Volumetric Zero-Variance-Based Path Guiding

Sebastian Herholz[1]
Derek Nowrouzezahrai [2]

Yangyang Zhao[2]
Hendrik P. A. Lensch[1]

Oskar Elek[3]
Jaroslav Křivánek[3]

[1]University of Tübingen
[2]McGill University Montreal
[3]Charles University Prague

MOTIVATION

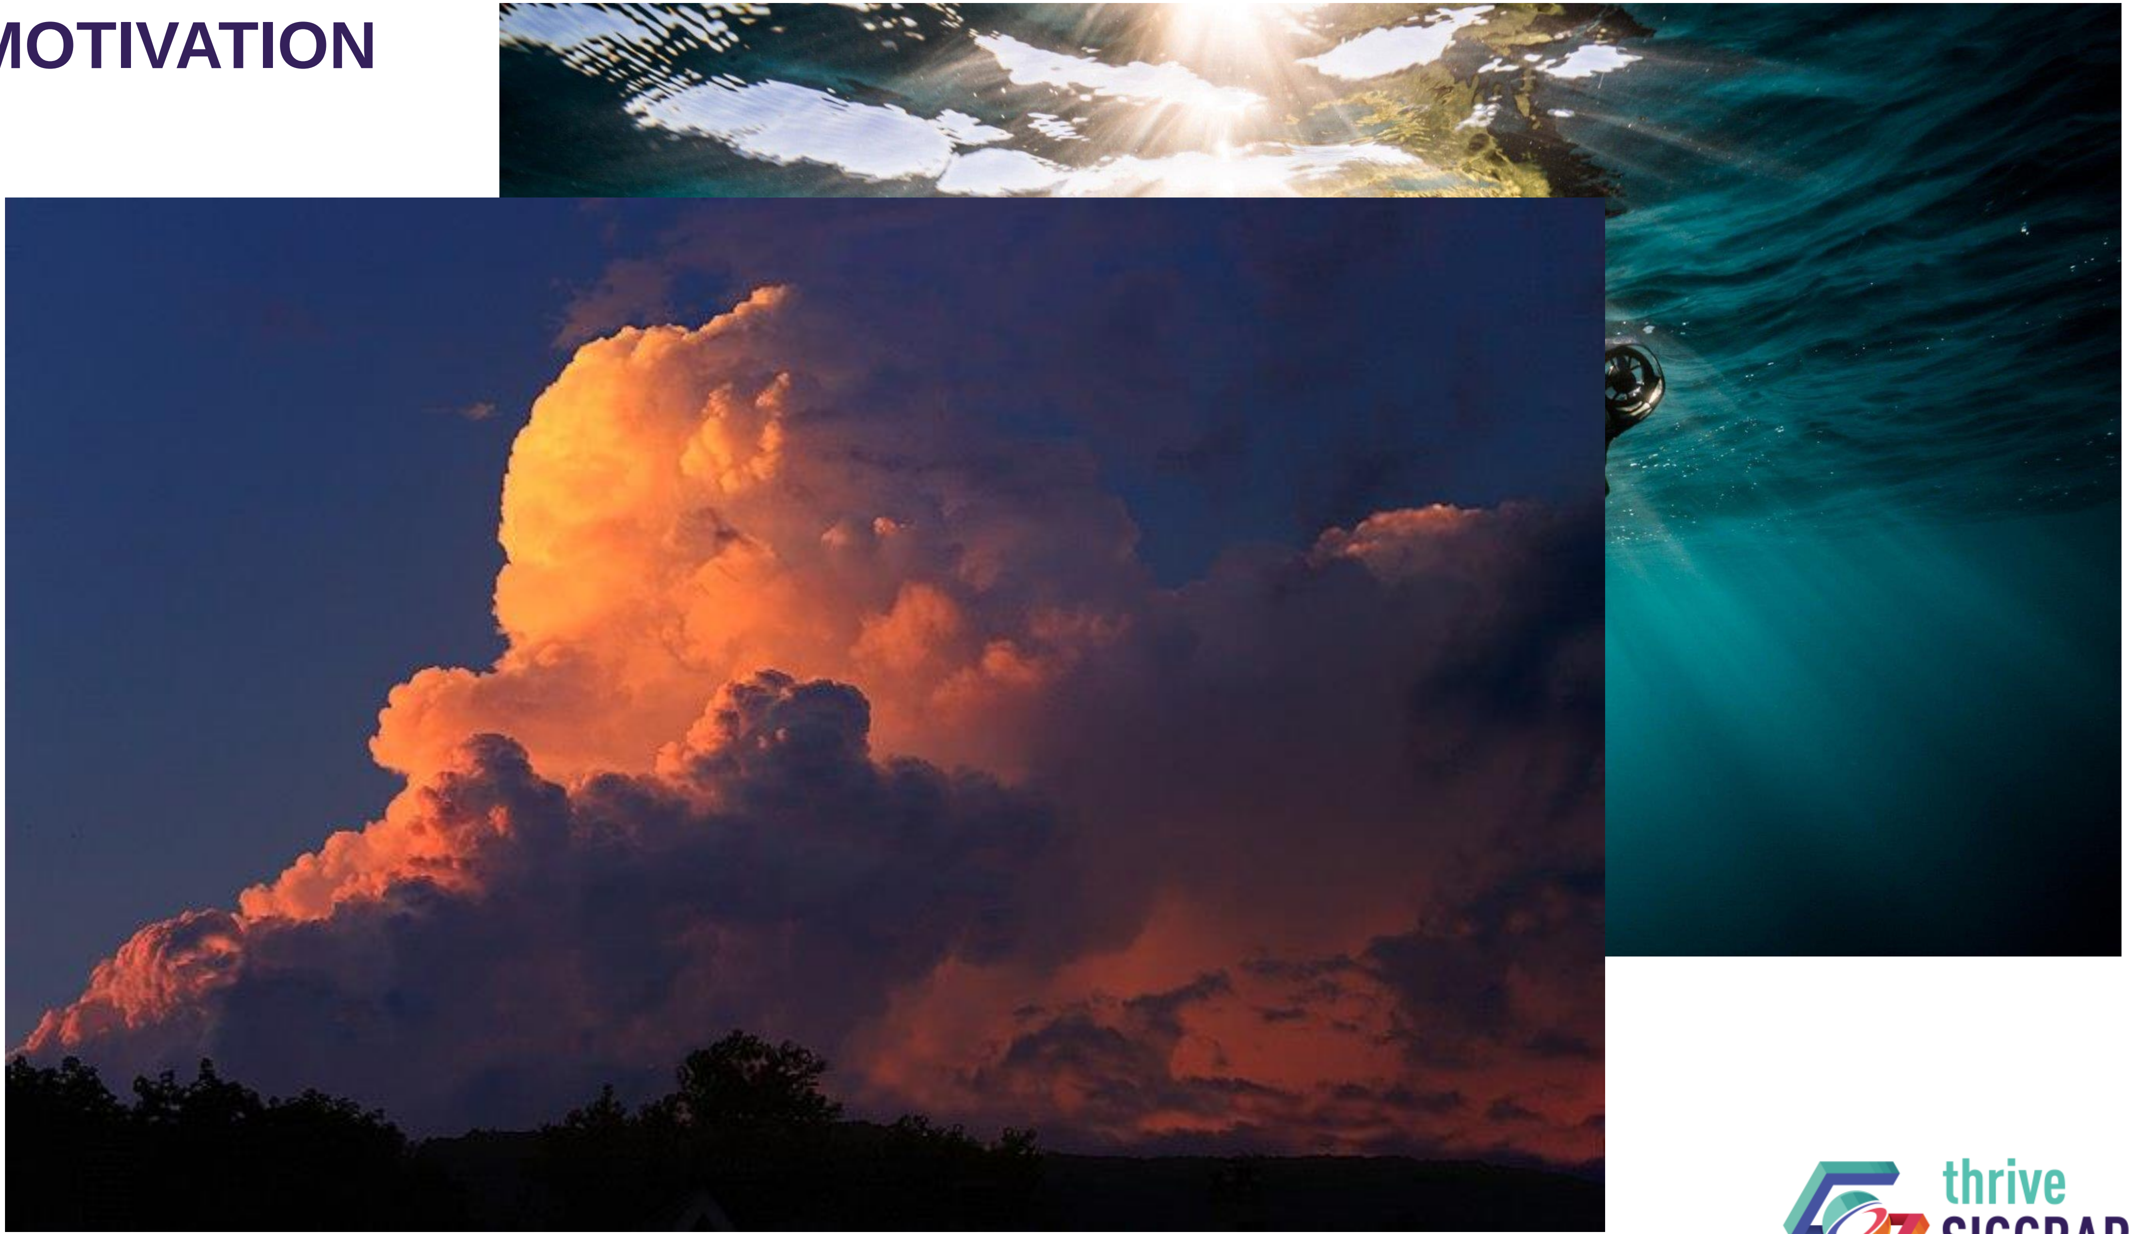

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

MOTIVATION

MOTIVATION

- A correct physically-based simulation of volumetric effects is crucial for rendering realistic scenes

- In the recent years, brute-force path tracing these effects started to become applicable in production environments ([Fong2017], [Novak2018])

- **Increased complexity of the light transport makes it still challenging**

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

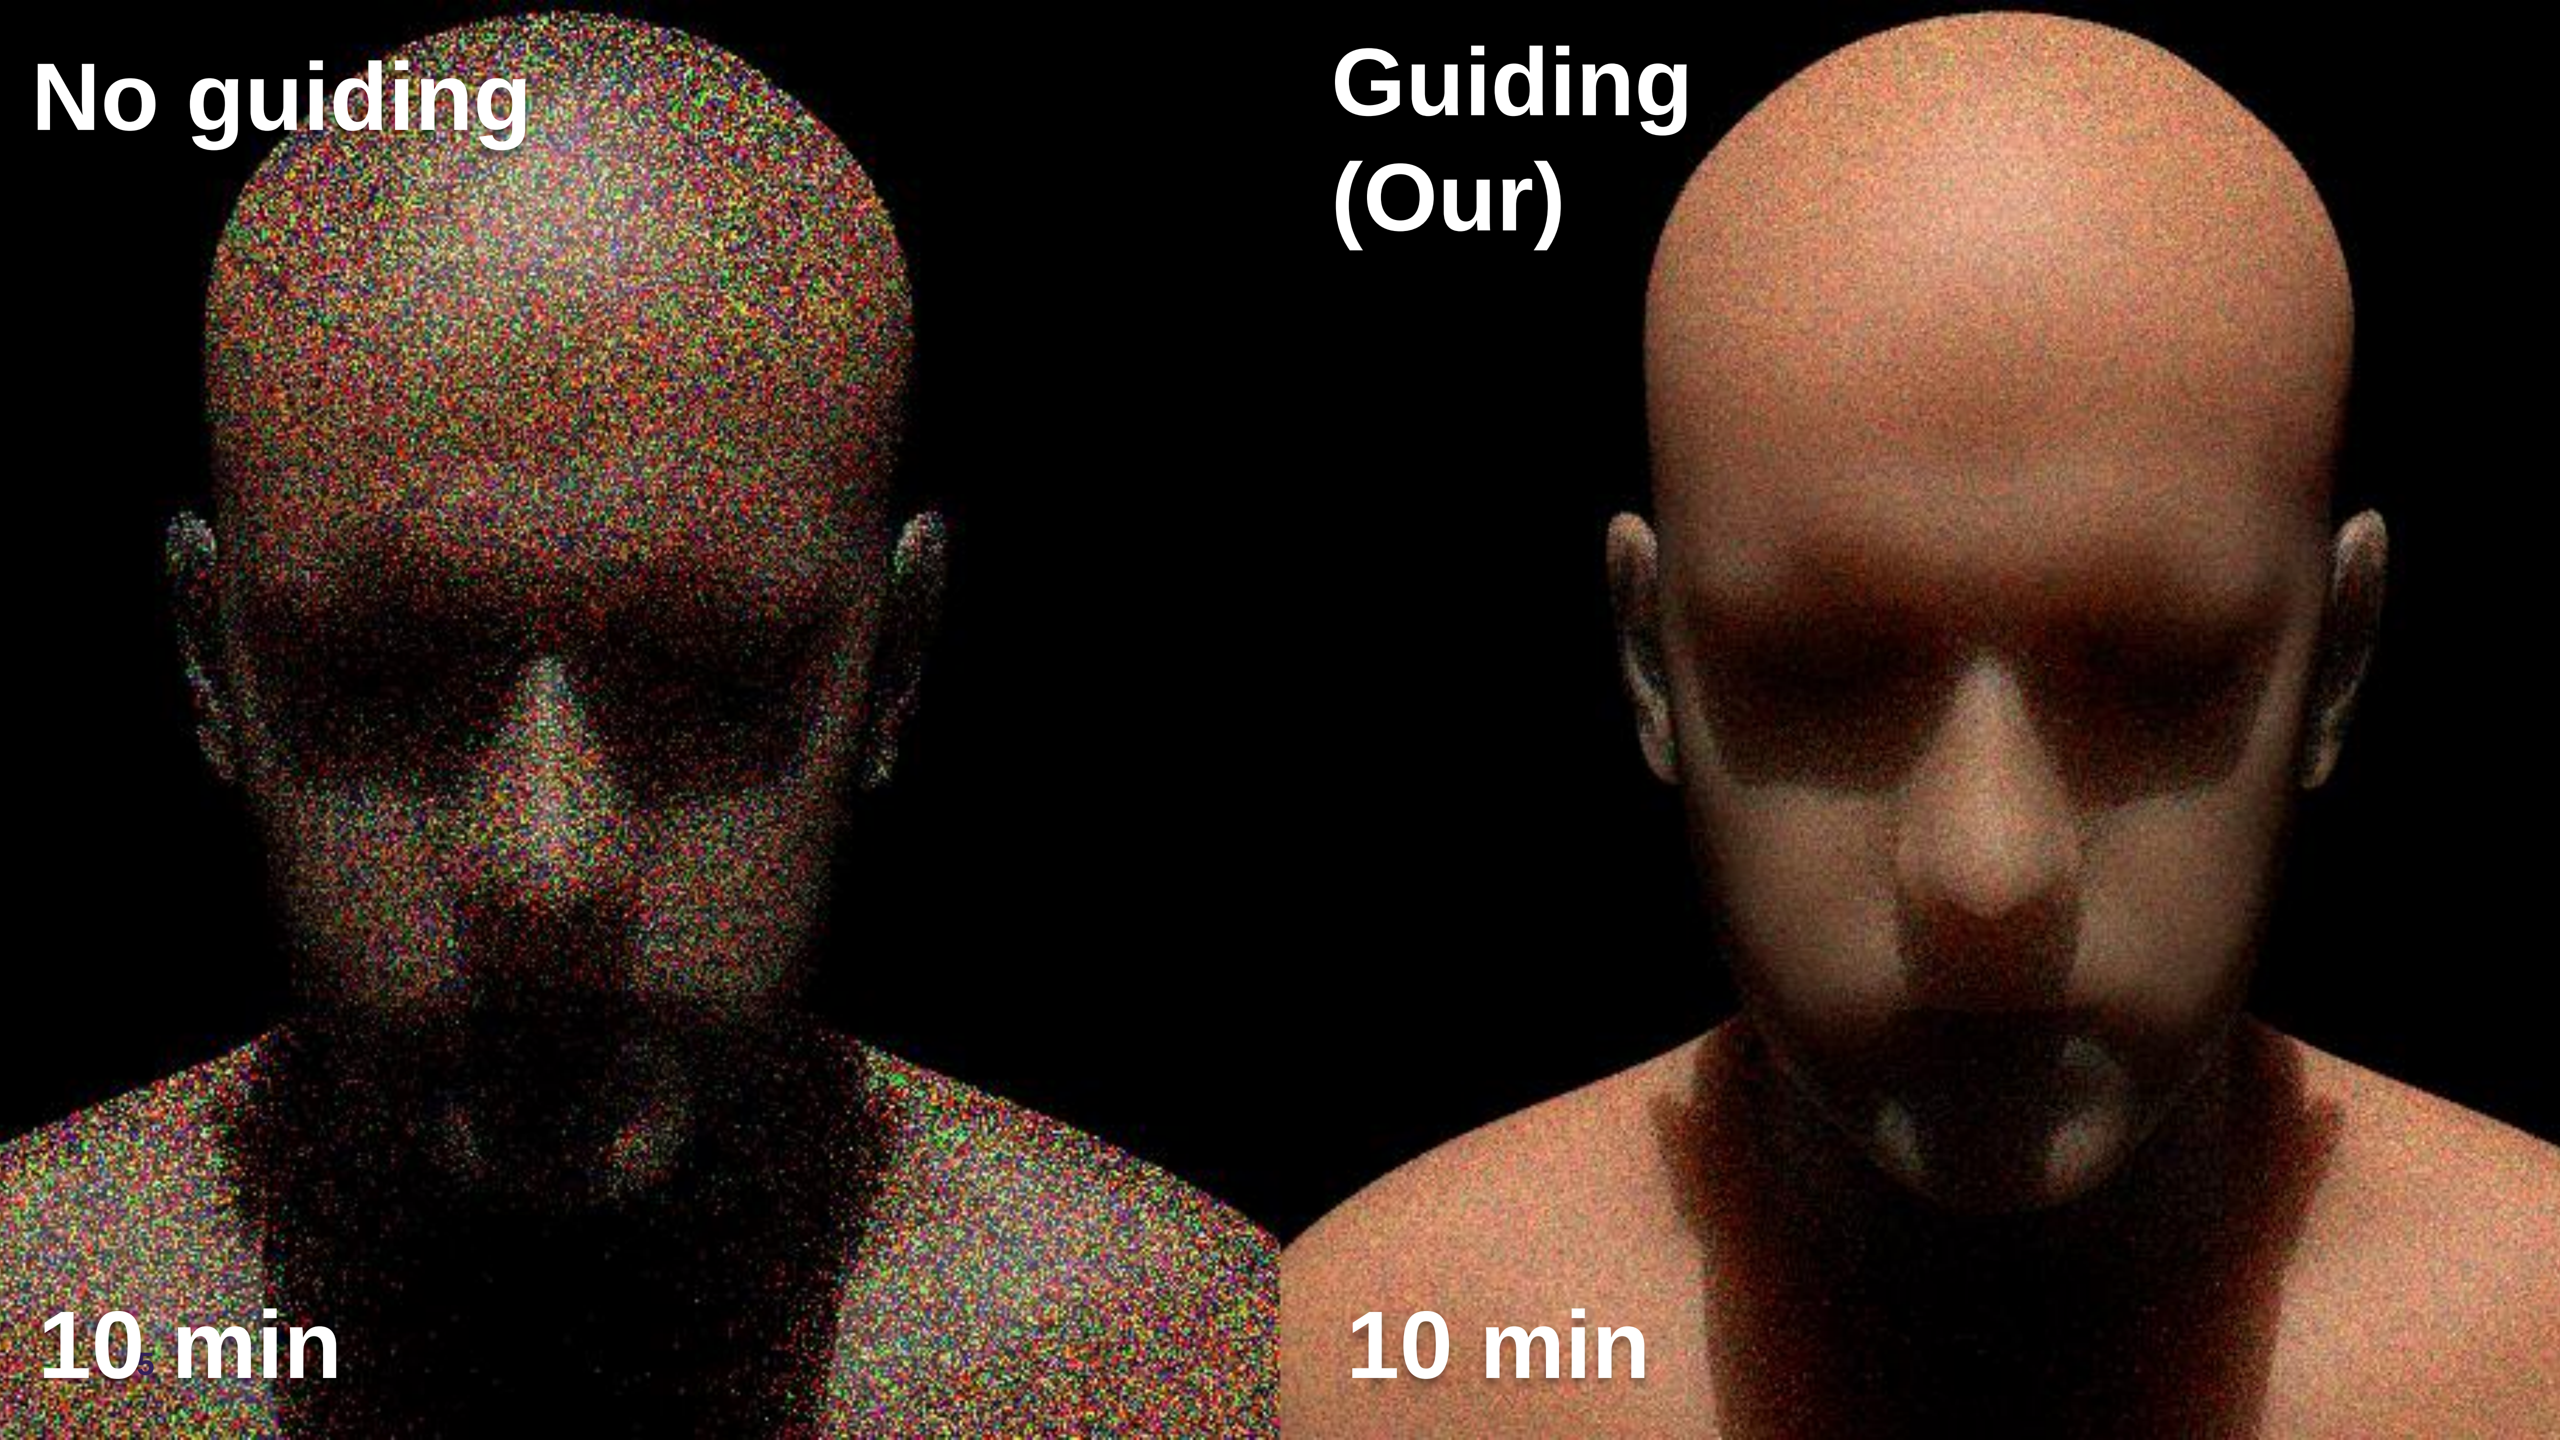
No guiding
Guiding
(Our)
10 min
10 min

VOLUMETRIC MONTE-CARLO PATH TRACING

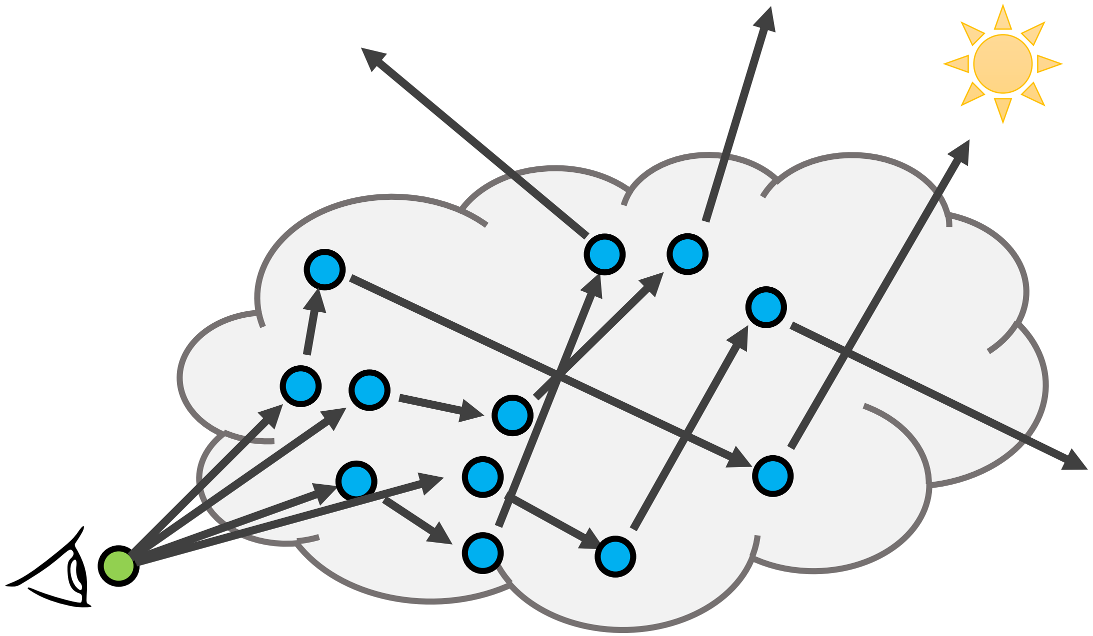

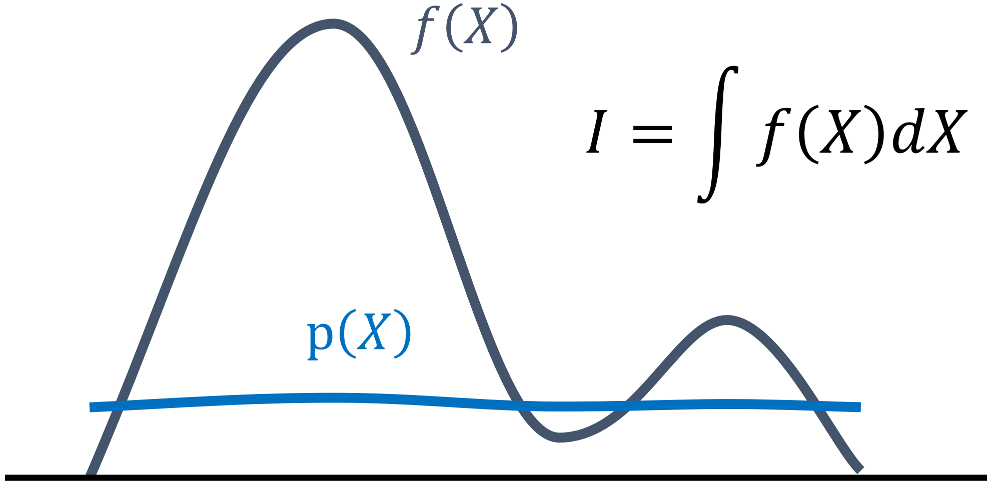

$$I = \int f(X)dX$$

- The variance is defined by how well we can generate random paths proportional to the volumetric light transport:

$$\sigma^2 = V\left[\frac{f(X)}{p(X)}\right]$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VOLUMETRIC MONTE-CARLO PATH TRACING: ZERO-VARIANCE

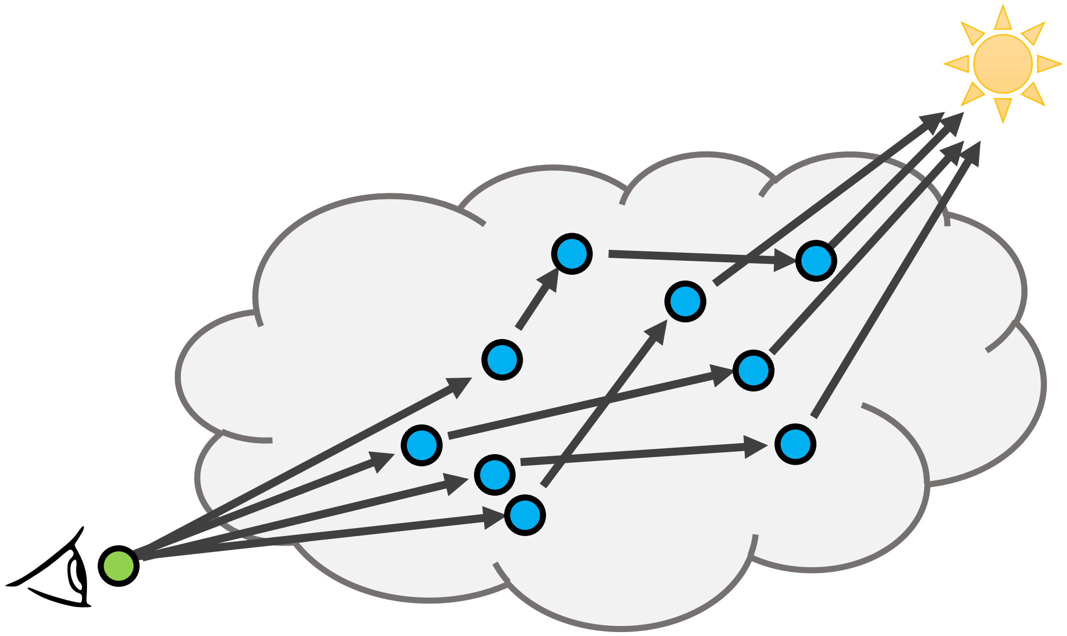

We need to know the shape of $f(X)$!

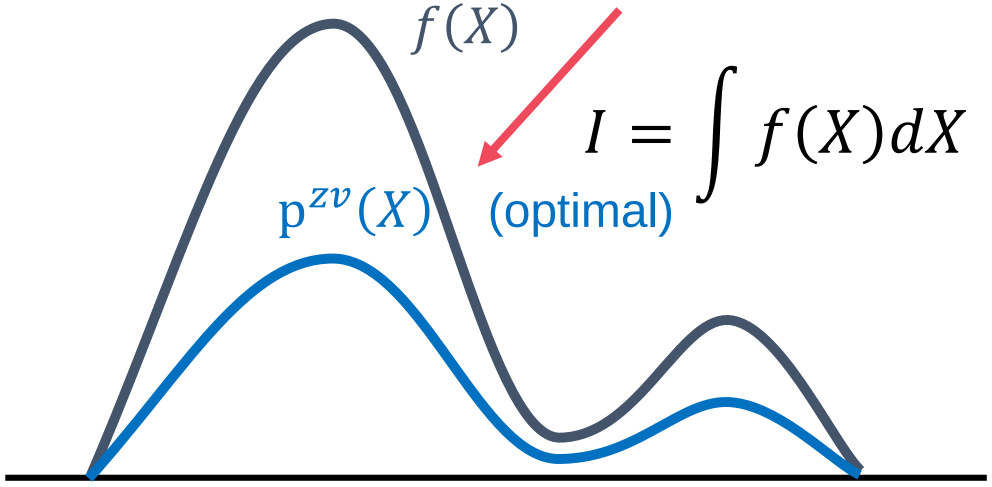

- If the PDF for all paths is proportional to the light transport function we would get a perfect **zero-variance** estimator:

$$\sigma^2 = V\left[\frac{f(X)}{p^{zv}(X)}\right] = 0$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

THE 4 SAMPLING DECISIONS:

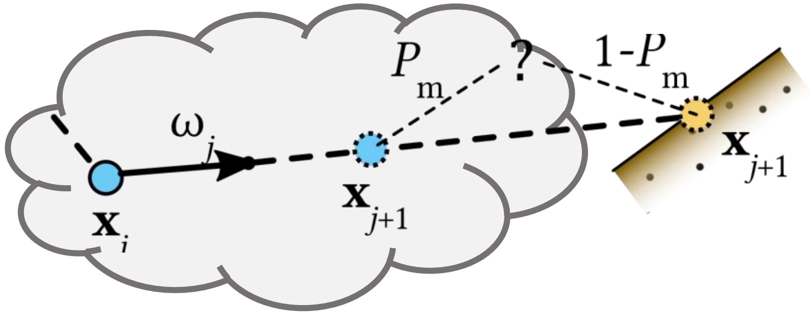

1. Scatter: $P_m(\boldsymbol{x}_j, \omega_j)$

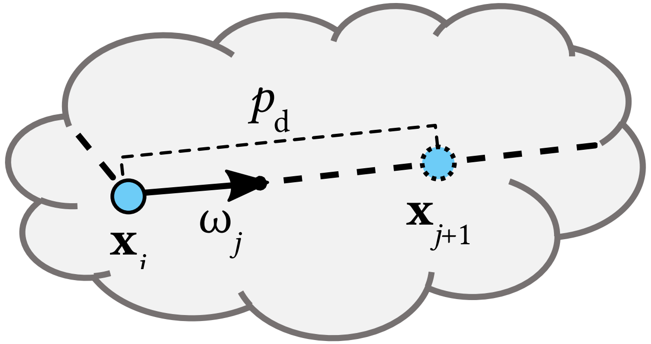

2. Distance: $p_d(d_{j+1}|\boldsymbol{x}_j, \omega_j)$

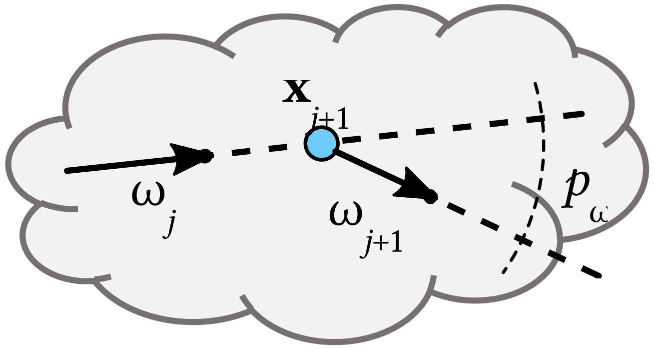

3. Direction: $p_\omega(\omega_{j+1}|\boldsymbol{x}_{j+1}, \omega_j)$

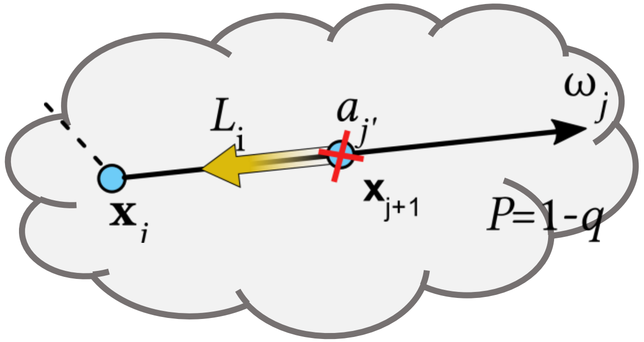

4. Termination: $P_{RR}(\boldsymbol{x}_j, \omega_{j-1})$

thrive SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

VOLUMETRIC RANDOM WALK - DECISIONS

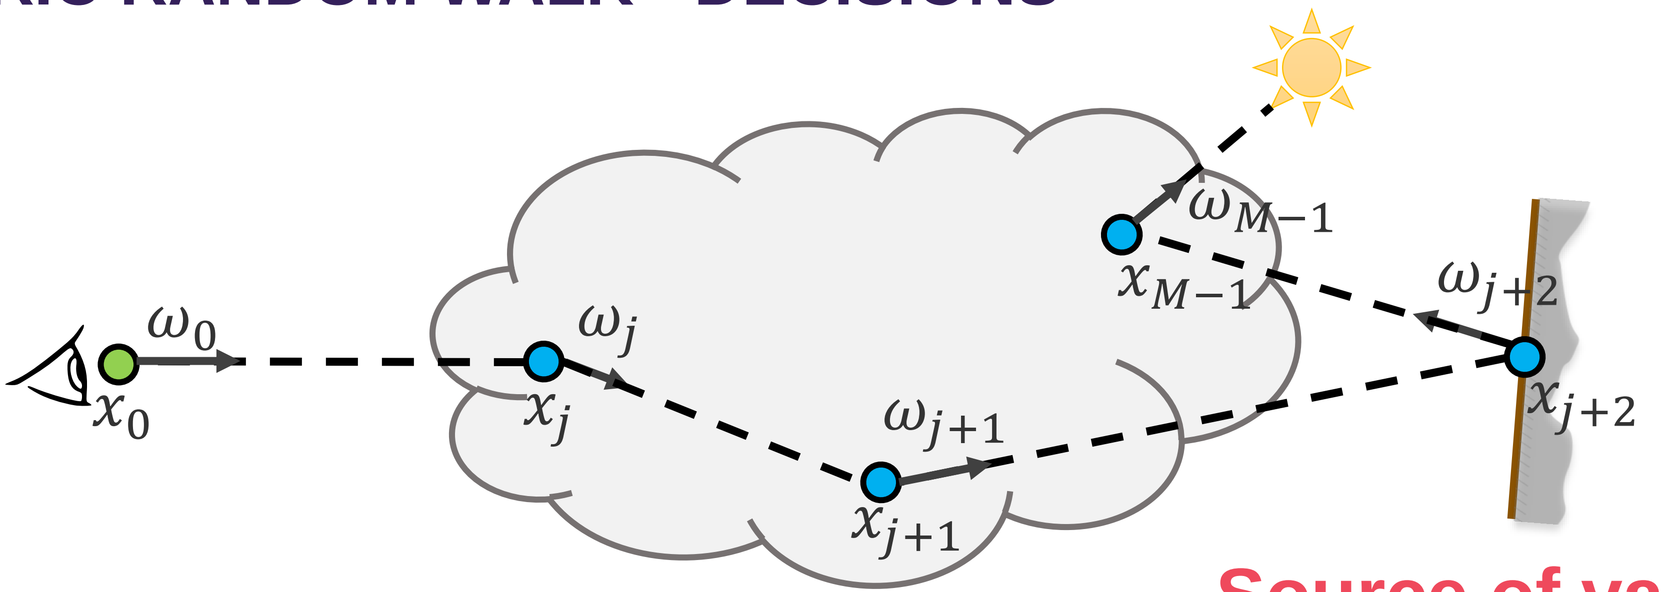

Source of variance

- Path PDF :

$$p(\boldsymbol{X}) = \boxed{\prod_{j=1}^{M-1} \underbrace{P_m(\dots) \cdot p_d(\dots) \cdot p_\omega(\dots) \cdot \left(1 - P_{RR}(\dots)\right)}_{\text{Path segment PDF}}}$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VOLUME RENDERING EQUATION

- Incident radiance (volume):

$$L(x,\omega) = \underbrace{T(\dots)\cdot L_o(\dots)}_{\text{Surface contribution}} + \underbrace{\int T(\dots)\cdot \sigma_s(\dots)\cdot \boxed{L_i(\dots)}\mathrm{d}d}_{\text{Volume contribution}}$$

- In-scattered radiance:

$$\boxed{L_i(\dots) = \int f(\dots)\cdot L(\dots)\mathrm{d}\omega'}$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VOLUME RENDERING EQUATION

- Incident radiance (volume):

Transmittance

$$L(x, \omega) = \boxed{T(\dots)} \cdot L_o(\dots) + \int \boxed{T(\dots) \cdot \sigma_s(\dots)} \cdot L_i(\dots)\mathrm{d}d$$

Known Local Quantities

- In-scattered radiance:

$$L_i(\dots) = \int \boxed{f(\dots)} \cdot L(\dots)\mathrm{d}\omega'$$

Phase function

VOLUME RENDERING EQUATION

surface radiance

In-scattered radiance

- Incident radiance (volume):

$$L(x, \omega) = T(\ldots) \cdot \boxed{L_o(\ldots)} + \int T(\ldots) \cdot \sigma_s(\ldots) \cdot \boxed{L_i(\ldots)} \mathrm{d}d$$

Unknown Light Transport Quantities

- In-scattered radiance:

$$L_i(\ldots) = \int f(\ldots) \cdot \boxed{L(\ldots)} \mathrm{d}\omega'$$

Incident radiance

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

STANDARD SAMPLING

1+2 Scatter and Distance:

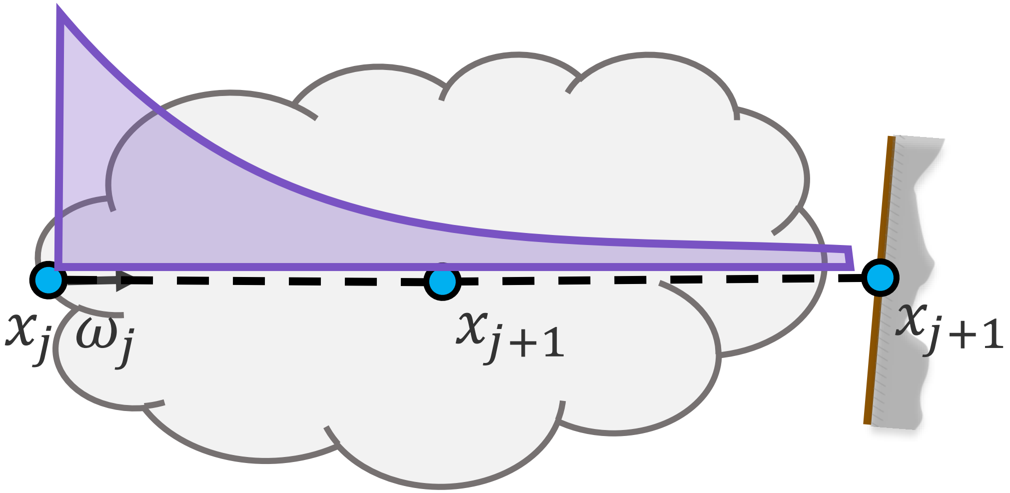

- Based on the transmittance

3 Direction:

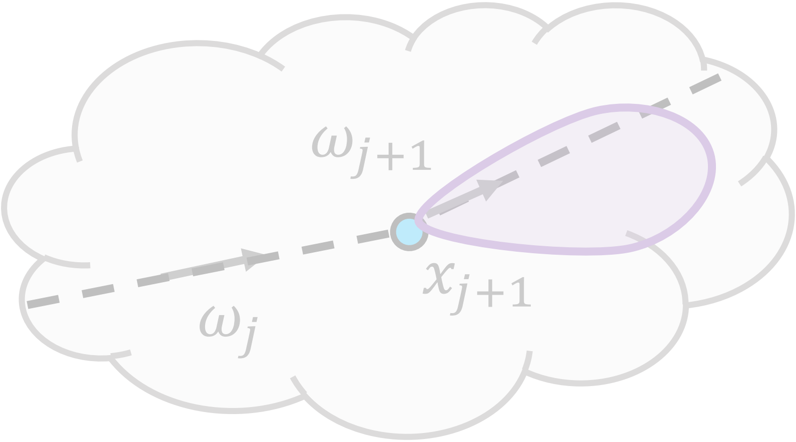

- Based on the phase function

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

STANDARD SAMPLING

4 Termination:

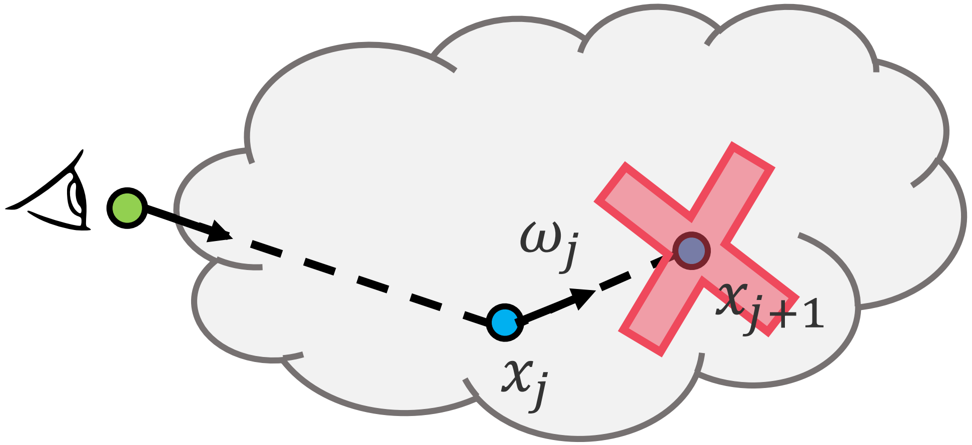

• Based on local albedo or throughput

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

CHALLENGES FOR VOLUME SAMPLING

Why do we need volumetric path guiding?

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

LIGHT SHAFTS

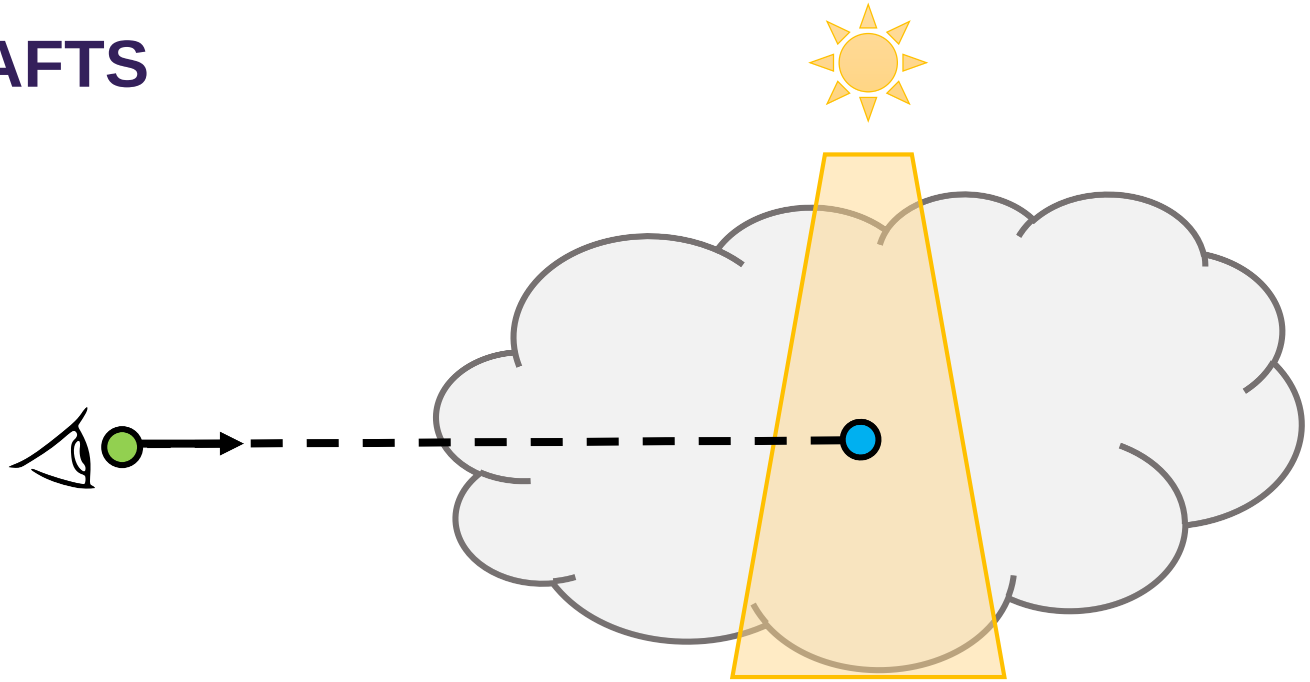

- Light shafts:
  - We need to scatter inside the light shaft.

