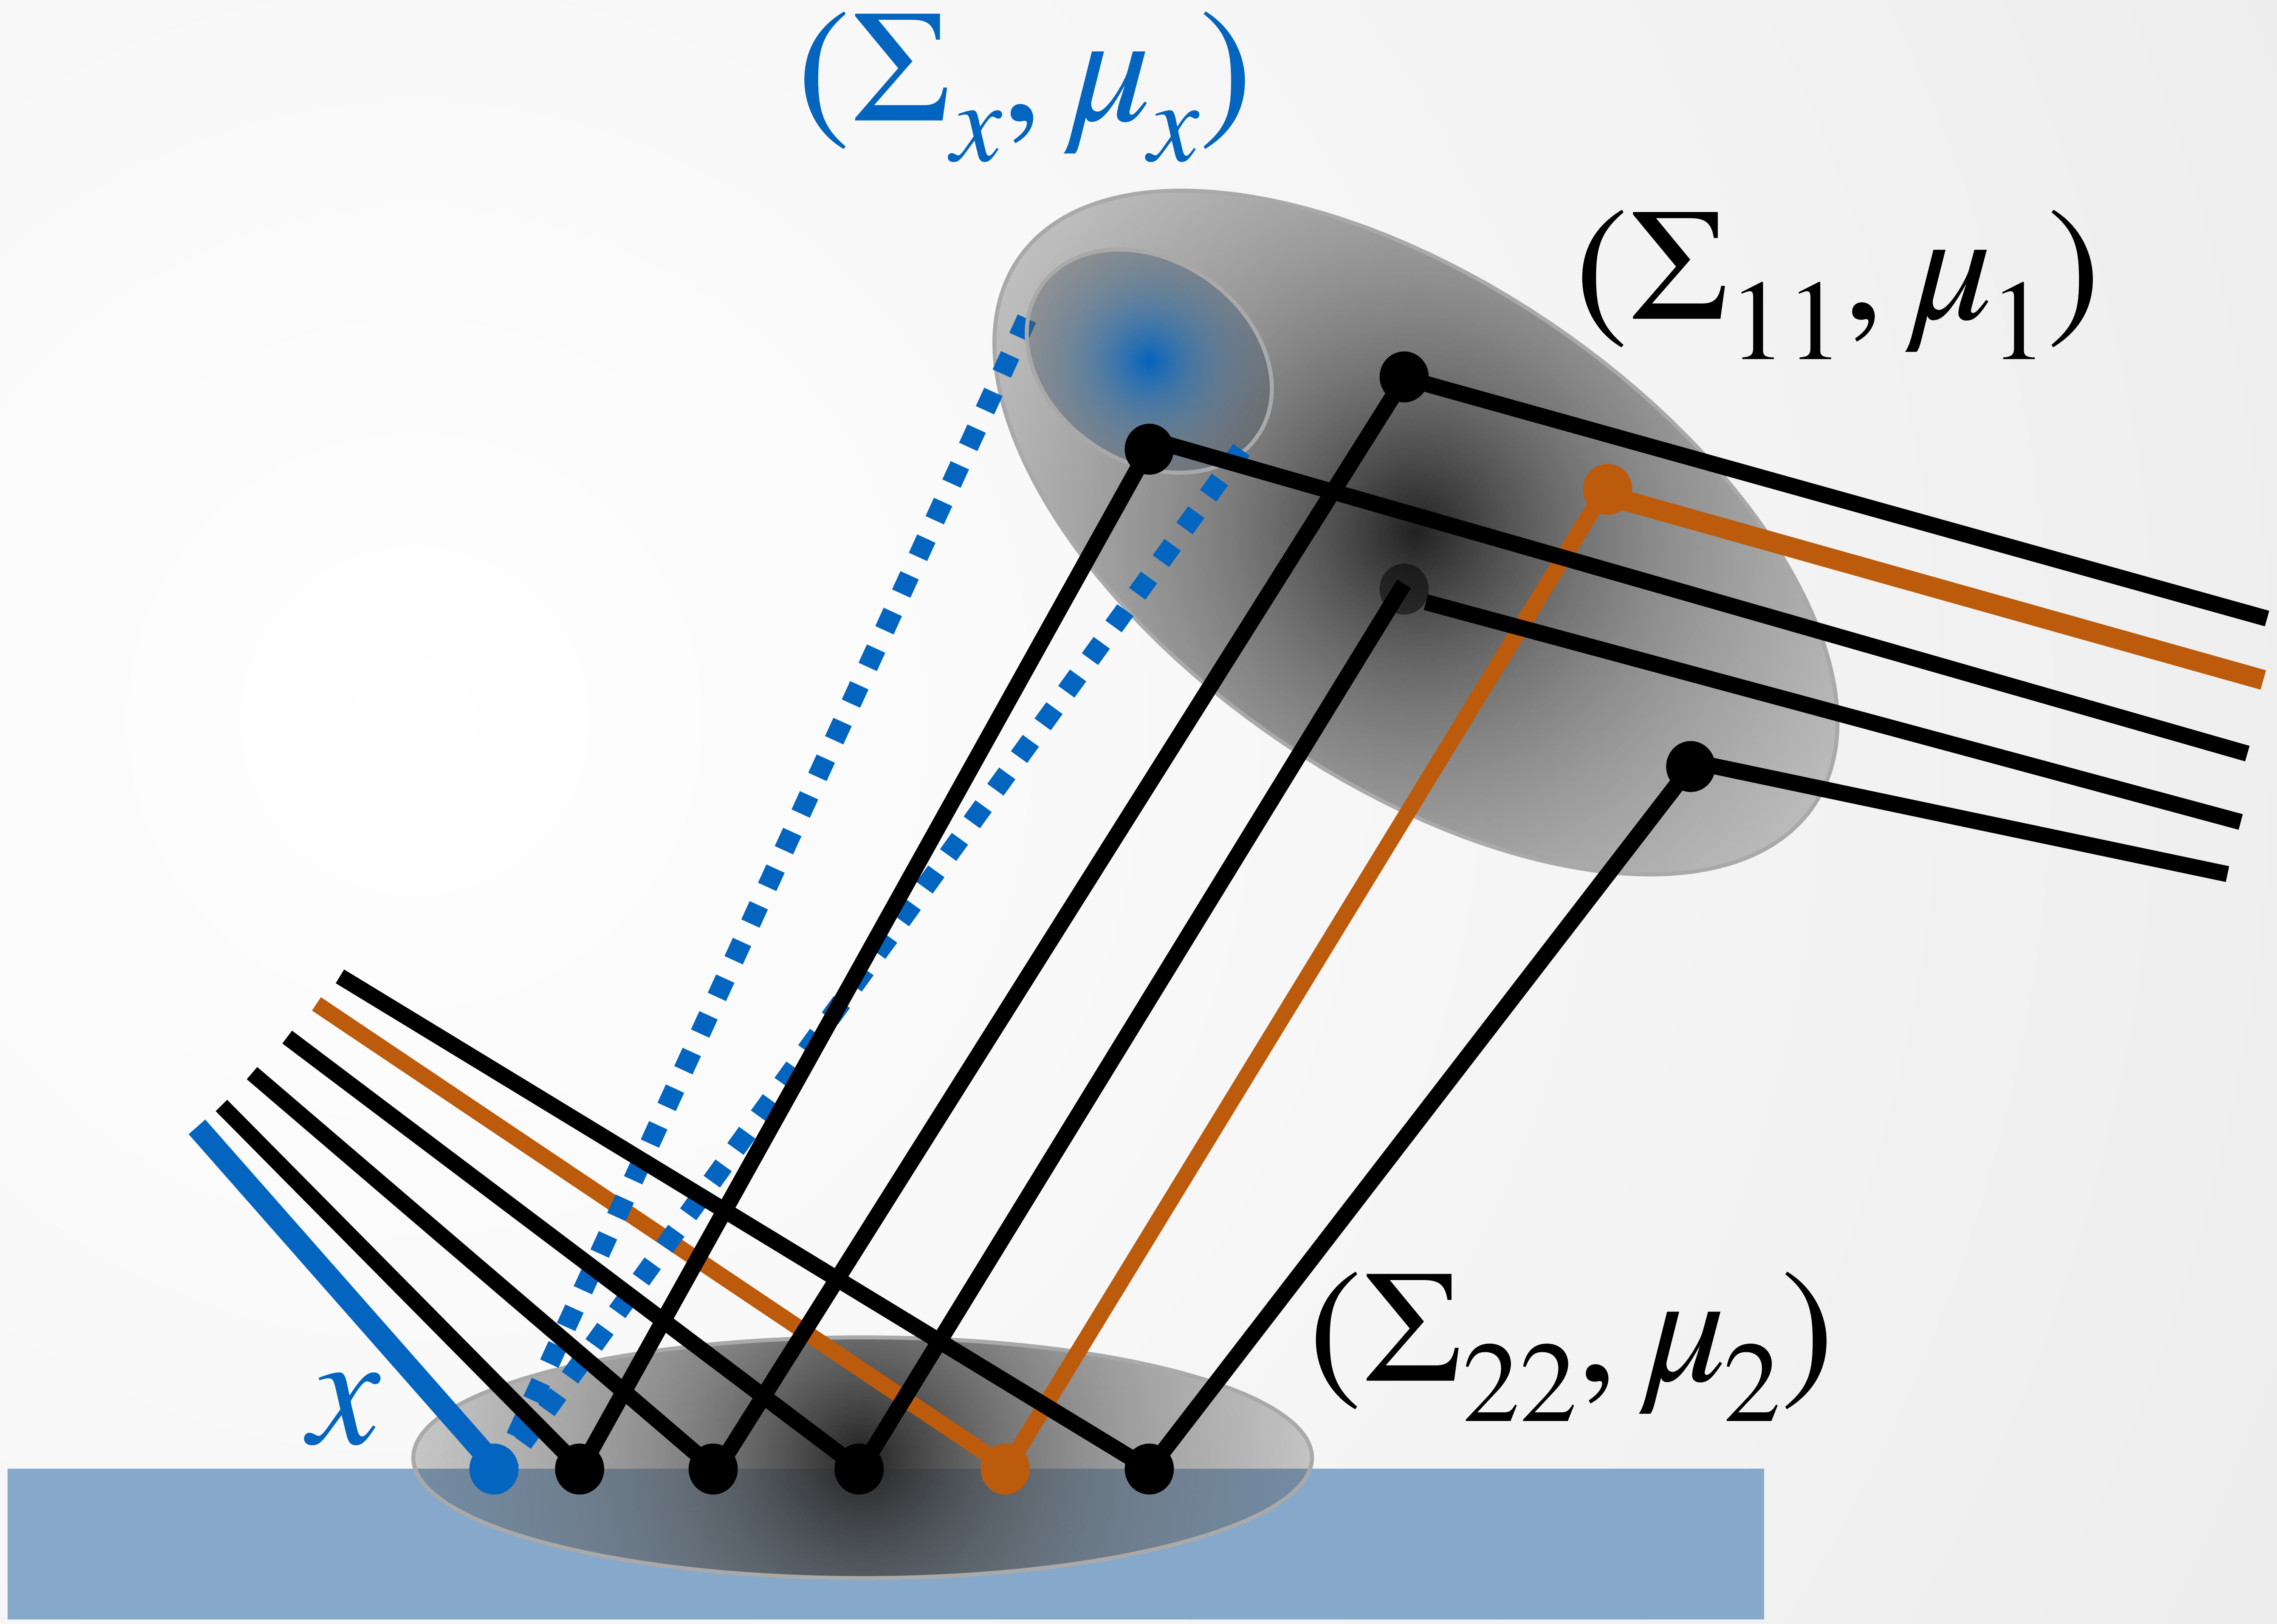

Conditional Gaussians

$$\mu_x = \mu_1 + \Sigma_{12}\Sigma_{22}^{-1}(x - \mu_2)$$

$$\Sigma_x = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21}$$

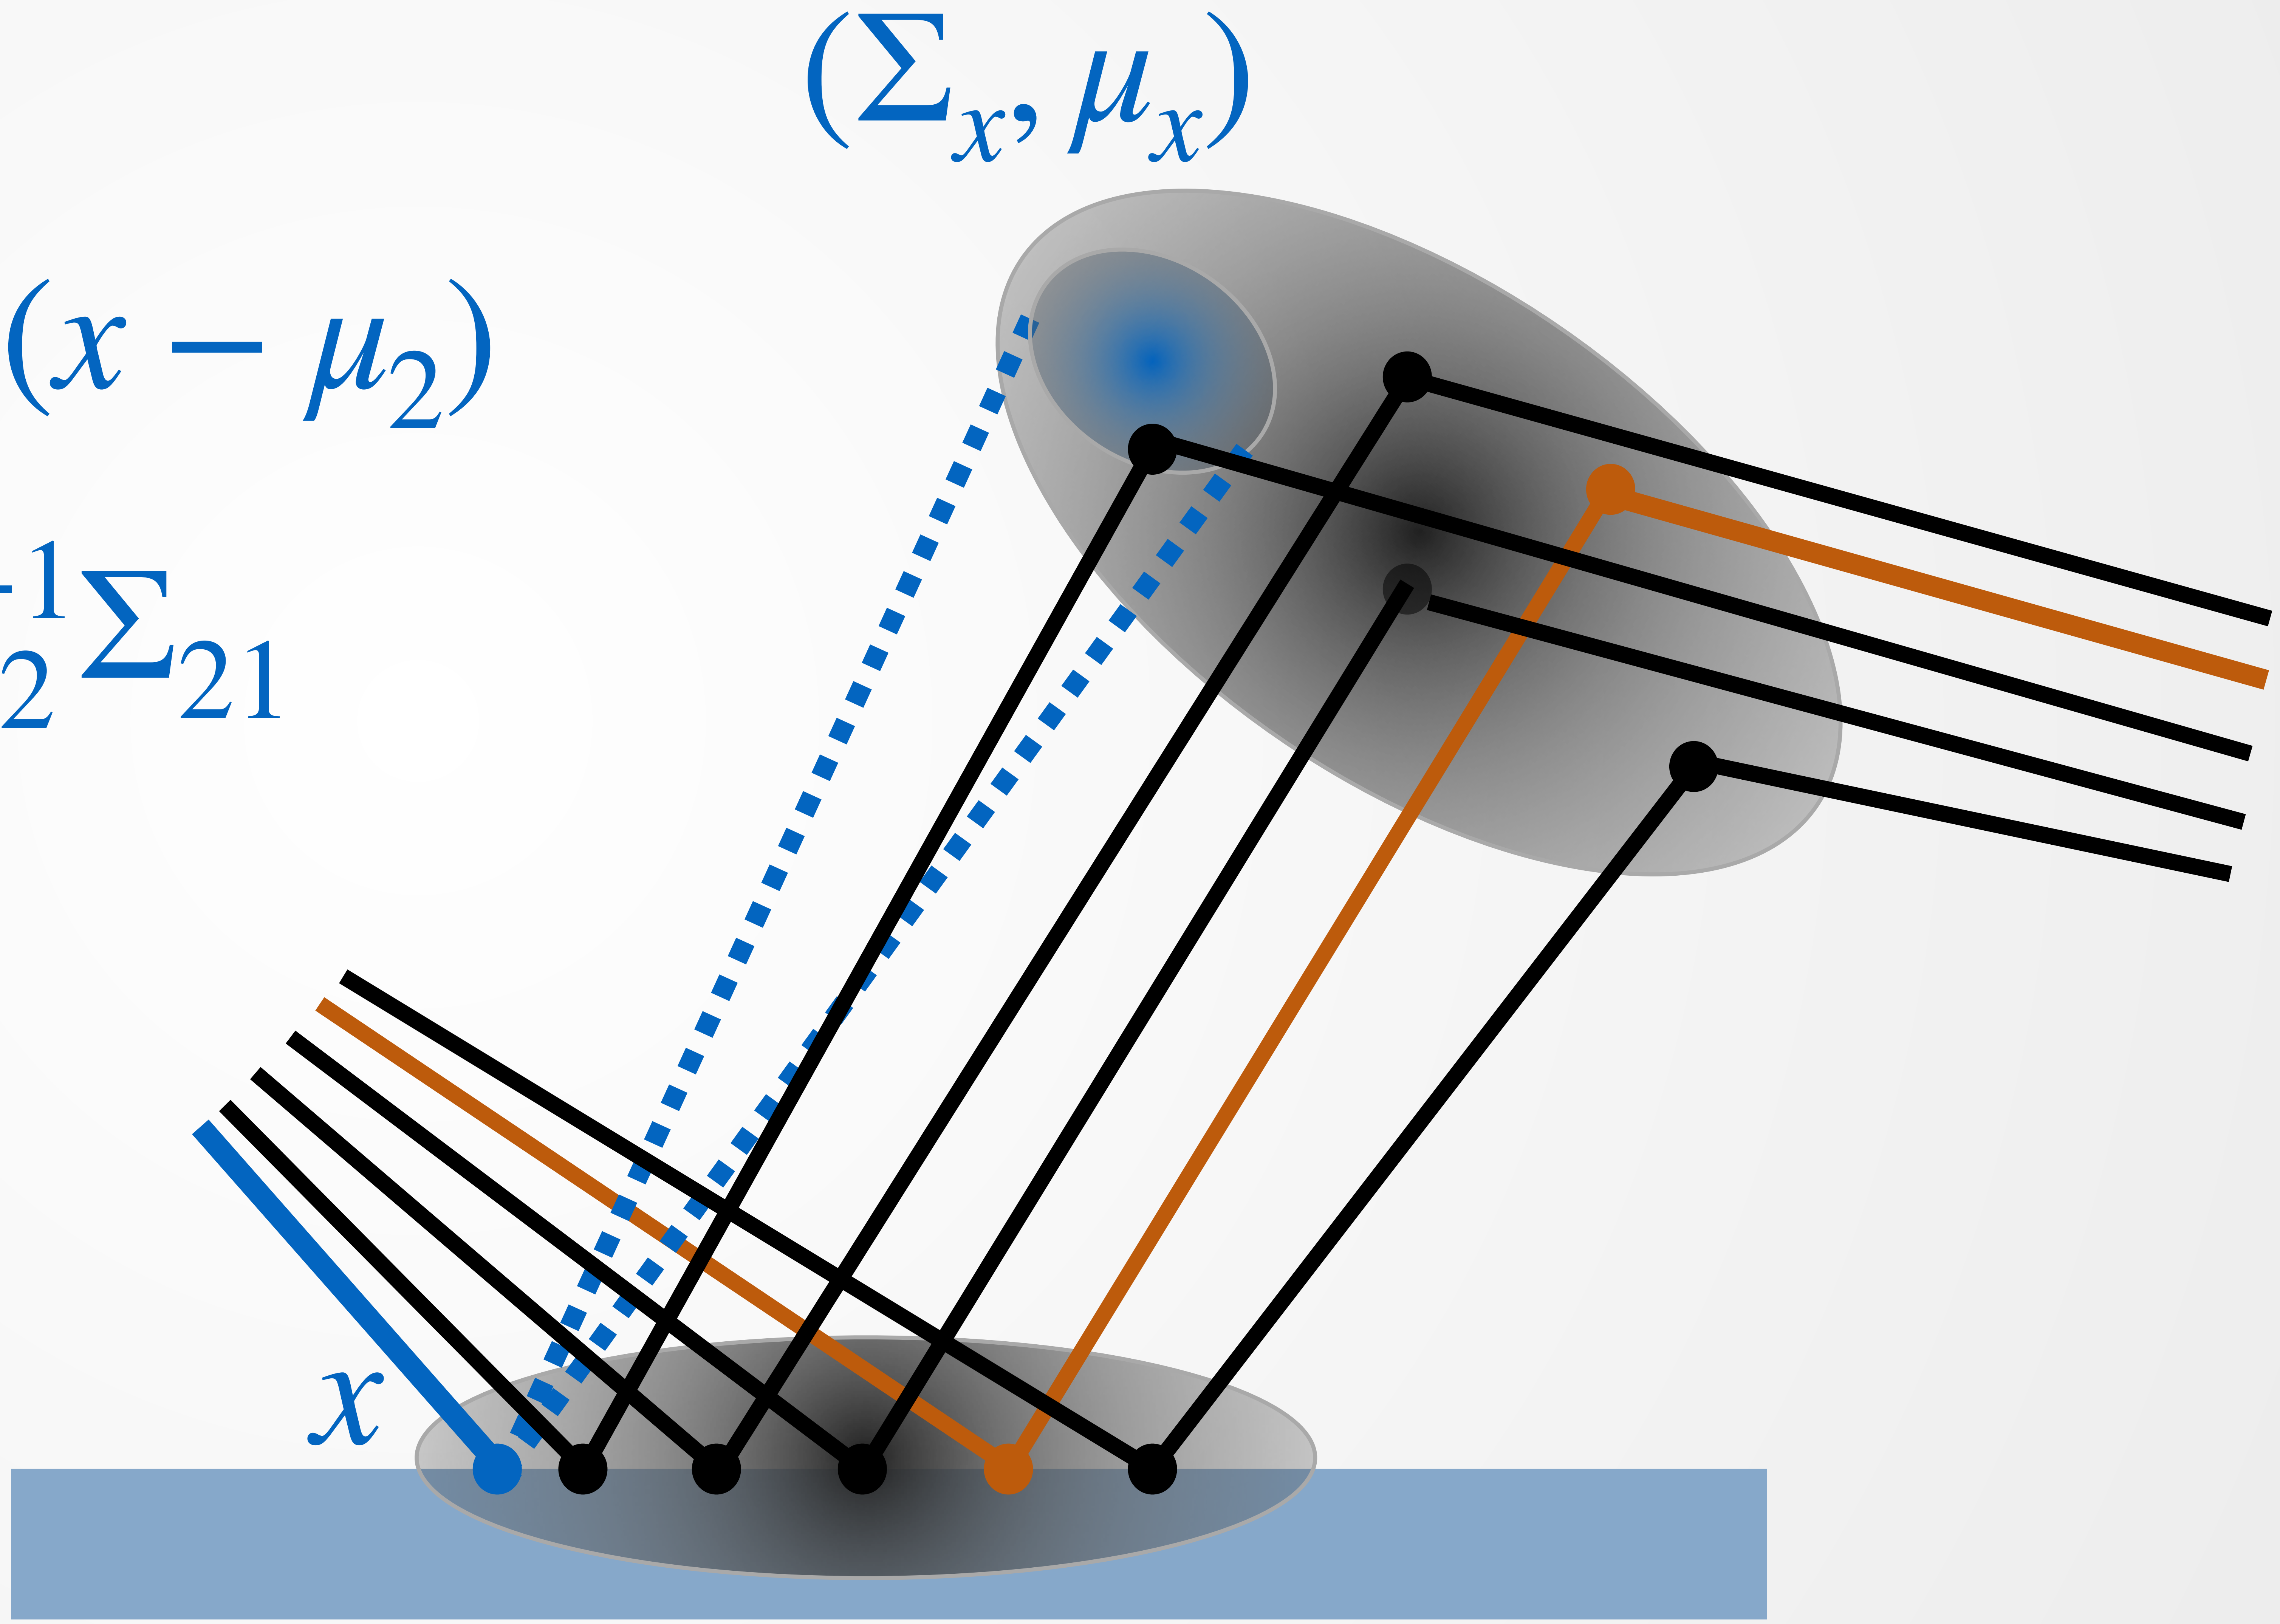

Selecting Guide Paths

- DBOR: DeCoro et al. [2010] / Zirr et al. [2018]
- Framebuffer cascade with histogram
- Samples split according to throughput

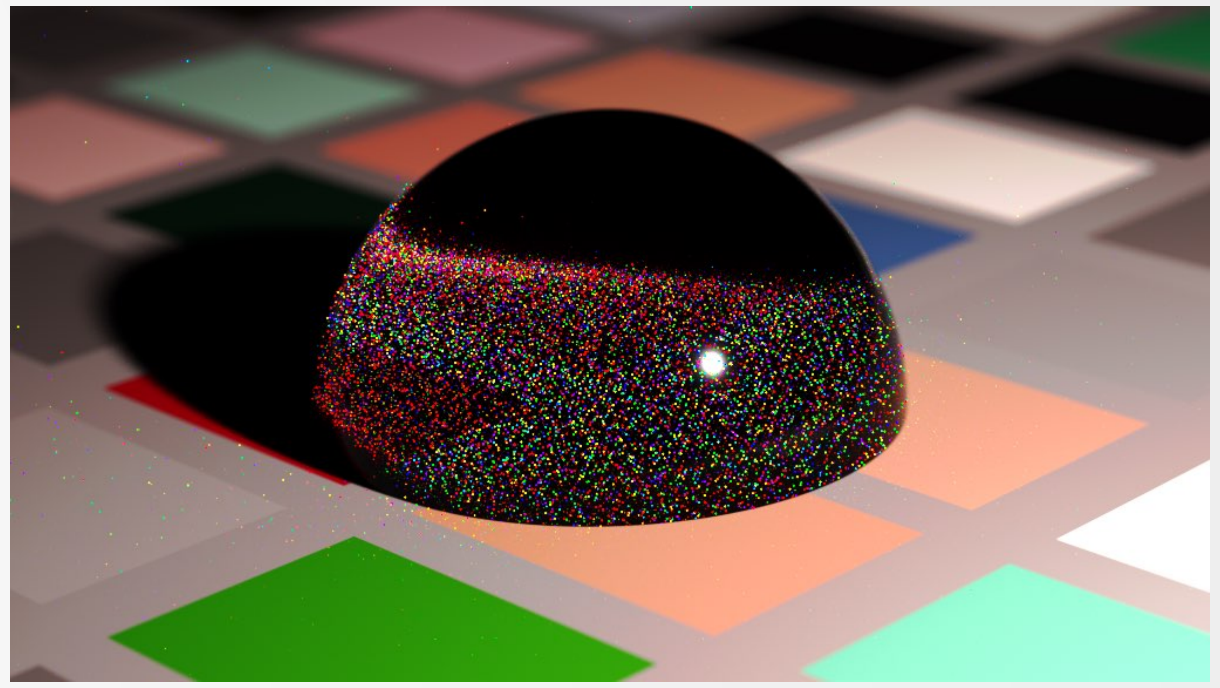

$C \in [0, 8^1]$ $C \in [8^1, 8^2]$ $C \in [8^2, 8^3]$ $C \in [8^3, \infty)$

Remaining outliers

- Either: stop adding, clear the framebuffer and restart rendering.
  - Outliers possible because of yet unexplored lighting features or gaps between guide paths.
- Or: keep adding, don't clear the framebuffer
  - Outliers will remain in the framebuffer and won't converge away in reasonable time.

