Beyond Points and Beams

Higher-Dimensional Photon Samples for Volumetric Light Transport

Benedikt Bitterli   Wojciech Jarosz

Dartmouth College

DARTMOUTH VISUAL COMPUTING LAB

30 JULY–3 AUGUST Los Angeles
SIGGRAPH2017

Motivation

Photon Mapping

Photon Mapping

- Volumetric Photon Mapping
  - [Jensen & Christensen ’98]

Photon Mapping

- Volumetric Photon Mapping
  - [Jensen & Christensen ’98]
- Beam Radiance Estimate
  - [Jarosz et al. ’08]

Photon Mapping

- Volumetric Photon Mapping
  - [Jensen & Christensen ’98]
- Beam Radiance Estimate
  - [Jarosz et al. ’08]
- Photon Beams
  - [Jarosz et al. ’11]

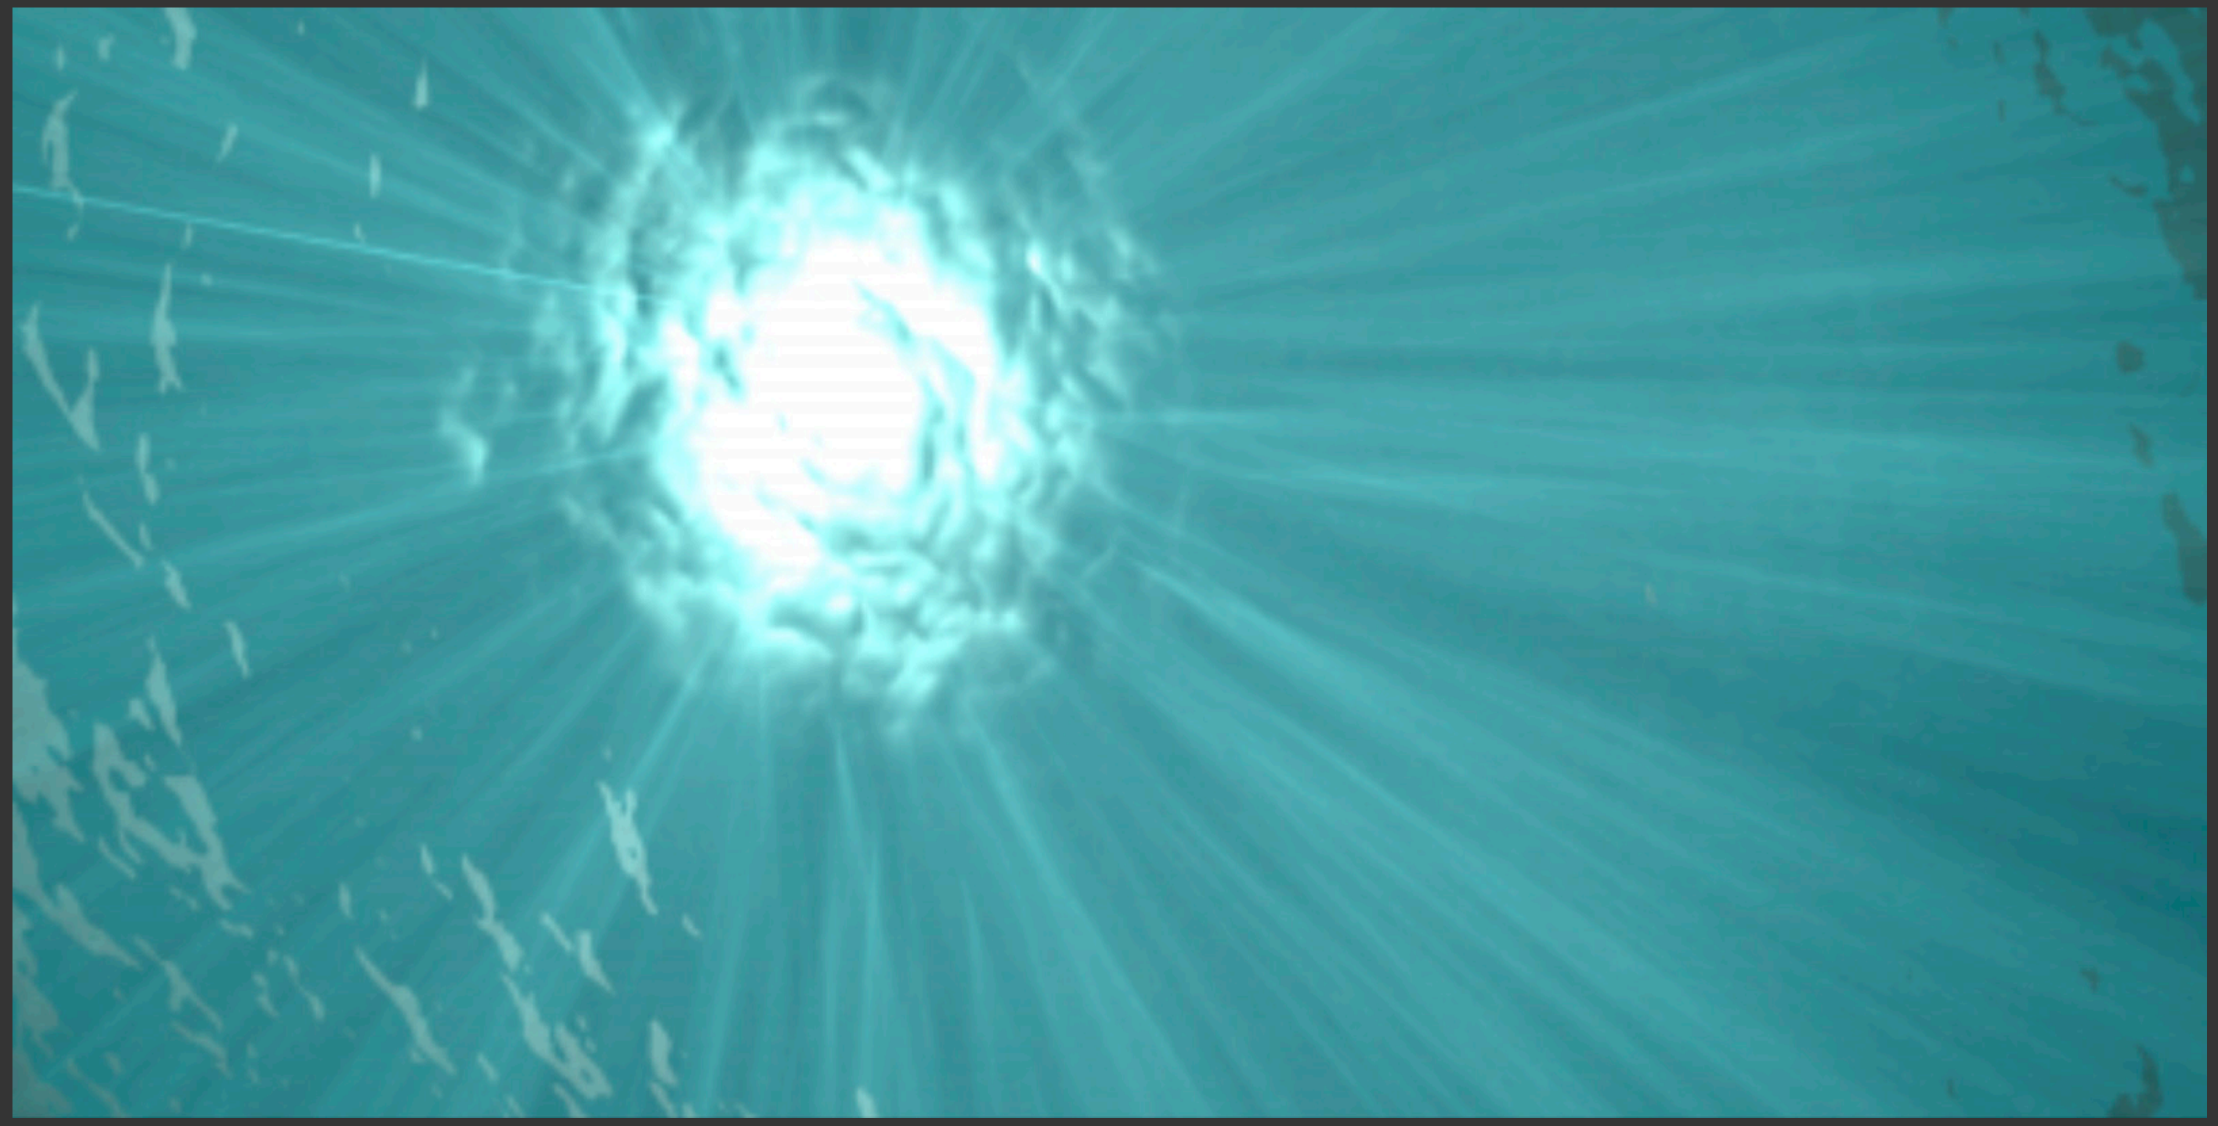

Photon Mapping

- Volumetric Photon Mapping
  - [Jensen & Christensen ’98]
- Beam Radiance Estimate
  - [Jarosz et al. ’08]
- Photon Beams
  - [Jarosz et al. ’11]
- Analysis & MIS with unbiased methods
  - [Křivánek et al. 2014]

Contributions

Contributions

- Generalized theory of density estimation

  Predict *infinite* collection of new, unbiased estimators

Contributions

- Generalized theory of density estimation

  Predict *infinite* collection of new, unbiased estimators

- Theoretical error analysis

  These estimators outperform prior work in theory

Contributions

- Generalized theory of density estimation

  Predict *infinite* collection of new, unbiased estimators
- Theoretical error analysis

  These estimators outperform prior work in theory
- Practical implementations

  …and they also do better in practice

Motivation

Disclaimer

Disclaimer

- Dense Paper

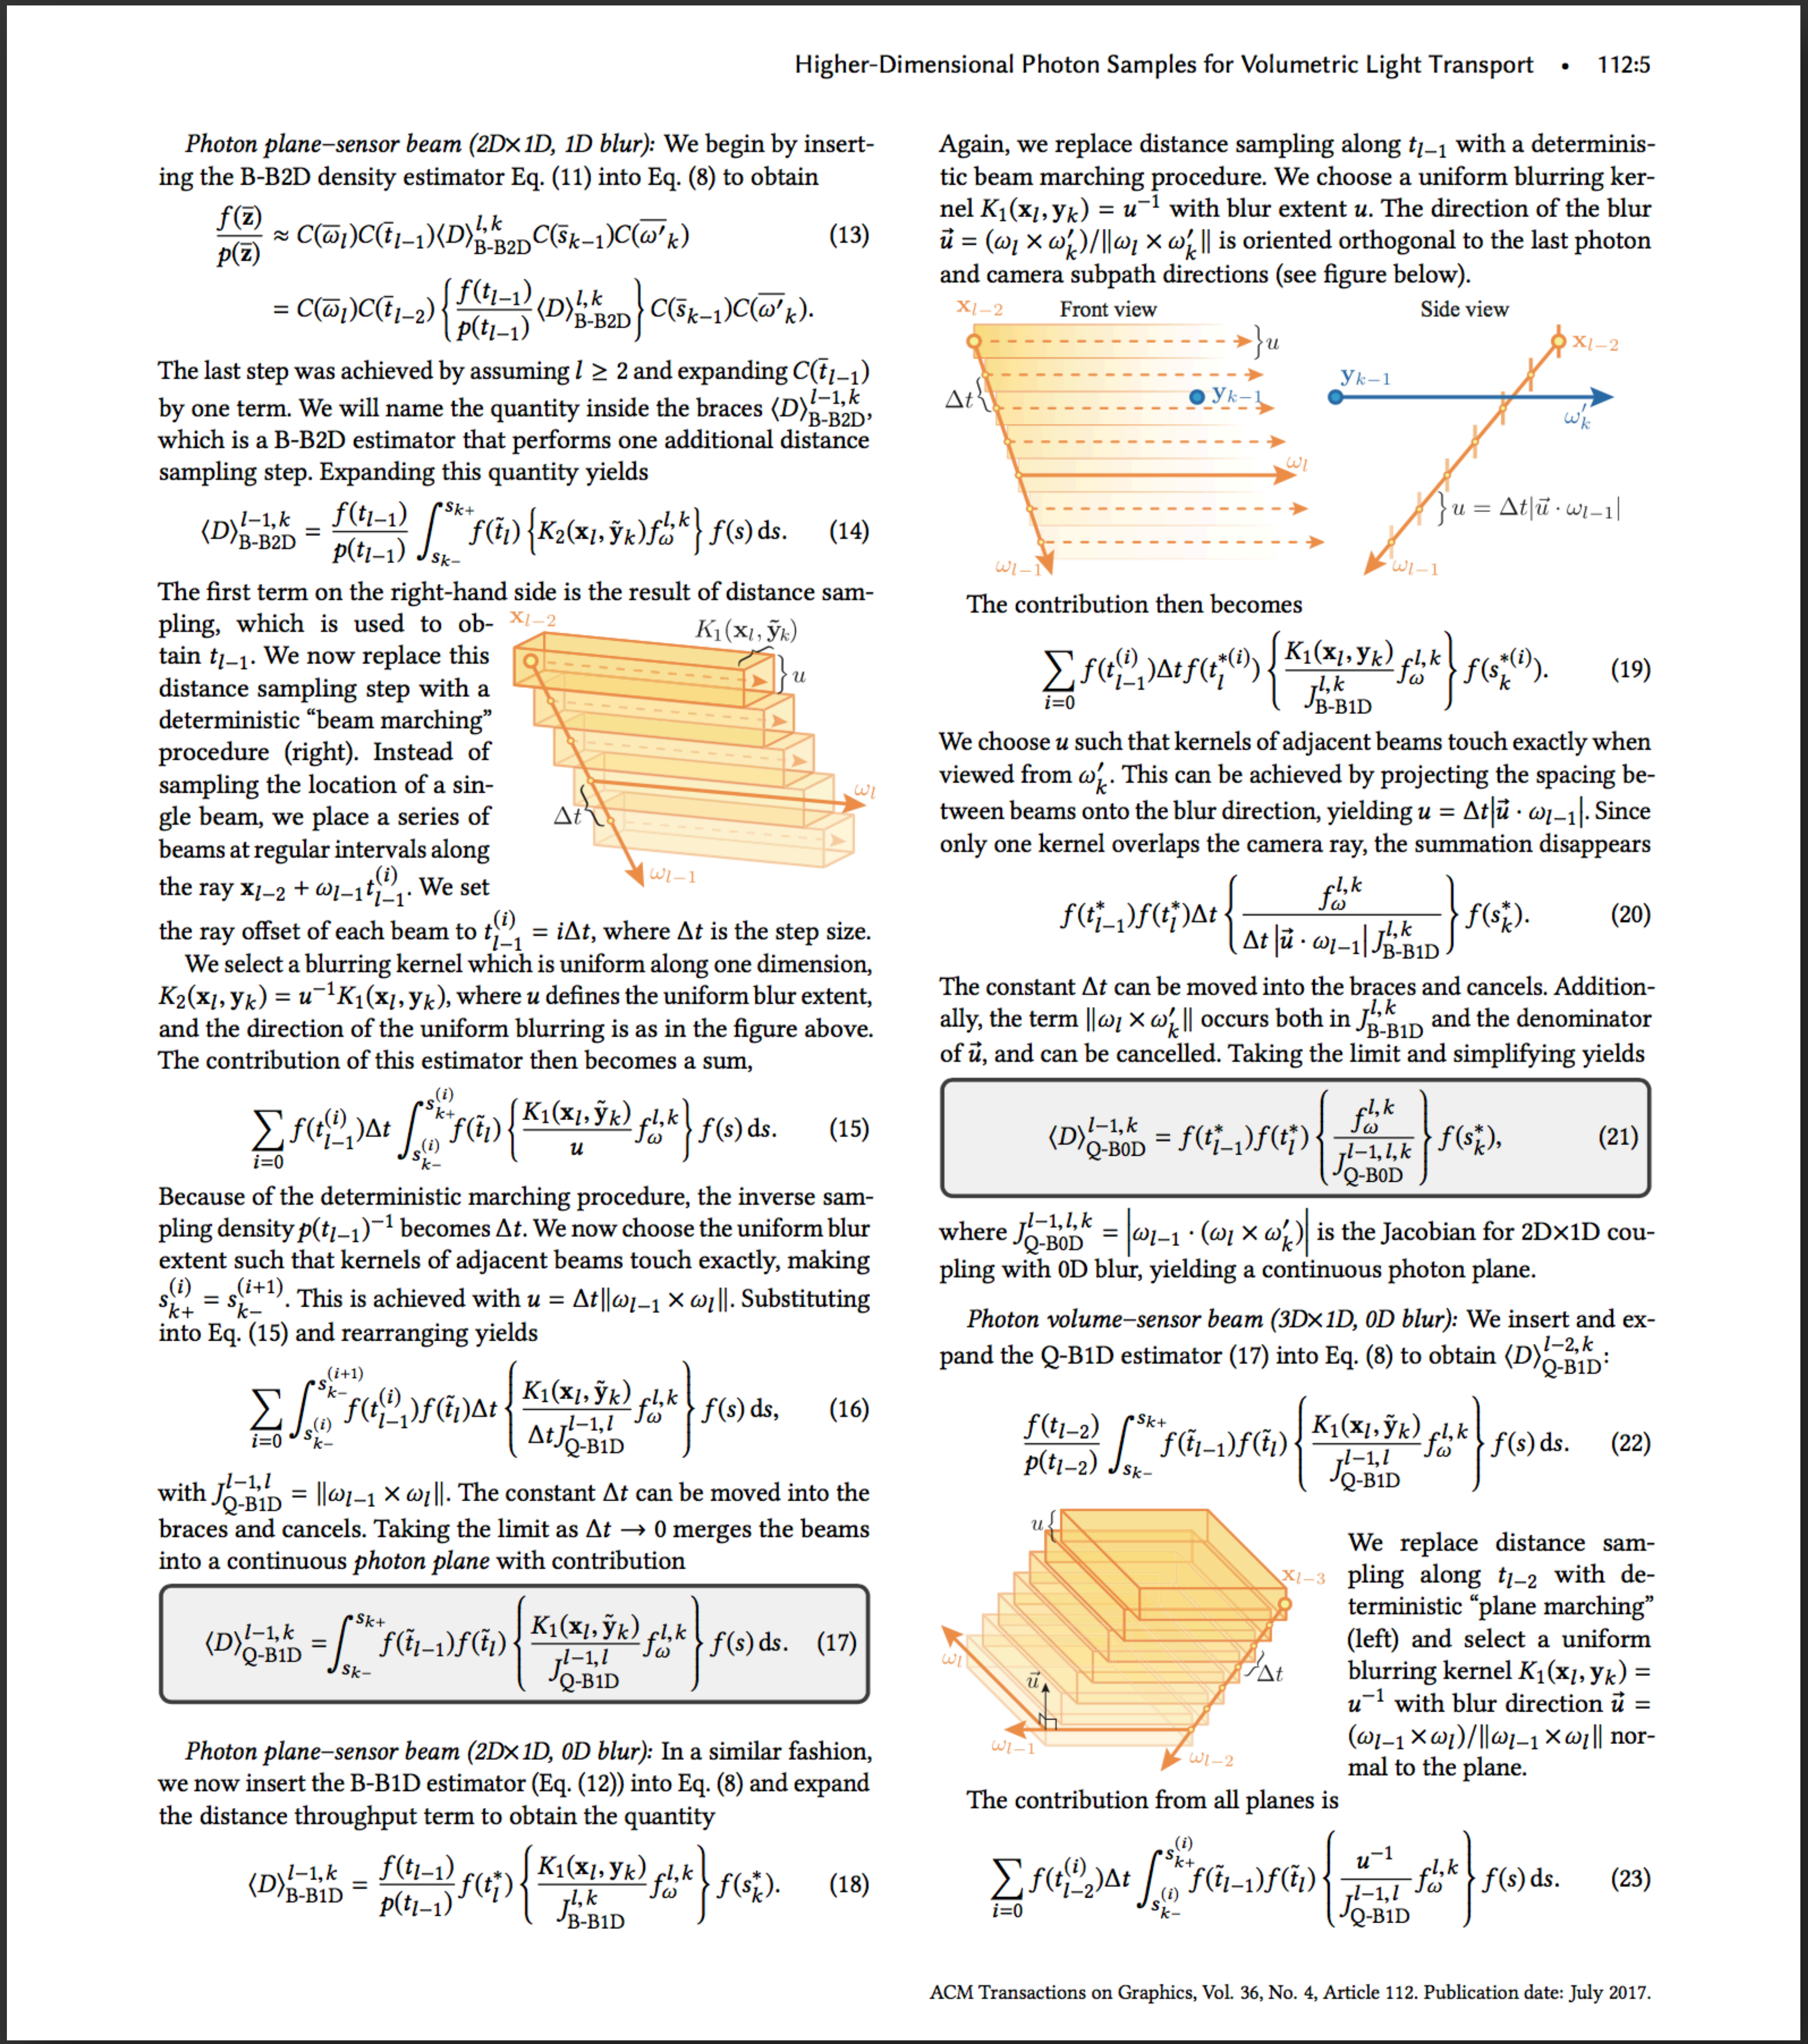
Higher-Dimensional Photon Samples for Volumetric Light Transport • 112:5

Photon plane–sensor beam (2D×1D, 1D blur): We begin by inserting the B-B2D density estimator Eq. (11) into Eq. (8) to obtain

$$\frac{f(\bar{\mathbf{z}})}{p(\bar{\mathbf{z}})} \approx C(\overline{\omega}_l)C(\bar{t}_{l-1})\langle D\rangle^{l,k}_{\text{B-B2D}}C(\bar{s}_{k-1})C(\overline{\omega'}_k) = C(\overline{\omega}_l)C(\bar{t}_{l-2})\left\{\frac{f(t_{l-1})}{p(t_{l-1})}\langle D\rangle^{l,k}_{\text{B-B2D}}\right\}C(\bar{s}_{k-1})C(\overline{\omega'}_k). \quad (13)$$

The last step was achieved by assuming $l \geq 2$ and expanding $C(\bar{t}_{l-1})$ by one term. We will name the quantity inside the braces $\langle D\rangle^{l-1,k}_{\text{B-B2D}}$, which is a B-B2D estimator that performs one additional distance sampling step. Expanding this quantity yields

$$\langle D\rangle^{l-1,k}_{\text{B-B2D}} = \frac{f(t_{l-1})}{p(t_{l-1})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_l)\left\{K_2(\mathbf{x}_l, \tilde{\mathbf{y}}_k)f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (14)$$

The first term on the right-hand side is the result of distance sampling, which is used to obtain t_{l-1}. We now replace this distance sampling step with a deterministic "beam marching" procedure (right). Instead of sampling the location of a single beam, we place a series of beams at regular intervals along the ray $\mathbf{x}_{l-2} + \omega_{l-1}t^{(i)}_{l-1}$. We set the ray offset of each beam to $t^{(i)}_{l-1} = i\Delta t$, where Δt is the step size.

We select a blurring kernel which is uniform along one dimension, $K_2(\mathbf{x}_l, \mathbf{y}_k) = u^{-1}K_1(\mathbf{x}_l, \mathbf{y}_k)$, where u defines the uniform blur extent, and the direction of the uniform blurring is as in the figure above. The contribution of this estimator then becomes a sum,

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t \int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l, \tilde{\mathbf{y}}_k)}{u}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (15)$$

Because of the deterministic marching procedure, the inverse sampling density $p(t_{l-1})^{-1}$ becomes Δt. We now choose the uniform blur extent such that kernels of adjacent beams touch exactly, making $s^{(i)}_{k+} = s^{(i+1)}_{k-}$. This is achieved with $u = \Delta t\|\omega_{l-1} \times \omega_l\|$. Substituting into Eq. (15) and rearranging yields

$$\sum_{i=0}\int_{s^{(i)}_{k-}}^{s^{(i+1)}_{k-}} f(t^{(i)}_{l-1})f(\tilde{t}_l)\Delta t\left\{\frac{K_1(\mathbf{x}_l, \tilde{\mathbf{y}}_k)}{\Delta t J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s, \quad (16)$$

with $J^{l-1,l}_{\text{Q-B1D}} = \|\omega_{l-1} \times \omega_l\|$. The constant Δt can be moved into the braces and cancels. Taking the limit as $\Delta t \to 0$ merges the beams into a continuous *photon plane* with contribution

$$\langle D\rangle^{l-1,k}_{\text{Q-B1D}} = \int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l, \tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (17)$$

Photon plane–sensor beam (2D×1D, 0D blur): In a similar fashion, we now insert the B-B1D estimator (Eq. (12)) into Eq. (8) and expand the distance throughput term to obtain the quantity

$$\langle D\rangle^{l-1,k}_{\text{B-B1D}} = \frac{f(t_{l-1})}{p(t_{l-1})}f(t^*_l)\left\{\frac{K_1(\mathbf{x}_l, \mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_{\omega}\right\}f(s^*_k). \quad (18)$$

Again, we replace distance sampling along t_{l-1} with a deterministic beam marching procedure. We choose a uniform blurring kernel $K_1(\mathbf{x}_l, \mathbf{y}_k) = u^{-1}$ with blur extent u. The direction of the blur $\vec{u} = (\omega_l \times \omega'_k)/\|\omega_l \times \omega'_k\|$ is oriented orthogonal to the last photon and camera subpath directions (see figure below).

The contribution then becomes

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t f(t^{*(i)}_l)\left\{\frac{K_1(\mathbf{x}_l, \mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_{\omega}\right\}f(s^{*(i)}_k). \quad (19)$$

We choose u such that kernels of adjacent beams touch exactly when viewed from ω'_k. This can be achieved by projecting the spacing between beams onto the blur direction, yielding $u = \Delta t|\vec{u} \cdot \omega_{l-1}|$. Since only one kernel overlaps the camera ray, the summation disappears

$$f(t^*_{l-1})f(t^*_l)\Delta t\left\{\frac{f^{l,k}_{\omega}}{\Delta t\left|\vec{u} \cdot \omega_{l-1}\right|J^{l,k}_{\text{B-B1D}}}\right\}f(s^*_k). \quad (20)$$

The constant Δt can be moved into the braces and cancels. Additionally, the term $\|\omega_l \times \omega'_k\|$ occurs both in $J^{l,k}_{\text{B-B1D}}$ and the denominator of $\vec{u}$, and can be cancelled. Taking the limit and simplifying yields

$$\langle D\rangle^{l-1,k}_{\text{Q-B0D}} = f(t^*_{l-1})f(t^*_l)\left\{\frac{f^{l,k}_{\omega}}{J^{l-1,l,k}_{\text{Q-B0D}}}\right\}f(s^*_k), \quad (21)$$

where $J^{l-1,l,k}_{\text{Q-B0D}} = \left|\omega_{l-1} \cdot (\omega_l \times \omega'_k)\right|$ is the Jacobian for 2D×1D coupling with 0D blur, yielding a continuous photon plane.

Photon volume–sensor beam (3D×1D, 0D blur): We insert and expand the Q-B1D estimator (17) into Eq. (8) to obtain $\langle D\rangle^{l-2,k}_{\text{Q-B1D}}$:

$$\frac{f(t_{l-2})}{p(t_{l-2})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l, \tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (22)$$

We replace distance sampling along t_{l-2} with deterministic "plane marching" (left) and select a uniform blurring kernel $K_1(\mathbf{x}_l, \mathbf{y}_k) = u^{-1}$ with blur direction $\vec{u} = (\omega_{l-1} \times \omega_l)/\|\omega_{l-1} \times \omega_l\|$ normal to the plane.

The contribution from all planes is

$$\sum_{i=0} f(t^{(i)}_{l-2})\Delta t \int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{u^{-1}}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (23)$$

ACM Transactions on Graphics, Vol. 36, No. 4, Article 112. Publication date: July 2017.

Disclaimer

- Dense Paper
- Light Talk

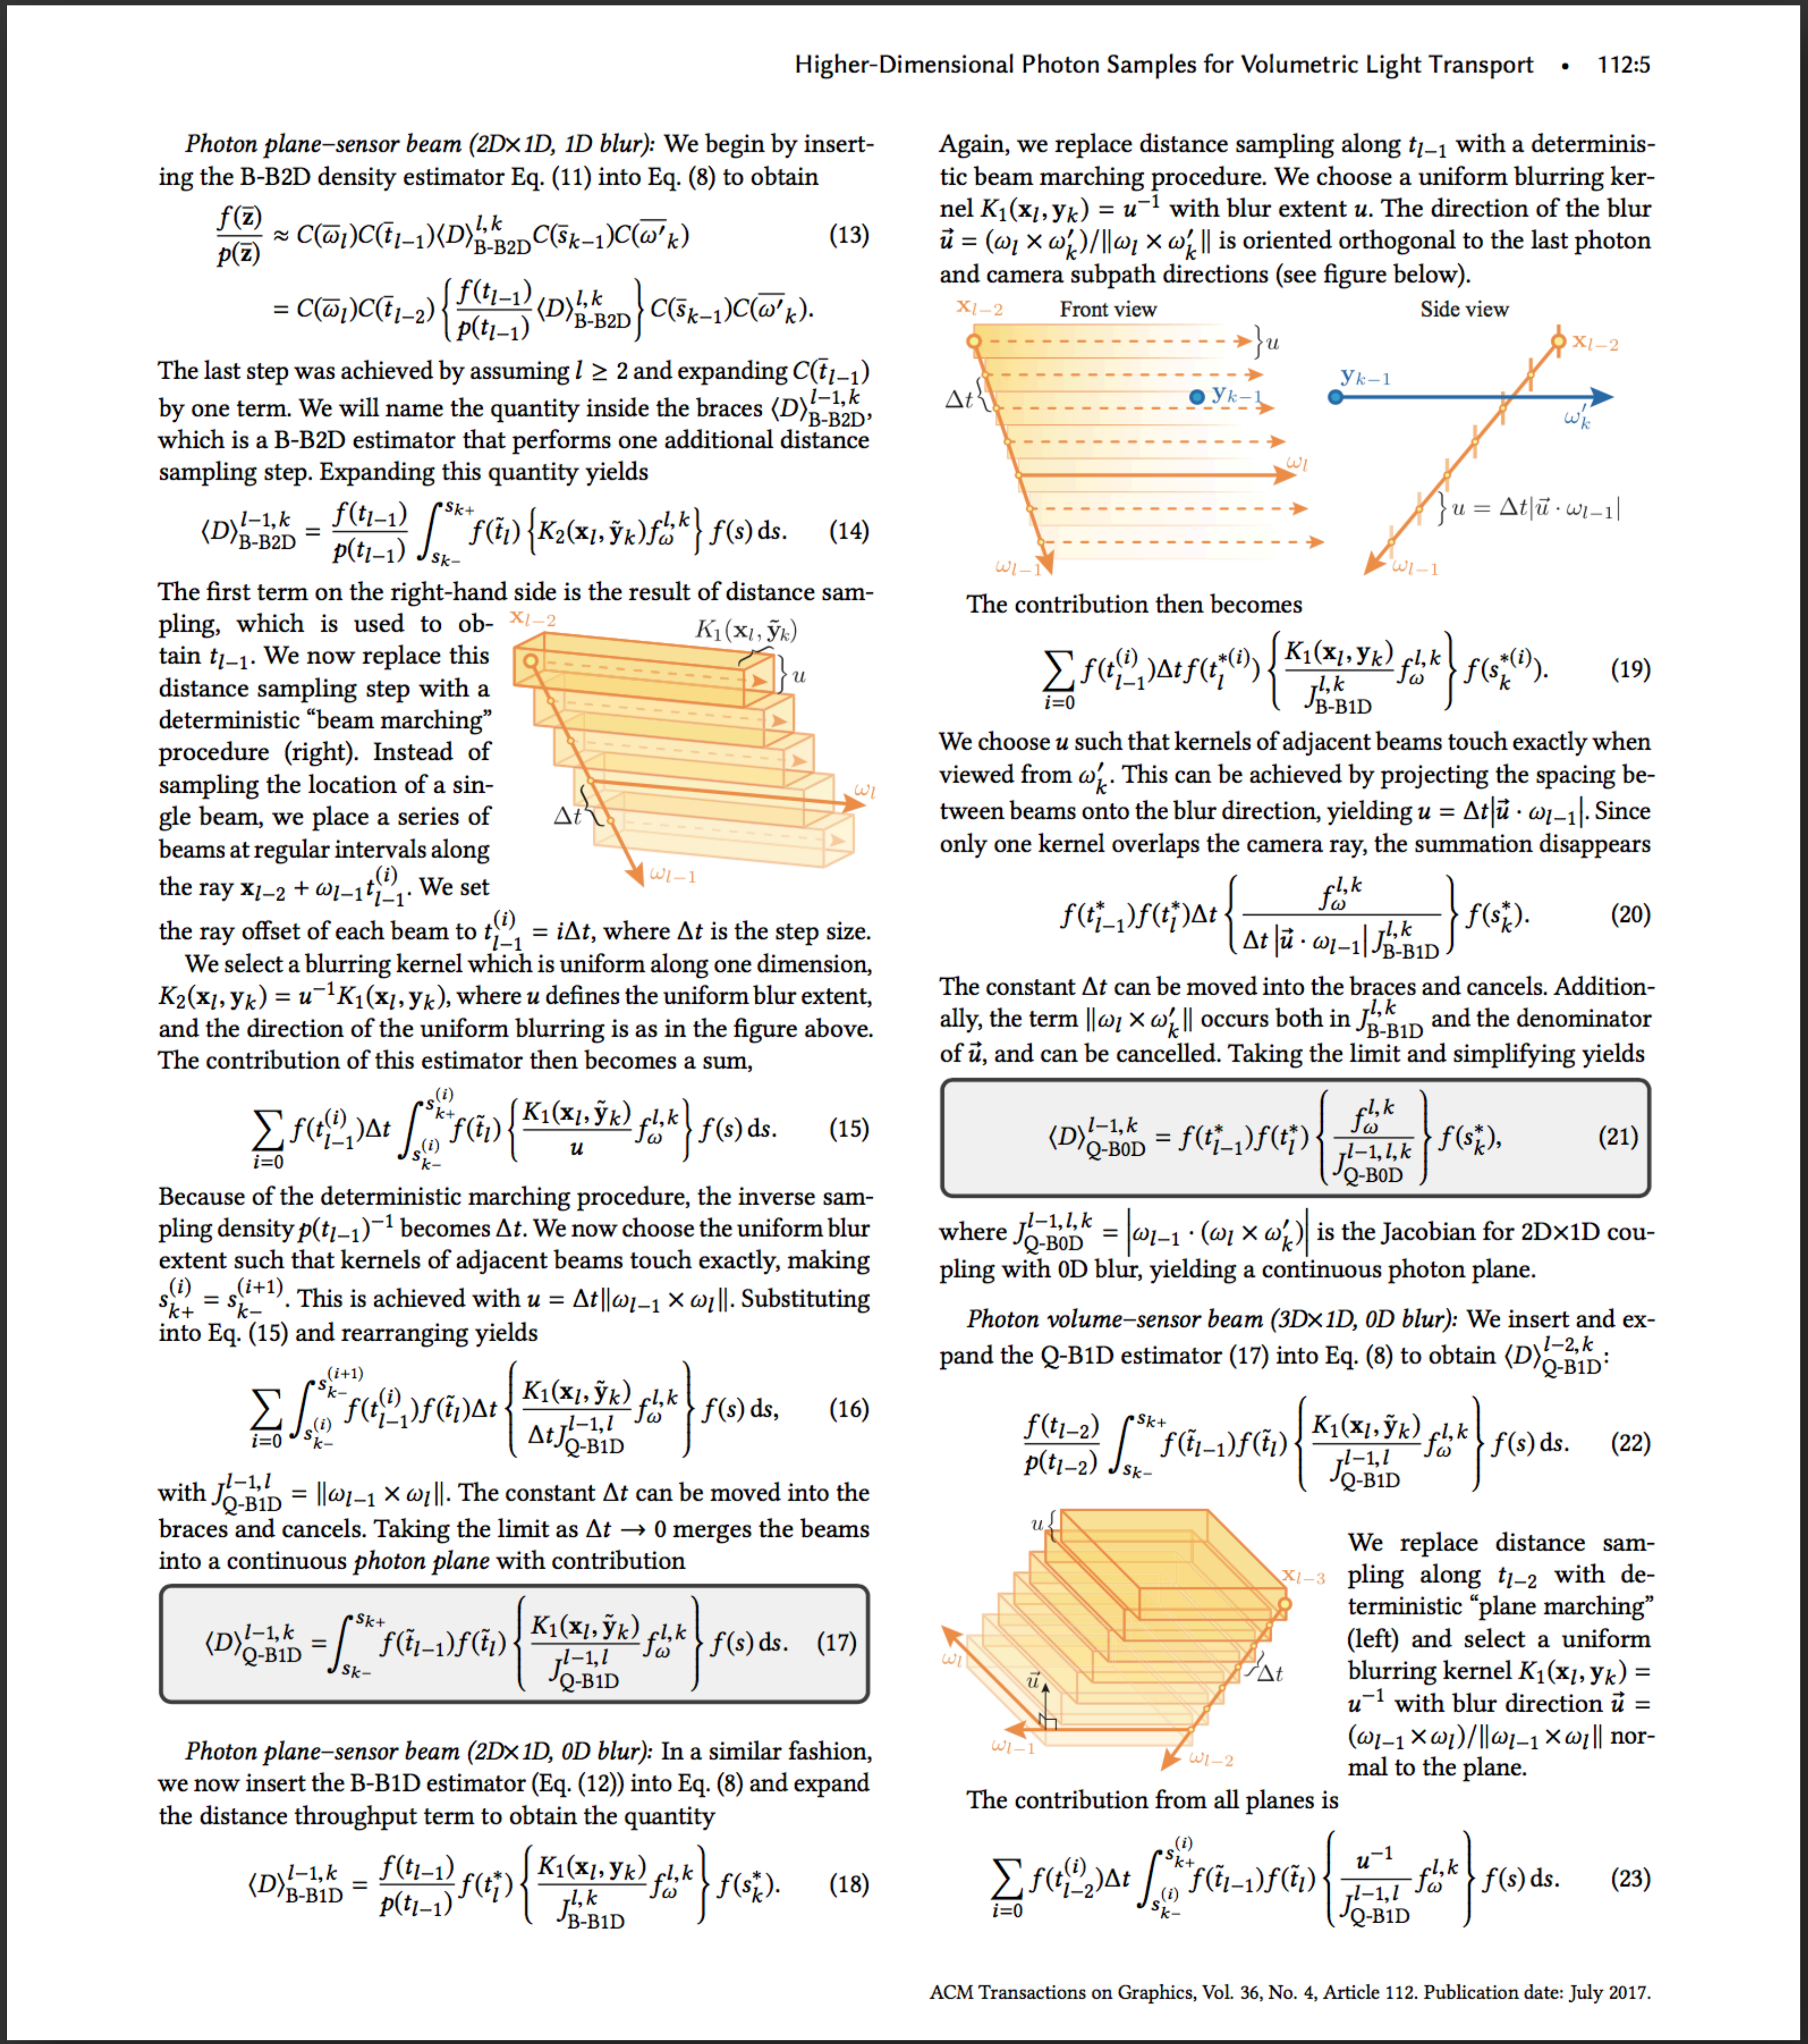

Higher-Dimensional Photon Samples for Volumetric Light Transport • 112:5

Photon plane–sensor beam (2D×1D, 1D blur): We begin by inserting the B-B2D density estimator Eq. (11) into Eq. (8) to obtain

$$\frac{f(\bar{\mathbf{z}})}{p(\bar{\mathbf{z}})} \approx C(\overline{\omega}_l)C(\bar{t}_{l-1})\langle D\rangle^{l,k}_{\text{B-B2D}}C(\bar{s}_{k-1})C(\overline{\omega'}_k) \quad (13)$$
$$= C(\overline{\omega}_l)C(\bar{t}_{l-2})\left\{\frac{f(t_{l-1})}{p(t_{l-1})}\langle D\rangle^{l,k}_{\text{B-B2D}}\right\}C(\bar{s}_{k-1})C(\overline{\omega'}_k).$$

The last step was achieved by assuming $l \geq 2$ and expanding $C(\bar{t}_{l-1})$ by one term. We will name the quantity inside the braces $\langle D\rangle^{l-1,k}_{\text{B-B2D}}$, which is a B-B2D estimator that performs one additional distance sampling step. Expanding this quantity yields

$$\langle D\rangle^{l-1,k}_{\text{B-B2D}} = \frac{f(t_{l-1})}{p(t_{l-1})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_l)\left\{K_2(\mathbf{x}_l,\tilde{\mathbf{y}}_k)f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \quad (14)$$

The first term on the right-hand side is the result of distance sampling, which is used to obtain t_{l-1}. We now replace this distance sampling step with a deterministic "beam marching" procedure (right). Instead of sampling the location of a single beam, we place a series of beams at regular intervals along the ray $\mathbf{x}_{l-2} + \omega_{l-1}t^{(i)}_{l-1}$. We set the ray offset of each beam to $t^{(i)}_{l-1} = i\Delta t$, where Δt is the step size.

We select a blurring kernel which is uniform along one dimension, $K_2(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}K_1(\mathbf{x}_l,\mathbf{y}_k)$, where u defines the uniform blur extent, and the direction of the uniform blurring is as in the figure above. The contribution of this estimator then becomes a sum,

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{u}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \quad (15)$$

Because of the deterministic marching procedure, the inverse sampling density $p(t_{l-1})^{-1}$ becomes Δt. We now choose the uniform blur extent such that kernels of adjacent beams touch exactly, making $s^{(i)}_{k+} = s^{(i+1)}_{k-}$. This is achieved with $u = \Delta t\|\omega_{l-1}\times\omega_l\|$. Substituting into Eq. (15) and rearranging yields

$$\sum_{i=0}\int_{s^{(i)}_{k-}}^{s^{(i+1)}_{k-}} f(t^{(i)}_{l-1})f(\tilde{t}_l)\Delta t\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{\Delta t J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s, \quad (16)$$

with $J^{l-1,l}_{\text{Q-B1D}} = \|\omega_{l-1}\times\omega_l\|$. The constant Δt can be moved into the braces and cancels. Taking the limit as $\Delta t \to 0$ merges the beams into a continuous *photon plane* with contribution

$$\langle D\rangle^{l-1,k}_{\text{Q-B1D}} = \int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \quad (17)$$

Photon plane–sensor beam (2D×1D, 0D blur): In a similar fashion, we now insert the B-B1D estimator (Eq. (12)) into Eq. (8) and expand the distance throughput term to obtain the quantity

$$\langle D\rangle^{l-1,k}_{\text{B-B1D}} = \frac{f(t_{l-1})}{p(t_{l-1})}f(t^*_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_\omega\right\}f(s^*_k). \quad (18)$$

Again, we replace distance sampling along t_{l-1} with a deterministic beam marching procedure. We choose a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur extent u. The direction of the blur $\vec{u} = (\omega_l\times\omega'_k)/\|\omega_l\times\omega'_k\|$ is oriented orthogonal to the last photon and camera subpath directions (see figure below).

Front view Side view

The contribution then becomes

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t f(t^{*(i)}_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_\omega\right\}f(s^{*(i)}_k). \quad (19)$$

We choose u such that kernels of adjacent beams touch exactly when viewed from ω'_k. This can be achieved by projecting the spacing between beams onto the blur direction, yielding $u = \Delta t|\vec{u}\cdot\omega_{l-1}|$. Since only one kernel overlaps the camera ray, the summation disappears

$$f(t^*_{l-1})f(t^*_l)\Delta t\left\{\frac{f^{l,k}_\omega}{\Delta t\,|\vec{u}\cdot\omega_{l-1}|\,J^{l,k}_{\text{B-B1D}}}\right\}f(s^*_k). \quad (20)$$

The constant Δt can be moved into the braces and cancels. Additionally, the term $\|\omega_l\times\omega'_k\|$ occurs both in $J^{l,k}_{\text{B-B1D}}$ and the denominator of $\vec{u}$, and can be cancelled. Taking the limit and simplifying yields

$$\langle D\rangle^{l-1,k}_{\text{Q-B0D}} = f(t^*_{l-1})f(t^*_l)\left\{\frac{f^{l,k}_\omega}{J^{l-1,l,k}_{\text{Q-B0D}}}\right\}f(s^*_k), \quad (21)$$

where $J^{l-1,l,k}_{\text{Q-B0D}} = \left|\omega_{l-1}\cdot(\omega_l\times\omega'_k)\right|$ is the Jacobian for 2D×1D coupling with 0D blur, yielding a continuous photon plane.

Photon volume–sensor beam (3D×1D, 0D blur): We insert and expand the Q-B1D estimator (17) into Eq. (8) to obtain $\langle D\rangle^{l-2,k}_{\text{Q-B1D}}$:

$$\frac{f(t_{l-2})}{p(t_{l-2})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \quad (22)$$

We replace distance sampling along t_{l-2} with deterministic "plane marching" (left) and select a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur direction $\vec{u} = (\omega_{l-1}\times\omega_l)/\|\omega_{l-1}\times\omega_l\|$ normal to the plane.

The contribution from all planes is

$$\sum_{i=0} f(t^{(i)}_{l-2})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{u^{-1}}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_\omega\right\}f(s)\,\mathrm{d}s. \quad (23)$$

ACM Transactions on Graphics, Vol. 36, No. 4, Article 112. Publication date: July 2017.

Disclaimer

- Dense Paper
- Light Talk
- Details: See paper

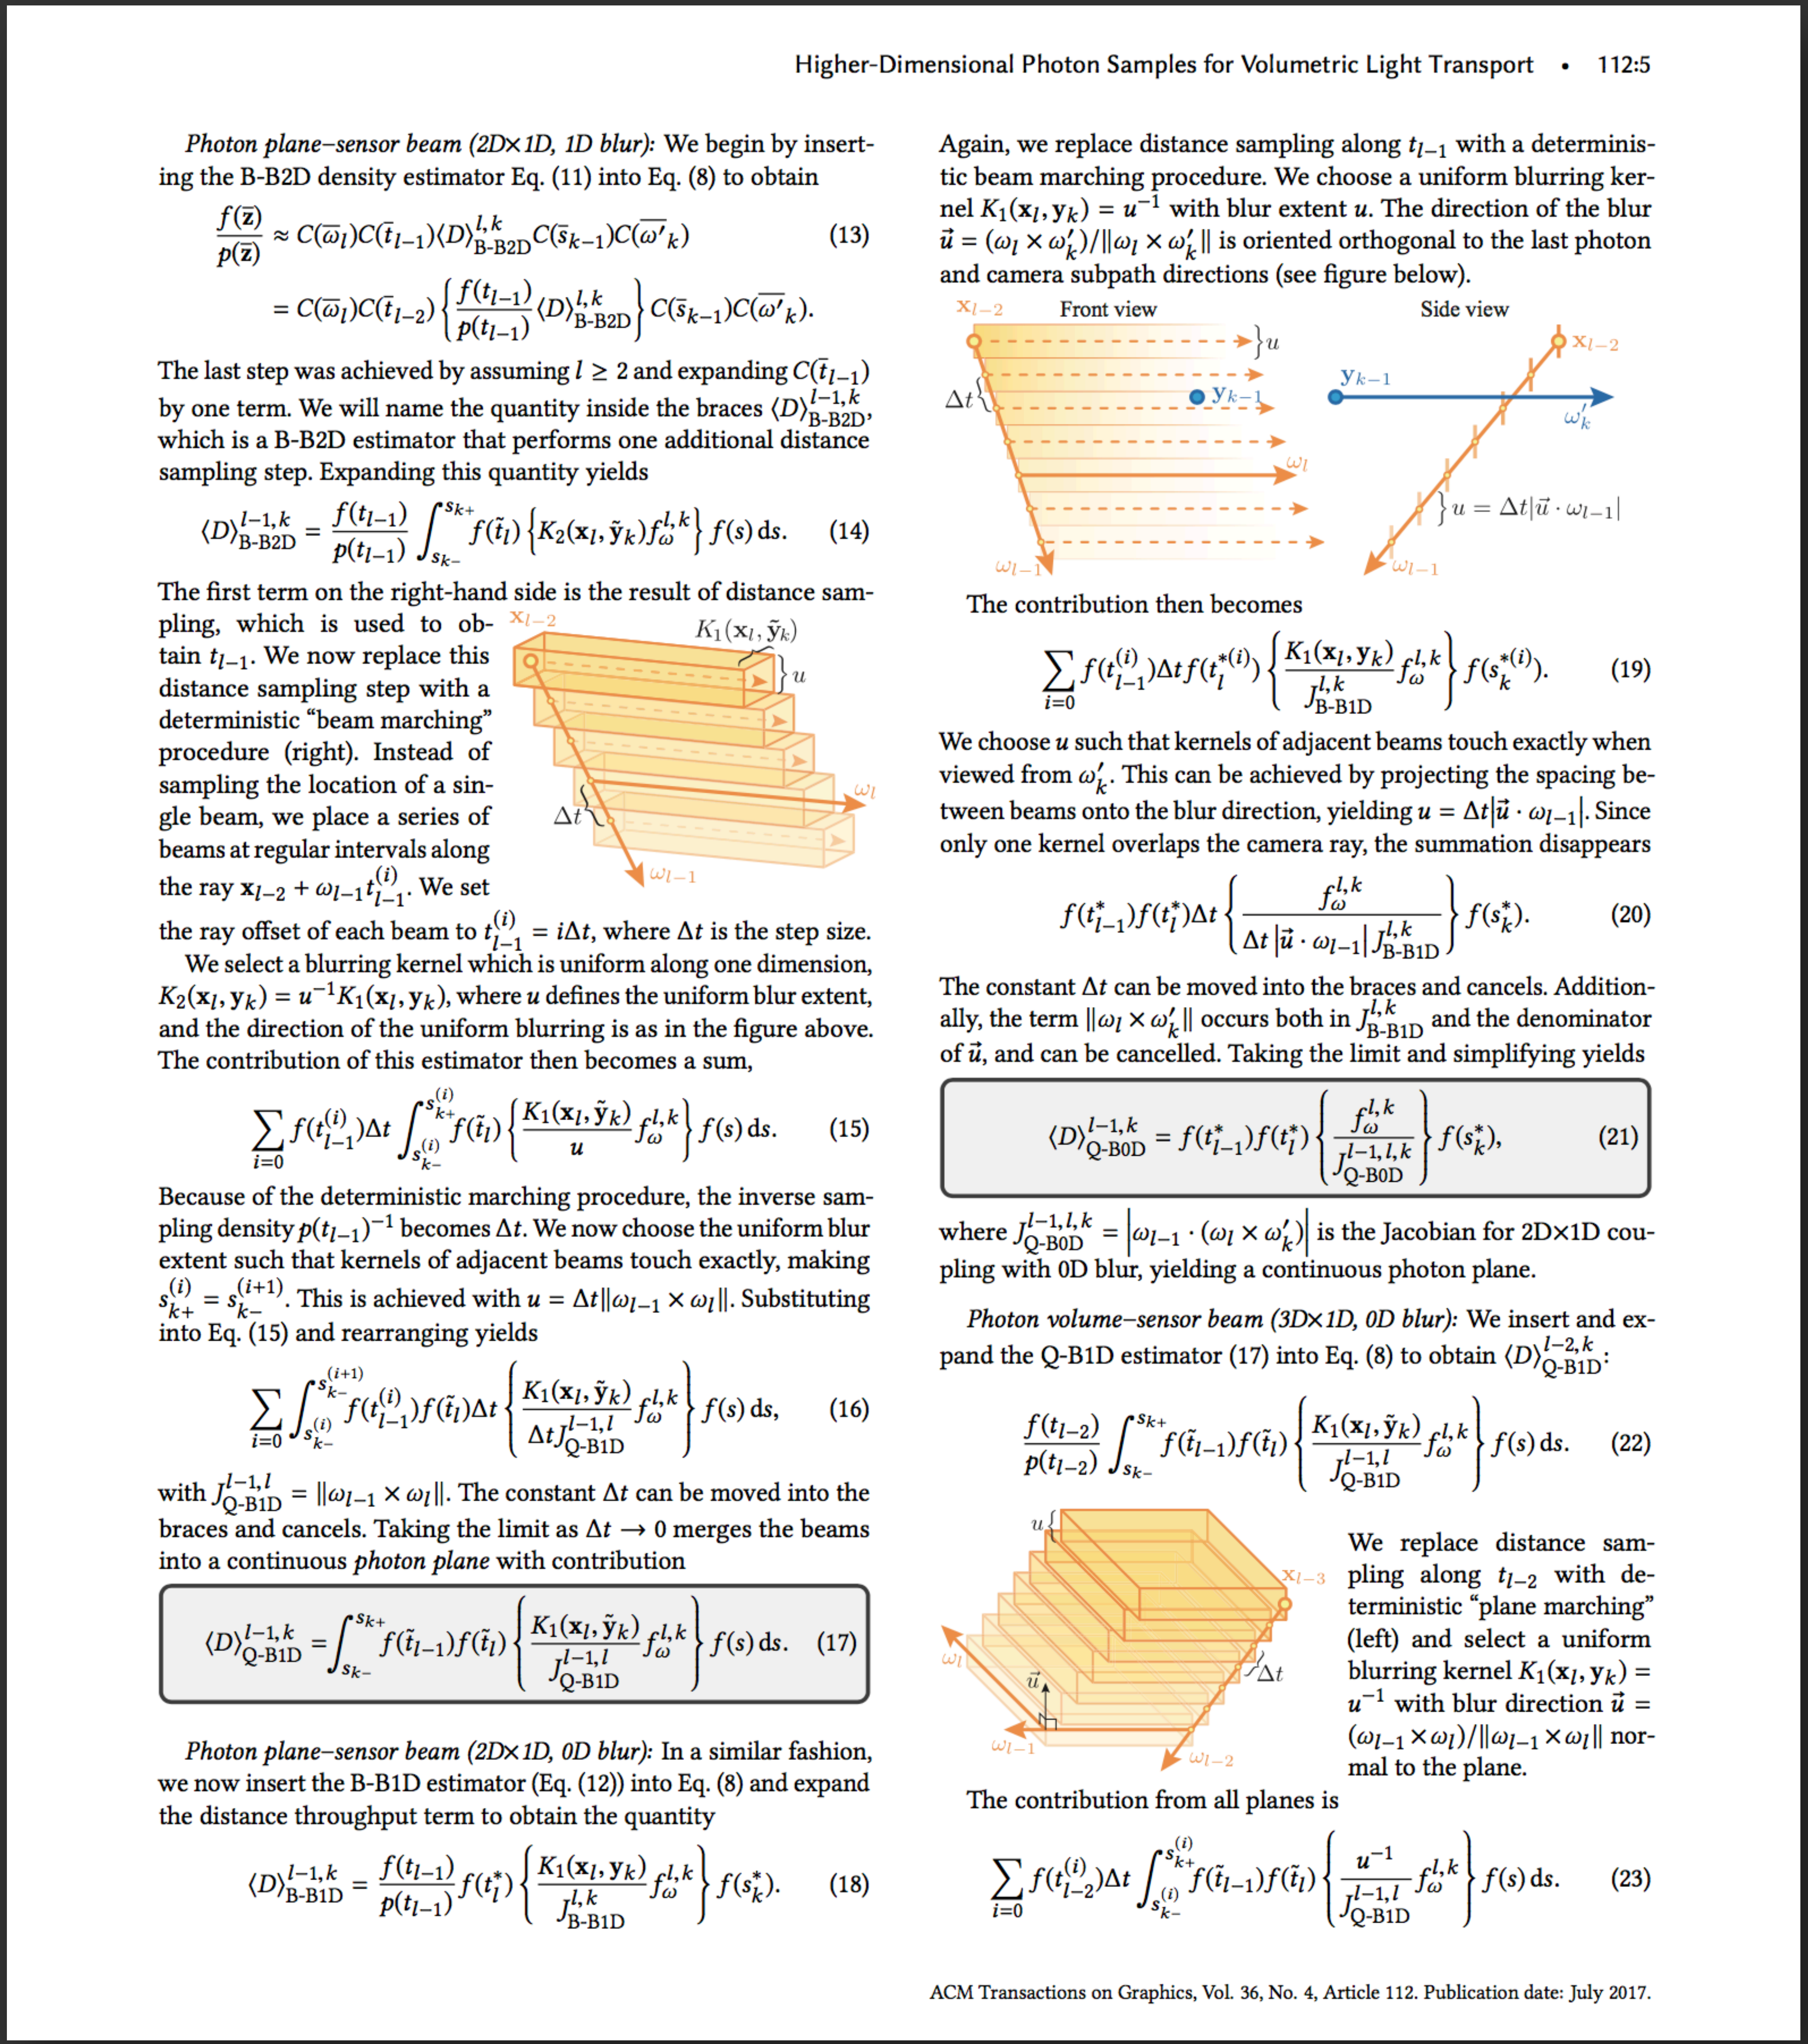

Photon plane–sensor beam (2D×1D, 1D blur): We begin by inserting the B-B2D density estimator Eq. (11) into Eq. (8) to obtain

$$\frac{f(\bar{\mathbf{z}})}{p(\bar{\mathbf{z}})} \approx C(\overline{\omega}_l)C(\bar{t}_{l-1})\langle D\rangle^{l,k}_{\text{B-B2D}}C(\bar{s}_{k-1})C(\overline{\omega'}_k) \quad (13)$$

$$= C(\overline{\omega}_l)C(\bar{t}_{l-2})\left\{\frac{f(t_{l-1})}{p(t_{l-1})}\langle D\rangle^{l,k}_{\text{B-B2D}}\right\}C(\bar{s}_{k-1})C(\overline{\omega'}_k).$$

The last step was achieved by assuming $l \geq 2$ and expanding $C(\bar{t}_{l-1})$ by one term. We will name the quantity inside the braces $\langle D\rangle^{l-1,k}_{\text{B-B2D}}$, which is a B-B2D estimator that performs one additional distance sampling step. Expanding this quantity yields

$$\langle D\rangle^{l-1,k}_{\text{B-B2D}} = \frac{f(t_{l-1})}{p(t_{l-1})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_l)\left\{K_2(\mathbf{x}_l,\tilde{\mathbf{y}}_k)f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (14)$$

The first term on the right-hand side is the result of distance sampling, which is used to obtain t_{l-1}. We now replace this distance sampling step with a deterministic "beam marching" procedure (right). Instead of sampling the location of a single beam, we place a series of beams at regular intervals along the ray $\mathbf{x}_{l-2} + \omega_{l-1}t^{(i)}_{l-1}$. We set the ray offset of each beam to $t^{(i)}_{l-1} = i\Delta t$, where Δt is the step size.

We select a blurring kernel which is uniform along one dimension, $K_2(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}K_1(\mathbf{x}_l,\mathbf{y}_k)$, where u defines the uniform blur extent, and the direction of the uniform blurring is as in the figure above. The contribution of this estimator then becomes a sum,

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{u}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (15)$$

Because of the deterministic marching procedure, the inverse sampling density $p(t_{l-1})^{-1}$ becomes Δt. We now choose the uniform blur extent such that kernels of adjacent beams touch exactly, making $s^{(i)}_{k+} = s^{(i+1)}_{k-}$. This is achieved with $u = \Delta t\|\omega_{l-1}\times\omega_l\|$. Substituting into Eq. (15) and rearranging yields

$$\sum_{i=0}\int_{s^{(i)}_{k-}}^{s^{(i+1)}_{k-}} f(t^{(i)}_{l-1})f(\tilde{t}_l)\Delta t\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{\Delta t J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s, \quad (16)$$

with $J^{l-1,l}_{\text{Q-B1D}} = \|\omega_{l-1}\times\omega_l\|$. The constant Δt can be moved into the braces and cancels. Taking the limit as $\Delta t \to 0$ merges the beams into a continuous *photon plane* with contribution

$$\langle D\rangle^{l-1,k}_{\text{Q-B1D}} = \int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (17)$$

Photon plane–sensor beam (2D×1D, 0D blur): In a similar fashion, we now insert the B-B1D estimator (Eq. (12)) into Eq. (8) and expand the distance throughput term to obtain the quantity

$$\langle D\rangle^{l-1,k}_{\text{B-B1D}} = \frac{f(t_{l-1})}{p(t_{l-1})}f(t^*_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_{\omega}\right\}f(s^*_k). \quad (18)$$

Again, we replace distance sampling along t_{l-1} with a deterministic beam marching procedure. We choose a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur extent u. The direction of the blur $\vec{u} = (\omega_l\times\omega'_k)/\|\omega_l\times\omega'_k\|$ is oriented orthogonal to the last photon and camera subpath directions (see figure below).

The contribution then becomes

$$\sum_{i=0} f(t^{(i)}_{l-1})\Delta t f(t^{*(i)}_l)\left\{\frac{K_1(\mathbf{x}_l,\mathbf{y}_k)}{J^{l,k}_{\text{B-B1D}}}f^{l,k}_{\omega}\right\}f(s^{*(i)}_k). \quad (19)$$

We choose u such that kernels of adjacent beams touch exactly when viewed from ω'_k. This can be achieved by projecting the spacing between beams onto the blur direction, yielding $u = \Delta t|\vec{u}\cdot\omega_{l-1}|$. Since only one kernel overlaps the camera ray, the summation disappears

$$f(t^*_{l-1})f(t^*_l)\Delta t\left\{\frac{f^{l,k}_{\omega}}{\Delta t\,|\vec{u}\cdot\omega_{l-1}|\,J^{l,k}_{\text{B-B1D}}}\right\}f(s^*_k). \quad (20)$$

The constant Δt can be moved into the braces and cancels. Additionally, the term $\|\omega_l\times\omega'_k\|$ occurs both in $J^{l,k}_{\text{B-B1D}}$ and the denominator of $\vec{u}$, and can be cancelled. Taking the limit and simplifying yields

$$\langle D\rangle^{l-1,k}_{\text{Q-B0D}} = f(t^*_{l-1})f(t^*_l)\left\{\frac{f^{l,k}_{\omega}}{J^{l-1,l,k}_{\text{Q-B0D}}}\right\}f(s^*_k), \quad (21)$$

where $J^{l-1,l,k}_{\text{Q-B0D}} = \left|\omega_{l-1}\cdot(\omega_l\times\omega'_k)\right|$ is the Jacobian for 2D×1D coupling with 0D blur, yielding a continuous photon plane.

Photon volume–sensor beam (3D×1D, 0D blur): We insert and expand the Q-B1D estimator (17) into Eq. (8) to obtain $\langle D\rangle^{l-2,k}_{\text{Q-B1D}}$:

$$\frac{f(t_{l-2})}{p(t_{l-2})}\int_{s_{k-}}^{s_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{K_1(\mathbf{x}_l,\tilde{\mathbf{y}}_k)}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (22)$$

We replace distance sampling along t_{l-2} with deterministic "plane marching" (left) and select a uniform blurring kernel $K_1(\mathbf{x}_l,\mathbf{y}_k) = u^{-1}$ with blur direction $\vec{u} = (\omega_{l-1}\times\omega_l)/\|\omega_{l-1}\times\omega_l\|$ normal to the plane.

The contribution from all planes is

$$\sum_{i=0} f(t^{(i)}_{l-2})\Delta t\int_{s^{(i)}_{k-}}^{s^{(i)}_{k+}} f(\tilde{t}_{l-1})f(\tilde{t}_l)\left\{\frac{u^{-1}}{J^{l-1,l}_{\text{Q-B1D}}}f^{l,k}_{\omega}\right\}f(s)\,\mathrm{d}s. \quad (23)$$

Generalized Theory

Background

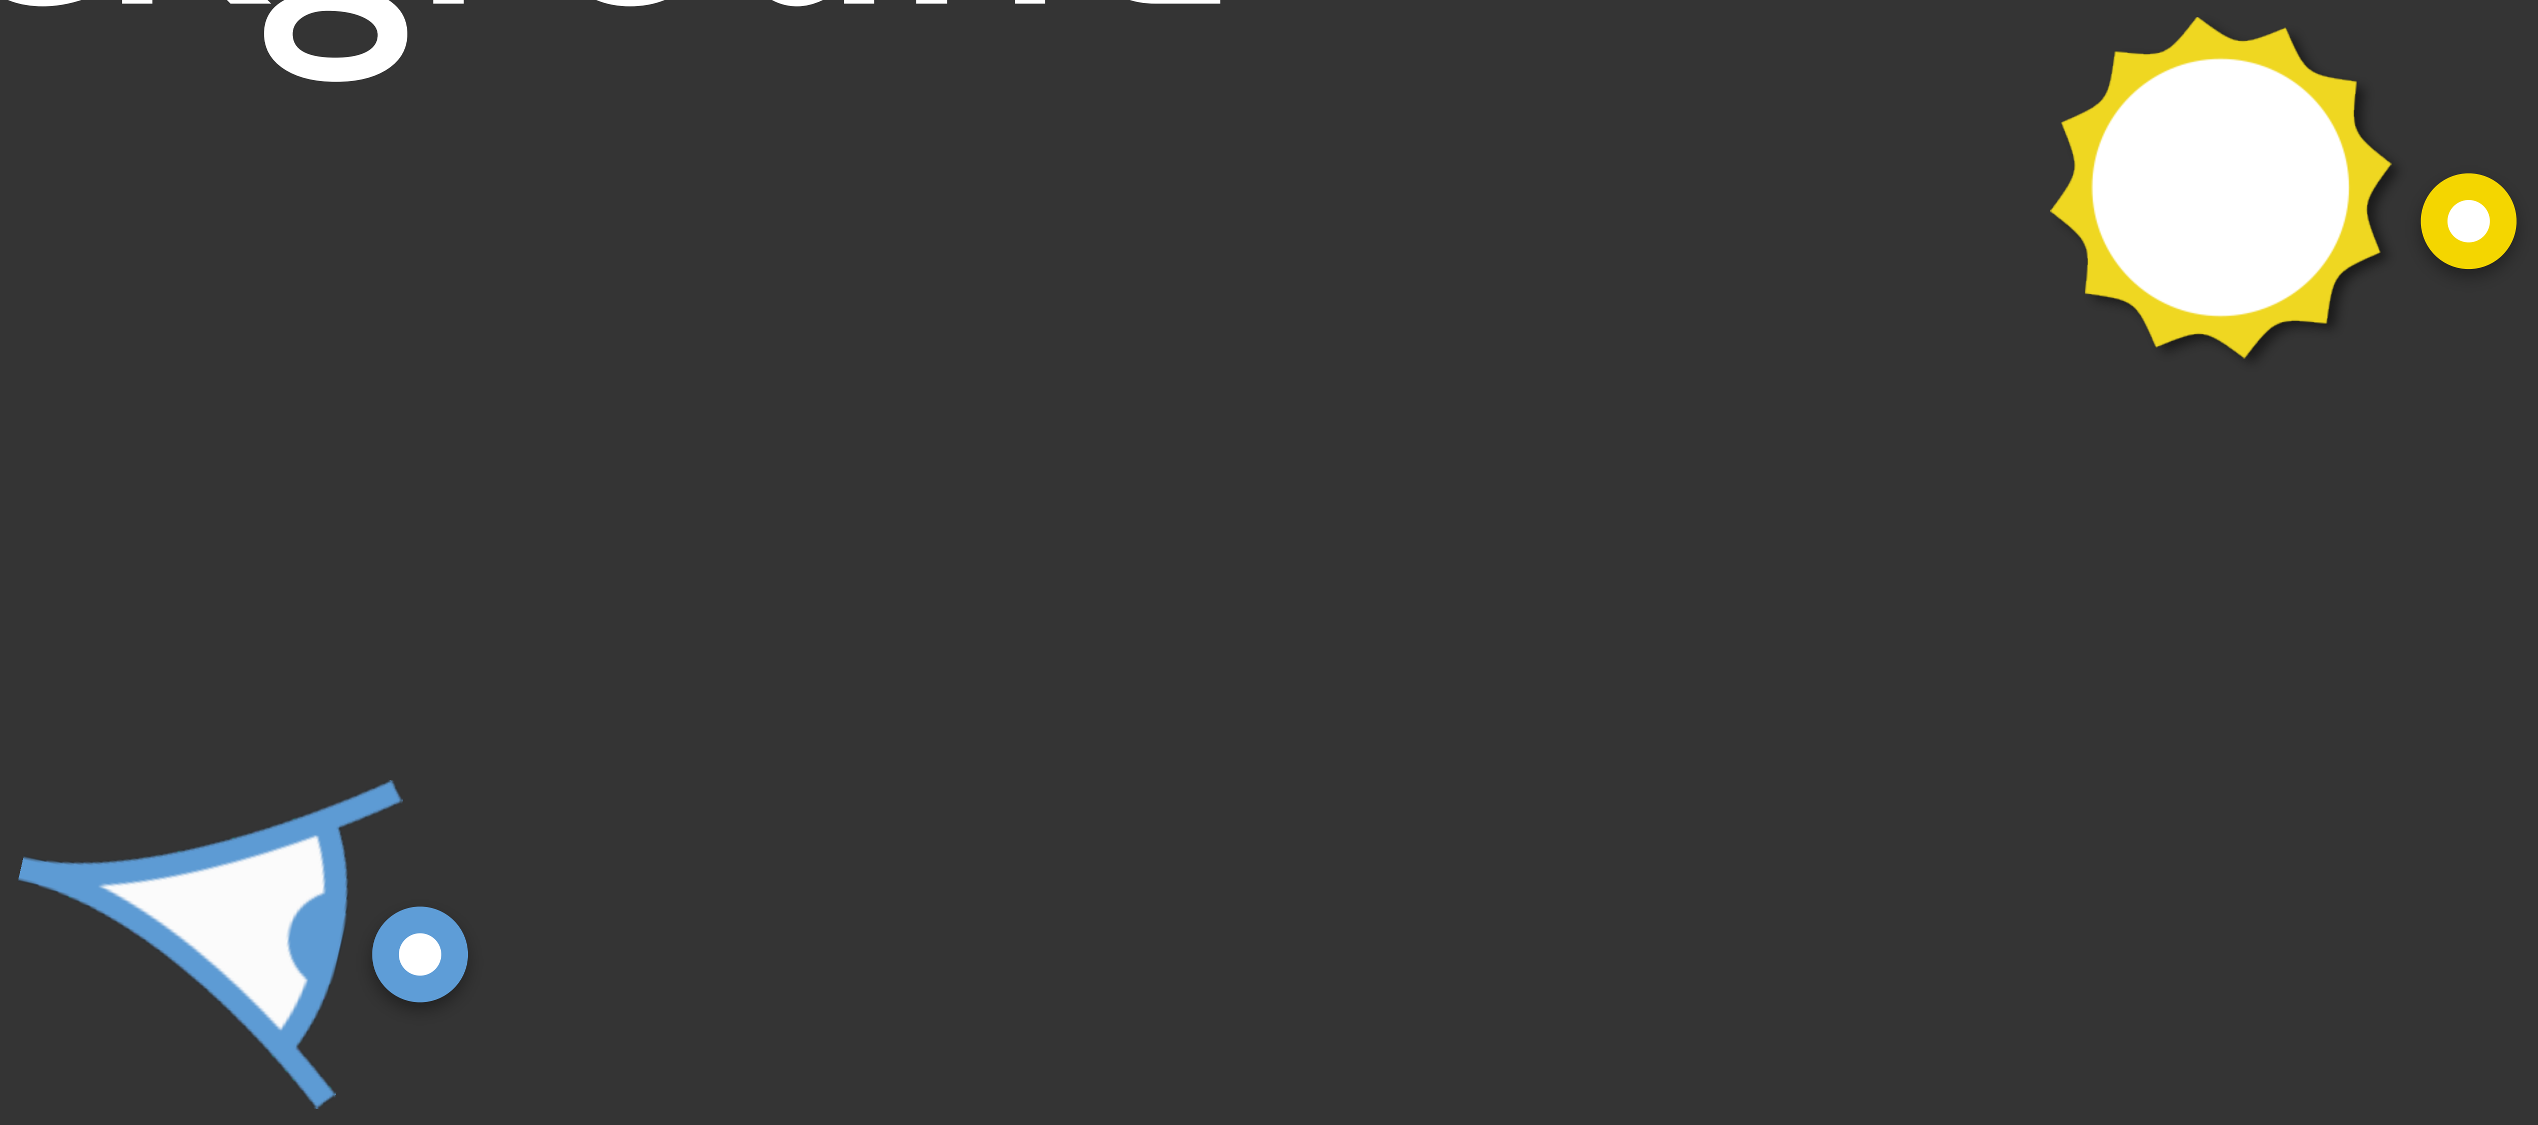

Background

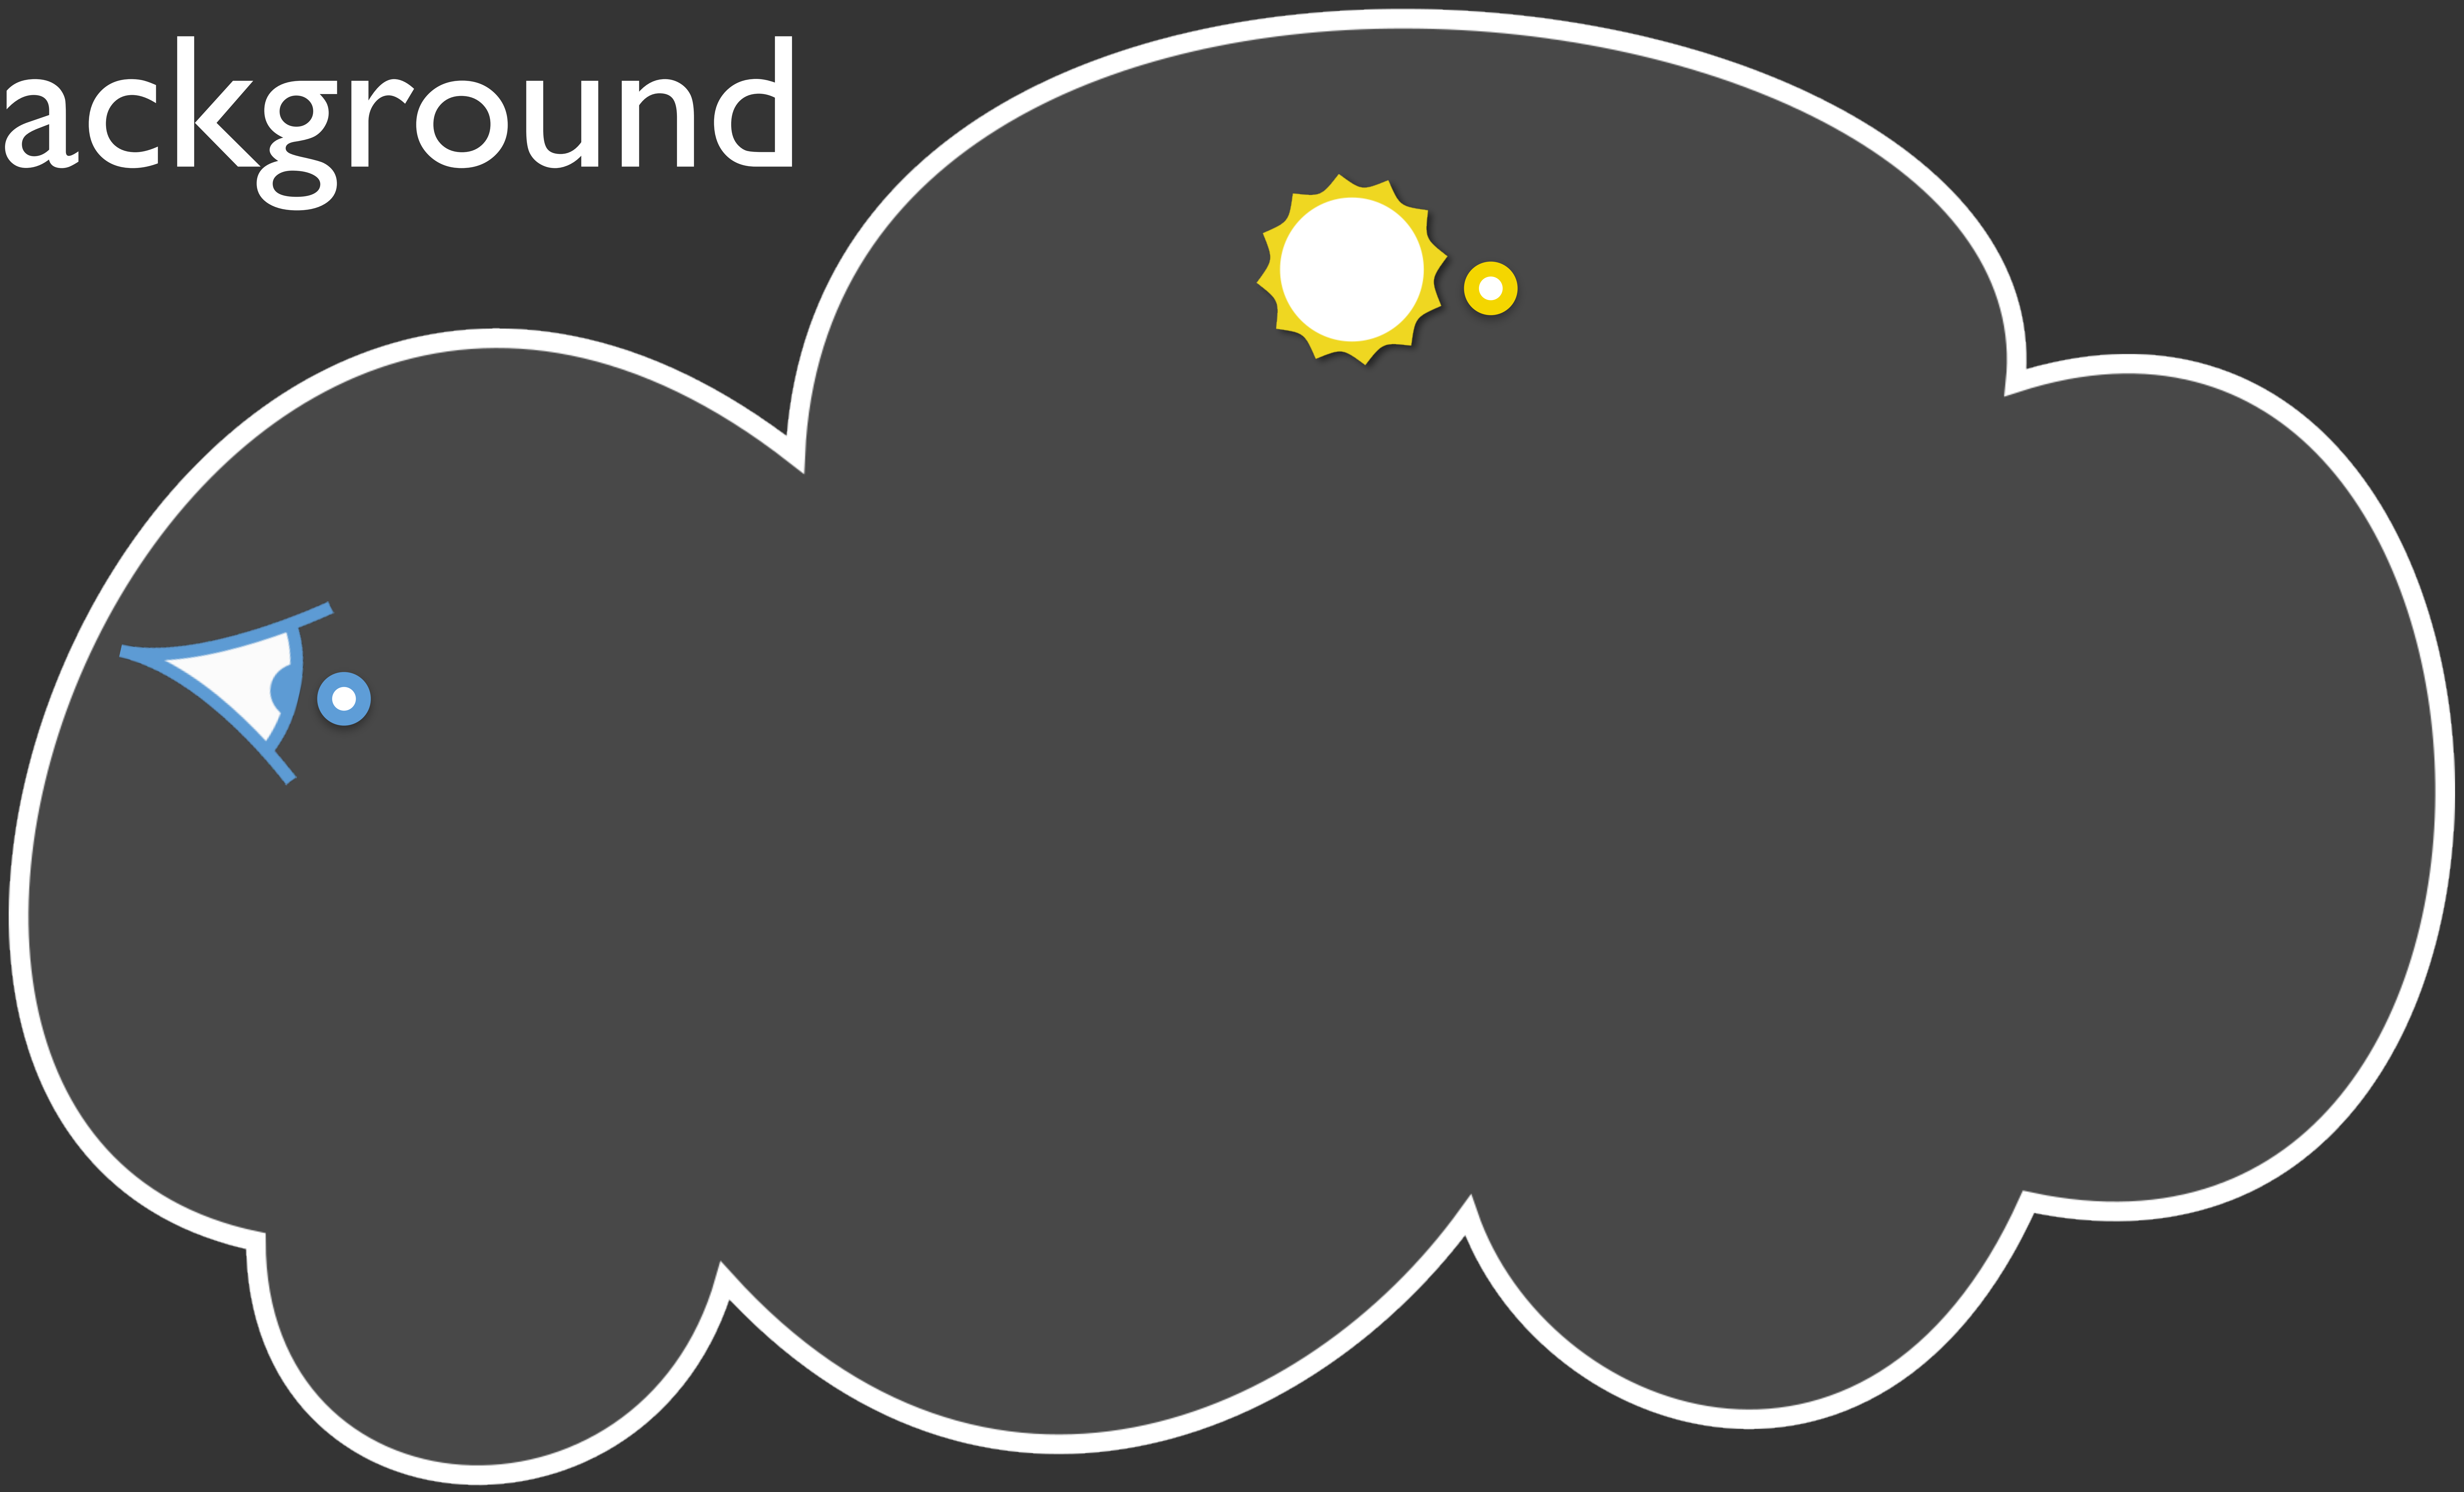

Background

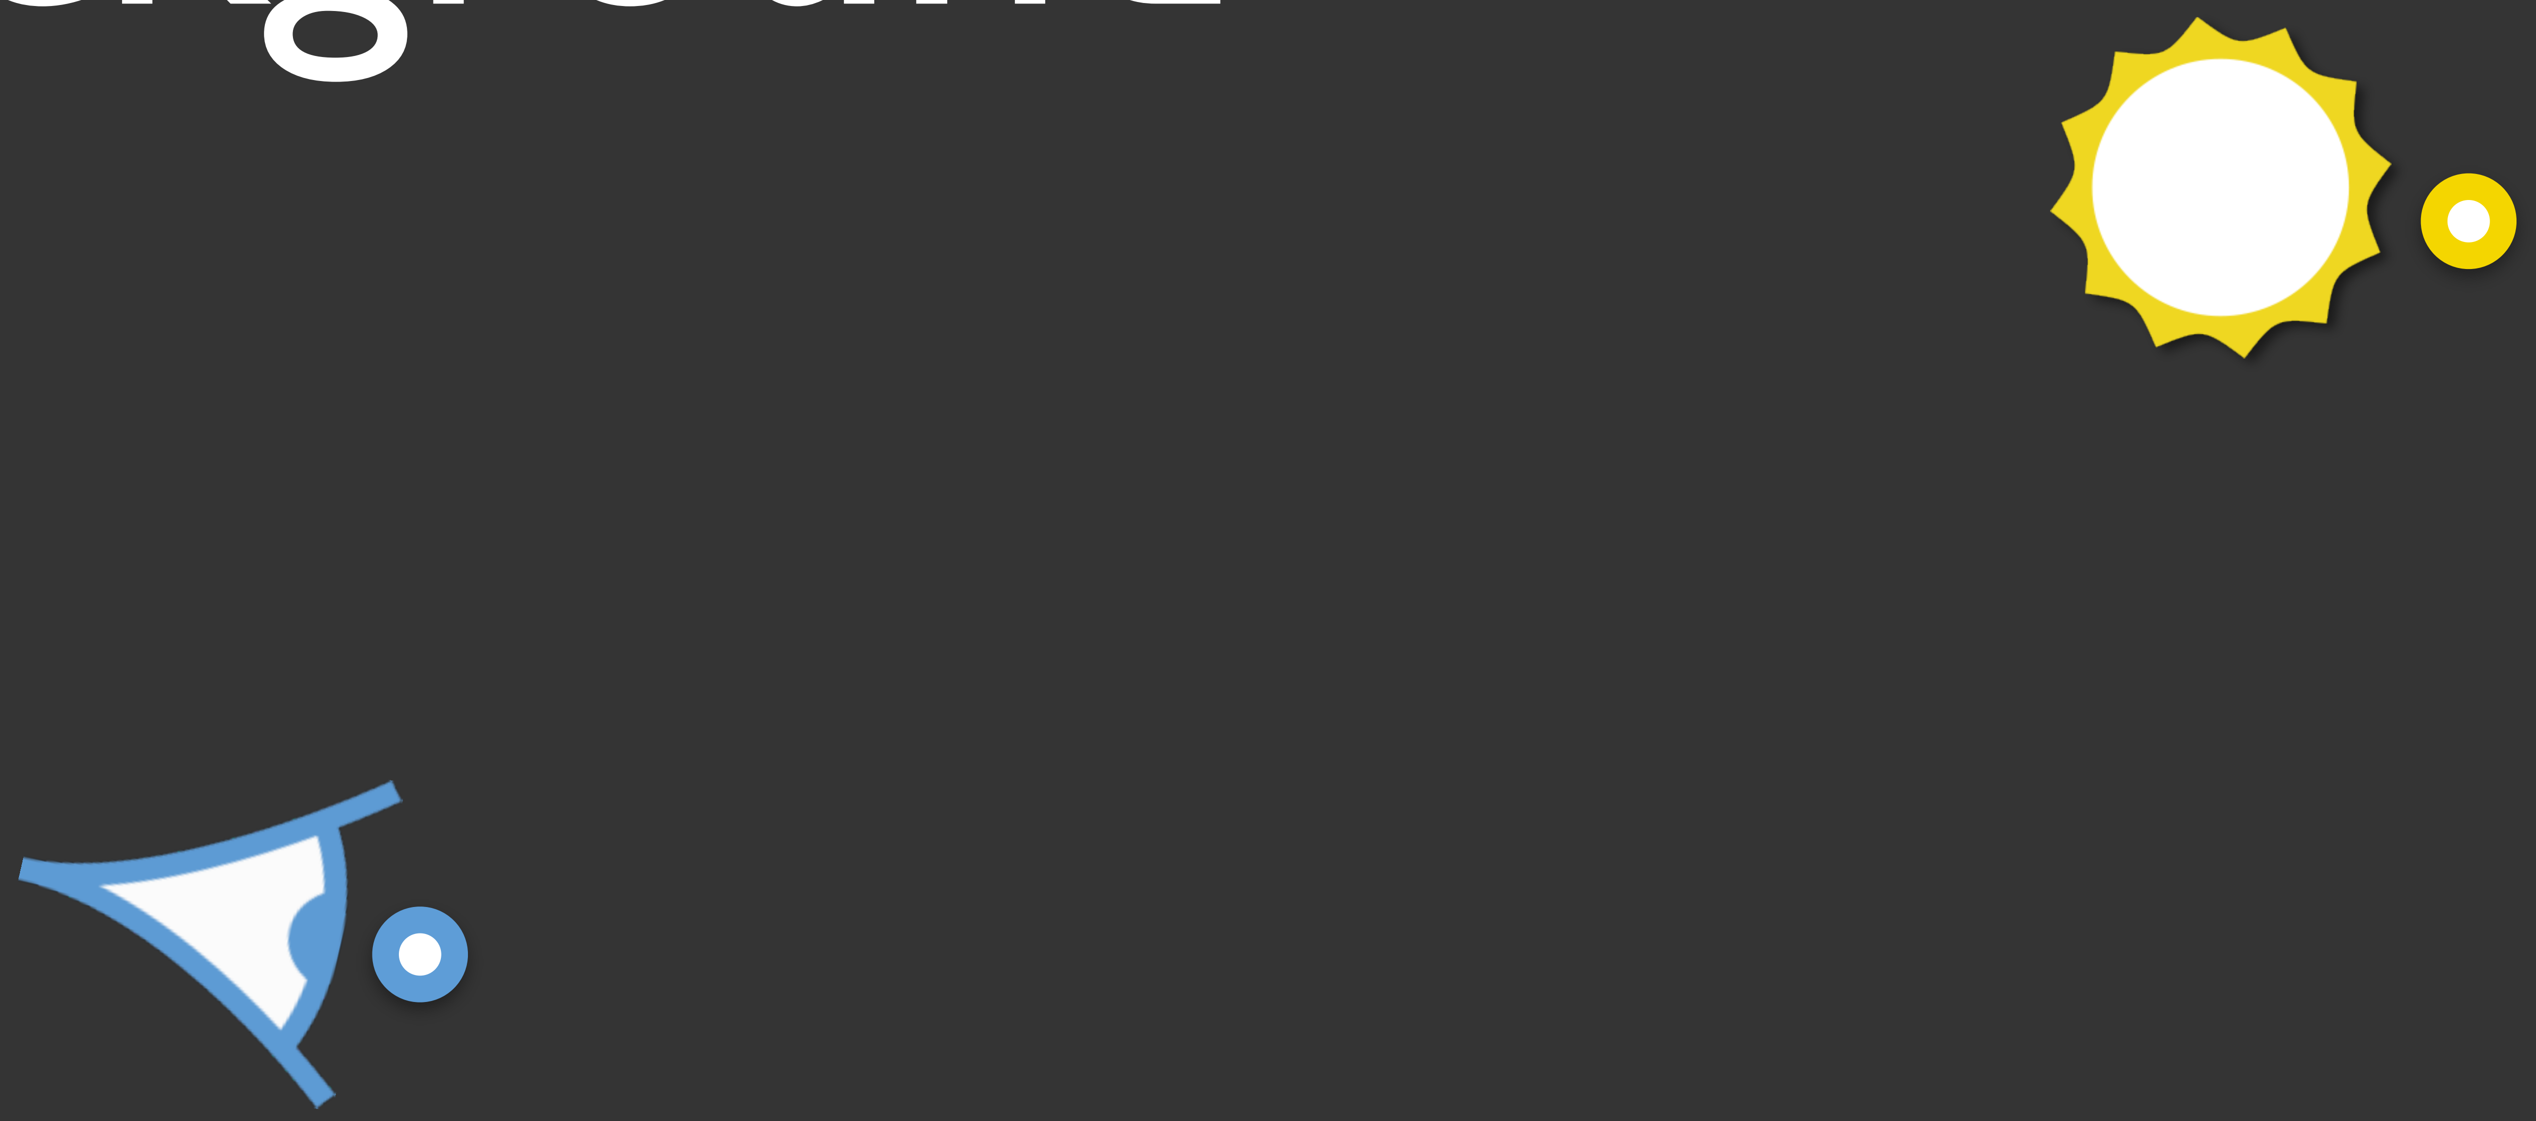

Background

Background

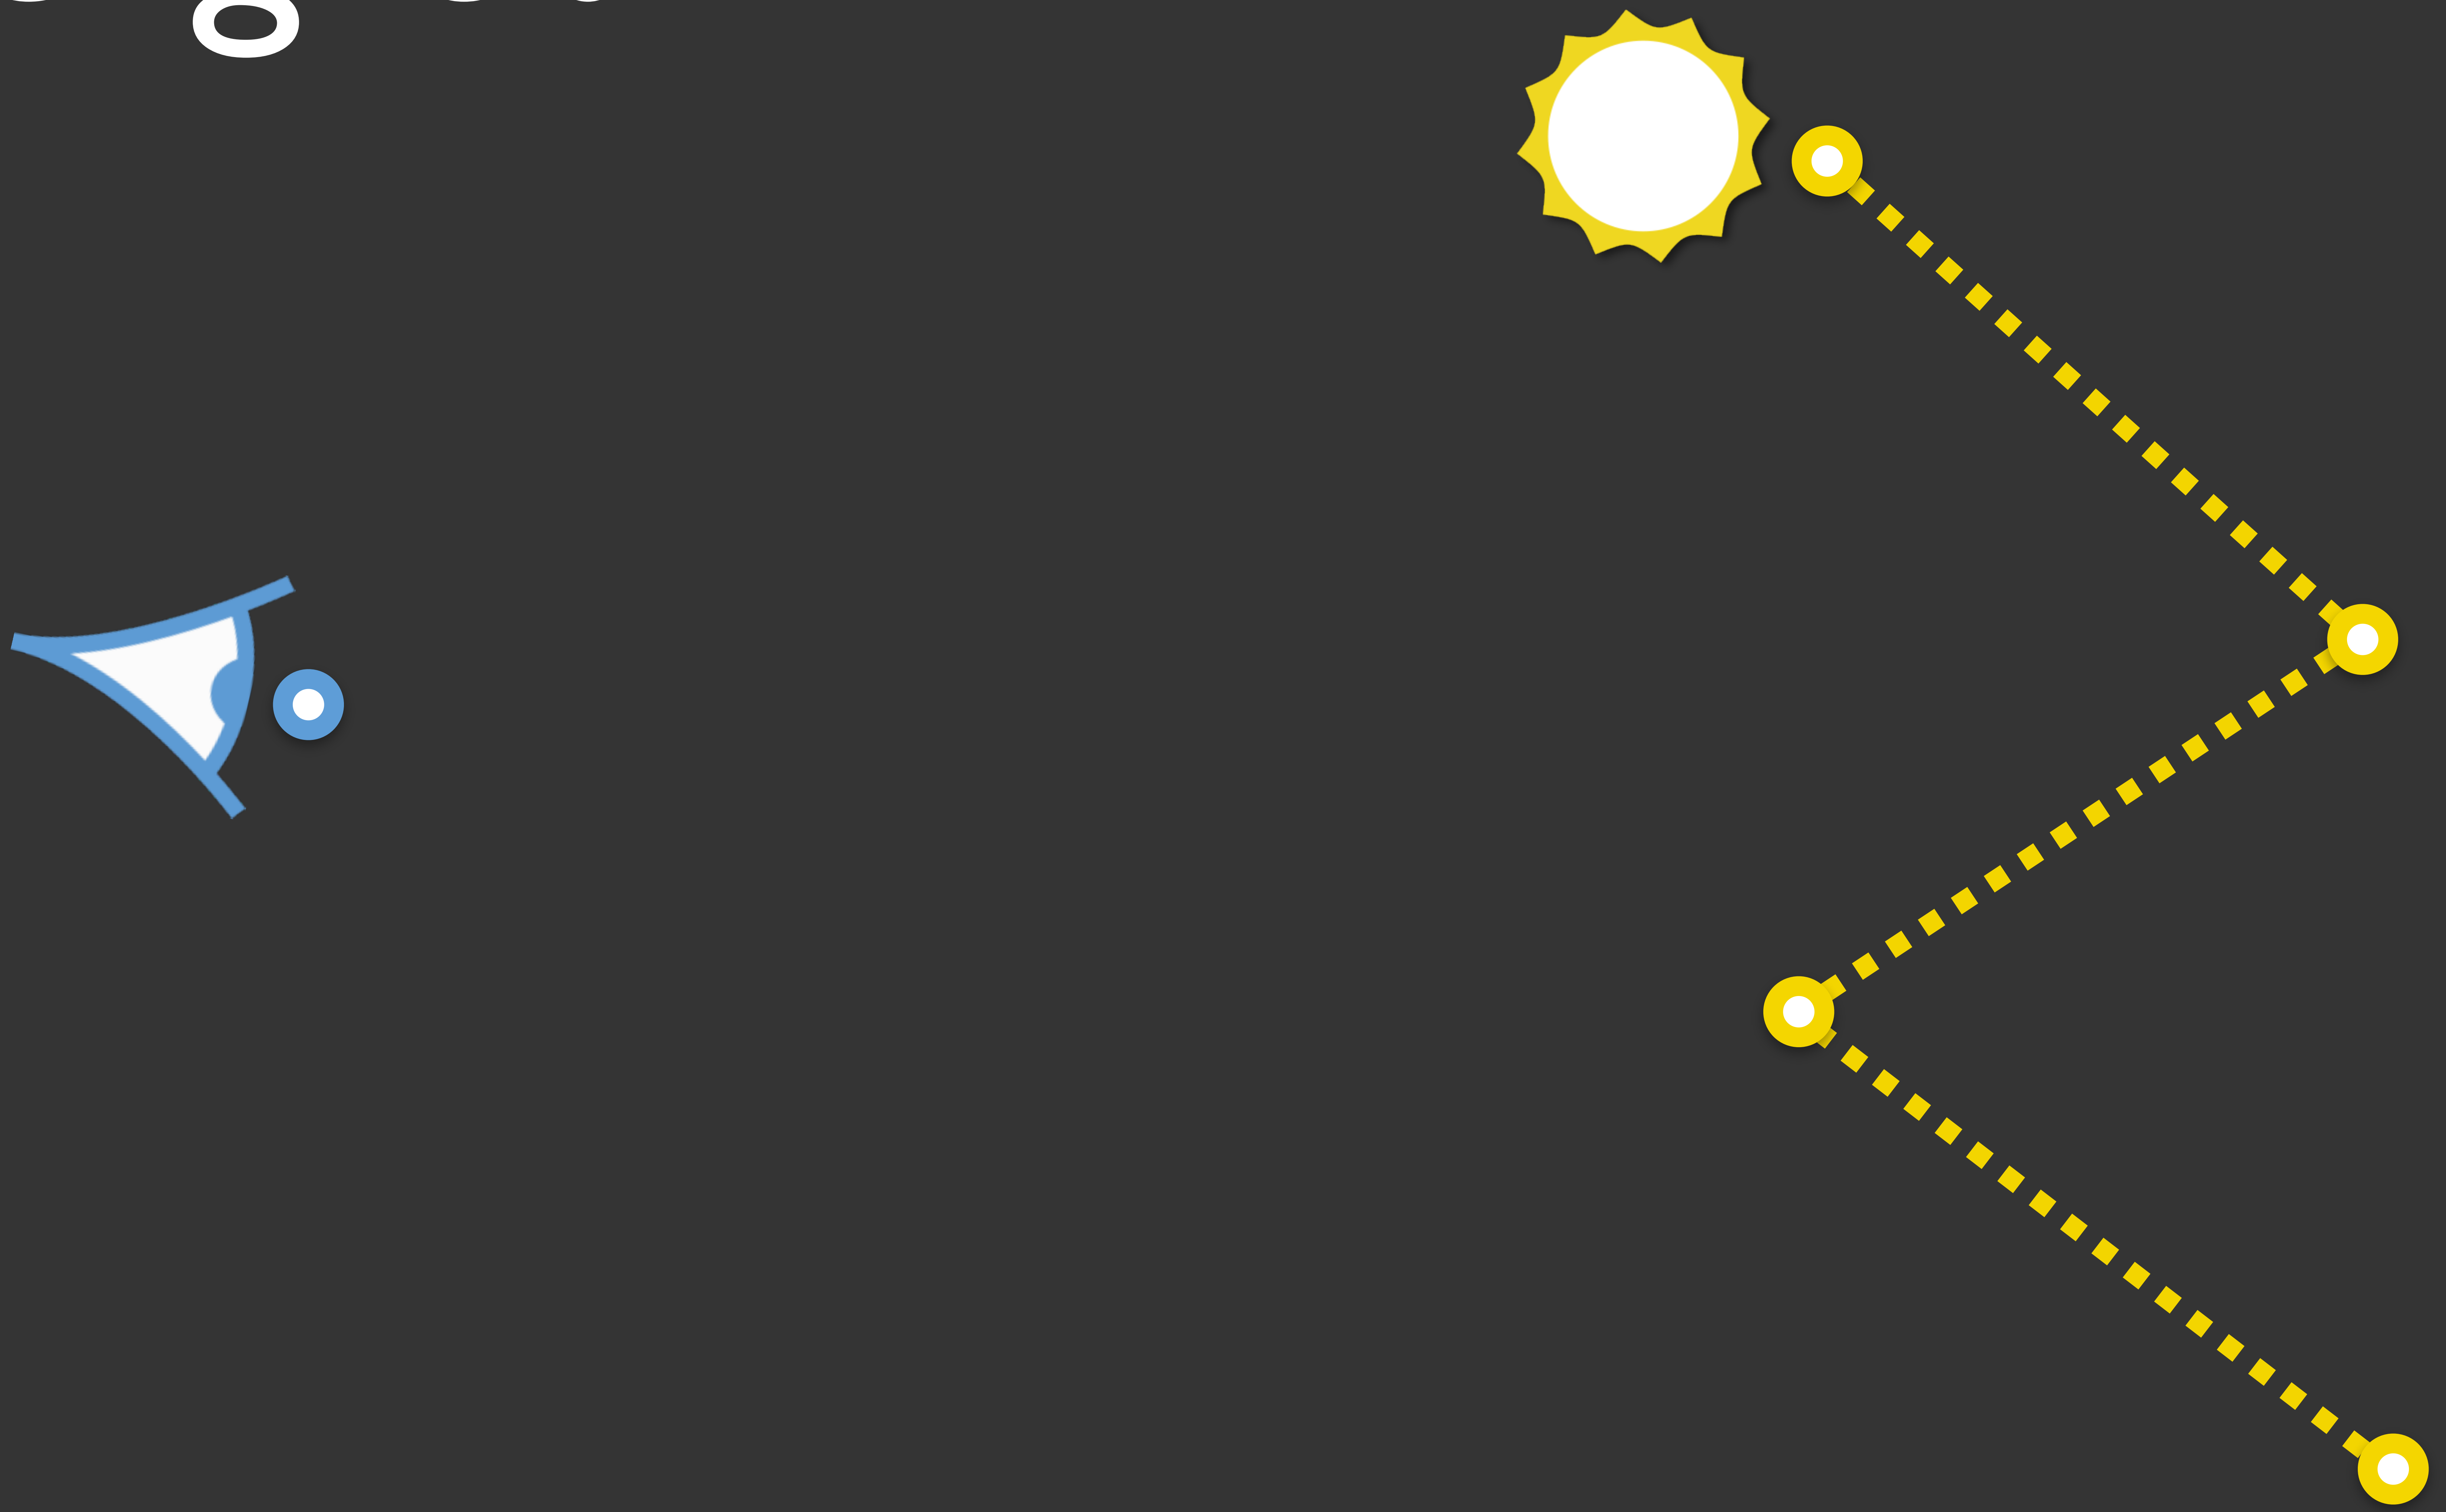

Background

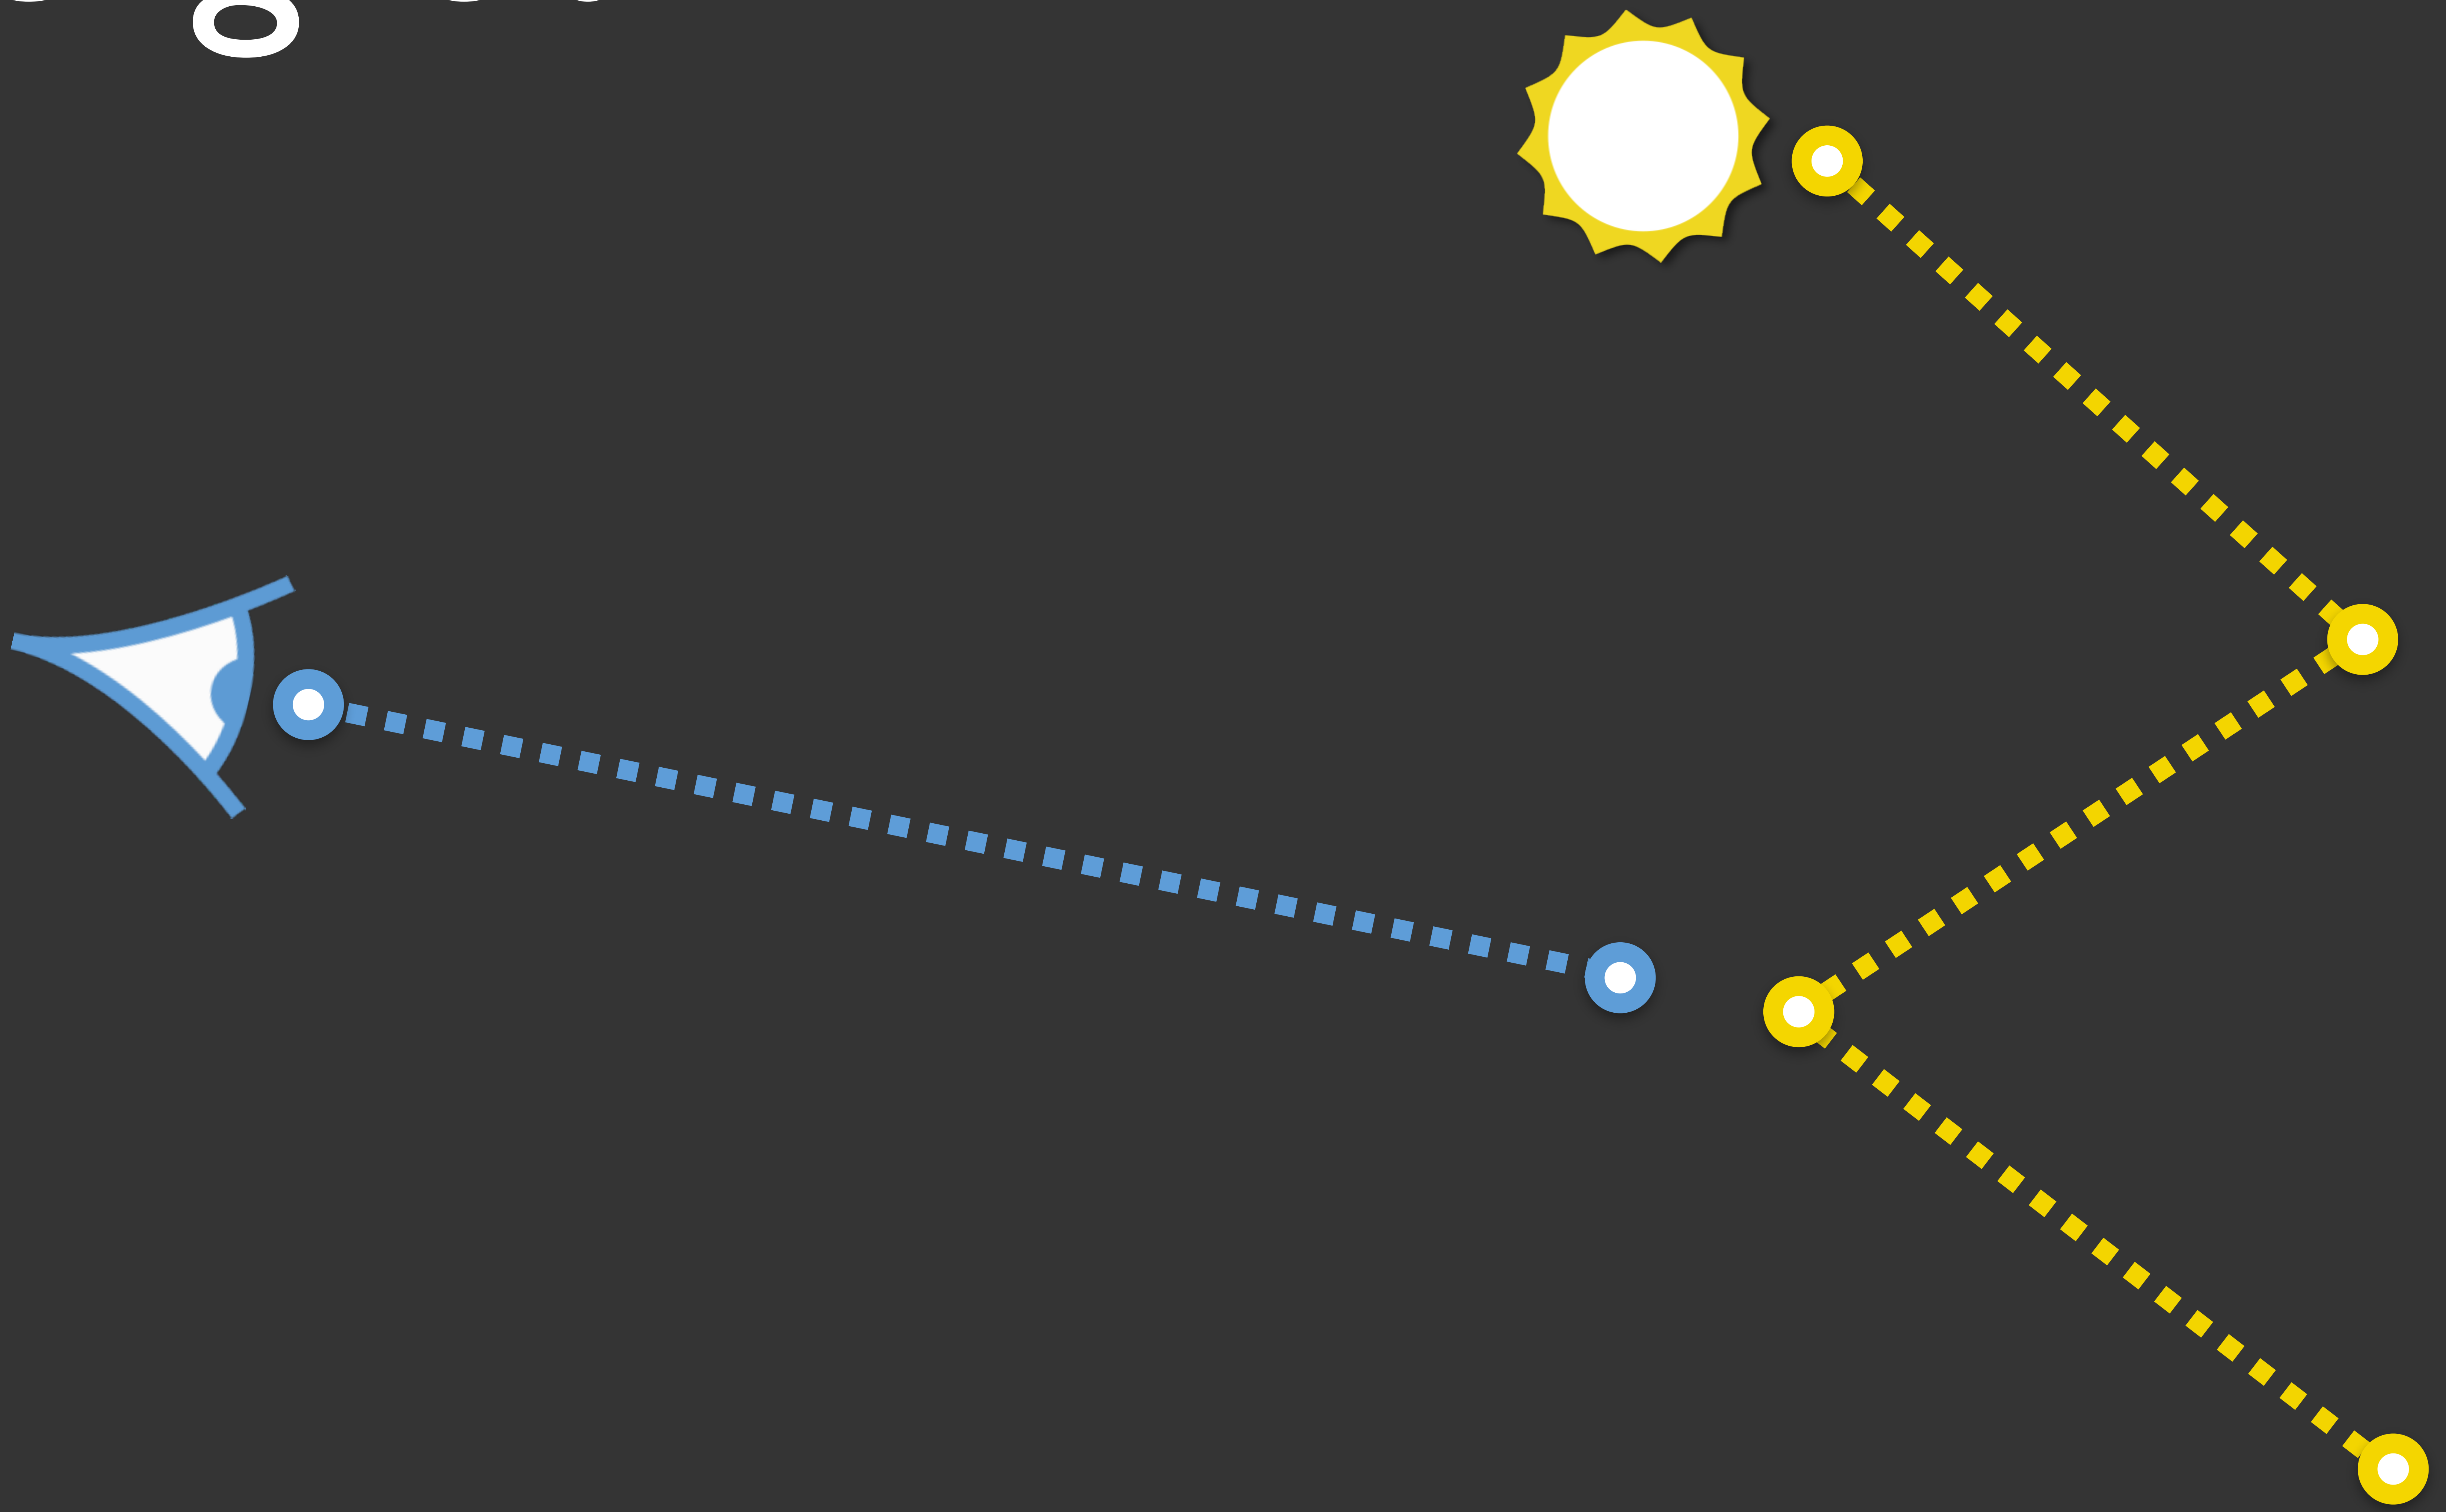

Background

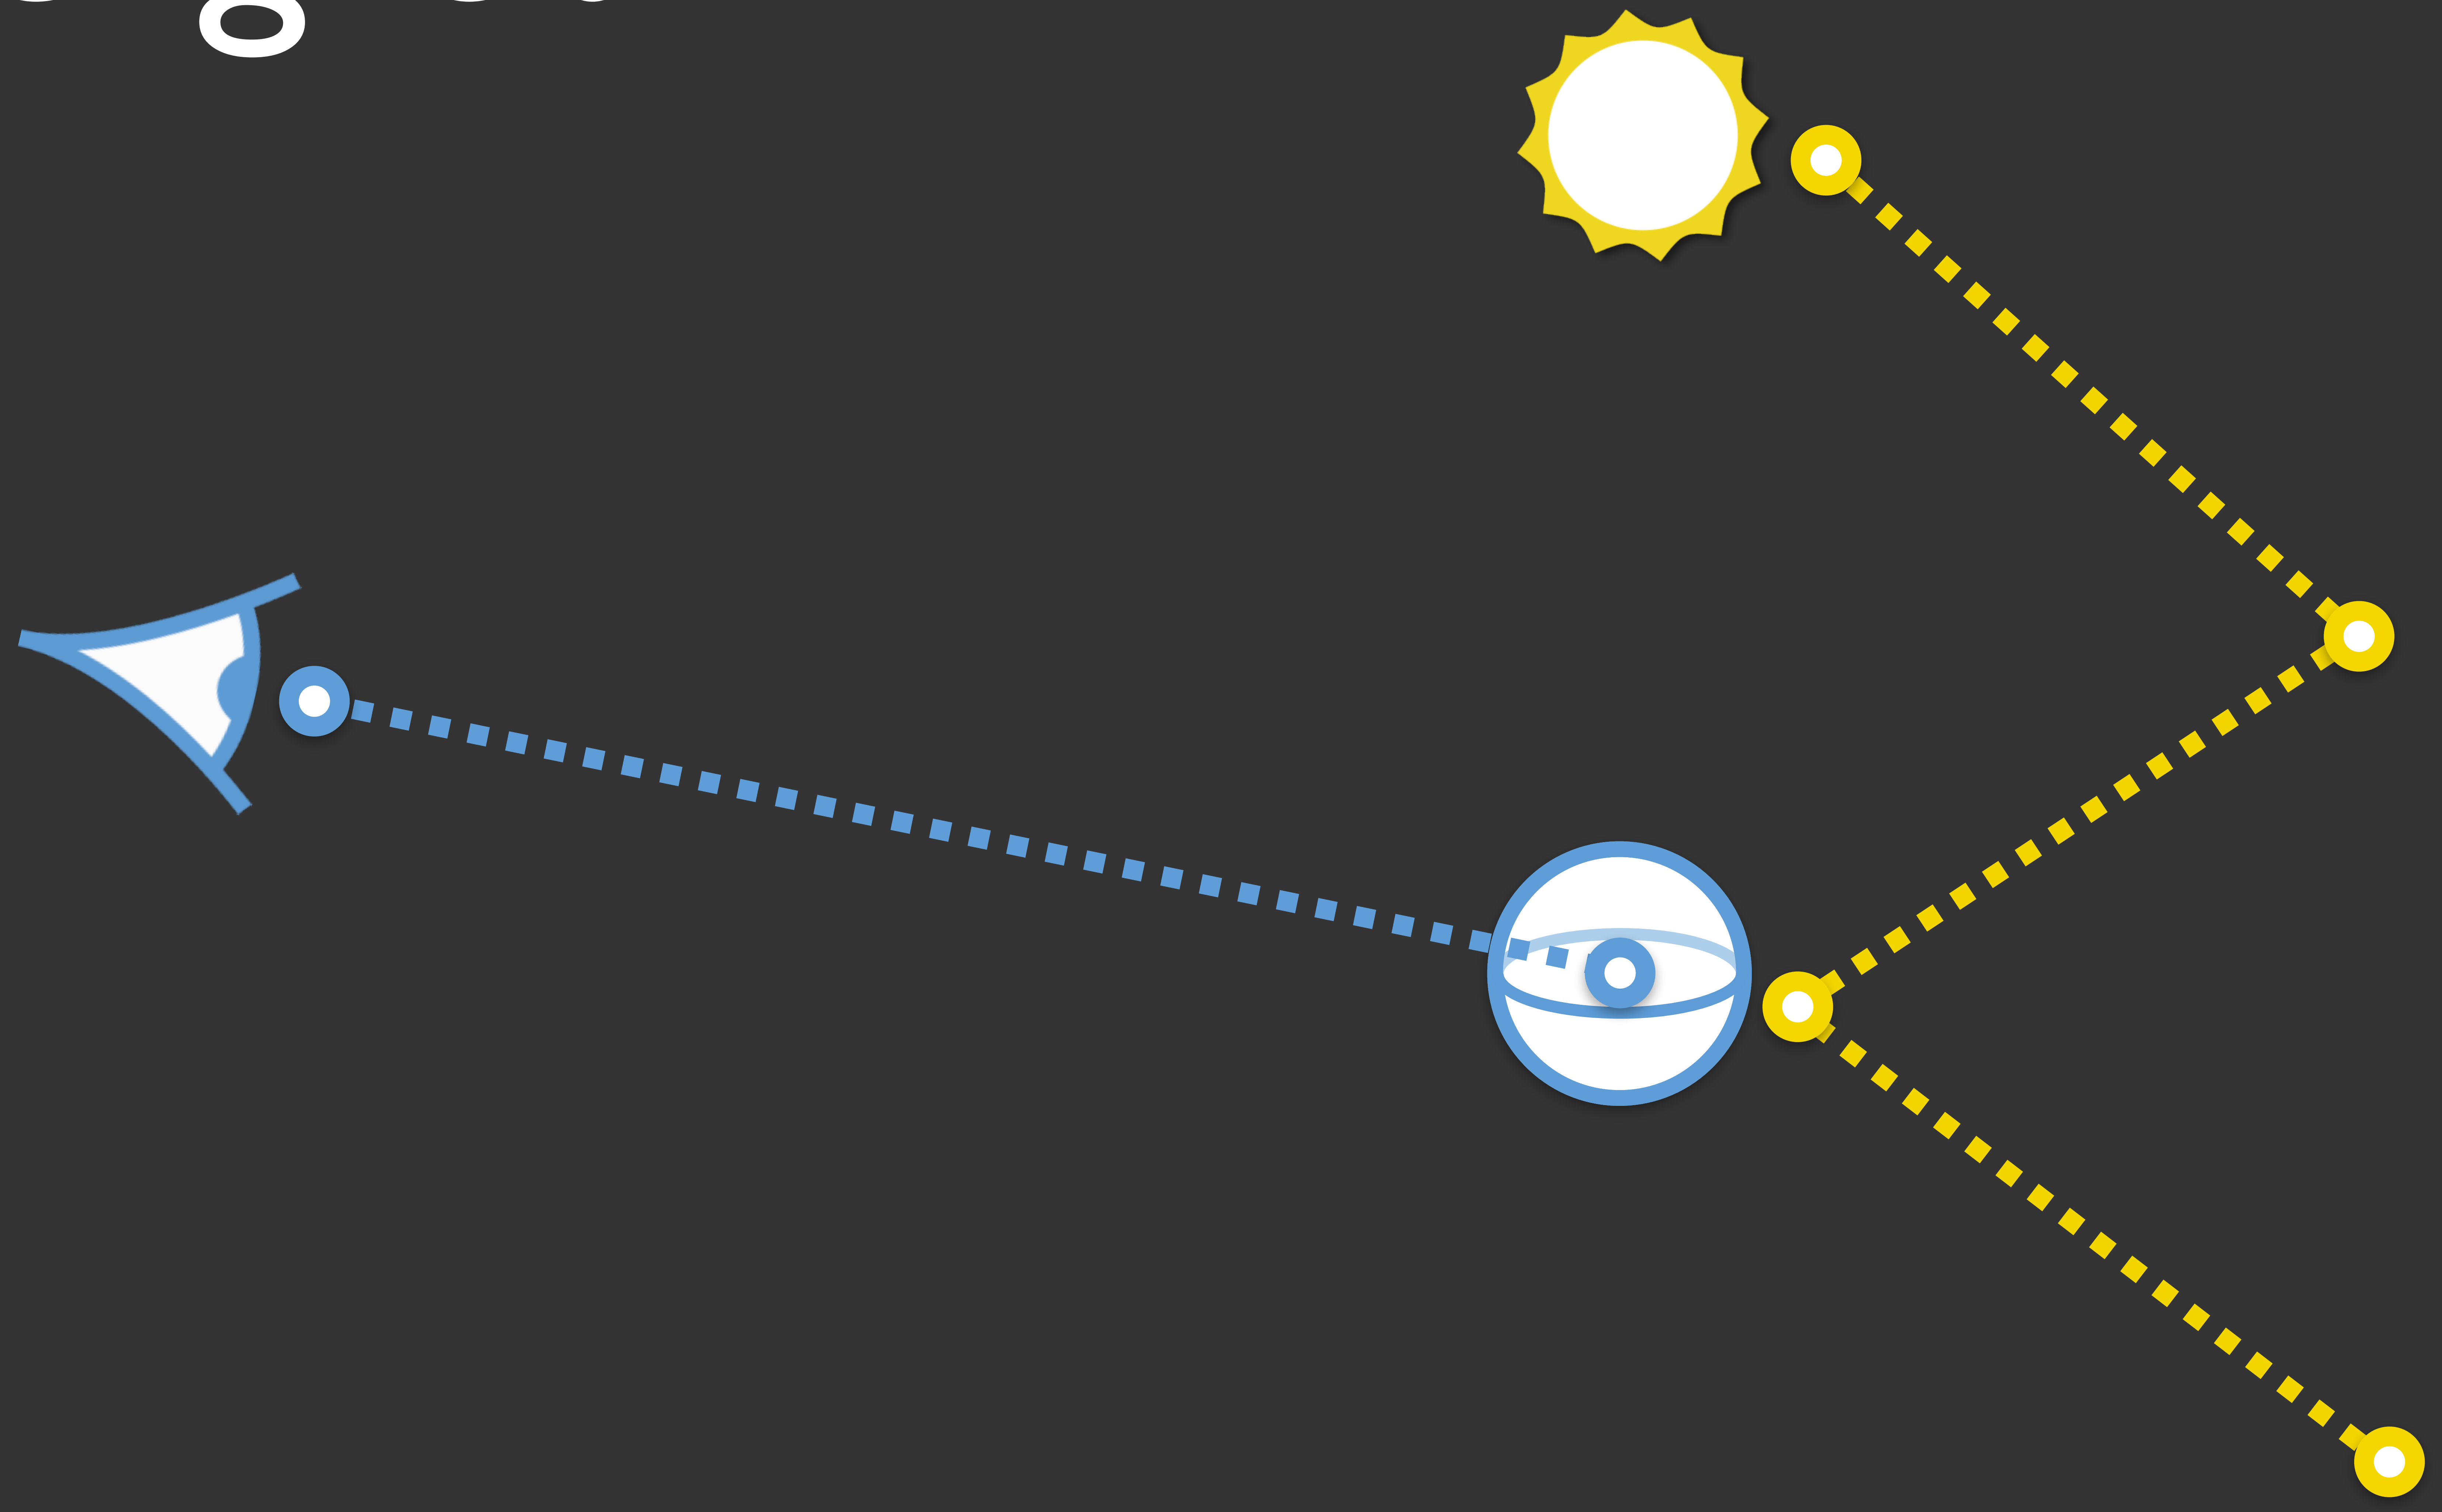

Background

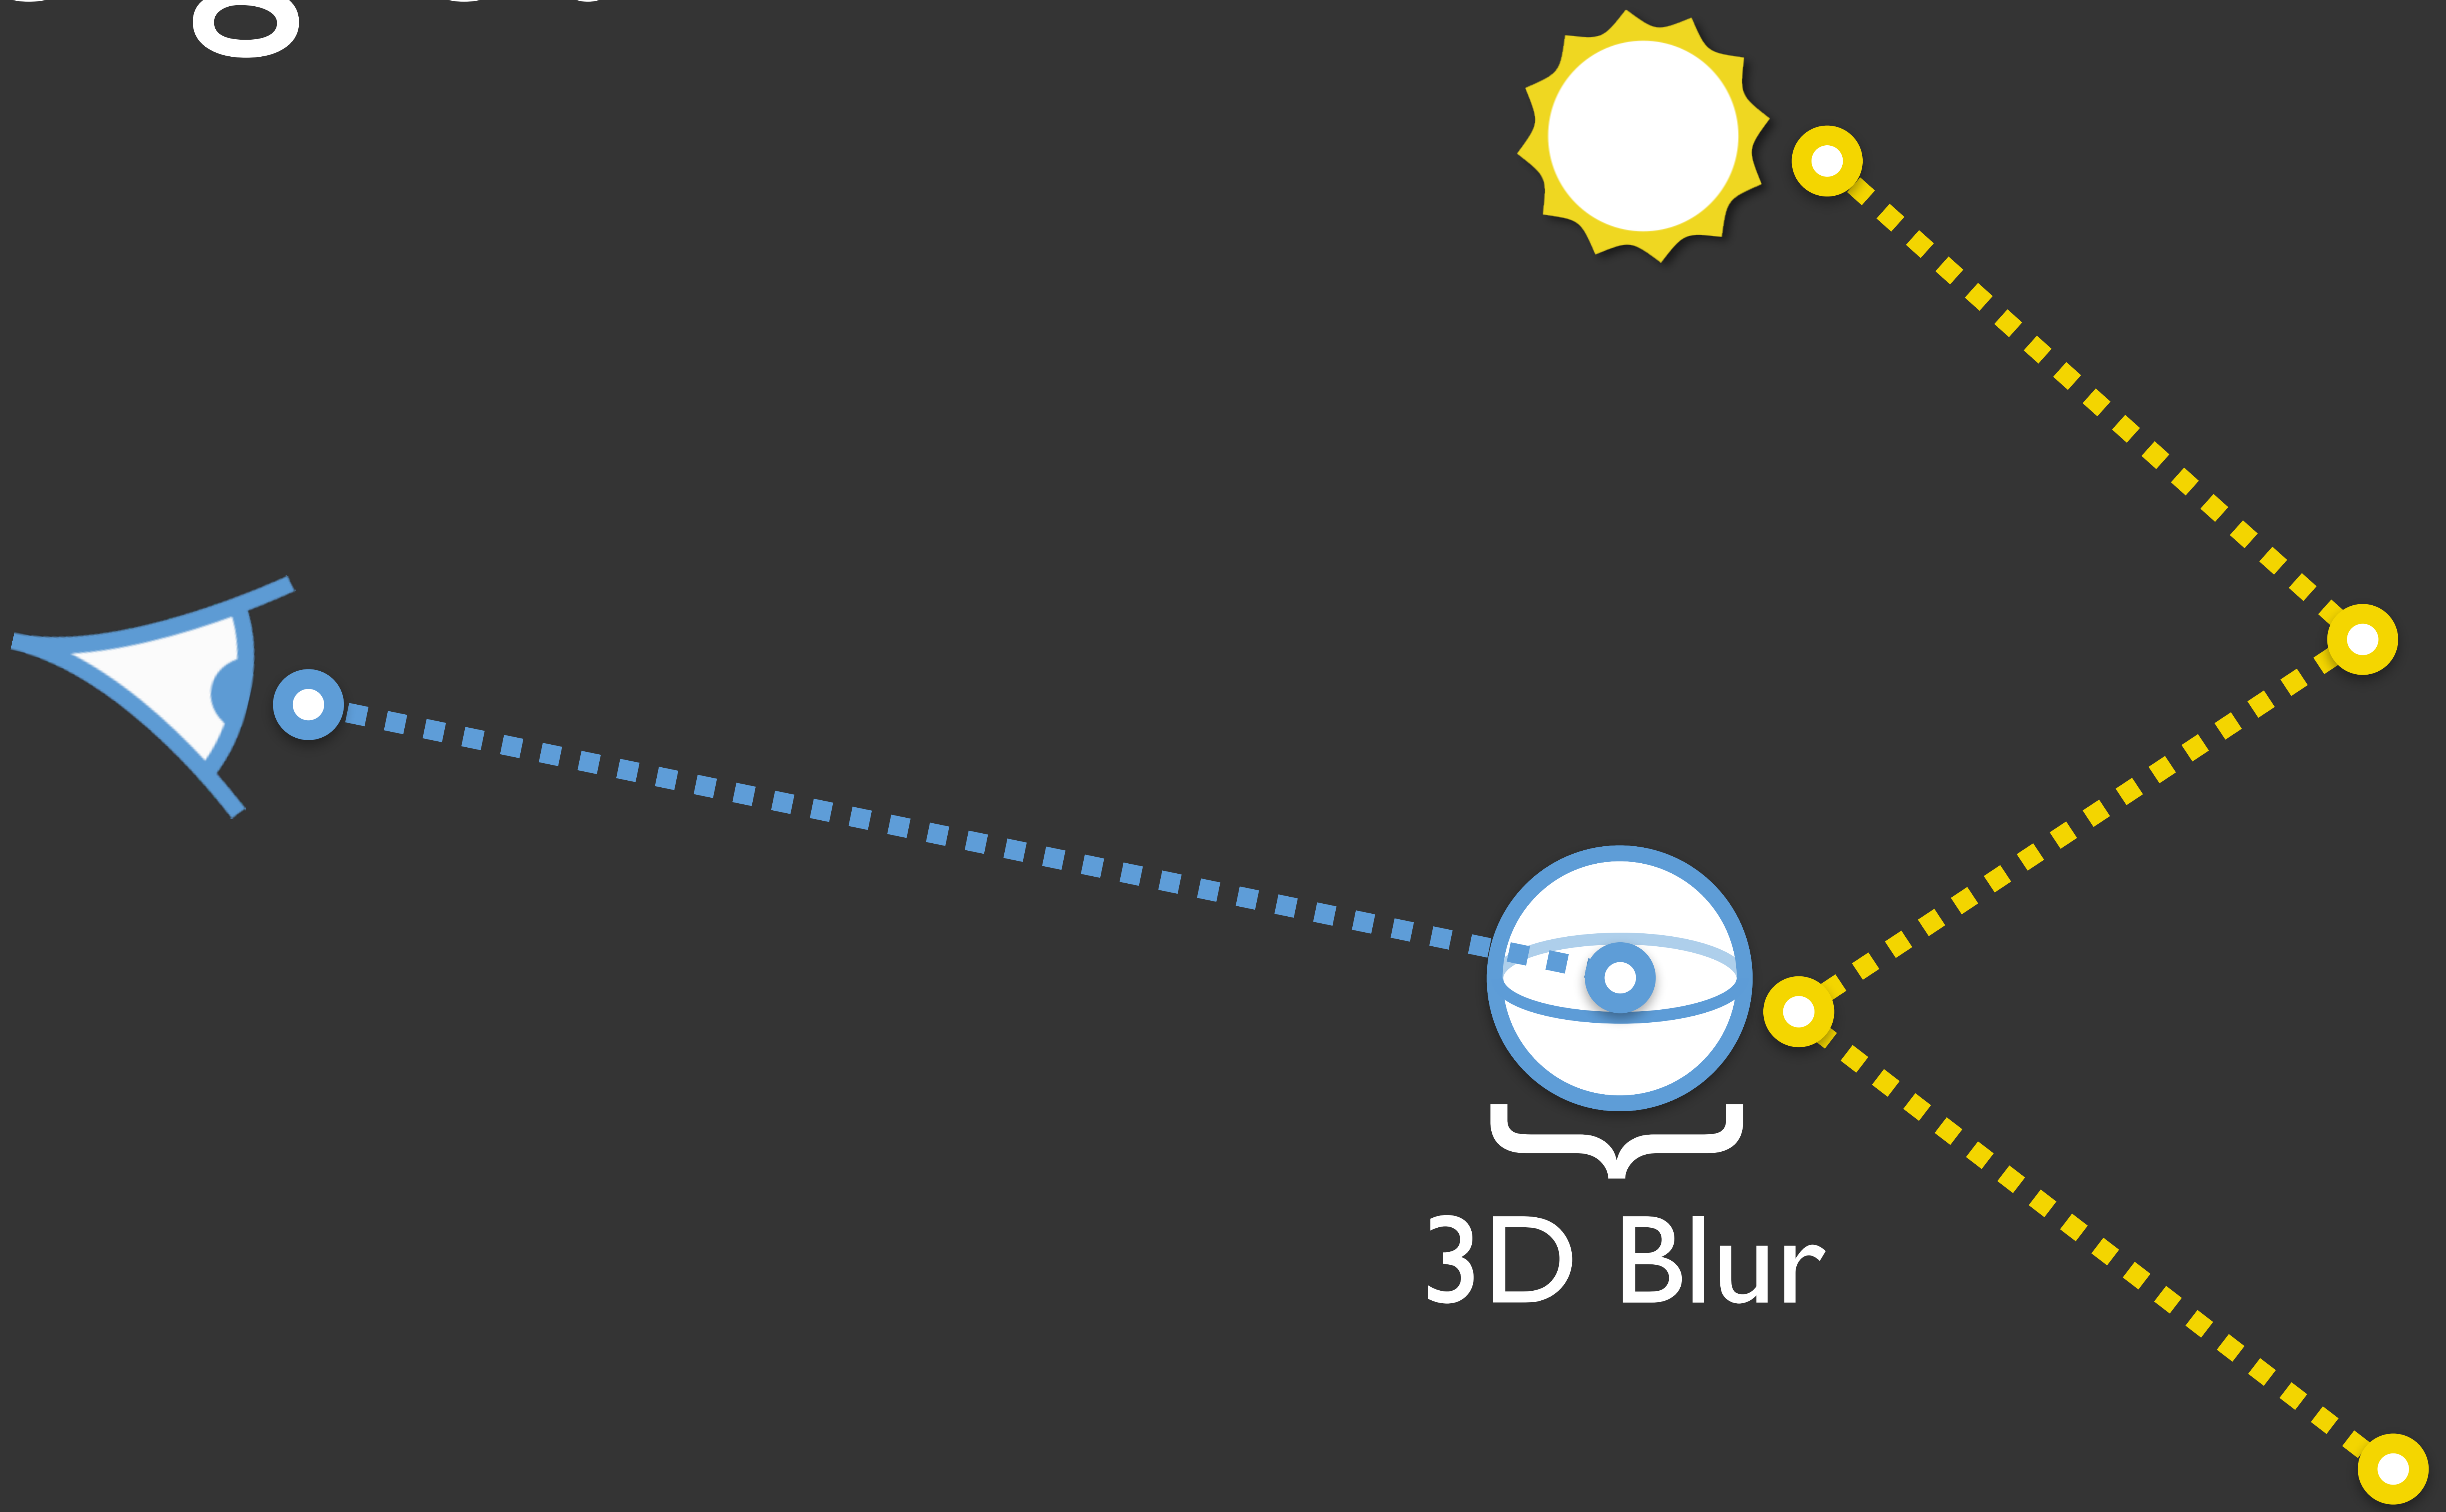

Background

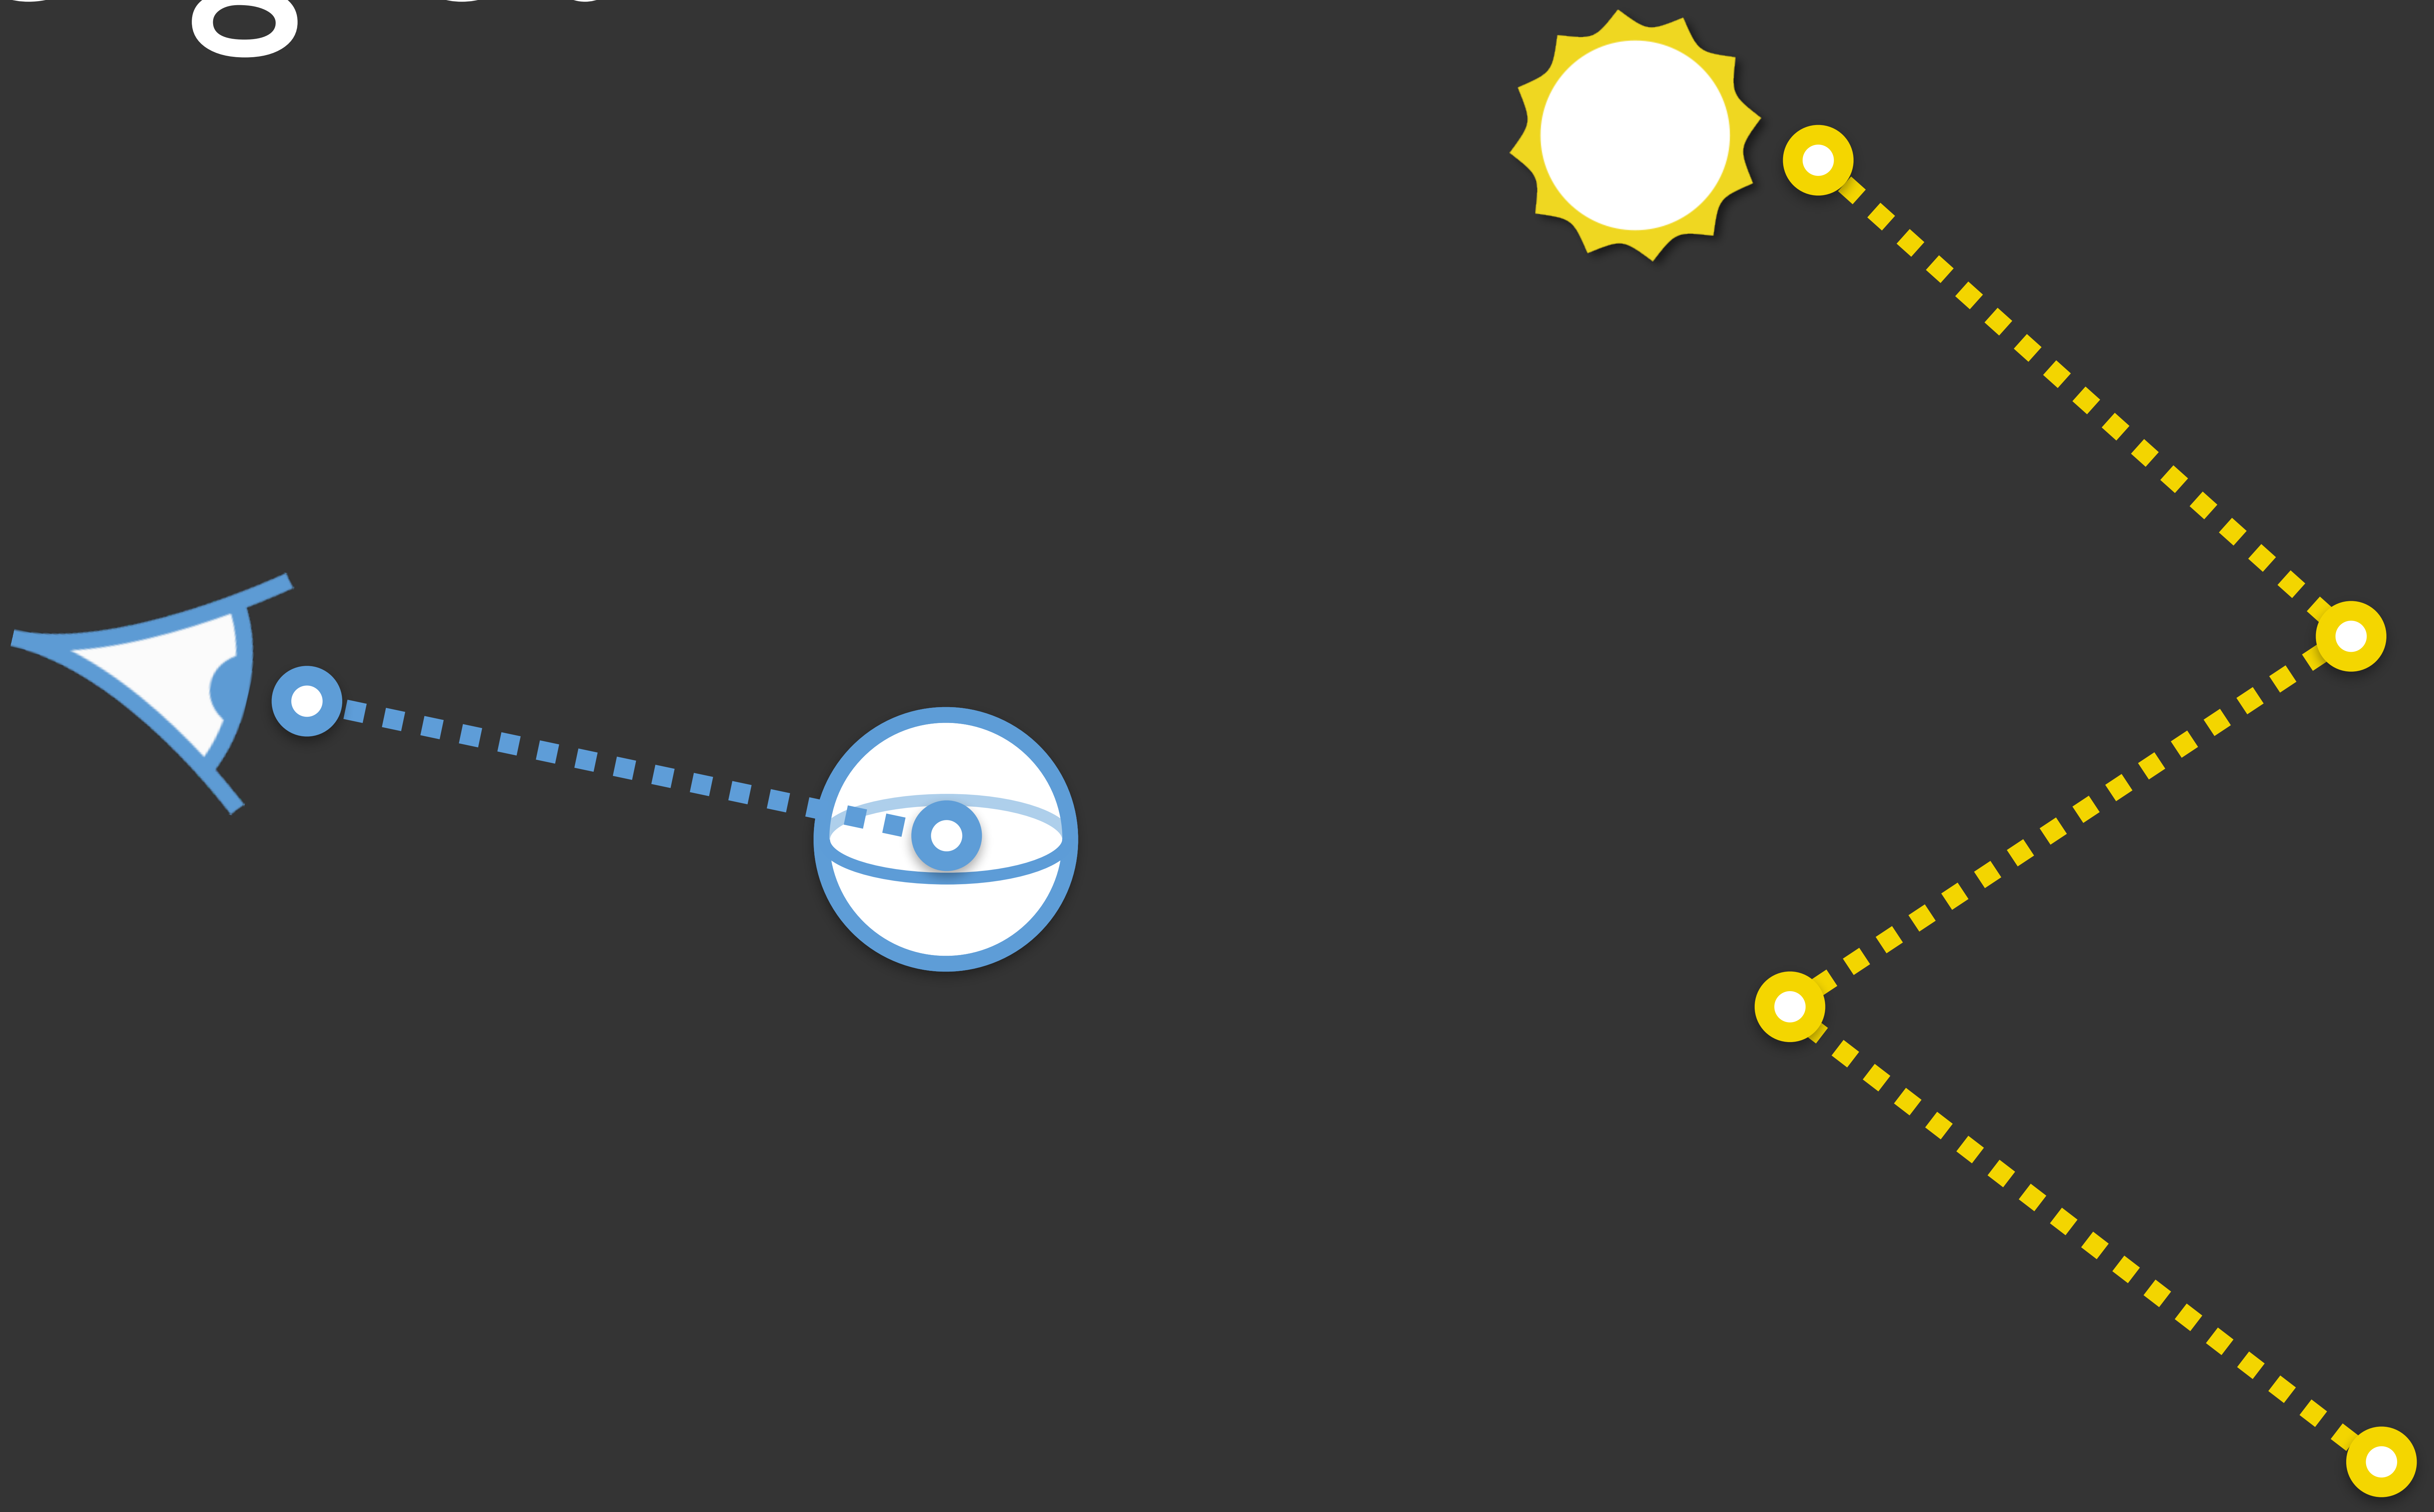

Background

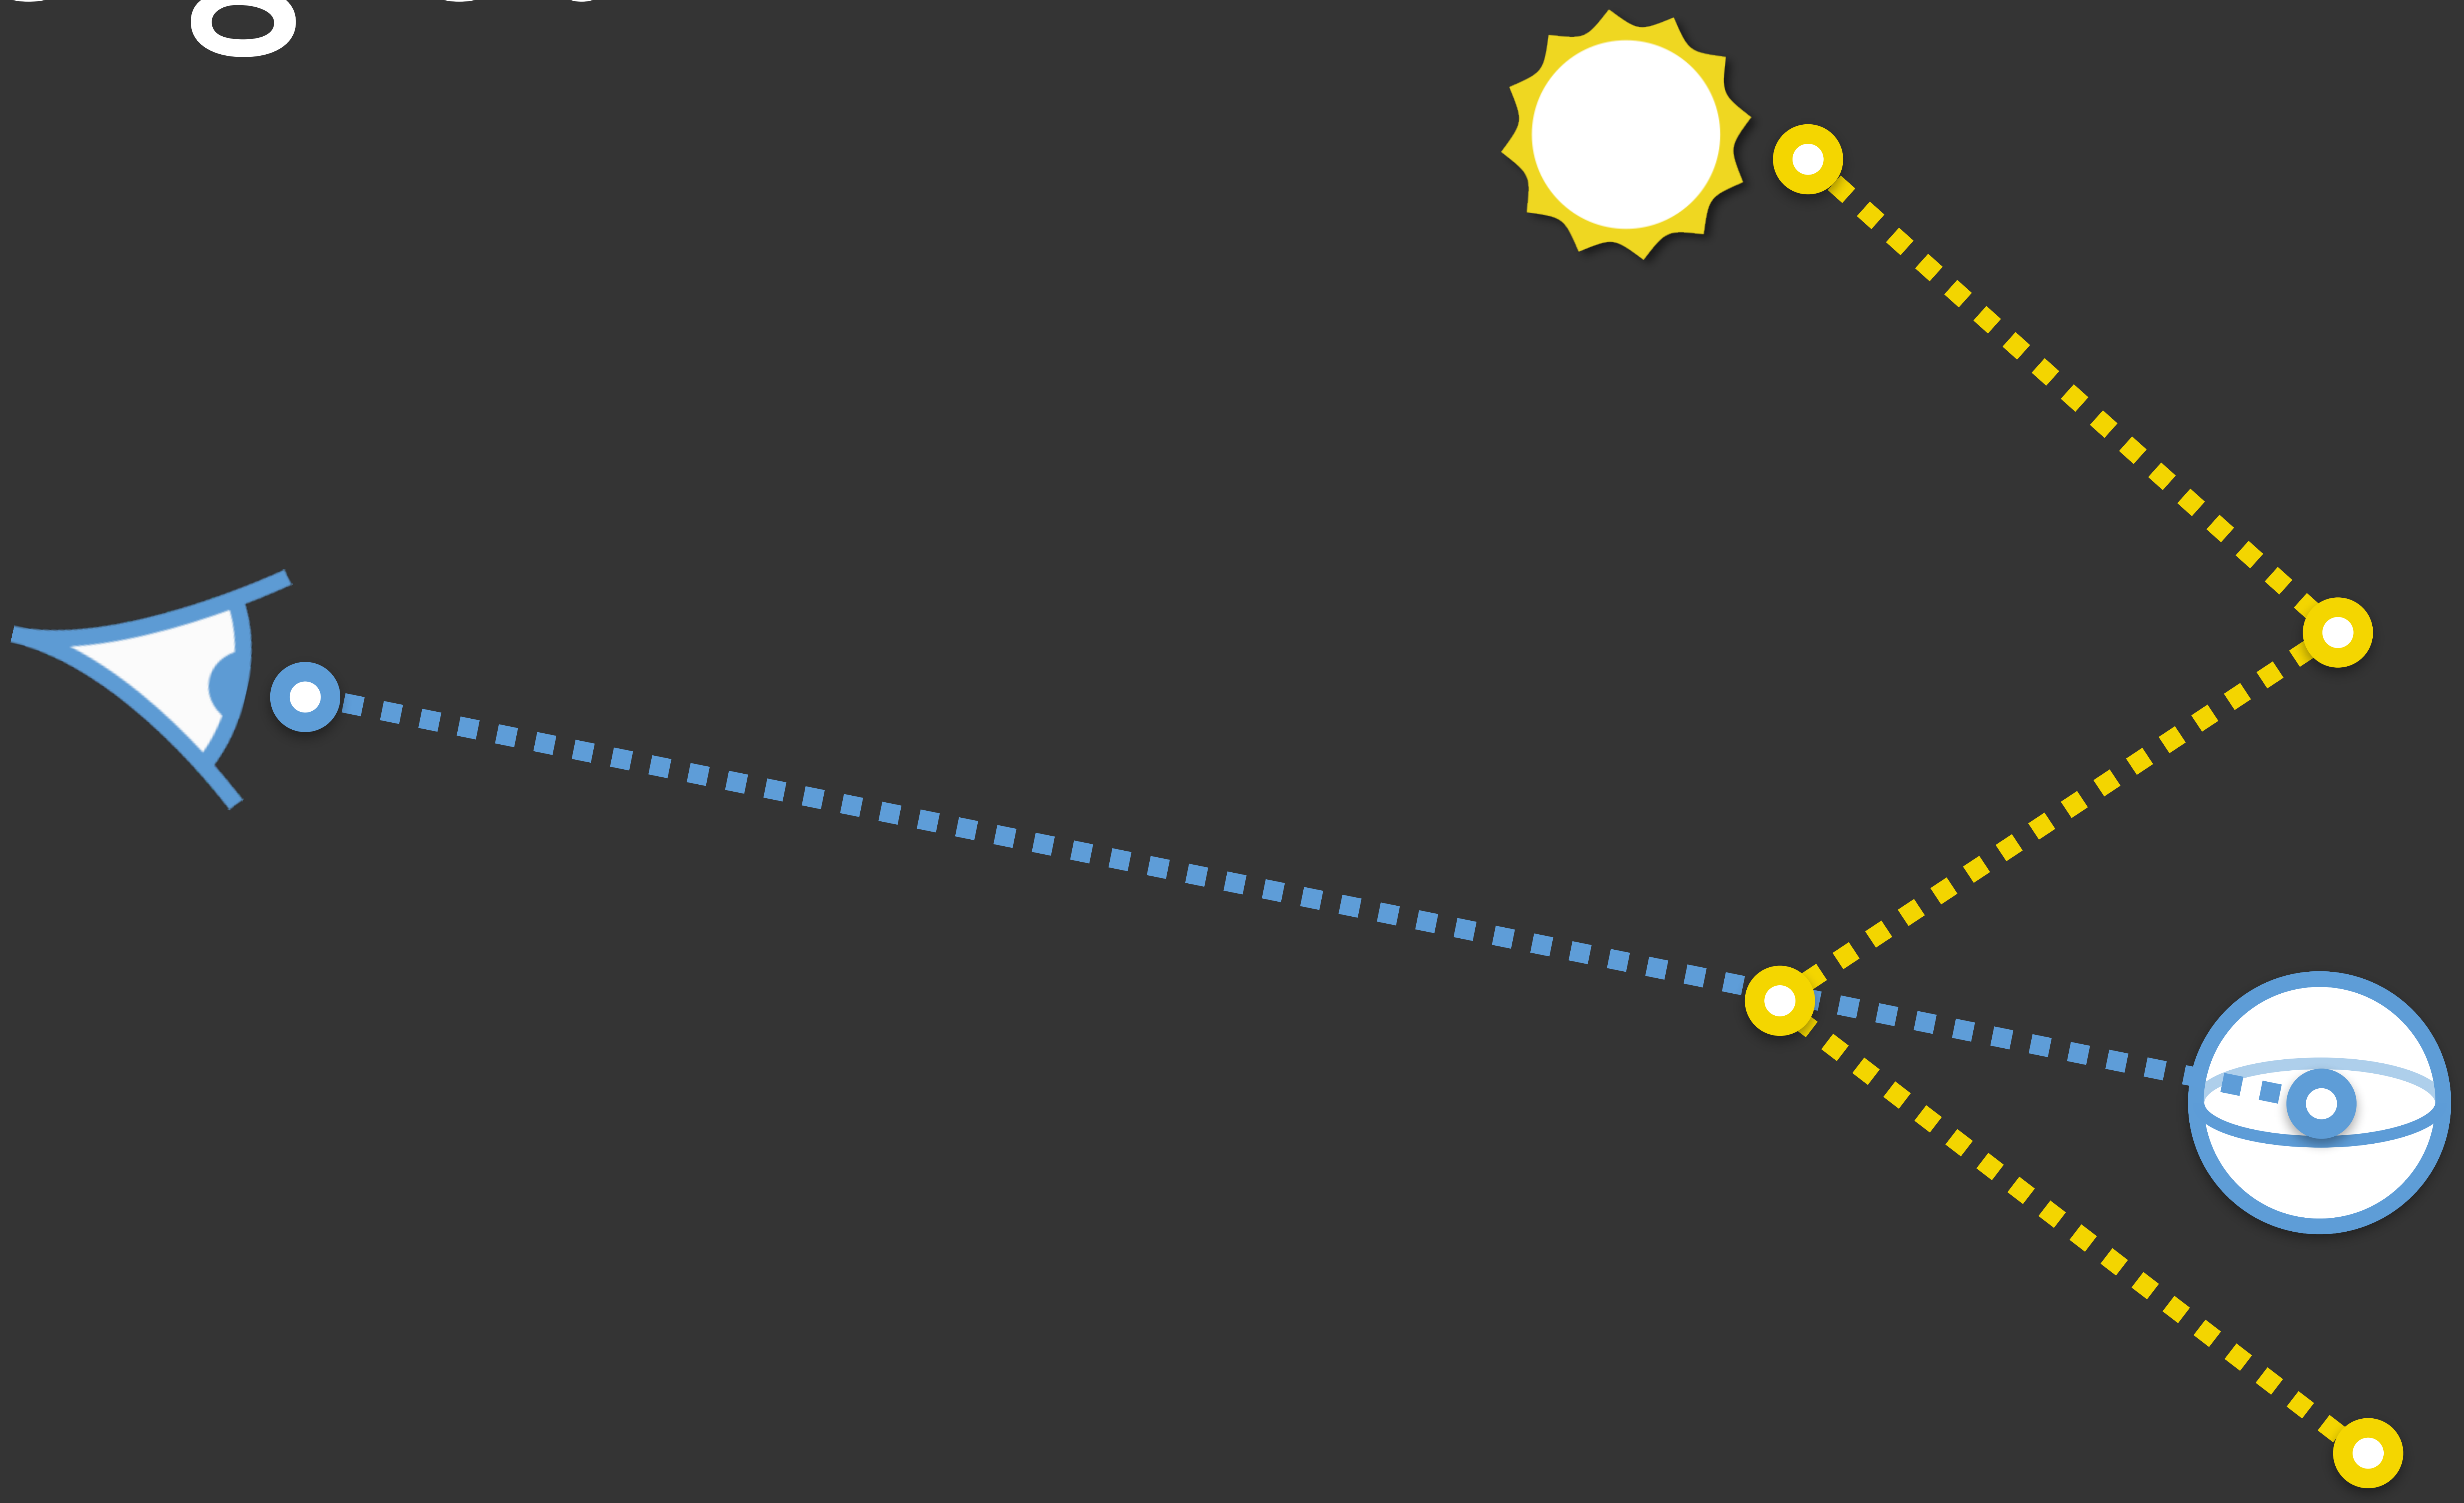

Ray Marching

[Jensen & Christensen ’98]