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

LIGHT SHAFTS

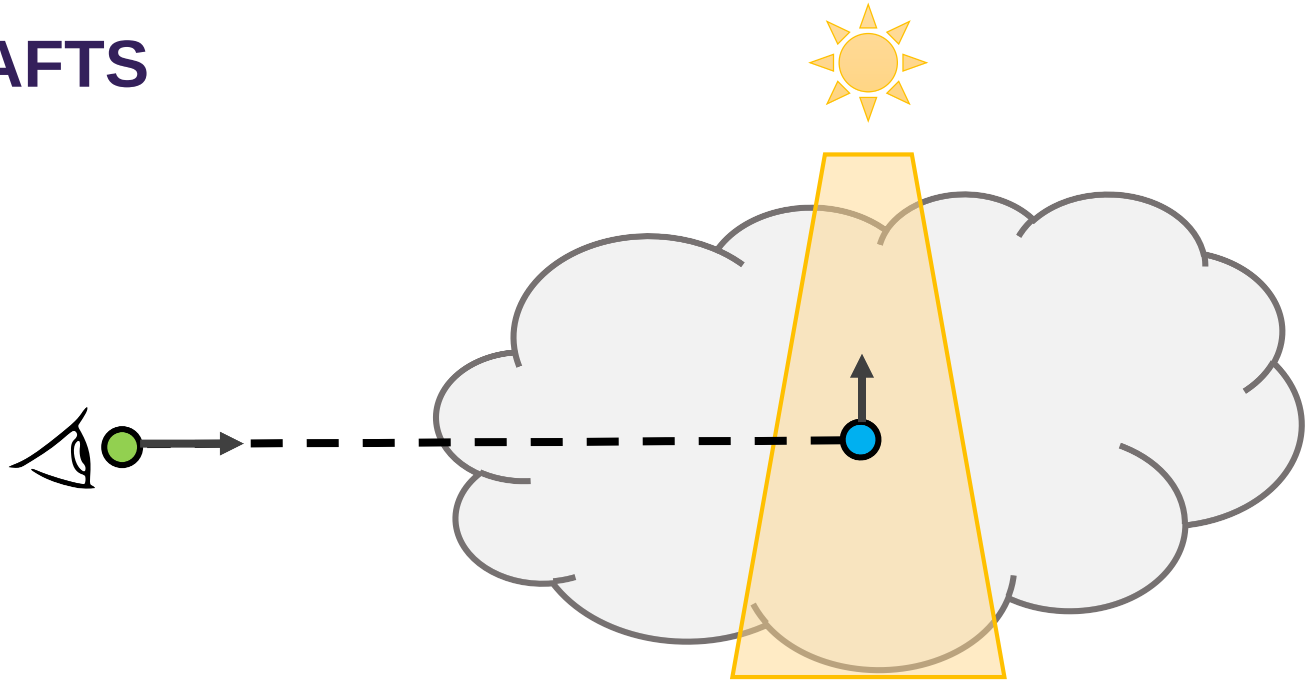

- Light shafts:
  - We need to scatter inside the light shaft.
  - We need to follow the direction of the light shaft.

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

LIGHT SHAFTS

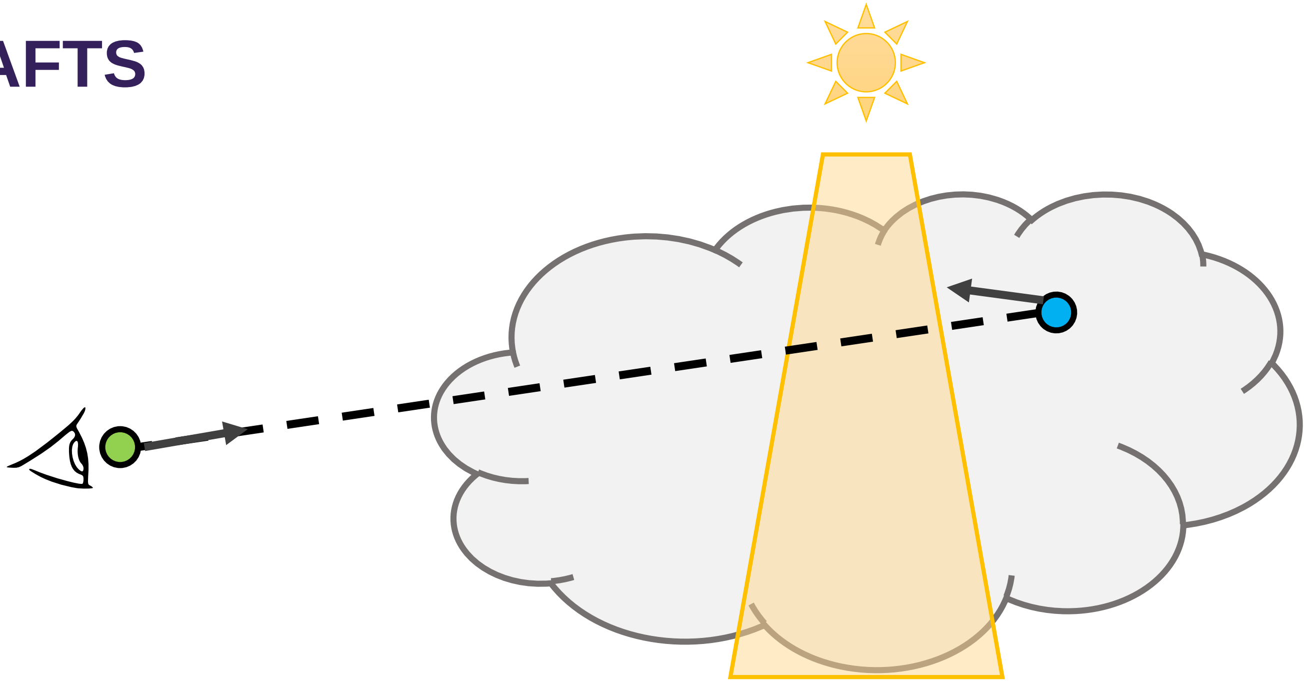

- Light shafts:
  - We need to scatter inside the light shaft.
  - We need to follow the direction of the light shaft.
  - We need to scatter towards the light shaft.

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

No guiding (1024 spp)

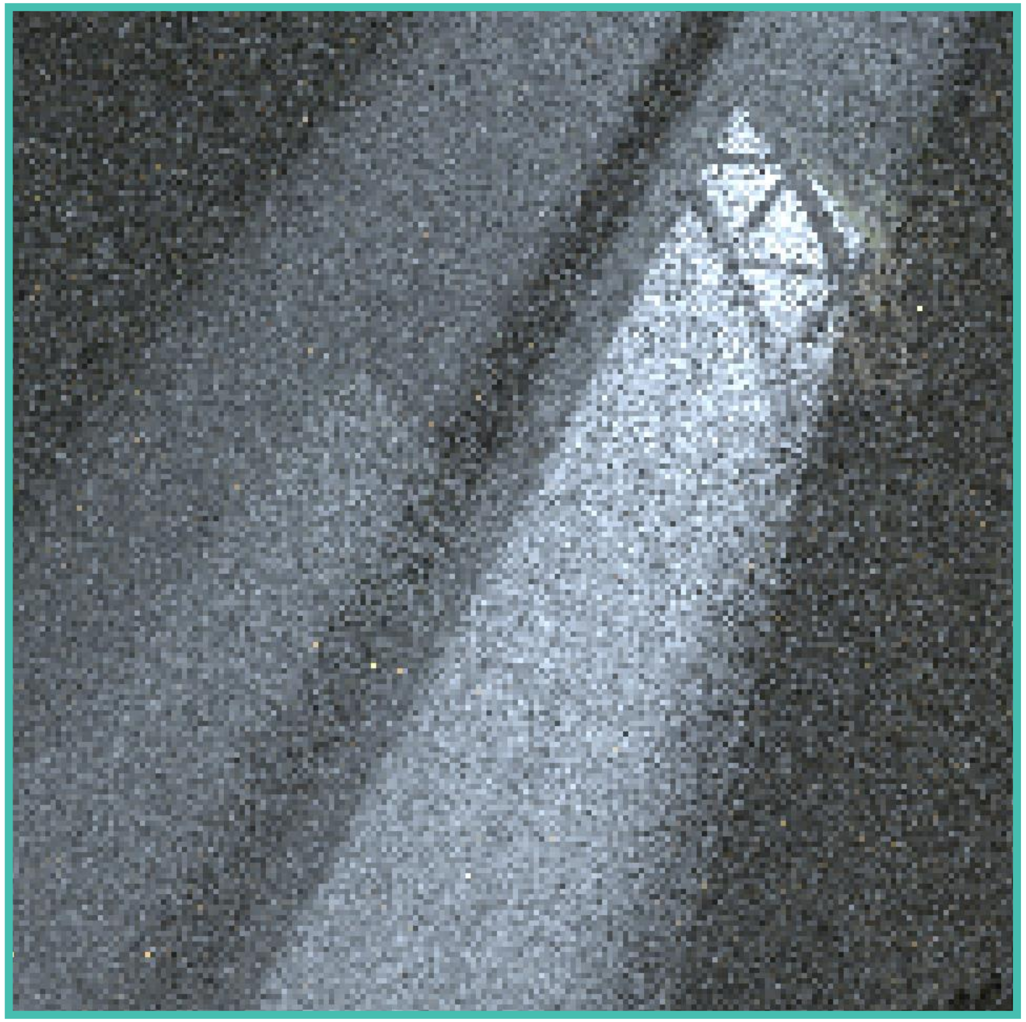

Our guiding (1024 spp)

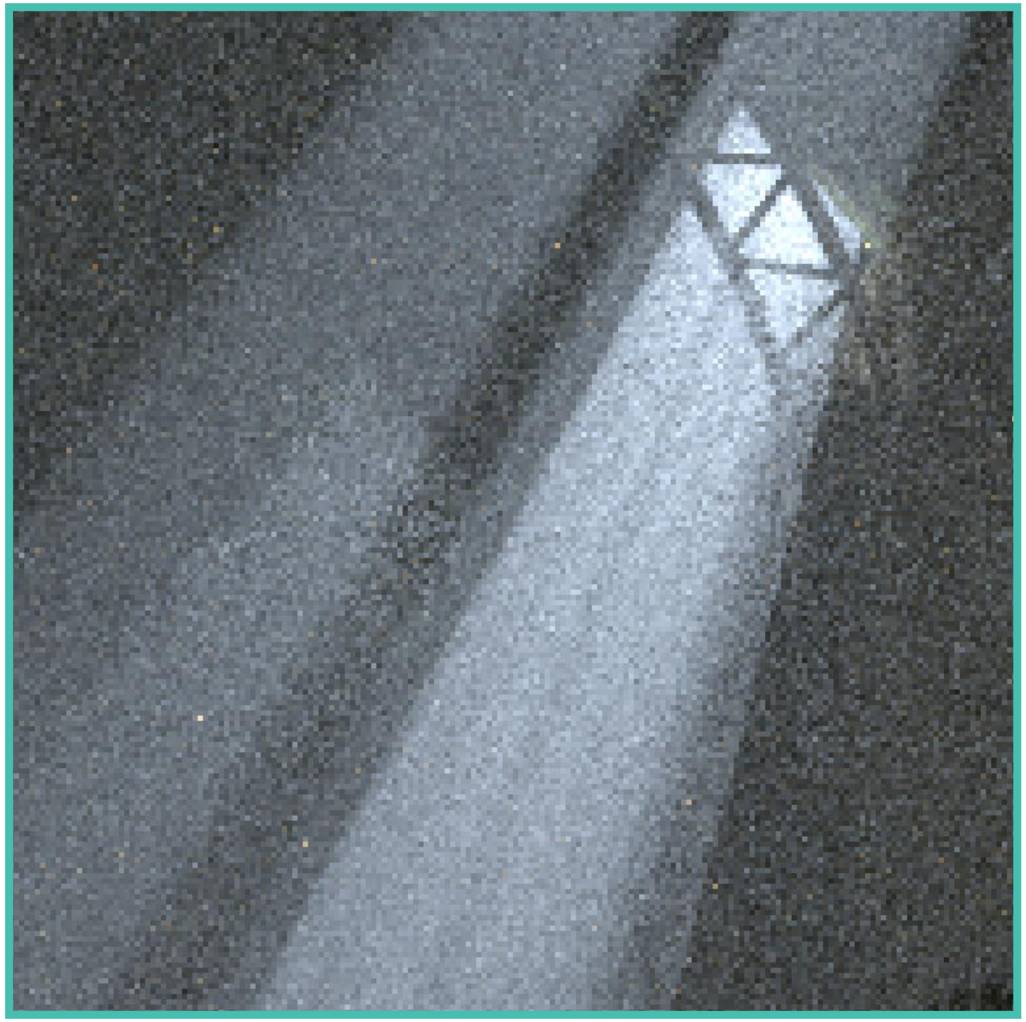

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

SUB-SURFACE-SCATTERING

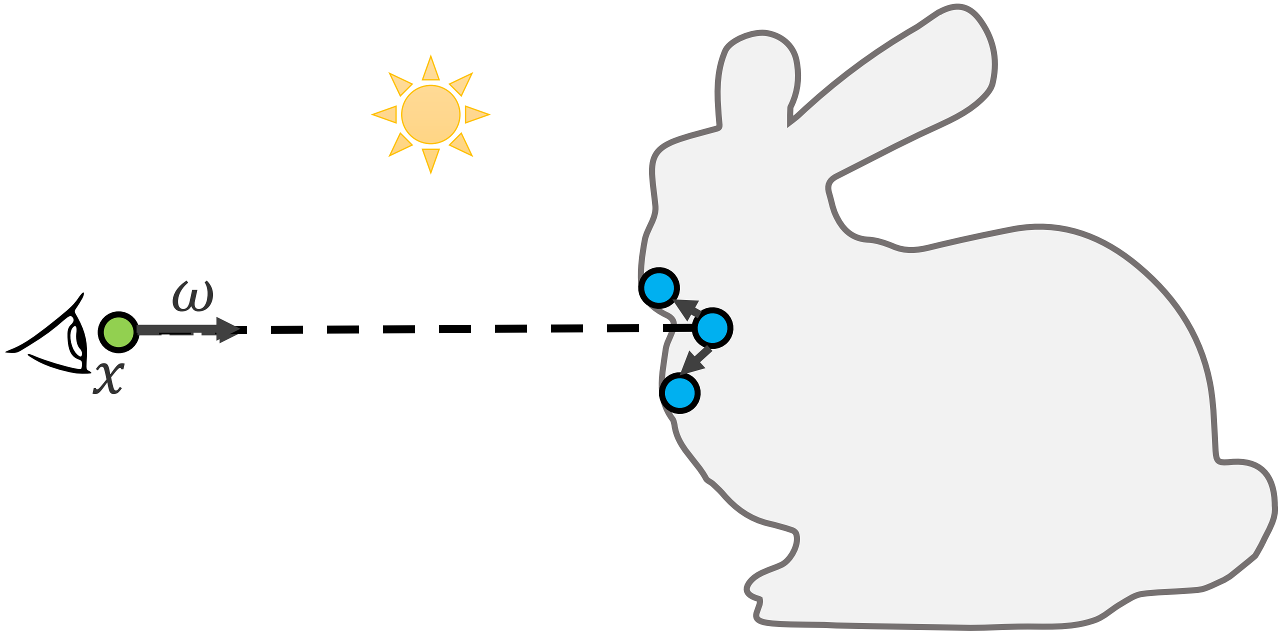

- Sub-Surface-Scattering:
  - We ‘often’ need stay close to the surface

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

SUB-SURFACE-SCATTERING

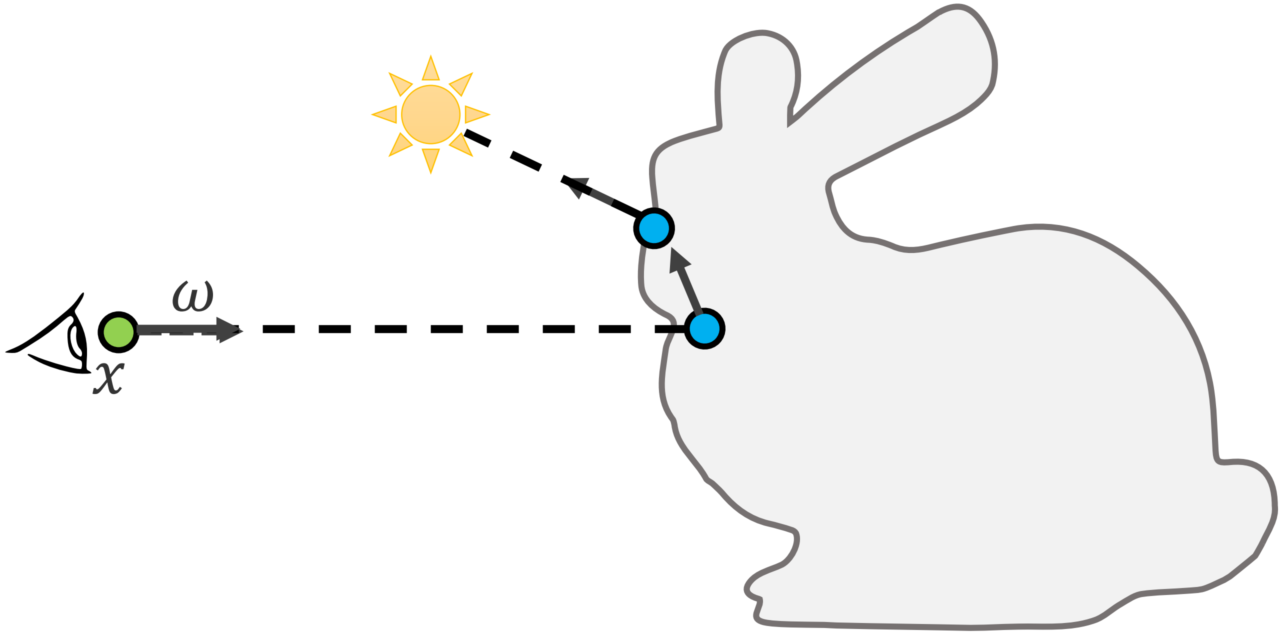

- Sub-Surface-Scattering:
  - We ‘often’ need to stay close to the surface
  - We need to leave the object with the right direction

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

No guiding (256 spp)

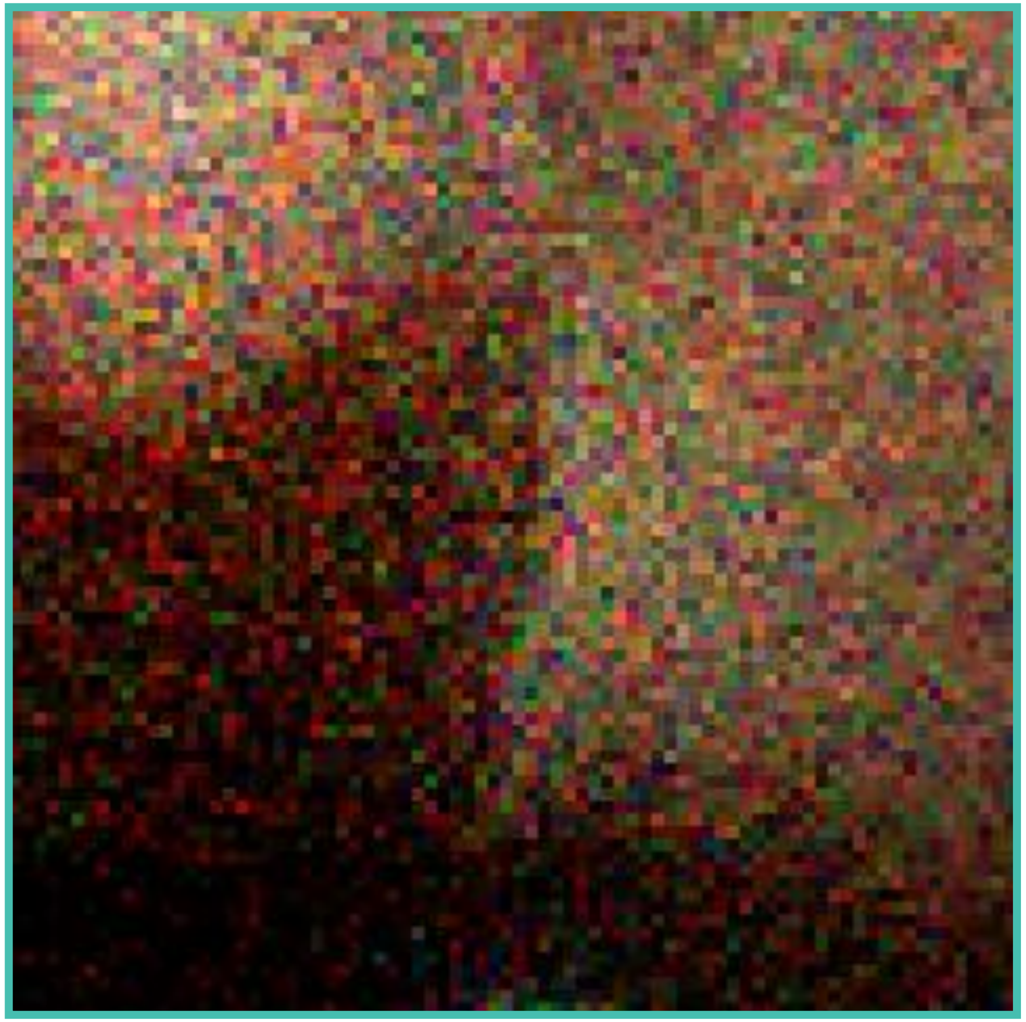

Our guiding (256 spp)

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

DENSE MEDIA

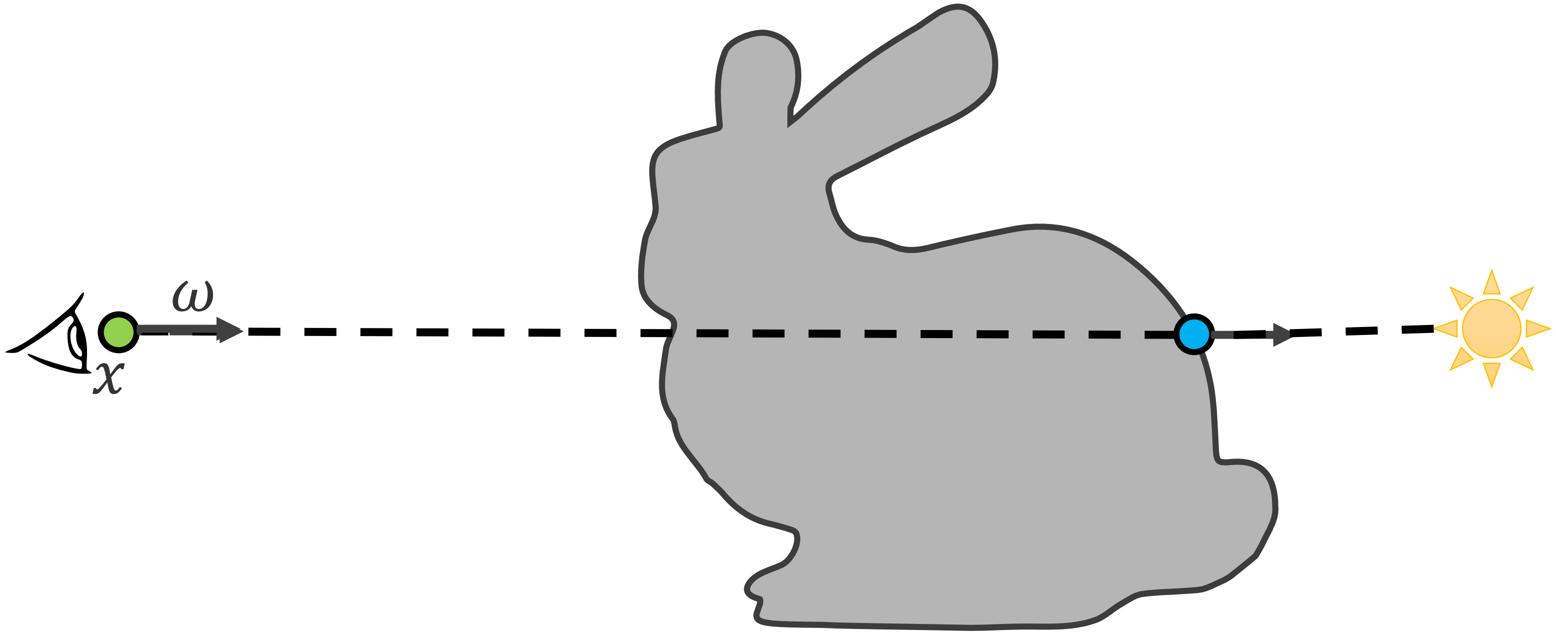

- Dense media (back illuminated):
  - We may need to **'avoid'** generating a scattering event even if the transmittance is low (e.g. strong light source behind the volume).

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

No guiding (256 spp)

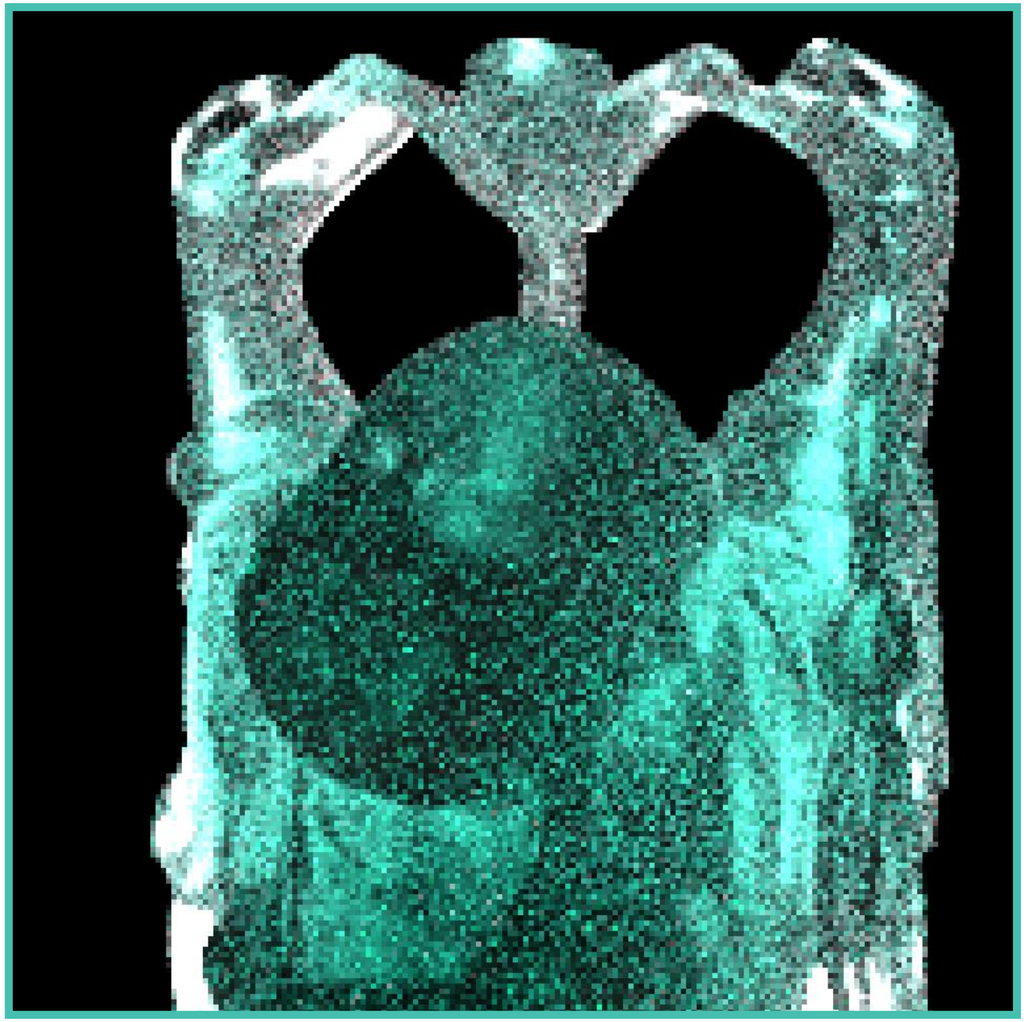

Our guiding (256 spp)

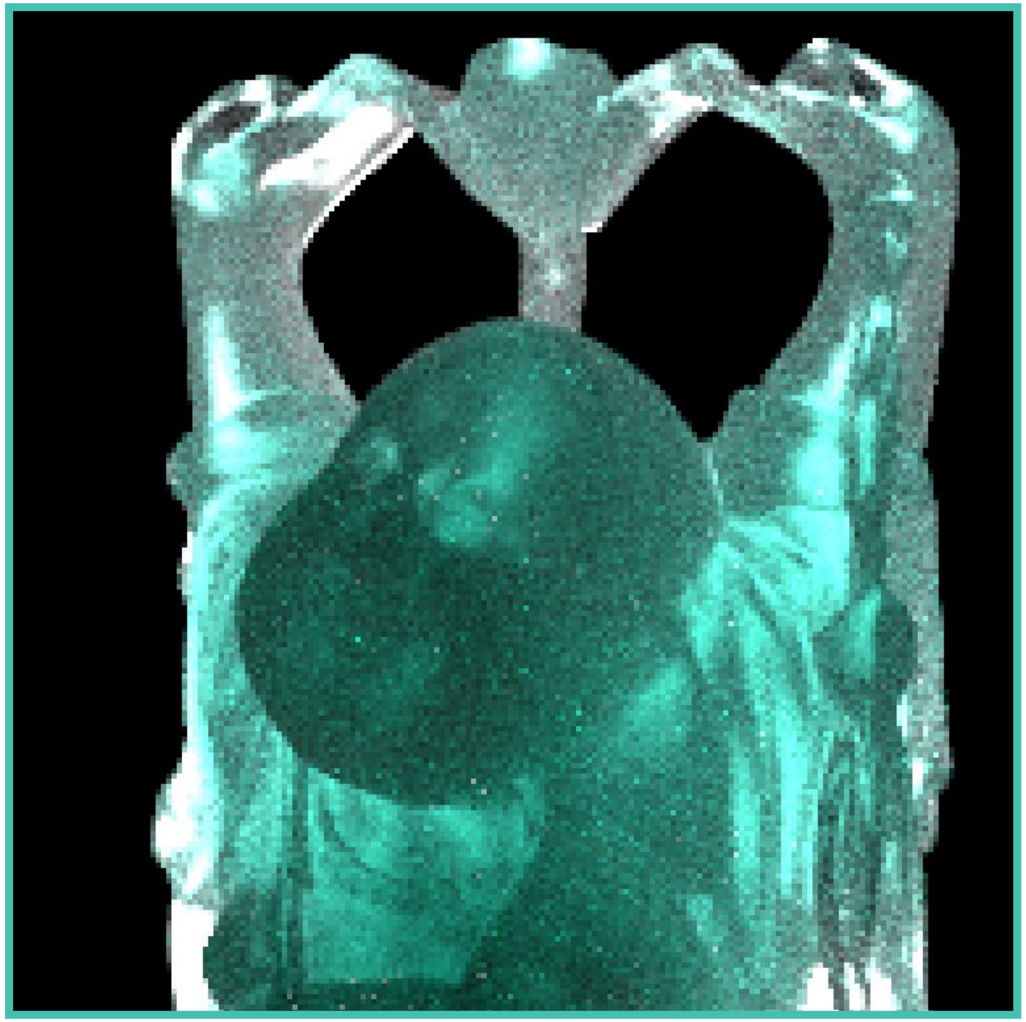

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

NON-DENSE MEDIA

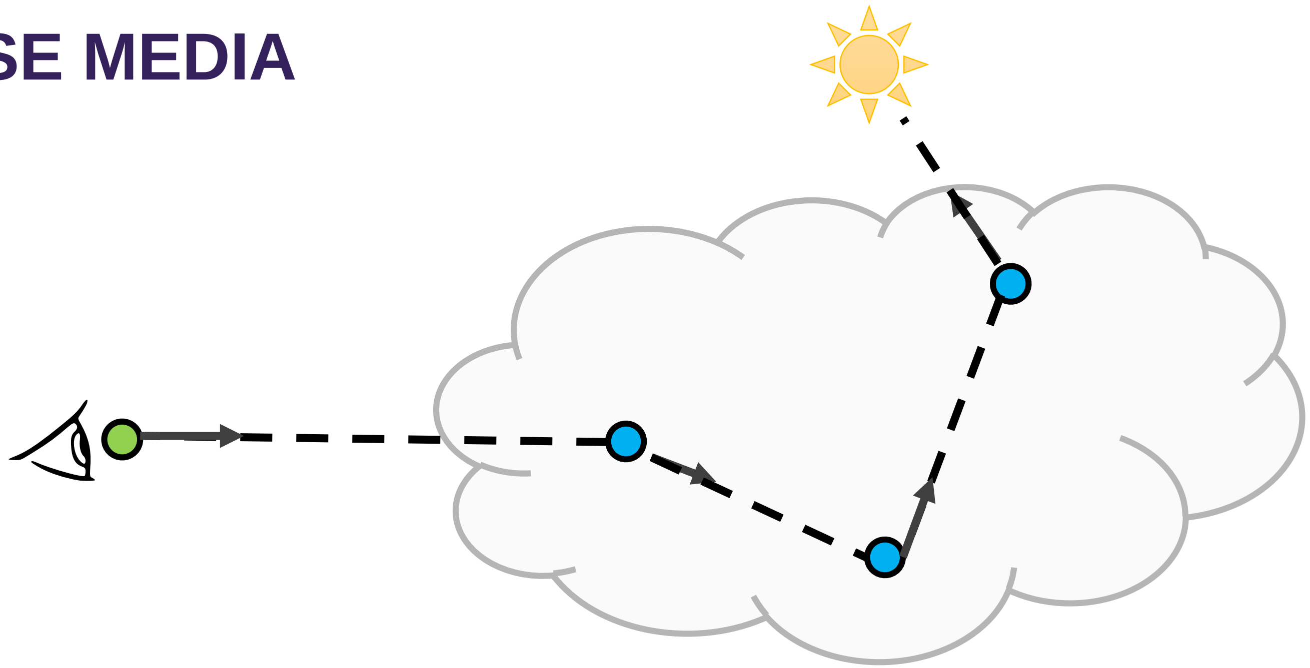

- Non-dense media (no back illumination):
  - We may need to '**force**' a scattering event even if the transmittance is high.

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

No guiding (256 spp)

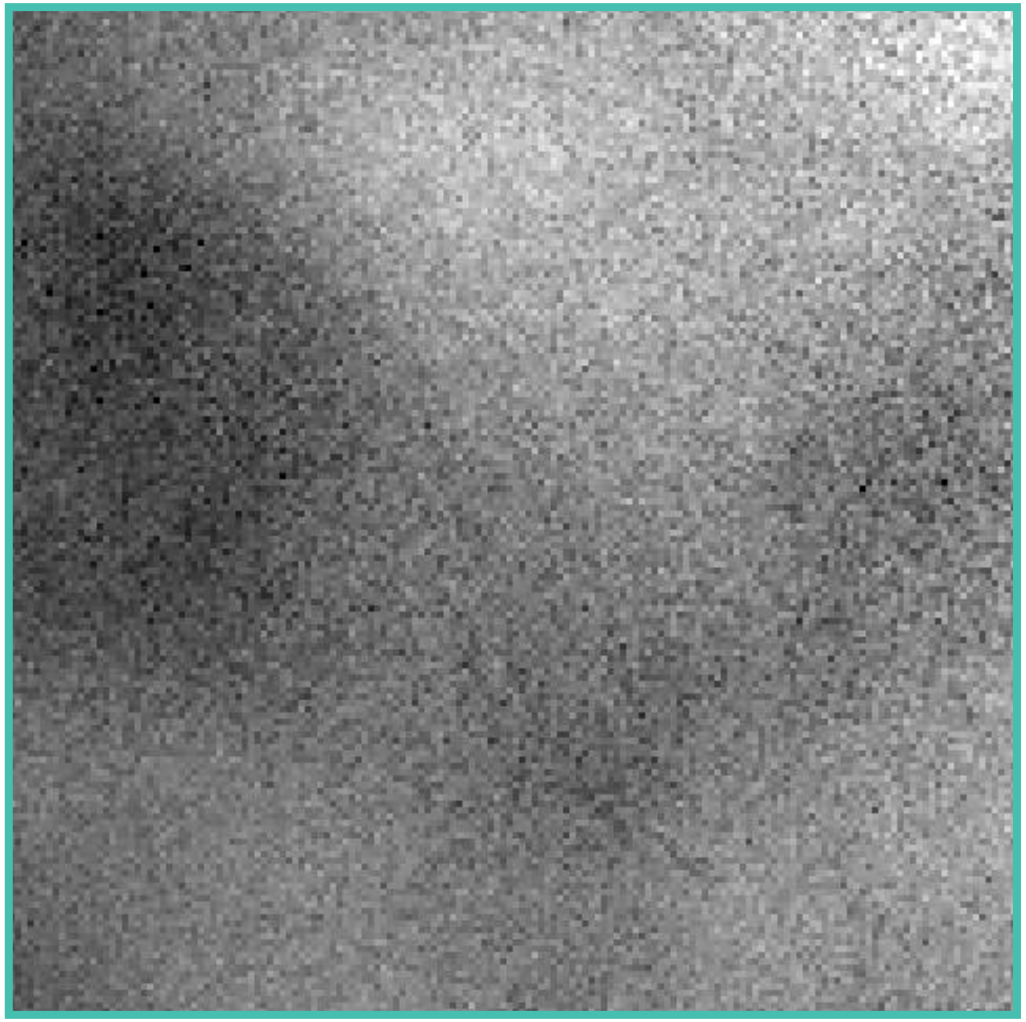

Our guiding (256 spp)

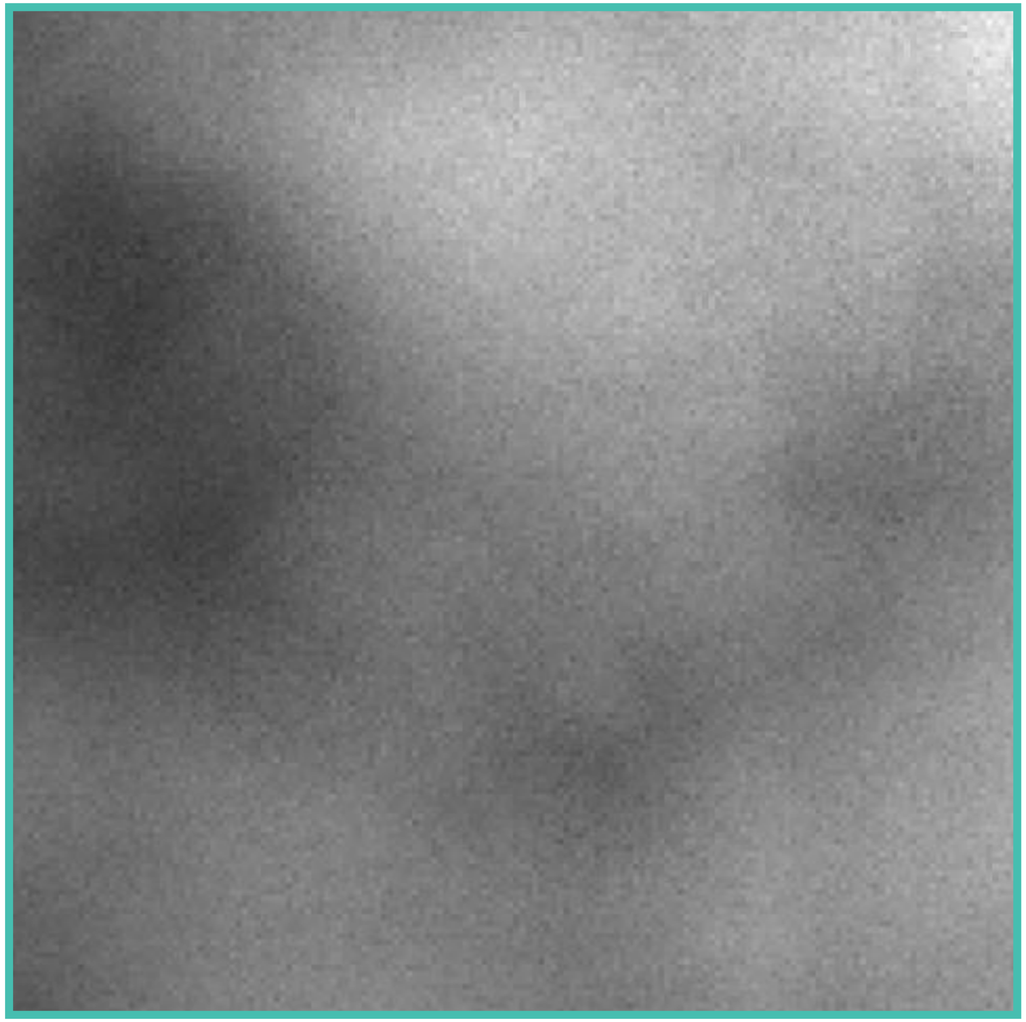

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

SPECIALIZED SOLUTIONS: SHORTCOMINGS

- Many individual solutions:
  - Equiangular Sampling: [Kulla2012]
  - Joint-Importance Sampling: [Georgiev2012]
  - Zero-Variance Dwivedi Sampling: [Krivanek2014] [Meng2016]
  - Directional (illumination-based) guiding: [Pegoraro2008] [Bashford2012]
- Only considering sub-sets or special cases:
  - Surface-bounded volumes
  - Homogenous or isotropic volumes
  - Single scattering
- **None of the current methods importance samples the full volumetric light transport!**