Dining Room - 30min

reference

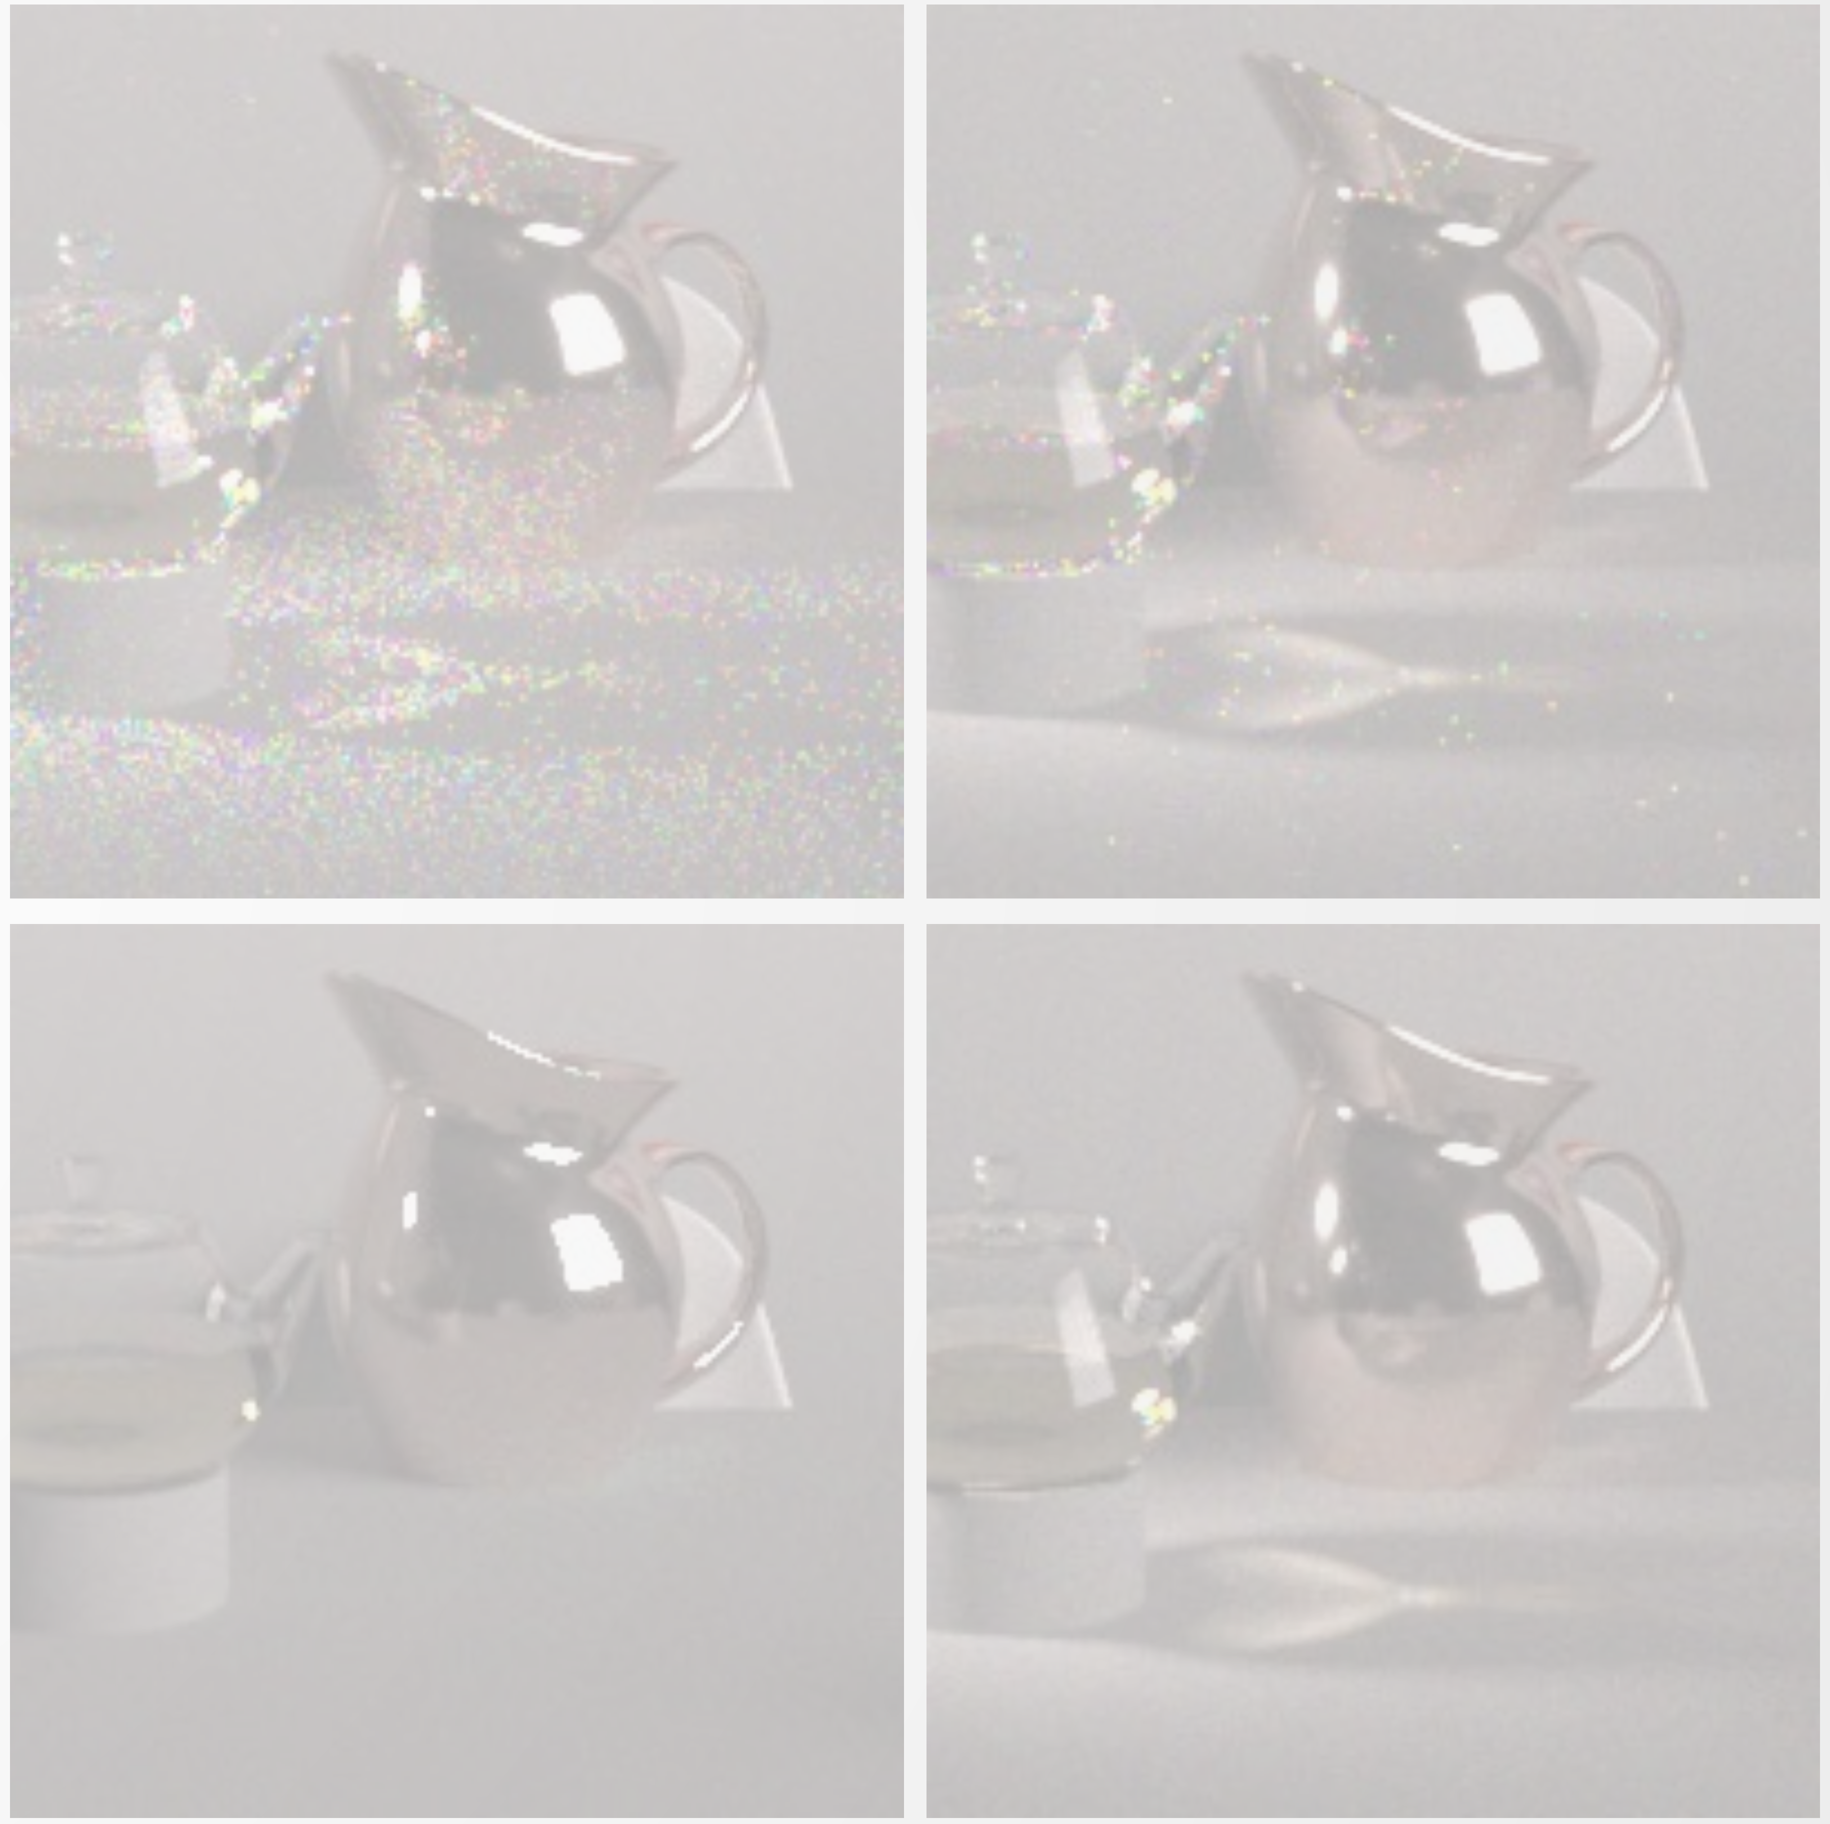

Dining Room - 30min

path tracing

RMSE 0.0435

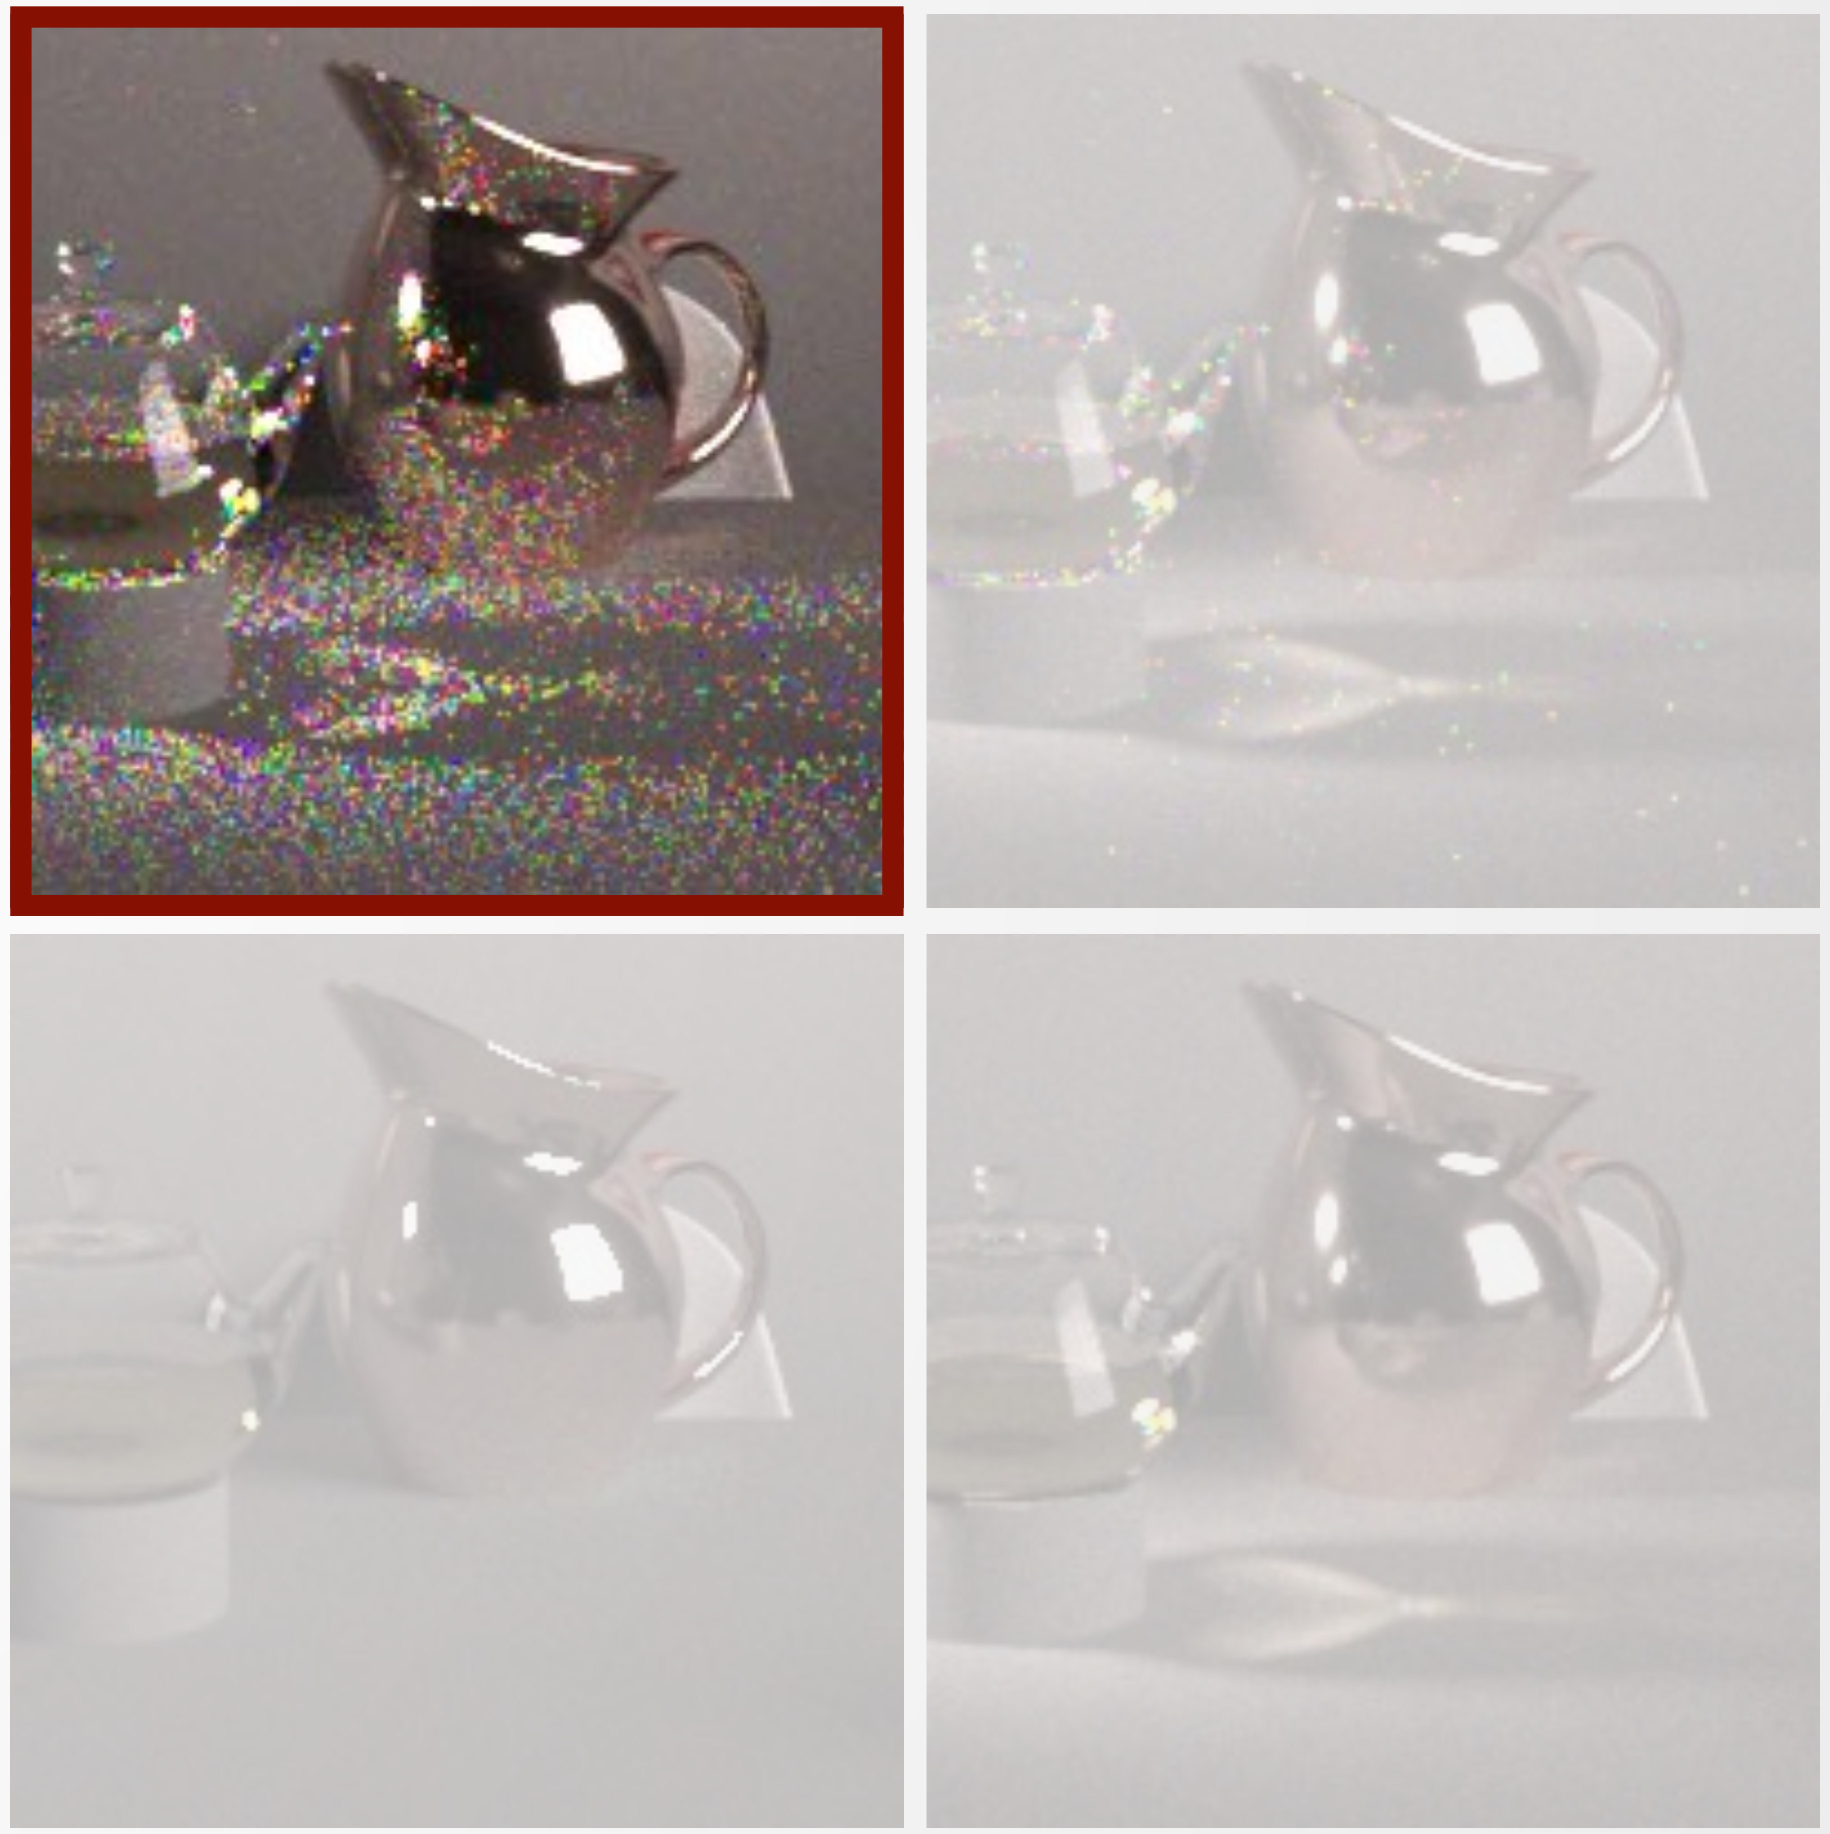

KIT Karlsruhe Institute of Technology

Dining Room - 30min

guided path tracing

RMSE 0.0435 RMSE 0.0345

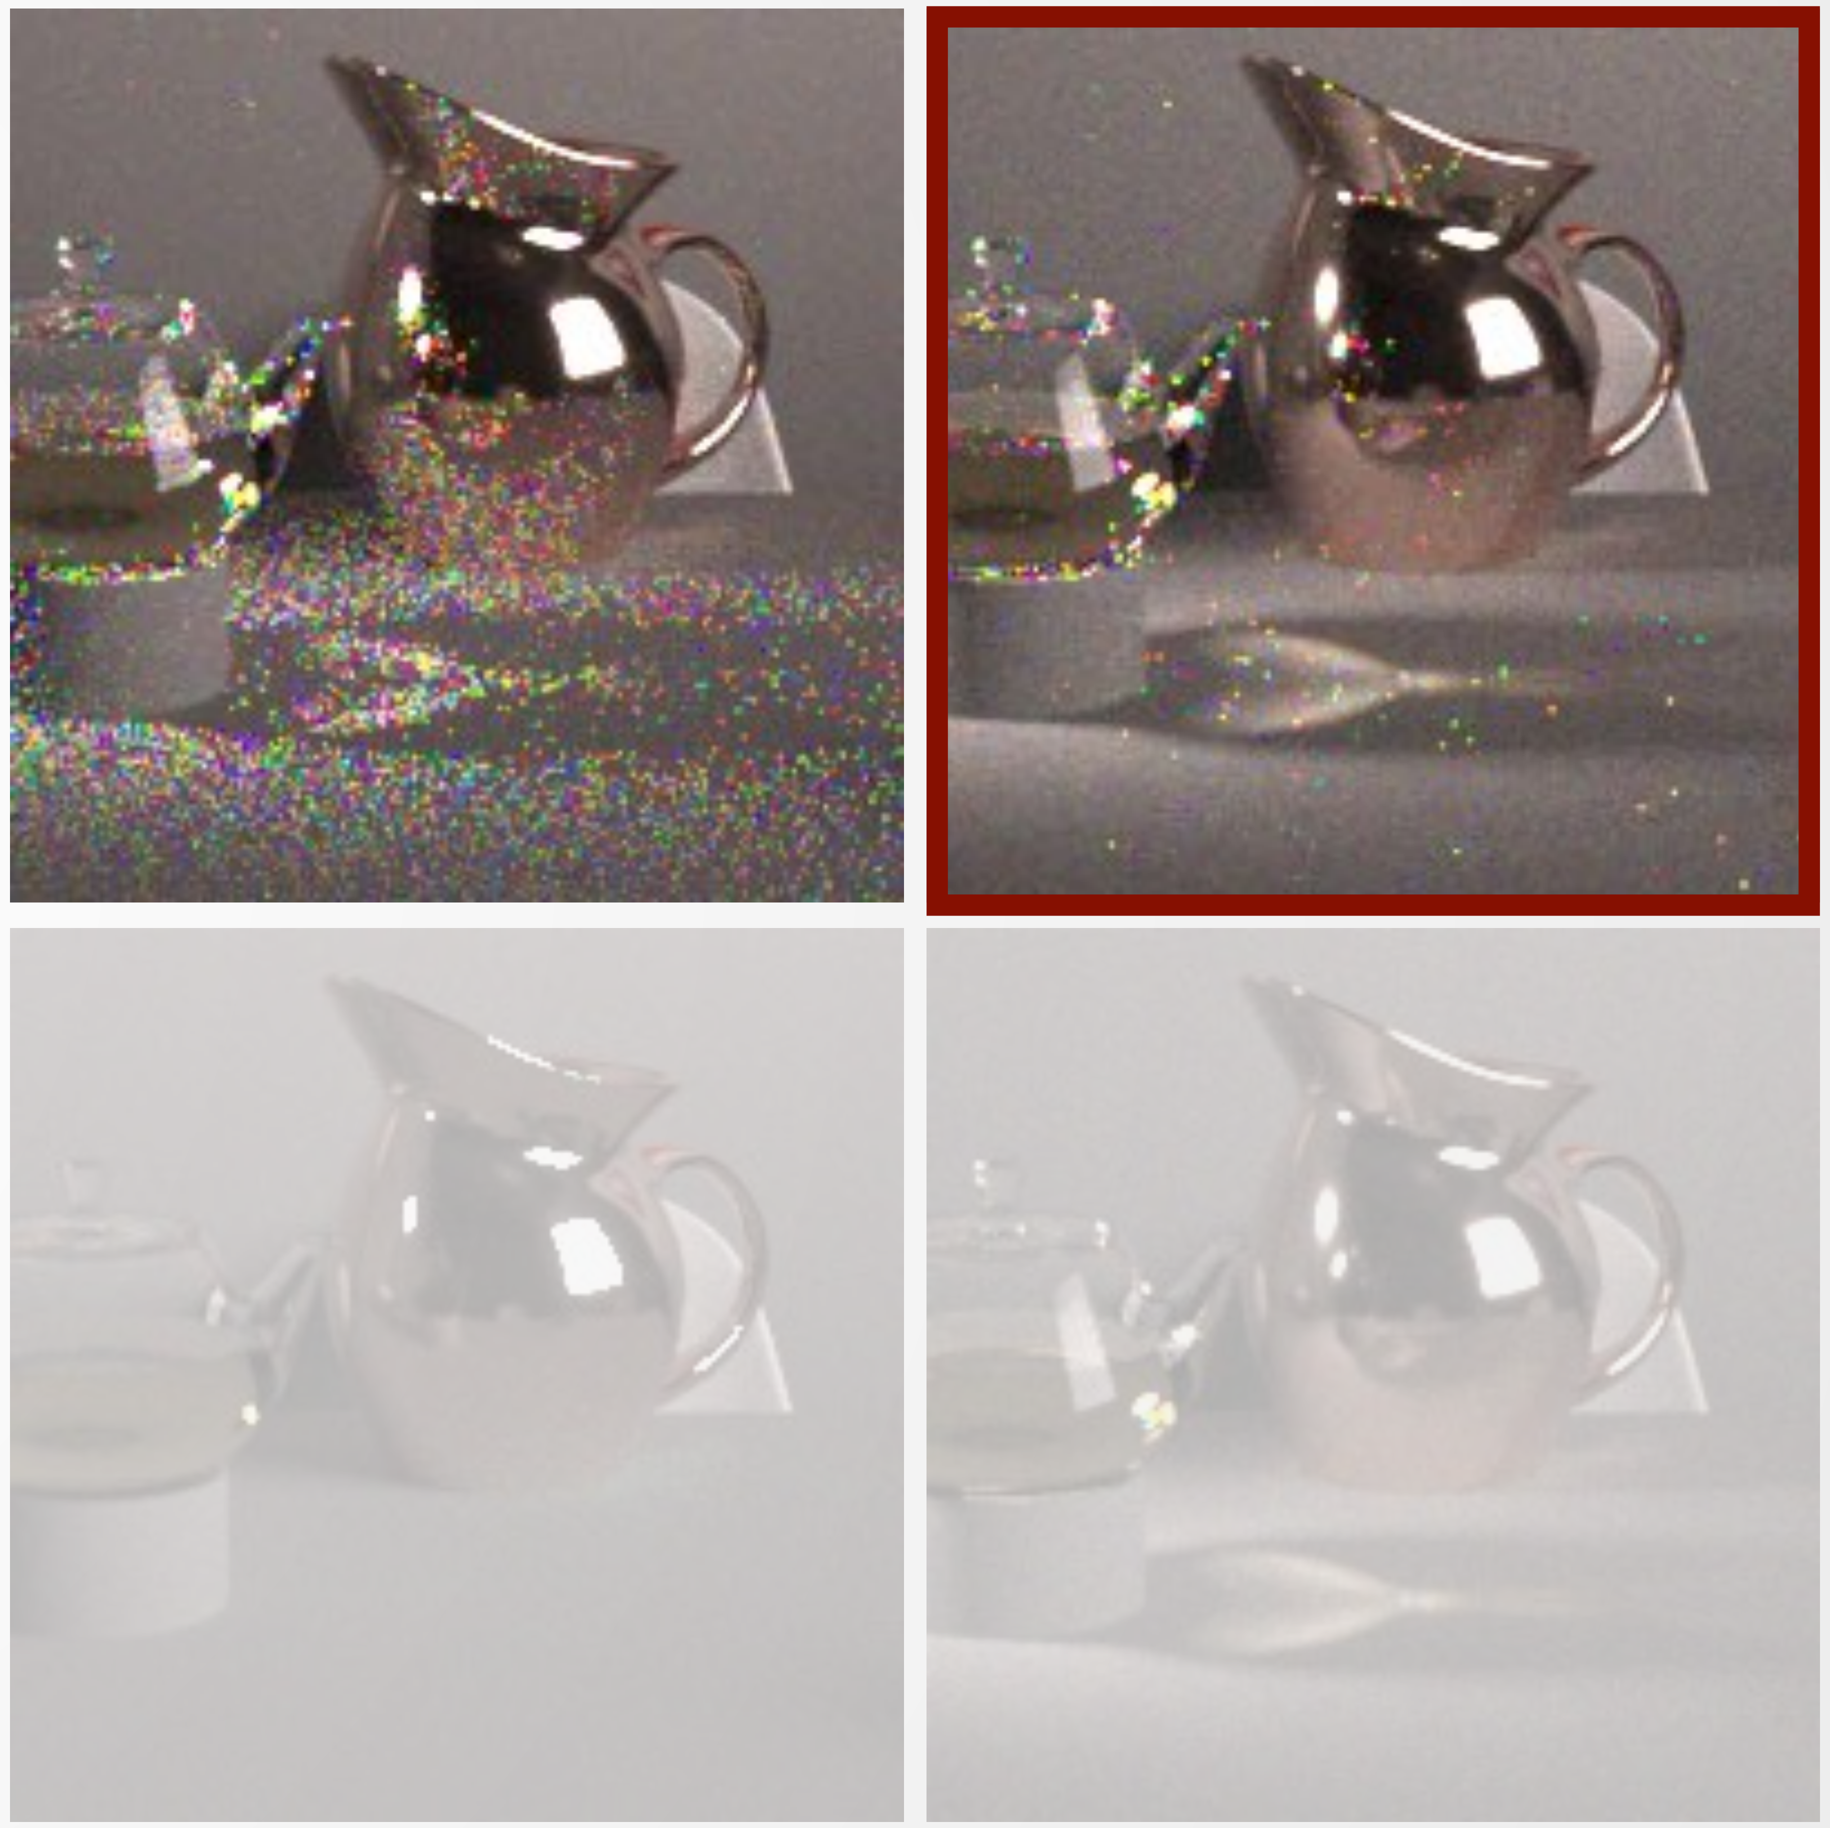

Dining Room - 30min

guided path tracing + DBOR

RMSE 0.0435

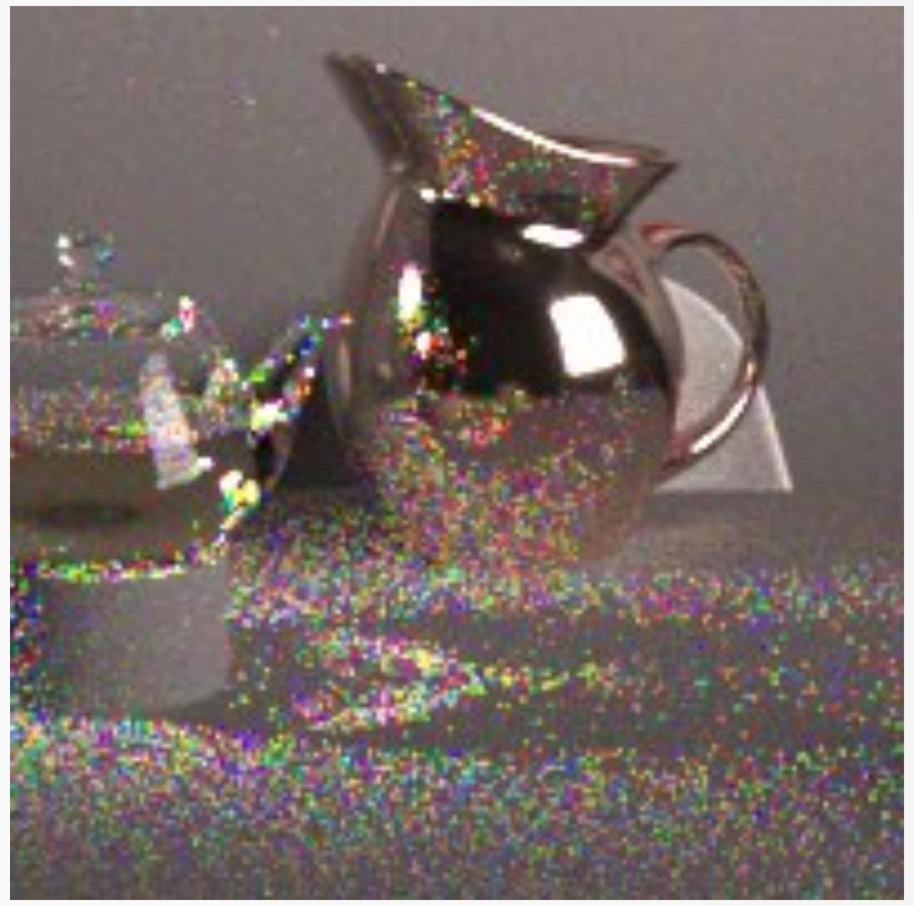