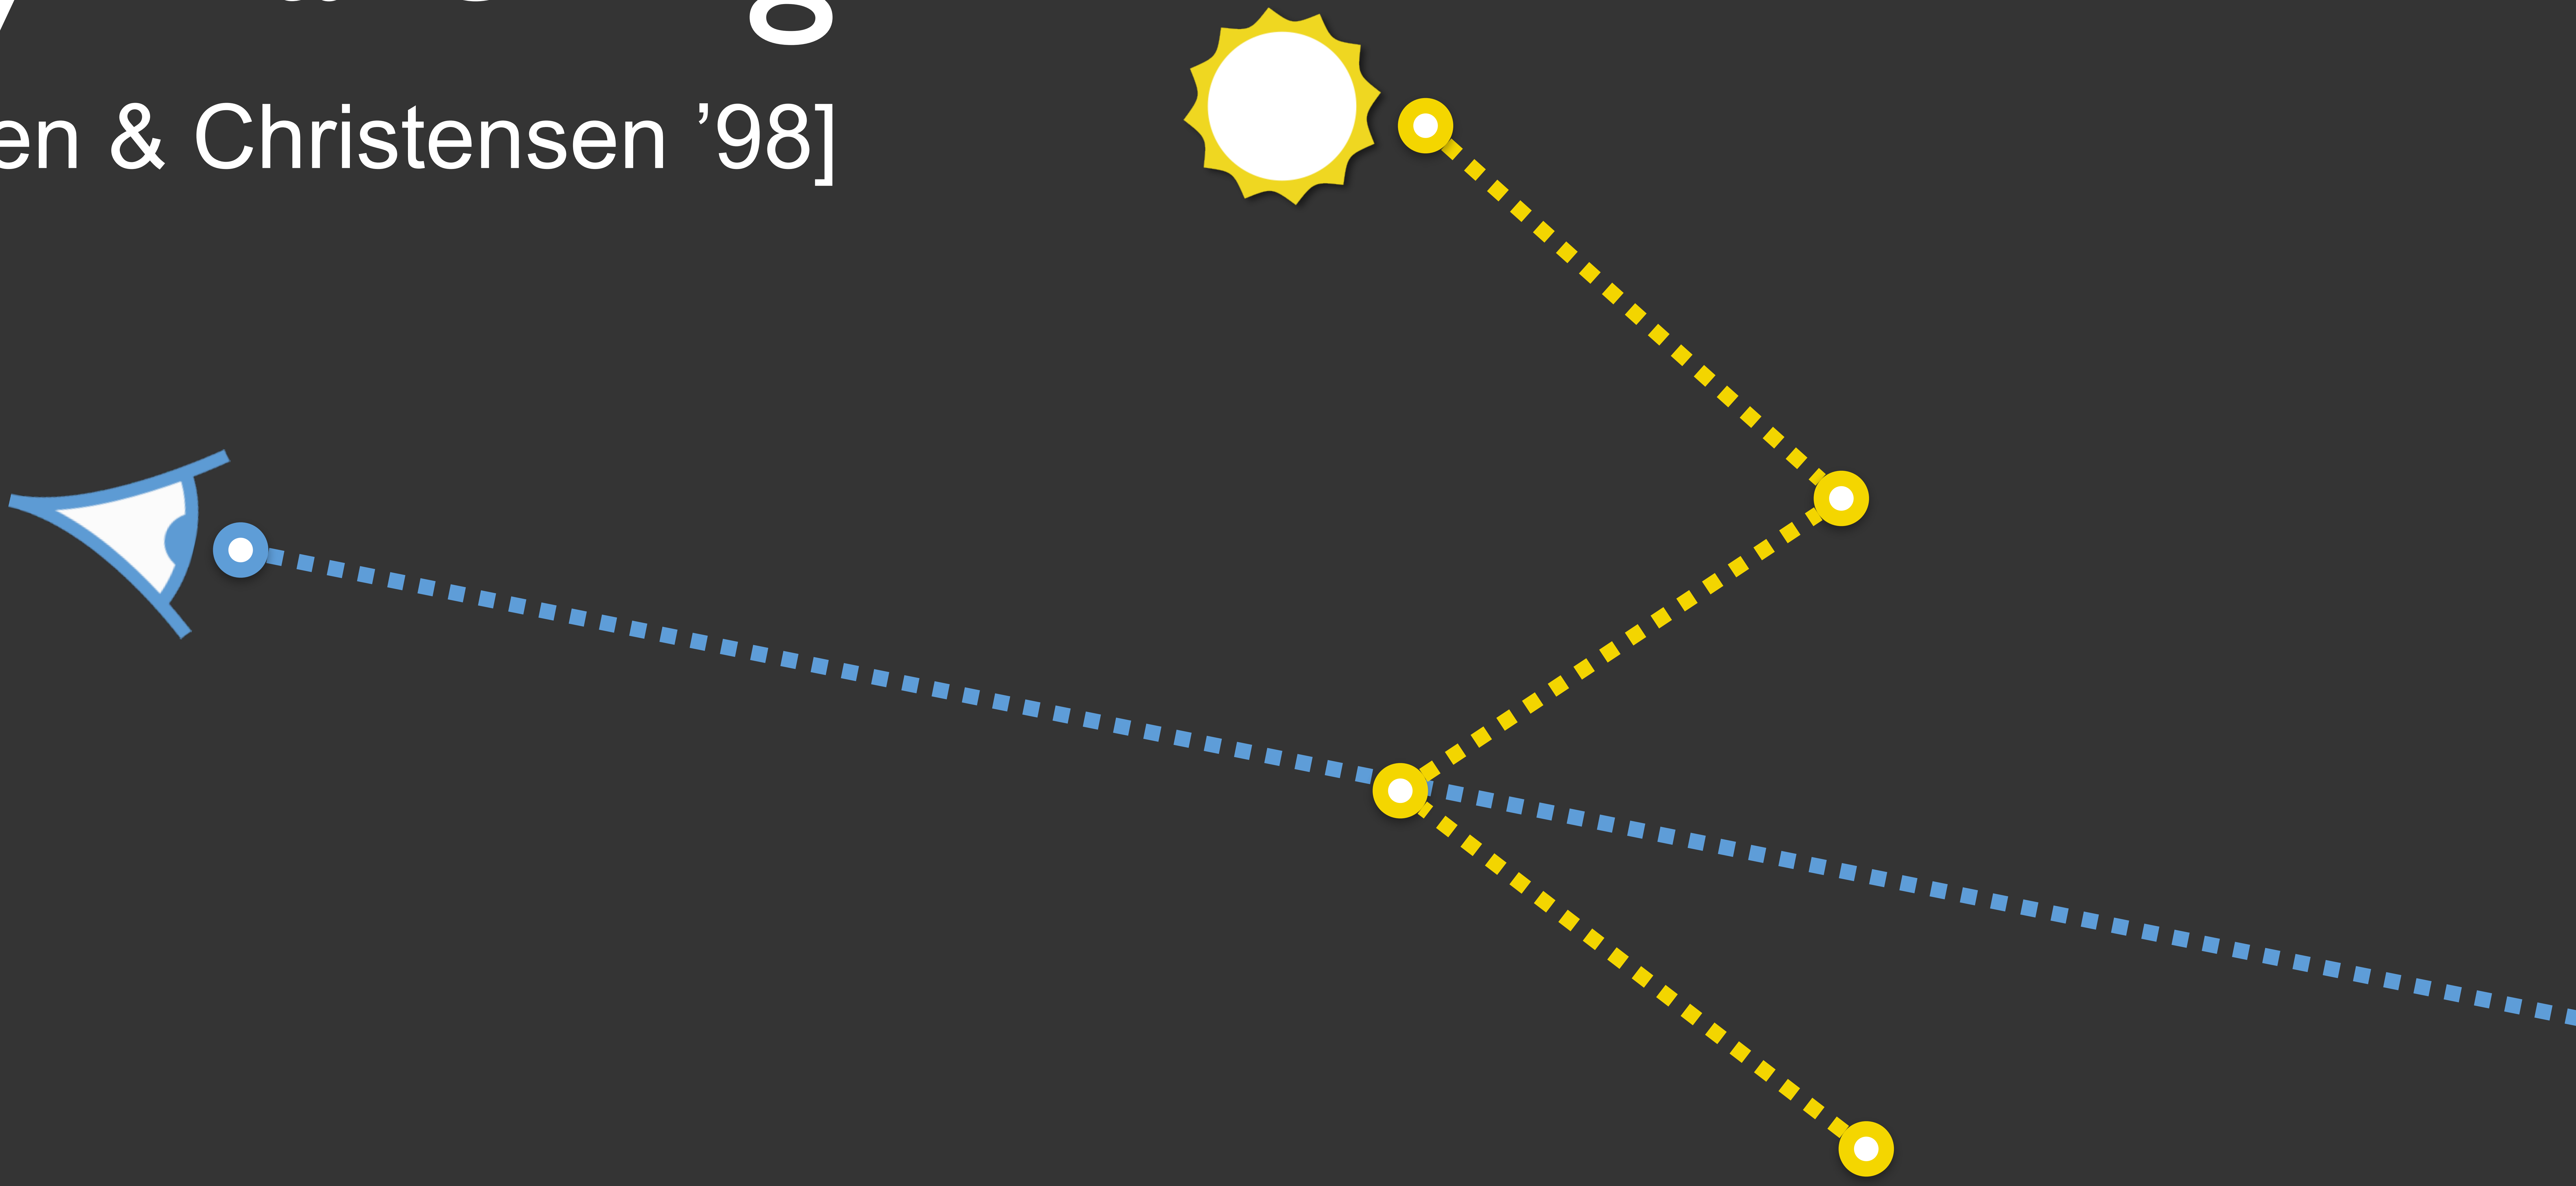

Ray Marching

[Jensen & Christensen ’98]

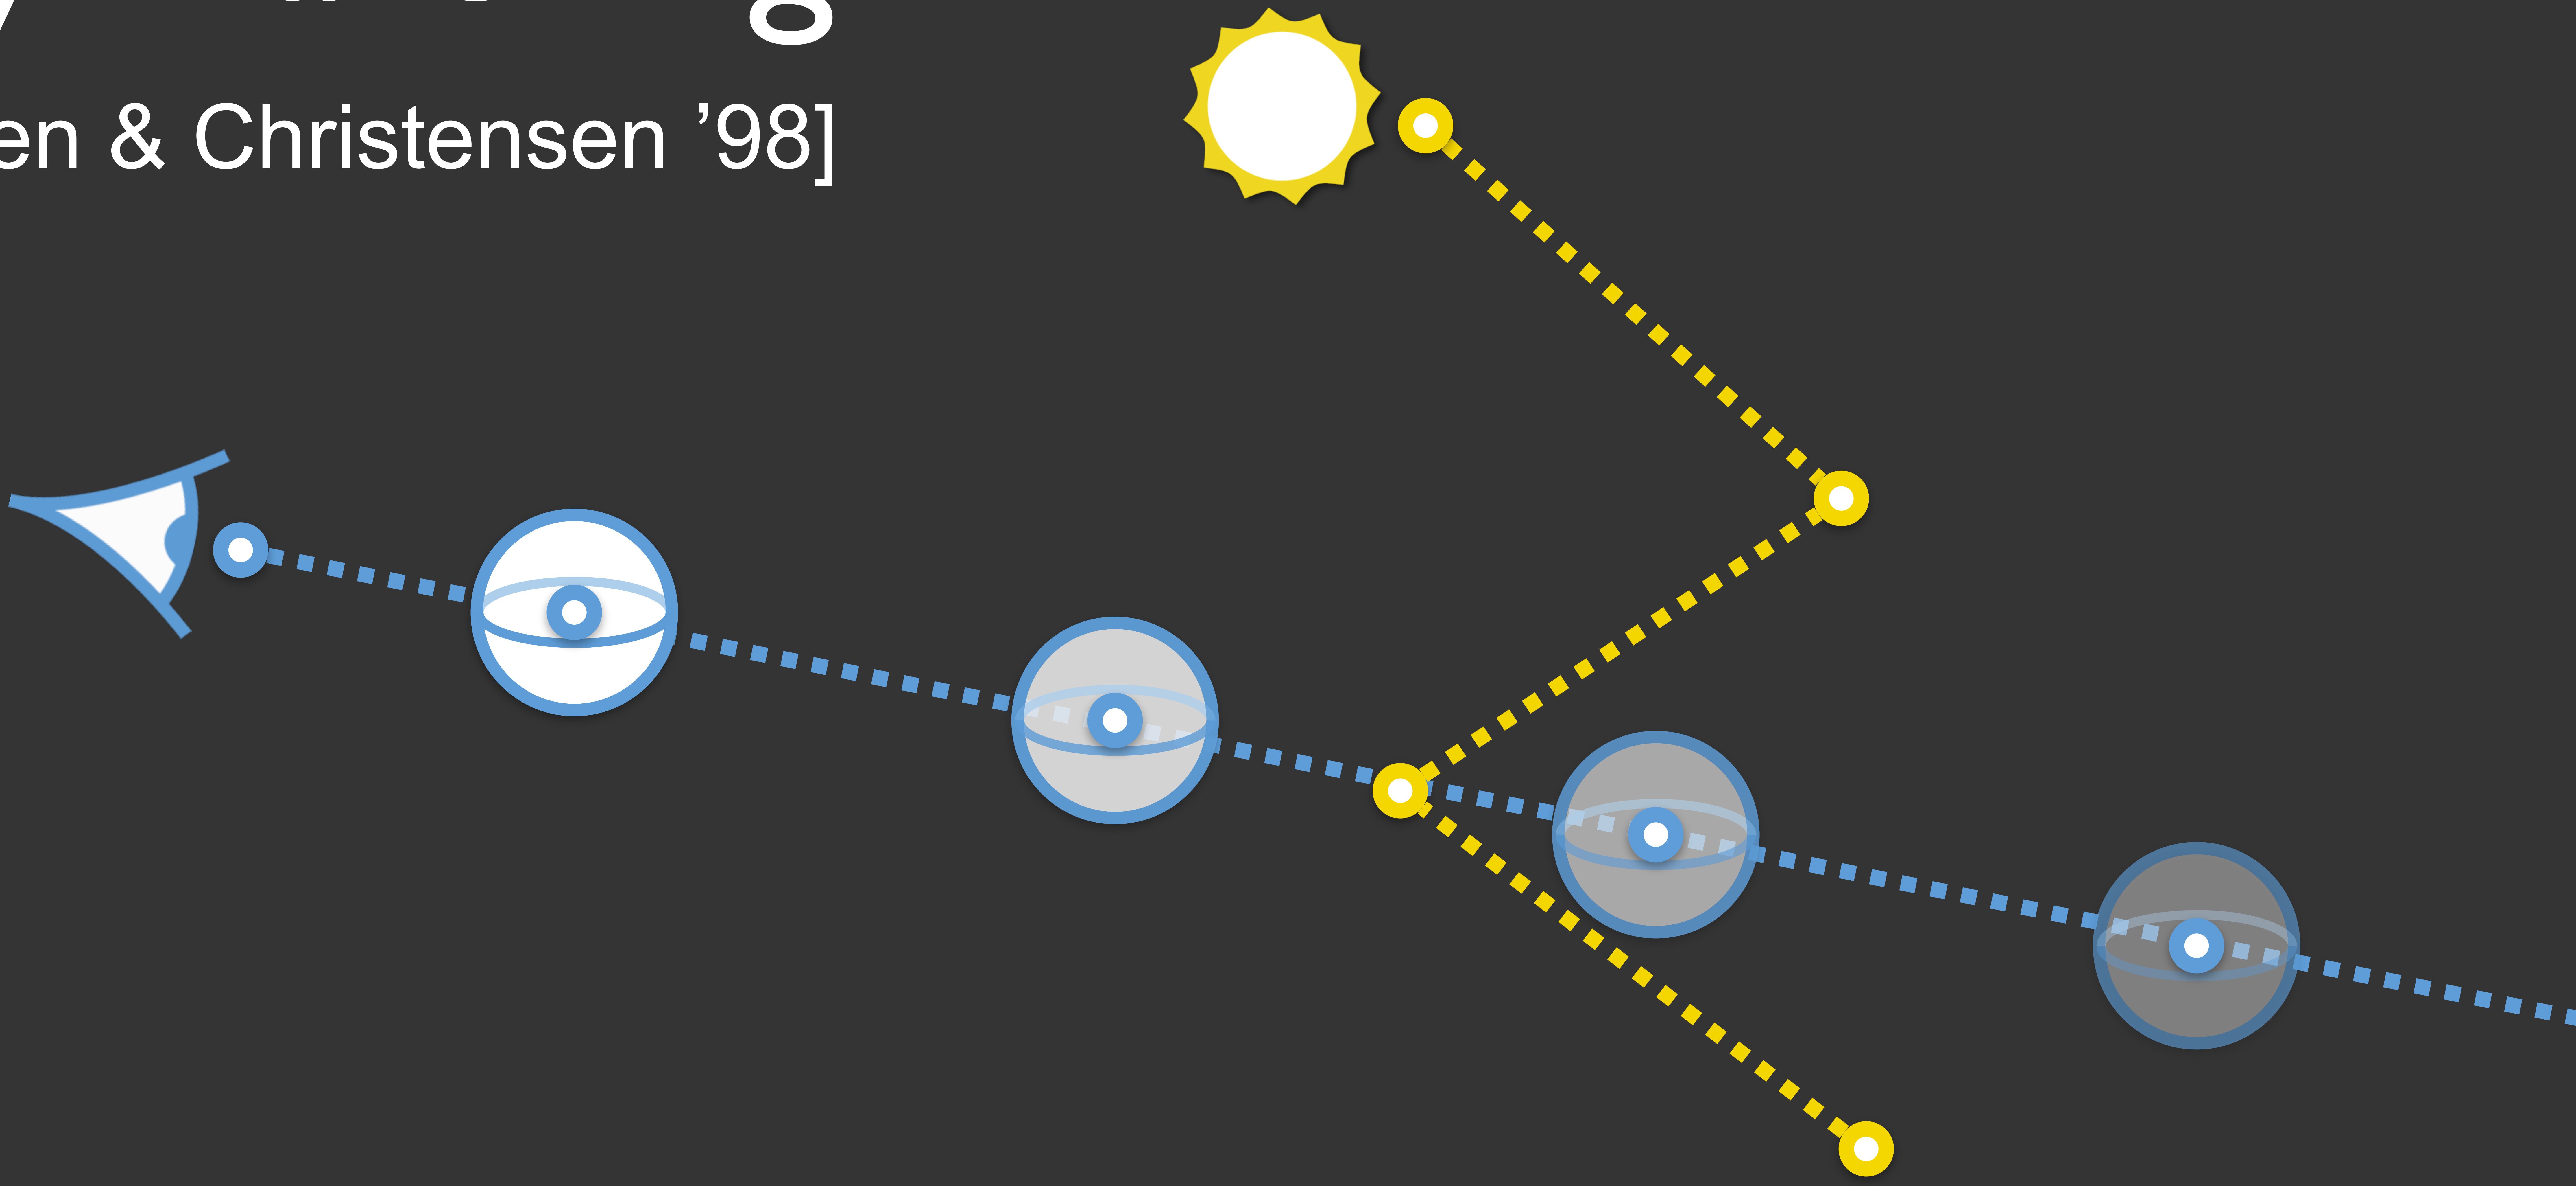

Ray Marching

[Jensen & Christensen ’98]

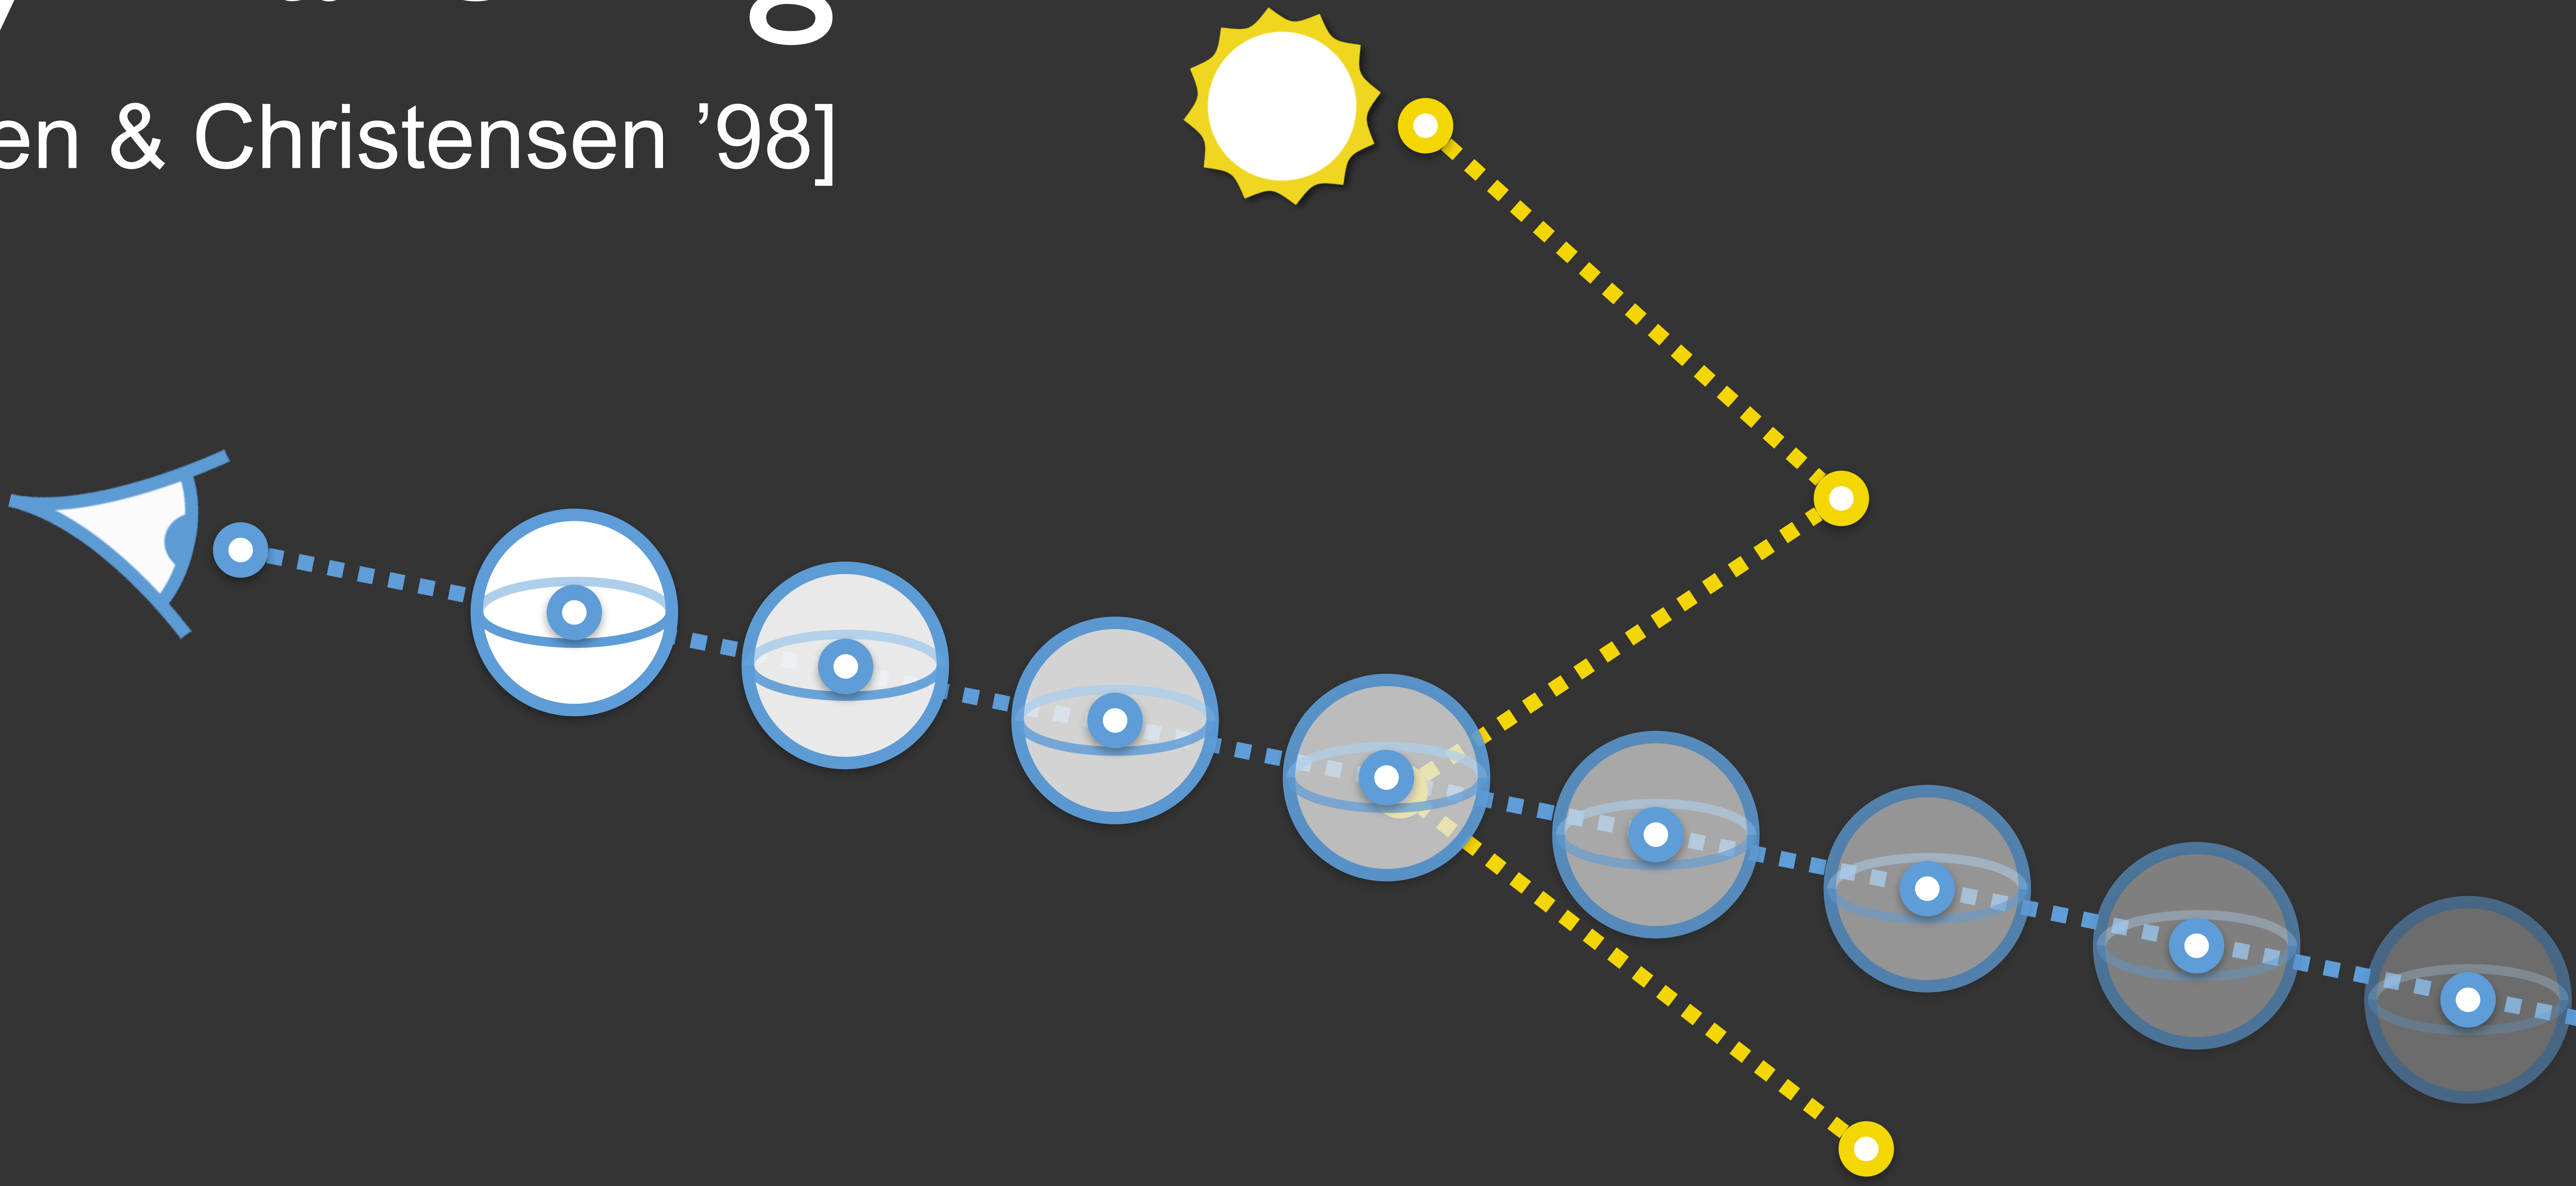

Beam Radiance Estimate

[Jarosz et al. ’08]

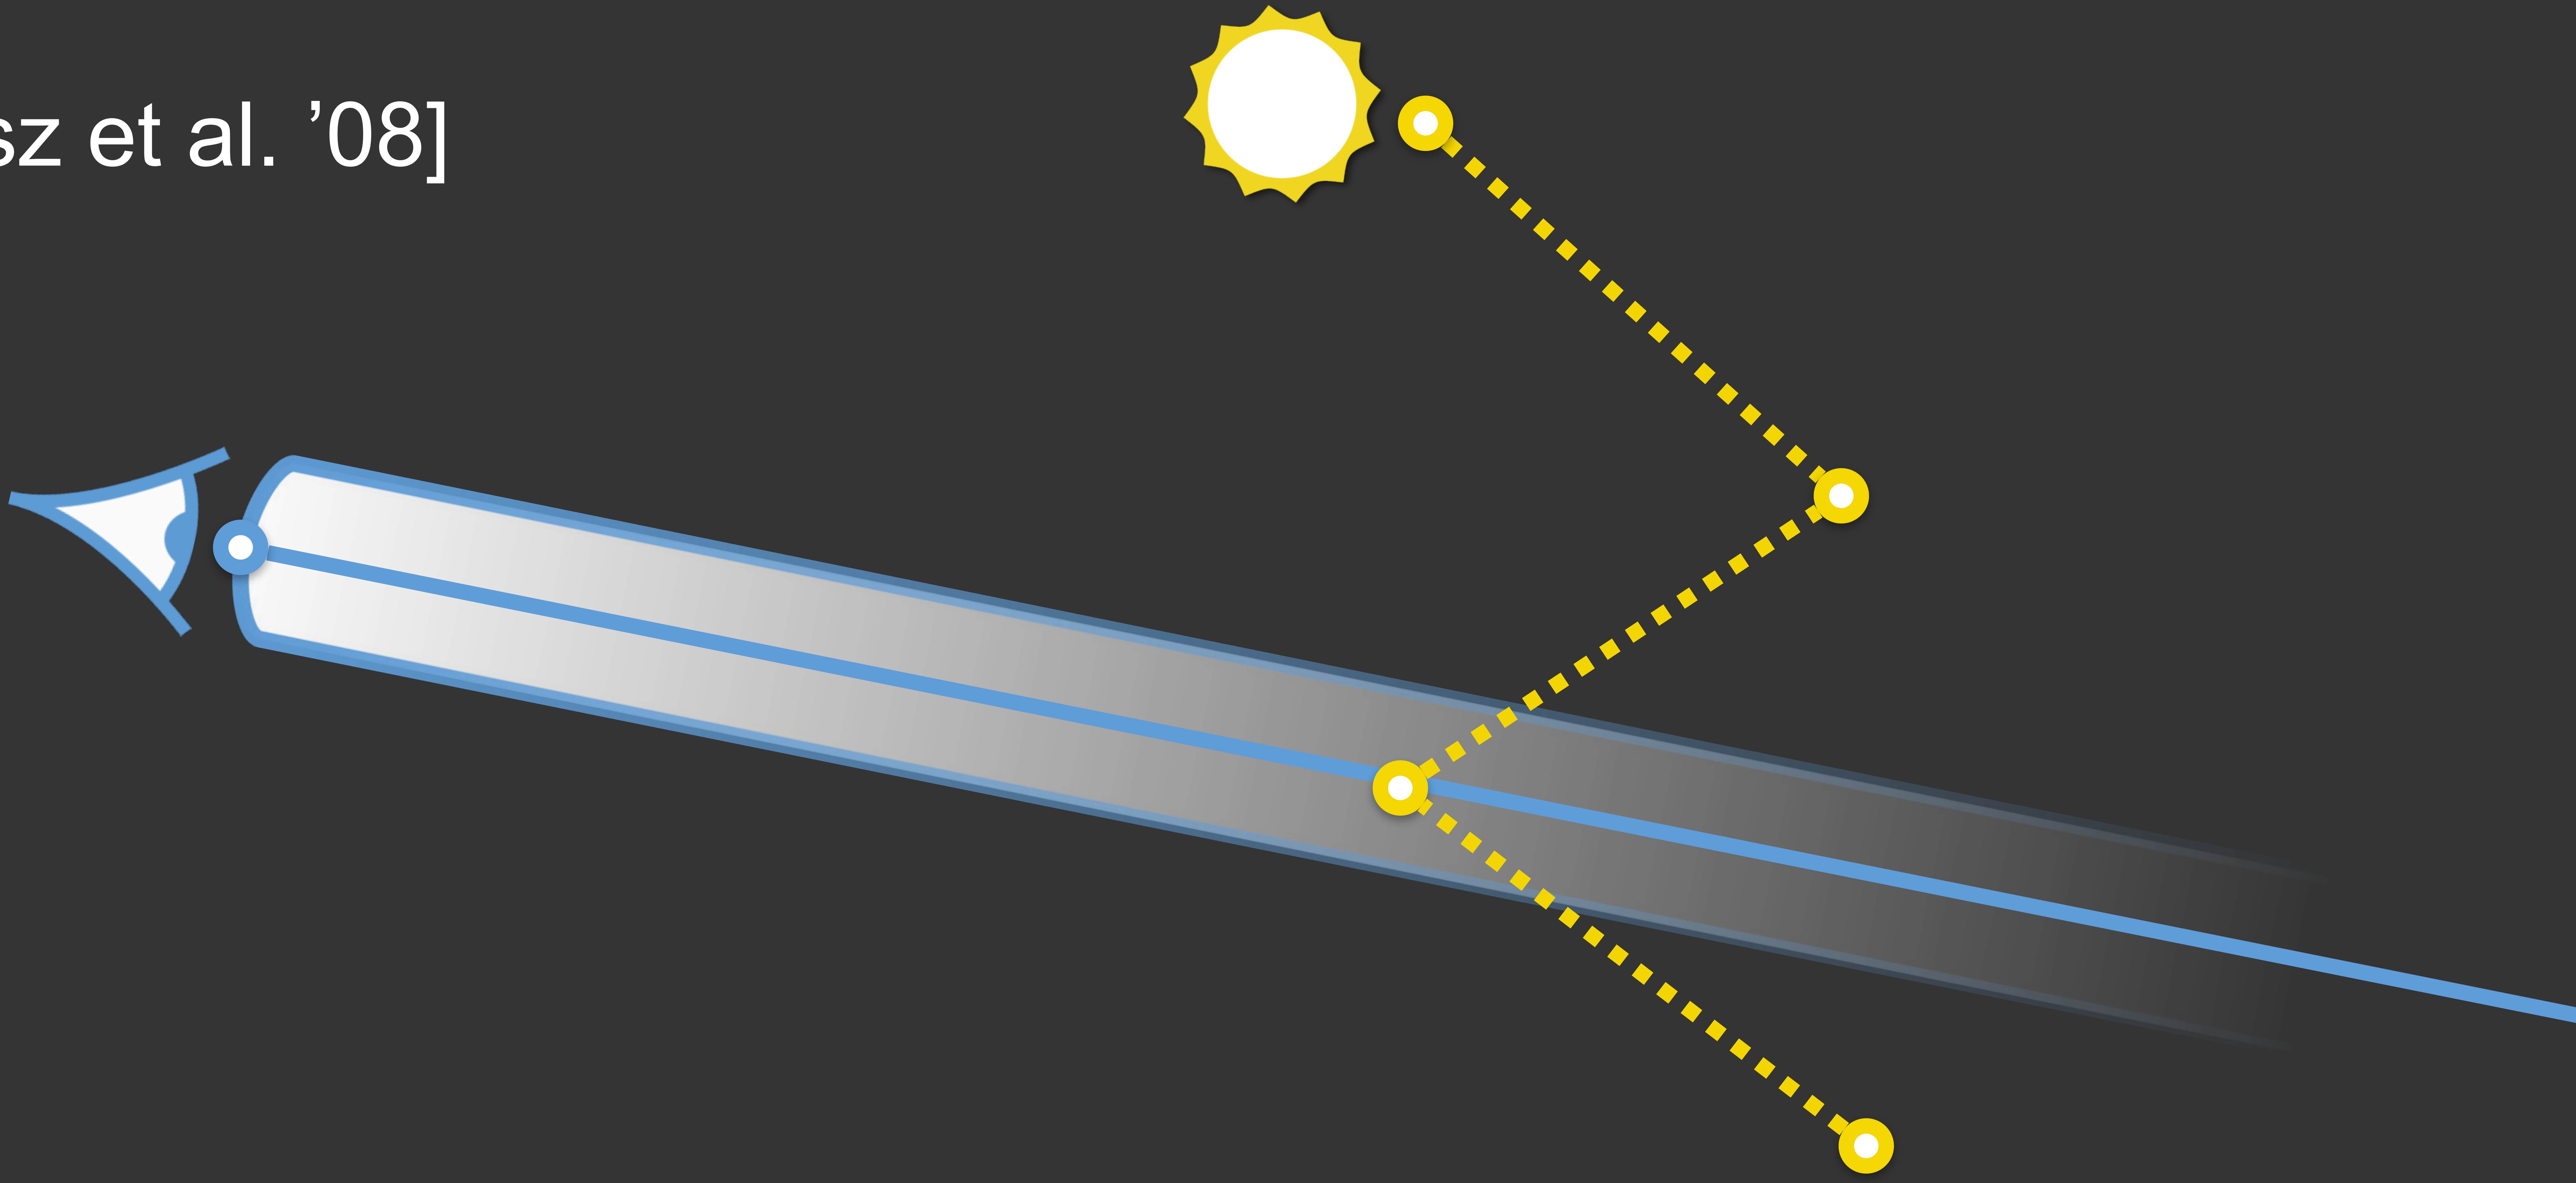

Beam Radiance Estimate

[Jarosz et al. ’08]

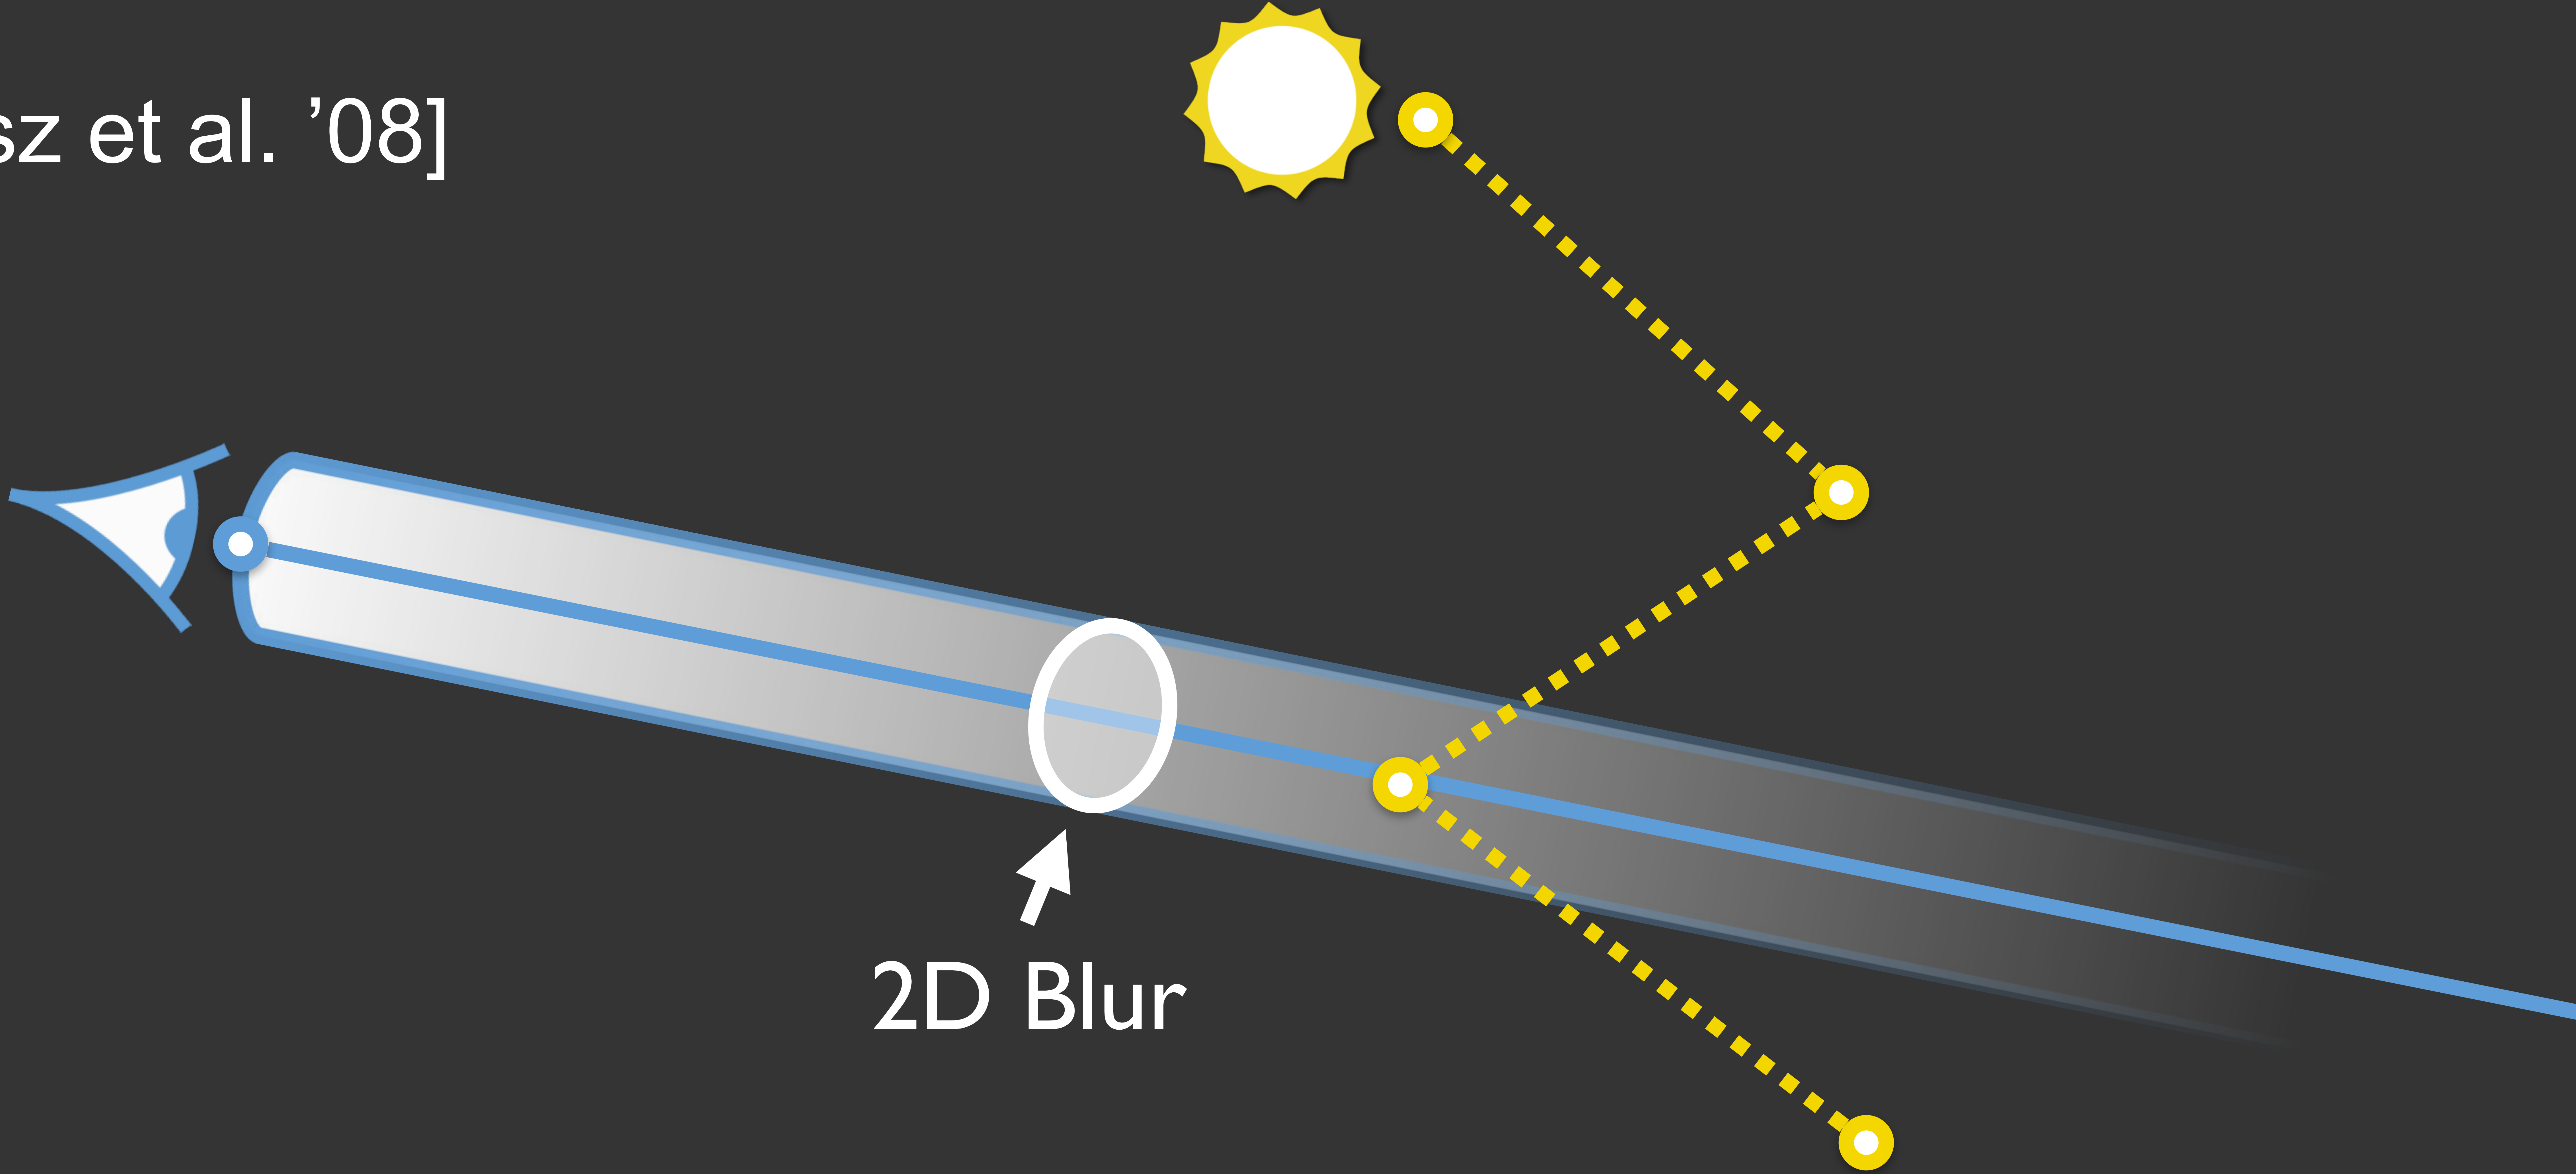

Background

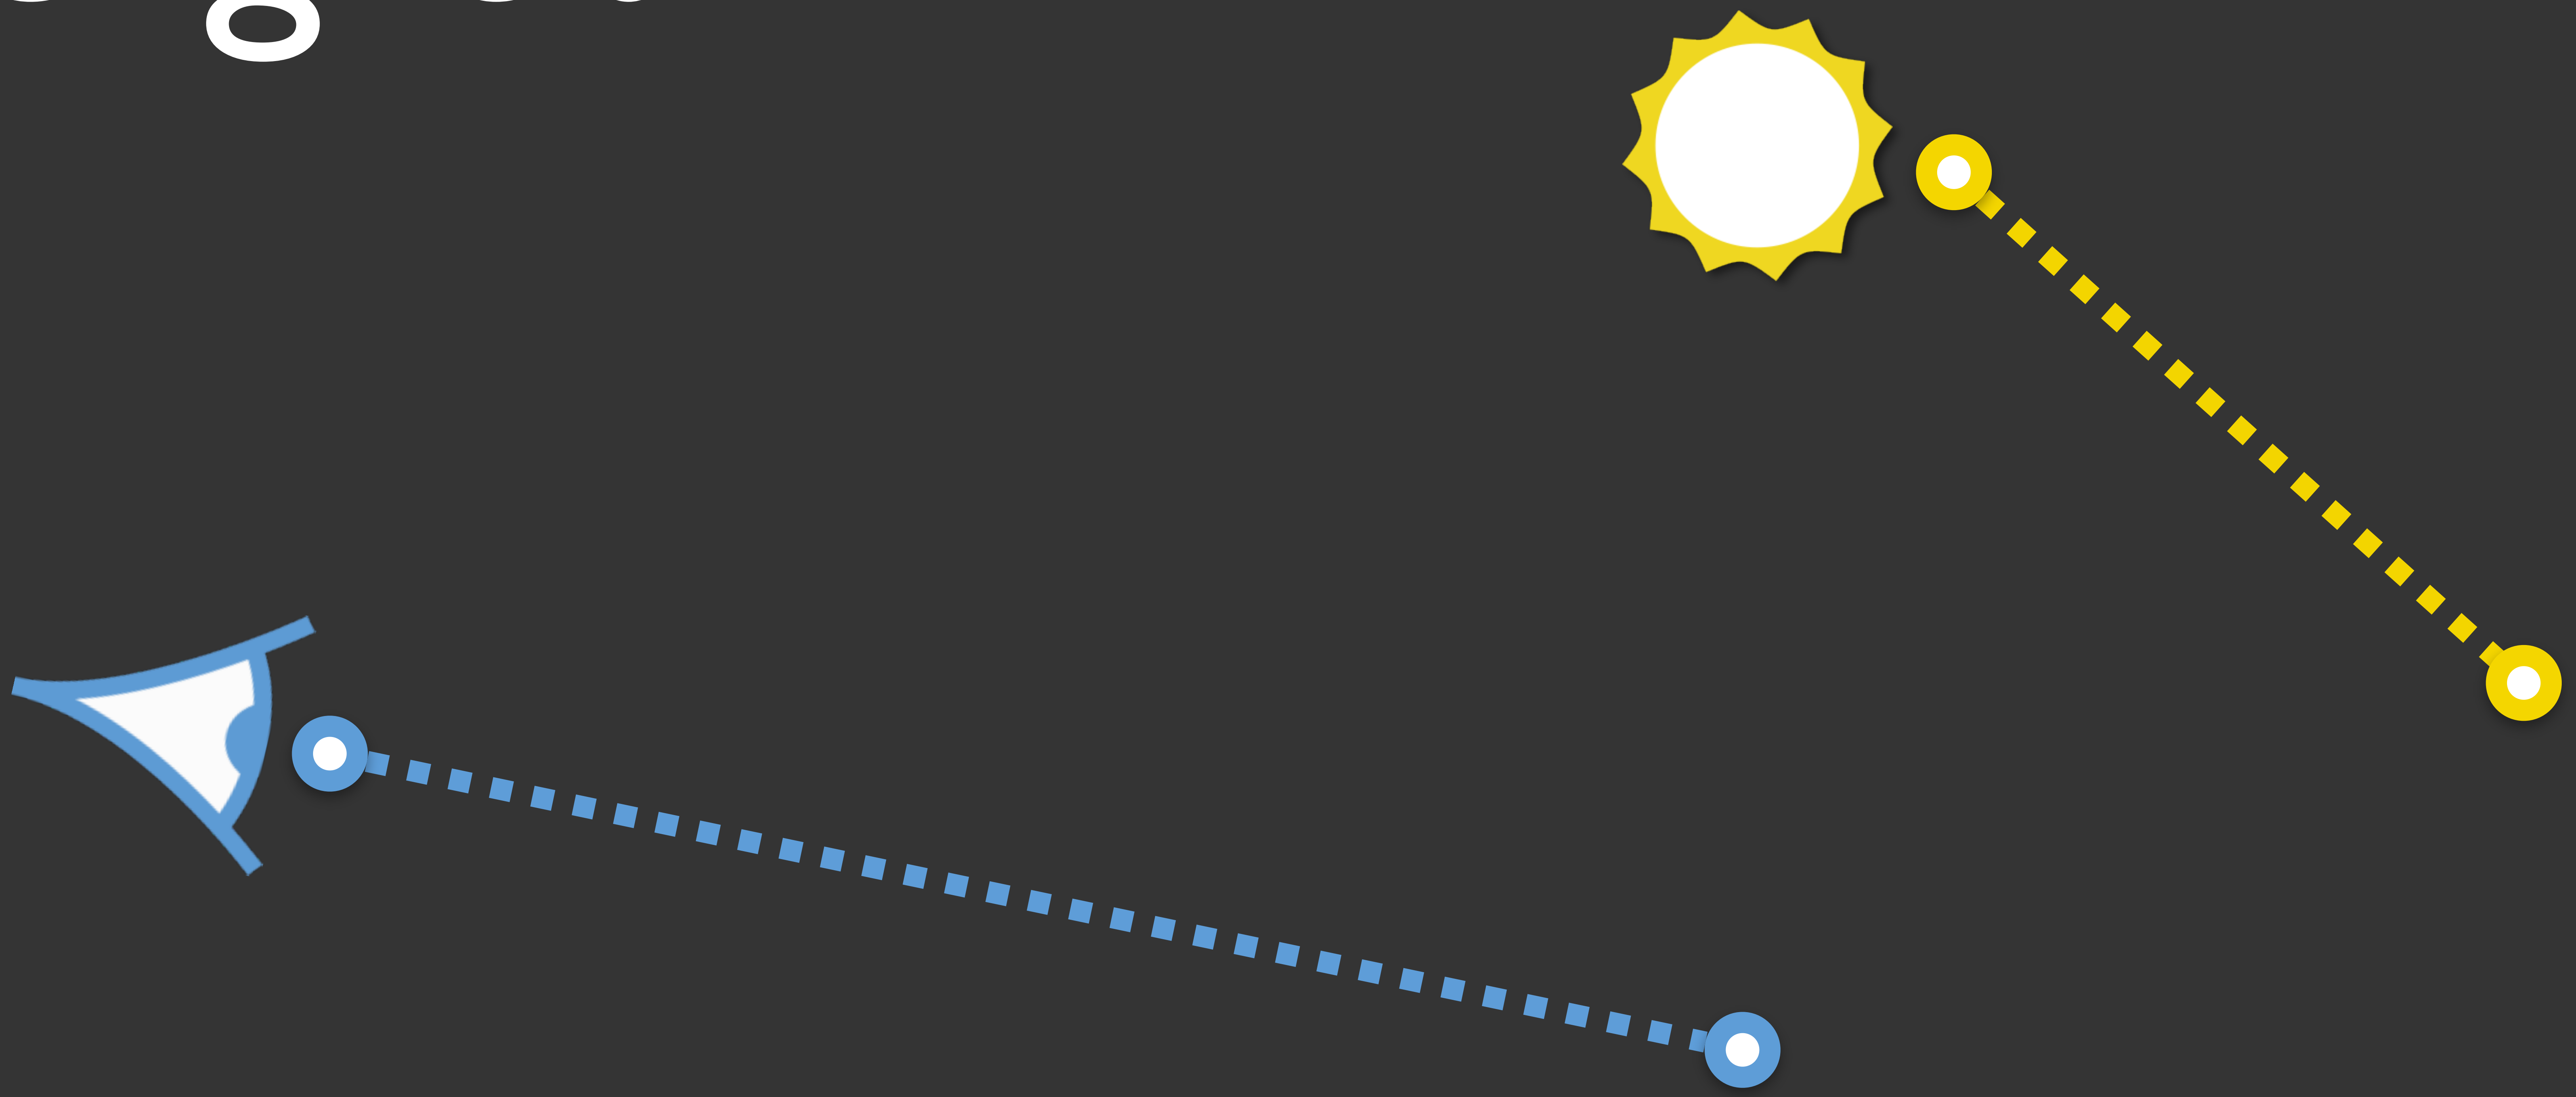

Background

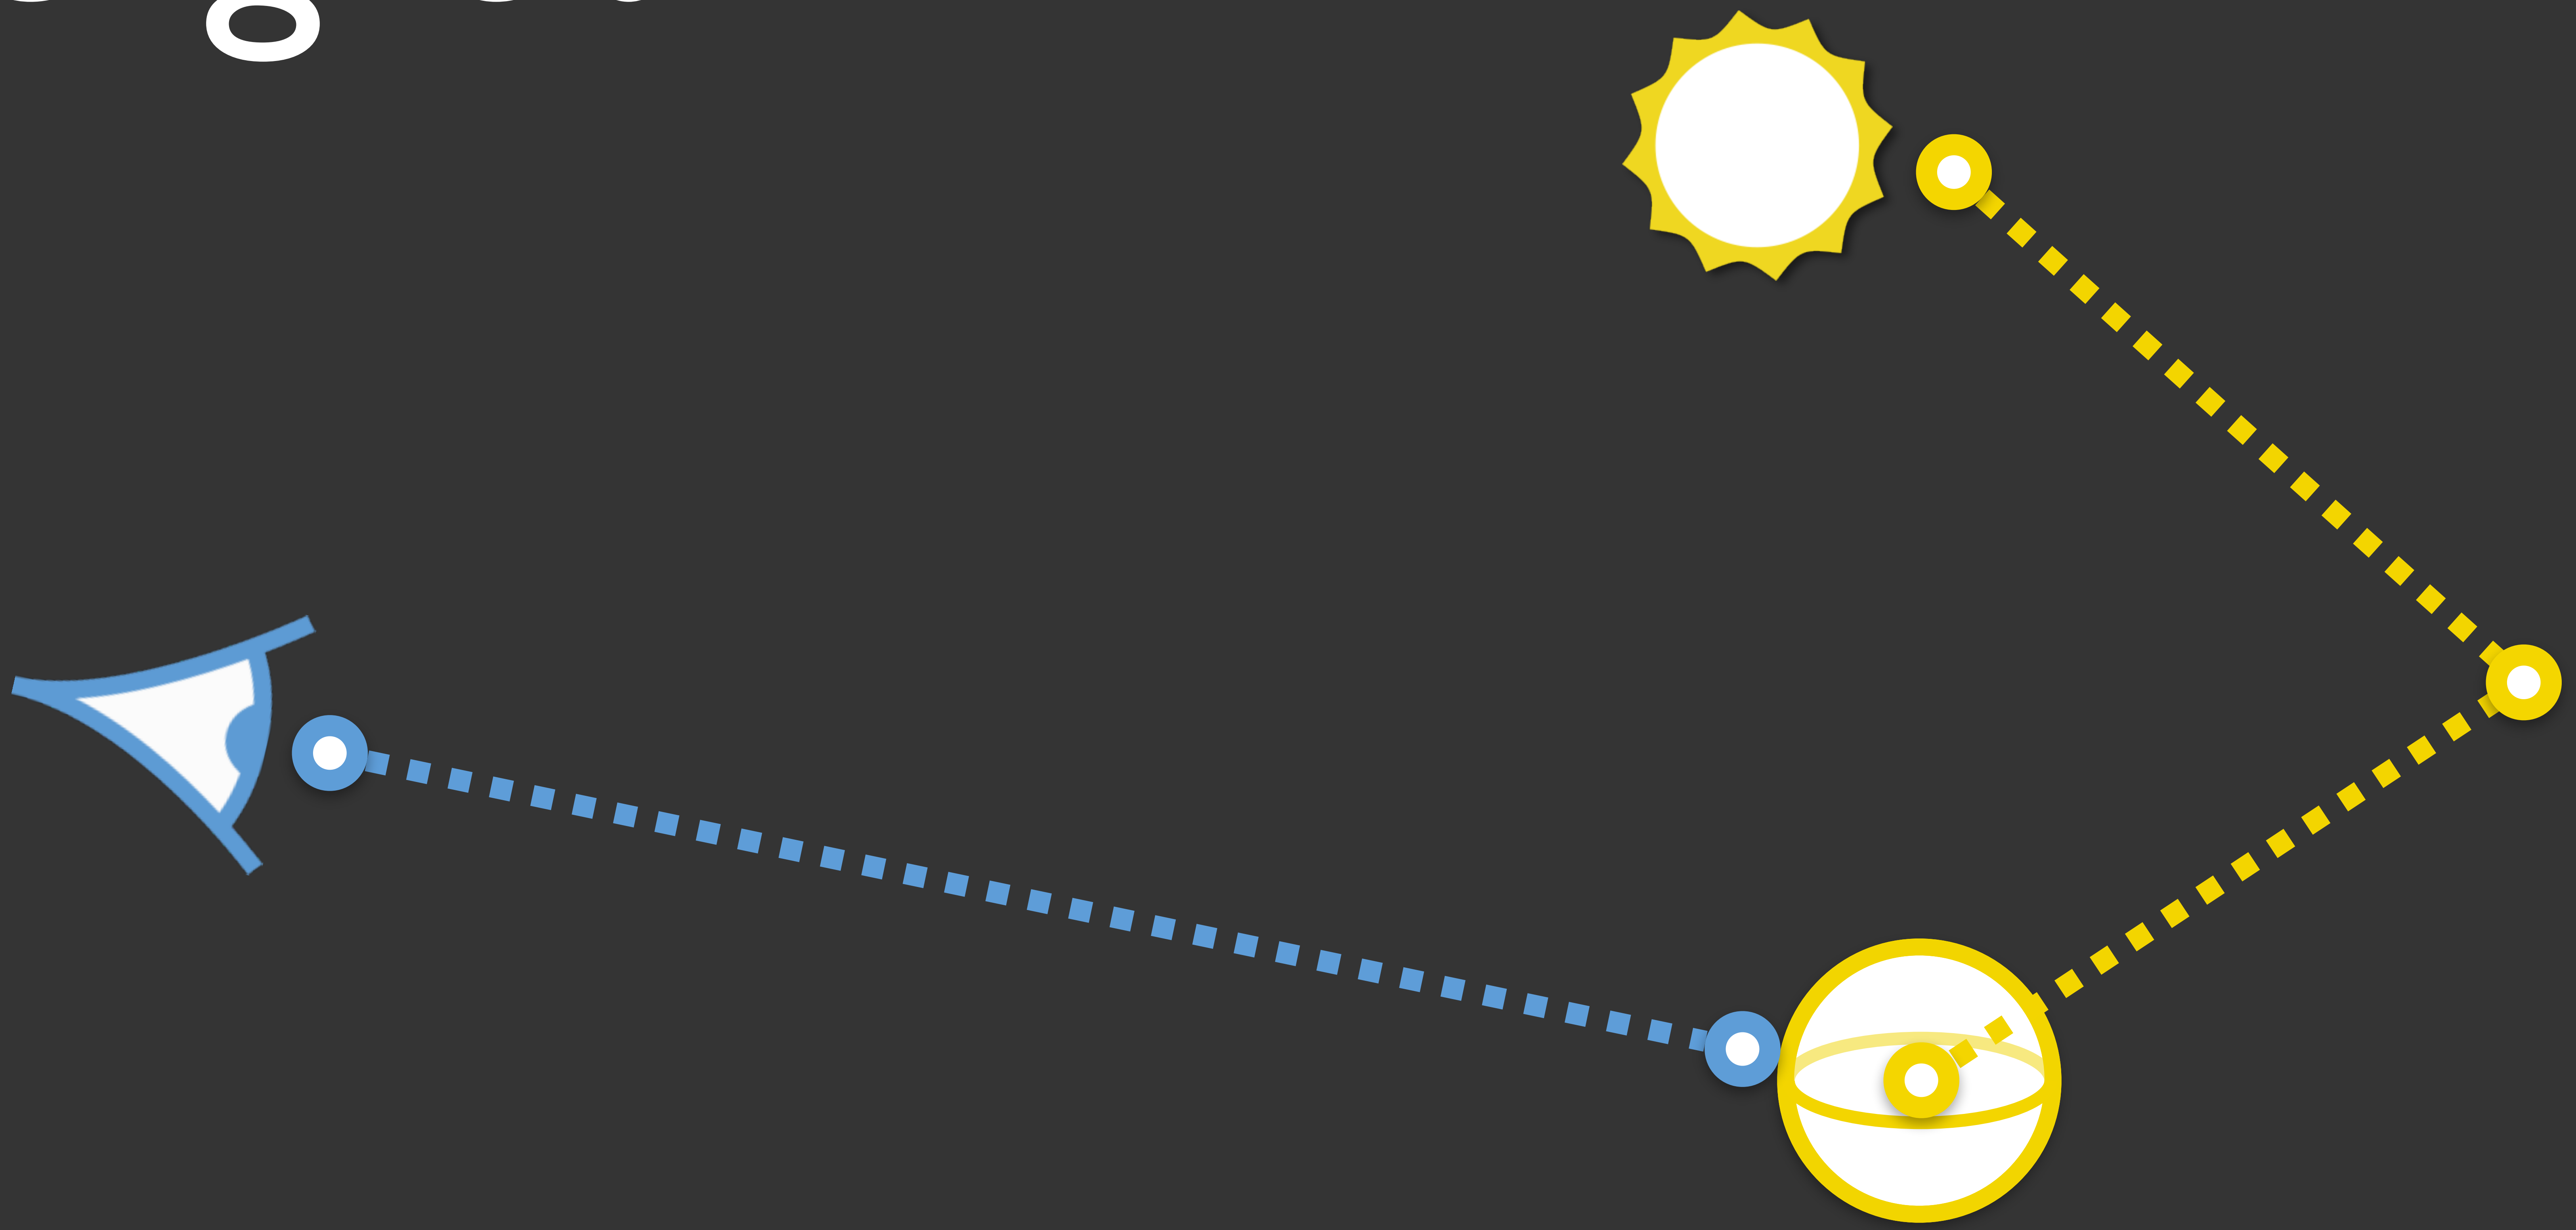

Background

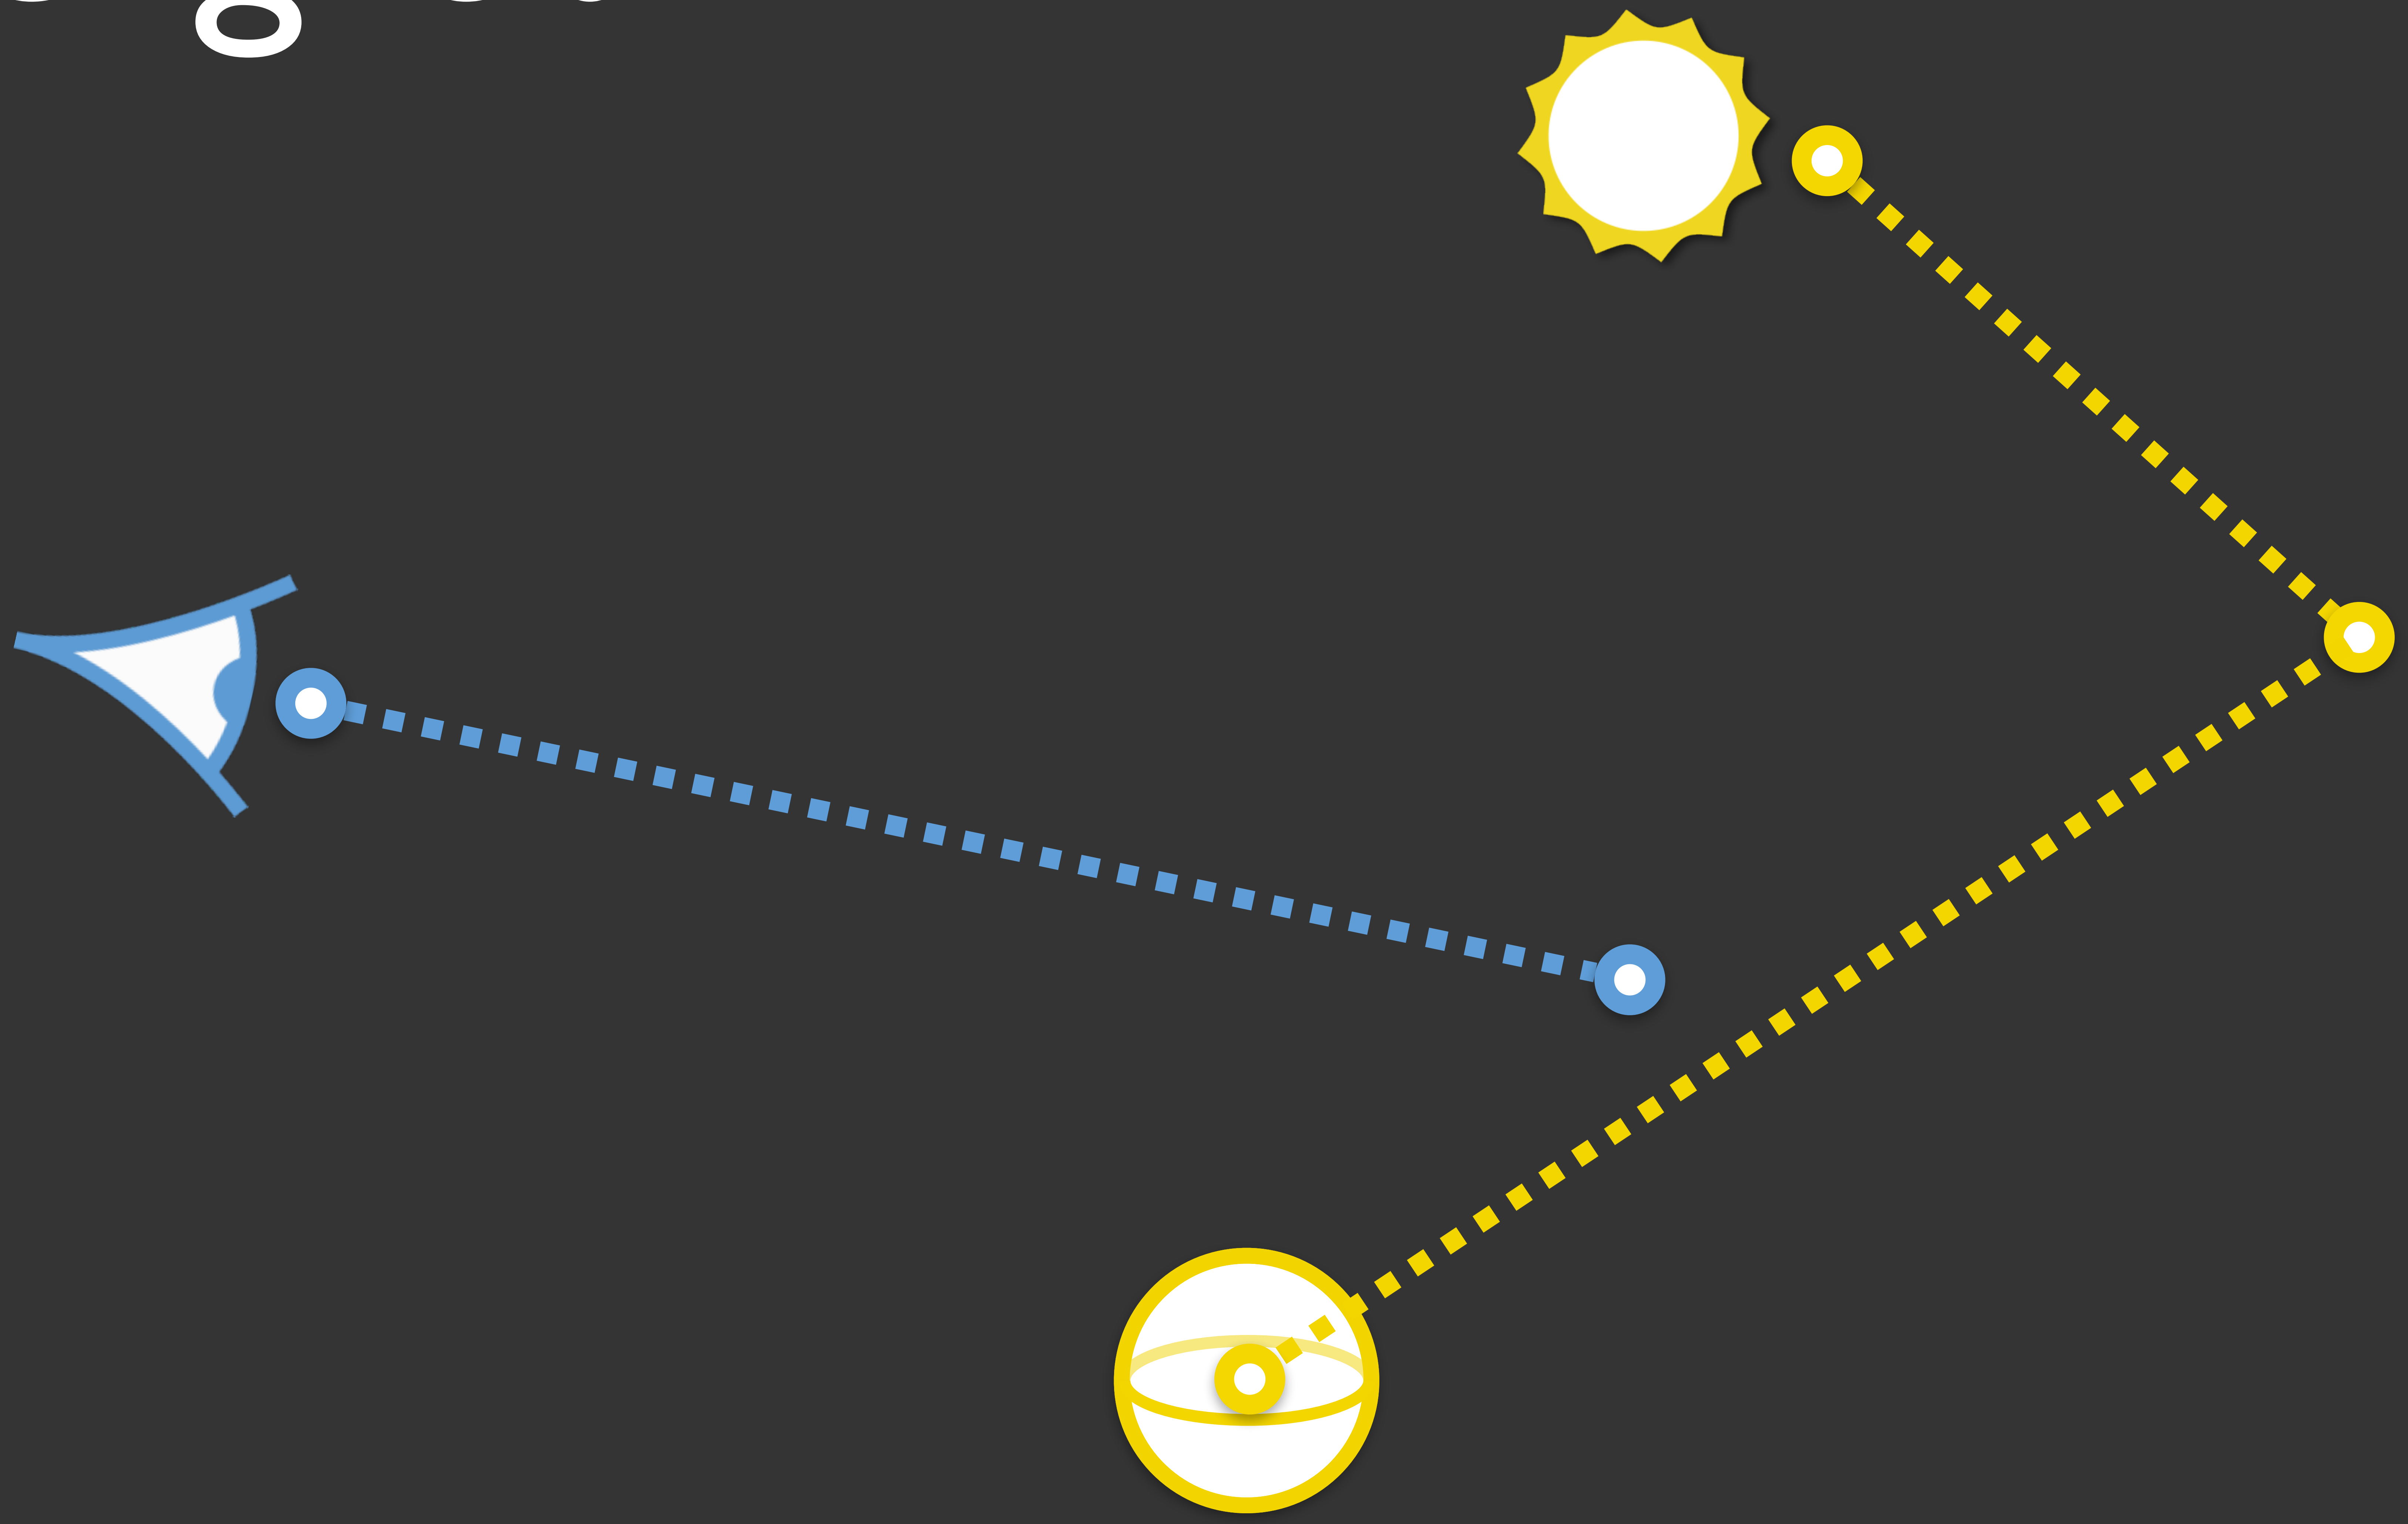

Background

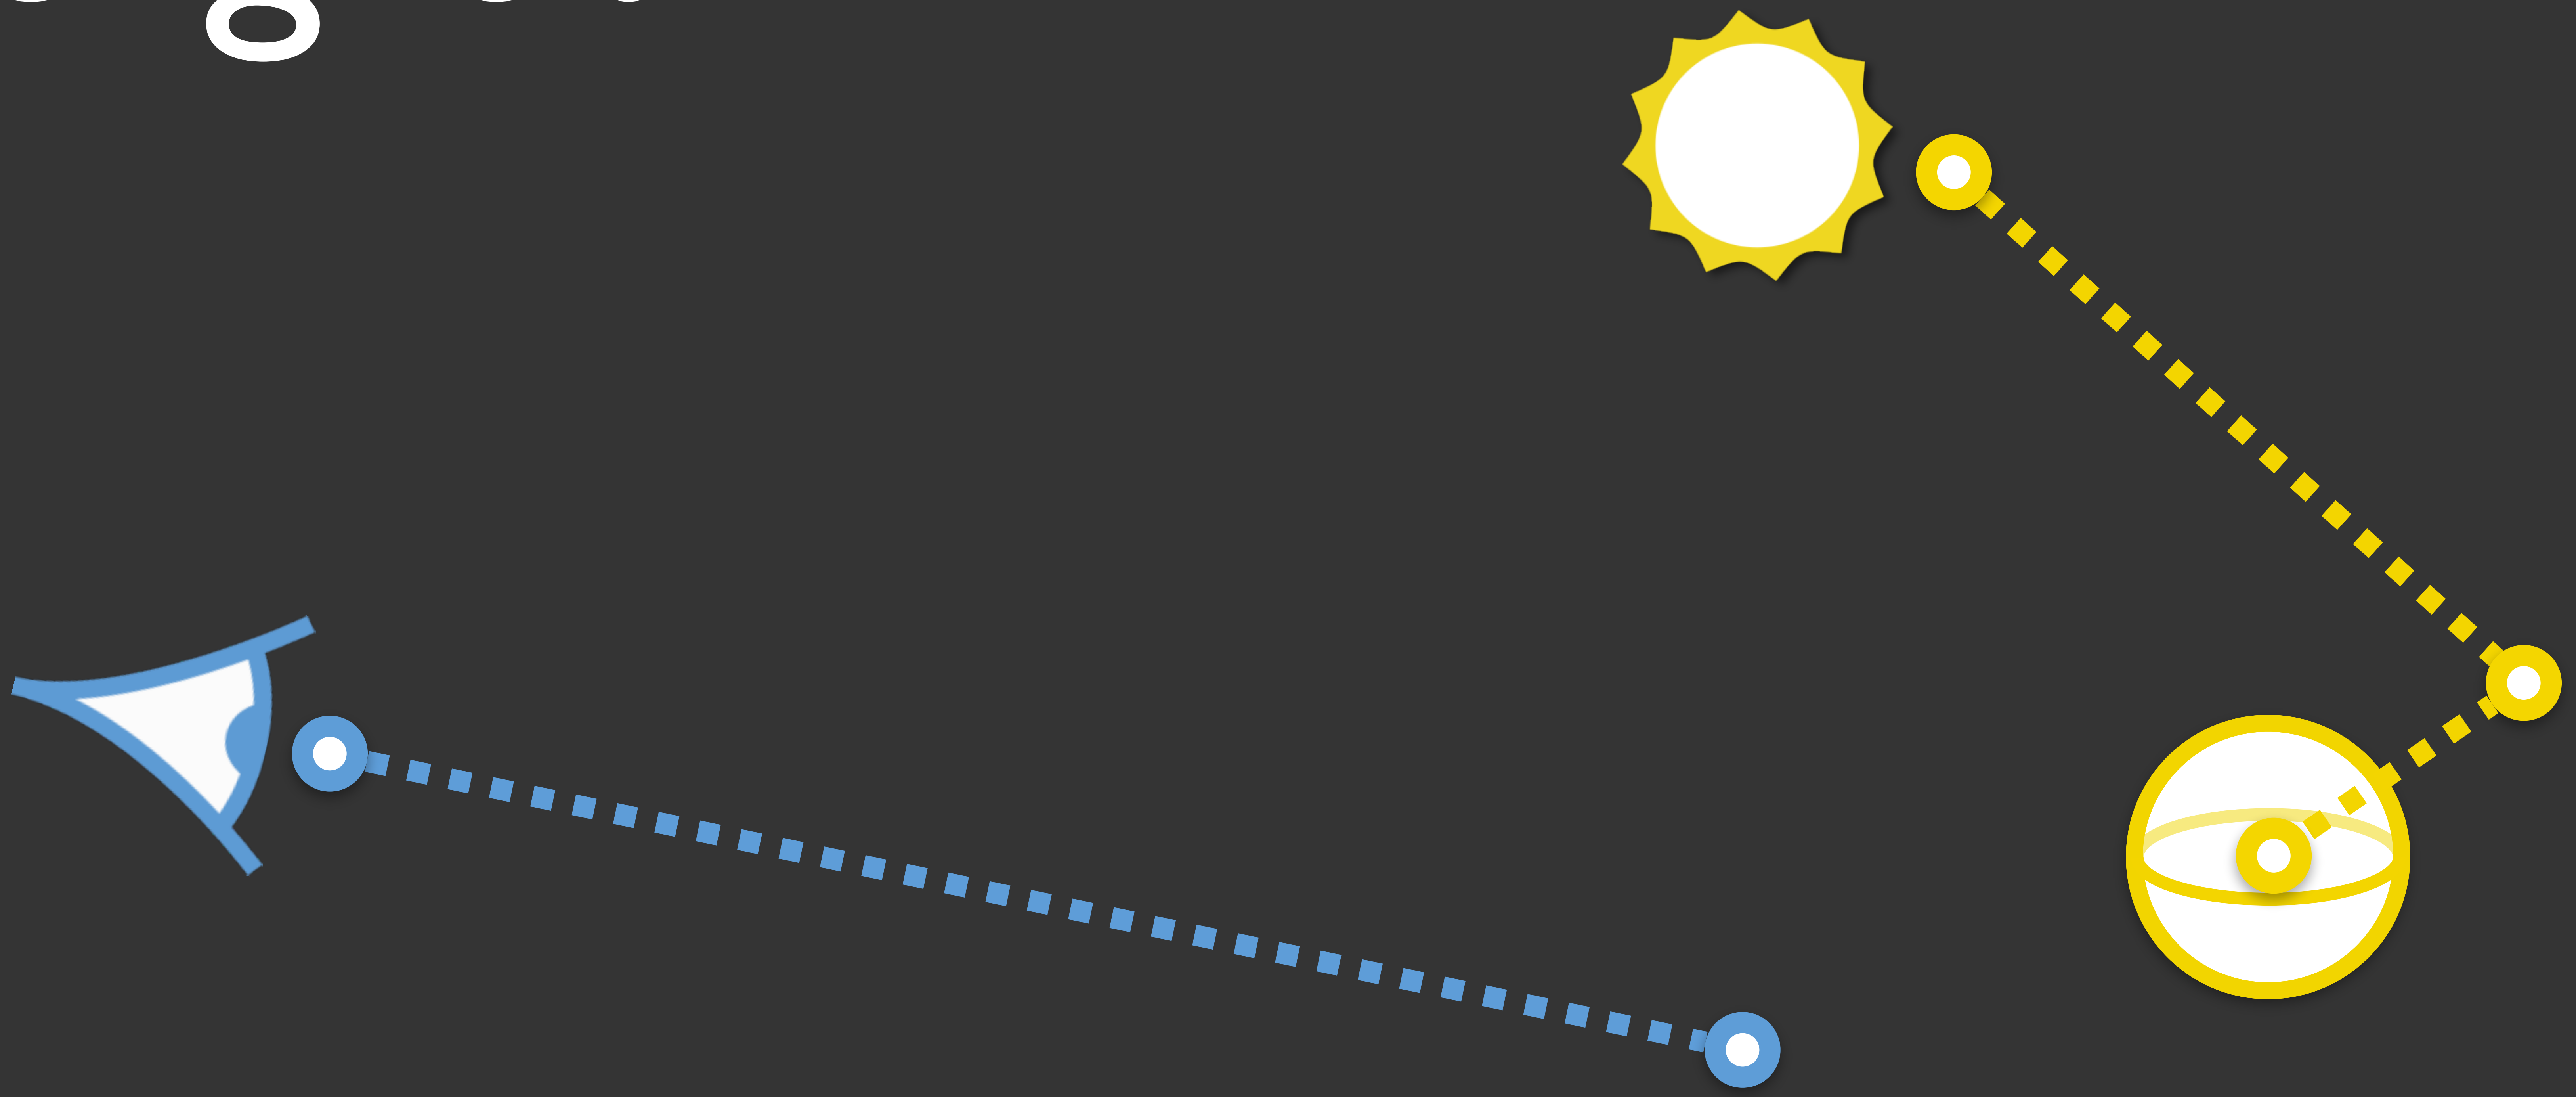

Photon Marching

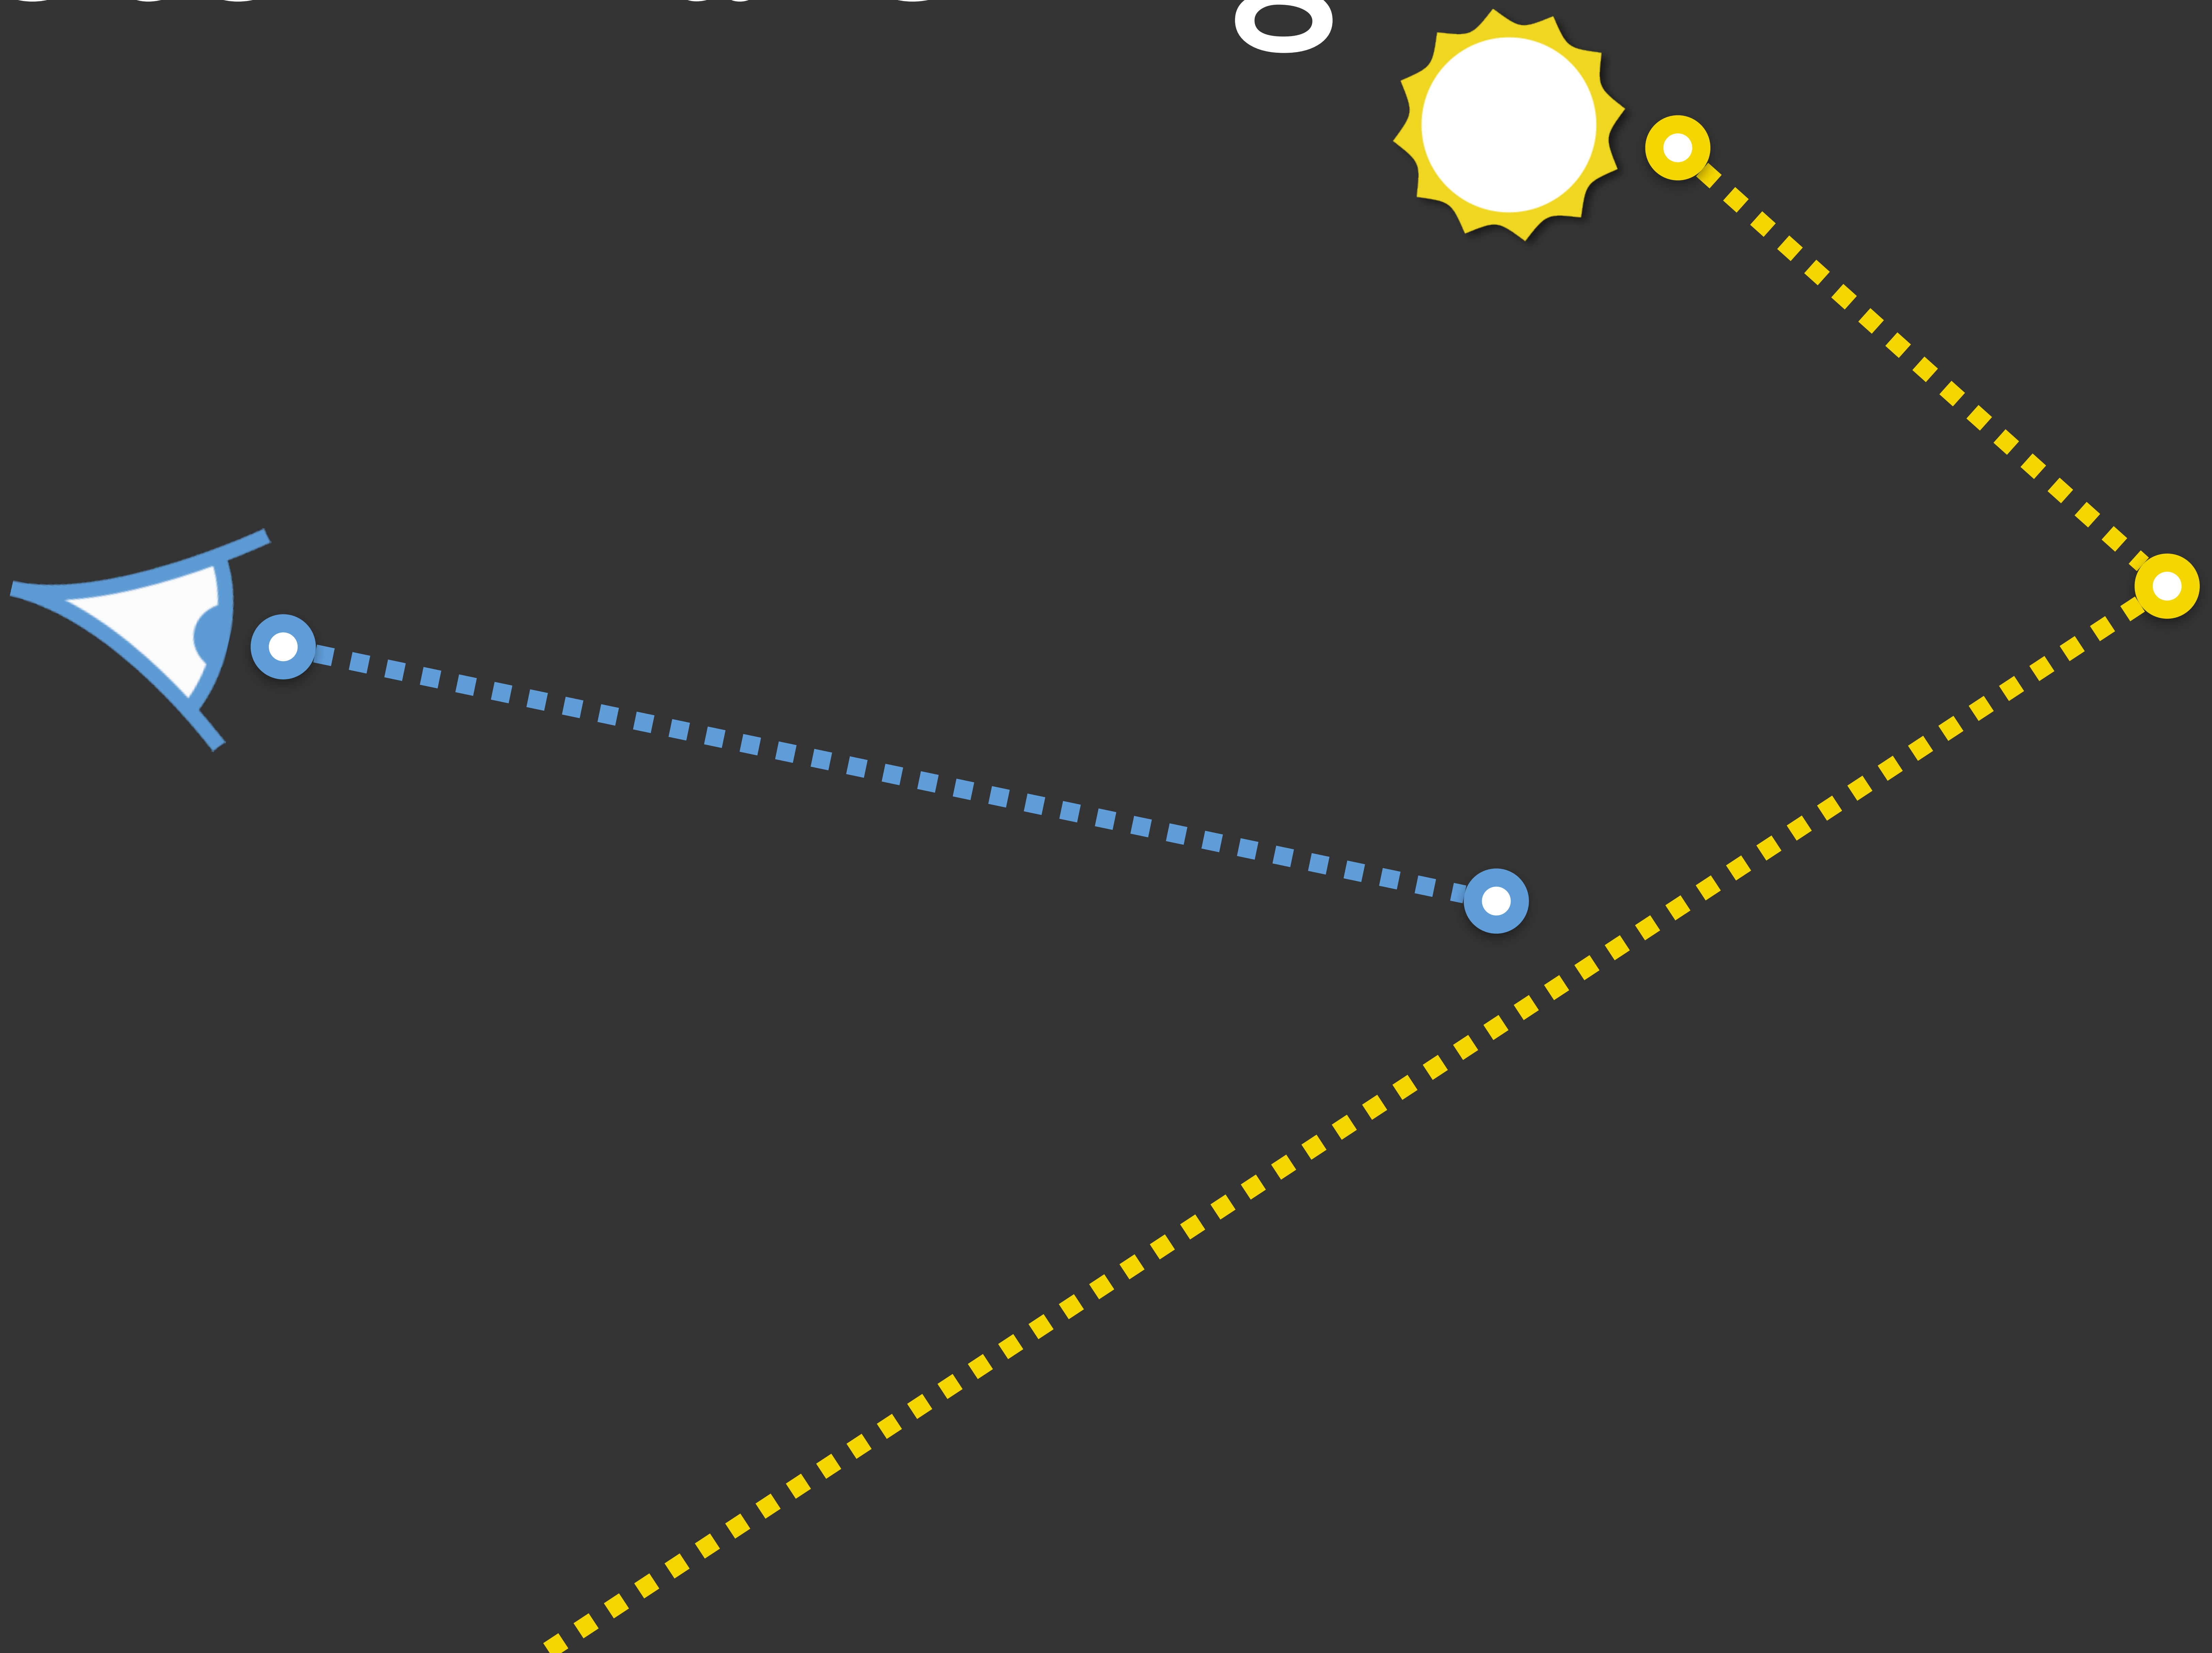

Photon Marching

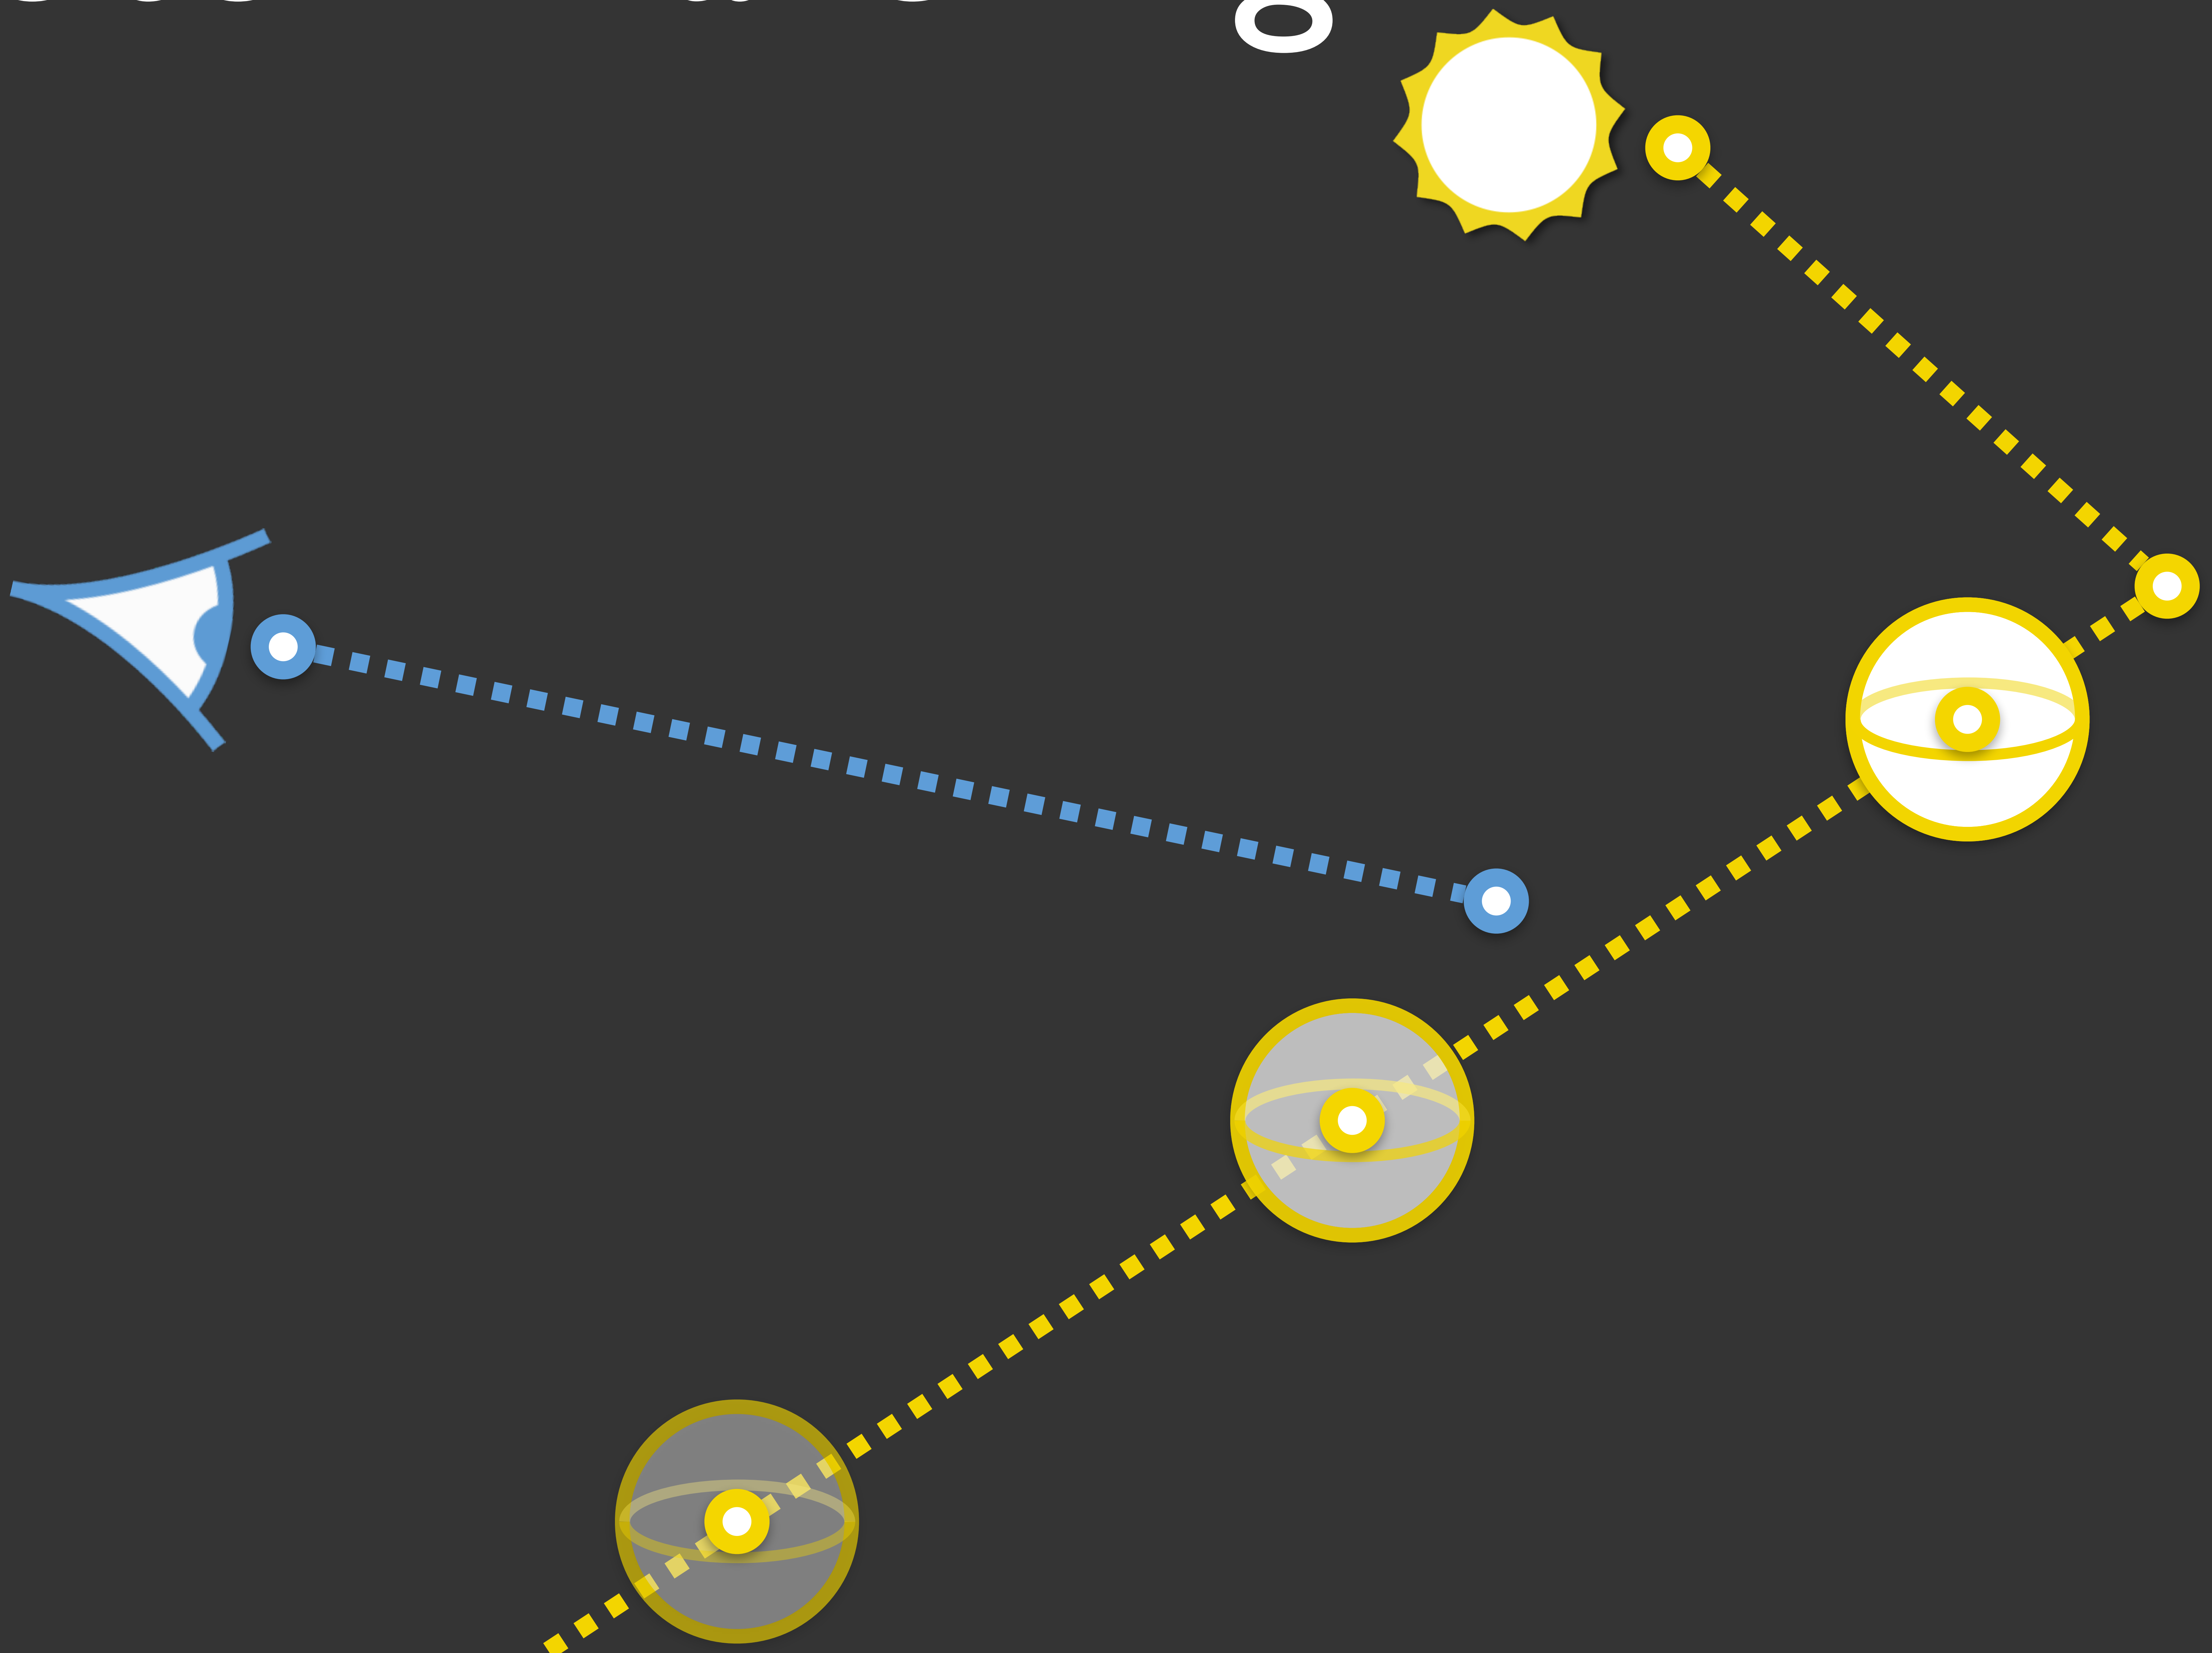

Photon Marching

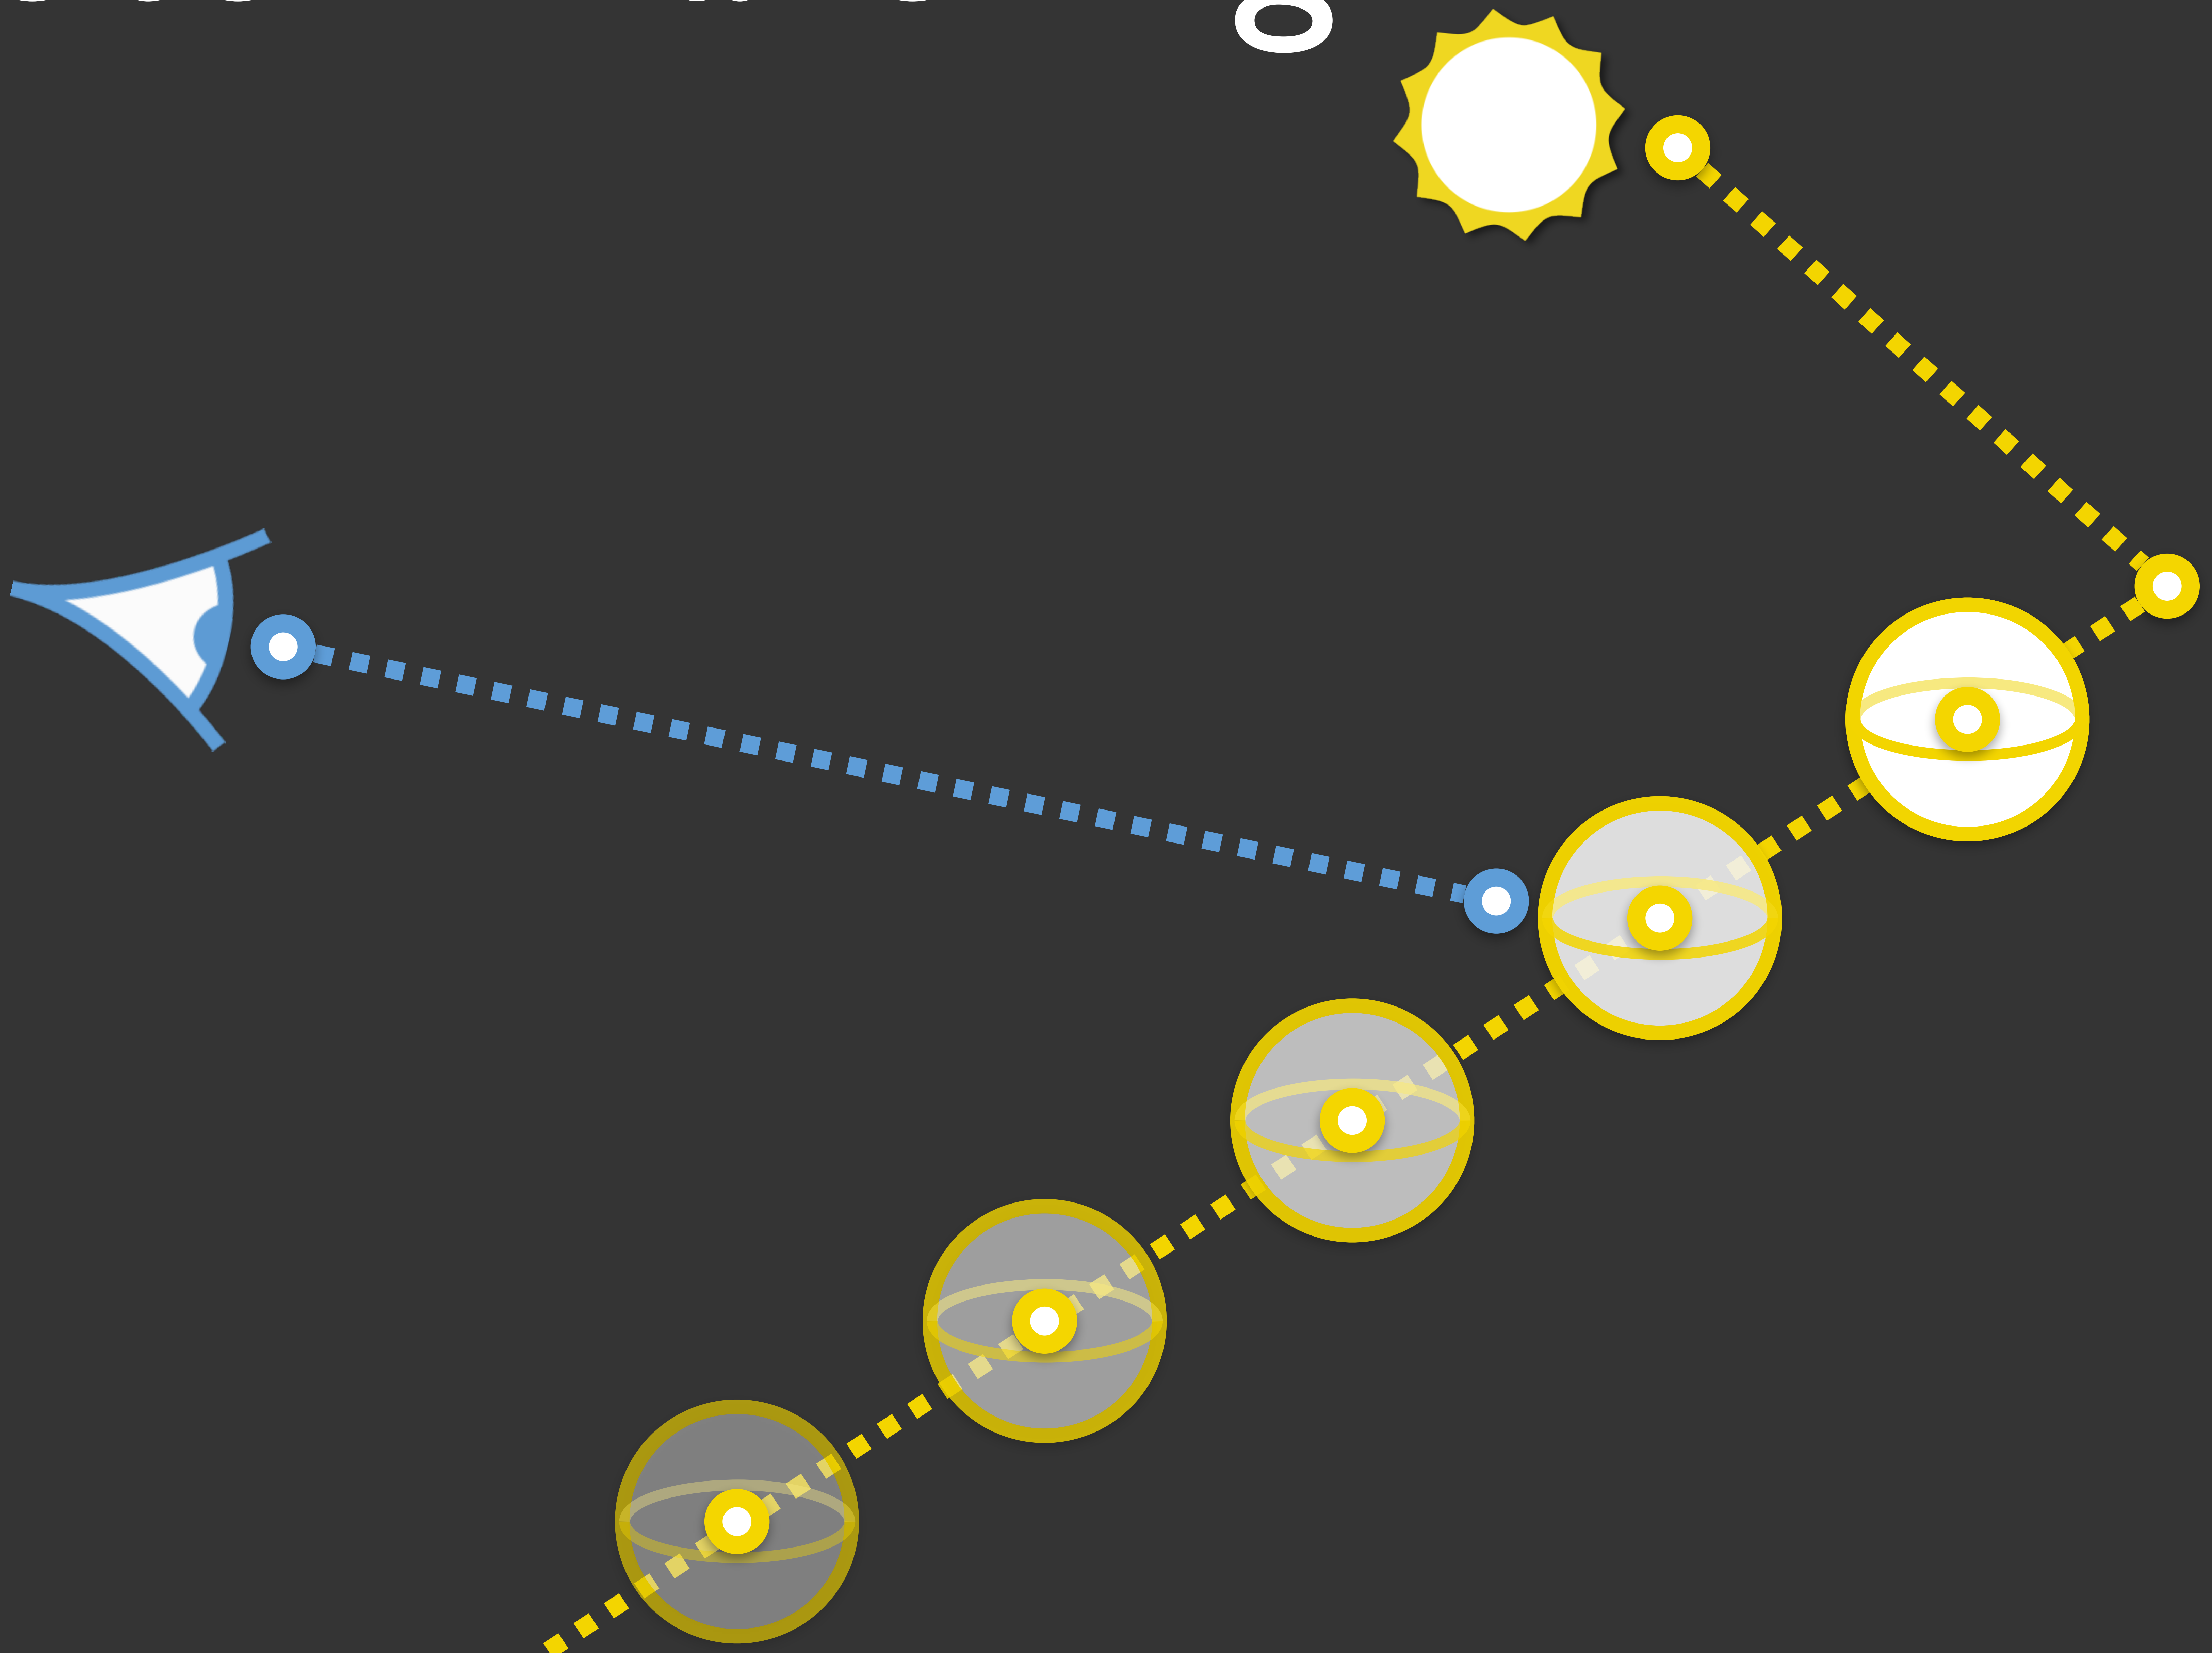

Photon Beams

[Jarosz et al. '11]

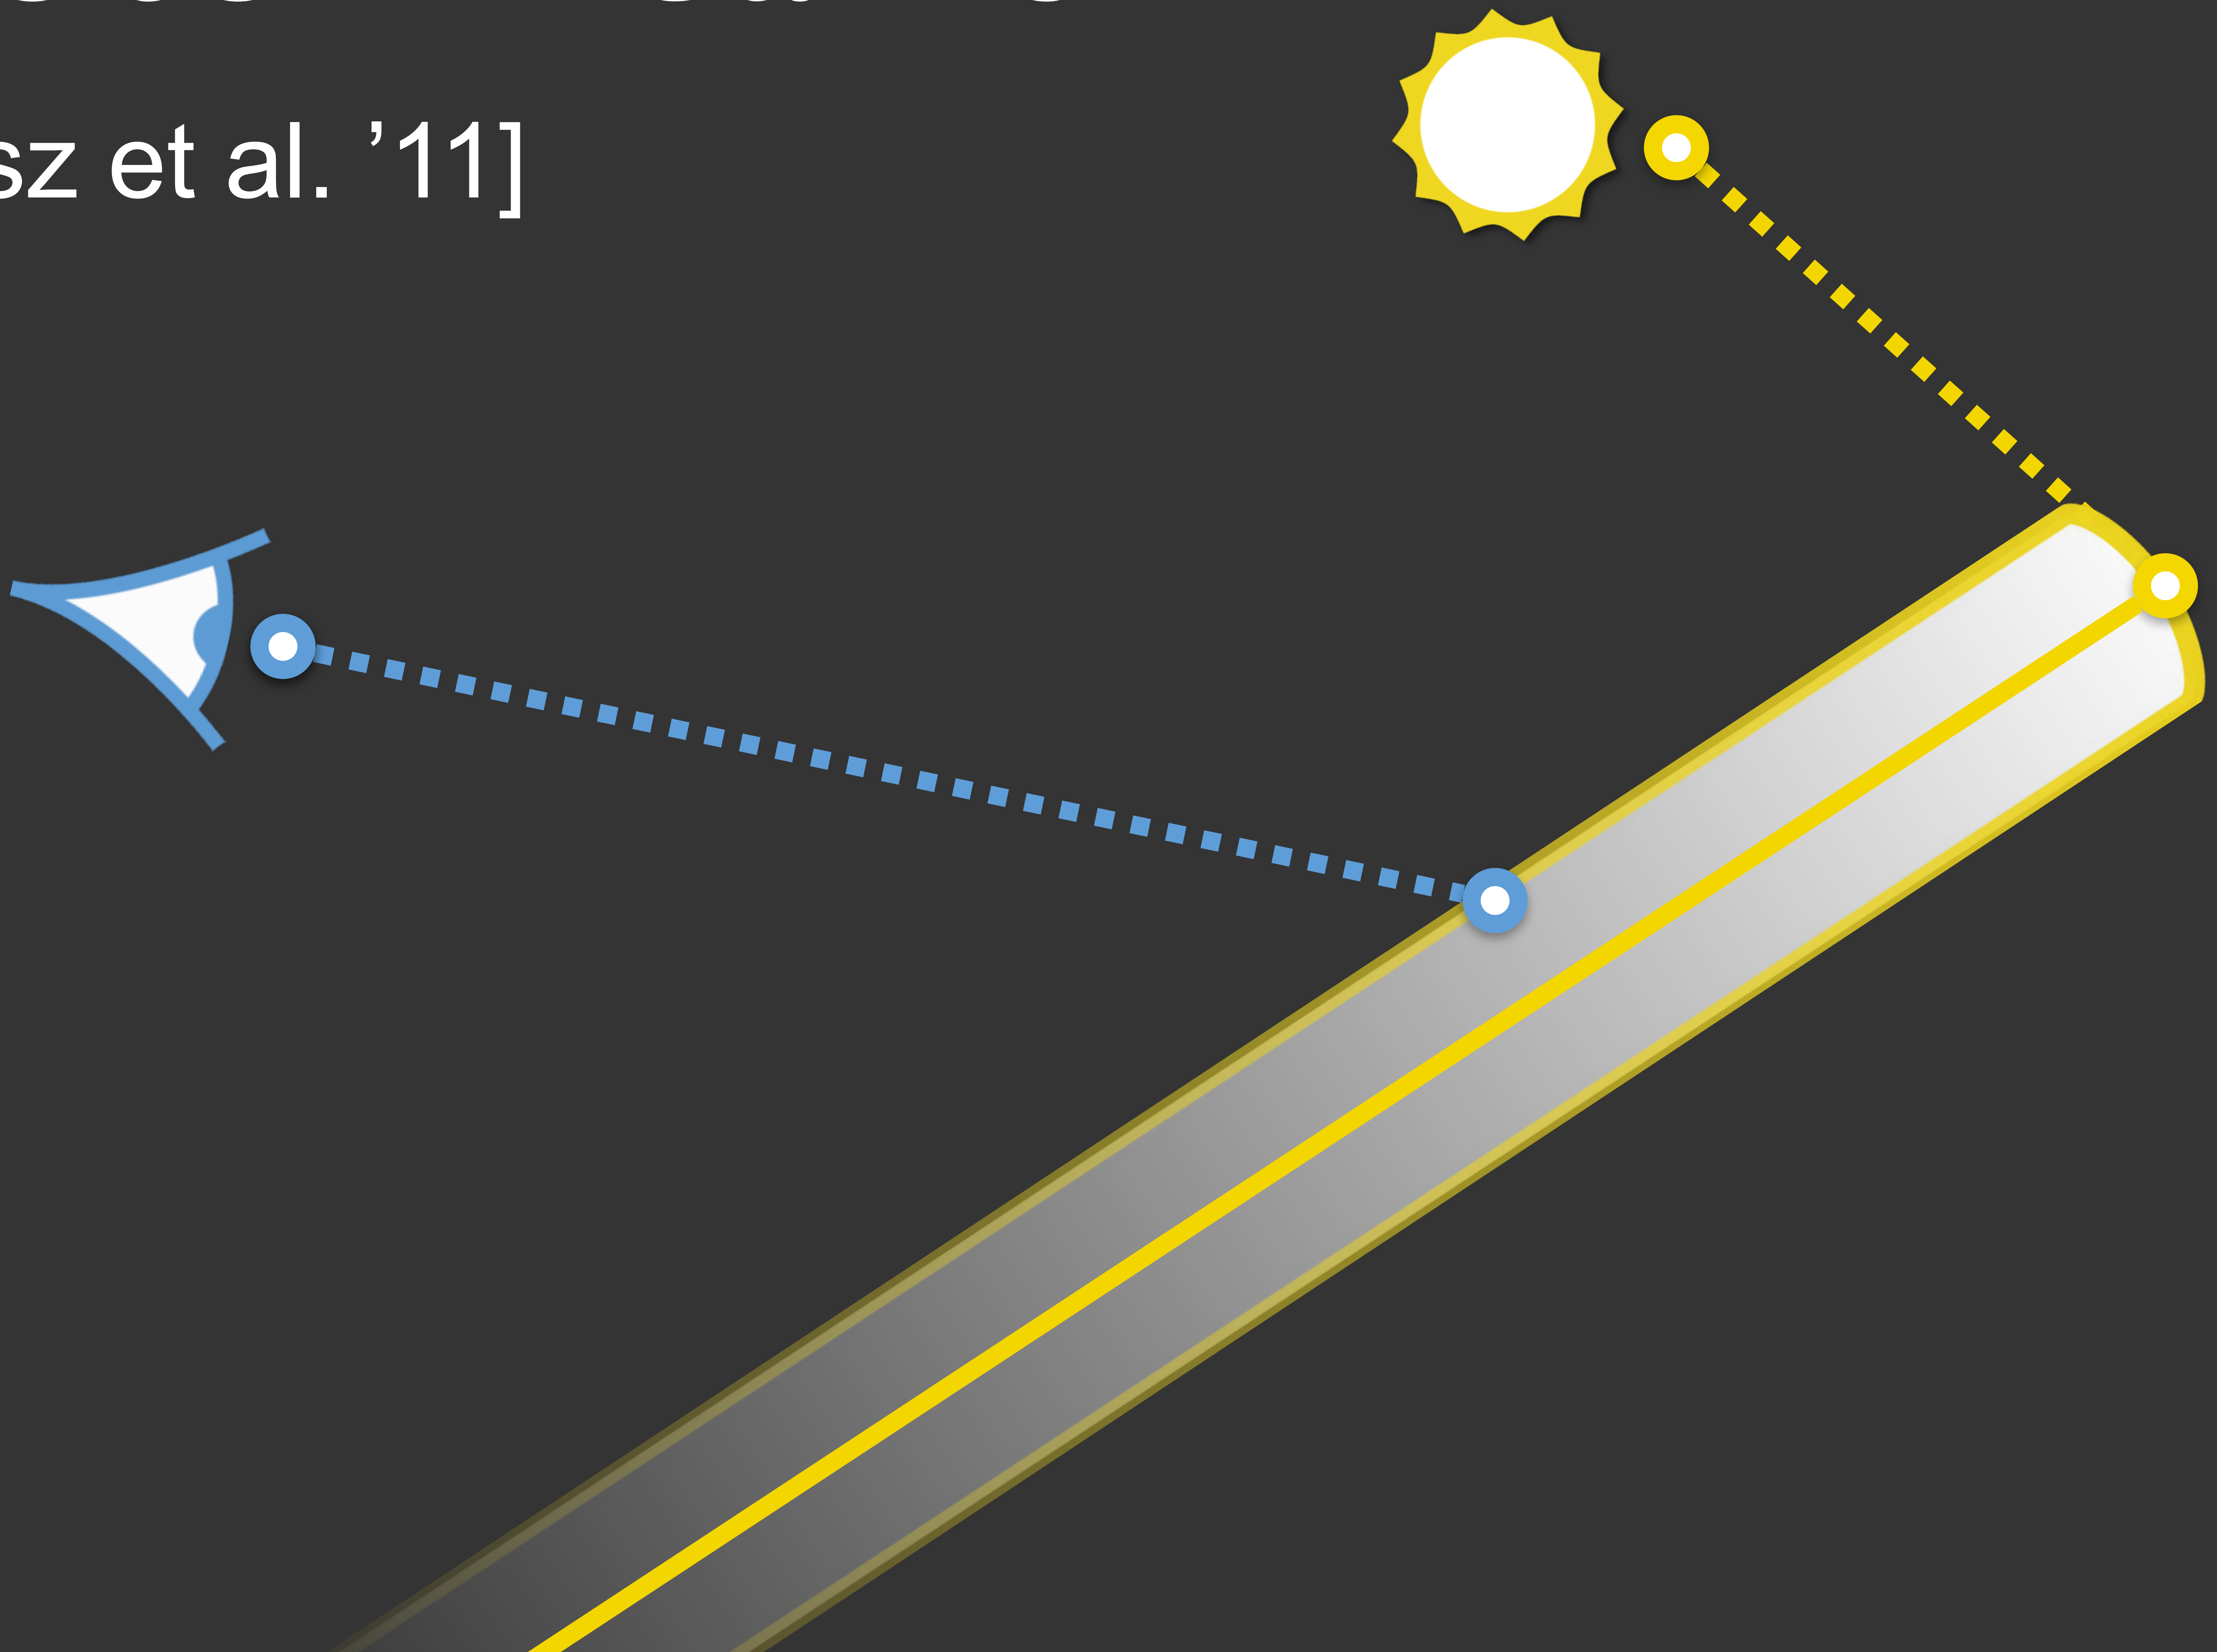

Photon Beams

[Jarosz et al. ’11]

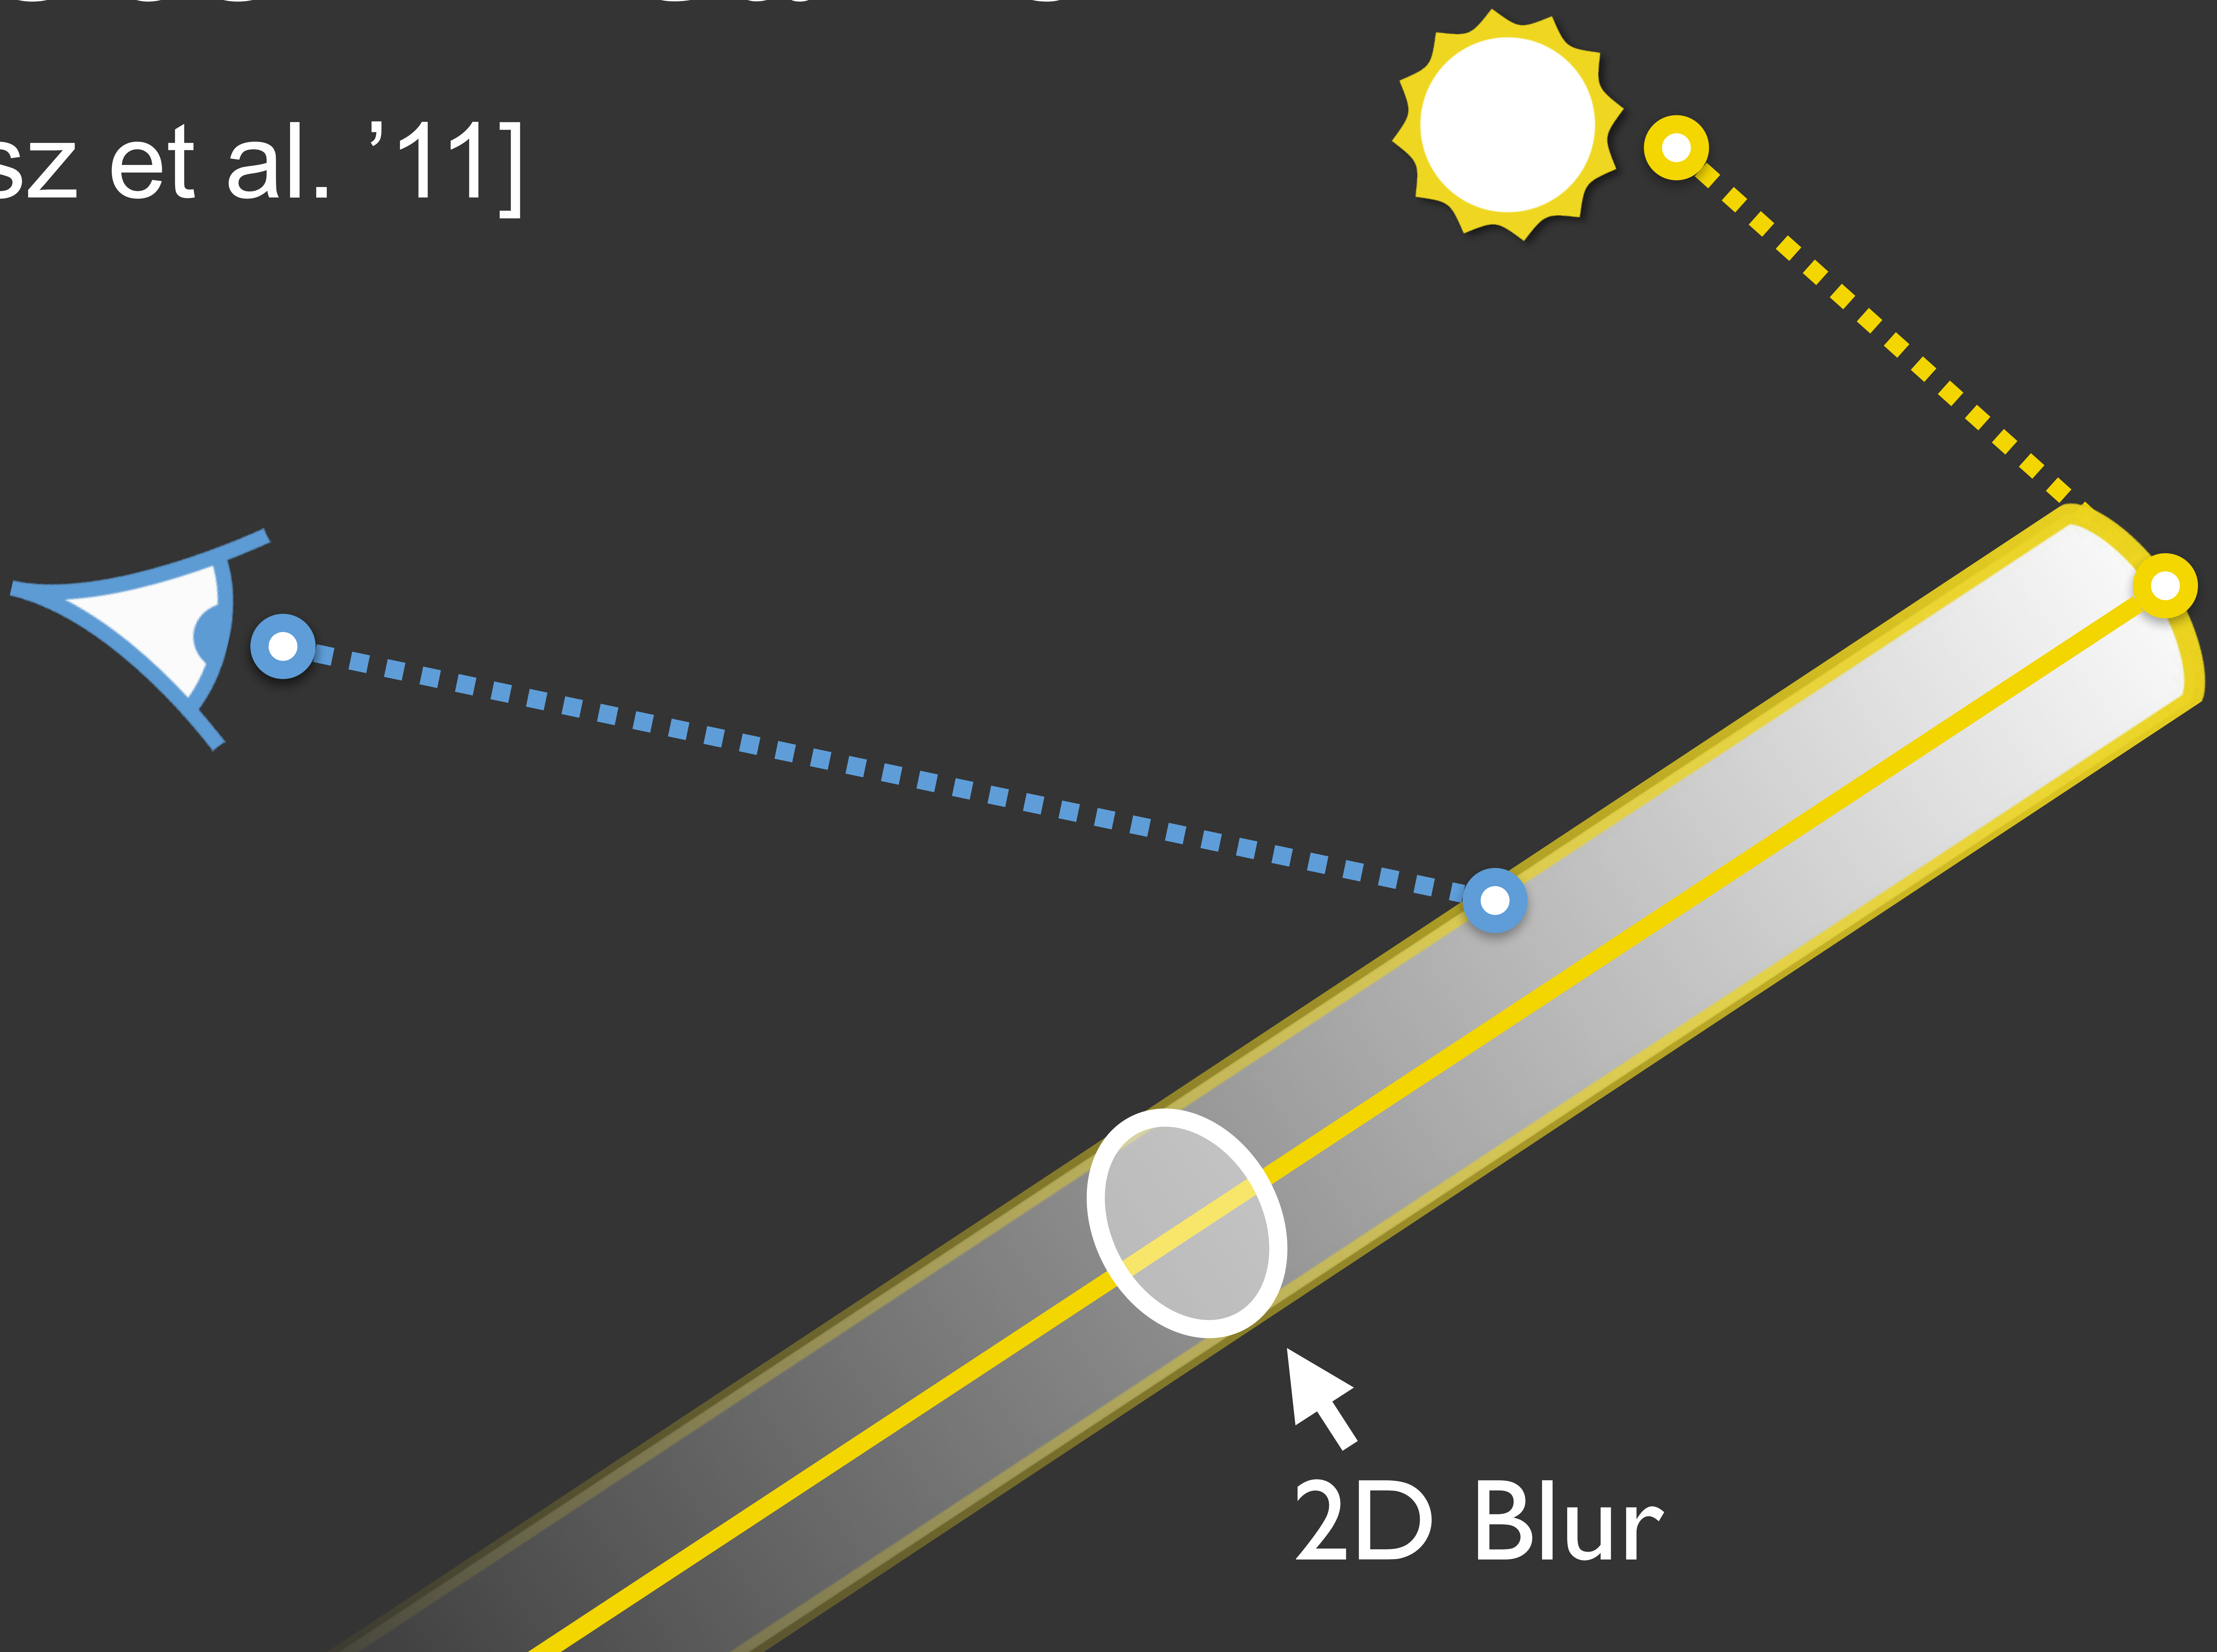

Photon Beams

[Jarosz et al. ’11]

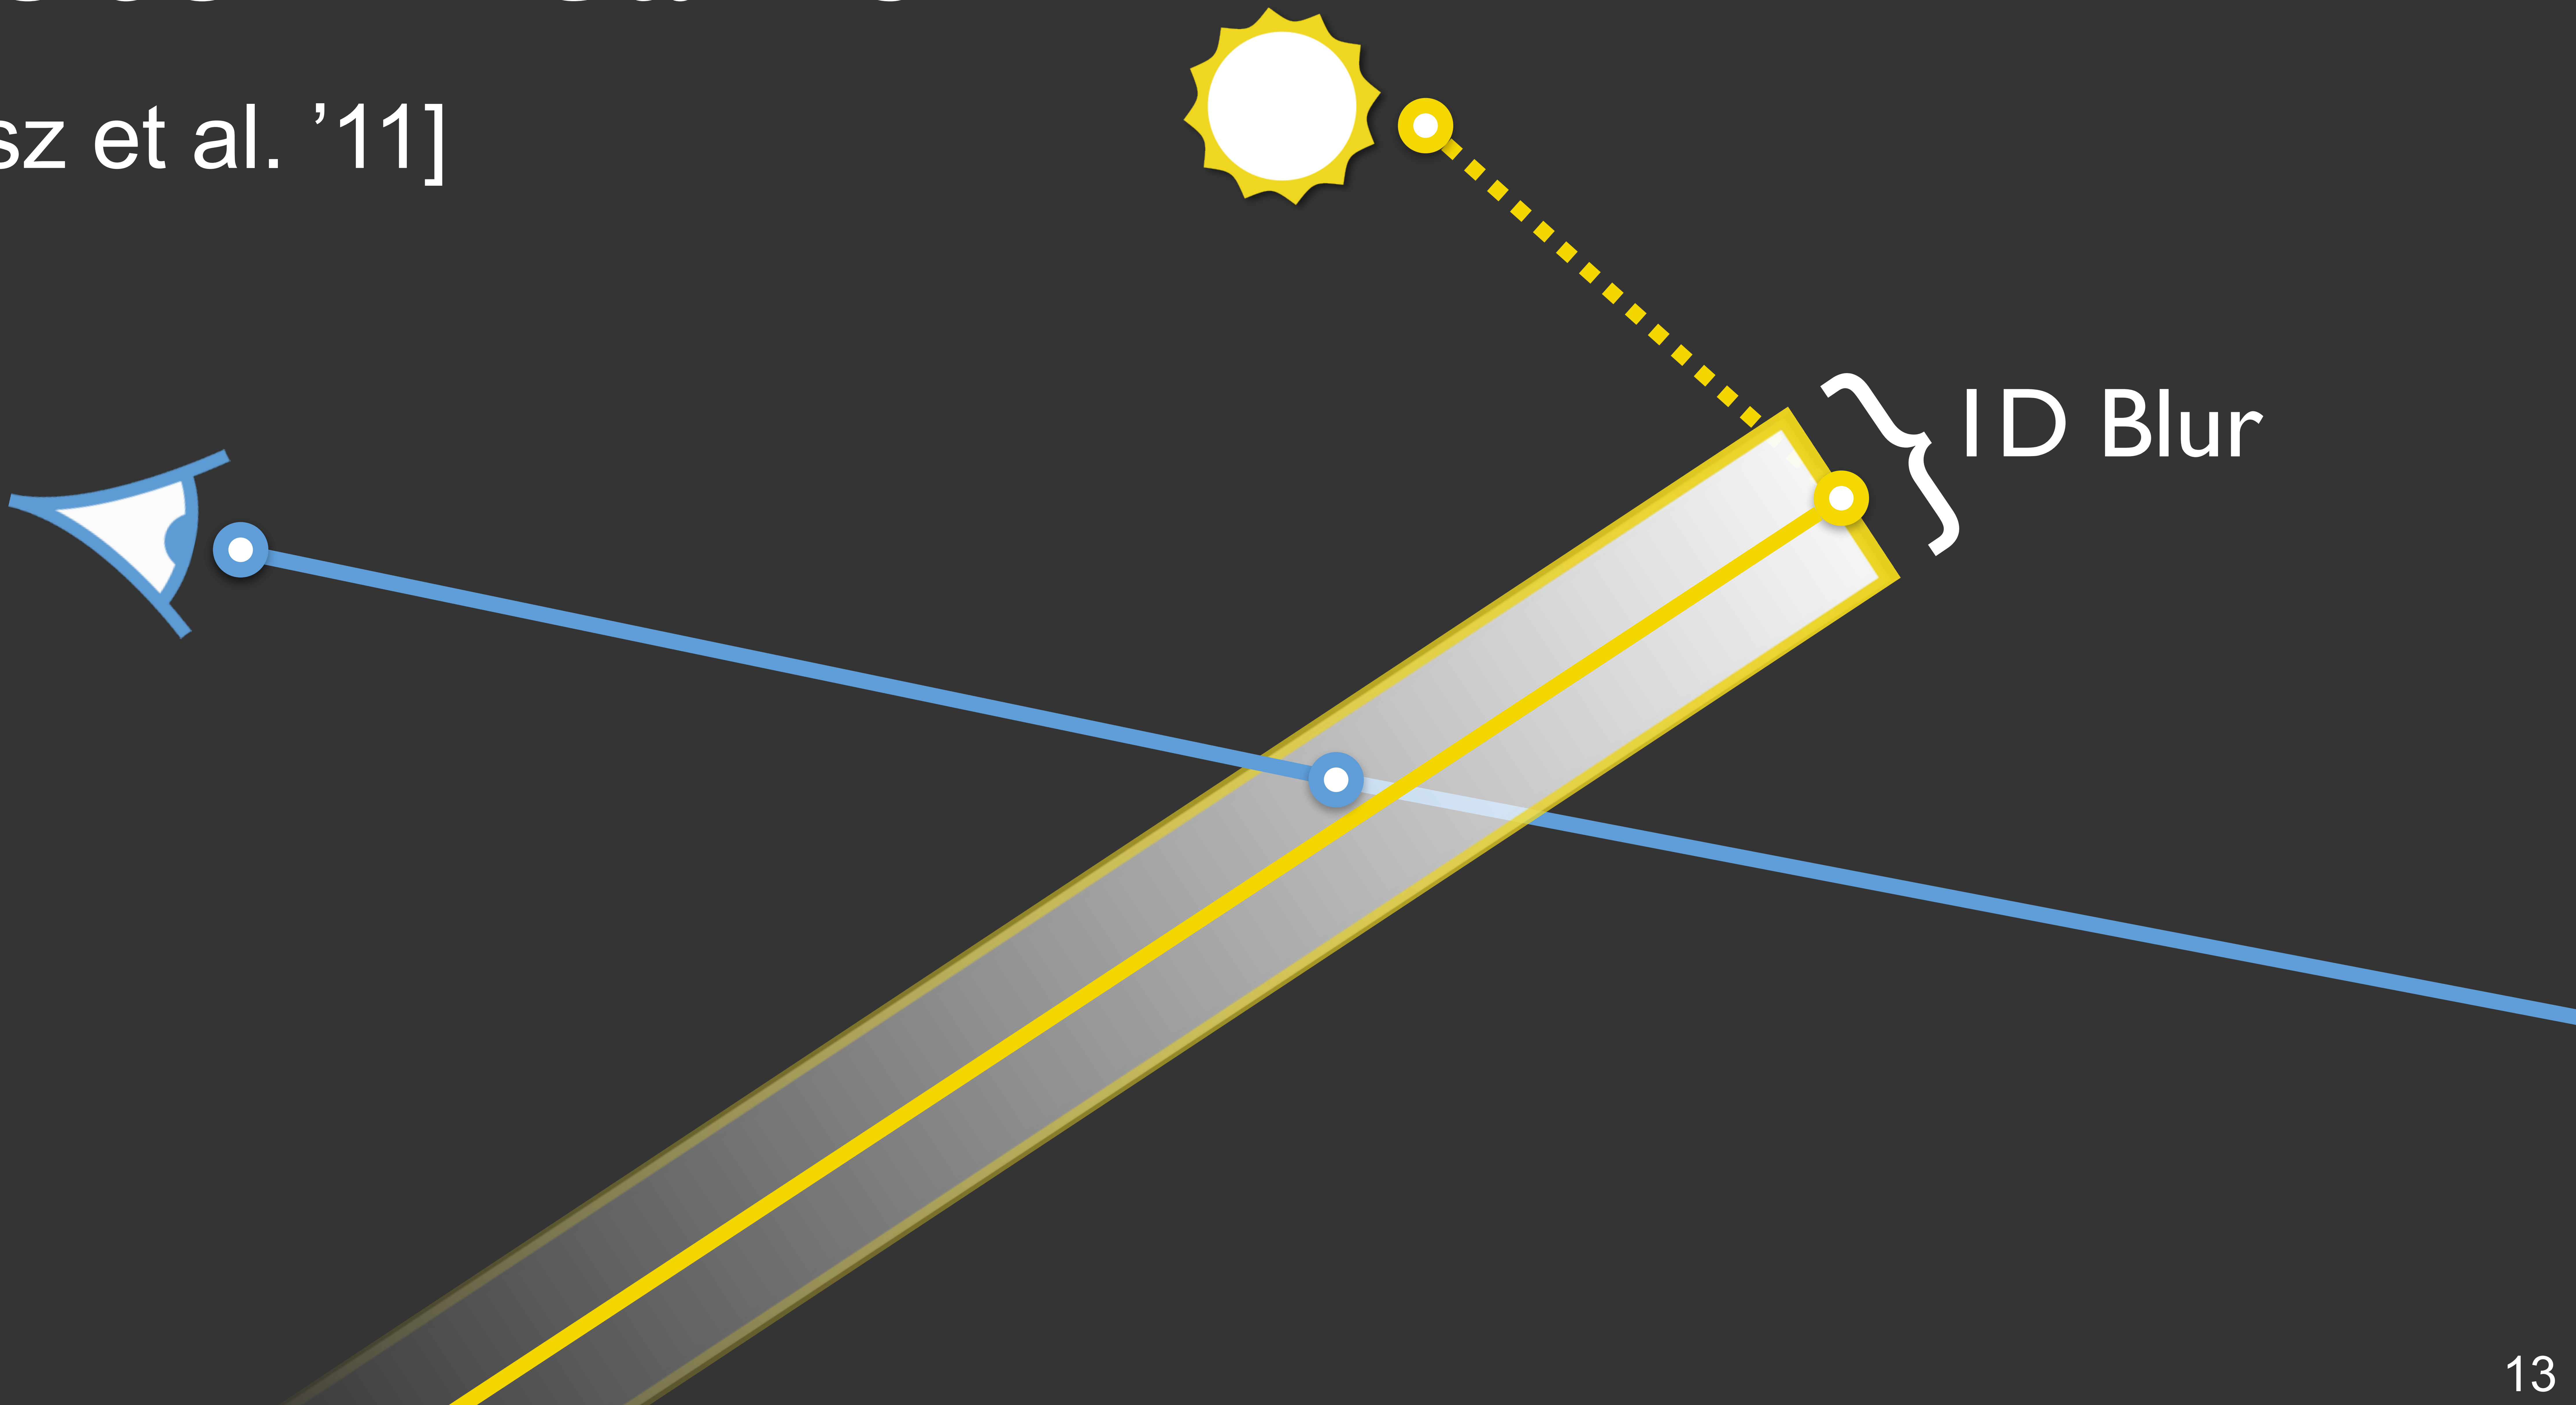

Photon Estimators

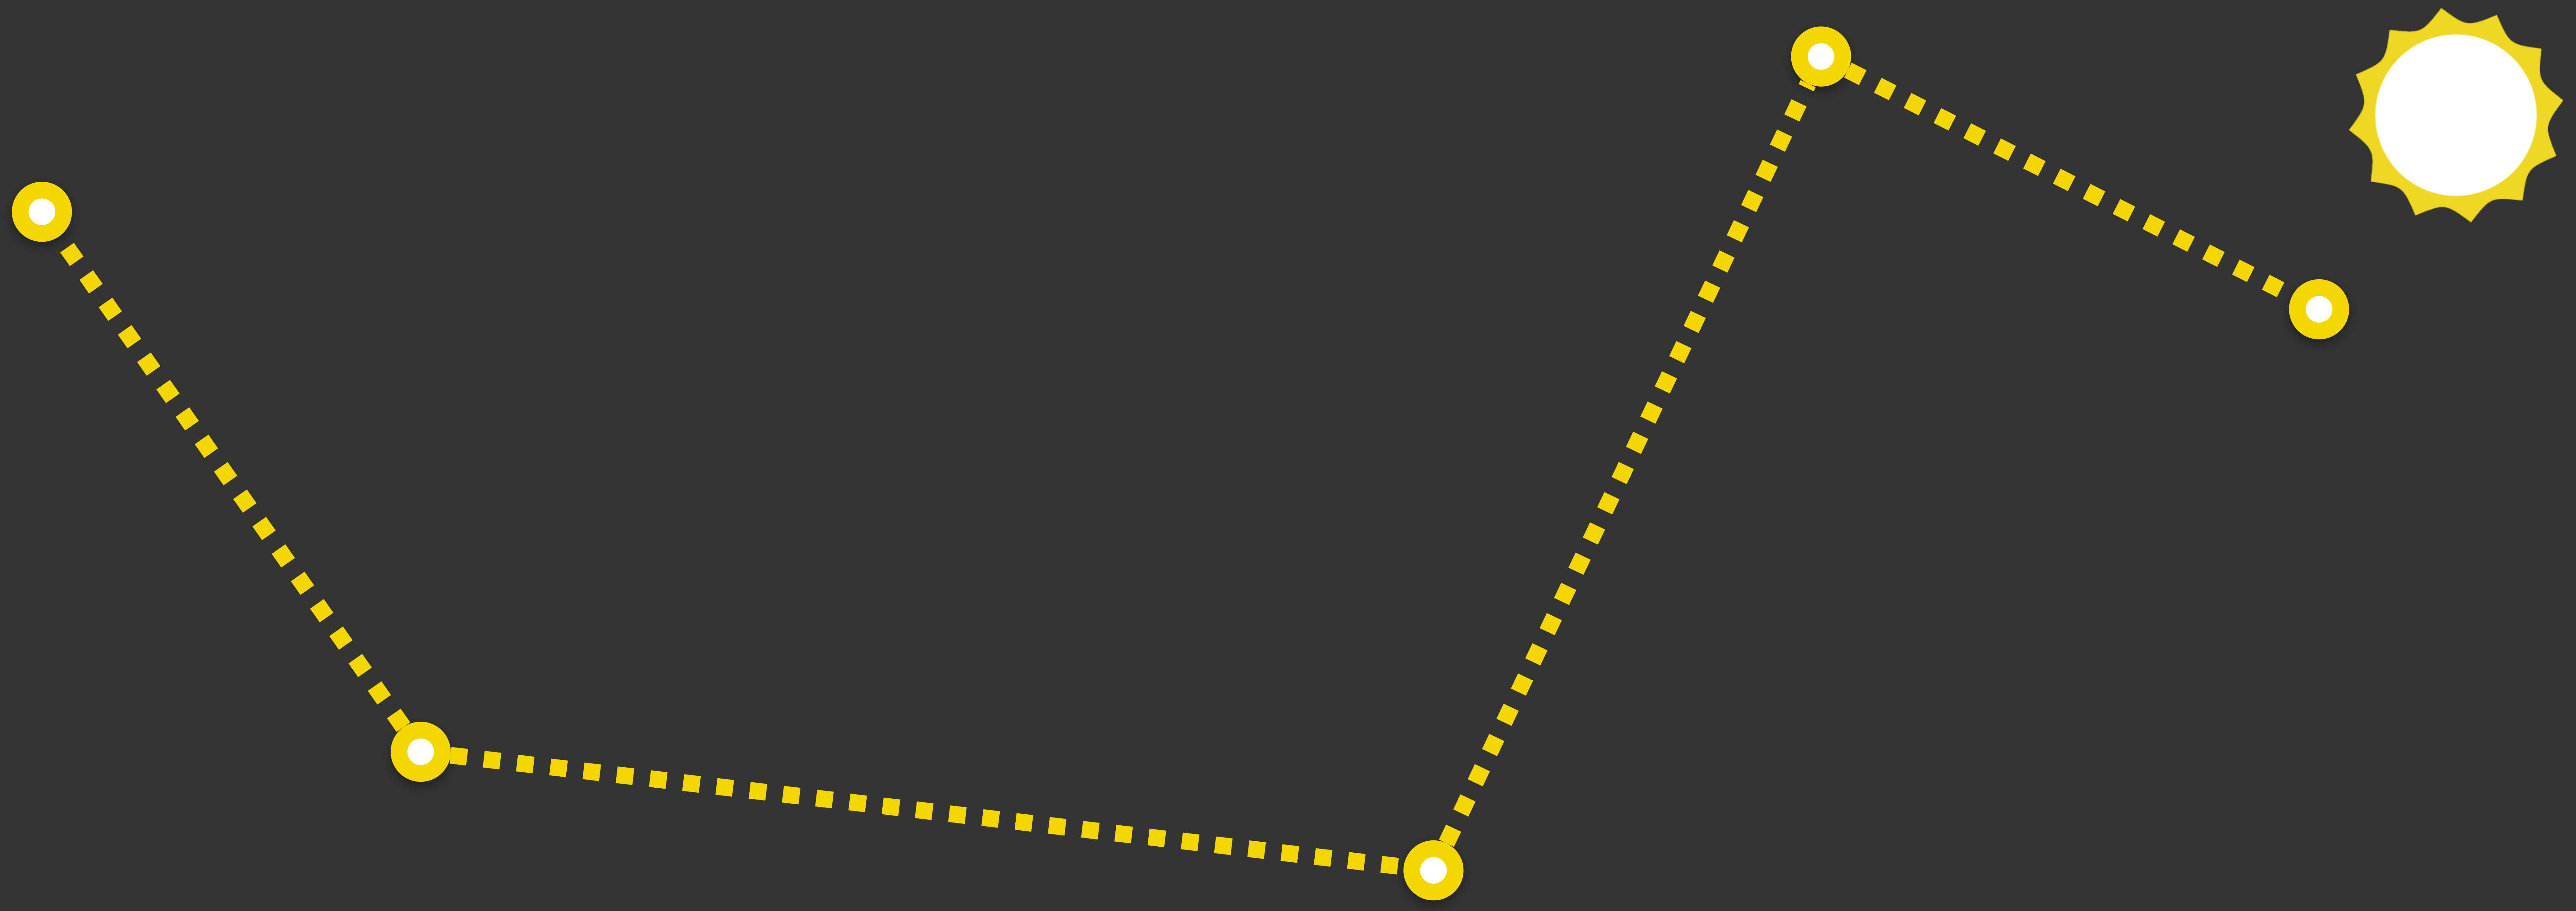

Photon Estimators

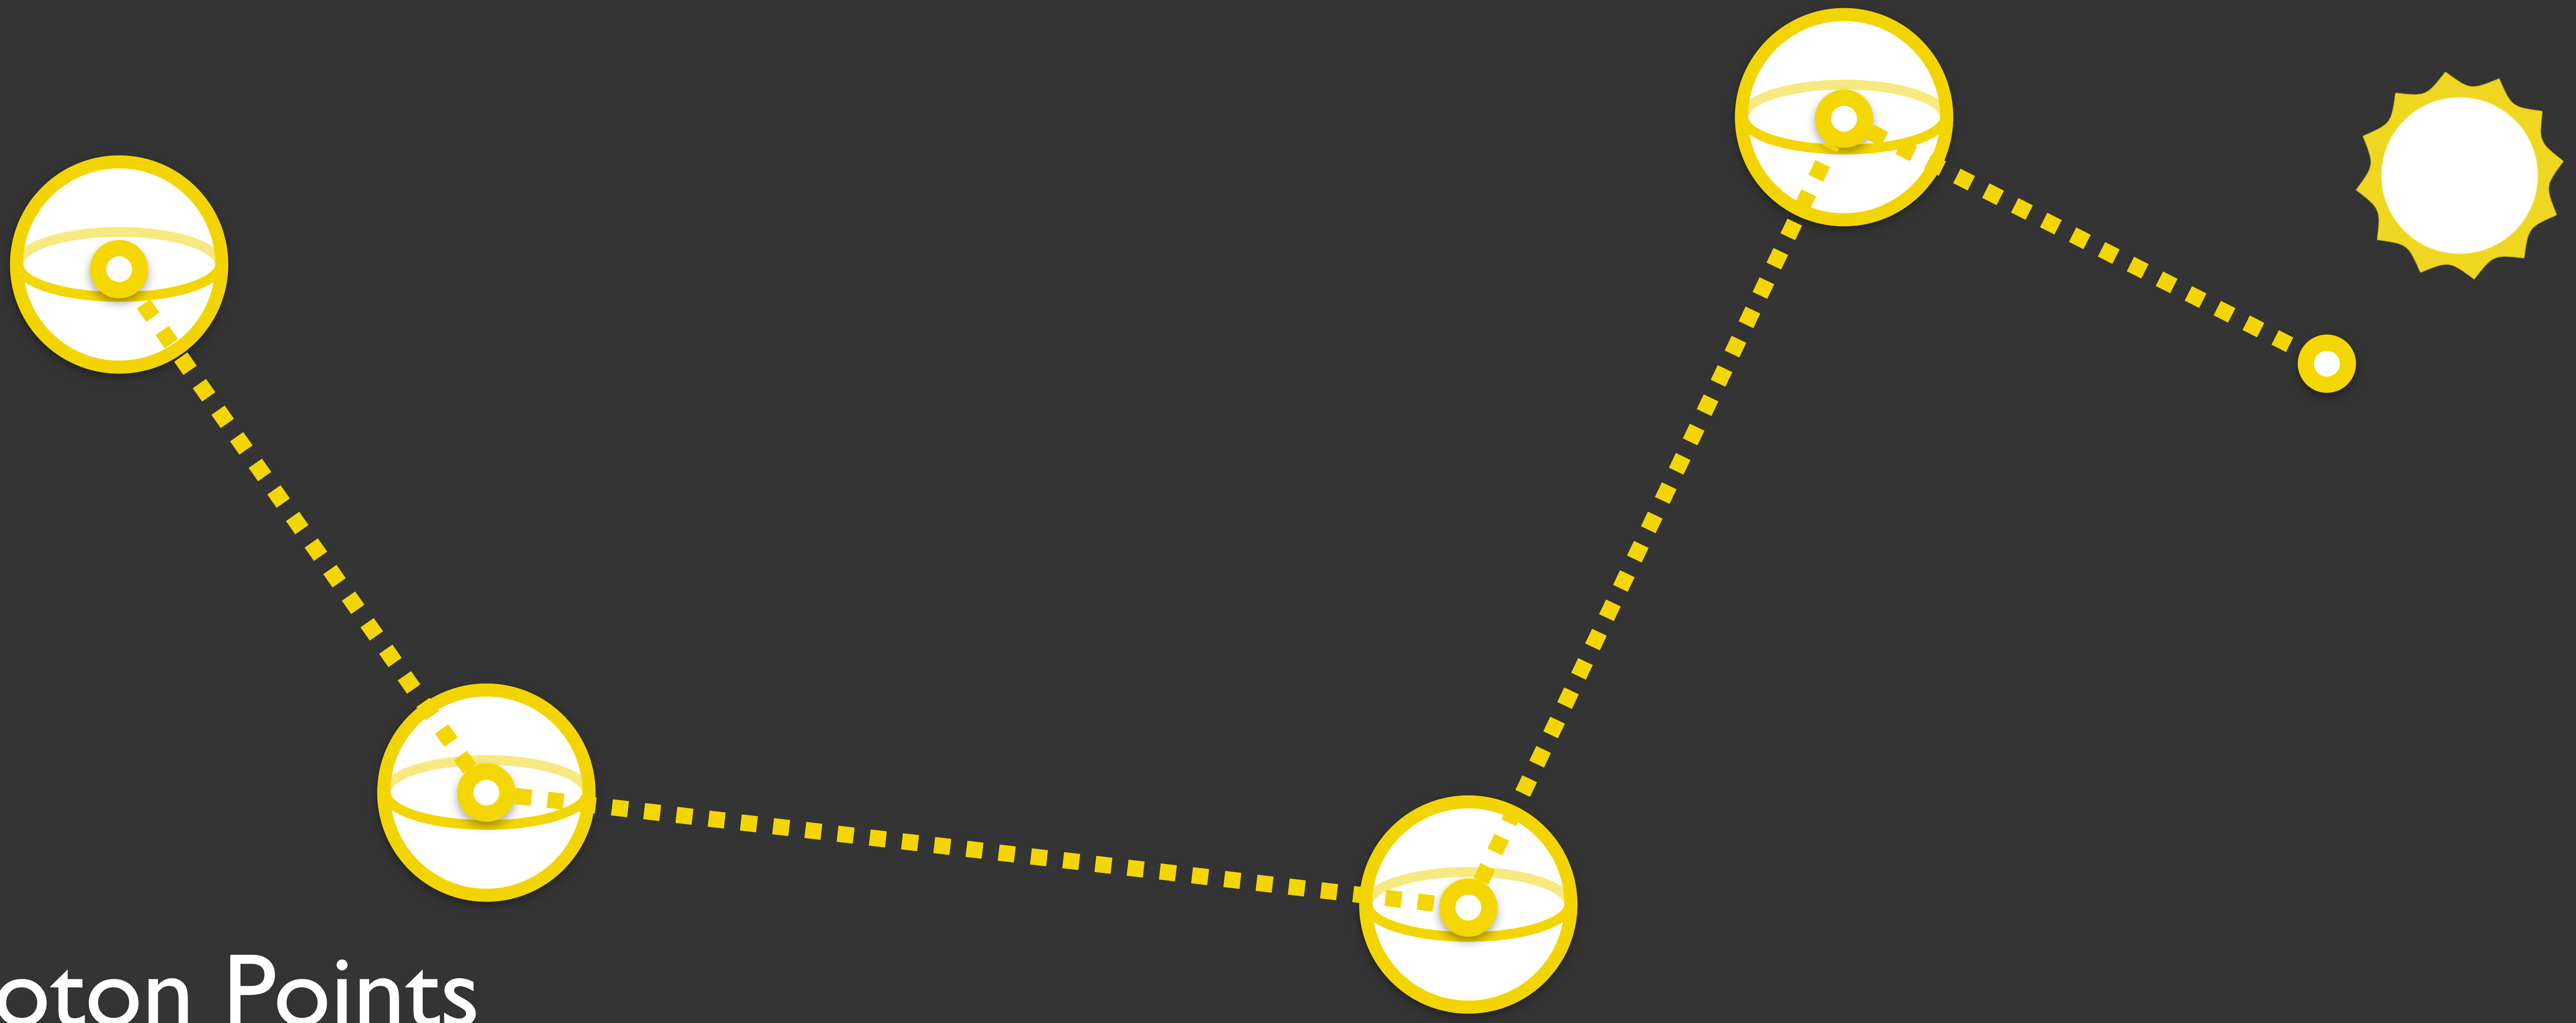

Photon Points

[Jensen & Christensen ’98]

