

Robust Adaptive Photon Tracing using Photon Path Visibility

Toshiya Hachisuka^{1,2} Henrik Wann Jensen²

¹Aarhus University ²UC San Diego

ACM Transaction on Graphics, Volume 30, Issue 5, 2011









Light source





Window



Eye







Invisible paths = wasted computation


Metropolis Light Transport



Metropolis Light Transport



Reference




Metropolis Light Transport [100 min]

Specular-Diffuse-Specular

- Fundamental limitation of MC integration [Veach 98]
 - Missing “reflections of caustics from a point light”
 - Applicable to all local path sampling methods
 - True for small area lights in practice [Hachisuka et al. 08]


Specular-Diffuse-Specular



Specular-Diffuse-Specular



Specular-Diffuse-Specular



Specular-Diffuse-Specular




Specular-Diffuse-Specular



Progressive Photon Mapping





Progressive Photon Mapping [100 min]



Metropolis Light Transport [100 min]

Inefficient Case



Progressive Photon Mapping [100 min]

Ideal

- Can we combine these two algorithms?
 - **MLT**: Efficient for difficult lighting scenarios
 - **PPM**: Robust to complex types of light paths

Contributions

MLT + PPM

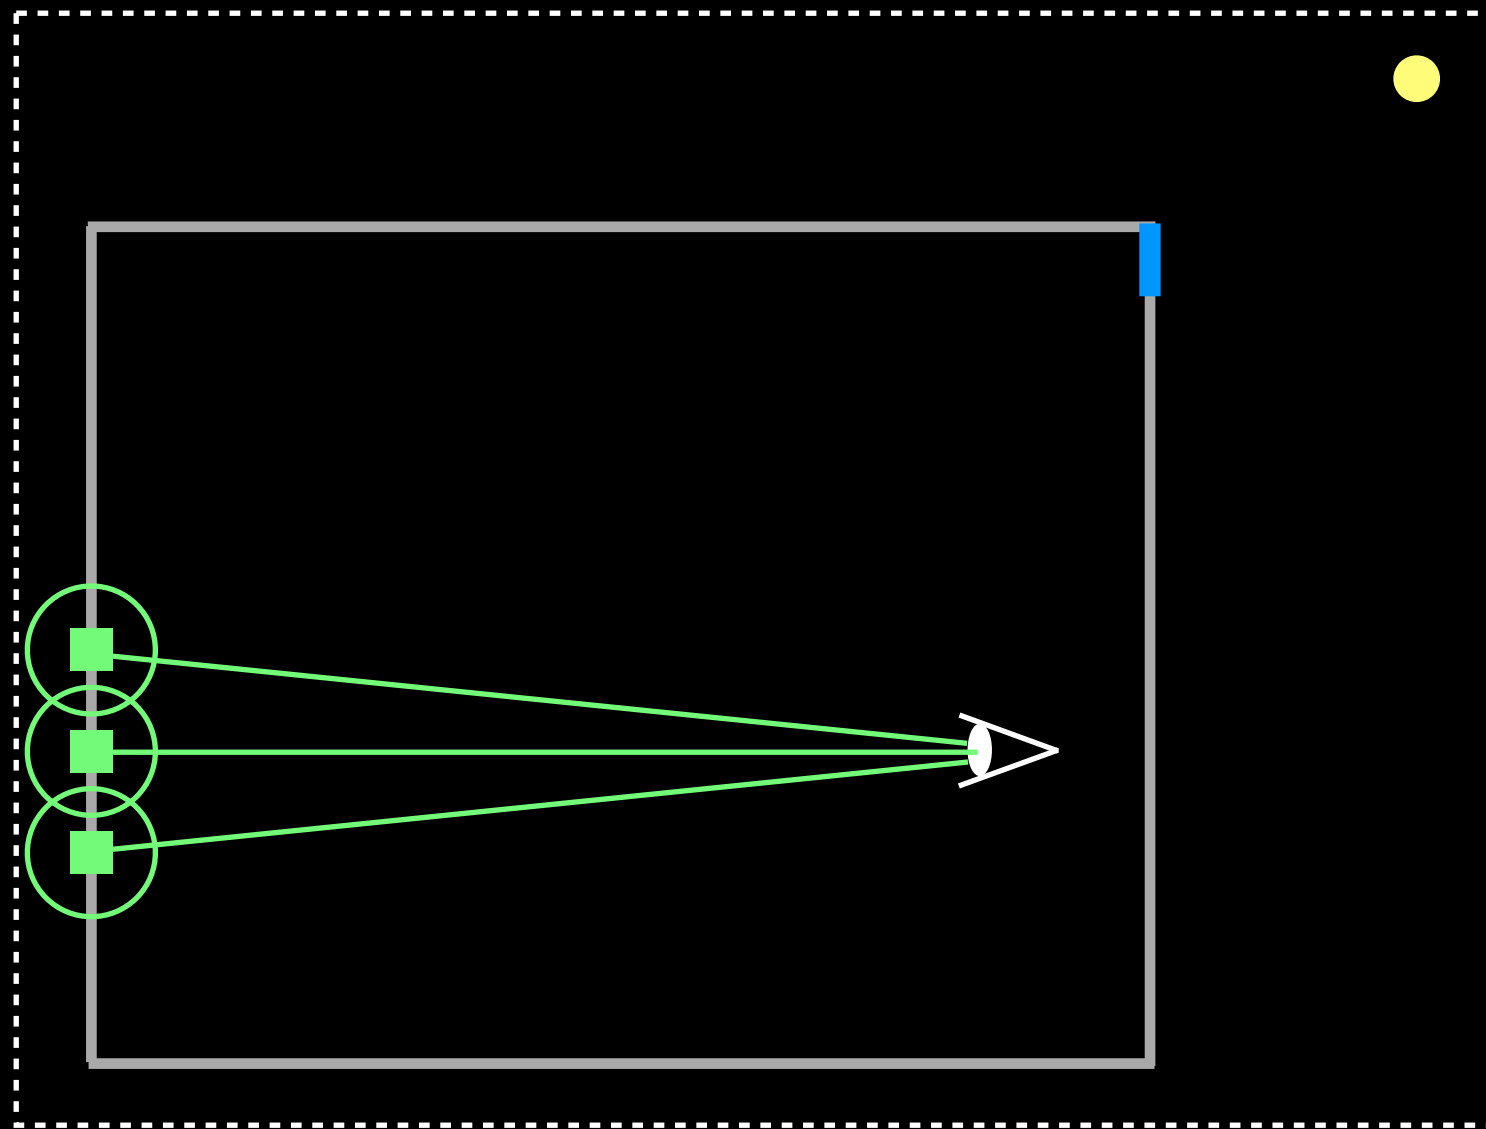
Contributions

MLT + PPM

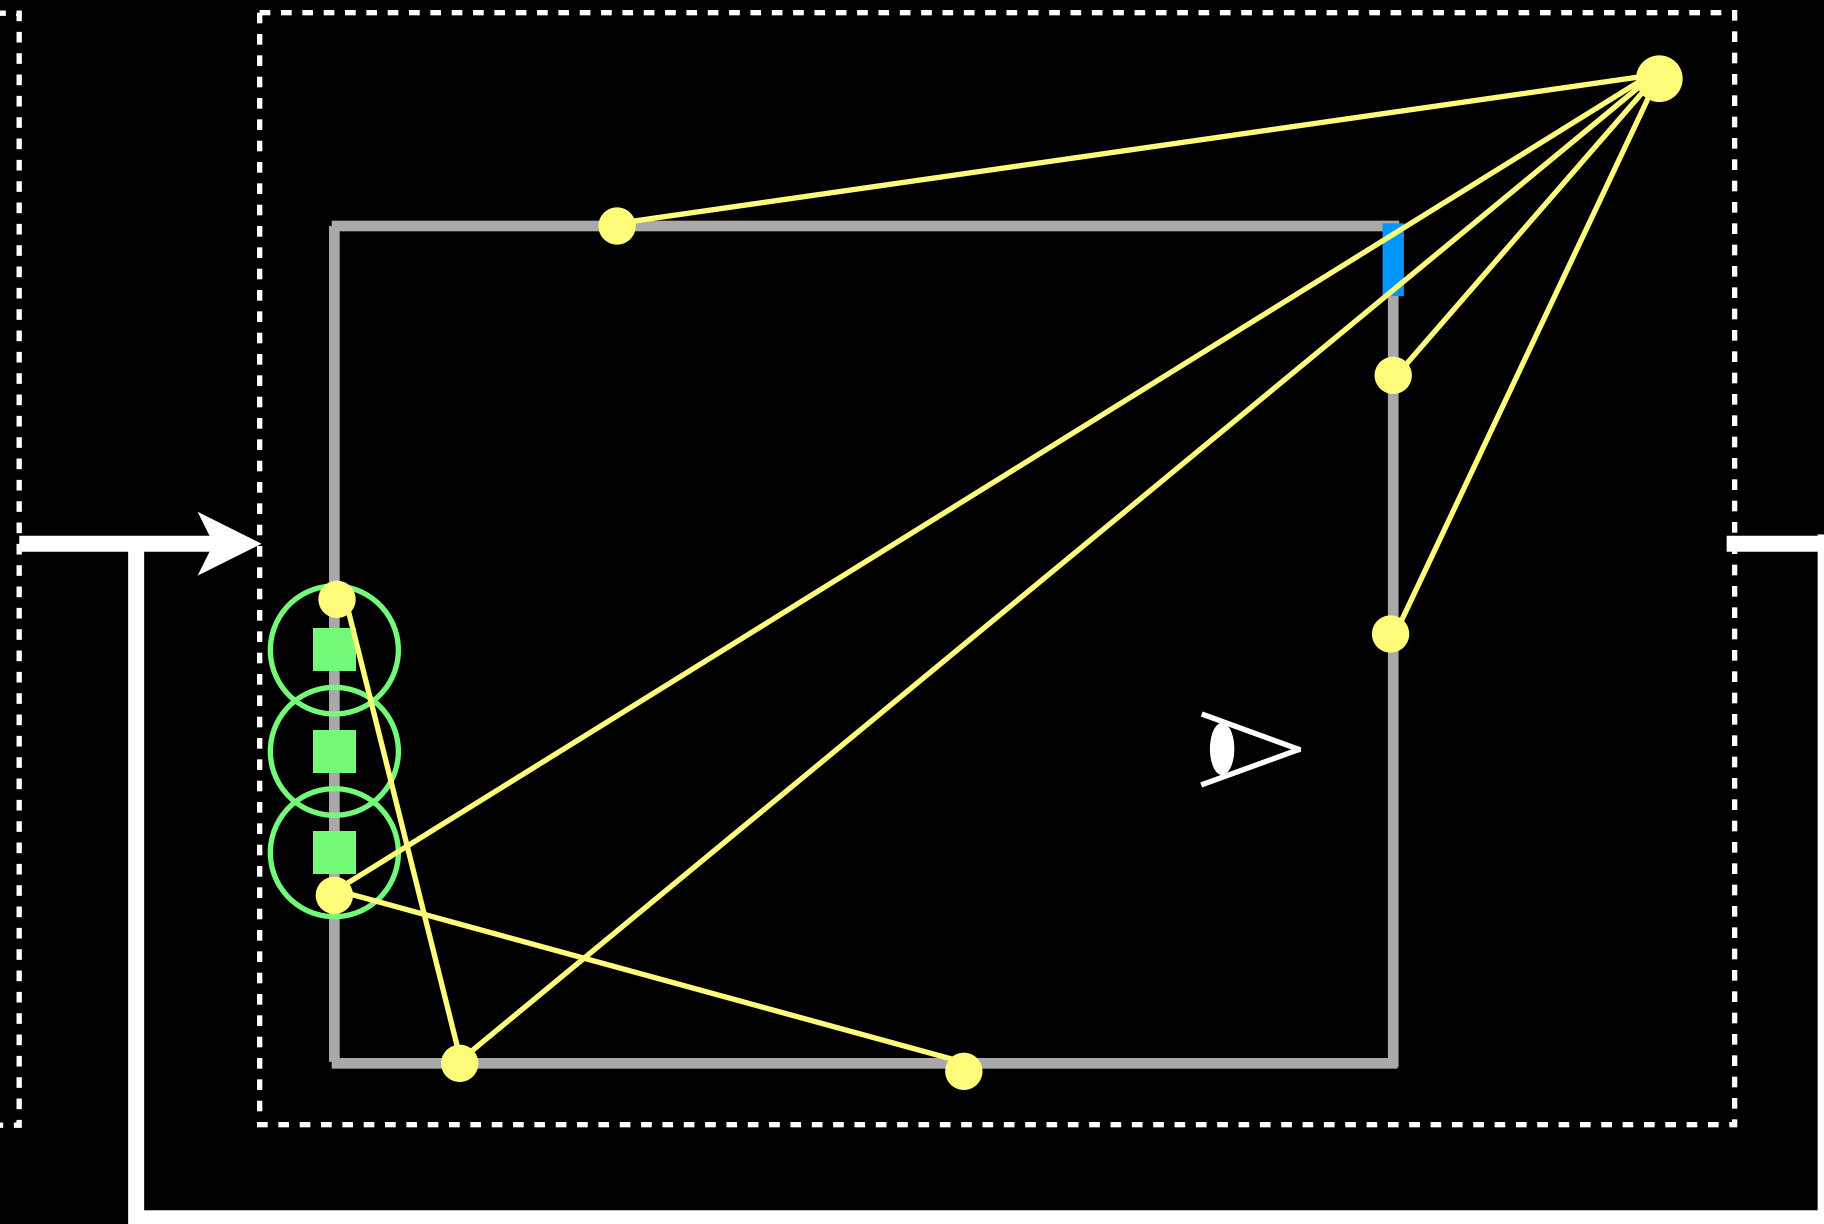
Simple, fast, general, and easy to use

Method

Progressive Photon Mapping




Eye pass




Photon pass

Eye Pass




Eye Pass



Eye Pass



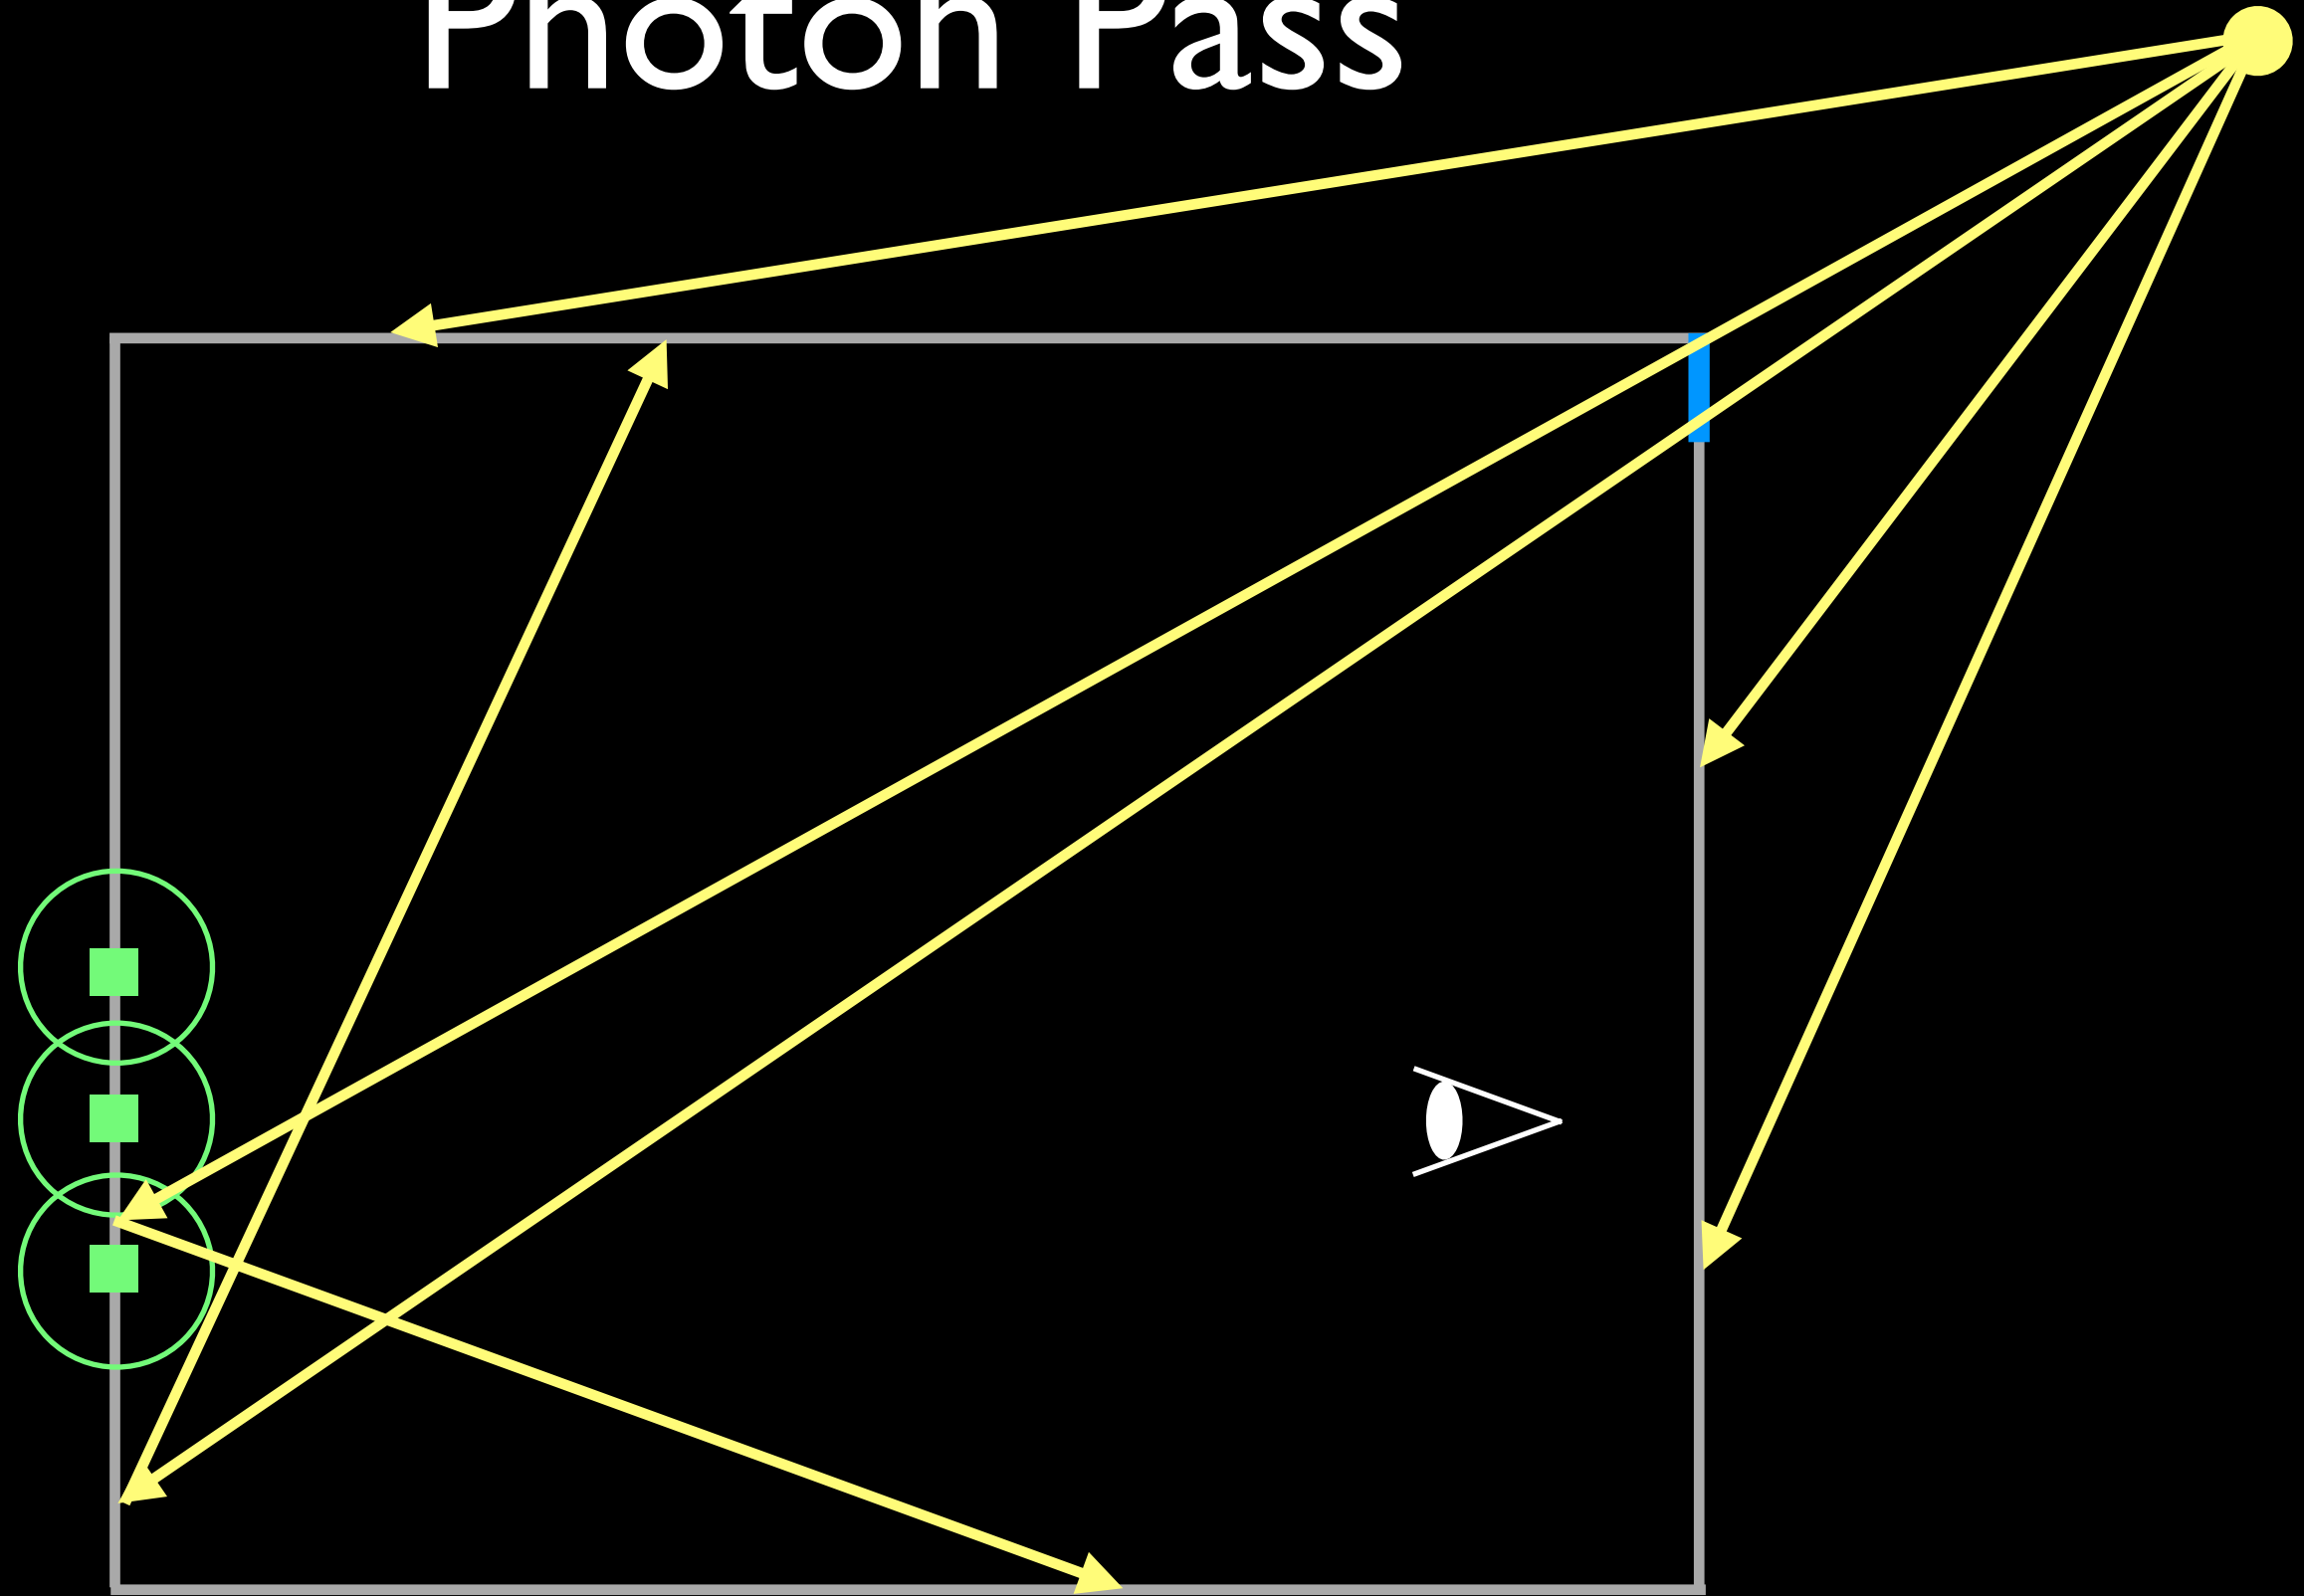
Visible points




Photon Pass




Photon Pass



Photon Pass

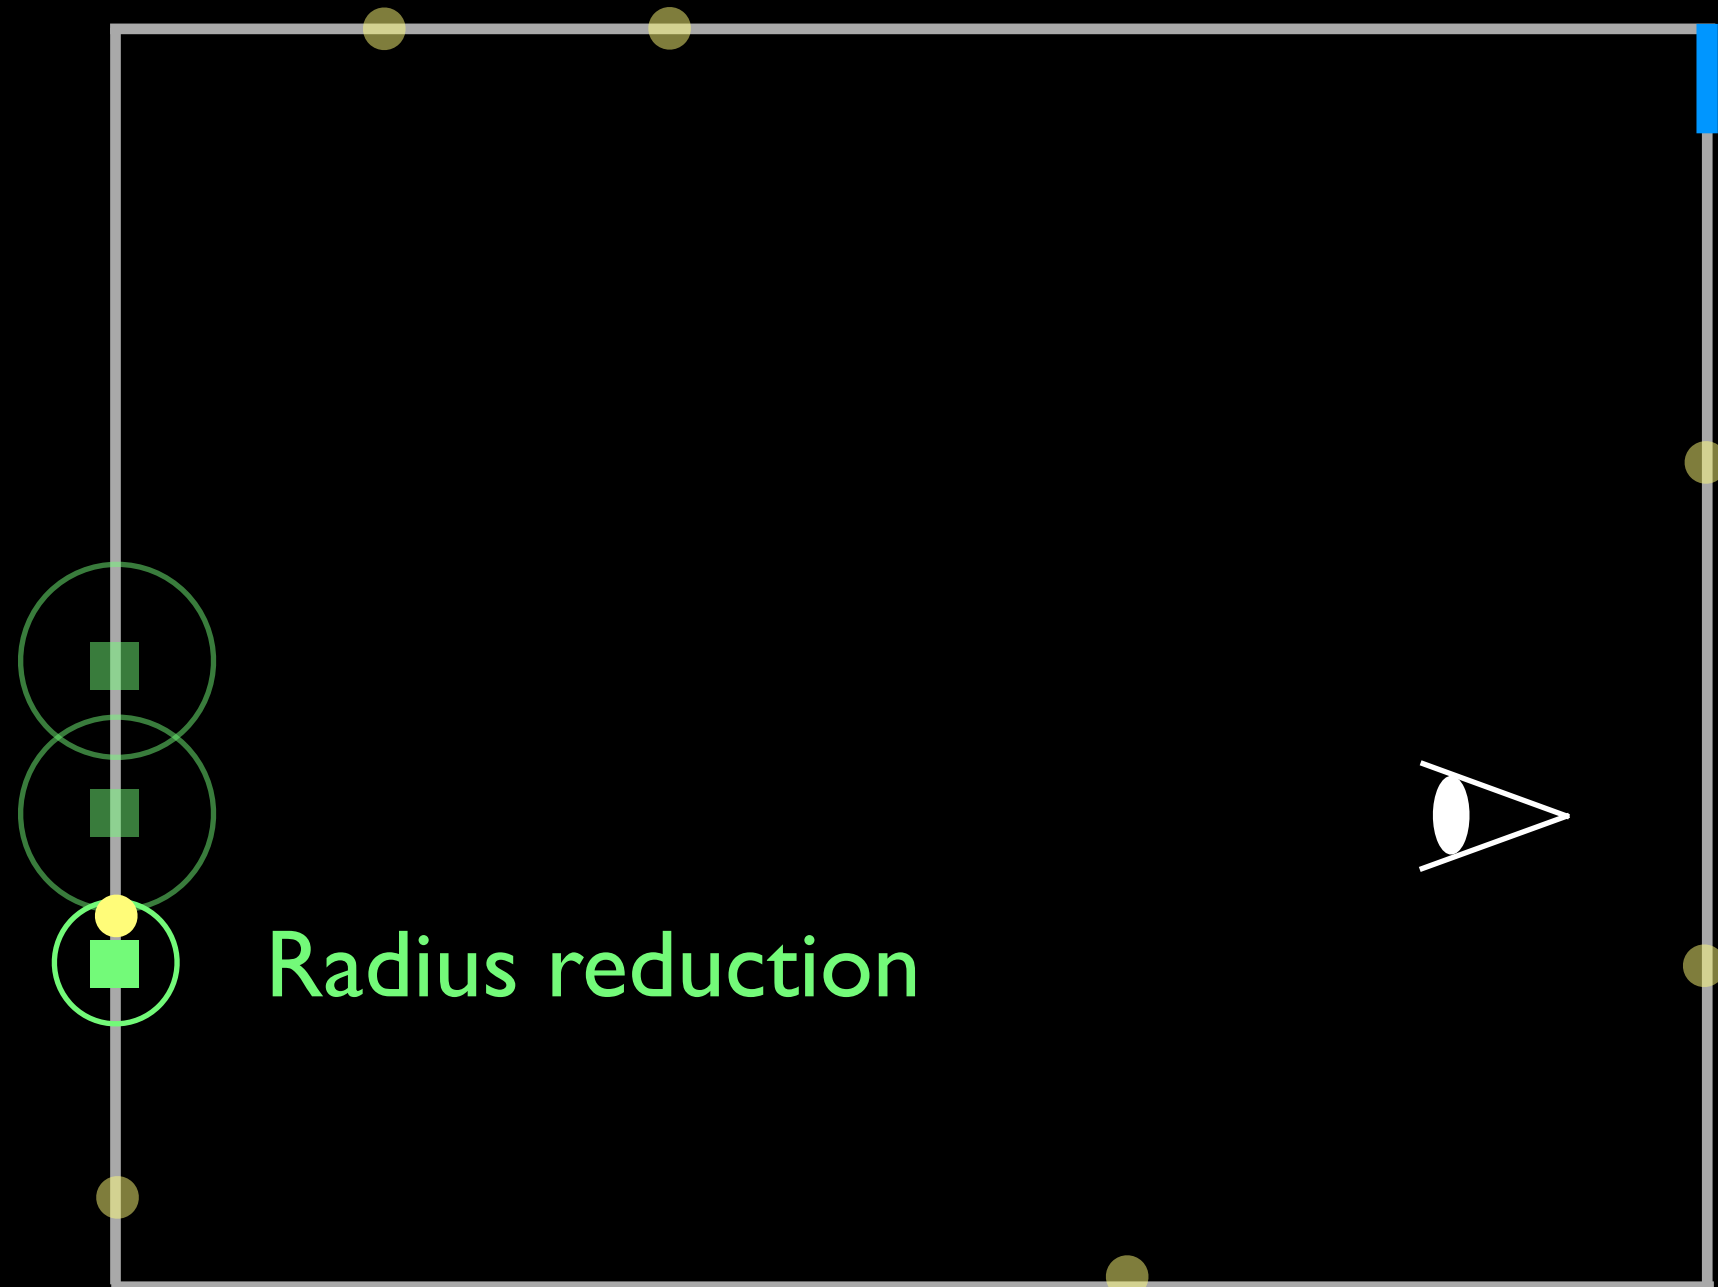


Photon Pass



Range query

Photon Pass

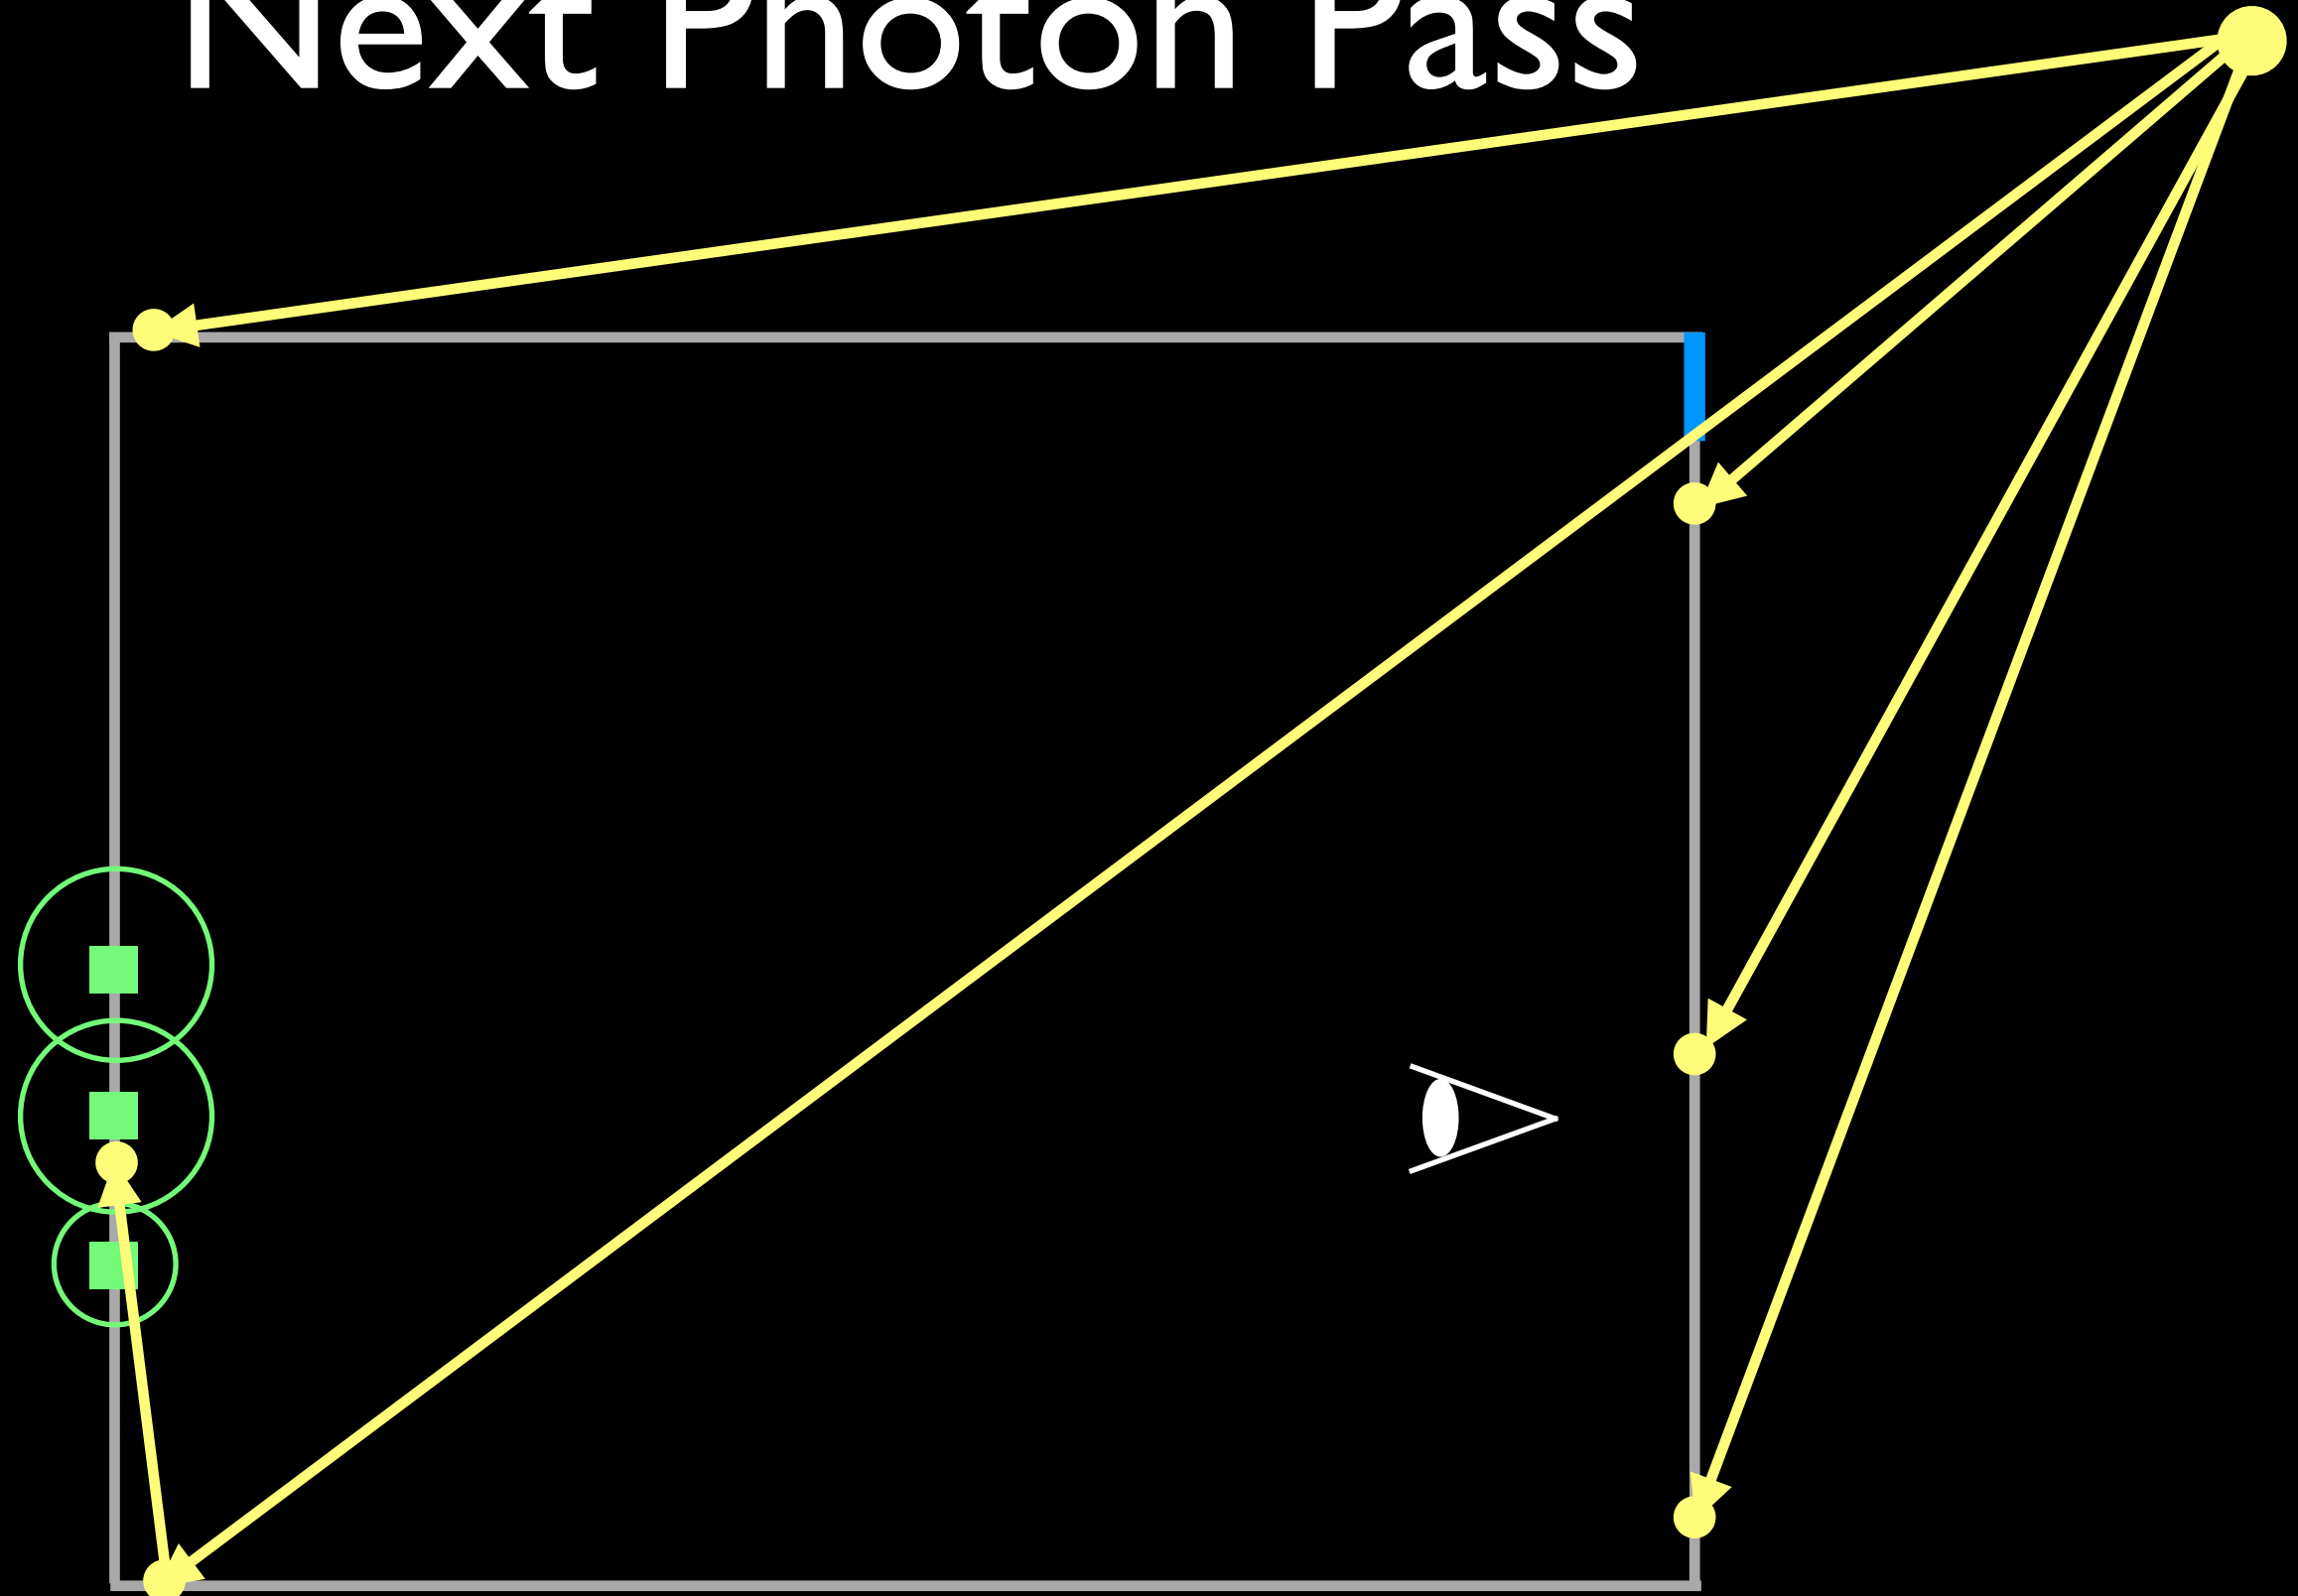


Radius reduction


Next Photon Pass




Next Photon Pass



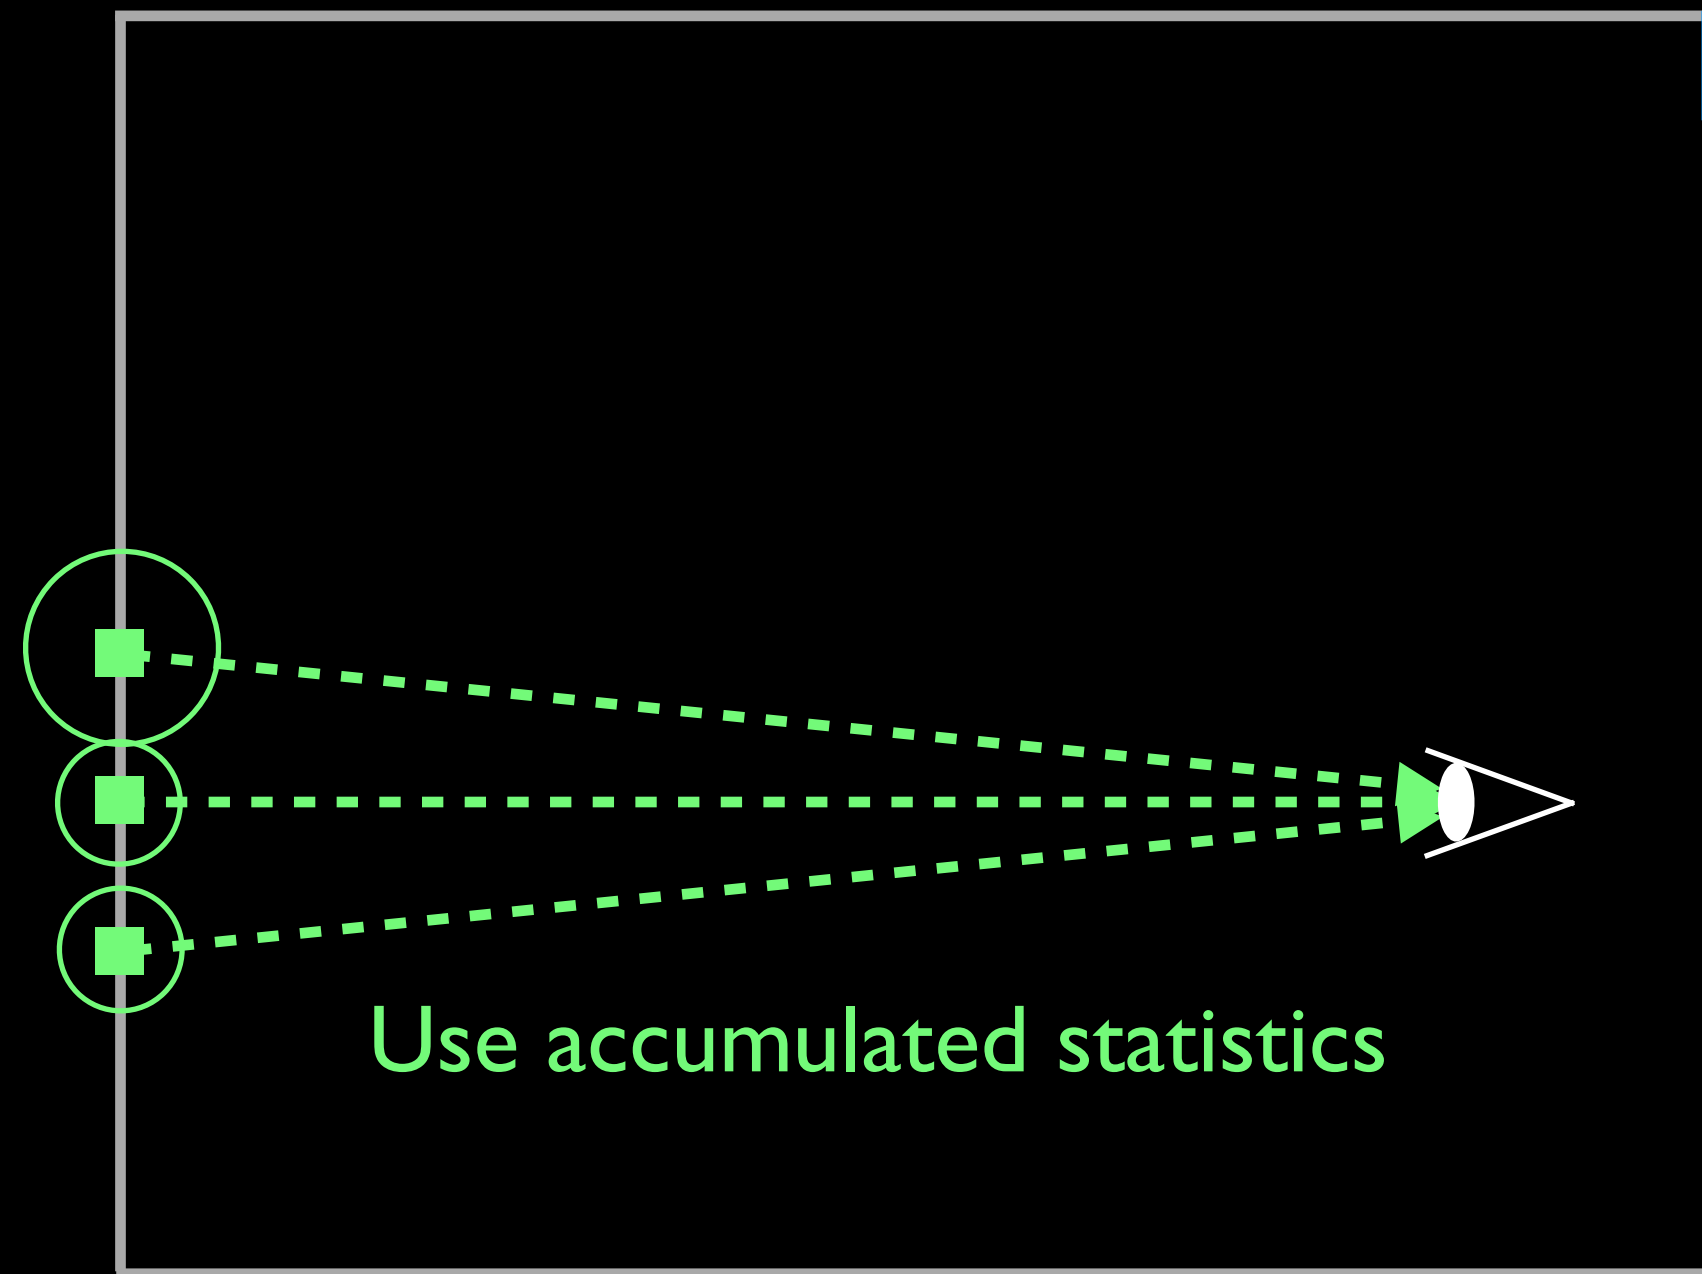
Next Photon Pass



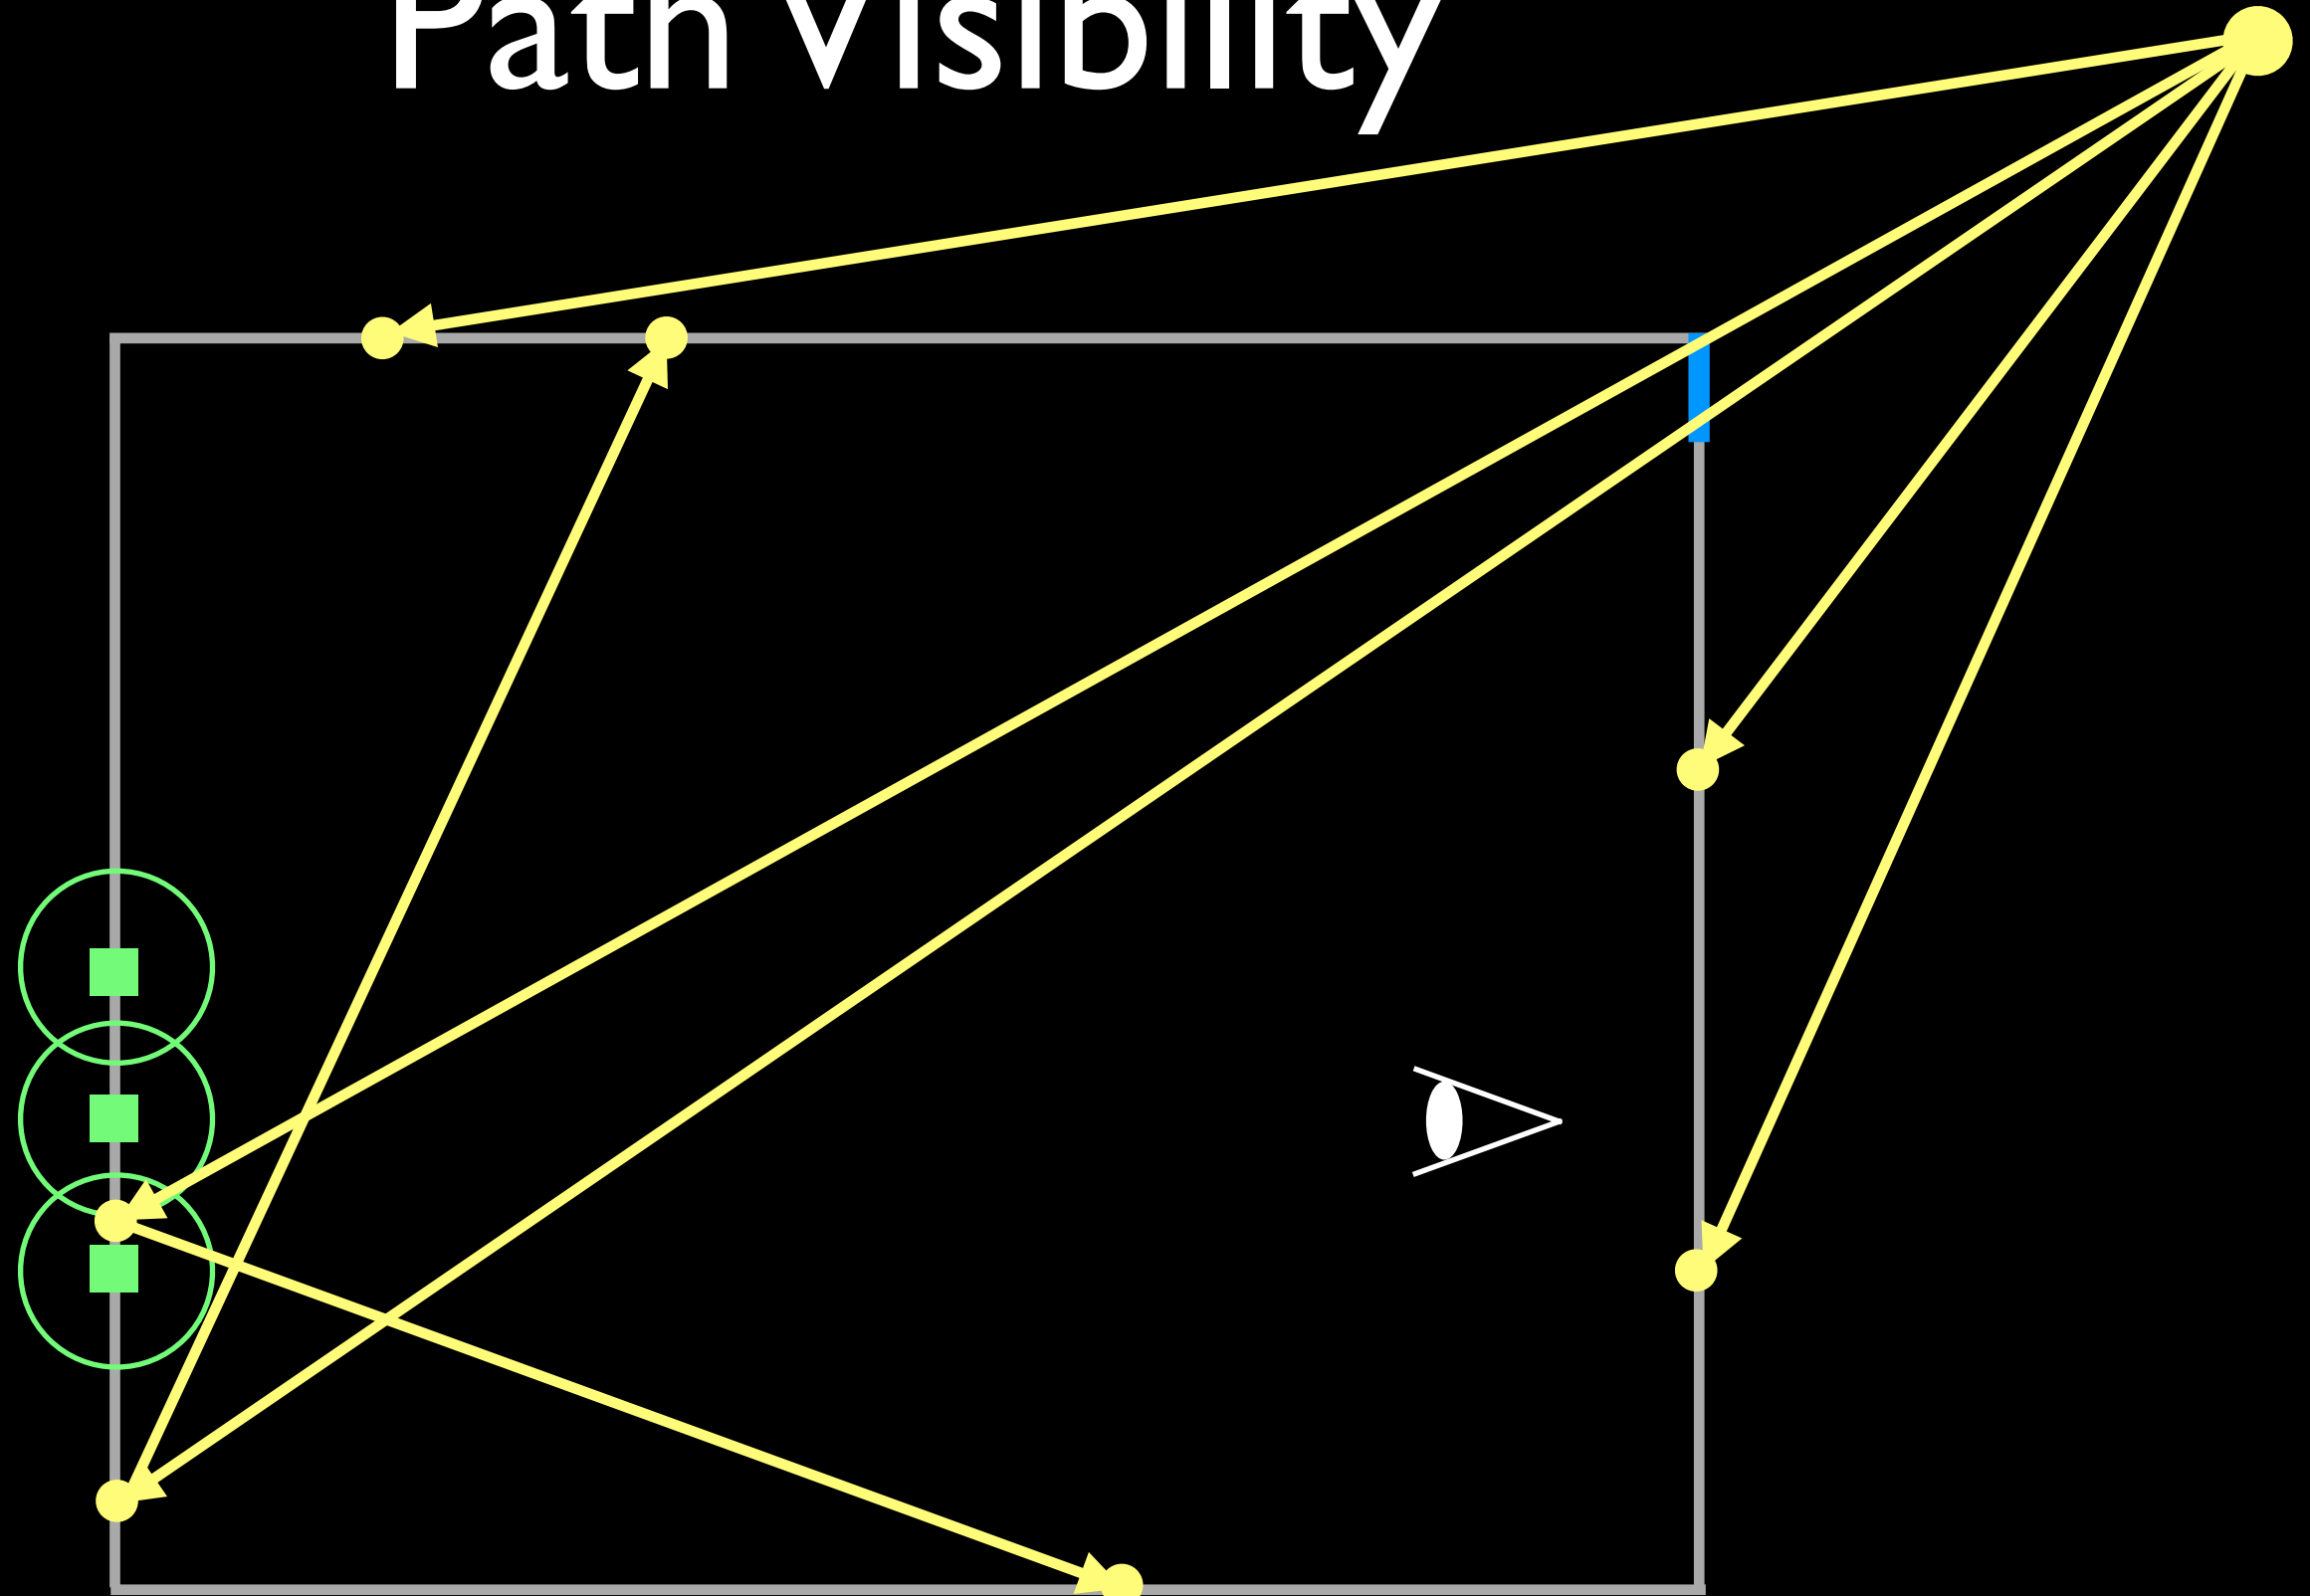
Next Photon Pass



Rendering



Path Visibility




Path Visibility




Key Observation

- We can determine whether a **photon path** is visible or not
 - Because PPM stores **visible points** from the eye
 - Contributed to at least one **visible point** = visible

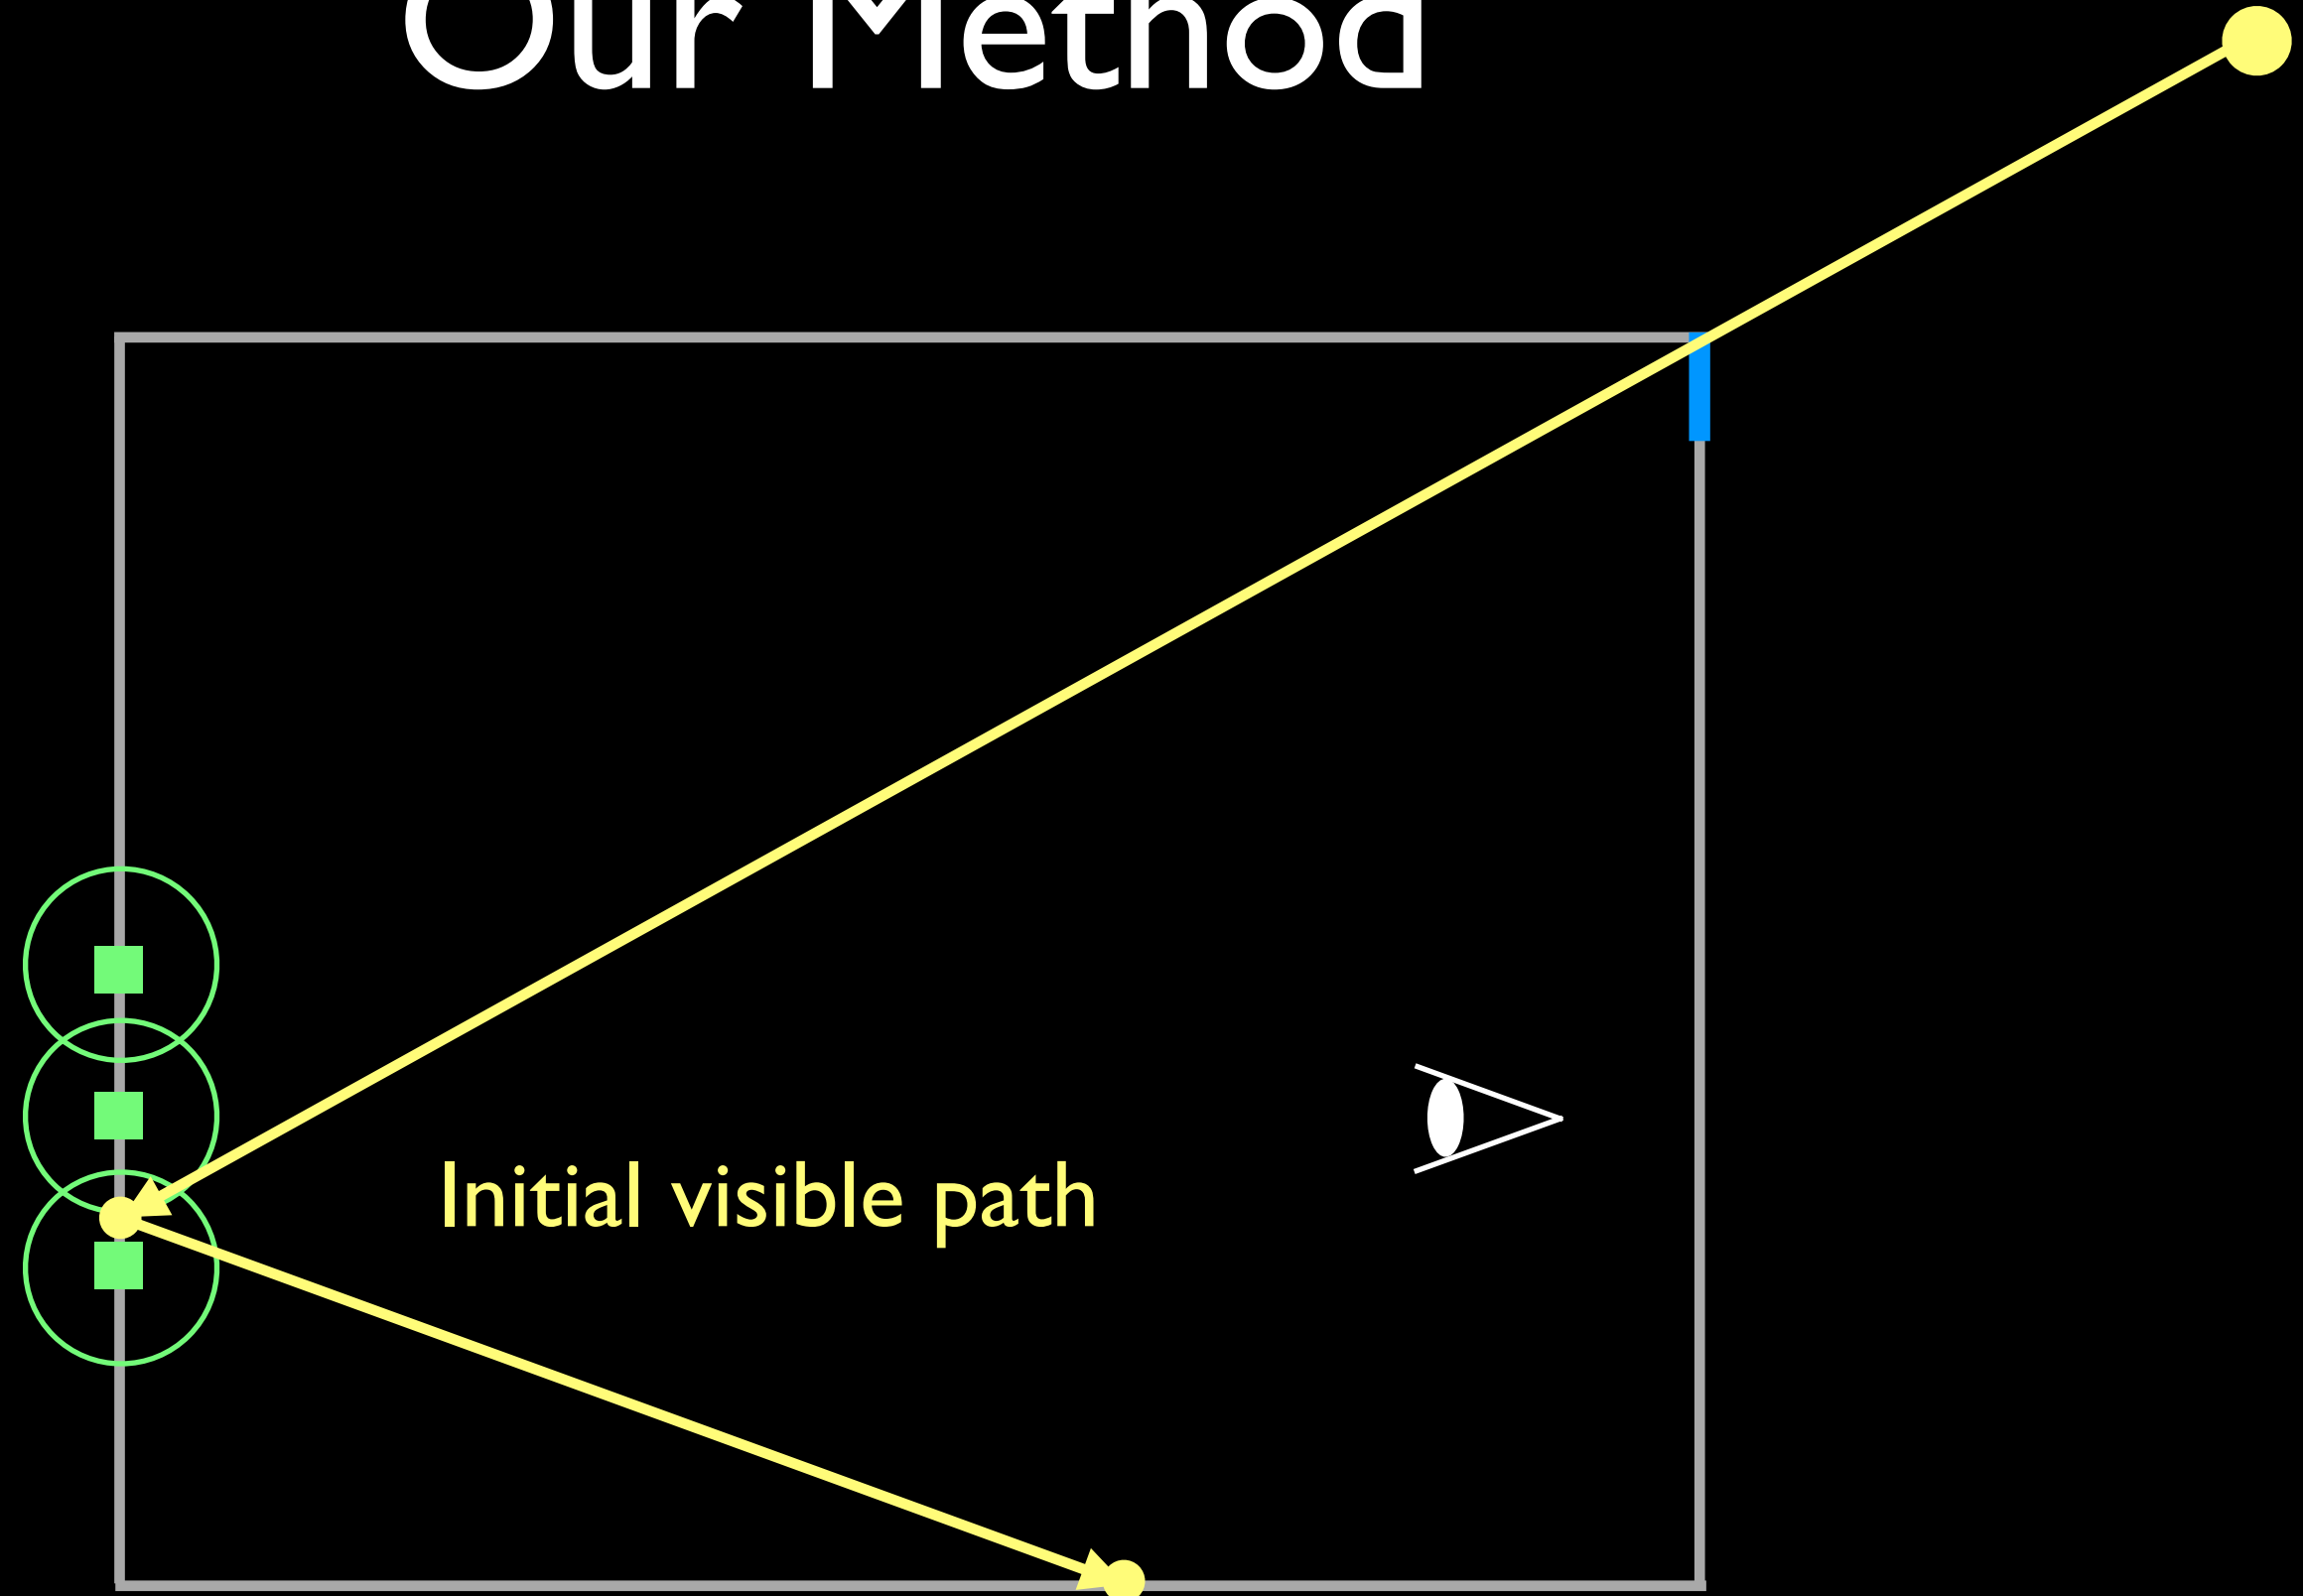
Path Visibility



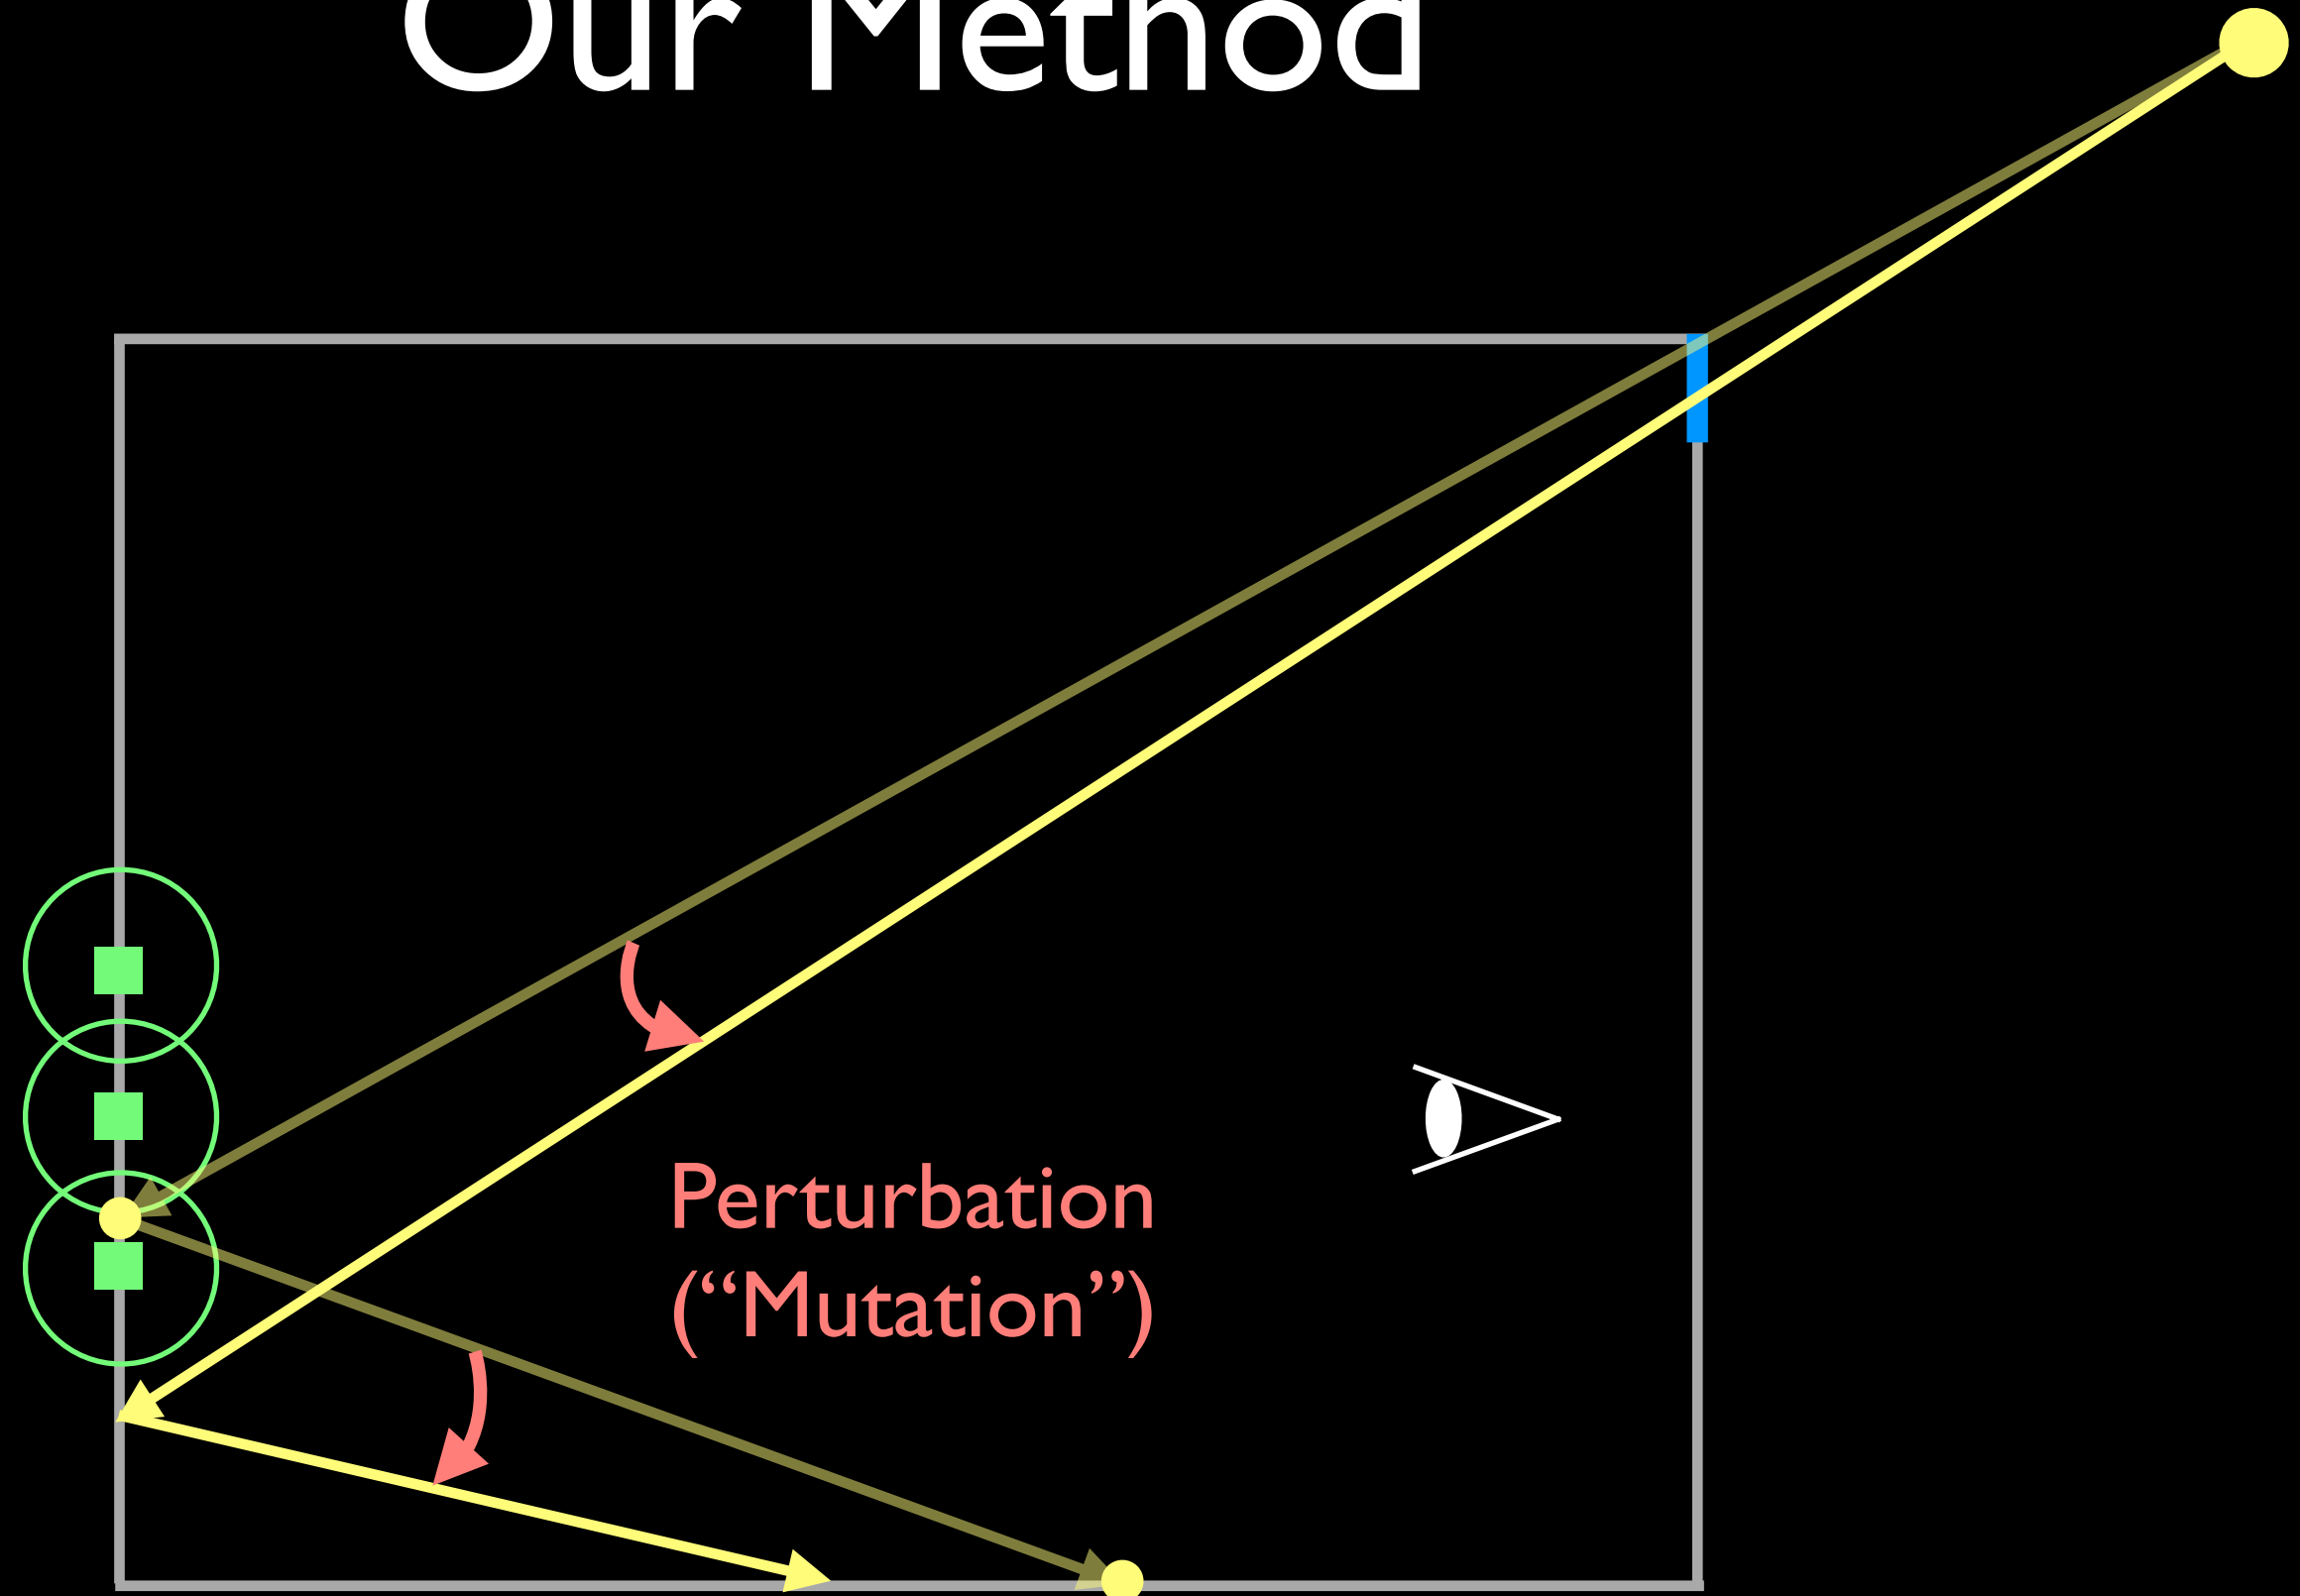
Our Method



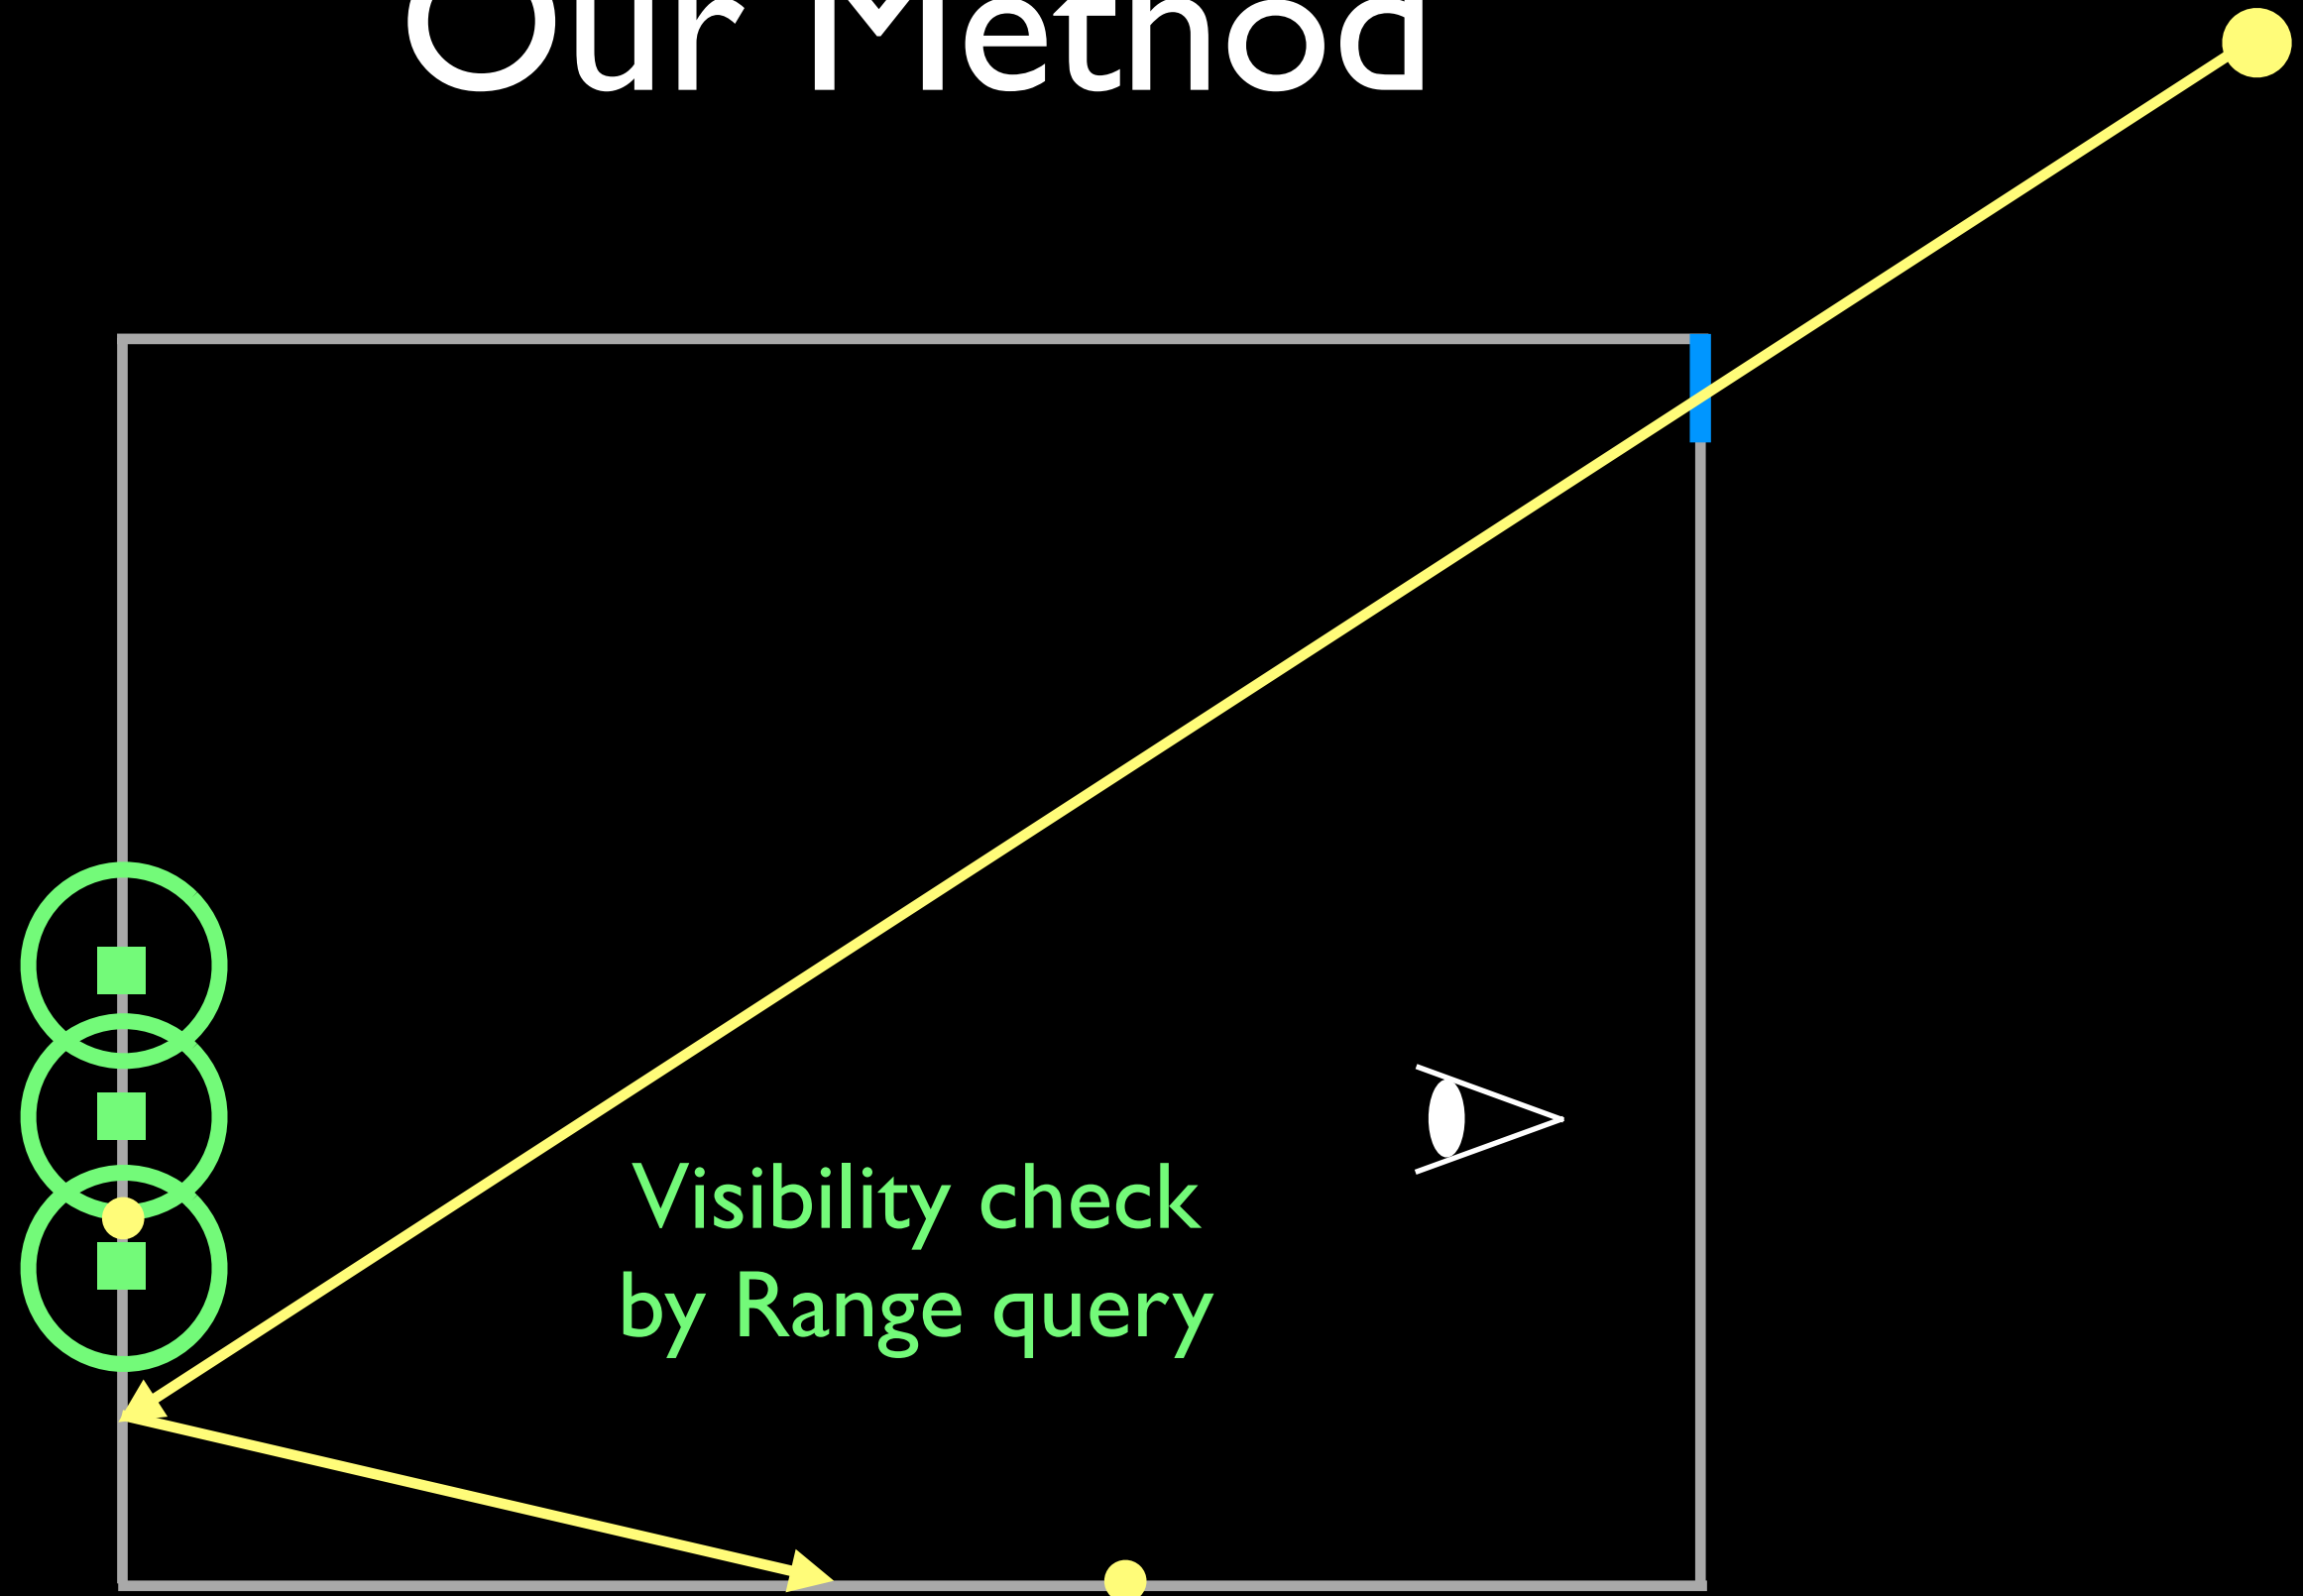
Our Method



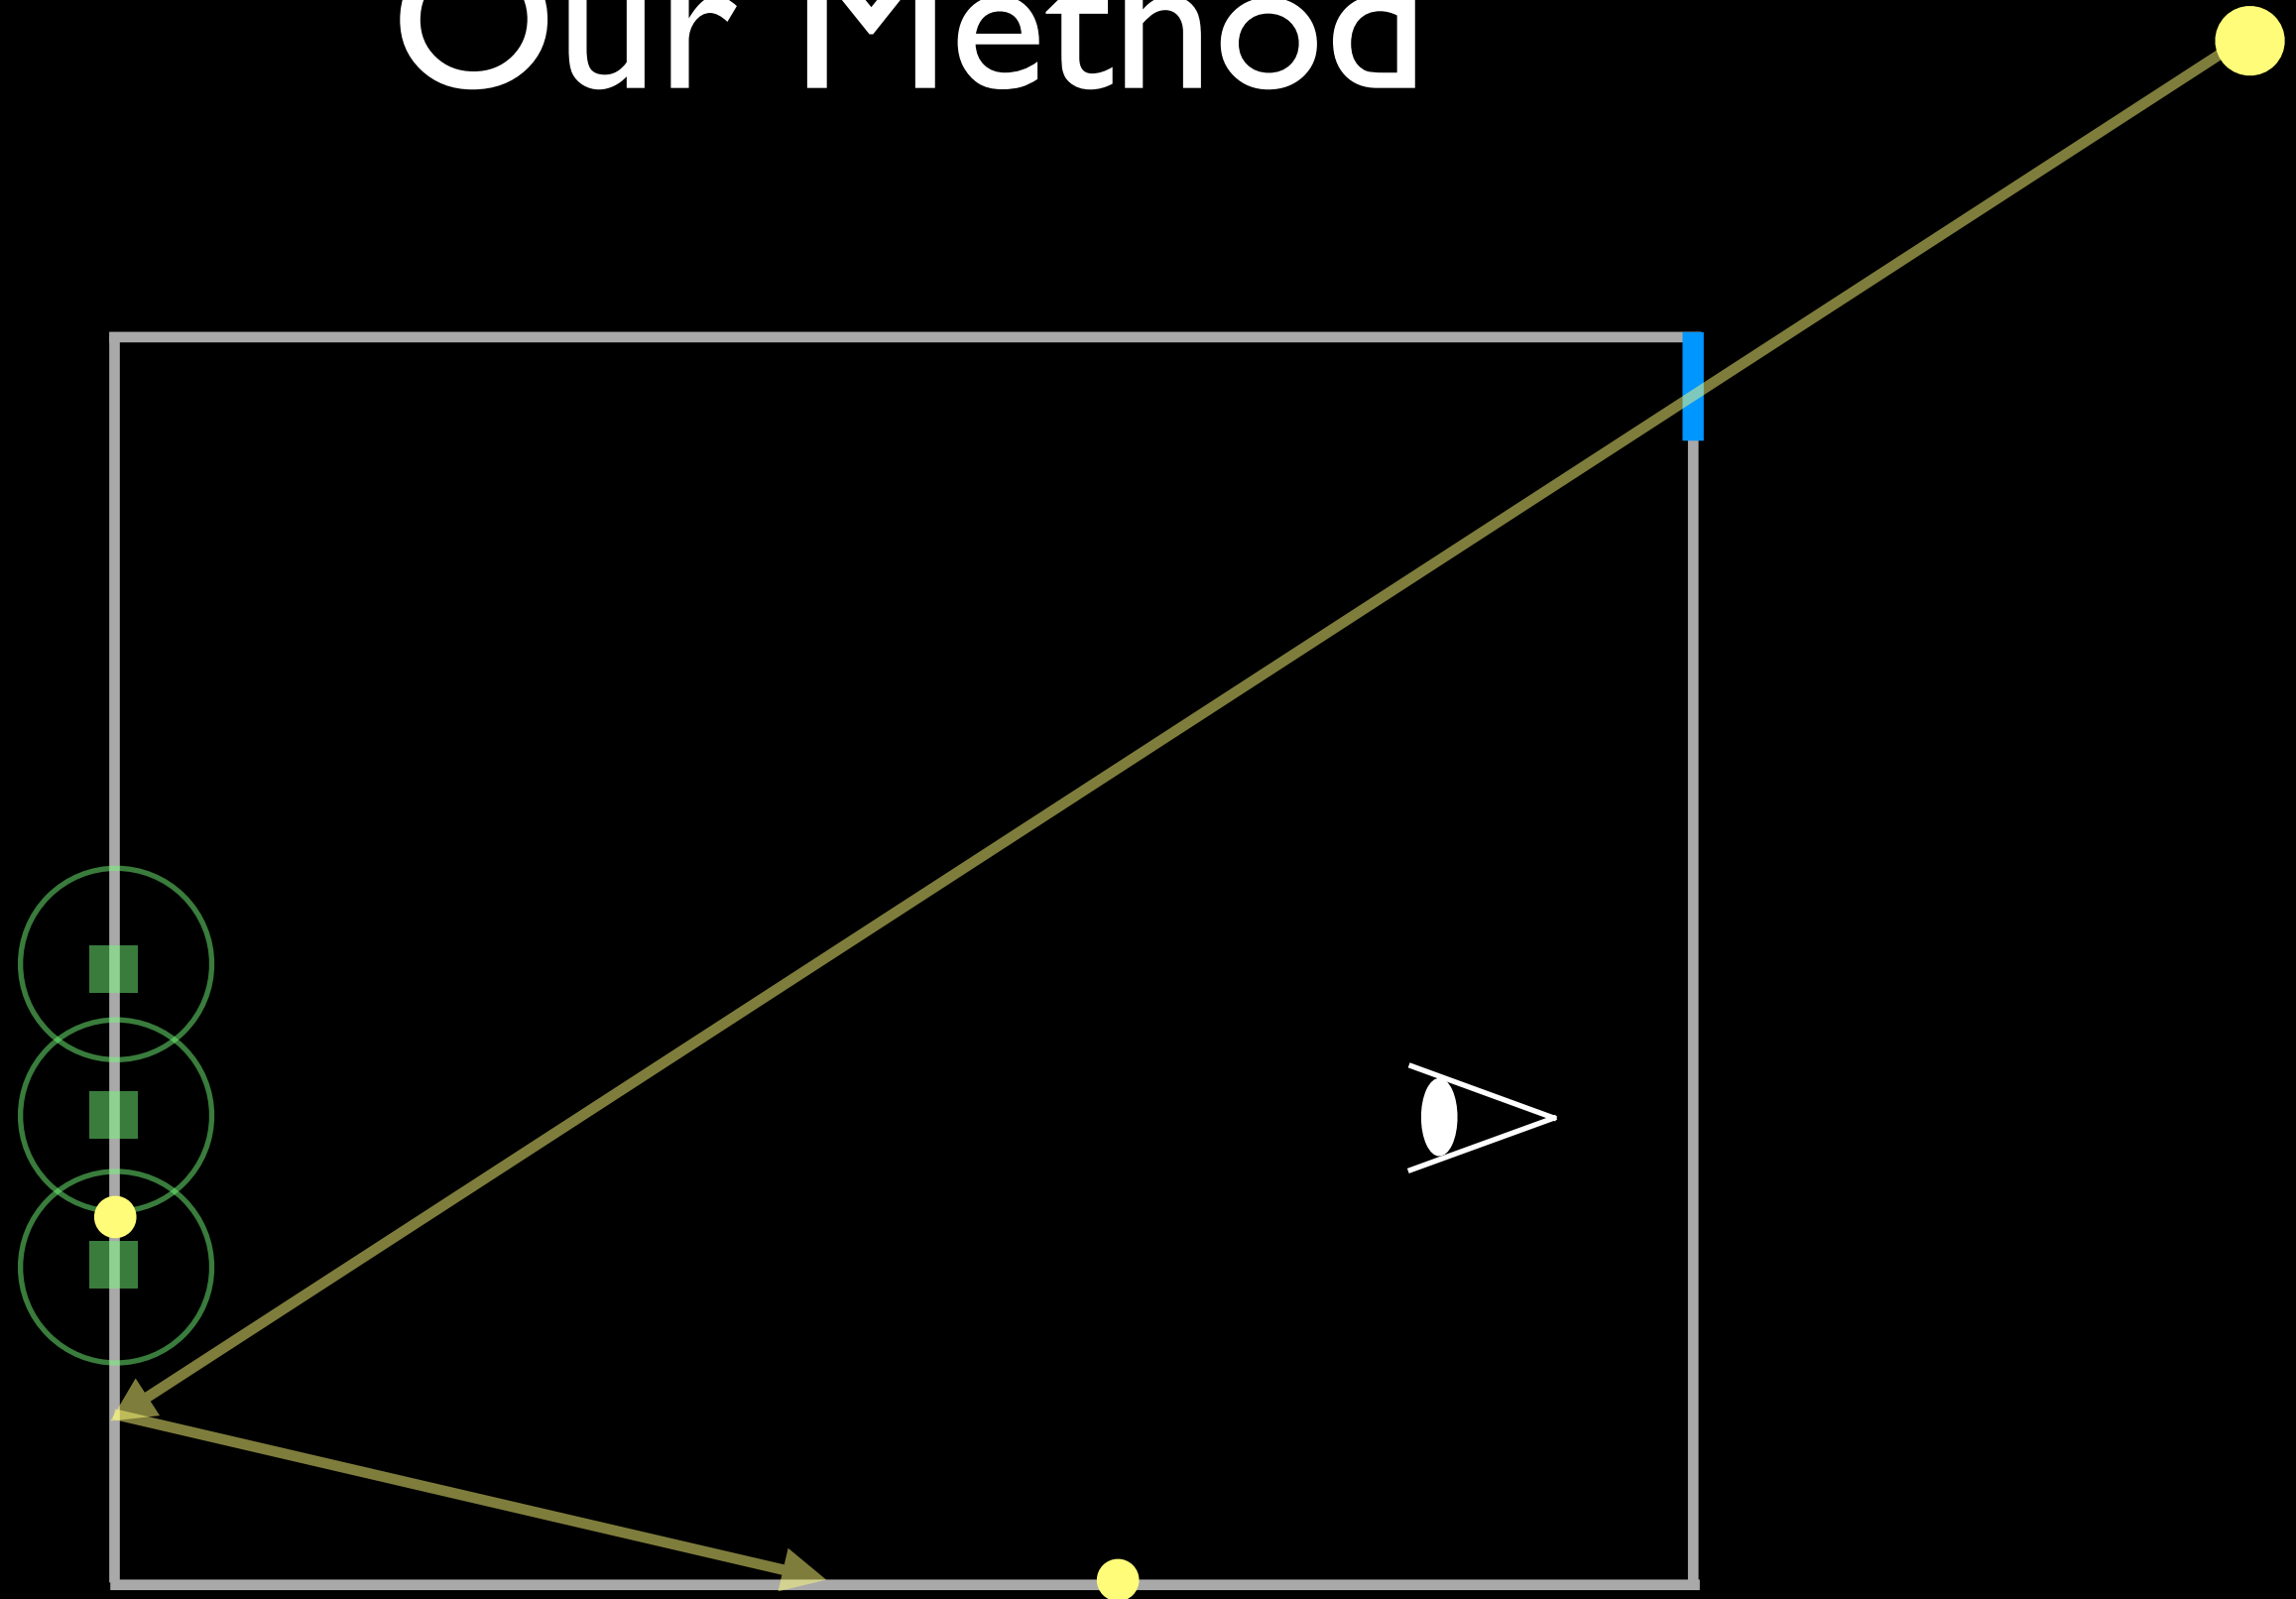
Our Method



Our Method



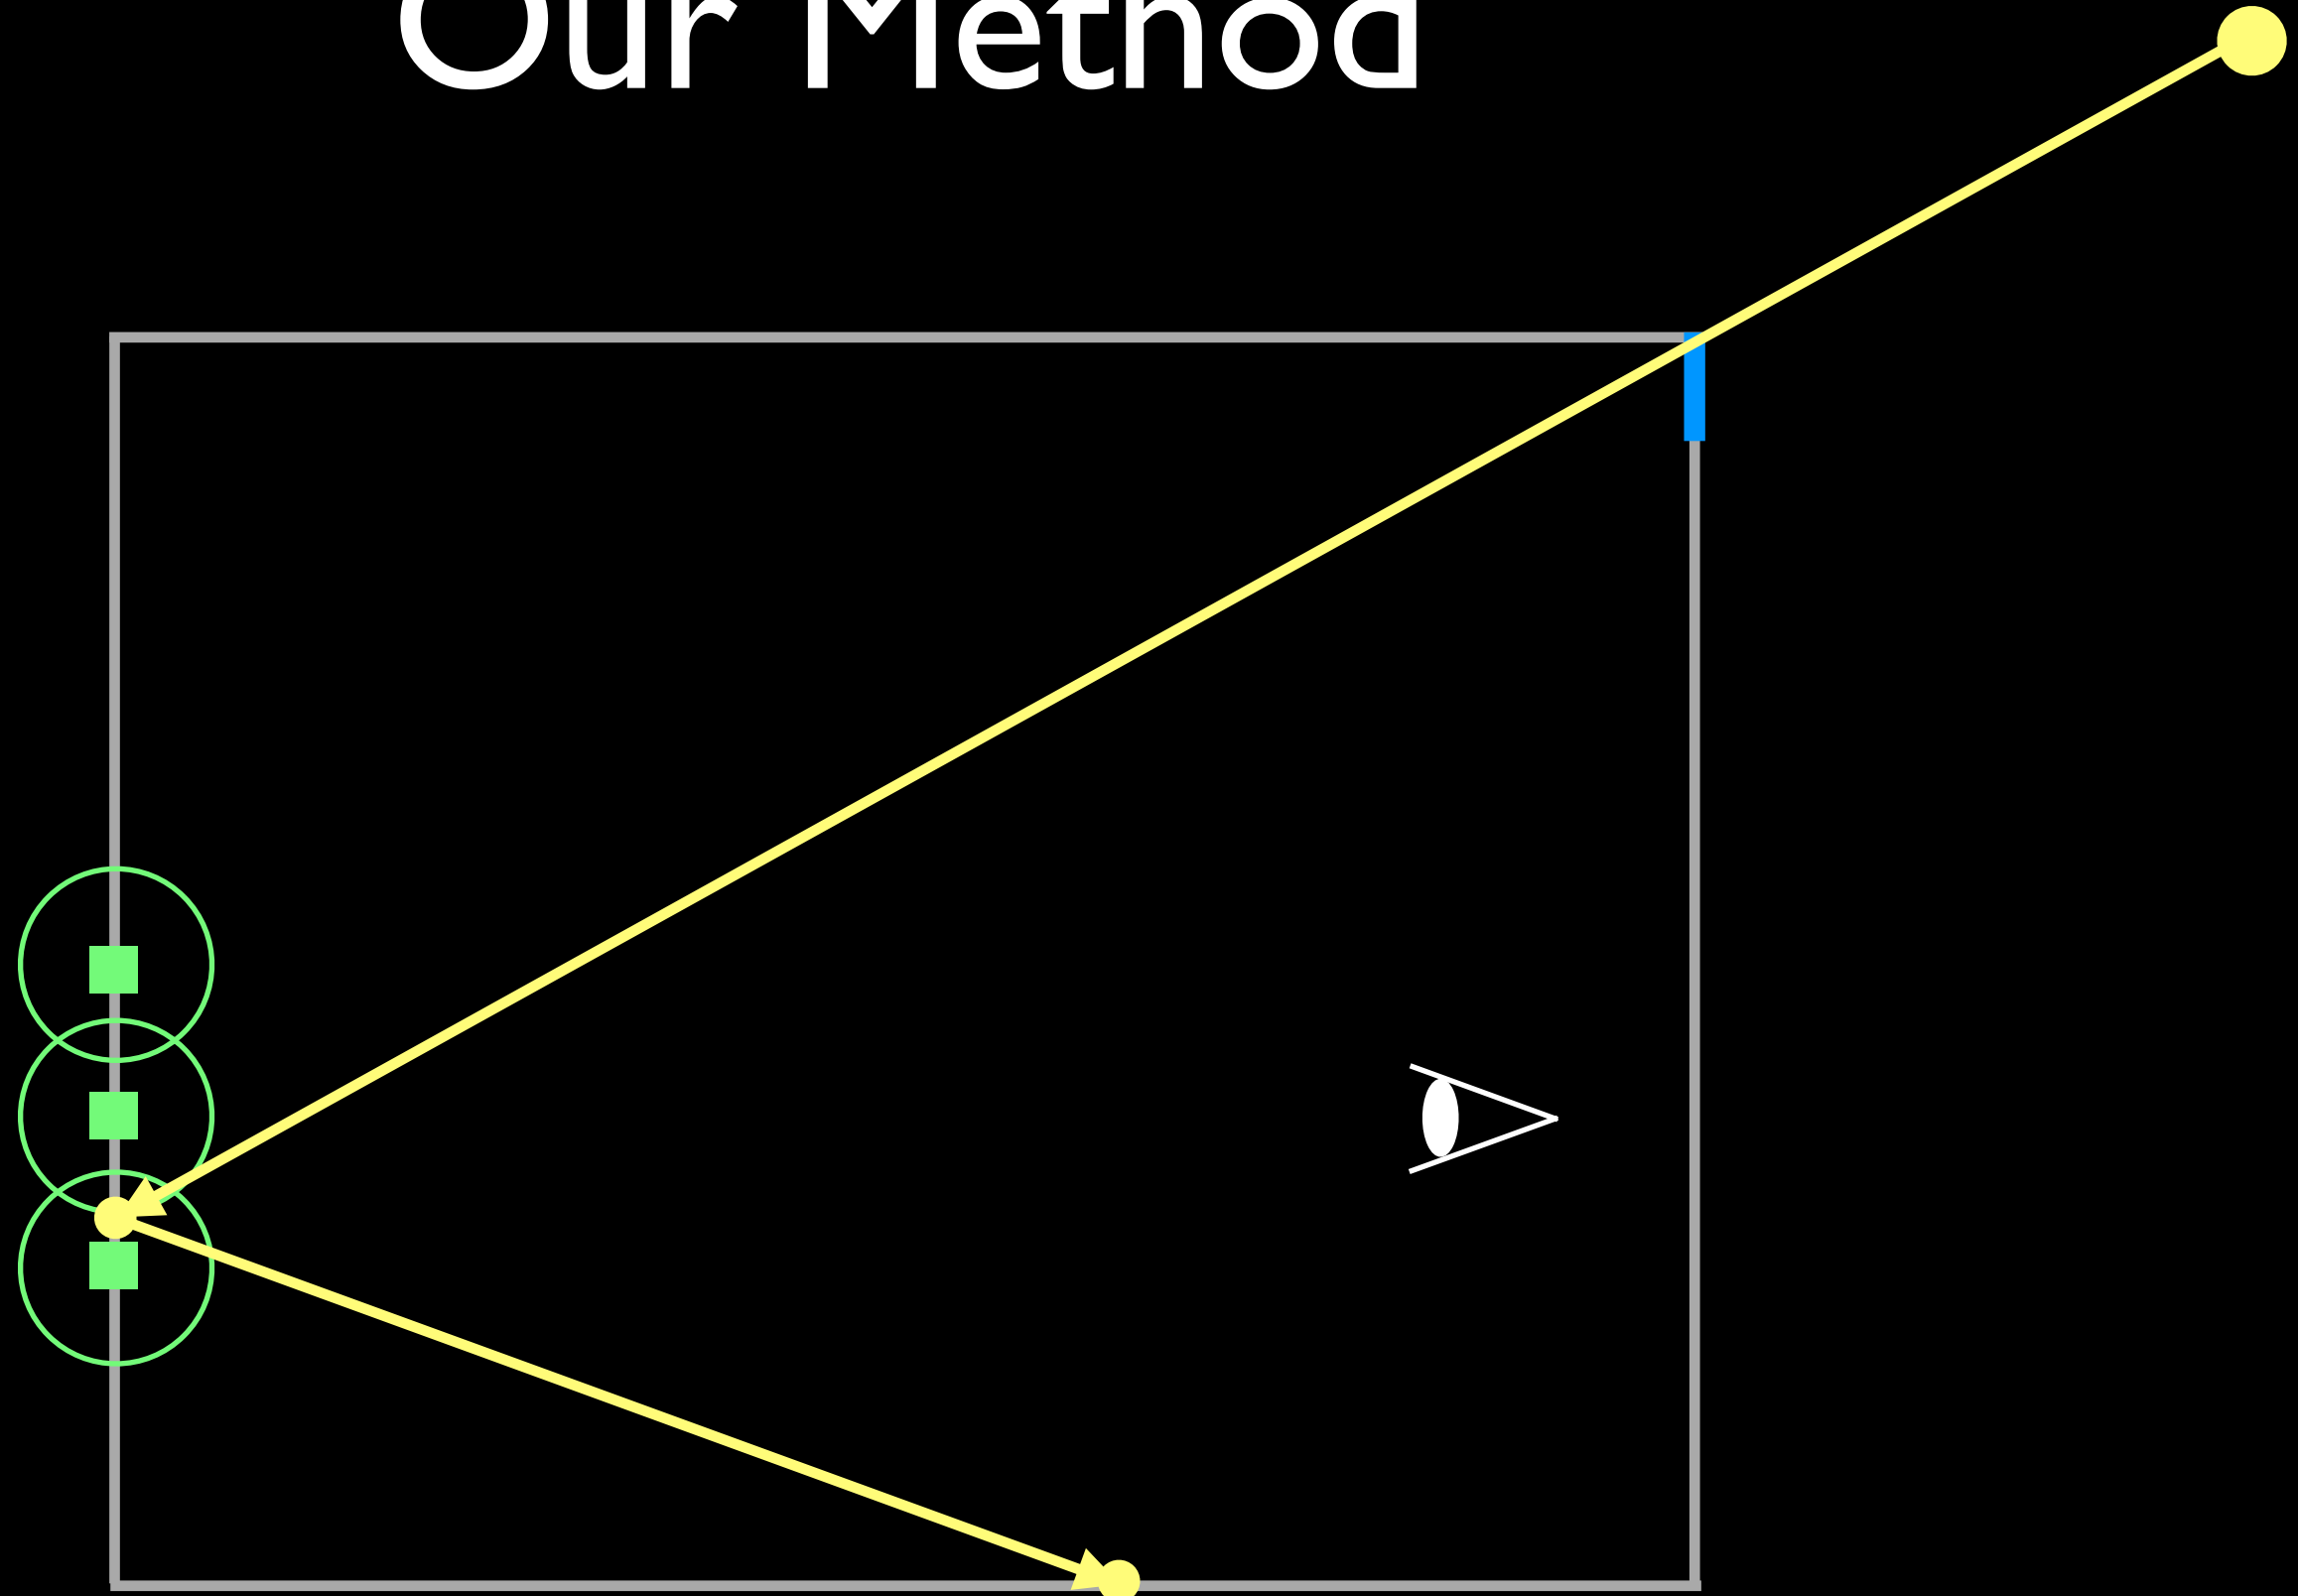
Our Method



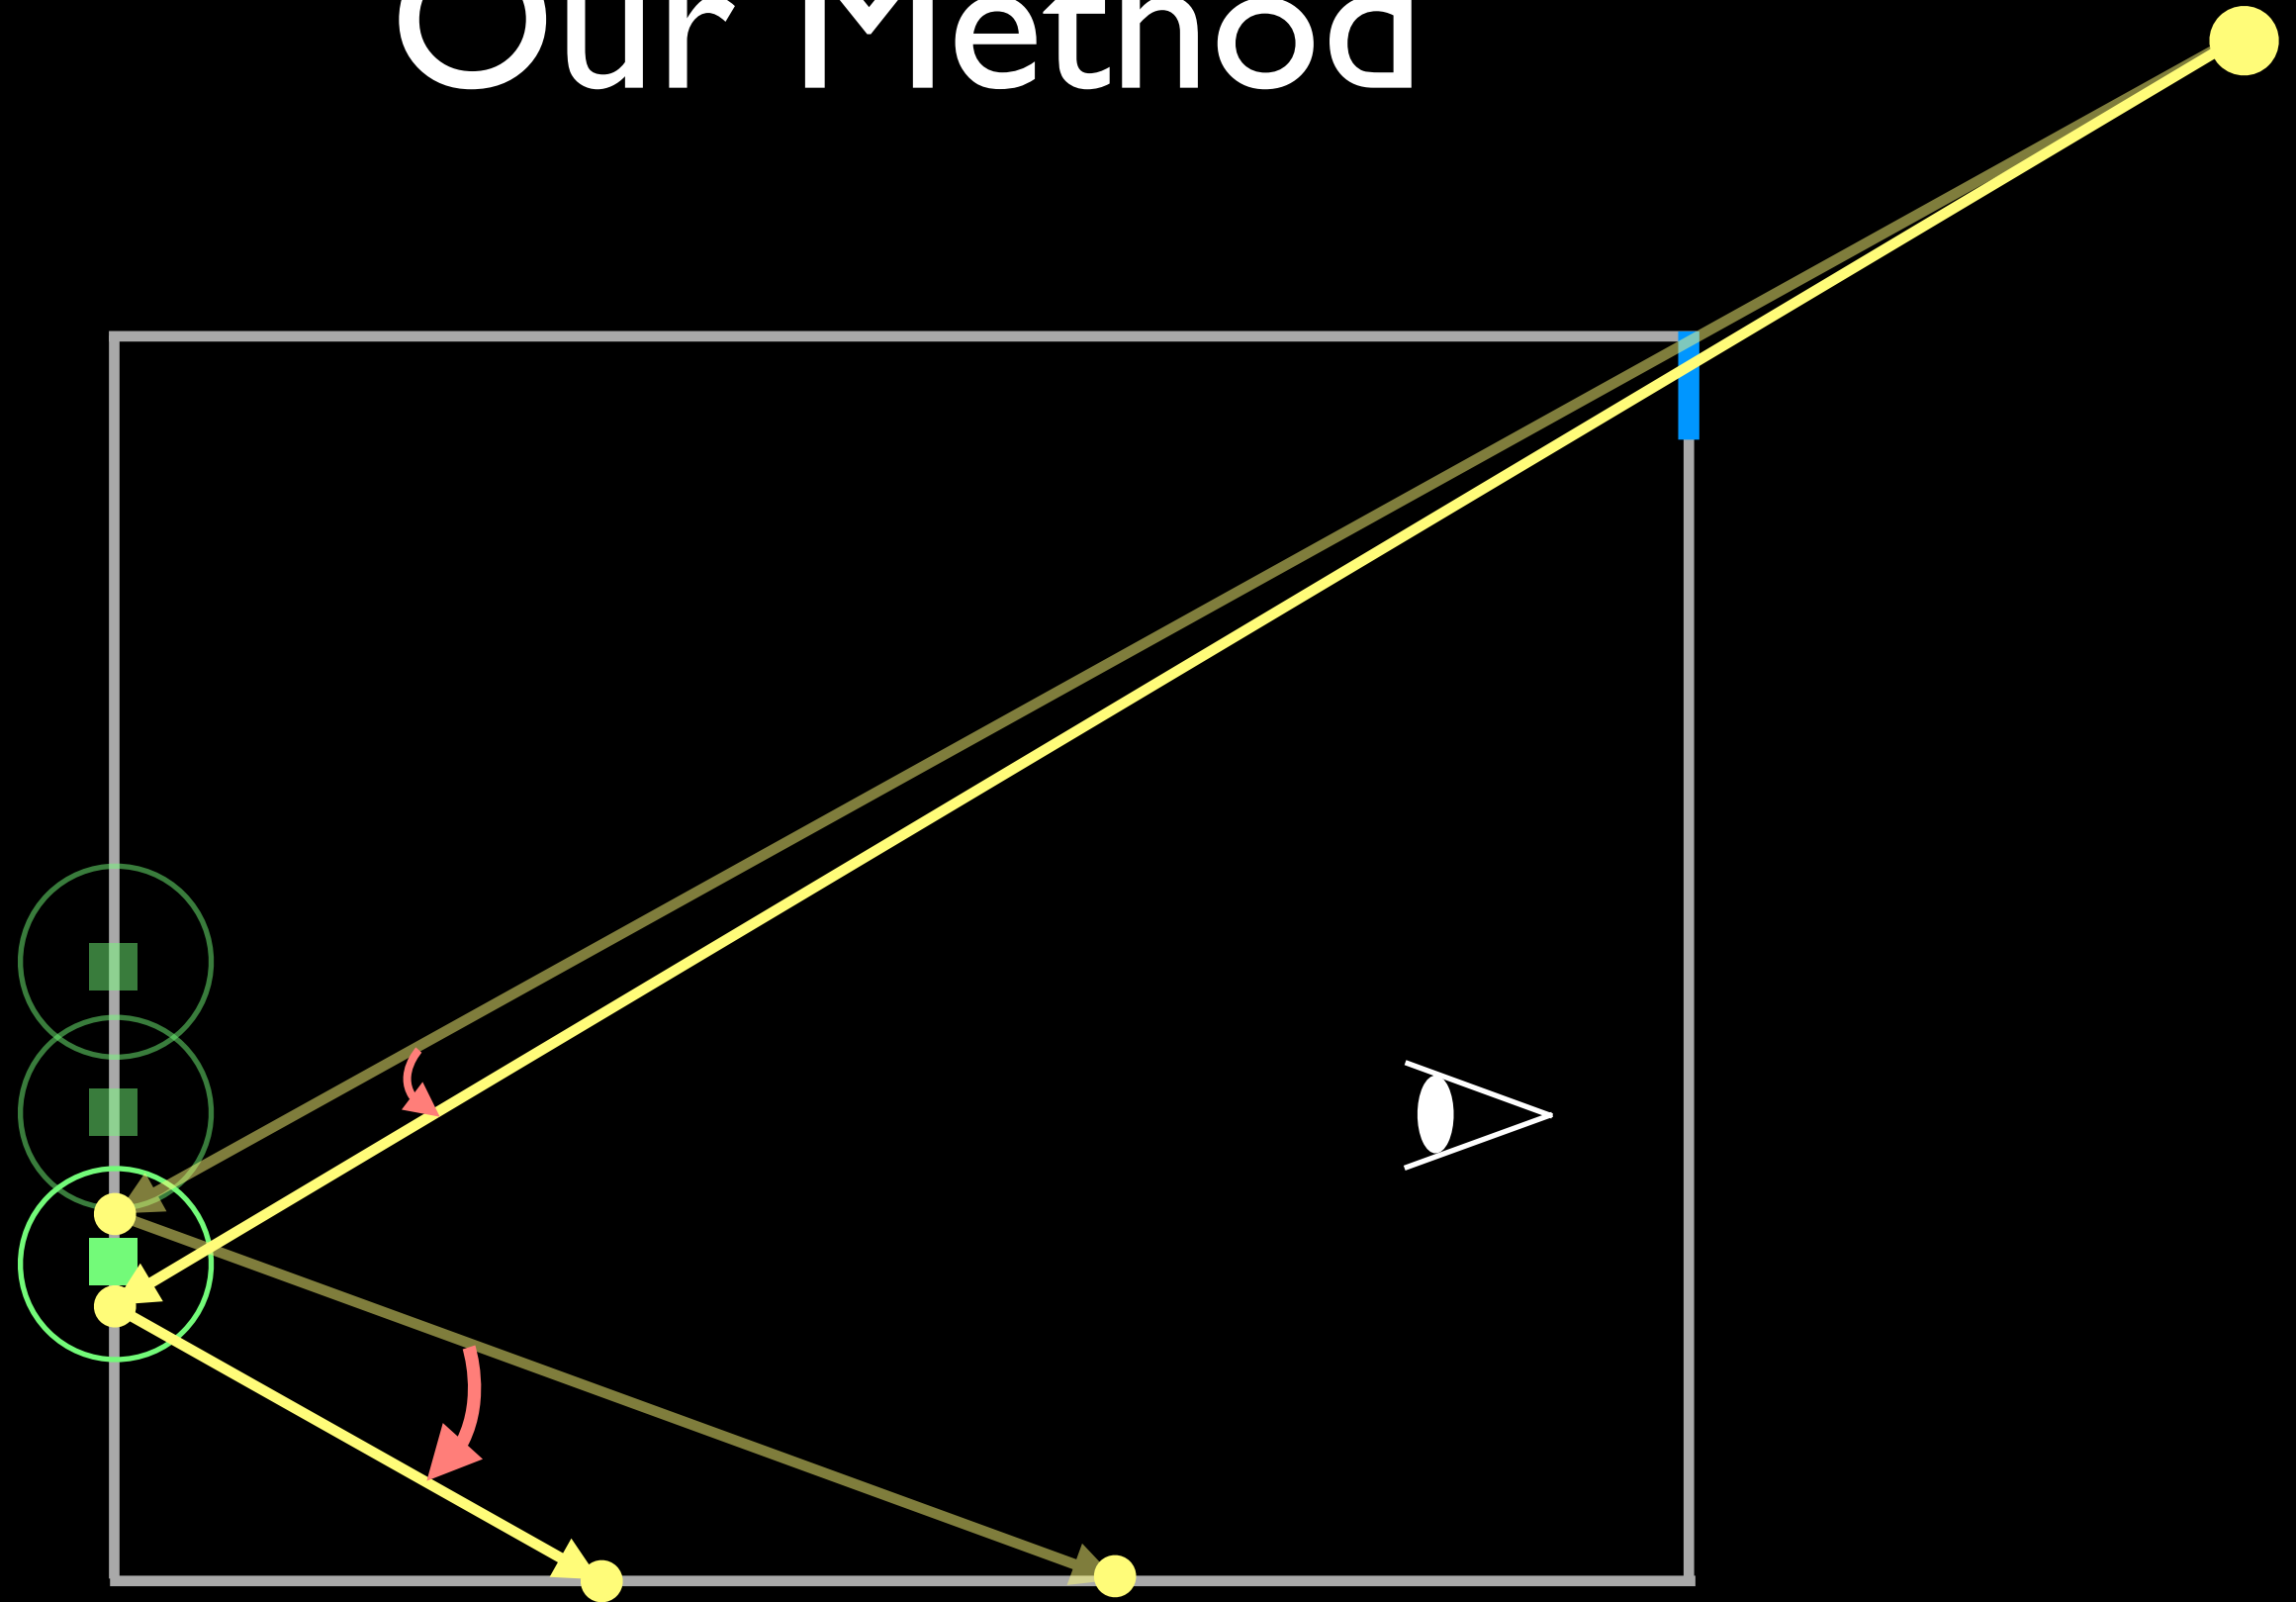
Our Method



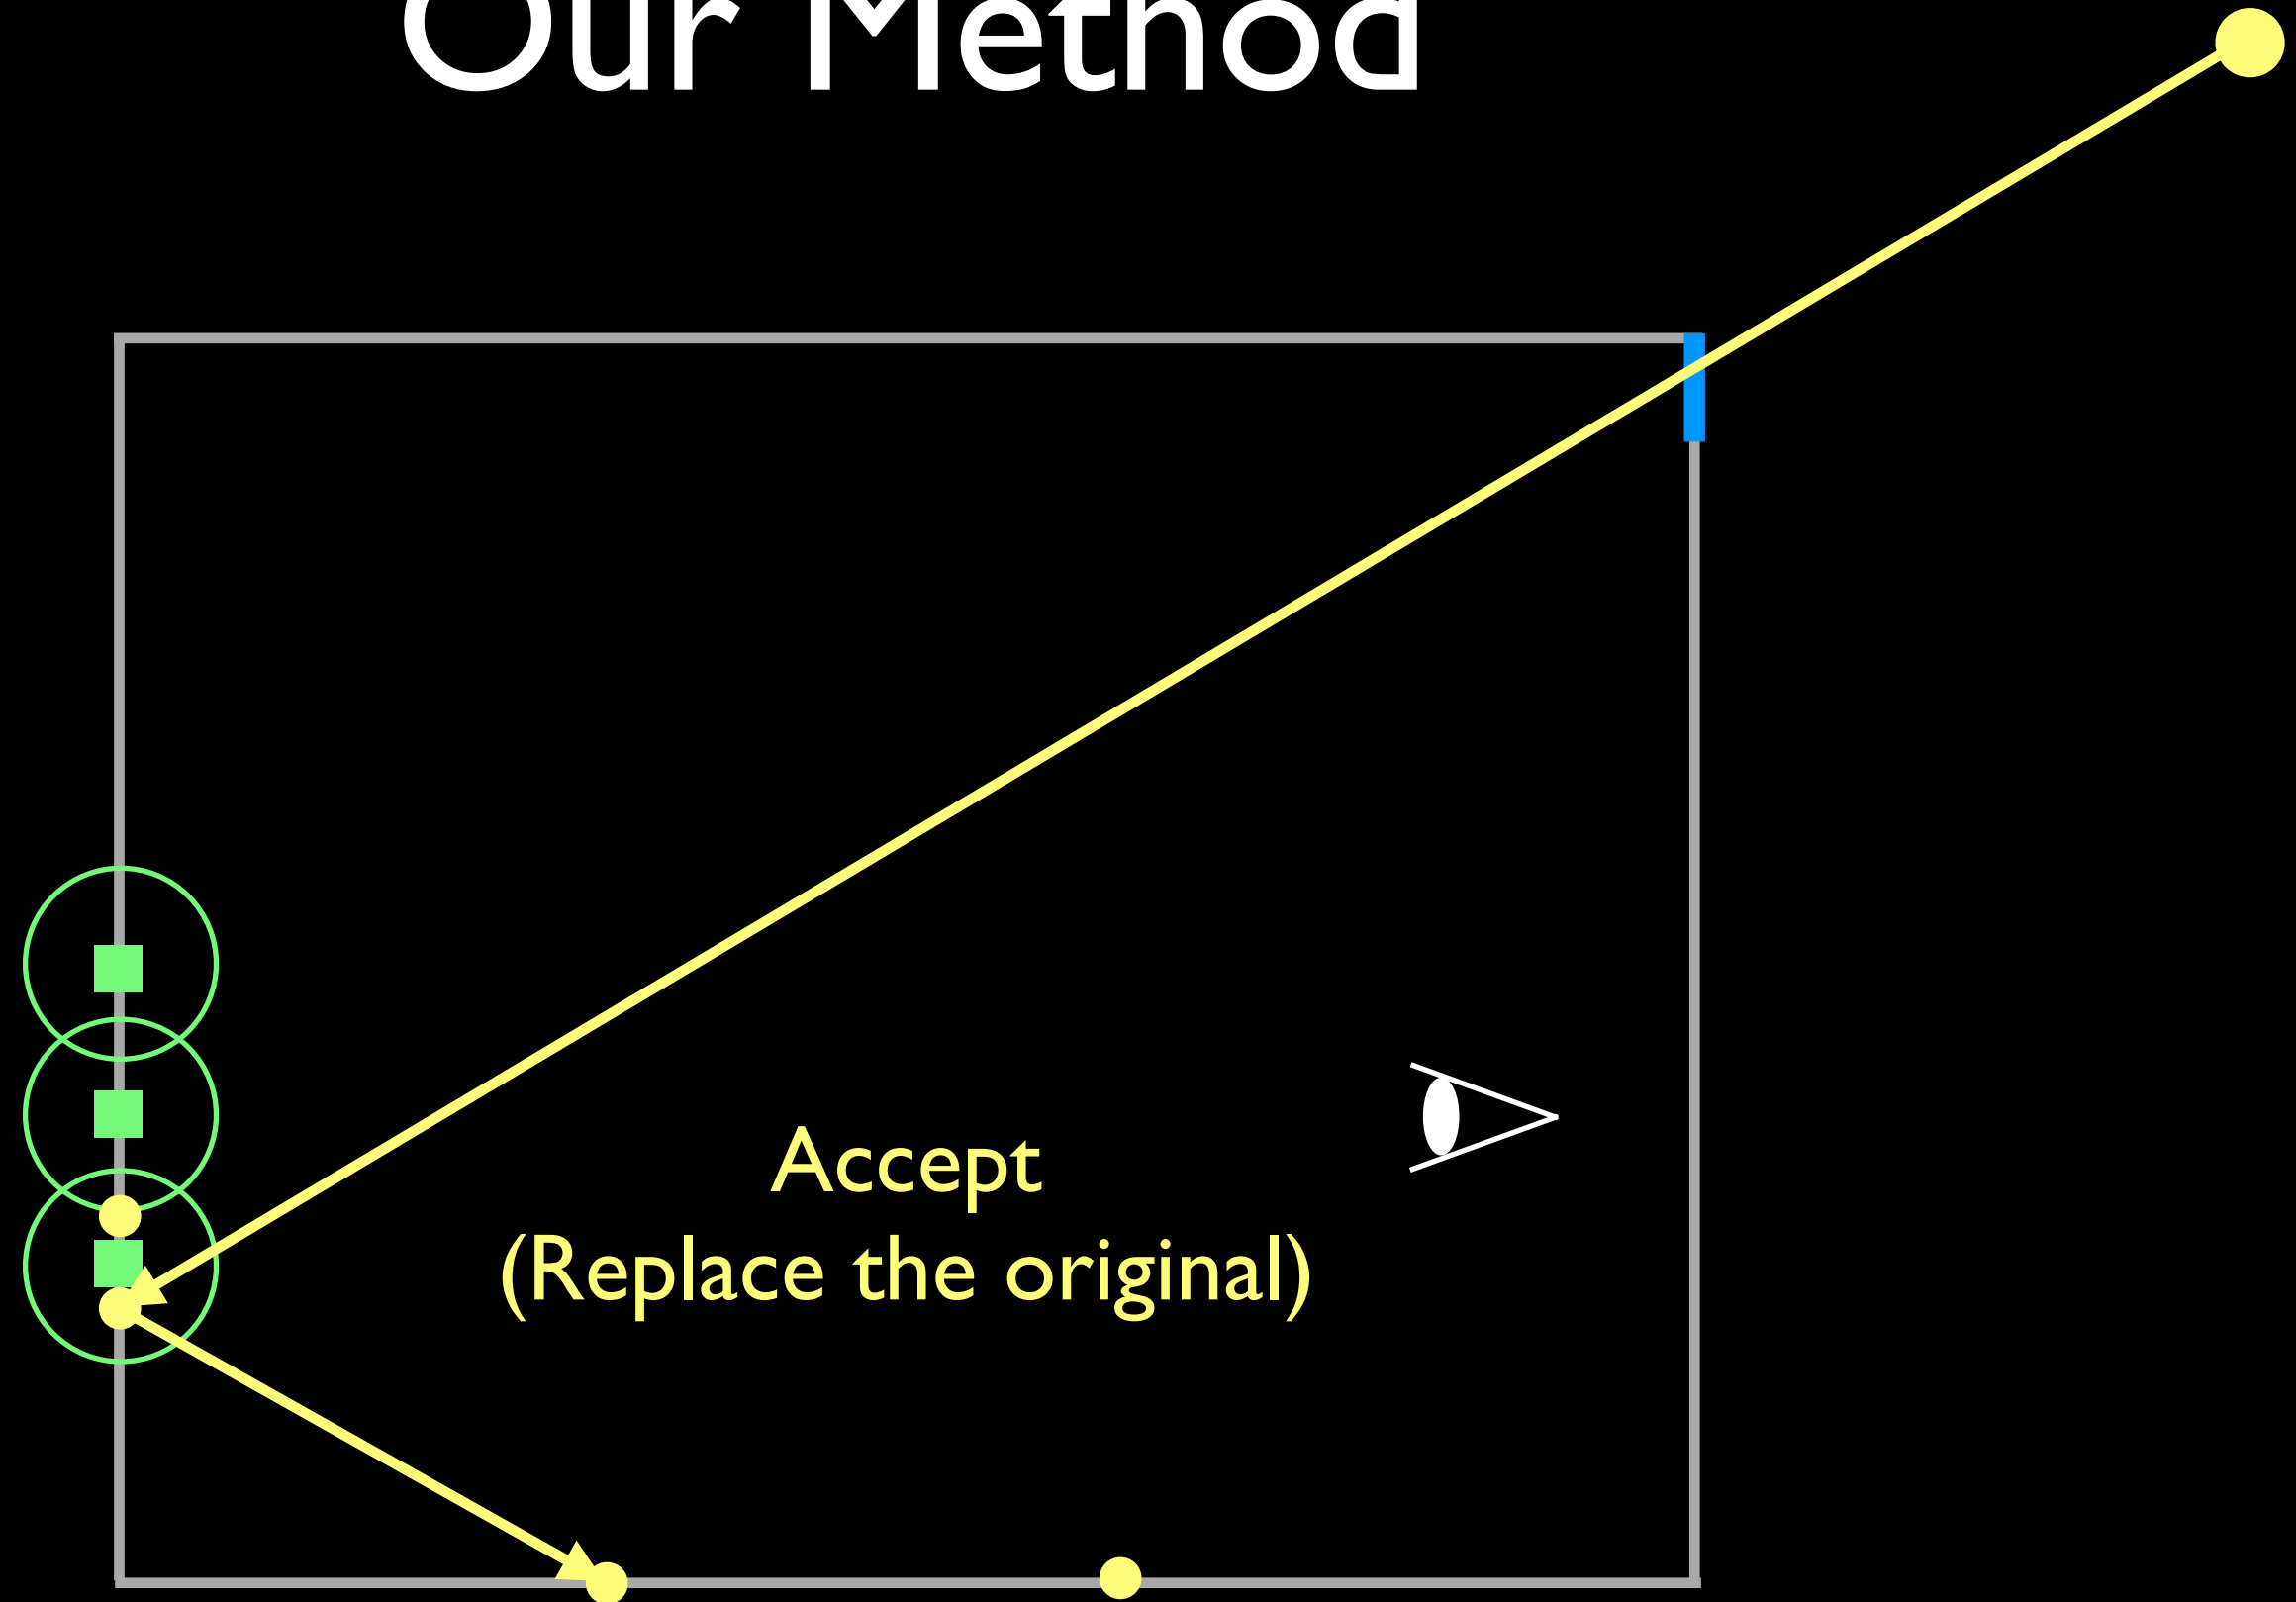
Our Method



Our Method




Our Method




Our Method