RMSE 0.0345

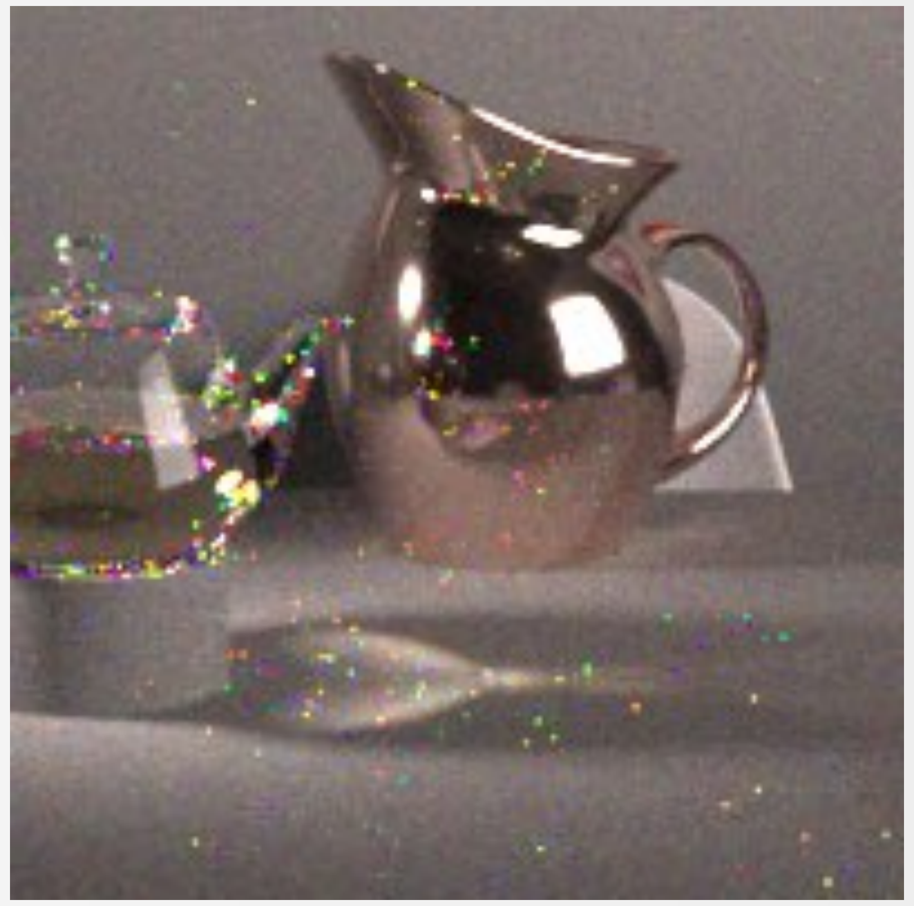

RMSE 0.0231

Dining Room - 30min

path tracing + DBOR

RMSE 0.0435

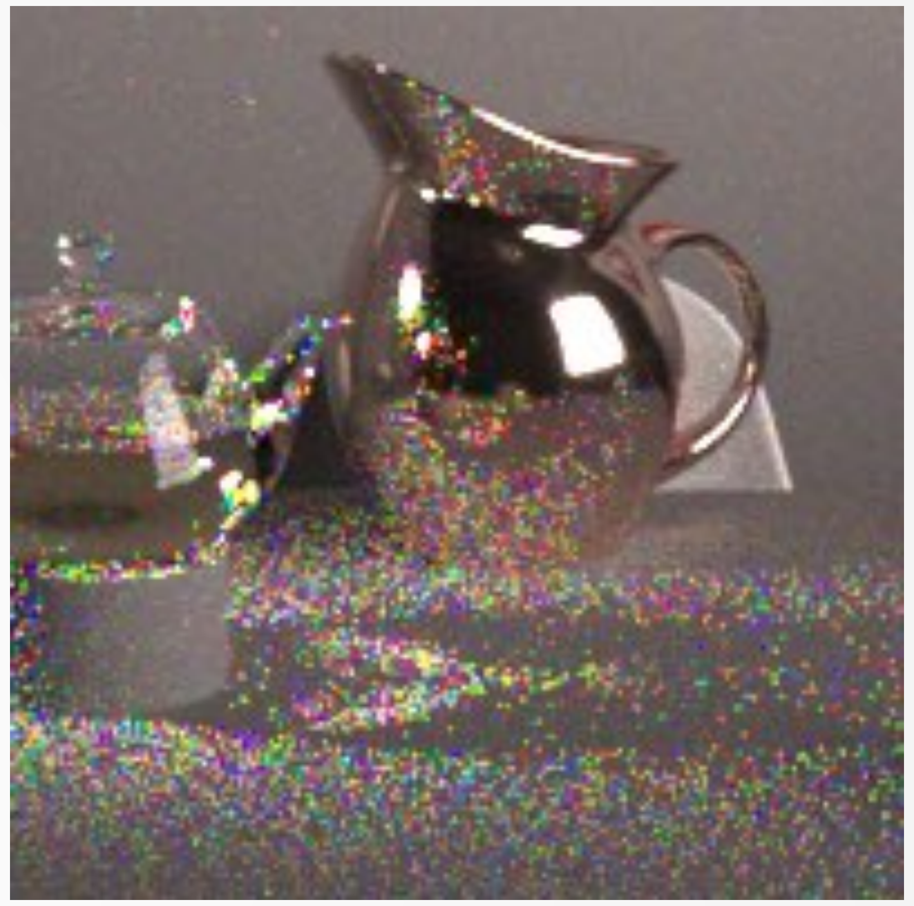

RMSE 0.0345

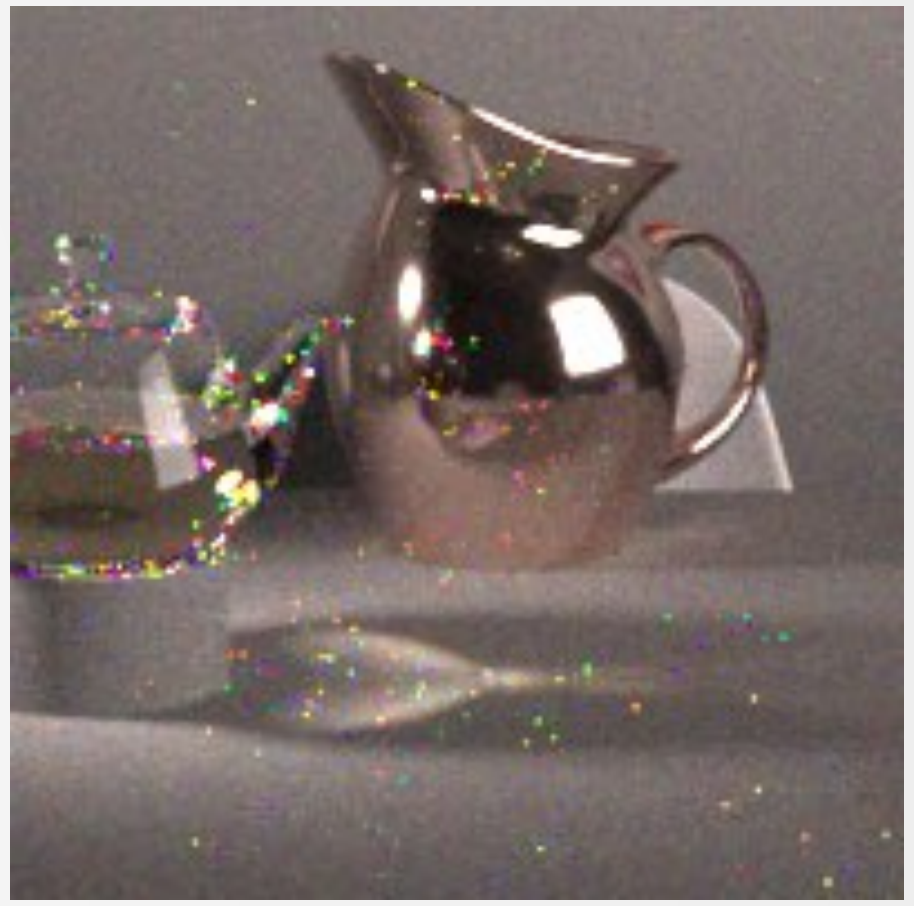

RMSE 0.1170

RMSE 0.0231

Pool - 30min

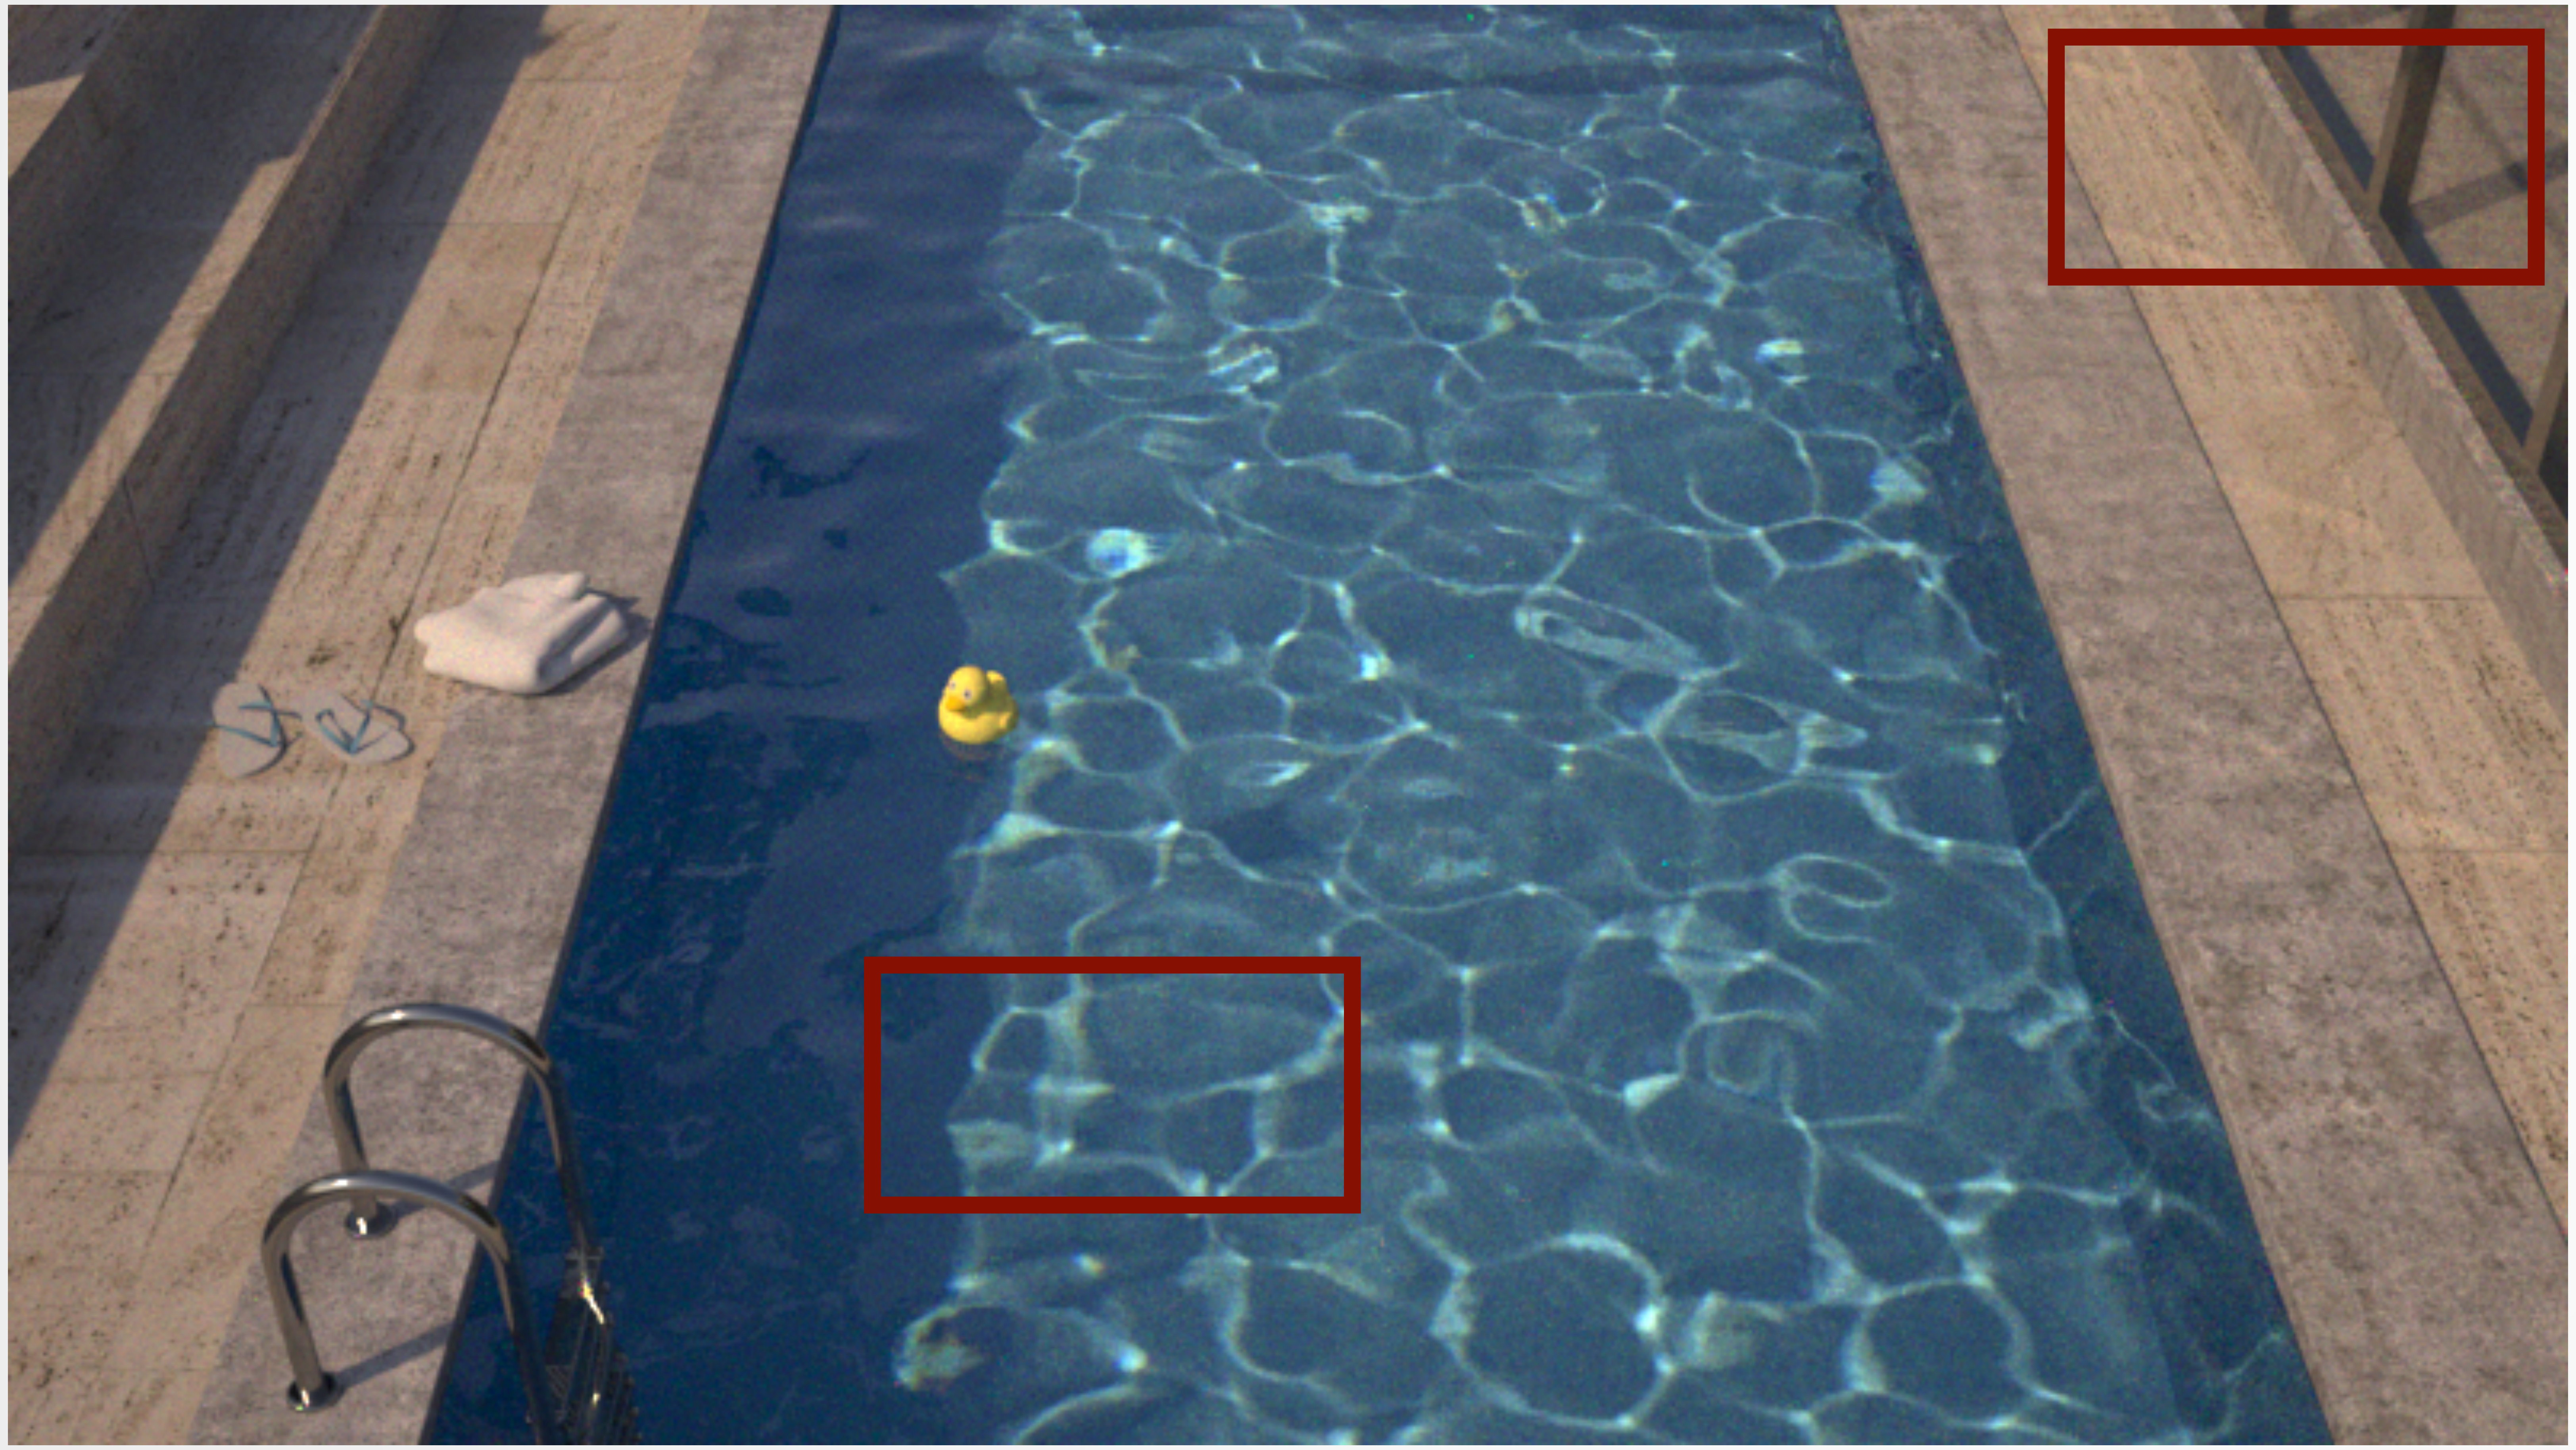

guided path tracing + DBOR

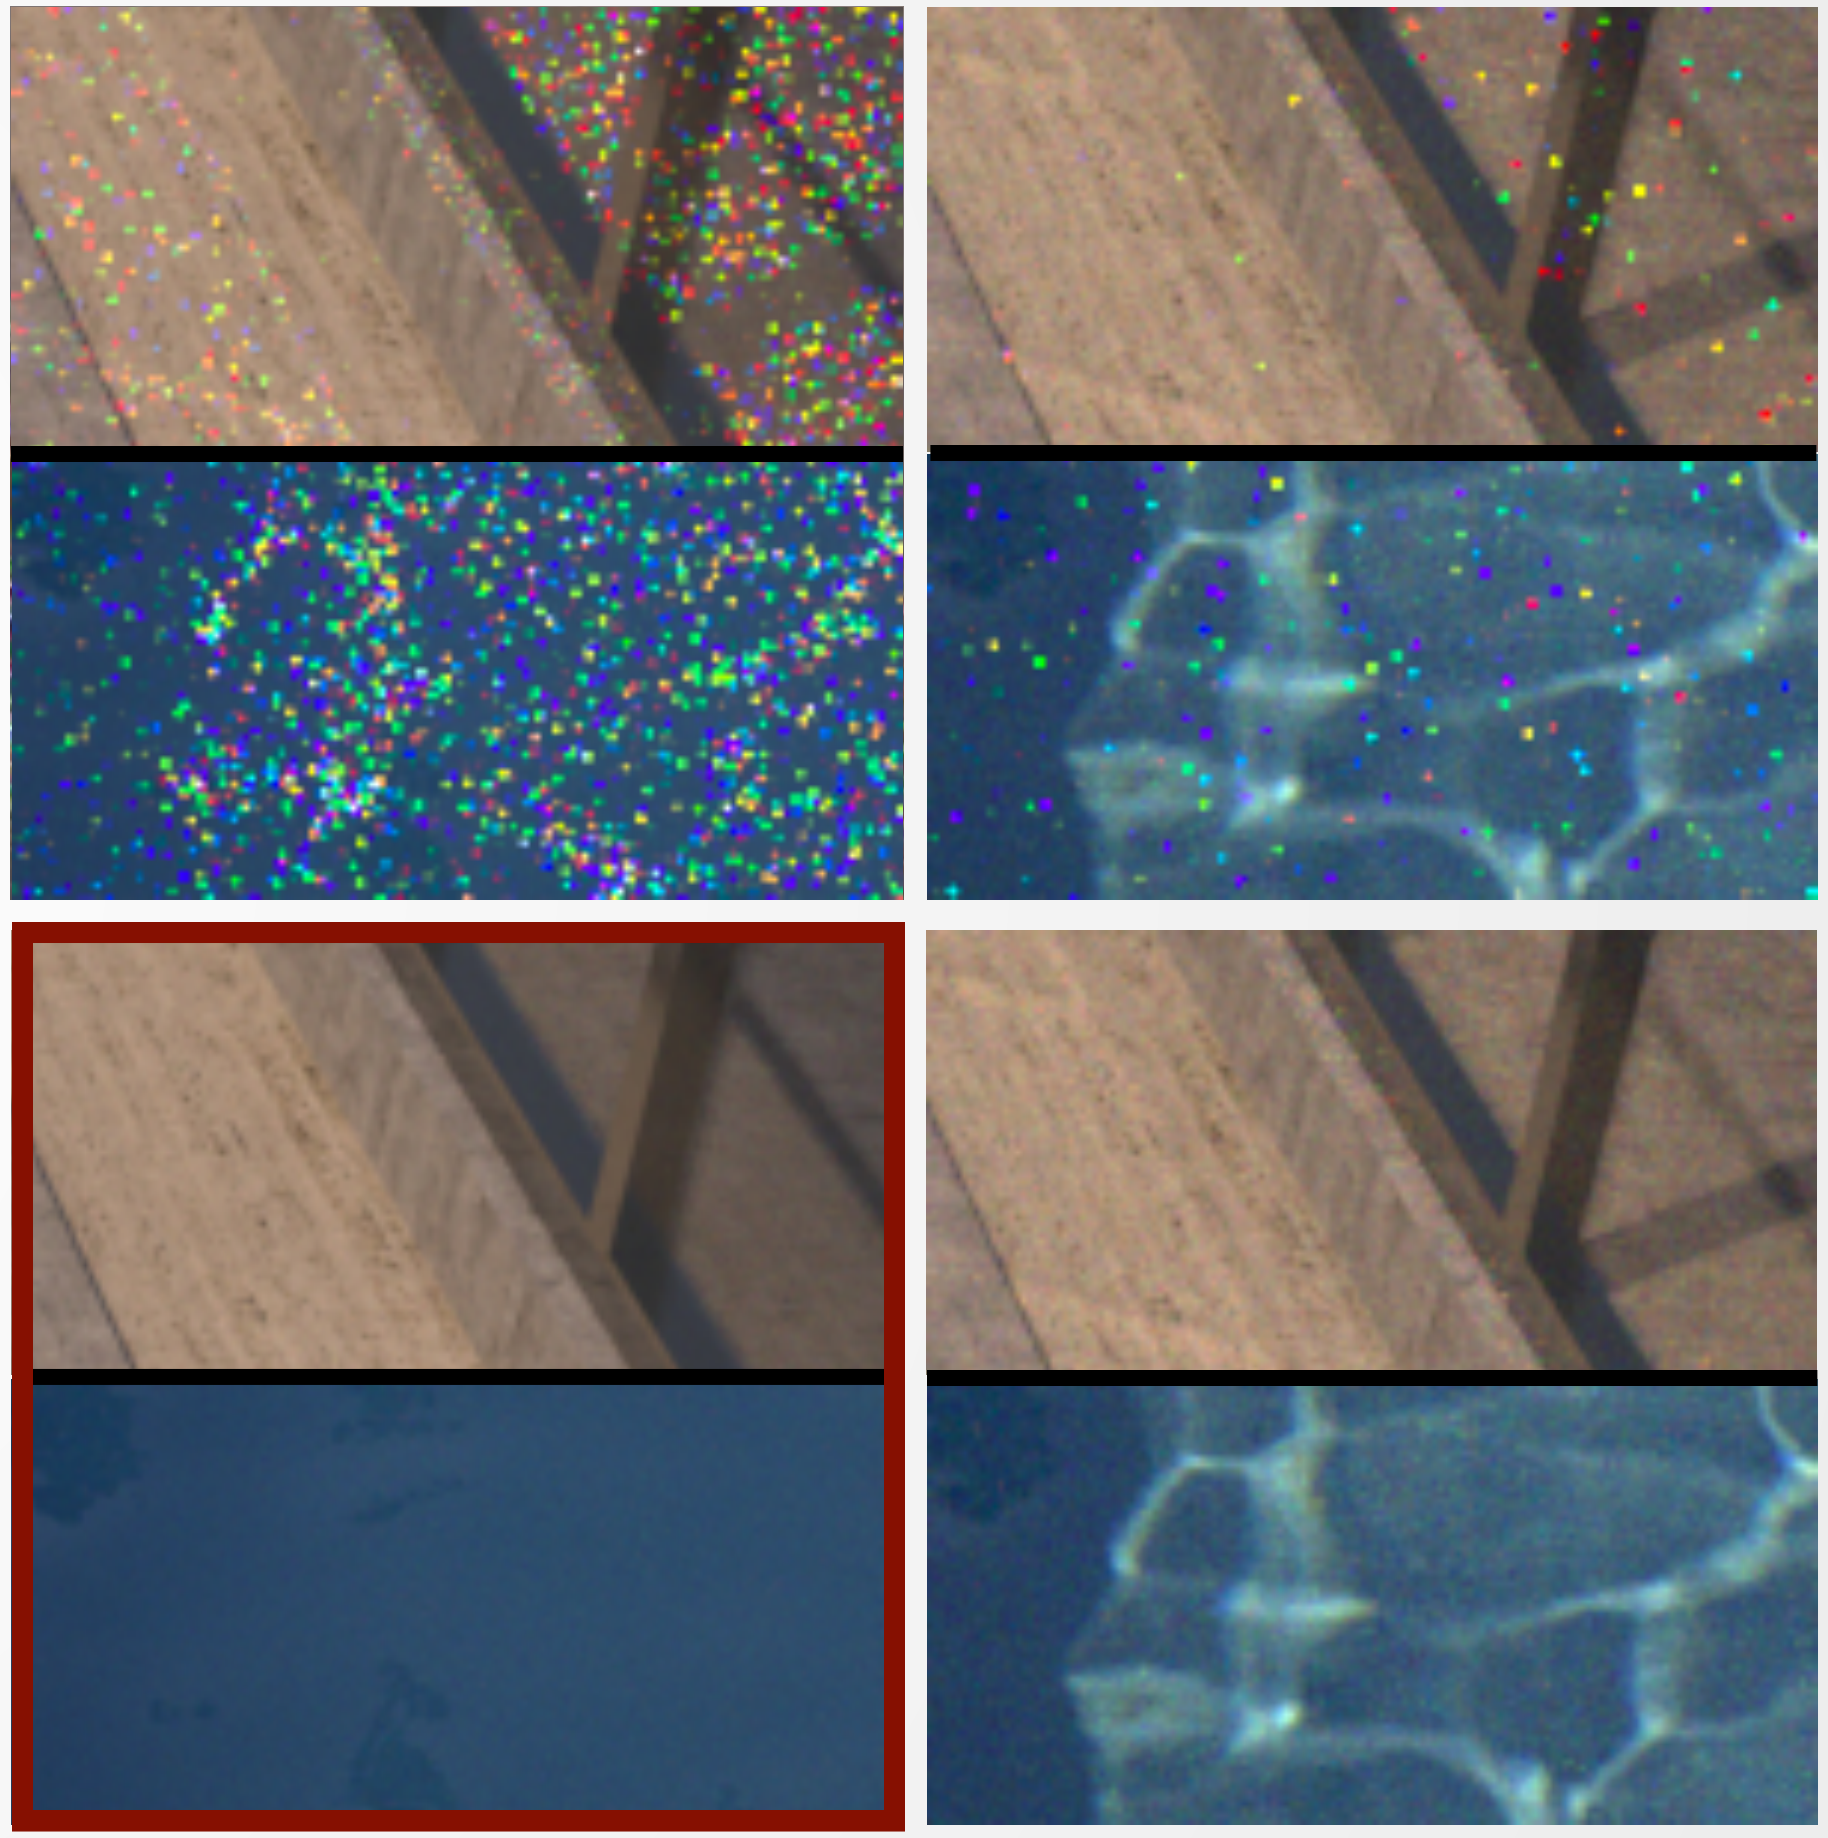

RMSE 0.2550
RMSE 0.3162