Photon Estimators

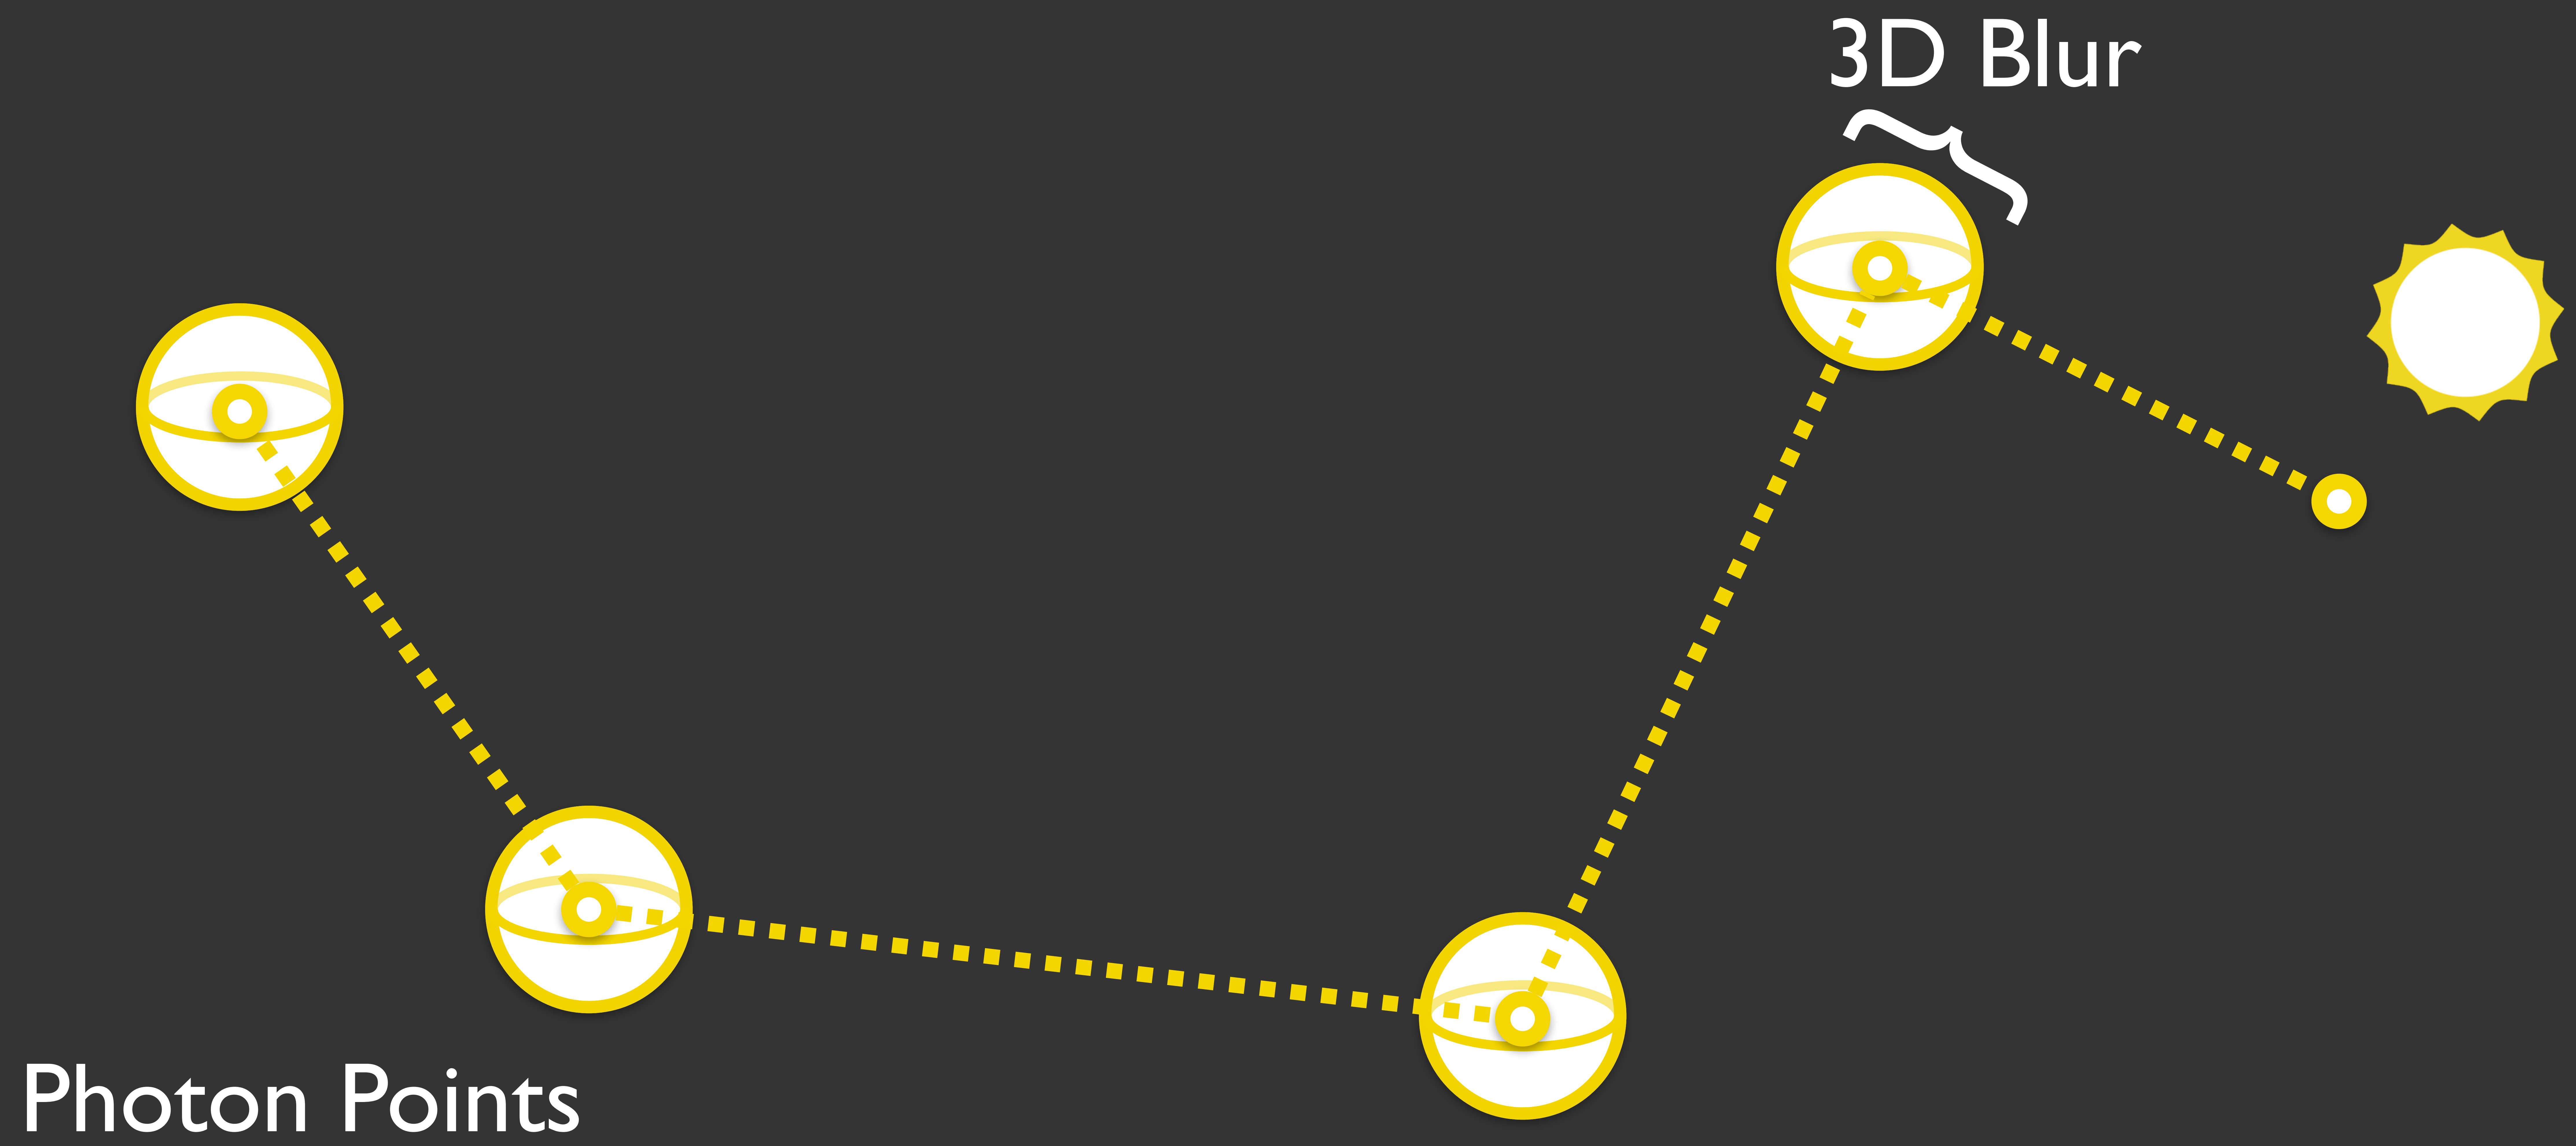

[Jensen & Christensen ’98]

Photon Estimators

Photon Beams

[Jarosz et al. '11]

Photon Estimators

"Long" Beams

[Jarosz et al. '11]

Photon Estimators

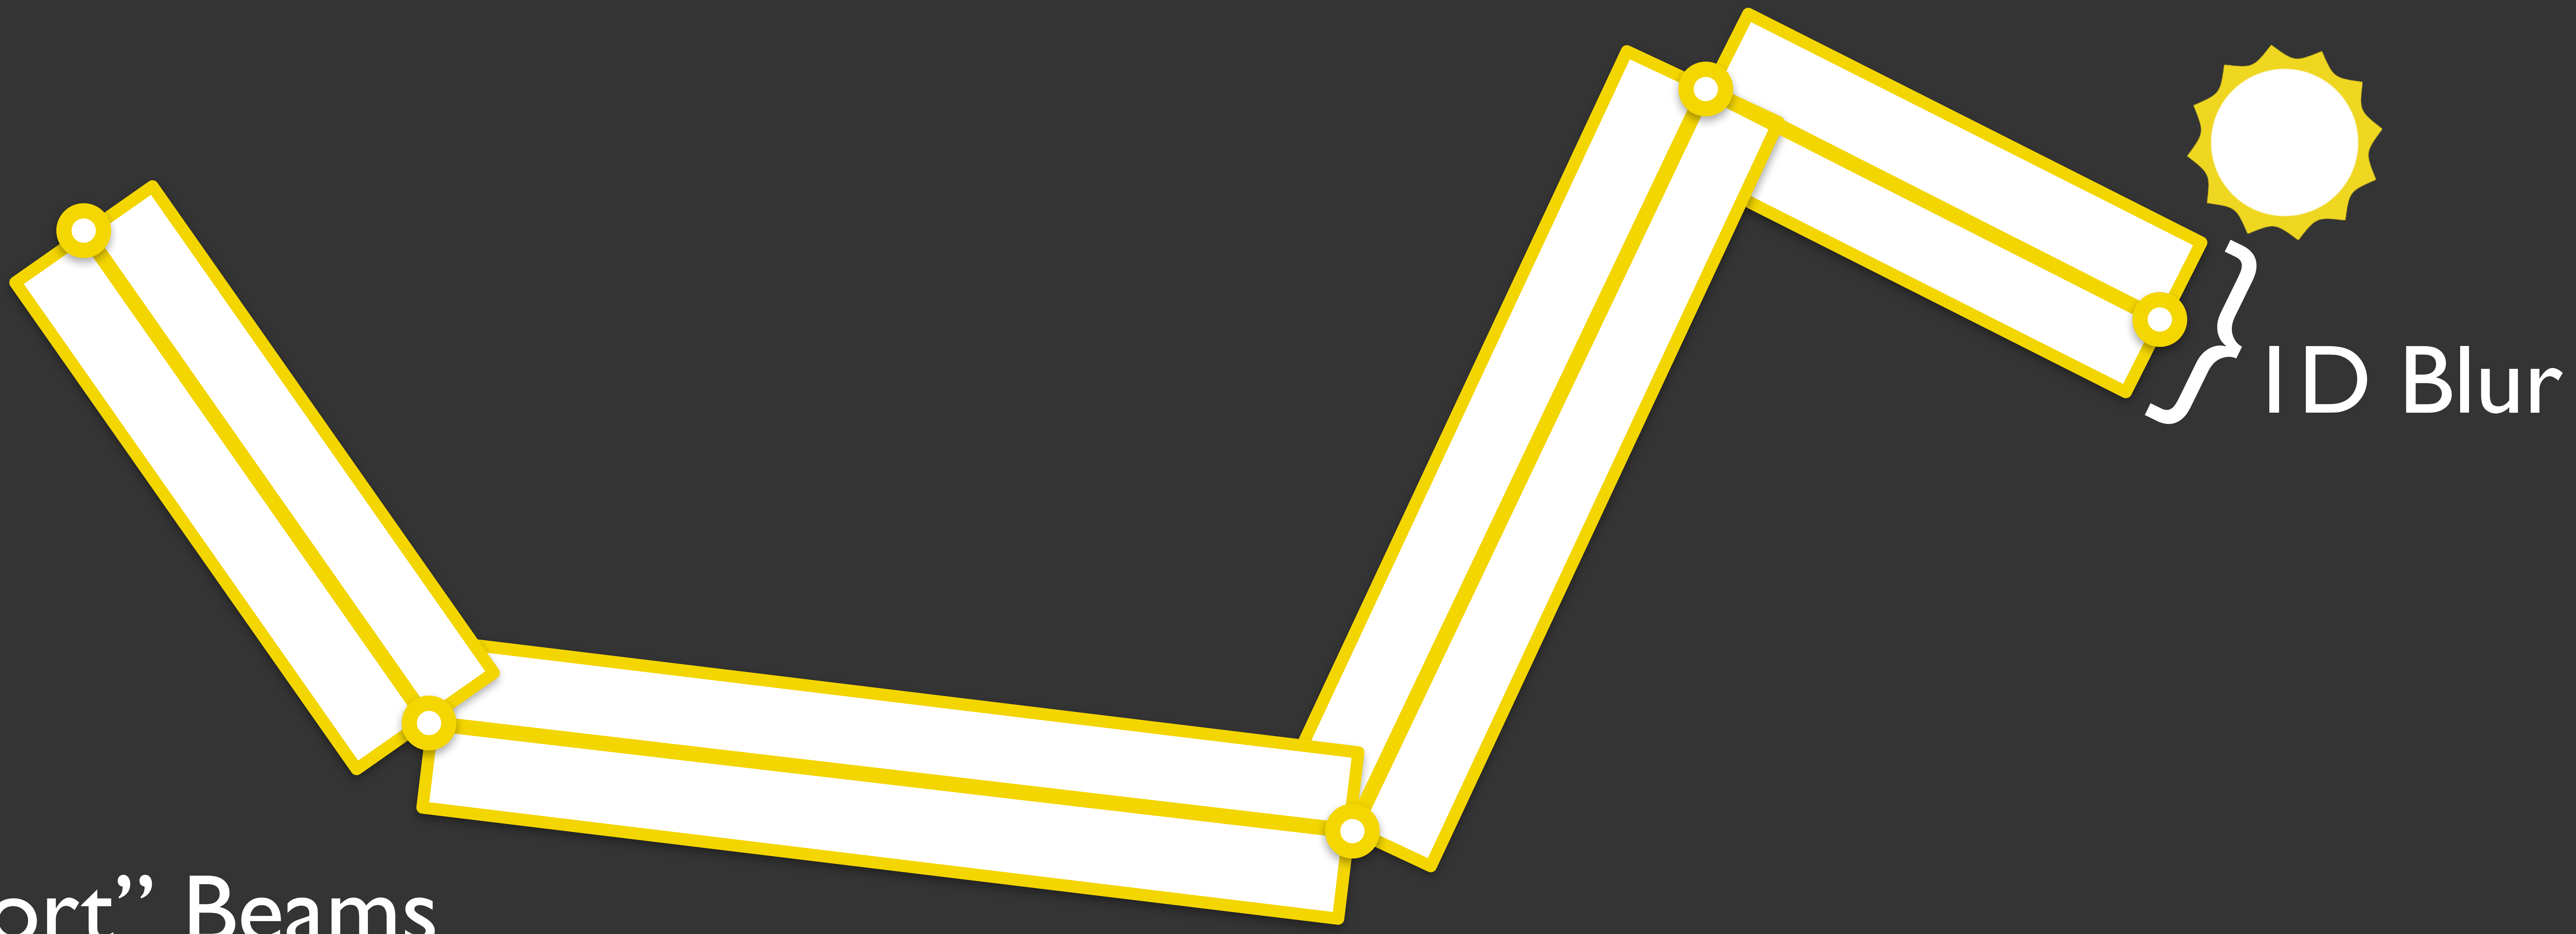

"Short" Beams

[Jarosz et al. '11]

Summary

- “Marching” is a mechanism to obtain new photons

Photon Estimators

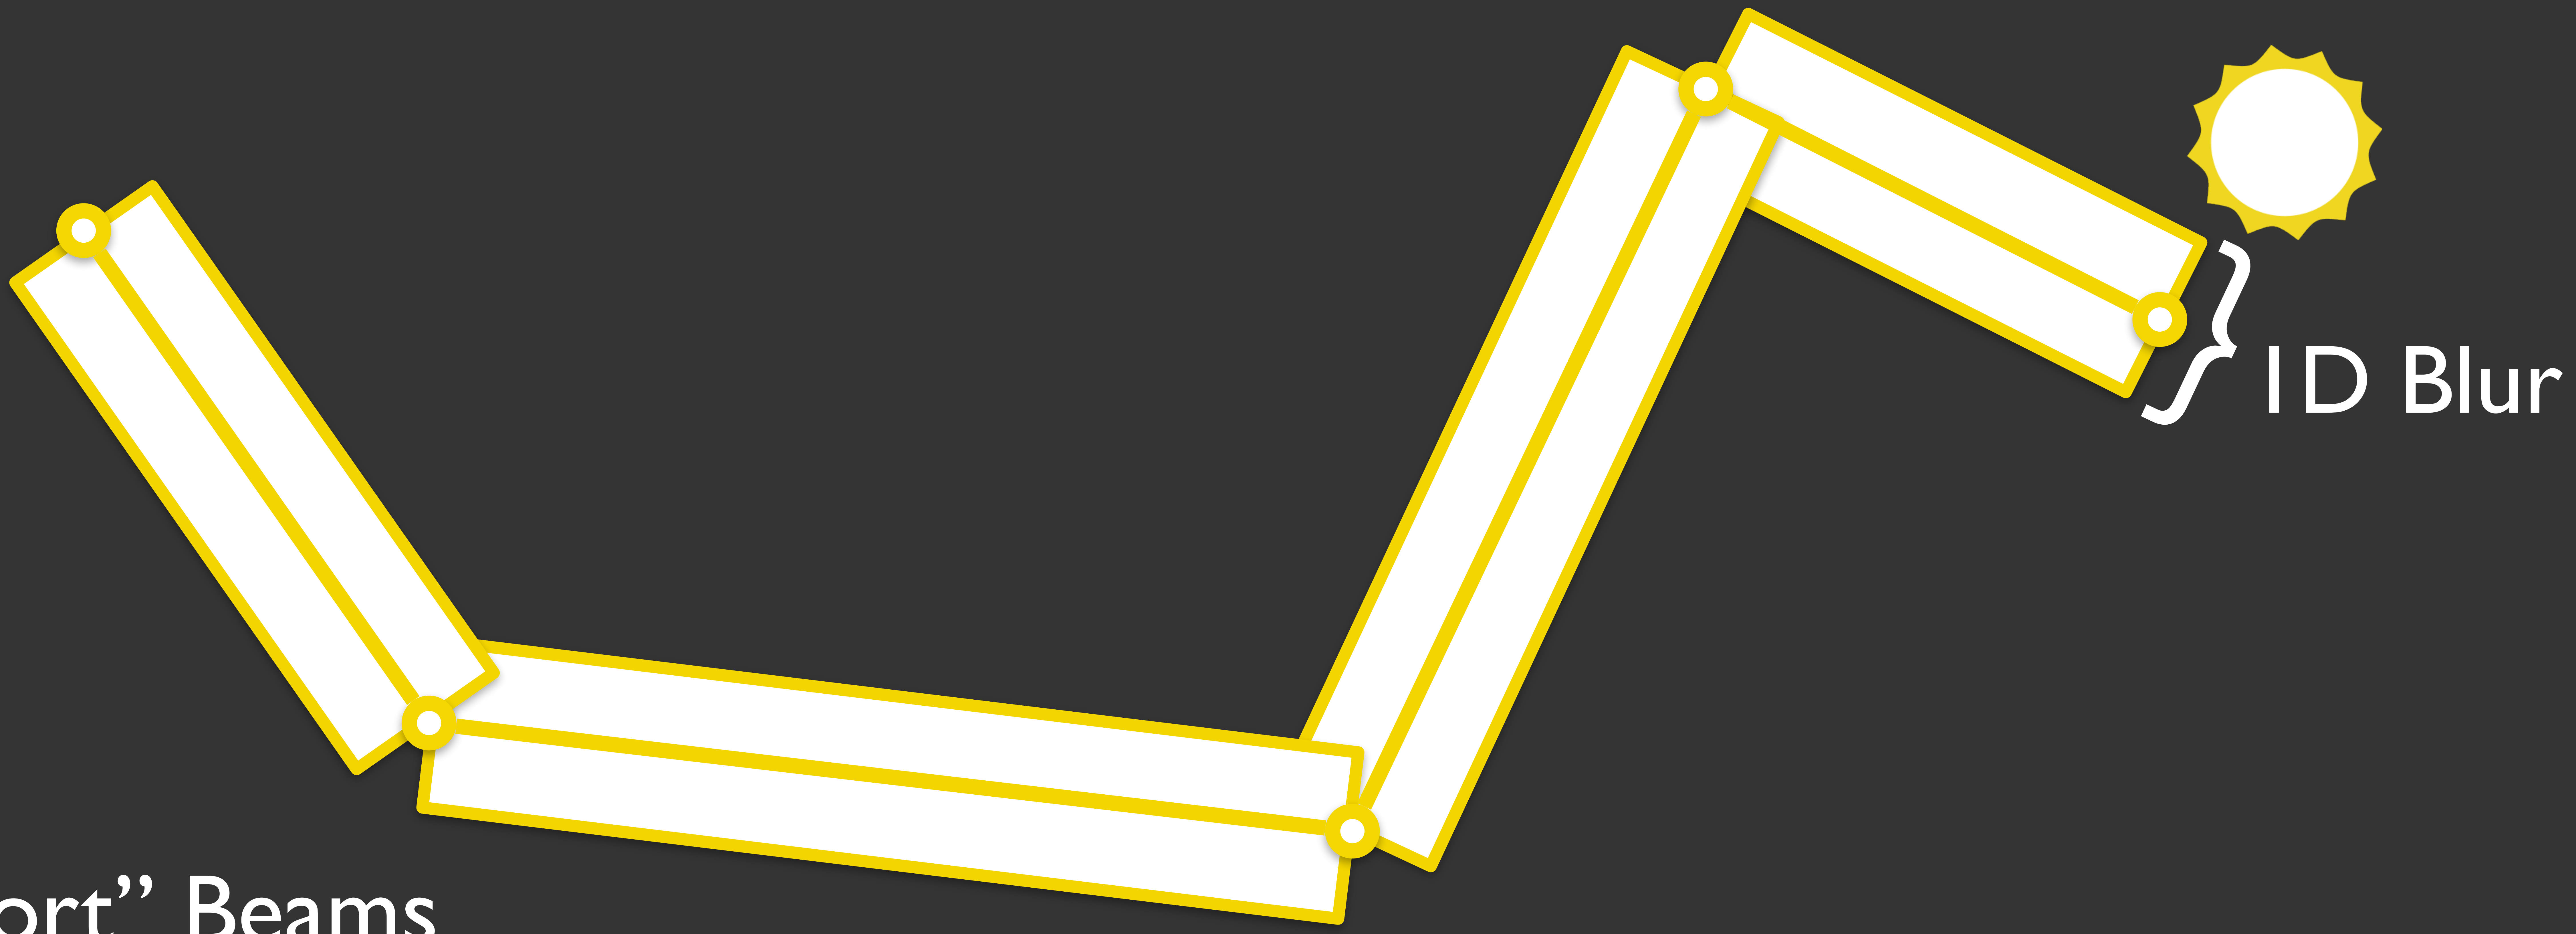

"Short" Beams

[Křivánek et al. 2014]

Photon Estimators

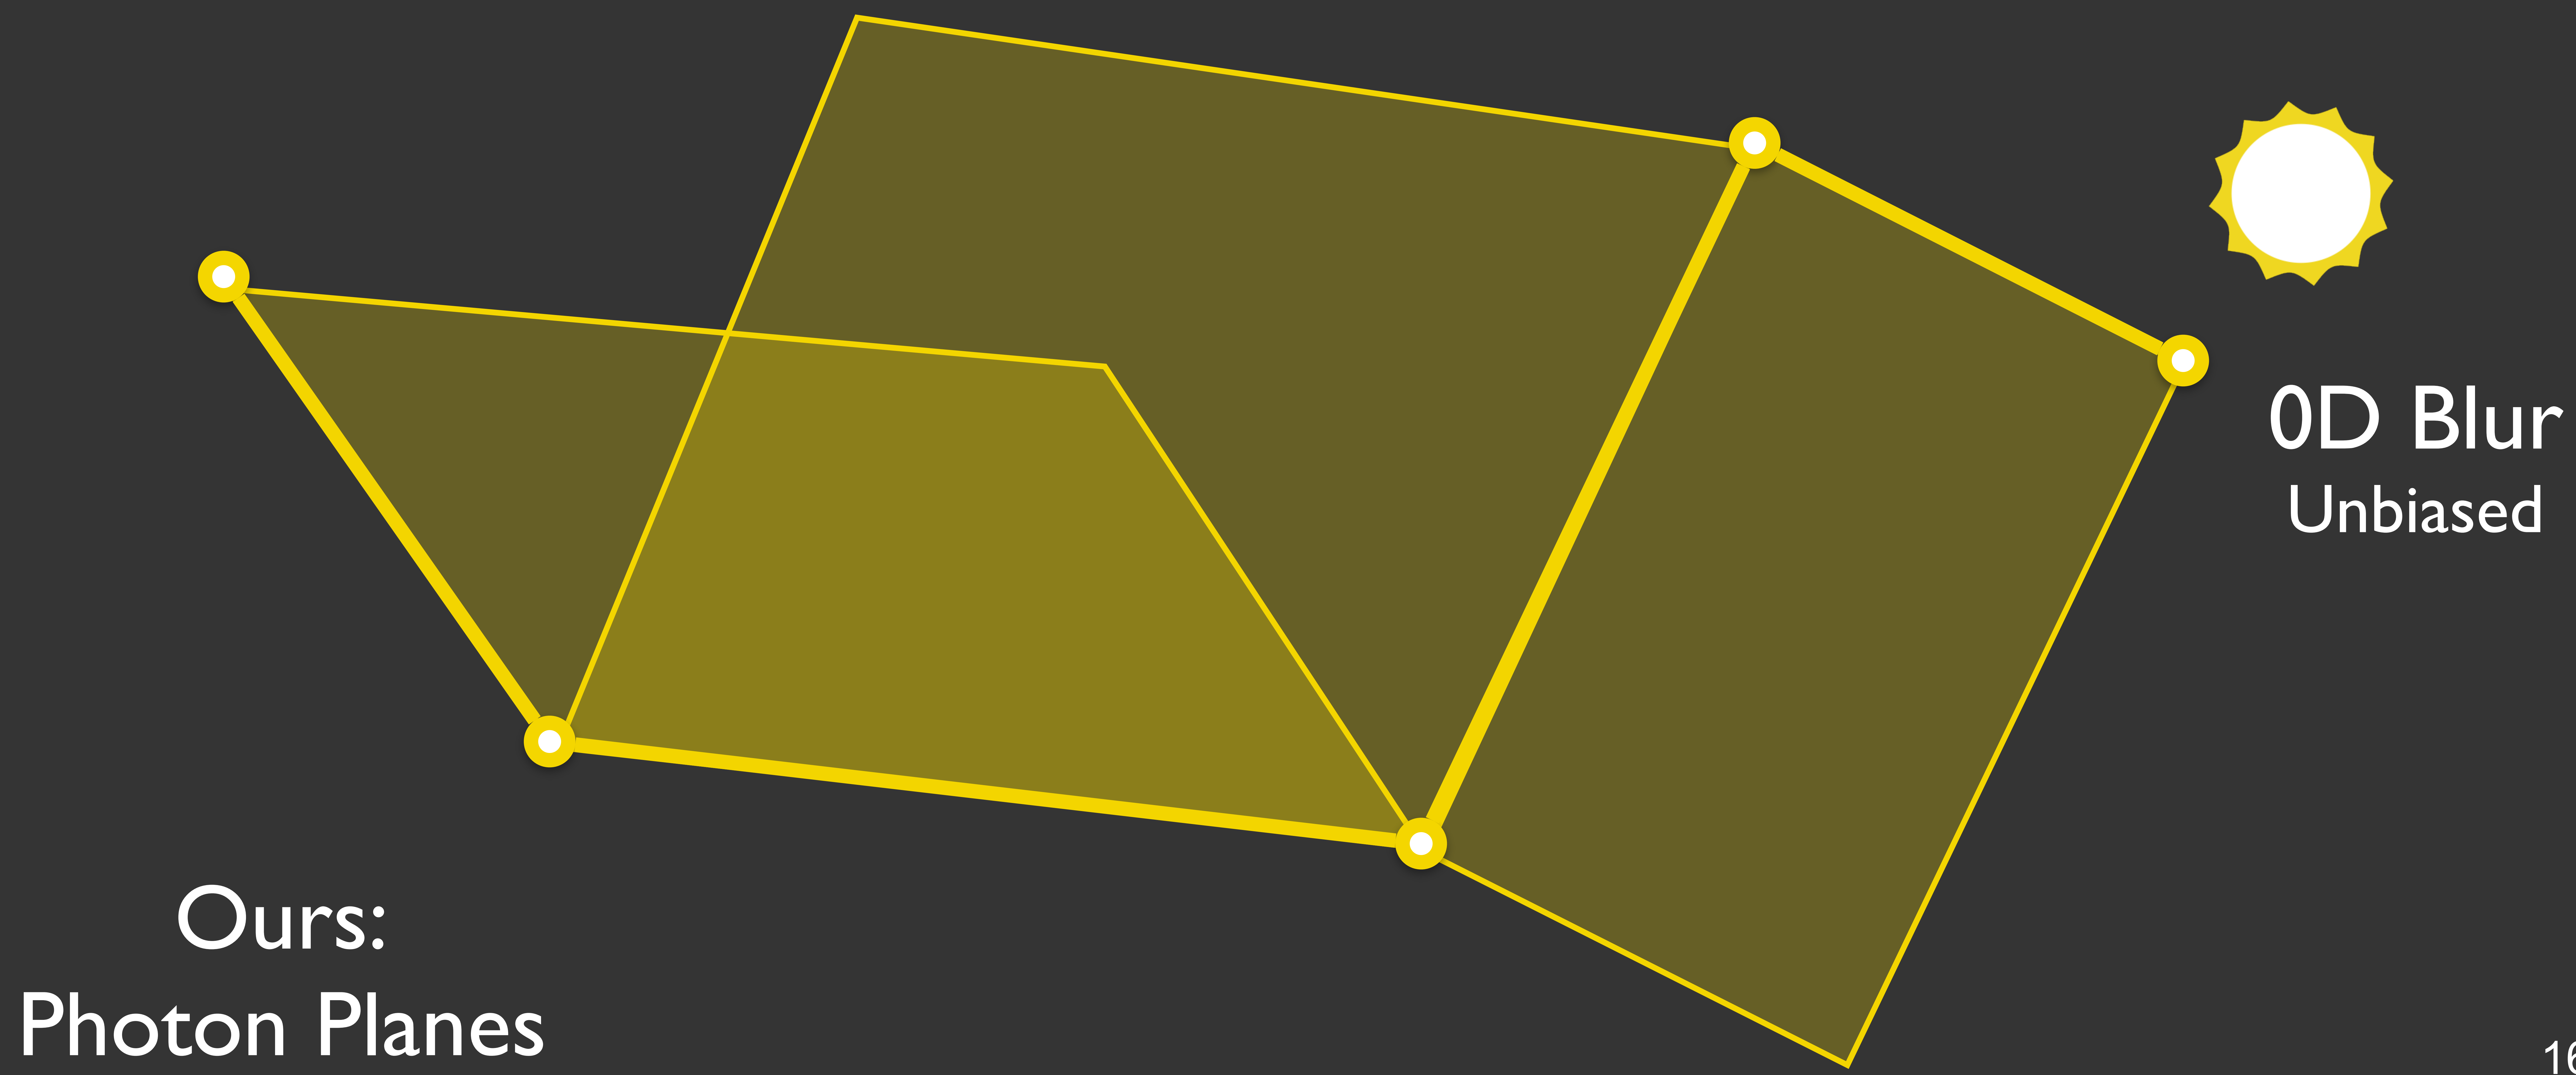

Photon Estimators

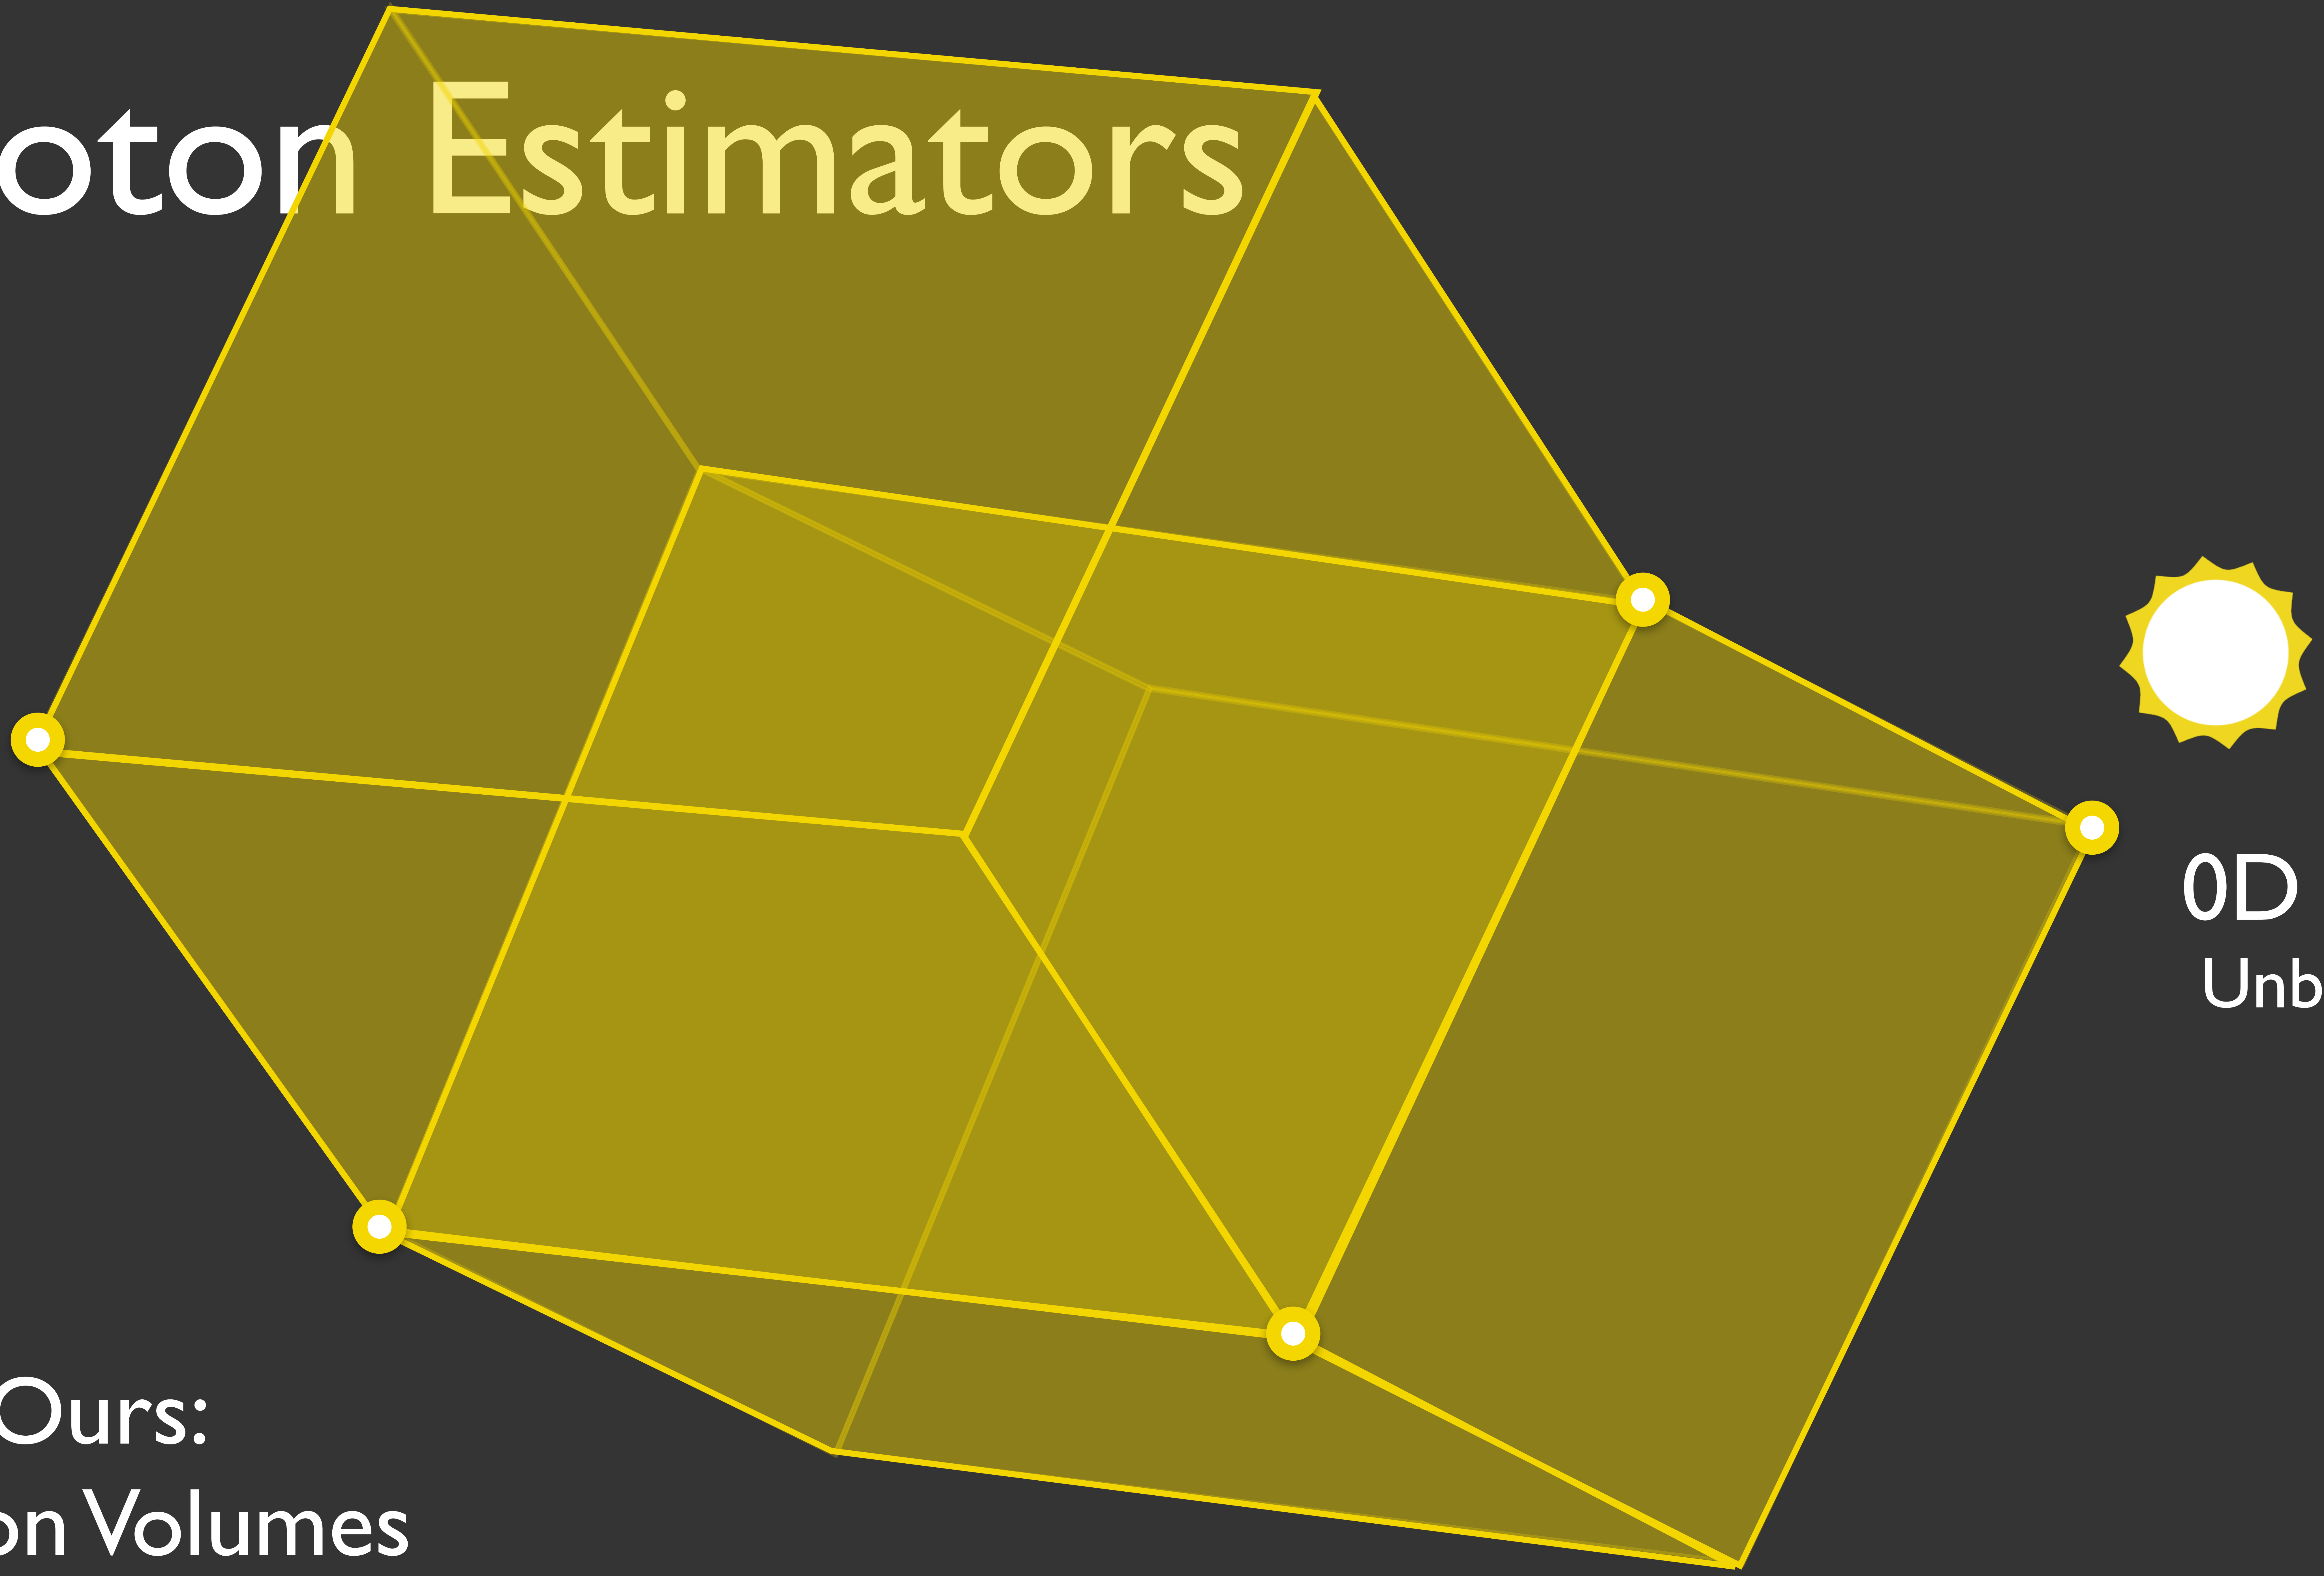

0D Blur
Unbiased

Ours:
Photon Volumes

Summary

- “Marching” is a mechanism to obtain new photons

Summary

- “Marching” is a mechanism to obtain new photons
- Observation:

  “Marching” replaces transmittance estimators

Transmittance Estimators

Transmittance Estimators

- Originate in *neutron transport*

Transmittance Estimators

- Originate in *neutron transport*
- Closely linked to photons

[Křivánek et al. 2014]

Transmittance Estimators

Transmittance Estimators

Transmittance Estimators

Transmittance Estimators

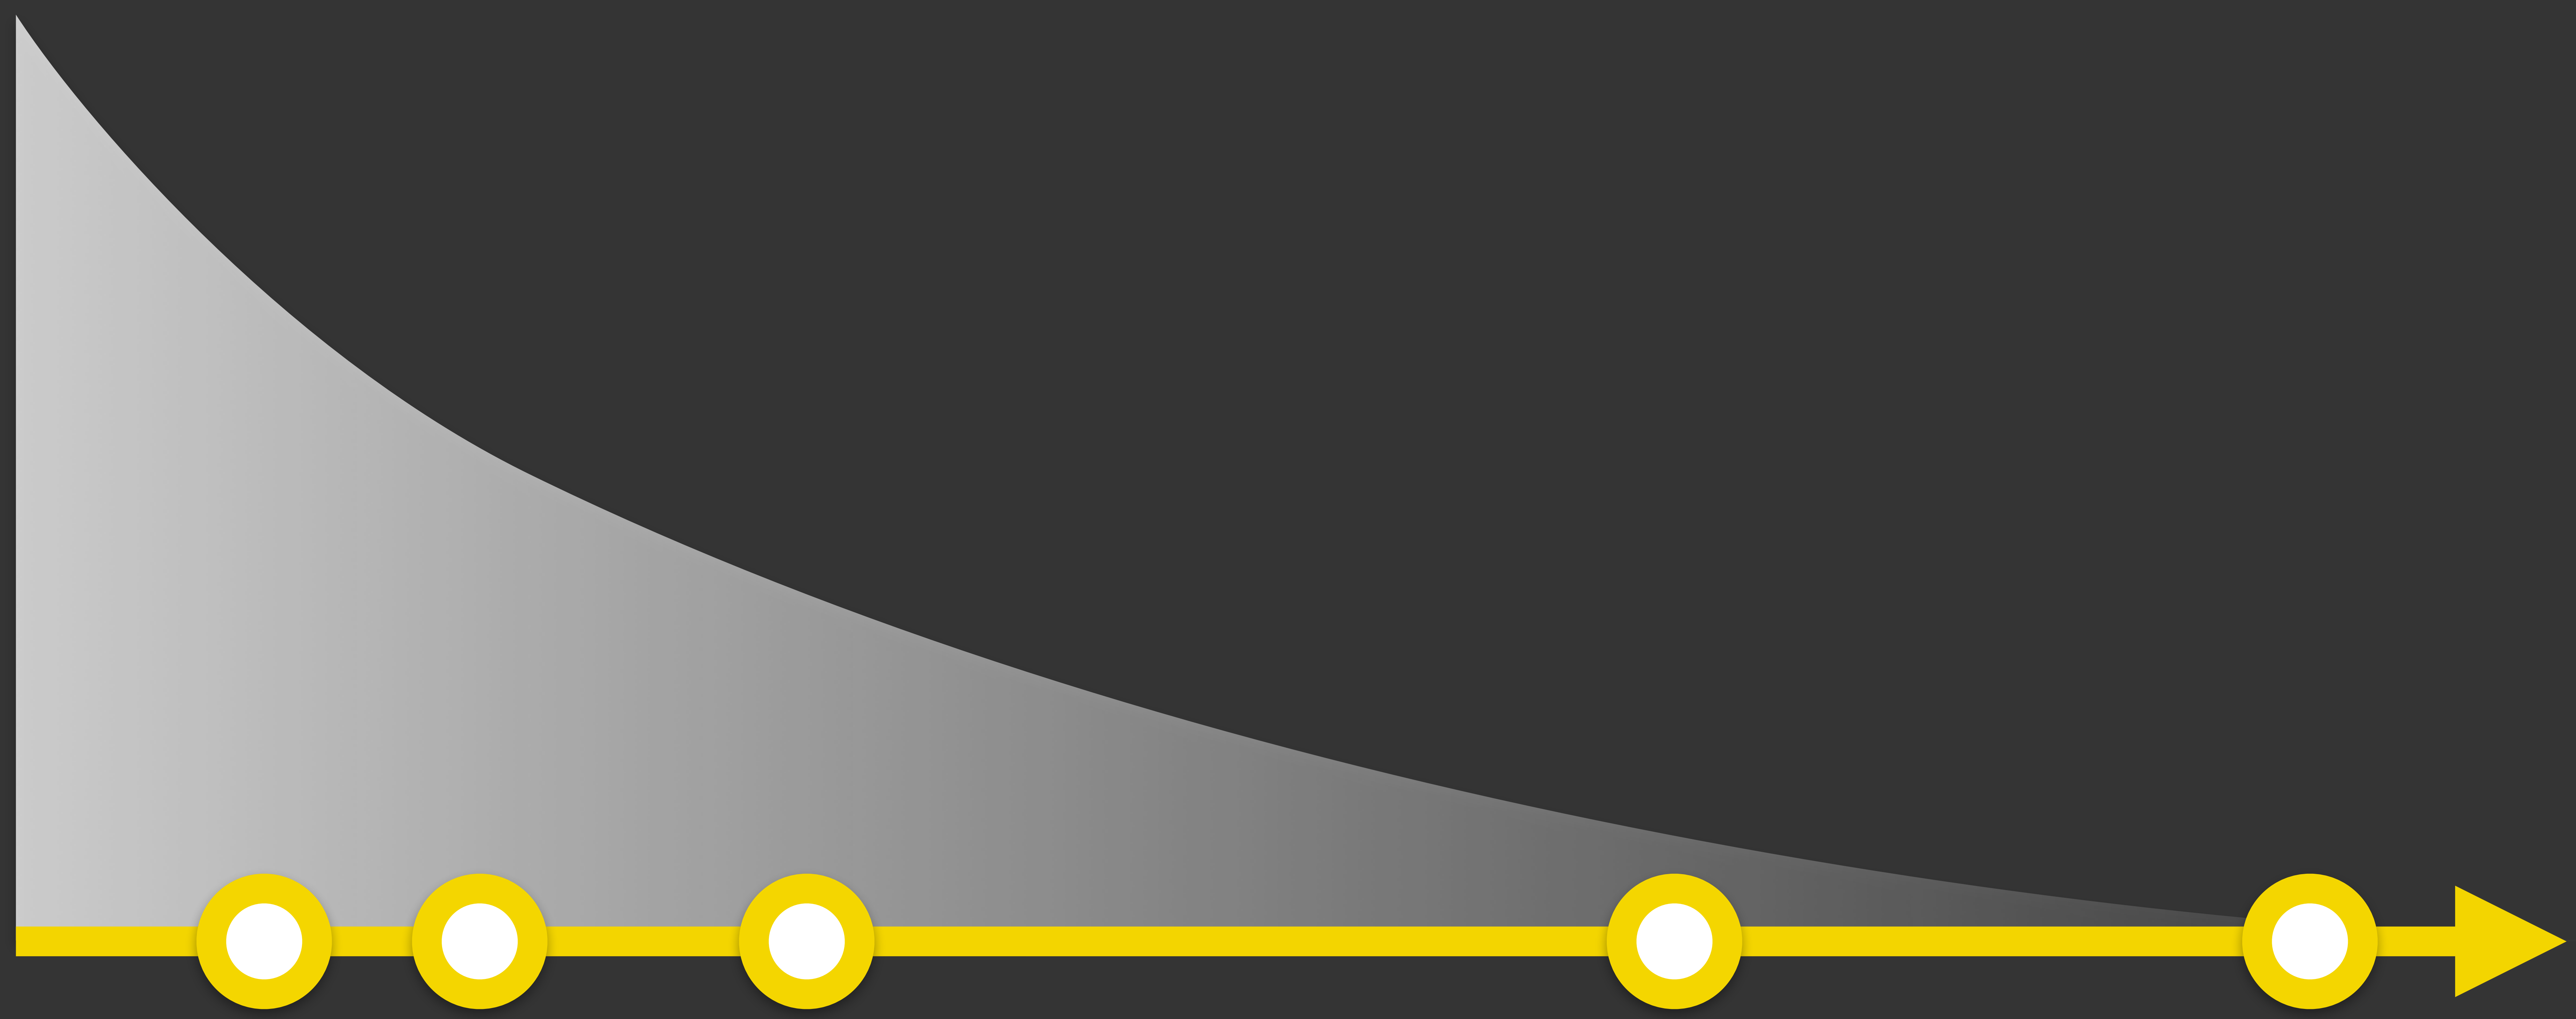

Transmittance Estimators

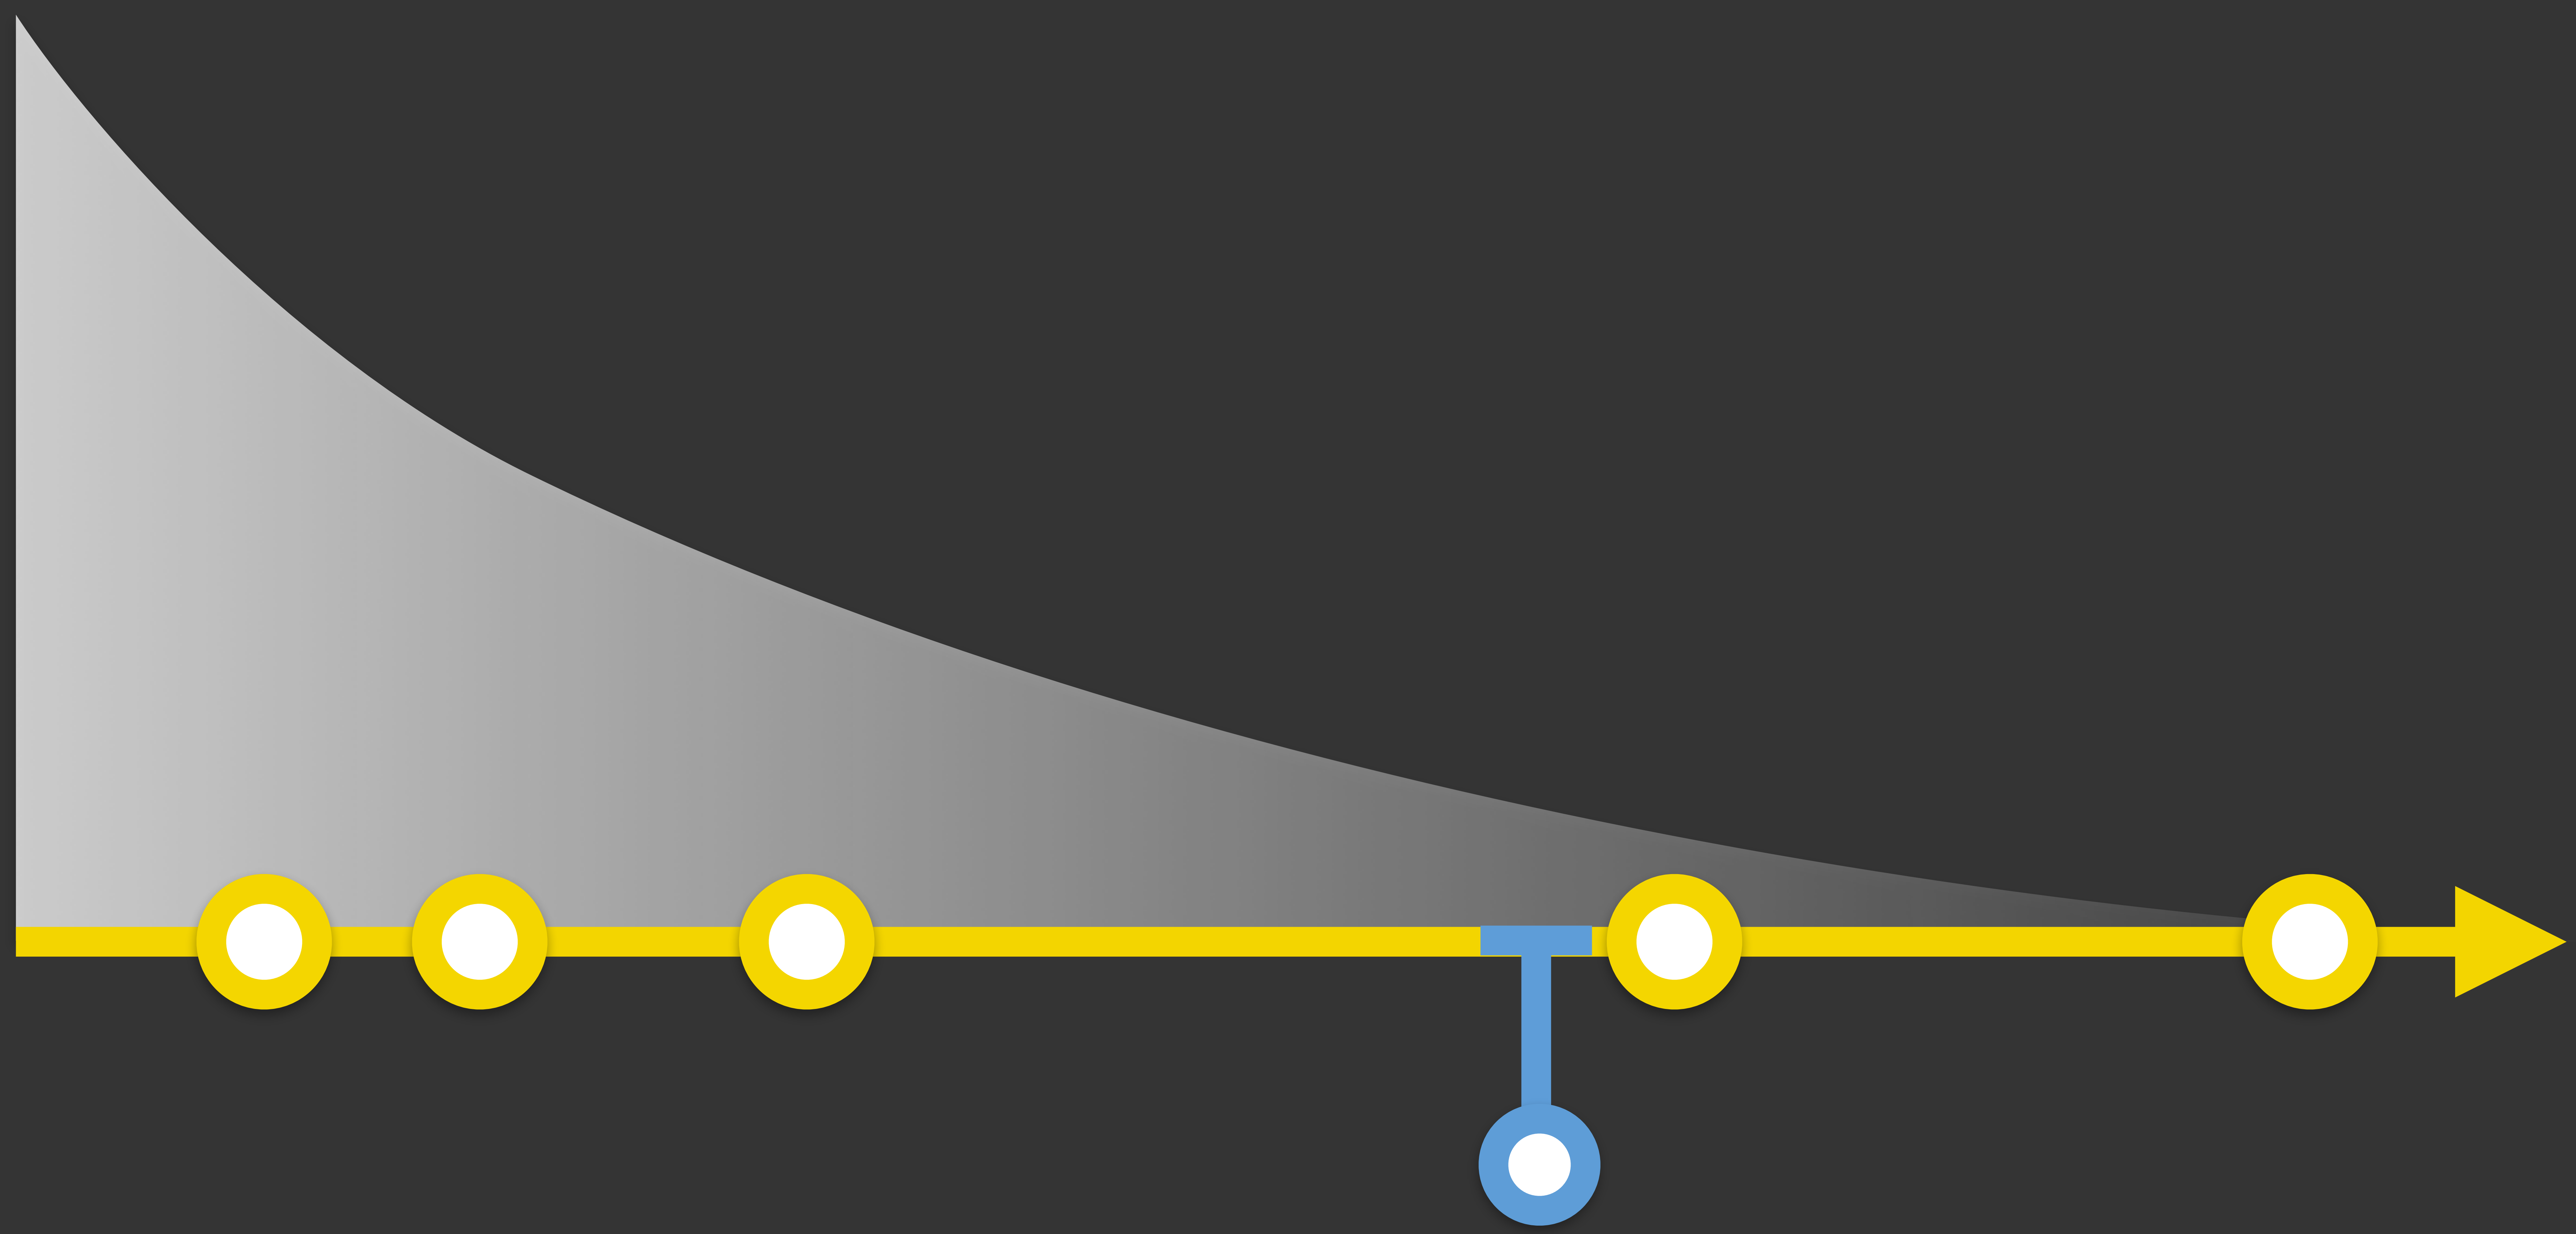

Collision Estimator

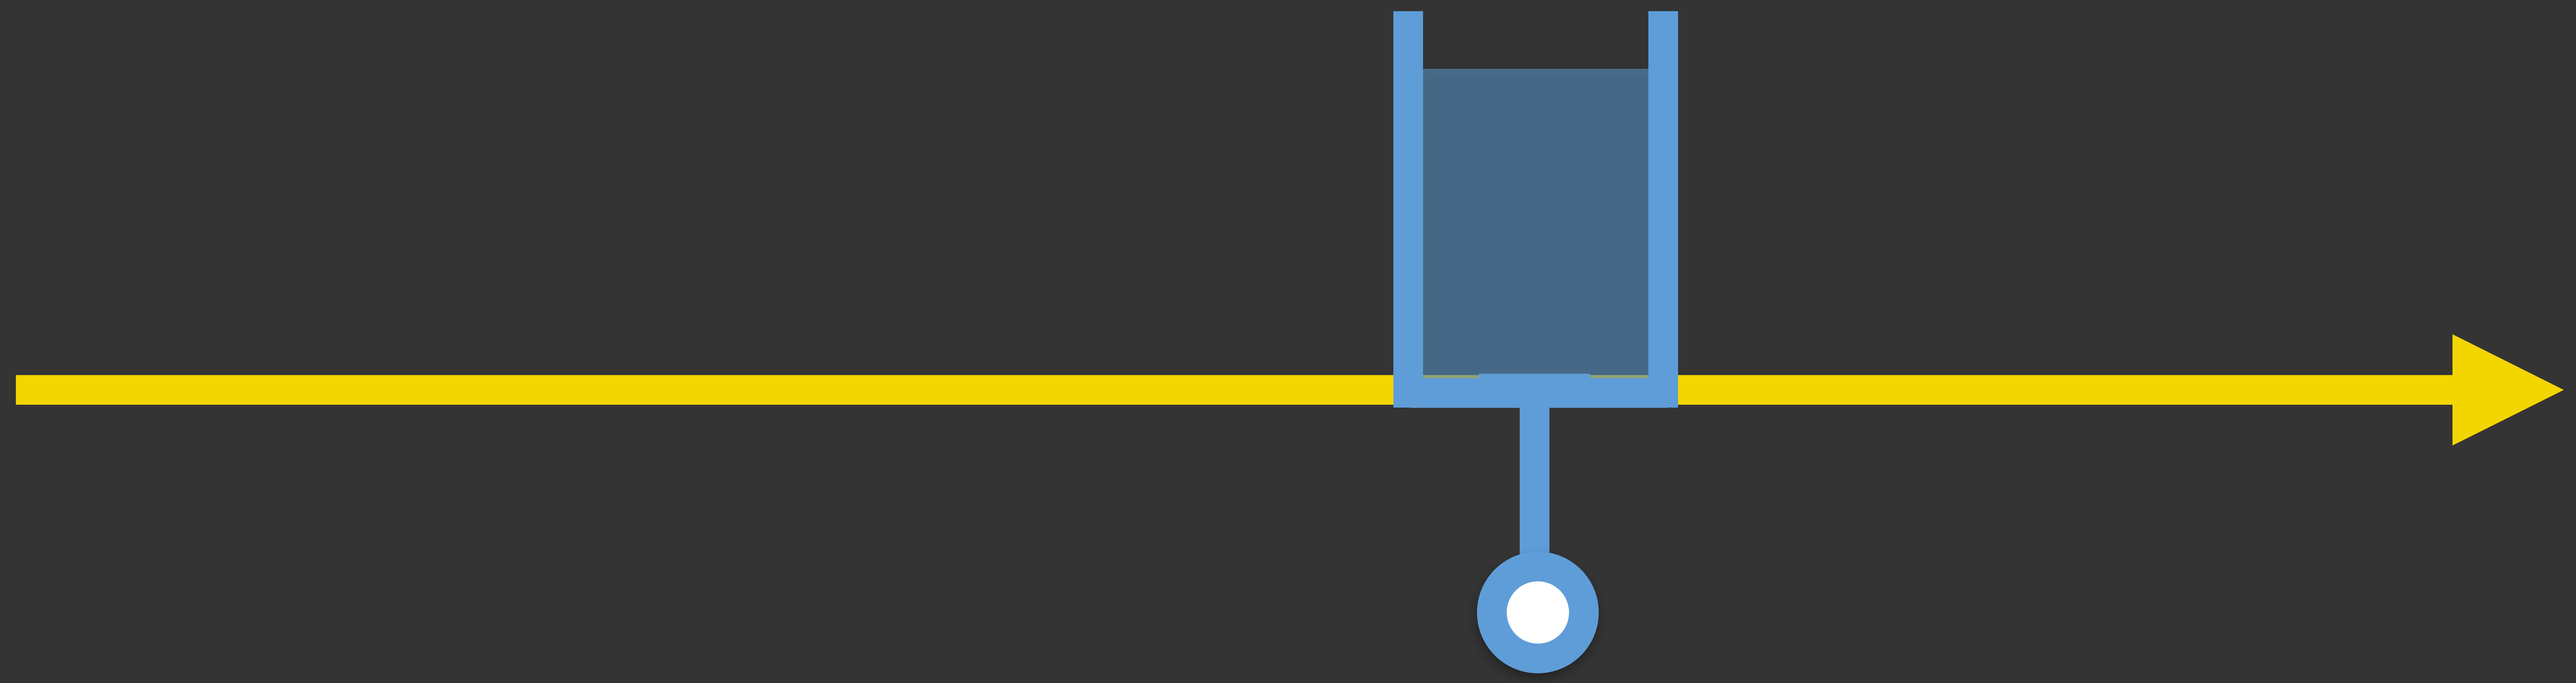

Collision Estimator

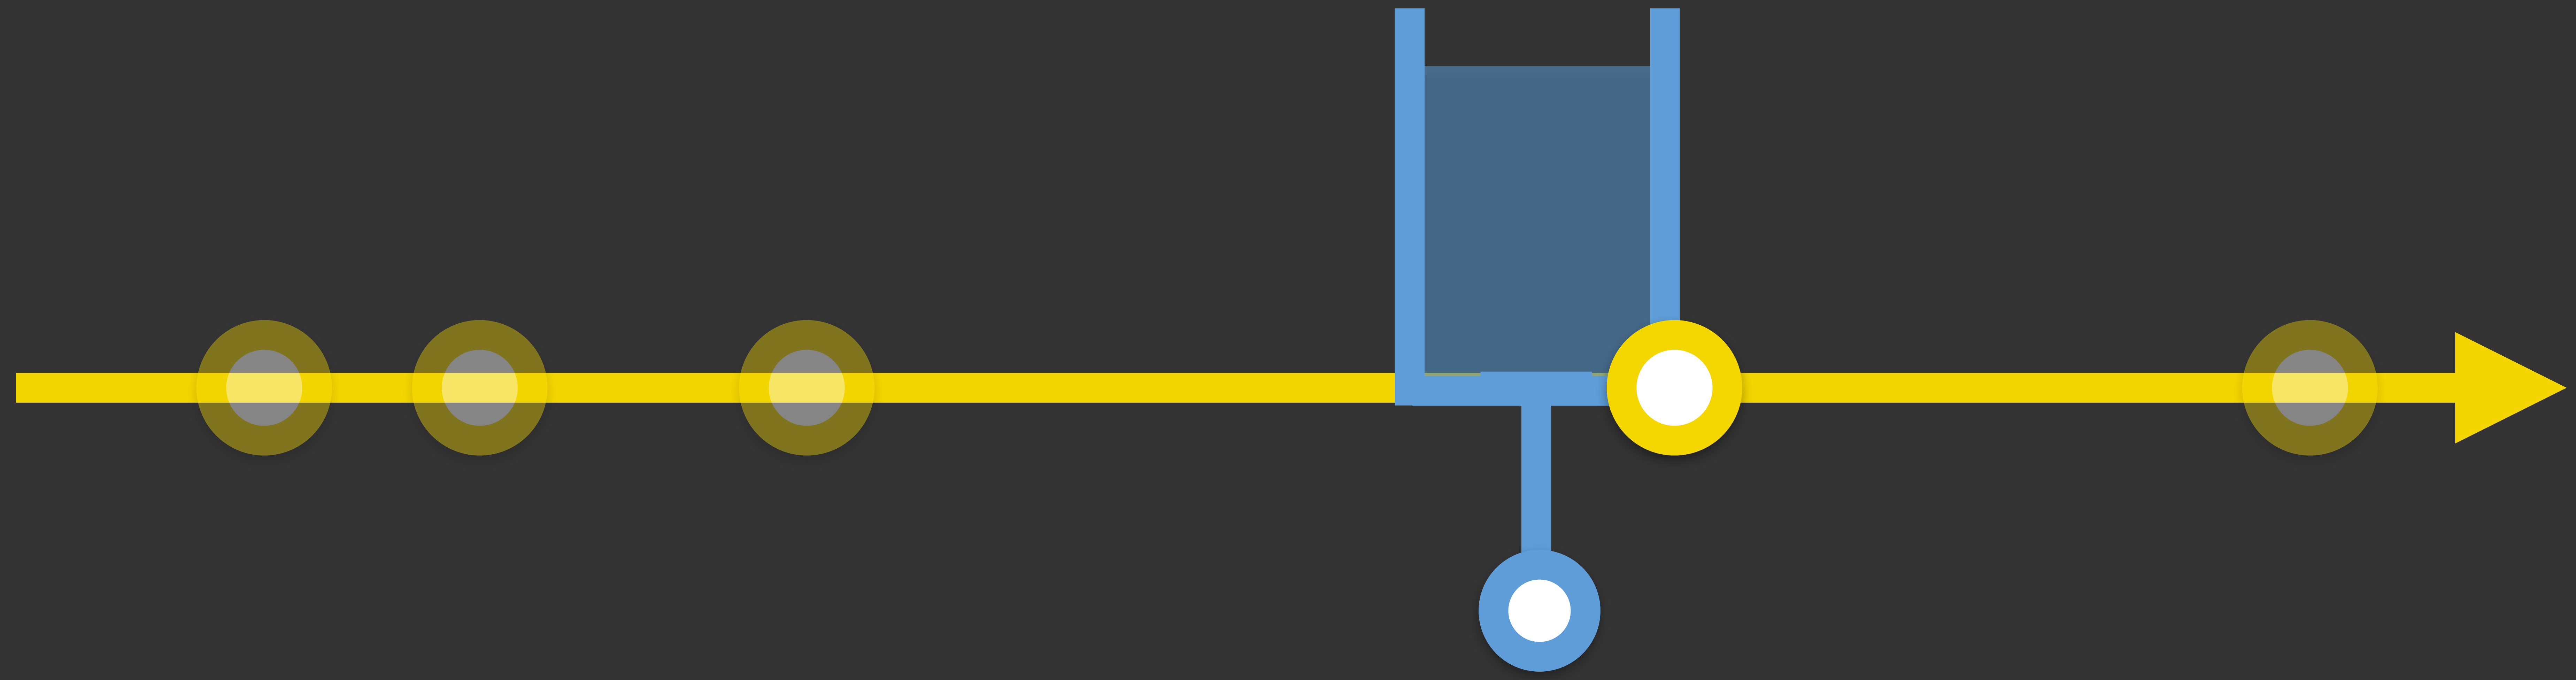

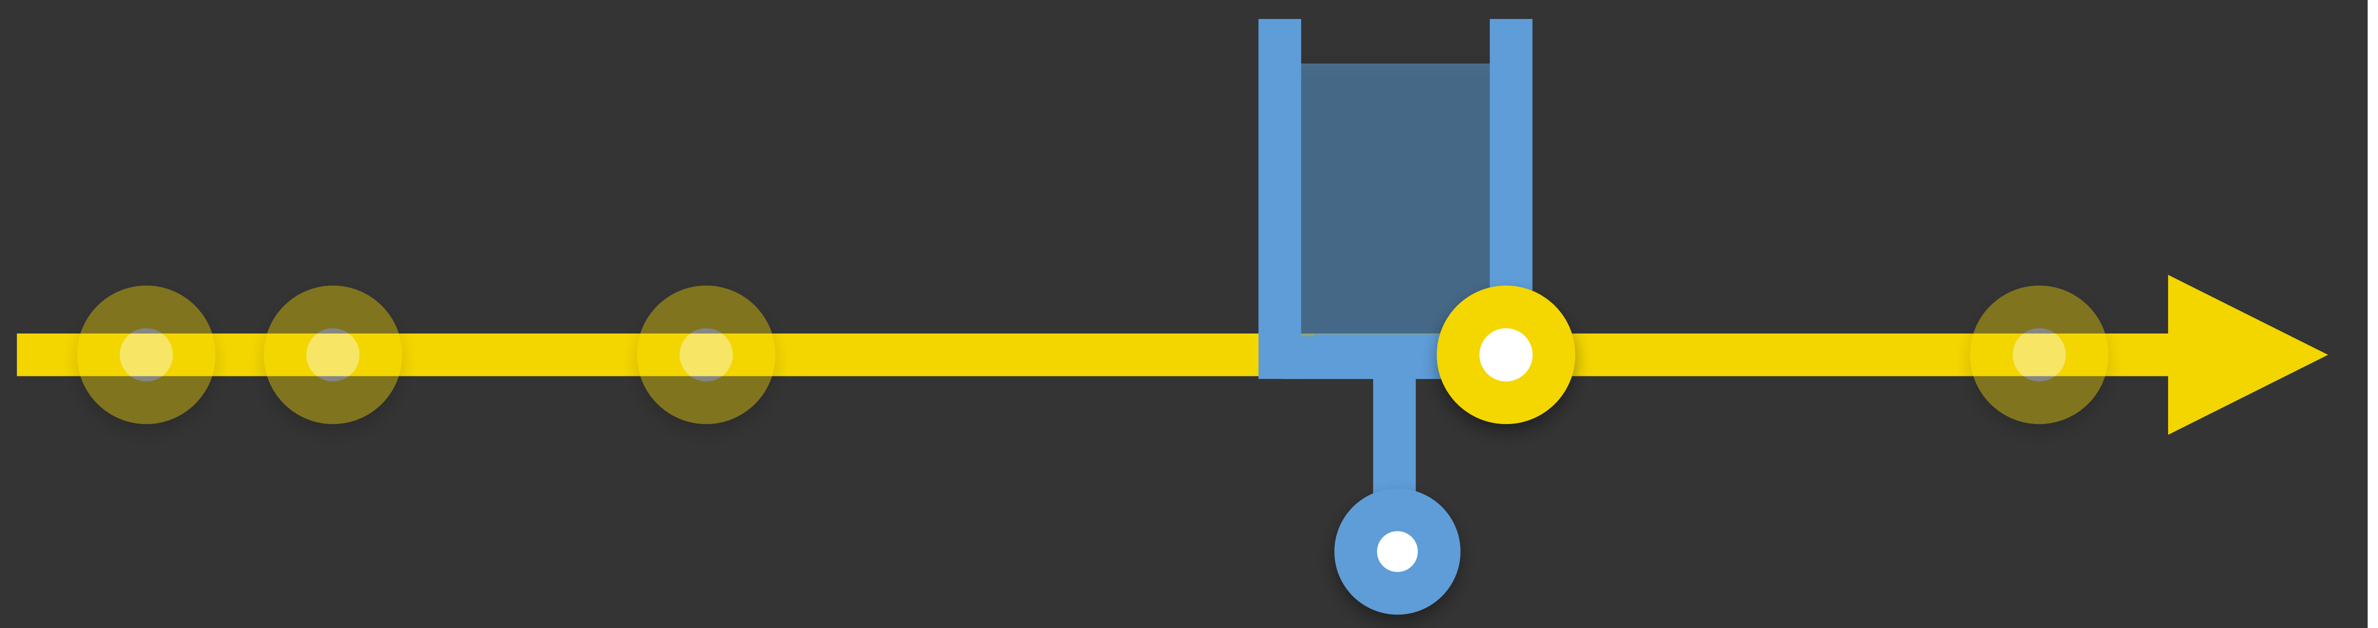

Collision Estimator

Neutron Transport

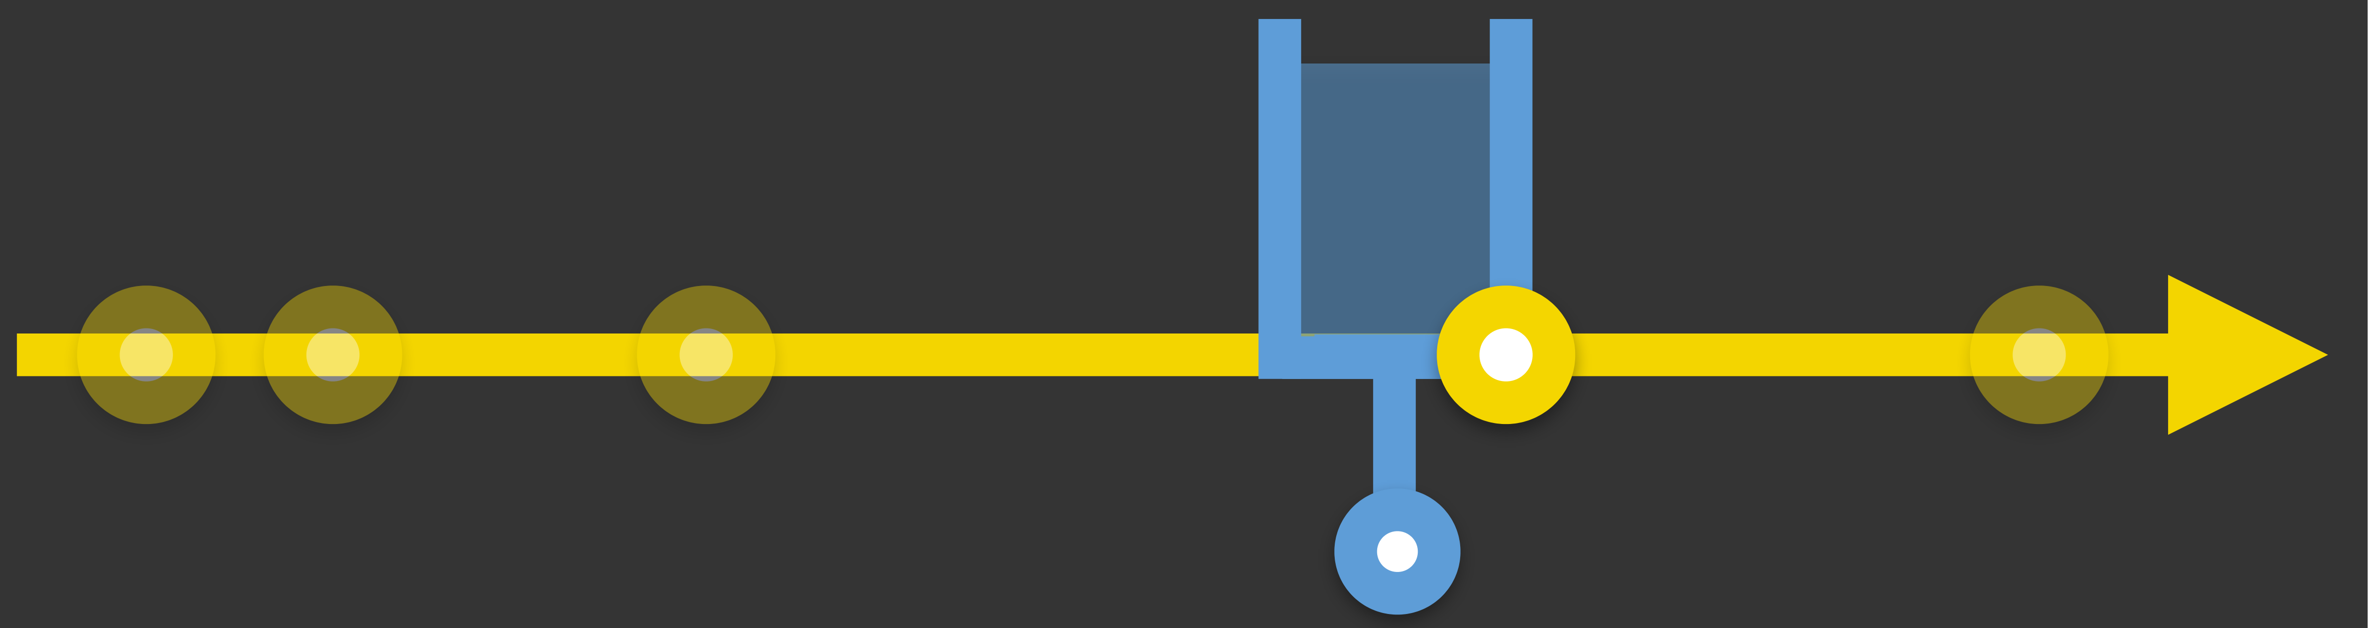

Collision Estimator

Neutron Transport

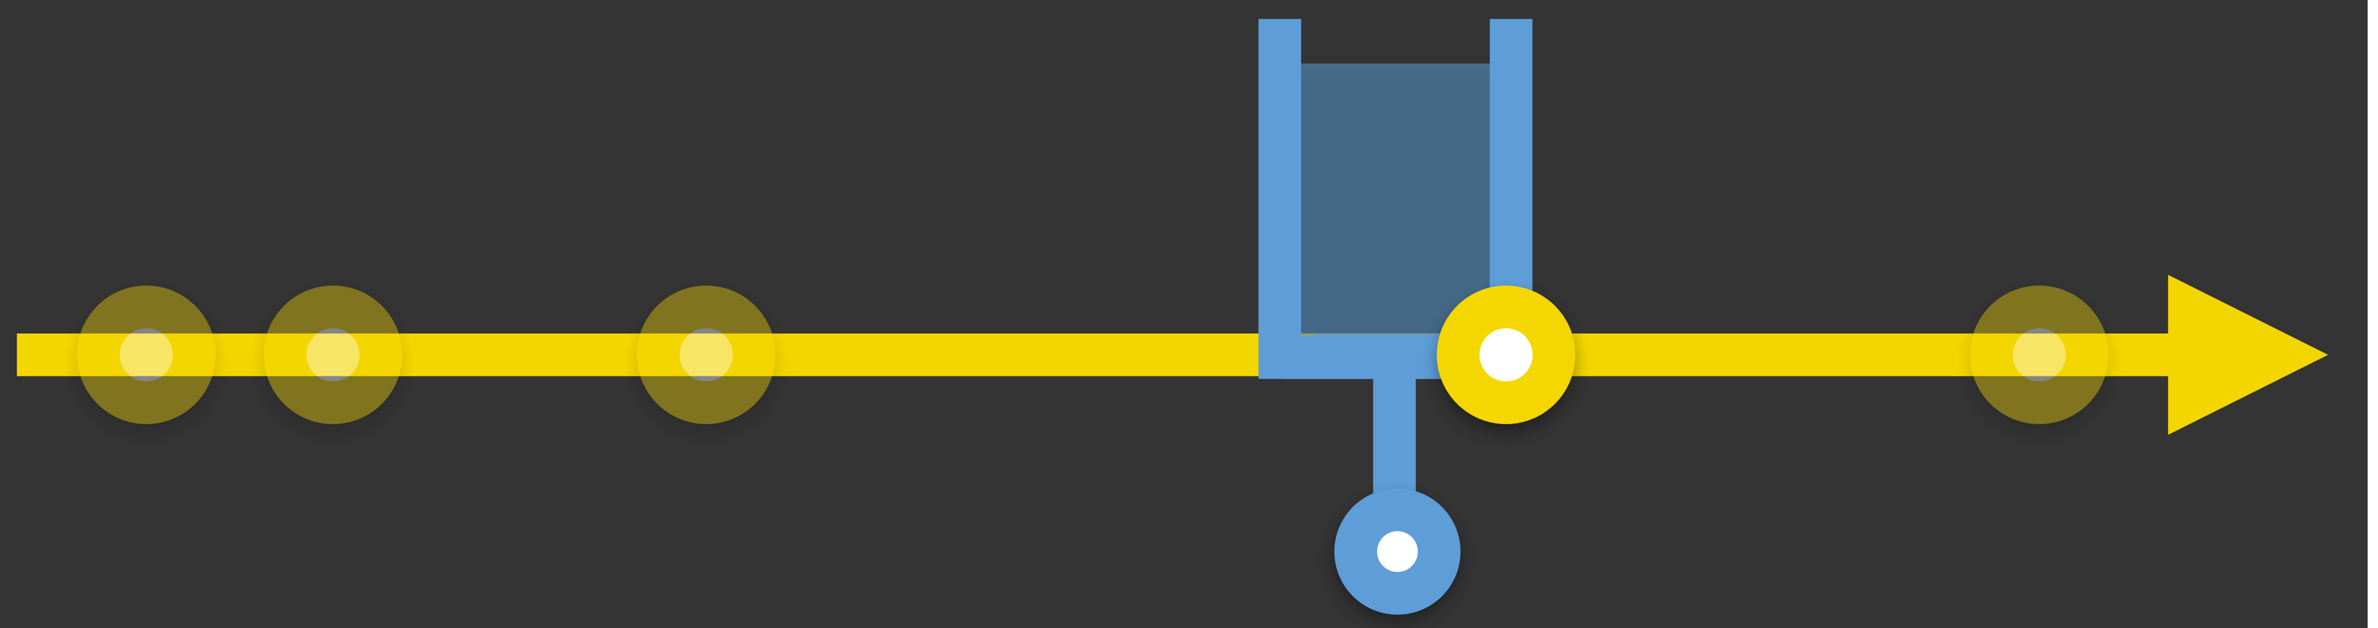

Collision Estimator

Photon Mapping

Neutron Transport

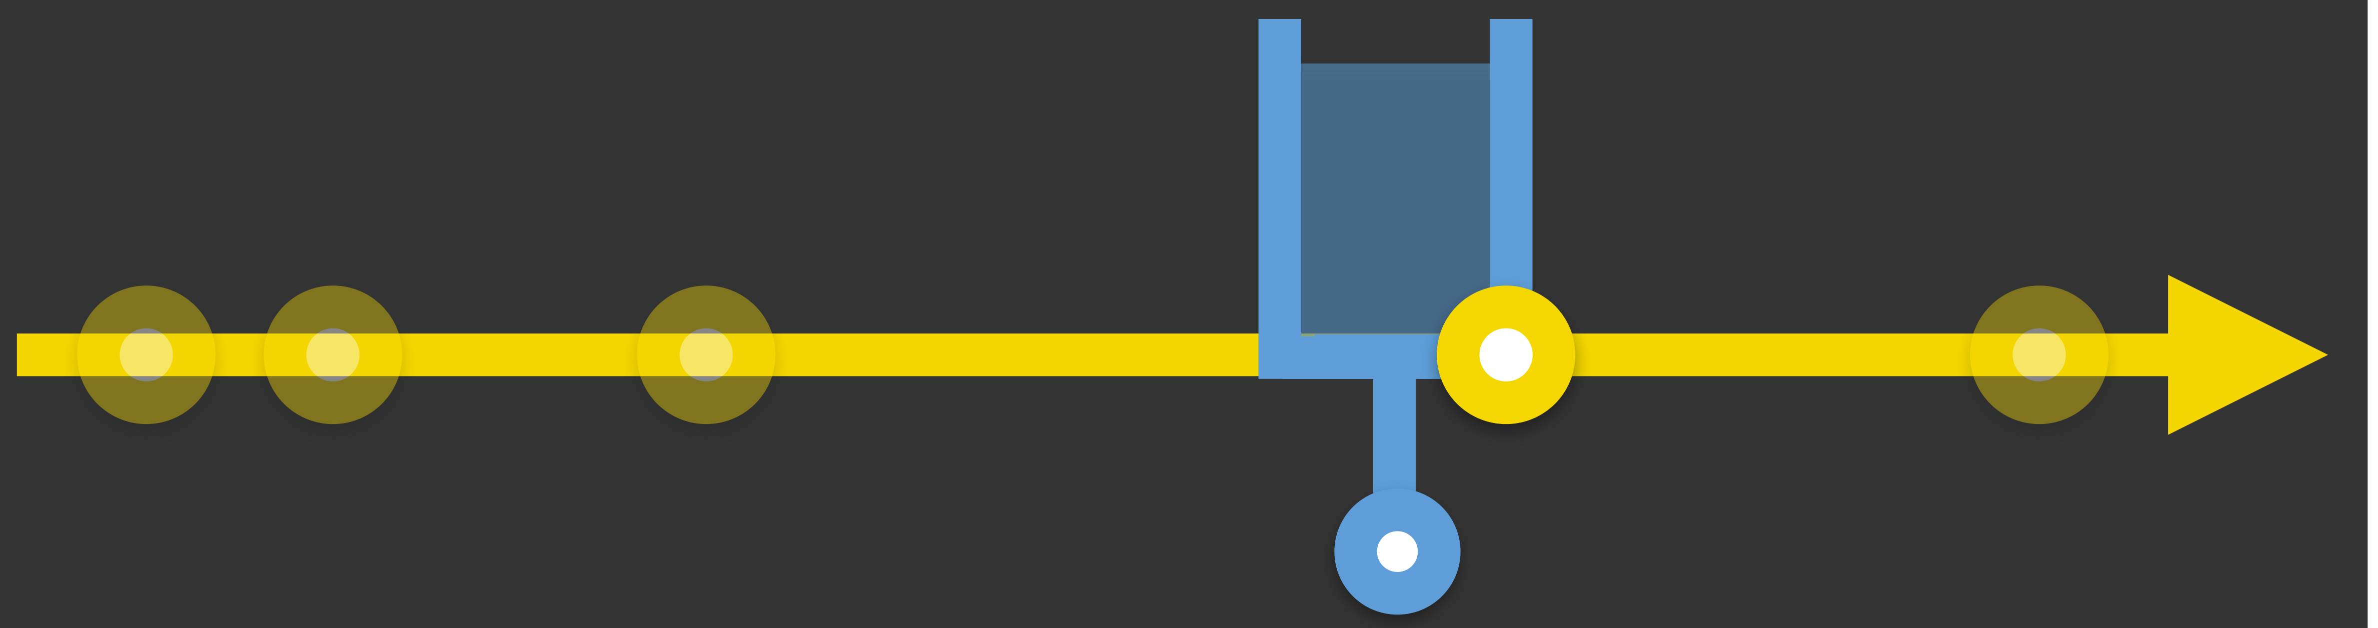

Collision Estimator

Photon Mapping

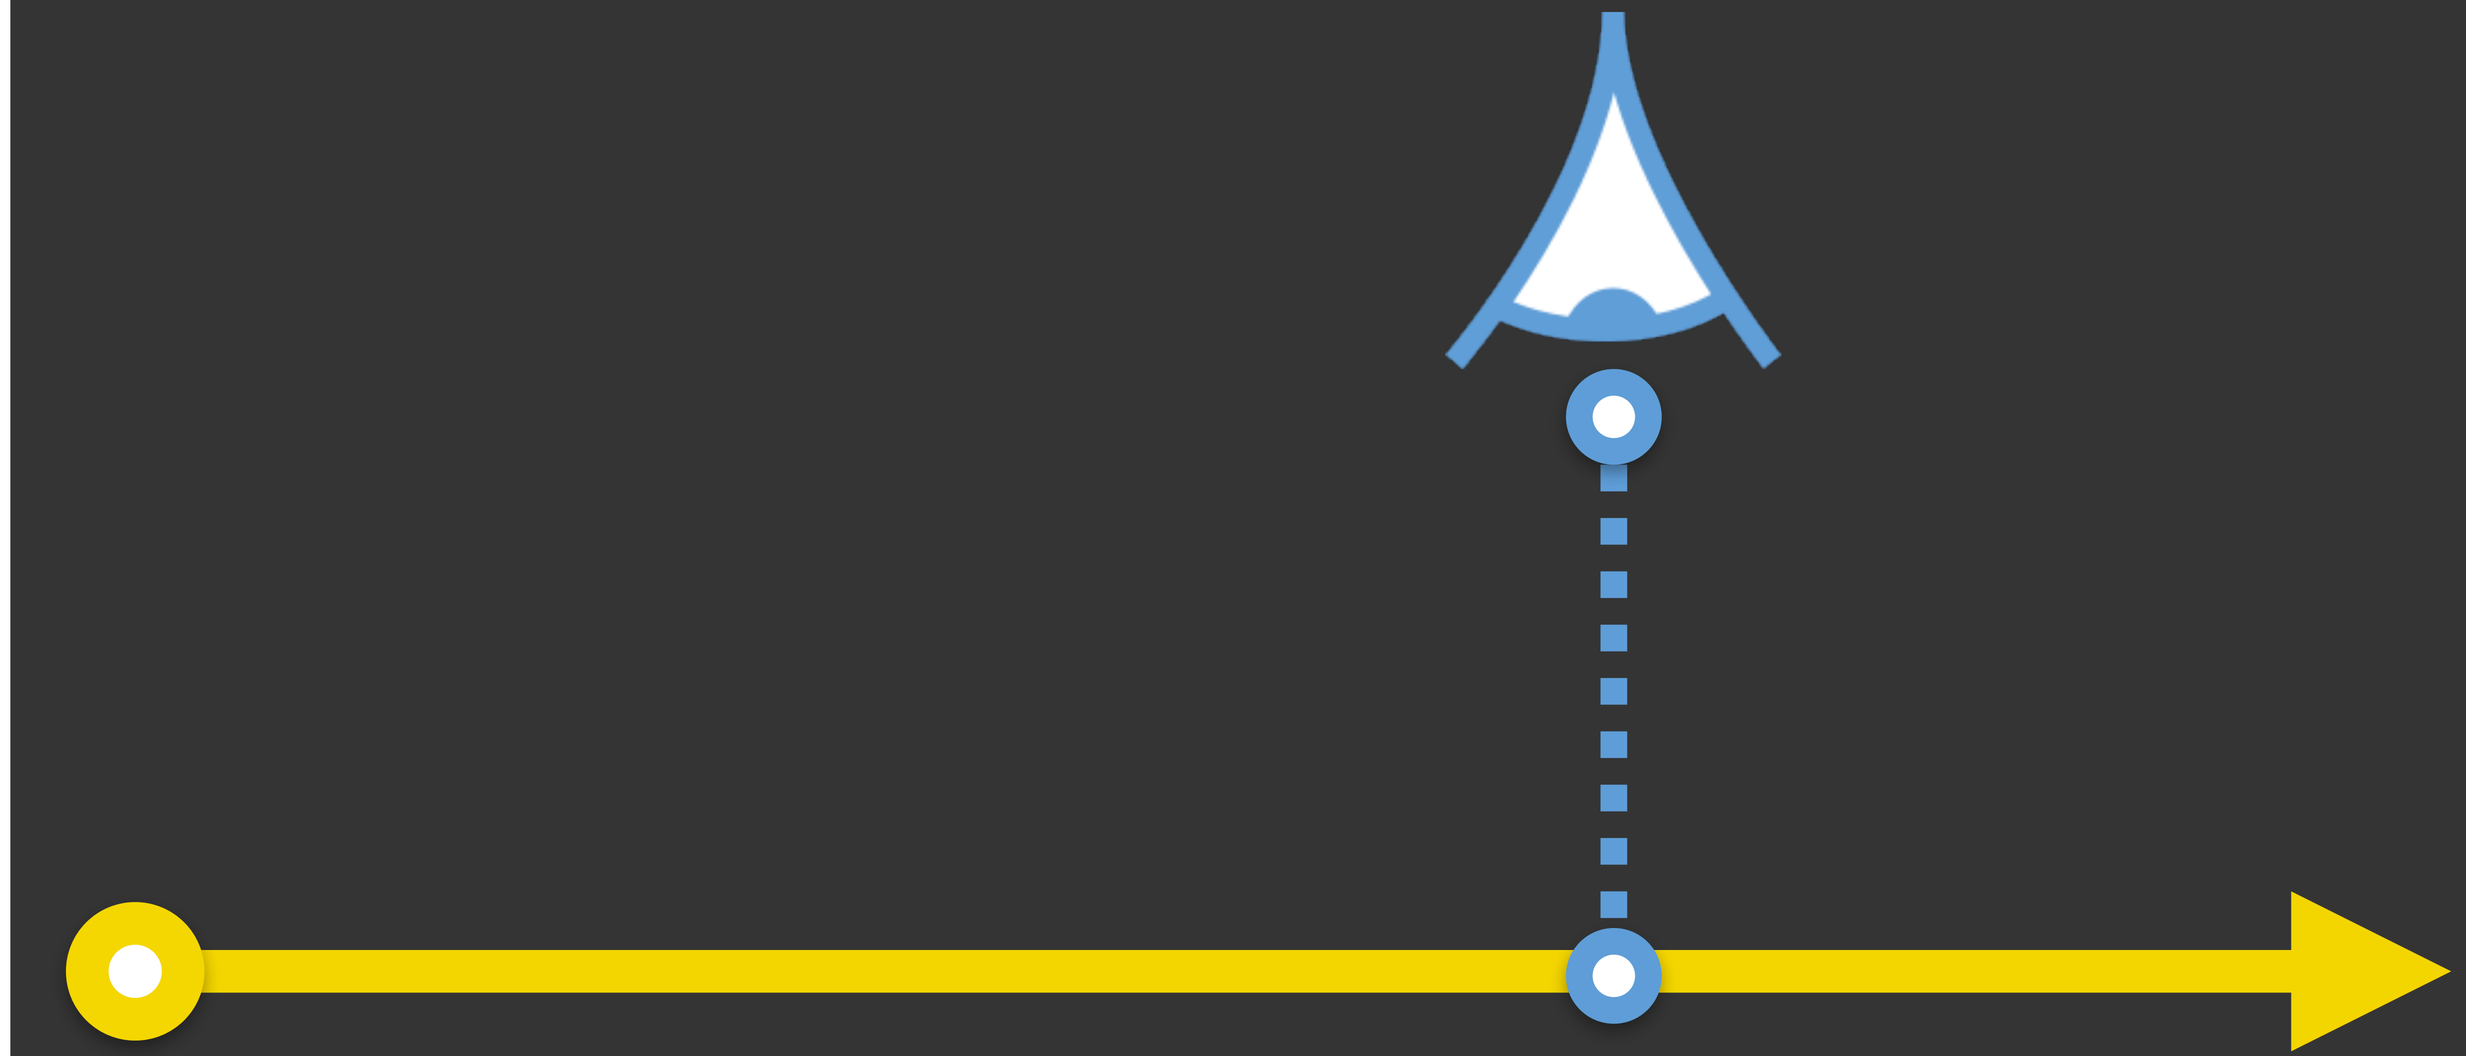

Neutron Transport

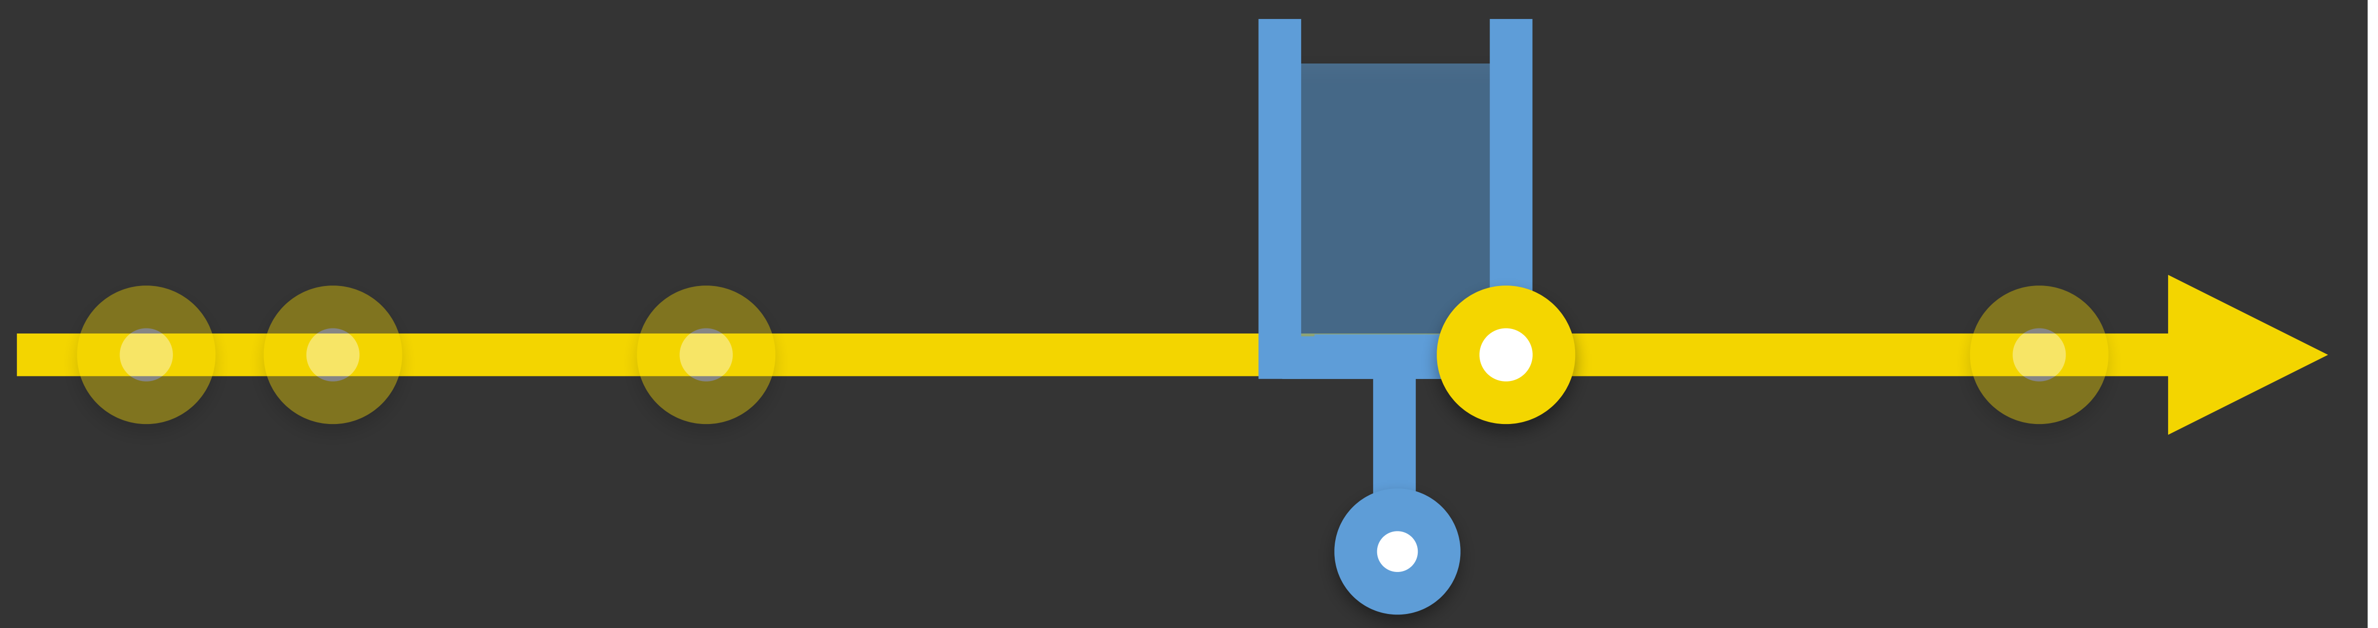

Collision Estimator

Photon Mapping

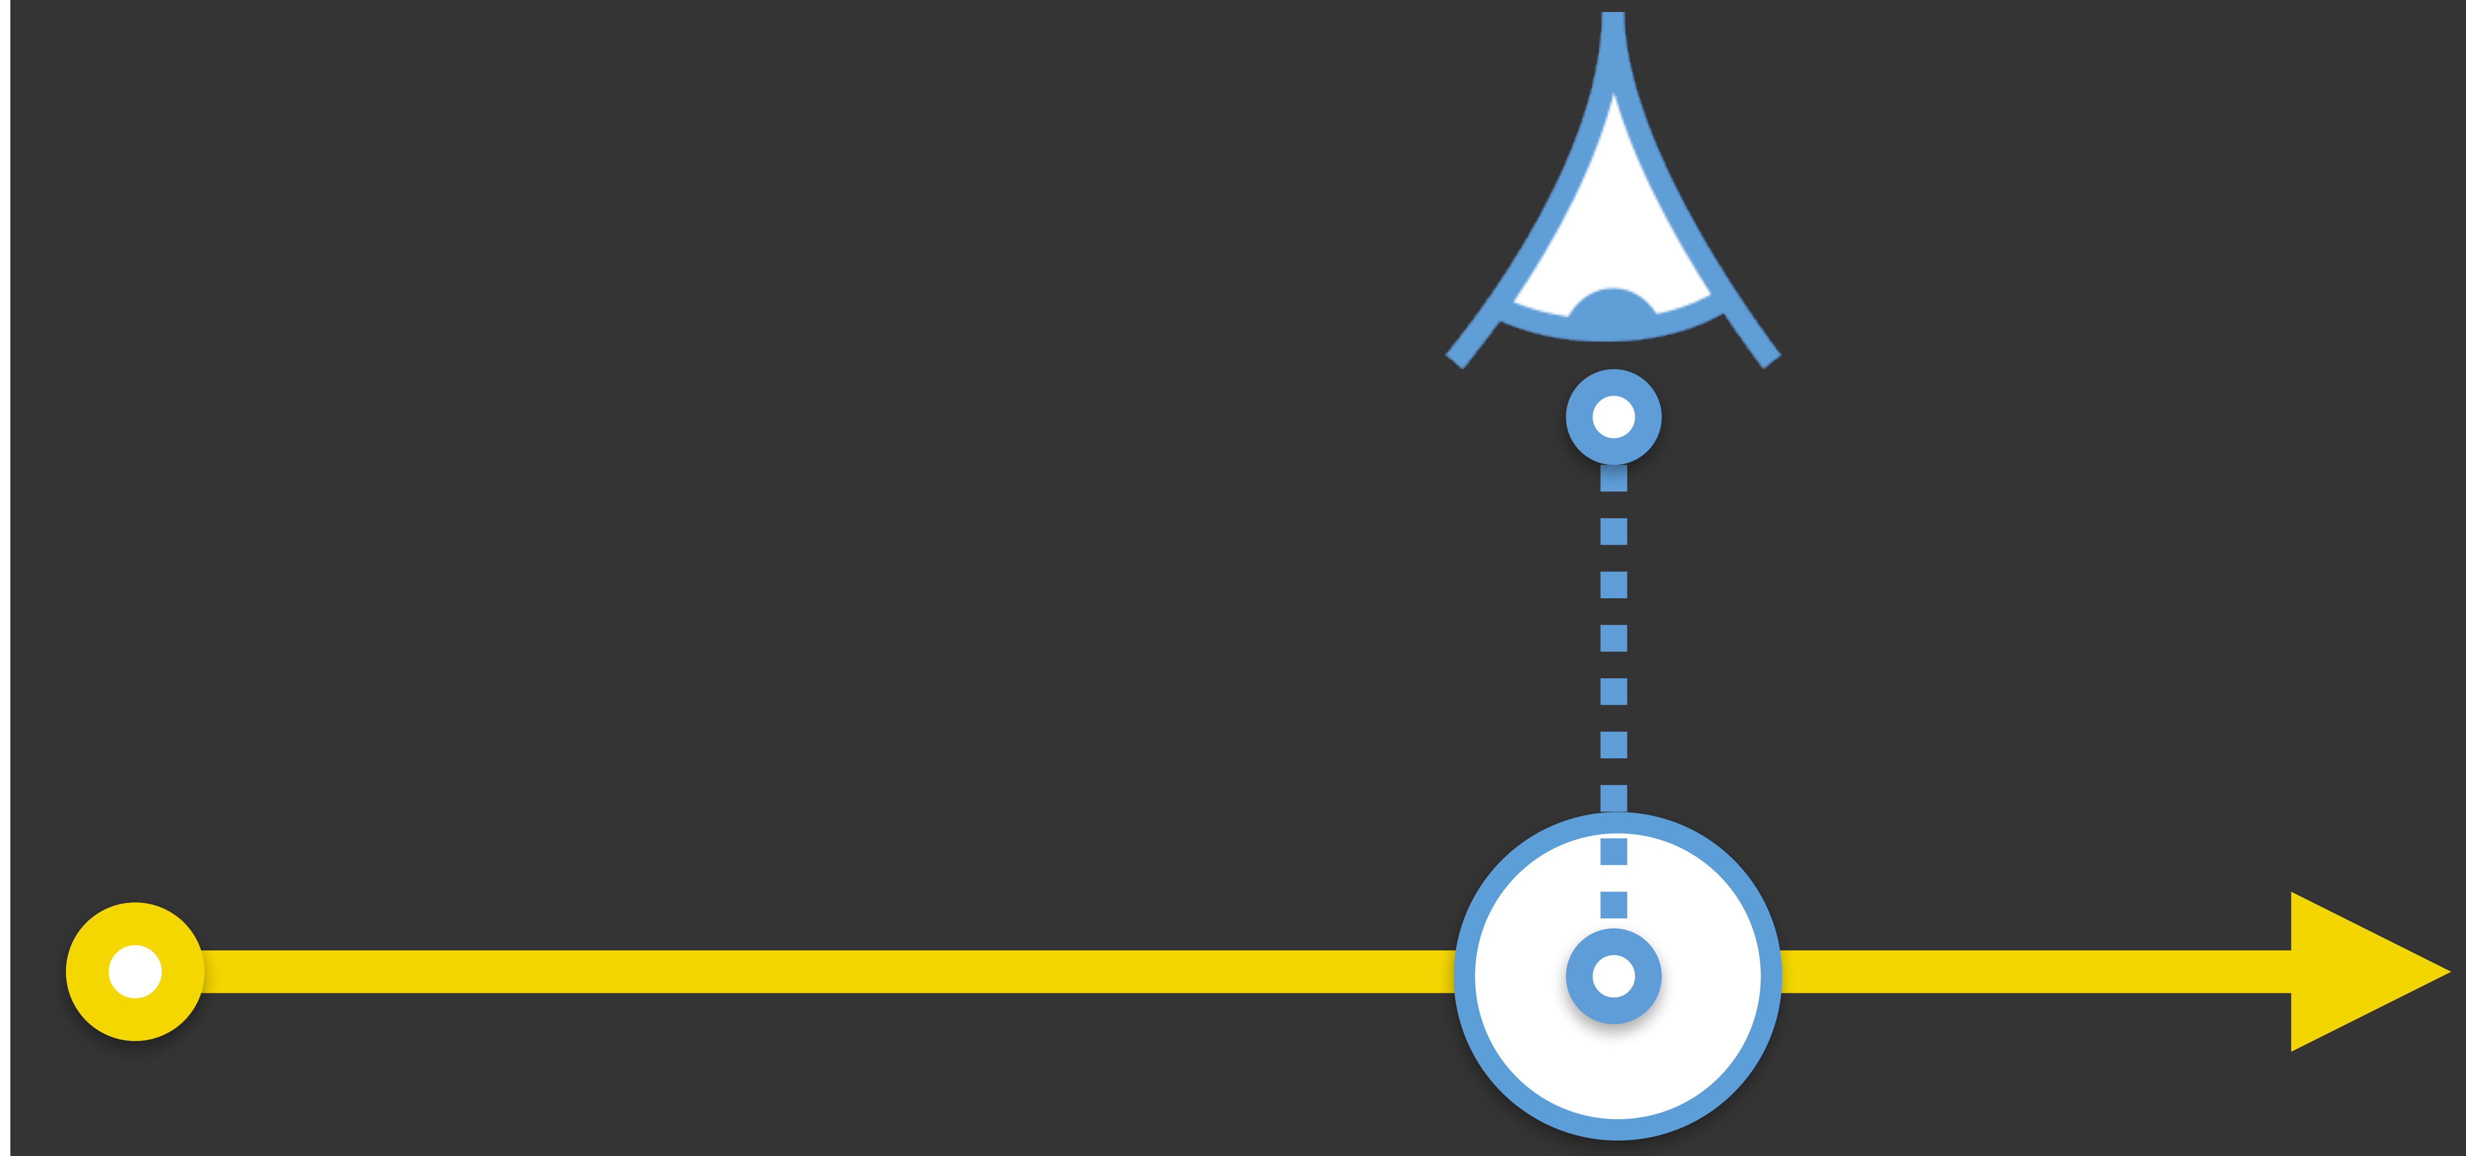

Neutron Transport

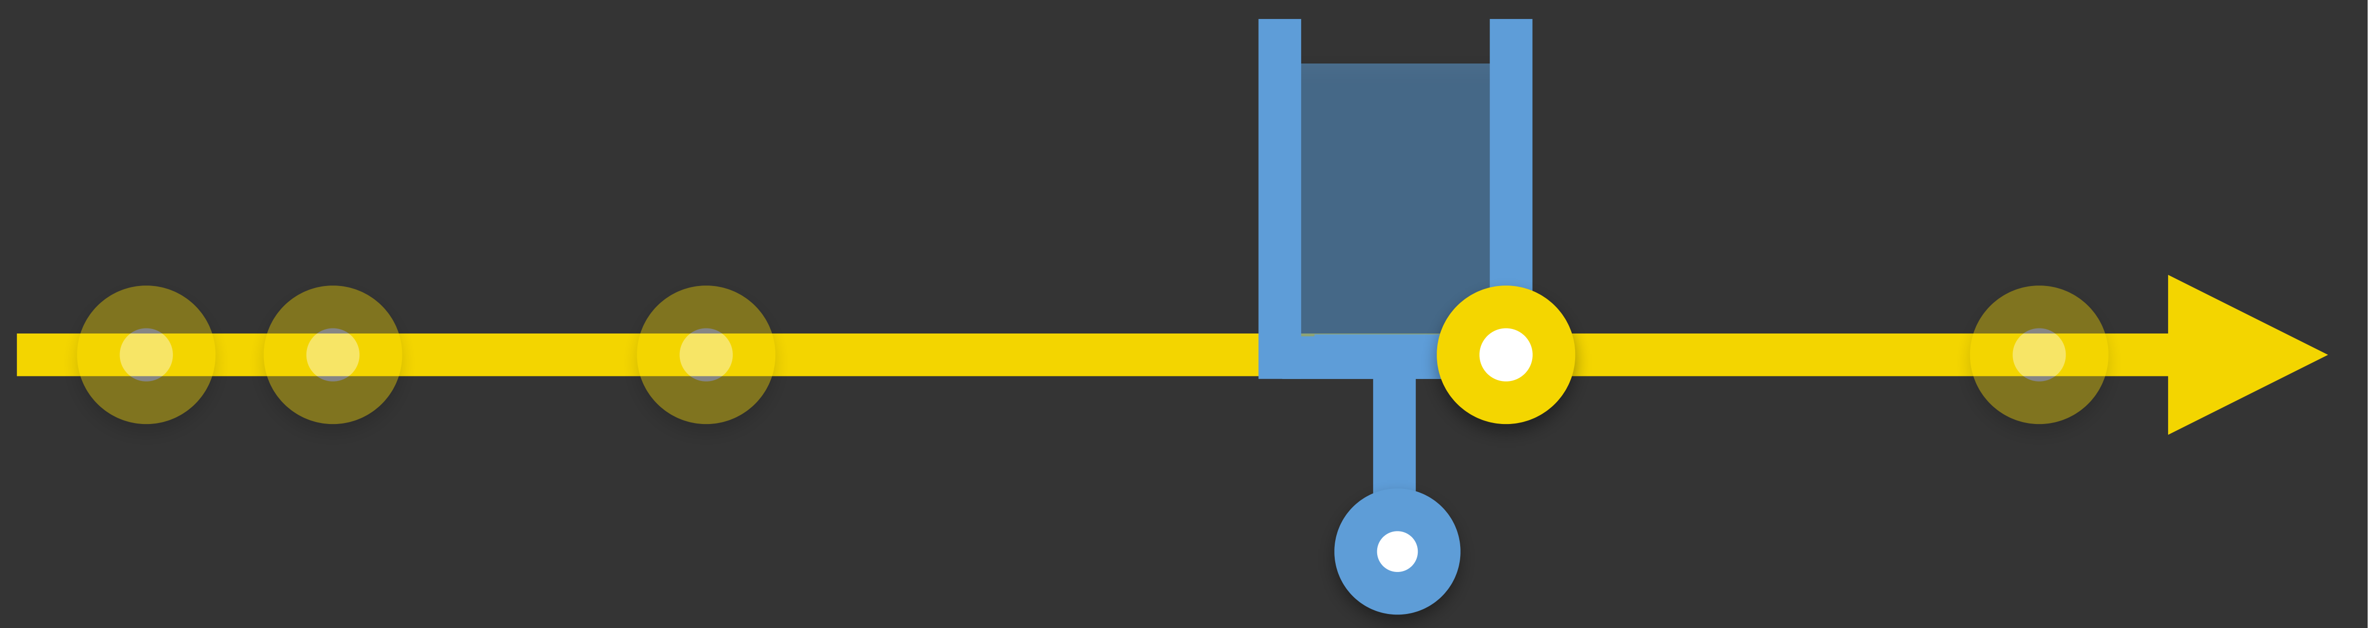

Collision Estimator

Photon Mapping

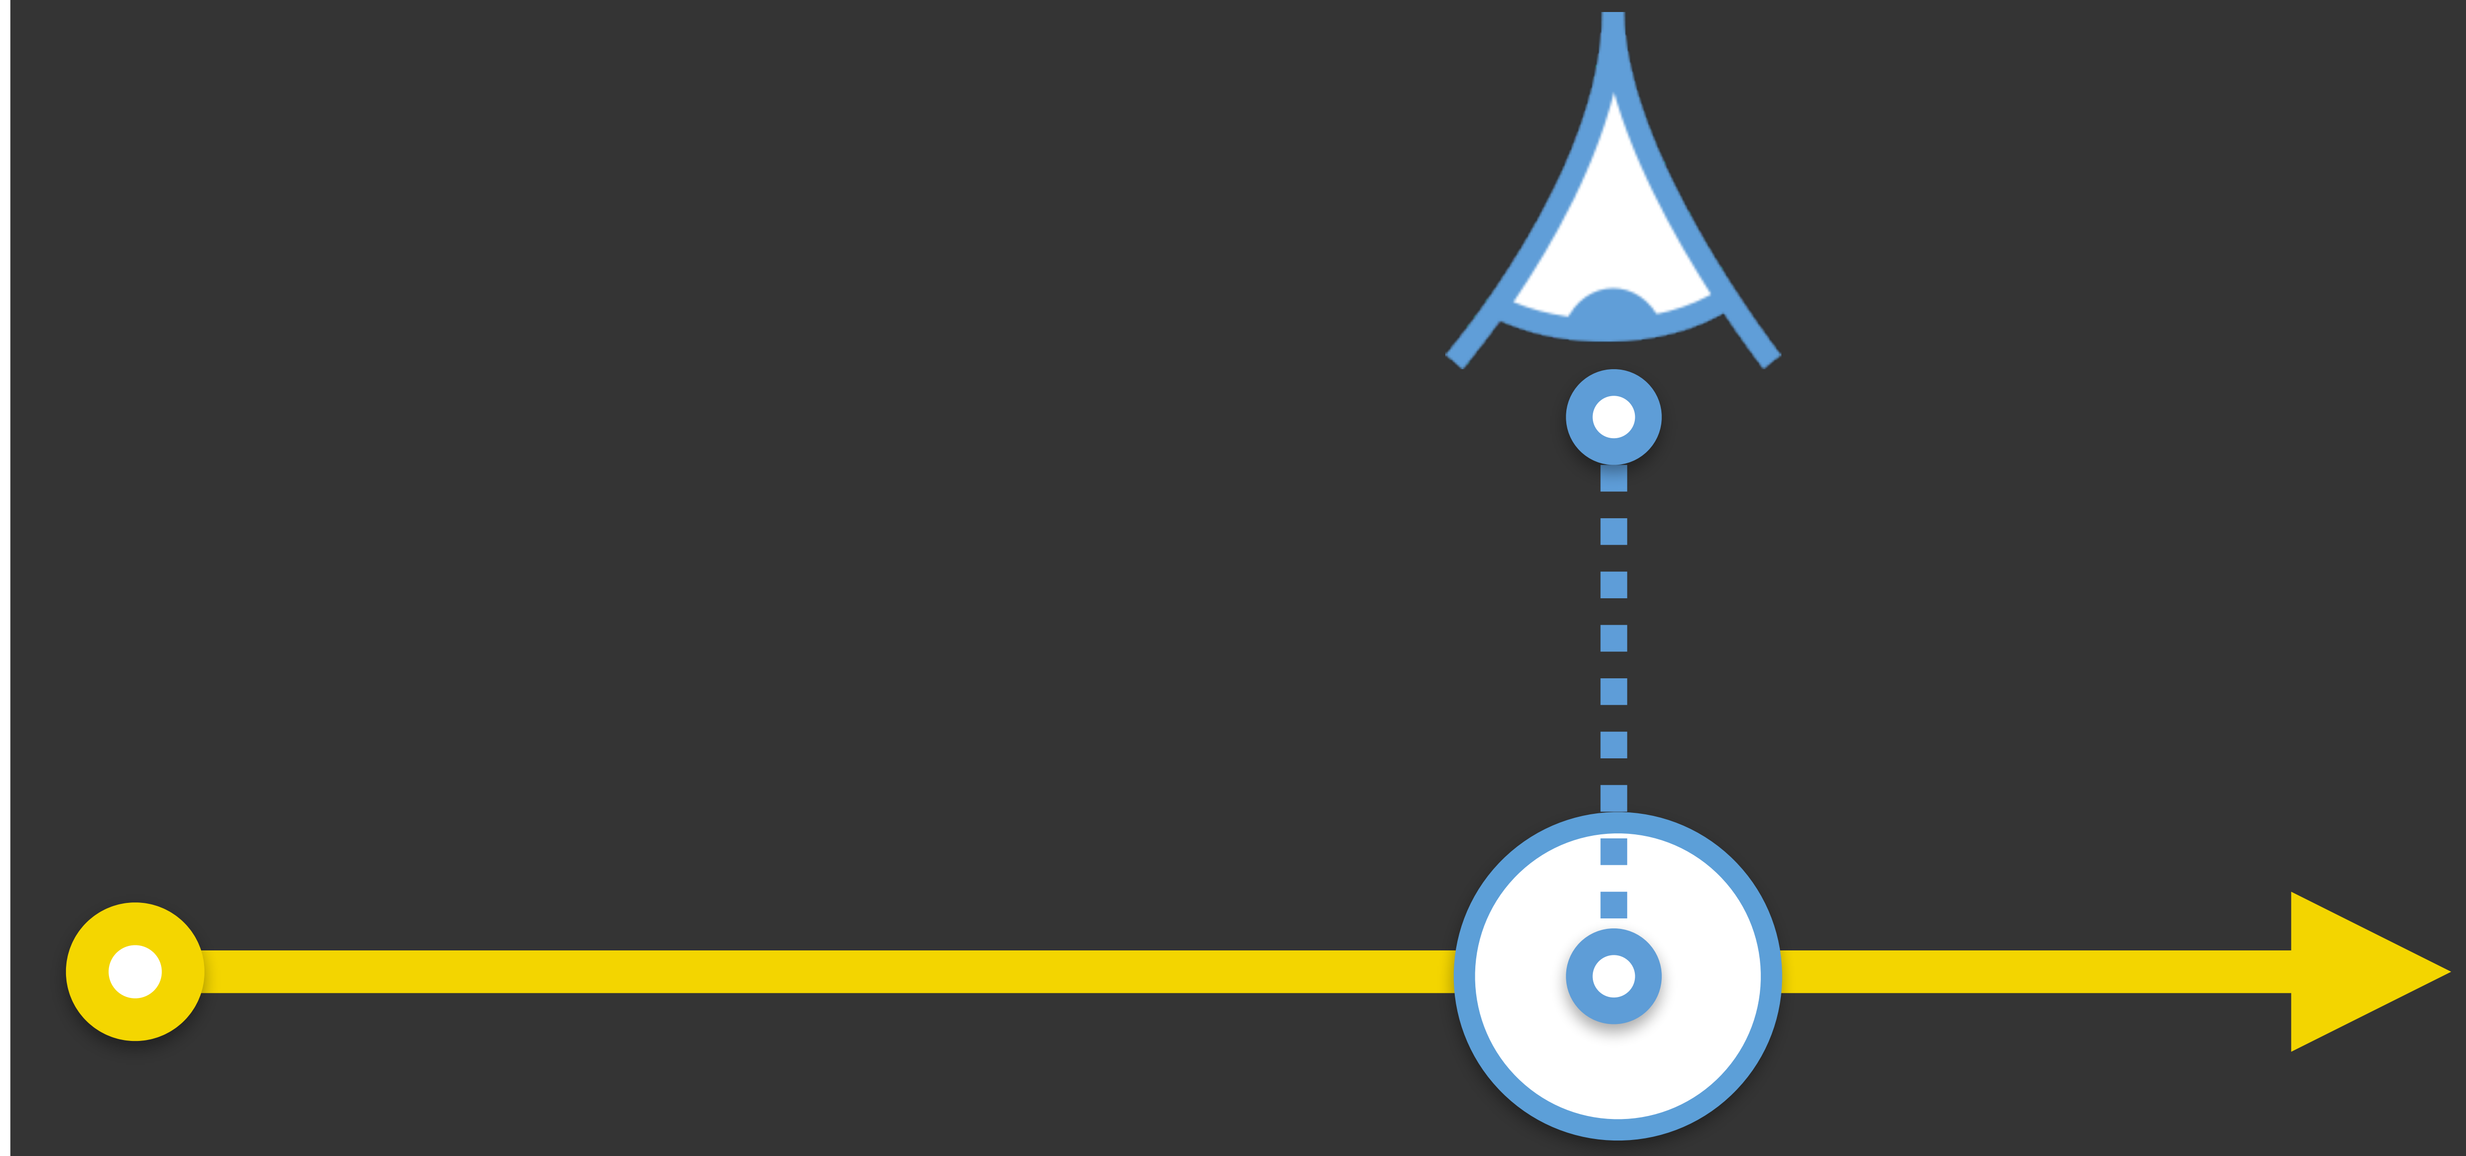

Photon Points

Neutron Transport

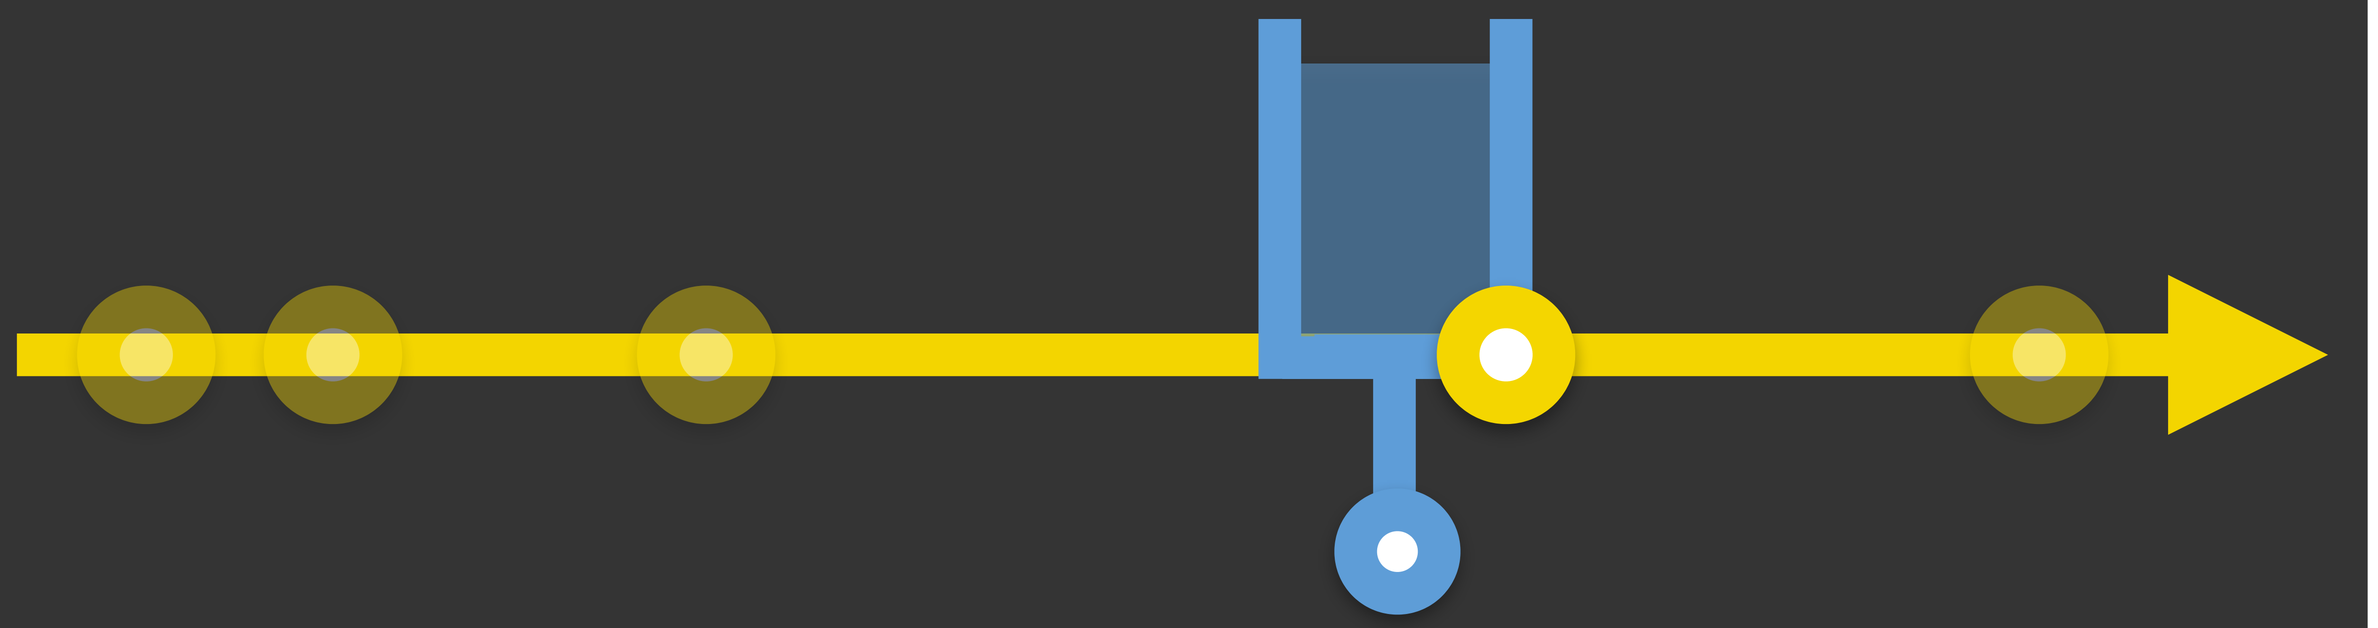

Collision Estimator

Photon Mapping

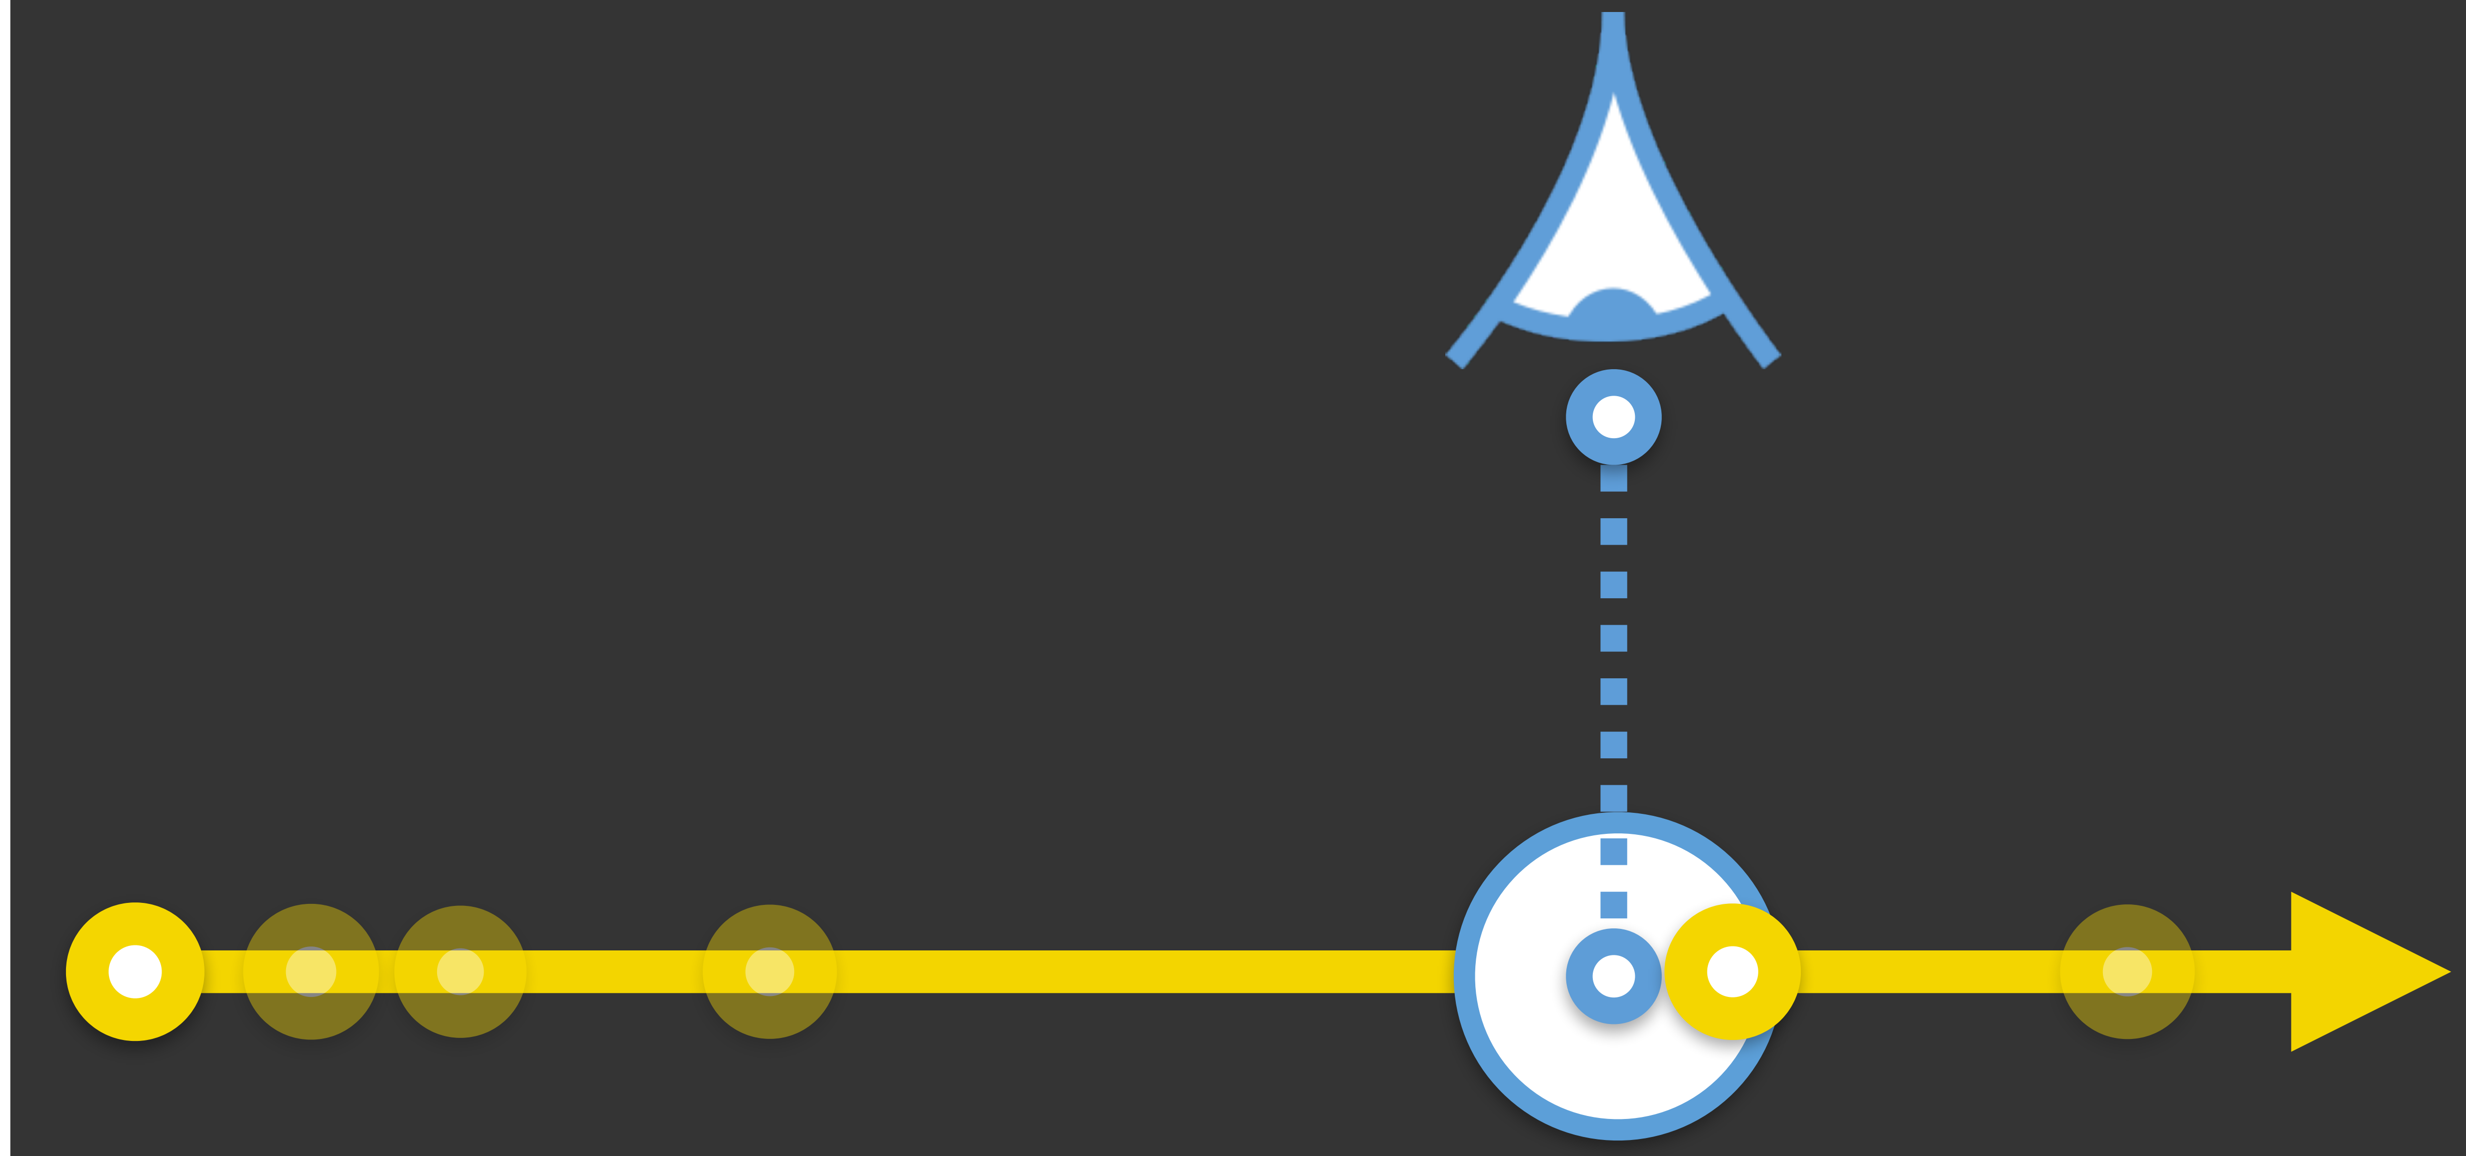

Photon Points

Neutron Transport

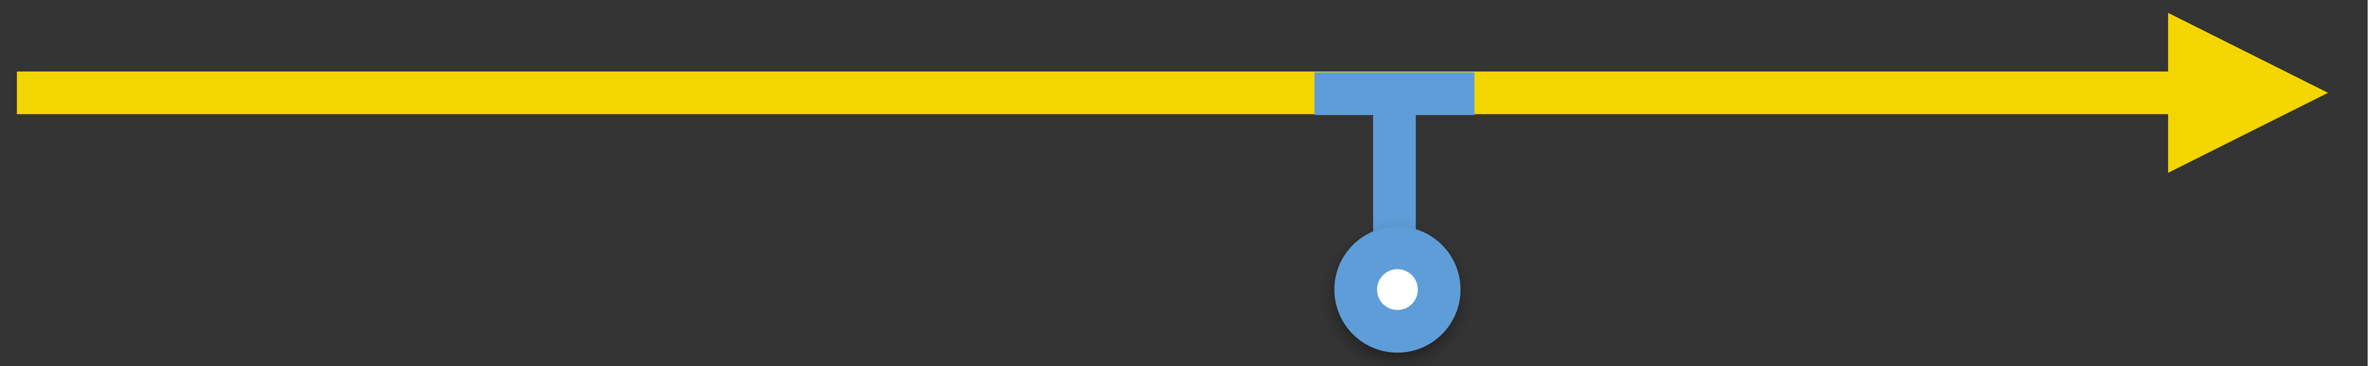

Photon Mapping

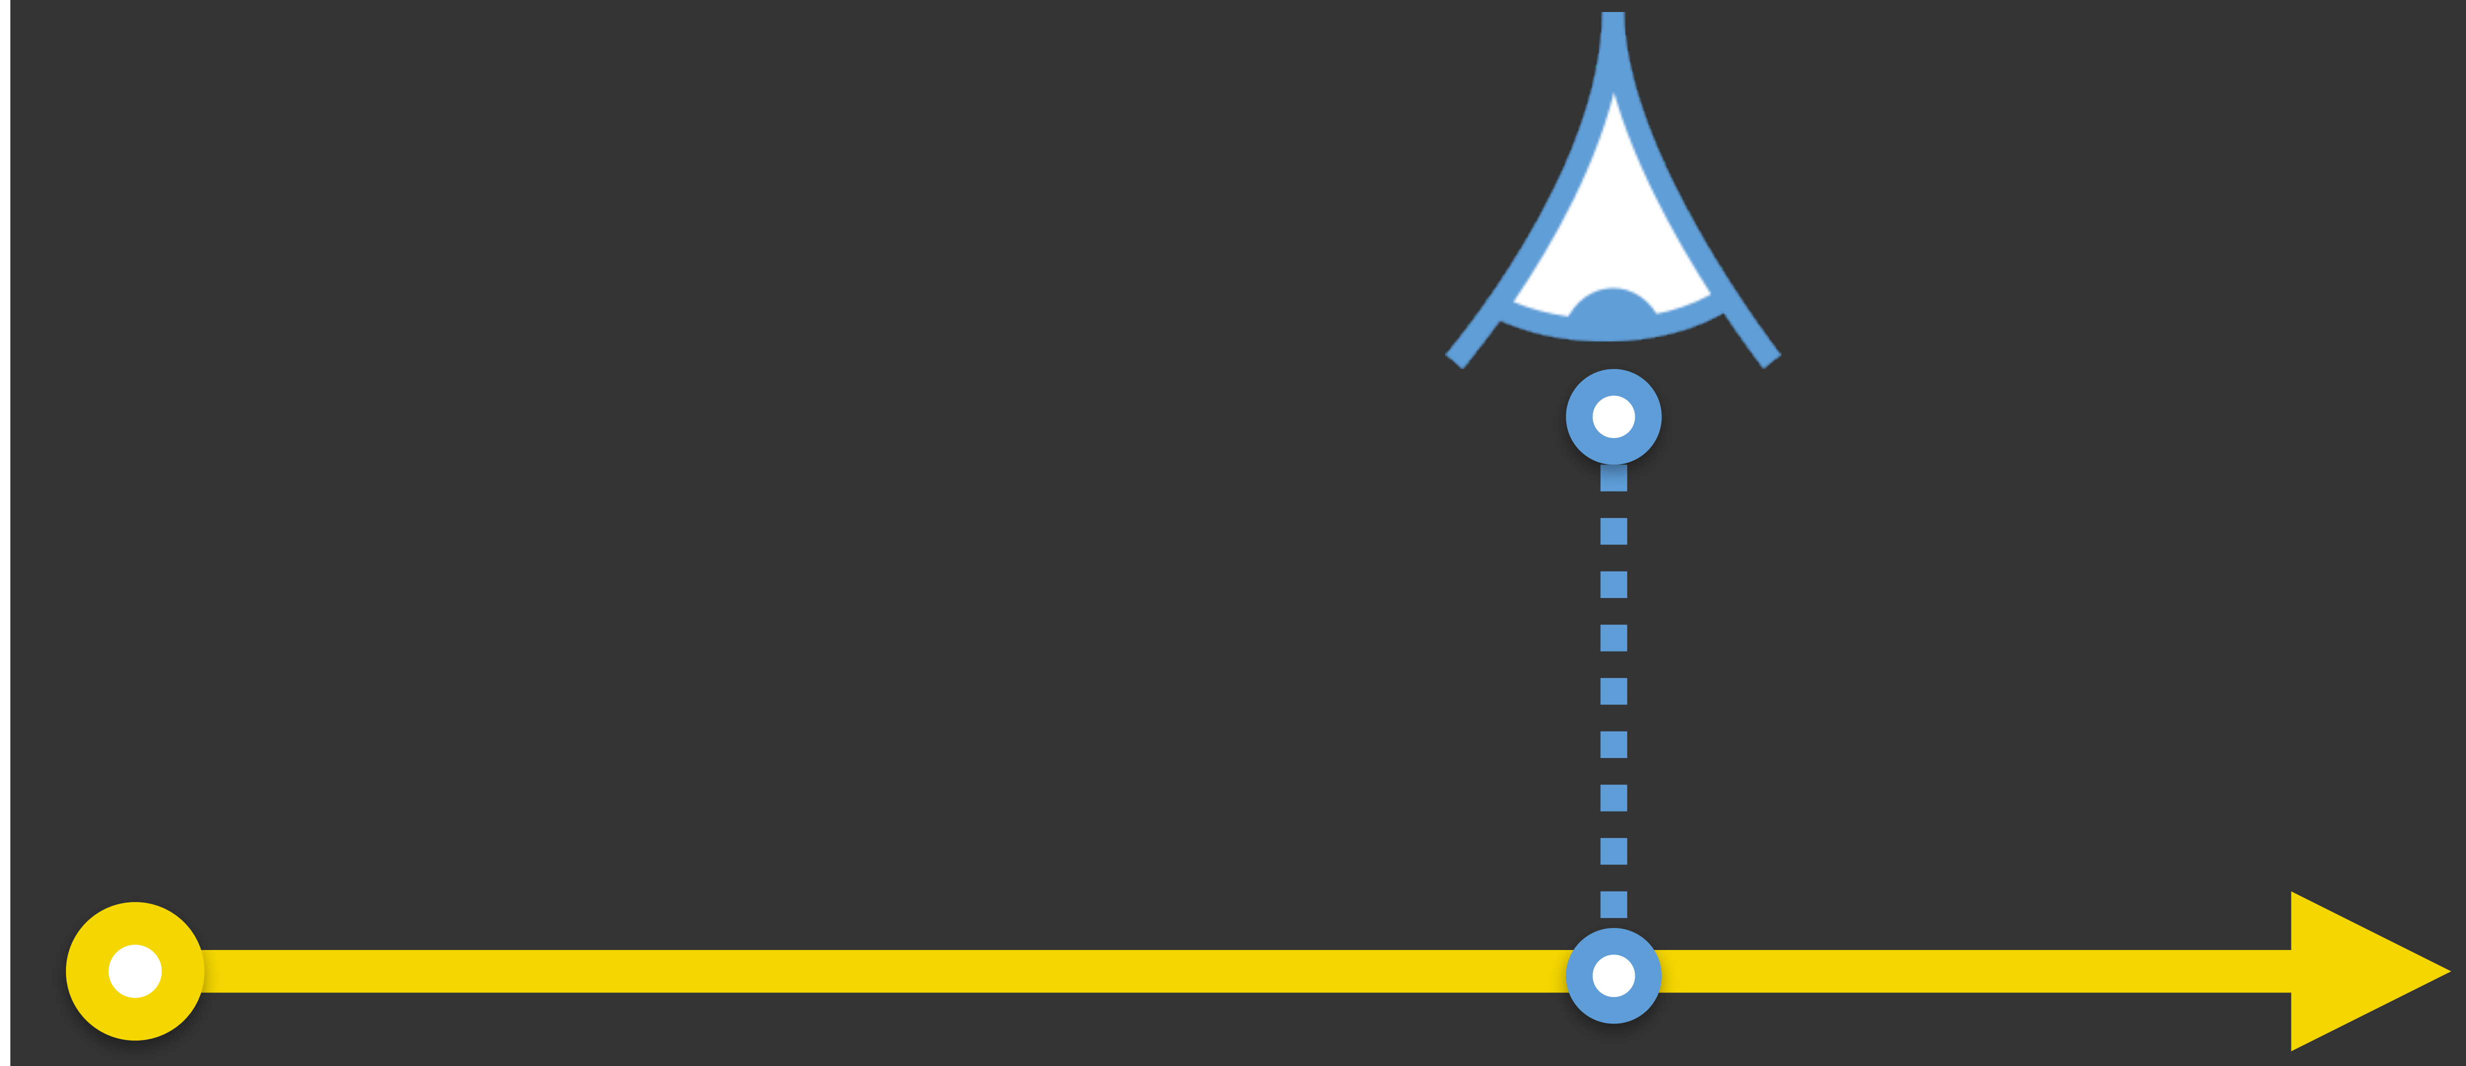

Photon Points

Neutron Transport

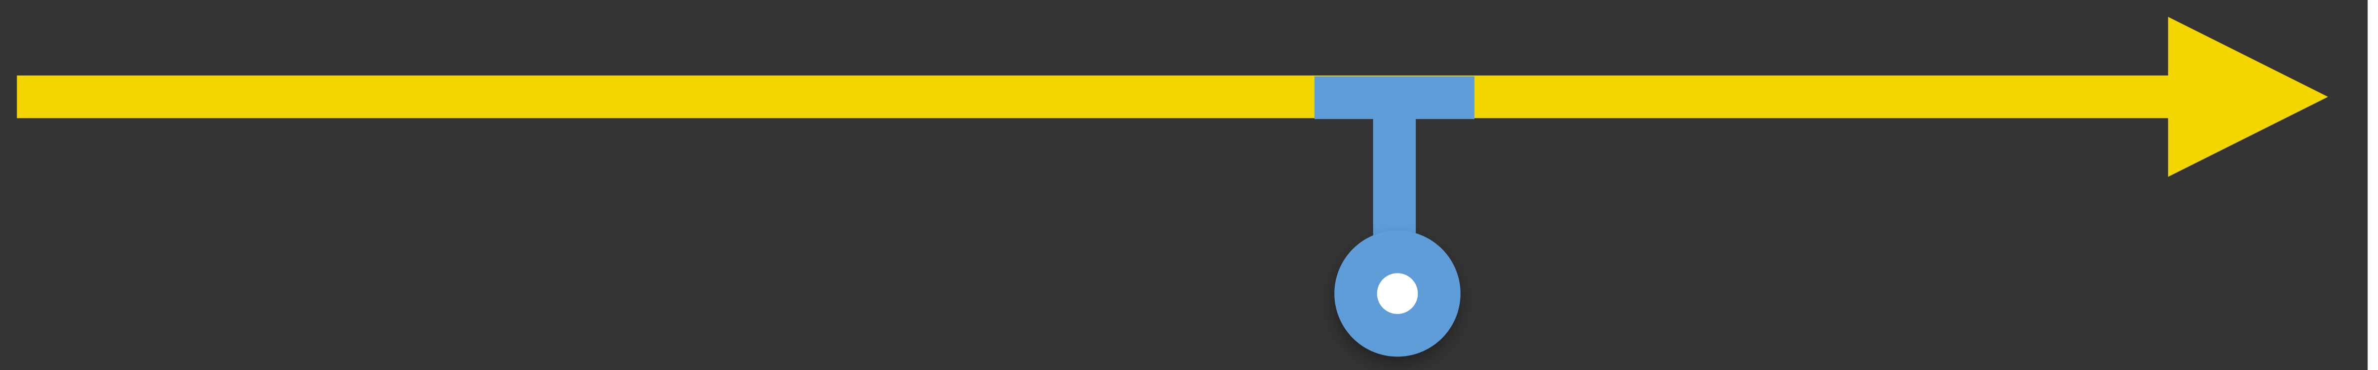

Photon Mapping

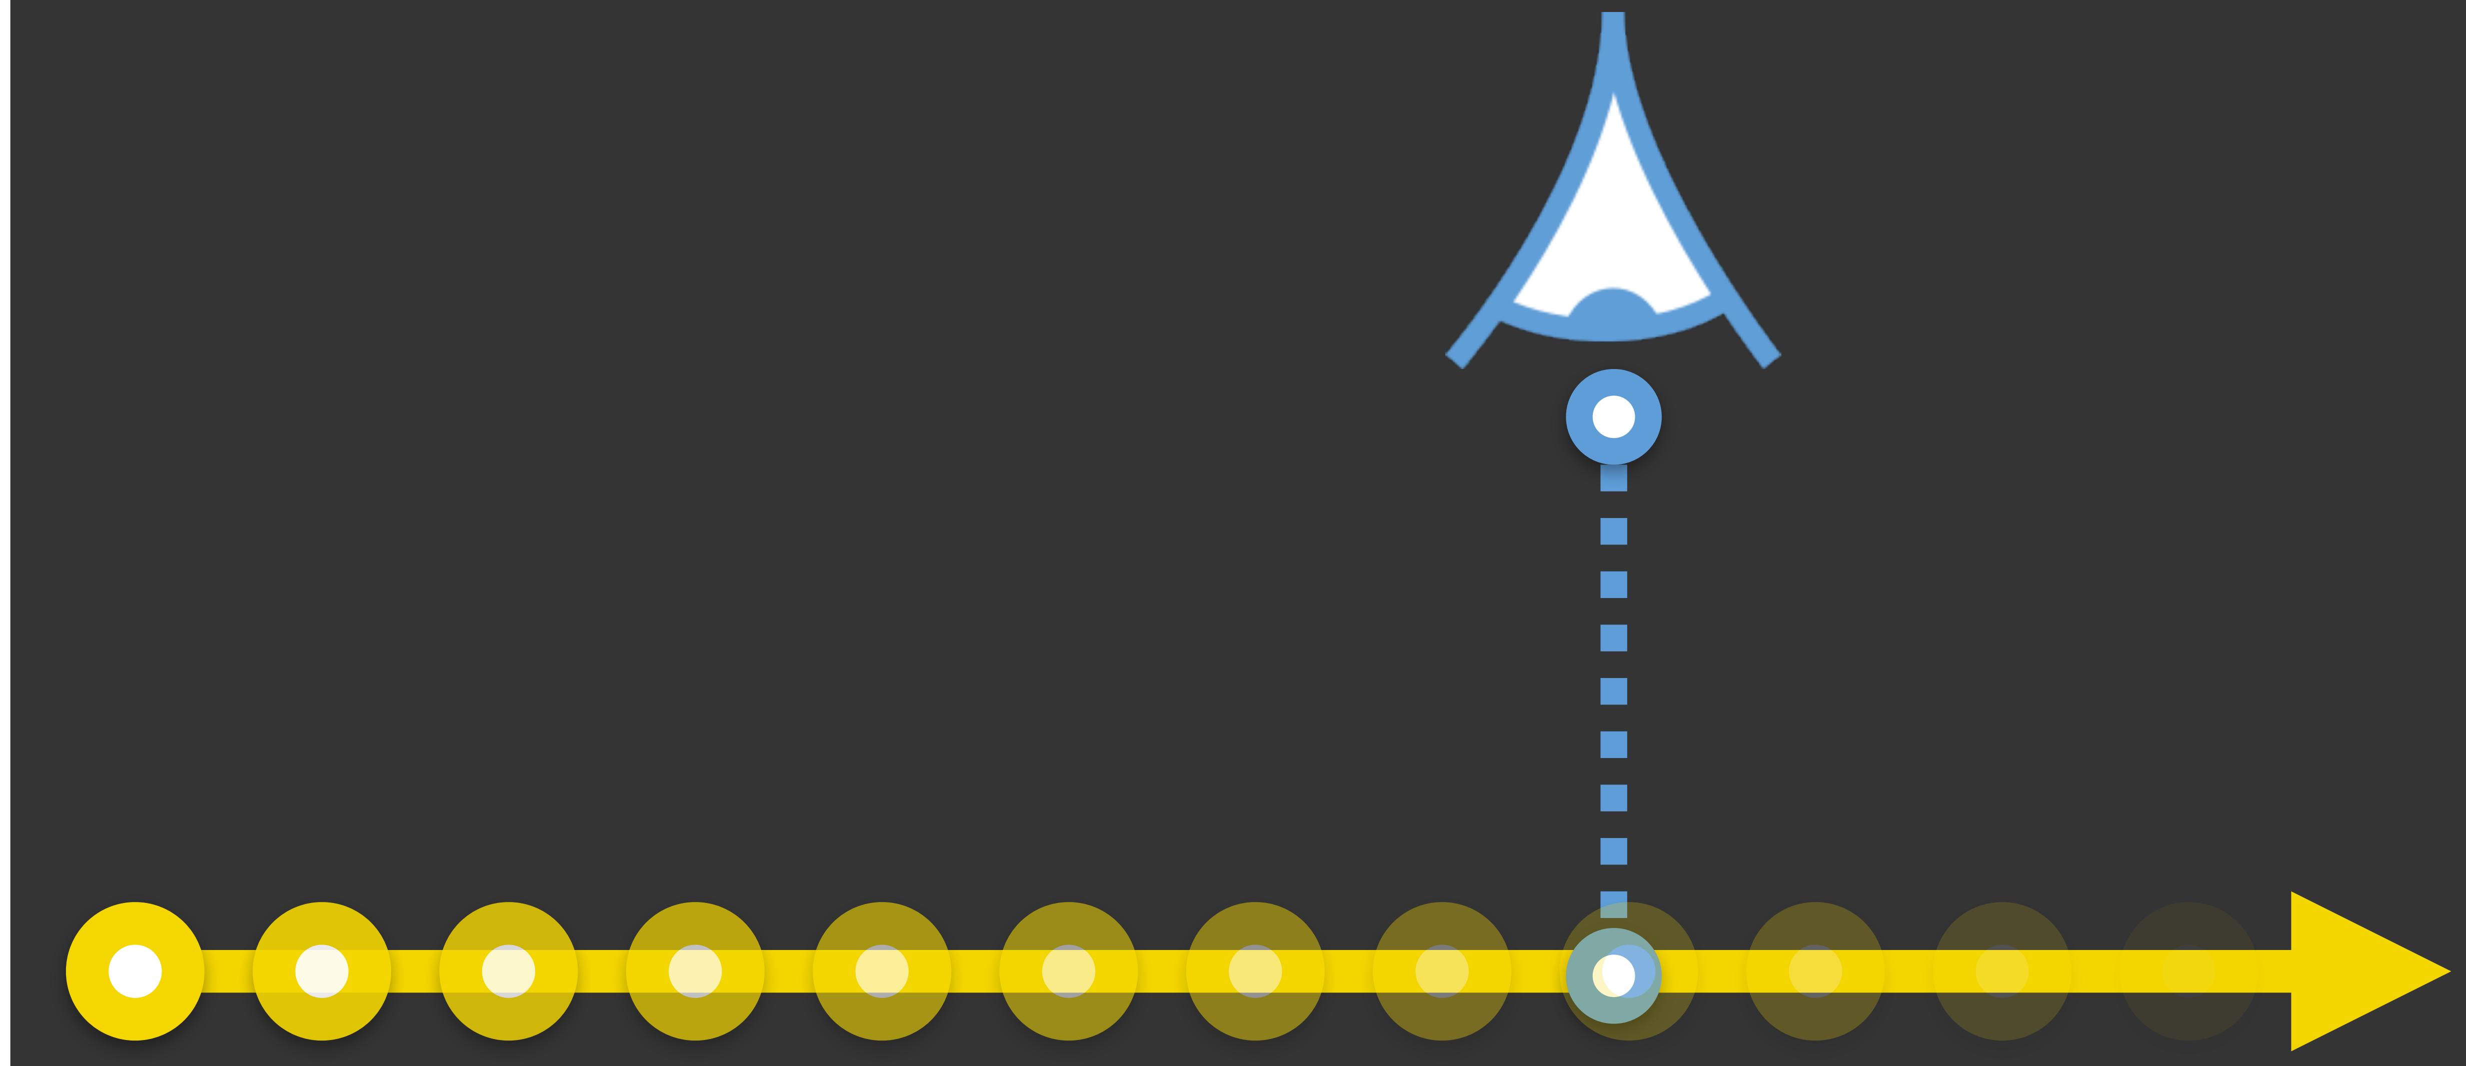

Photon Points

Neutron Transport

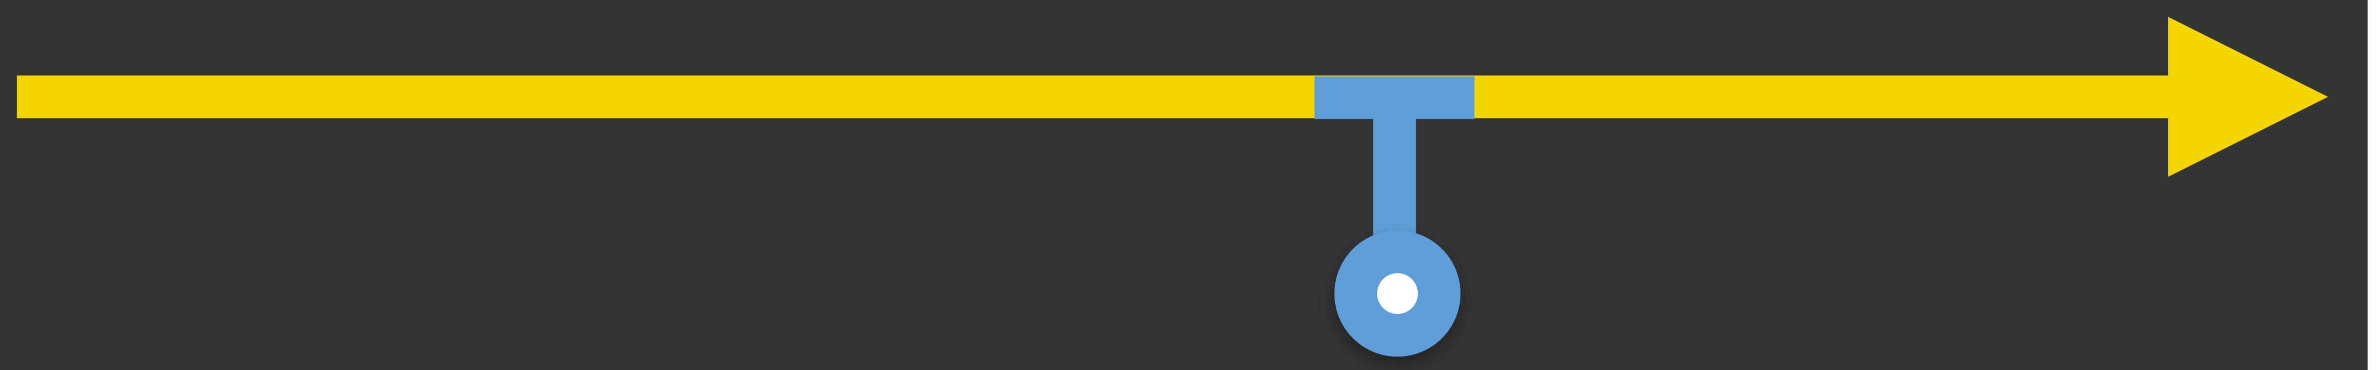

Photon Mapping

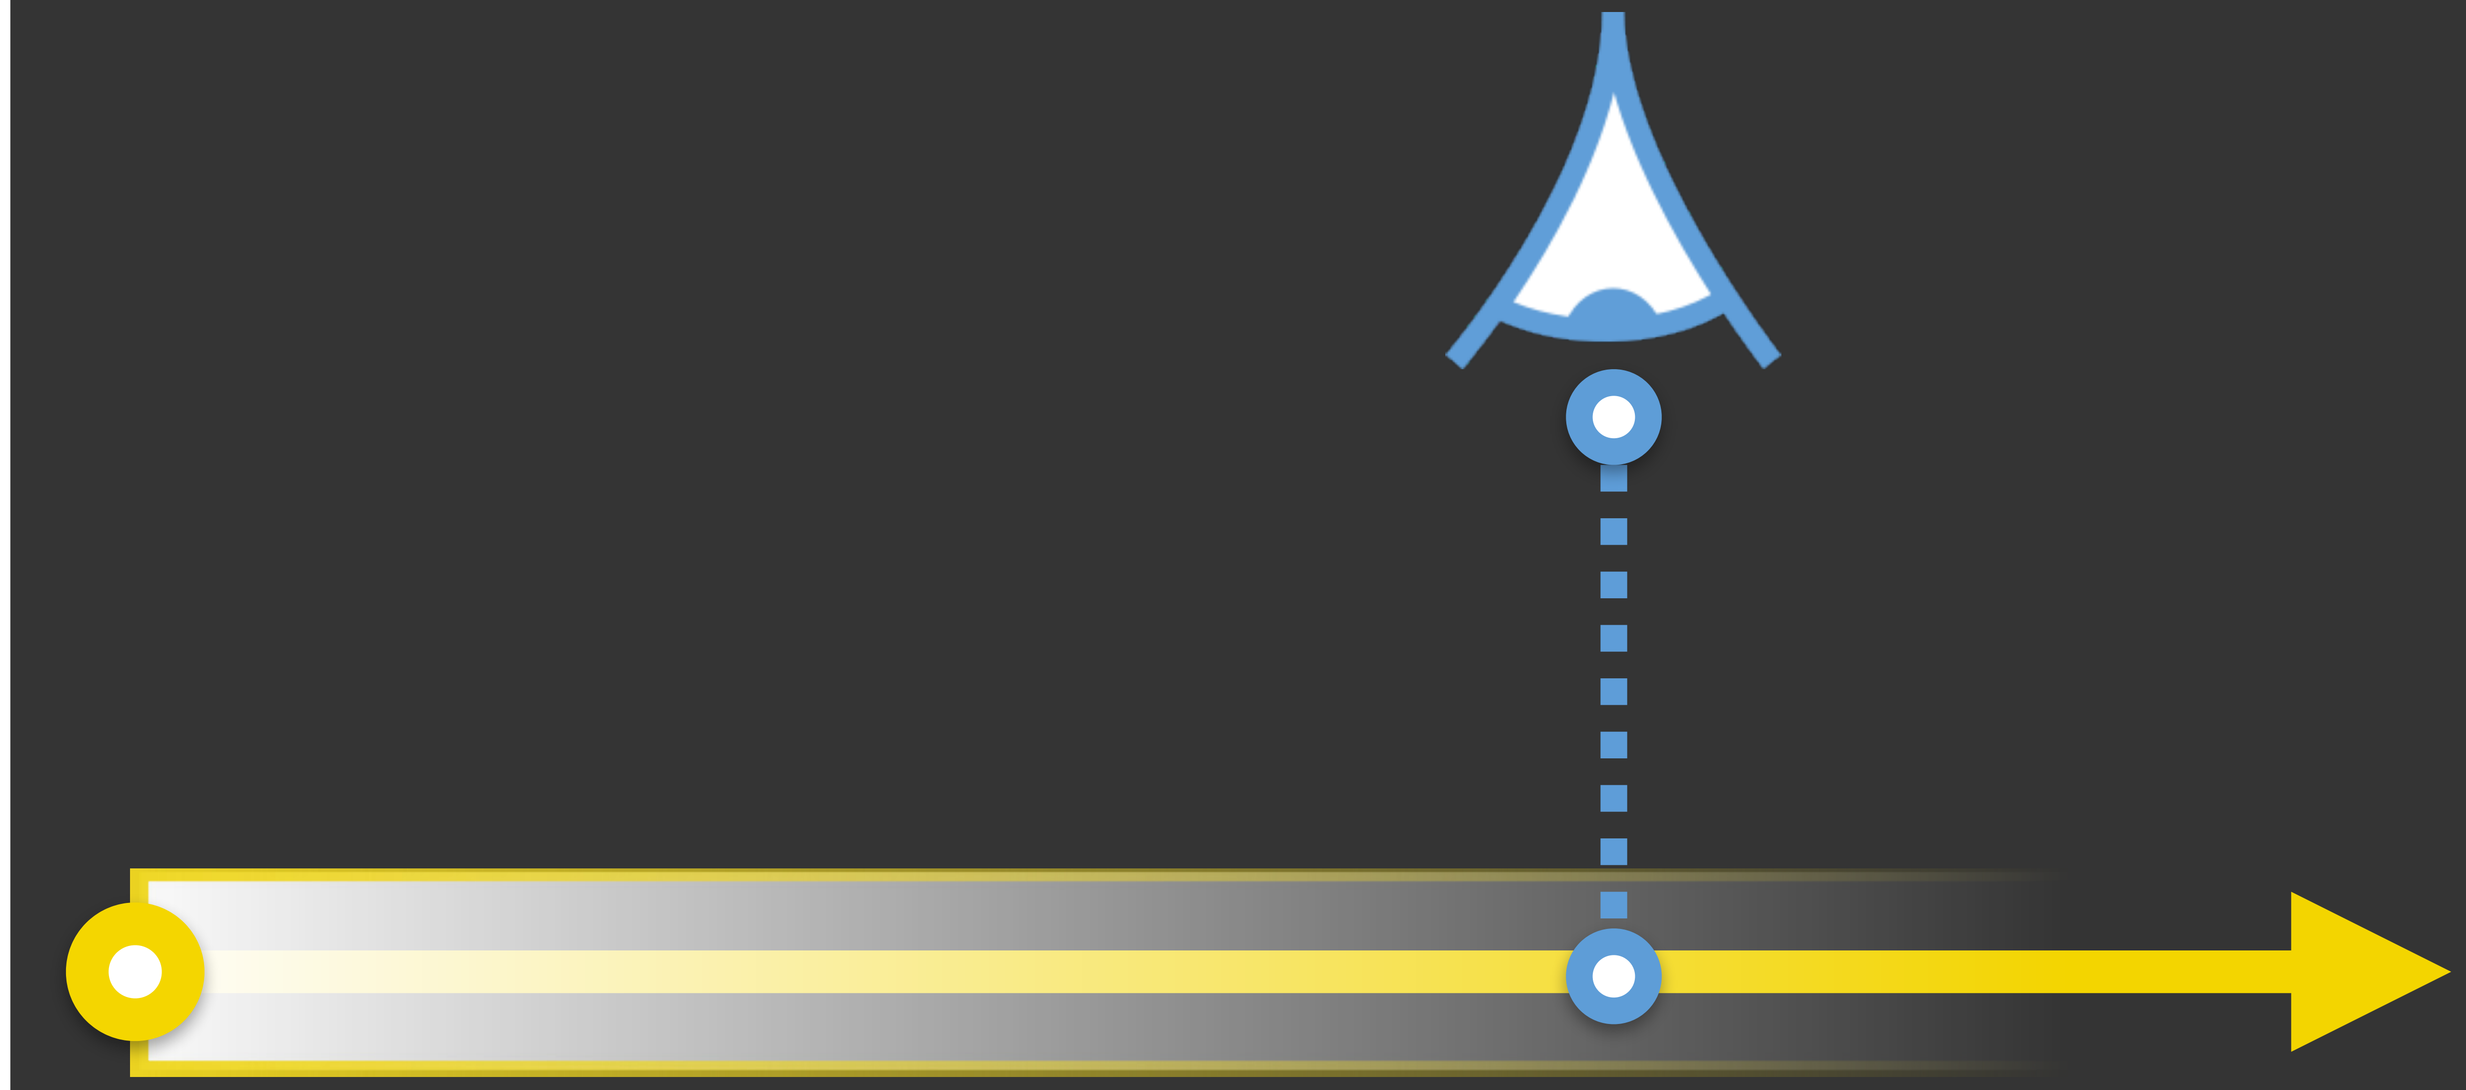

Long Beams

Neutron Transport

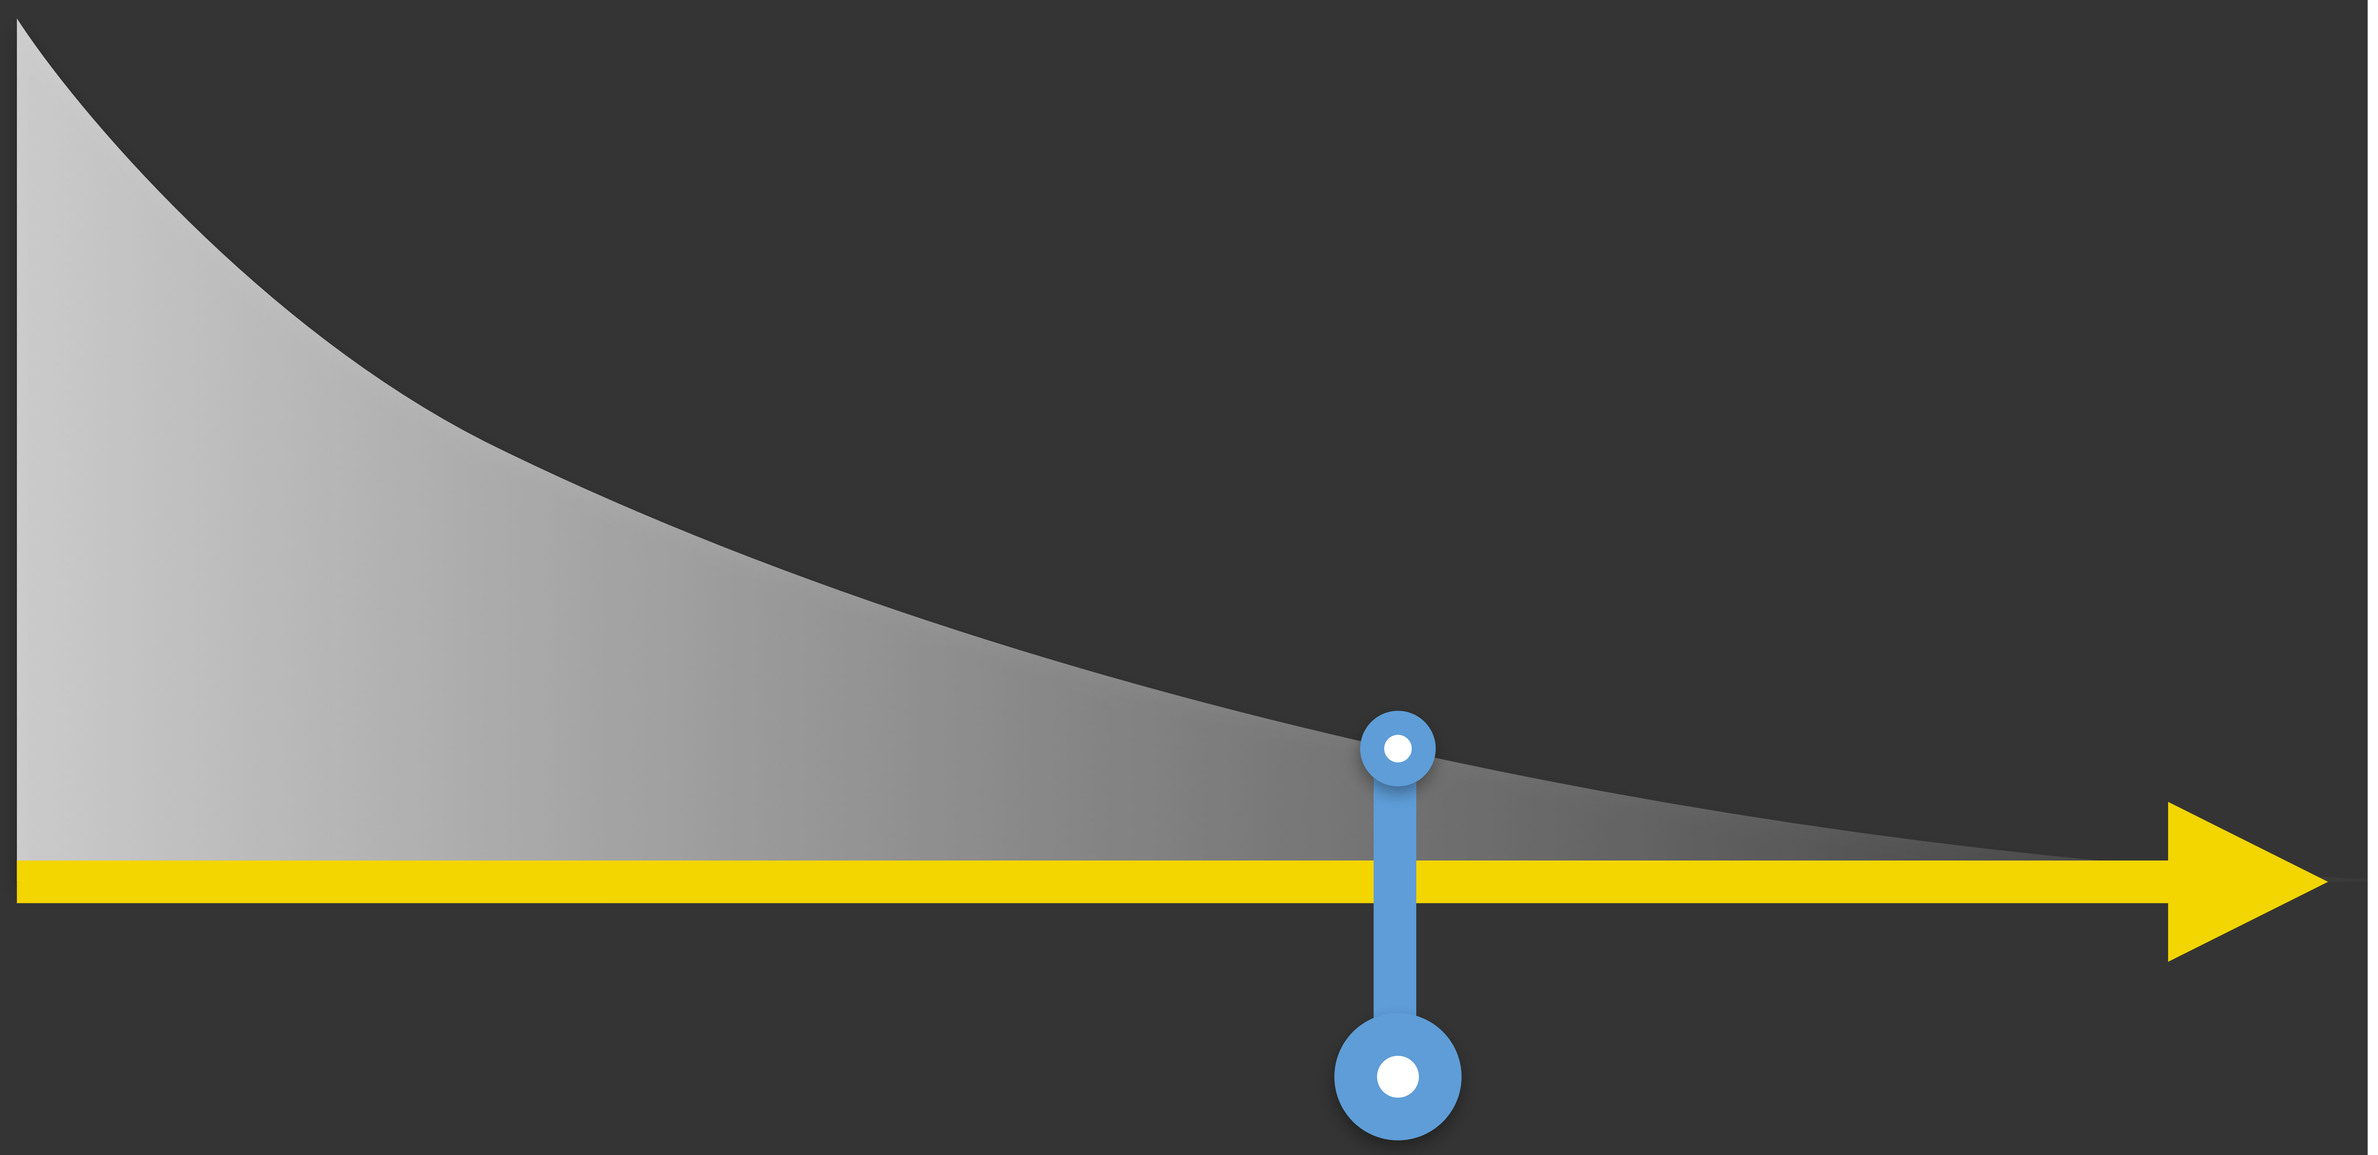

Expected Value Estimator

Photon Mapping

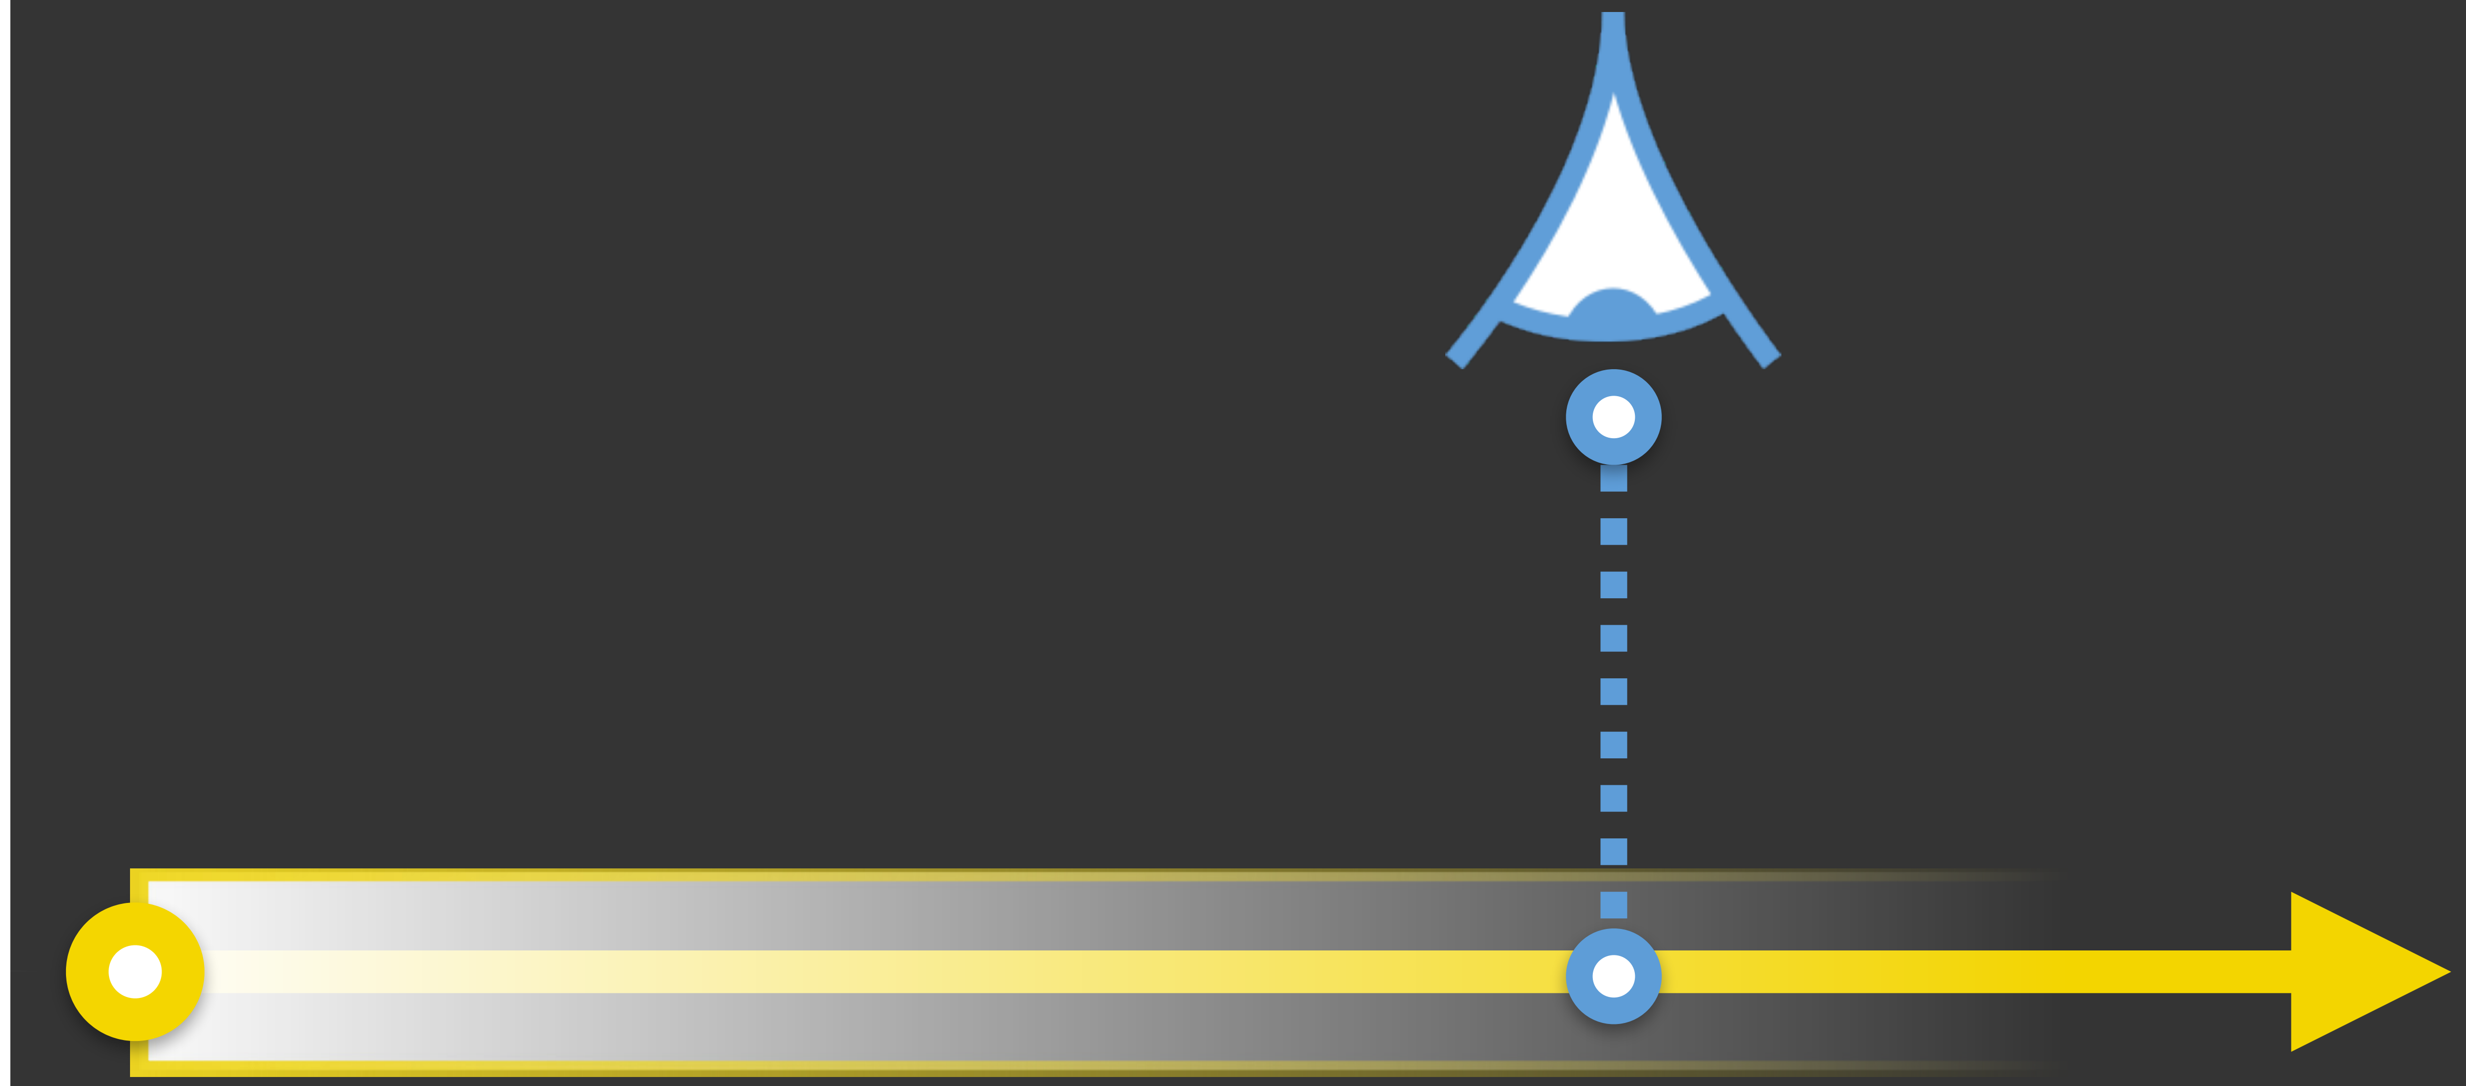

Long Beams

Neutron Transport

Photon Mapping

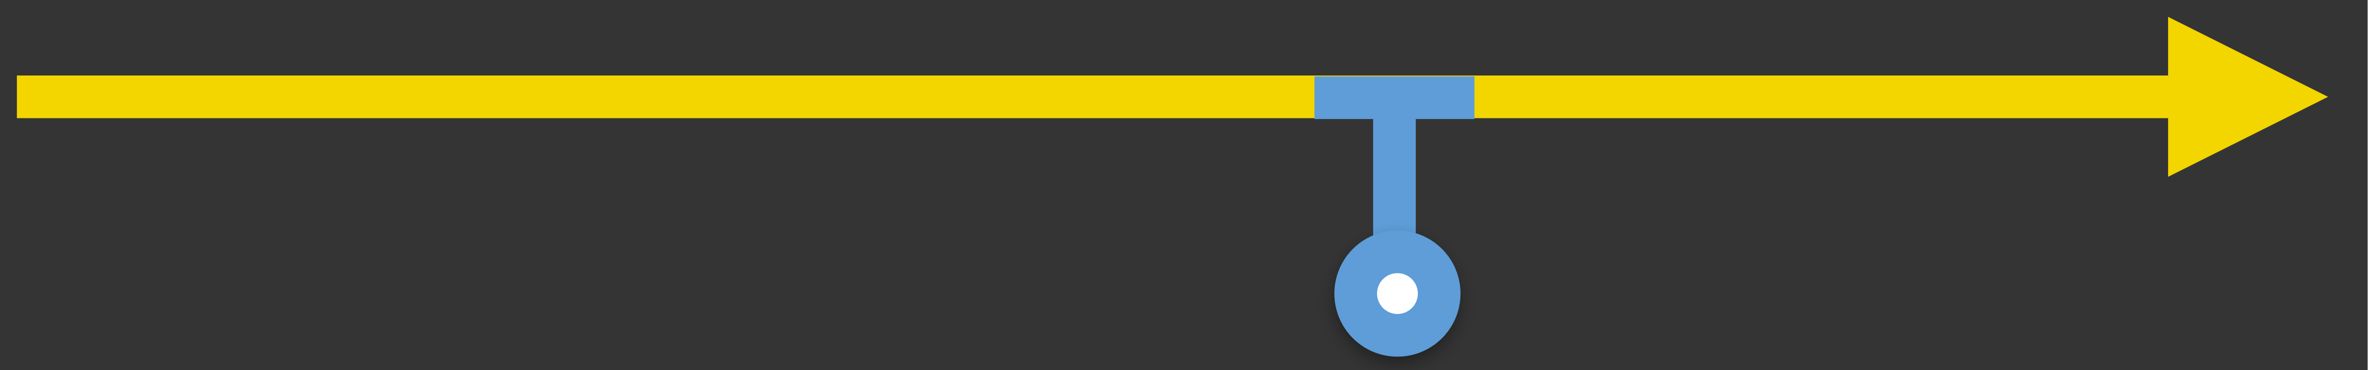

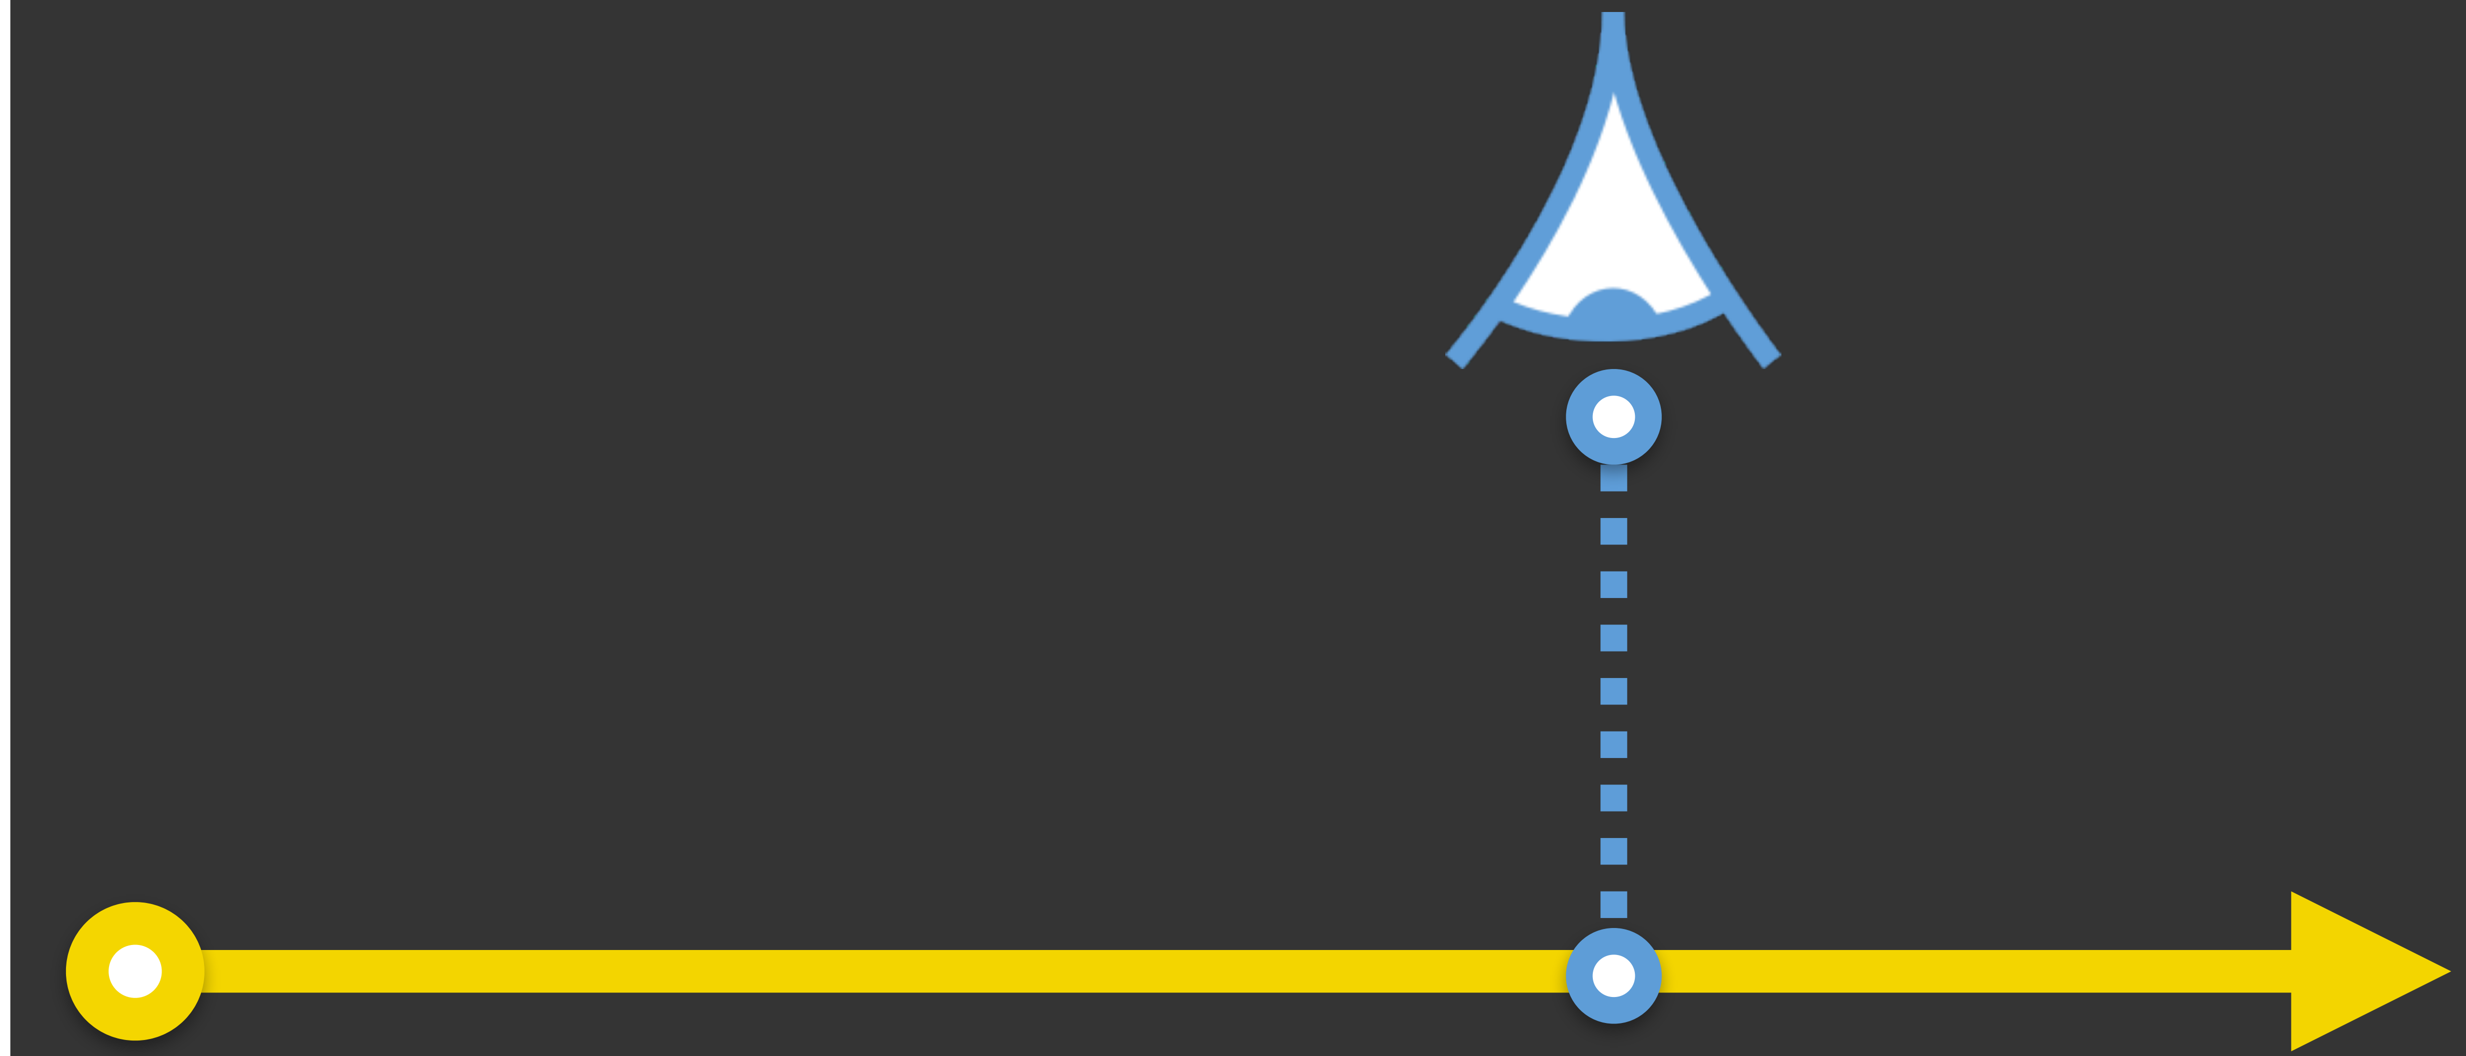

Track-Length Estimator

Short Beams

Neutron Transport

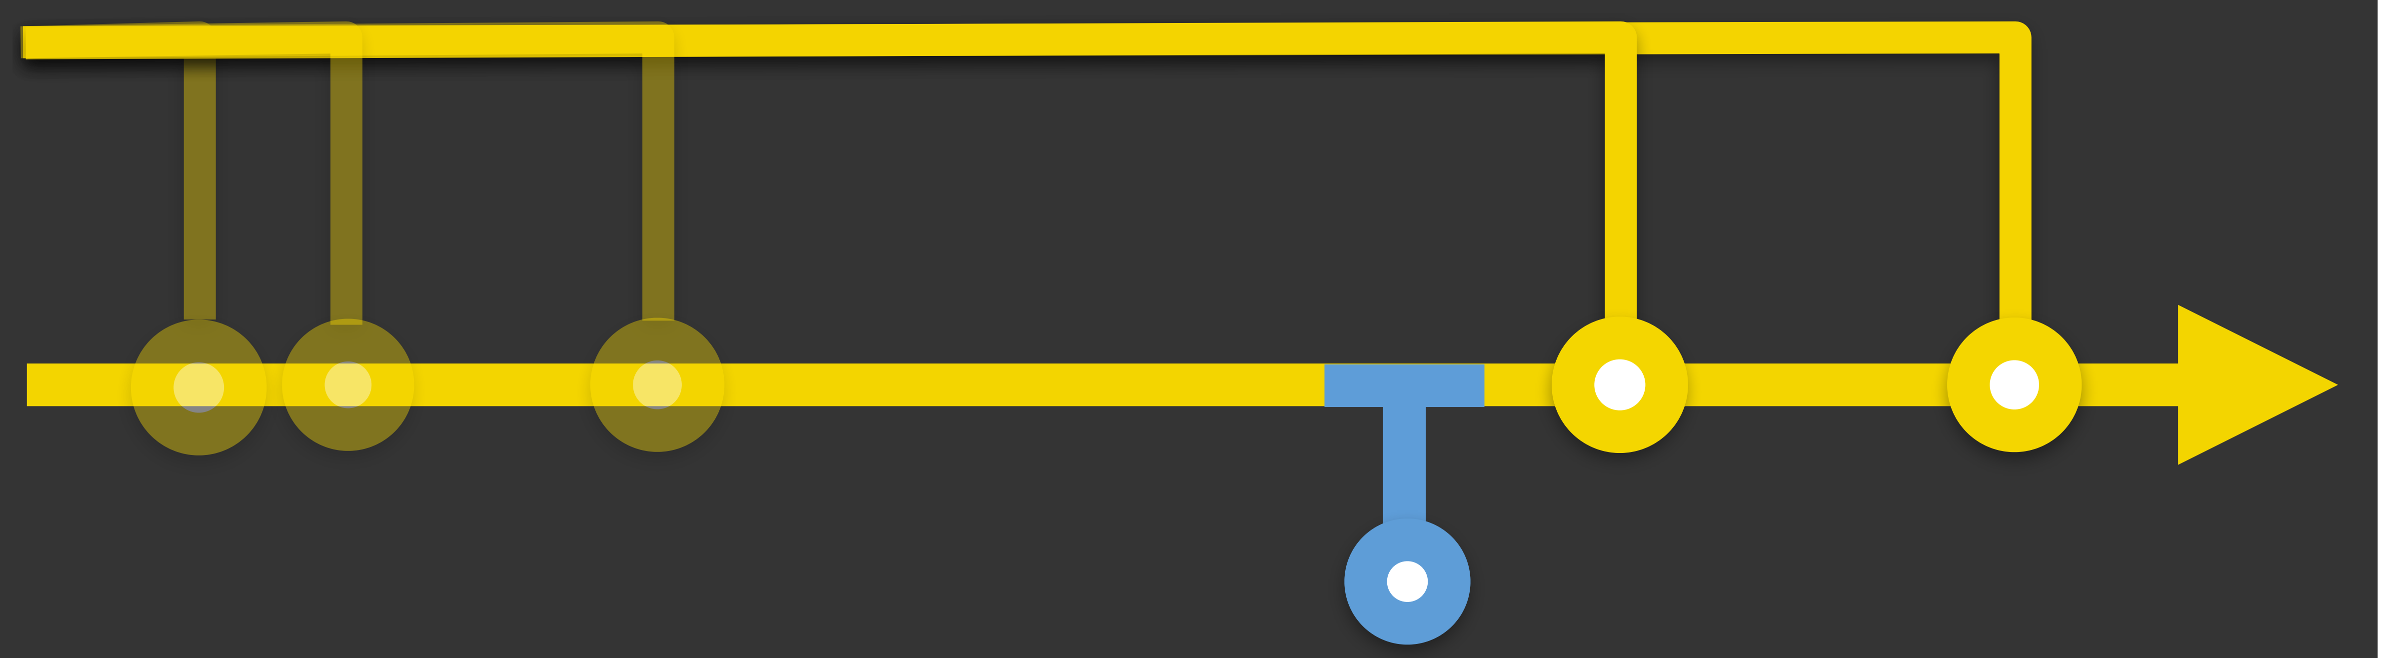

Track-Length Estimator

Photon Mapping

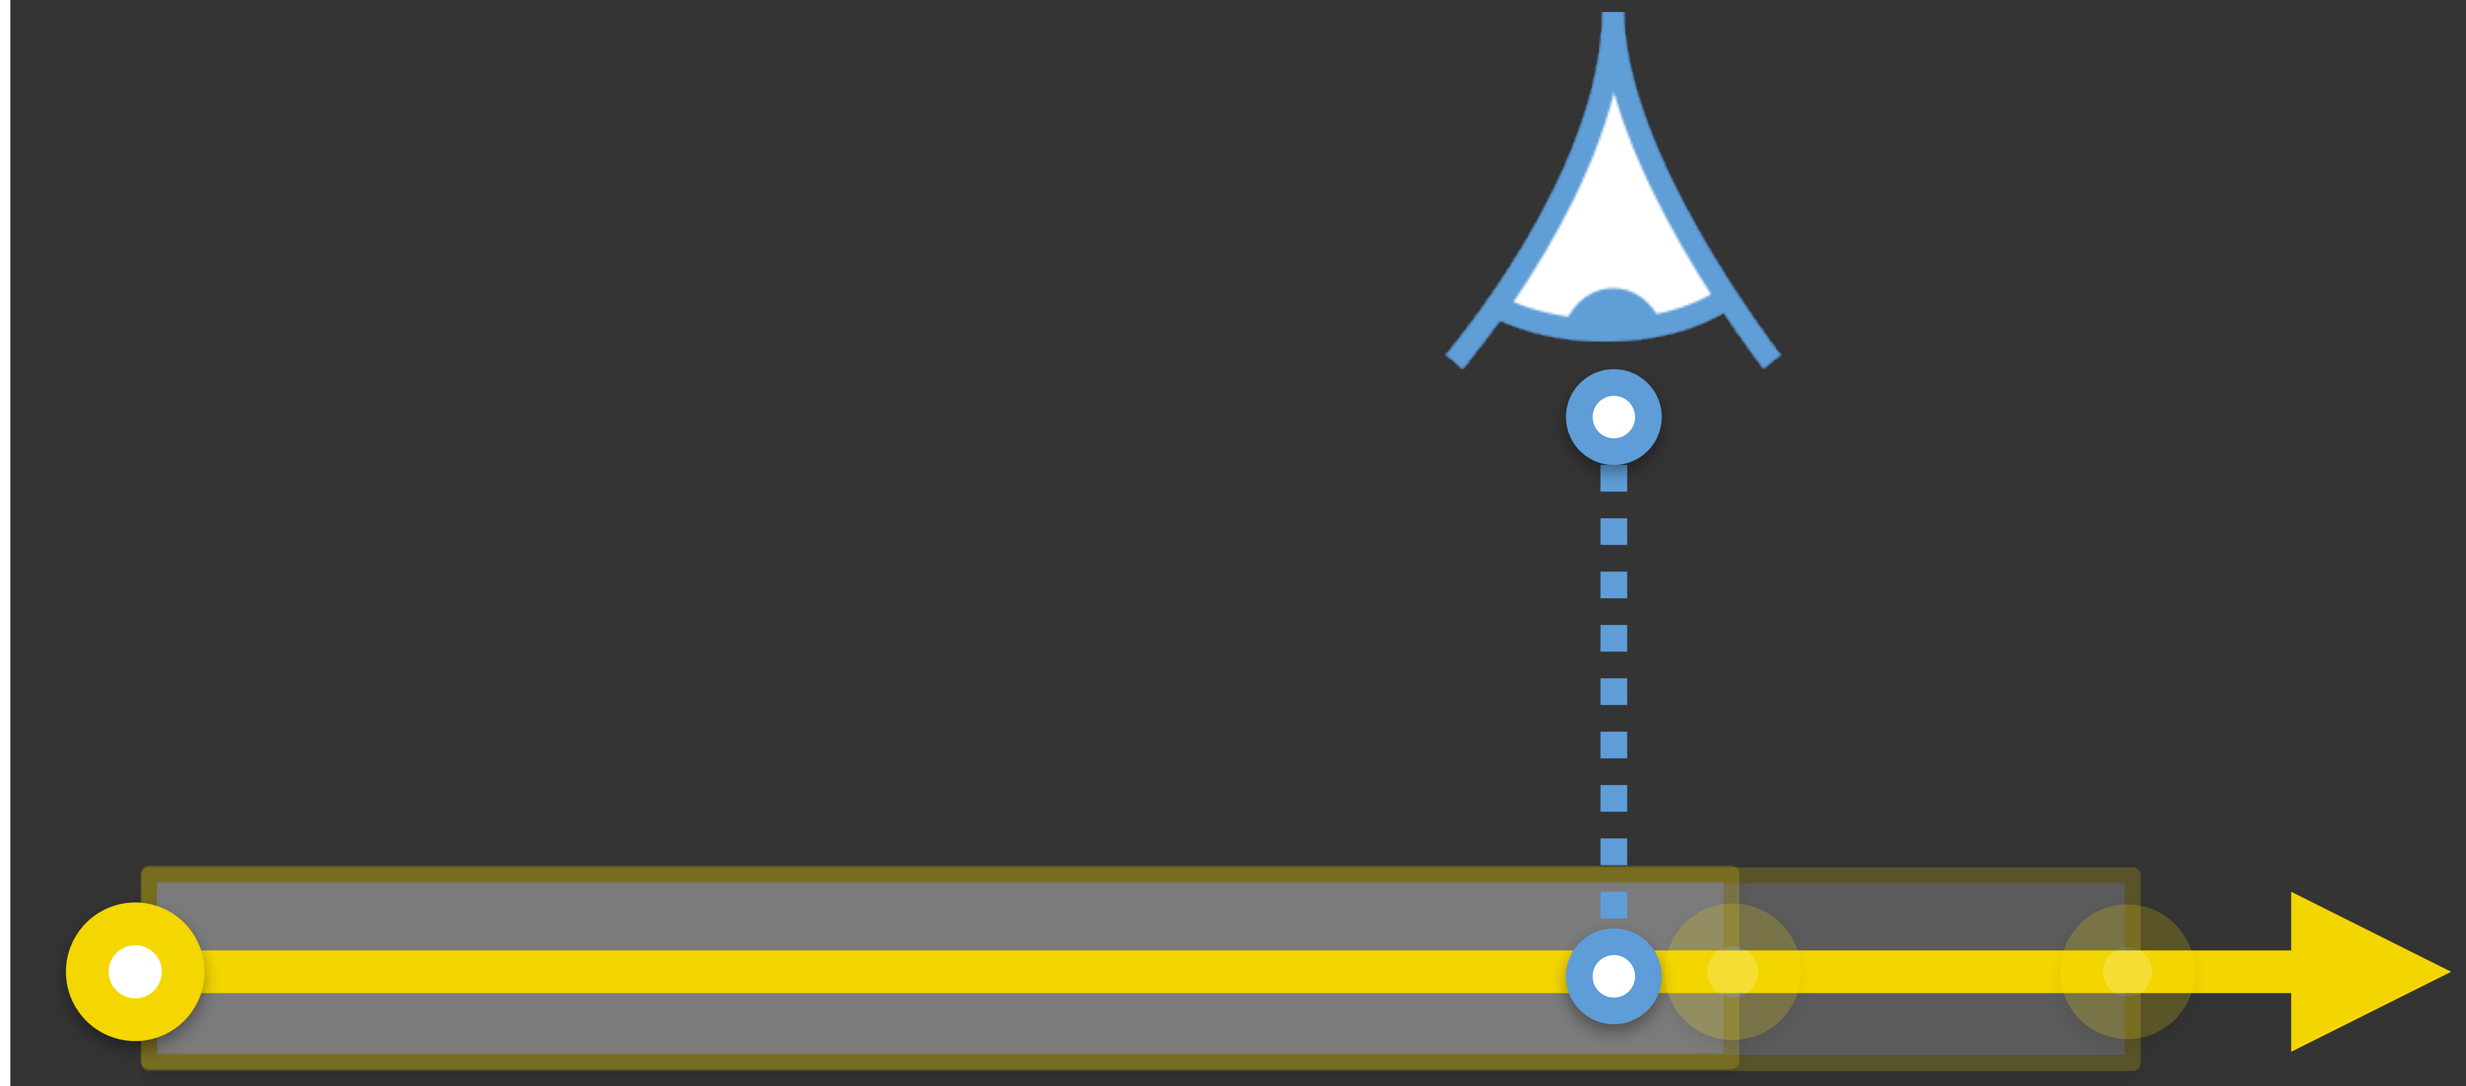

Short Beams

Beyond Points and Beams

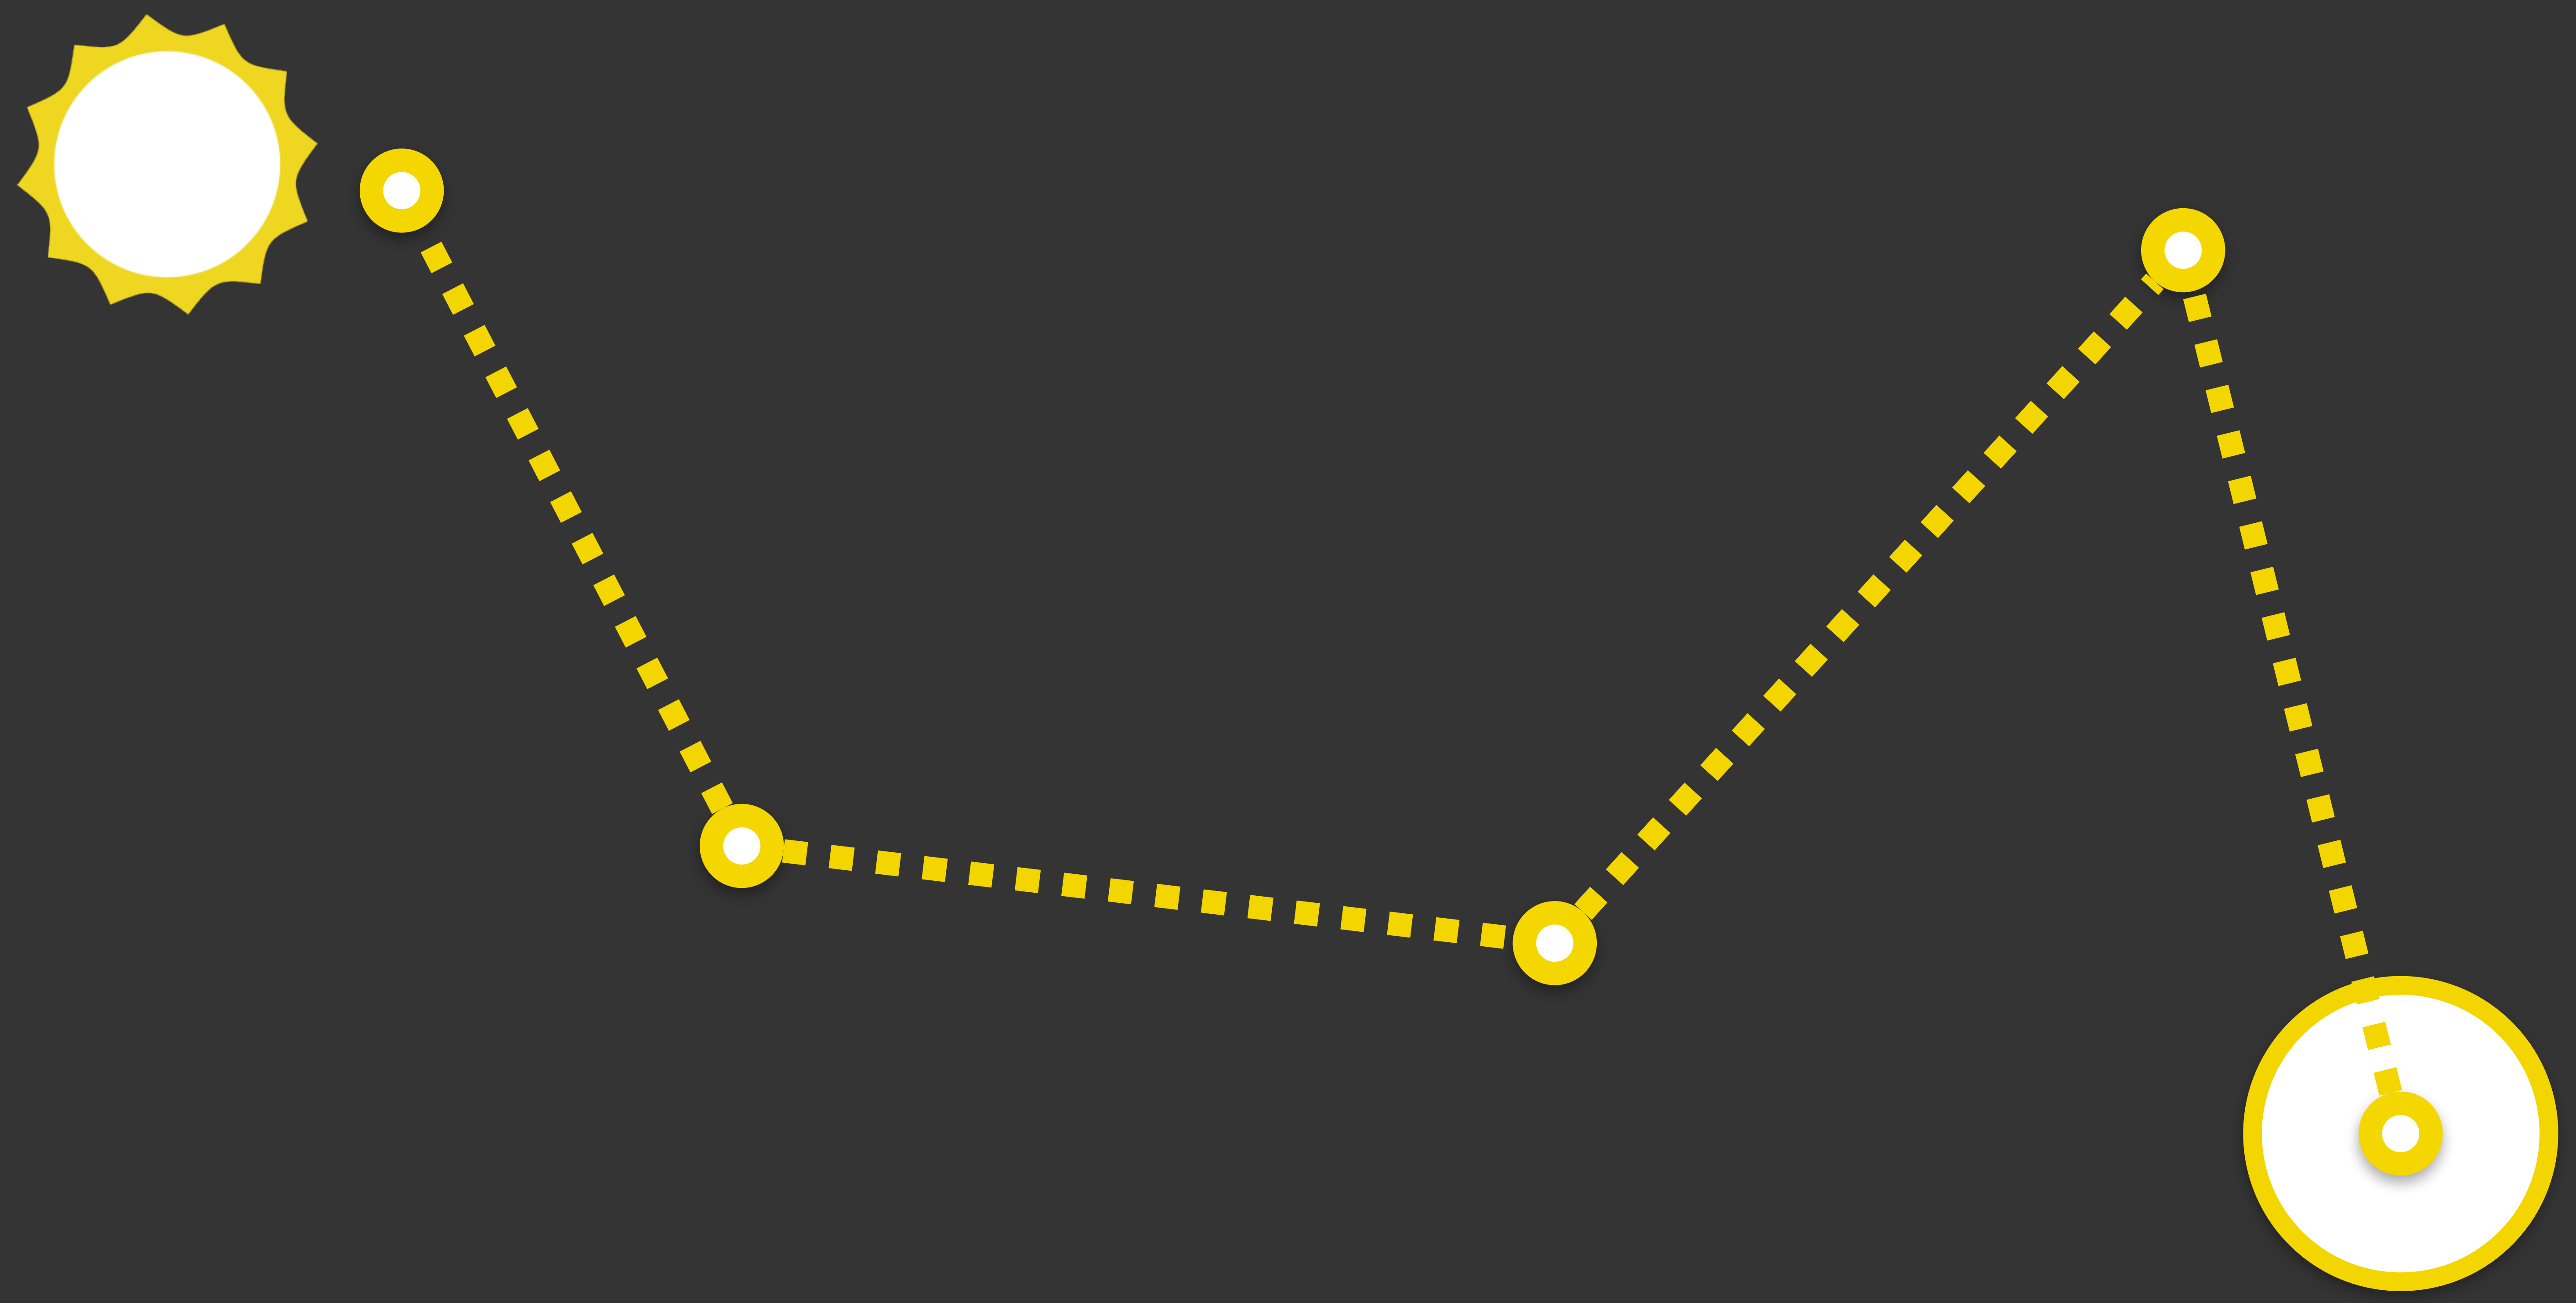

Beyond Points and Beams

Beyond Points and Beams

- “Marching”: Replace one collision estimator with…

Beyond Points and Beams

- "Marching": Replace one collision estimator with…

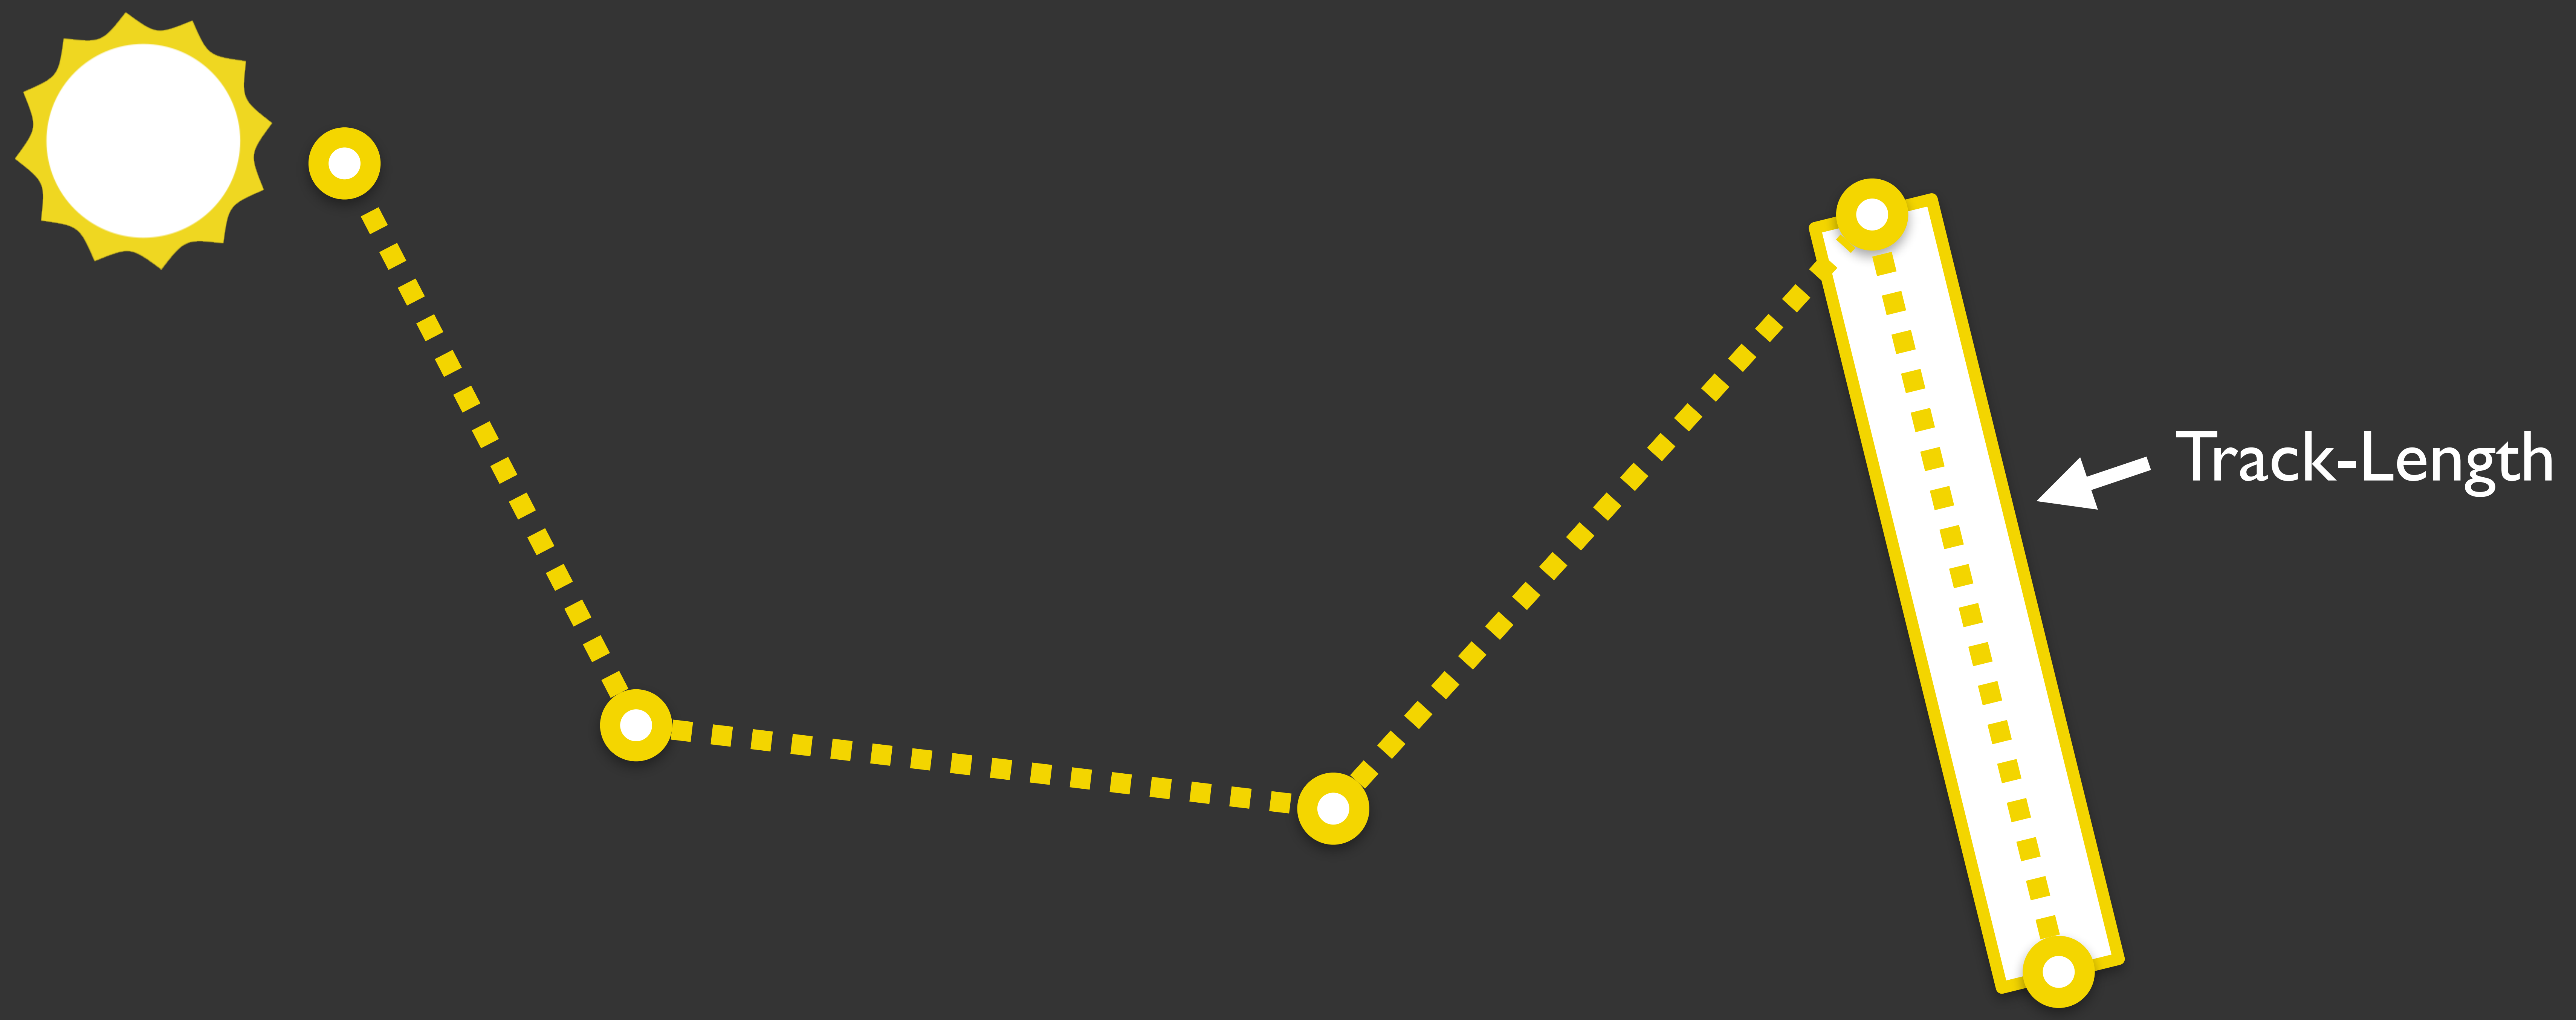

Beyond Points and Beams

- “Marching”: Replace one collision estimator with…

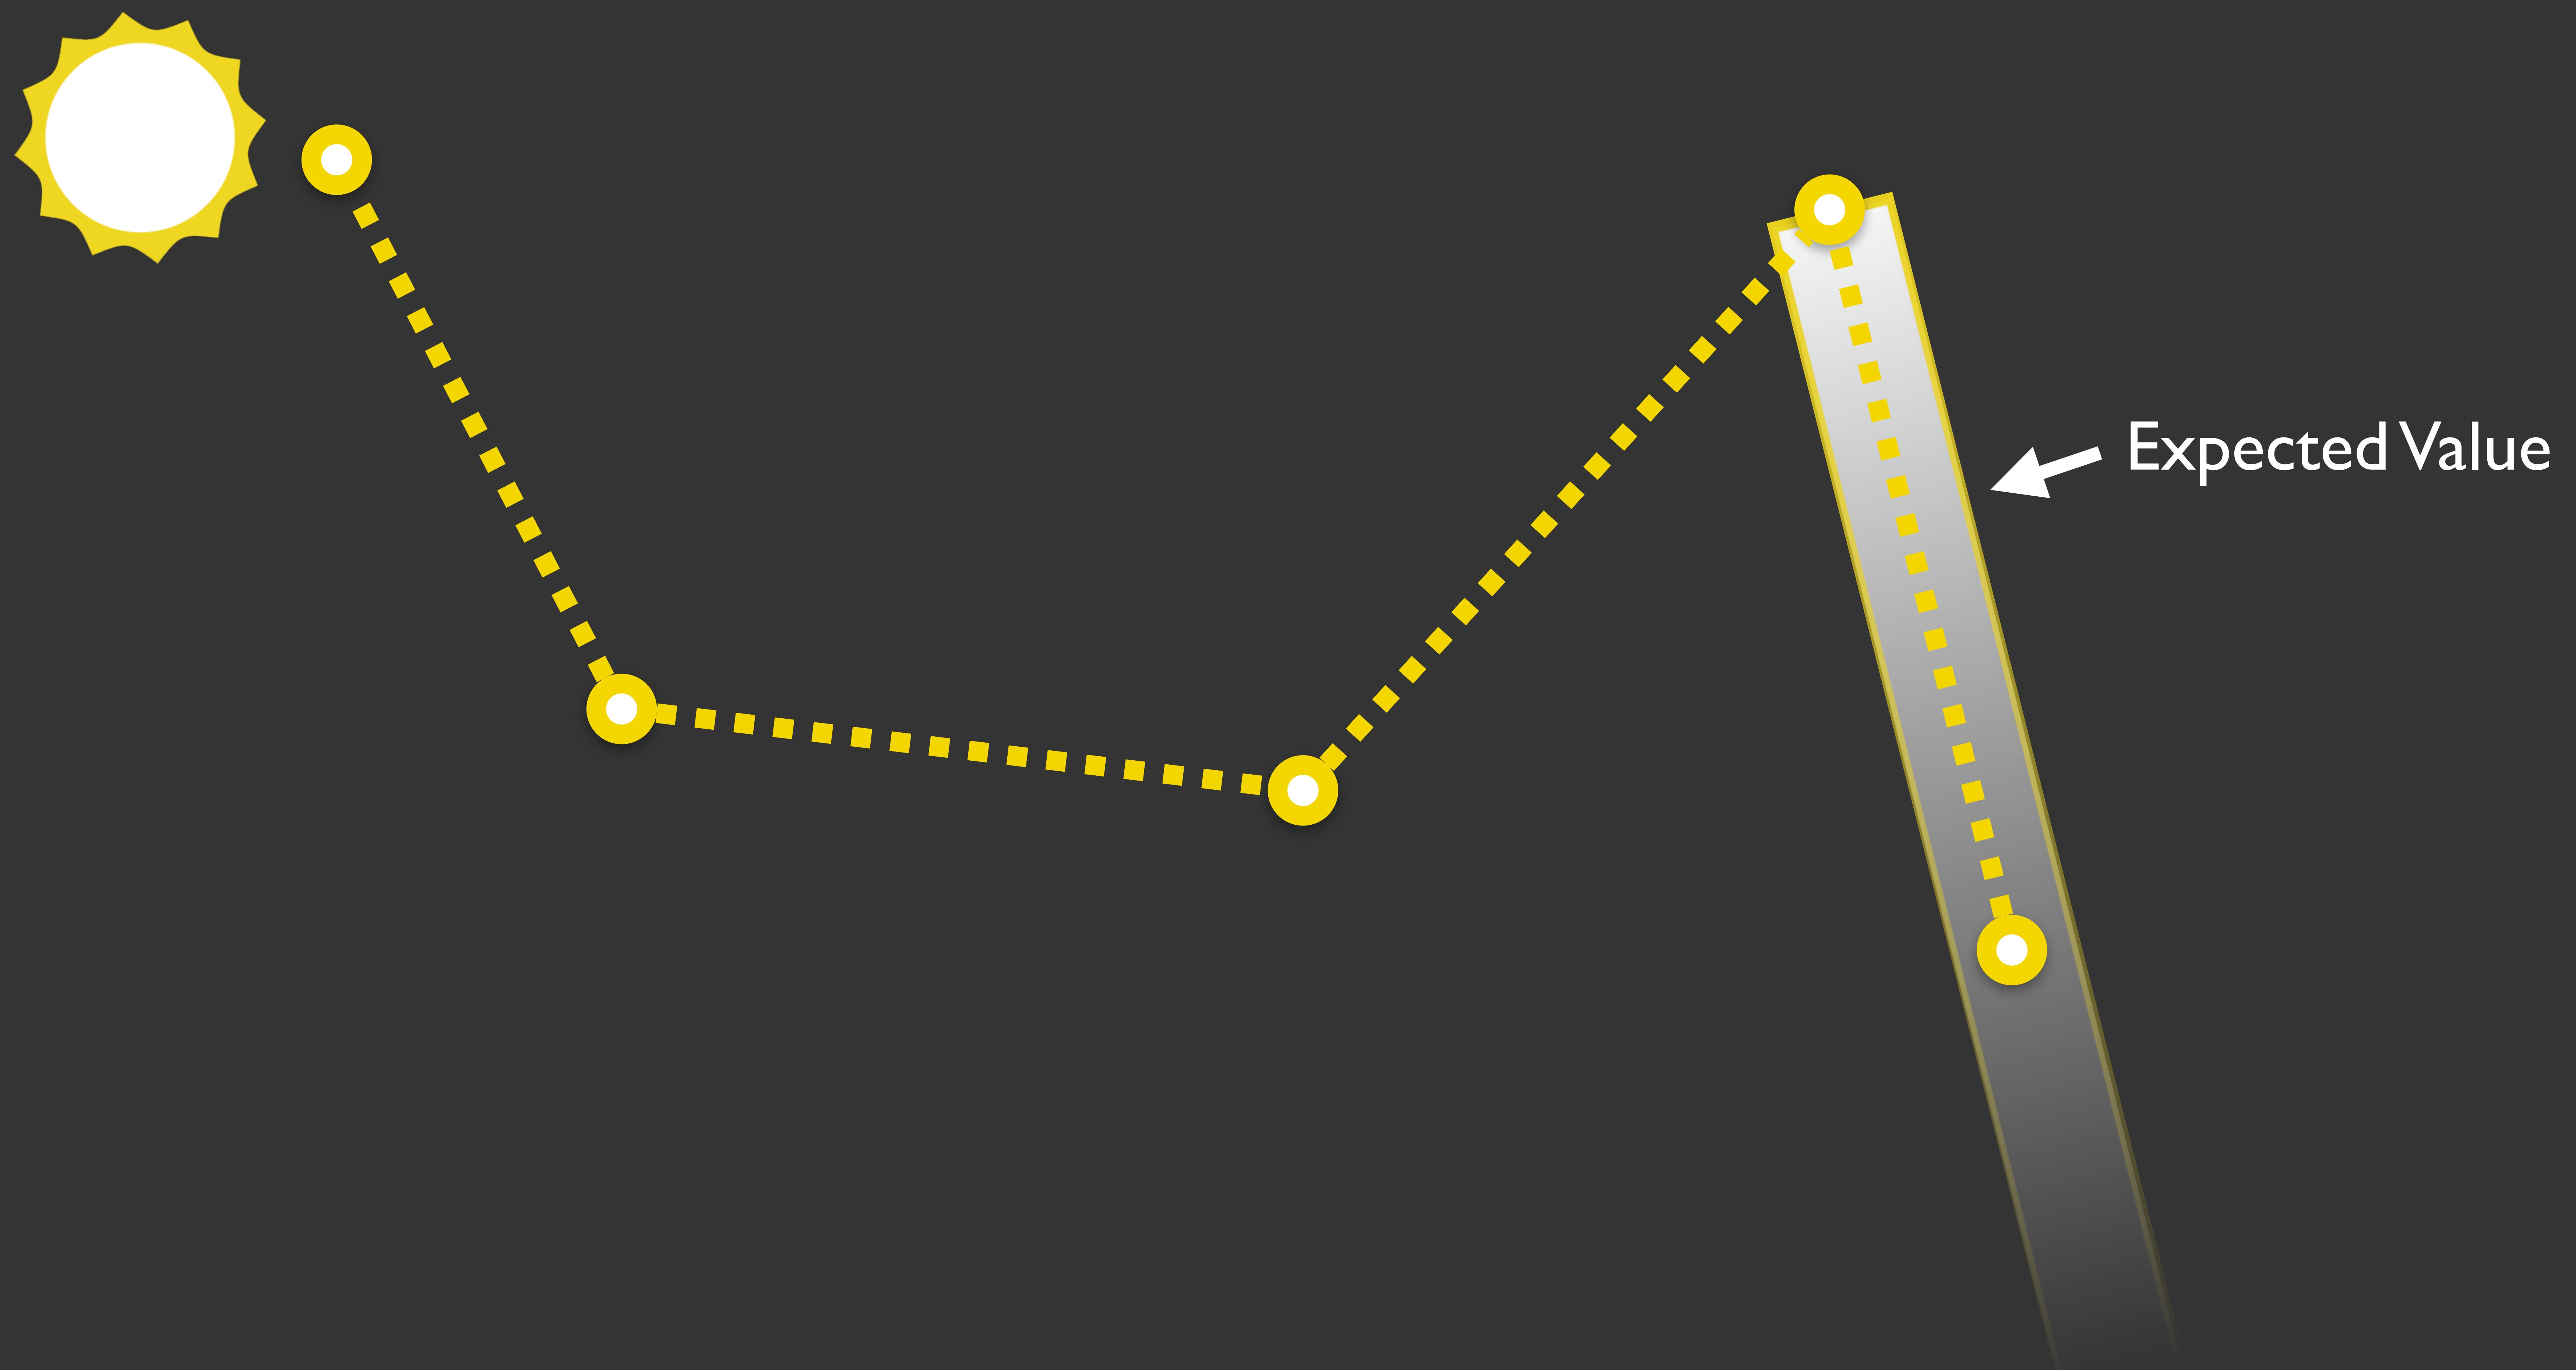

Beyond Points and Beams

- “Marching”: Replace one collision estimator with…

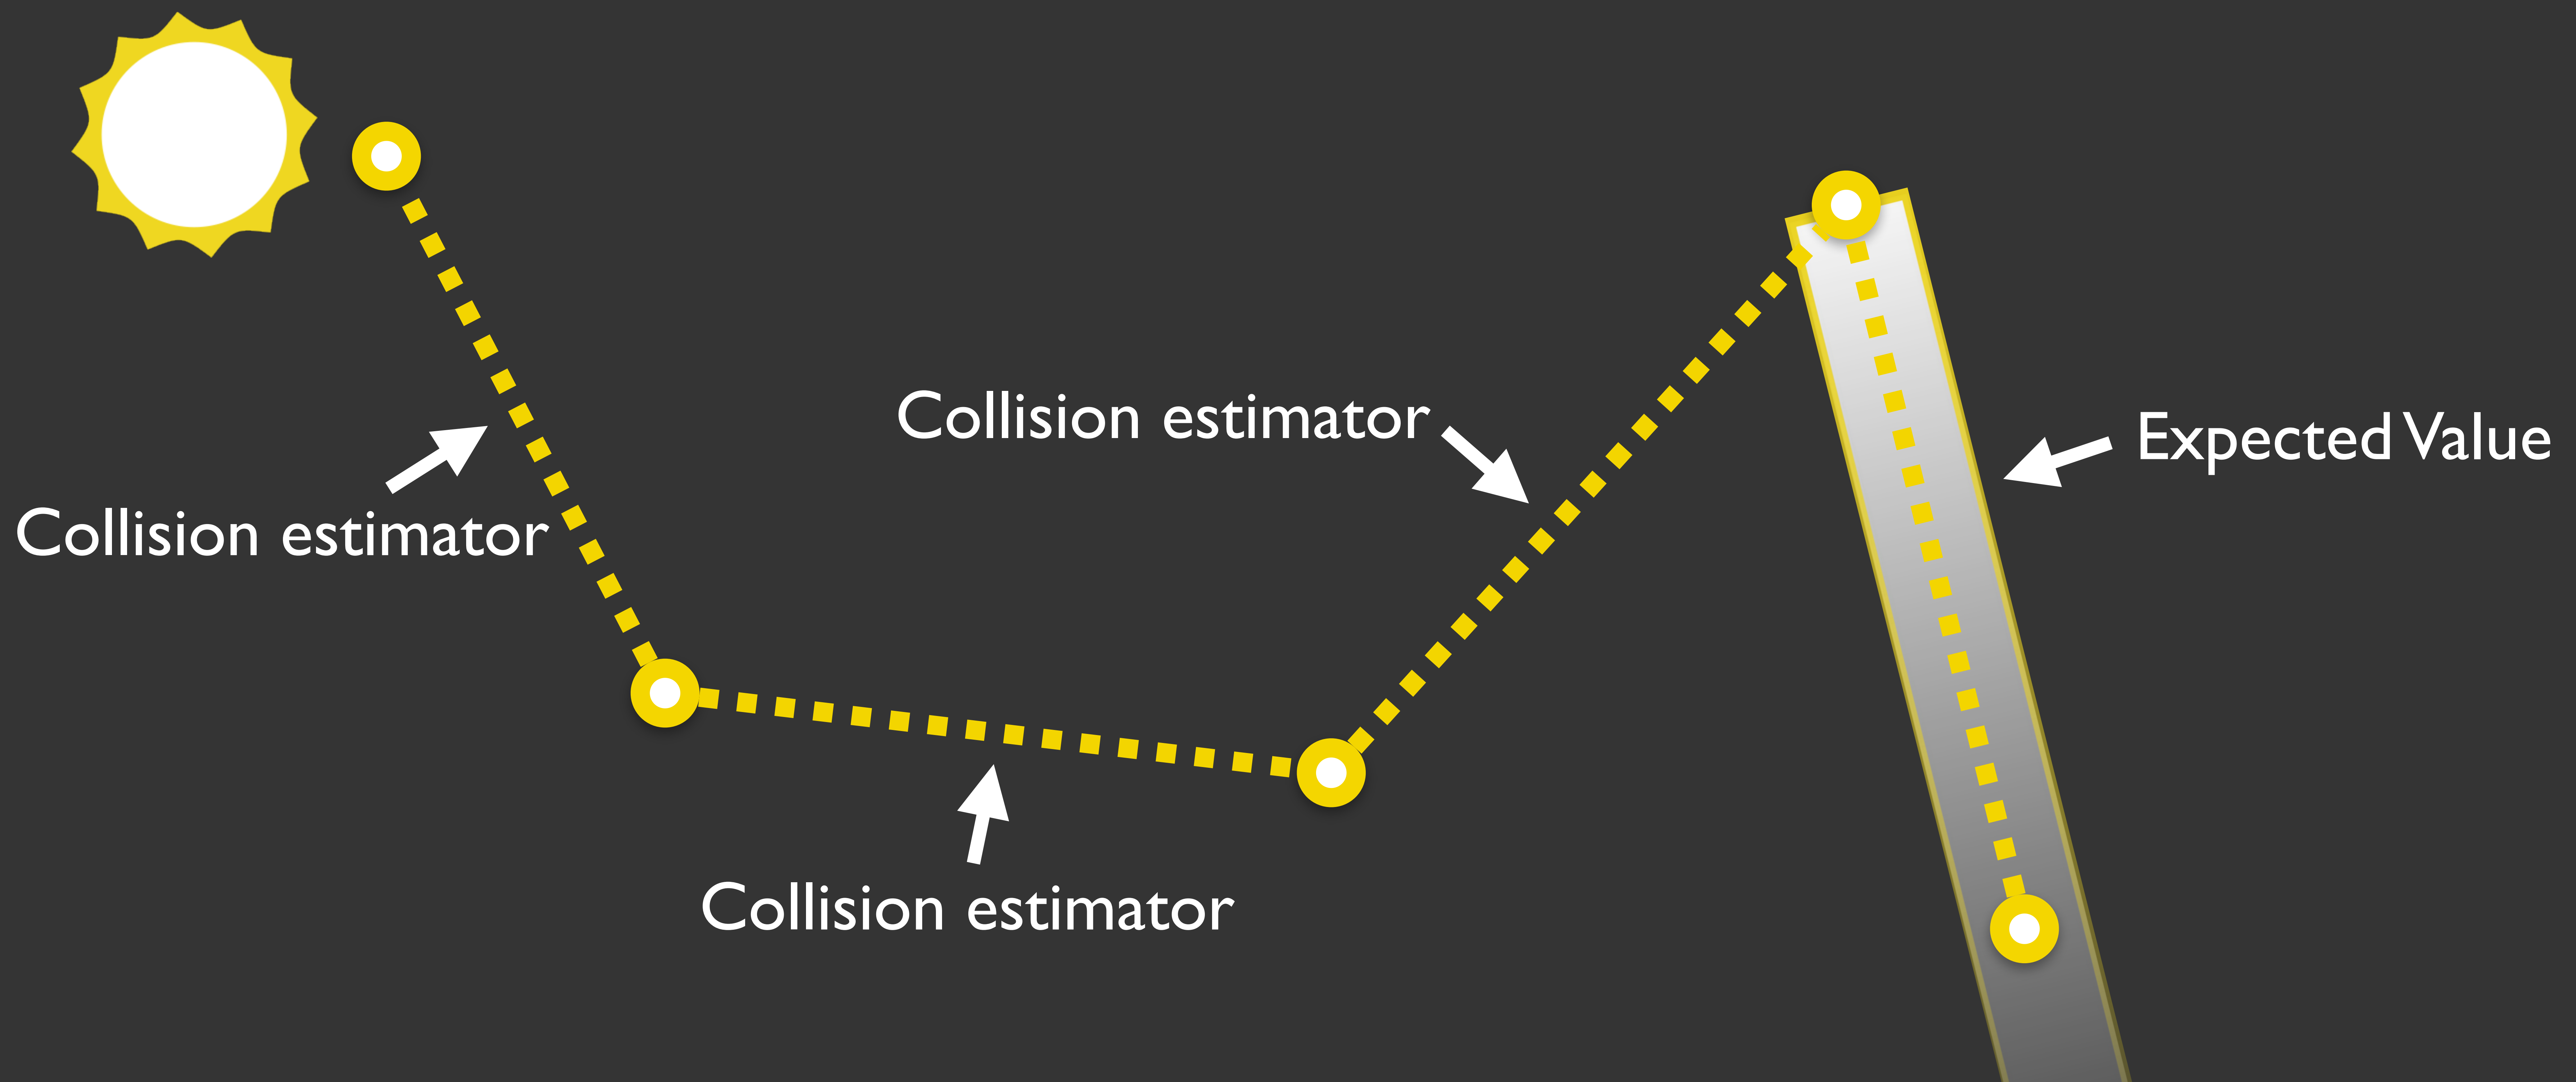

- But: Collision estimator on *every segment!*

Photon Beams

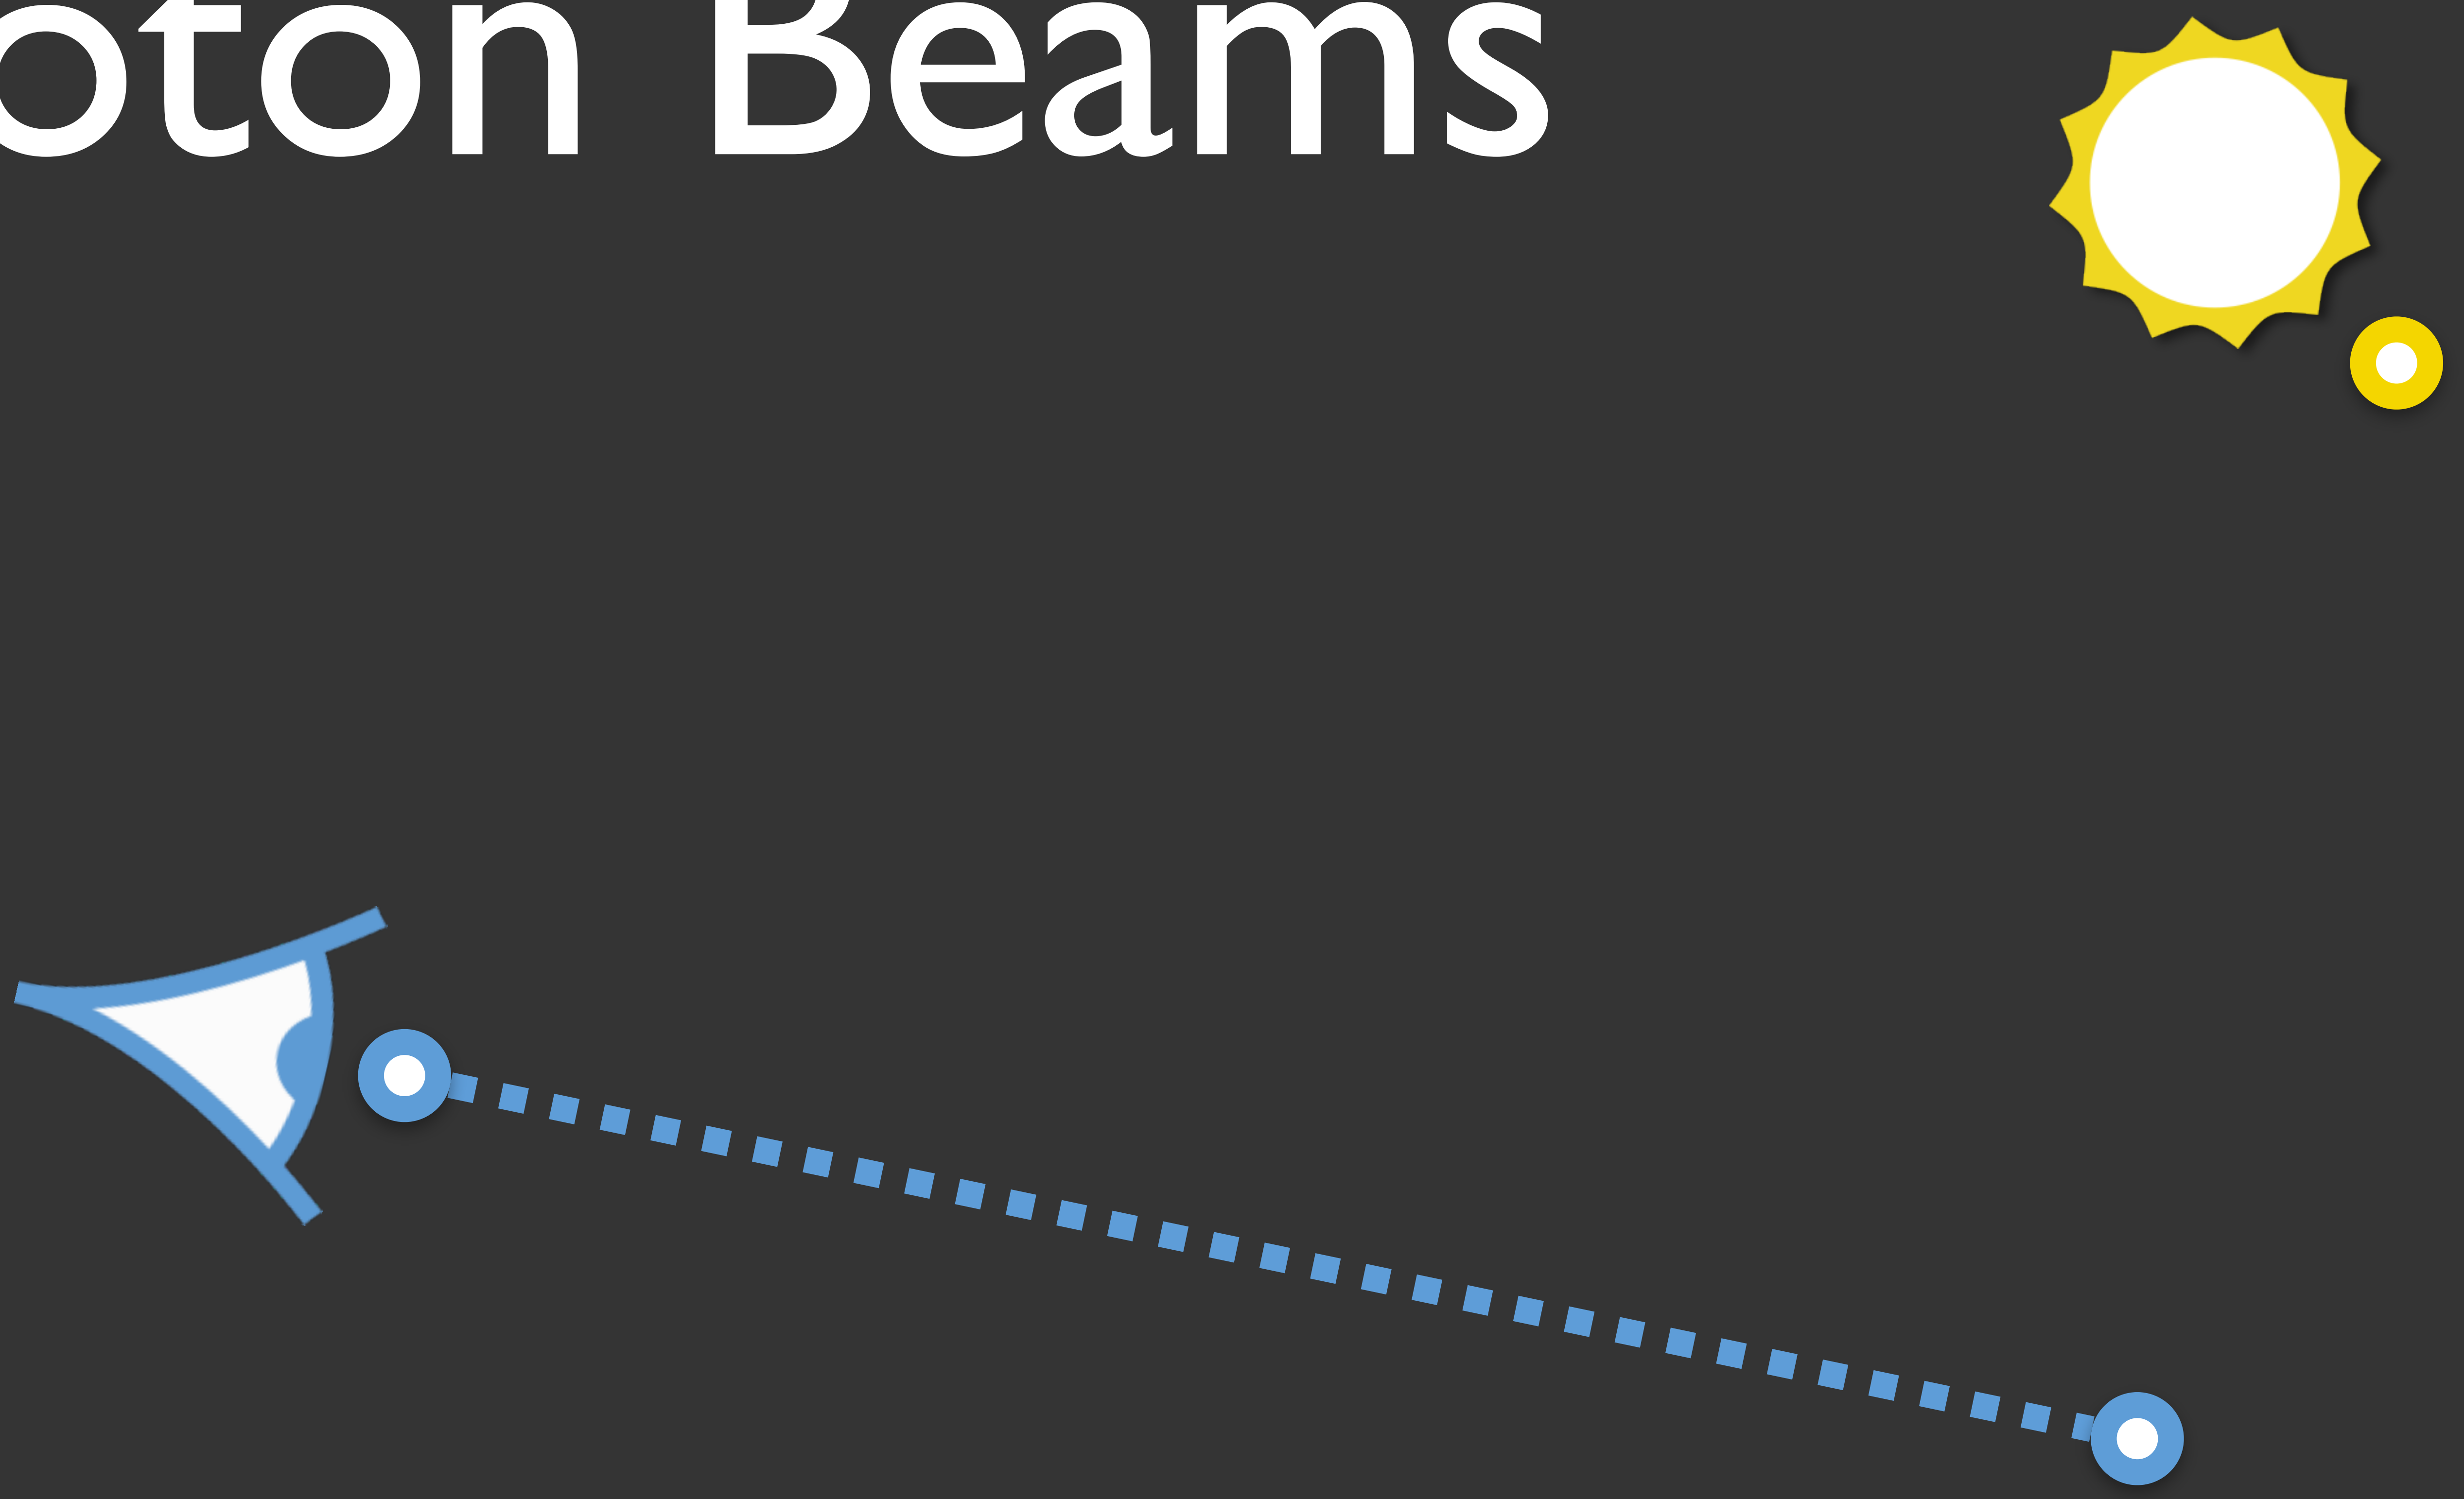

Photon Beams

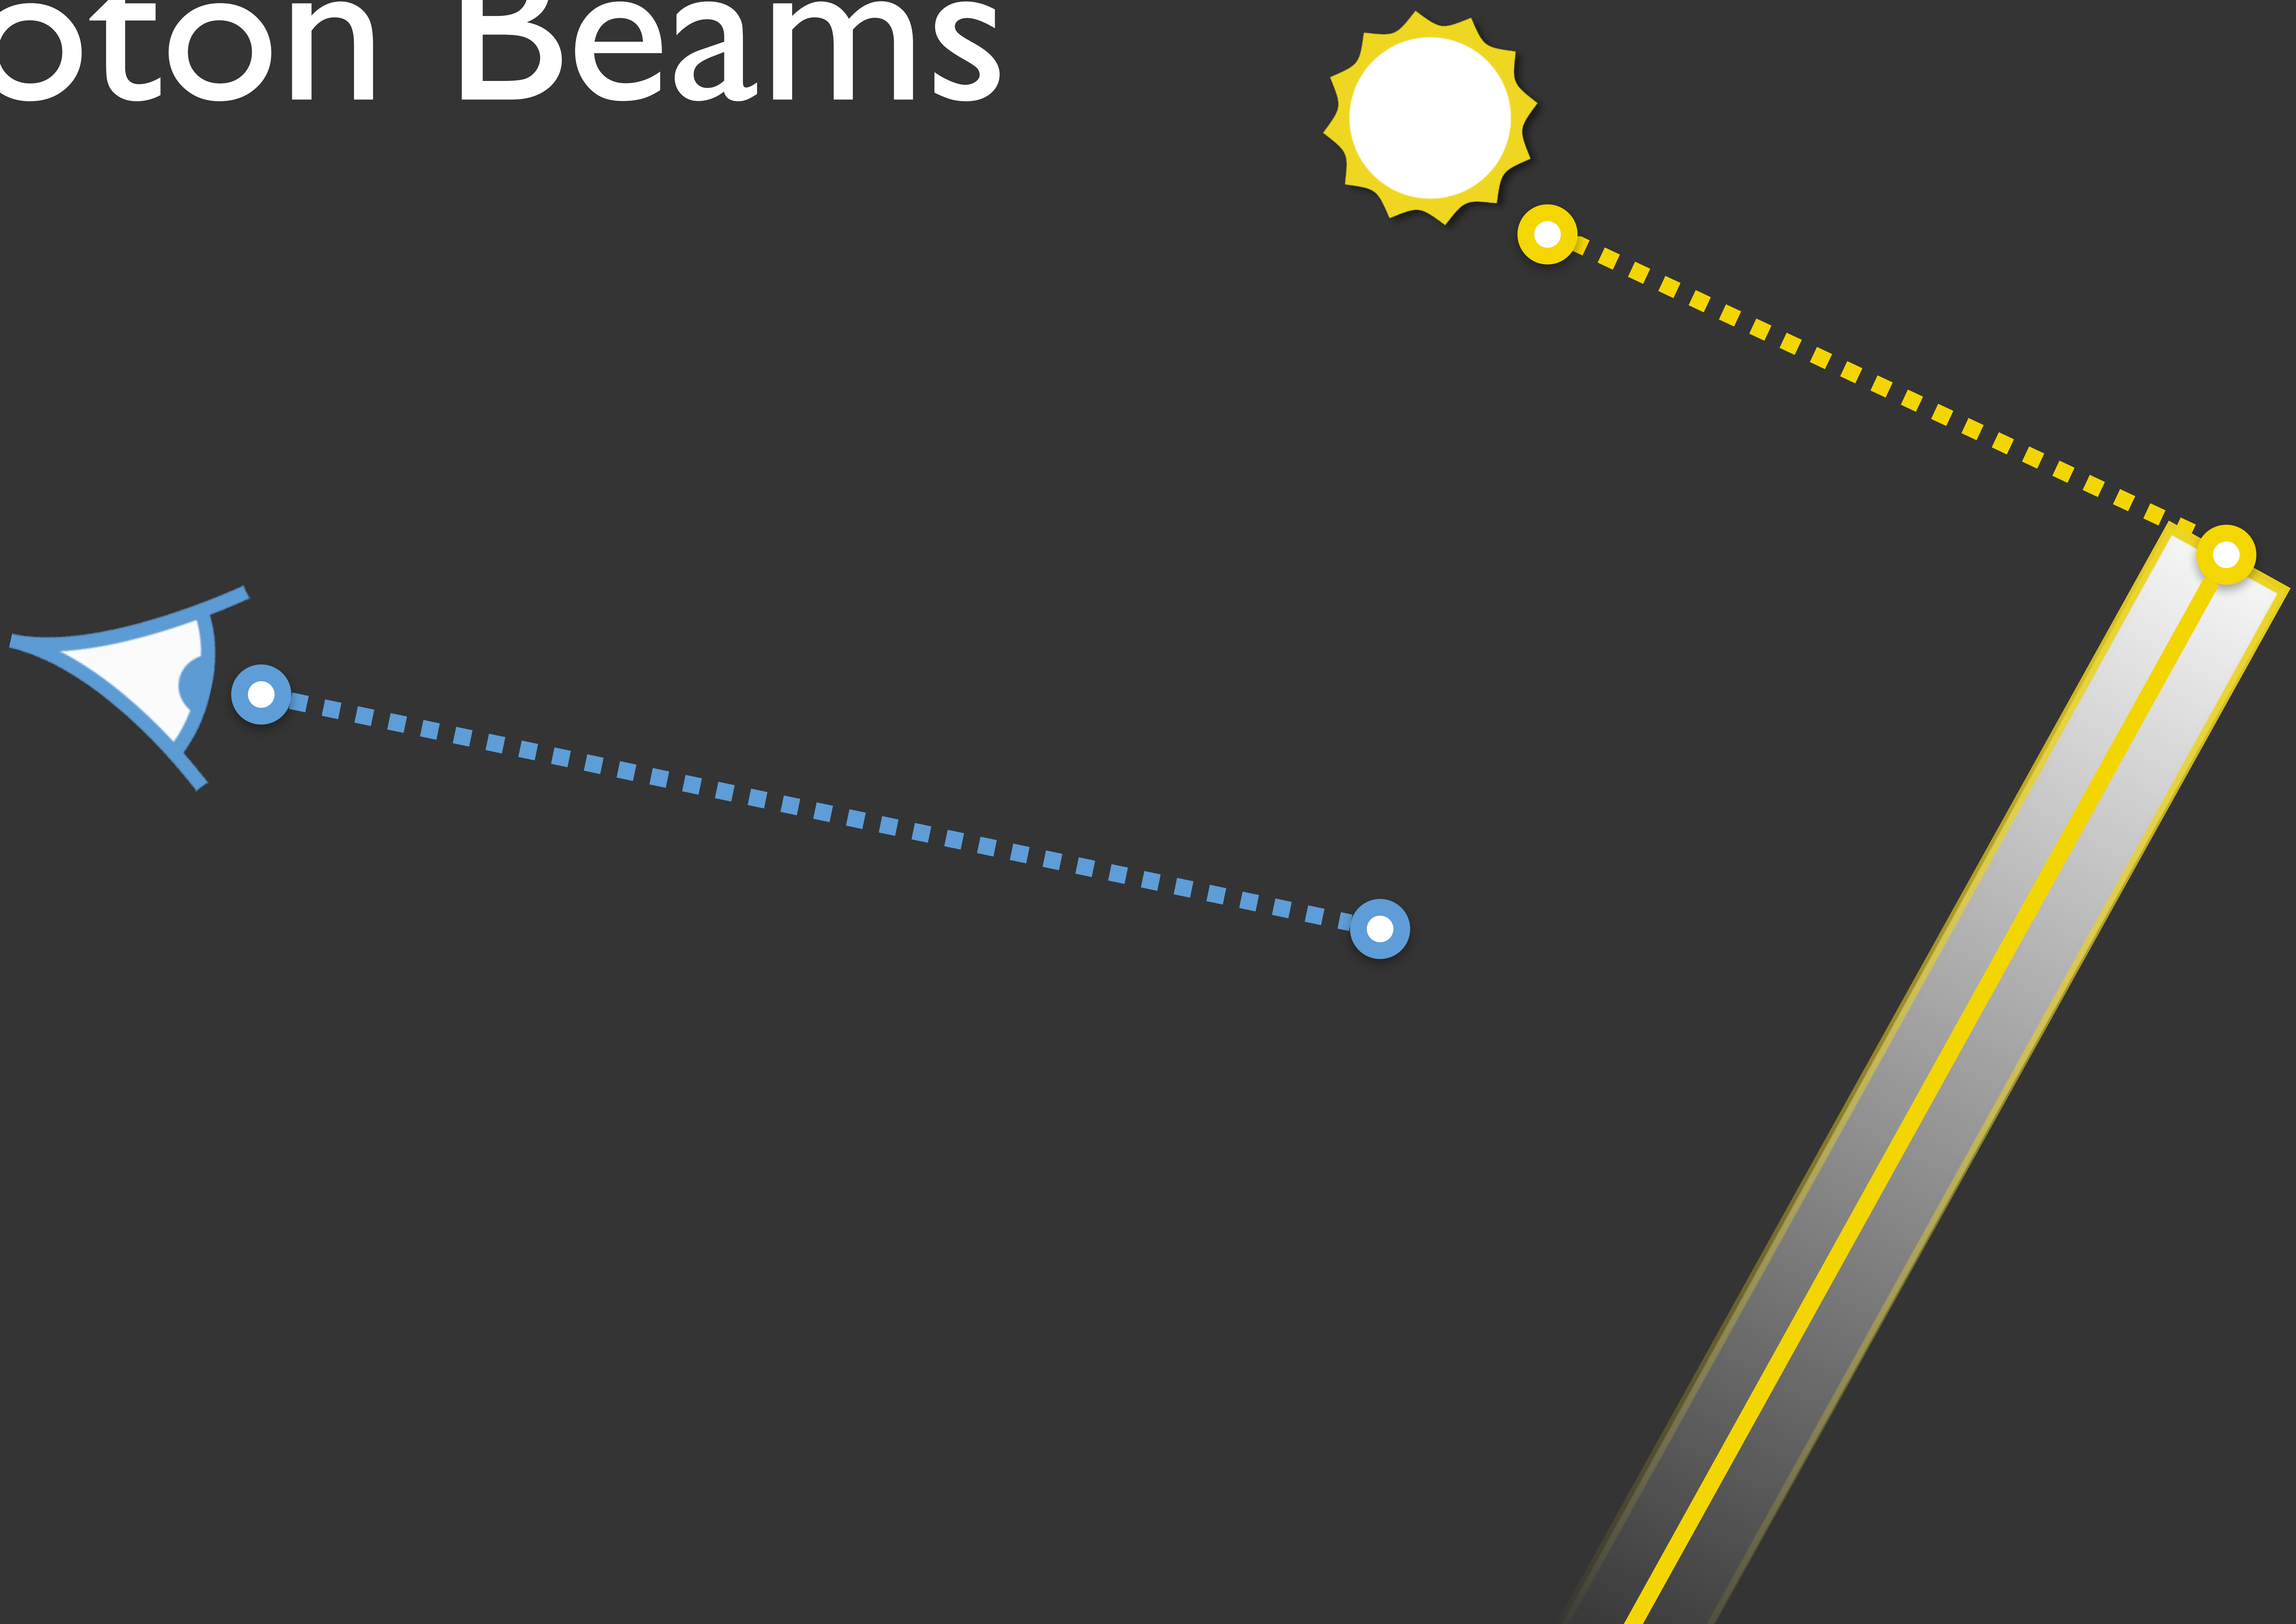

Photon Beams

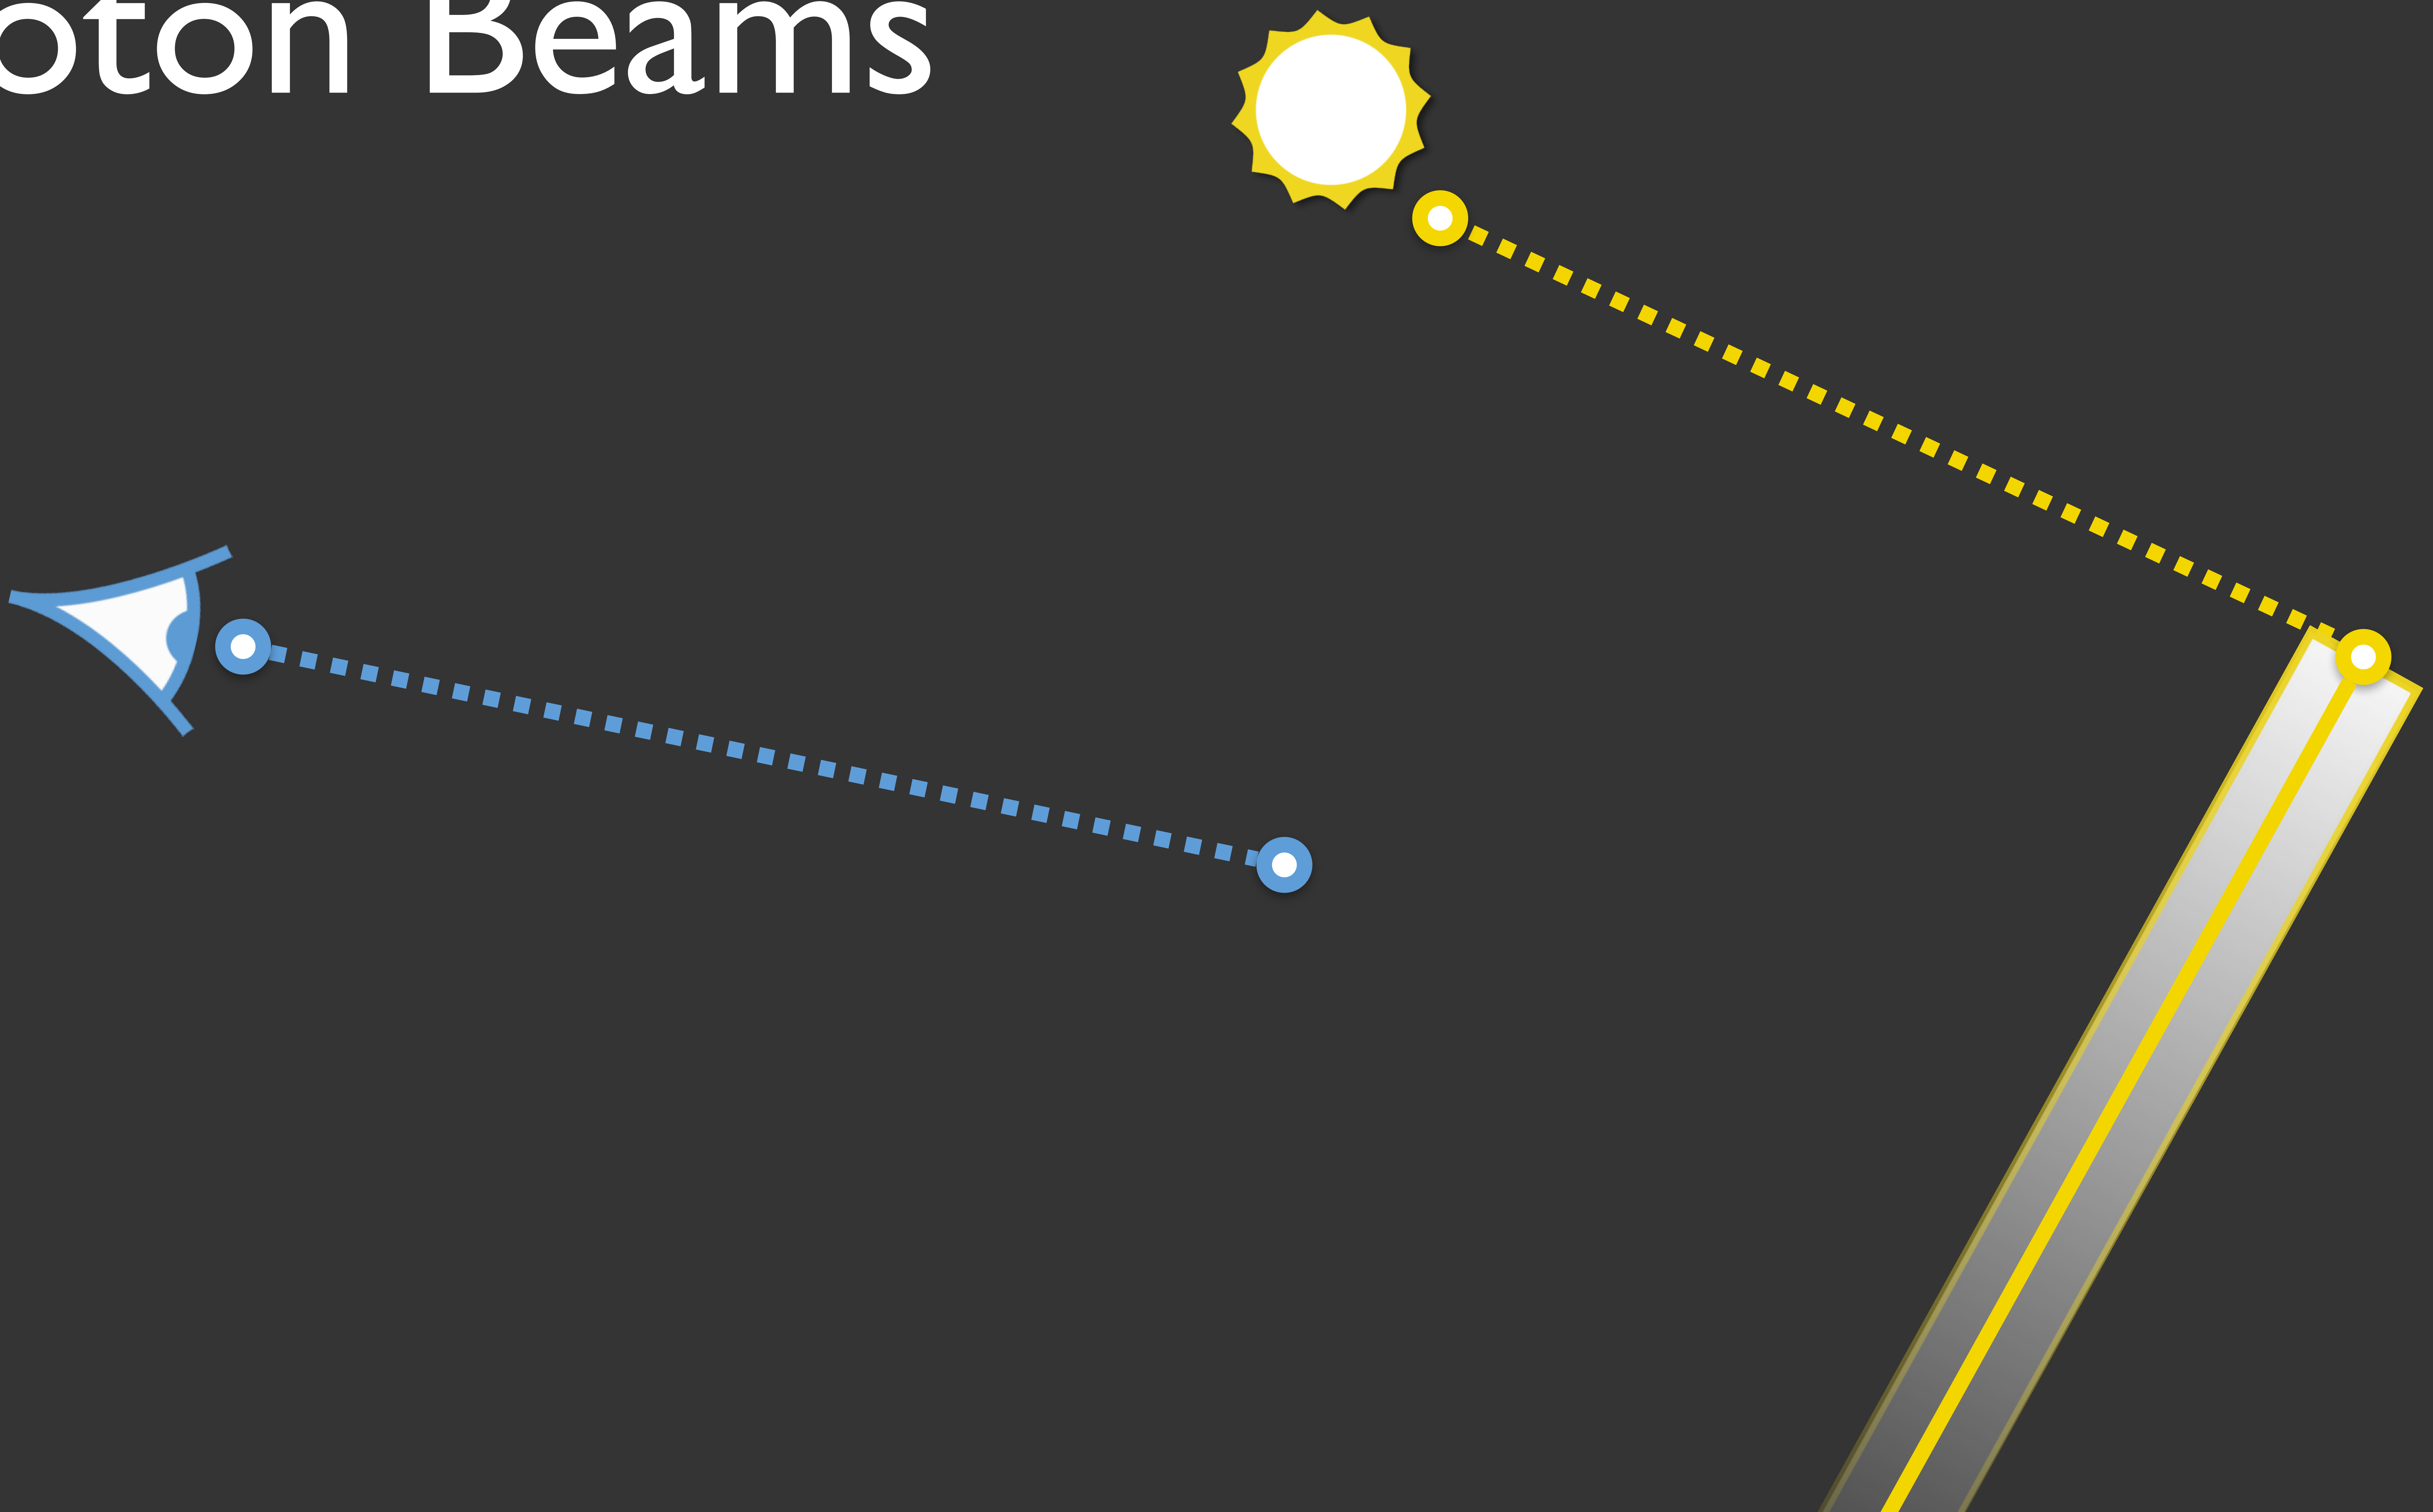

Photon Beams

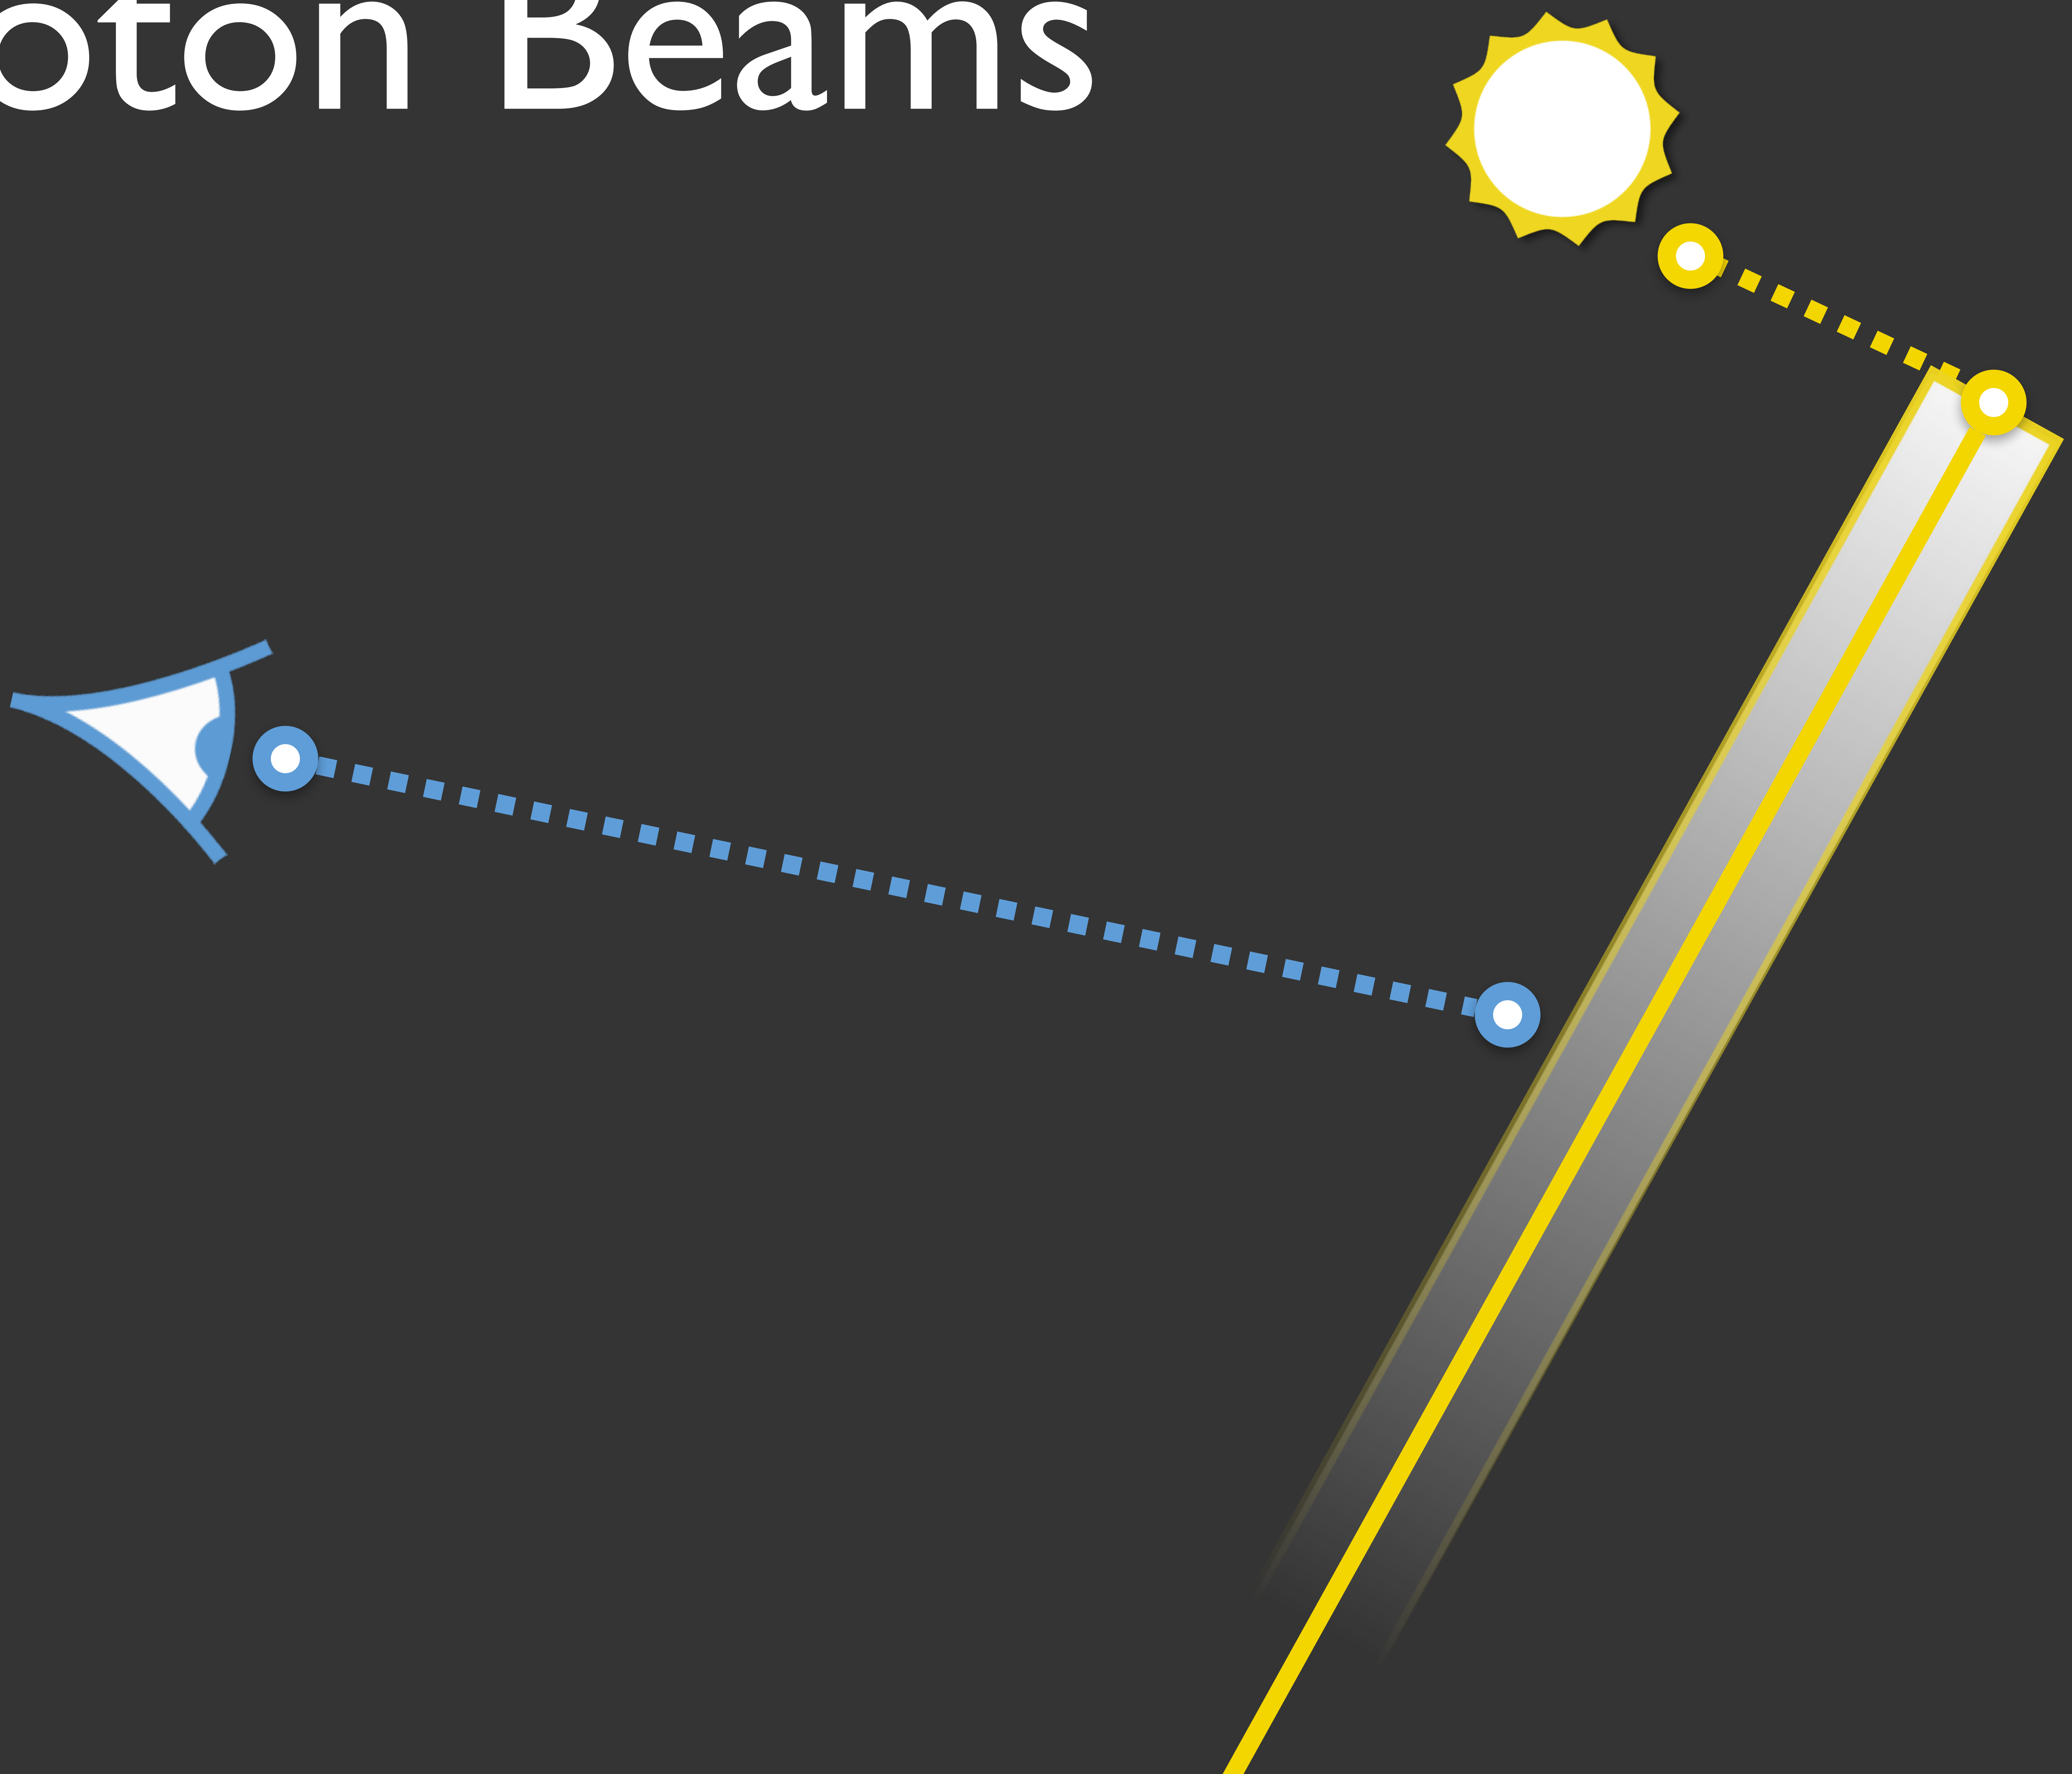

Photon Beams

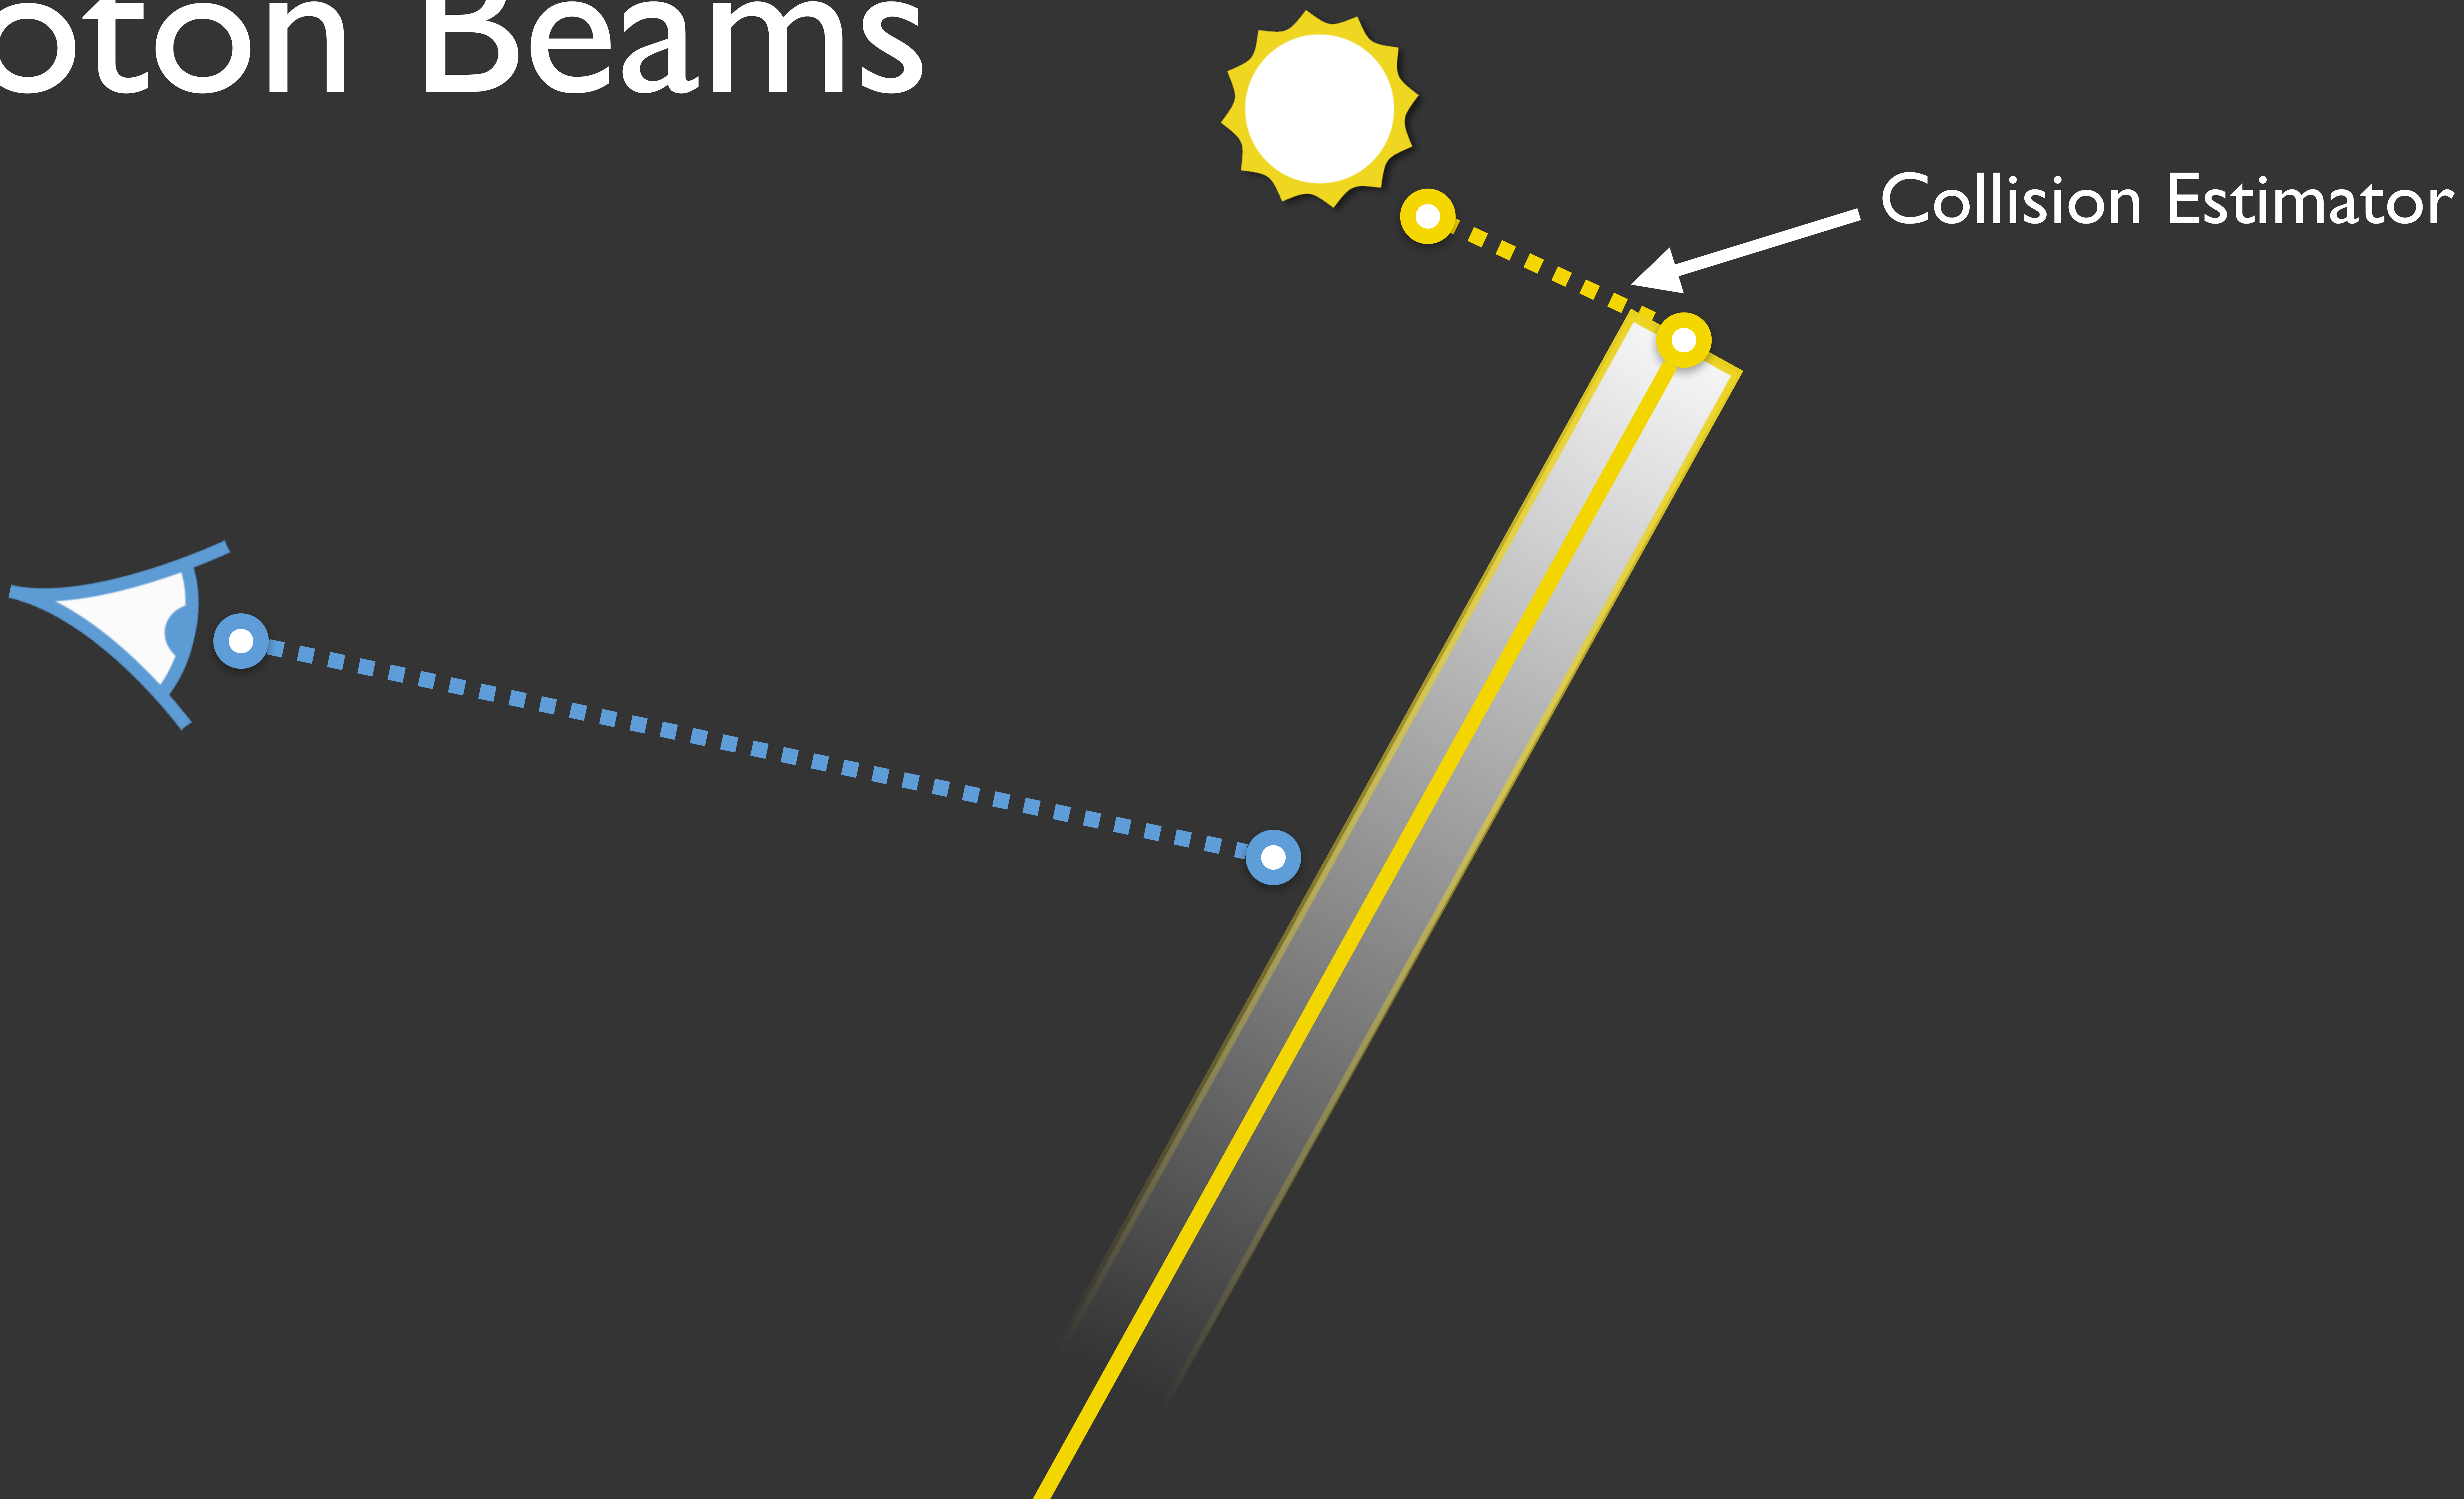

Beam Marching

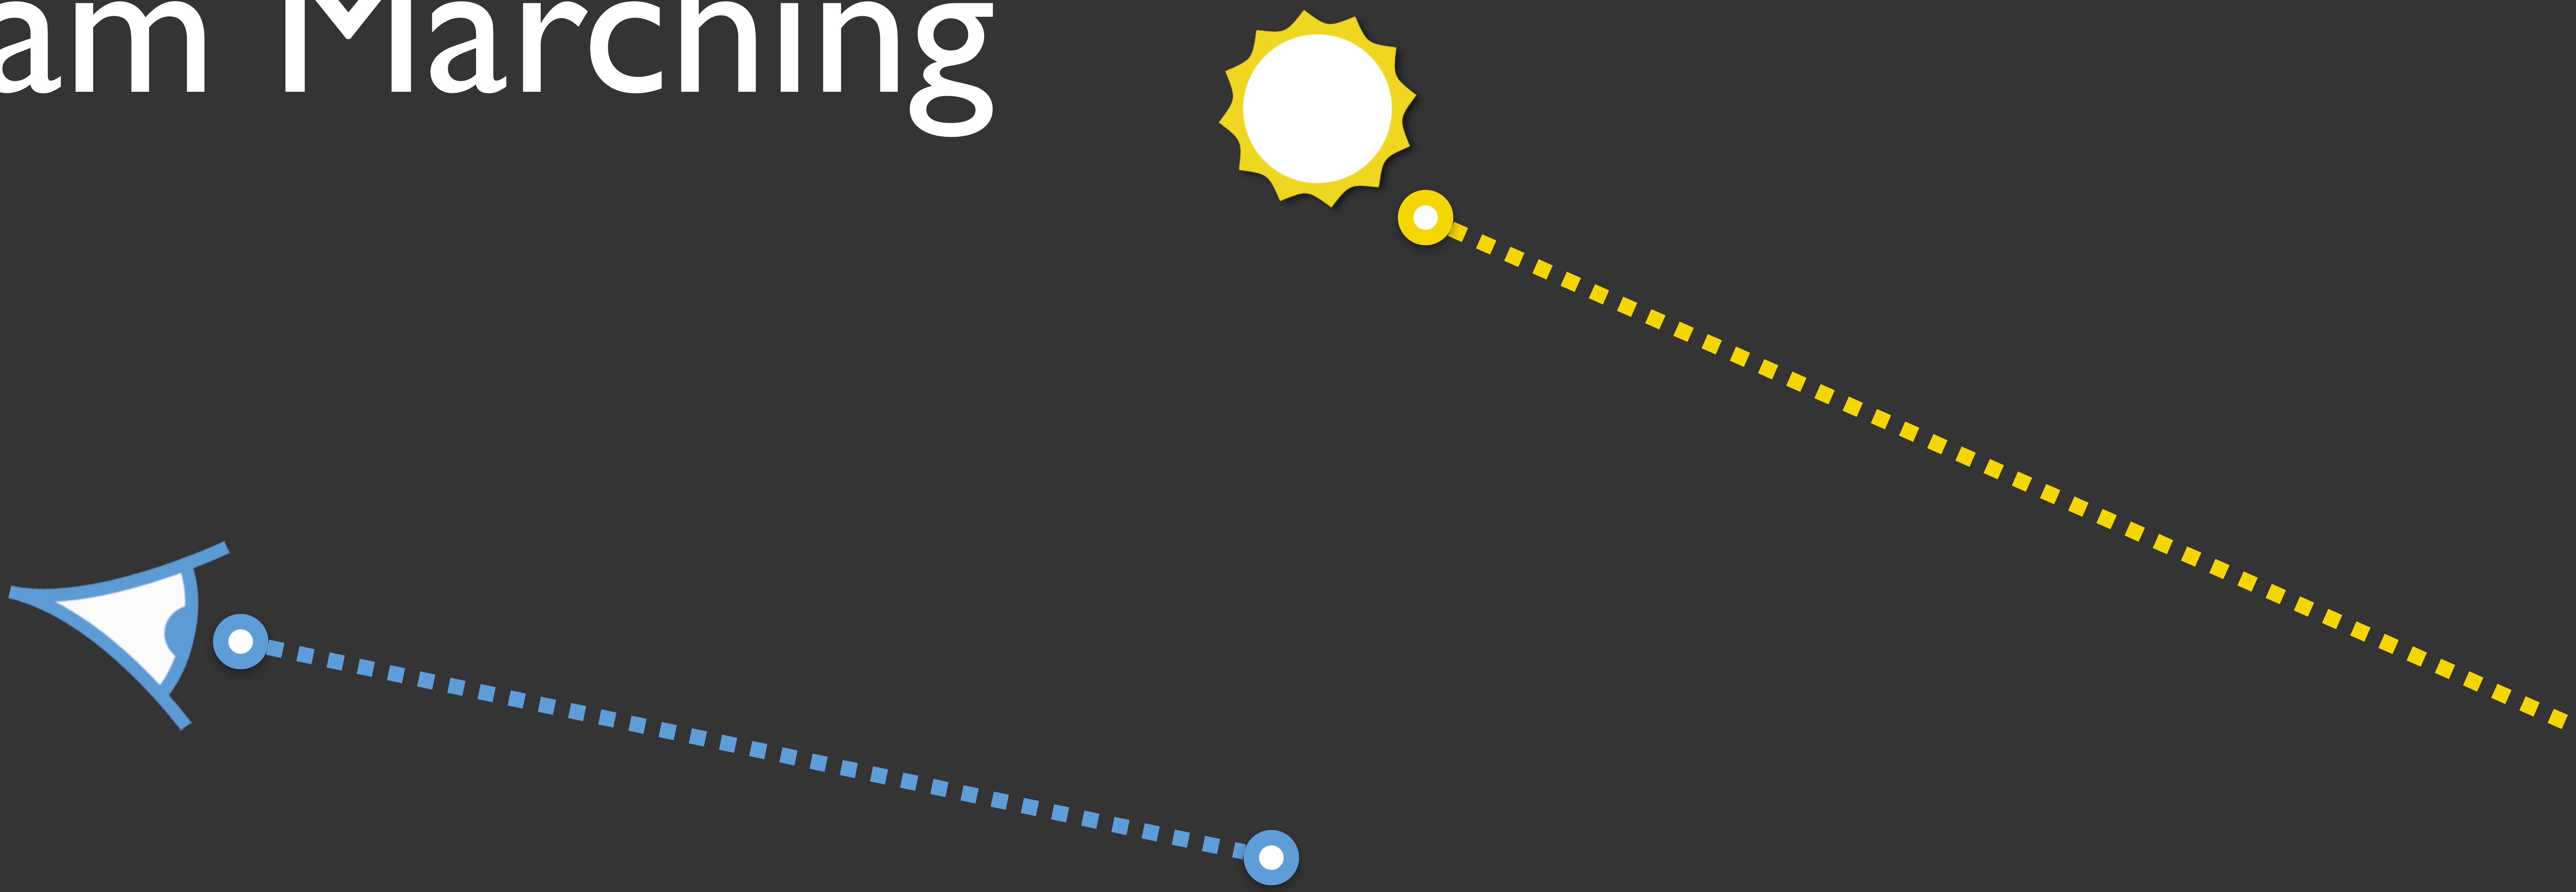

Beam Marching

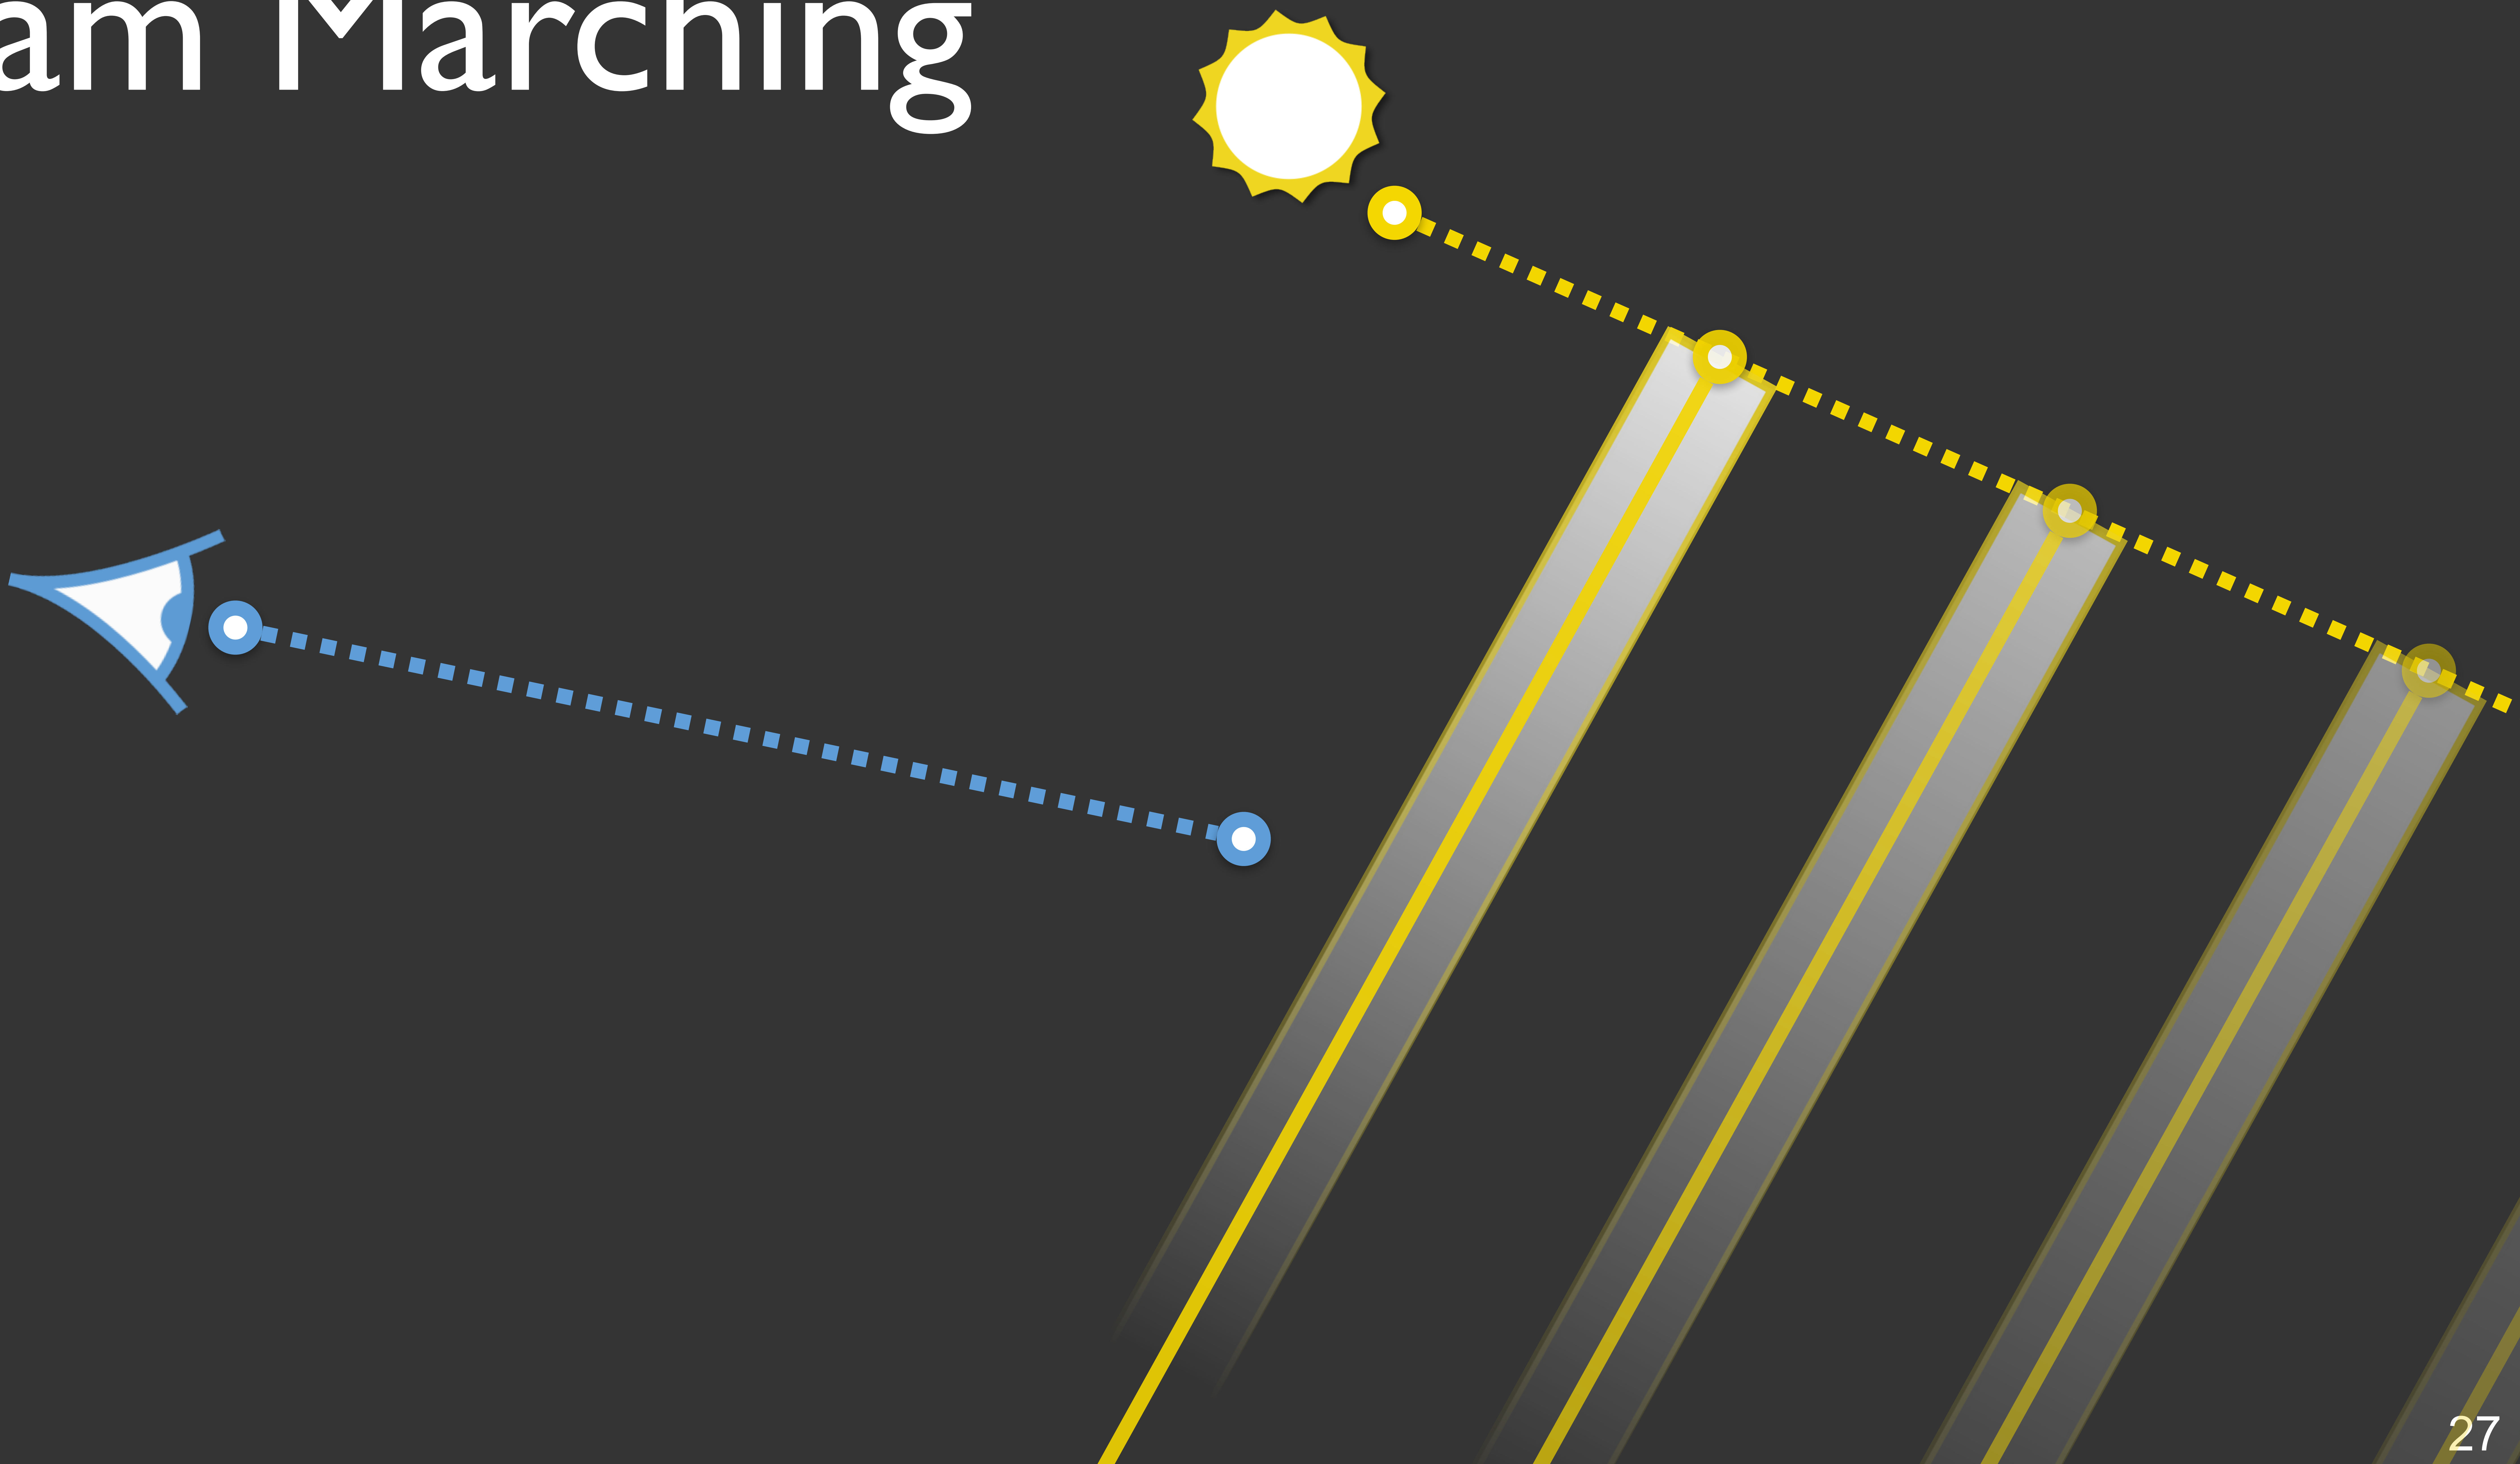

Beam Marching

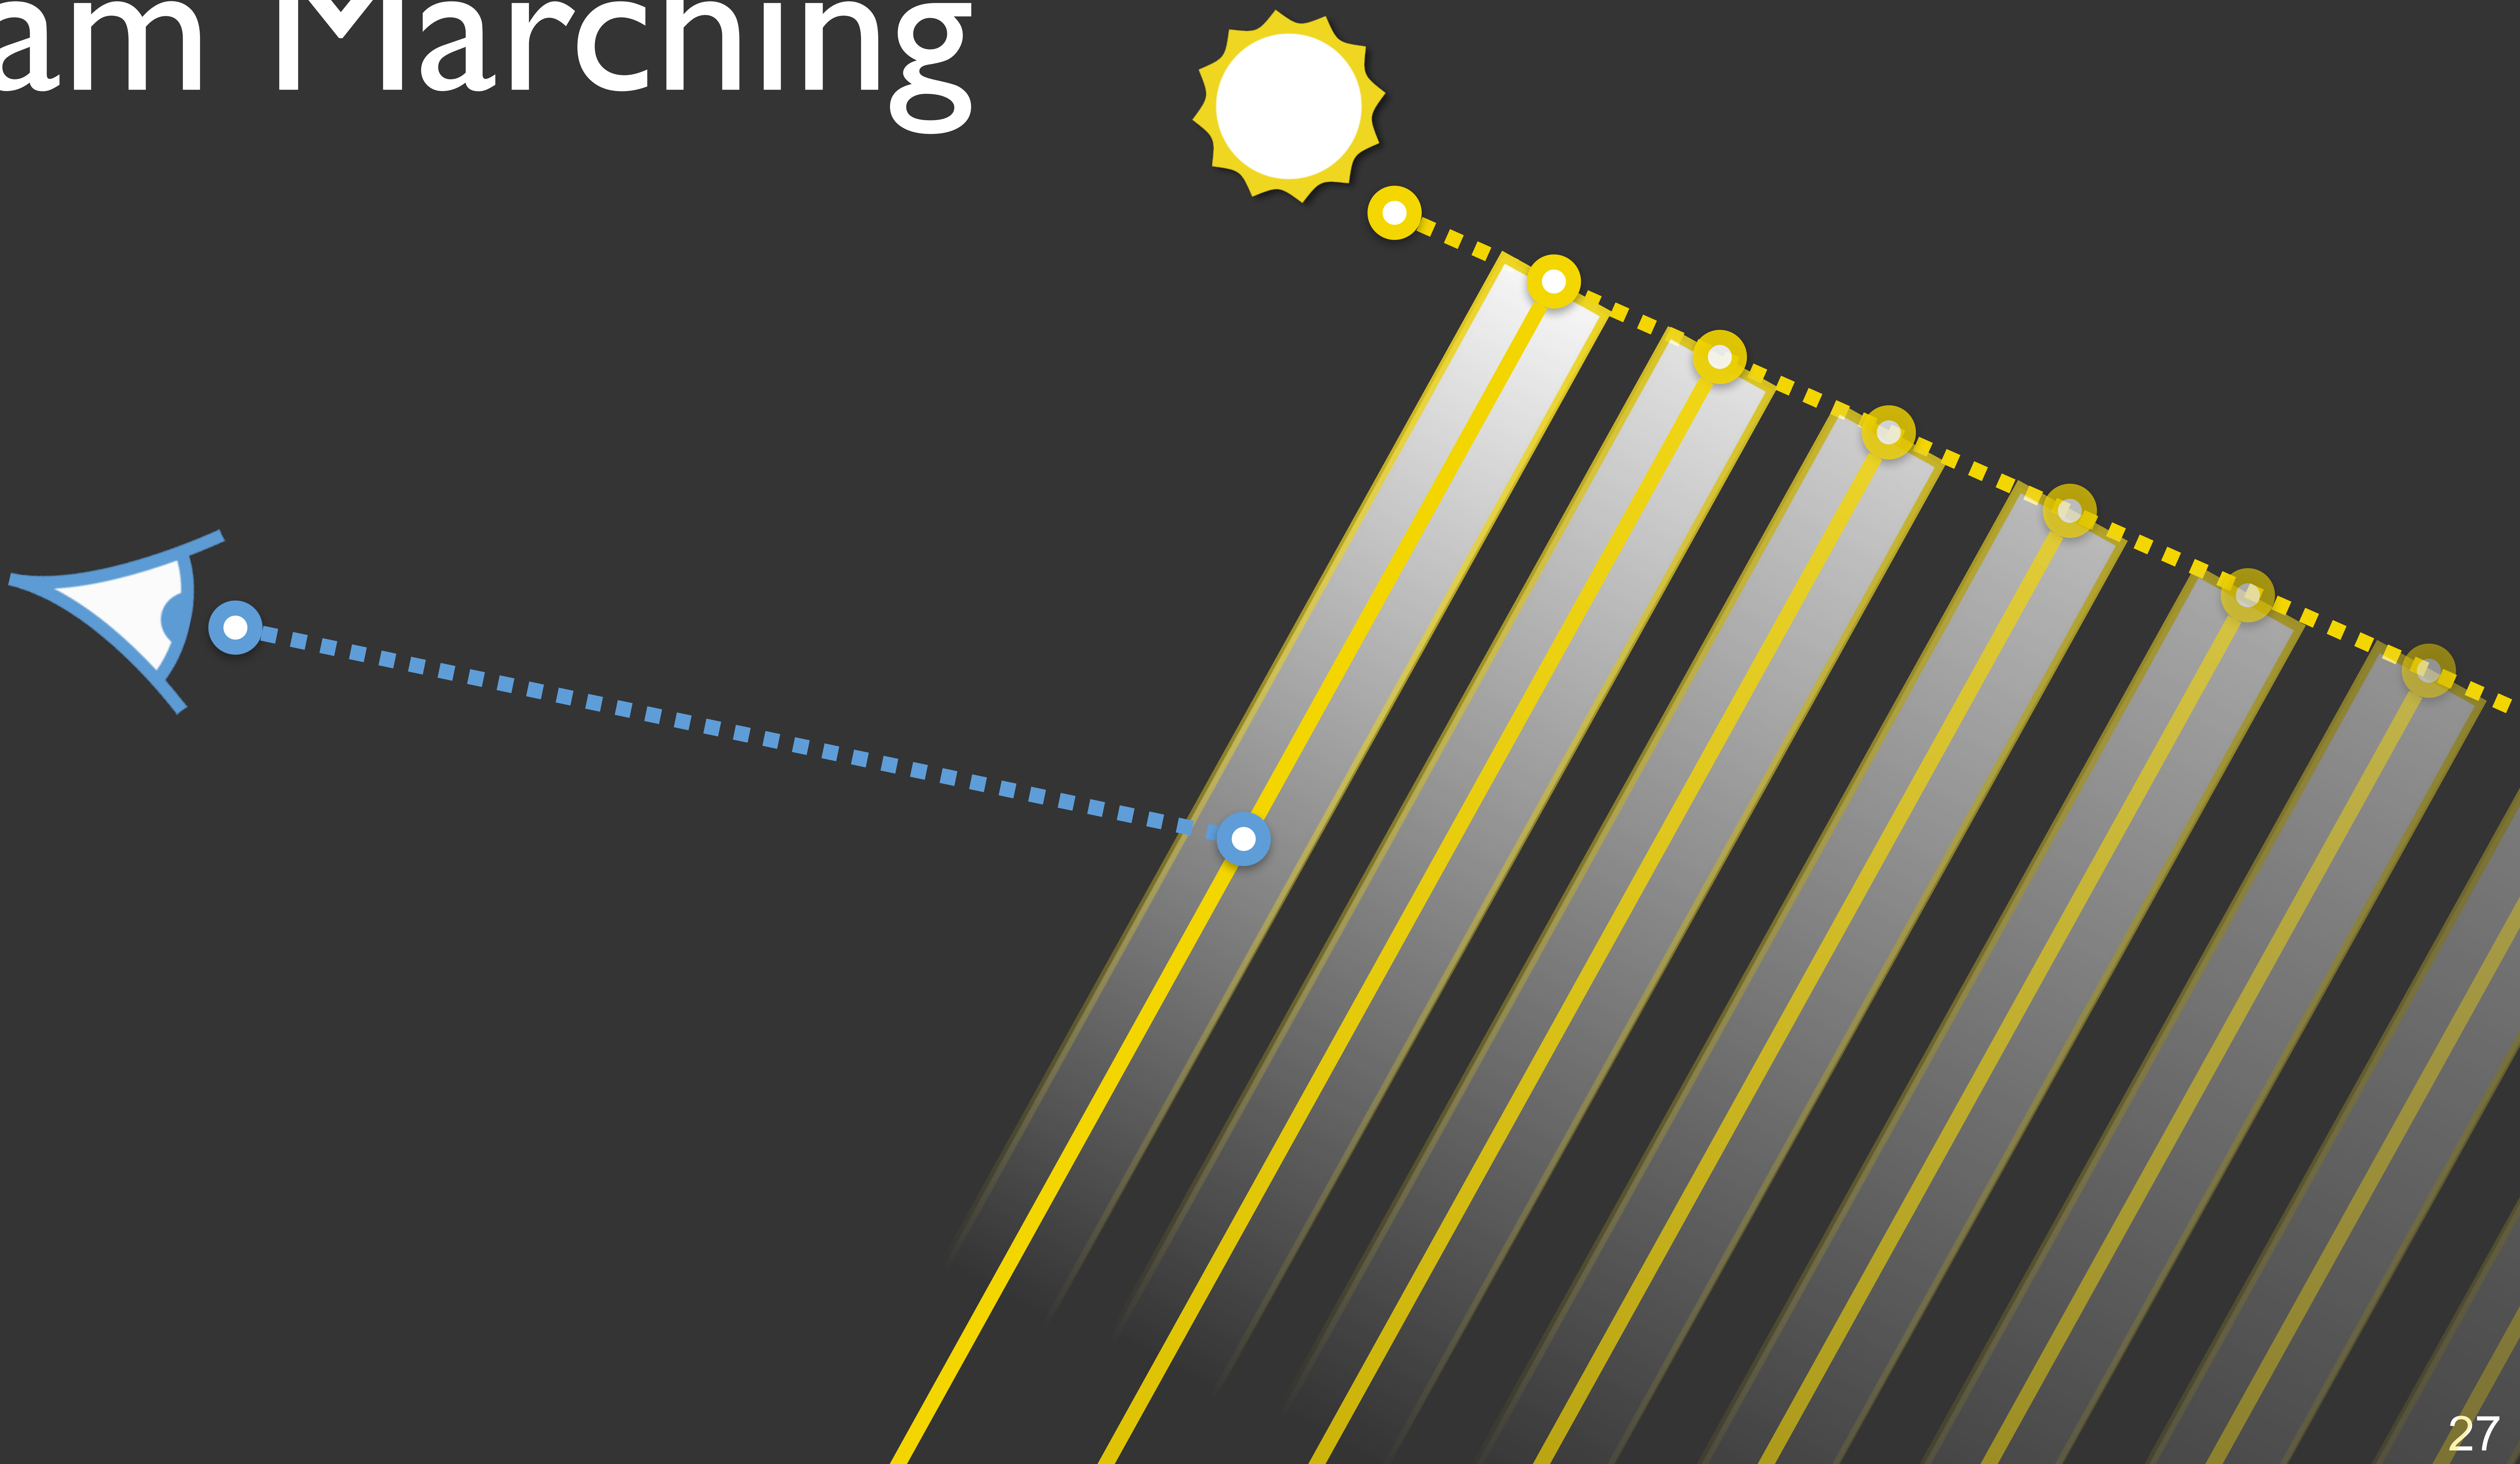

Photon Plane

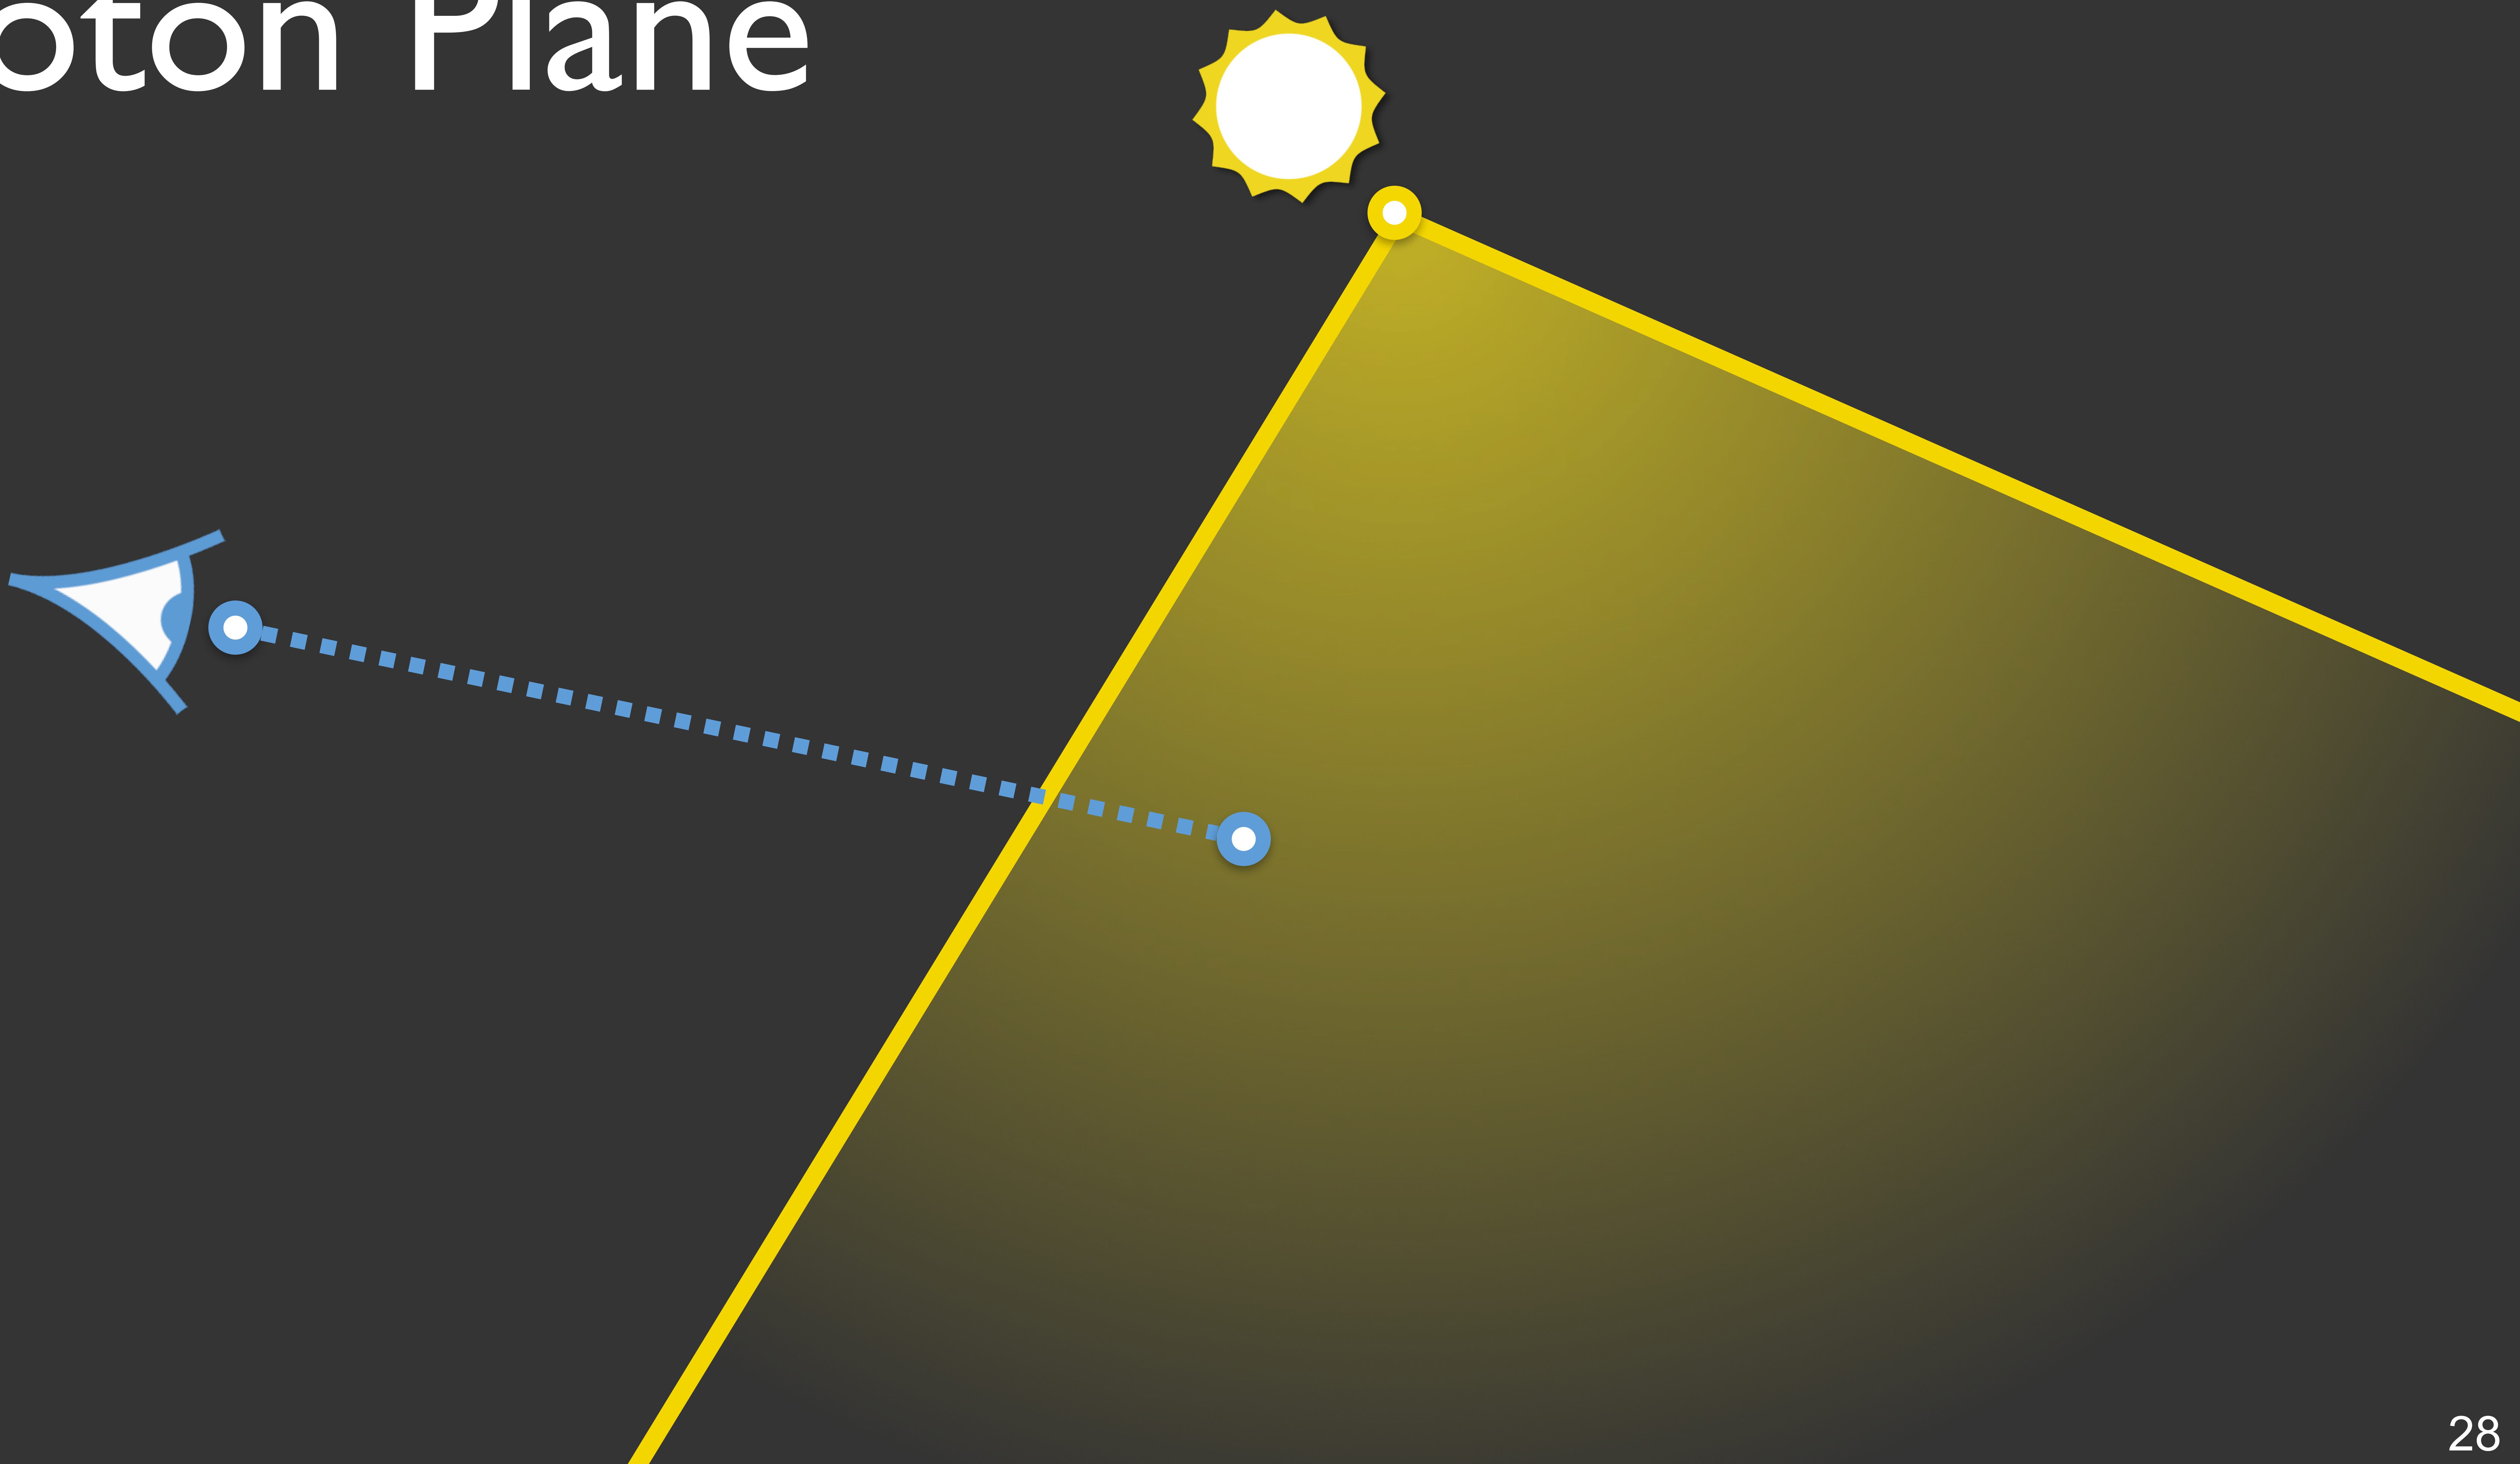

Photon Planes

- Plane geometry depends on estimators used

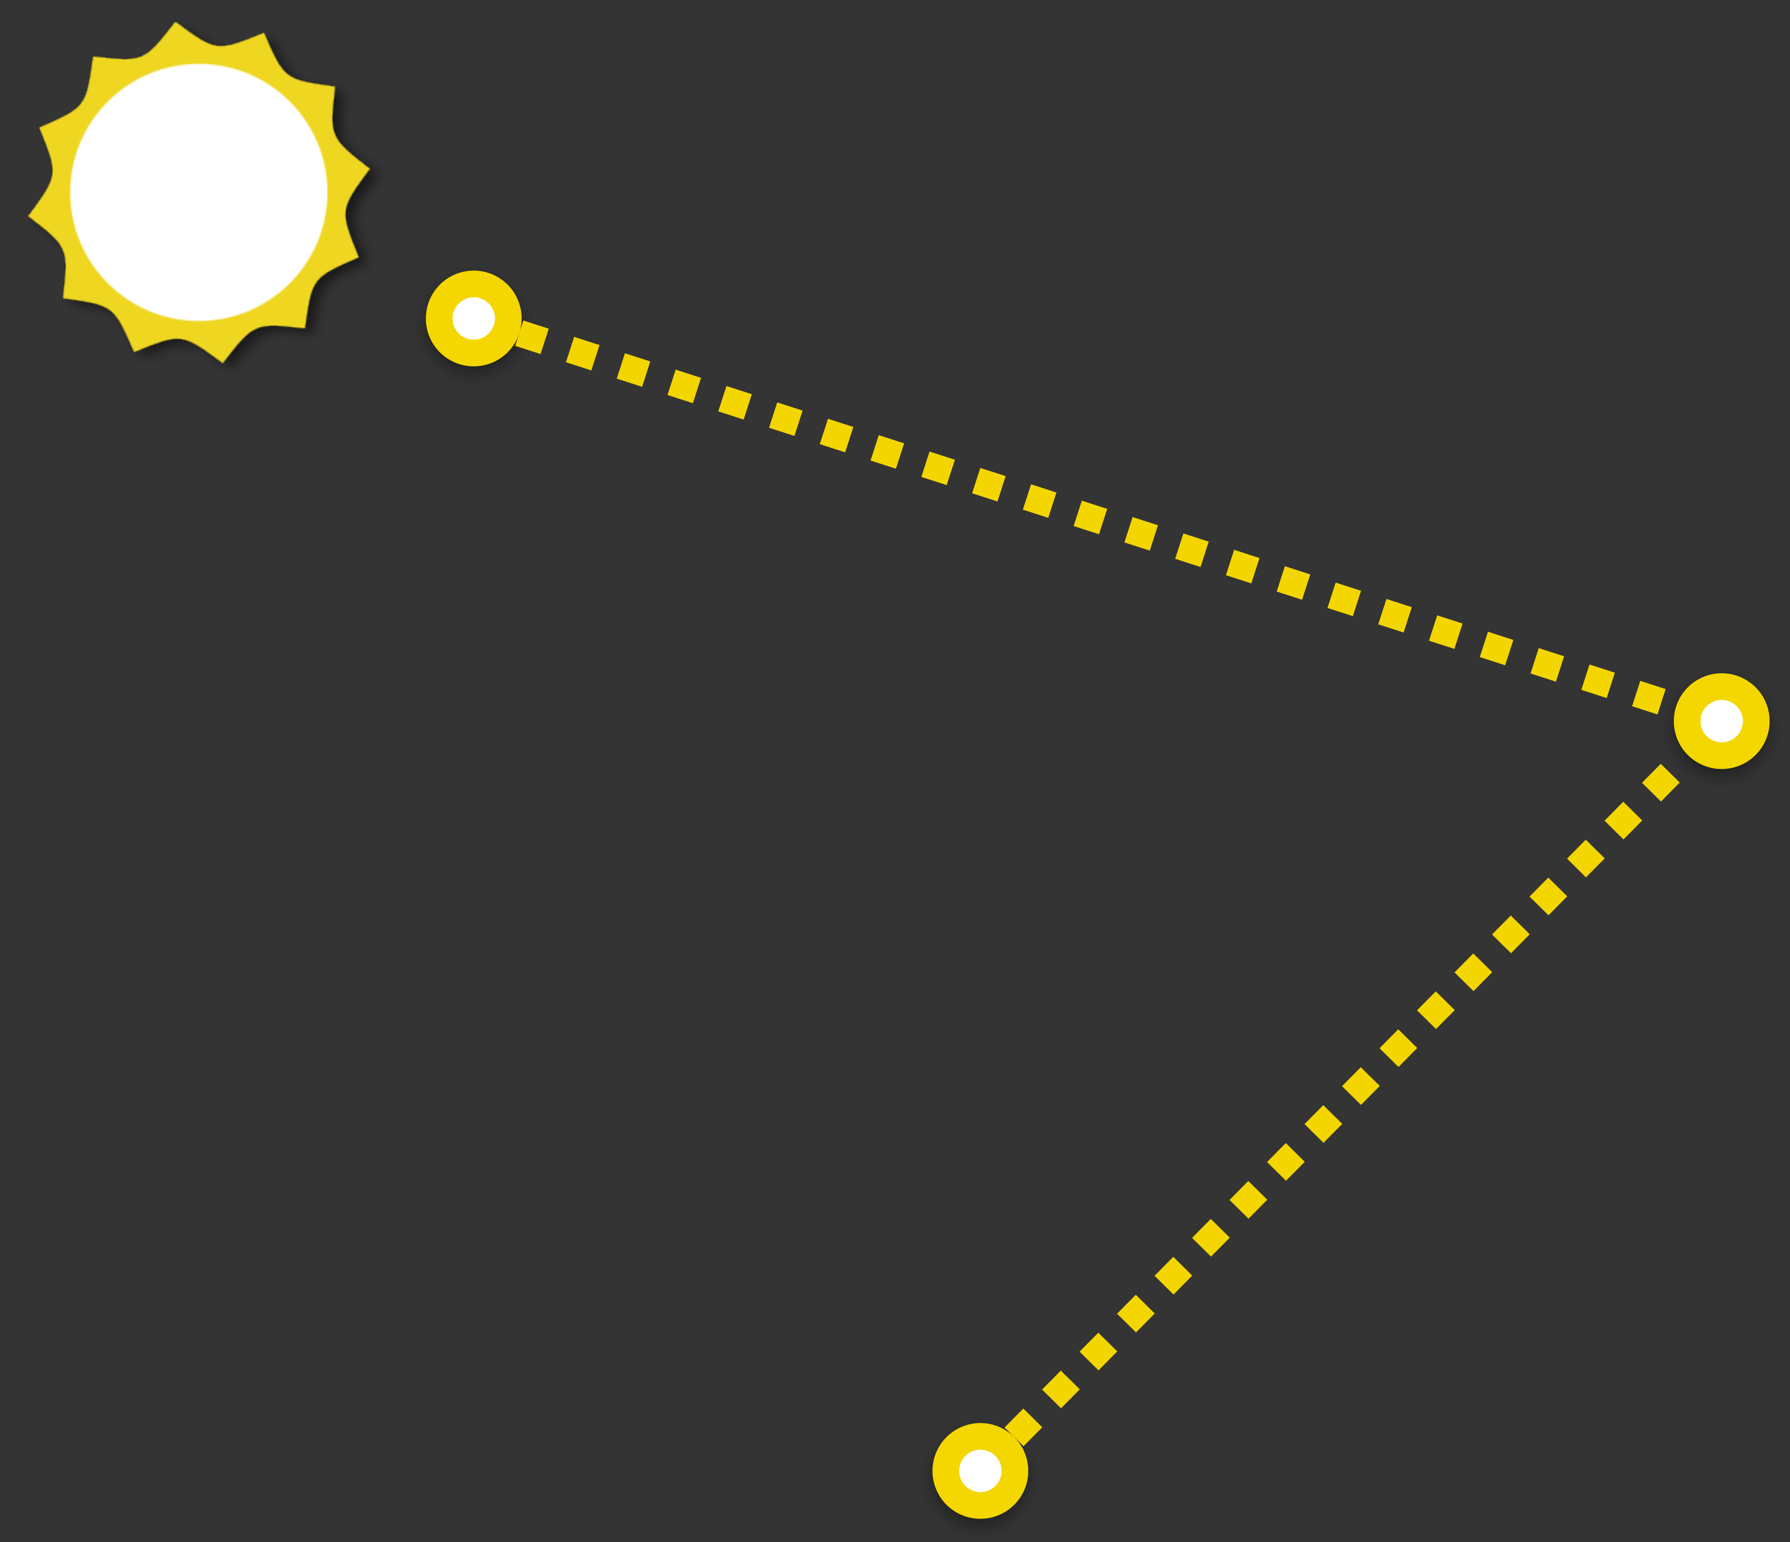

Photon Planes

- Plane geometry depends on estimators used

“Short” Plane

Track-Length × Track-Length

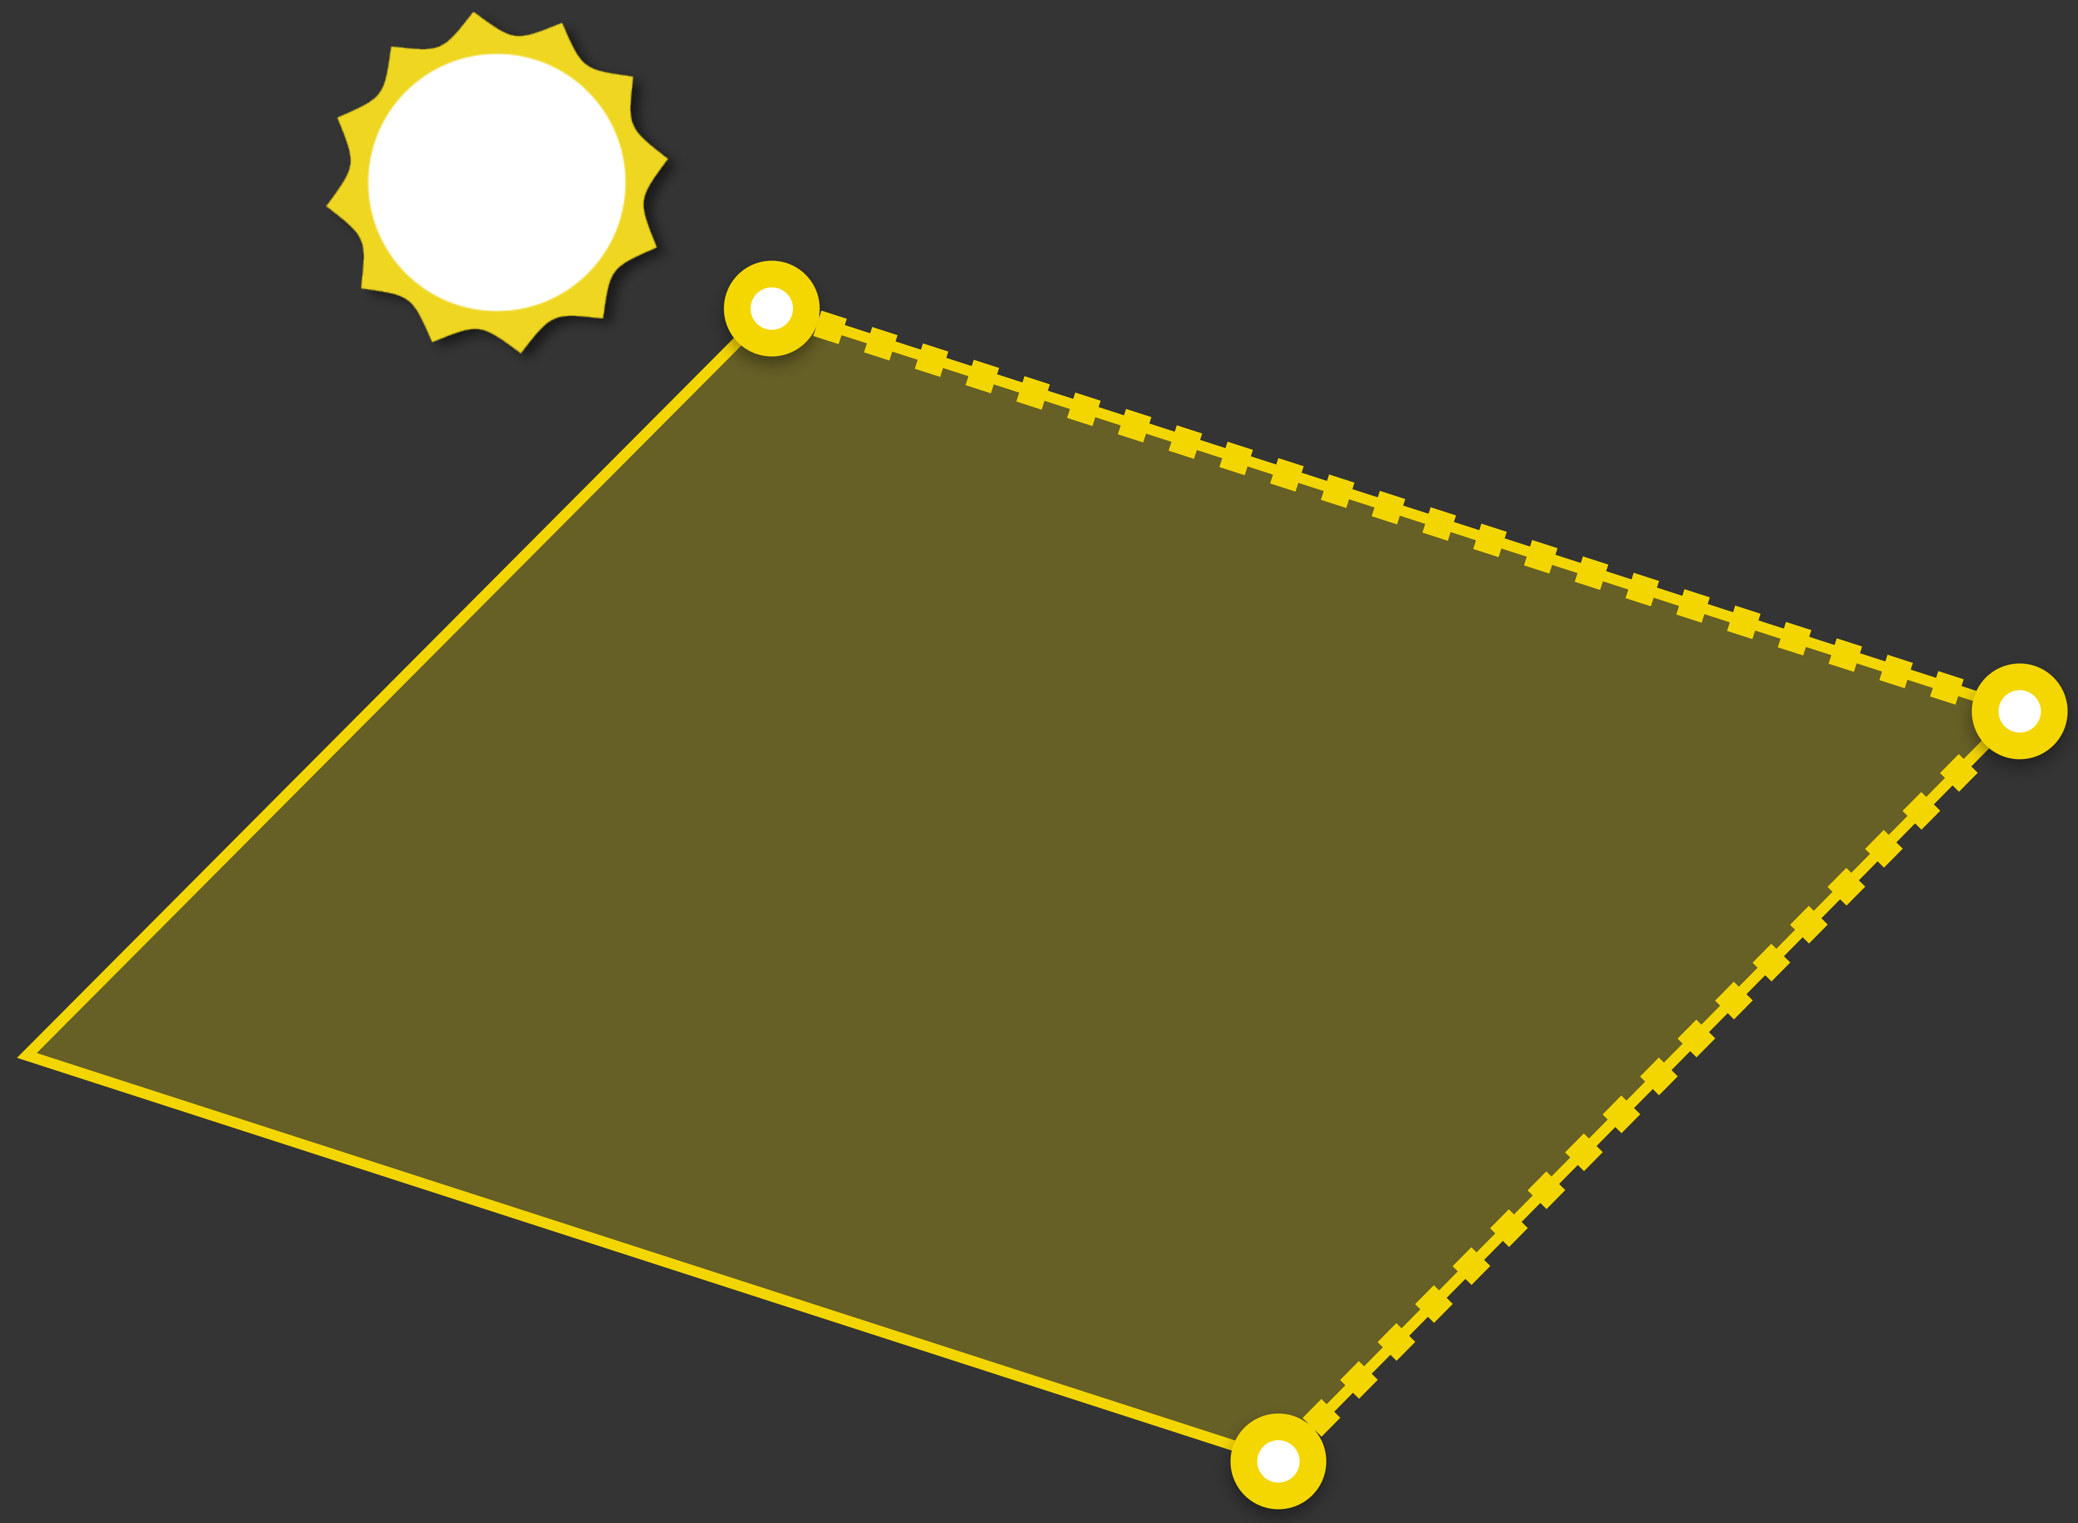

Photon Planes

- Plane geometry depends on estimators used

"Long" Plane

Expected × Expected

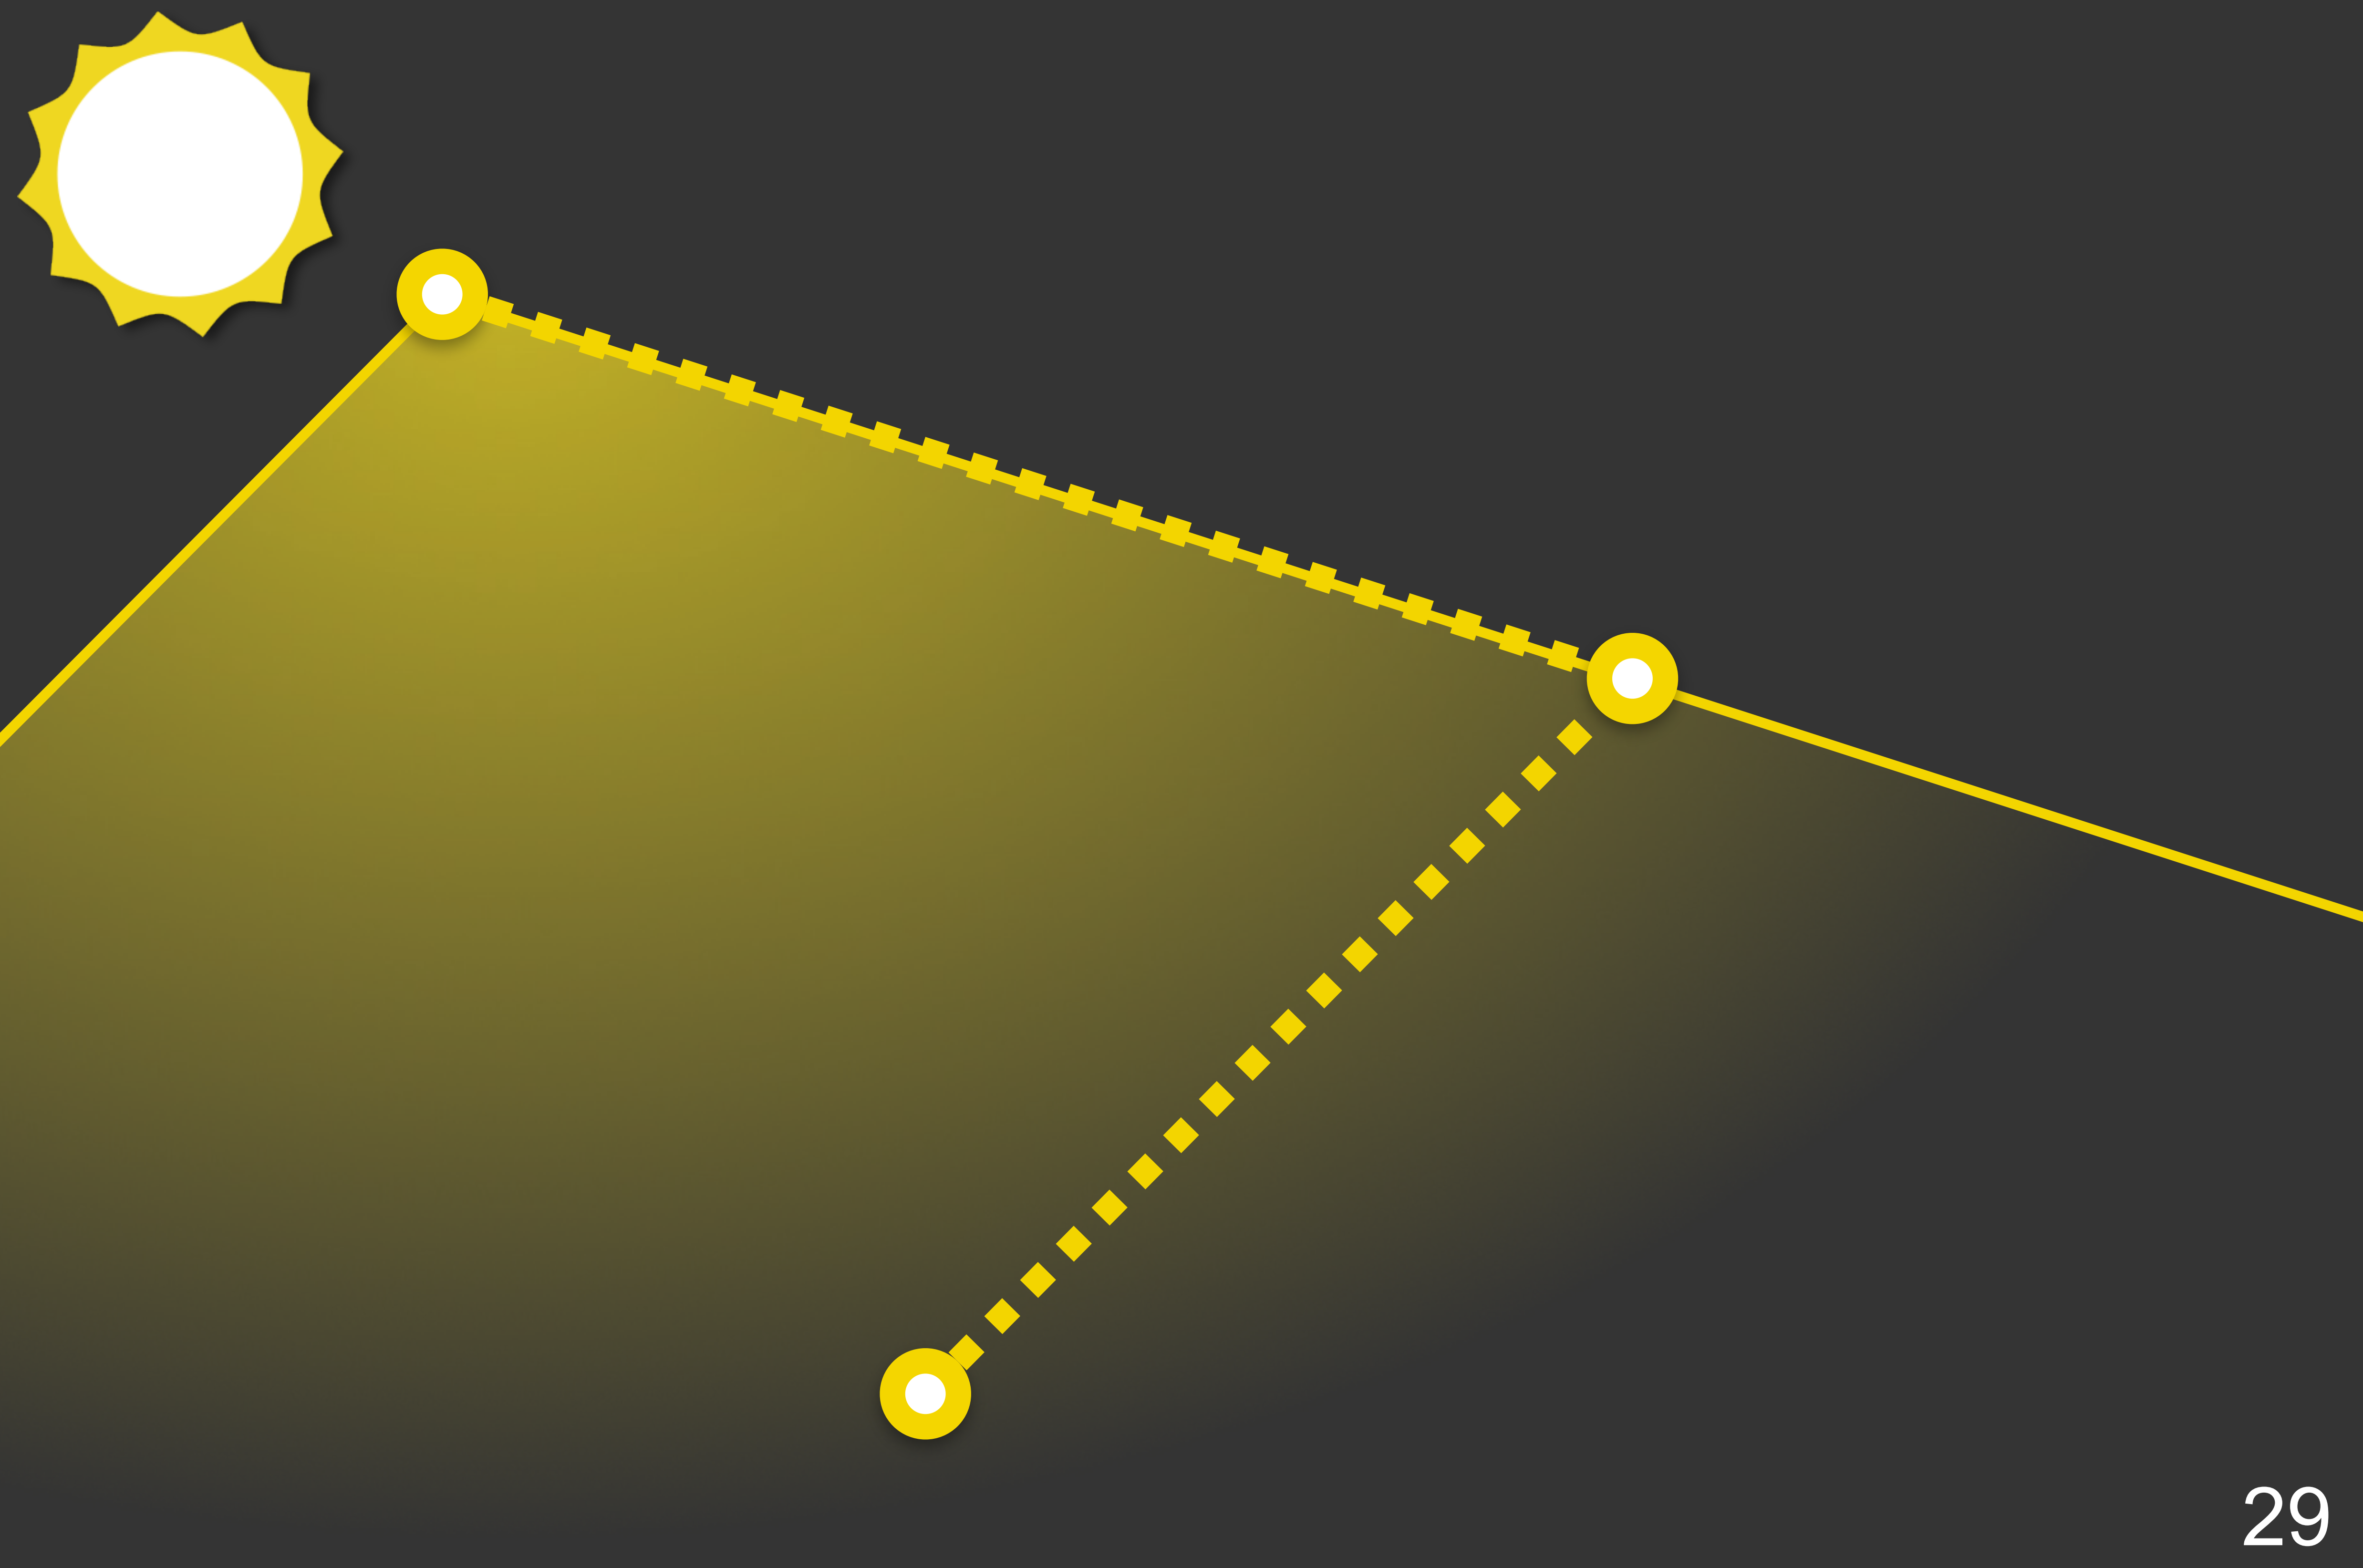

Photon Planes

- Plane geometry depends on estimators used

“Long-Short” Plane

Expected × Track-Length

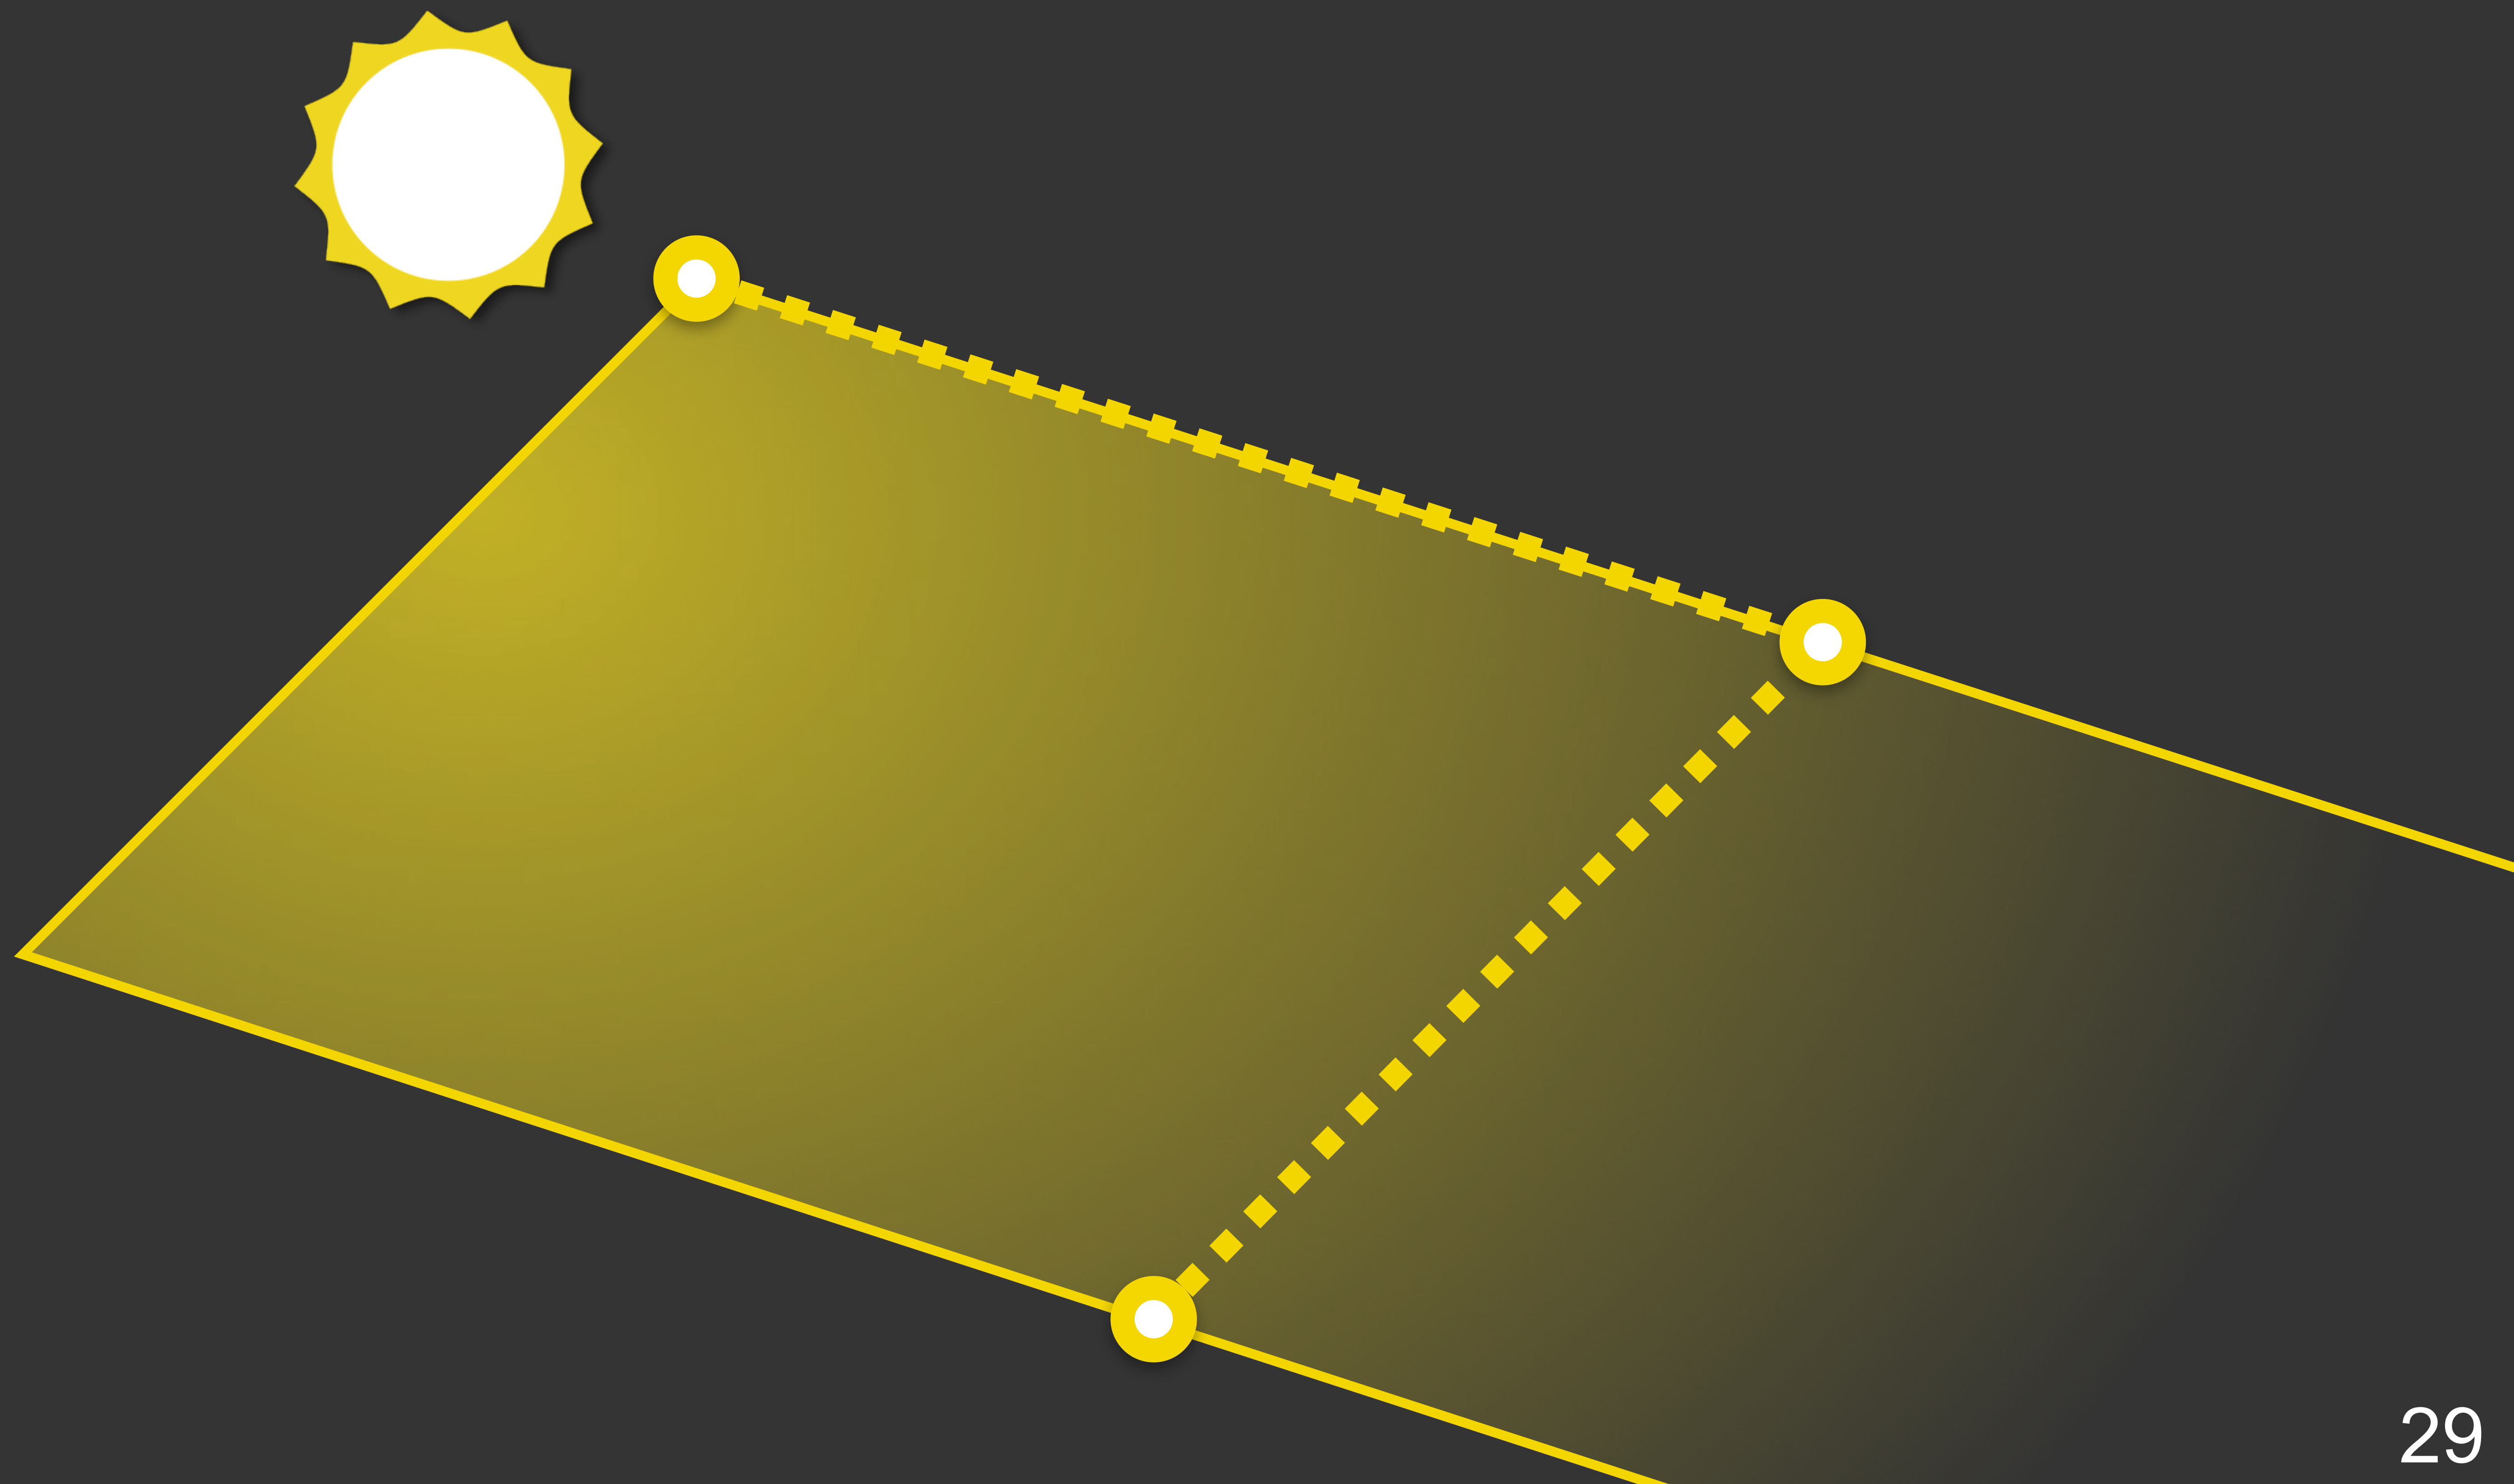

Photon Planes

- Plane geometry depends on estimators used

“Short-Long” Plane

Track-Length × Expected

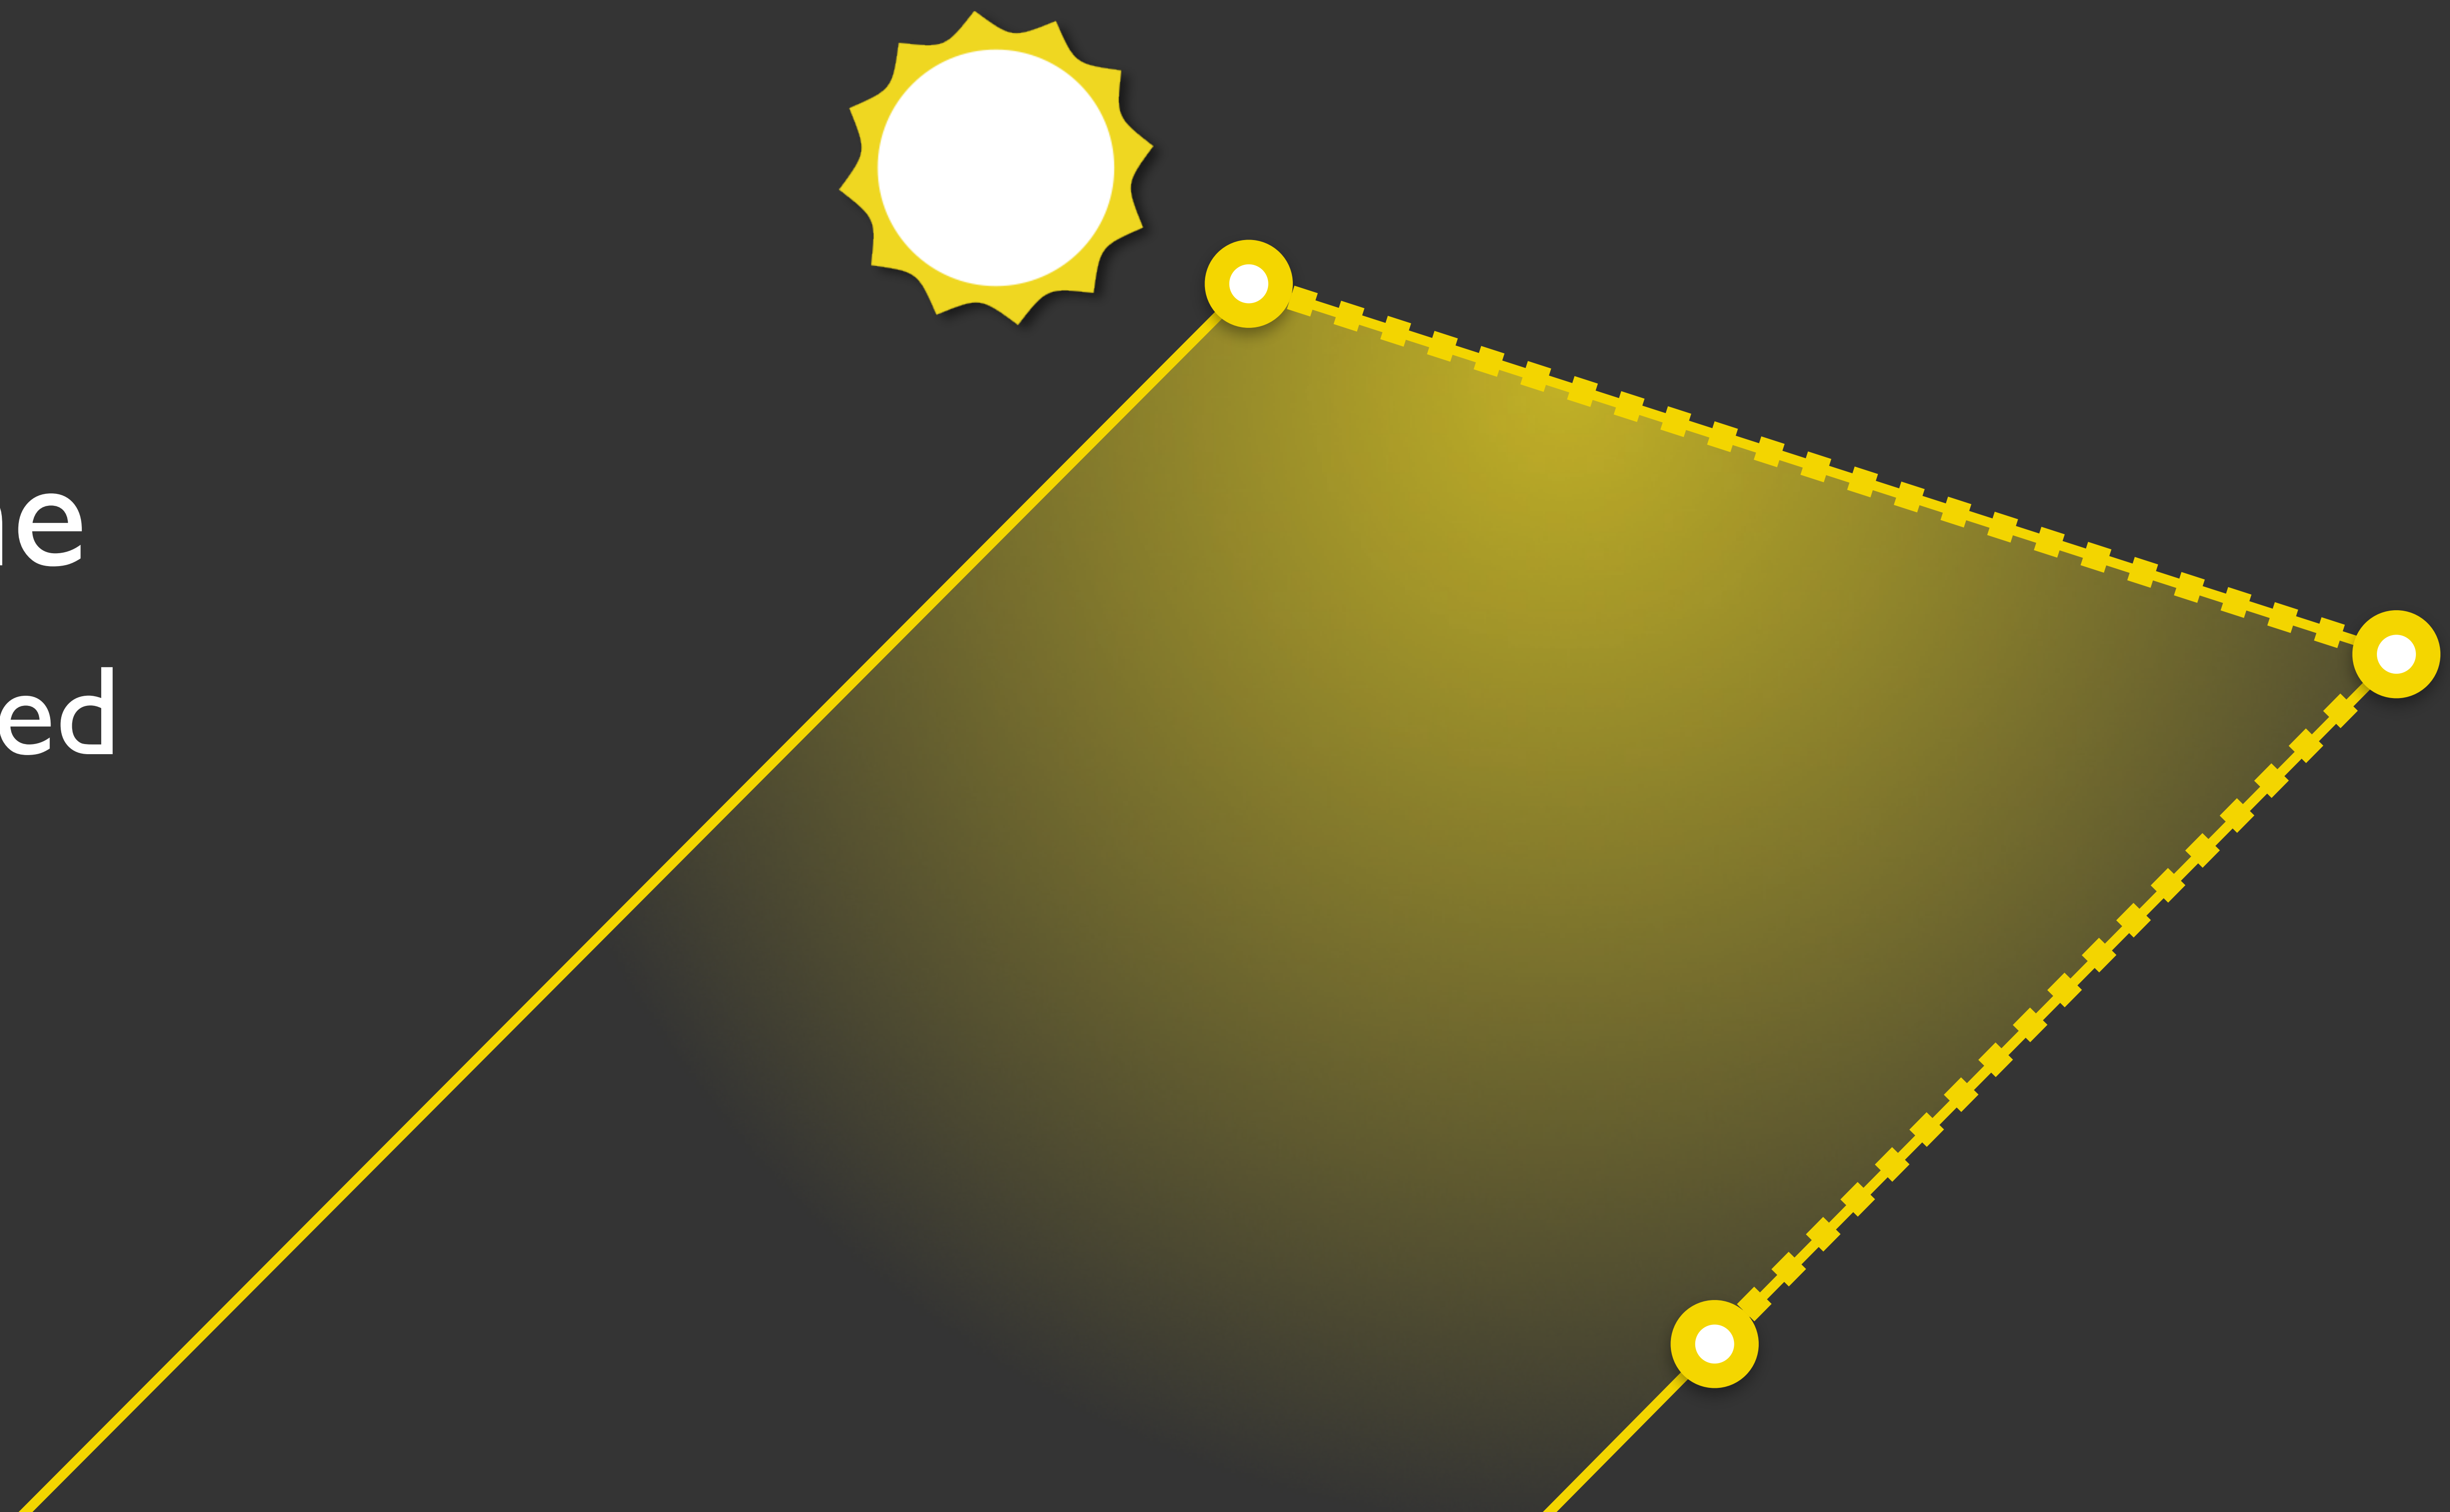

Beyond Points and Beams

- We can keep repeating this!

Plane Marching

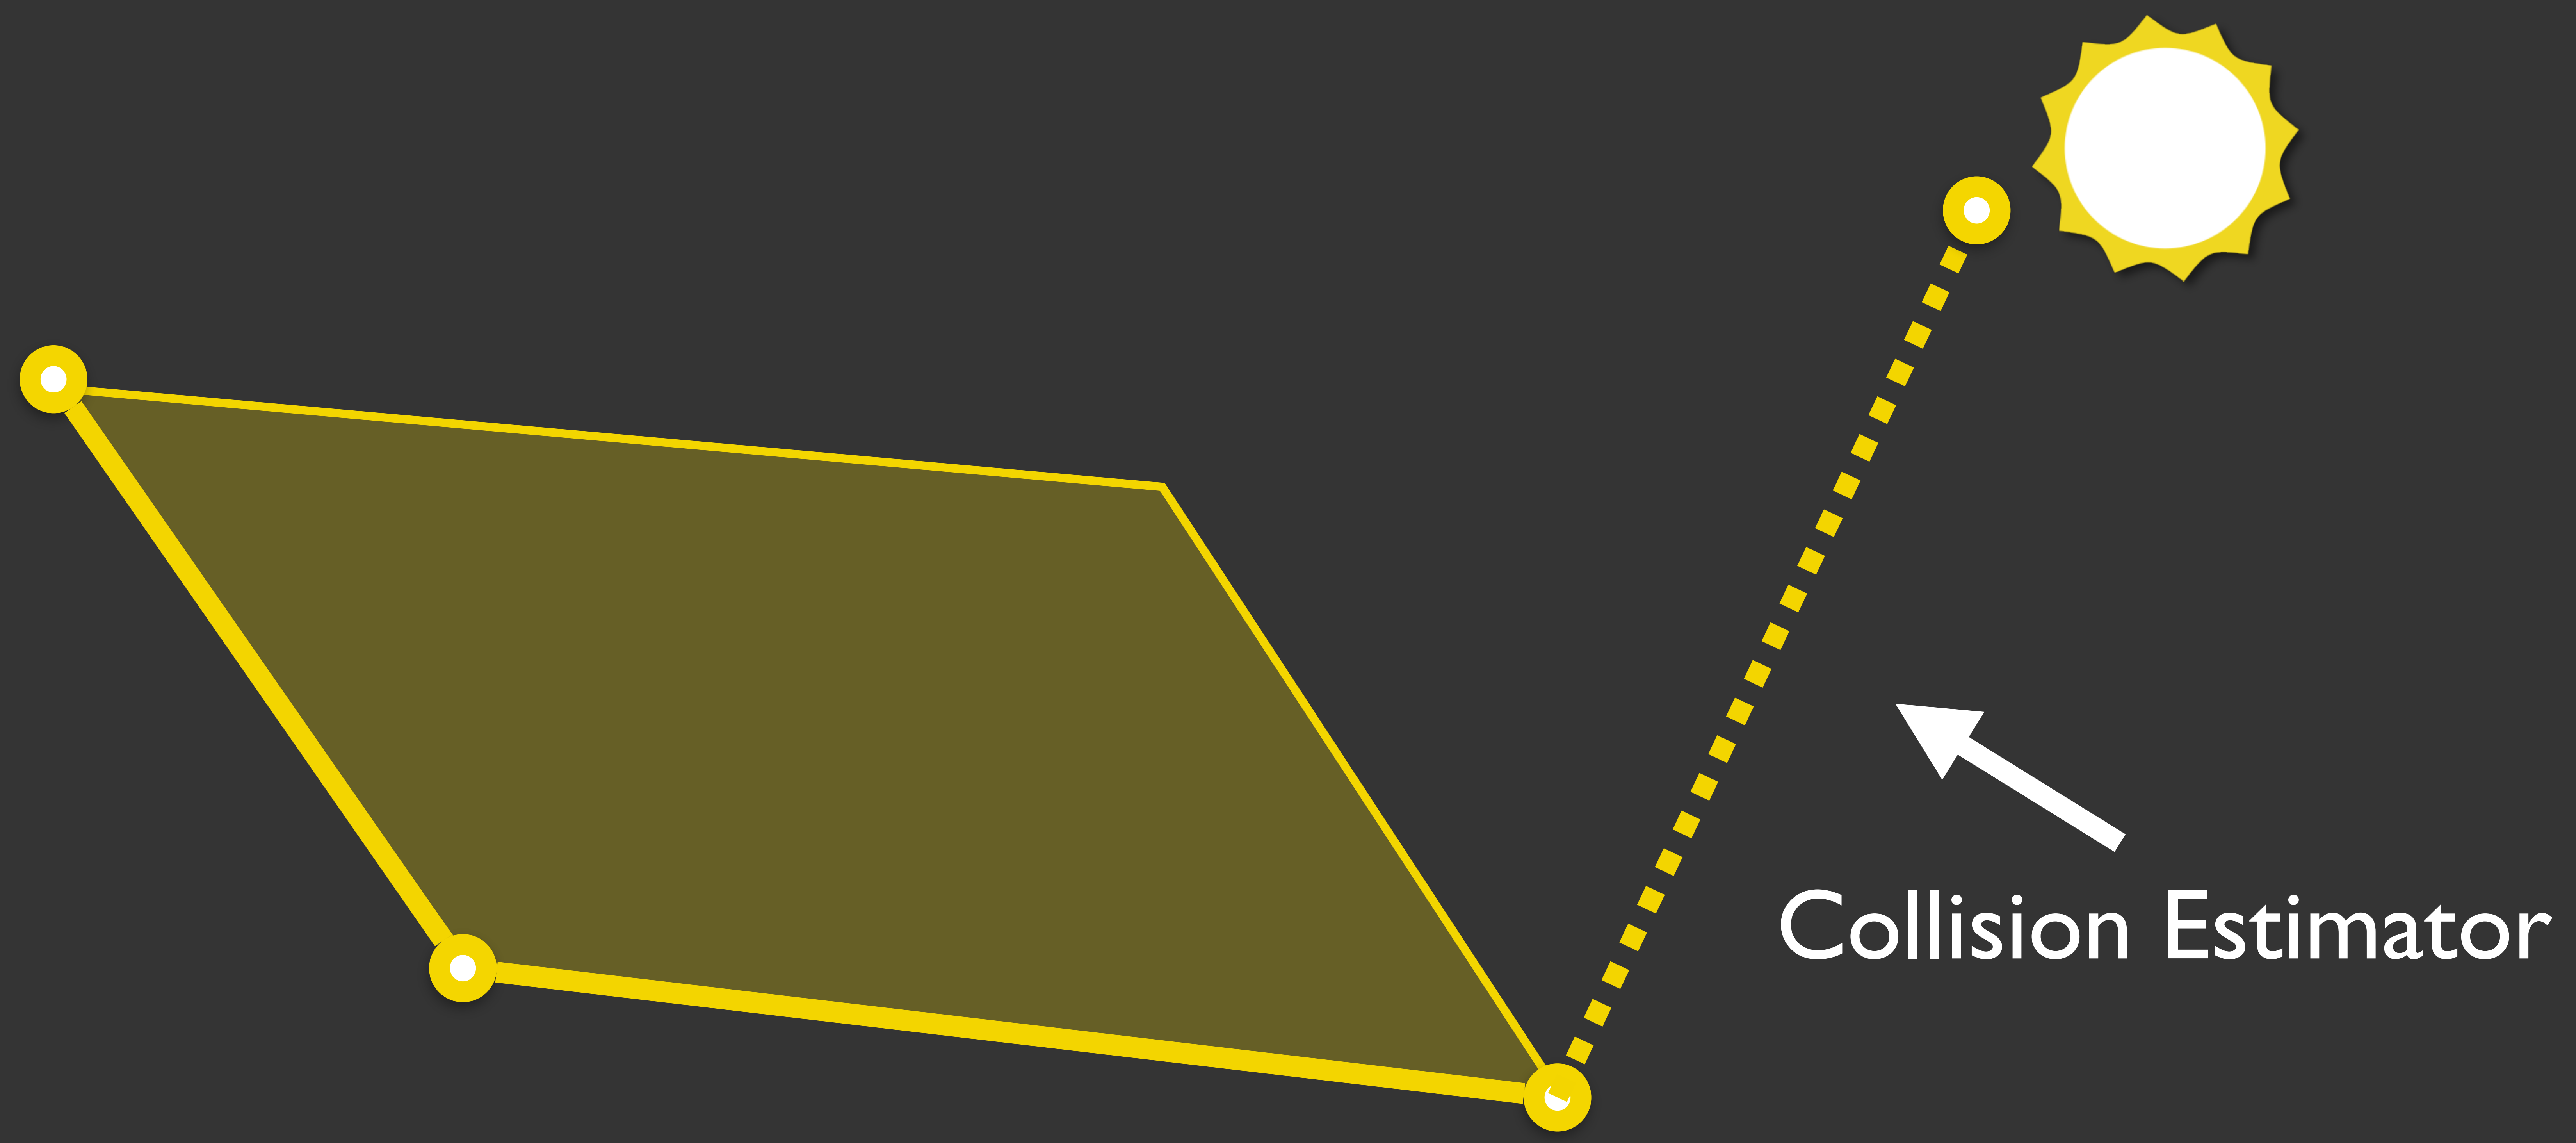

Plane Marching

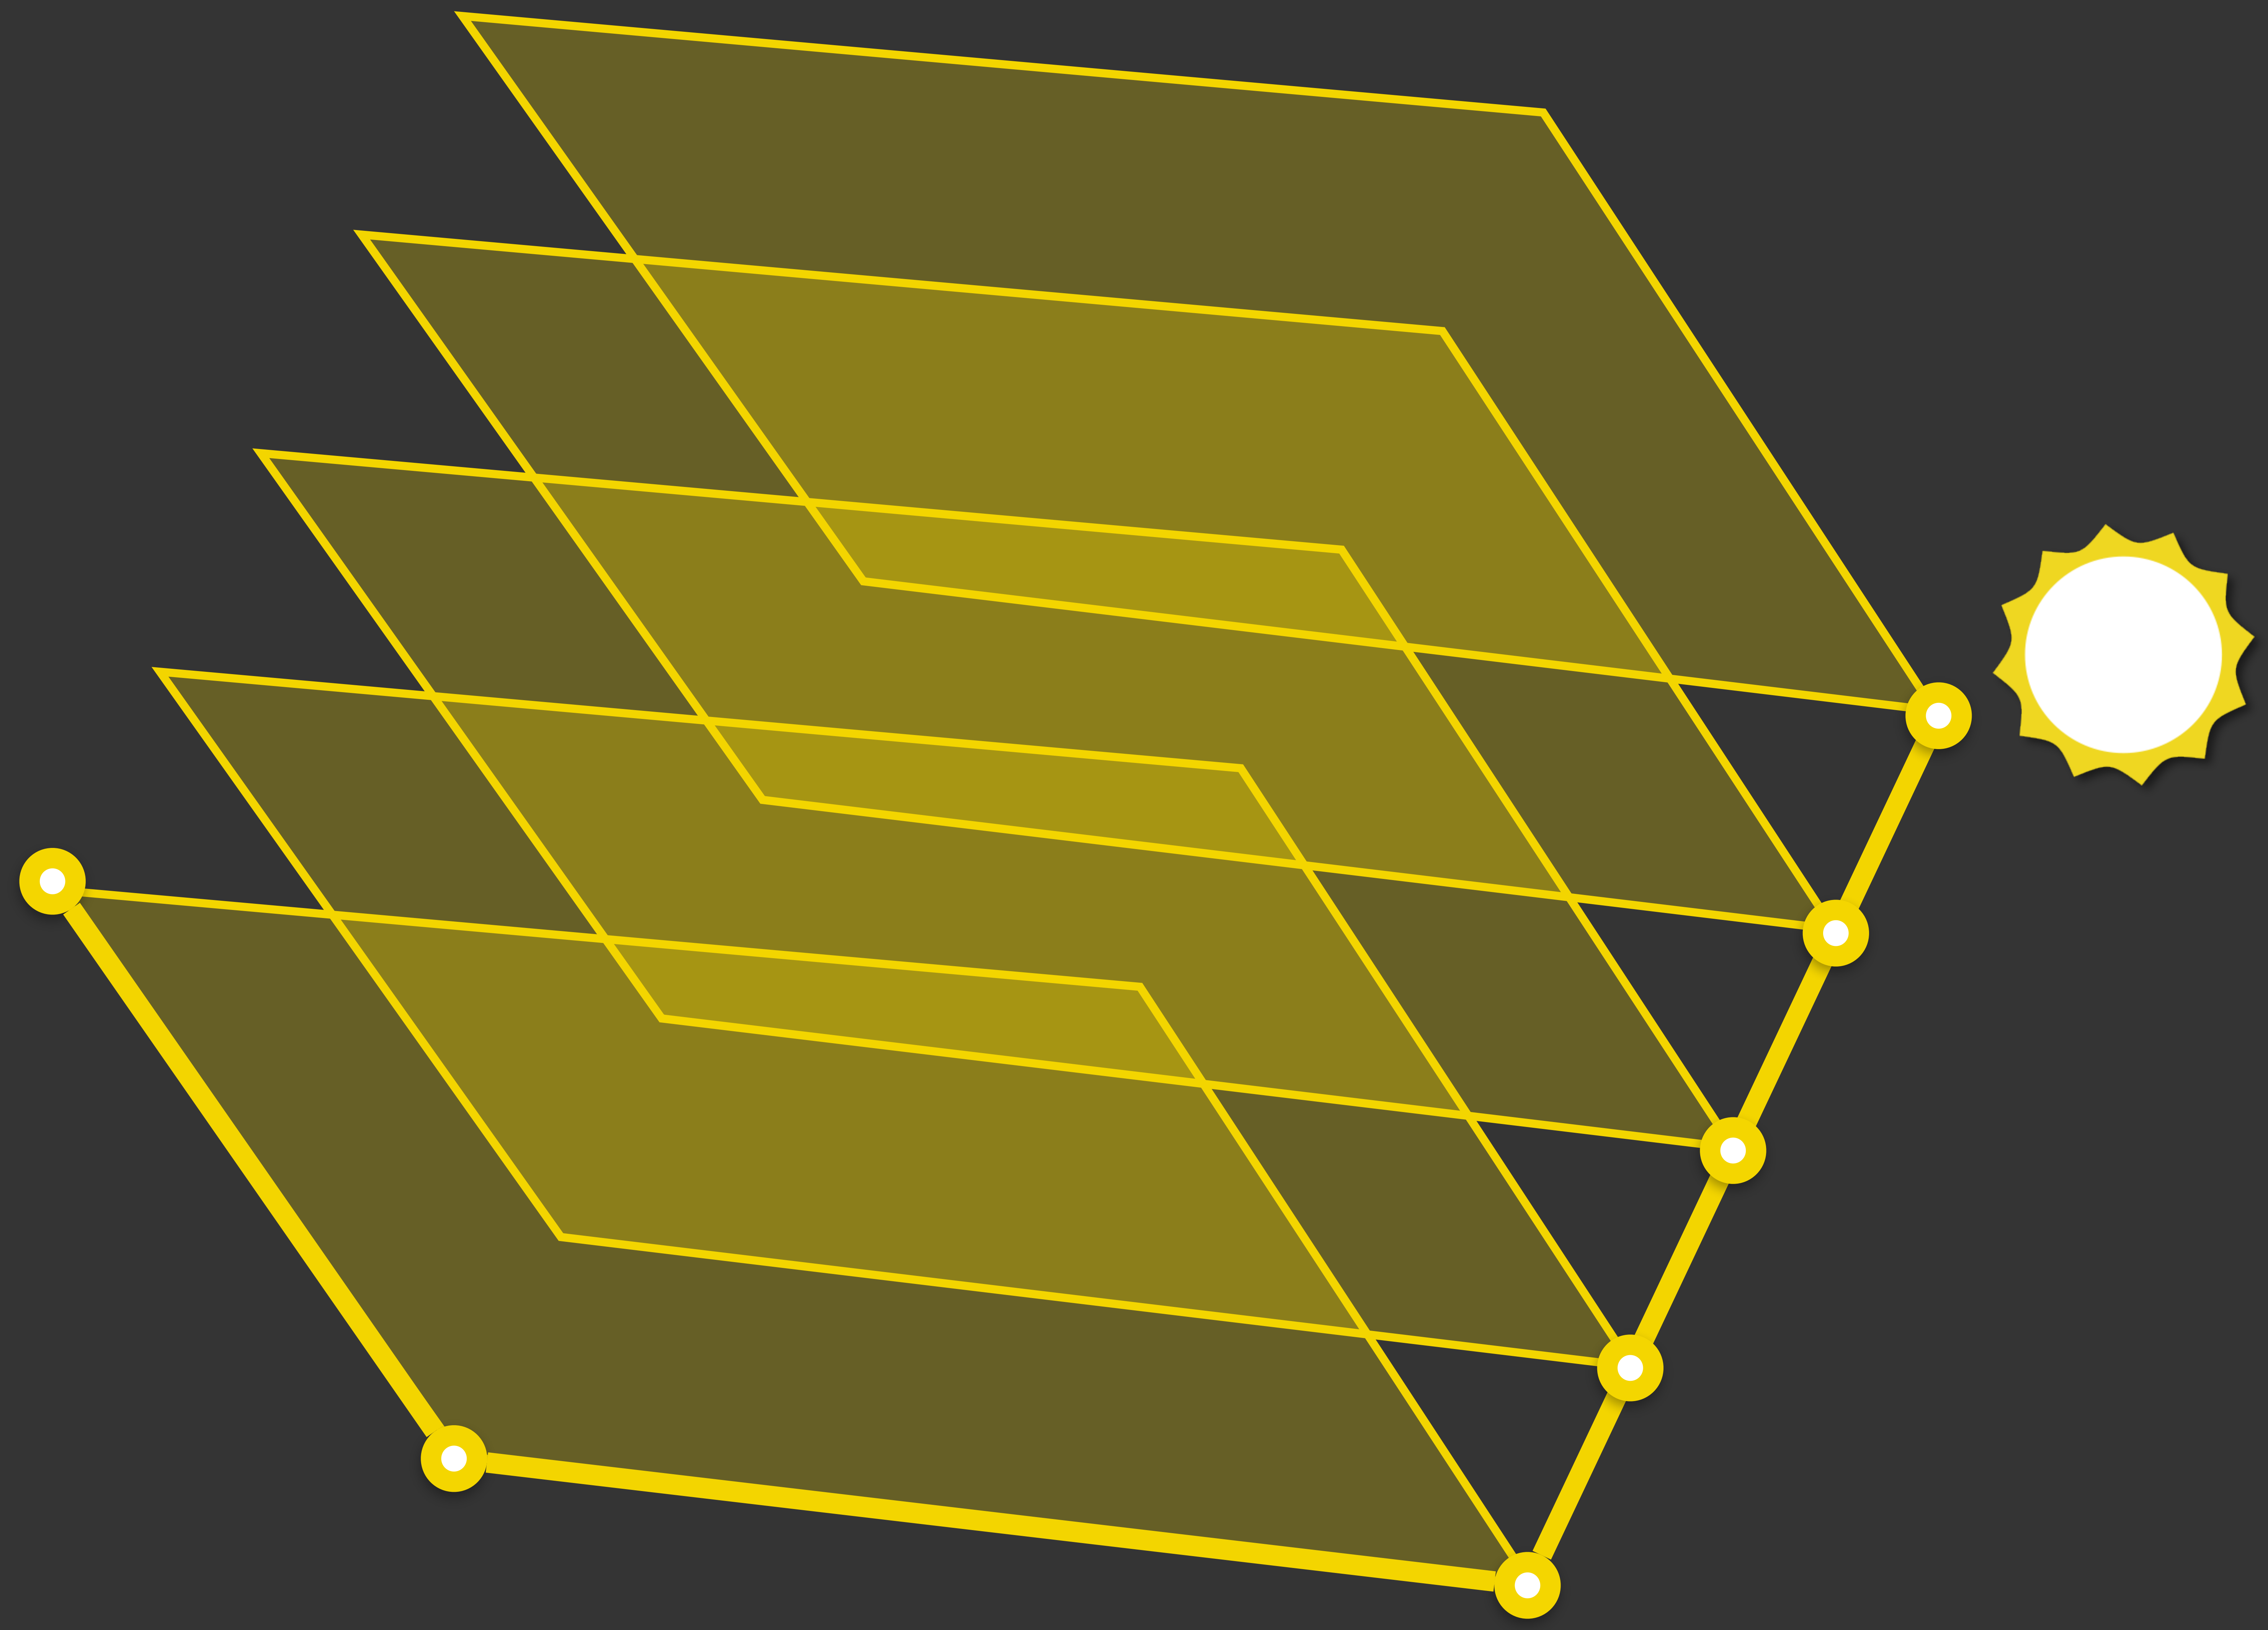

Photon Volume

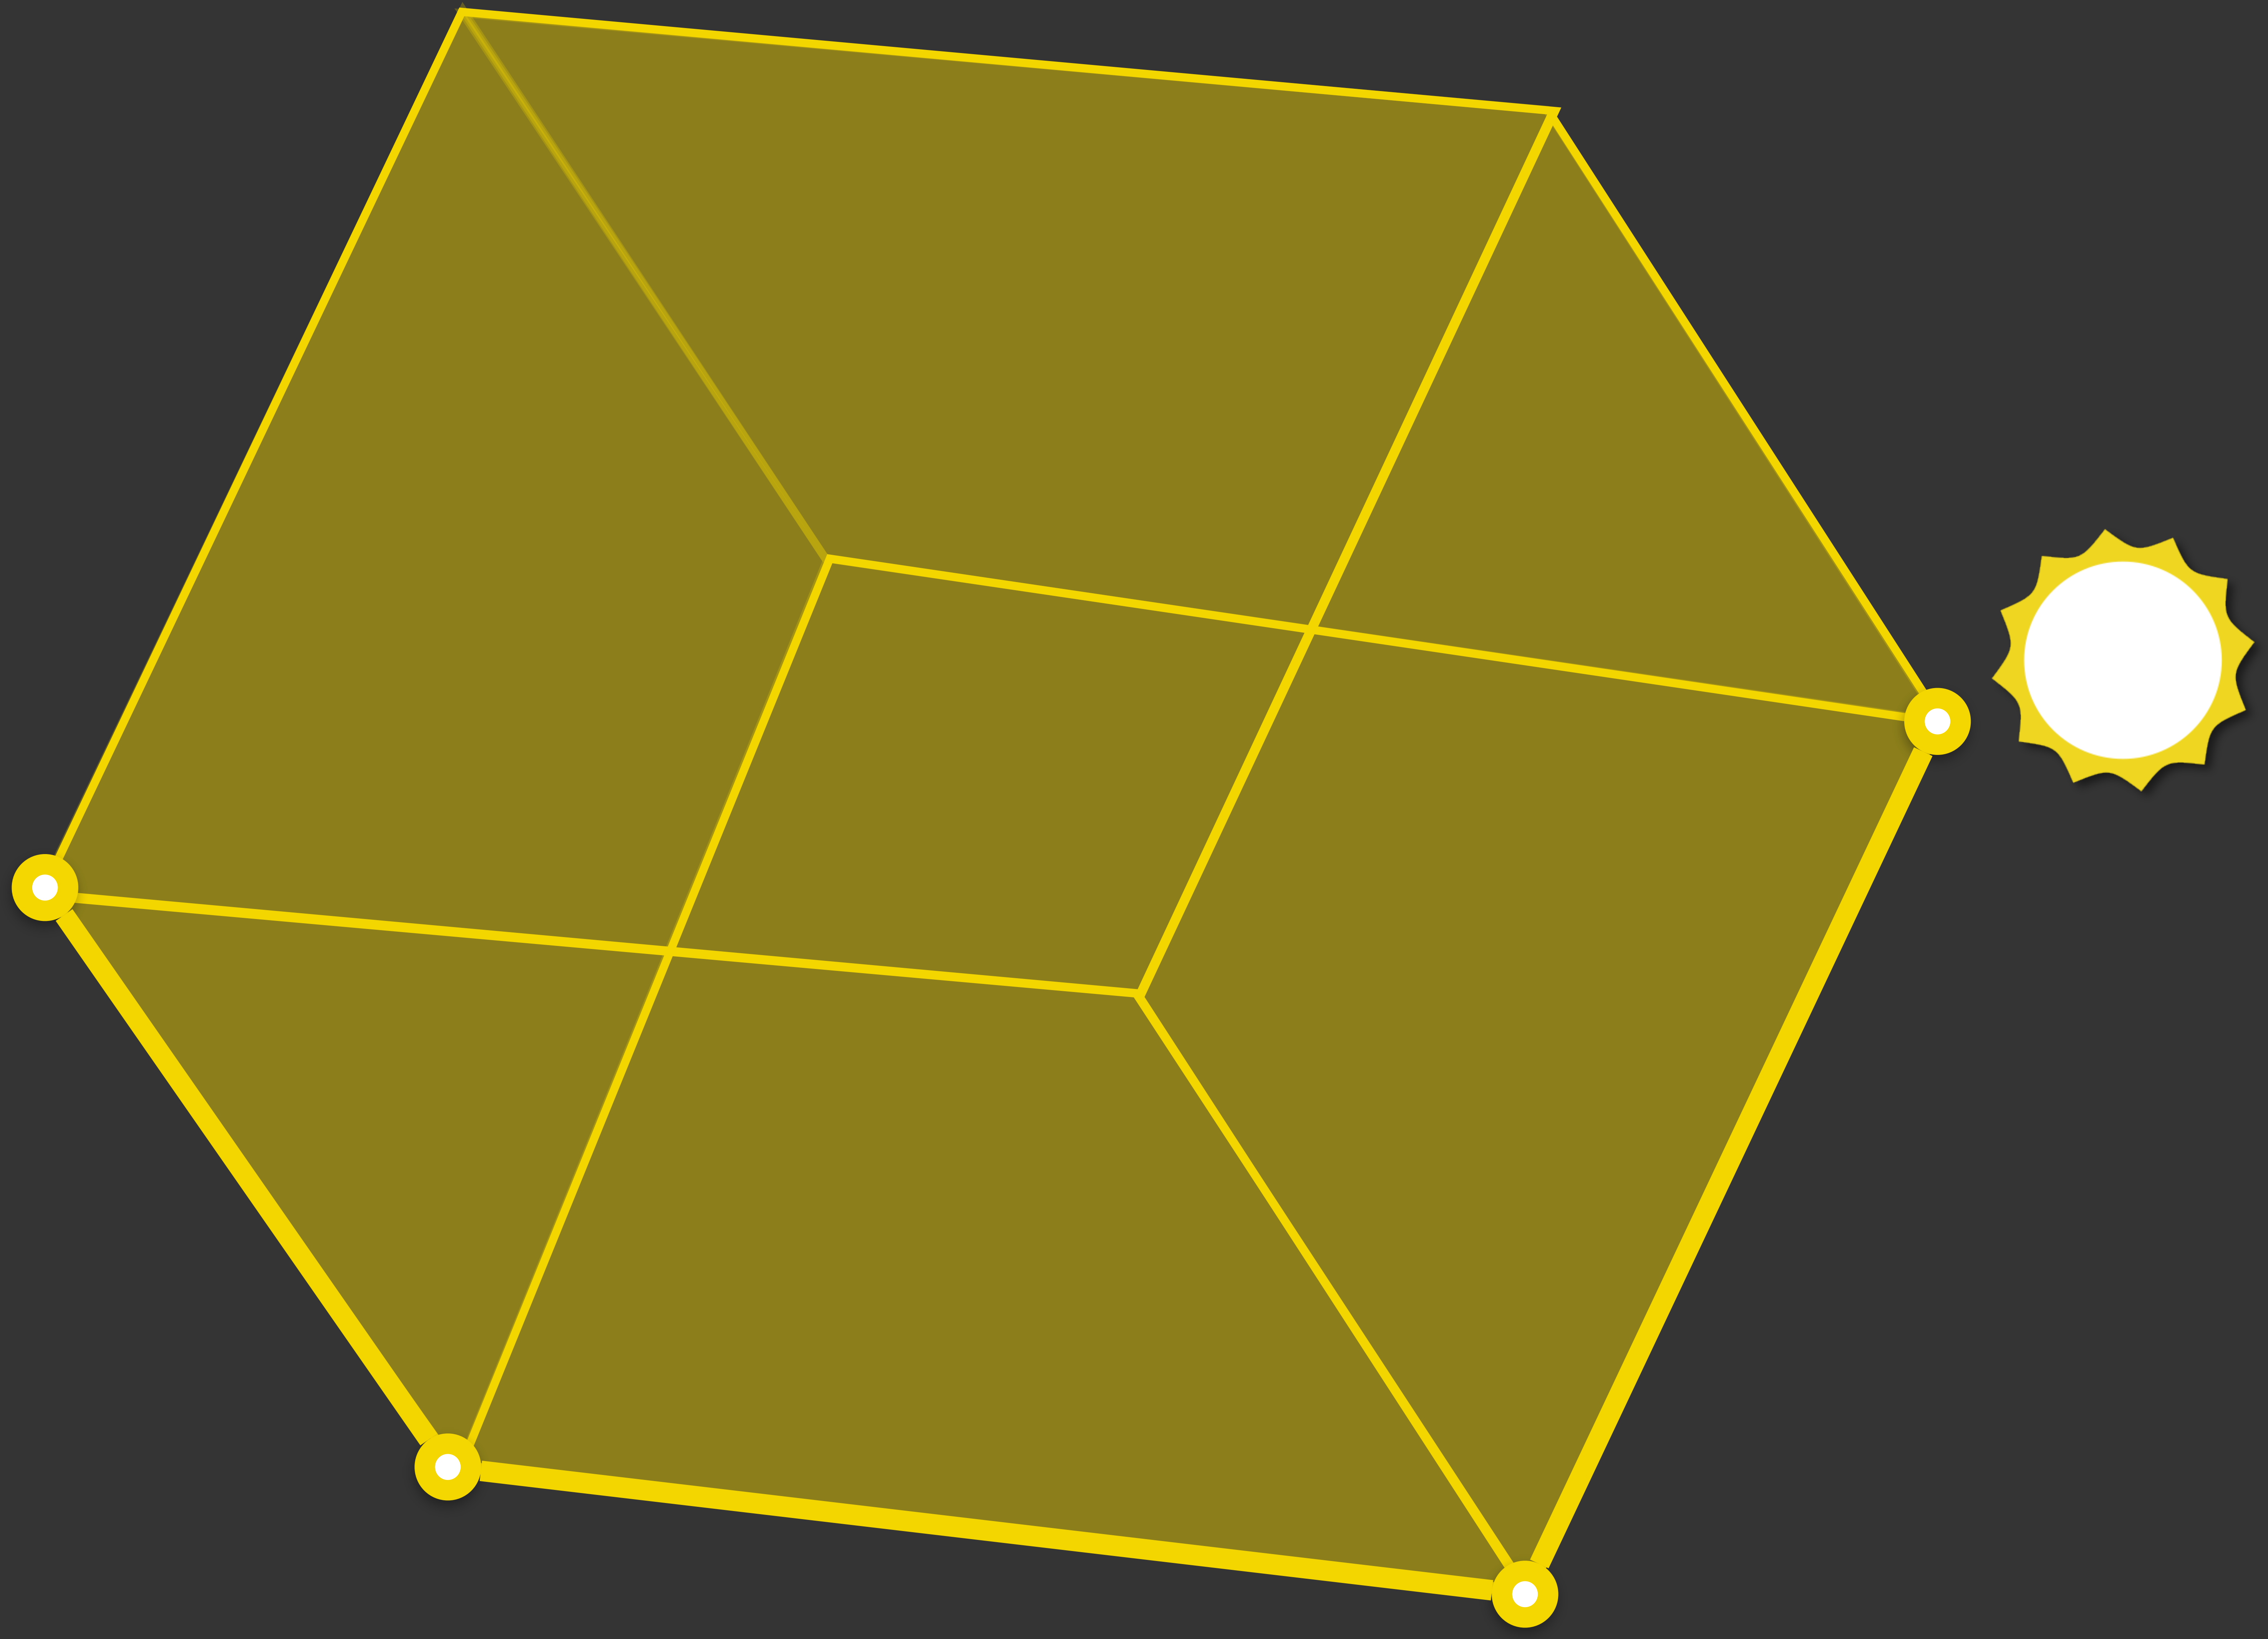

About Marching

About Marching

- Need careful photon arrangement to obtain limit

About Marching

- Need careful photon arrangement to obtain limit
- Arrangement introduces Jacobian term
  - Represents photon “squishing” and “stretching”

About Marching

- Need careful photon arrangement to obtain limit
- Arrangement introduces Jacobian term
  - Represents photon “squishing” and “stretching”
- Details: See paper

About Bias

About Bias

- Replacing distance sampling decreases bias

About Bias

- Replacing distance sampling decreases bias

Photon Points	3D Blur
Beam Radiance Estimate	2D Blur
Photon Beams	1D Blur
Photon Planes	0D Blur
Photon Volumes	0D Blur

About Bias

- Replacing distance sampling decreases bias
- Planes and beyond: *Unbiased*

	Photon Points	3D Blur
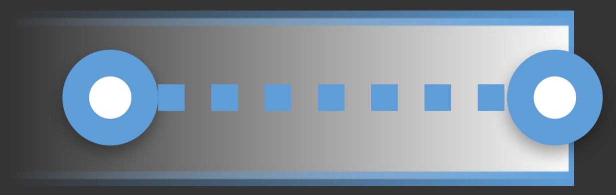	Beam Radiance Estimate	2D Blur
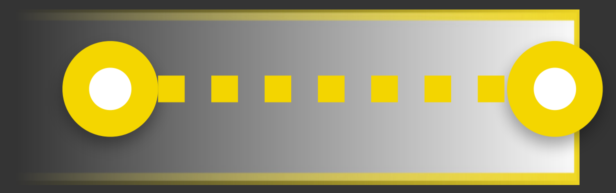	Photon Beams	1D Blur
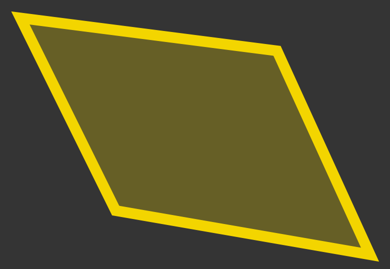	Photon Planes	0D Blur
	Photon Volumes	0D Blur

About Bias

- Replacing distance sampling decreases bias
- Planes and beyond: *Unbiased*

About Bias

- Replacing distance sampling decreases bias
- Planes and beyond: *Unbiased*
- But: Bias $\leftrightarrow$ Variance tradeoff

About Bias

- Replacing distance sampling decreases bias
- Planes and beyond: *Unbiased*
- But: Bias $\leftrightarrow$ Variance tradeoff

- In paper: Planes (0D Blur)
  Planes (1D Blur)

Summary

Summary

- Previous work:

  Replace one collision with track-length/expected value

Summary

- Previous work:

  Replace one collision with track-length/expected value

- Our work:

  Repeat this process along preceding segments

Summary

- Previous work:

  Replace one collision with track-length/expected value

- Our work:

  Repeat this process along preceding segments

- Can do this for both photons and cameras

Summary

- Previous work:

  Replace one collision with track-length/expected value

- Our work:

  Repeat this process along preceding segments

- Can do this for both photons and cameras
- These new estimators are *unbiased*

Error Analysis

Error Analysis

- Analytic bias & variance of 27 different photons

Error Analysis Setup

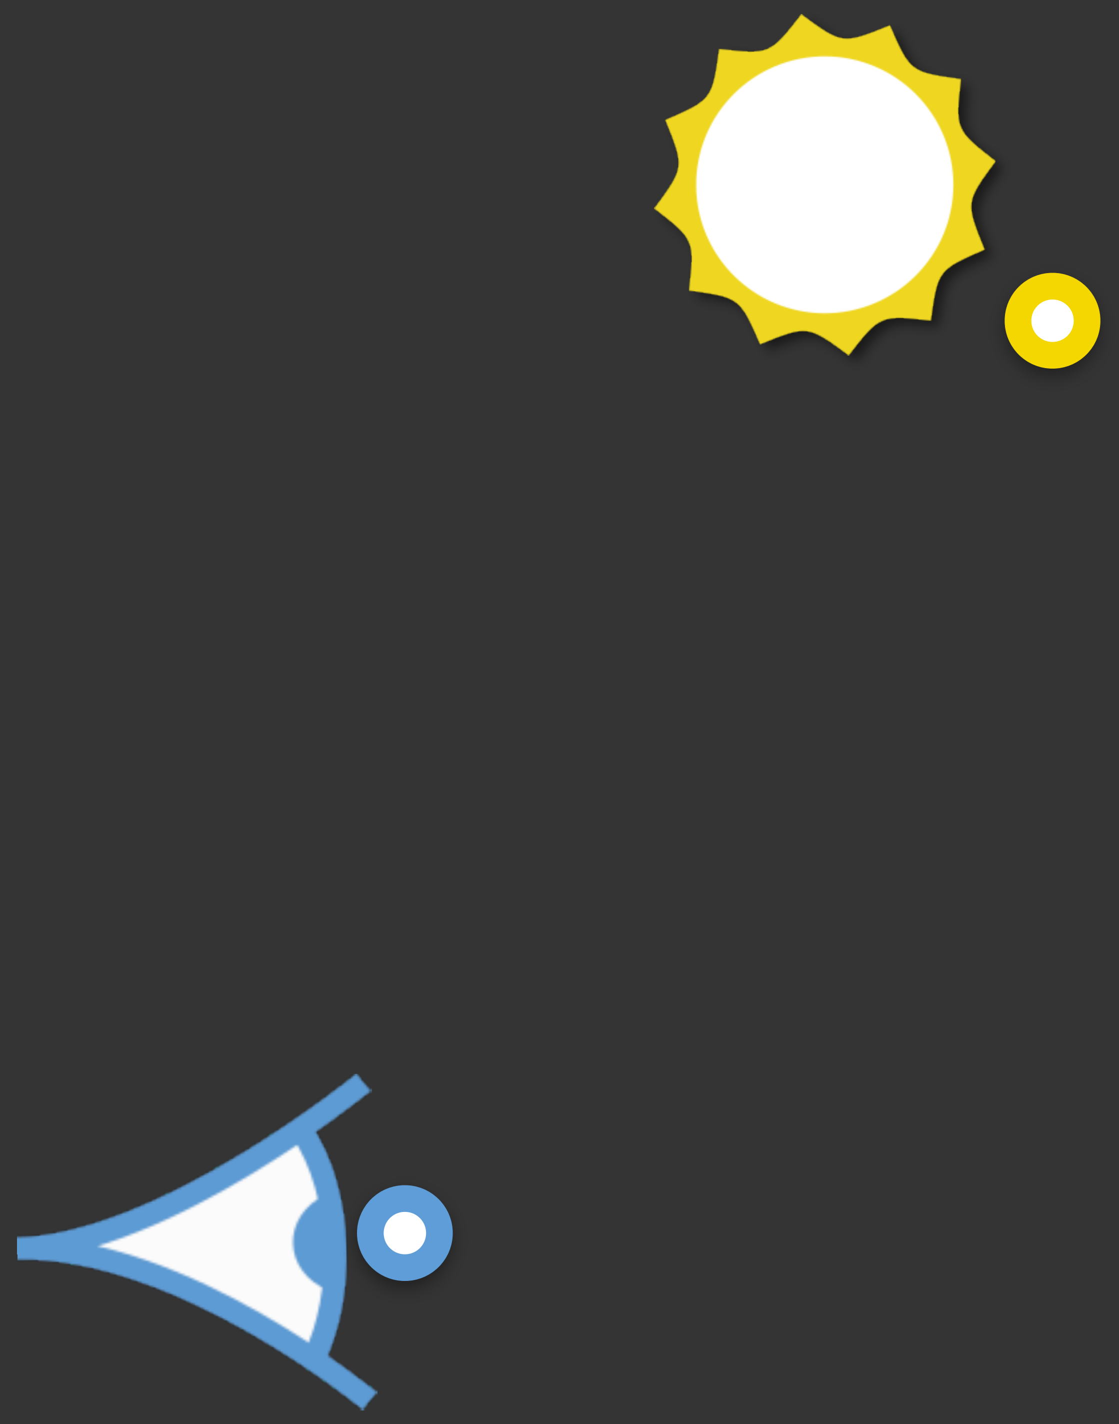

Error Analysis Setup

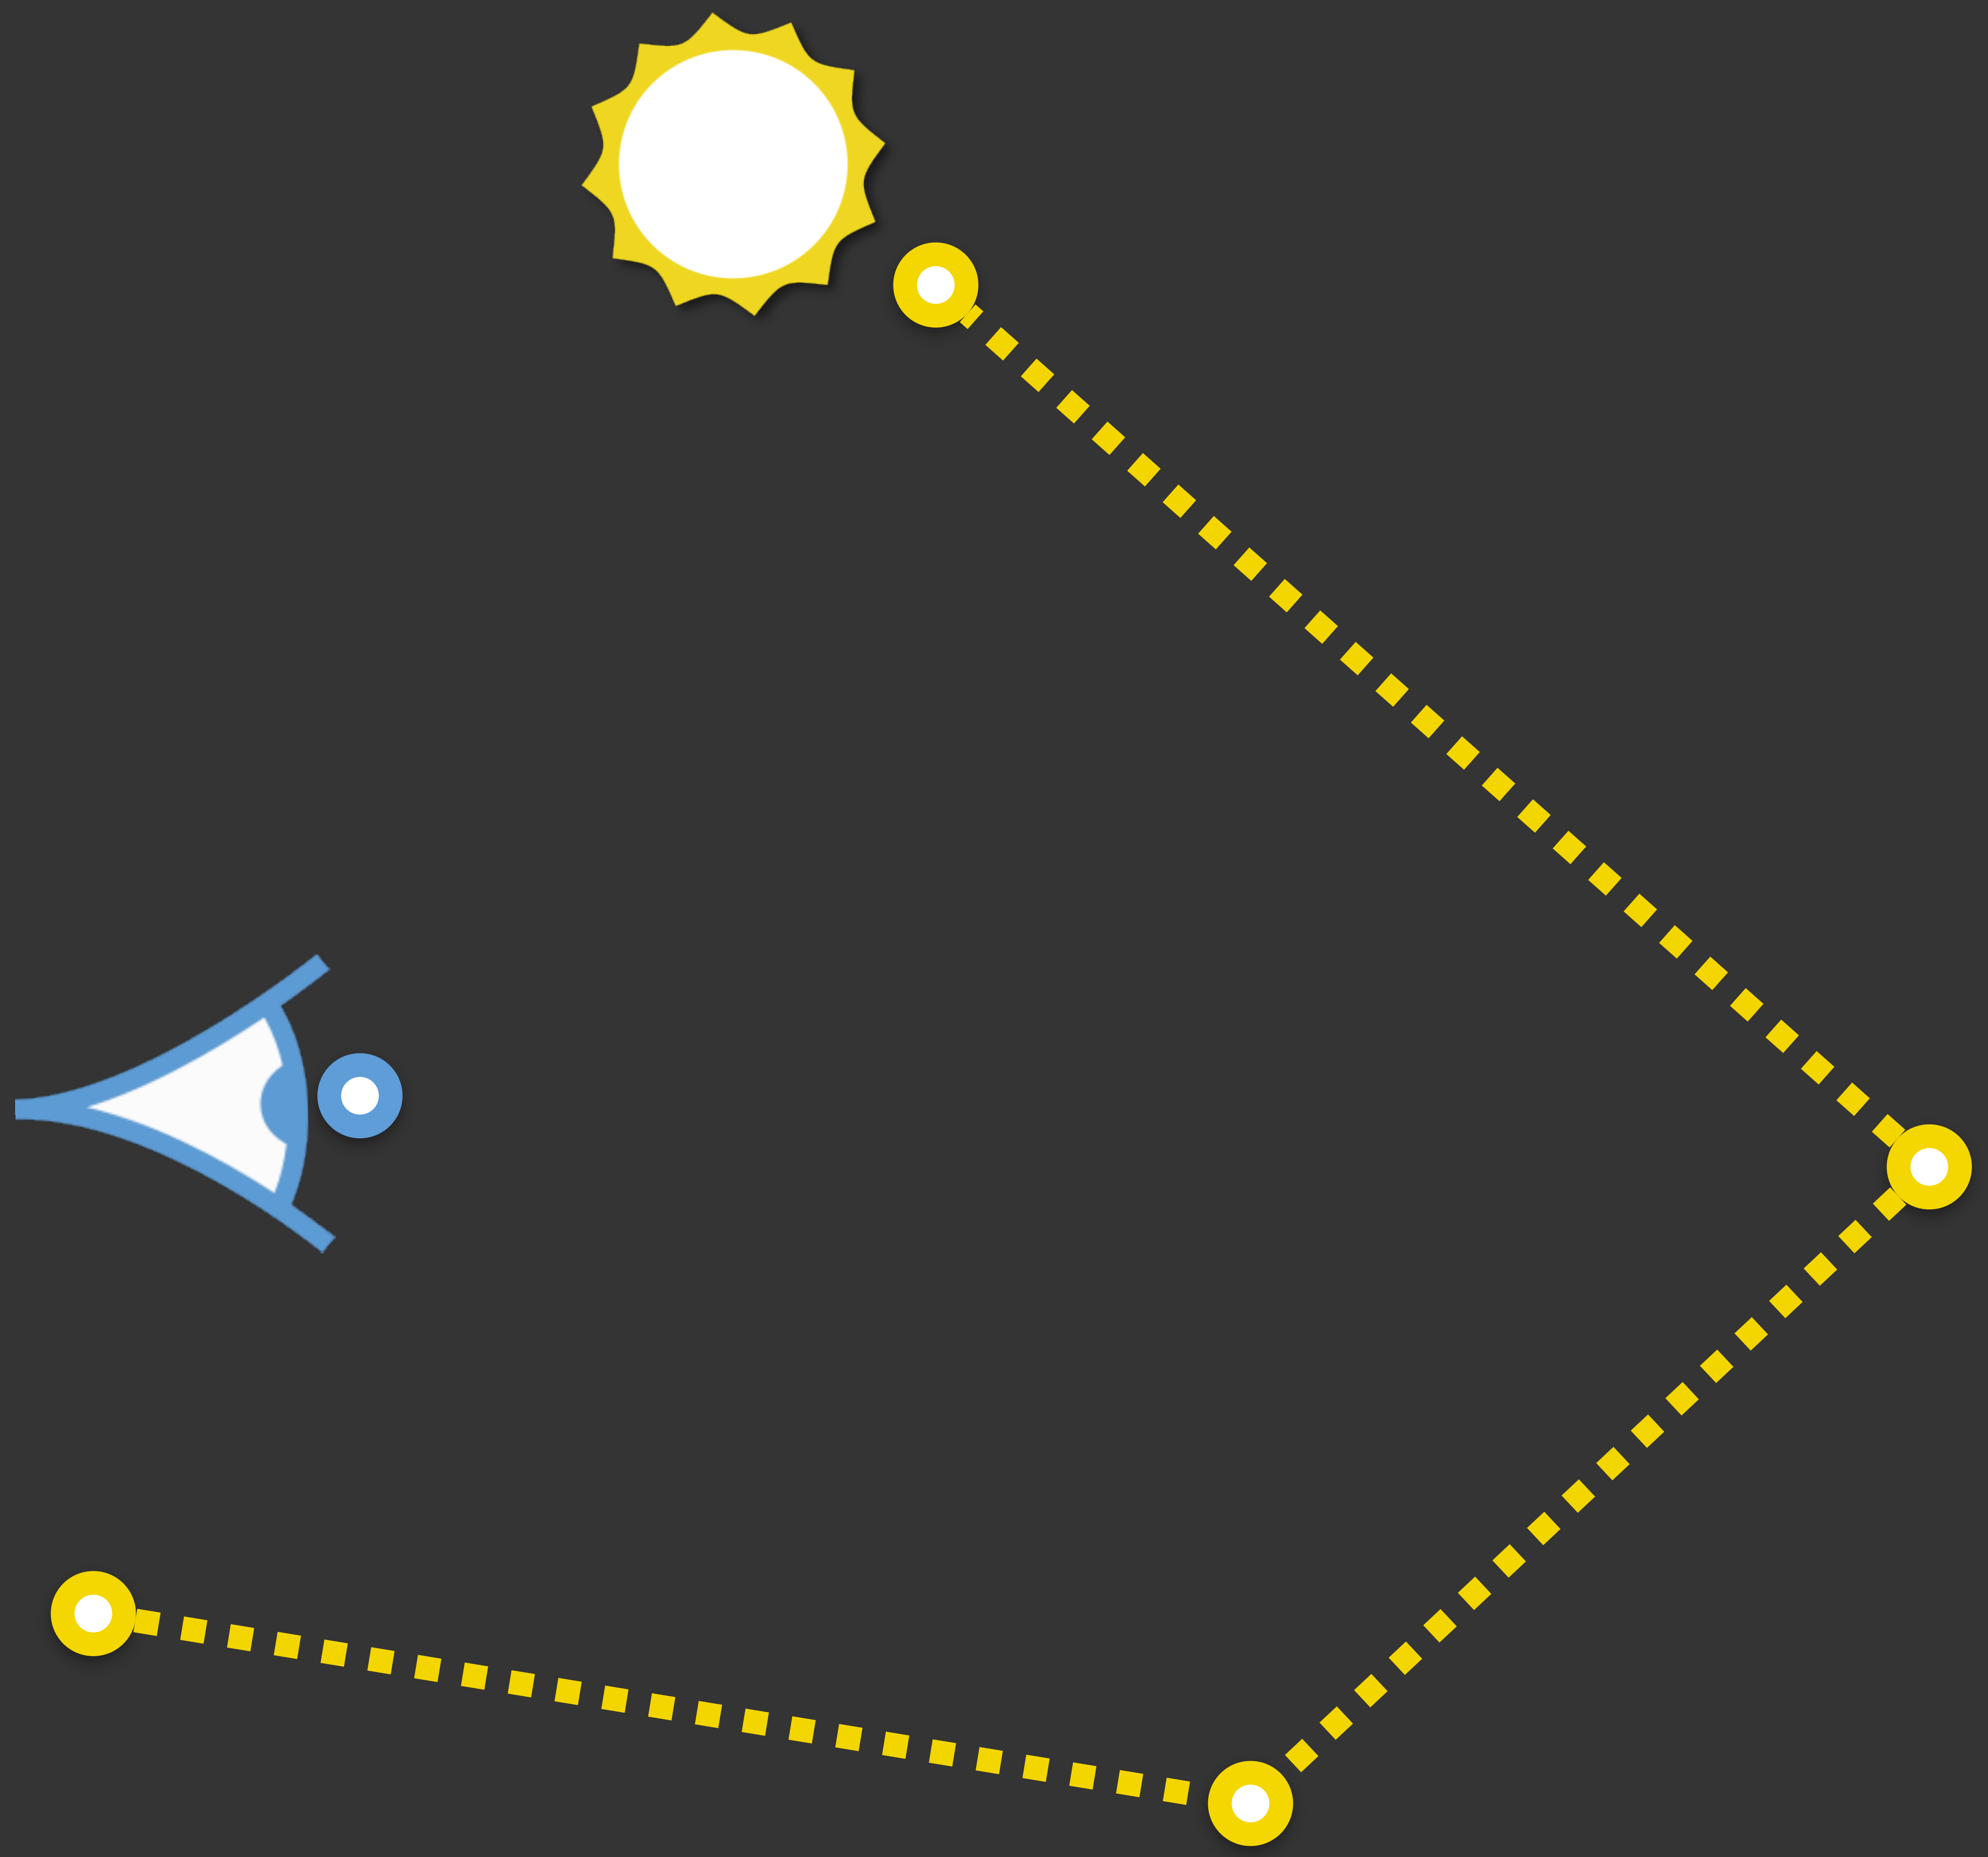

Error Analysis Setup

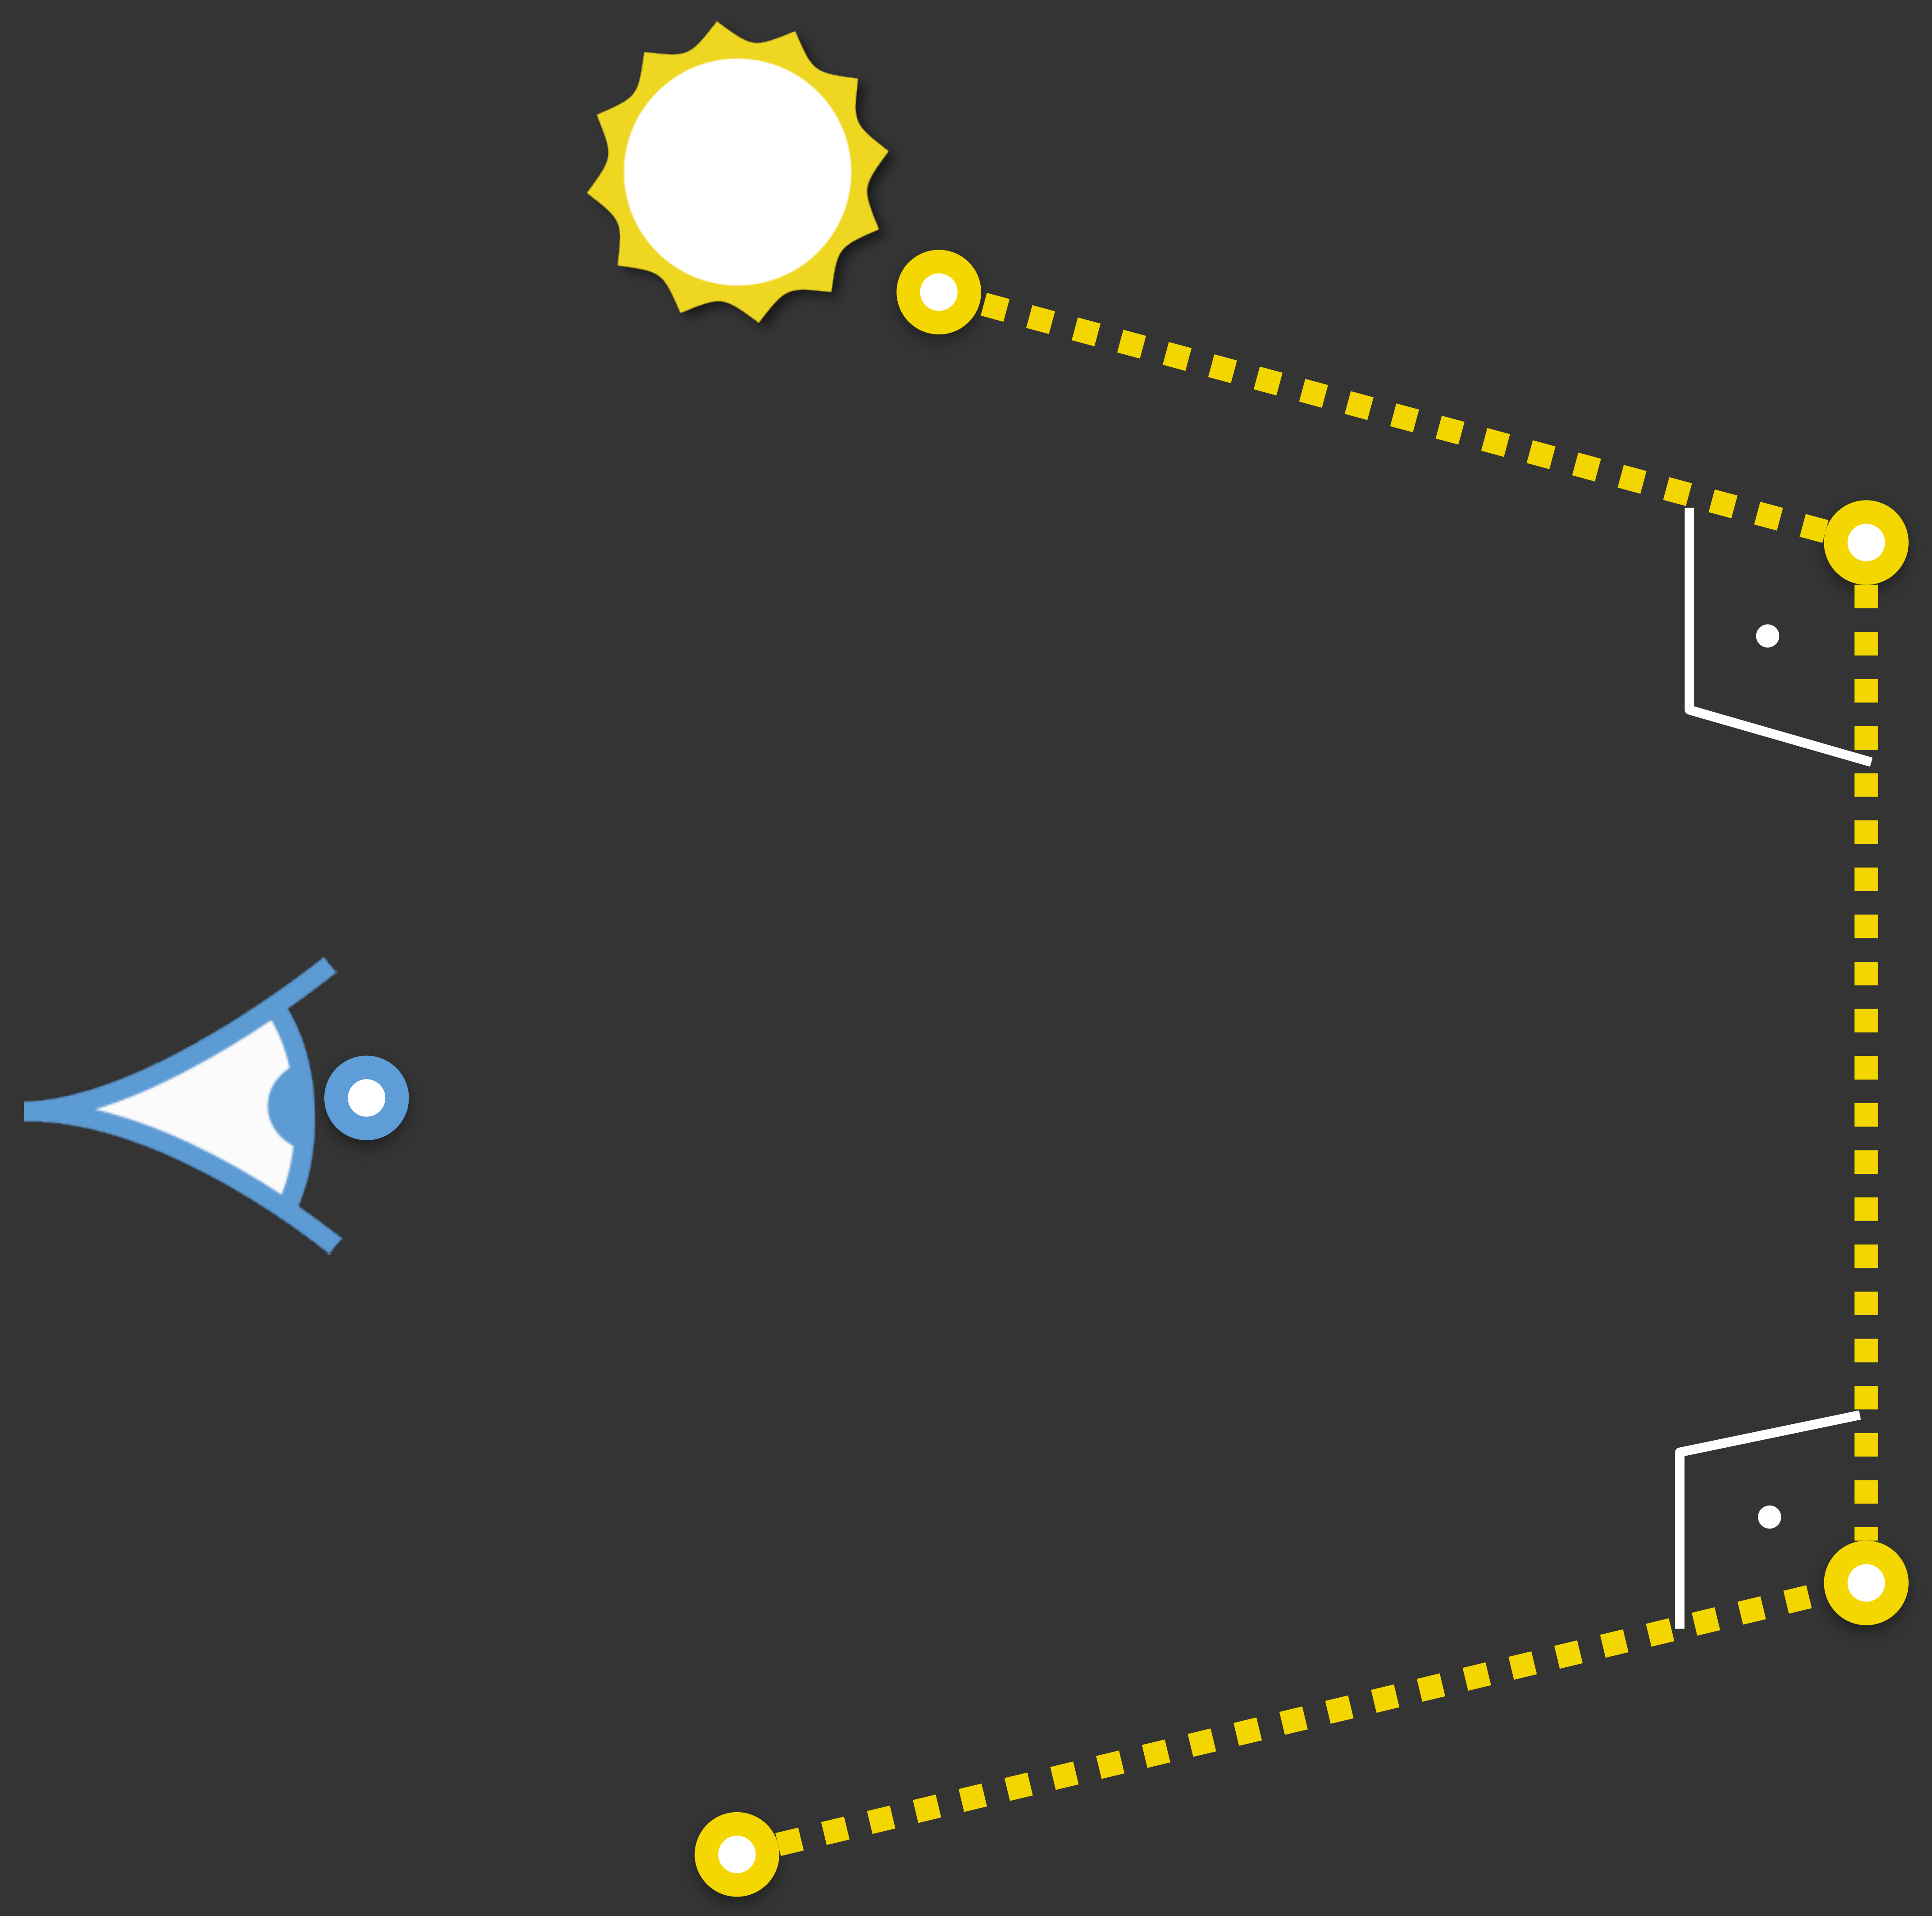

Error Analysis Setup

Křivánek et al. [2014]

