PATH GUIDING BY MACHINE LEARNING

Jaroslav Křivánek
Charles University – Render Legion | Chaos Group

LIGHT TRANSPORT

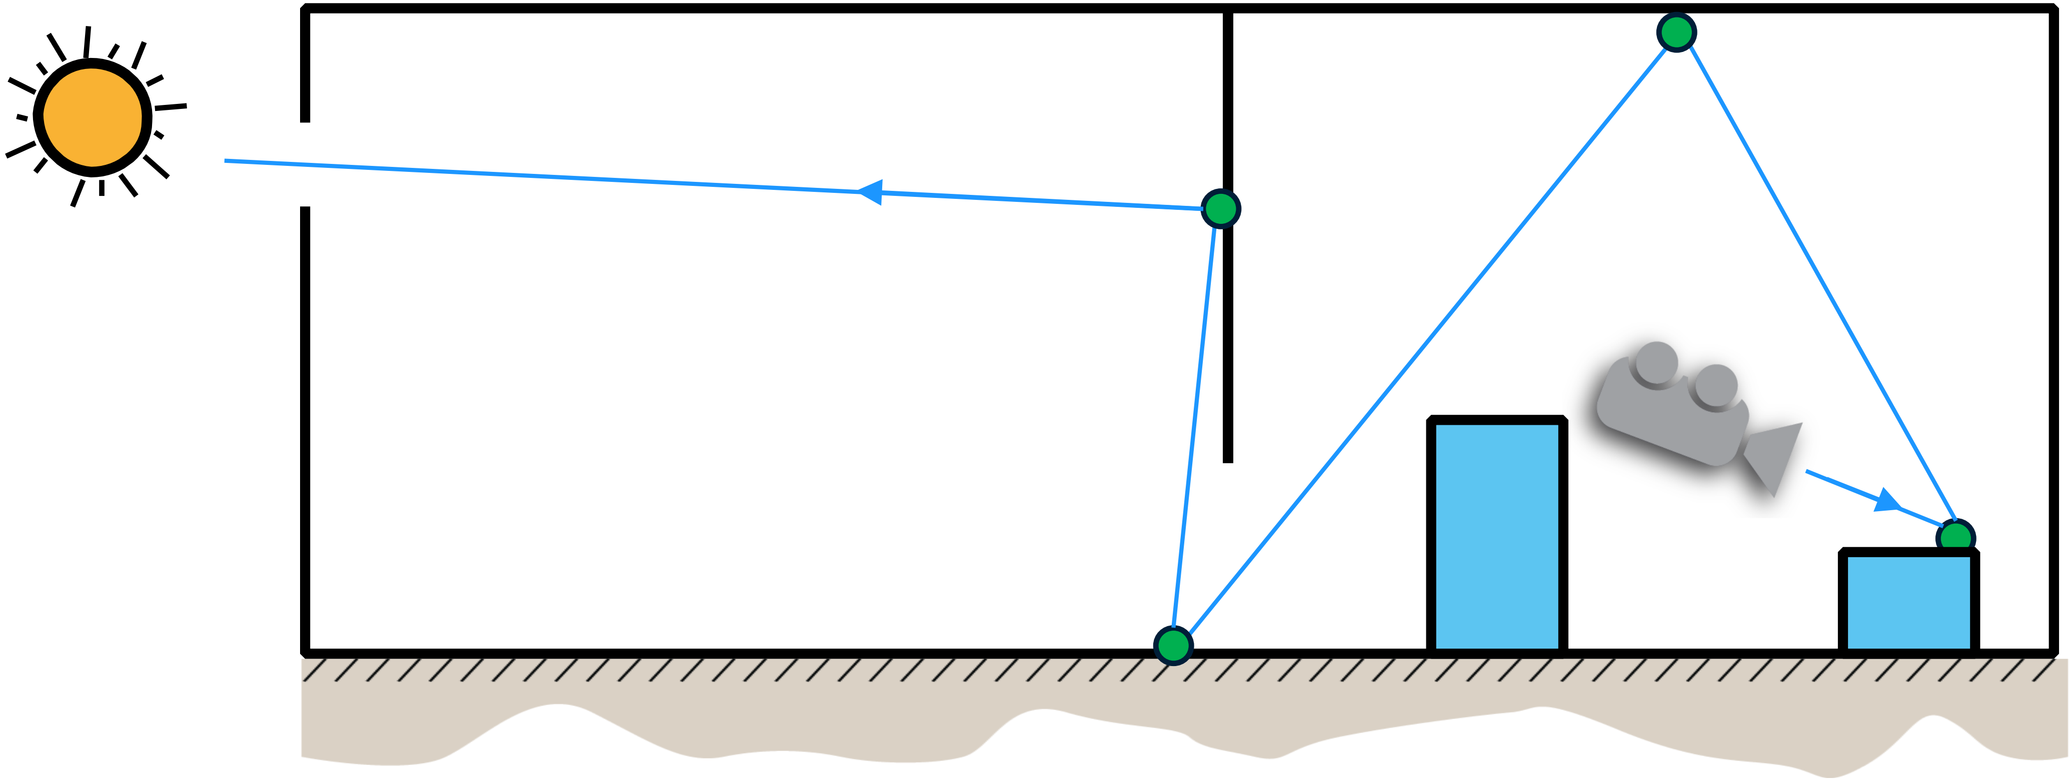

TODAYS' RENDERING IS OLD NEWS

- From **Matt Pharr**'s editorial to ACM TOG special issue on production rendering [Pharr 2018]:

> ***"Today … renderers are … based on … path tracing. Introduced … by Jim Kajiya (1986)."***
>
> *"Many advancements were made … including*
> *- more effective light sampling algorithms (Shirley et al. **1996**),*
> *- high-quality sampling patterns (Kollig and Keller **2002**), and*
> *- multiple importance sampling (Veach and Guibas **1995**),"*
>
> *"… the core ray tracing [got] more efficient (Wald et al. **2001**)."*

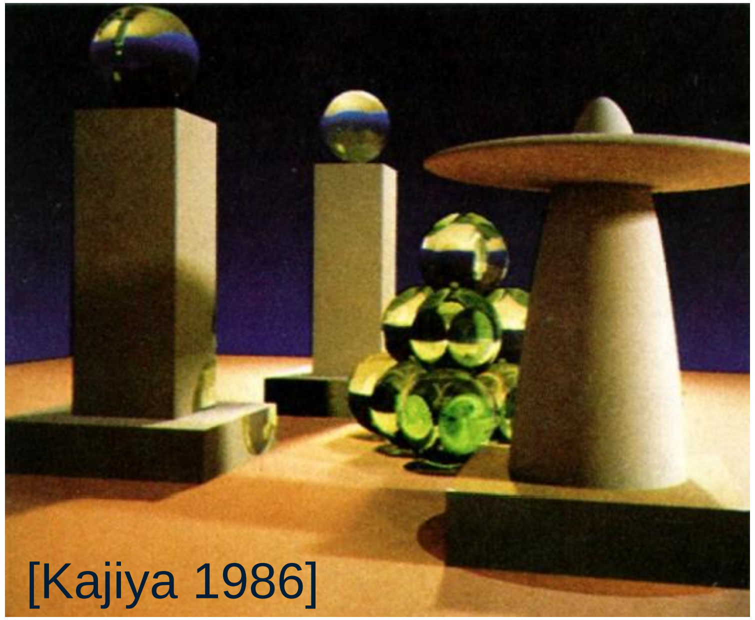

[Kajiya 1986]

ADVANCED LIGHT TRANSPORT

- Why are advanced light transport algorithms not used in practice?

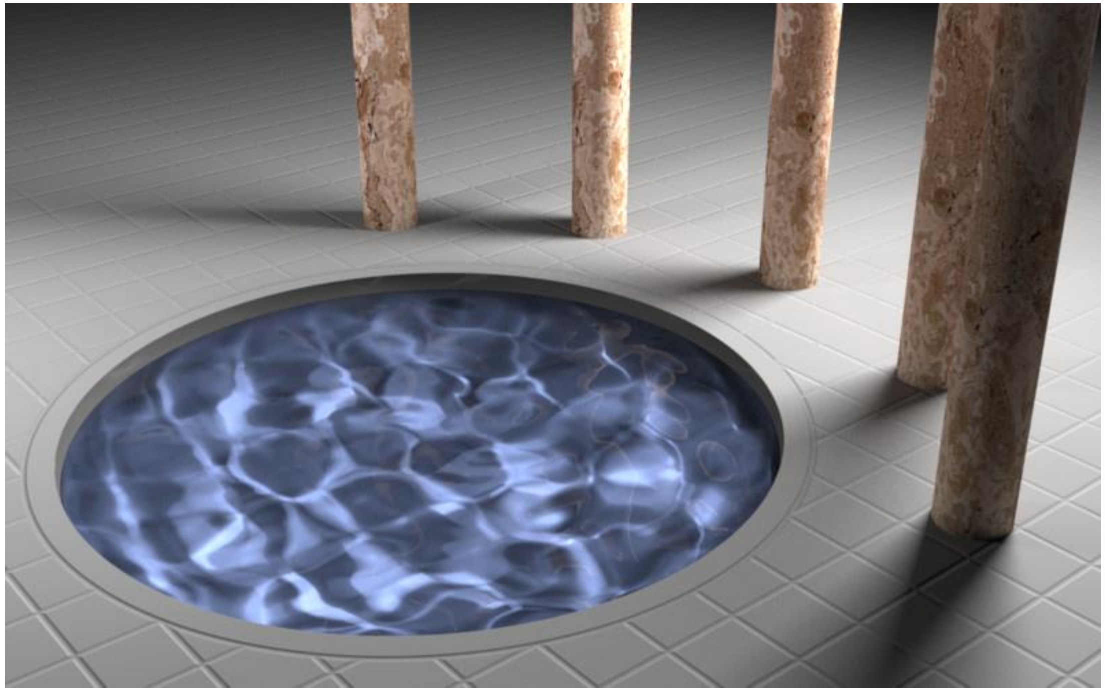

Metropolis Light Transport [Veach and Guibas 1997]

A GOOD LIGHT TRANSPORT ALGORITHM ...

- **… has to be**
  - Fast in common scenes
  - Robust & reliable
  - Easy-to-use (no parameters)
  - Interactive & progressive
  - …

THE GOOD ALGORITHM CHECKLIST

- Fast in common scenes
- Robust & reliable
- Easy-to-use (no parameters)
- Interactive & progressive

PATH TRACING

- [Kajiya 1986, Veach and Guibas 1995, Shirley 1996,…]

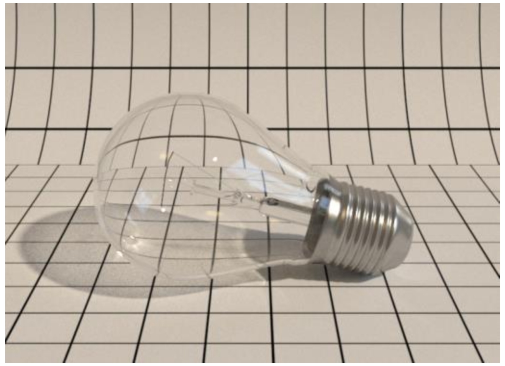

Reference rendering
(VCM)

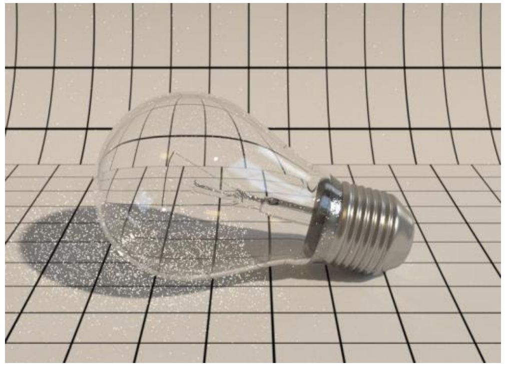

Path tracing
(no clamping)

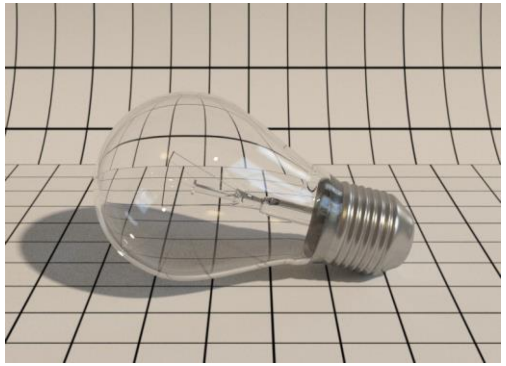

Path tracing
(with clamping)

- Fast in common scenes
- ~~Robust & reliable~~
- Easy-to-use (no parameters)
- Interactive & progressive

THE LIGHT TRANSPORT CHALLENGE

Algorithm that can renders this at least as fast as a path tracer…

… and it can also render this.

BIDIR / VCM

BIDIR / VCM

- [Lafortune and Willems 1993, Veach and Guibas 1995]
- [Georgiev et al. 2012, Hachisuka et al. 2012]

- VCM = Photon mapping + Bidir
- **“Brute-force robustness”** – Overhead

- ~~**Fast in common scenes**~~
- **Robust & reliable**
- Easy-to-use (no parameters)
- Interactive & progressive

METROPOLIS LIGHT TRANSPORT

MLT + Manifold exploration [Jakob and Marschner 2012]

Reference

METROPOLIS LIGHT TRANSPORT

- [Veach and Guibas 1997, …]

- Uneven convergence, temporal instability

- ~~**Fast in common scenes**~~
- ~~**Robust & reliable**~~
- Easy-to-use (no parameters)
- Interactive & progressive

DESIGNING THE ULTIMATE PRACTICAL ALGORITHM

THE ULTIMATE LIGHT TRANSPORT ALGORITHM

- Start off from PT
  - because it ticks most of the boxes

- Fast in common scenes
- Robust & reliable
- Easy-to-use (no parameters)
- Interactive & progressive

- Address its problems

- Root of the problem: **lack of information** in sampling decisions

Path sampling in unidirectional path tracing

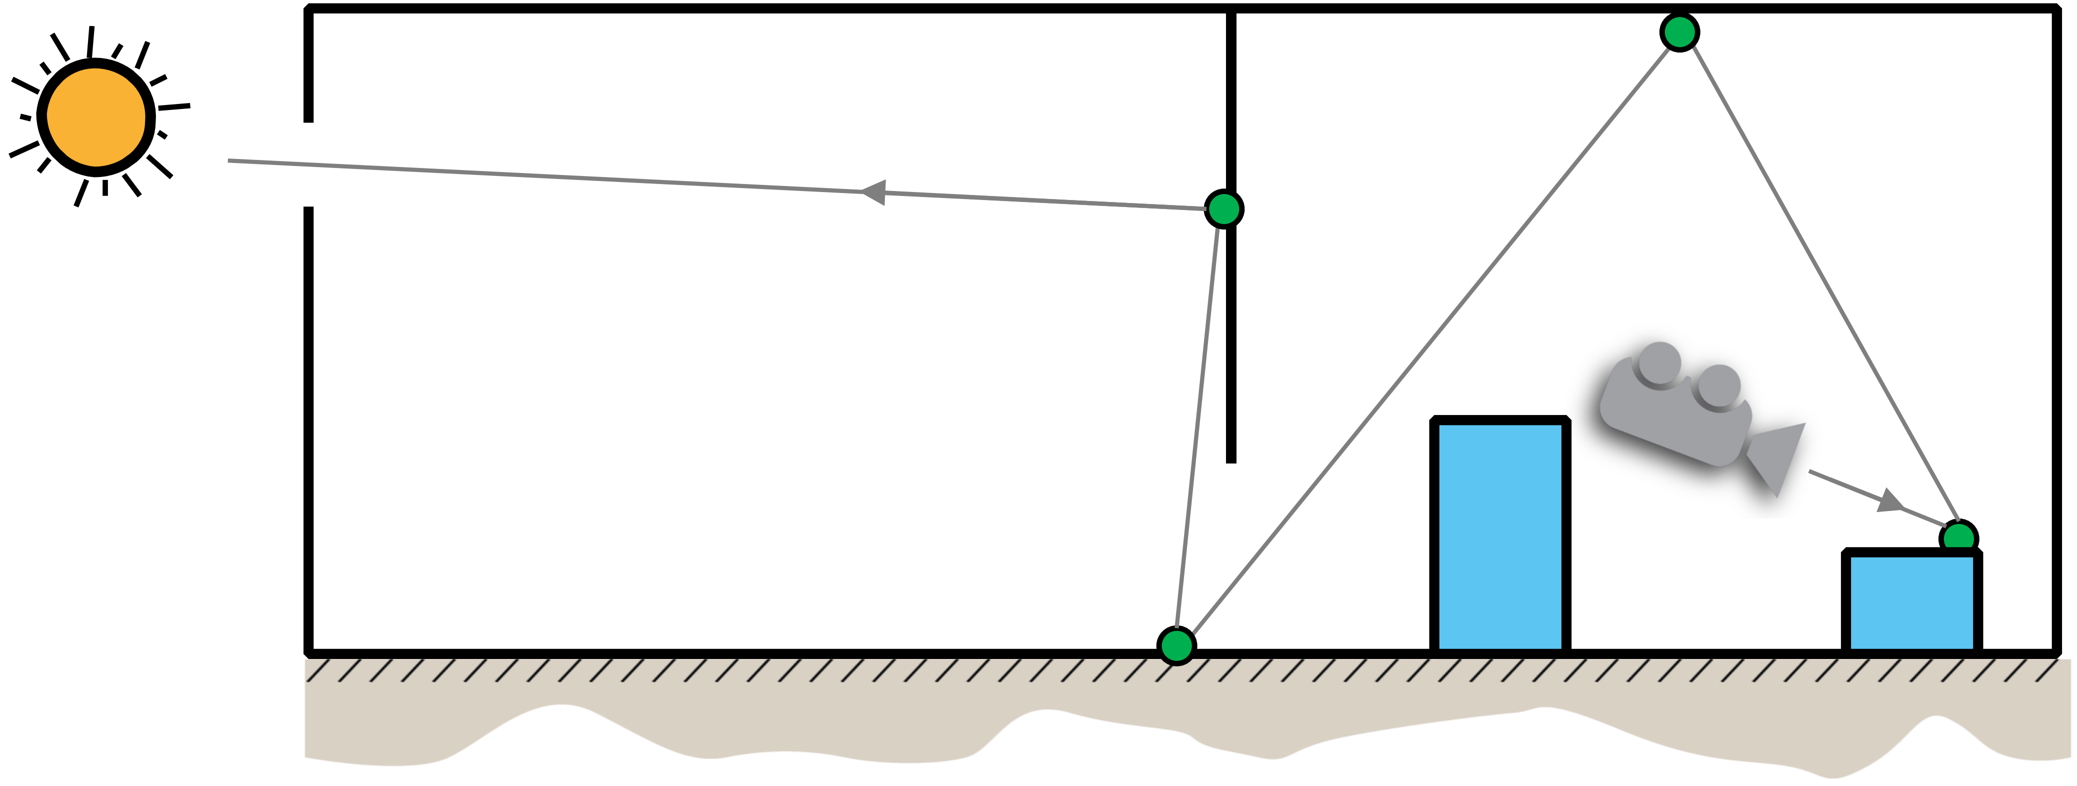

Directional sampling

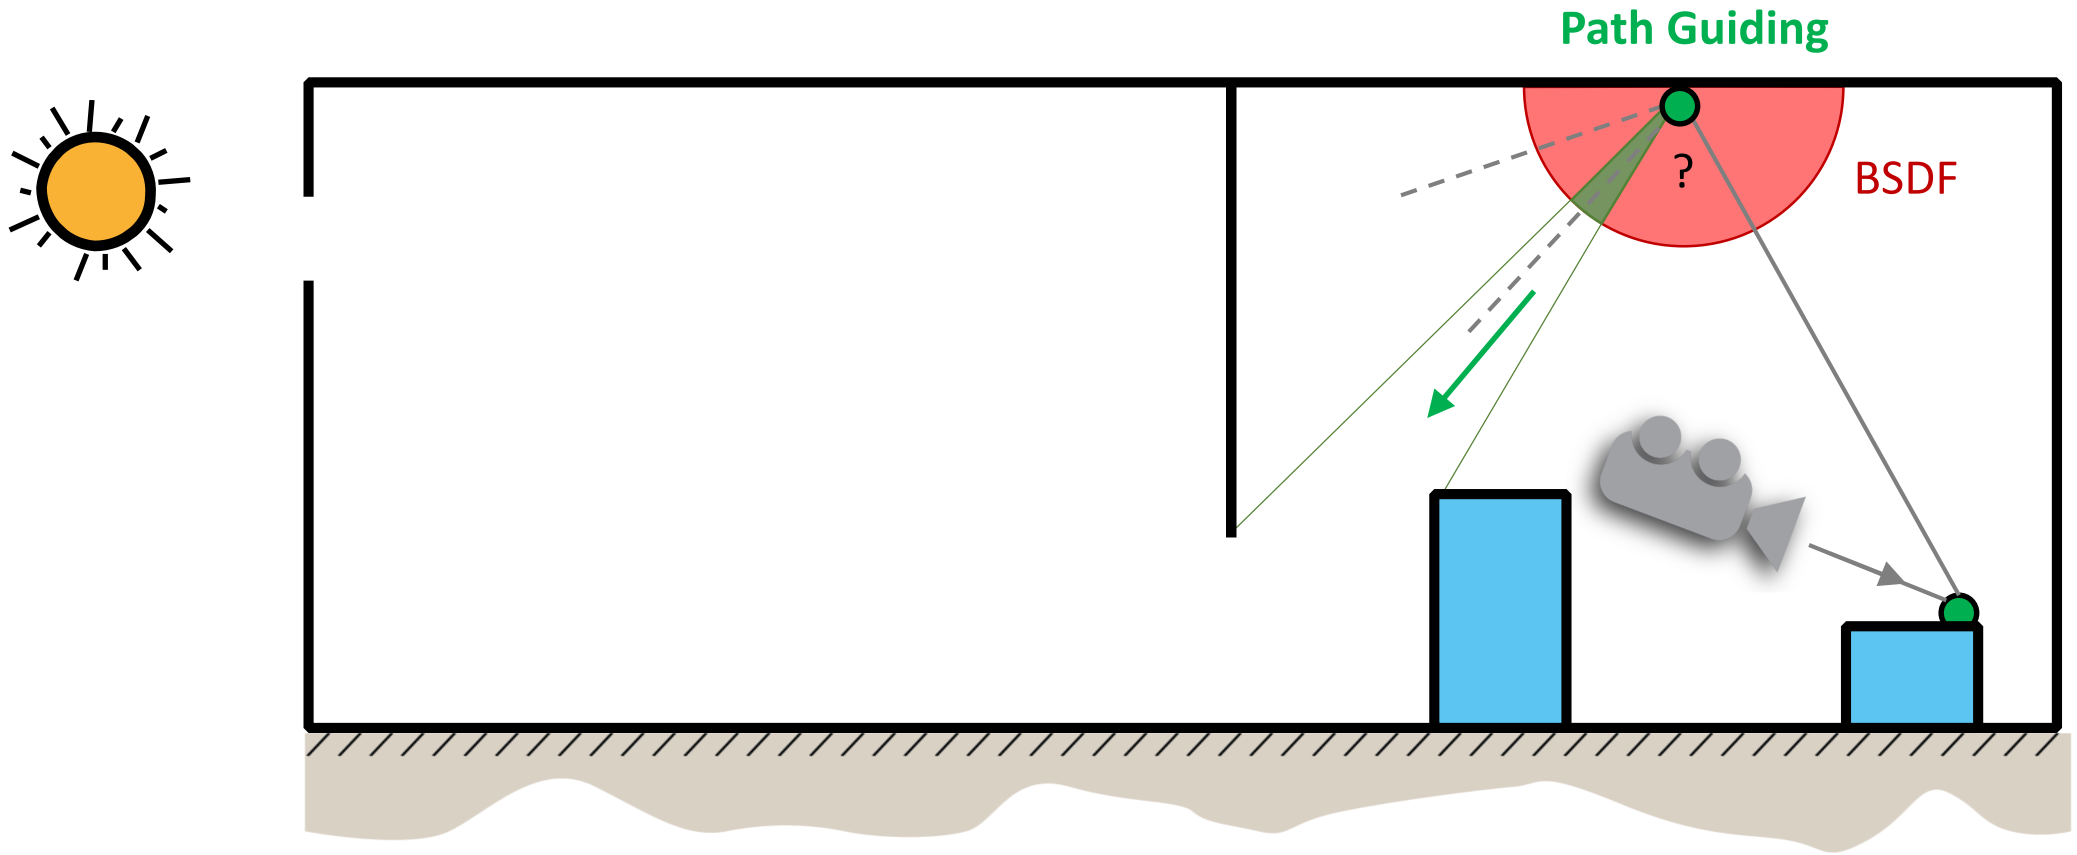

Path termination (Russian roulette)

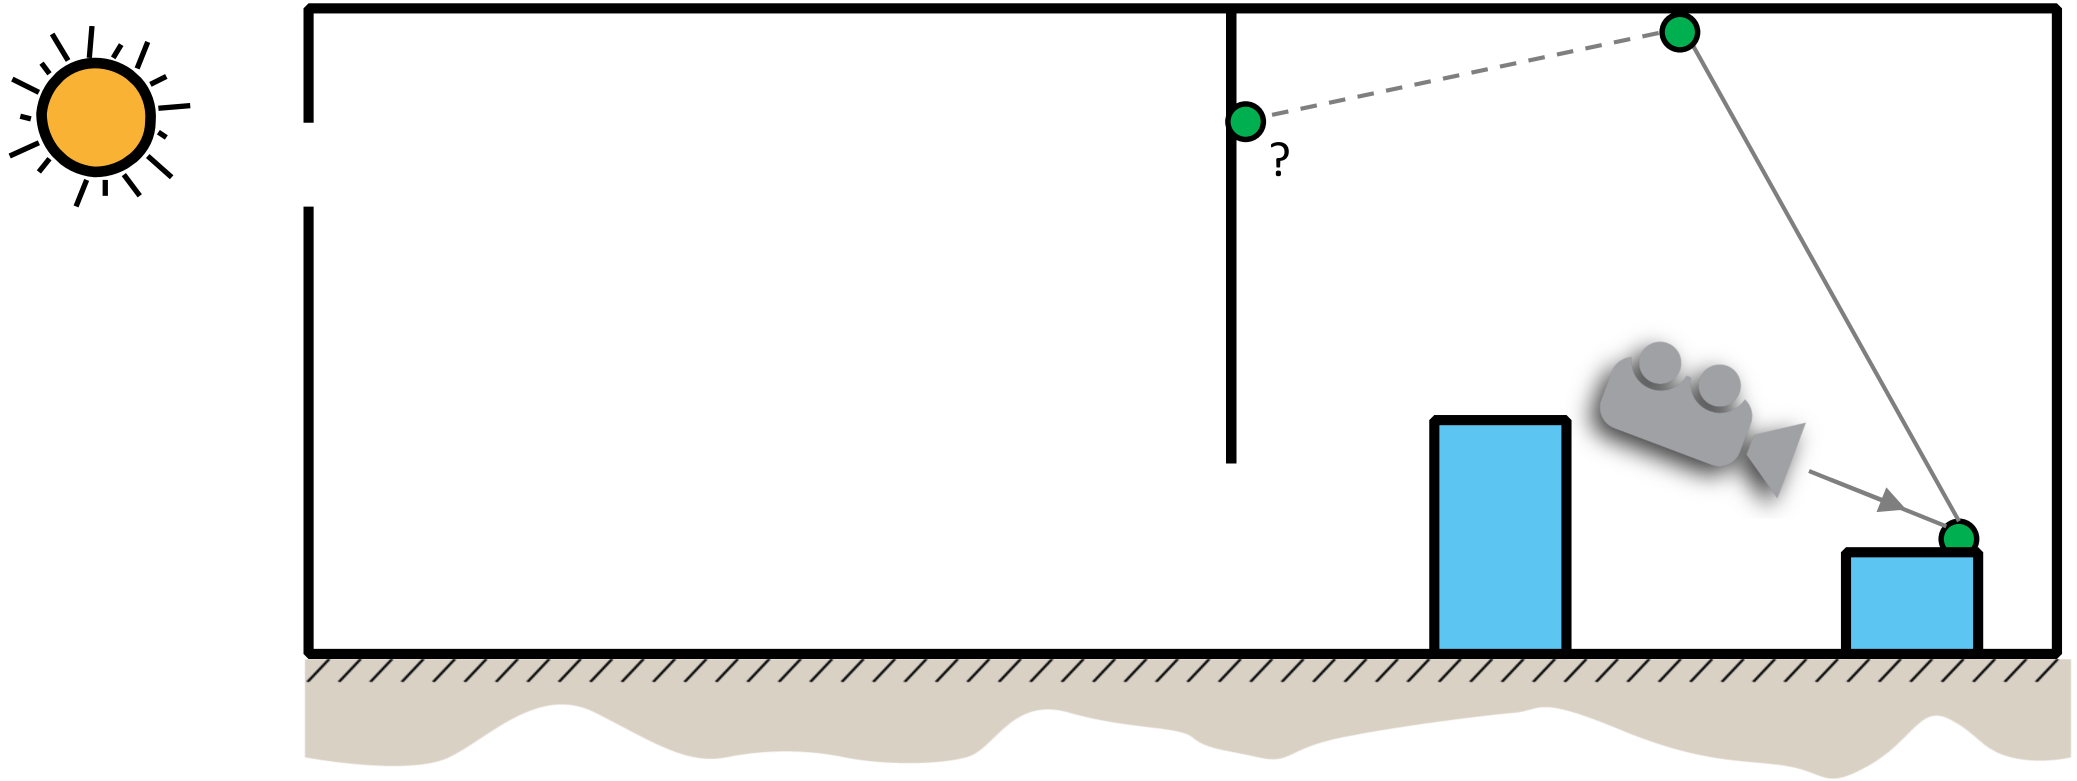

Path termination (Russian roulette)

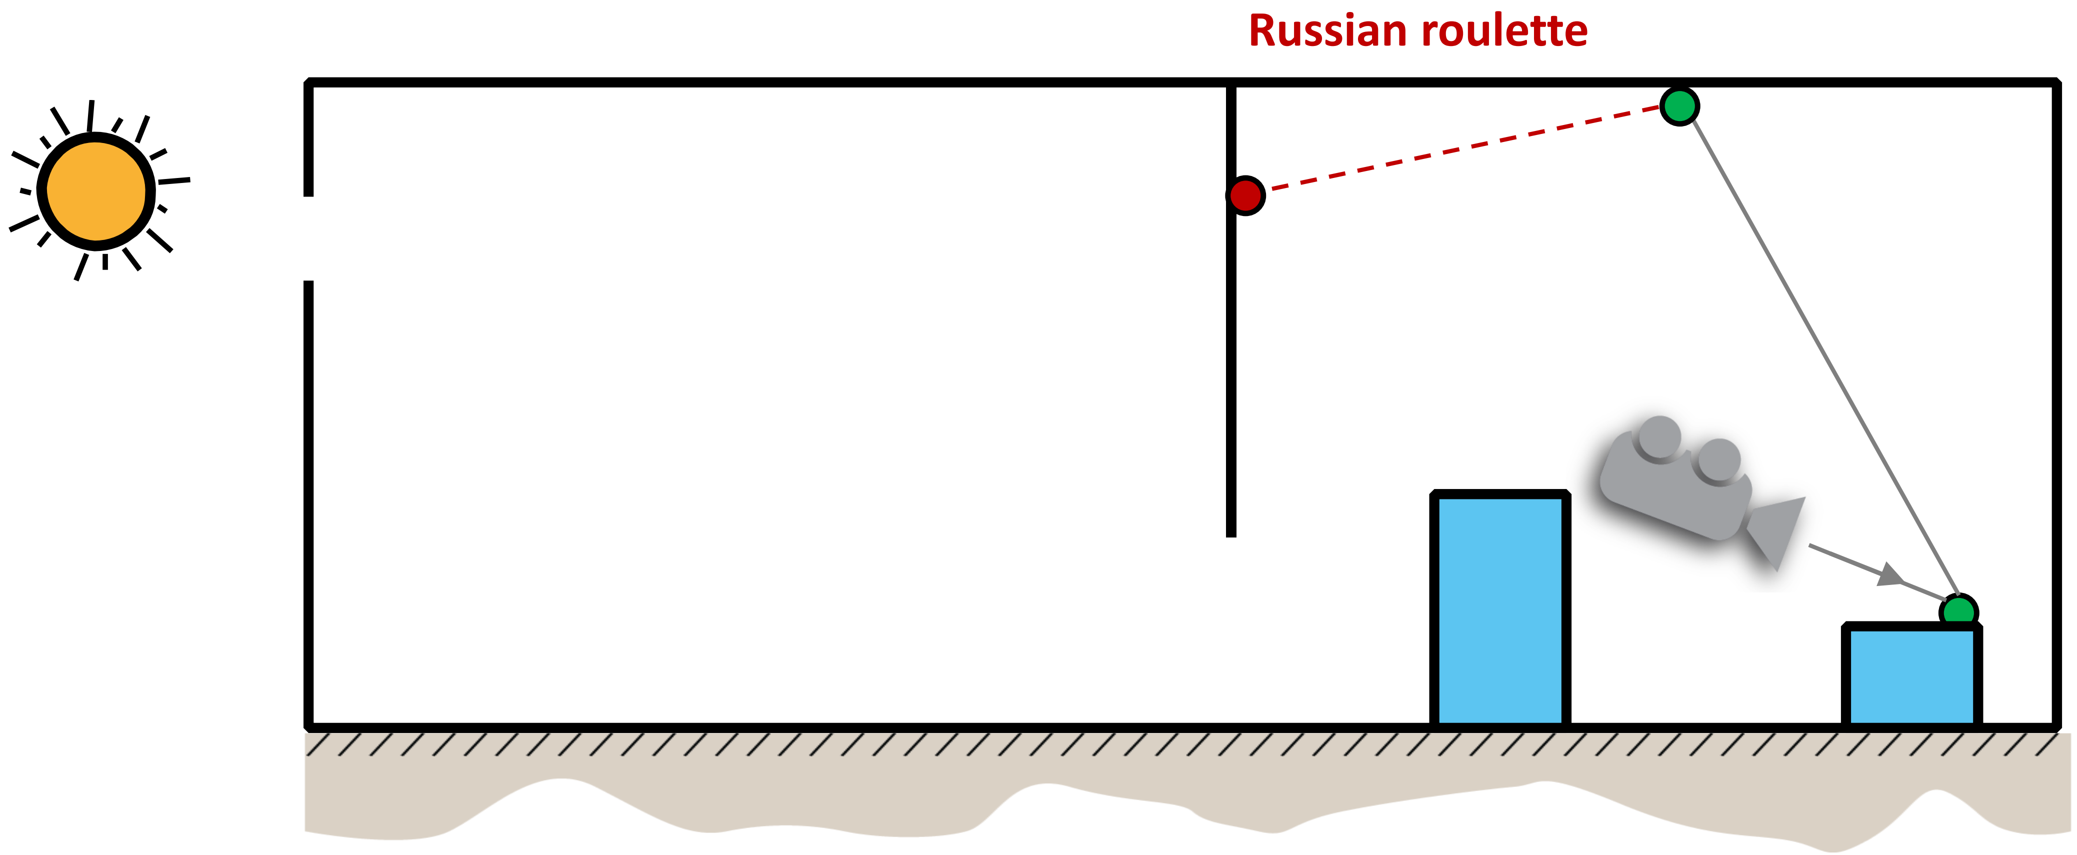

Splitting

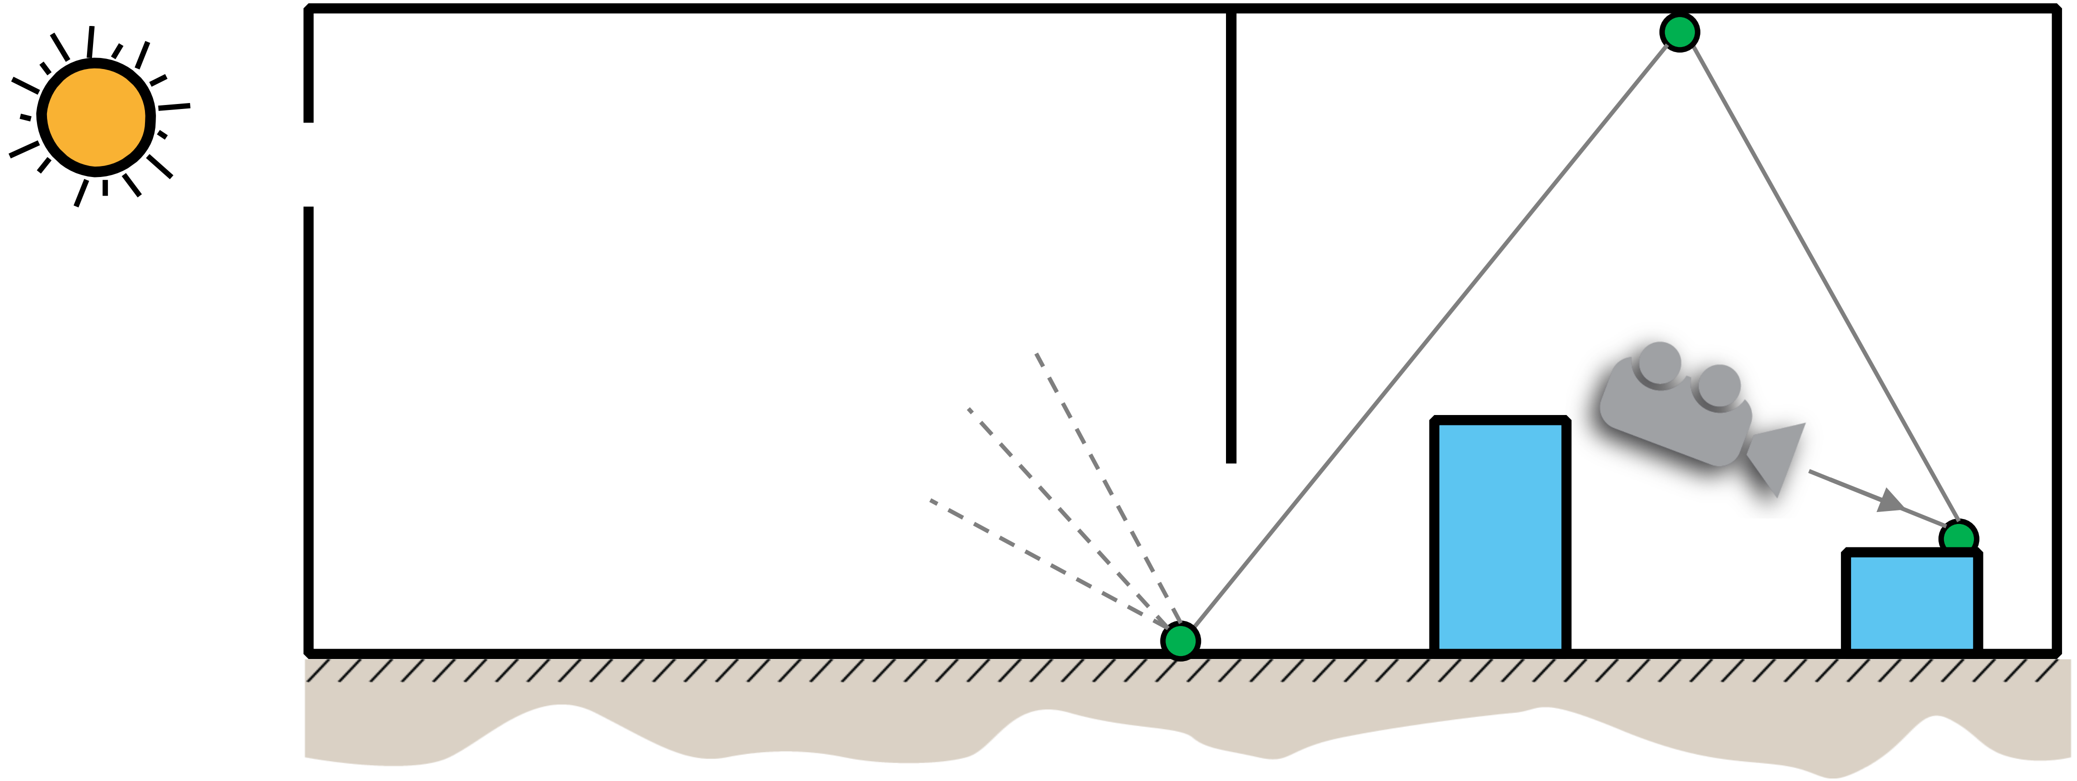

SOLUTION IDEA

- Give path tracing extra information
- Chicken-and-egg problem
- **Adaptive sampling**
- How to make it **robust** when there's so much **uncertainty**? – **Machine learning methods**

OVERVIEW

- Path guiding through **online mixture model training** [Vorba et al. 2014]
  - Guided Russian roulette and splitting [Vorba and Křivánek 2016]
  - Path guiding in volumes [TOG, conditionally accepted]
- Robust adaptive direct illumination through **online Bayesian regression** [Vévoda et al. 2018]