RMSE 0.1238
RMSE 0.0499

Dragon - 10h

SIGGRAPH ASIA 2018 TOKYO

guided path tracing + DBOR

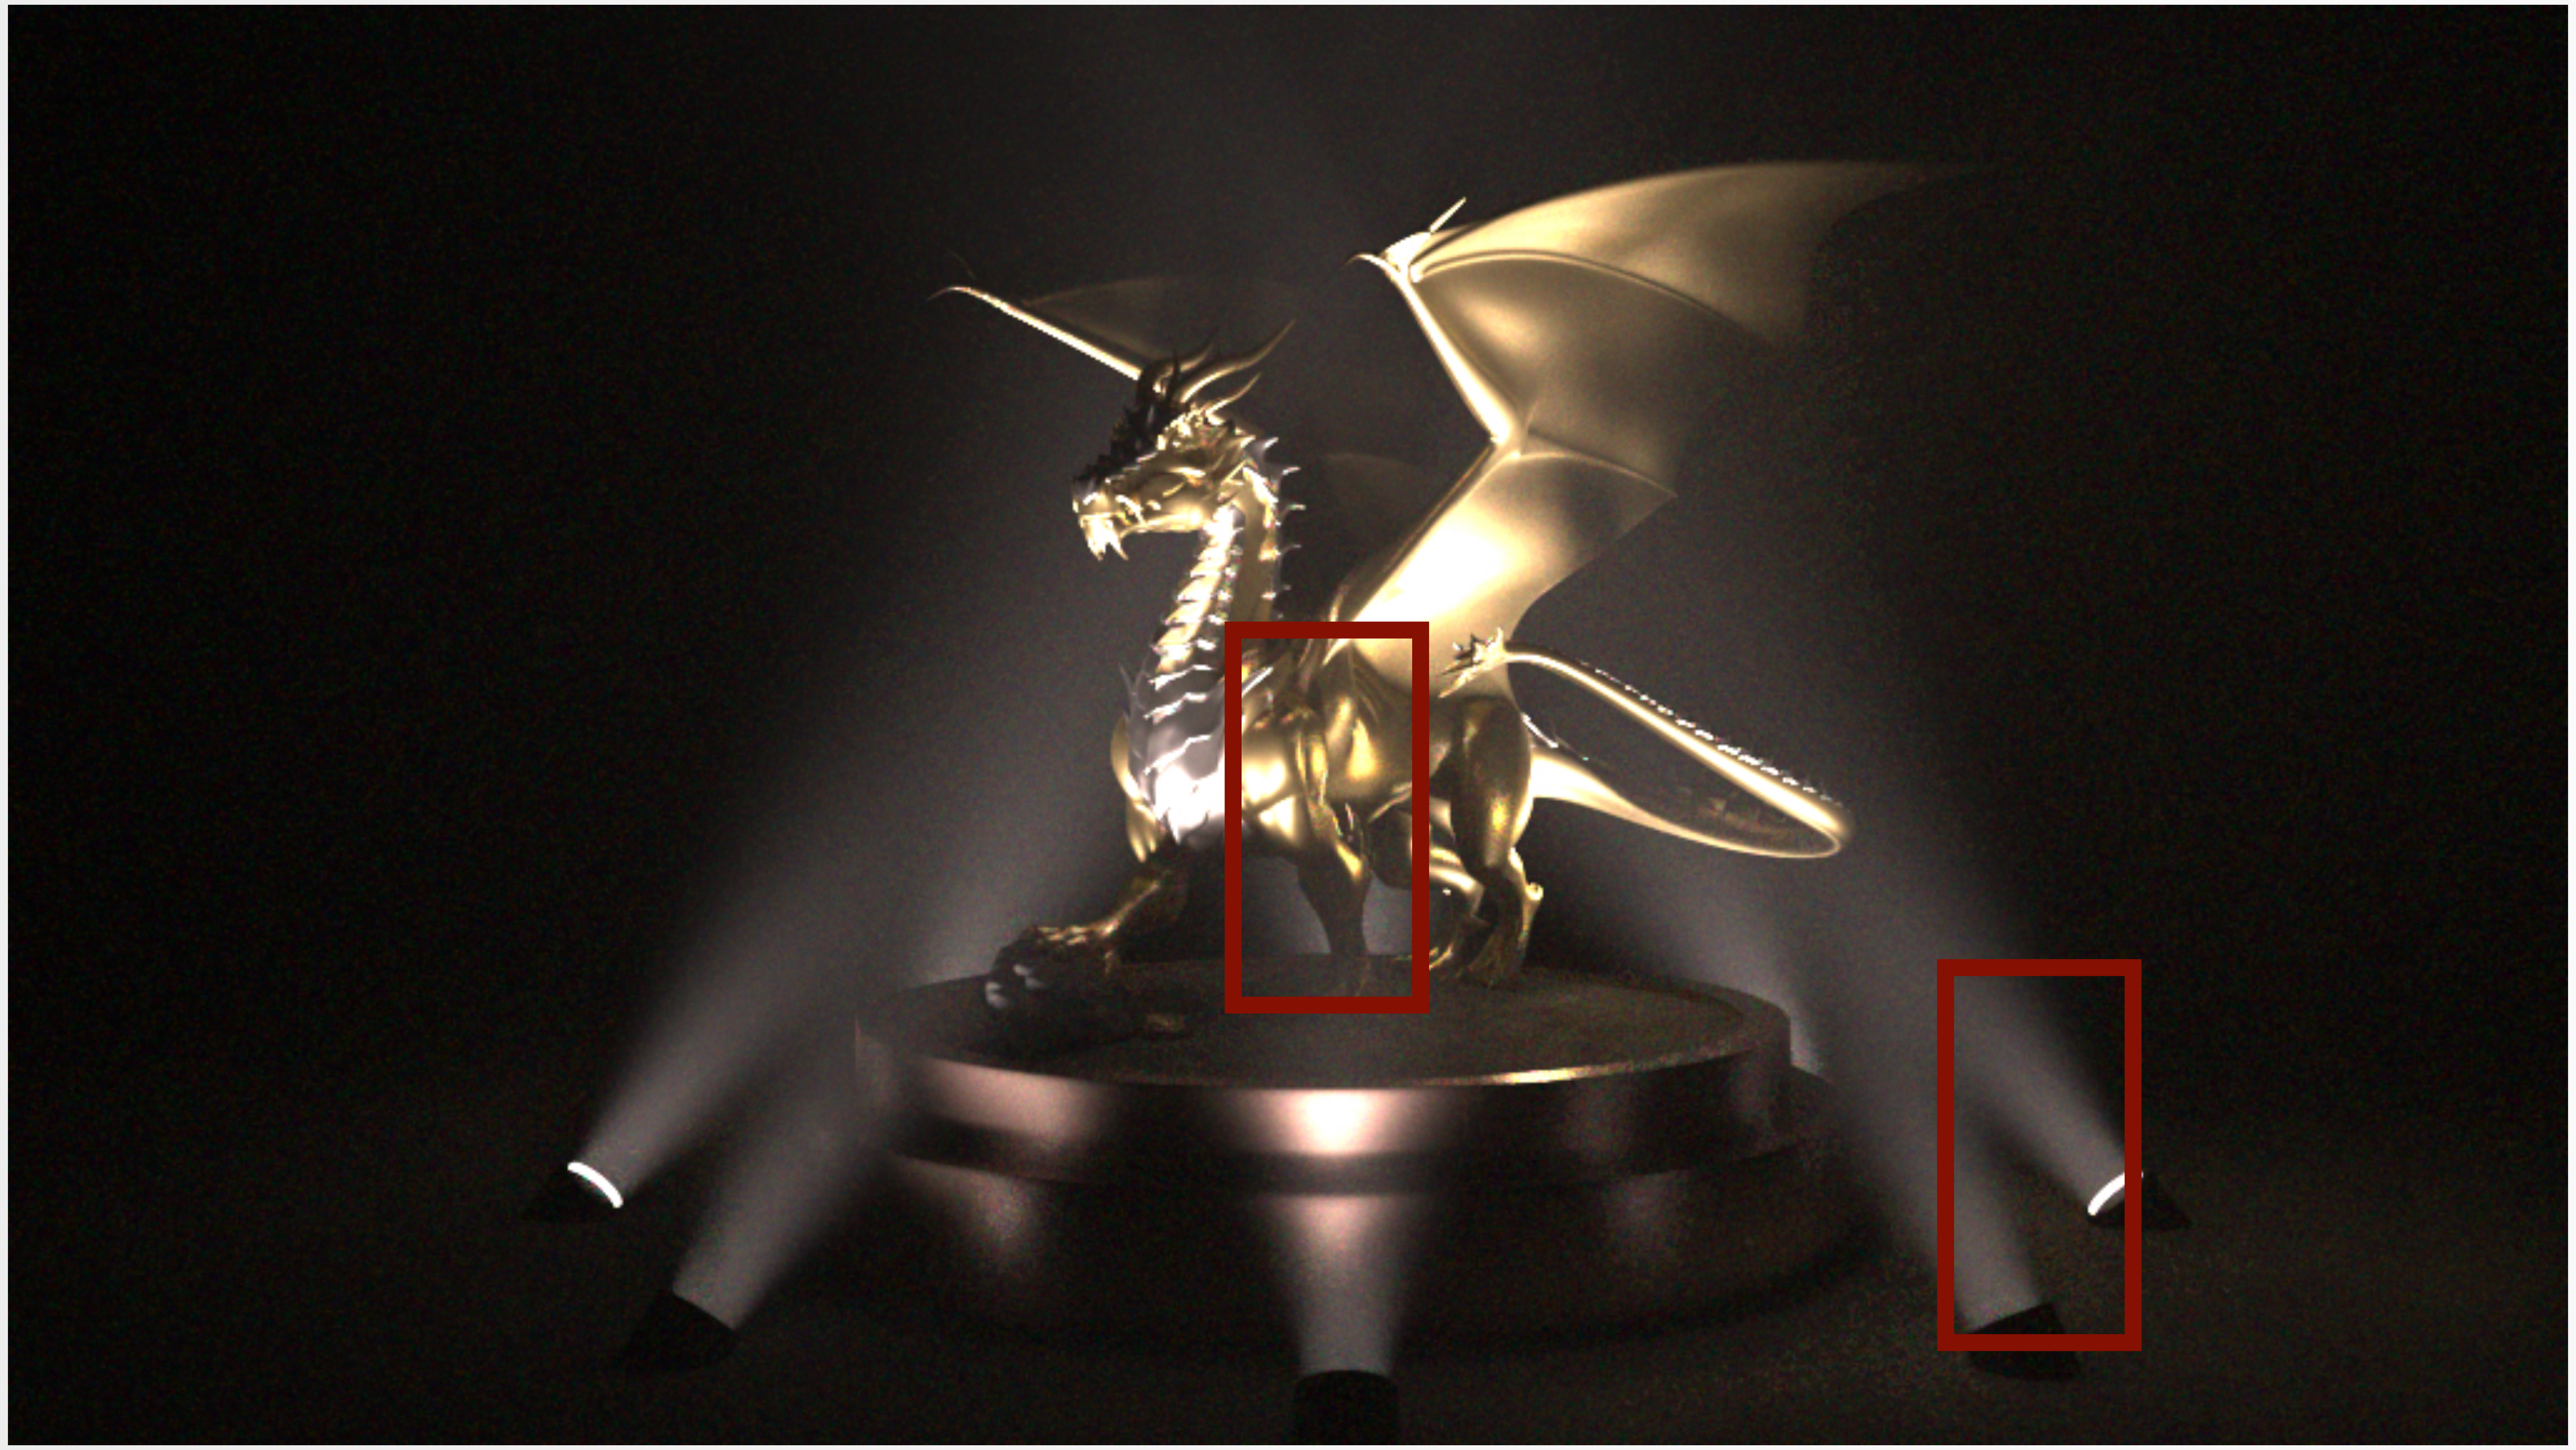

RMSE 0.5012 **RMSE 0.2121**

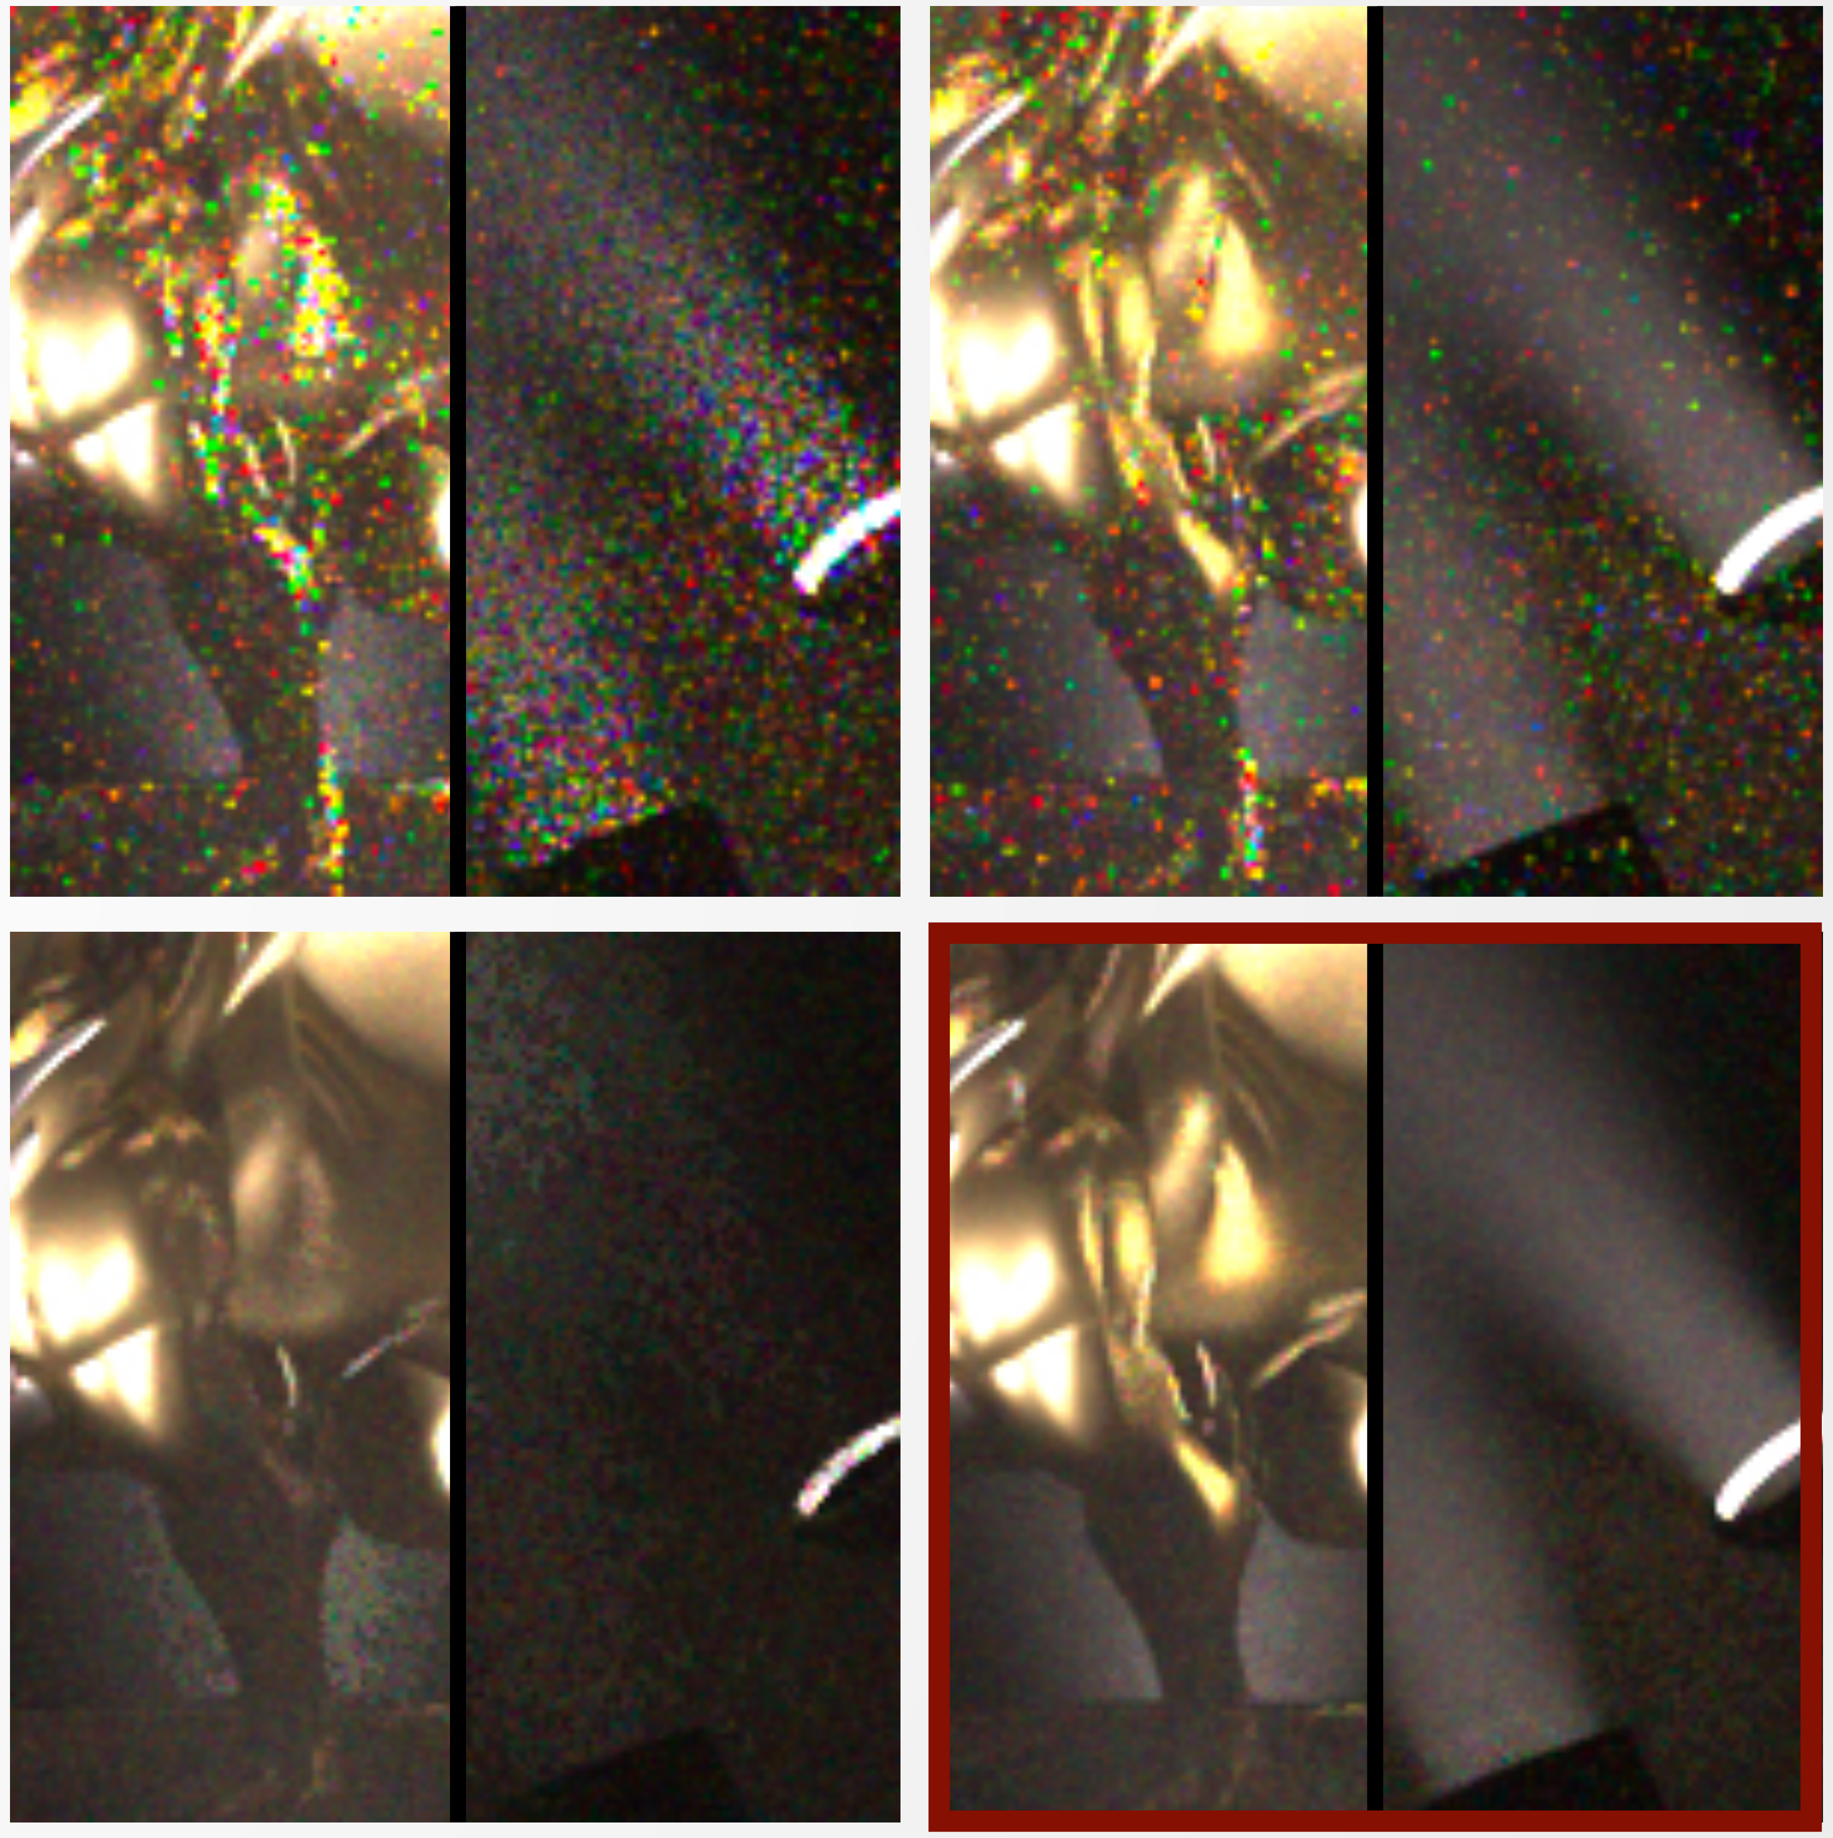

RMSE 0.6143 **RMSE 0.2198**

Dining Room - 30min

reference

ours

removed by DBOR

path tracing

guided path tracing

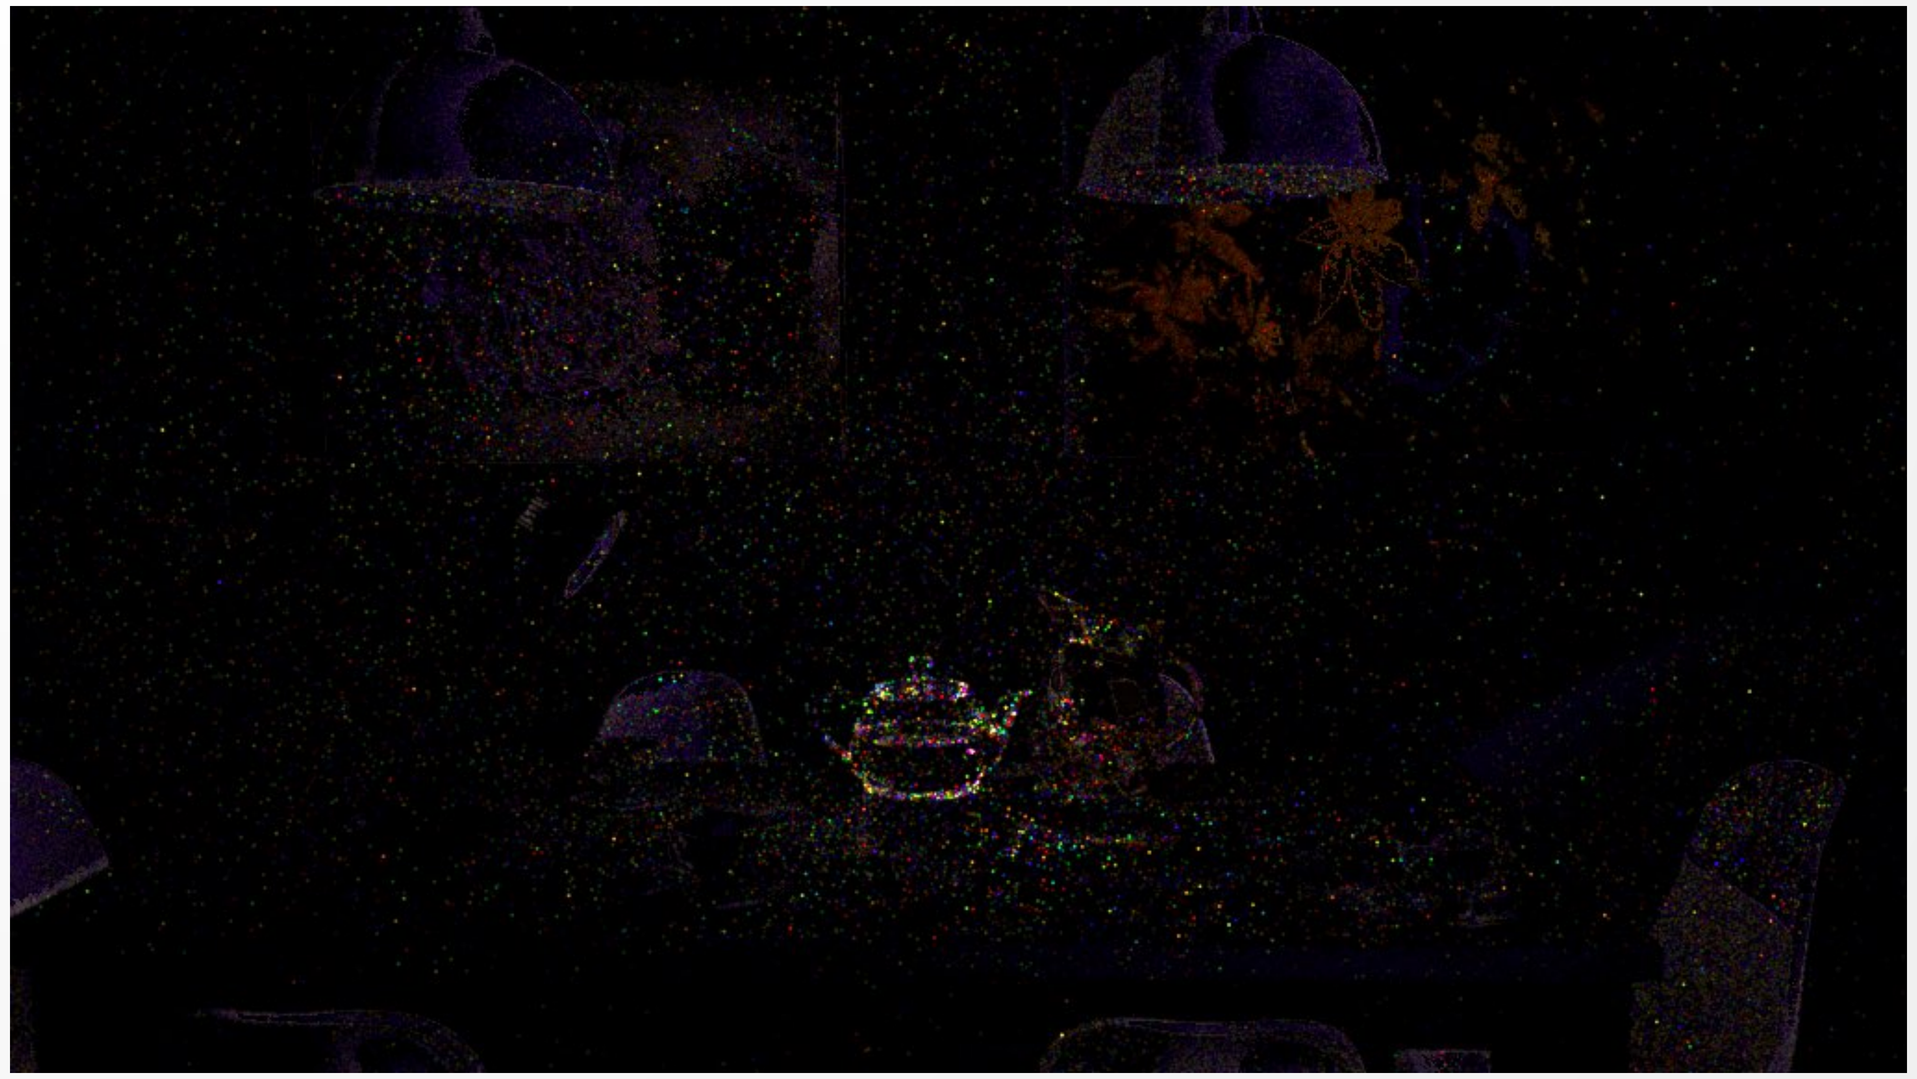

Pool - 30min

reference

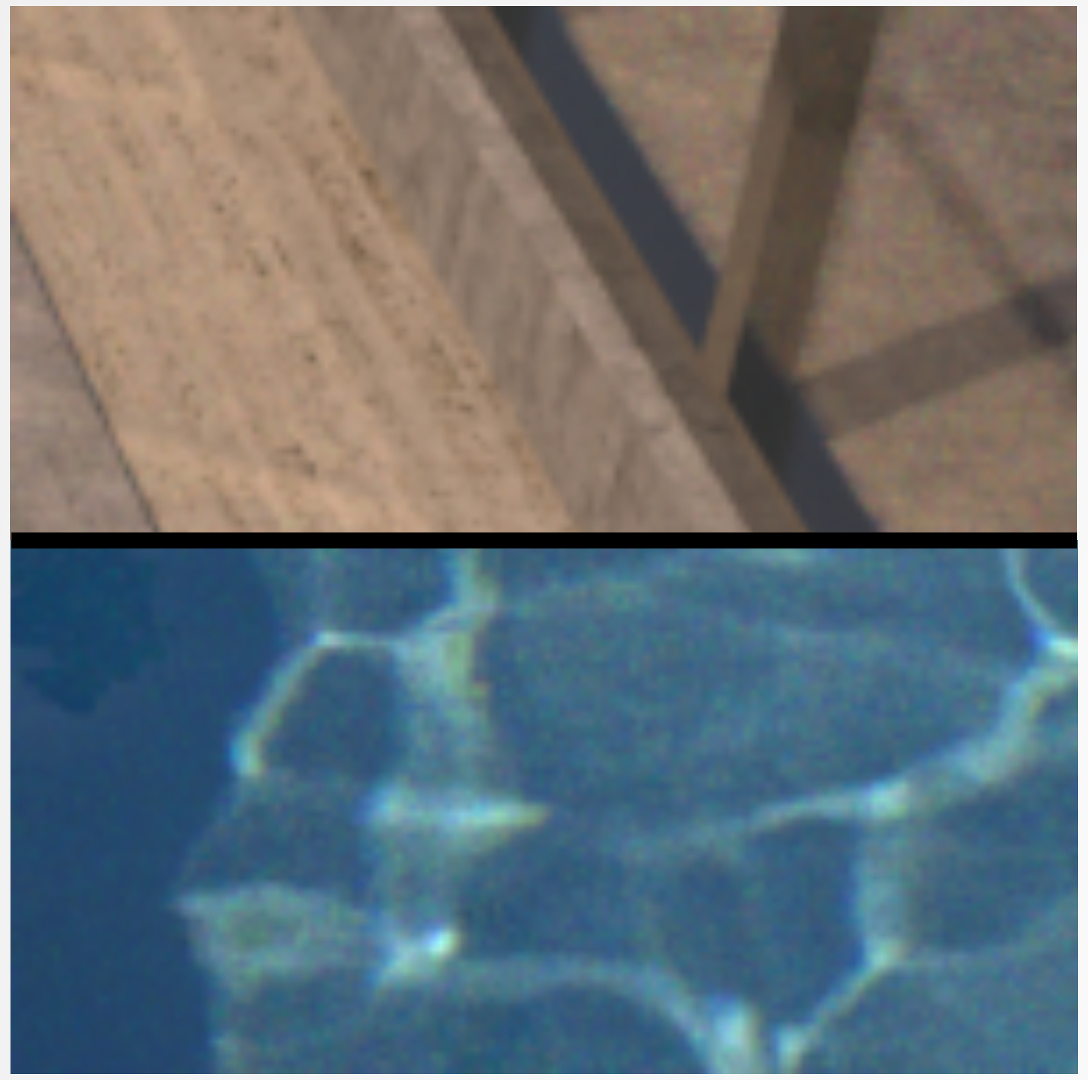

ours

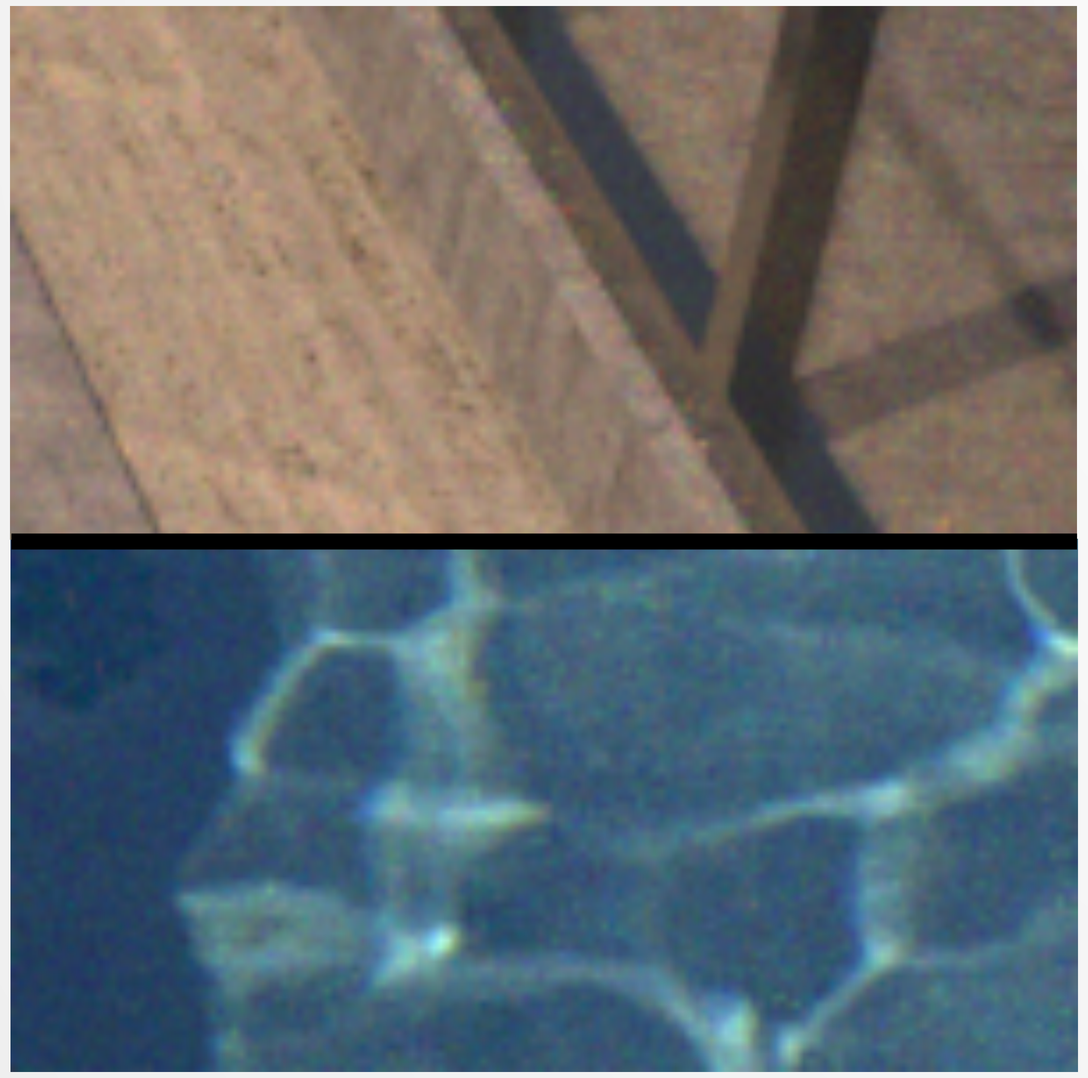

removed by DBOR

path tracing

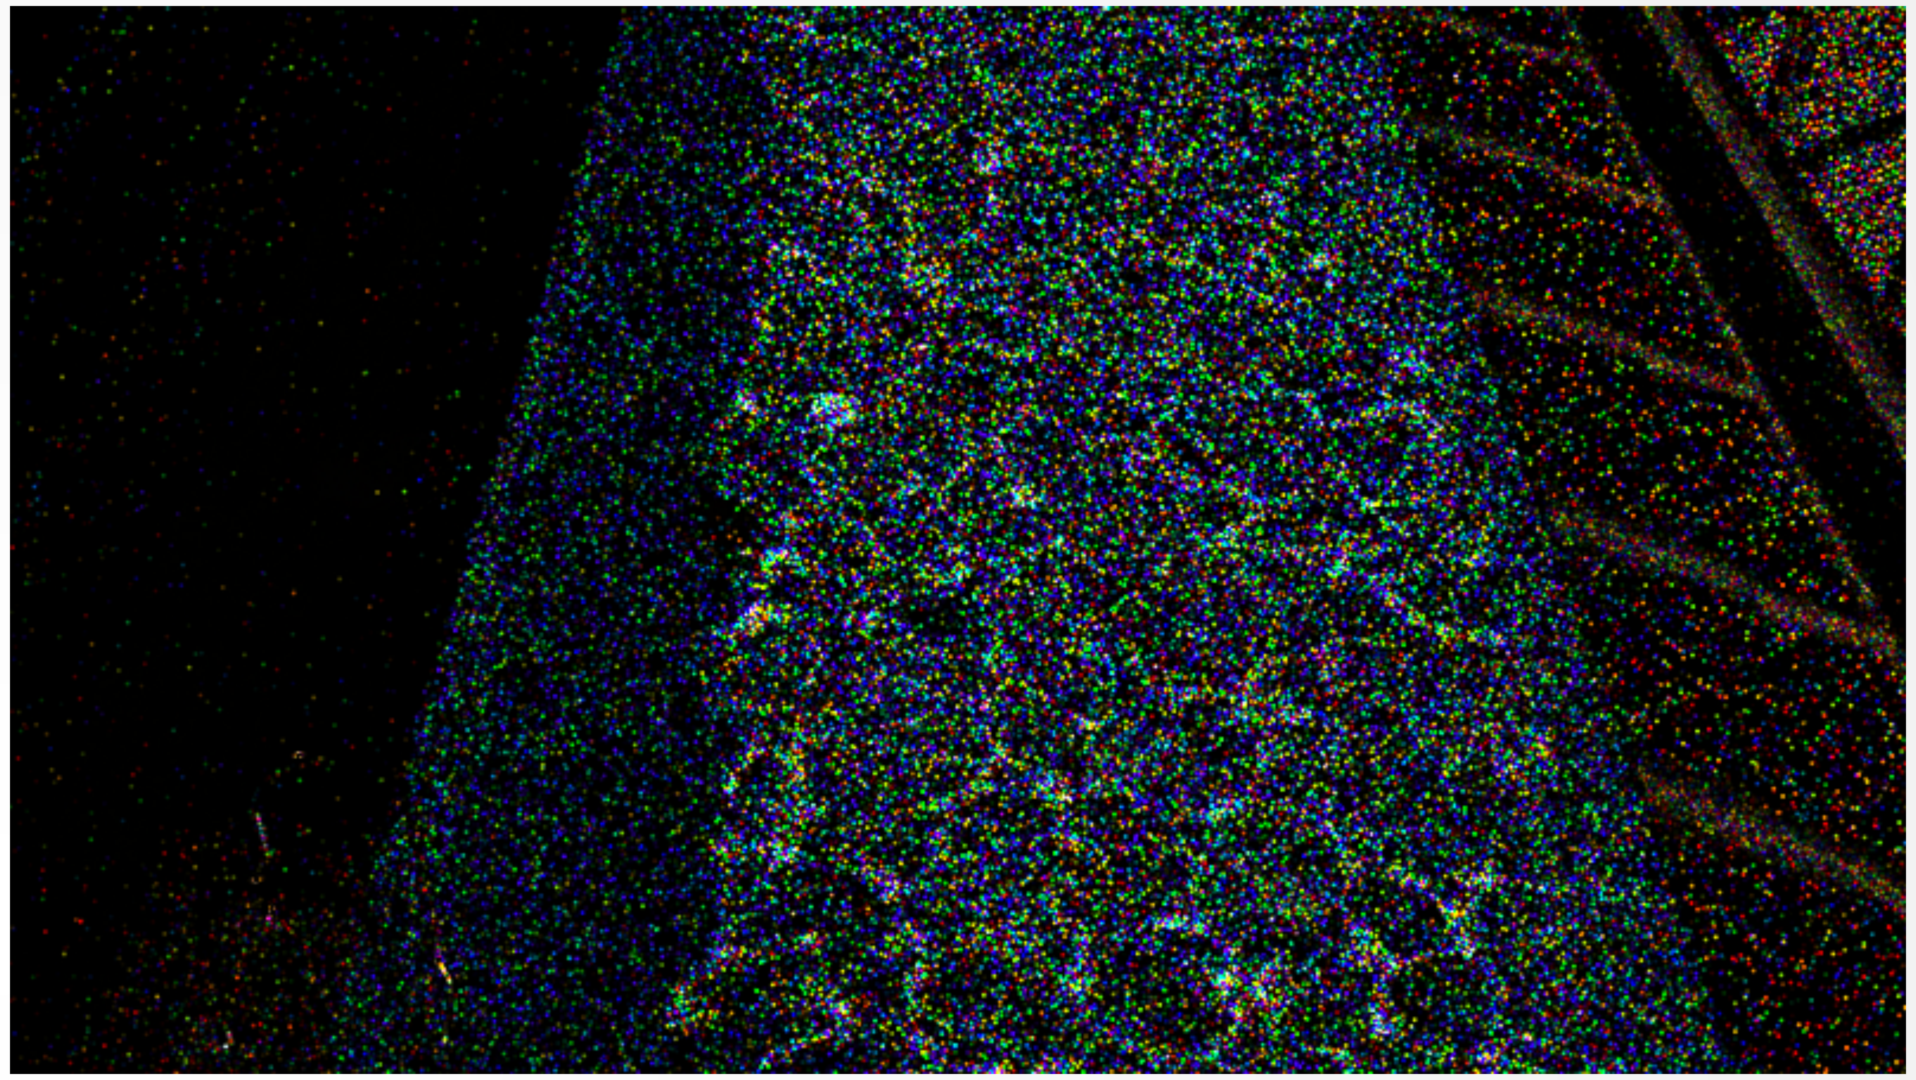

guided path tracing

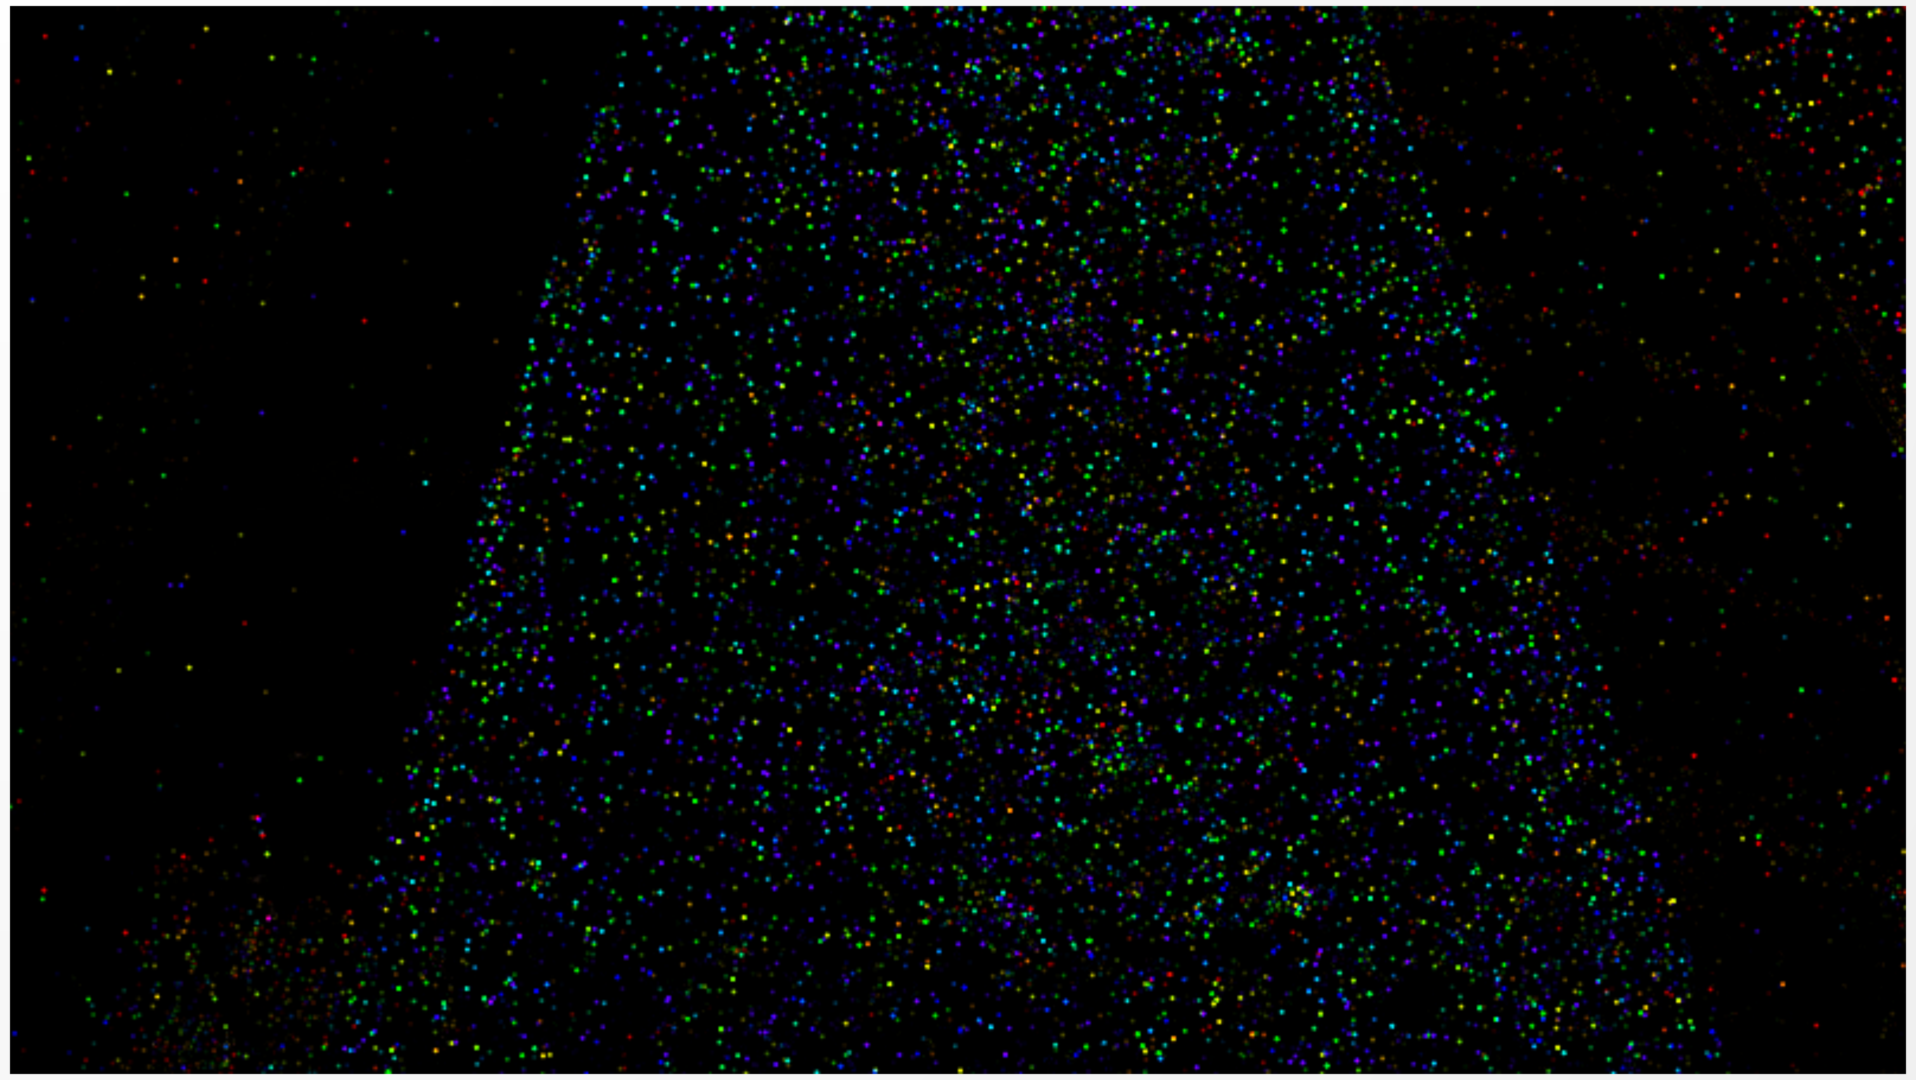

Dragon - 10h

reference

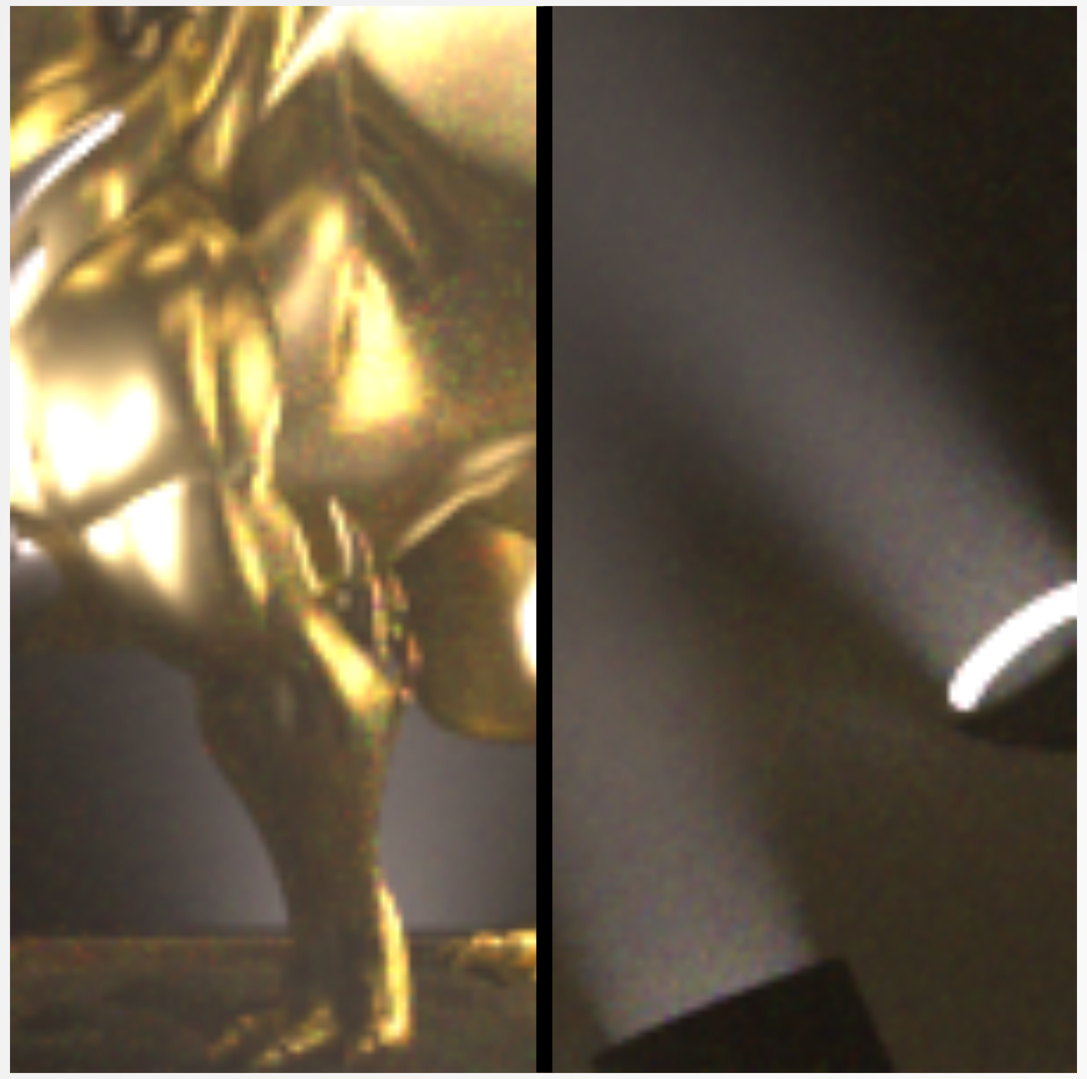

ours

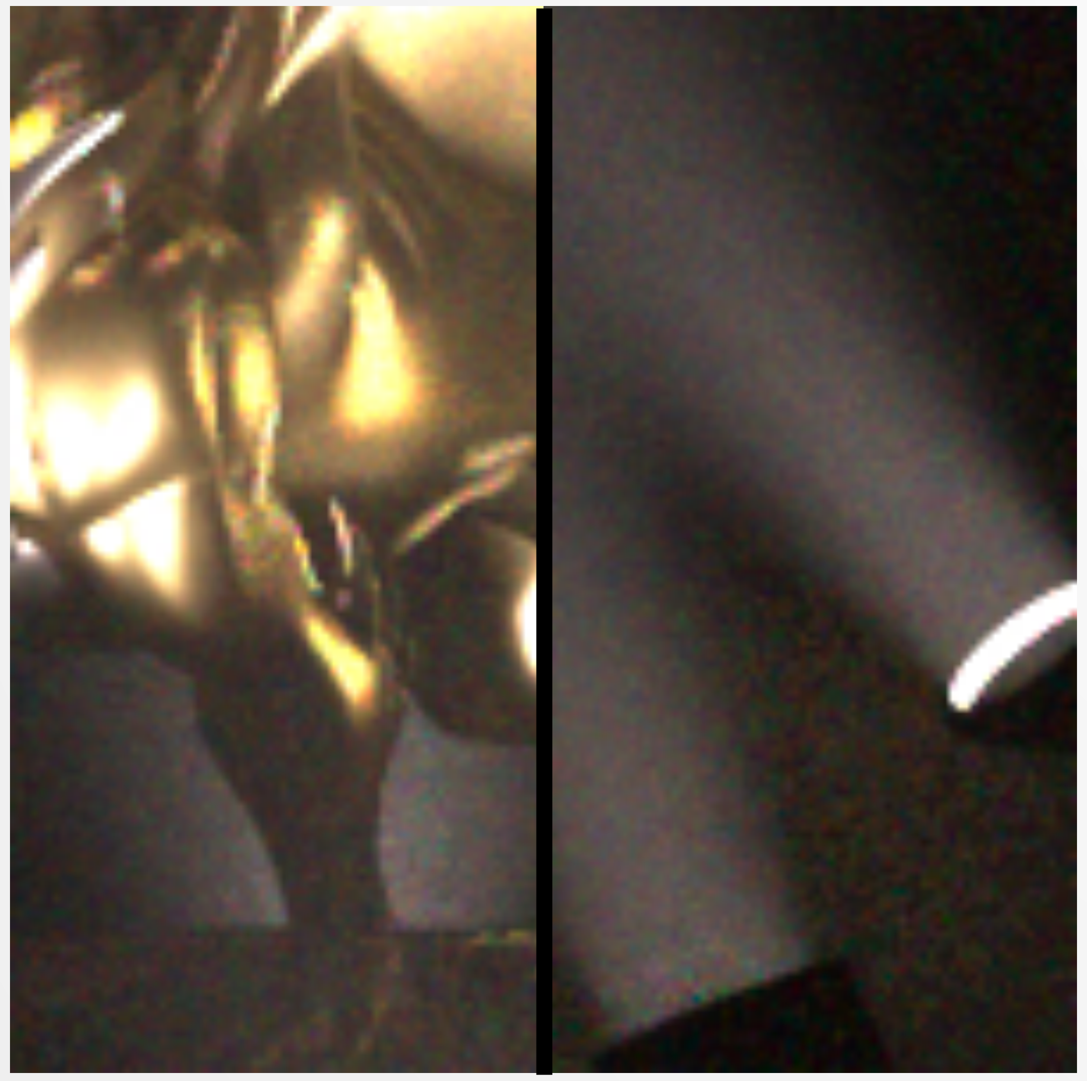

removed by DBOR

path tracing

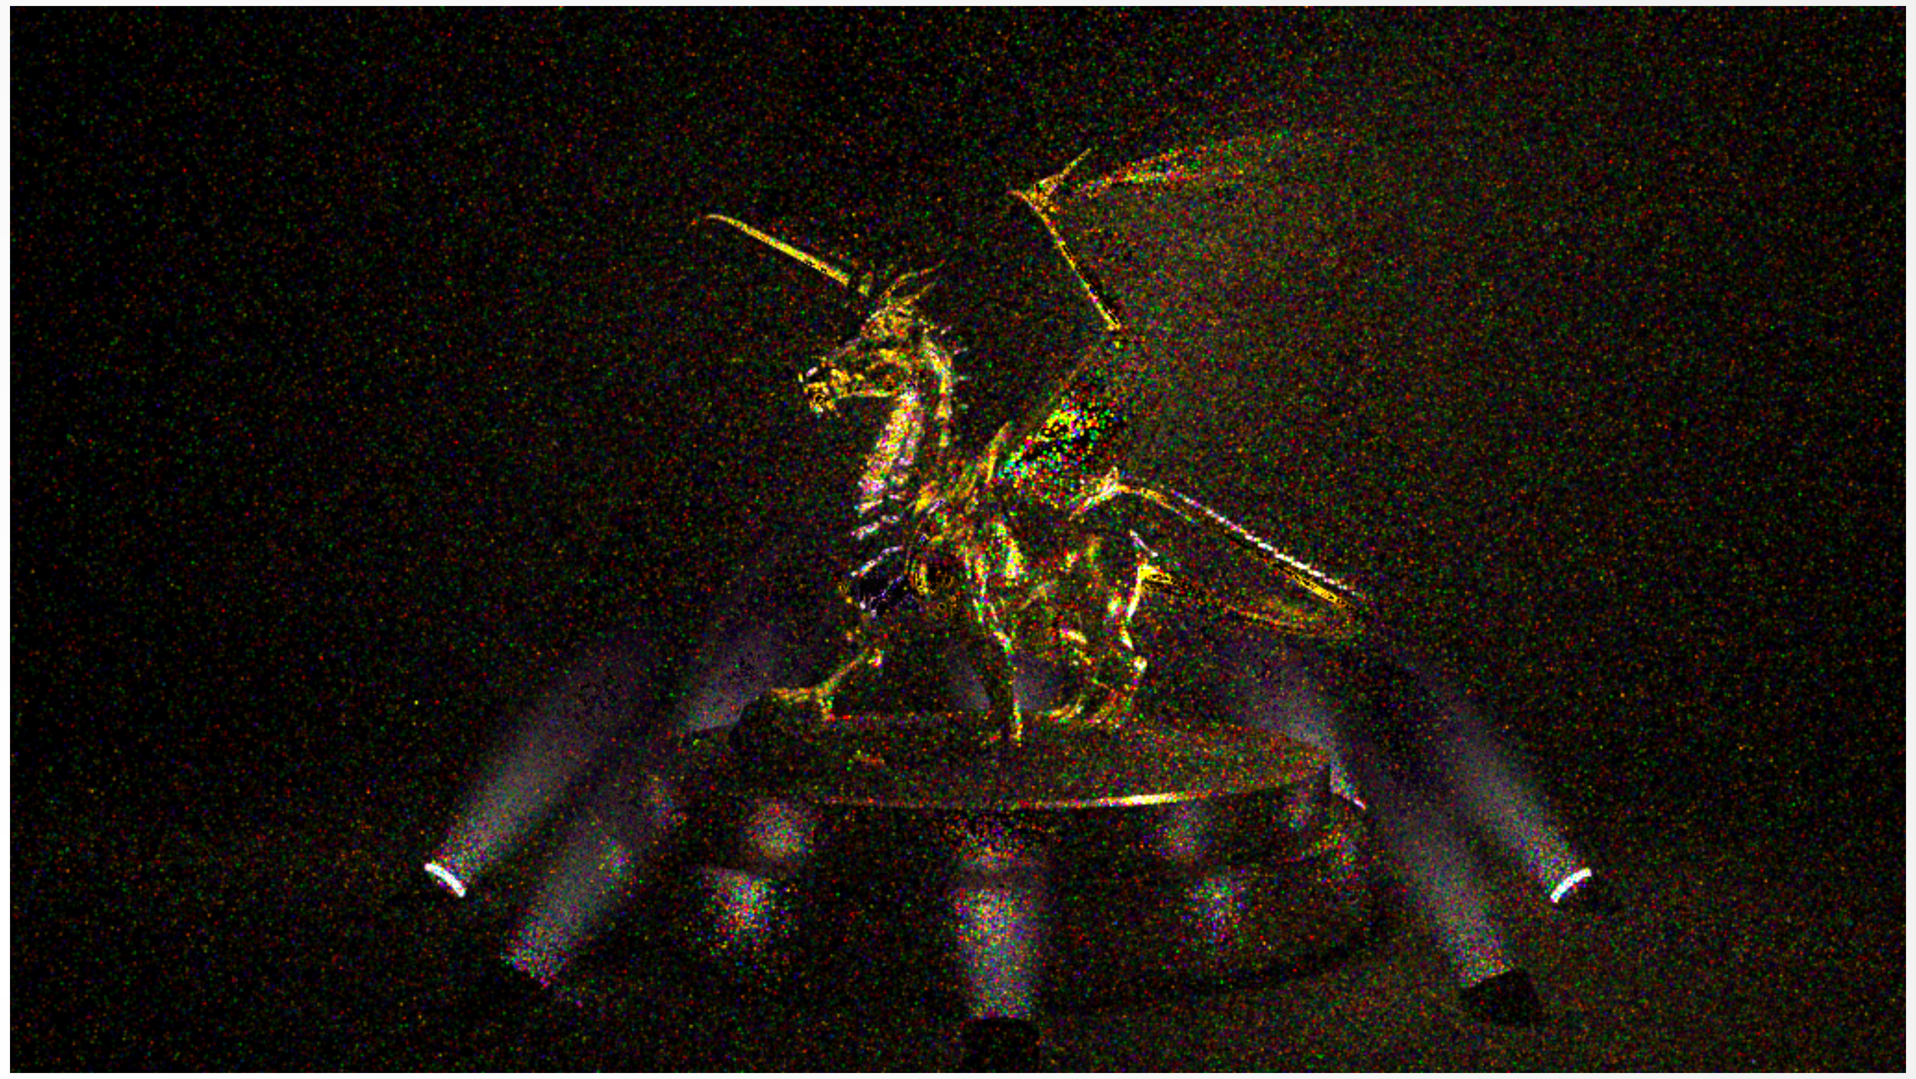

guided path tracing

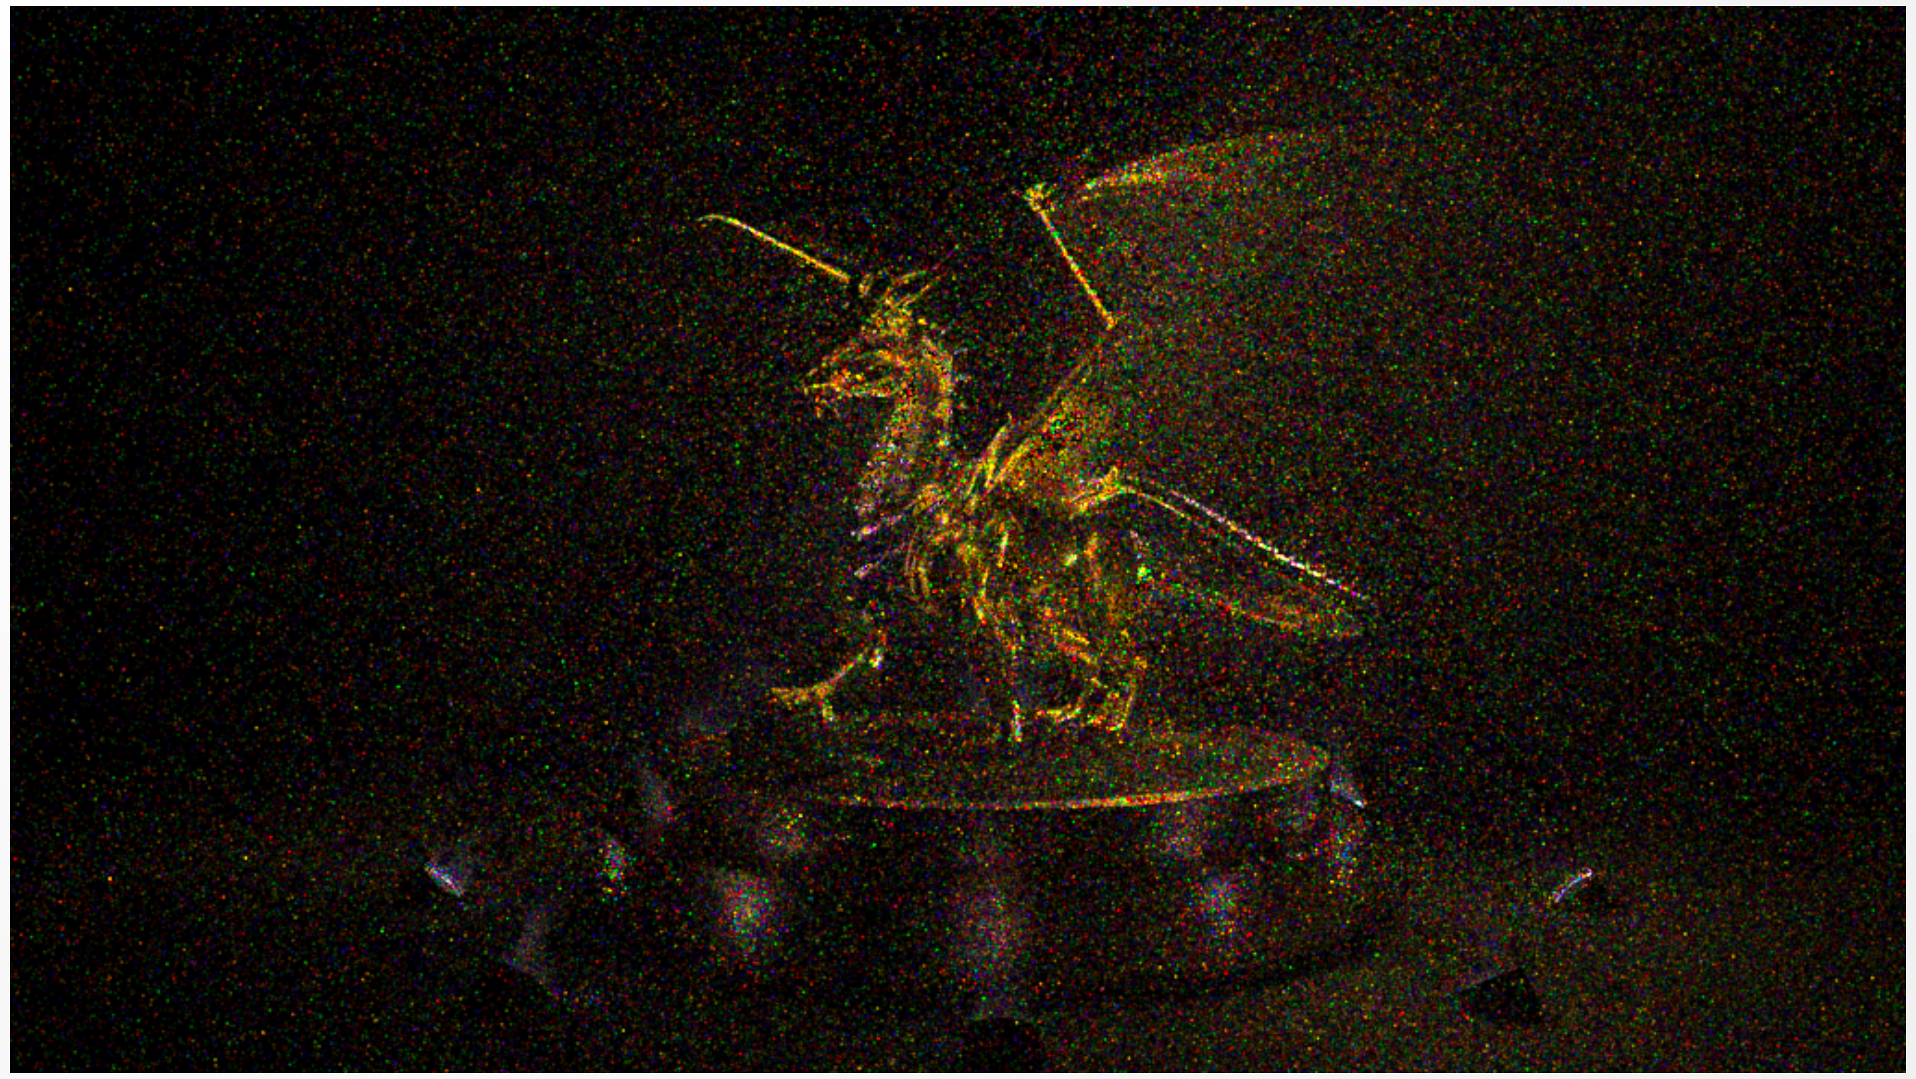

Convergence

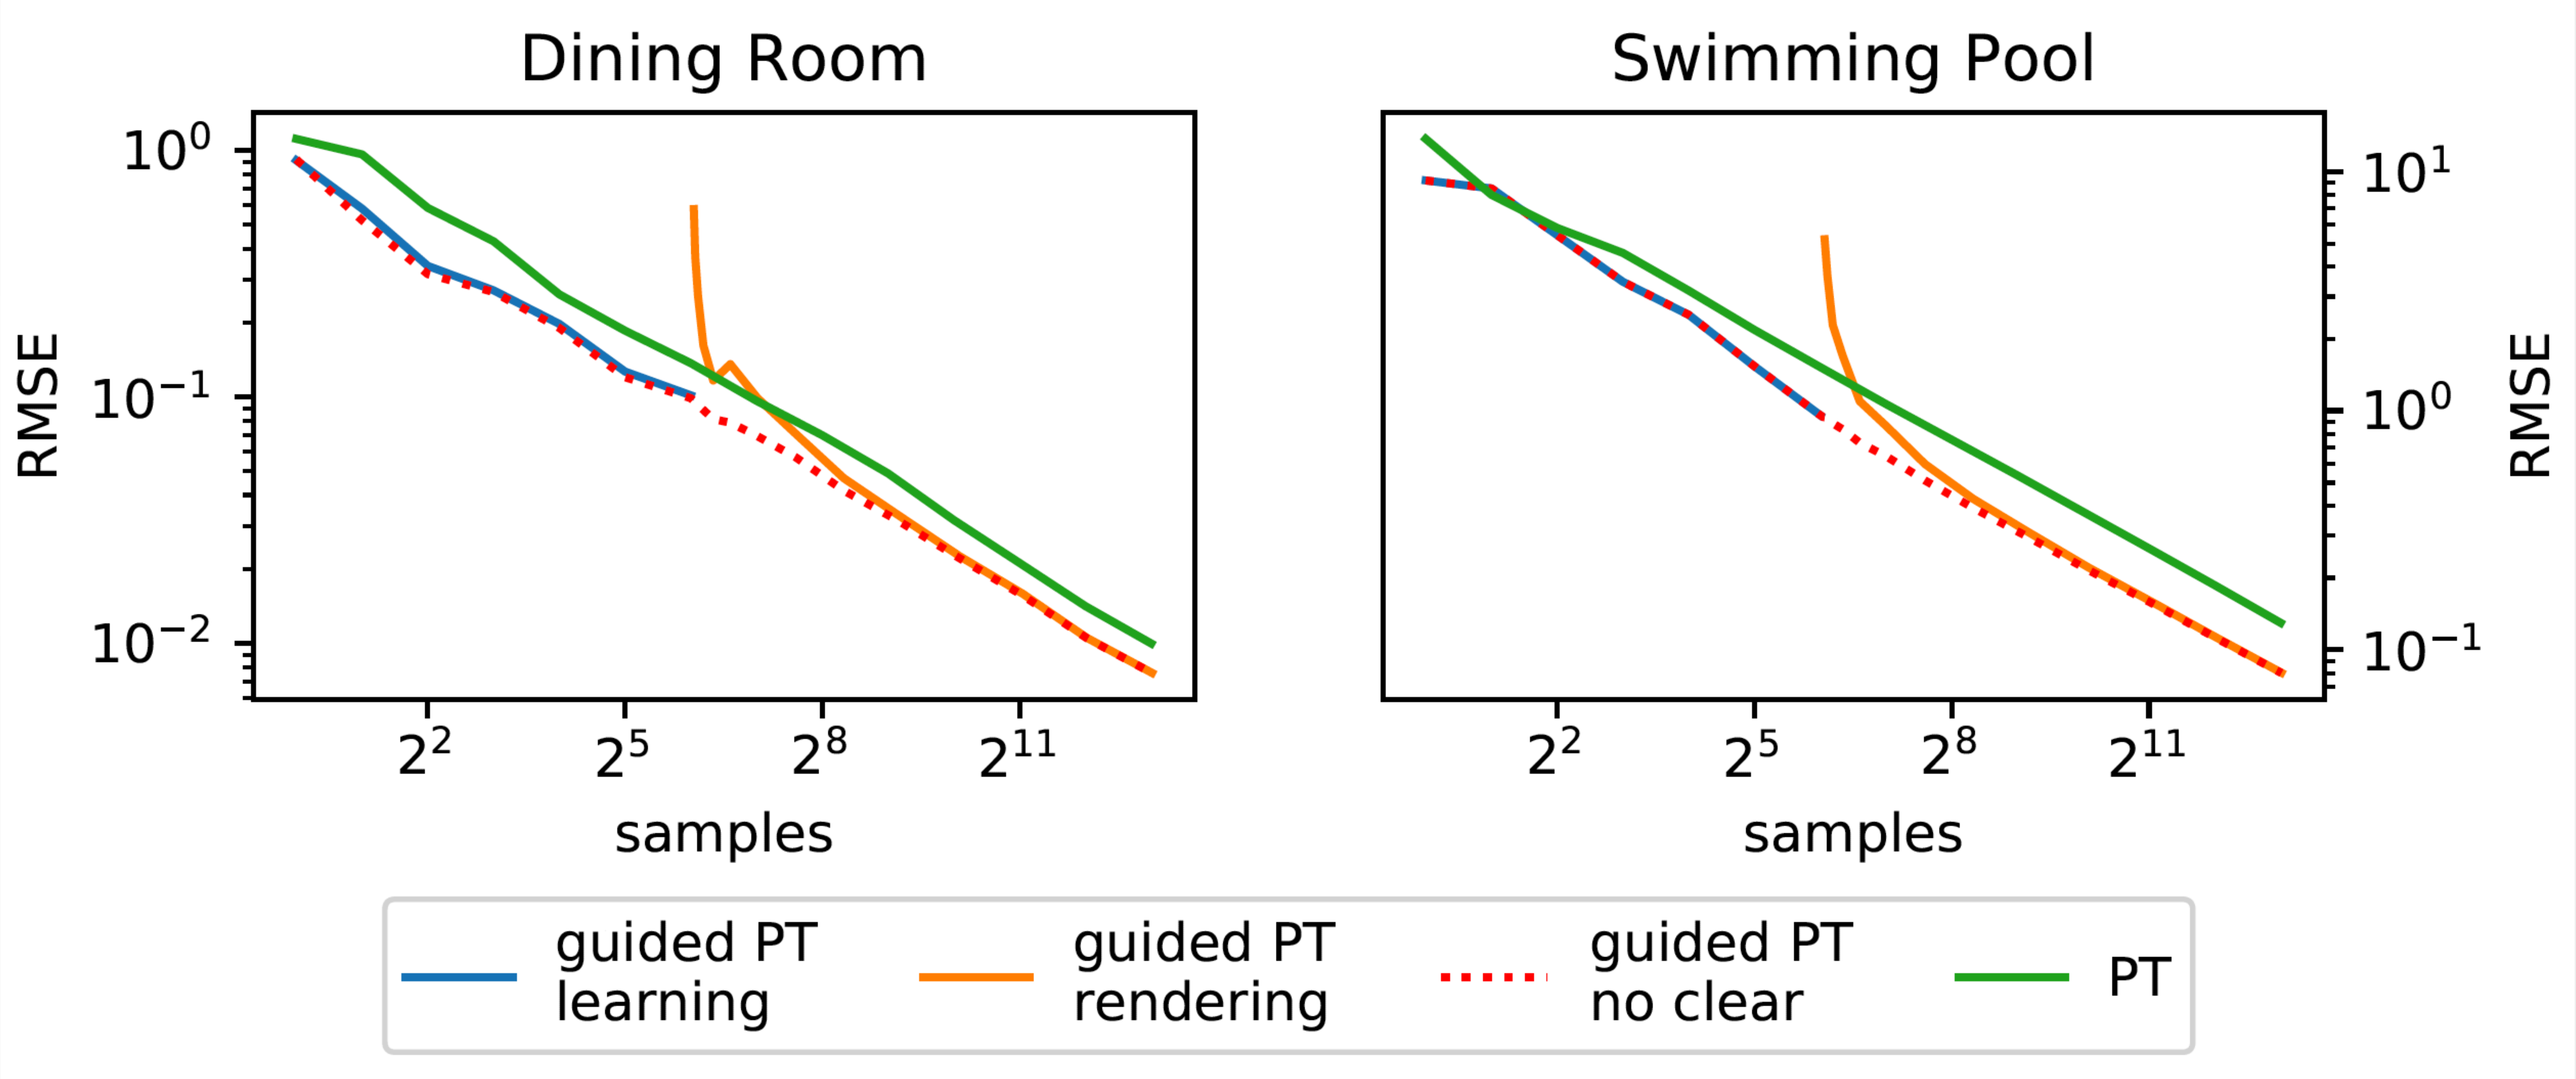

Difference

Reference

Guided path tracing + DBOR

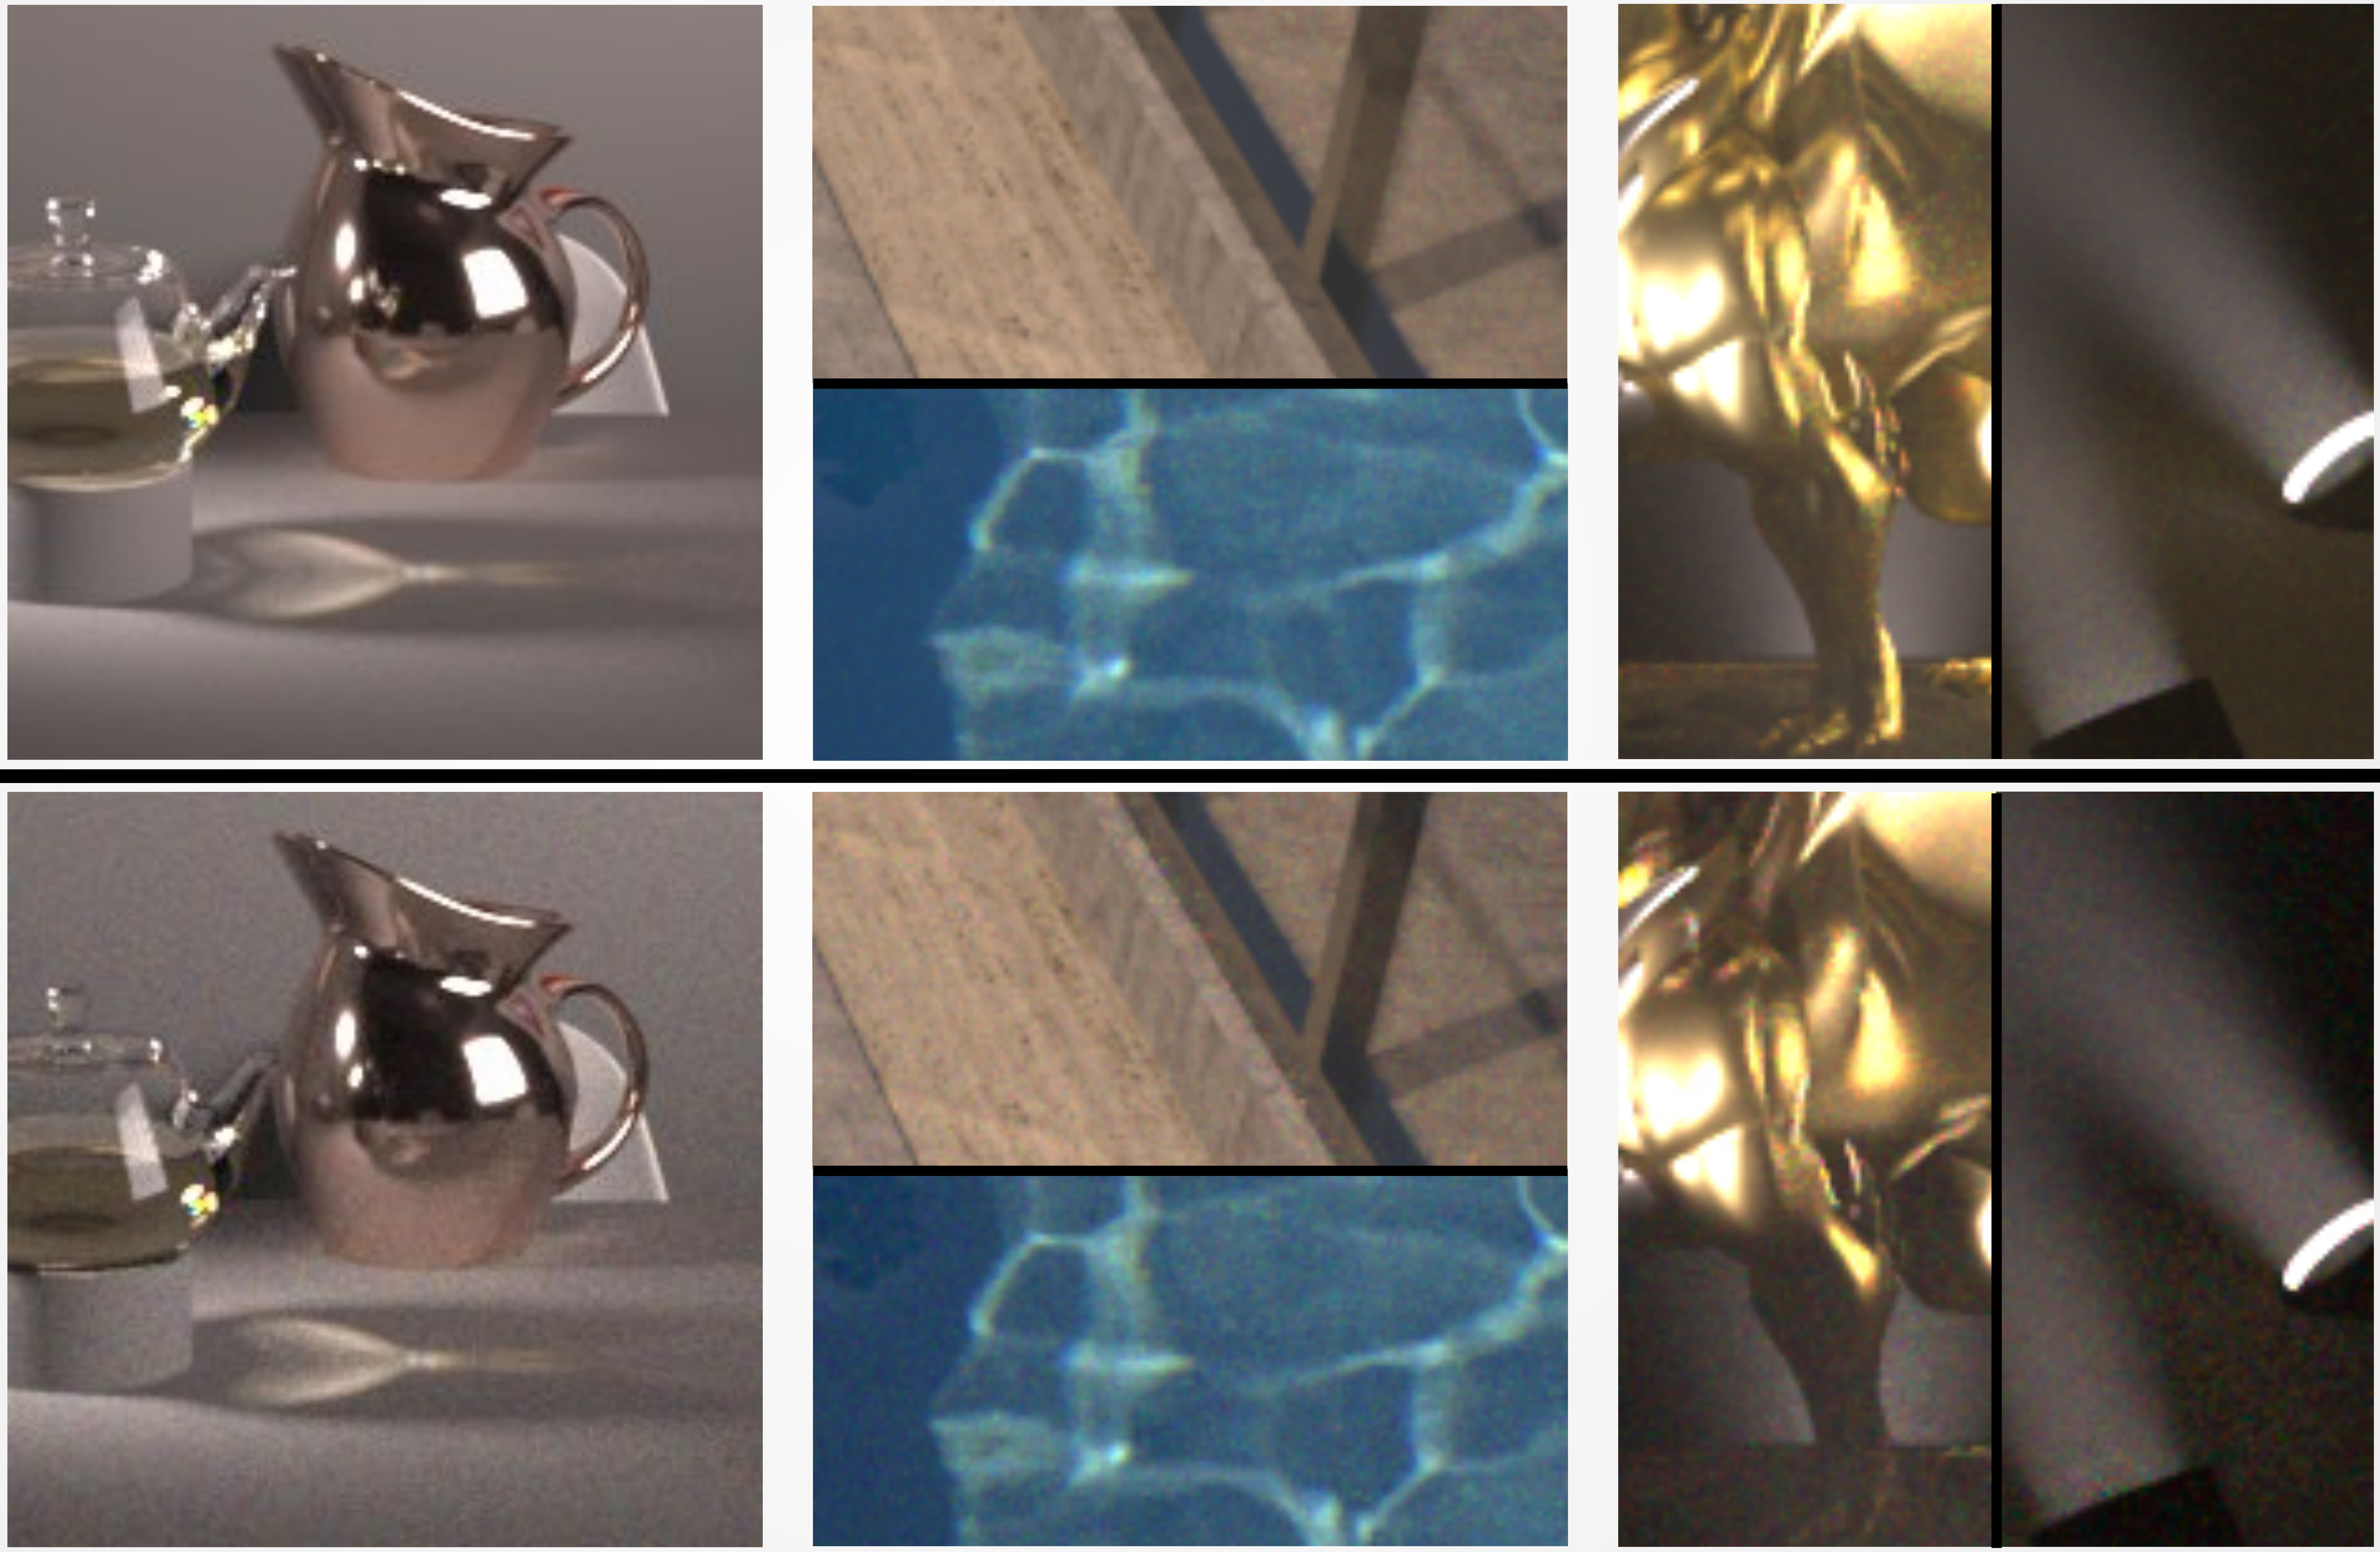

KIT Karlsruhe Institute of Technology

Difference

Scene		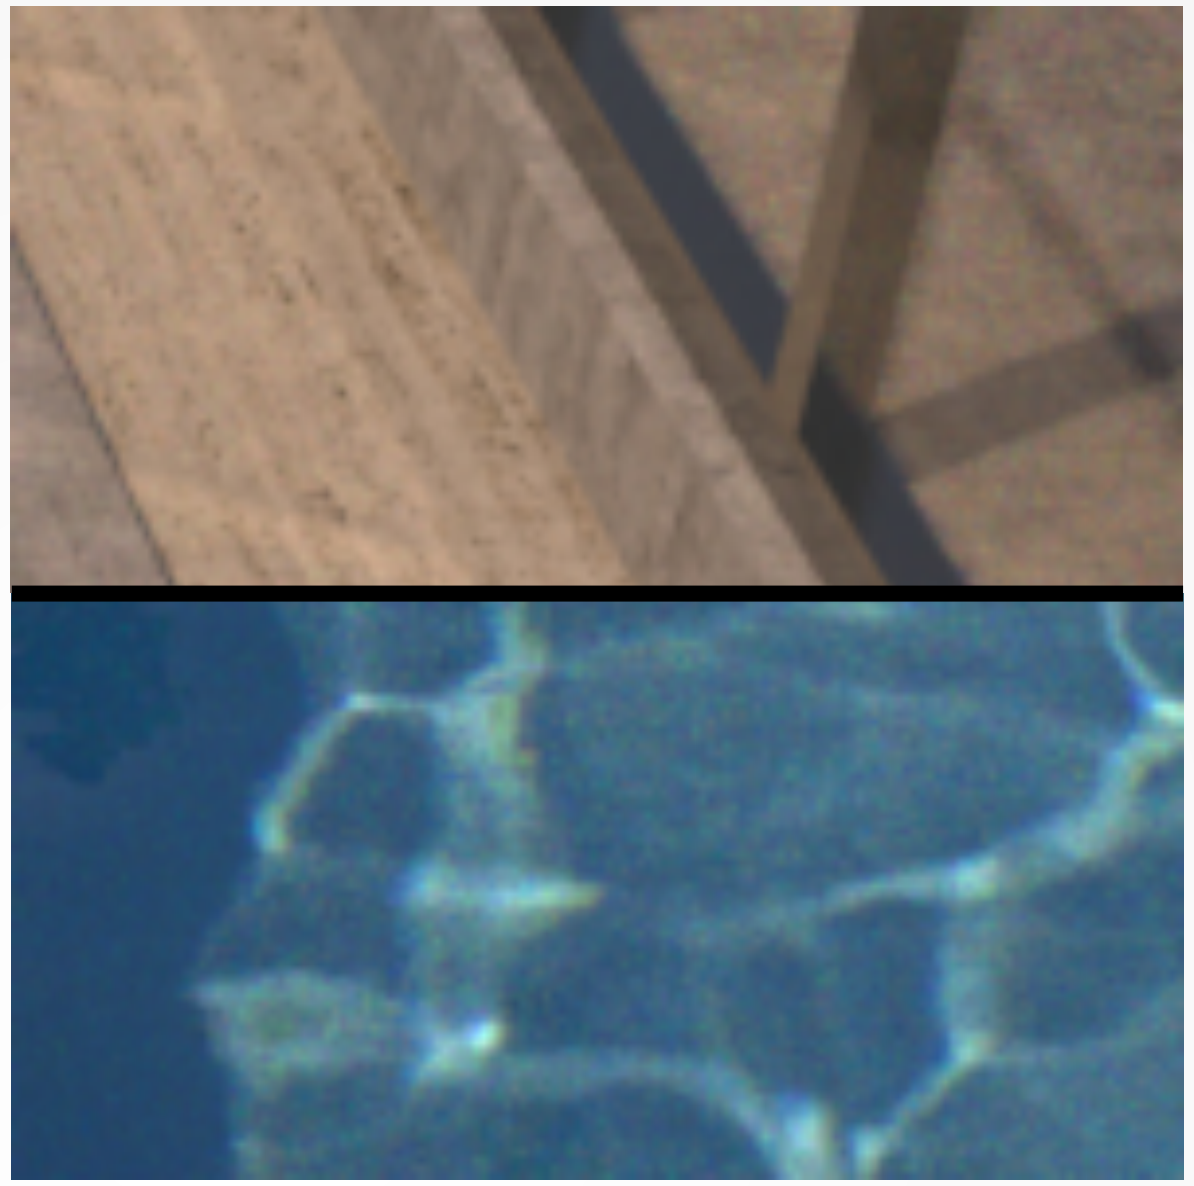	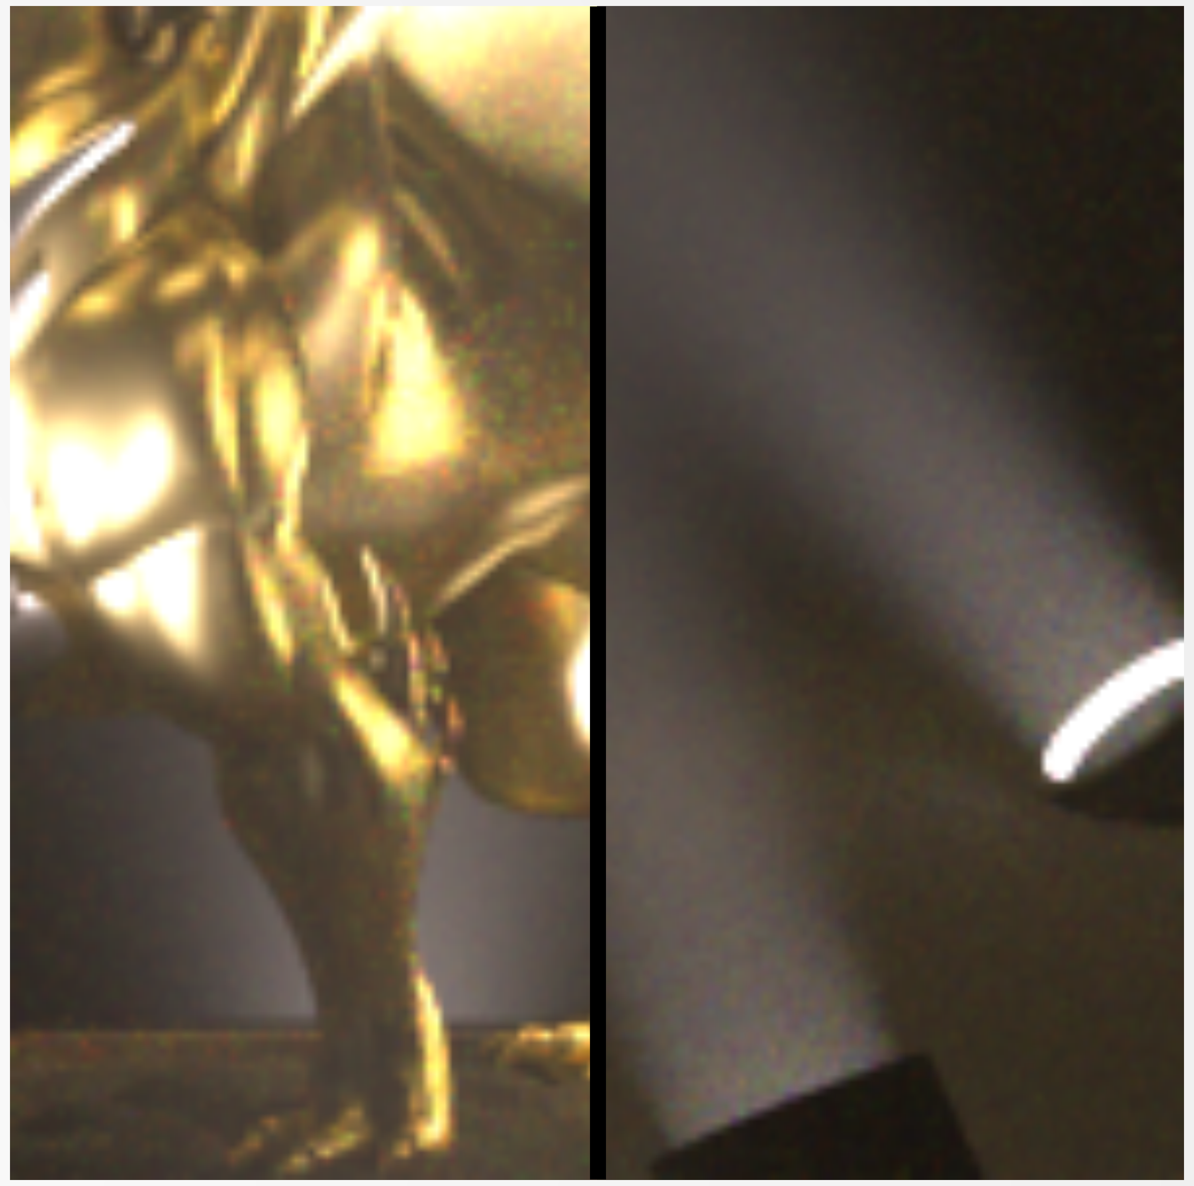
Learning time	3min	3min	3min
Number of guide paths	33k	104k (≈300MB)	69k