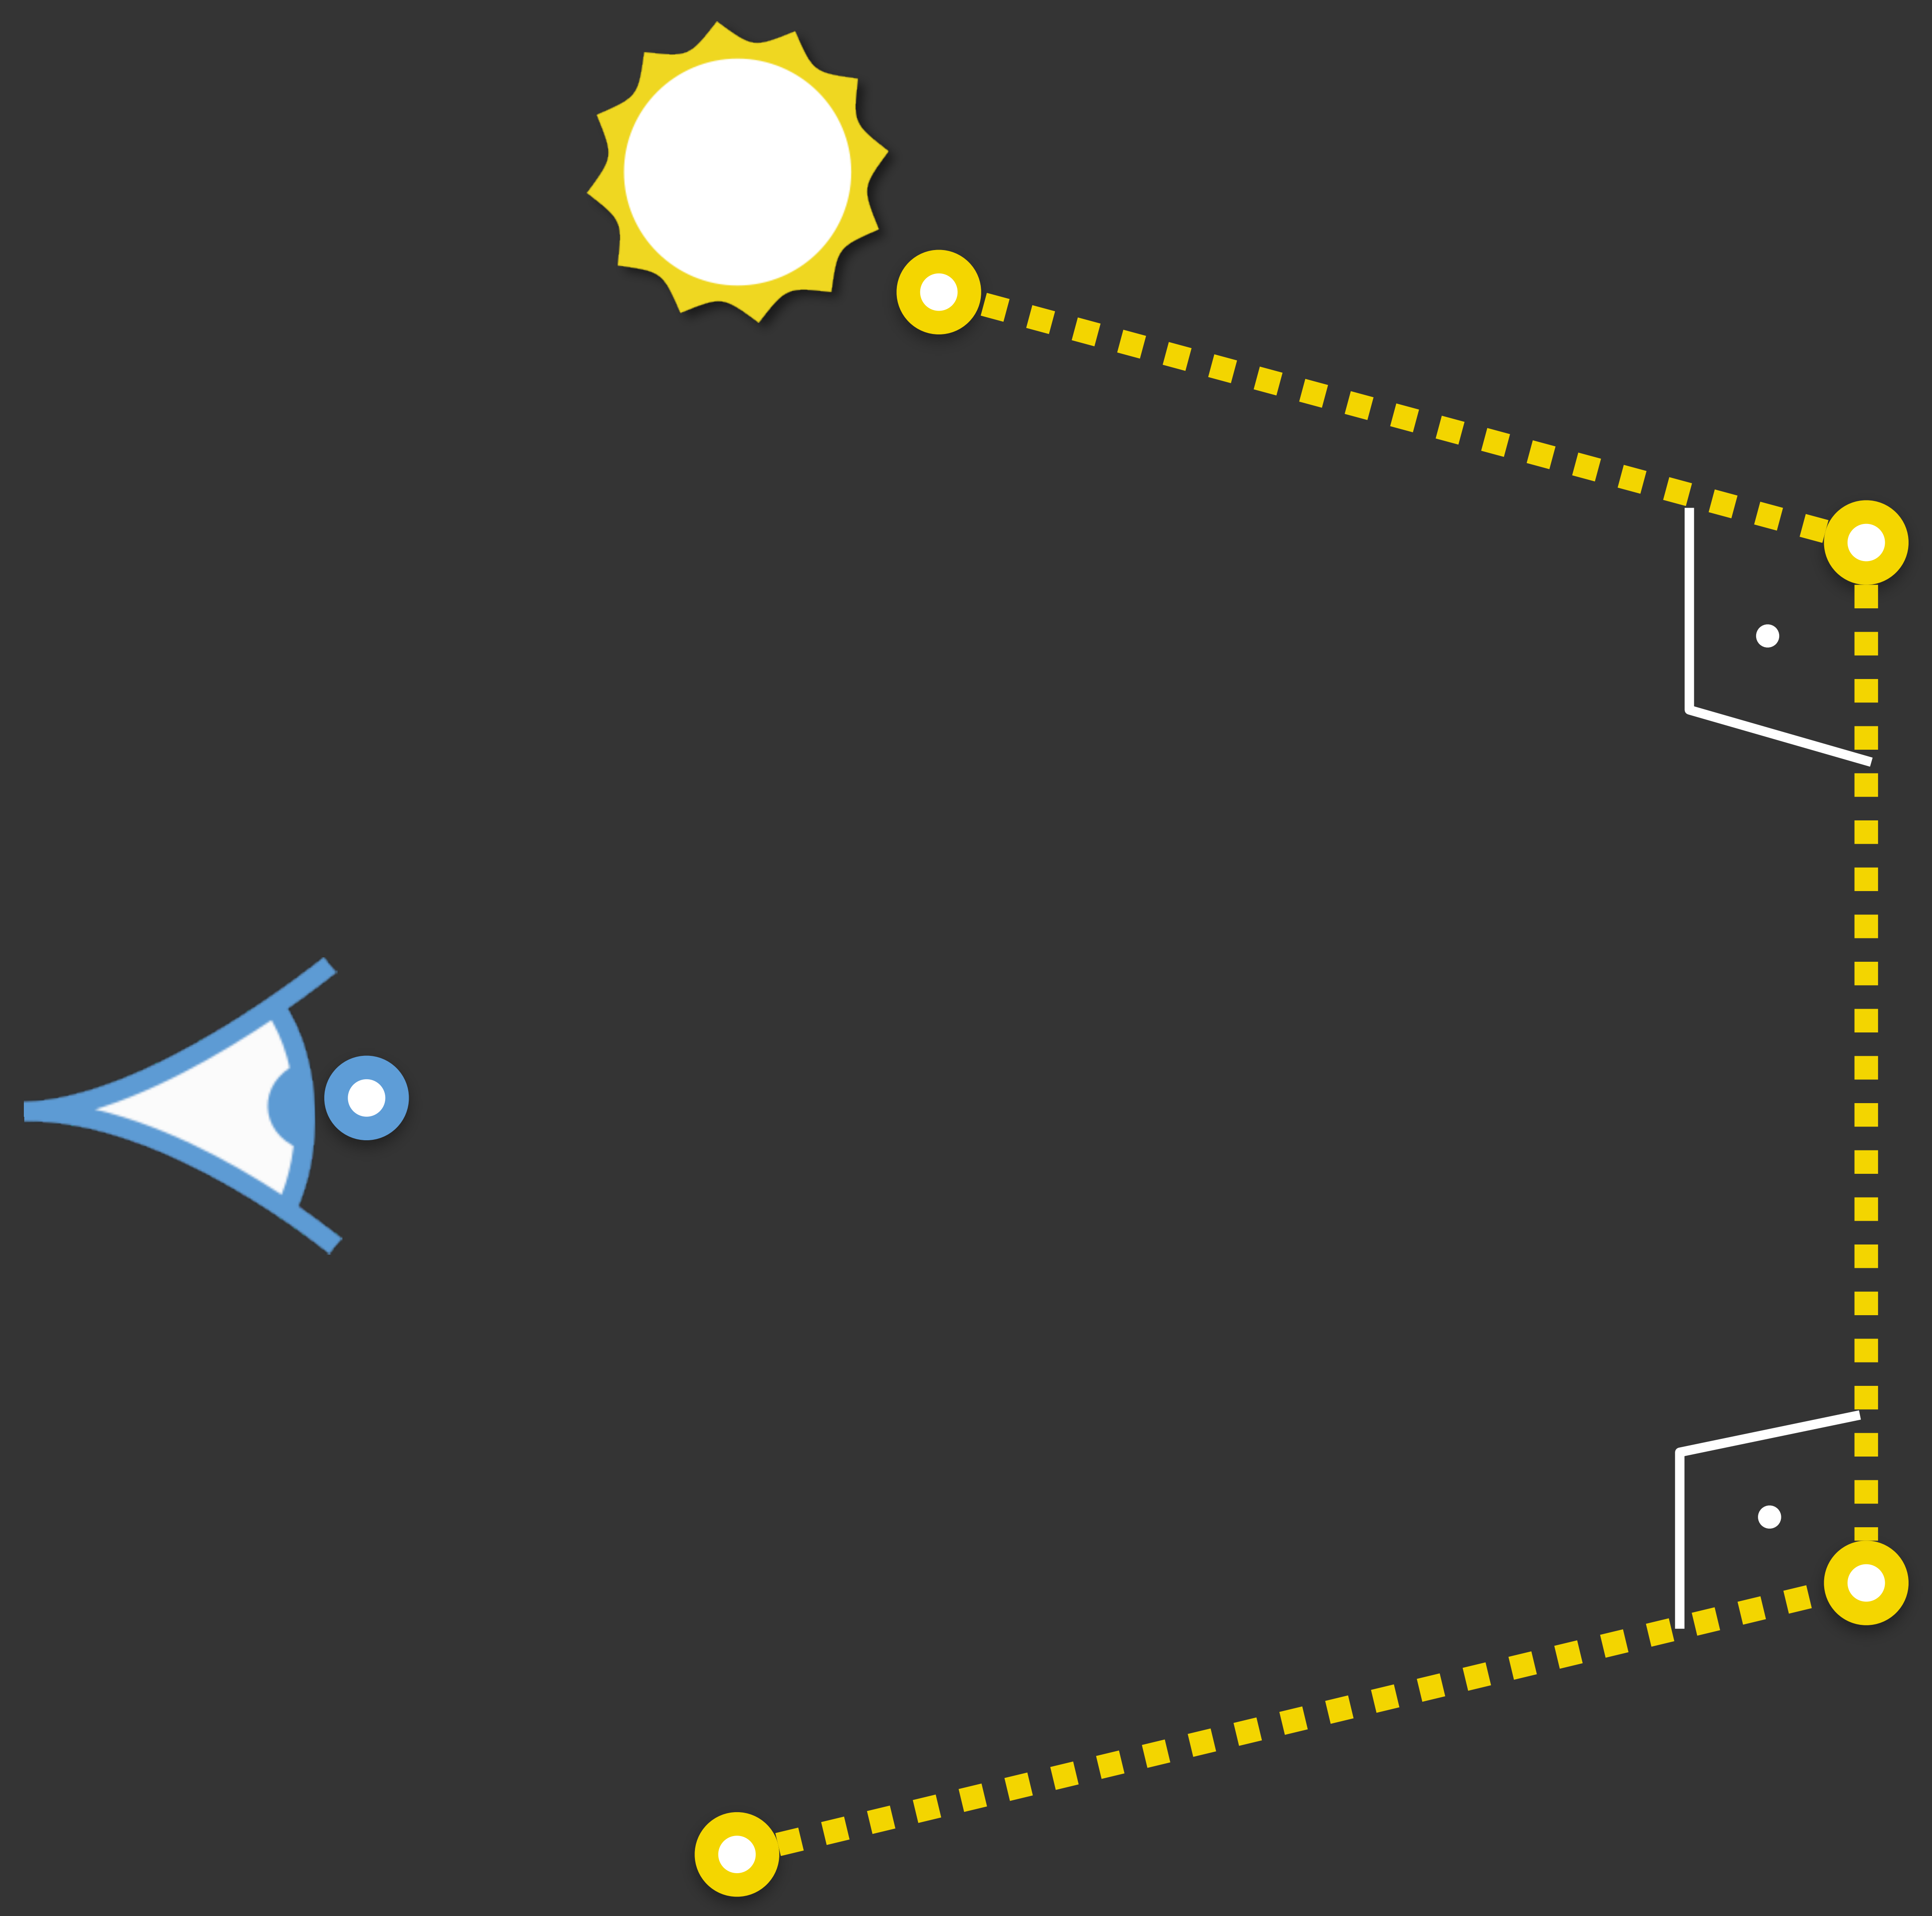

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

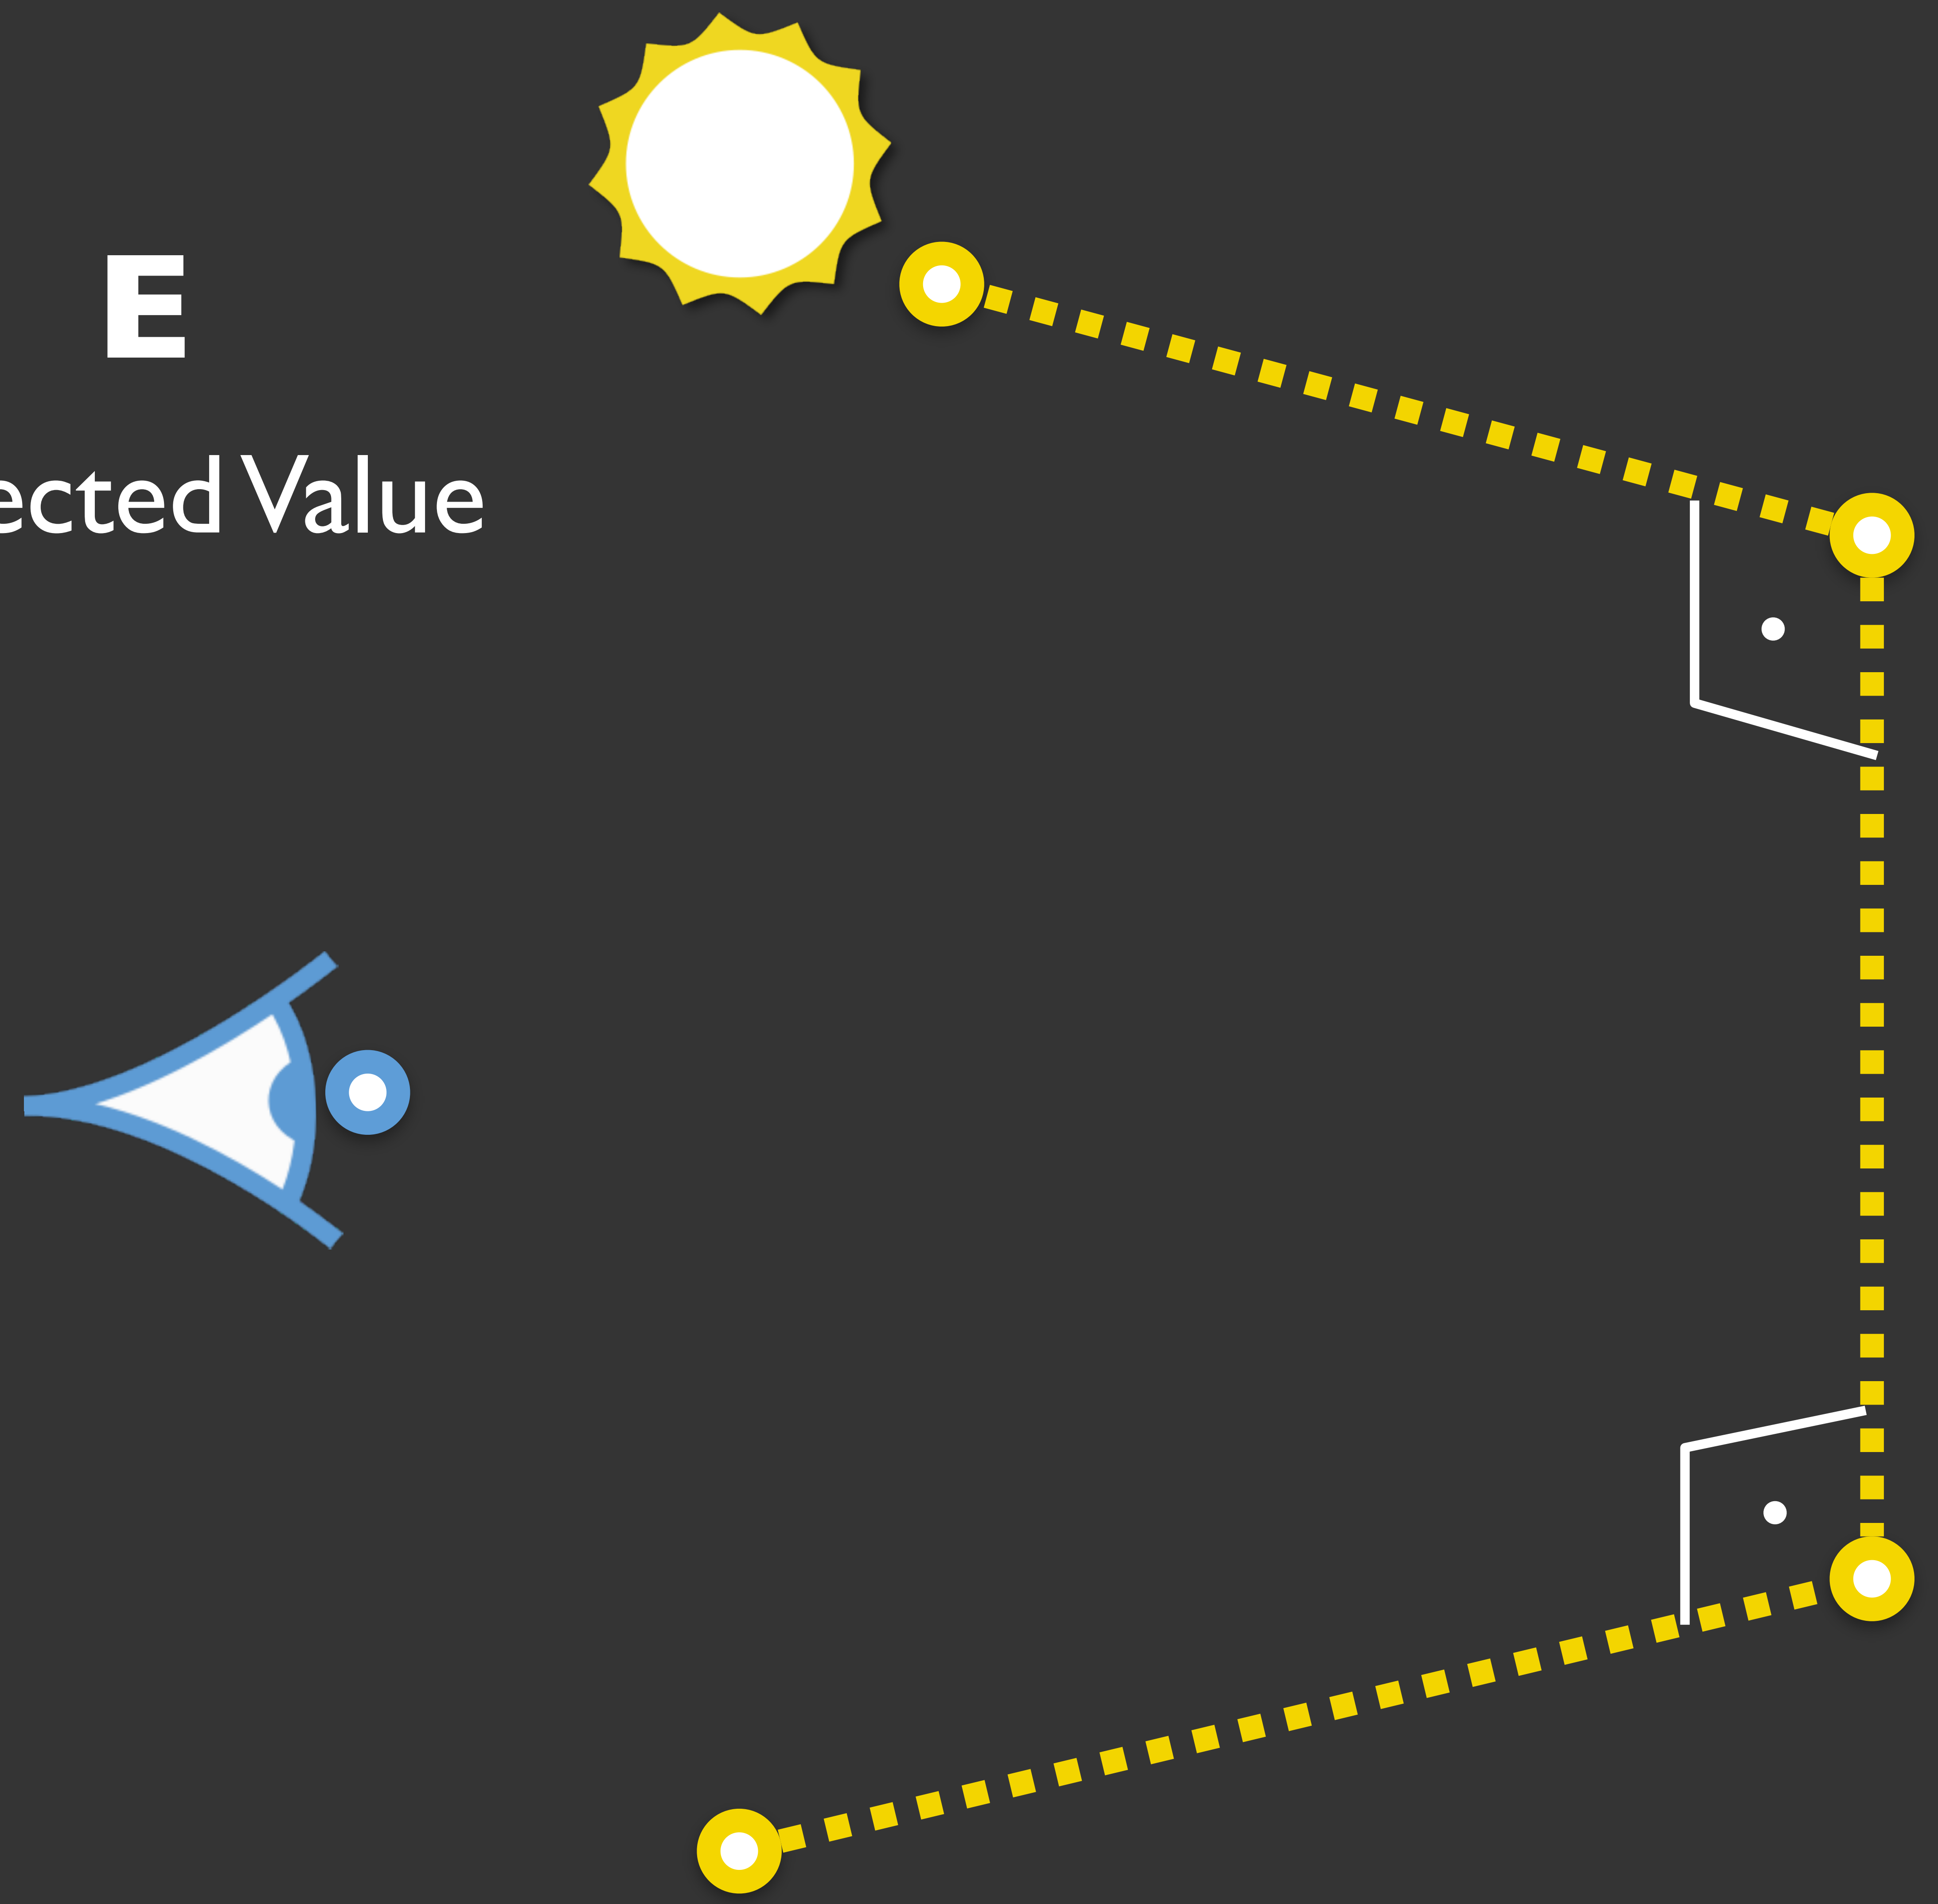

Error Analysis Setup

C Collision **T** Track-Length **E** Expected Value

C

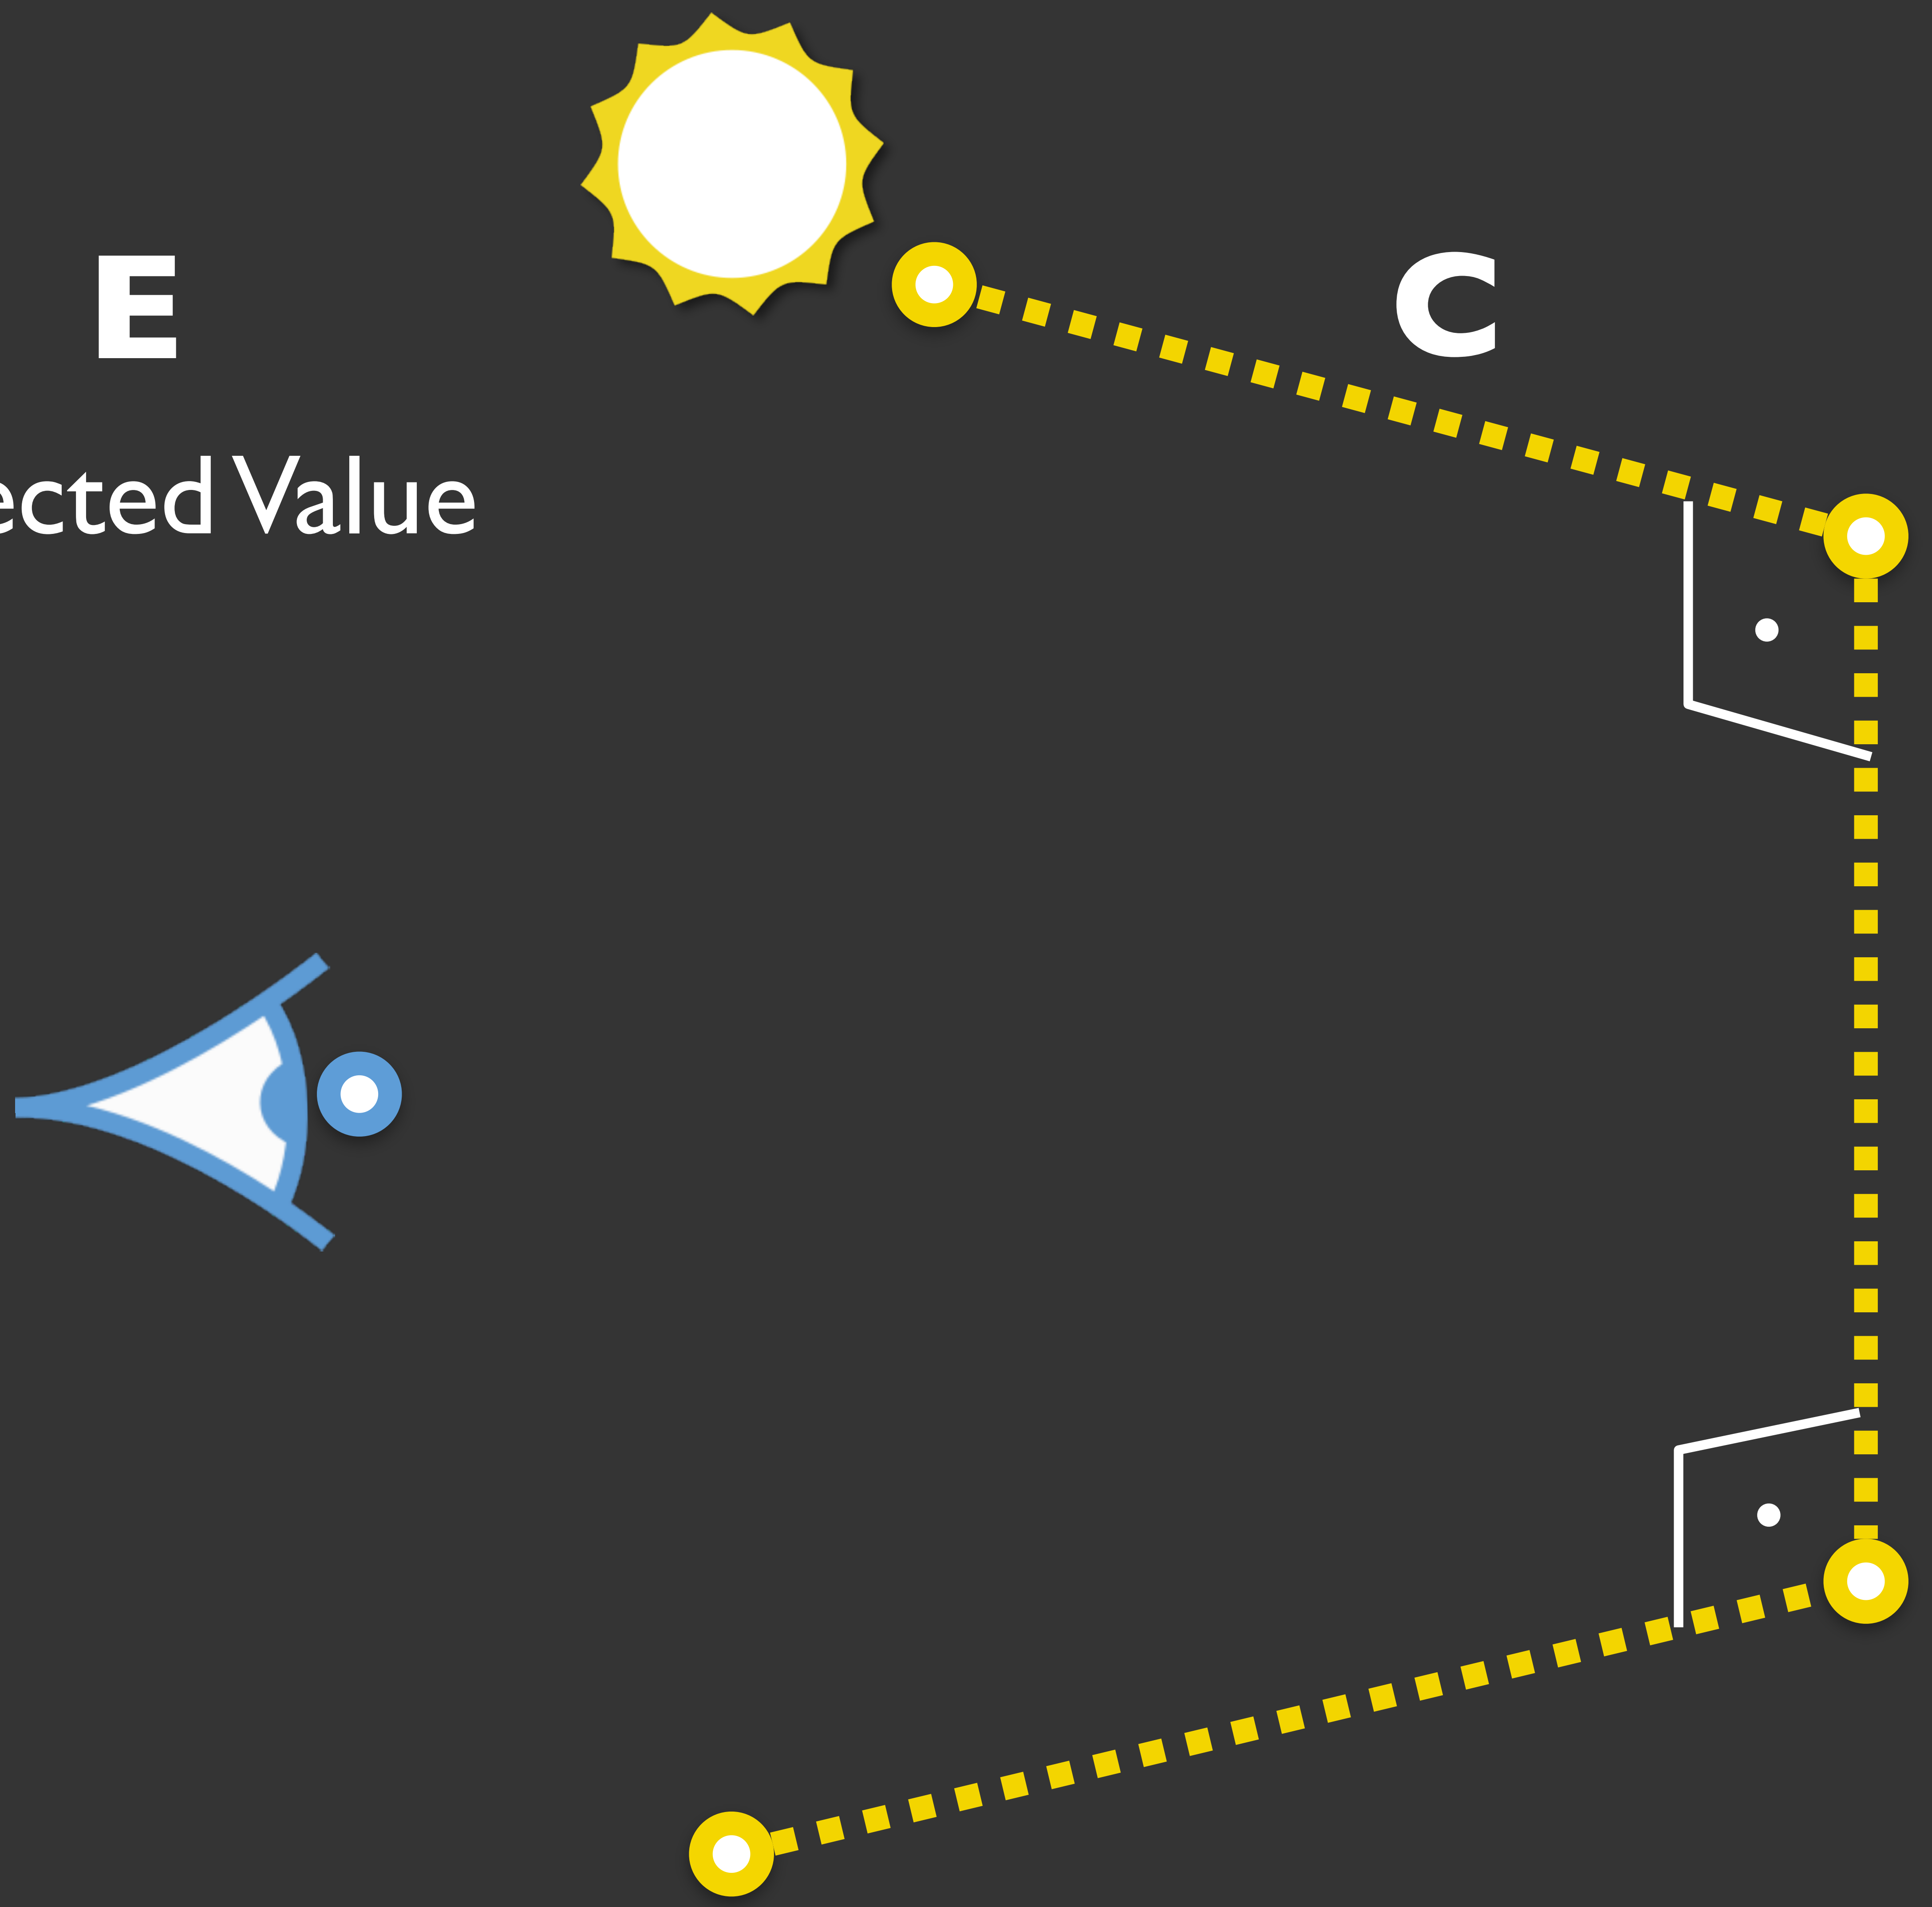

Error Analysis Setup

C Collision

T Track-Length

E Expected Value

CC

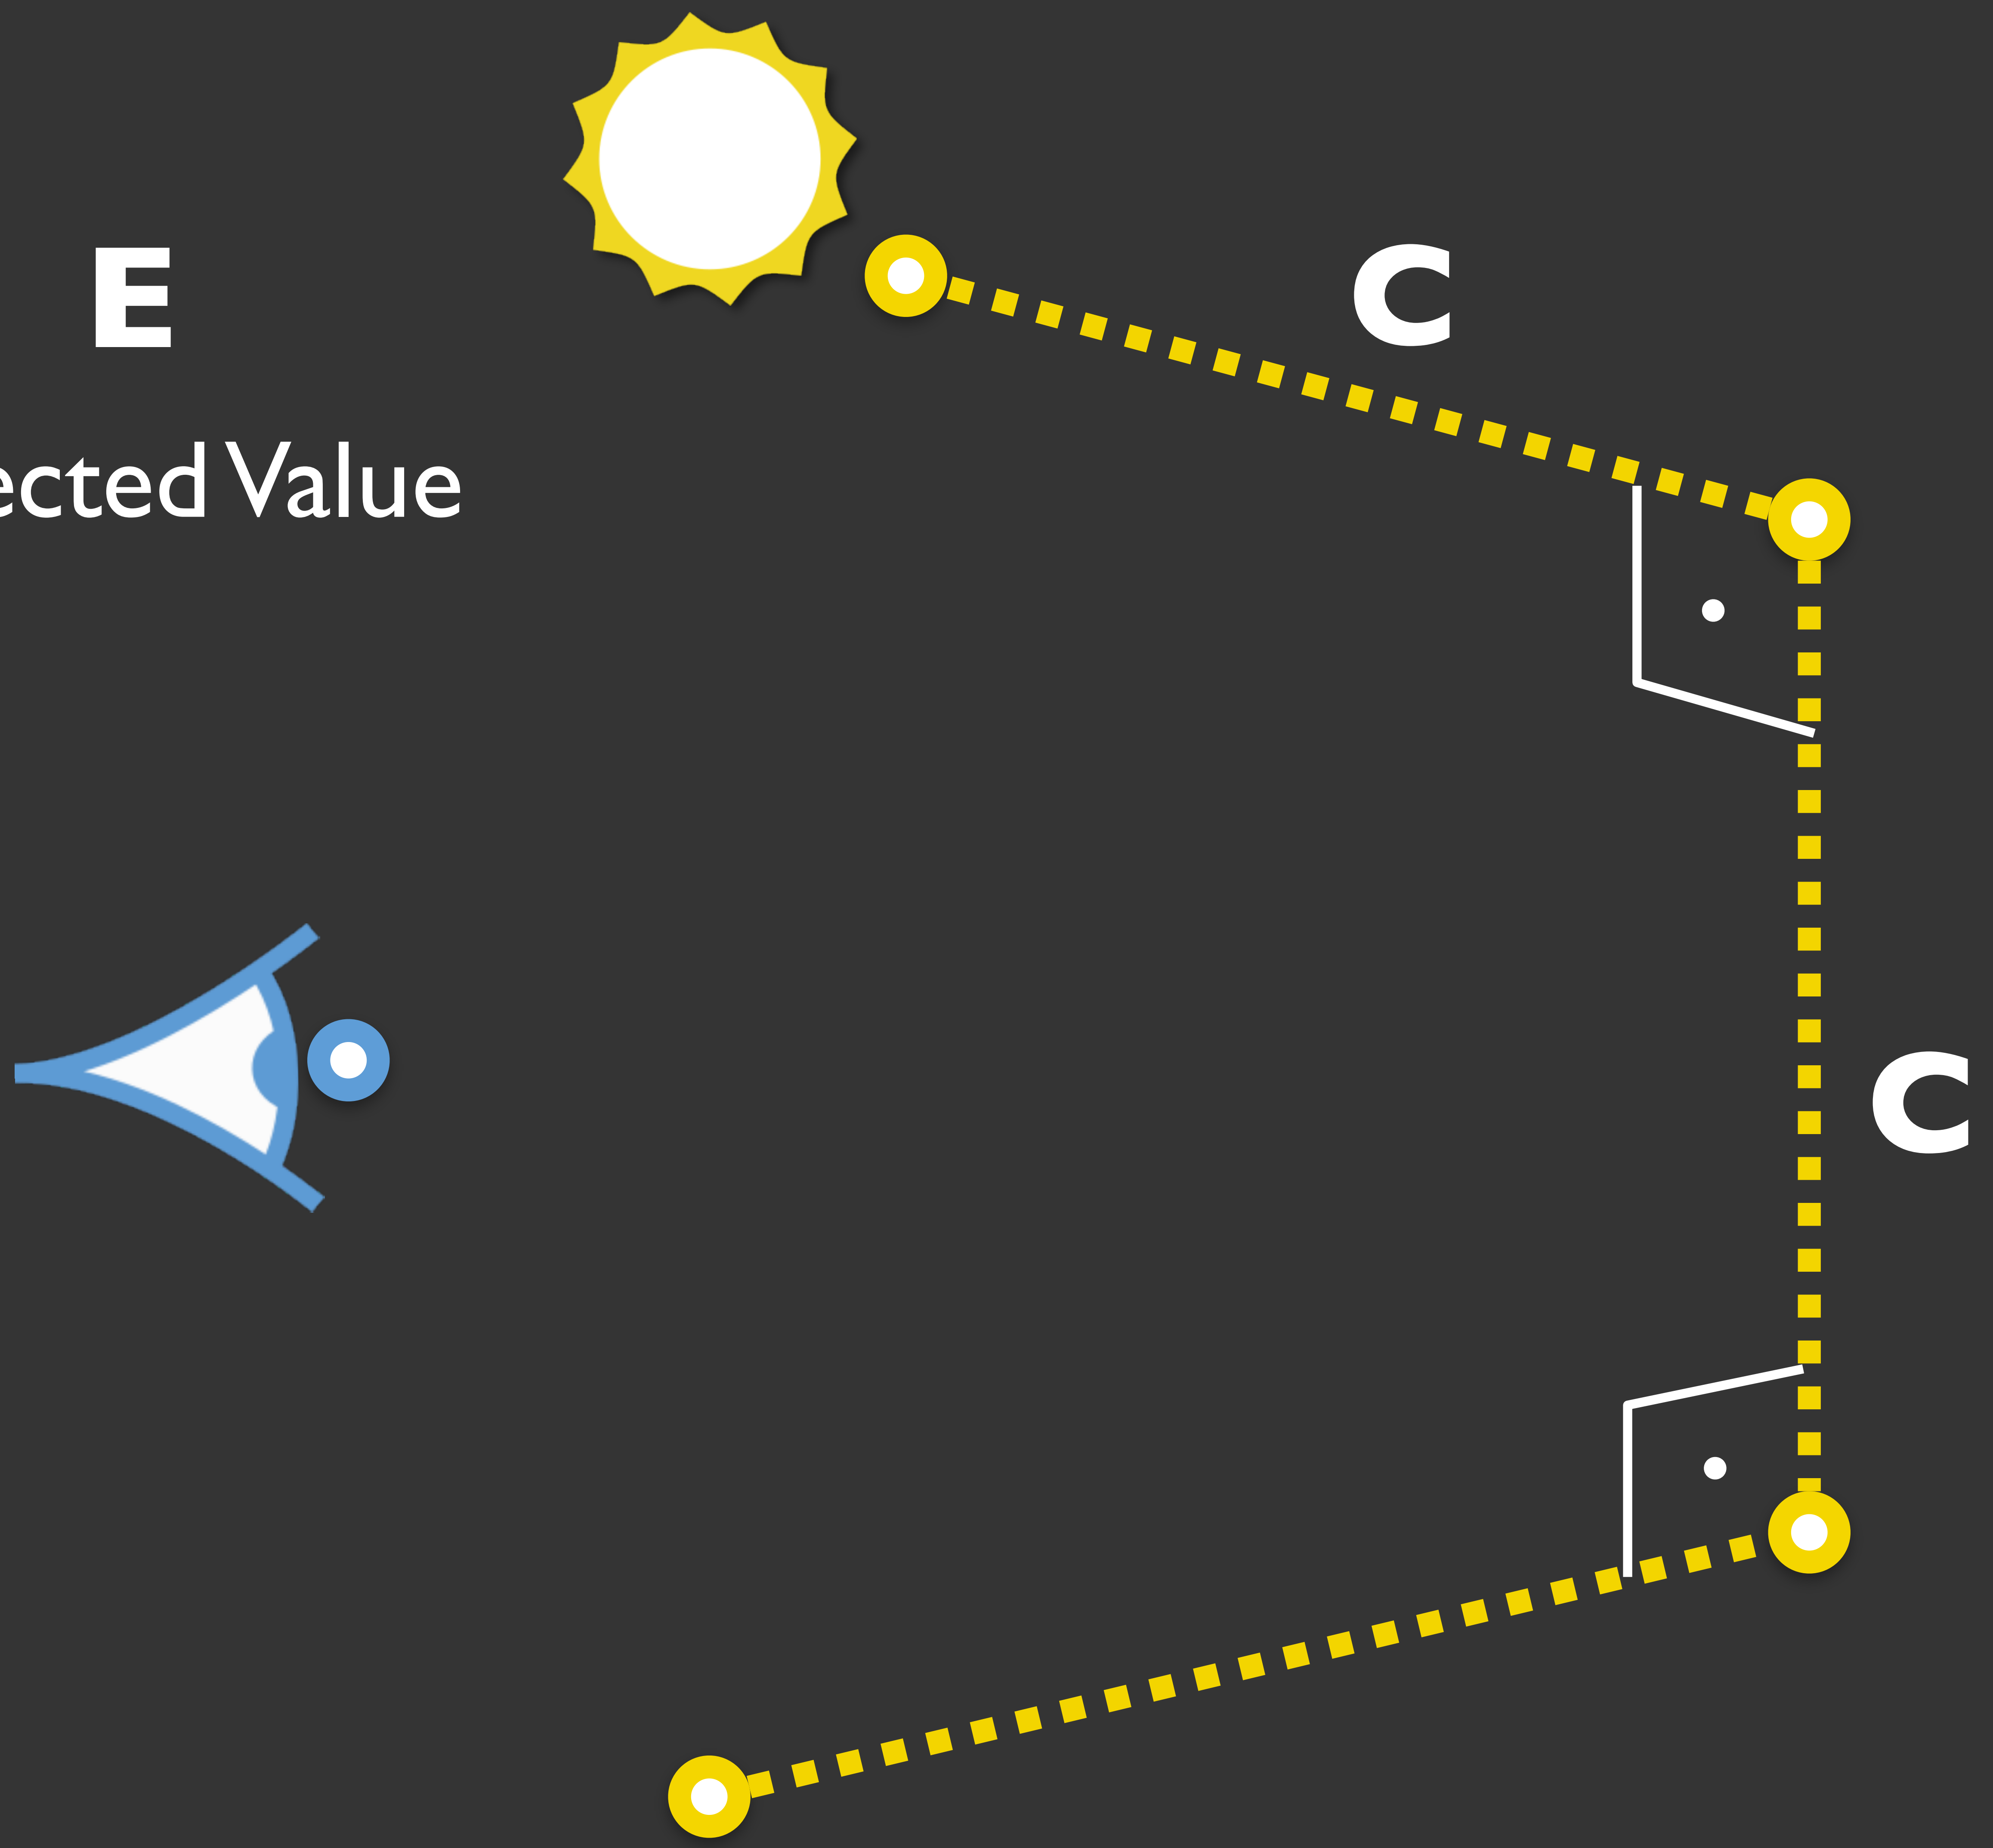

Error Analysis Setup

C Collision

T Track-Length

E Expected Value

CCC

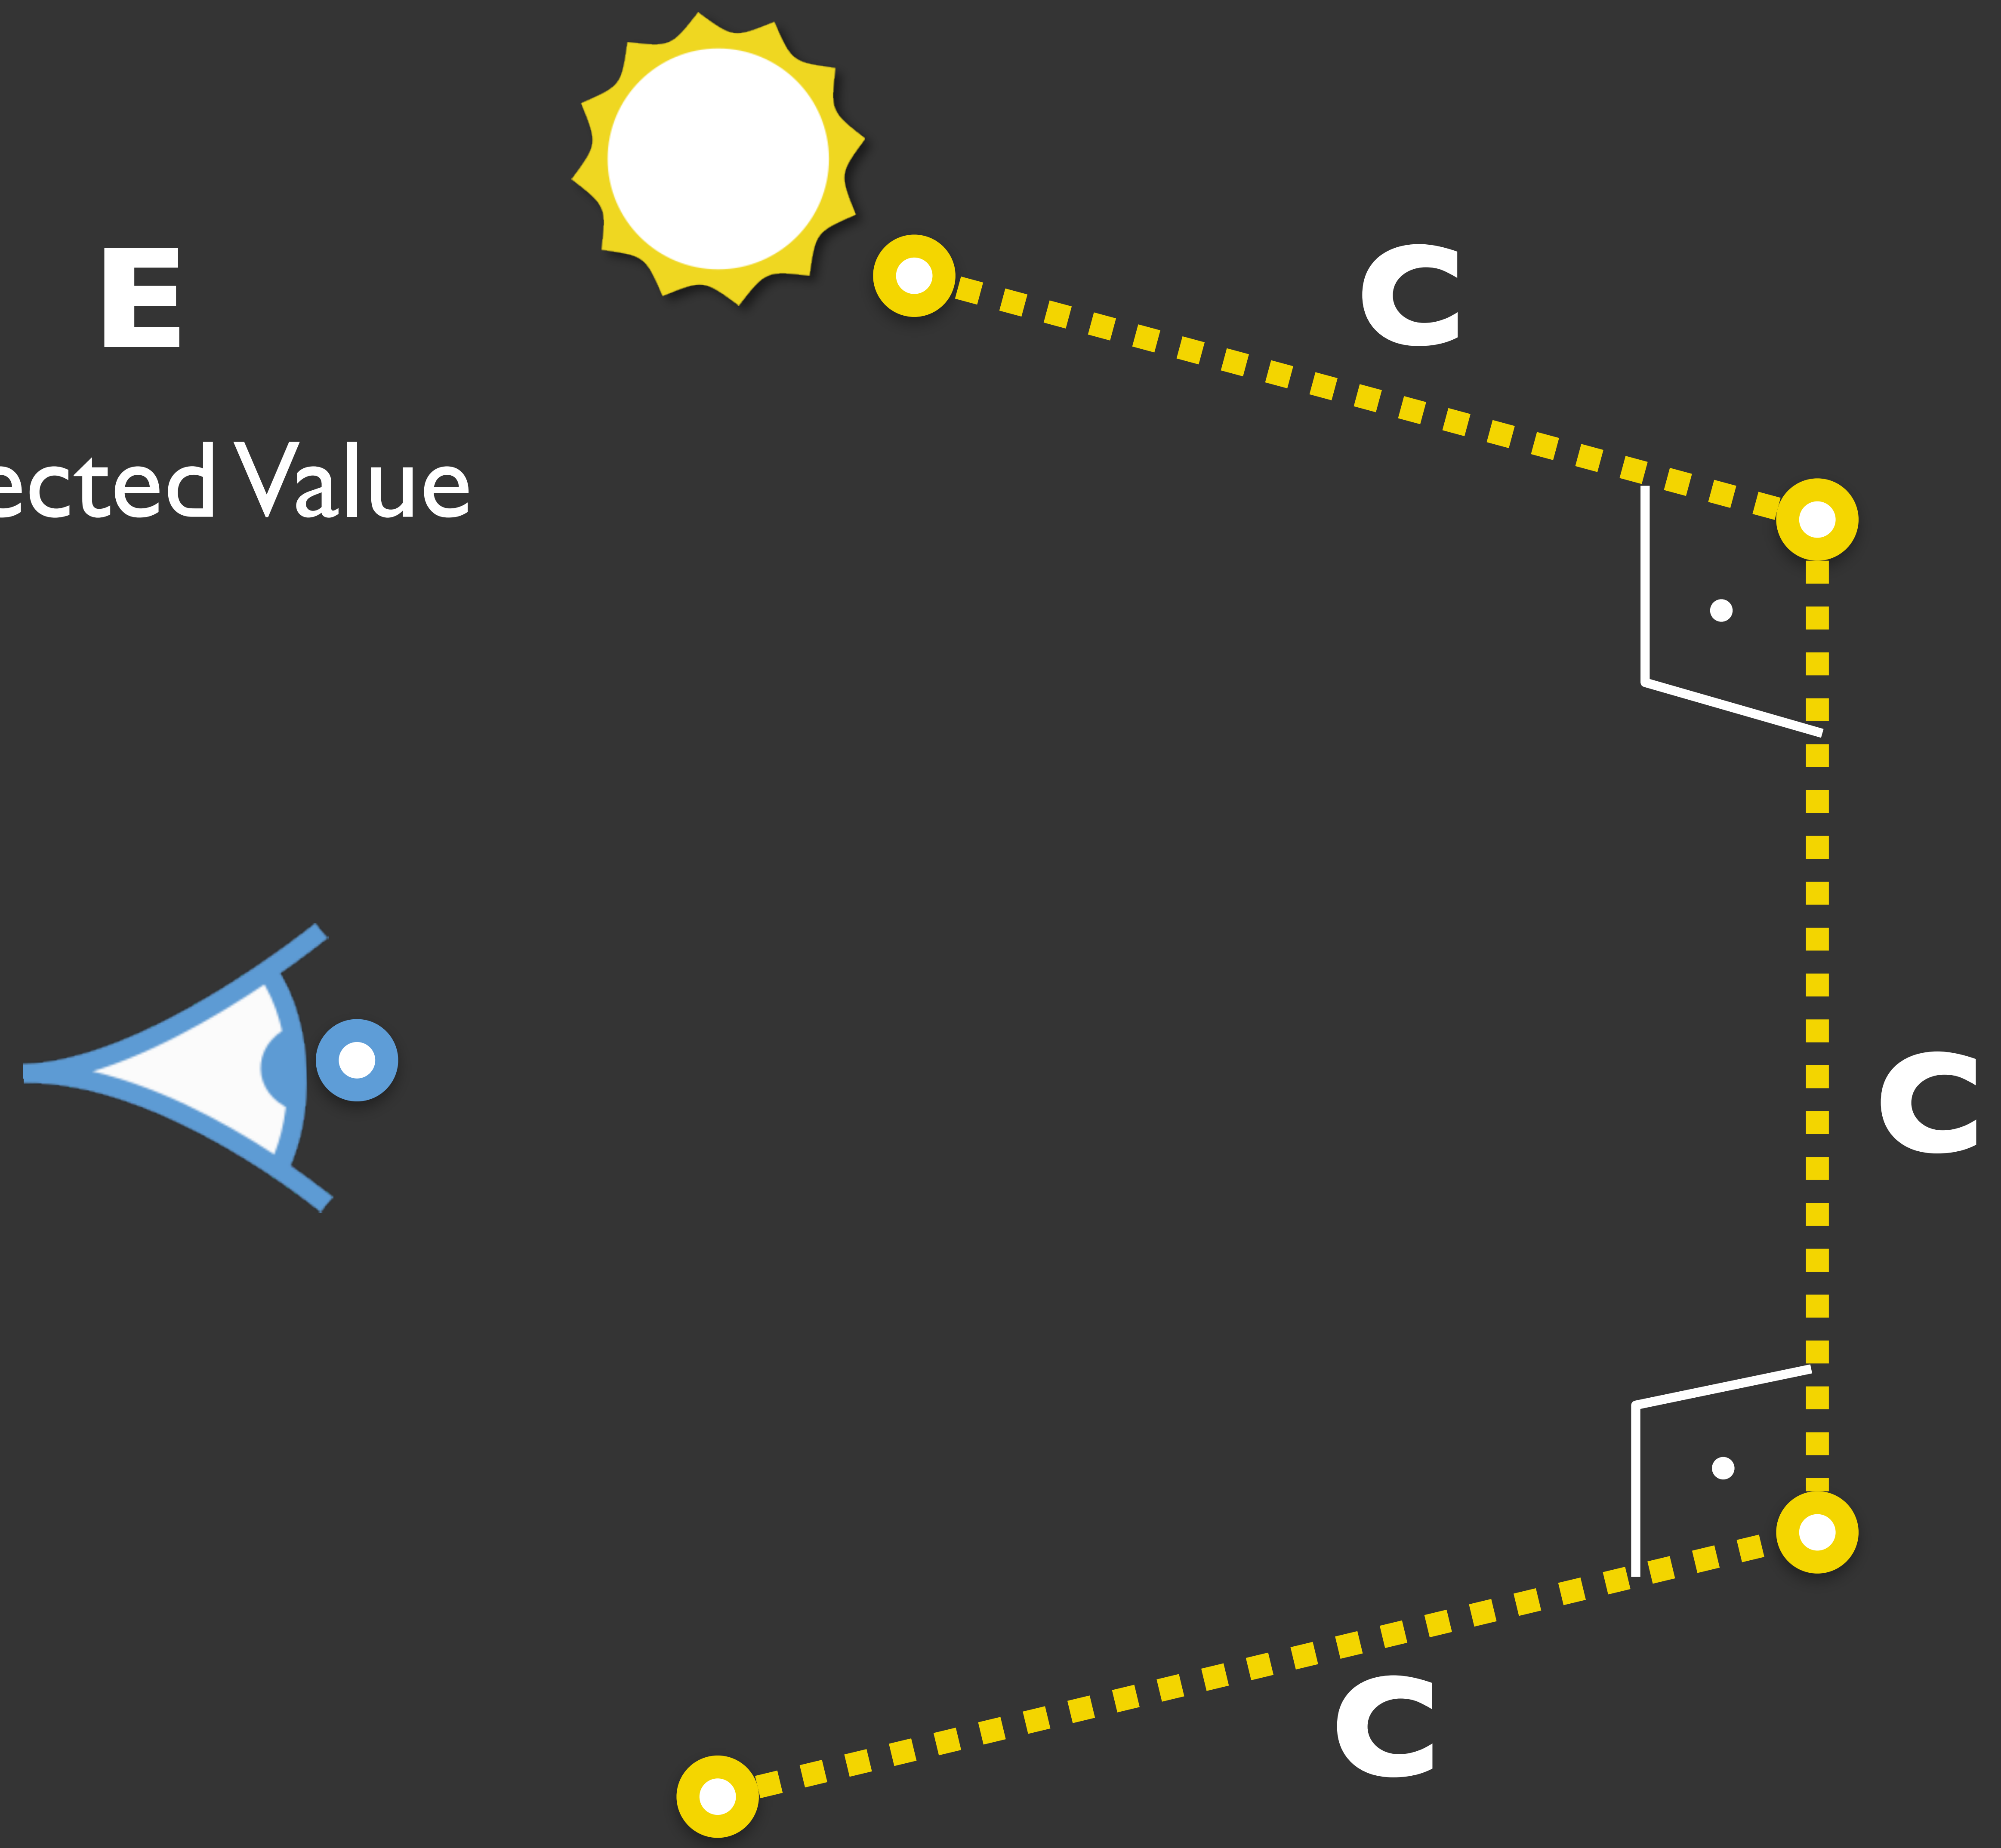

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CCC
Photon Points

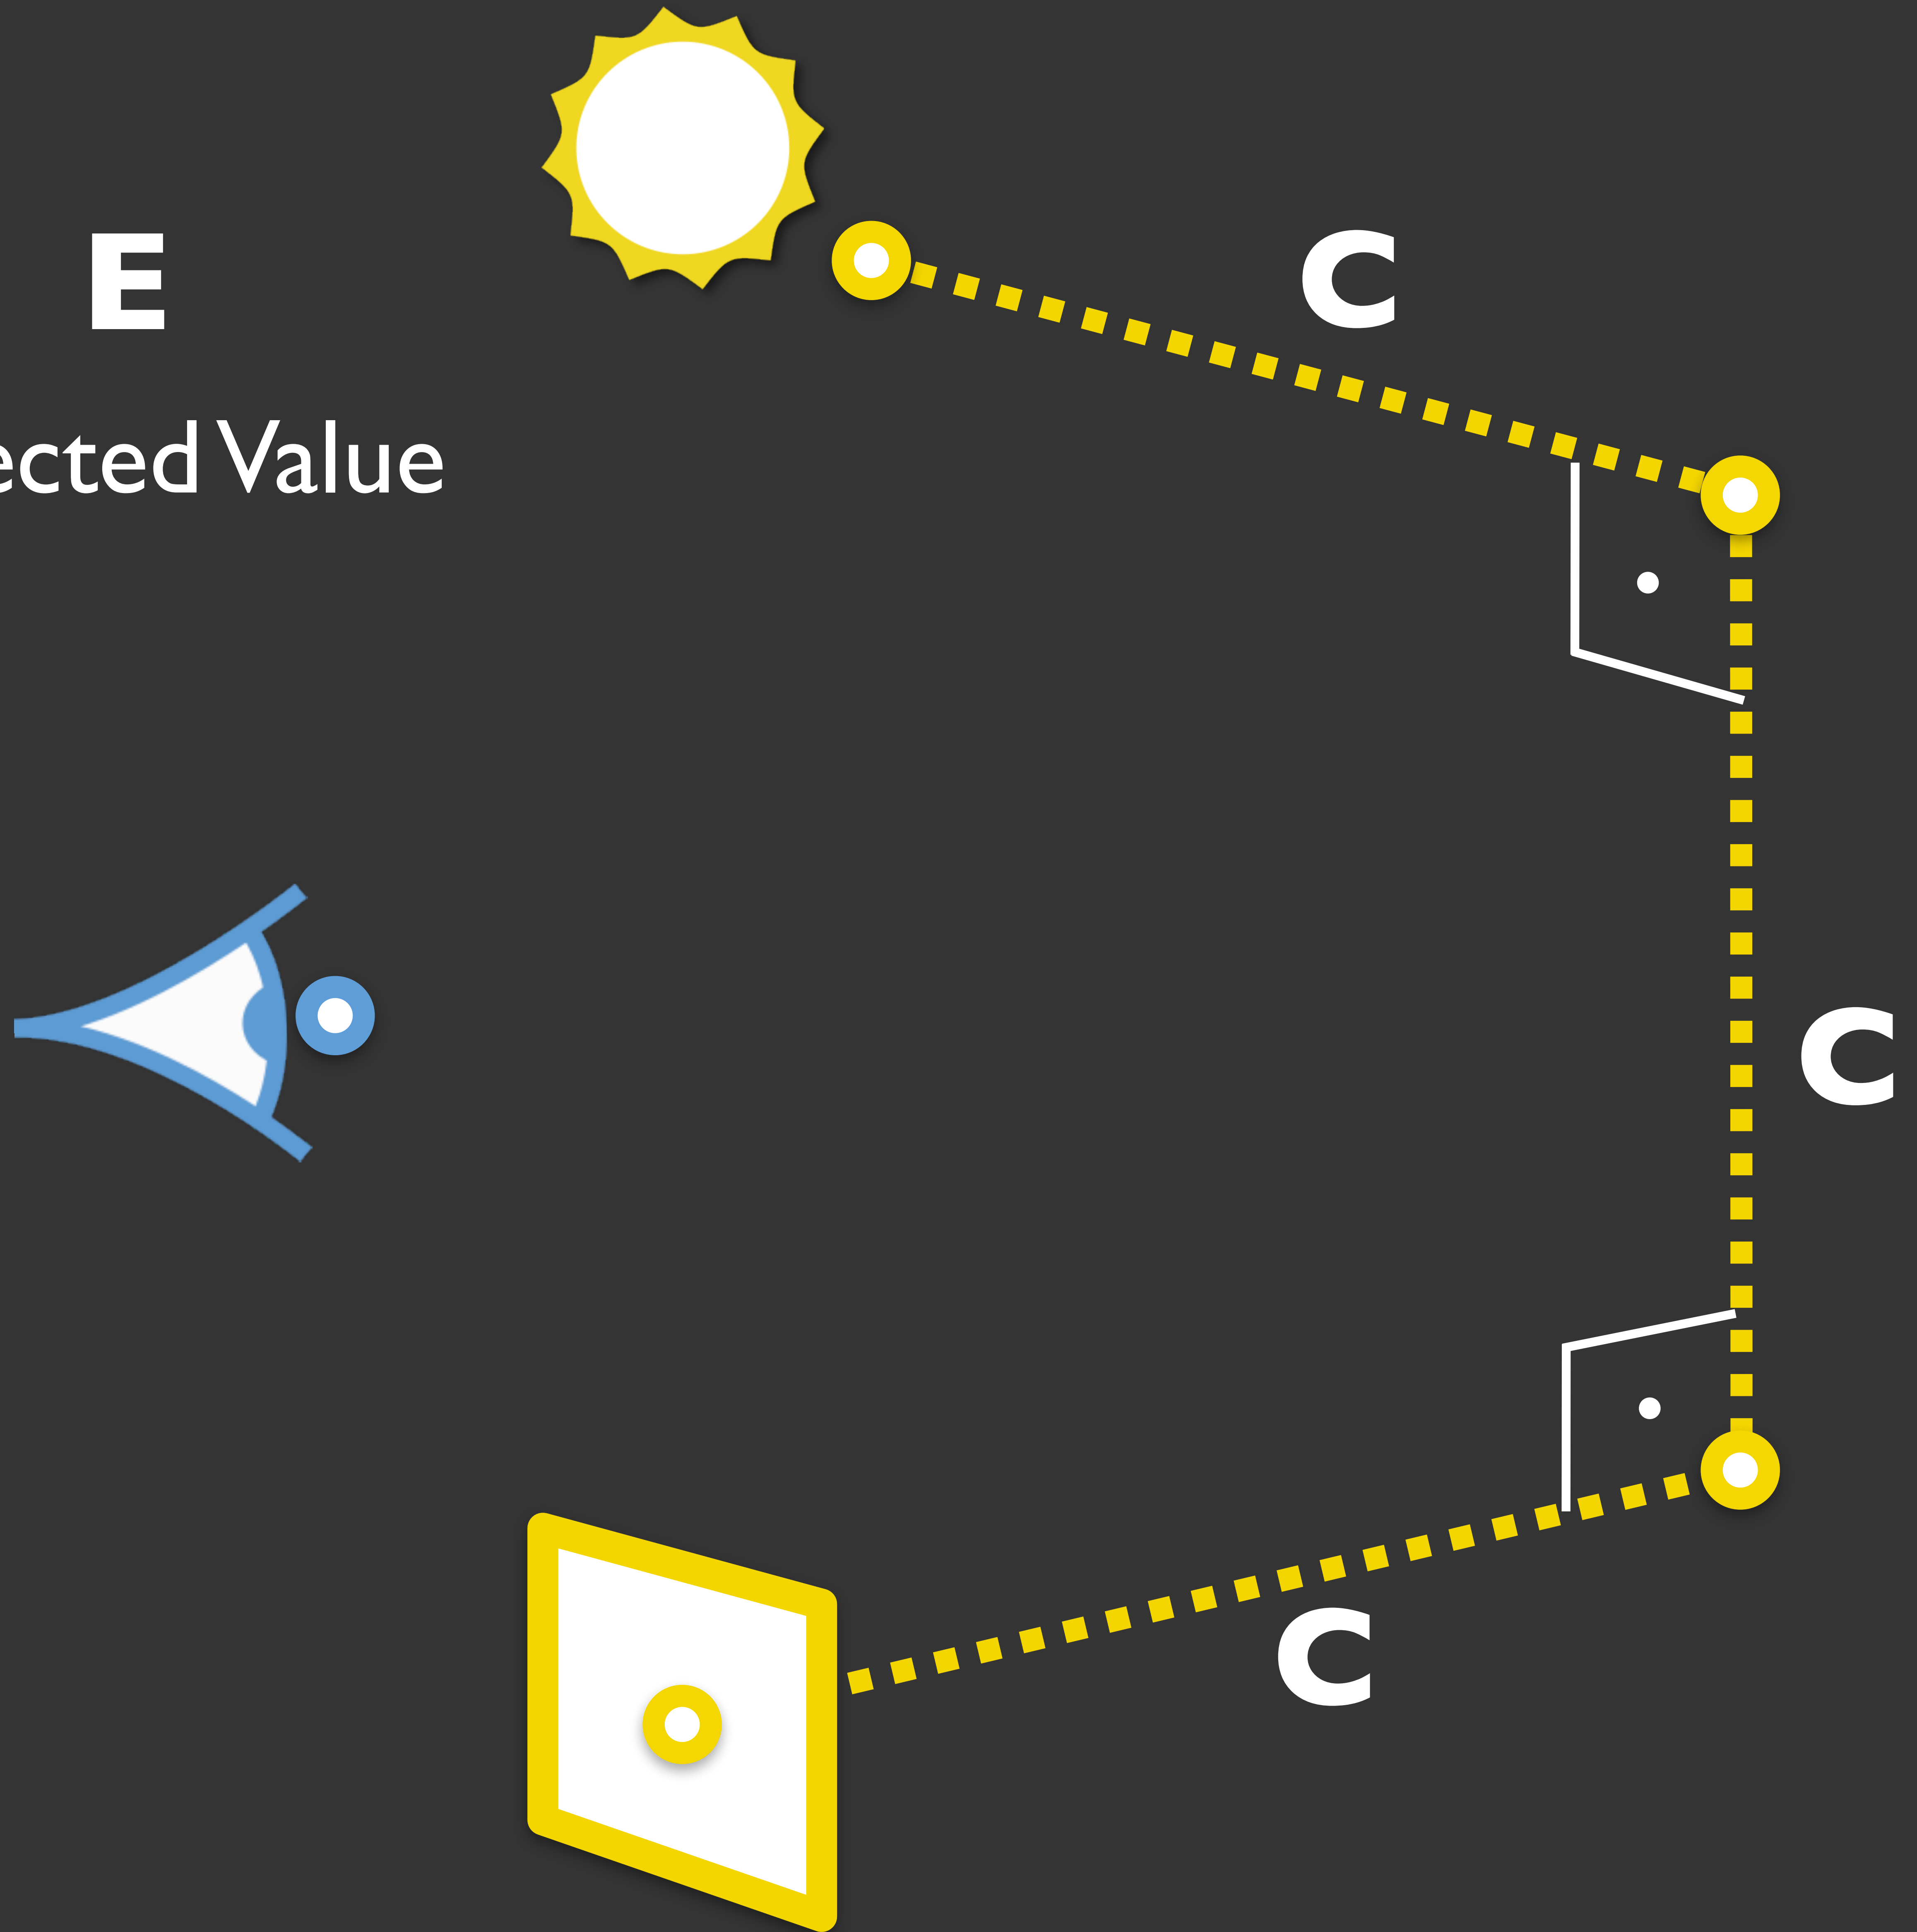

Error Analysis Setup

c
Collision

T
Track-Length

E
Expected Value

ccc
Photon Points

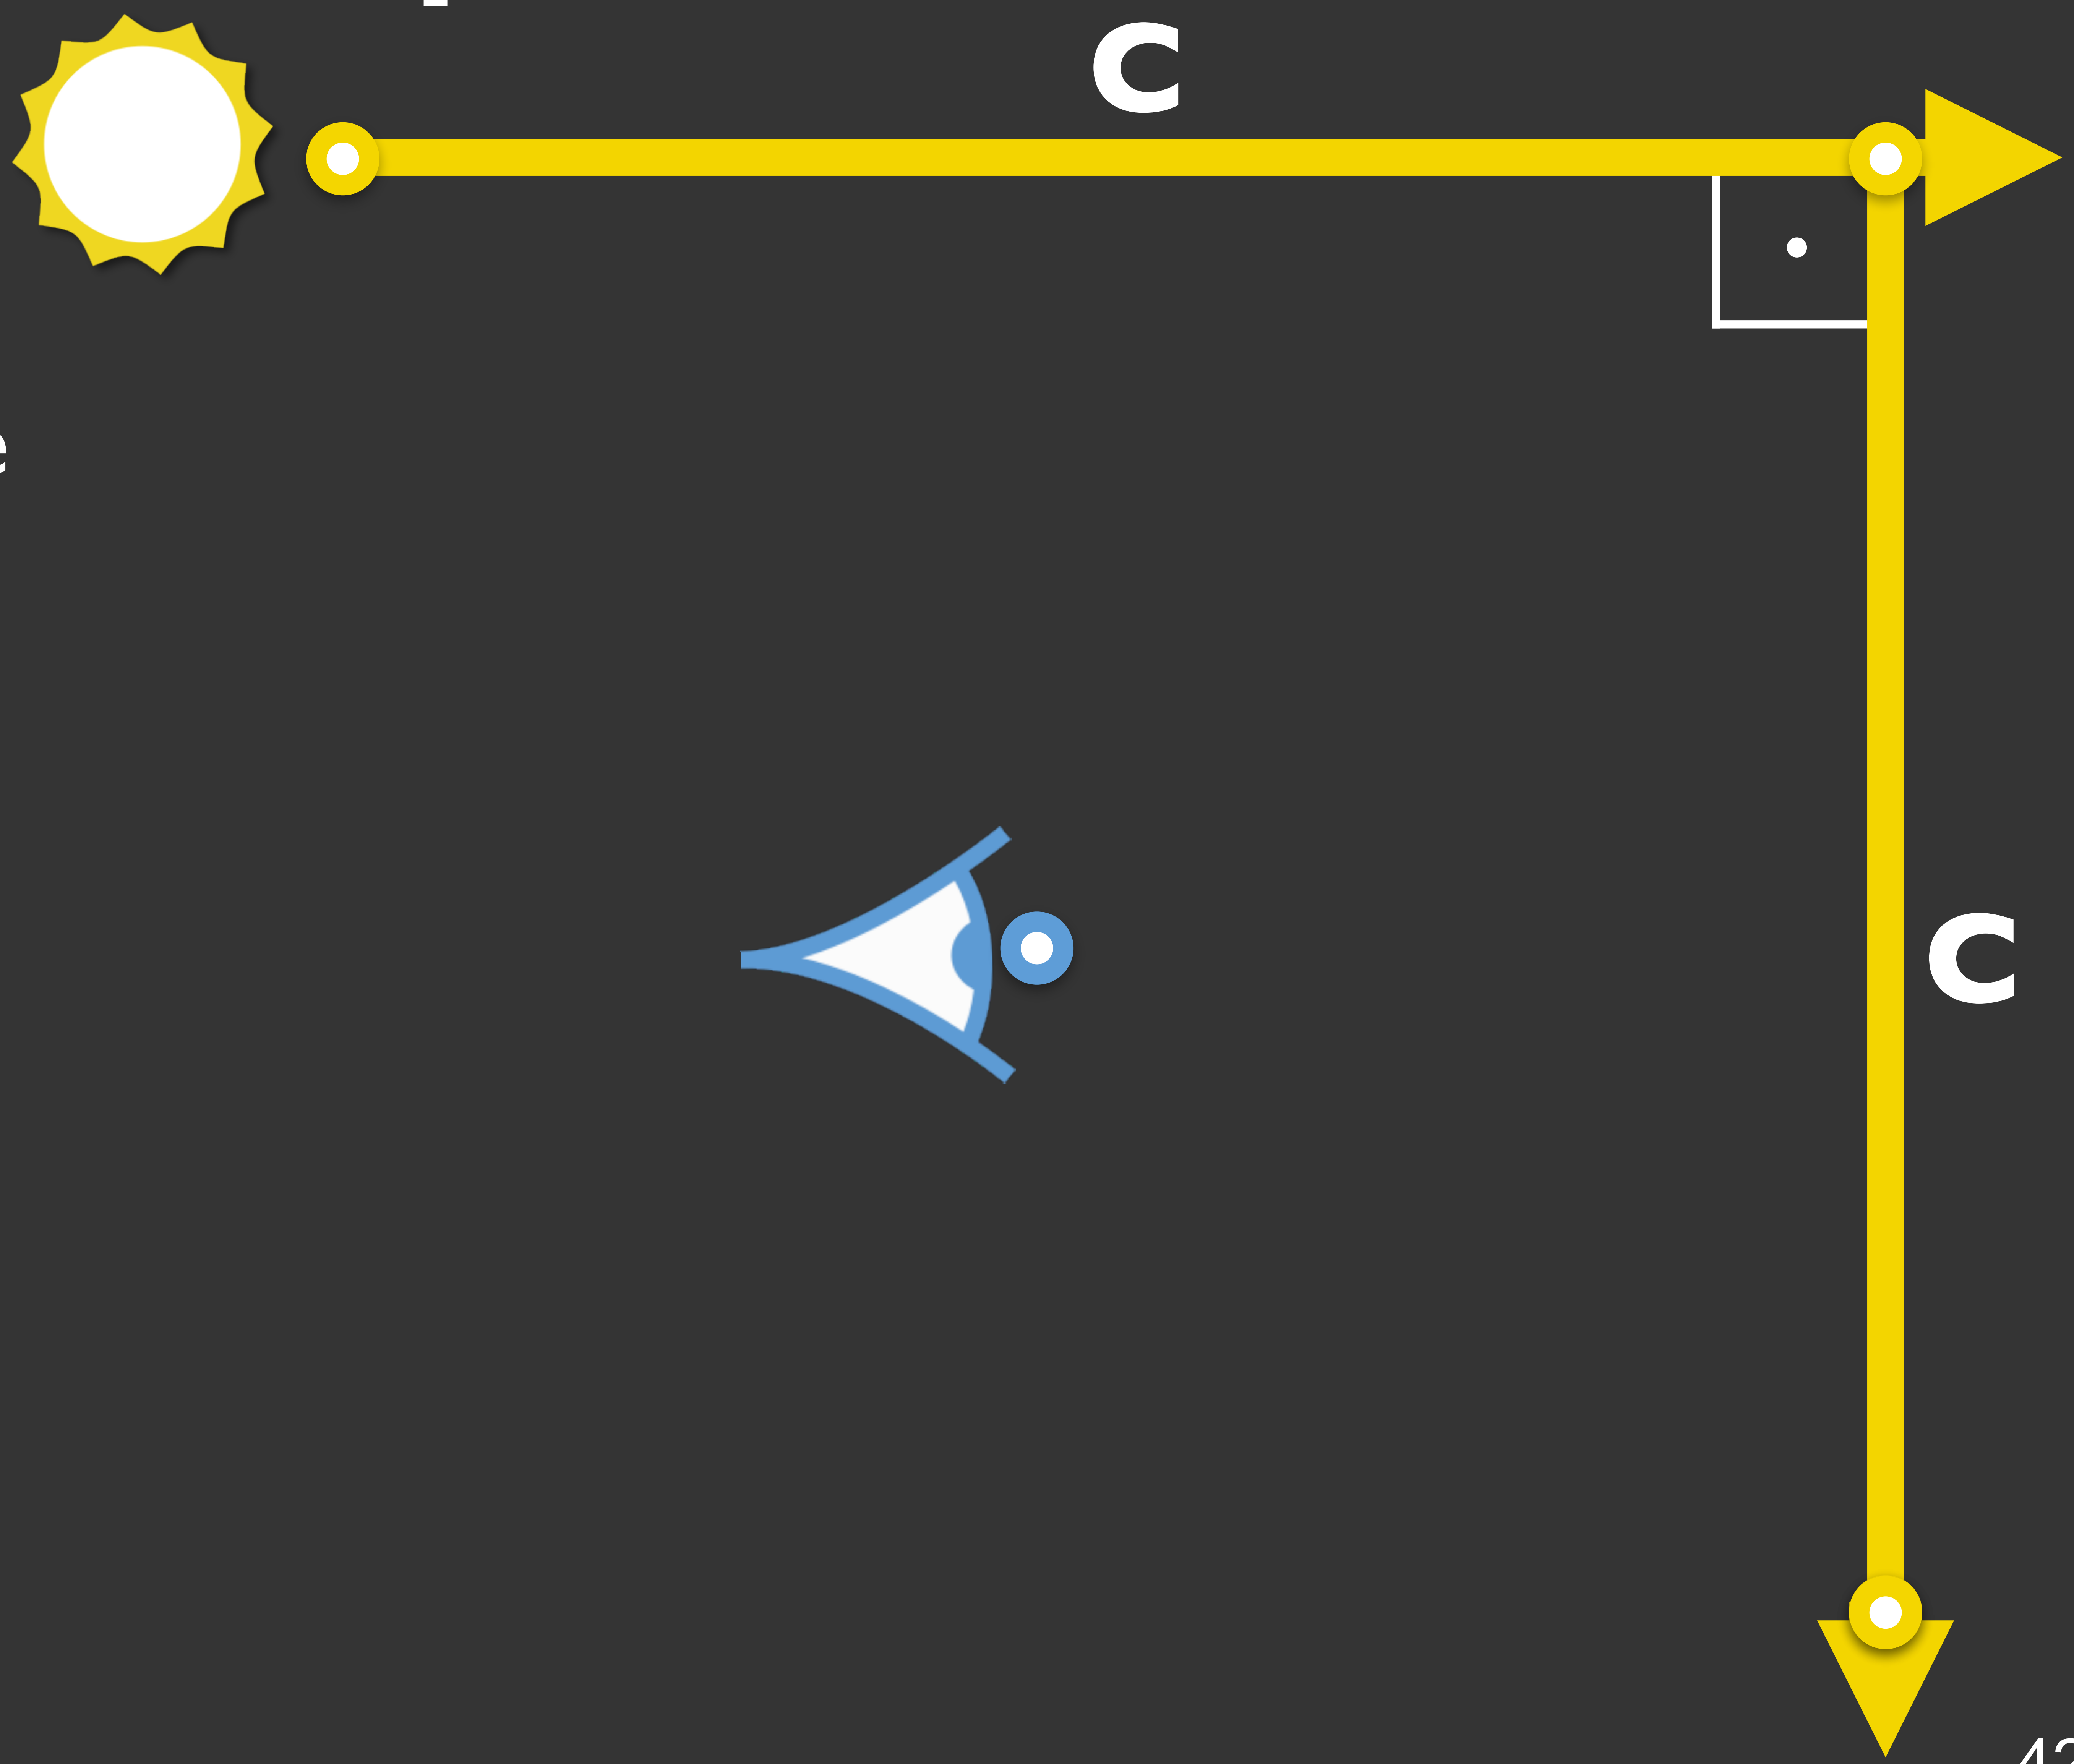

Error Analysis Setup

c

Collision

T

Track-Length

E

Expected Value

ccc

Photon Points

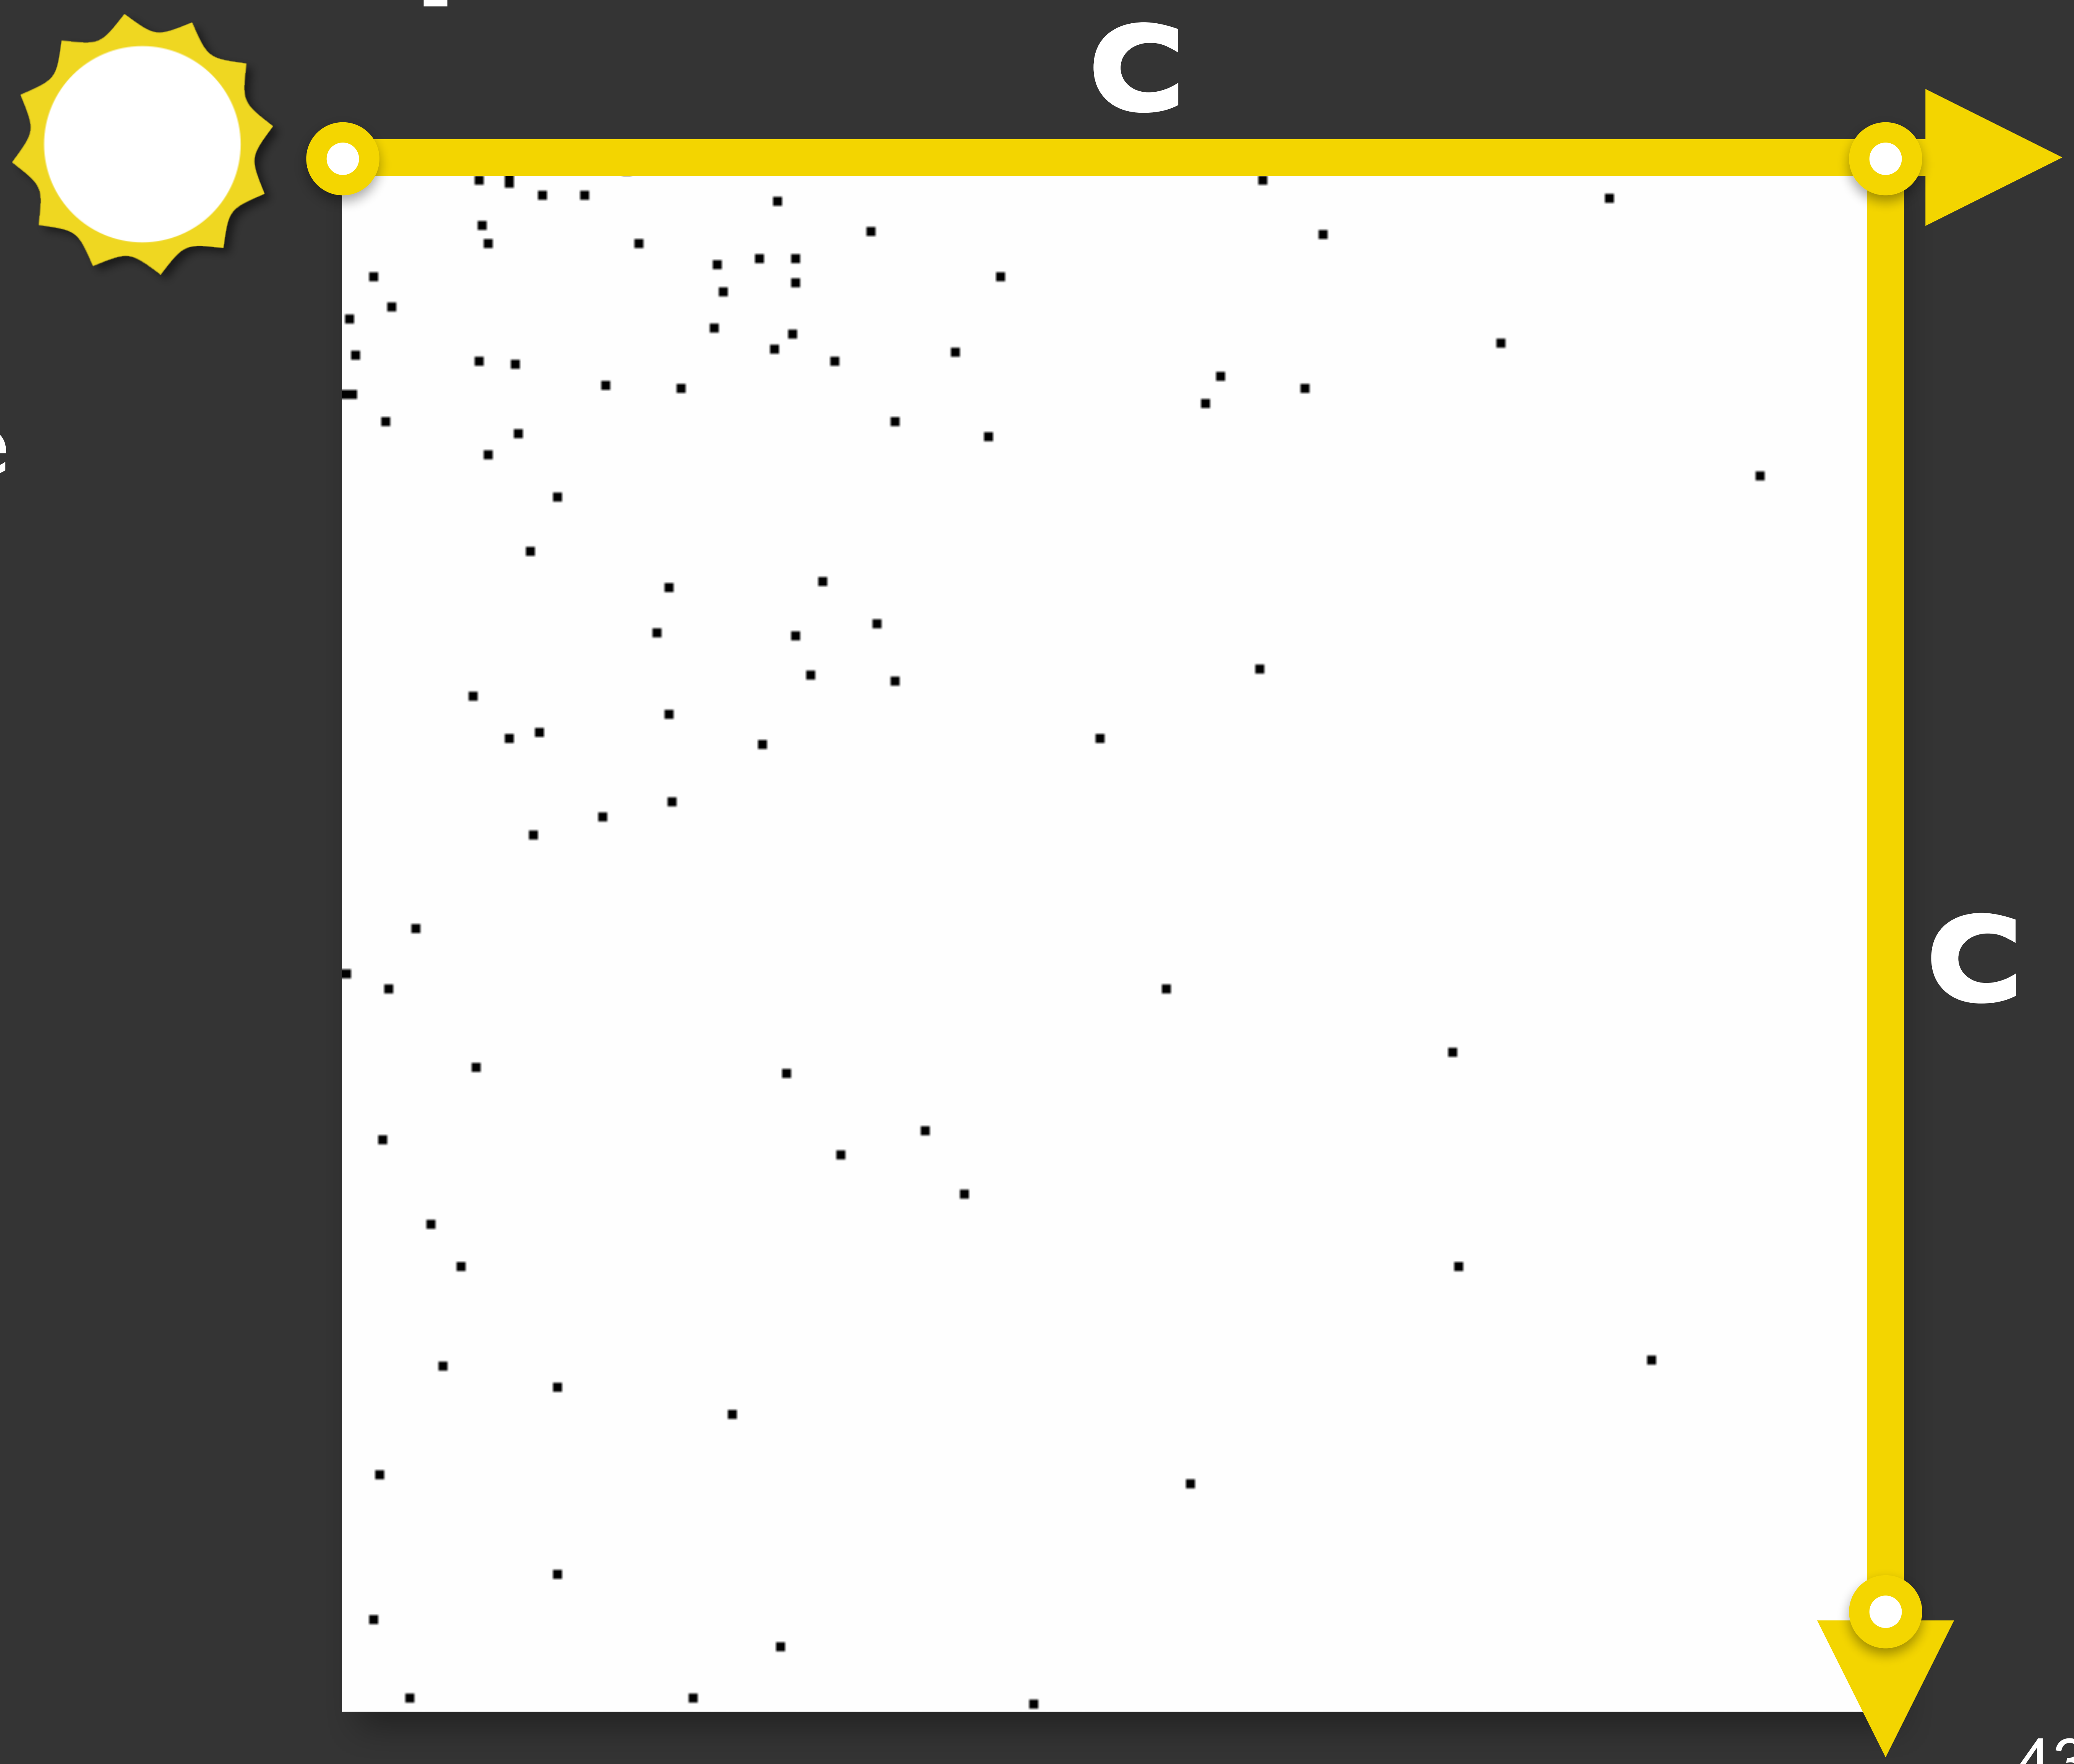

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CCE
Long Beams

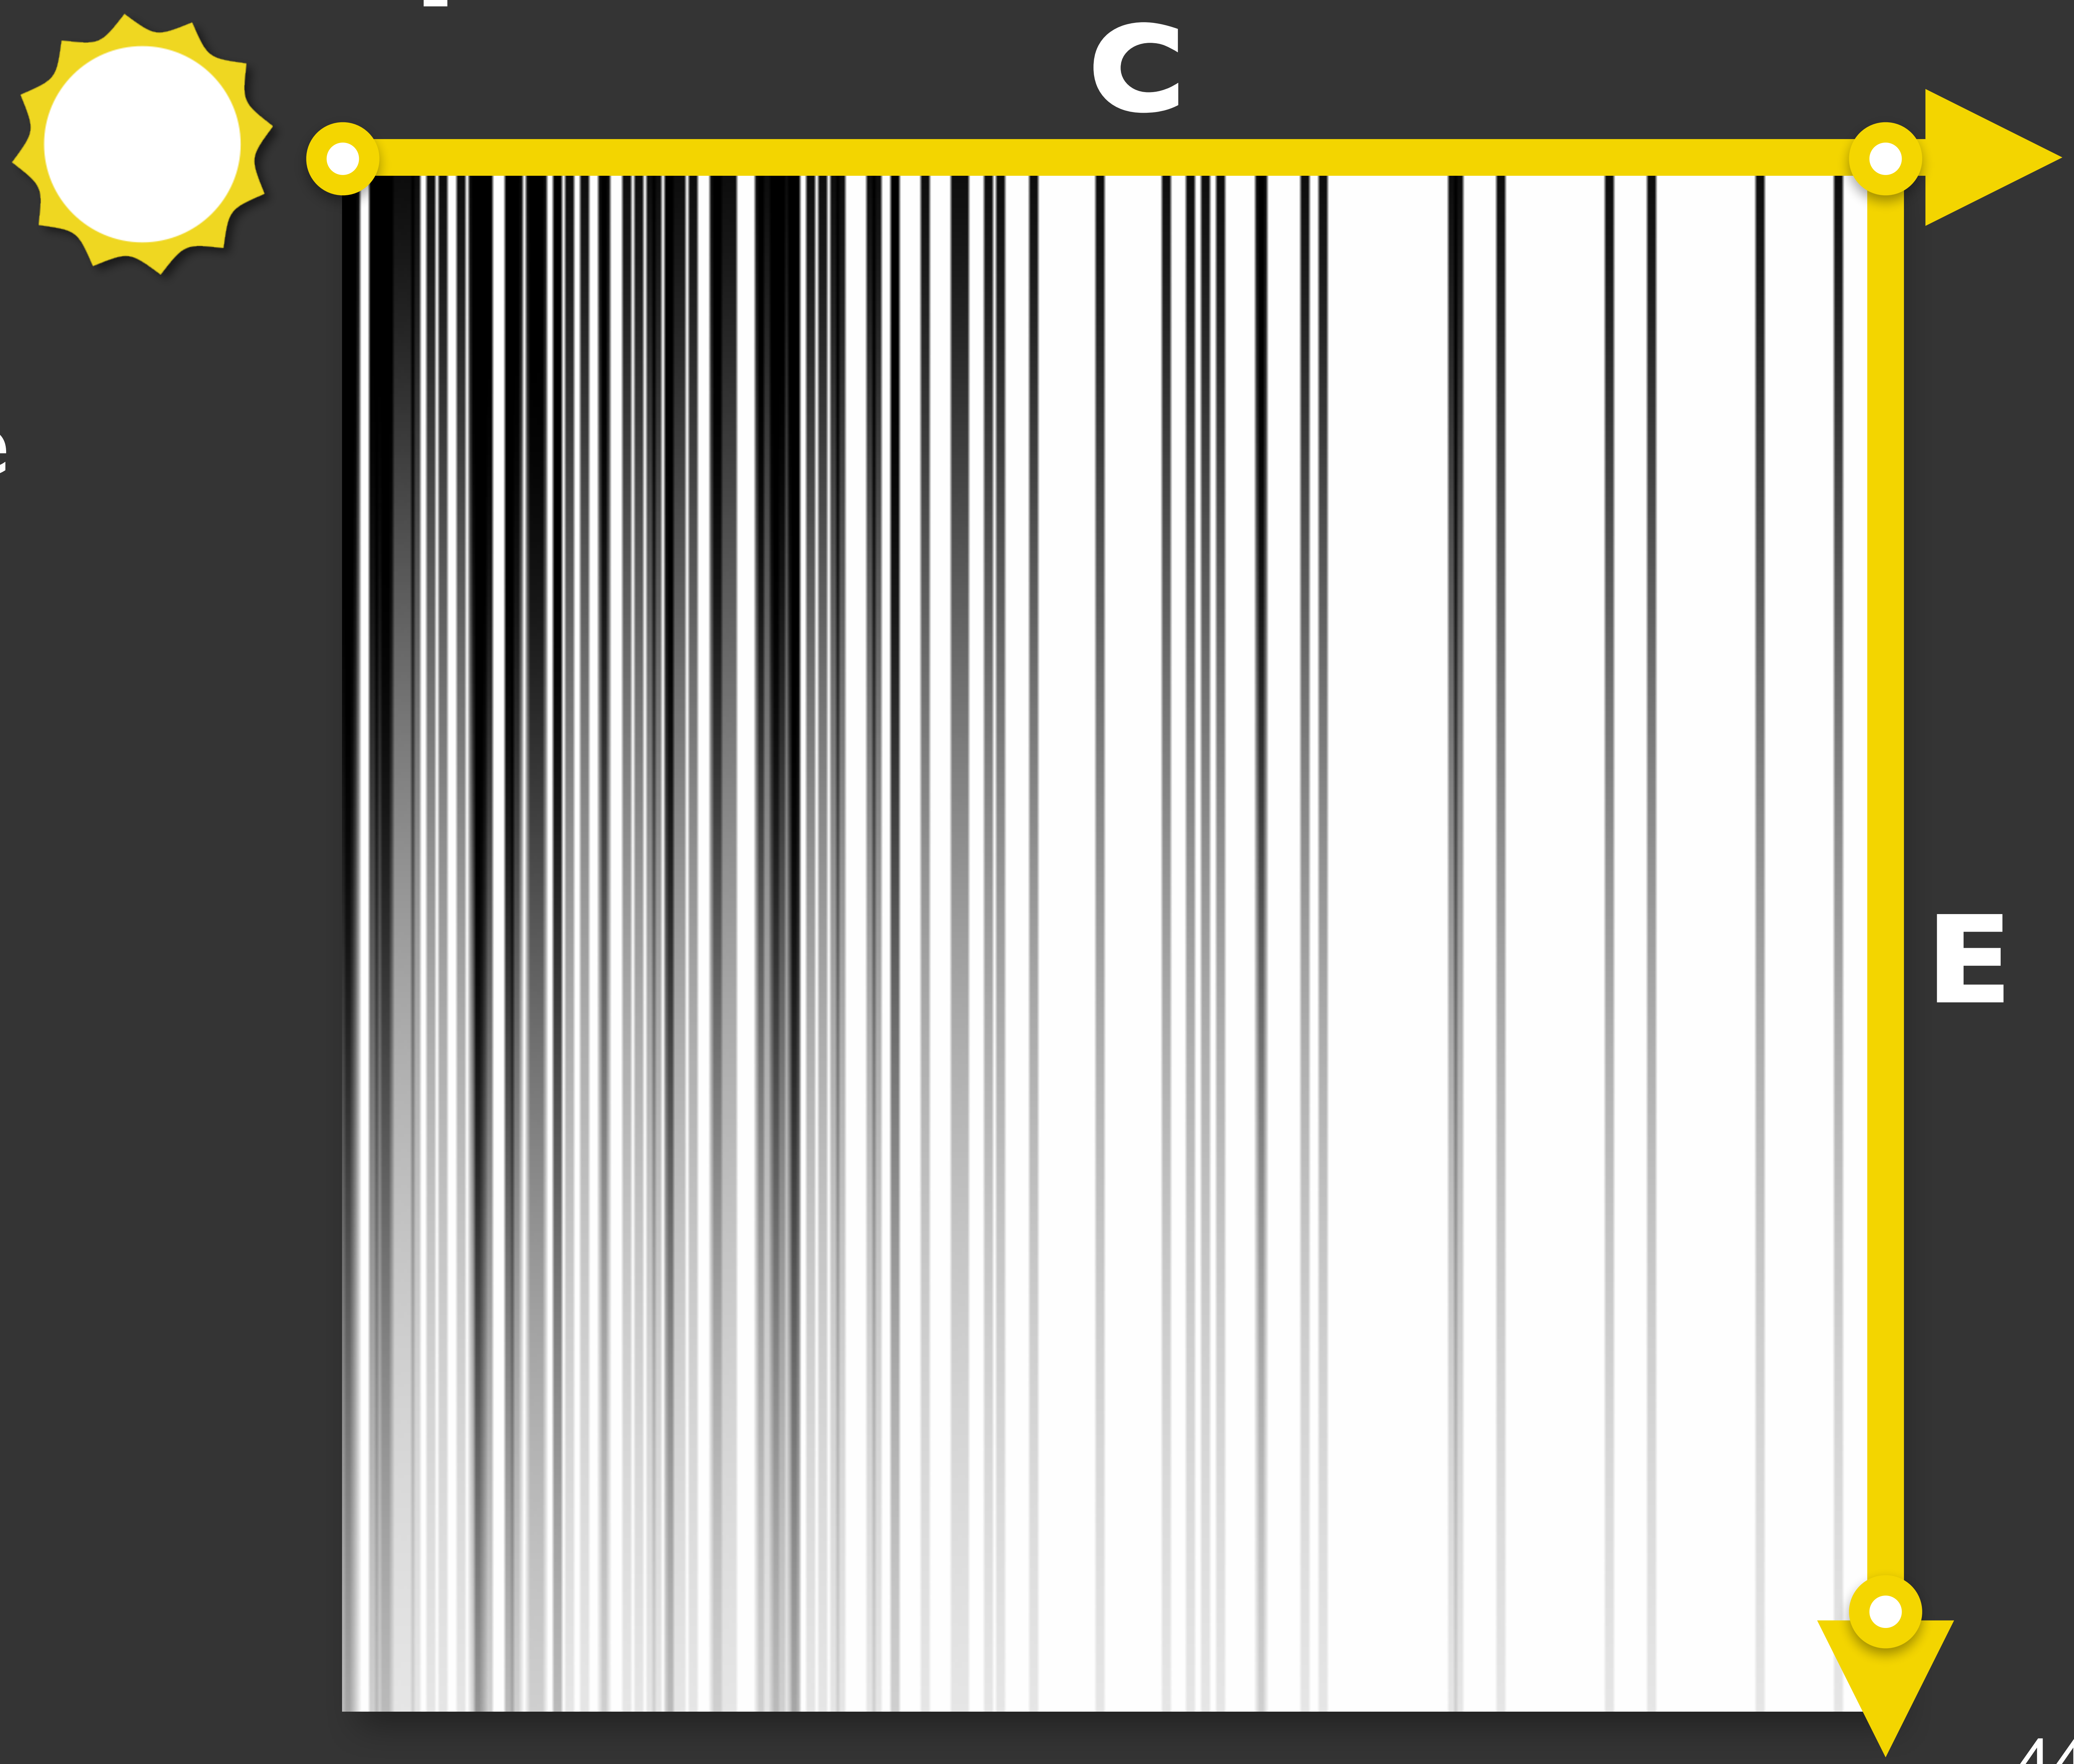

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CCE
Long Beams

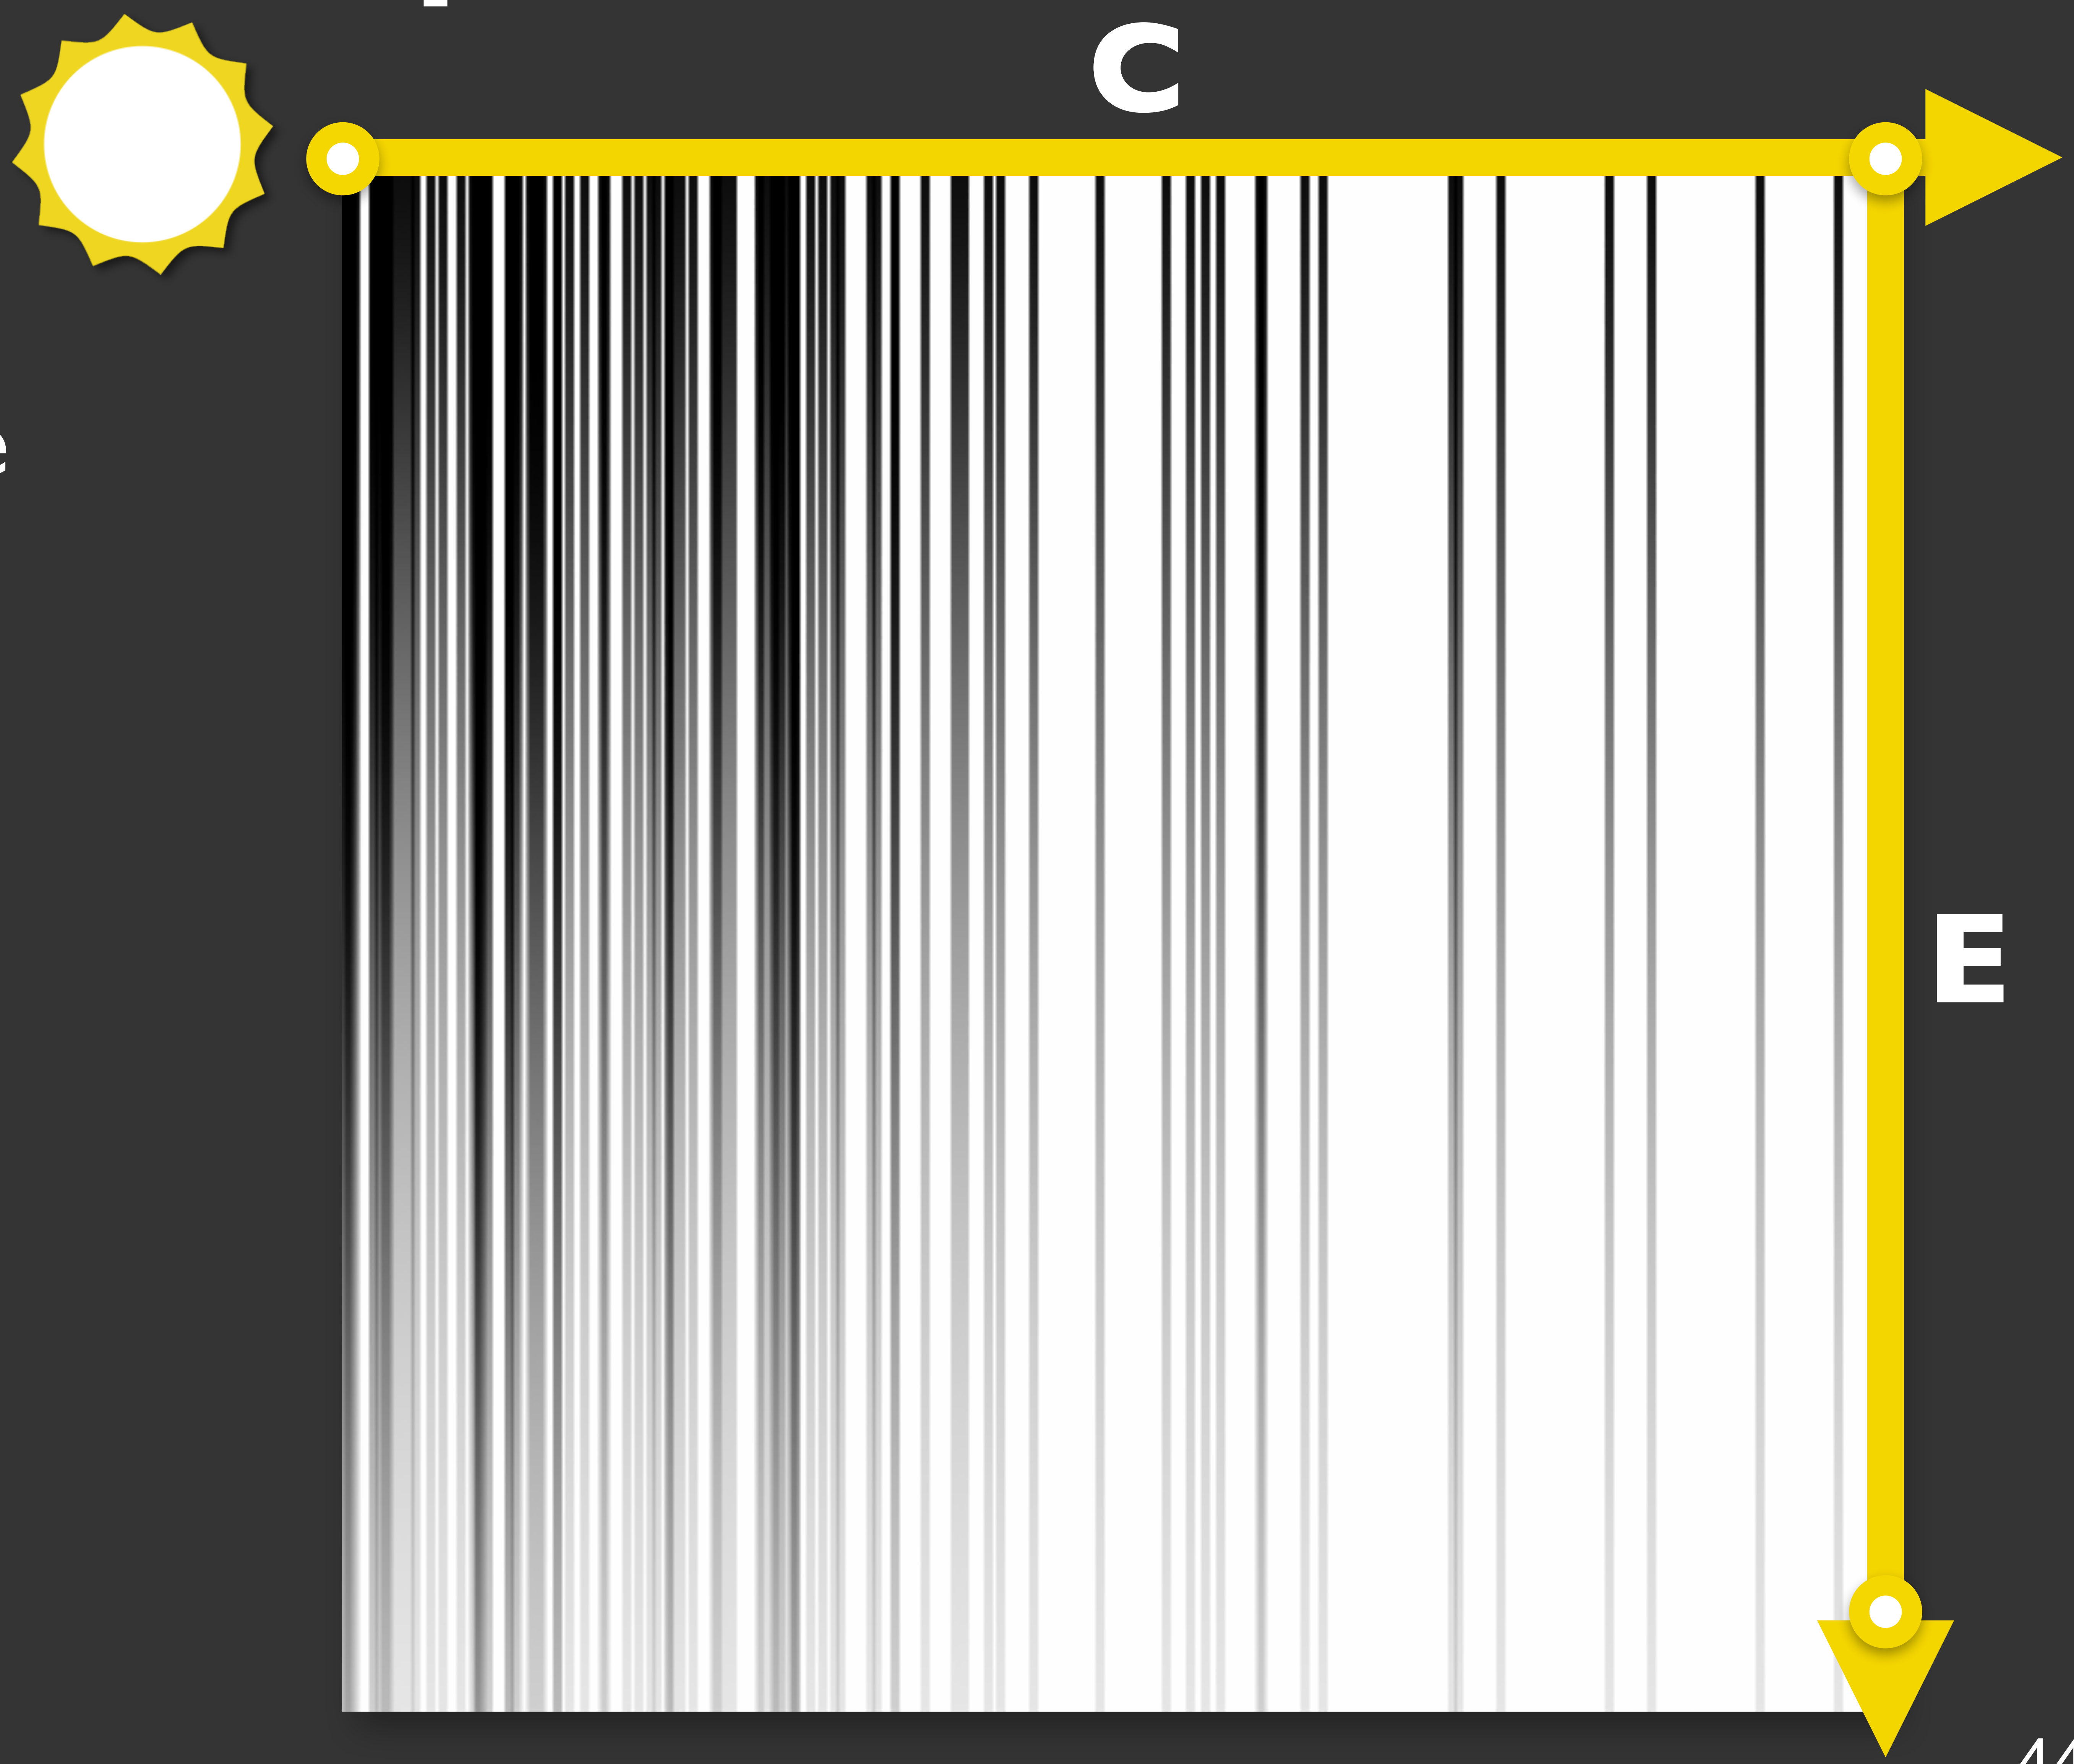

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CCT
Short Beams

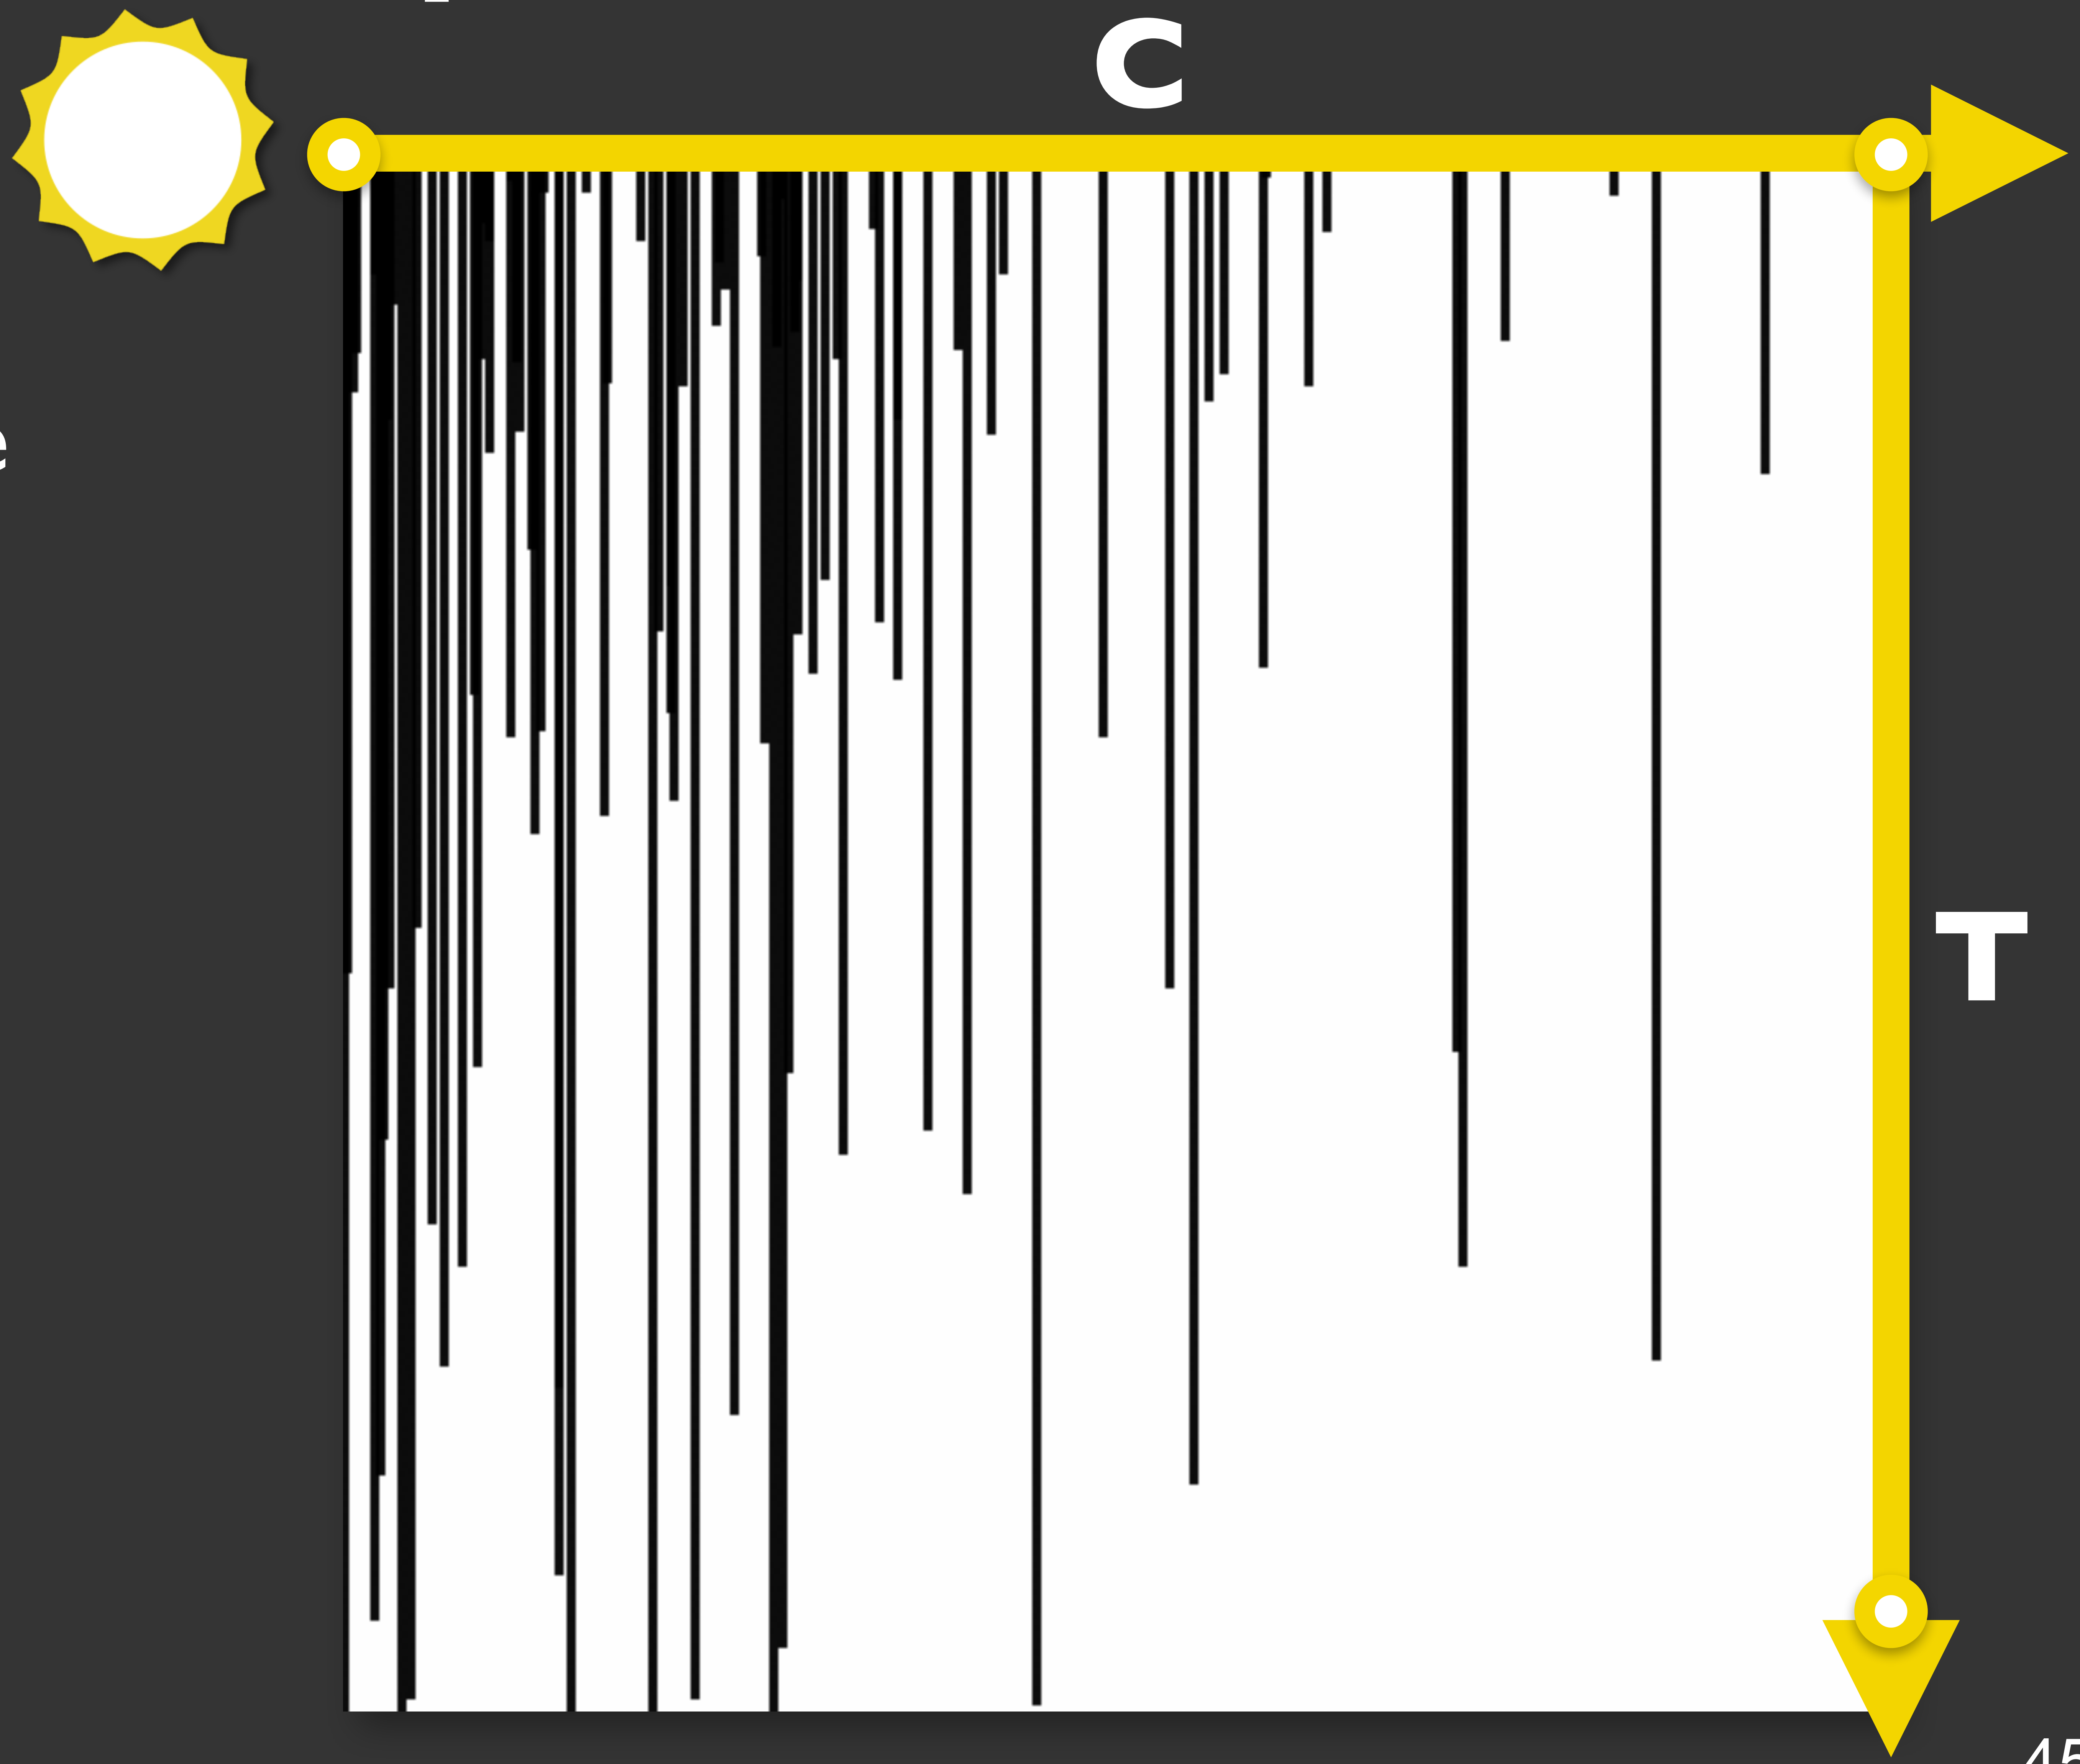

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CCT
Short Beams

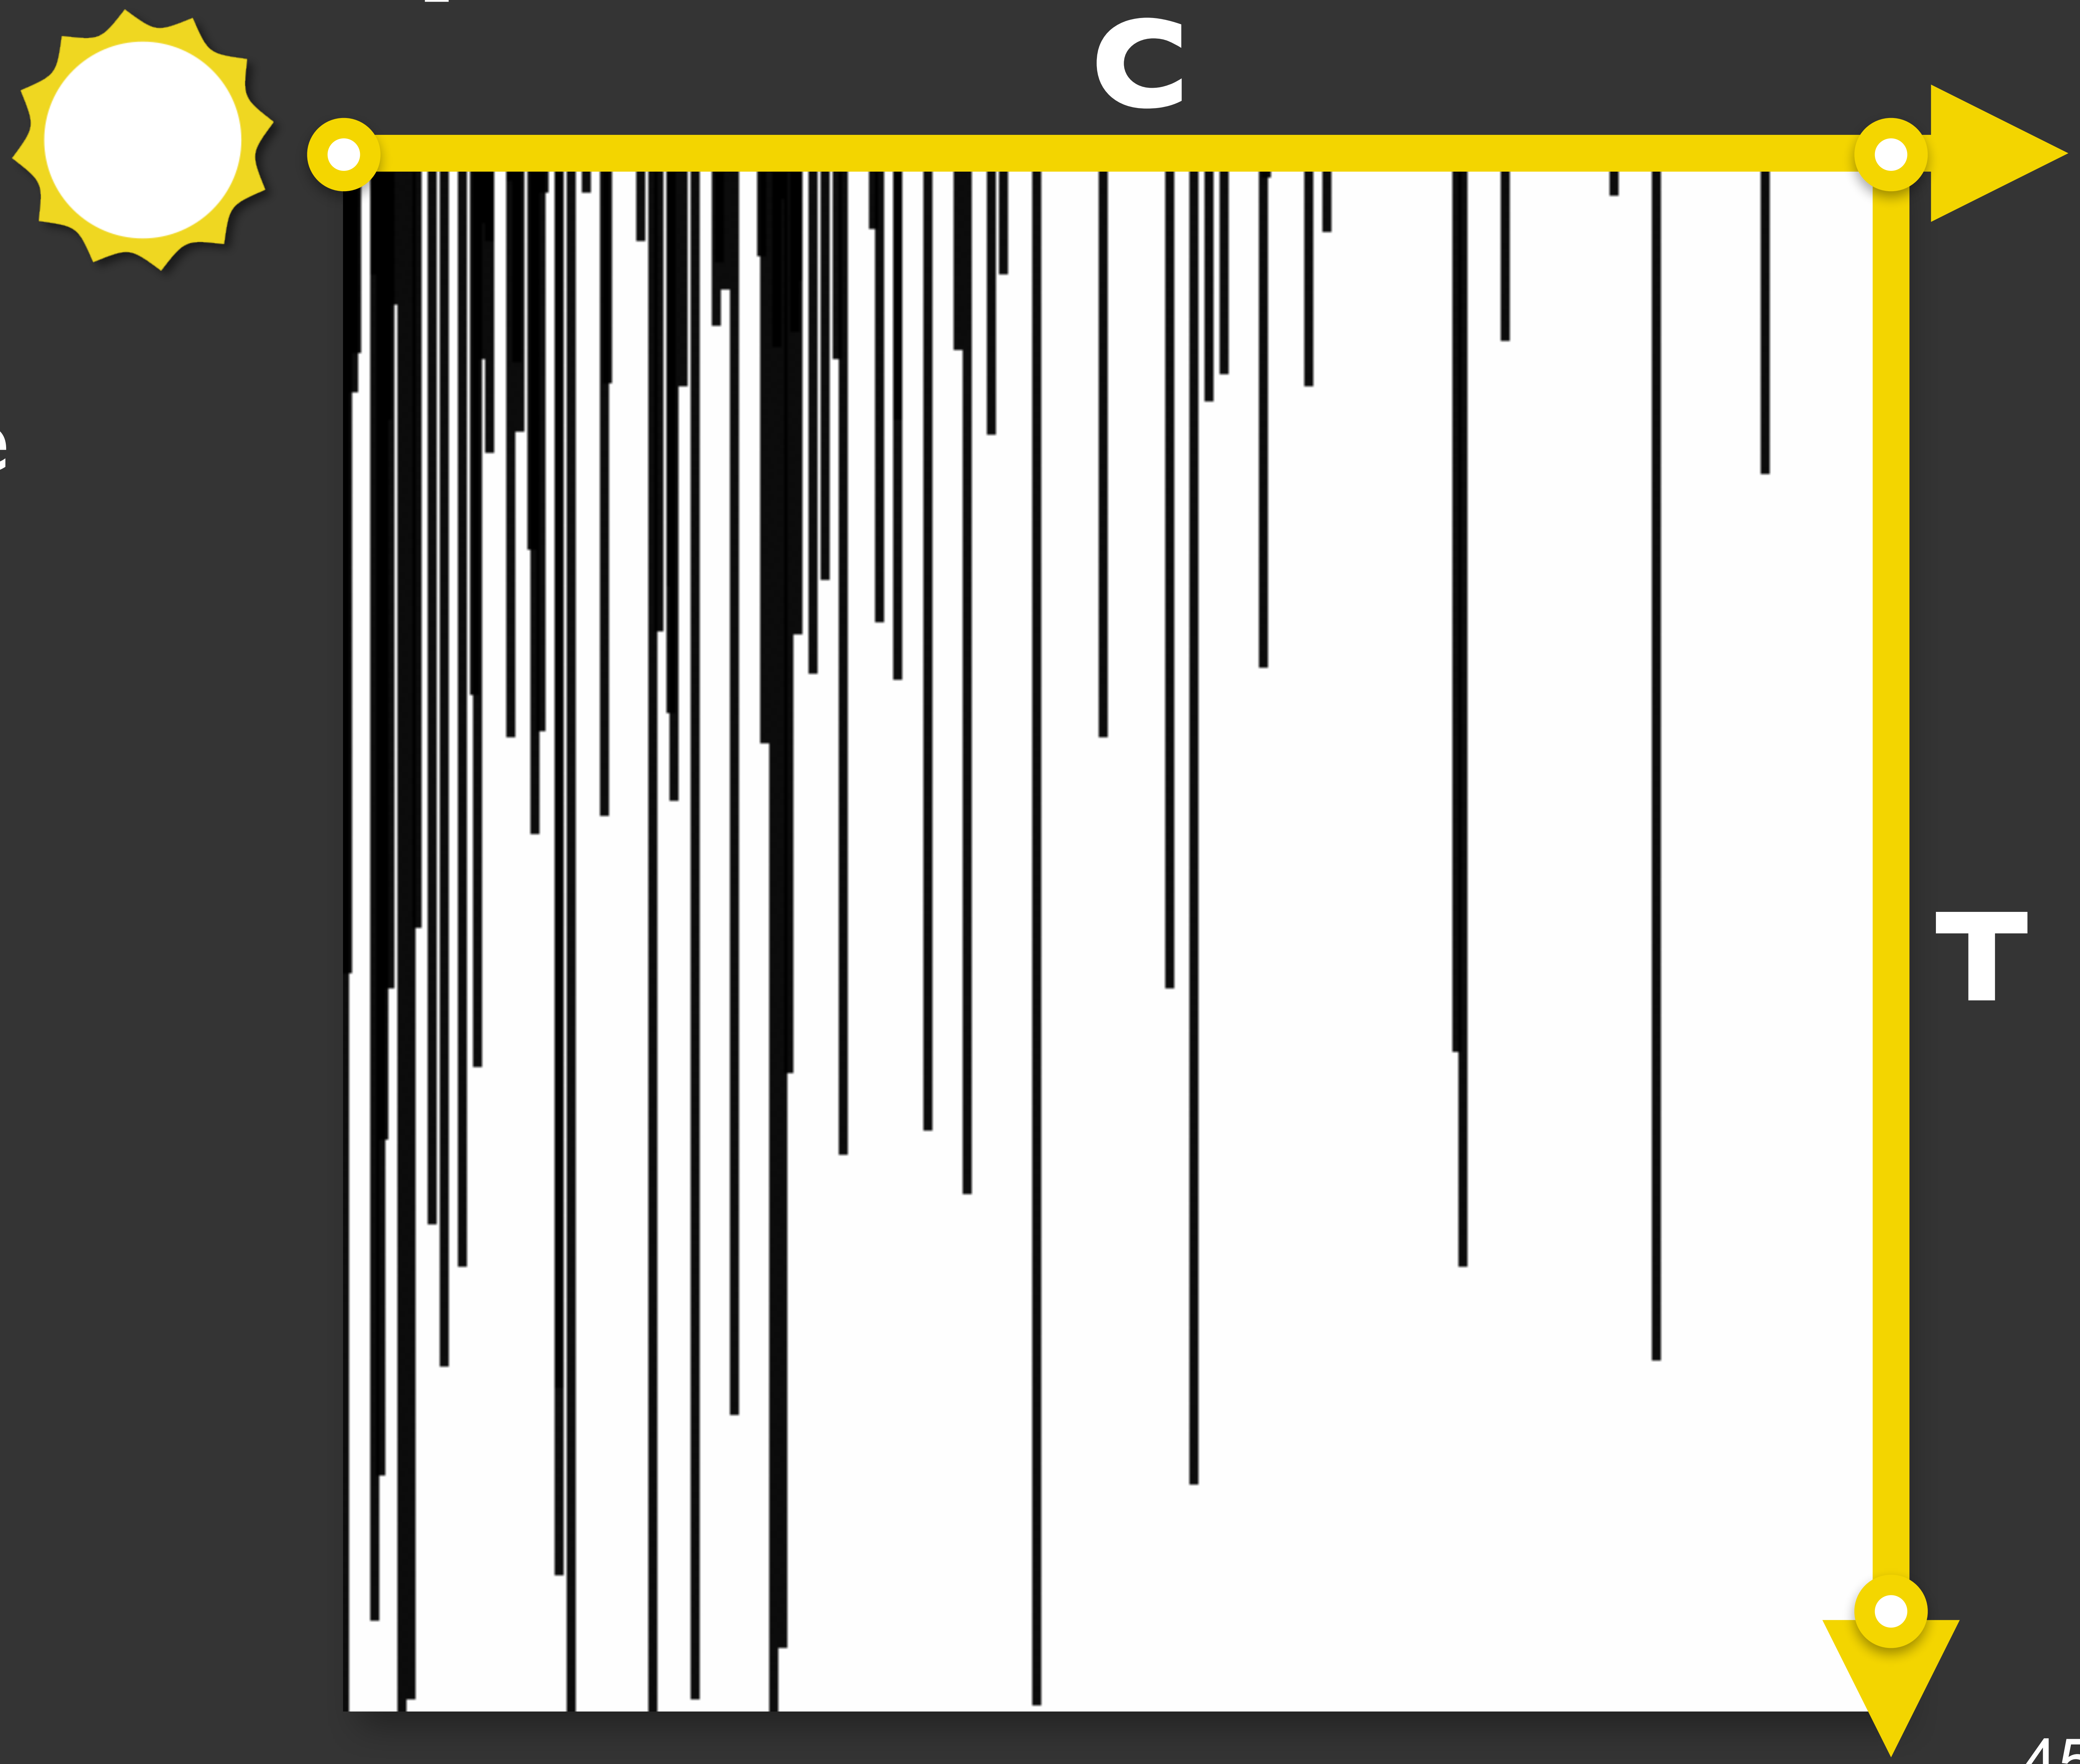

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CTT
Short Planes

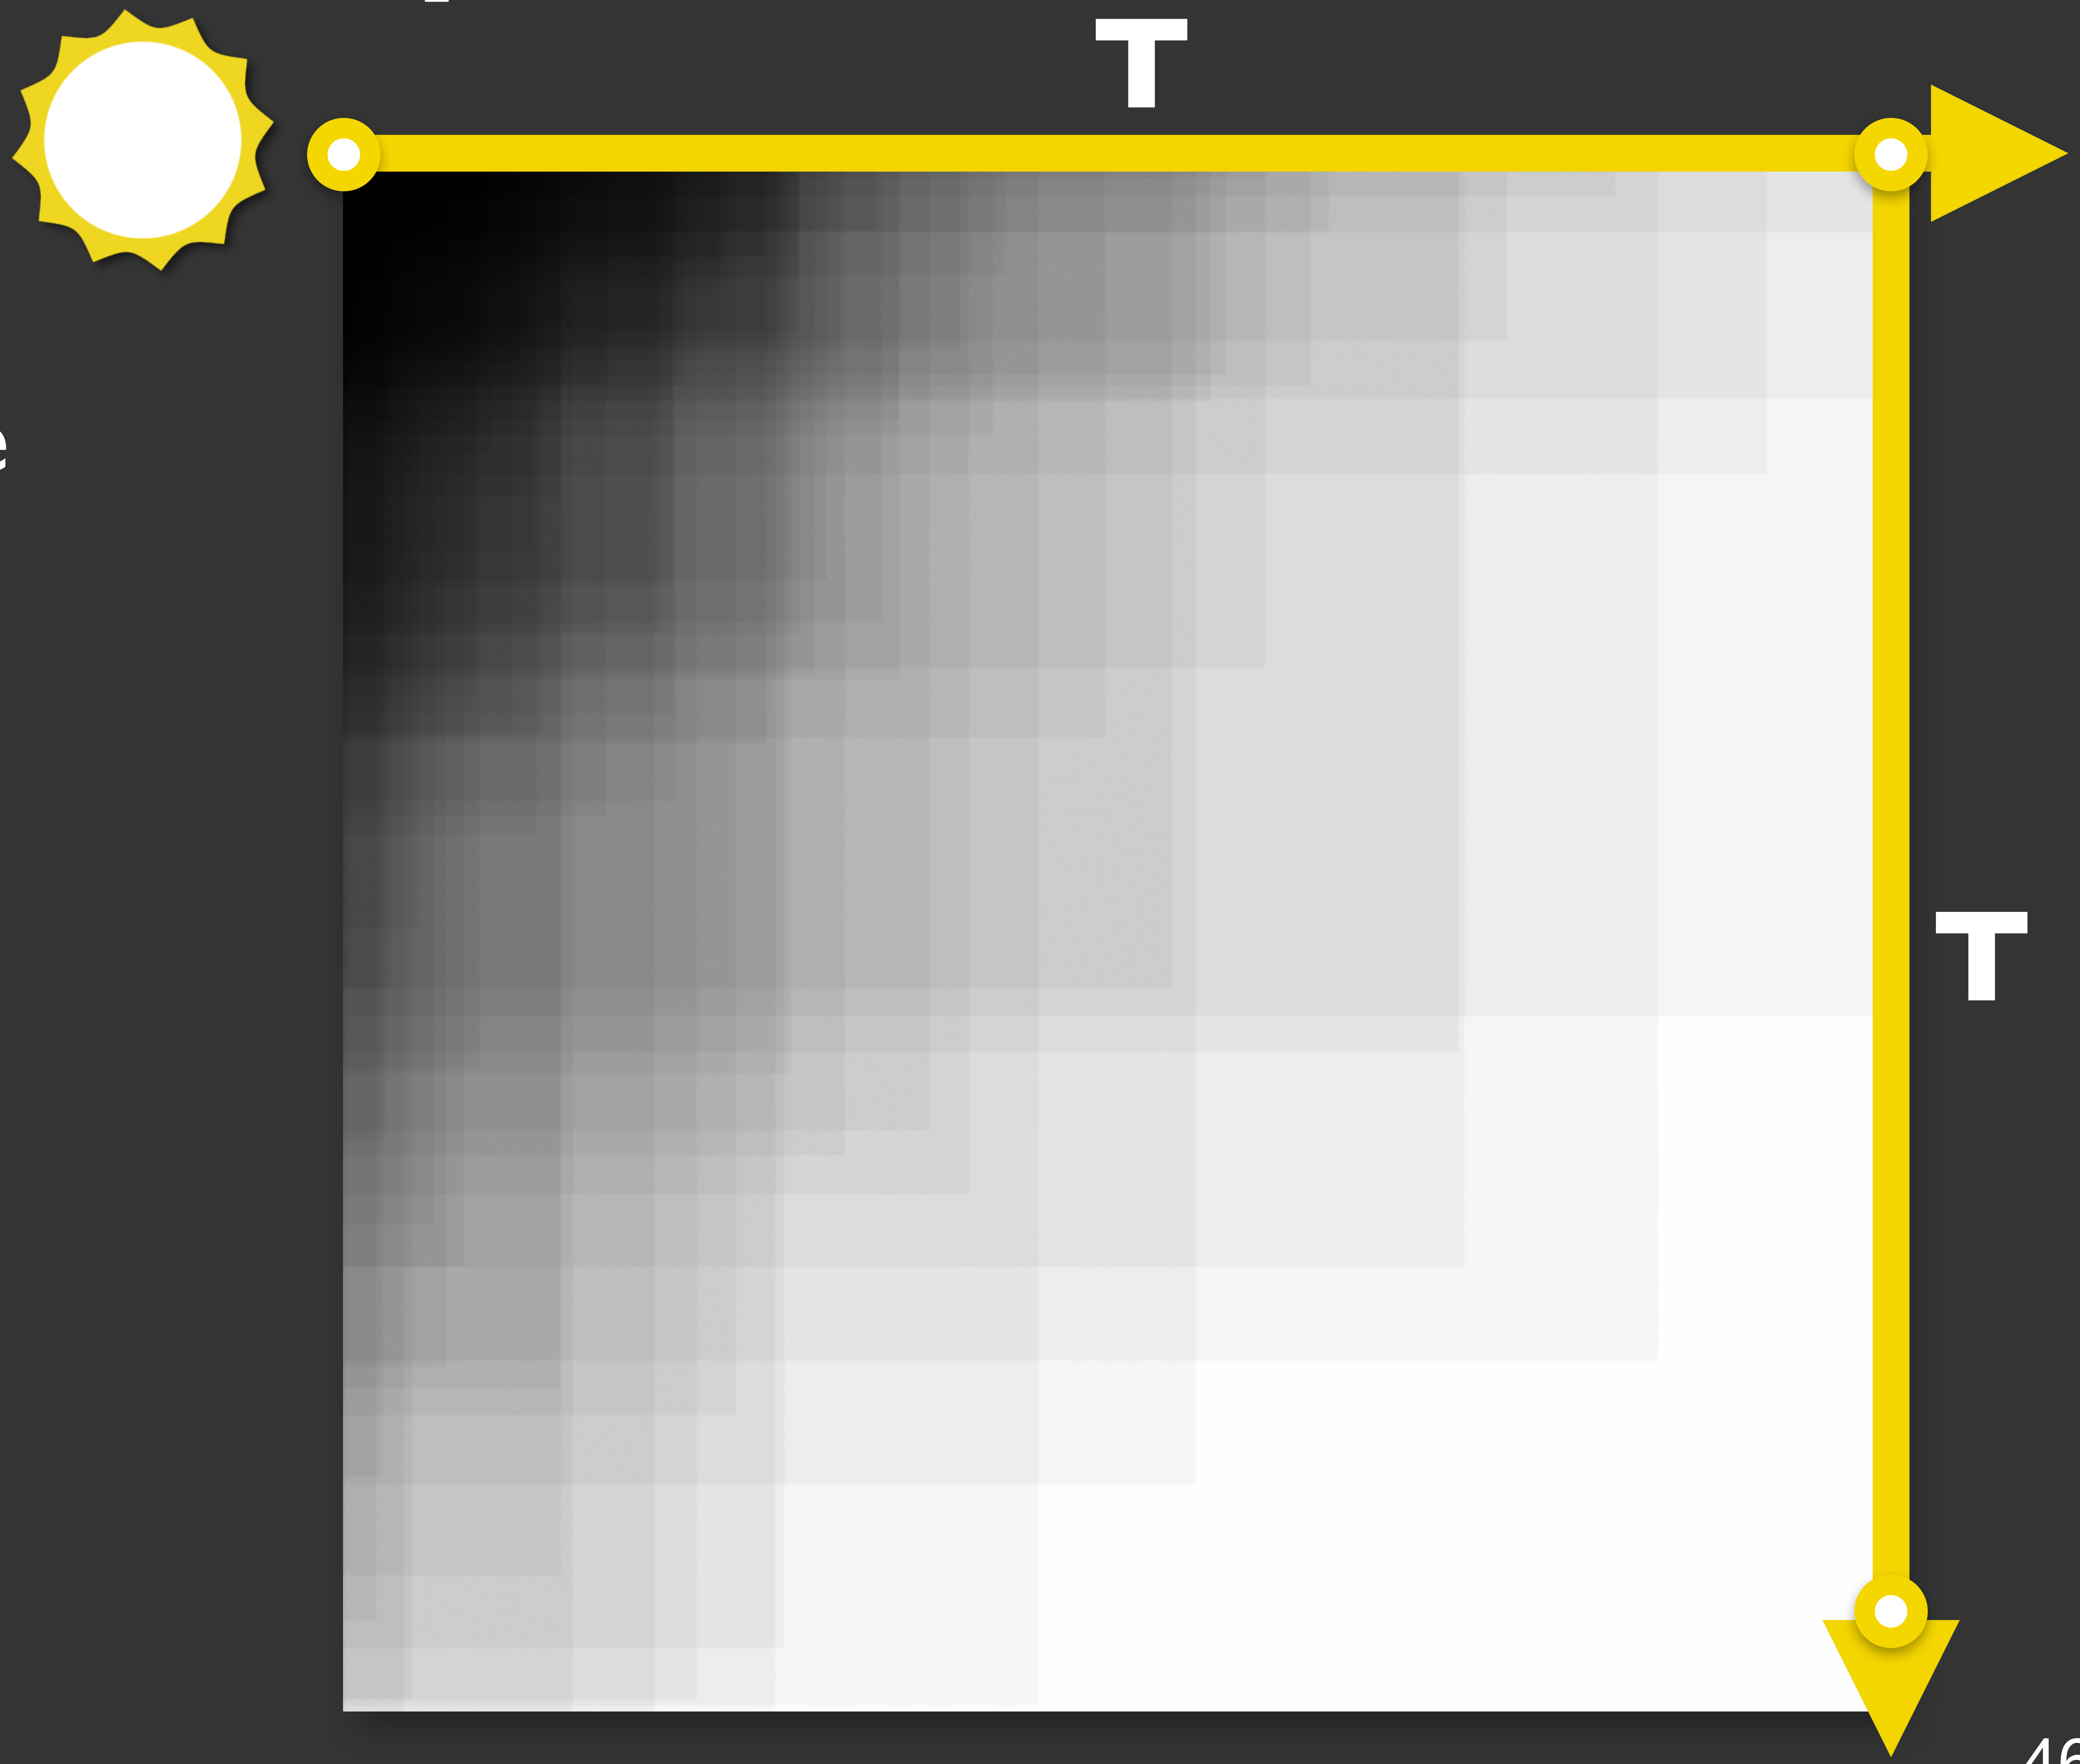

Error Analysis Setup

C
Collision

T
Track-Length

E
Expected Value

CTT
Short Planes

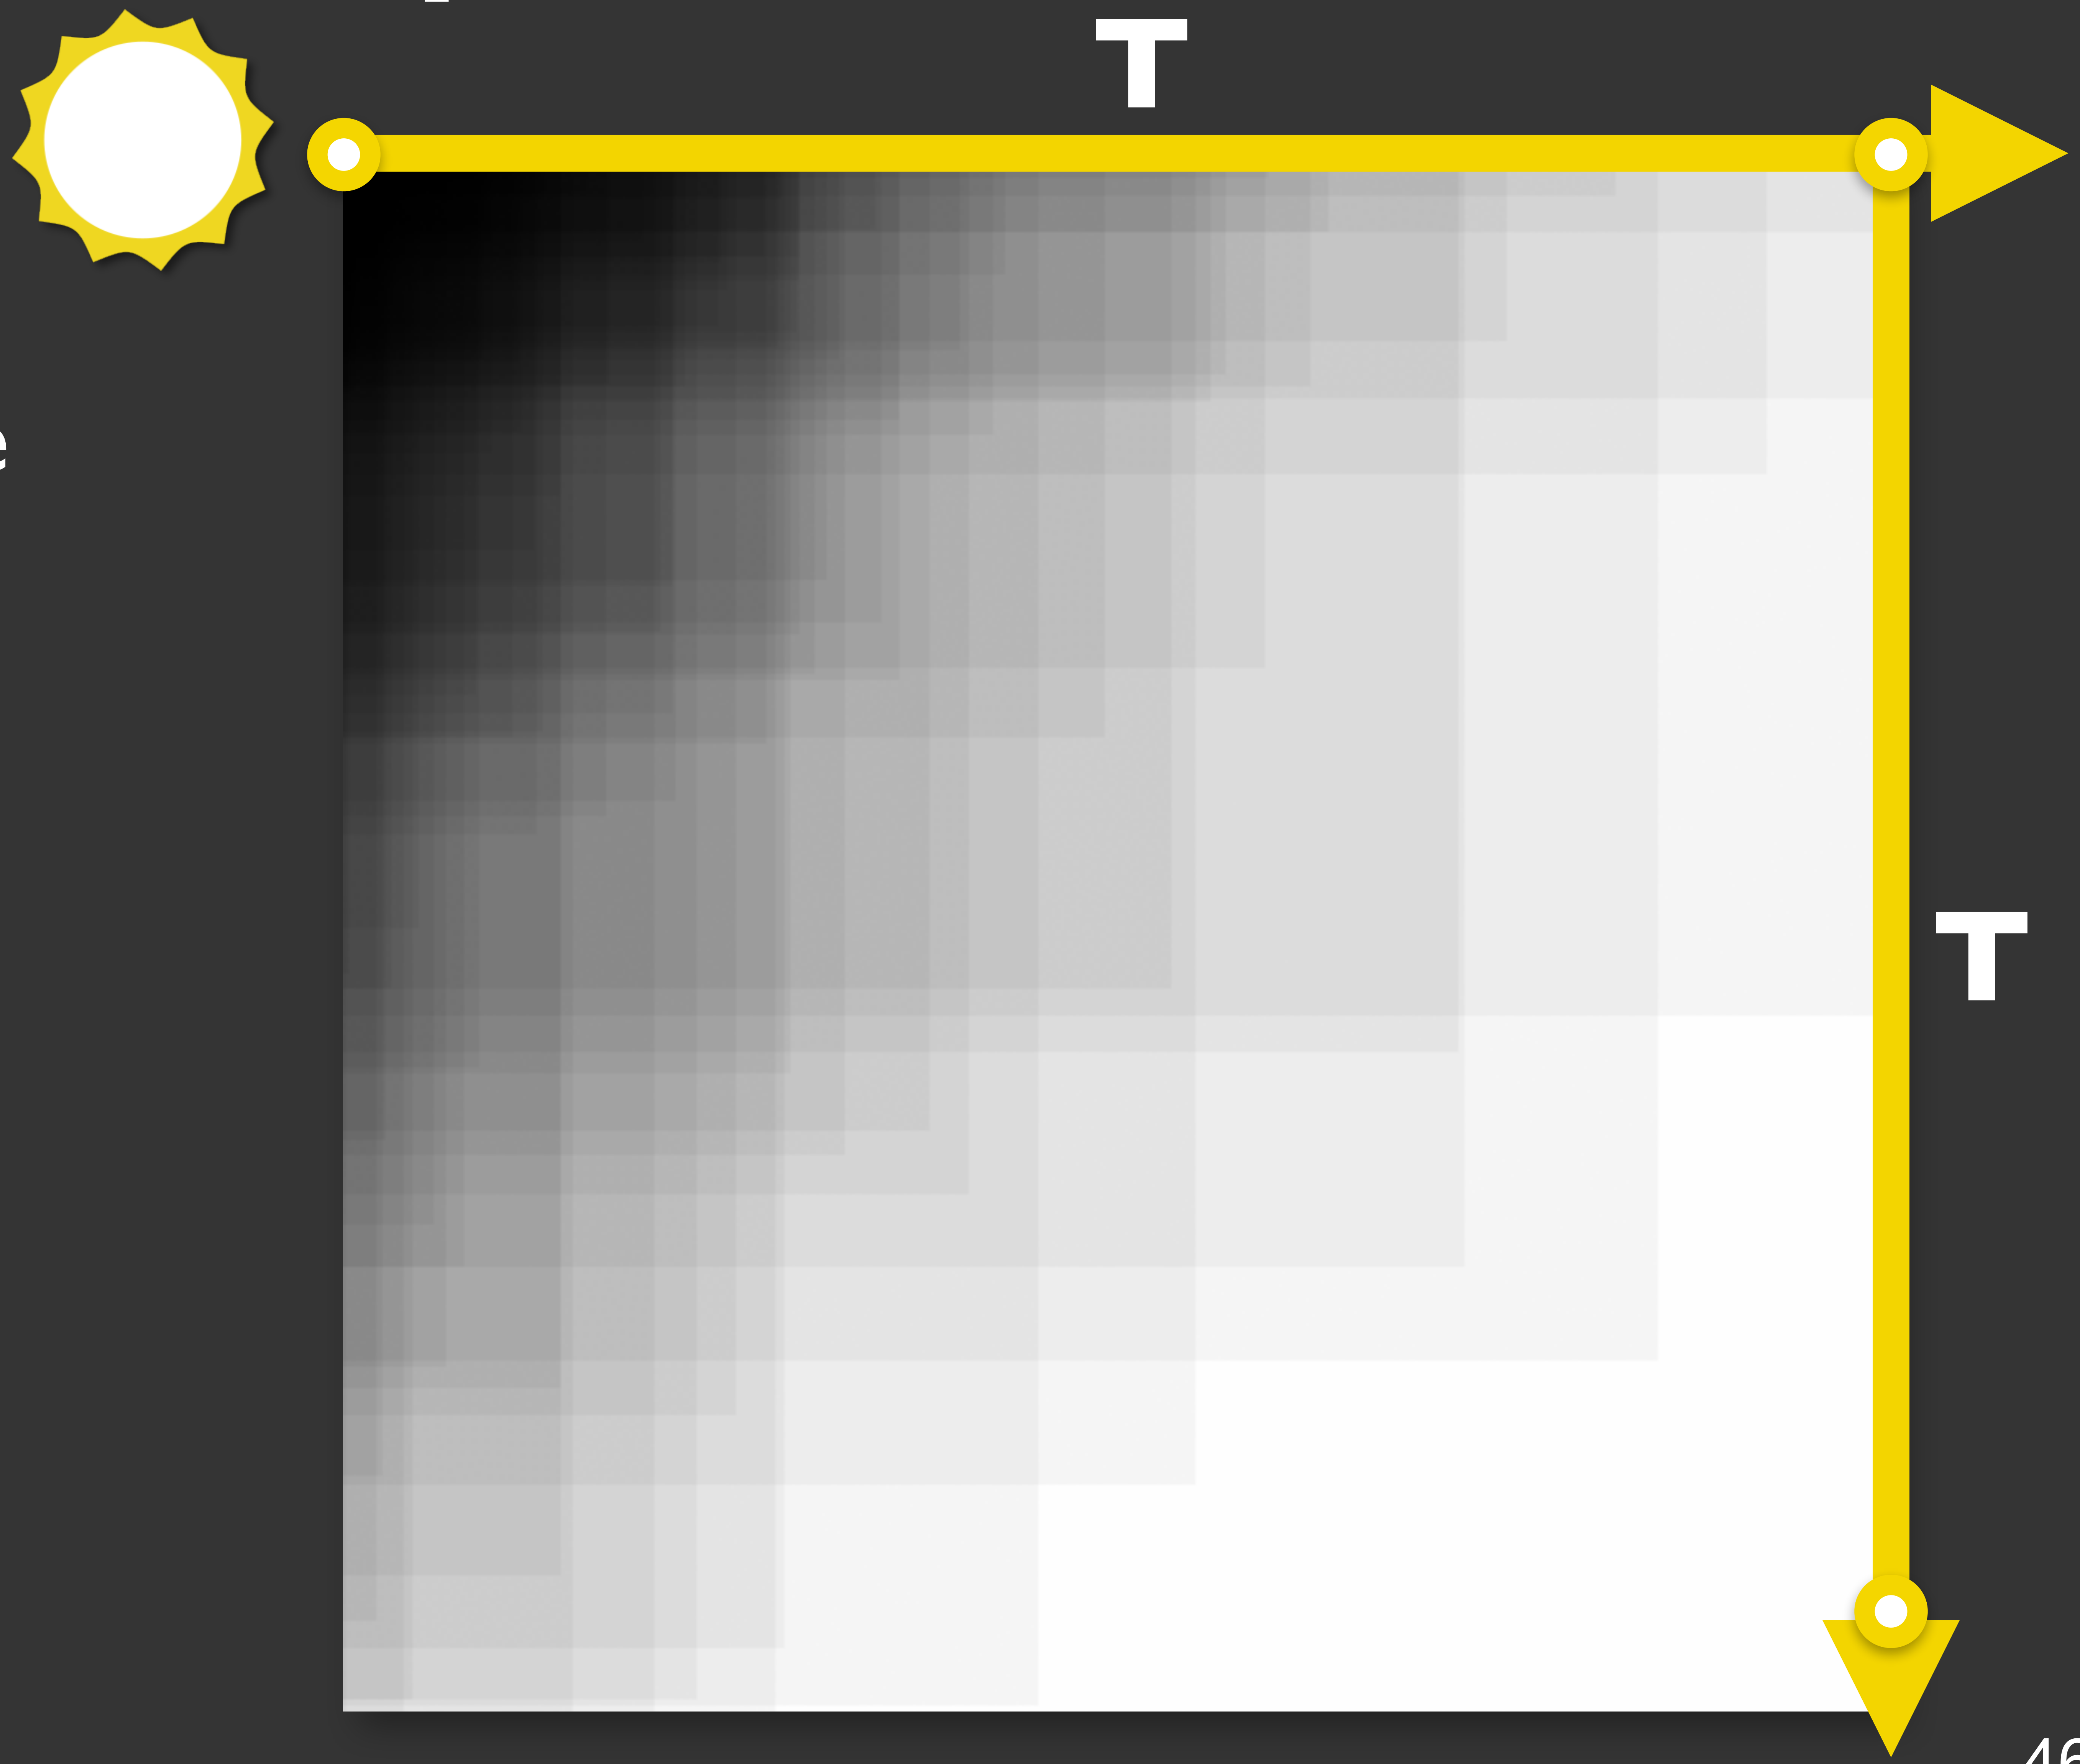

Error Analysis Results

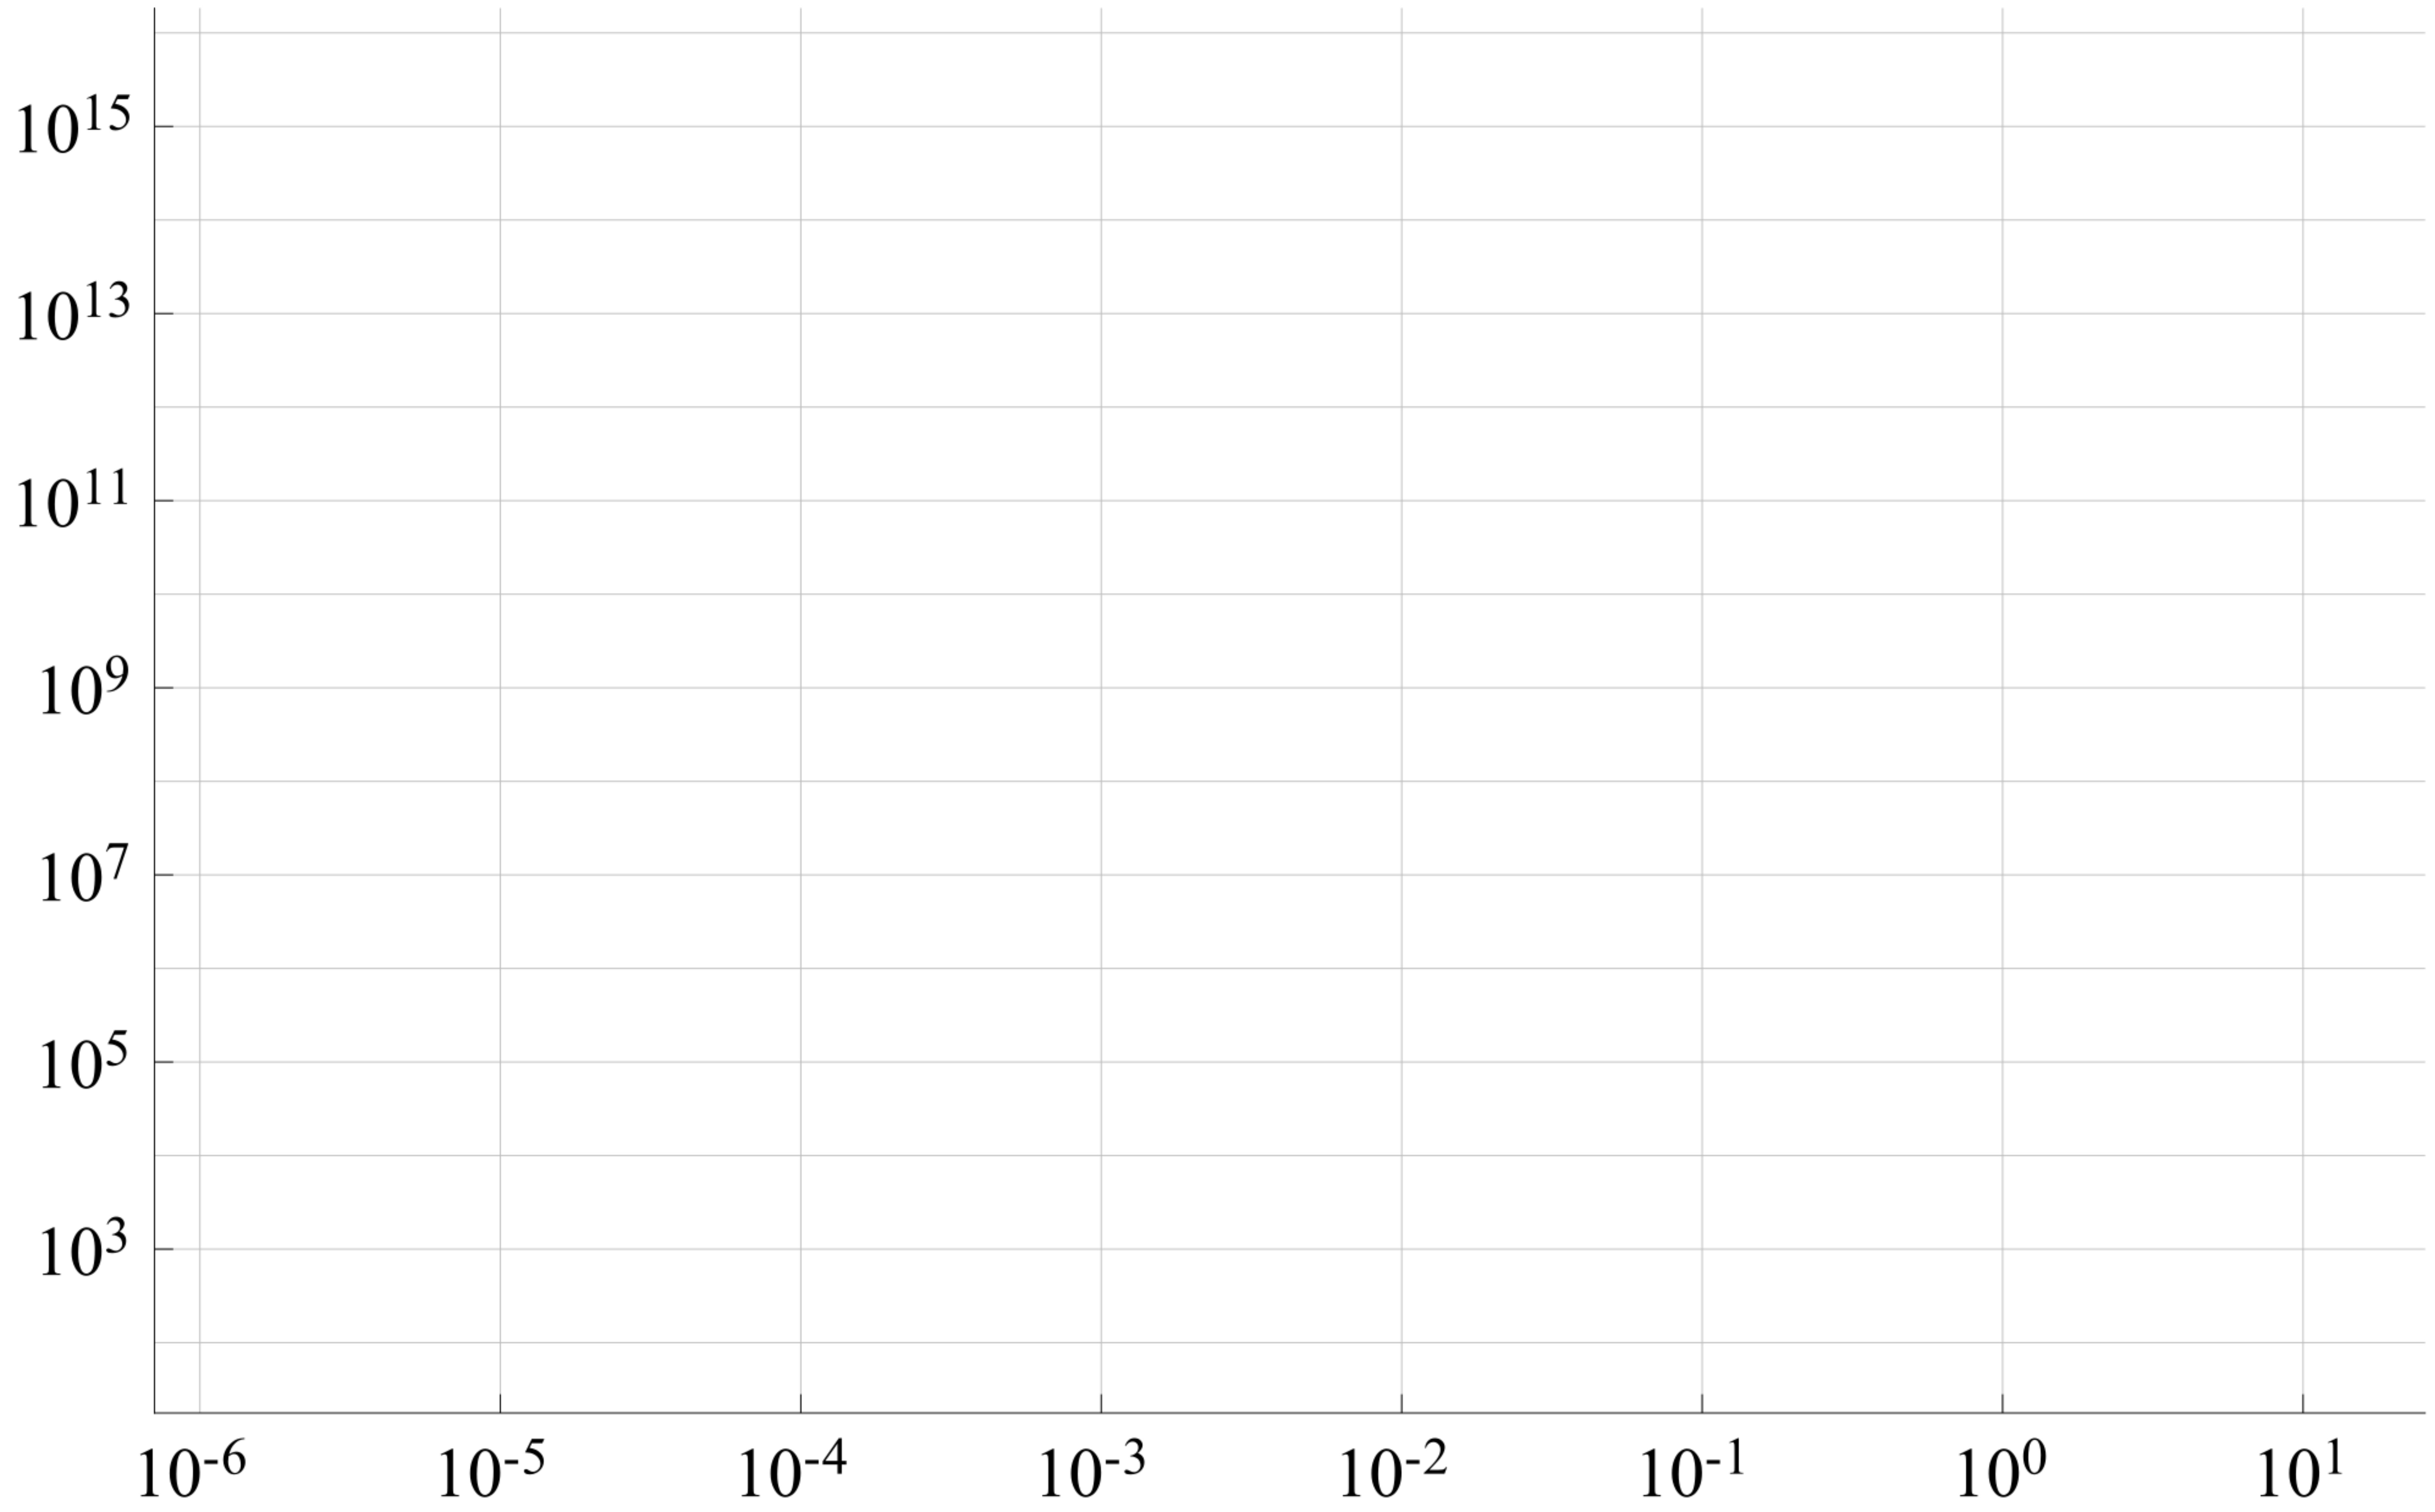

10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}

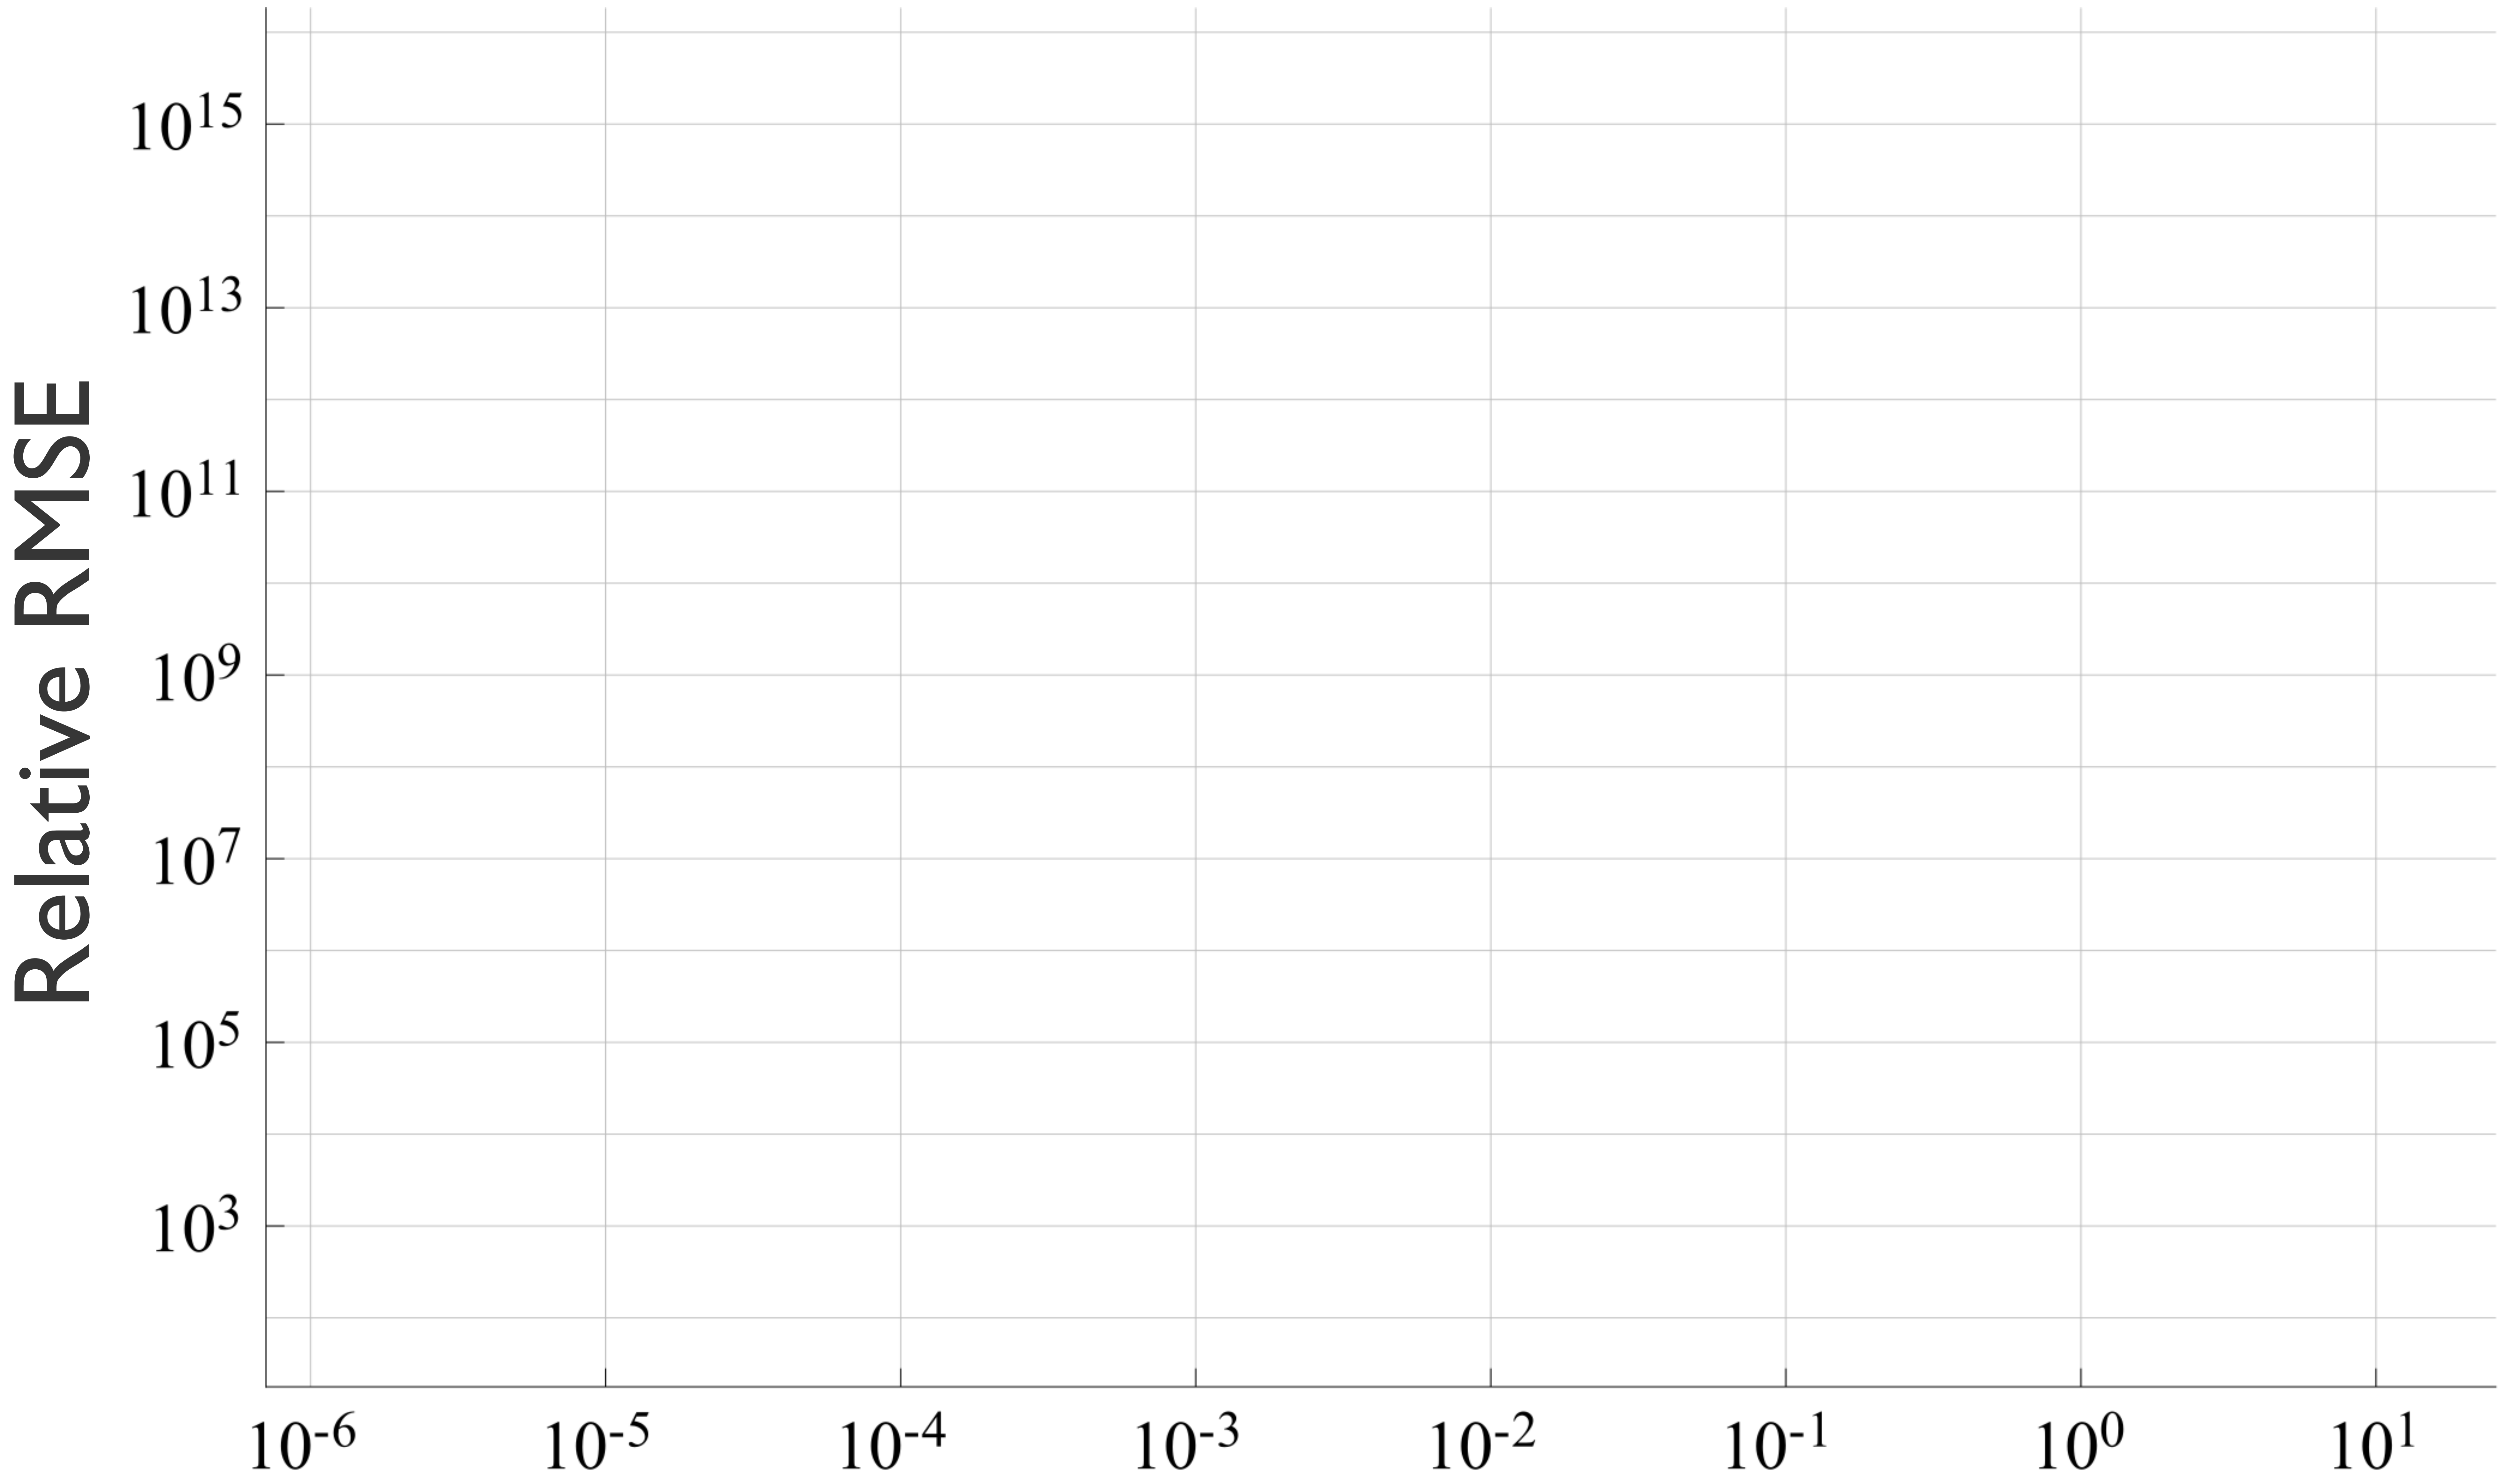
Relative RMSE
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}

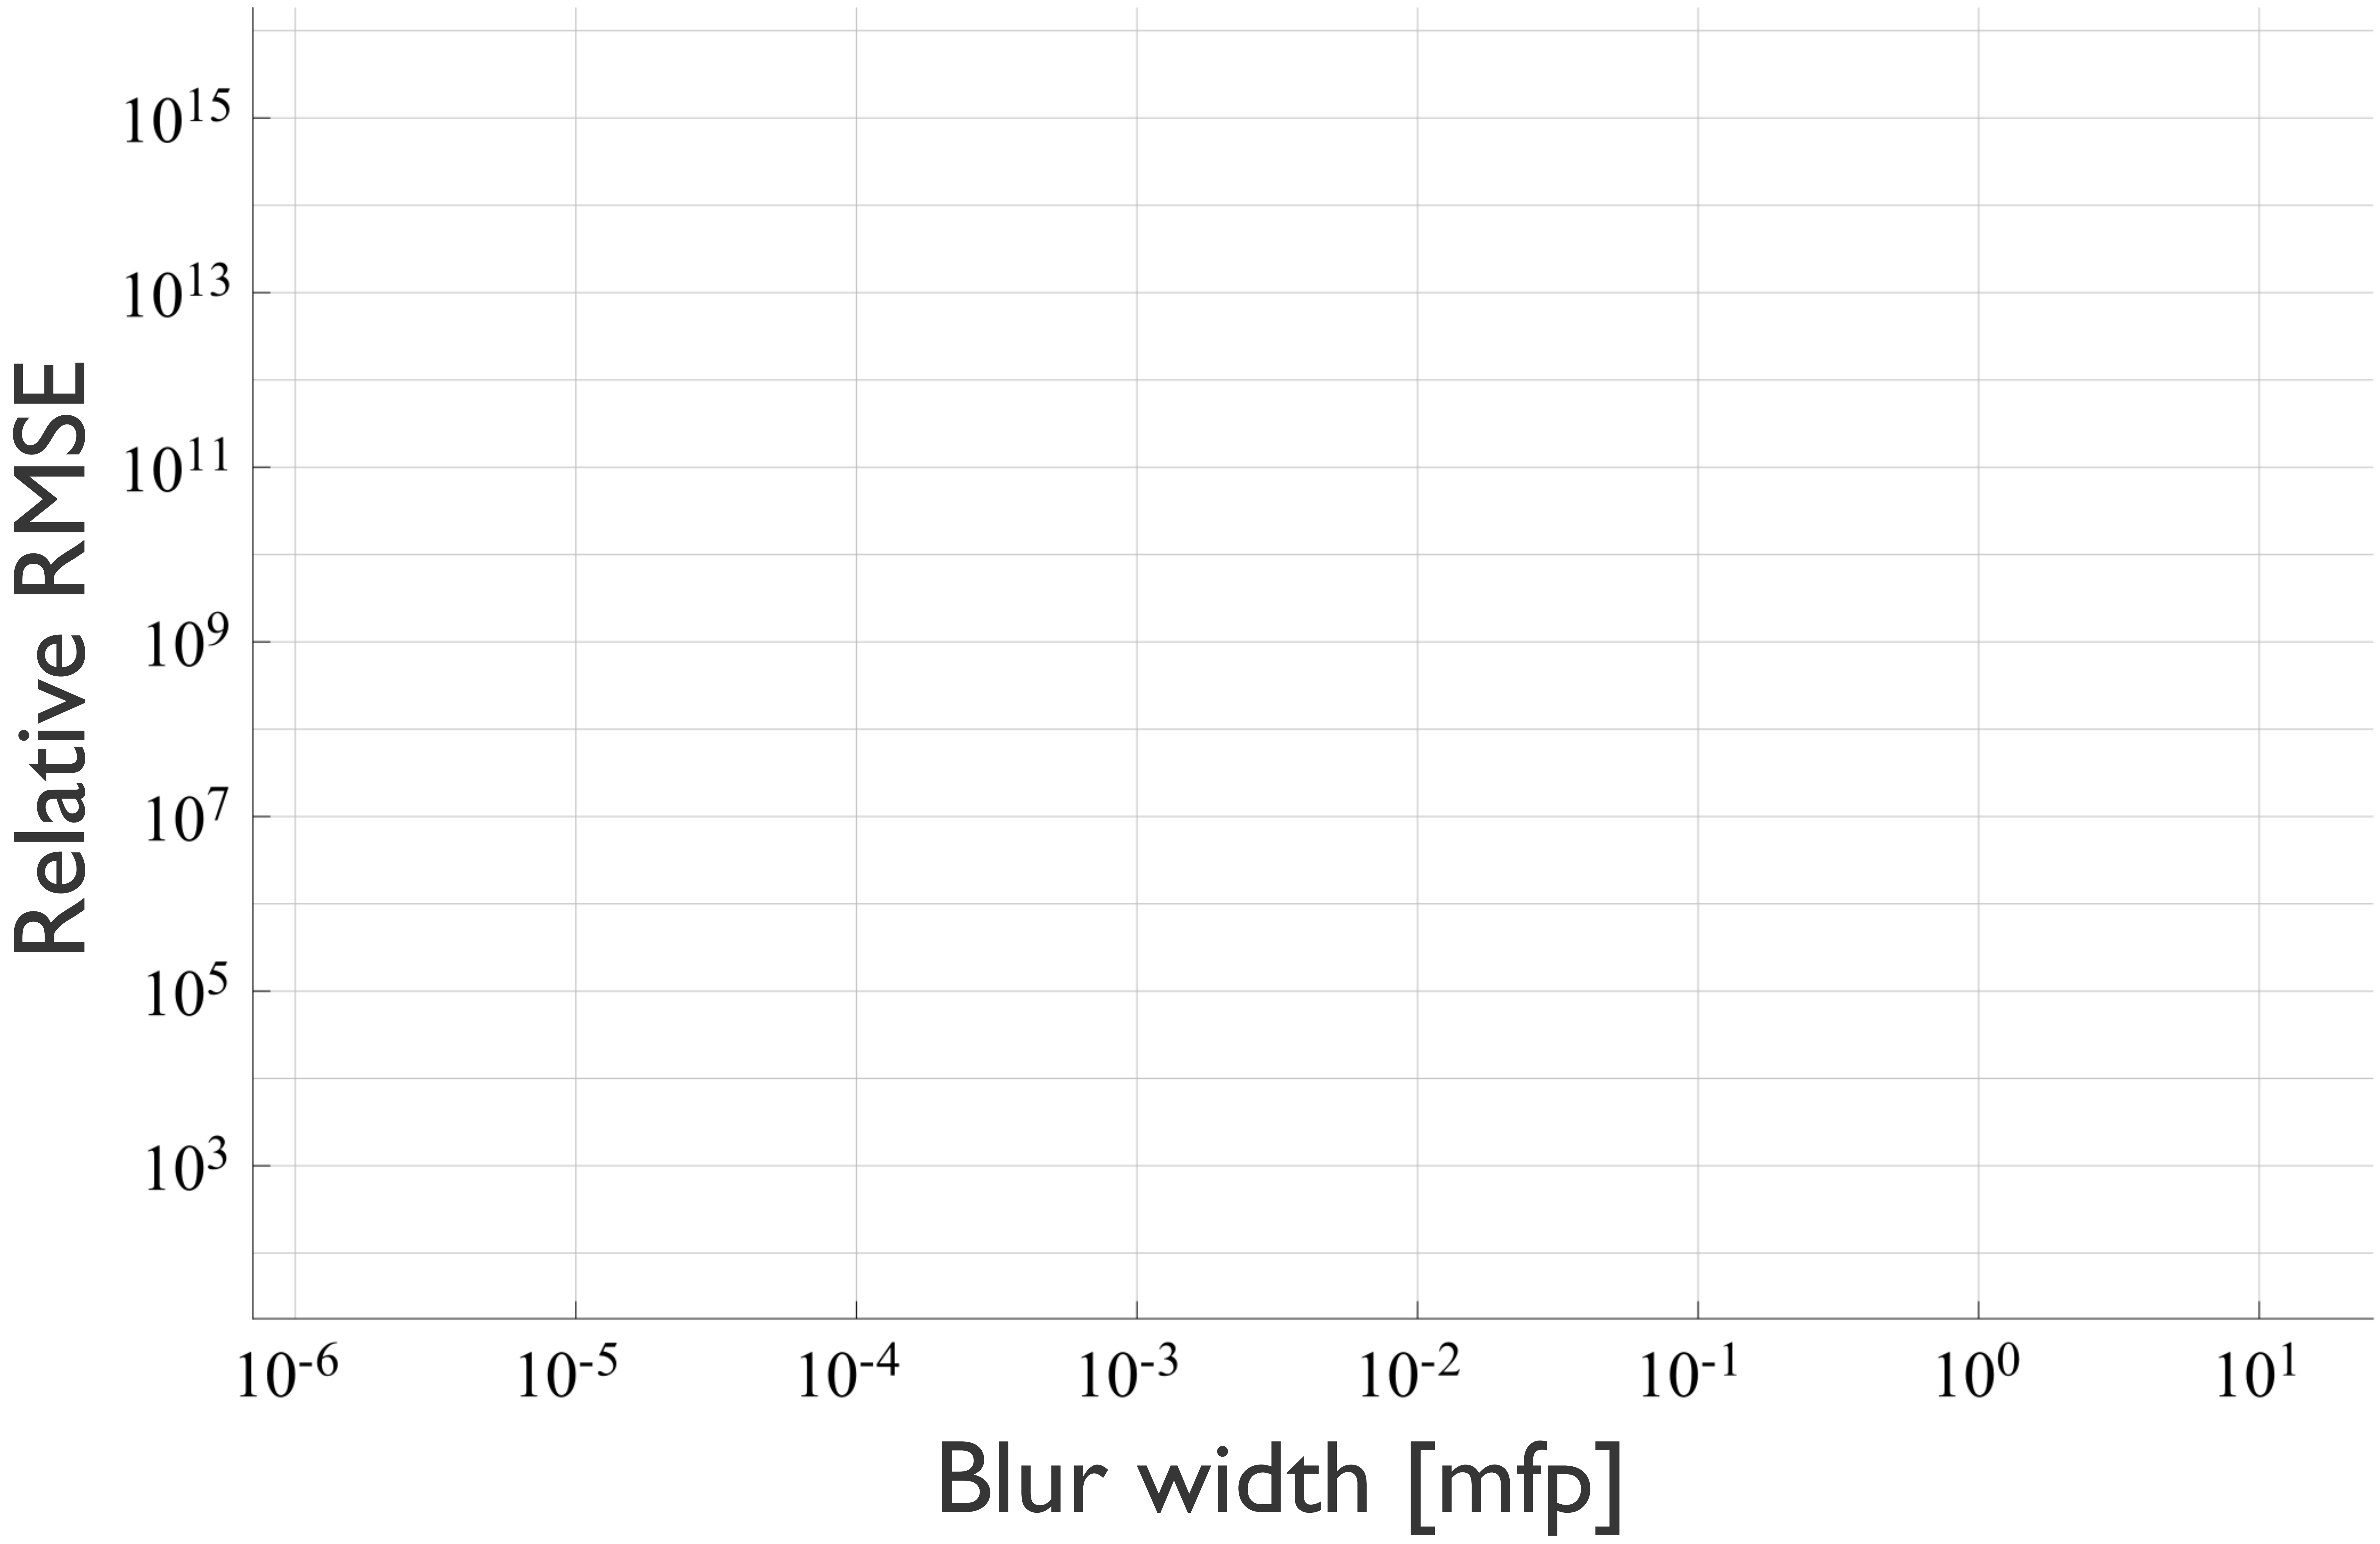

Relative RMSE
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
Blur width [mfp]

