Volumetric Zero-Variance-Based Path Guiding

Sebastian Herholz[1] Yangyang Zhao[2] Oskar Elek[3]
Derek Nowrouzezahrai [2] Hendrik P. A. Lensch[1] Jaroslav Křivánek[3]

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

MOTIVATION

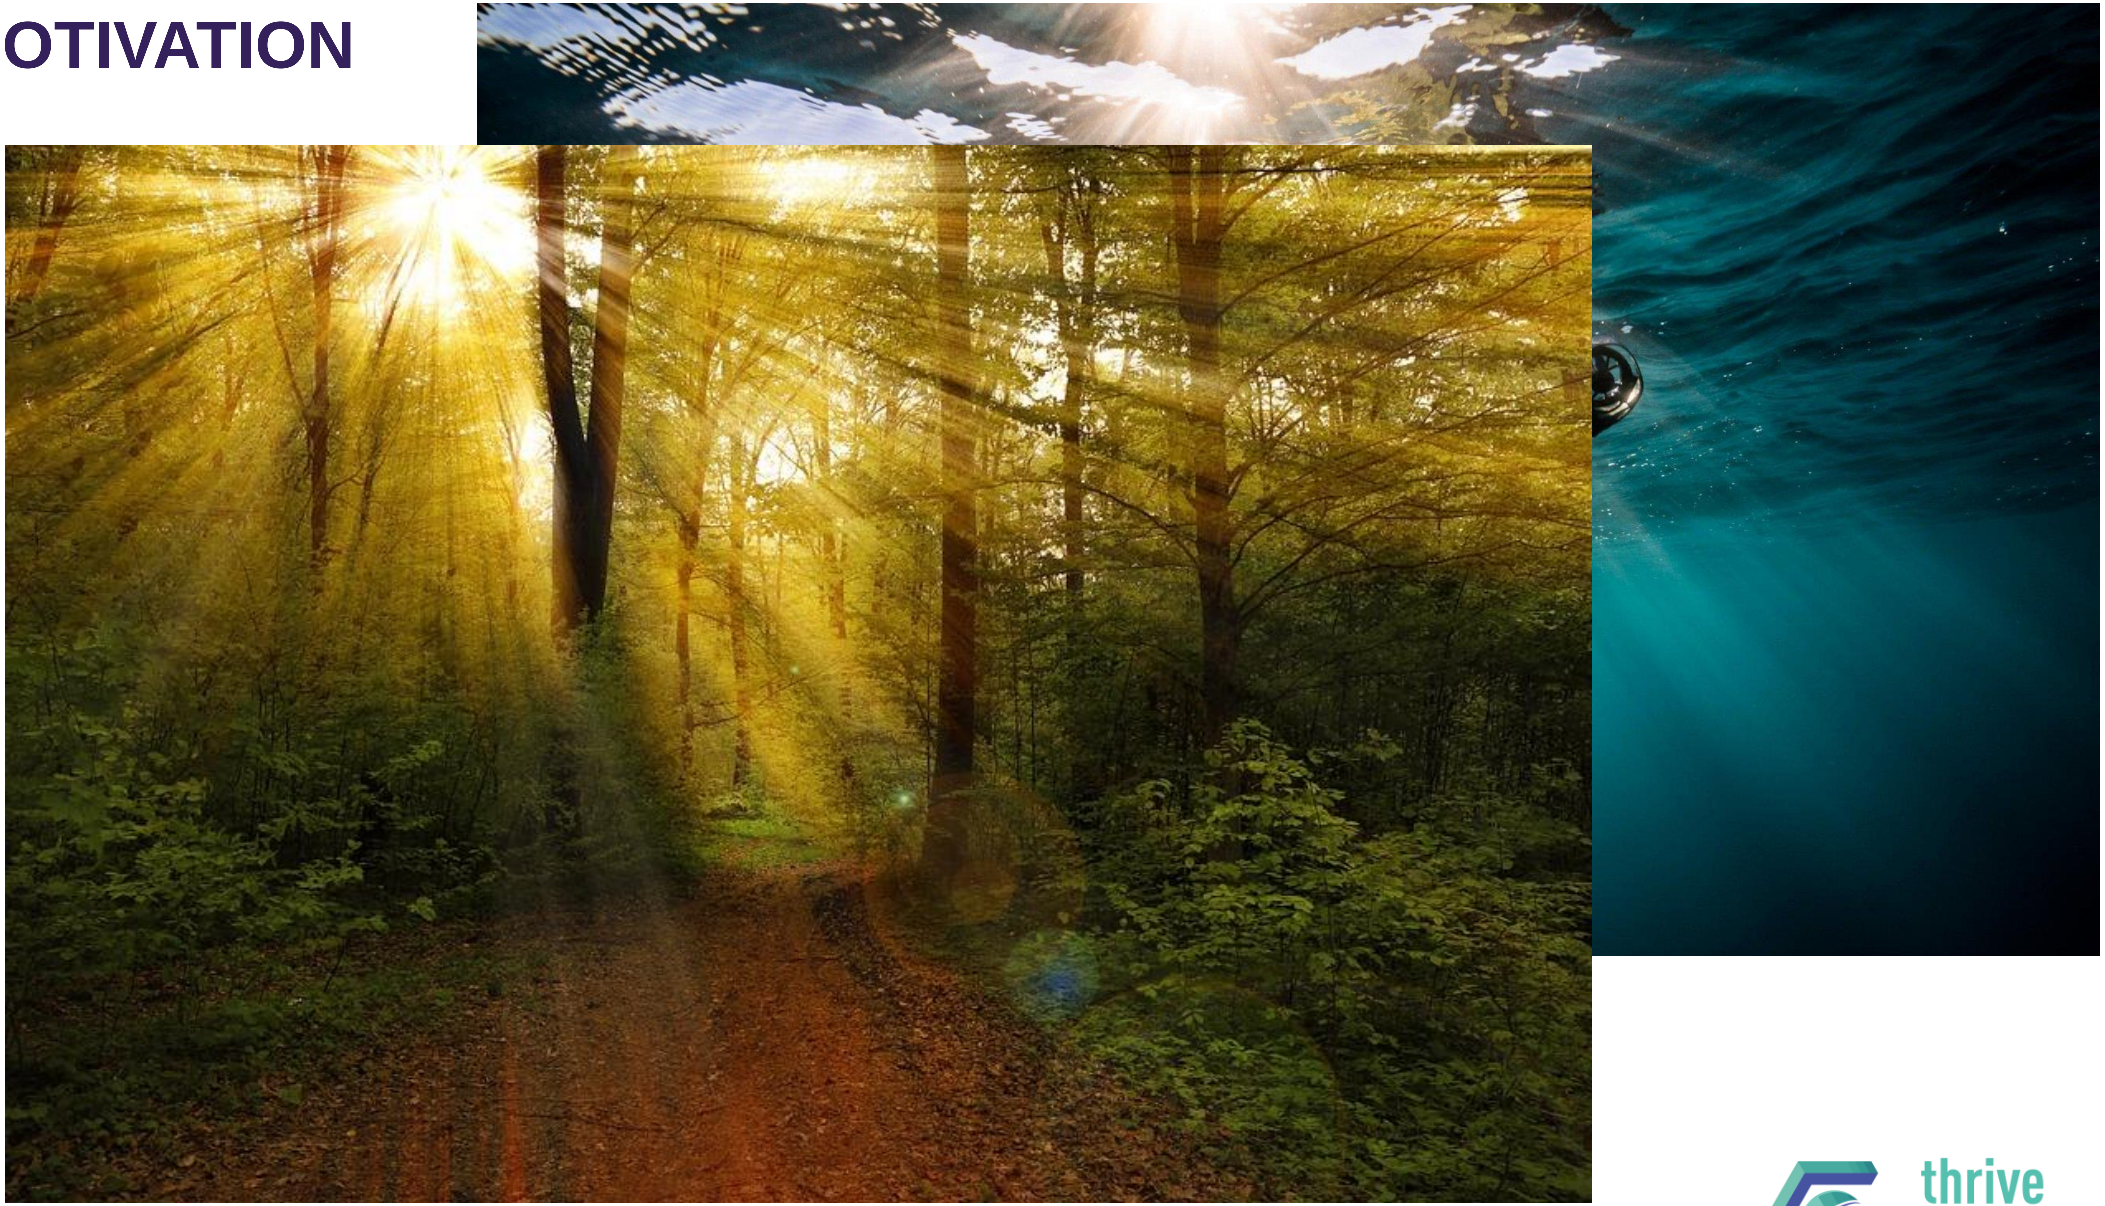

MOTIVATION

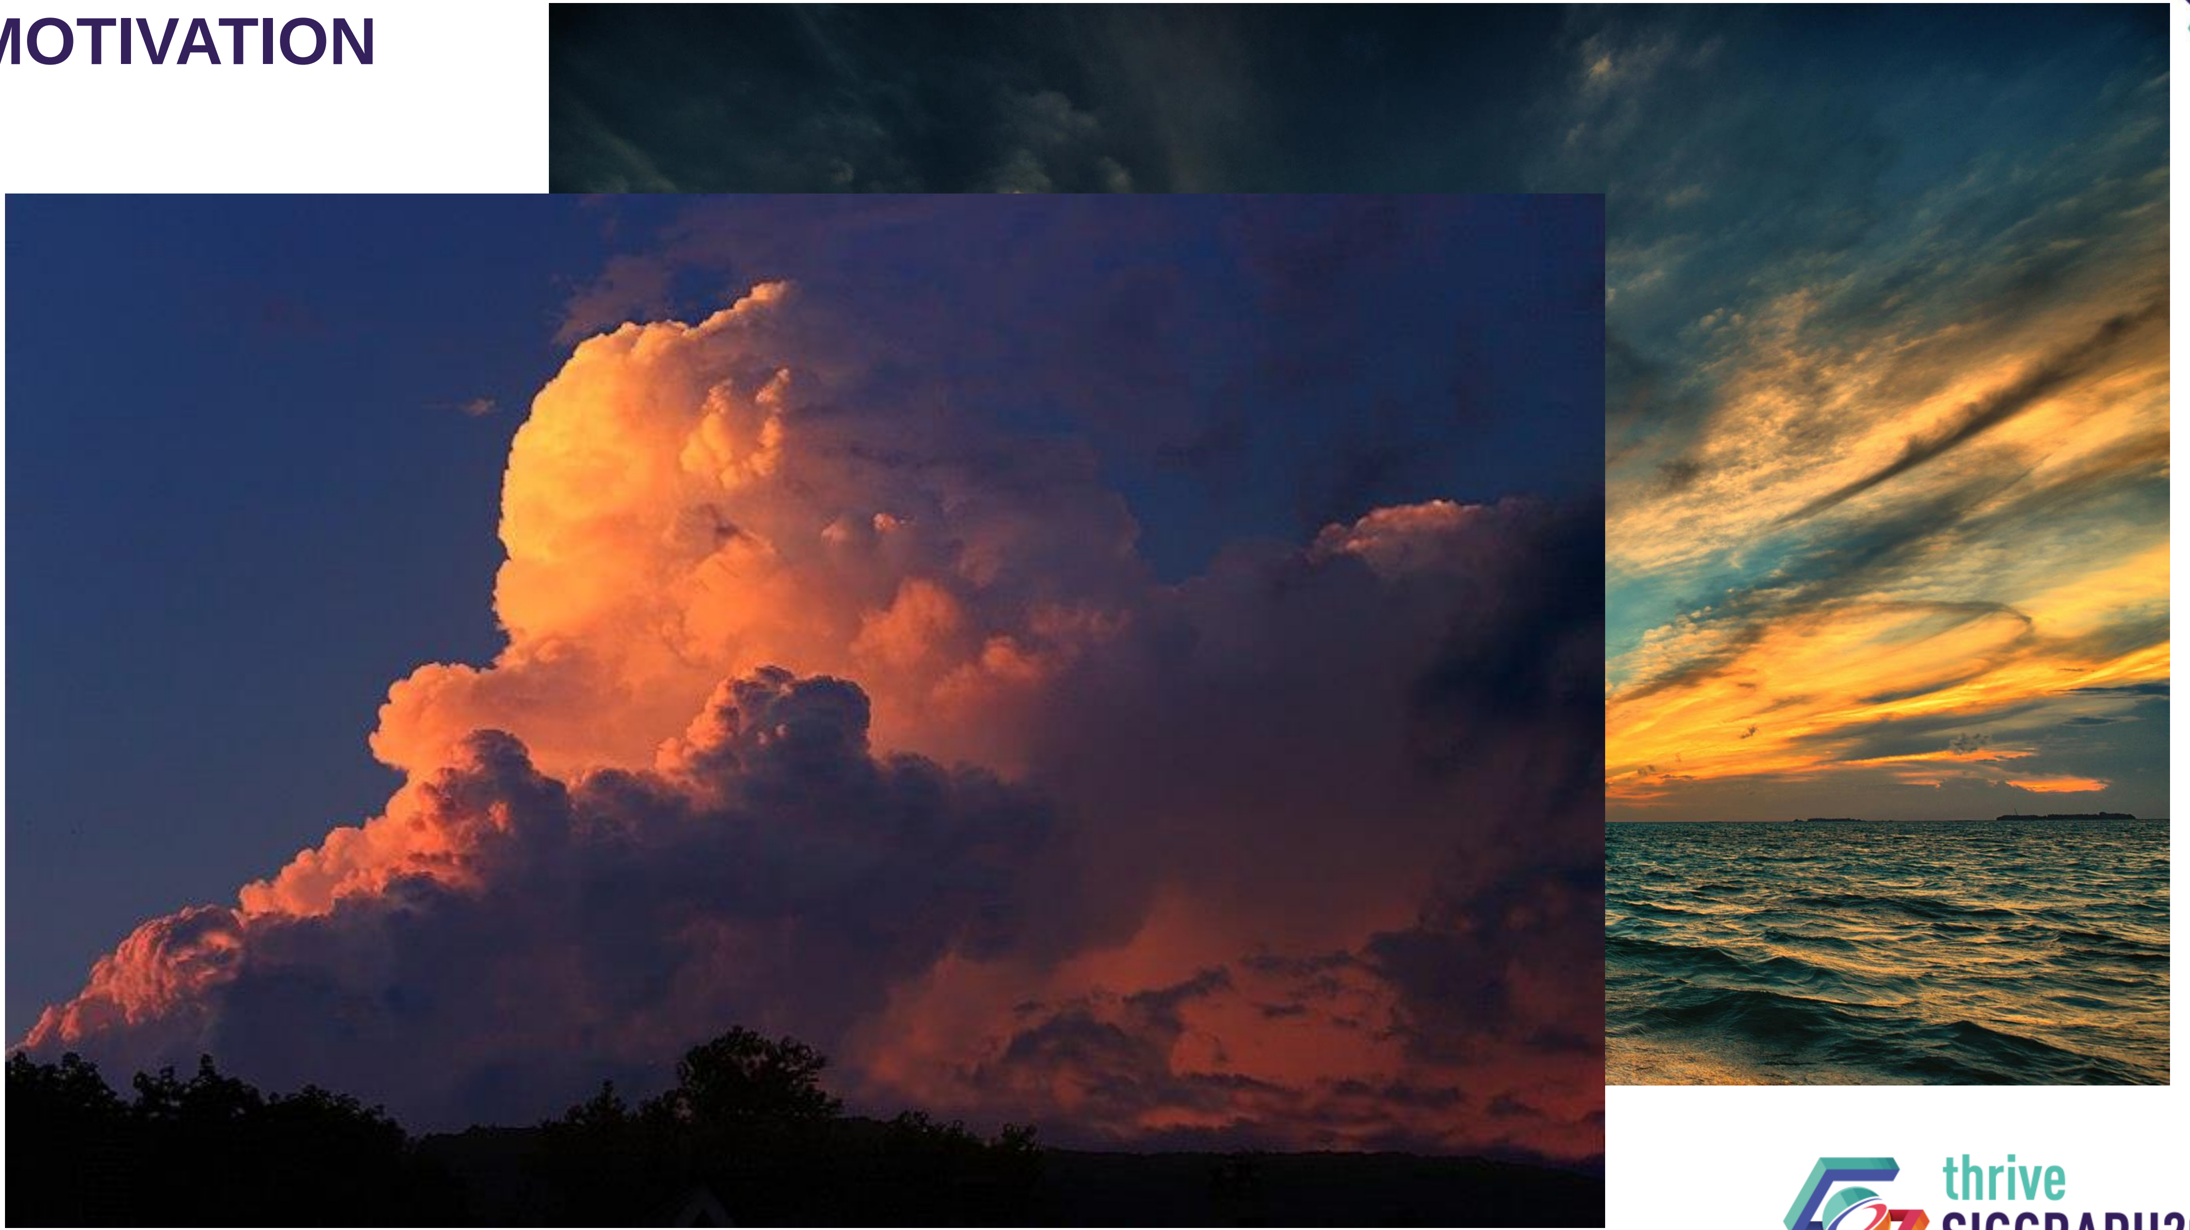

MOTIVATION

MOTIVATION

- Introduction in Volumetric Light transport

- Volumetric Path tracing
  - Samling decisions

Volumetric Path Guiding

MONTE-CARLO

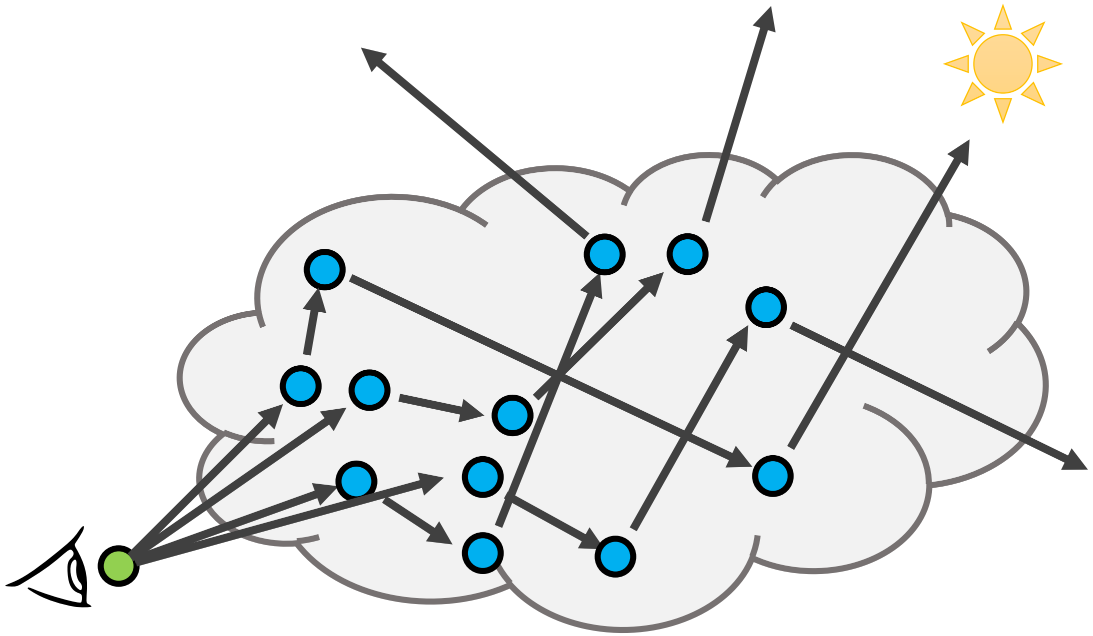

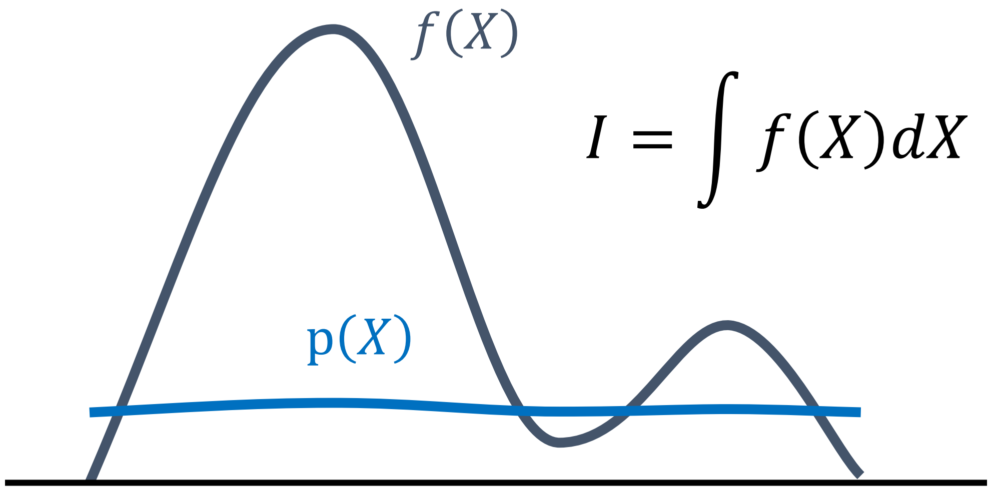

$$I = \int f(X)dX$$

- Estimator:

$$\hat{I}(X_1, \ldots, X_N) = \frac{1}{N} \sum \frac{f(X_i)}{p(X_i)}$$

- Variance:

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

ZERO VARIANCE MONTE-CARLO

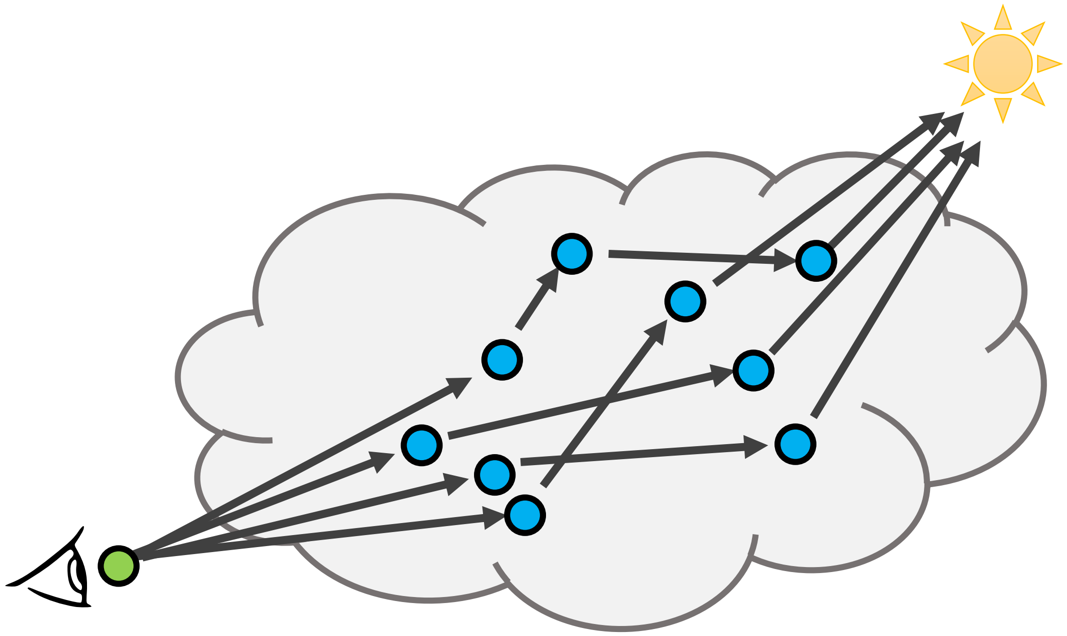

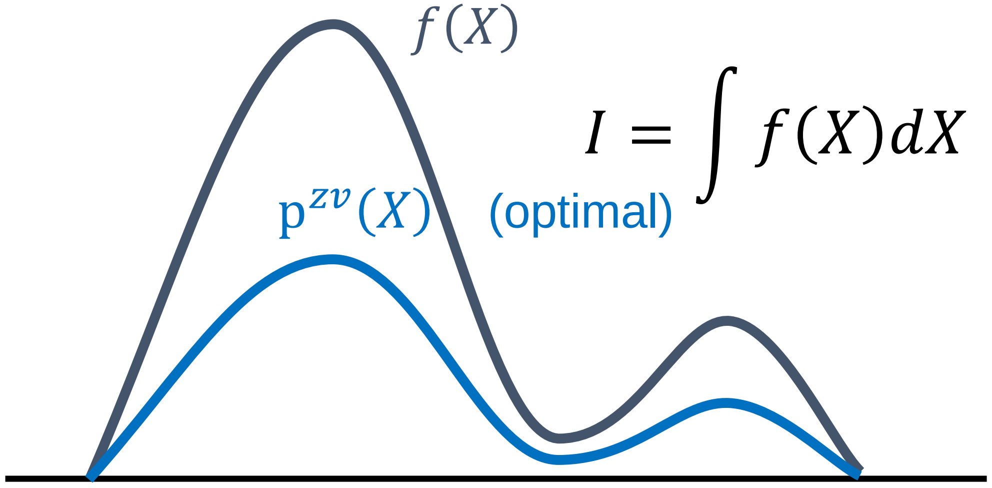

- Estimator:

$$\hat{I}(X_1, \ldots, X_N) = \frac{1}{N} \sum \frac{f(X_i)}{p^{zv}(X_i)} = c = I$$

- Zero-Variance:

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

THE 4 SAMPLING DECISIONS: SCATTER

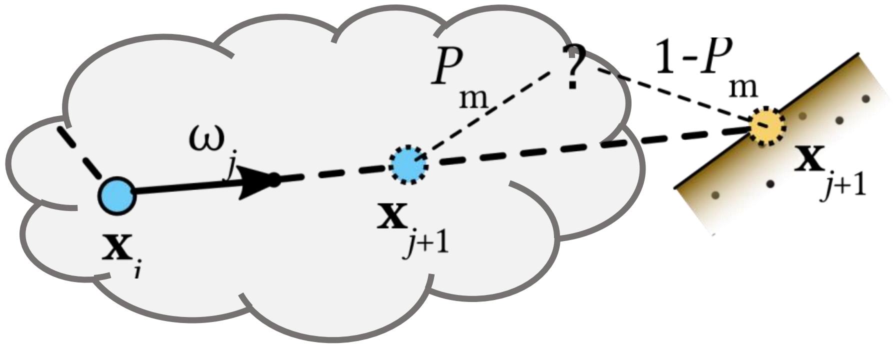

$$P_m(\boldsymbol{x}_j, \omega_j)$$

- Scatter:
  - Is the next path vertex inside or behind the volume?
  - Scatter probability: $P_m(\boldsymbol{x}_j, \omega_j)$

THE 4 SAMPLING DECISIONS: DISTANCE

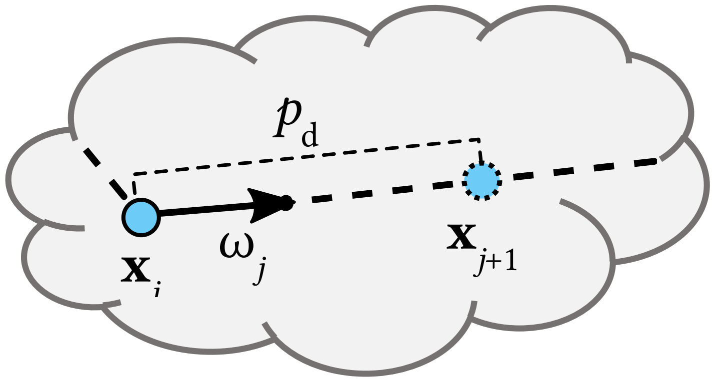

$$p_d(d_{j+1}|\boldsymbol{x}_j, \omega_j)$$

- Distance:
  - The distance (d_{j+1}) the next scattering occurs
  - Distance PDF: $p_d(d_{j+1}|\boldsymbol{x}_j, \omega_j)$

THE 4 SAMPLING DECISIONS: DIRECTION

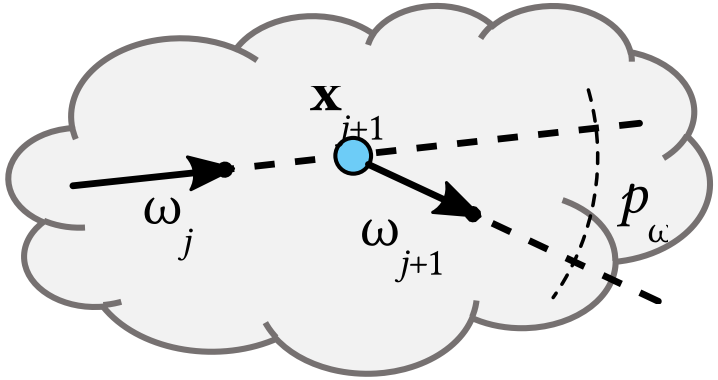

$$p_\omega(\omega_{j+1}|\boldsymbol{x}_{j+1}, \omega_j)$$

- Direction:
  - In which direction (ω_{j+1}) should the path continue?
  - Directional PDF: $p_\omega(\omega_{j+1}|\boldsymbol{x}_{j+1}, \omega_j)$

THE 4 SAMPLING DECISIONS: TERMINATION

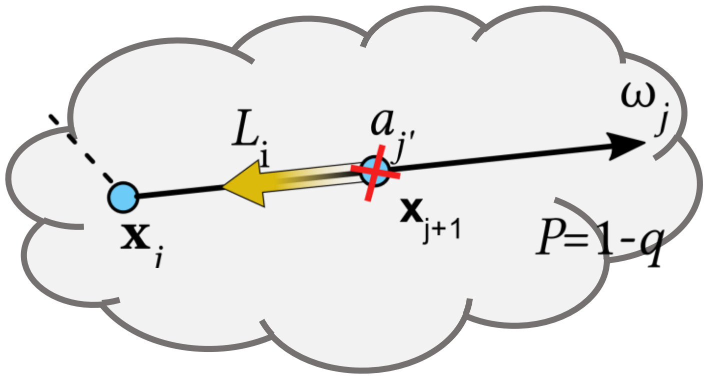

$$P_{RR}(\boldsymbol{x}_j, \omega_{j-1})$$

- Russian Roulette Termination:
  - Should we continue generating the random path/walk?
  - Termination probability PDF: $P_{RR}(\boldsymbol{x}_j, \omega_{j-1})$

VOLUMETRIC RANDOM WALK - DECISIONS

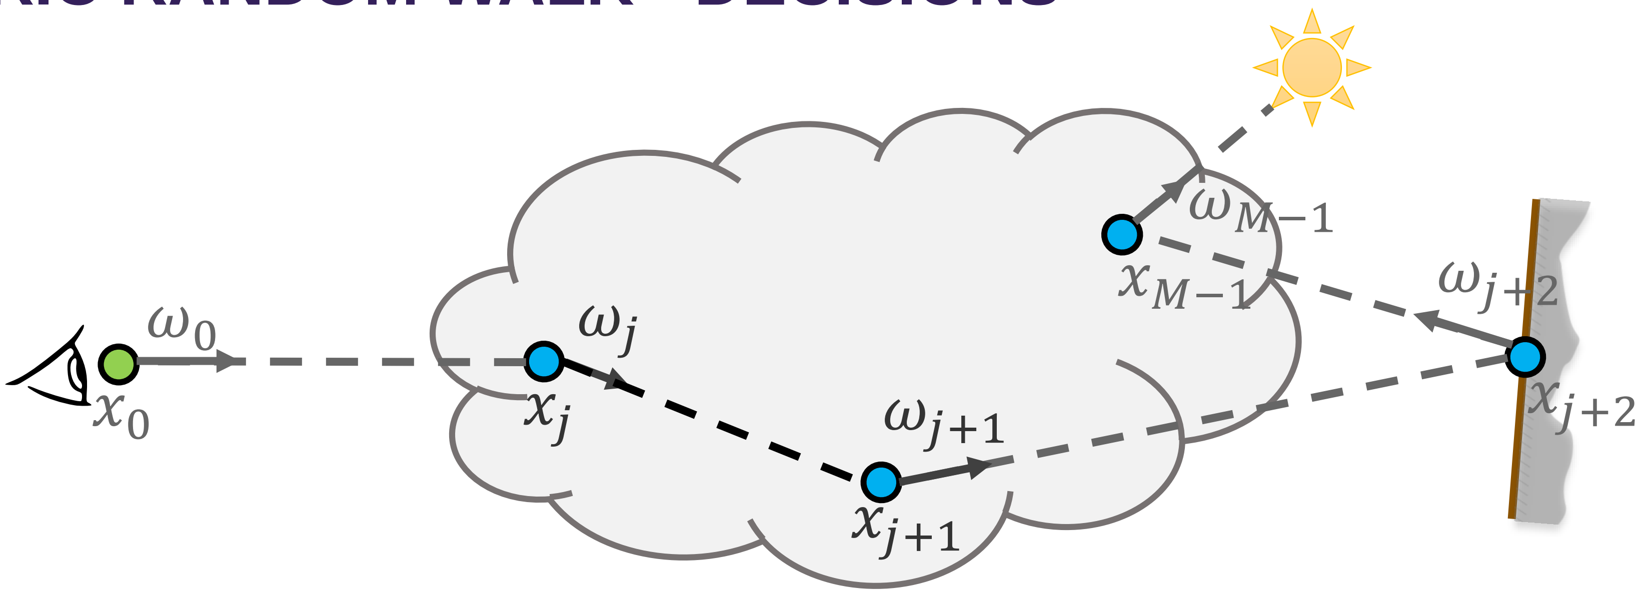

- Path-segment PDF:

$$p(\boldsymbol{x}_{j+1}, \omega_{j+1} | \boldsymbol{x}_j, \omega_j) = P_m(\ldots) \cdot p_d(\ldots) \cdot p_\omega(\ldots) \cdot (1 - P_{RR}(\ldots))$$

VOLUMETRIC RANDOM WALK - DECISIONS

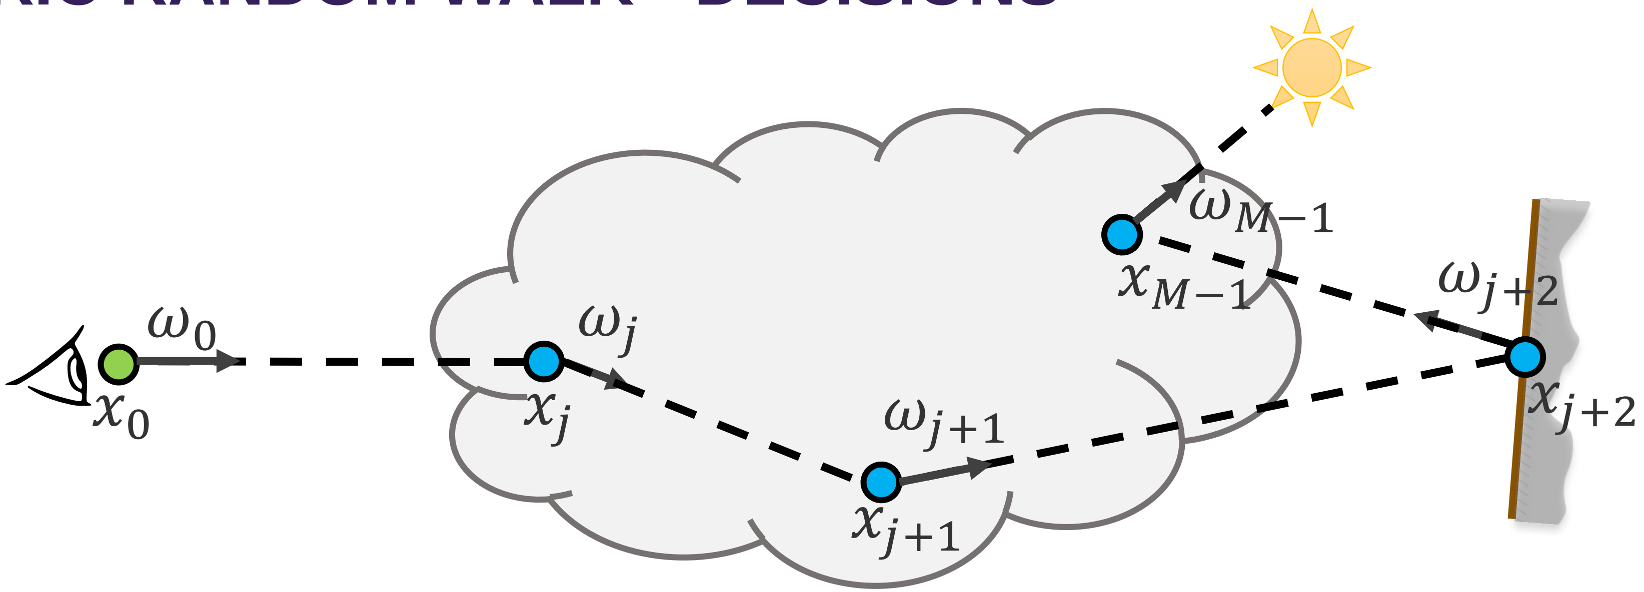

- Path-segment PDF:

$$p(\boldsymbol{x}_{j+1}, \omega_{j+1} | \boldsymbol{x}_j, \omega_j) = P_m(\dots) \cdot p_d(\dots) \cdot p_\omega(\dots) \cdot (1 - P_{RR}(\dots))$$

- Path PDF:

$$p(\boldsymbol{X}) = \prod_{j=1}^{M-1} p(\boldsymbol{x}_{j+1}, \omega_{j+1} | \boldsymbol{x}_j, \omega_j)$$

VOLUMETRIC RANDOM WALK - DECISIONS

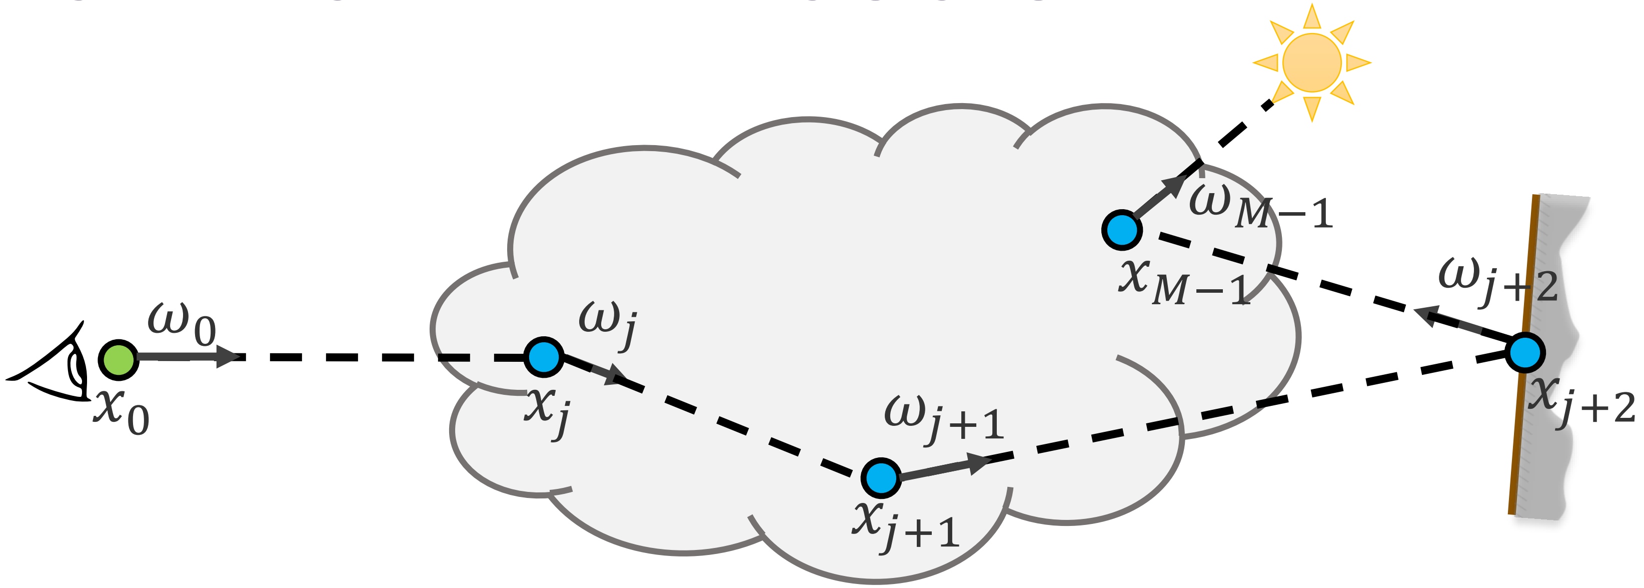

- Path-segment PDF:

$$p(\boldsymbol{x}_{j+1}, \omega_{j+1} | \boldsymbol{x}_j, \omega_j) = P_m(\dots) \cdot p_d(\dots) \cdot p_\omega(\dots) \cdot (1 - P_{RR}(\dots))$$

- Path PDF:

$$p(\boldsymbol{X}) = \prod_{j=1}^{M-1} p(\boldsymbol{x}_{j+1}, \omega_{j+1} | \boldsymbol{x}_j, \omega_j)$$

Source of variance

VOLUME RENDERING EQUATION

- Incident radiance:

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

- In-scattered radiance:

$$L_i(\mathrm{x}_d,\omega) = \int f(\omega,\omega') \cdot L(\mathrm{x_d},\omega')\mathrm{d}\omega'$$

VOLUME RENDERING EQUATION

- Incident radiance (volume):

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

Known Local Quantities

- In-scattered radiance:

$$L_i(\mathrm{x}_d,\omega) = \int \boxed{f(\omega,\omega')} \cdot L(\mathrm{x}_\mathrm{d},\omega')\mathrm{d}\omega'$$

VOLUME RENDERING EQUATION

- Incident radiance (volume):

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

Known Local Quantities

Unknown Light Transport Quantities

- In-scattered radiance:

$$L_i(\mathrm{x}_d,\omega) = \int f(\omega,\omega') \cdot L(\mathrm{x_d},\omega') \mathrm{d}\omega'$$

CHALLENGES FOR VOLUME SAMPLING

LIGHT SHAFTS

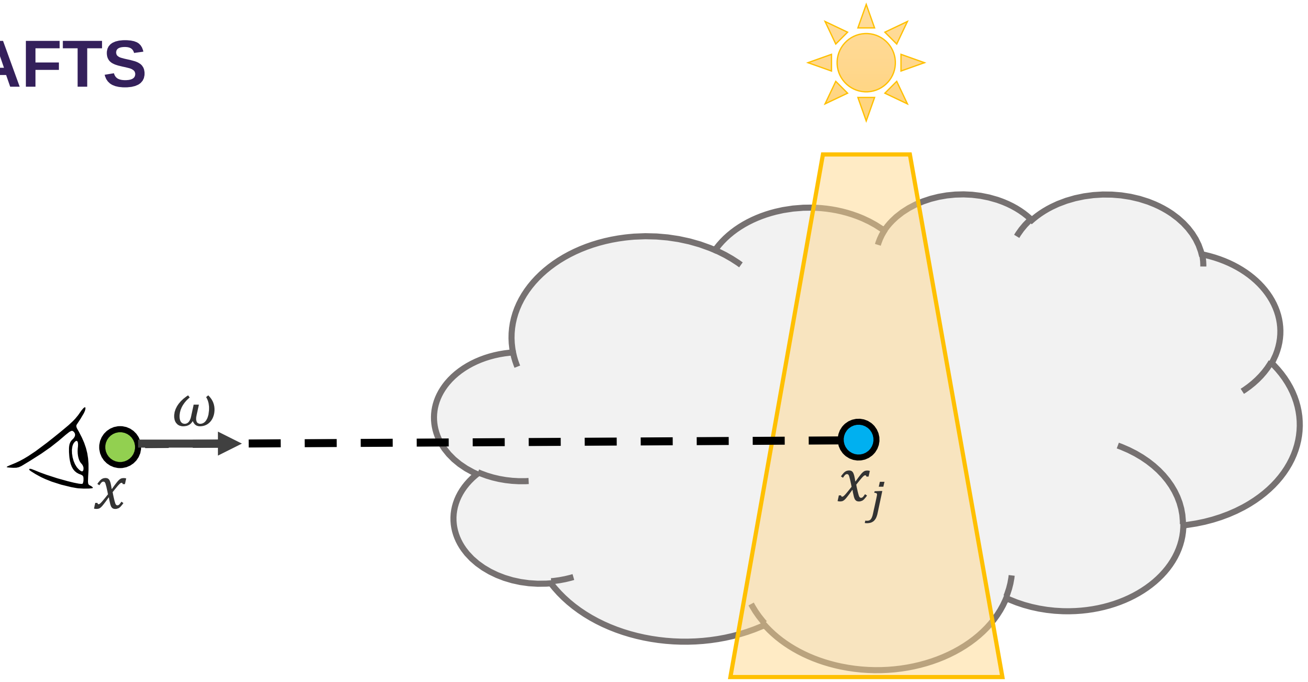

- Light shafts:
  - We need to scatter inside the light shaft.
  - We need to follow the direction of the light shaft.
  - We need to scatter towards the light shaft.

LIGHT SHAFTS

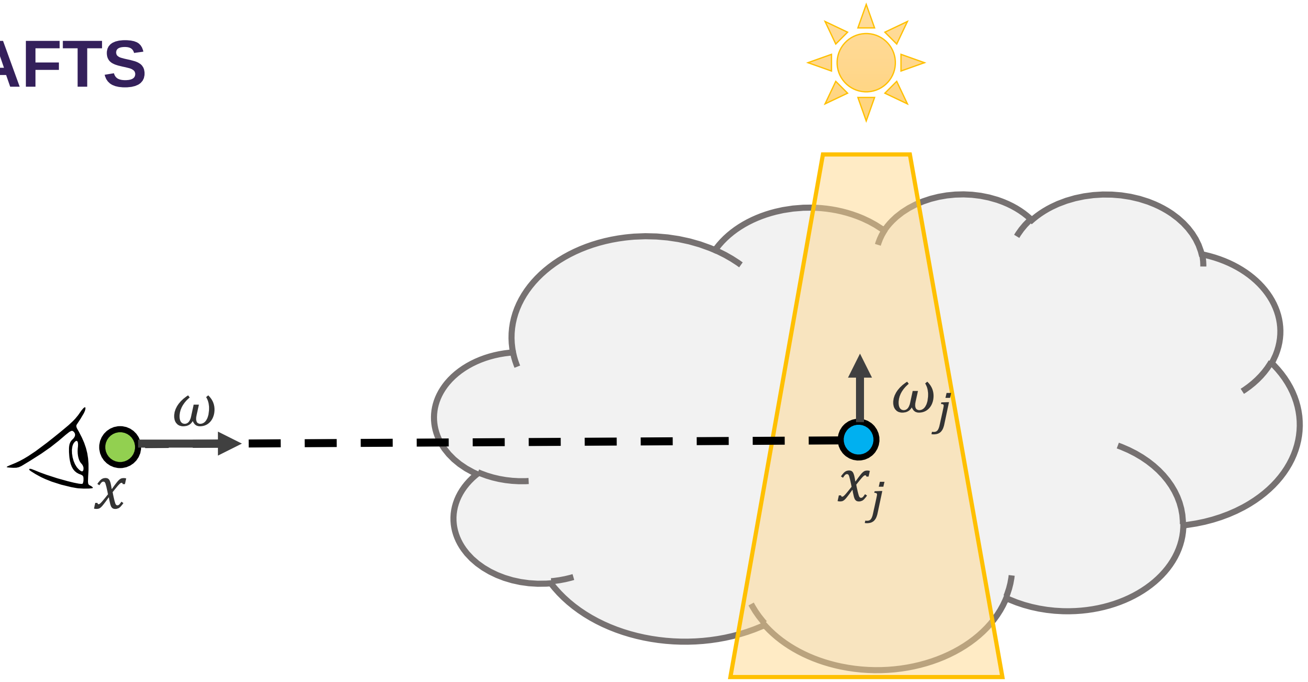

- Light shafts:
  - We need to scatter inside the light shaft.
  - We need to follow the direction of the light shaft.
  - We need to scatter towards the light shaft.

LIGHT SHAFTS

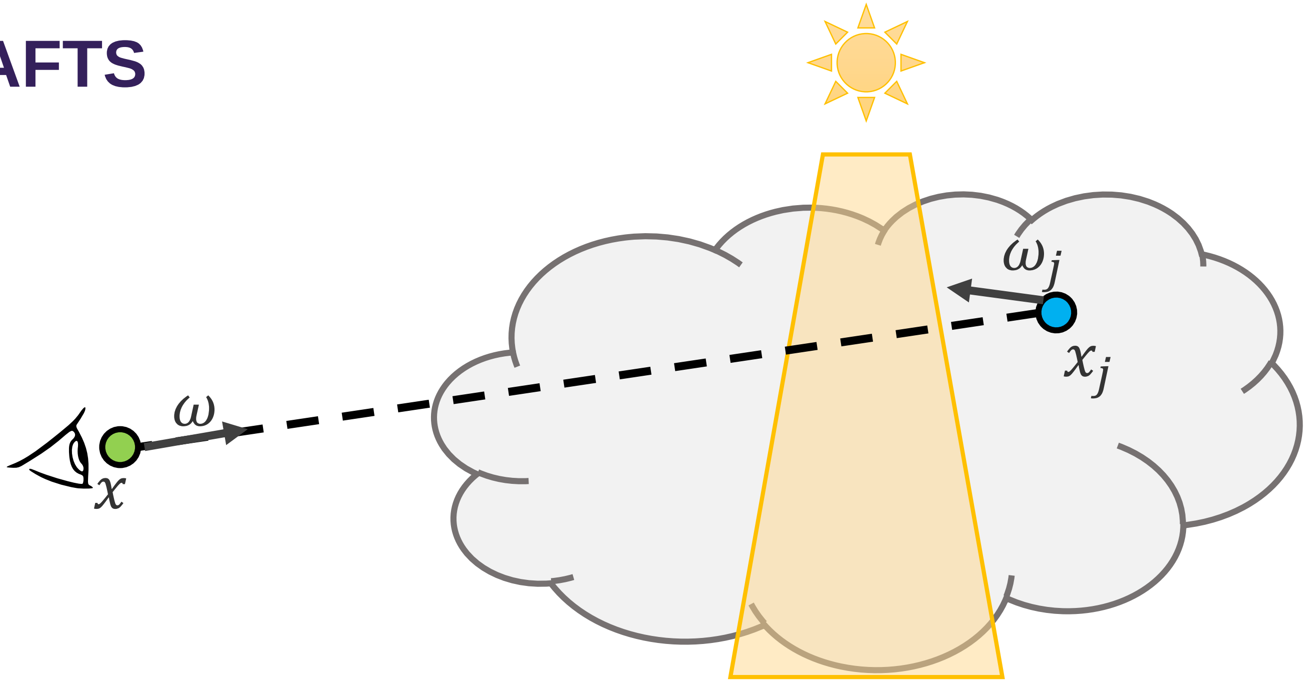

- Light shafts:
  - We need to scatter inside the light shaft.
  - We need to follow the direction of the light shaft.
  - We need to scatter towards the light shaft.
- Specialized solutions:

SUB-SURFACE-SCATTERING

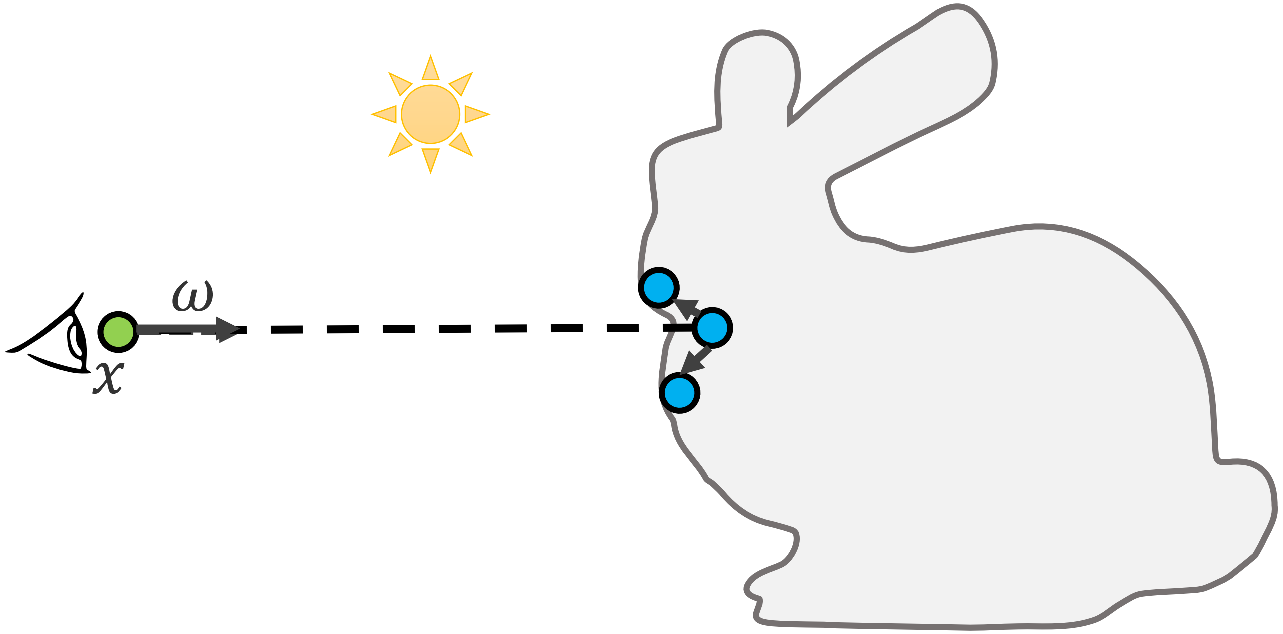

- Sub-Surface-Scattering:
  - We ‘often’ need stay close to the surface

SUB-SURFACE-SCATTERING

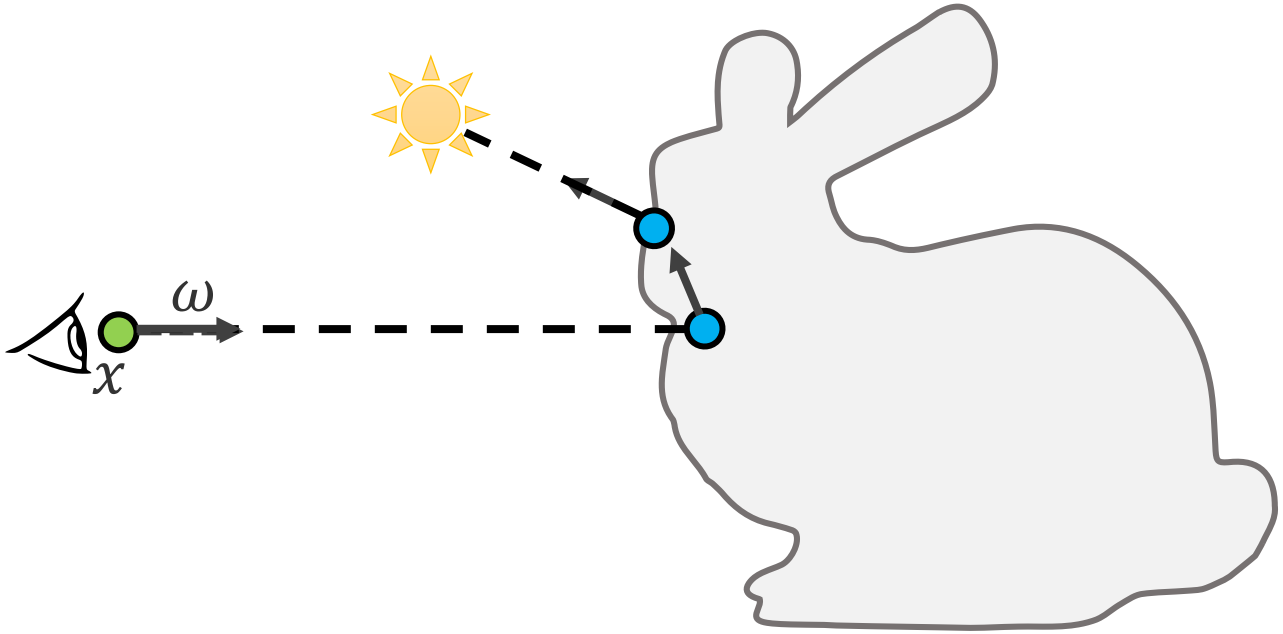

- Sub-Surface-Scattering:
  - We ‘often’ need stay close to the surface
  - We need to leave the object with the right direction
- Specialized solutions:

DENSE MEDIA

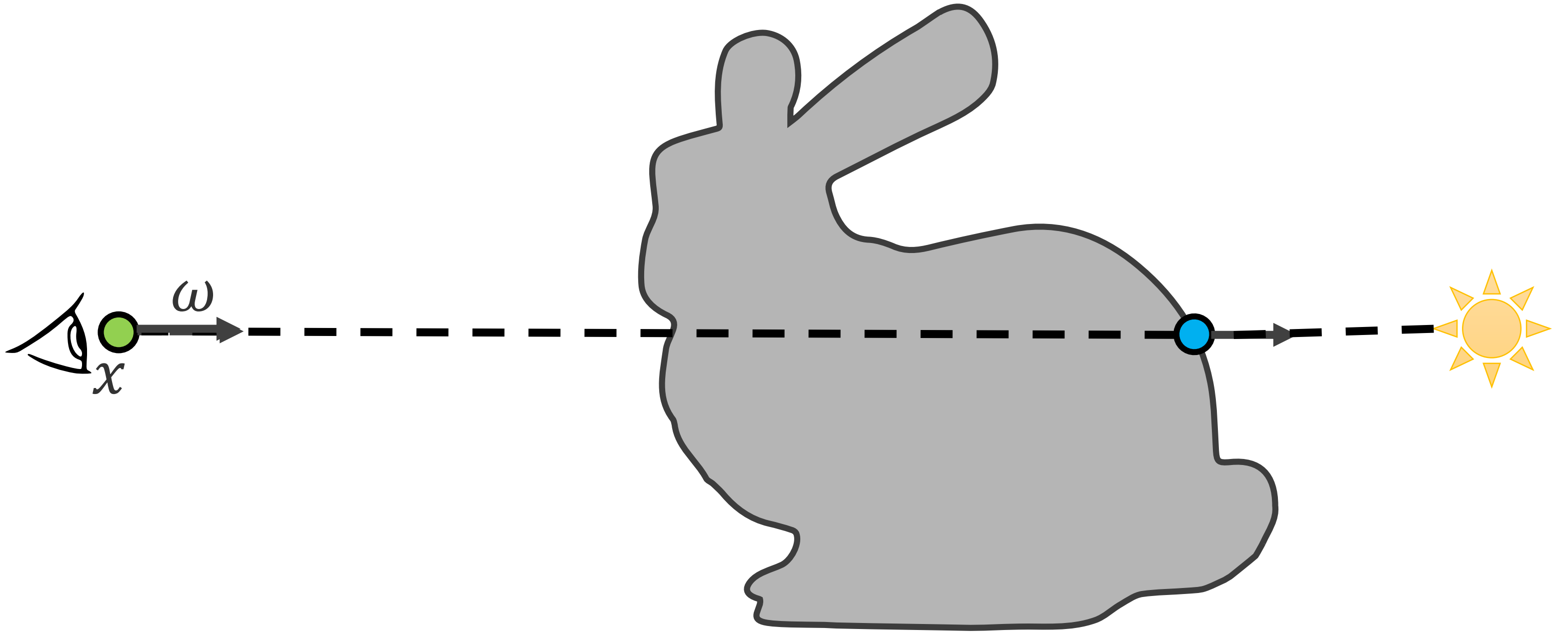

- Dense media:
  - We may need to **'avoid'** generating a scattering event even if the transmittance is low (e.g. strong light source behind the volume).
- Specialized solutions:

NON-DENSE MEDIA

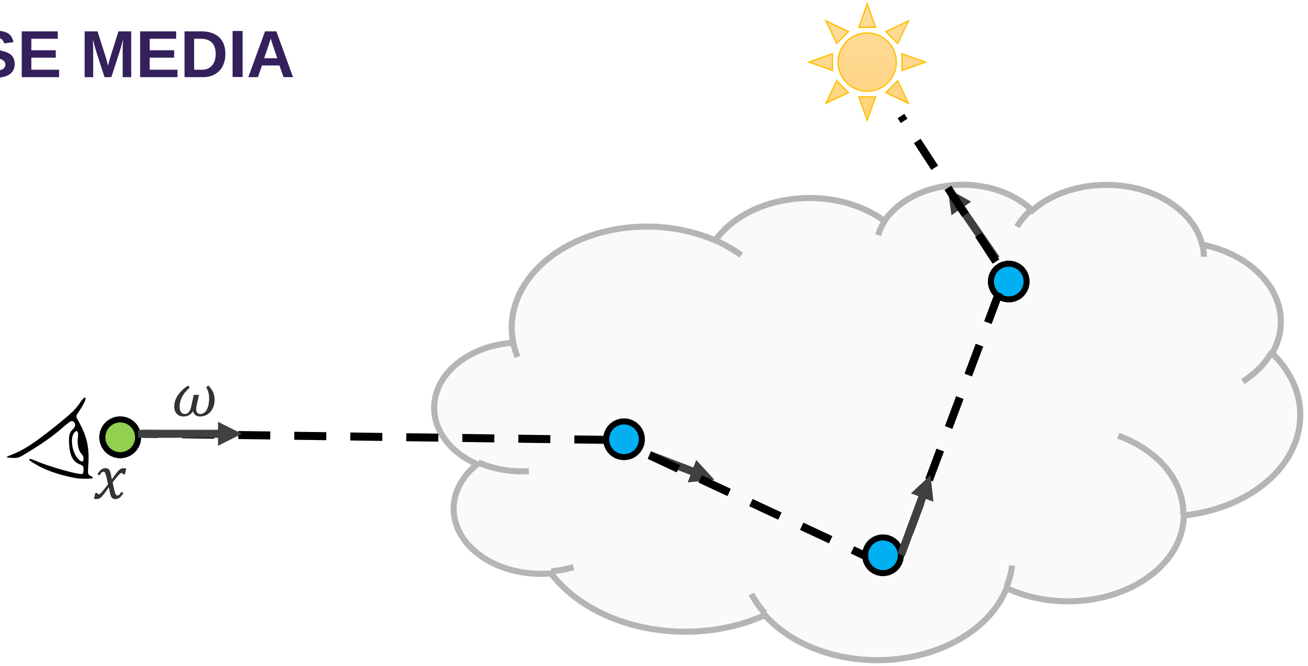

- Non-dense media:
  - We may need to '**force**' a scattering event even if the transmittance is high (e.g. no contribution from behind the volume).
- Specialized solutions:

SPECIALIZED SOLUTIONS: SHORTCOMMINGS

- Many individual solutions/algorithms:
  - Complicates the rendering code

- Only considering special cases:
  - Surface-bounded volumes
  - Homogenous or isotropic volumes
  - Single scattering

- Not intuitive (for an artist) to decided which feature helps when.

ZERO-VARIANCE RANDOM WALK THEORY

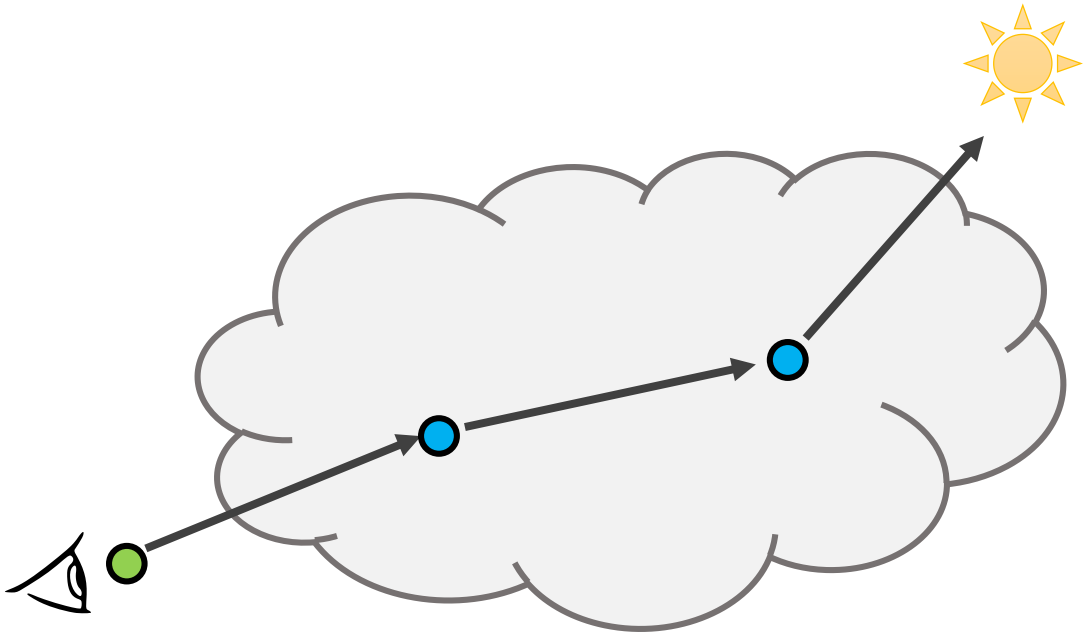

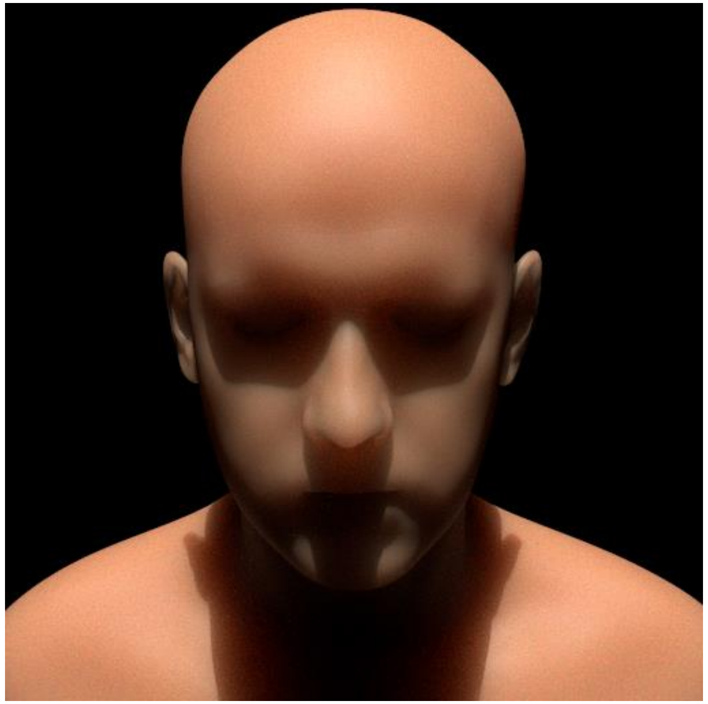

- **Theoretical** framework for the optimal segment PDF
- **All** 4 local decision have to be optimal:

$$p^{zv}(\ldots) = P_m^{zv}(\ldots) \cdot p_d^{zv}(\ldots) \cdot p_\omega^{zv}(\ldots) \cdot (1 - P_{RR}^{zv}(\ldots))$$

ZERO-VARIANCE PDF EXAMPLES

- Opt. distance PDF:

$$p_d^{zv}(d_{j+1}|\boldsymbol{x}_j, \omega_j) \propto T(\boldsymbol{x}_j, \boldsymbol{x}_{j+1}) \cdot \sigma_s(\boldsymbol{x}_{j+1}) \cdot \boxed{L_i\,(\boldsymbol{x}_{j+1}, \omega_j)}$$

Unknown Light Transport Quantities

- Opt. direction PDF:

$$p_\omega^{zv}(\omega_{j+1}|\boldsymbol{x}_{j+1}, \omega_j) \propto f(x_{j+1}, \omega_j, \omega_{j+1}) \cdot \boxed{L(x_{j+1}, \omega_{j+1})}$$

FUN FACT: STD. VOLUME SAMPLING AND ZERO-VARIANCE

- Std. volume sampling resolves to a zero-variance estimator if:

$$L(\boldsymbol{x}, \omega) = const \qquad L_i(\boldsymbol{x}, \omega) = const \qquad \forall\ \boldsymbol{x},\ \omega$$

- Its variance depends on the deviation of the **actual** volumetric light transport to this assumption!

- Consequence:

Any conservative guiding towards the actual VLT results in a variance reduction !!!

ZERO-VARIANCE-BASED VOLUMETRIC PATH GUIDING

ZV-BASED VOLUMETRIC PATH GUIDING: GOALS

- Consider the **complete** volumetric light transport
- No prior assumptions or special cases
- Leverage success of local surface guiding methods
  - Extend the concept to volumes

ZV-BASED VOLUMETRIC PATH GUIDING: CONTRIBUTIONS

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

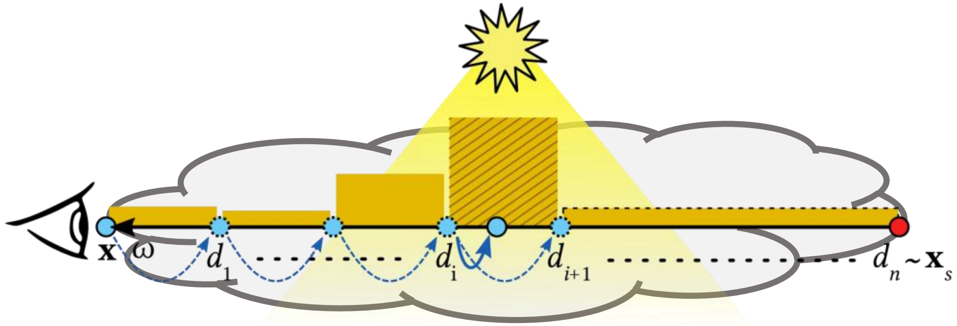

  - 3 Guided product directional sampling:

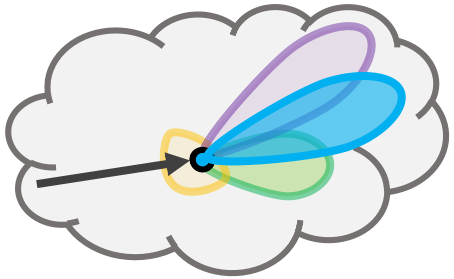

  - 4 Guided Russian roulette and Splitting:

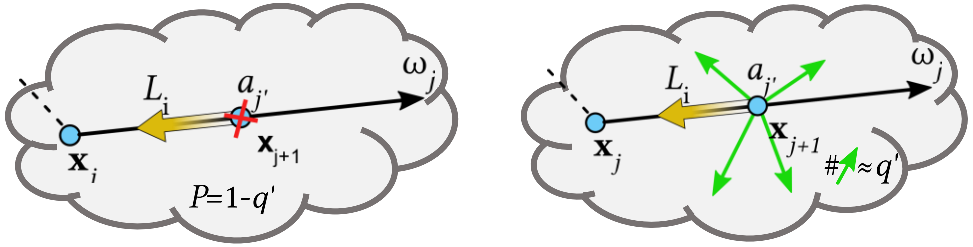


Standard Sampling
45 min


Our Guided Sampling
45 min

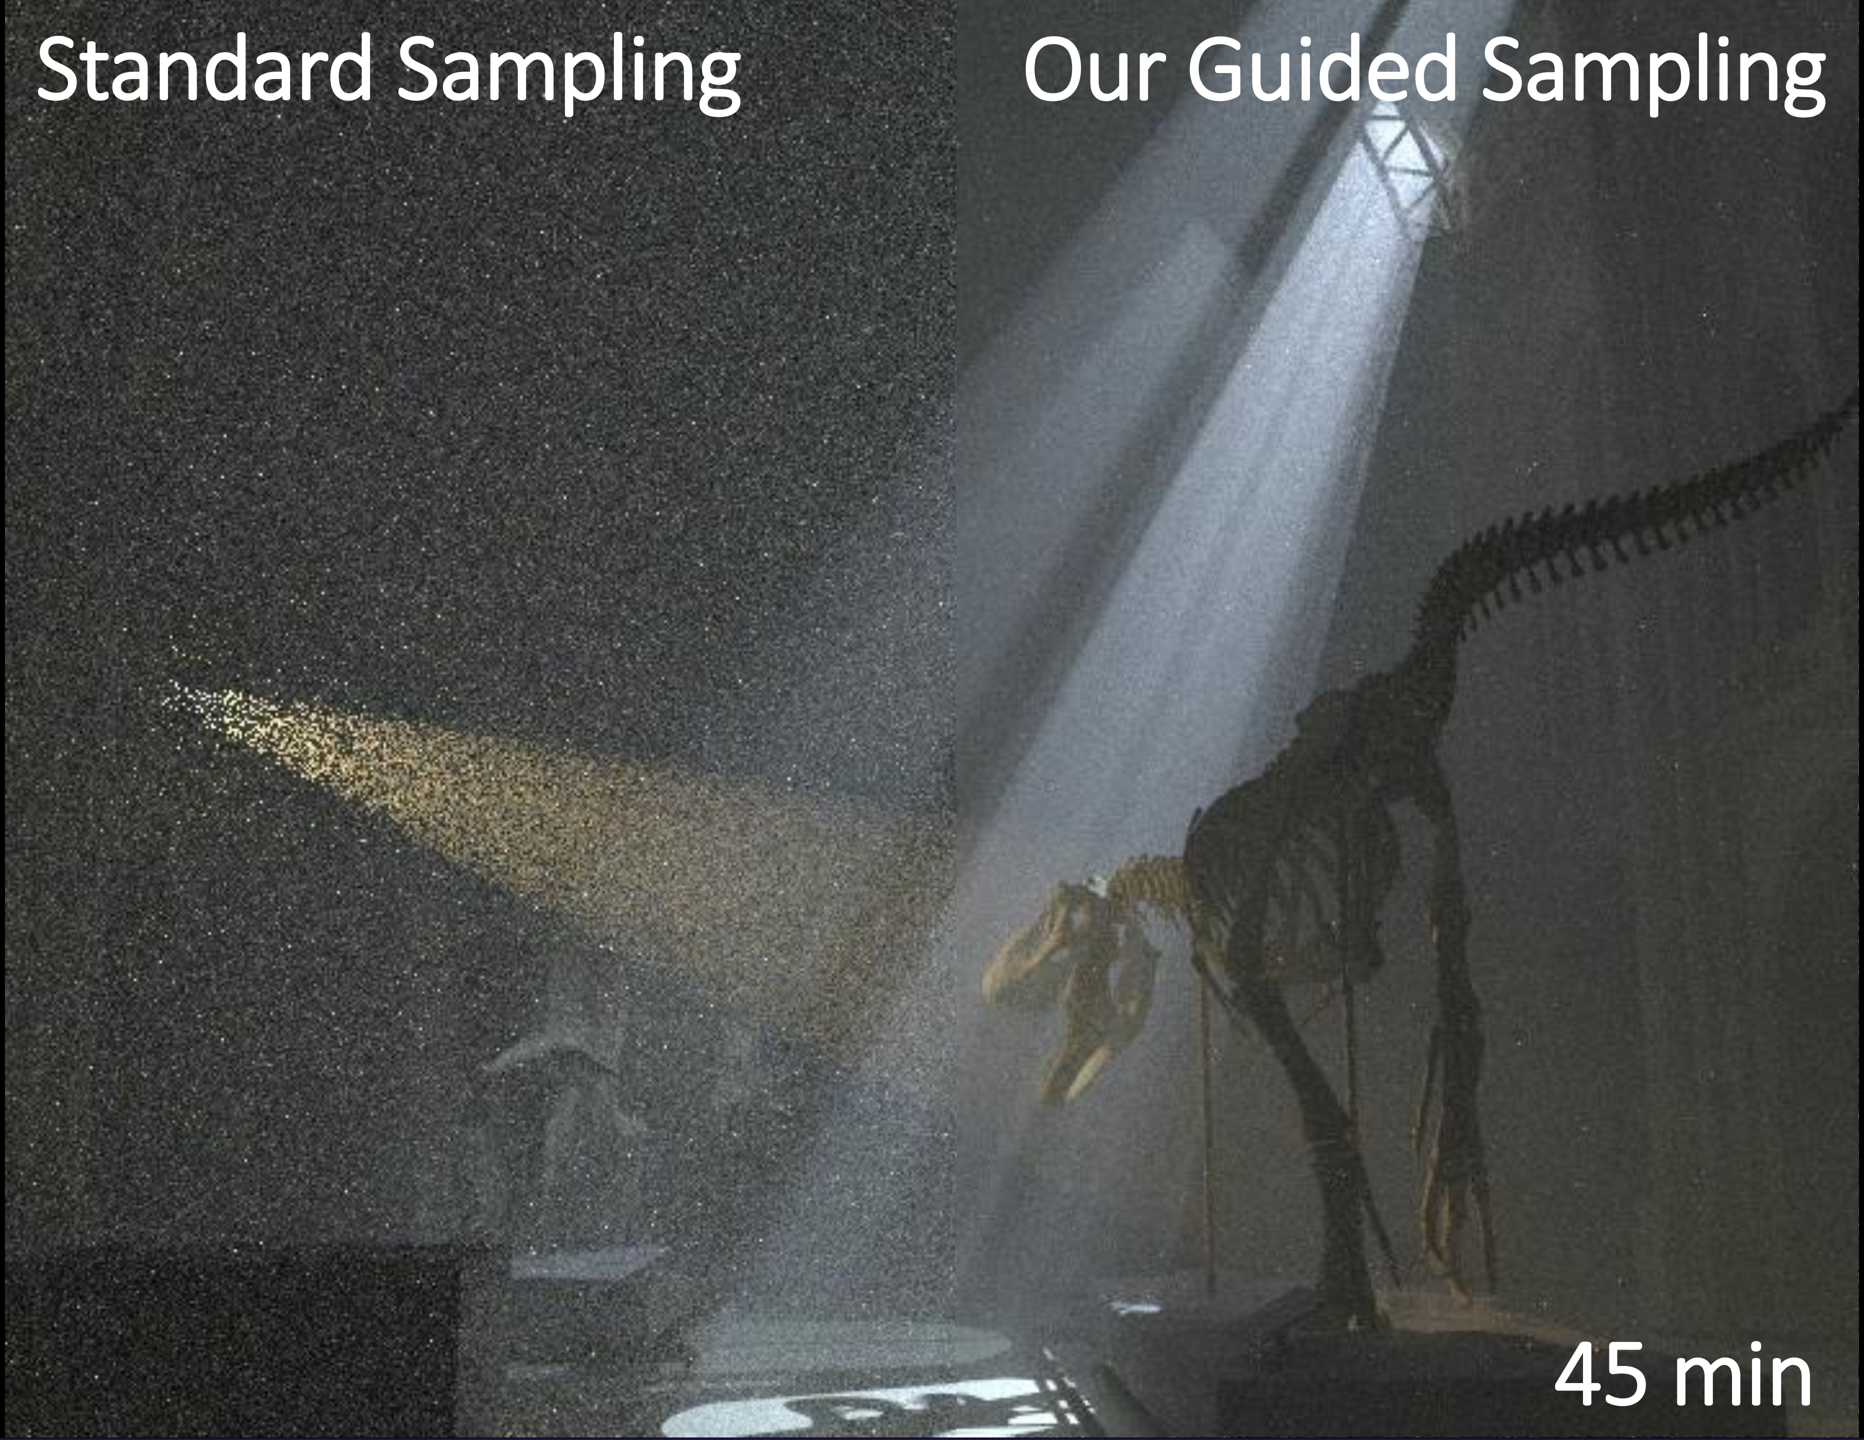
Standard Sampling
Our Guided Sampling
45 min

VOLUME RADIANCE ESTIMATES

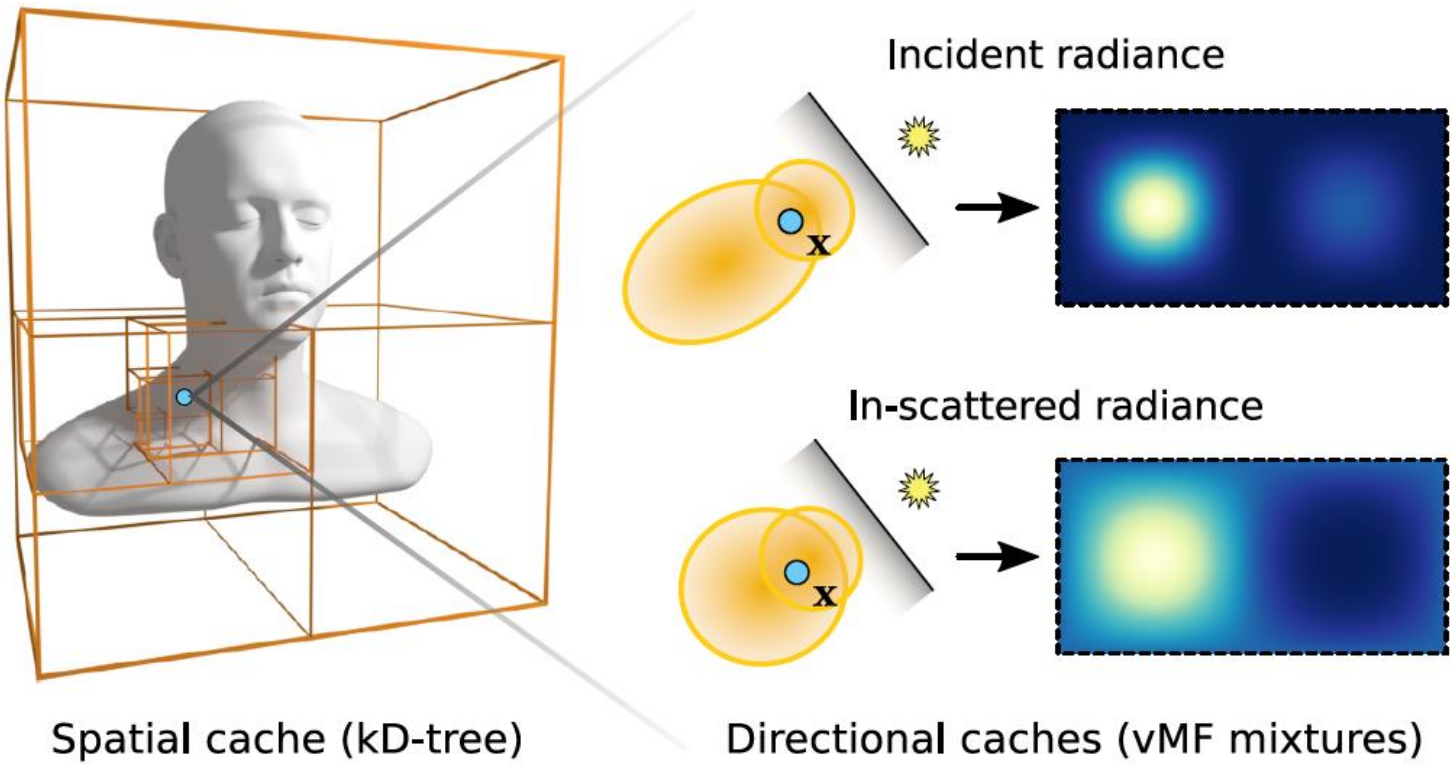

- Pre-processing step to fit estimates from photons (50M)
- Spatial caches via BSP-tree: max. 2K photons per node

VON MISES FISHER MIXTURE MODEL (VMM)

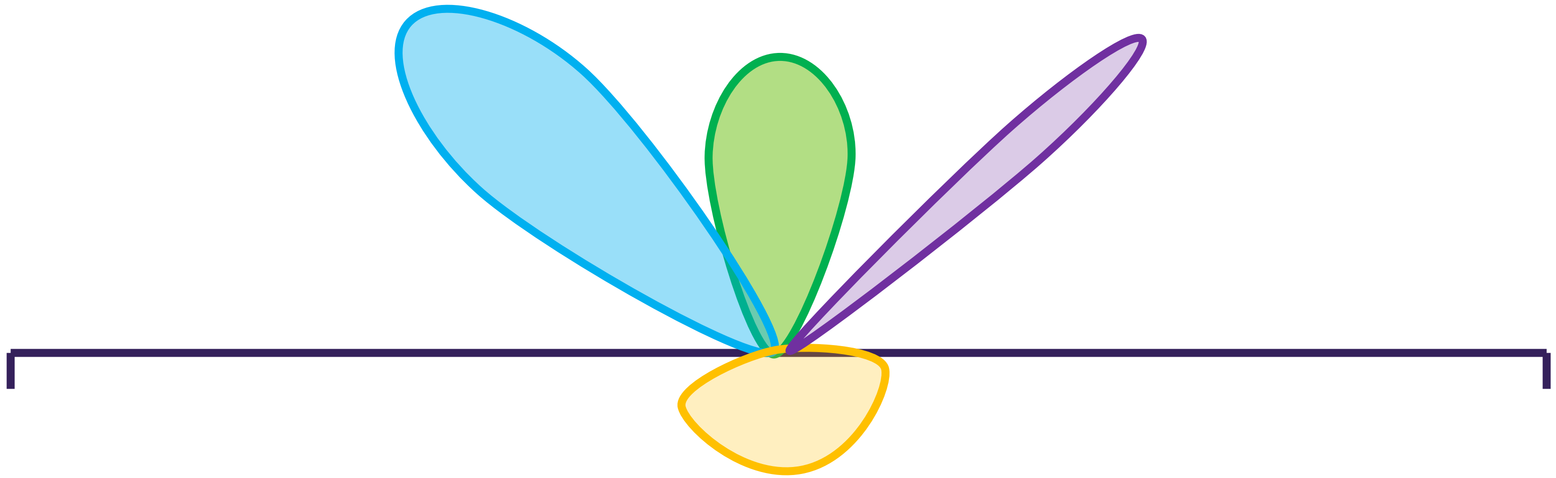

$$V(\omega|\Theta) = \sum^{K} \pi_i v(\omega|\mu_i, \kappa_i)$$

$$\Theta = \{\pi_0 \ldots, \mu_0 \ldots, \kappa_0 \ldots\}$$

VON MISES FISHER MIXTURE MODEL (VMM)

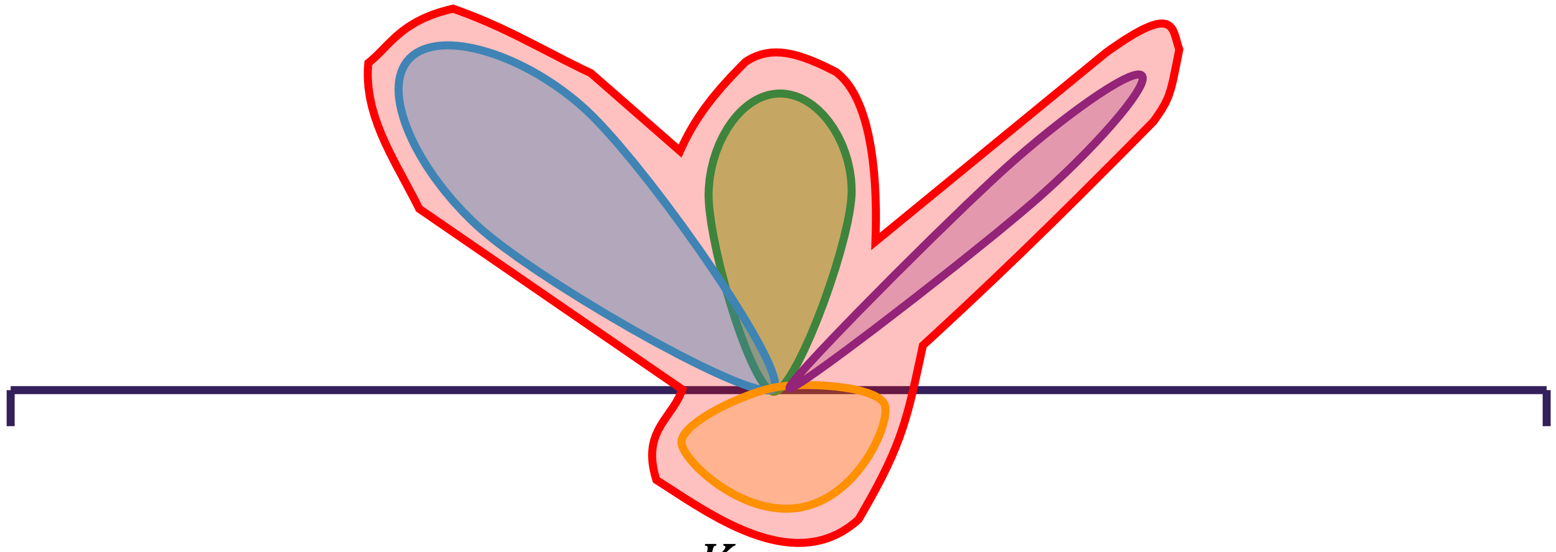

$$V(\omega|\Theta) = \sum^{K} \pi_i v(\omega|\mu_i, \kappa_i)$$

$$\Theta = \{\pi_0 \ldots, \mu_0 \ldots, \kappa_0 \ldots\}$$

VMM: PRODUCT

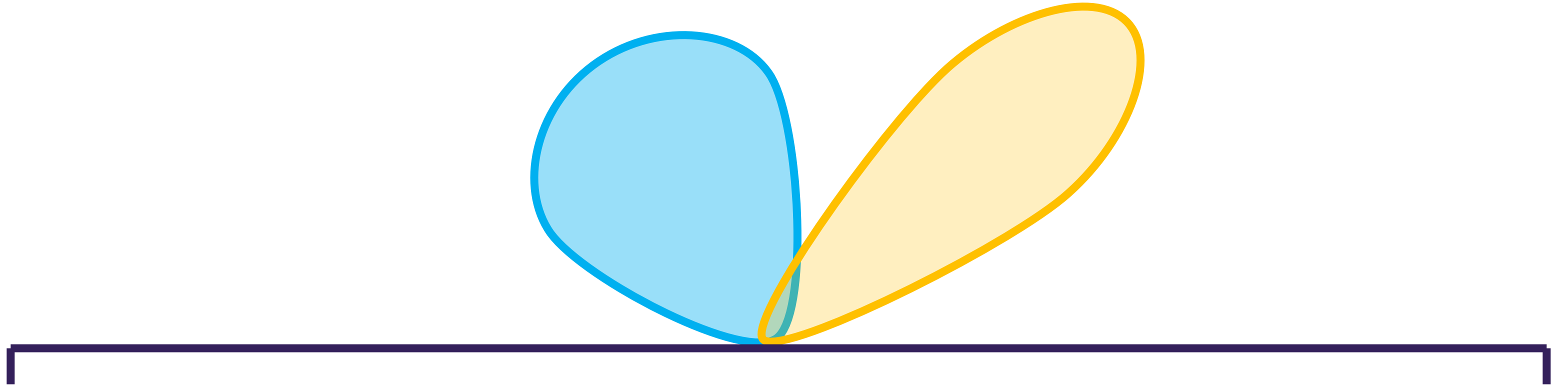

$$\pi_i v(\omega|\mu_i,\kappa_i) \cdot \pi_j v(\omega|\mu_j,\kappa_j) = \pi_{ij} v(\omega|\mu_{ij},\kappa_{ij})$$

VMM: PRODUCT

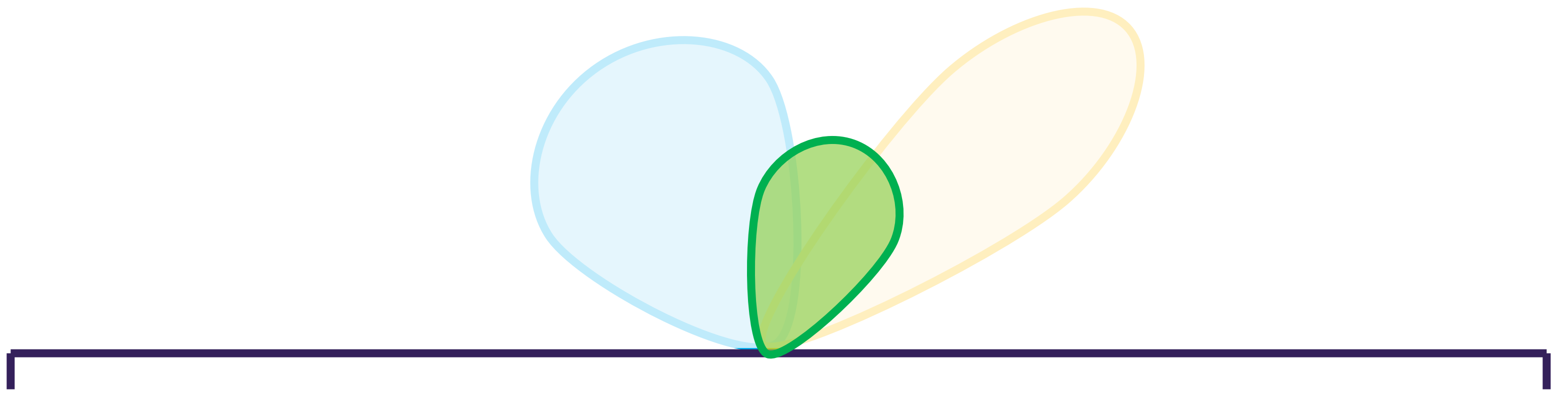

$$\pi_i v(\omega|\mu_i,\kappa_i) \cdot \pi_j v(\omega|\mu_j,\kappa_j) = \pi_{ij} v(\omega|\mu_{ij},\kappa_{ij})$$

VMM: CONVOLUTION

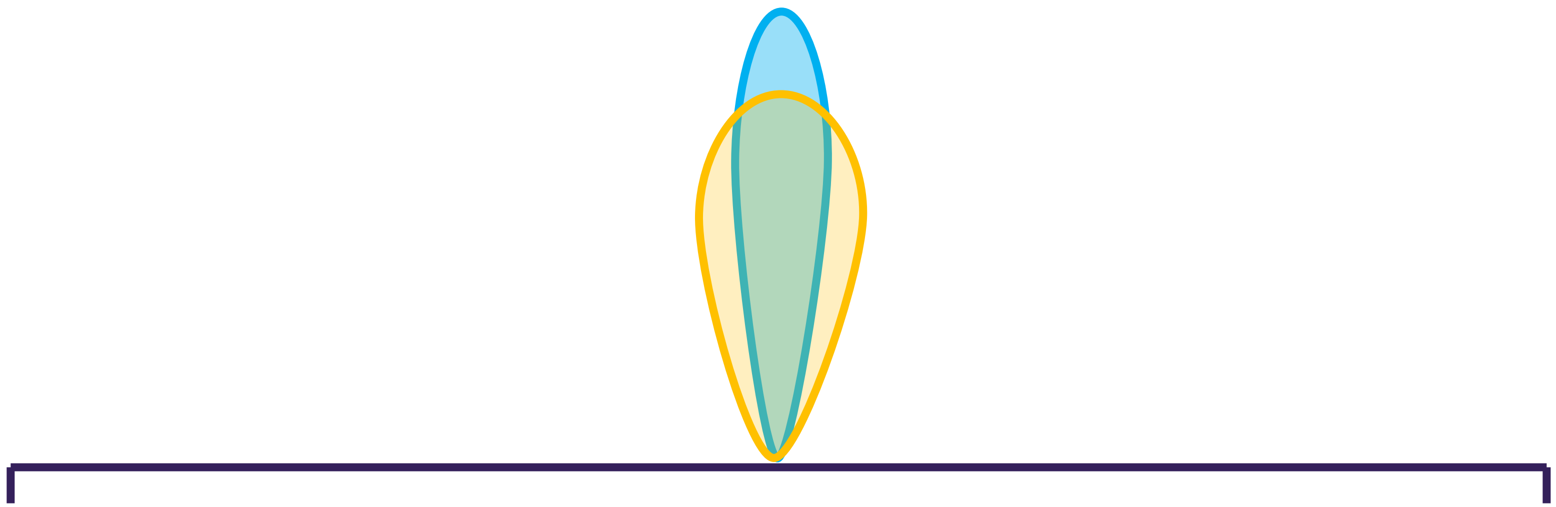

$$v_i(\ldots) * v_j(\ldots) = v_{ij}(\ldots)$$

VMM: CONVOLUTION

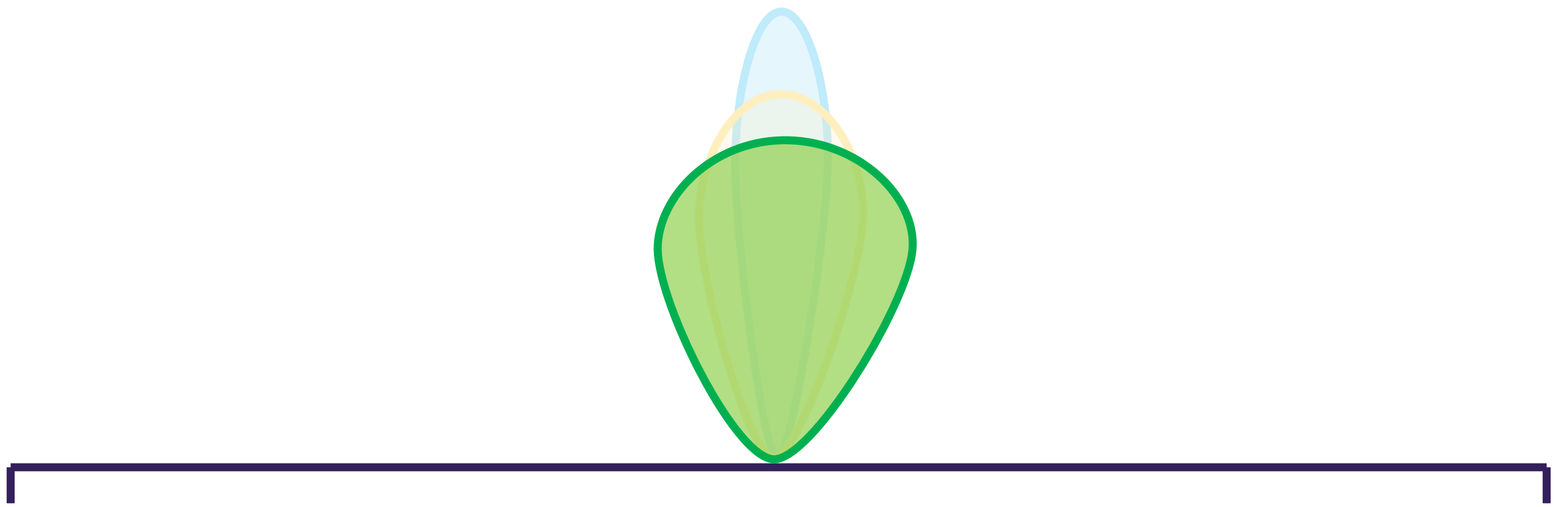

$$v_i(\ldots) * v_j(\ldots) = v_{ij}(\ldots)$$

INCIDENT RADIANCE ESTIMATES

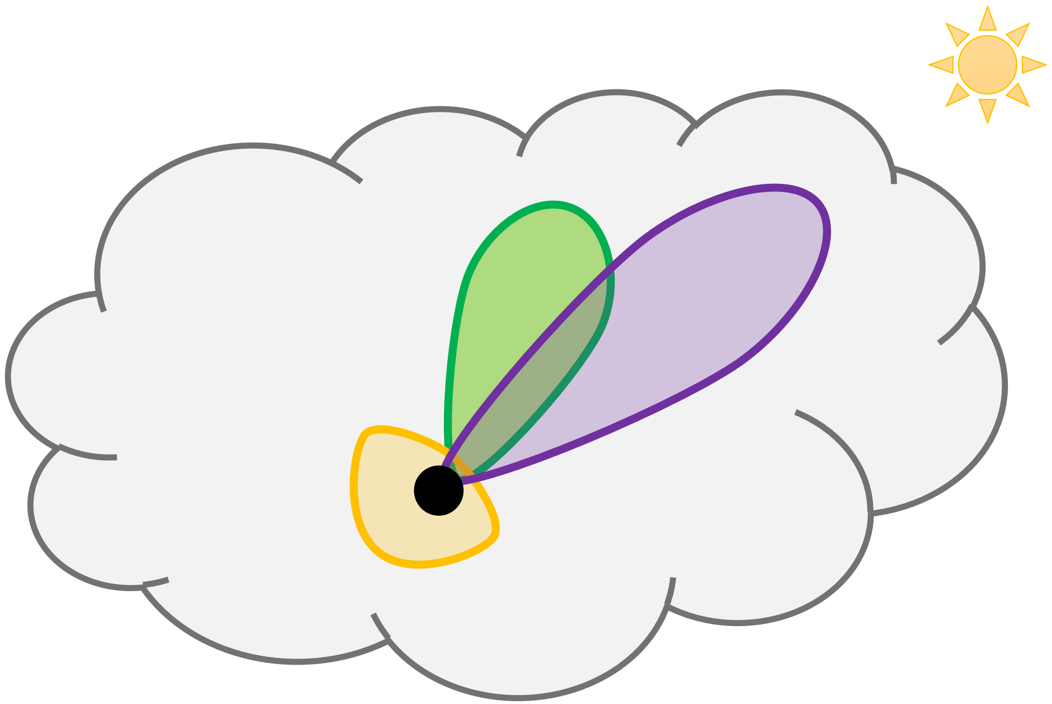

- Scaled Incident Radiance Distribution:

$$\tilde{L}(x, \omega) = \Phi(x) \cdot V(\omega | \Theta(x))$$

- Fluence:

$$\Phi(x) = \int_S L(x, \omega')d\omega'$$

INCIDENT RADIANCE ESTIMATES

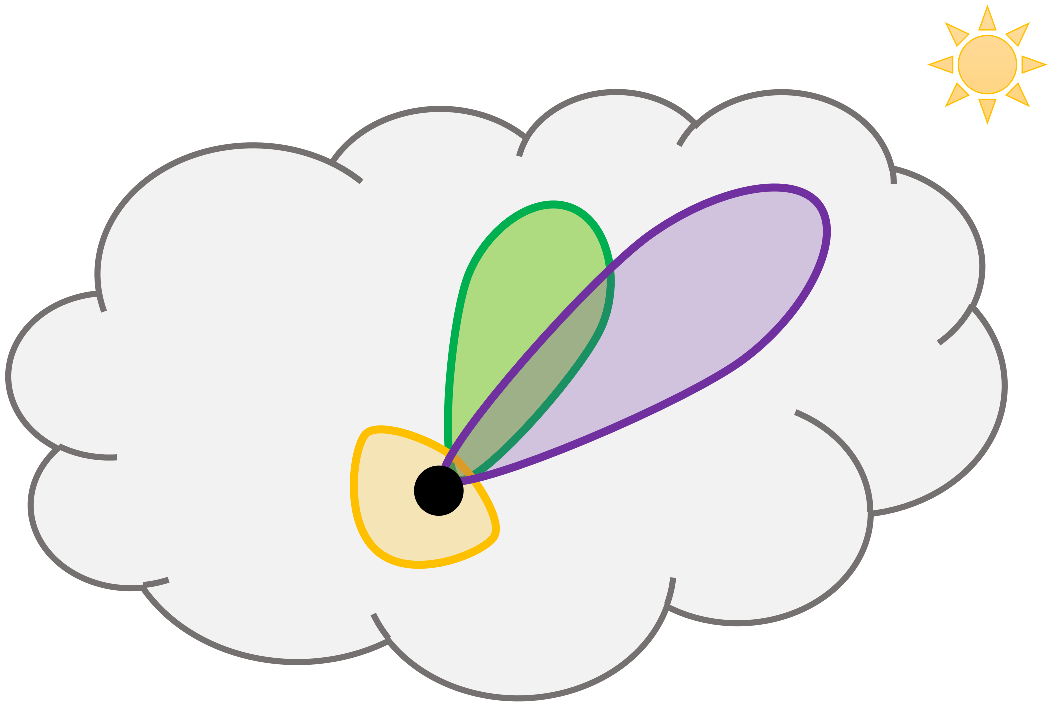

- Scaled Incident Radiance Distribution:

$$\tilde{L}(x, \omega) = \Phi(x) \cdot V(\omega|\Theta(x))$$

- Fluence:

$$\Phi(x) = \int_S L(x, \omega')d\omega'$$

INCIDENT RADIANCE ESTIMATES

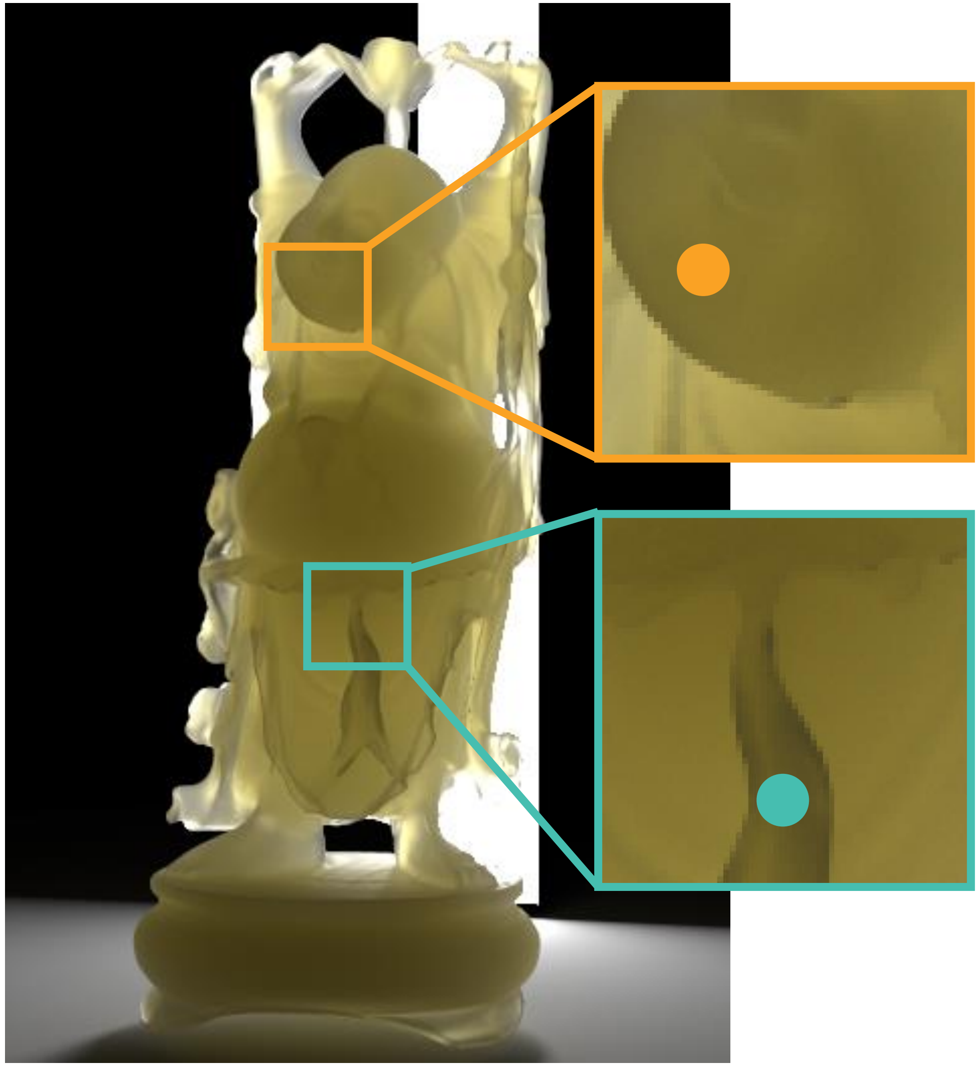

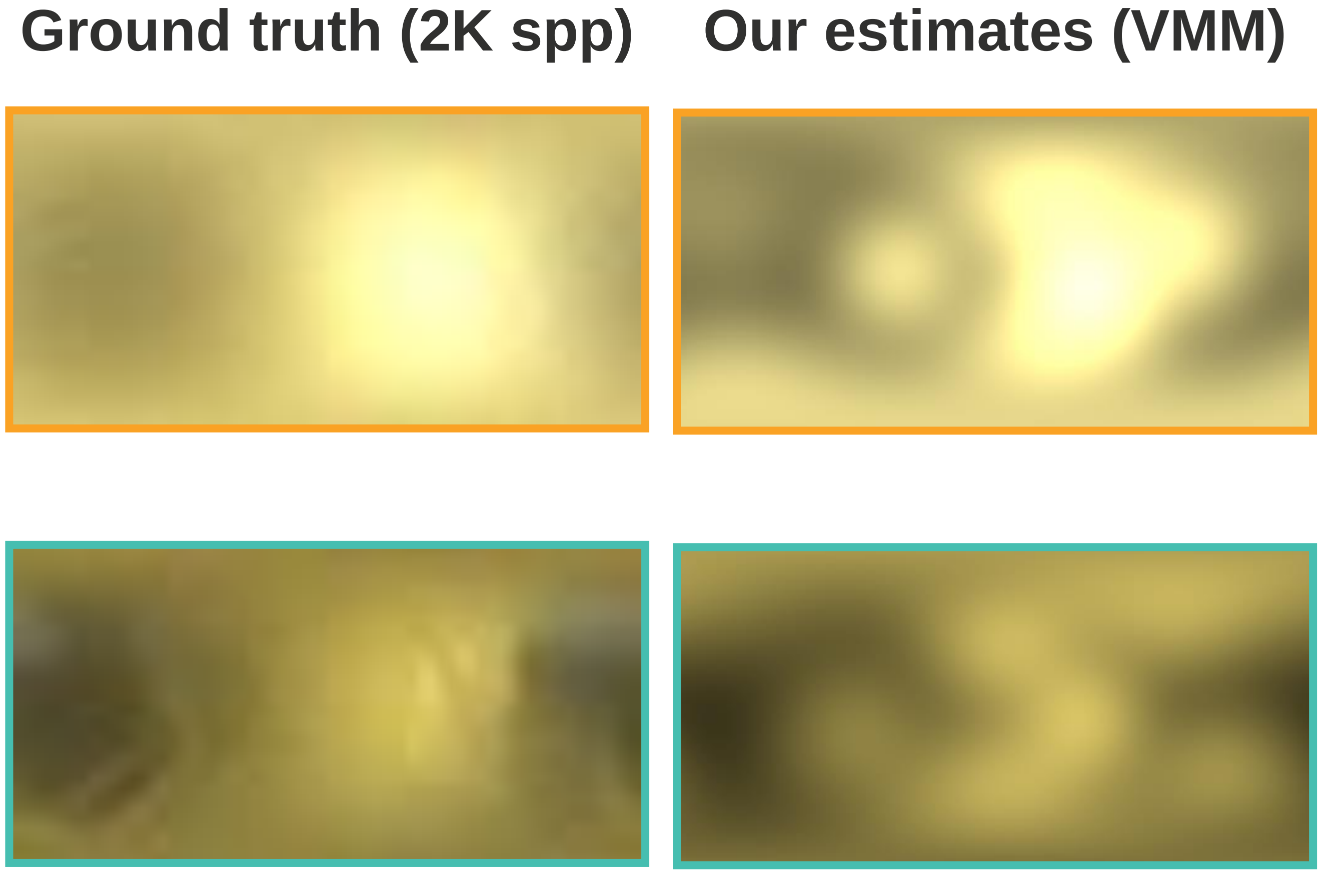

IN-SCATTERED RADIANCE ESTIMATES

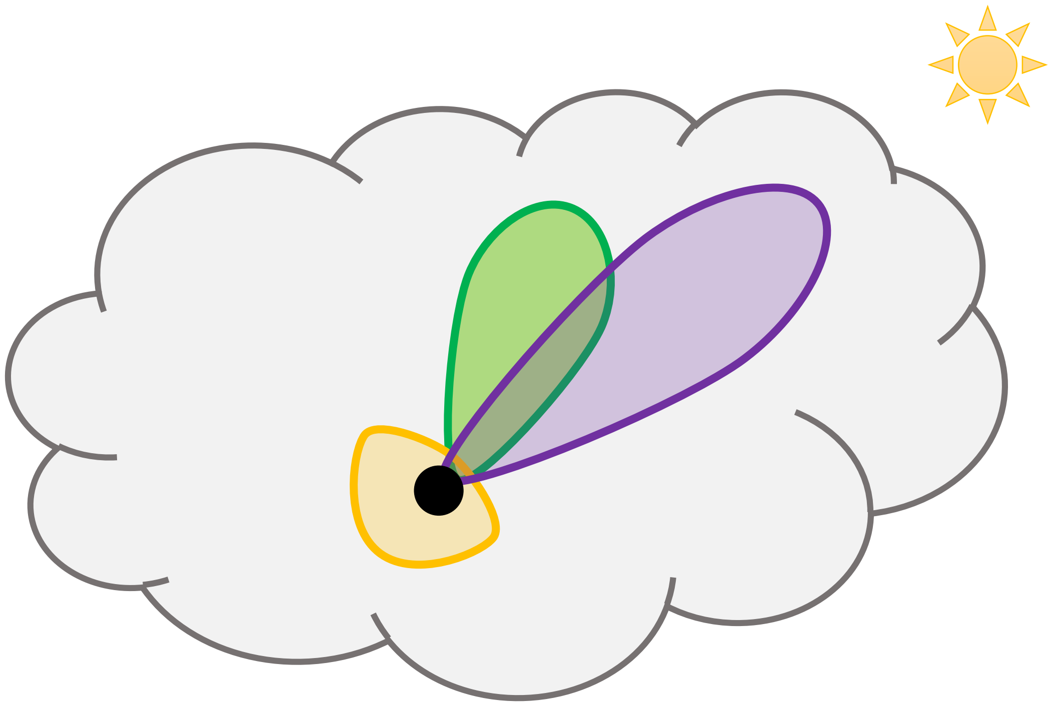

- Convolution between L and phase function:

IN-SCATTERED RADIANCE ESTIMATES

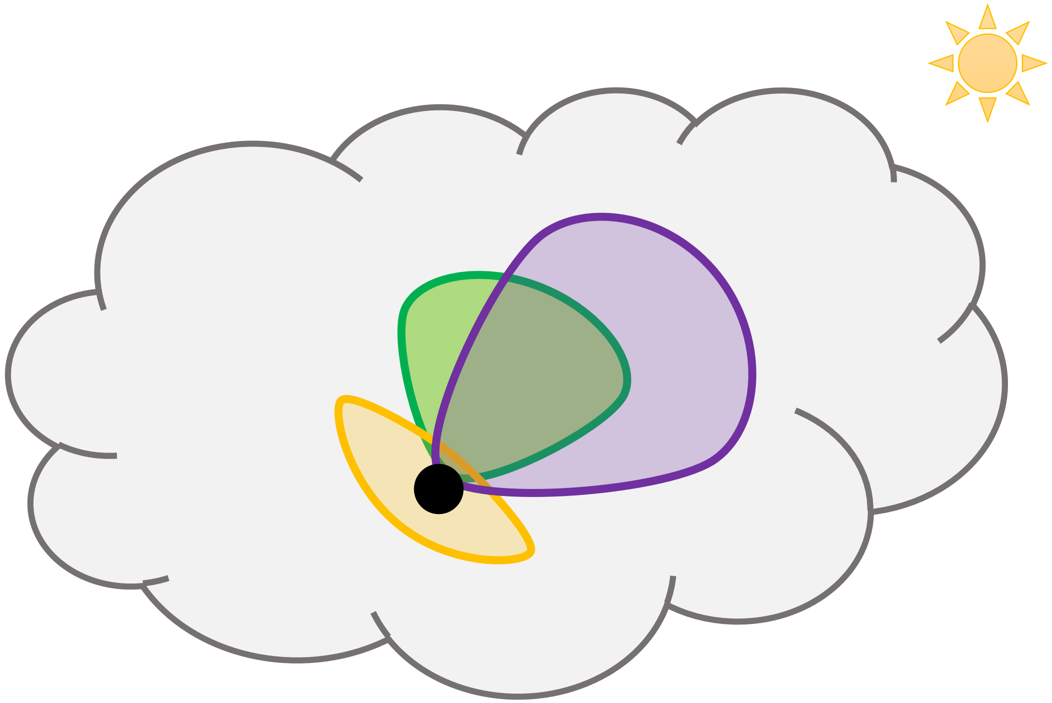

- Convolution between L and phase function:

$$V_{L_i}(\omega) = (V_f * V_L)(\omega)$$

IN-SCATTERED RADIANCE ESTIMATES

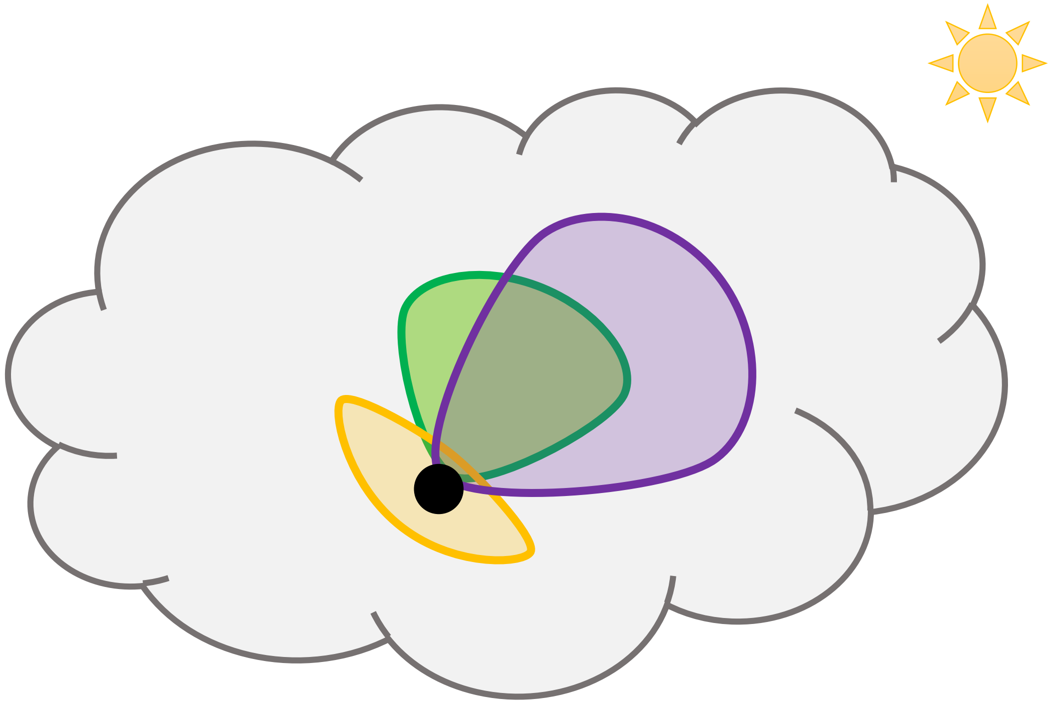

- Convolution between L and phase function:

$$V_{L_i}(\omega) = (V_f * V_L)(\omega)$$

$$\tilde{L}_i(x,\omega) = \Phi(x) \cdot V_{L_i}(\omega | \Theta_{L_i}(x))$$

IN-SCATTERED RADIANCE ESTIMATES

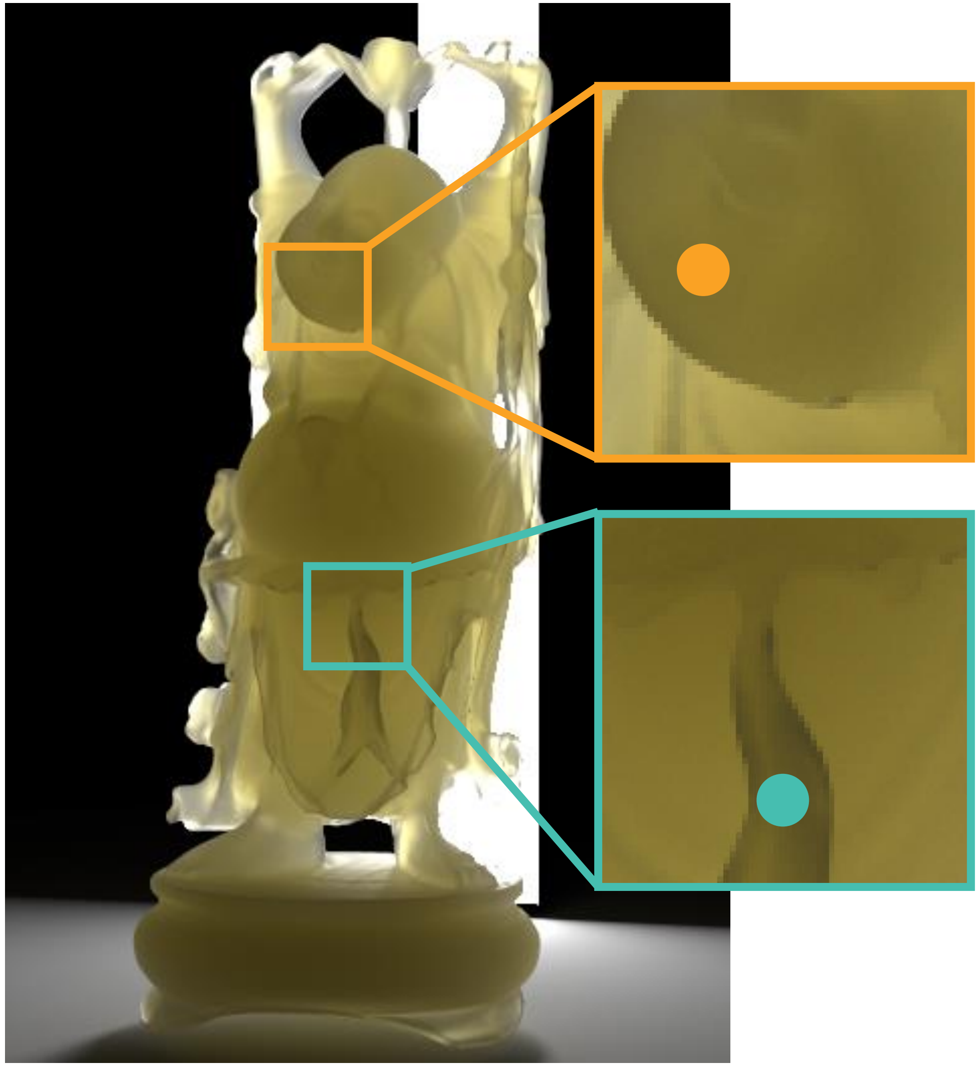

Ground truth (2K spp) | Our estimates (VMM)

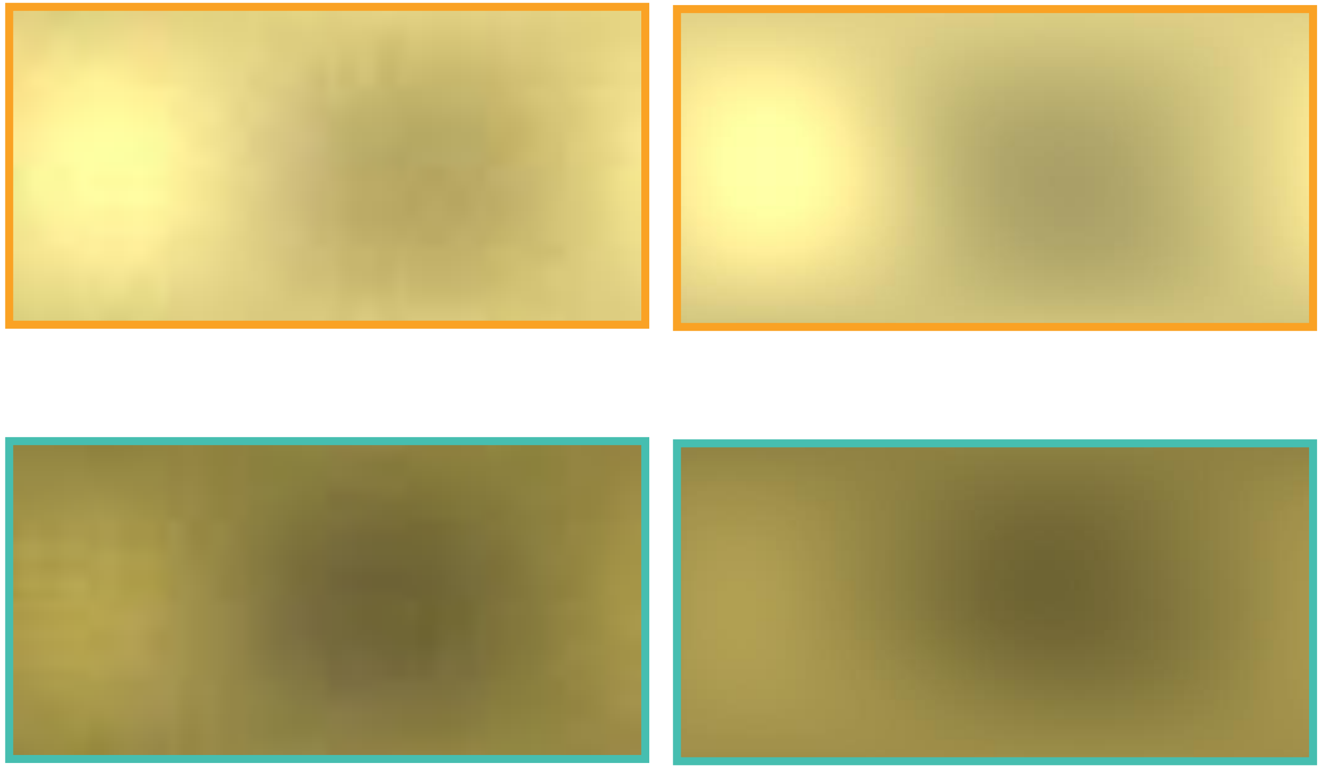

GUIDED SAMPLING DECISIONS

GUIDED PRODUCT DISTANCE SAMPLING

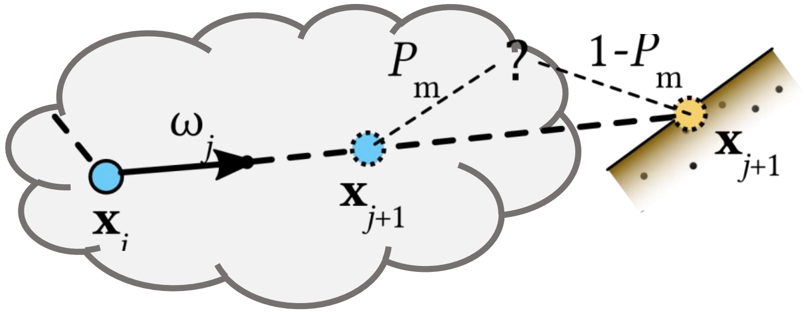

1. Scatter decision

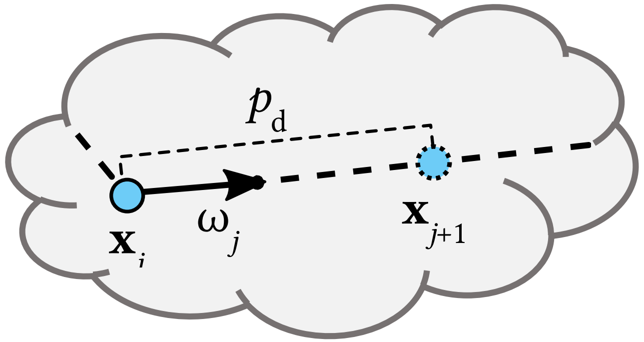

2. Scatter distance

GUIDED PRODUCT DISTANCE SAMPLING

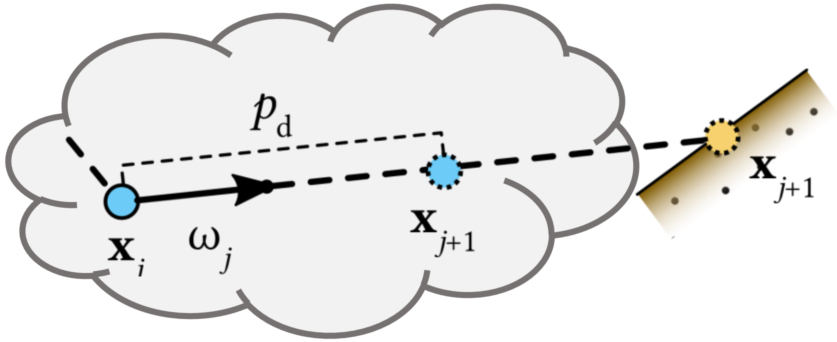

1+2. Event distance

Traditional distance sampling

- Optimal ZV-PDF:

$$p_d^{zv}(d_{j+1}|\boldsymbol{x}_j,\omega_j) = \frac{T(\boldsymbol{x}_j,\boldsymbol{x}_{j+1})\cdot\sigma_s(\boldsymbol{x}_{j+1})\cdot L_i(\boldsymbol{x}_{j+1},\omega_j)}{L(\boldsymbol{x}_j,\omega_j)}$$

Our estimates

- Our guided PDF:

$$\tilde{p}_d^{zv}(d_{j+1}|\boldsymbol{x}_j,\omega_j) = \frac{T(\boldsymbol{x}_j,\boldsymbol{x}_{j+1})\cdot\sigma_s(\boldsymbol{x}_{j+1})\cdot\tilde{L}_i(\boldsymbol{x}_{j+1},\omega_j)}{\tilde{L}(\boldsymbol{x}_j,\omega_j)}$$

INCREMENTAL GUIDED DISTANCE SAMPLING

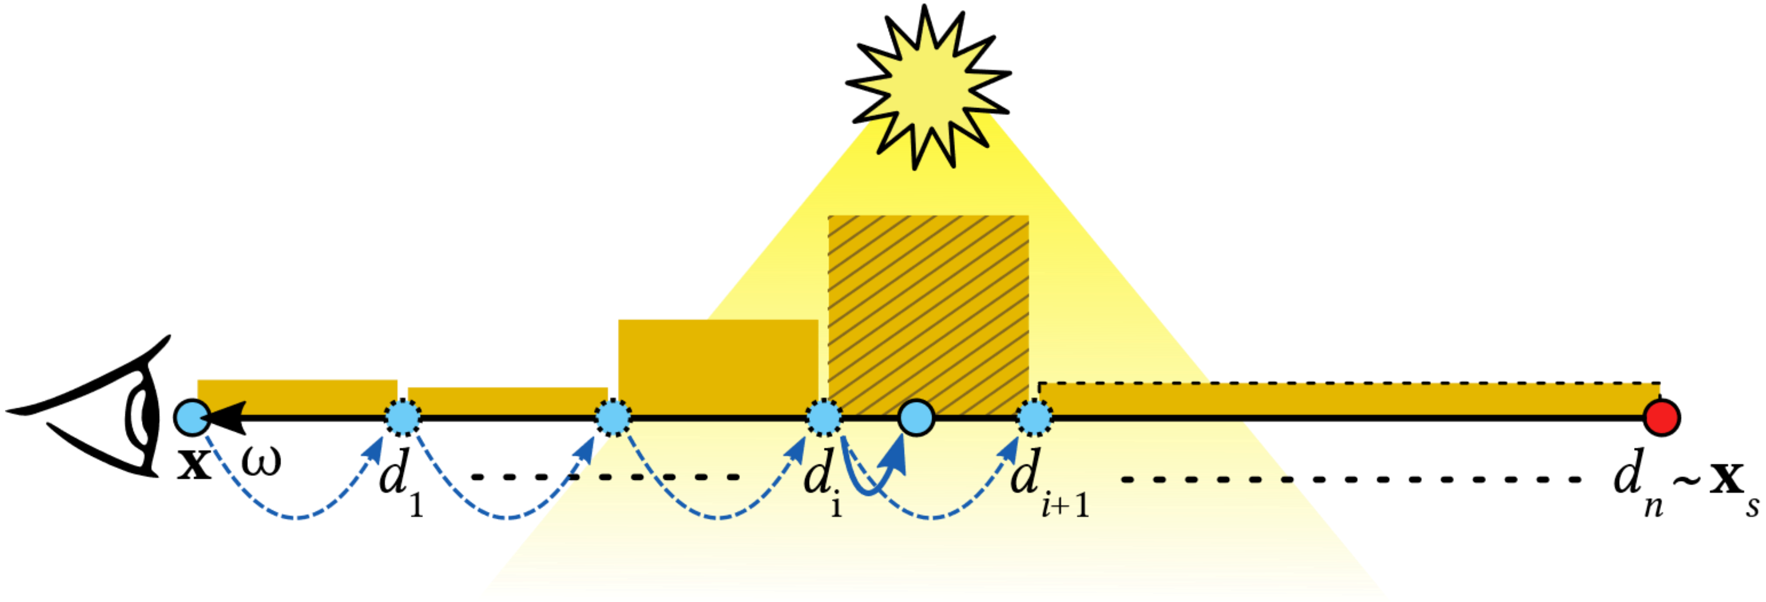

- Incremental approach:
  - At each step make a local decision, if we scatter inside the bin.

INCREMENTAL GUIDED DISTANCE SAMPLING

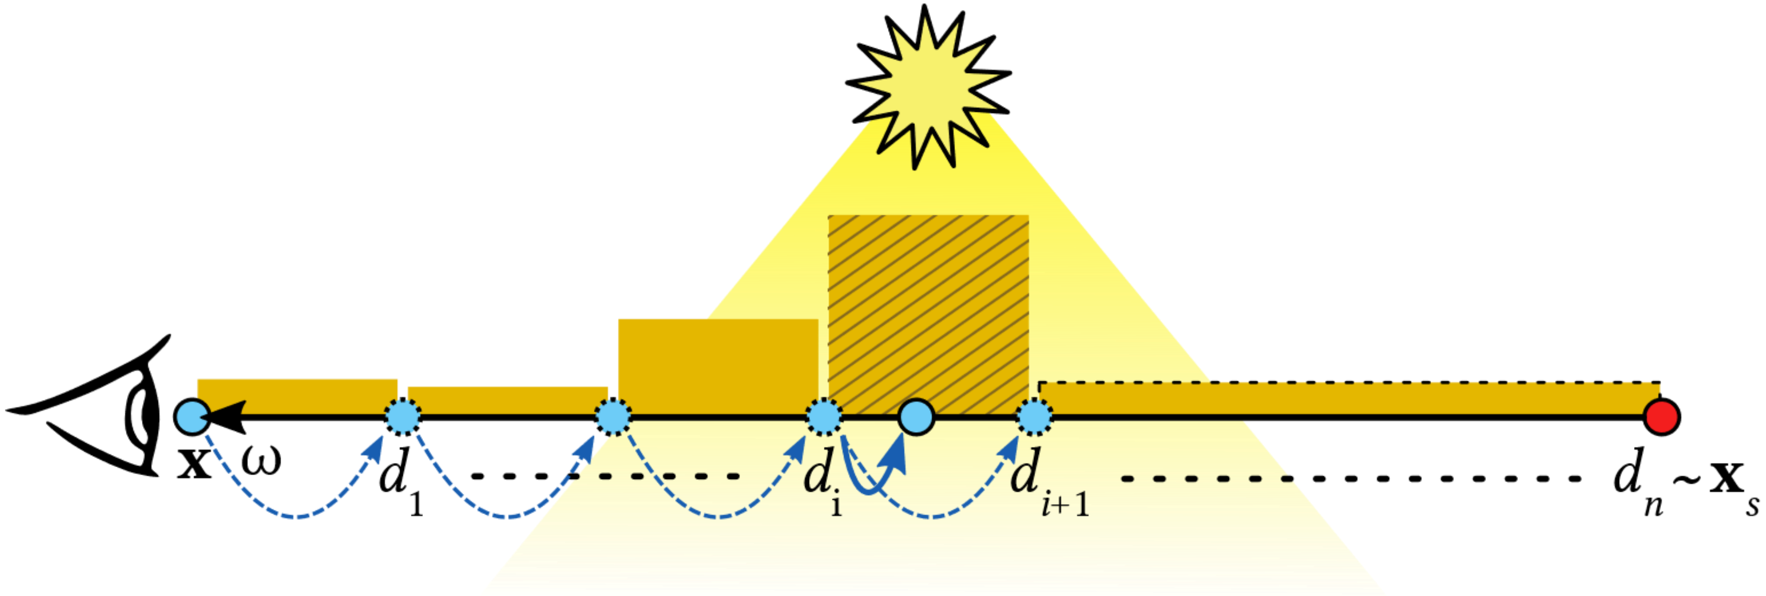

- Local bin scatter probability:

$$P_i(D \leq d_{i+1}) \approx \frac{1 - T(d_i, d_{i+1})}{\sigma_t(d_i)} \cdot \frac{\sigma_s(d_i) \cdot \tilde{L}_i(d_i)}{\tilde{L}(d_i, \omega)}$$

INCREMENTAL GUIDED DISTANCE SAMPLING

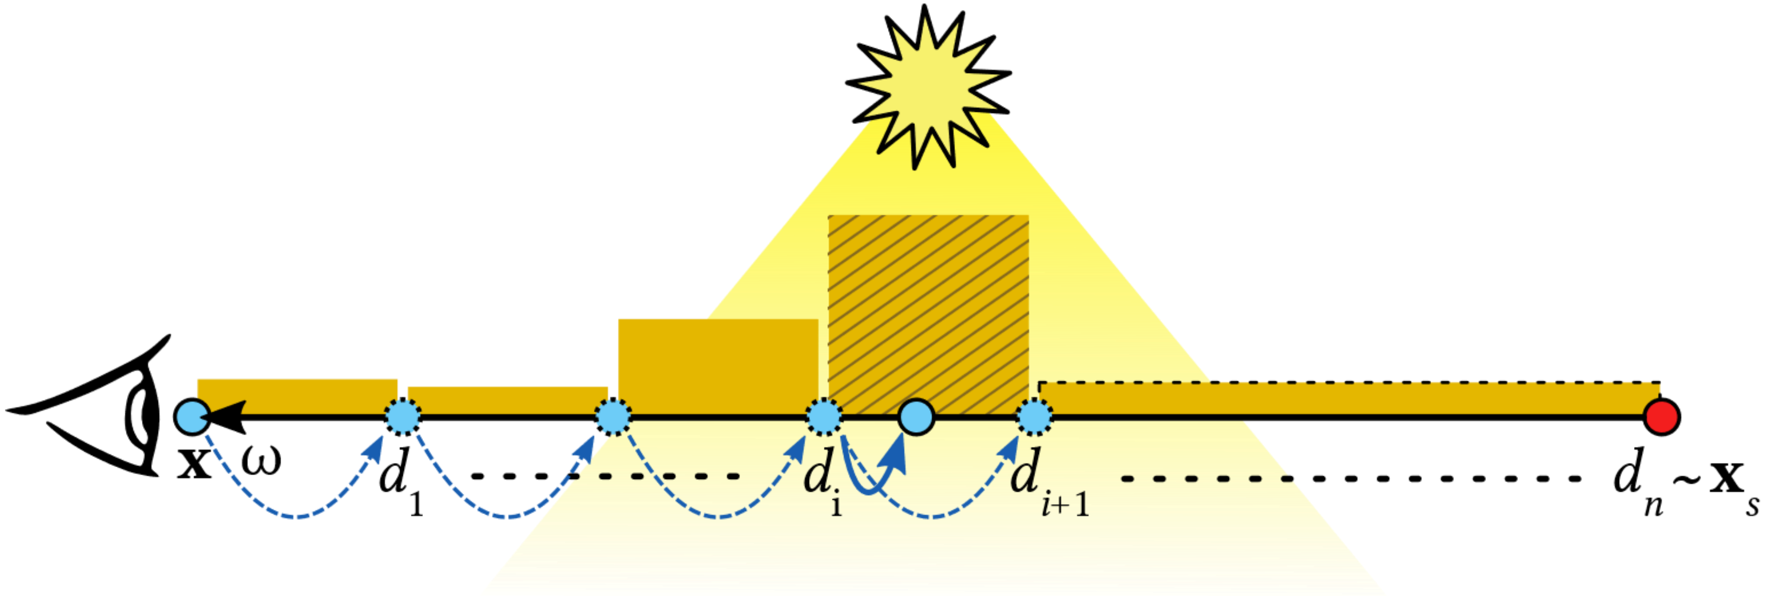

- We only need to step until the scattering event happens

No guiding (256 spp)

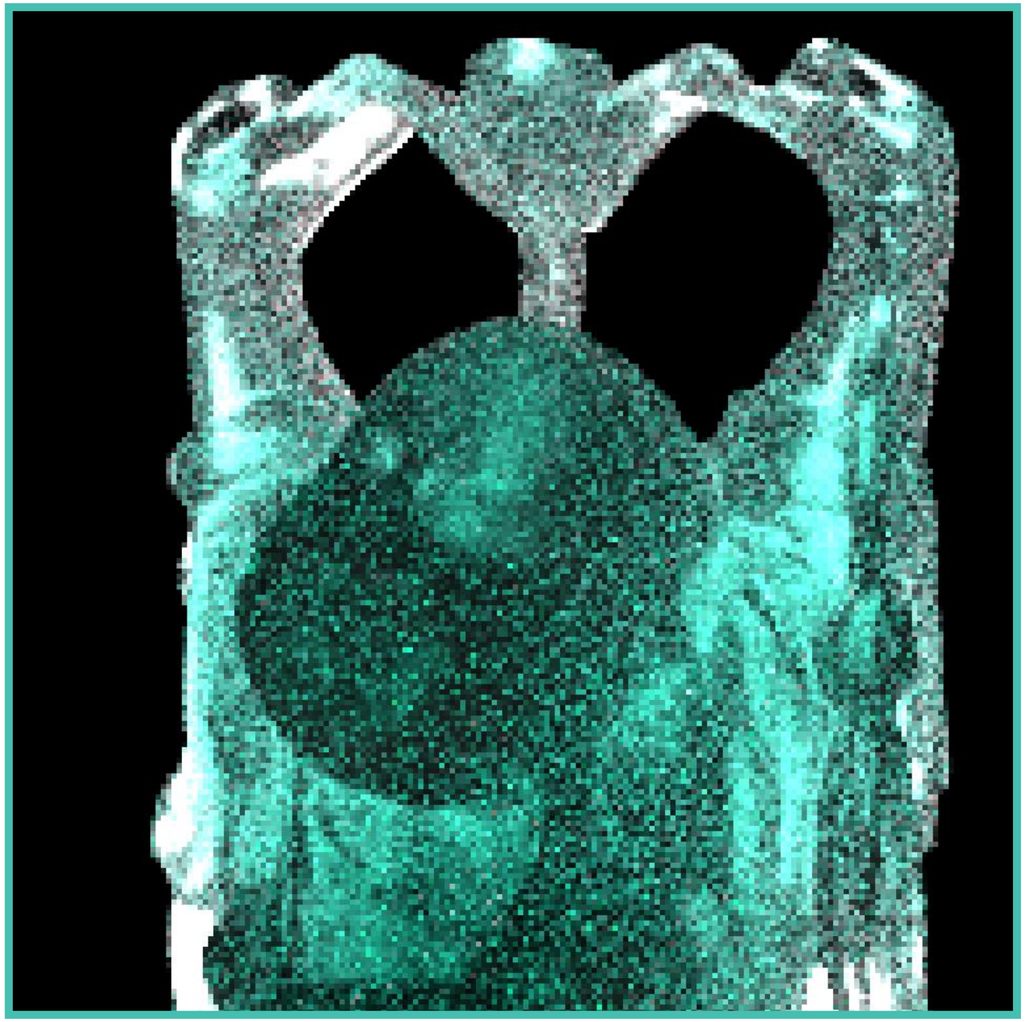

Distance guiding (256 spp)

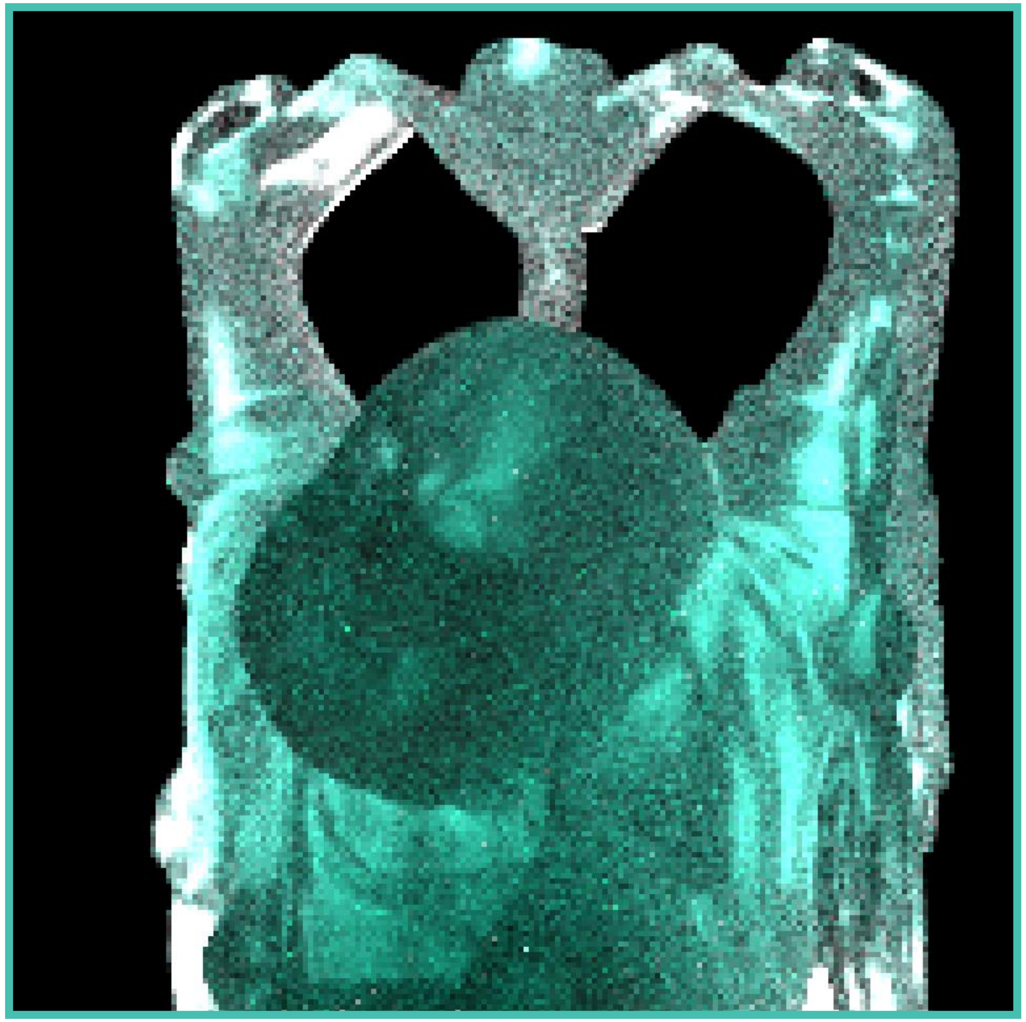

No guiding (1024 spp)

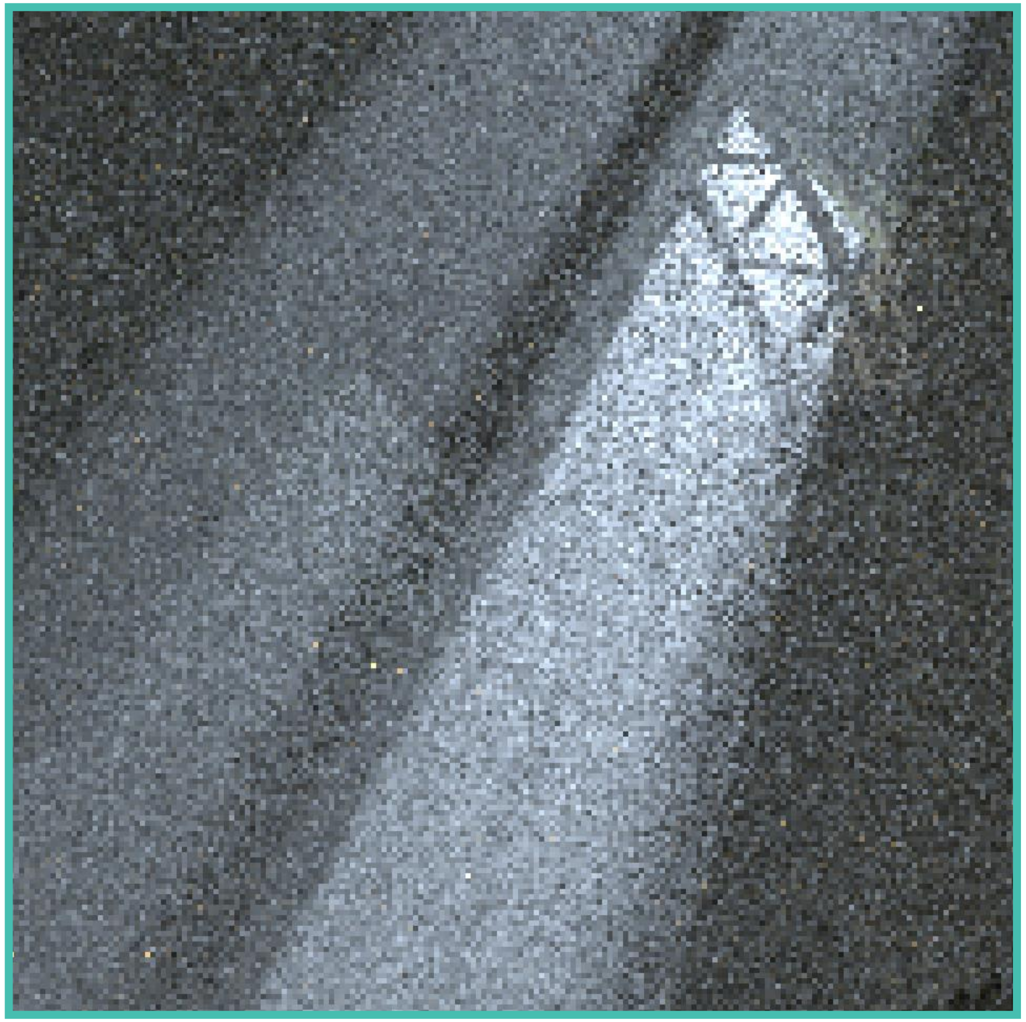

Distance guiding (1024 spp)

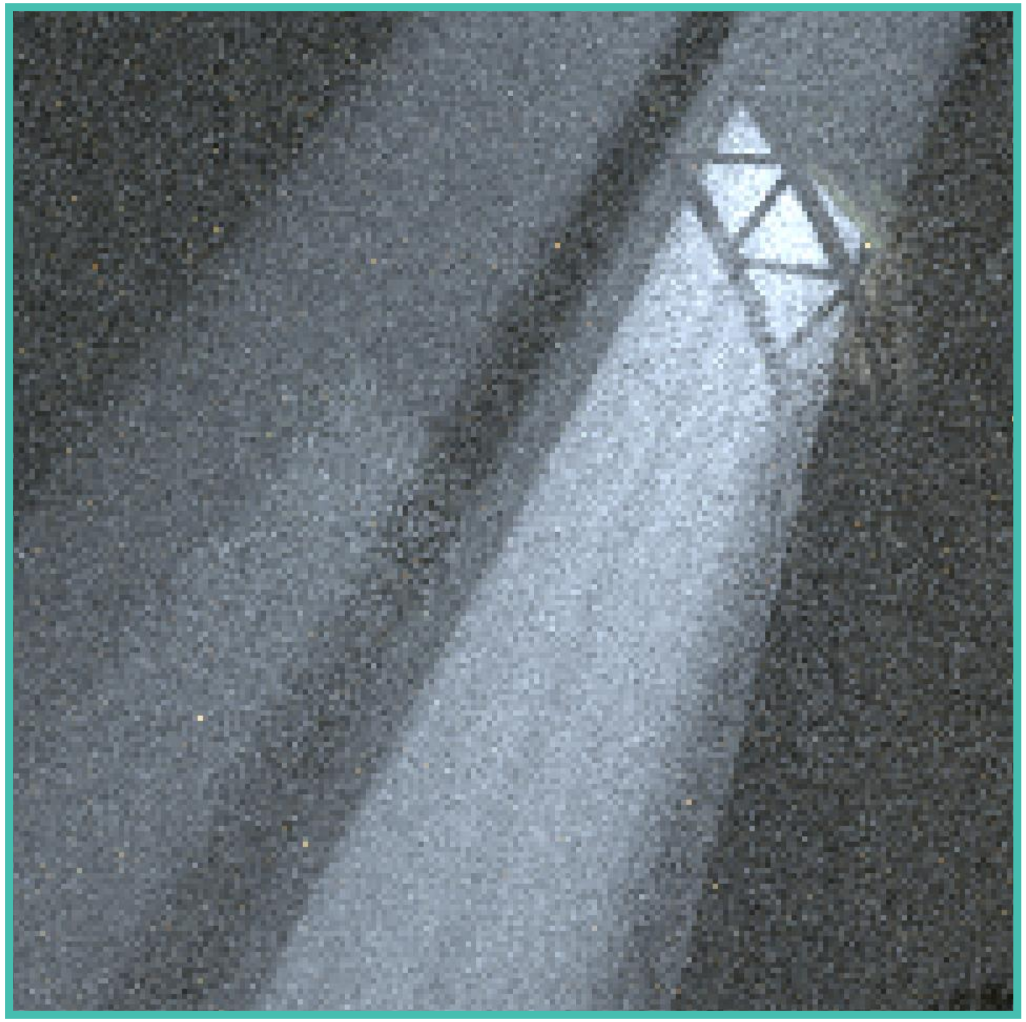


45 min

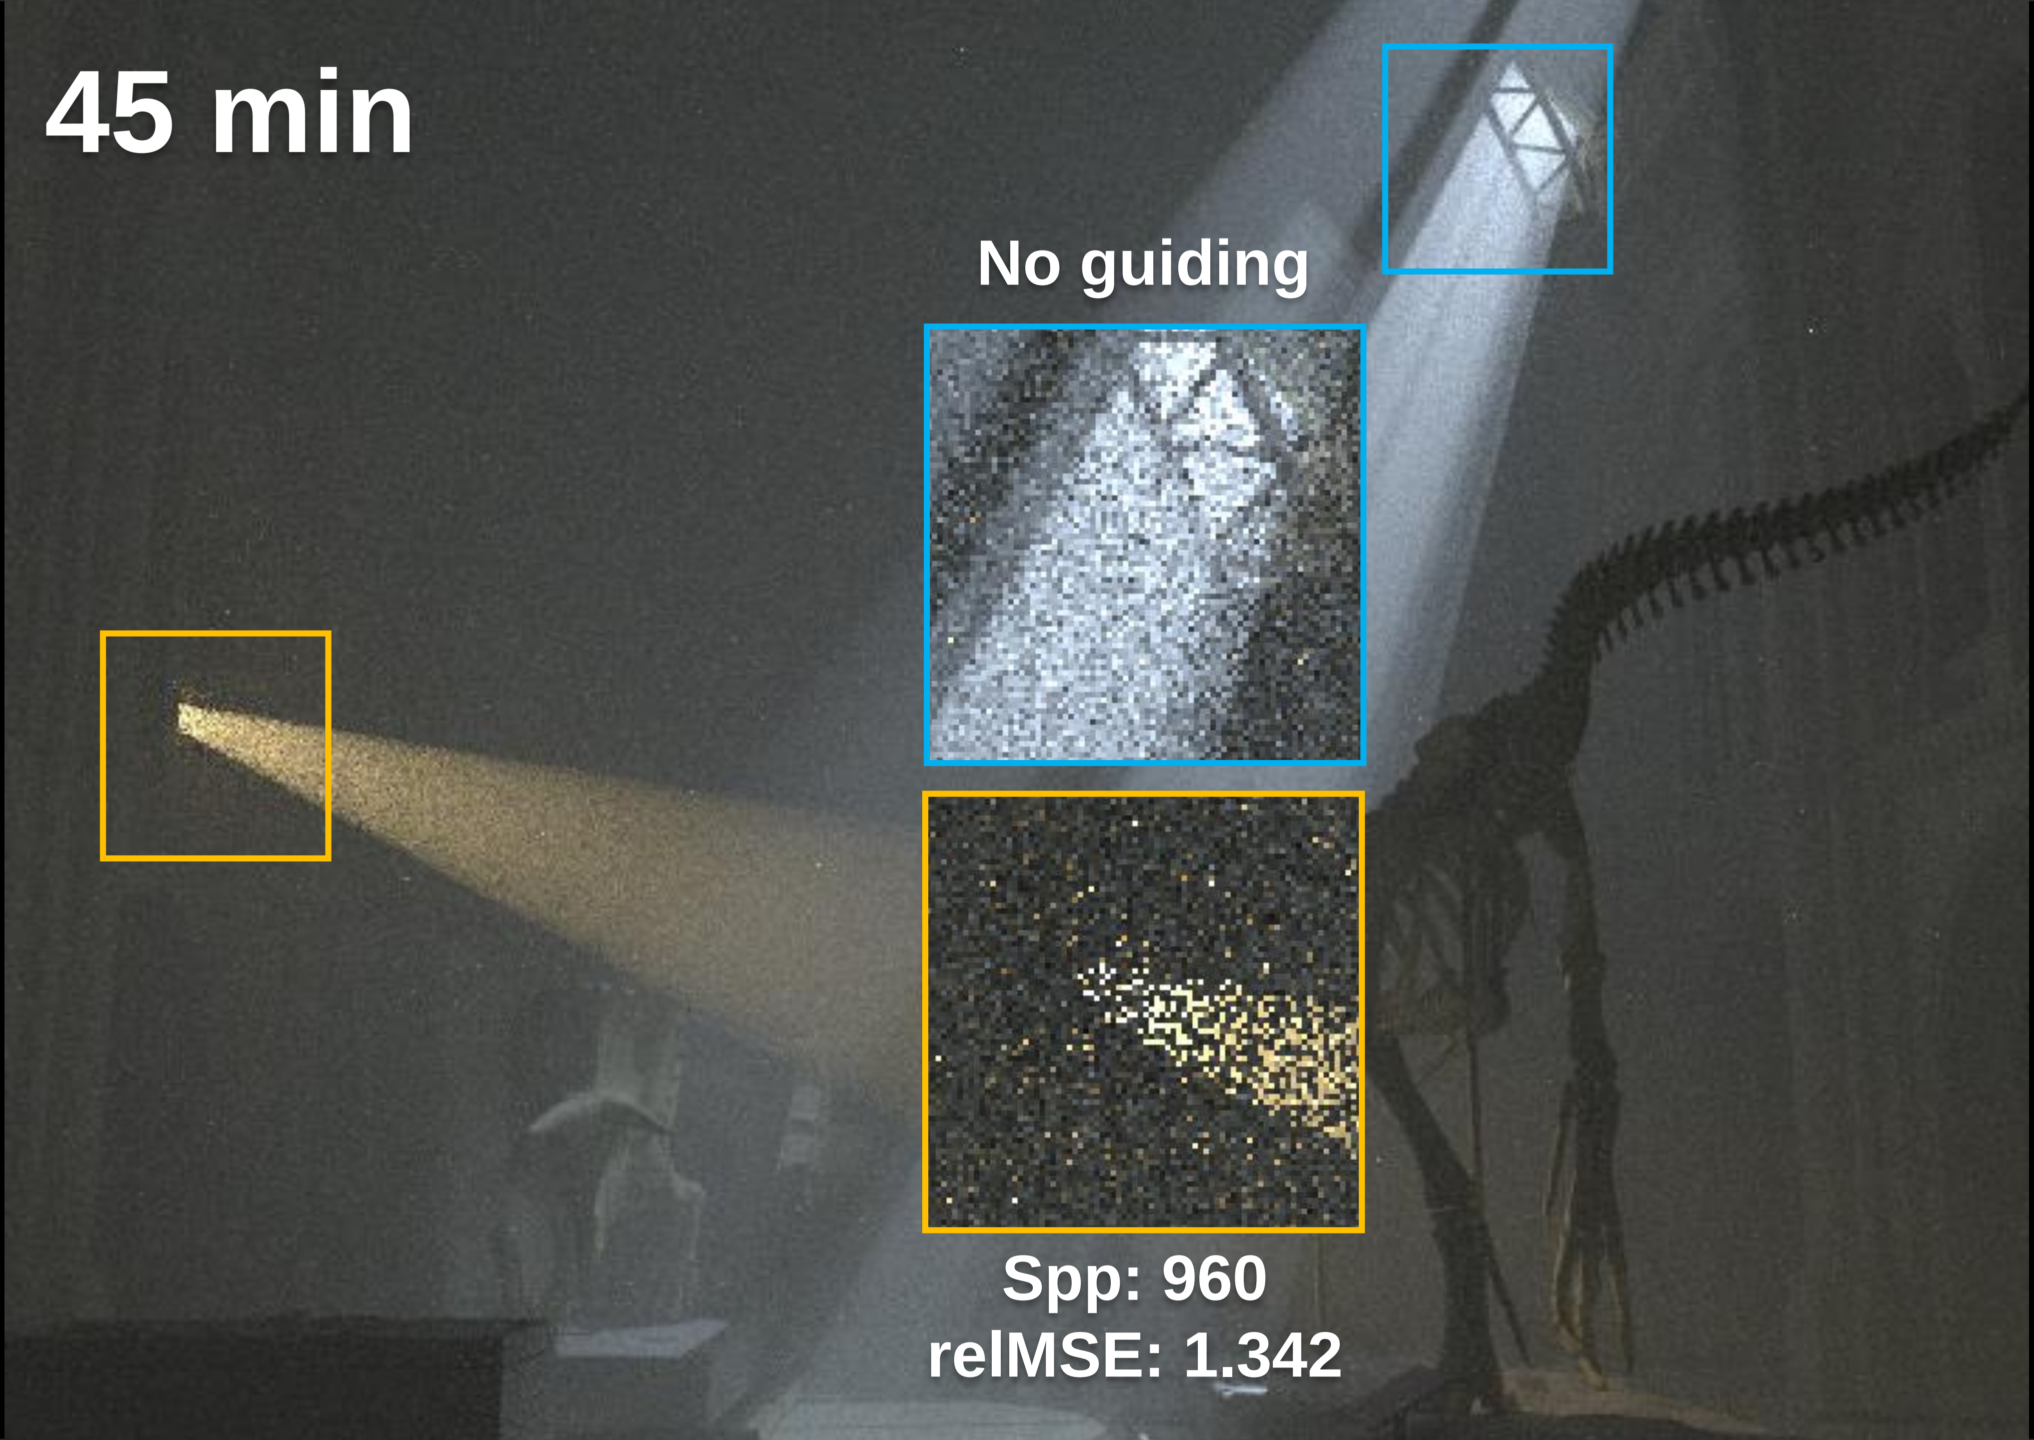
45 min
No guiding
Spp: 960
relMSE: 1.342

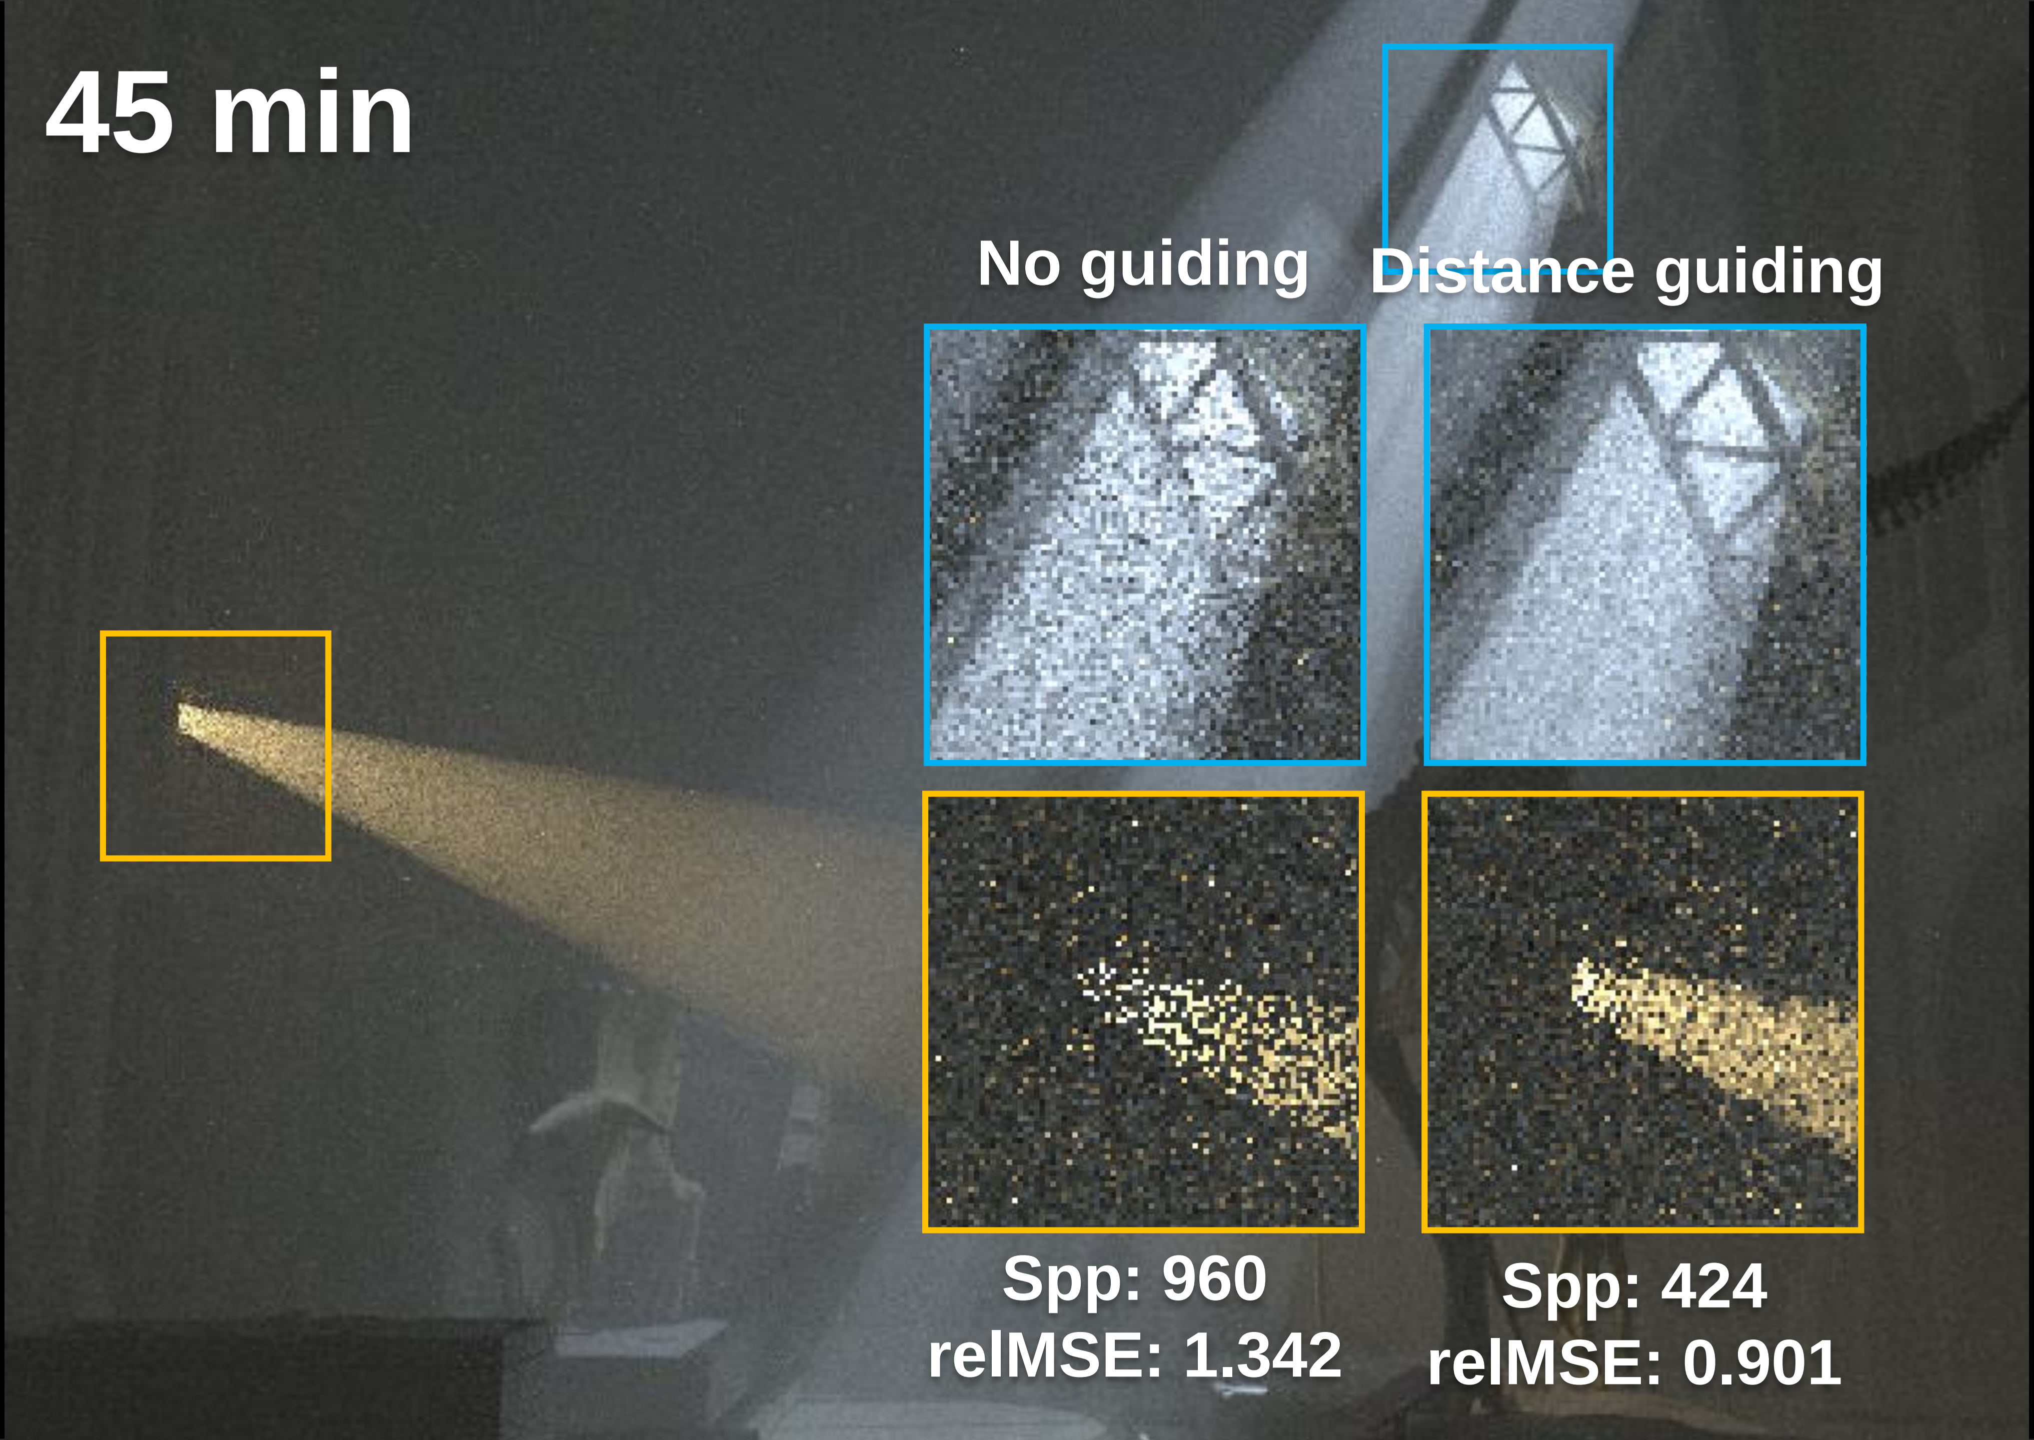
45 min
No guiding
Distance guiding
Spp: 960
relMSE: 1.342
Spp: 424
relMSE: 0.901

GUIDED PRODUCT DIRECTIONAL SAMPLING

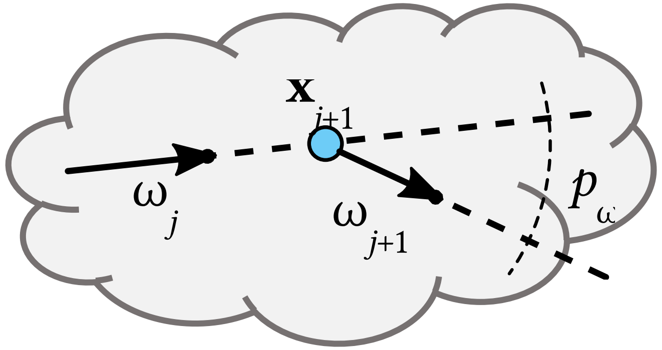

3. Scatter direction

Traditional directional sampling

- Optimal ZV-PDF:

$$p_{\omega}^{zv}(\omega_{j+1}|\boldsymbol{x}_{j+1},\omega_j) \propto \boxed{f(x_{j+1},\omega_j,\omega_{j+1})} \cdot L(x_{j+1},\omega_{j+1})$$

Our estimates

- Our guided PDF:

$$\tilde{p}_{\omega}^{zv}(\omega_{j+1}|\boldsymbol{x}_{j+1},\omega_j) \propto \boxed{\tilde{f}(x_{j+1},\omega_j,\omega_{j+1})} \cdot \boxed{\tilde{L}(x_{j+1},\omega_{j+1})}$$

GUIDED PRODUCT DIRECTIONAL SAMPLING

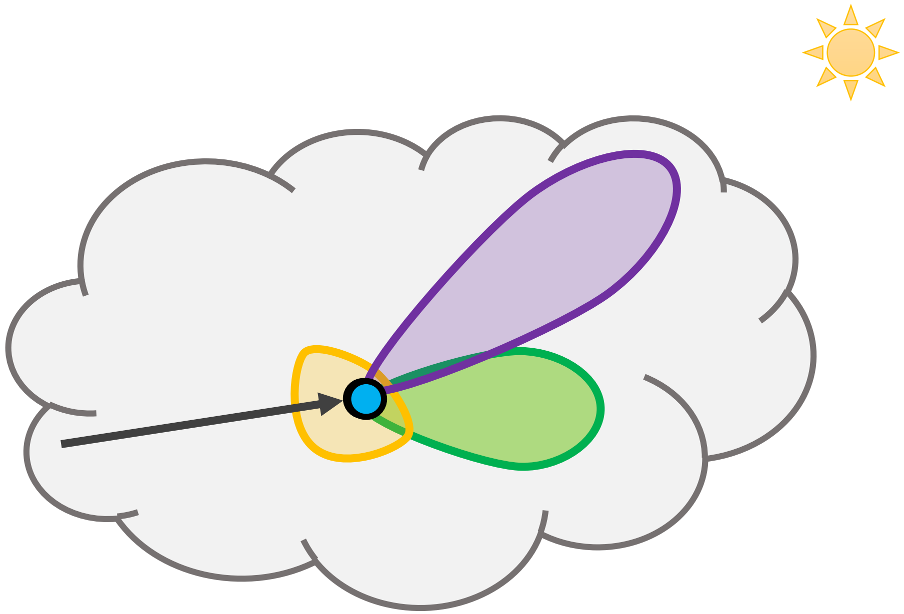

- Product between two VMM is a VMM:

$$V_f(\omega)V_L(\omega) = \left(V_f \otimes V_L\right)(\omega) = V_\otimes(\omega)$$

GUIDED PRODUCT DIRECTIONAL SAMPLING

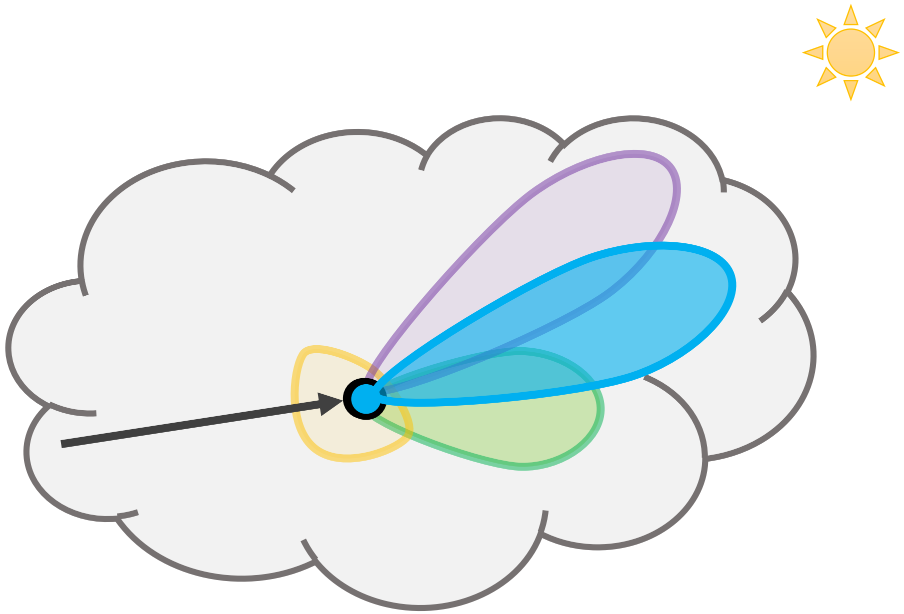

- Product between two VMM is a VMM:

$$V_f(\omega)V_L(\omega) = \left(V_f \otimes V_L\right)(\omega) = V_\otimes(\omega)$$

GUIDED PRODUCT DIRECTIONAL SAMPLING

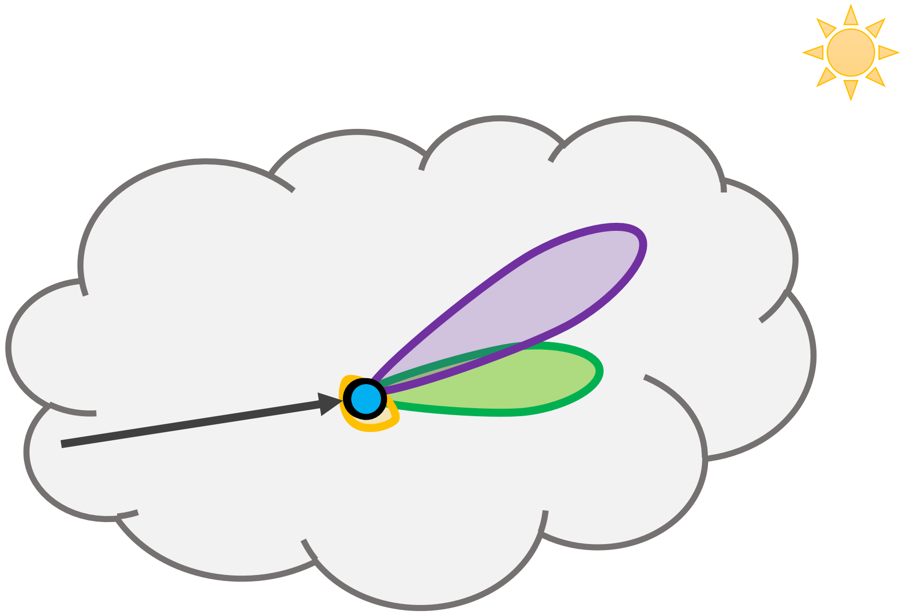

- Product between two VMM is a VMM:

$$V_f(\omega)V_L(\omega) = \left(V_f \otimes V_L\right)(\omega) = V_\otimes(\omega)$$

30 min

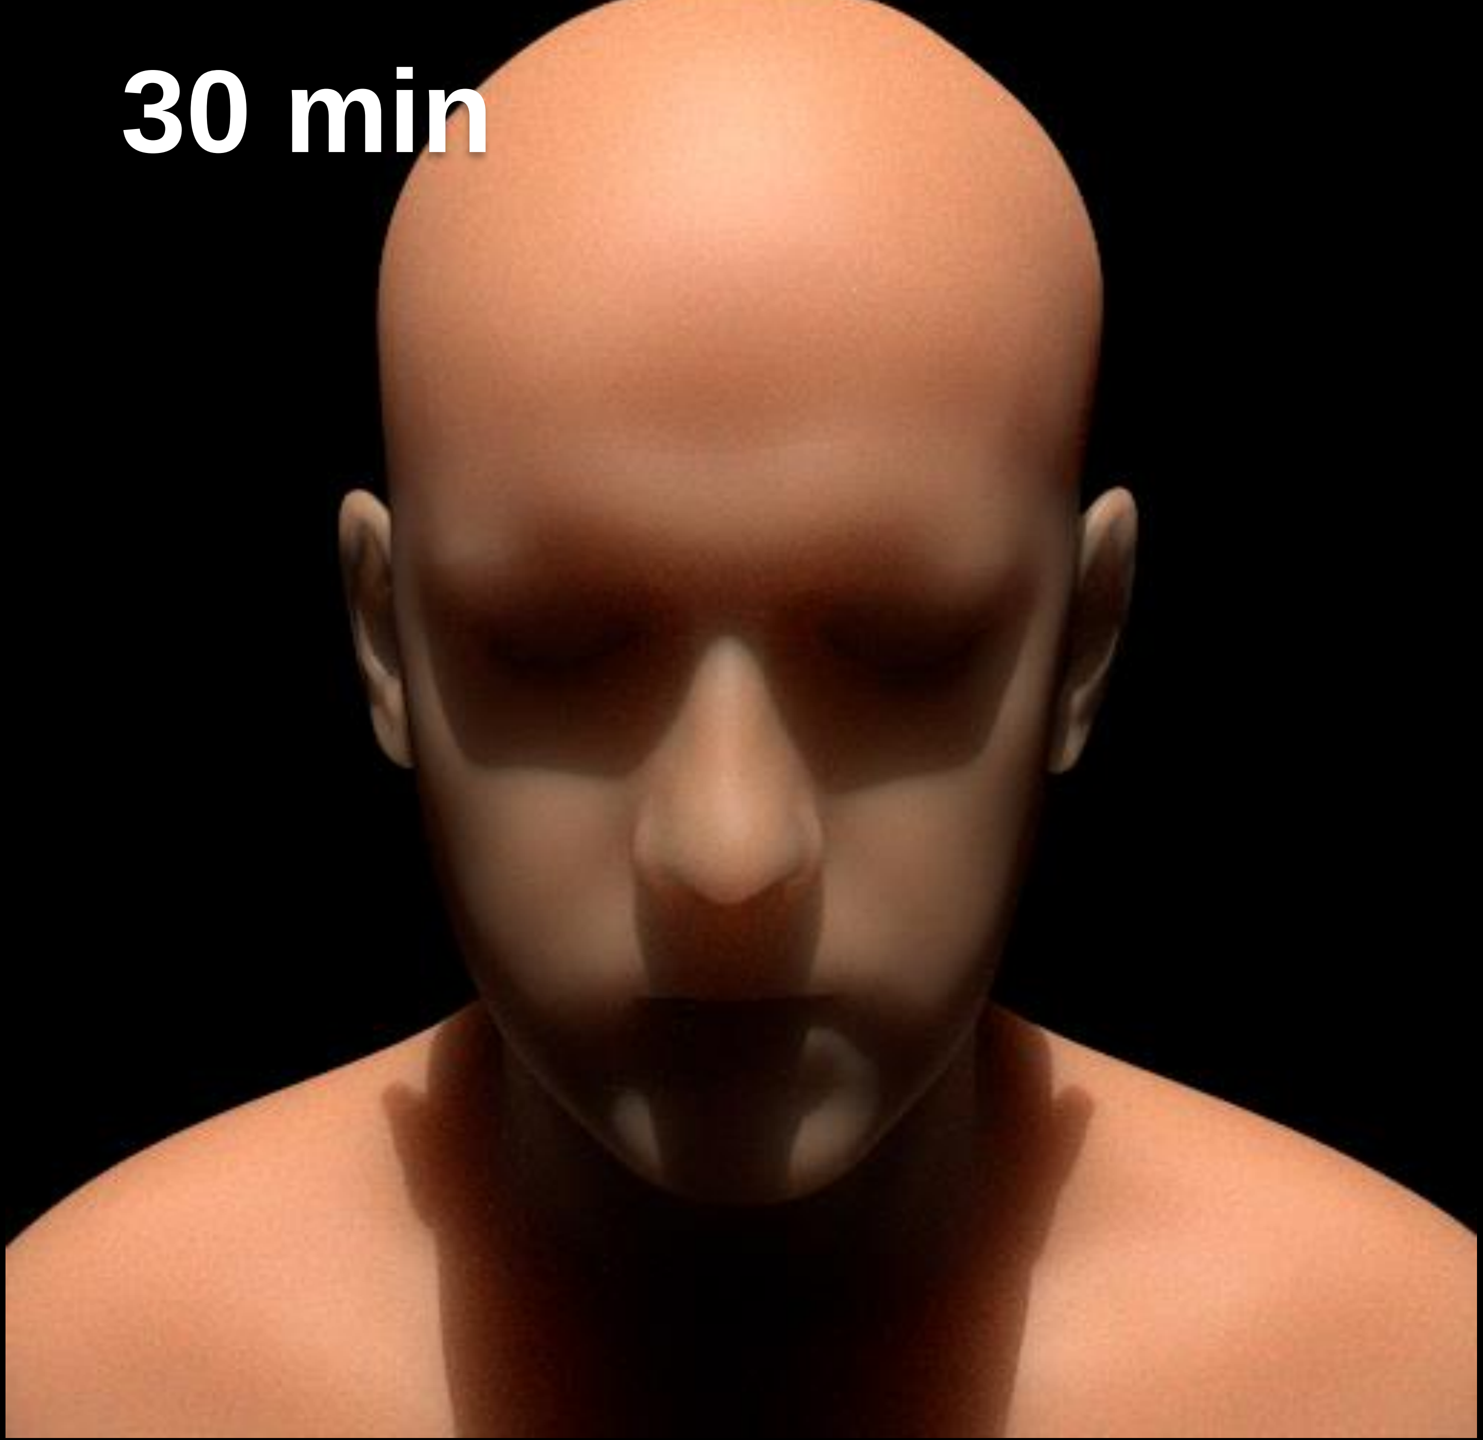

30 min

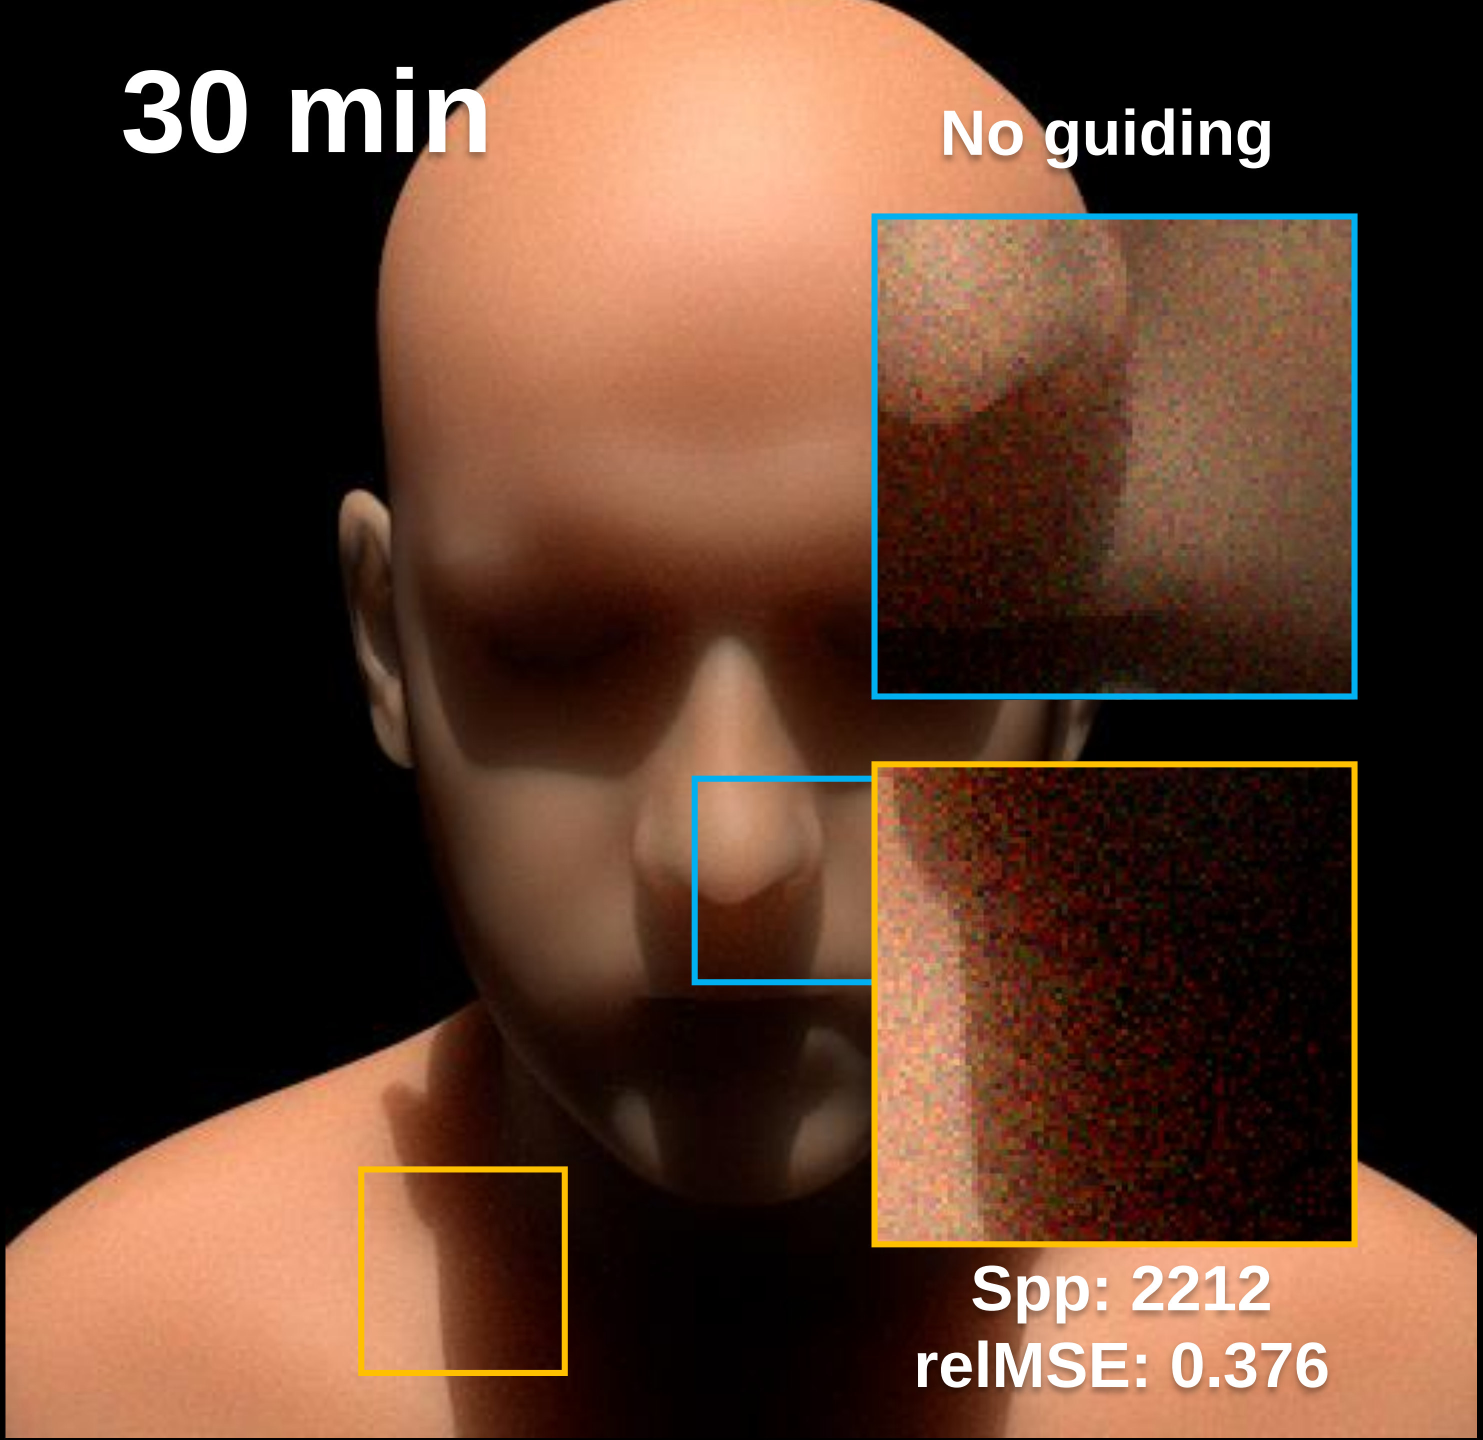

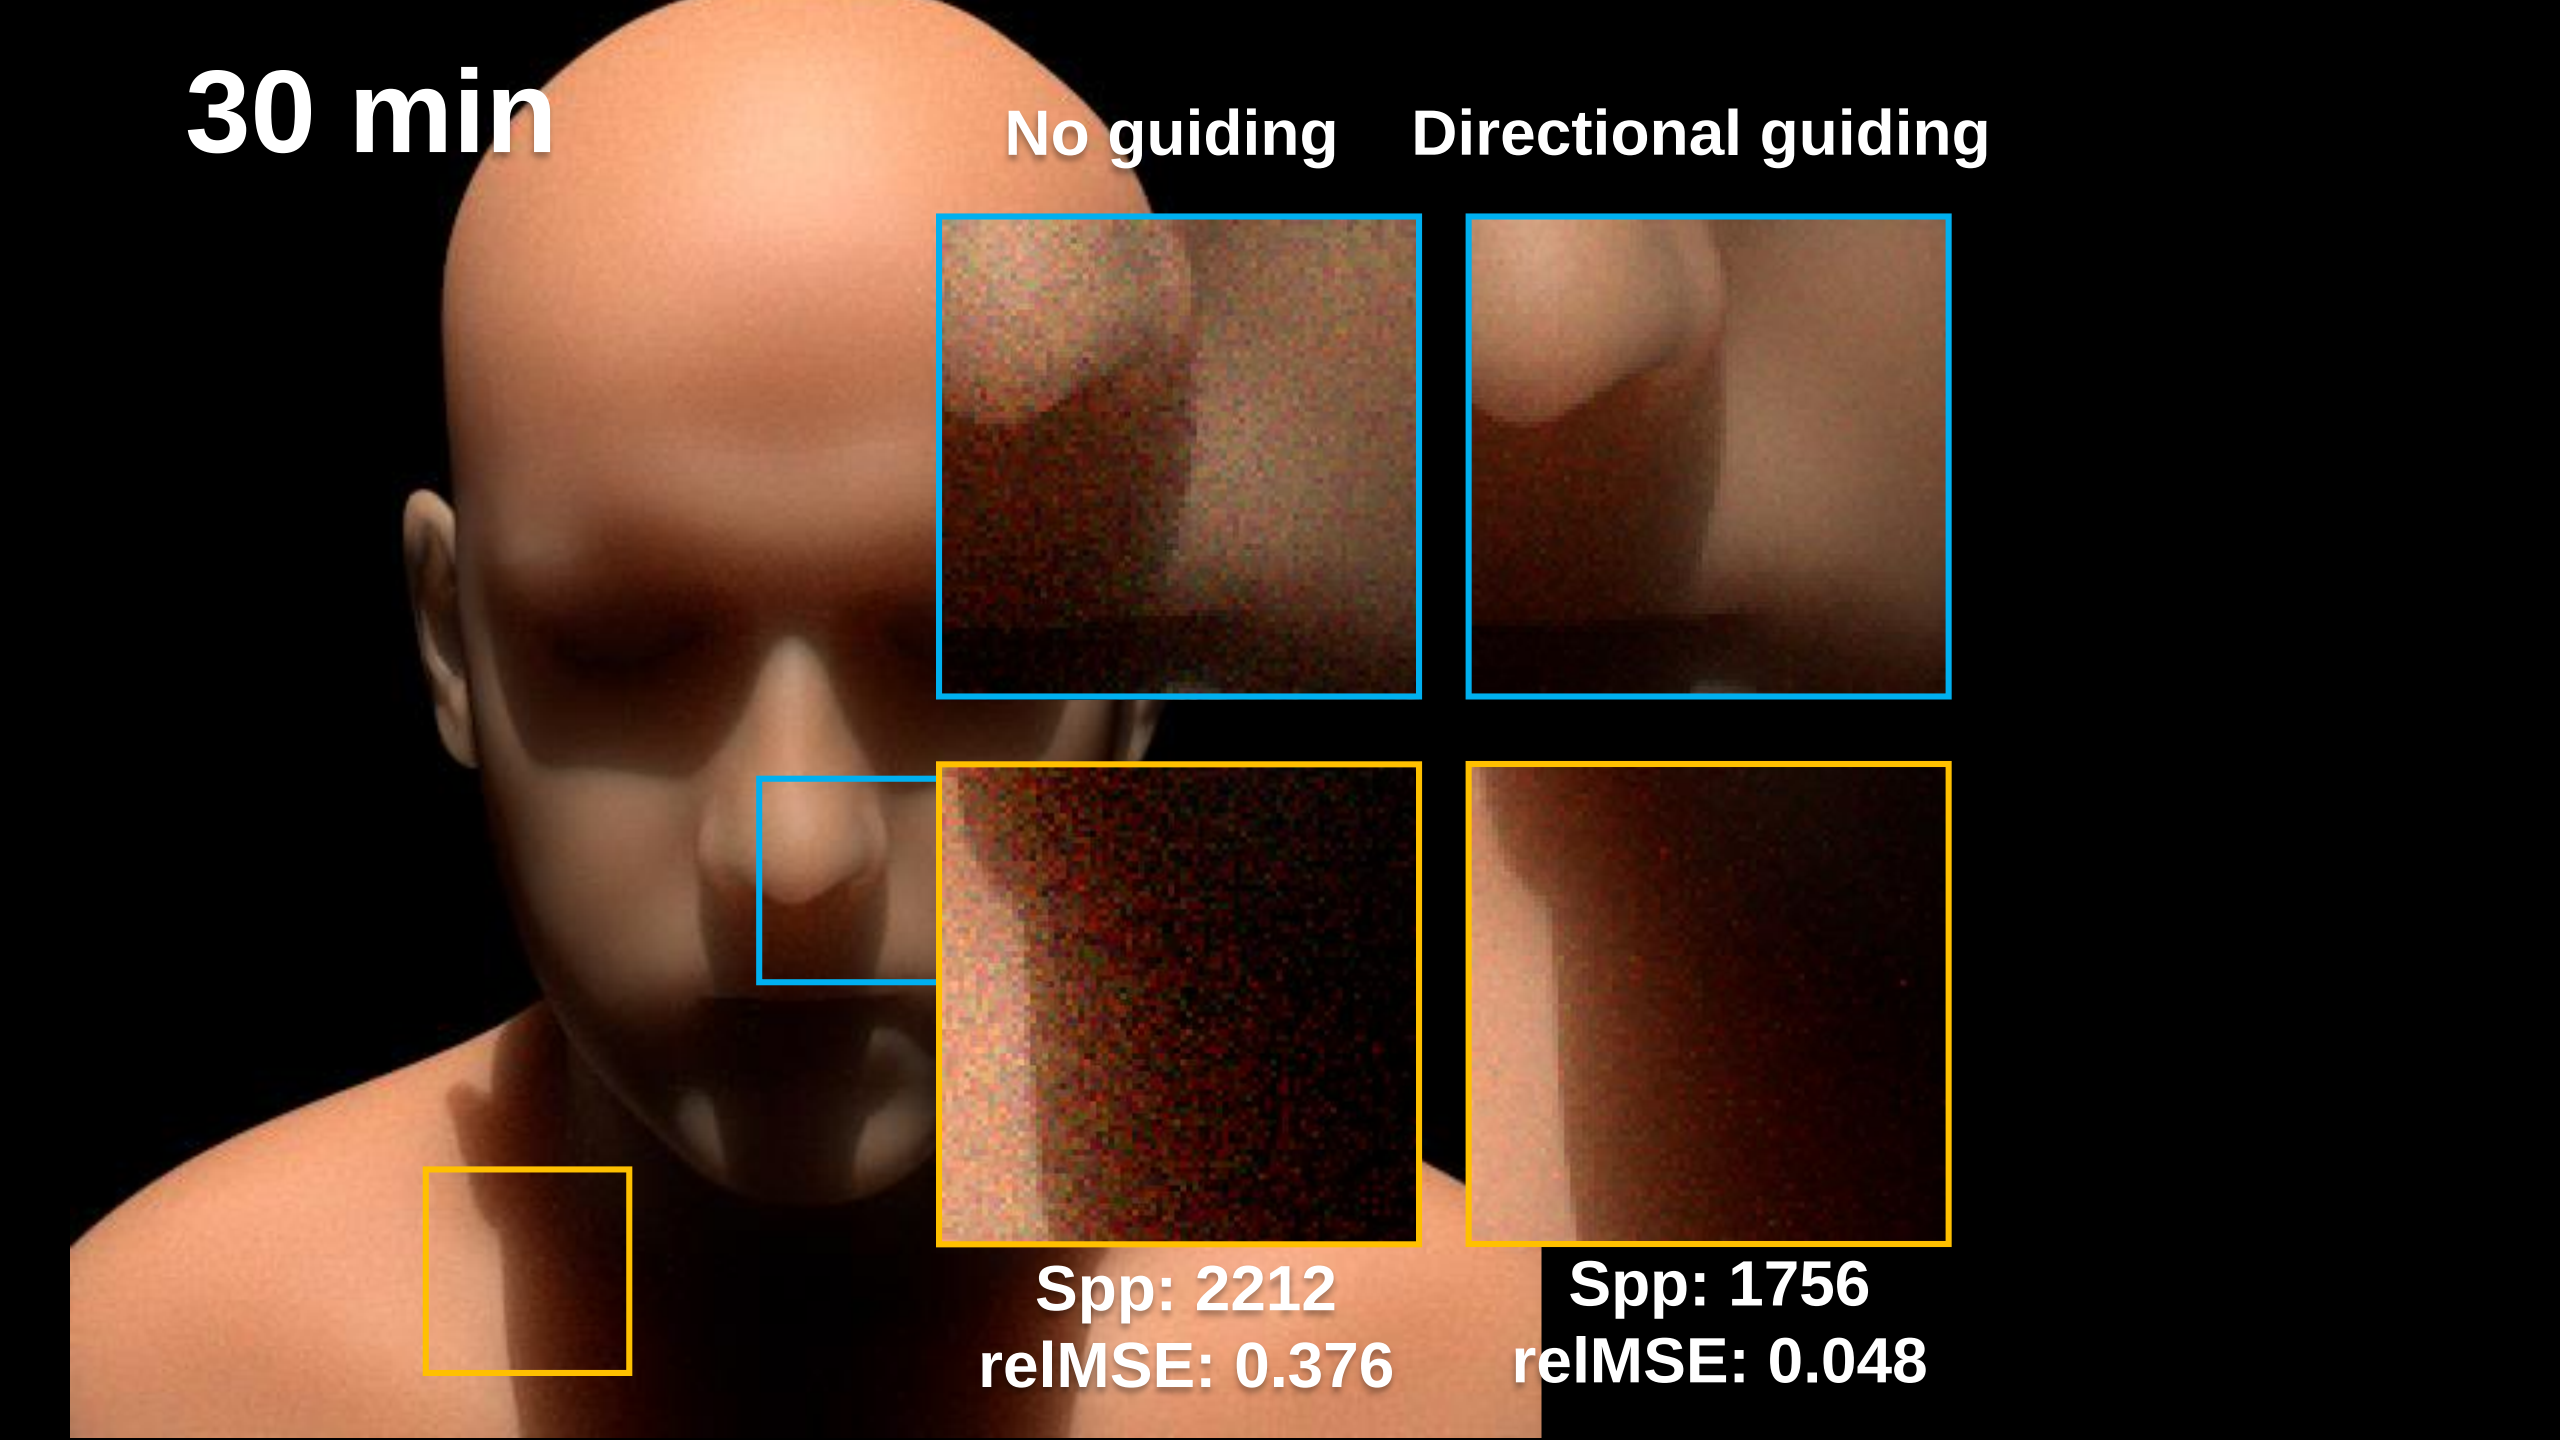
30 min
No guiding
Directional guiding
Spp: 2212
relMSE: 0.376
Spp: 1756
relMSE: 0.048

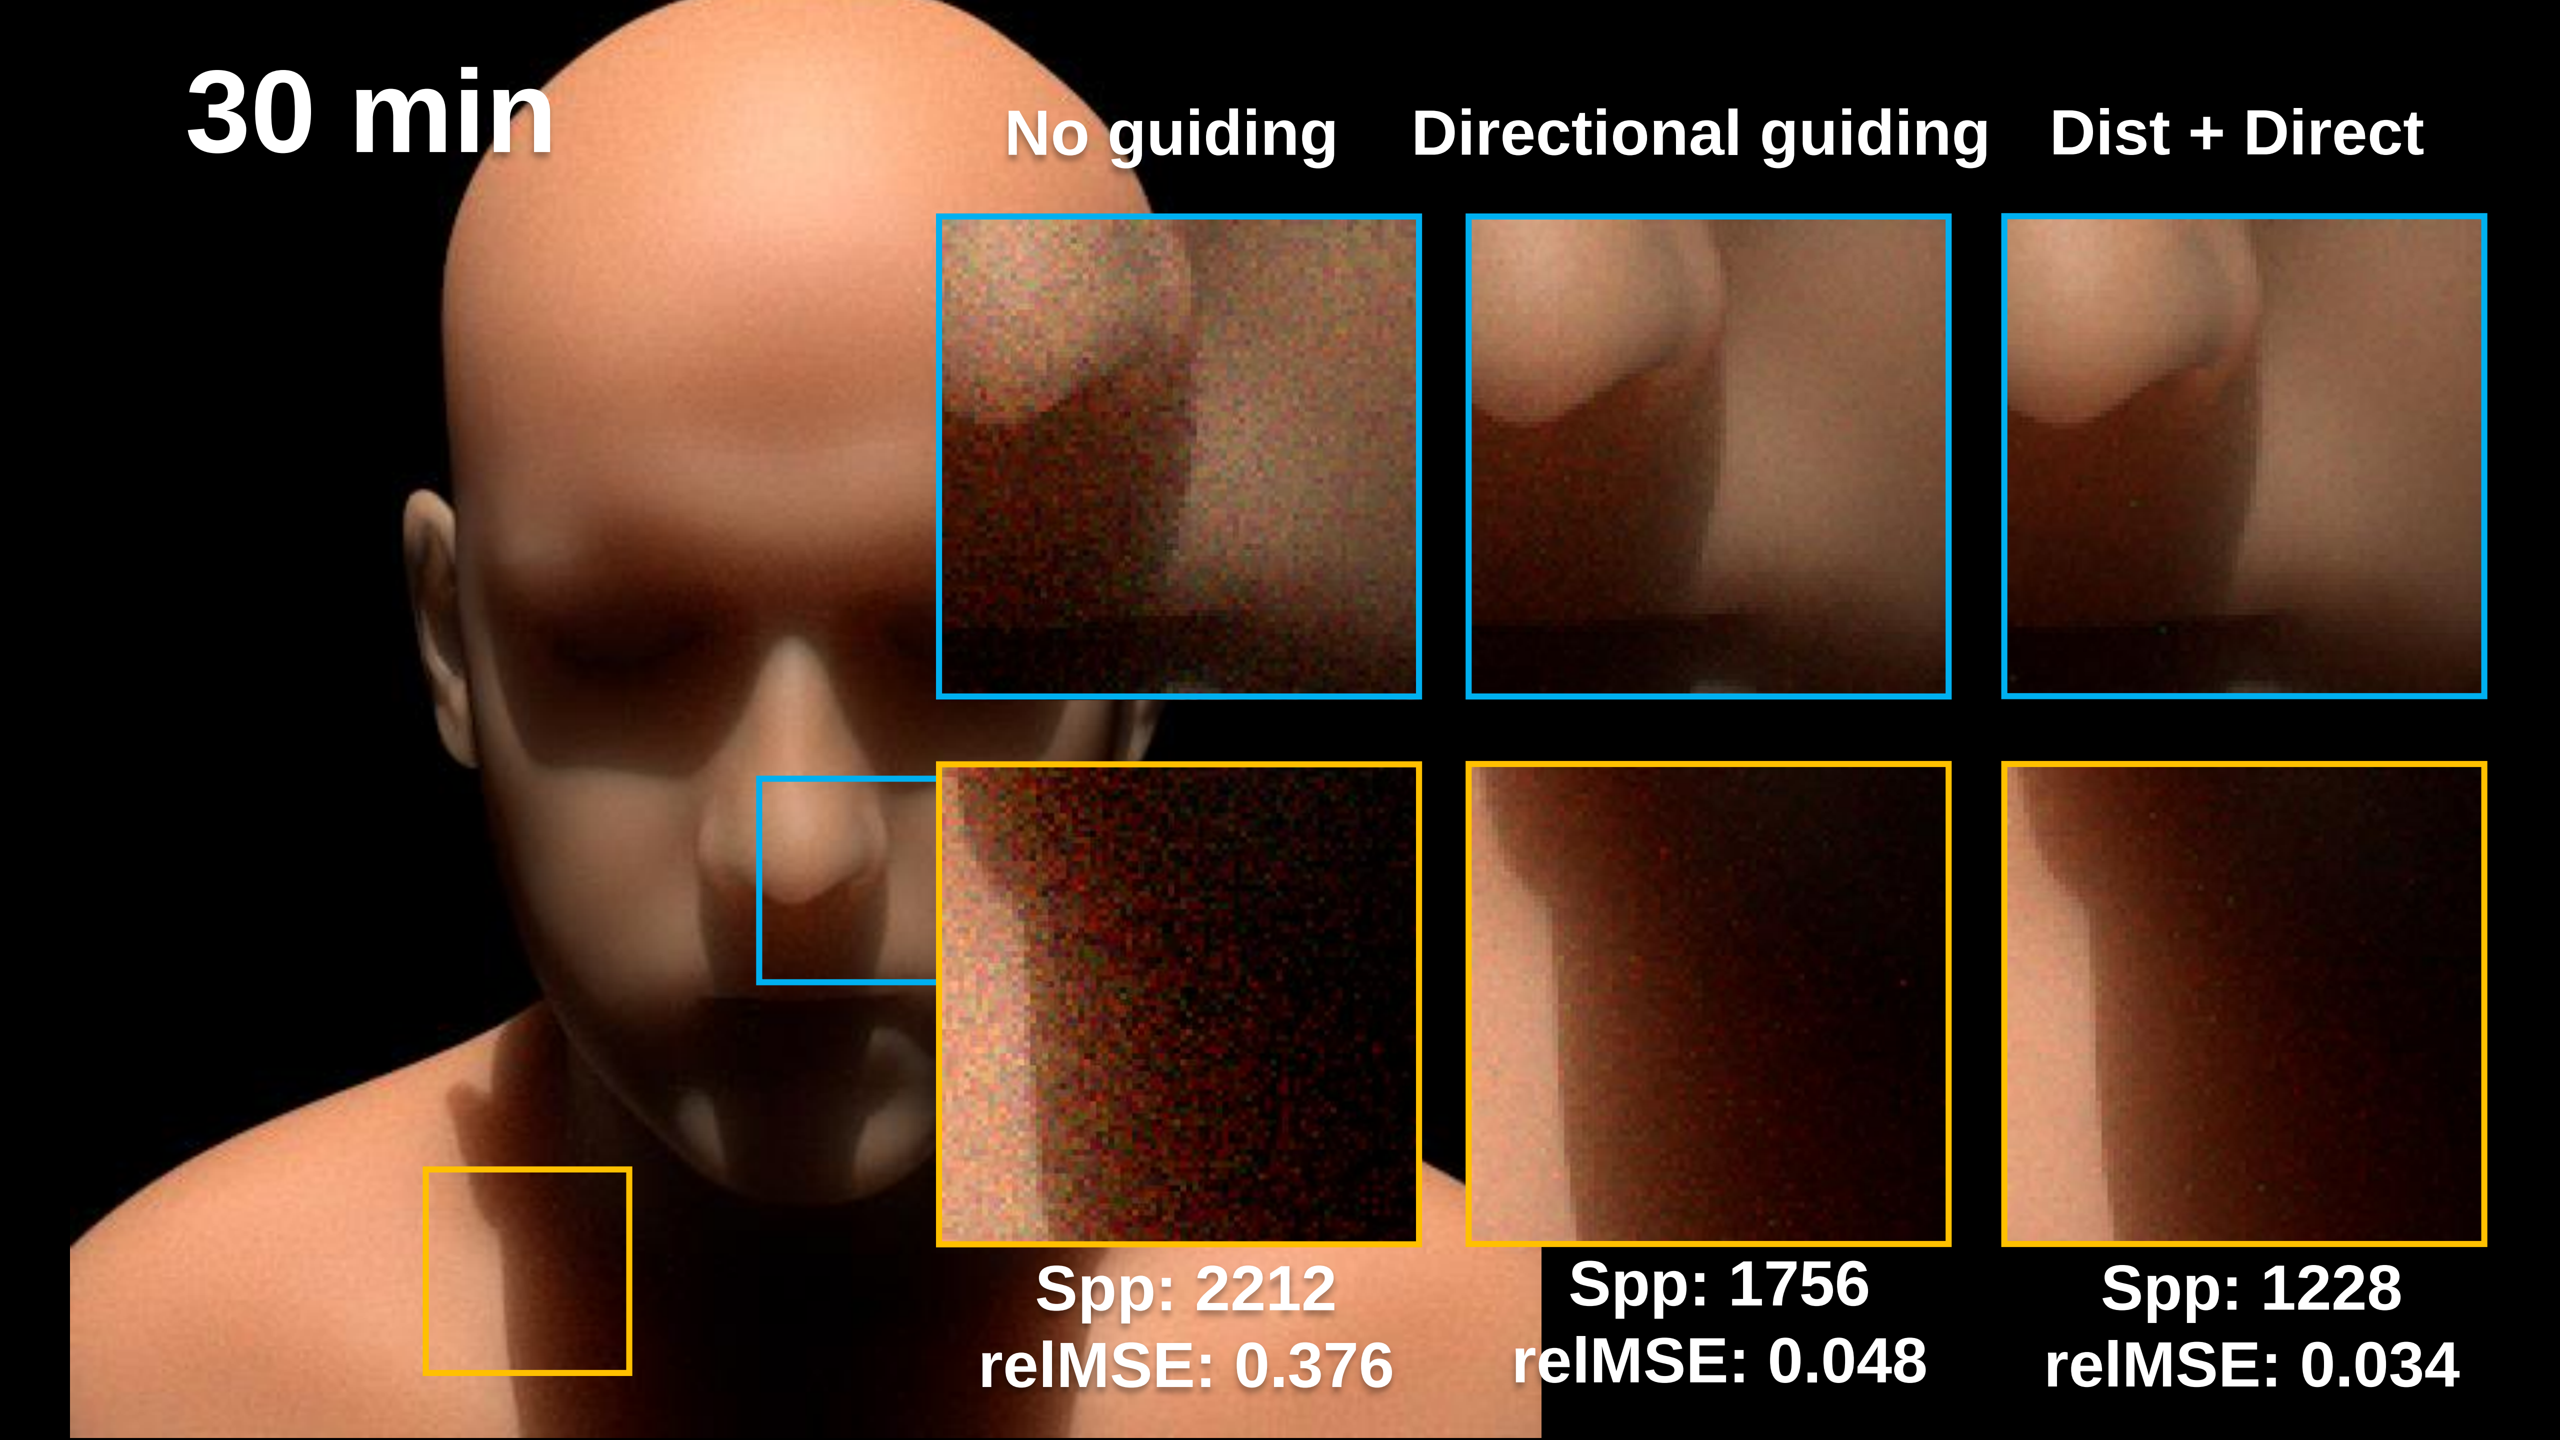
30 min
No guiding
Directional guiding
Dist + Direct
Spp: 2212
relMSE: 0.376
Spp: 1756
relMSE: 0.048
Spp: 1228
relMSE: 0.034

IMPORTANCE OF THE PRODUCT FOR DIRECTIONAL GUIDING

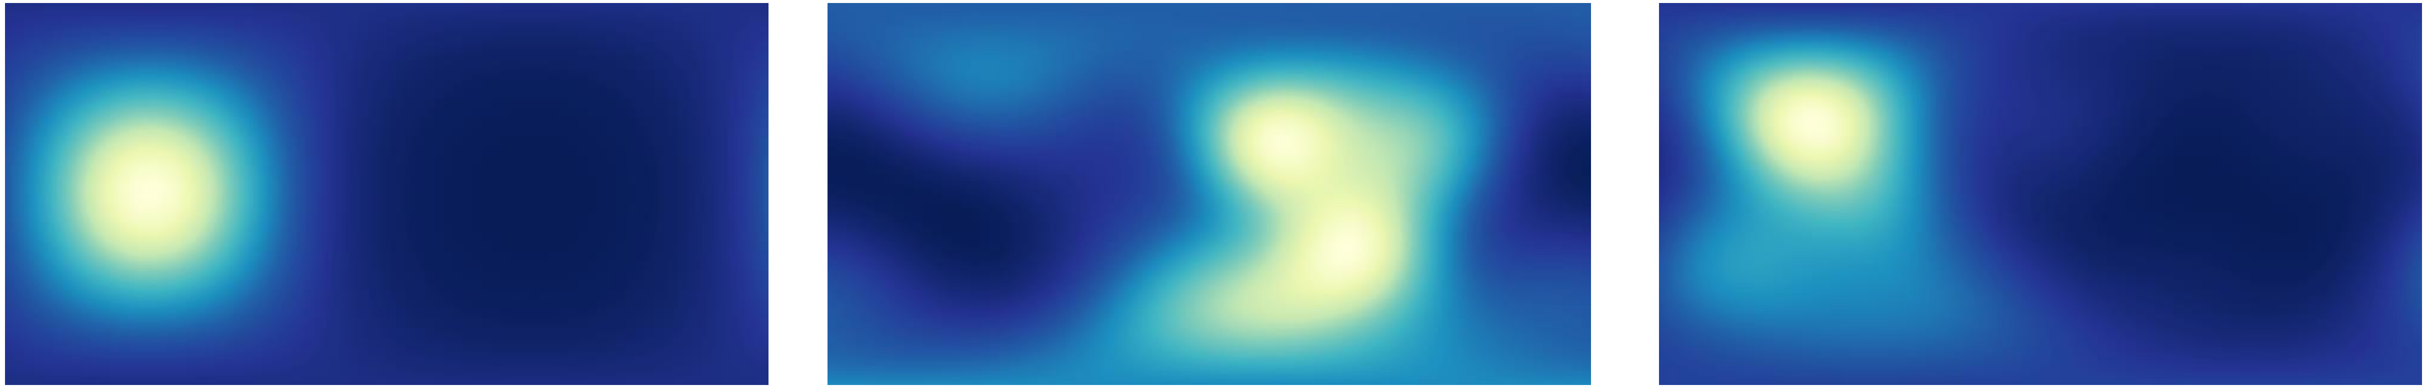

- Phase function PDF:

$$p_\omega^f(\ldots) \propto f(\ldots)$$

- Incident Radiance PDF:

$$p_\omega^L(\ldots) \propto L(x_{j+1}, \omega_{j+1})$$

- Mixture PDF:

$$p_\omega^{guide}(\ldots) = \alpha \cdot p_\omega^f(\ldots) + (1-\alpha) \cdot p_\omega^L(\ldots)$$

IMPORTANCE OF THE PRODUCT FOR DIRECTIONAL GUIDING

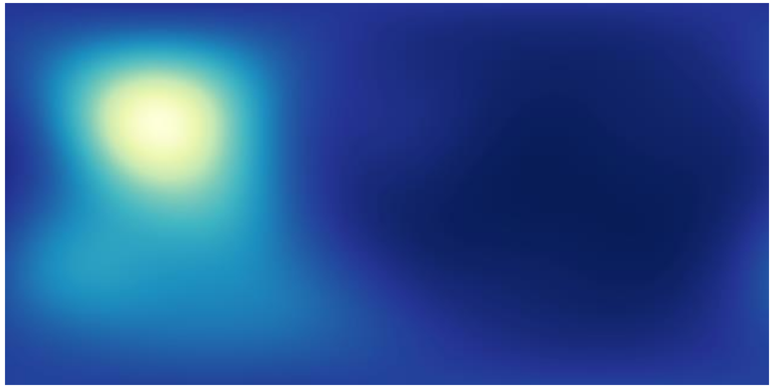

$\neq$

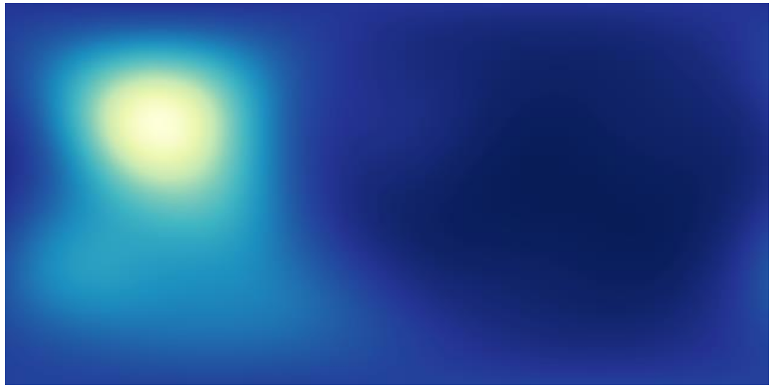

- Mixture PDF:

$$p_\omega^{mix}(\dots) = \alpha \cdot p_\omega^f(\dots) + (1-\alpha) \cdot p_\omega^L(\dots)$$

- Product PDF:

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

GUIDED RUSSIAN ROULETTE AND SPLITTING

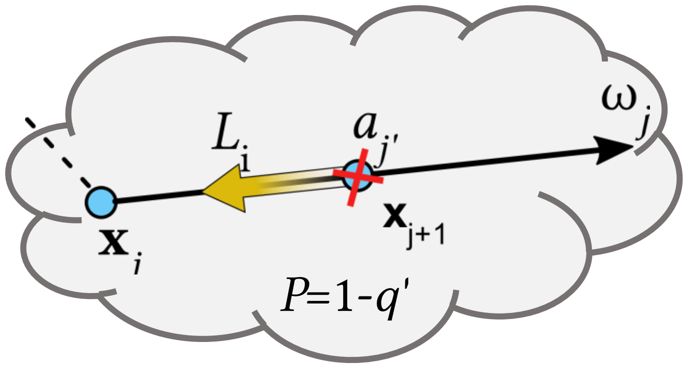

4a. Termination

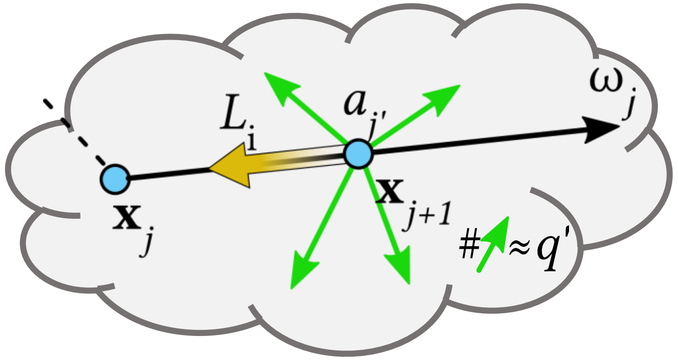

4b. Splitting

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

GUIDED RUSSIAN ROULETTE AND SPLITTING

Directional

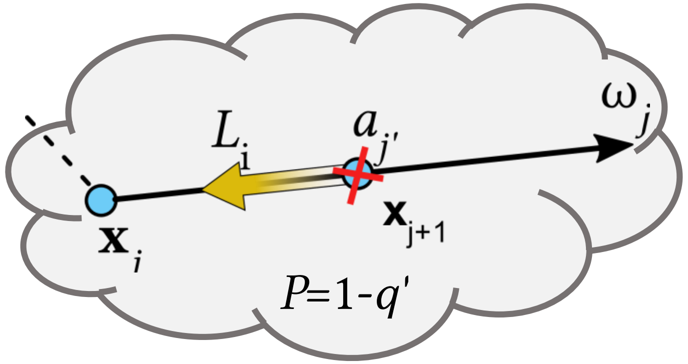

4a. Termination

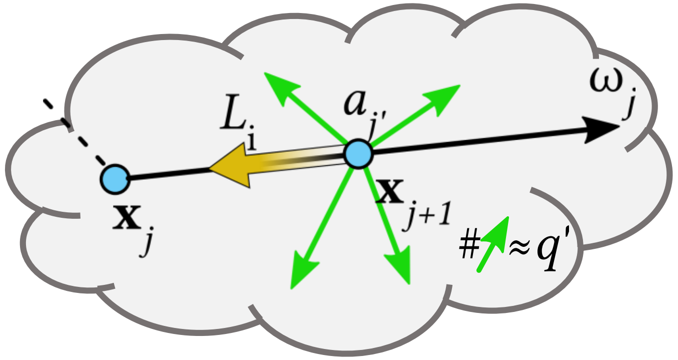

4b. Splitting

Distance

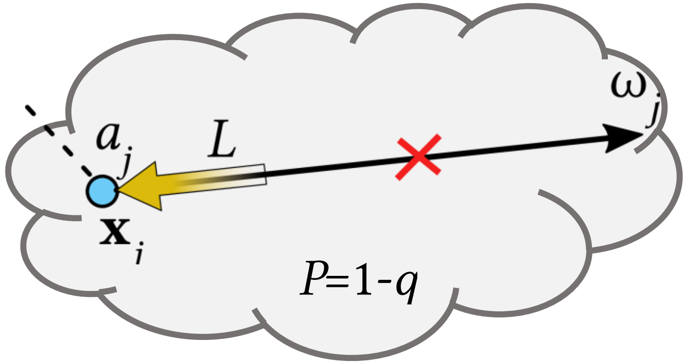

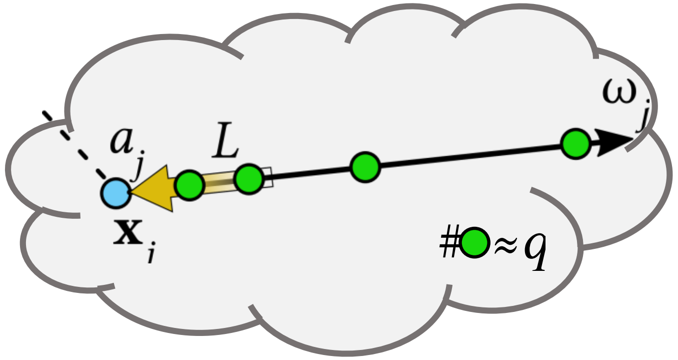

GUIDED RUSSIAN ROULETTE AND SPLITTING

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

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

GUIDED RUSSIAN ROULETTE AND SPLITTING

Path contribution

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

Reference solution

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

VOLUME RENDERING EQUATION

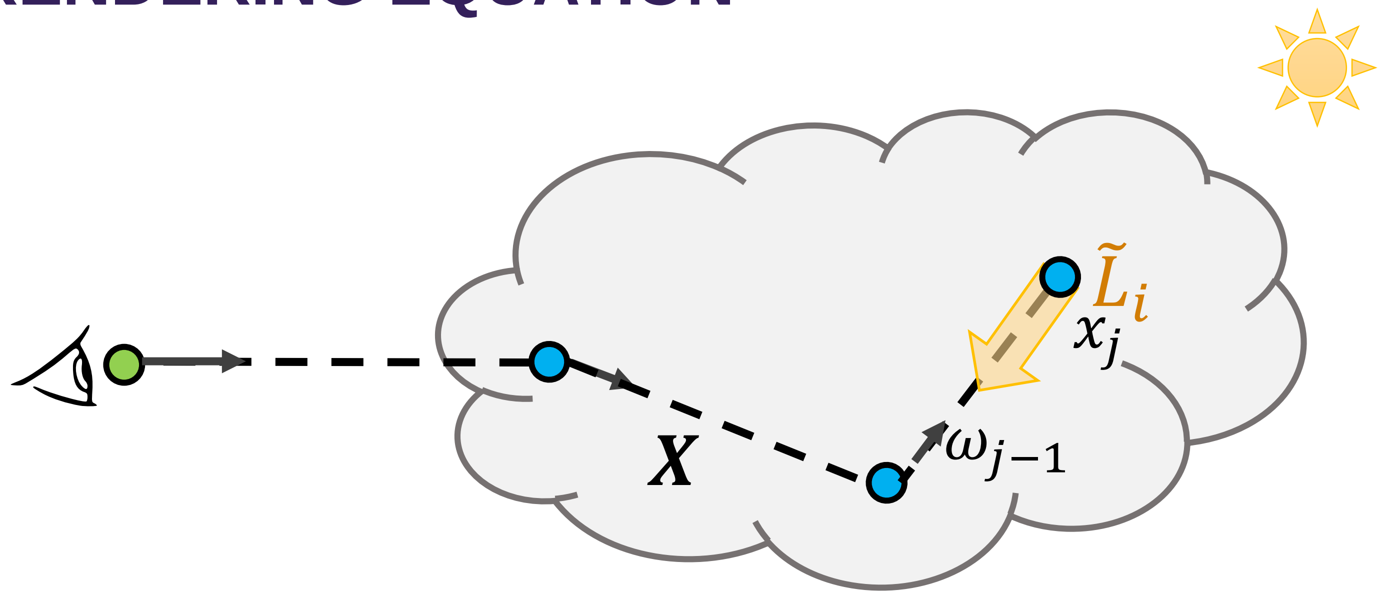

Path throughput: $f(\boldsymbol{X})/p(\boldsymbol{X})$

In-scattered radiance estimate

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

- See course notes or paper for more details

GUIDED RUSSIAN ROULETTE AND SPLITTING: PIXEL ESTIMATE

GUIDED RUSSIAN ROULETTE AND SPLITTING: PIXEL ESTIMATE

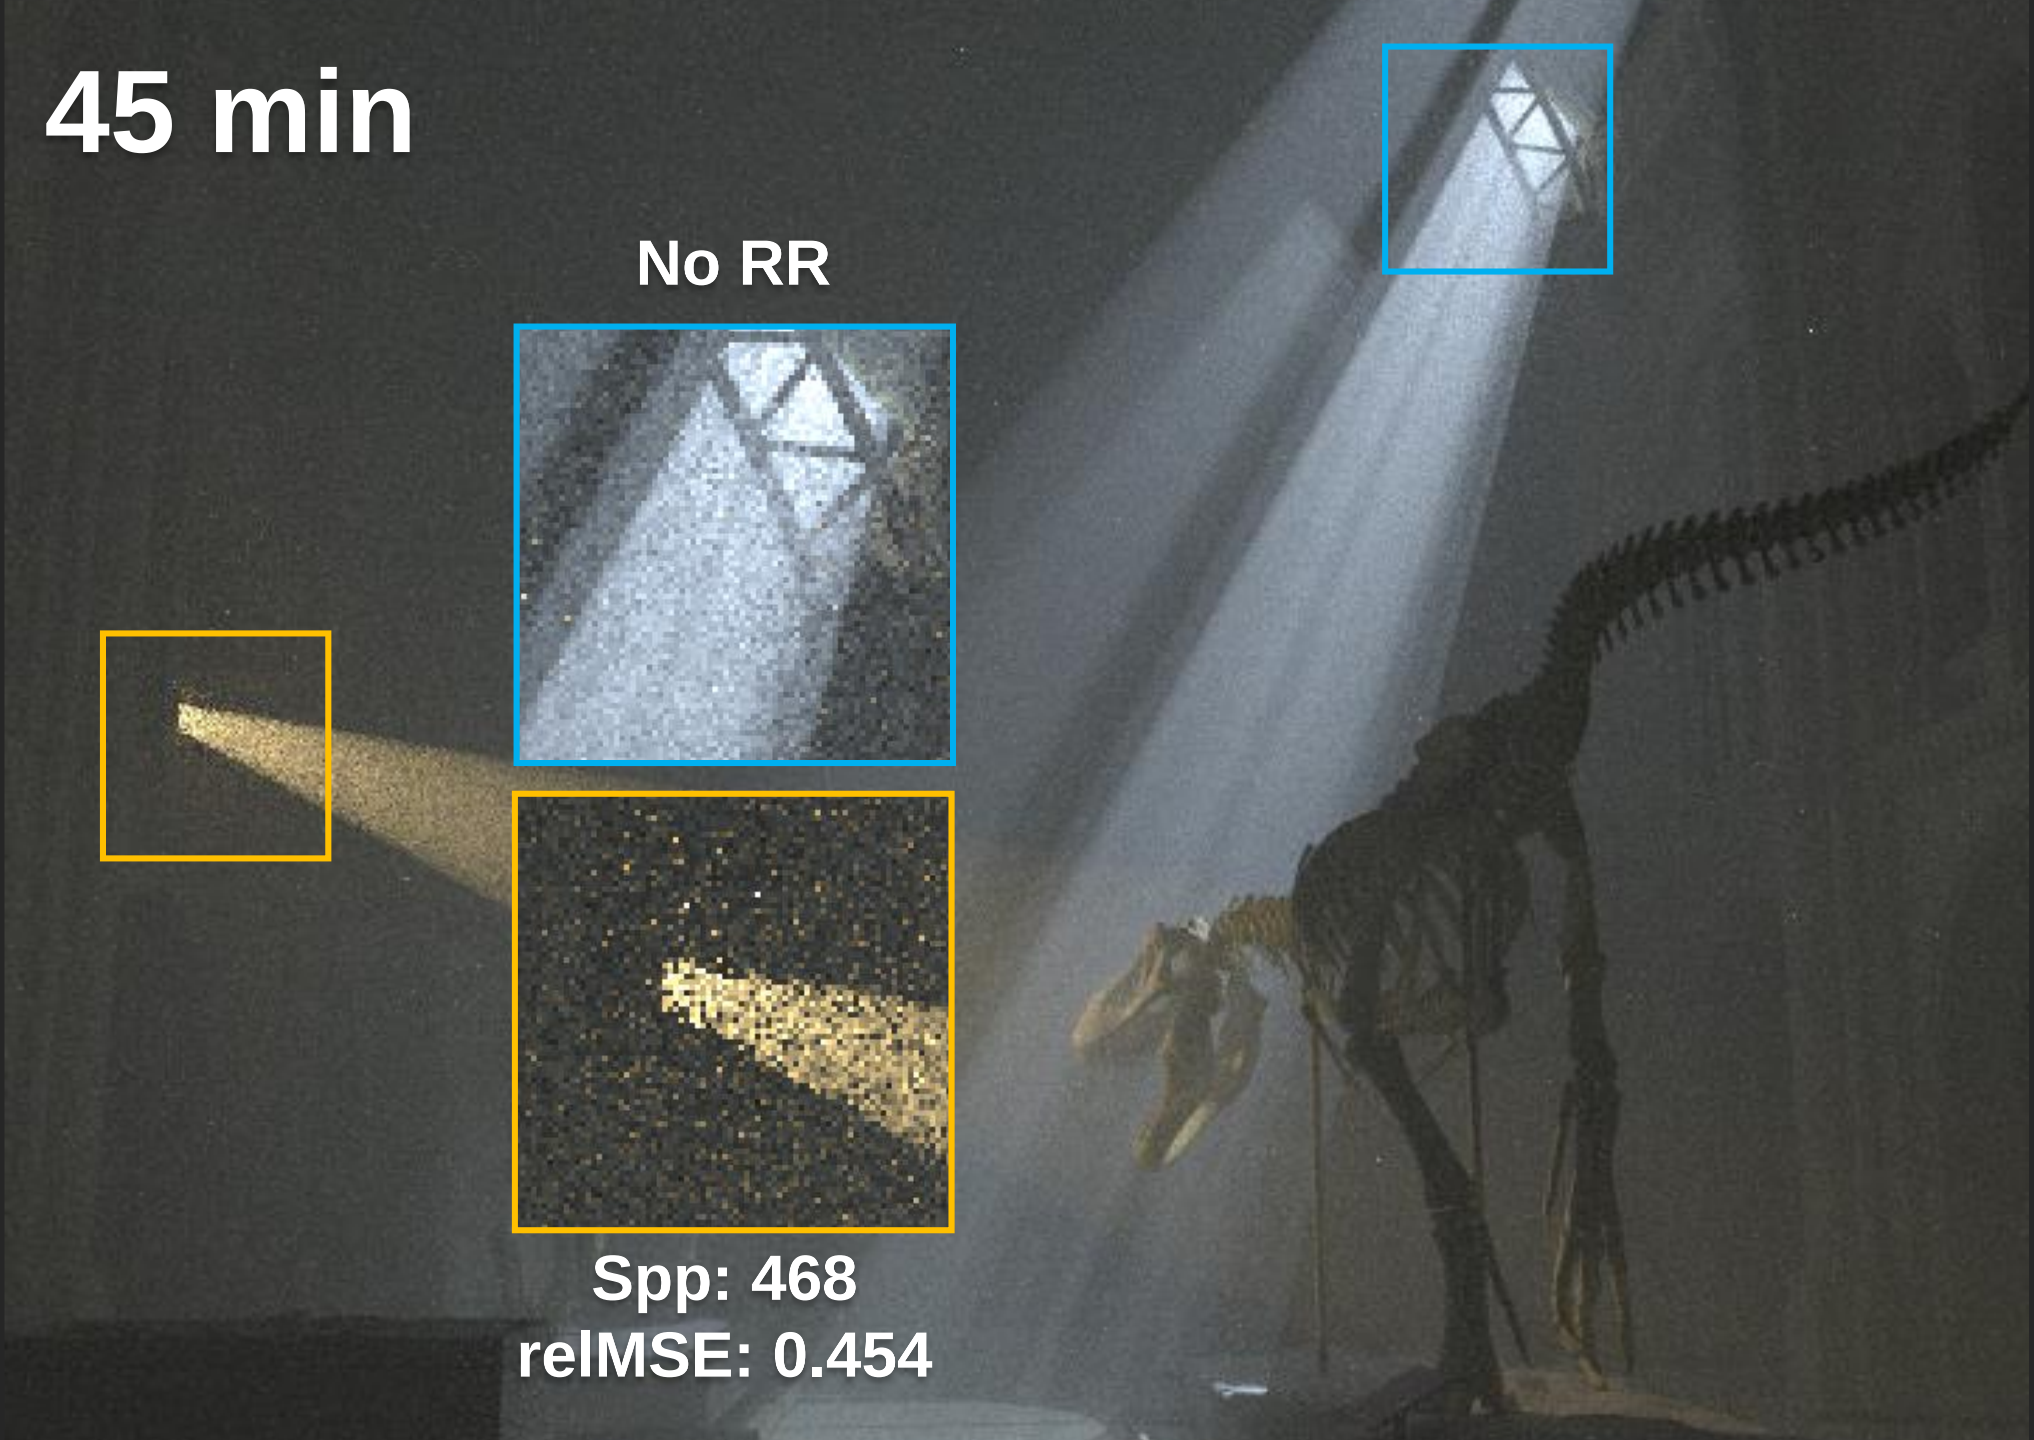
45 min
No RR
Spp: 468
relMSE: 0.454

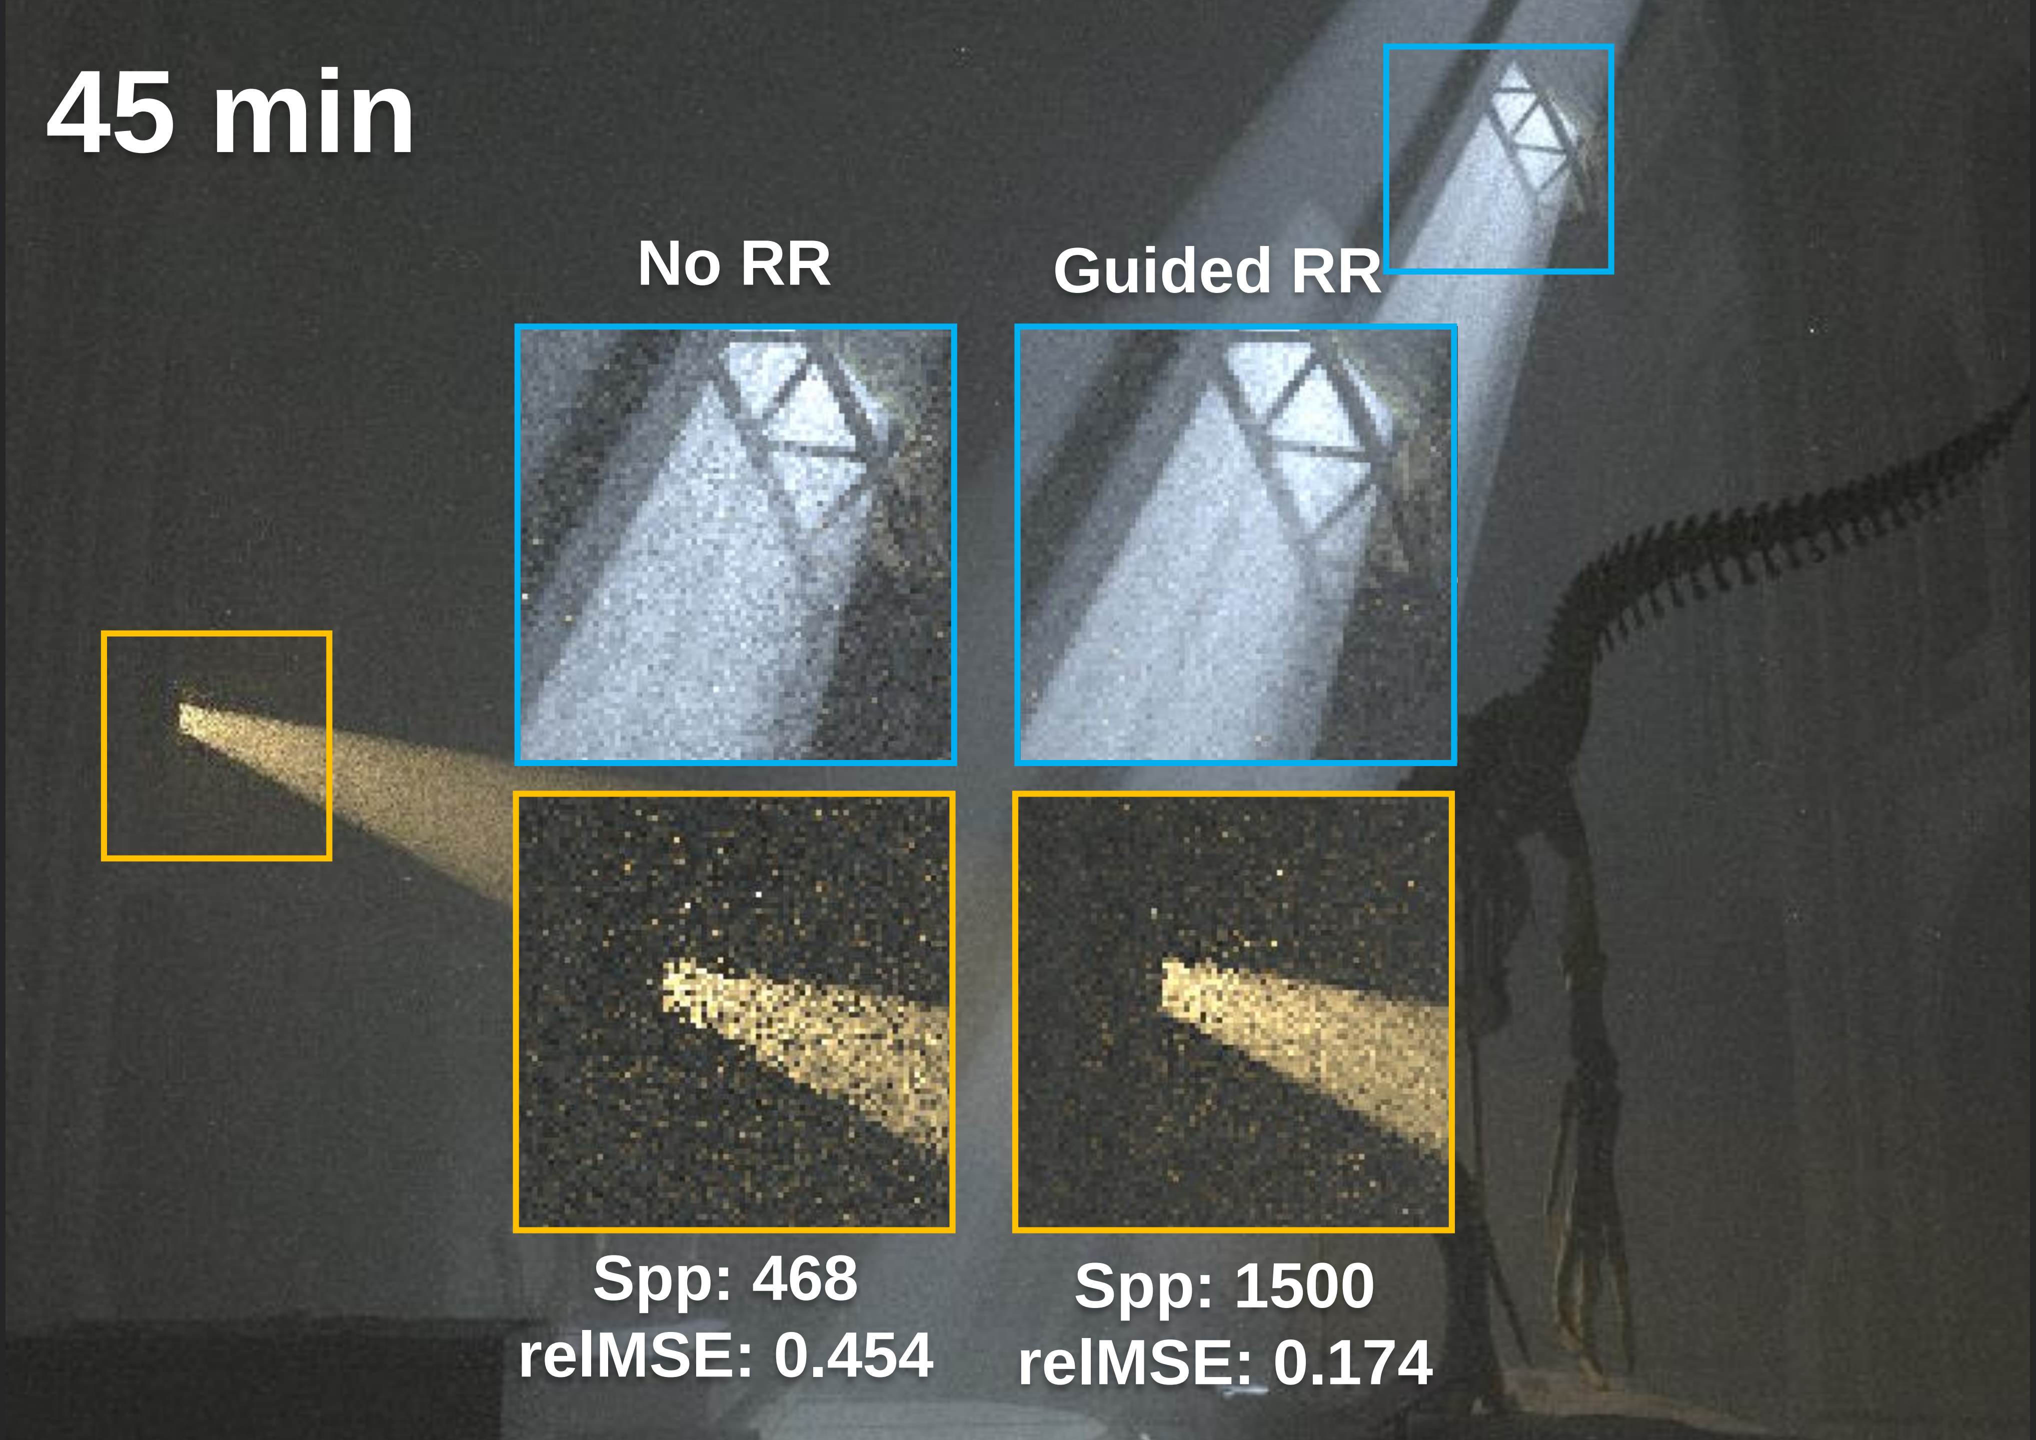
45 min
No RR
Guided RR
Spp: 468
relMSE: 0.454
Spp: 1500
relMSE: 0.174