GENERATIONS / VANCOUVER 12-16 AUGUST
SIGGRAPH2018

PATH GUIDING

Vorba et al. – ACM SIGGRAPH 2014

Previous work

- Jensen *[1995]*

photon tracing

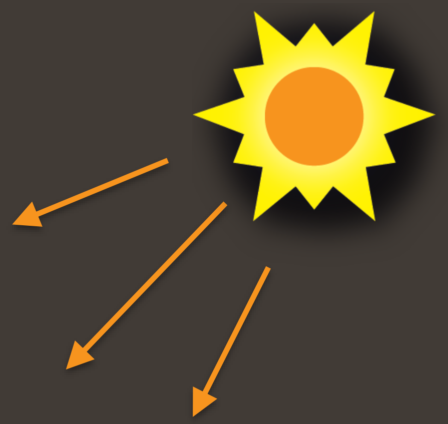

Previous work

- Jensen *[1995]*

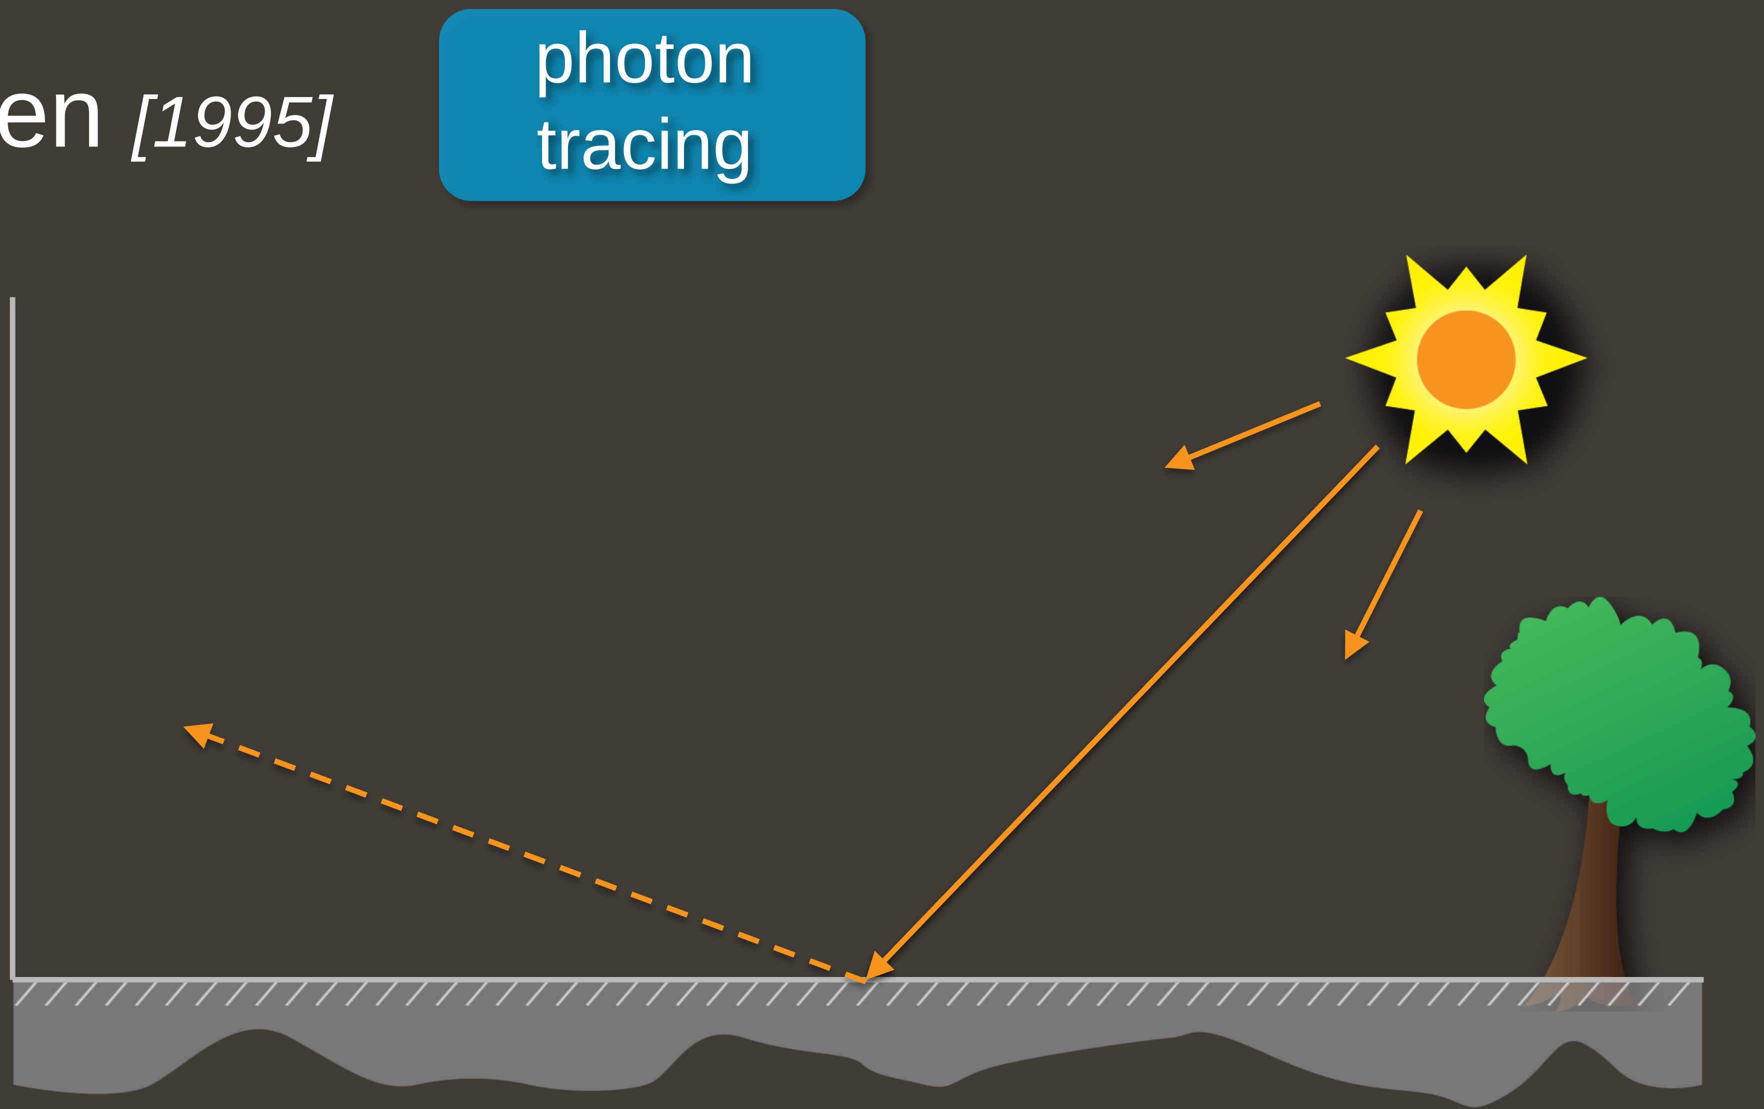

Previous work

- Jensen *[1995]*

photon tracing

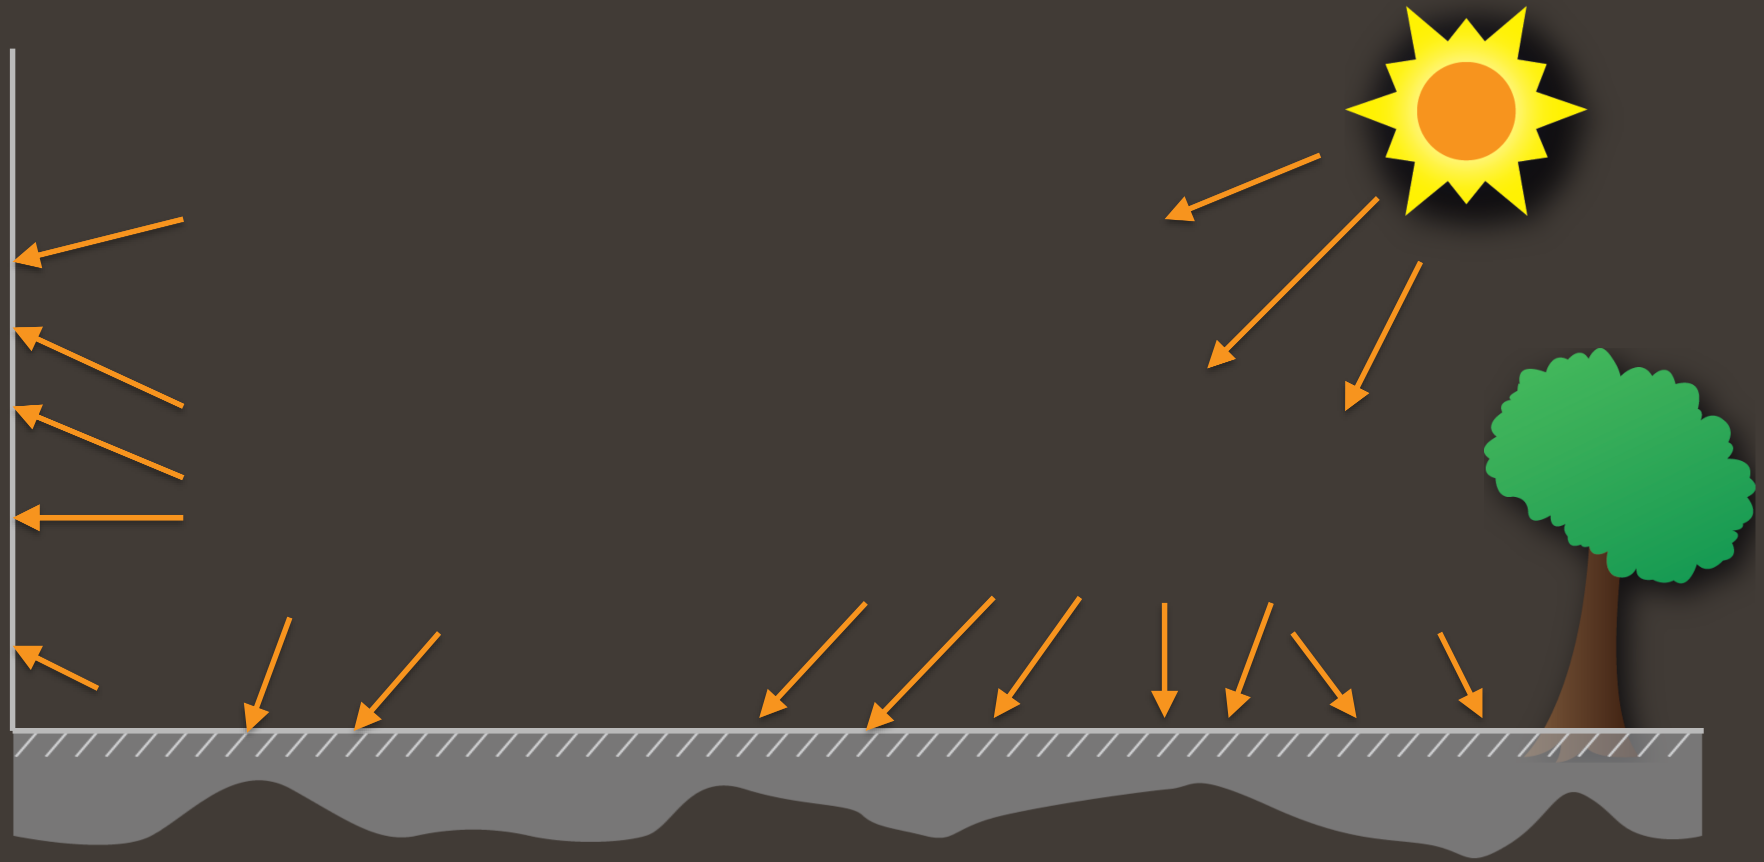

Previous work

- Jensen *[1995]*

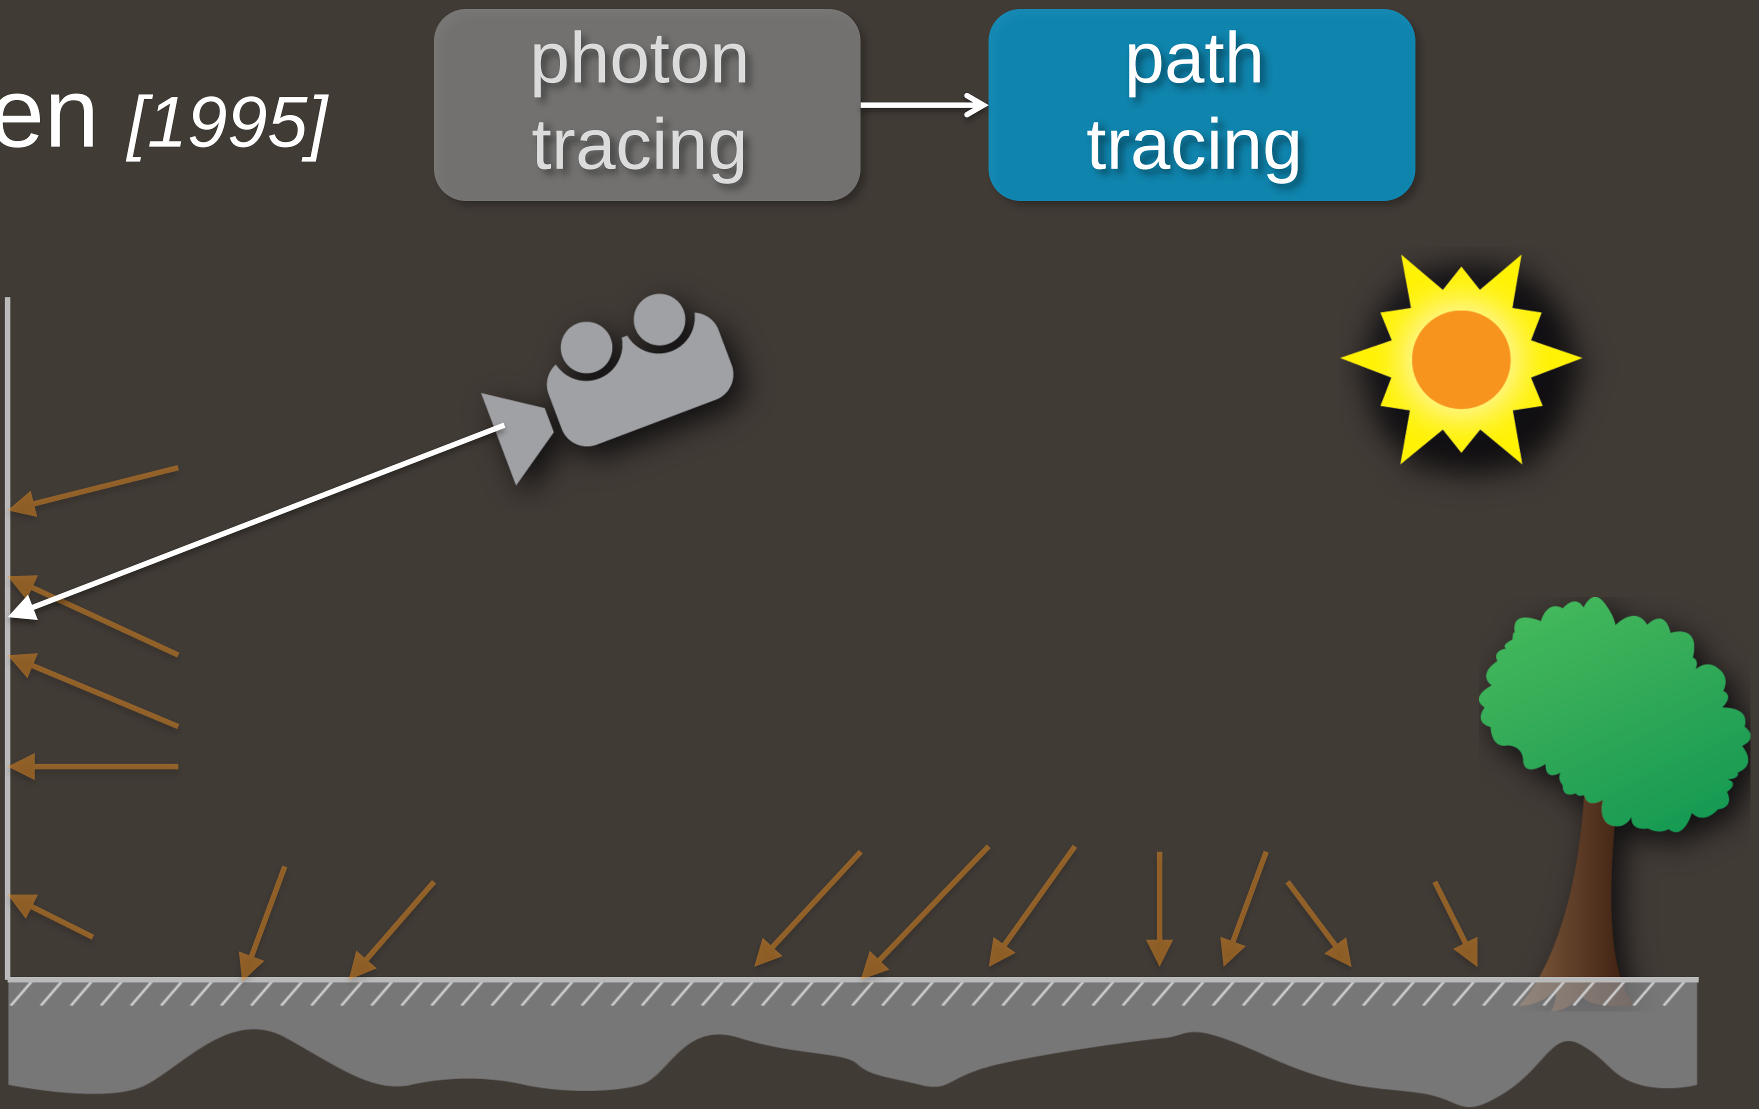

Previous work

- Jensen *[1995]*

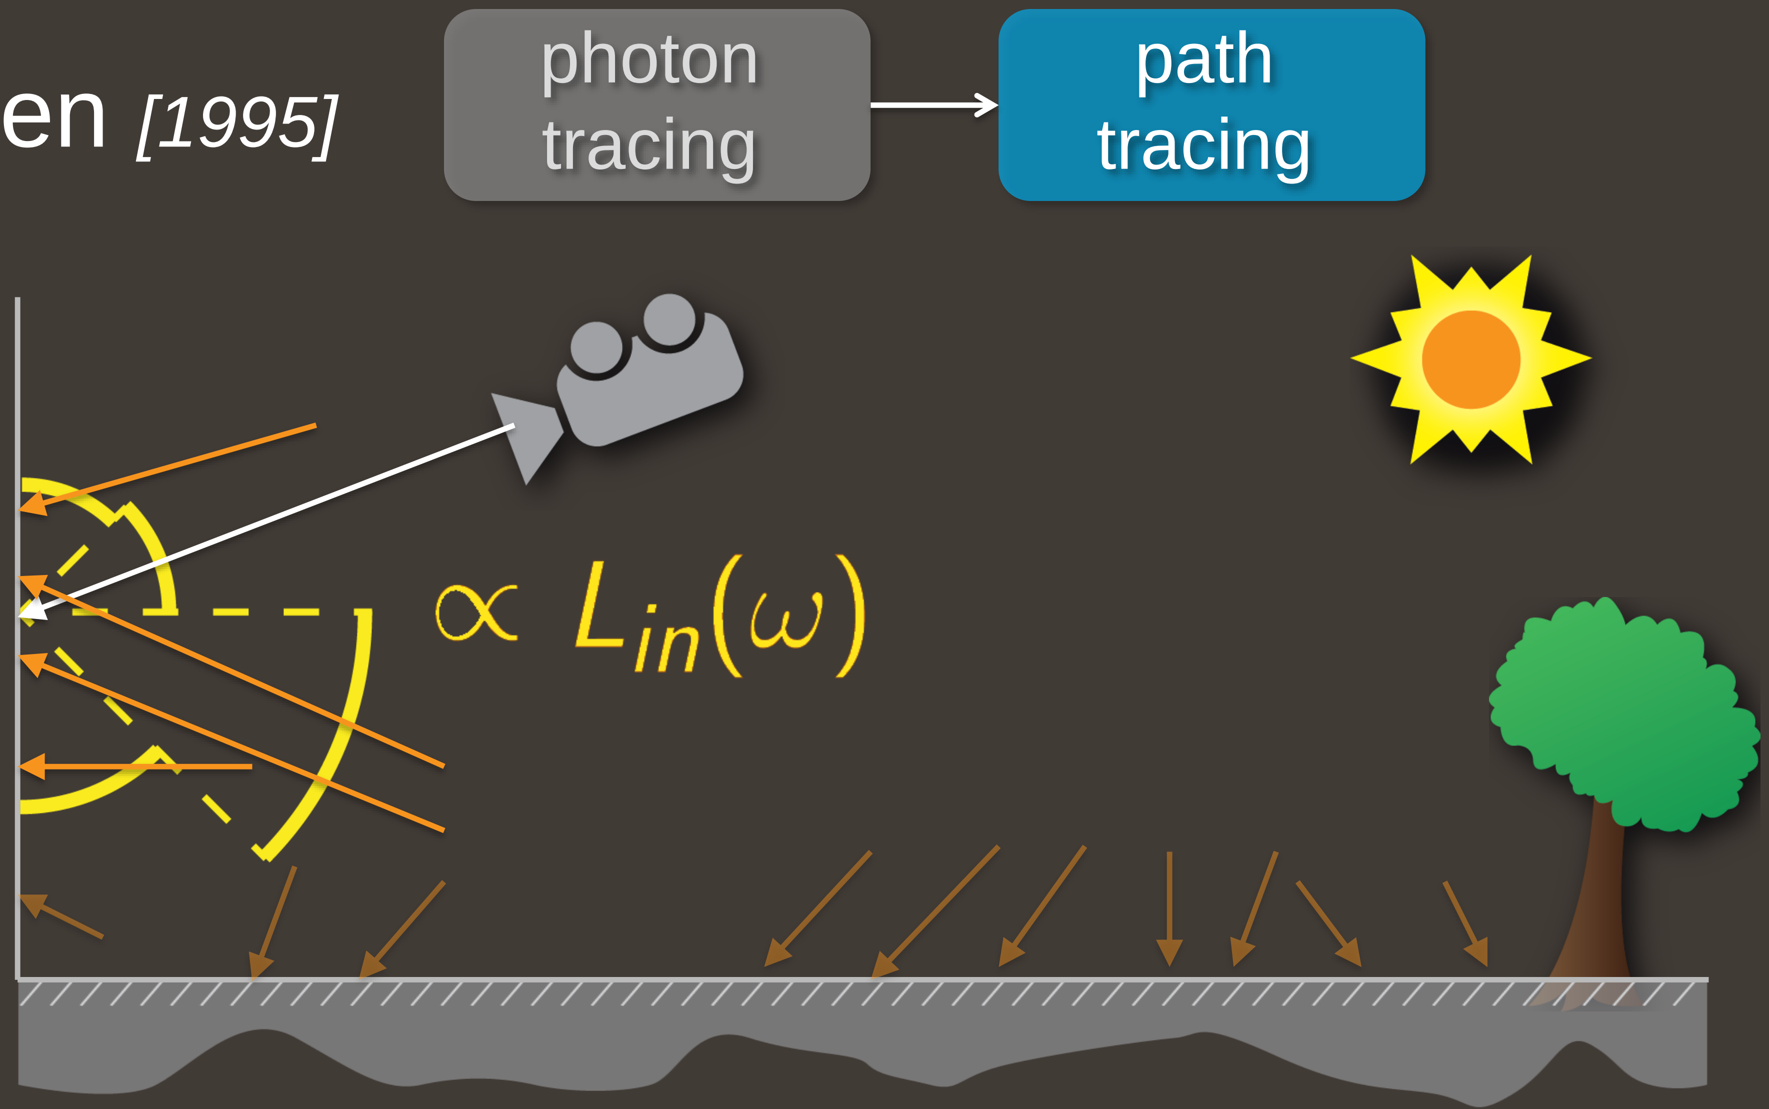

Previous work

- Jensen *[1995]*

photon tracing → path tracing

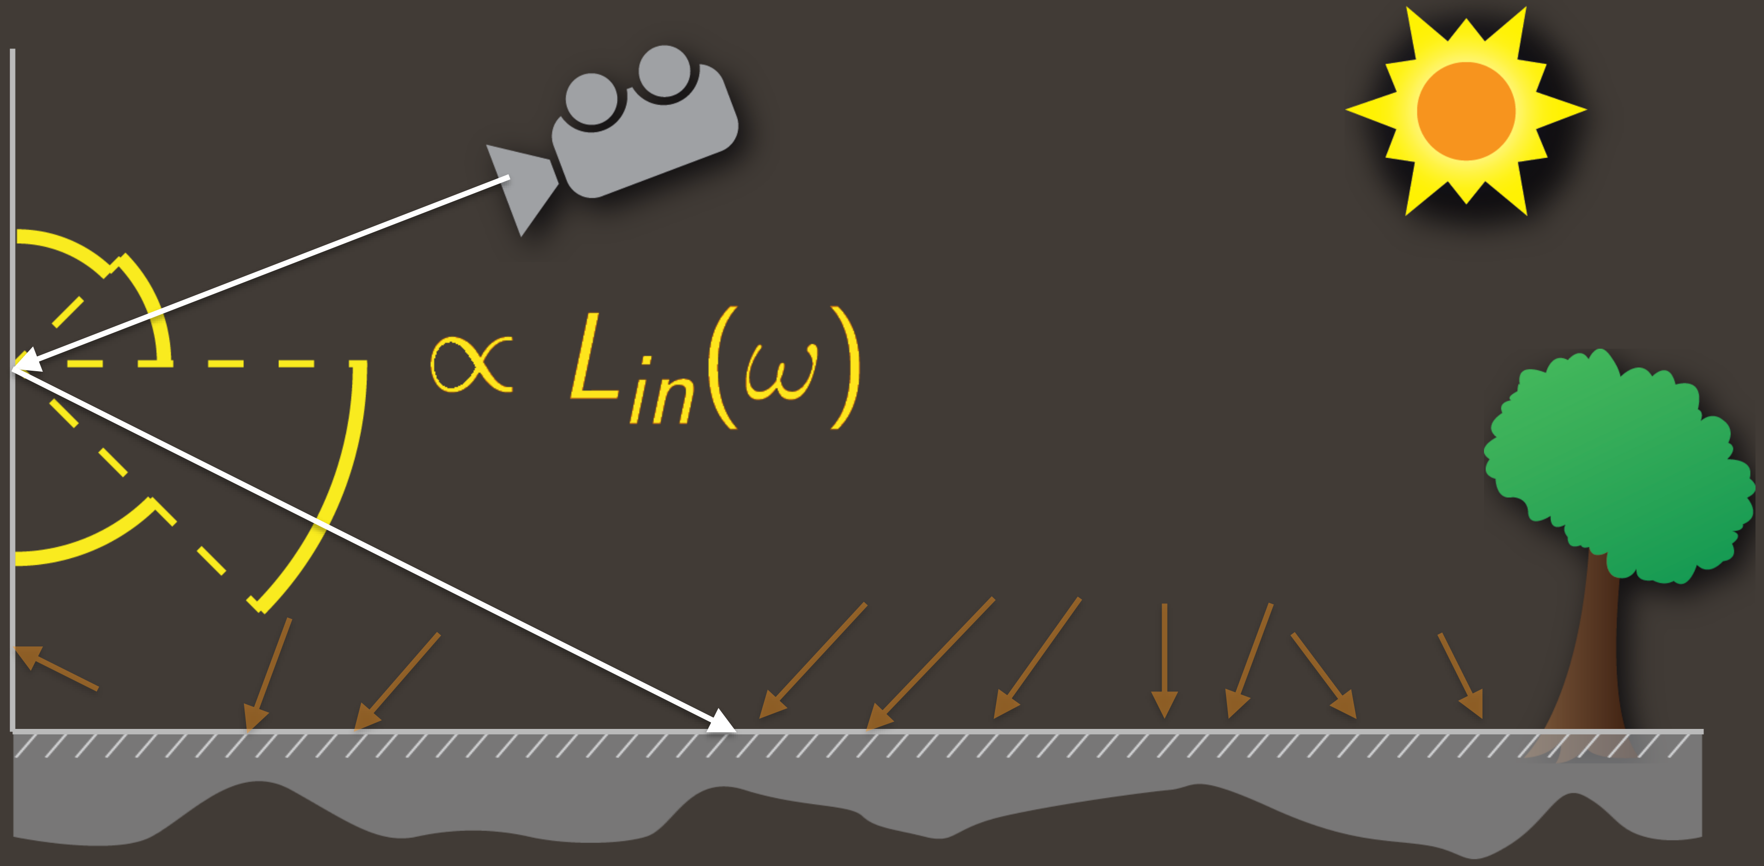

Previous work

- Jensen *[1995]*

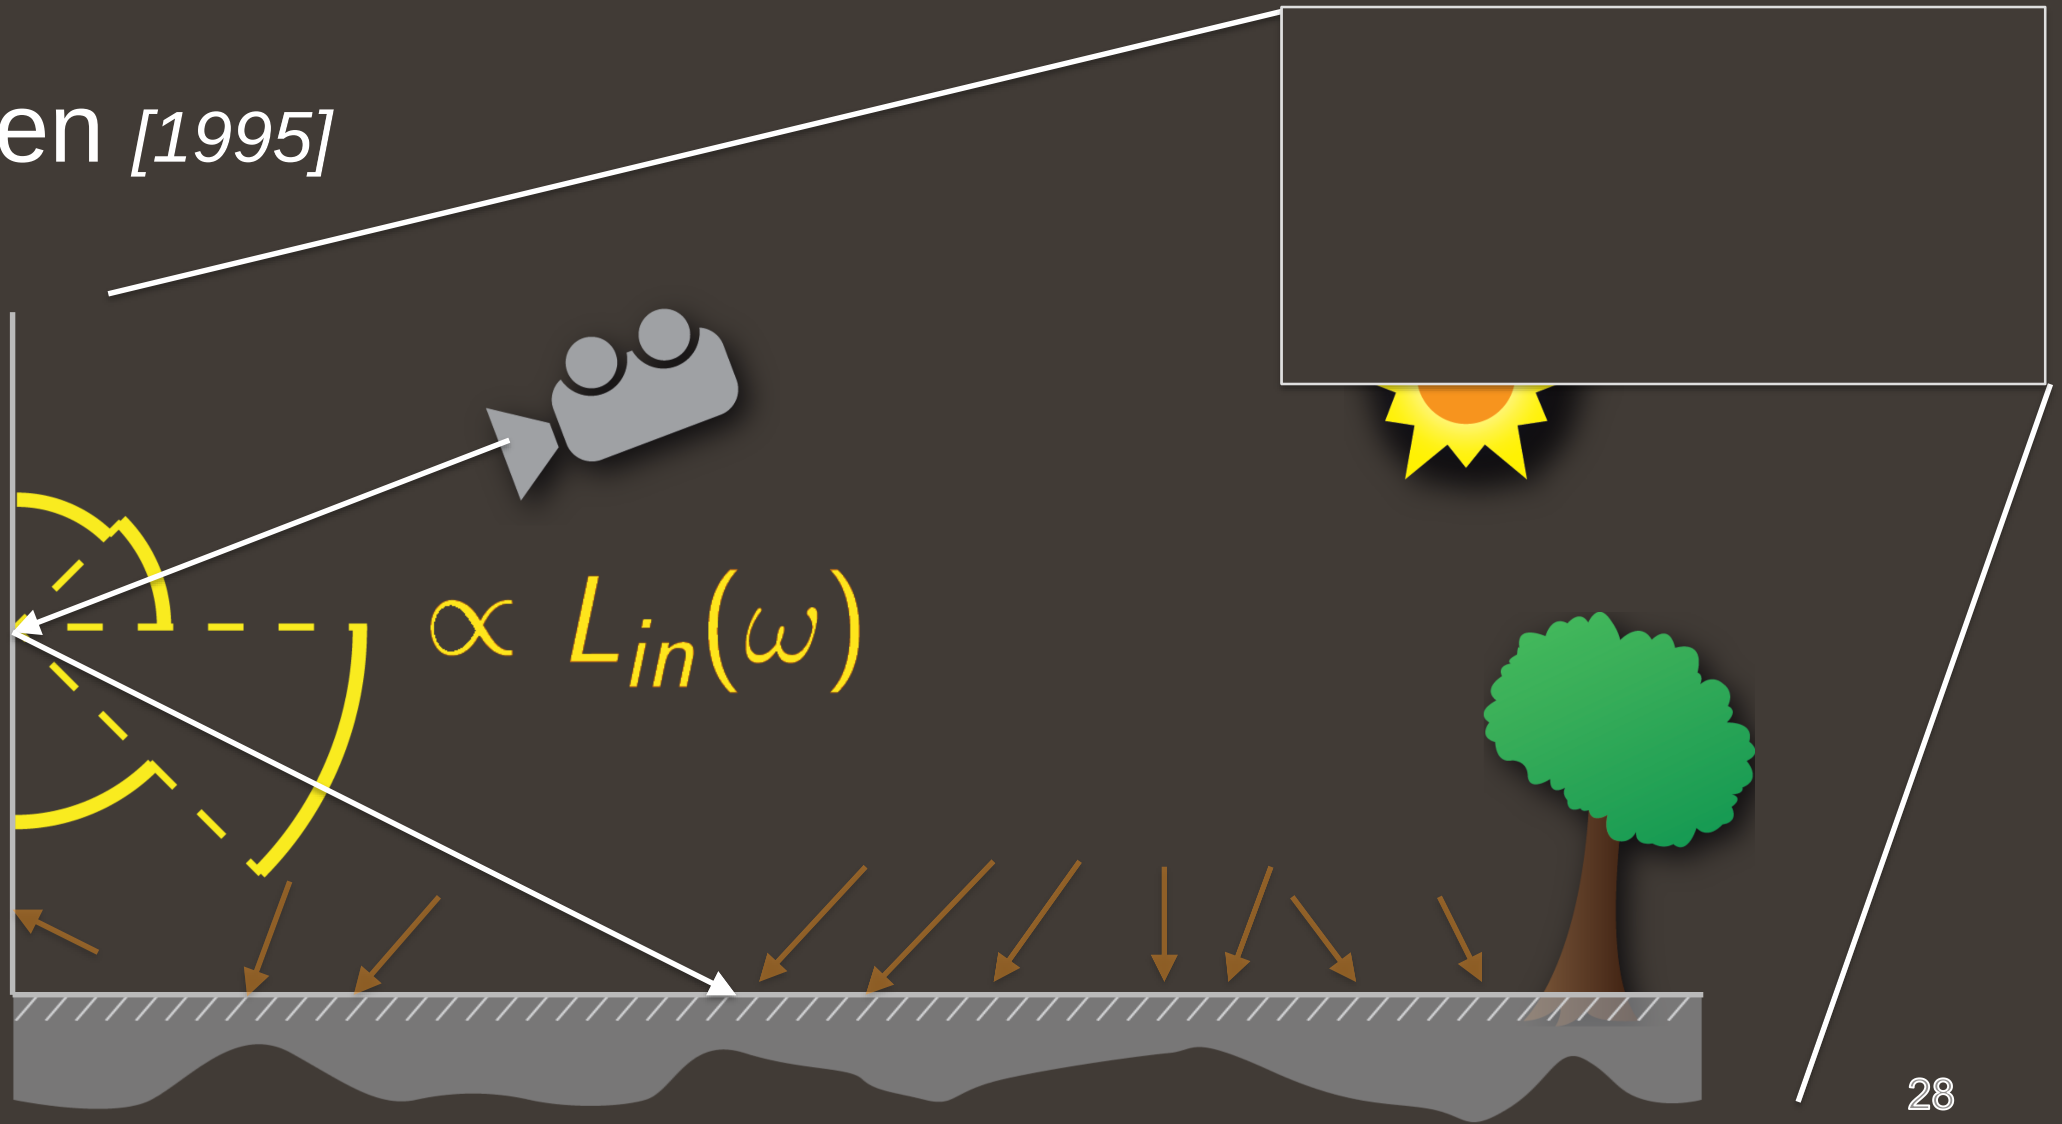

Previous work

- Jensen *[1995]*

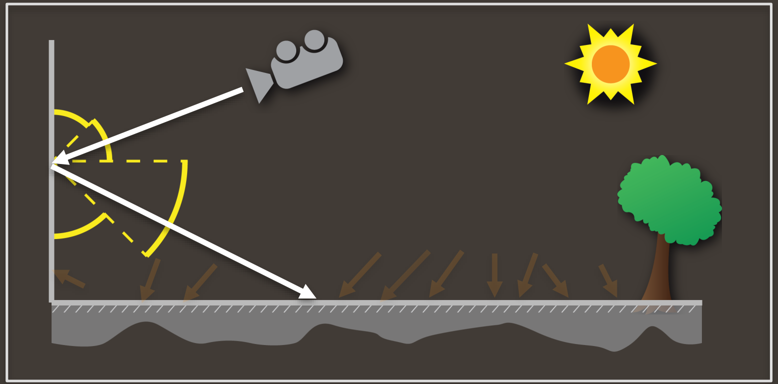

Previous work

- Jensen *[1995]*: reconstruction

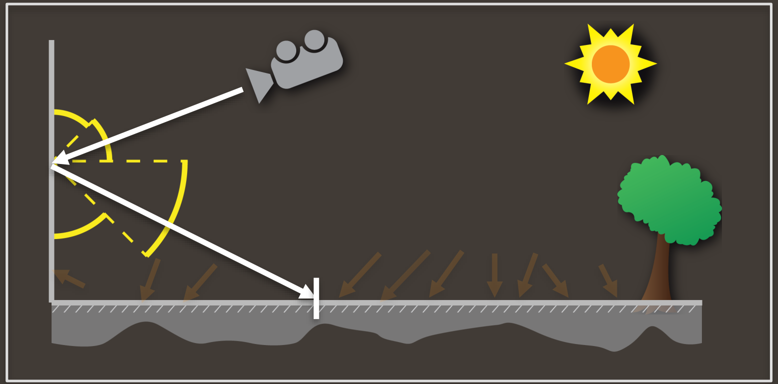

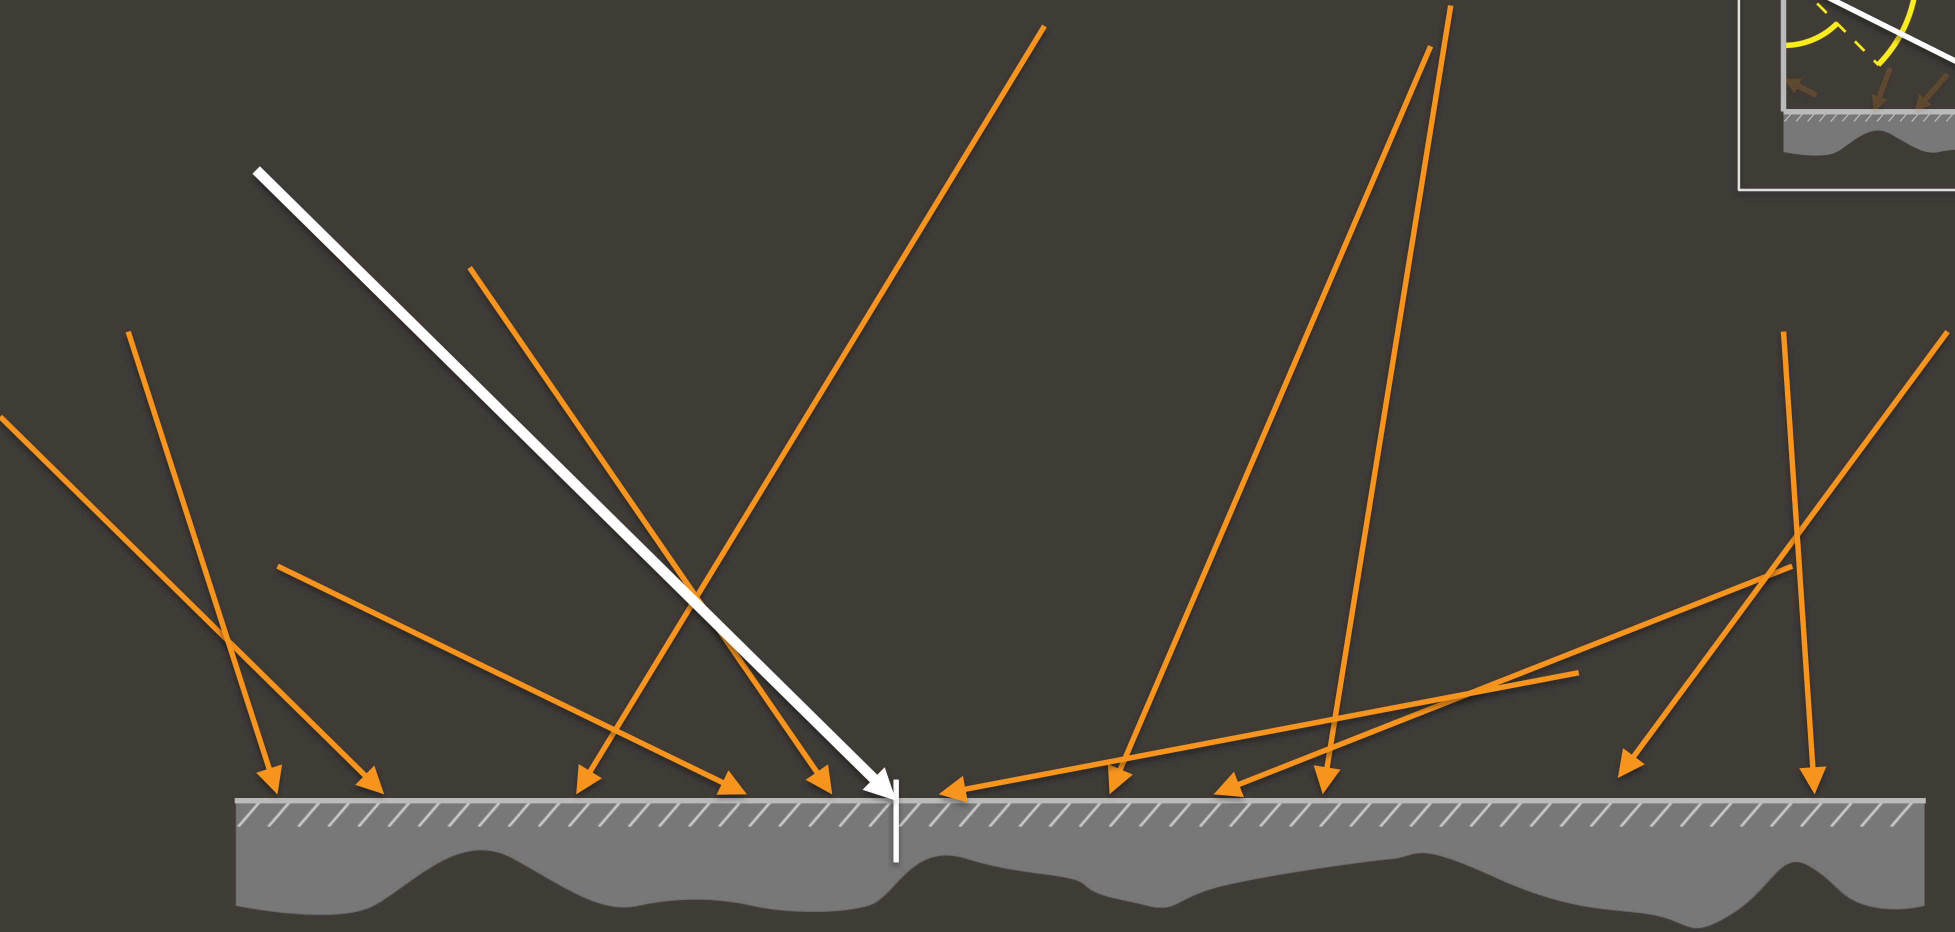

Previous work

- Jensen *[1995]*: reconstruction

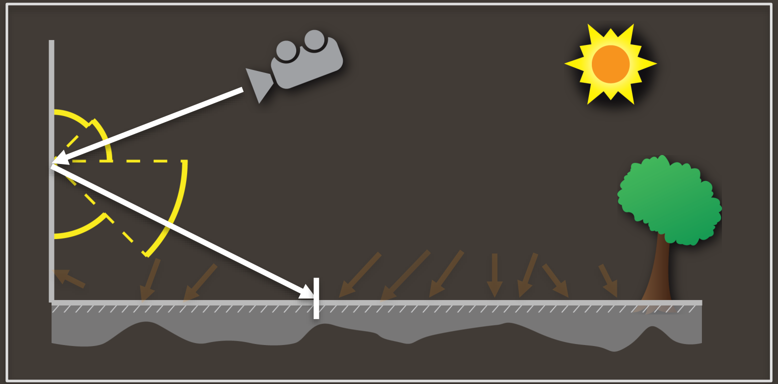

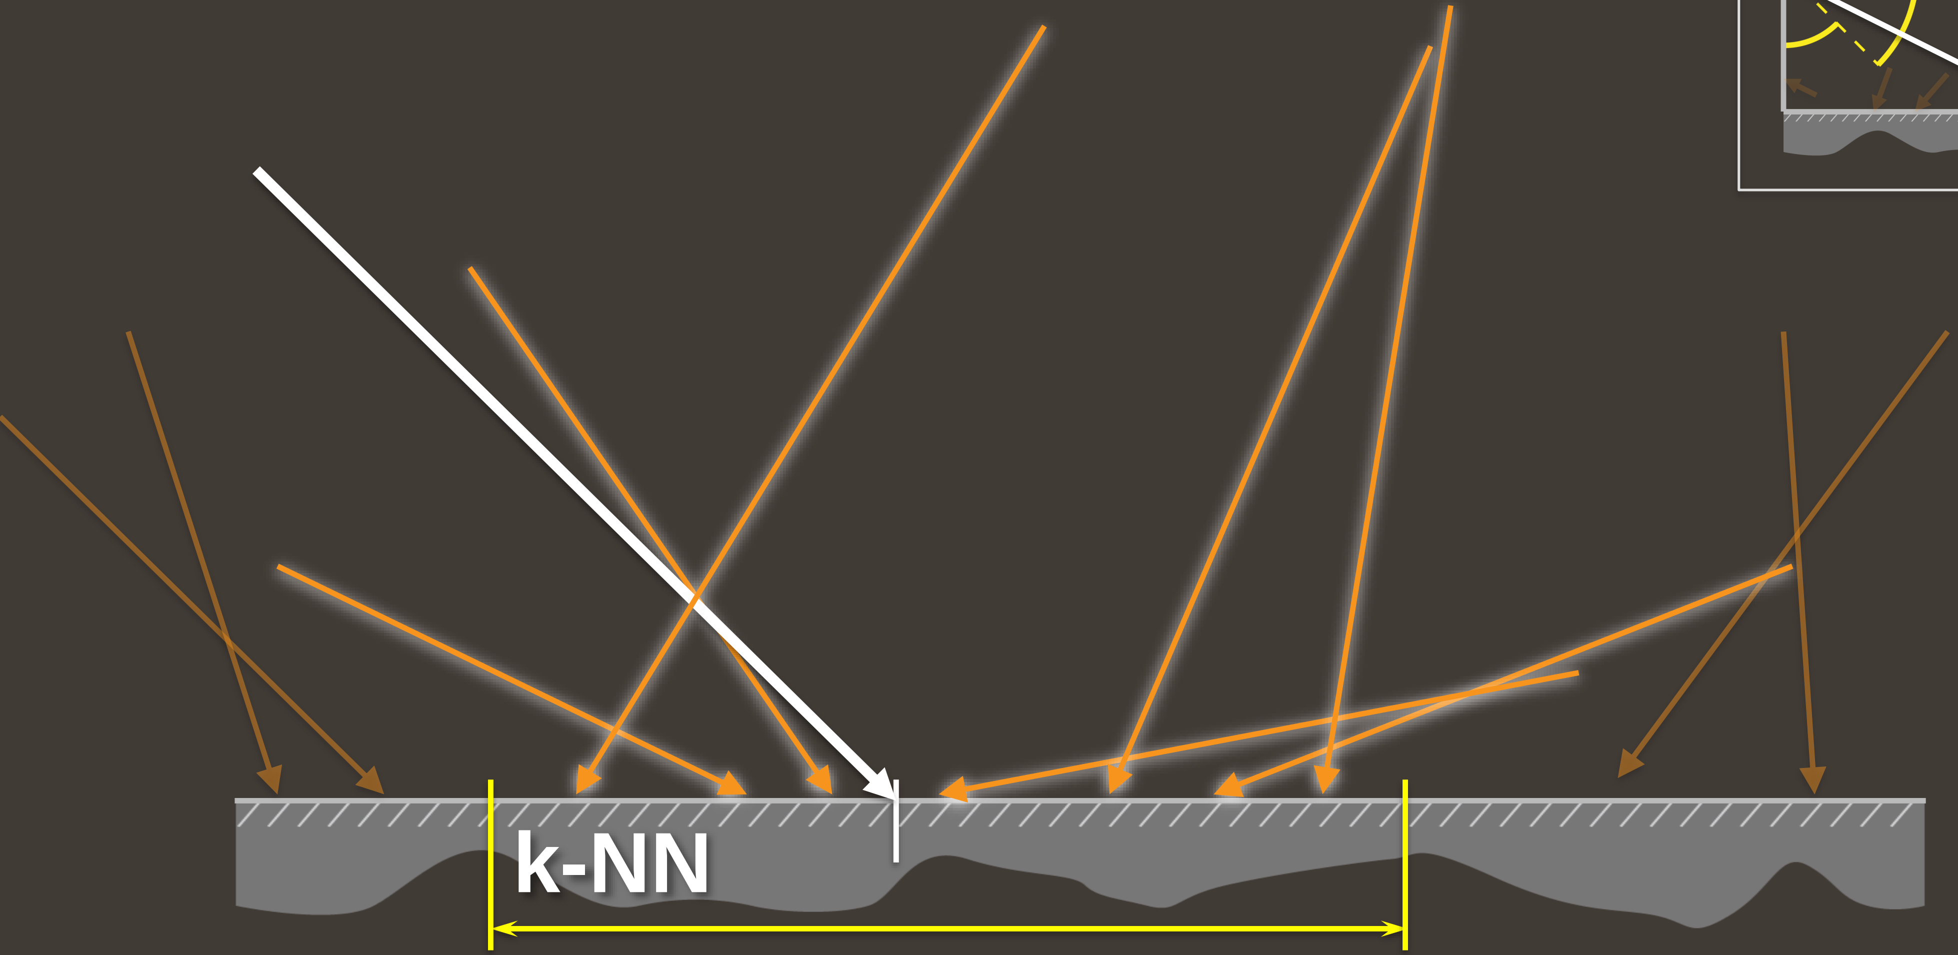

Previous work

- Jensen *[1995]*: reconstruction

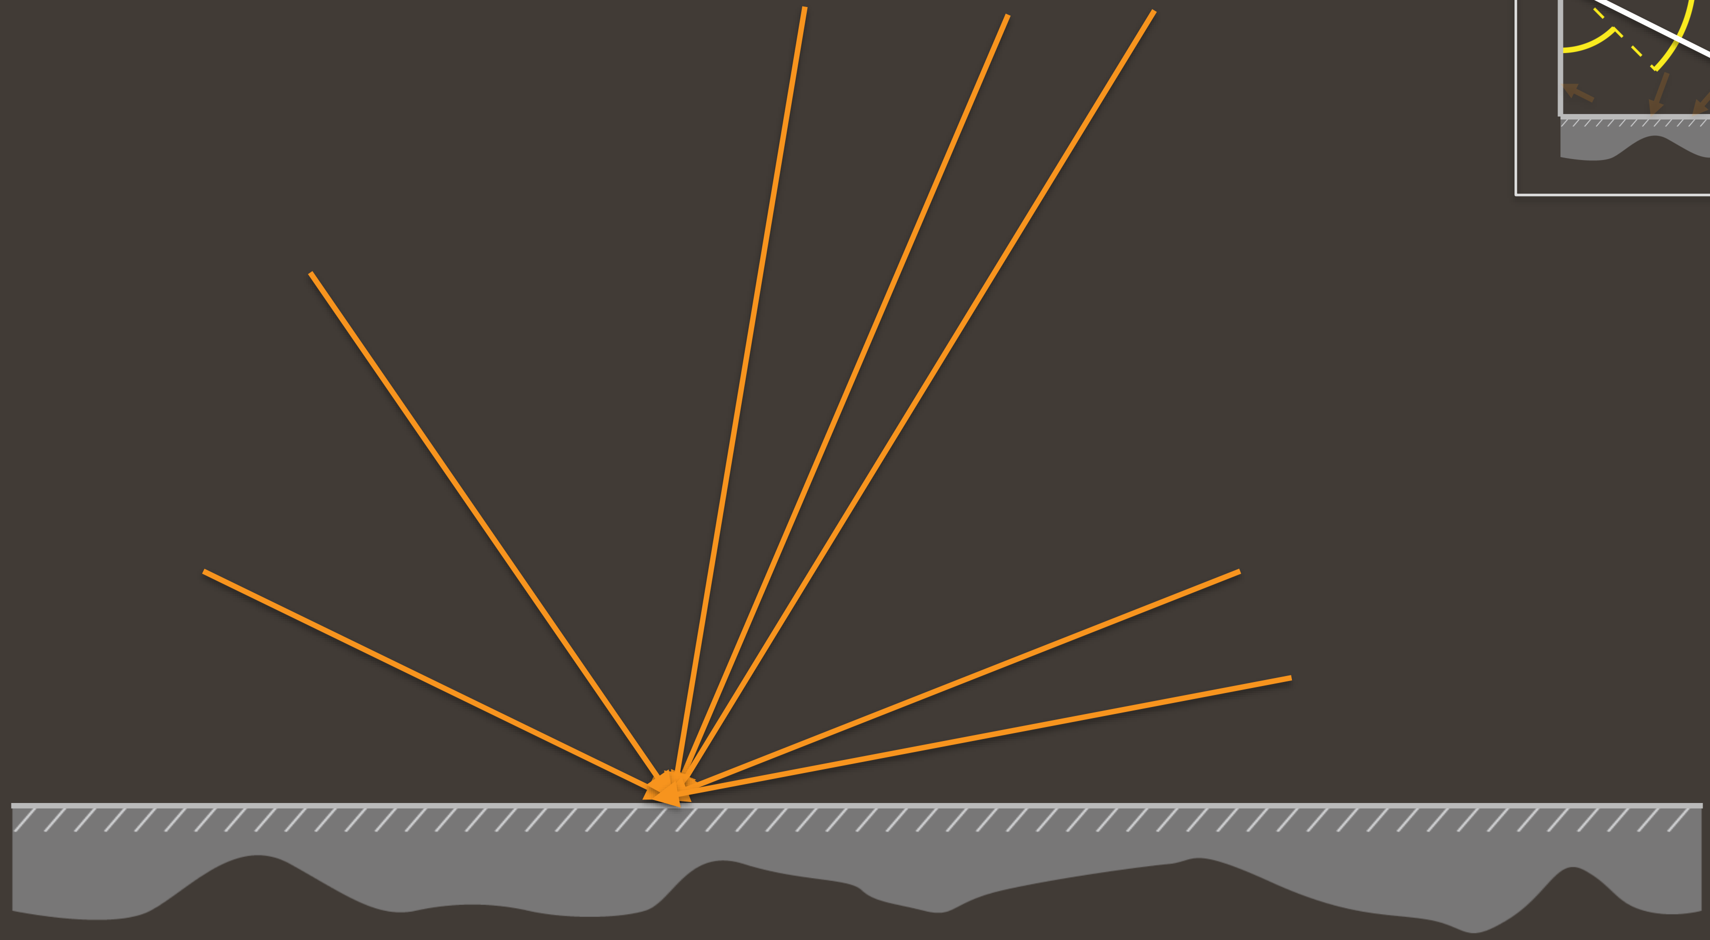

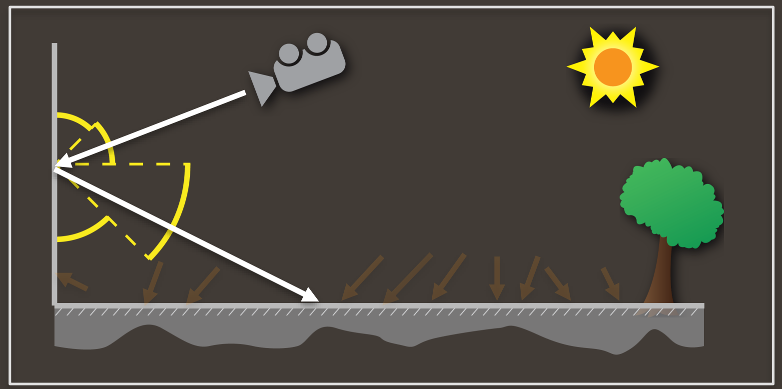

Previous work

- Jensen *[1995]*: reconstruction

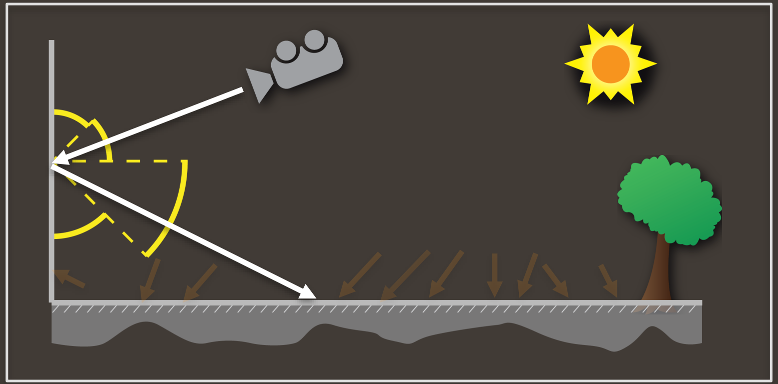

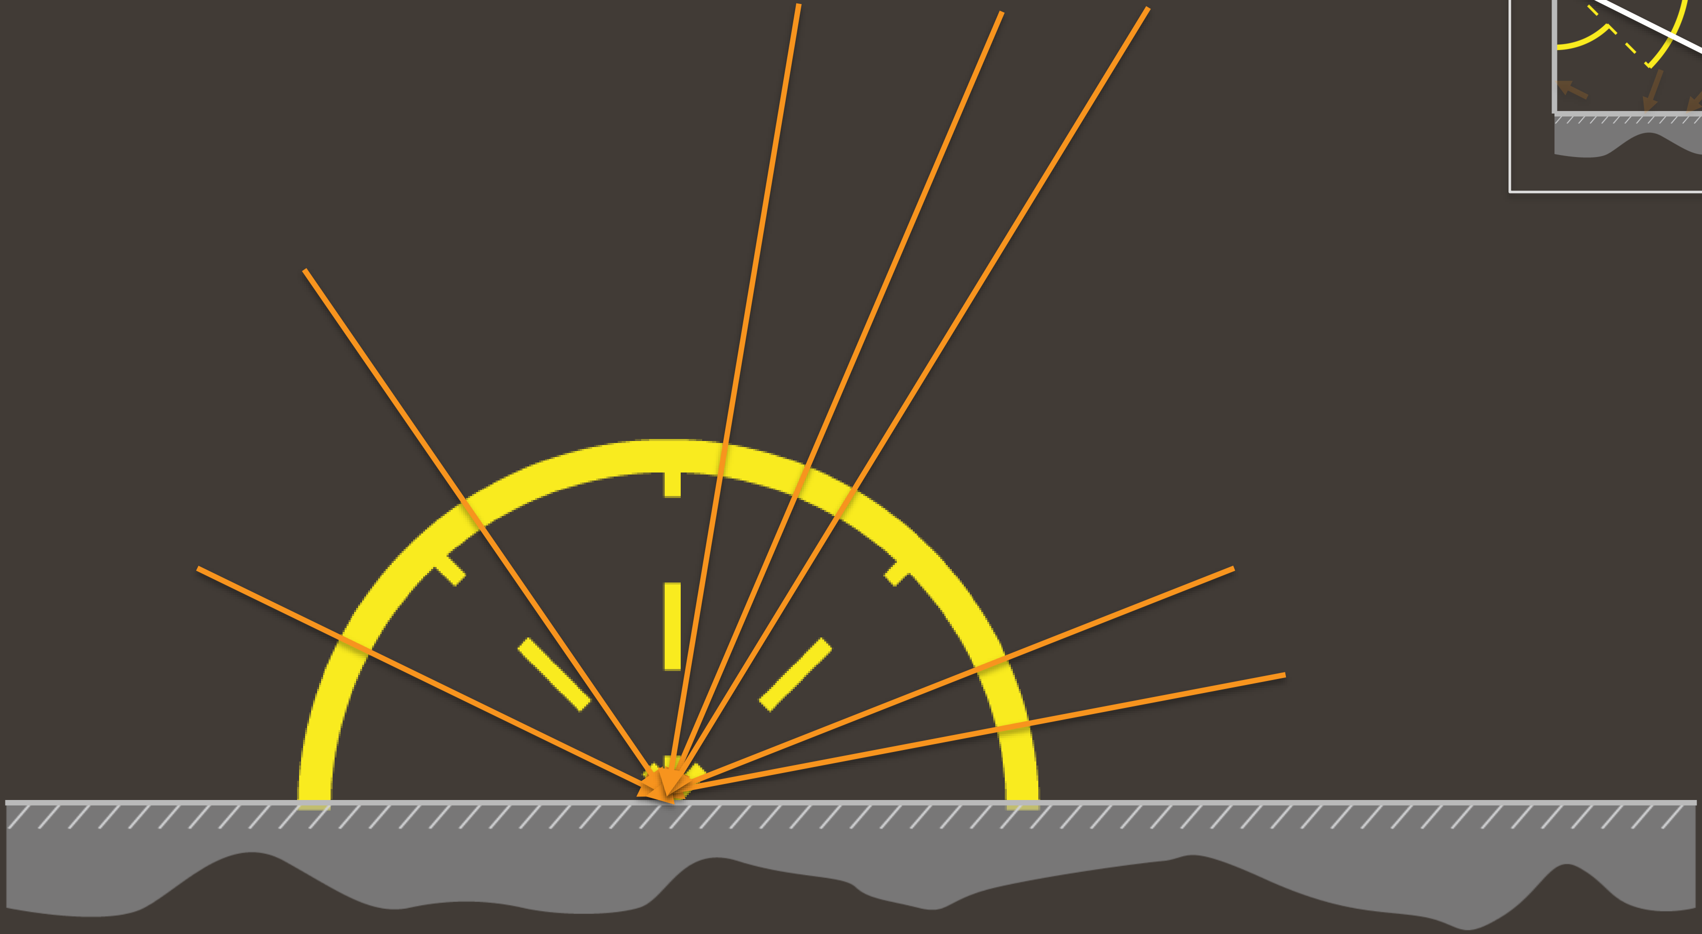

Previous work

- Jensen *[1995]*: reconstruction

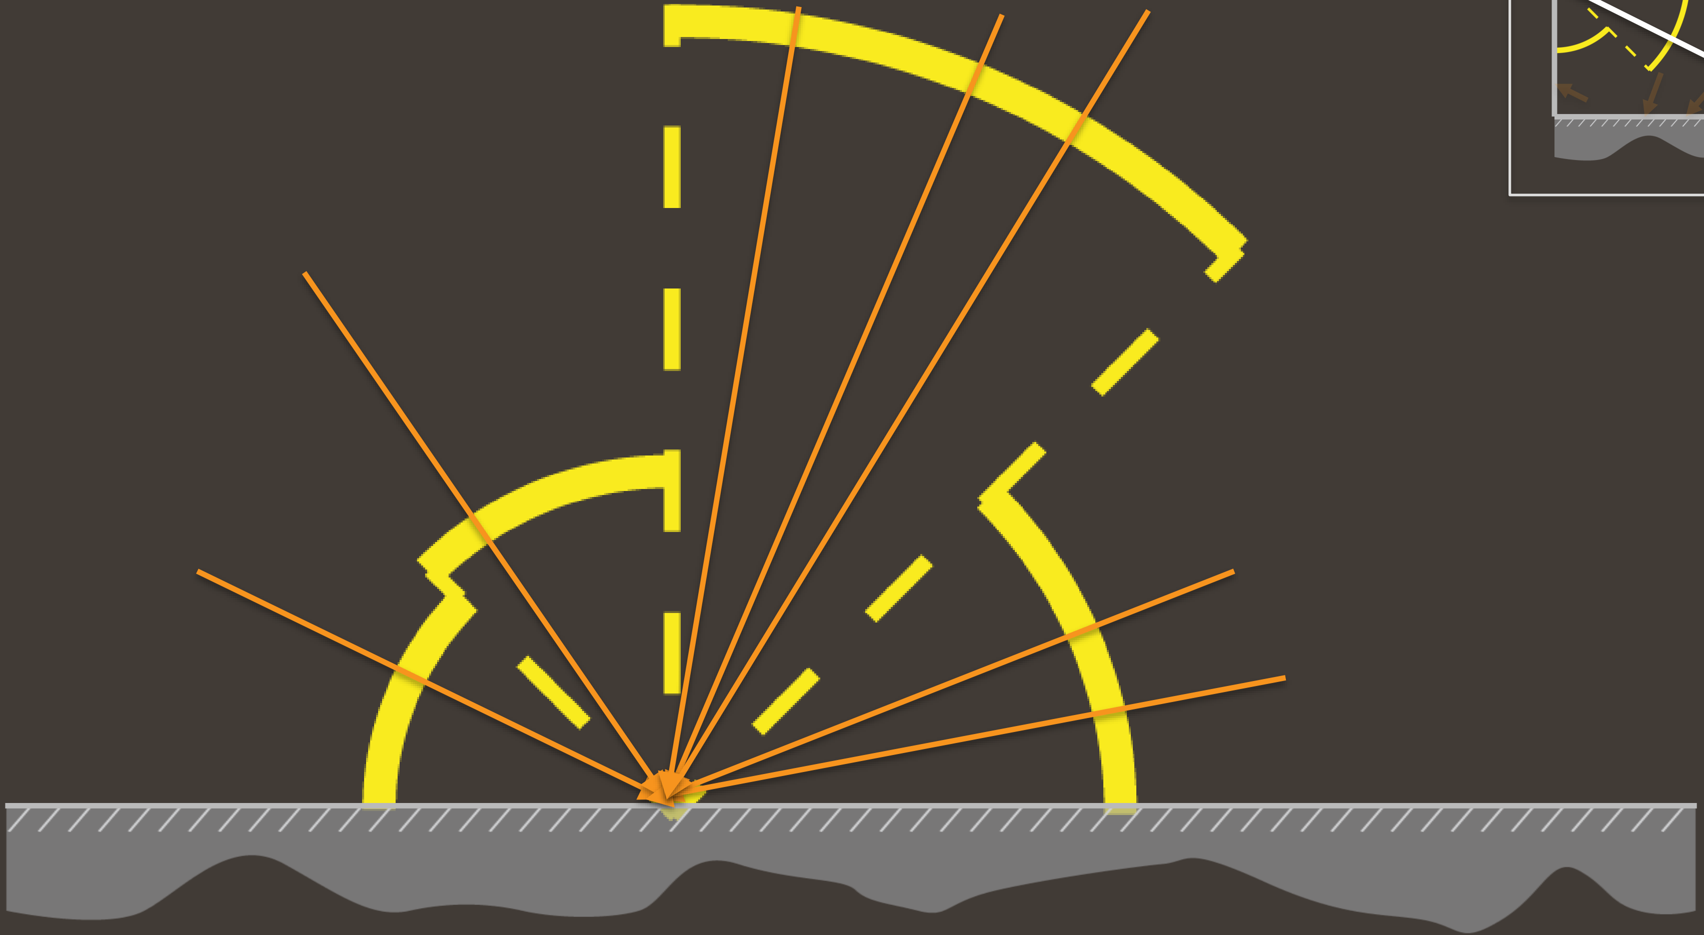

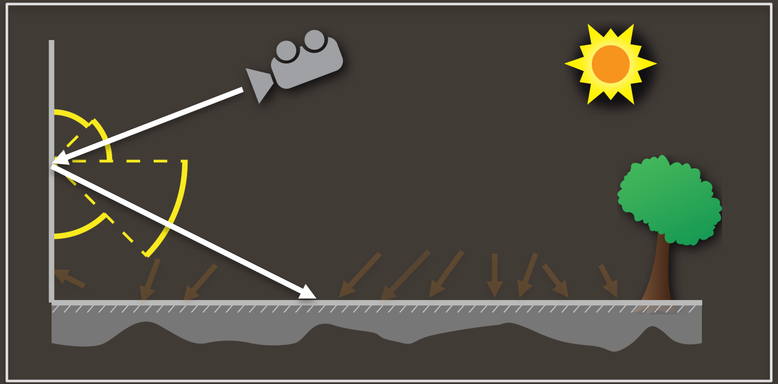

Limitations of previous work

- Bad approximation of $L_{in}(\omega)$ in complex scenes

Limitations of previous work

- Bad approximation of $L_{in}(\omega)$ in complex scenes

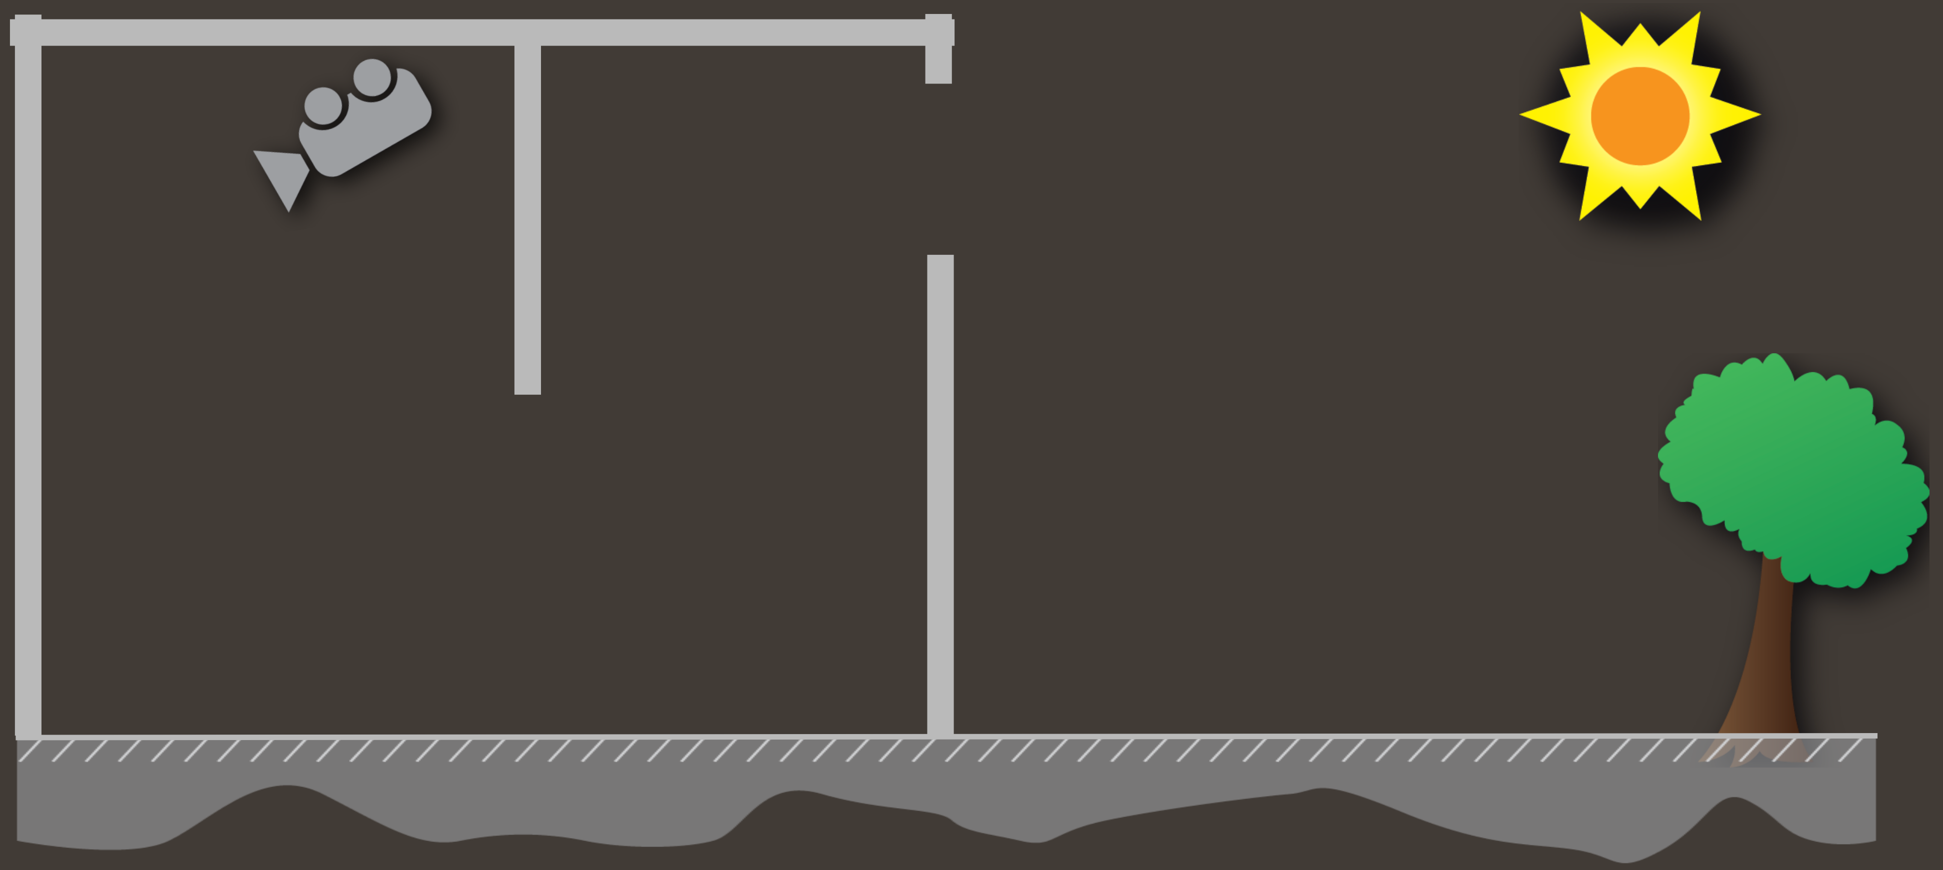

Limitations of previous work

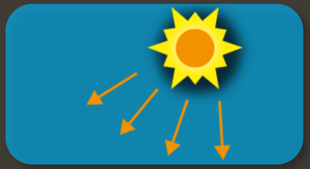

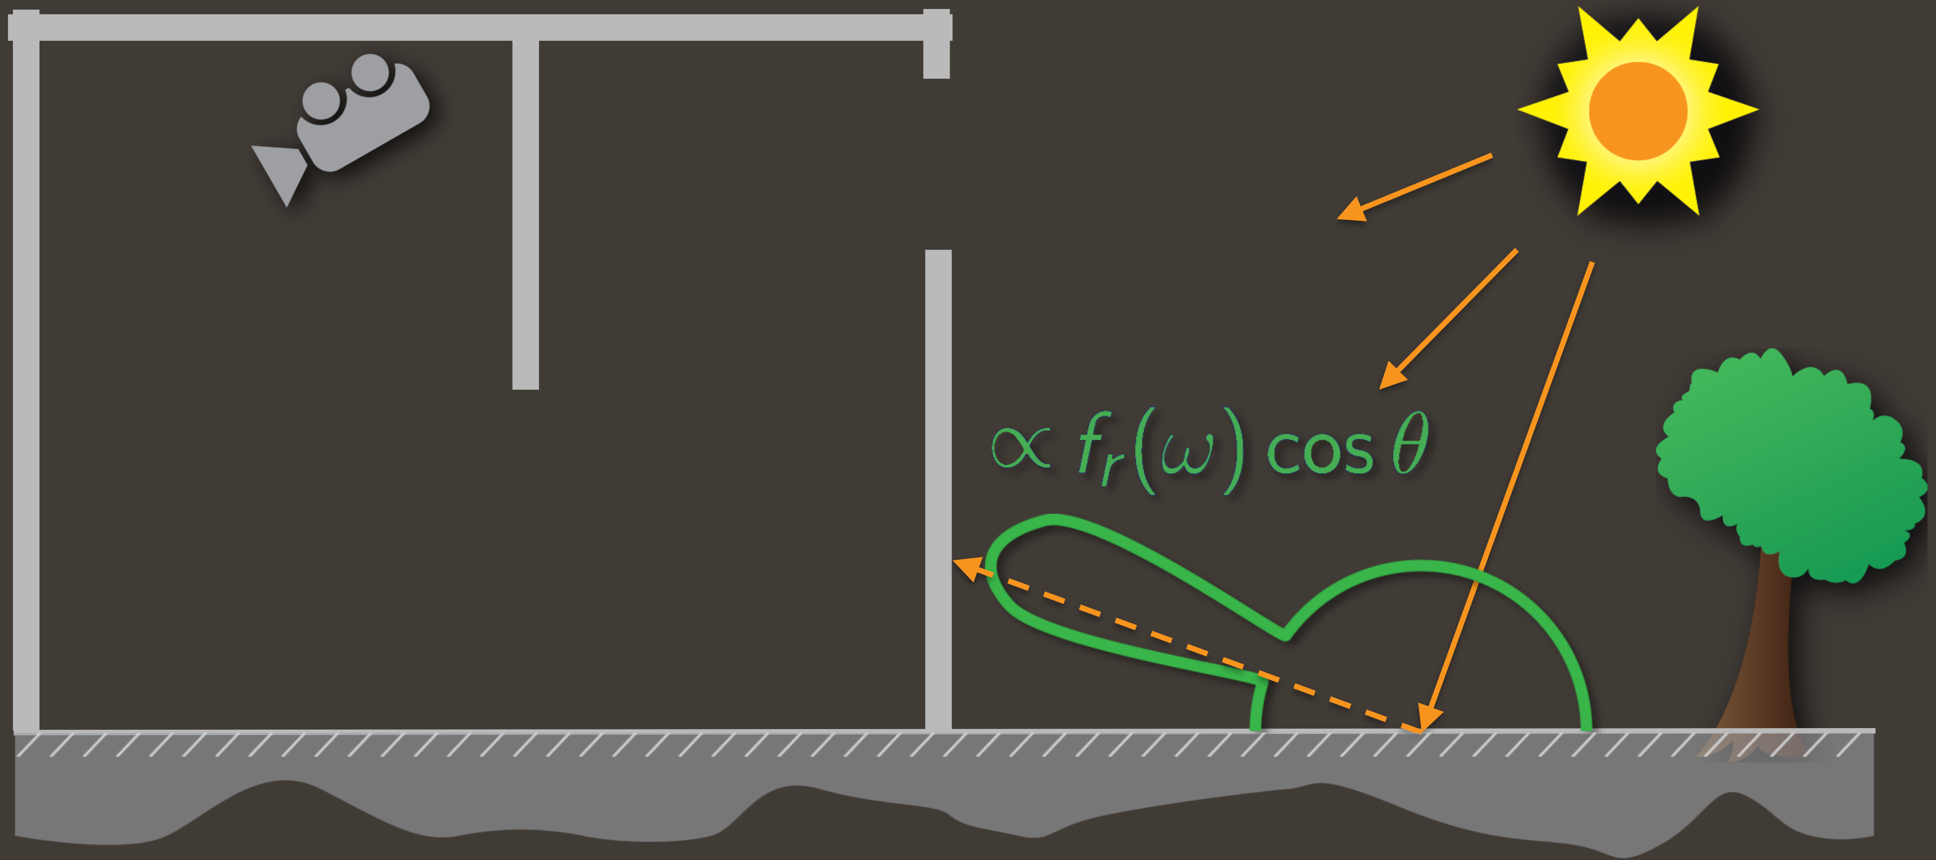

Limitations of previous work

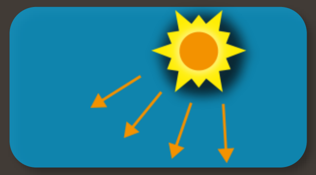

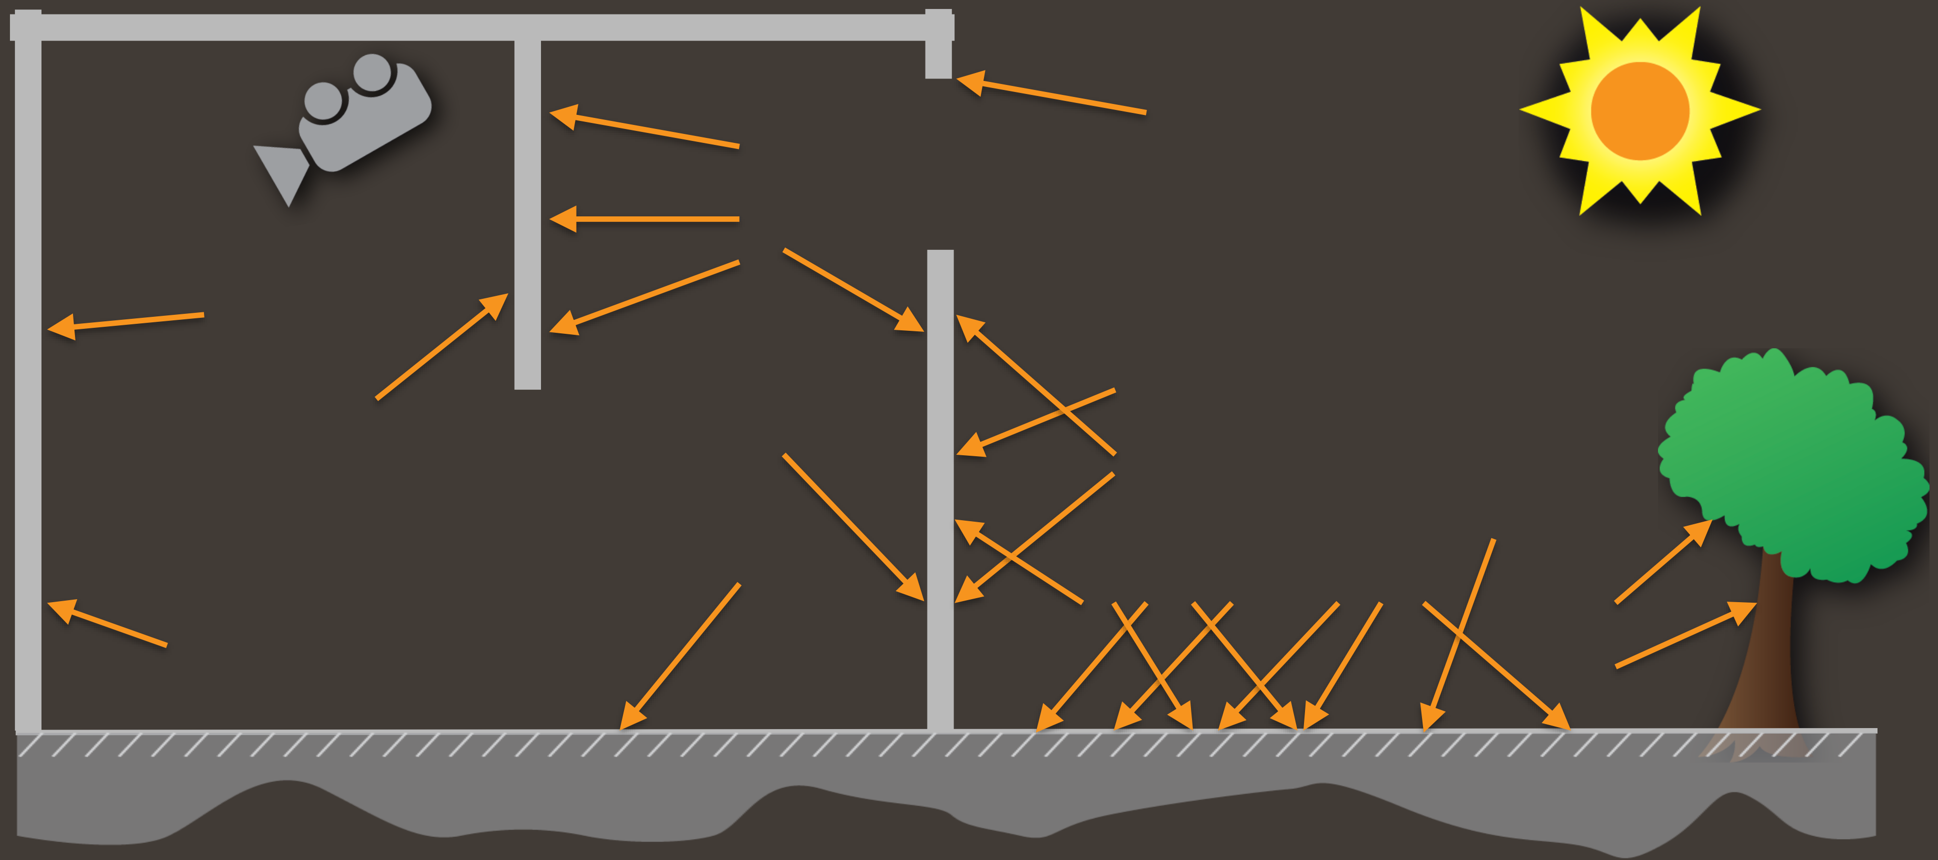

Limitations of previous work

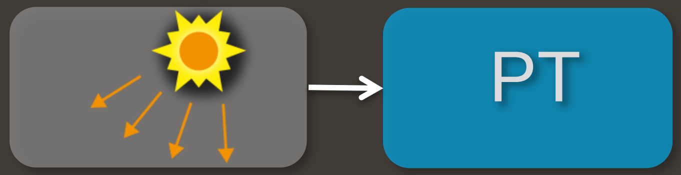

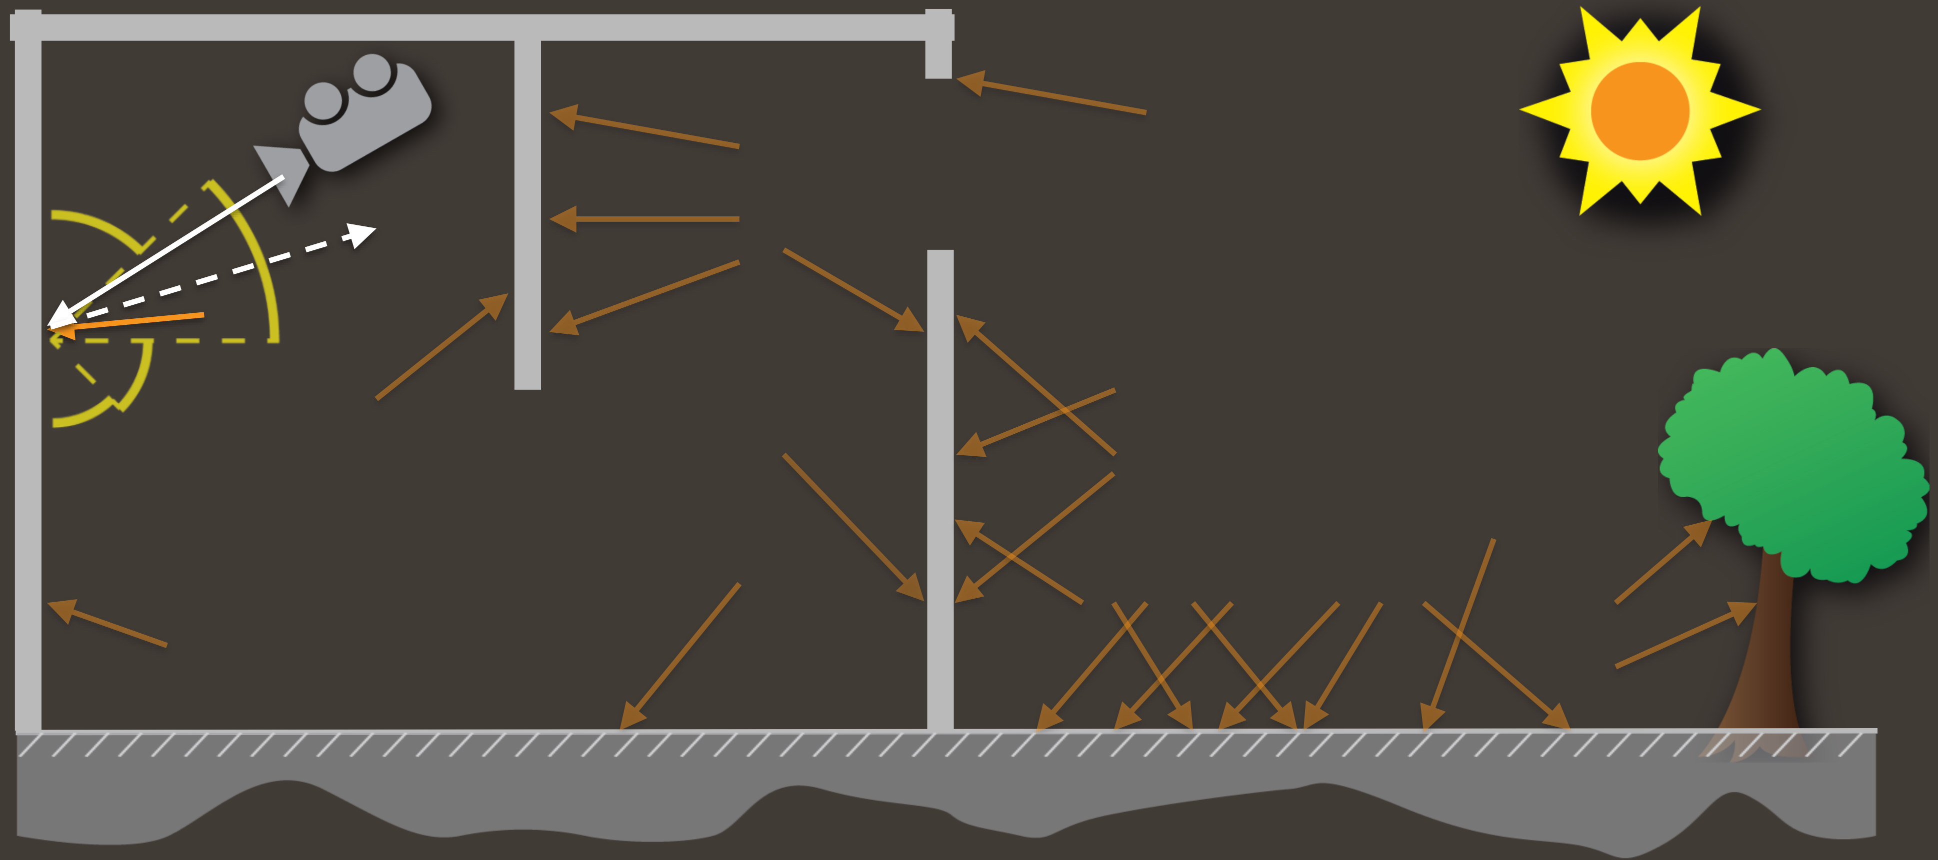

Limitations of previous work

Not enough memory!

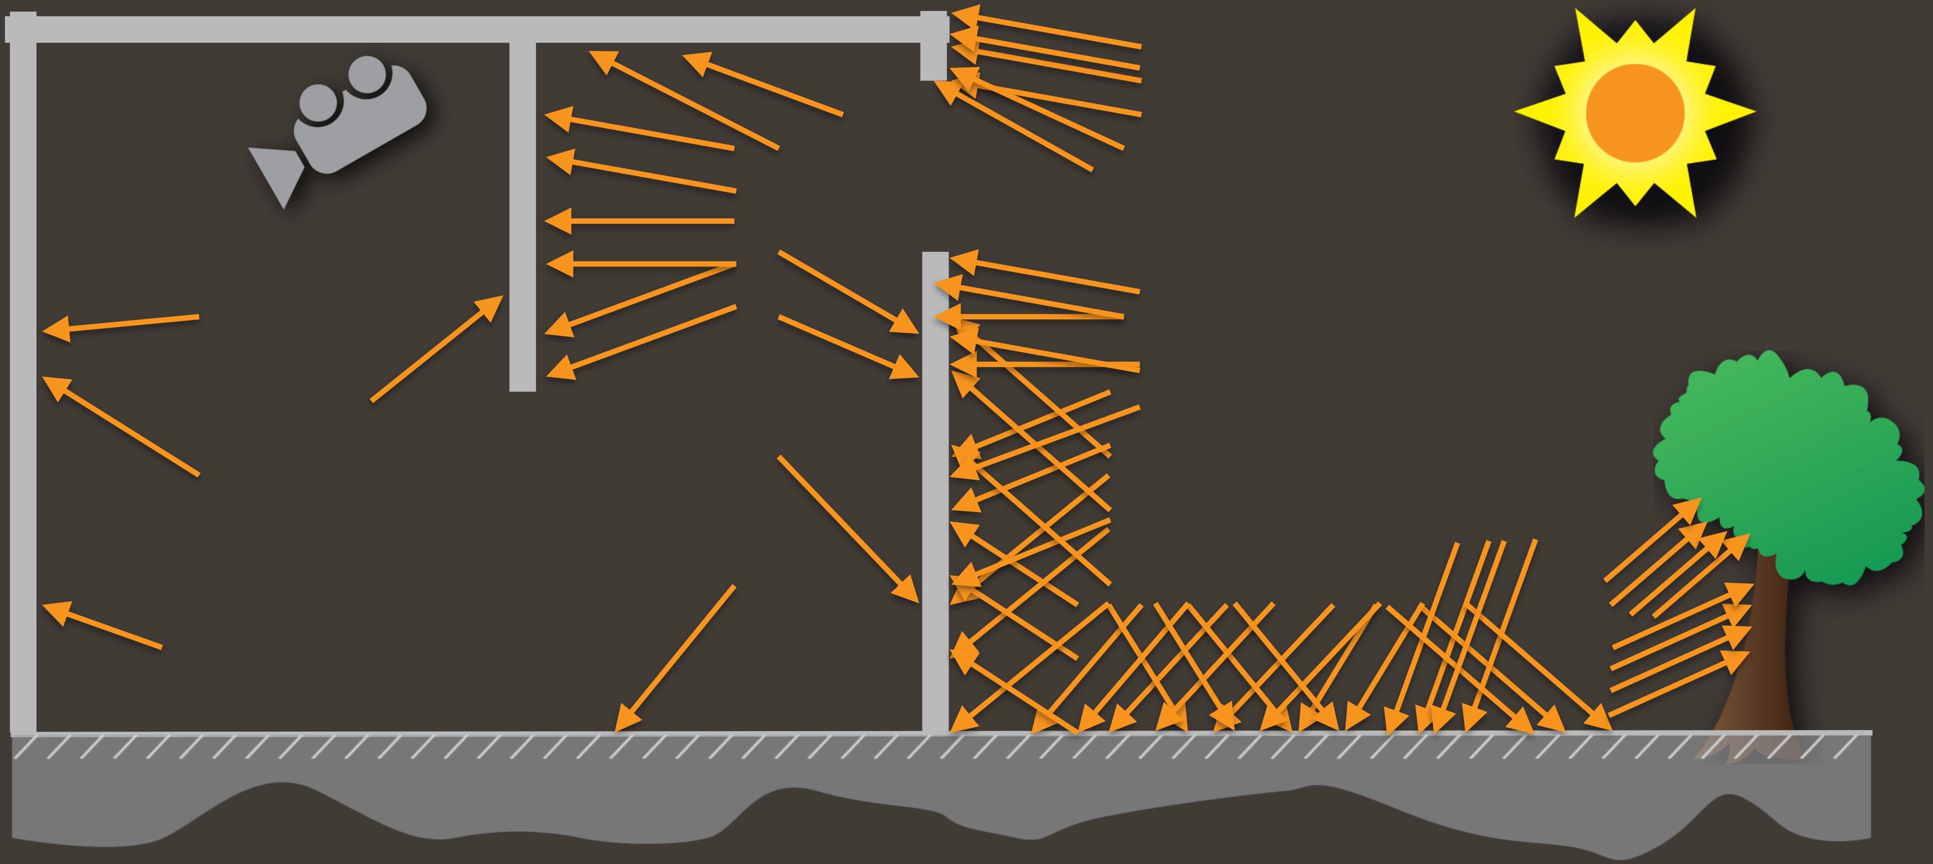

Our solution

GMM

- The Gaussian mixture model (GMM)

Our solution

GMM $\Rightarrow$ **on-line learning**

- The Gaussian mixture model (GMM)

Our solution

GMM $\Rightarrow$ **on-line learning** $\Rightarrow$ **constant memory**

- The Gaussian mixture model (GMM)

Overcoming the memory constraint

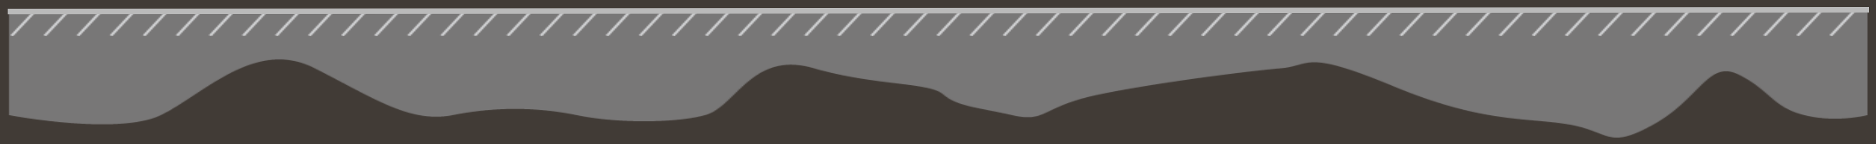

Overcoming the memory constraint

1st pass 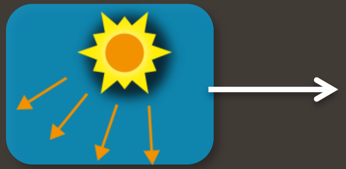

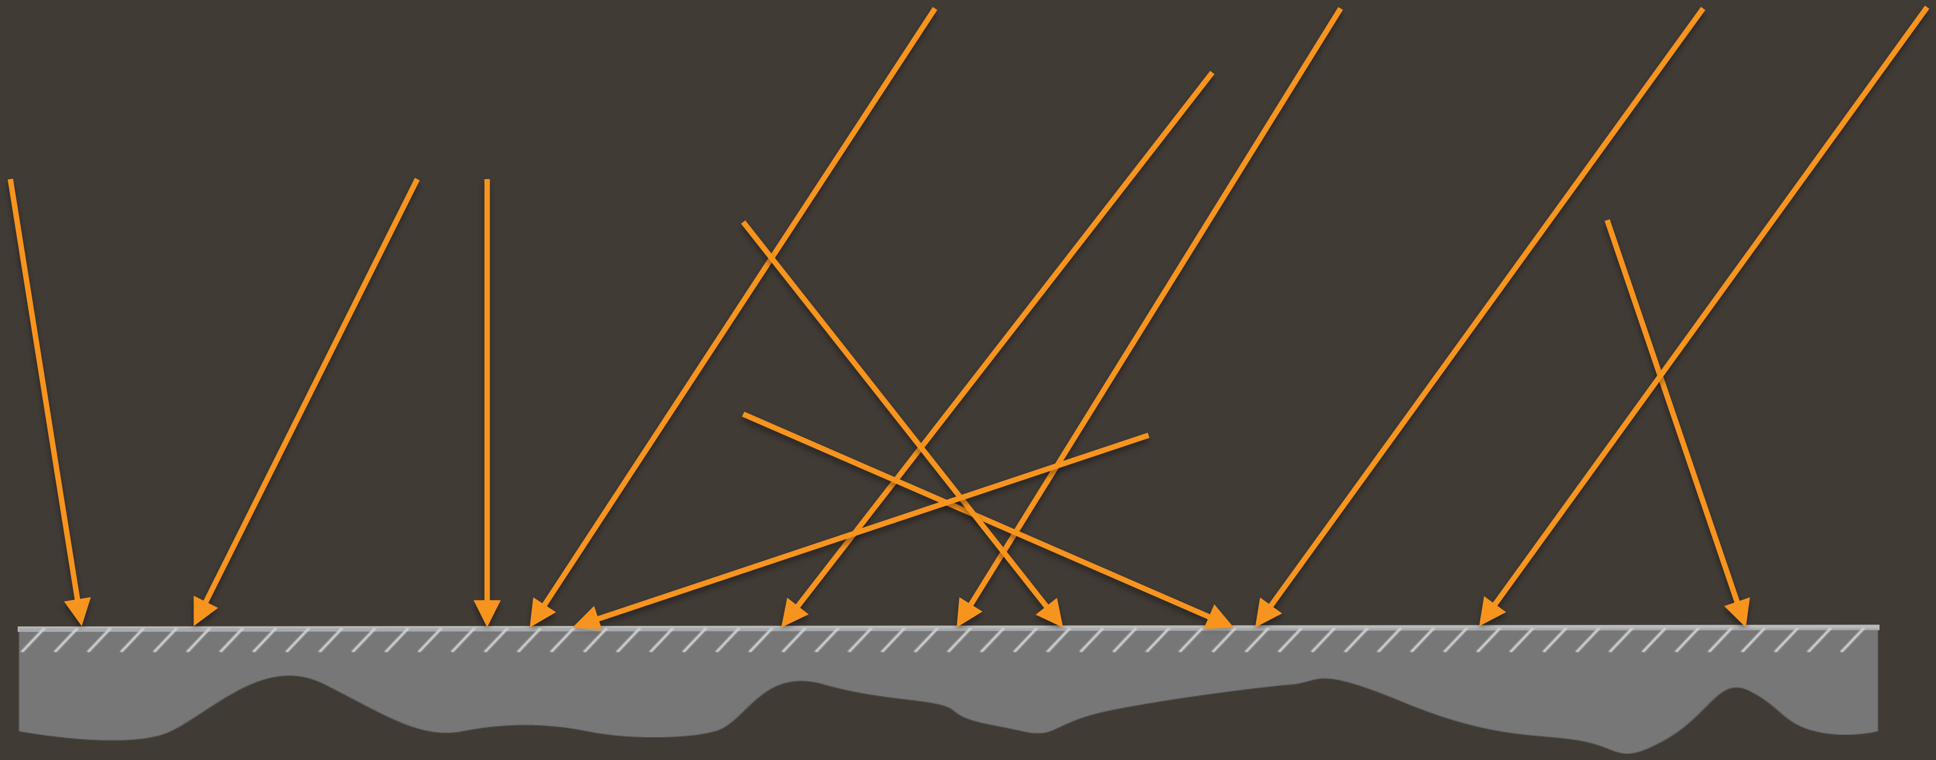

Overcoming the memory constraint

1st pass 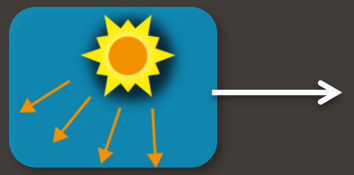

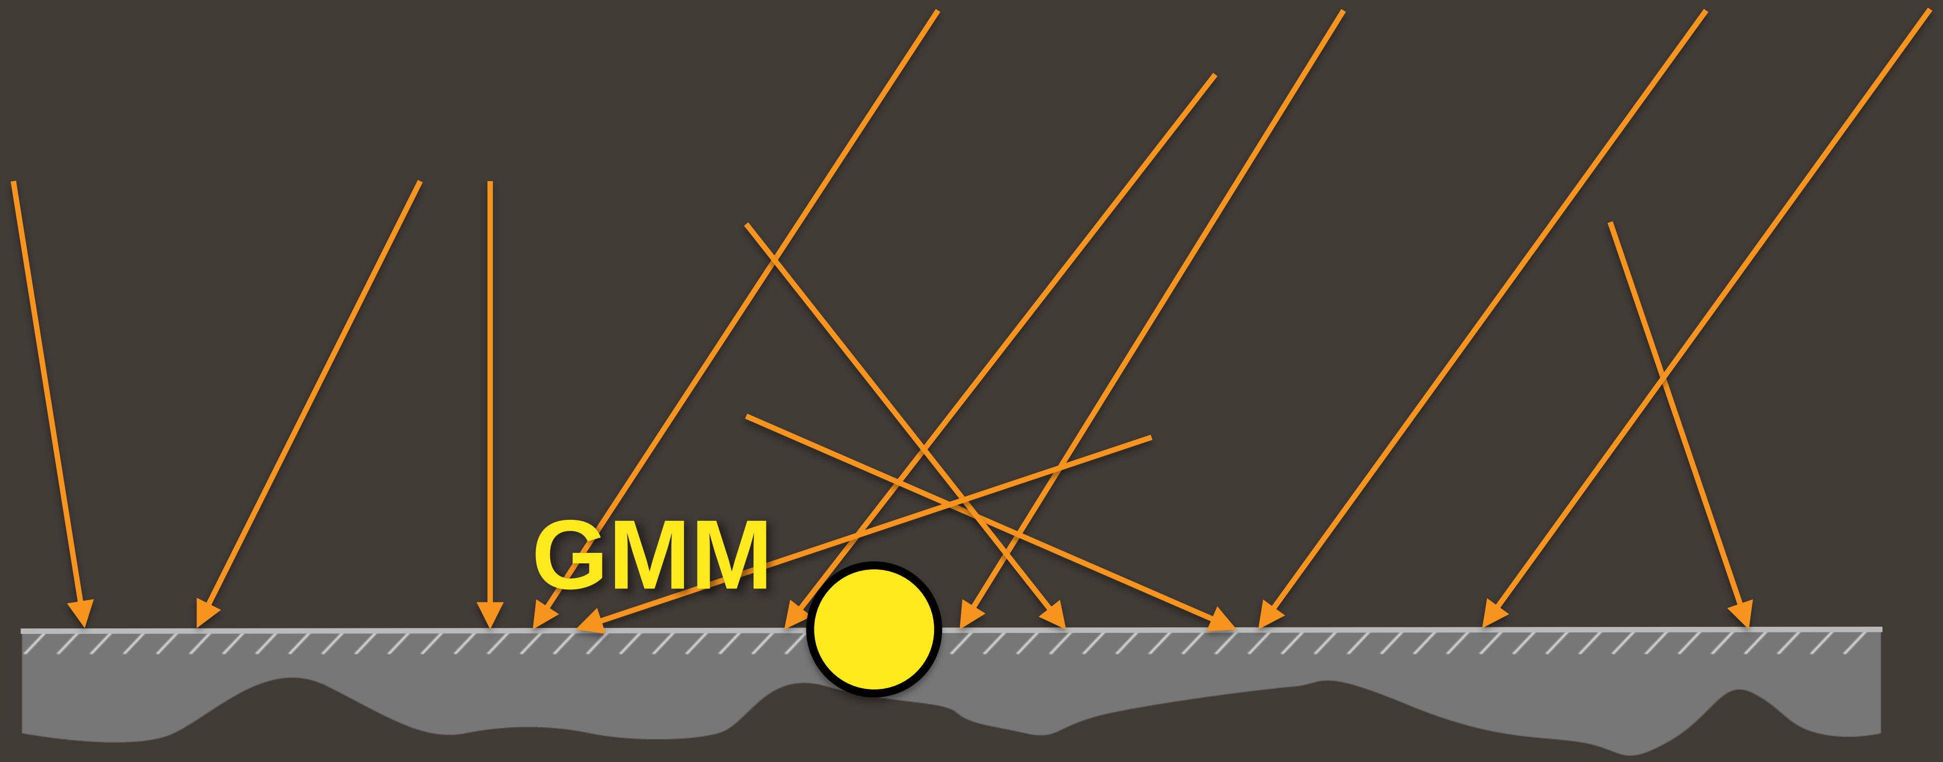

Overcoming the memory constraint

1st pass 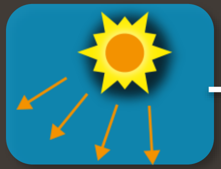 →

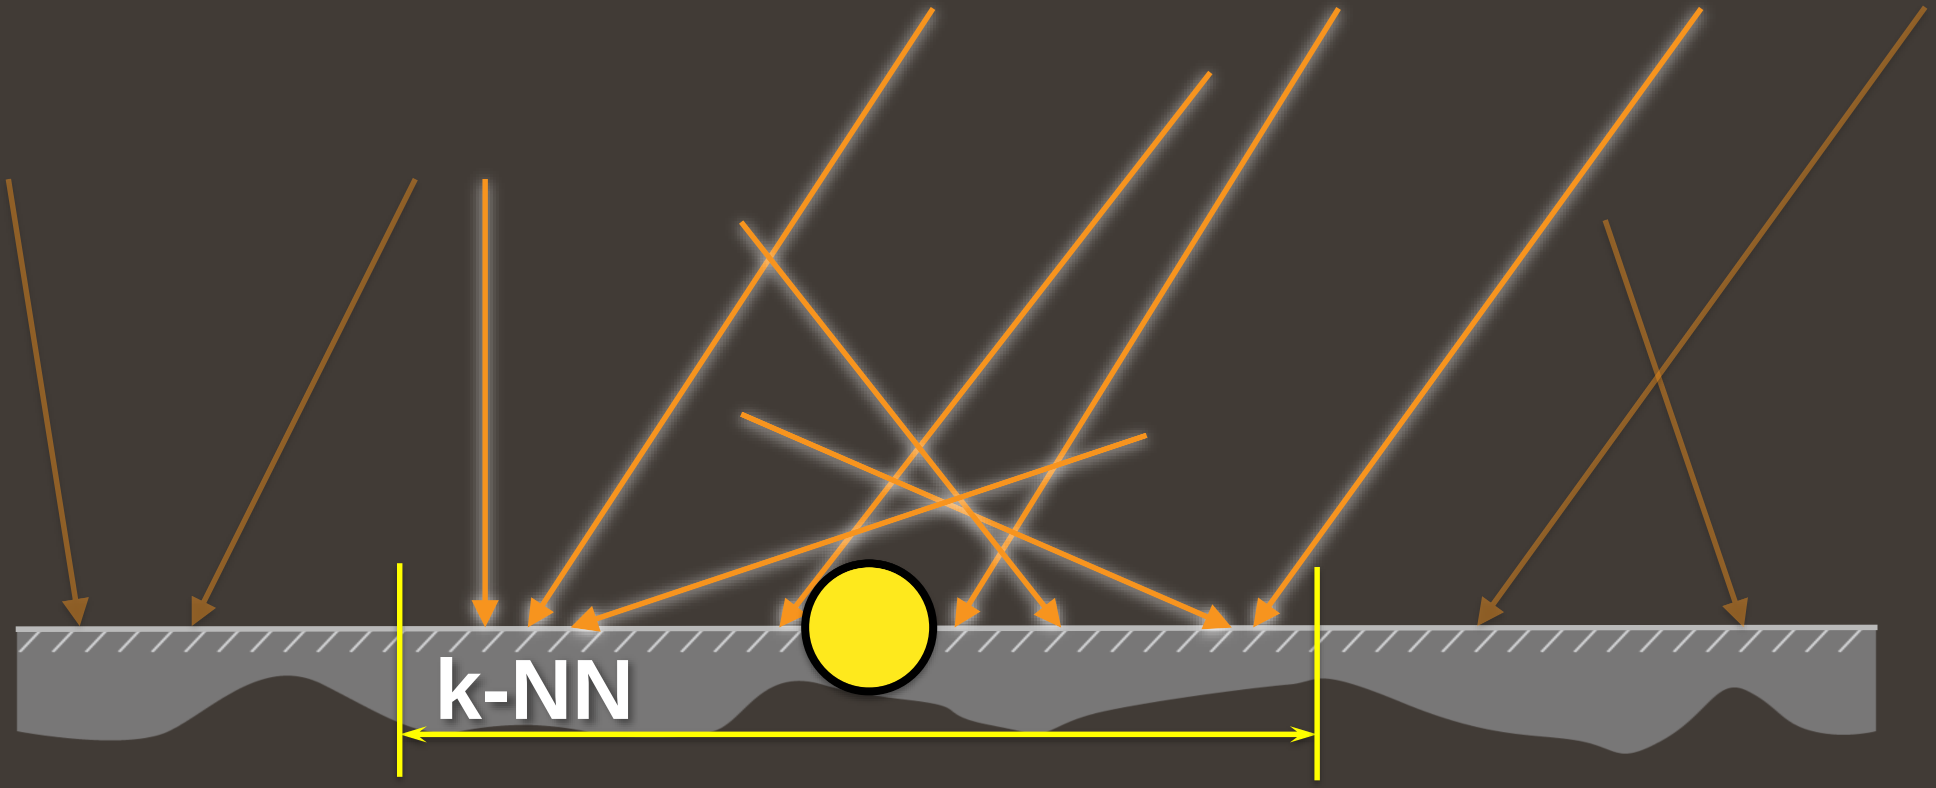

Overcoming the memory constraint

1st pass 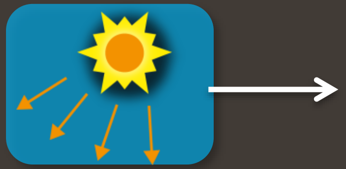

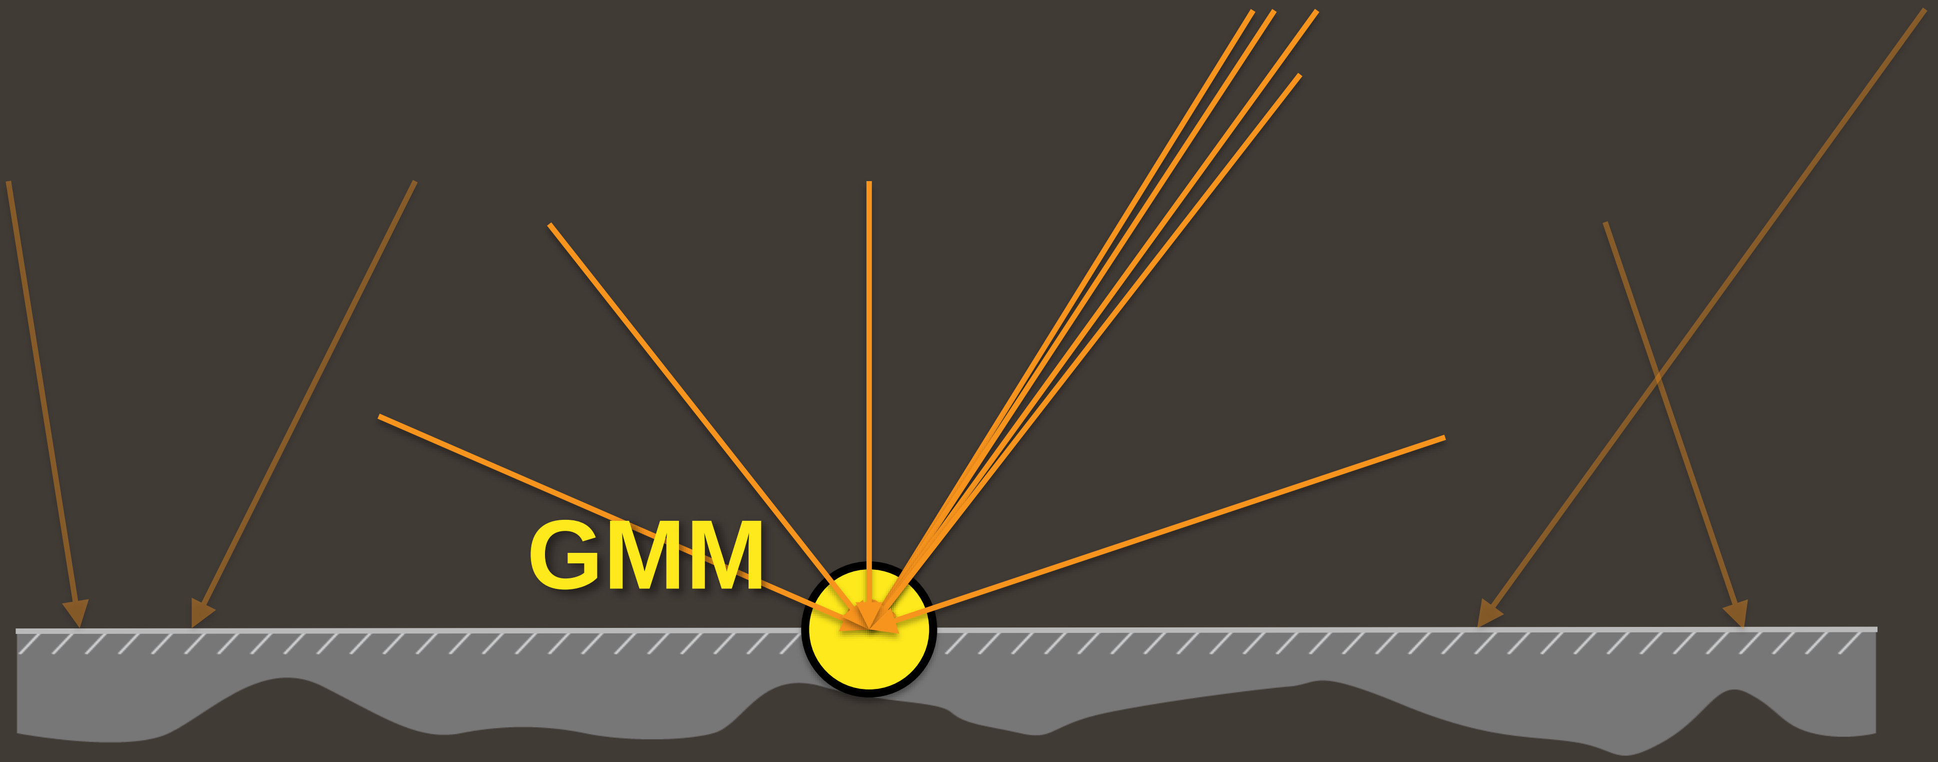

Overcoming the memory constraint

1st pass 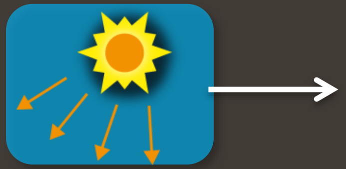

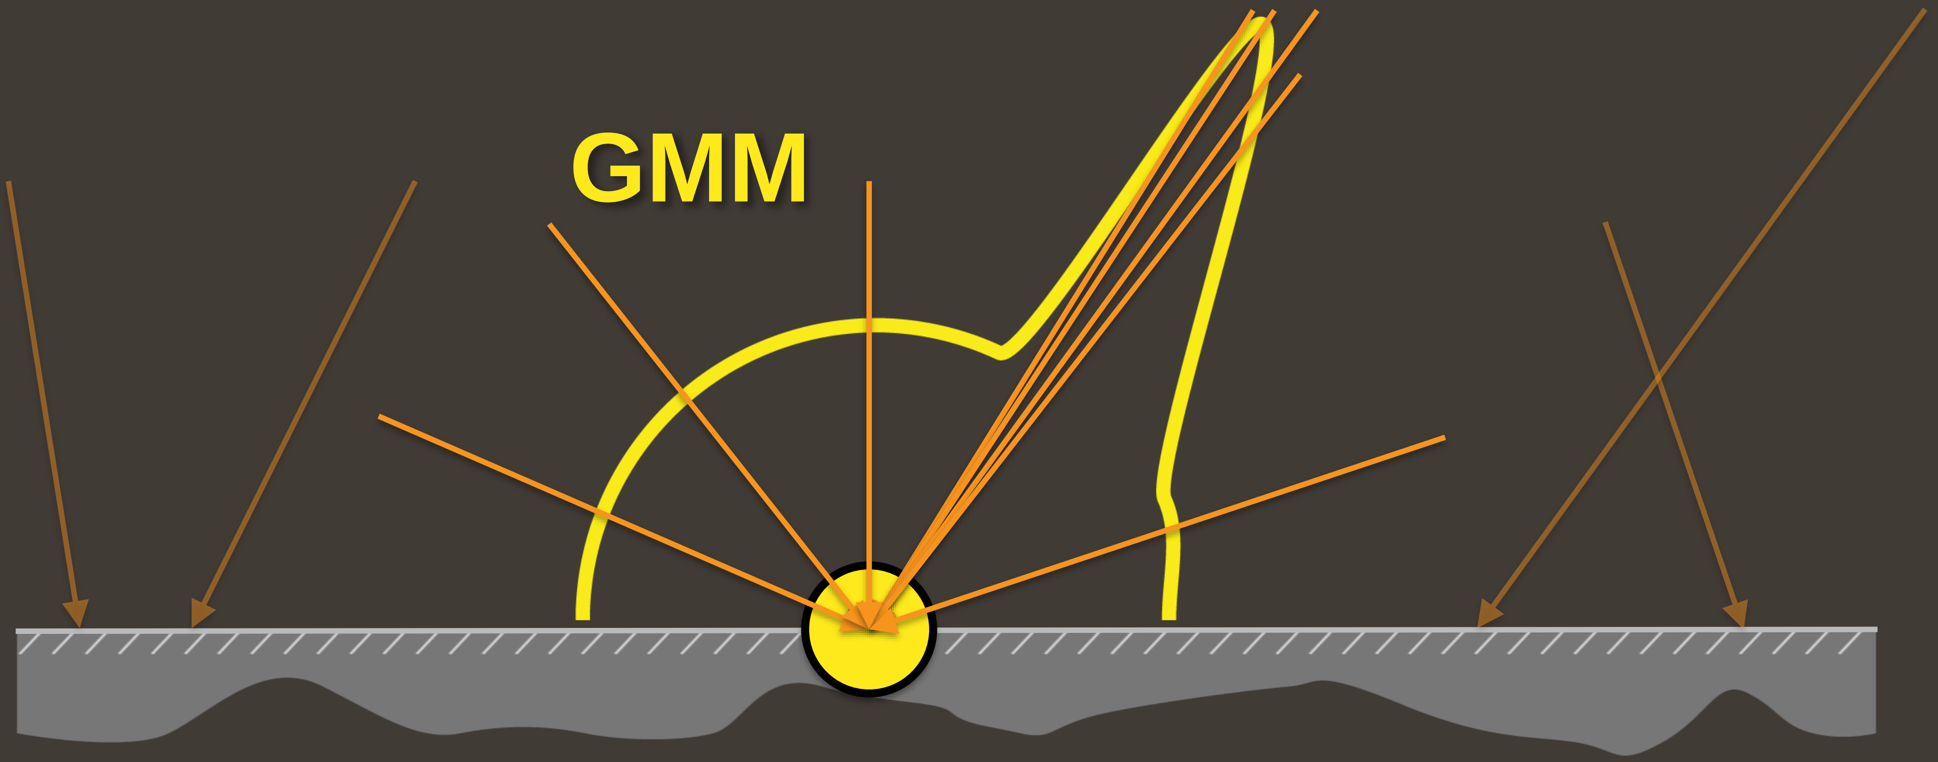

Overcoming the memory constraint

1st pass 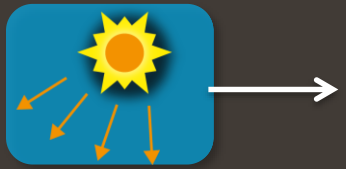

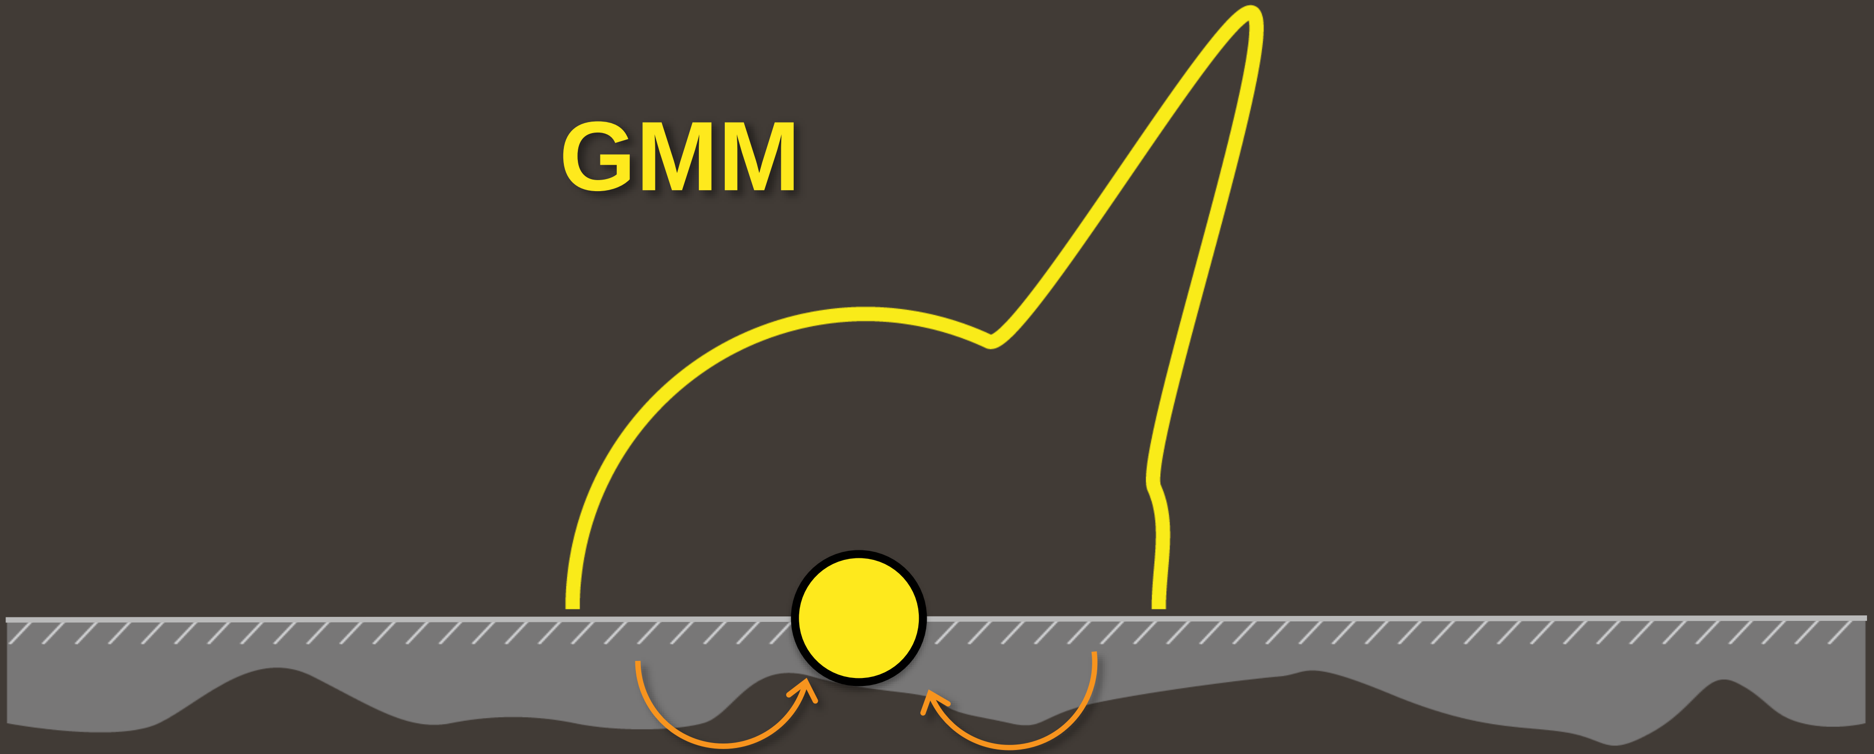

Overcoming the memory constraint

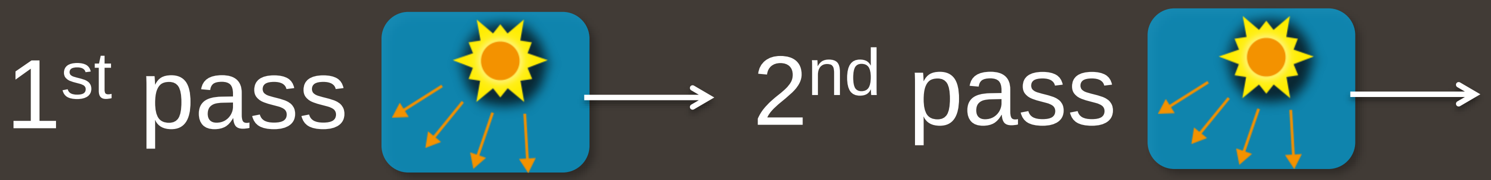

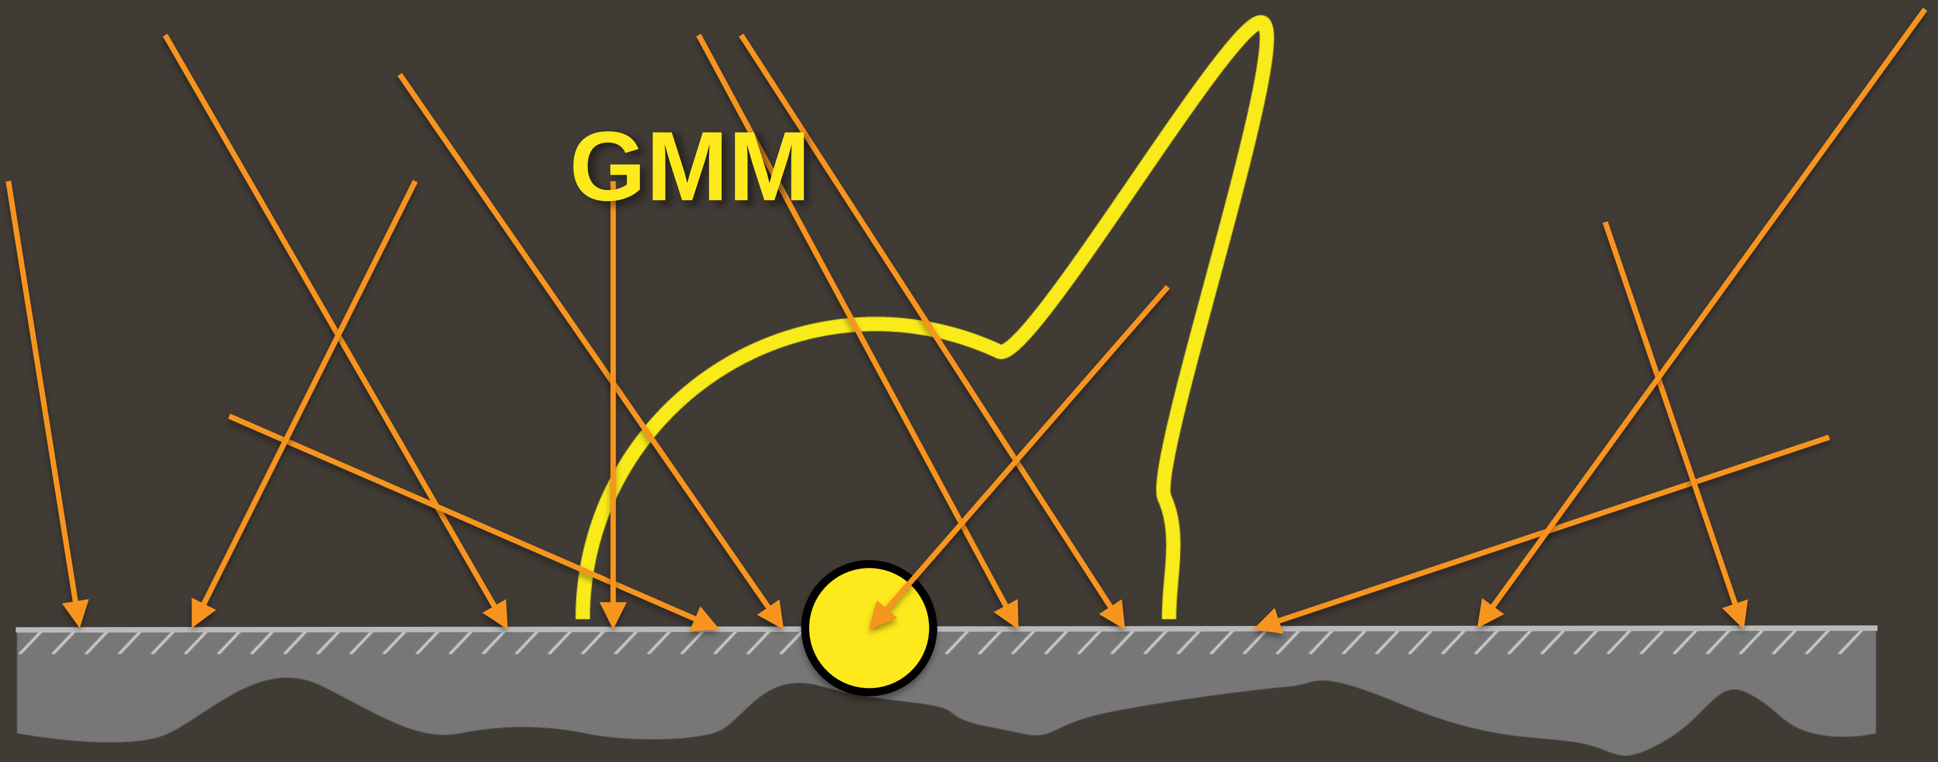

Overcoming the memory constraint

1st pass 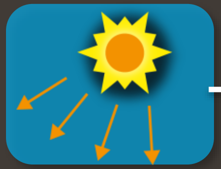 → 2nd pass 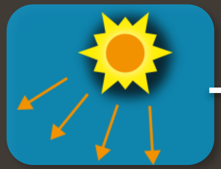 →

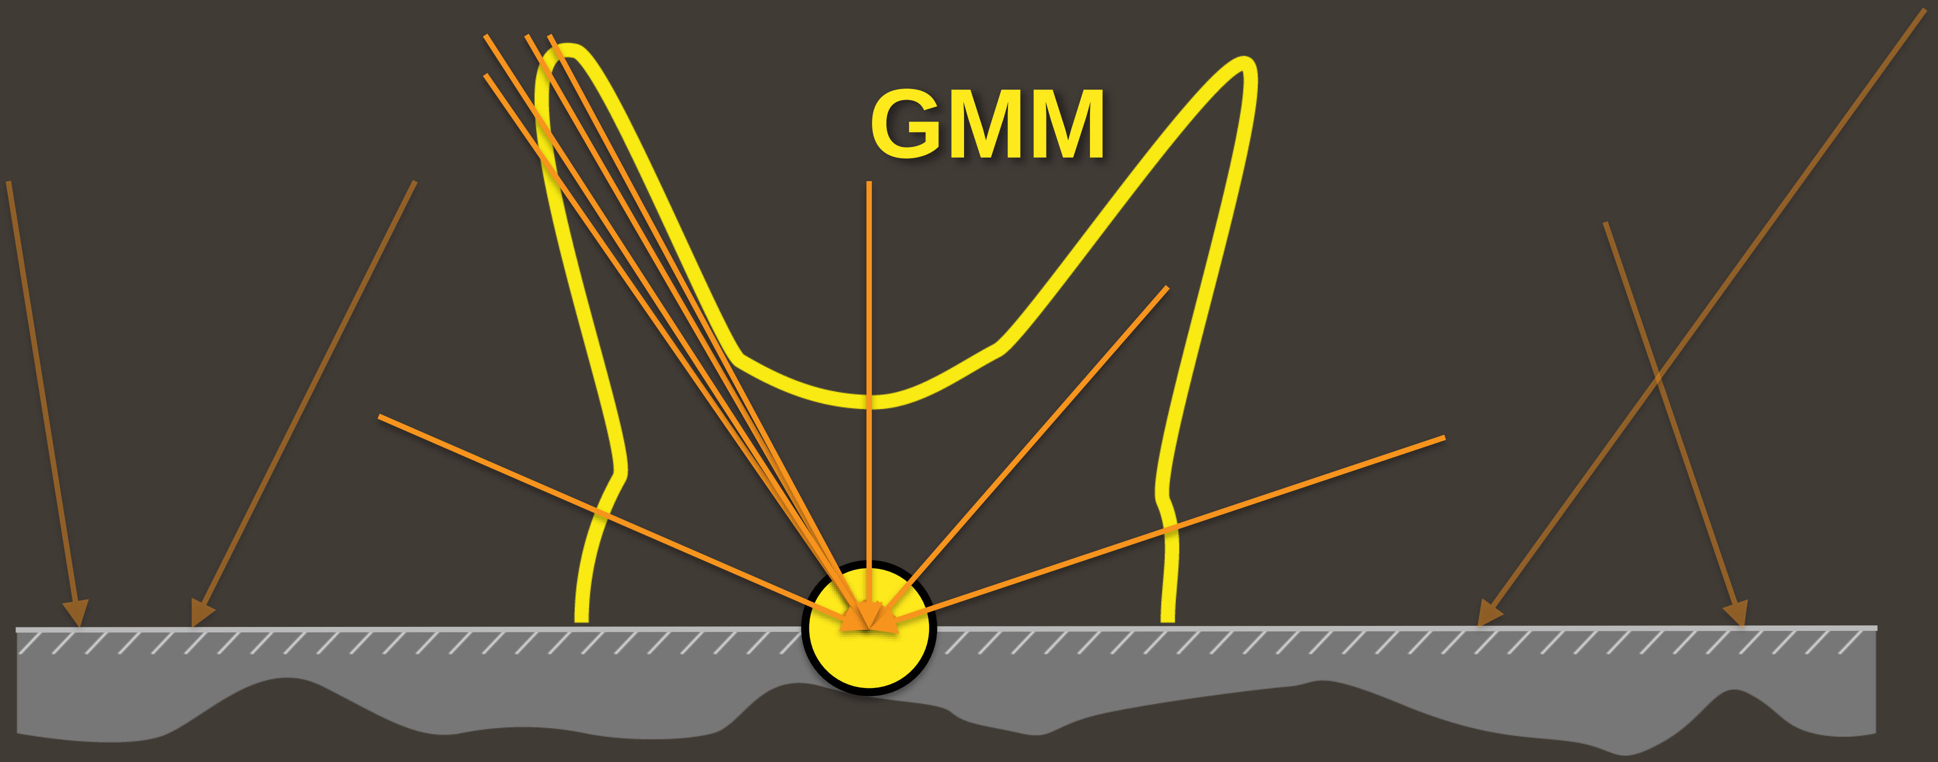

Overcoming the memory constraint

1^{st} pass 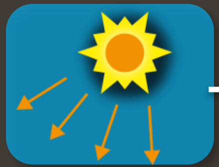 ⟶ 2^{nd} pass 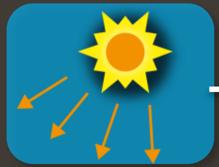 ⟶ 3^{rd} pass 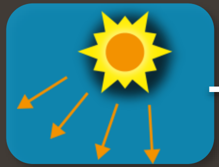 ⟶ ...

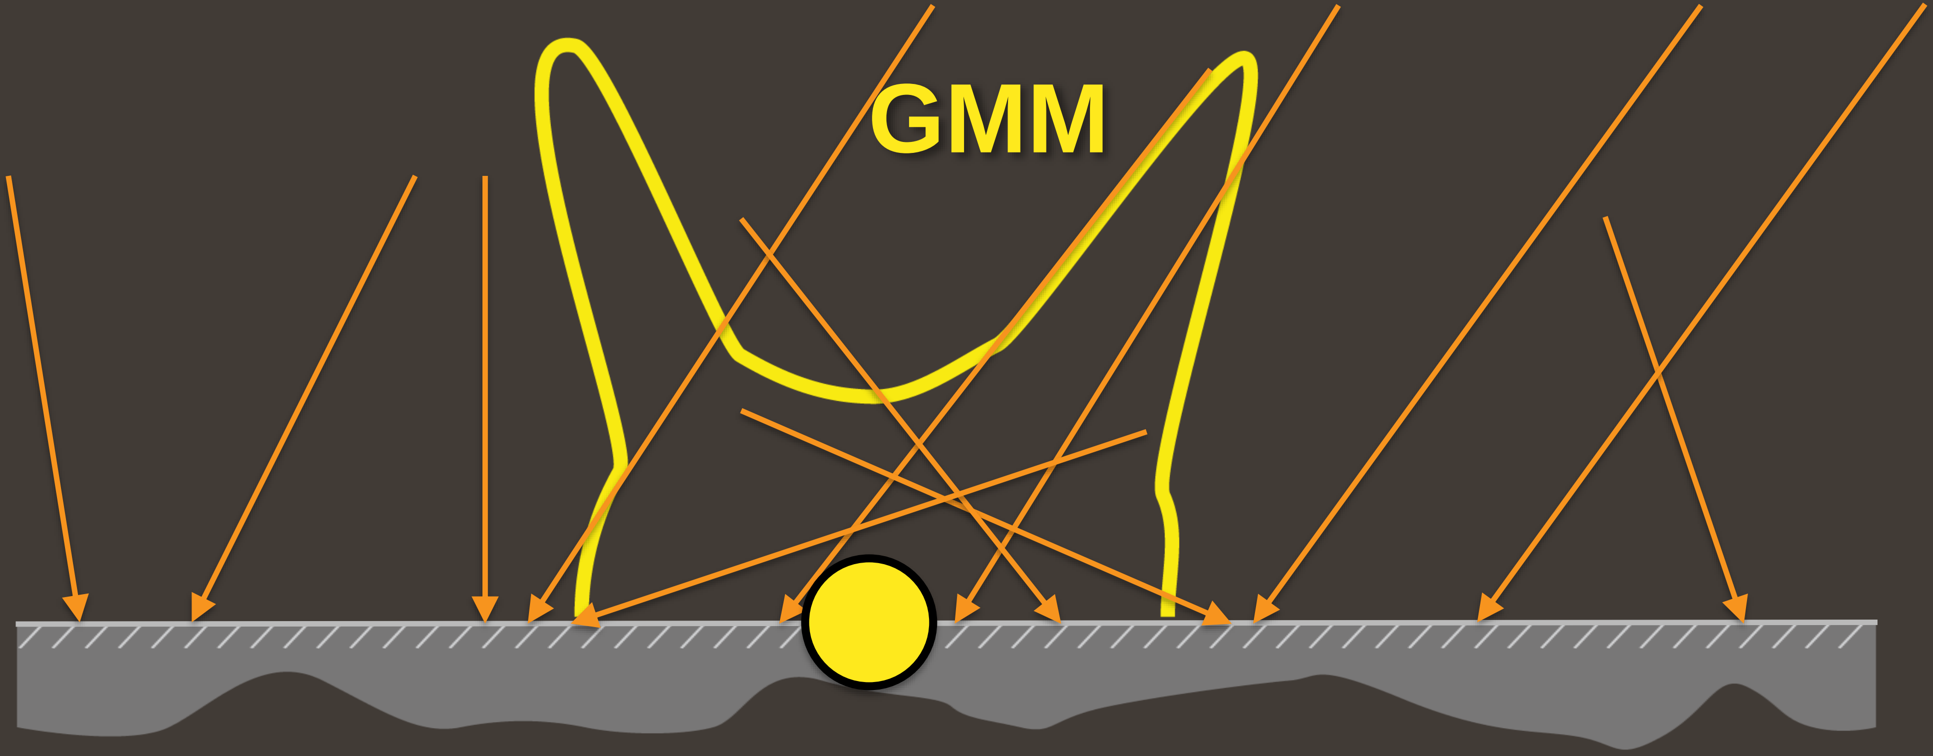

Overcoming the memory constraint

1st pass 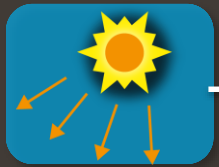 → 2nd pass 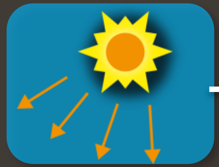 → 3rd pass 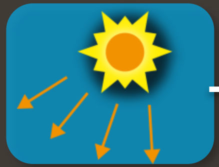 → ...

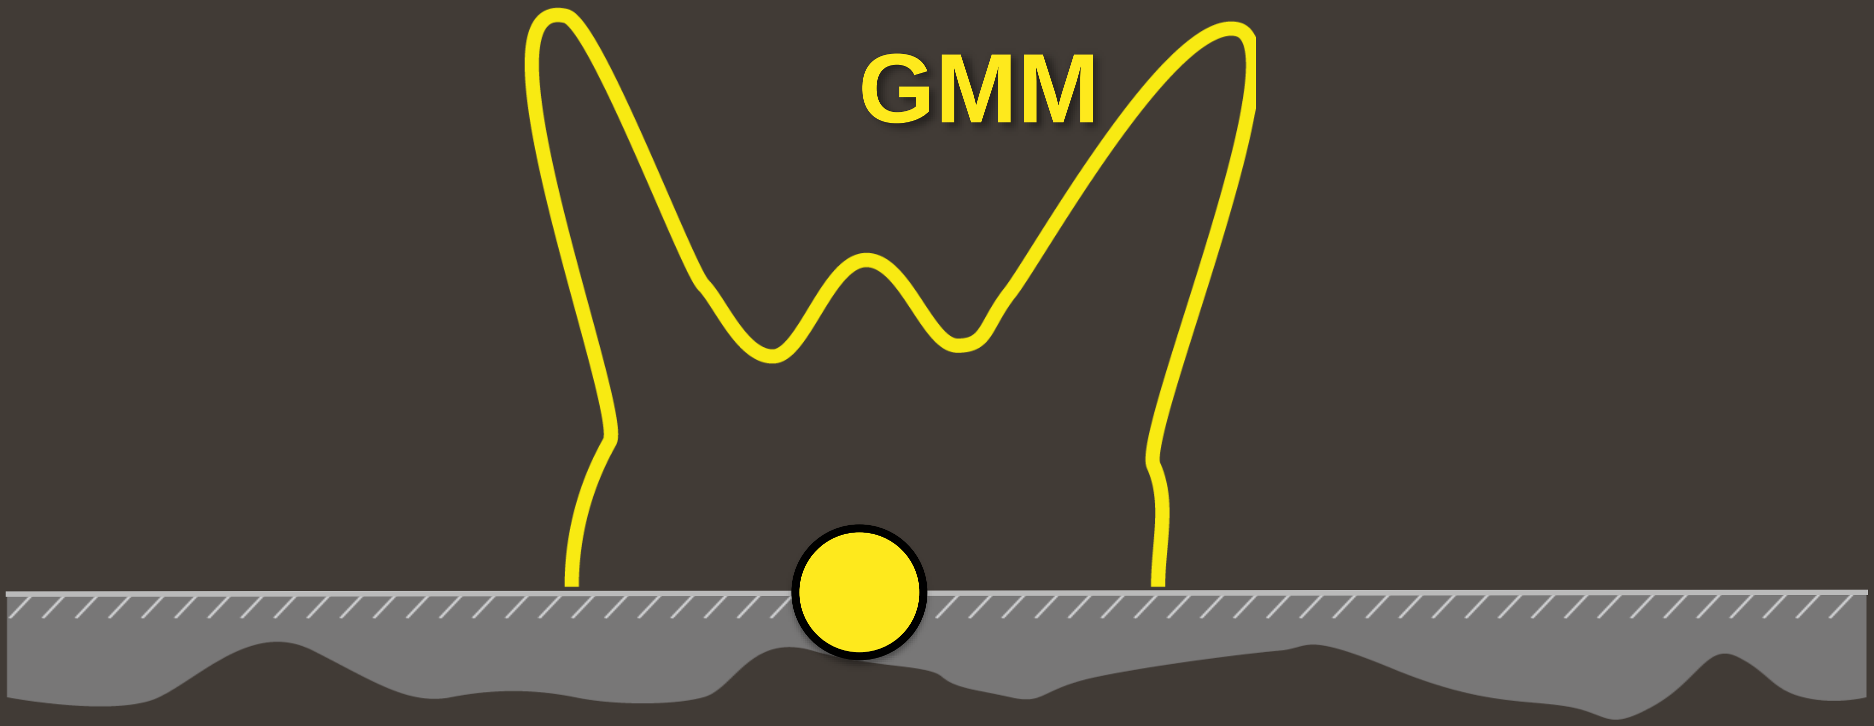

GM: superior estimate

On-line stepwise Expectation-Maximization

[Cappé & Moulines 2009]

Input: **an infinite stream of particles**

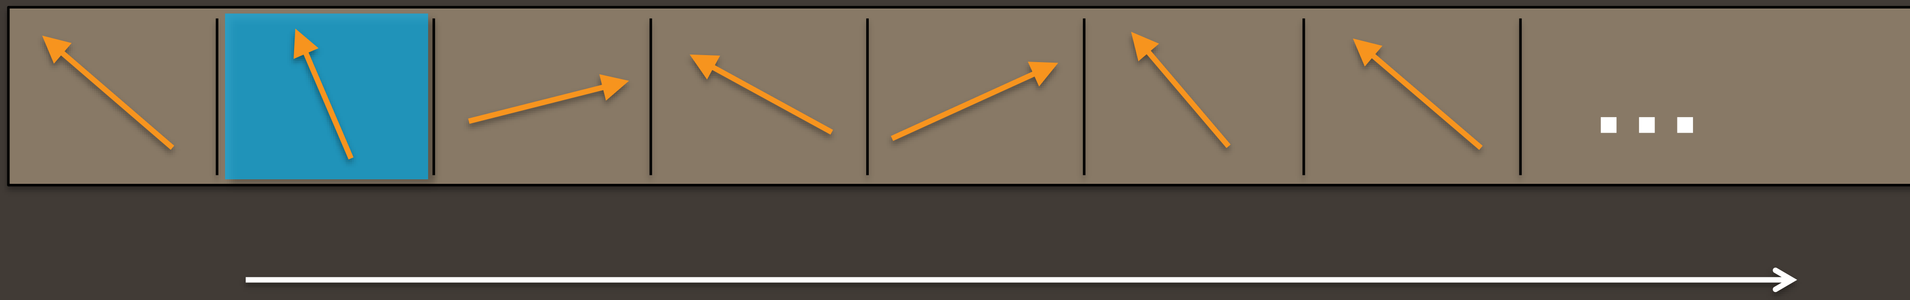

Method outline

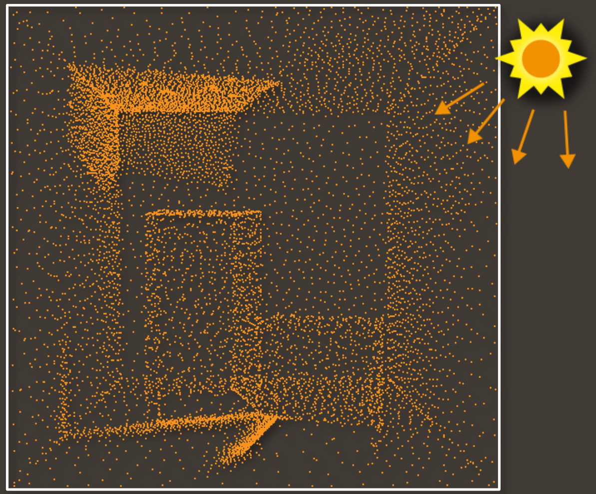

training

rendering

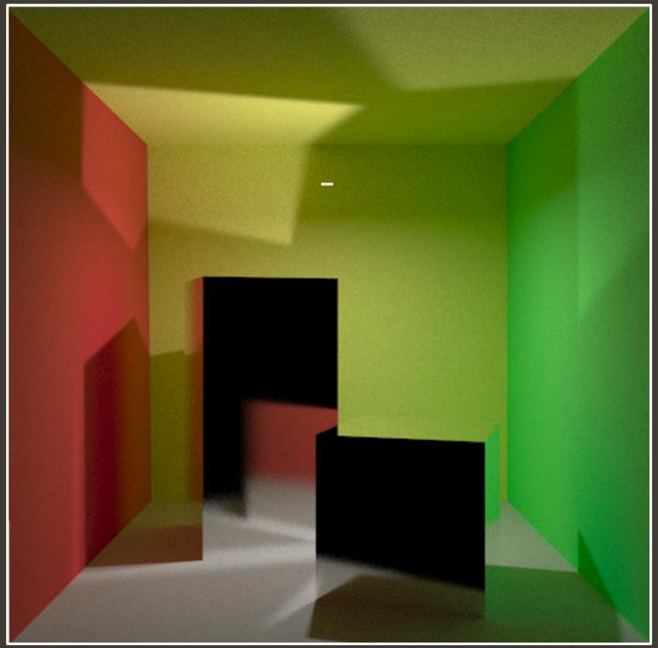

Guided path sampling

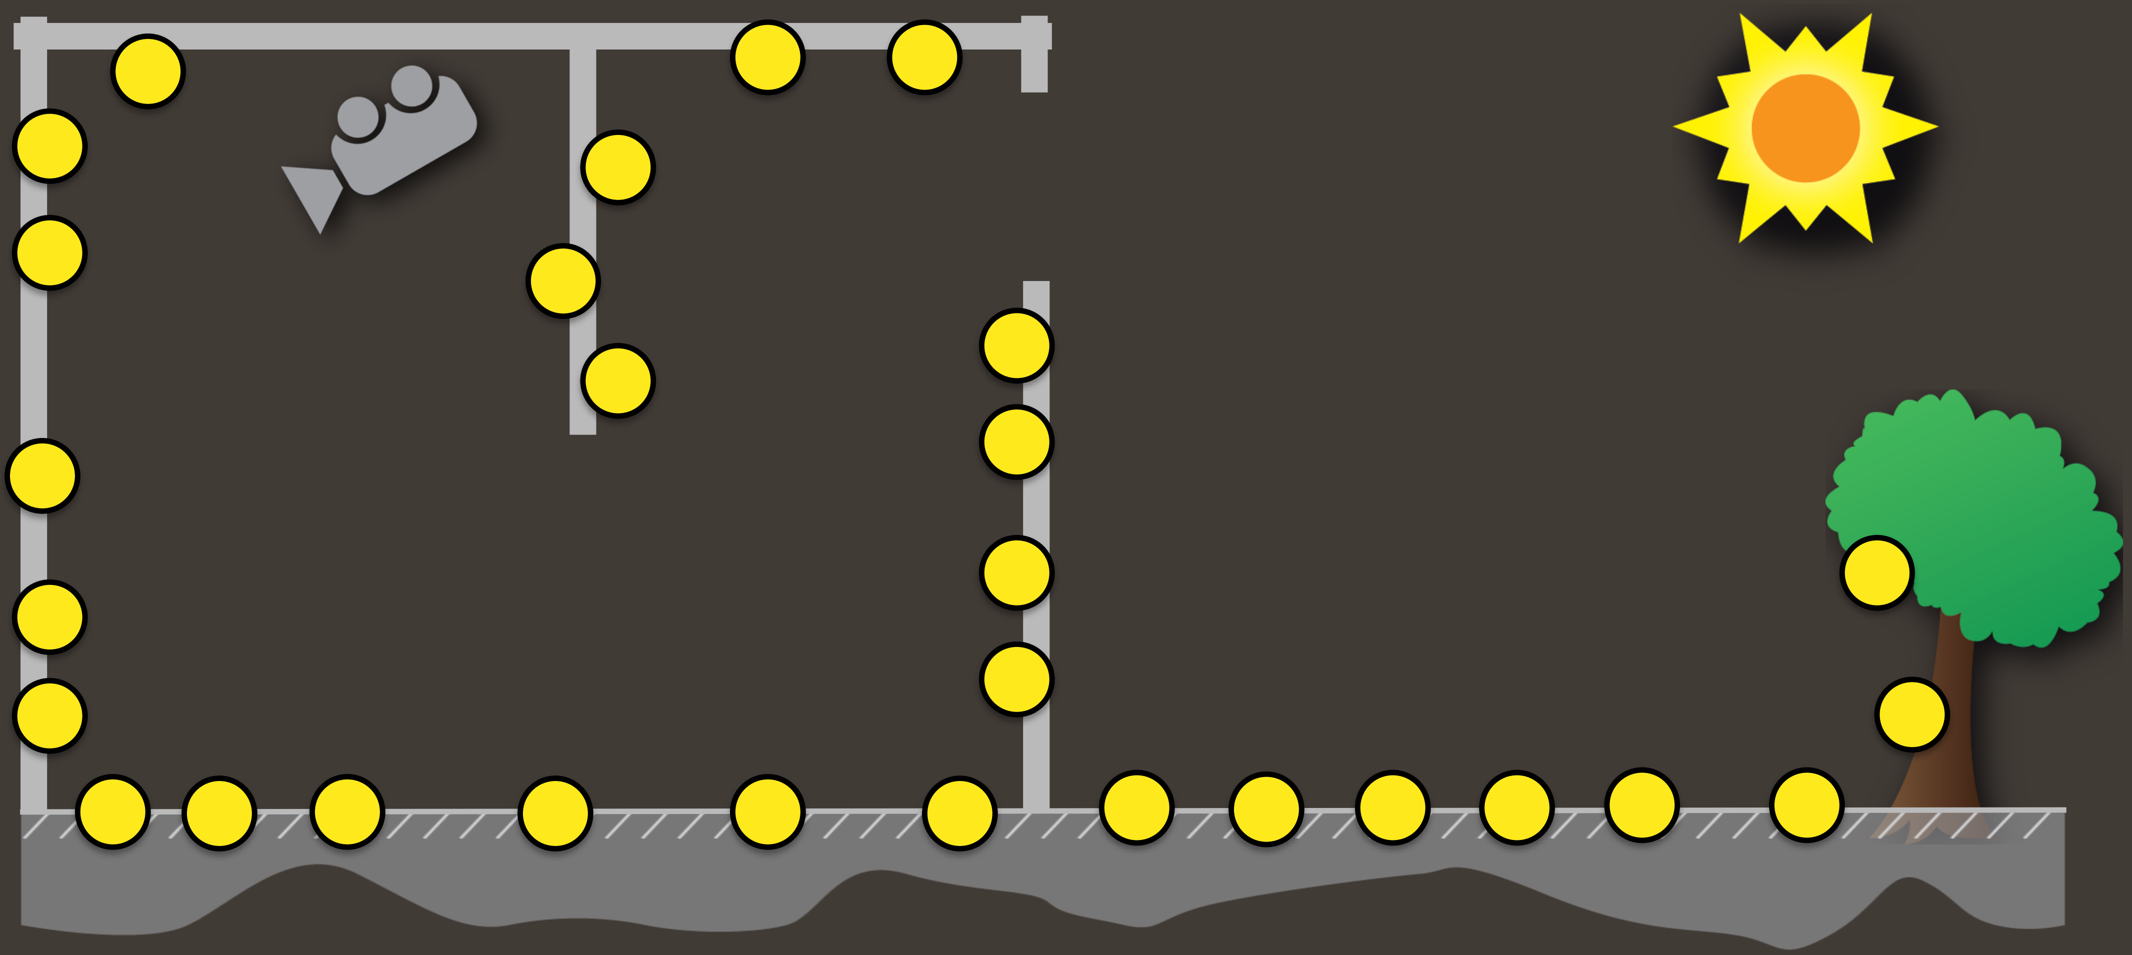

Guided path sampling

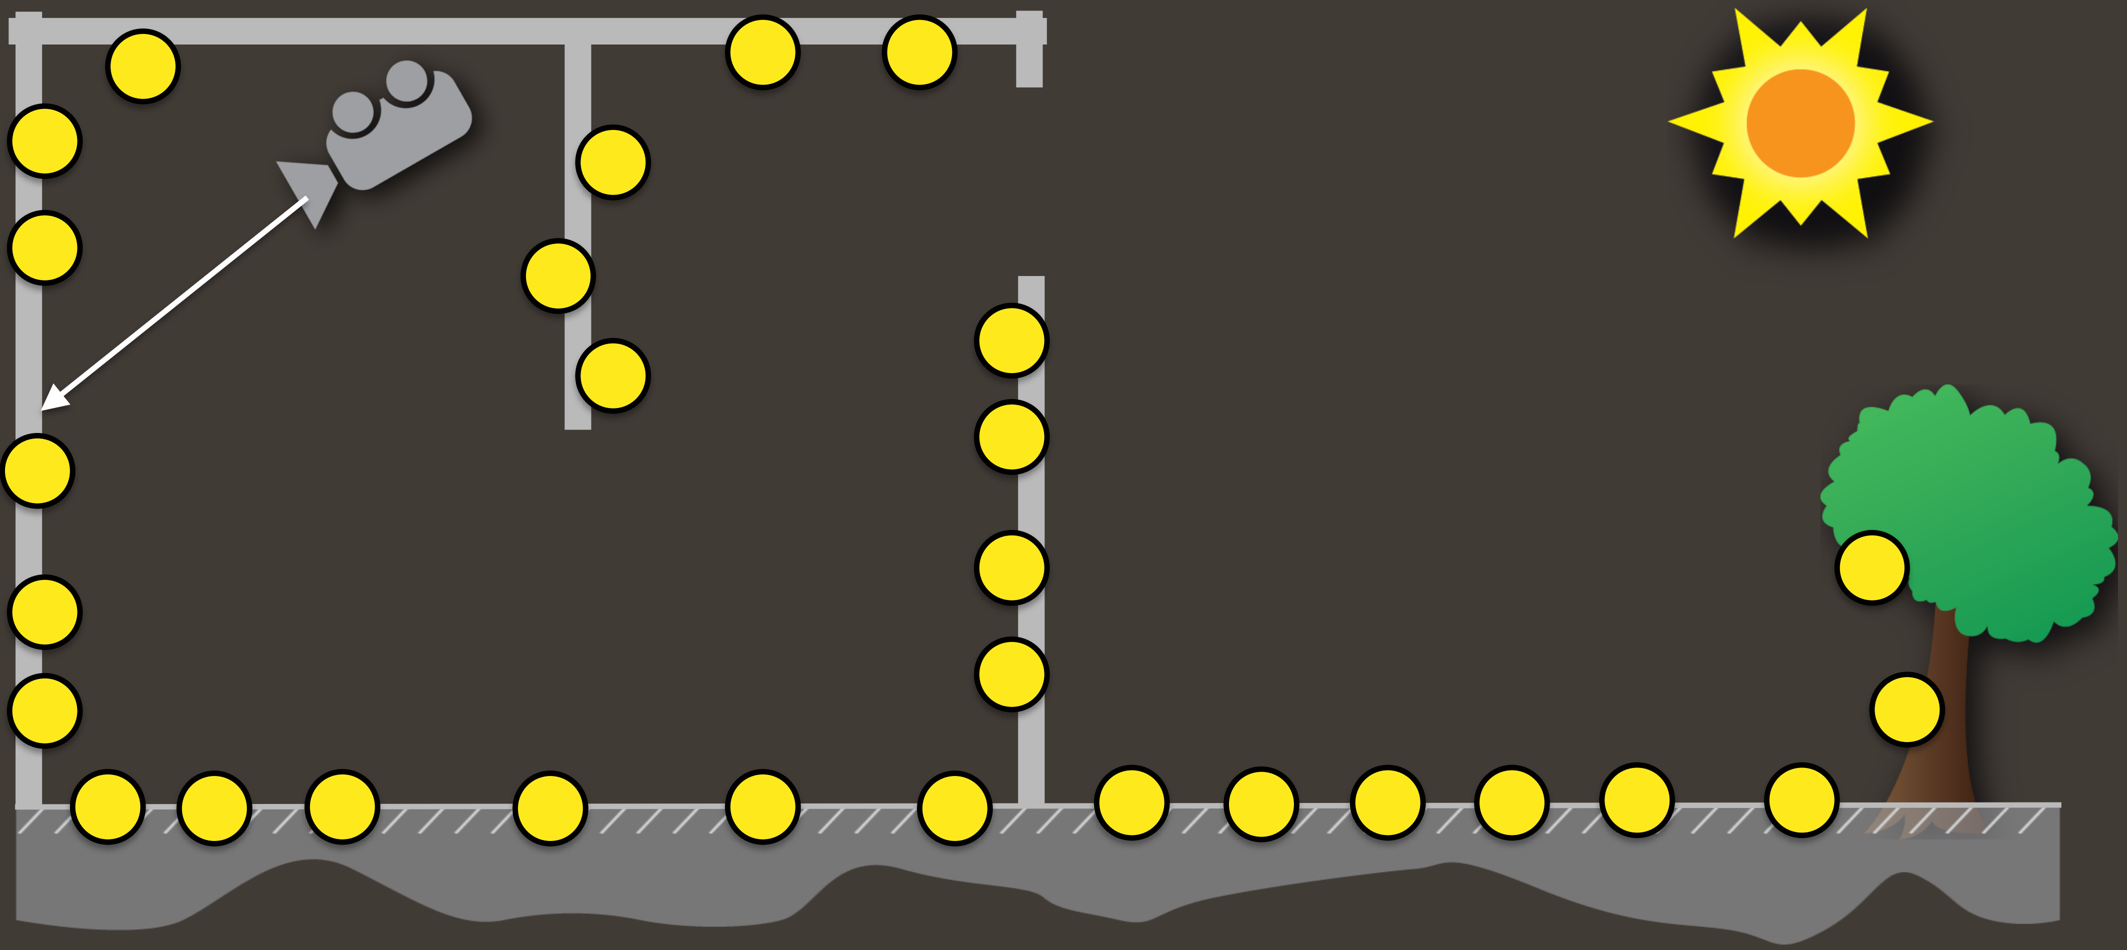

Guided path sampling

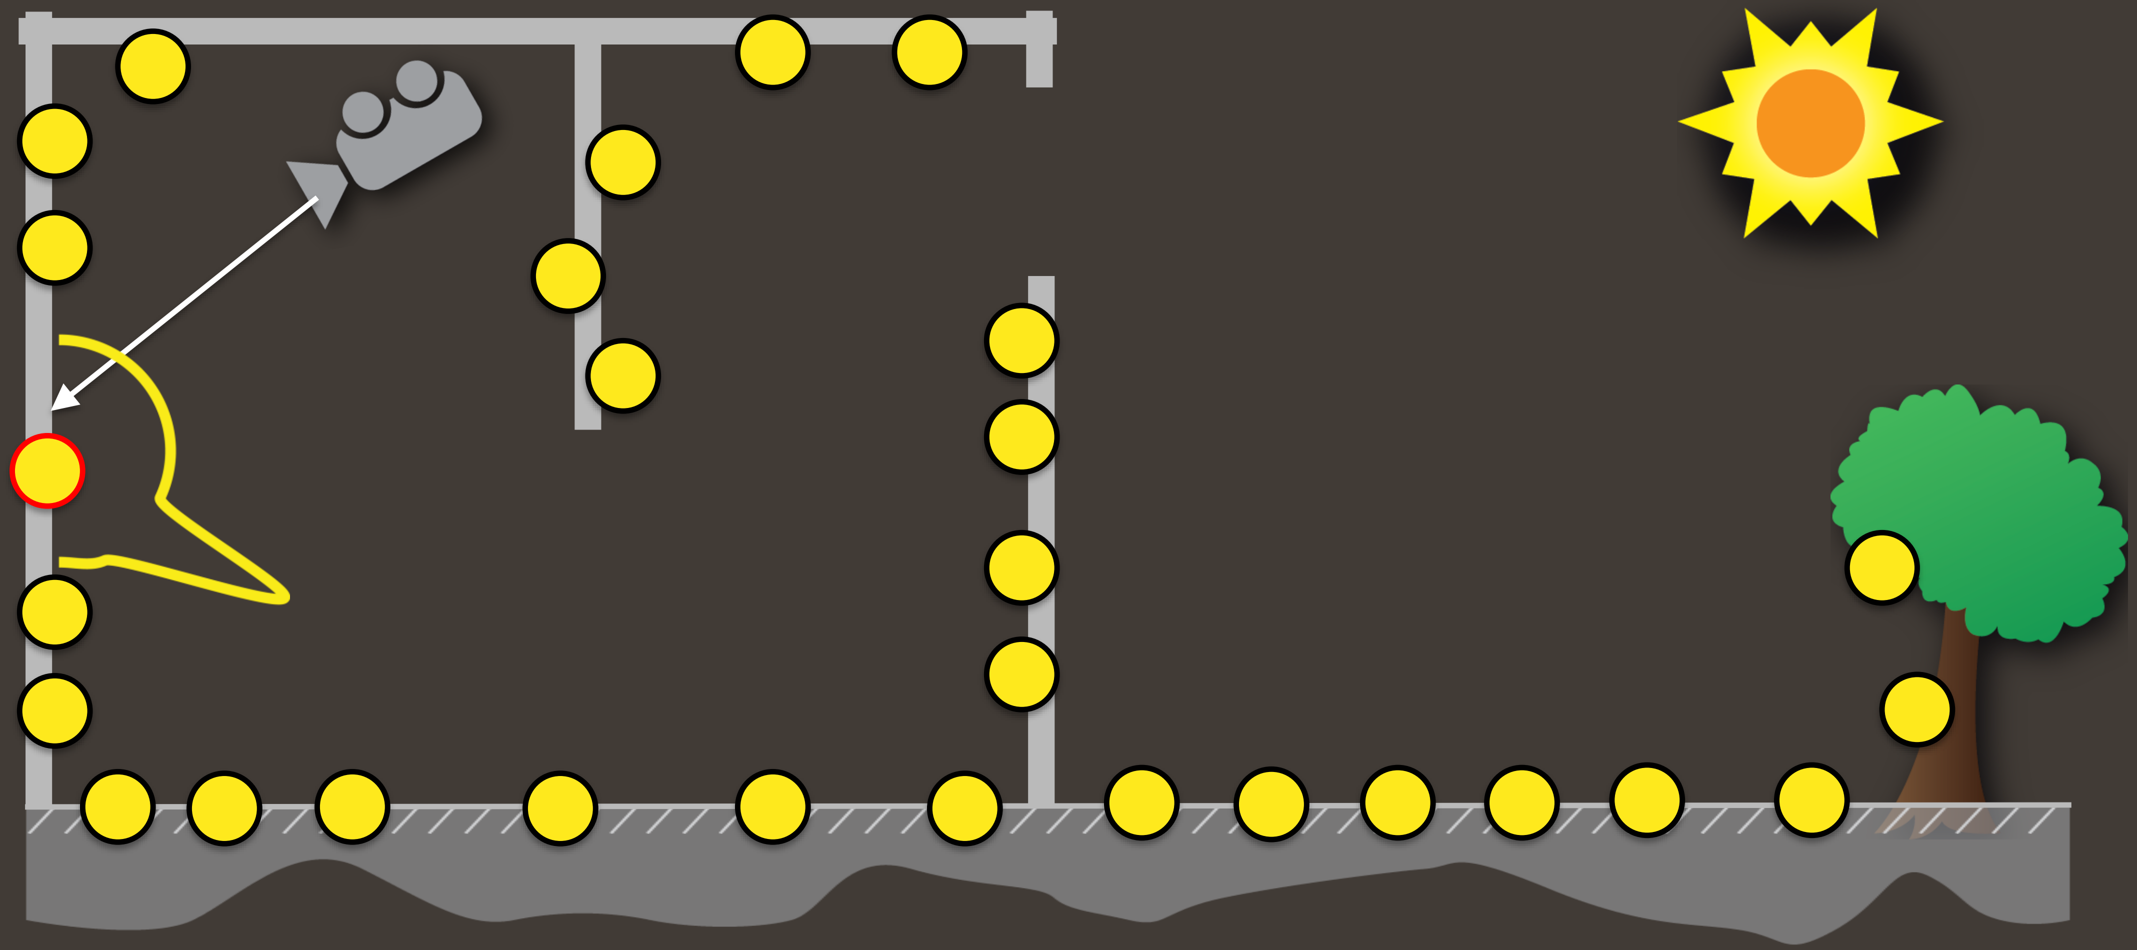

Guided path sampling

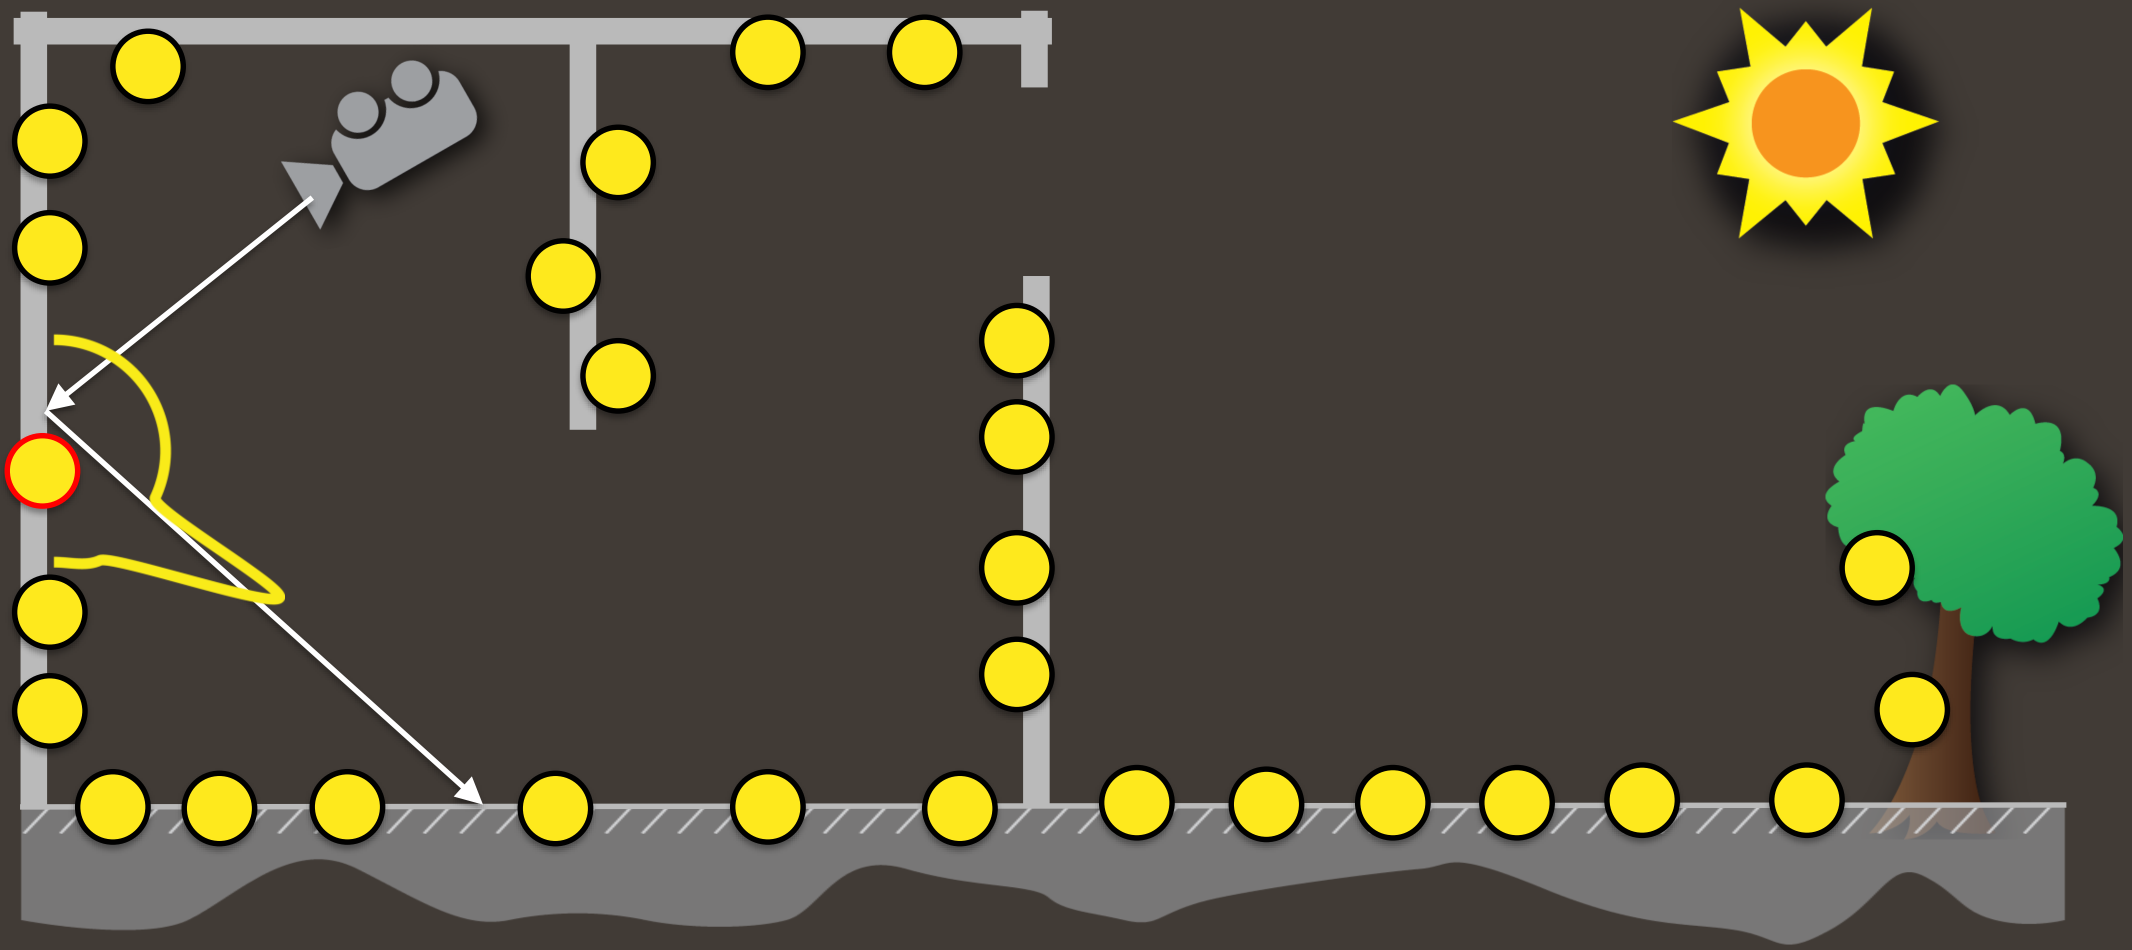

Guided path sampling

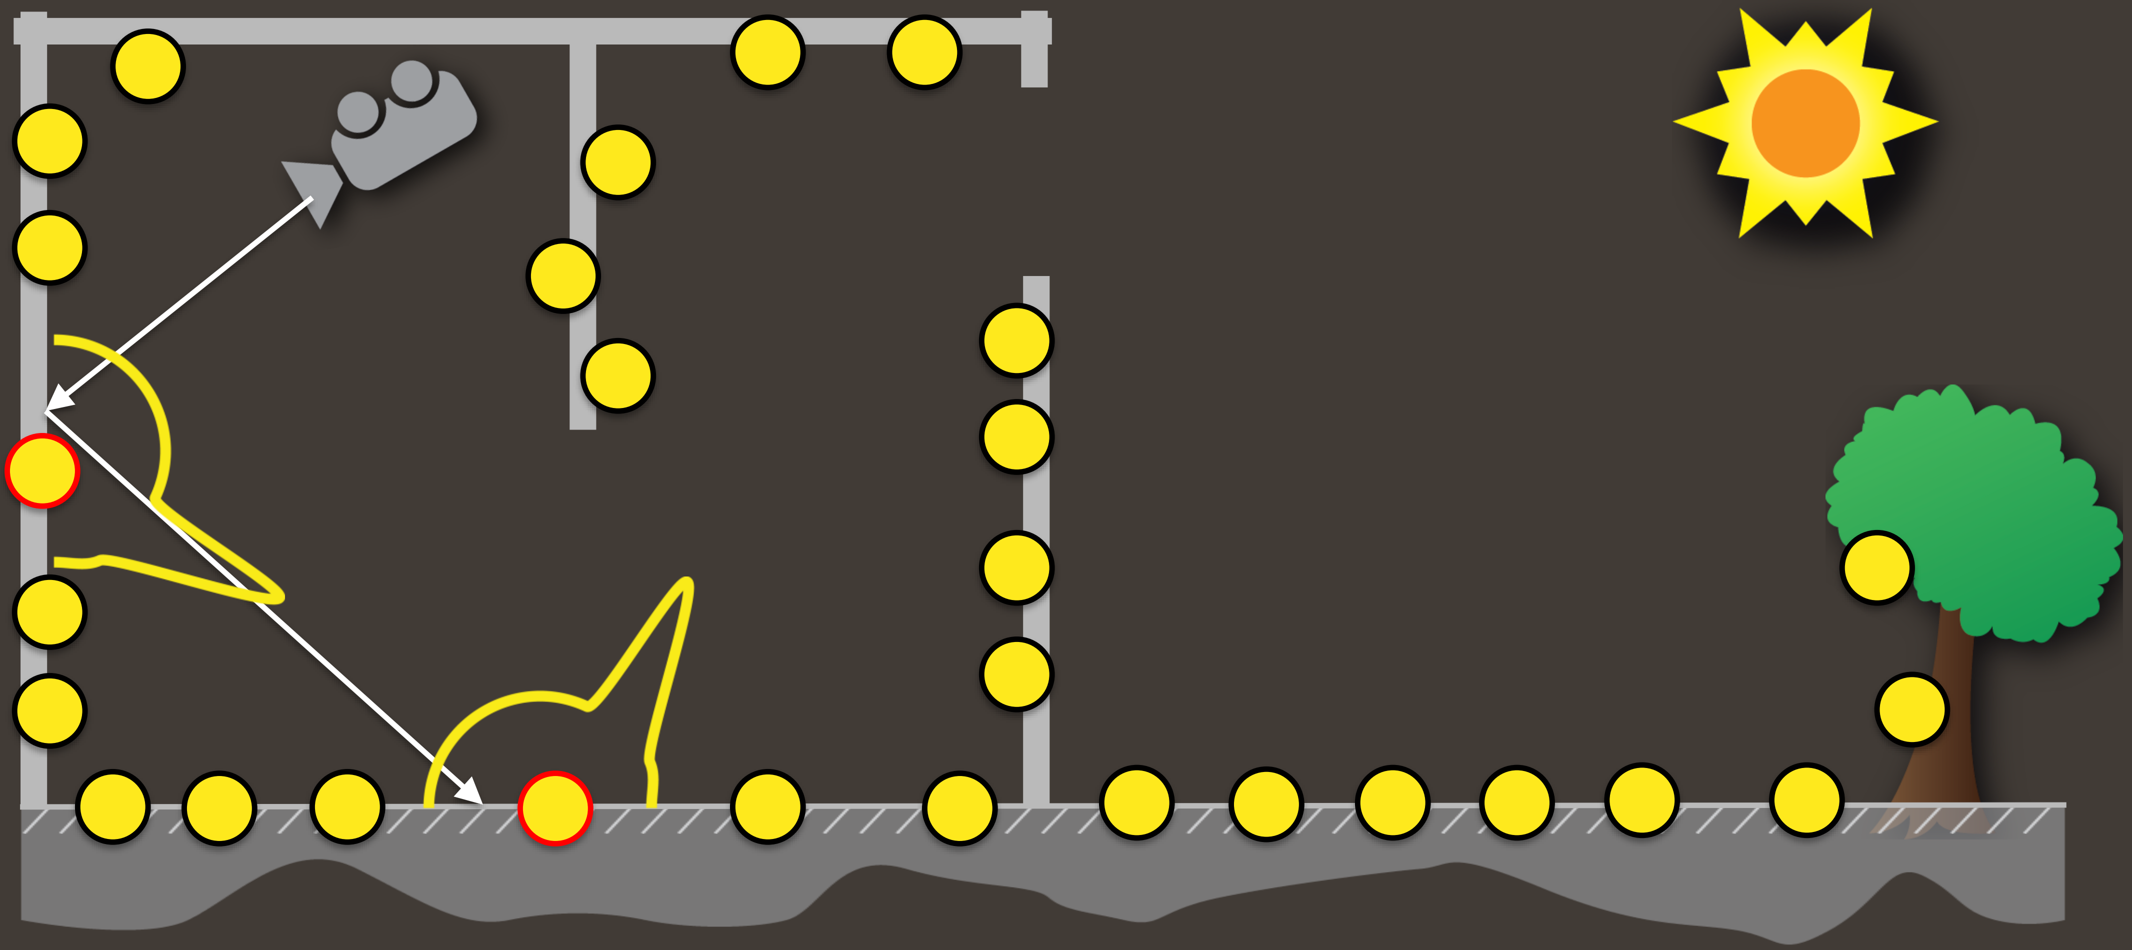

Guided path sampling

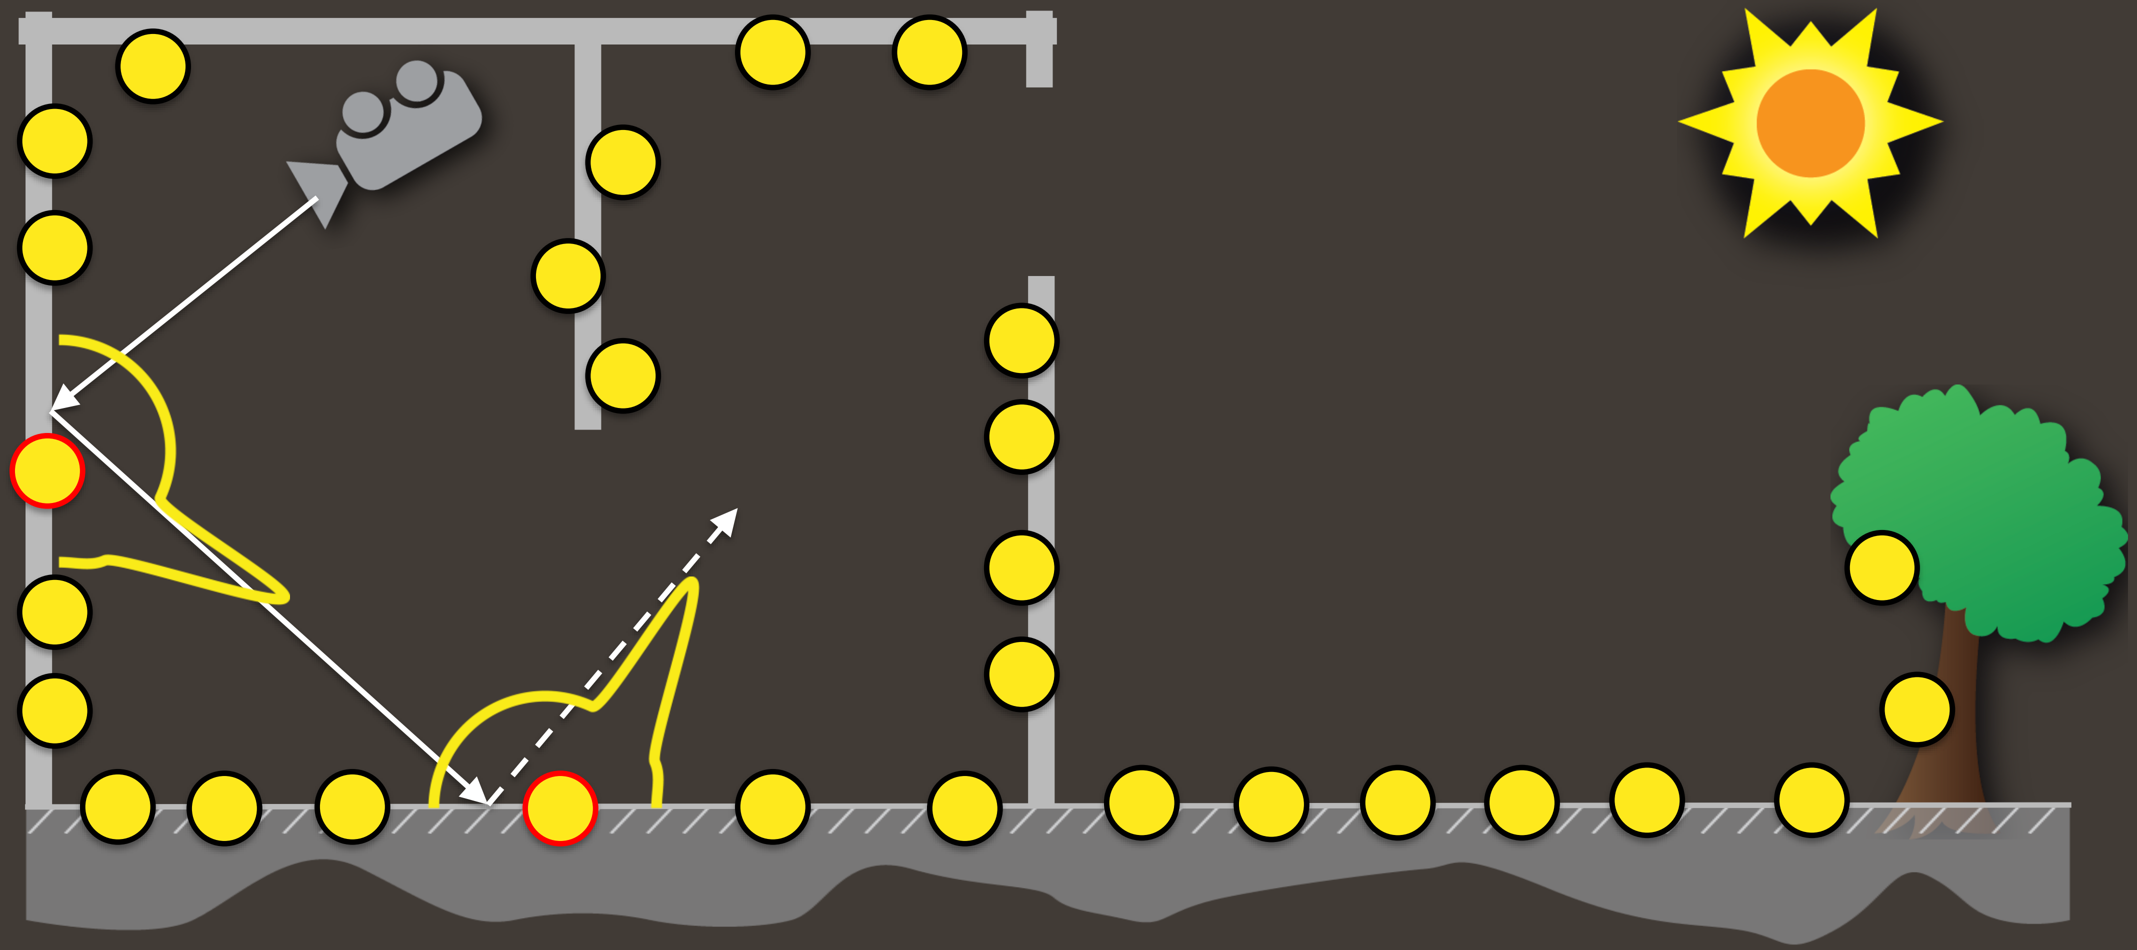

GUIDED PATH TERMINATION (RUSSIAN ROULETTE)

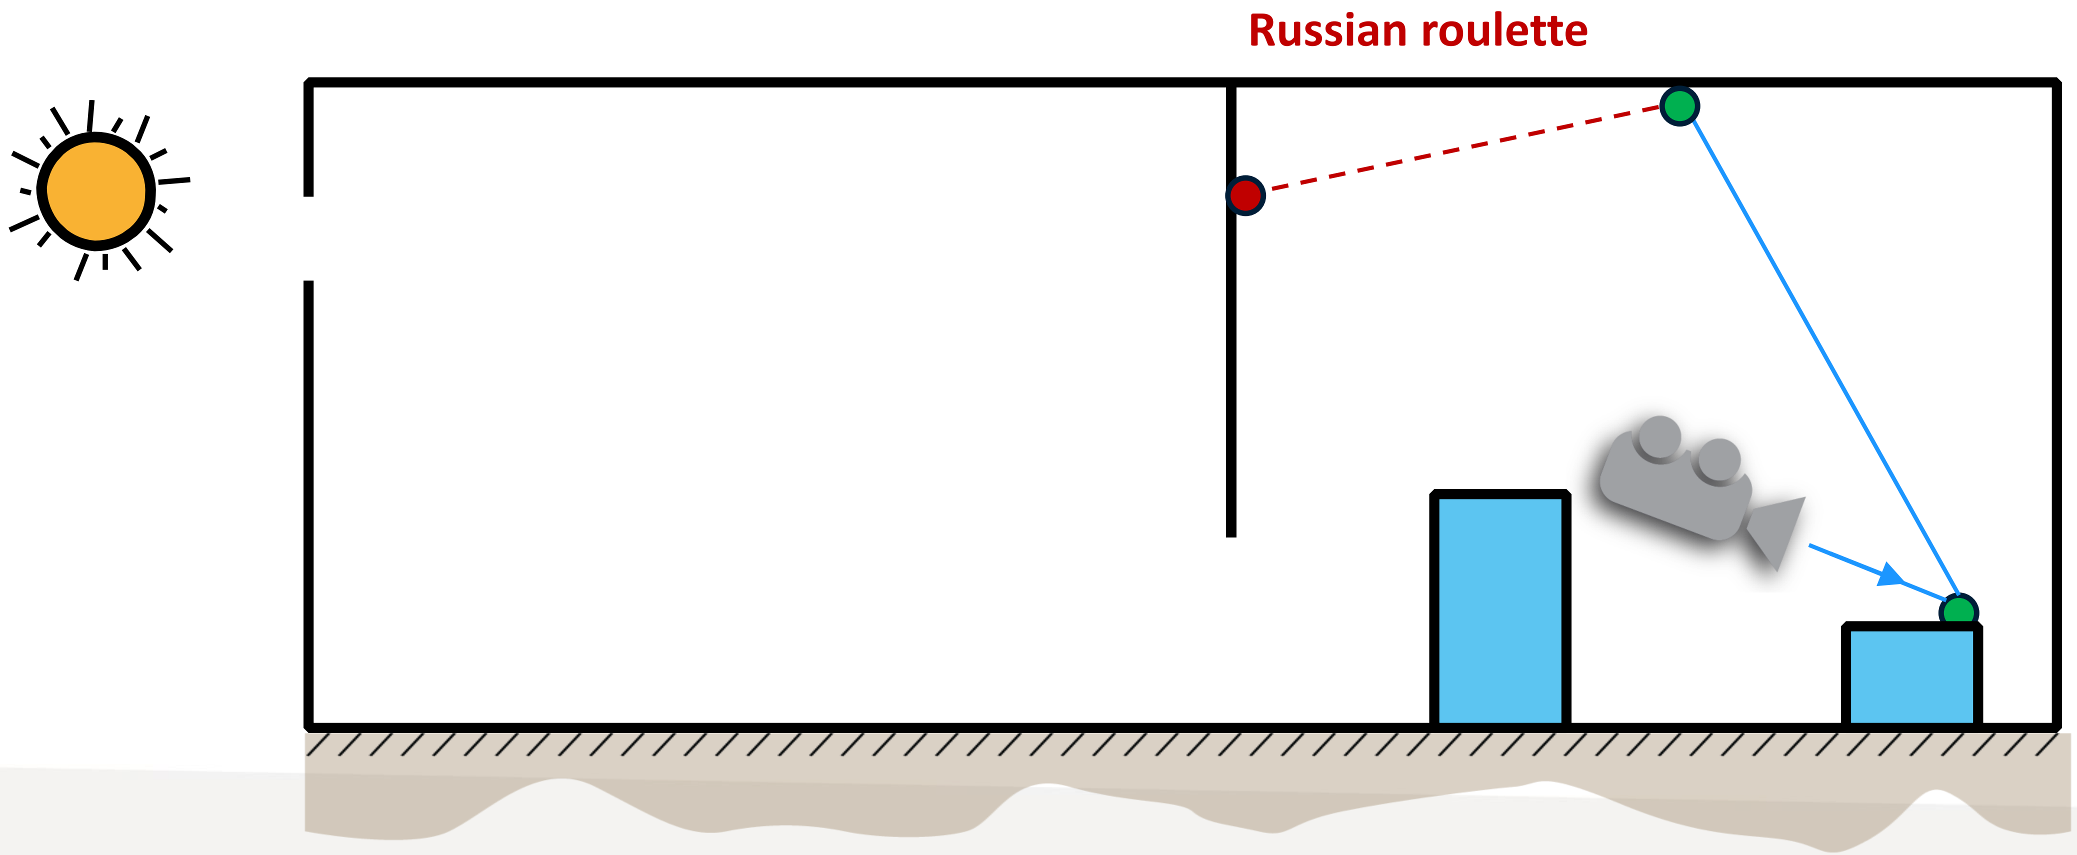

GUIDED SPLITTING

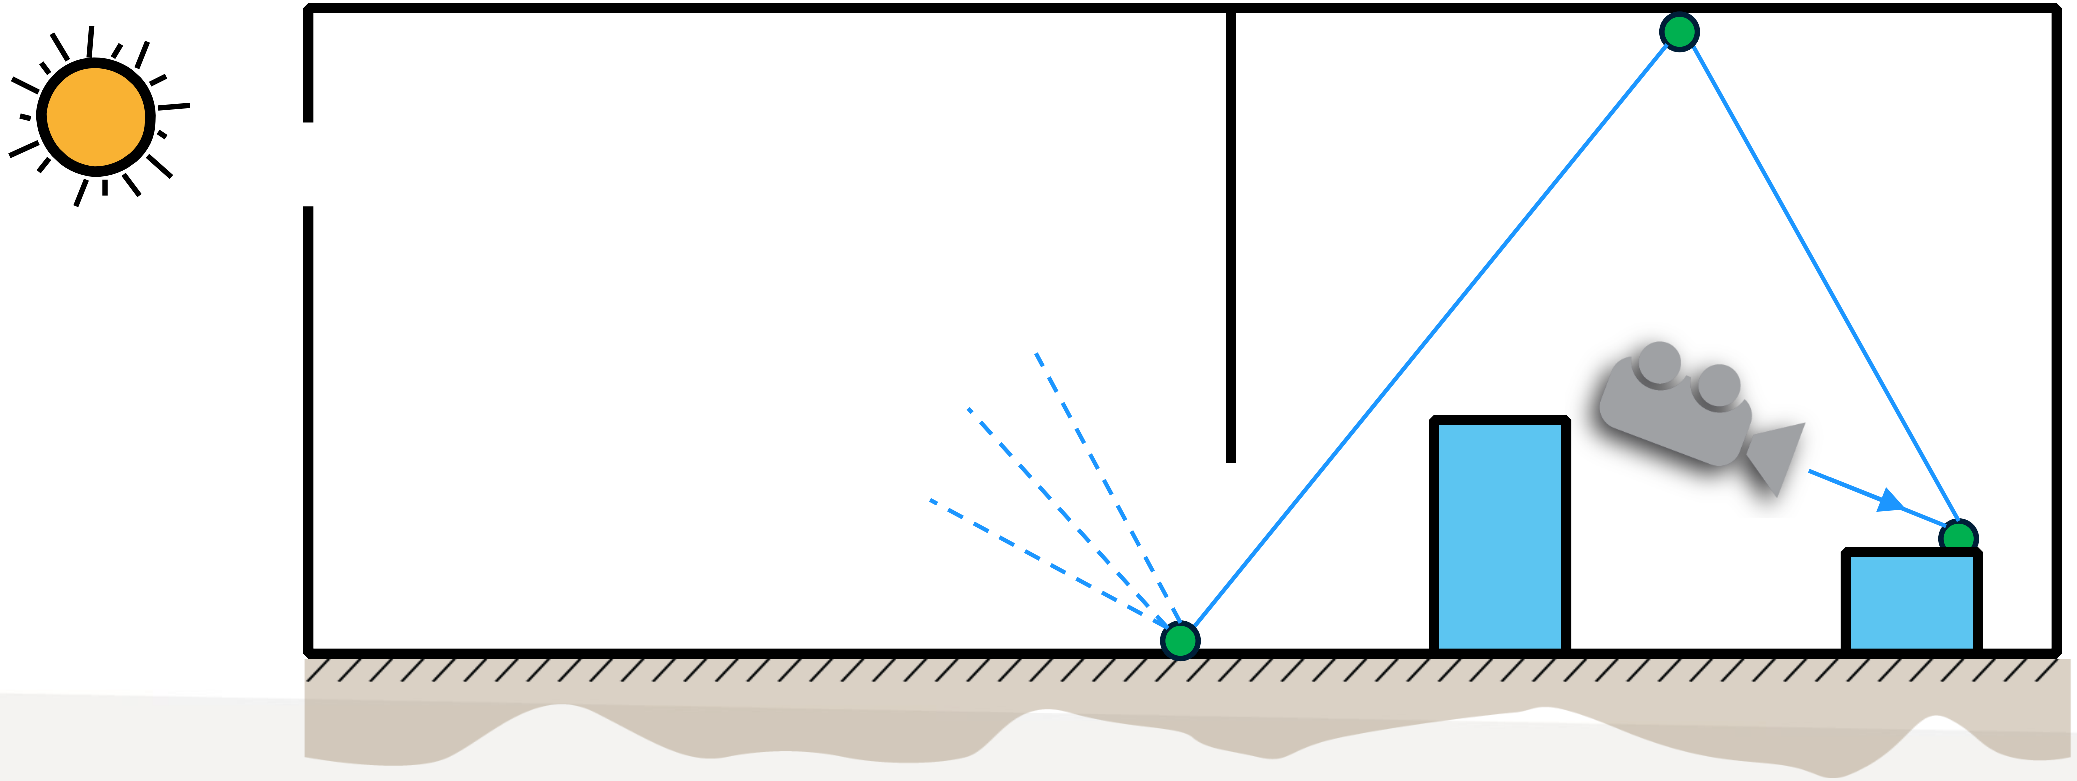

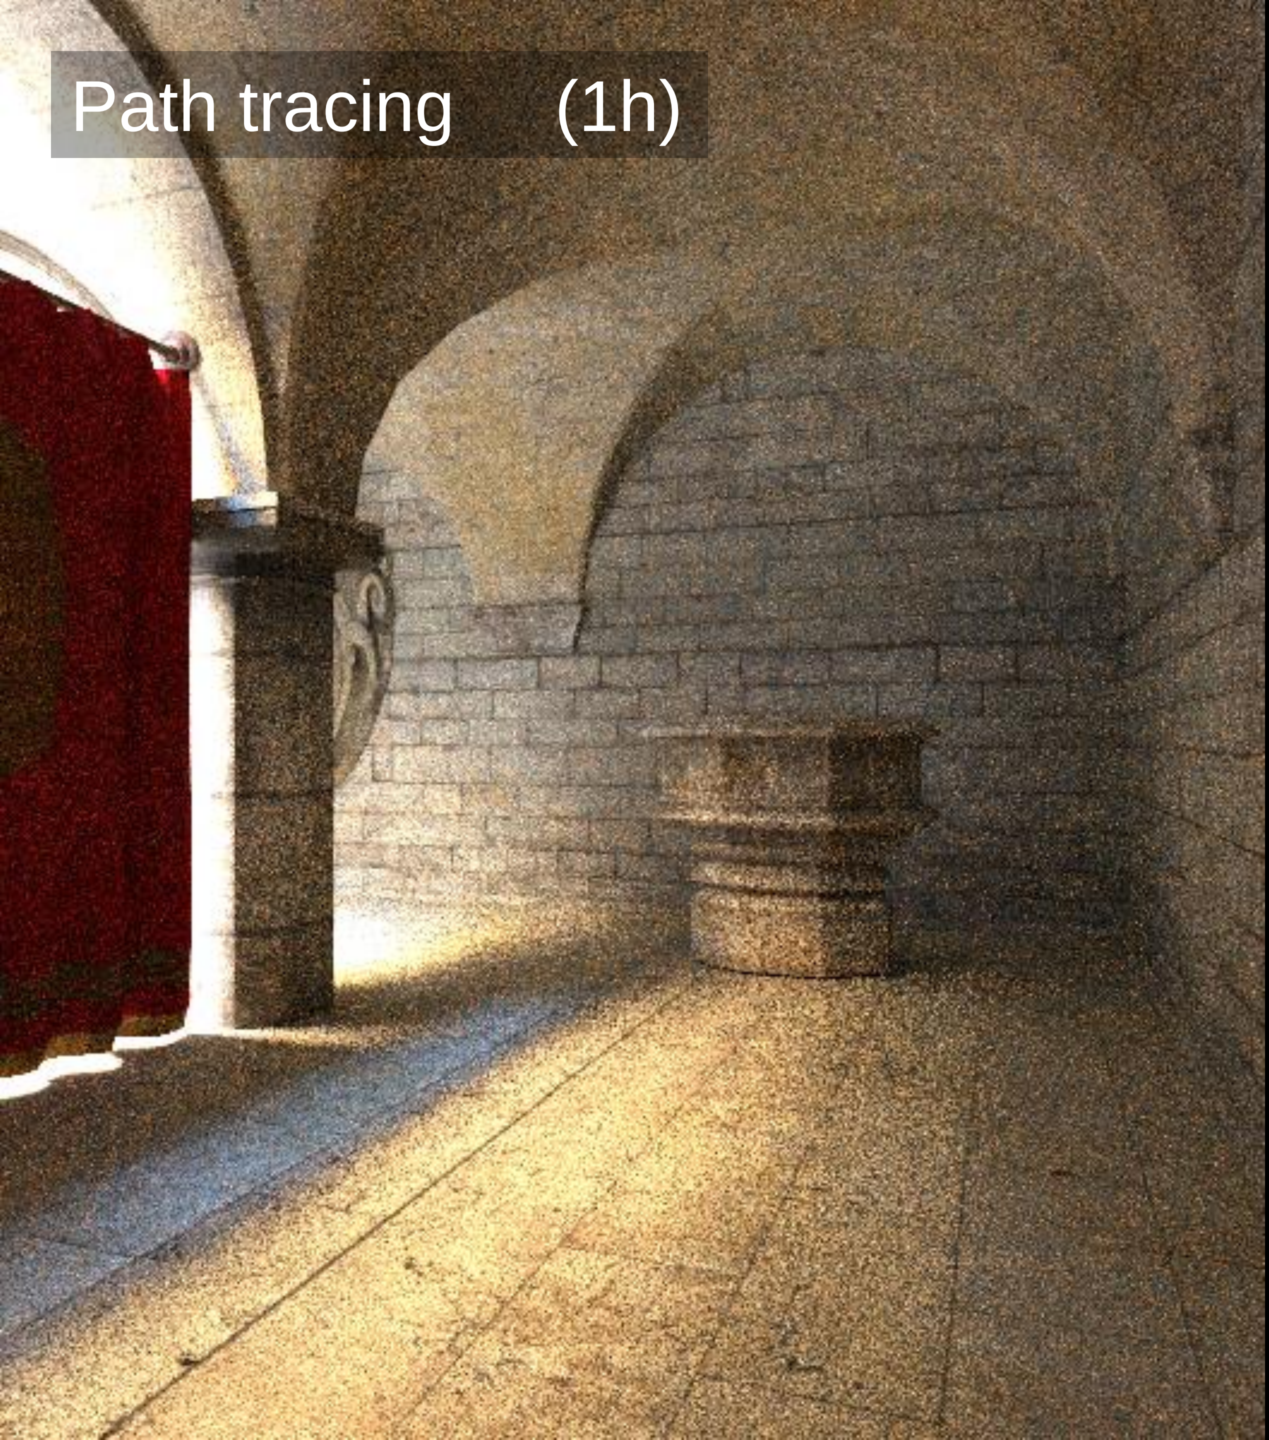
Path tracing (1h)

Guided path tracing (1h)

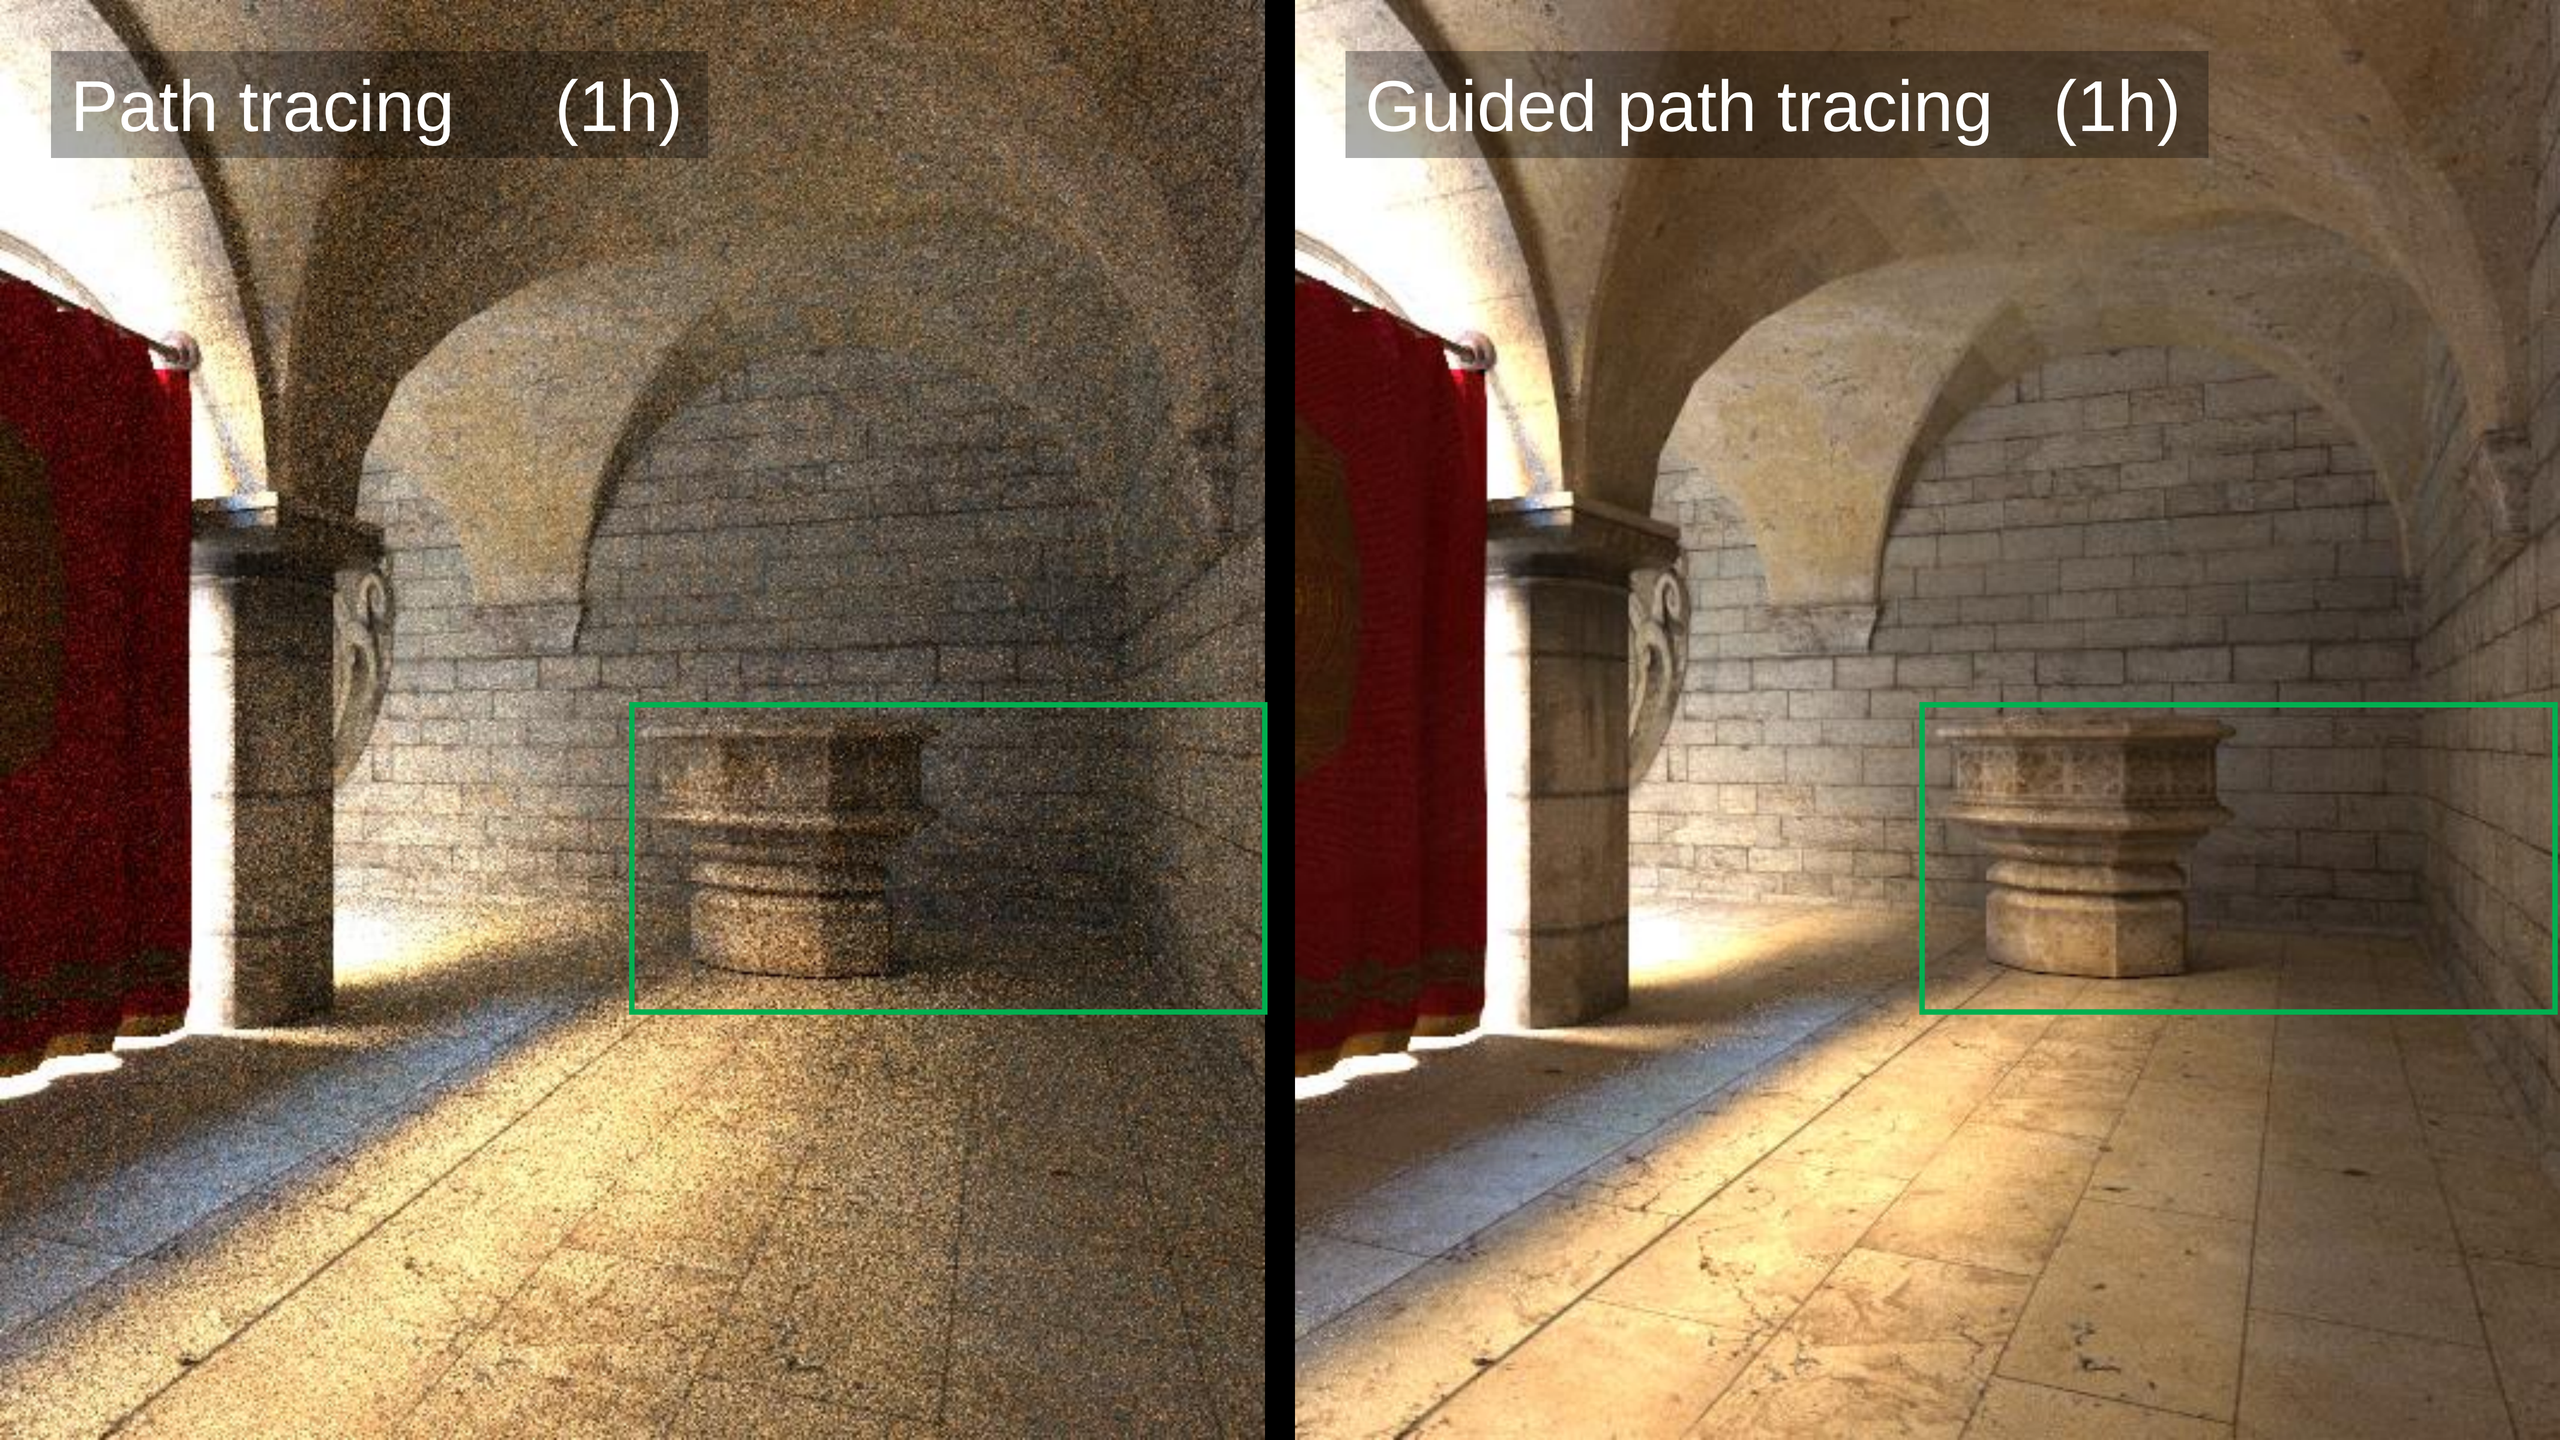
Path tracing (1h)
Guided path tracing (1h)

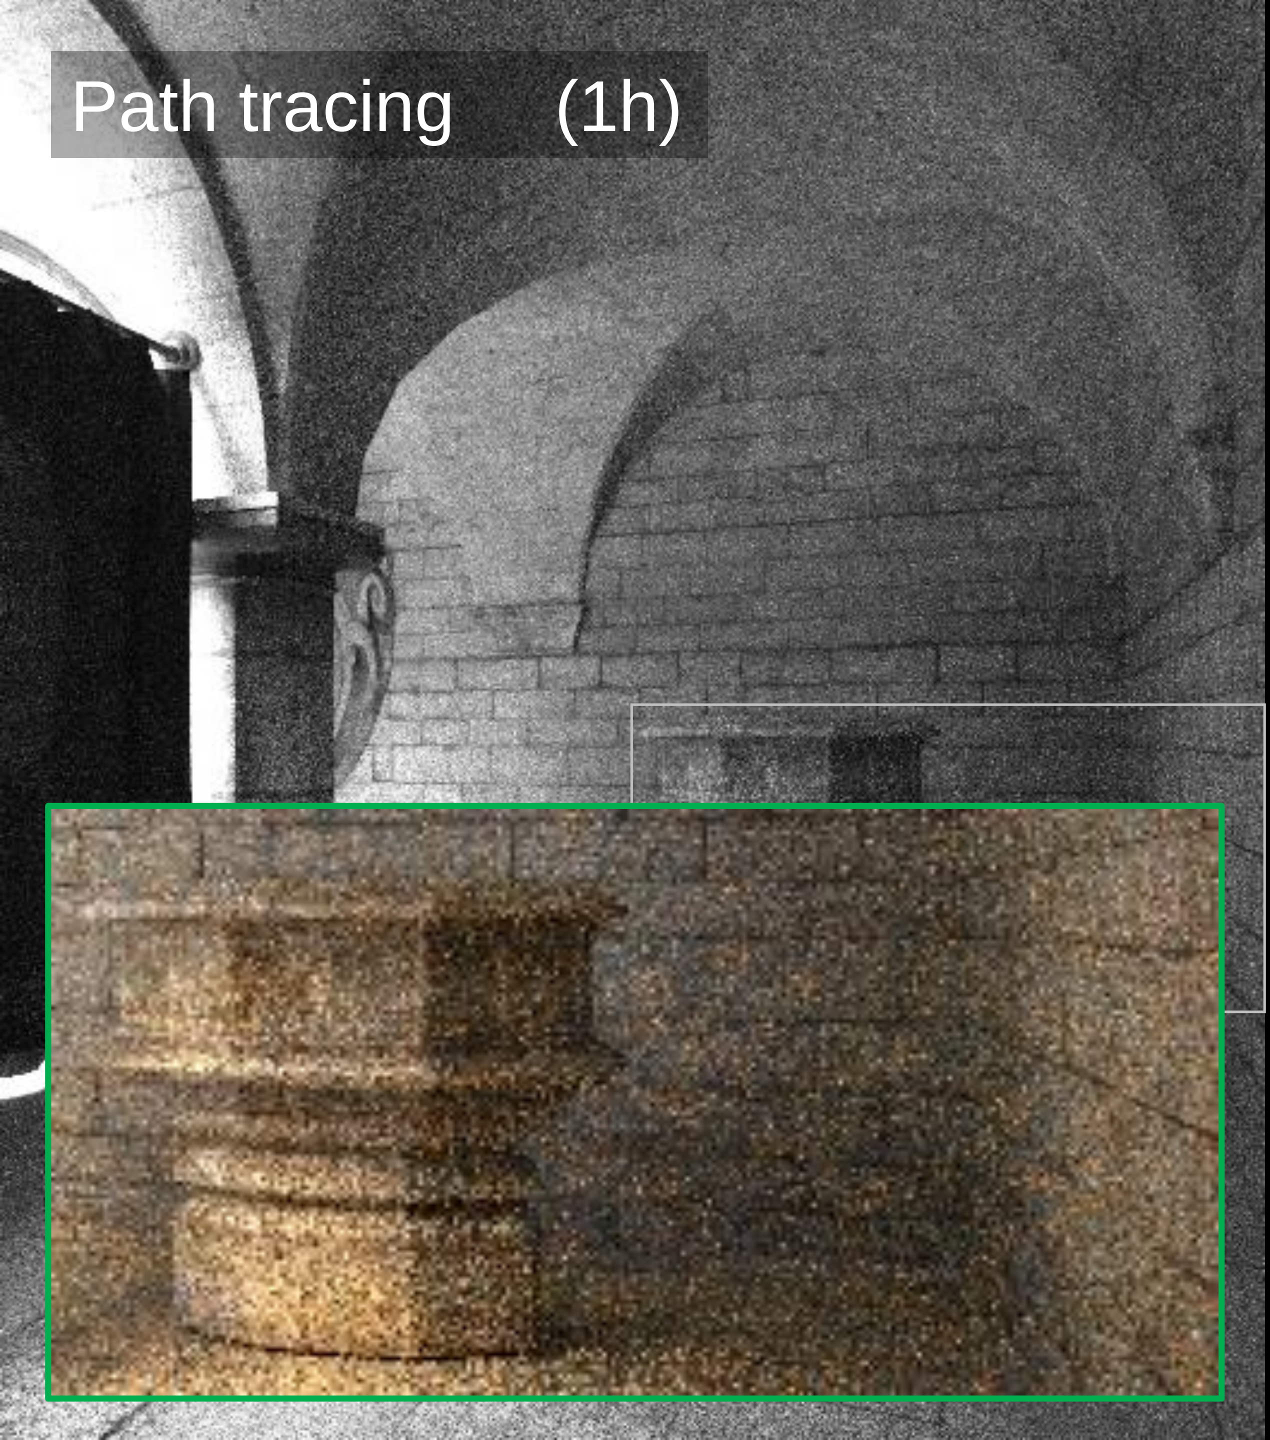
Path tracing (1h)
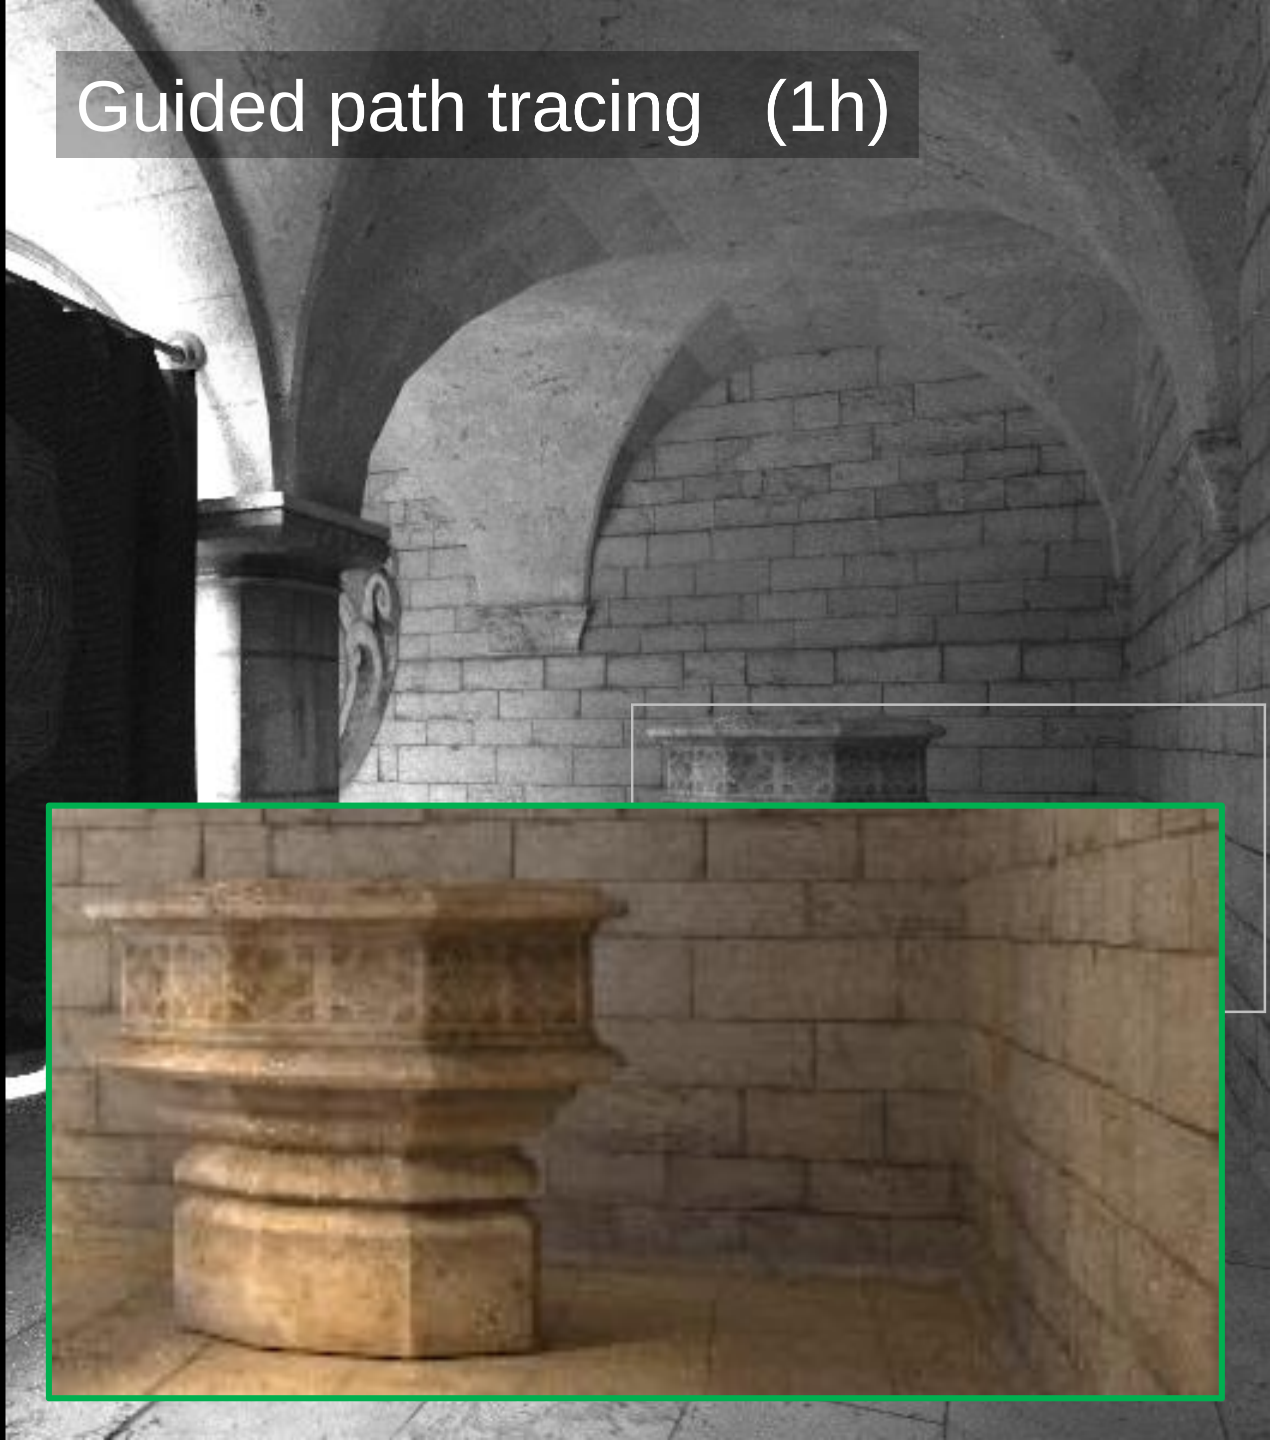
Guided path tracing (1h)

Reference

Path tracing

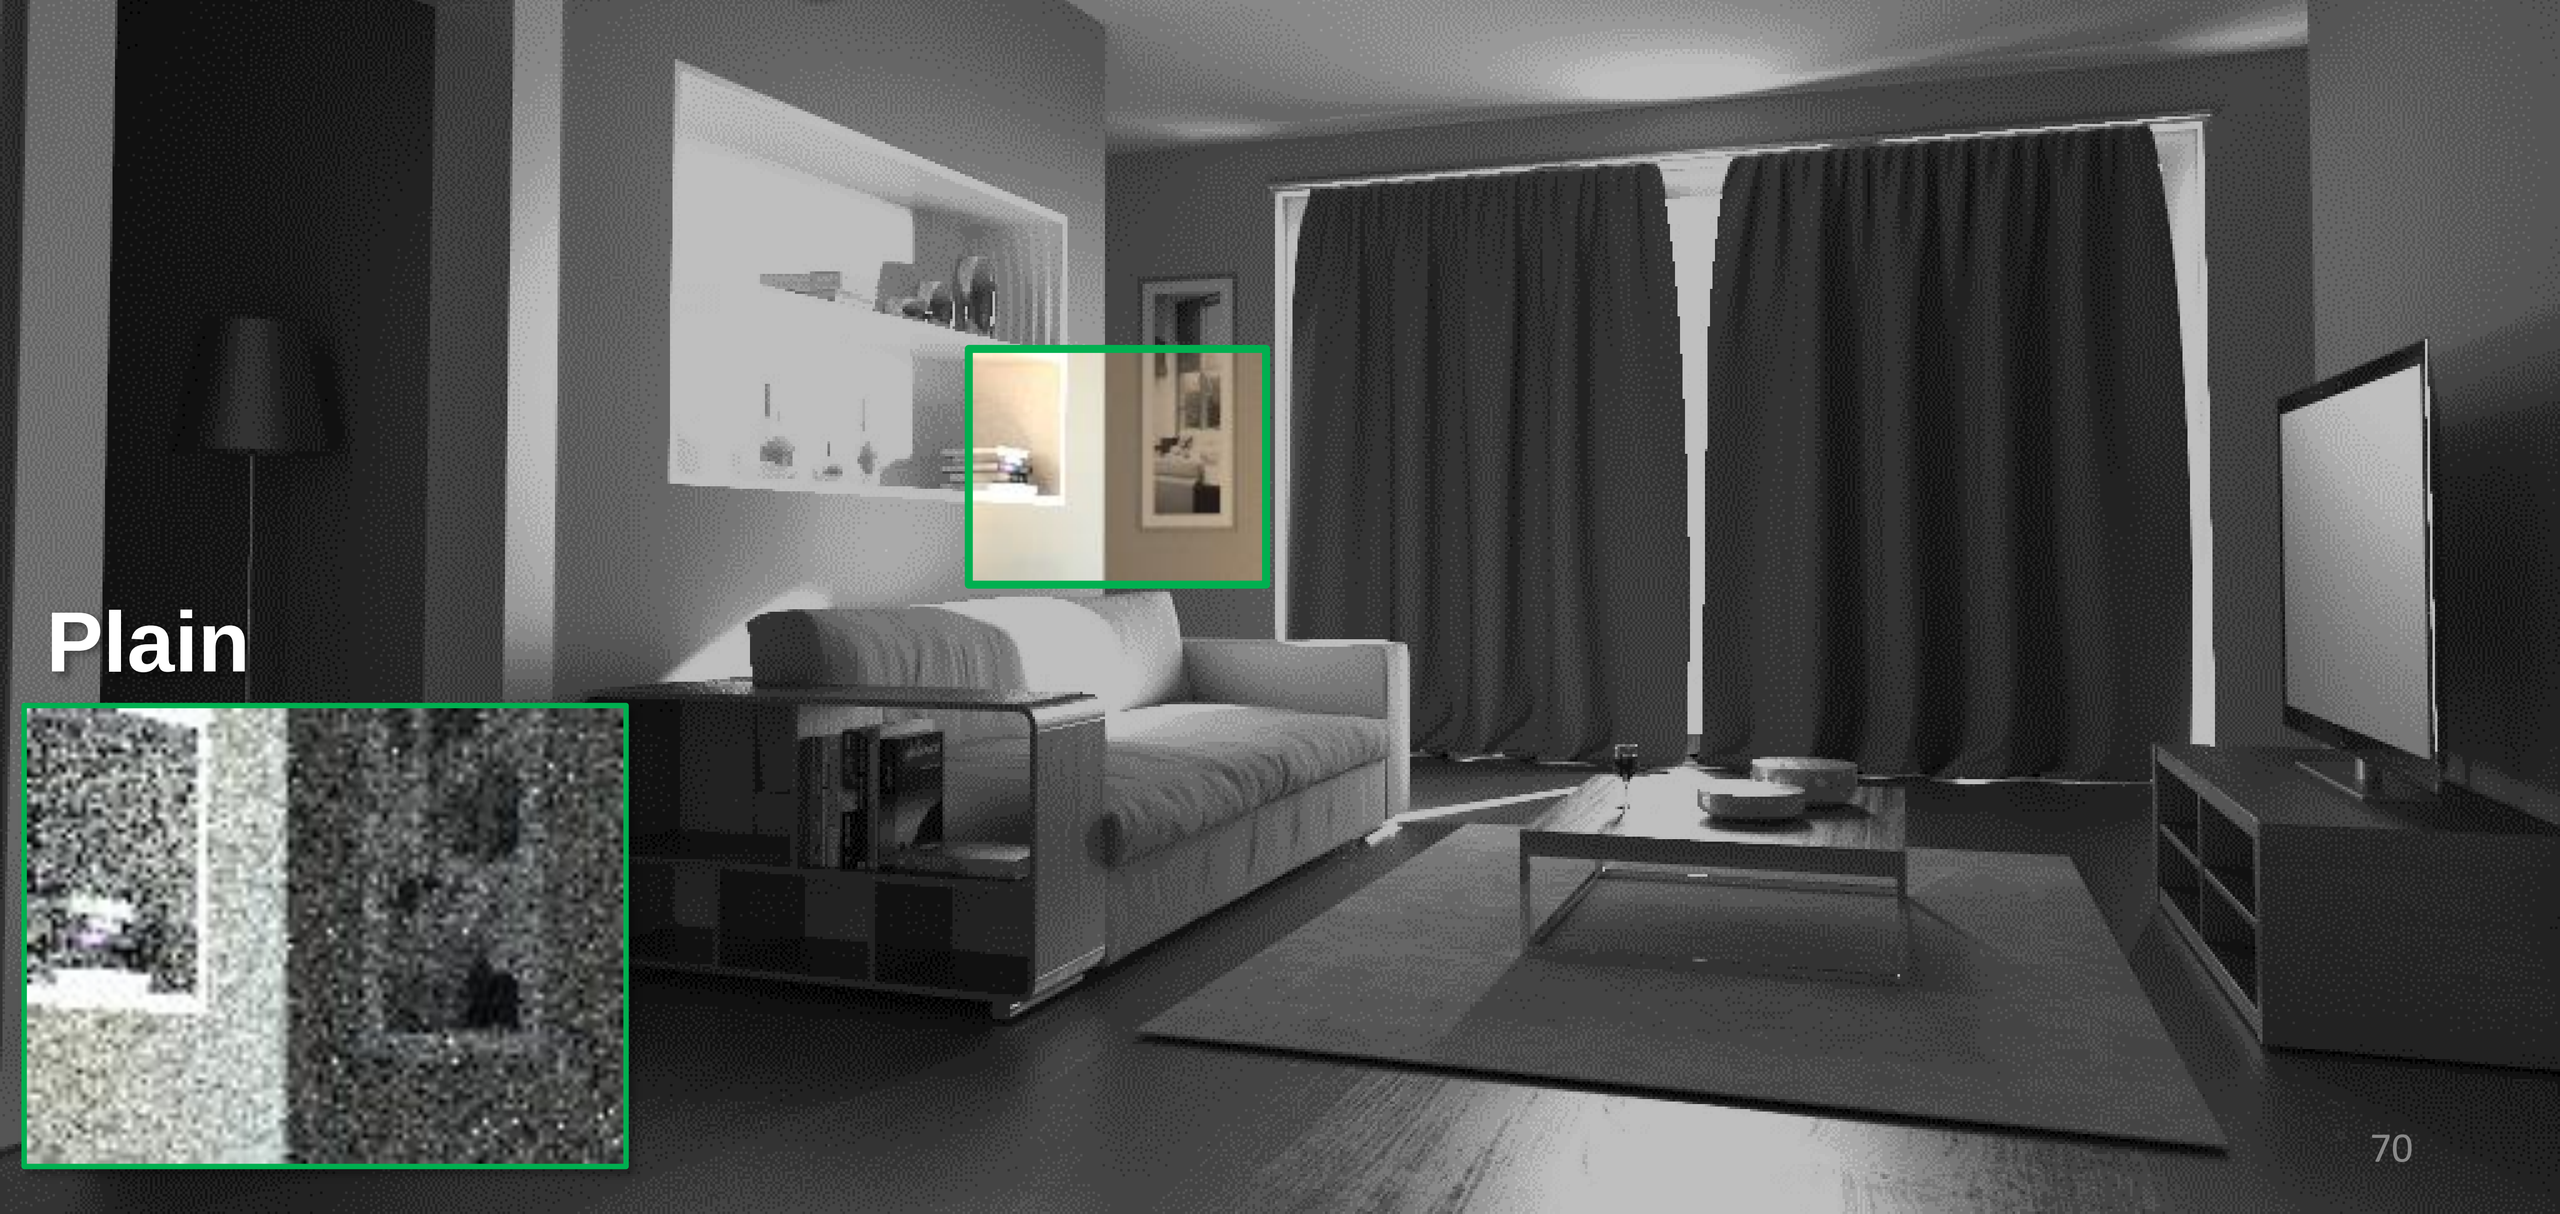

Path tracing

Path tracing

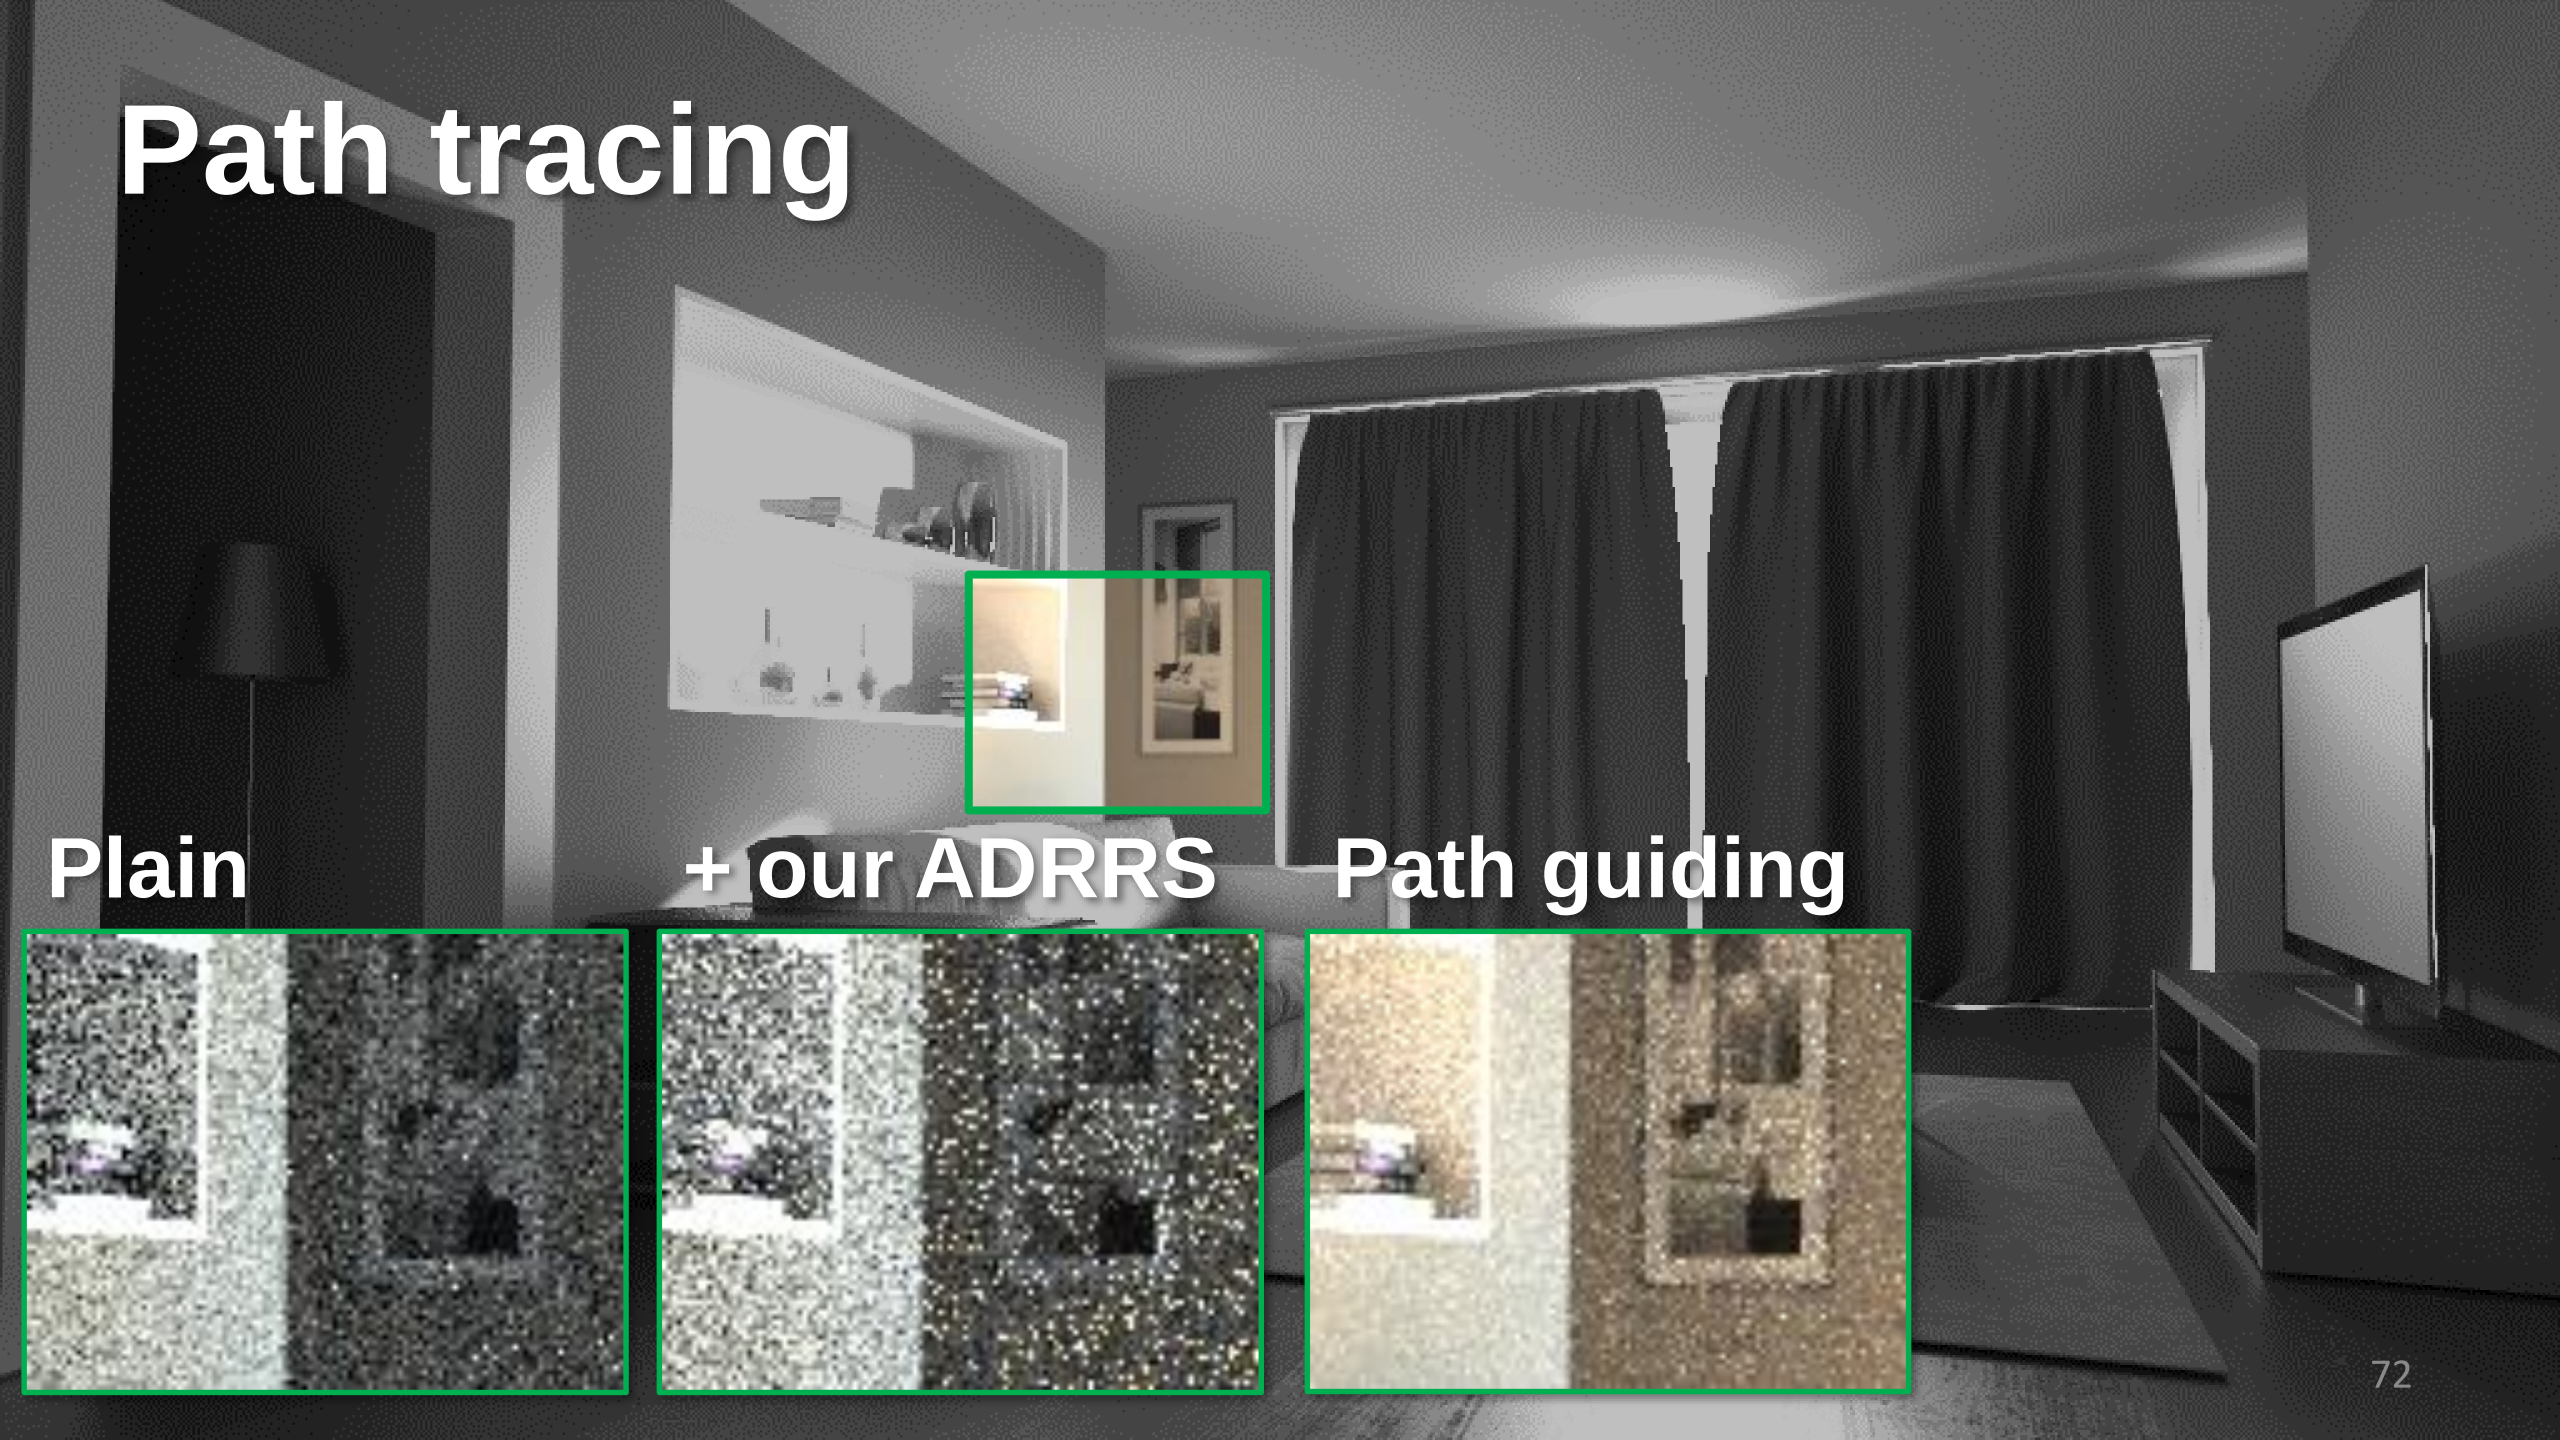

Path tracing

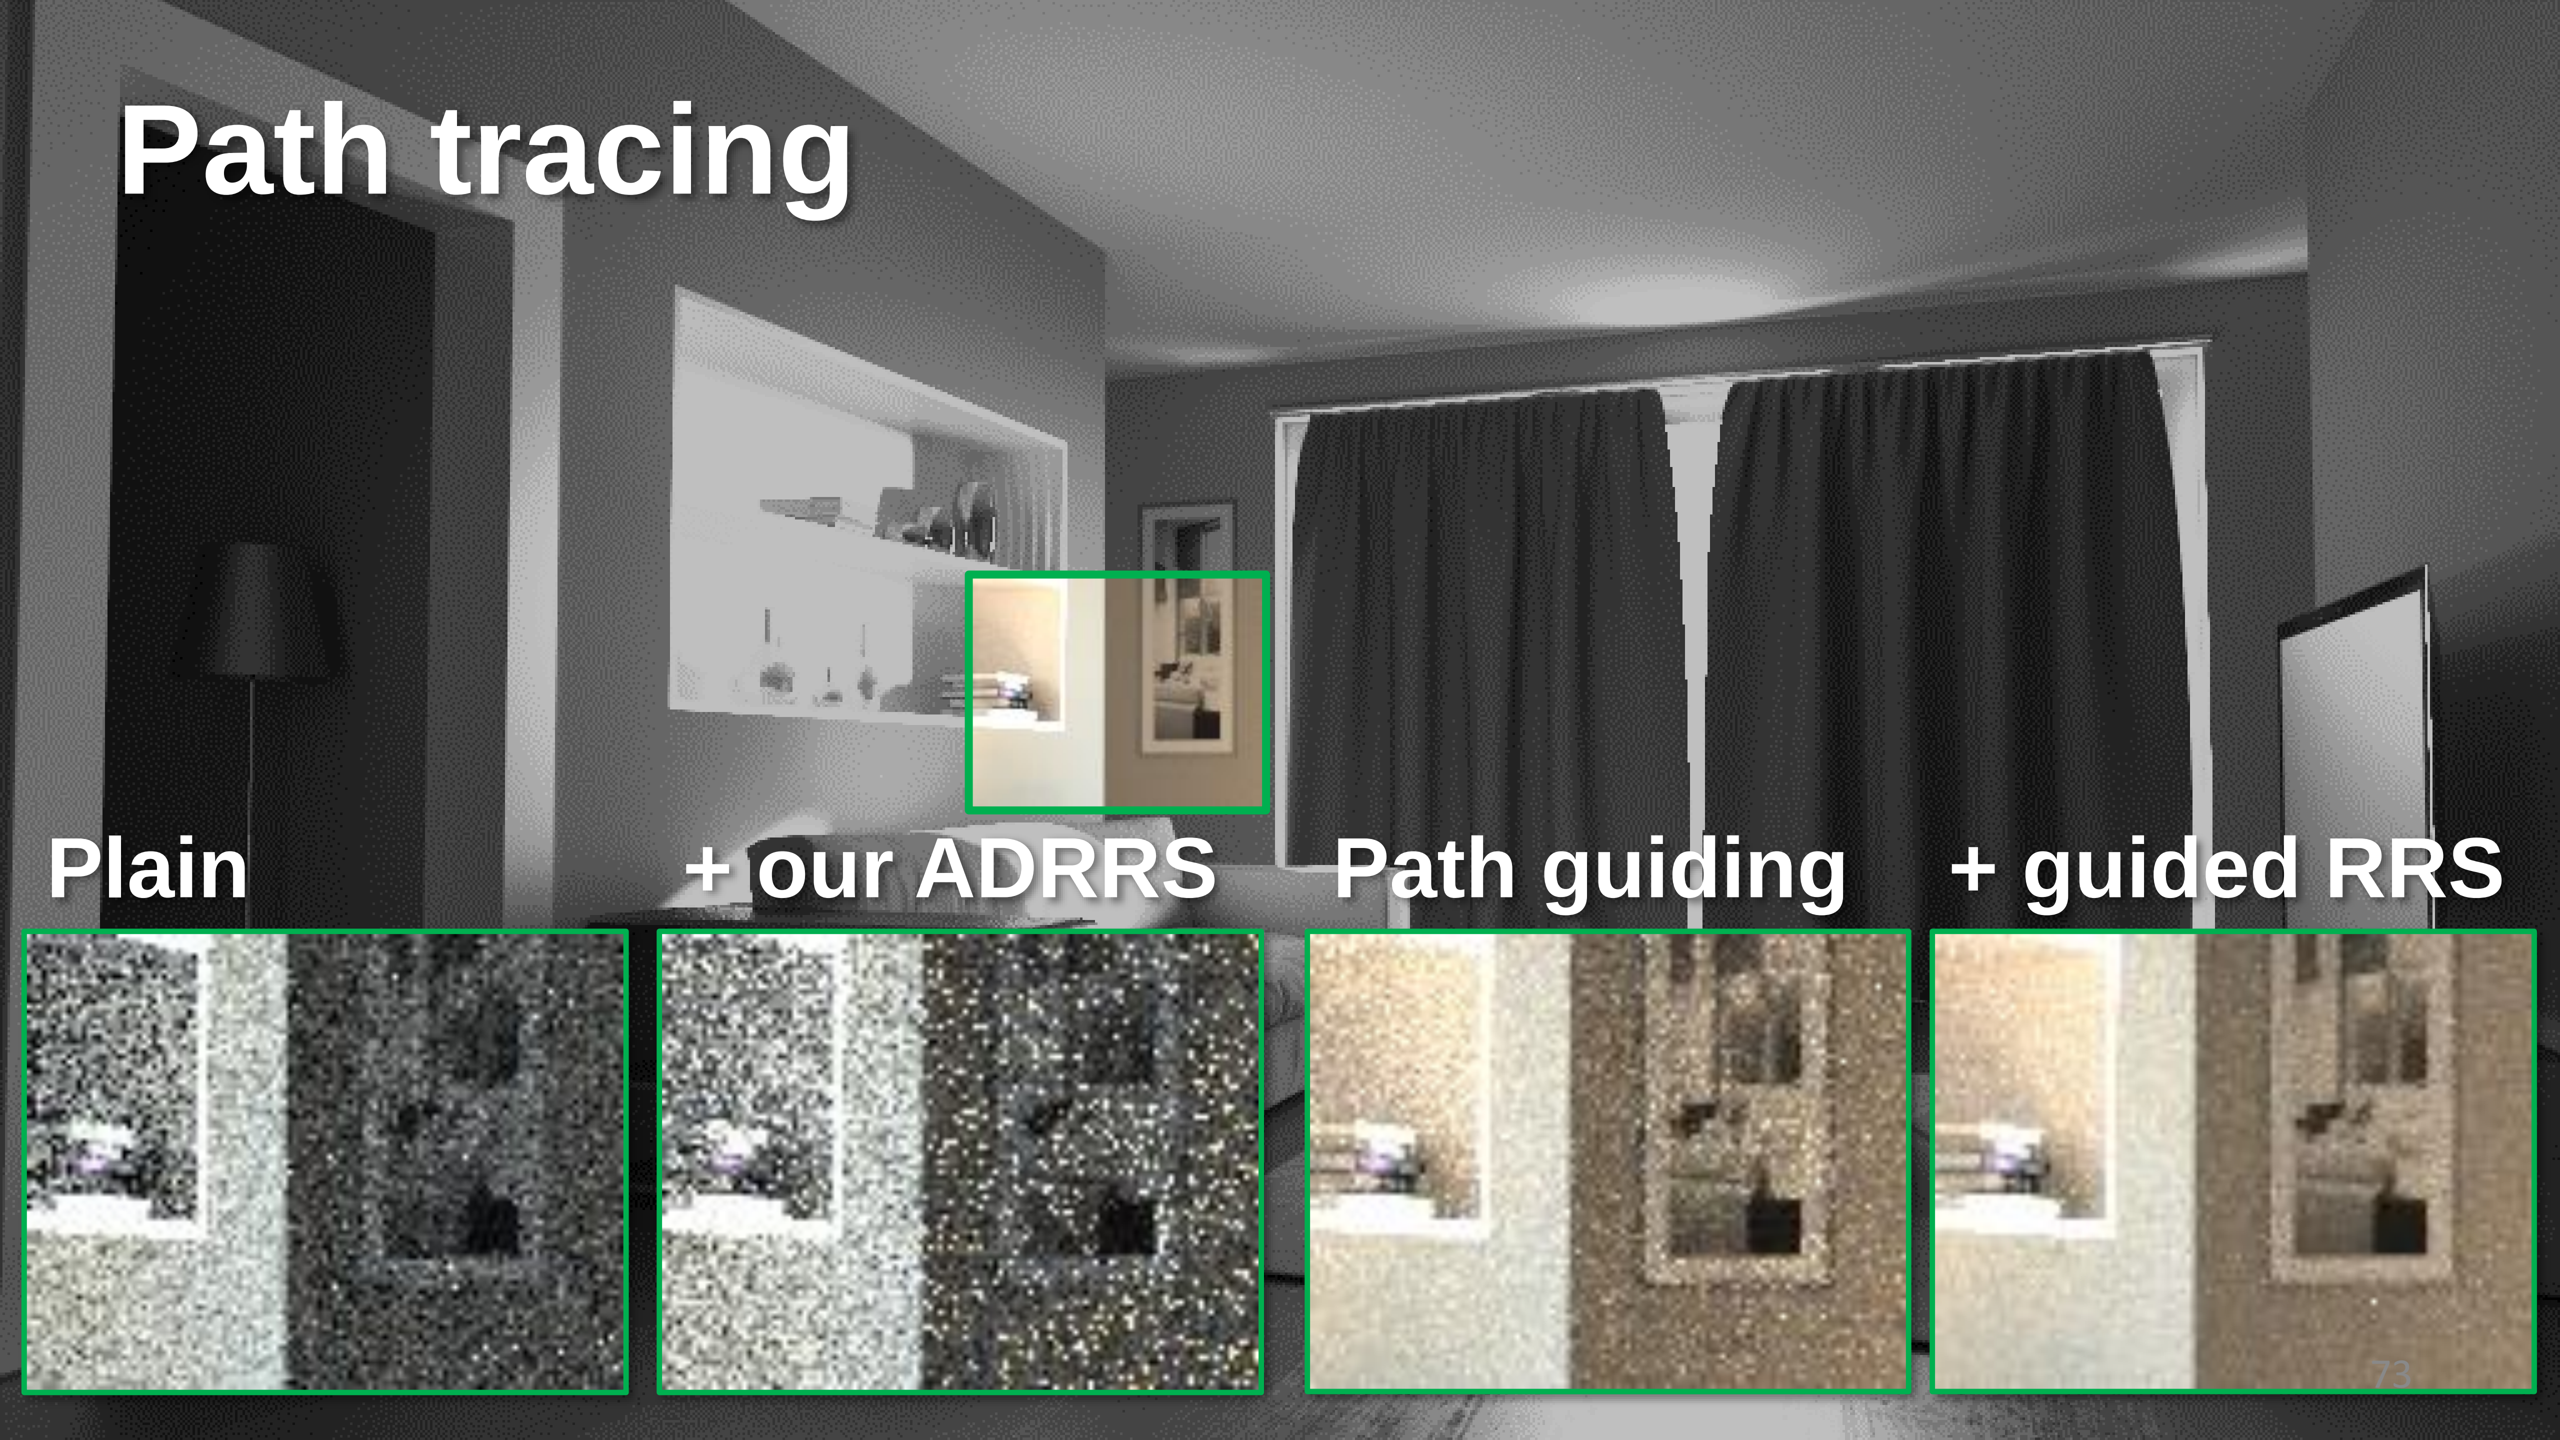

Complex Bidirectional Methods (VCM)

Guided path tracing can match complex methods

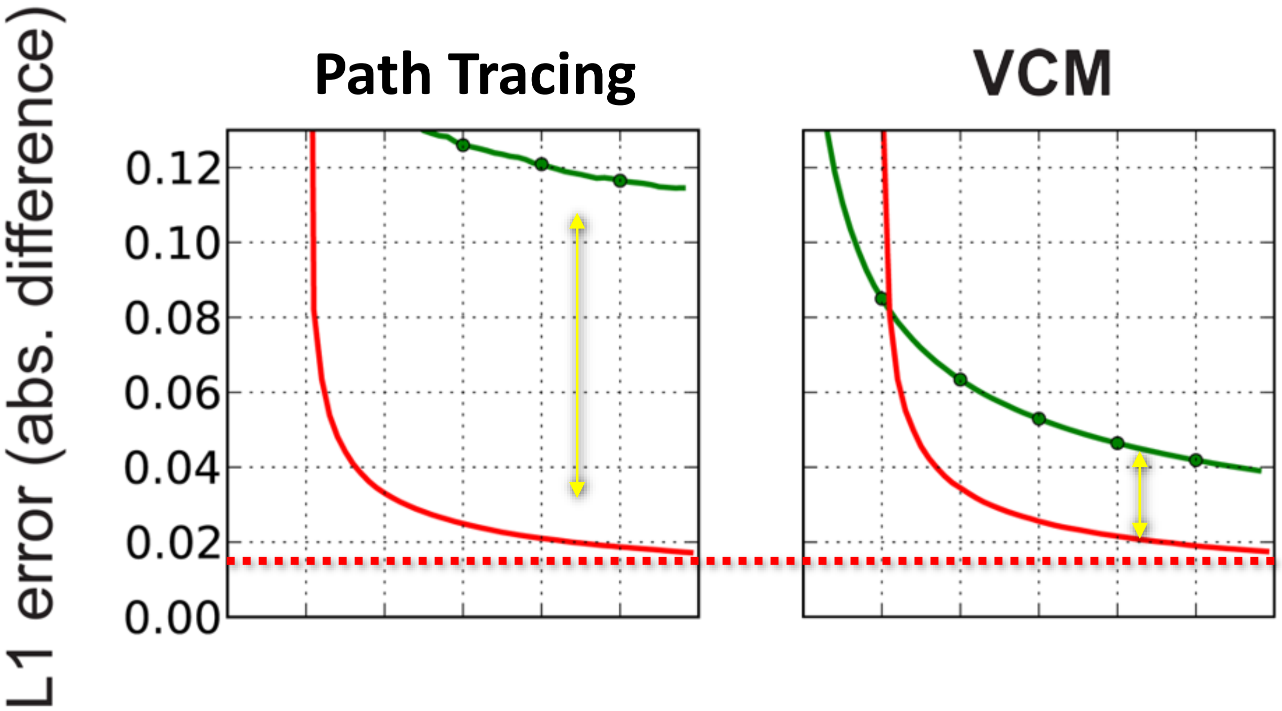

Practical Implication

- Providing path tracer with information makes it much more robust
    - Machine learning is the key (online step-wise EM formulation)
- Step towards a simpler ultimate algorithm
- Path guiding applicable in production

GUIDED VOLUMETRIC TRANSPORT

EBERHARD KARLS
UNIVERSITÄT
TÜBINGEN

Computer
Graphics
Charles
University

DiSTRO

Volume path guiding

- All events importance sampled
- Product sampling for collision distance

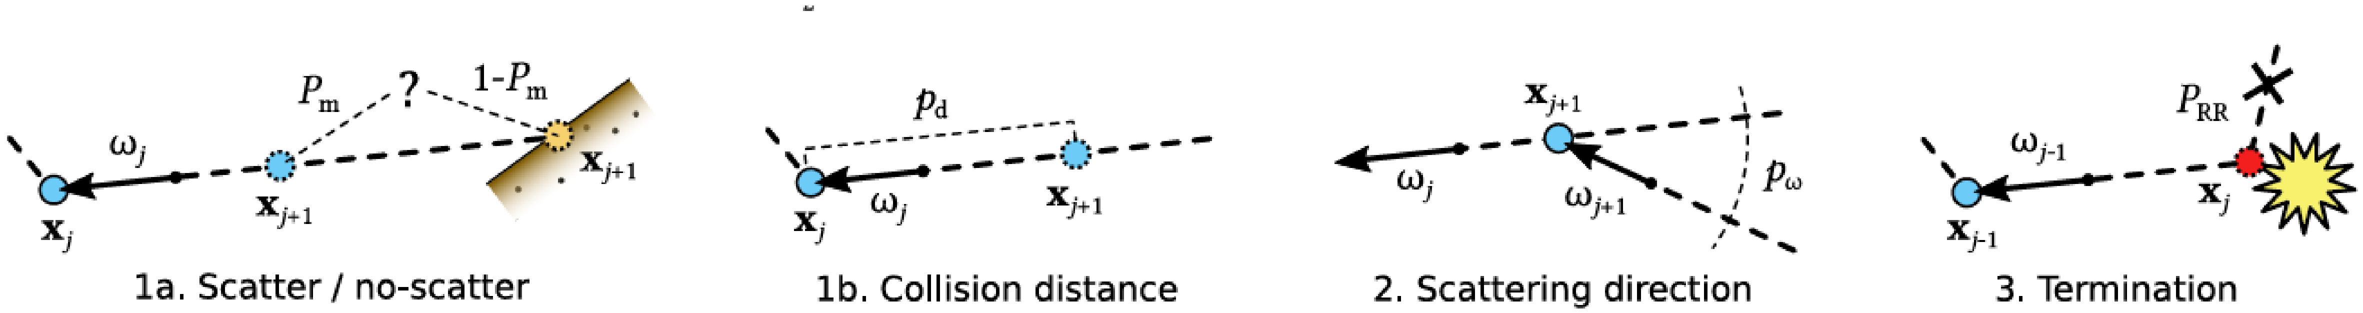

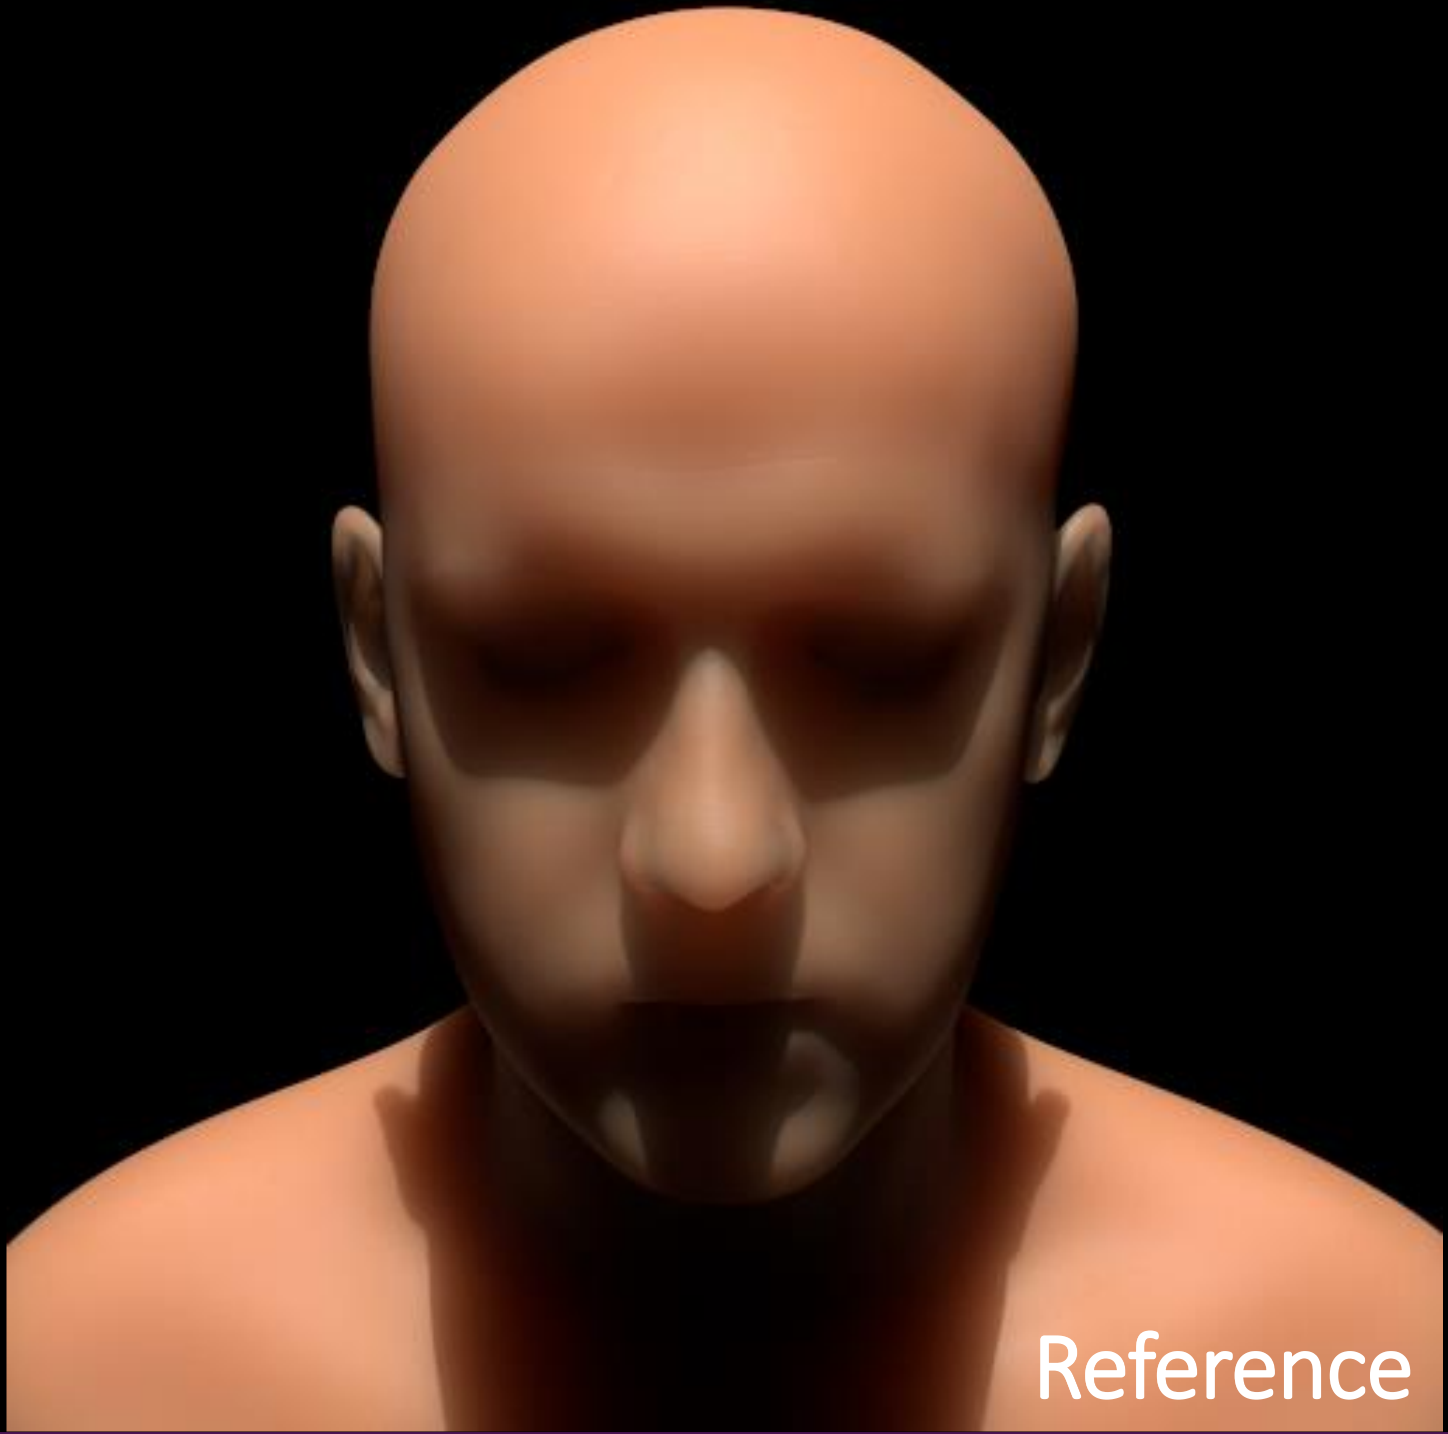
Reference

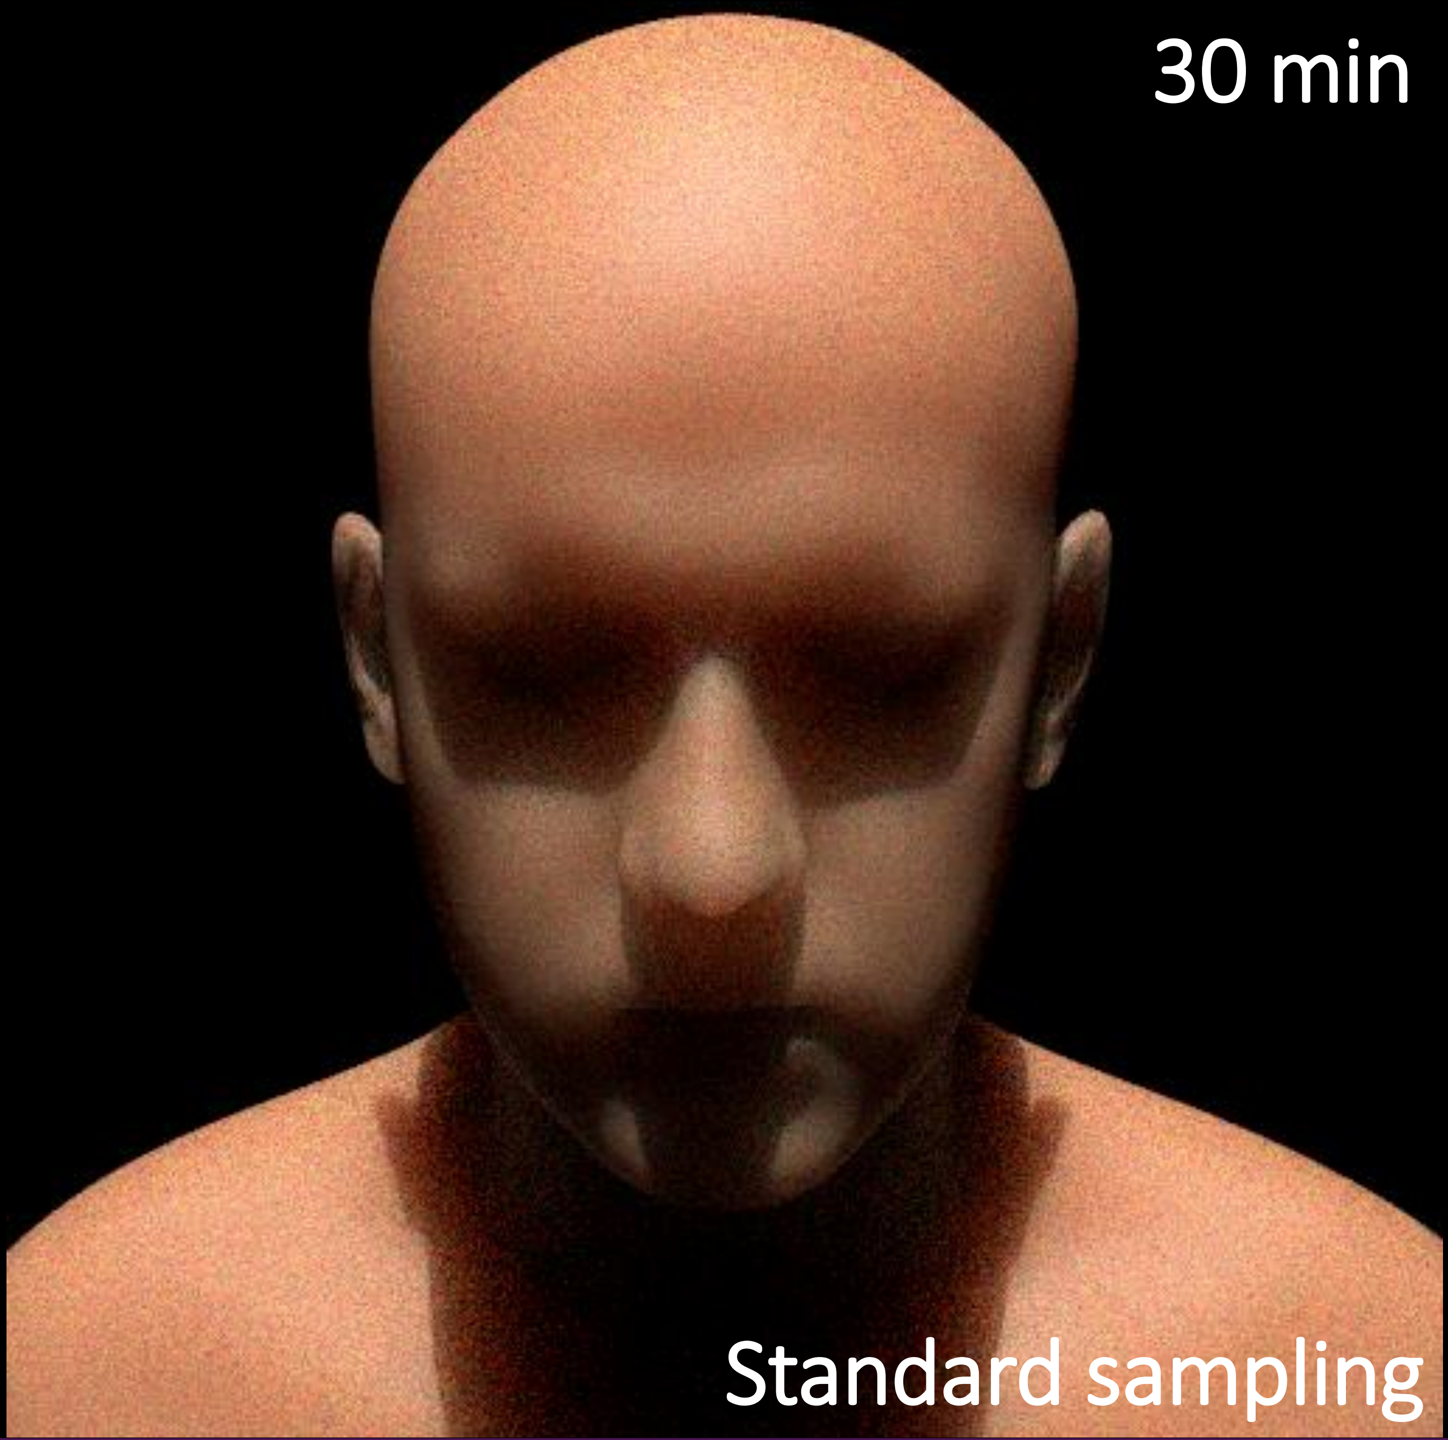
30 min
Standard sampling

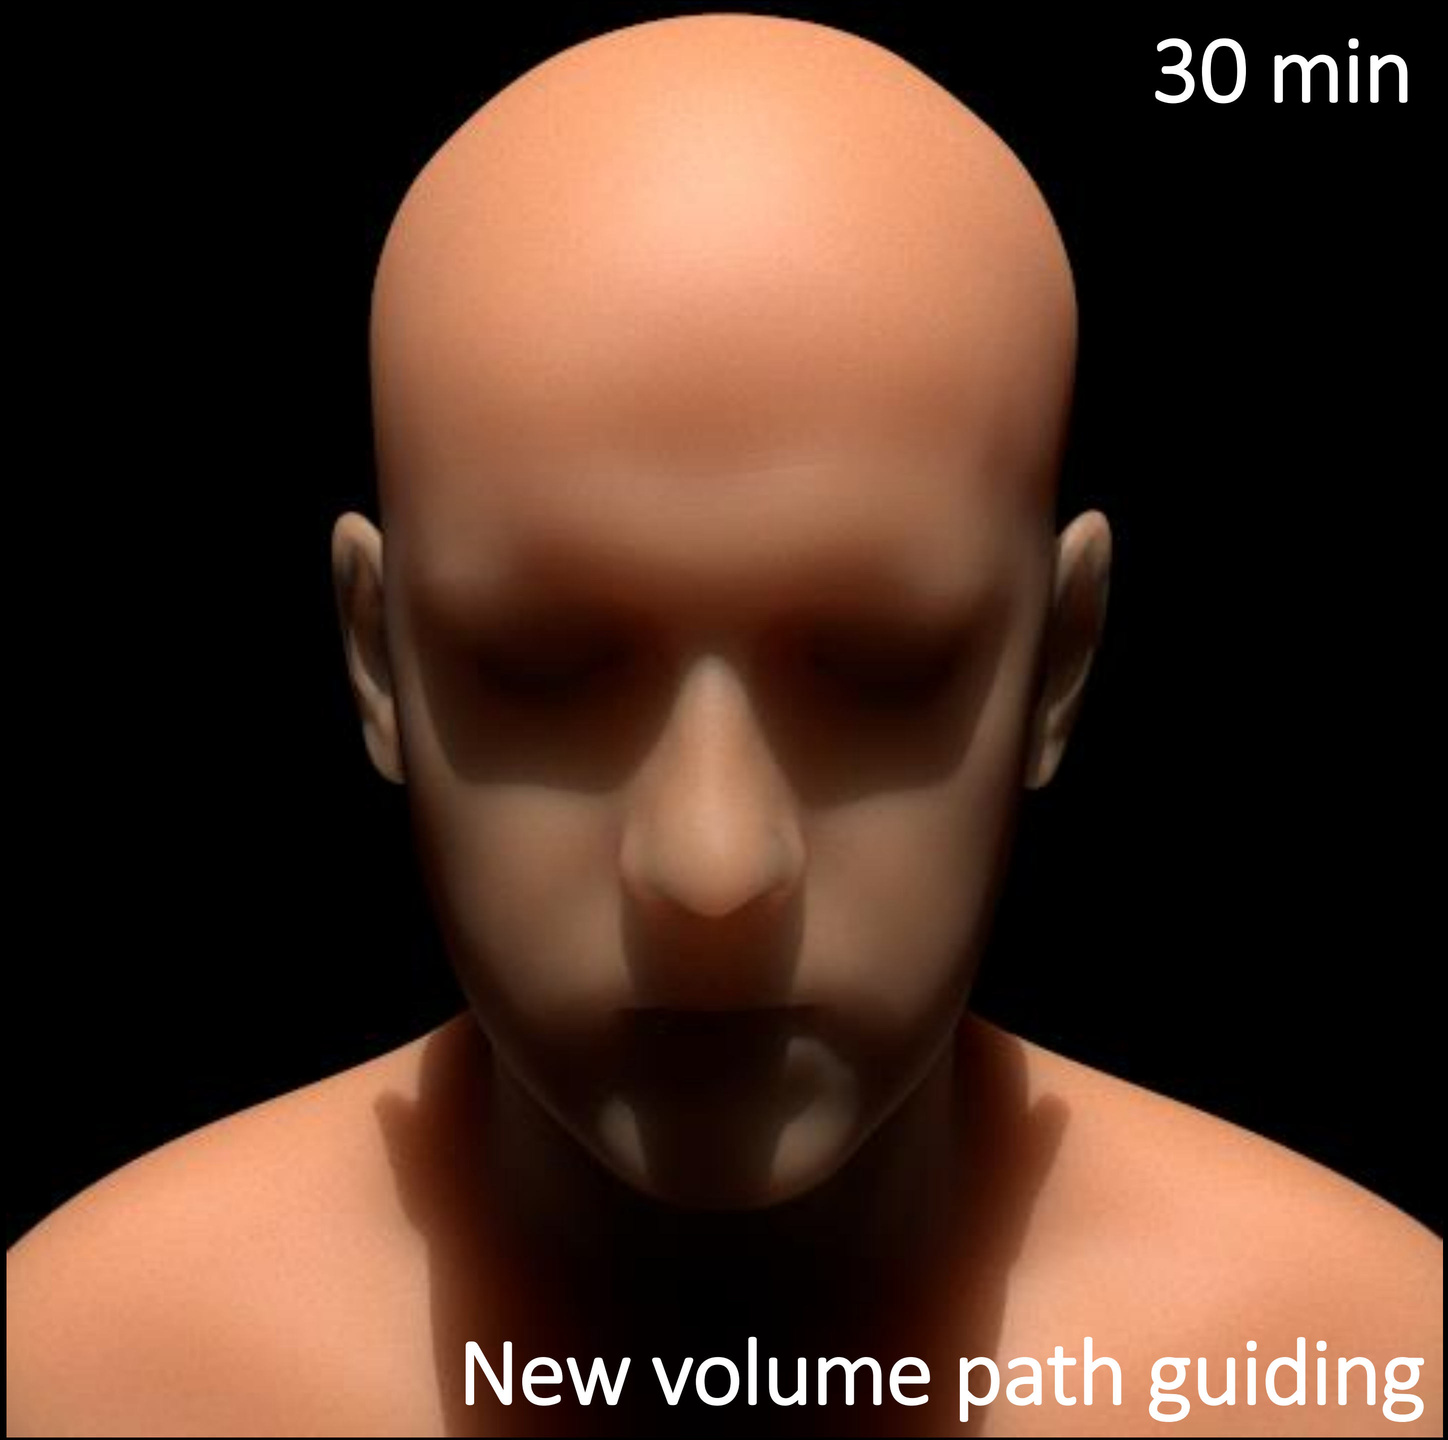
30 min
New volume path guiding

Standard sampling

SPP: 1580
relMSE: 6.458

Dist. + dir. guiding

SPP: 1288
relMSE: 1.354

RR + splitting

SPP: 1660
relMSE: 0.401

Reference


45 min
Standard sampling

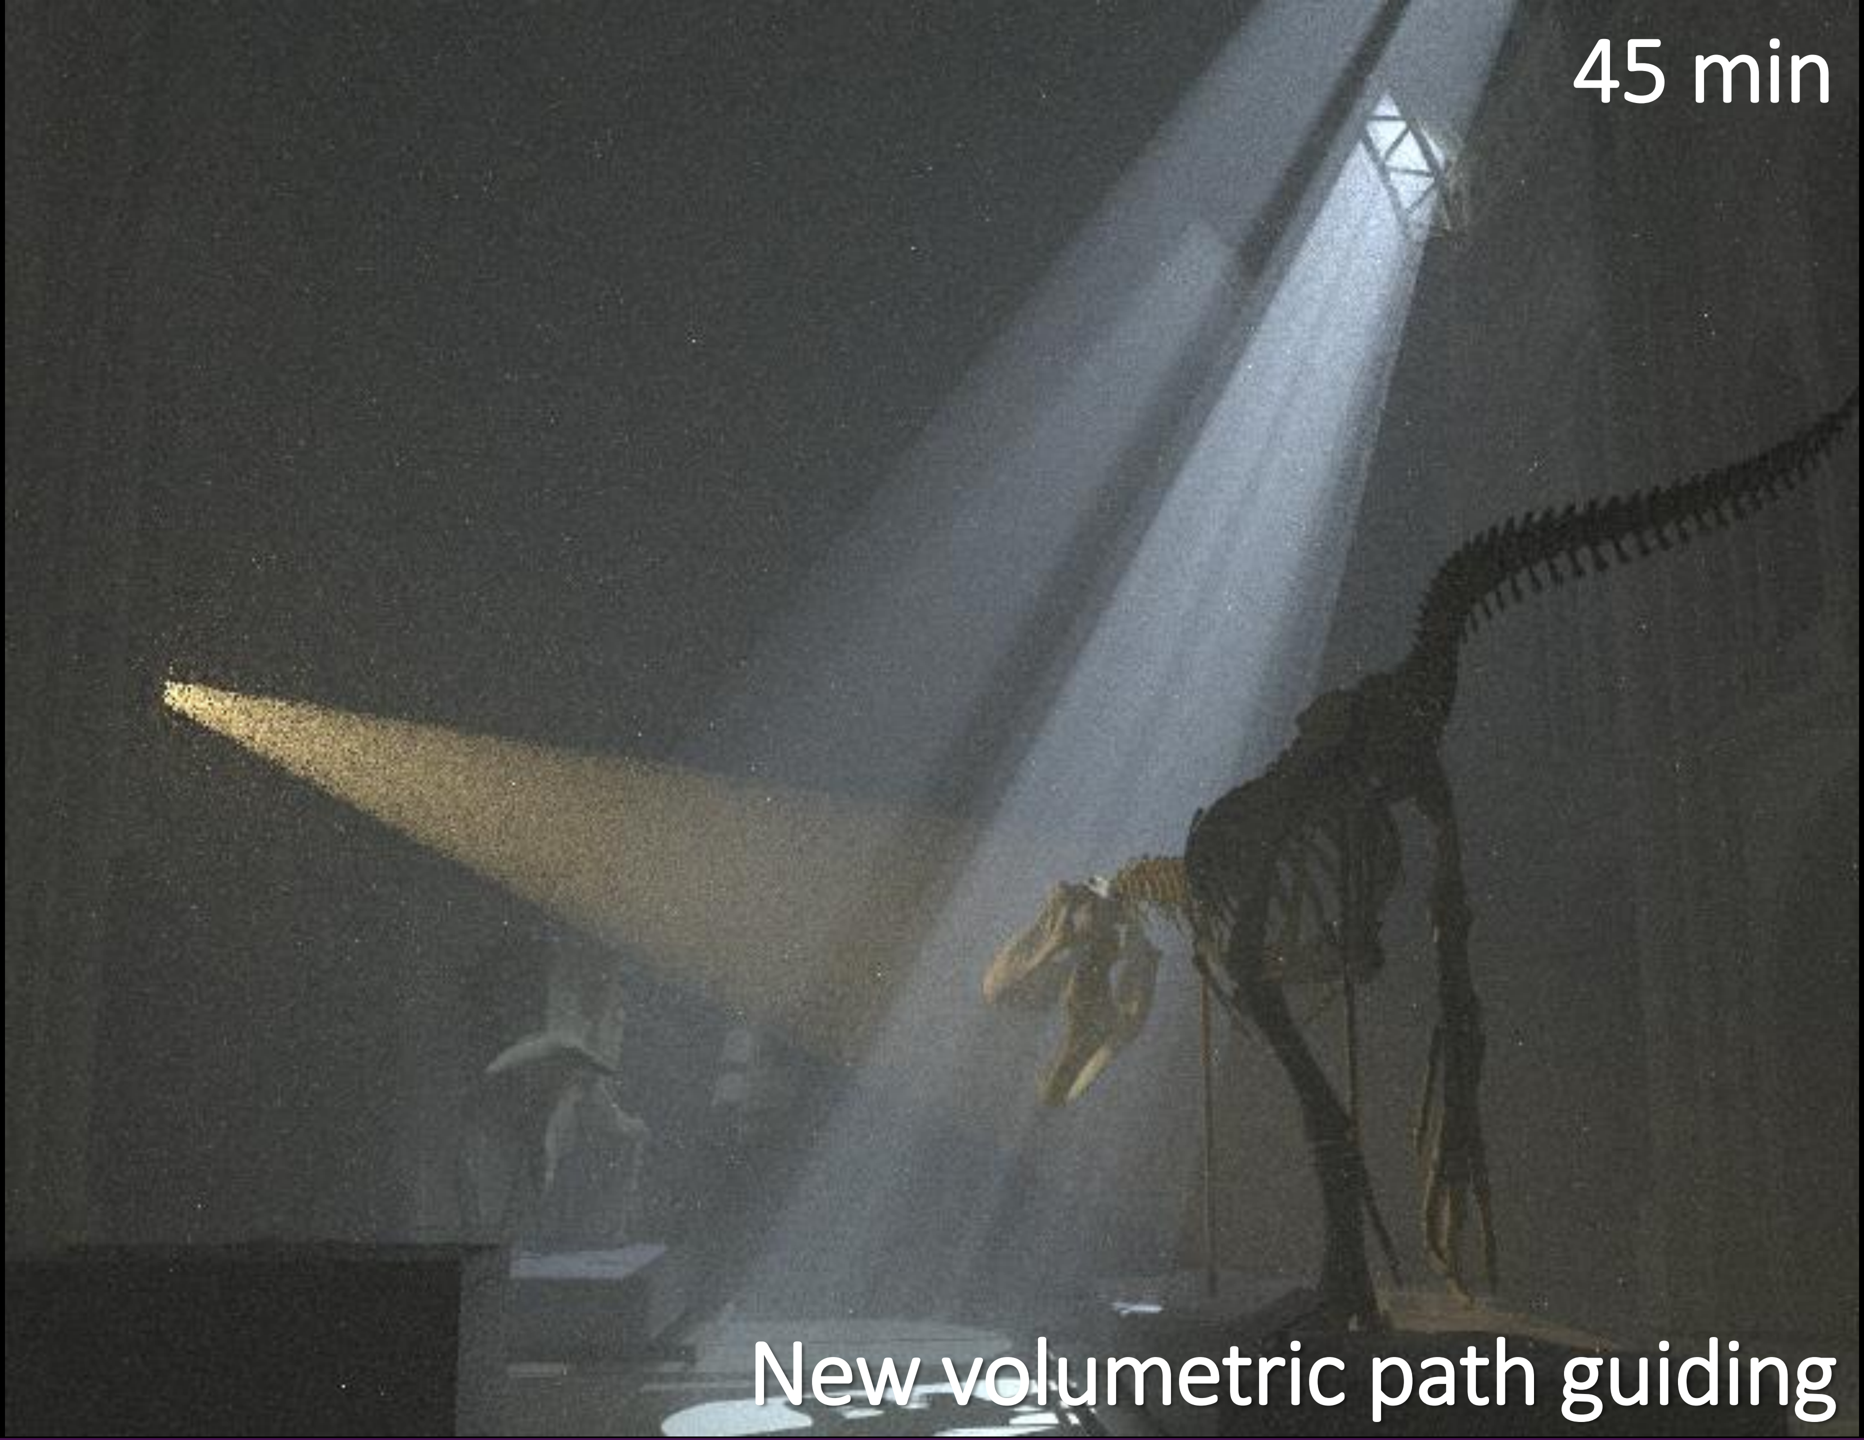
45 min
New volumetric path guiding

Standard sampling

SPP: 796
relMSE: 1.725

Dist. + dir. guiding

SPP: 392
relMSE: 0.747

RR + splitting

SPP: 1068
relMSE: 0.123

Bayesian online regression for adaptive direct illumination sampling

Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek

Render Legion, a.s.
Charles University, Prague

Computer Graphics Charles University

Direct + indirect illumination

Direct + indirect illumination

Non-adaptive sampling
[Wang et al. 2009]

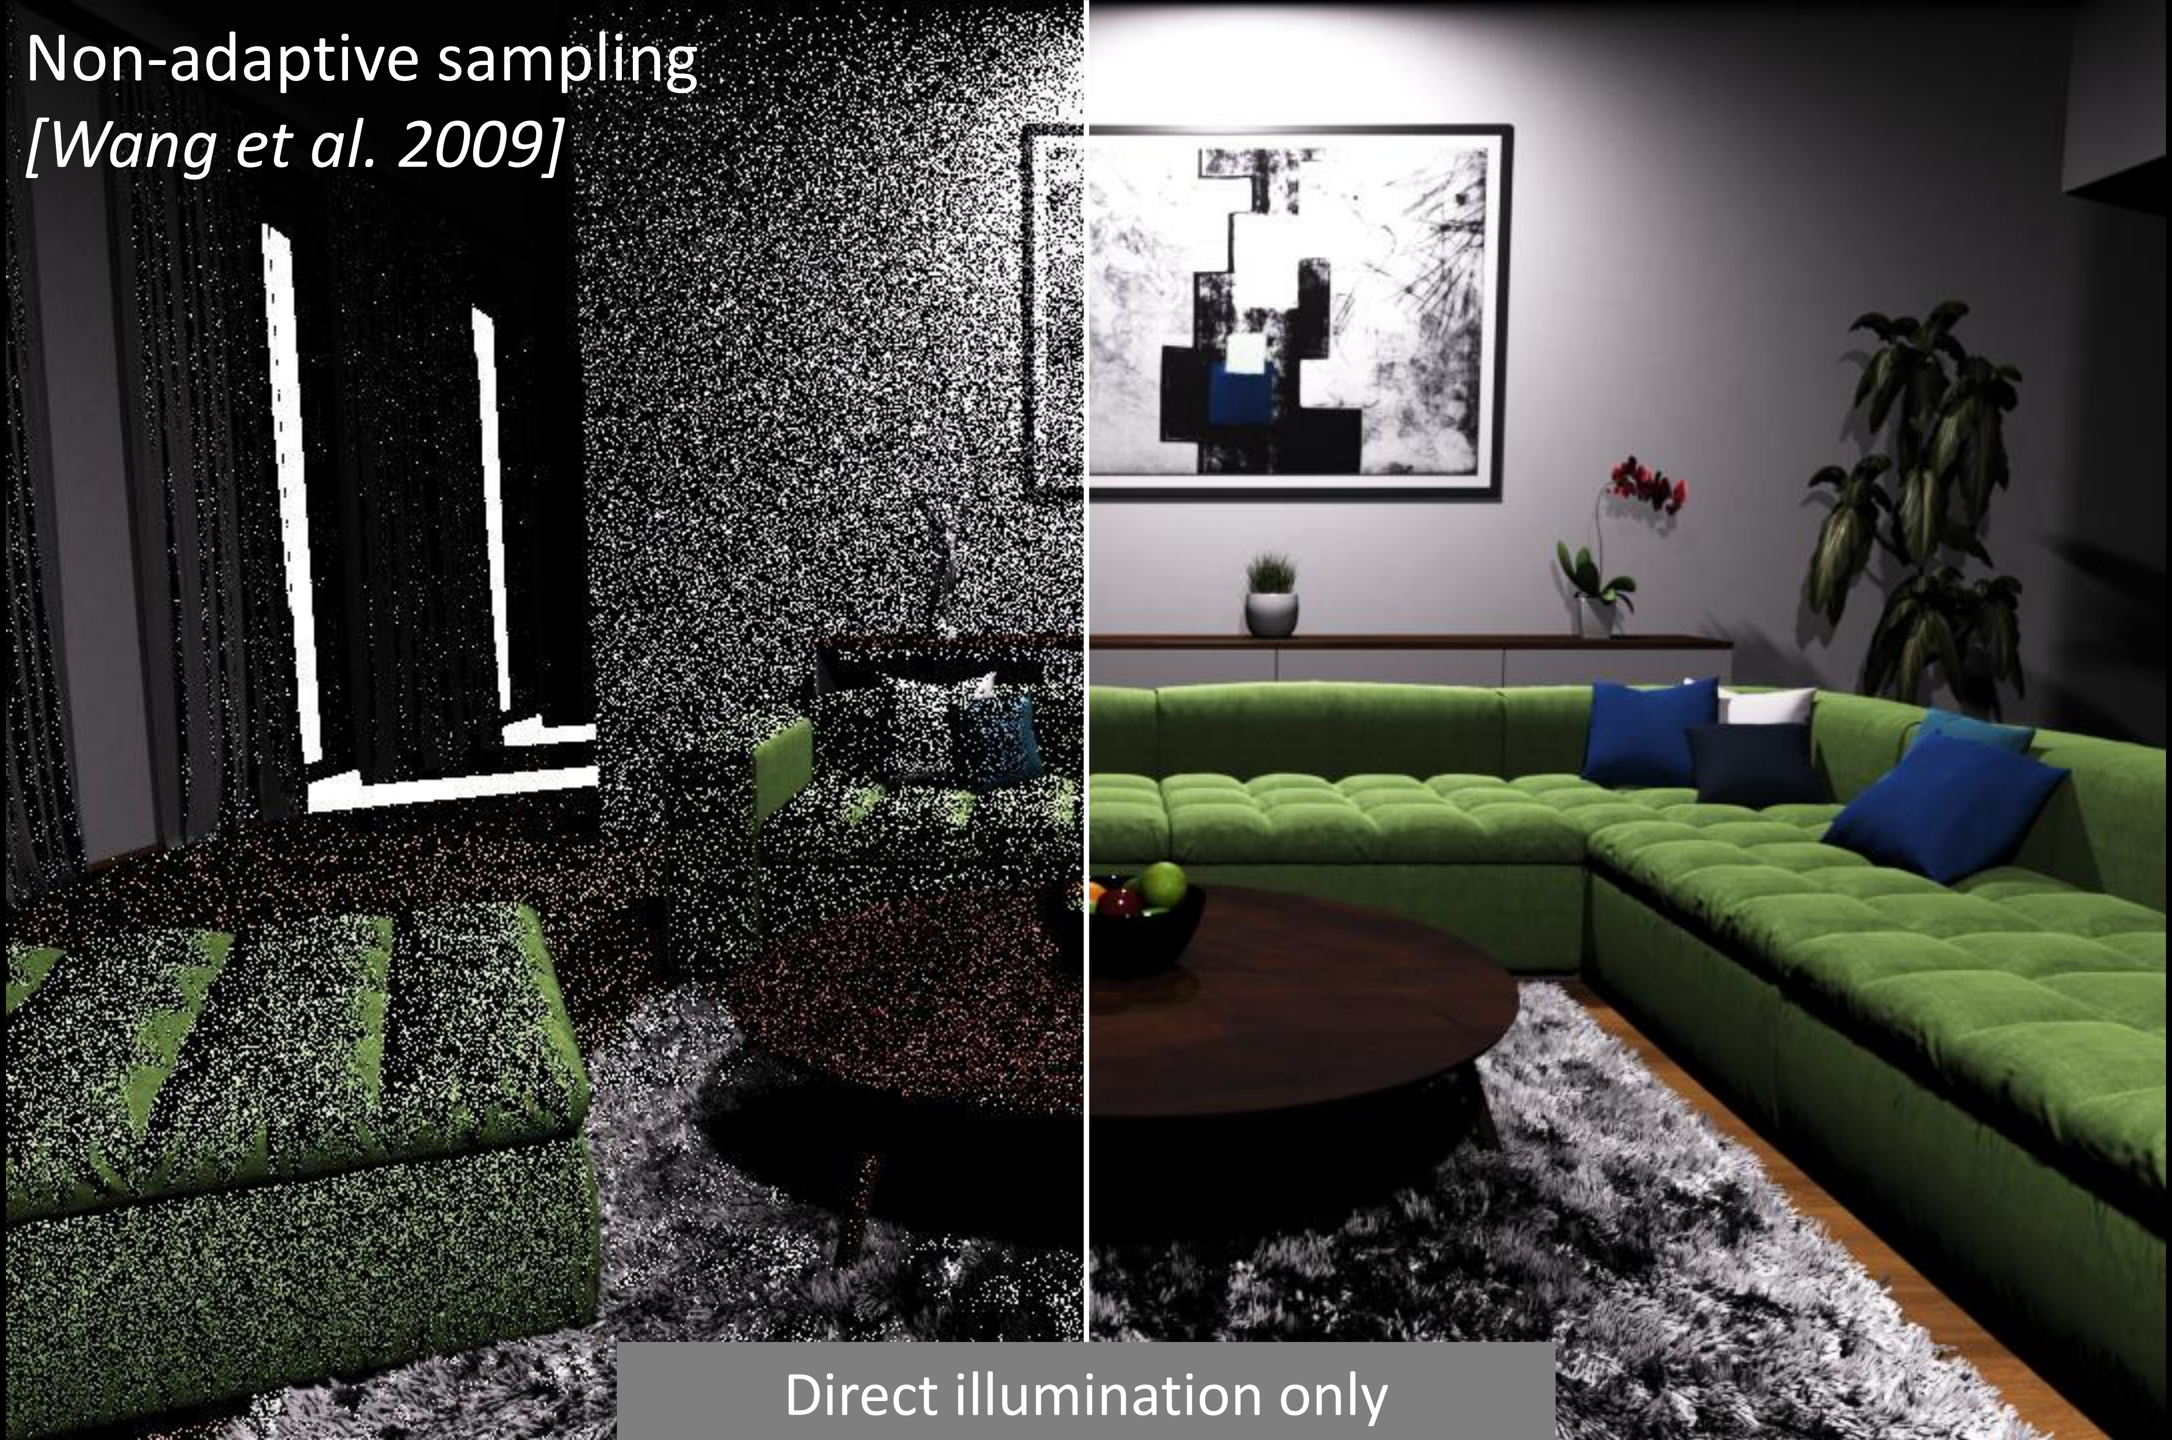

Direct illumination only

Non-adaptive sampling
[Wang et al. 2009]

Adaptive sampling
[Donikian et al. 2006]

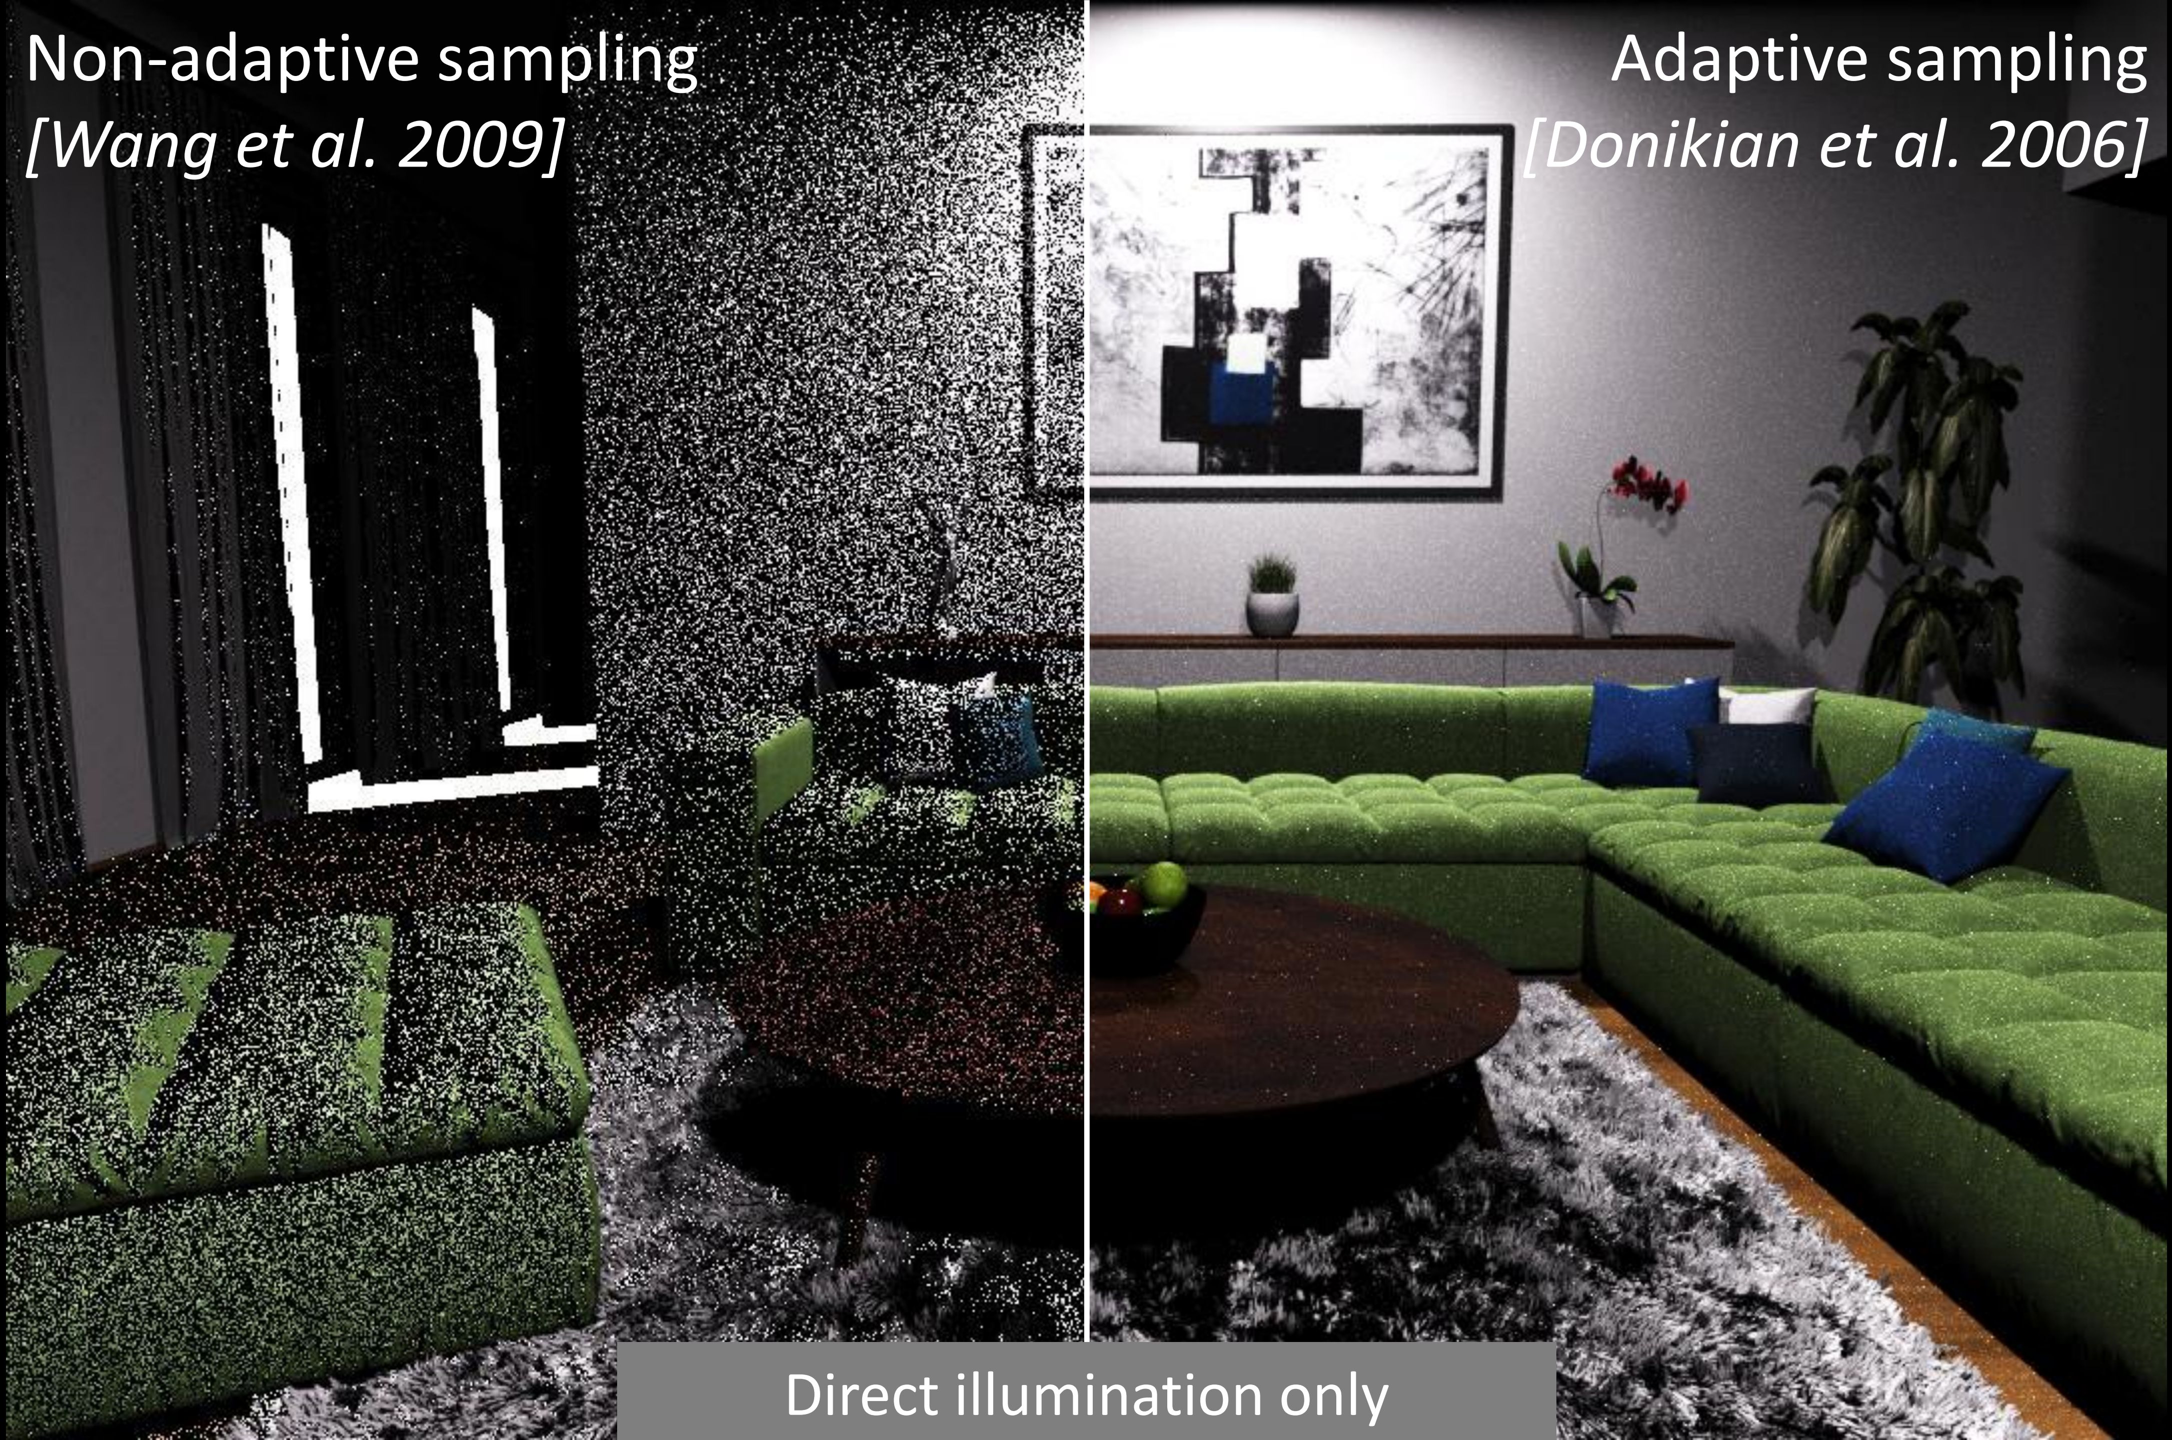

Direct illumination only

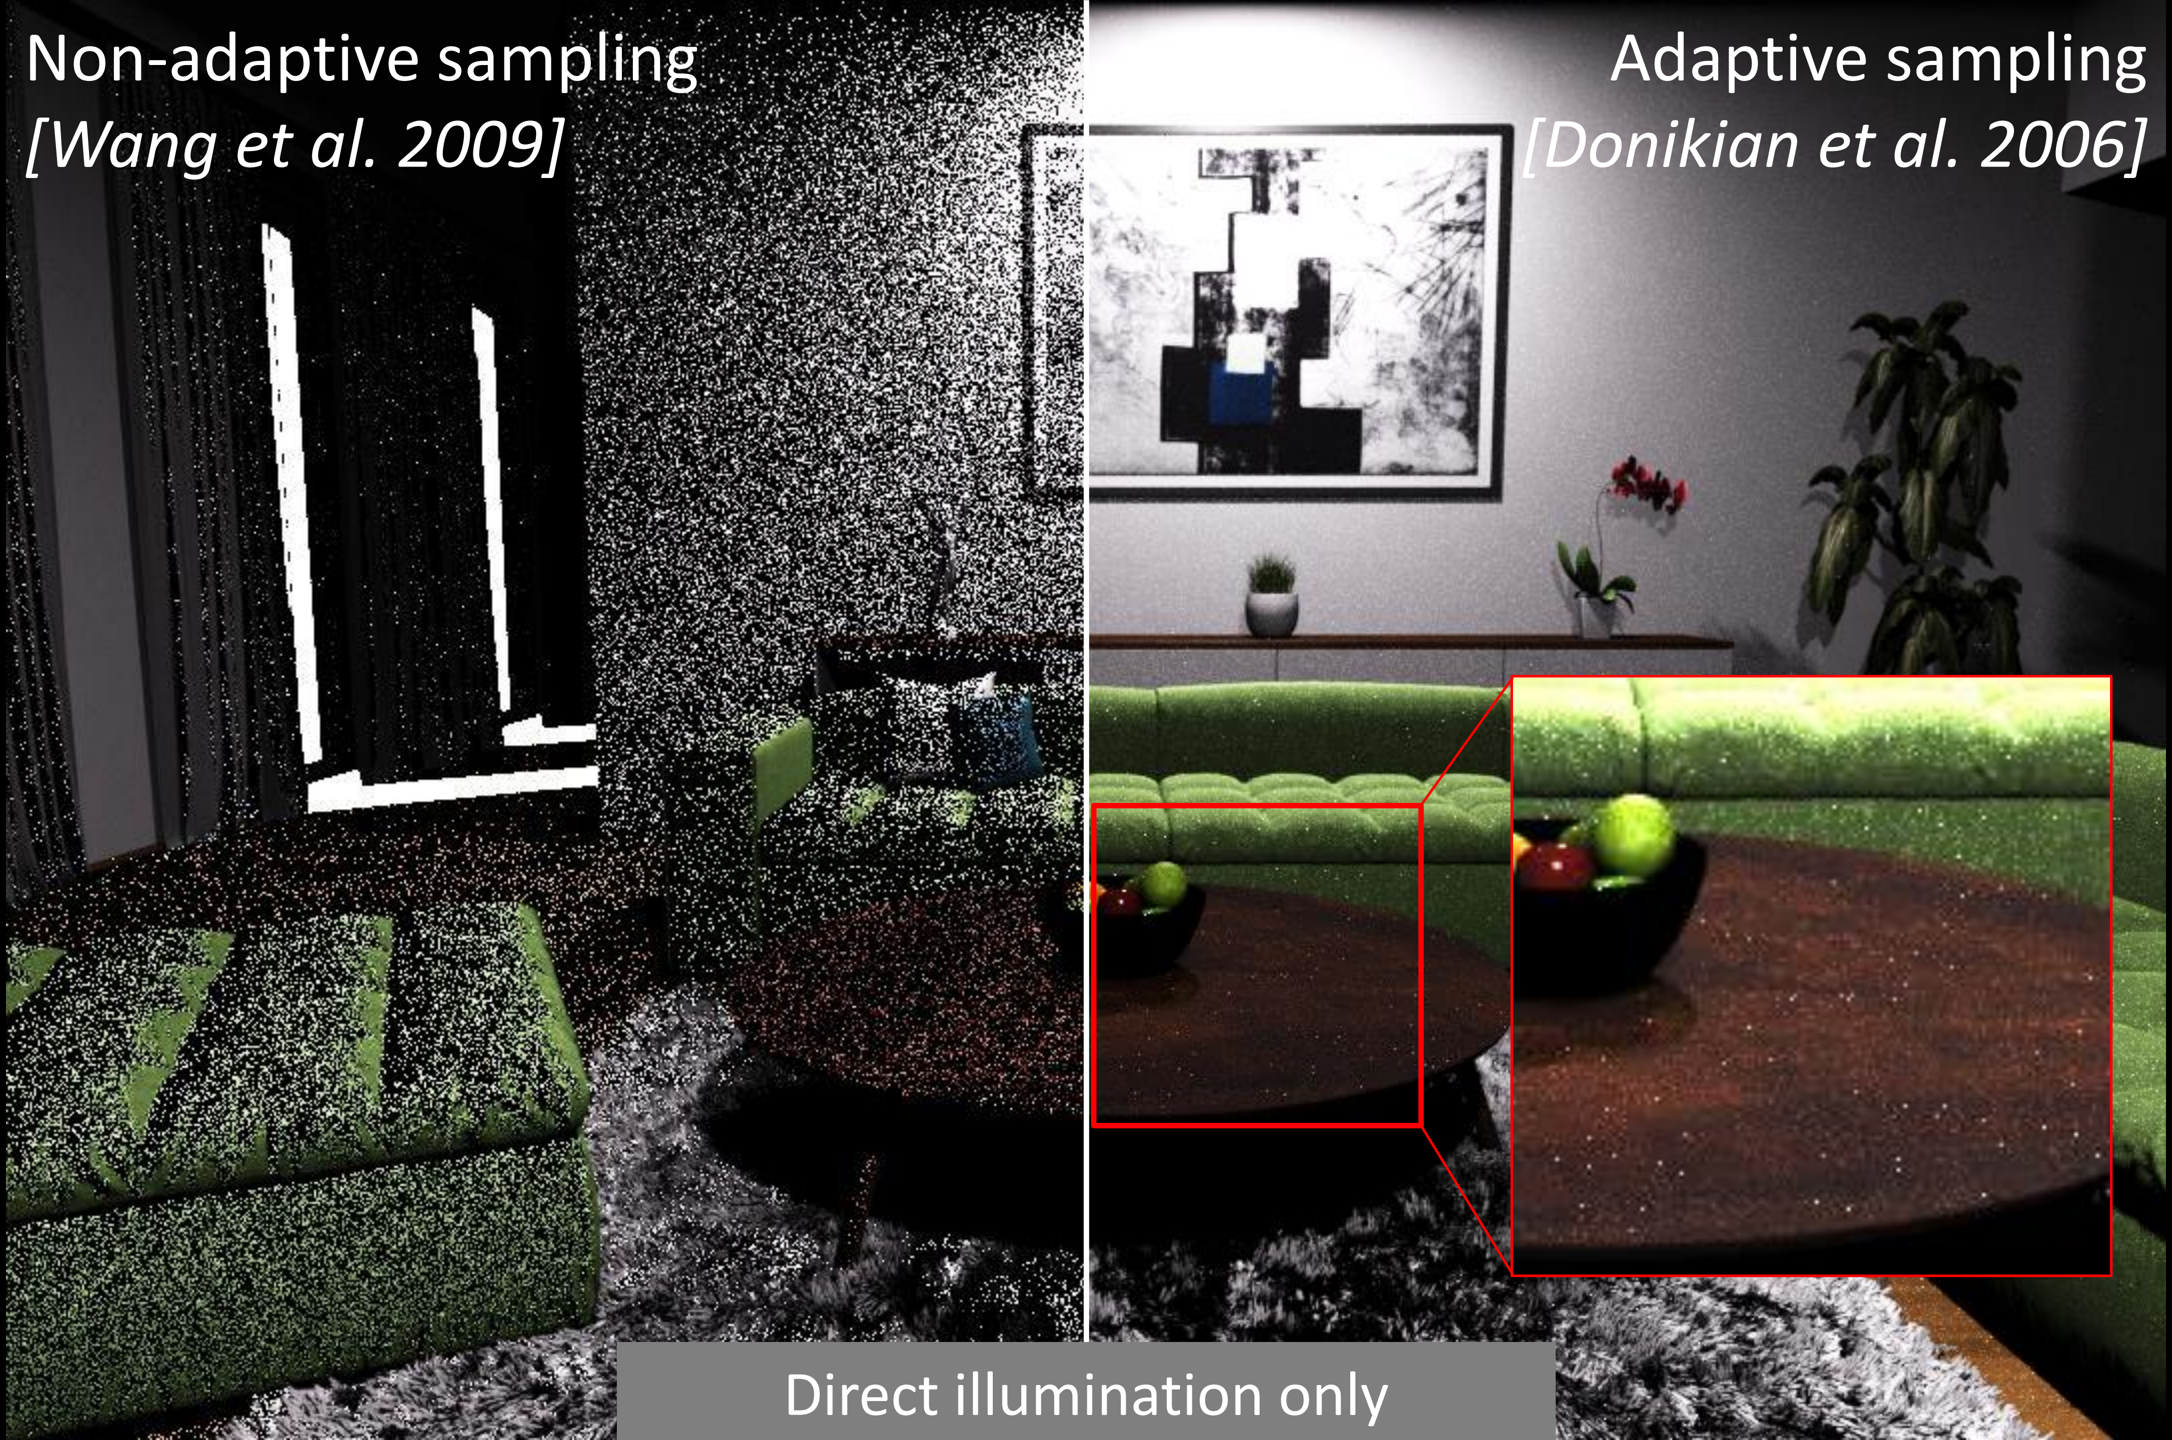
Non-adaptive sampling
[Wang et al. 2009]
Adaptive sampling
[Donikian et al. 2006]
Direct illumination only

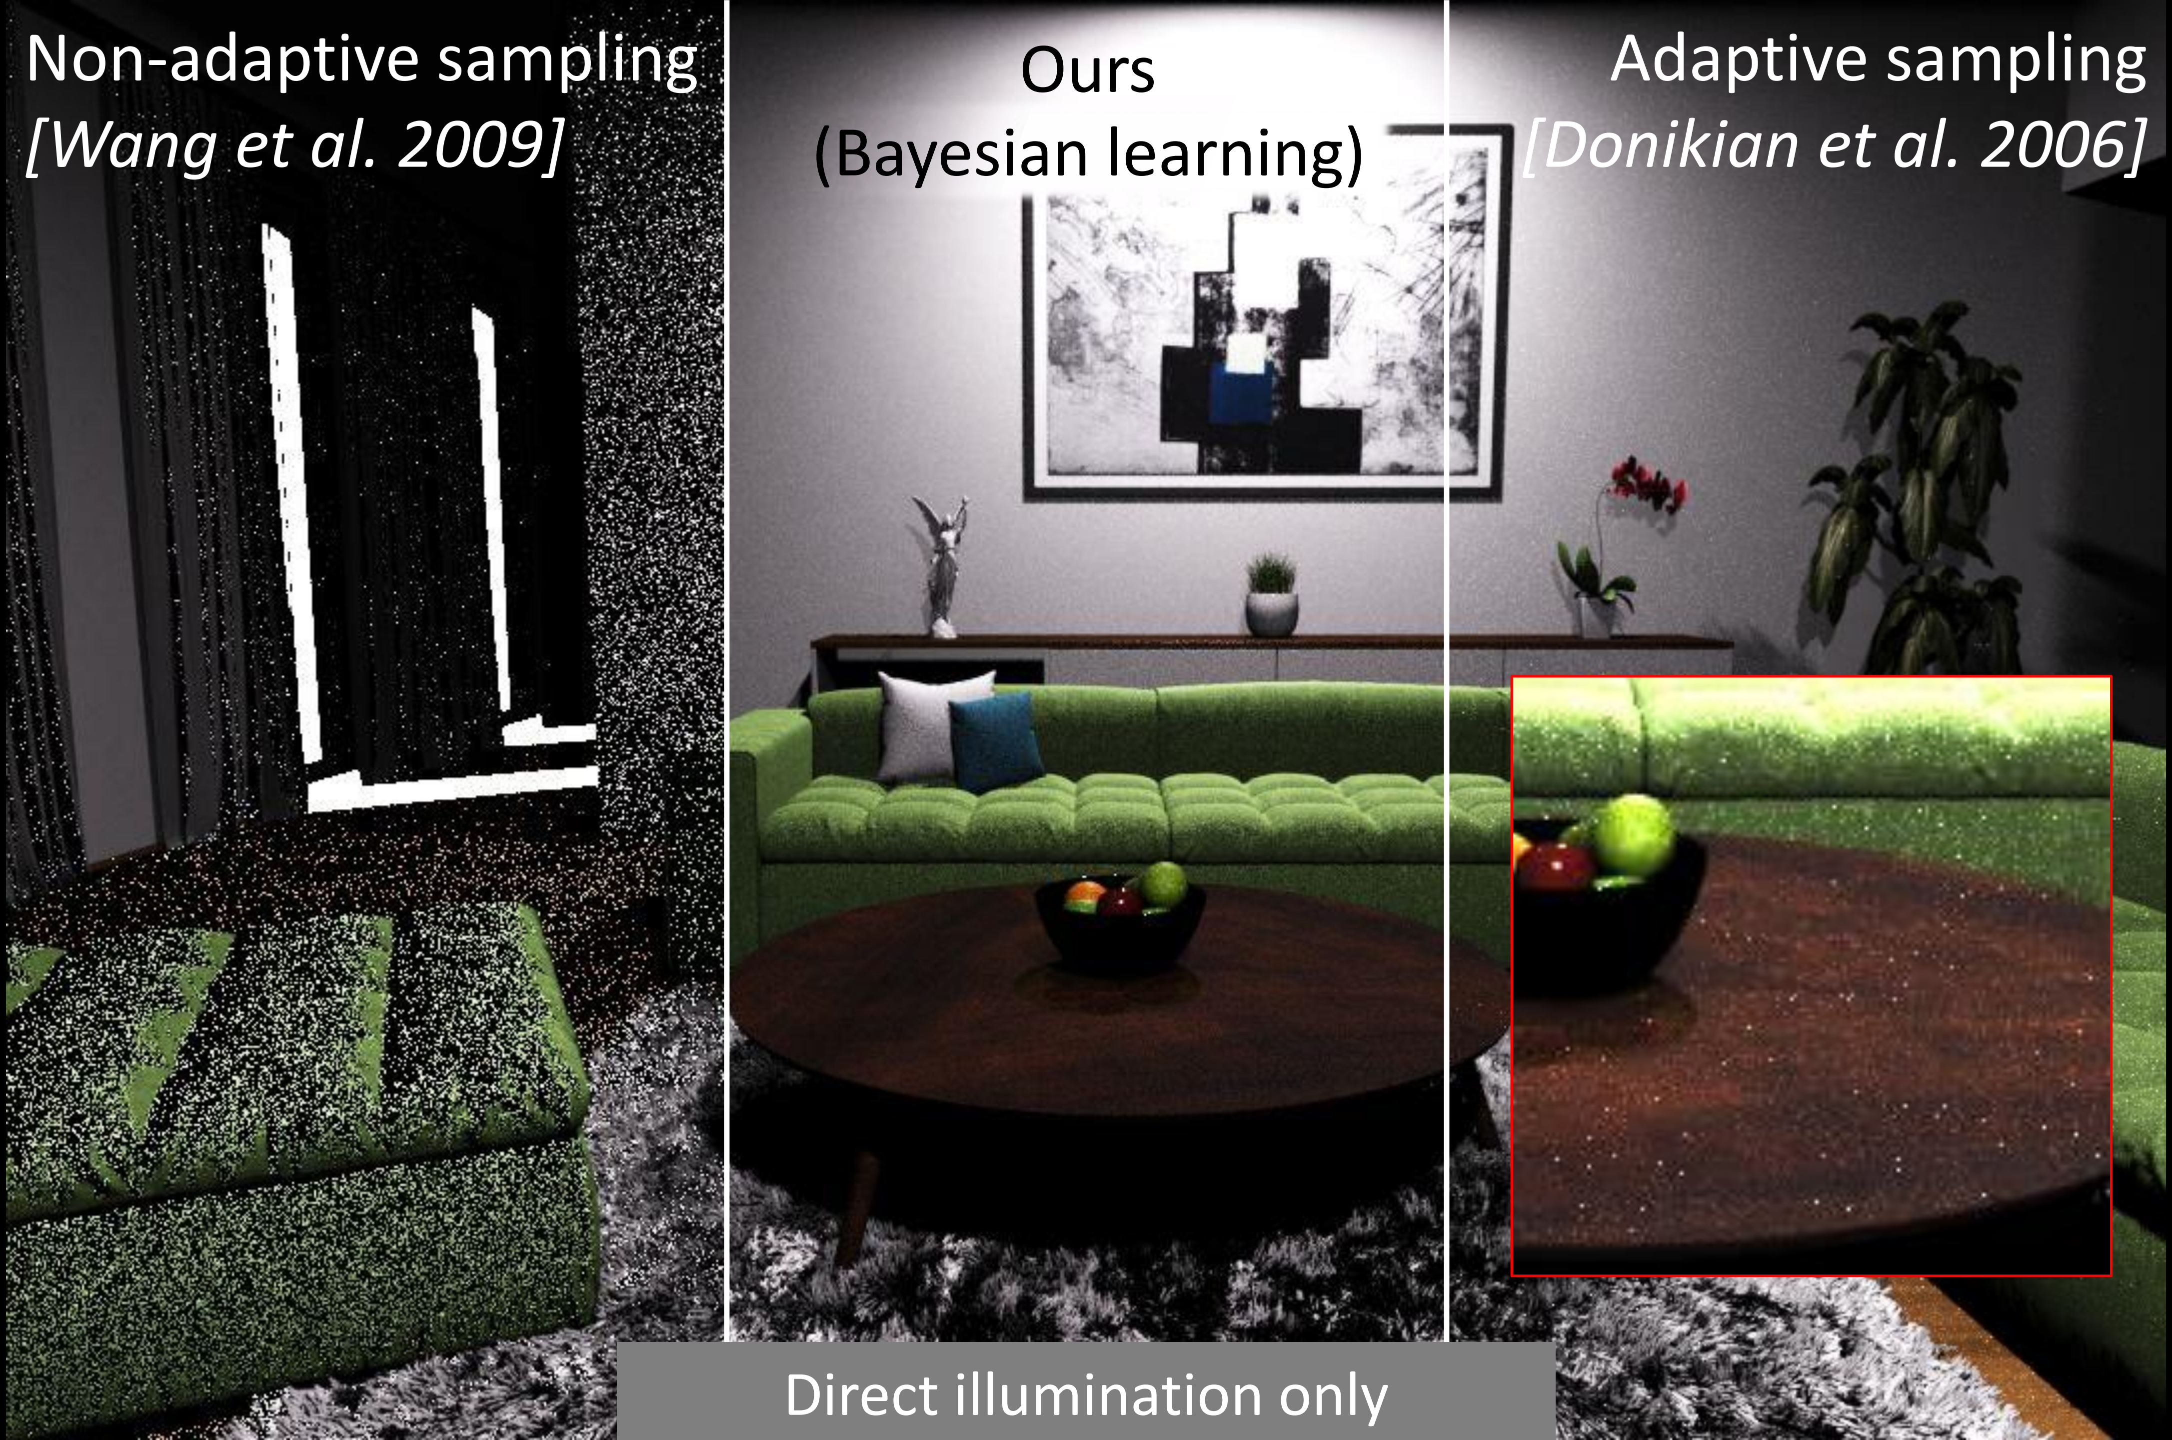
Non-adaptive sampling
[Wang et al. 2009]
Ours
(Bayesian learning)
Adaptive sampling
[Donikian et al. 2006]
Direct illumination only

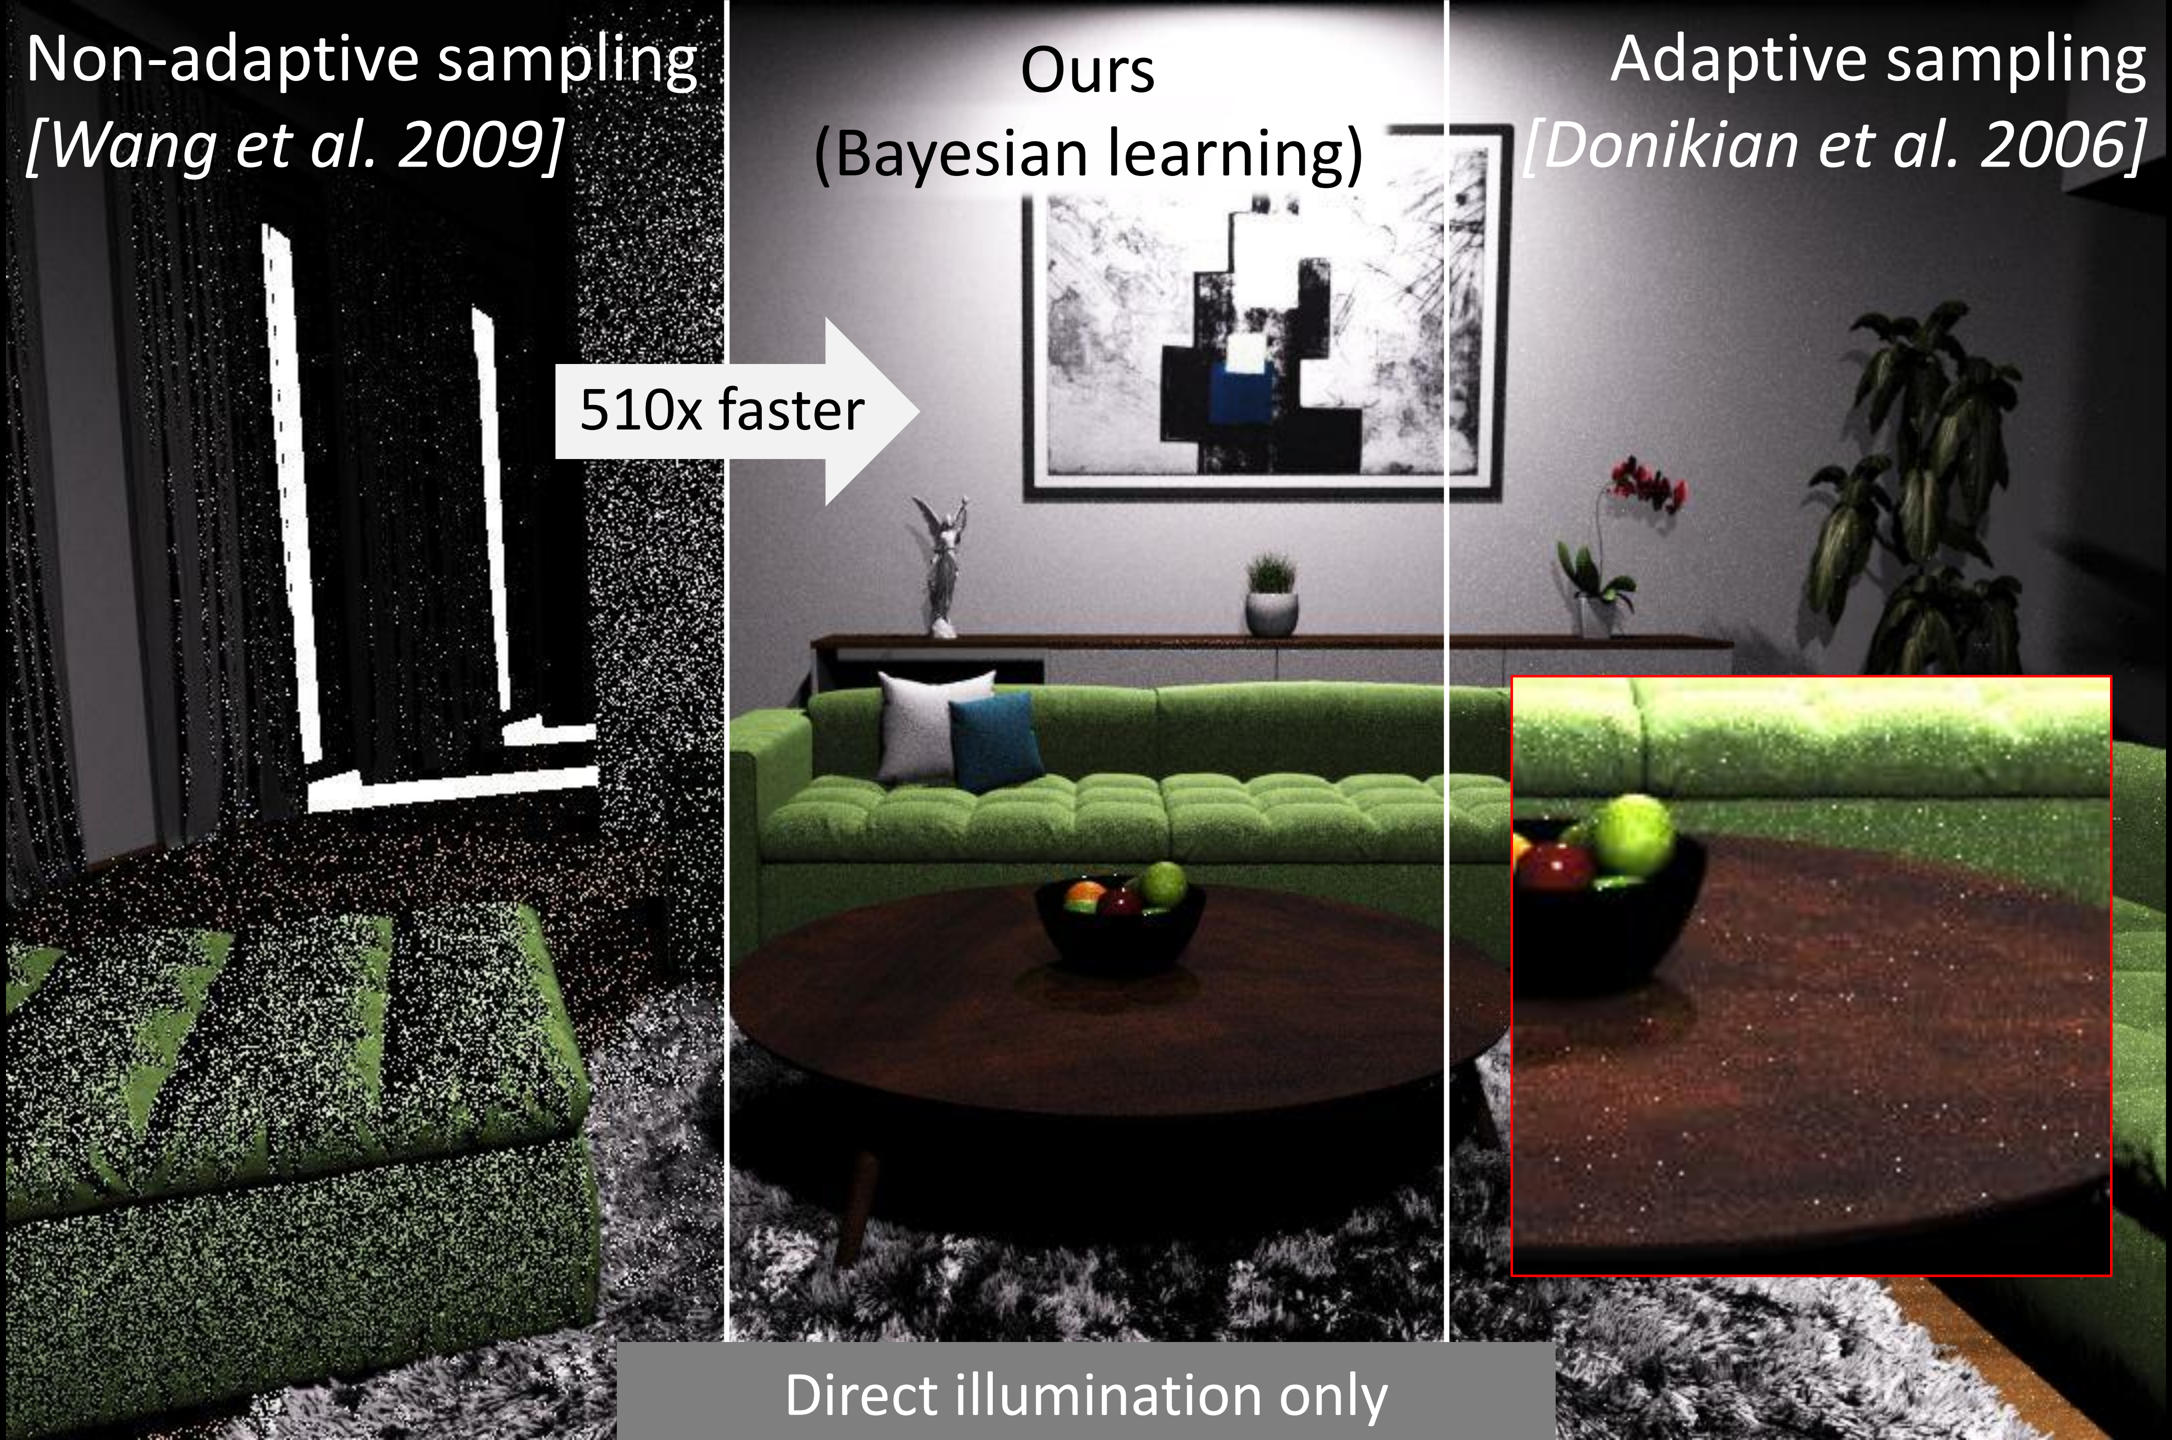
Non-adaptive sampling
[Wang et al. 2009]
Ours
(Bayesian learning)
Adaptive sampling
[Donikian et al. 2006]
510x faster
Direct illumination only

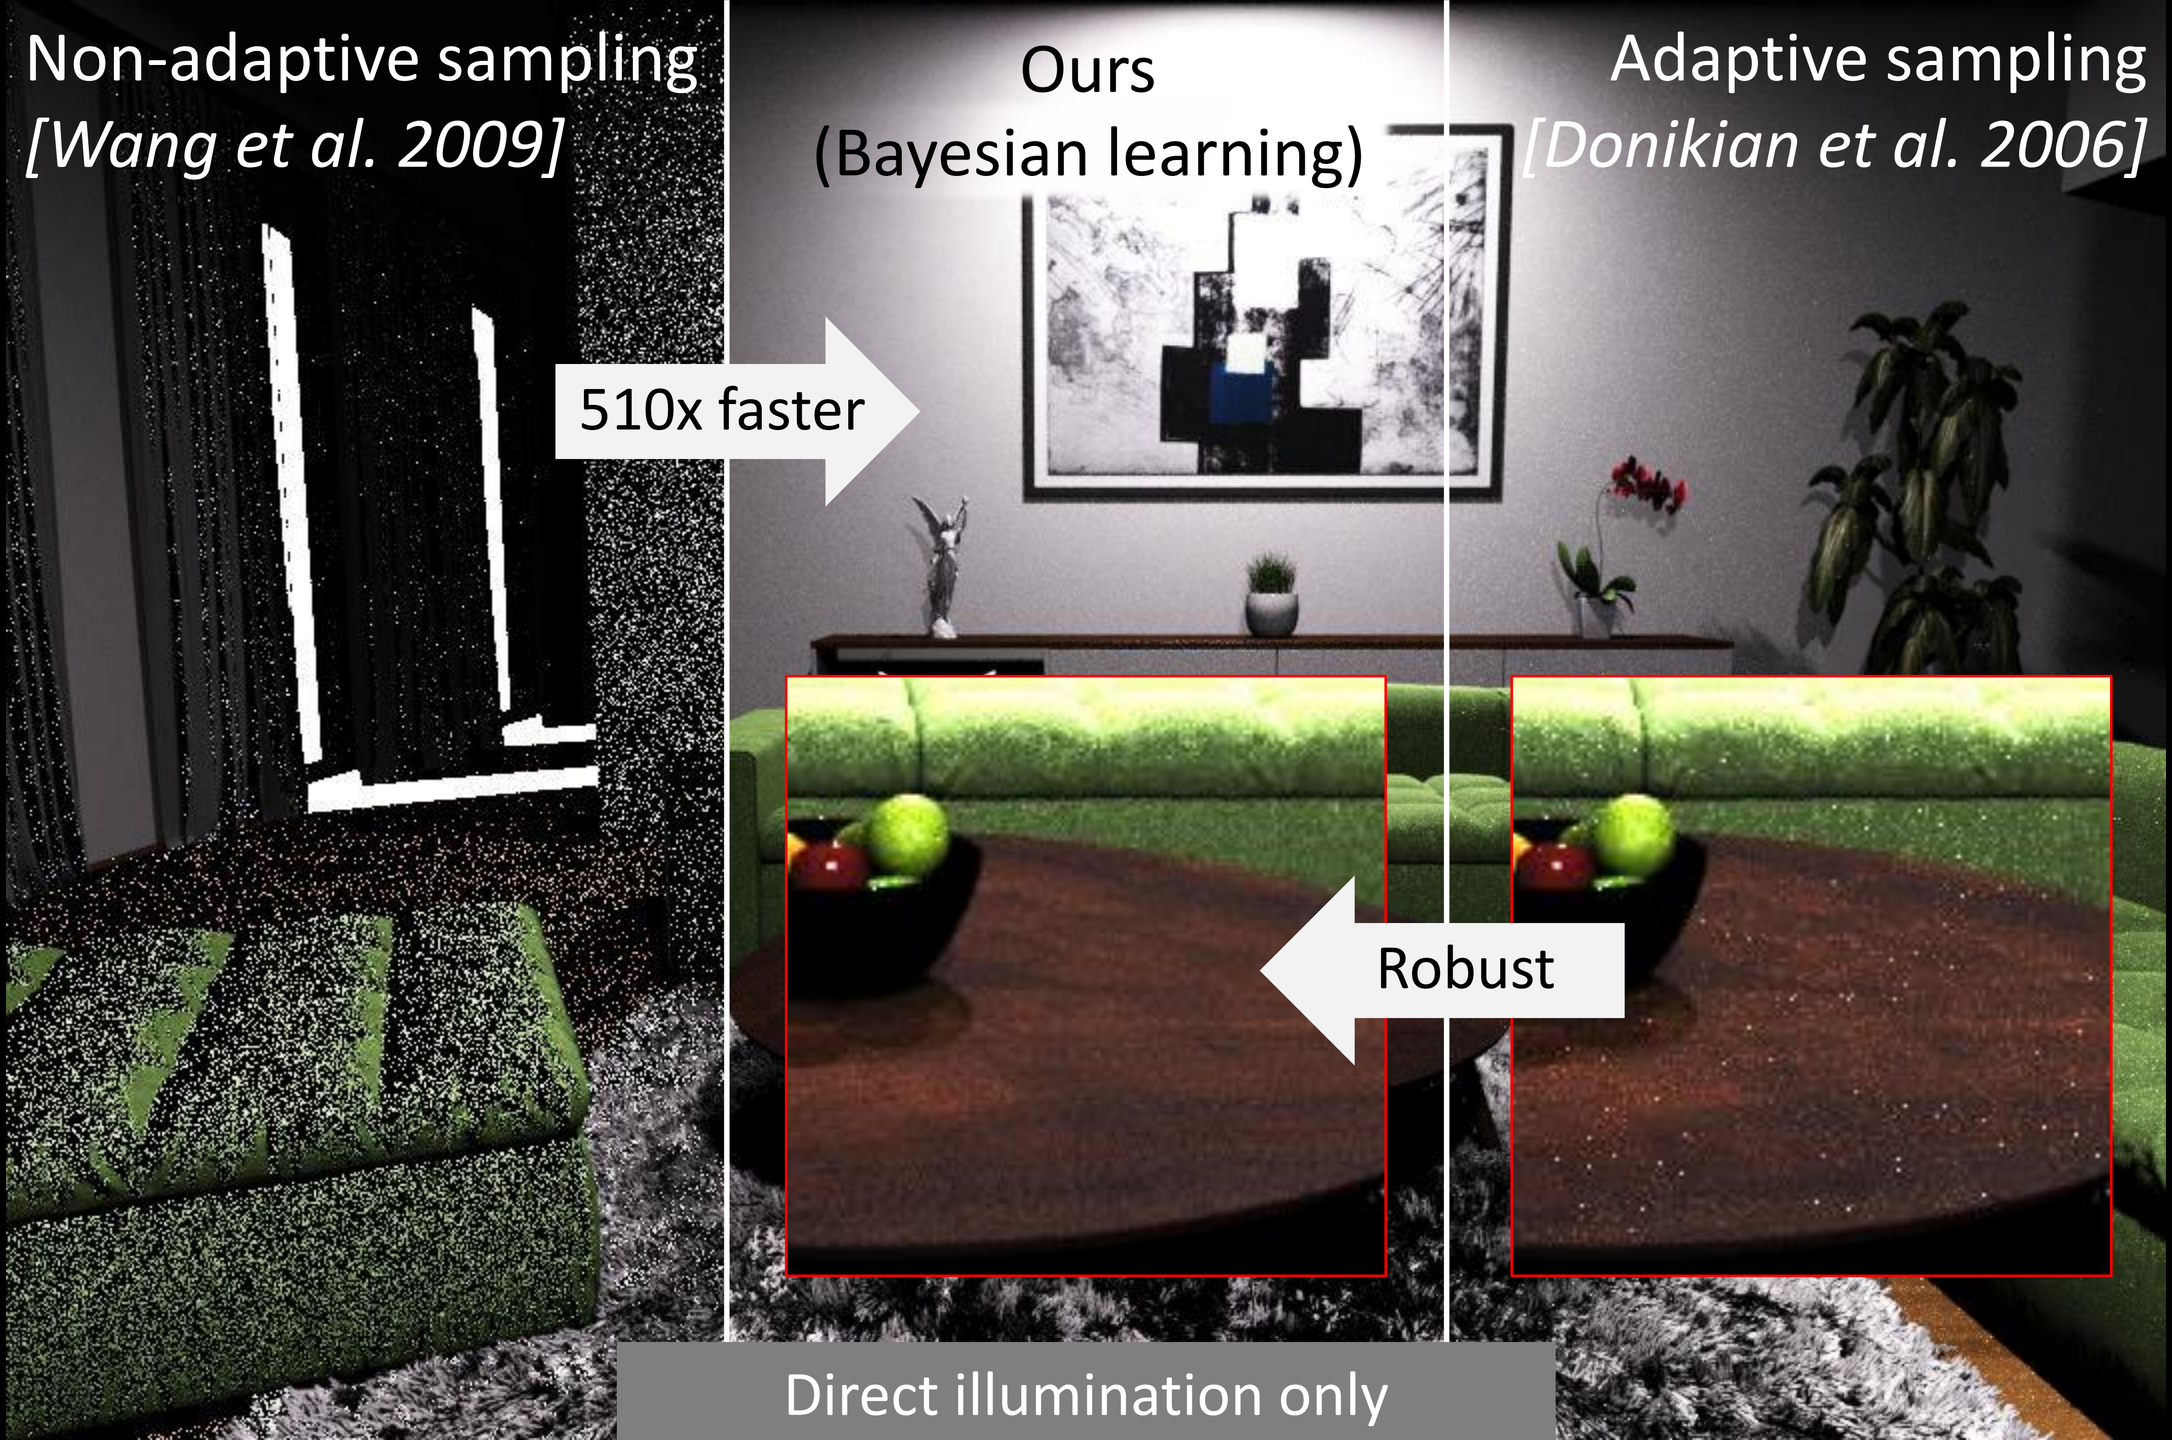
Non-adaptive sampling
[Wang et al. 2009]
Ours
(Bayesian learning)
Adaptive sampling
[Donikian et al. 2006]
510x faster
Robust
Direct illumination only

Previous work

Adaptive sampling

- General Monte Carlo
  - Vegas algorithm
    - [*Lepage 1980*]
  - Population MC
    - [*Cappé et al. 2004, ...*]
- Rendering
  - Image sampling
    - [*Mitchell 1987, ...*]
  - Indirect illumination (path guiding)
    - [*Dutre and Willems 1995, Jensen 1995, Lafortune et al. 1995, ...*]
    - [*Vorba et al. 2014, Muller et al. 2017*]
  - Direct illumination
    - [*Shirley et al. 1996, Donikian et al. 2006, Wang et al. 2009*]

Bayesian methods in rendering

- Filtering
  - NonLocal Bayes [*Boughida and Boubekeur 2017*]
- Global illumination
  - Bayesian Monte Carlo [*Brouilat et al. 2009, Marques et al. 2013*]
  - Path guiding [*Vorba et al. 2014*]

Background

Direct illumination

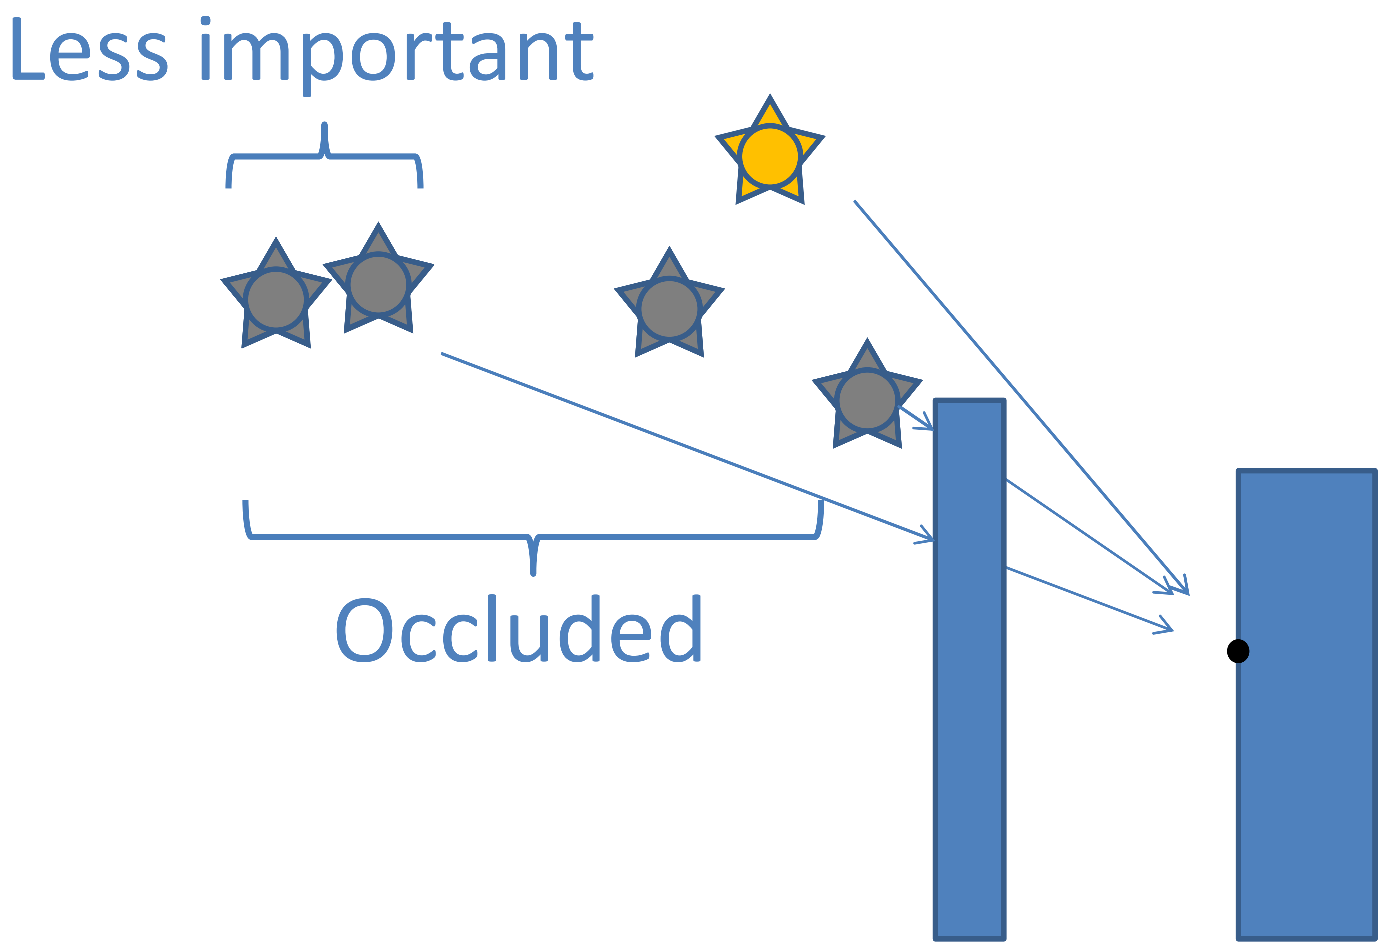

Clustering (Lightcuts)

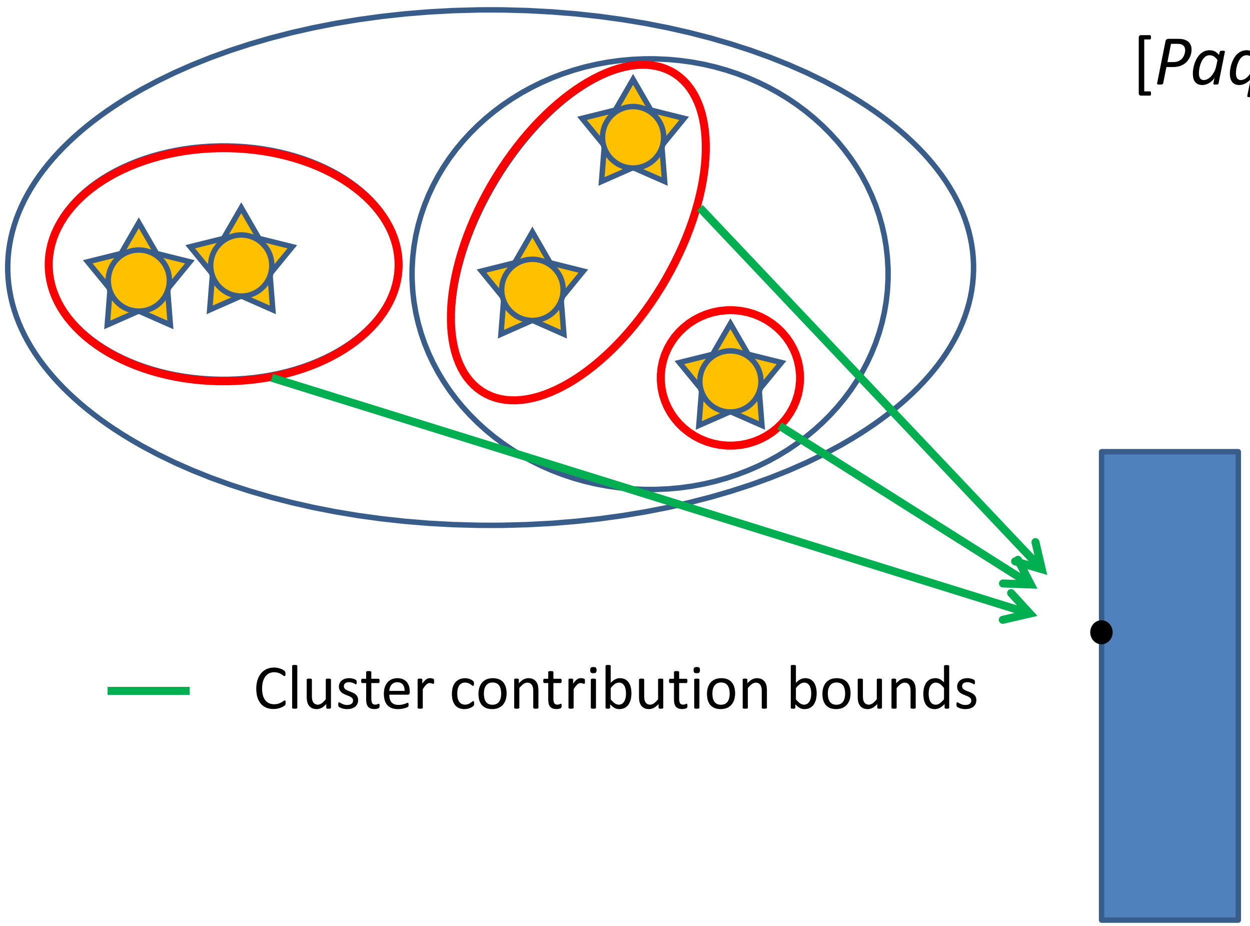

[Paquette et al. 1998,

Walter et al. 2006]

Cluster sampling

[*Wang and Akerlung 2009*]

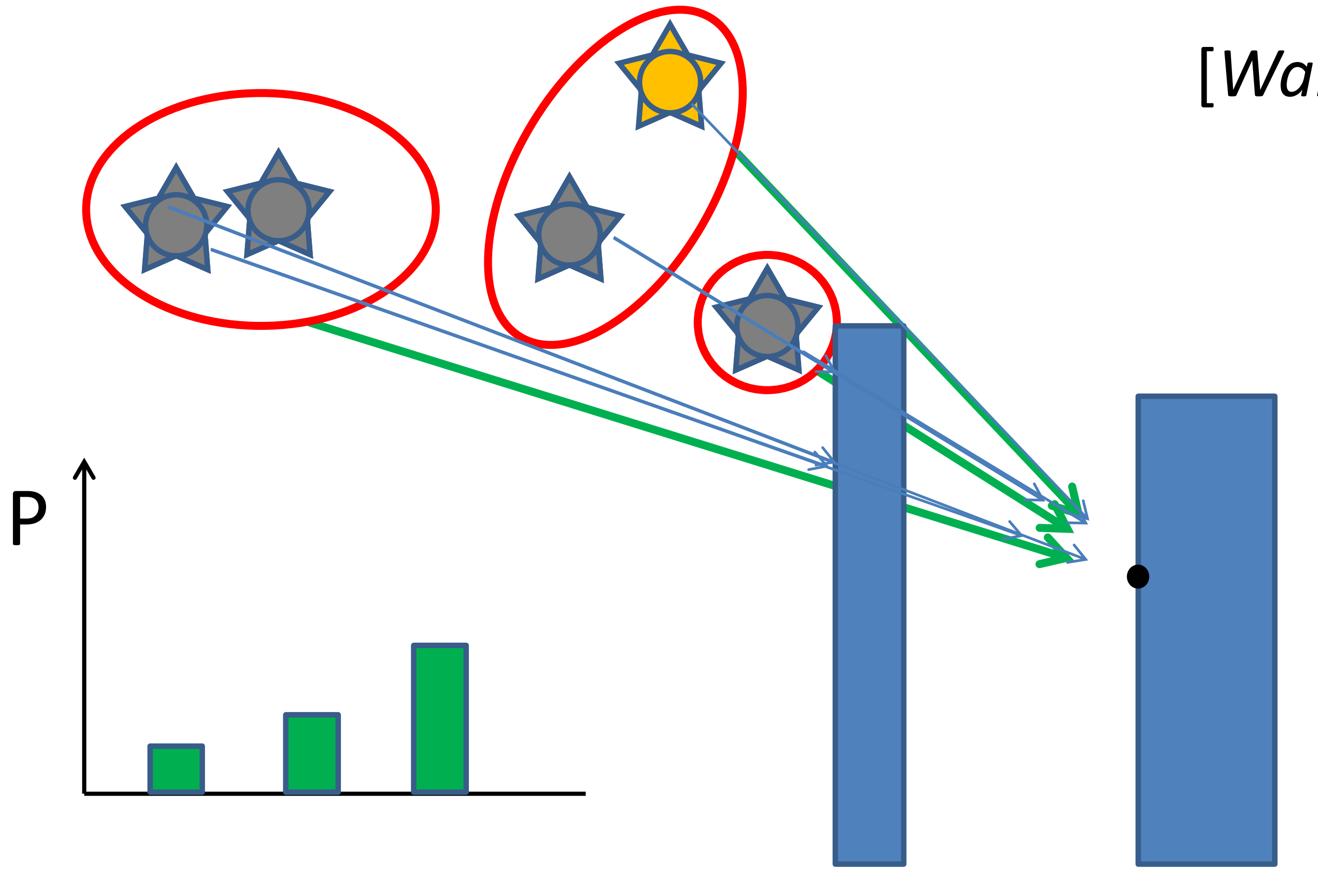

Adaptive light sampling

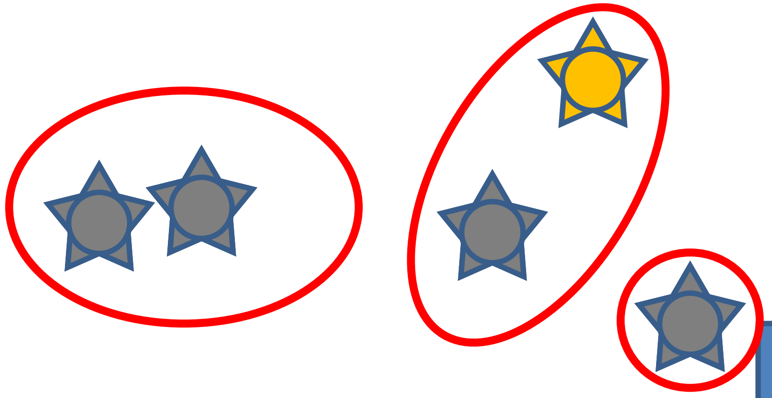

[*Donikian et al. 2006*]

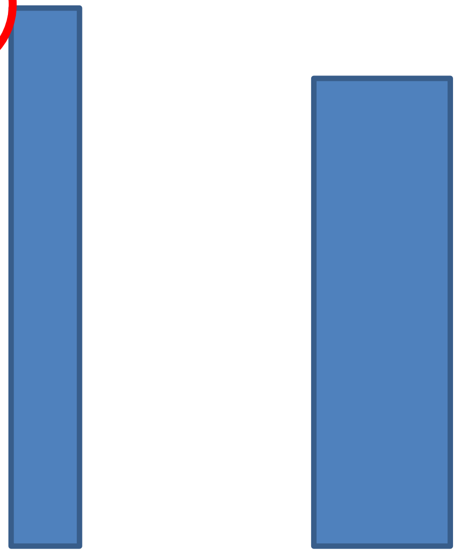

screen space

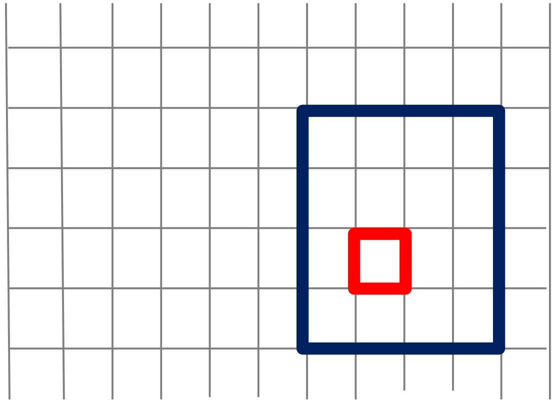

Ad-hoc combination

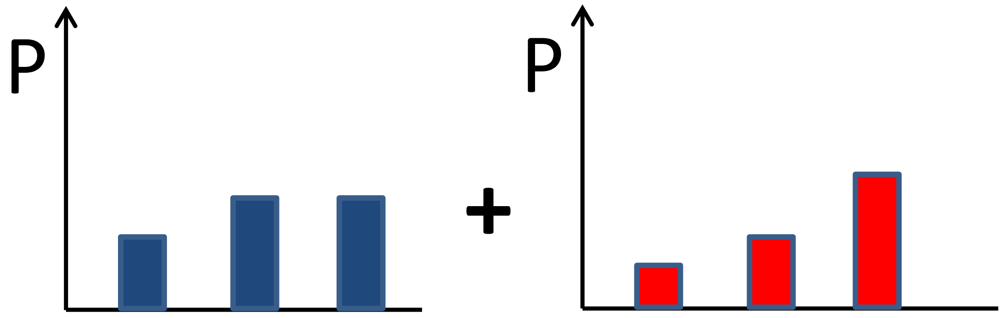

Problem summary

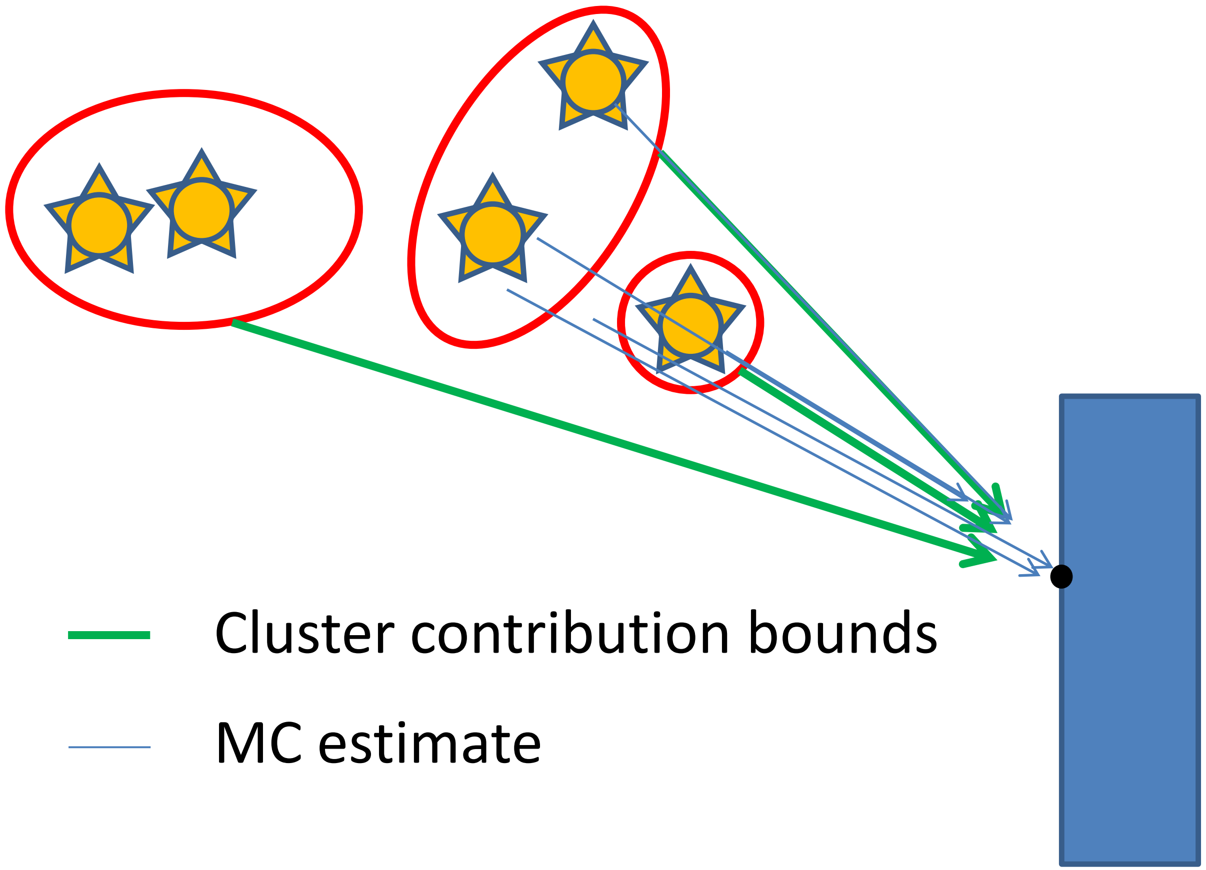

Our approach

Contributions

- Optimal sampling of clusters
- Adaptive sampling by Bayesian inference

Optimal cluster sampling

$$P(C) \propto \sqrt{\text{mean}^2 + \text{variance}}$$

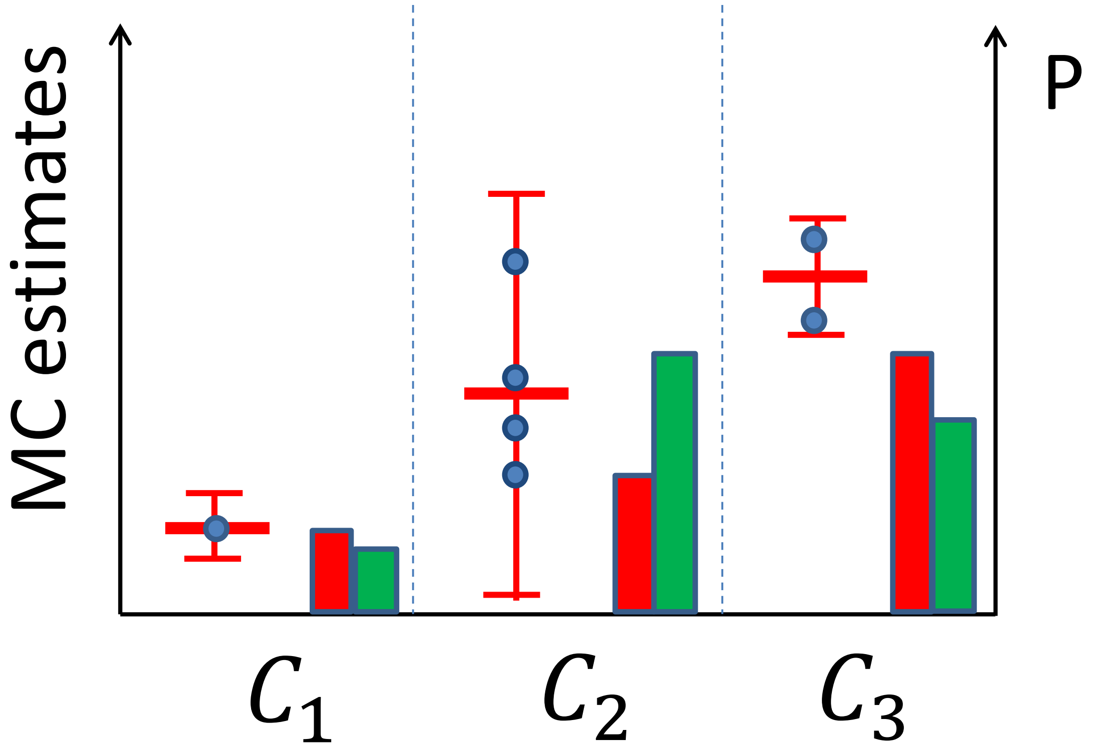

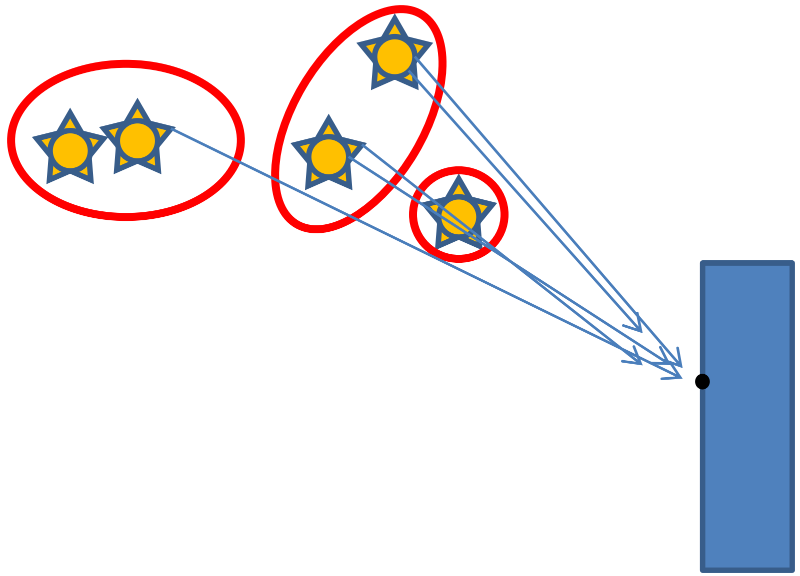

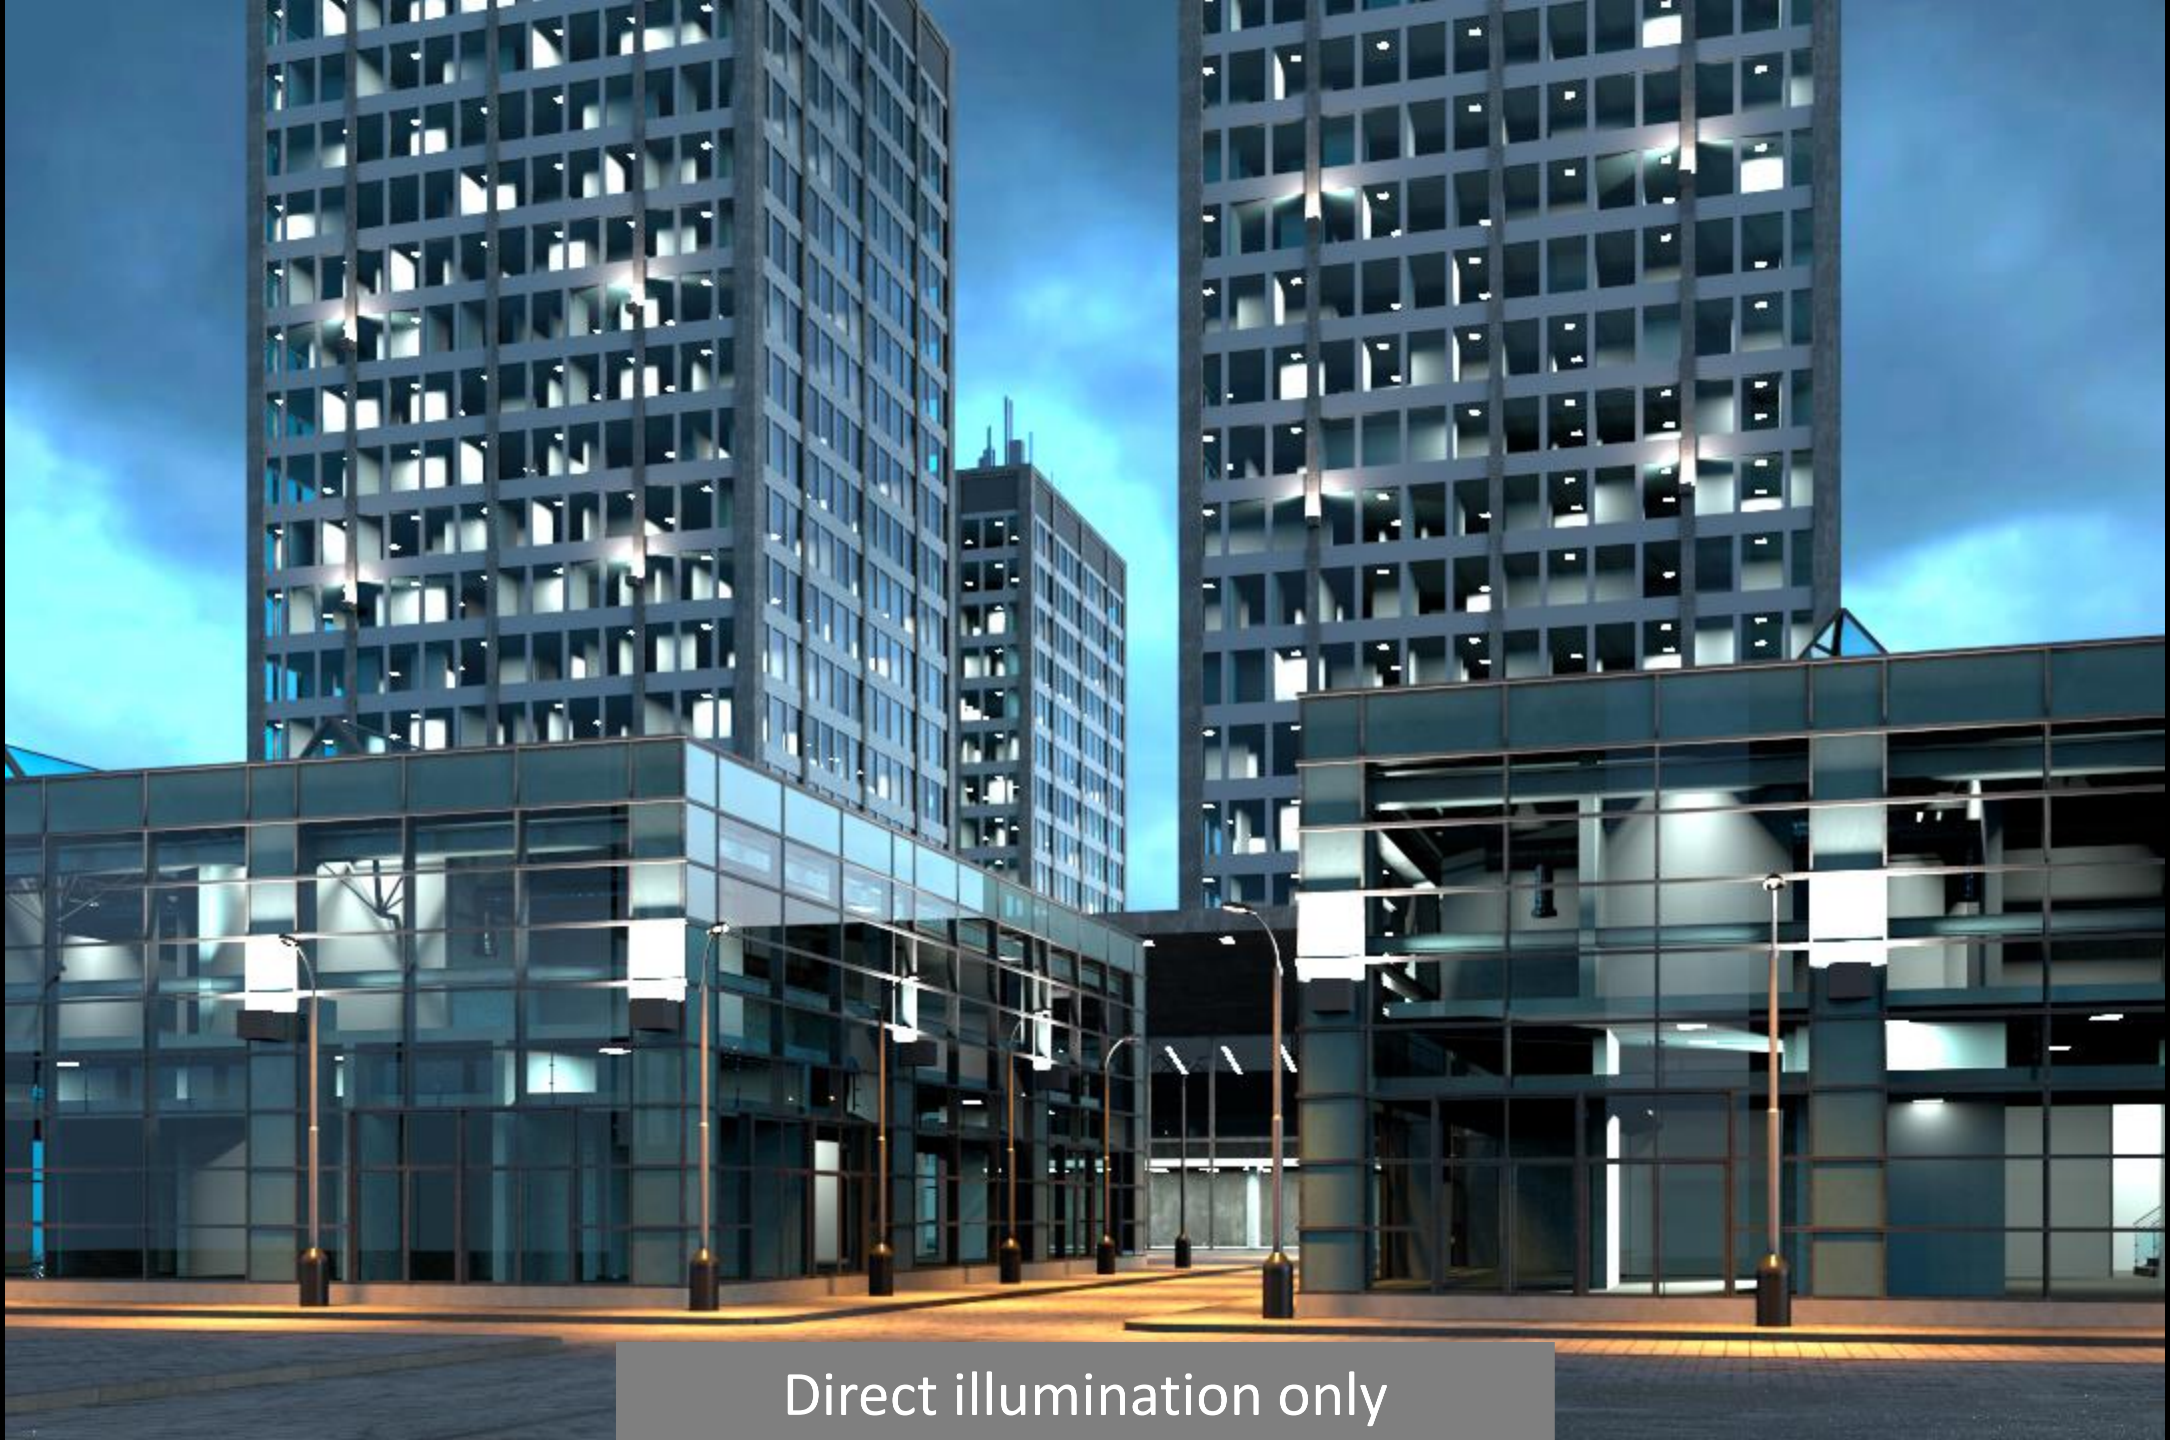

Direct illumination only

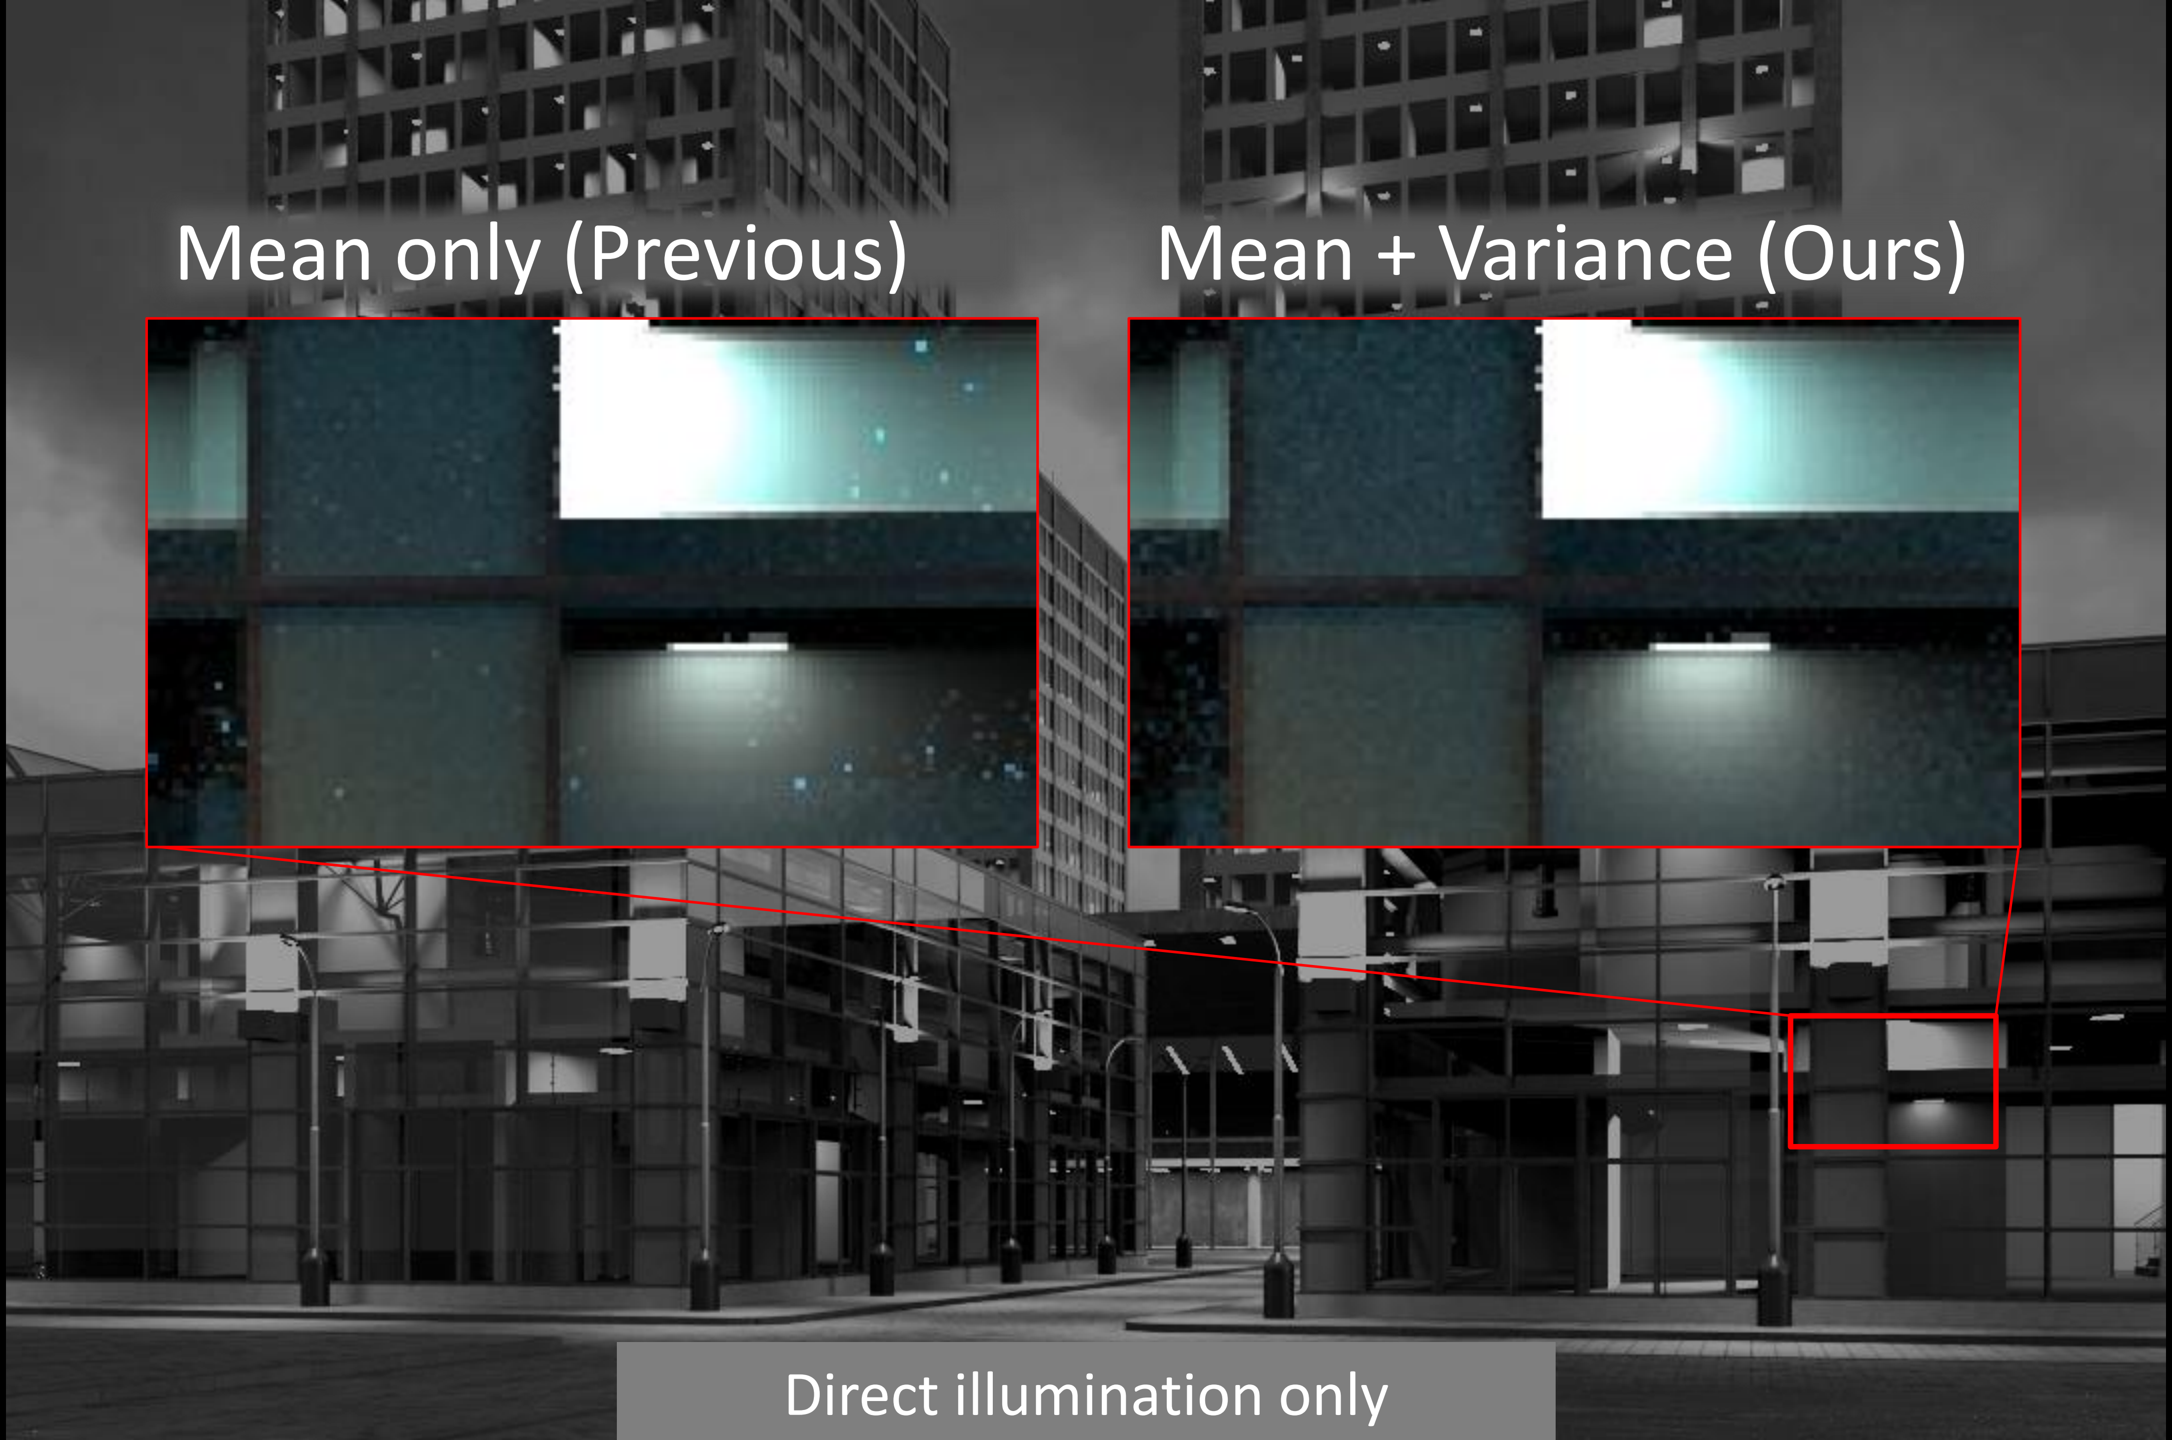
Mean only (Previous)
Mean + Variance (Ours)
Direct illumination only

Contributions

- Optimal sampling of clusters
- Adaptive sampling by Bayesian inference

Naive adaptive cluster sampling

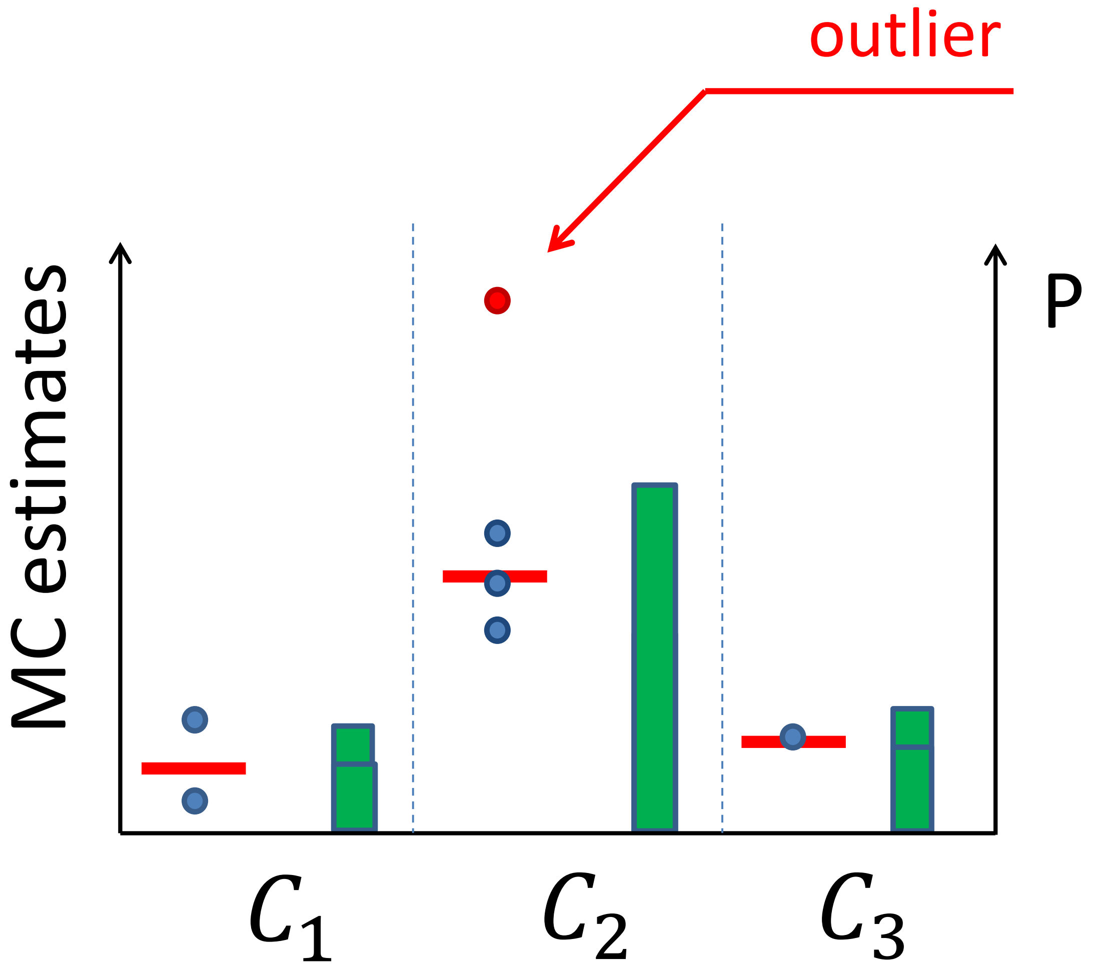

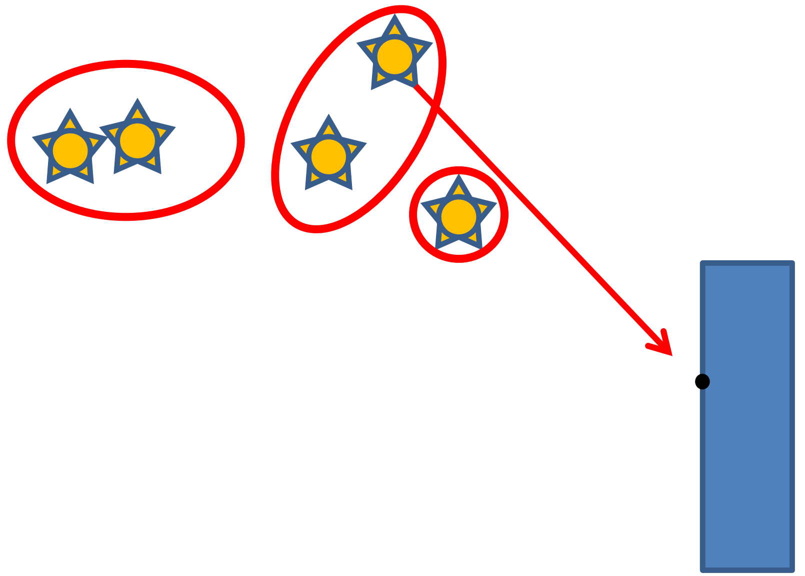

Bayes cluster adaptive sampling

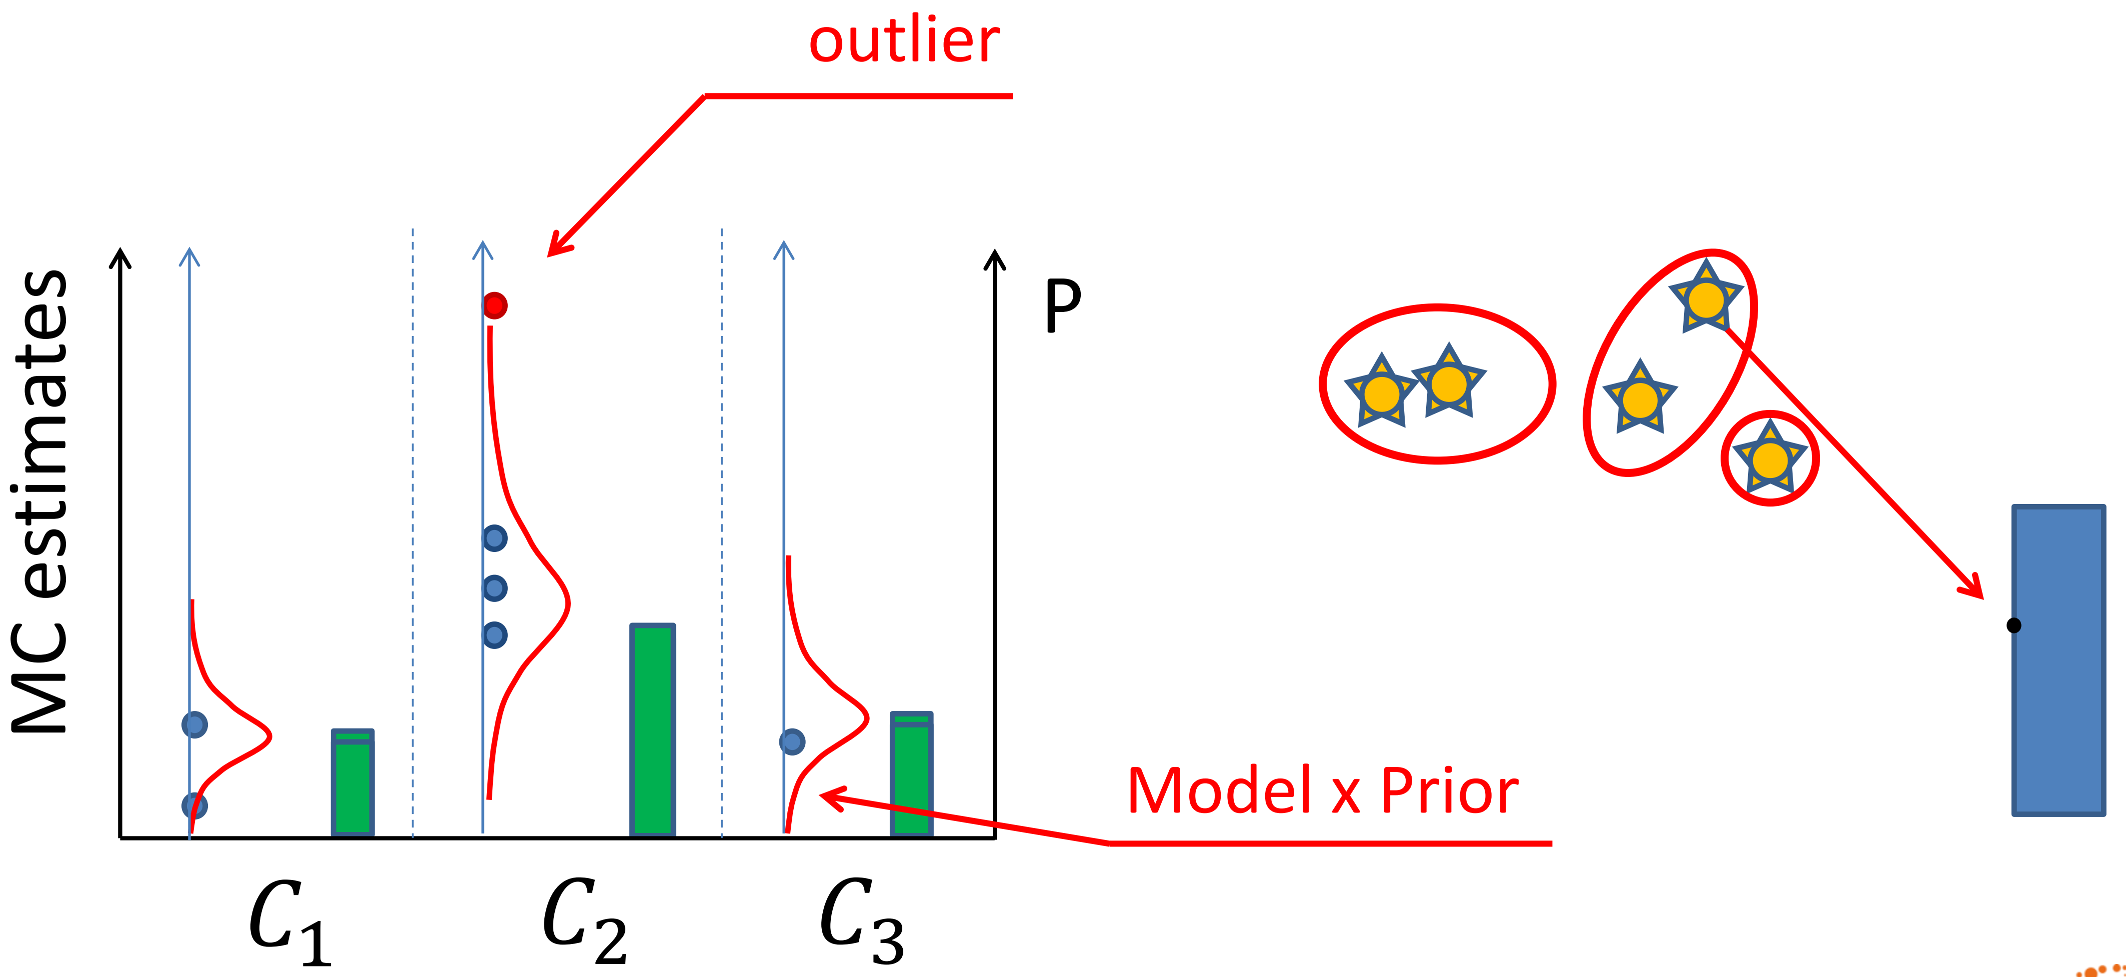

Cluster-region pairs

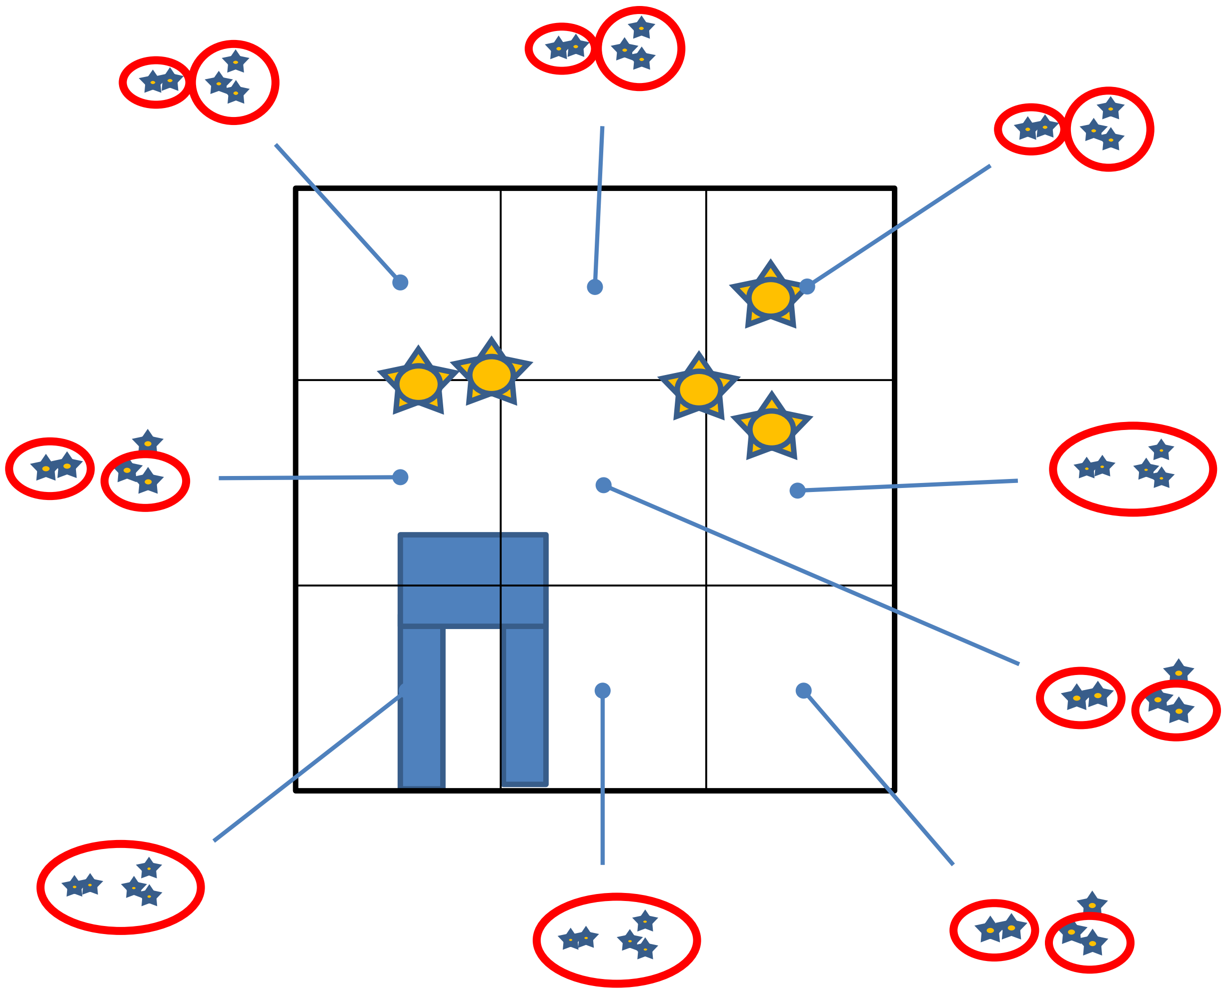

Cluster-Region data

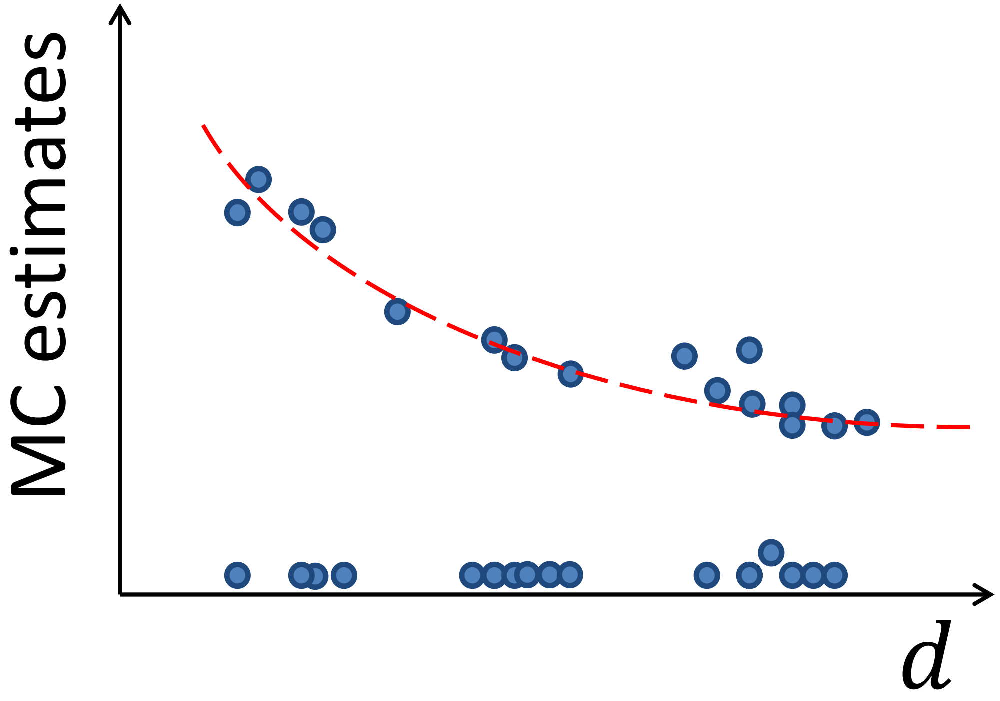

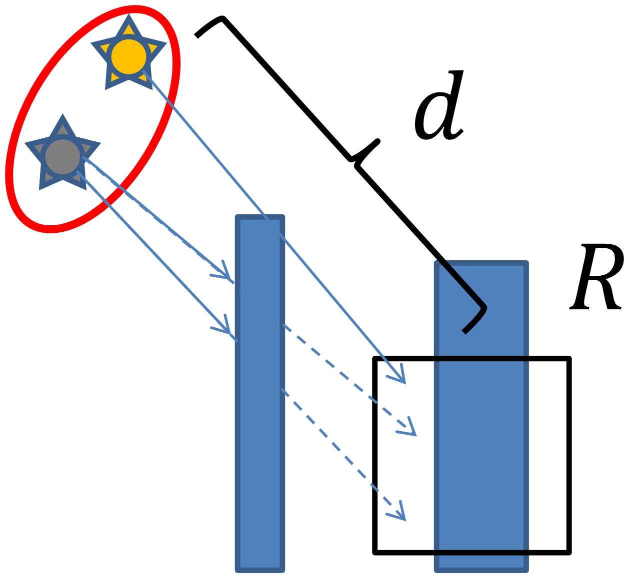

Regresion Data model

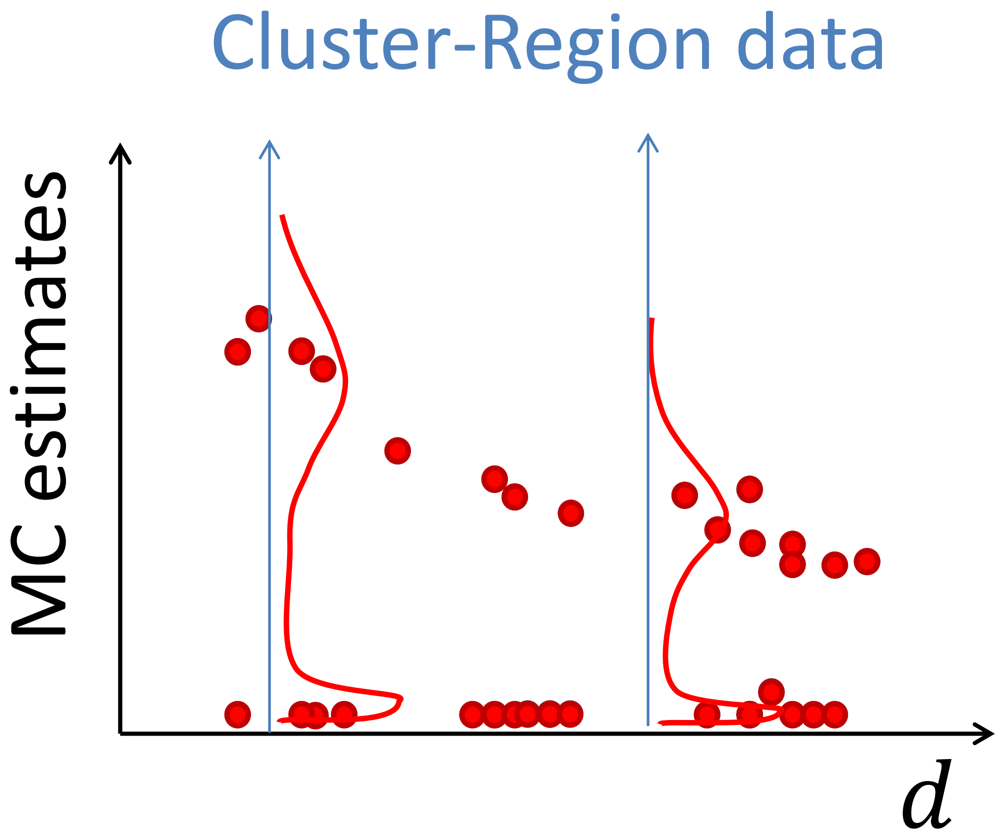

Parameters:

k, h - normal distr. parameters

p_0 - probability of occlusion

$$(1 - p_0) \times N(\text{est.} \mid \frac{k}{d^2}, \frac{h}{d^4})$$

$$p_0 \times \delta(\text{est.})$$

Conjugate prior

$$\underbrace{\textbf{posterior}}_{\text{Same functional form}} \propto \text{likelihood} \times \underbrace{\textbf{prior}}$$

Same functional form

Our (conjugate) Priors

$$\mathrm{p}_0 \sim \mathrm{Beta}(p_0 | \underbrace{\ \dots\ }_{\text{Hyperparameters}})$$

$$\mathrm{k, h} \sim \mathrm{Normal\ inverse\ gamma}(k, h | \underbrace{\mu_0}_{\text{Cluster contrib. estimate}}, \underbrace{\dots}_{\text{Hyperparameters}})$$

Hyperparameters

Cluster contrib. estimate

Algorithm summary

- Light preprocess (clustering)
- During each Next event estimation:
  - Obtain clustering (Cut) cached in a region
  - Compute distributions of estimates for each cluster in Cut
    -> mean, variance
  - Build distribution over clusters
  - Sample direct illumination
  - Record new data for sampled cluster

Results

Direct illumination only

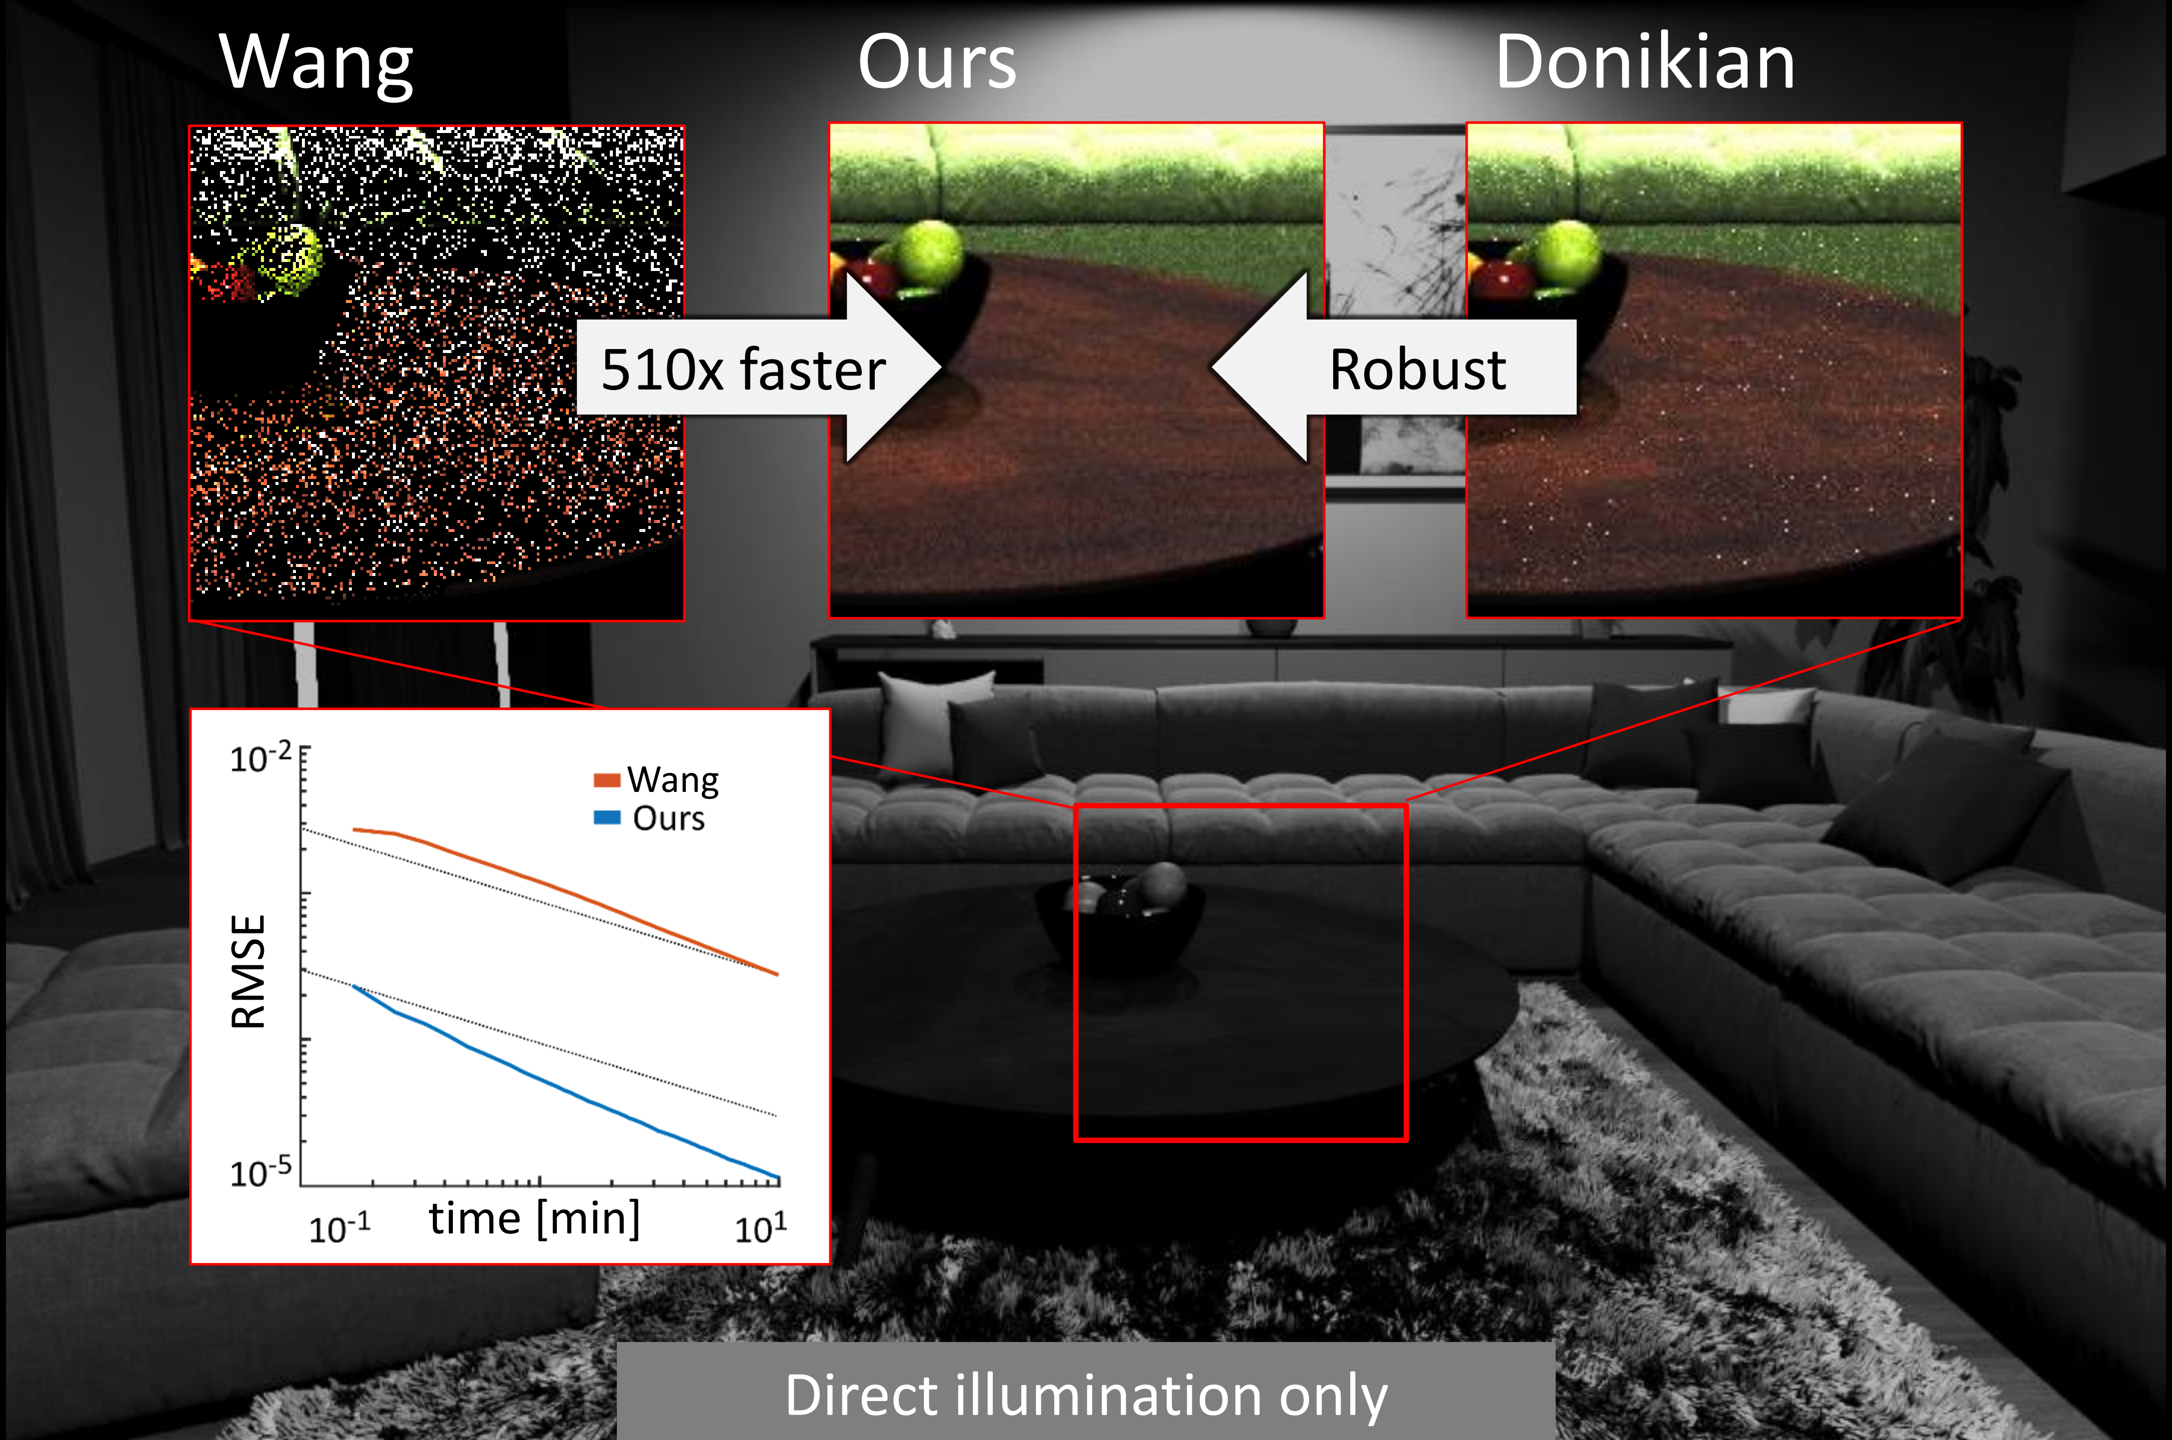

Direct illumination only

Direct + indirect illumination

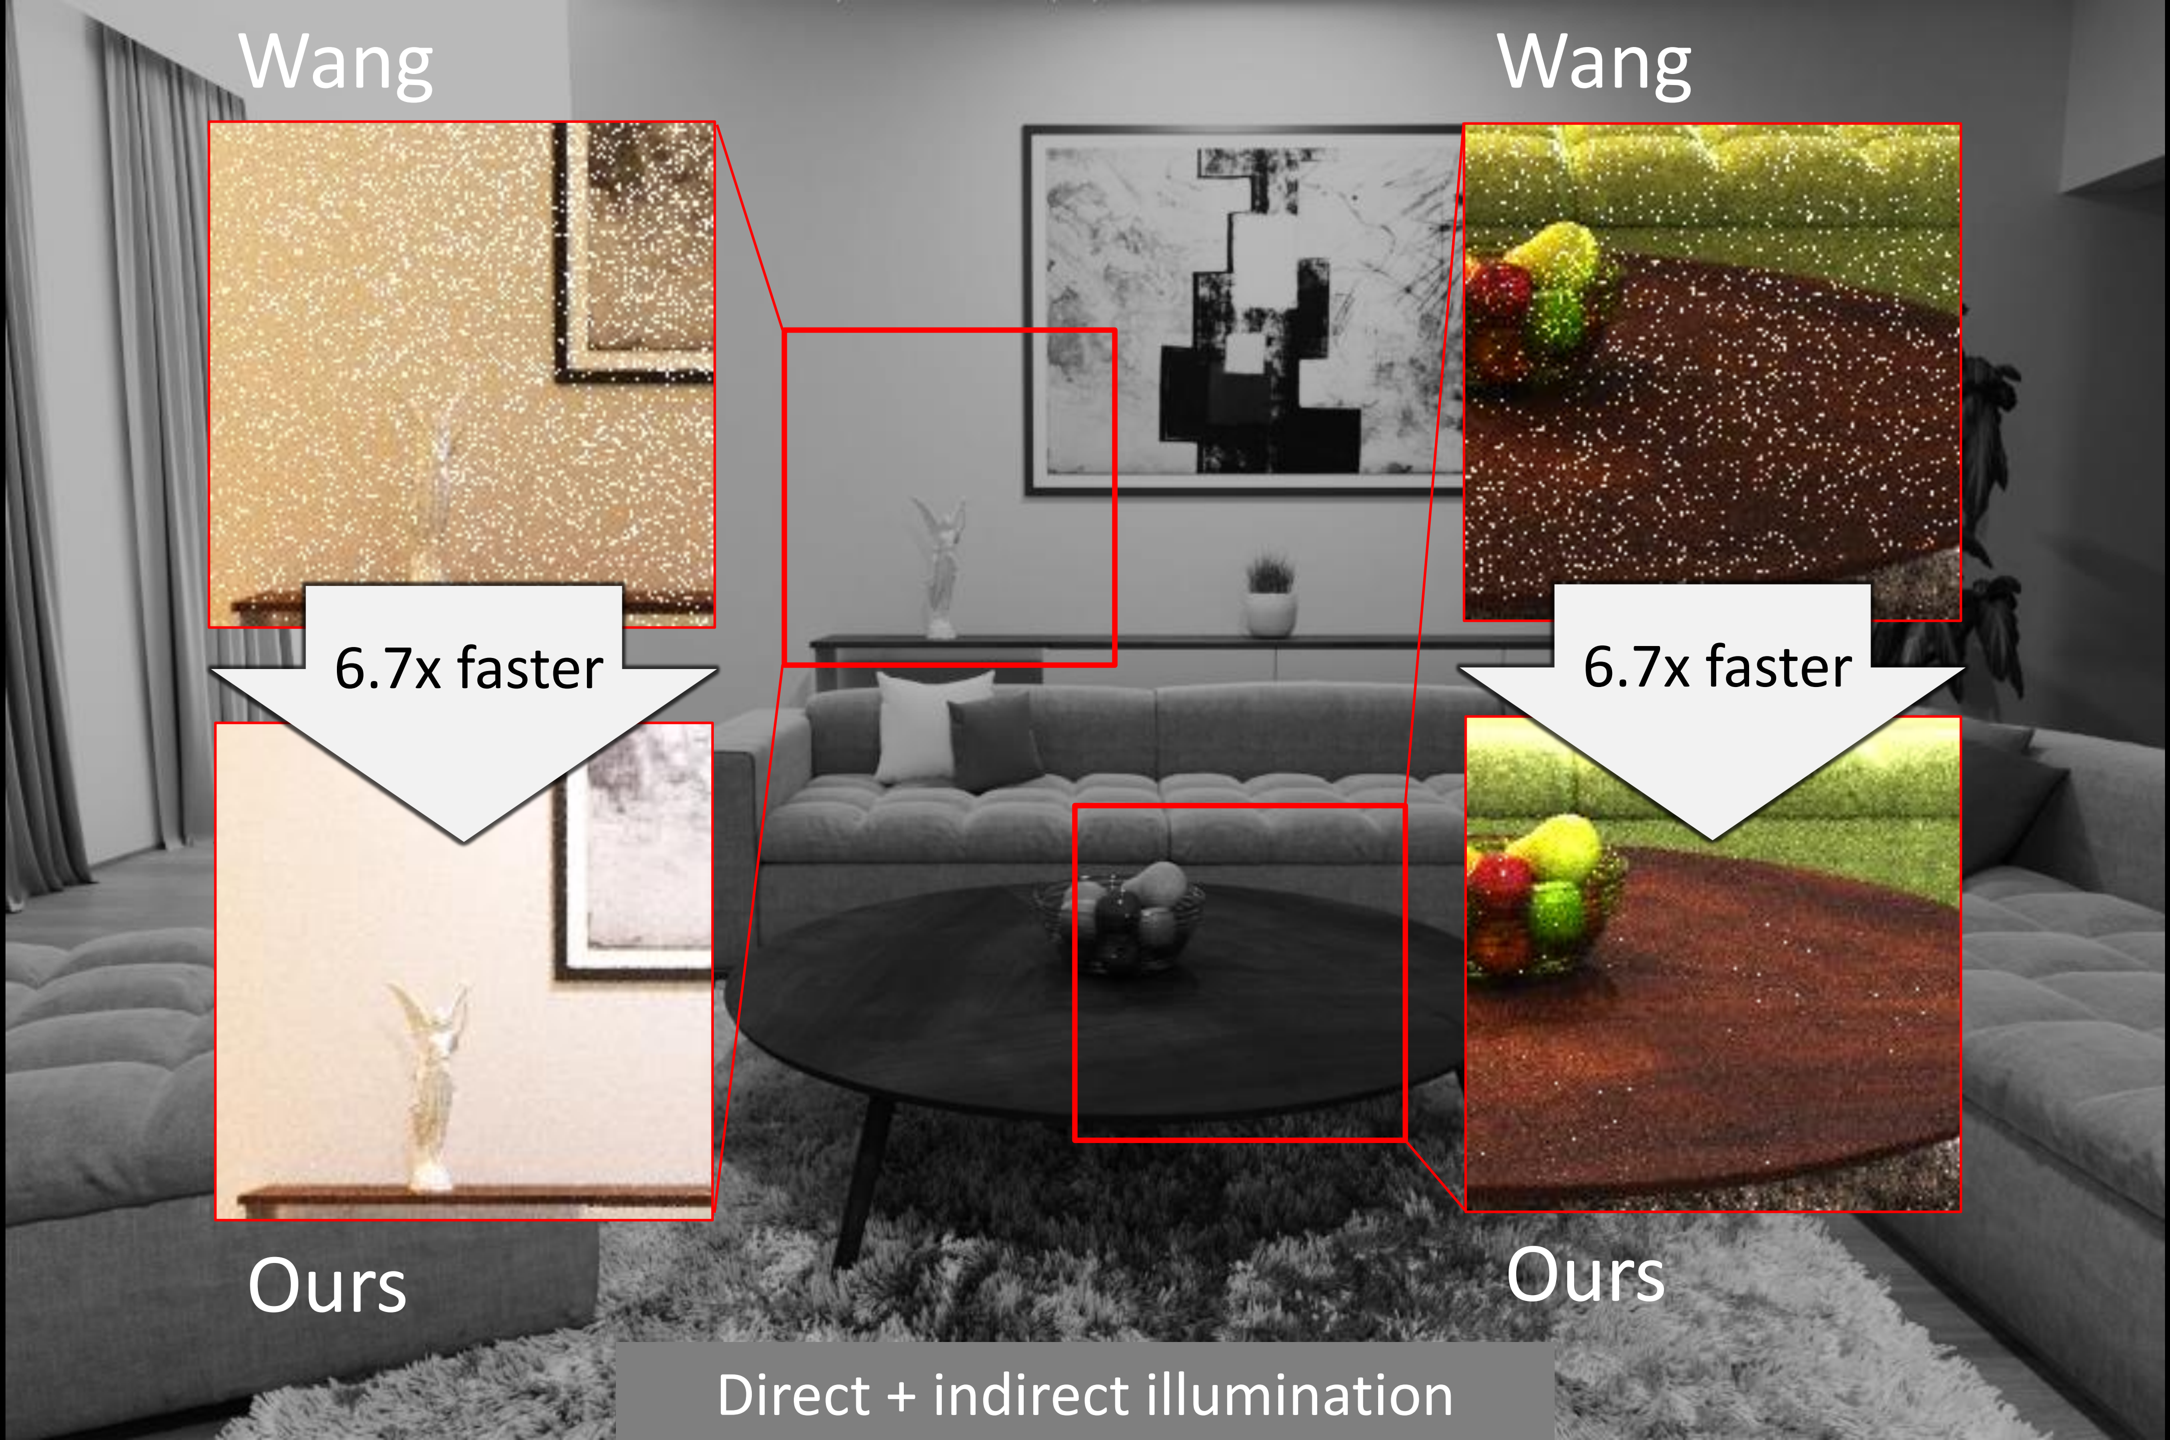
Wang
Wang
6.7x faster
6.7x faster
Ours
Ours
Direct + indirect illumination

Direct illumination only

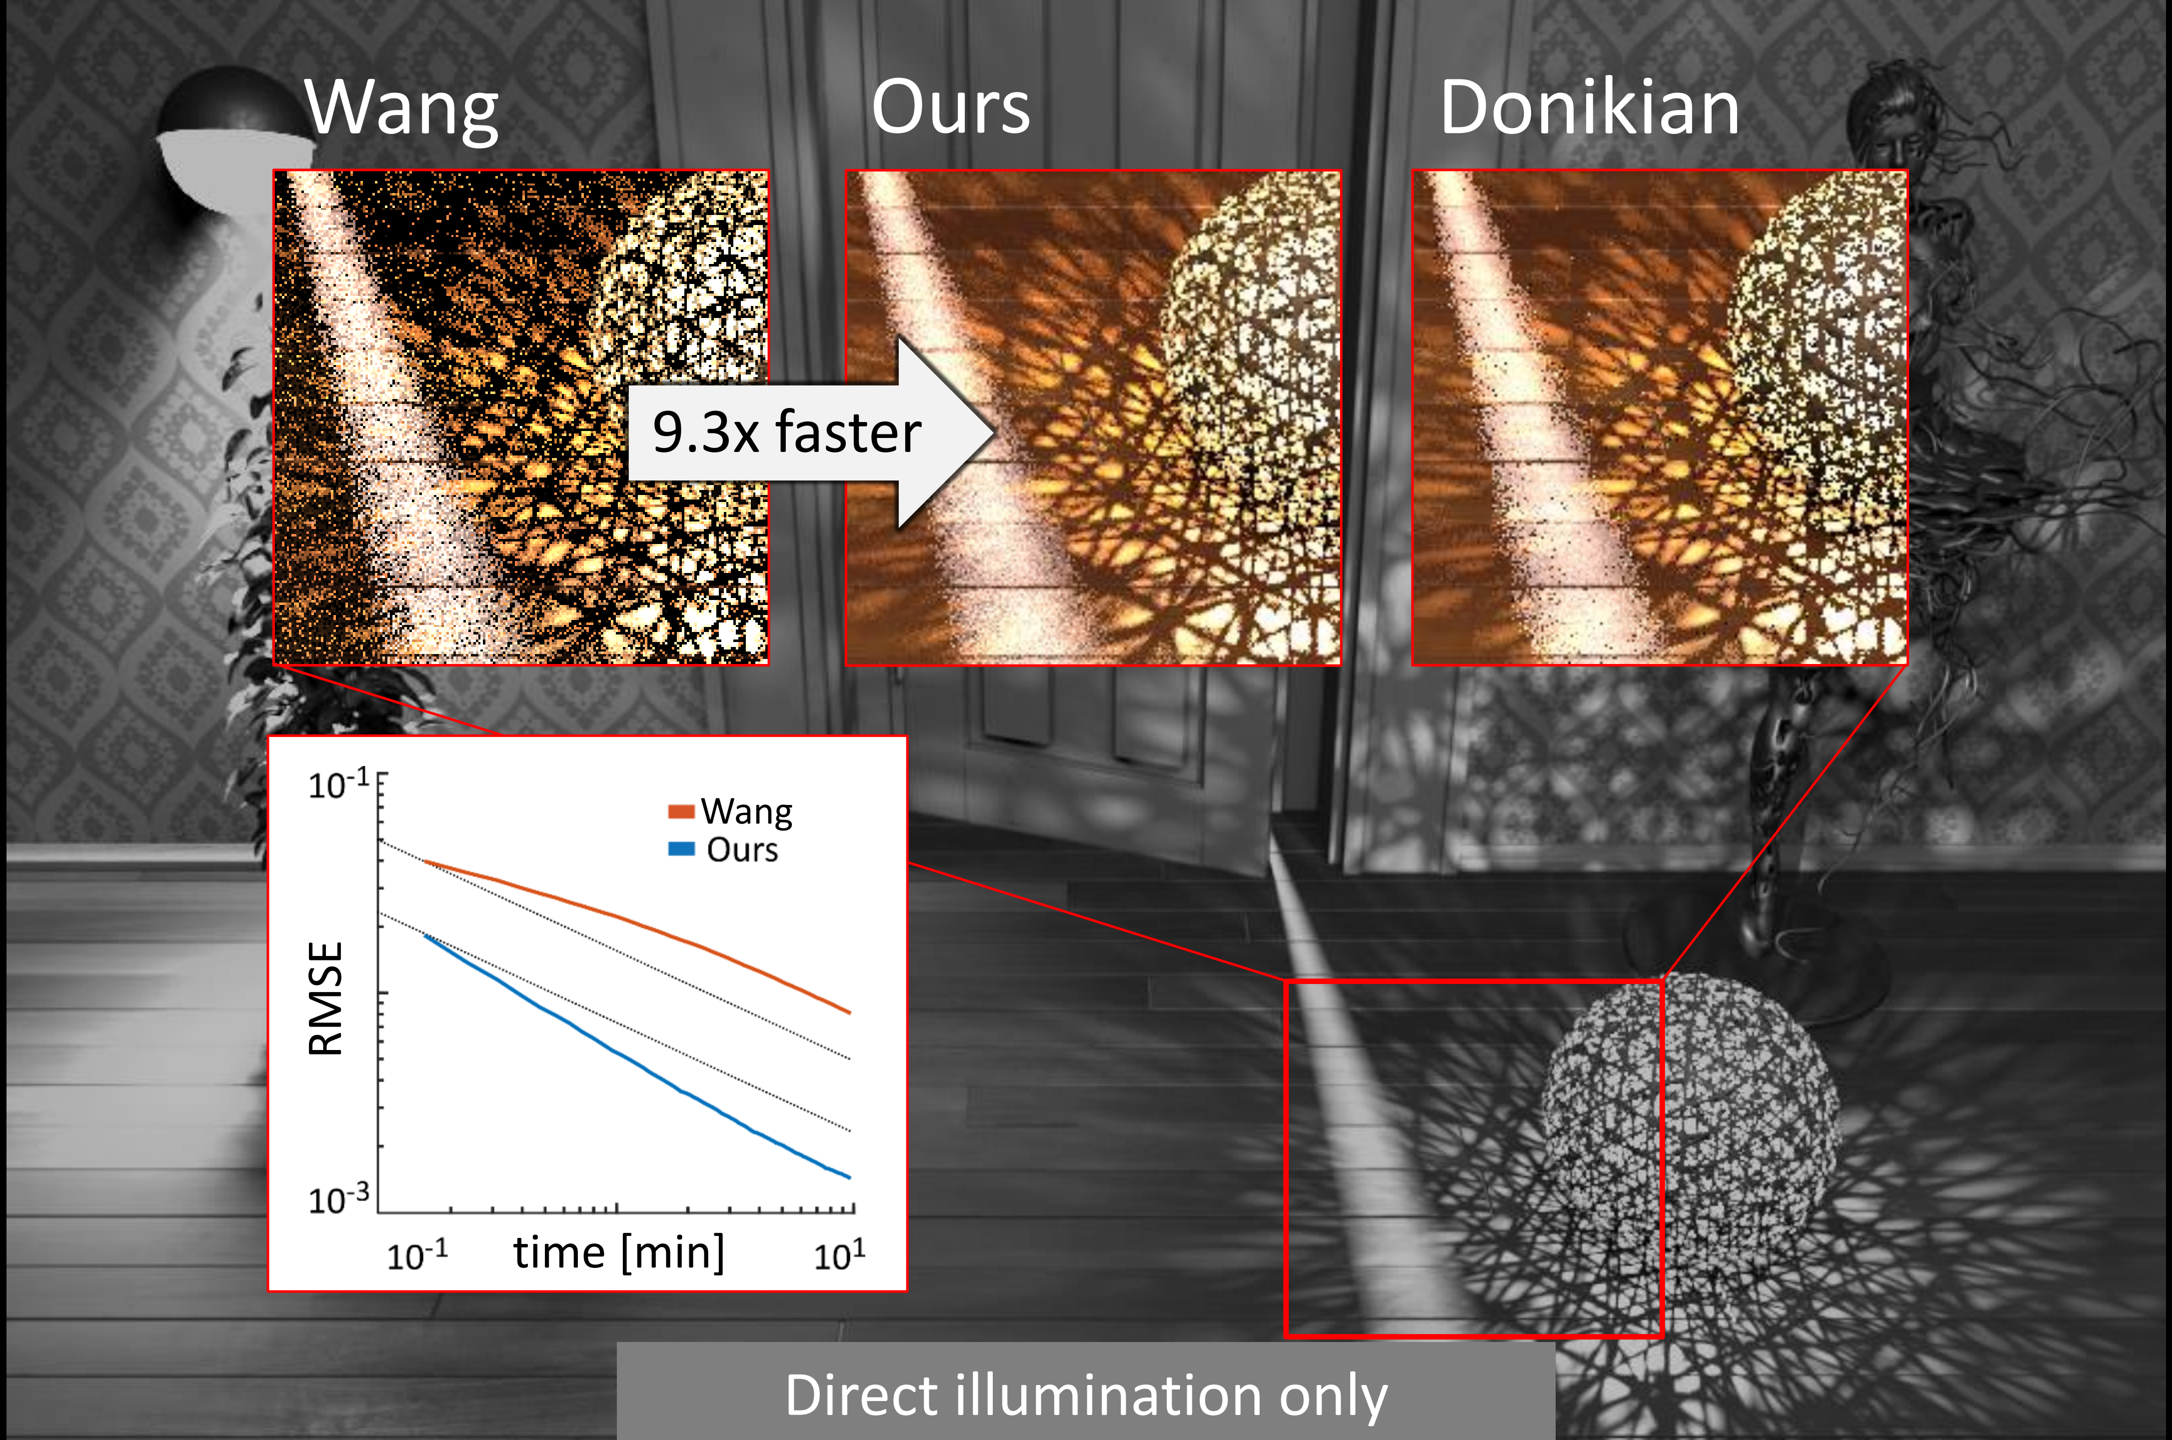

Wang
Ours
Donikian
9.3x faster
10^{-1}
10^{-3}
RMSE
Wang
Ours
10^{-1}
time [min]
10^{1}
Direct illumination only

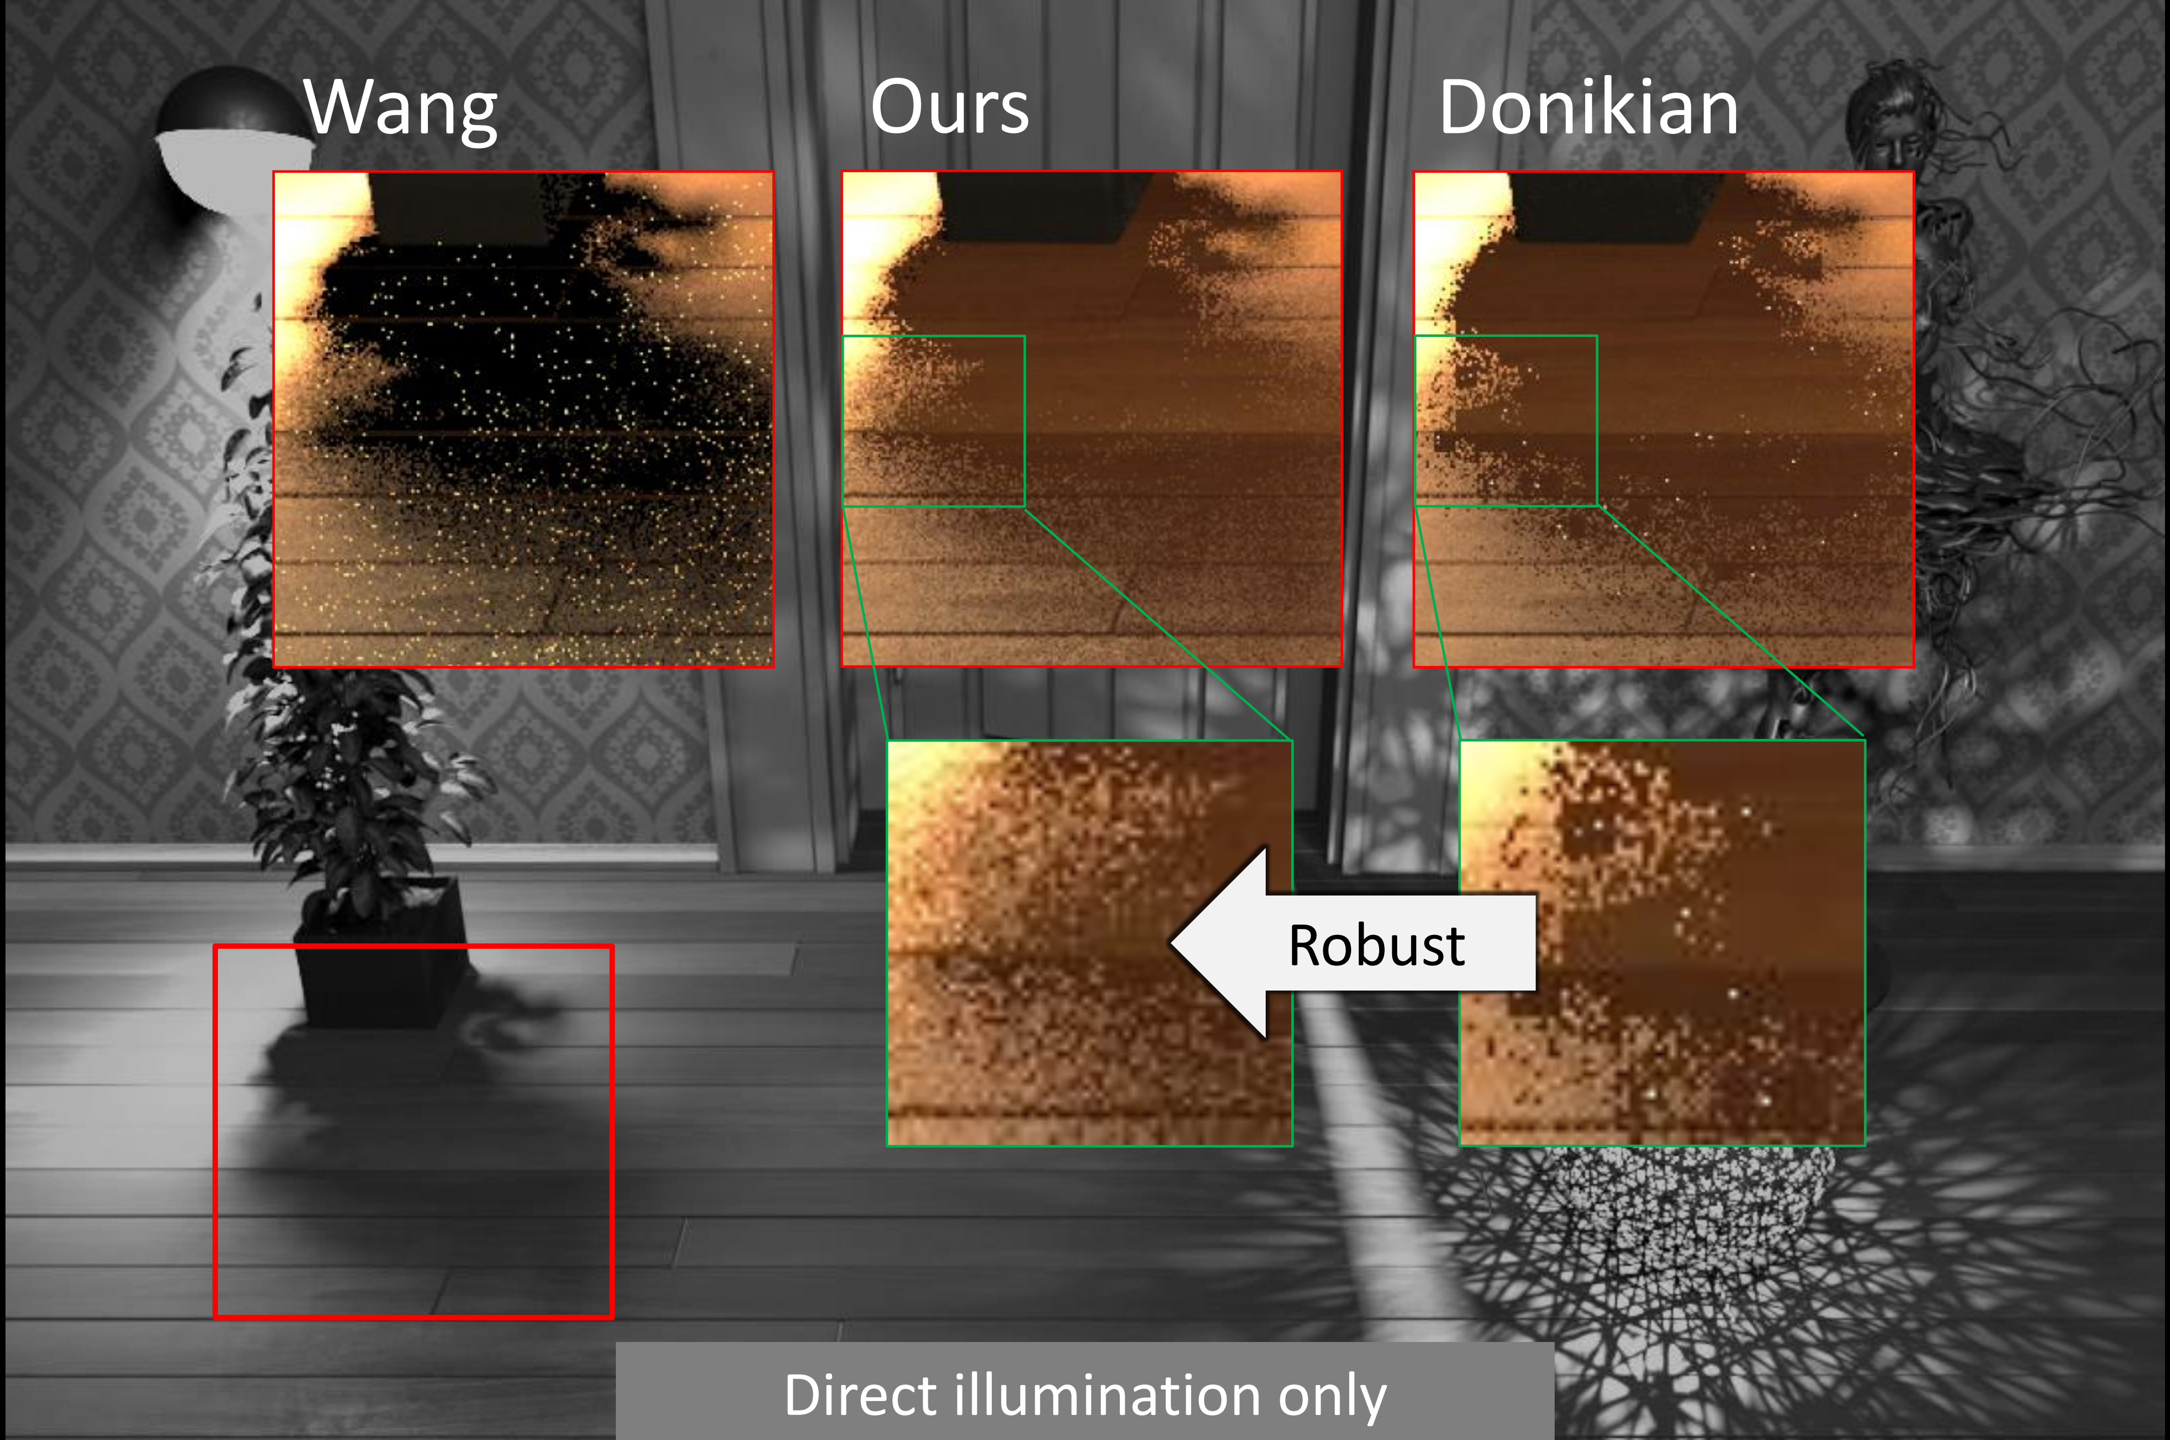
Wang
Ours
Donikian
Robust
Direct illumination only

Direct + indirect illumination

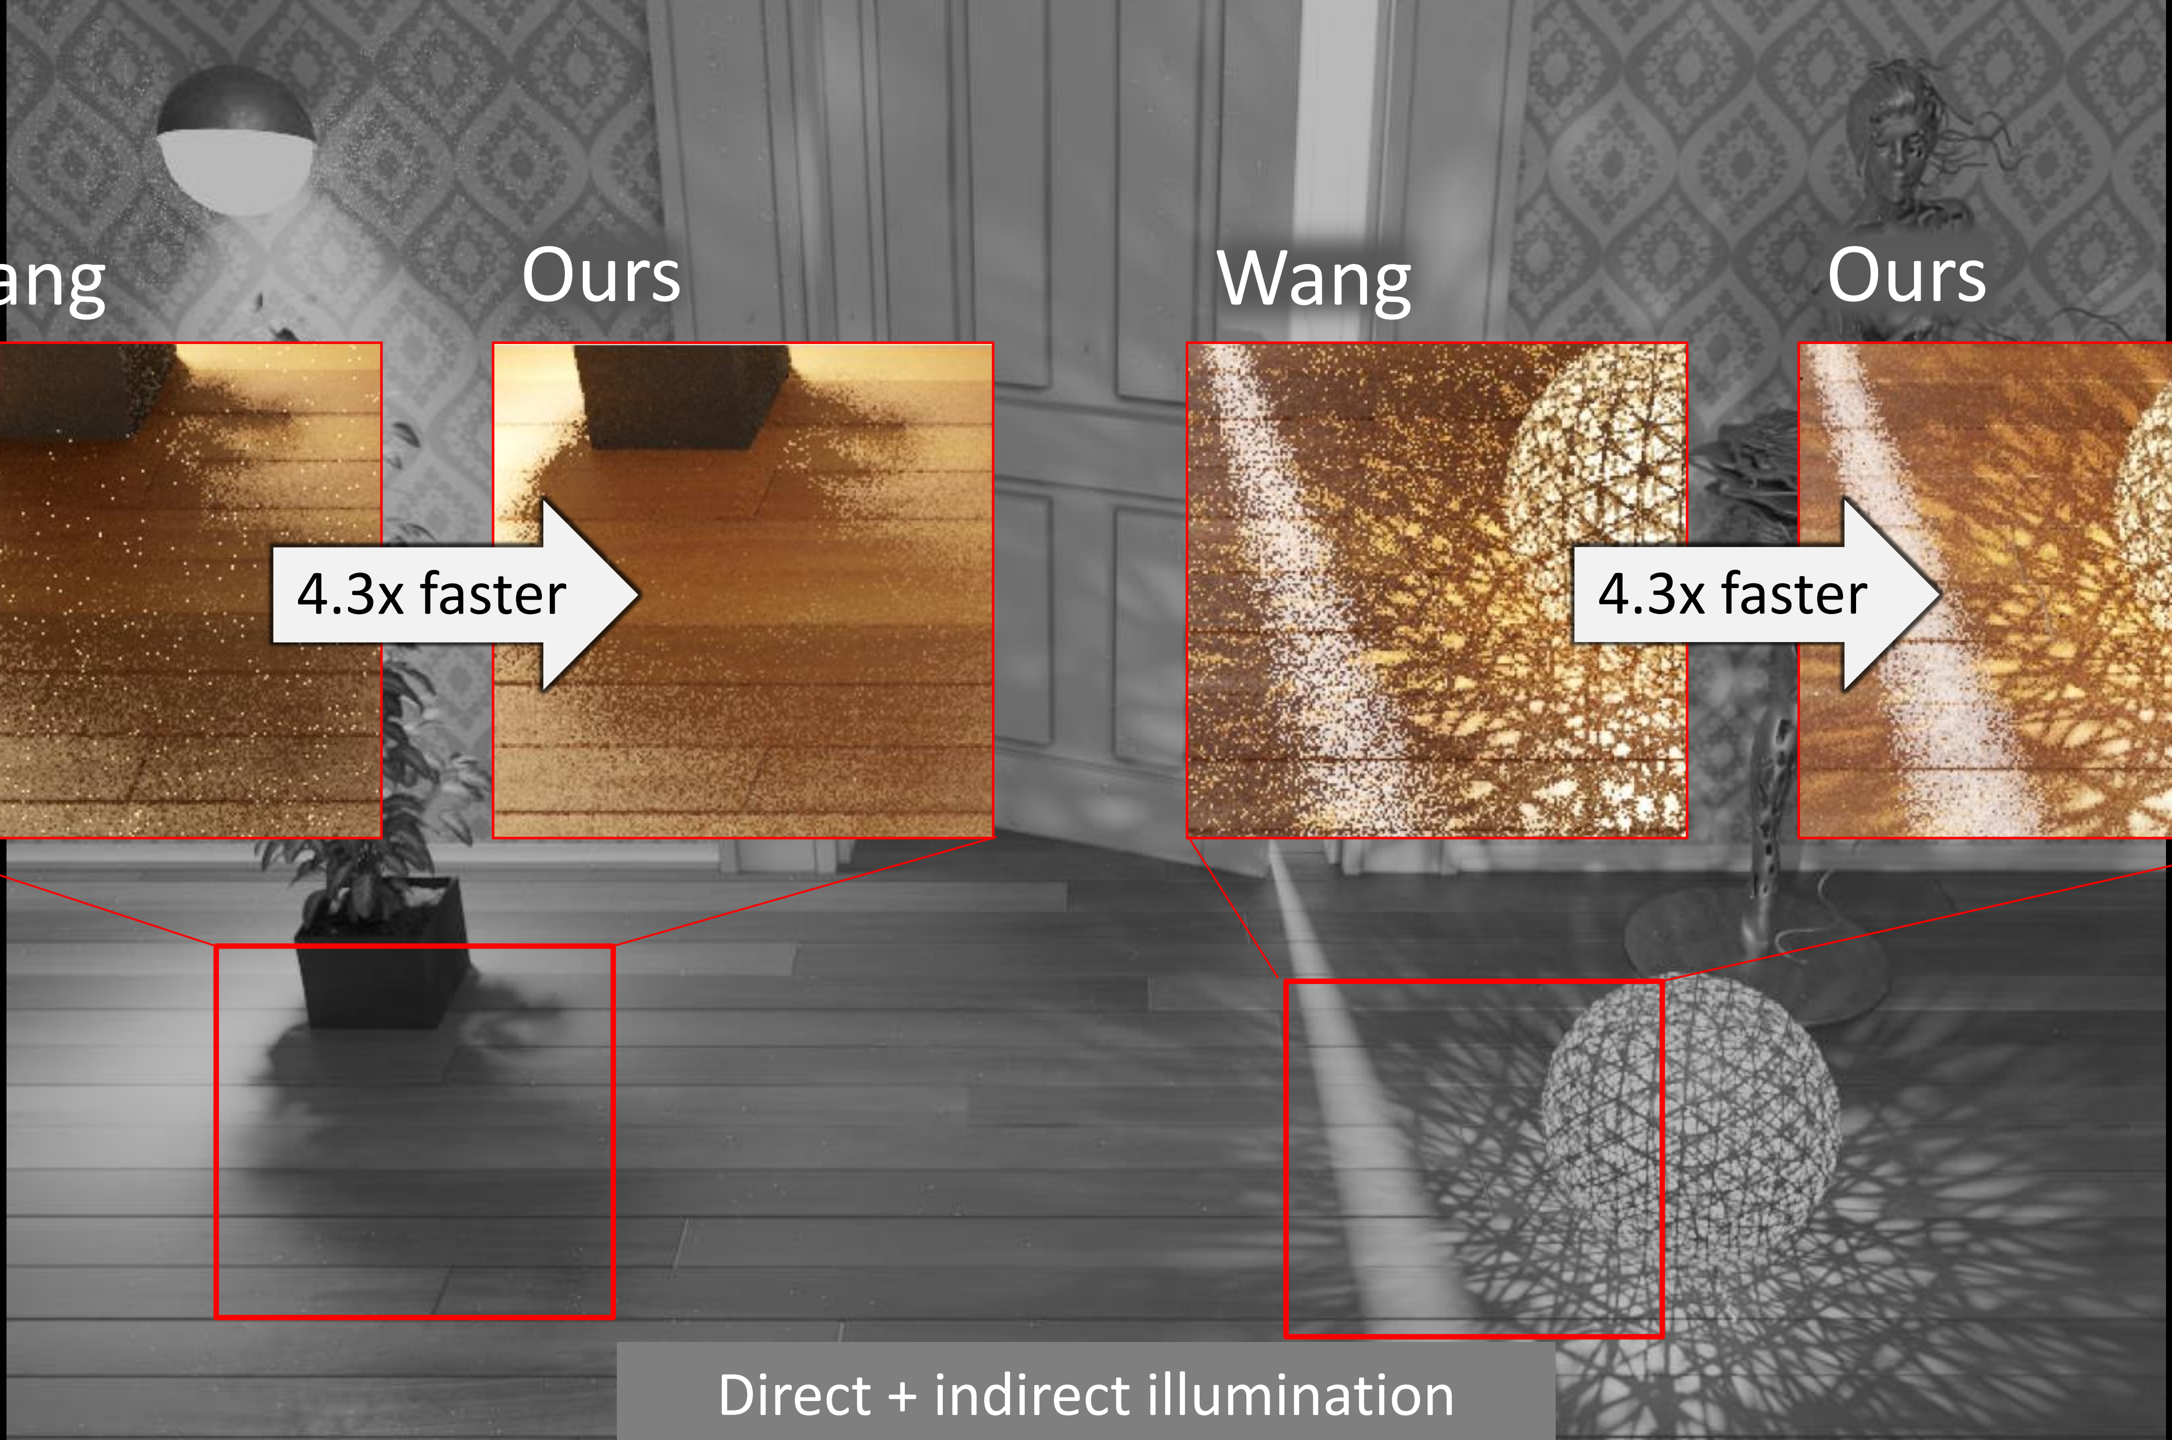
Wang
Ours
4.3x faster
Wang
Ours
4.3x faster
Direct + indirect illumination

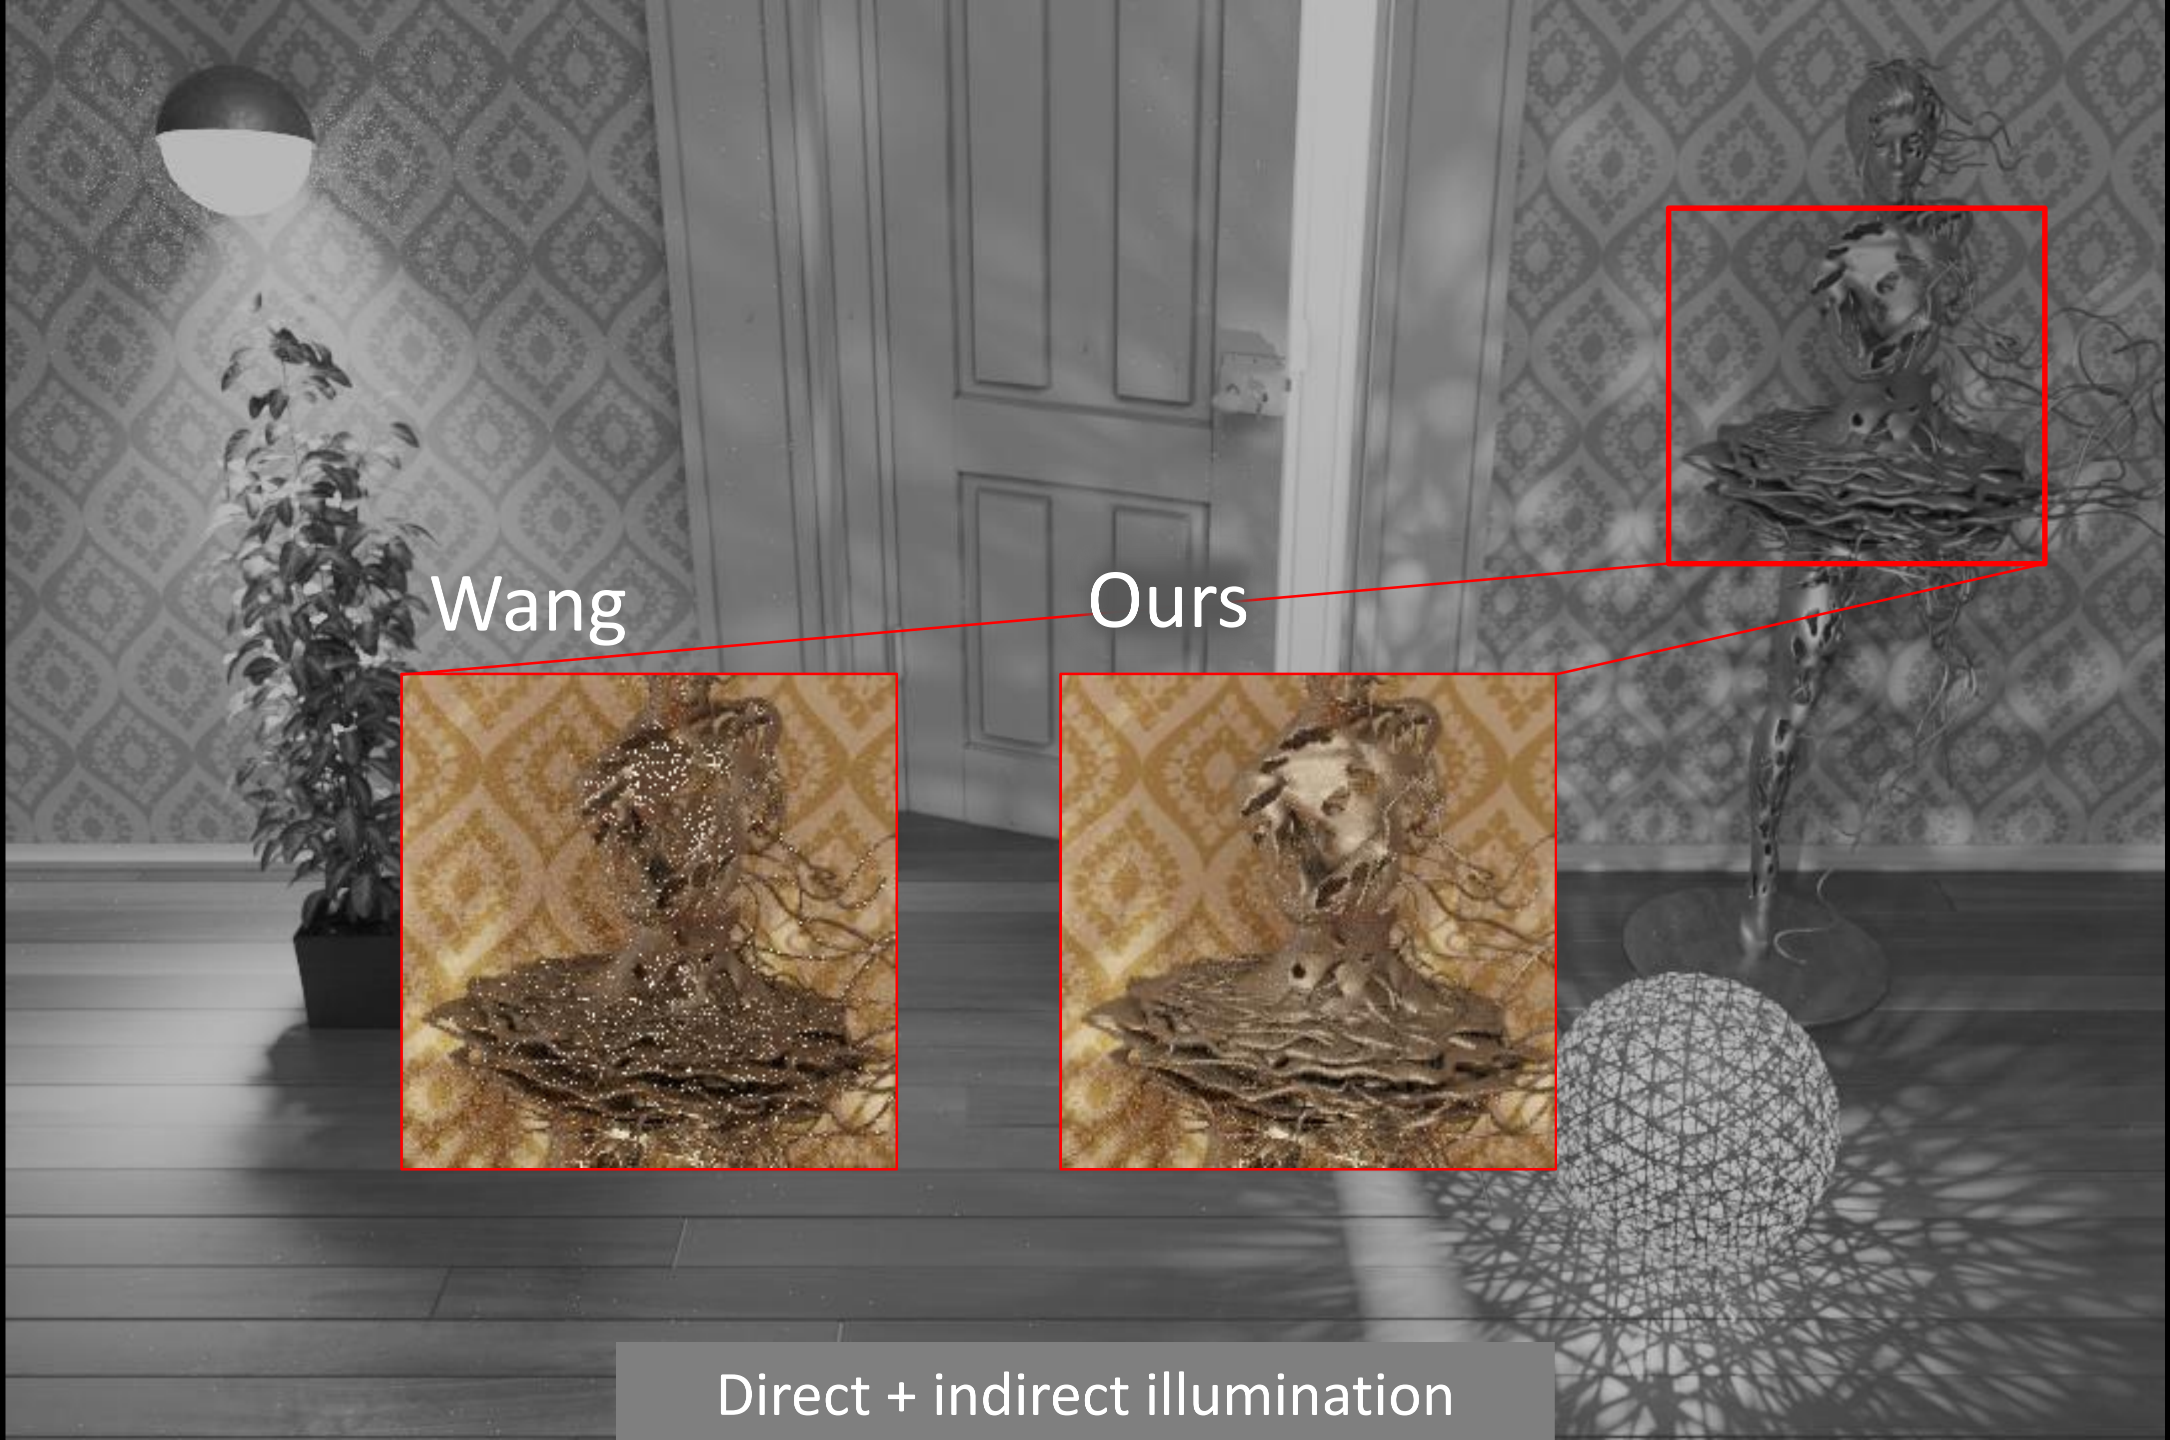

Wang
Ours
Direct + indirect illumination

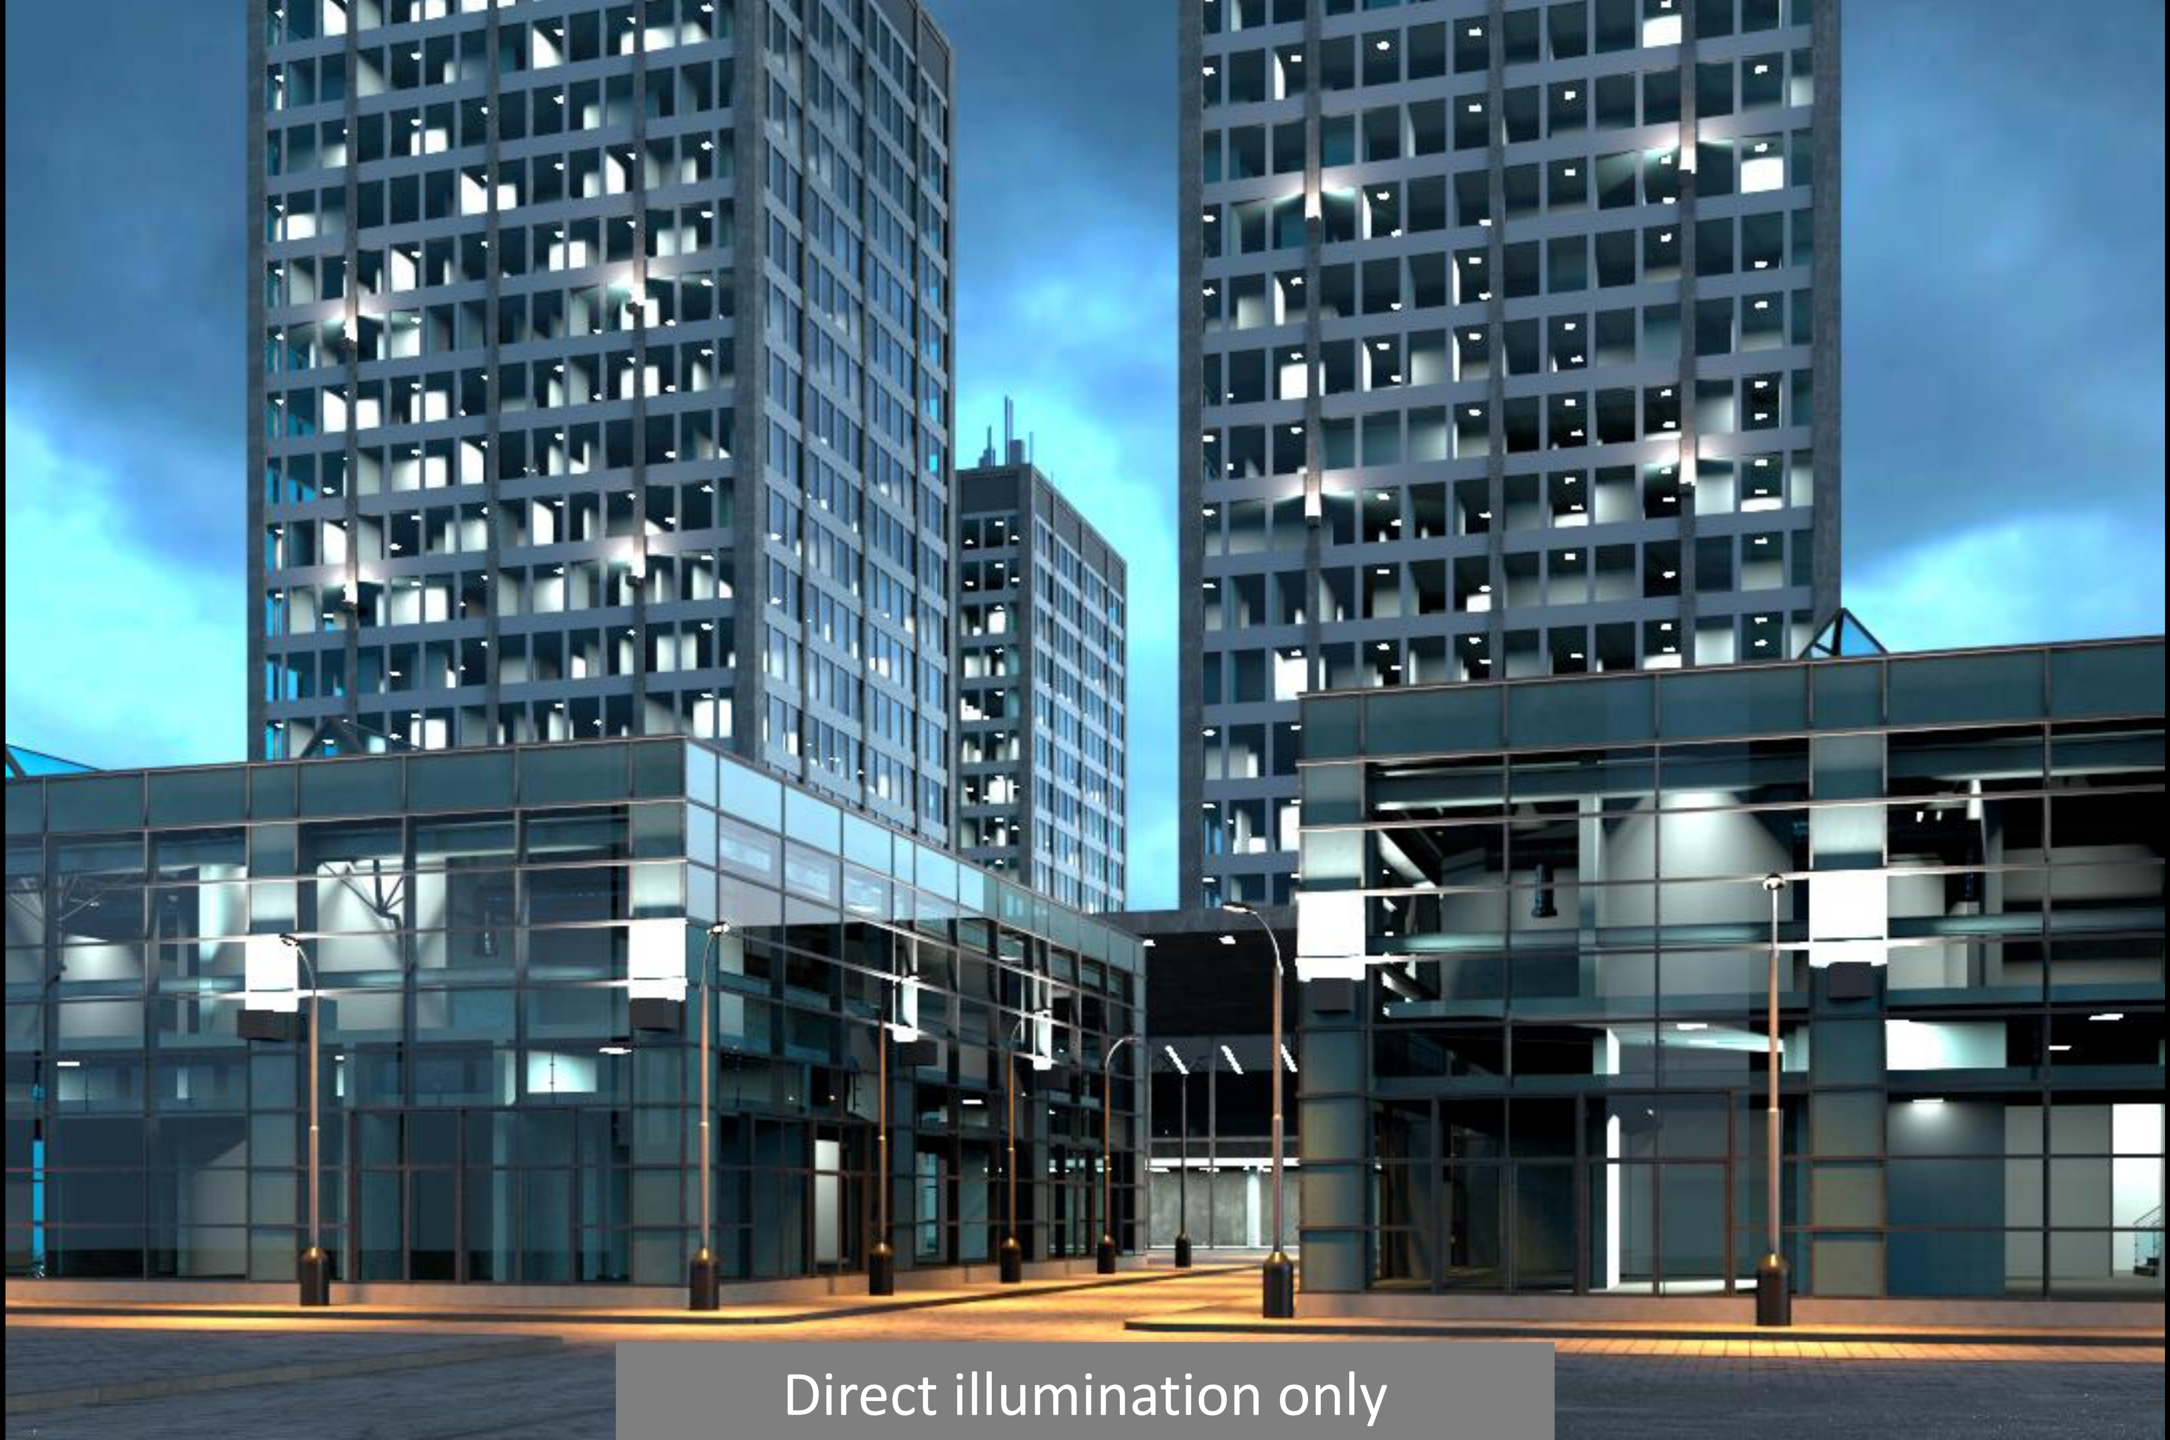

Direct illumination only

Wang

Ours (64)

No regression

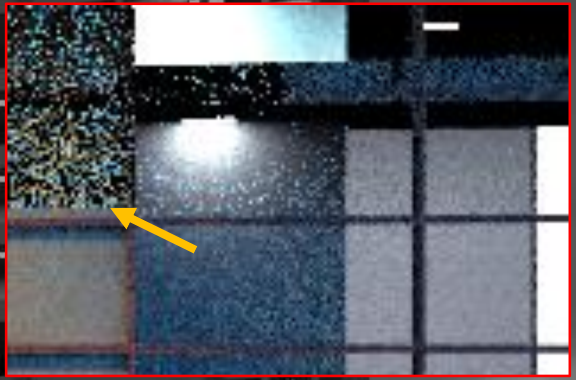

3.6x faster

$$(1 - p_0) \times N\left(\text{est.} \,\middle|\, \frac{k}{d^2}, \frac{h}{d^4}\right)$$
$$p_0 \quad \times \ \delta(\text{est.})$$

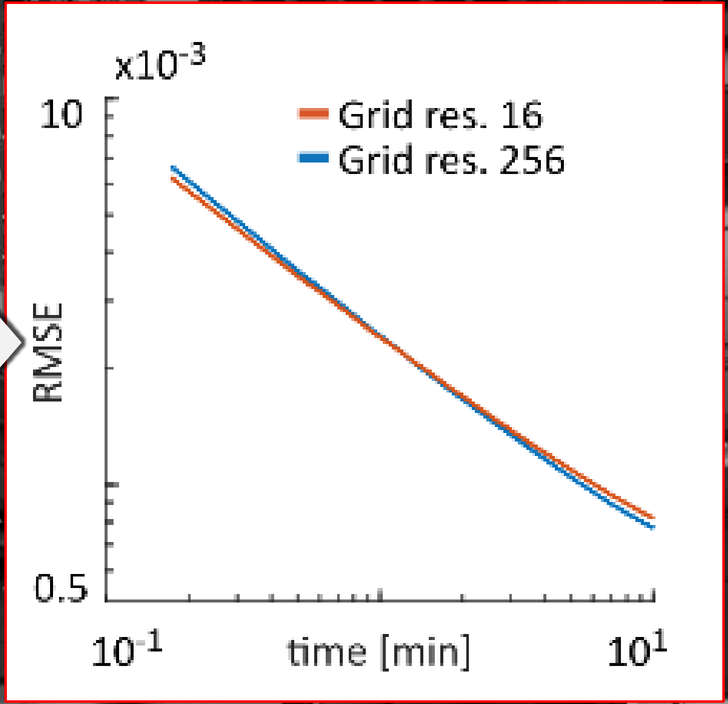

Direct illumination only

Contribution

- **Bayesian framework for robust adaptivity**
- Optimal cluster sampling
- Algorithm for direct illumination
  - Unbiased, adaptive, robust
  - Easy to integrate into a path tracer

CONCLUSION

CONCLUSION

- **Path guiding**
  - Makes complex bidirectional method unnecessary
  - Potential for wide adoption practice
- Machine learning methods = principled way to achieve **robust, online adaptive sampling**
  - Path guiding – **online learning of parametric mixture models**
  - Direct illumination sampling – **Bayesian online regression**
- **Online learning methods** compatible interactive rendering workflows & progressivity
- **Bayesian methodology** can provide the necessary **robustness**

FUTURE WORK

- Bayesian model selection
- Full Bayesian inference – Variational Bayes?
- Adaptive decision based on reinforcement learning
- Deep learning for light field reconstruction for path guiding
- Can this be that one missing piece to make MCMC methods useful in practice?

THANK YOU!

- **Acknowledgments**
  - Colleagues and students from Charles university & Corona
  - Funding: Czech Science Foundation (16-18964S), Charles University Grant Agency project GAUK 1172416, by the grant SVV-2017-260452

- While you may think that rendering is science, remember that first and foremost, **rendering is magic**.