Scalable Virtual Ray Lights Rendering for Participating Media

Nicolas Vibert Adrien Gruson Heine Stokholm Troels Mortensen
Wojciech Jarosz Toshiya Hachisuka Derek Nowrouzezahrai

McGill University Luxion VIA University College Dartmouth College The University of Tokyo

MOTIVATION: Surface and Volume interaction

MOTIVATION: Volume interaction only

VOLUMETRIC RENDERING

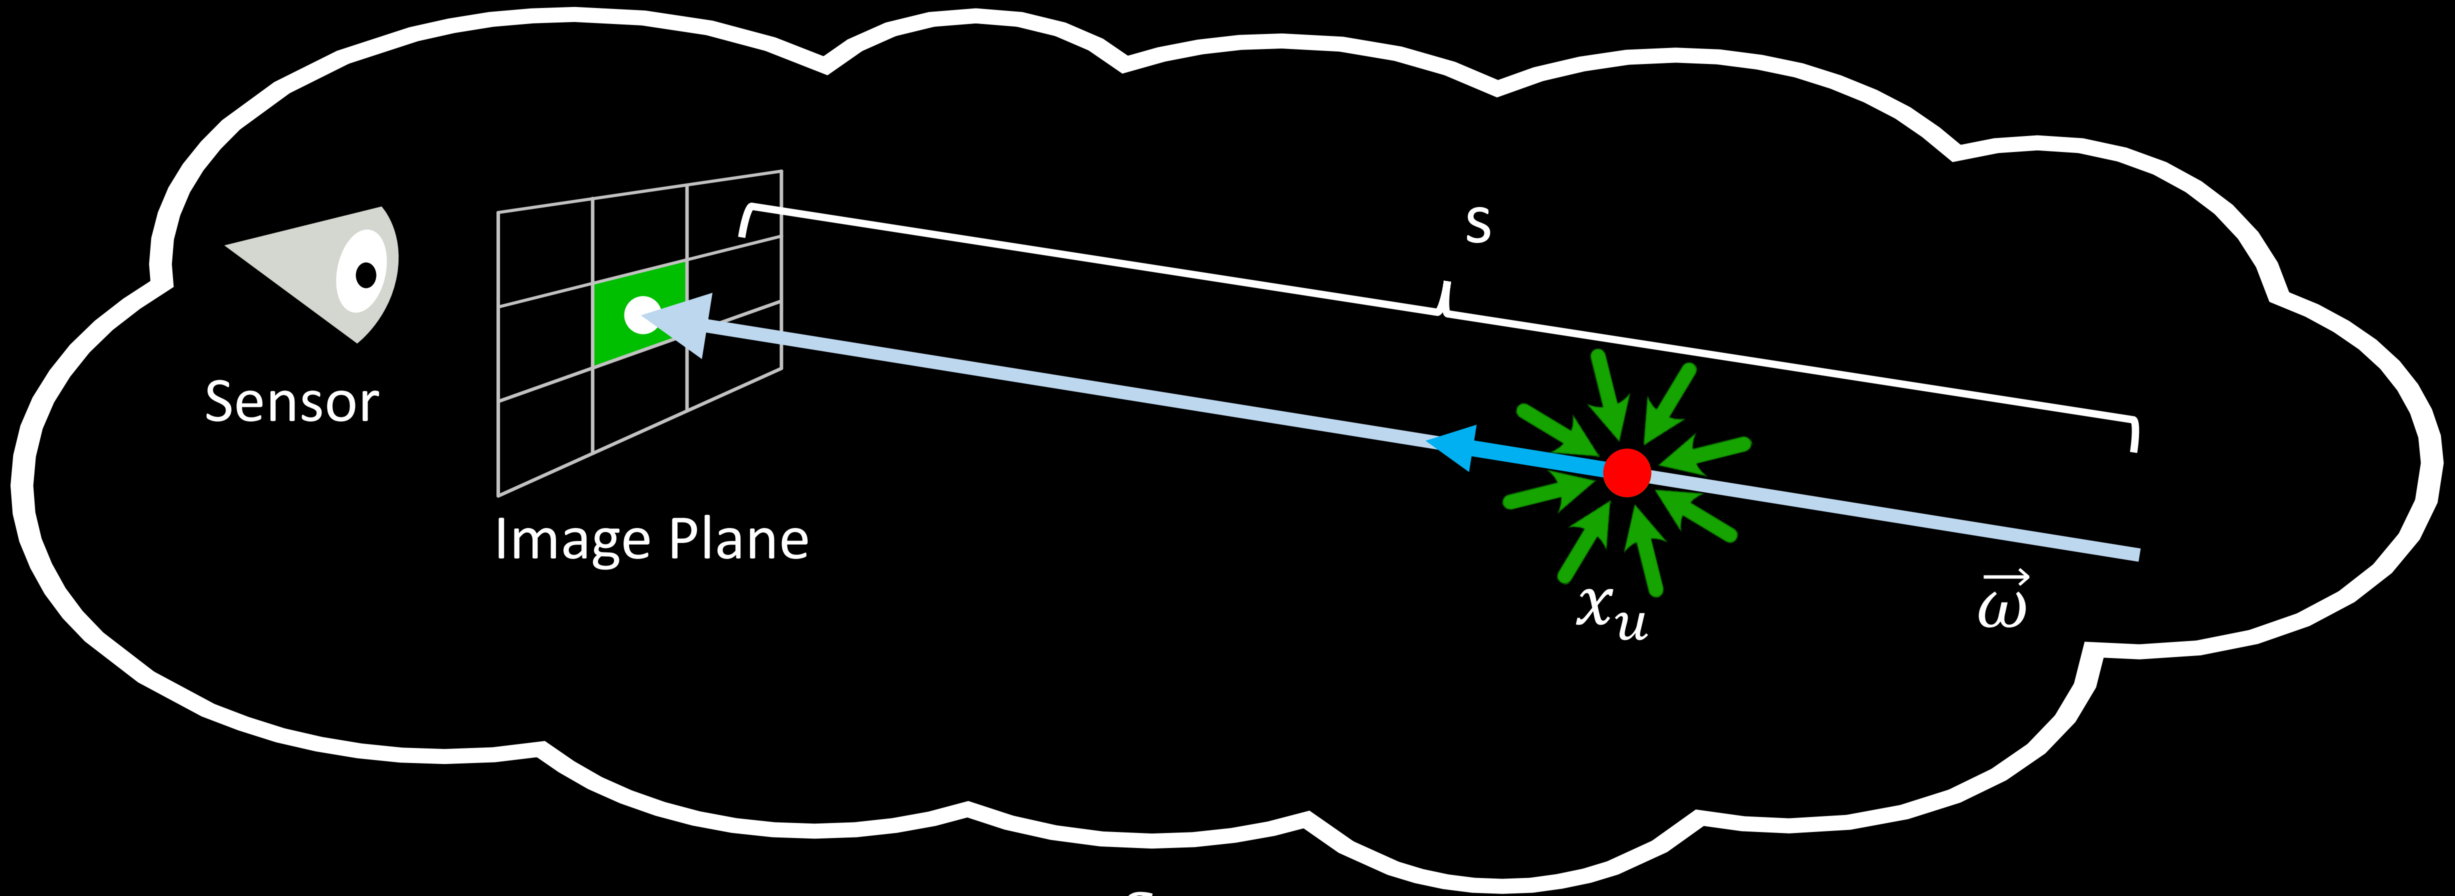

$$L_m(x, \vec{\omega}) = \int_0^s T_r(u) L_i(x_u, \vec{\omega}) du$$

VOLUMETRIC RENDERING: Rendering techniques

- Path tracing / Bidirectional path tracing

VOLUMETRIC RENDERING: Many techniques

- Path tracing / Bidirectional path tracing
- Density estimation:
  - Volumetric Photon Mapping
  - Photon Beam
  - Photon Planes
  - “Higher-order geometric primitives”

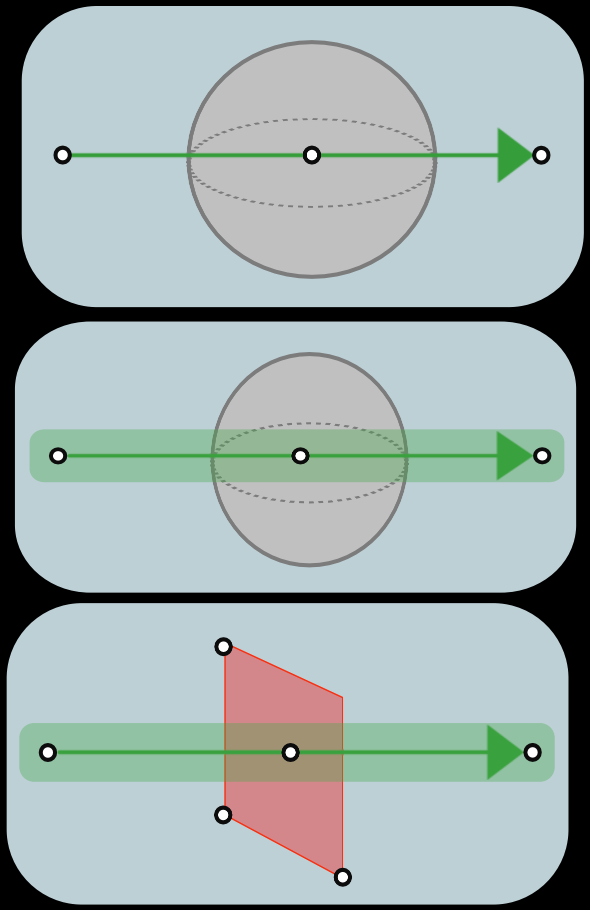

VOLUMETRIC RENDERING: Many techniques

- Path tracing / Bidirectional path tracing
- Density estimation:
  - Volumetric Photon Mapping
  - Photon Beam
  - Photon Planes
  - “Higher-order geometric primitives”
- Many lights:
  - Virtual point lights
  - Virtual spherical lights
  - Virtual ray lights
  - “Higher-order geometric primitives”

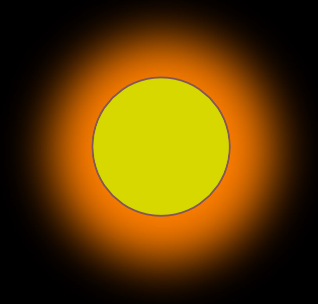

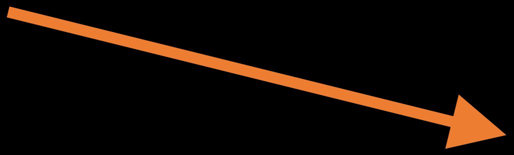

MANY LIGHTS

- Many-light techniques have been introduced in "instant radiosity" [Keller et al. 1997]
- Indirect illumination as a sum of direct illumination of virtual lights

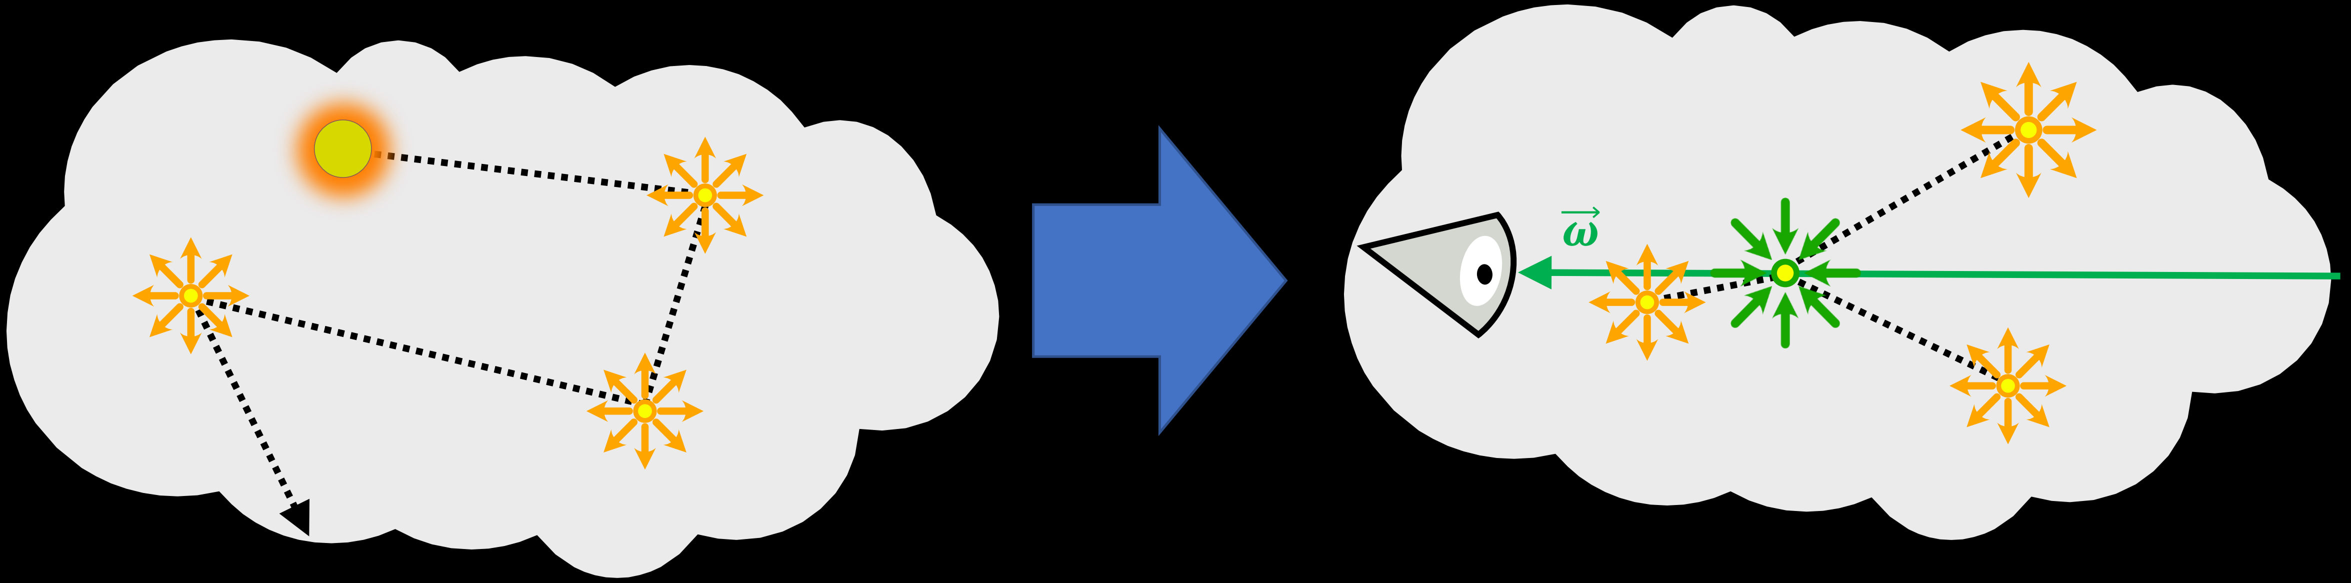

Virtual point light contribution

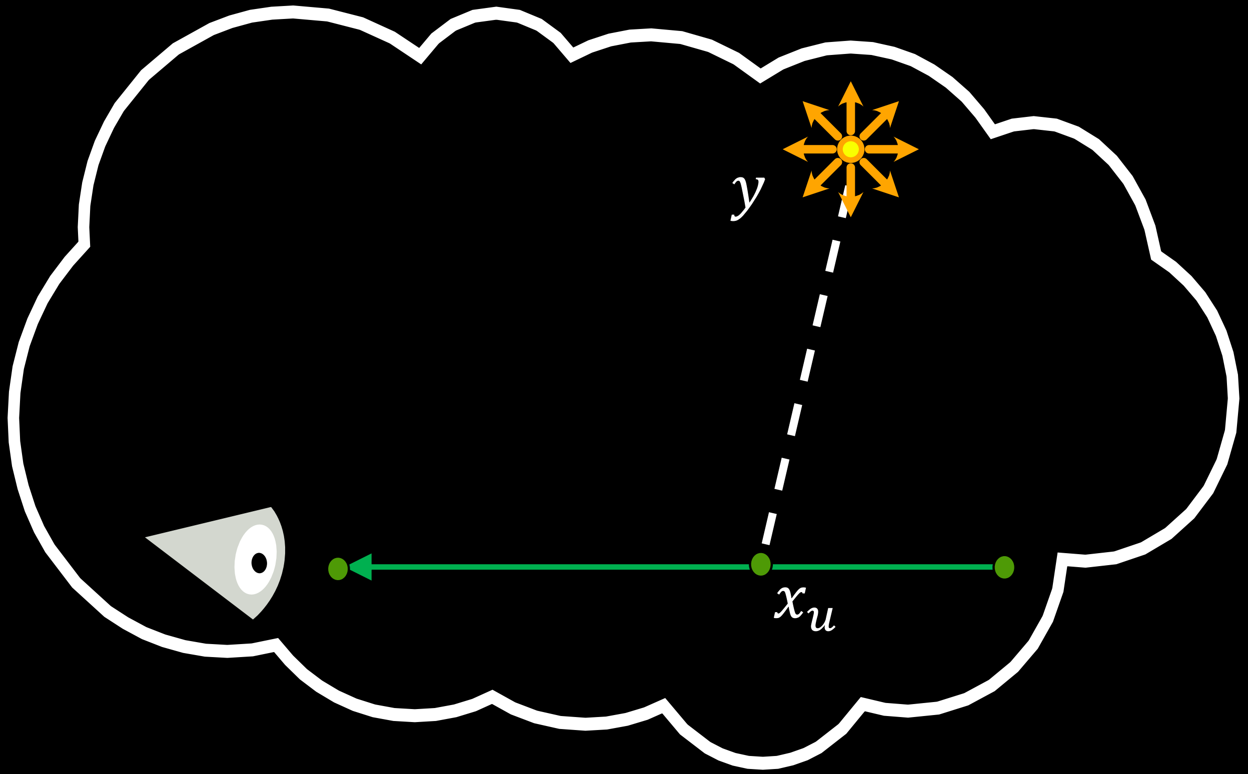

$$L_m^{VPL} = \frac{VPL(y, \mathrm{x}_u)}{|y - x_u|^2 p(\mathrm{x}_u)}$$

Virtual point light contribution

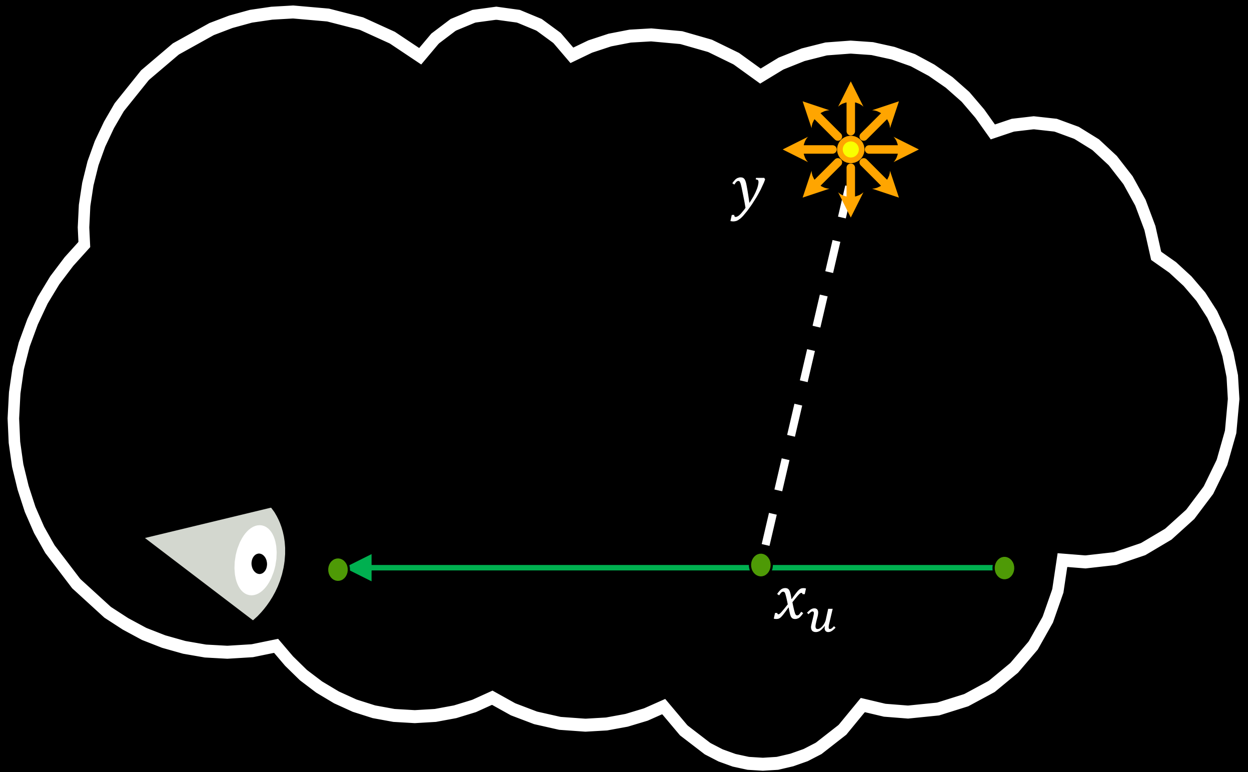

$$L_m^{VPL} = \frac{VPL(y, \mathrm{x}_u)}{|y - x_u|^2 p(\mathrm{x}_u)}$$

$$\propto \boxed{\frac{1}{|y - x_u|^2}}$$

MANY LIGHTS

- VPL vs. short VRL

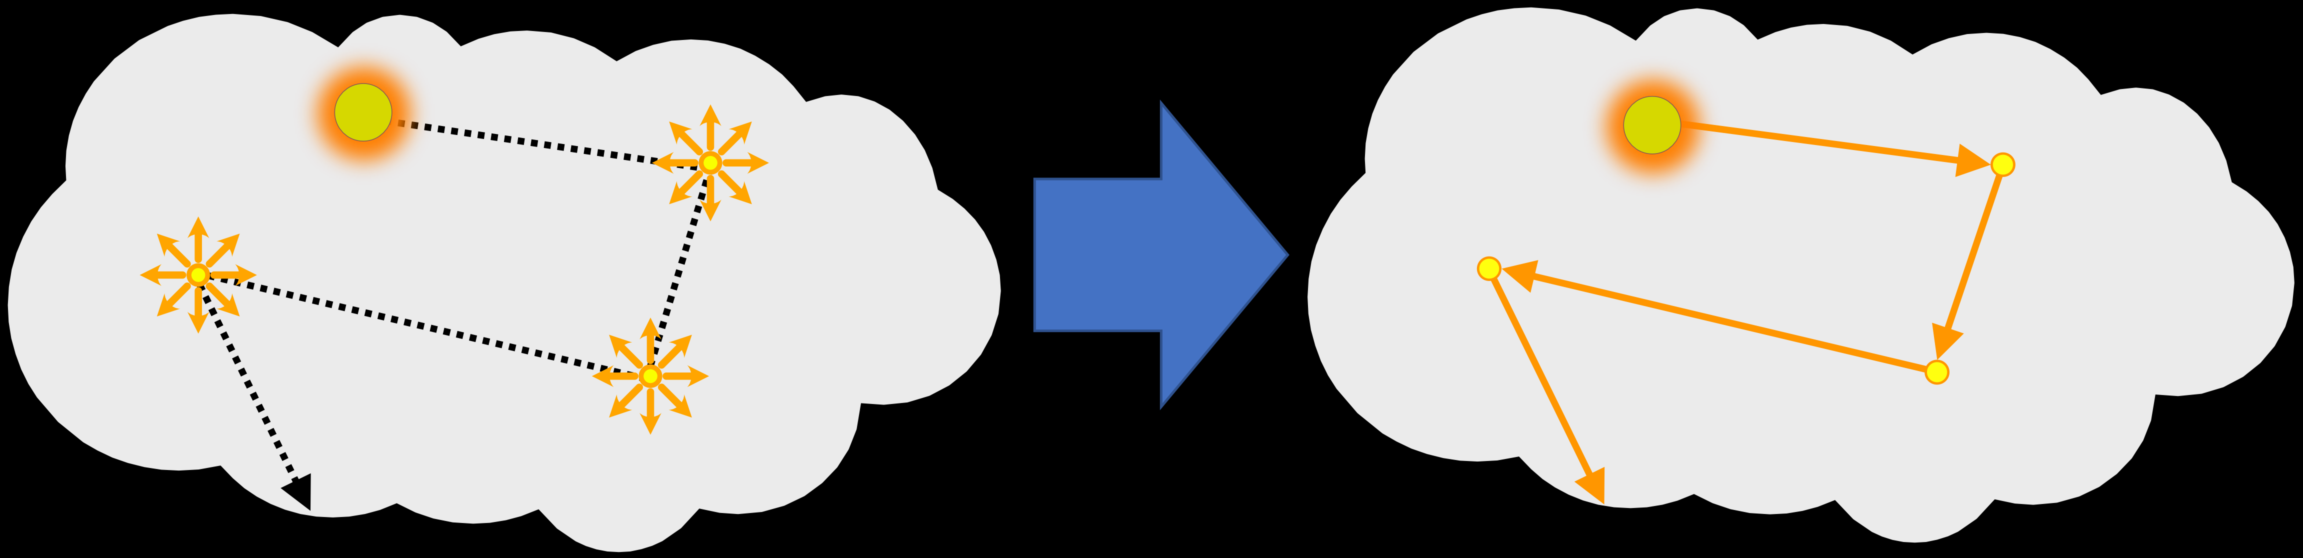

MANY LIGHTS: VRL

Virtual ray lights contribution

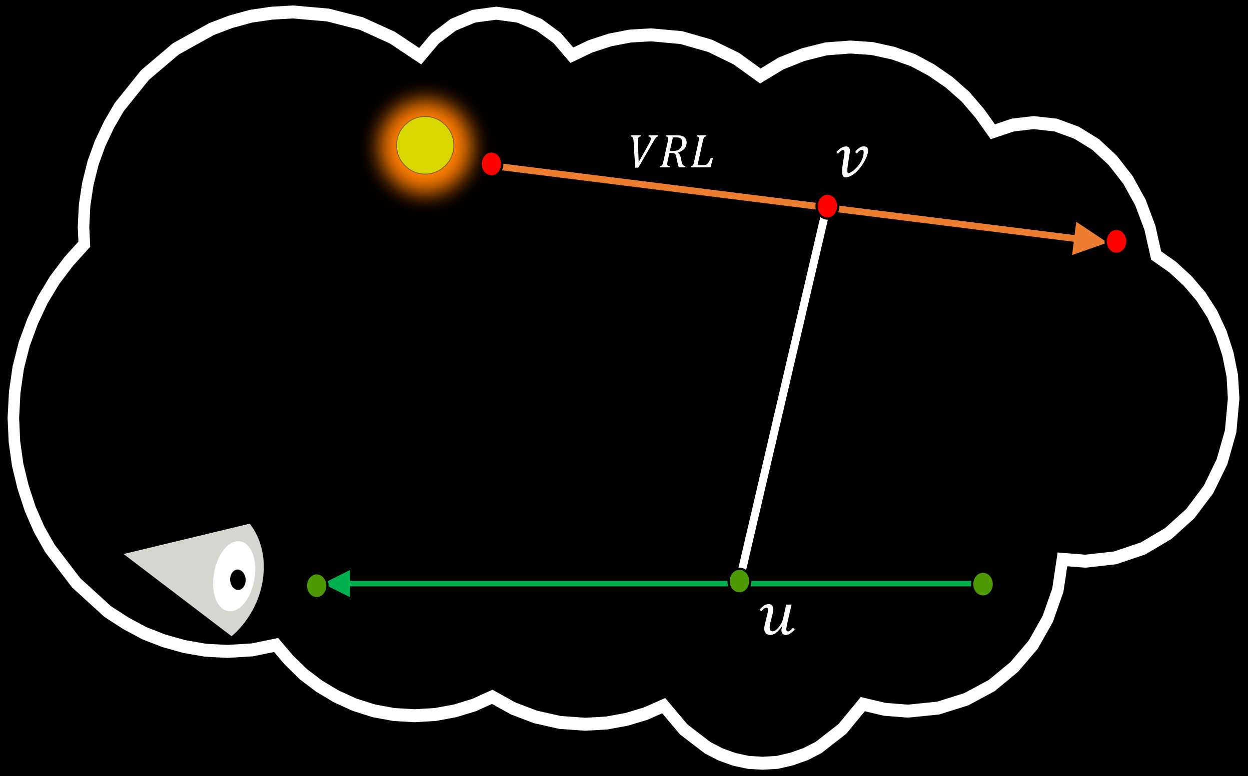

$$L_m^{VRL} = \int_0^s \int_0^t \frac{VRL(u,v)}{w(u,v)^2}\, dv\, du$$

MANY LIGHTS: VRL

Virtual ray lights contribution

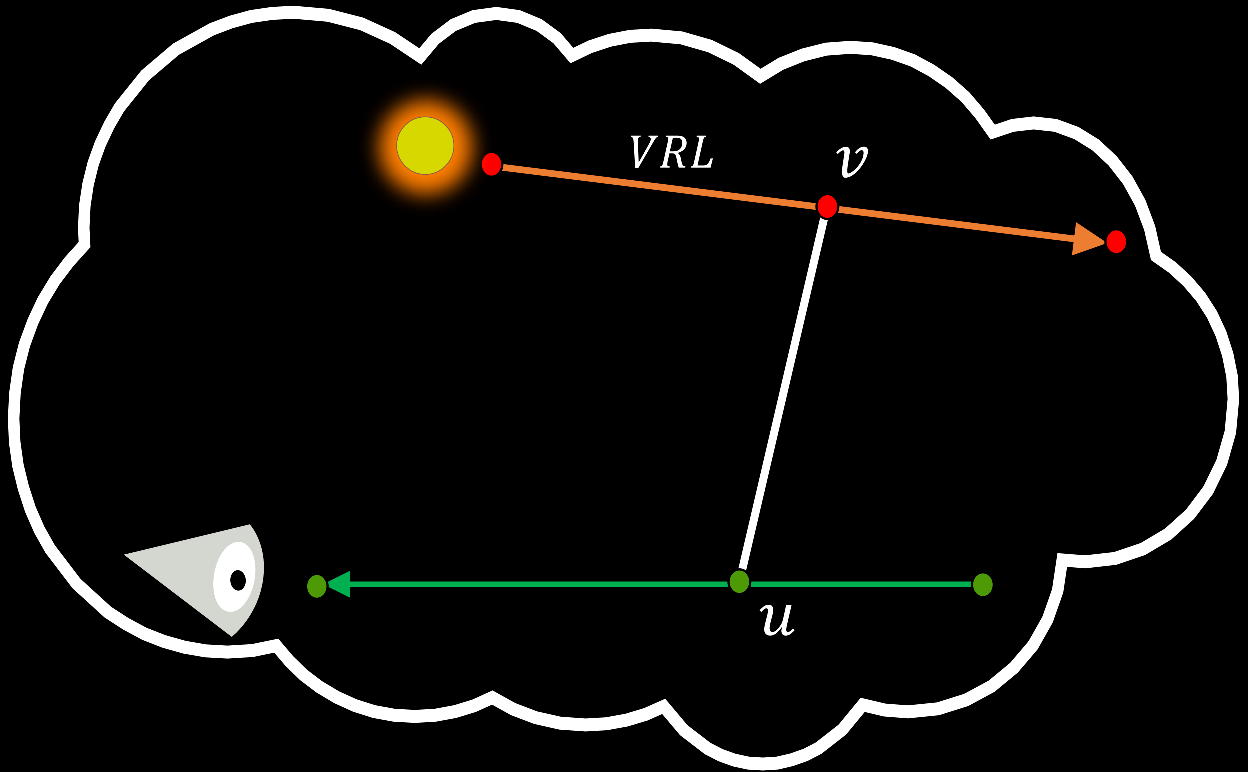

$$L_m^{VRL} = \frac{VRL(u,v)}{w(u,v)^2 p(u,v)}$$

$$p(u,v) \propto w(u,v)^{-2}$$

MANY LIGHTS: VRL vs. VPL

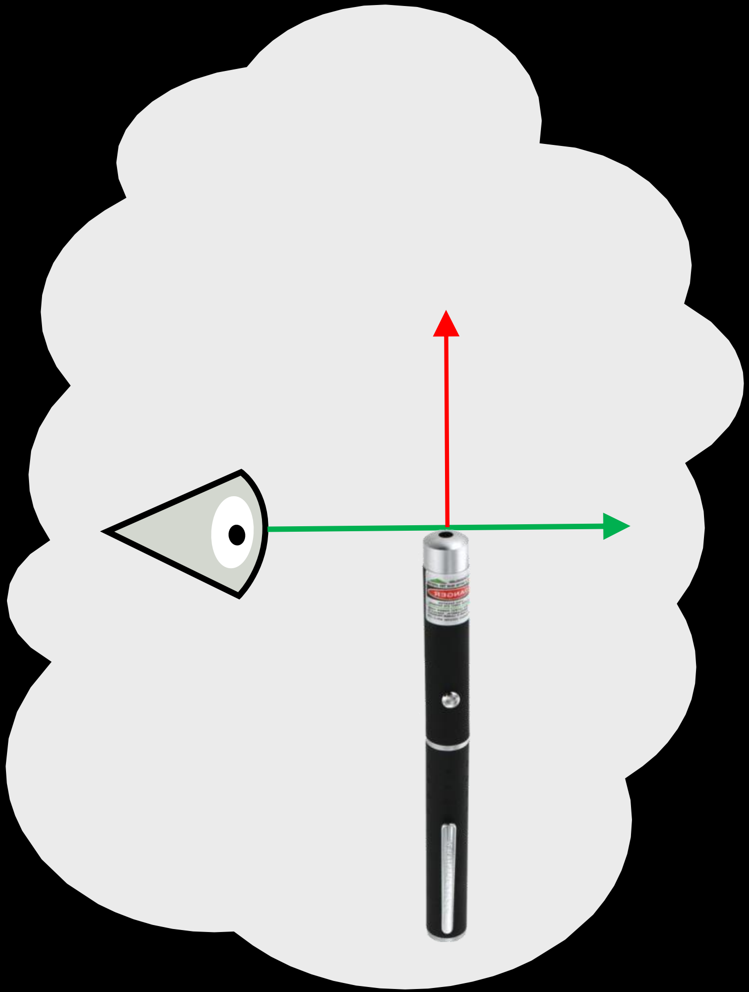

VRL

VPL

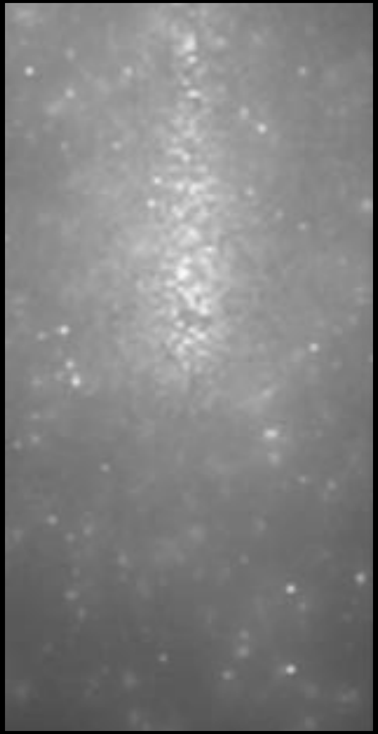

Equal rendering time

MANY LIGHTS

Realistic rendering :

- Unmanageable amount of virtual lights
- Cost linear with lights

MANY LIGHTS

Realistic rendering:
- Unmanageable amount of virtual lights
- Cost linear with lights

Aim:
- Sub-linear cost
- Scalable methods
  [Walter et al. 2005][Walter et al. 2006][Walter et al. 2012]
  [Hasan et al. 2007][Ou et al. 2011][Bus et al. 2015]

SCALABILITY

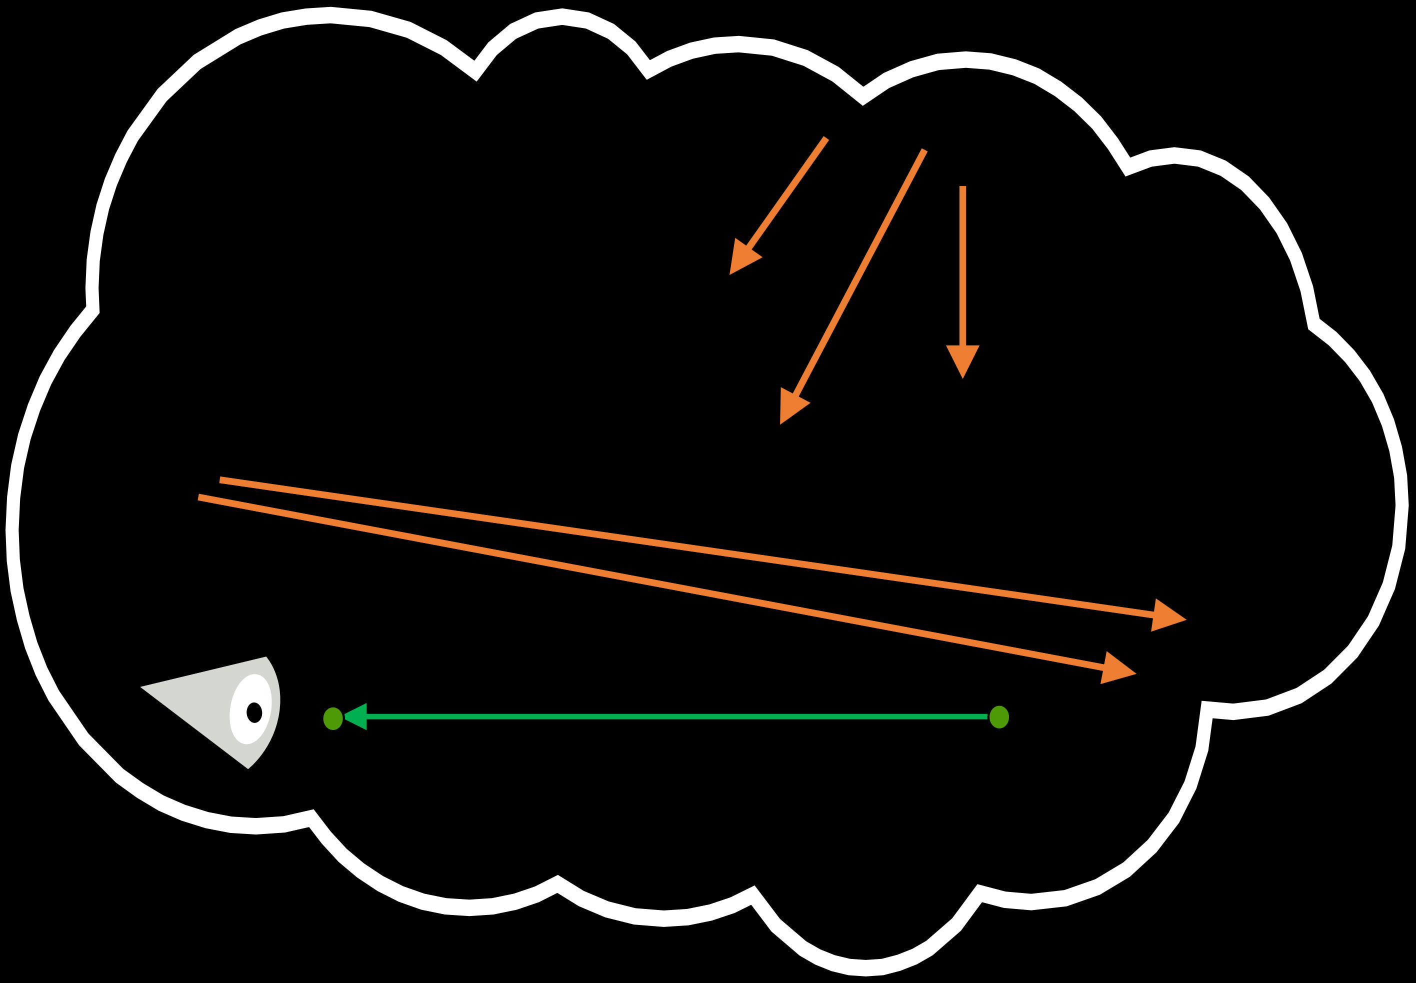

SCALABILITY

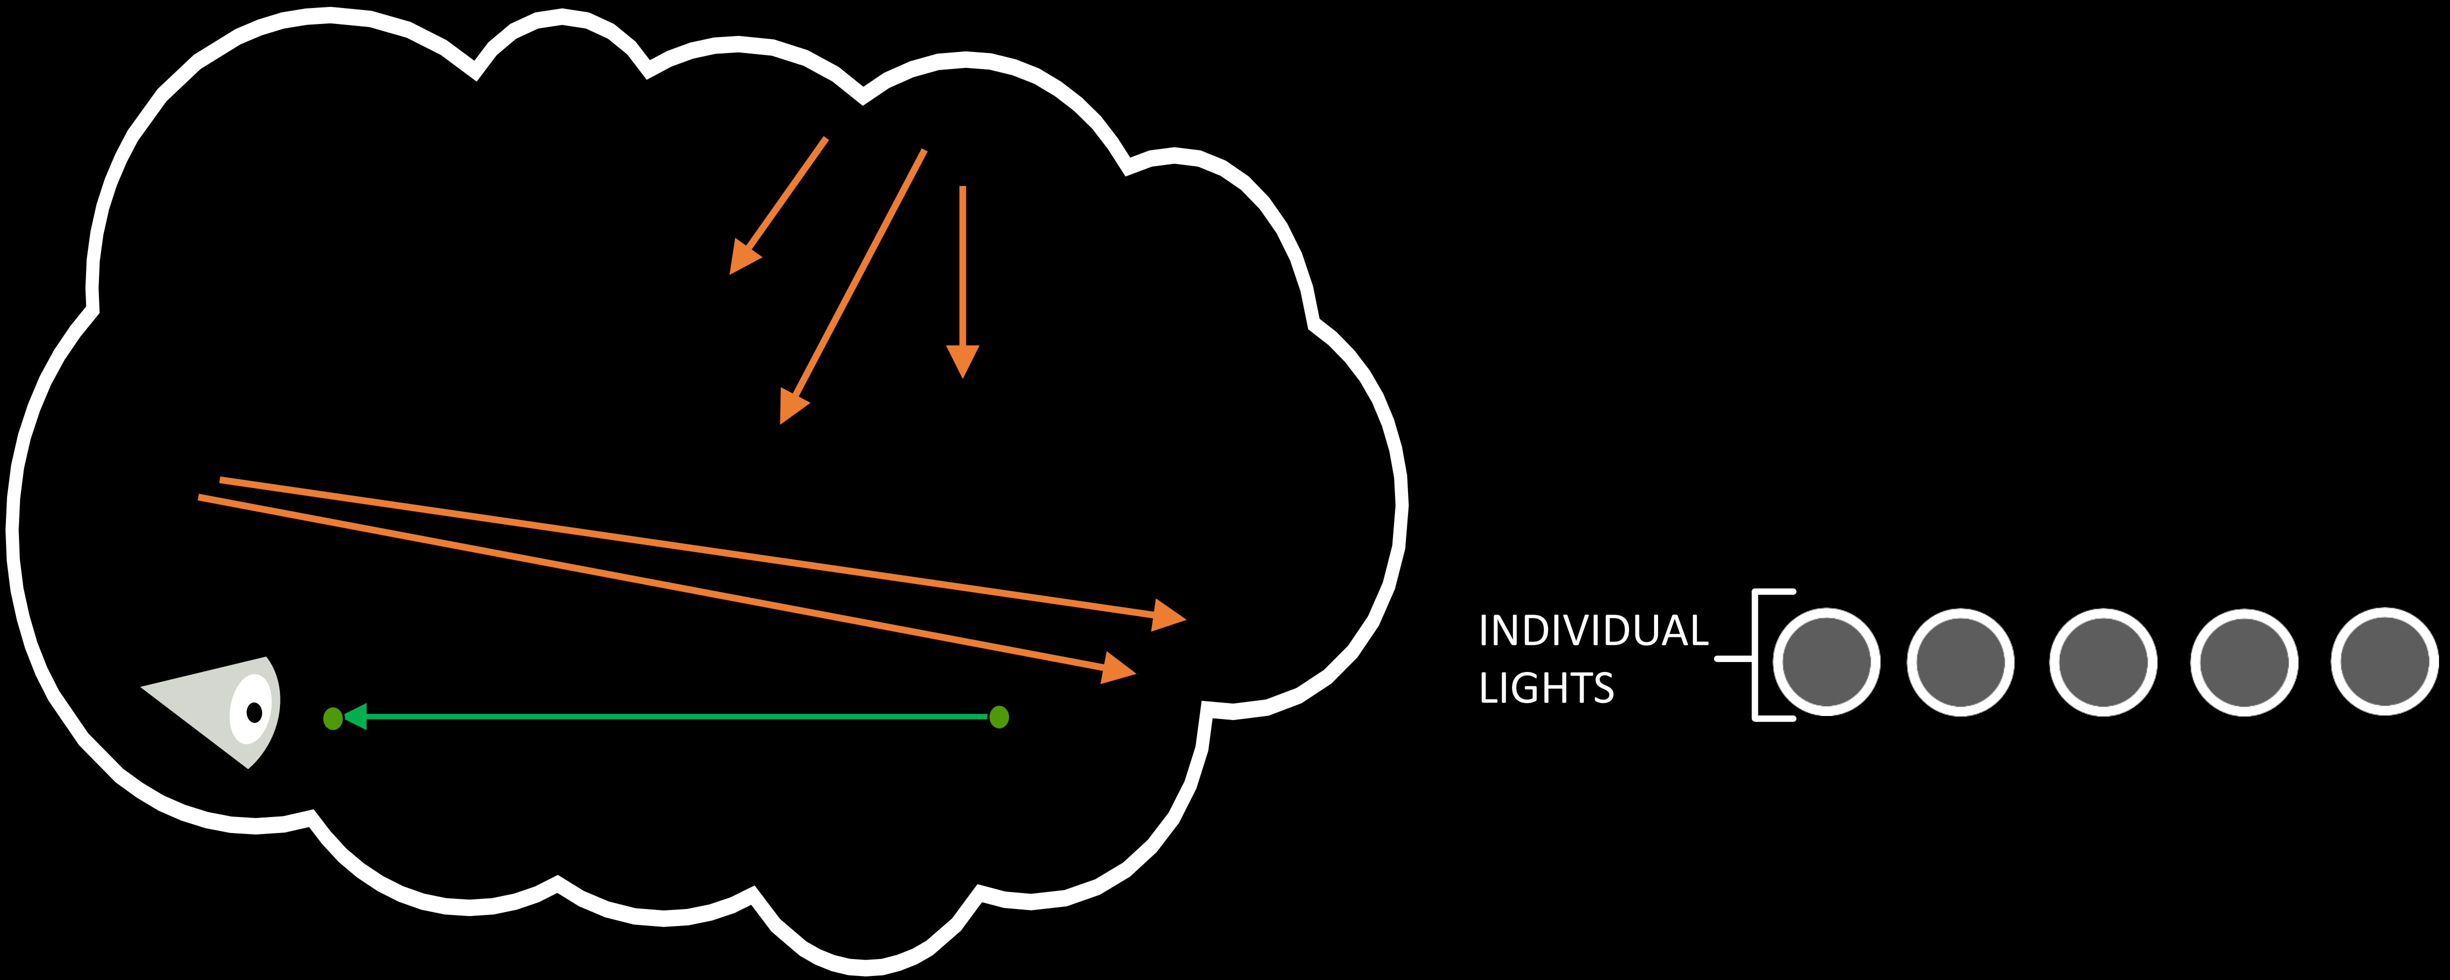

SCALABILITY

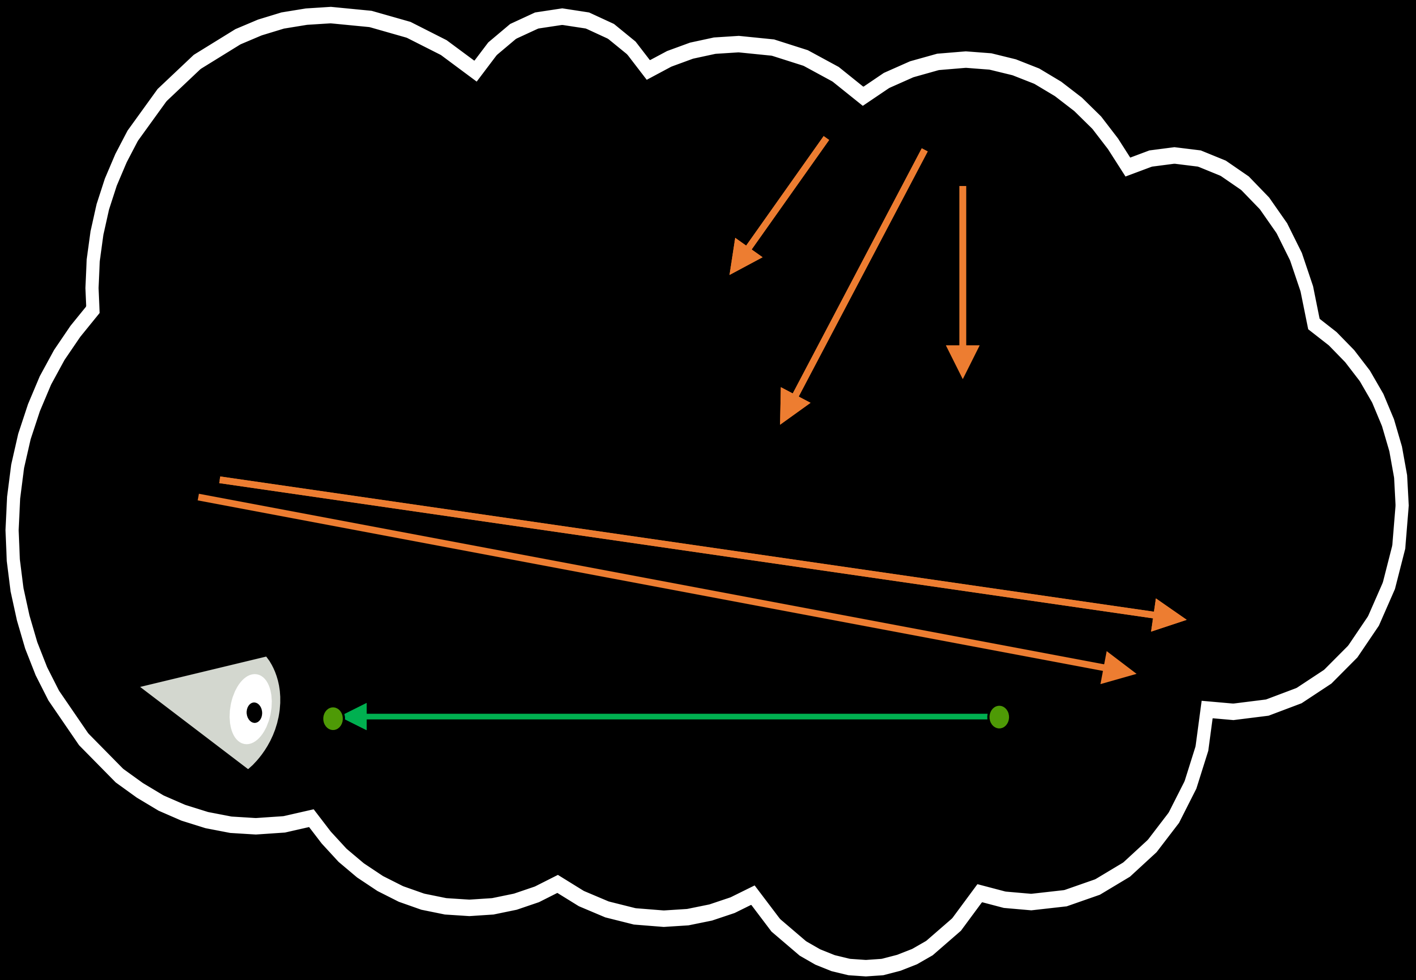

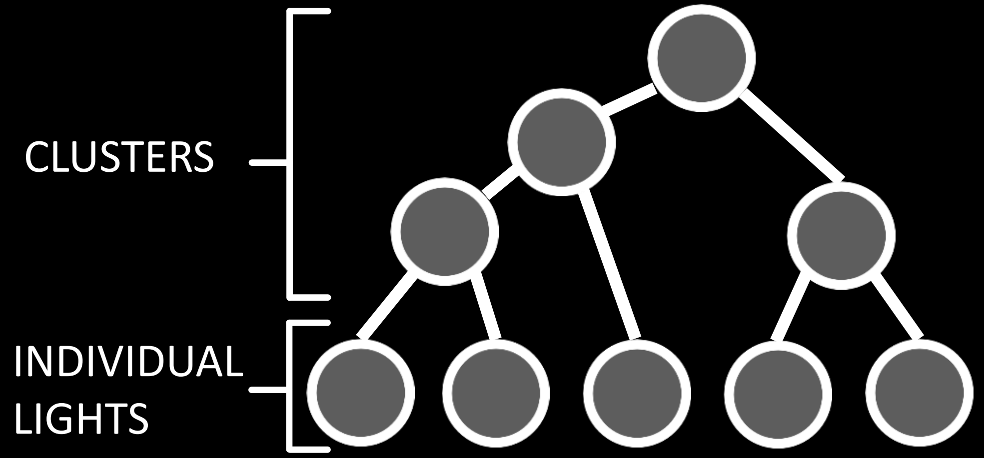

SCALABILITY

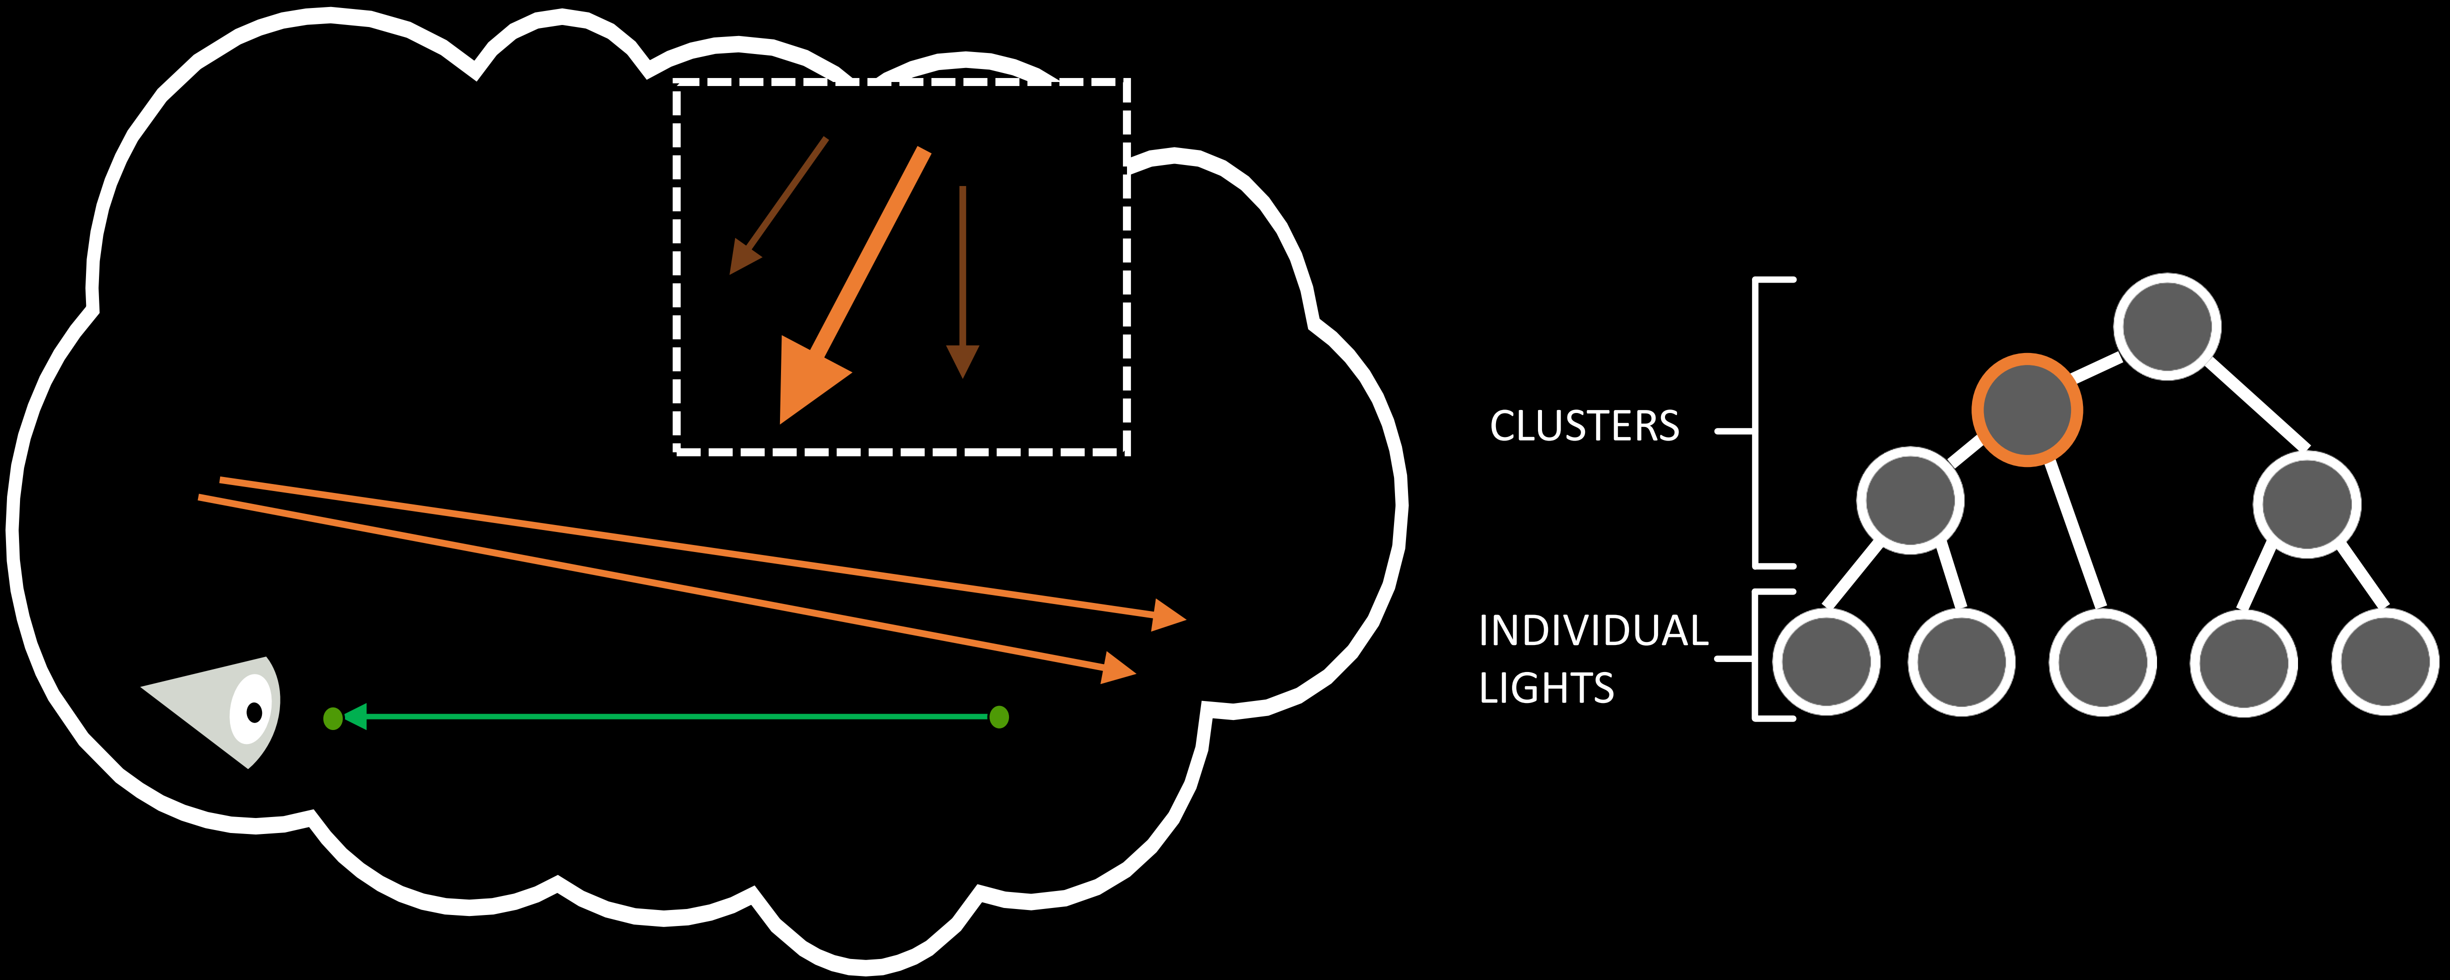

RELATED WORKS

Previous works have already explored a combination of VRLs with scalable techniques:

- Adaptive light-slice for virtual ray light [Frederickx al. 2015]
- Adaptive matrix column sampling and completion for rendering participating media [Huo et al. 2016]

Our solution: Upper bound

$$\frac{VRL(u,v)}{w(u,v)^2 p(u,v)} < \mathrm{B}$$

B=?

Our solution: Upper bound

$$\frac{\Phi f(u,v) T_r(u) T_r(w(u,v))}{w(u,v)^2 p(u,v)} < \mathrm{B}$$

$\Phi f(u,v)$: Constant within the VRL cluster

Our solution: Upper bound

$$\frac{\Phi f(u,v)T_r(u)T_r(w(u,v))}{w(u,v)^2 p(u,v)} < \mathrm{B}$$

$\Phi f(u,v)$: Constant within the VRL cluster

$T_r(u)$: Impossible to control, assuming worst case => 1

Our solution: Upper bound

$$\frac{\Phi f(u,v) T_r(u) T_r(w(u,v))}{w(u,v)^2 p(u,v)} < B$$

$\Phi f(u,v)$: Constant within the VRL cluster

$T_r(u)$: Impossible to control, assuming worst case => 1

$T_r(w(u,v))$: Based on the min distance

=> $h_{min} \leq w_k(u,v)$

=> $T_r(w\,(u,v)) < T_r(h_{min})$

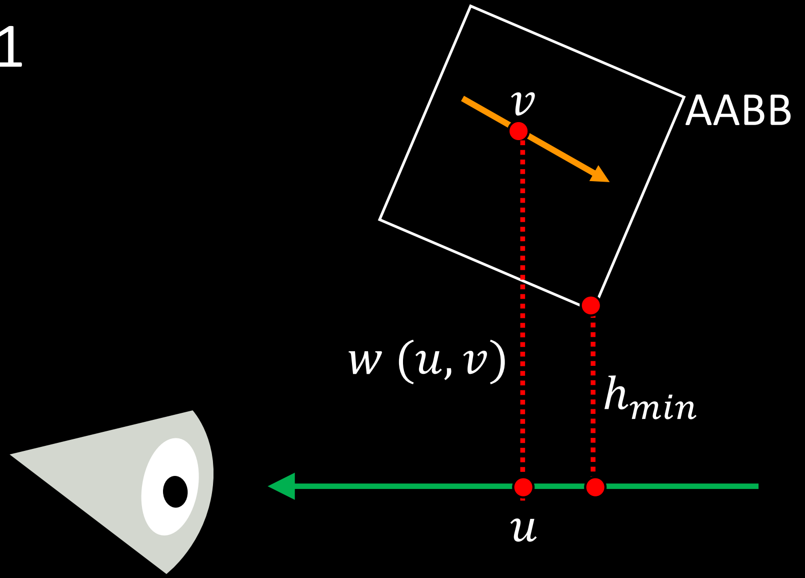

Our solution: Upper bound

$$\frac{\Phi f(u,v) T_r(u) T_r(w(u,v))}{w(u,v)^2 p(u,v)} < B$$

$\Phi f(u,v)$: Constant within the VRL cluster

$T_r(u)$: Impossible to control, assuming worst case => 1

$T_r(w(u,v))$: Based on the min distance

=> $h_{min} \leq w_k(u,v)$

=> $T_r(w\ (u,v)) < T_r(h_{min})$

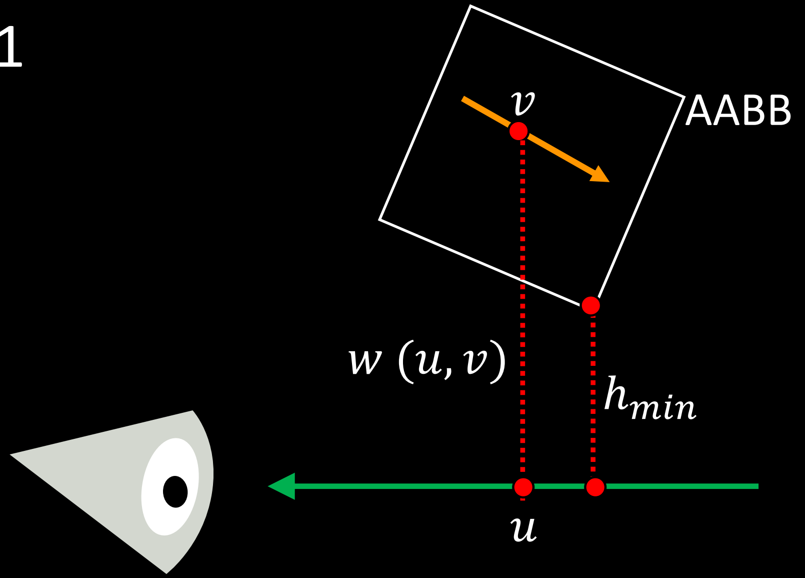

Our solution: Upper bound

$$w_k(u,v)^2 p(u,v)$$

$$= w_k(u,v)^2 p(u|v) p(v)$$

Our solution: Upper bound

$$w_k(u,v)^2 p(u,v)$$
$$= w_k(u,v)^2 p(u|v) p(v)$$

Worst case when VRL and sensor ray are parallel.

$p(v) < \frac{1}{L_{max}}$, with L_{max} the maximum length inside the cluster.

Our solution: Upper bound

$$w_k(u,v)^2 p(u,v)$$
$$= w_k(u,v)^2 p(u|v) p(v)$$

Worst case when VRL and sensor ray are parallel.

$p(v) < \frac{1}{L_{max}}$, with L_{max} the maximum length inside the cluster.

Our solution: Upper bound

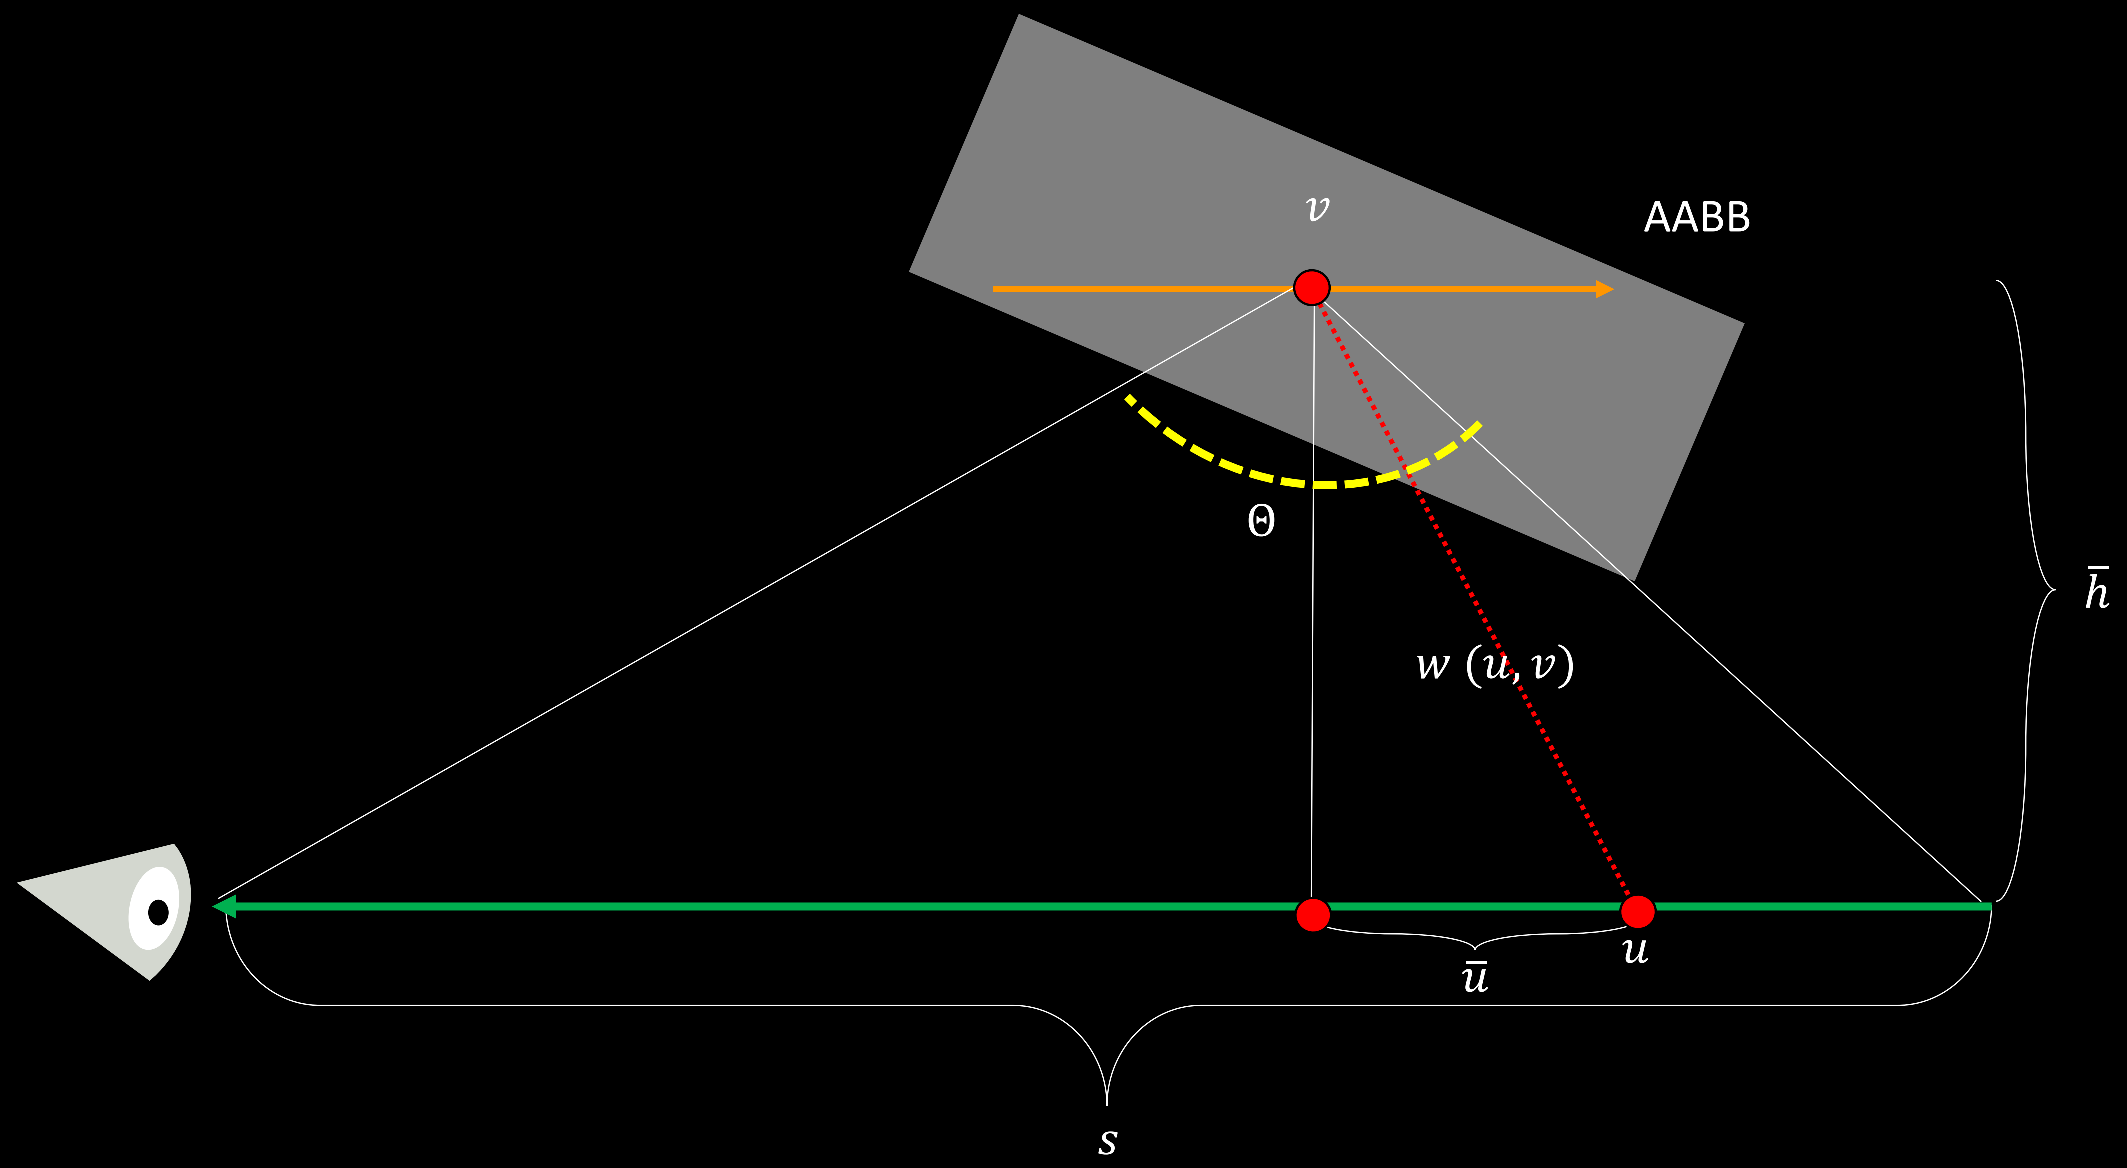

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}(\bar{h}^2+\bar{u}^2)}{\Theta(\bar{h}^2+\bar{u}^2)}$$

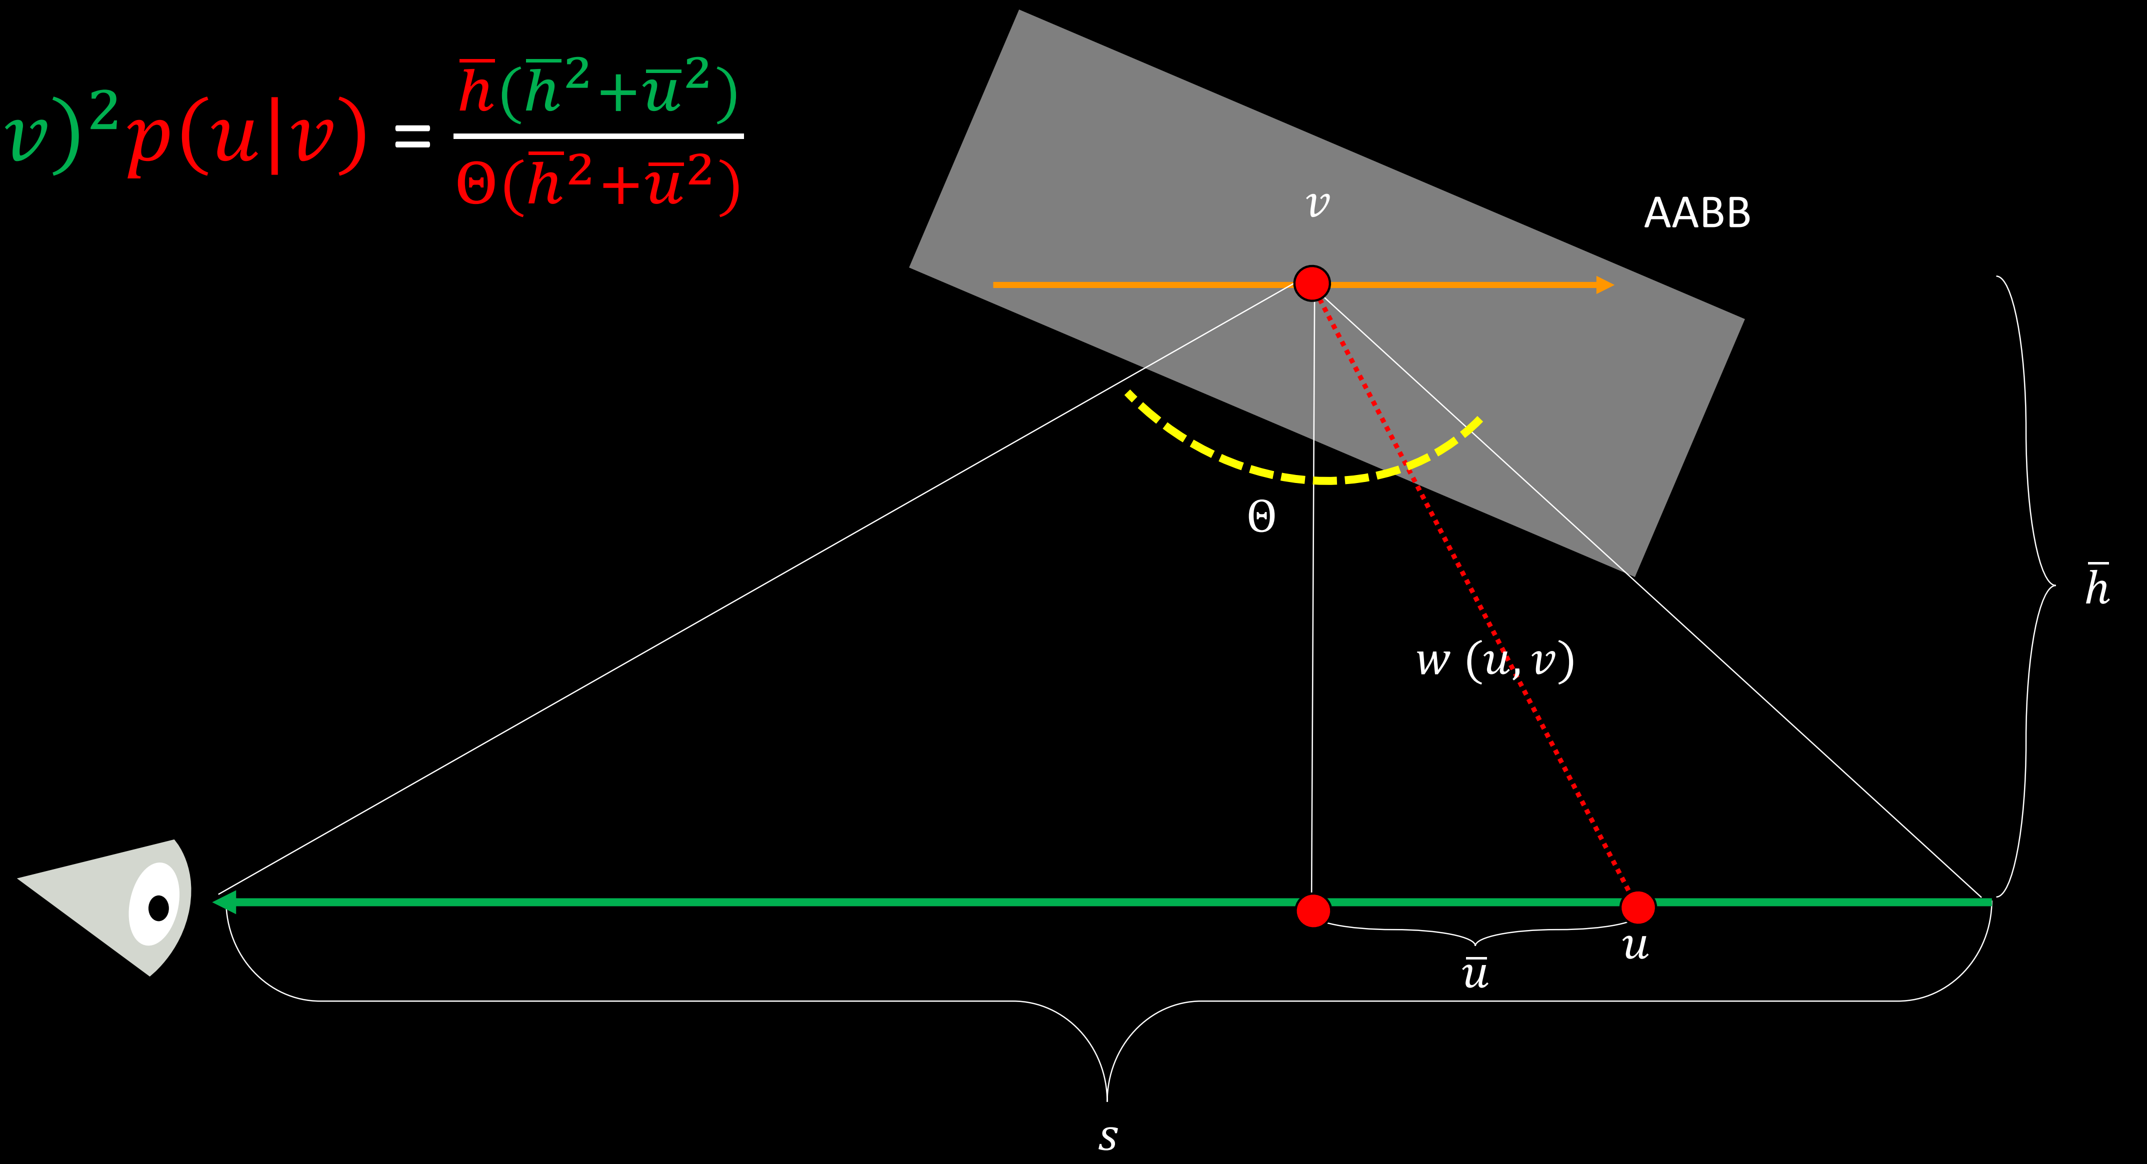

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}\cancel{(\bar{h}^2+\bar{u}^2)}}{\Theta\cancel{(\bar{h}^2+\bar{u}^2)}}$$

$$= \frac{\bar{h}}{\Theta}$$

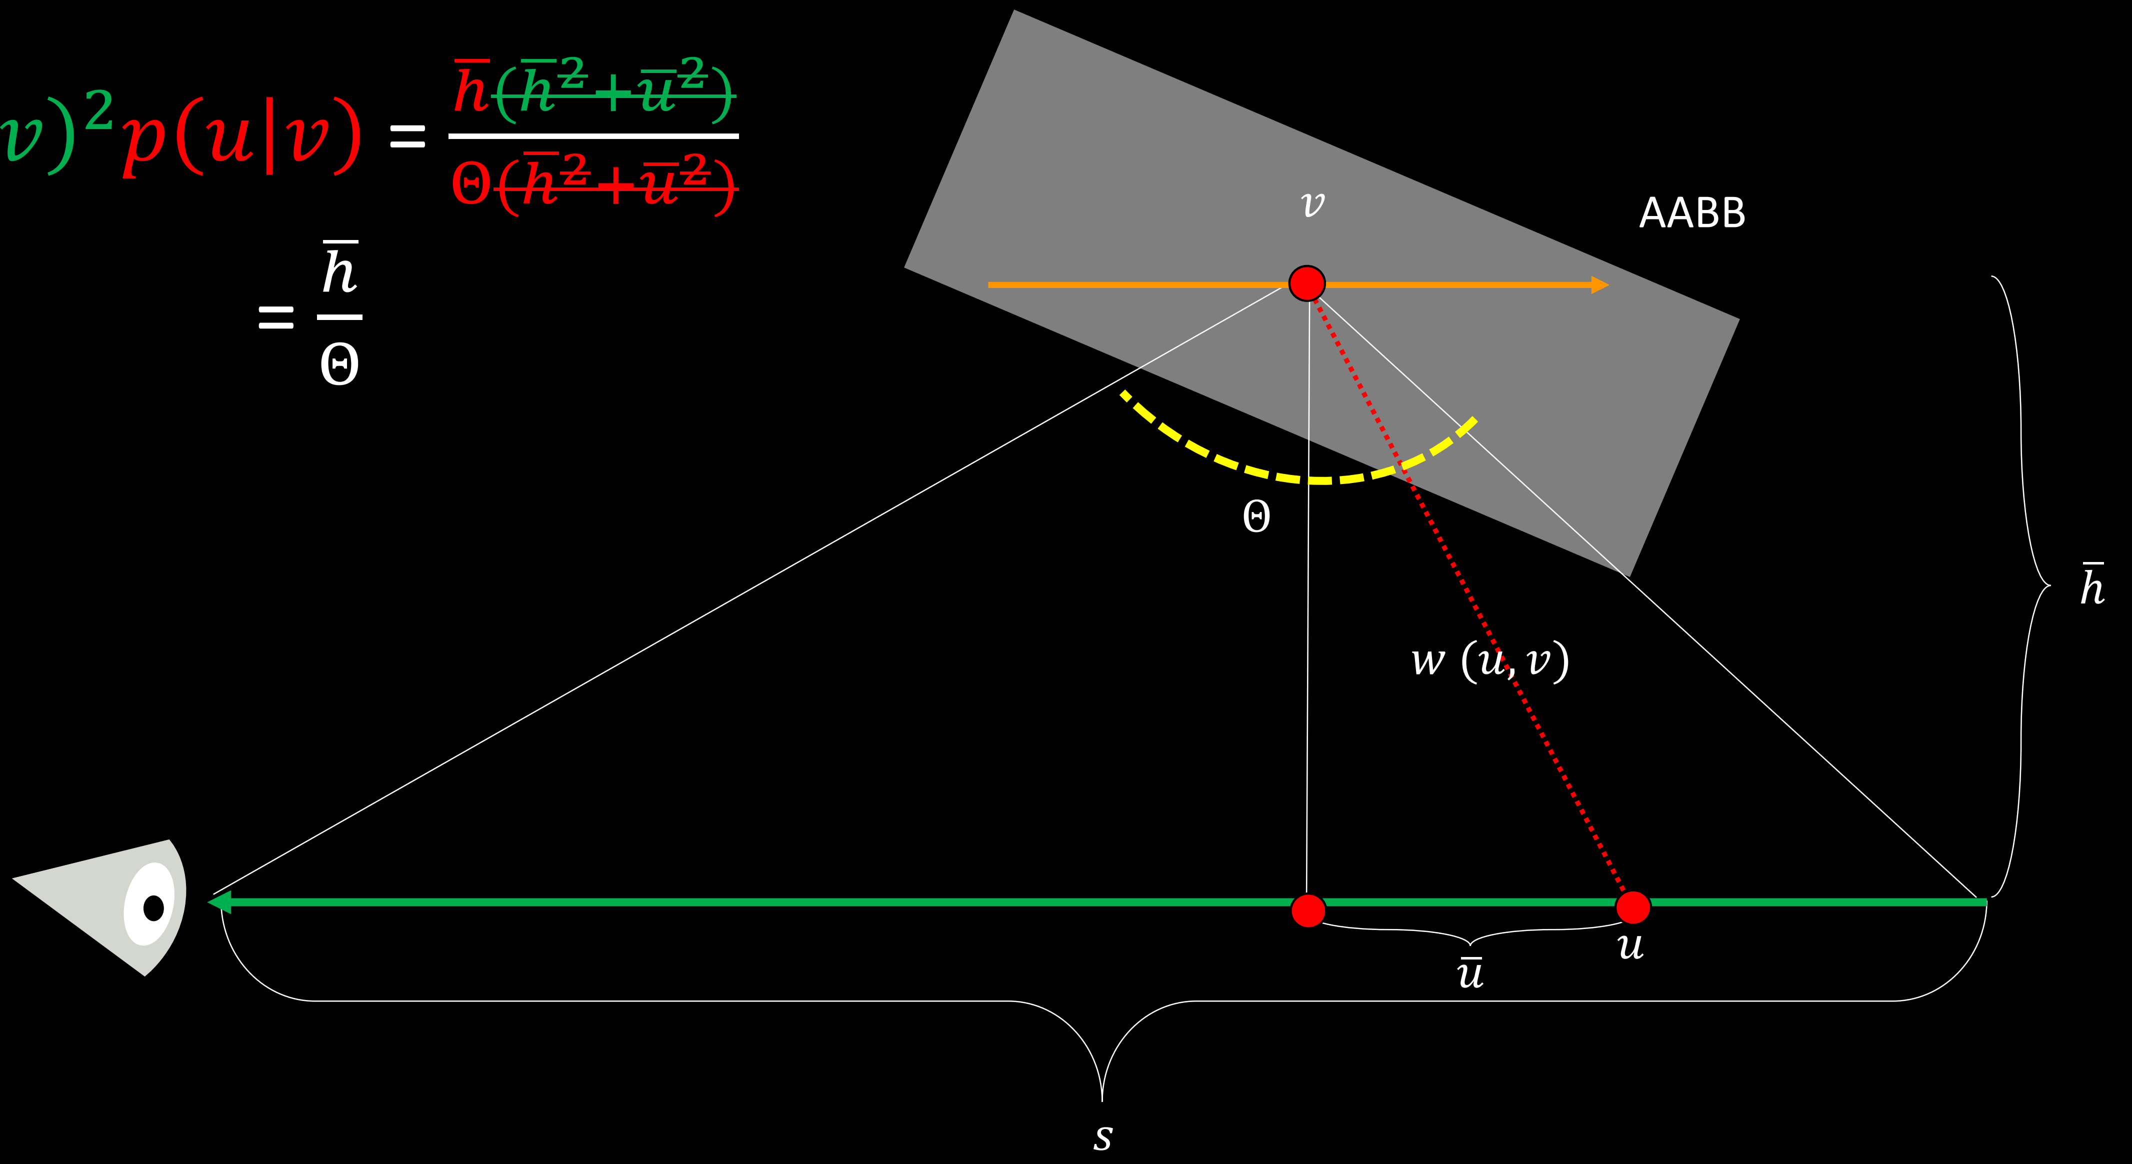

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}\cancel{(\bar{h}^2+\bar{u}^2)}}{\Theta\cancel{(\bar{h}^2+\bar{u}^2)}}$$

$$= \frac{\bar{h}}{\Theta}$$

$$\bar{h} > h_{min}$$

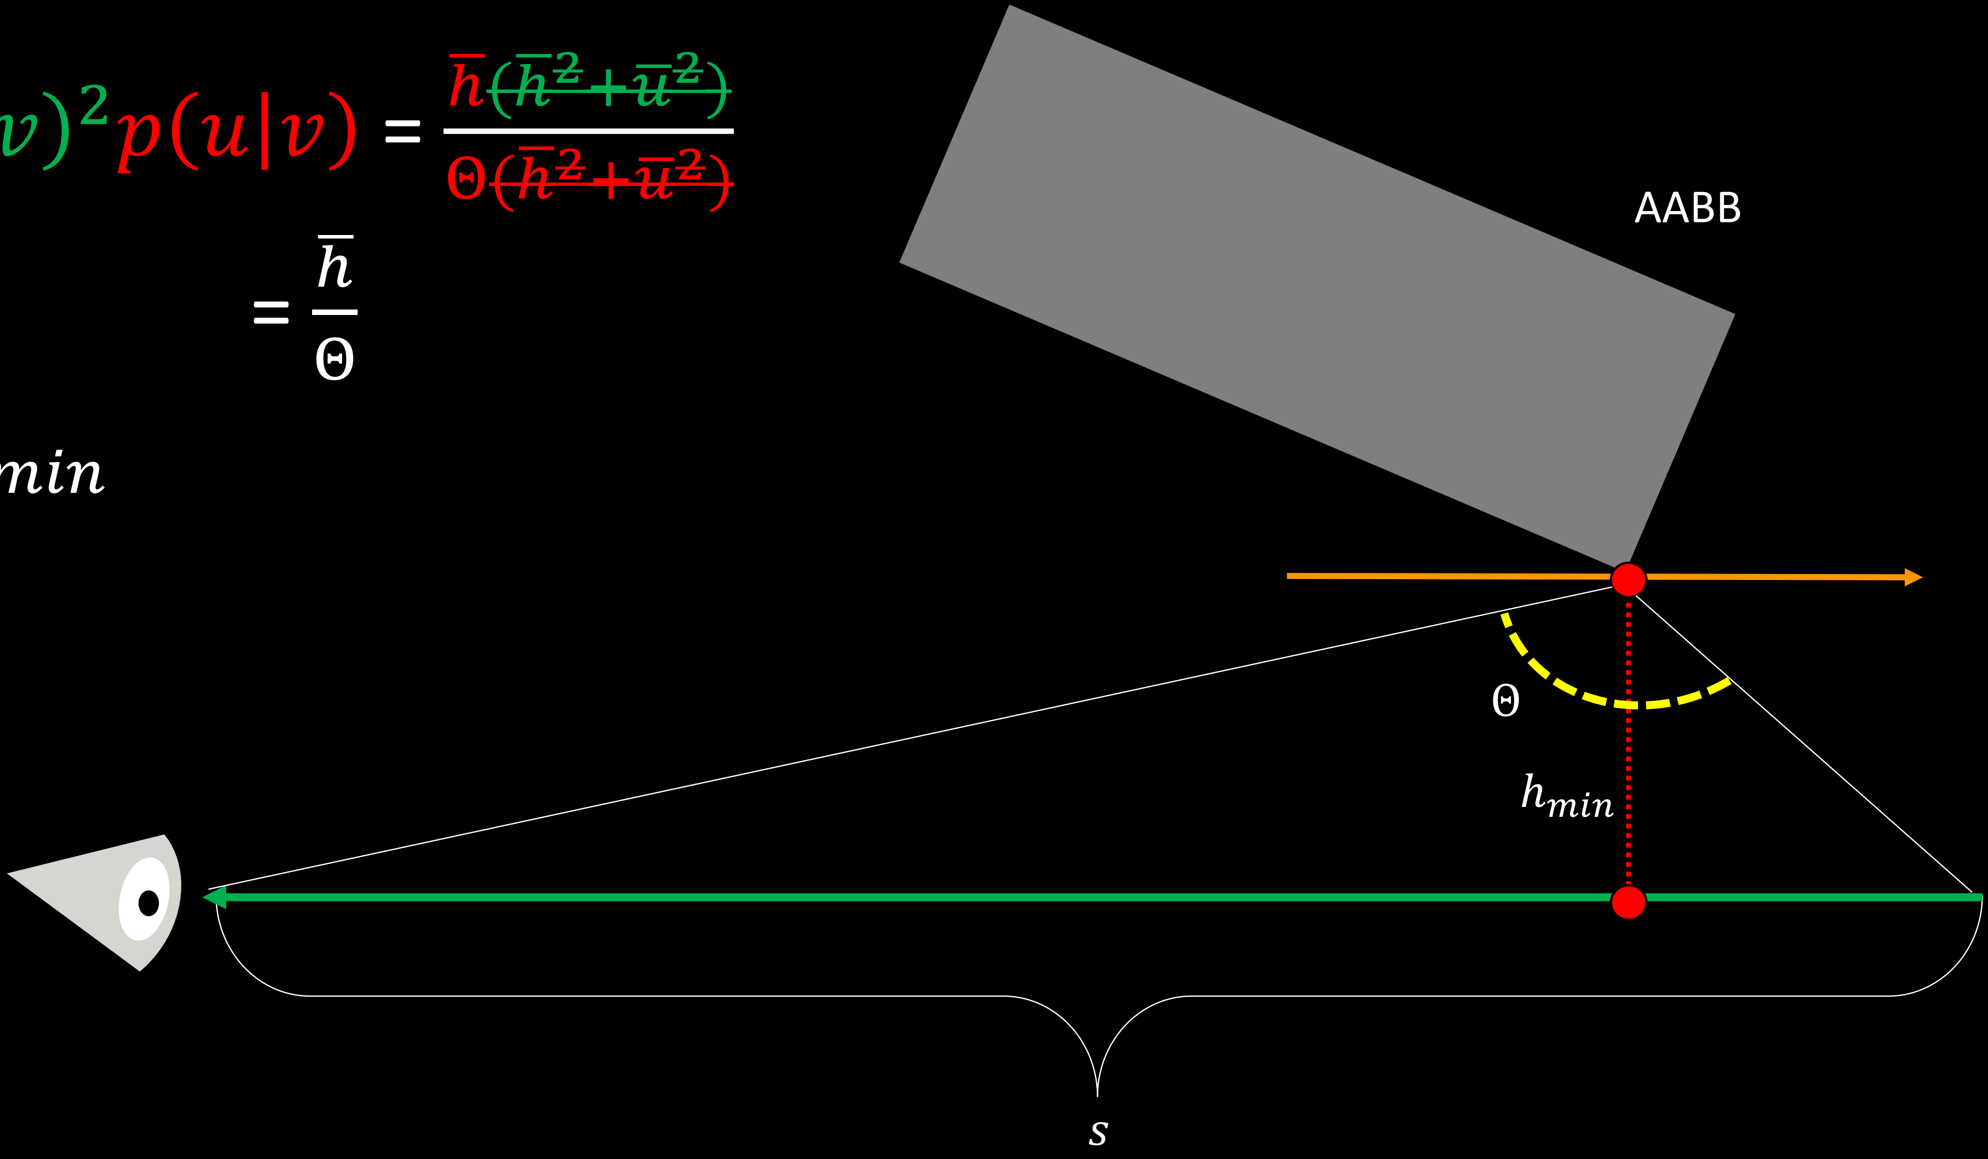

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}\,\not{(\bar{h}^2+\bar{u}^2)}}{\Theta\,\not{(\bar{h}^2+\bar{u}^2)}}$$

$$= \frac{\bar{h}}{\Theta}$$

$$\bar{h} > h_{min}$$

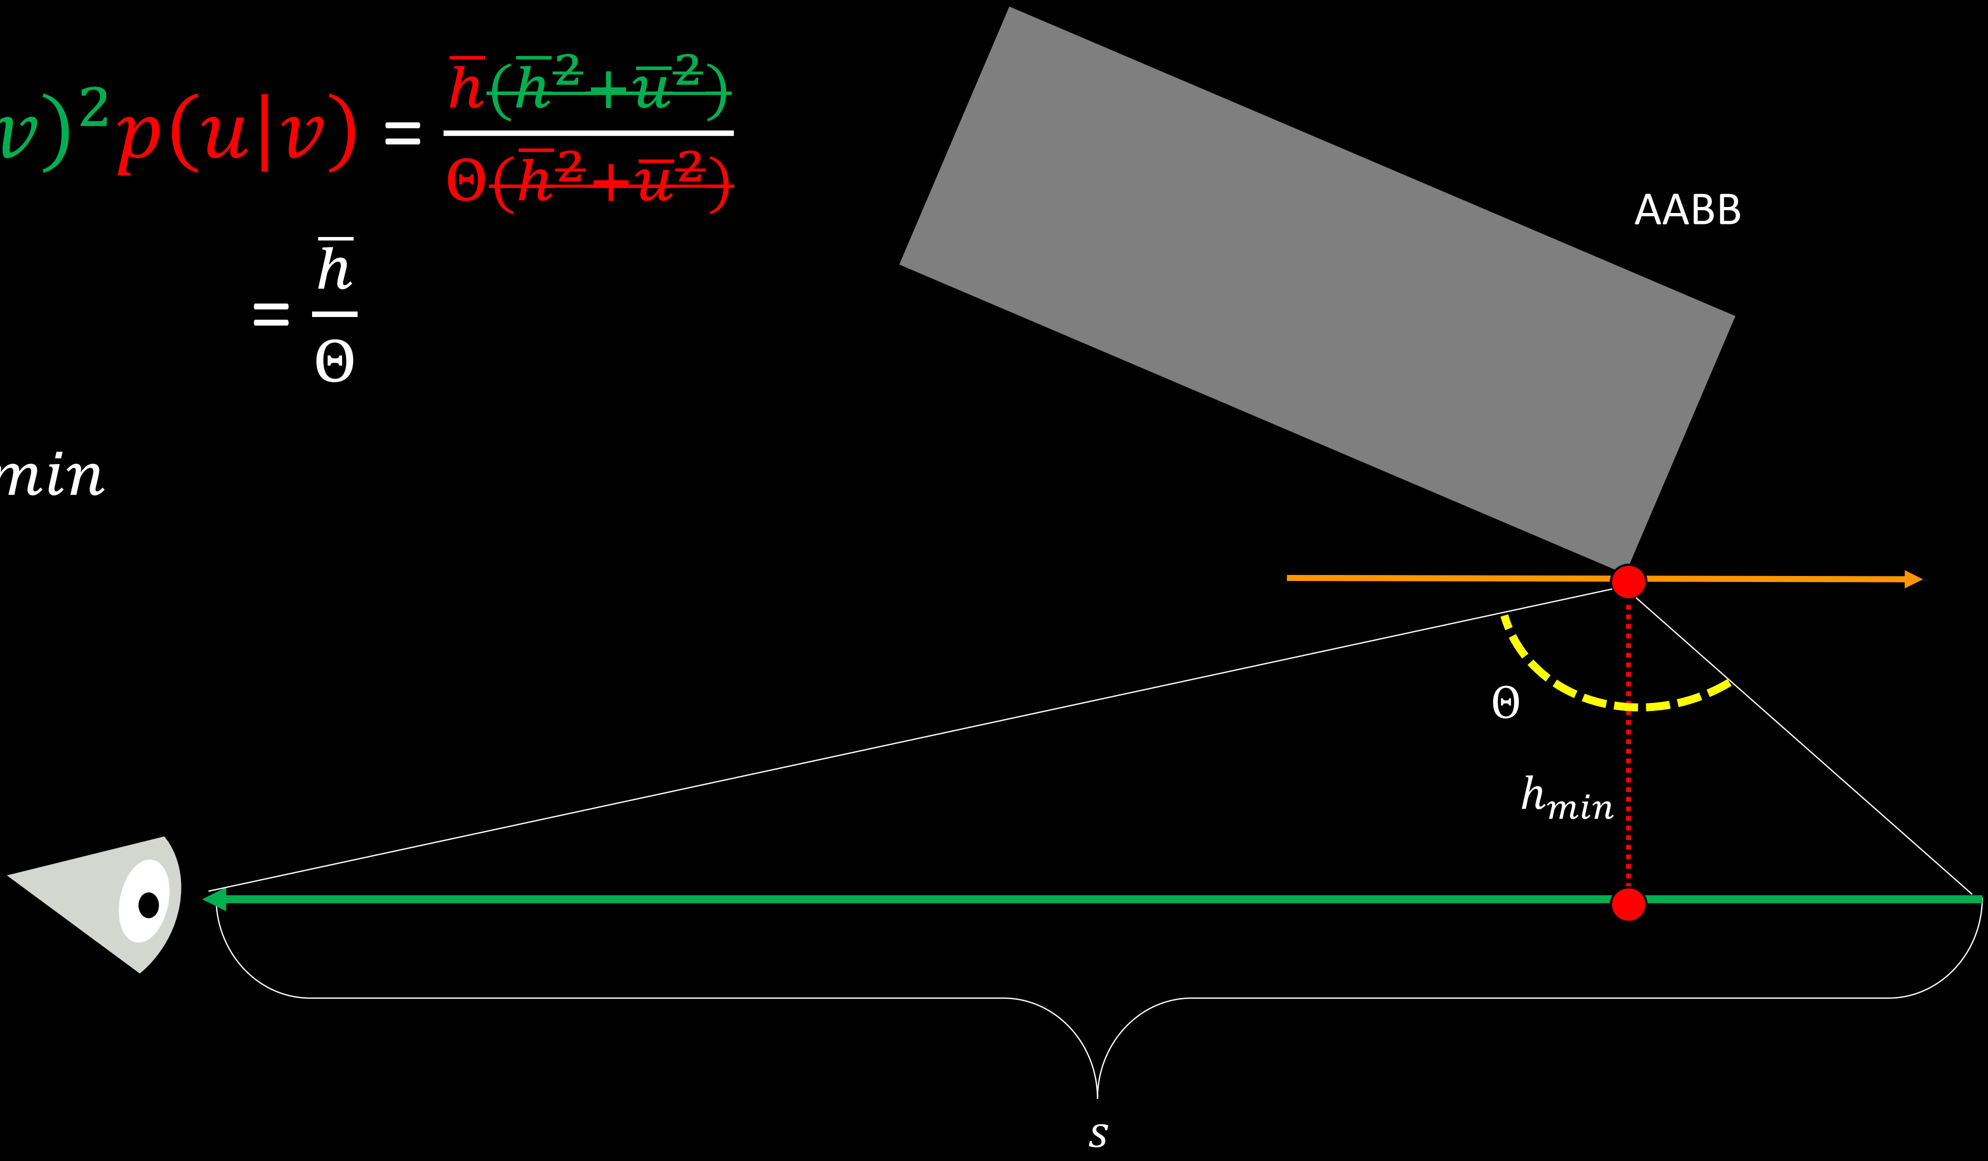

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}\,\cancel{(\bar{h}^2+\bar{u}^2)}}{\Theta\,\cancel{(\bar{h}^2+\bar{u}^2)}}$$

$$= \frac{\bar{h}}{\Theta}$$

$$\bar{h} > h_{min}$$

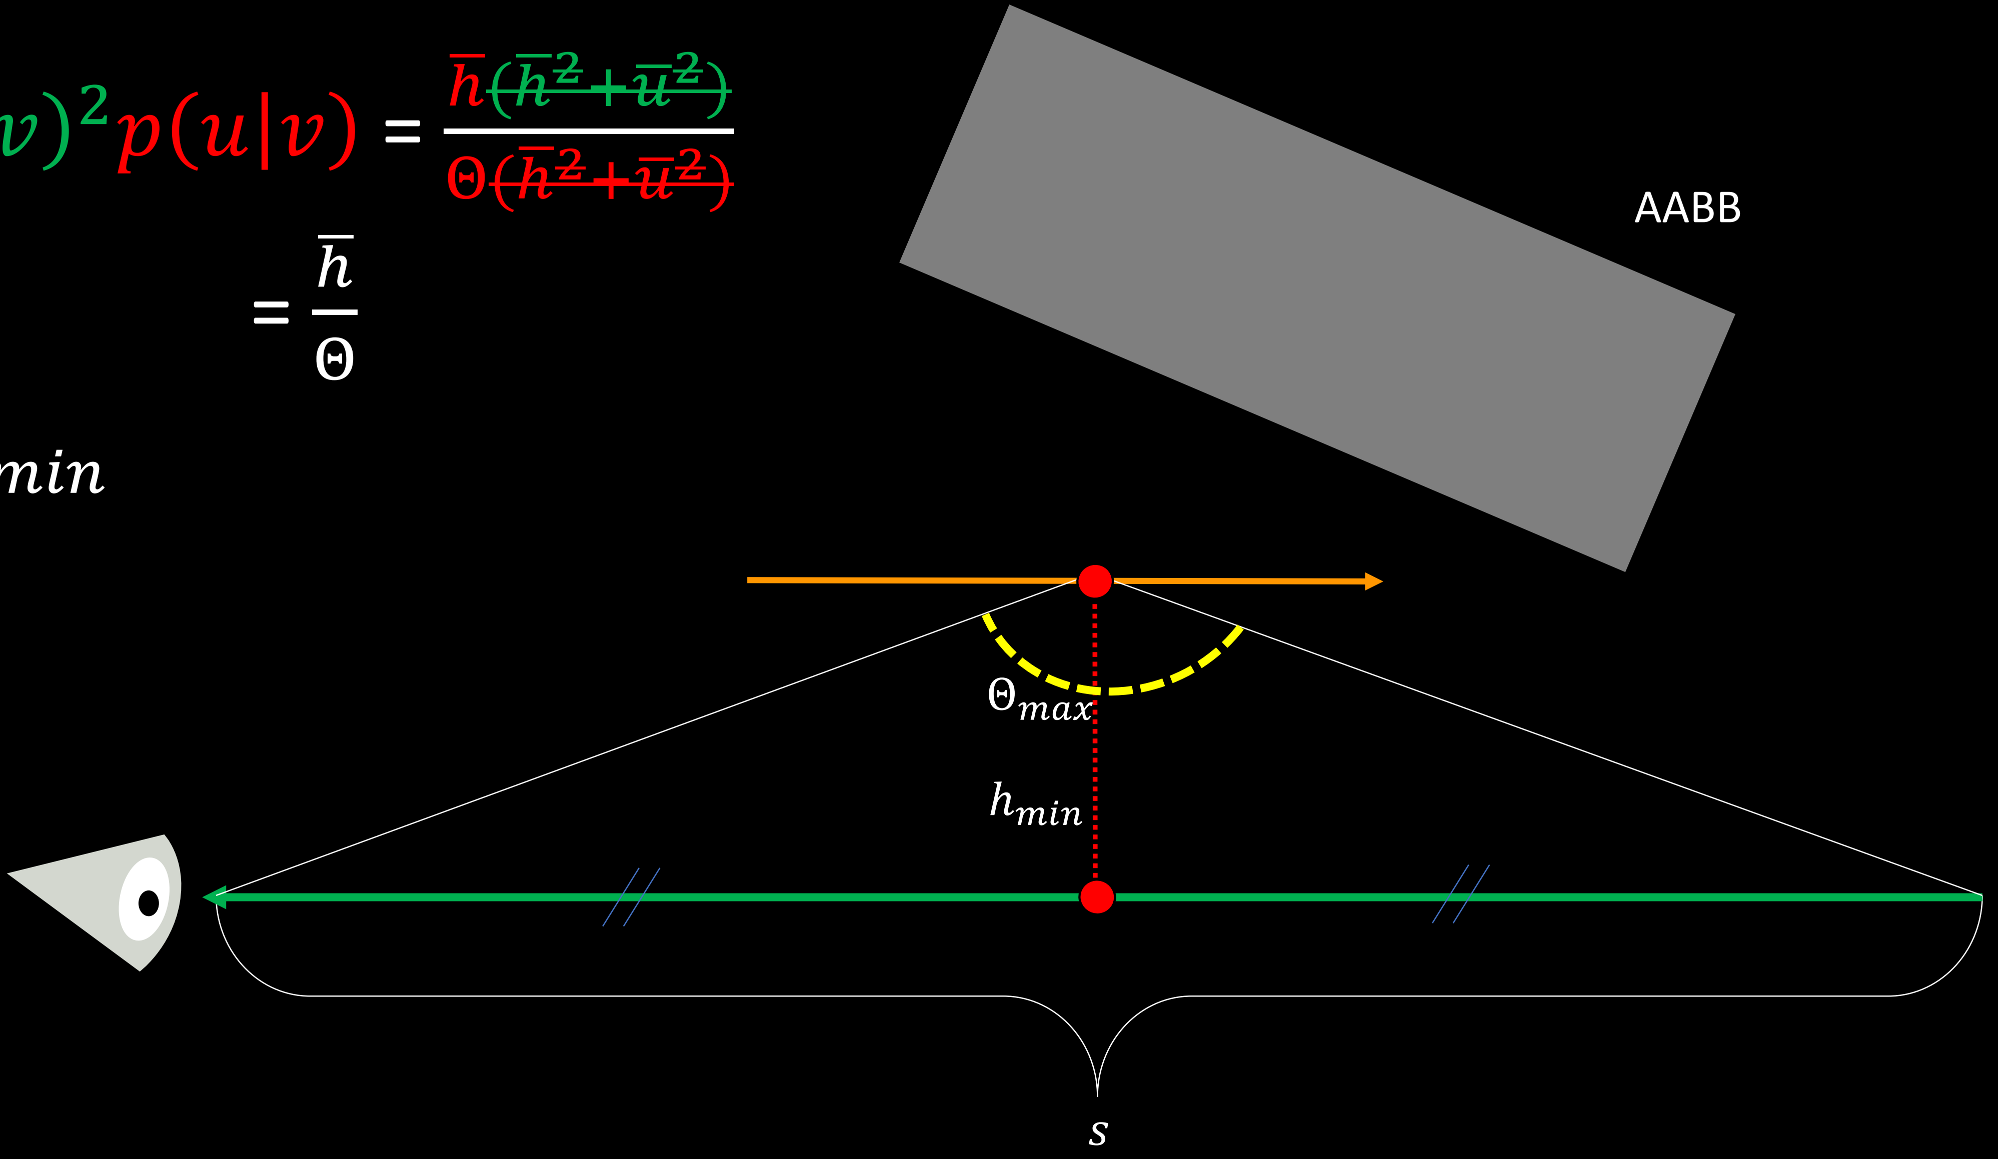

Our solution: Upper bound

$$w(u,v)^2 p(u|v) = \frac{\bar{h}\cancel{(\bar{h}^2+\bar{u}^2)}}{\Theta\cancel{(\bar{h}^2+\bar{u}^2)}}$$

$$= \frac{\bar{h}}{\Theta}$$

$\bar{h} > h_{min}$

$\Theta < \Theta_{max}$

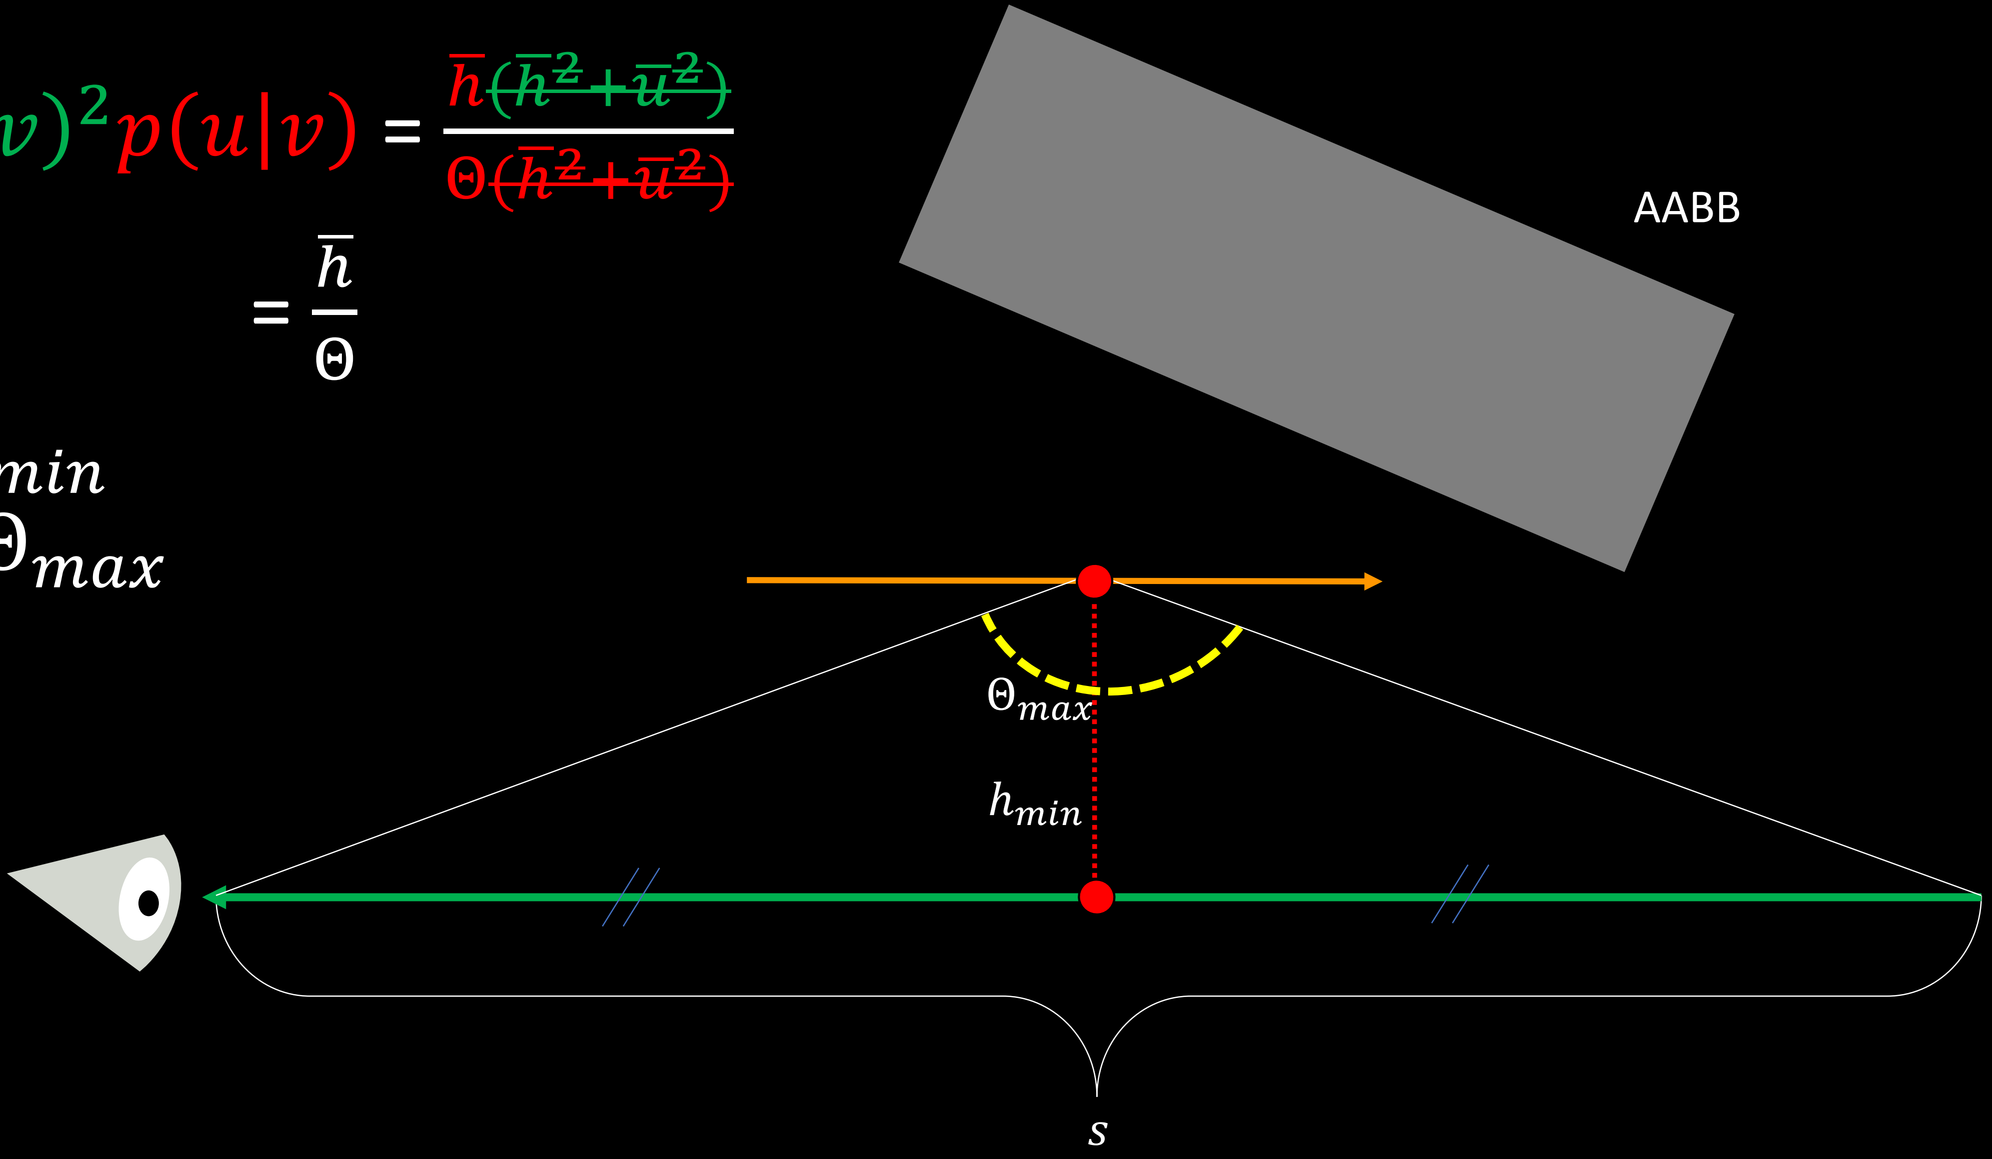

Our solution: Upper bound

$$B = \frac{\Phi f(u,v) T_r(\mathrm{h_{min}}) \Theta_{\mathrm{max}} \mathrm{L_{max}}}{h_{min}}$$

Our solution: Light tree

Agglomerative approach [Walter et al. 2008]:

- Not multithreaded
- Does not scale with high node overlapping $O(N^2)$

What we need:

- Fast/Parallelizable
- Agglomerative principal

Our solution: Light tree

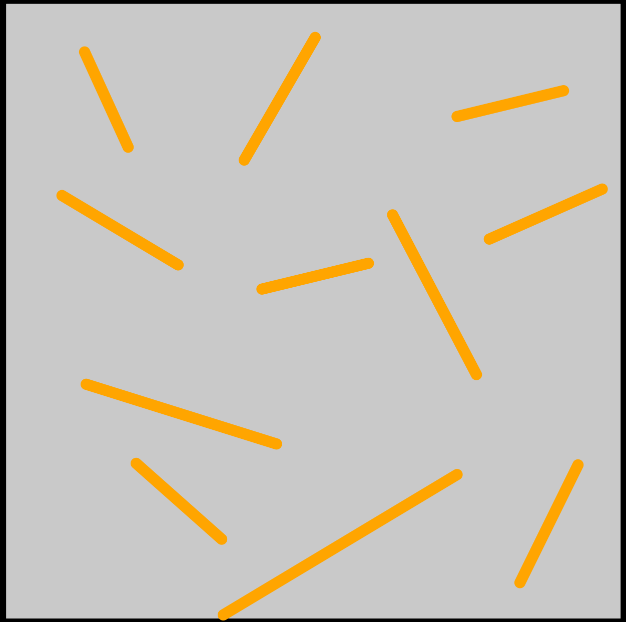

Our solution: Light tree

Step 1: Partition the space with sorting

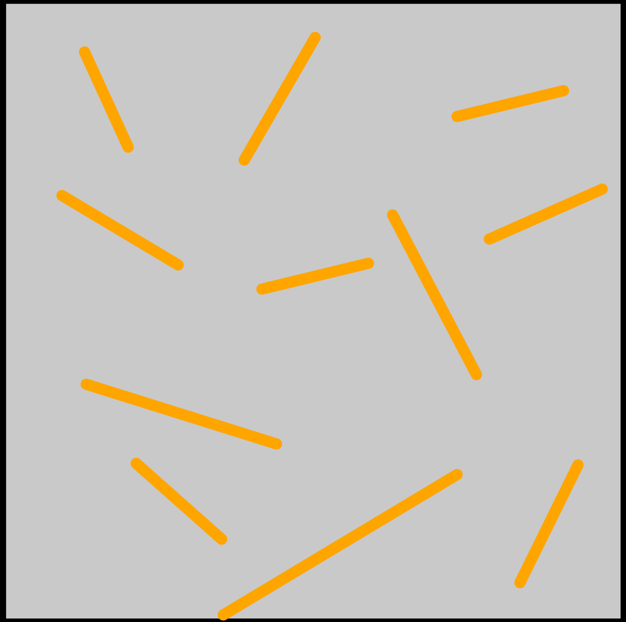

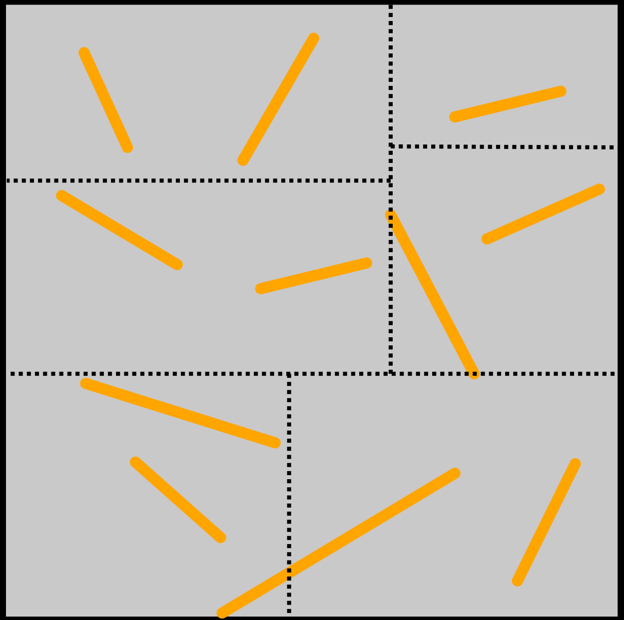

Our solution: Light tree

Step 2: Build local light tree that minimize the metric

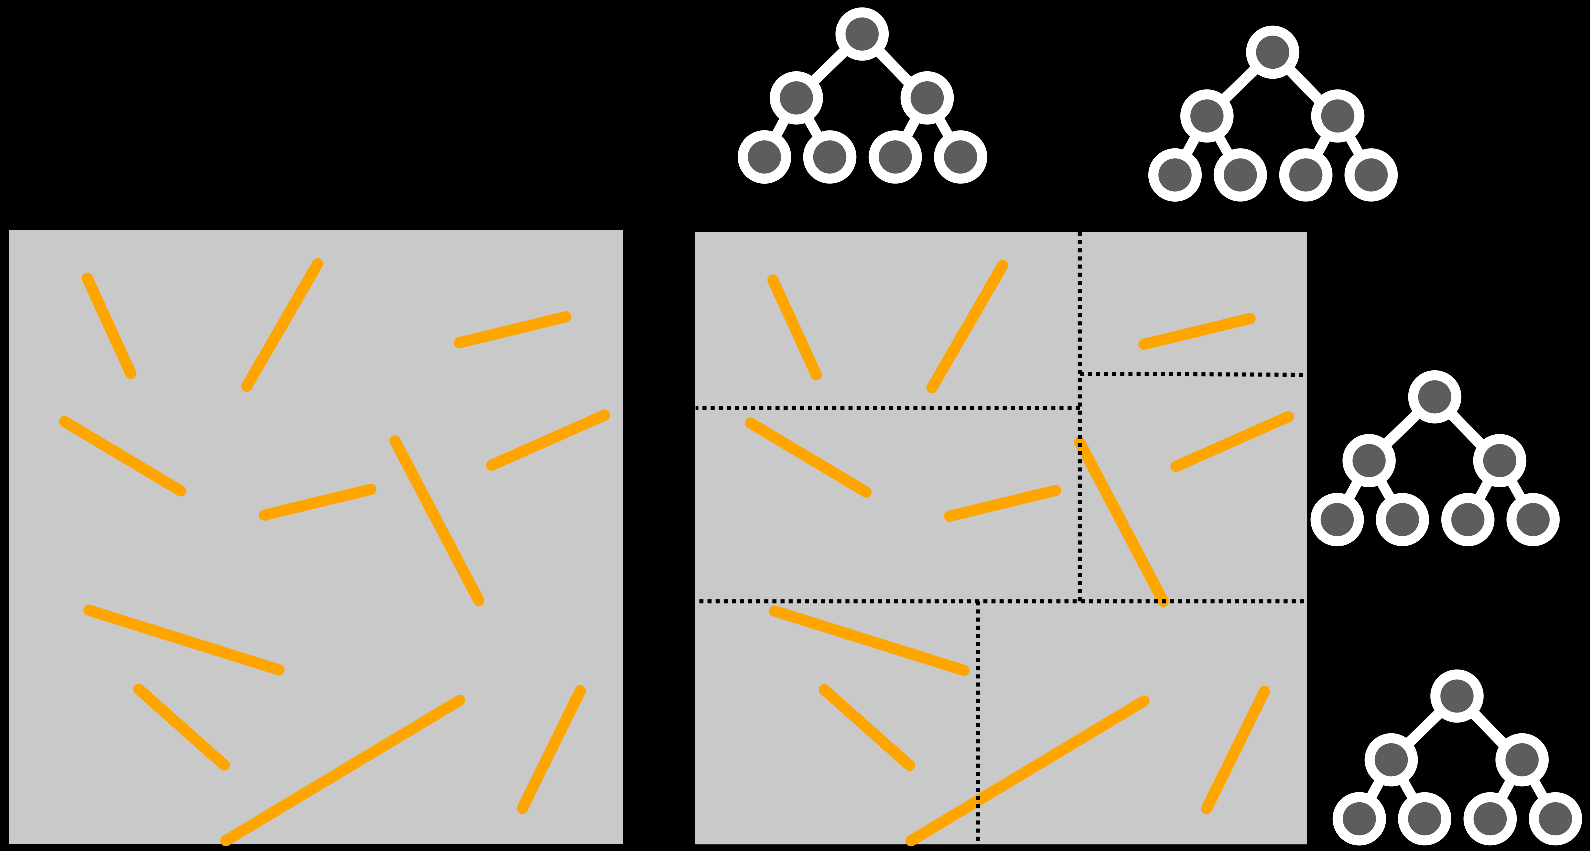

Our solution: Light tree

Step 3: Build final tree with agglomerative process

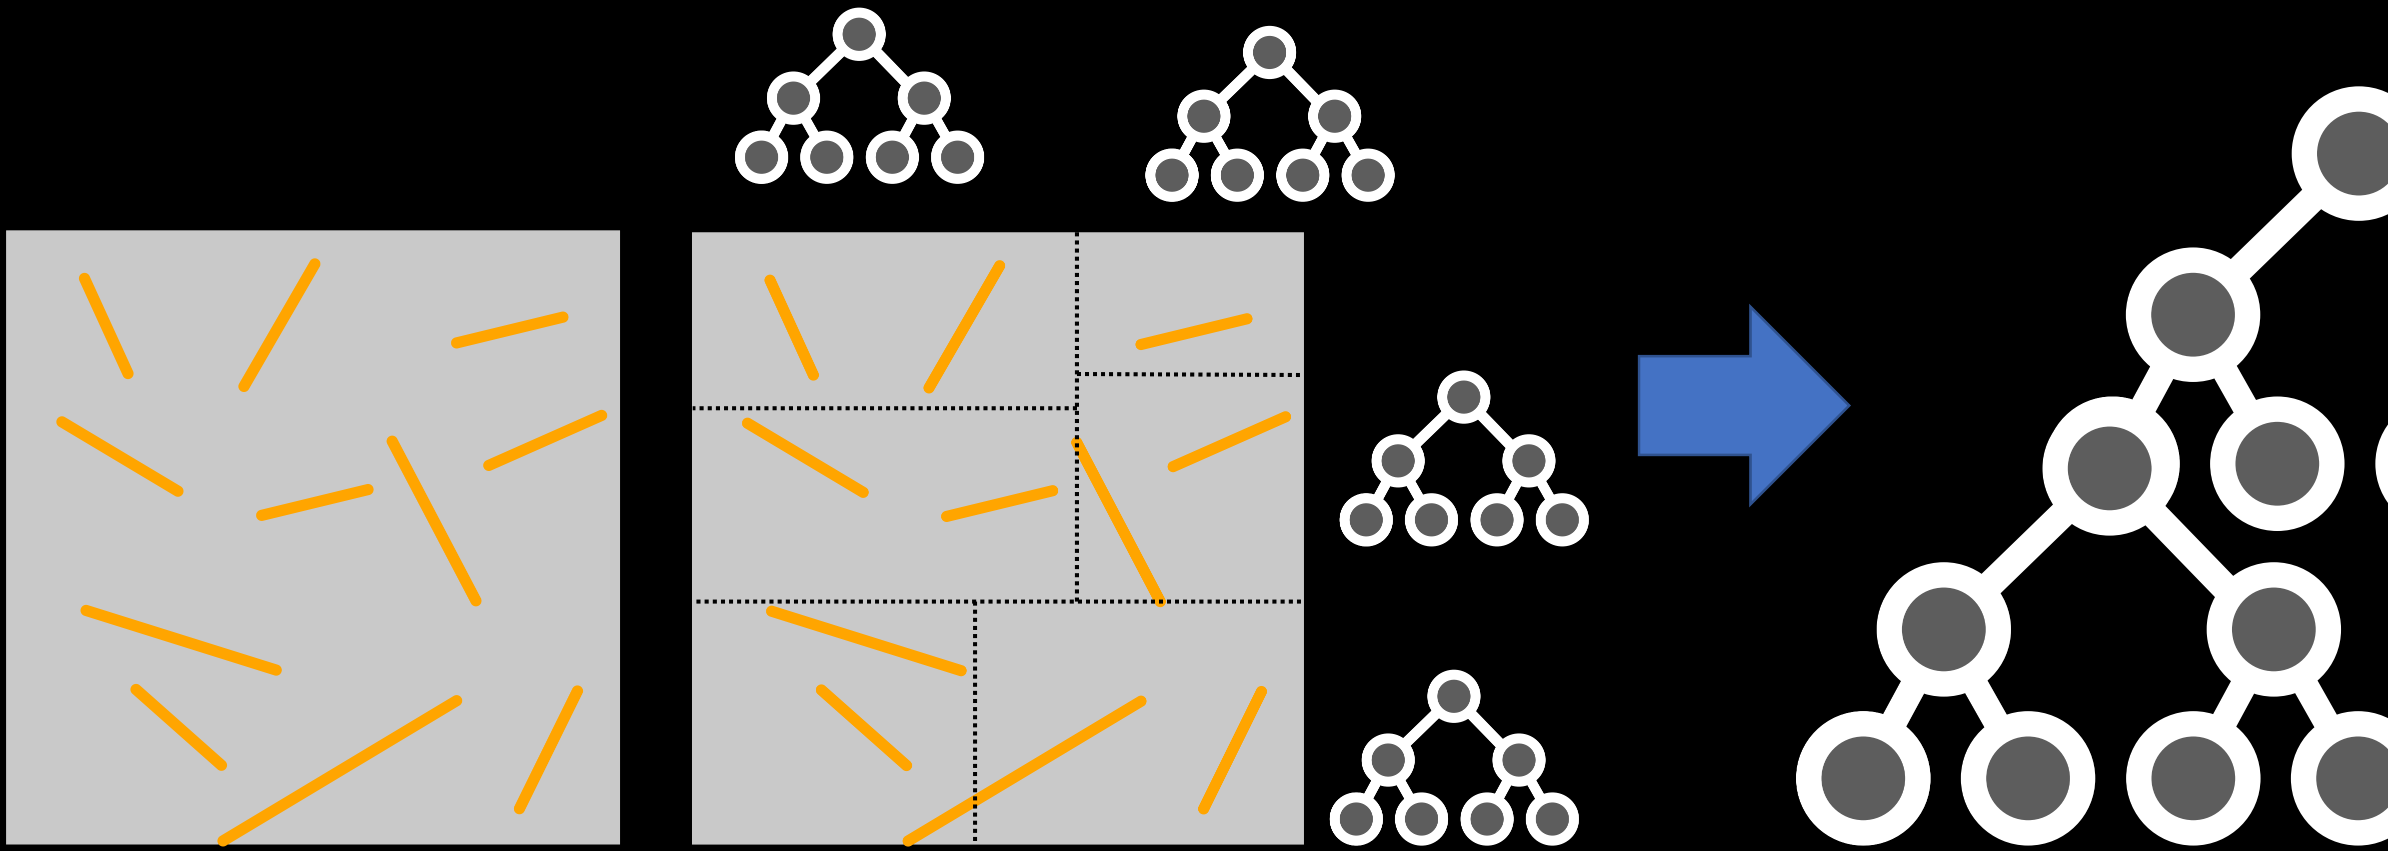

Results

- Equal time comparison
  - VPL with LC
  - VRL
- Two metrics
  - RMSE: sensitive to fireflies
  - SMAPE: robust to fireflies
- Isotropic medium, only medium-medium interactions

