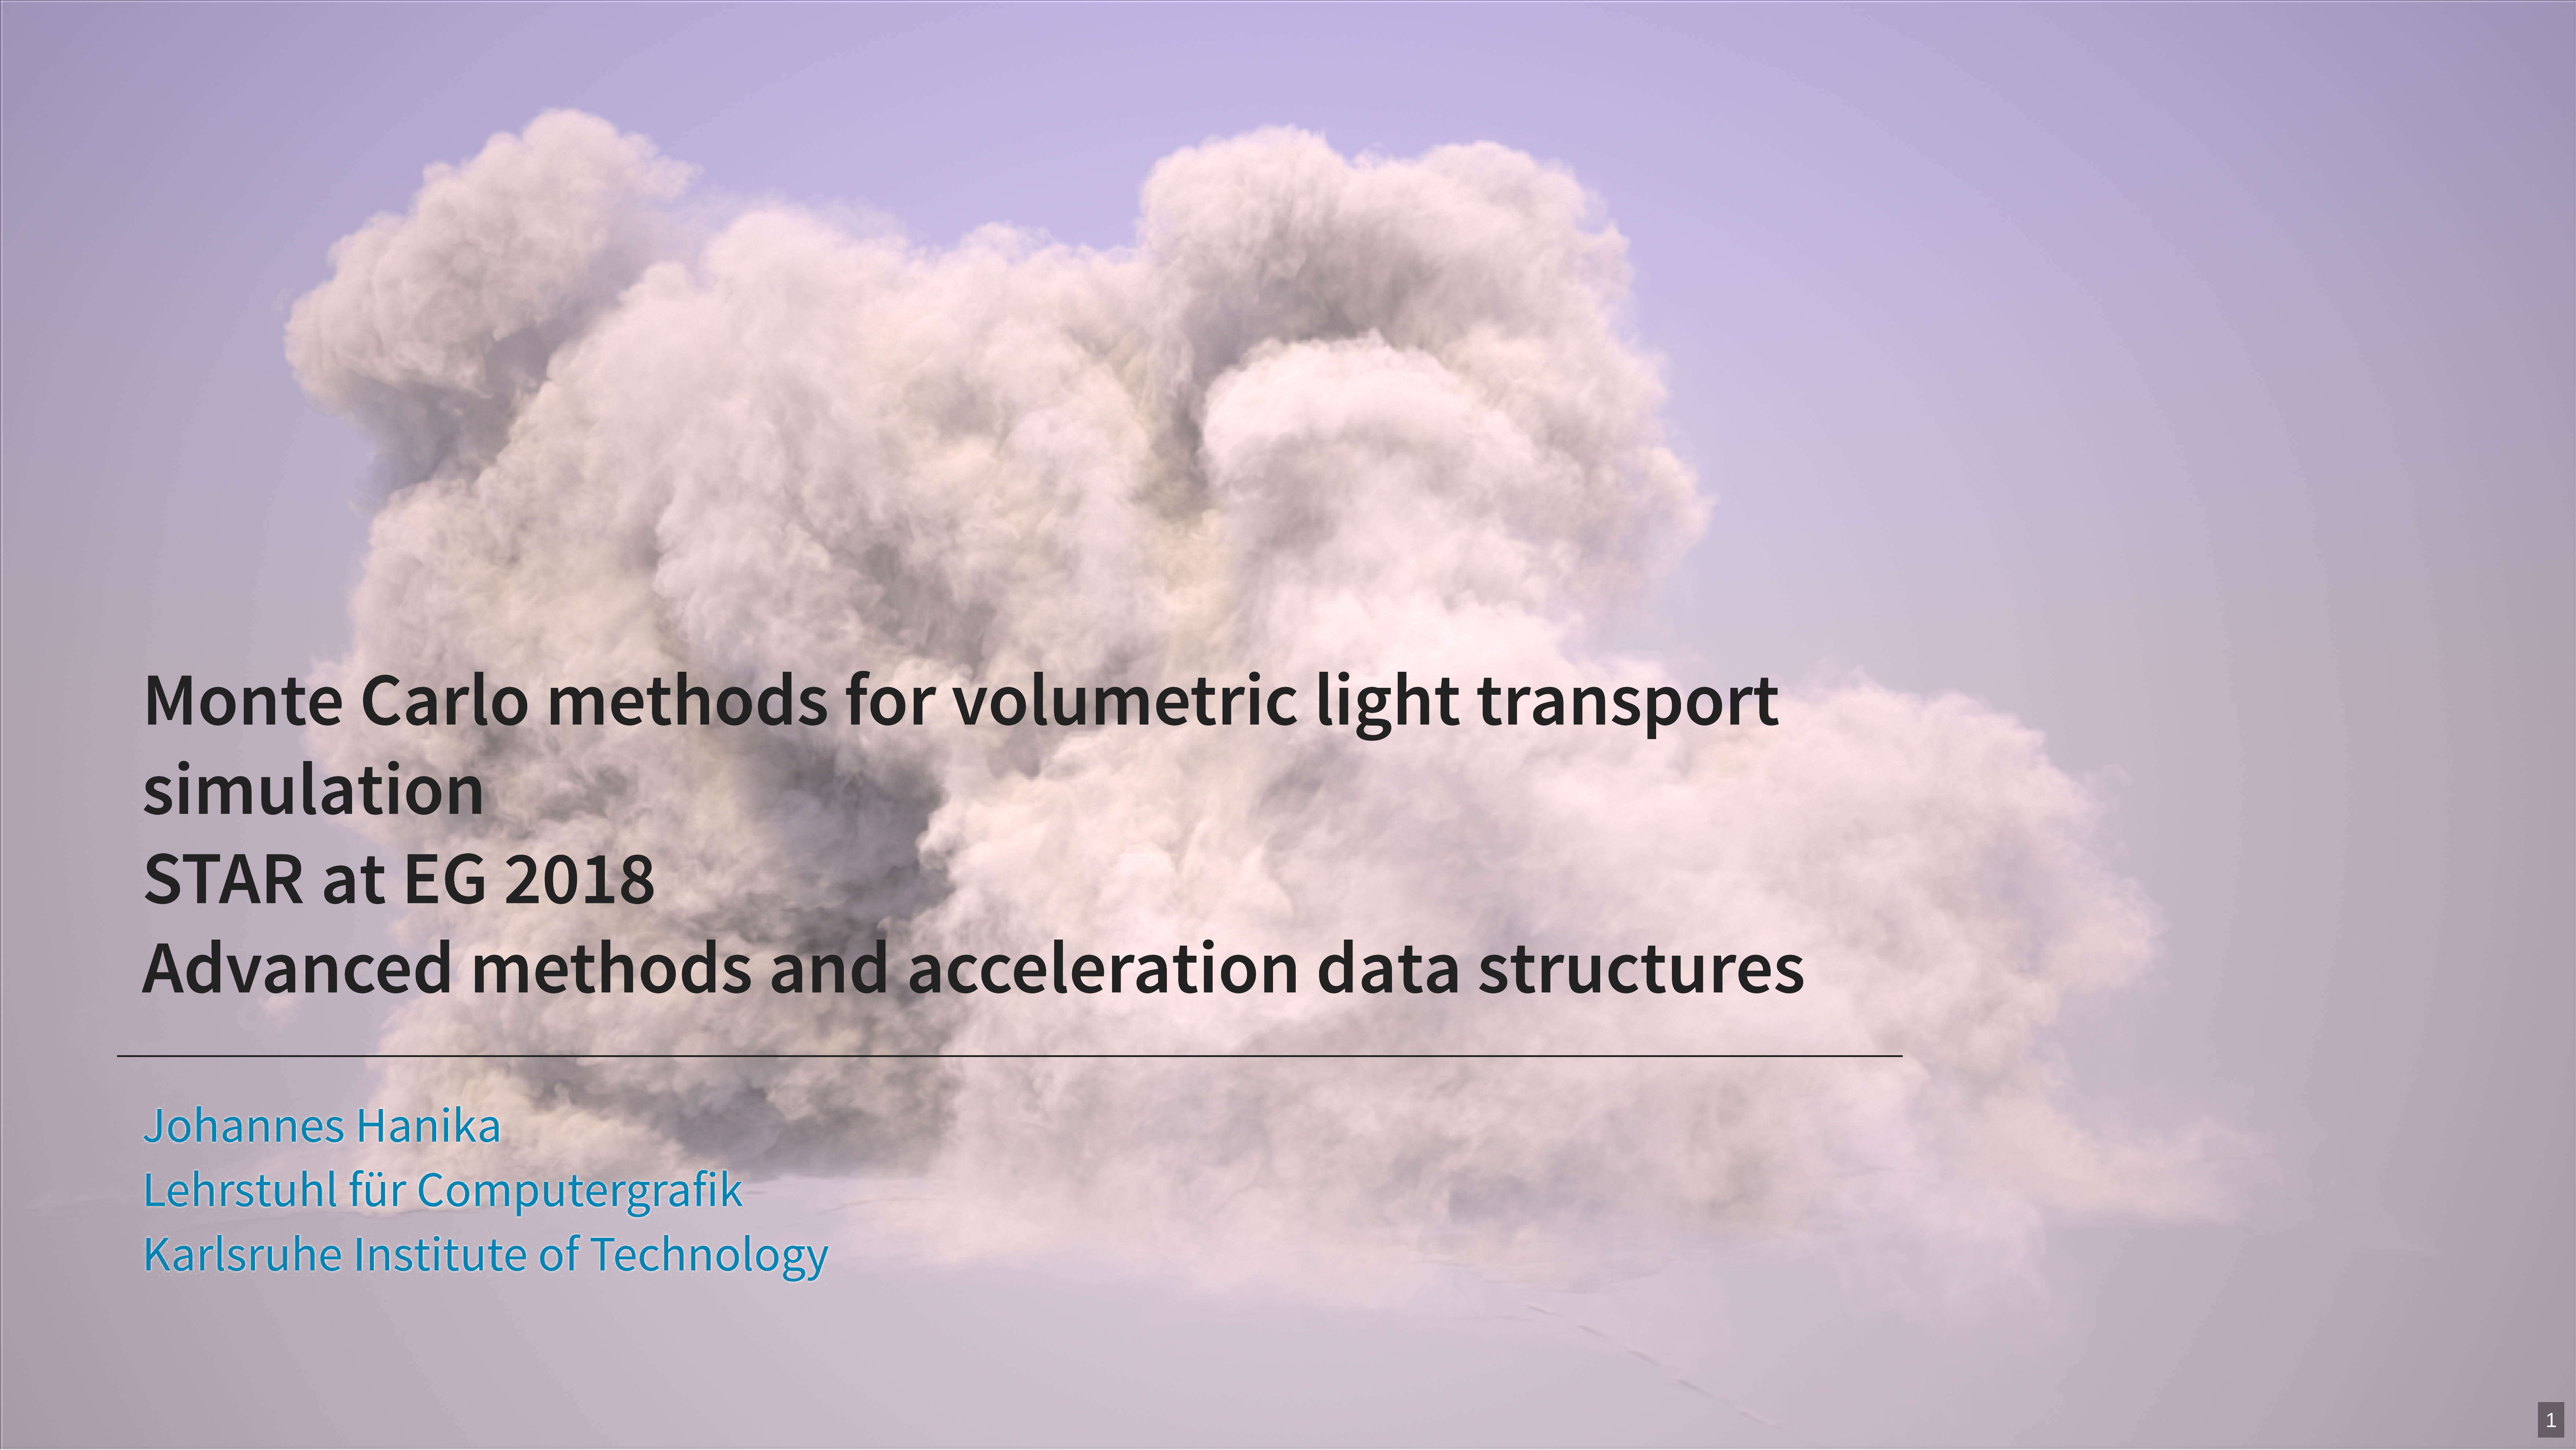

Monte Carlo methods for volumetric light transport simulation
STAR at EG 2018
Advanced methods and acceleration data structures

Johannes Hanika
Lehrstuhl für Computergrafik
Karlsruhe Institute of Technology

contents

- selection of advanced methods:
  - Dwivedi sampling / zero variance random walks
  - spectral tracking
- acceleration data structures
  - for regular tracking
  - inside these: null collision based
- emissive media