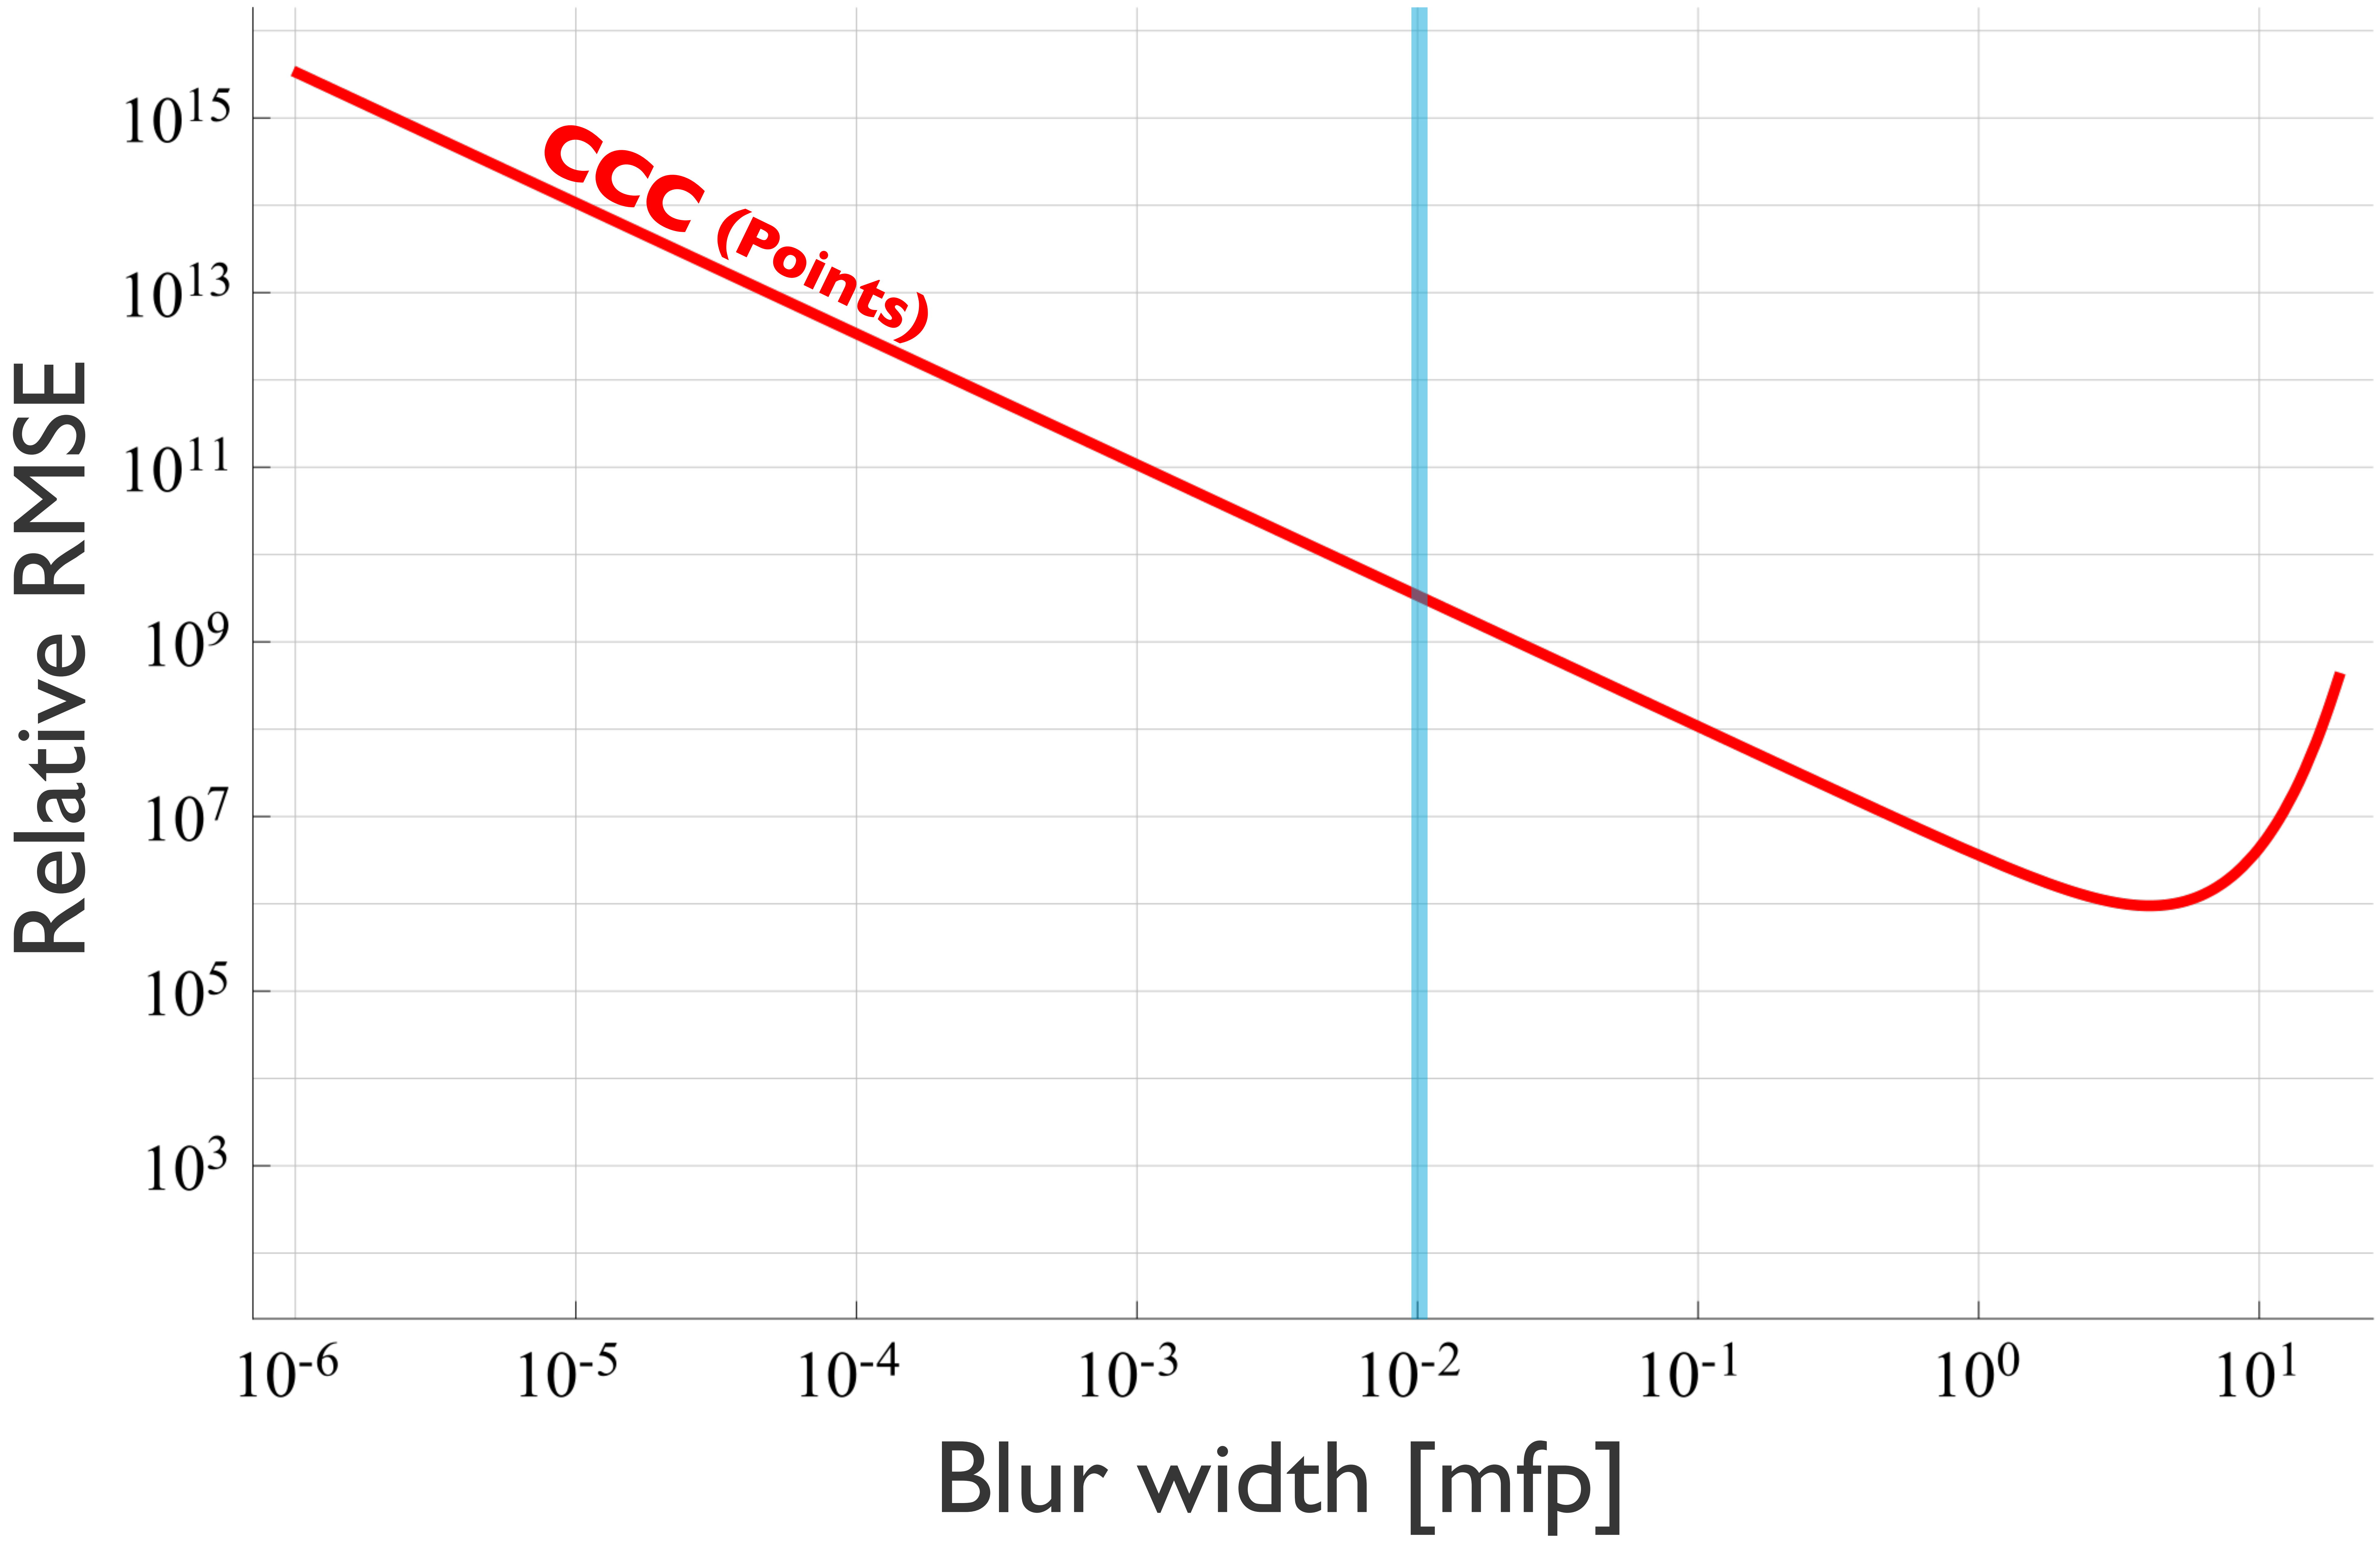
CCC (Points)
Relative RMSE
10^15
10^13
10^11
10^9
10^7
10^5
10^3
10^-6
10^-5
10^-4
10^-3
10^-2
10^-1
10^0
10^1
Blur width [mfp]

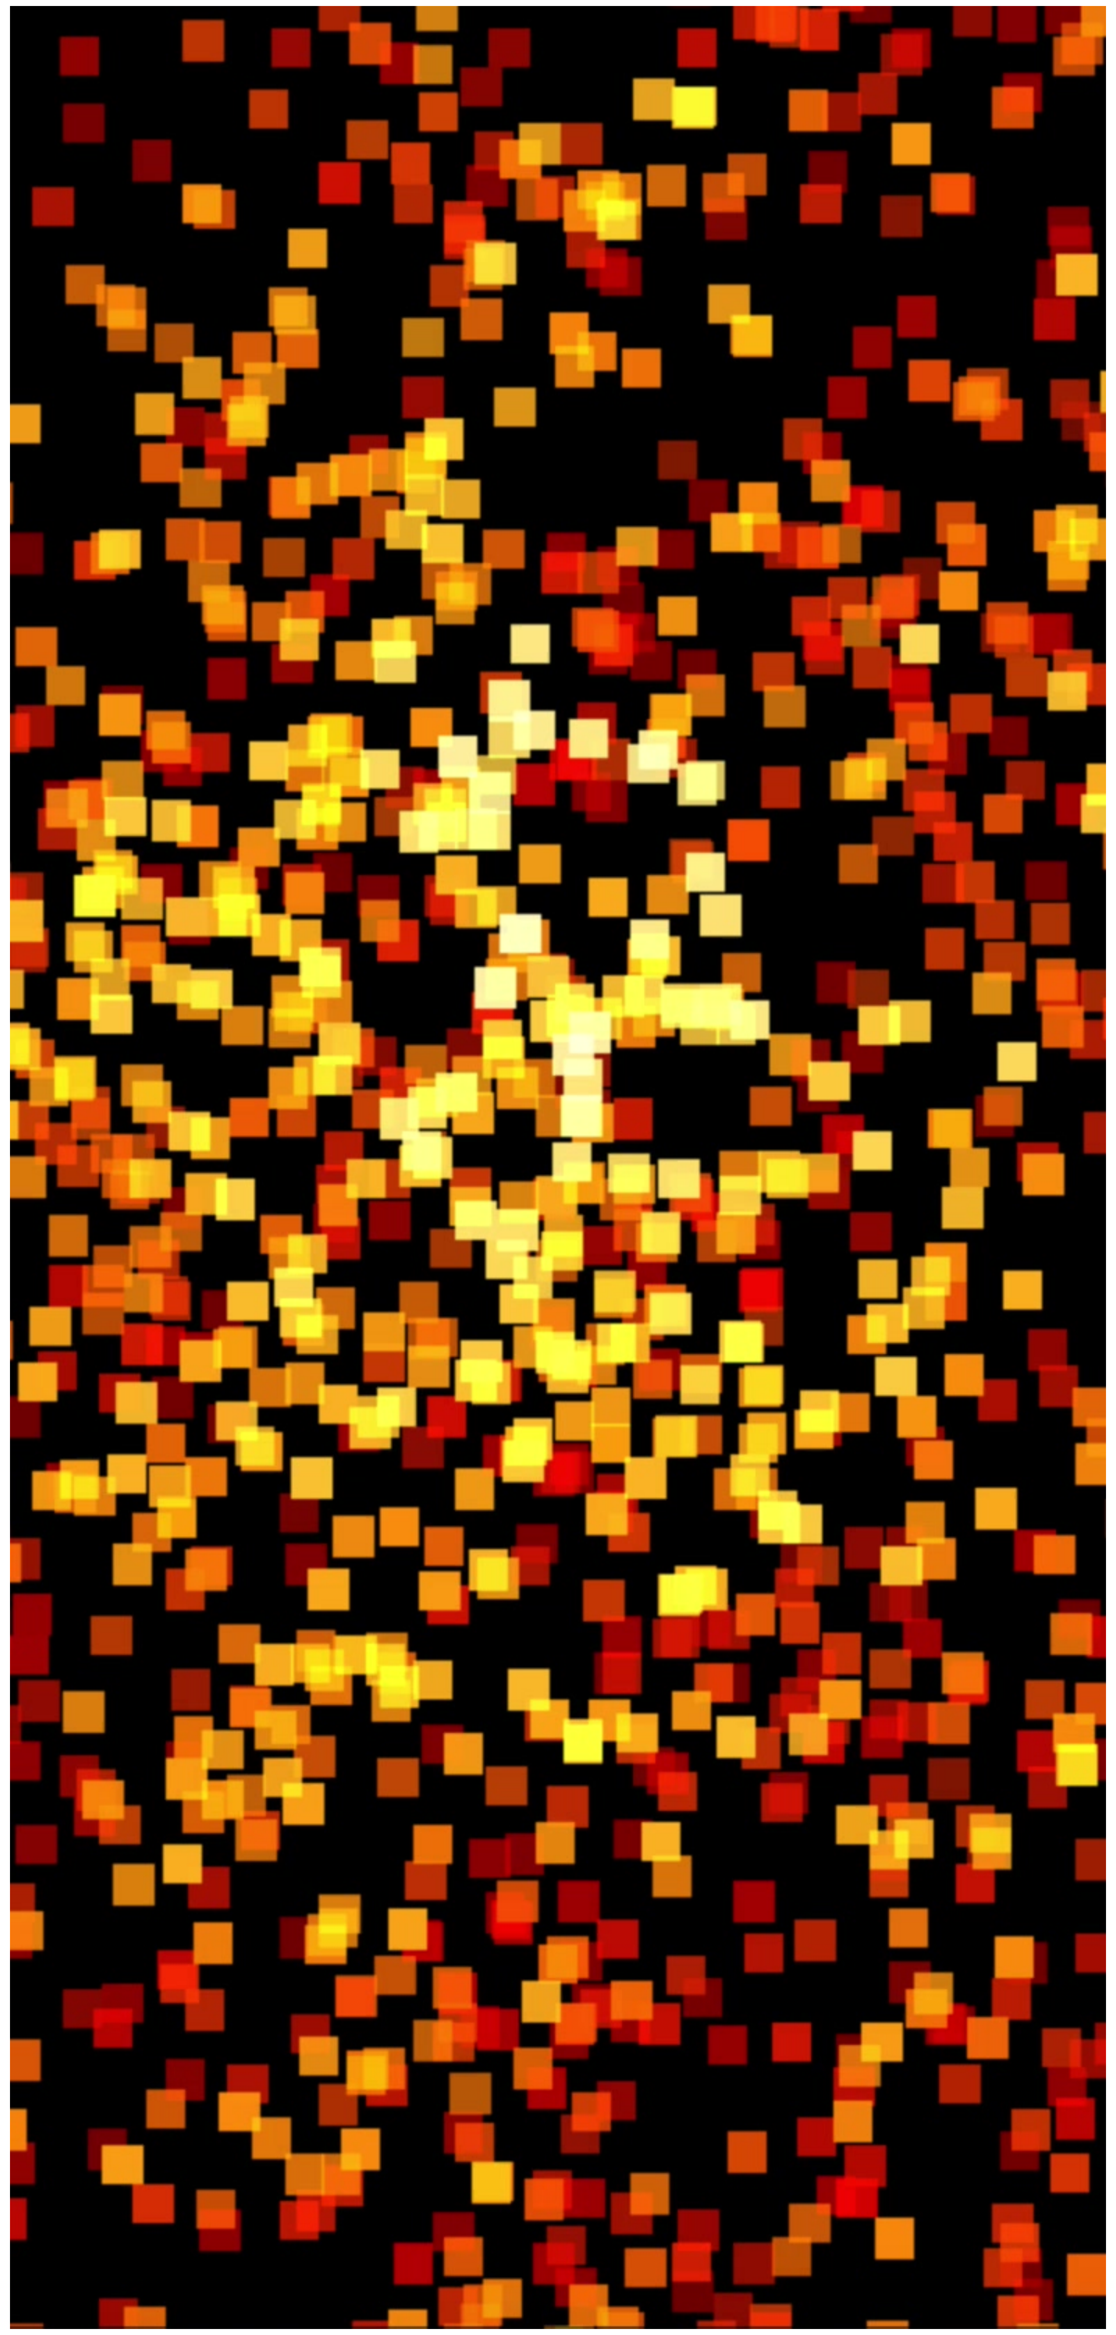


Relative RMSE
10^15
10^13
10^11
10^9
10^7
10^5
10^3
CCC (Points)
10^-6
10^-5
10^-4
10^-3
10^-2
10^-1
10^0
10^1
Blur width [mfp]

Relative RMSE
Blur width [mfp]
CCC (Points)
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}

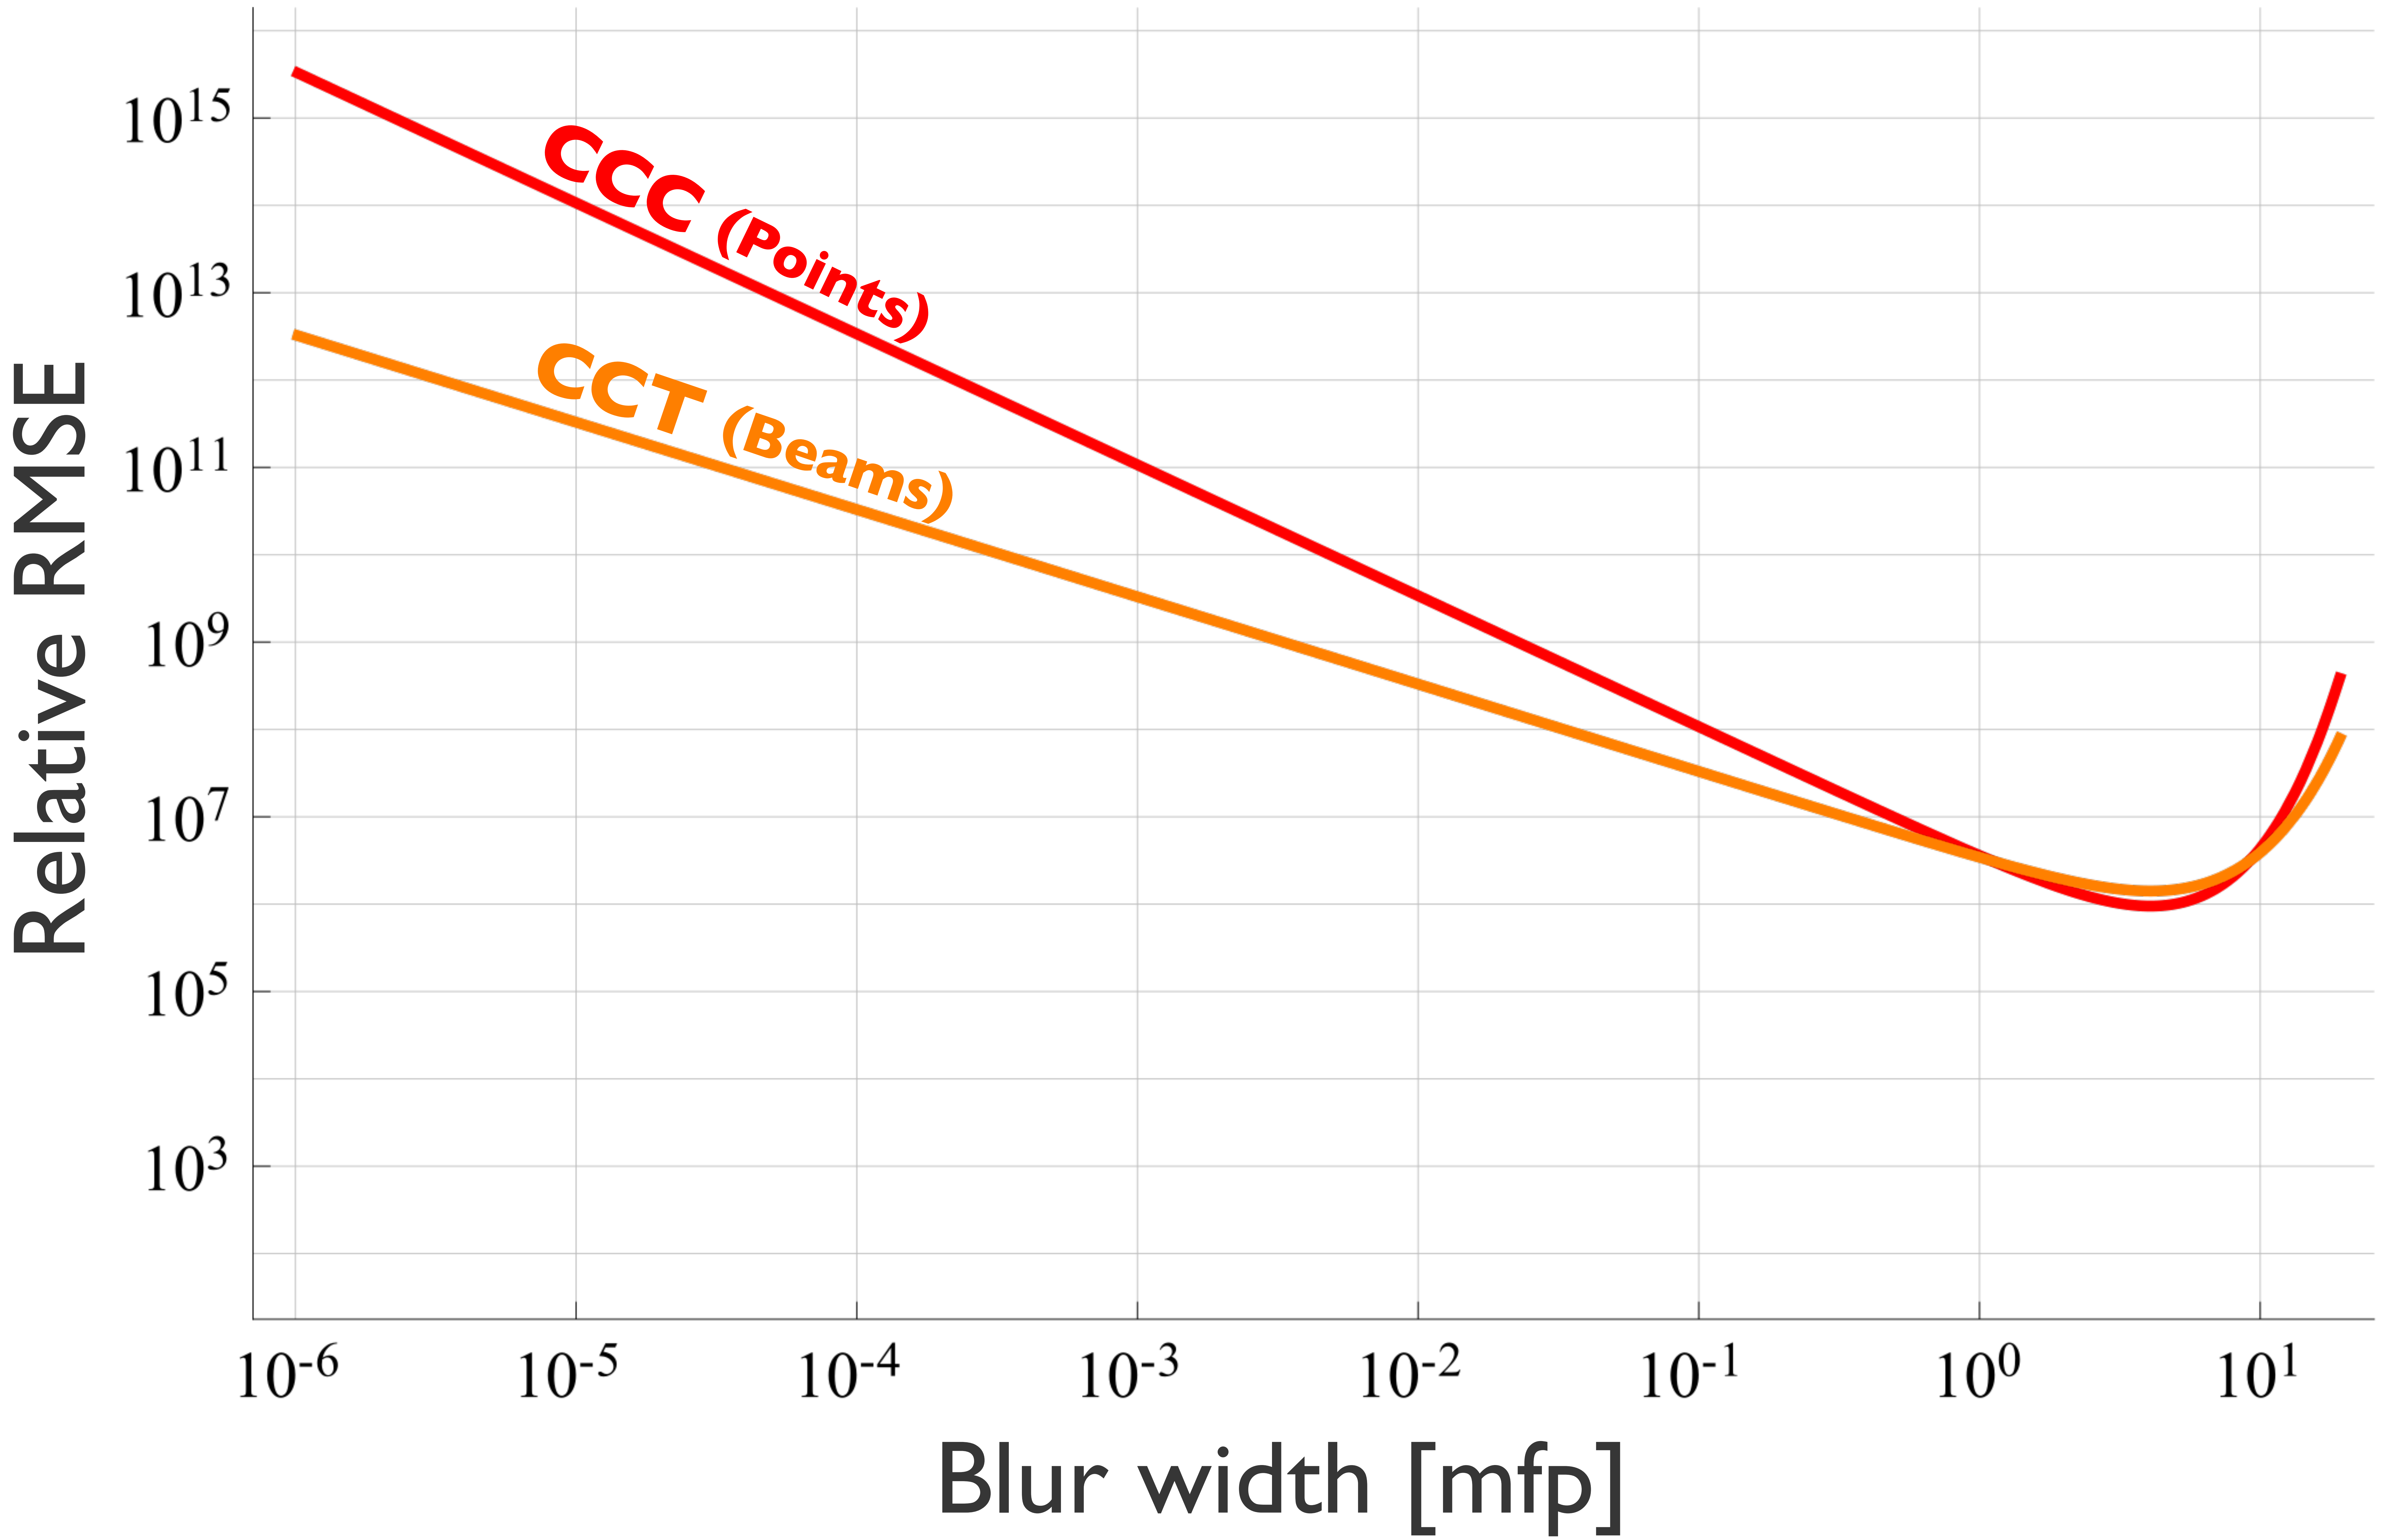

Relative RMSE
10^15
10^13
10^11
10^9
10^7
10^5
10^3
CCC (Points)
CCT (Beams)
10^-6
10^-5
10^-4
10^-3
10^-2
10^-1
10^0
10^1
Blur width [mfp]

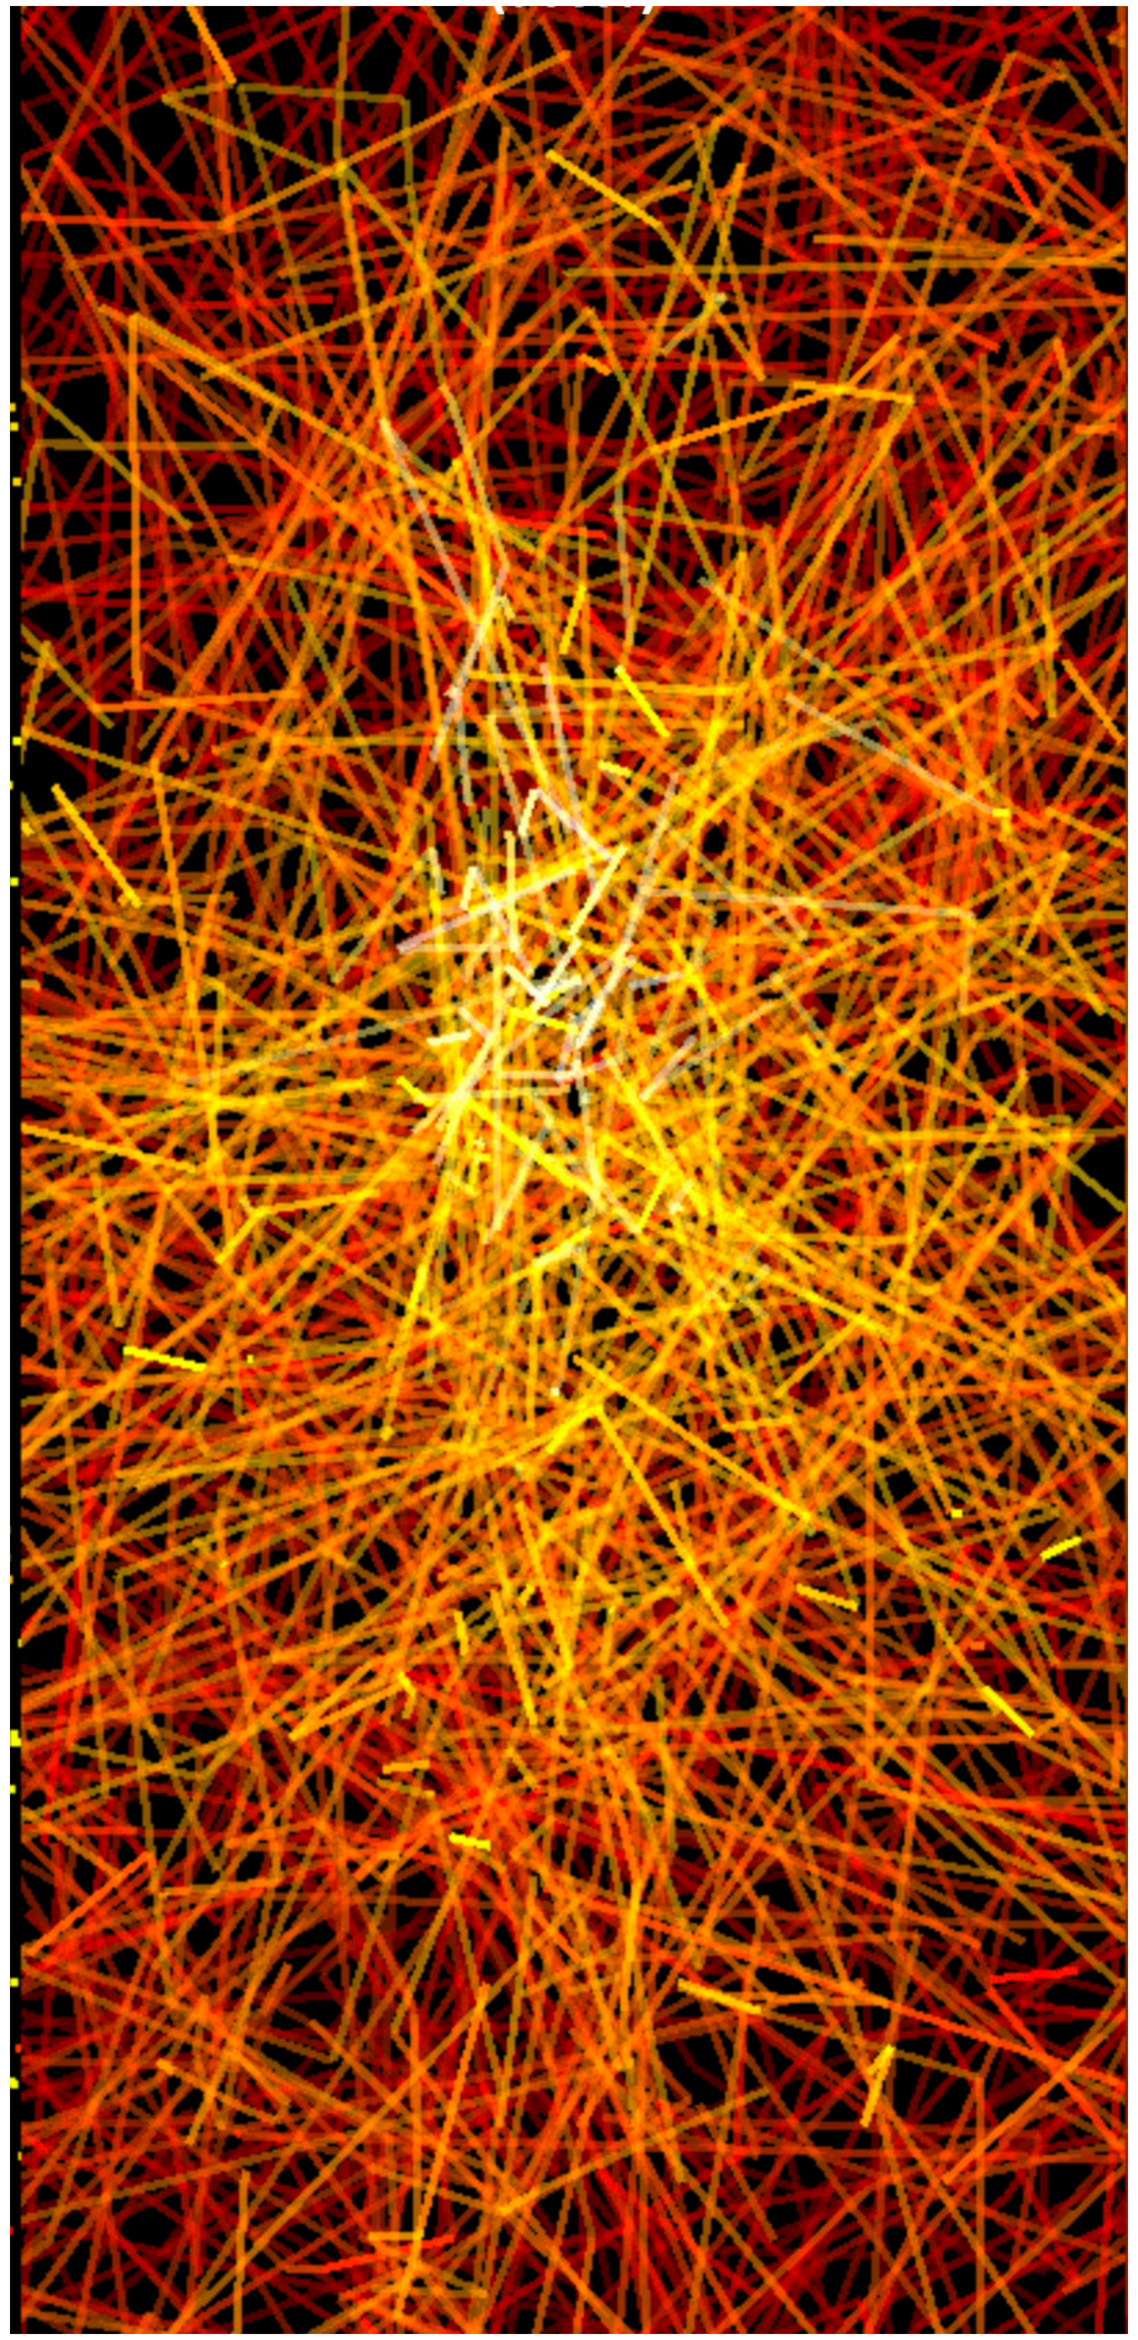

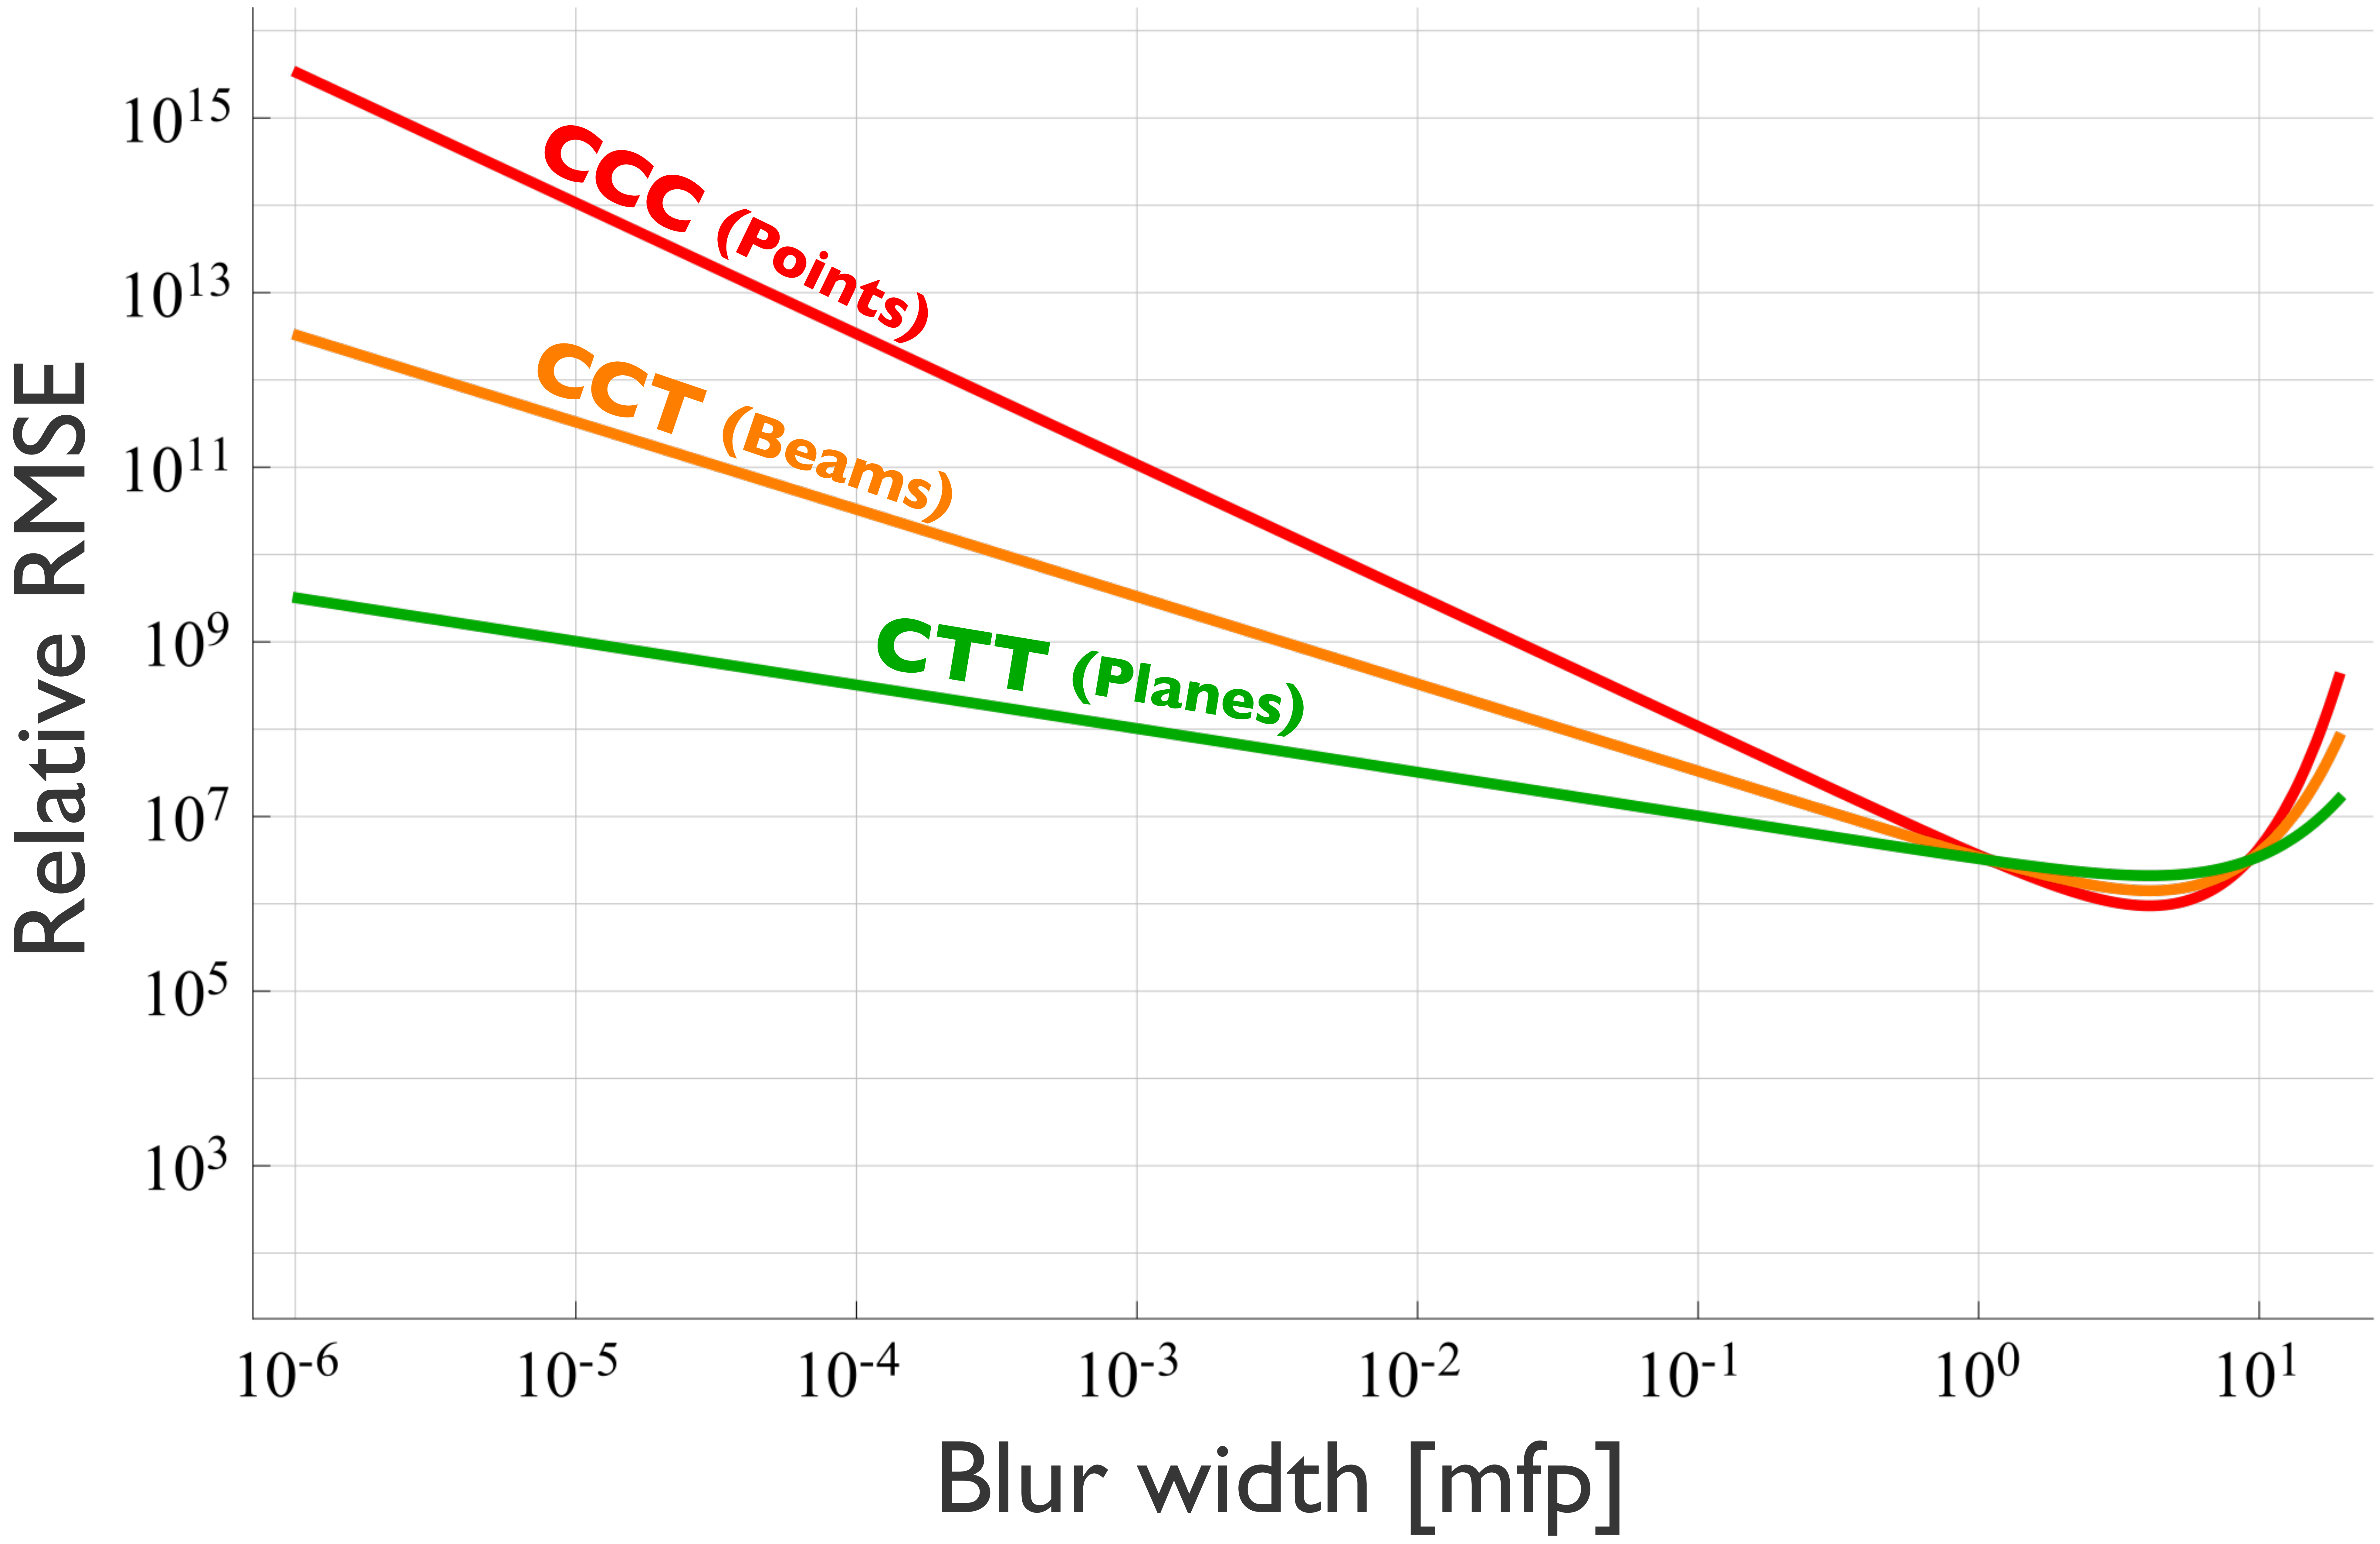

CCC (Points)
CCT (Beams)
CTT (Planes)
Relative RMSE
Blur width [mfp]

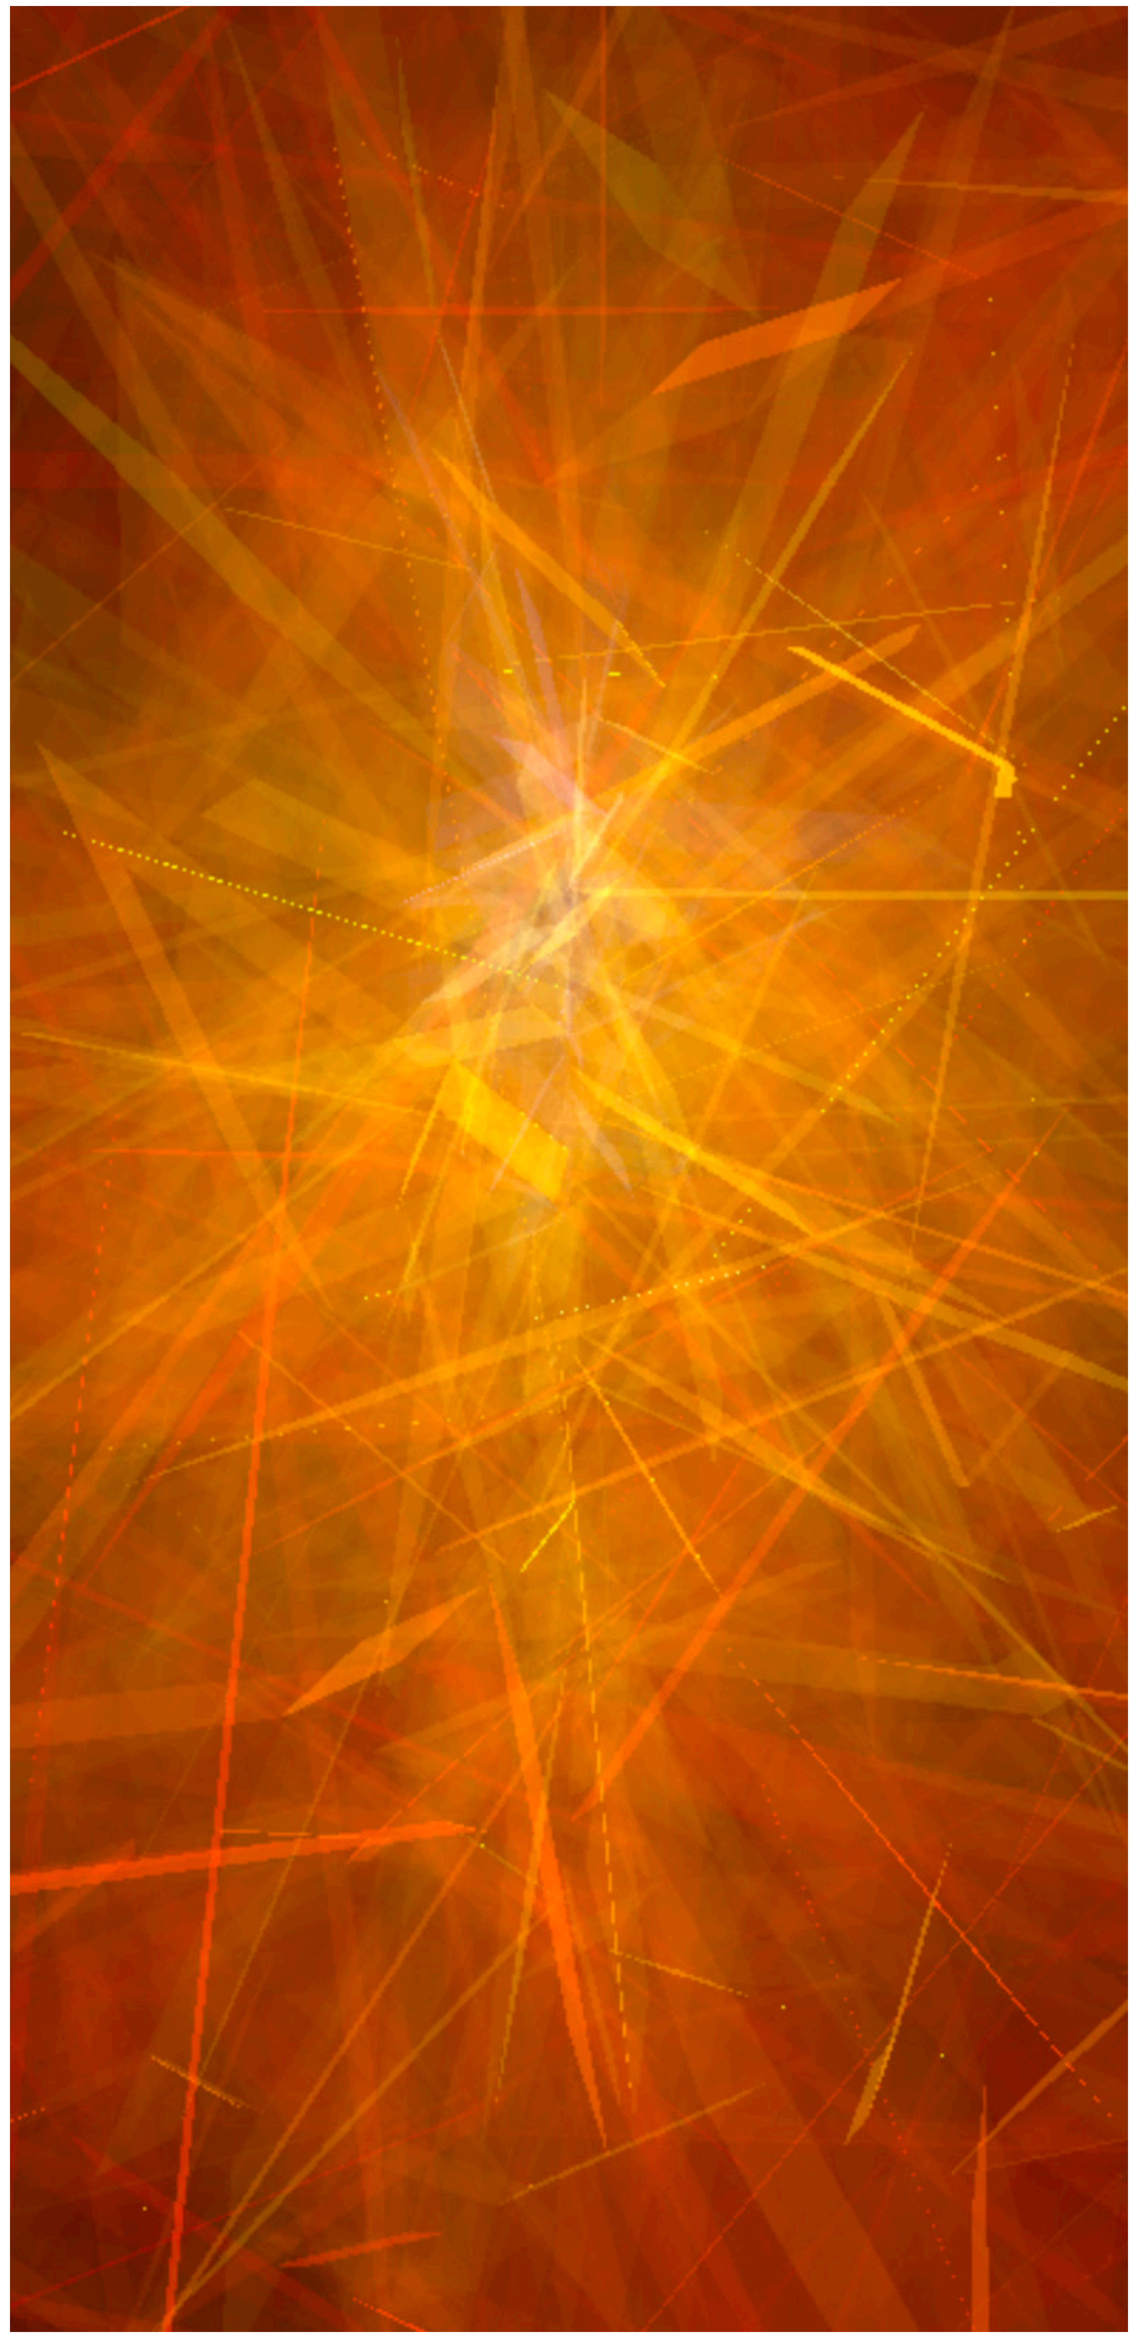

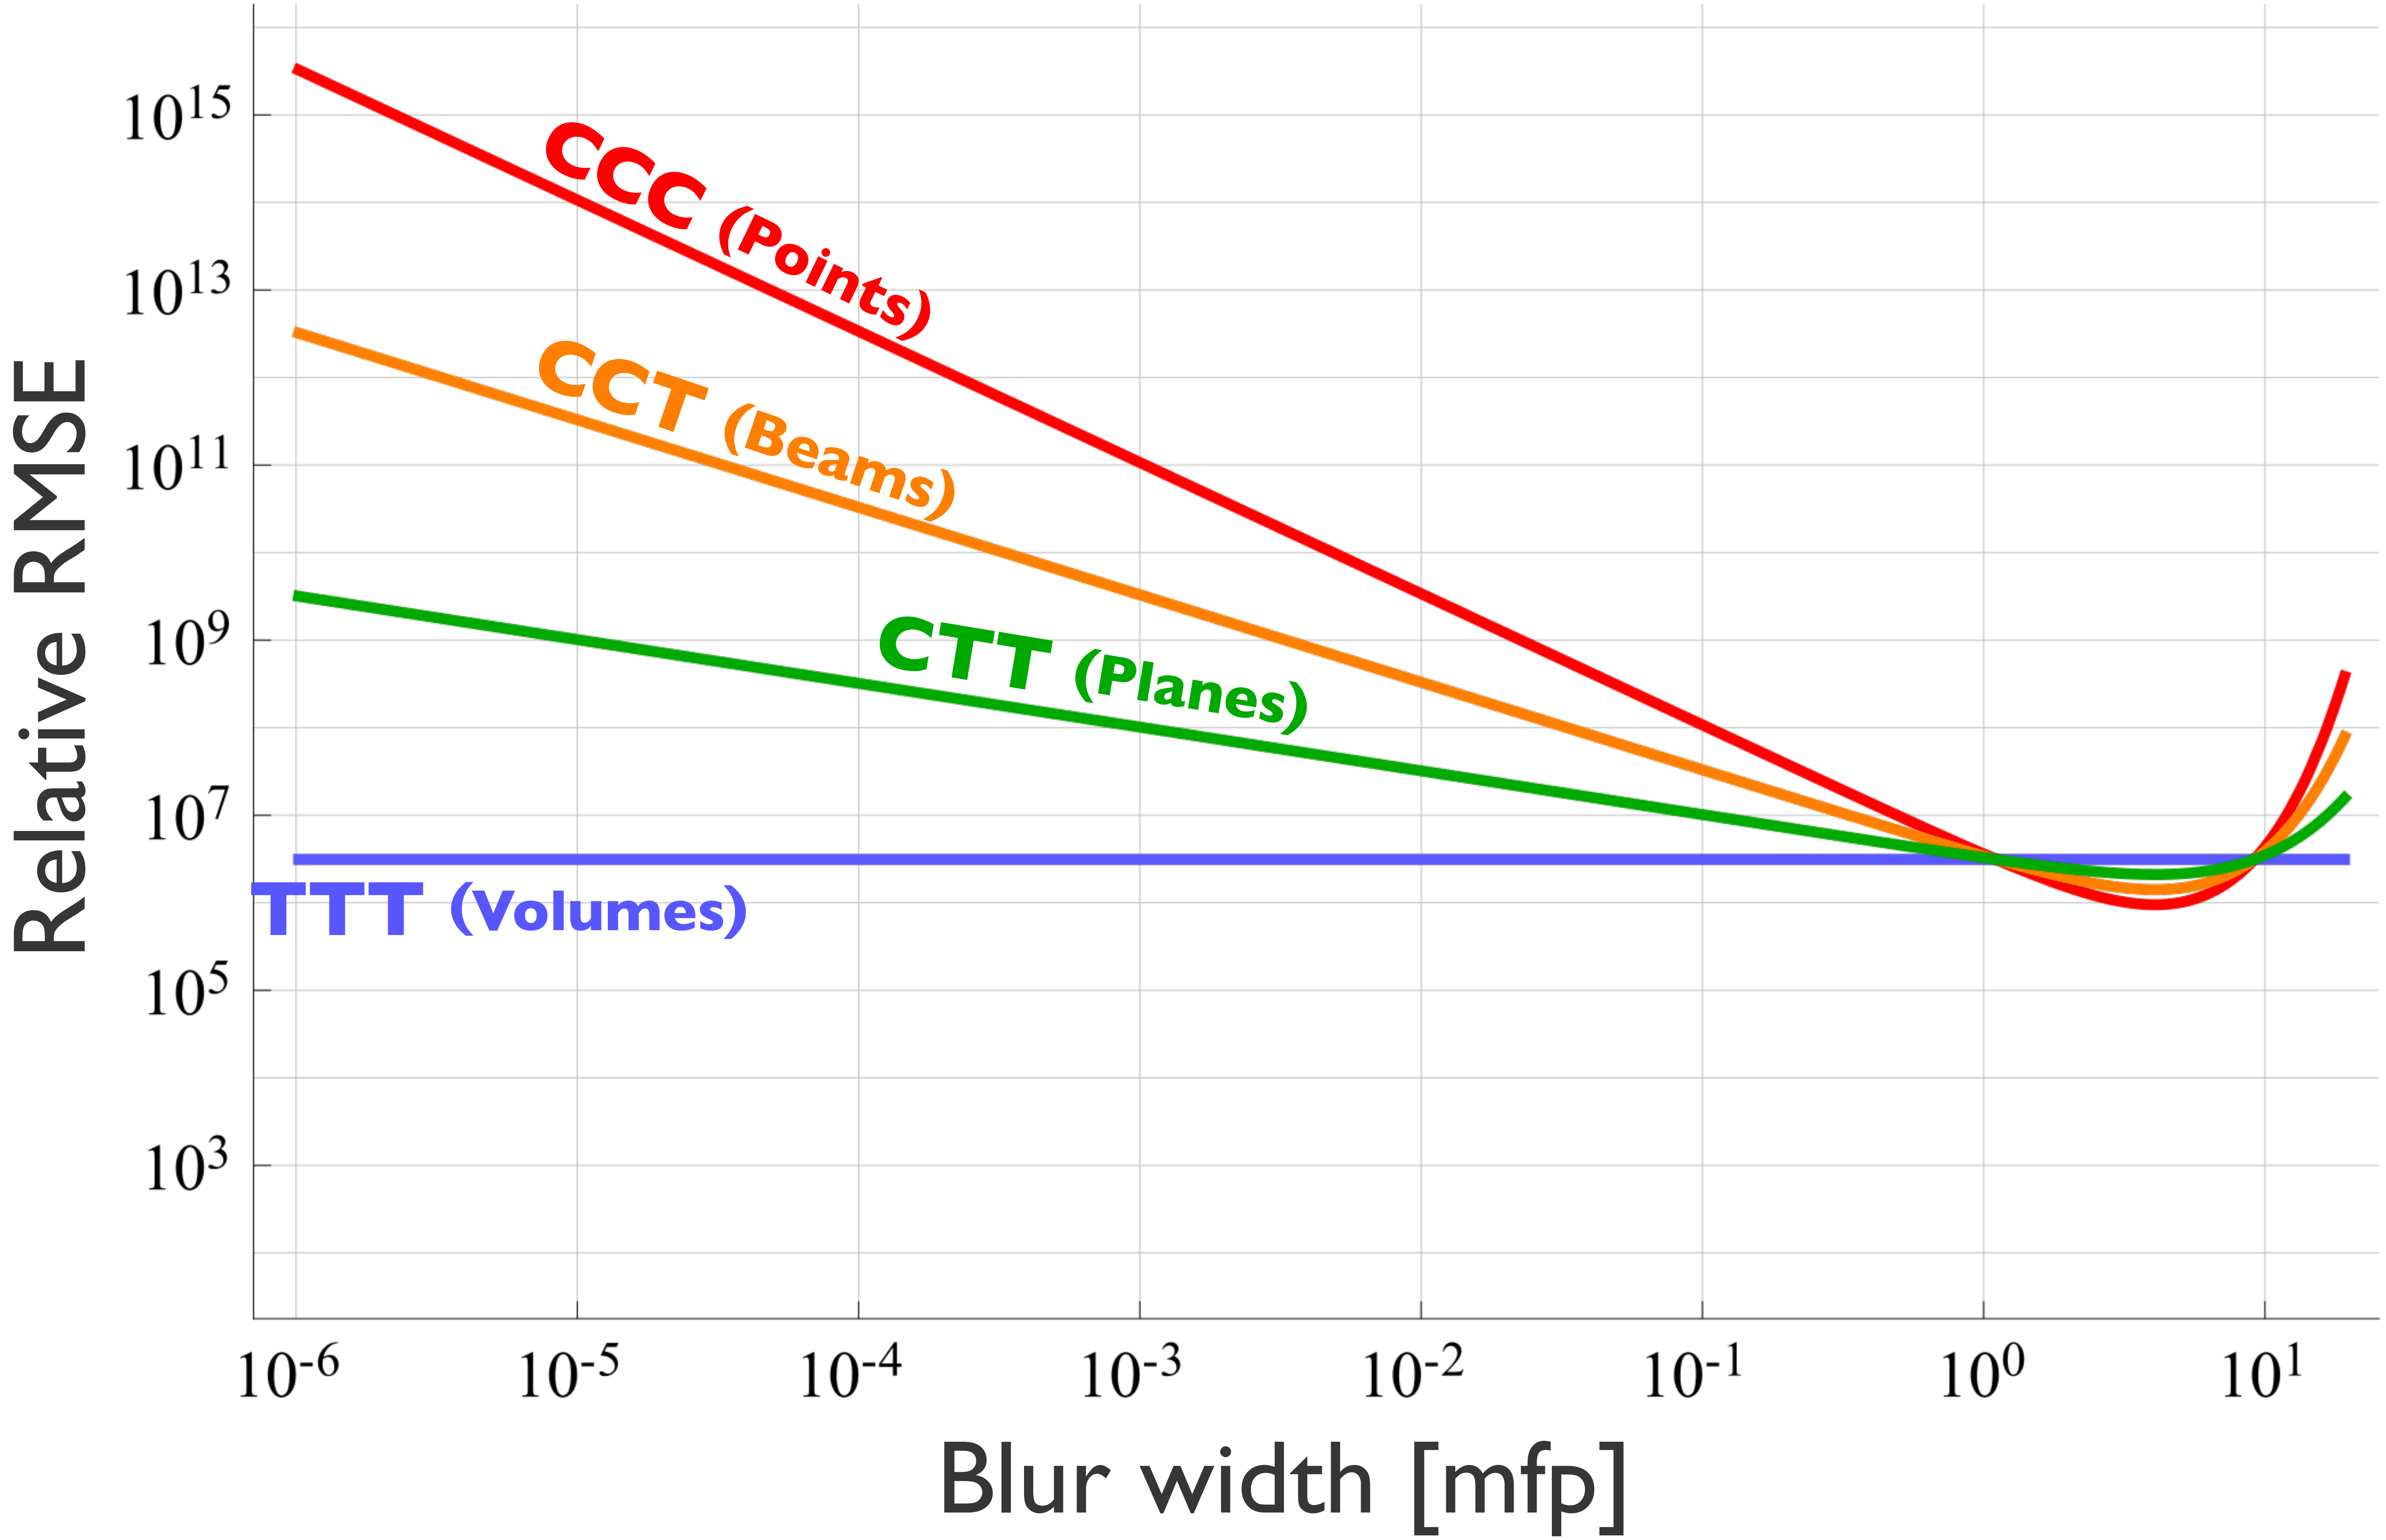

Relative RMSE
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
CCC (Points)
CCT (Beams)
CTT (Planes)
TTT (Volumes)
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
Blur width [mfp]

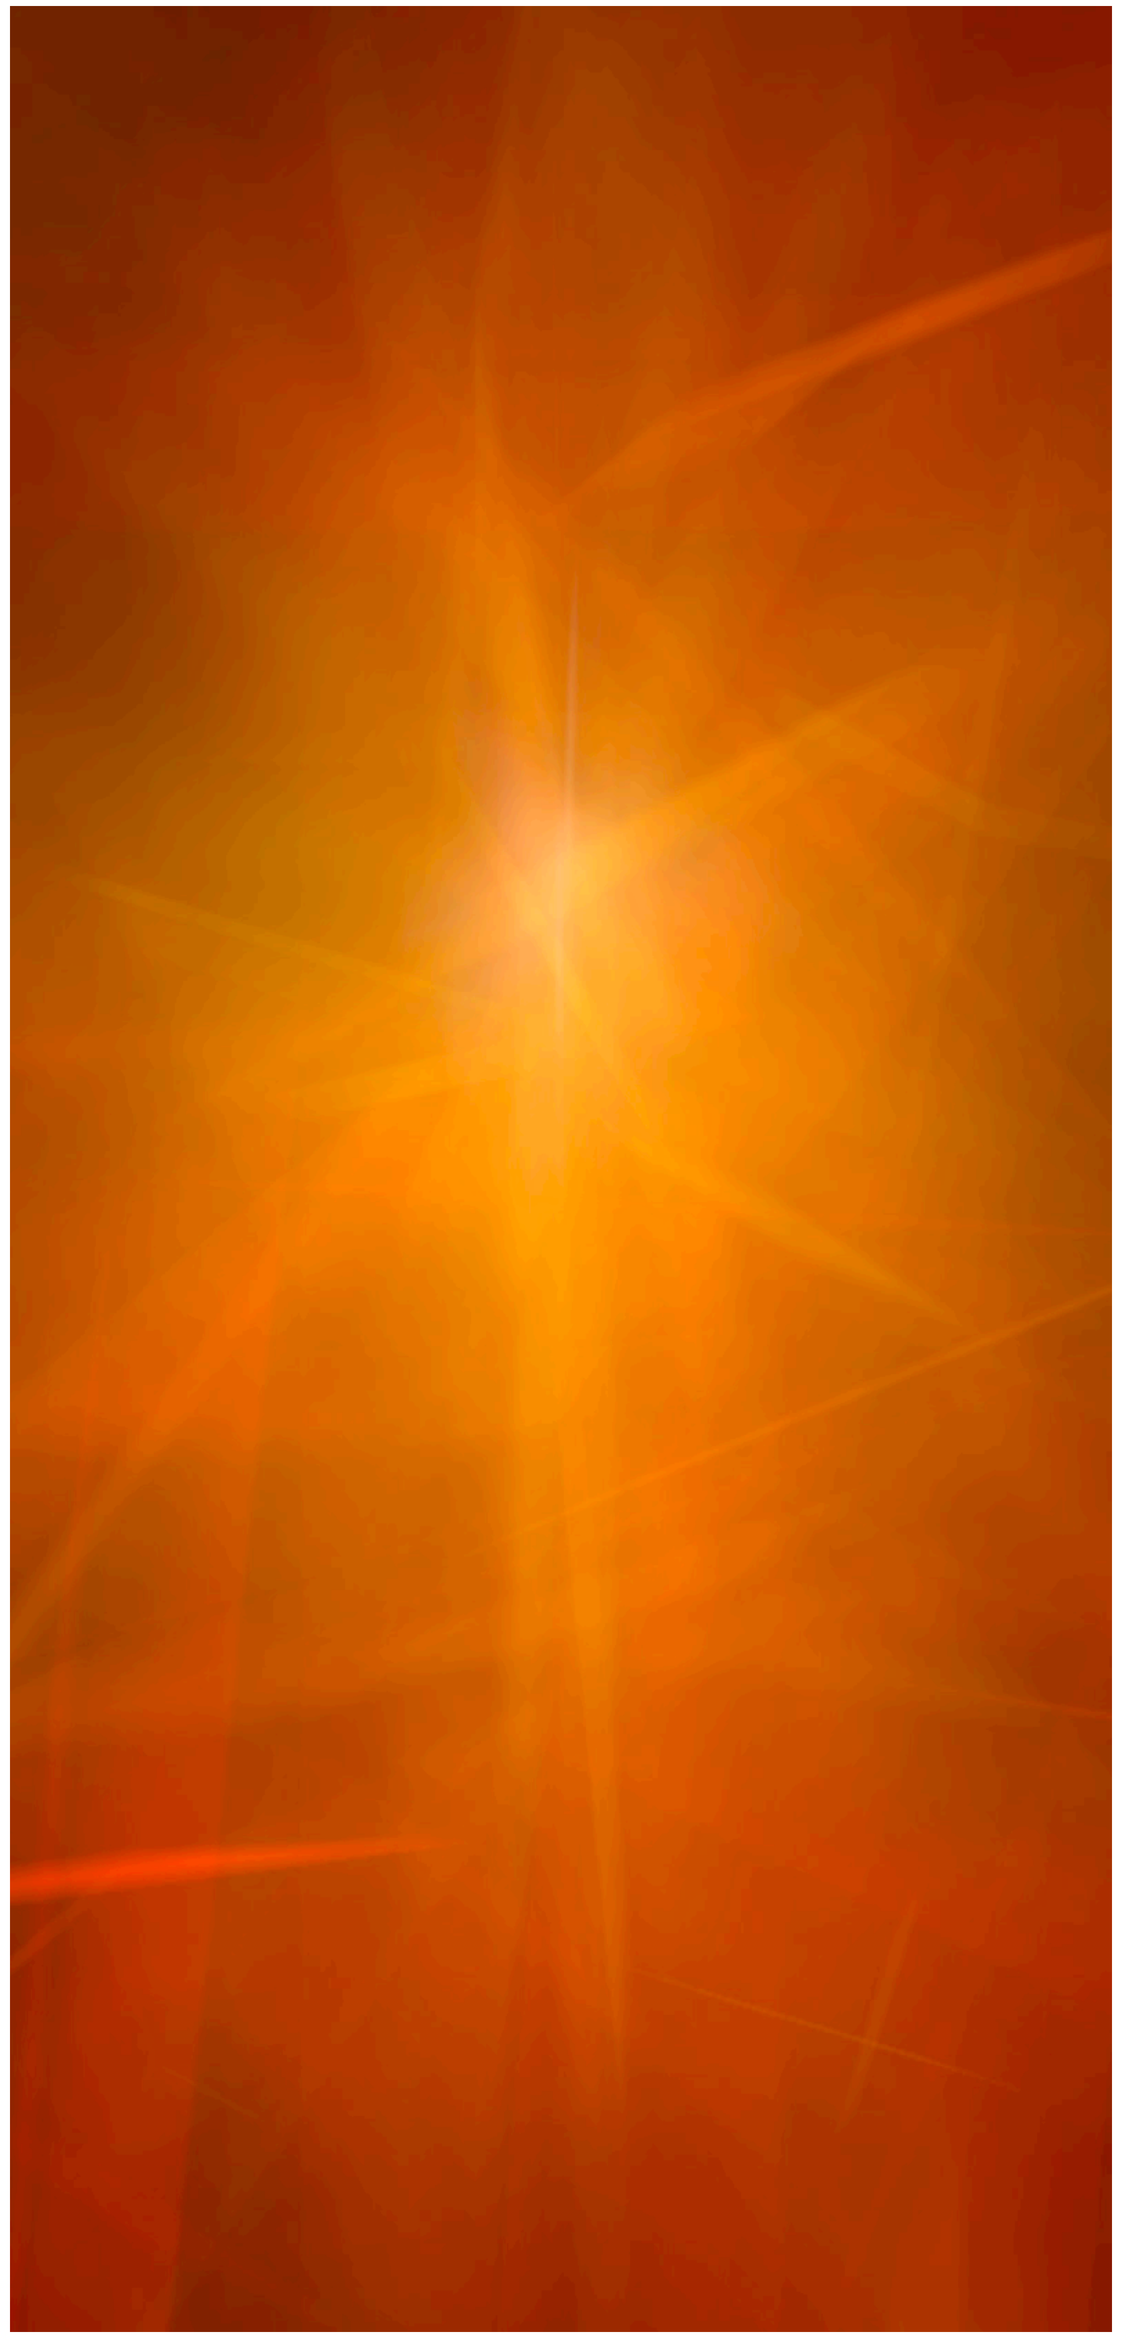

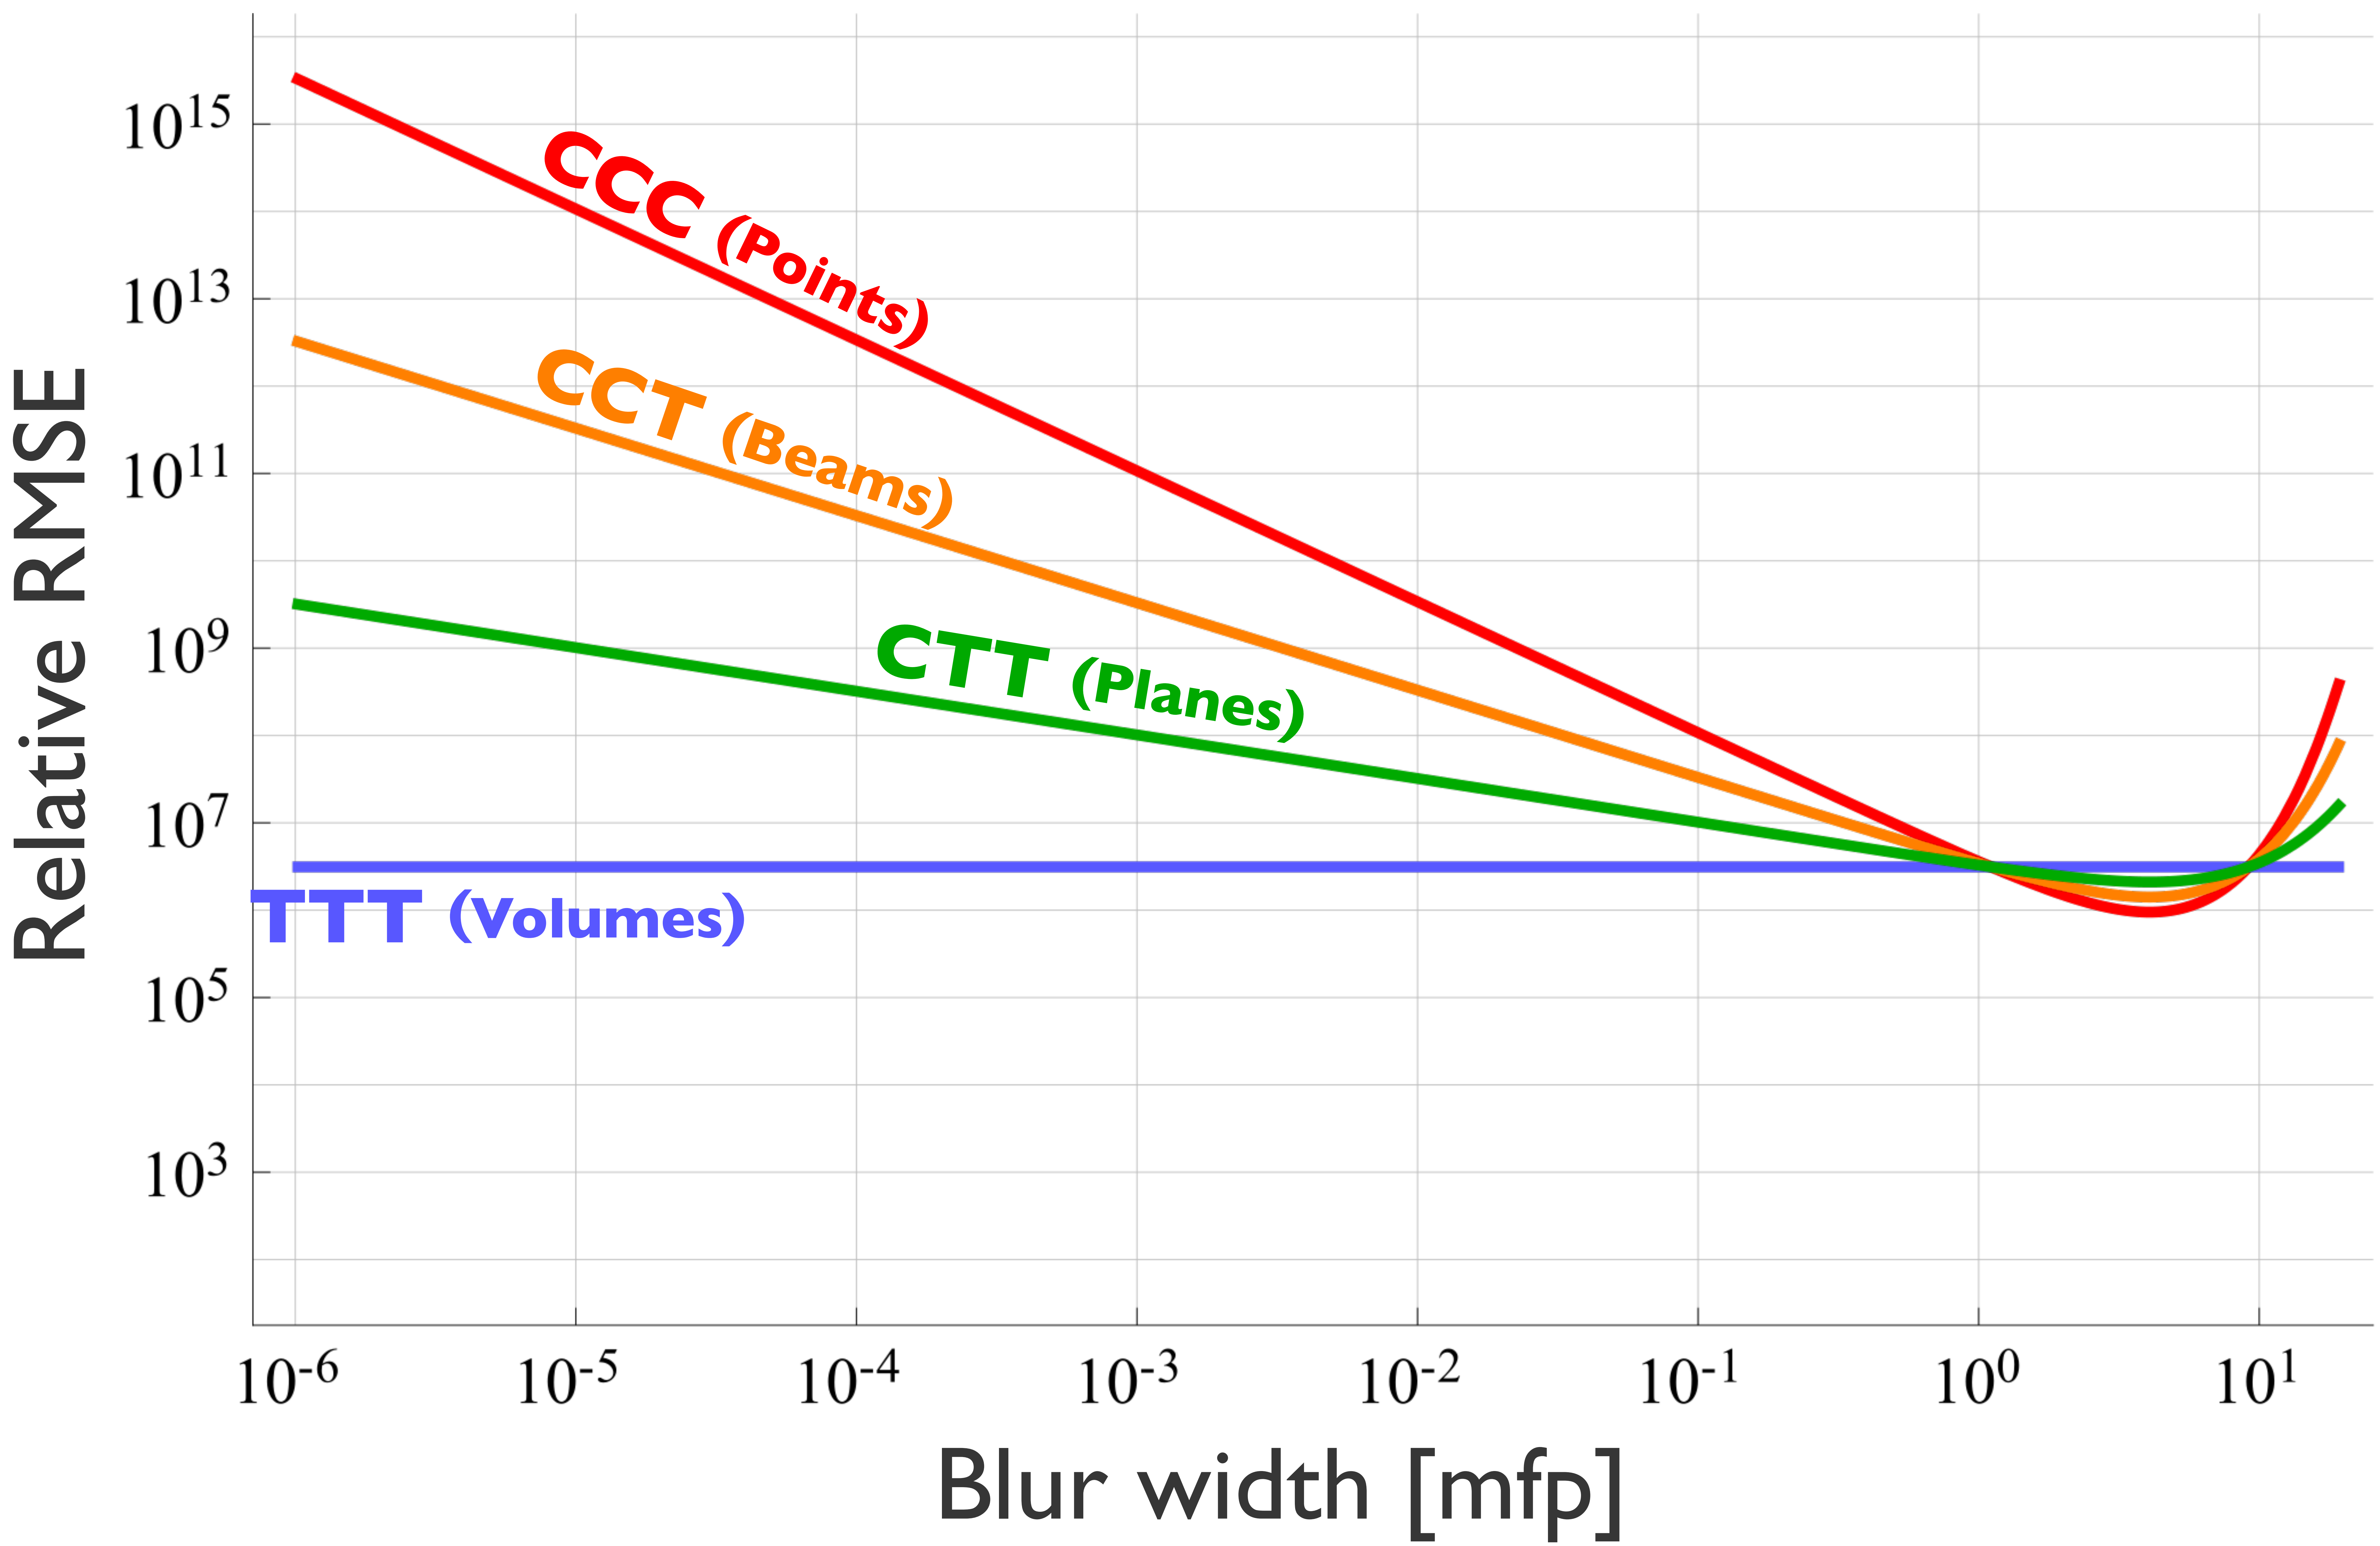

Replacing **C** improves error *asymptotically*

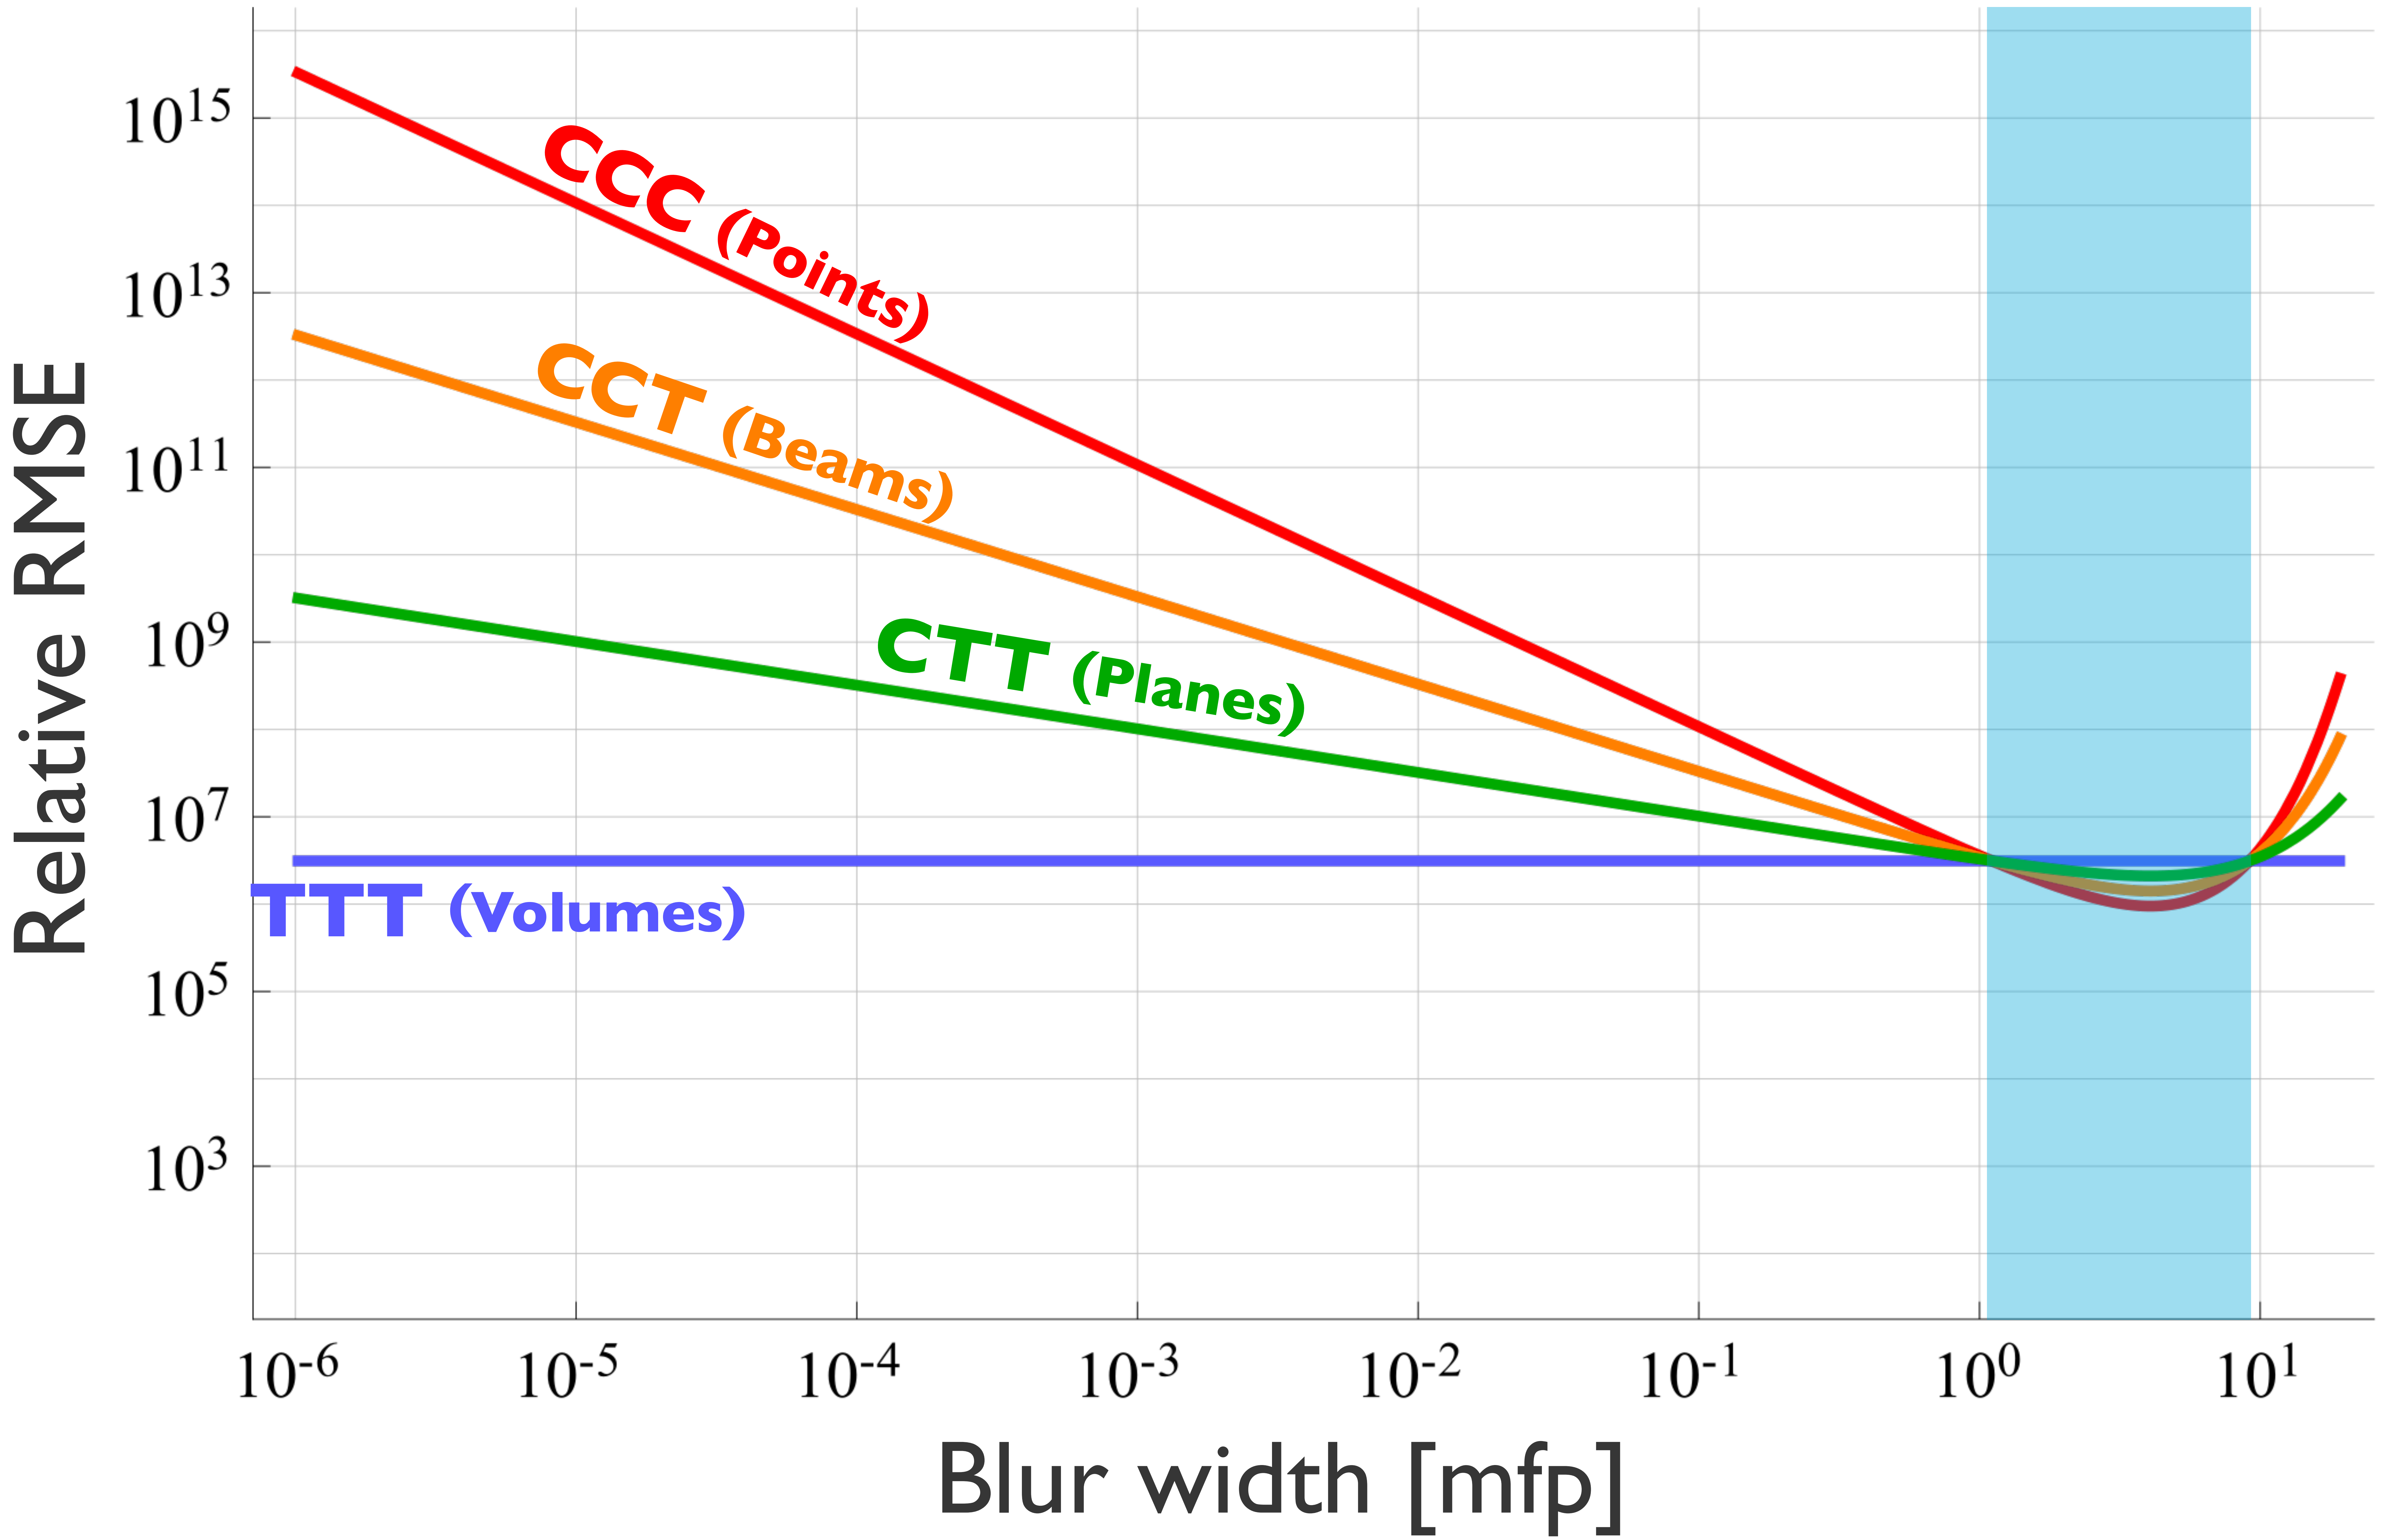

Replacing **C** with **T** almost always better

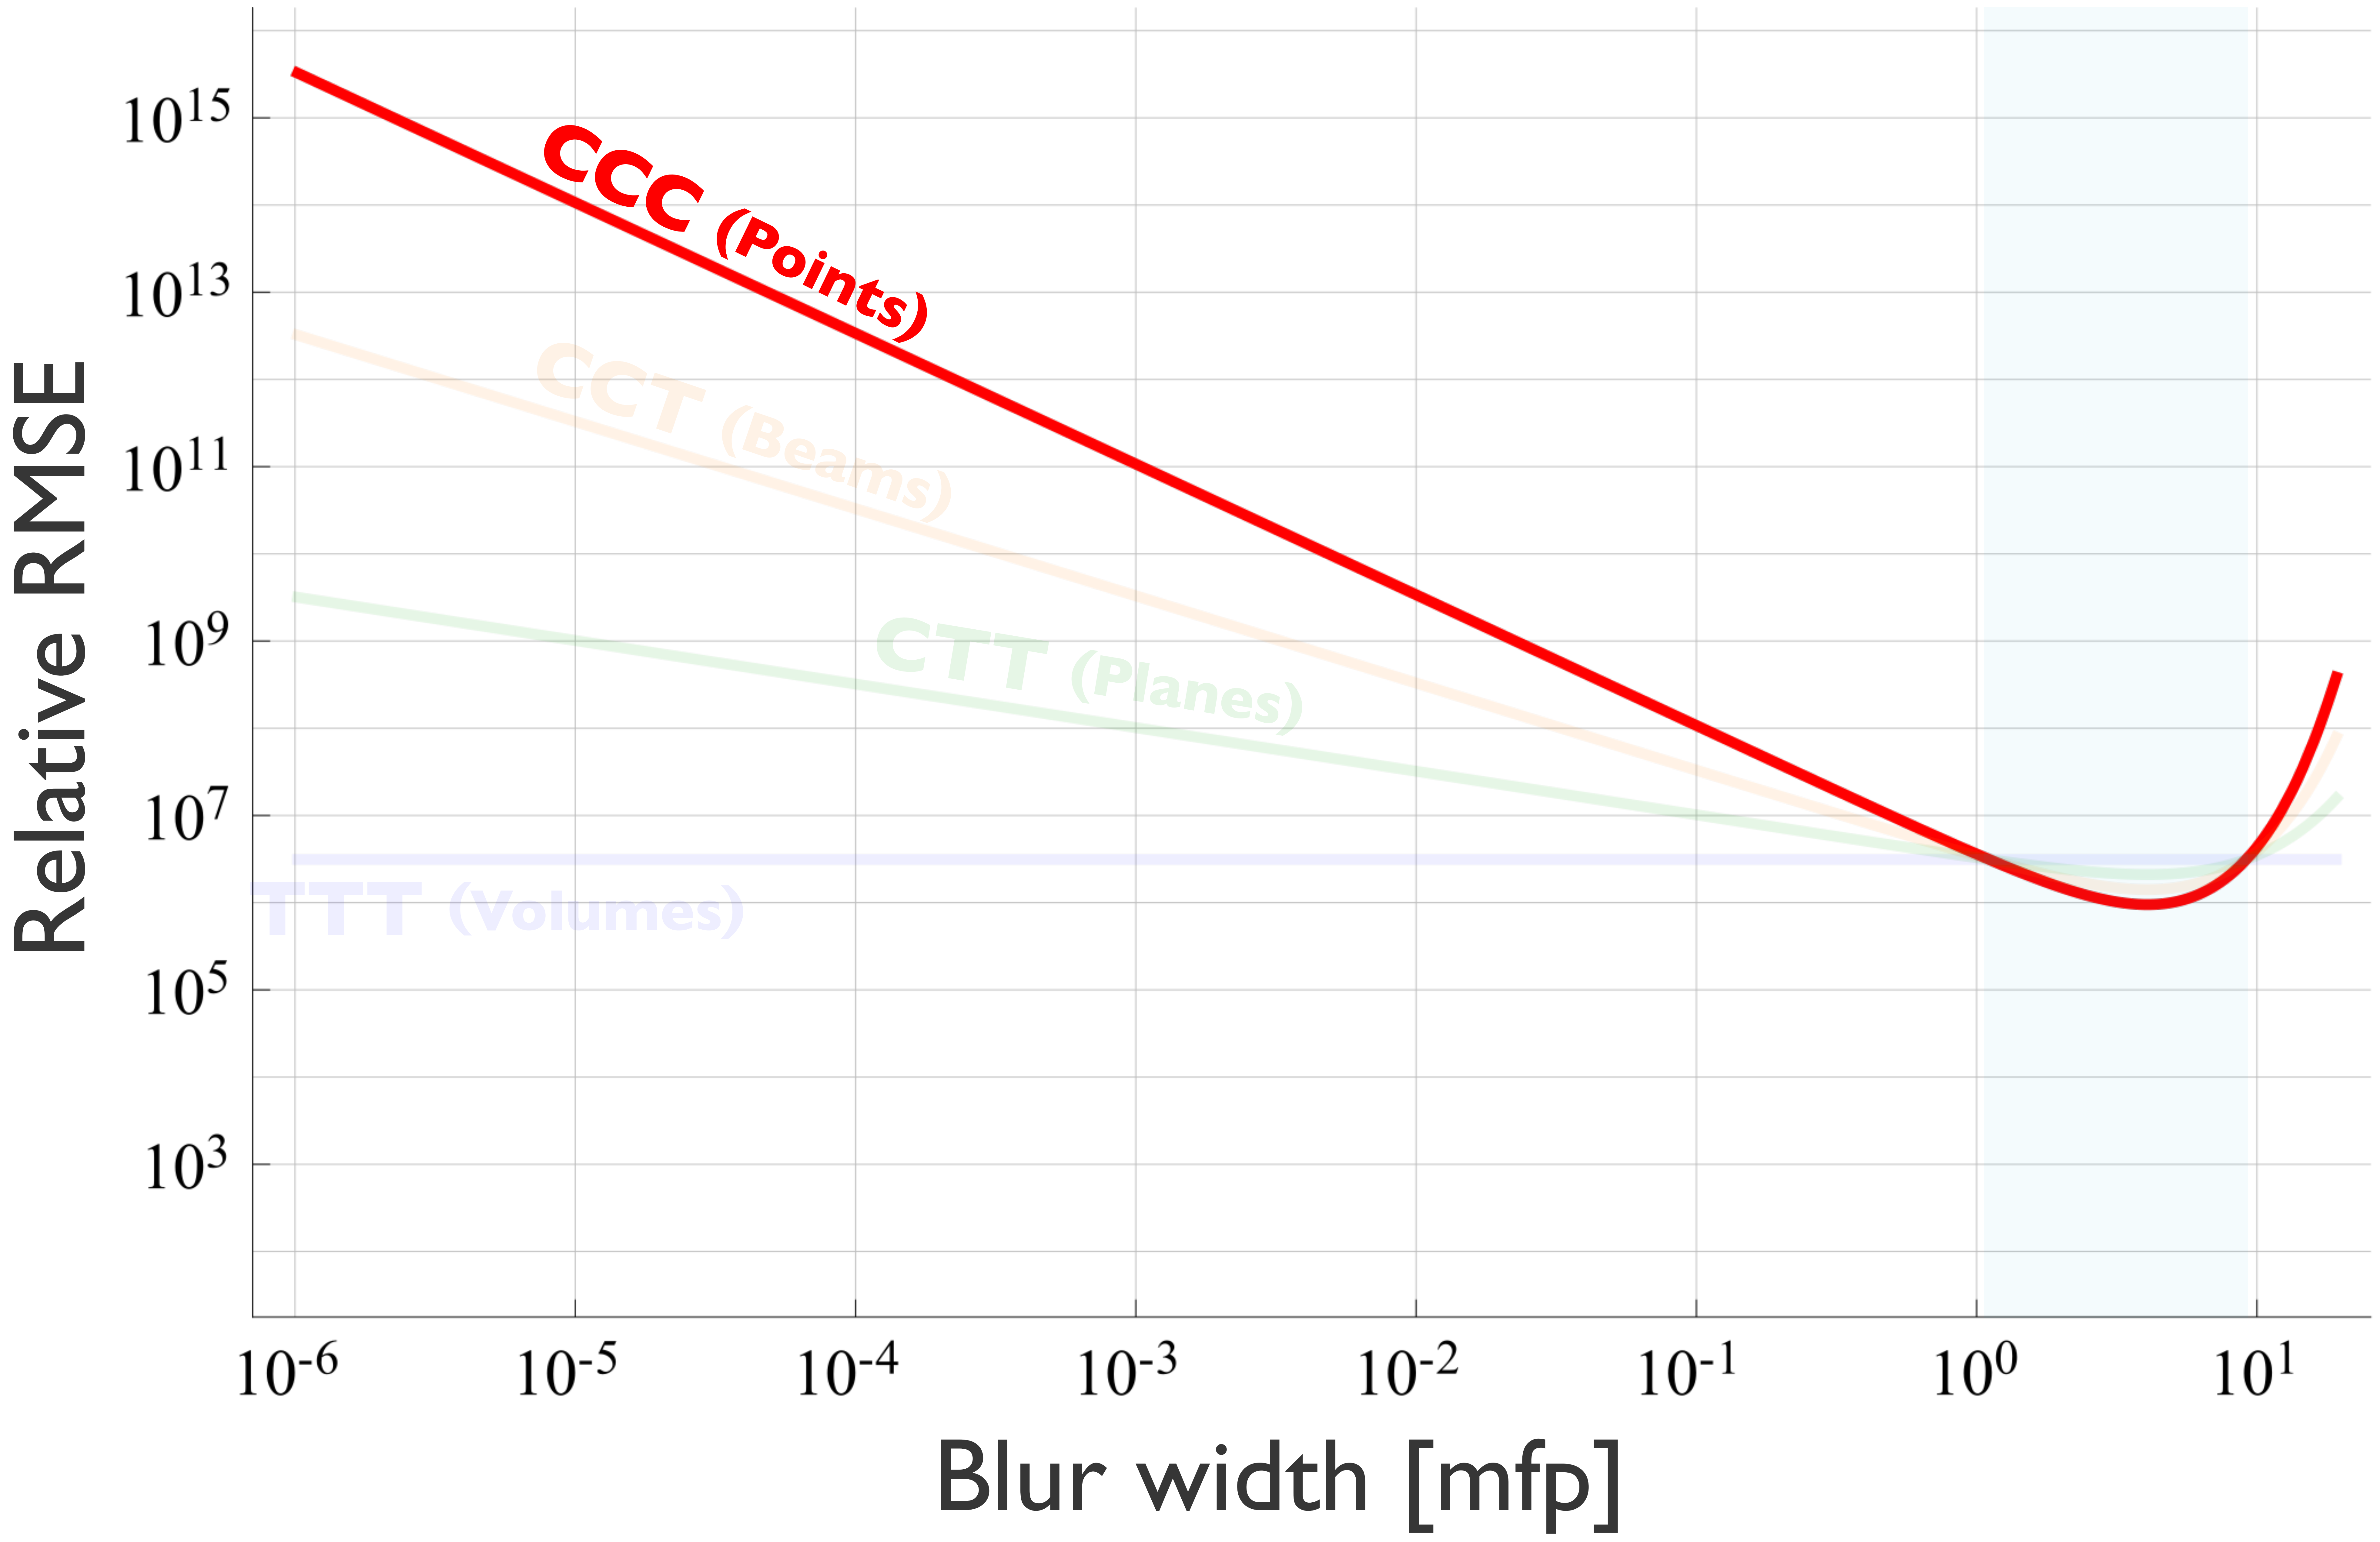
Relative RMSE
Blur width [mfp]
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
CCC (Points)
CCT (Beams)
CTT (Planes)
TTT (Volumes)

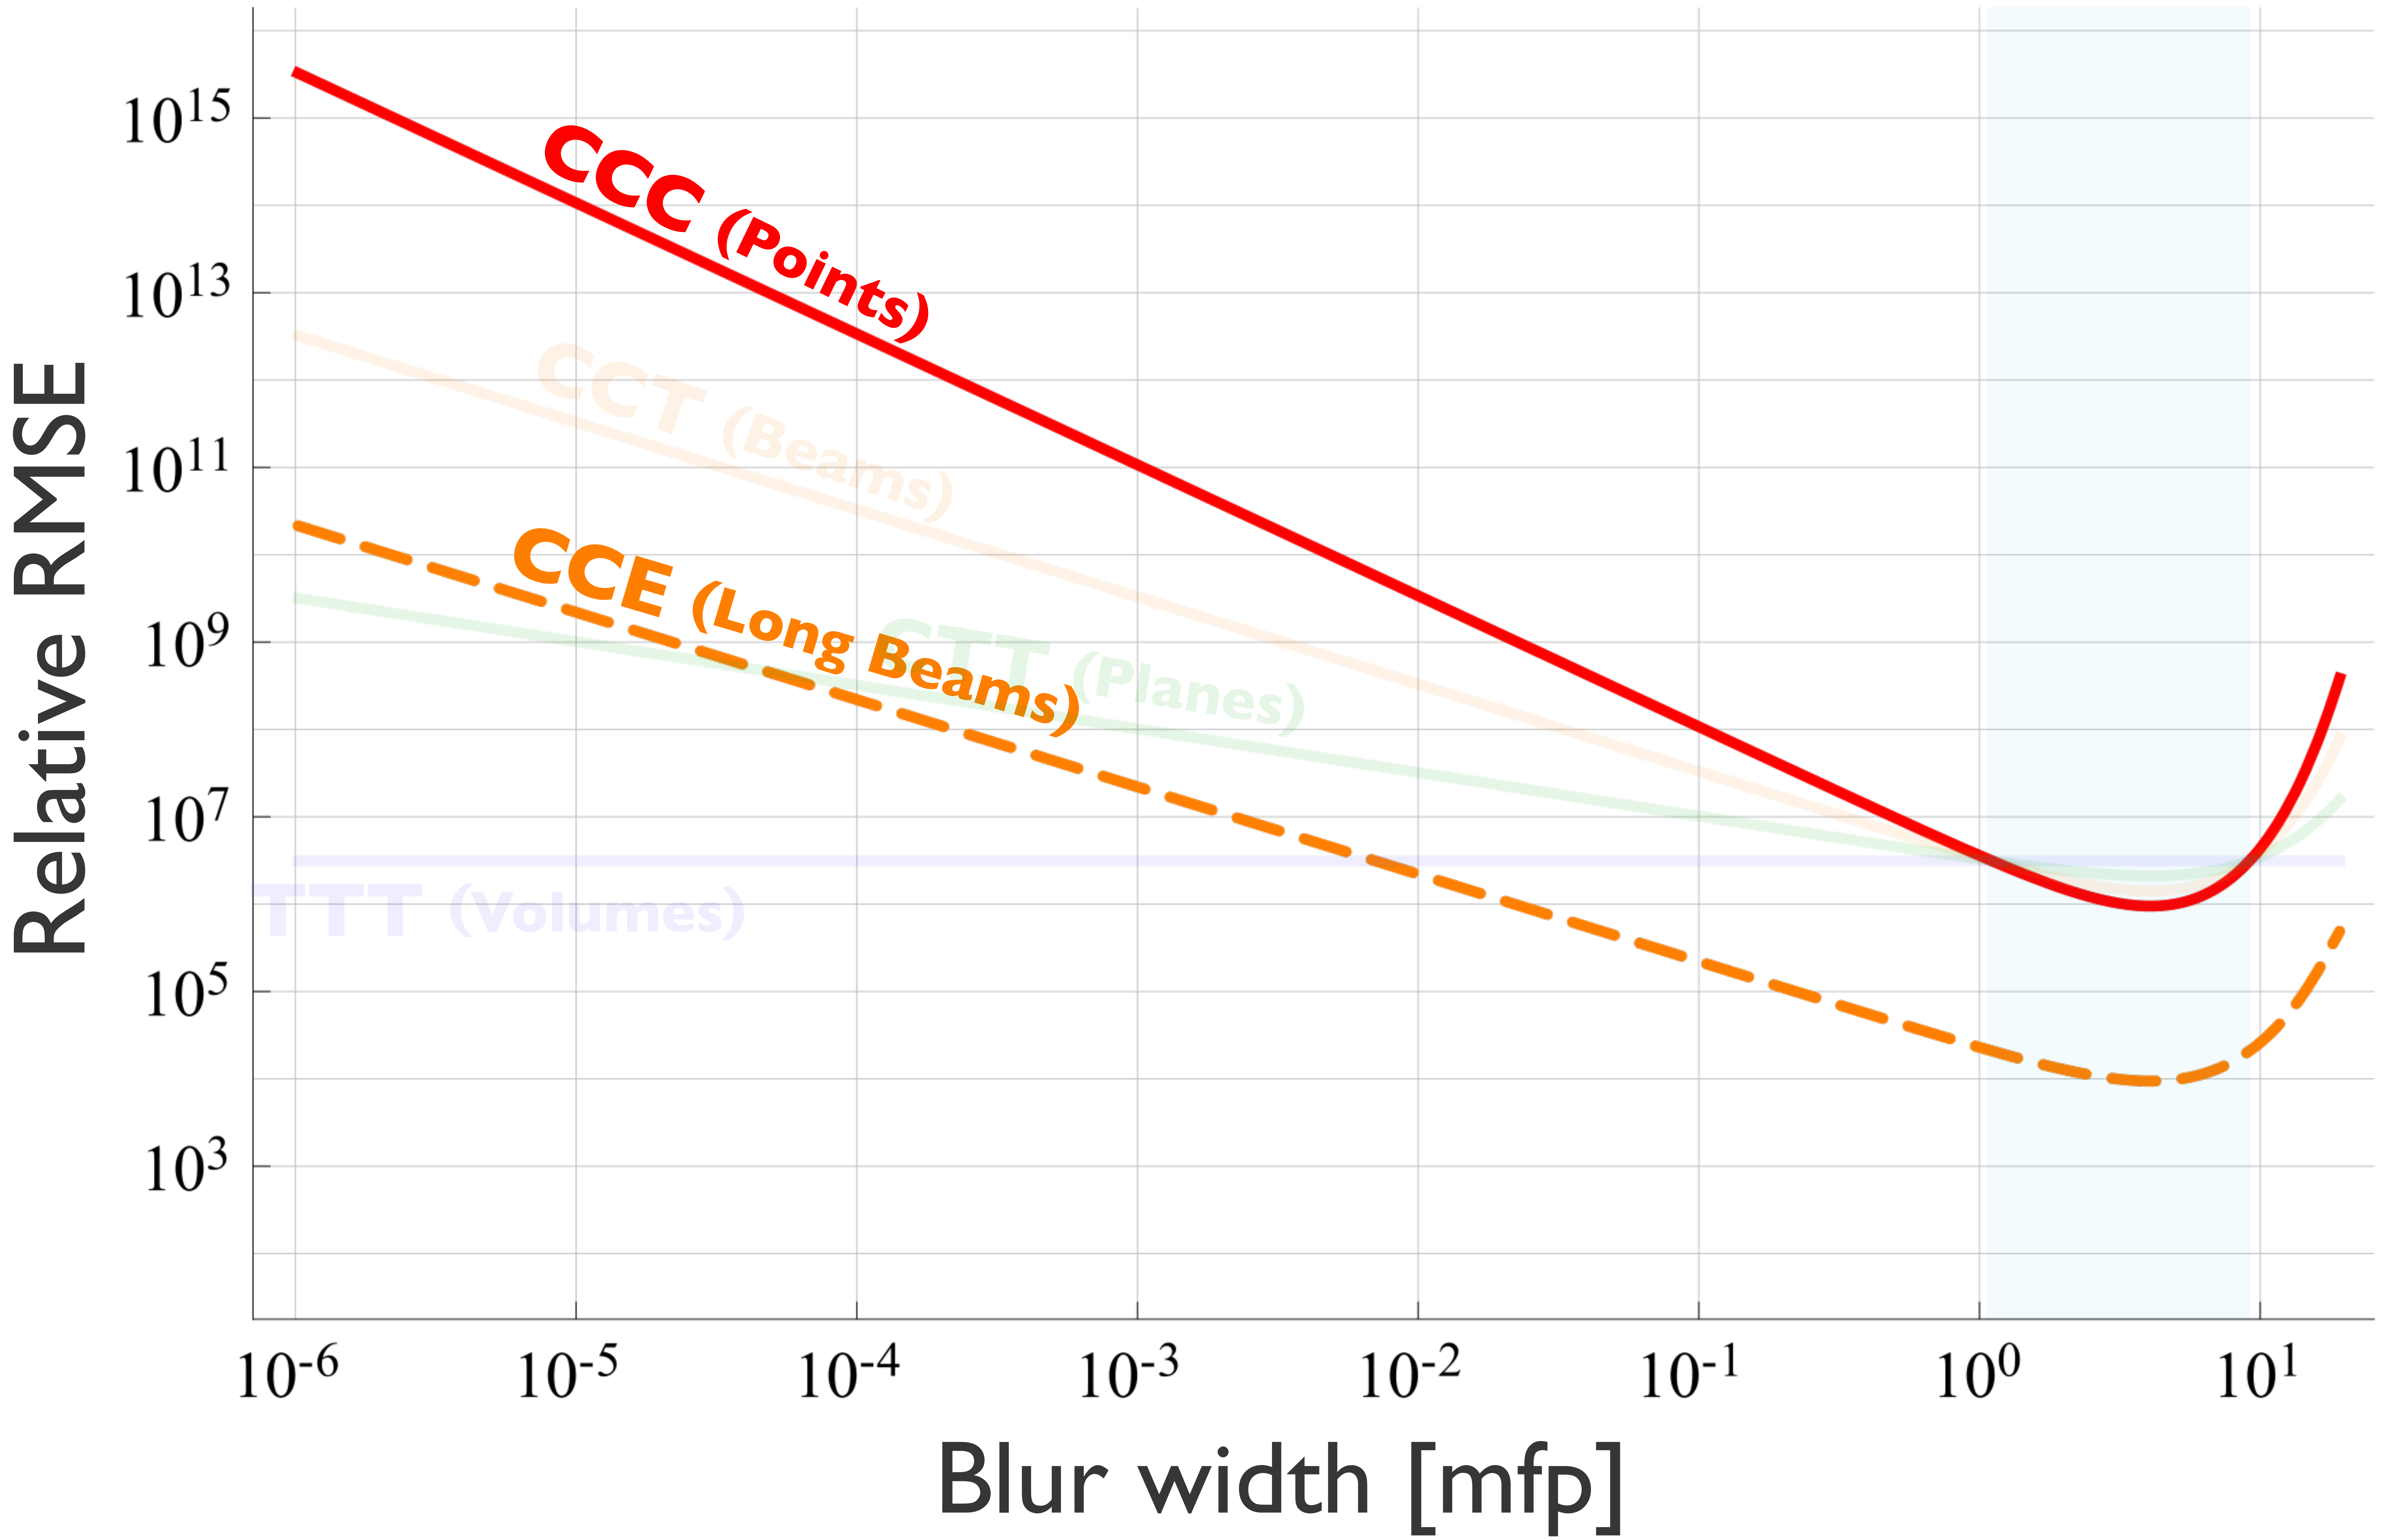
CCC (Points)
CCT (Beams)
CCE (Long Beams)
CTT (Planes)
TTT (Volumes)
Relative RMSE
Blur width [mfp]
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}

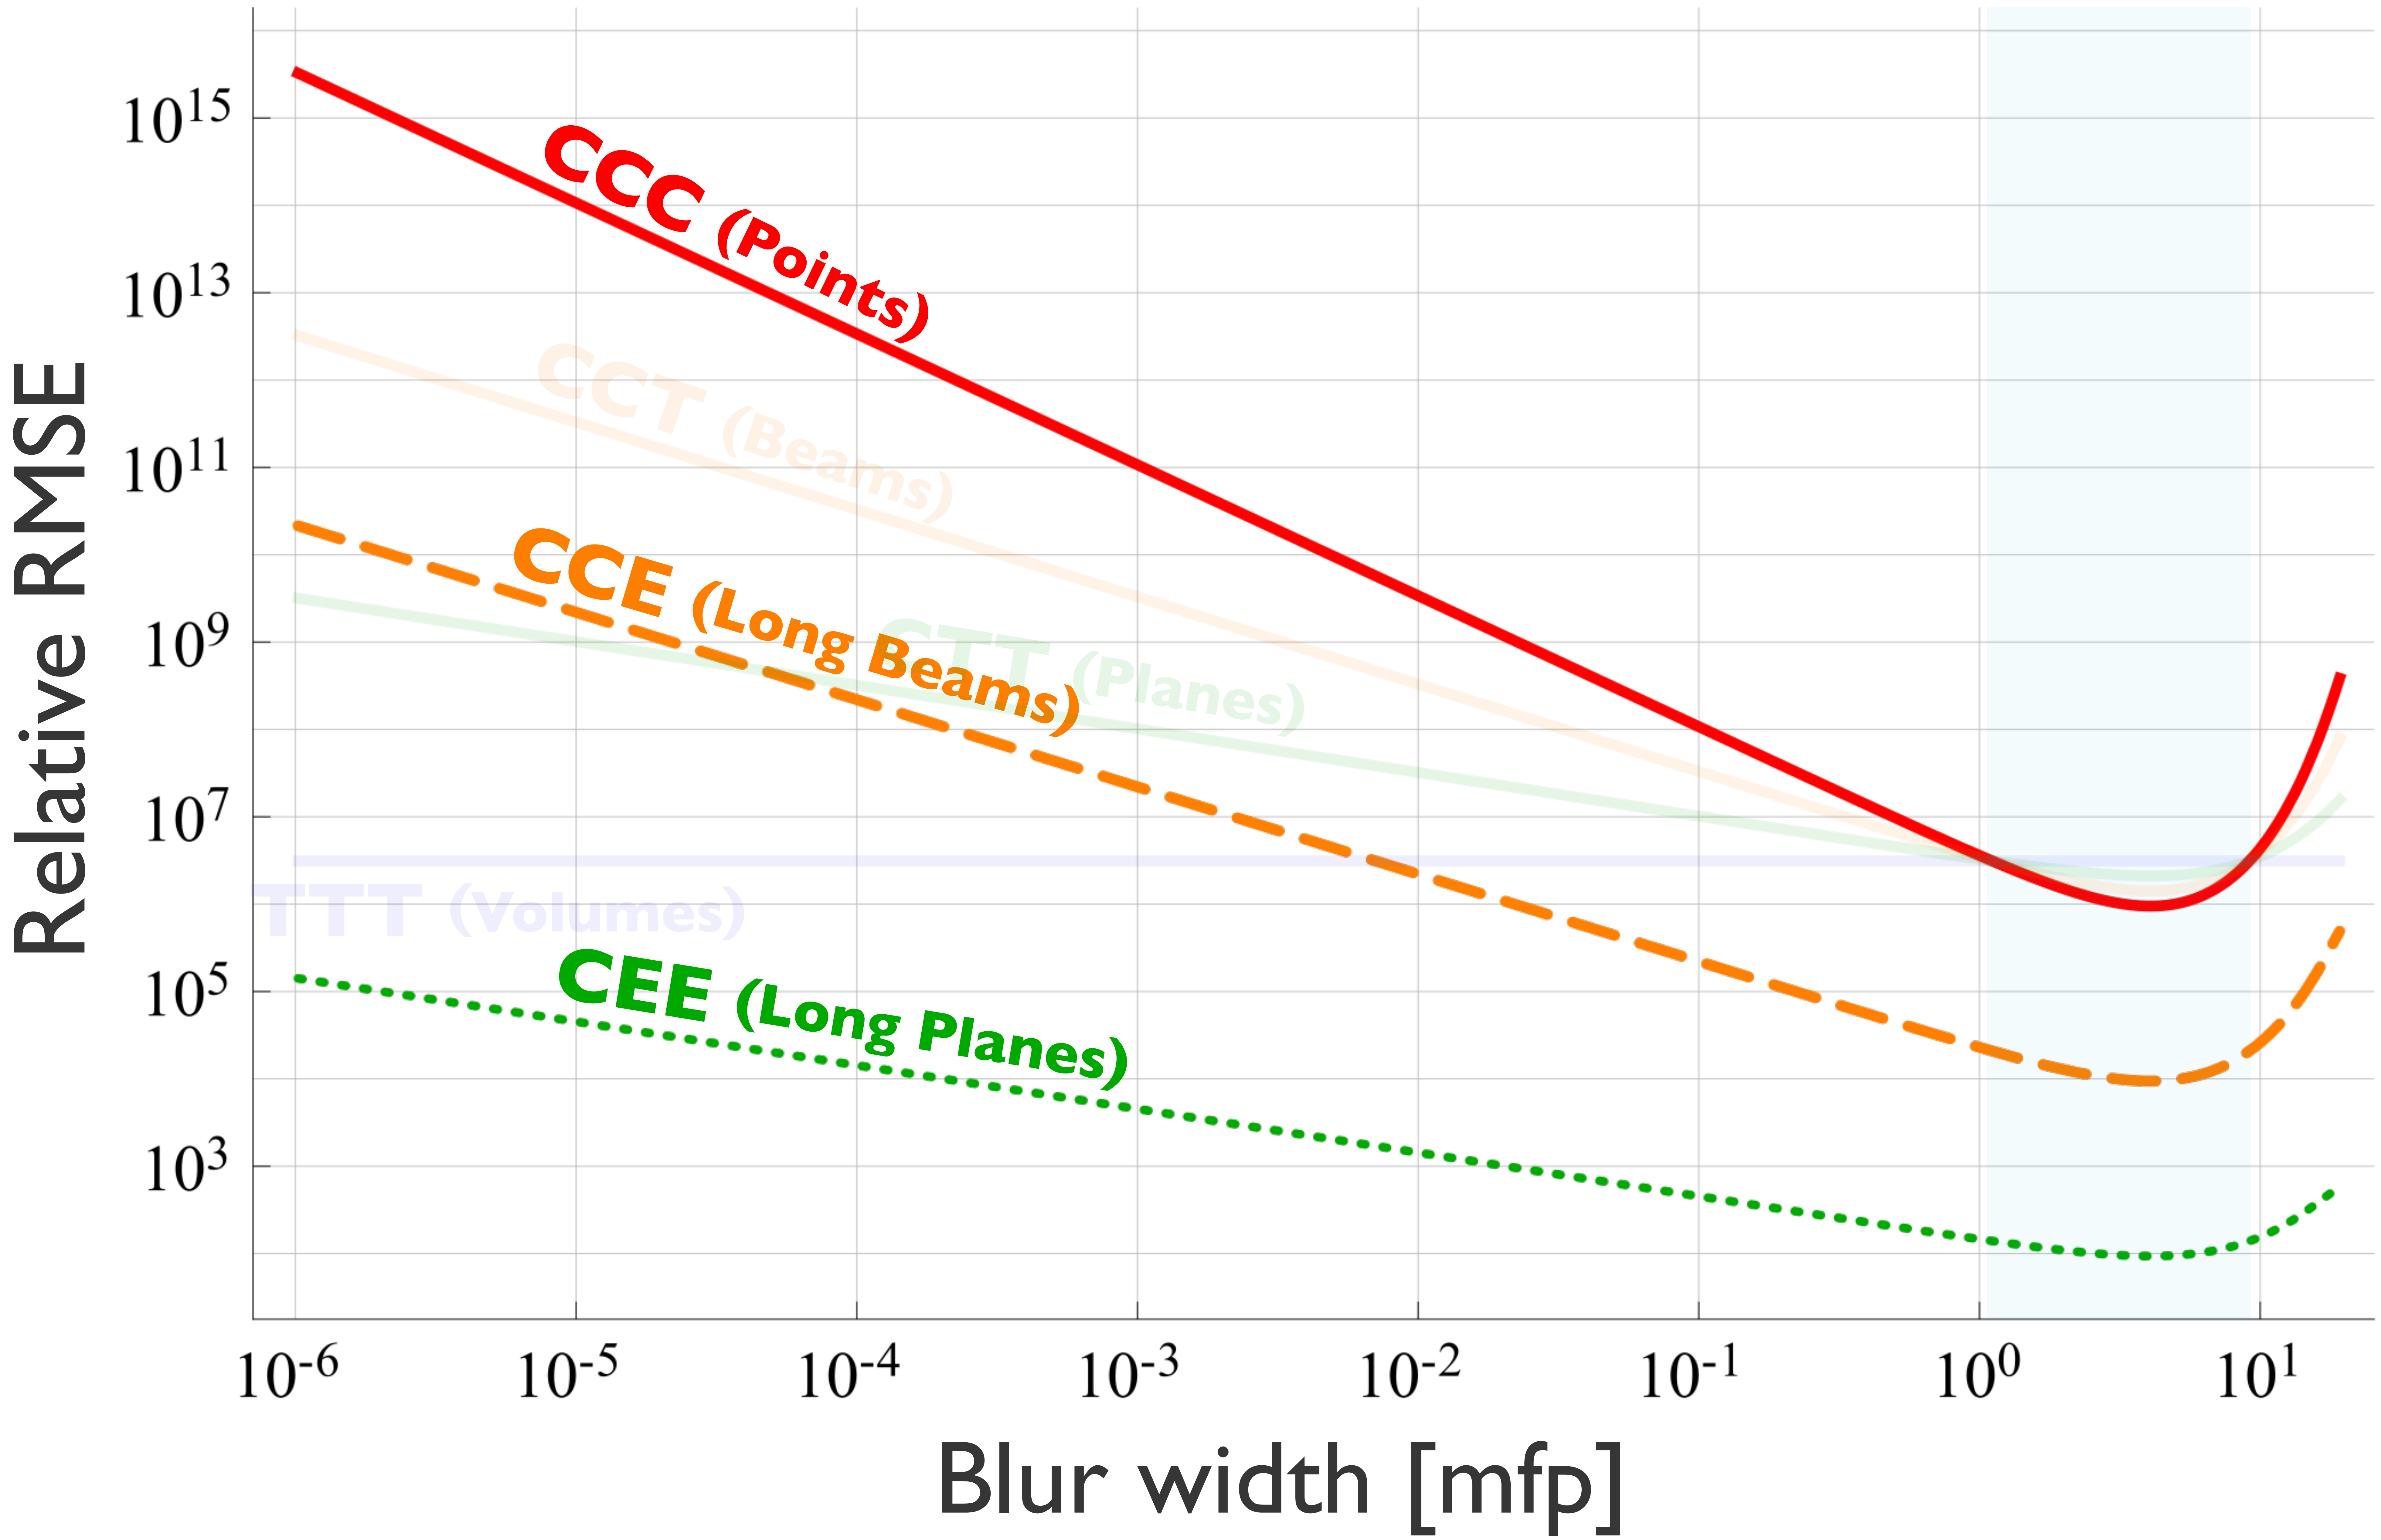

CCC (Points)
CCT (Beams)
CCE (Long Beams)
CTT (Planes)
TTT (Volumes)
CEE (Long Planes)
10^{15}
10^{13}
10^{11}
10^{9}
10^{7}
10^{5}
10^{3}
10^{-6}
10^{-5}
10^{-4}
10^{-3}
10^{-2}
10^{-1}
10^{0}
10^{1}
Relative RMSE
Blur width [mfp]

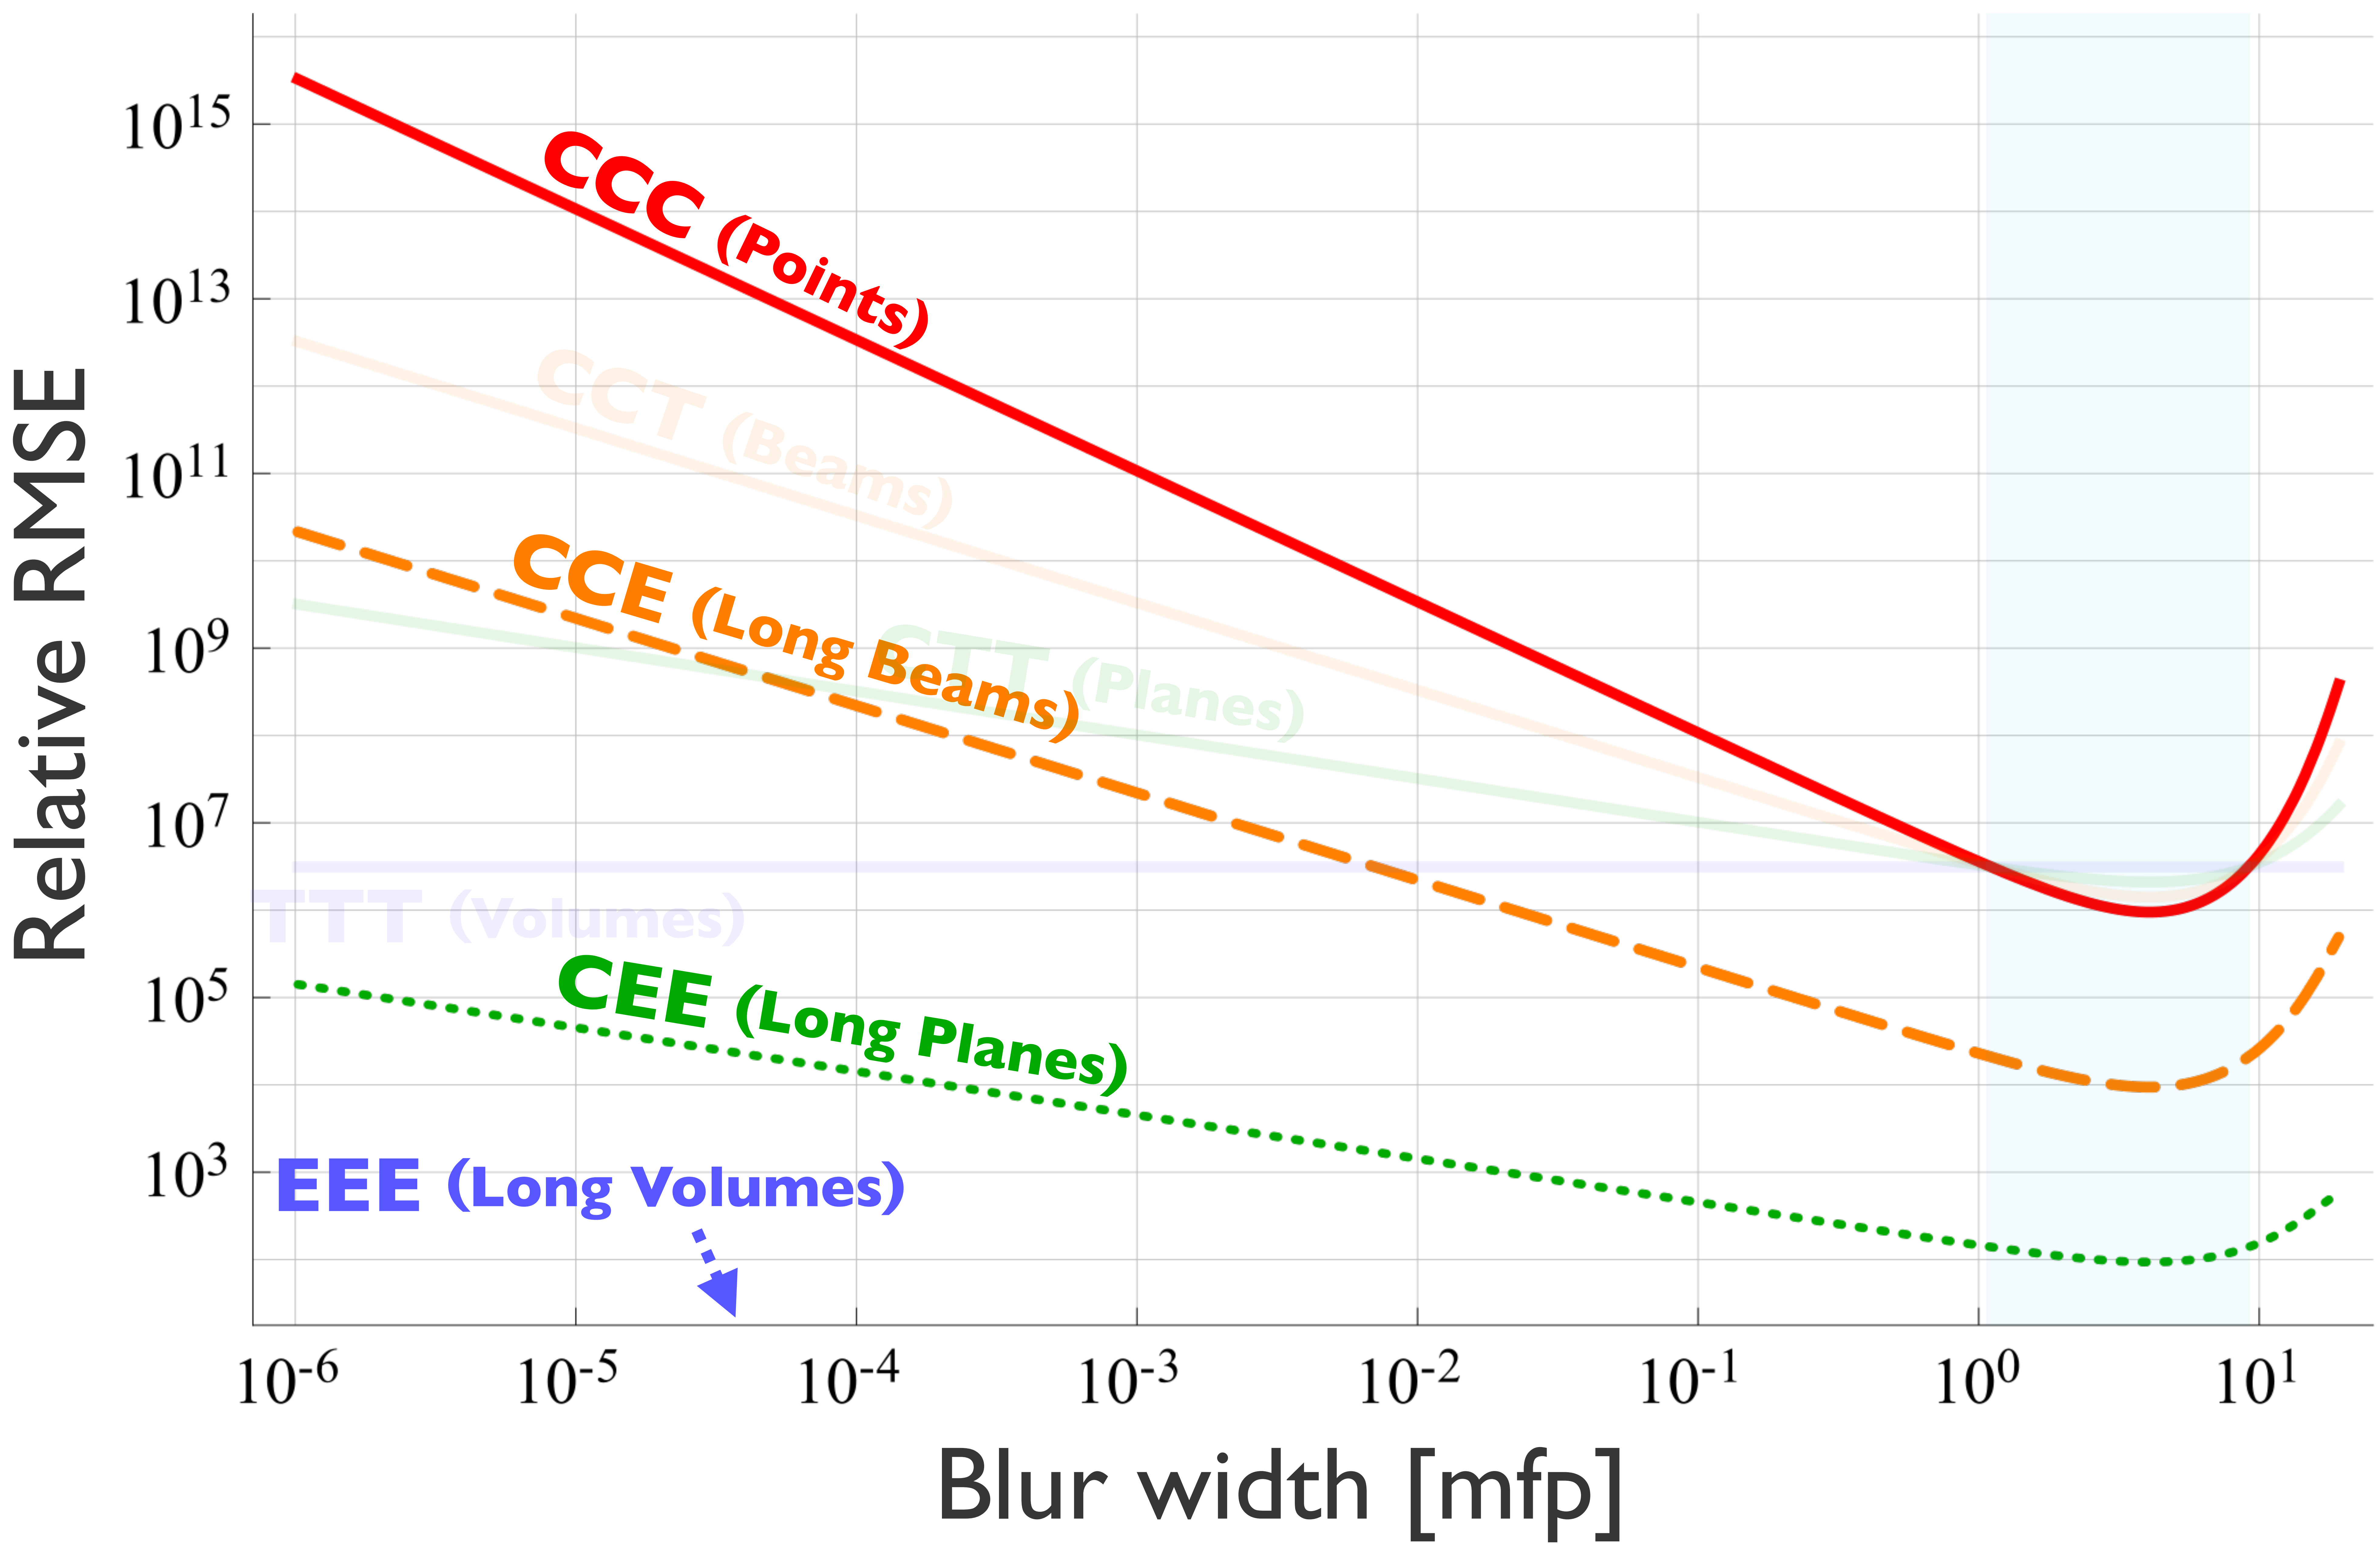
Relative RMSE
10^15
10^13
10^11
10^9
10^7
10^5
10^3
CCC (Points)
CCT (Beams)
CCE (Long Beams)
CTT (Planes)
TTT (Volumes)
CEE (Long Planes)
EEE (Long Volumes)
10^-6
10^-5
10^-4
10^-3
10^-2
10^-1
10^0
10^1
Blur width [mfp]

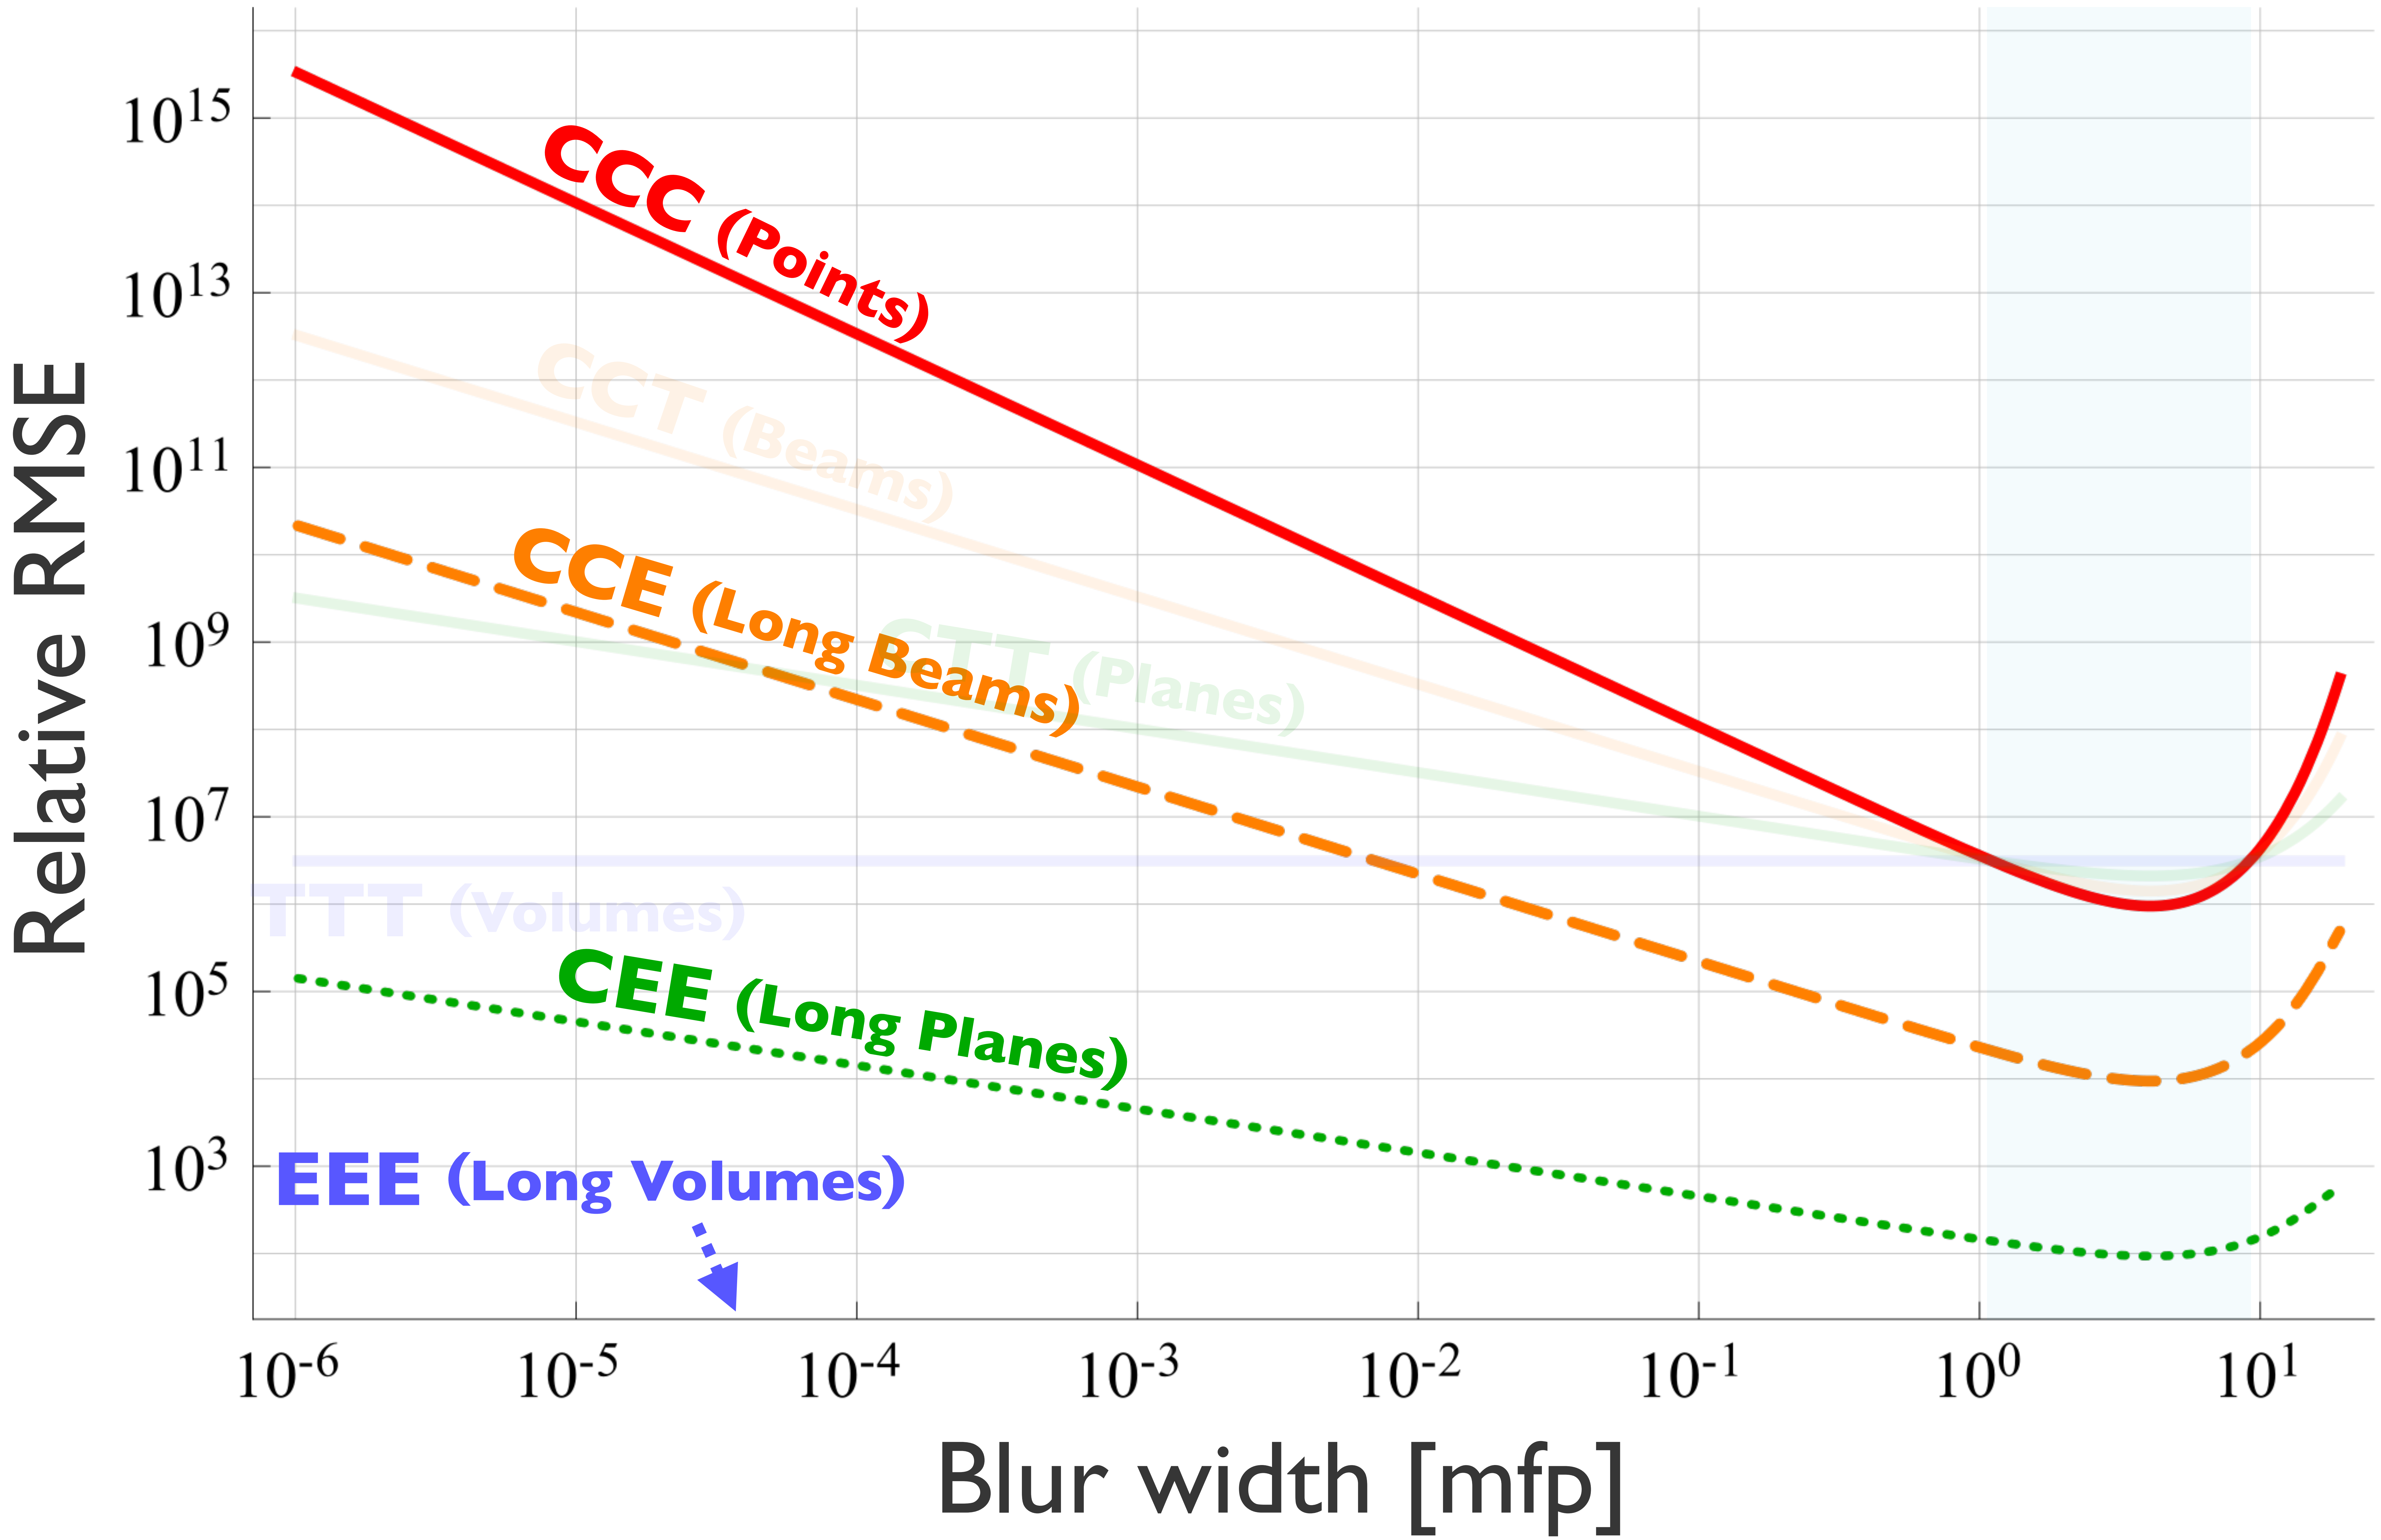

Replacing **C** with **E** always better

Results

Results

- Two implementations of our method

OpenGL Implementation

- CPU: Trace photon paths
- GPU: Rasterize photons

OpenGL Implementation

- CPU: Trace photon paths
- GPU: Rasterize photons
- Can do this in the browser!

OpenGL Implementation

- CPU: Trace photon paths
- GPU: Rasterize photons
- Can do this in the browser!