Results

Reference	VPL LC (1M) 344 secs RMSE: 9.01 SMAPE: 3.25	VRL LC (100k) 370 secs RMSE: 2.70 SMAPE: 4.31	VRL (10k) 323 secs RMSE: 5.46 SMAPE: 9.93
	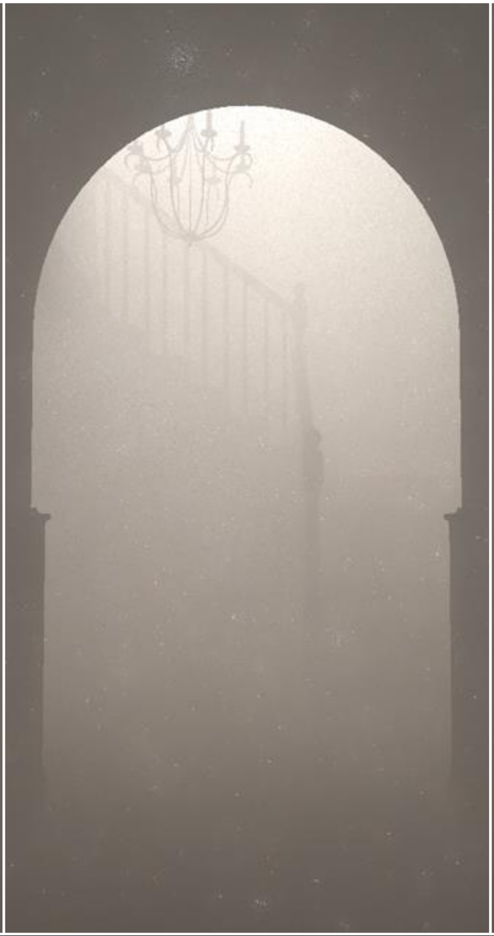		

Results

VPL LC (1M) - 147 secs RMSE: 0.33 SMAPE: 2.44	VRL LC (100k) - 121 secs RMSE: 0.01 SMAPE 1.88	VRL (10k) - 147 secs RMSE: 0.05 SMAPE 9.11

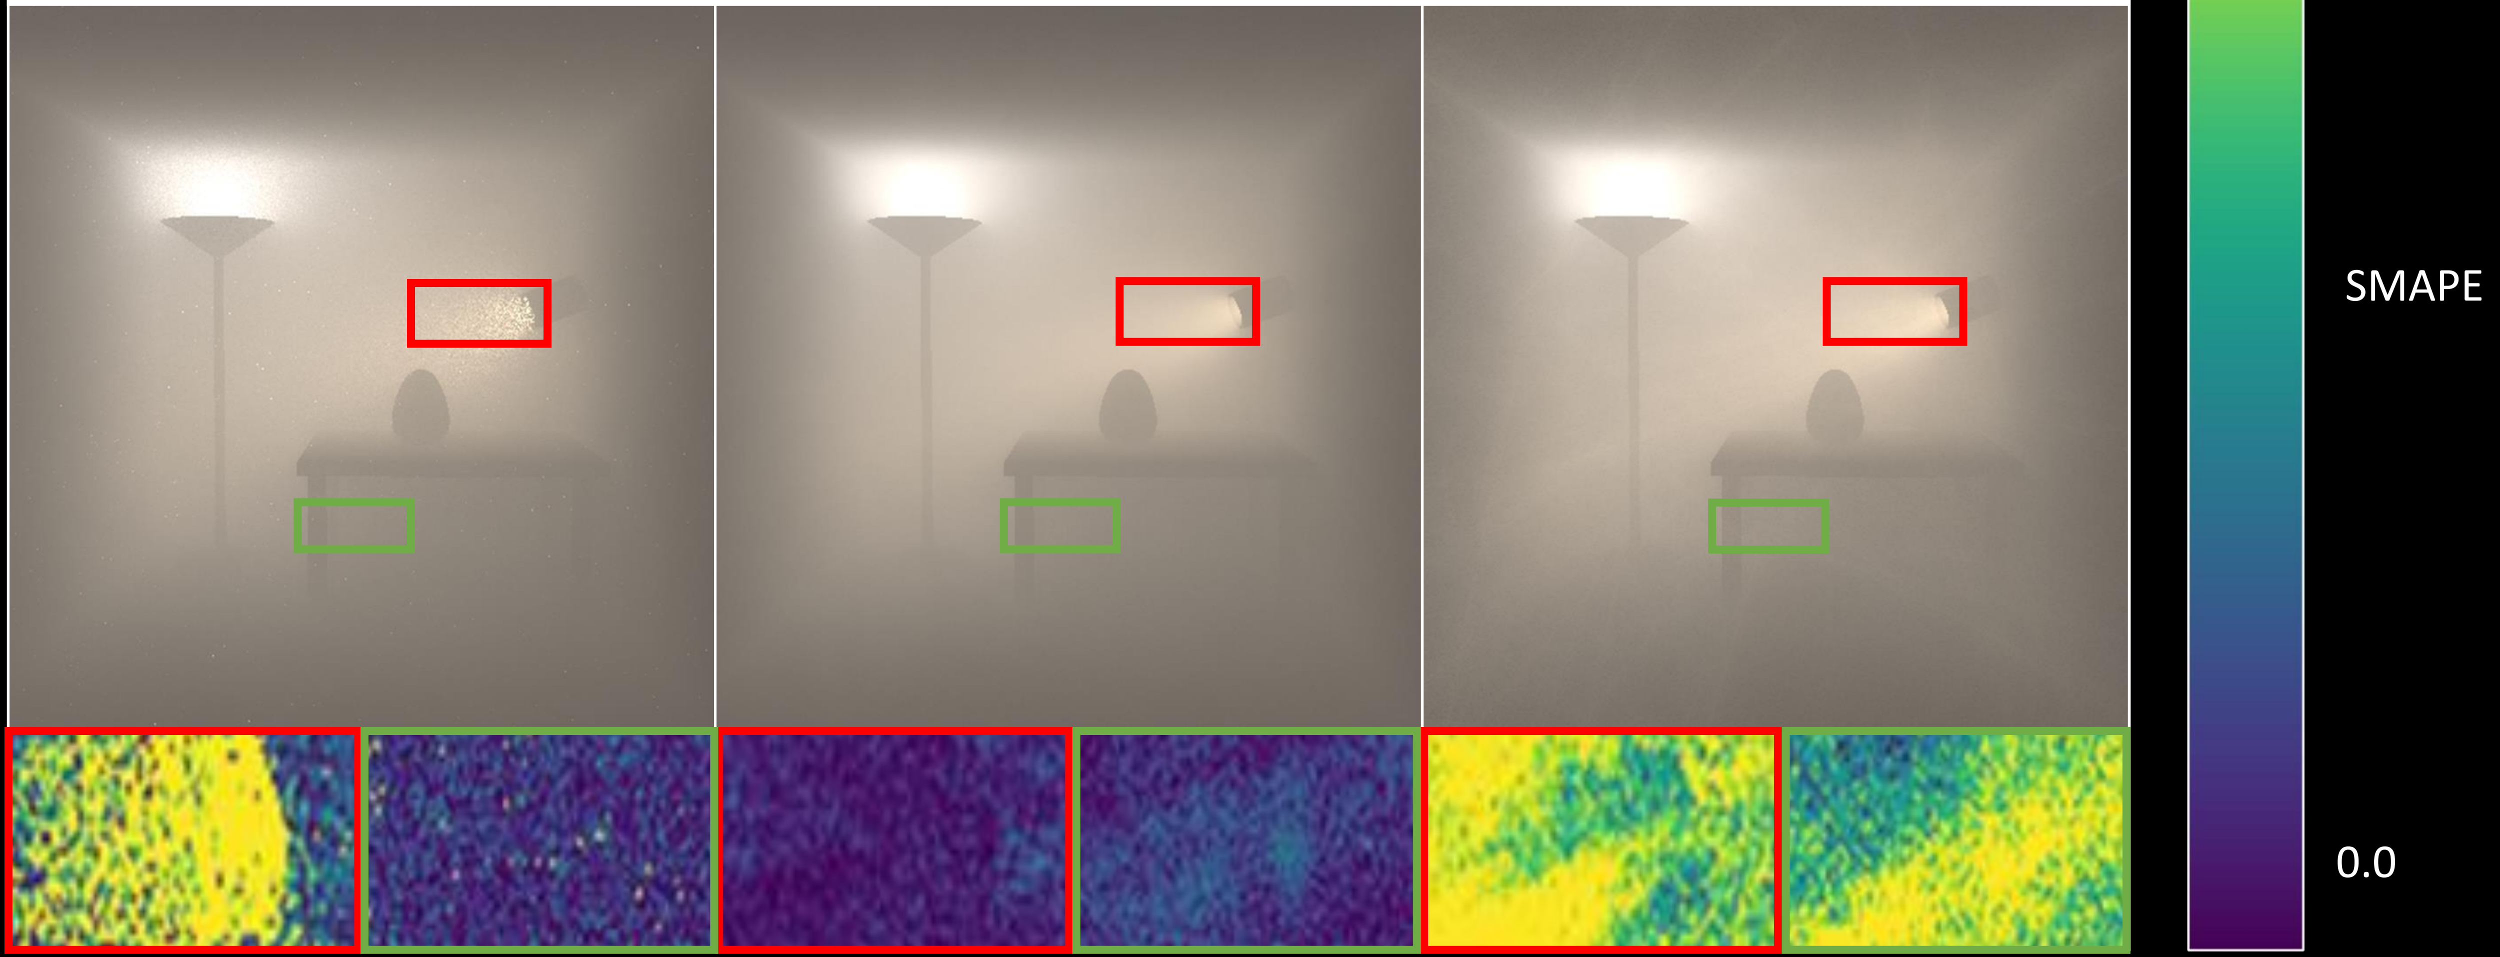

Results

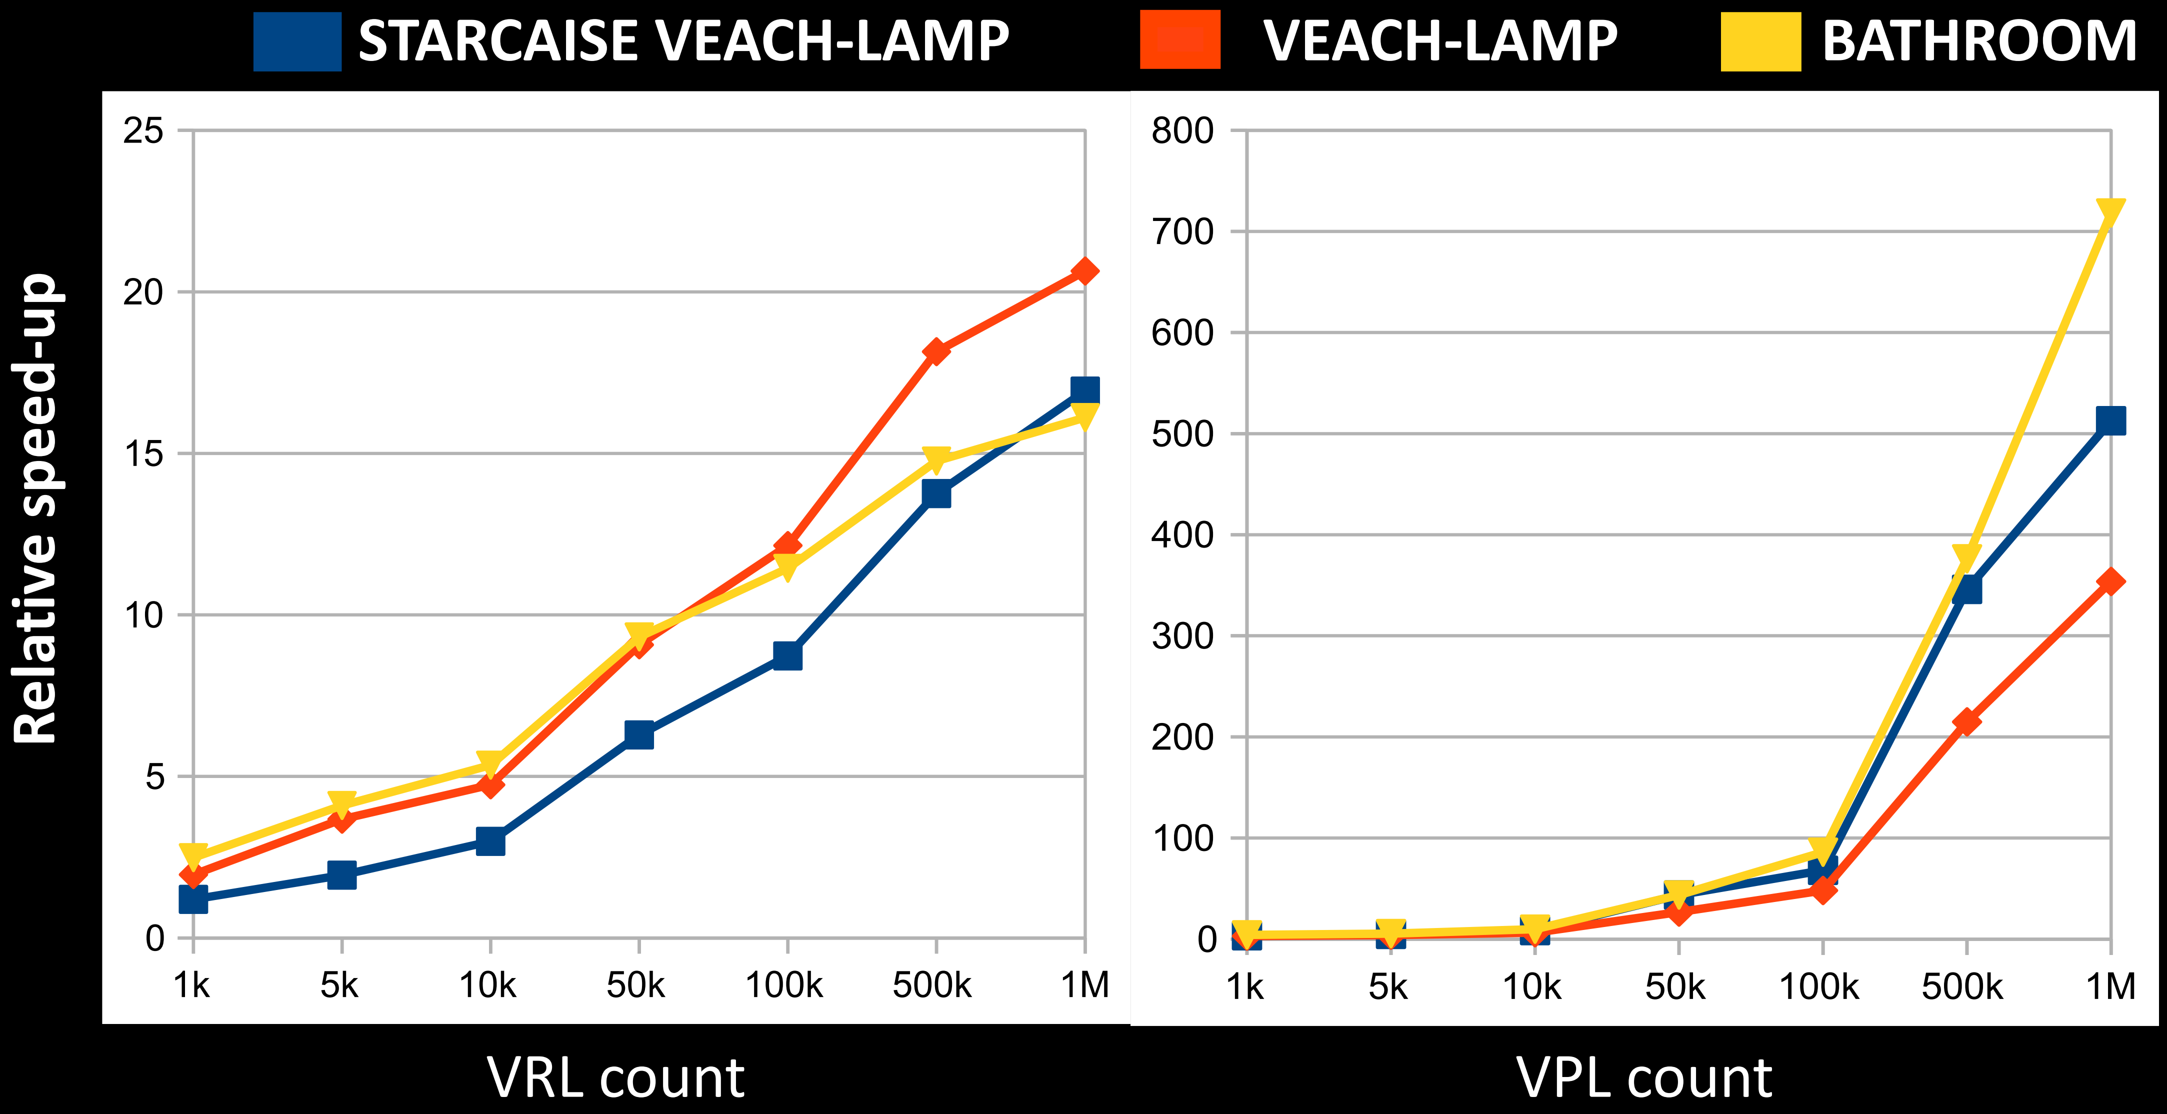
STARCAISE VEACH-LAMP
VEACH-LAMP
BATHROOM
Relative speed-up
25
20
15
10
5
0
1k
5k
10k
50k
100k
500k
1M
VRL count
800
700
600
500
400
300
200
100
0
1k
5k
10k
50k
100k
500k
1M
VPL count

Summary

Contributions:

- New bound for VRL cluster
- Efficient tree construction
- X10 Speedup

Questions ?