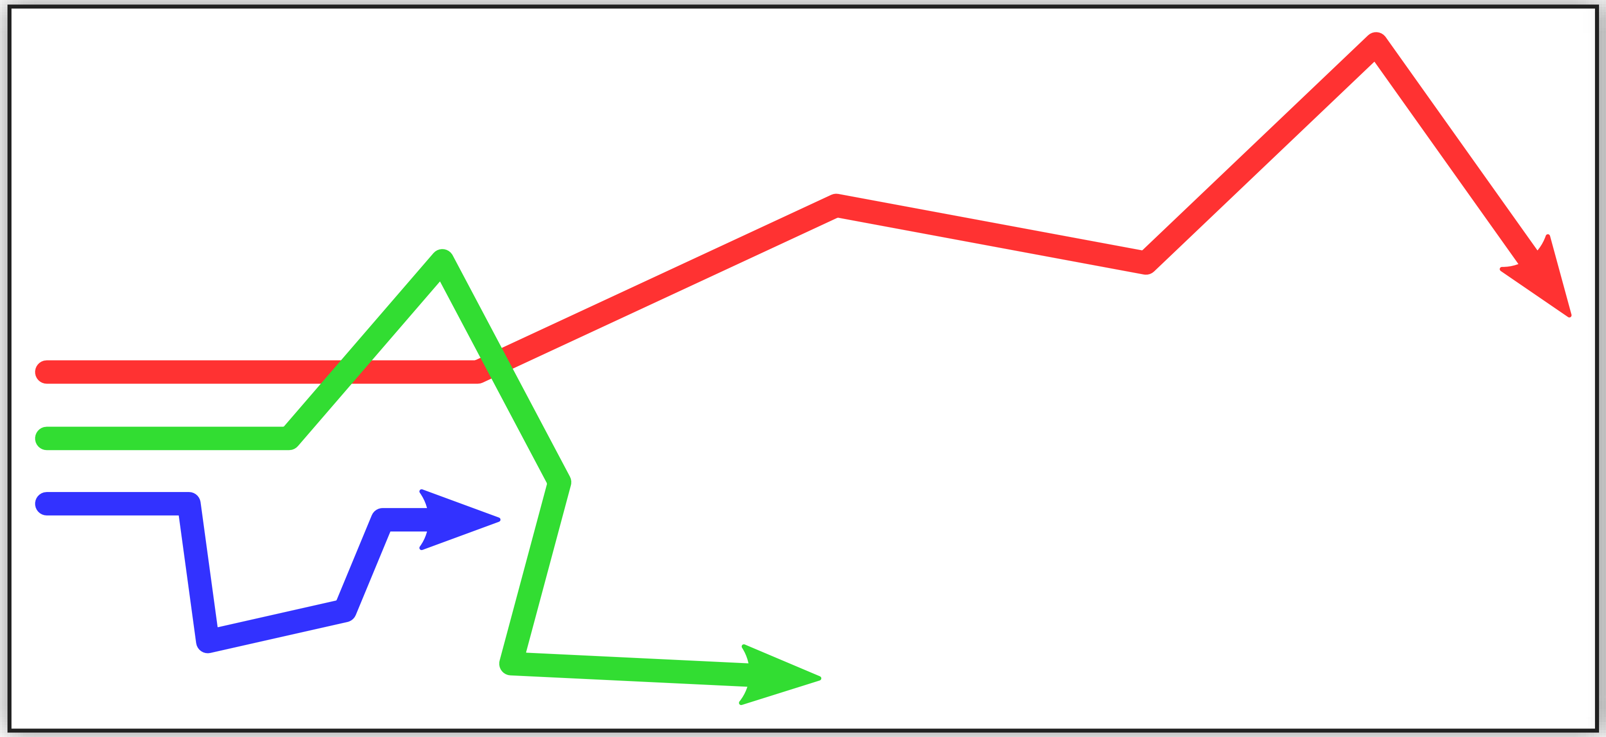
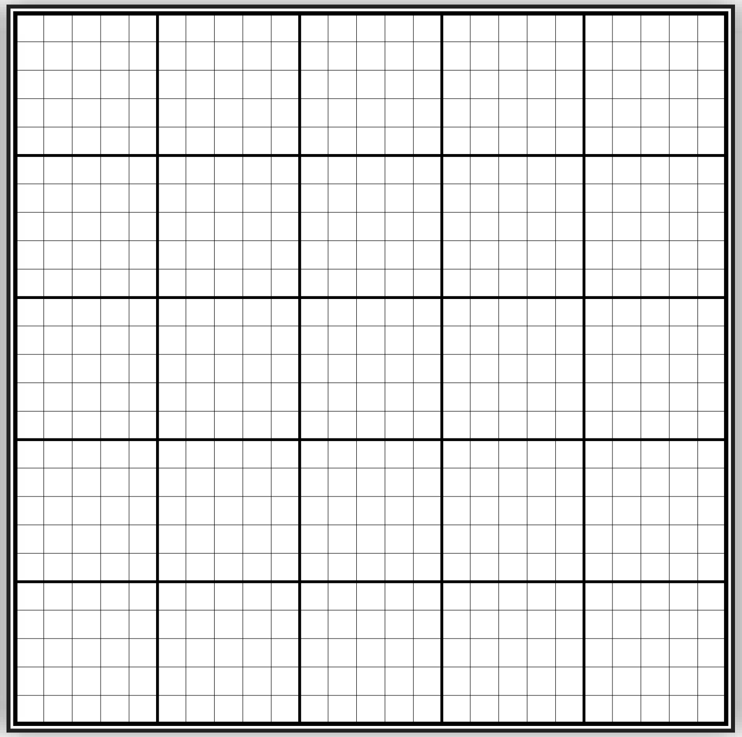
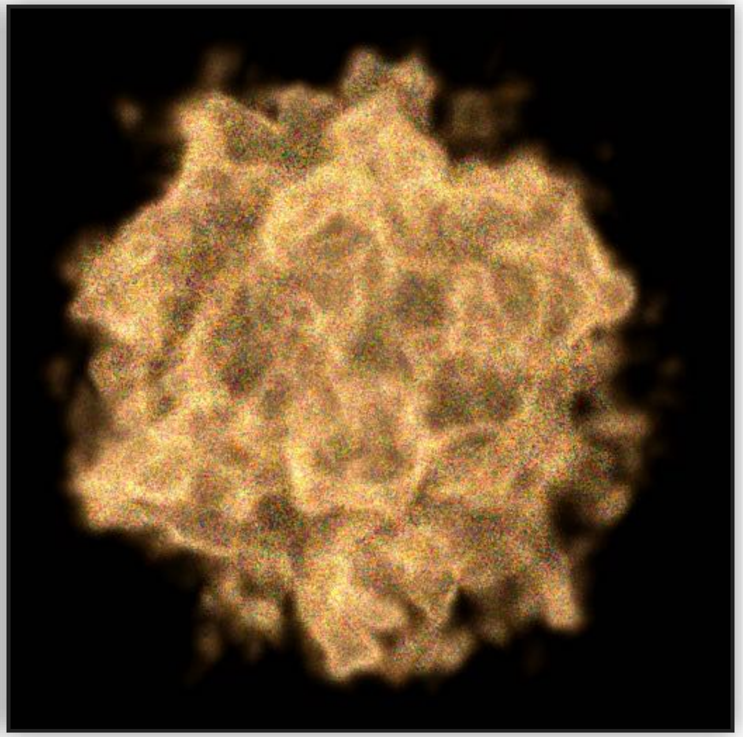

Dwivedi sampling

- regular path tracing random walk
  - tends to get lost inside a **volume bounded by a shape**

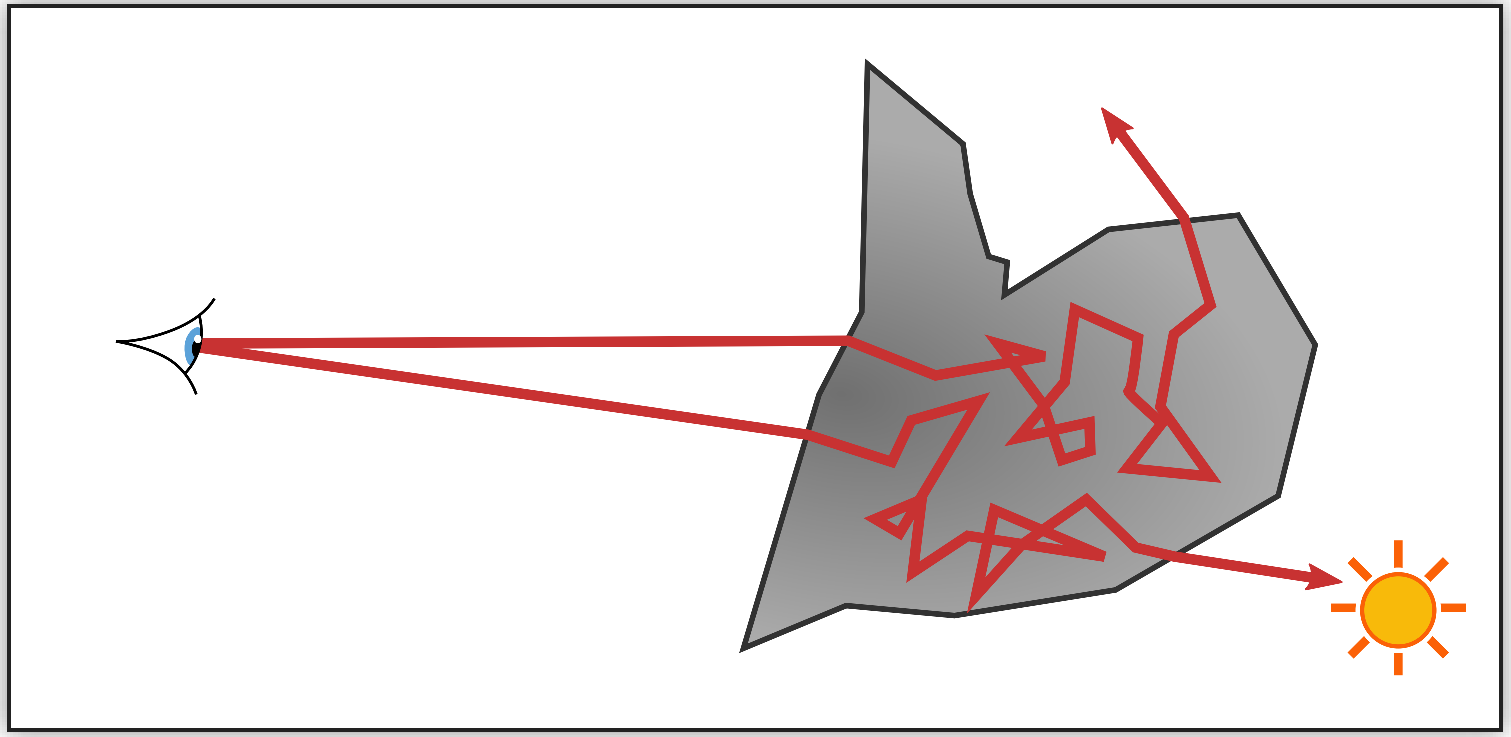

Dwivedi sampling

- random walk biased to exit **bounded volume** as quickly as possible
  - assumes constant illumination from the outside [Kd14]

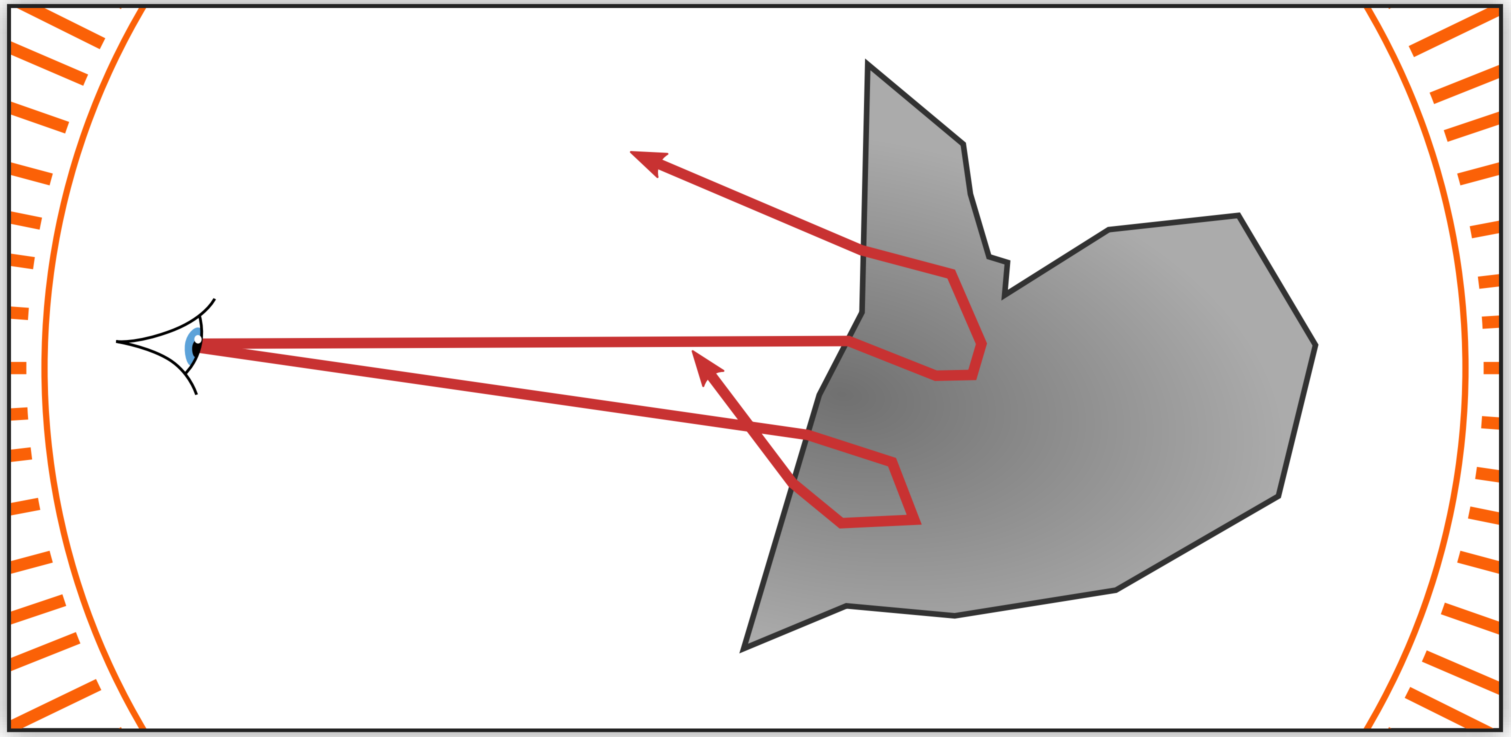

  - assume **homogeneous slab** with **isotropic** phase function
  - approximate closed-form solution of transport using this simplified setting
    - known as **zero-variance theory** (term may be a bit bold)

Dwivedi sampling

- random walk biased towards light source
  - aims to exit towards light [MHD16]

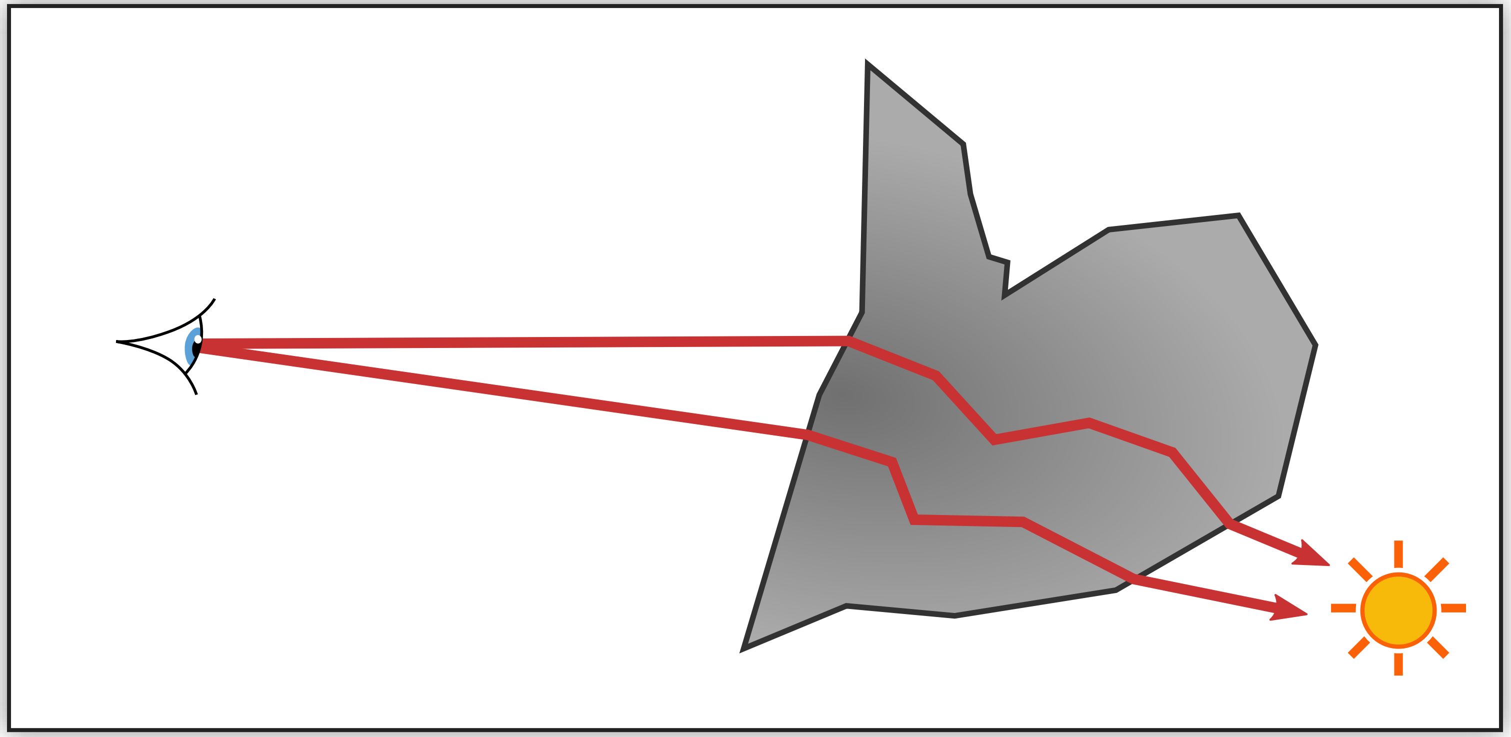

Dwivedi sampling

- random walk biased towards light source
  - aims to exit towards light [MHD16]