(a) *Path traced light transport in clouds*

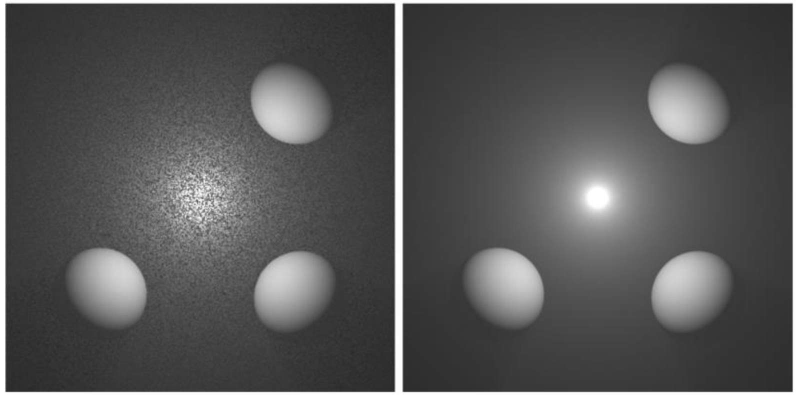

(a) Distance Sampling (b) Equi-angular Sampling

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

SPECIALIZED SOLUTIONS: SHORTCOMINGS

- Many individual solutions:
  - Equiangular Sampling: [Kulla2012]
  - Joint-Importance Sampling: [Georgiev2012]
  - Zero-Variance Dwivedi Sampling: [Krivanek2014] [Meng2016]
  - Directional (illumination-based) guiding: [Pegoraro2008] [Bashford2012]
- Only considering sub-sets or special cases:
  - Surface-bounded volumes
  - Homogenous or isotropic volumes
  - Single scattering
- **None of the current methods importance samples the full volumetric light transport!**

(a) *Path traced light transport in clouds*

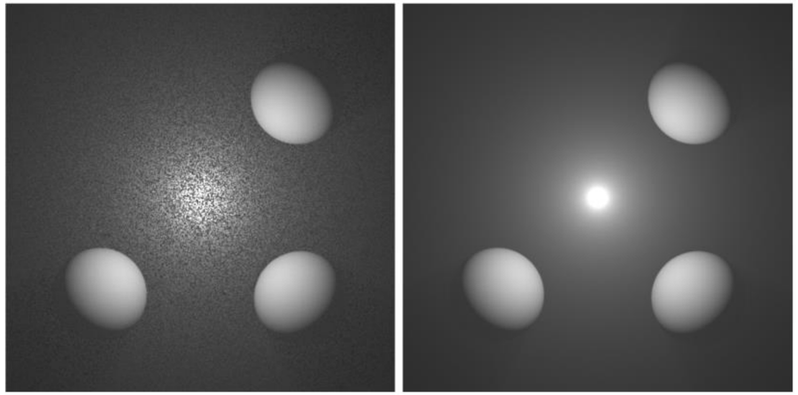

(a) Distance Sampling (b) Equi-angular Sampling

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

SPECIALIZED SOLUTIONS: SHORTCOMINGS

- Many individual solutions:
  - Equiangular Sampling: [Kulla2012]
  - Joint-Importance Sampling: [Georgiev2012]
  - Zero-Variance Dwivedi Sampling: [Krivanek2014] [Meng2016]
  - Directional (illumination-based) guiding: [Pegoraro2008] [Bashford2012]
- Only considering sub-sets or special cases:
  - Surface-bounded volumes
  - Homogenous or isotropic volumes
  - Single scattering
- **None of the current methods importance samples the full volumetric light transport!**

(a) *Path traced light transport in clouds*

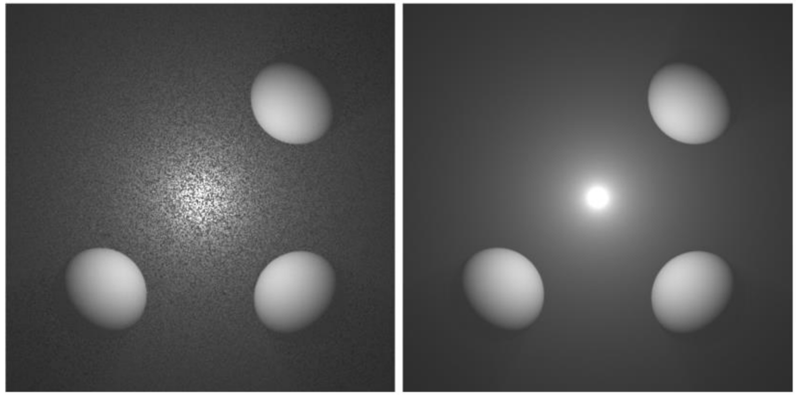

(a) Distance Sampling (b) Equi-angular Sampling

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

ZERO-VARIANCE-BASED VOLUMETRIC PATH GUIDING

TECH TALK: TUE: 30TH JULY
TIME: 9:00 AM
ROOM: 152

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

ZV-BASED VOLUMETRIC PATH GUIDING: GOALS

- Leverage recent success of local surface guiding methods: [Vorba2014], [Herholz2016], [Mueller2017]
  - Extend the concept to volumes
- Consider the **complete** volumetric light transport:
  - No prior assumptions or special cases
  - Guide based on the **optimal** zero-variance decisions [Hoogenboom 2008]
- Replace unknown quantities by estimates:

$$L(\boldsymbol{x},\omega) = \tilde{L}(\boldsymbol{x},\omega) \qquad L_i(\boldsymbol{x},\omega) = \tilde{L}_i(\boldsymbol{x},\omega)$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

ZV-BASED VOLUMETRIC PATH GUIDING: CONTRIBUTIONS

- Guiding **all** local sampling decisions:
  - 1+2 Guided product distance sampling:

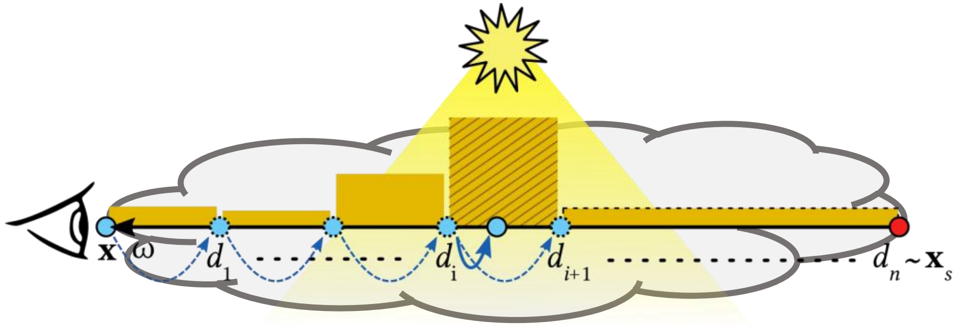

  - 3 Guided product directional sampling:

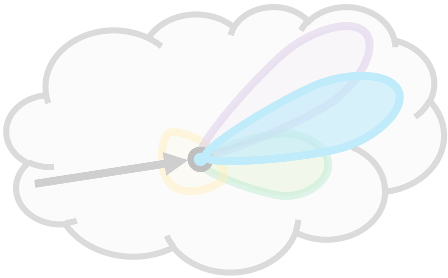

  - 4 Guided Russian roulette and Splitting:

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

ZV-BASED VOLUMETRIC PATH GUIDING: CONTRIBUTIONS

- Guiding **all** local sampling decisions:
  - 1+2 Guided product distance sampling:

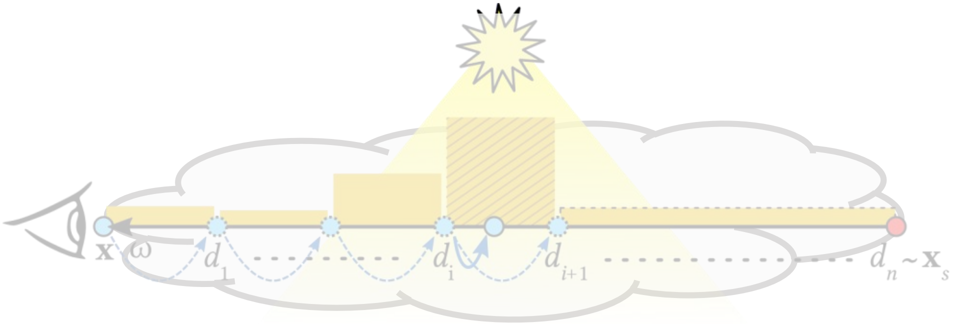

  - 3 Guided product directional sampling:

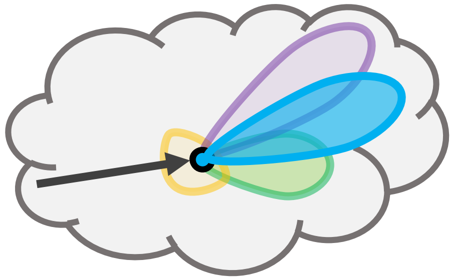

  - 4 Guided Russian roulette and Splitting:

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

ZV-BASED VOLUMETRIC PATH GUIDING: CONTRIBUTIONS

- Guiding **all** local sampling decisions:
  - 1+2 Guided product distance sampling:

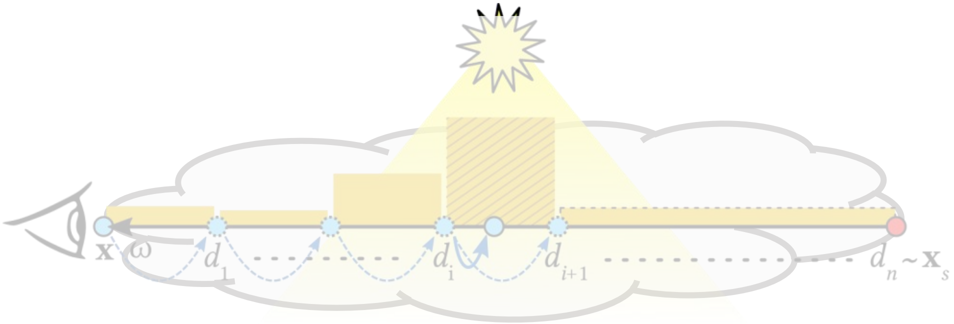

  - 3 Guided product directional sampling:

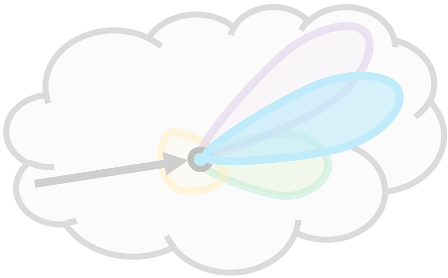

  - 4 Guided Russian roulette and Splitting:

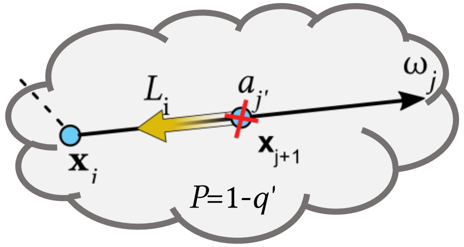

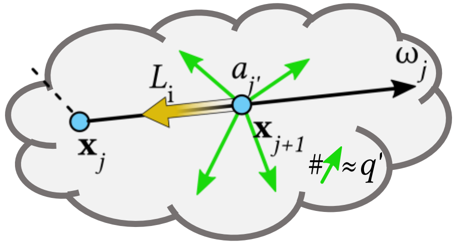

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VOLUME RADIANCE ESTIMATES

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

VOLUME RADIANCE ESTIMATES

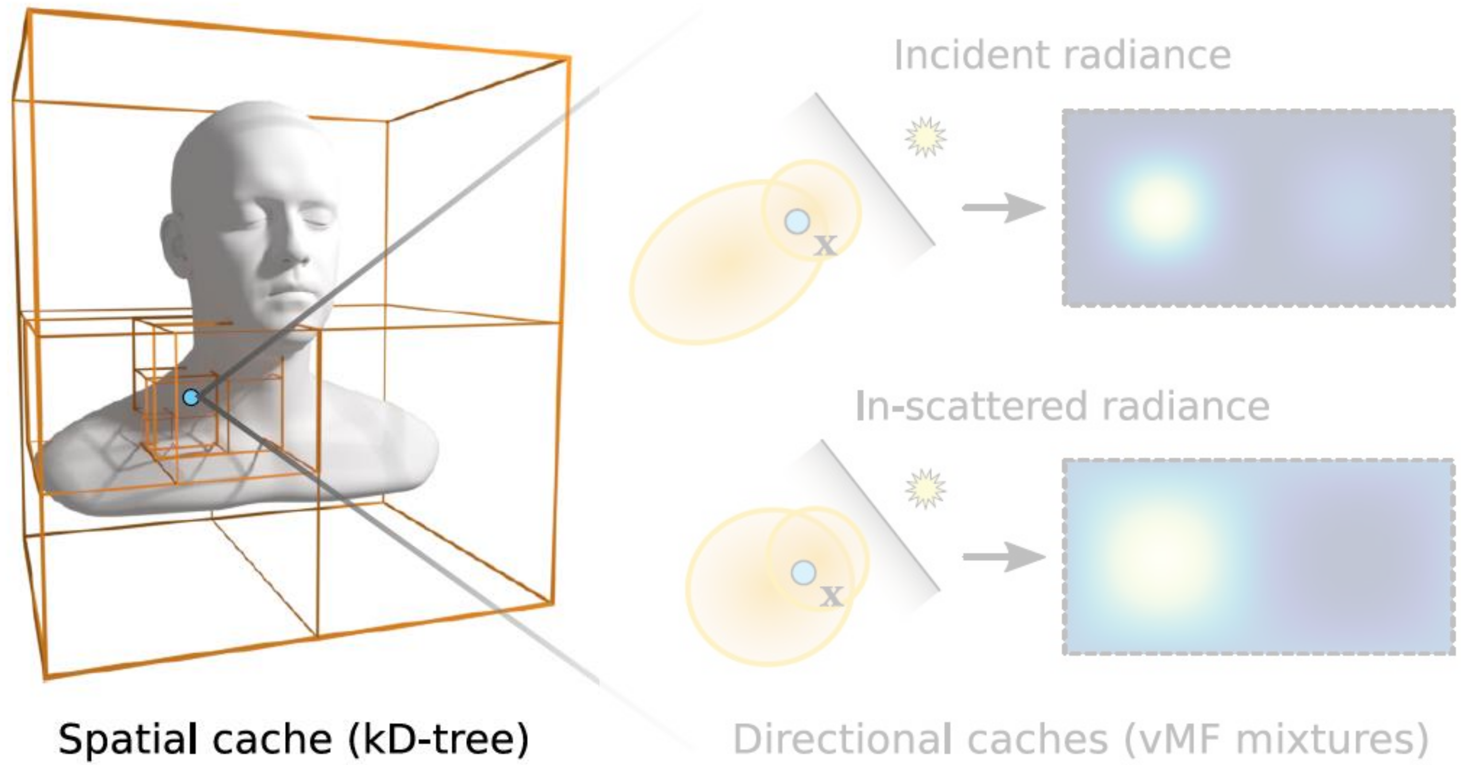

- Spatial caches via BSP-tree: max. 2K photons per node:
  - Similar 3D structure as PPG [Mueller2017]

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

VOLUME RADIANCE ESTIMATES

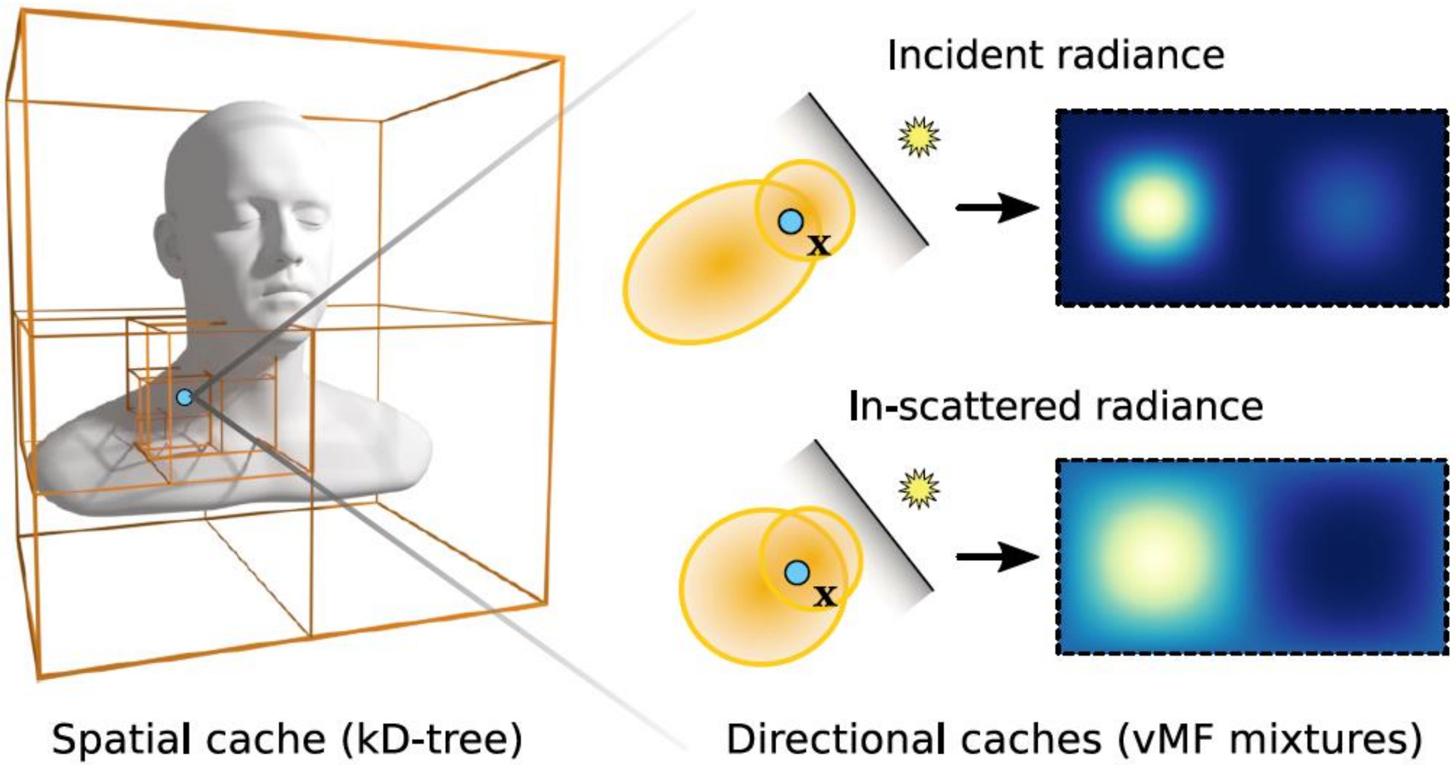

- Pre-processing step to fit estimates from photons (50M):
  - EM-fitting of von Mises-Fisher mixtures (similar to [Vorba2014]'s GMMs)

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VON MISES-FISHER MIXTURE MODELS

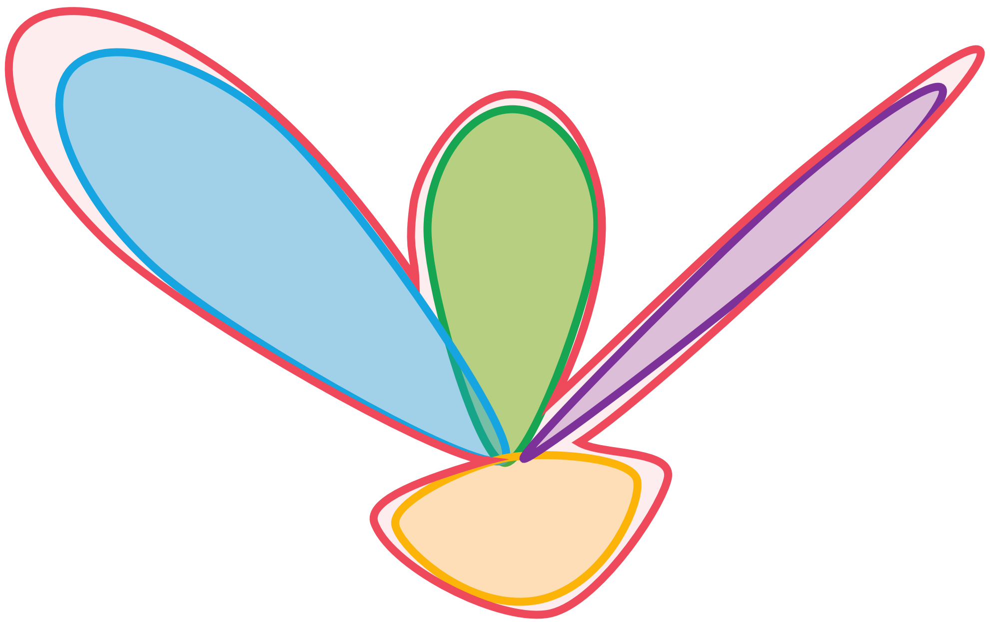

- Spherical Distribution:

$$V(\omega|\Theta) = \sum^{K} \pi_i v(\omega|\mu_i, \kappa_i)$$

- Features (closed-form):
  - Sampling
  - Convolution
  - Product

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

RADIANCE ESTIMATES

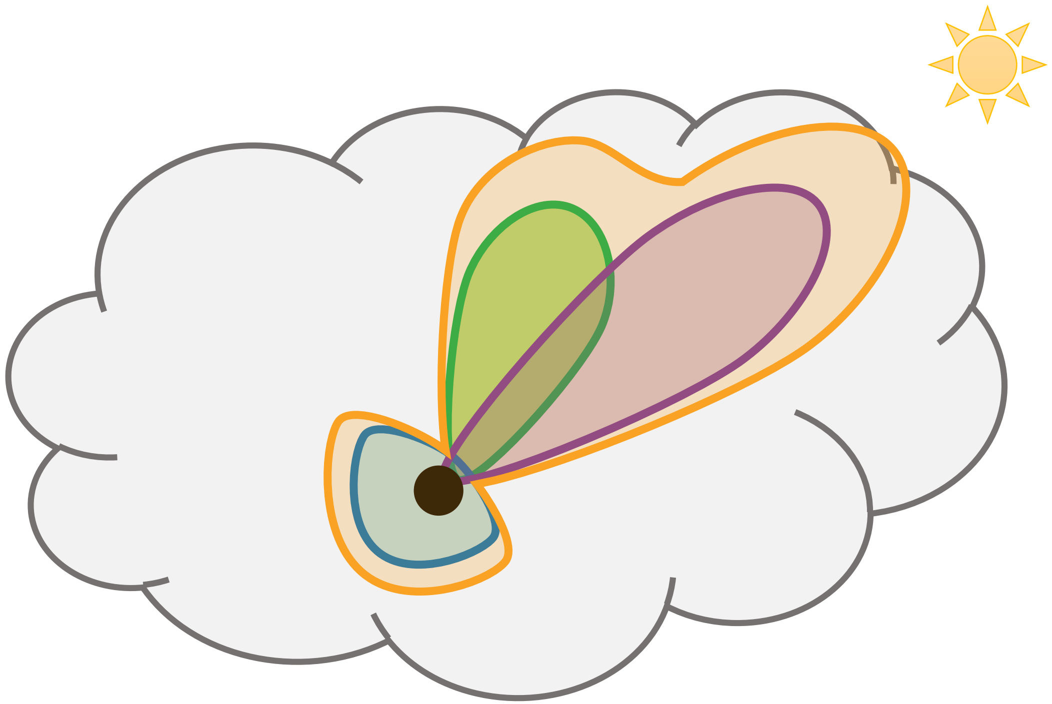

- Incident Radiance Distribution

- In-Scattered Radiance Distribution

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

INCIDENT RAD. TO IN-SCATTERED RAD. TRANSFORMATION

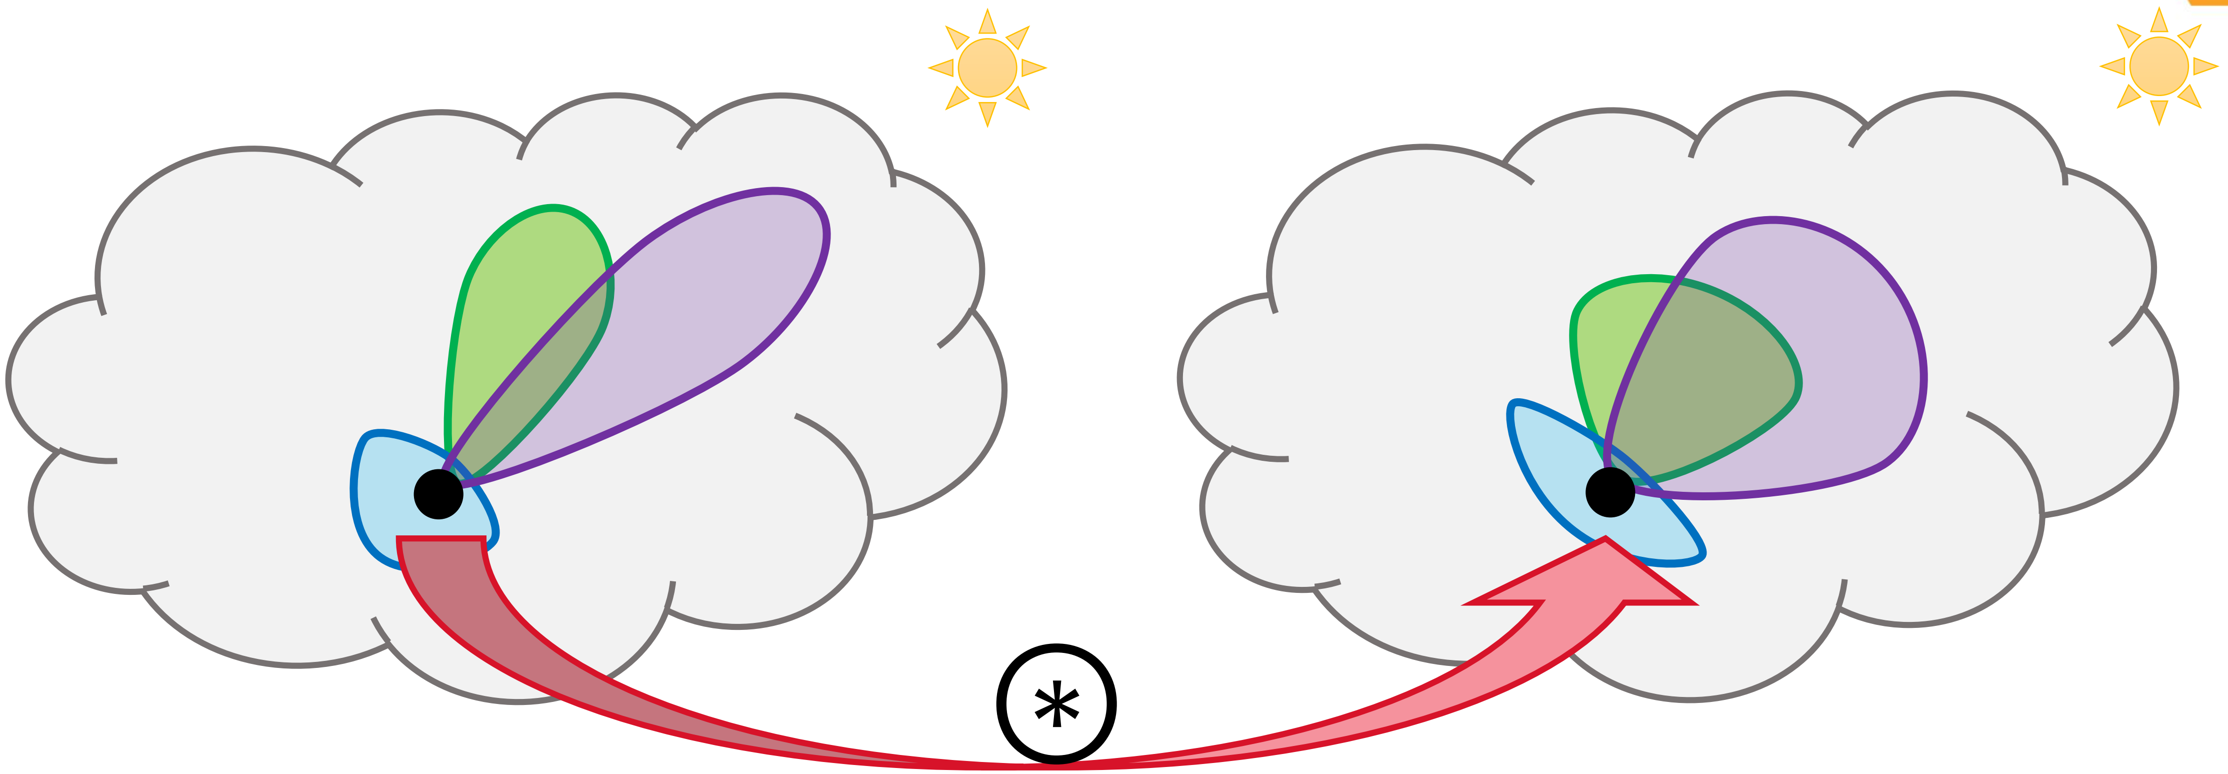

- Convolution between incident radiance L and the phase function f

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

INCIDENT RADIANCE ESTIMATES

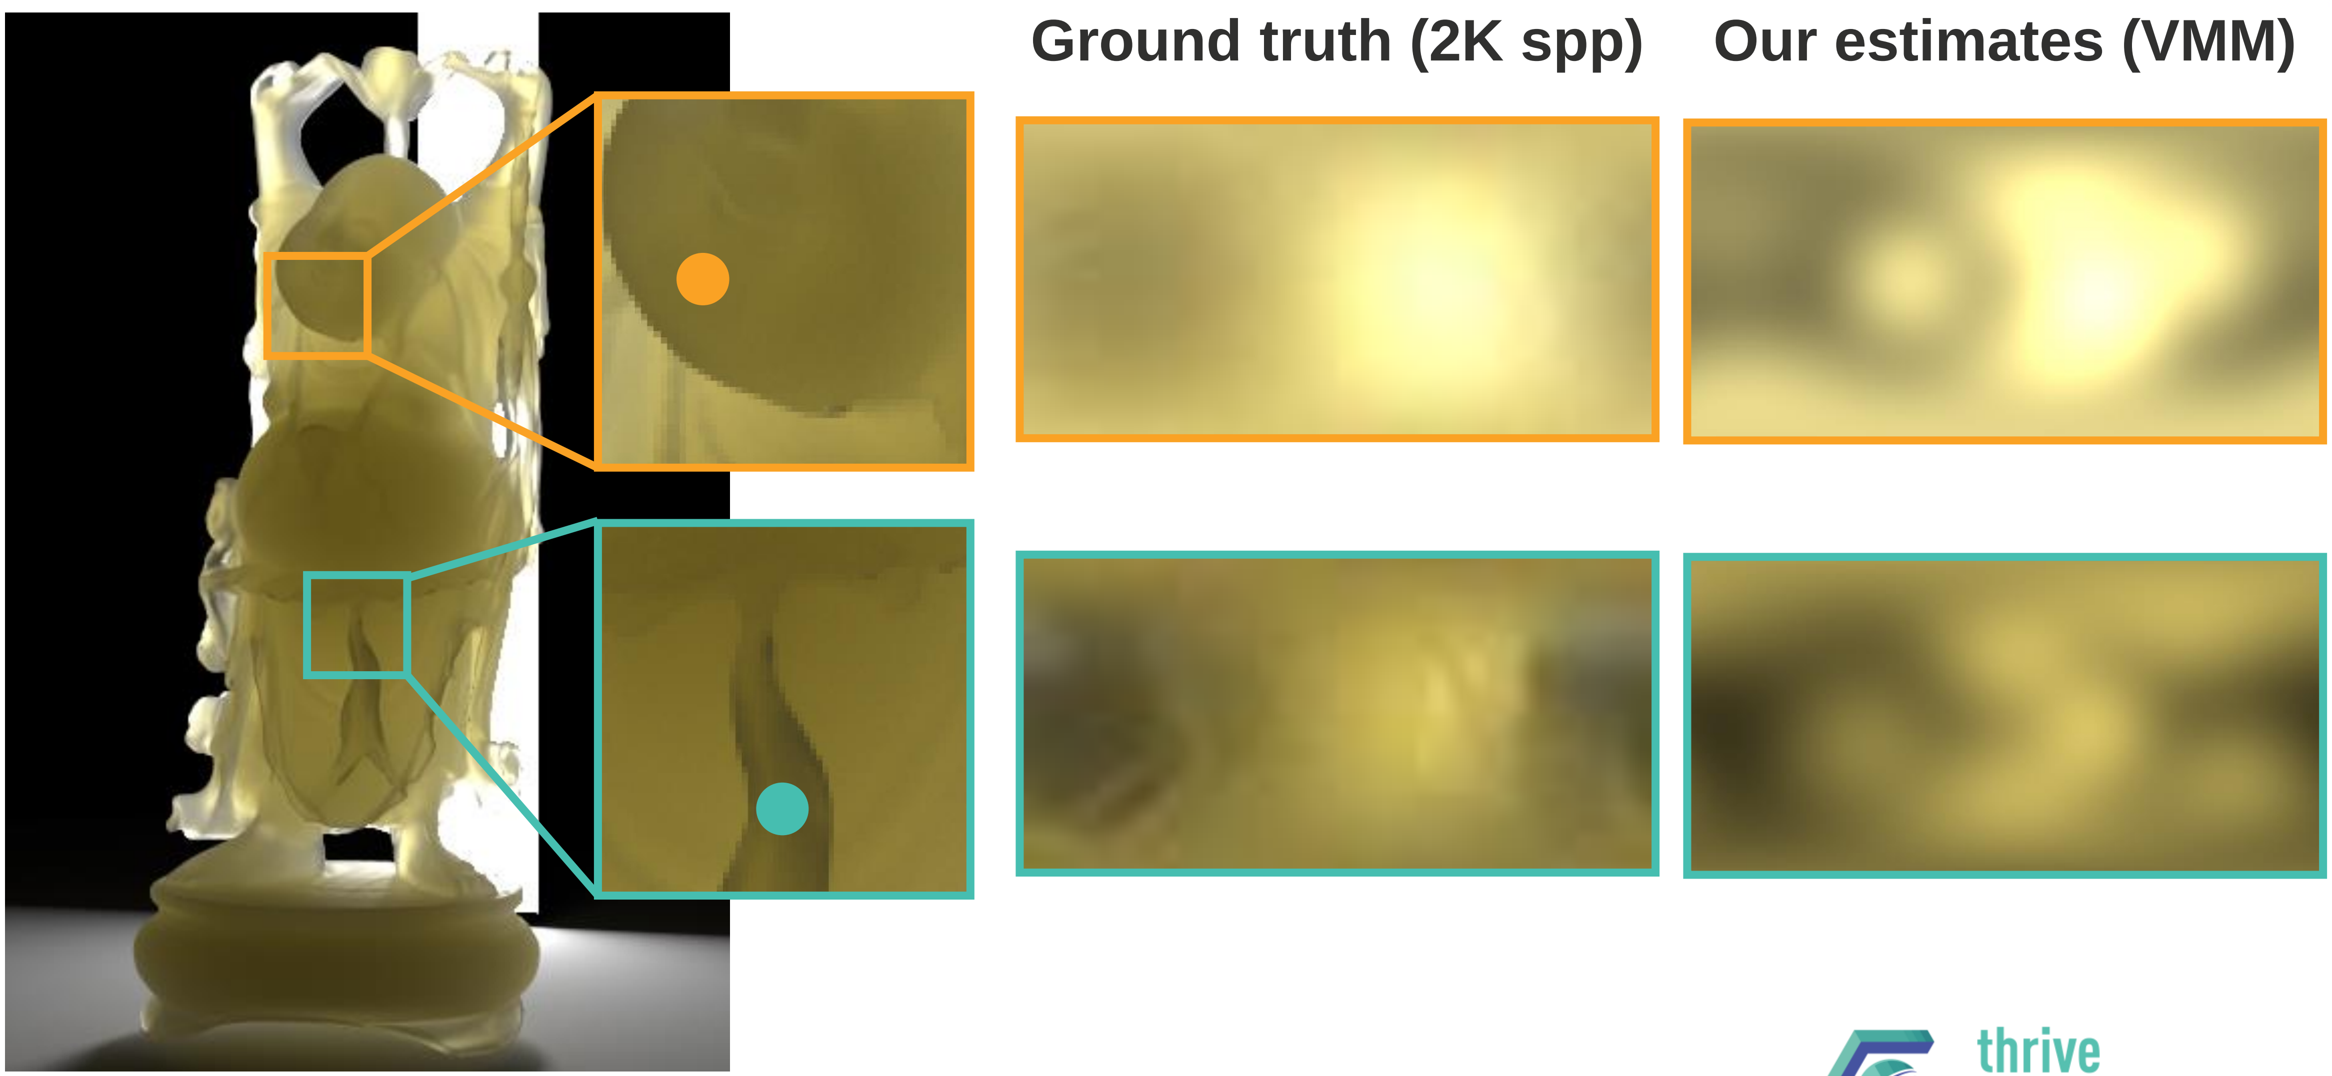

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

IN-SCATTERED RADIANCE ESTIMATES

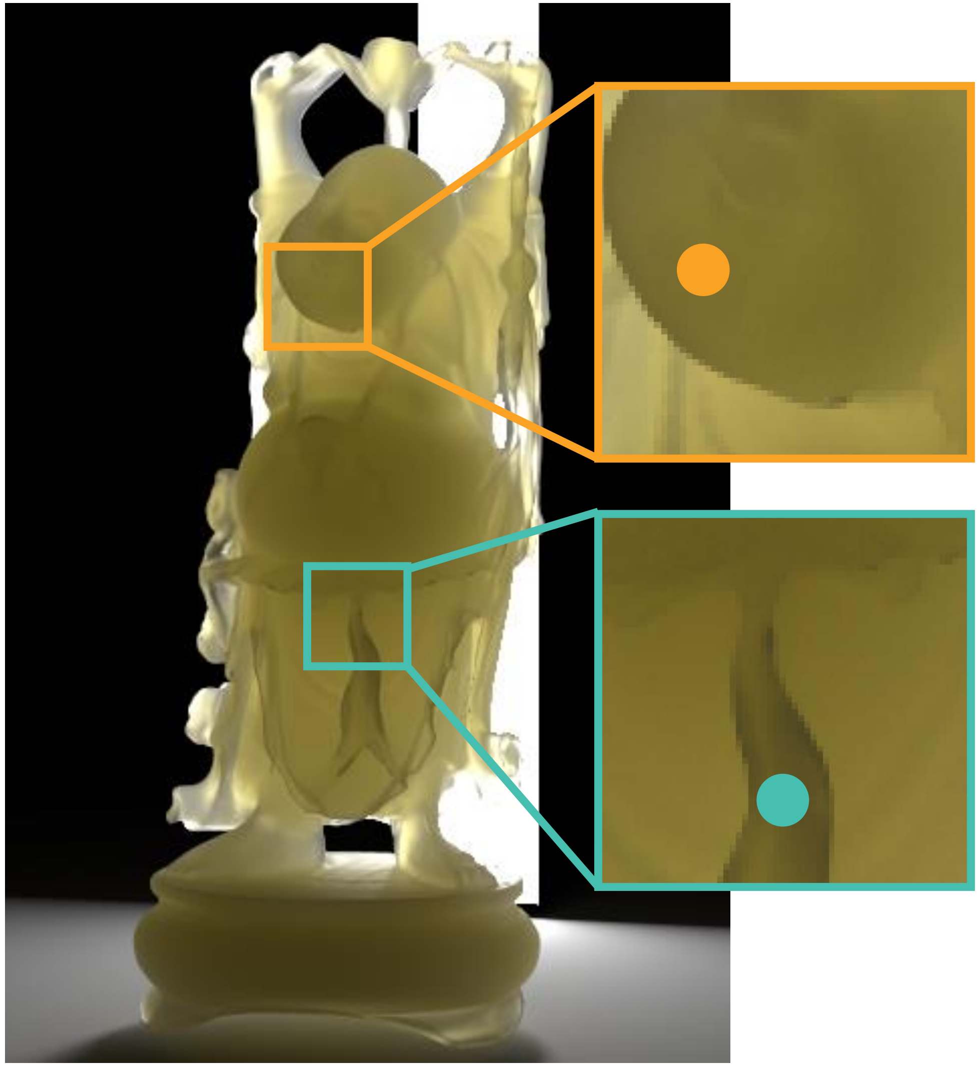

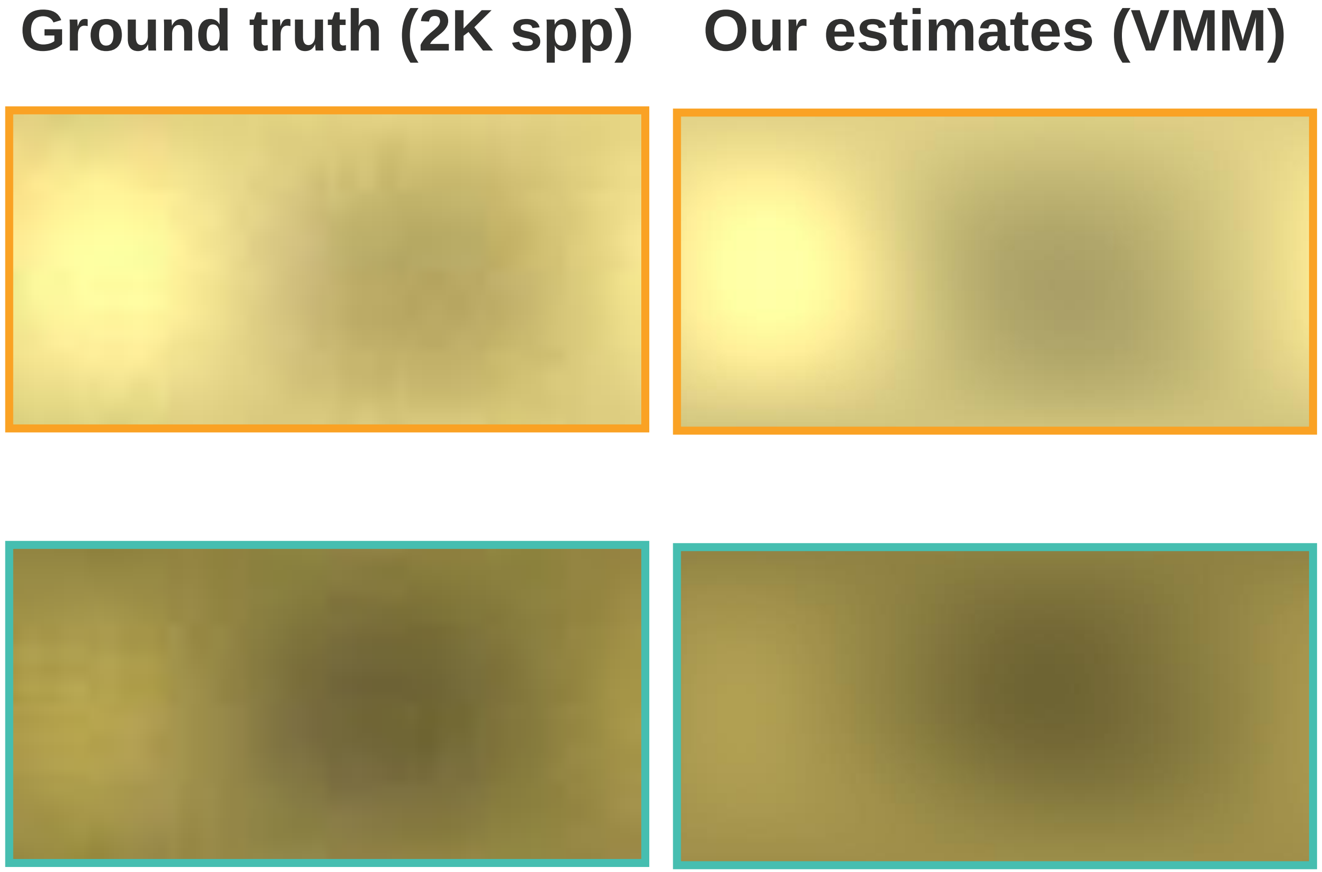

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED SAMPLING DECISIONS

DISTANCE SAMPLING

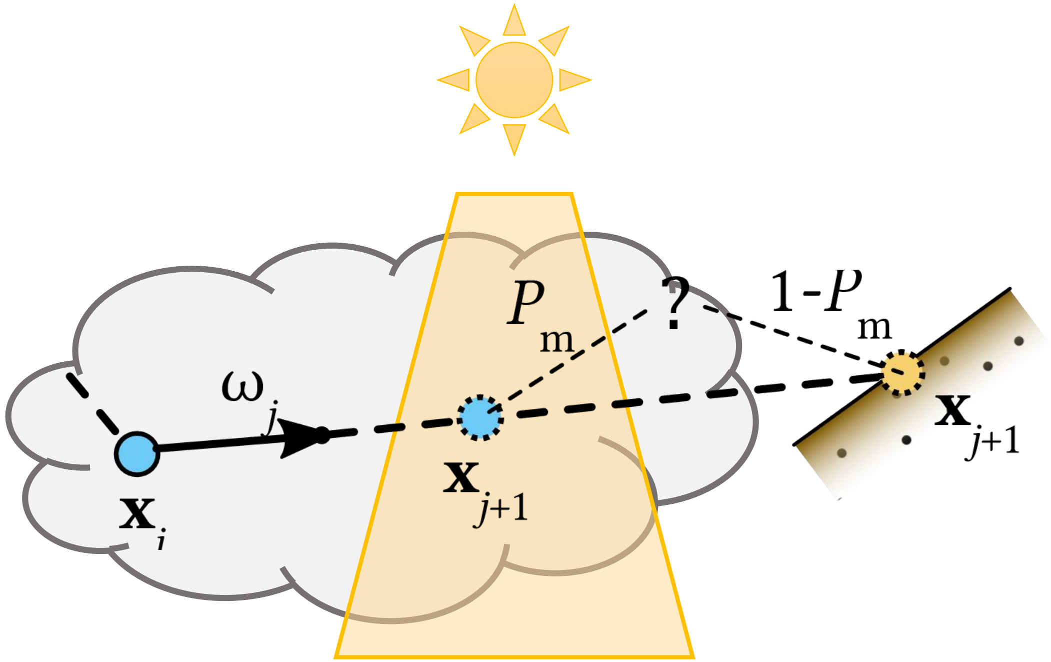

1. Volume or surface decision

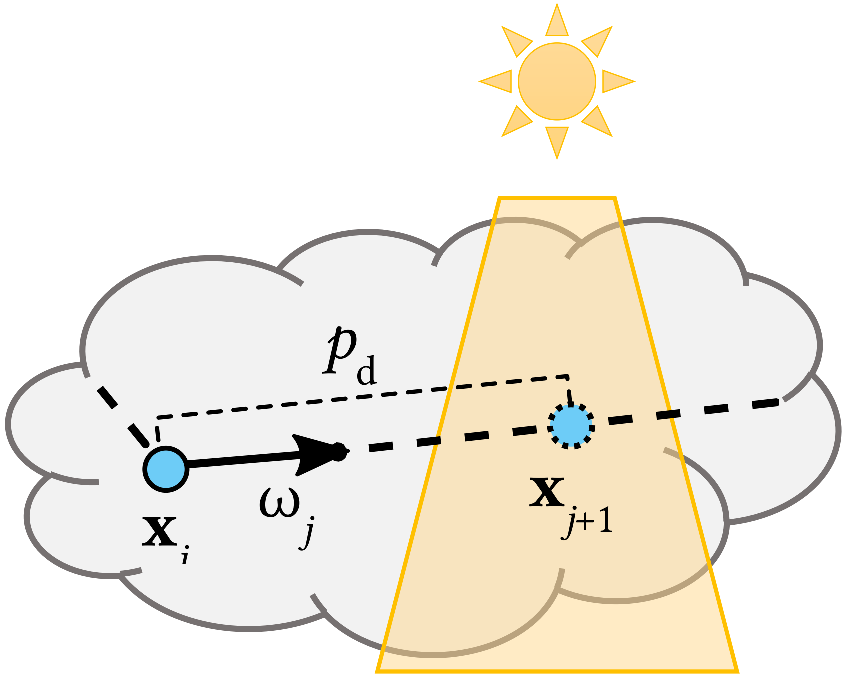

2. Scatter distance

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

DISTANCE SAMPLING

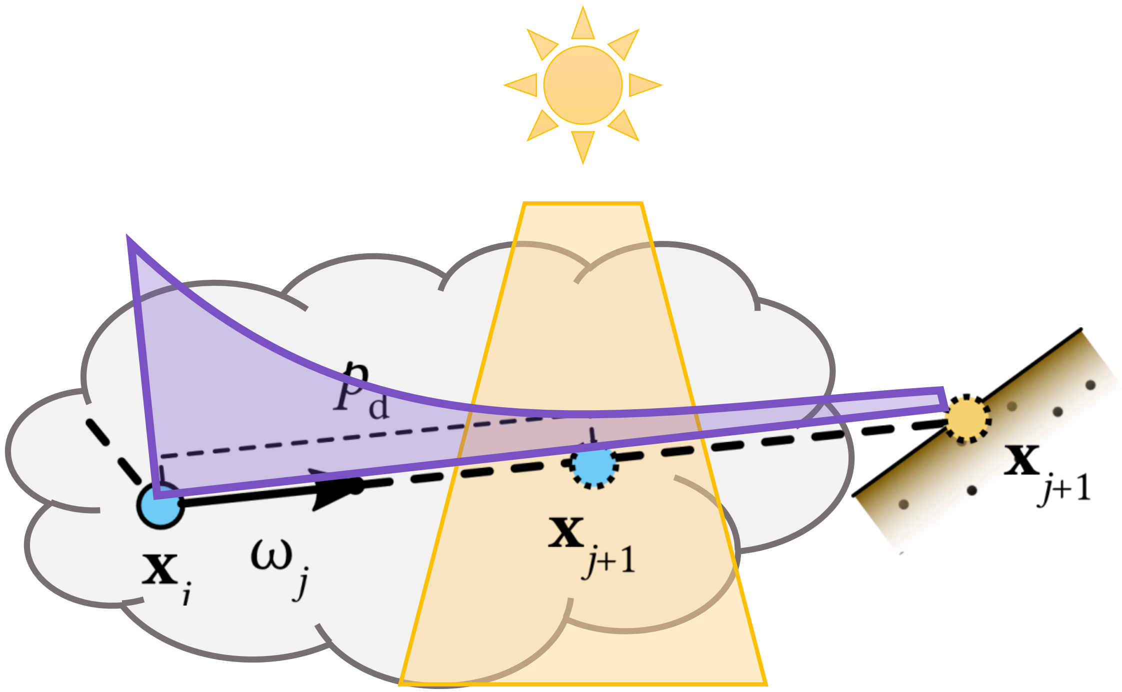

1+2. Event distance

- Standard distance PDF:

$$p_d^{std}(\dots) \propto T(\dots) \cdot \sigma_s(\dots)$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED PRODUCT DISTANCE SAMPLING

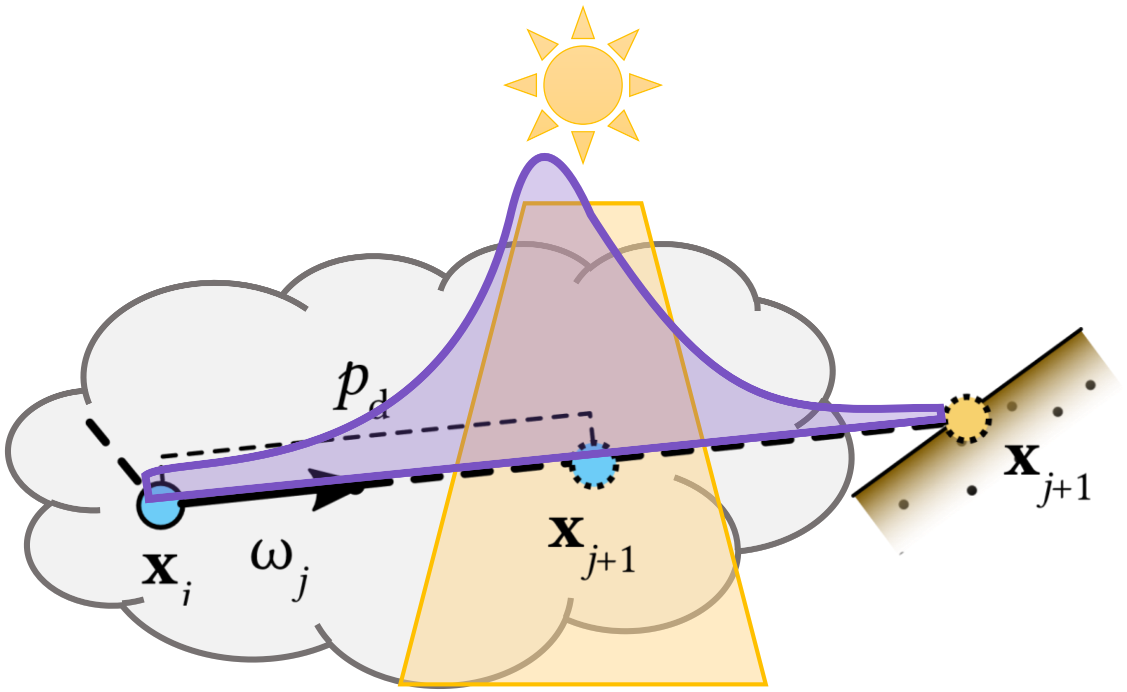

1+2. Event distance

- Our guided PDF:

Our estimates

$$\tilde{p}_d^{zv}(\ldots) = \frac{T(\ldots) \cdot \sigma_s(\ldots) \cdot \boxed{\tilde{L}_i(\ldots)}}{\boxed{\tilde{L}(\ldots)}}$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

NAÏVE TABULATED APPROACH

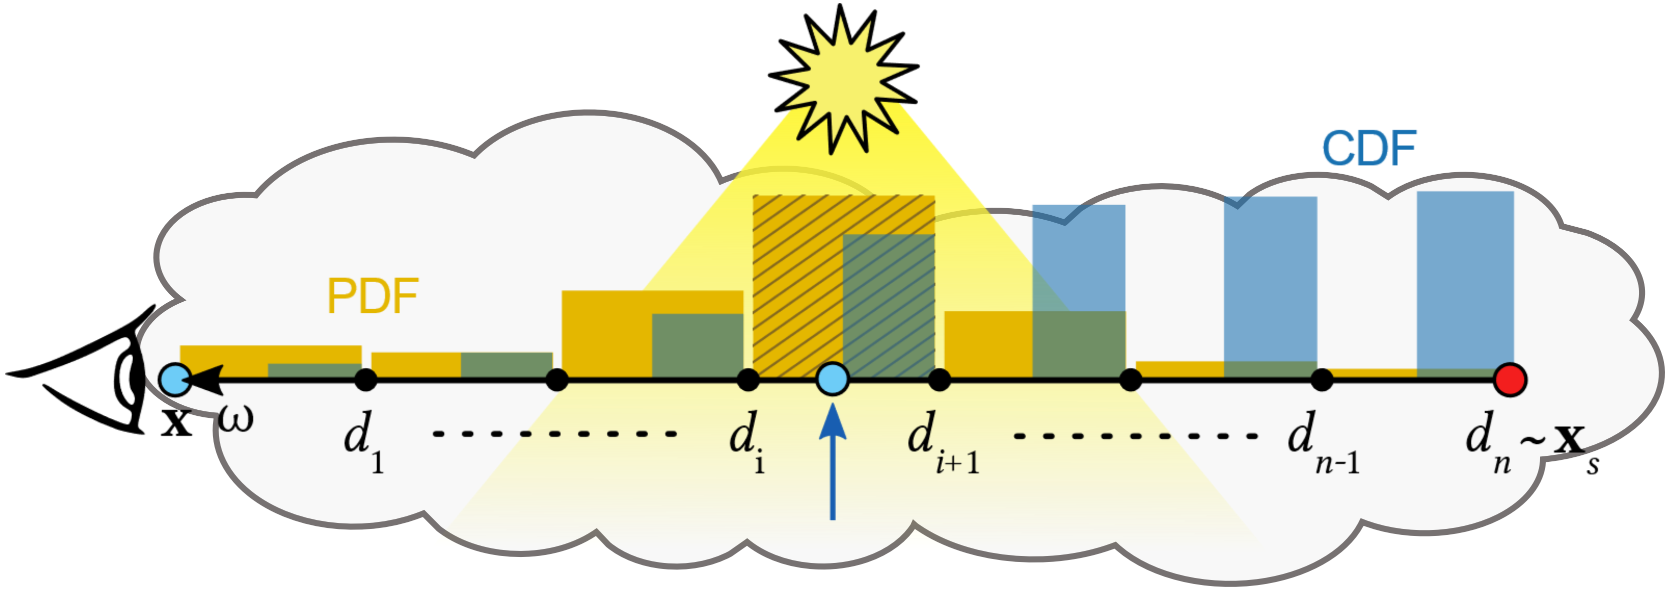

- Naïve tabulated approach:
  - Step through the **complete** volume and build a tabulated PDF
- Inefficient (large scenes dense media):
  - we always need to evaluated all bins first

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

OUR INCREMENTAL GUIDED PRODUCT DISTANCE SAMPLING

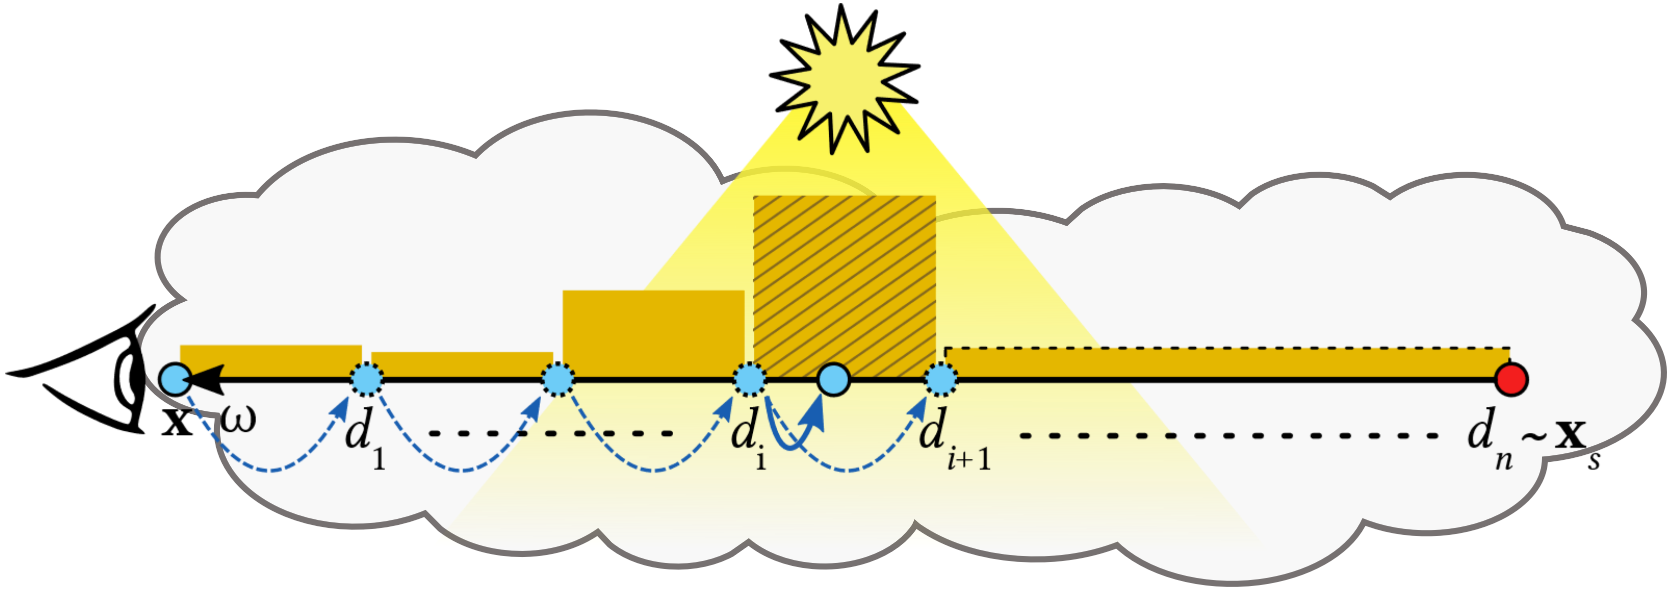

- Incremental approach:
  - At each step make a local decision, if we scatter inside the bin

- **We only need to step until the scattering event happens**

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

Full CDF (30min)

Spp: 548
Avg. steps: 18

Our incremental (30min)

Spp: 1140
Avg. steps: 4

thrive SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

45 min

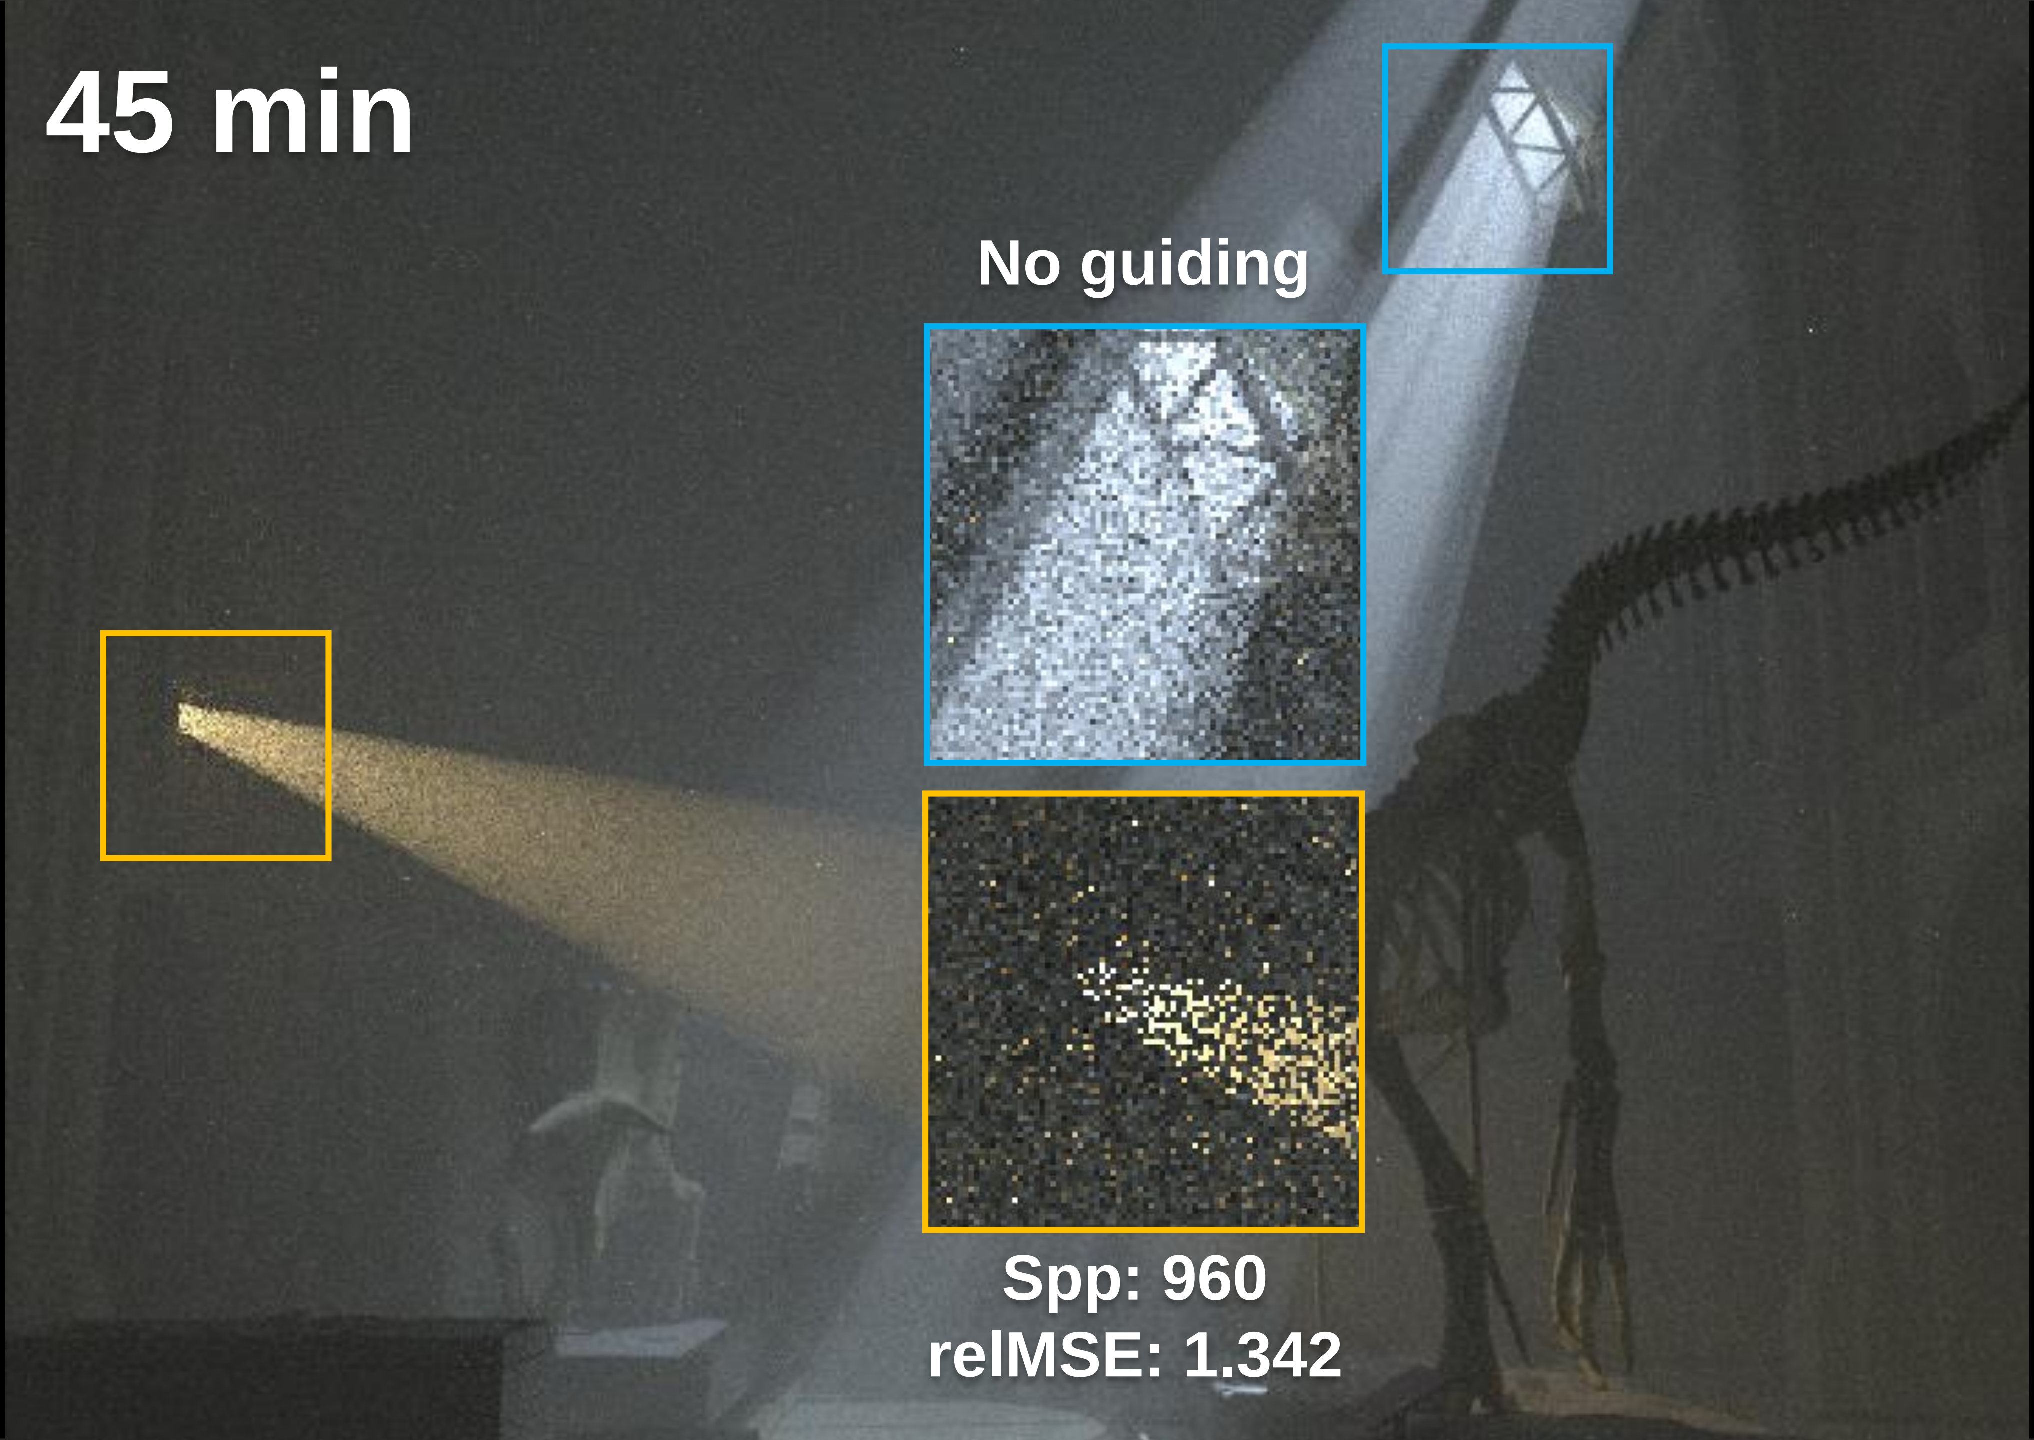
45 min
No guiding
Spp: 960
relMSE: 1.342

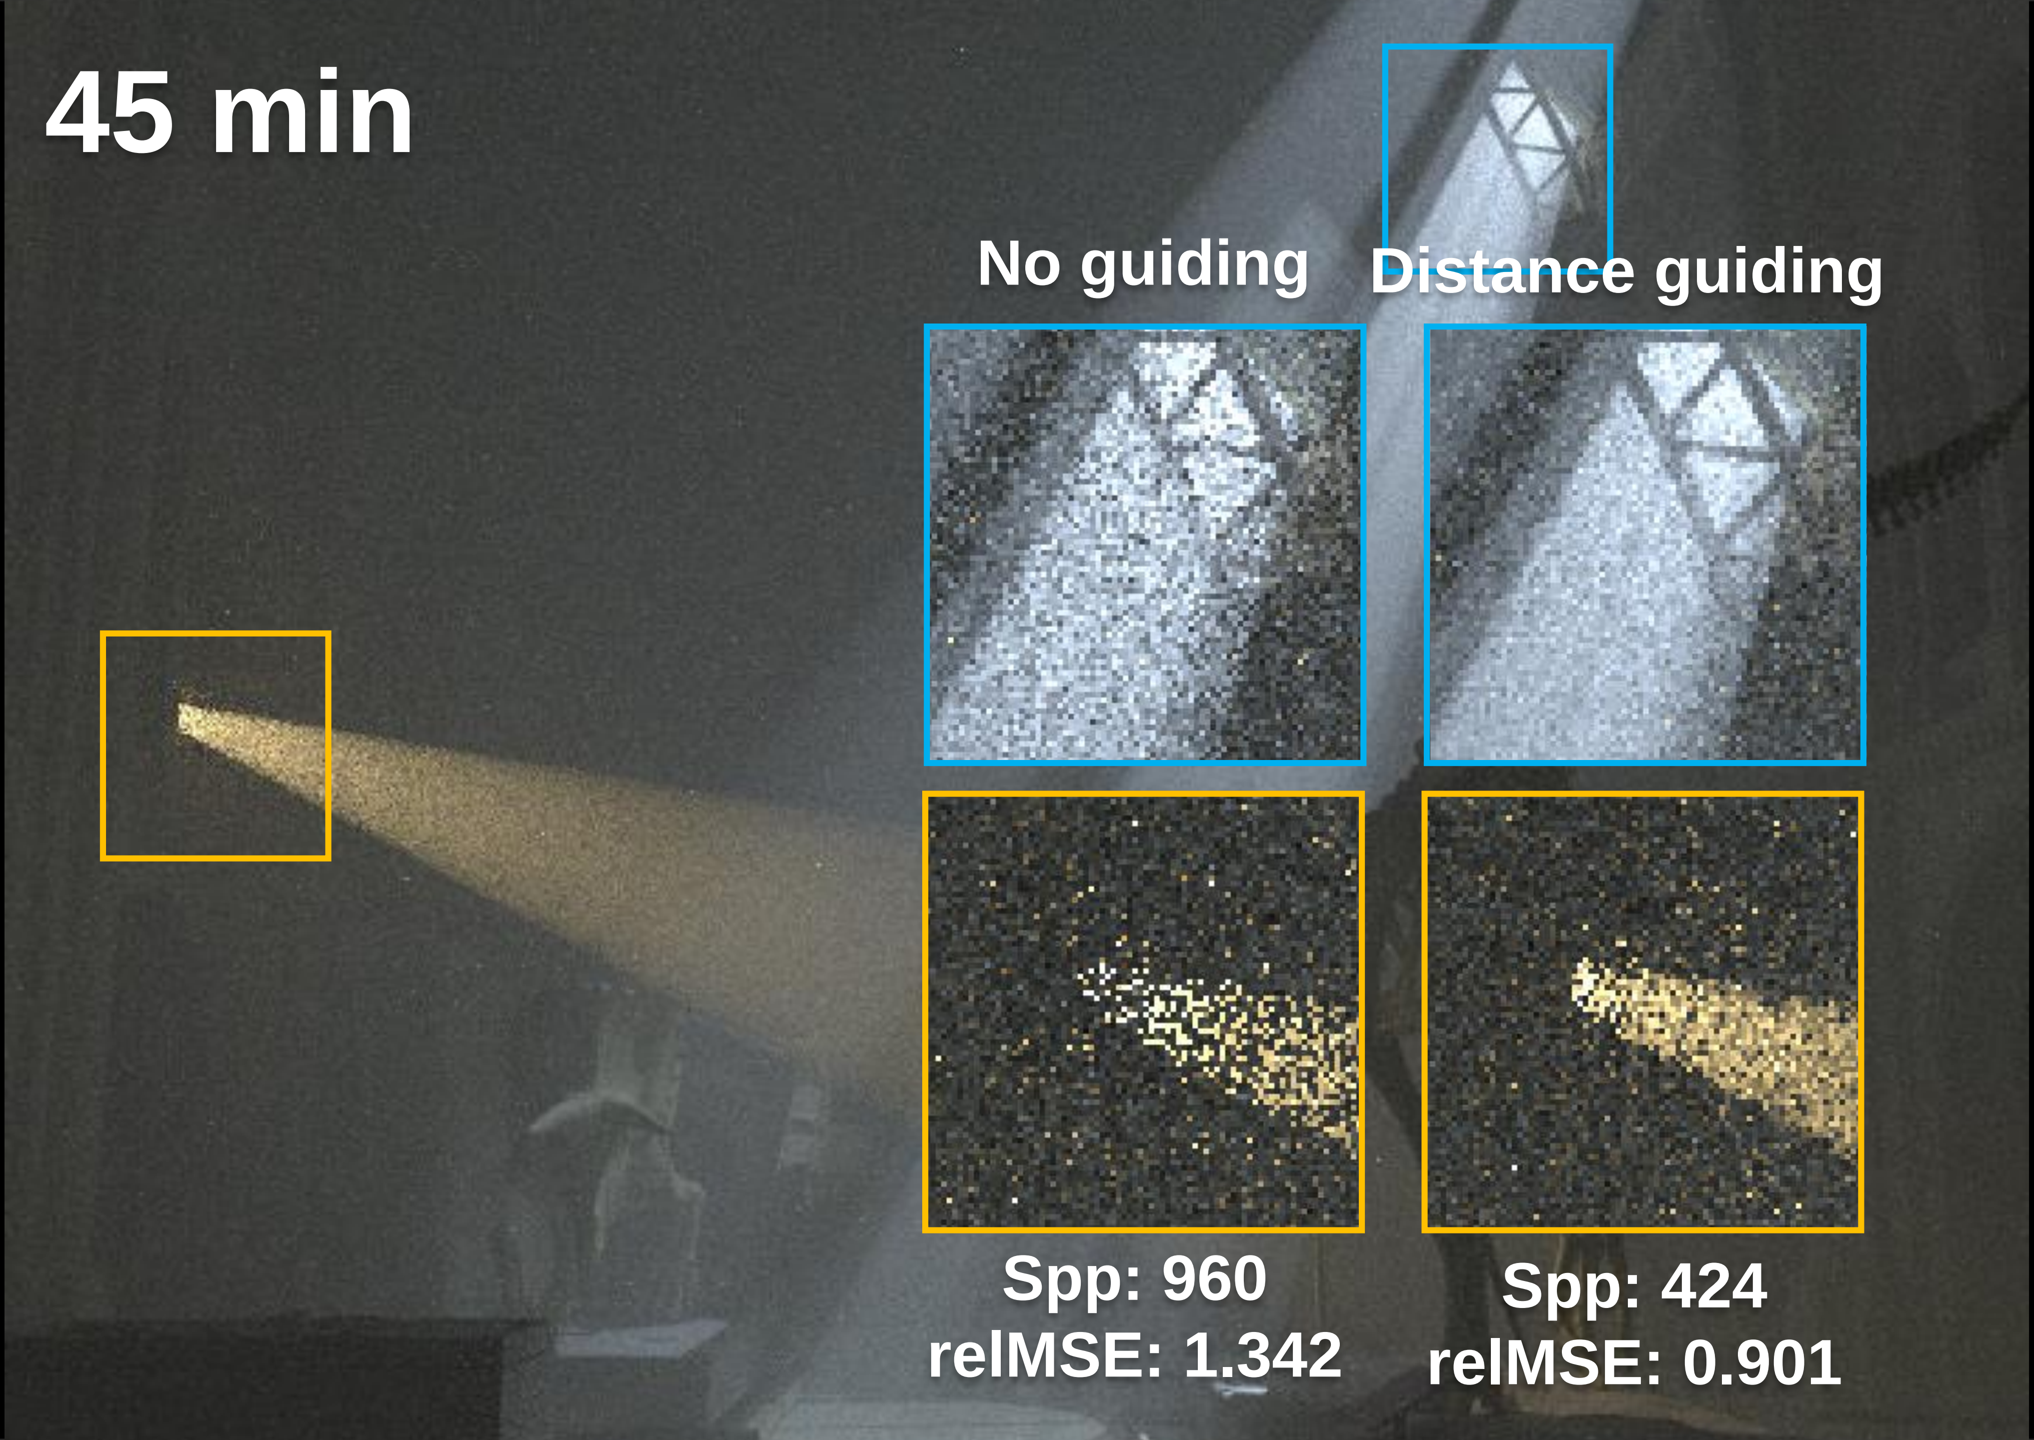
45 min
No guiding
Distance guiding
Spp: 960
relMSE: 1.342
Spp: 424
relMSE: 0.901

No guiding (256 spp)

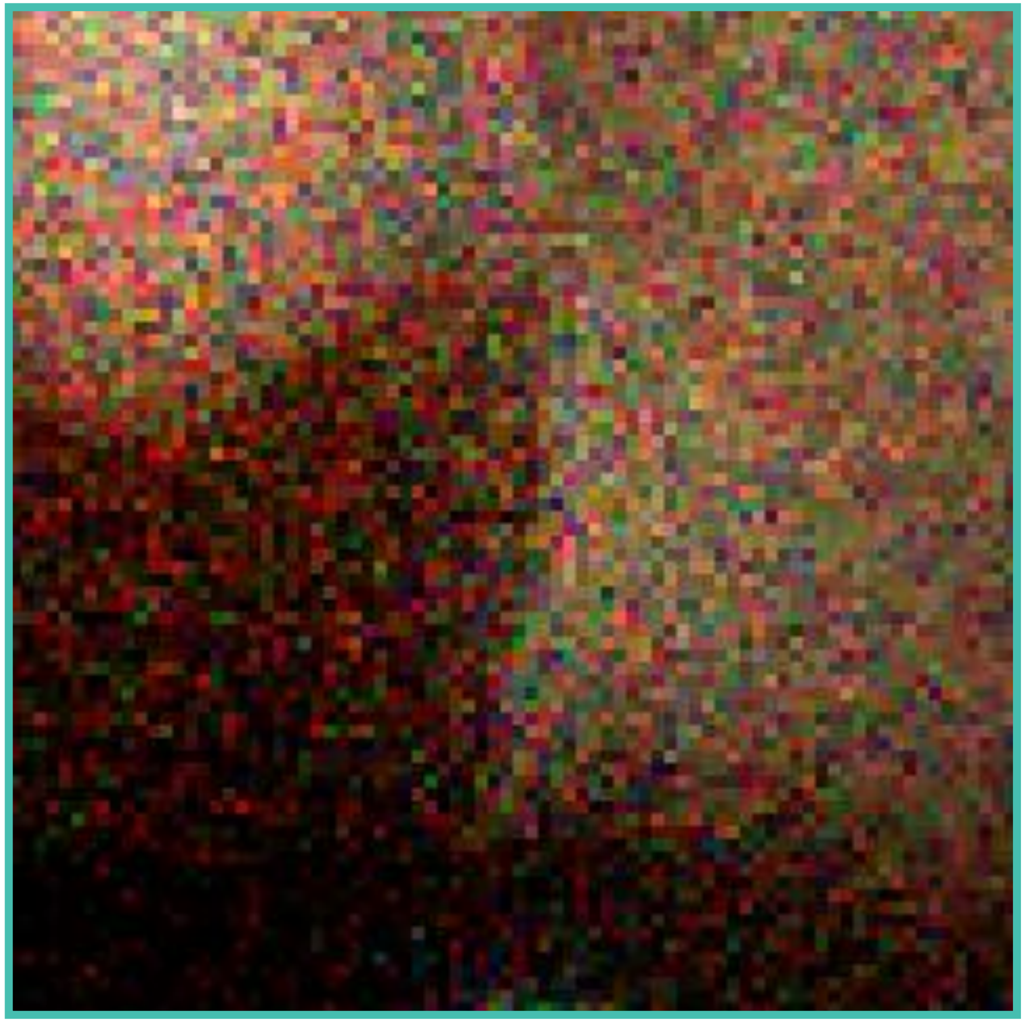

Distance guiding (256 spp)

- Here, distance sampling is not the main cause of variance!!