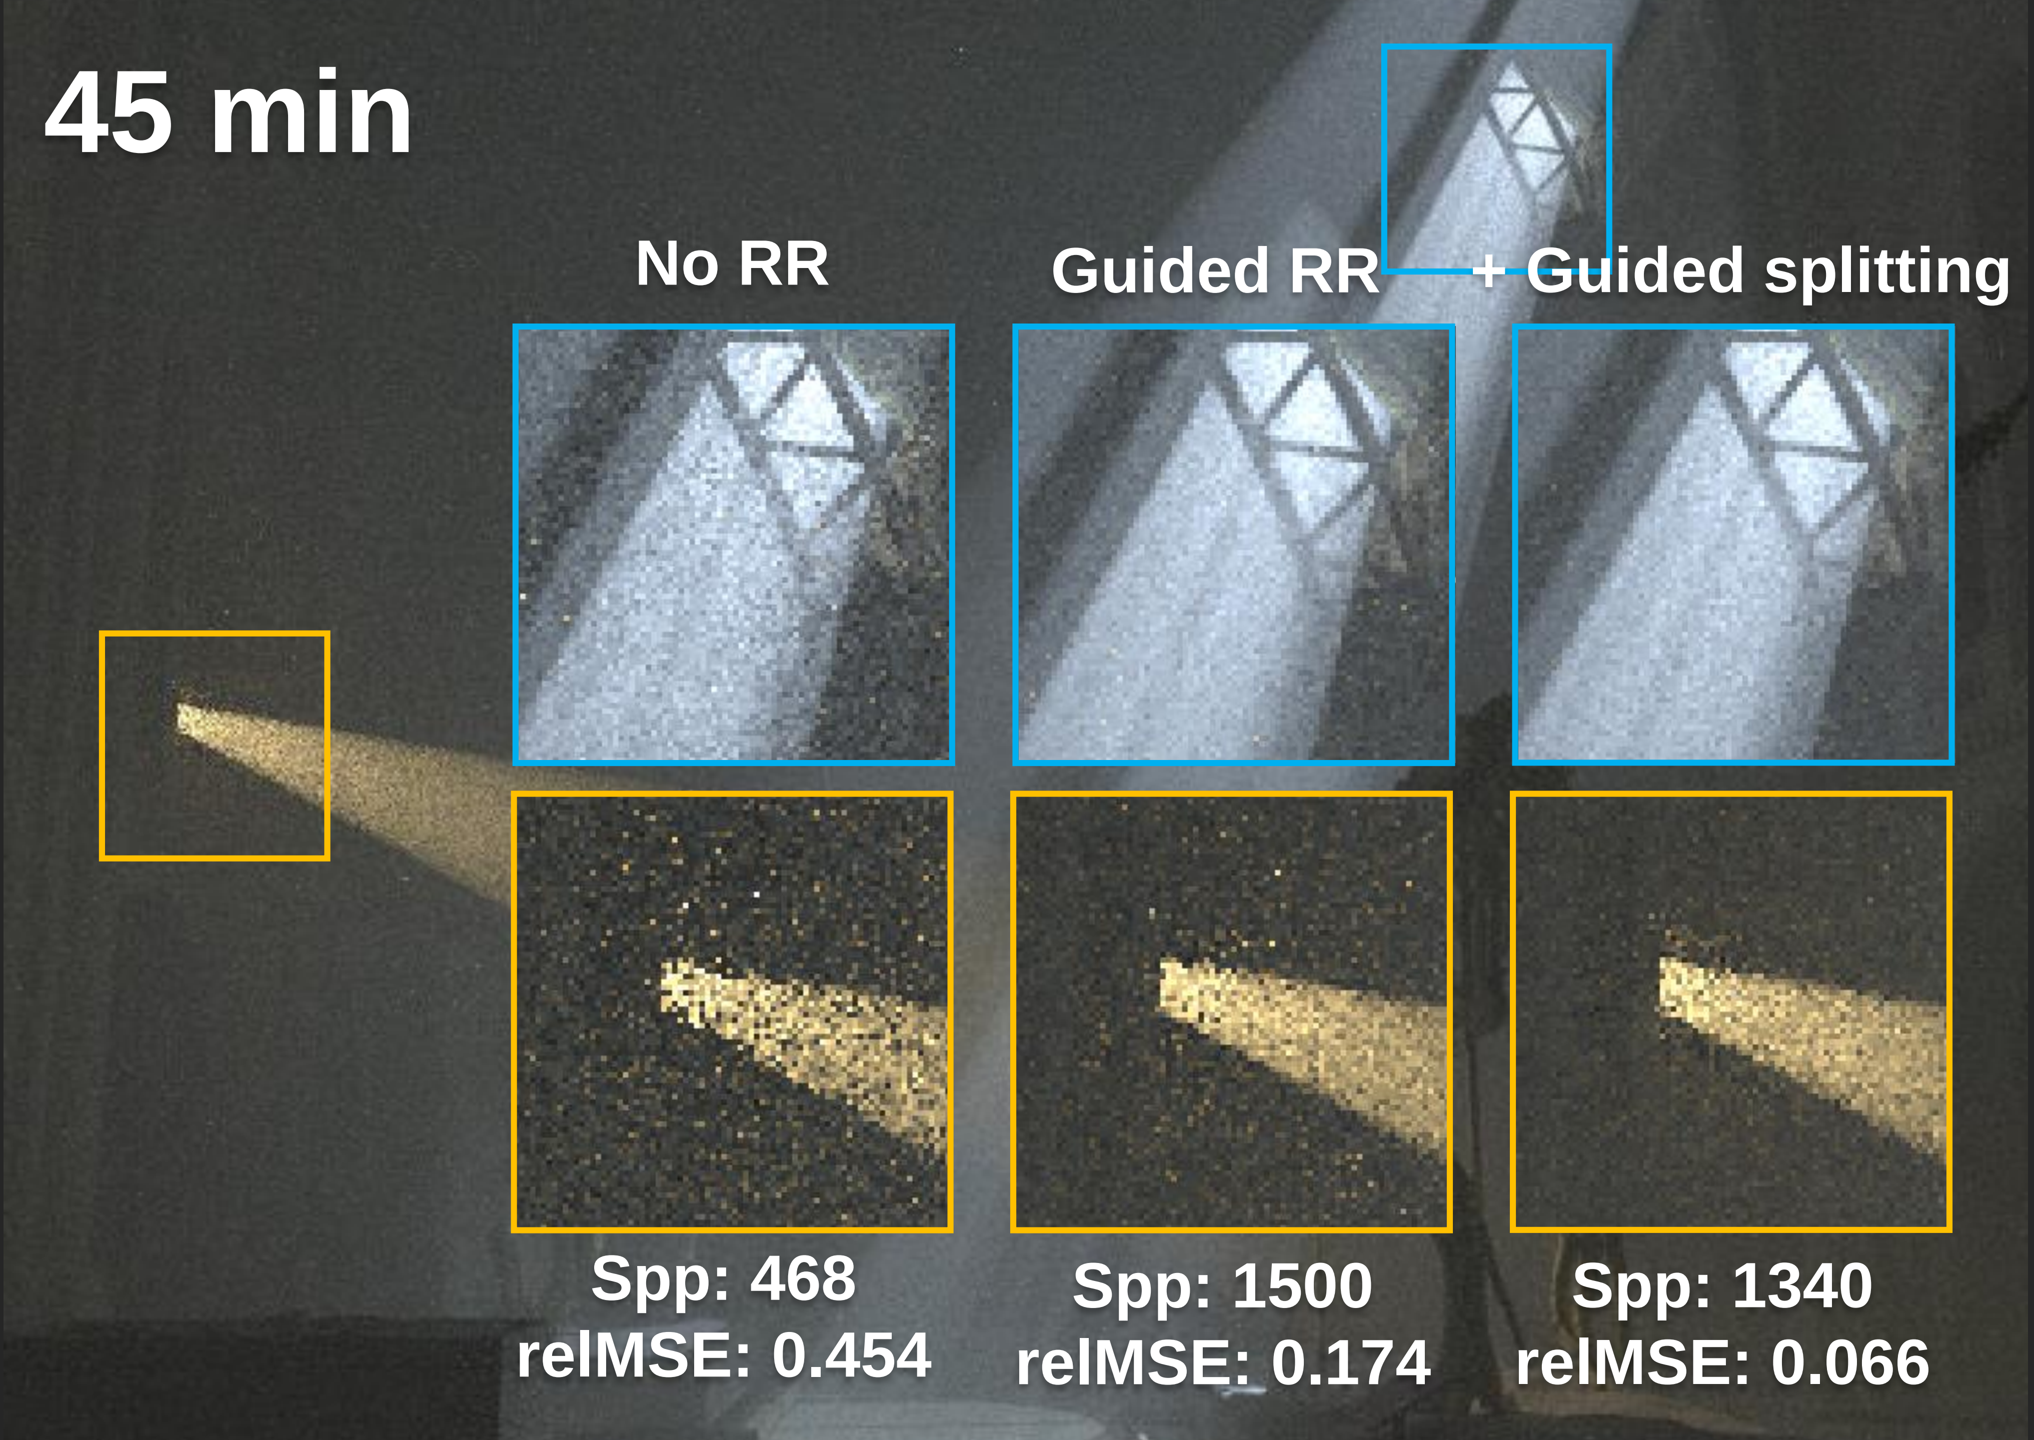
45 min
No RR
Guided RR
+ Guided splitting
Spp: 468
relMSE: 0.454
Spp: 1500
relMSE: 0.174
Spp: 1340
relMSE: 0.066

Guided RR

Spp: 1500
relMSE: 0.174

+ Guided splitting

Spp: 1340
relMSE: 0.066

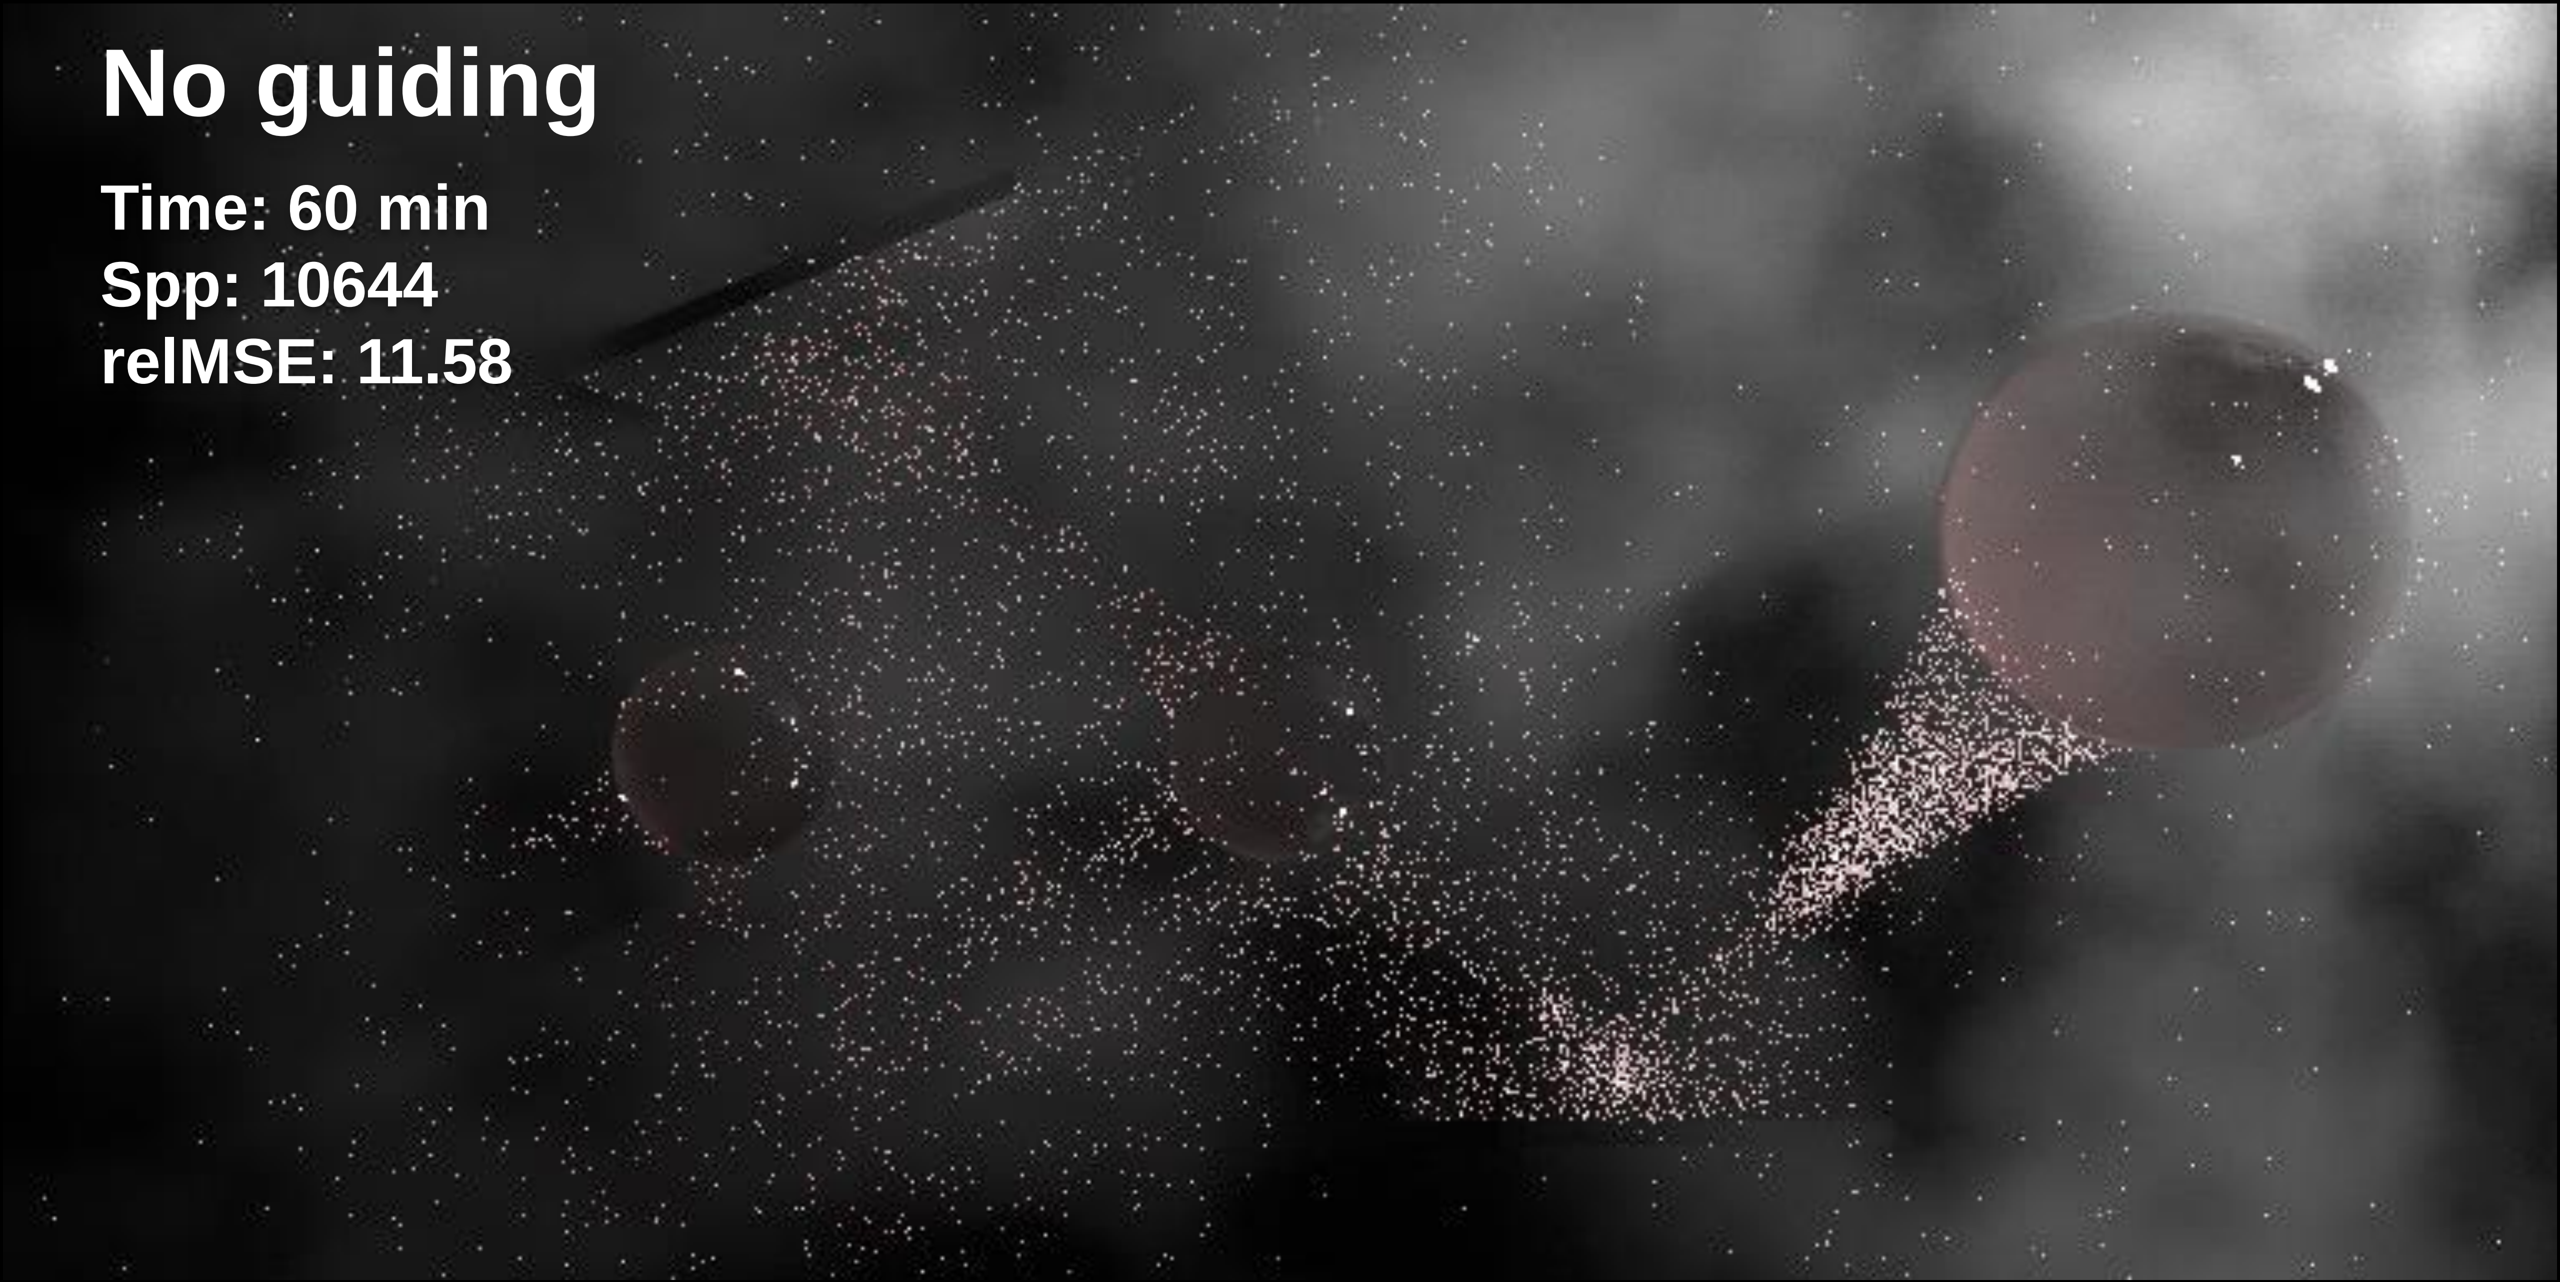
No guiding
Time: 60 min
Spp: 10644
relMSE: 11.58

Distance guiding

Time: 60 min
Spp: 4624
relMSE: 3.520

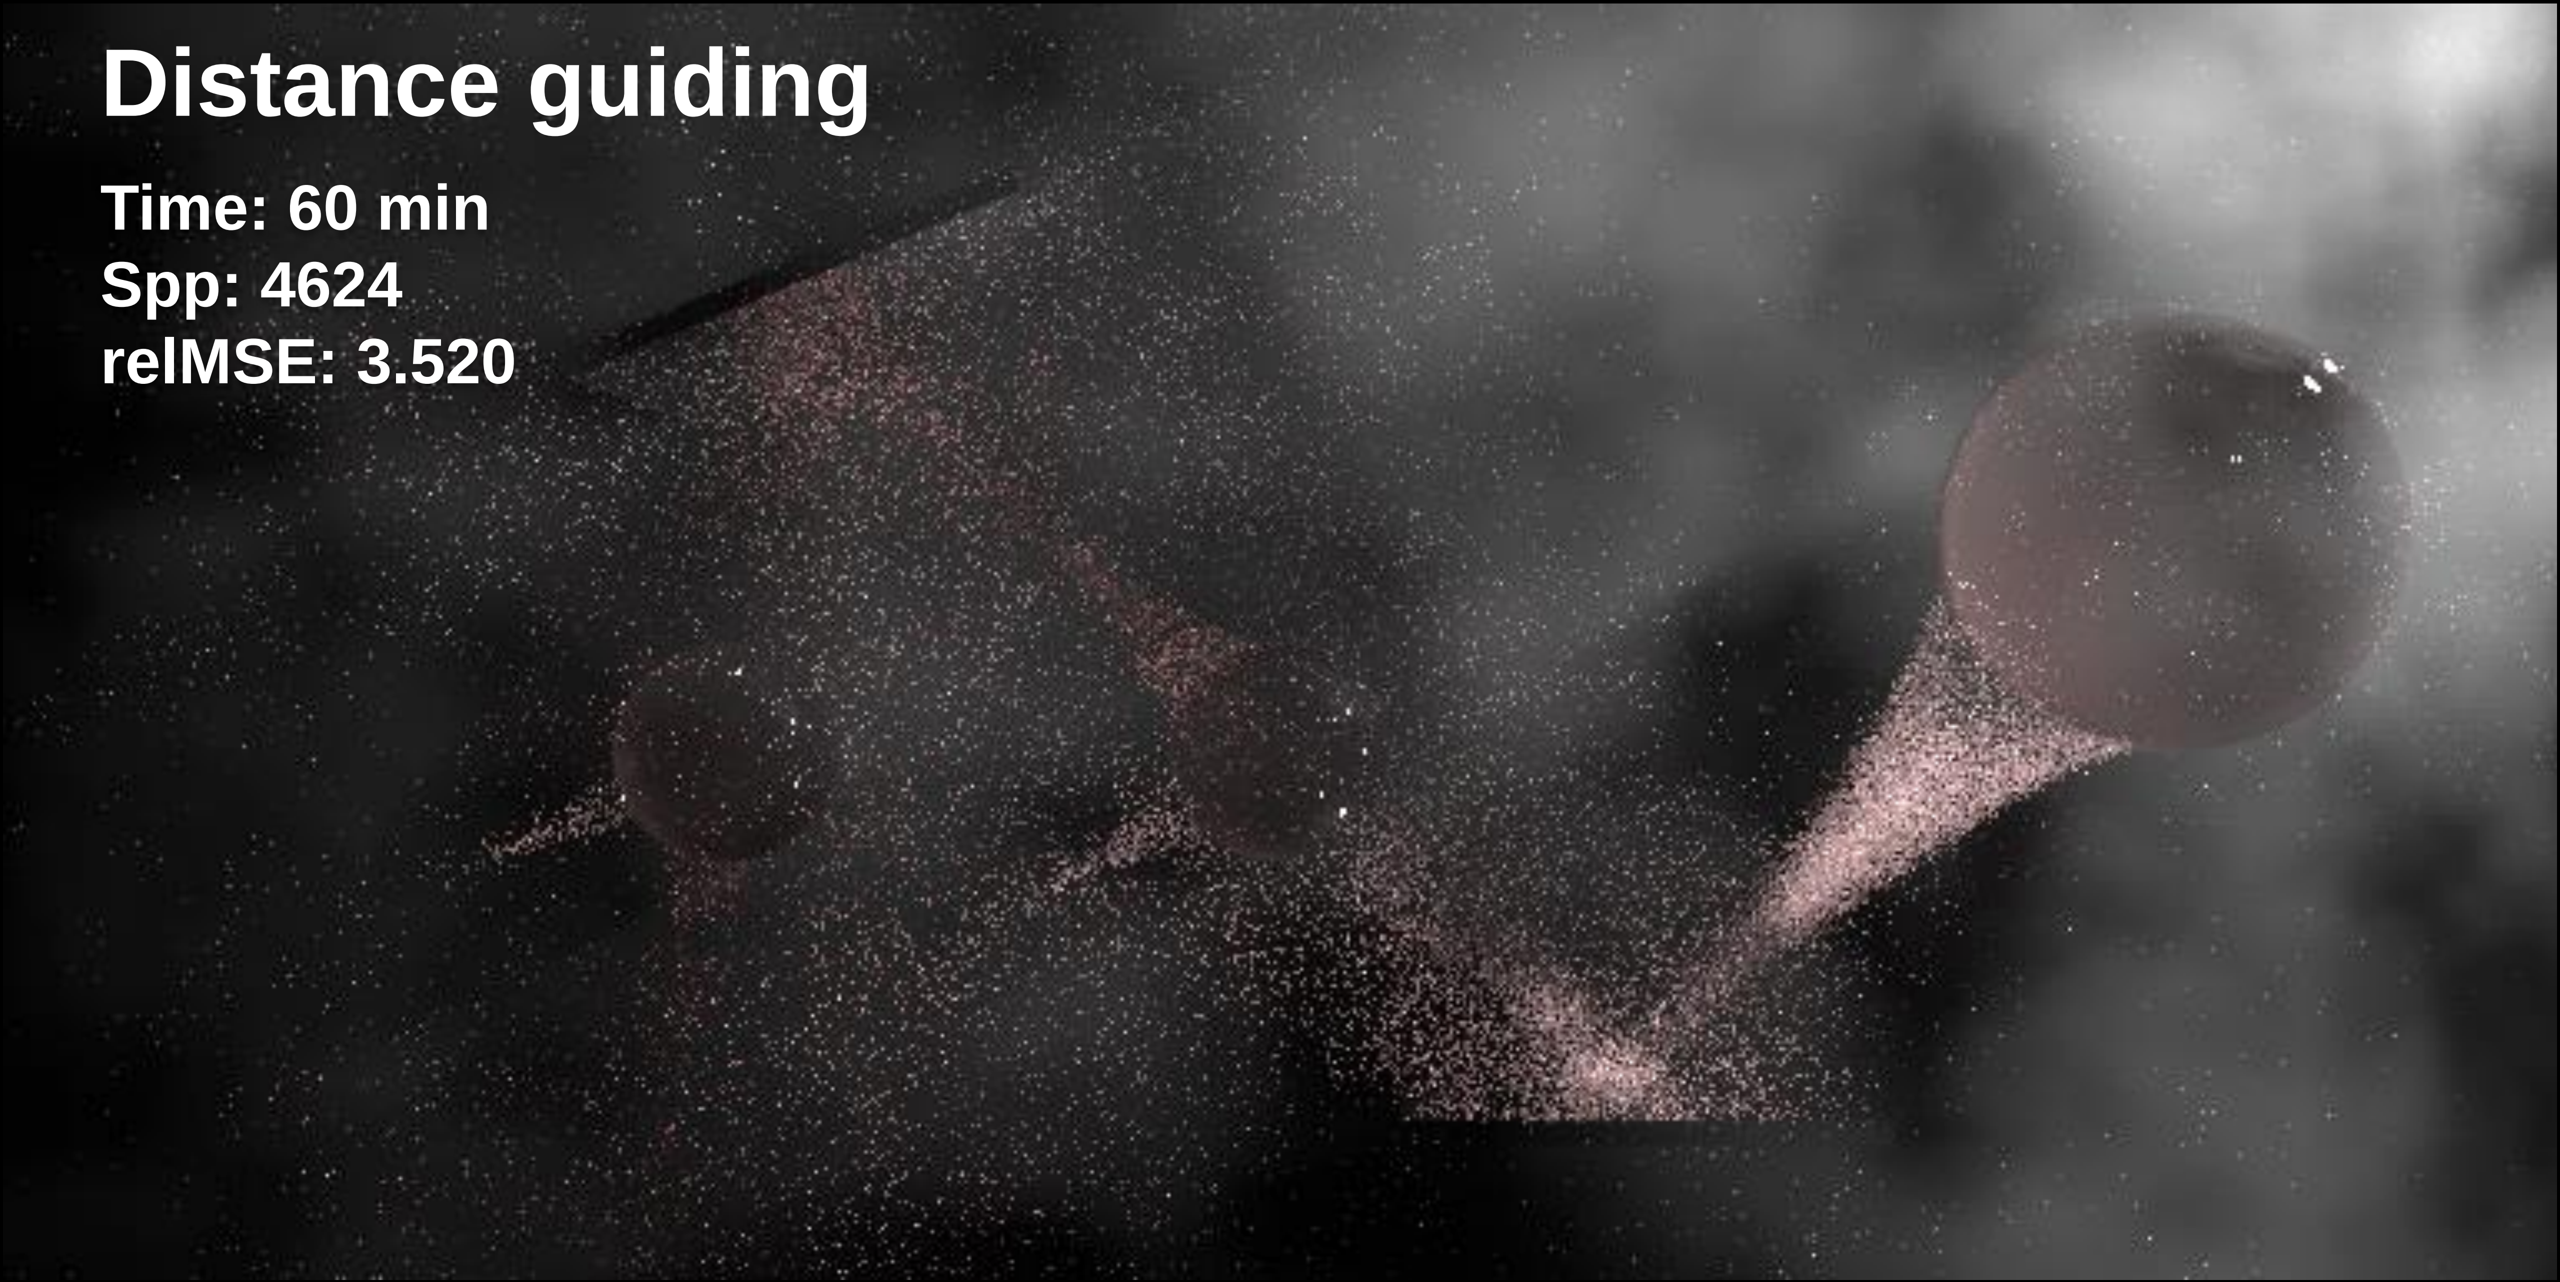

Distance + directional guiding

Time: 60 min
Spp: 4448
relMSE: 0.468


Distance + directional guiding + GRRS
Time: 60 min
Spp: 3796
relMSE: 0.321

ADDITIONAL RESULTS