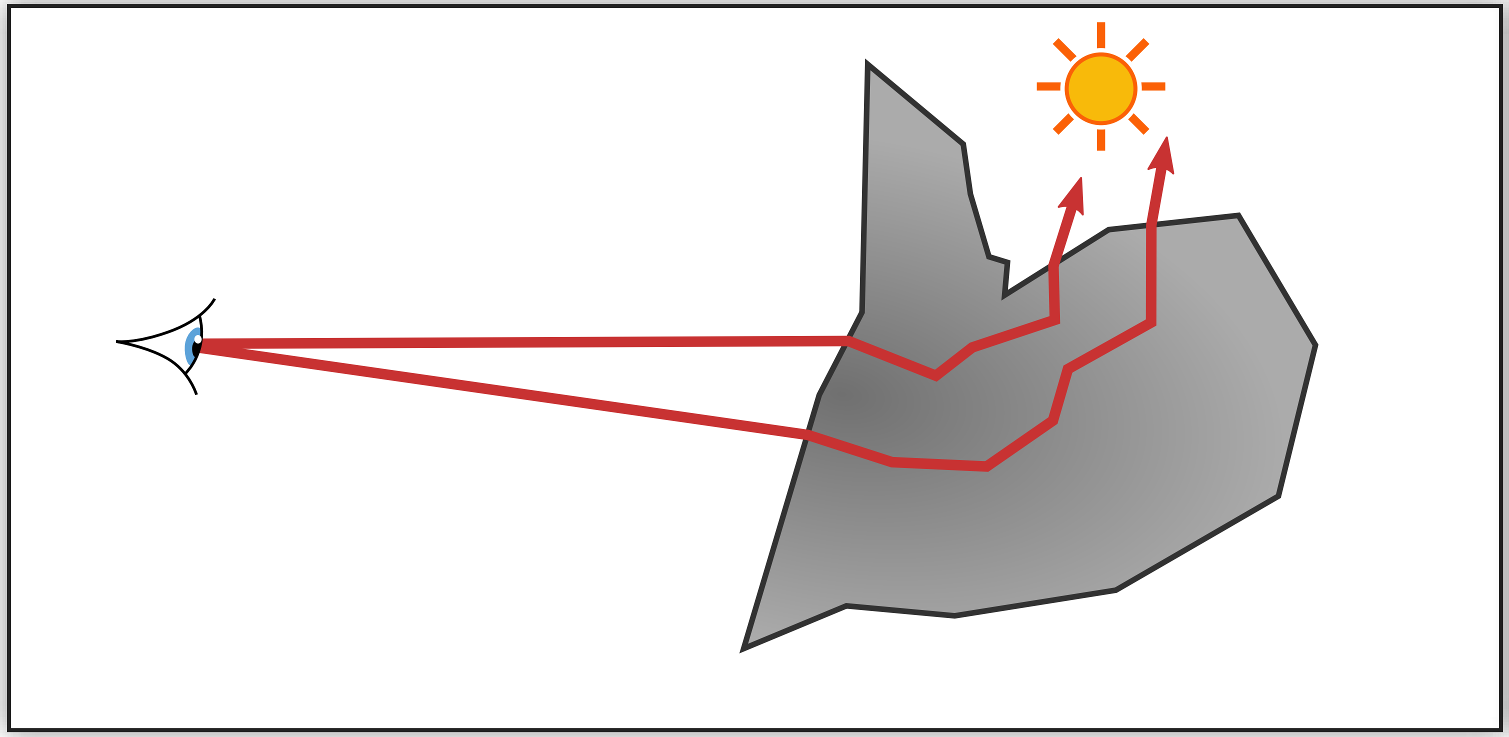
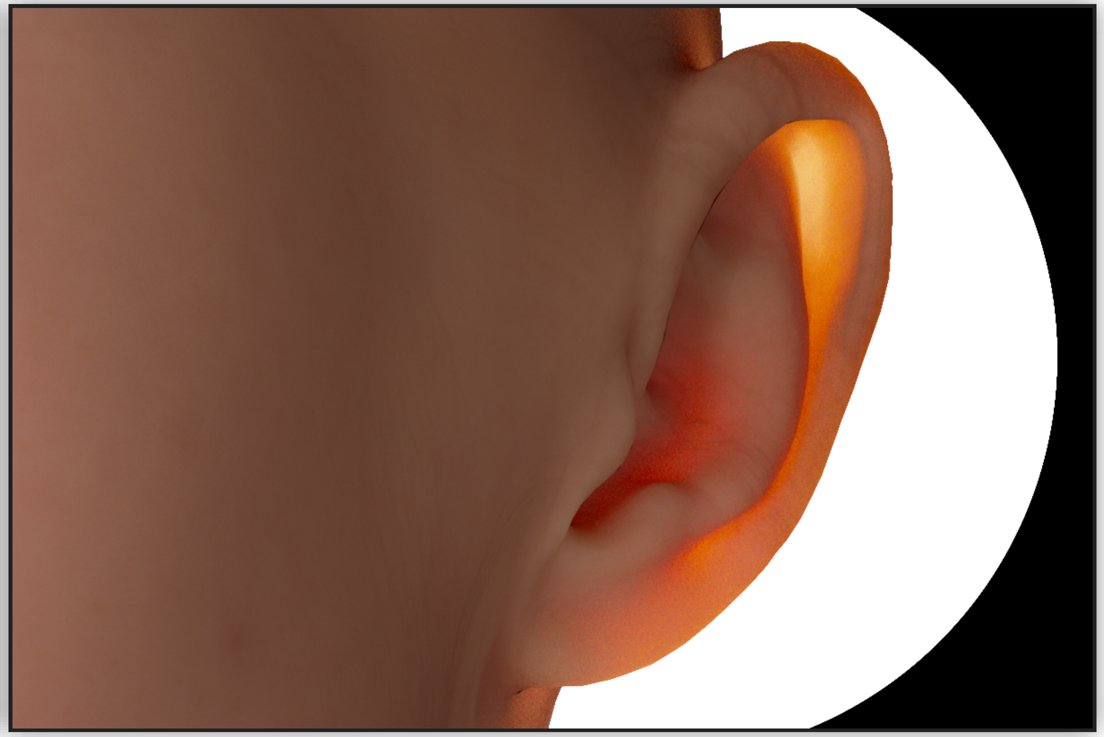

  - achieved by **biasing the PDF** to sample direction and distance
  - estimator remains unbiased!

spectral tracking

- another problem with skin: chromatic media
  - collision coefficients μ depend on wavelength λ
  - for instance free flight distance much longer for long wavelengths:

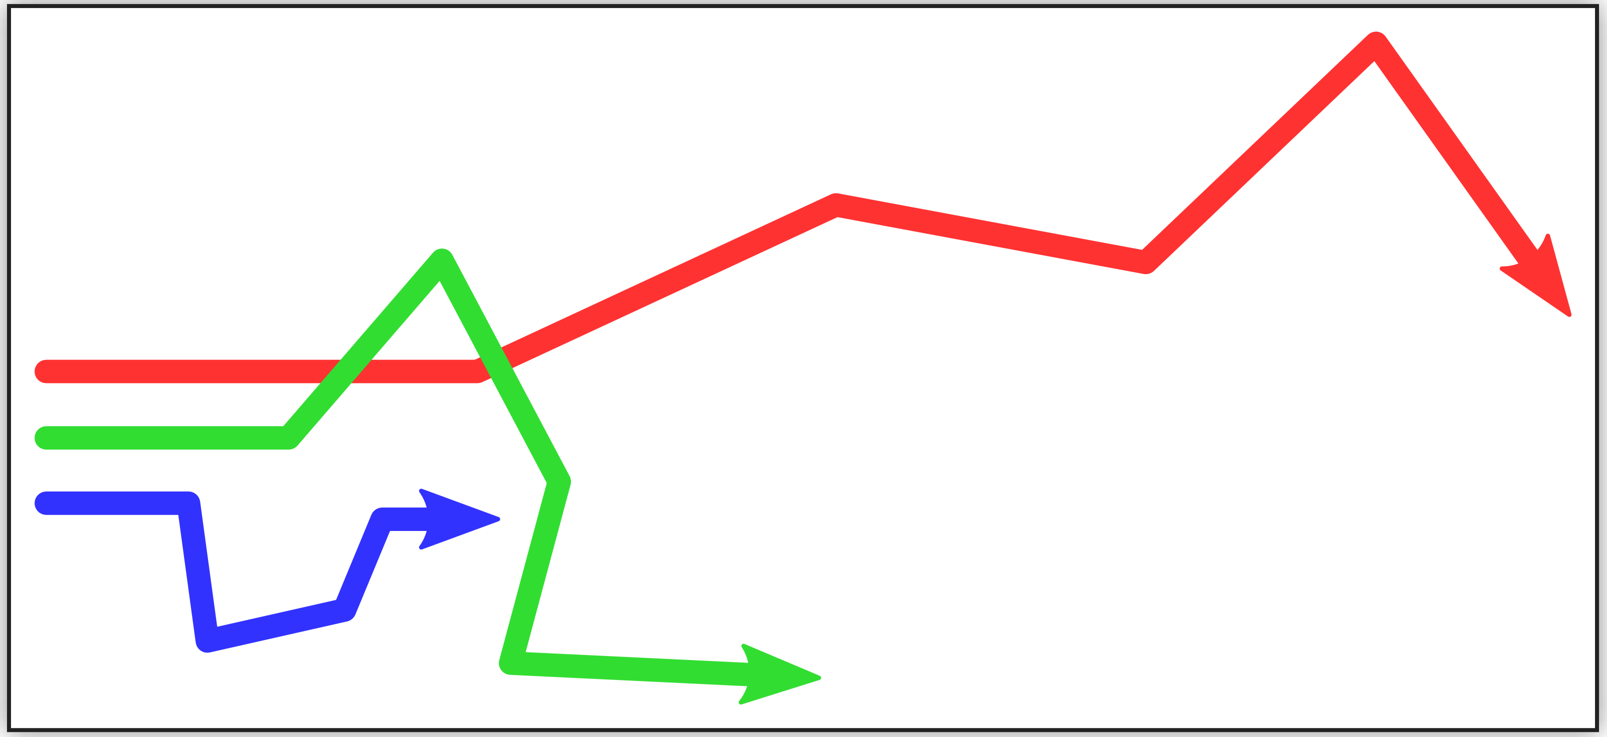

  - makes path invalid for different wavelength?
  - can we still exploit coherence?

spectral tracking via MIS

- **hero wavelength** sampling [WND∗14]

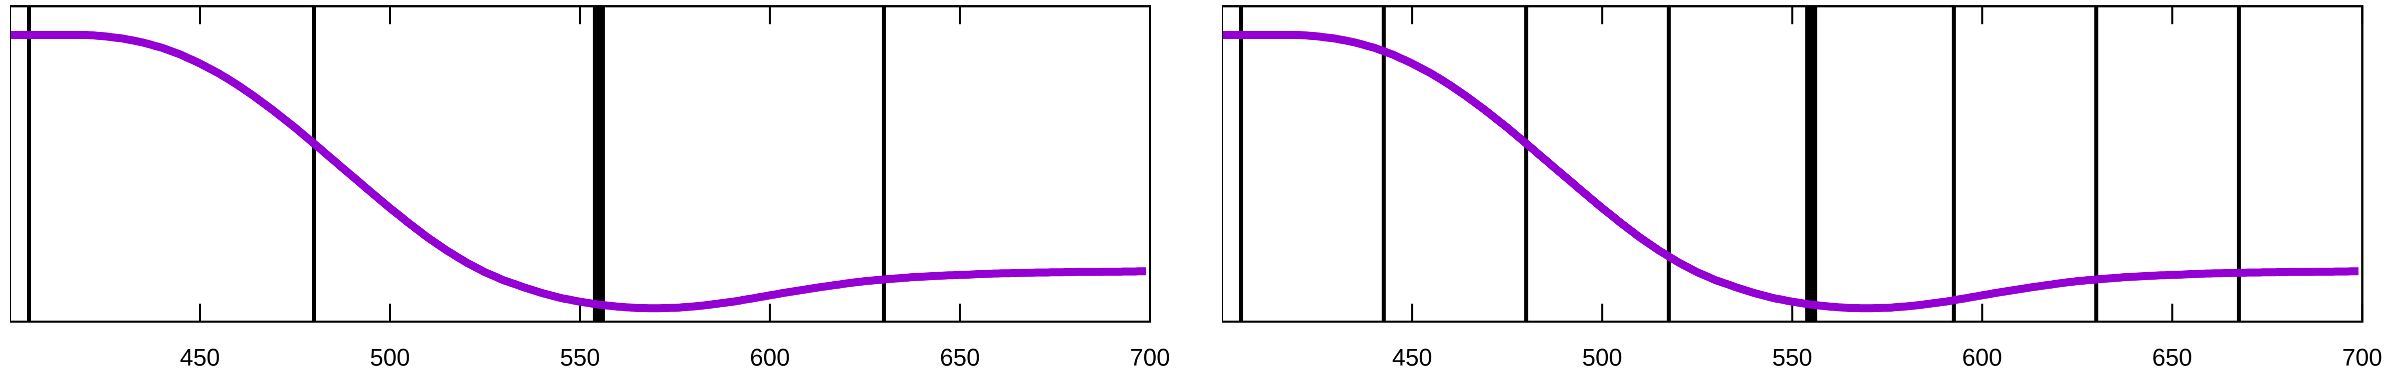

  - sample perfectly for one single wavelength λ_0
  - evaluate path for a stratified set of wavelengths λ_i at the same time
  - optimally weighted combination via **MIS (balance heuristic)**
    - limited to **regular tracking** because it requires explicit evaluation of PDF