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

DIRECTIONAL SAMPLING

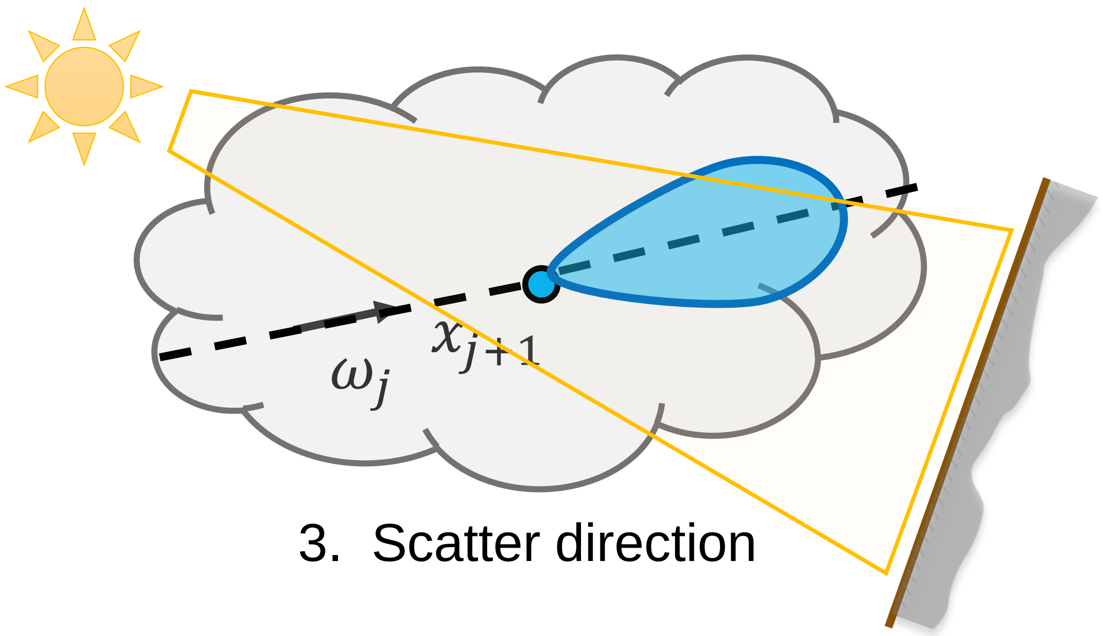

3. Scatter direction

- Standard PDF:

$$p_{\omega}^{std}(\dots) \propto \tilde{f}(\dots)$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED PRODUCT DIRECTIONAL SAMPLING

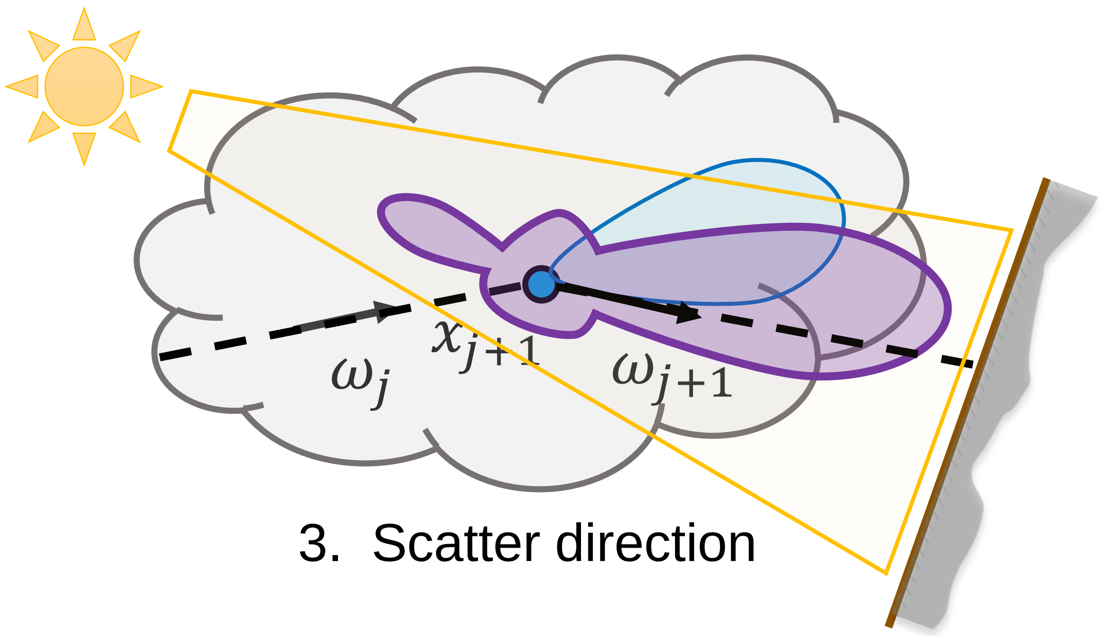

3. Scatter direction

- Our guided PDF:

Our estimates

$$\tilde{p}_{\omega}^{zv}(\dots) \propto \boxed{\tilde{f}(\dots)} \cdot \boxed{\tilde{L}(\dots)}$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

OUR GUIDED PRODUCT DIRECTIONAL SAMPLING

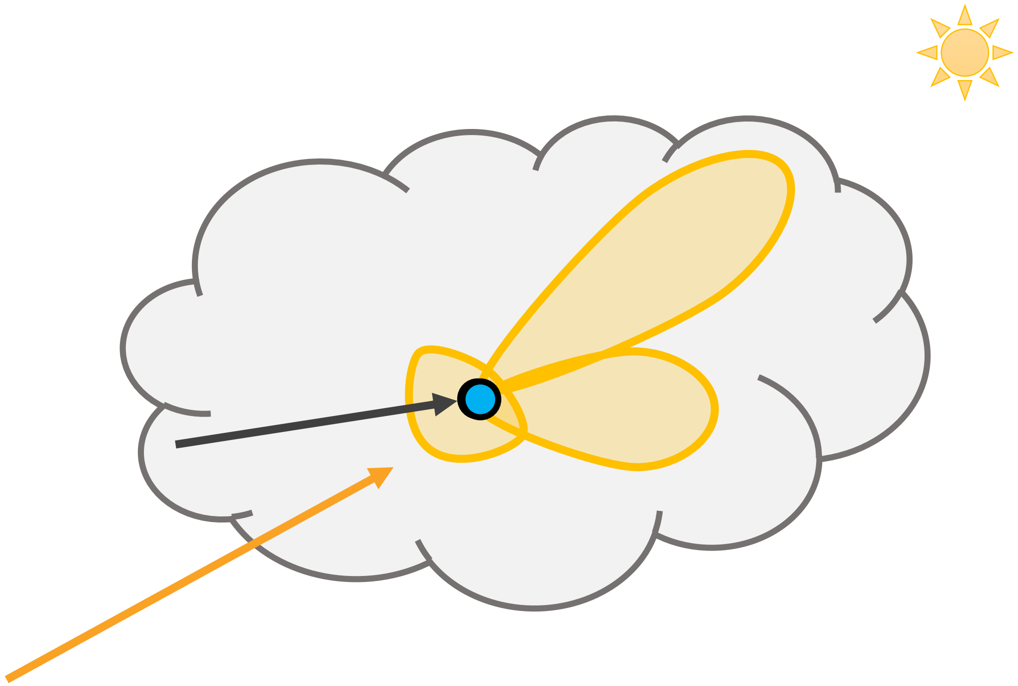

Incident radiance VMM

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

OUR GUIDED PRODUCT DIRECTIONAL SAMPLING

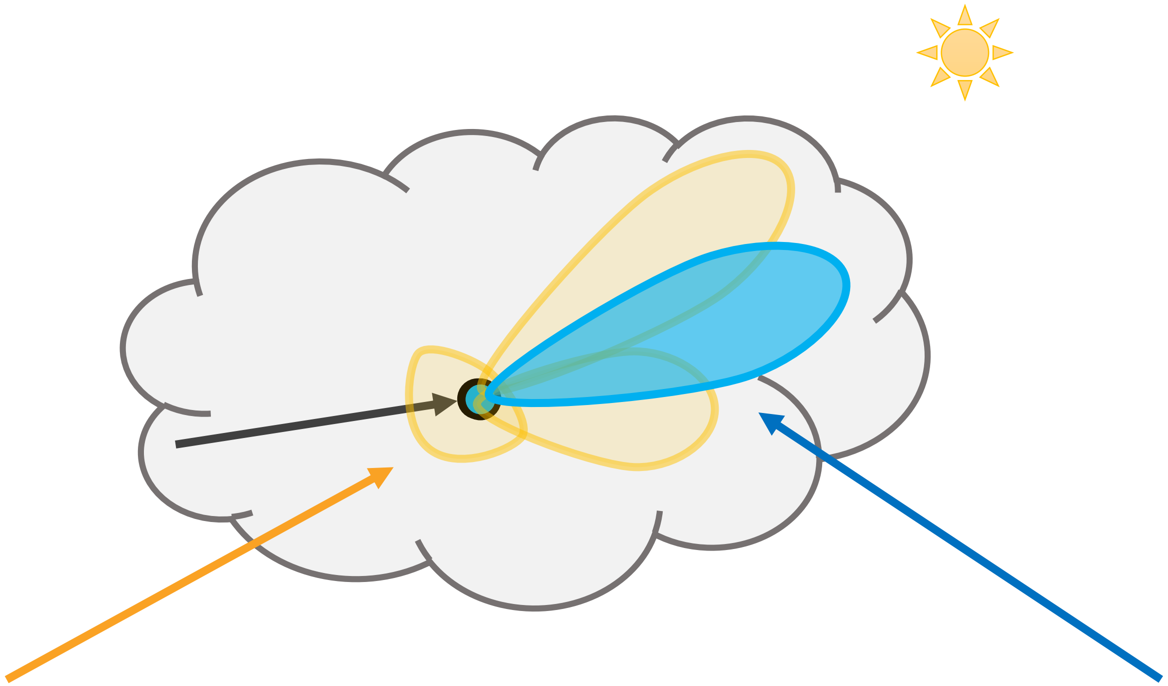

Incident radiance VMM

Phase function VMM

OUR GUIDED PRODUCT DIRECTIONAL SAMPLING

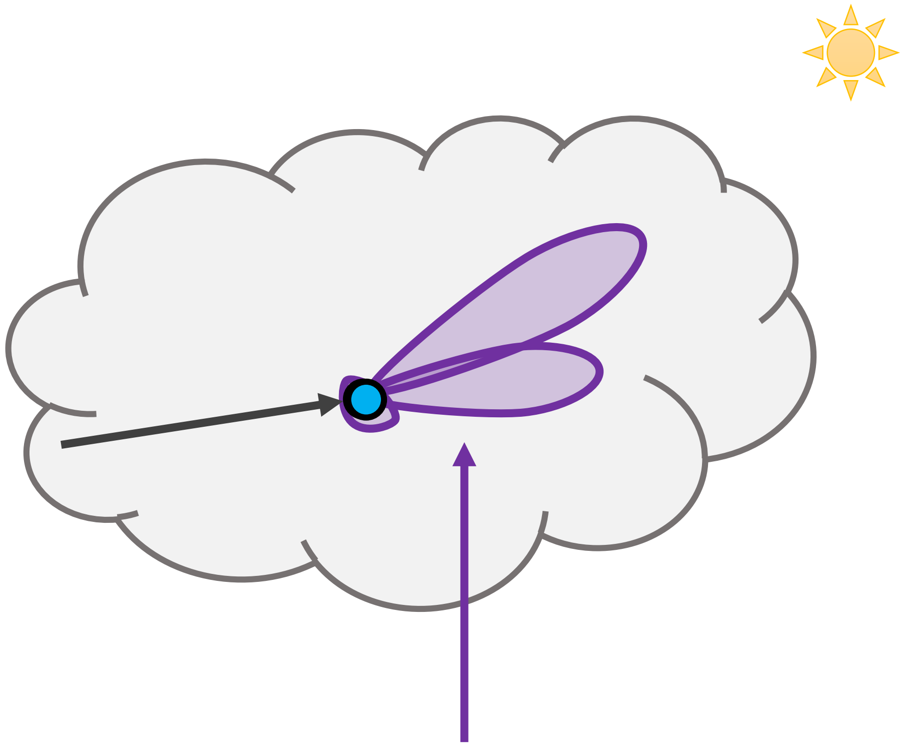

Product sampling VMM

30 min

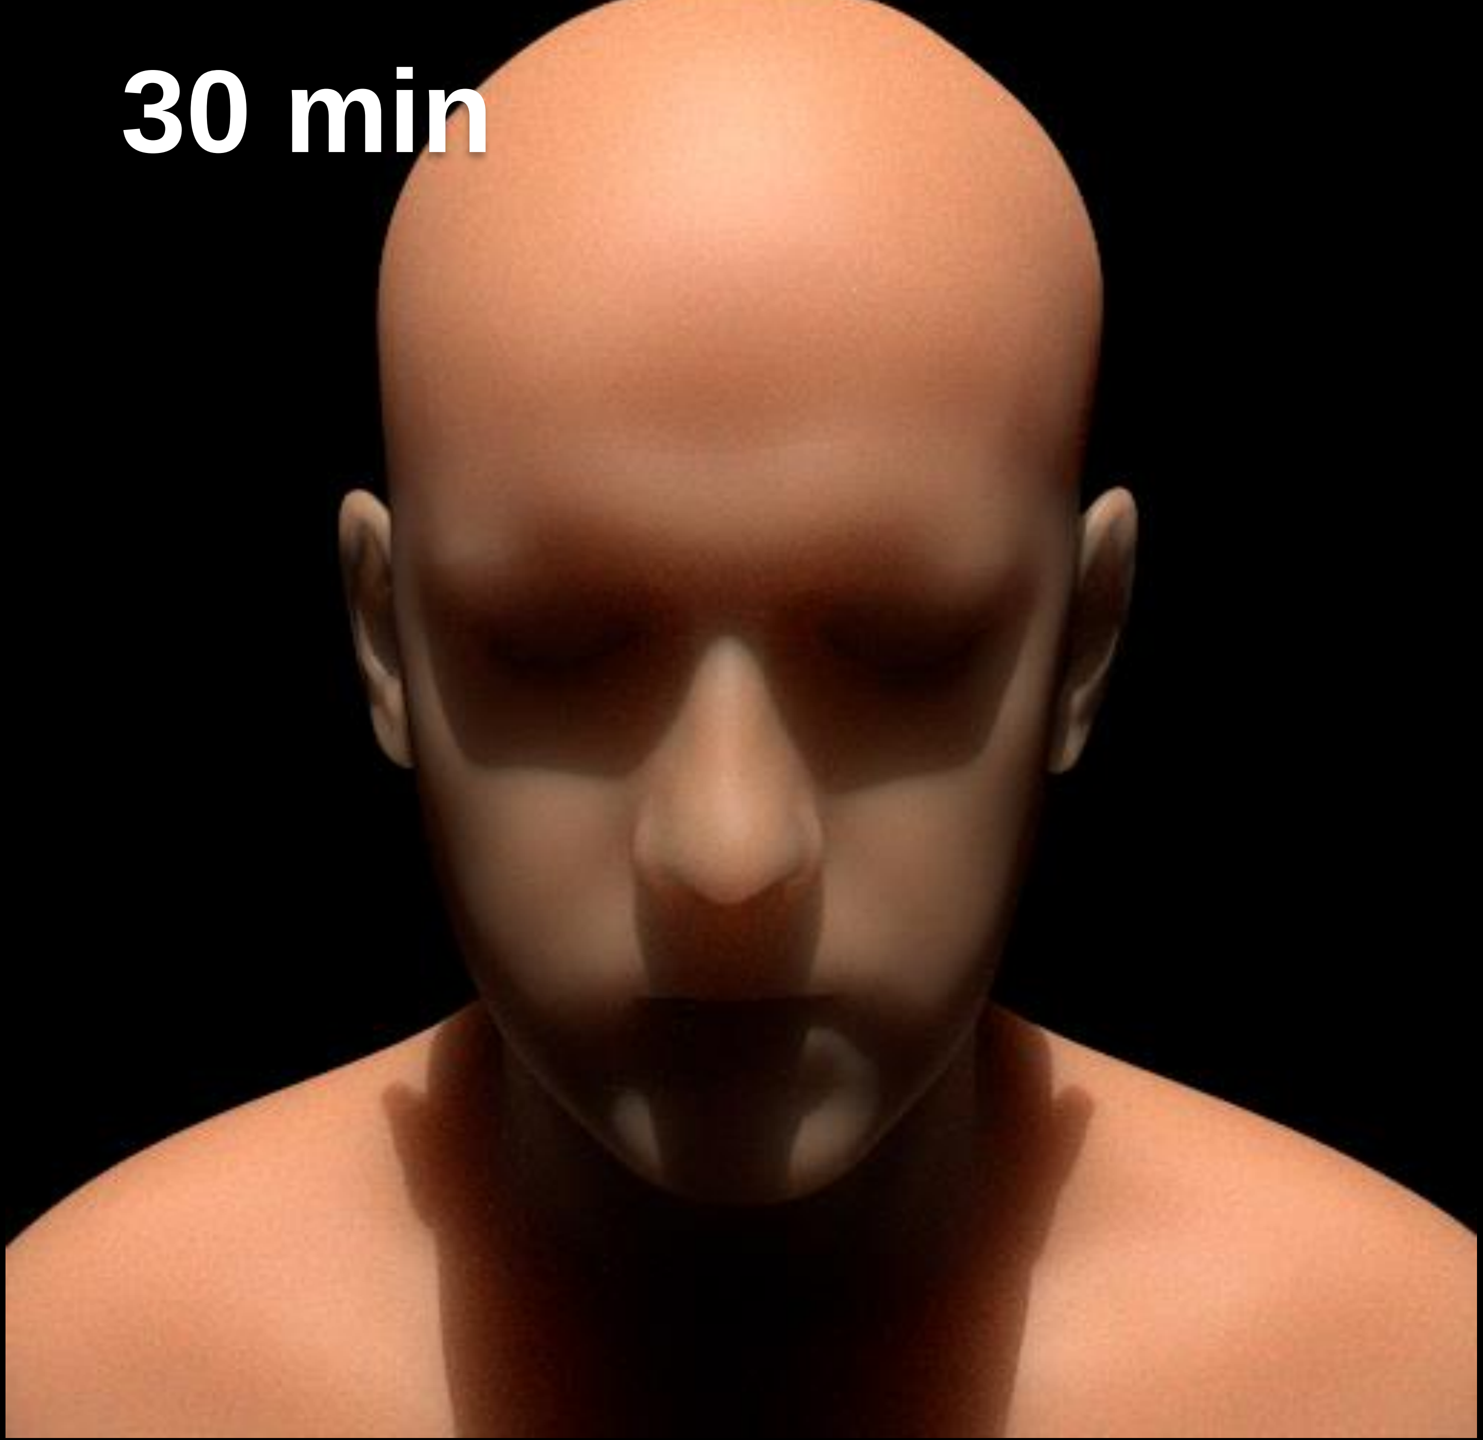

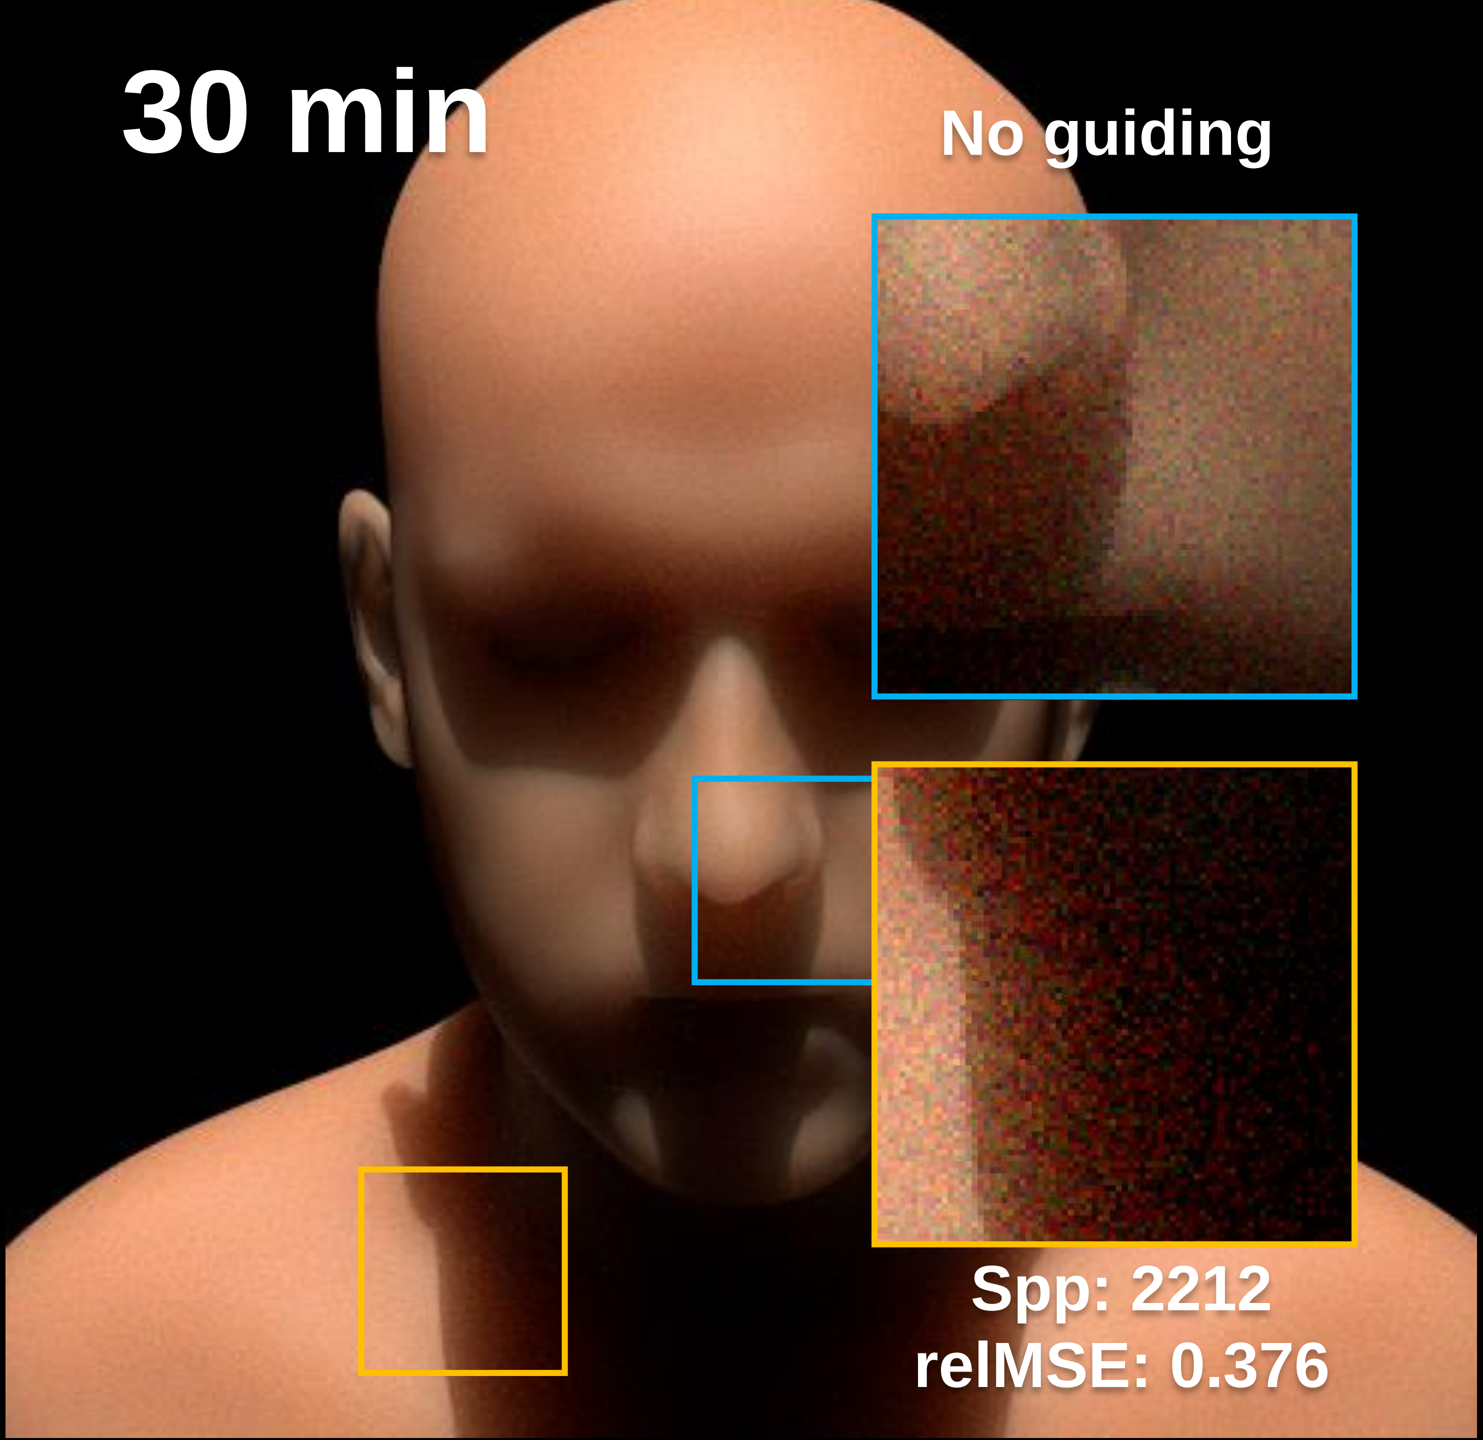
30 min
No guiding
Spp: 2212
relMSE: 0.376

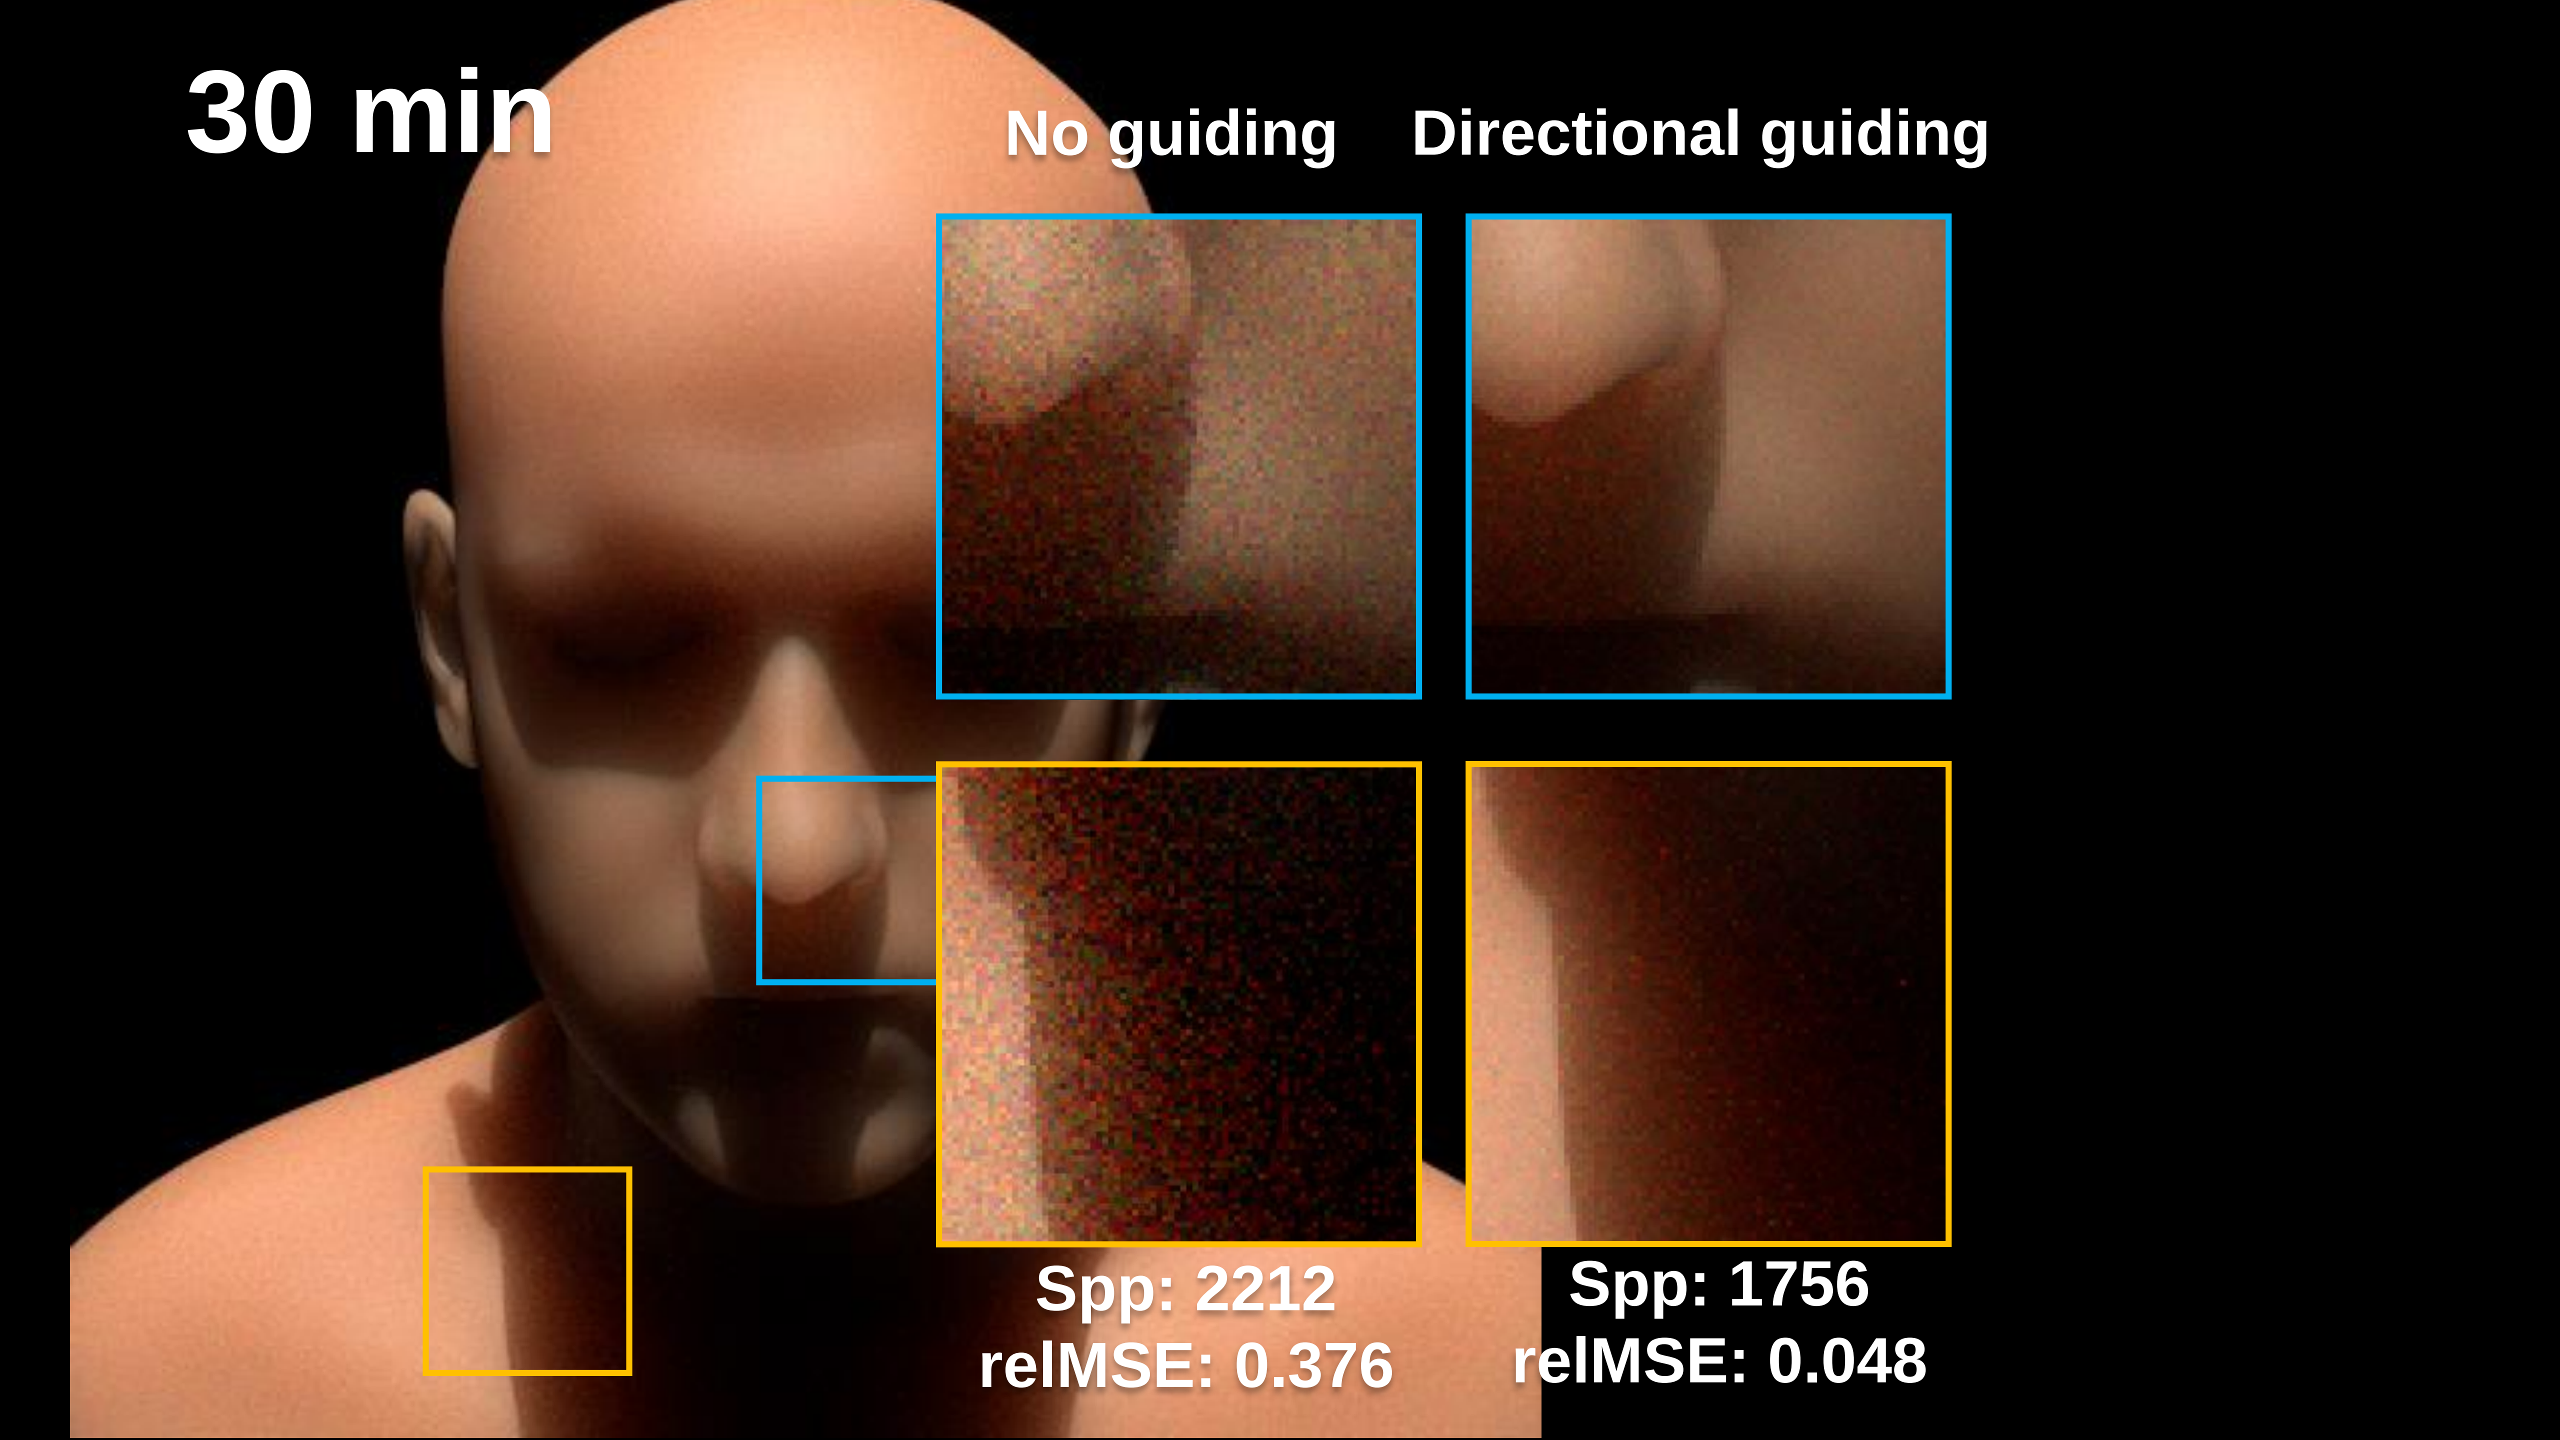
30 min
No guiding
Directional guiding
Spp: 2212
relMSE: 0.376
Spp: 1756
relMSE: 0.048

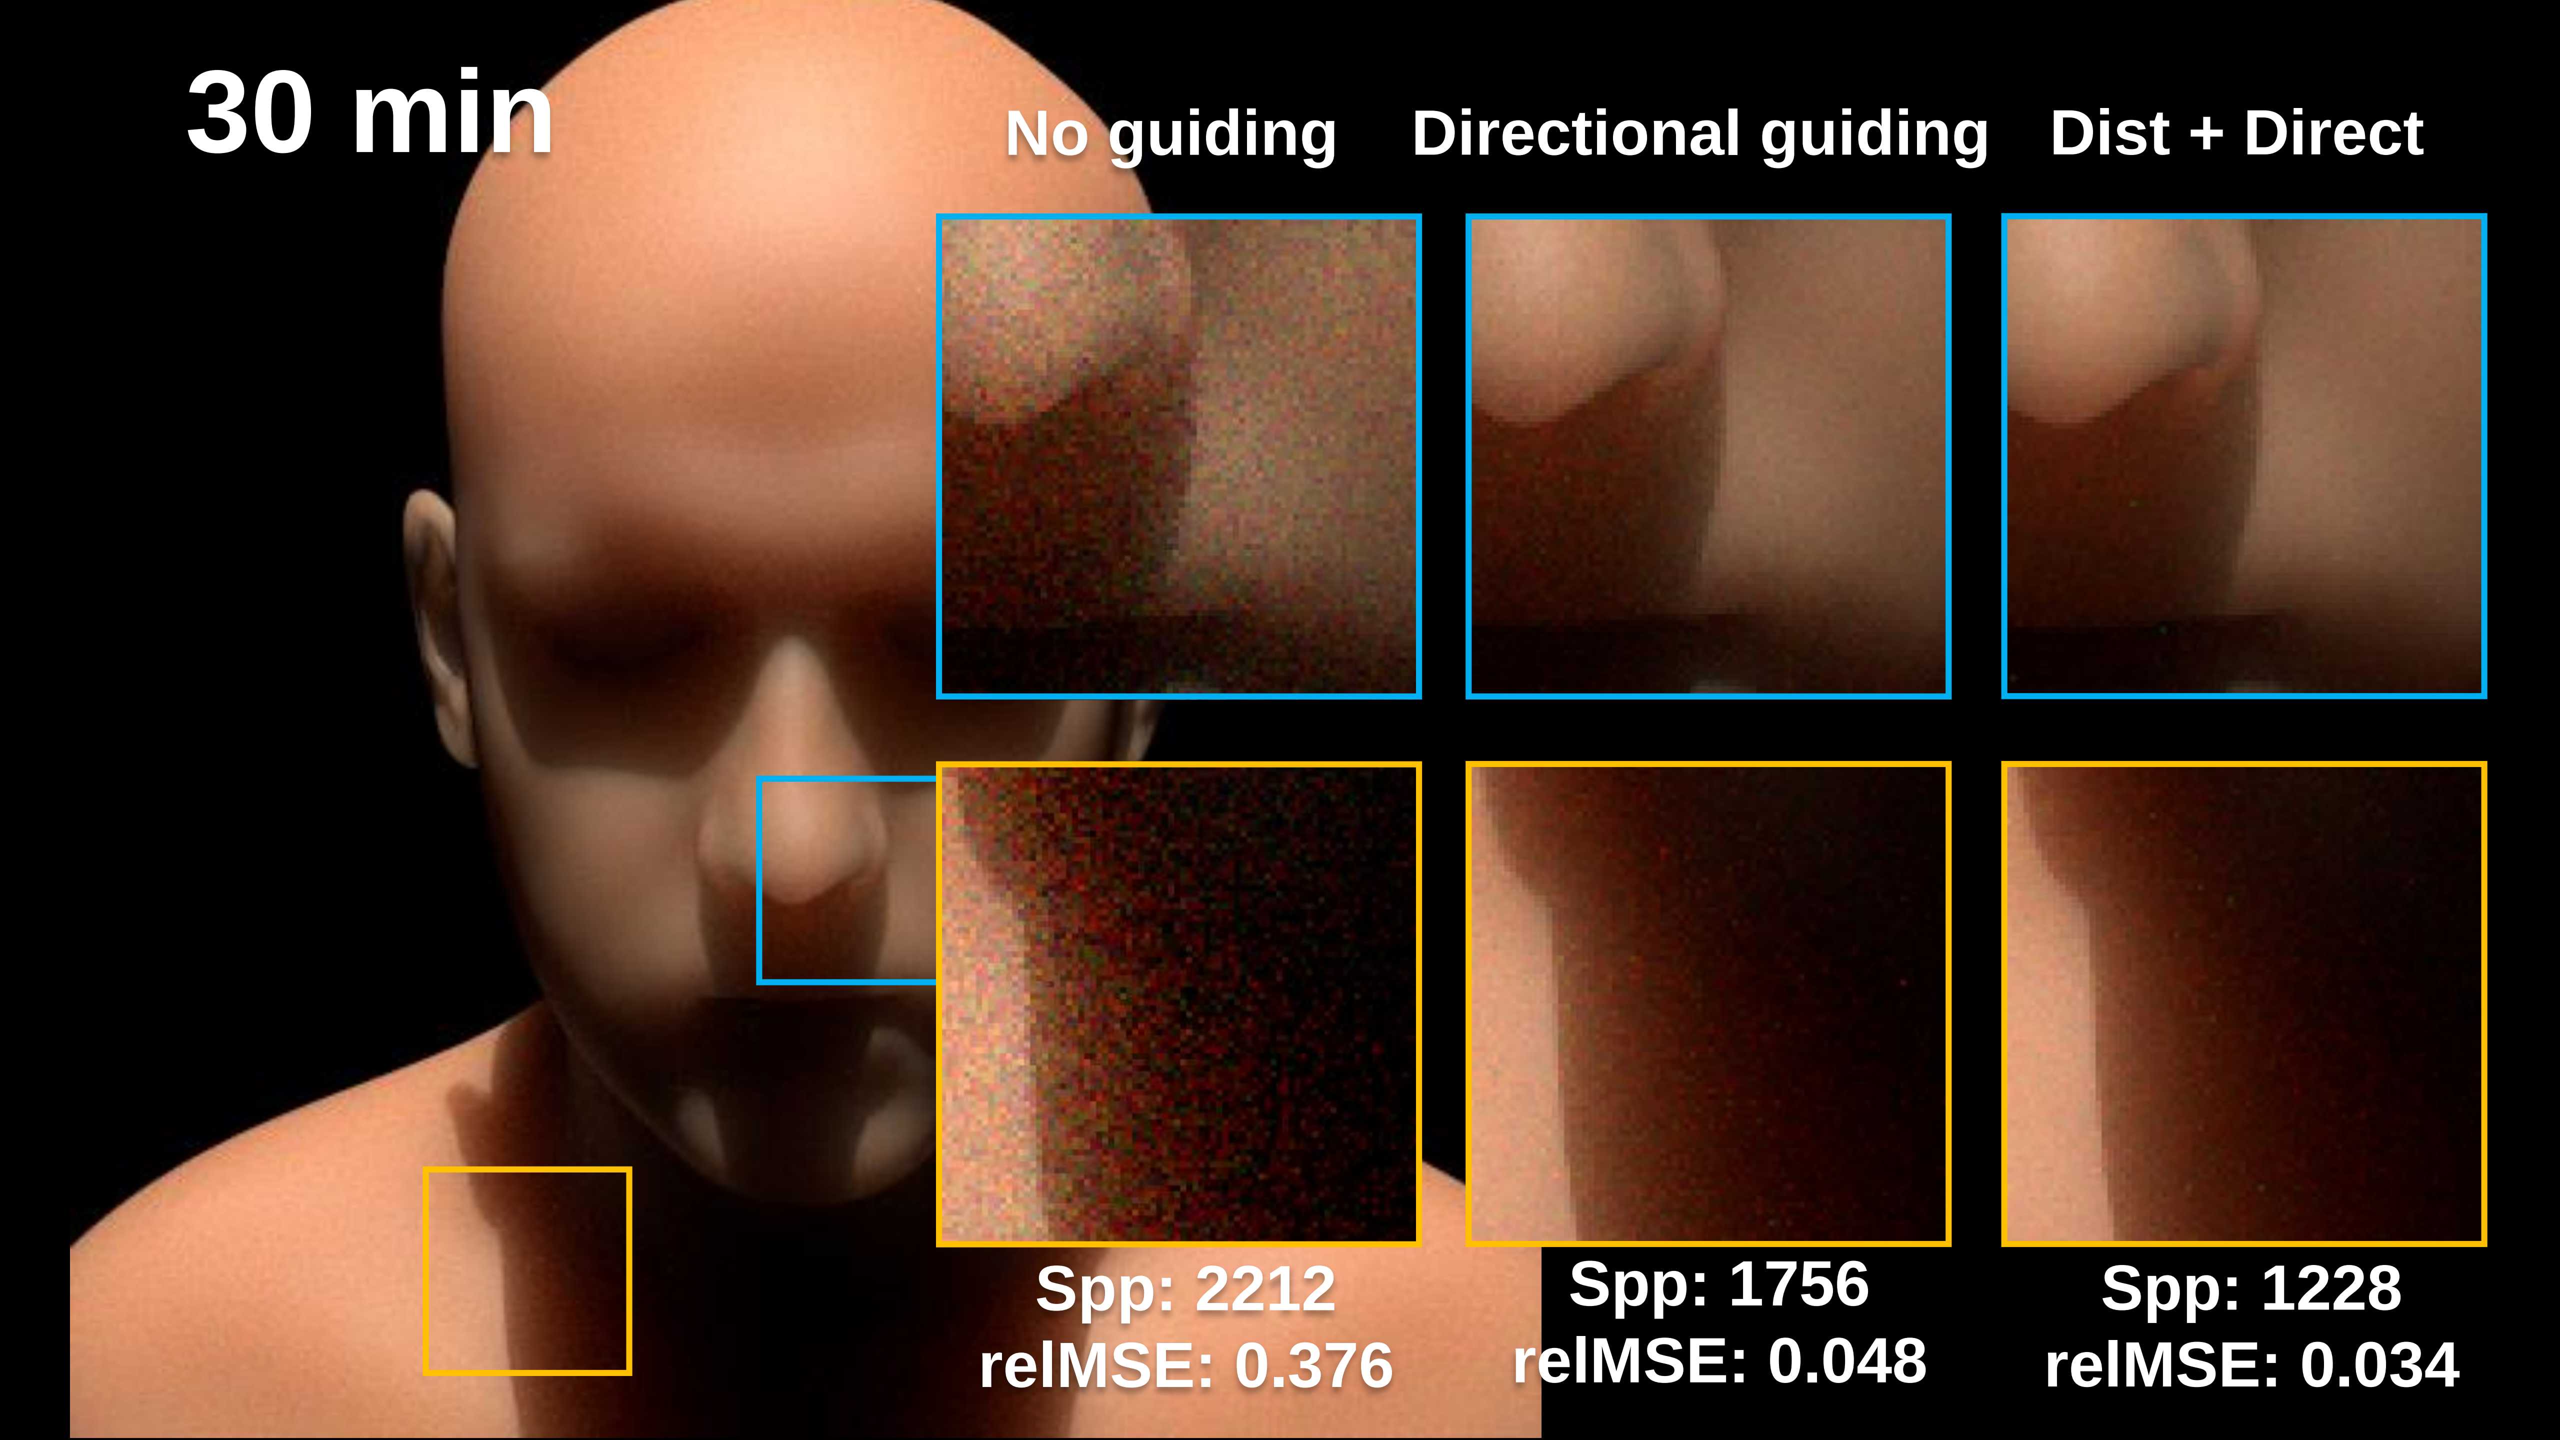
30 min
No guiding
Directional guiding
Dist + Direct
Spp: 2212
relMSE: 0.376
Spp: 1756
relMSE: 0.048
Spp: 1228
relMSE: 0.034