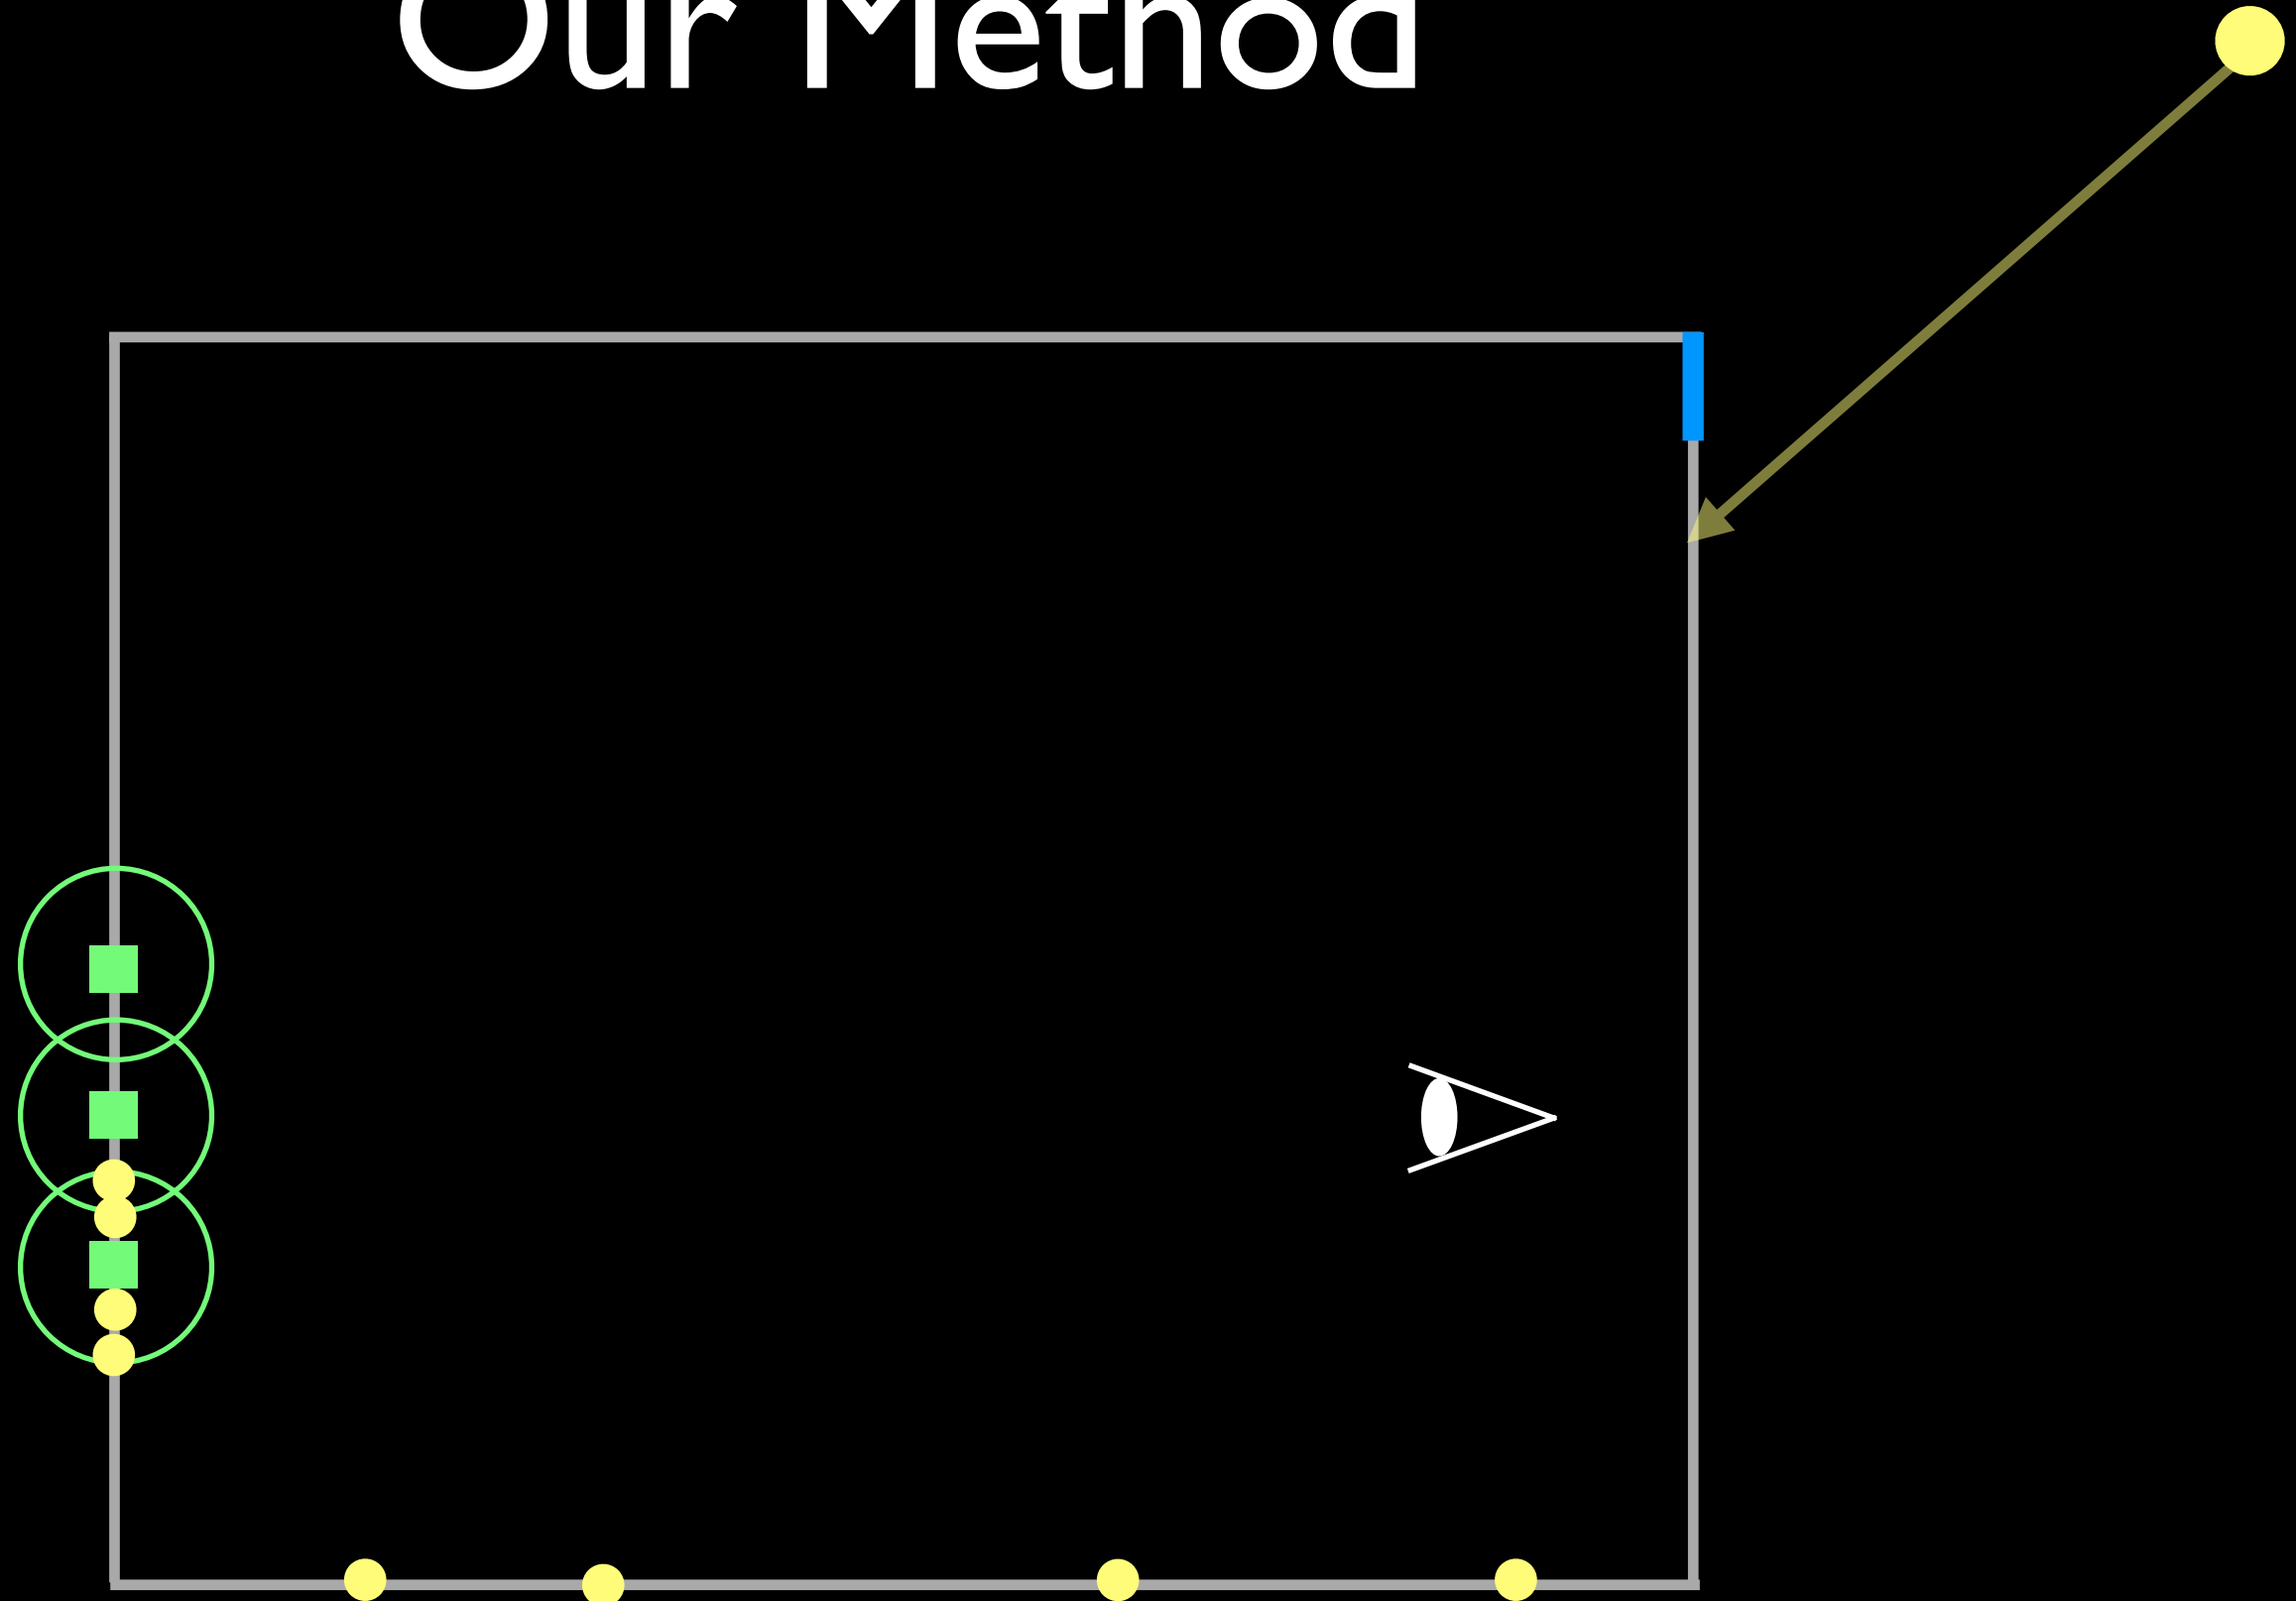
Our Method



Our Method



Our Method

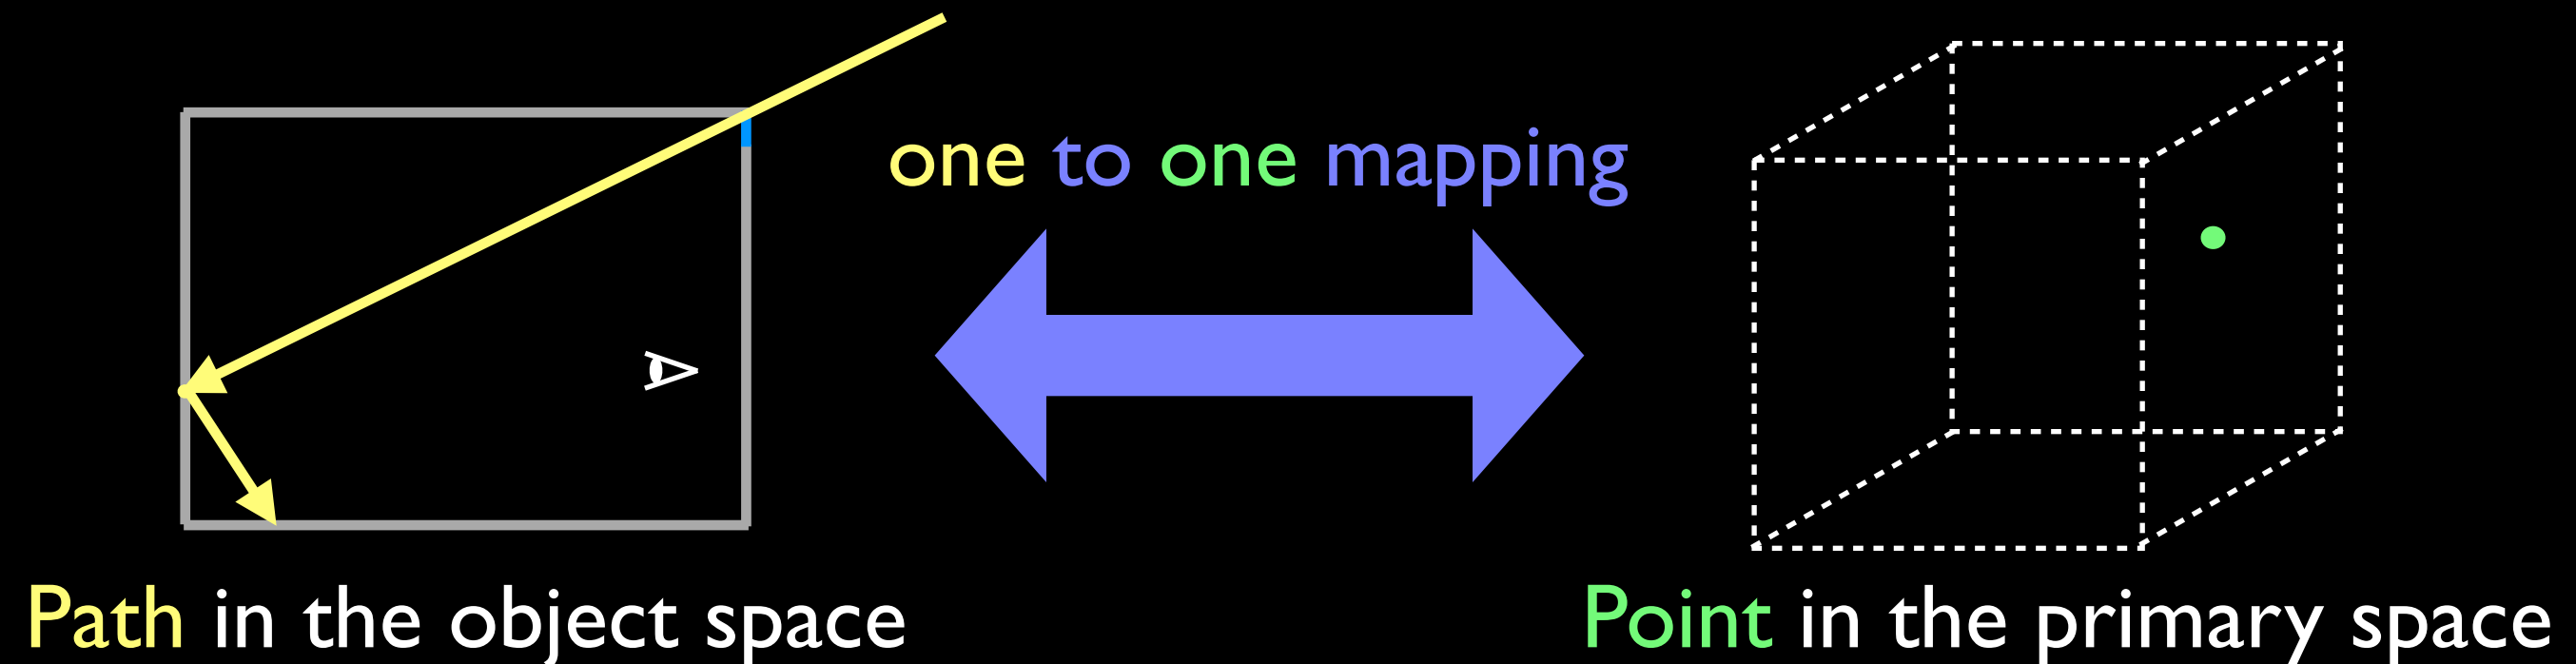


Sampling using Path Visibility

- Three technical components
 - Mapping to the primary space
 - Definition of the visibility function
 - Sampling via Markov chain Monte Carlo

Primary Space


- Mapping a path to a point in hypercube [Kelemen et al. 2002]
- Path = vector of random numbers $\vec{u} = (\xi_1, \dots, \xi_N) \in (0, 1)^N$




Visibility Function

- Binary function in the primary space

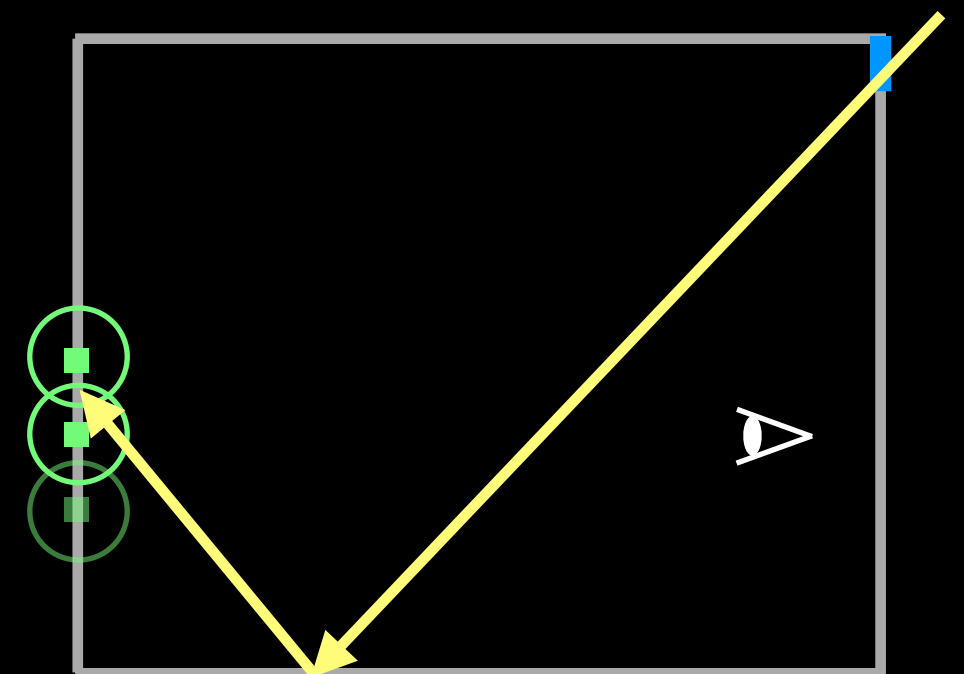
$$V(\vec{u}) = \begin{cases} 0 & \text{No contribution} \\ 1 & \text{Contributed to at least one measurement point} \end{cases}$$



Visible



Not visible



Visible

Markov Chain Monte Carlo

- Importance sampling the visibility function
 - Based on perturbation & accept/reject
 - Sampling visible paths only = Sampling only points with $V(\vec{u}) = 1$

Two Issues

- Markov chain Monte Carlo tends to...
 - be sensitive with parameter tuning
 - get trapped in small regions of the image

Problem I: Parameter Tuning

- Amount of perturbation (mutation size) affects the efficiency
- Scene dependent and unintuitive to tweak



Too small



Appropriate



Too large

Solution I: Adaptive MCMC

- Self-tuning Markov chain Monte Carlo methods
 - Relatively new techniques in statistics [Andrieu and Thomas 08]
 - Uses all the past samples to tweak parameters



Hand-tuned



Adaptive


Problem2: Insufficient Exploration

- Markov chain can get trapped to a small region in the domain
- “Sampling one of the windows too many times”




Solution2: Replica Exchange

- Sampling multiple distributions simultaneously
- Chain moves across distributions (= *exchange*)
- Bridging peaks via “easy” distributions



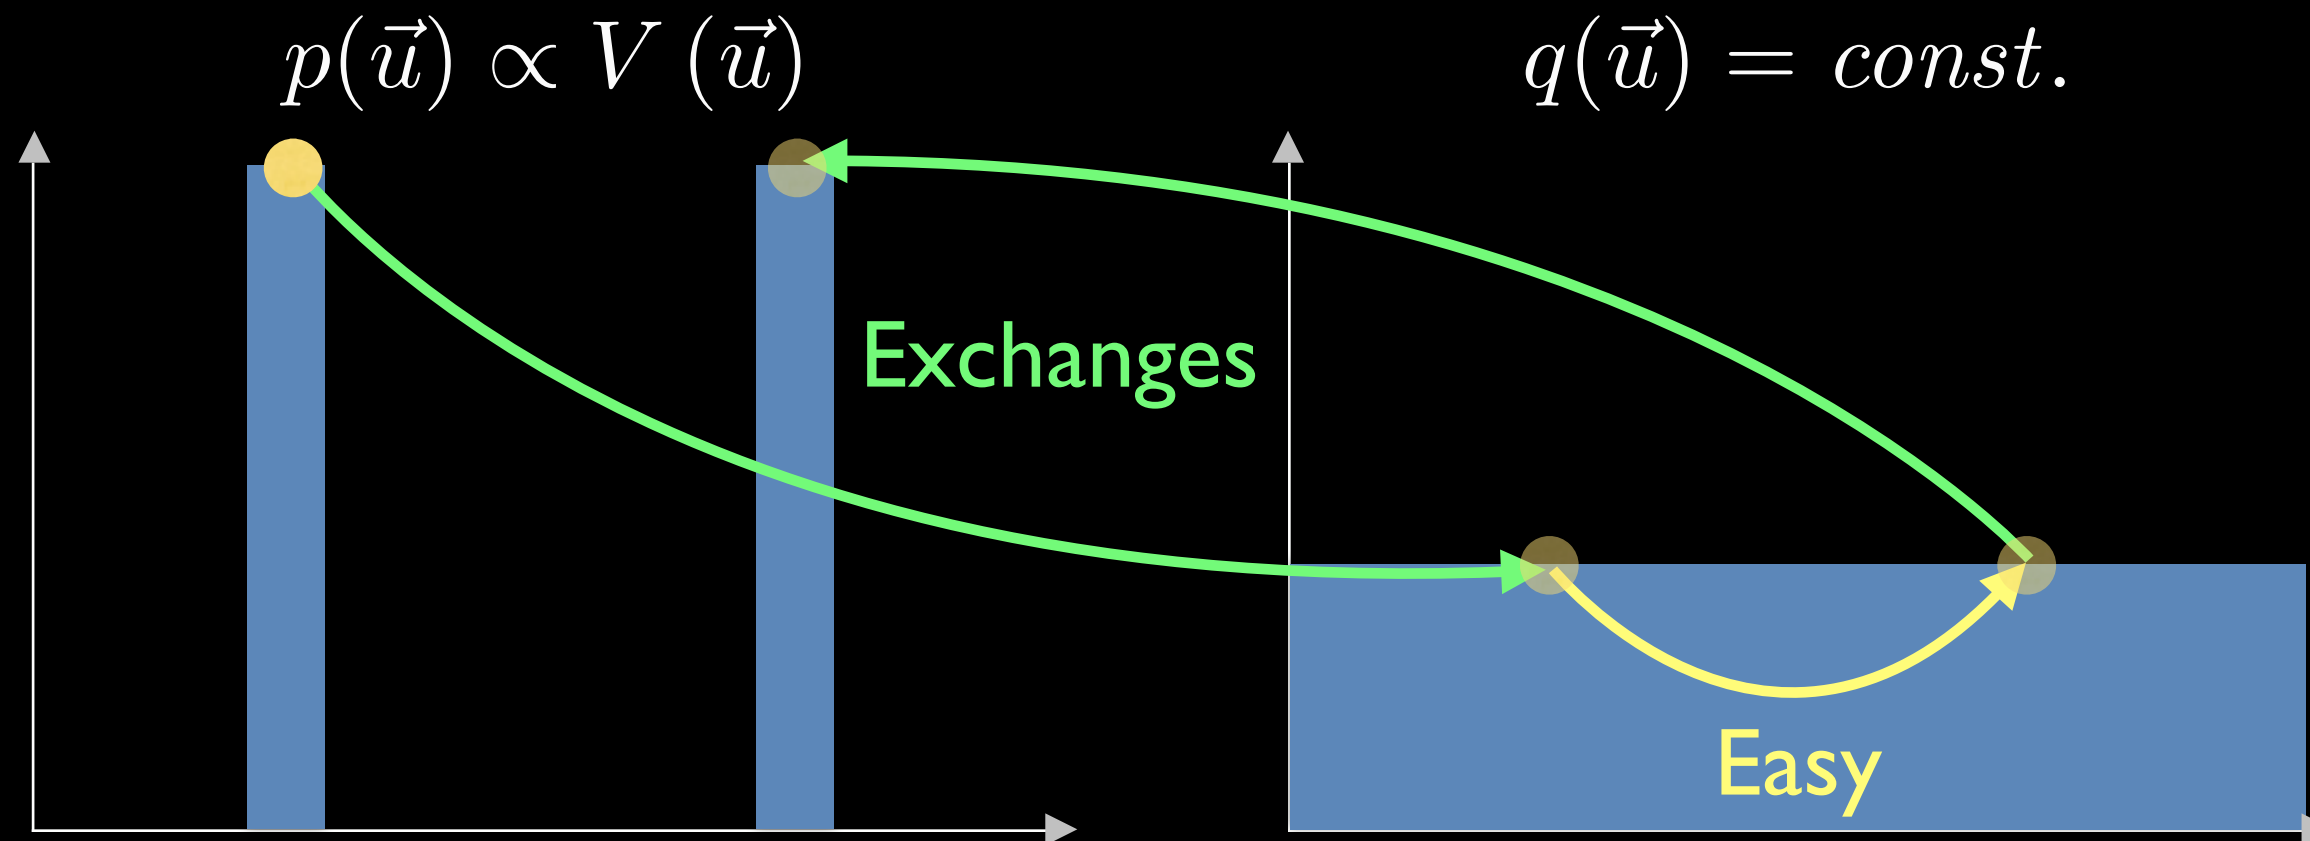
Solution2: Replica Exchange

- Sampling multiple distributions simultaneously
- Chain moves across distributions (= *exchange*)
- Bridging peaks via “easy” distributions



Solution2: Replica Exchange

- Sampling multiple distributions simultaneously
- Chain moves across distributions (= *exchange*)
- Bridging peaks via “easy” distributions



Results

Implementation

- Simple extension over a standard PPM code
 - Just enable generation of a path from a random vector
 - Successfully used as an assignment in a graphics course





1 min

15 min





30 min


60 min

Uniform



Our method











Zoom 0

Zoom 1





Zoom 2

Zoom 3


Uniform




Our method



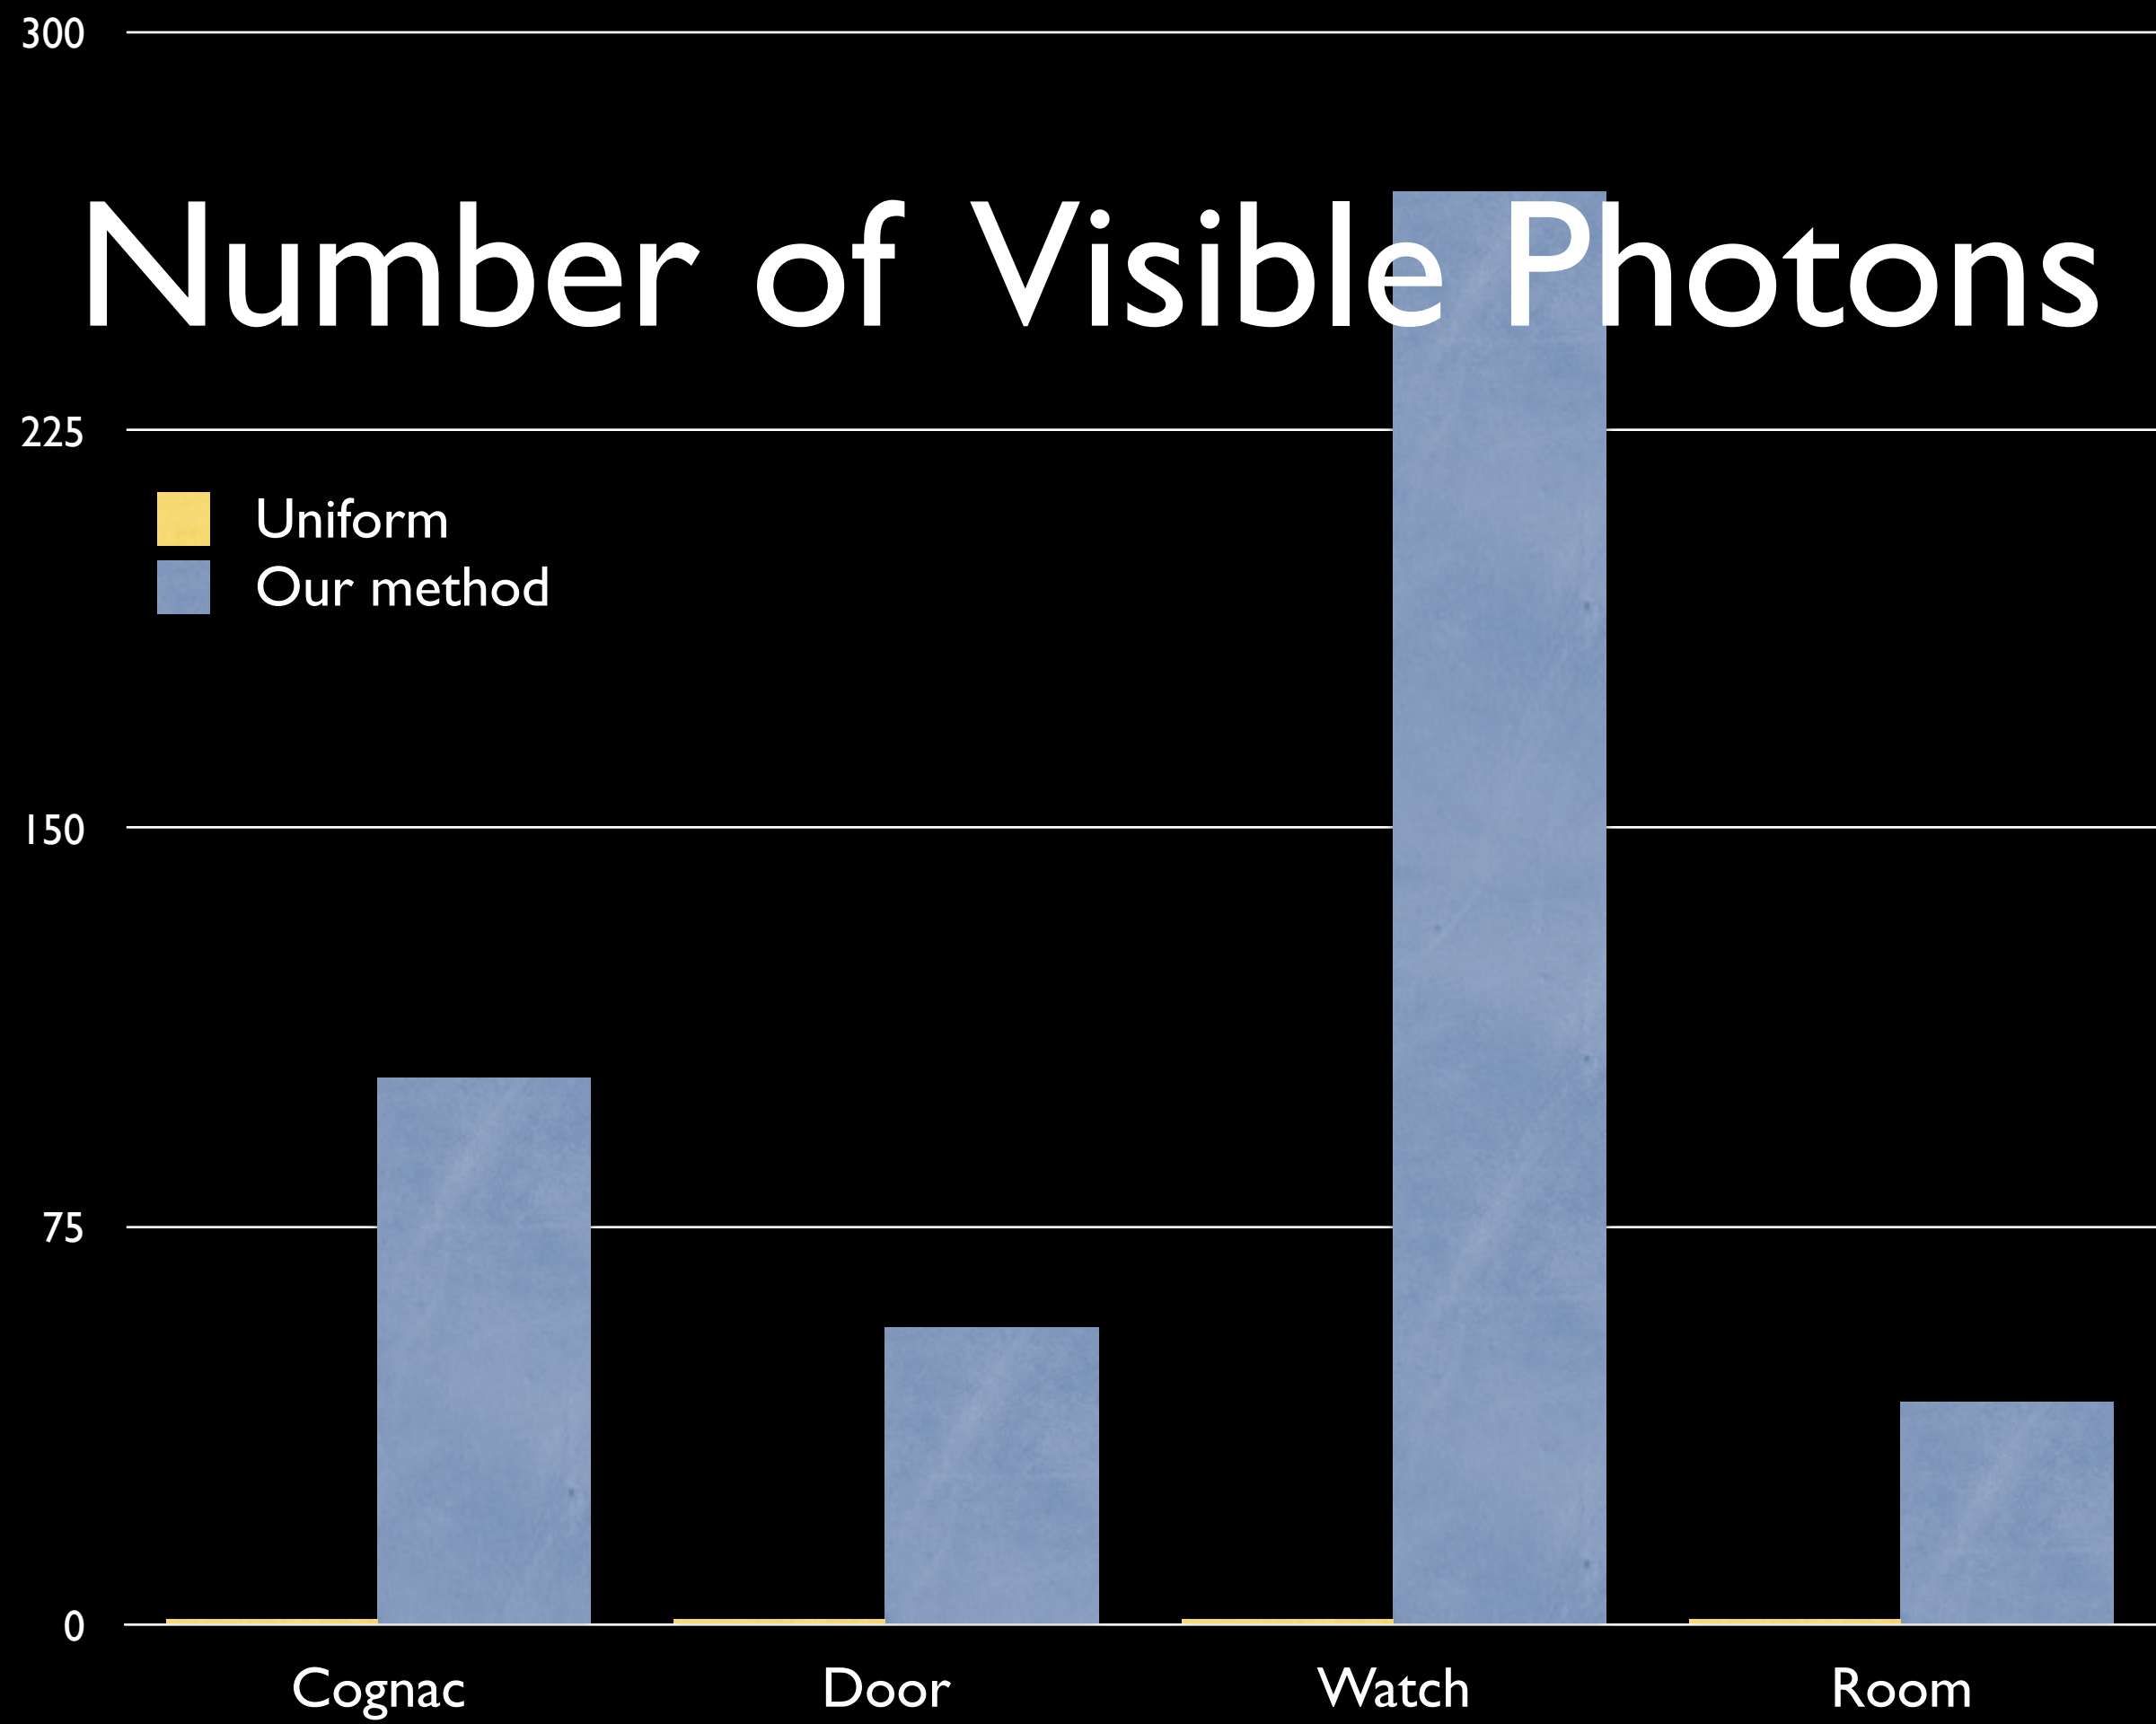
(Equal time comparisons)





Number of Visible Photons


Number of Visible Photons [M photons]



“Easy” Scenes

- Does not hurt rendering times of “easy” scenes


Our method
Conventional



(Equal time comparisons)

Limitations

- Does not resolve noise due to BRDFs



Conventional

Our method

Limitations

- ~~Does not resolve noise due to BRDFs~~



See our paper for a solution:

"A Path Space Extension for Robust Light Transport Simulation"
(a.k.a. "Vertex Connection and Merging" by Georgiev et al.)

Conventional

Our method

Limitations

- No proof of convergence

Limitations

- ~~No proof of convergence~~
- Kaplanyan and Dachsbacher proved convergence on essentially the same problem:
“Path Space Regularization for Holistic and Robust Light Transport”,
Eurographics 2013


Conclusions

- Adaptive photon tracing based on visibility
 - Samples only visible photon paths
 - Completely parameter-free
 - Efficient, simple, and **practical**

Conclusions

- Adaptive photon tracing based on visibility
 - Samples only visible photon paths
 - Completely parameter-free
 - Efficient, simple, and practical
 - You should have already tested it two years ago (if not, do so tonight)

Thank You



- Youichi Kimura (Studio Azurite) for providing us the room model and some inspiration images
- Marko Dabrovic for the Sibenik Cathedral model
- VC-ISTI for the Buddha model