$$\frac{f(\bar{\mathbf{x}}, \lambda_i)}{\sum_j p(\bar{\mathbf{x}}, \lambda_j)}$$

spectral tracking via MIS

image comparison 64spp

- skin material with 1 wavelength

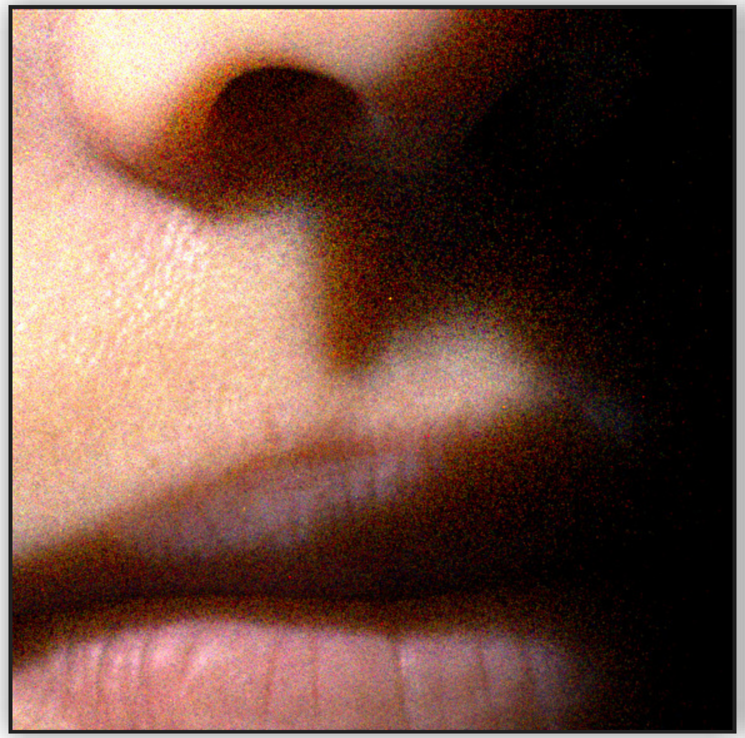
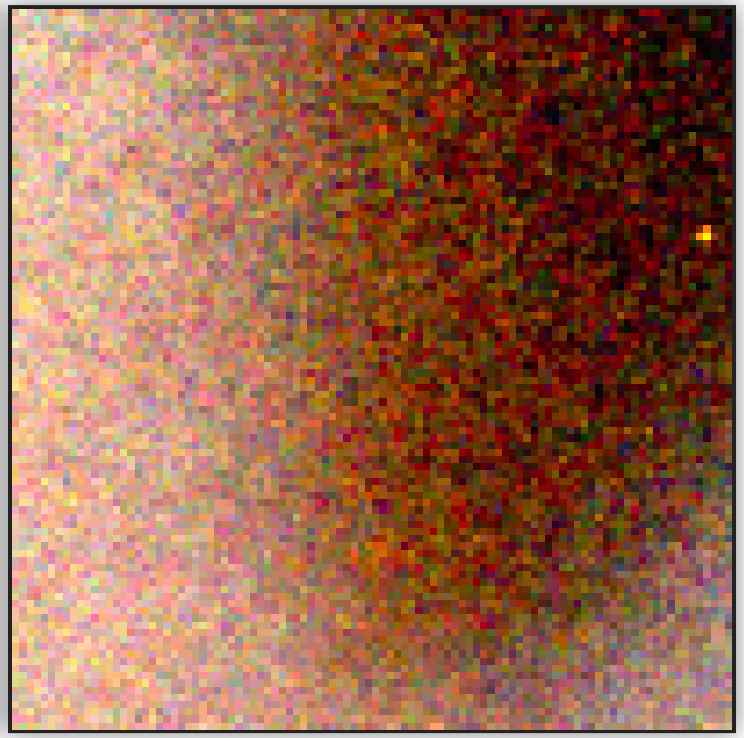

spectral tracking via MIS

image comparison 64spp

- skin material with 4 wavelengths (SSE)

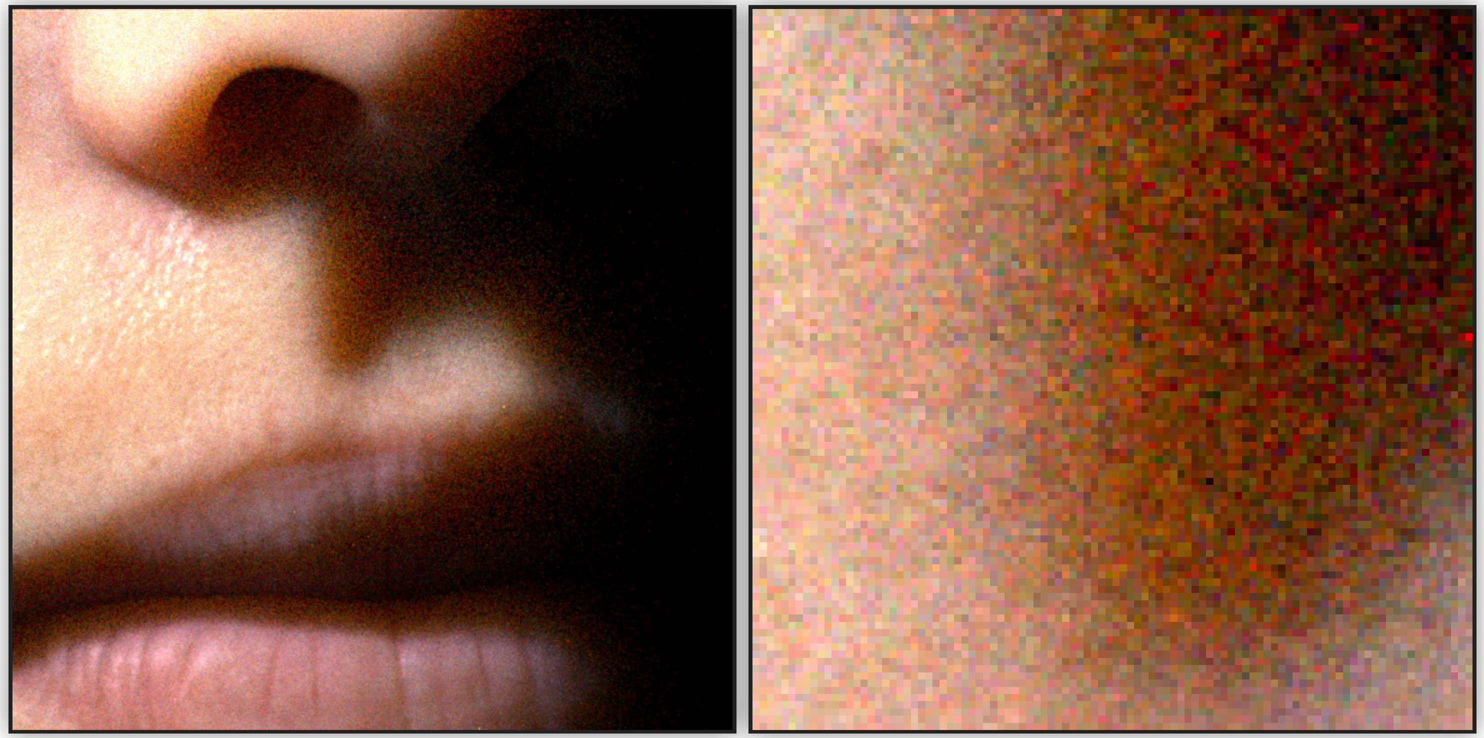

spectral tracking via MIS

image comparison 64spp

- skin material with 8 wavelengths (AVX)

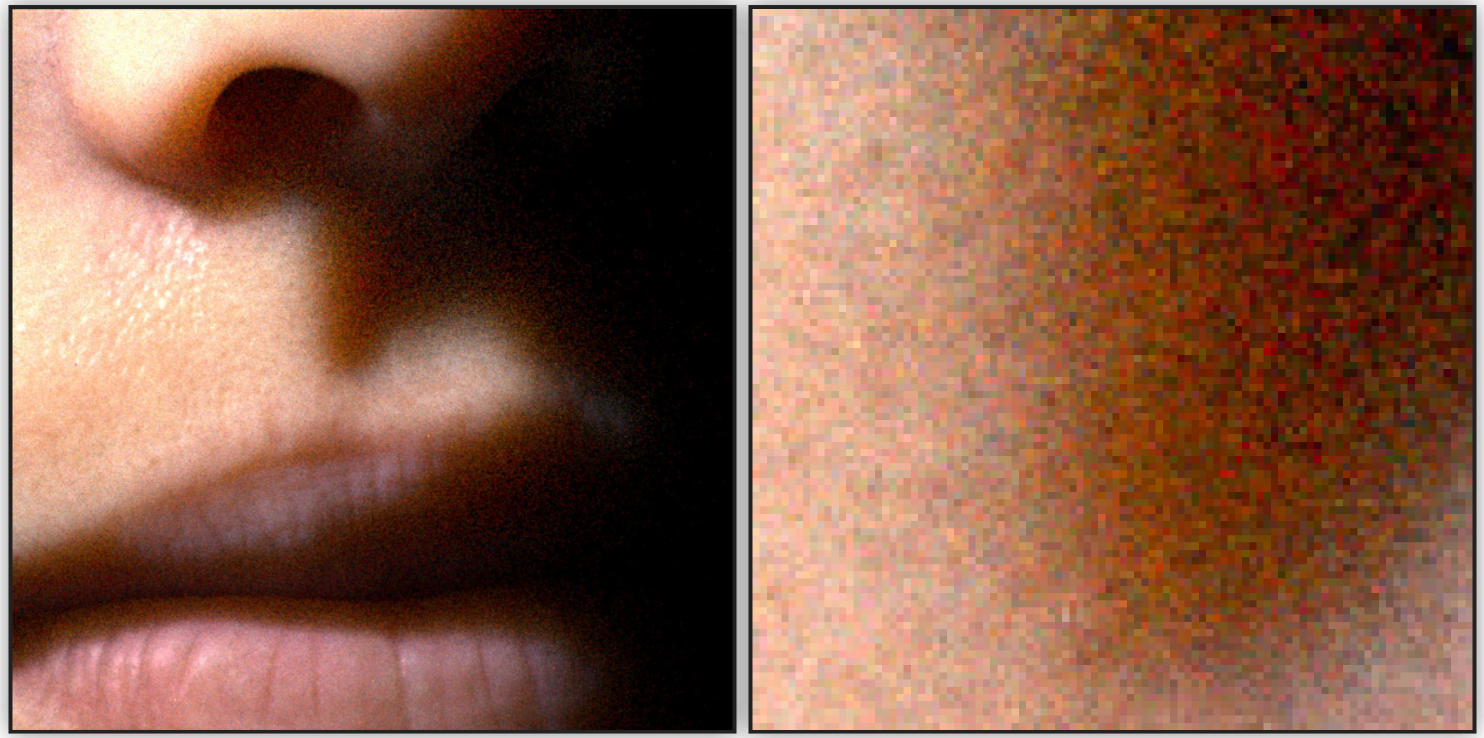

spectral tracking without PDF [KHLN17]

- sample by common majorant $\bar{\mu}$
- how do decide for null collision, scattering, or absorption?
- probability according to $\mu_n(\lambda), \mu_s(\lambda), \mu_a(\lambda)$
  - pick by maximum over λ_i
  - pick by average weighted by spectral path throughput history

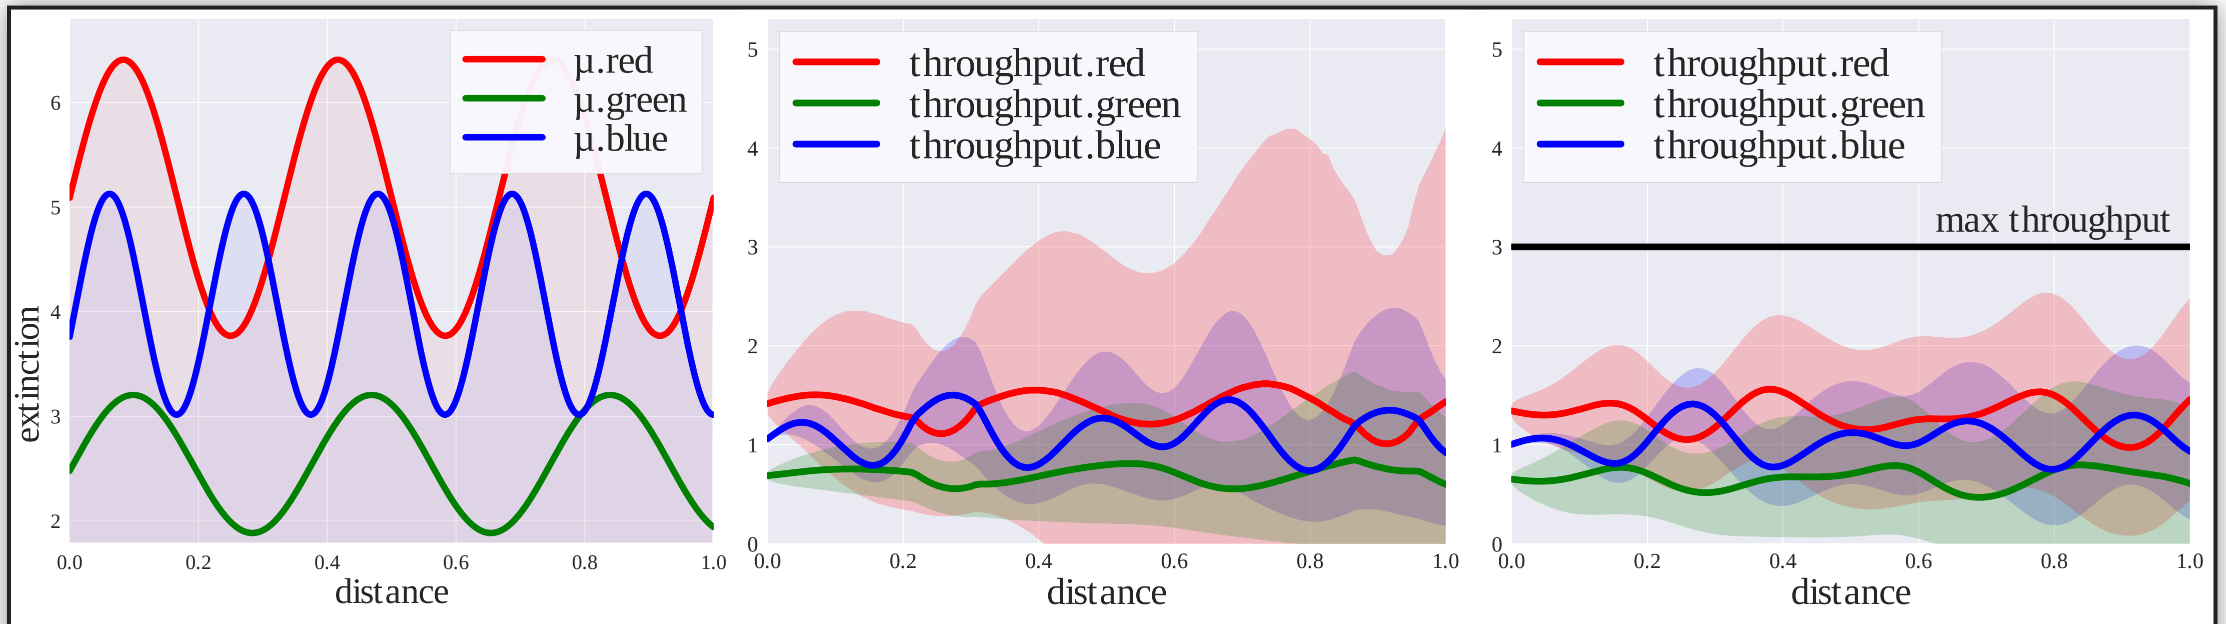

RGB medium Maximum History-aware avg.