IMPORTANCE OF THE PRODUCT FOR DENSE ANISOTROPIC MEDIA

No Guiding (256 spp)

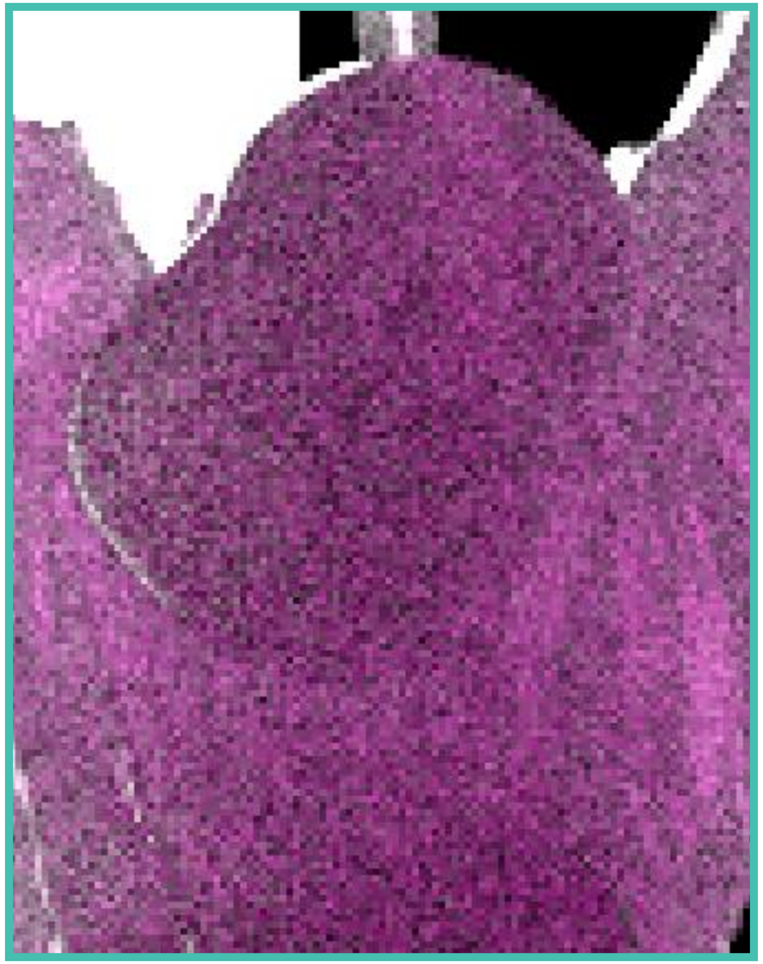

Product Guiding (256 spp)

Illum Guiding (256 spp)

GUIDED RUSSIAN ROULETTE AND SPLITTING

Directional

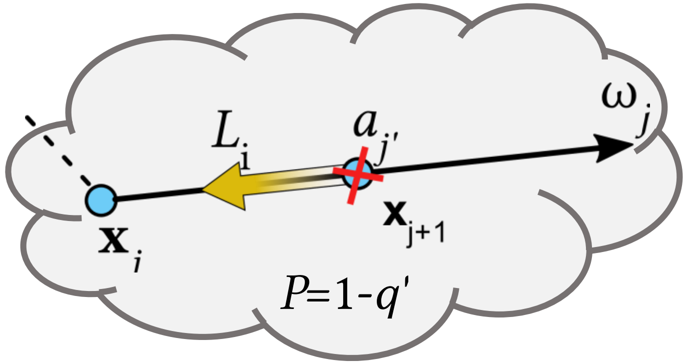

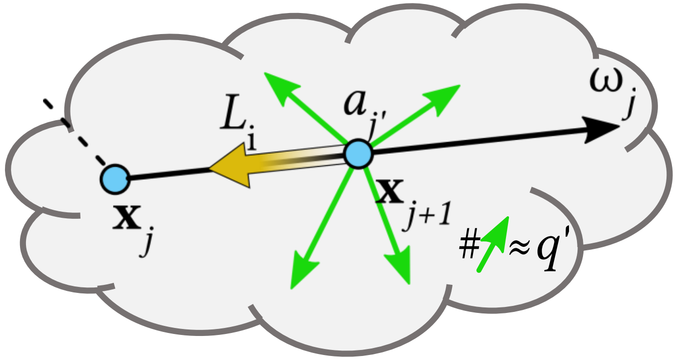

4a. Termination

4b. Splitting

Distance

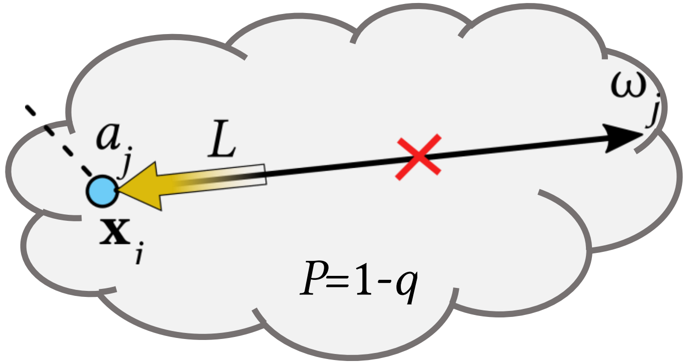

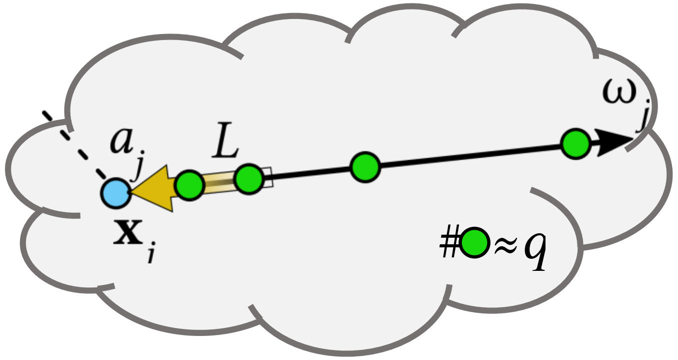

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED RUSSIAN ROULETTE AND SPLITTING

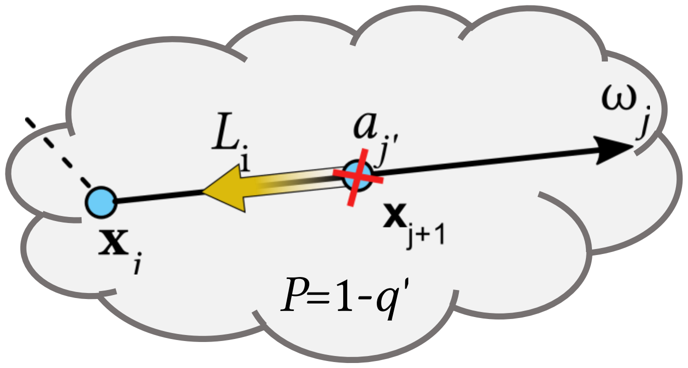

4a. Termination

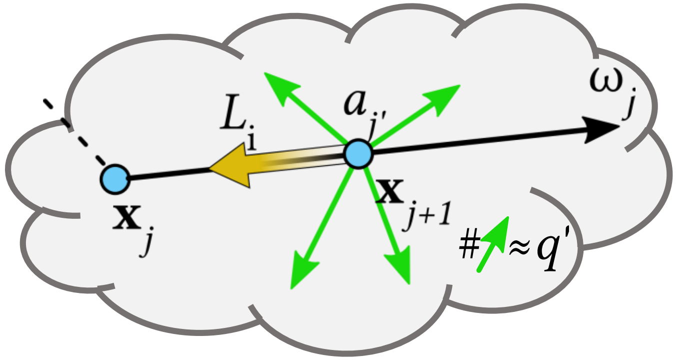

4b. Splitting

- Post-sampling compensation strategies:
  - Identify, if we **did** a **sub-optimal** sampling decision
  - Terminate: to increase performance
  - Split: bound/reduce sample variance

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED RUSSIAN ROULETTE AND SPLITTING

Path contribution

survival prob / splitting factor → $$q = \frac{E[X]}{I}$$

Reference solution

- Path contribution: $E[X]$
  - The expected contribution if we continue the path
- Reference solution: I
  - the final pixel value

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED RUSSIAN ROULETTE AND SPLITTING

Path contribution

survival prob / splitting factor

$$q = \frac{E[X]}{I} = 1$$

Zero-Variance Estimator

Reference solution

- Path contribution: $E[X]$
  - The expected contribution if we continue the path
- Reference solution: I
  - the final pixel value

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED RUSSIAN ROULETTE AND SPLITTING

Path contribution

survival prob / splitting factor $\rightarrow$ $$q = \frac{E[X]}{I}$$

Reference solution

- If $q \leq 1$: Russian Roulette
  - Terminates low contributing paths
  - Survival probability: q
- If $q > 1$: Splitting
  - Splits an under sampled paths with a potential high contribution (q times)

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

ESTIMATED PATH CONTRIBUTION

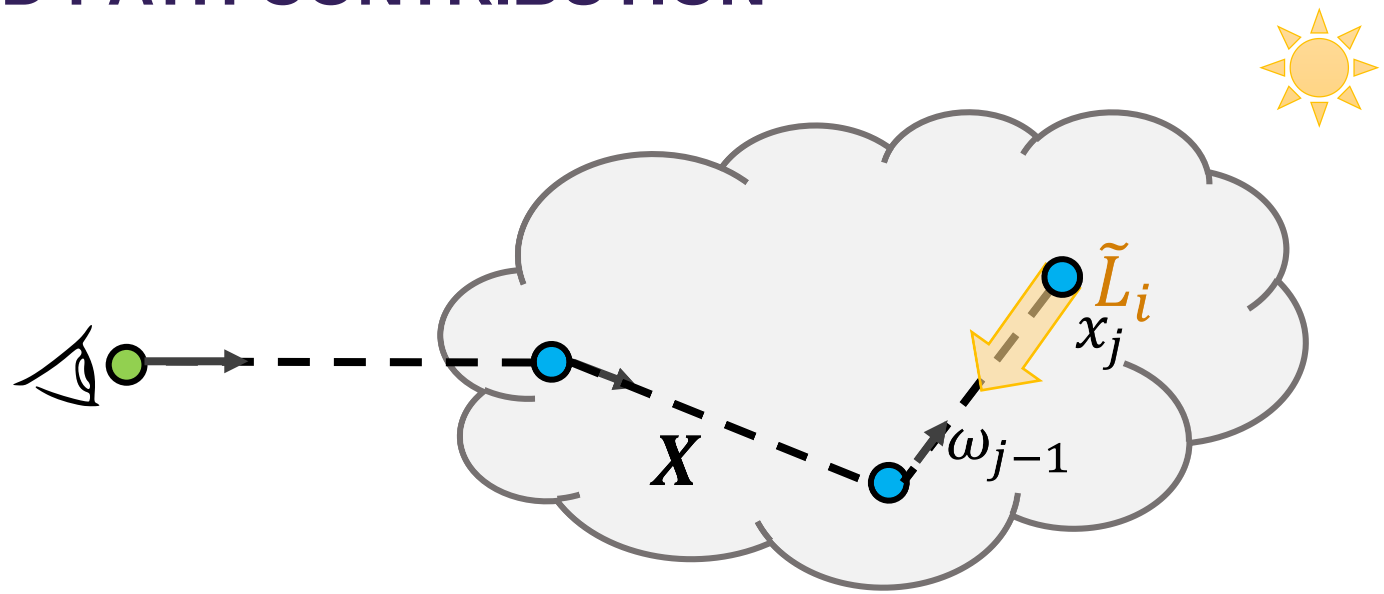

Path throughput: $f(\boldsymbol{X})/p(\boldsymbol{X})$

In-scattered radiance estimate

$$E[\boldsymbol{X}] = a'(\boldsymbol{X}) \cdot \tilde{L}_i\left(x_j, \omega_{j-1}\right)$$

- See course notes or paper for more details

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

GUIDED RUSSIAN ROULETTE AND SPLITTING: PIXEL ESTIMATE

- Ray marched cache to integrate: $T \cdot \sigma_s \cdot \tilde{L}_i$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

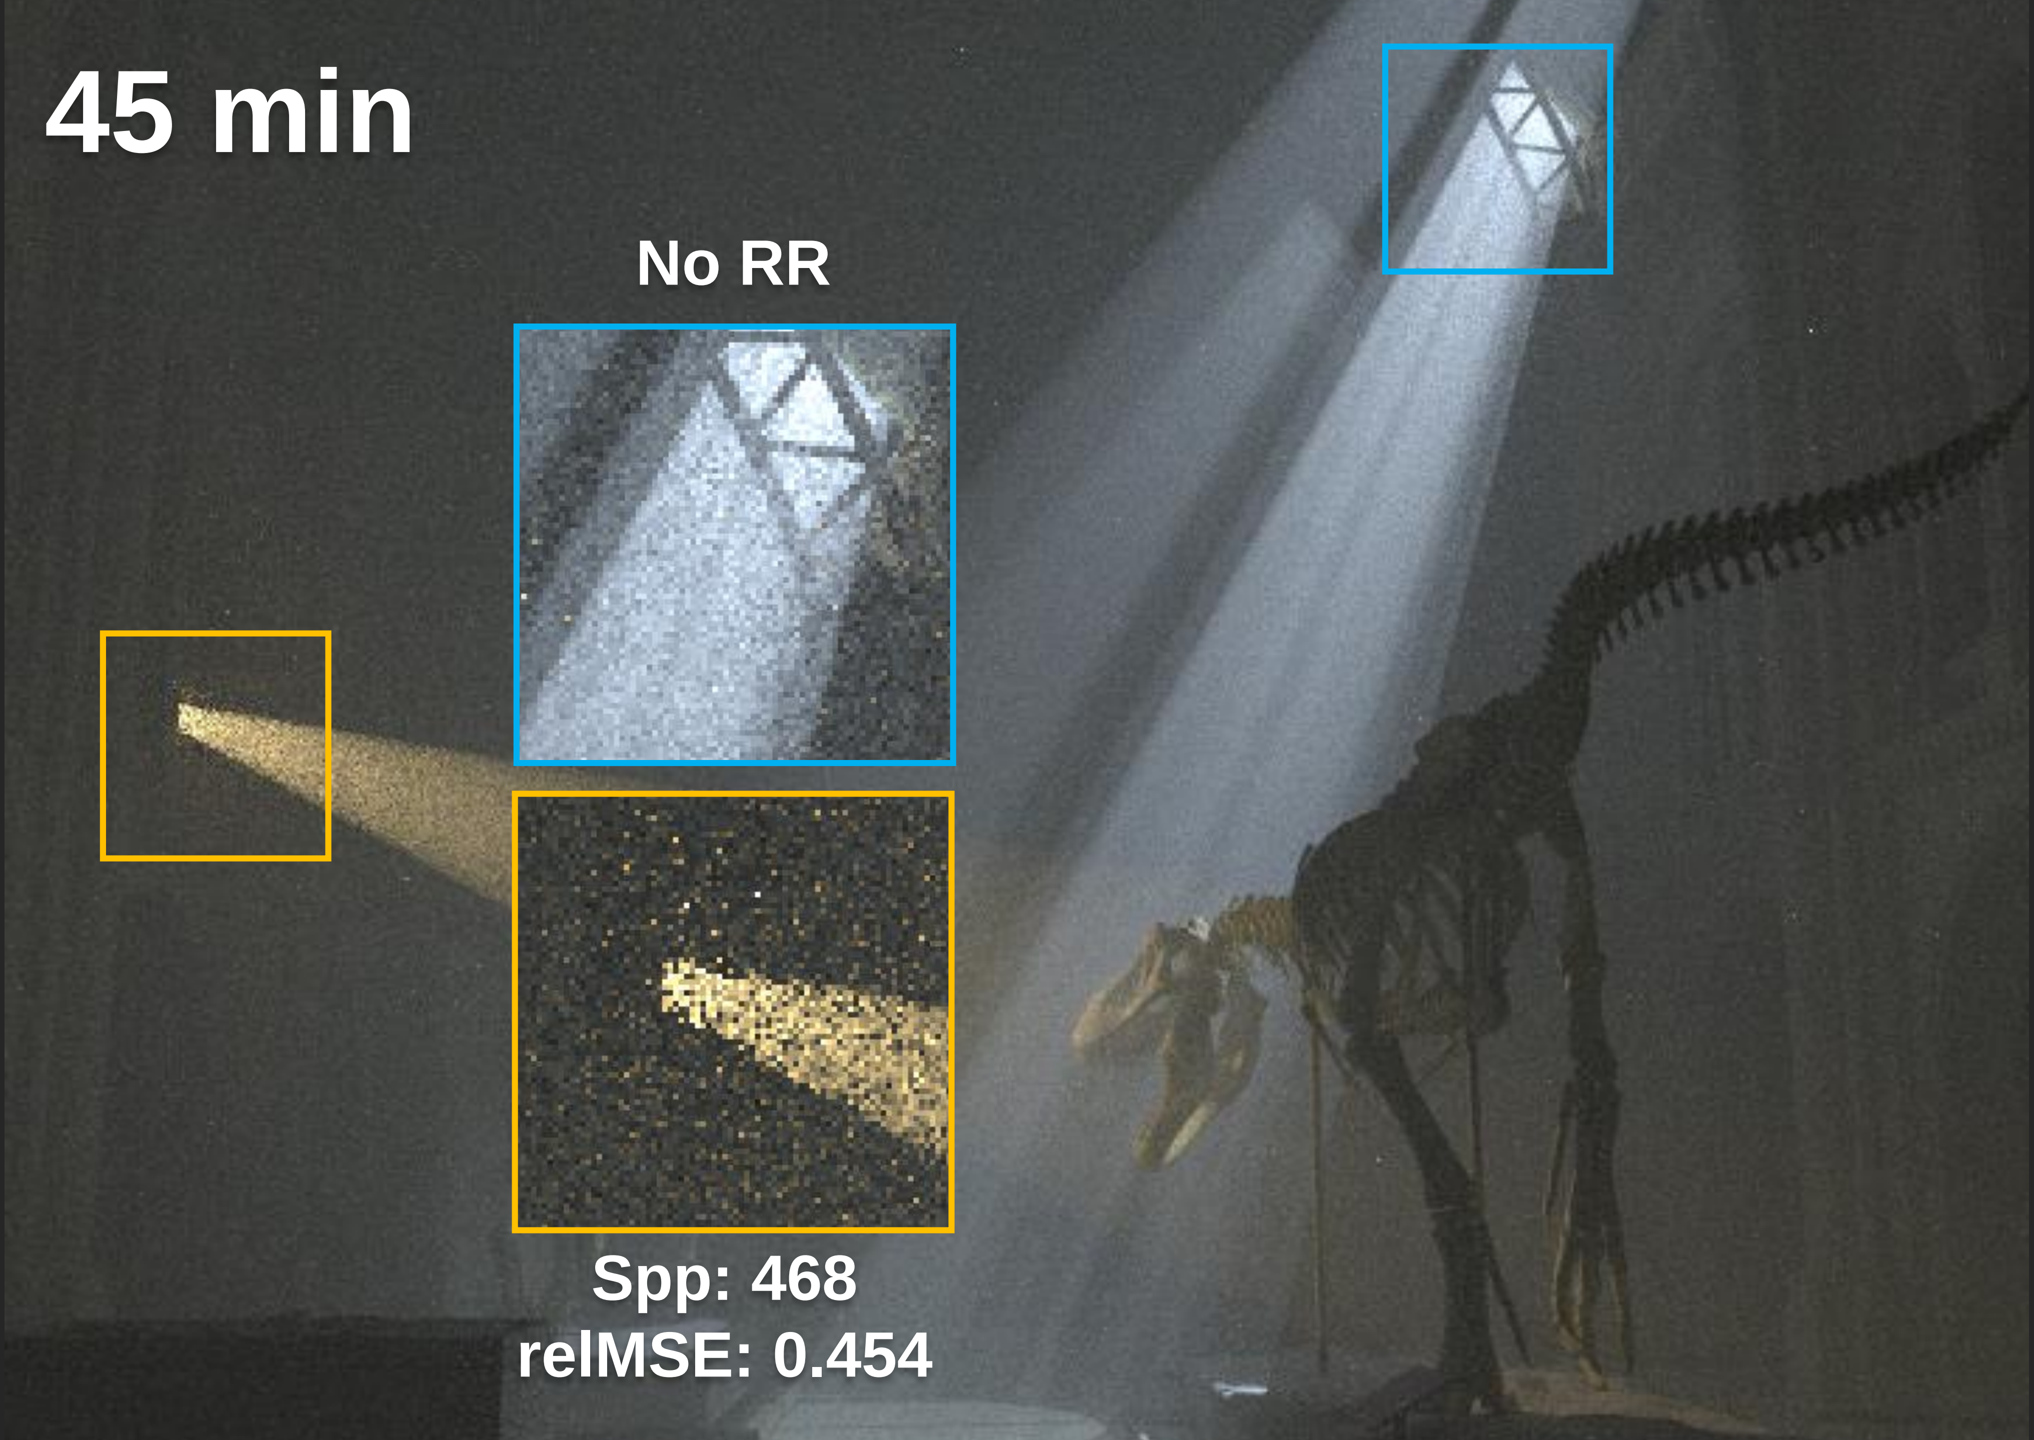
45 min
No RR
Spp: 468
relMSE: 0.454

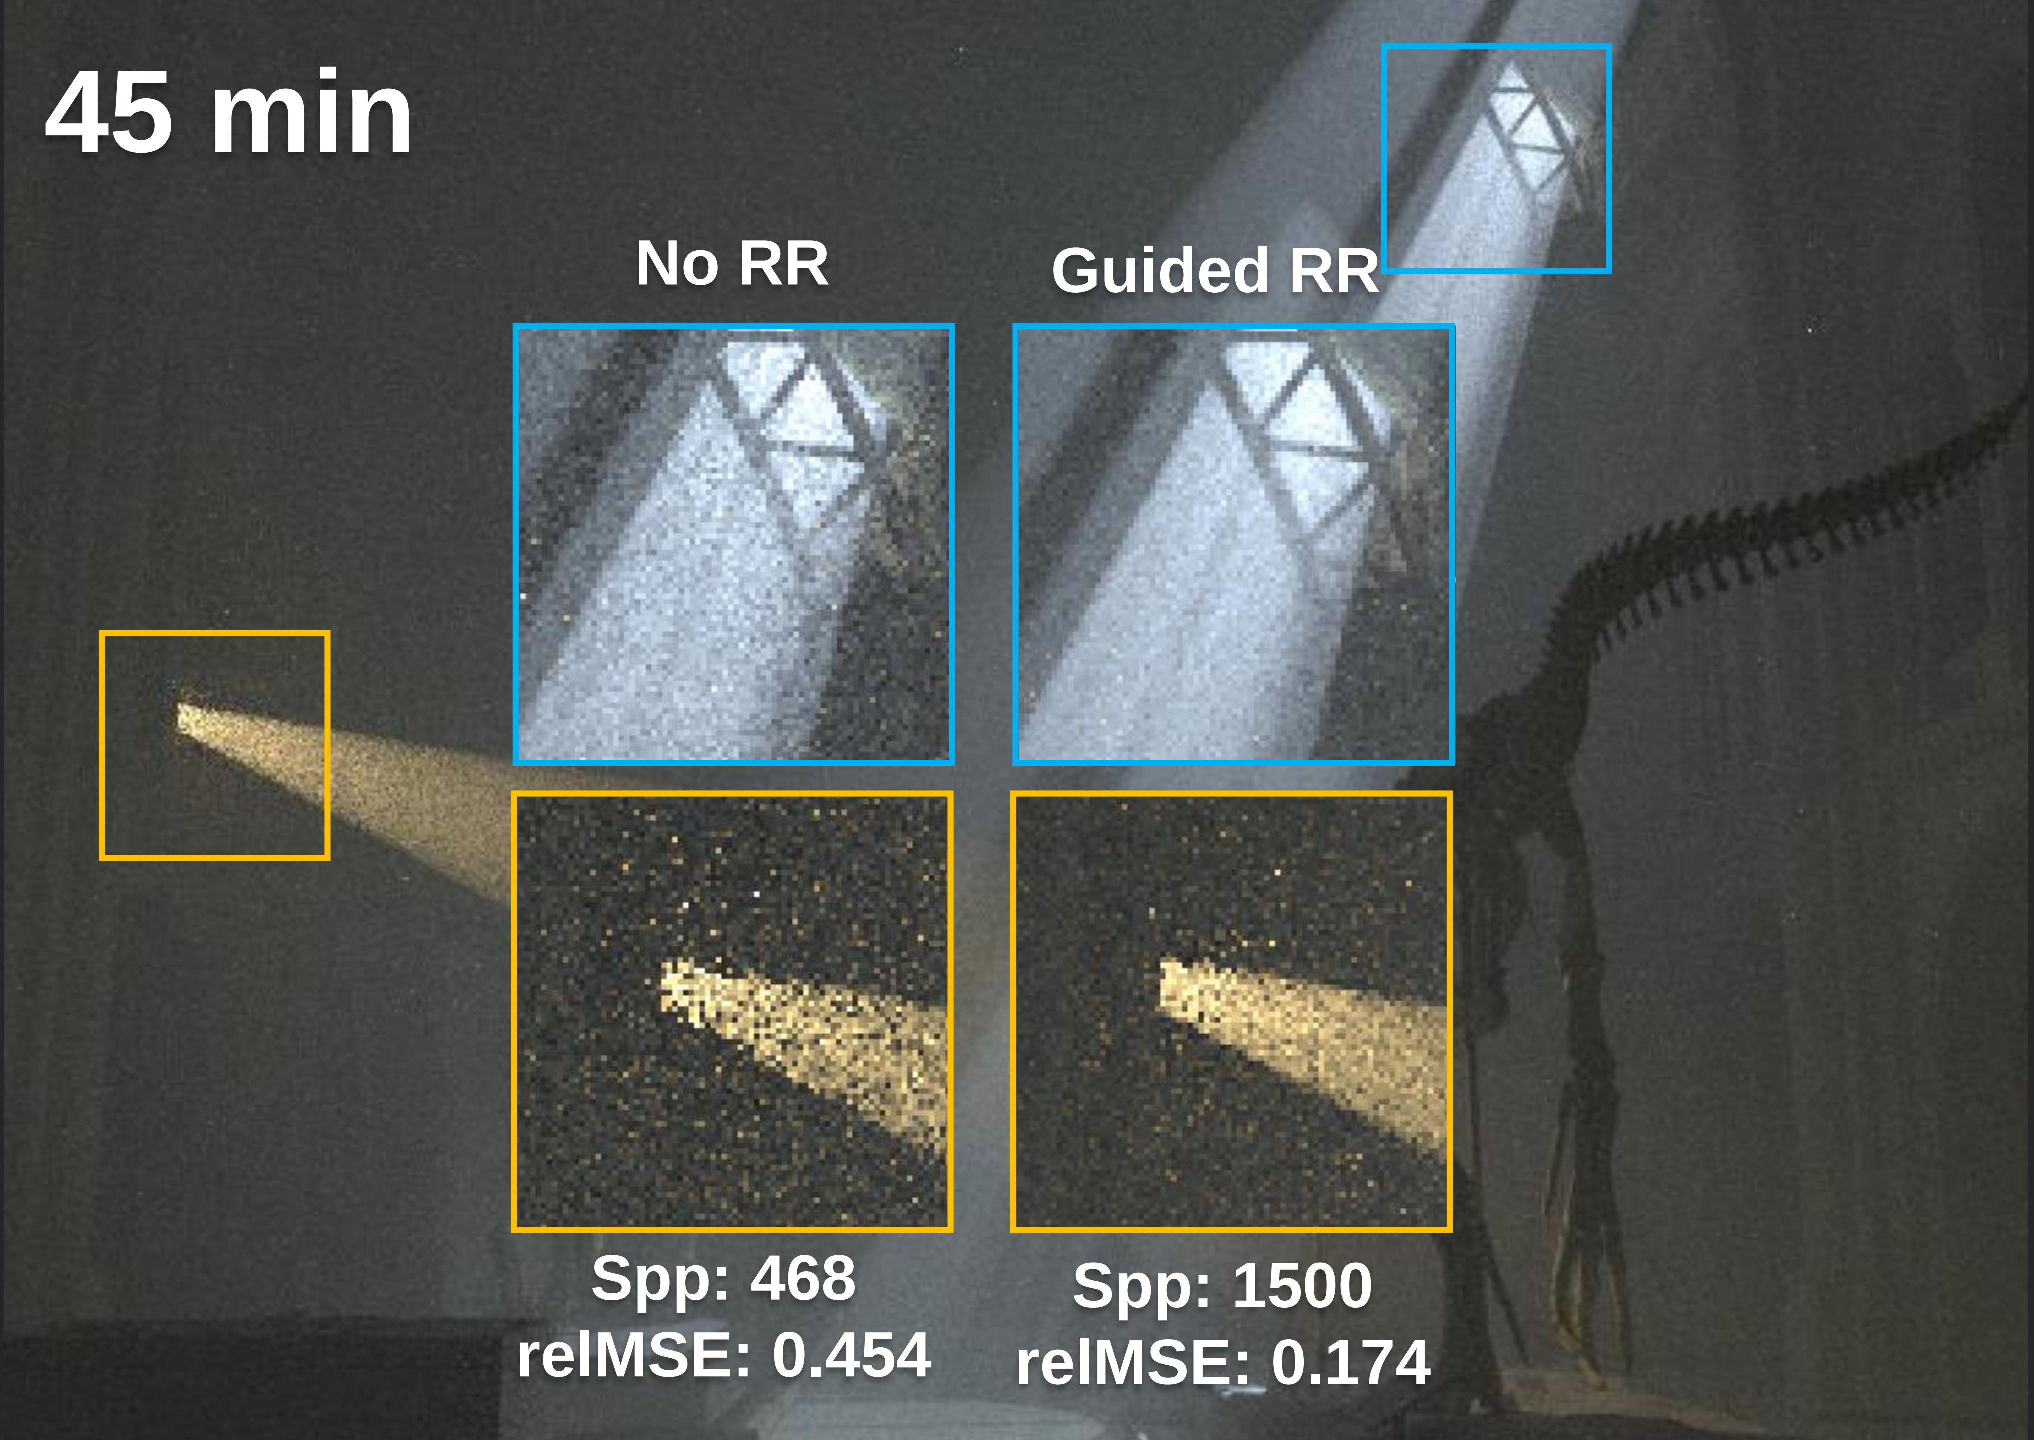
45 min
No RR
Guided RR
Spp: 468
relMSE: 0.454
Spp: 1500
relMSE: 0.174

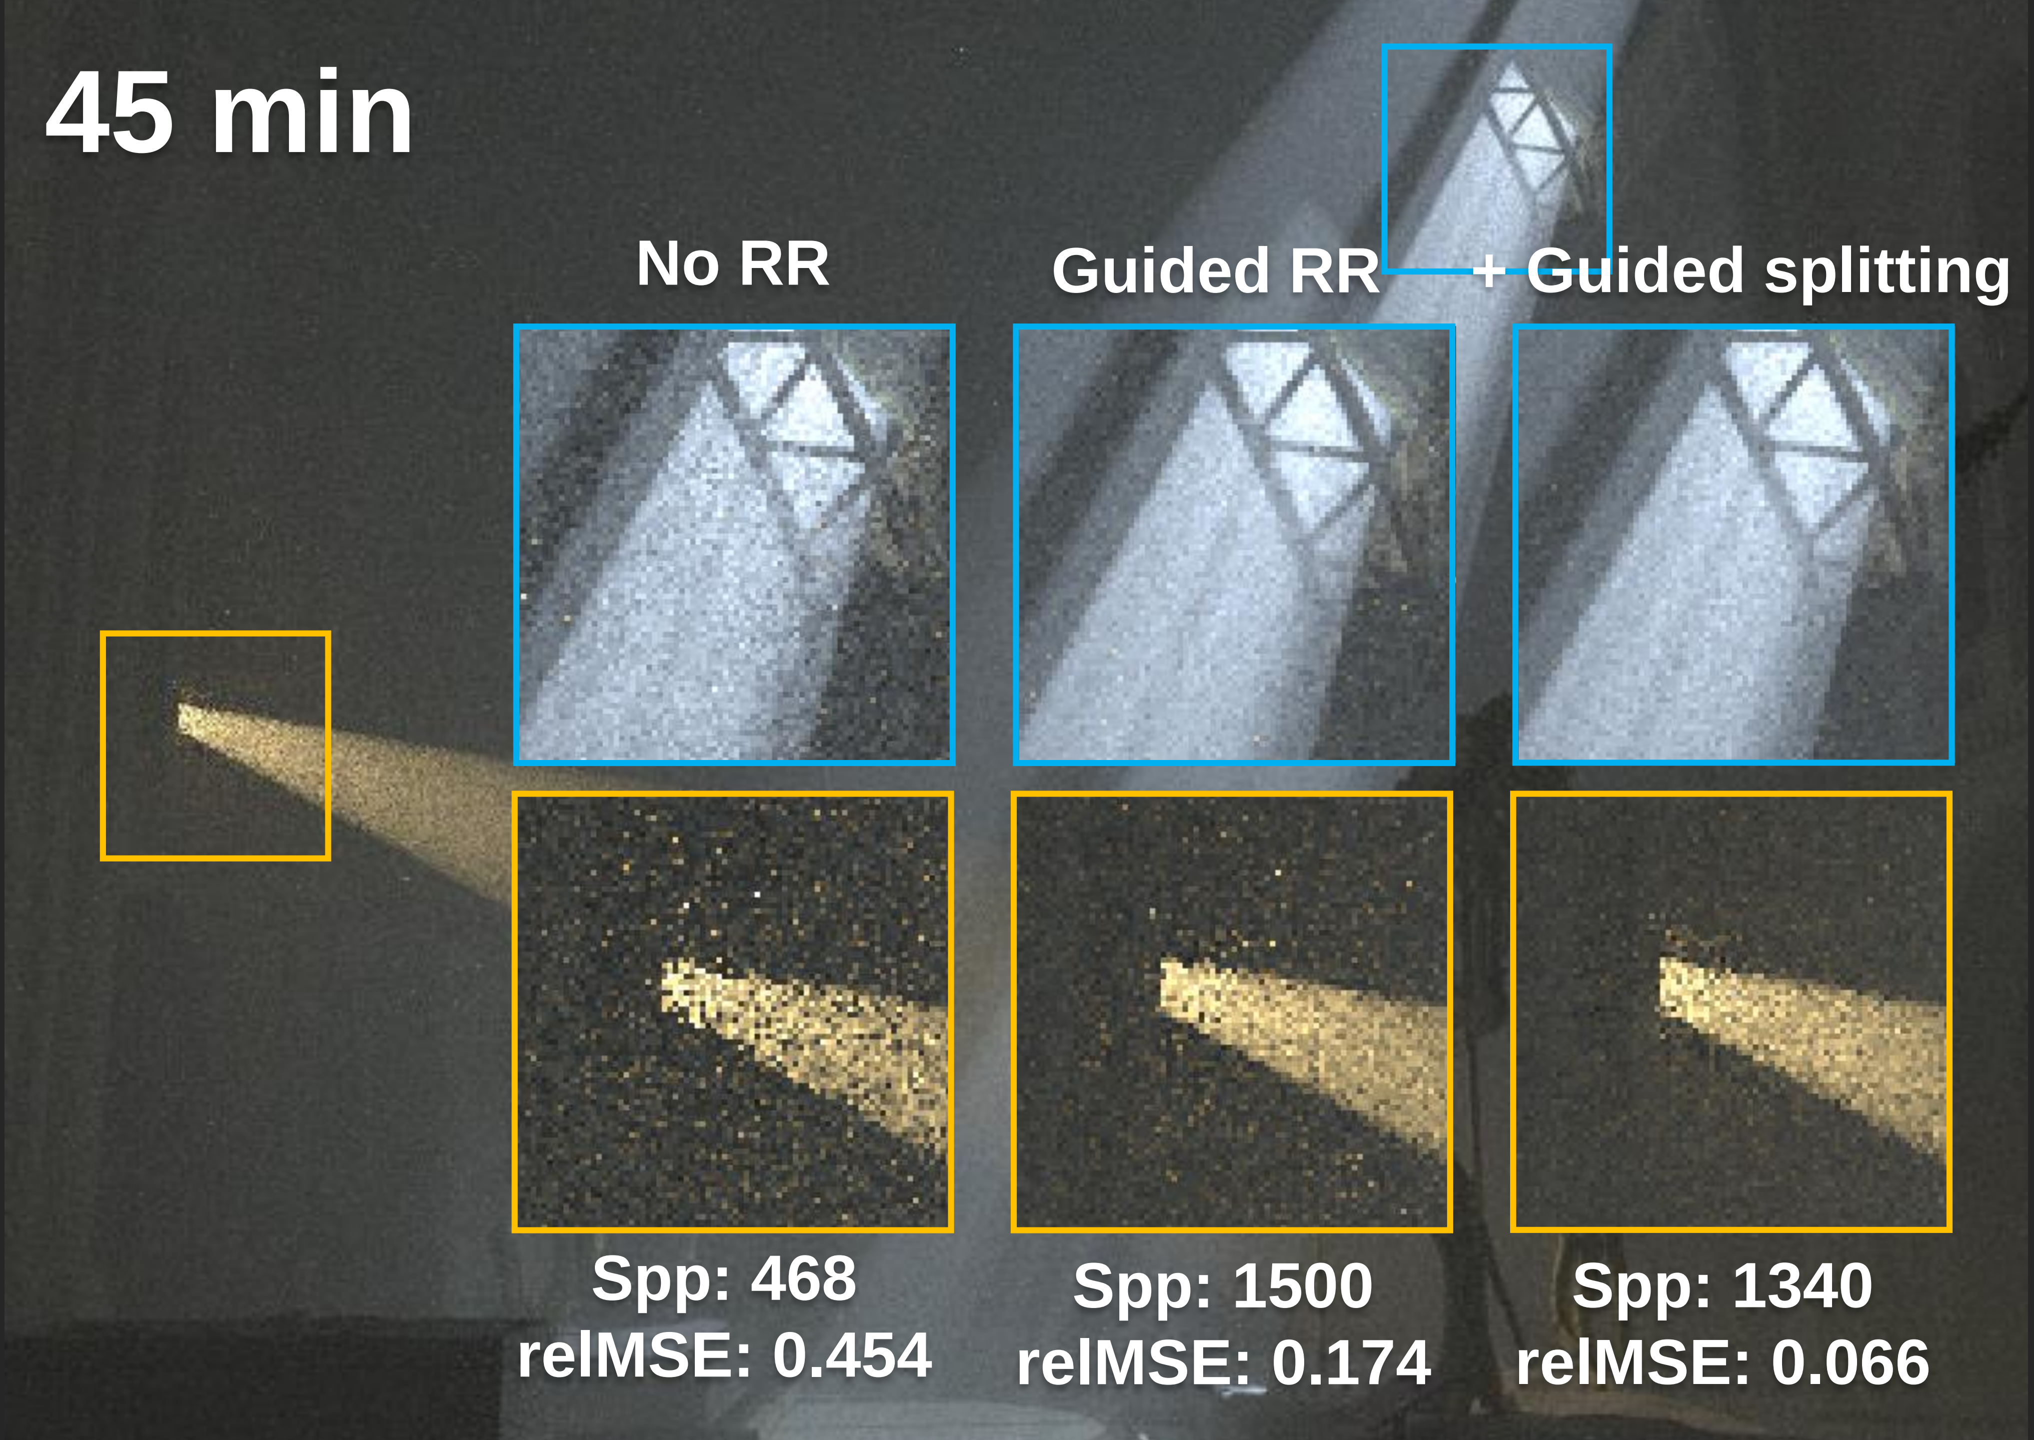

45 min
No RR
Guided RR
+ Guided splitting
Spp: 468
relMSE: 0.454
Spp: 1500
relMSE: 0.174
Spp: 1340
relMSE: 0.066