Pool DBOR Cascade

path tracing

guided path tracing

Guiding Behaviour

Guided path tracing - iteration 1

Guiding Behaviour

Guided path tracing - iteration 2

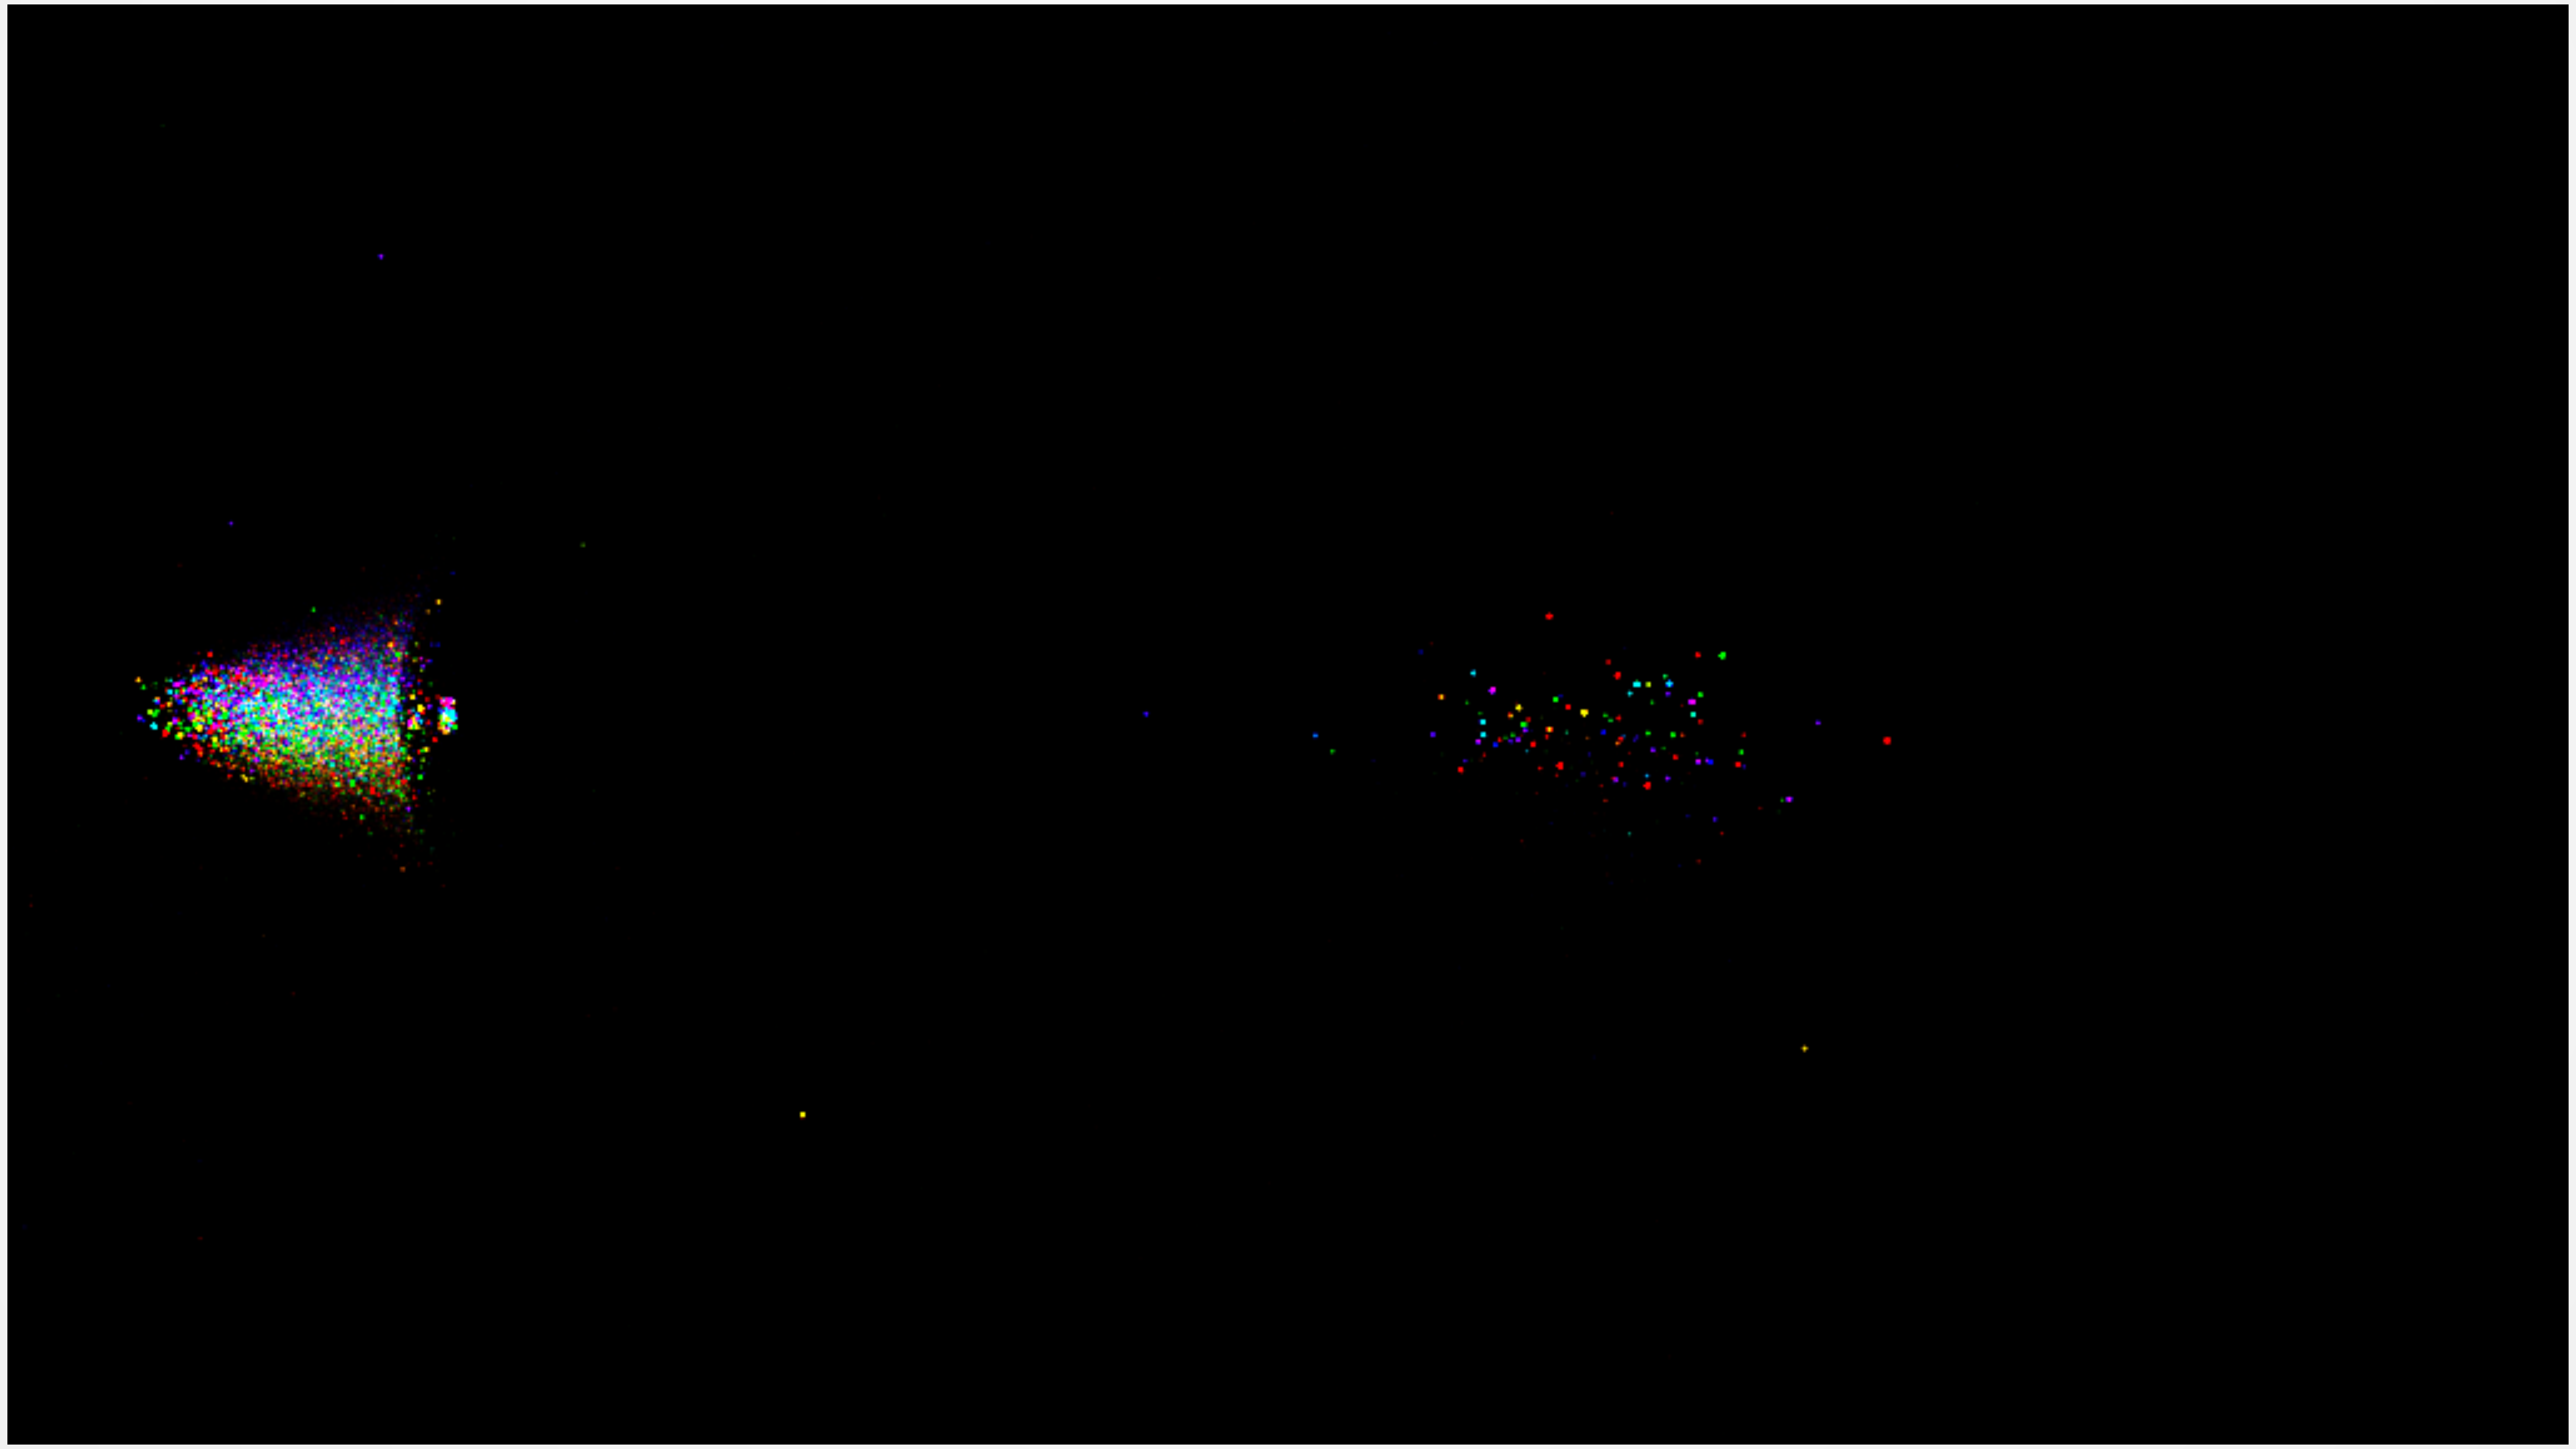

Guiding Behaviour

Guided path tracing - iteration 3

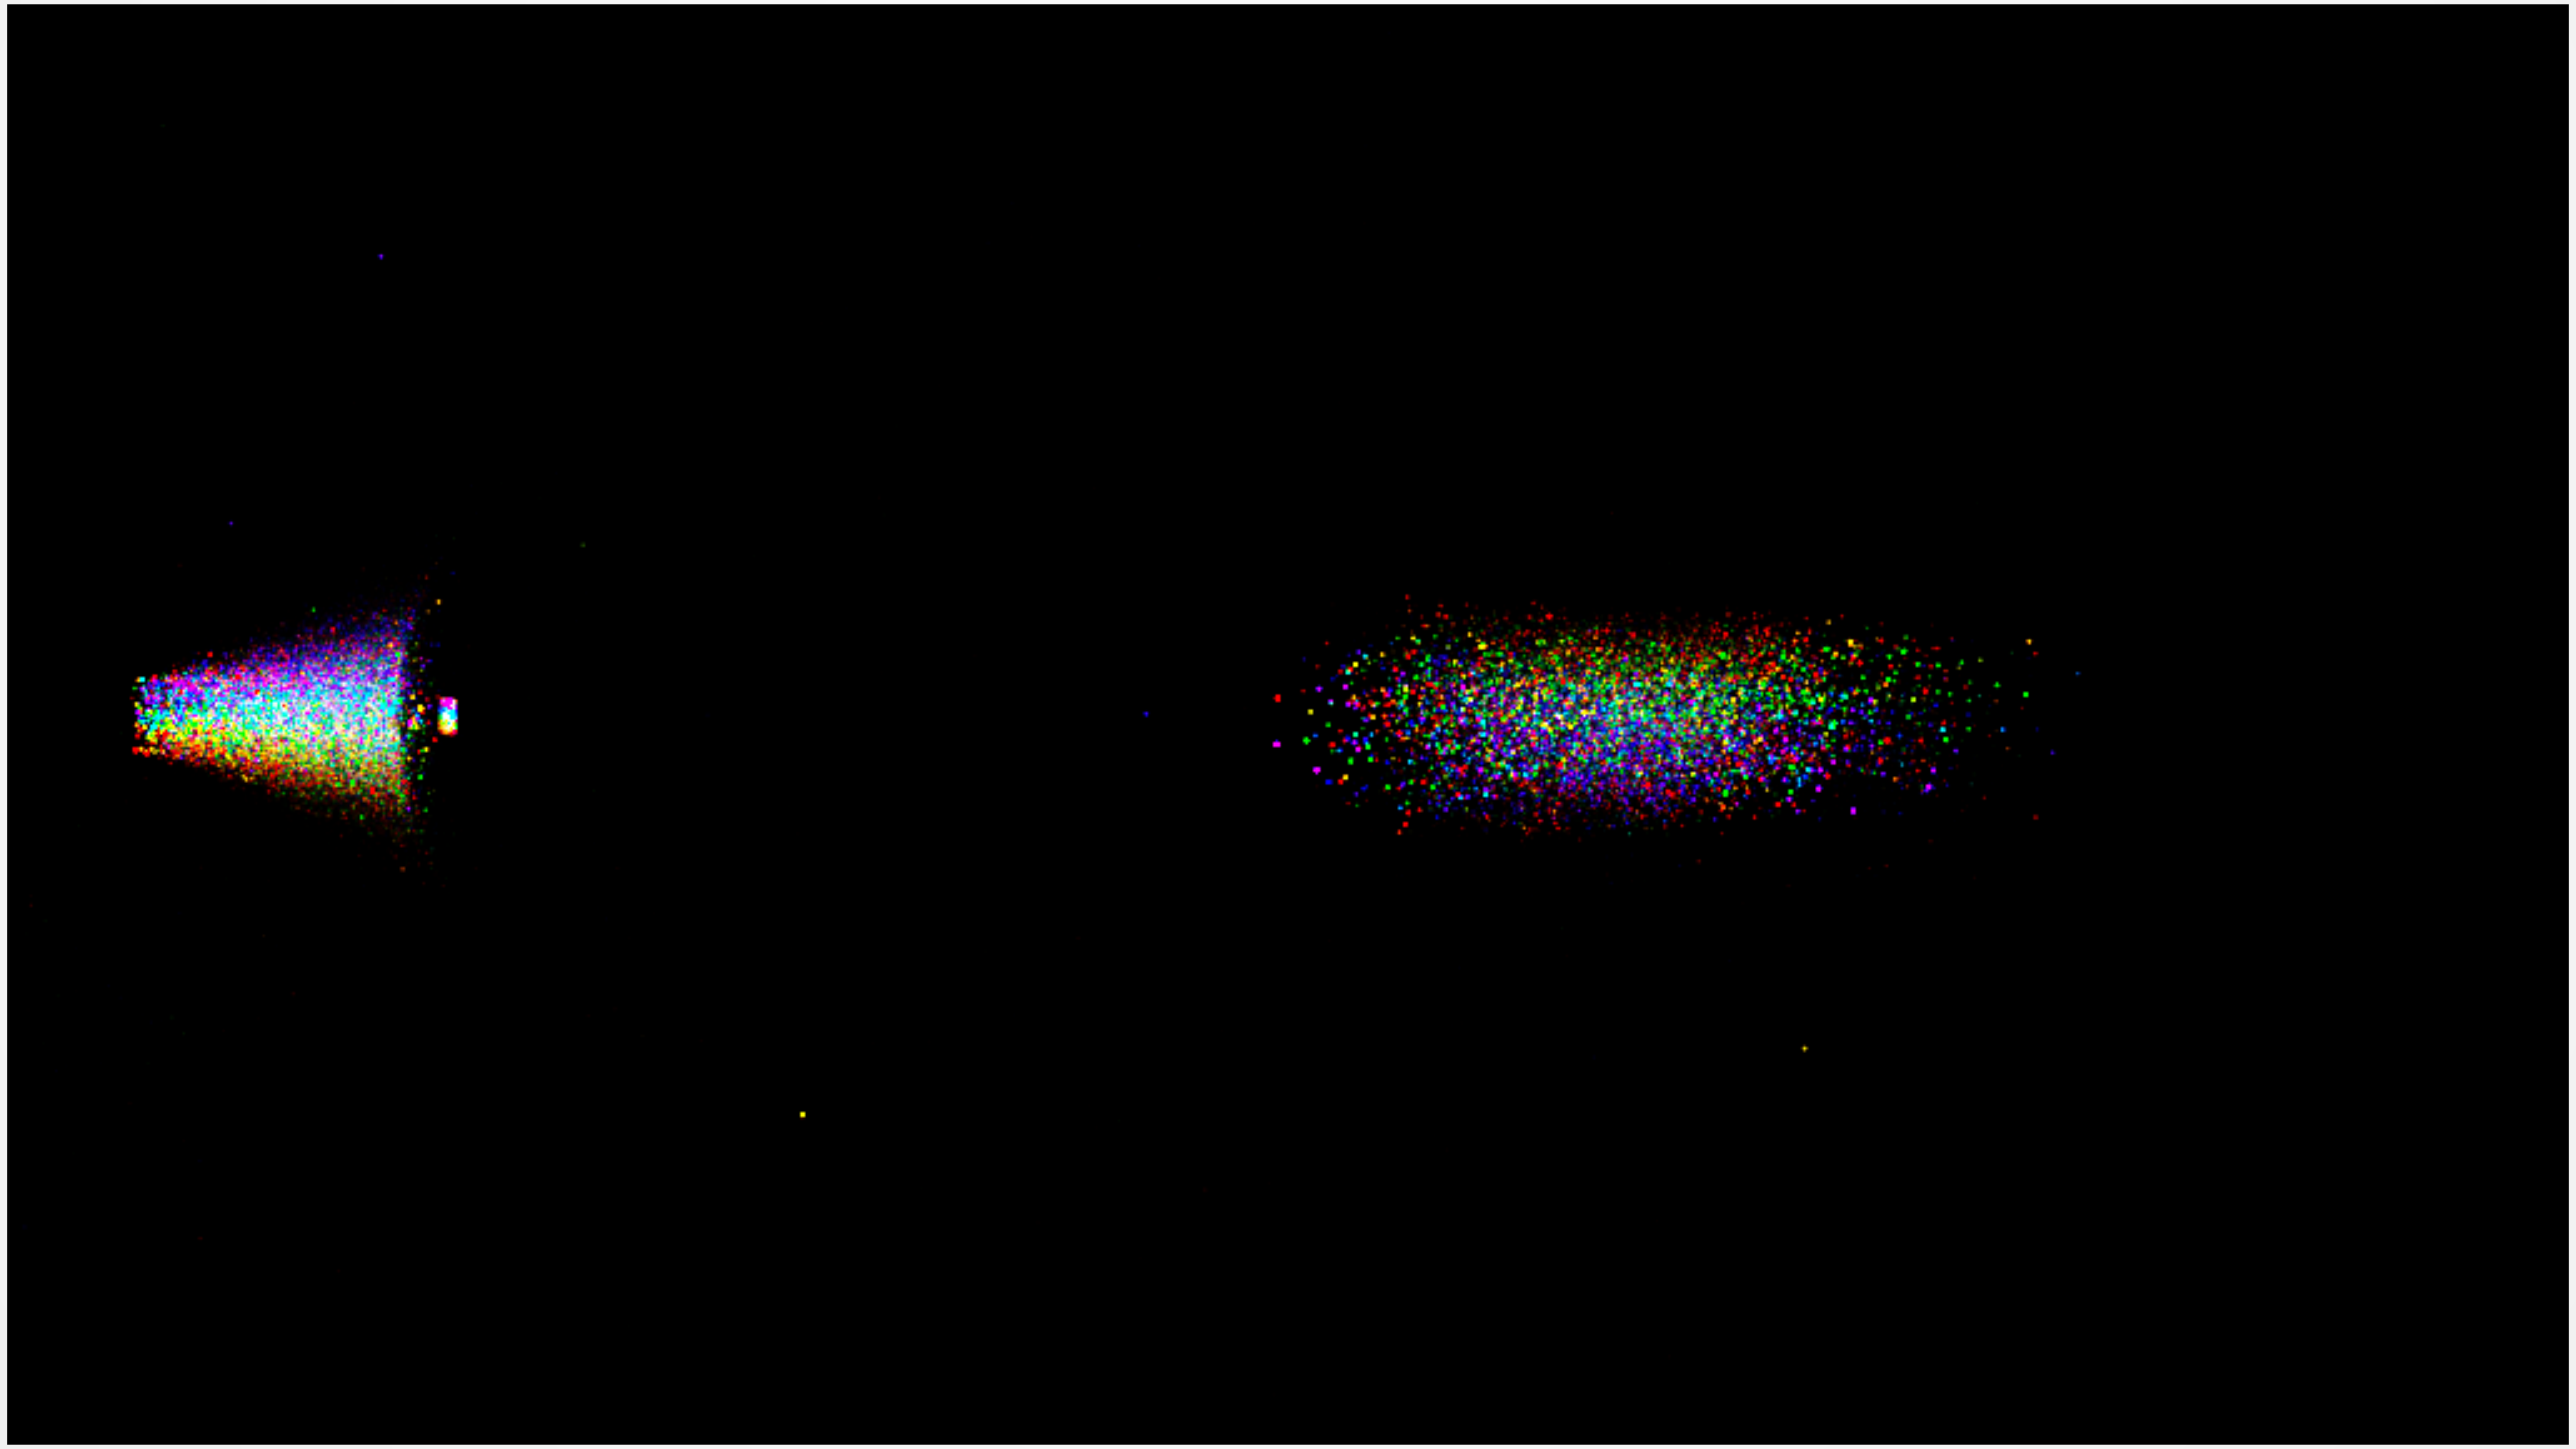

Guiding Behaviour

Guided path tracing - iteration 4

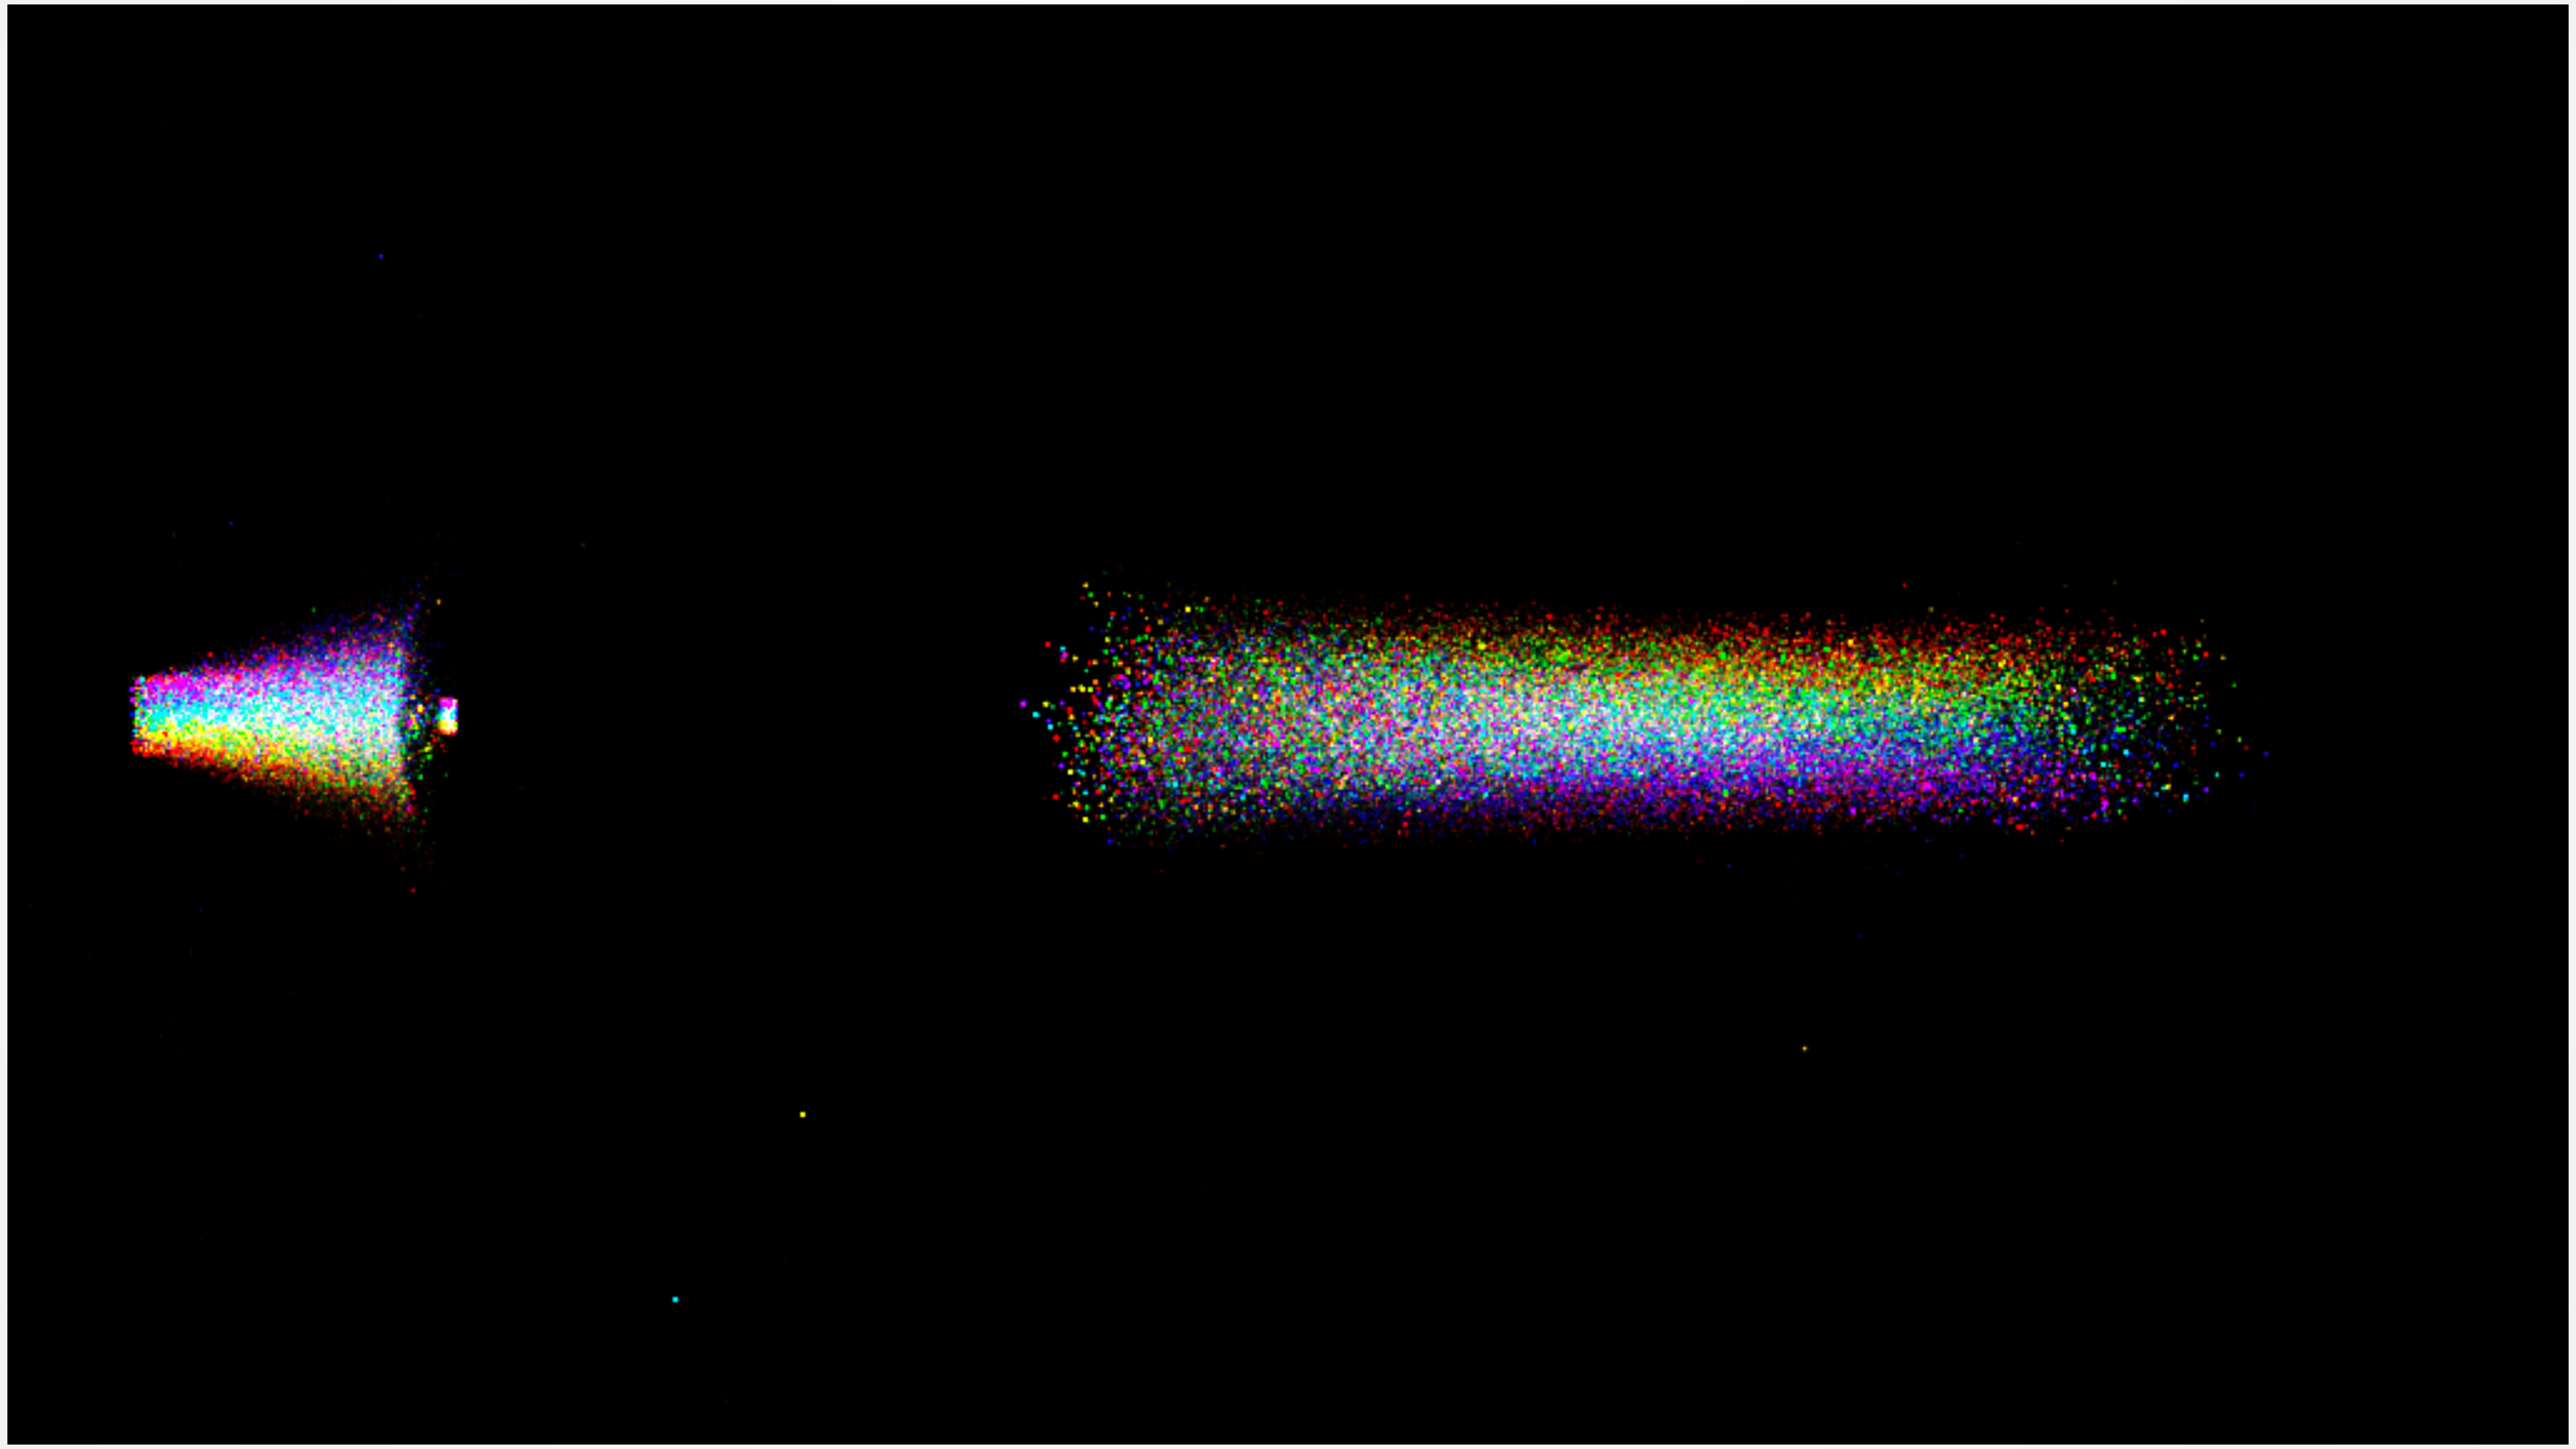

Guiding Behaviour

Guided path tracing - iteration 8

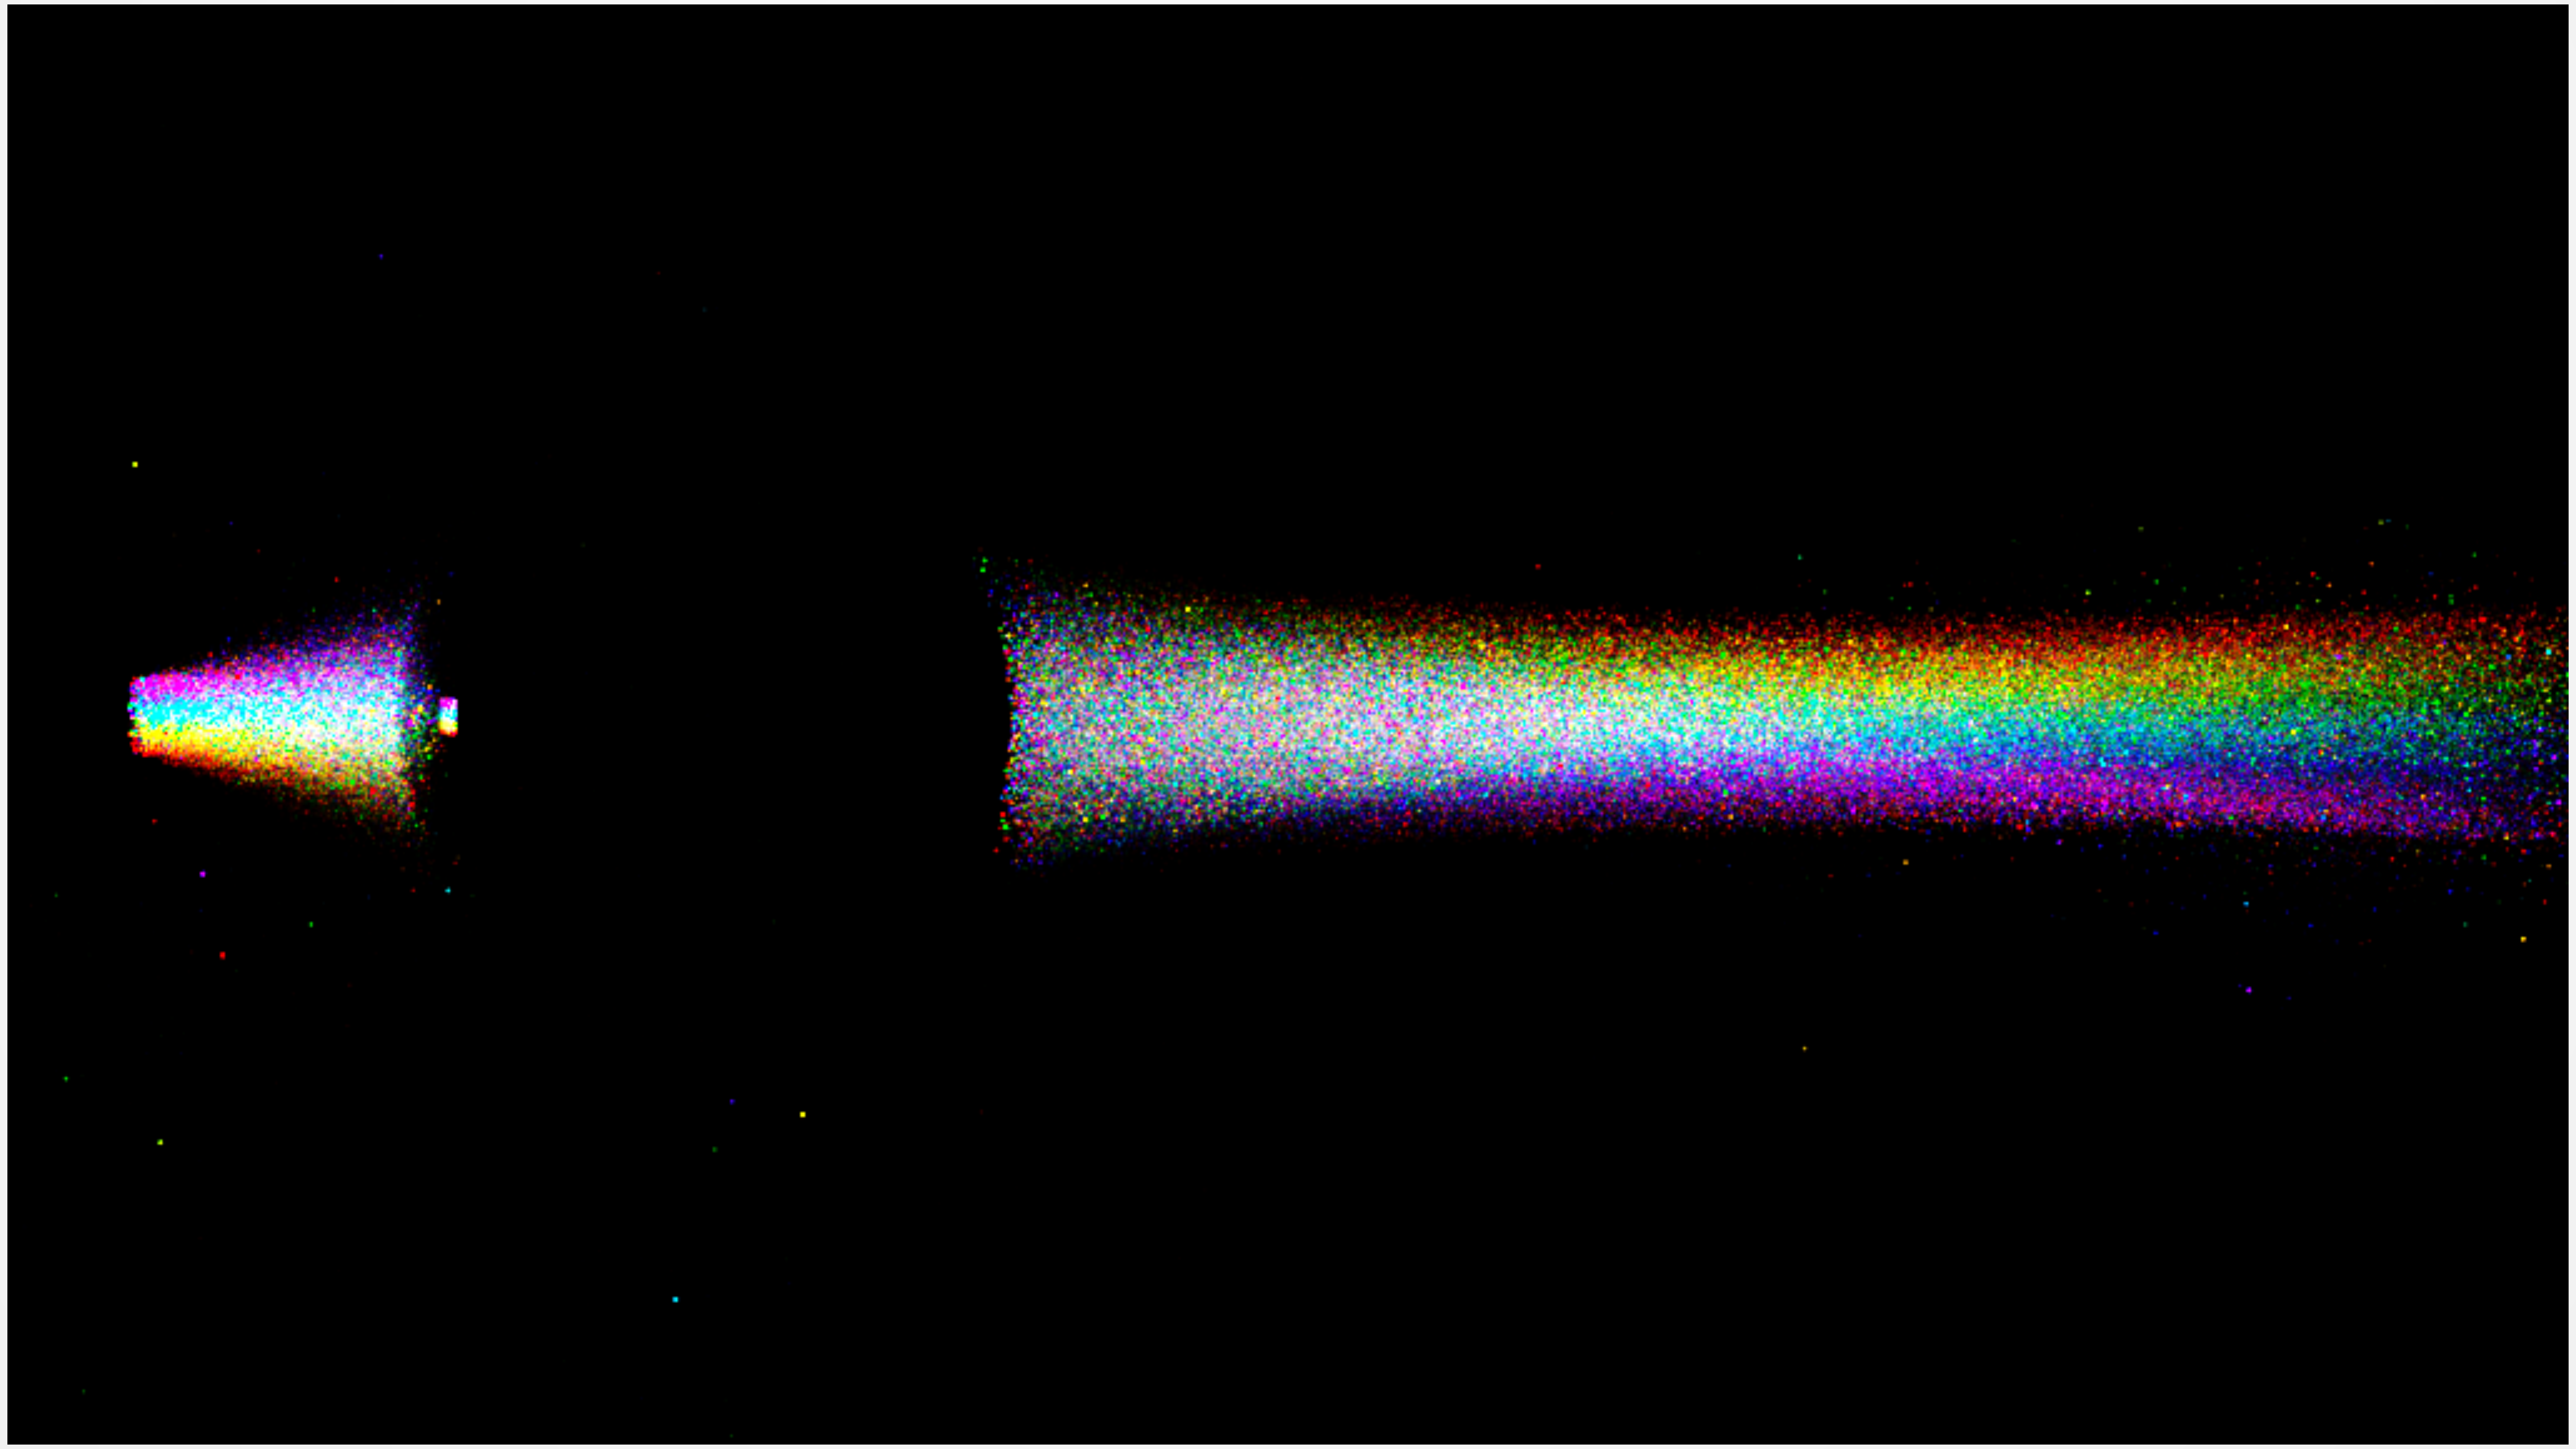

Guiding Behaviour

Guided path tracing - iteration 16

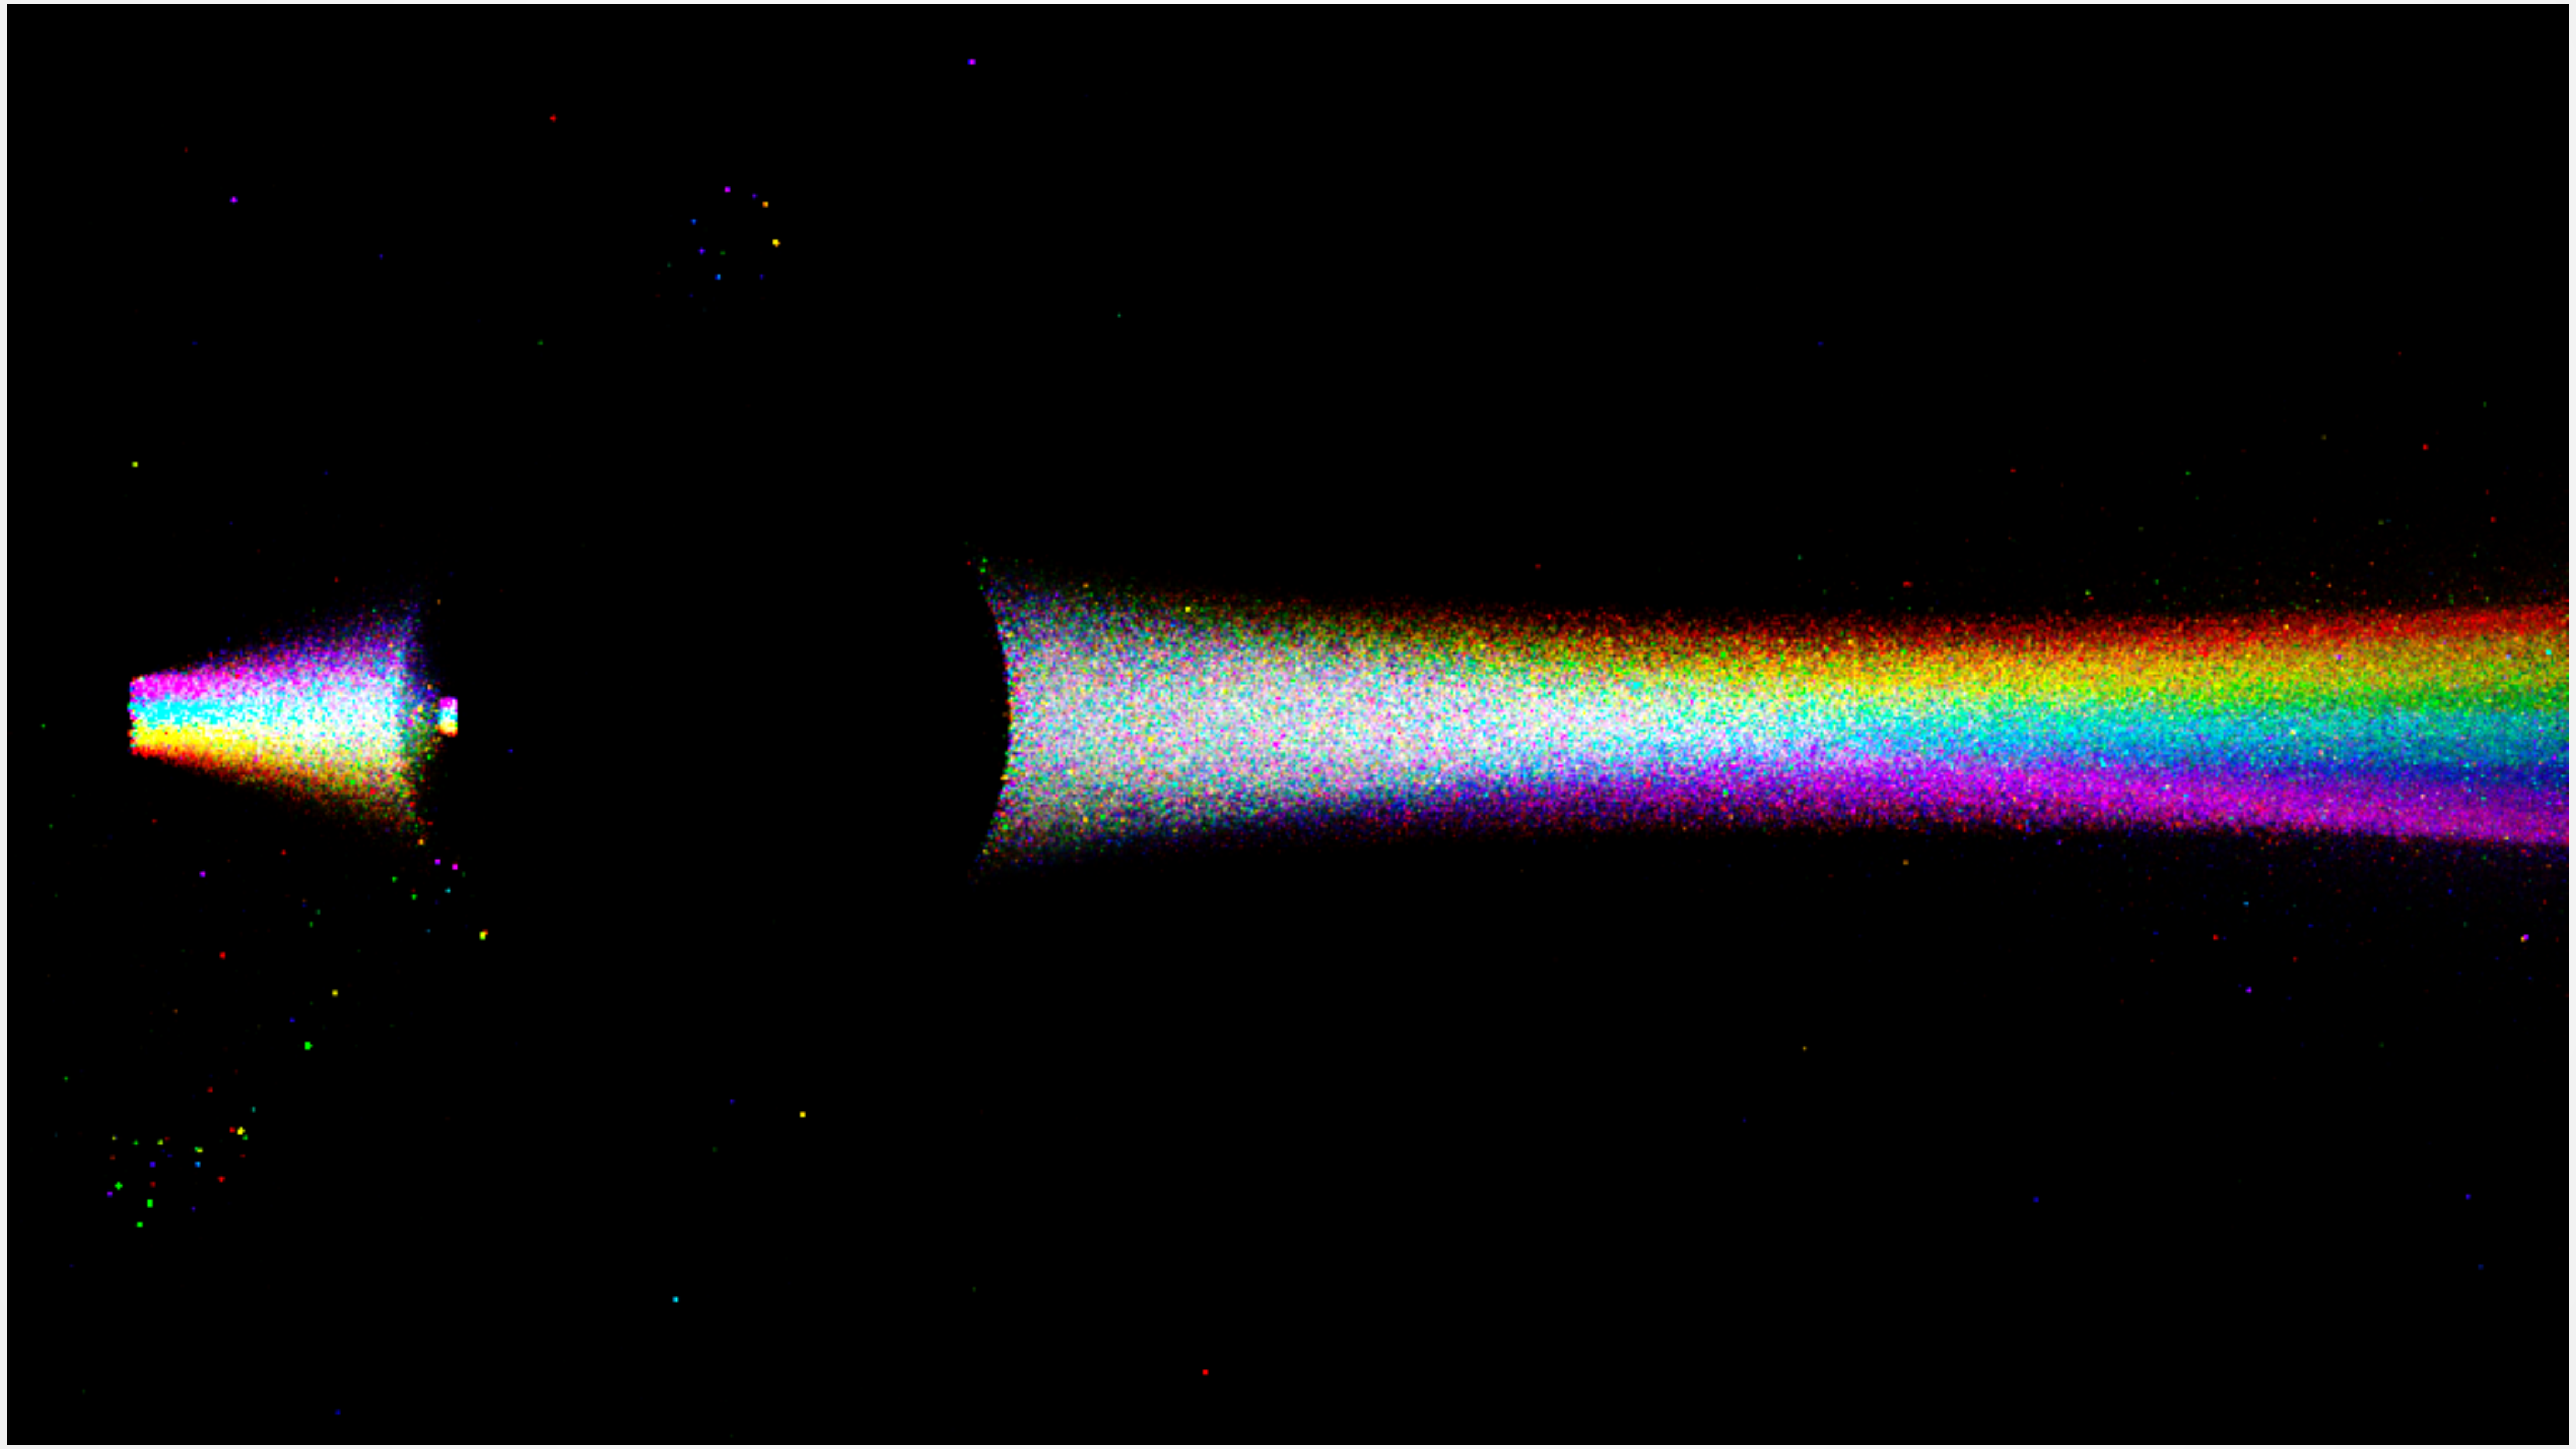

Limitations & Future Work

- Temporal stability is challenging

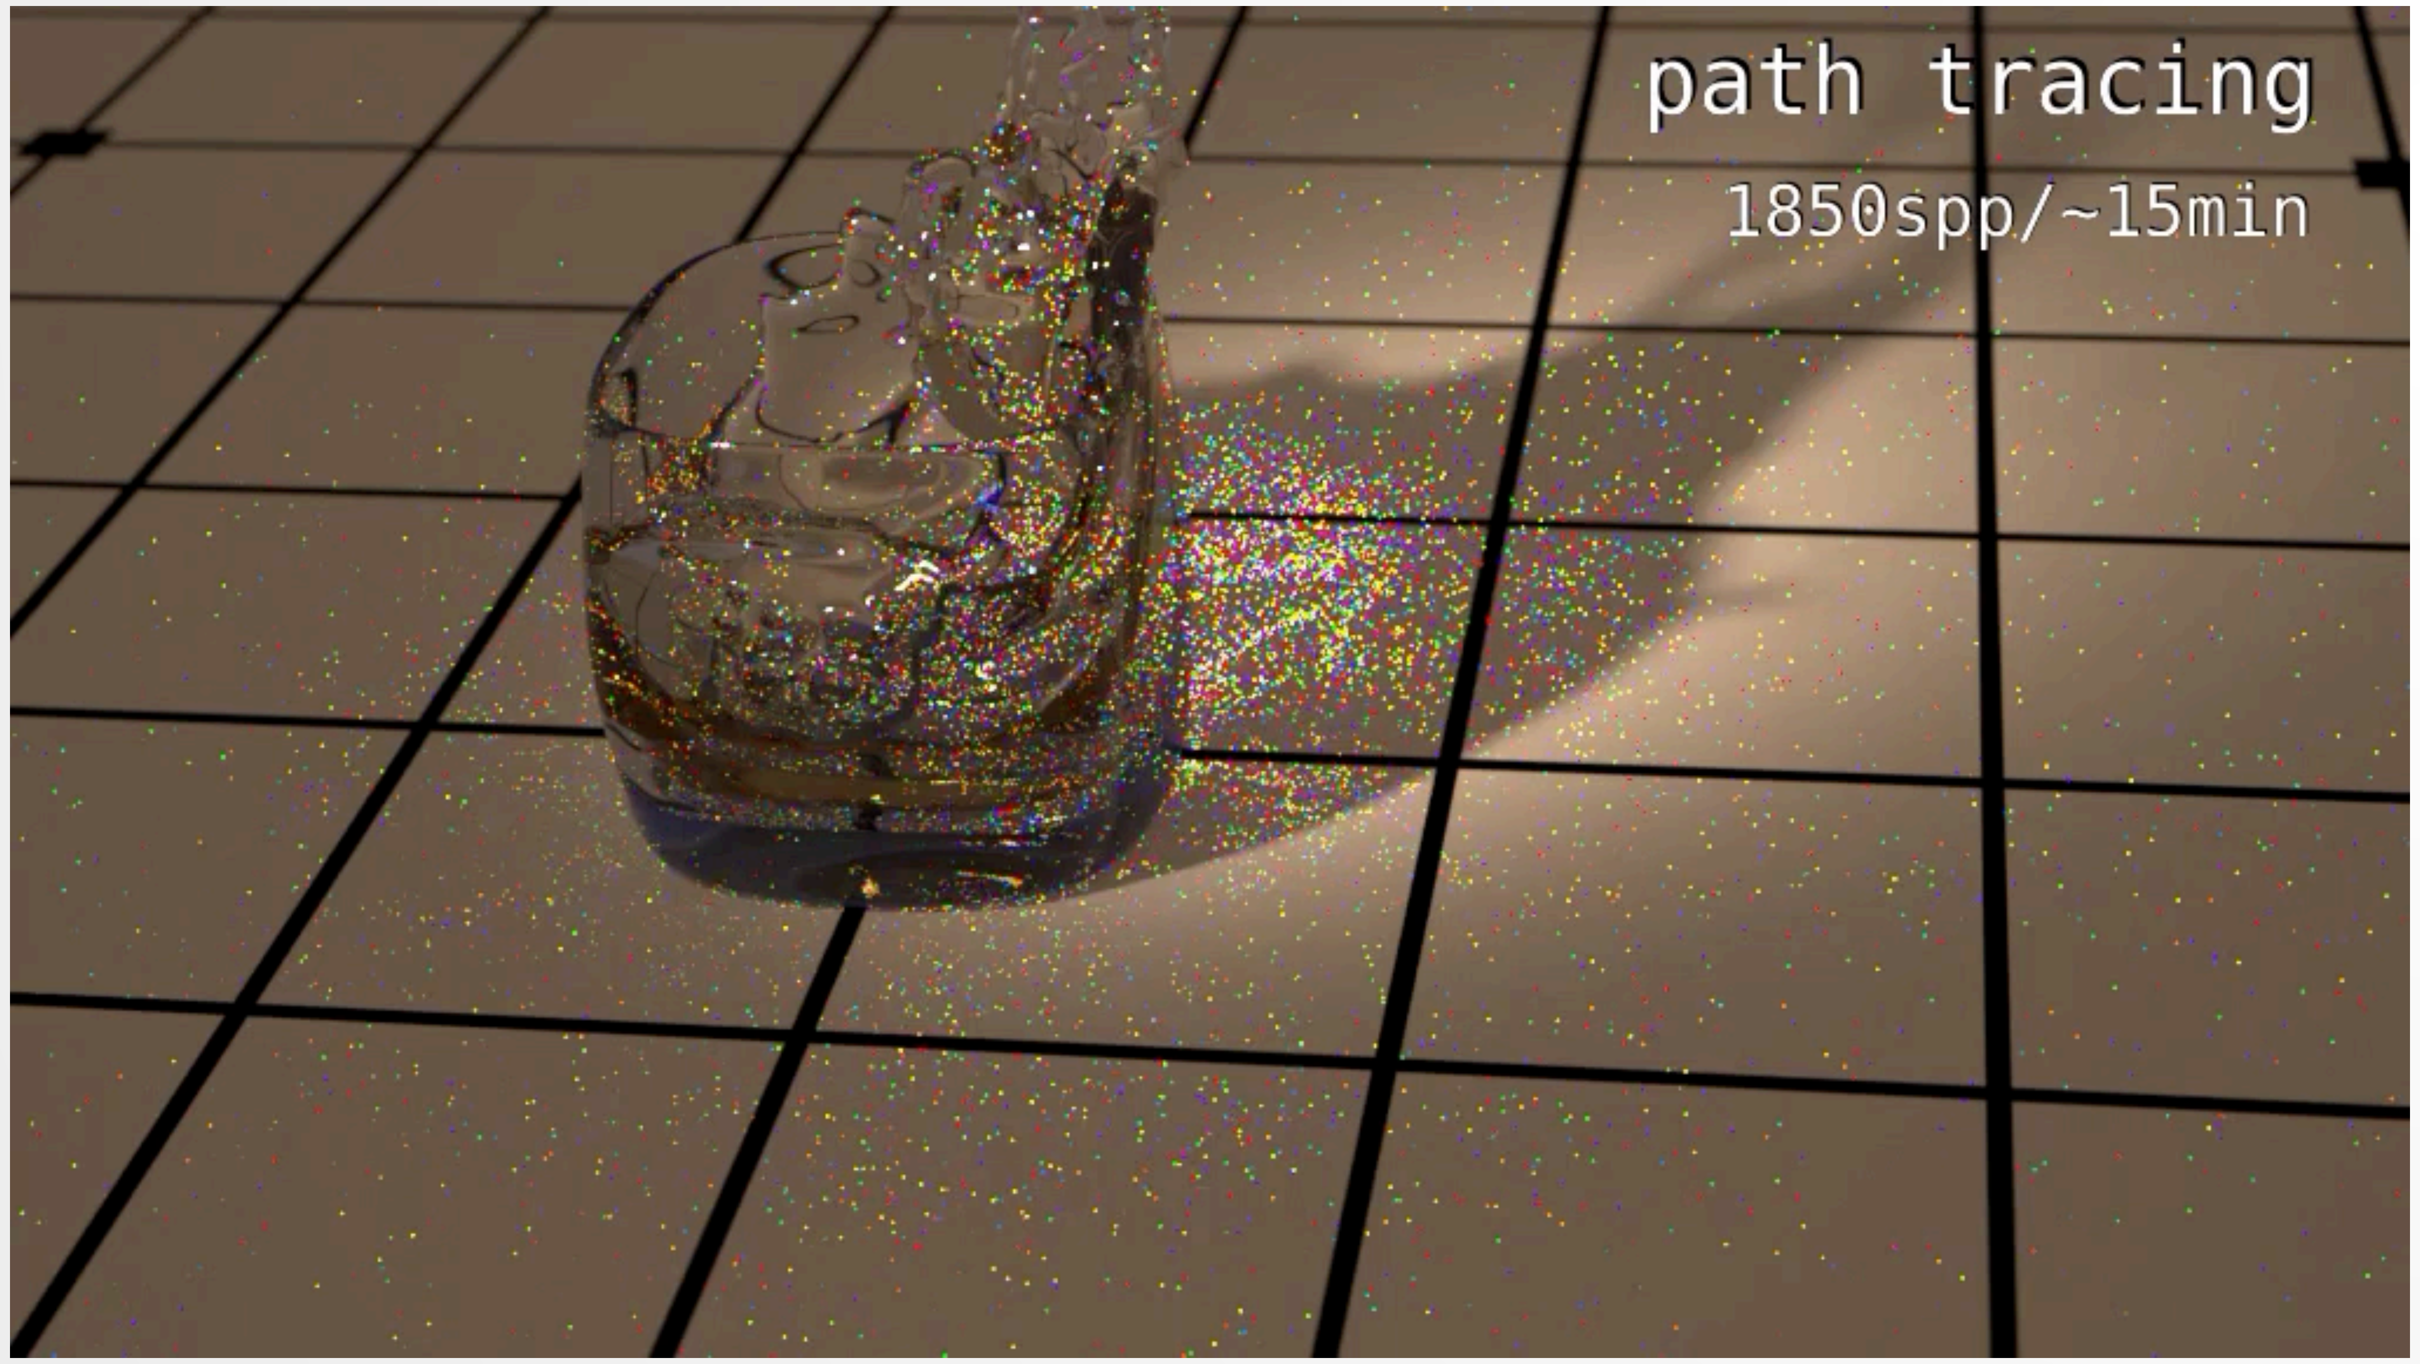

Limitations & Future Work

- Temporal stability is challenging

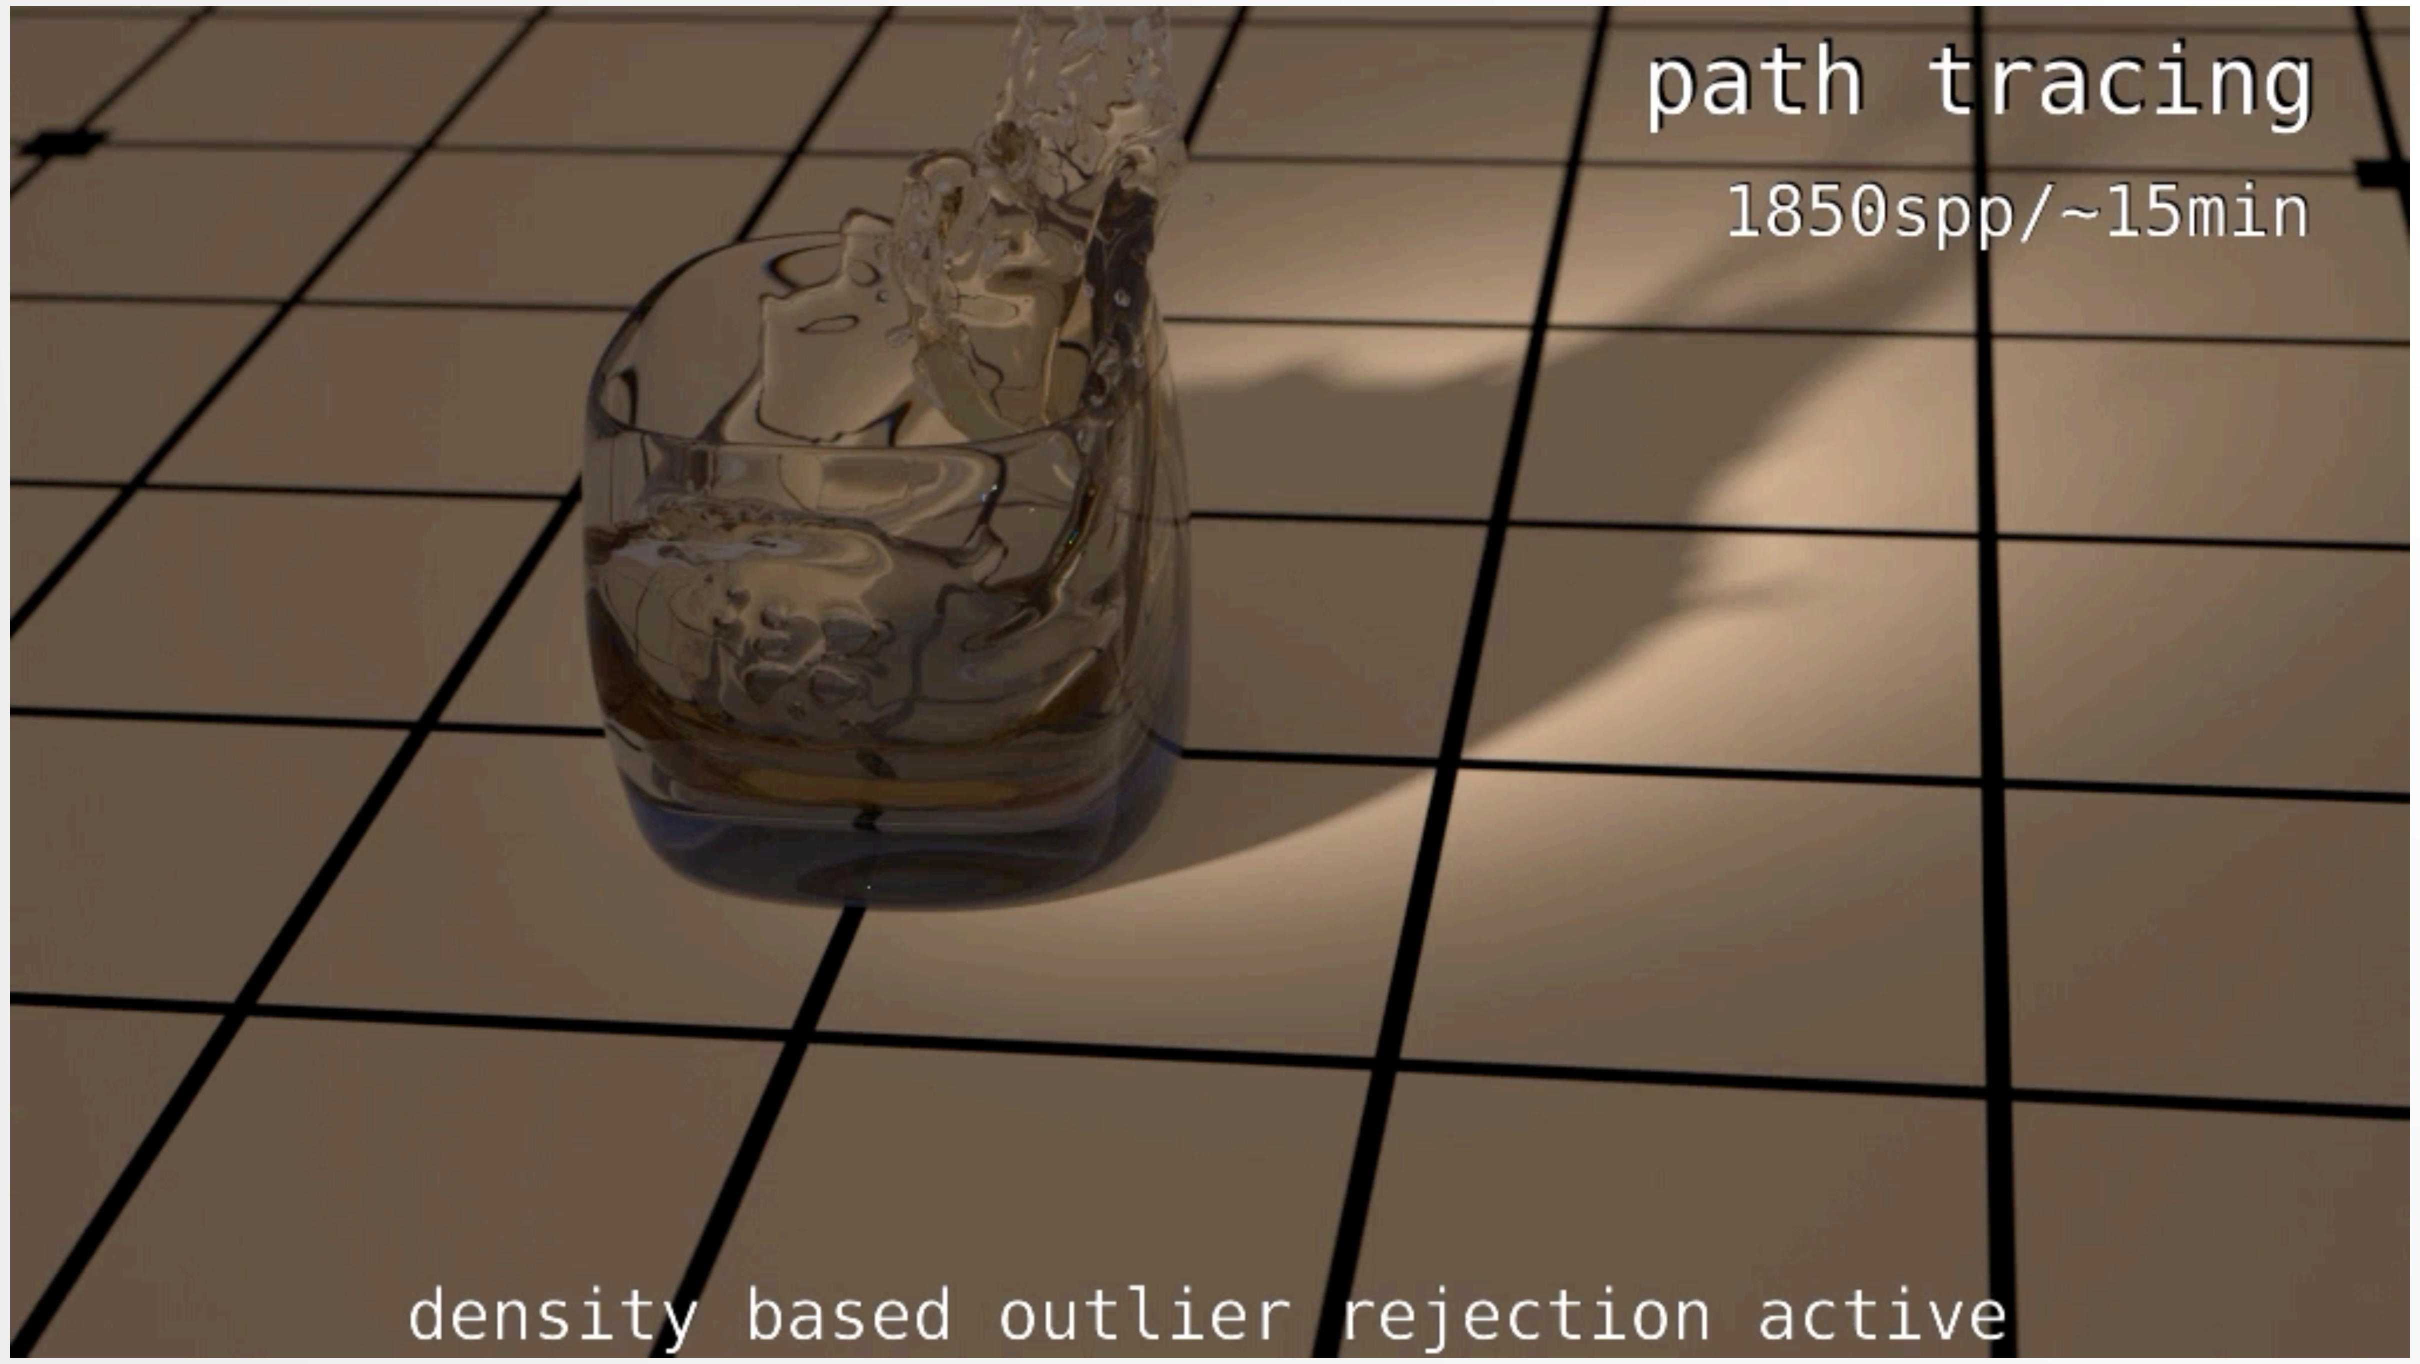

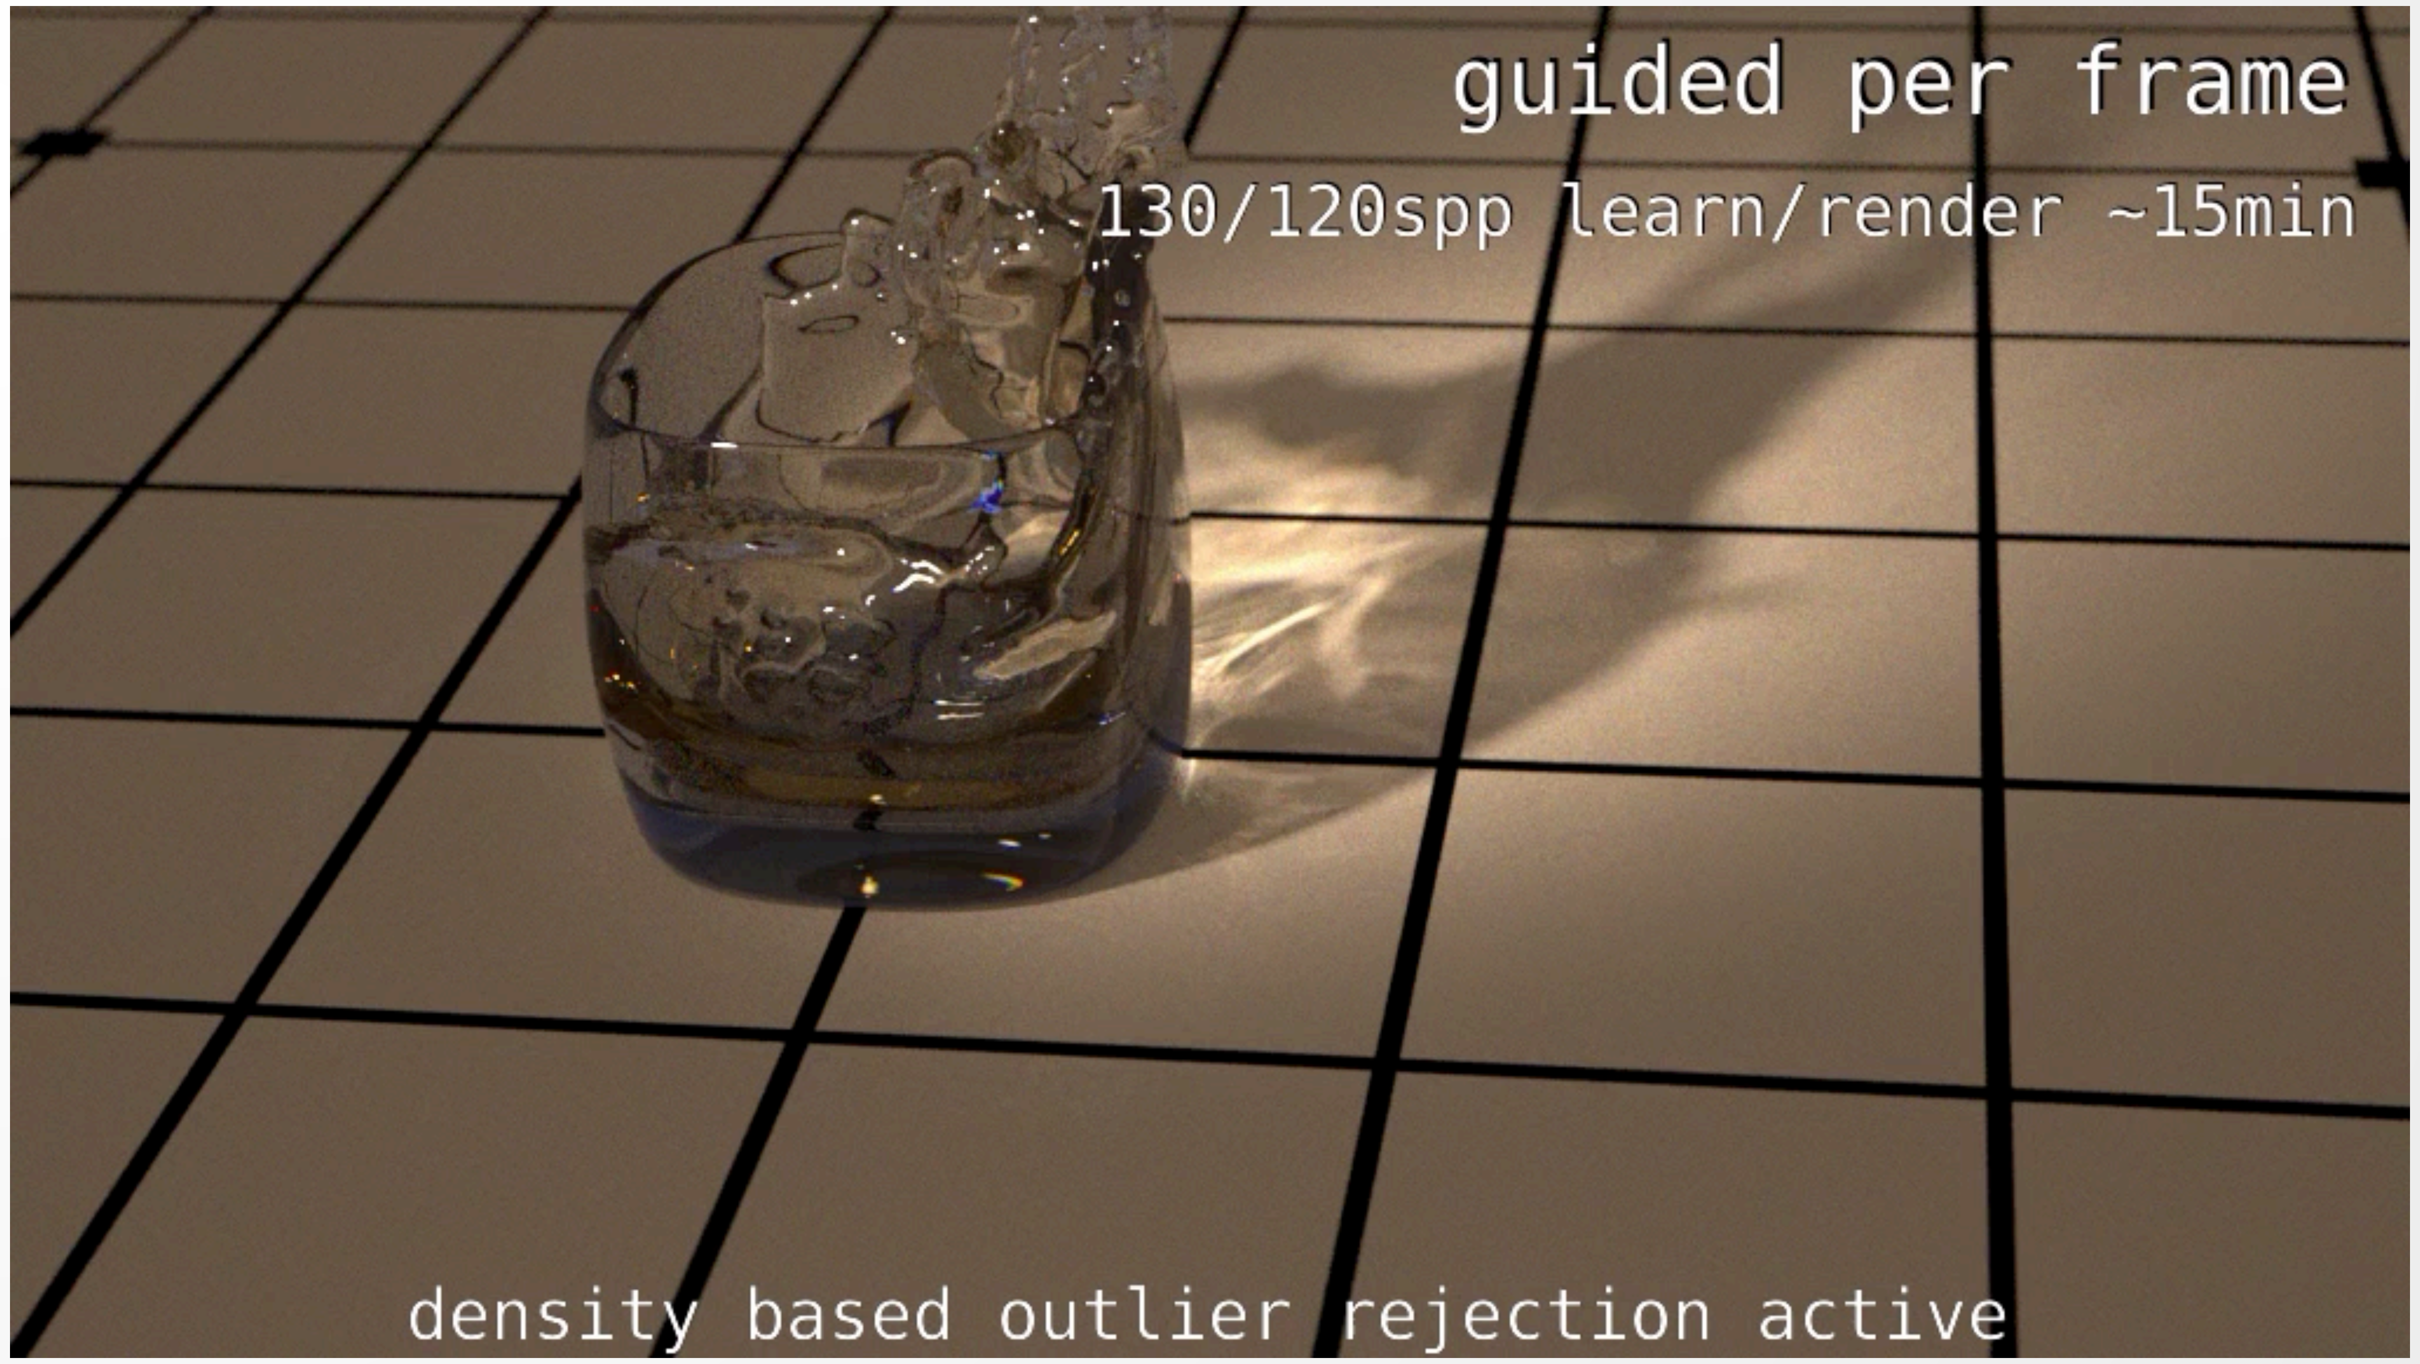

Guide path sampling