- Scene:

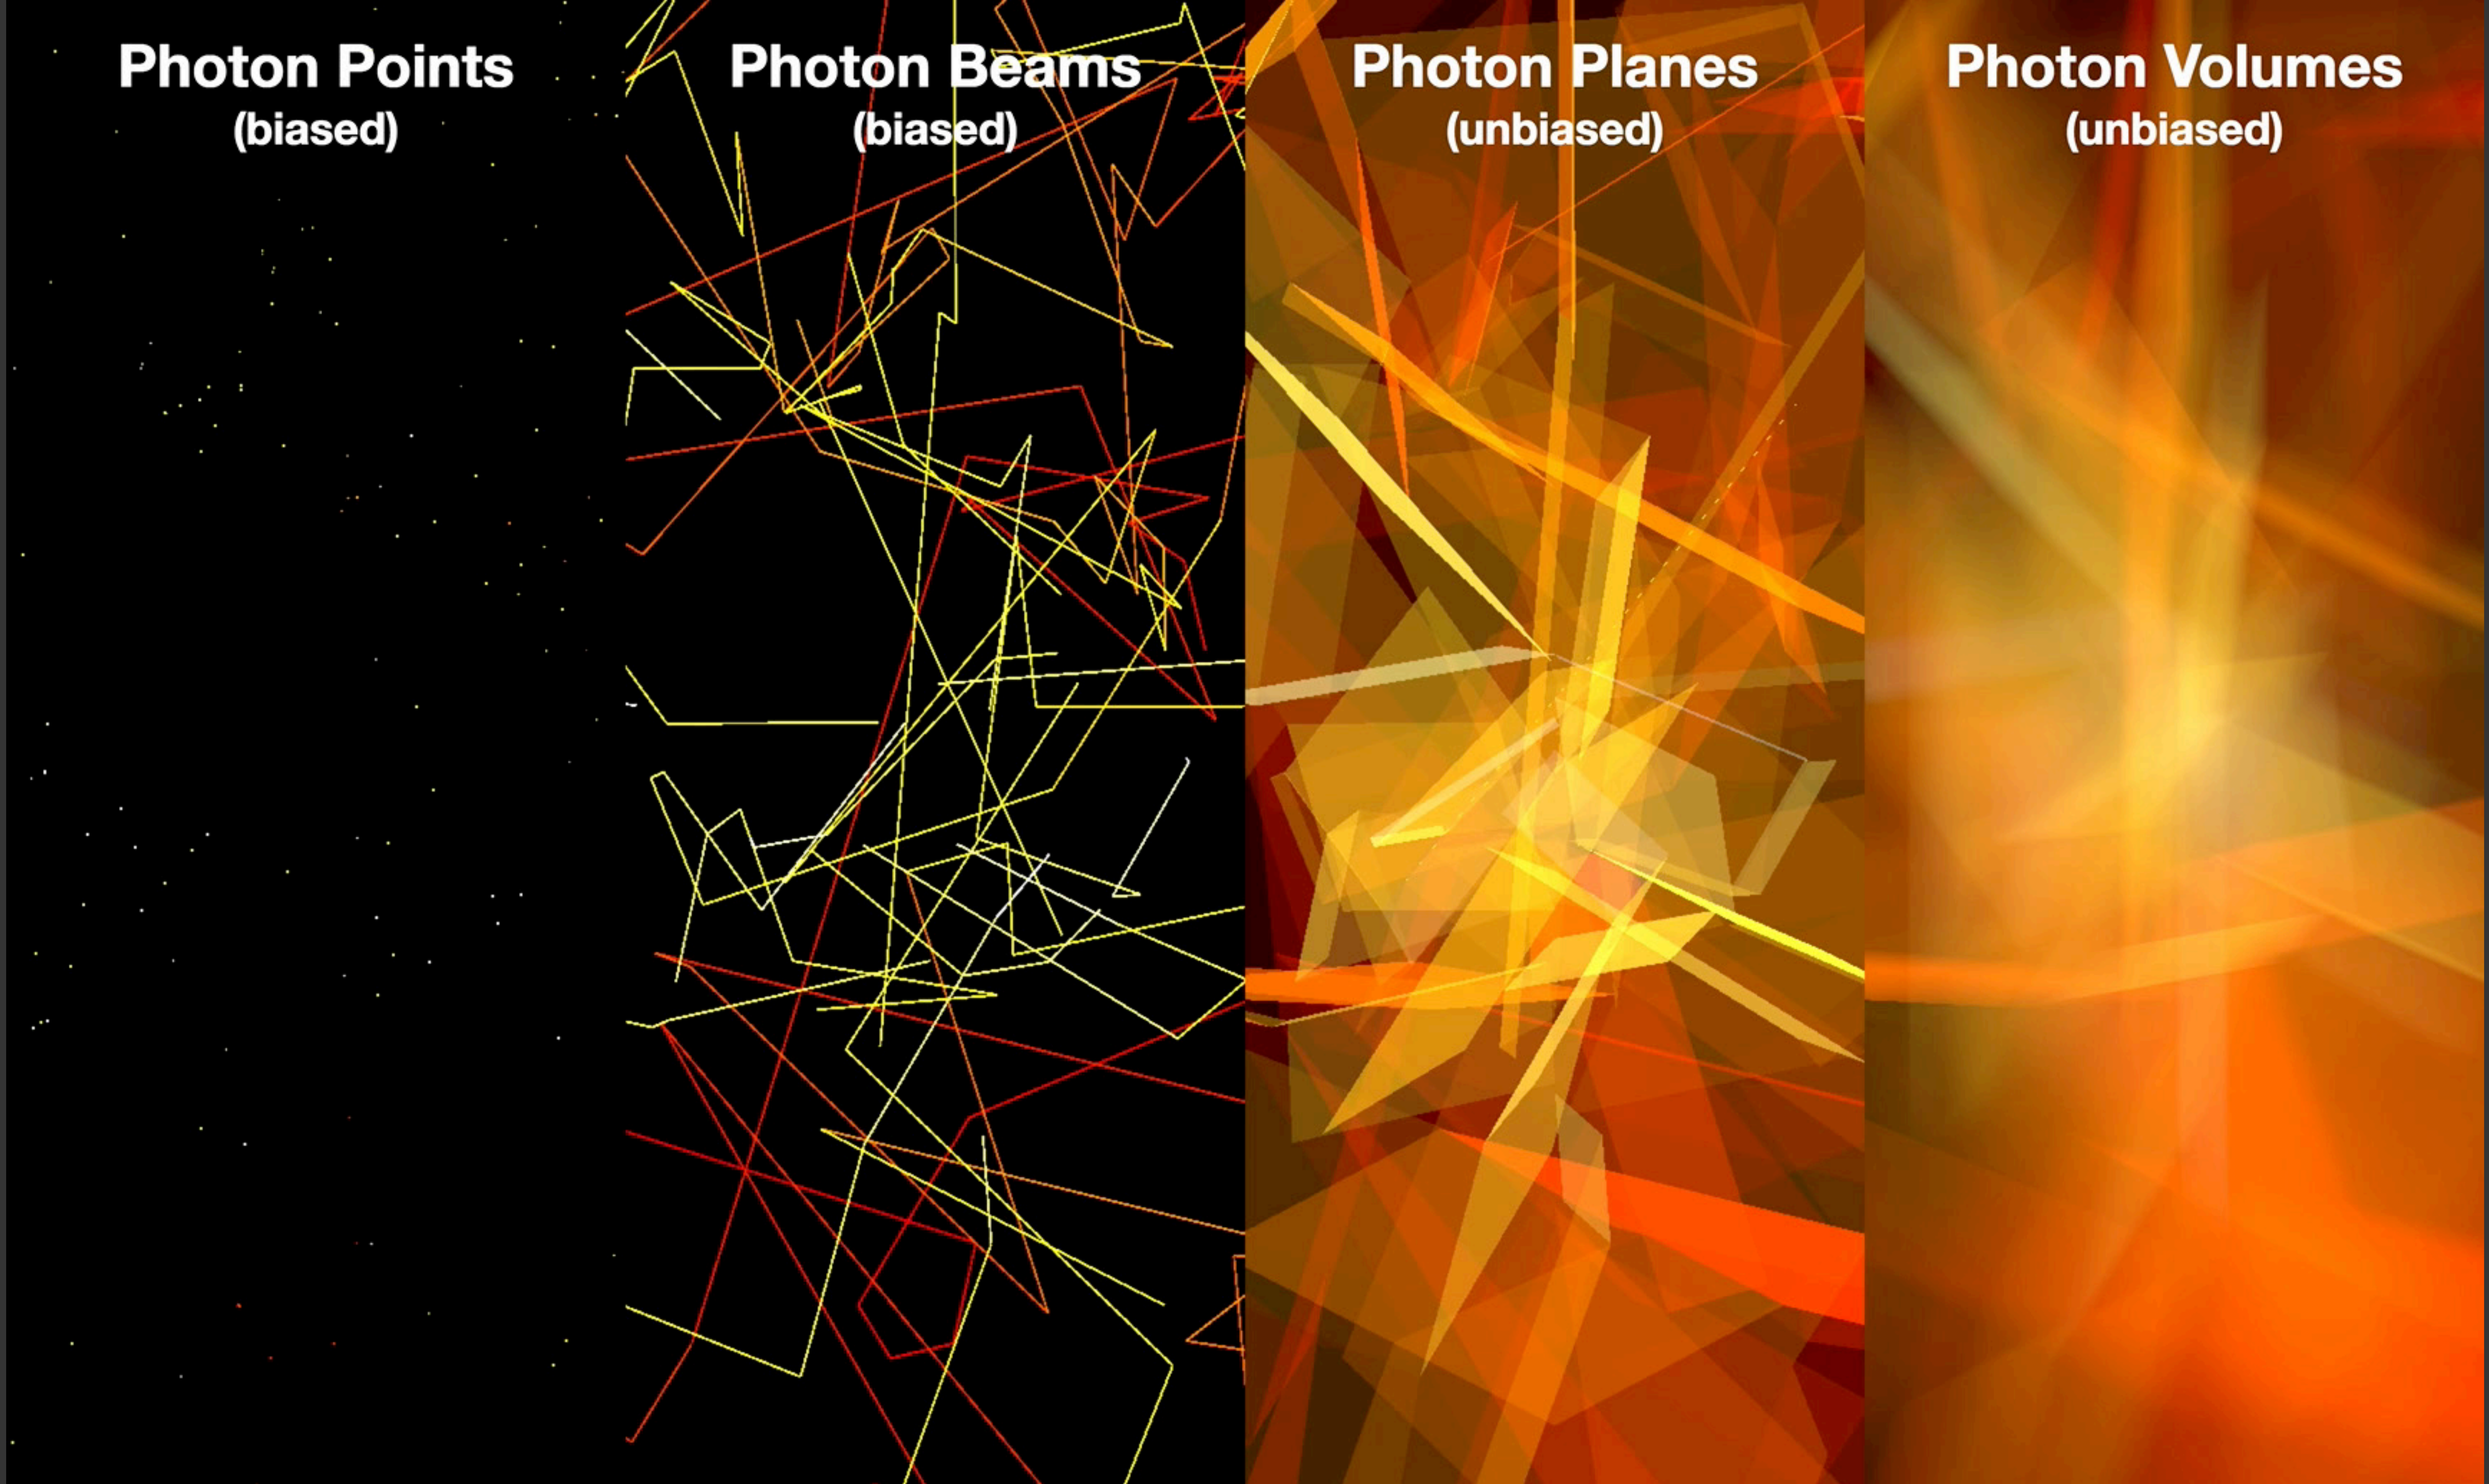
Photon Points
(biased)
Photon Beams
(biased)
Photon Planes
(unbiased)
Photon Volumes
(unbiased)

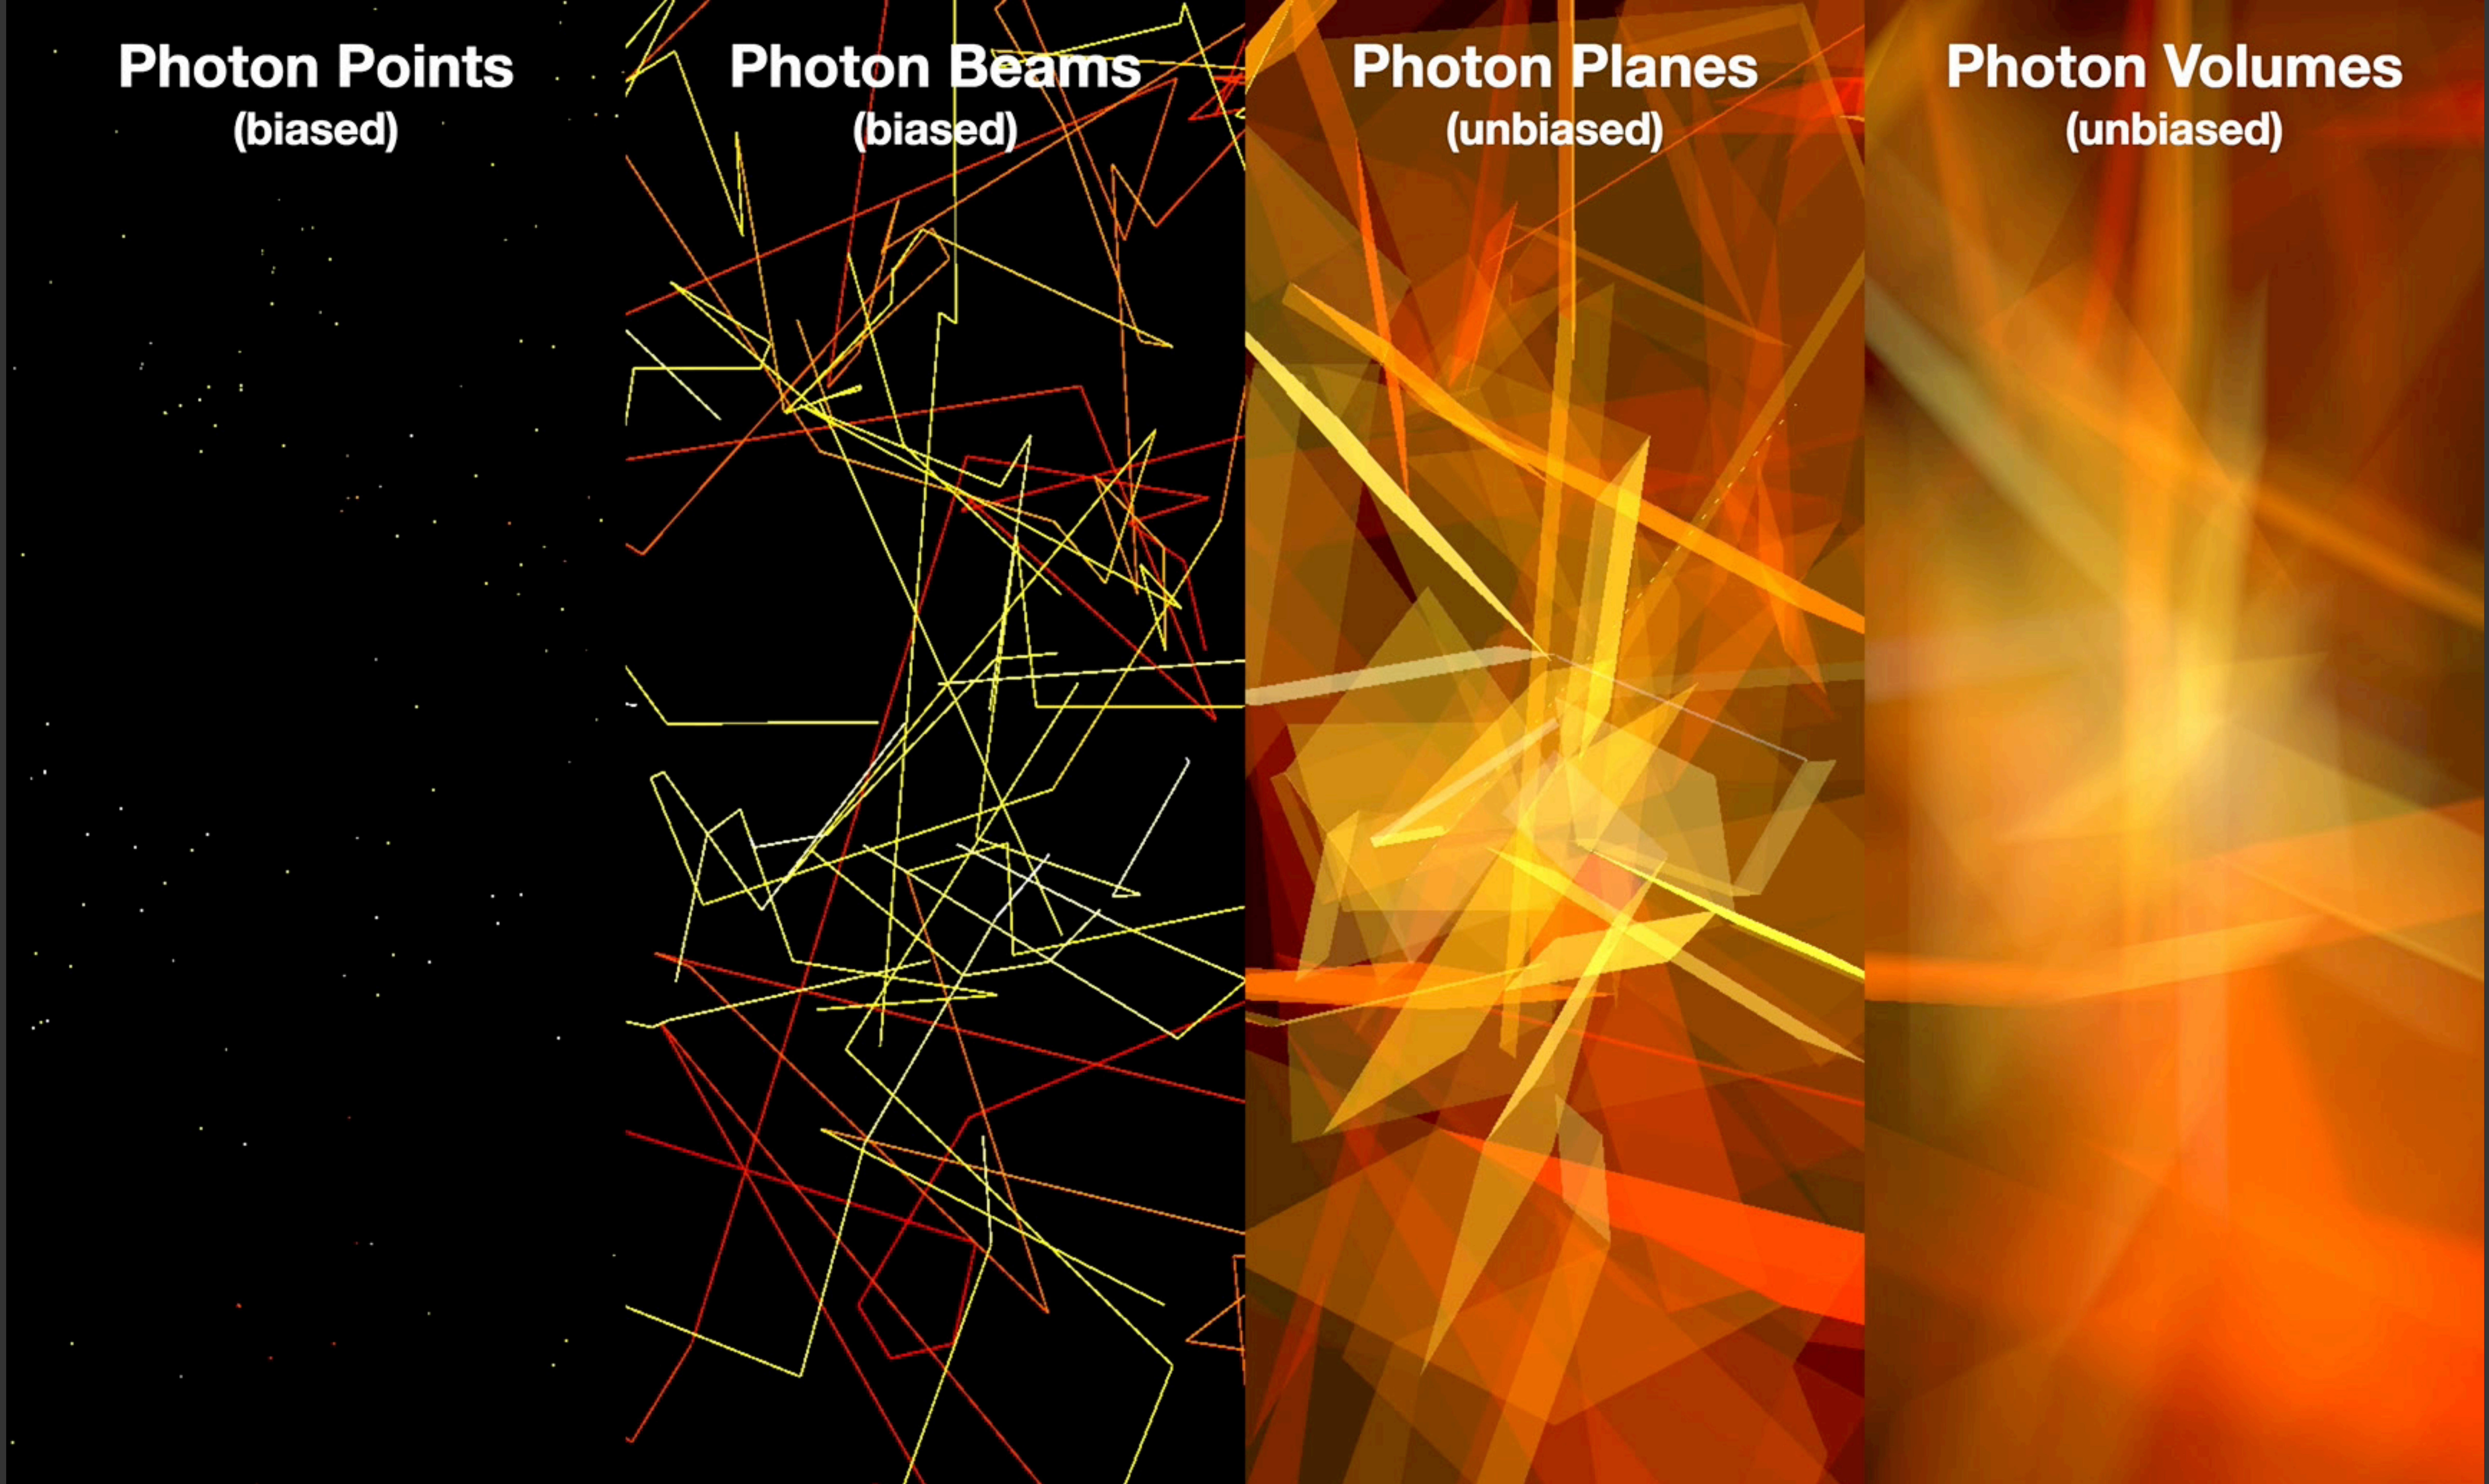
Photon Points
(biased)
Photon Beams
(biased)
Photon Planes
(unbiased)
Photon Volumes
(unbiased)

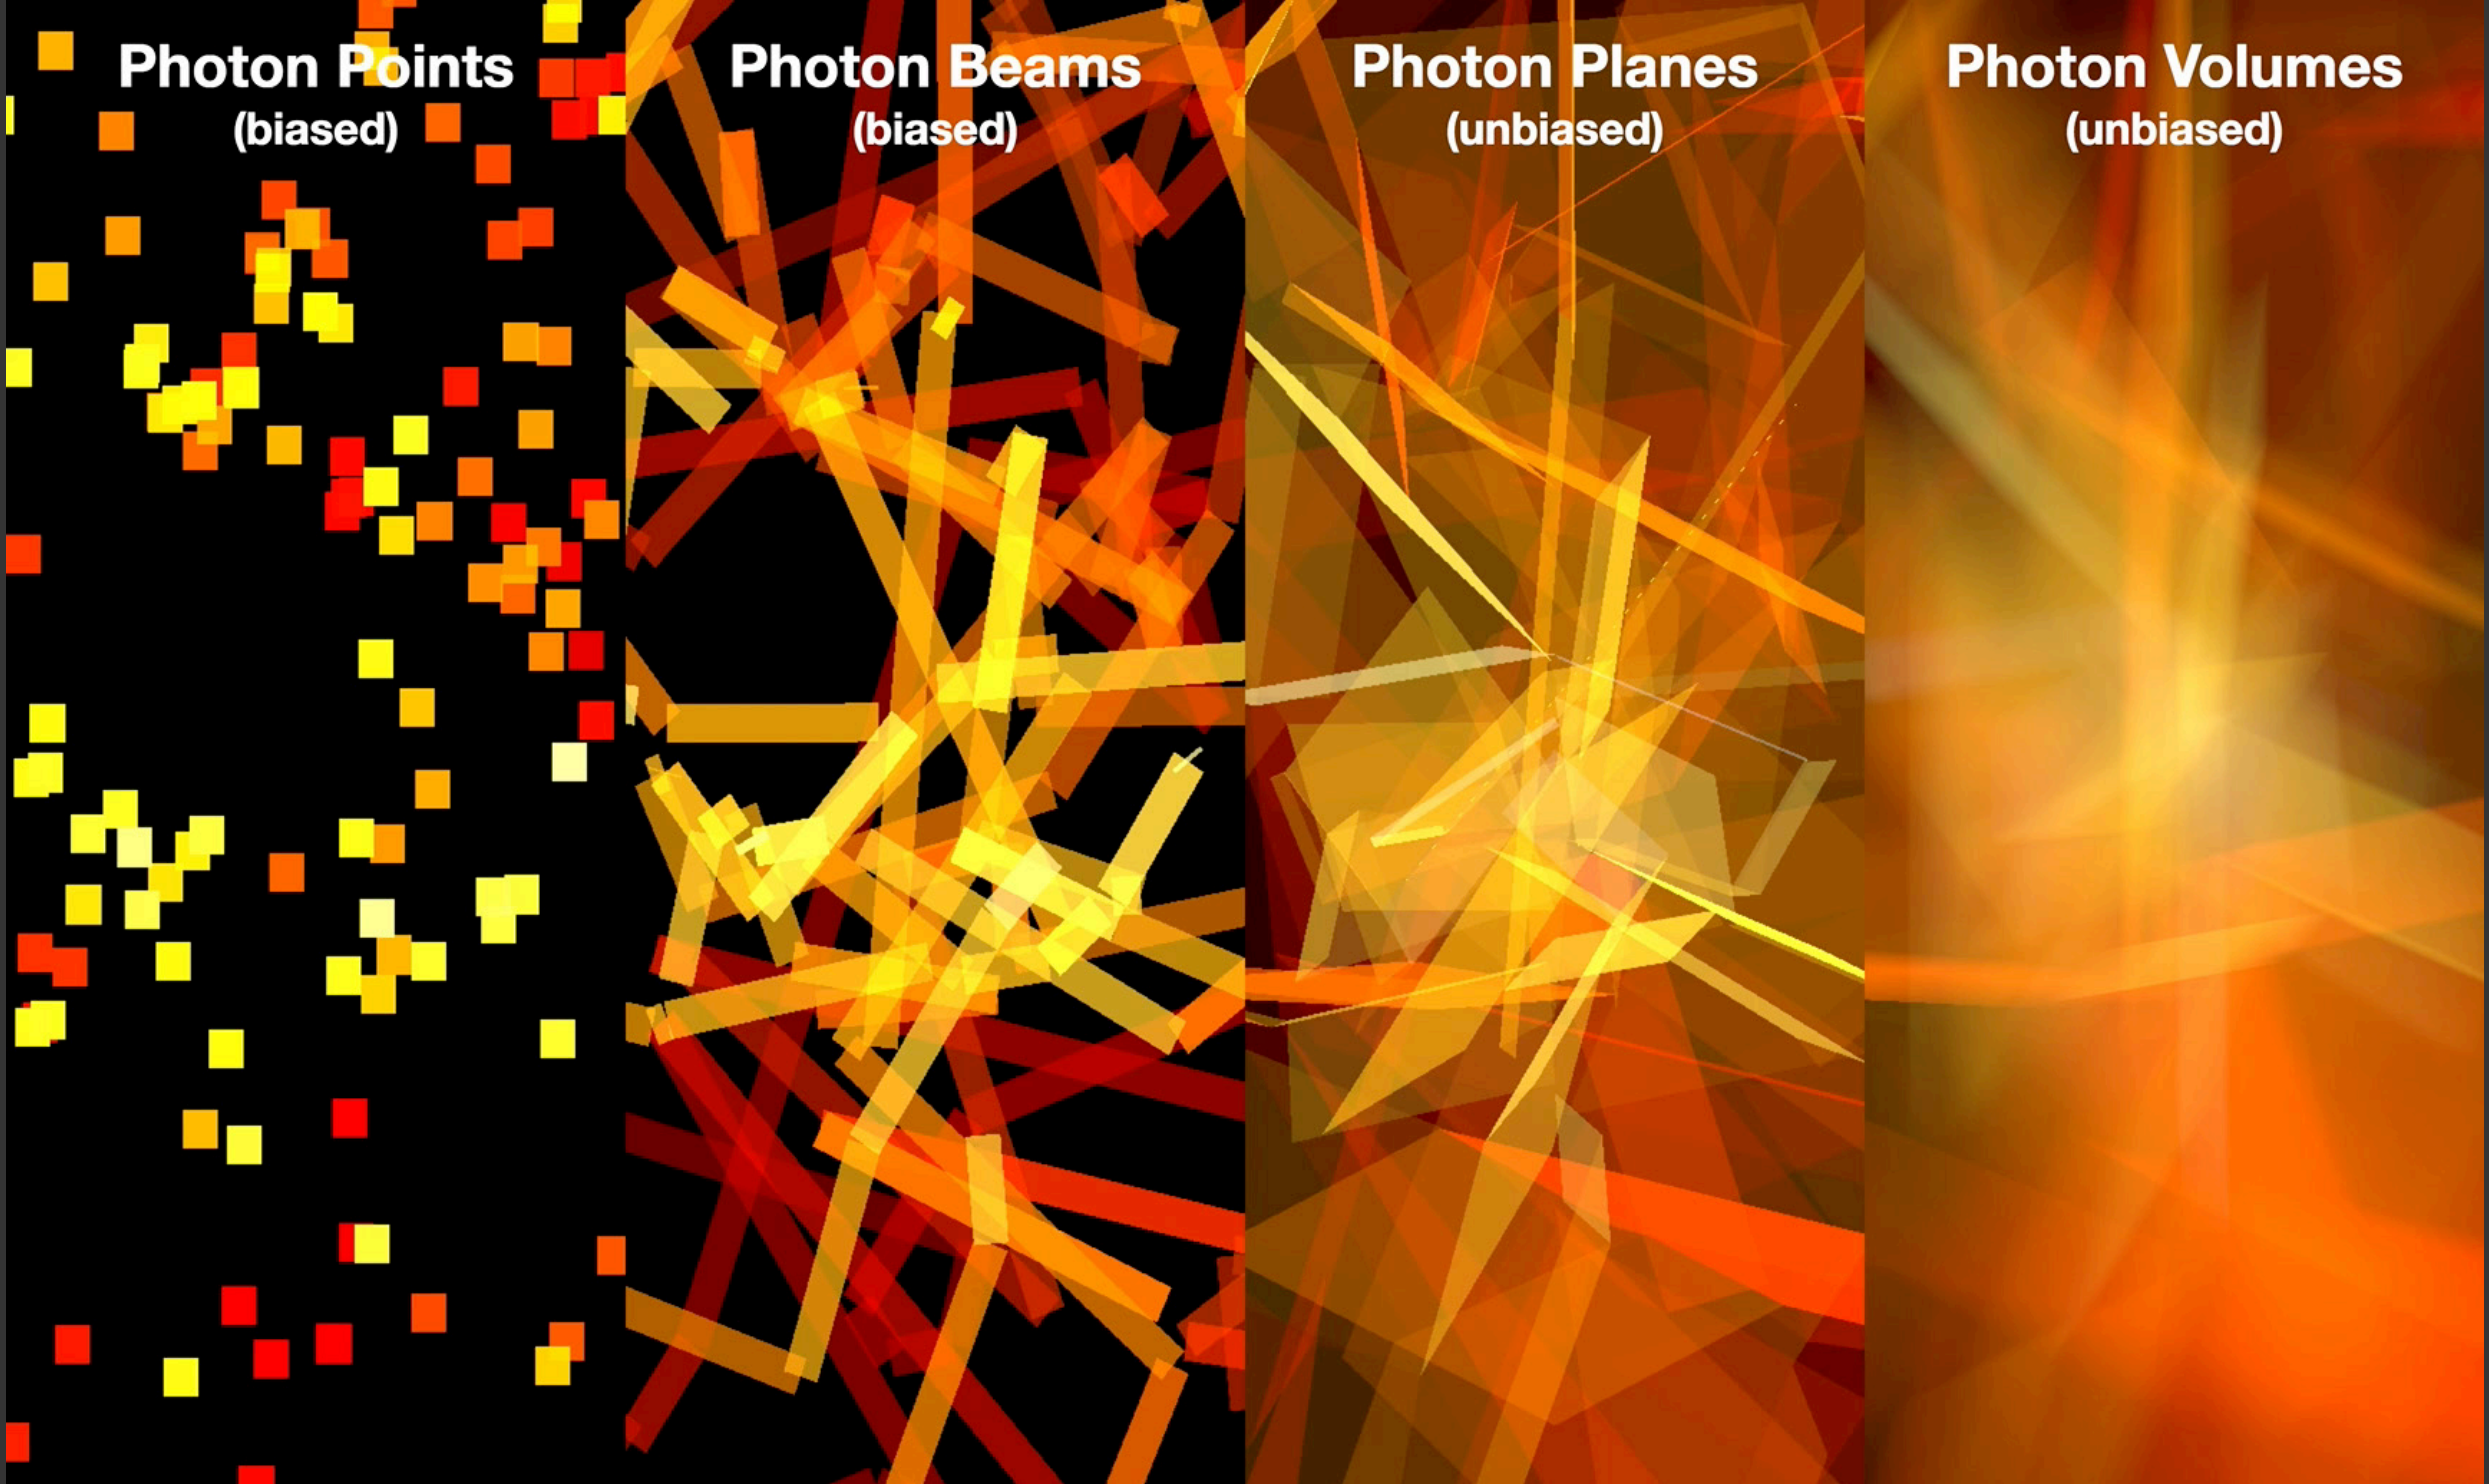
Photon Points
(biased)
Photon Beams
(biased)
Photon Planes
(unbiased)
Photon Volumes
(unbiased)

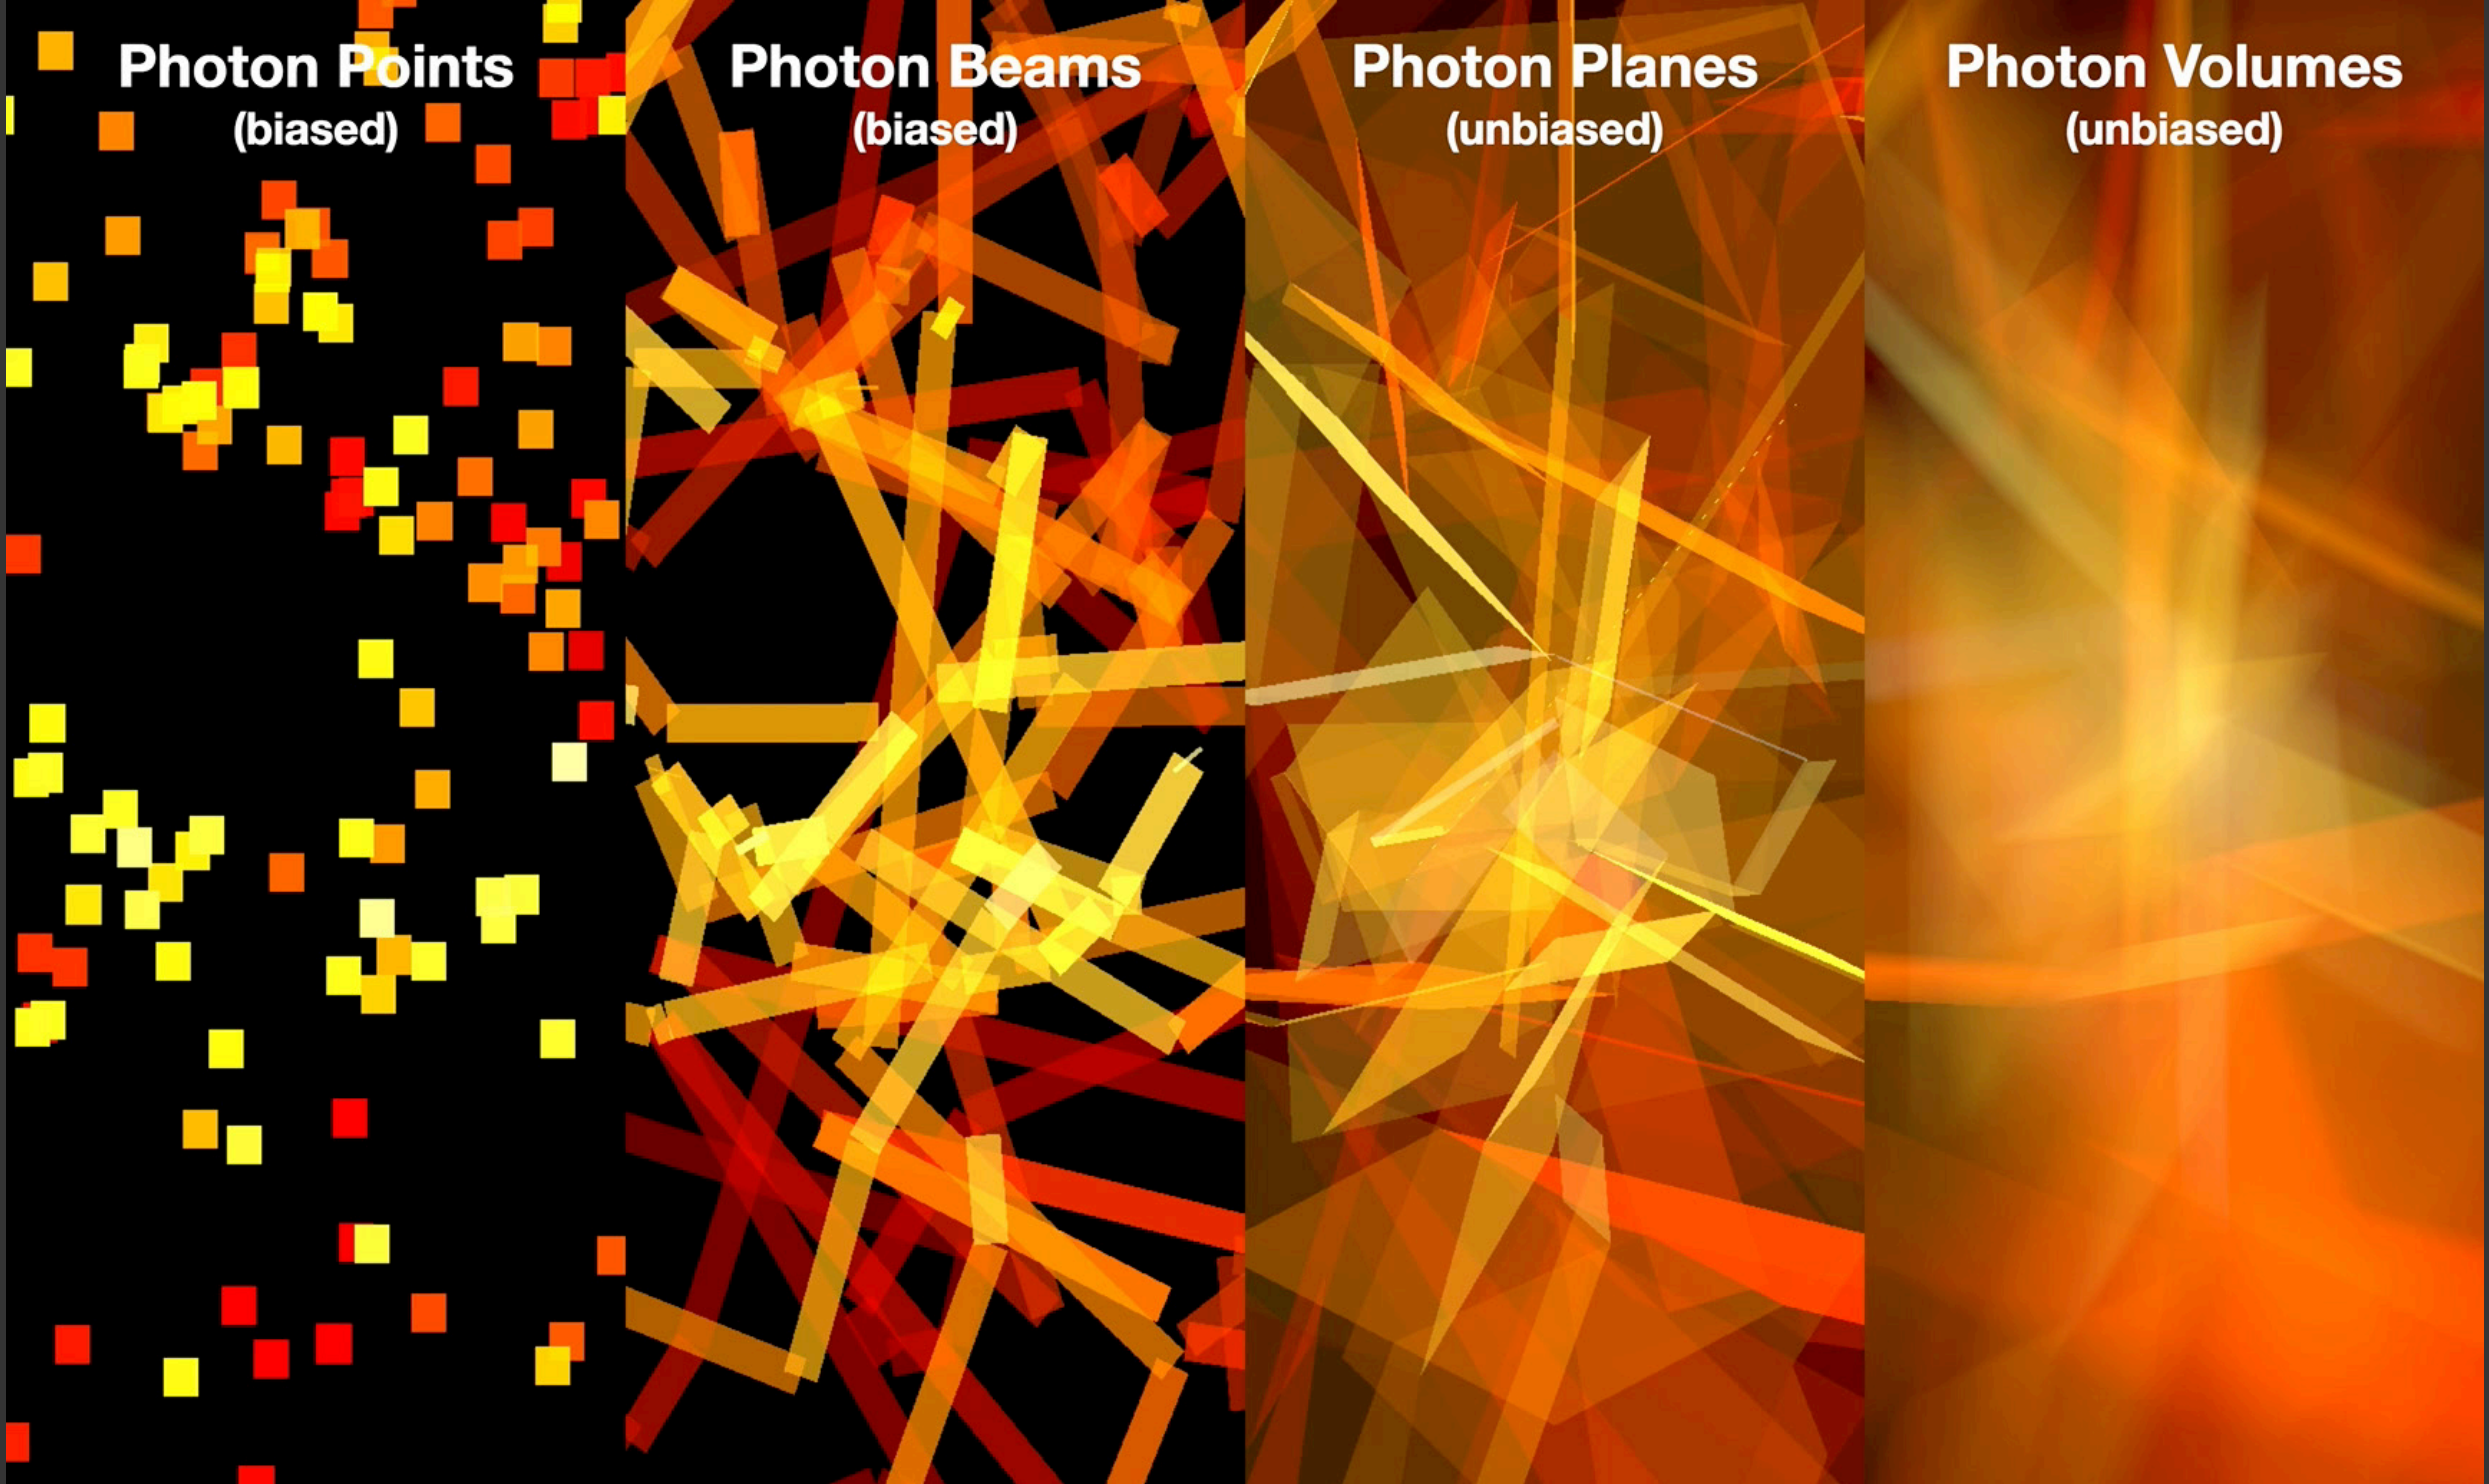
Photon Points
(biased)
Photon Beams
(biased)
Photon Planes
(unbiased)
Photon Volumes
(unbiased)

Raytracing Implementation

Raytracing Implementation

- Two-pass renderer:

Raytracing Implementation

- Two-pass renderer:

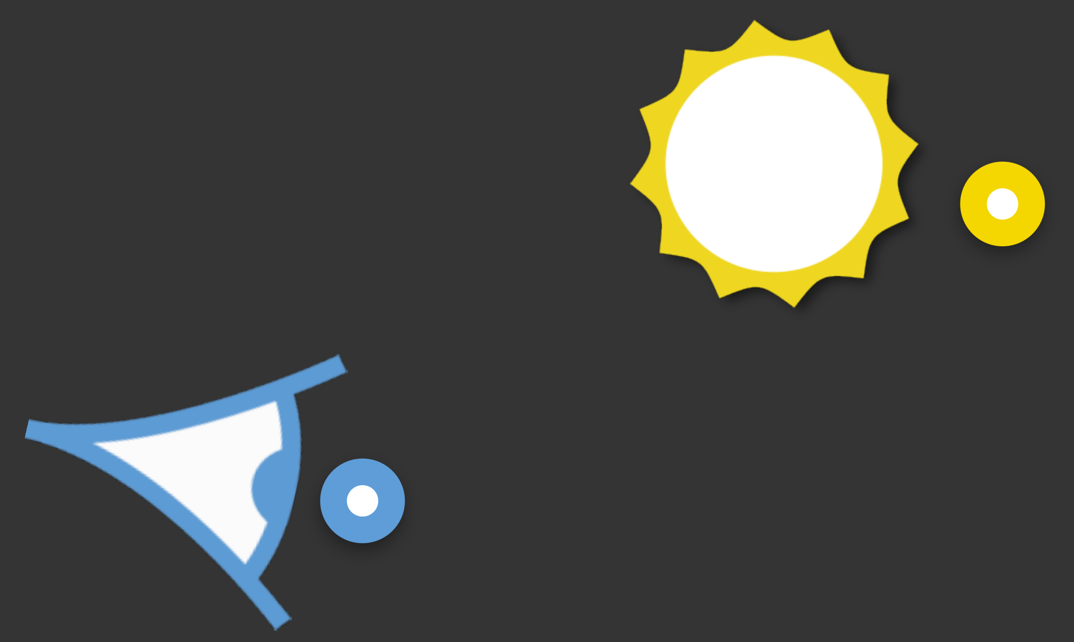

Raytracing Implementation

- Two-pass renderer:

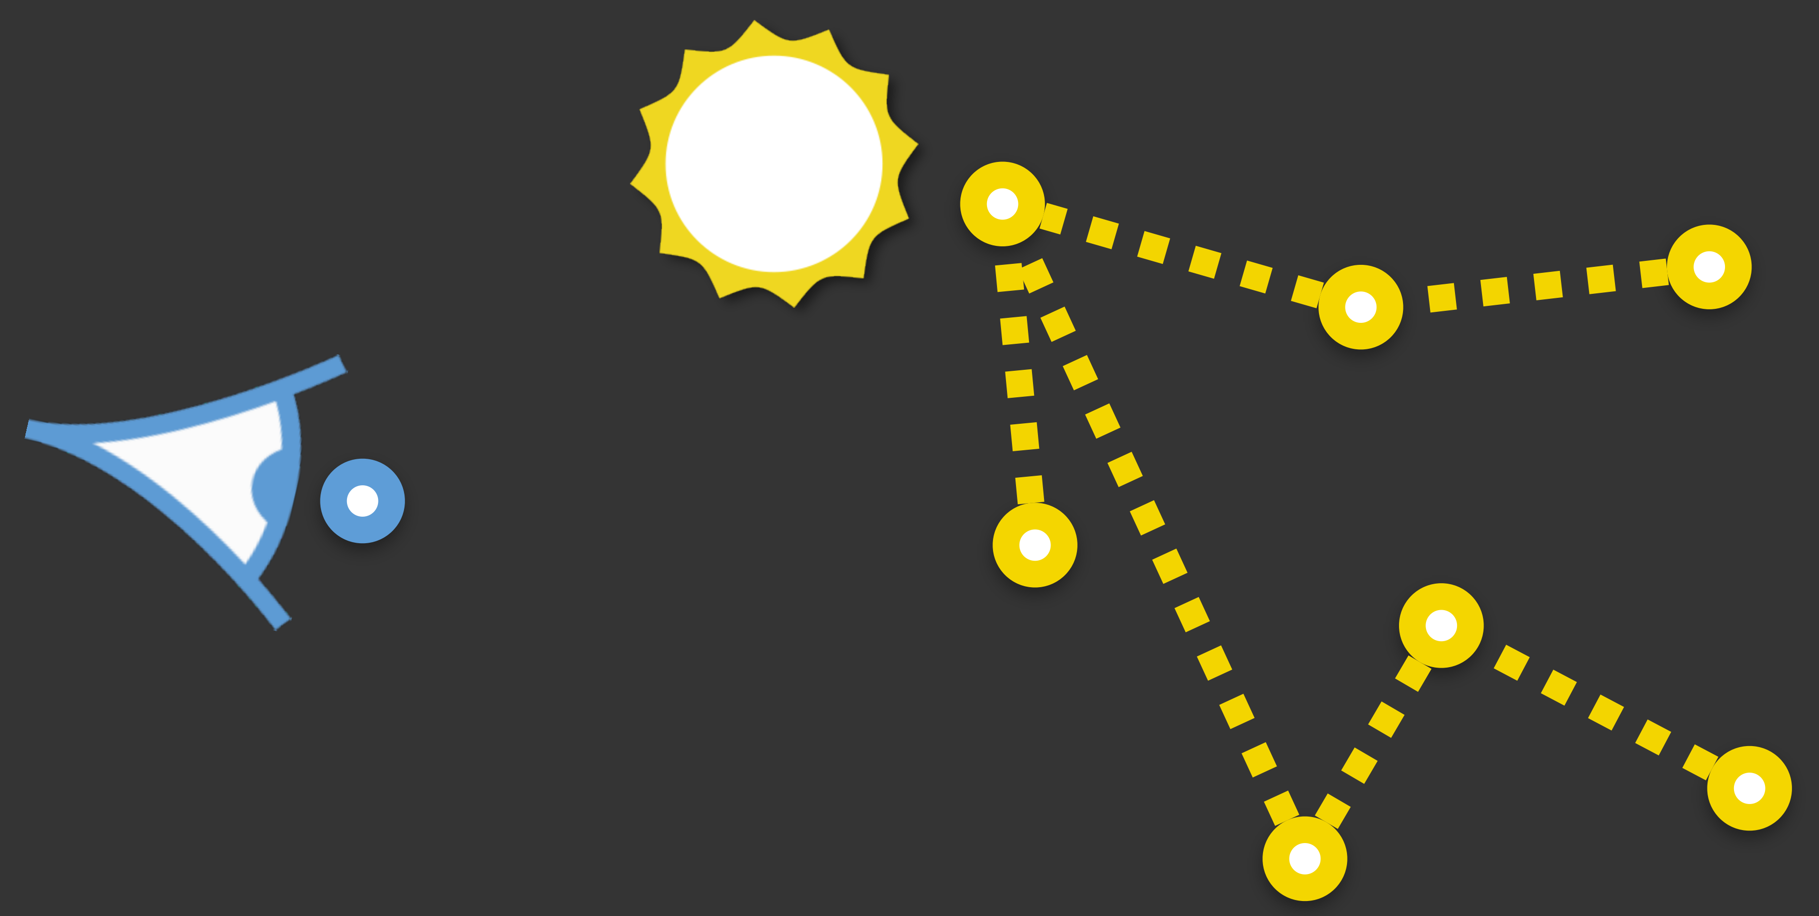

Raytracing Implementation

- Two-pass renderer:

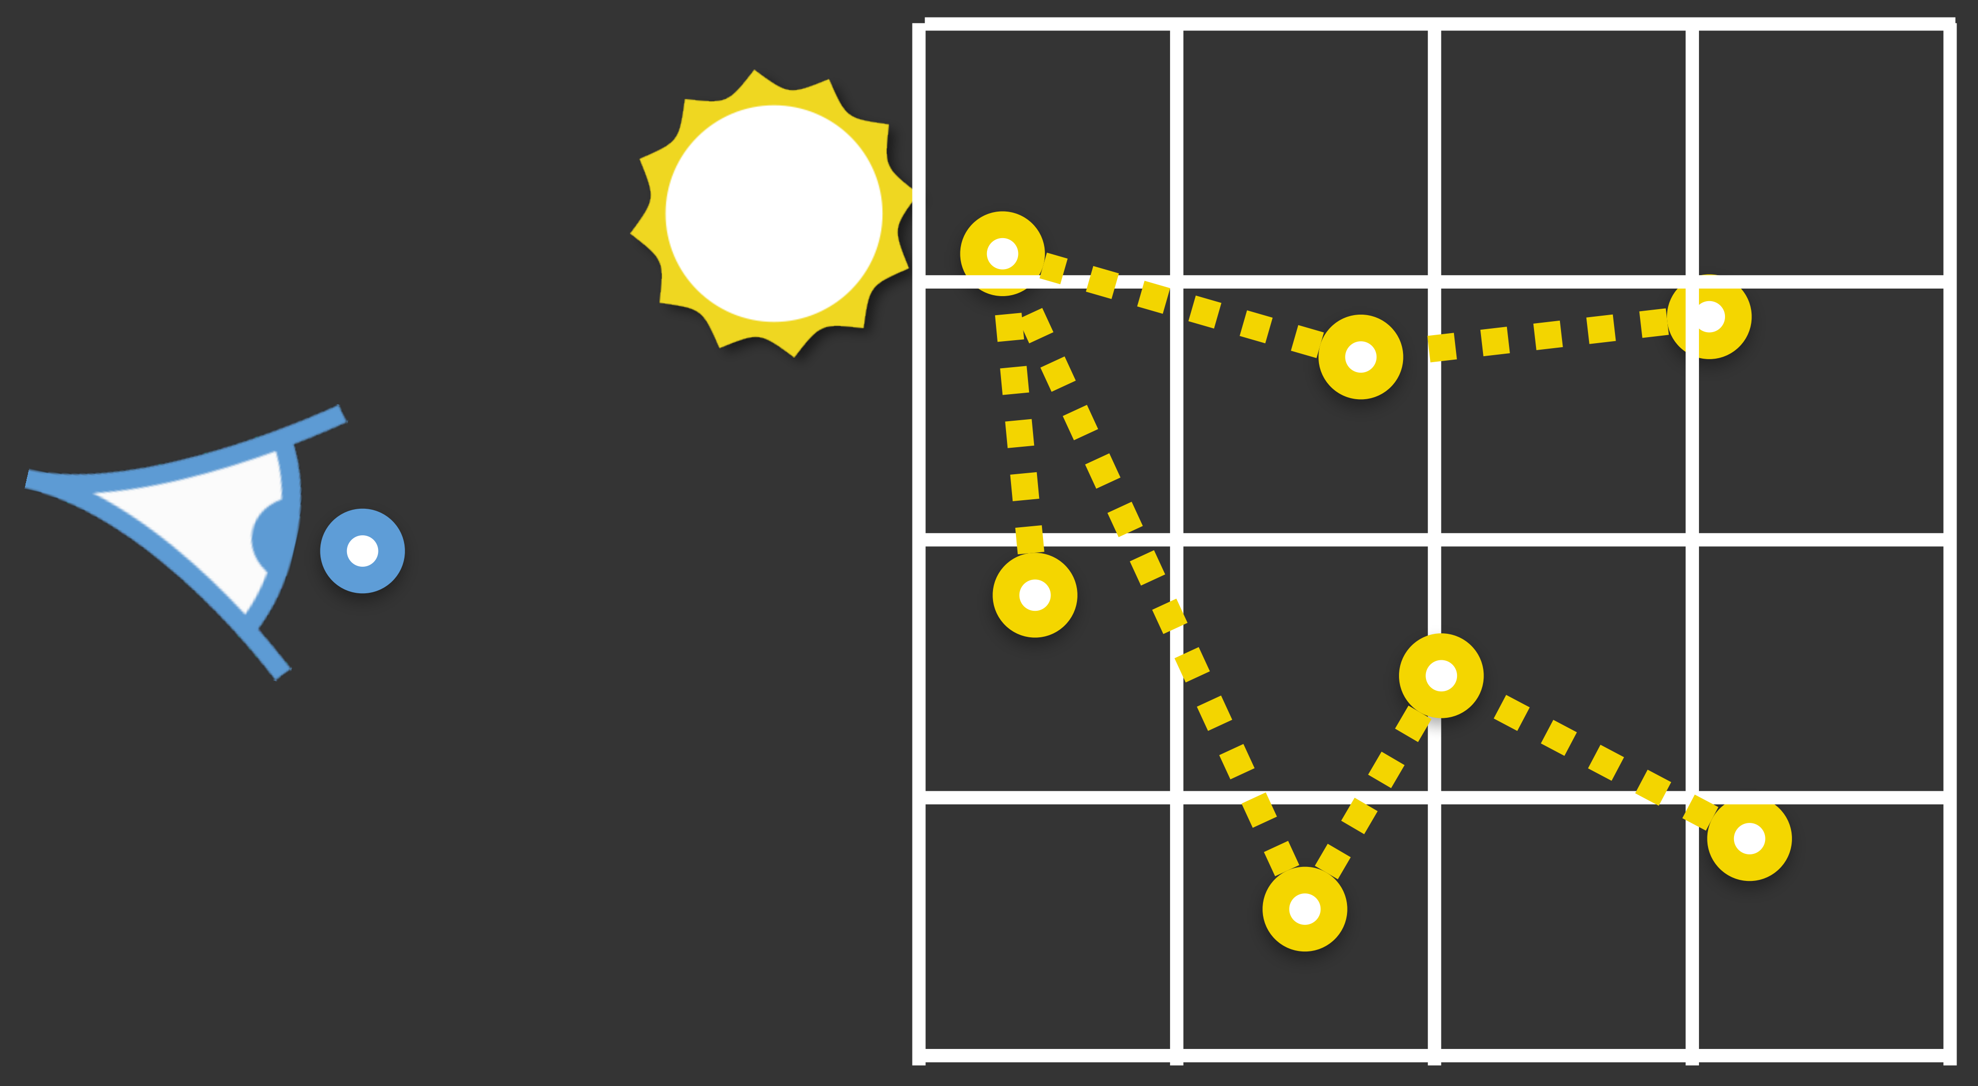

Raytracing Implementation

- Two-pass renderer:

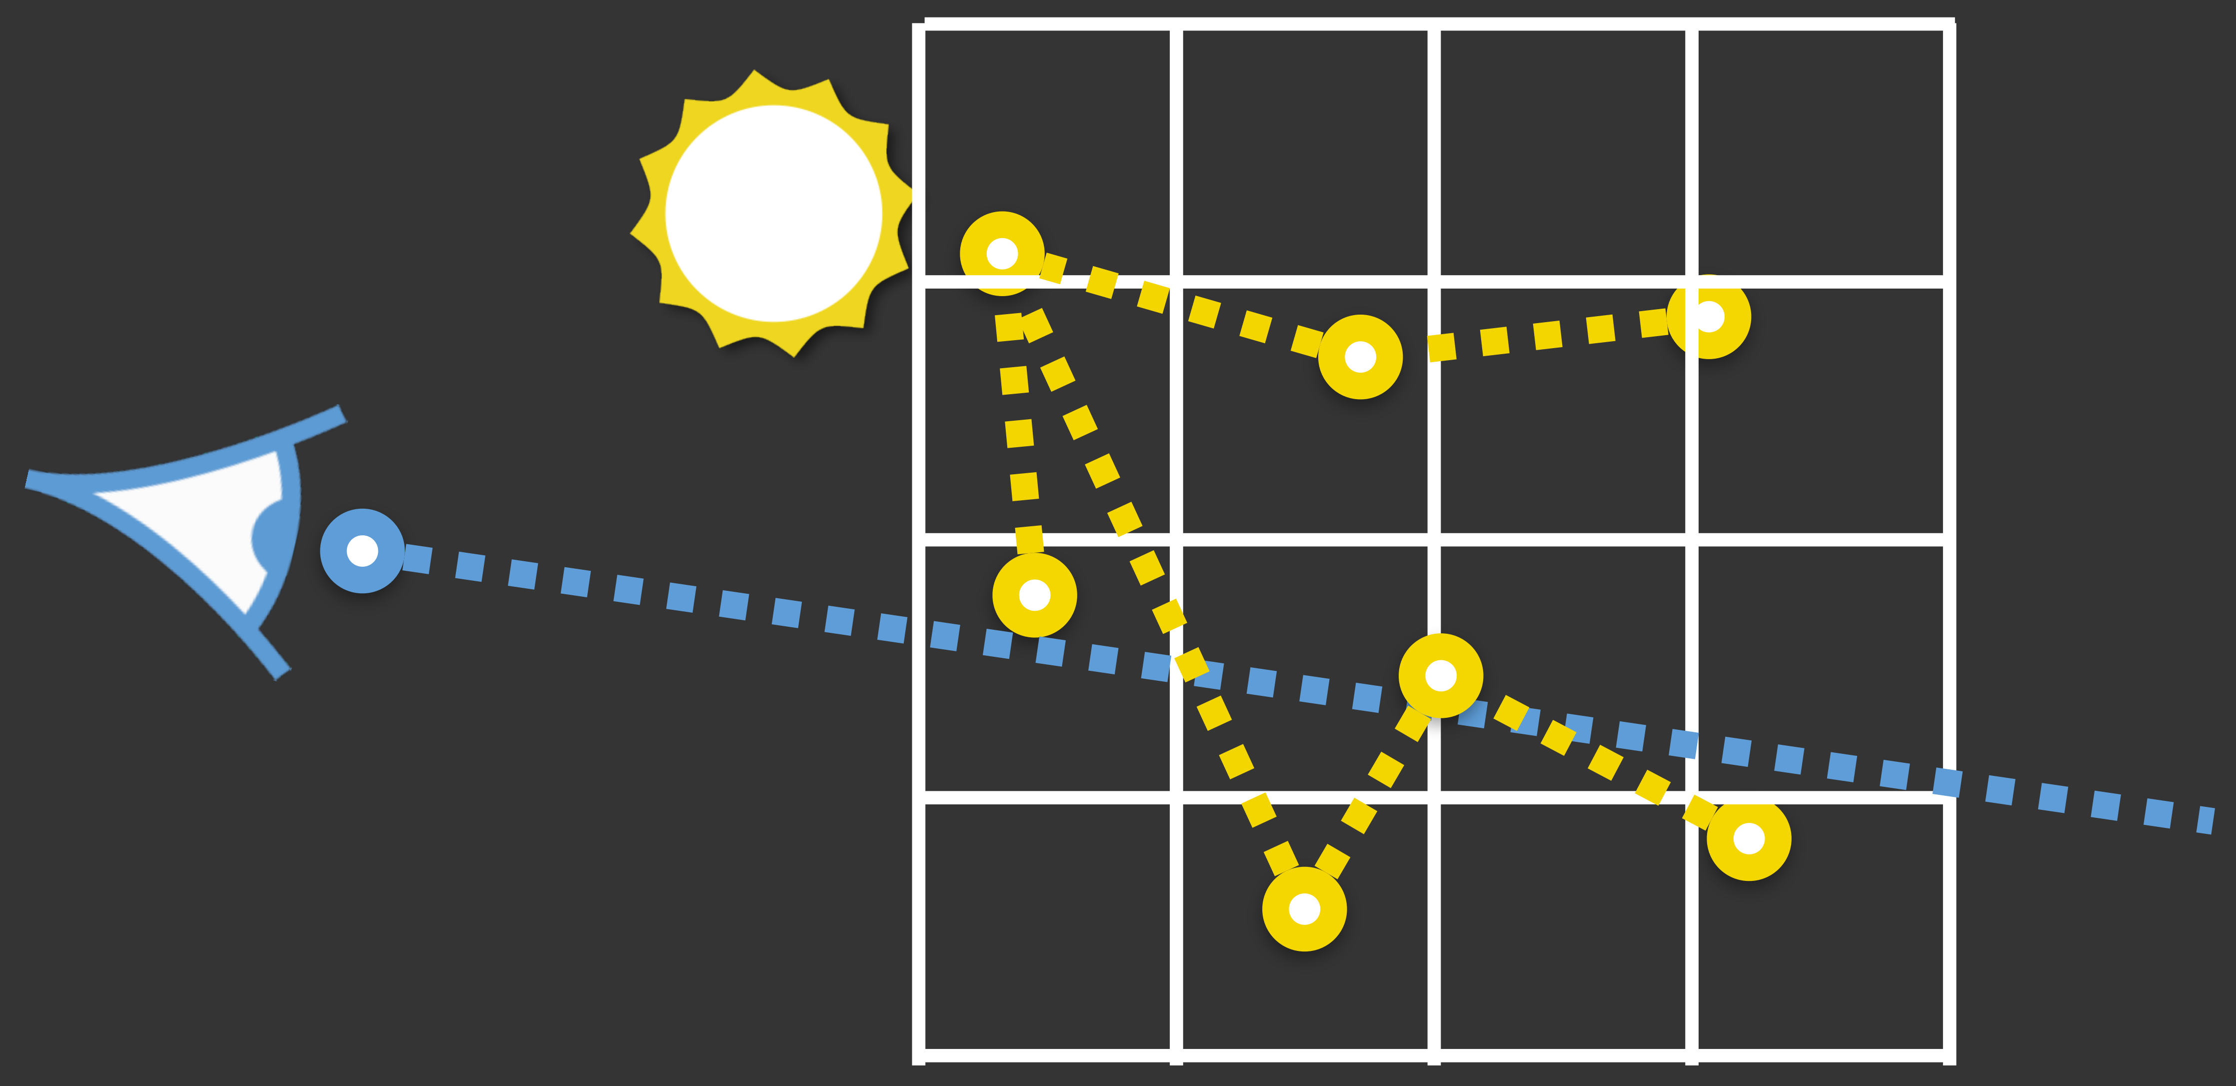

Raytracing Implementation

- Two-pass renderer:

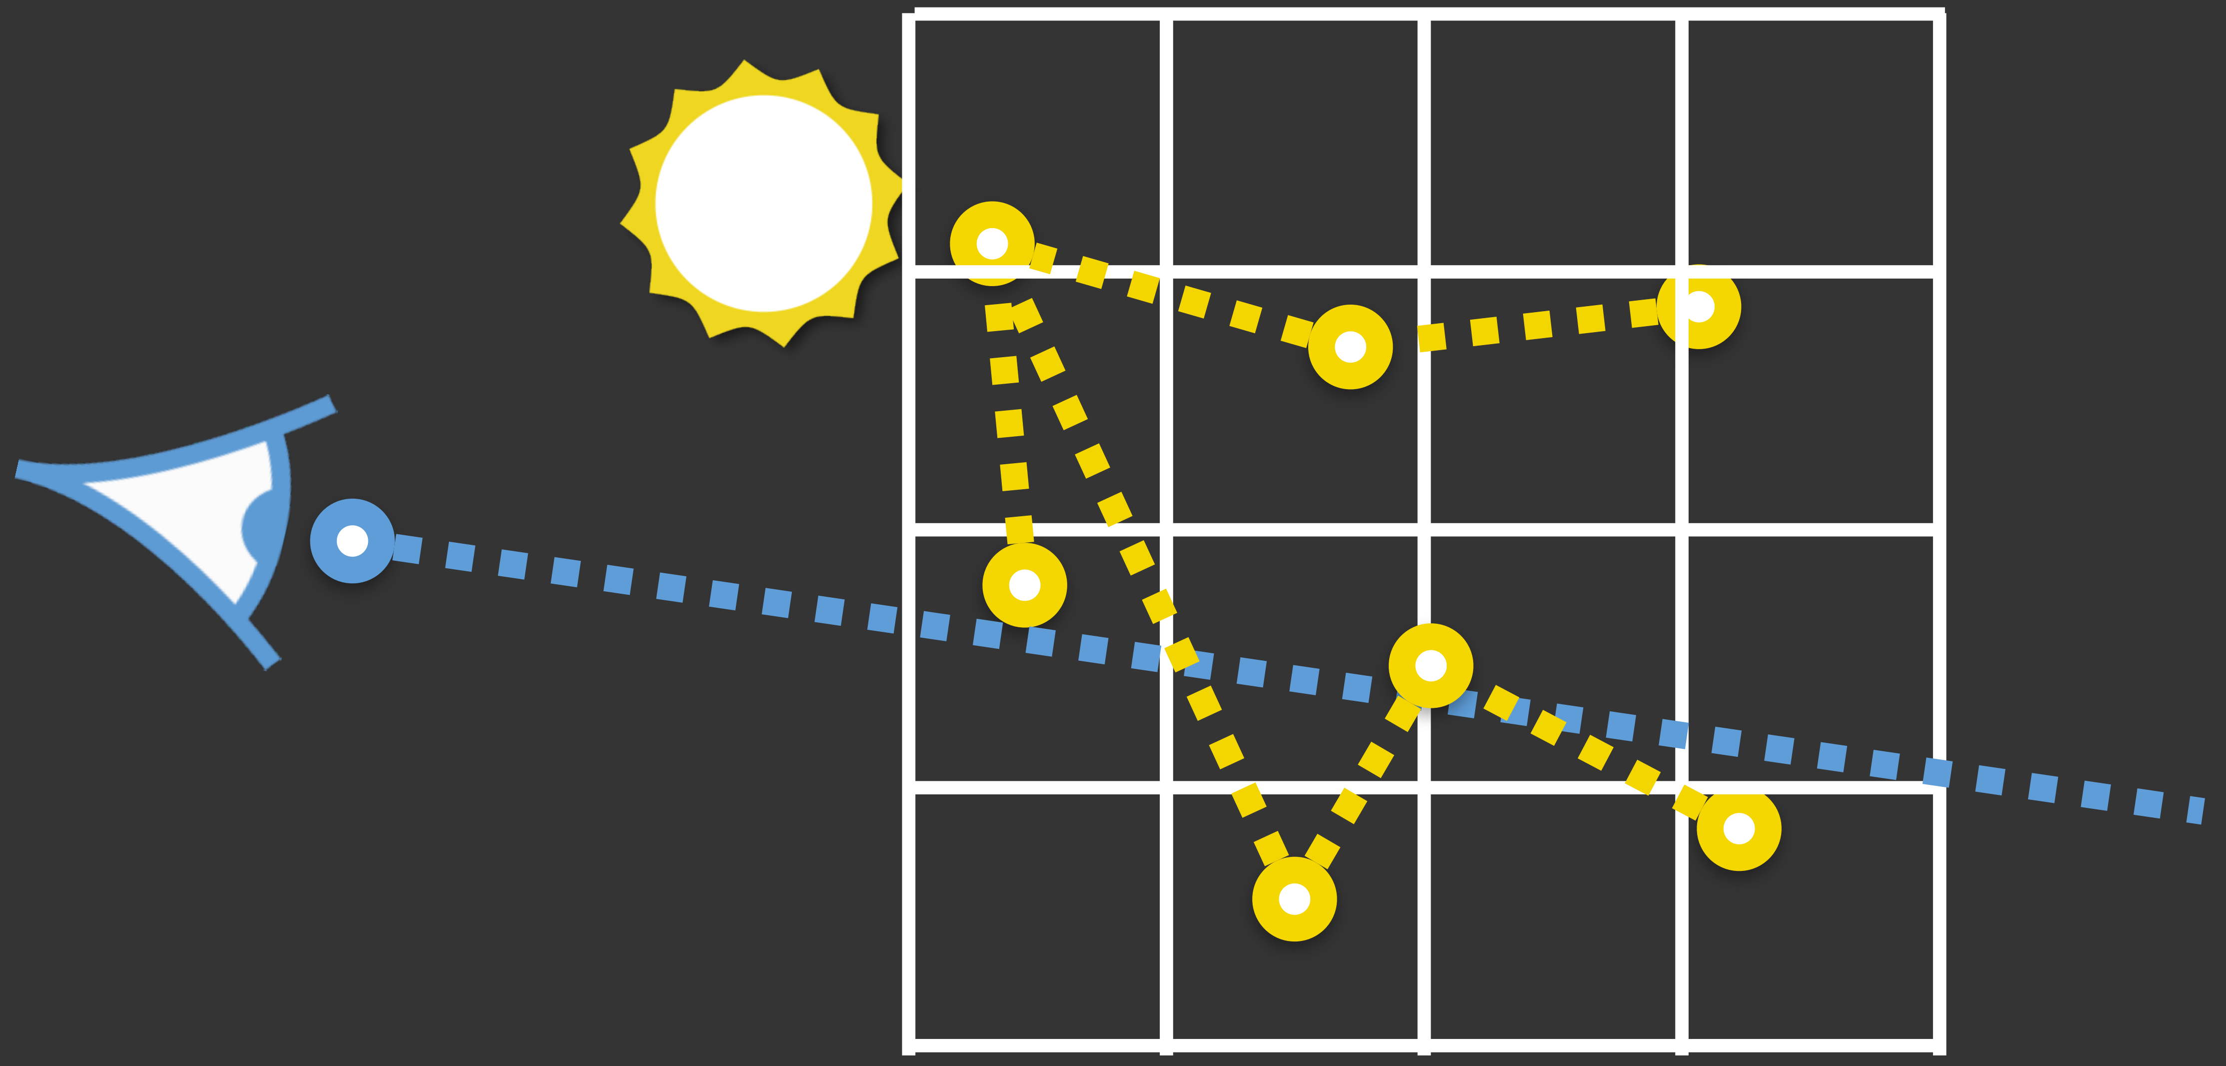

- Acceleration: Details in paper

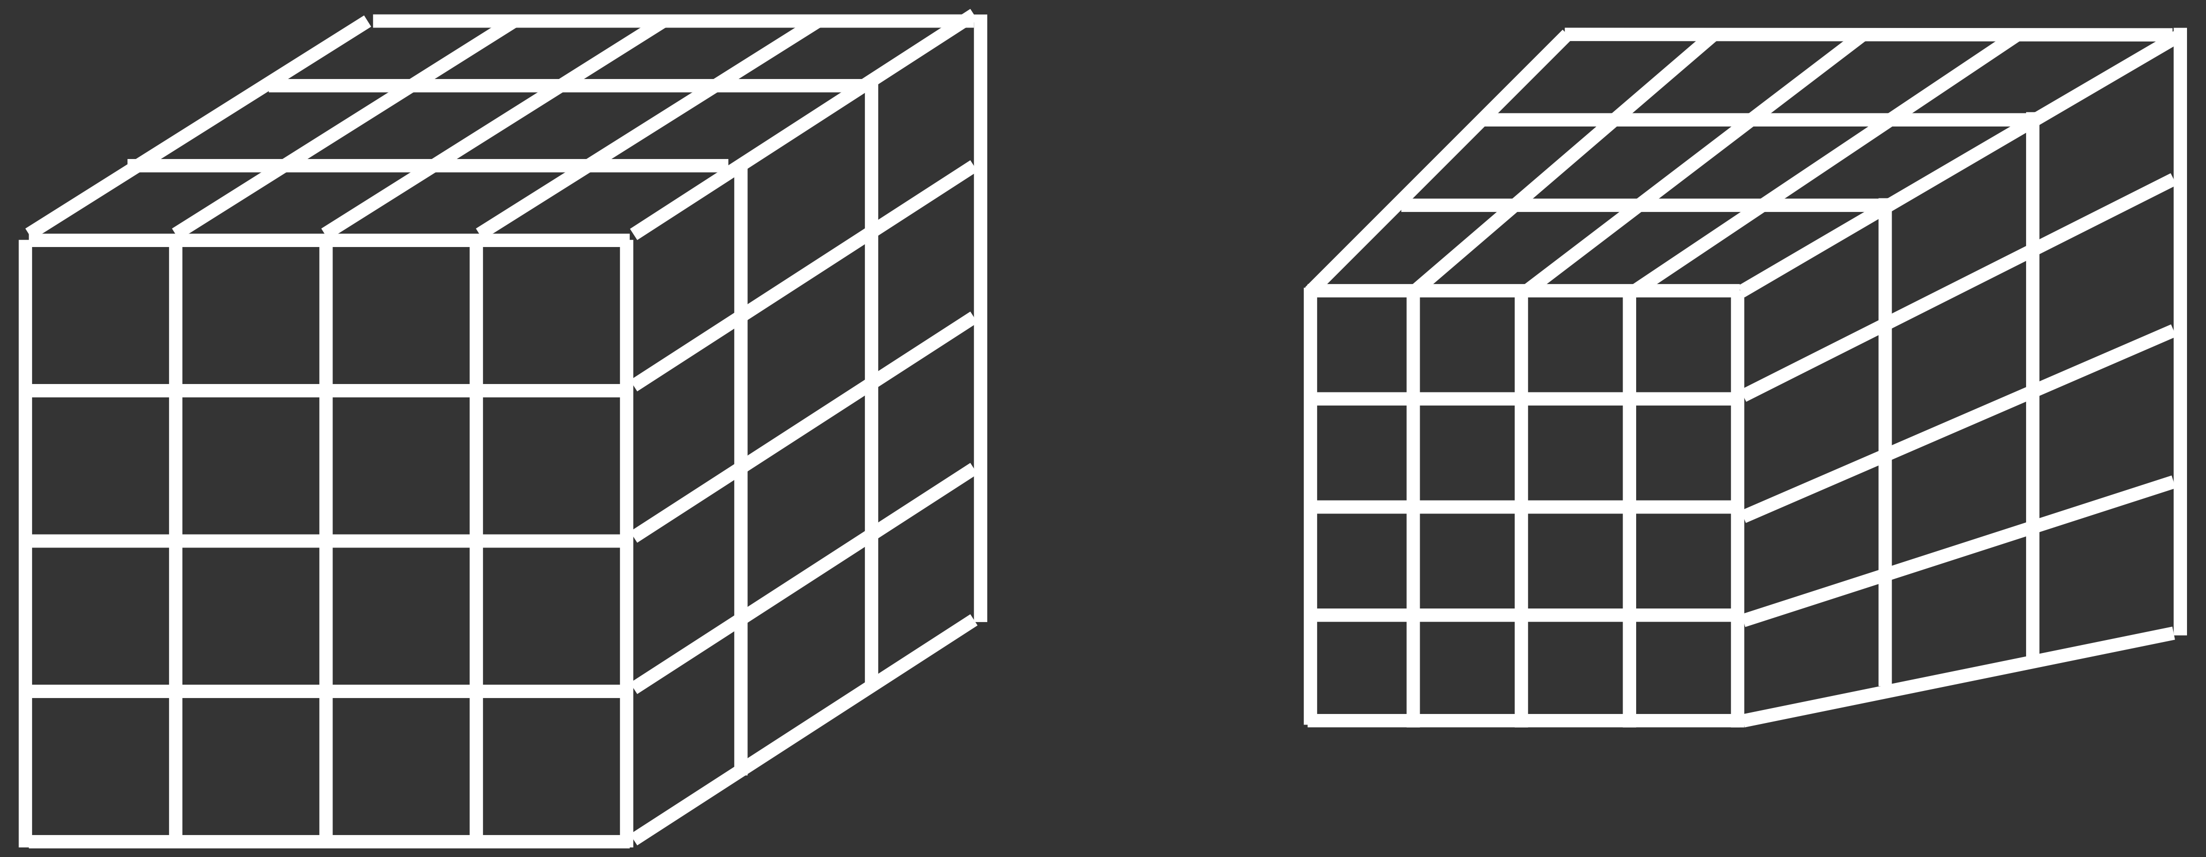

Evaluation

- Test bench of 7 scenes



Evaluation

Evaluation

- Qualitative comparison: Equal time renders (5m)

Evaluation

- Qualitative comparison: Equal time renders (5m)

- Quantitative comparison:

Var [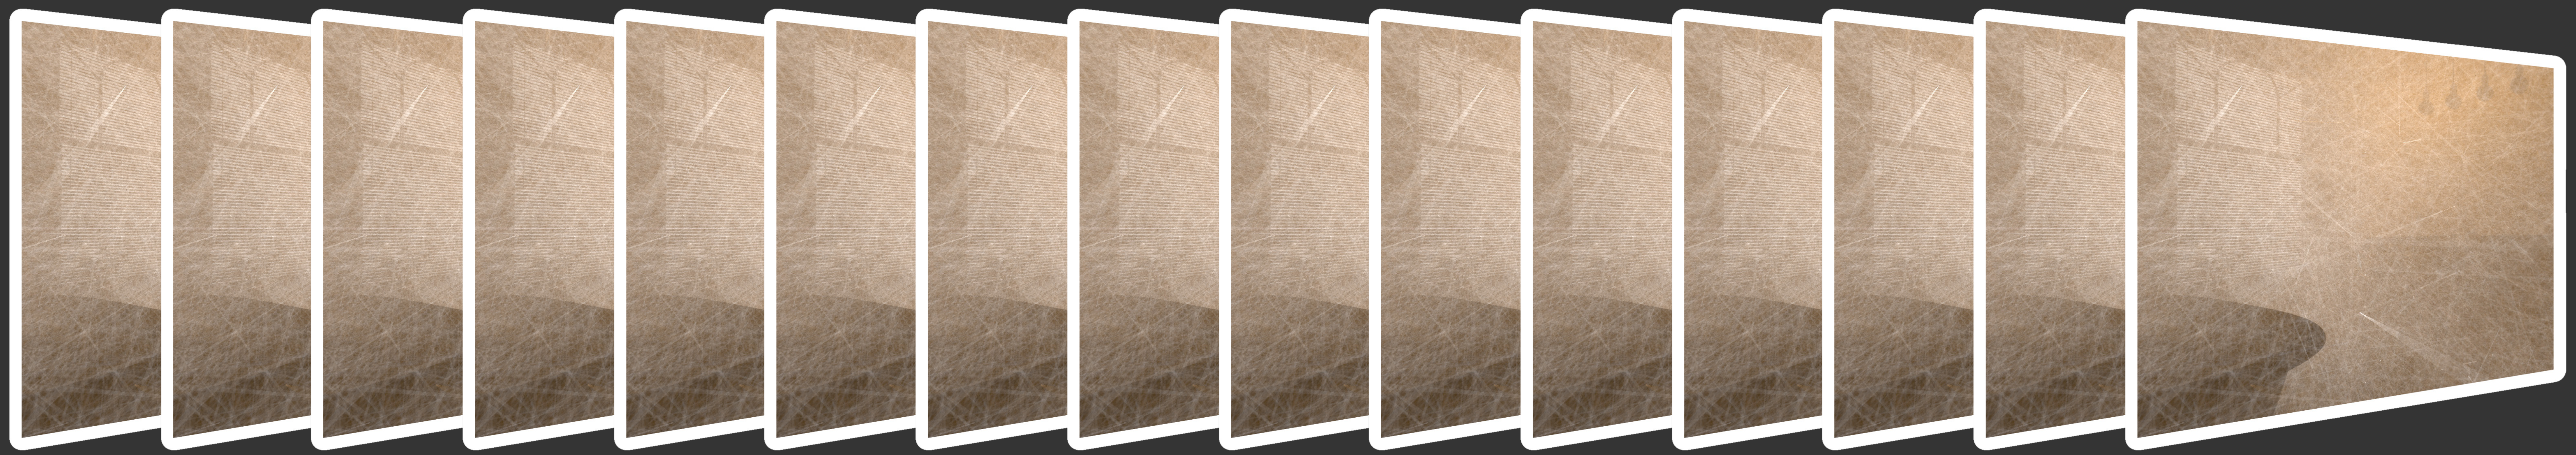]

Evaluation

- Qualitative comparison: Equal time renders (5m)

- Quantitative comparison:

Var [

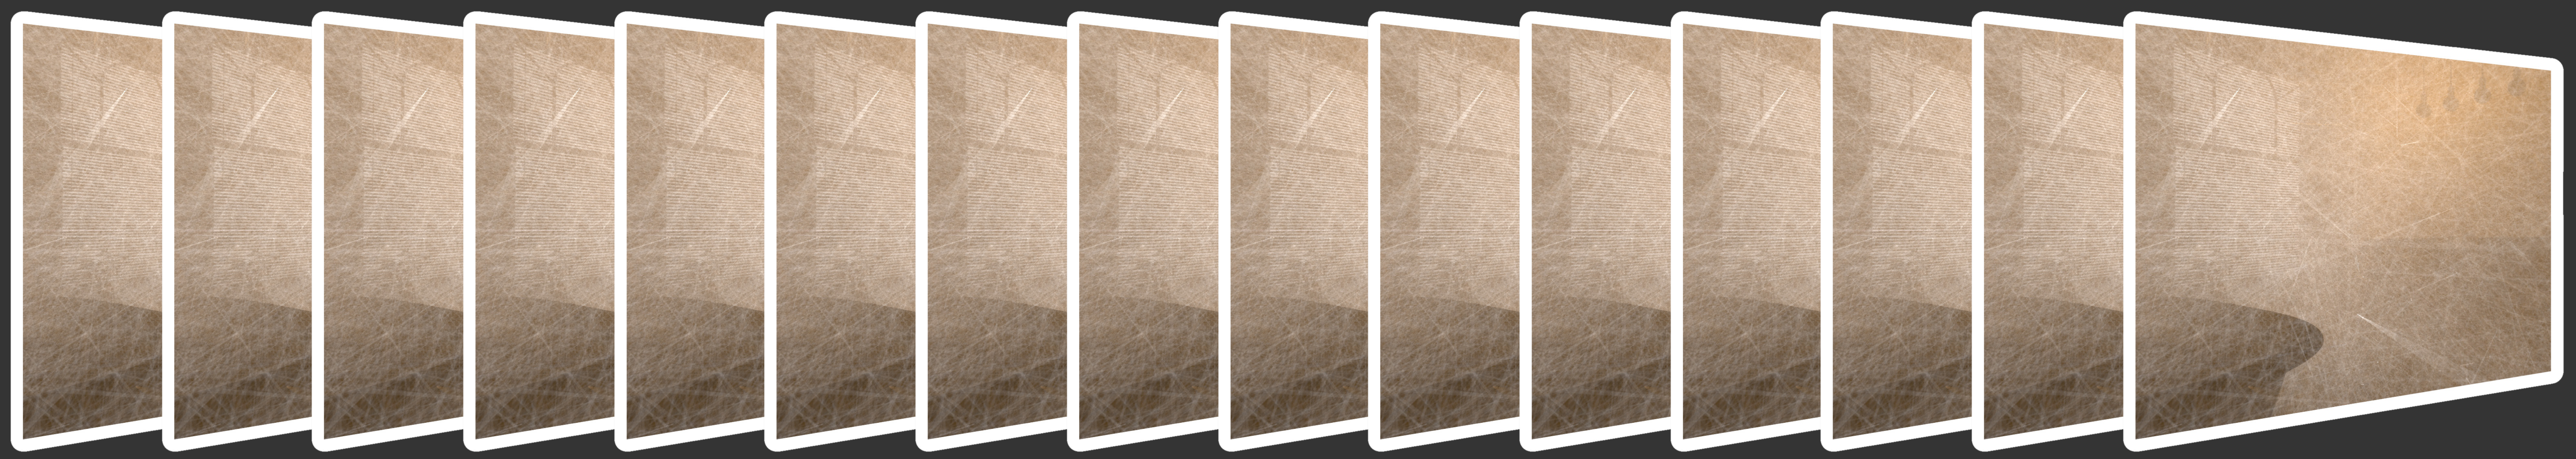

]

- Speedup: Ratio of variance


Full Light Transport

Photon Beams (1D blur)

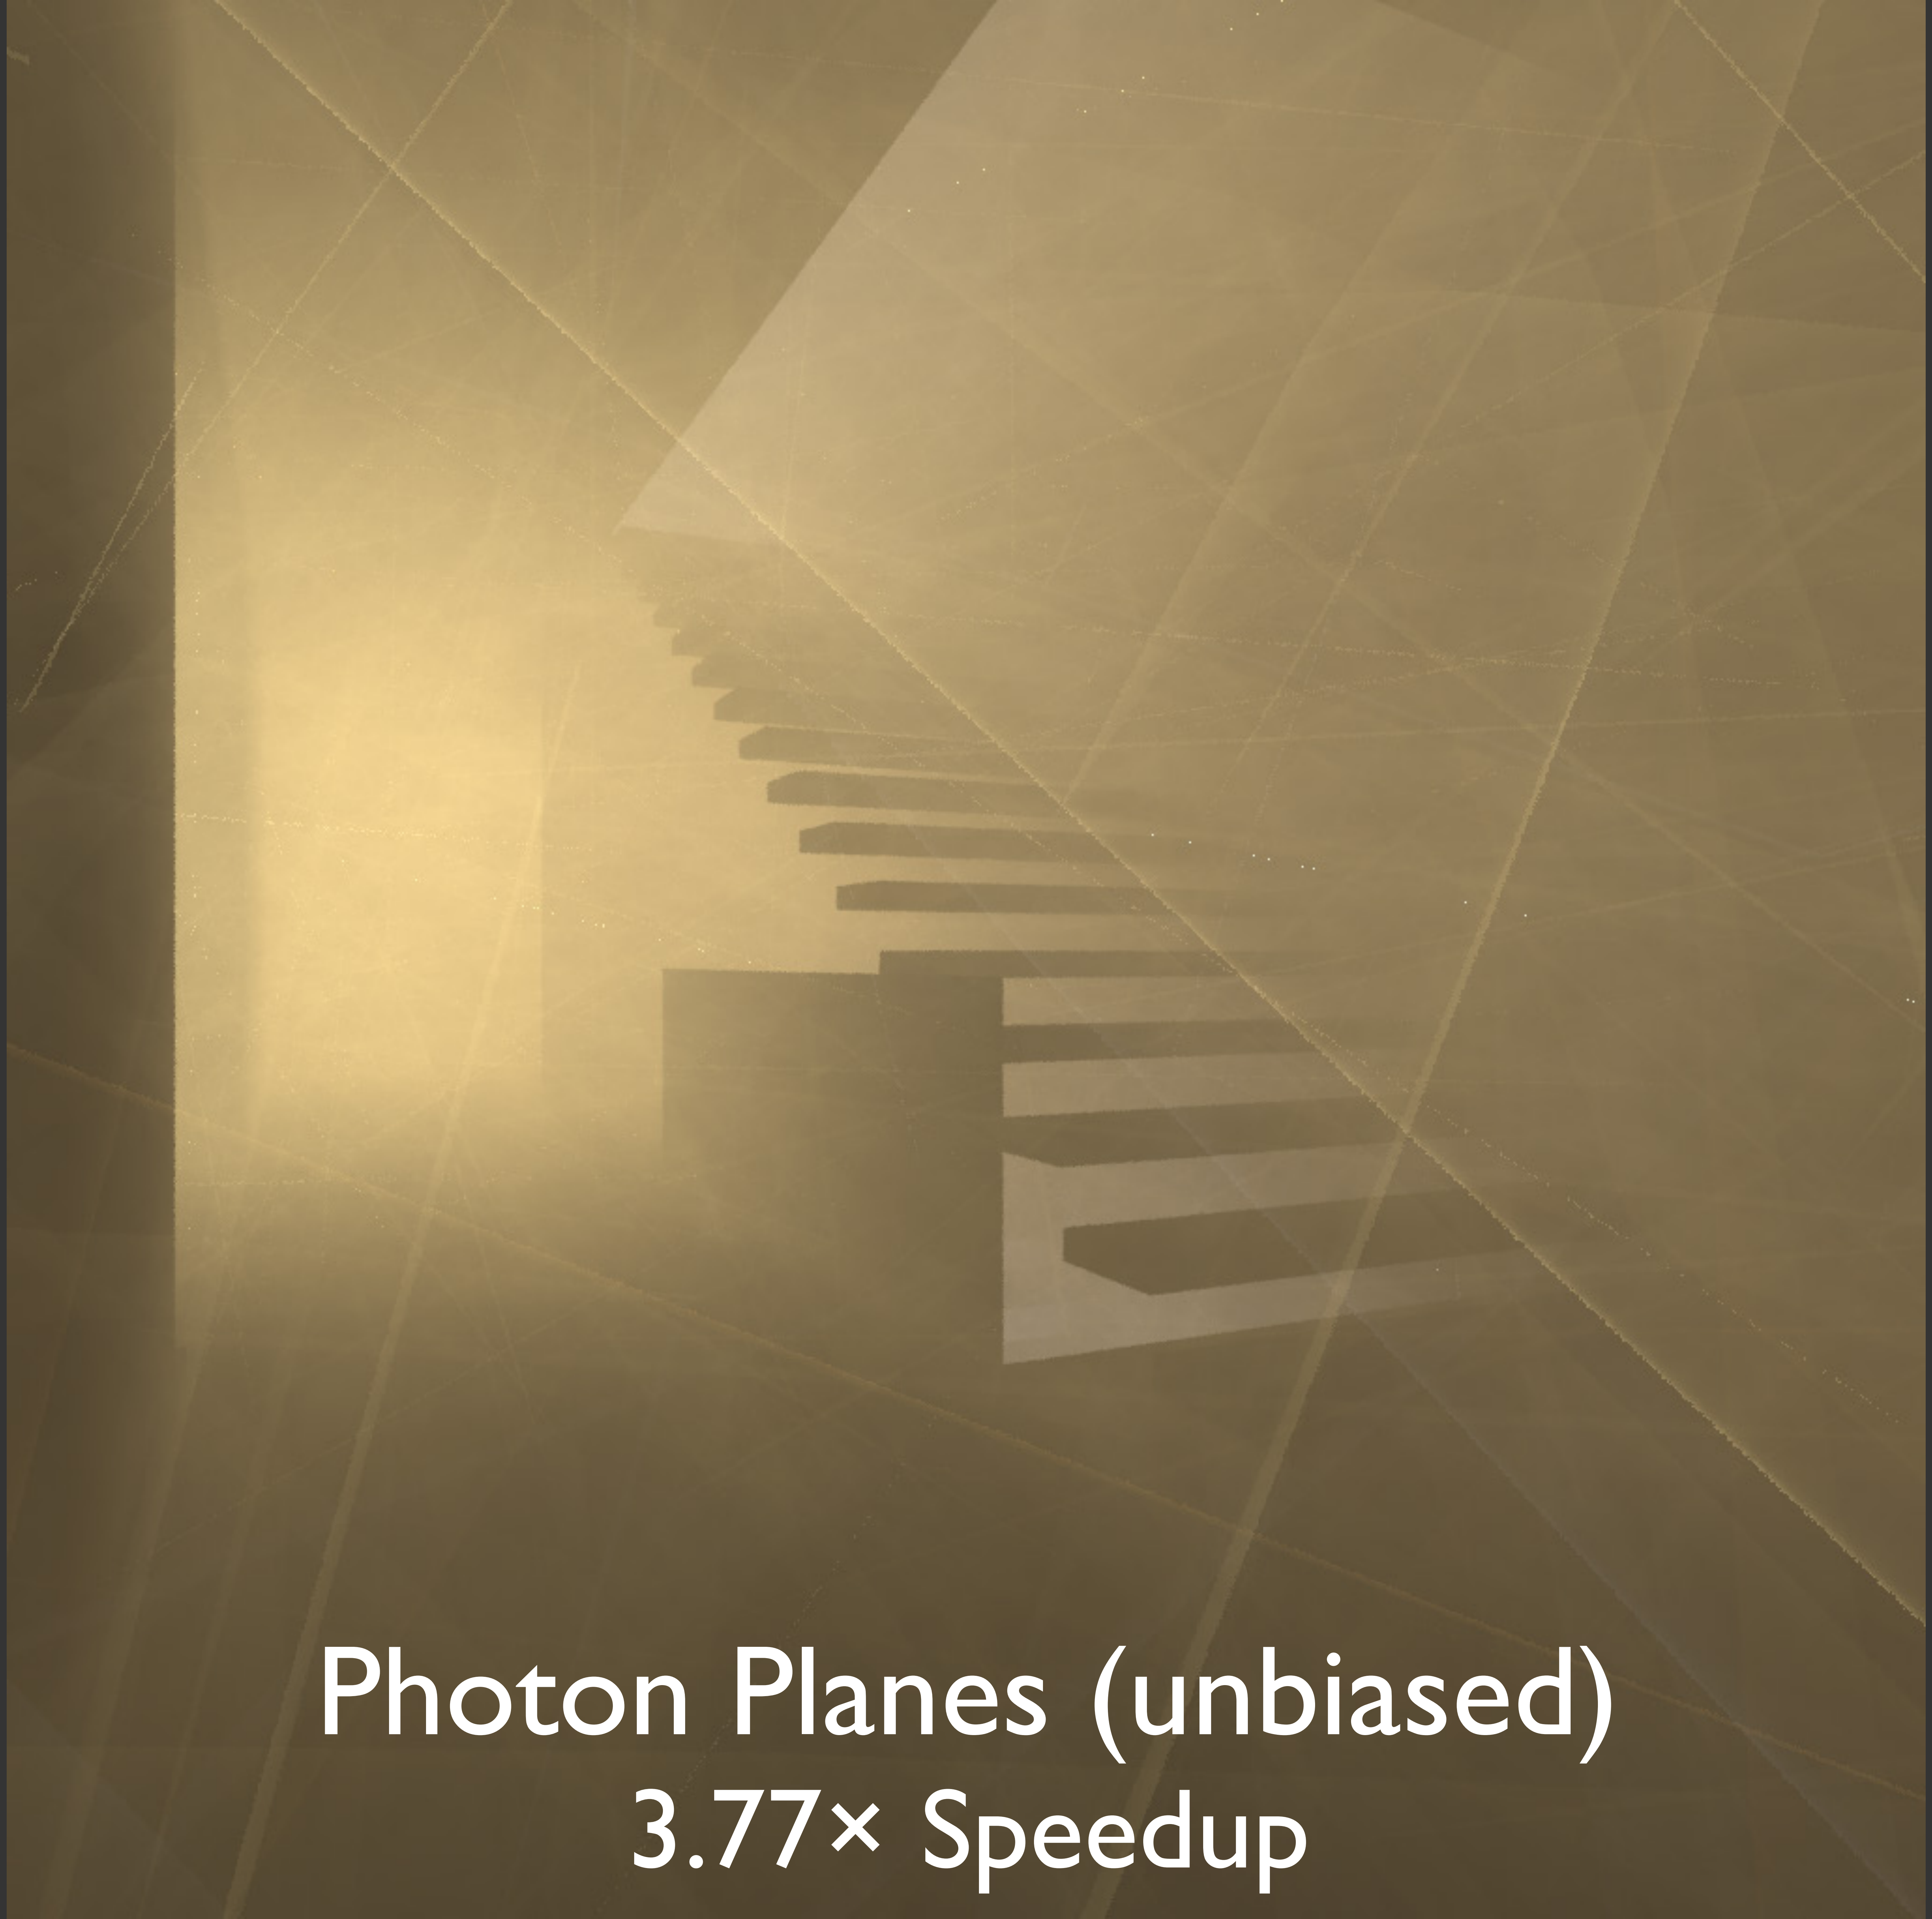
Photon Planes (unbiased)
3.77× Speedup

Photon Beams (1D blur)

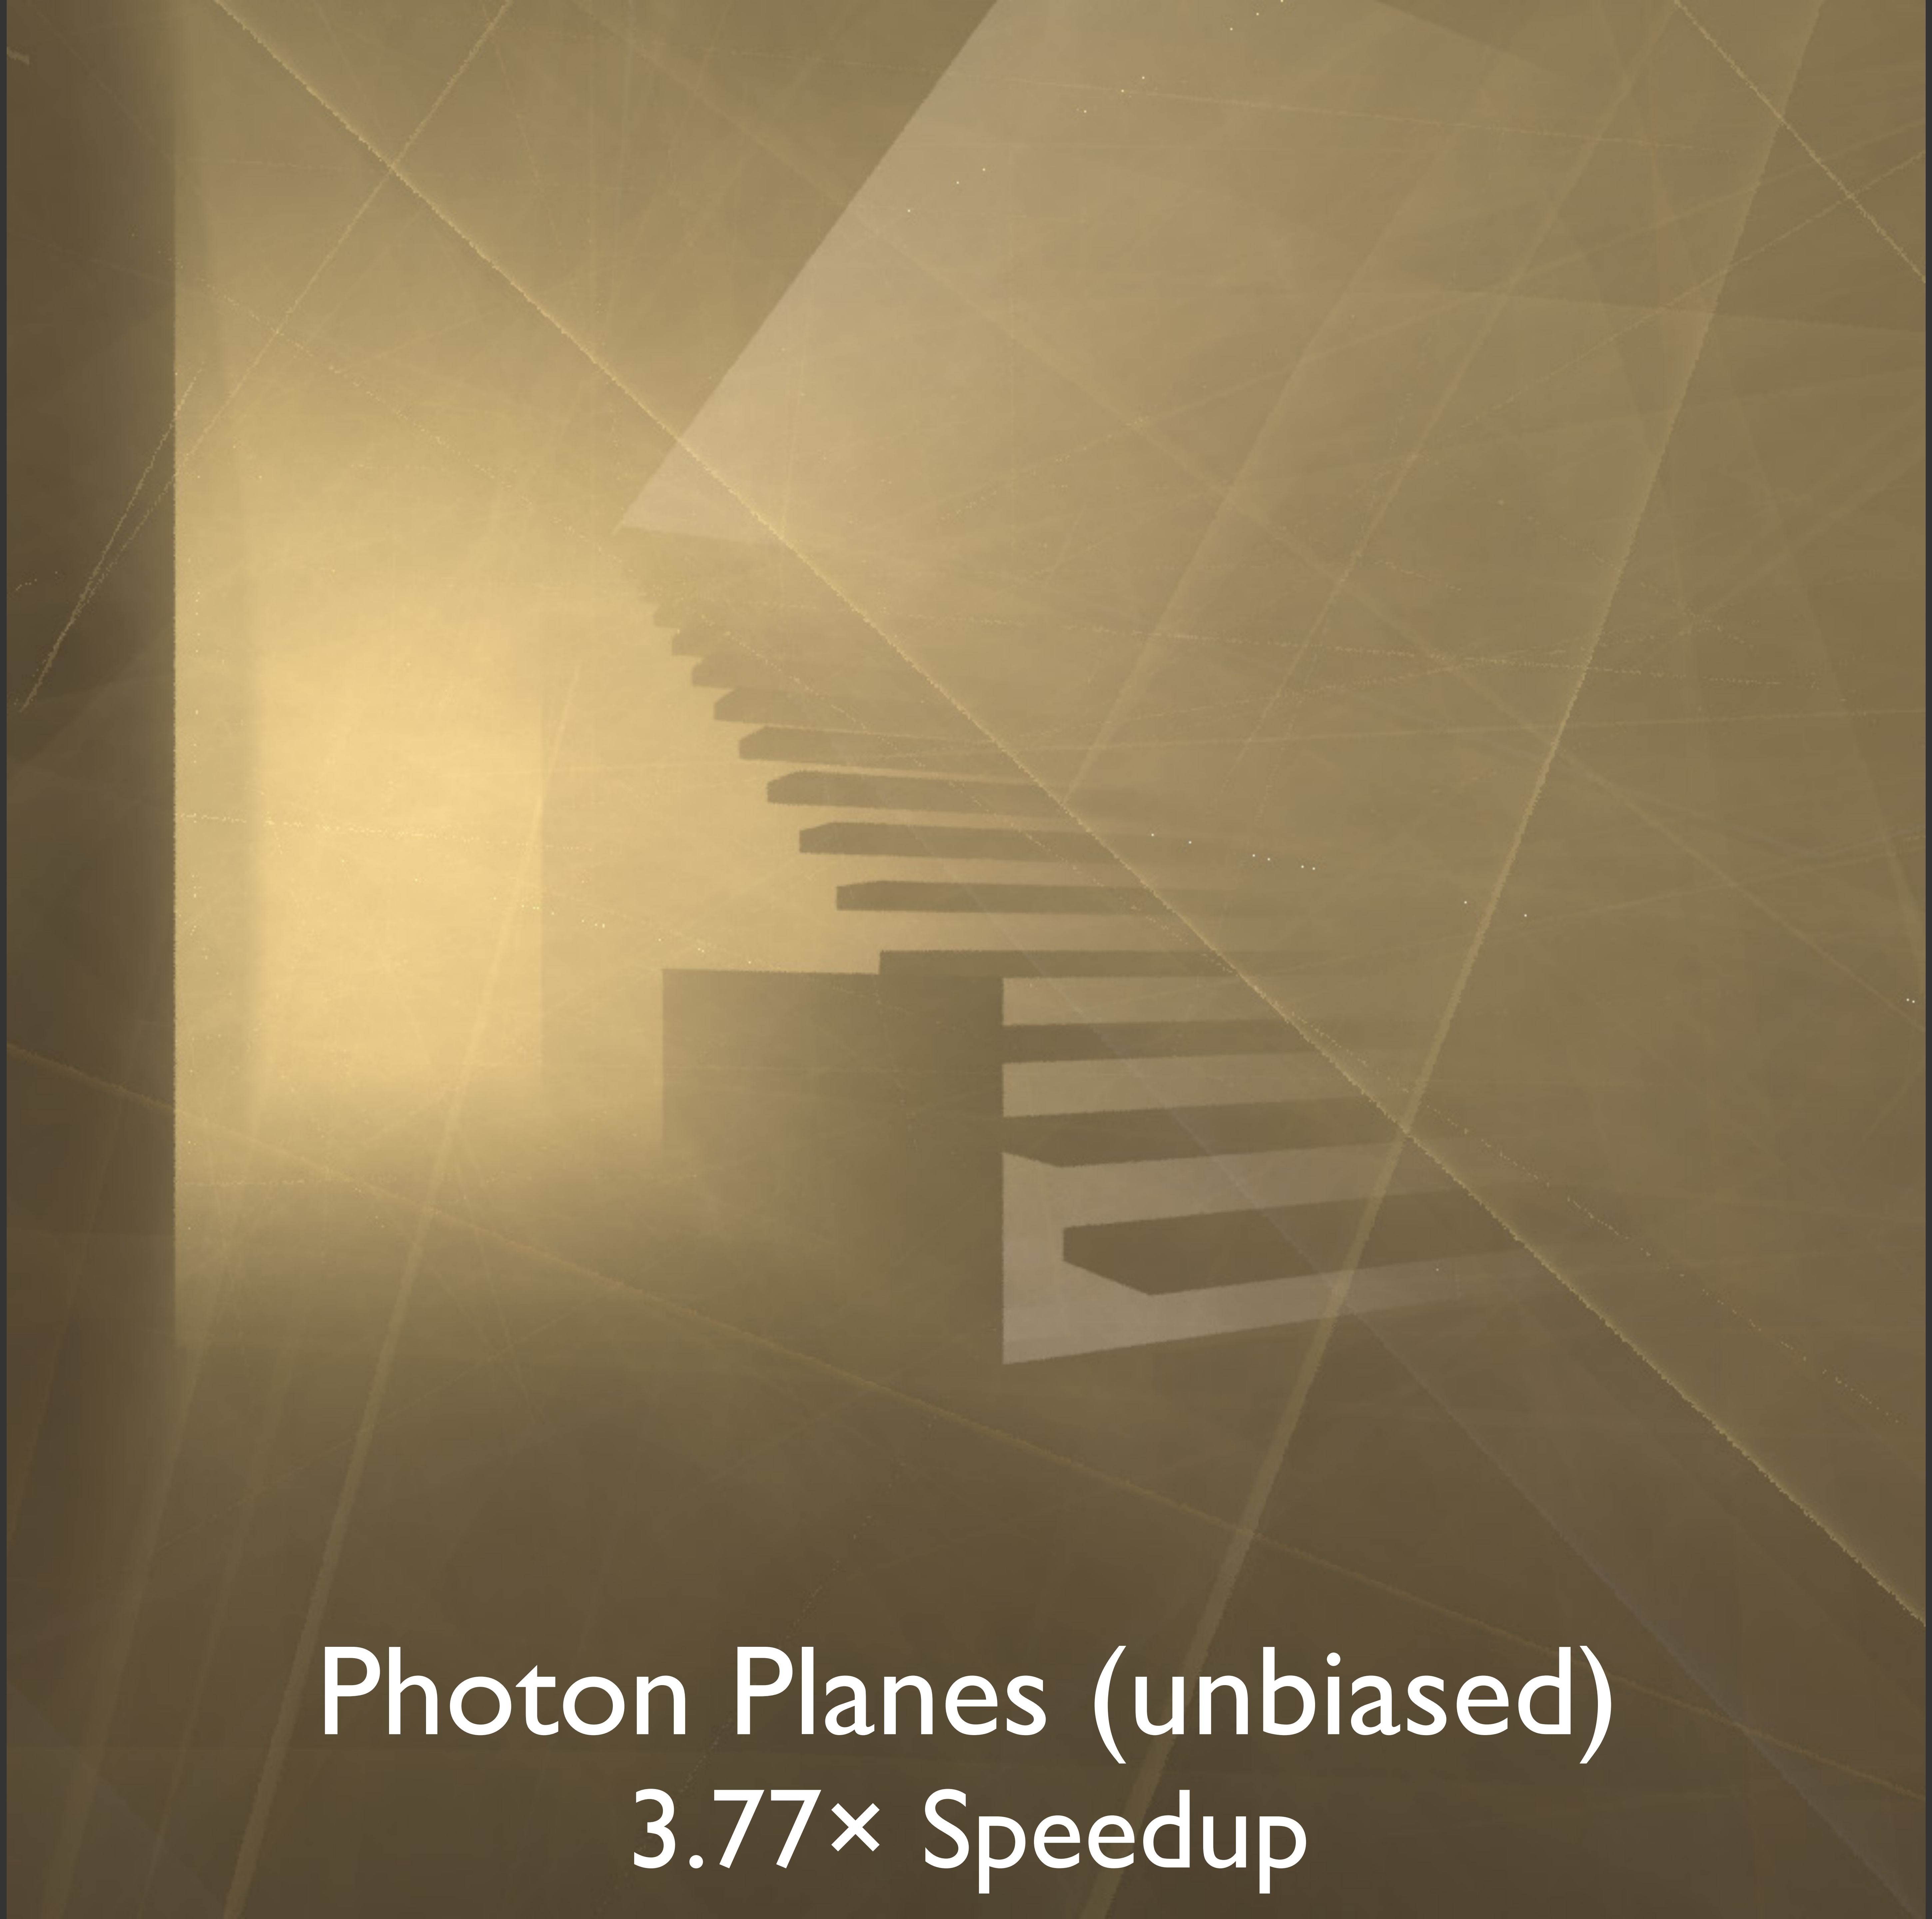

Photon Planes (unbiased)
3.77× Speedup


Photon Beams (1D blur)

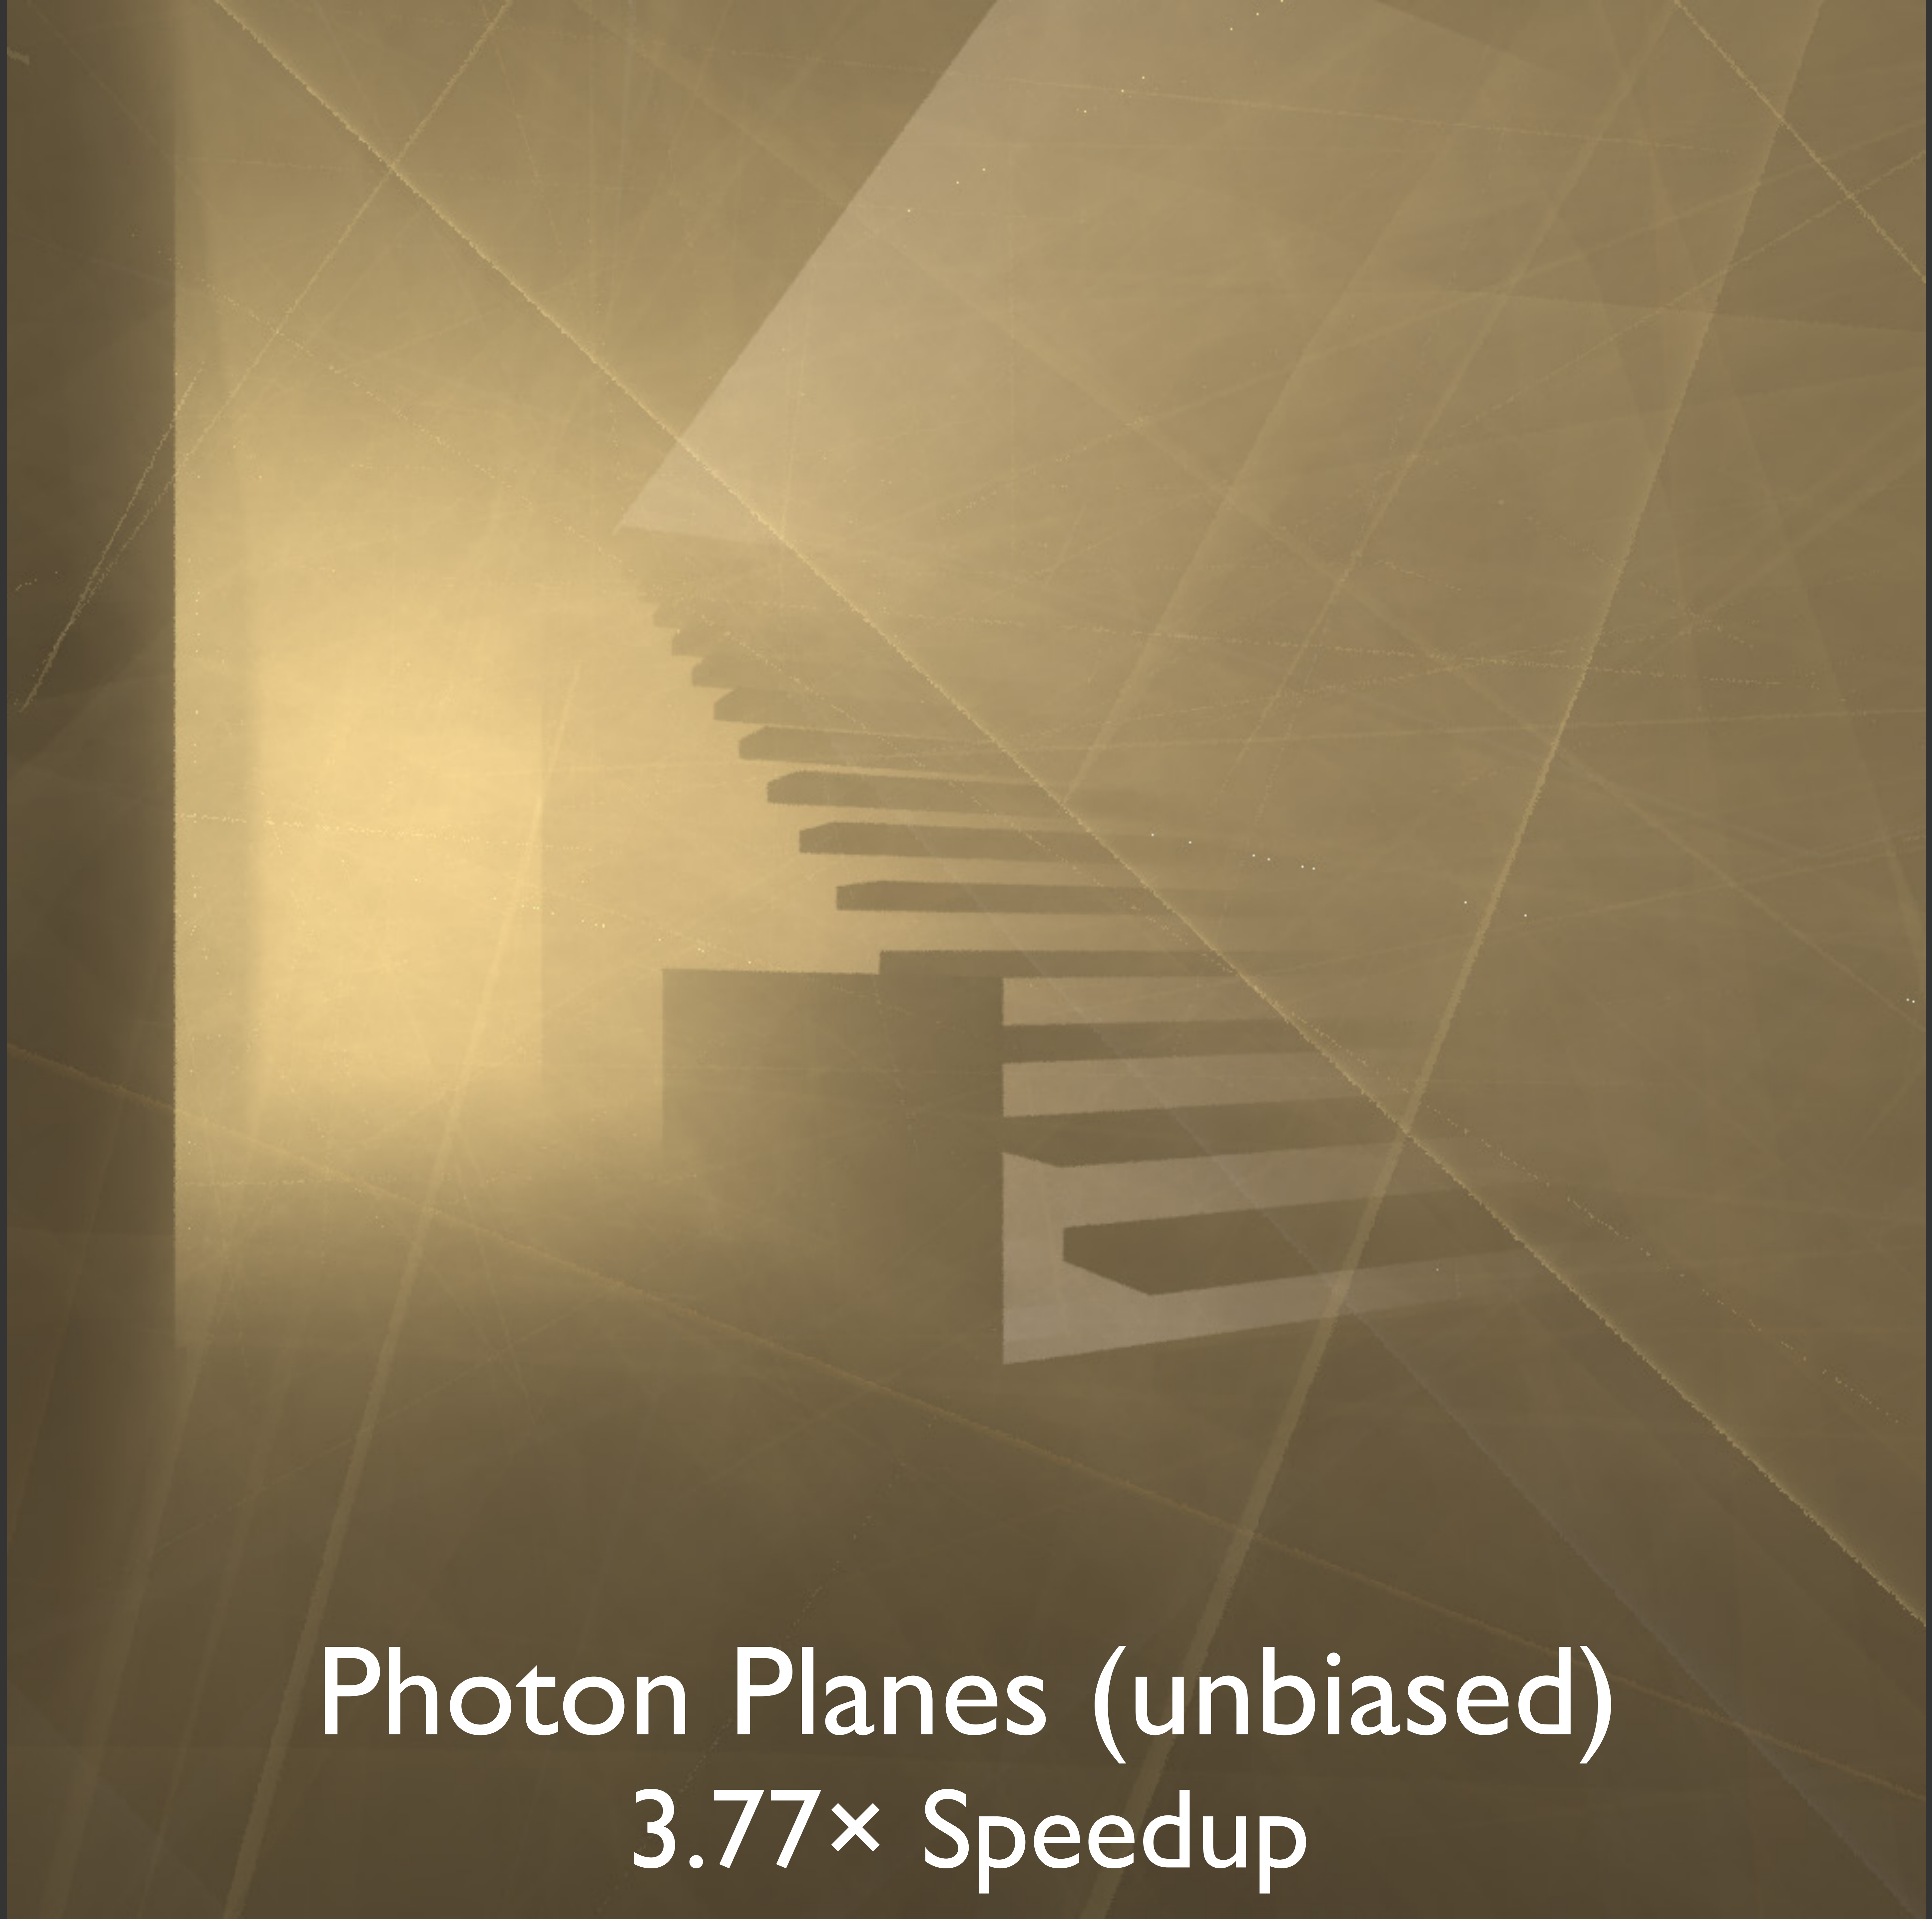

Photon Planes (unbiased)

3.77× Speedup

Photon Planes (1D blur)

14.14× Speedup


Photon Beams (1D blur)

Photon Planes (1D blur)

14.14× Speedup


Photon Beams (1D blur)

Photon Planes (1D blur)
14.14× Speedup

Photon Beams (1D blur)


Photon Planes (1D blur)
14.14× Speedup


Full Light Transport


Photon Beams (1D blur)

Photon Planes (unbiased)

2.85× Speedup


Photon Beams (1D blur)


Photon Planes (unbiased)
2.85× Speedup


Photon Planes (1D blur)
20.70× Speedup

Photon Beams (1D blur)

Photon Planes (1D blur)

20.70× Speedup

Conclusion

Conclusion

- Generalize prior density estimators

Conclusion

- Generalize prior density estimators
- Replace distance sampling with **T/E**

Conclusion

- Generalize prior density estimators
- Replace distance sampling with **T/E**
- Asymptotic error improvement

Conclusion

- Generalize prior density estimators
- Replace distance sampling with **T/E**
- Asymptotic error improvement
- In practice, 2 - 40× variance improvement

Limitations

Limitations

- No heterogeneity for short planes
  - Trivial for long planes

Limitations

- No heterogeneity for short planes
  - Trivial for long planes
- No new estimators for short paths
  - Still need beams for single scattering

Limitations

- No heterogeneity for short planes
  - Trivial for long planes
- No new estimators for short paths
  - Still need beams for single scattering

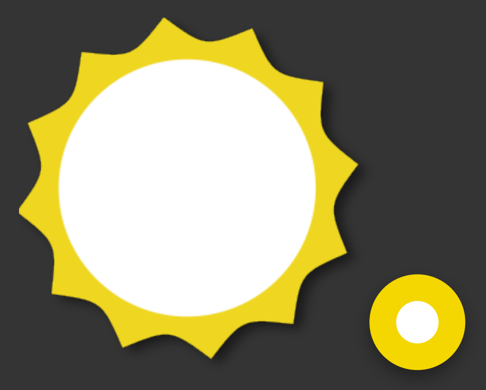

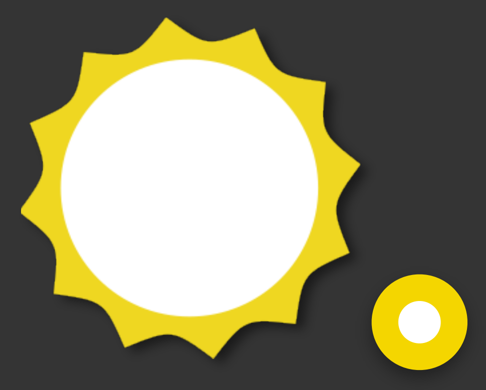

Limitations

- No heterogeneity for short planes
  - Trivial for long planes
- No new estimators for short paths
  - Still need beams for single scattering

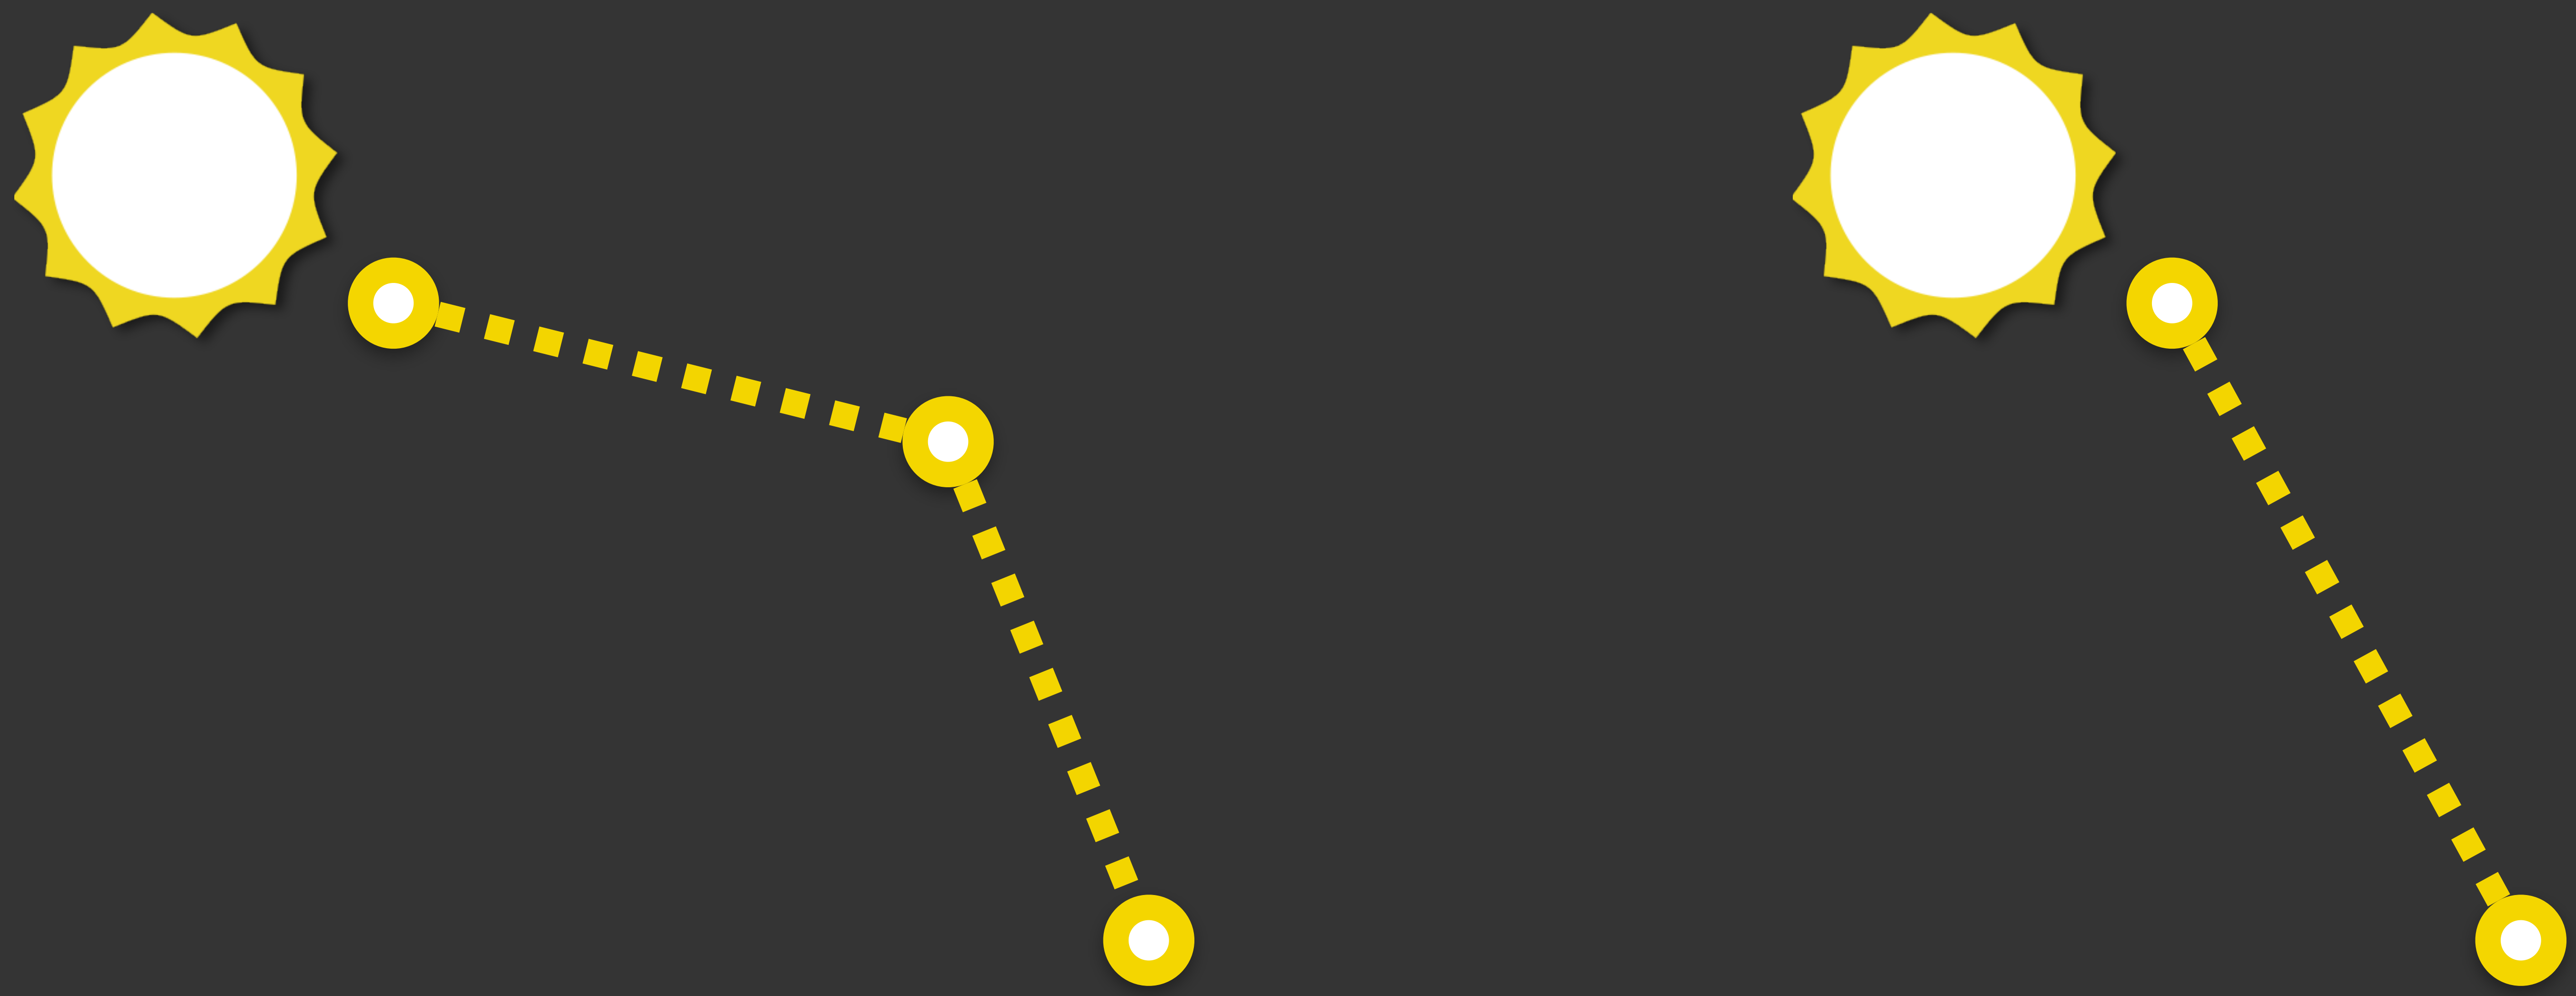

Limitations

- No heterogeneity for short planes
  - Trivial for long planes
- No new estimators for short paths
  - Still need beams for single scattering

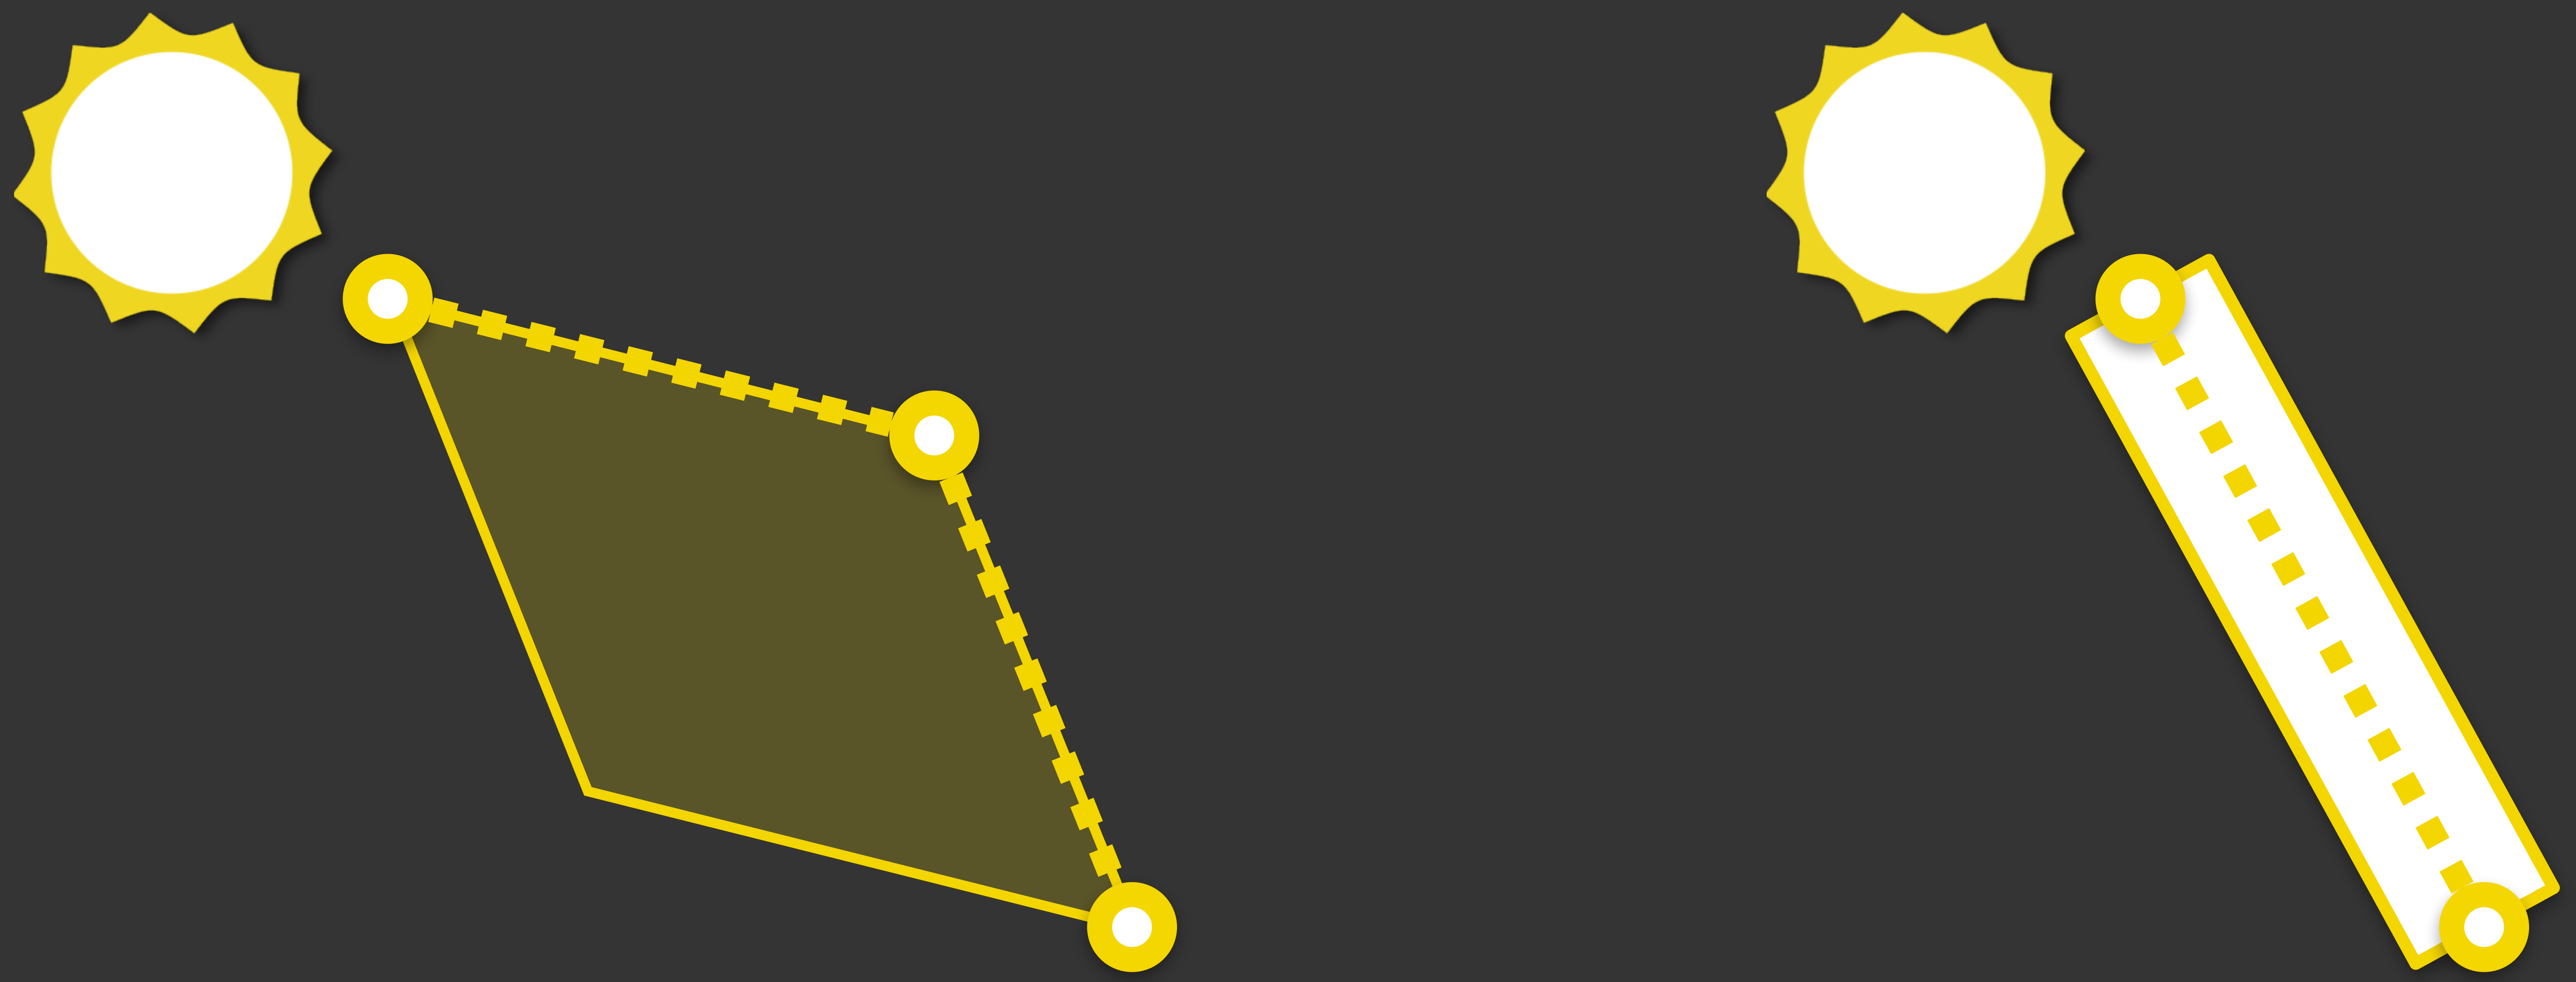

Future Work

Future Work

- New volume photons lead to new surface photons

Future Work

- New volume photons lead to new surface photons

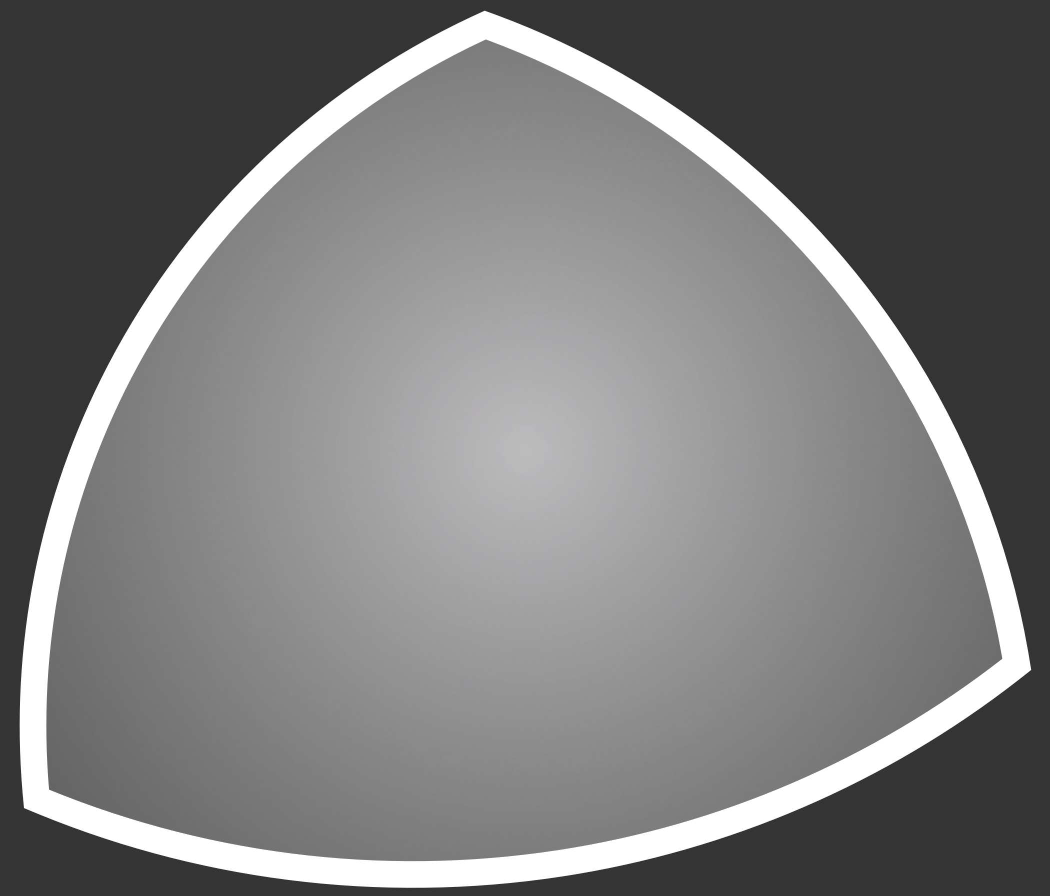

Future Work

- New volume photons lead to new surface photons

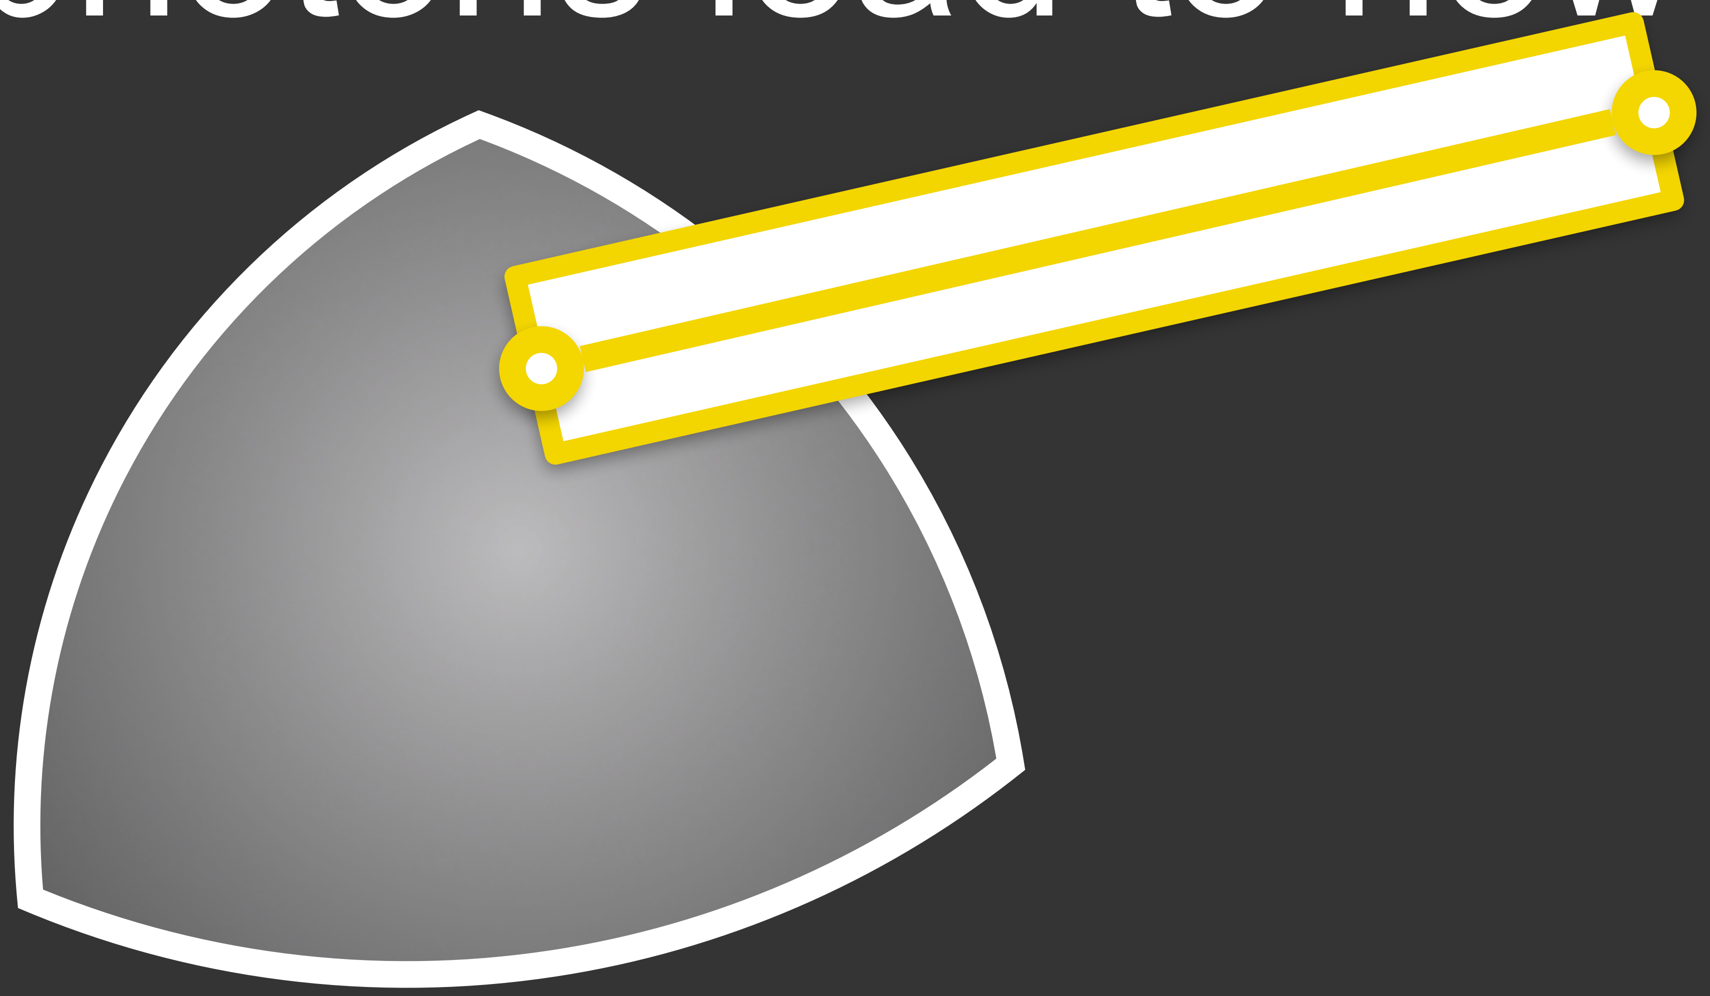

Future Work

- New volume photons lead to new surface photons

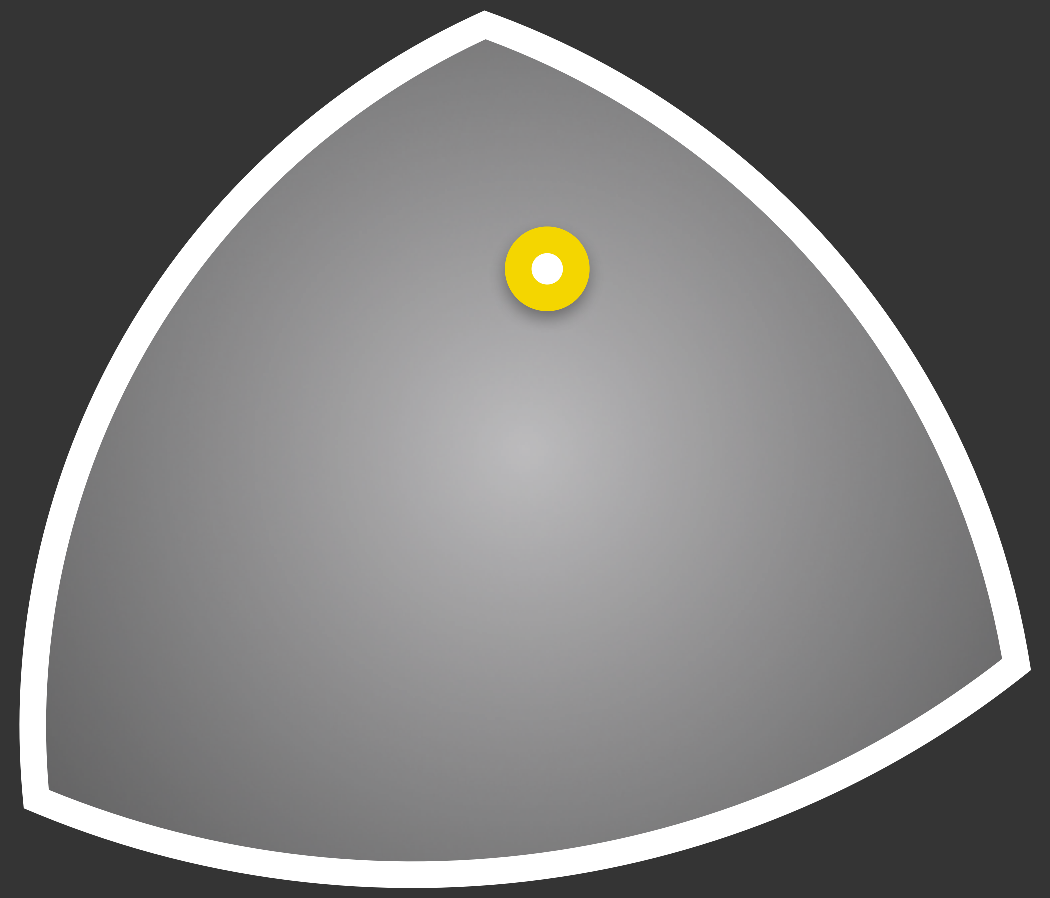

Future Work

- New volume photons lead to new surface photons

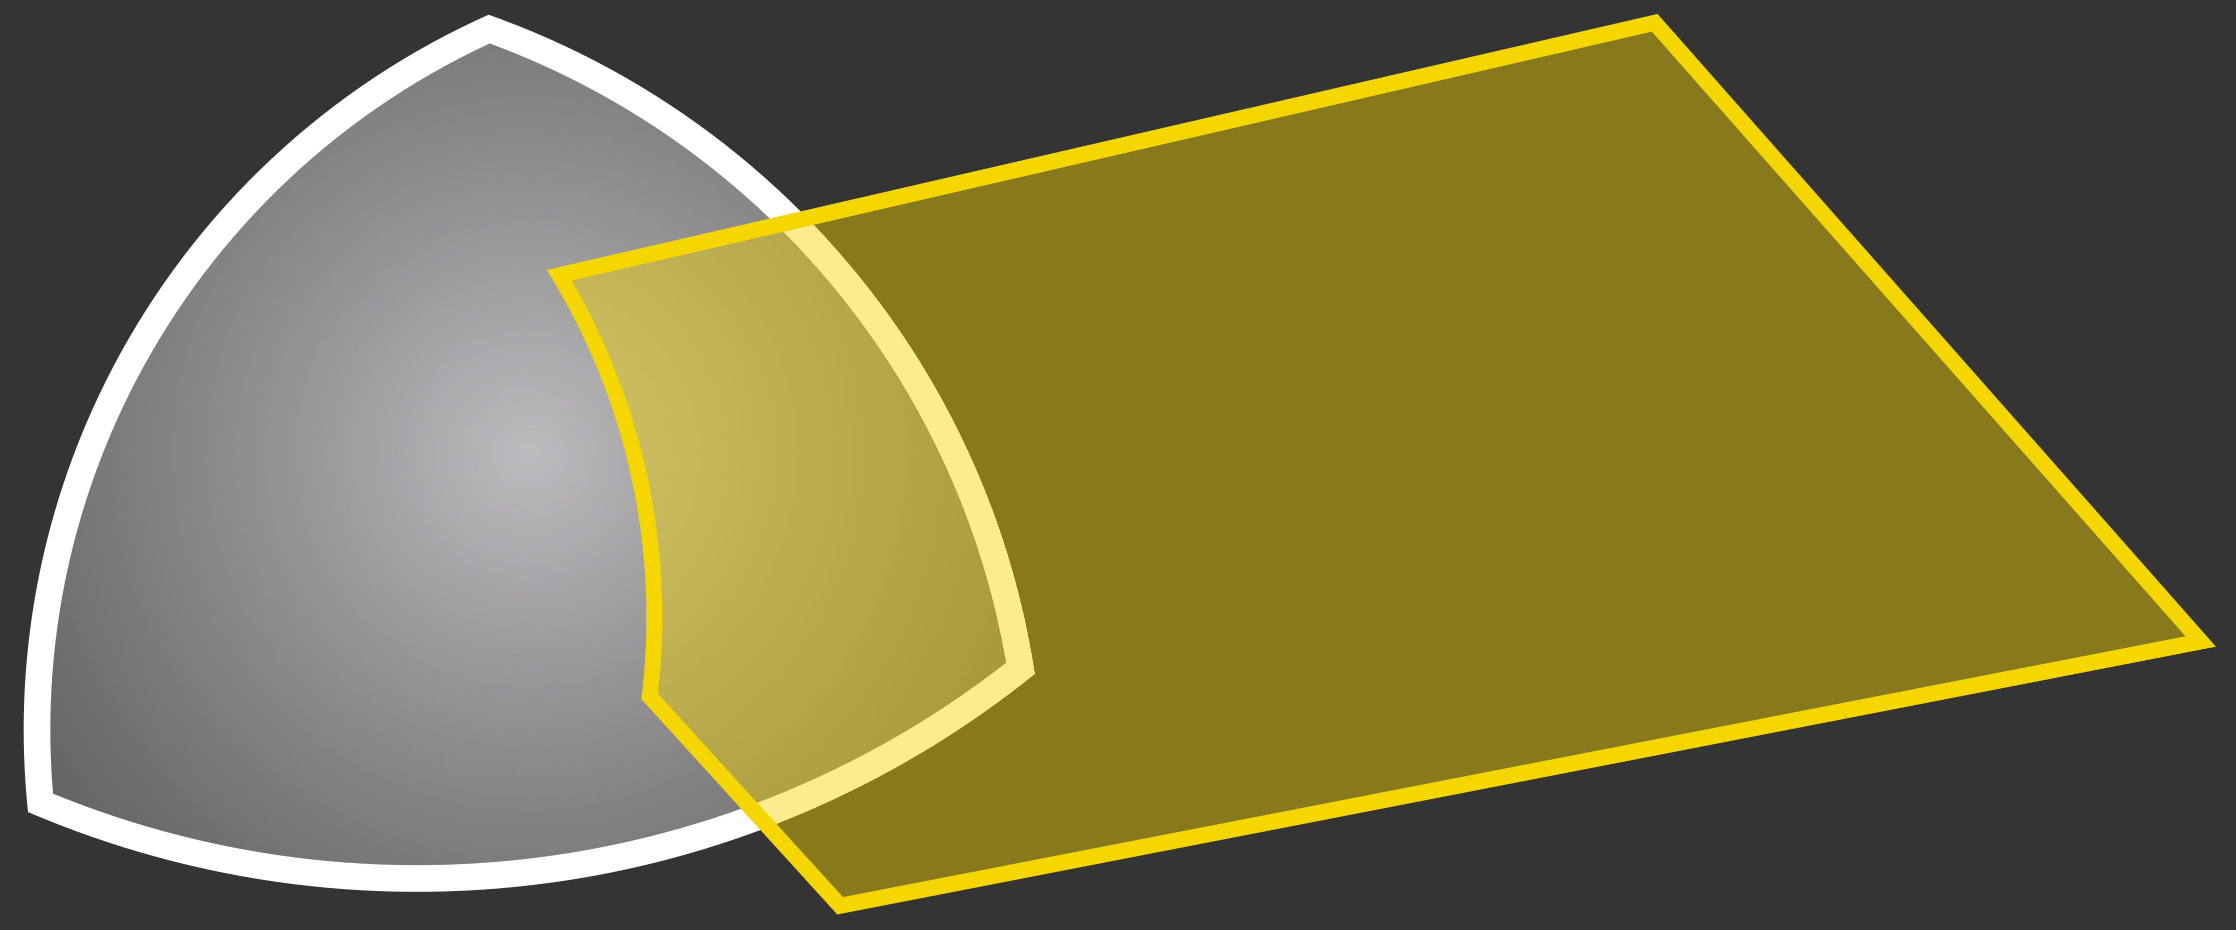

Future Work

- New volume photons lead to new surface photons

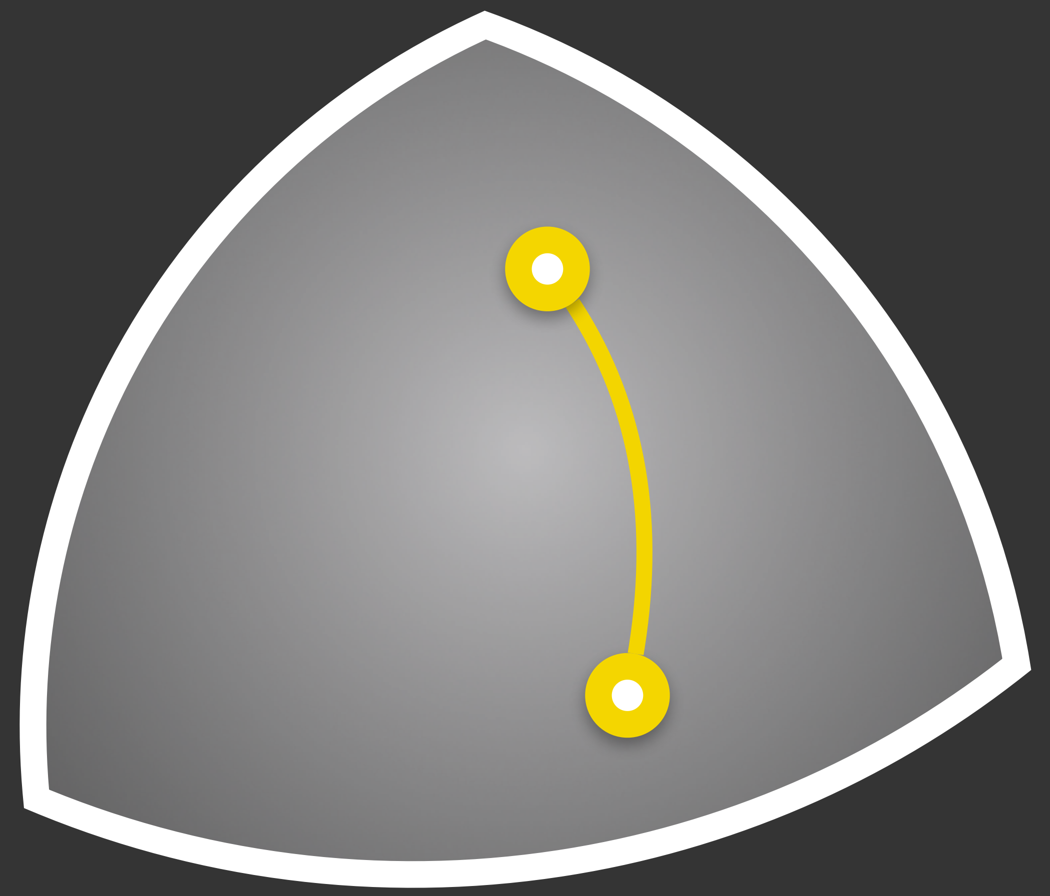

Future Work

- New volume photons lead to new surface photons

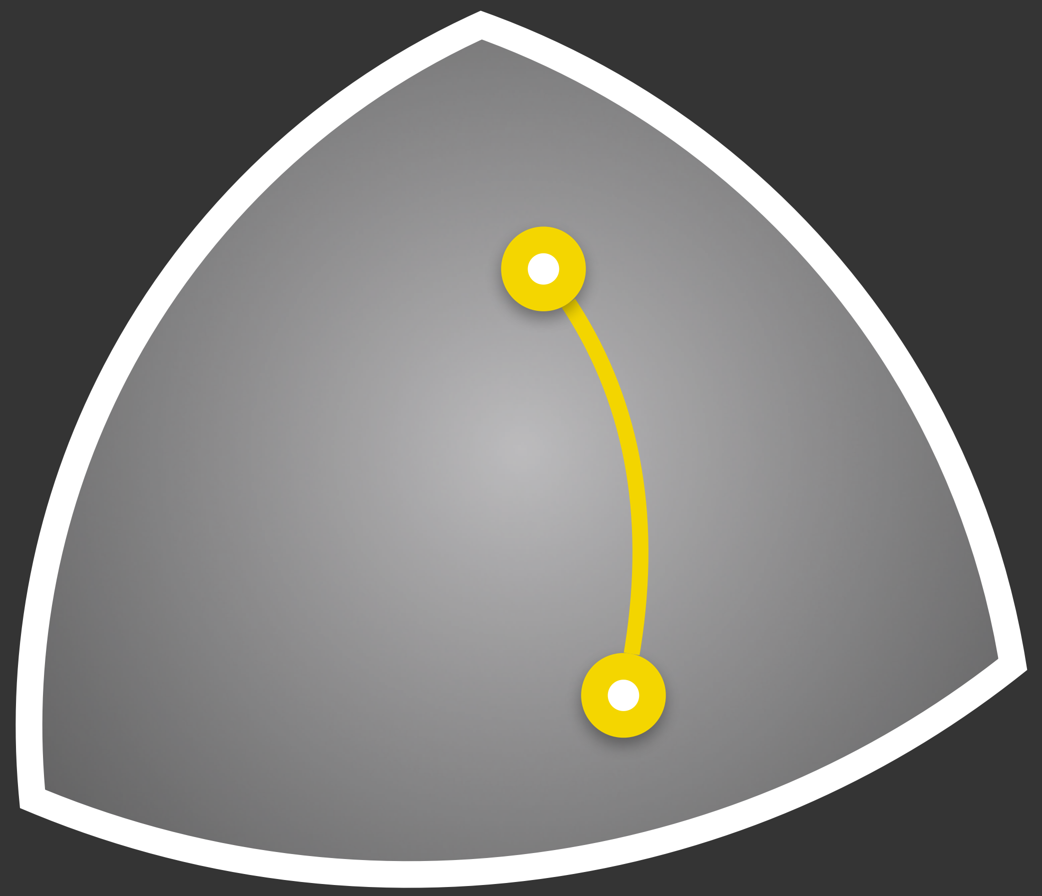

- What about phase functions?
  - “Photon spinning”: Photon rings, cones, cylinders…

Thanks!

- Try our WebGL Demo!

benedikt-bitterli.me/photon-planes