spectral tracking without PDF [KHLN17]

- sample by common majorant $\bar{\mu}$
- how do decide for null collision, scattering, or absorption?
- probability according to $\mu_n(\lambda), \mu_s(\lambda), \mu_a(\lambda)$
  - pick by maximum over λ_i
  - pick by average weighted by spectral path throughput history
  - results in different noise patterns:

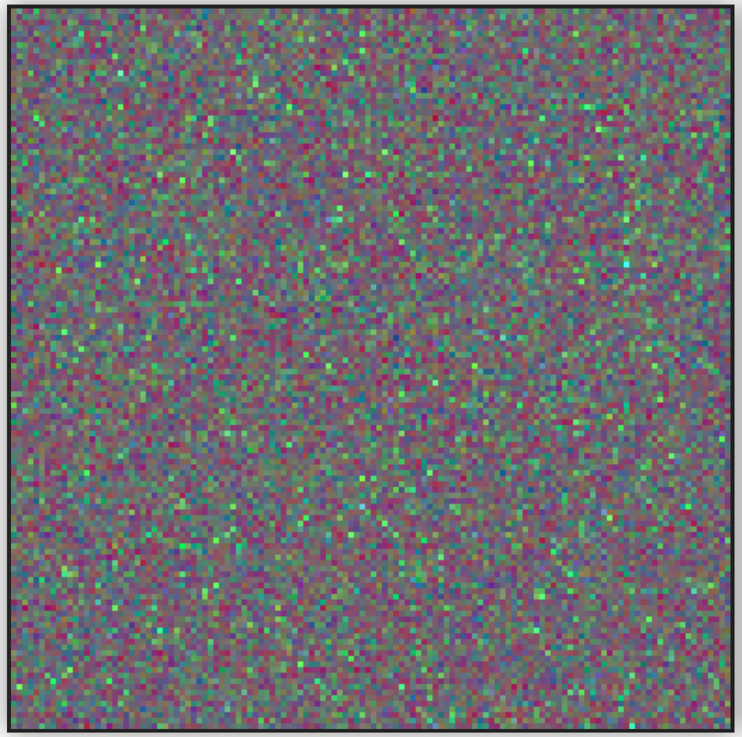
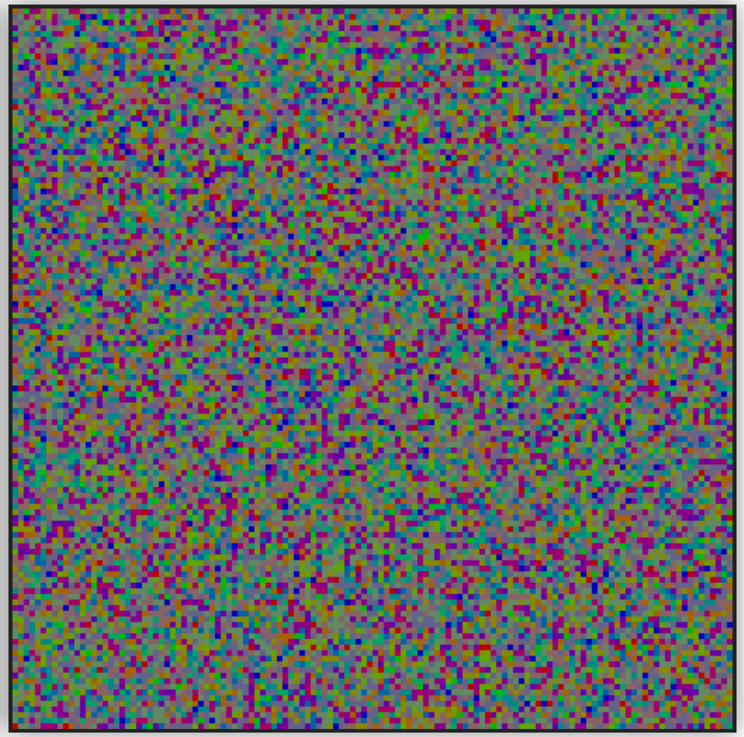

speed!

low variance estimators are important

- but also, in volumes most of the run time is memory fetching

acceleration data structures

- grid, super voxels [SKTM11], kd-tree [YIC∗11], adaptive blocks

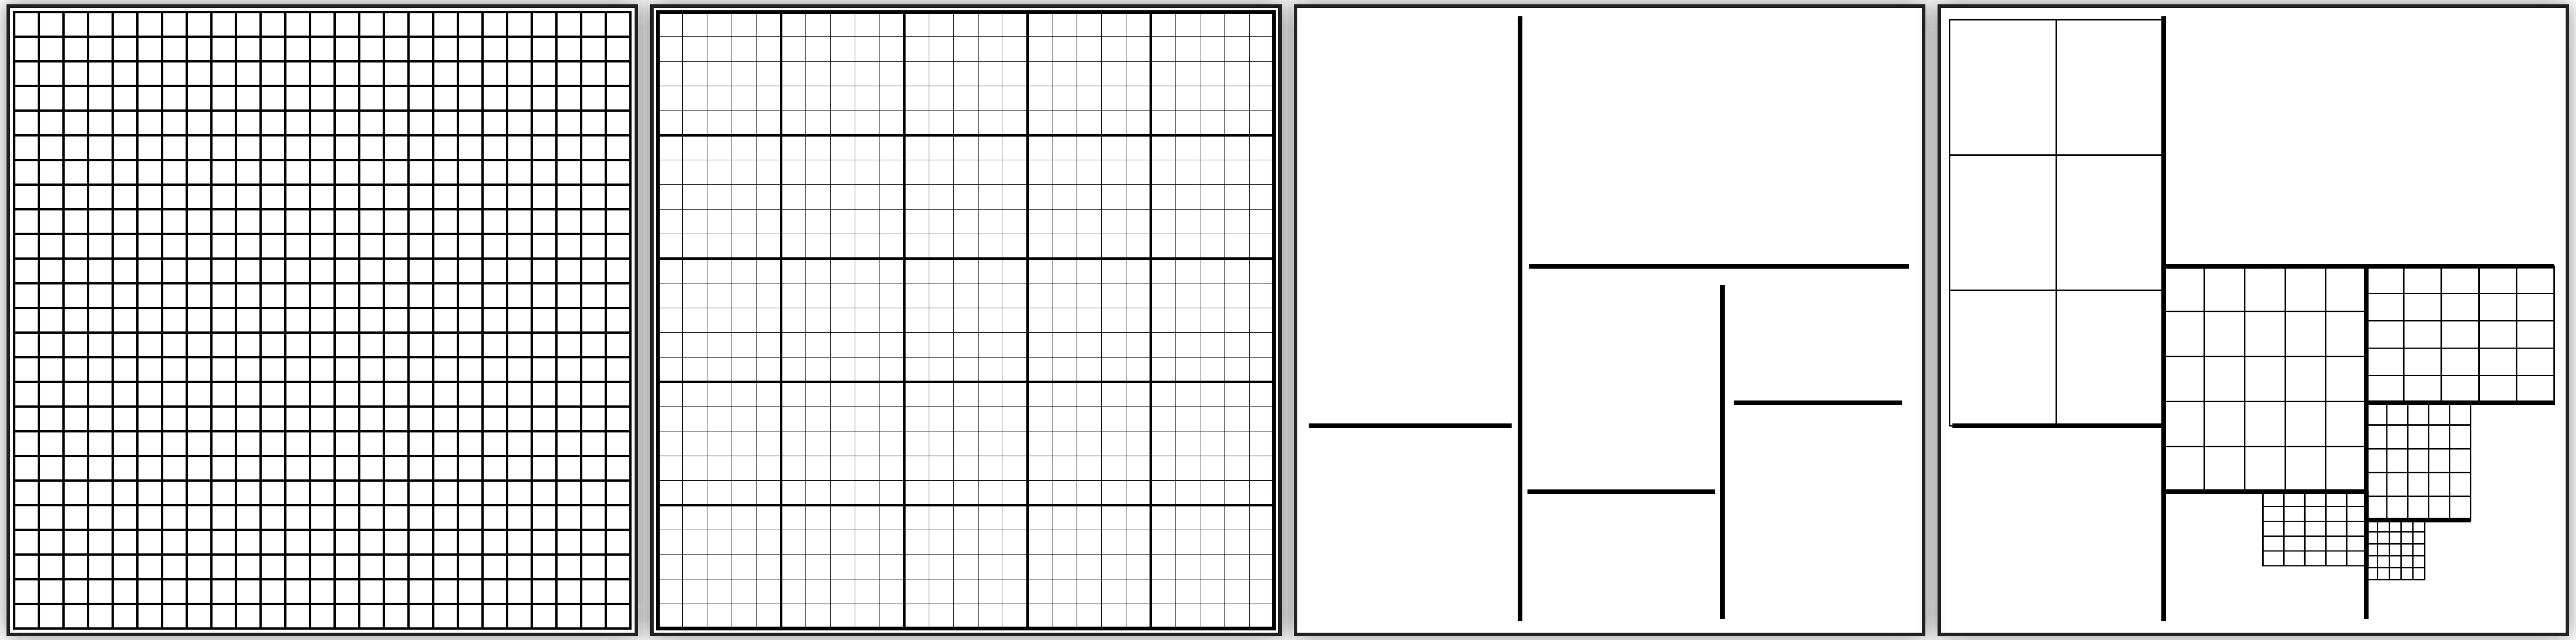

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- two-level modelling (super voxel, kd nodes) store majorants $\bar{\mu}$ in coarse blocks

acceleration data structures

- grid, super voxels [SKTM11], kd-tree [YIC∗11], adaptive blocks

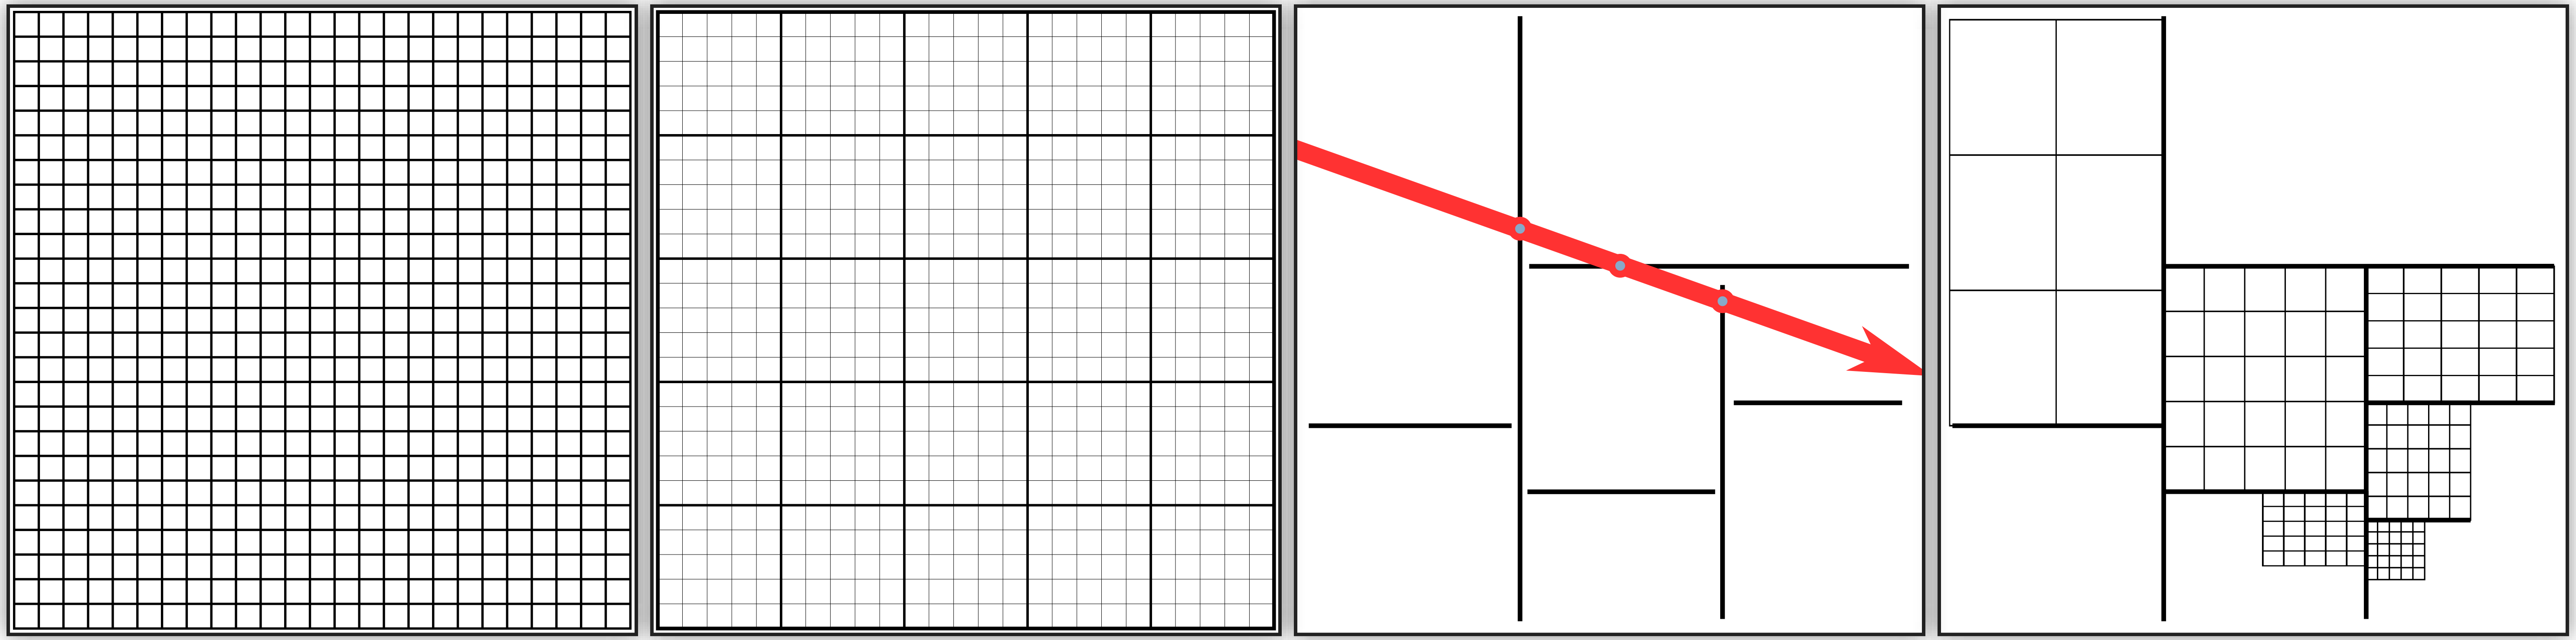

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- two-level modelling (super voxel, kd nodes) store majorants $\bar{\mu}$ in coarse blocks
  - perform **regular tracking** on coarse blocks [SKTM11]

acceleration data structures

- grid, super voxels [SKTM11], kd-tree [YIC∗11], adaptive blocks

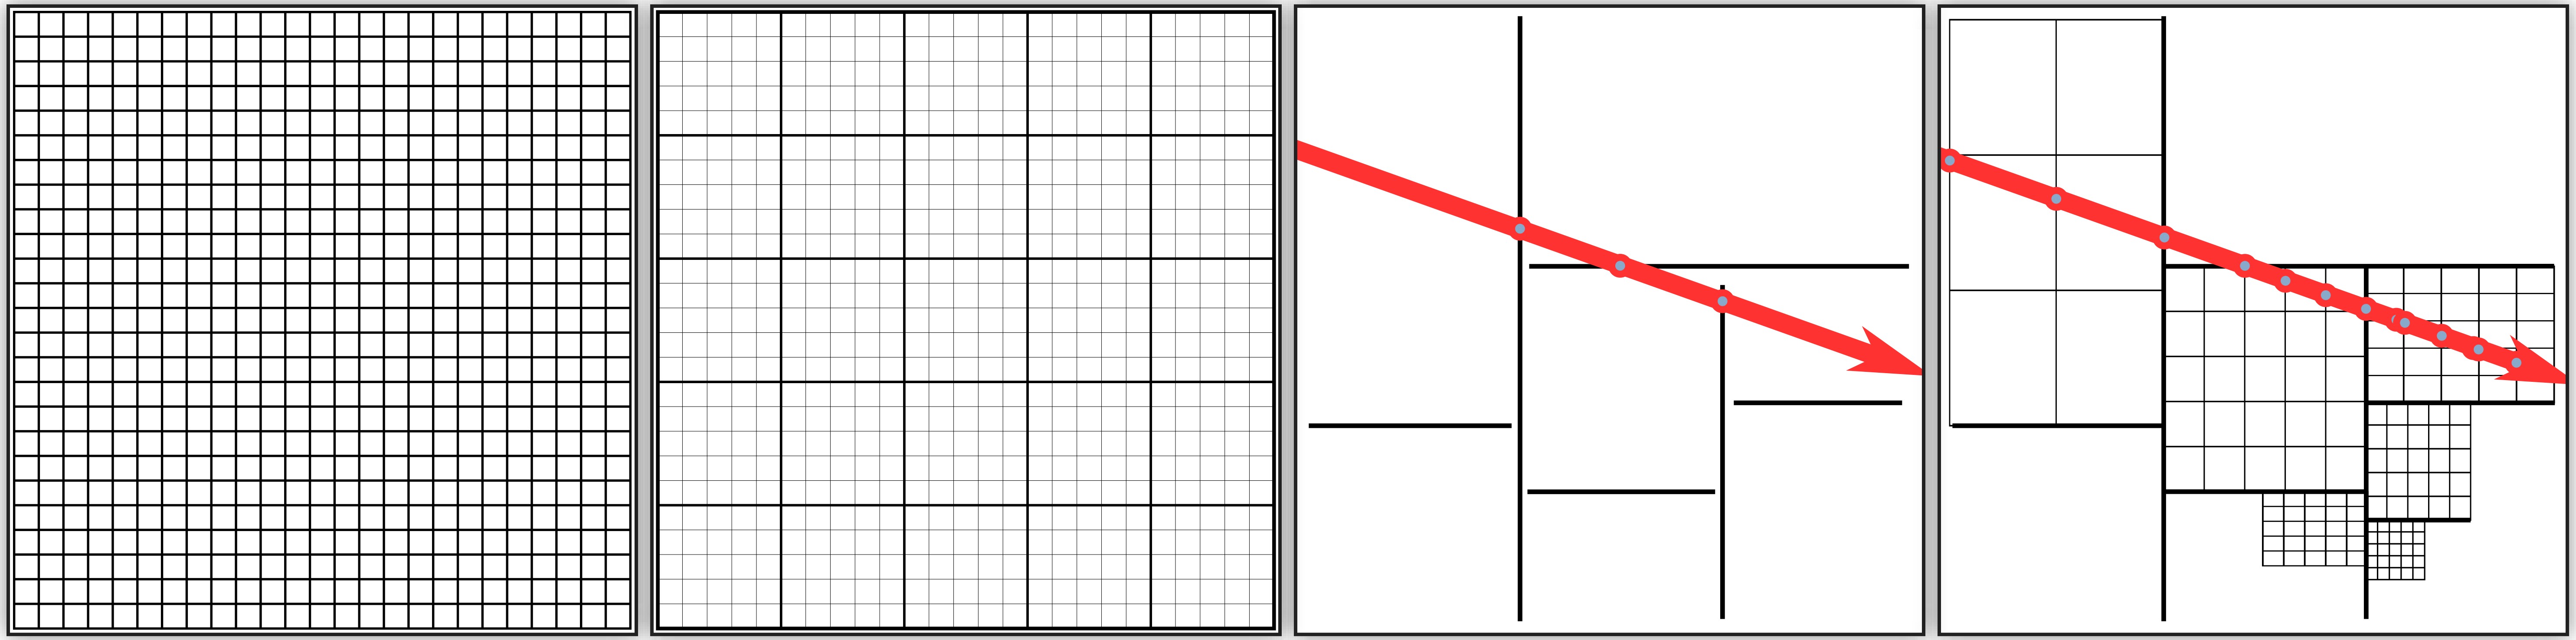

- adaptivity driven by
  - pixel footprint / camera tessellation
  - heterogeneity / variation
- two-level modelling (super voxel, kd nodes) store majorants $\bar{\mu}$ in coarse blocks
  - perform **regular tracking** on coarse blocks [SKTM11]
  - access $\mu_s(\lambda)$, $\mu_a(\lambda)$ on fine levels to sample collision type

acceleration data structures

regular tracking

- needs to step through *every* voxel, bad for fine tessellations
- well chosen tessellation is a big advantage!

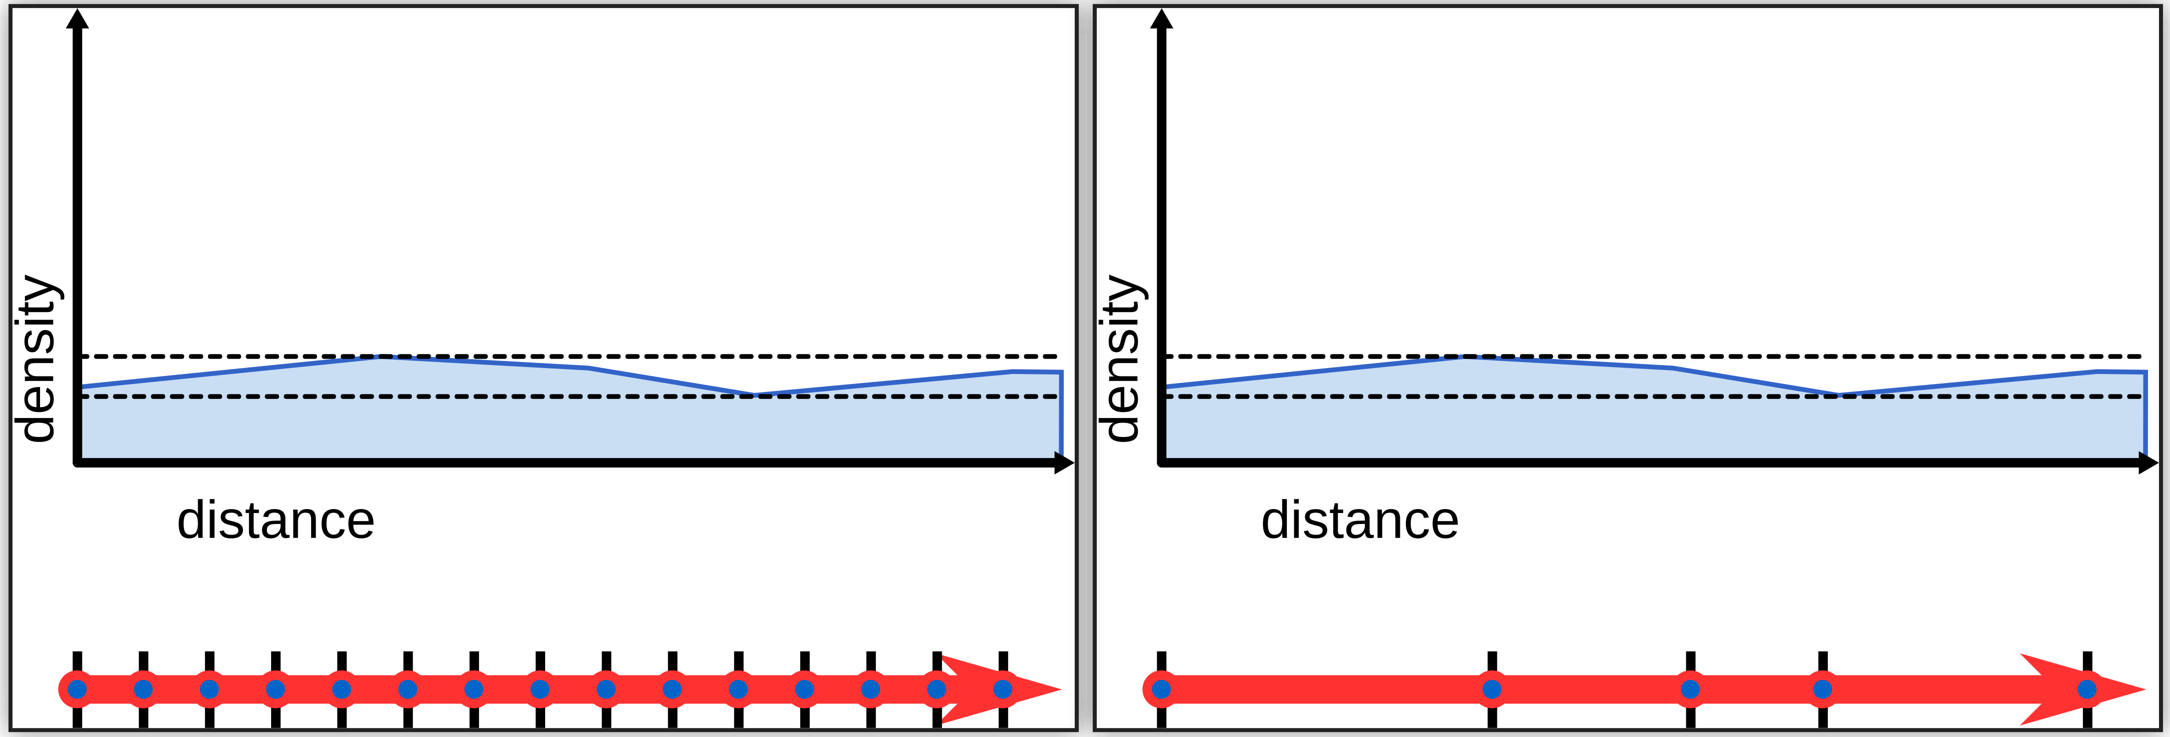

acceleration data structures

null collision-based tracking

- is independent of tessellation and is efficient in thin media (few events)
- high number of events in dense media, regardless of tessellation!

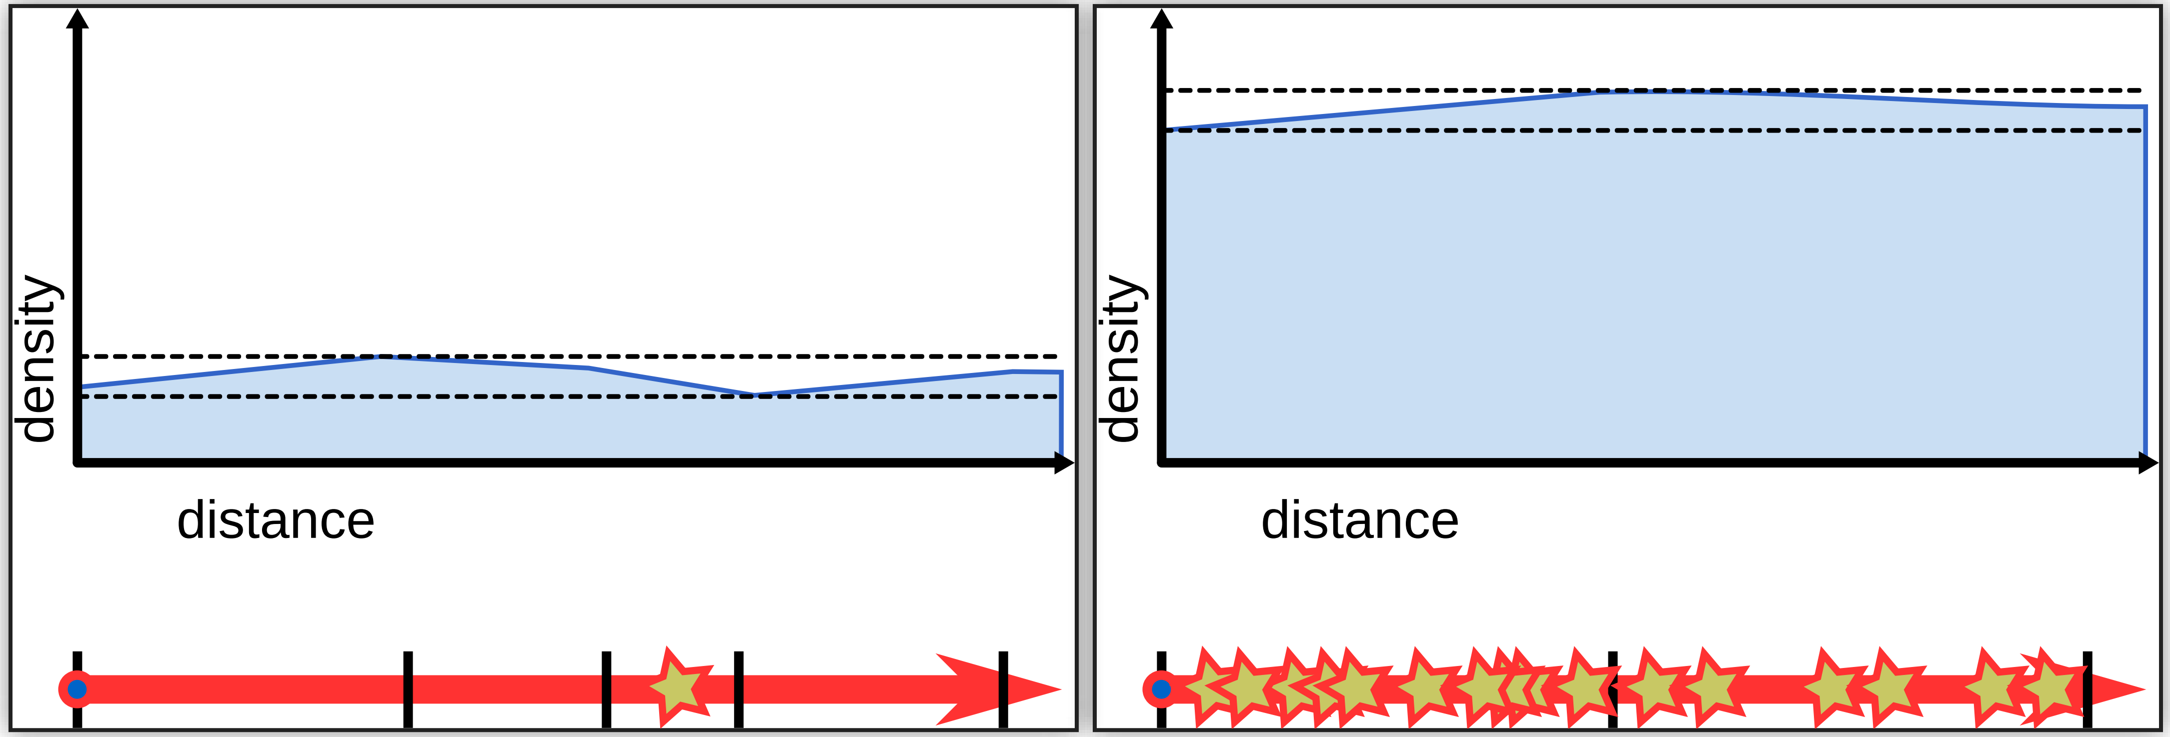

- accessing the memory within the same voxel is still expensive
- alleviated by decomposition tracking [KHLN17]
  - separate μ into sum of coarse and fine, to sample distance pick shortest (and early out!)
  - also profits full regular tracking

emissive media

thin/dense media make a difference

- no event inside the medium means we cannot pick up emission:

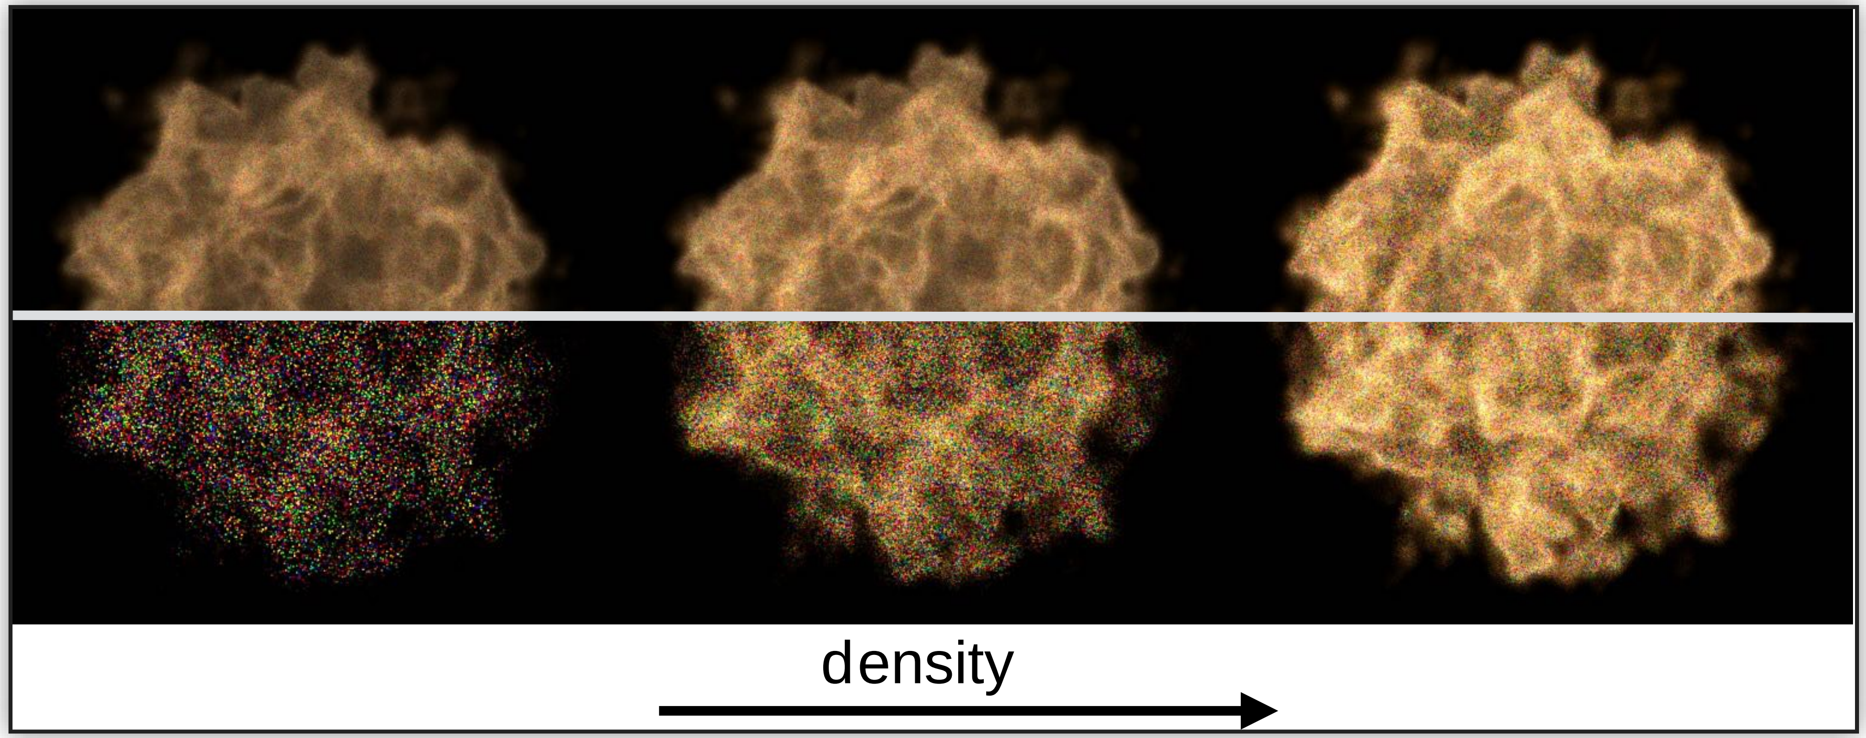

emissive media

thin/dense media make a difference

- following the idea of beams, collect emission along a ray

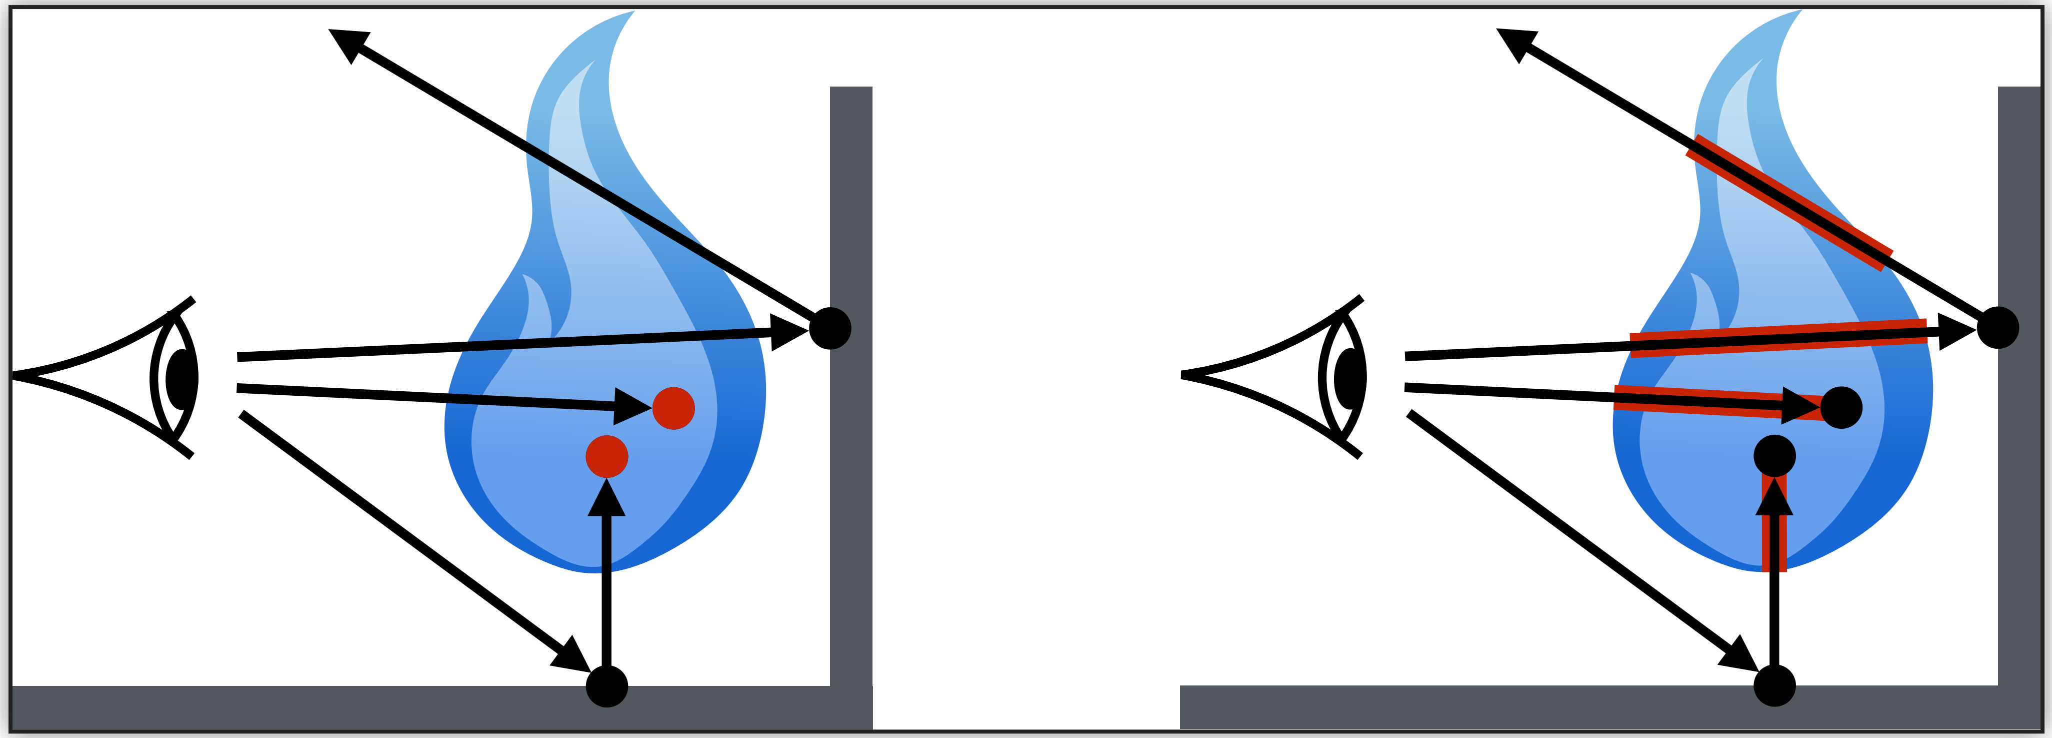

- particularly well suited for regular tracking, touching all voxels anyways

emissive media

thin/dense media make a difference

- direct application of MIS with NEE [VH13] **introduces noise**:

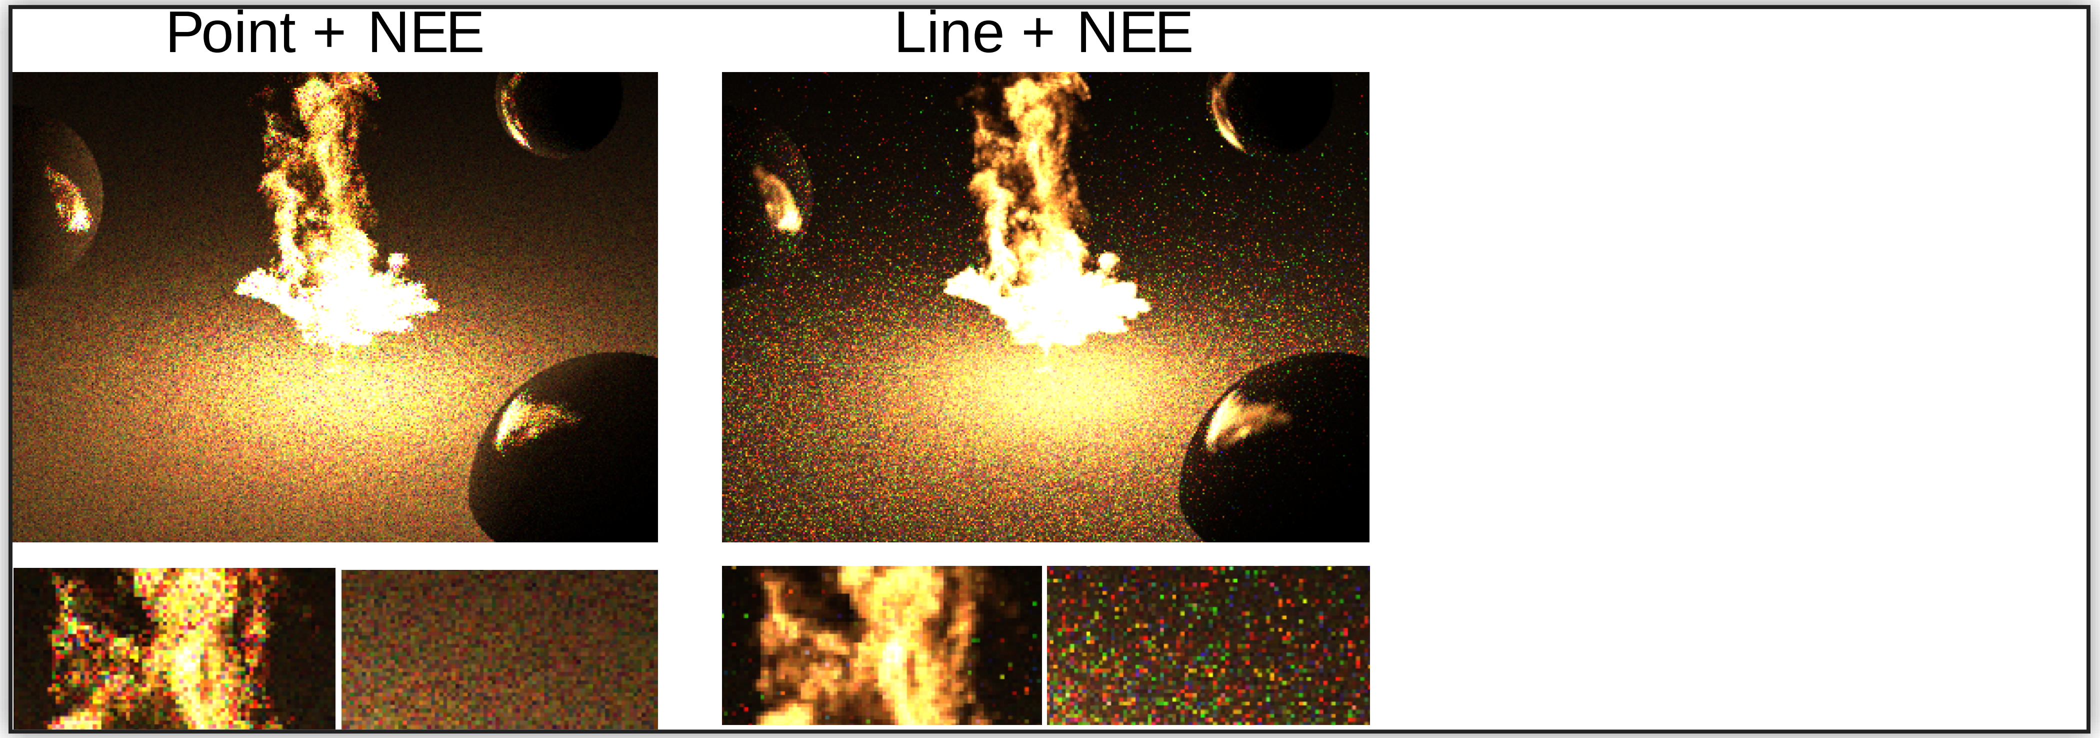

- reason: NEE cannot create paths with end point outside the medium
  - forward scattering PDF is poor, however, and now it picks up line emision!

emissive media

thin/dense media make a difference

- need to teach next event estimation about line emission [SHZD17]:

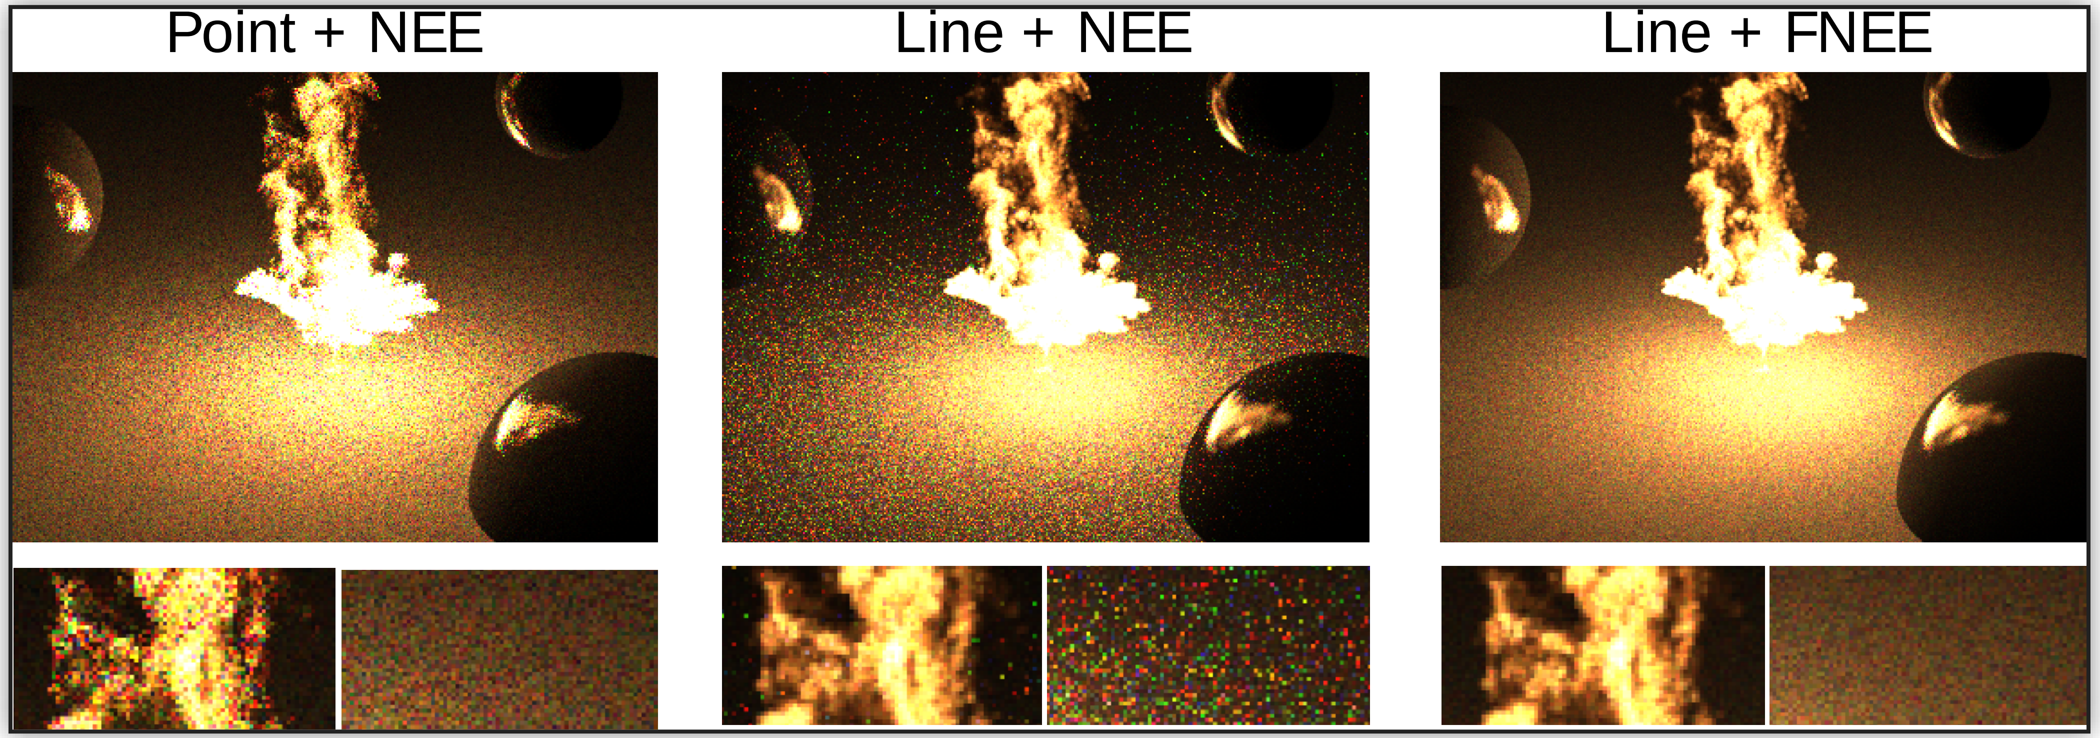

end of content

up next:

summary and open research problems

summary

free flight distance sampling

- woodcock/delta tracking

transmittance estimation

- track-length
- residual ratio
- free flight versions

path sampling

- path space formulation
- summary of advanced methods

acceleration structures

- for regular tracking
- for null collisions (bottom-level)

open research problems

null collision algorithms and MIS

- missing link to integrate into powerful framework
  - for instance combine with equi-angular sampling
- can we estimate the PDF?
  - expectation and division do not commute!

$$X = \frac{f(\bar{\mathbf{x}})}{p(\bar{\mathbf{x}})}$$

open research problems

leverage recent advances in machine learning

- special purpose denoising
  - including a volume prior?
- path guiding for volumes?
  - importance sampling for multiple vertices?

open research problems

joint handling of surfaces and geometry

- still often surface transport is handled separately
  - makes inclusion of all interreflections hard
  - custom-cut algorithms increase maintenance cost

open research problems

generalisation to correlated scatterers

- core assumption of exponential path length: uncorrelated particles!
  - particle repulsion such as in cell growth is very correlated
  - really, no collision can be found inside the current particle (min distance)
  - some existing work

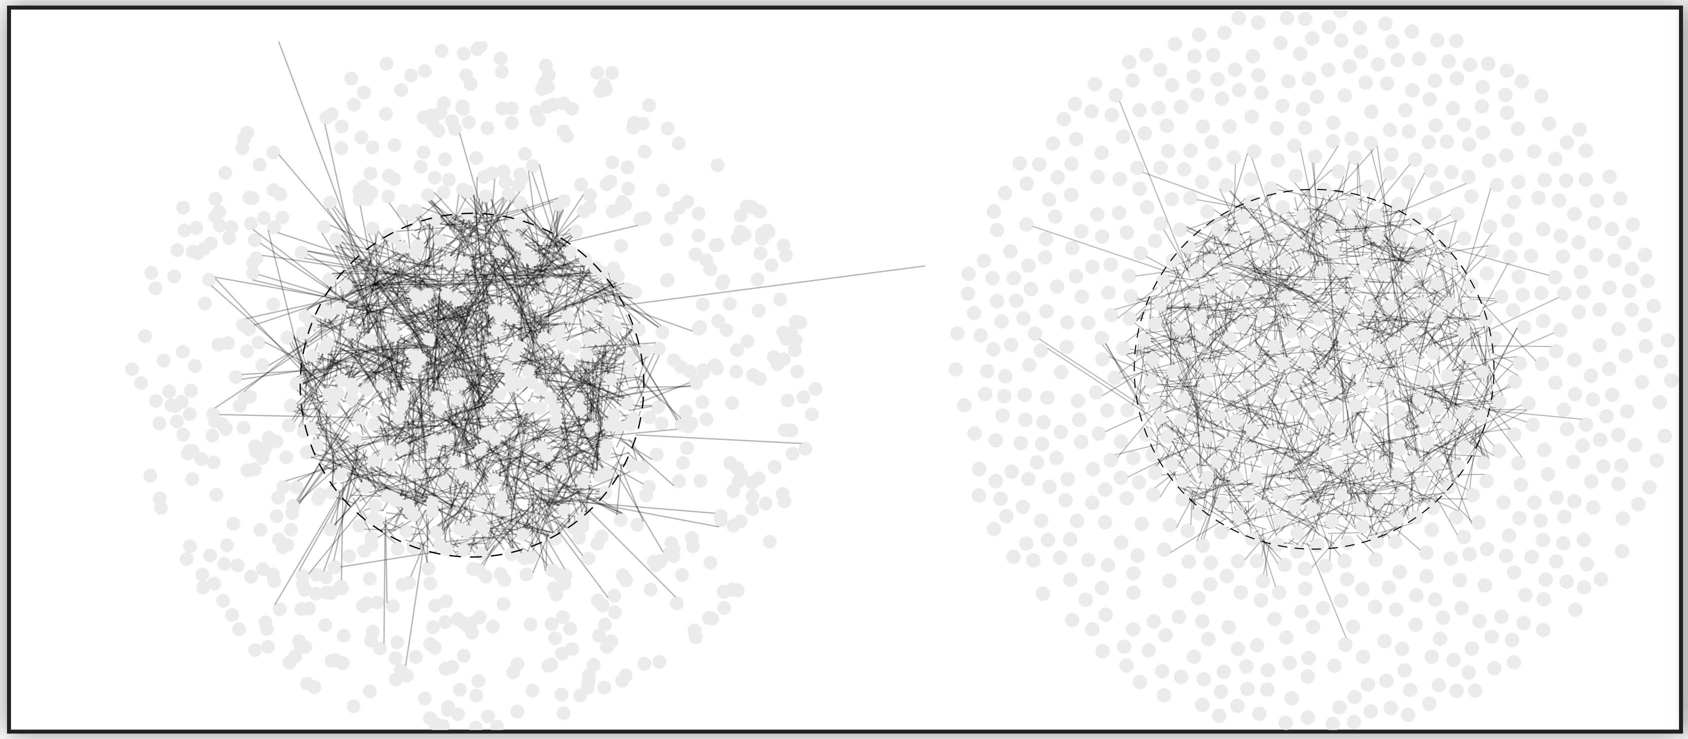

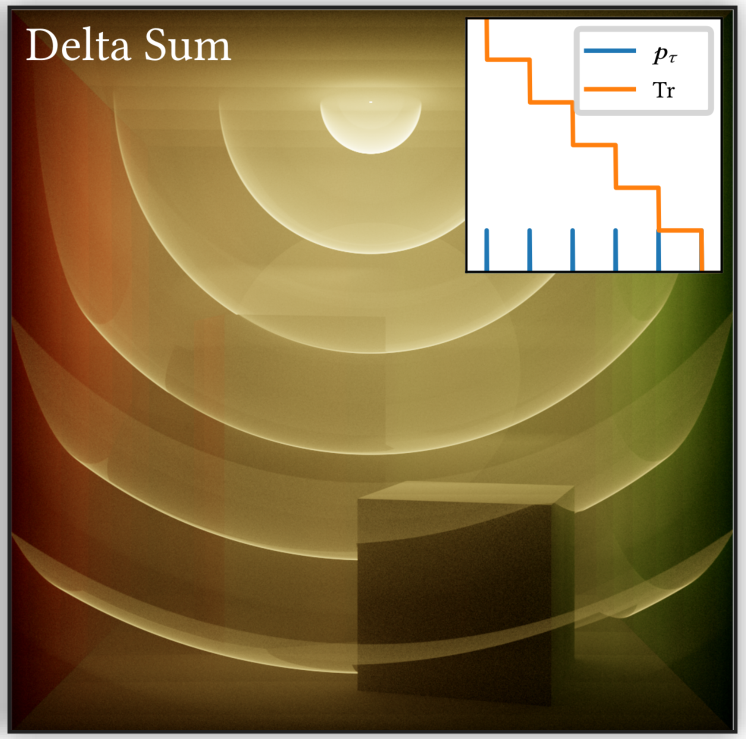

[d'Eon 2018, Bitterli et al. 2018]

thank you!

any questions?

acknowledgements:

- Peter Kutz
  - for tracing down many of early delta tracking papers
- Jaroslav Křivánek & reviewers
  - for feedback on the paper draft
- Maurizio Nitti
  - for help w/ fast forward and illustrations