Guided RR

Spp: 1500
relMSE: 0.174

+ Guided splitting

Spp: 1340
relMSE: 0.066

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

No guiding

Time: 60 min

Spp: 10644

relMSE: 11.58

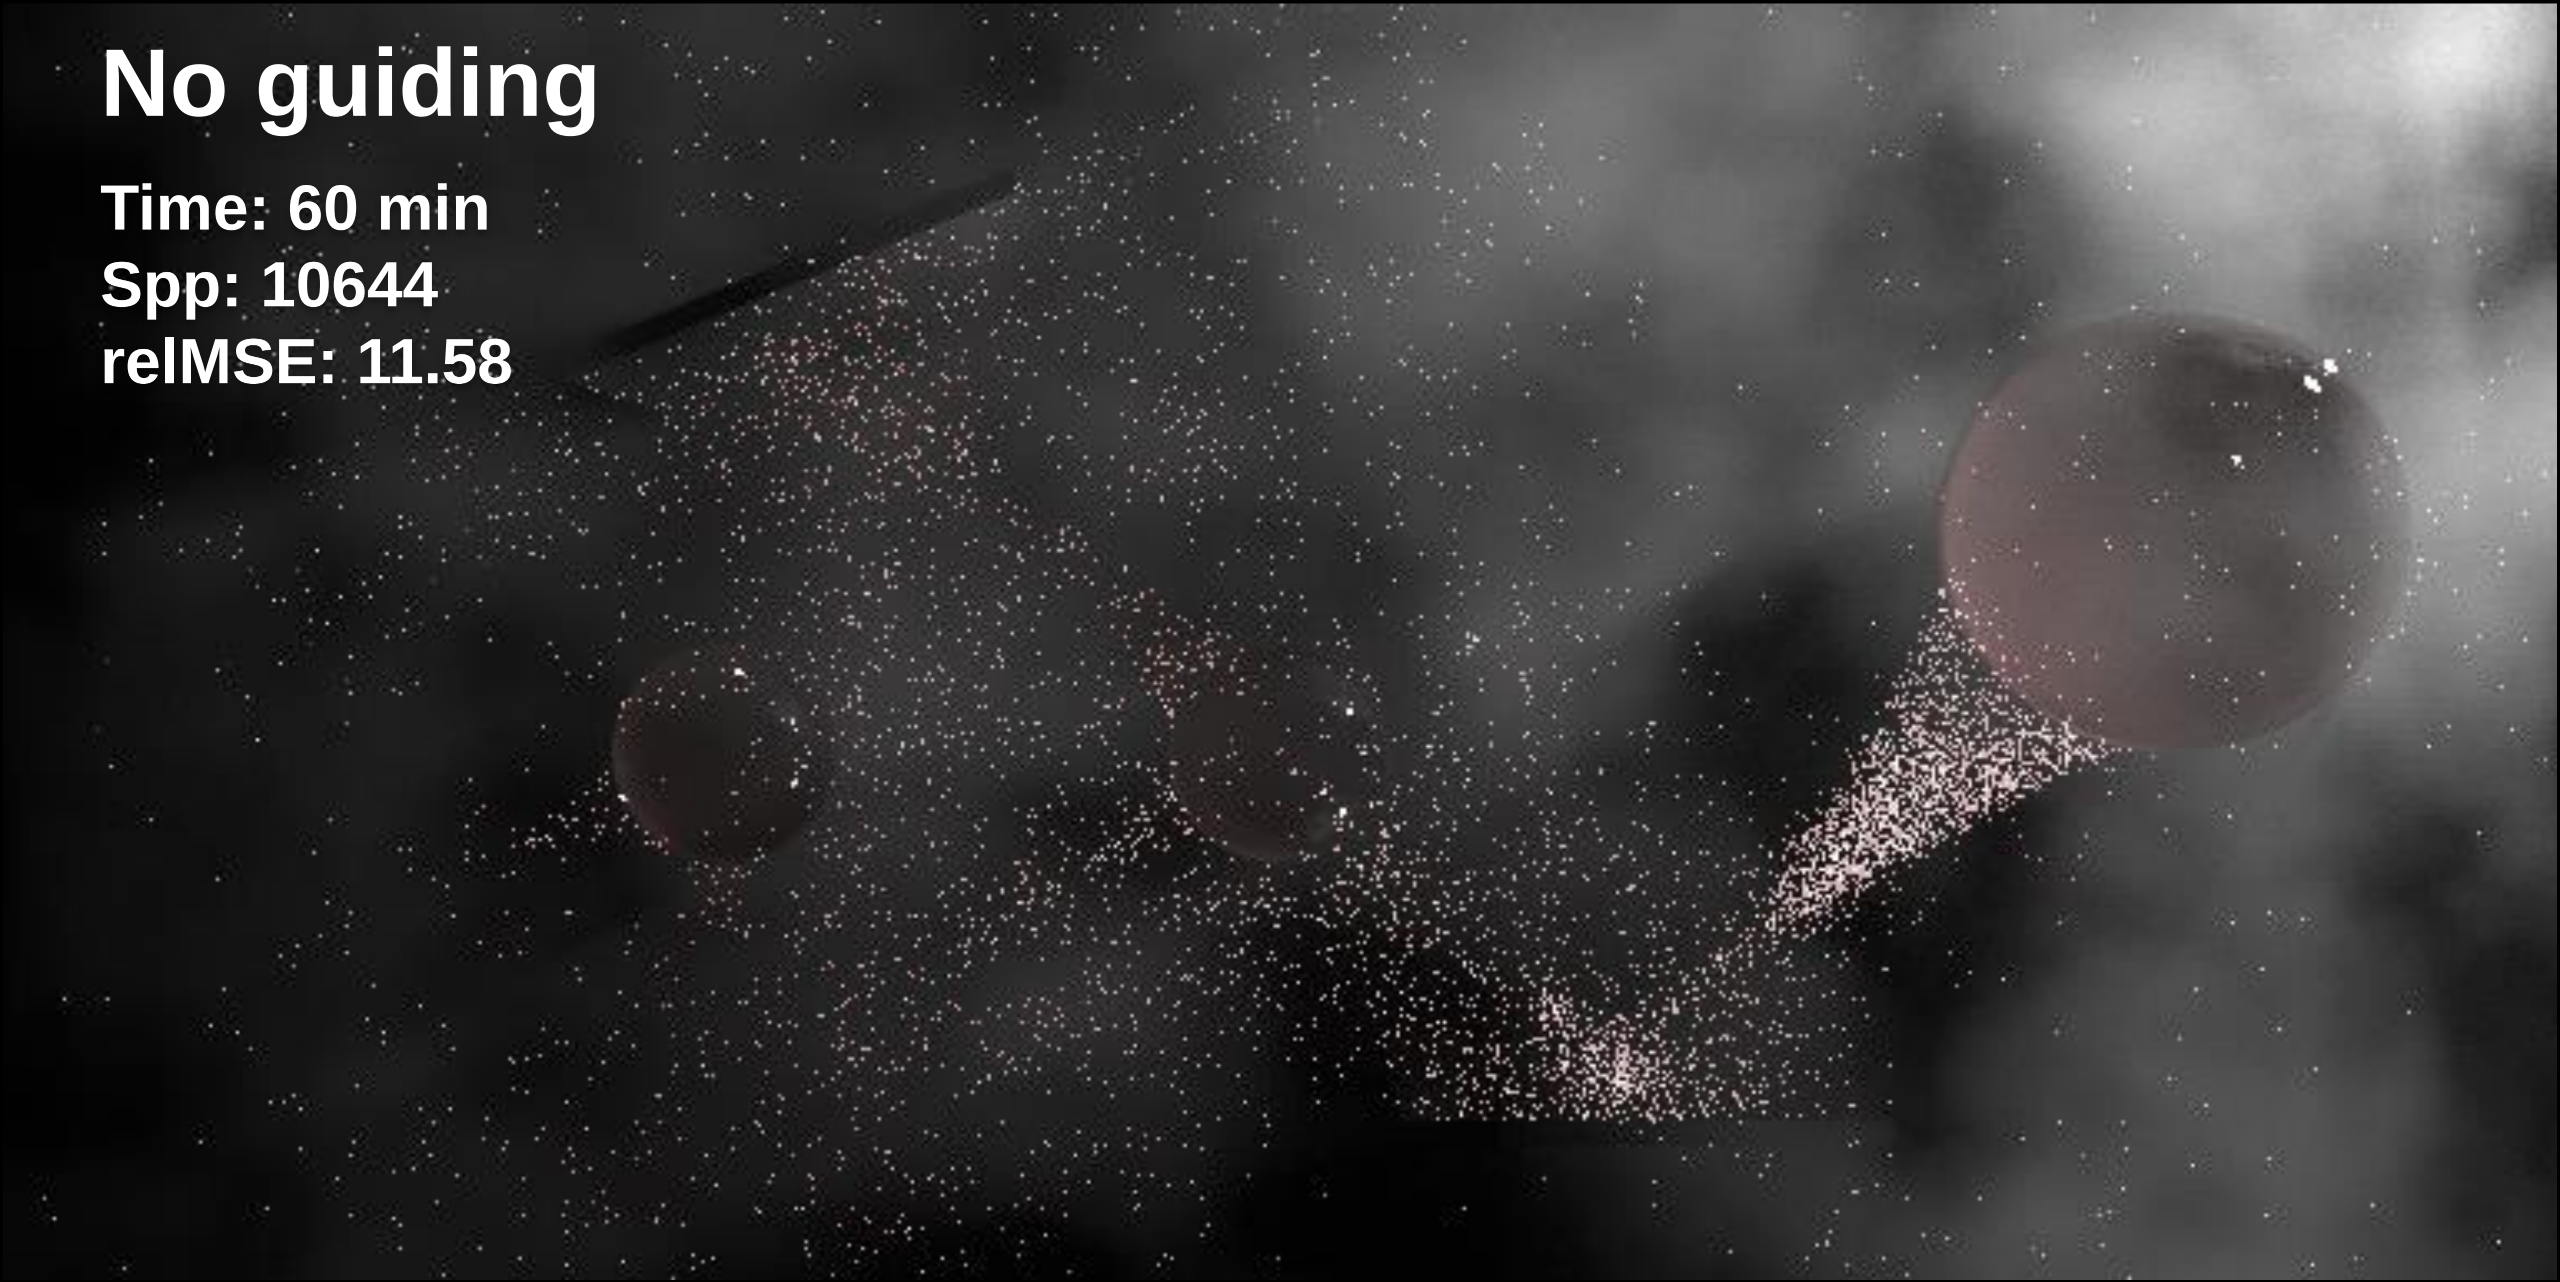

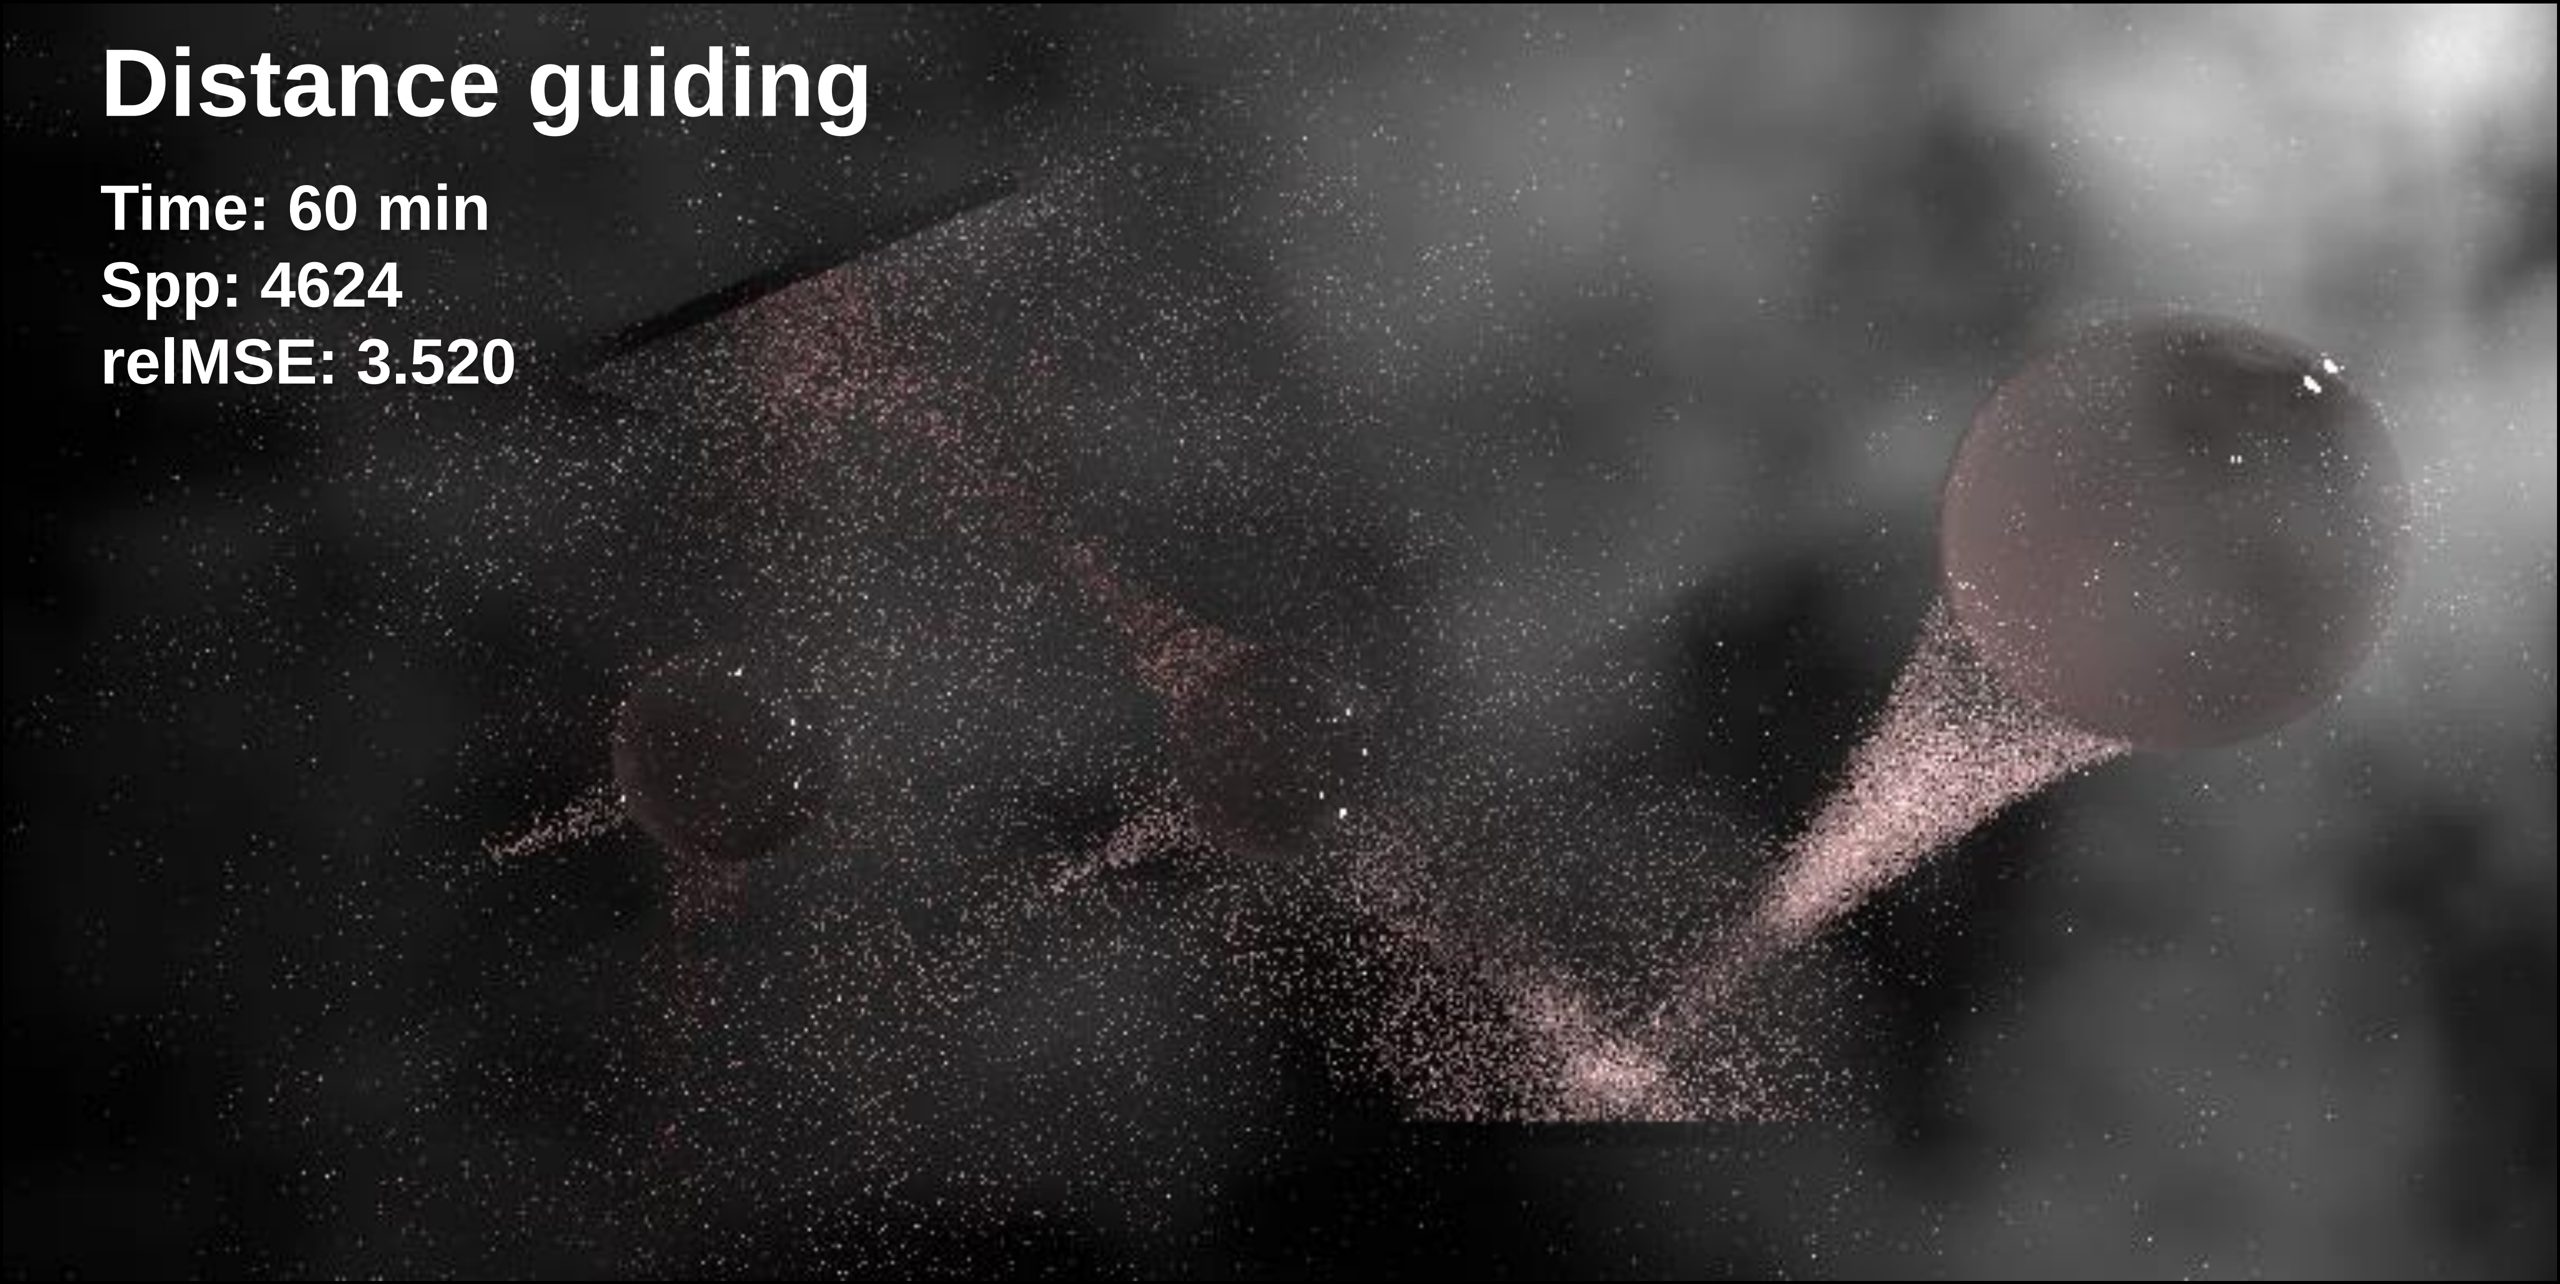
Distance guiding
Time: 60 min
Spp: 4624
relMSE: 3.520

Distance + directional guiding

Time: 60 min
Spp: 4448
relMSE: 0.468

Distance + directional guiding + GRRS

Time: 60 min
Spp: 3796
relMSE: 0.321

OPEN PROBLEMS AND LIMITATIONS

Open Challenges to make it bullet proof

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

SPATIAL CACHE STRUCTURE

- Naive approach to define the resolution:
  - Heuristic based on sample numbers
  - Takes time or many samples to model/separate fine features (e.g. thin shafts or caustics) (PPG by [Mueller2017] has the same problem)
  - Influences the performance of some sampling methods (e.g. distance)
- **Ideal structure should adjust to the light transport characteristics**

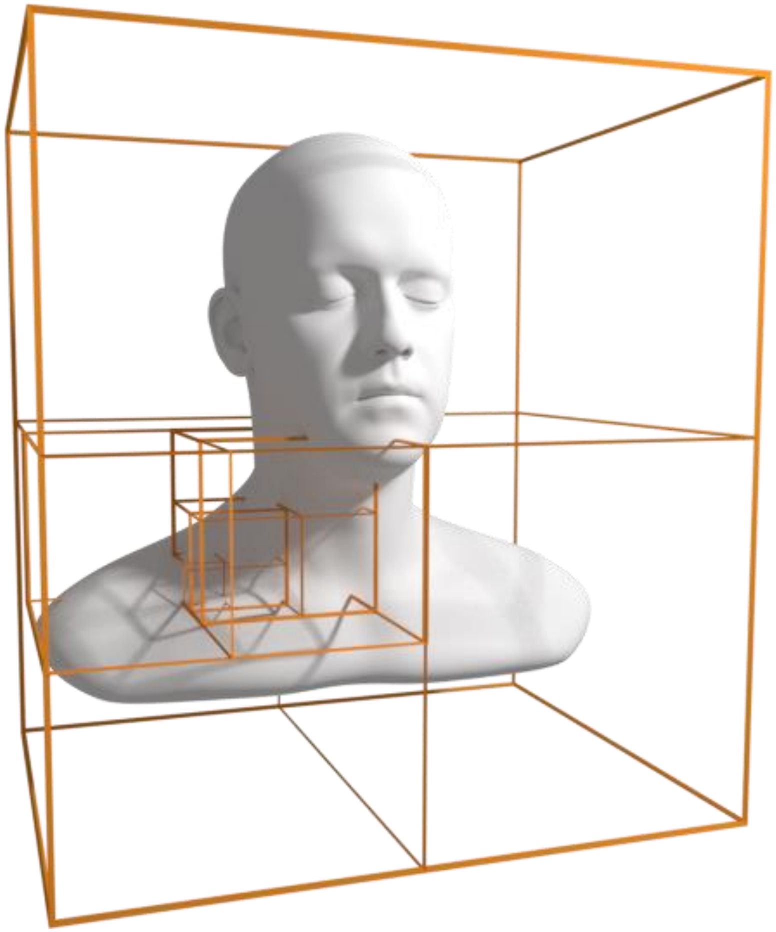

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

INSUFFICIENT CACHE SIZES

- Shared problem with other 3D caching based guiding approaches (e.g. [Vorba2014], [Mueller2017], …)
- By merging the samples of a spatial cache we blur the directional distribution
- Can lead to incorrect estimates of L and L_i

PRODUCTION CHALLENGES

How can we get good guiding estimates?

VMM PHASE FUNCTION FITTING: PRE-PROCESSING STEP

- Pre-processing step:
  - Fitting a VMM for each phase function

- Fitting up to $K = 4$ components
  - Details in the course notes

- **Open Challenge:**
  - **Procedural phase functions or procedural mixtures?**
  - **How to deal with changing mean cosines (roughening/mollification)?**

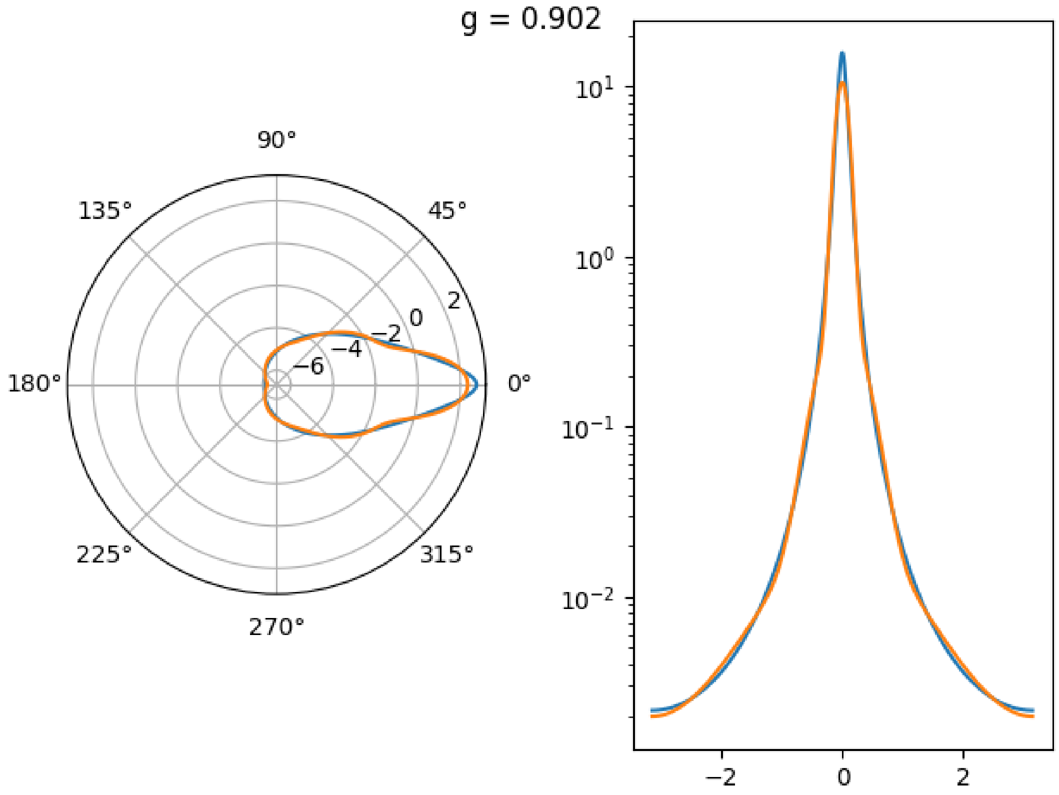

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VMM PHASE FUNCTION FITTING: PRE-PROCESSING STEP

- Pre-processing step:
  - Fitting a VMM for each phase function

- Fitting up to $K = 4$ components
  - Details in the course notes

- **Open Challenge:**
  - **Procedural phase functions or procedural mixtures?**
  - **How to deal with changing mean cosines (roughening/mollification)?**

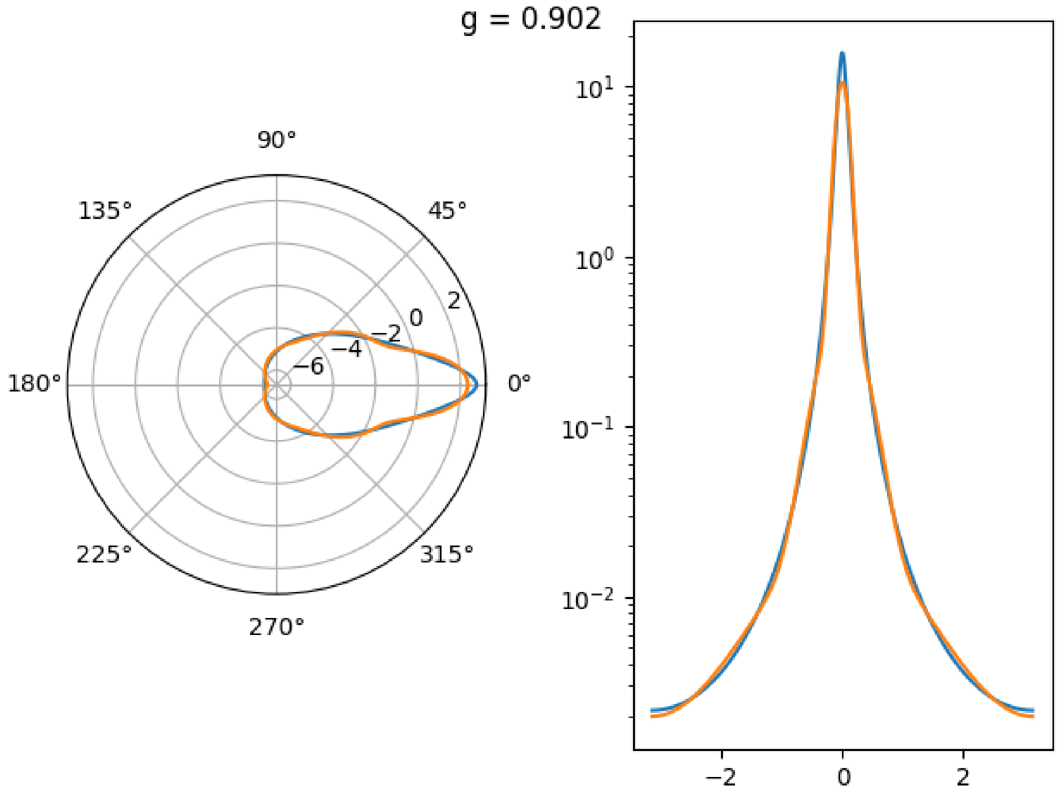

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

FITTING/LEARNING THE INCIDENT RADIANCE MIXTURES

Photon-based pre-processing ([Herholz2019][Vorba2014])

Pros:

- Photons directly represents the light transport
- Spatial distribution corresponds to important features (light shaft)
- Number of traced photons can be fixed
- No additional fitting during rendering

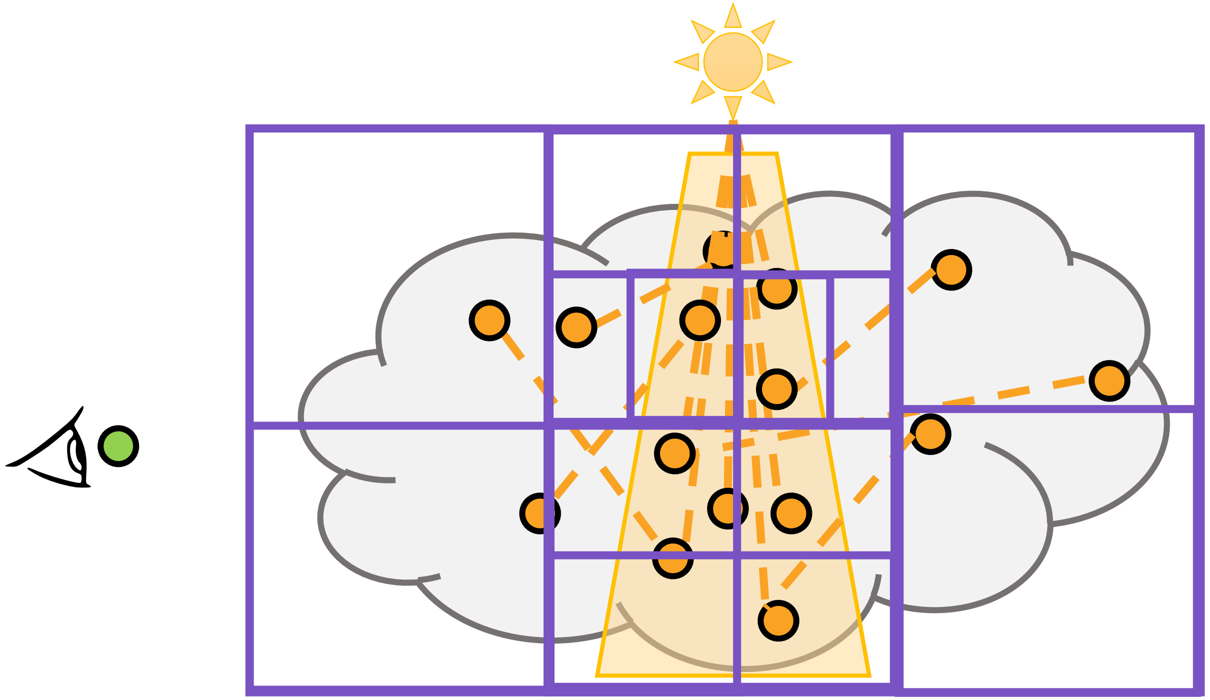

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

FITTING/LEARNING THE INCIDENT RADIANCE MIXTURES

Photon-based pre-processing ([Herholz2019][Vorba2014])

Cons:

- Pre-processing step:
  - Long time to first render iteration
- Complex scenes need bidirectional learning:
  - Ping-Pong style [Vorba2014]
- It is not the ideal solution for artists in production

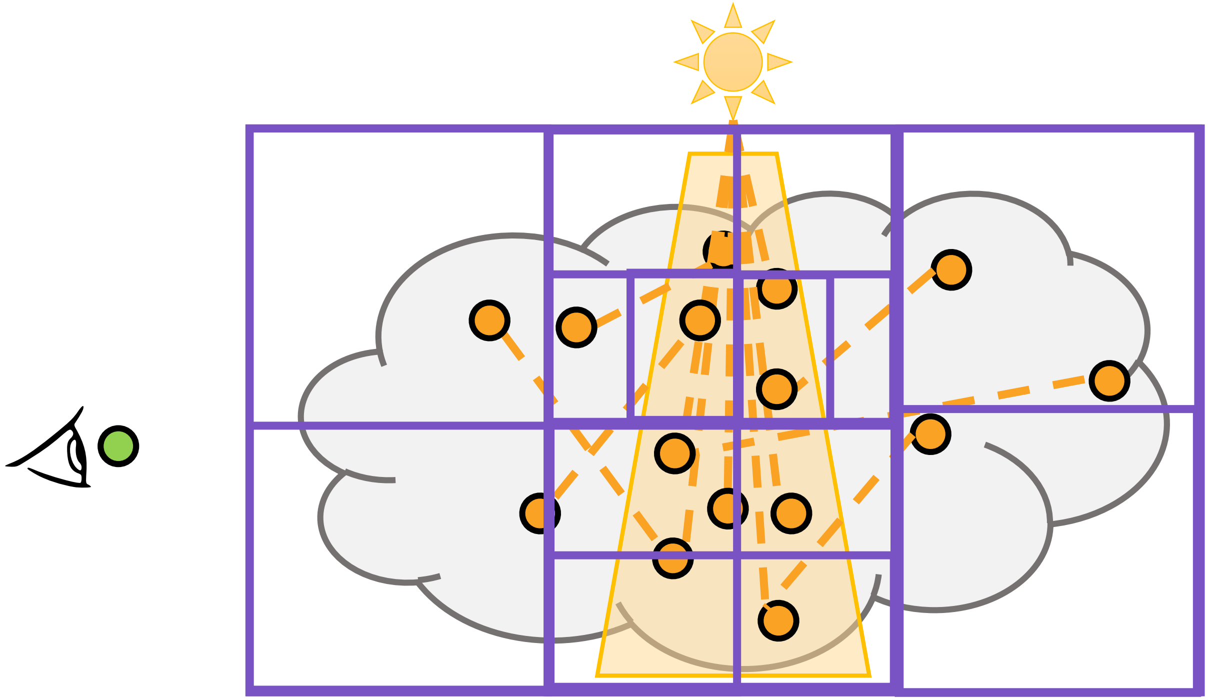

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

FITTING/LEARNING THE INCIDENT RADIANCE MIXTURES

Progressive learning (PPG-style [Mueller2017])

Pros:

- **First experiments show promising results**
- No pre-processing
- Refines spatial and directional distributions in each iteration
- Sample data gets more reliable

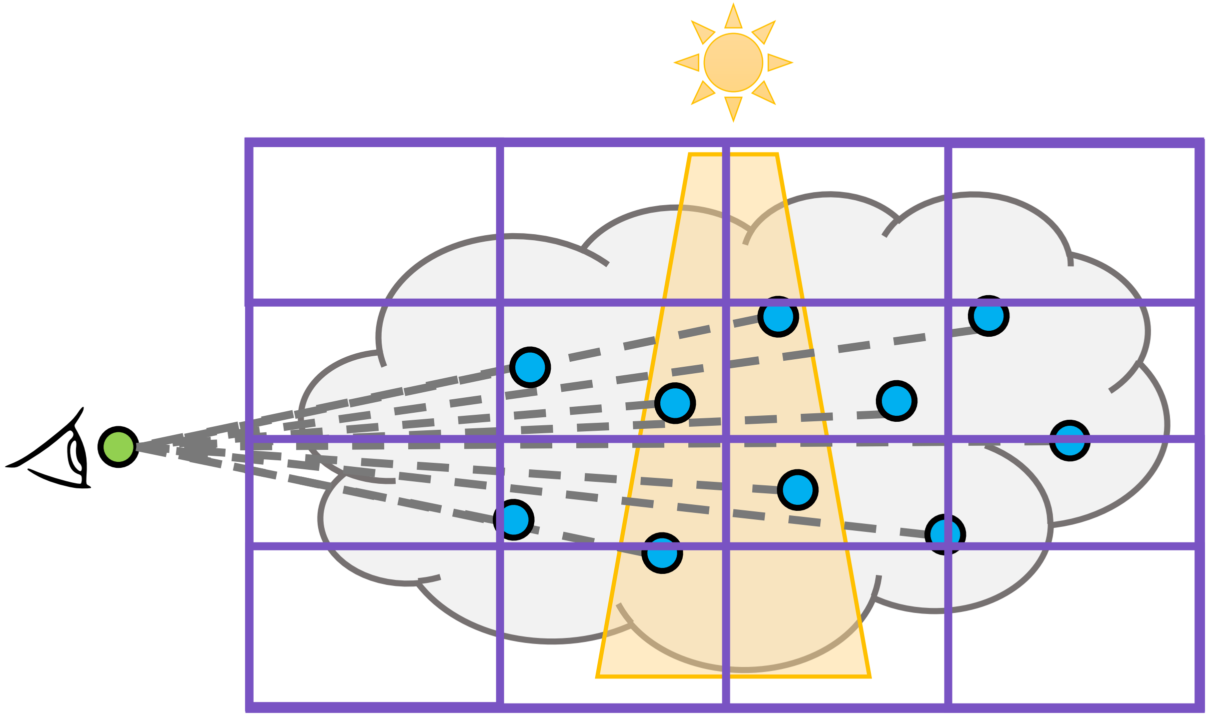

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

FITTING/LEARNING THE INCIDENT RADIANCE MIXTURES

Progressive learning (PPG-style [Mueller2017])

Cons (open challenges):

- Sample count grows exponential:
  - Memory and fitting time increases

  - **Shorter update rates ?**
  - **Online fitting for mixtures?**

- Spatial structure adapts slowly to LT:
  - Important for distance sampling

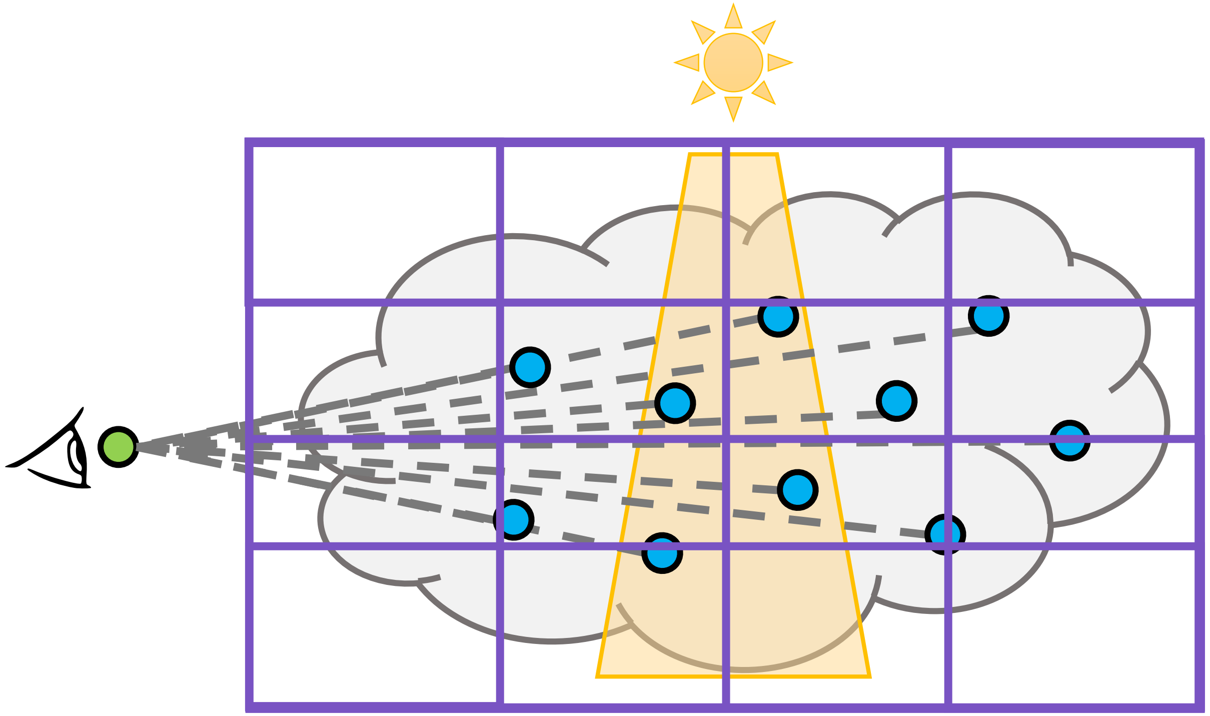

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

CONCLUSION

- **Even approximate local sampling decisions lead to a good approximation of the globally optimal guiding distribution** (and thus significantly reducing MC variance)
- **Converges to optimal ZV estimator in the hypothetical limit** (i.e., if the approximations were perfect)
- **Solely based one guiding structure for all decisions** (incident radiance VMMs)

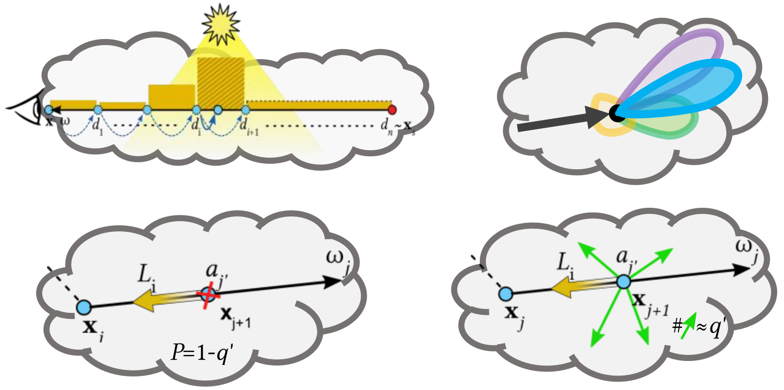

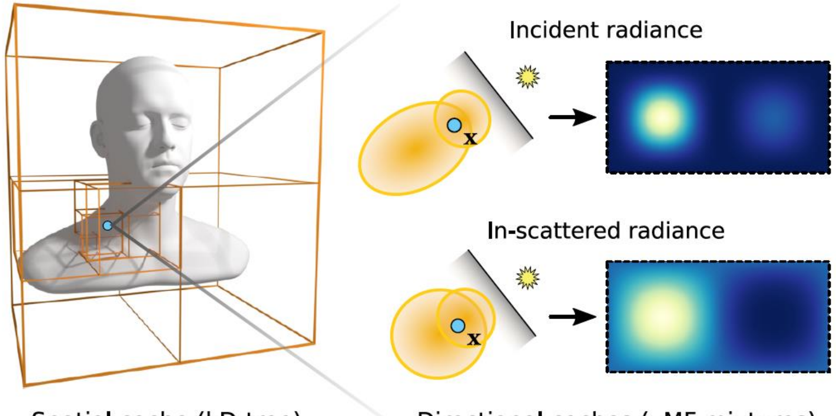

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

THANK YOU

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

REFERENCES

- [Fong2017]: “Production volume rendering”
- [Novak2018]: “Monte Carlo methods for volumetric light transport simulation”
- [Kulla2012]: “Importance sampling techniques for path tracing in participating media”
- [Georgiev2012]: “importance sampling of low-order volumetric scattering”
- [Krivanek2014]: “A zero-variance-based sampling scheme for Monte Carlo subsurface scattering”
- [Meng2016]: “Improving the Dwivedi sampling scheme”
- [Vorba2014]: “Online learning of parametric mixture models for light transport simulation”
- [Vorba2016]: “Adjoint-driven Russian roulette and splitting in light transport simulation”
- [Herholz2016]: “Product importance sampling for light transport path guiding”
- [Koerner2016]: “Subdivision next-event estimation for path-traced subsurface scattering”
- [Mueller2017]: “Practical path guiding for efficient light-transport simulation”
- [Hoogenboom2008]: “Zero-varianceMonte Carlo schemes revisited”
- [Pegoraro2008]:” Sequential Monte Carlo integration in low-anisotropy participating media”
- [Bashford2012]: “A significance cache for accelerating global illumination”