- Pick guide path X_j randomly using weights w_j
- Then in iteration k: $p_g^k(X) = \sum_l w_j^k p_g^k(X \mid X_l)$
- Choosing $w_j^k = \dfrac{C^k(X_j)}{\sum_l C^k(X_l)}$

  with $C^k(X_i) = \dfrac{f(X_i)}{\alpha p_u(X_i) + (1-\alpha) p_g^k(X_i)}$

circular dependency

Guide path sampling

- Pick guide path X_j randomly using weights w_j
- Then in iteration k: $p_g^k(X) = \sum_l w_j^k p_g^k(X \mid X_l)$
- Therefore $w_j^k = \dfrac{C^{k-1}(X_j)}{\sum_l C^{k-1}(X_l)}$

oscillation

with $C^k(X_i) = \dfrac{f(X_i)}{\alpha p_u(X_i) + (1-\alpha) p_g^k(X_i)}$

Guide path sampling

- Pick guide path X_j randomly using weights w_j
- Then in iteration k: $p_g^k(X) = \sum_l w_j^k p_g^k(X \mid X_l)$
- Therefore $w_j^k = t \cdot \frac{C^{k-1}(X_j)}{\sum_l C^{k-1}(X_l)} + (1-t) \cdot w_j^{k-1}$

with $C^k(X_i) = \frac{f(X_i)}{\alpha p_u(X_i) + (1-\alpha) p_g^k(X_i)}$