ADJOINT ESTIMATE ACCURACY

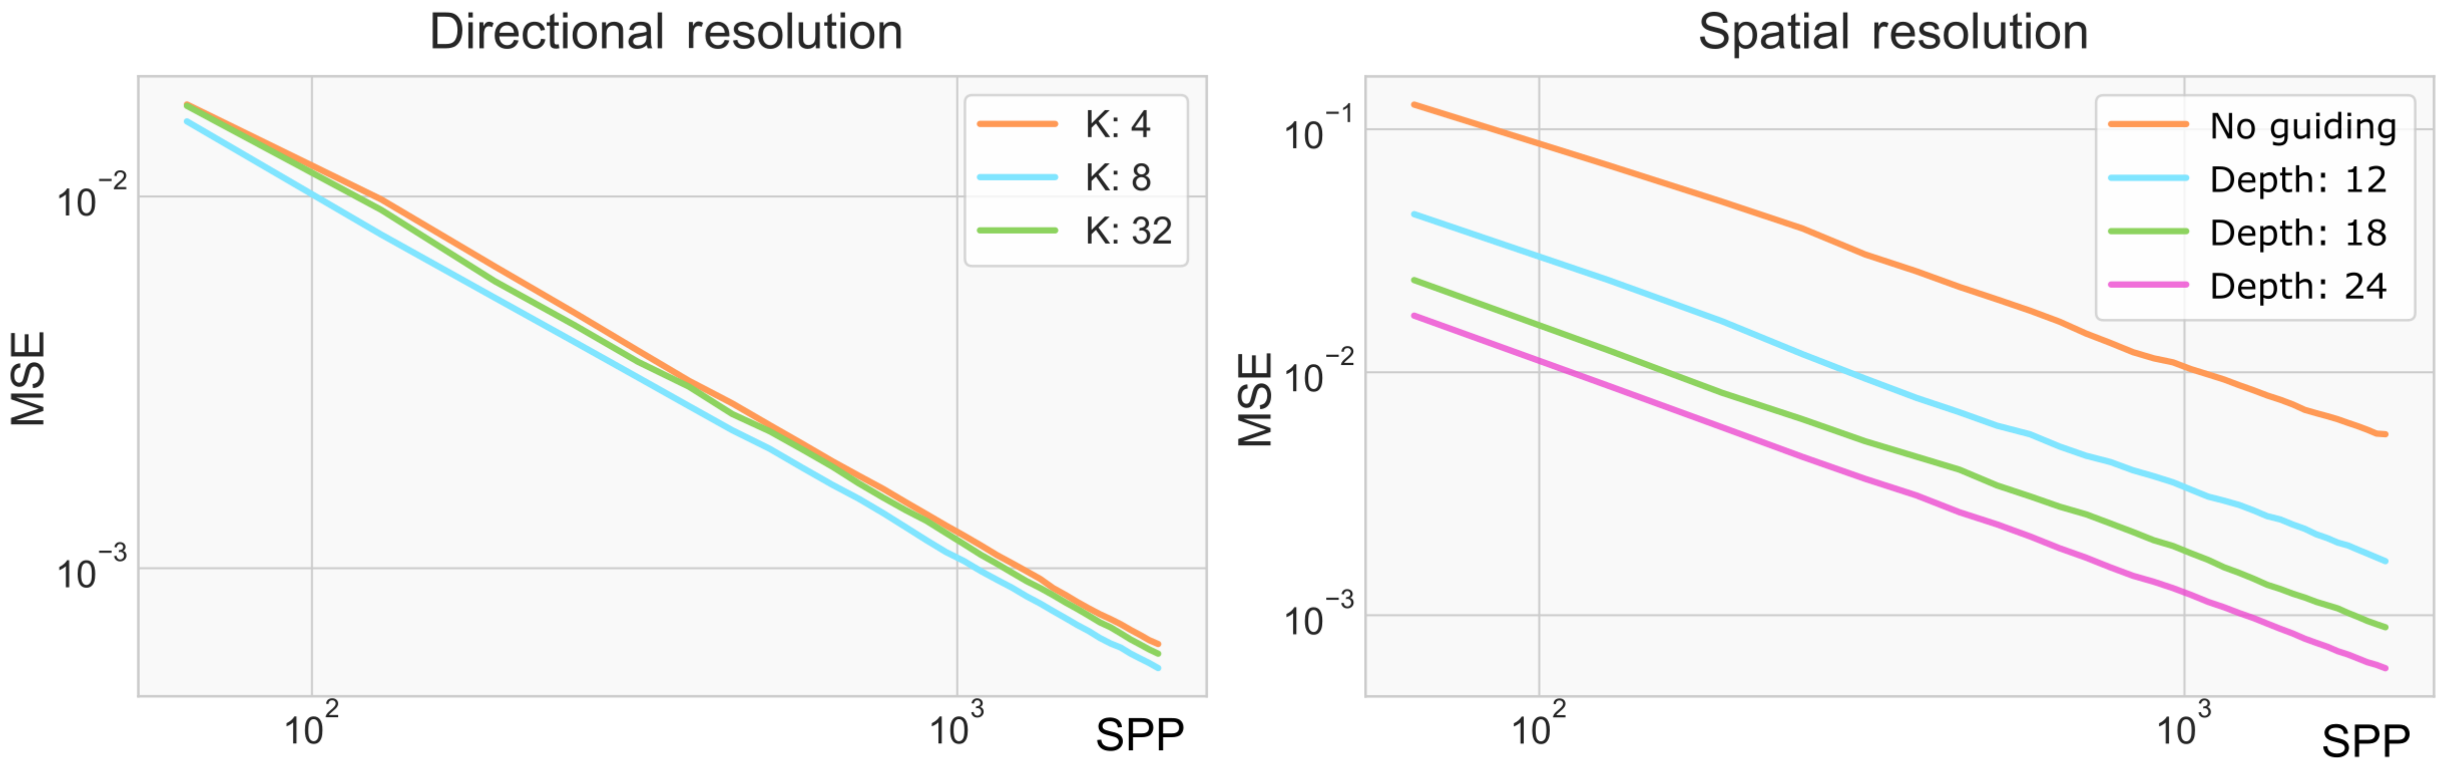

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

COMPARING AGAINST EQUIANGULAR SAMPLING

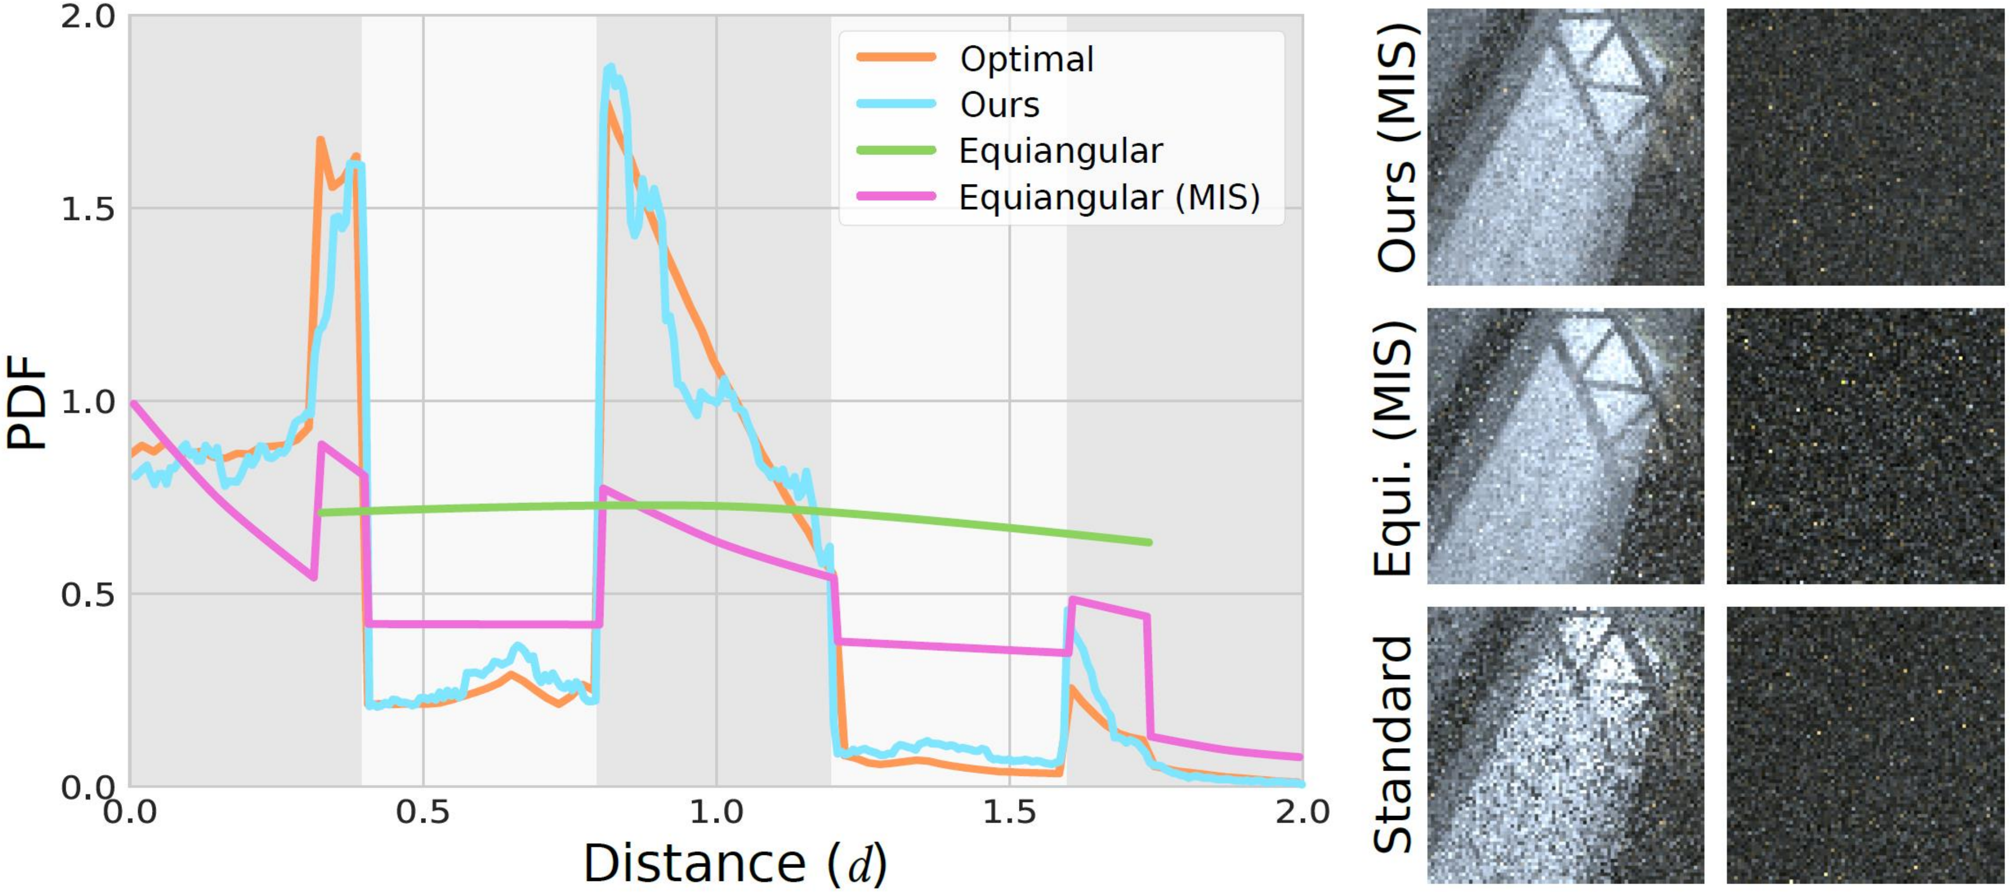

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

BUMPY SPHERE

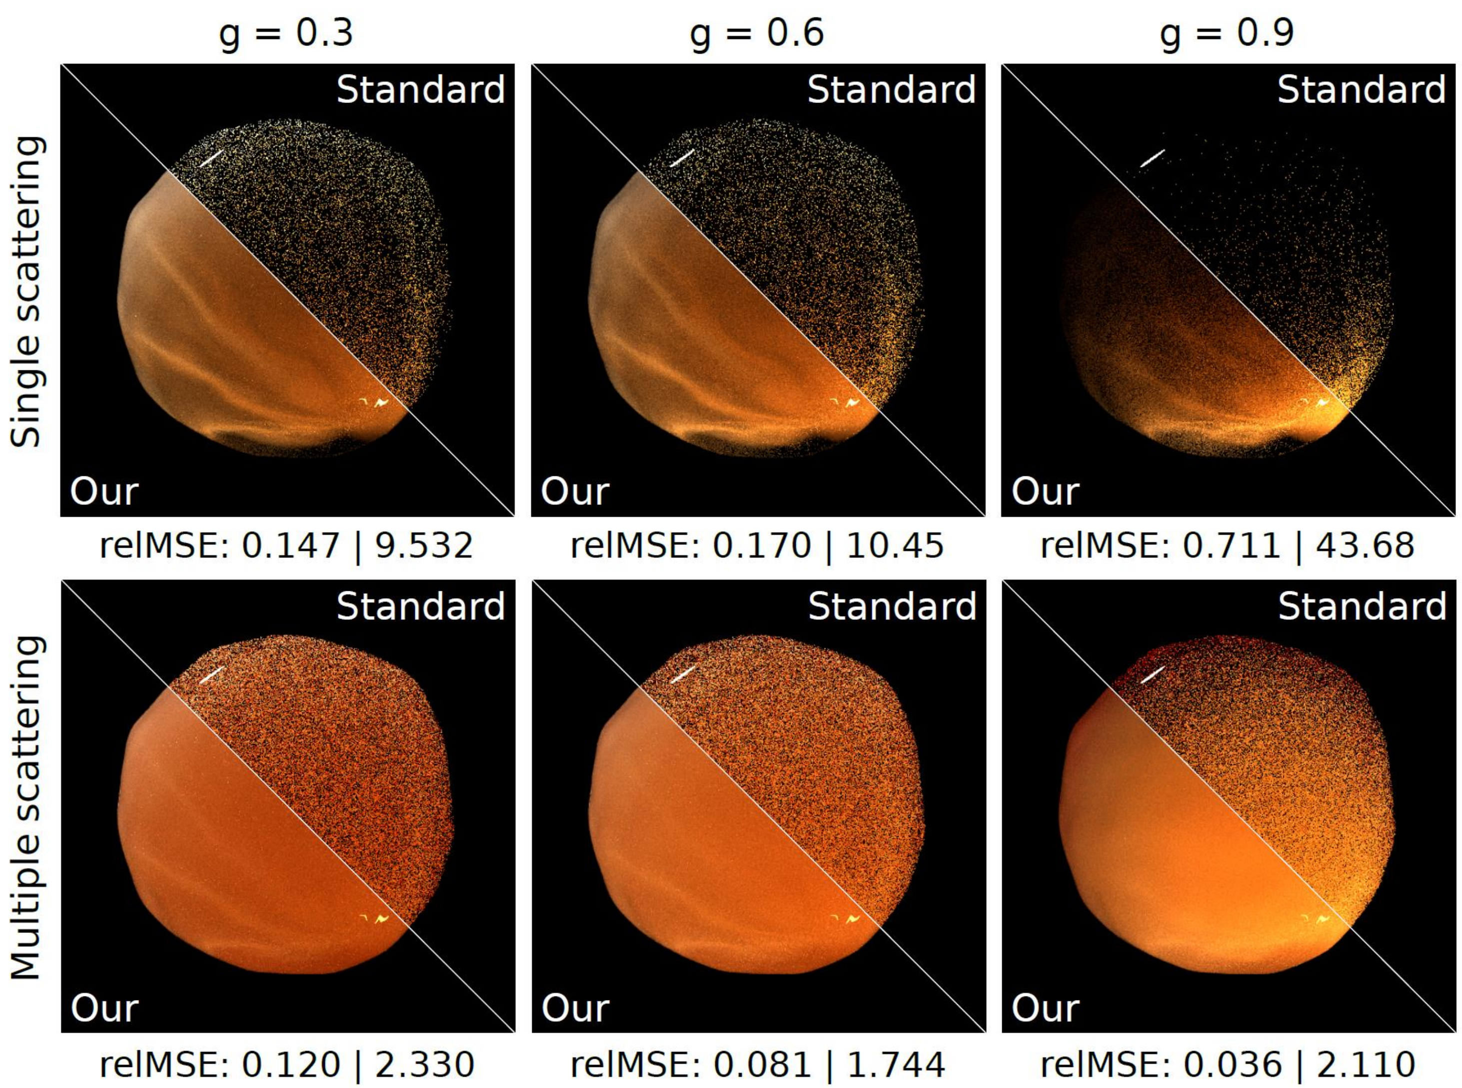

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

INCREMENTAL GUIDED DISTANCE SAMPLING

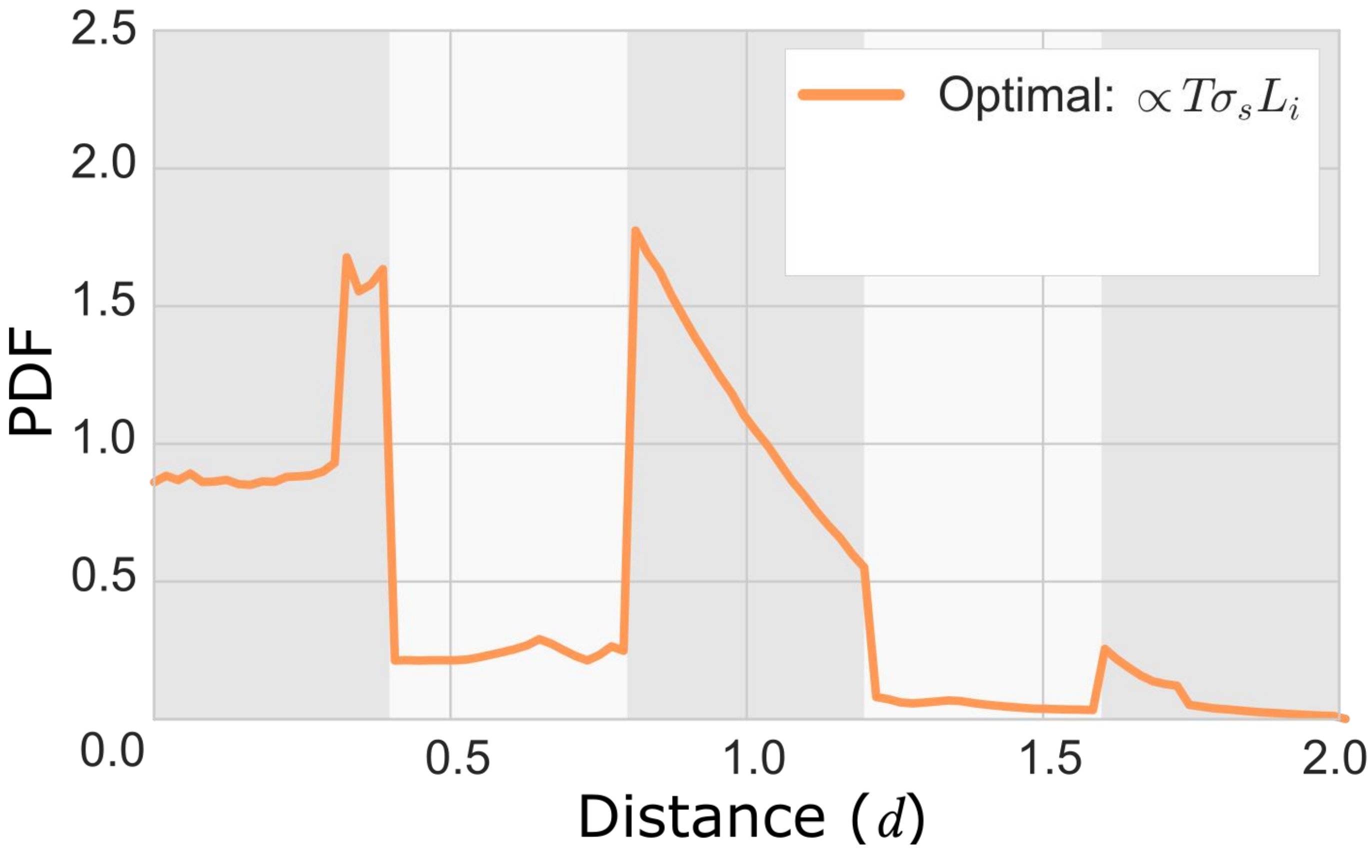

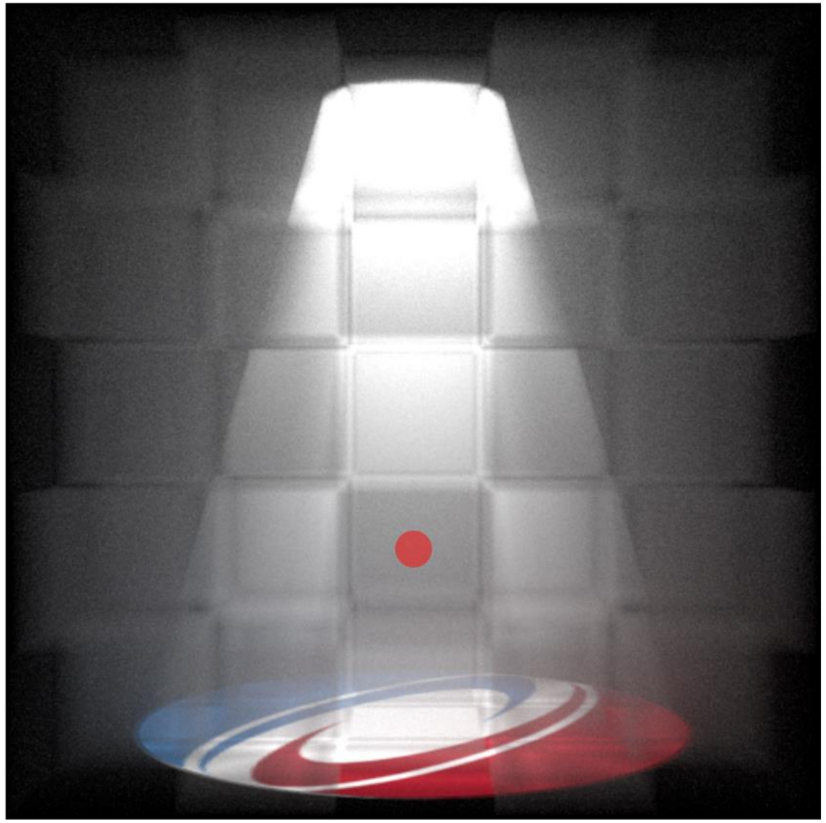

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

INCREMENTAL GUIDED DISTANCE SAMPLING

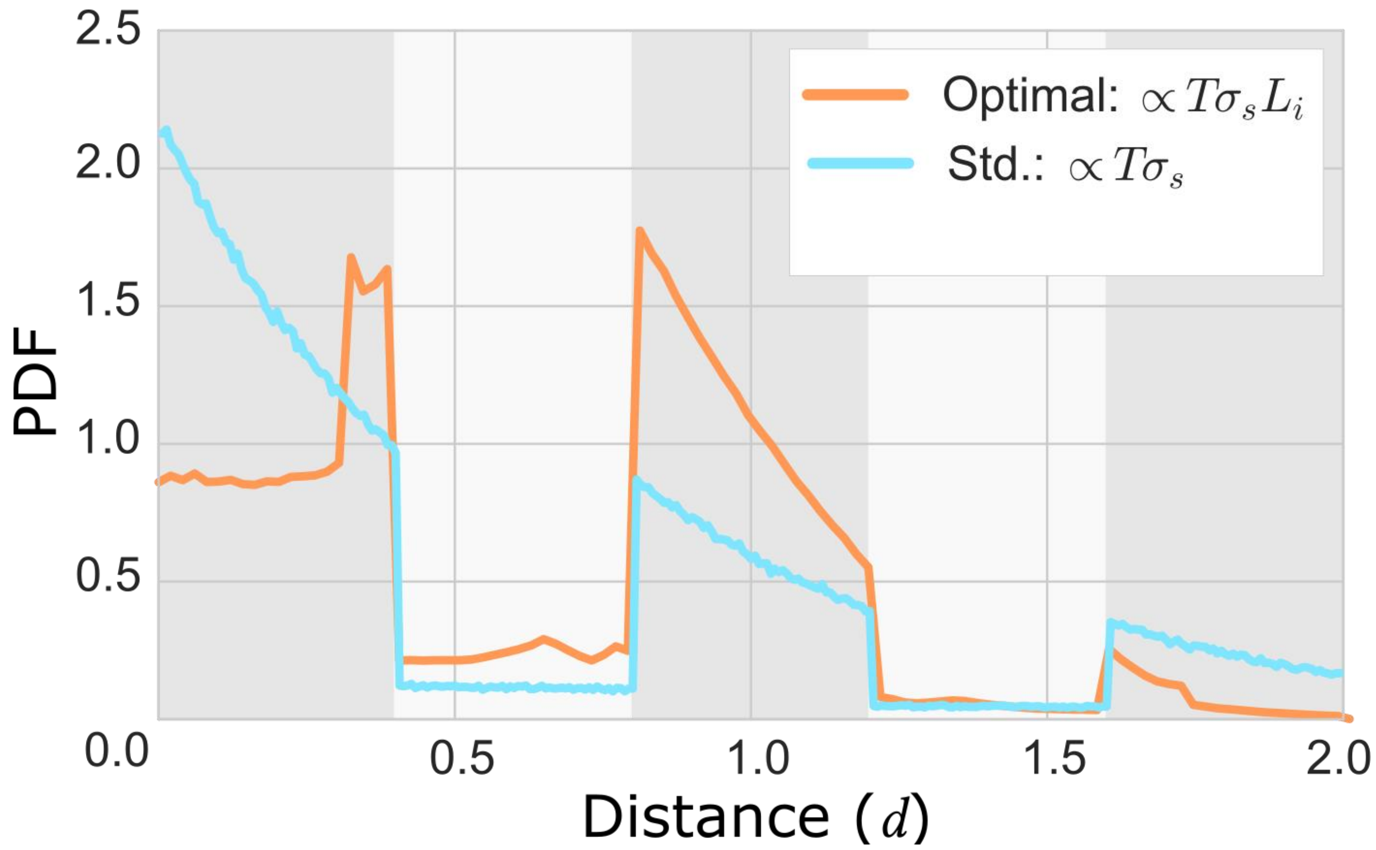

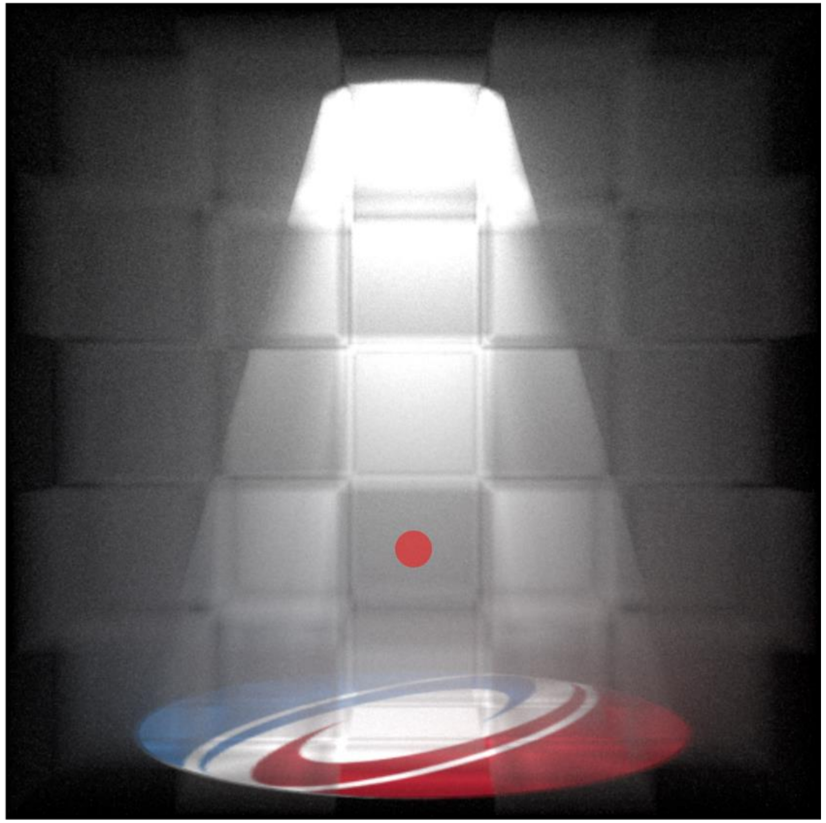

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

INCREMENTAL GUIDED DISTANCE SAMPLING

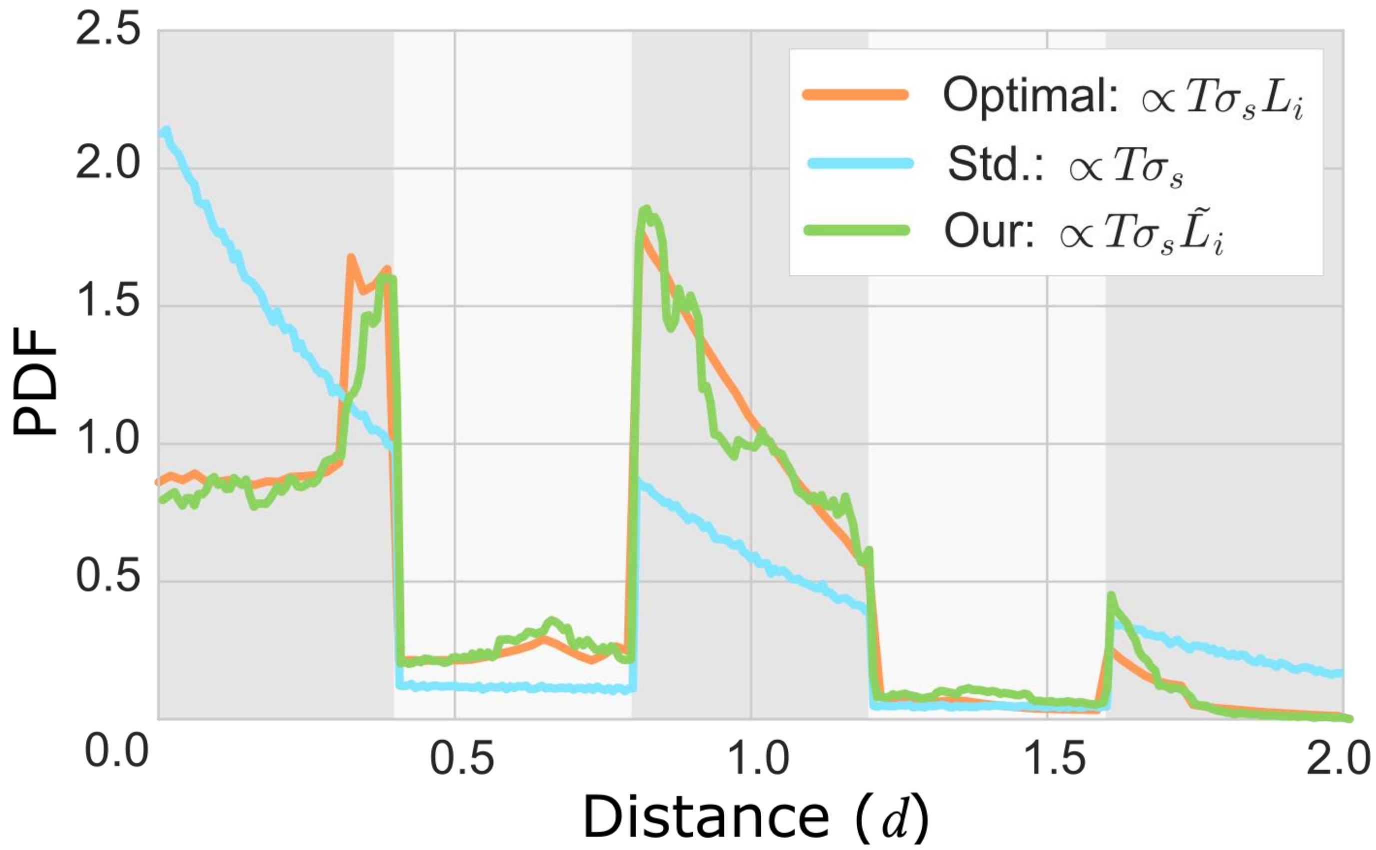

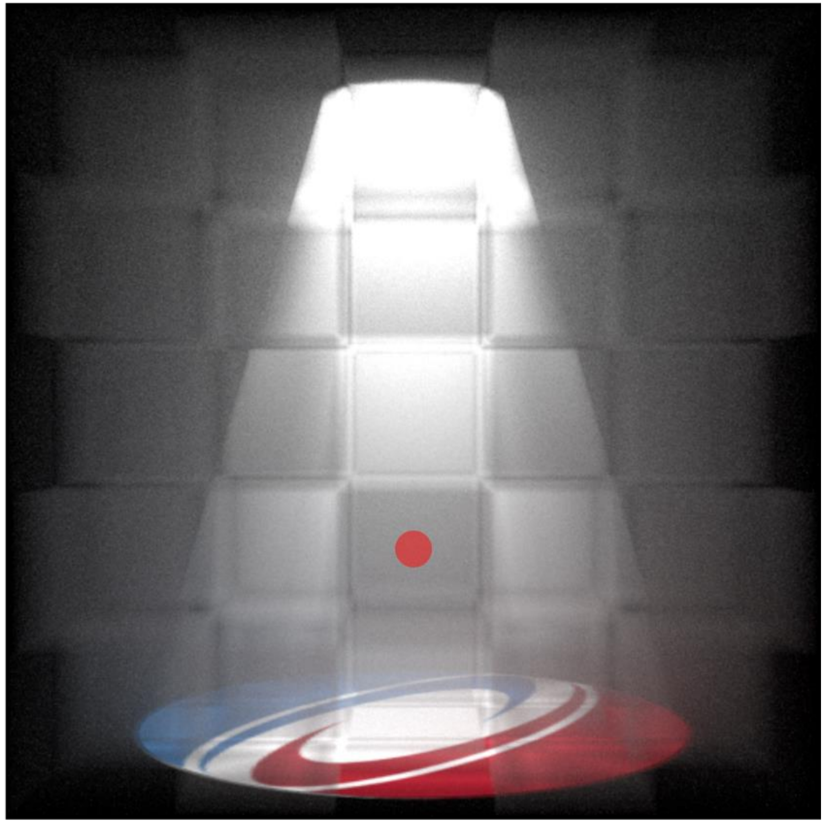

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VMM PHASE FUNCTION FITTING: PRE-PROCESSING STEP

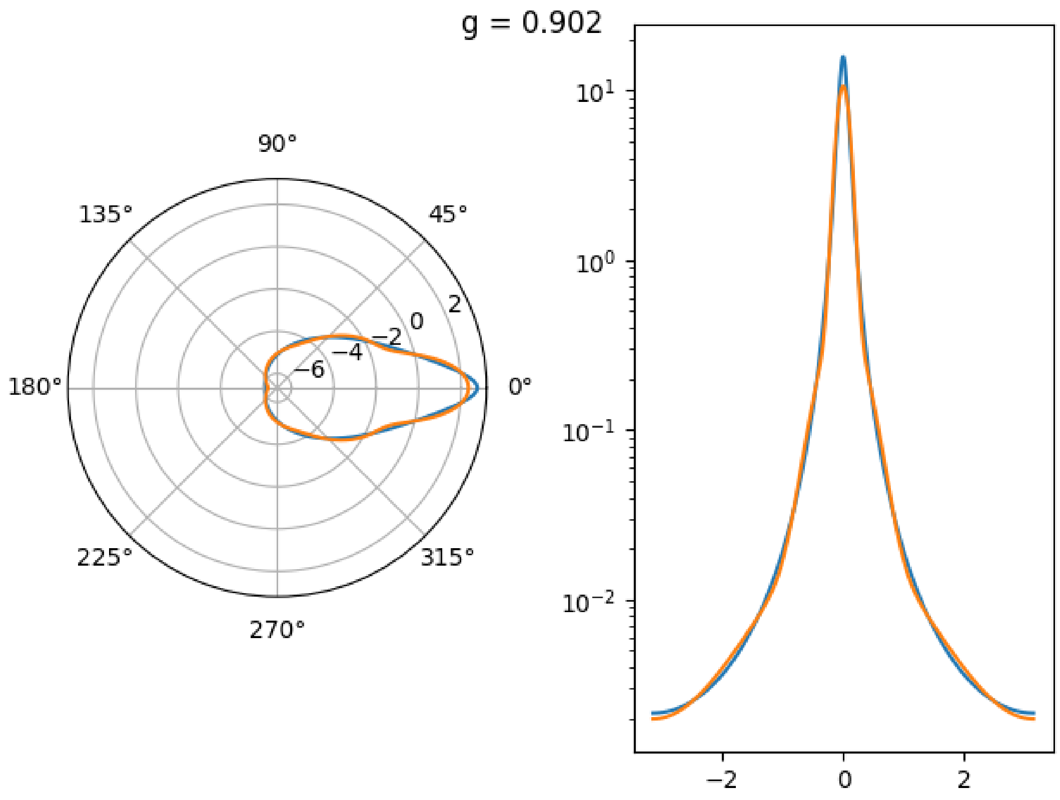

- Using up to $K = 4$ components
- Optimization Problem:

$$\underset{\Theta_f}{\arg\min} \sum_{n=1}^{N} \left[\mathcal{L}_{\log}(f(\omega_n, \dots), V(\omega_n, \Theta_f))\right]^2$$

$$\mathcal{L}_{\log}(d, m) = \log(d + \epsilon) - \log(m + \epsilon)$$

$$\epsilon = (1e - 4) \cdot \max(d_1, \dots, d_n)$$

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

VMM PHASE FUNCTION FITTING

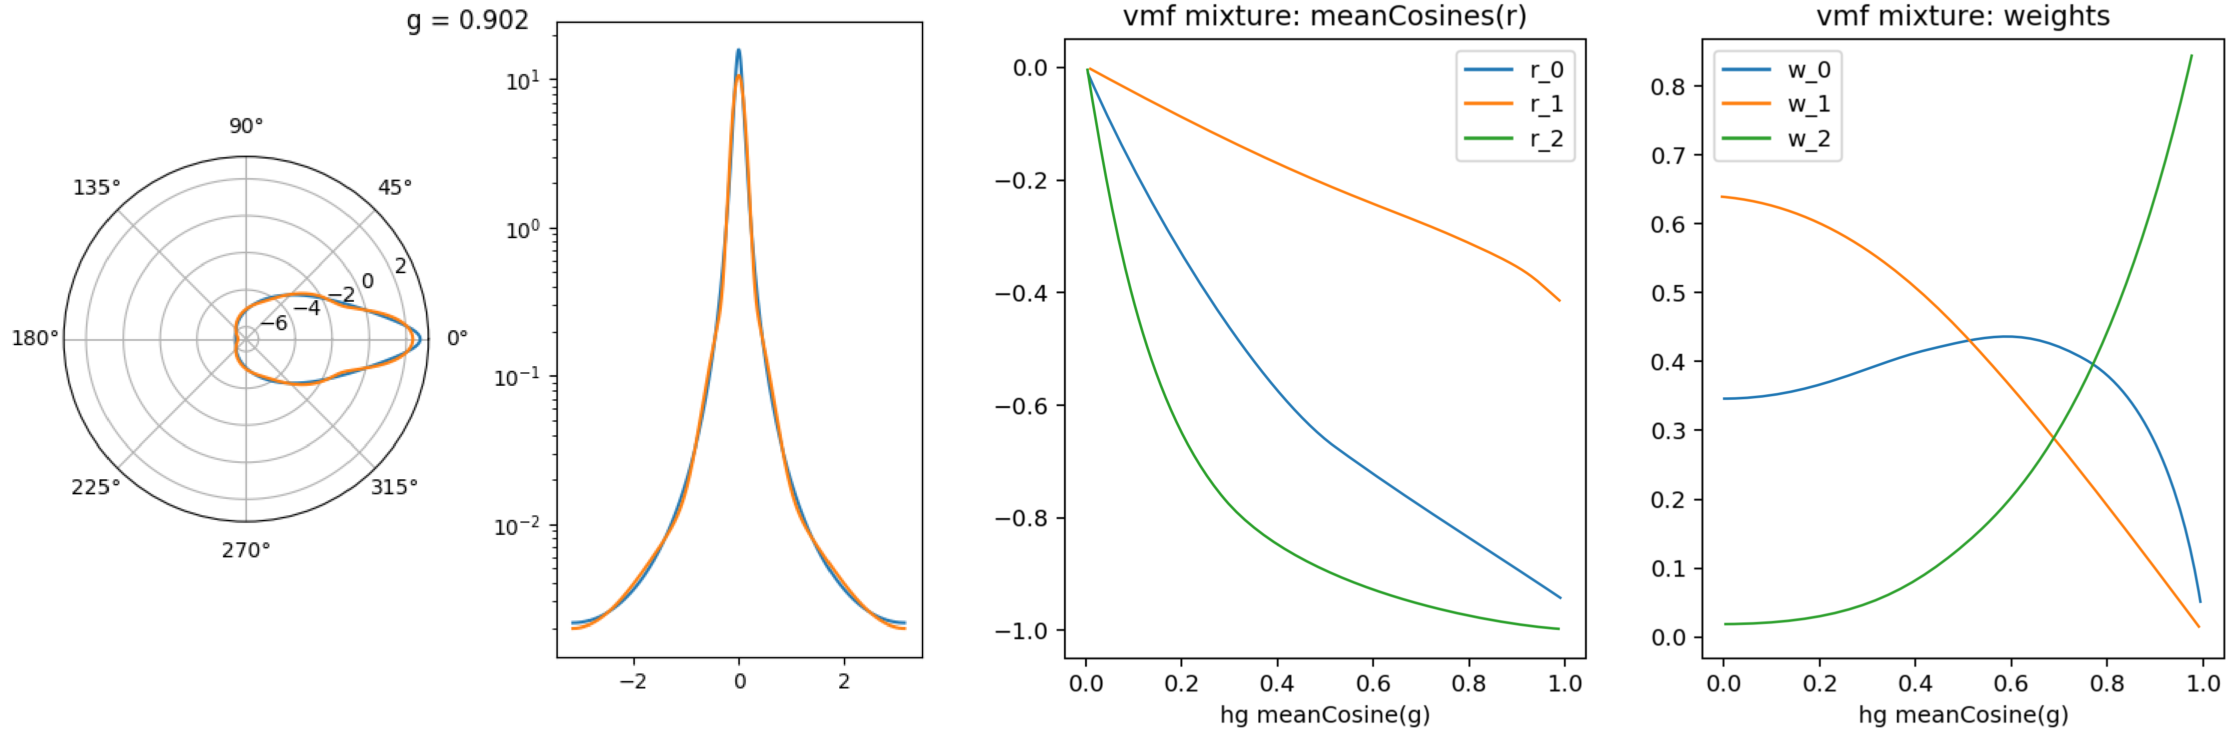

- Manifold representation of the VMM parameters for an HG phase function model for different mean cosines

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST