Residual Ratio Tracking

for estimating attenuation in
heterogeneous participating media

Jan Novák **Andrew Selle** **Wojciech Jarosz**

Walt Disney Animation Studios

Disney Research

Motivation

The motivation for developing the techniques that we present today was the need to render scenes with heterogeneous participating media, such as these. Our goal was to synthesize those with a production path tracer, where each path sample needs to be relatively cheap, but preferably unbiased, so that we can rely on averaging of samples.

Motivation

SIGGRAPH ASIA 2014 SHENZHEN

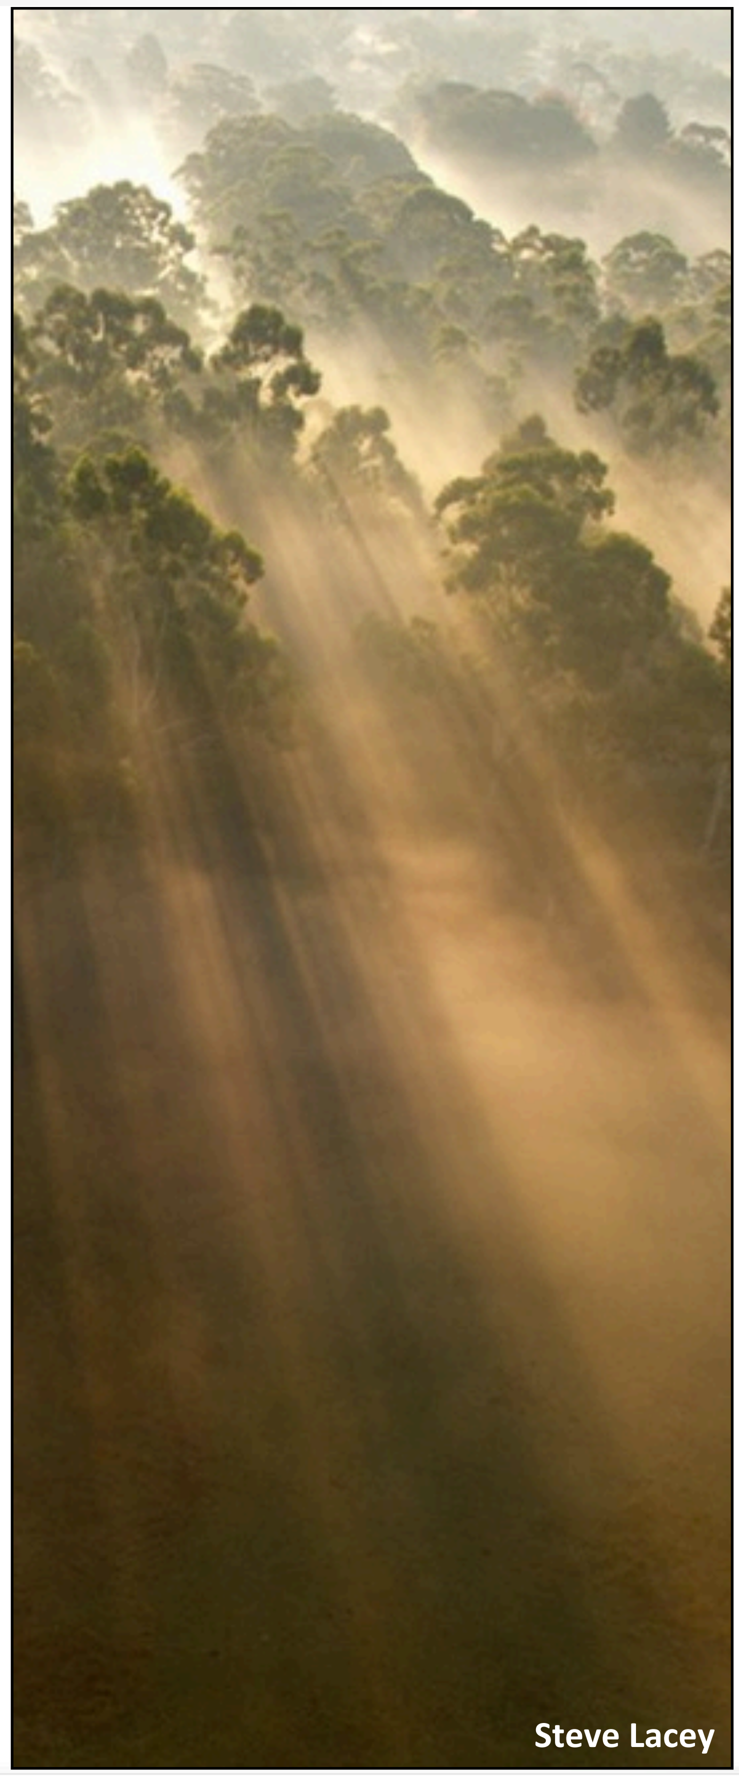

The motivation for developing the techniques that we present today was the need to render scenes with heterogeneous participating media, such as these. Our goal was to synthesize those with a production path tracer, where each path sample needs to be relatively cheap, but preferably unbiased, so that we can rely on averaging of samples.

Light Transport in Participating Media

Media that we are interested in consist of particles that interact with light. In graphics, we rarely model the particles explicitly. Instead, we describe the medium by coefficients that characterize the absorption and scattering of light. Adding these together, yields the extinction coefficient that represents how much light interacts with the medium in general per unit flight distance.

Light Transport in Participating Media

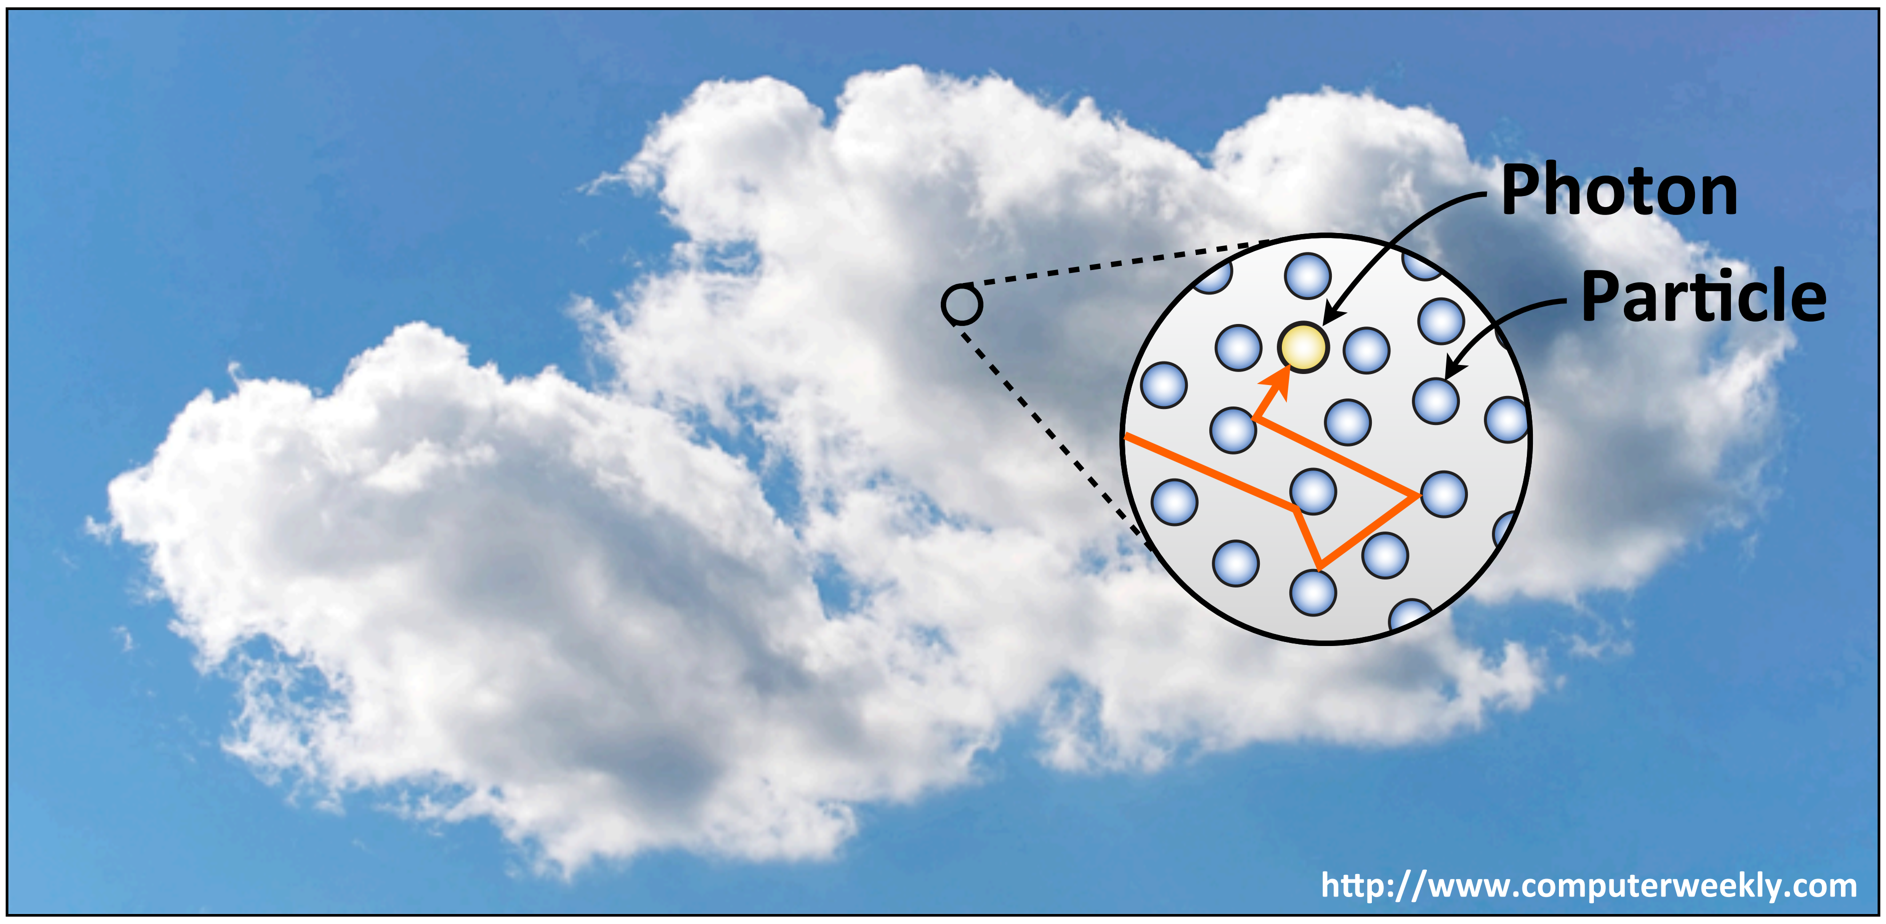

Media that we are interested in consist of particles that interact with light. In graphics, we rarely model the particles explicitly. Instead, we describe the medium by coefficients that characterize the absorption and scattering of light. Adding these together, yields the extinction coefficient that represents how much light interacts with the medium in general per unit flight distance.

Light Transport in Participating Media

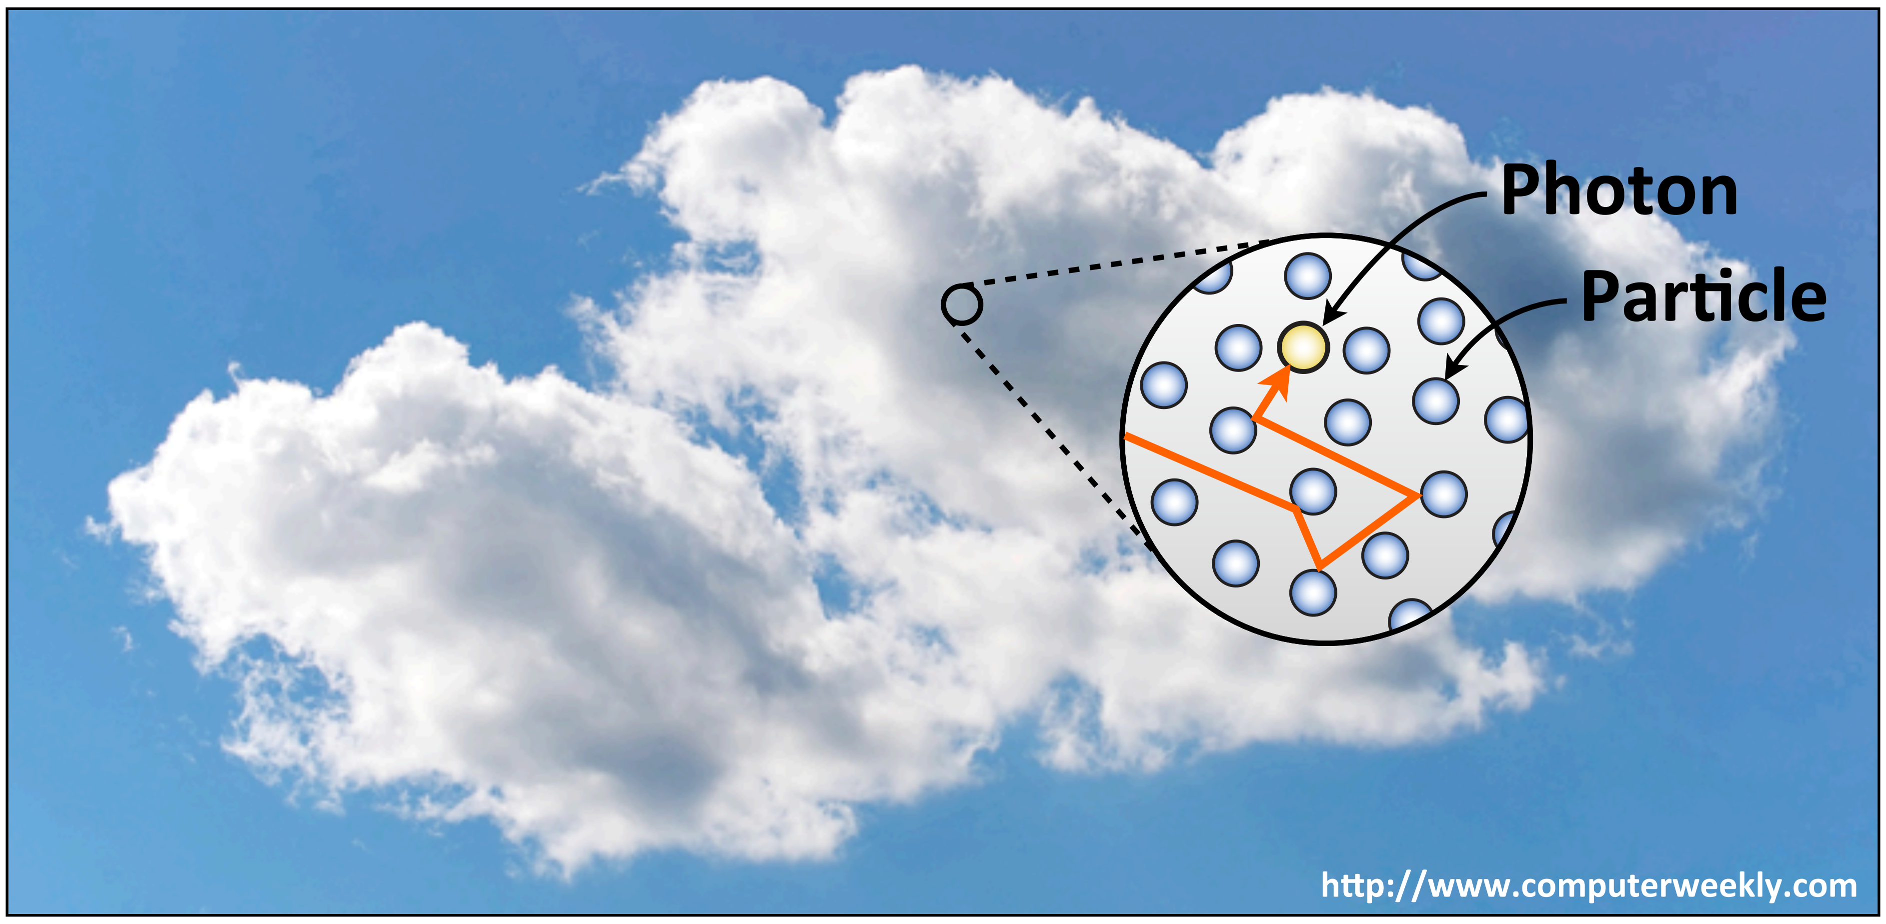

$$\mu_a(x)$$

Absorption coefficient

Media that we are interested in consist of particles that interact with light. In graphics, we rarely model the particles explicitly. Instead, we describe the medium by coefficients that characterize the absorption and scattering of light. Adding these together, yields the extinction coefficient that represents how much light interacts with the medium in general per unit flight distance.

Light Transport in Participating Media

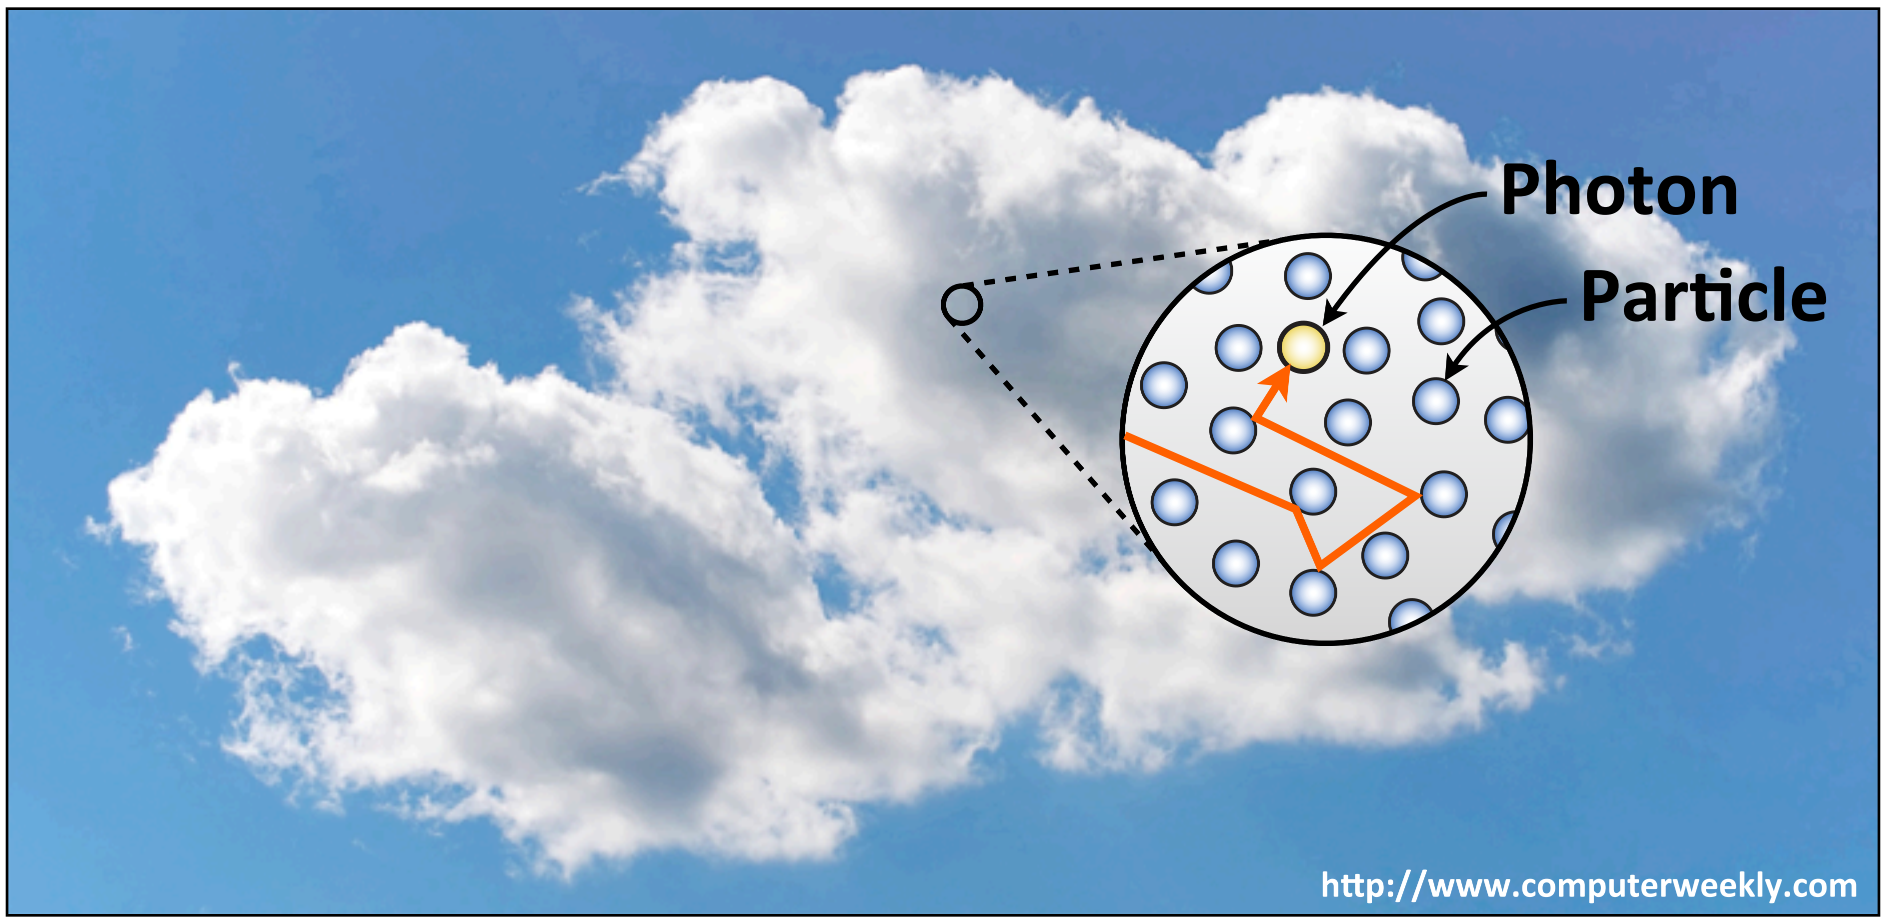

$\mu_a(x)$ $\mu_s(x)$

Absorption coefficient **Scattering coefficient**

Media that we are interested in consist of particles that interact with light. In graphics, we rarely model the particles explicitly. Instead, we describe the medium by coefficients that characterize the absorption and scattering of light. Adding these together, yields the extinction coefficient that represents how much light interacts with the medium in general per unit flight distance.

Light Transport in Participating Media

SIGGRAPH ASIA 2014 SHENZHEN

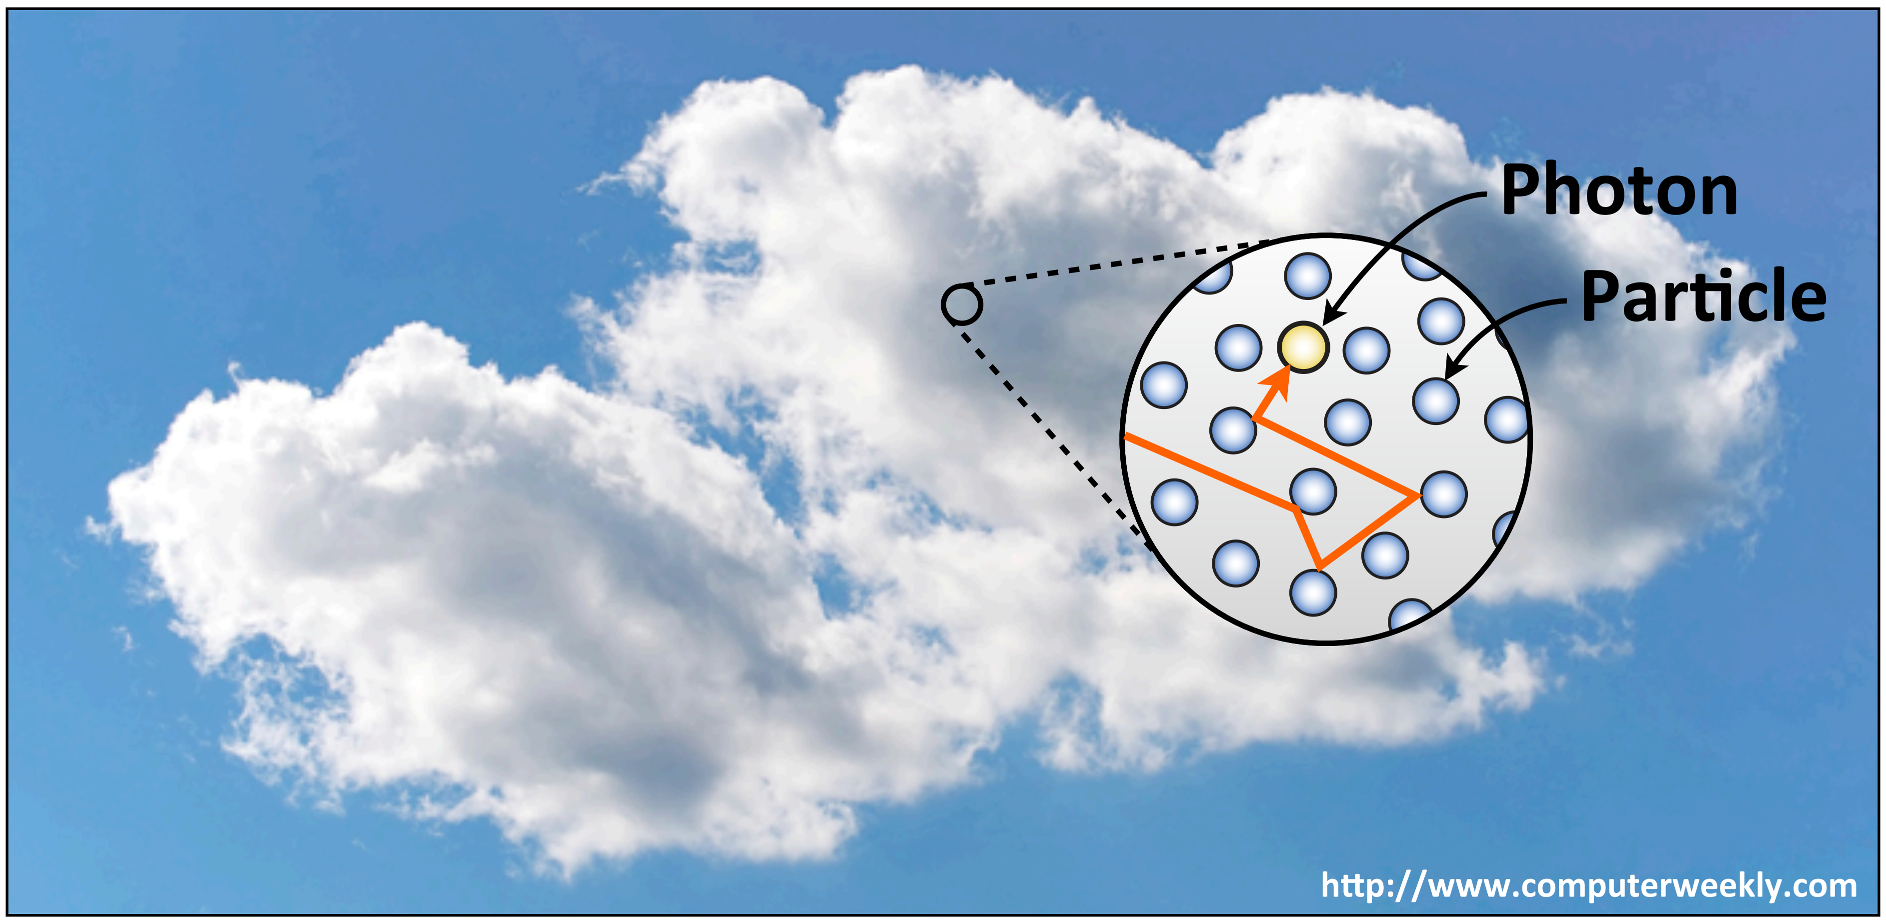

$$\mu(x) = \mu_a(x) + \mu_s(x)$$

Extinction coefficient **Absorption coefficient** **Scattering coefficient**

"How much light interacts with matter (per unit flight distance)."

Media that we are interested in consist of particles that interact with light. In graphics, we rarely model the particles explicitly. Instead, we describe the medium by coefficients that characterize the absorption and scattering of light. Adding these together, yields the extinction coefficient that represents how much light interacts with the medium in general per unit flight distance.

Free-path sampling

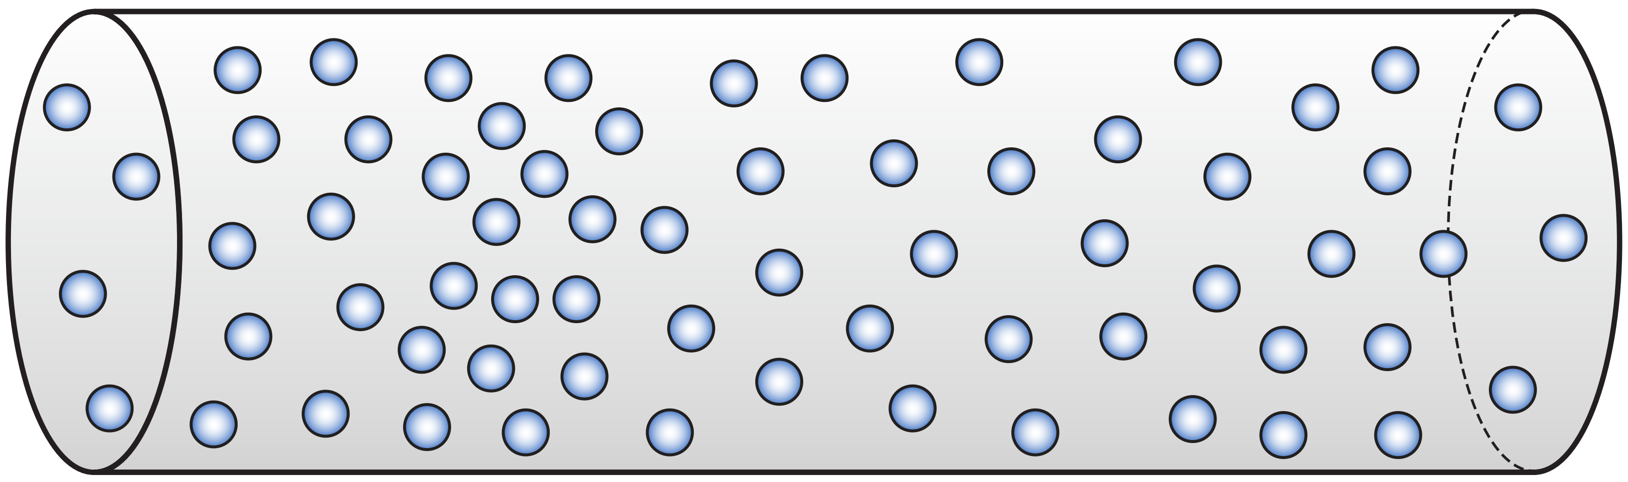

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Free-path sampling

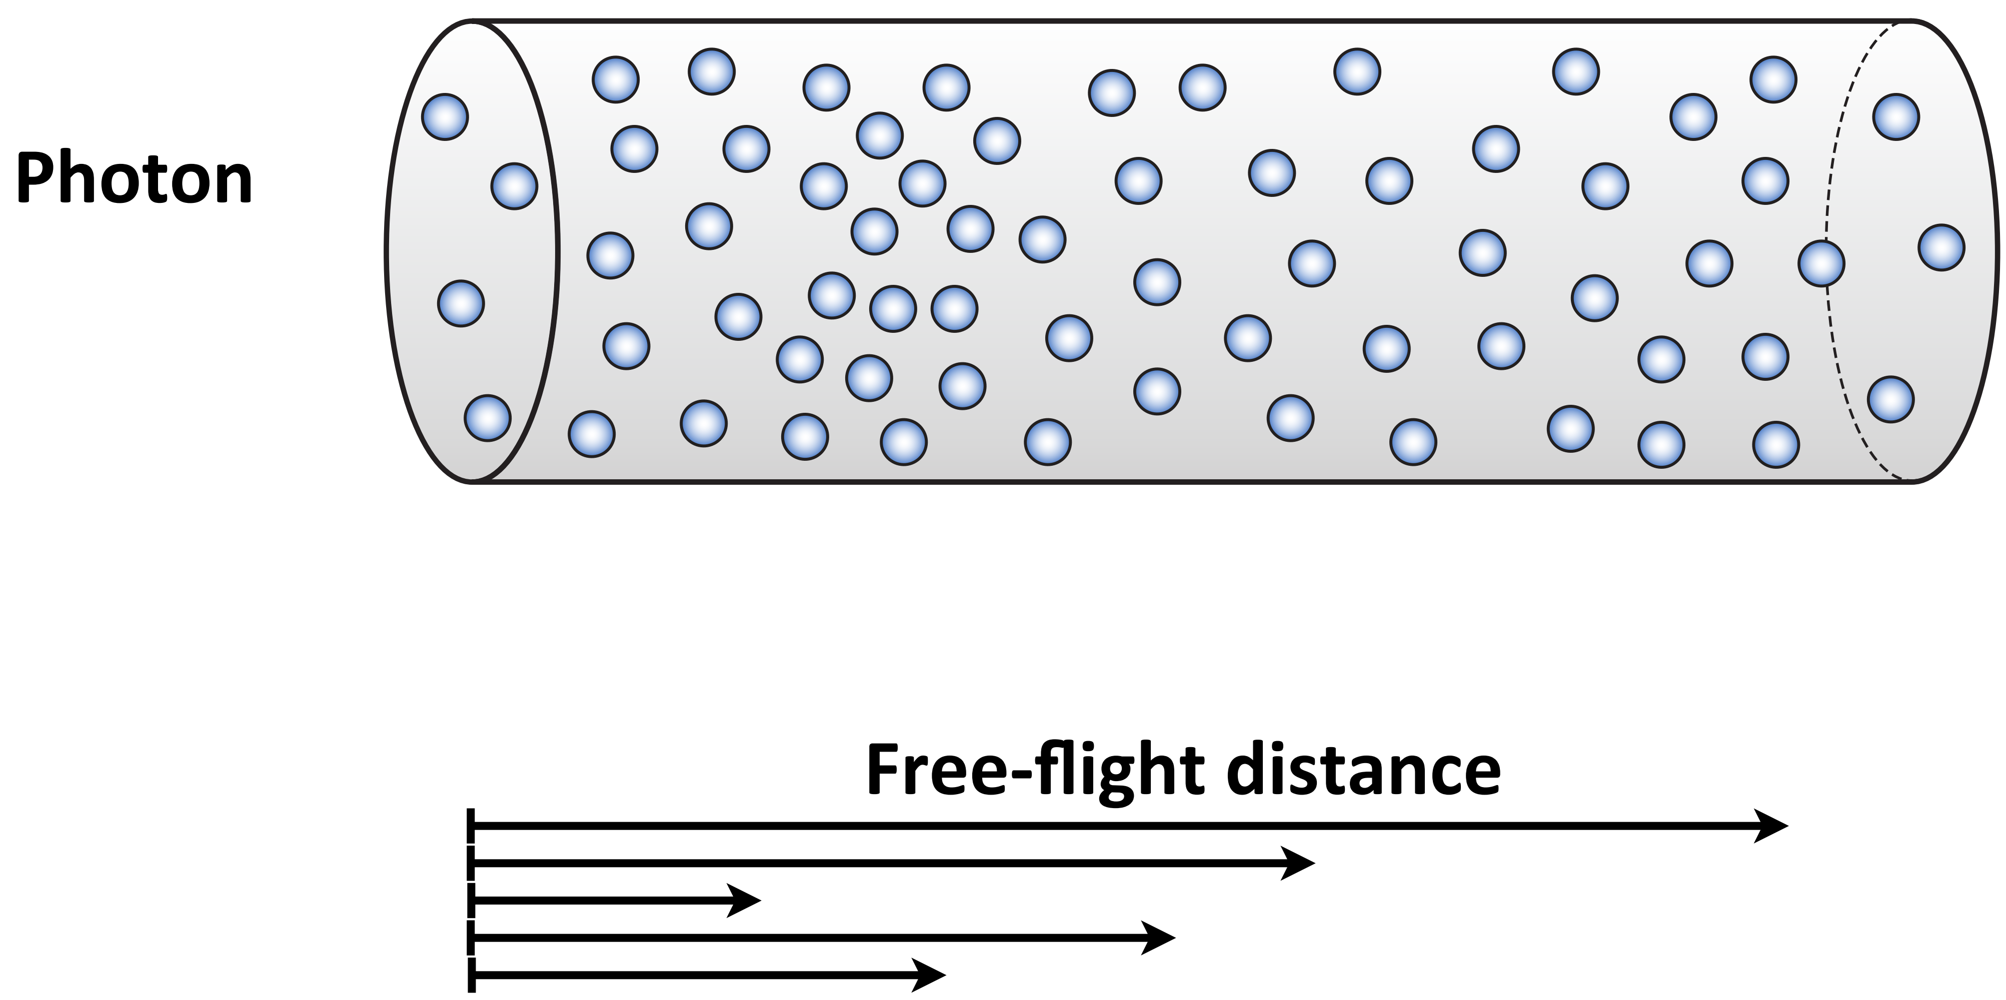

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Attenuation (Transmittance)

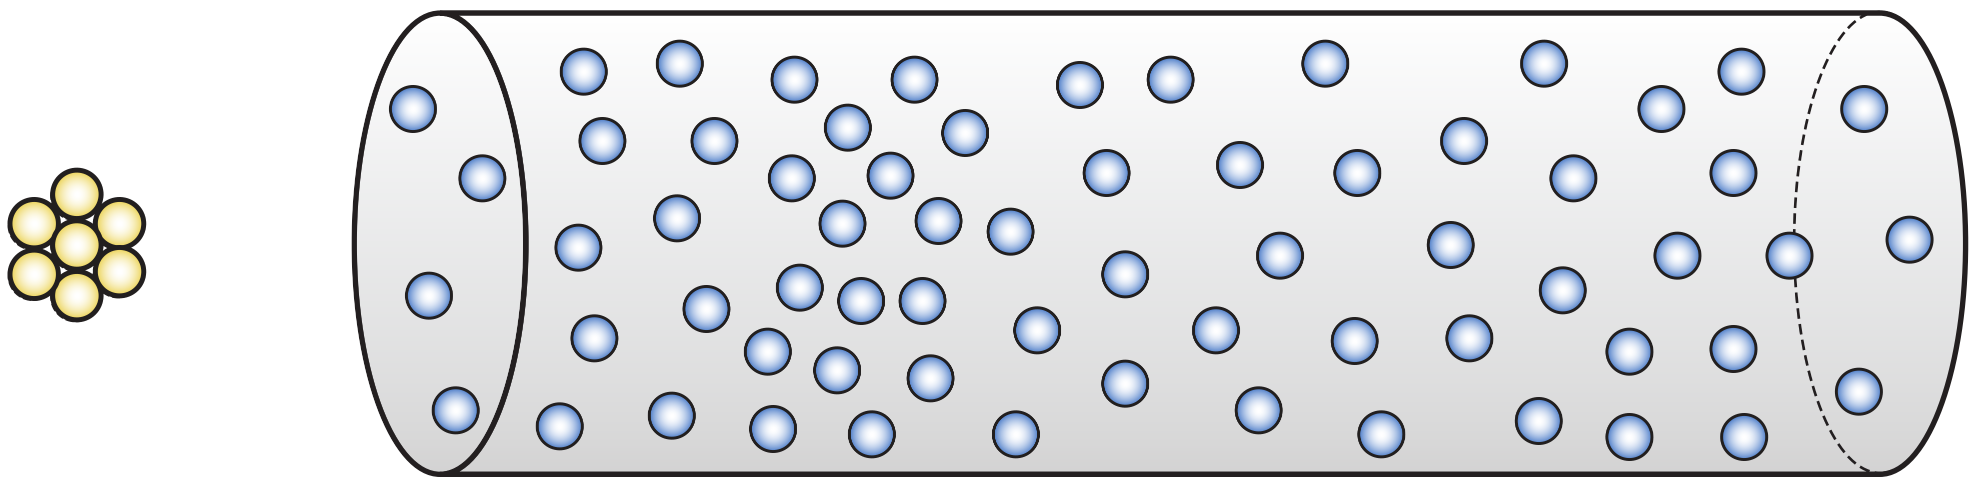

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Attenuation (Transmittance)

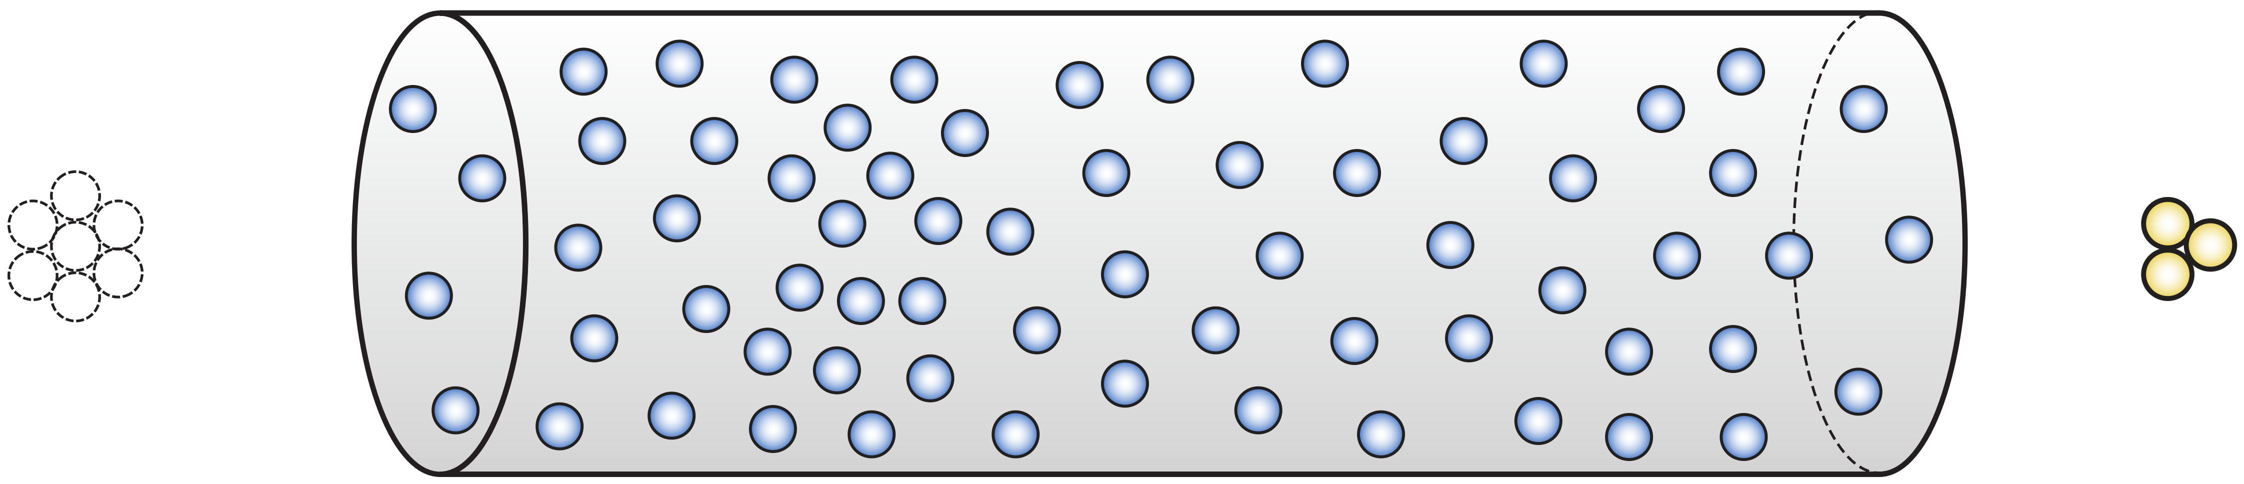

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Attenuation (Transmittance)

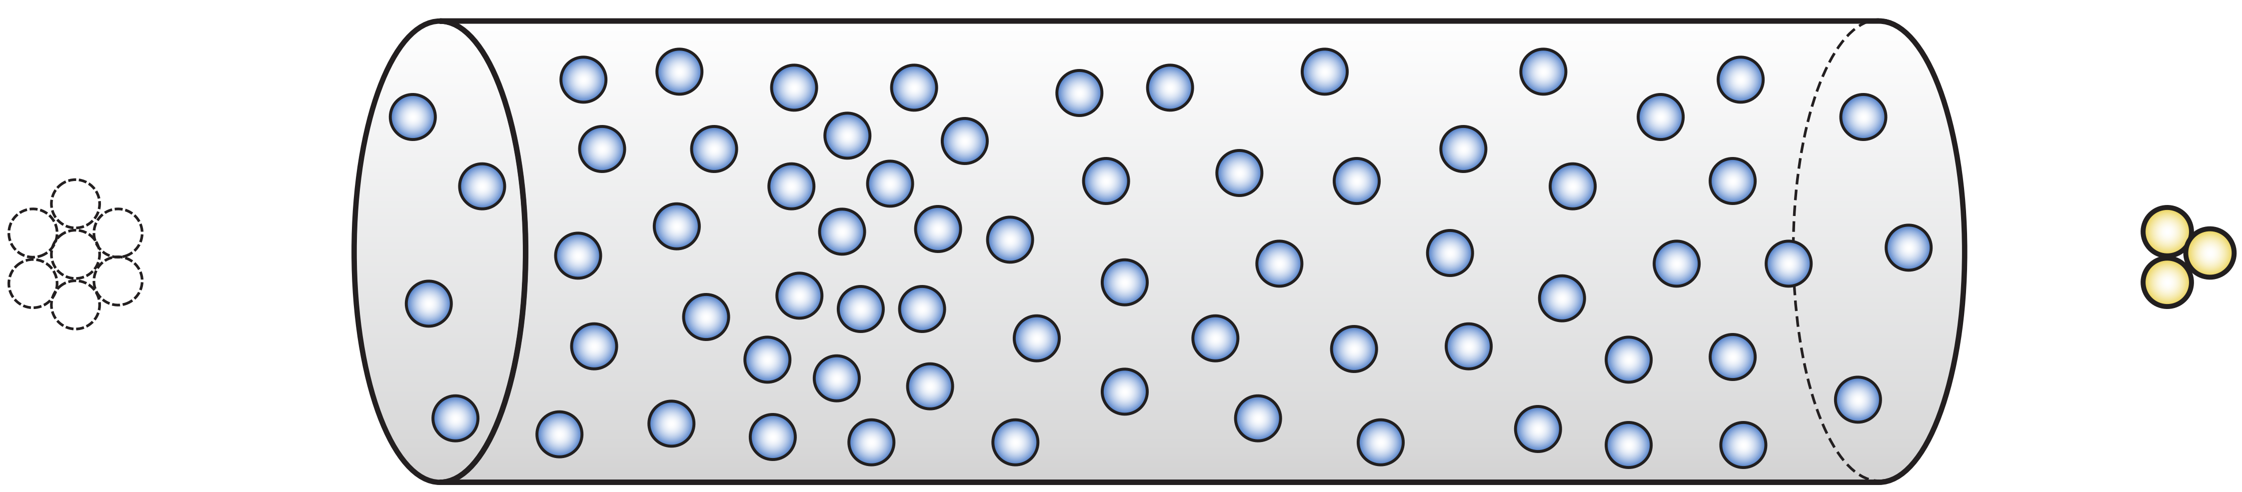

$$T(d) = \exp\left[-\int_0^d \mu(x)\mathrm{d}x\right]$$

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Attenuation (Transmittance)

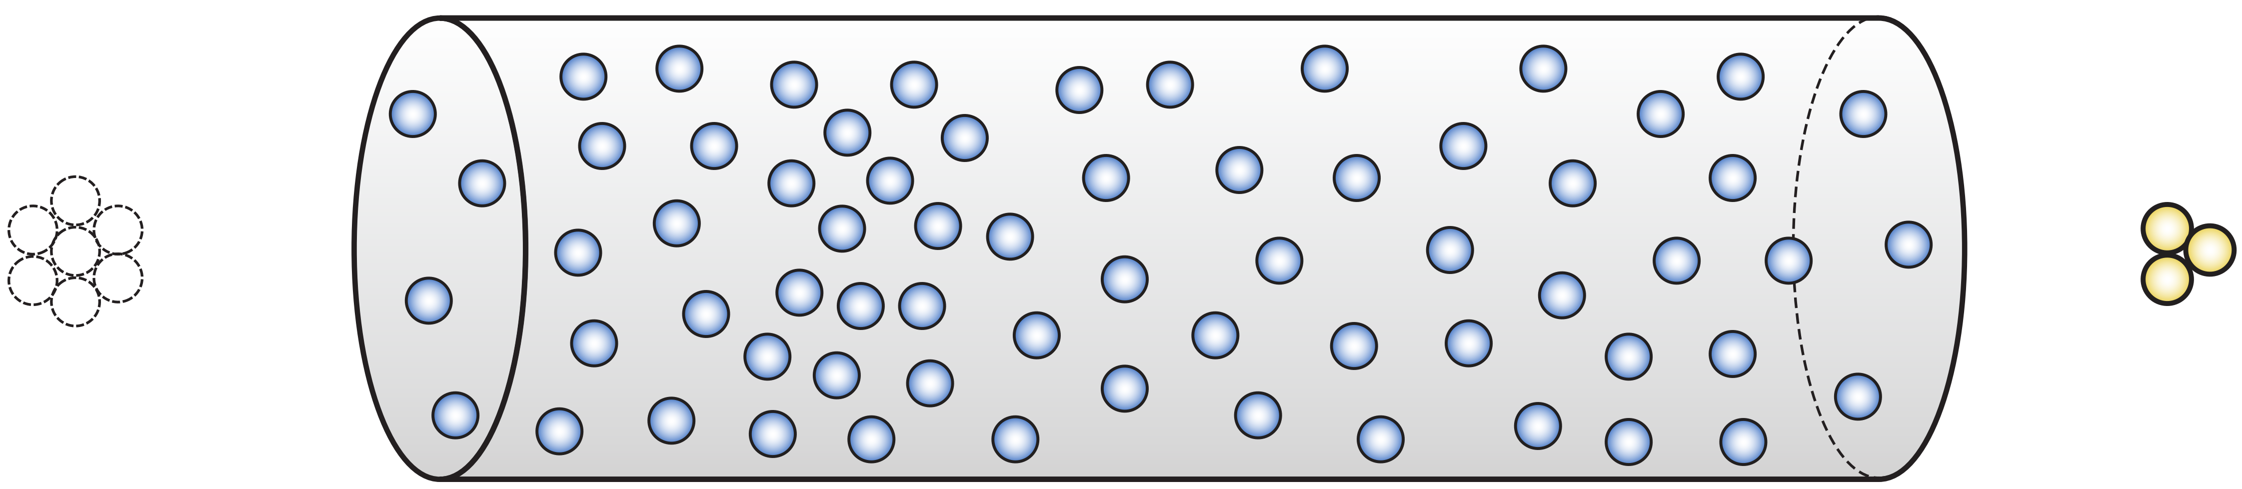

$$T(d) = \exp\left[-\int_0^d \mu(x)\mathrm{d}x\right]$$

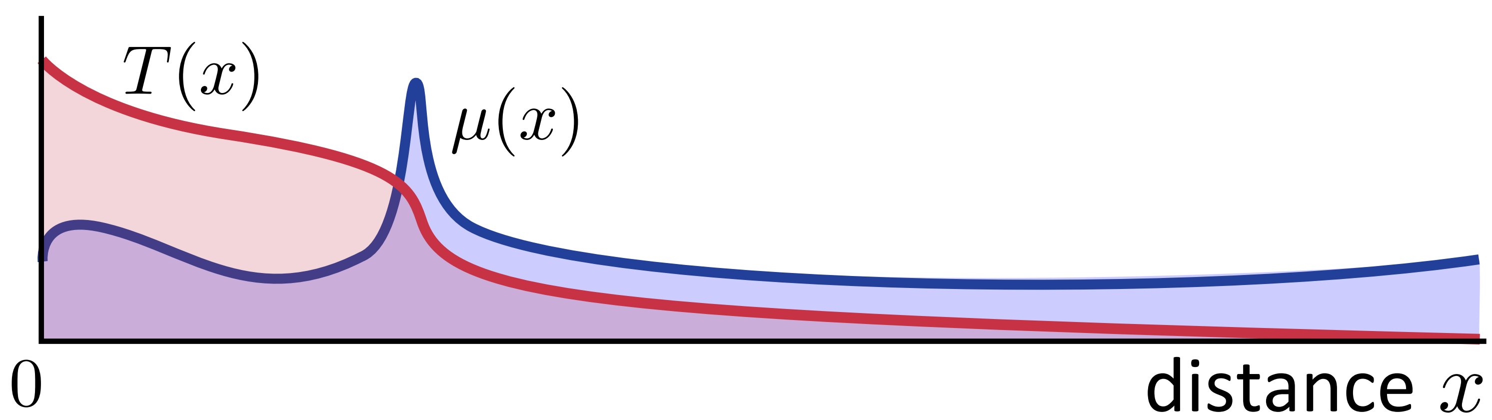

Having these coefficients, we can probabilistically simulate how far a photon travels before interacting with a particle. This is called free-path sampling (or free-flight distance sampling). We can also estimate how much light is transmitted between any two points without being absorbed or out-scattered. The transmittance through the medium is described by an exponentiated integral of the negative extinction coefficient. Here we plot the extinction coefficient, and the corresponding transmittance function along the ray.

Rendering Volumes

Efficient free-path sampling and transmittance estimation is key to many volume-rendering algorithms, here we have for instance a volumetric path tracer. A lot of research has been recently devoted to how to sample global illumination in volumes efficiently, and also how to build the free-paths quickly (the black segments here). But relatively little has been done to efficiently estimate the transmittance,... all the blue segments in this illustration. So this is what we focus on in this talk, but before I mention our contributions, let me first review existing techniques.

Rendering Volumes

SIGGRAPH ASIA 2014 SHENZHEN

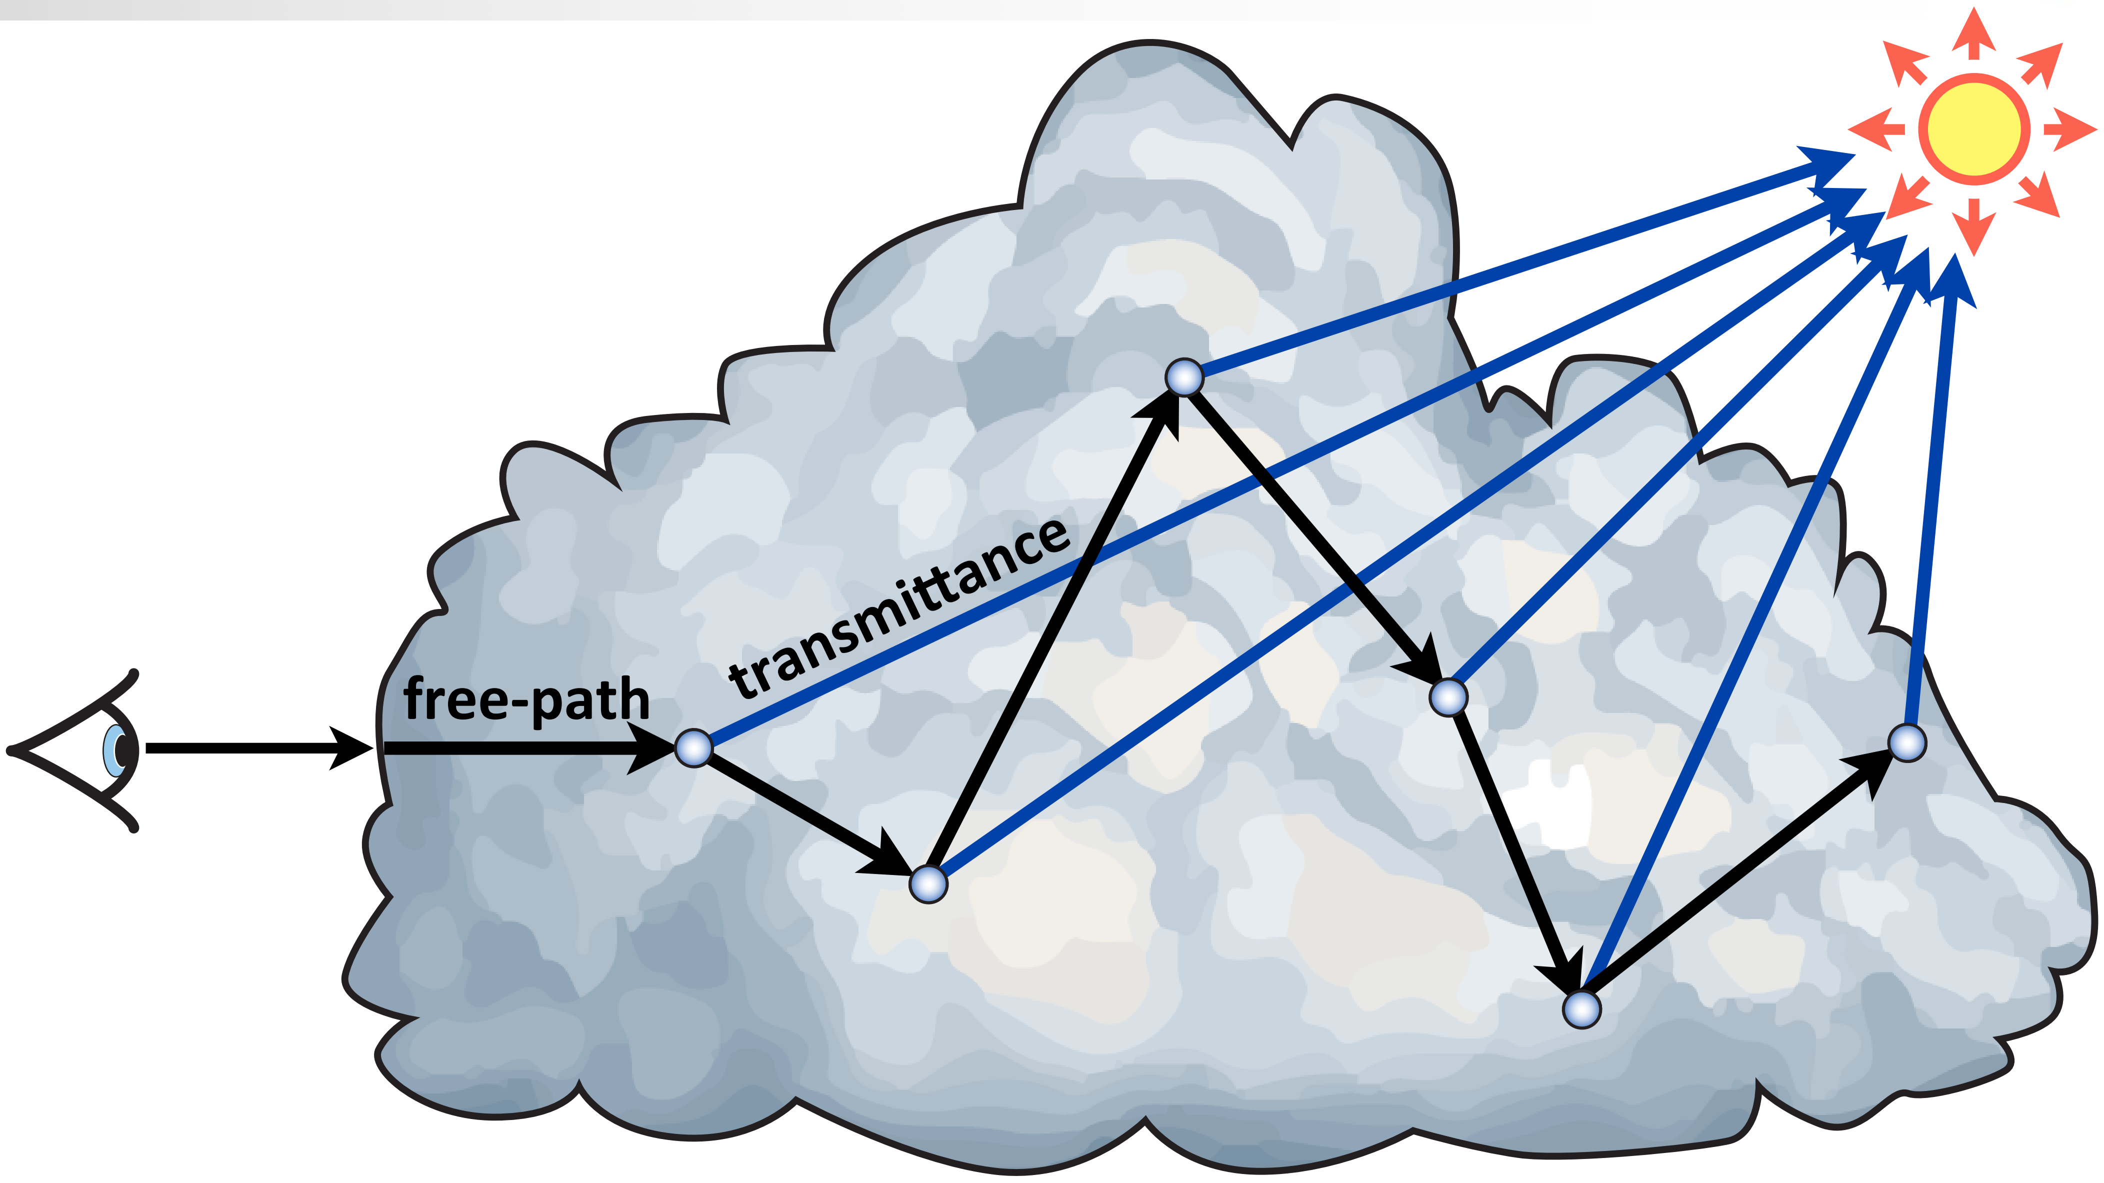

Efficient free-path sampling and transmittance estimation is key to many volume-rendering algorithms, here we have for instance a volumetric path tracer. A lot of research has been recently devoted to how to sample global illumination in volumes efficiently, and also how to build the free-paths quickly (the black segments here). But relatively little has been done to efficiently estimate the transmittance,... all the blue segments in this illustration. So this is what we focus on in this talk, but before I mention our contributions, let me first review existing techniques.

Rendering Volumes

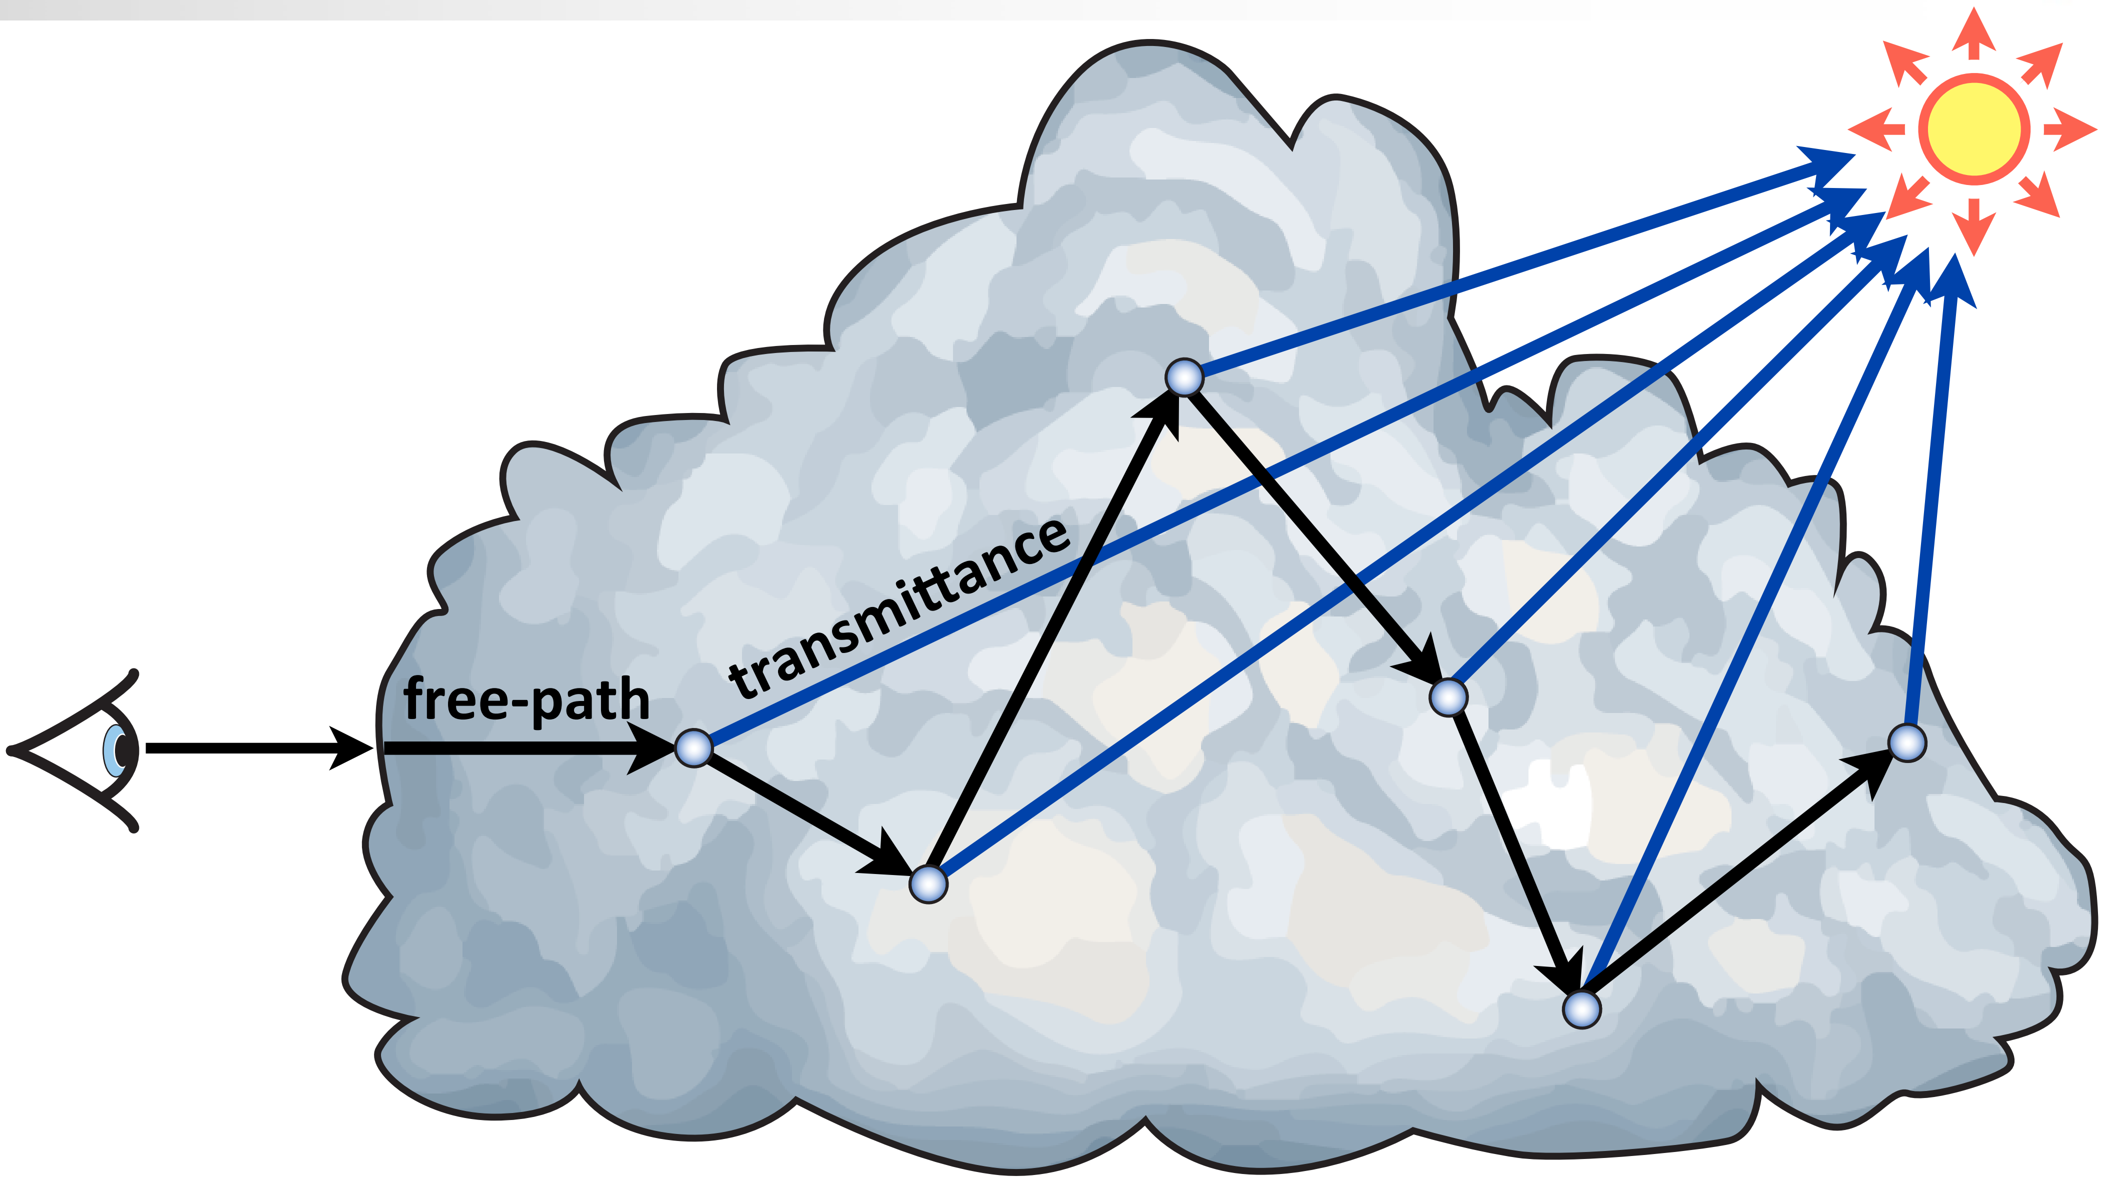

(Bidirectional) Path Tracing
Photon Mapping
Many-Light Rendering
[Jarosz et al. 2011],[Novak et al. 2012],
[Georgiev et al. 2013],[Krivanek et al. 2014]

Efficient free-path sampling and transmittance estimation is key to many volume-rendering algorithms, here we have for instance a volumetric path tracer. A lot of research has been recently devoted to how to sample global illumination in volumes efficiently, and also how to build the free-paths quickly (the black segments here). But relatively little has been done to efficiently estimate the transmittance,... all the blue segments in this illustration. So this is what we focus on in this talk, but before I mention our contributions, let me first review existing techniques.

Rendering Volumes

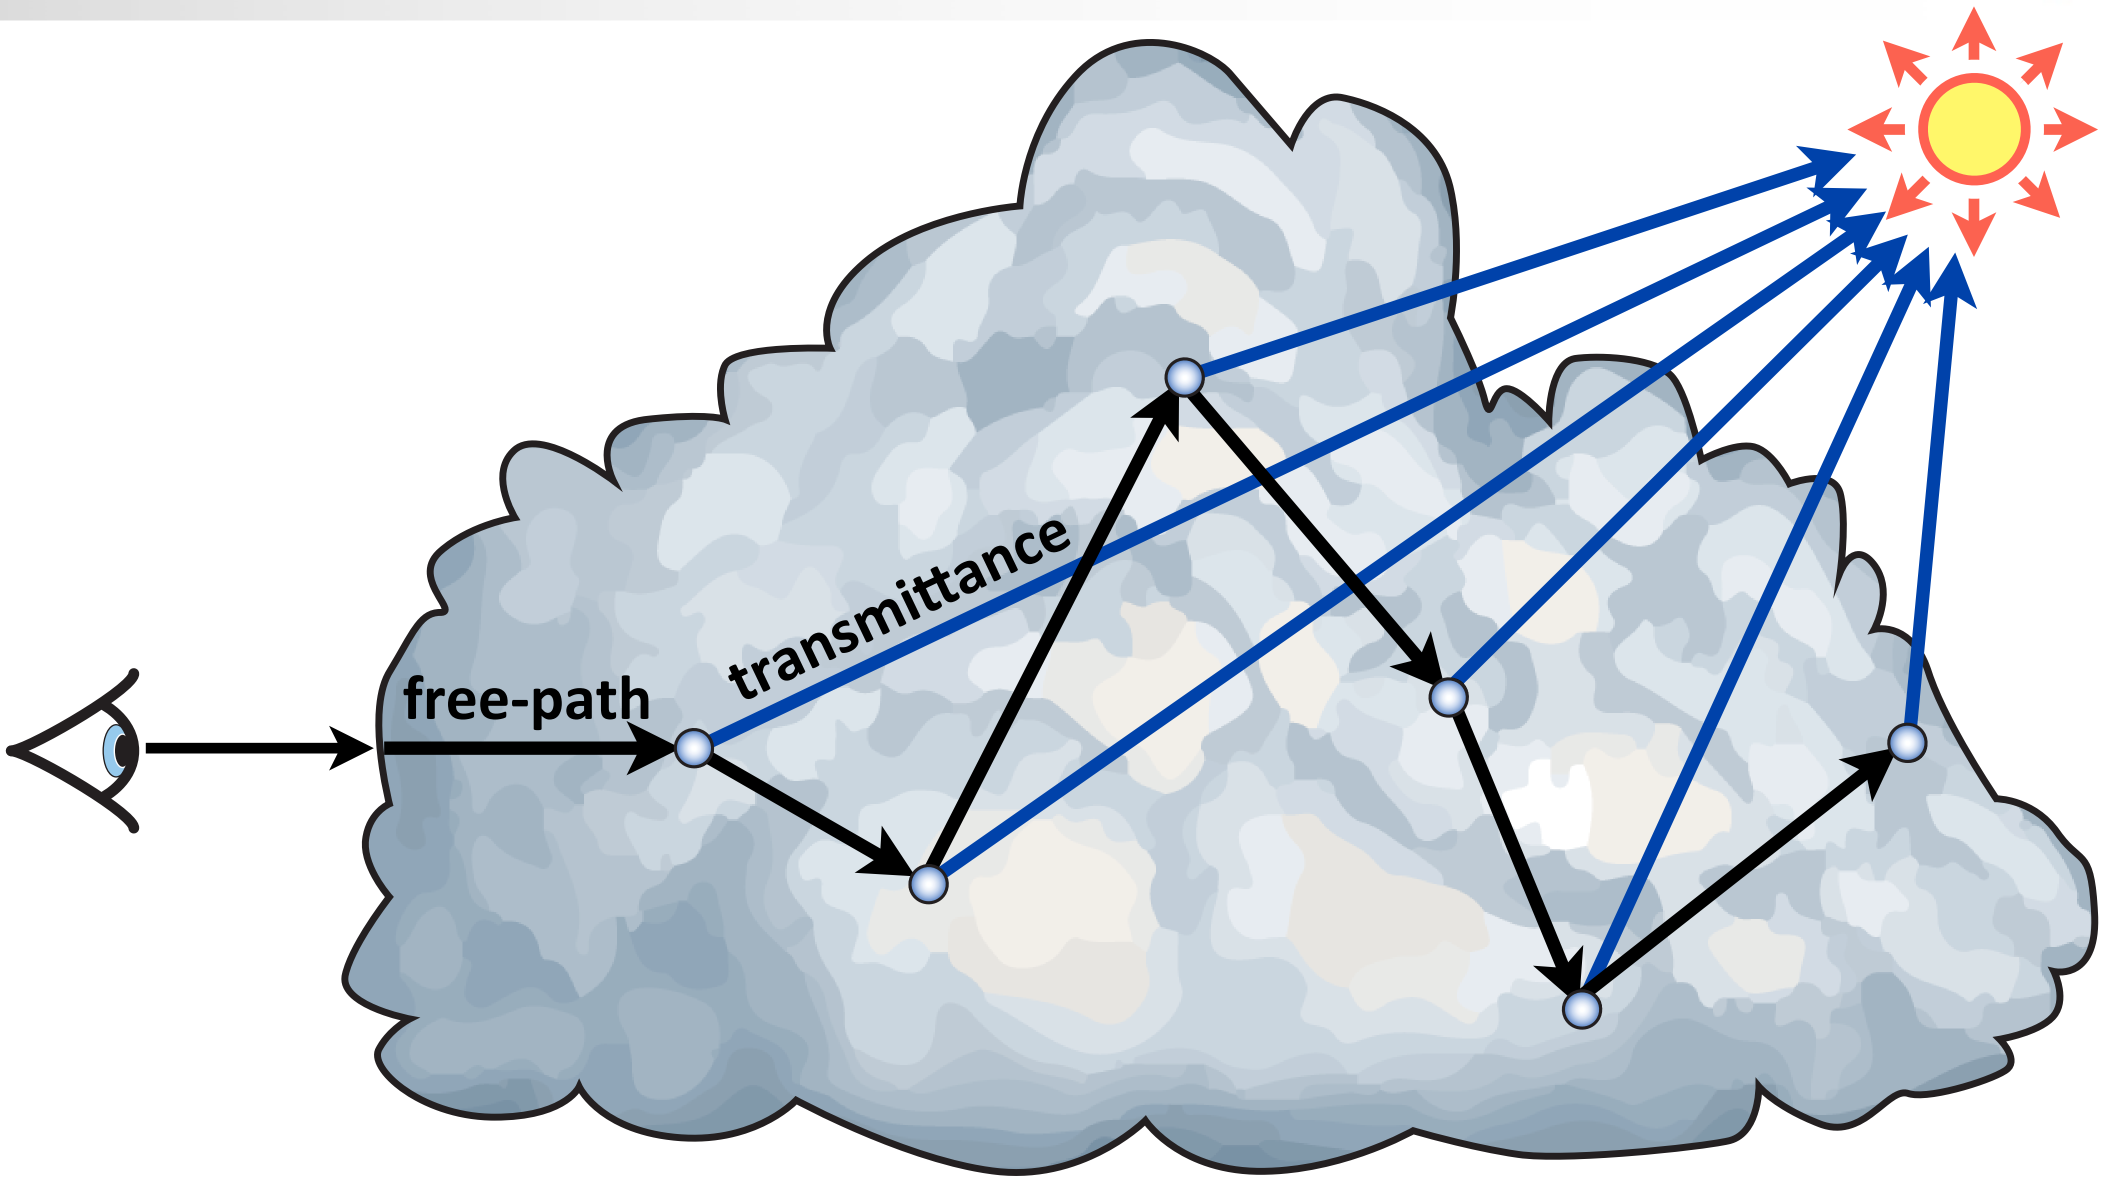

(Bidirectional) Path Tracing
Photon Mapping
Many-Light Rendering
[Jarosz et al. 2011],[Novak et al. 2012],
[Georgiev et al. 2013],[Krivanek et al. 2014]

Free-path sampling
[Yue et al. 2010]
[Szirmay-Kalos et al. 2011]

Efficient free-path sampling and transmittance estimation is key to many volume-rendering algorithms, here we have for instance a volumetric path tracer. A lot of research has been recently devoted to how to sample global illumination in volumes efficiently, and also how to build the free-paths quickly (the black segments here). But relatively little has been done to efficiently estimate the transmittance,... all the blue segments in this illustration. So this is what we focus on in this talk, but before I mention our contributions, let me first review existing techniques.

Prior Art

If the medium is homogeneous, we can evaluate transmittance exactly and sample free paths analytically.
This is unfortunately possibly only for homogeneous, or very simple volumes.

If the volume is heterogeneous, but can be represented as voxels, we can step through the voxels and use the analytic techniques. Unfortunately, the stepping can quickly become too expensive.

To make the computation tractable, we often ignore the boundaries and simply march with a constant step size. The drawback of such ray-marching is that it provides only approximate results, it is indeed BIASED.

In our case, we need a technique that is unbiased, such as Delta tracking, so that we can rely on the error being averaged out by taking more samples. Since Delta tracking forms the basis for our algorithms, I'll describe it in greater detail.

Prior Art

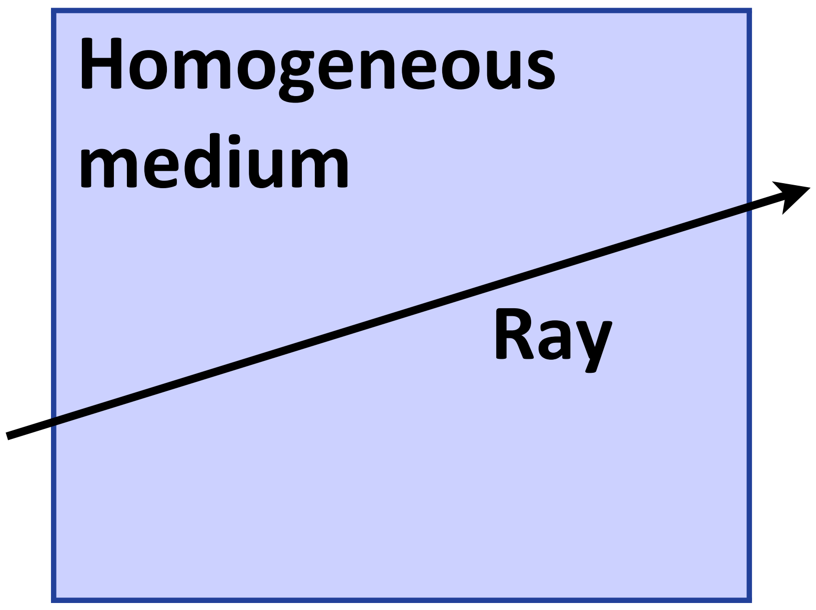

- **Analytic Techniques**
  - only for homogeneous or simple media
  - exact, efficient
    limited applicability

If the medium is homogeneous, we can evaluate transmittance exactly and sample free paths analytically.
This is unfortunately possibly only for homogeneous, or very simple volumes.

If the volume is heterogeneous, but can be represented as voxels, we can step through the voxels and use the analytic techniques. Unfortunately, the stepping can quickly become too expensive.

To make the computation tractable, we often ignore the boundaries and simply march with a constant step size. The drawback of such ray-marching is that it provides only approximate results, it is indeed BIASED.

In our case, we need a technique that is unbiased, such as Delta tracking, so that we can rely on the error being averaged out by taking more samples. Since Delta tracking forms the basis for our algorithms, I'll describe it in greater detail.

Prior Art

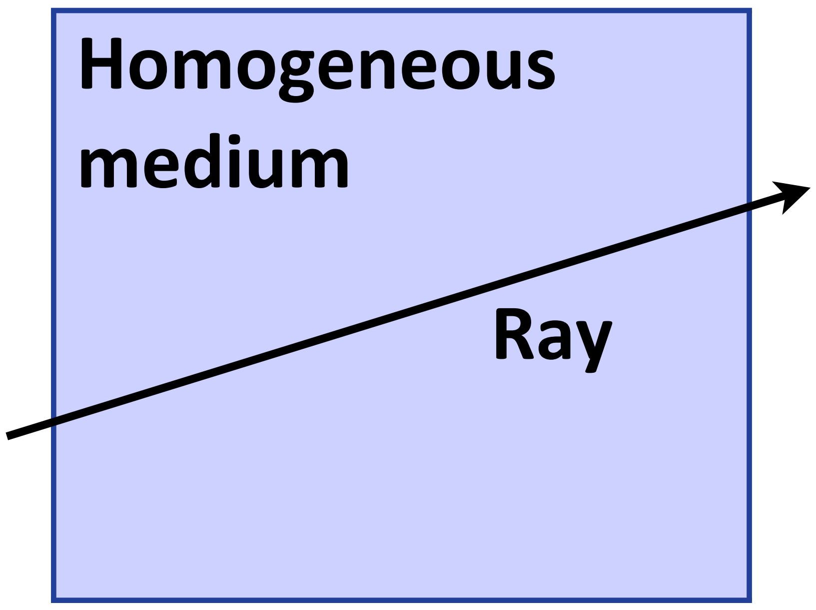

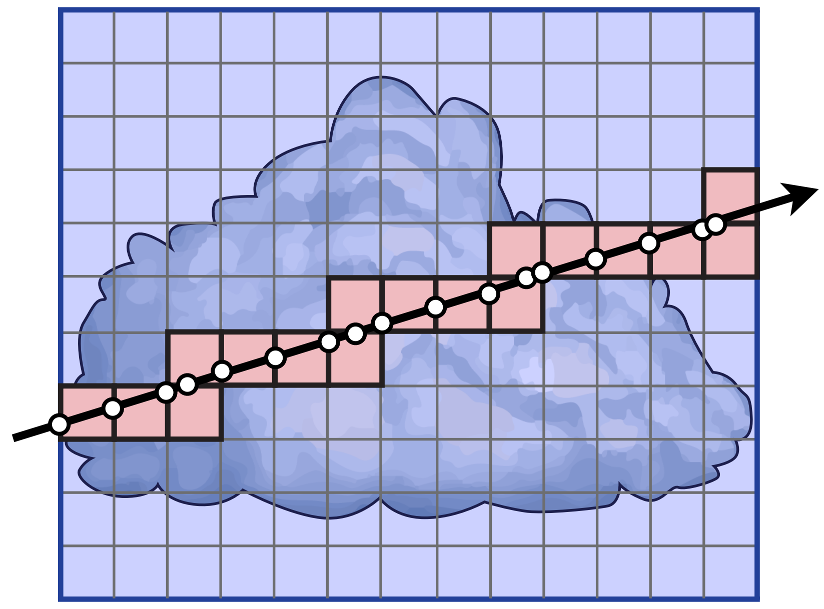

- **Analytic Techniques**
  - only for homogeneous or simple media
  - exact, efficient
    limited applicability
- **Regular Tracking**
  - set of homogeneous regions
  - exact, semi-analytic
    expensive

If the medium is homogeneous, we can evaluate transmittance exactly and sample free paths analytically.
This is unfortunately possibly only for homogeneous, or very simple volumes.

If the volume is heterogeneous, but can be represented as voxels, we can step through the voxels and use the analytic techniques. Unfortunately, the stepping can quickly become too expensive.

To make the computation tractable, we often ignore the boundaries and simply march with a constant step size. The drawback of such ray-marching is that it provides only approximate results, it is indeed BIASED.

In our case, we need a technique that is unbiased, such as Delta tracking, so that we can rely on the error being averaged out by taking more samples. Since Delta tracking forms the basis for our algorithms, I'll describe it in greater detail.

Prior Art

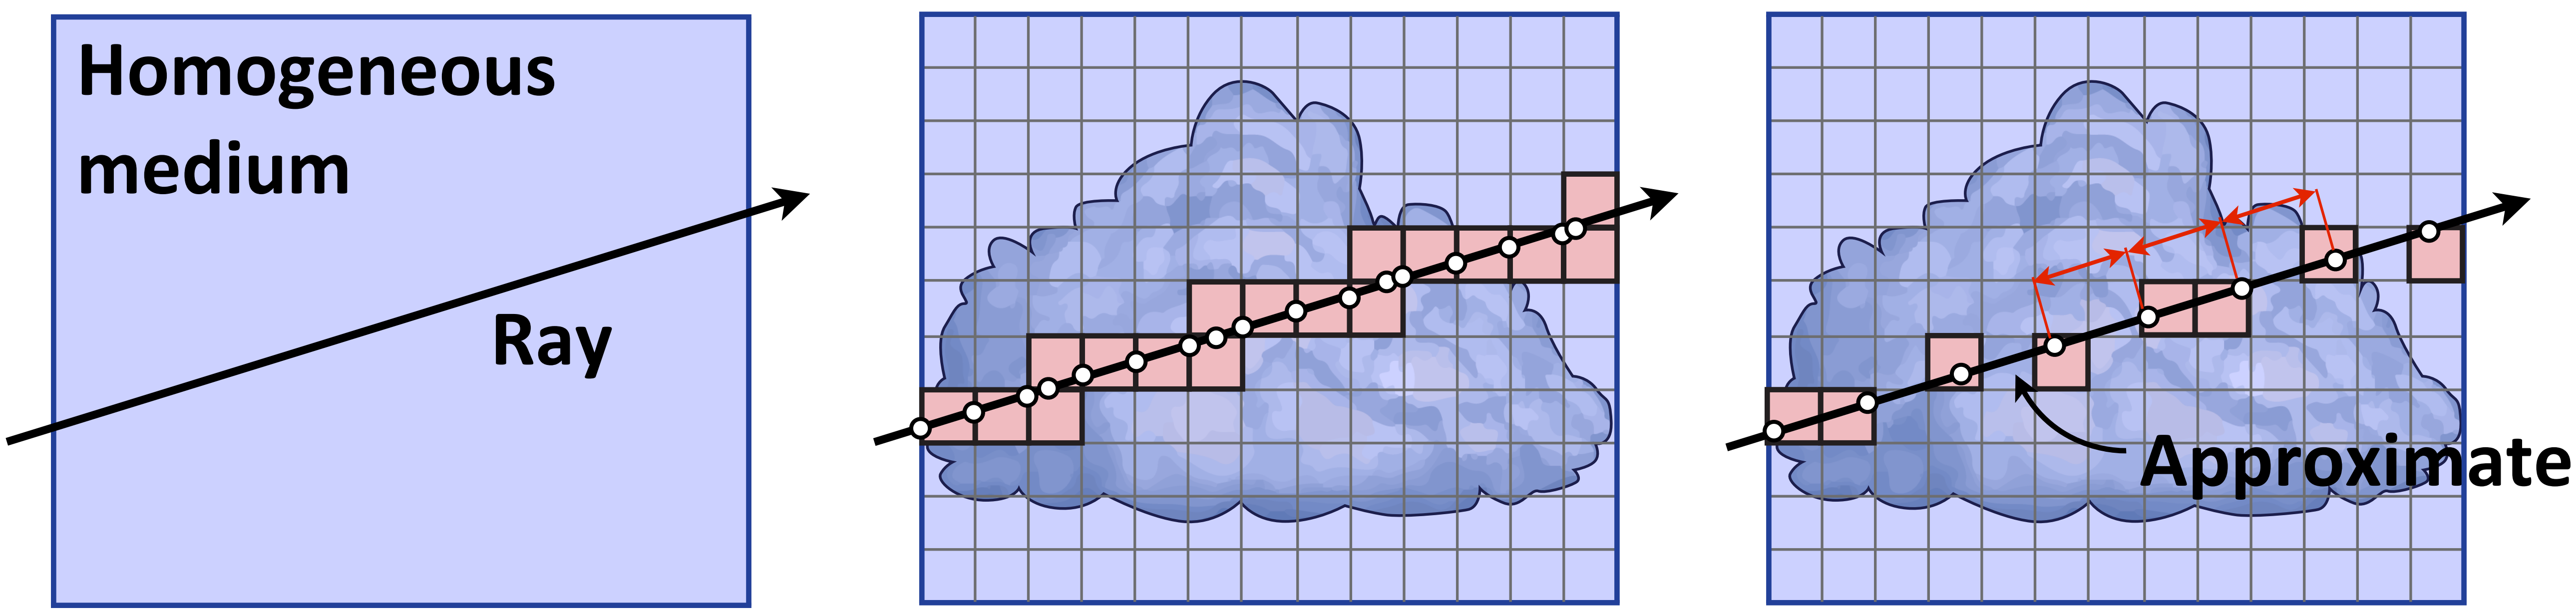

- **Analytic Techniques**
  - only for homogeneous or simple media
  - exact, efficient **limited applicability**
- **Regular Tracking**
  - set of homogeneous regions
  - exact, semi-analytic **expensive**
- **Ray Marching**
  - Perlin and Hoffert [1989]
  - heterogeneous media
  - approximate, **biased**

If the medium is homogeneous, we can evaluate transmittance exactly and sample free paths analytically.
This is unfortunately possibly only for homogeneous, or very simple volumes.

If the volume is heterogeneous, but can be represented as voxels, we can step through the voxels and use the analytic techniques. Unfortunately, the stepping can quickly become too expensive.

To make the computation tractable, we often ignore the boundaries and simply march with a constant step size. The drawback of such ray-marching is that it provides only approximate results, it is indeed BIASED.

In our case, we need a technique that is unbiased, such as Delta tracking, so that we can rely on the error being averaged out by taking more samples. Since Delta tracking forms the basis for our algorithms, I'll describe it in greater detail.

Prior Art

SIGGRAPH ASIA 2014 SHENZHEN

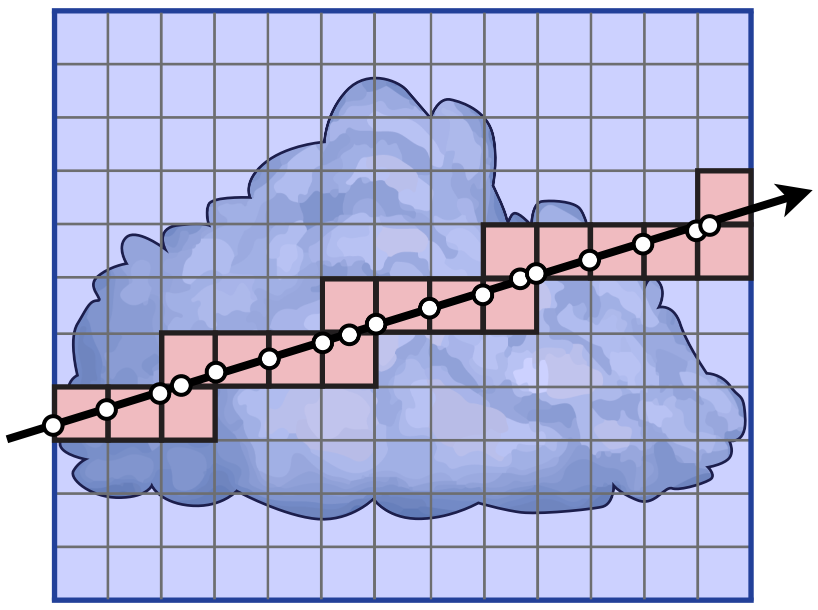

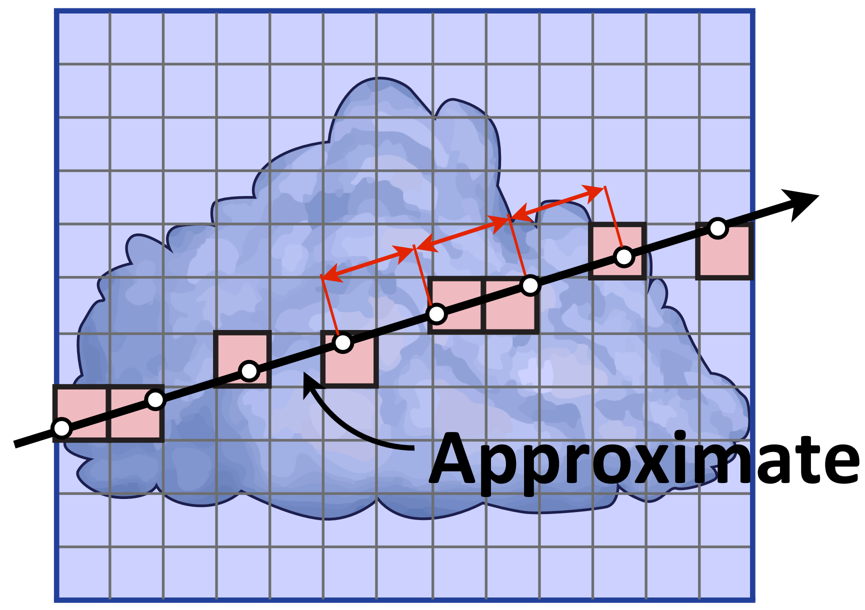

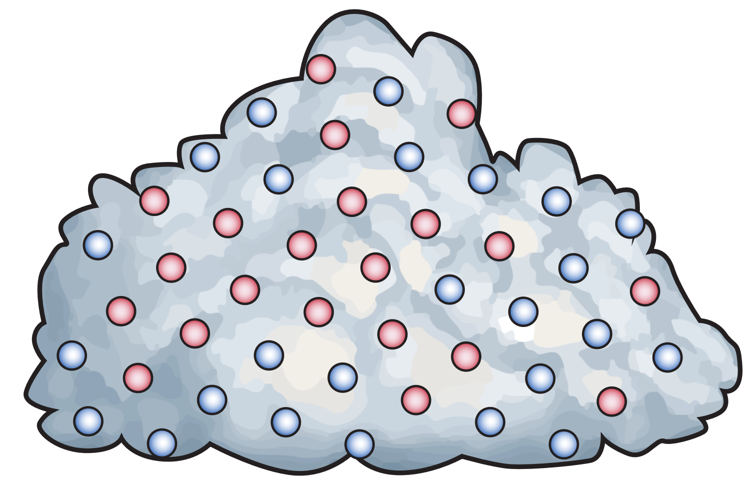

- **Regular Tracking**
  - set of homogeneous regions
  - exact, semi-analytic **expensive**
- **Ray Marching**
  - Perlin and Hoffert [1989]
  - heterogeneous media
  - approximate, **biased**
- **Delta Tracking**
  - von Neumann [1951]
  - Woodcock et al. [1956]
  - Skullerud [1968]
  - Galtier et al. [2013]

If the medium is homogeneous, we can evaluate transmittance exactly and sample free paths analytically.
This is unfortunately possibly only for homogeneous, or very simple volumes.

If the volume is heterogeneous, but can be represented as voxels, we can step through the voxels and use the analytic techniques. Unfortunately, the stepping can quickly become too expensive.

To make the computation tractable, we often ignore the boundaries and simply march with a constant step size. The drawback of such ray-marching is that it provides only approximate results, it is indeed BIASED.

In our case, we need a technique that is unbiased, such as Delta tracking, so that we can rely on the error being averaged out by taking more samples. Since Delta tracking forms the basis for our algorithms, I'll describe it in greater detail.

Delta Tracking

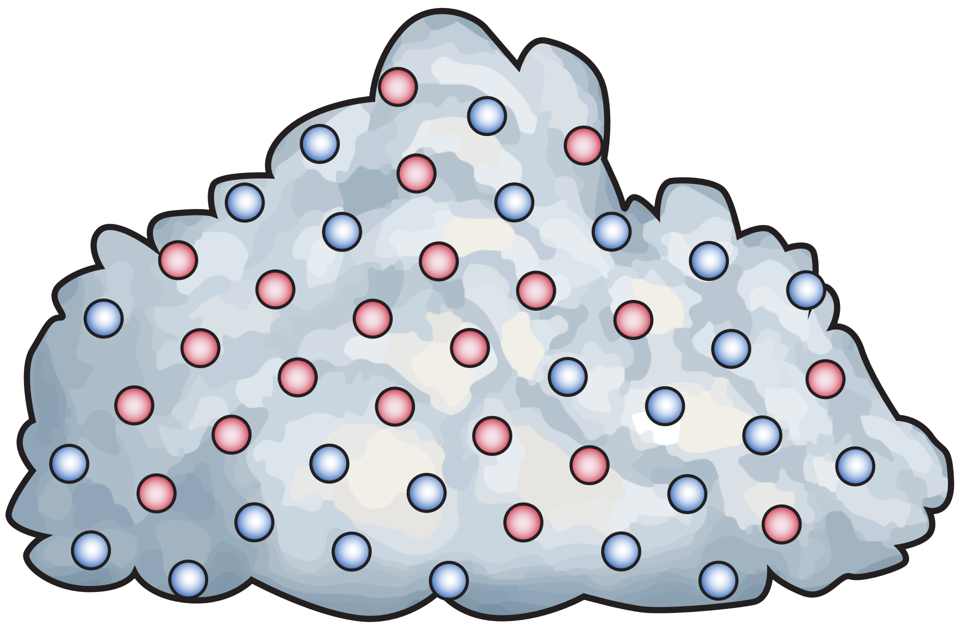

The technique is also known as Woodcock tracking, pseudo scattering, or the null-collision algorithm.

The fundamental idea of the algorithm is to homogenize the medium, so that we can sample free-paths analytically.

This is achieved by adding special particles that have albedo 1 and perfectly forward scattering phase function so they have effectively no impact on photon trajectories. We will refer to them as fictitious particles.

Prentending that such particles exist greatly simplifies the free-flight distance sampling.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

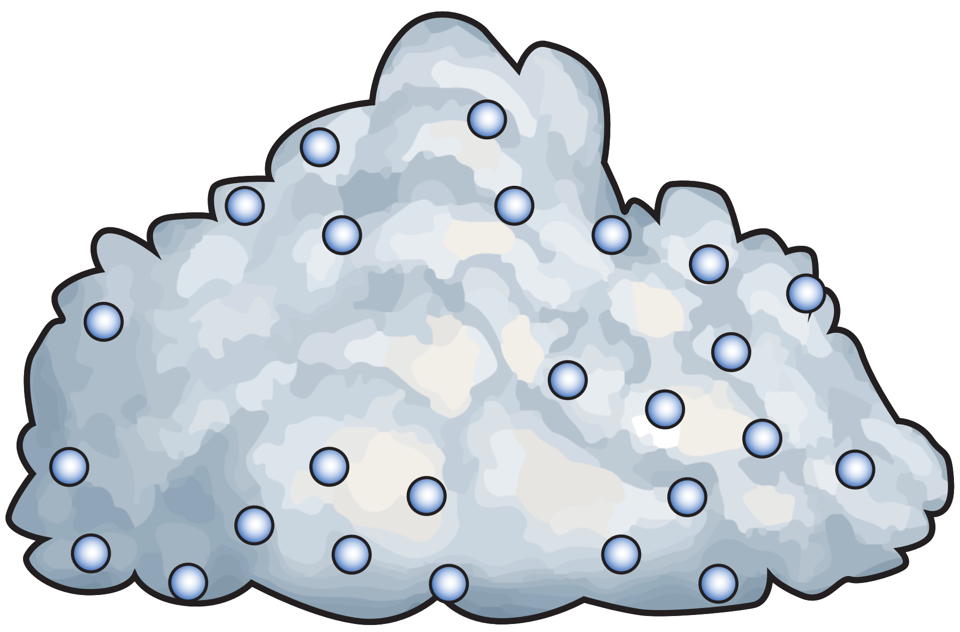

The technique is also known as Woodcock tracking, pseudo scattering, or the null-collision algorithm.

The fundamental idea of the algorithm is to homogenize the medium, so that we can sample free-paths analytically.

This is achieved by adding special particles that have albedo 1 and perfectly forward scattering phase function so they have effectively no impact on photon trajectories. We will refer to them as fictitious particles.

Prentending that such particles exist greatly simplifies the free-flight distance sampling.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

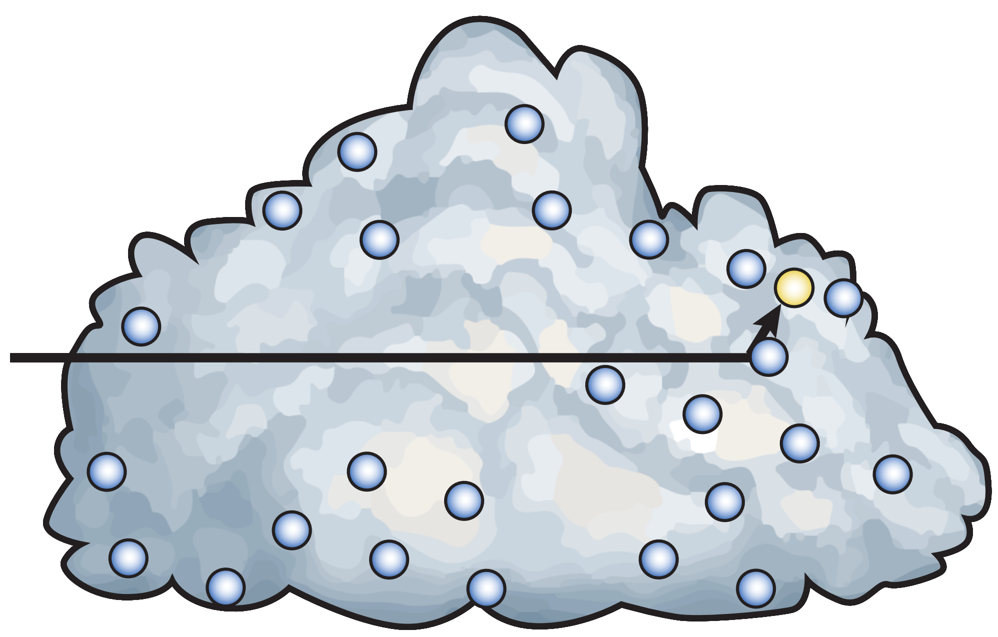

The technique is also known as Woodcock tracking, pseudo scattering, or the null-collision algorithm.

The fundamental idea of the algorithm is to homogenize the medium, so that we can sample free-paths analytically.

This is achieved by adding special particles that have albedo 1 and perfectly forward scattering phase function so they have effectively no impact on photon trajectories. We will refer to them as fictitious particles.

Prentending that such particles exist greatly simplifies the free-flight distance sampling.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

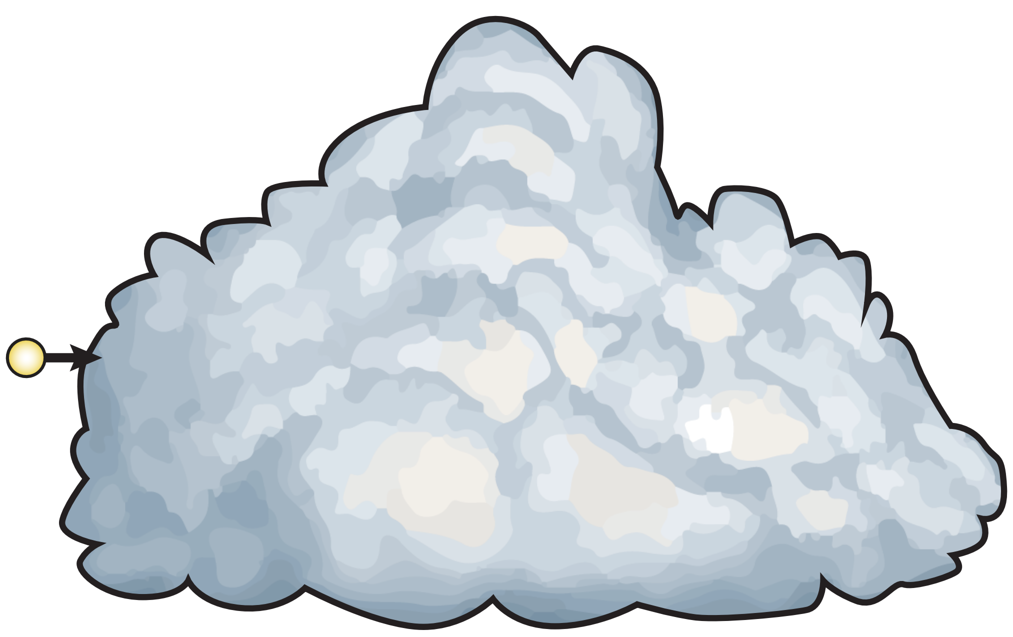

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

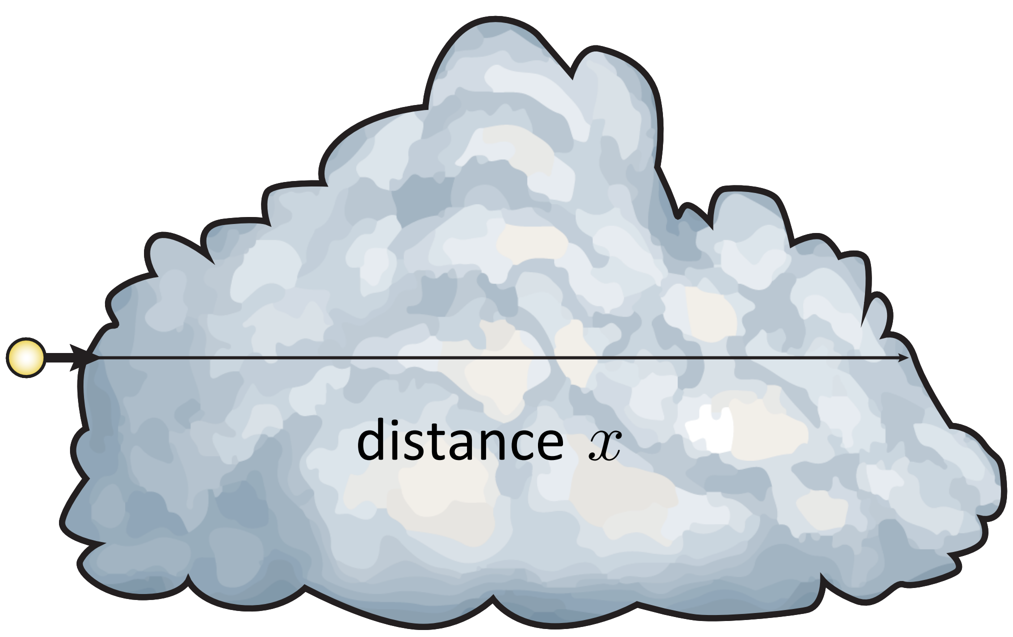

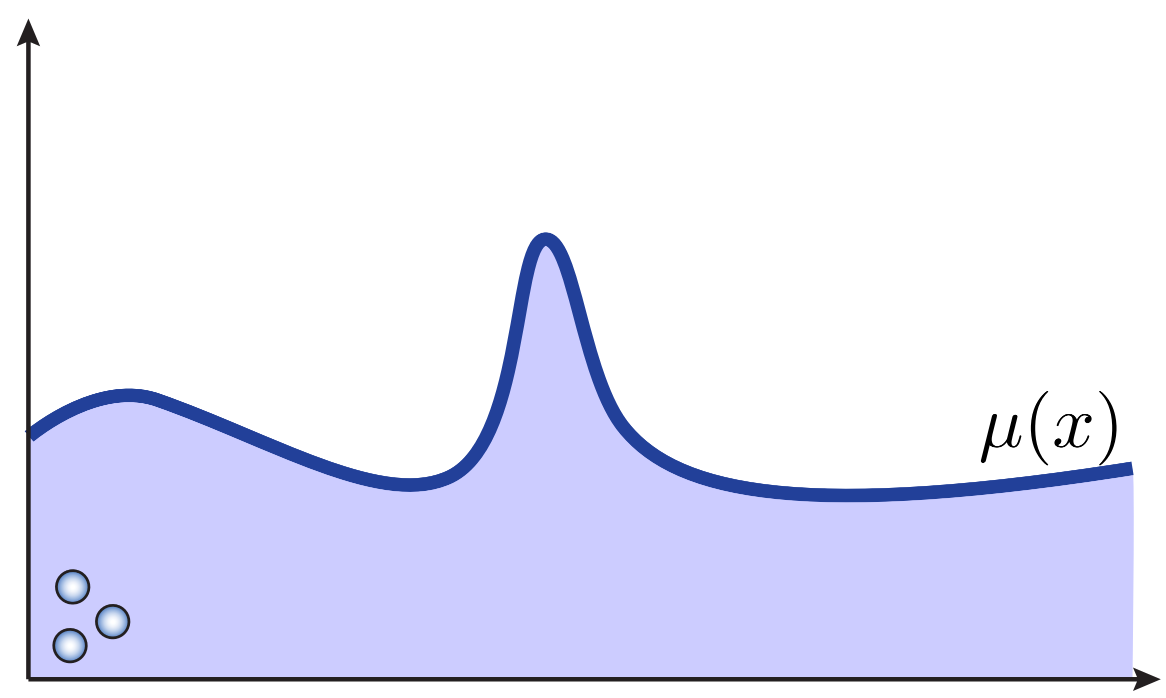

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

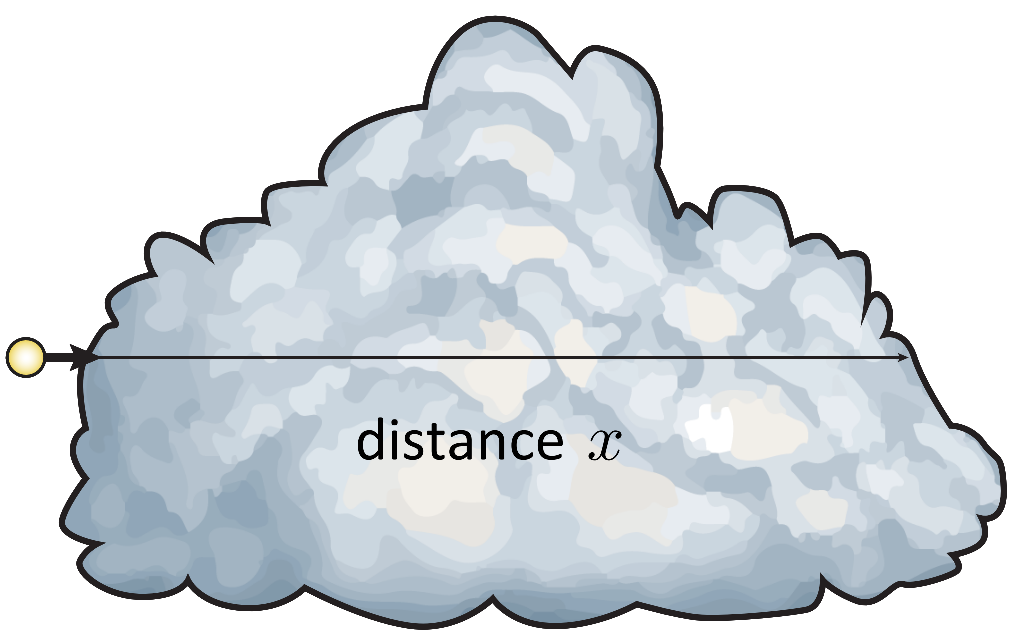

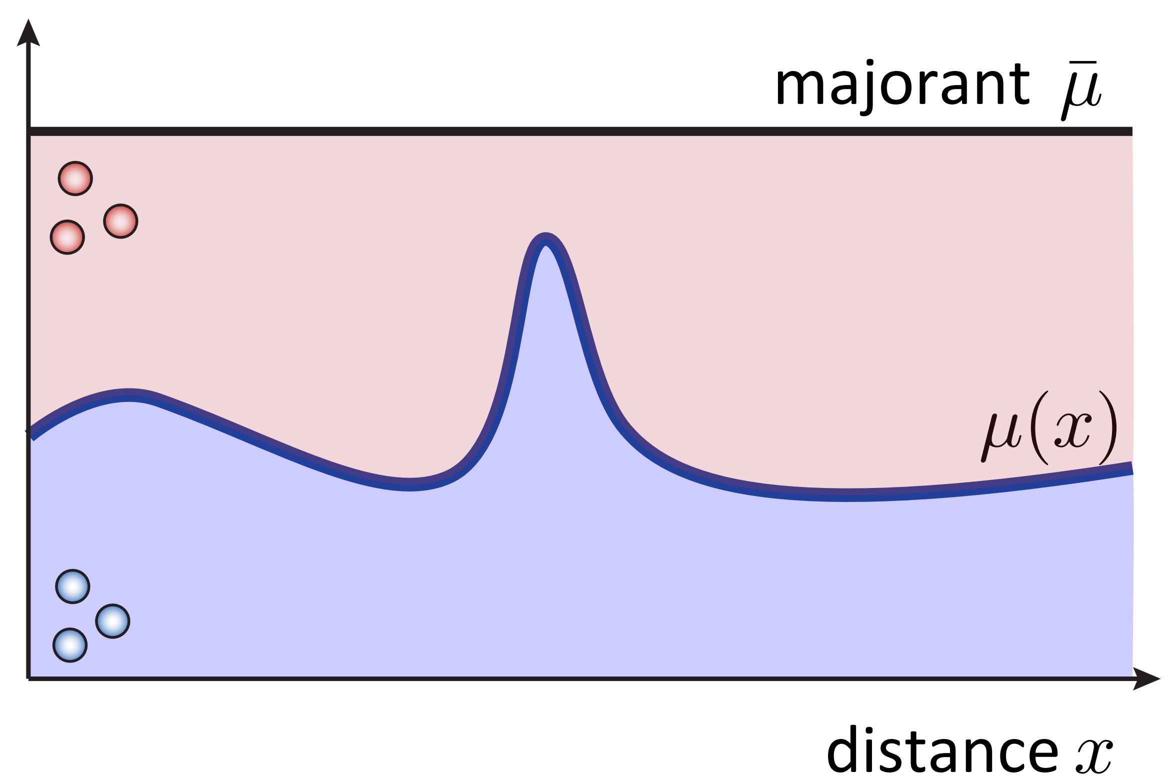

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

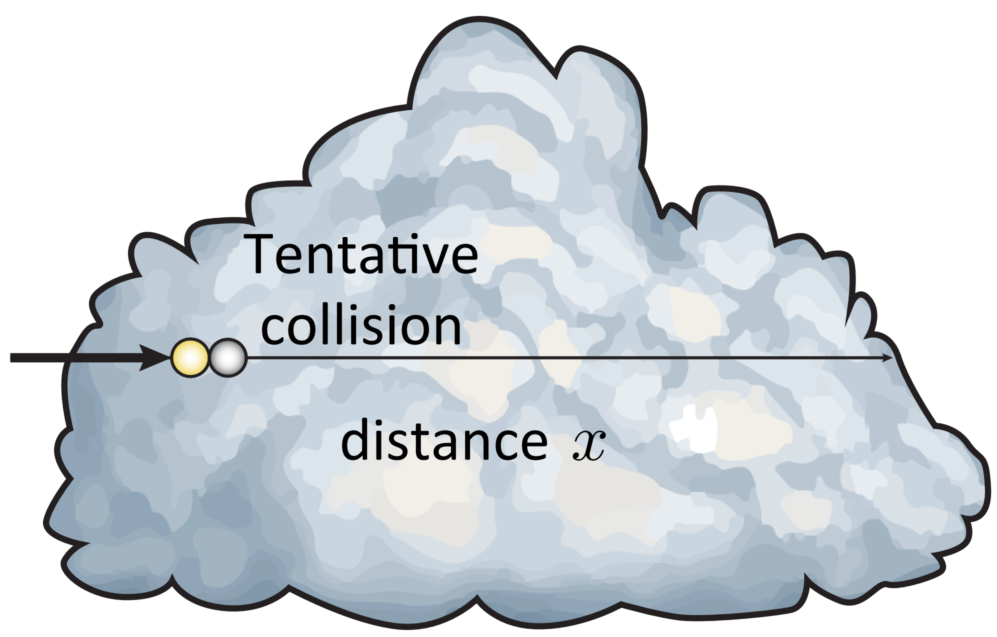

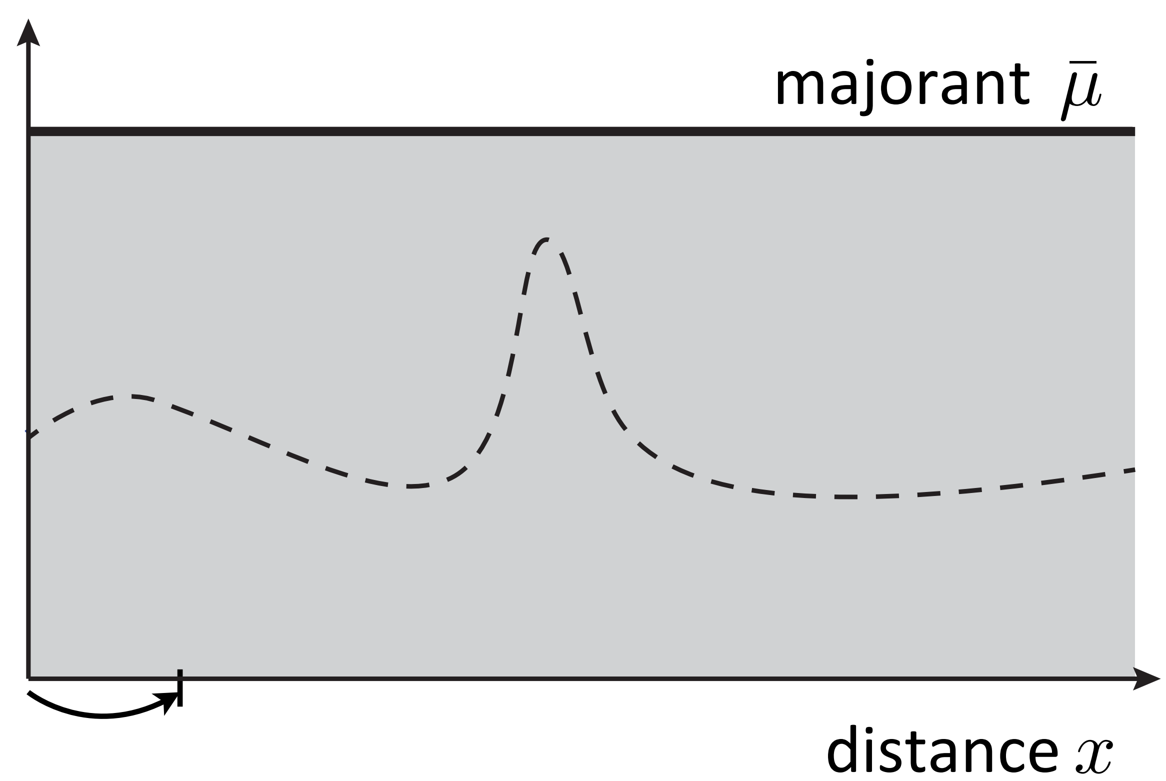

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

$$P_r = \frac{\mu(x)}{\bar{\mu}} \qquad P_f = 1 - P_r$$

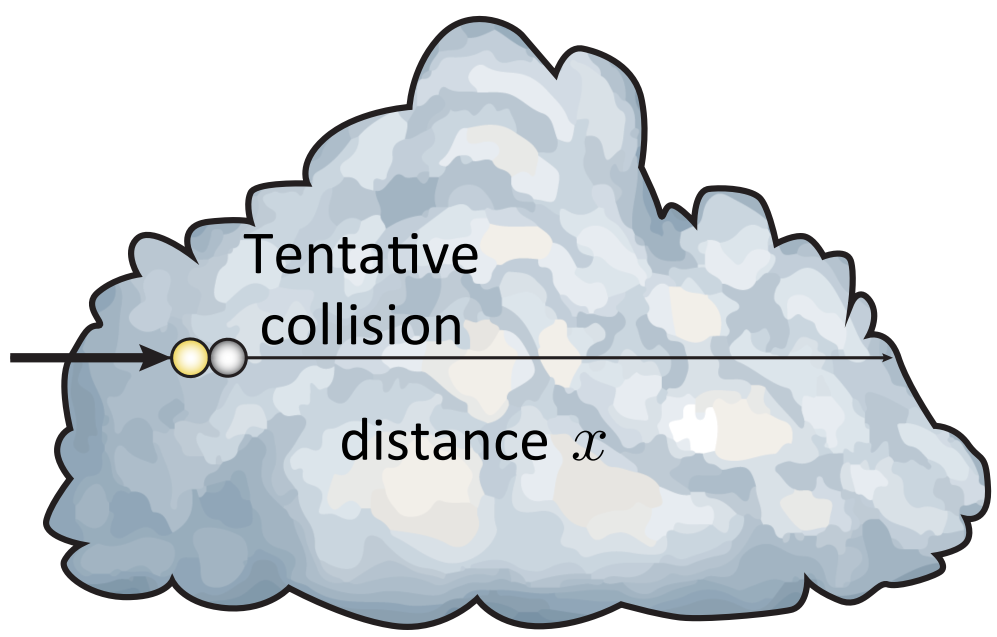

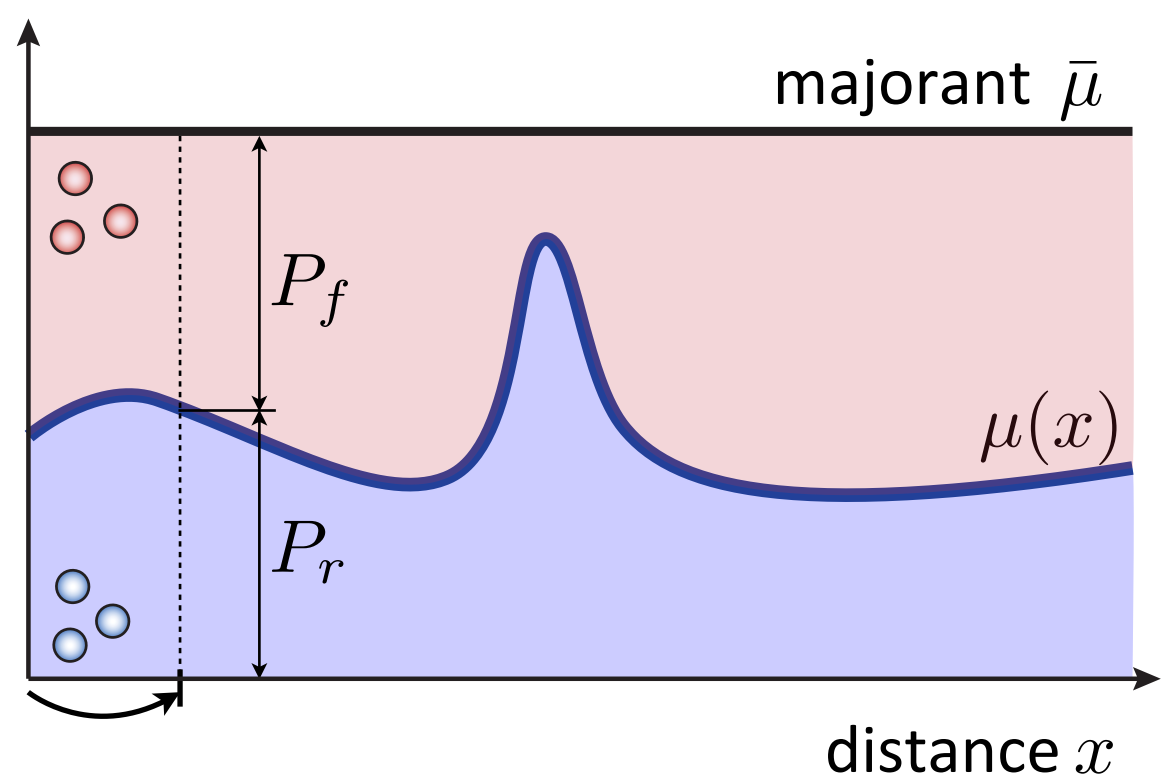

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

$$P_r = \frac{\mu(x)}{\bar{\mu}} \qquad P_f = 1 - P_r$$

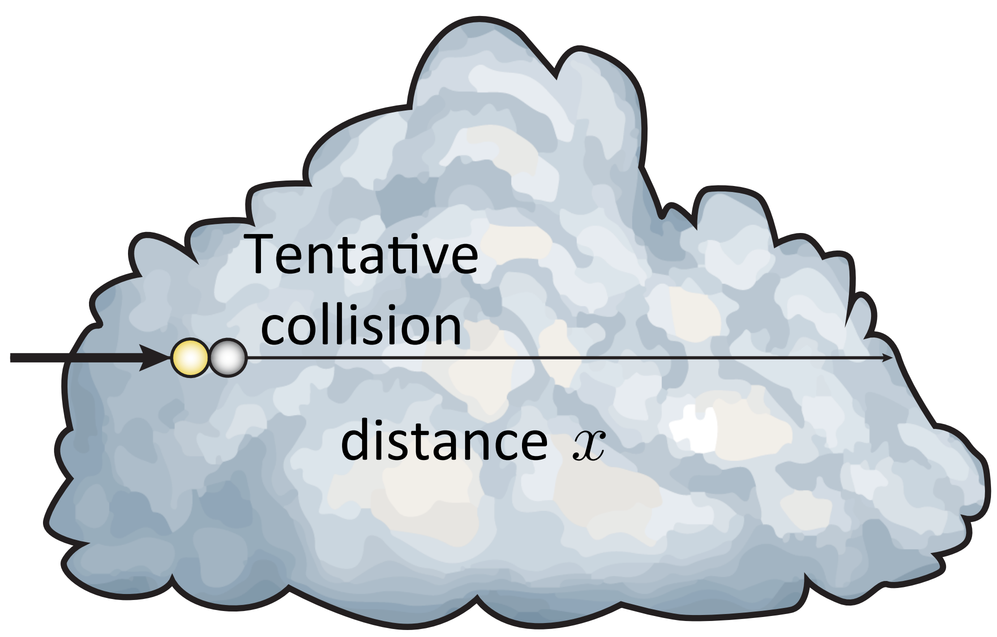

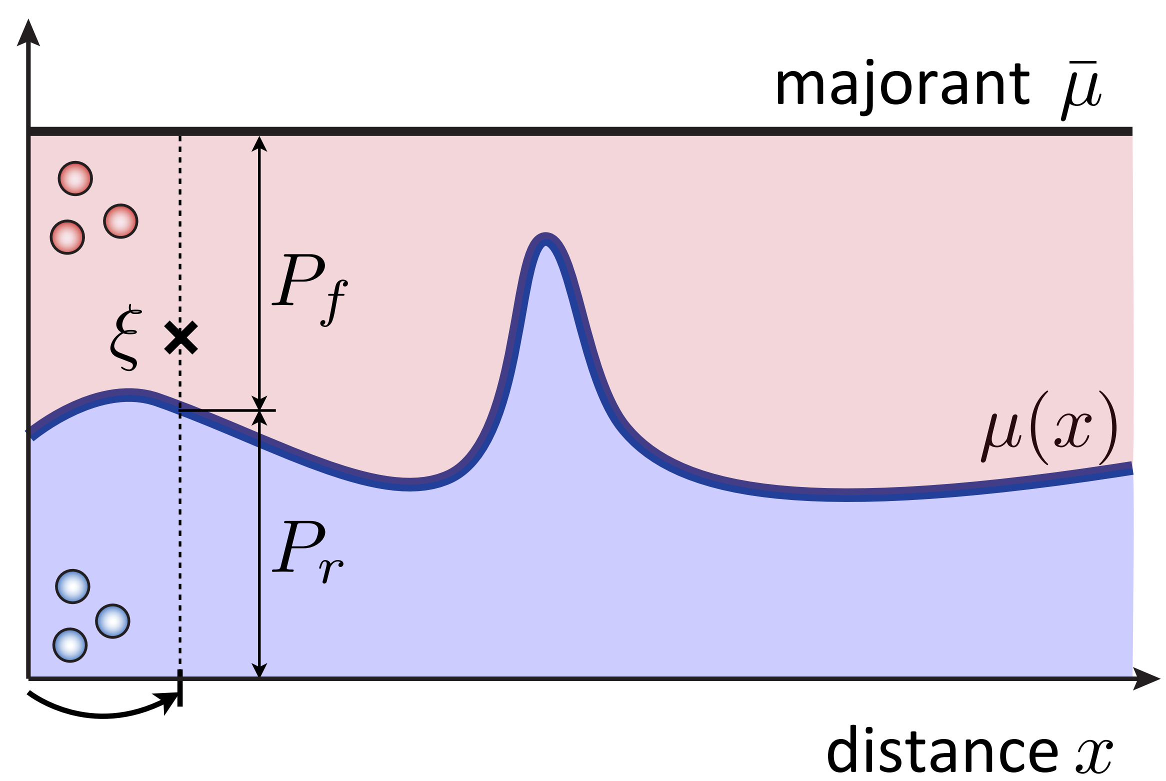

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

$$P_r = \frac{\mu(x)}{\bar{\mu}} \qquad P_f = 1 - P_r$$

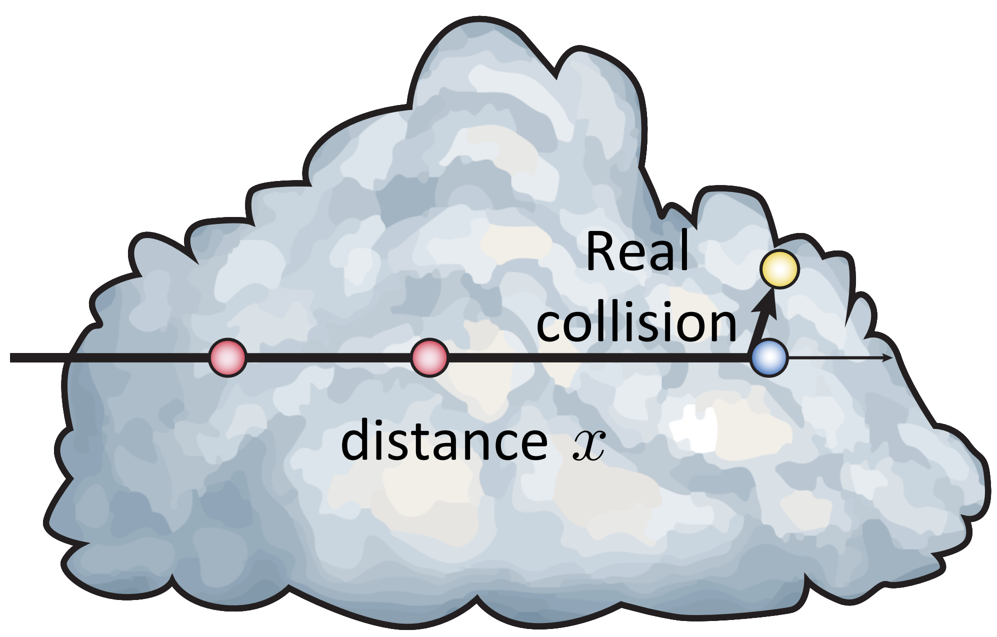

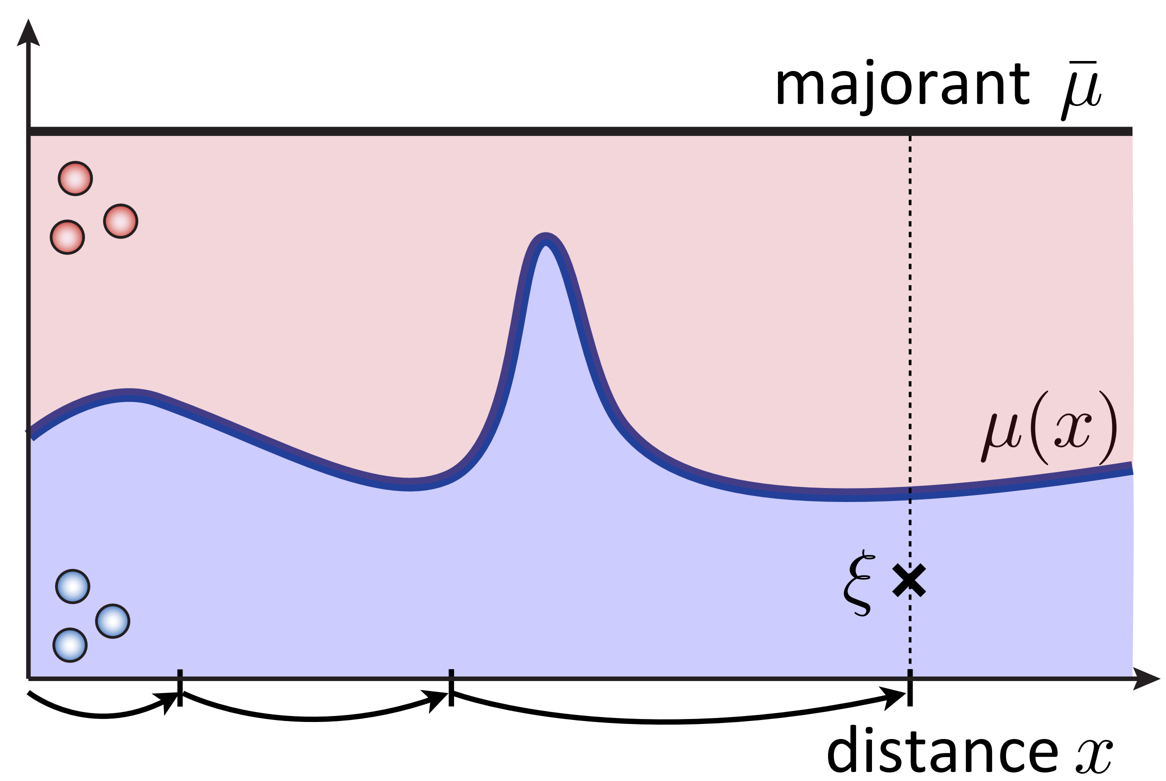

Let's say we want to probabilistically sample free-flight distance along a ray with the following extinction function.

We first add fictitious particles along the ray, but always just the right amount so that when we combine the real and the fictitious extinction, we get a constant value, often referred to as the majorant extinction.

The fact that the combined medium is homogeneous allows the algorithm to ANALYTICALLY sample free-paths
creating a so-called tentative collision.

Next, we need to decide, whether the tentative collision is a real one, or whether it involved a fictitious particle.

This is done probabilistically where the probabilities are set to the relative concentrations of real and fictitious particles.

In this case, the algorithm probabilistically classified the collision as fictitious, so it will continue generating new tentative collisions until is classified as real.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

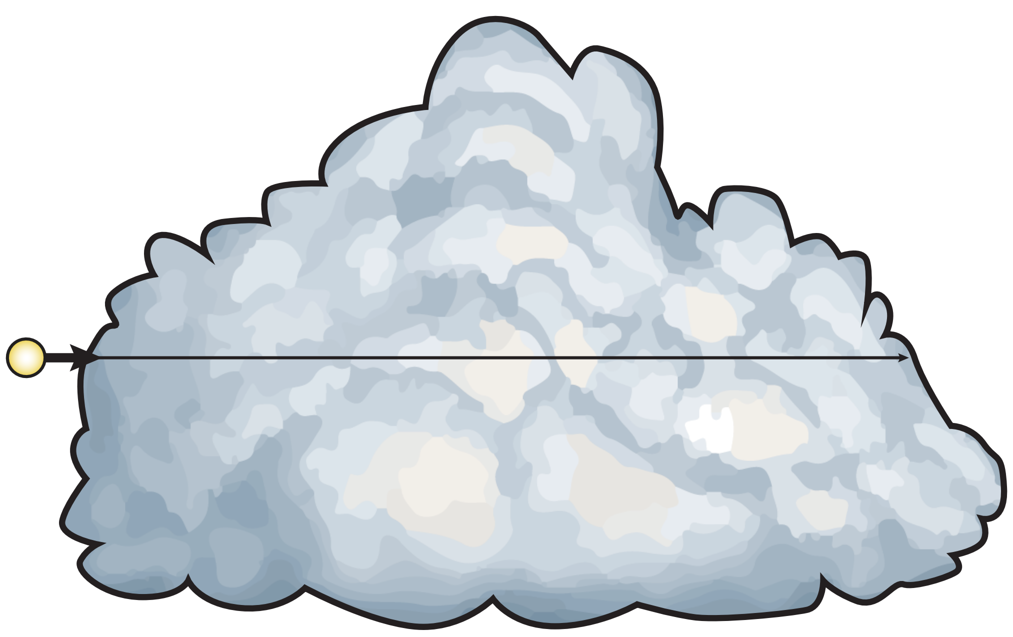

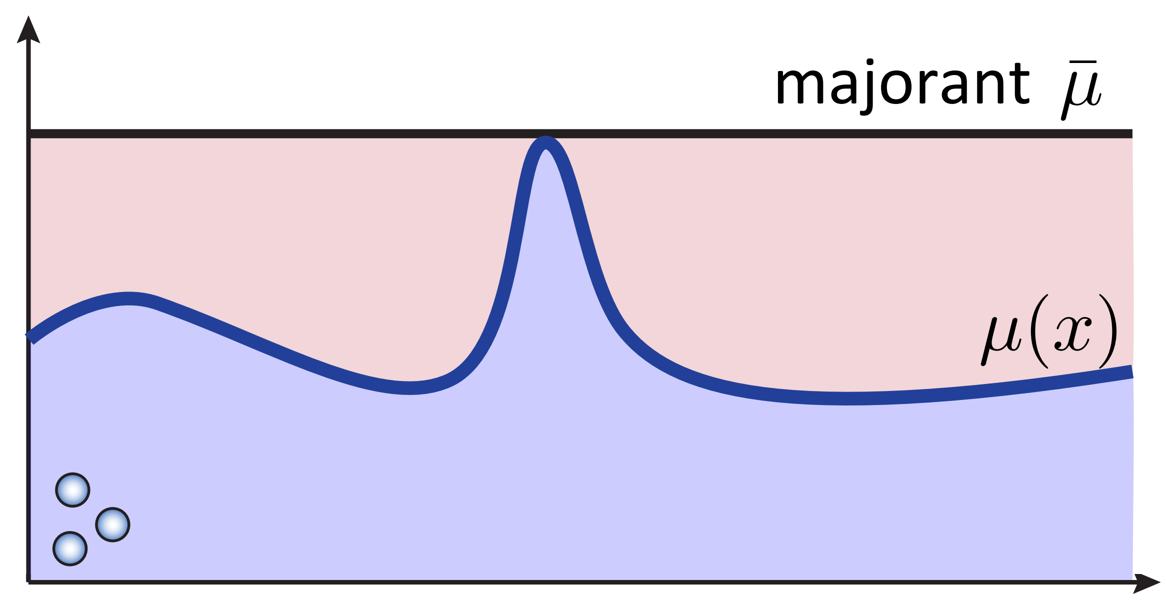

Please note, that the amount of fictitious particles impacts the efficiency of the algorithm.

When the majorant tightly bounds the extinction function, the algorithm is relatively efficient.

But when there are many fictitious particles, there will be many fictitious collisions,
and generating the free path is going to be expensive.

This is unfortunately quite common in practice, where the majorants are precomputed for entire volumes, or regions of the volume, and so it often happens that the majorant does not tightly bound the extinction along the ray.

Delta Tracking

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

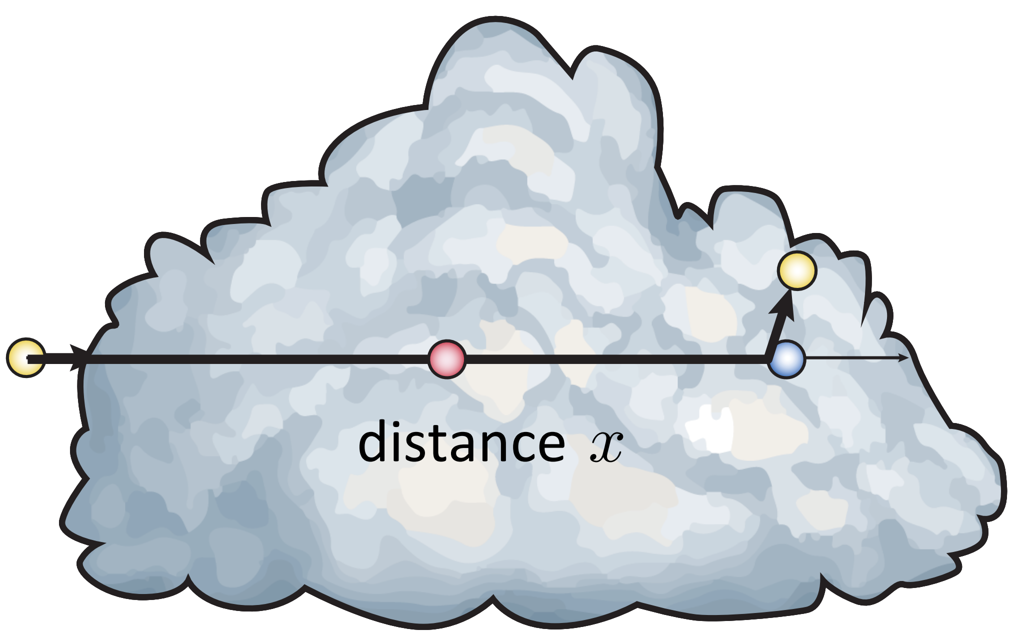

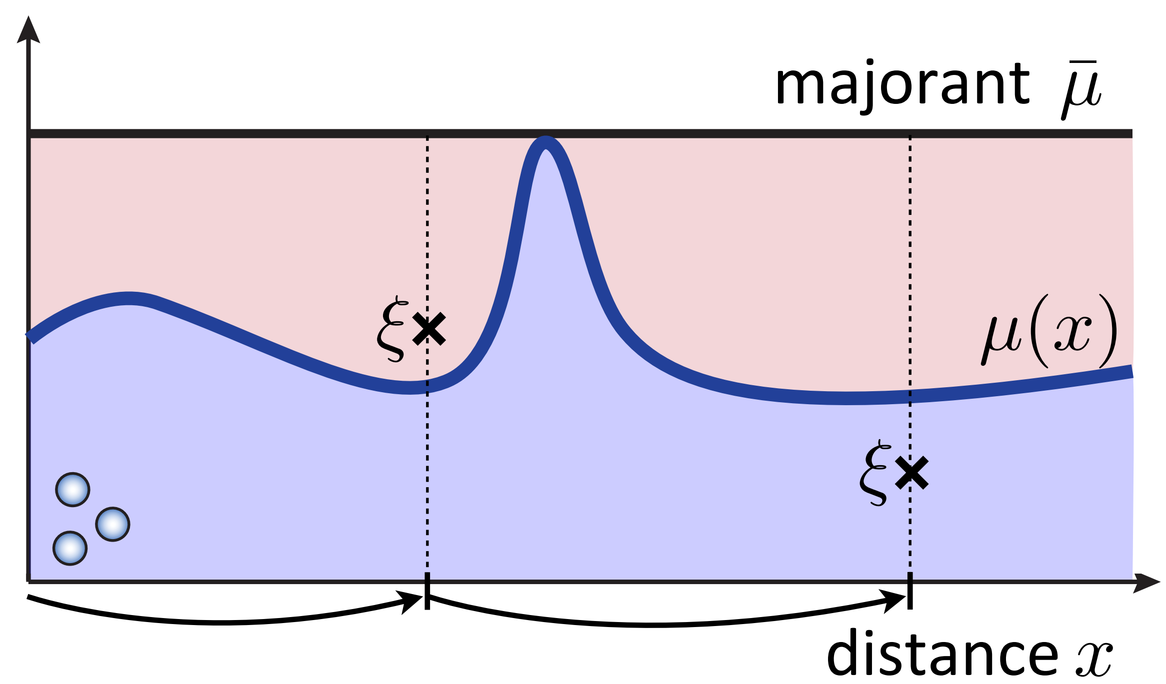

Please note, that the amount of fictitious particles impacts the efficiency of the algorithm.

When the majorant tightly bounds the extinction function, the algorithm is relatively efficient.

But when there are many fictitious particles, there will be many fictitious collisions,
and generating the free path is going to be expensive.

This is unfortunately quite common in practice, where the majorants are precomputed for entire volumes, or regions of the volume, and so it often happens that the majorant does not tightly bound the extinction along the ray.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

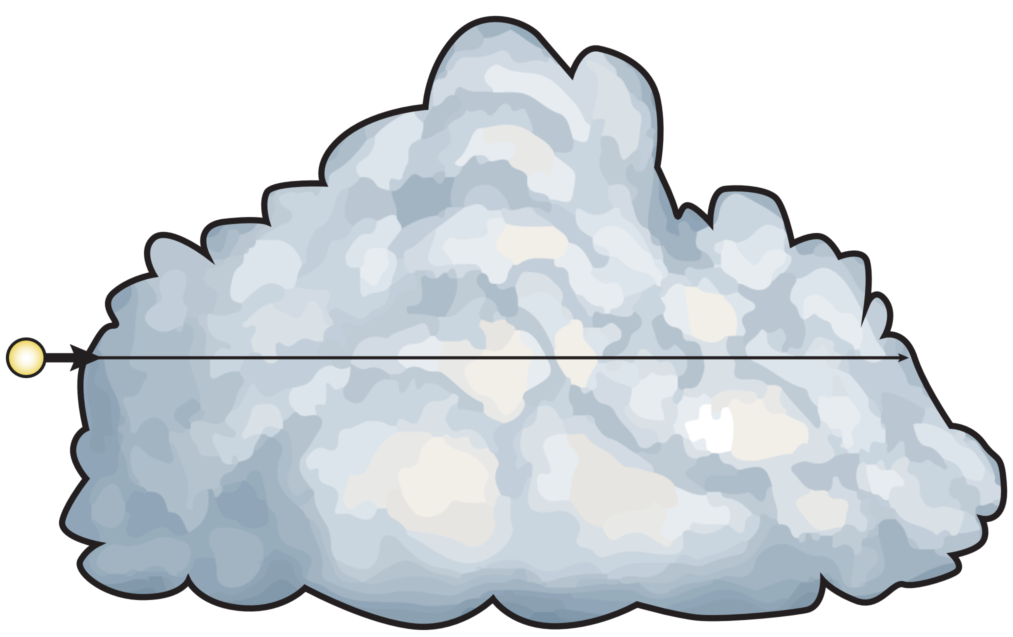

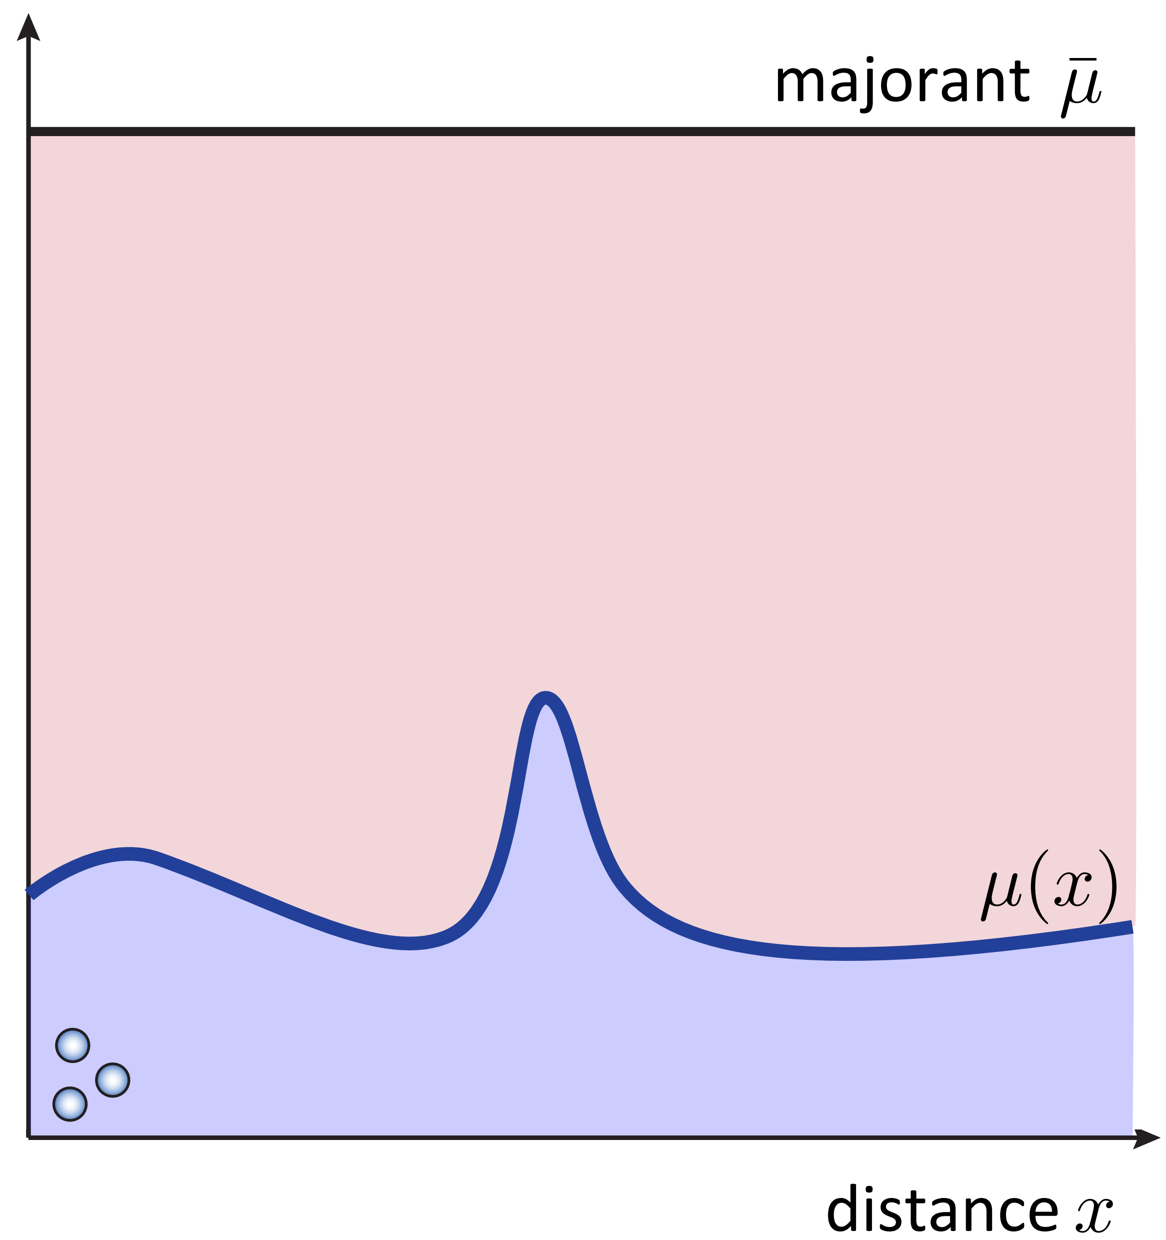

Please note, that the amount of fictitious particles impacts the efficiency of the algorithm.

When the majorant tightly bounds the extinction function, the algorithm is relatively efficient.

But when there are many fictitious particles, there will be many fictitious collisions,
and generating the free path is going to be expensive.

This is unfortunately quite common in practice, where the majorants are precomputed for entire volumes, or regions of the volume, and so it often happens that the majorant does not tightly bound the extinction along the ray.

Delta Tracking

SIGGRAPH ASIA 2014 SHENZHEN

- A.k.a. **Woodcock** tracking, pseudo-scattering, null-collision algorithm
- Add **fictitious particles** that:
  - simplify distance sampling, but
  - do not impact the light transport

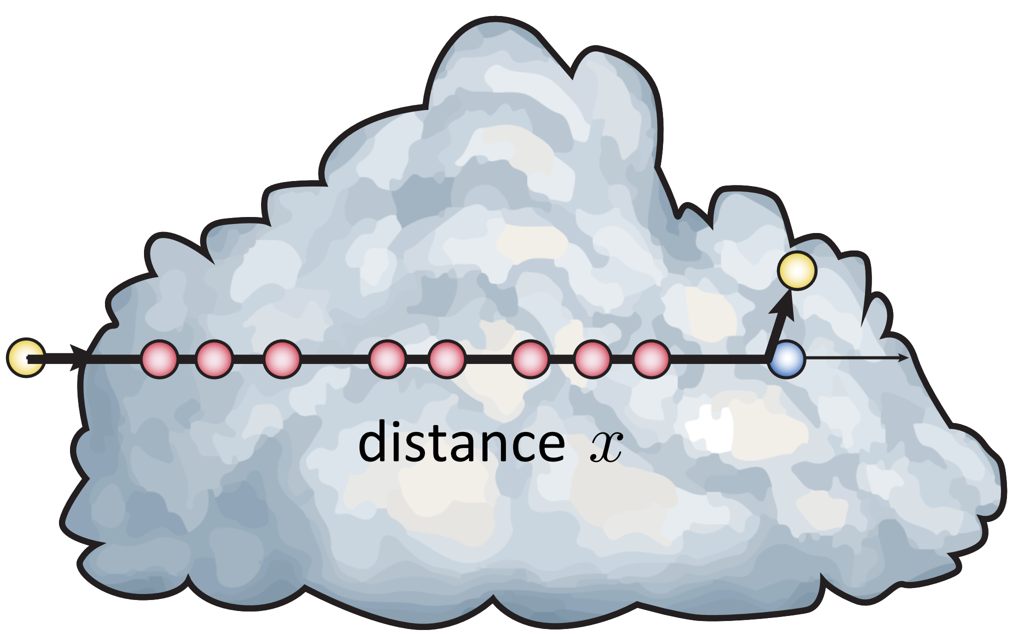

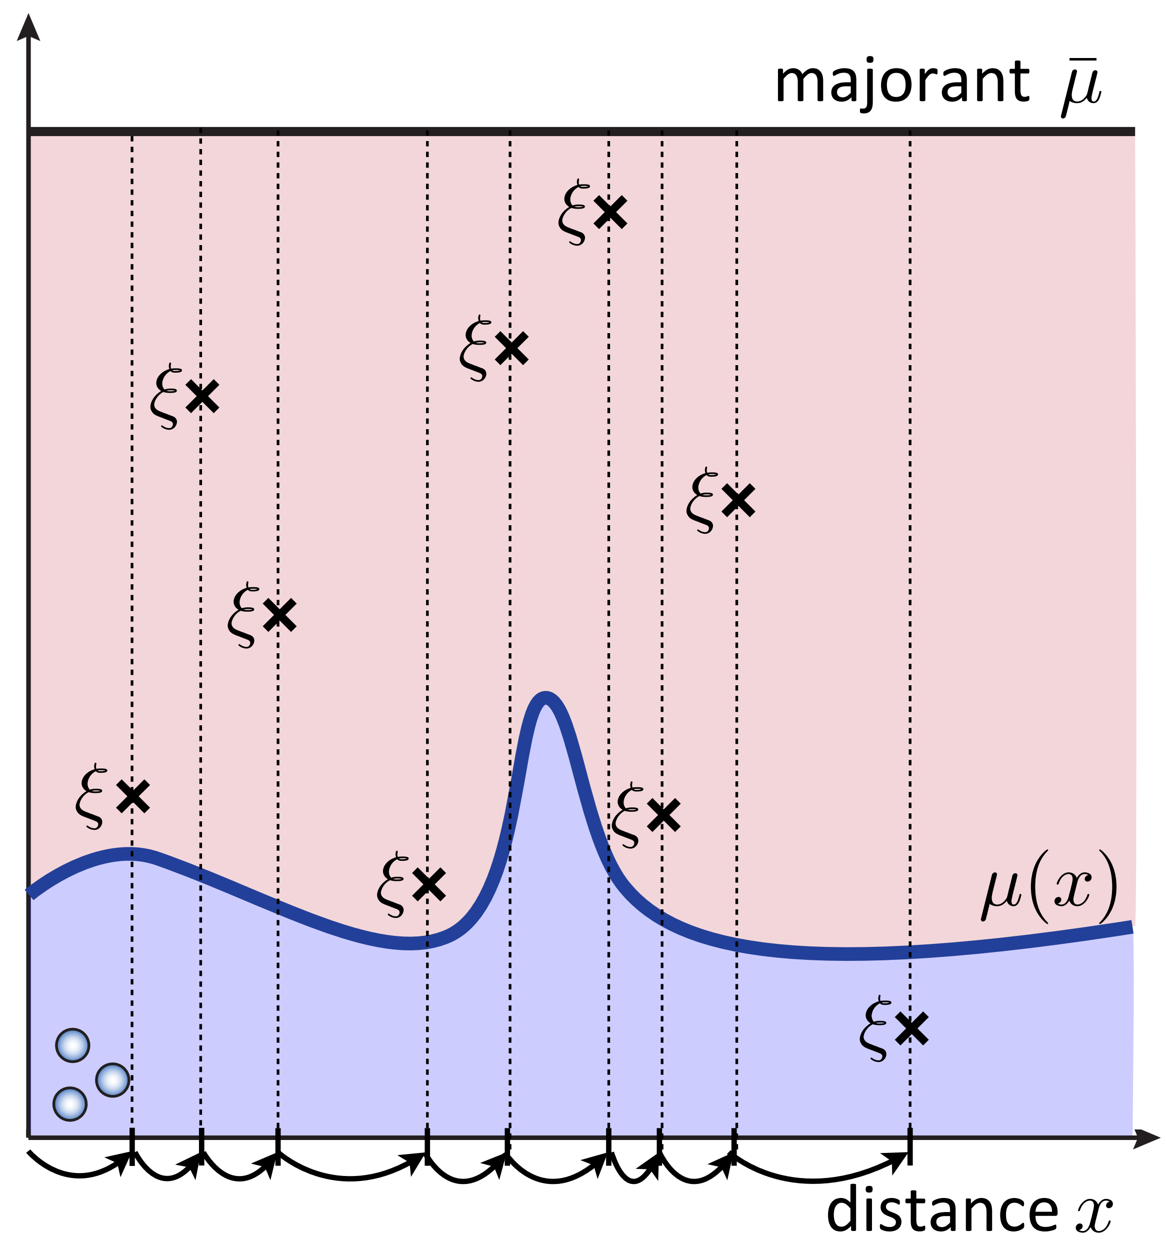

Please note, that the amount of fictitious particles impacts the efficiency of the algorithm.

When the majorant tightly bounds the extinction function, the algorithm is relatively efficient.

But when there are many fictitious particles, there will be many fictitious collisions,
and generating the free path is going to be expensive.

This is unfortunately quite common in practice, where the majorants are precomputed for entire volumes, or regions of the volume, and so it often happens that the majorant does not tightly bound the extinction along the ray.

Delta Tracking & Transmittance

SIGGRAPH ASIA 2014 SHENZHEN

In addition to sampling paths, the tracking can also be used for estimating transmittance between points. The idea is to start the tracking from one point and see if it reaches the second one. This happens when only fictitious collisions occur, then transmittance is estimated as 1.

Shall there be a real collision somewhere along the line, the tracking is terminated, and transmittance set to 0. You see that the estimation is binary, and will suffer from high variance.

Delta Tracking & Transmittance

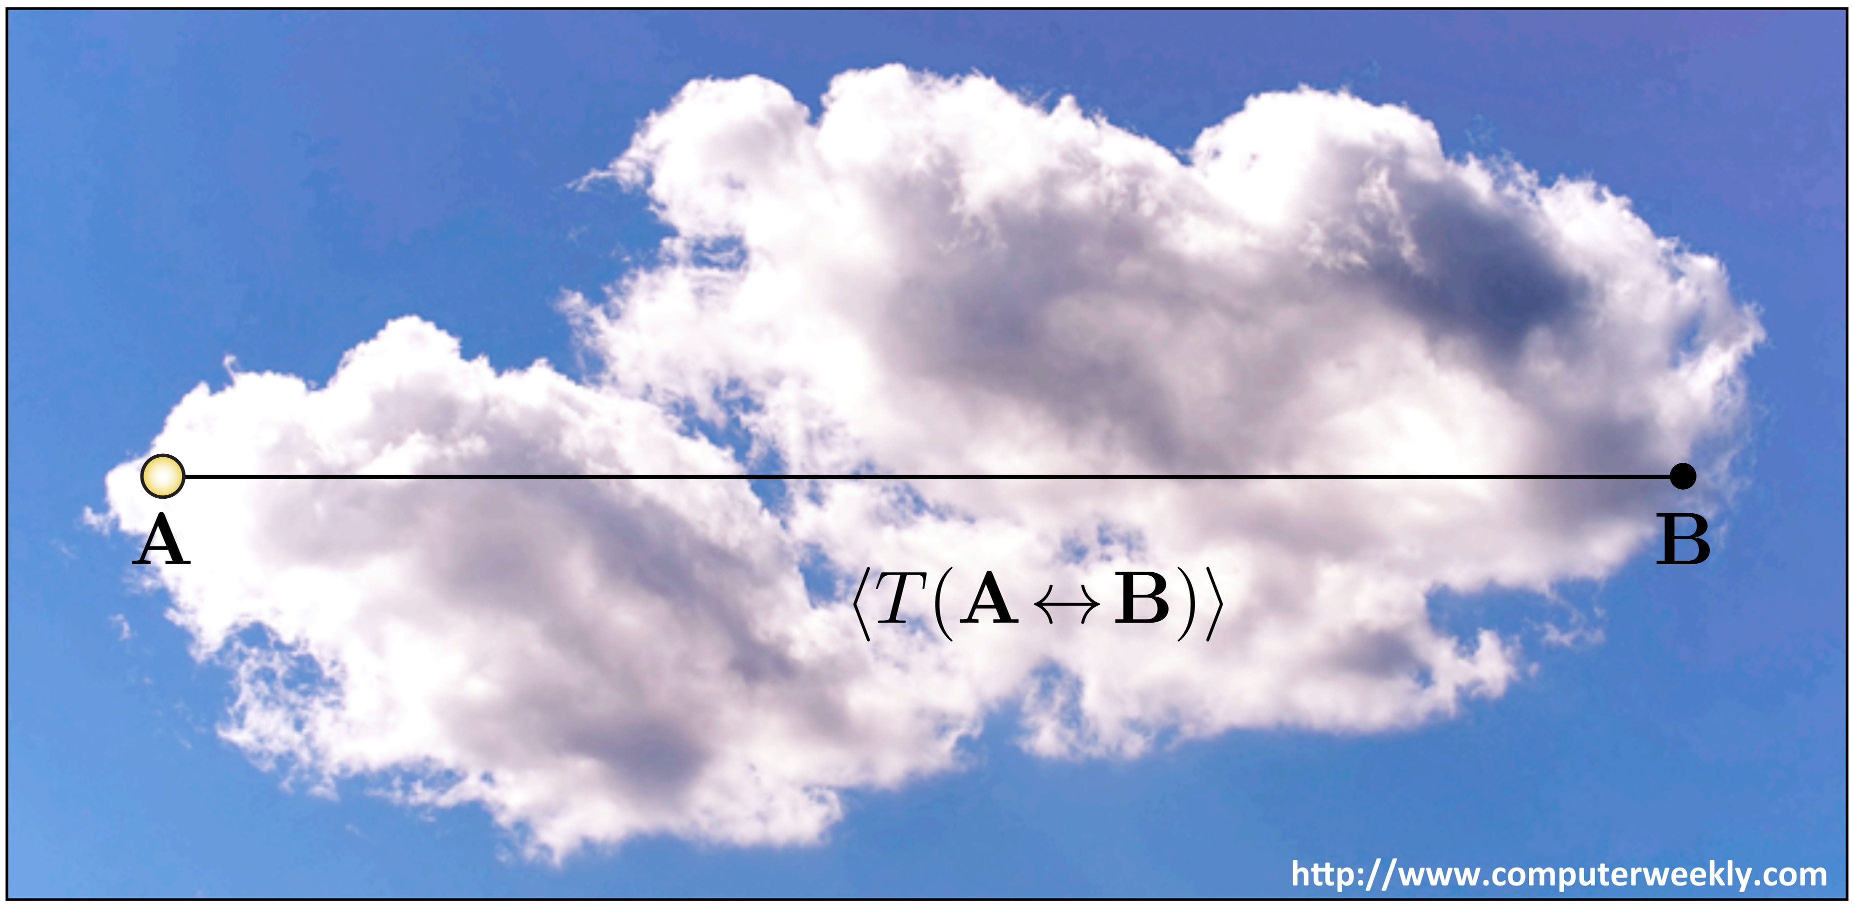

In addition to sampling paths, the tracking can also be used for estimating transmittance between points. The idea is to start the tracking from one point and see if it reaches the second one. This happens when only fictitious collisions occur, then transmittance is estimated as 1.

Shall there be a real collision somewhere along the line, the tracking is terminated, and transmittance set to 0. You see that the estimation is binary, and will suffer from high variance.

Delta Tracking & Transmittance

SIGGRAPH ASIA 2014 SHENZHEN

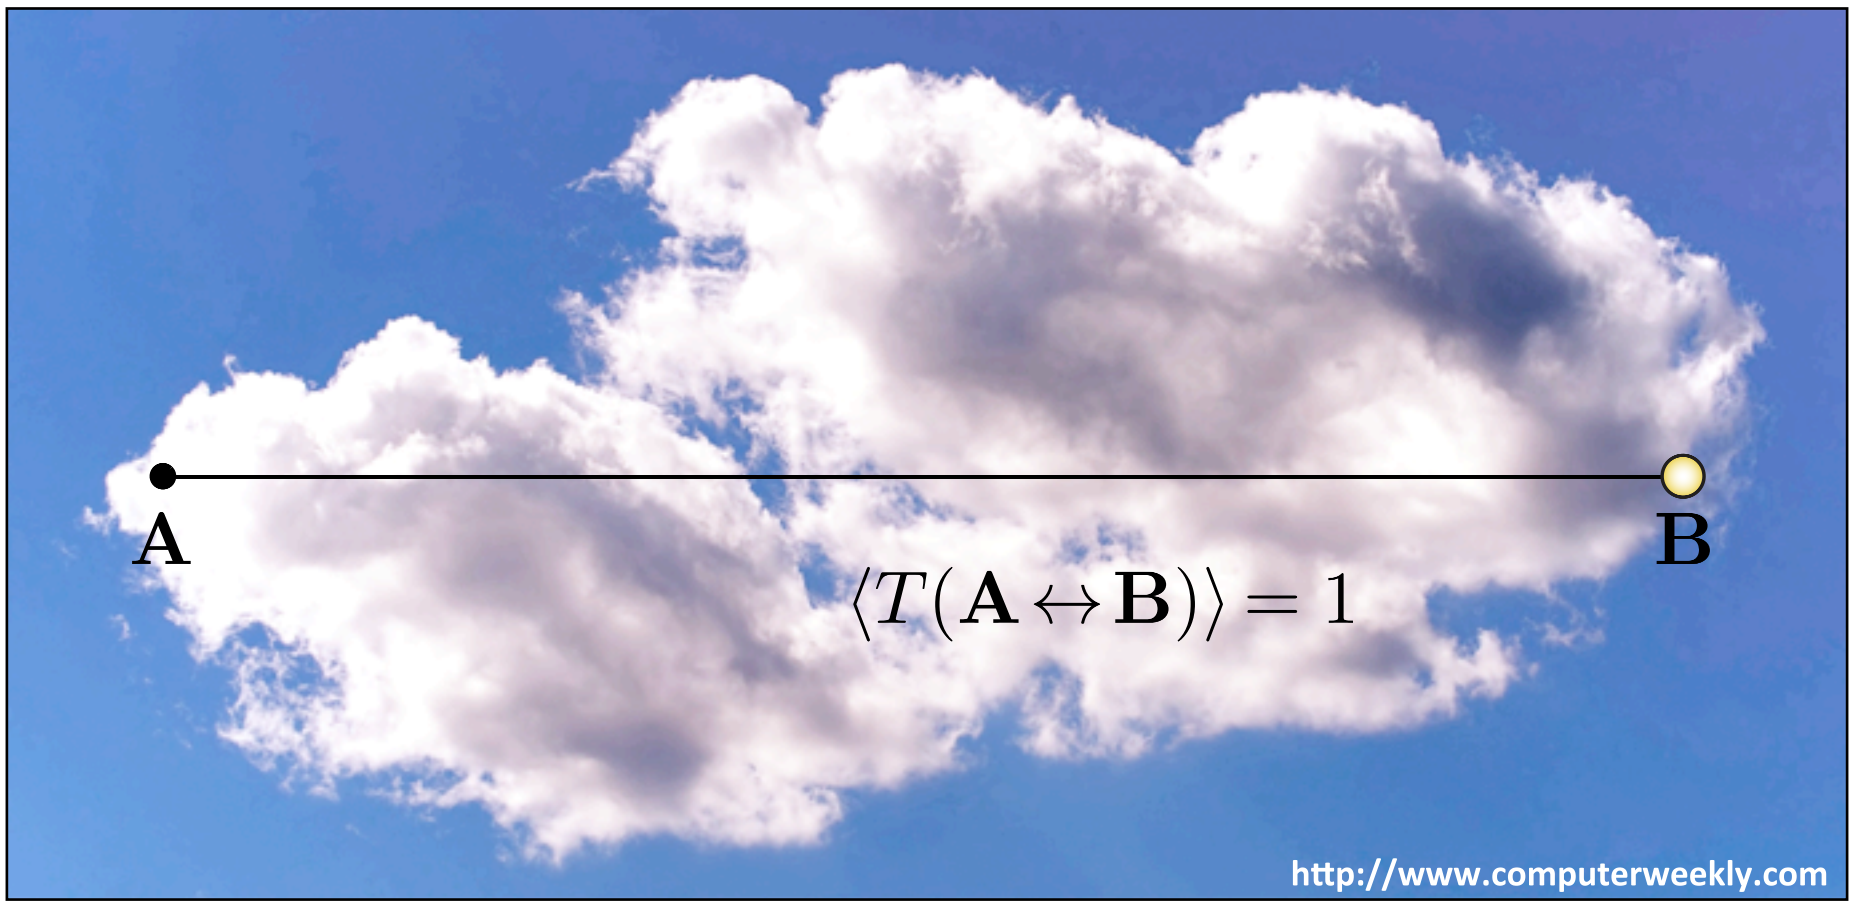

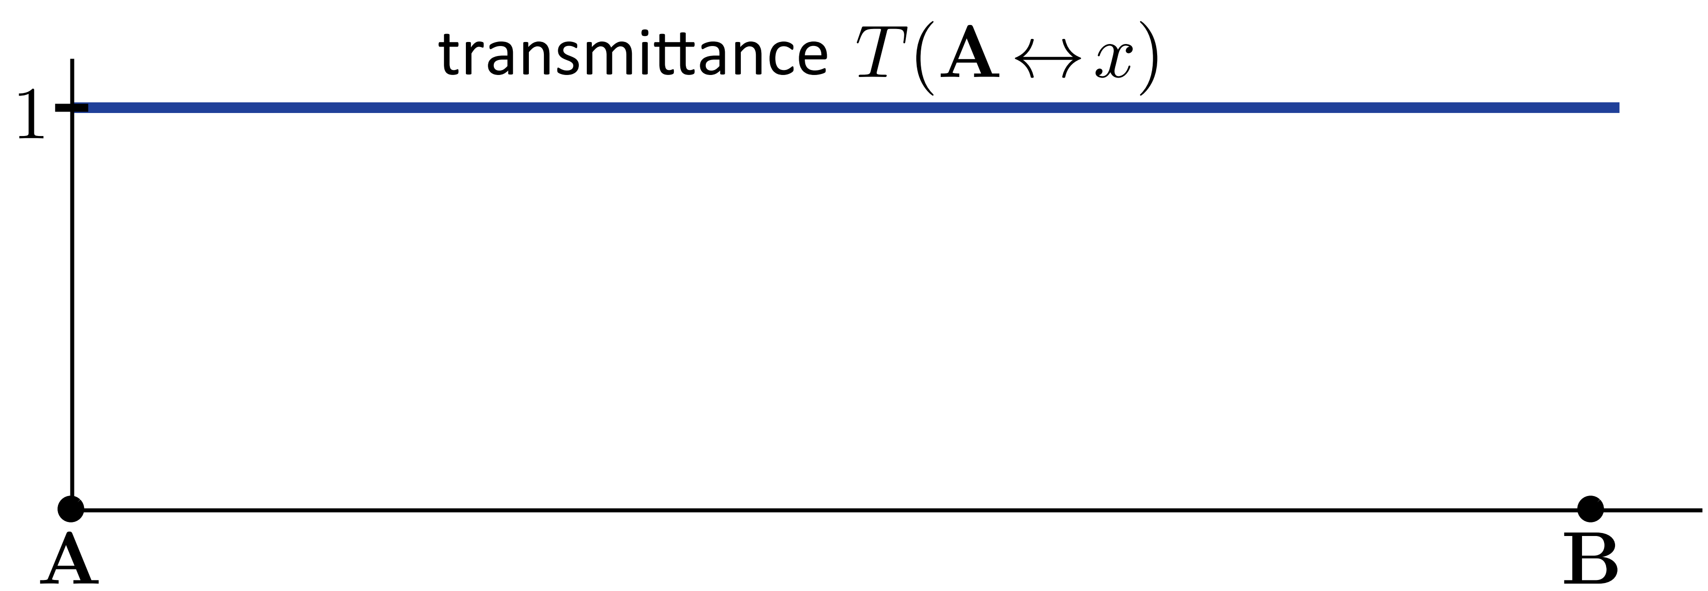

In addition to sampling paths, the tracking can also be used for estimating transmittance between points. The idea is to start the tracking from one point and see if it reaches the second one. This happens when only fictitious collisions occur, then transmittance is estimated as 1.

Shall there be a real collision somewhere along the line, the tracking is terminated, and transmittance set to 0. You see that the estimation is binary, and will suffer from high variance.

Delta Tracking & Transmittance

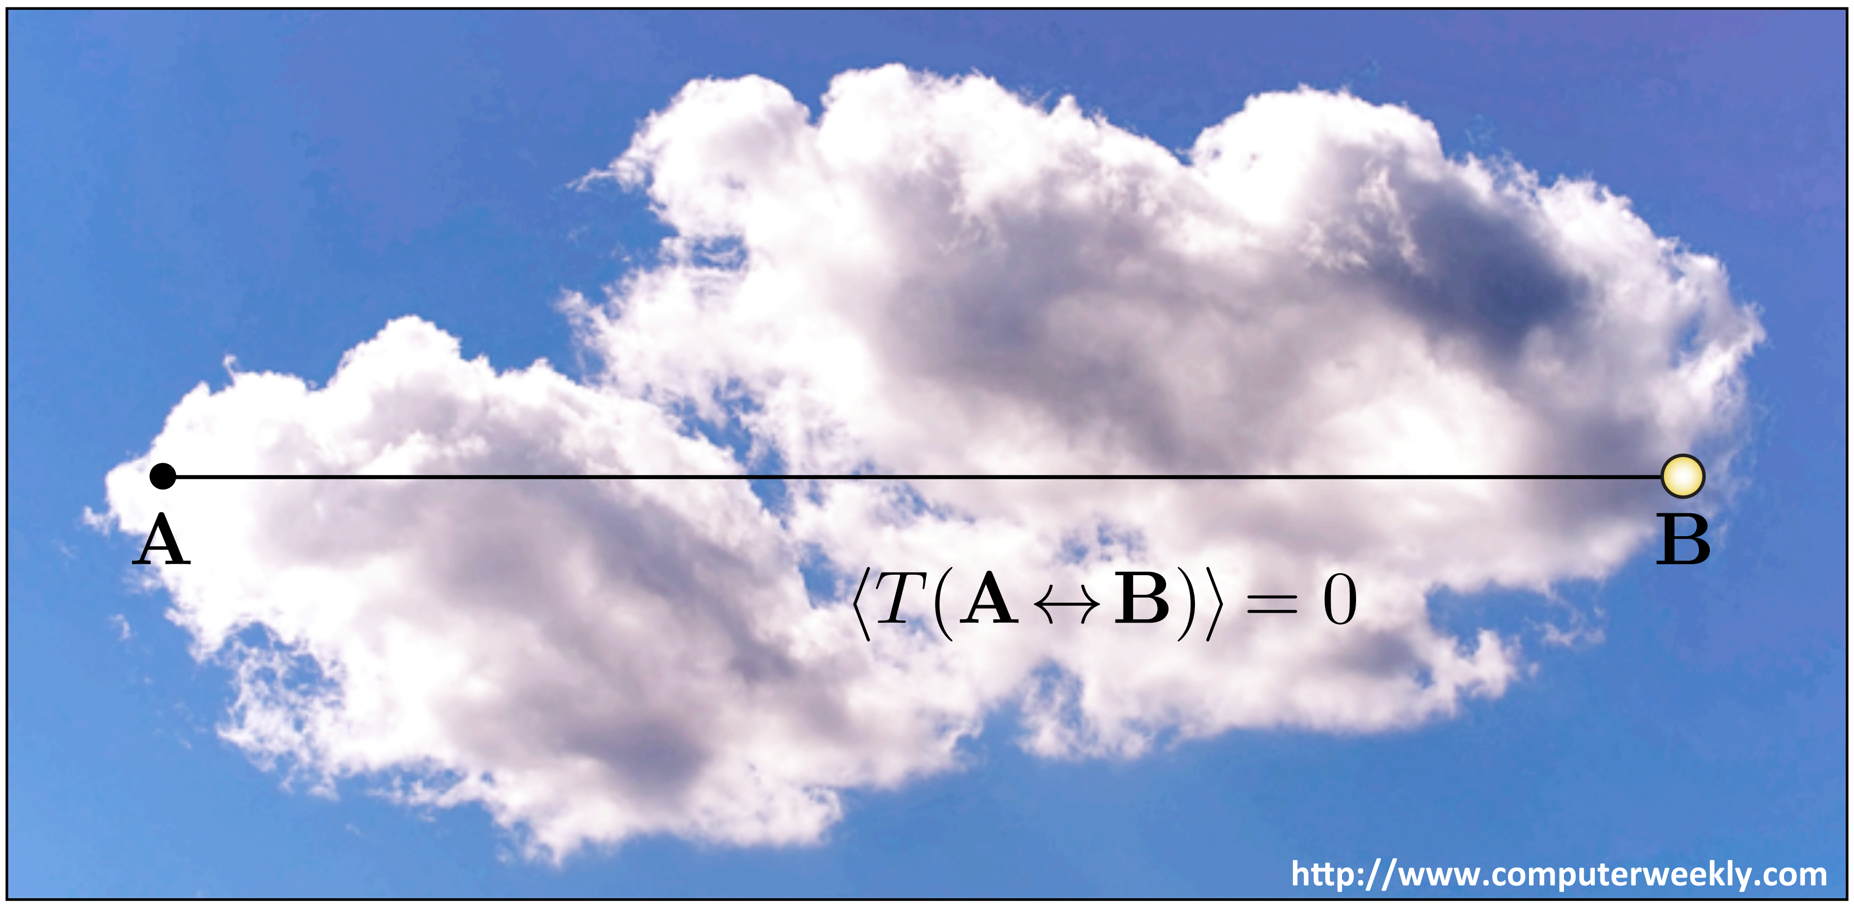

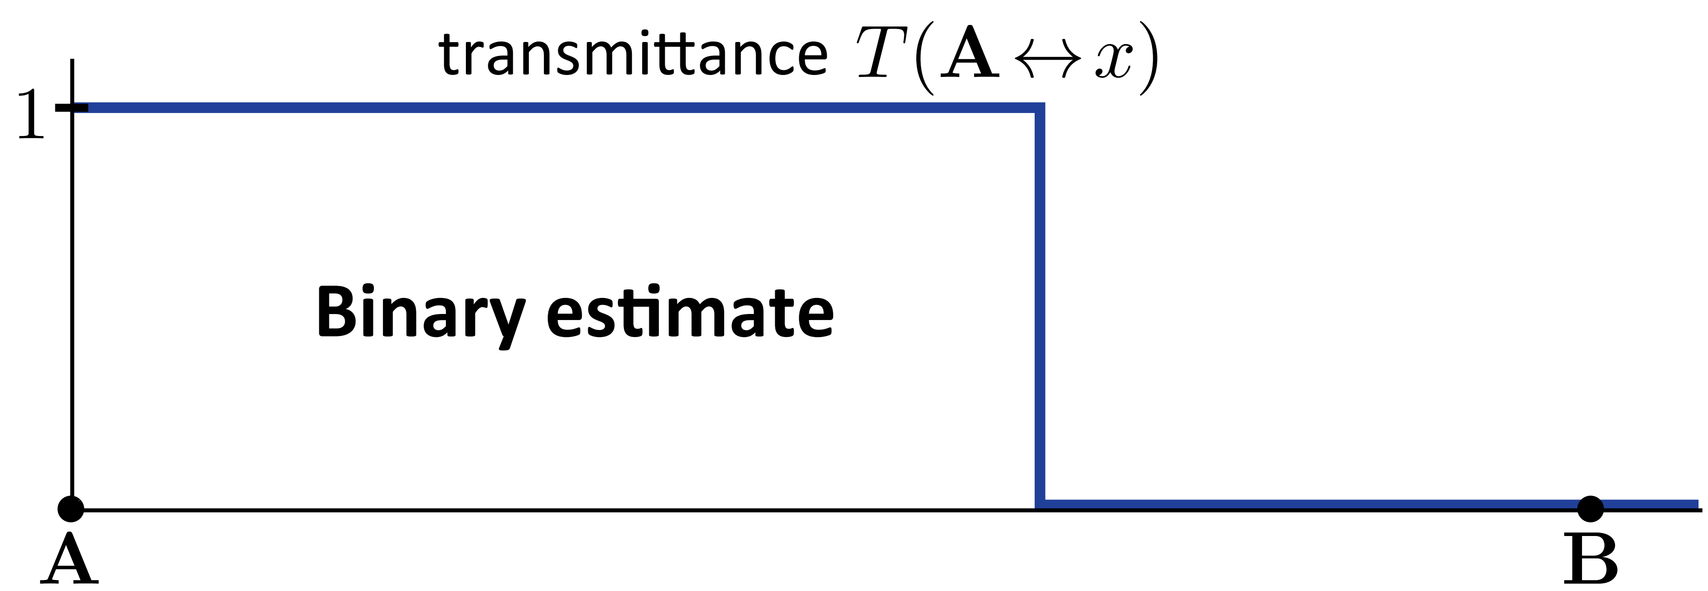

In addition to sampling paths, the tracking can also be used for estimating transmittance between points. The idea is to start the tracking from one point and see if it reaches the second one. This happens when only fictitious collisions occur, then transmittance is estimated as 1.

Shall there be a real collision somewhere along the line, the tracking is terminated, and transmittance set to 0. You see that the estimation is binary, and will suffer from high variance.

Delta Tracking with Loose Majorants

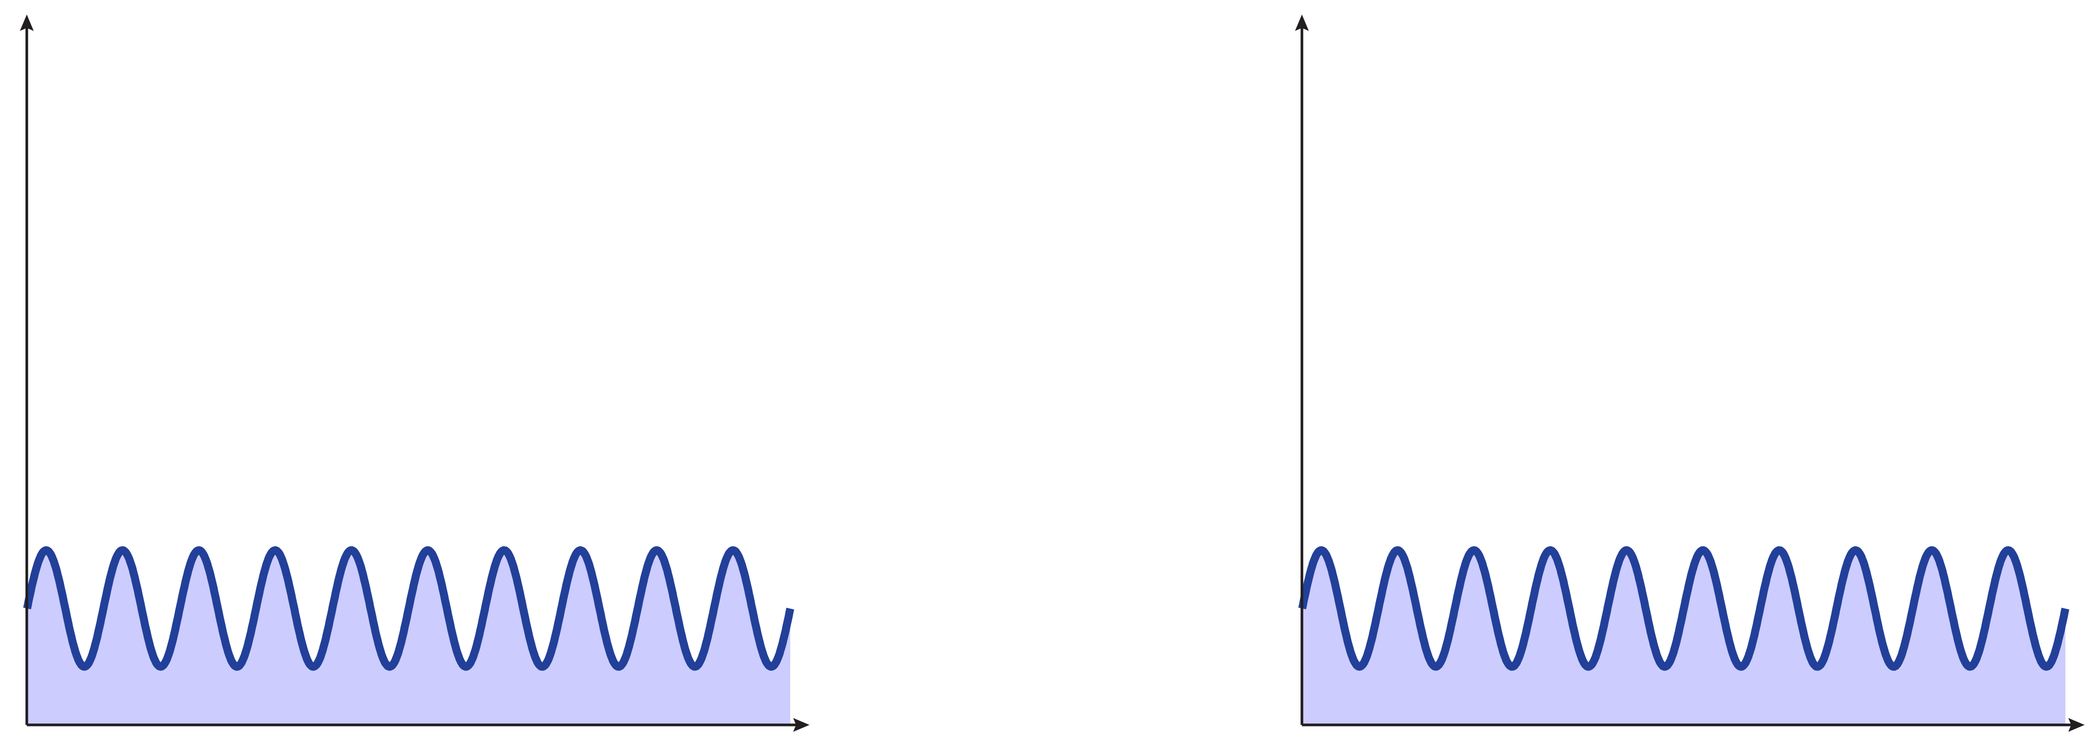

Things get even worse when the majorant is loose, as I mentioned previously, loose majorants increase the cost of the tracking significantly.

You can see this on the right-hand side, where the tracking takes many steps, but yields the same binary estimate.

The question now is,...   can we do better? Can we somehow utilize the steps more intelligently than deducing just a binary answer?

Delta Tracking with Loose Majorants

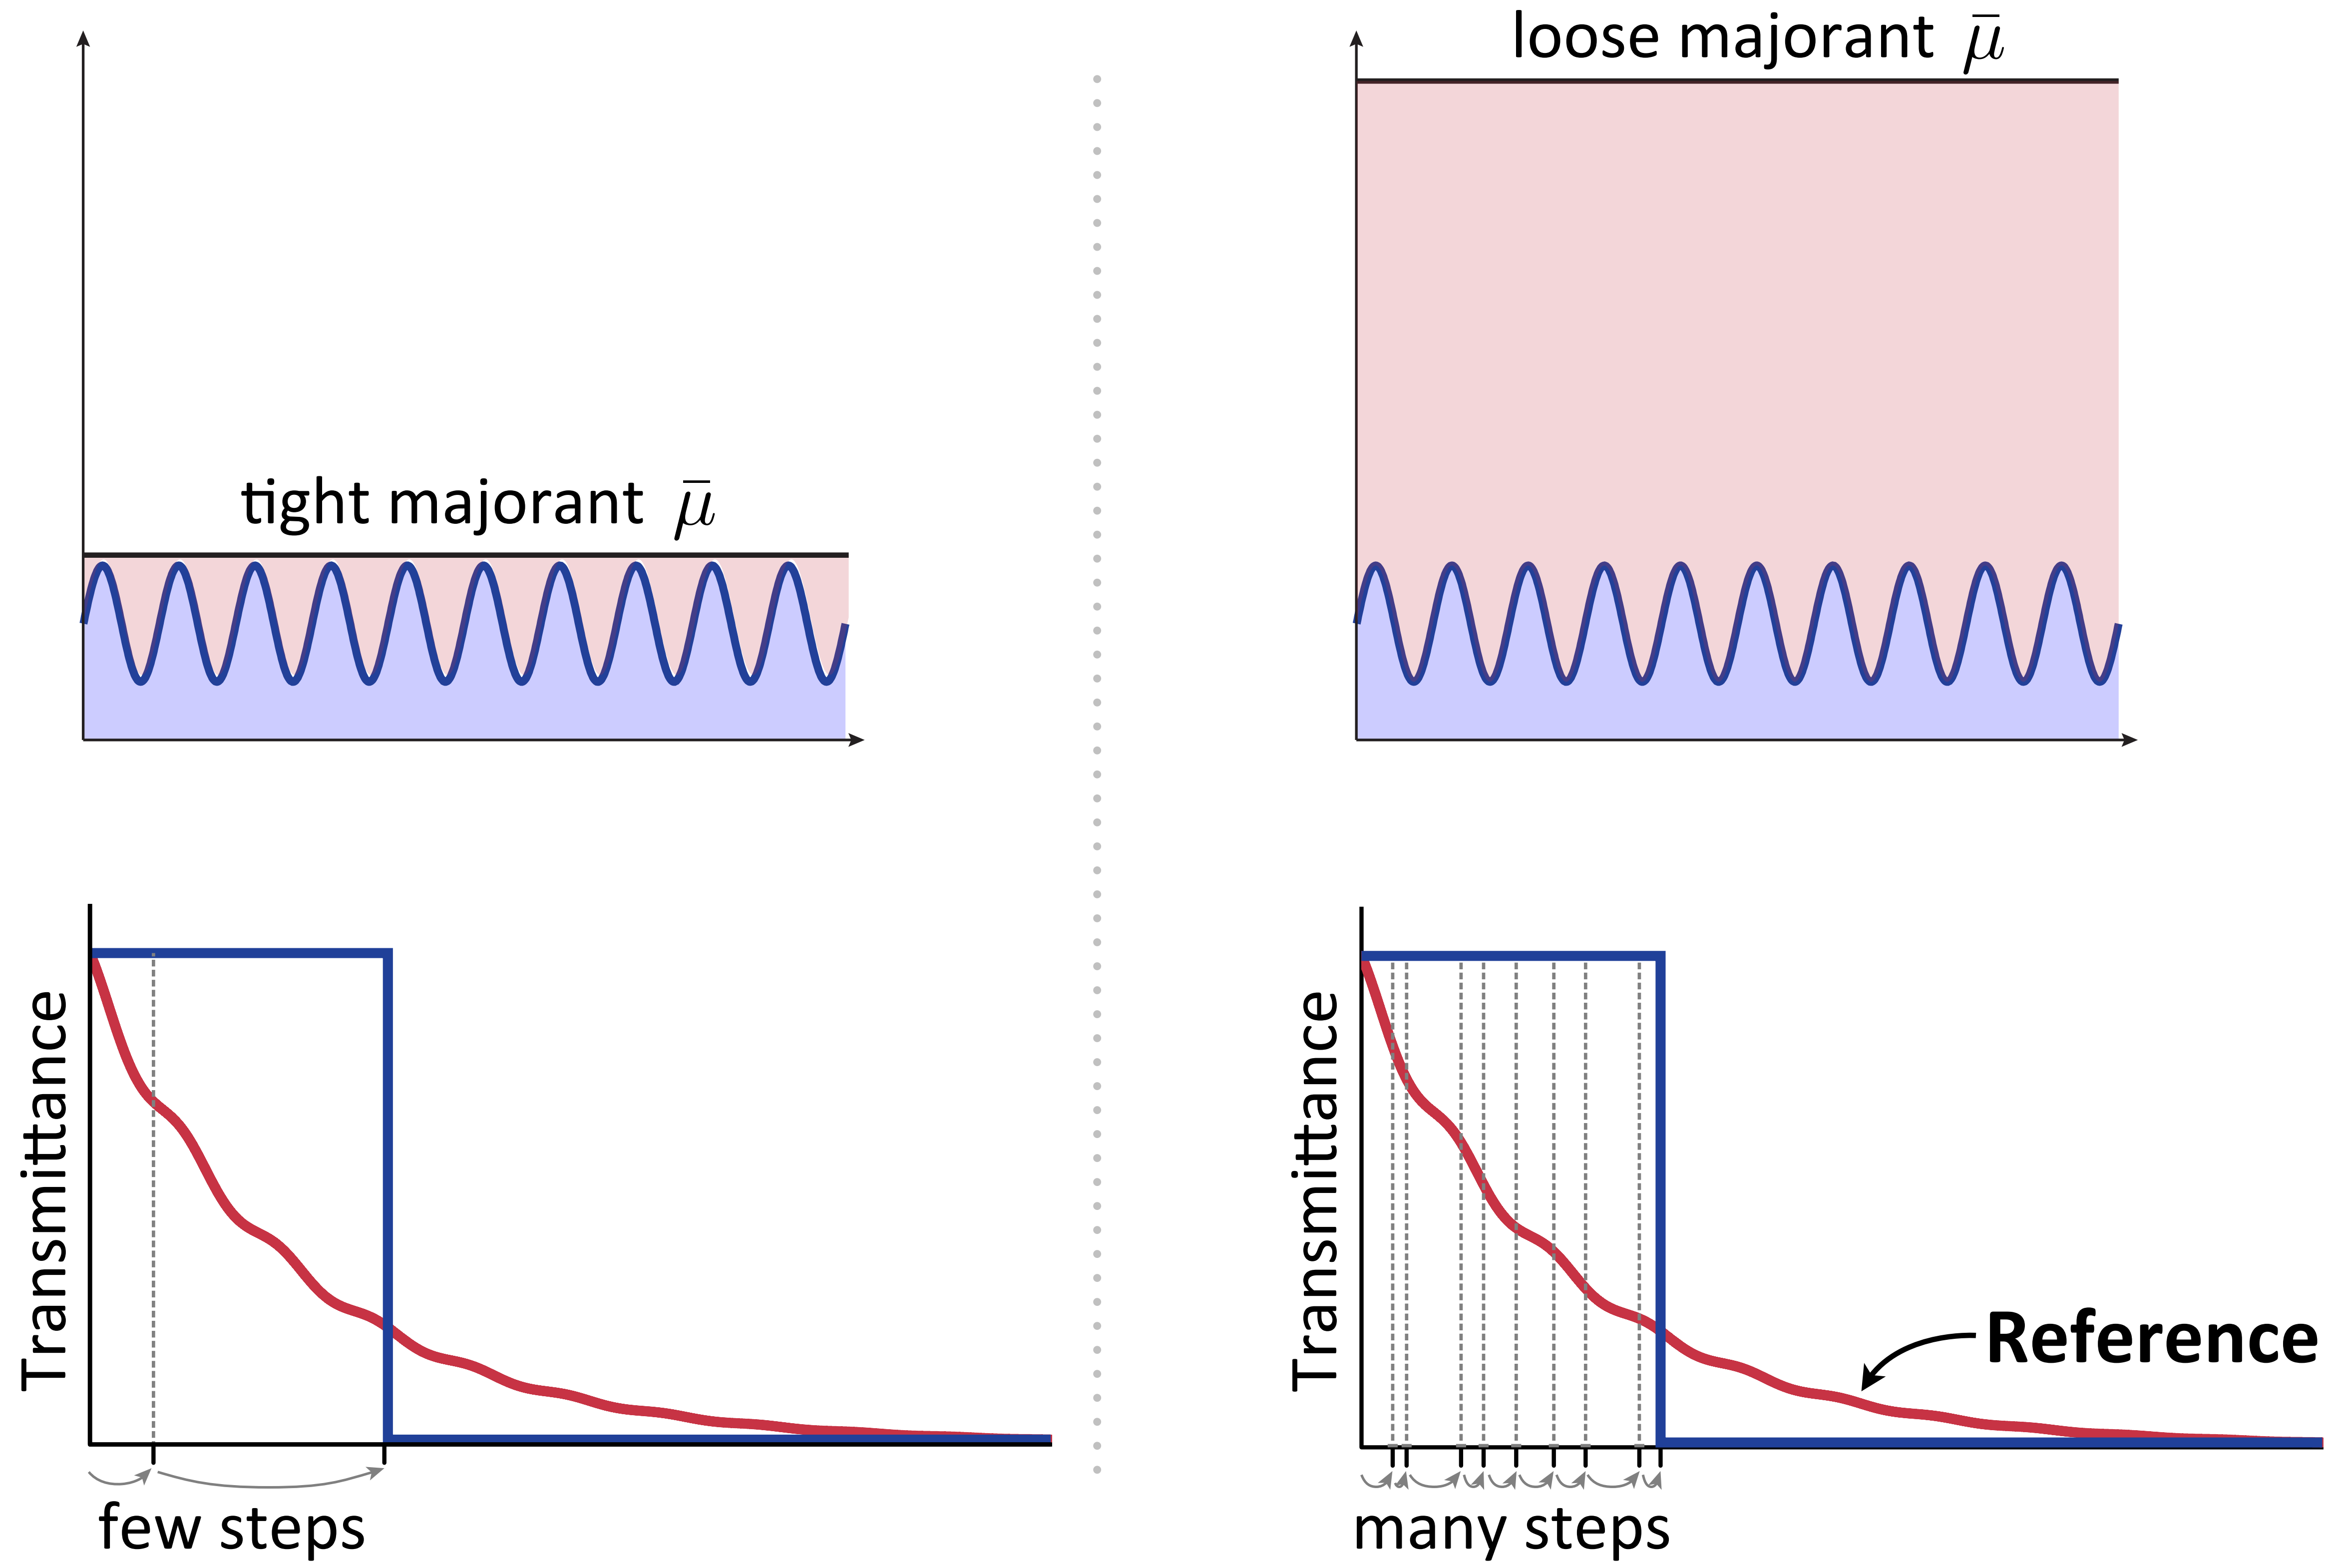

Things get even worse when the majorant is loose, as I mentioned previously, loose majorants increase the cost of the tracking significantly.

You can see this on the right-hand side, where the tracking takes many steps, but yields the same binary estimate.

The question now is,... can we do better? Can we somehow utilize the steps more intelligently than deducing just a binary answer?

Delta Tracking with Loose Majorants

SIGGRAPH ASIA 2014 SHENZHEN

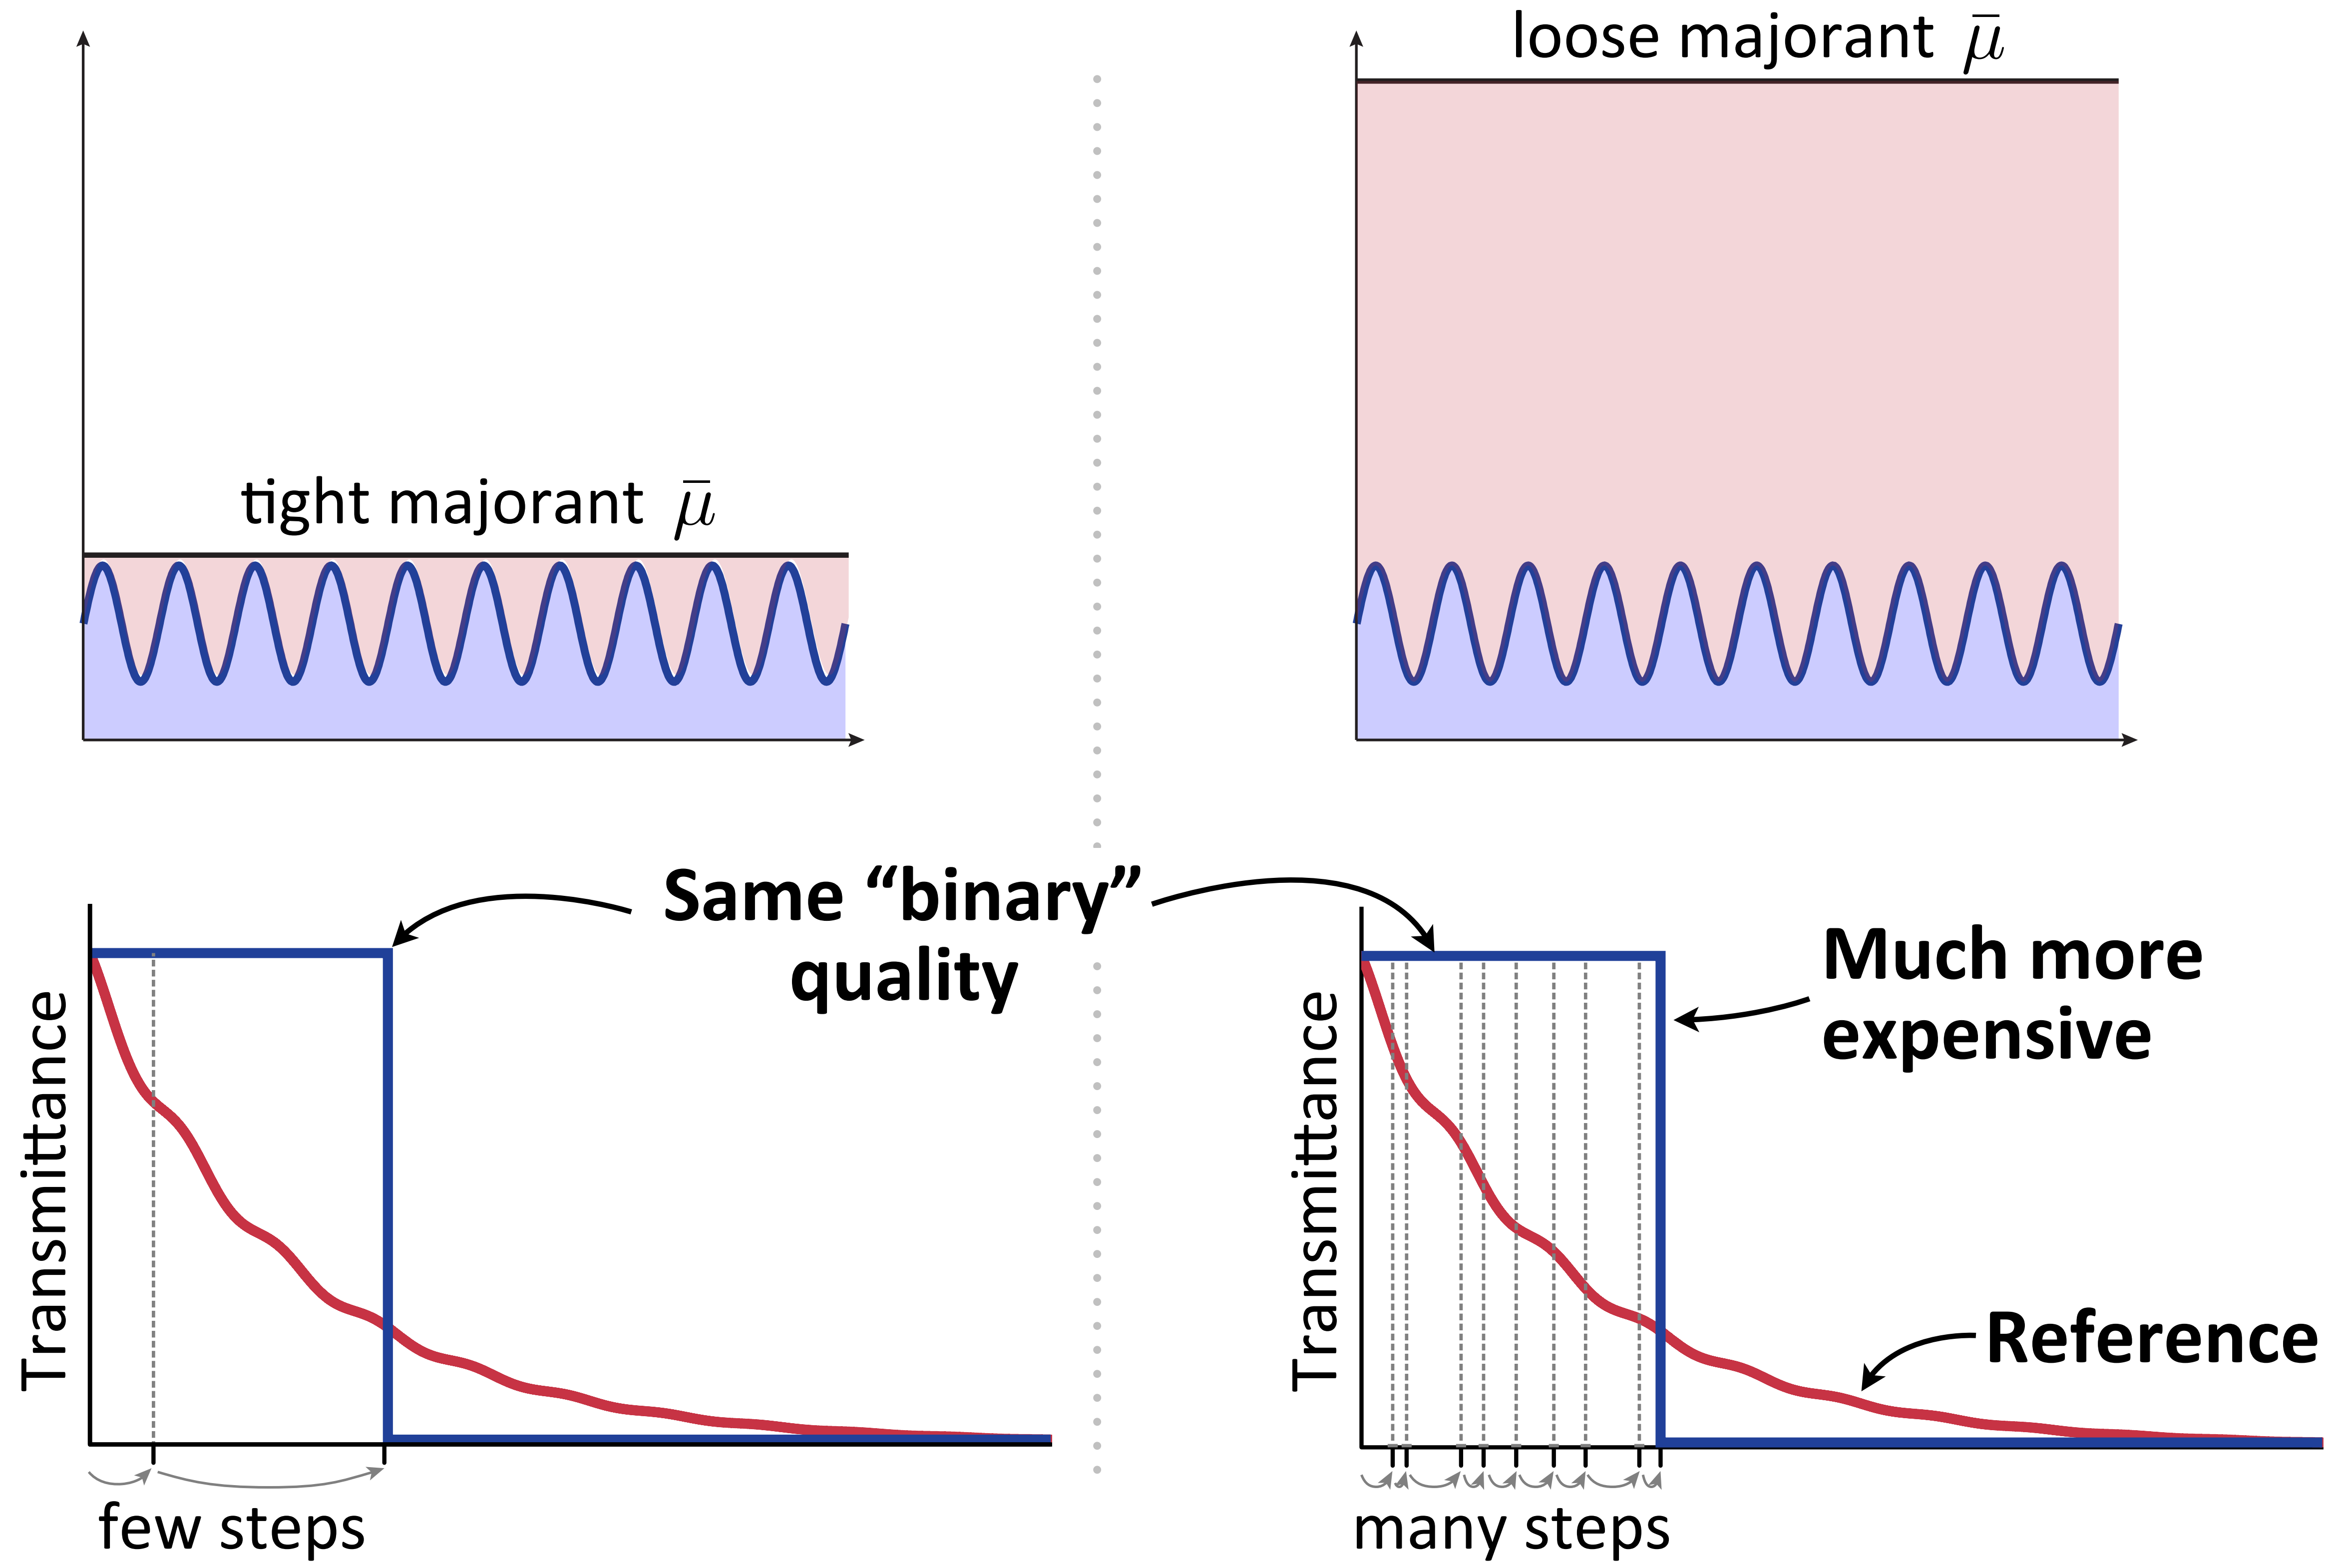

Things get even worse when the majorant is loose, as I mentioned previously, loose majorants increase the cost of the tracking significantly.

You can see this on the right-hand side, where the tracking takes many steps, but yields the same binary estimate.

The question now is,... can we do better? Can we somehow utilize the steps more intelligently than deducing just a binary answer?

Ratio Tracking

“from binary to piecewise constant”

The first technique that I’ll present aims at this problem.

We call it ratio tracking and it’s goal is to replace the binary estimate with a piece-wise constant approximation, without increasing the cost significantly.

Ratio Tracking

Let’s quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

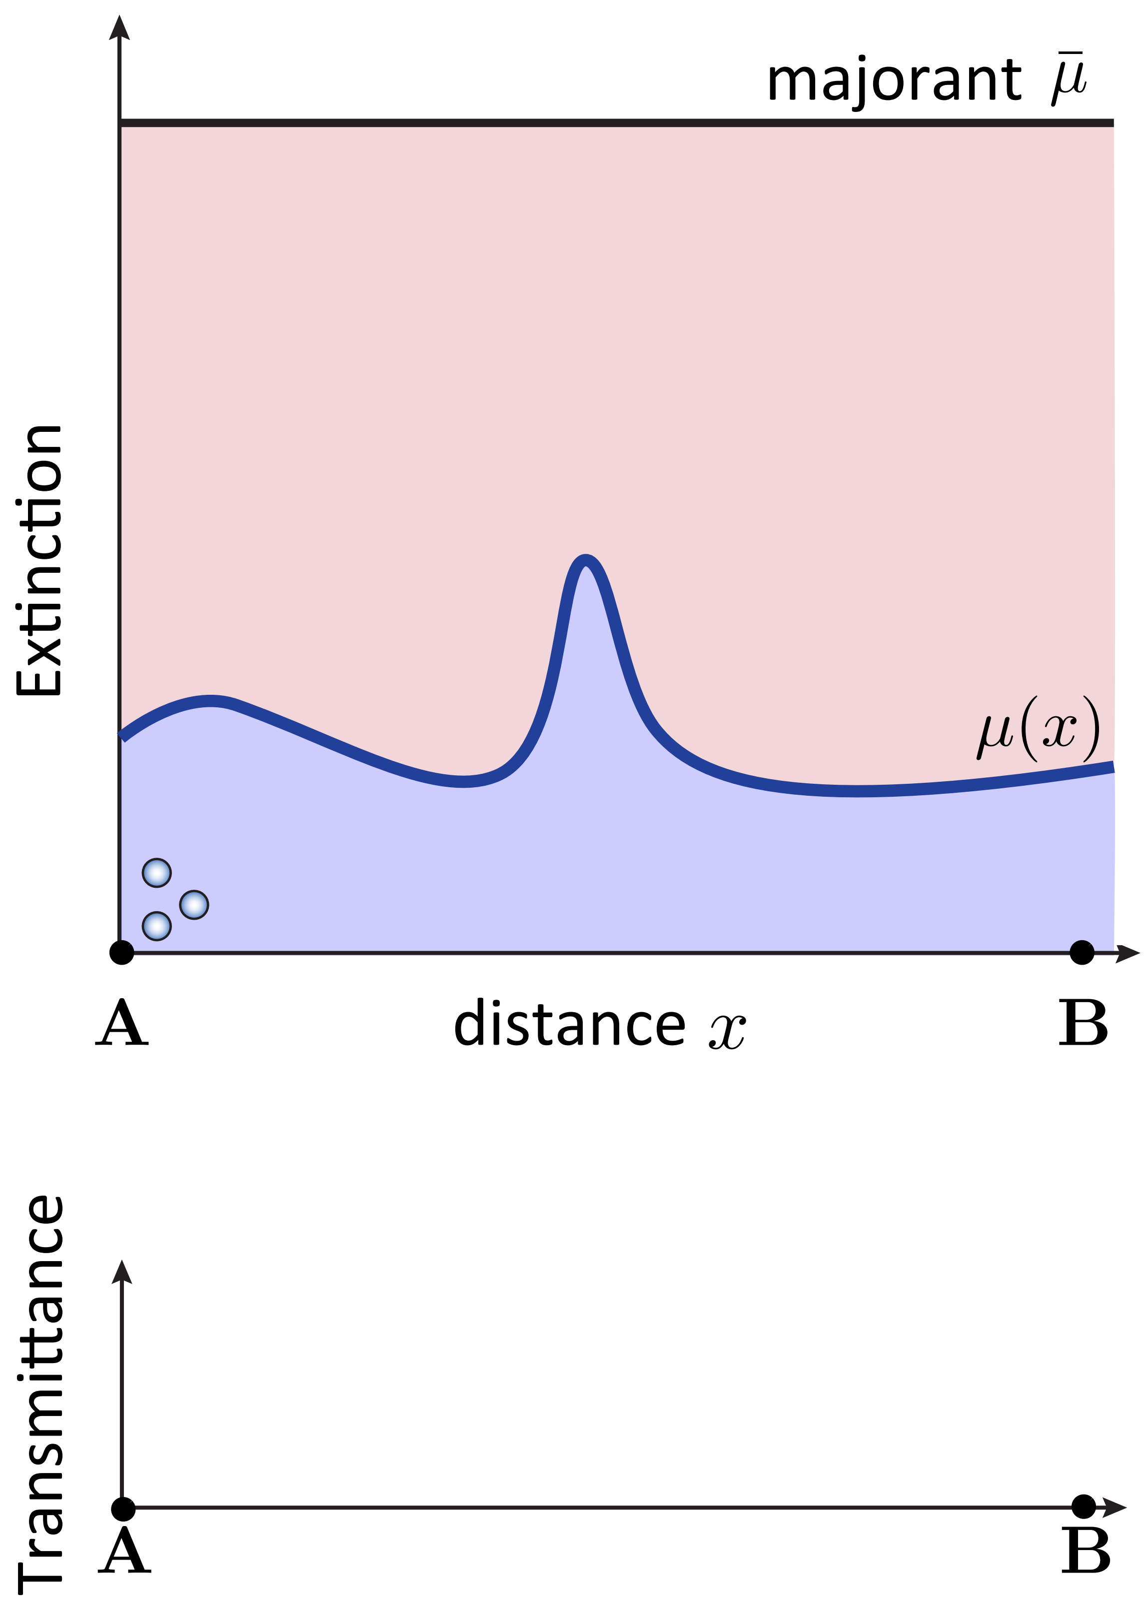

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

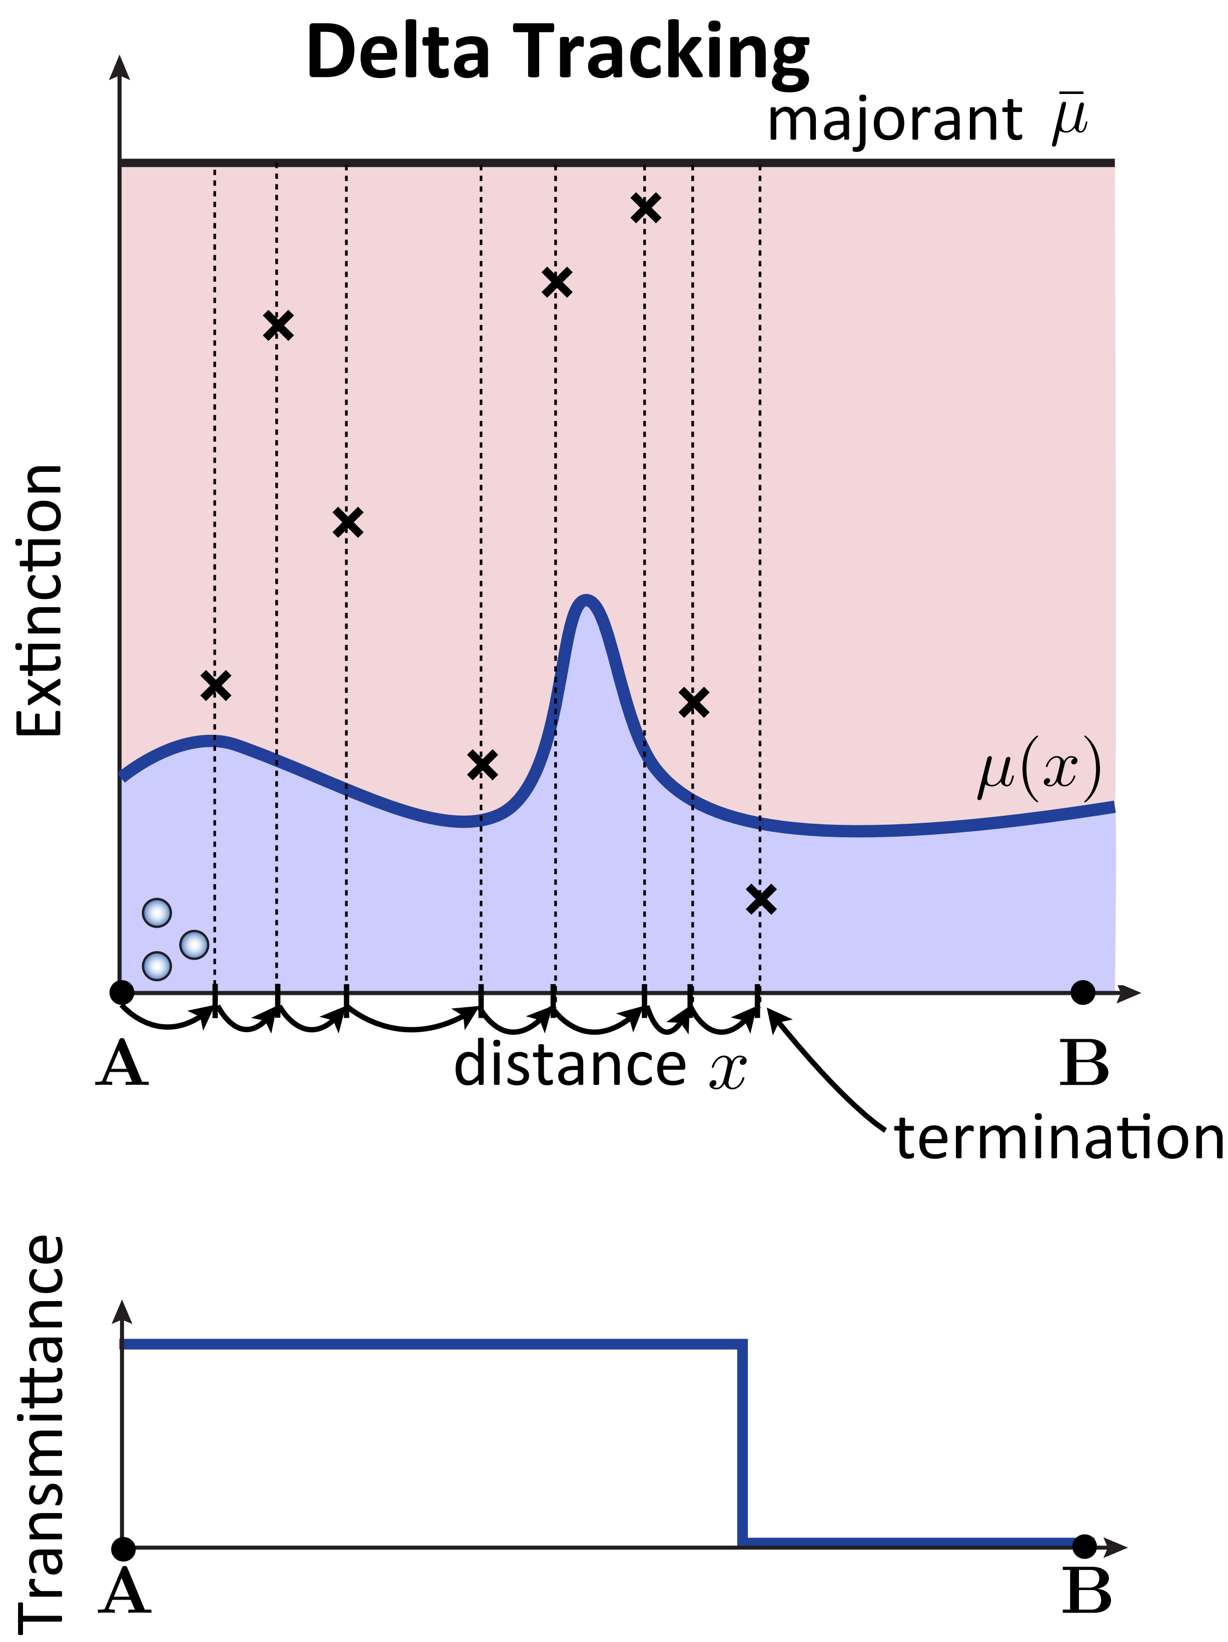

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

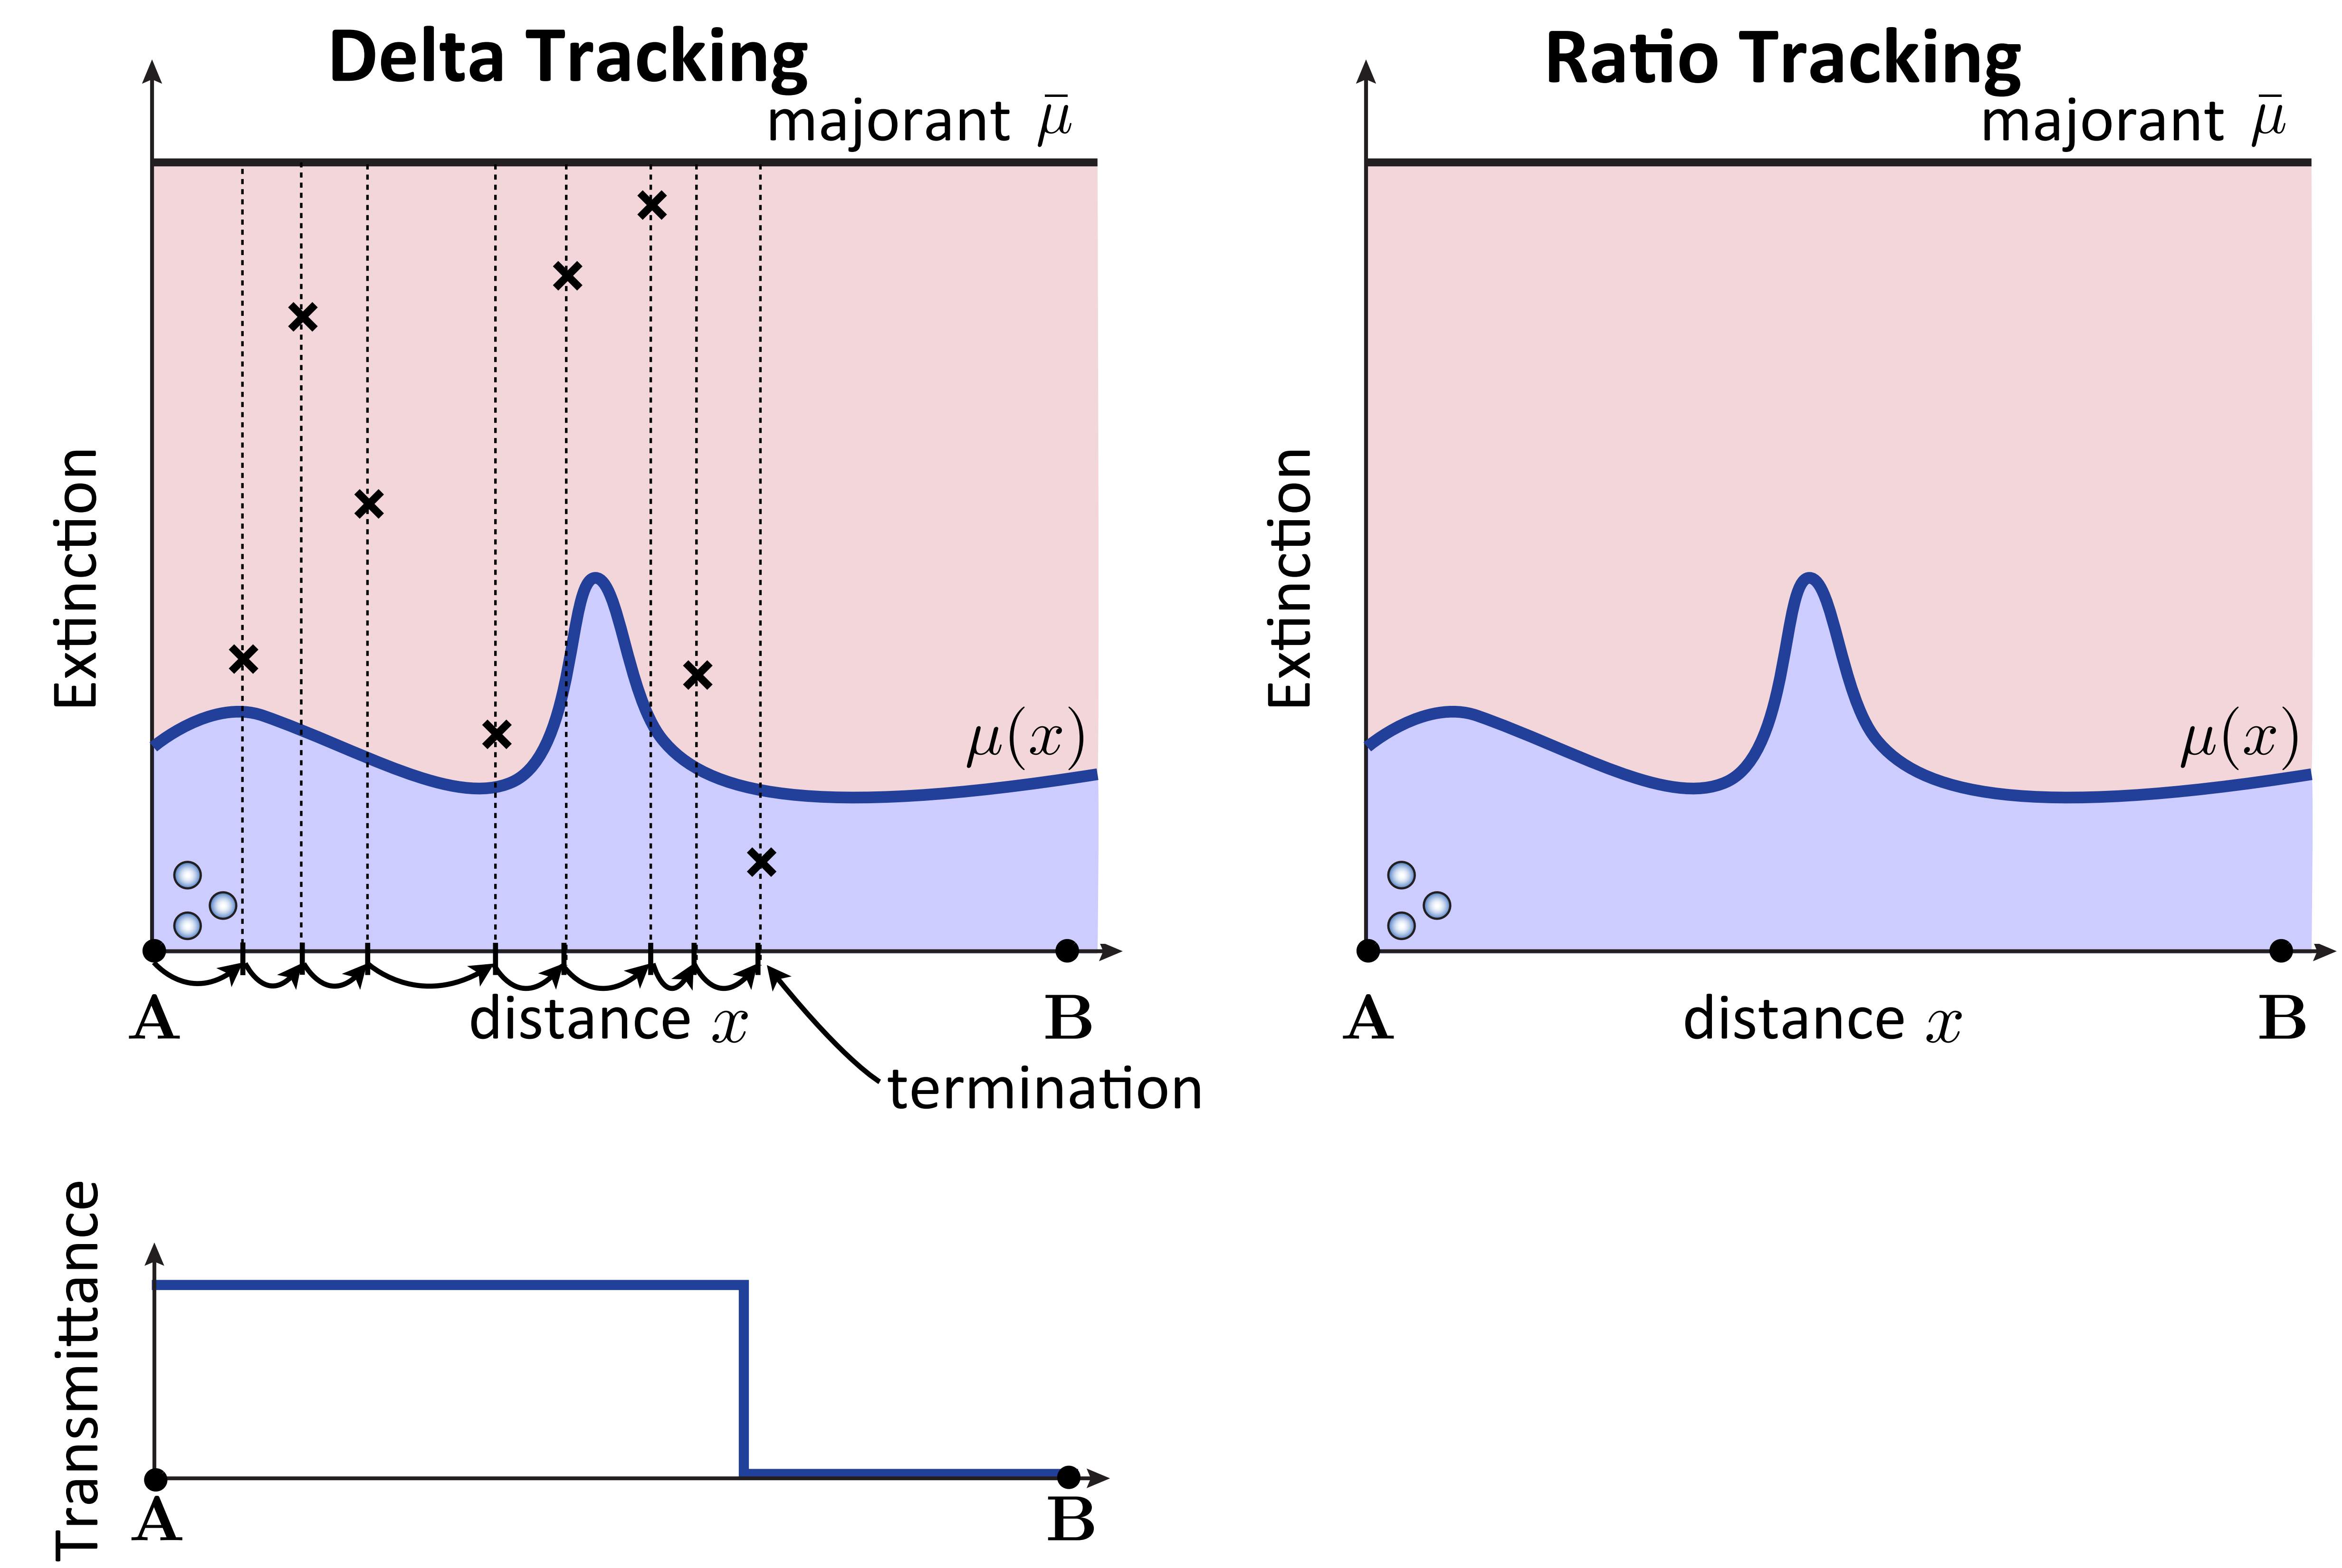

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

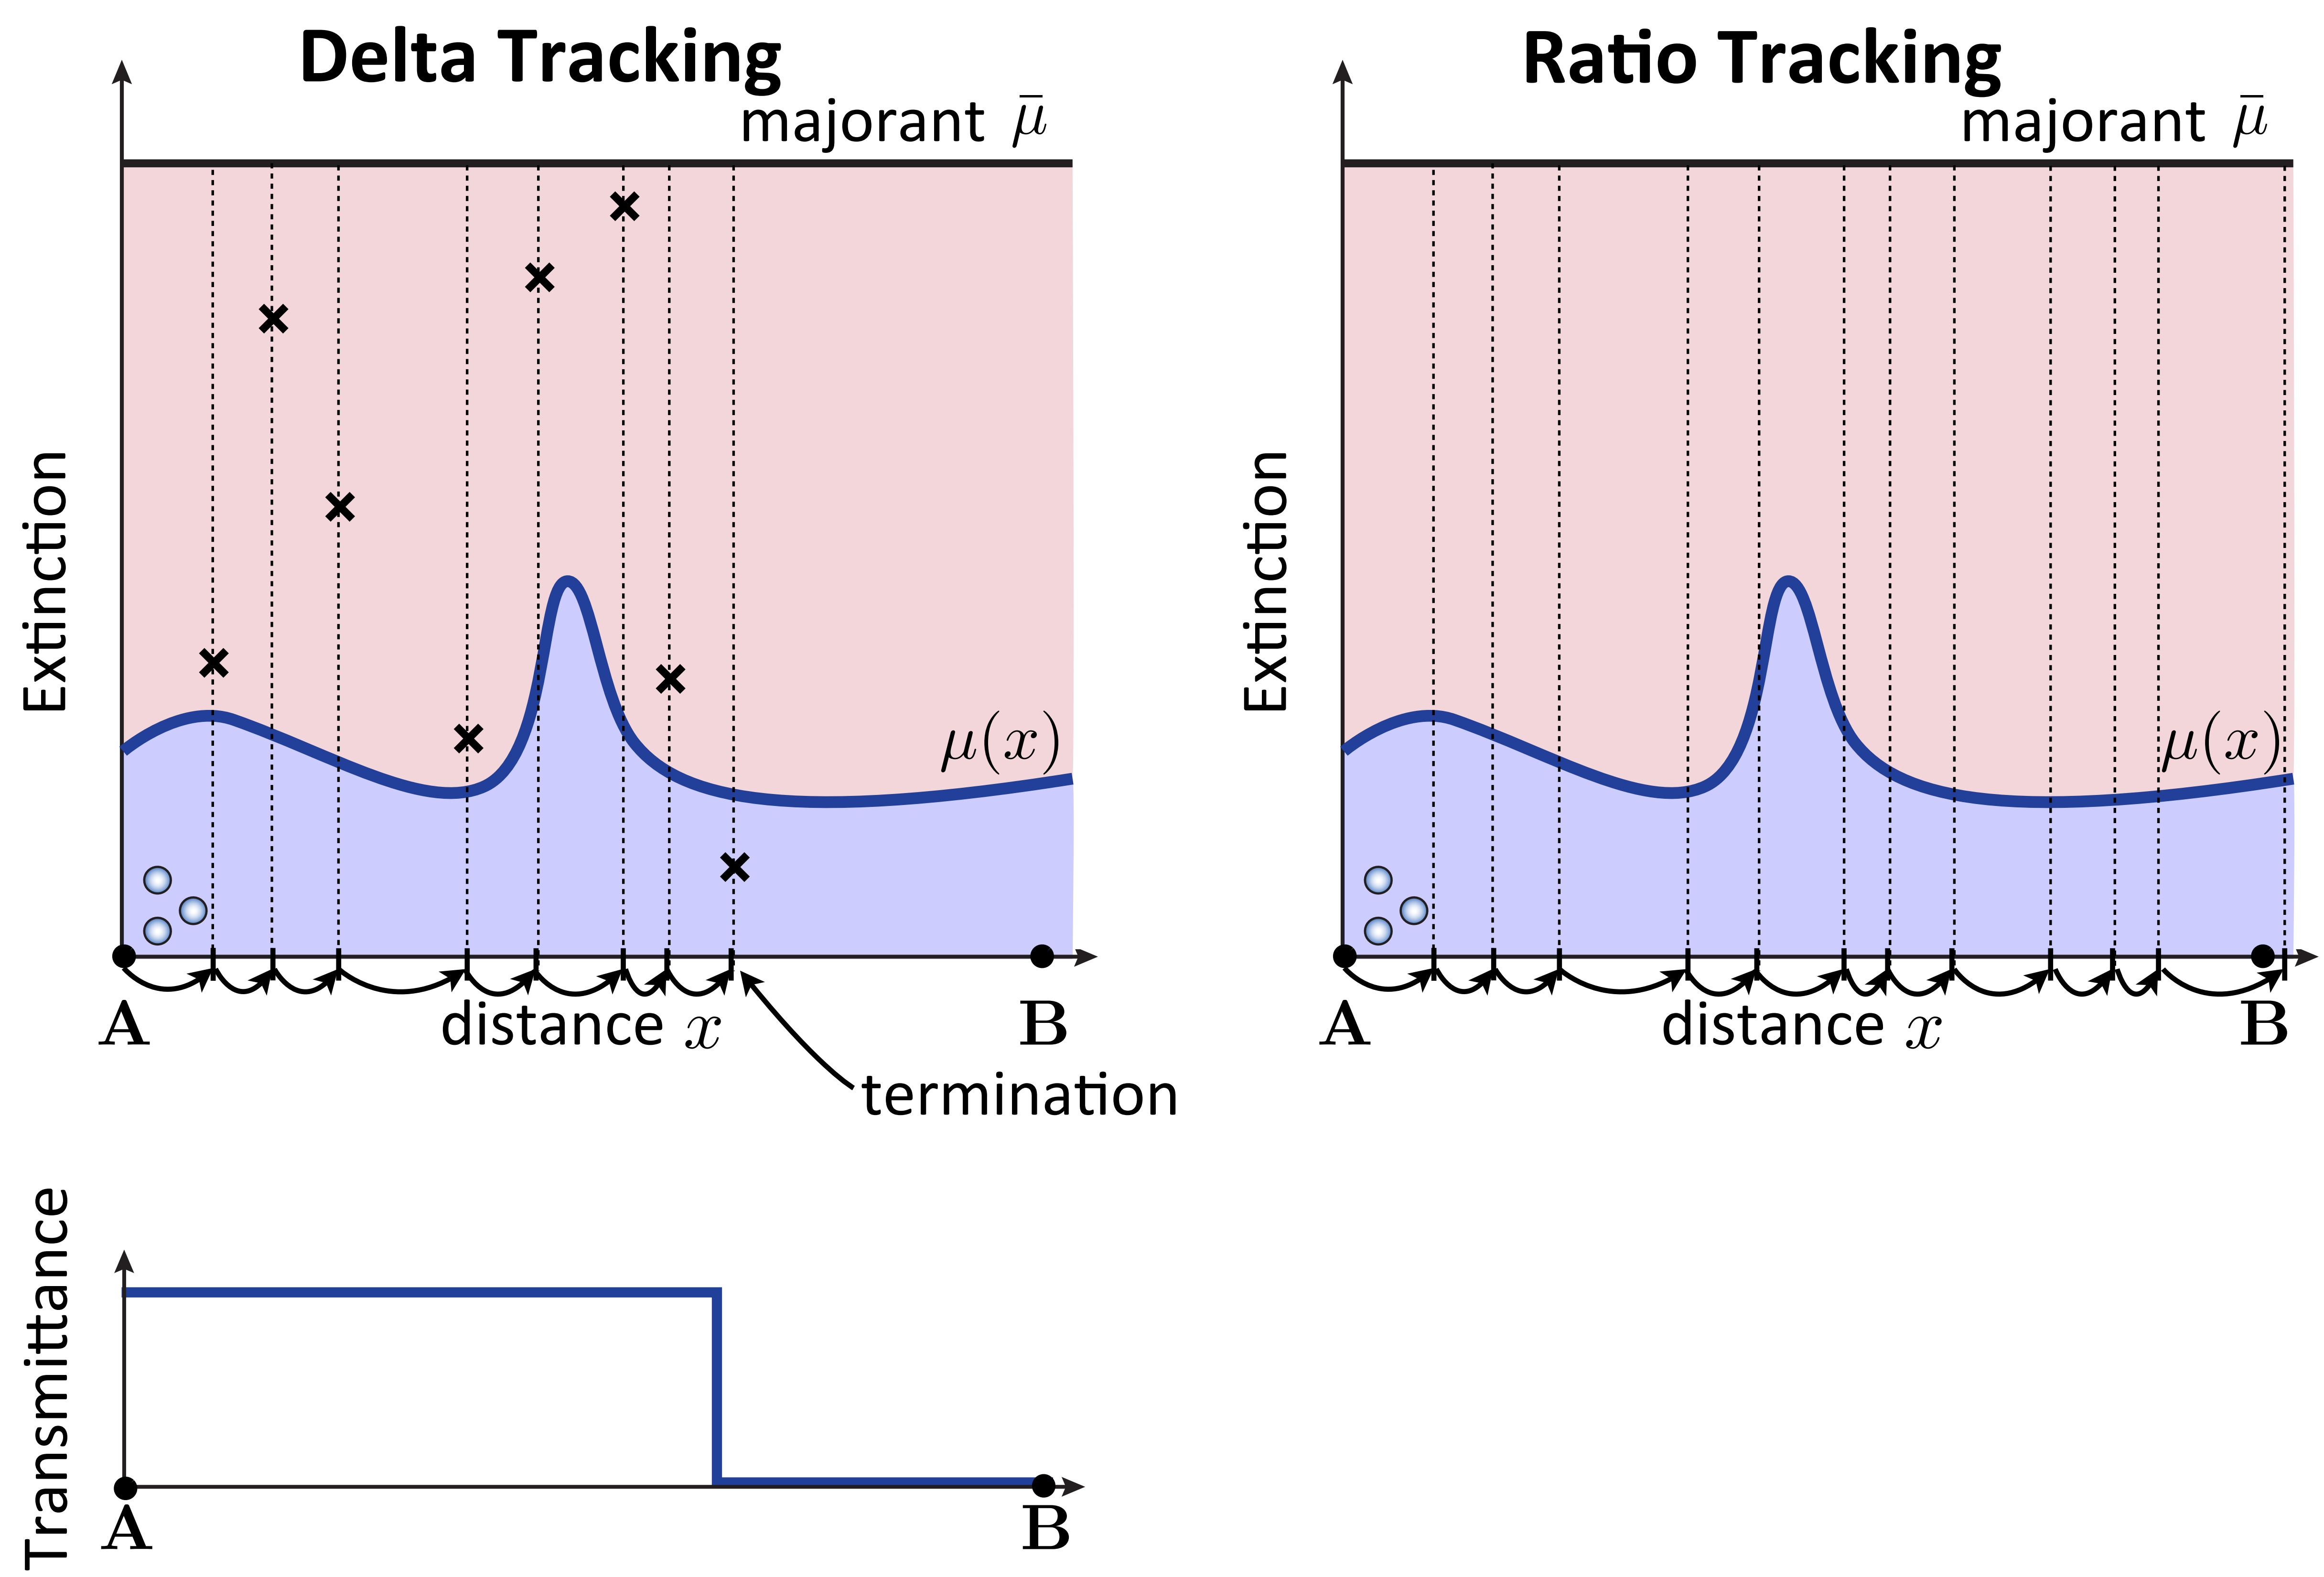

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

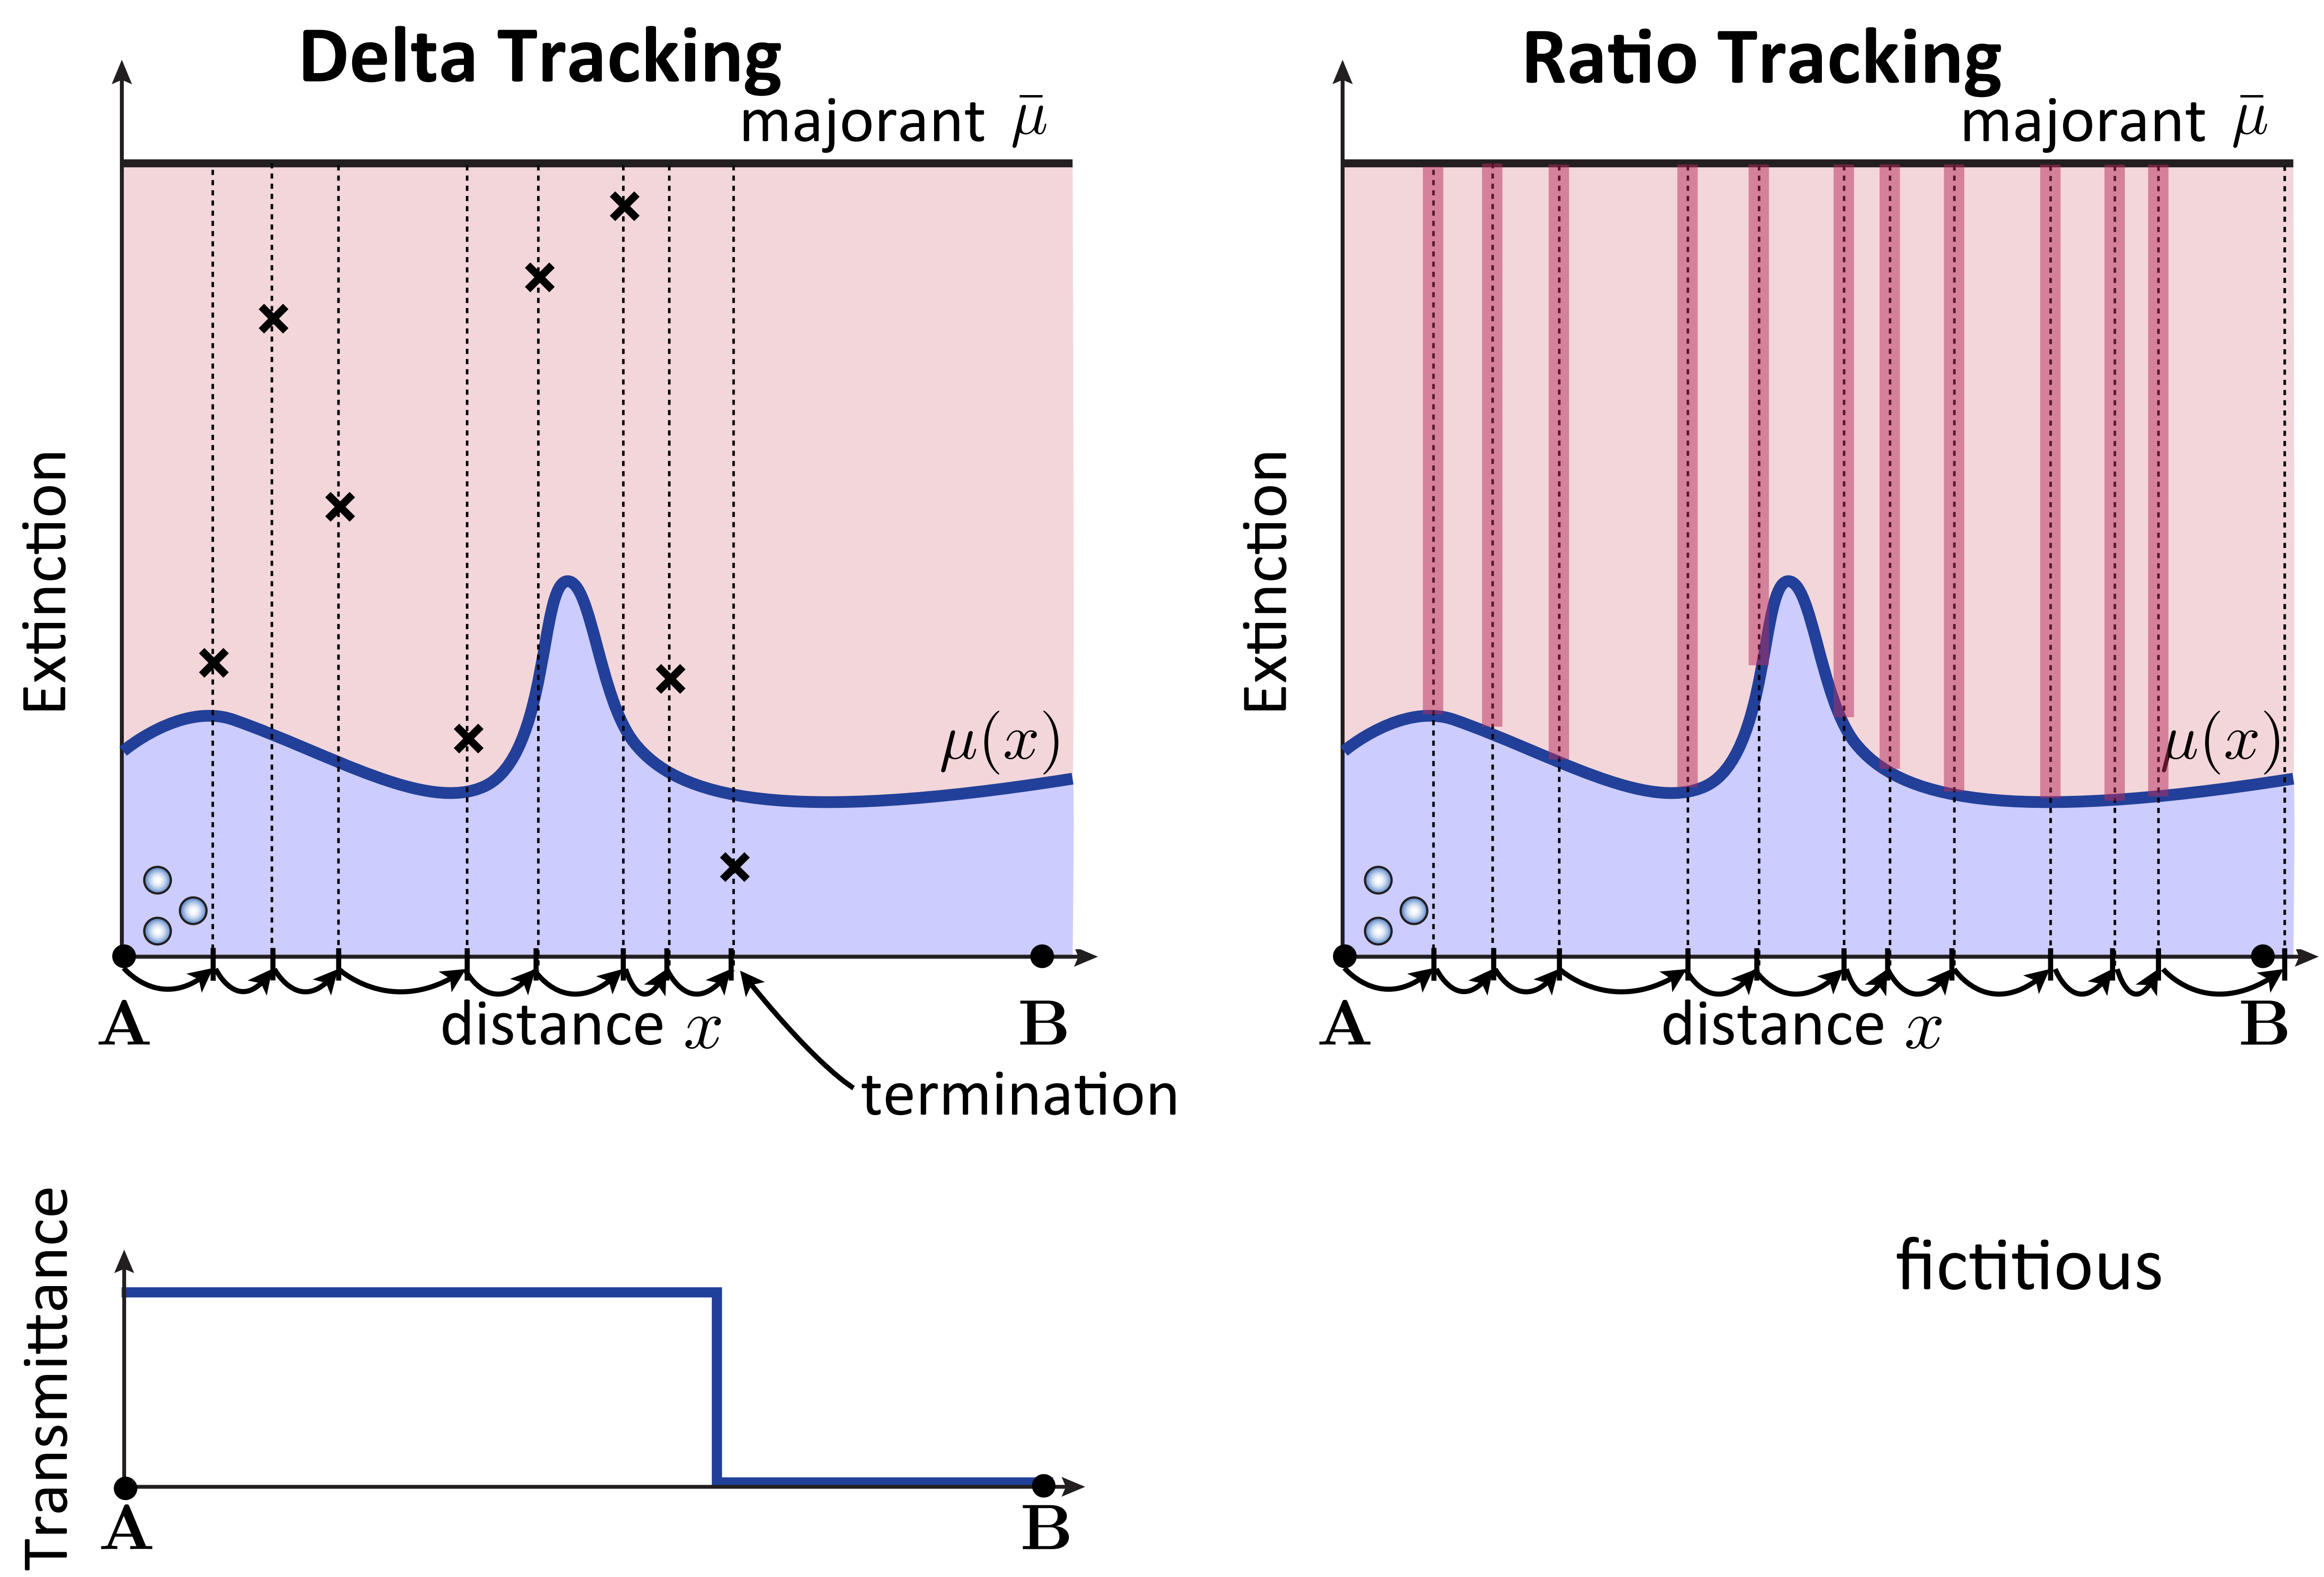

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

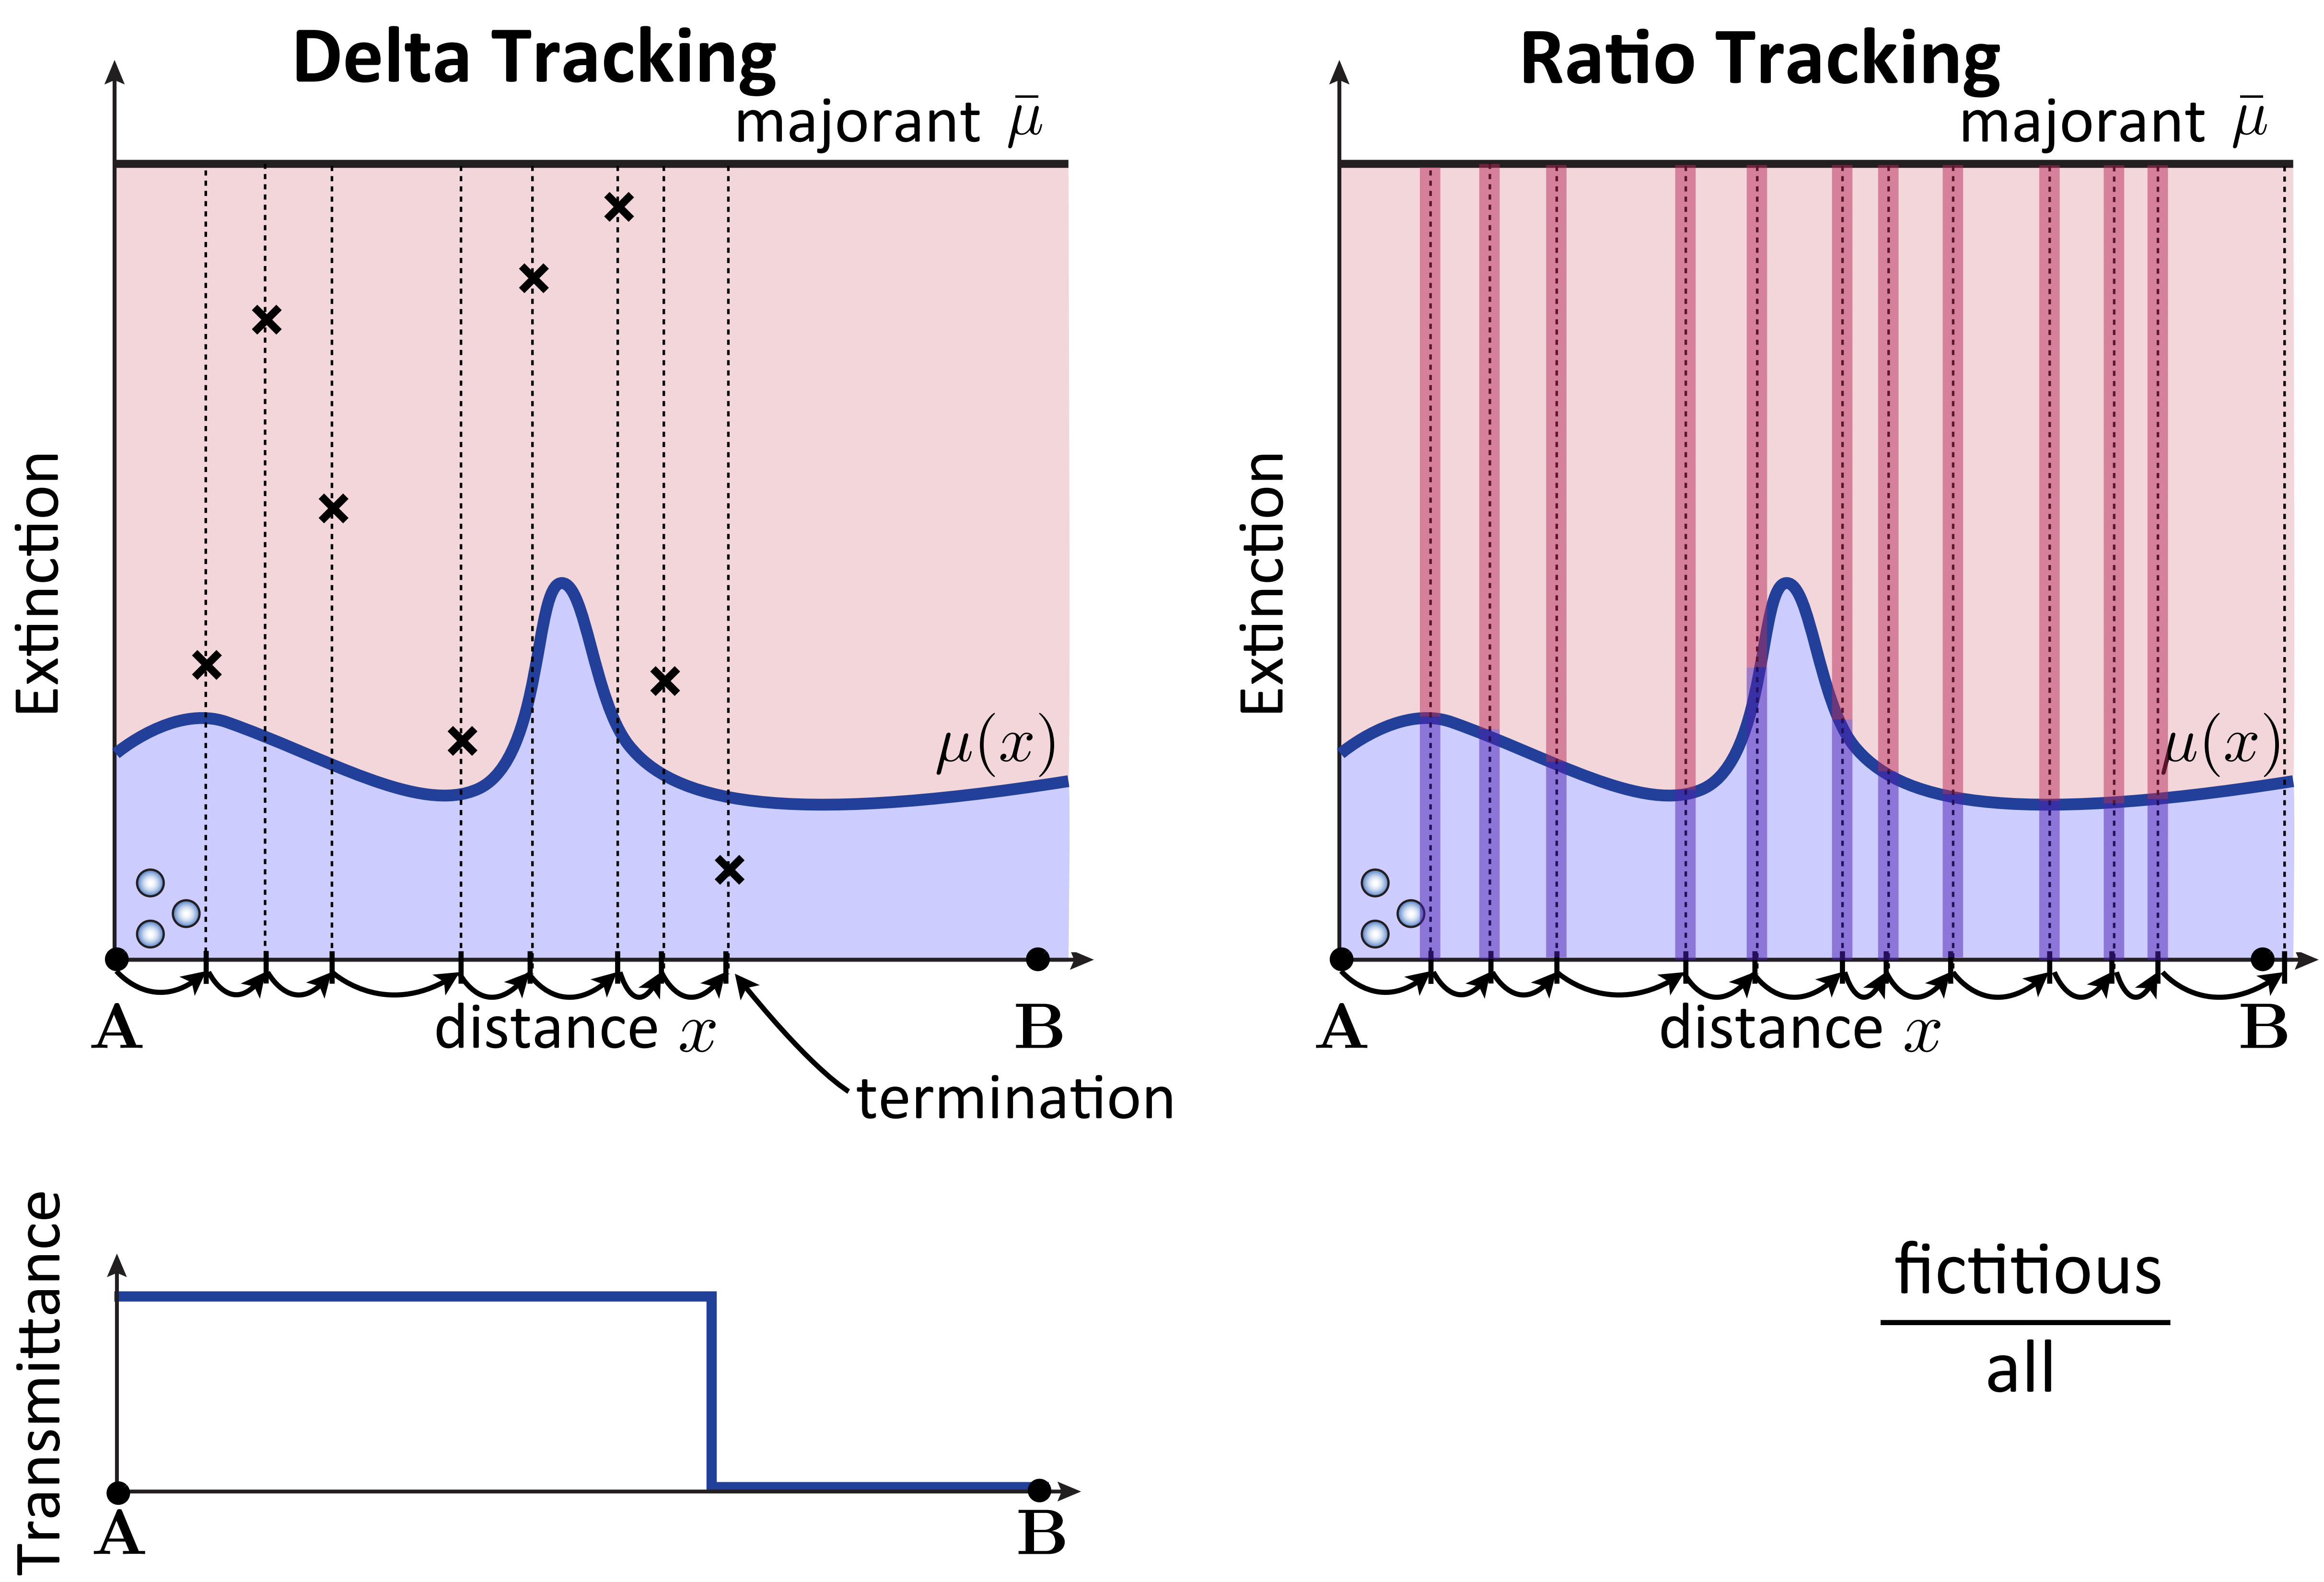

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

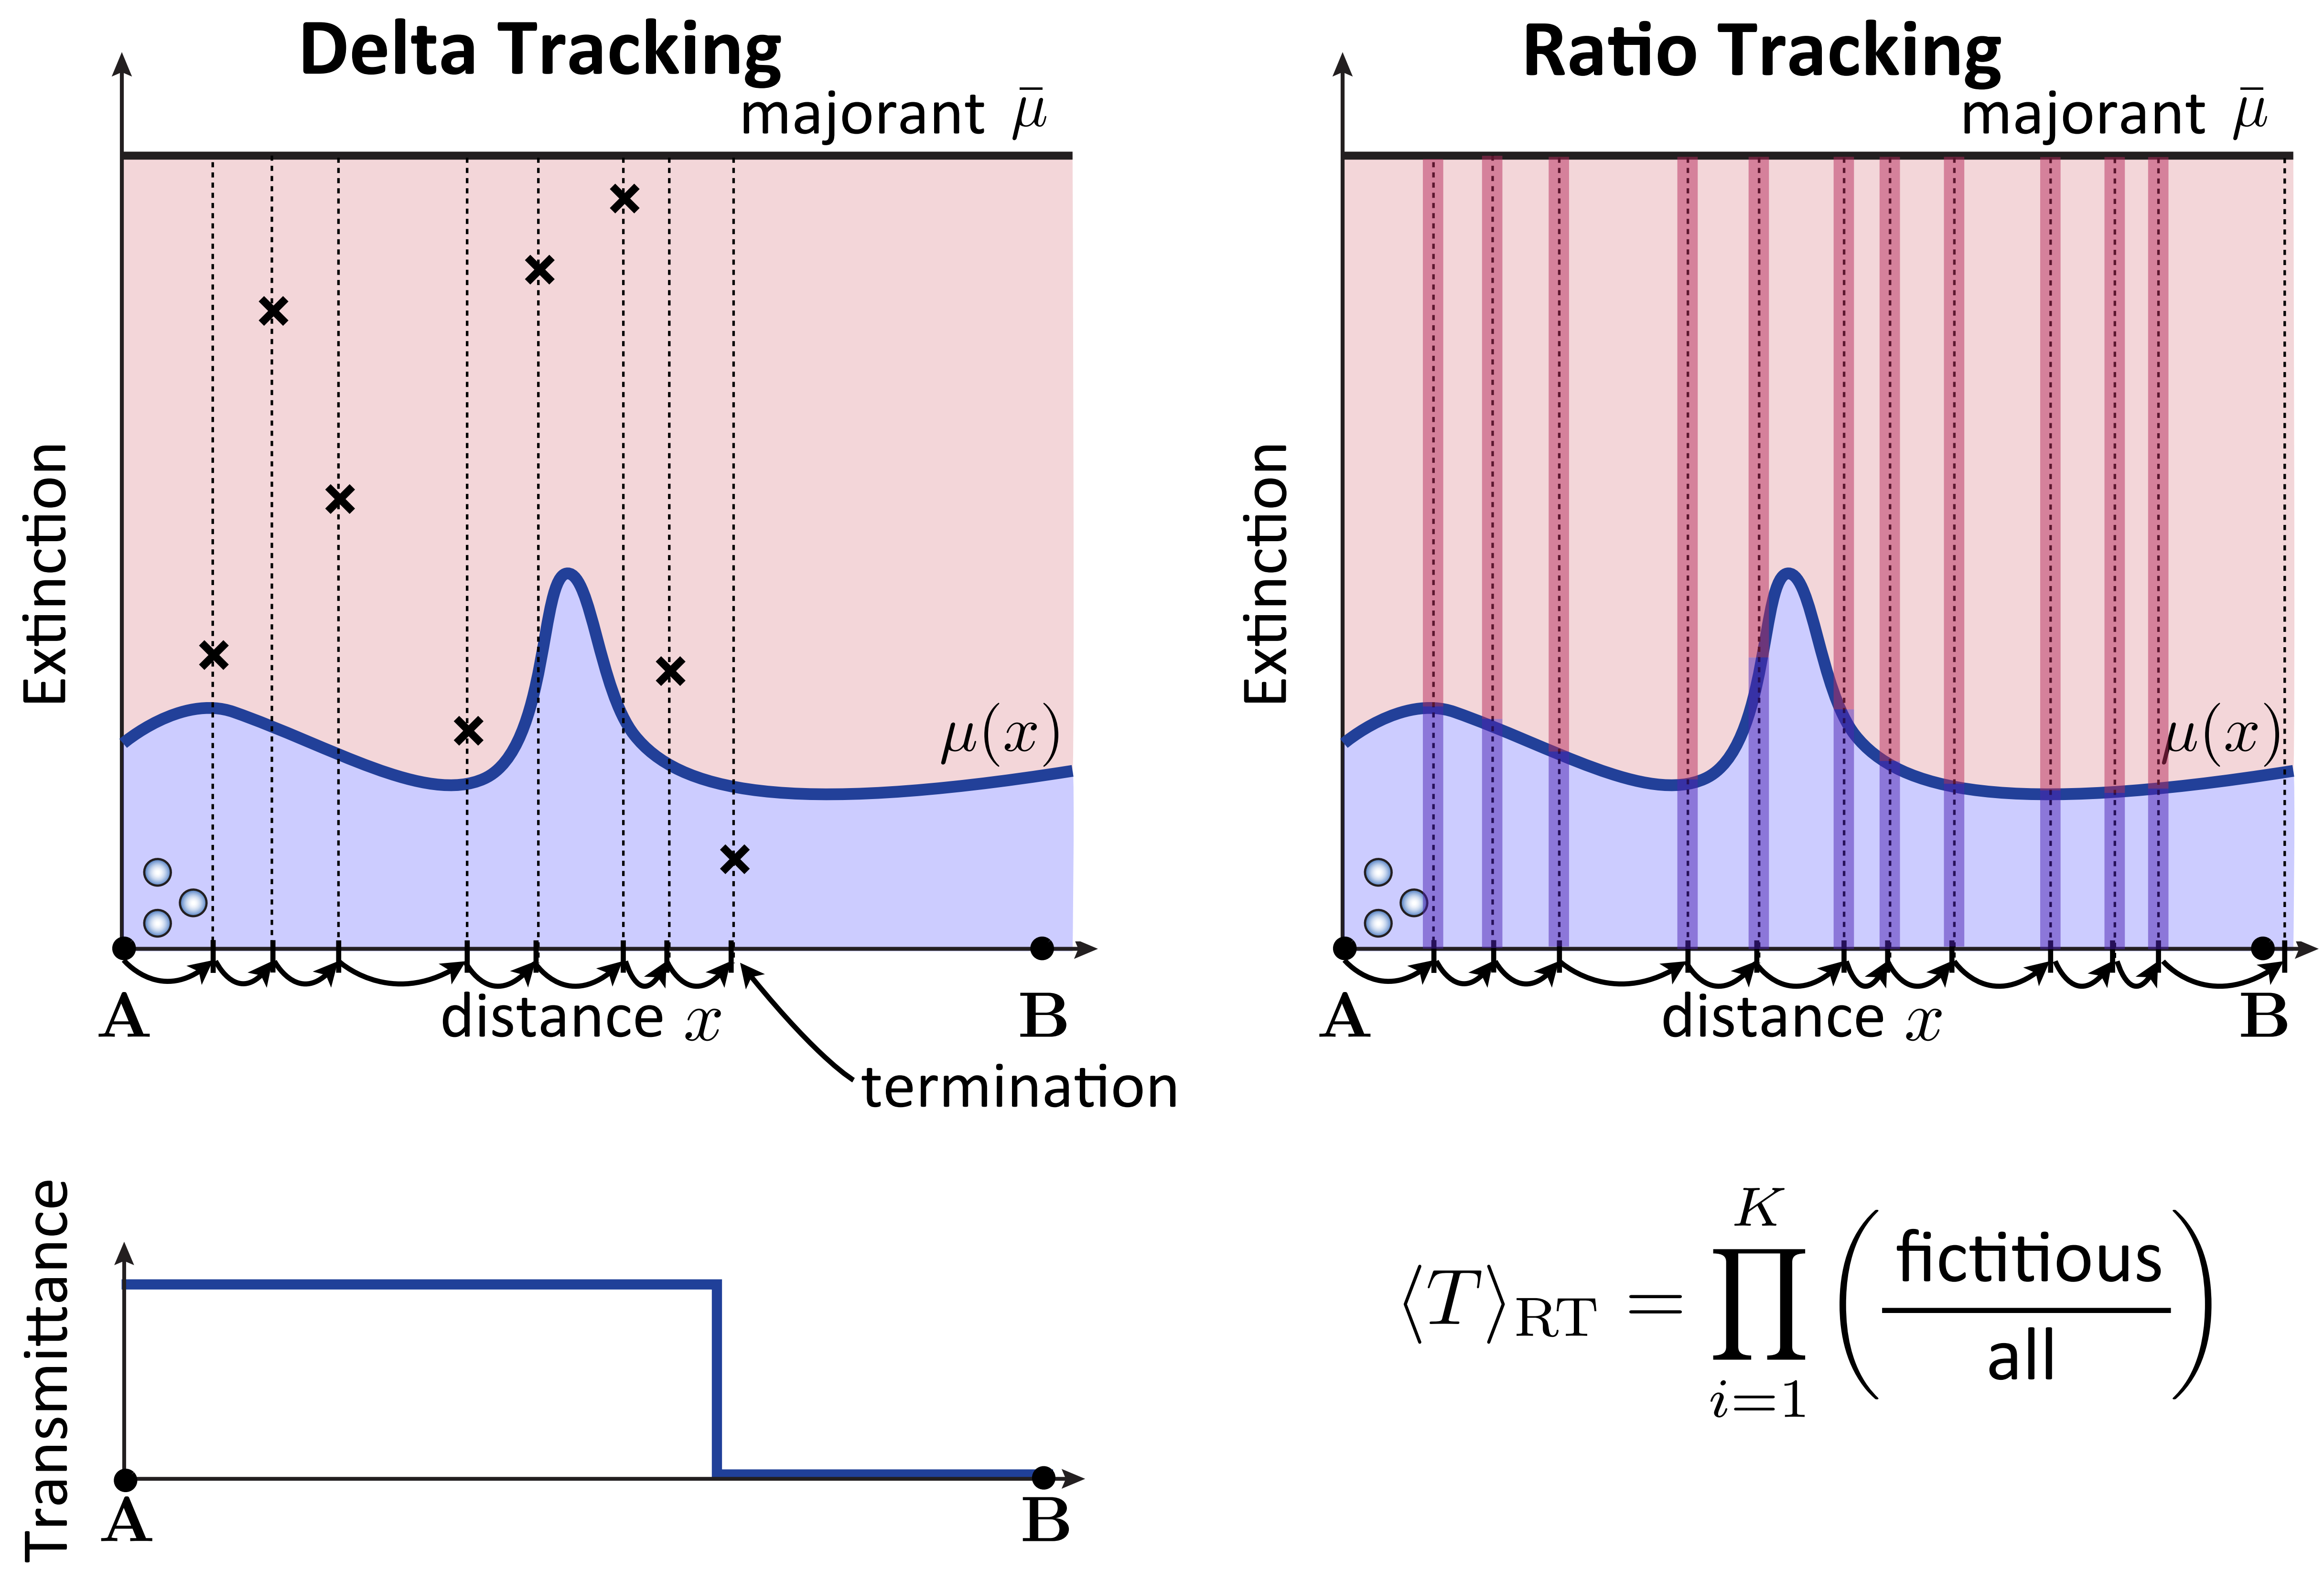

$$\langle T \rangle_{\mathrm{RT}} = \prod_{i=1}^{K} \left(\frac{\text{fictitious}}{\text{all}} \right)$$

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

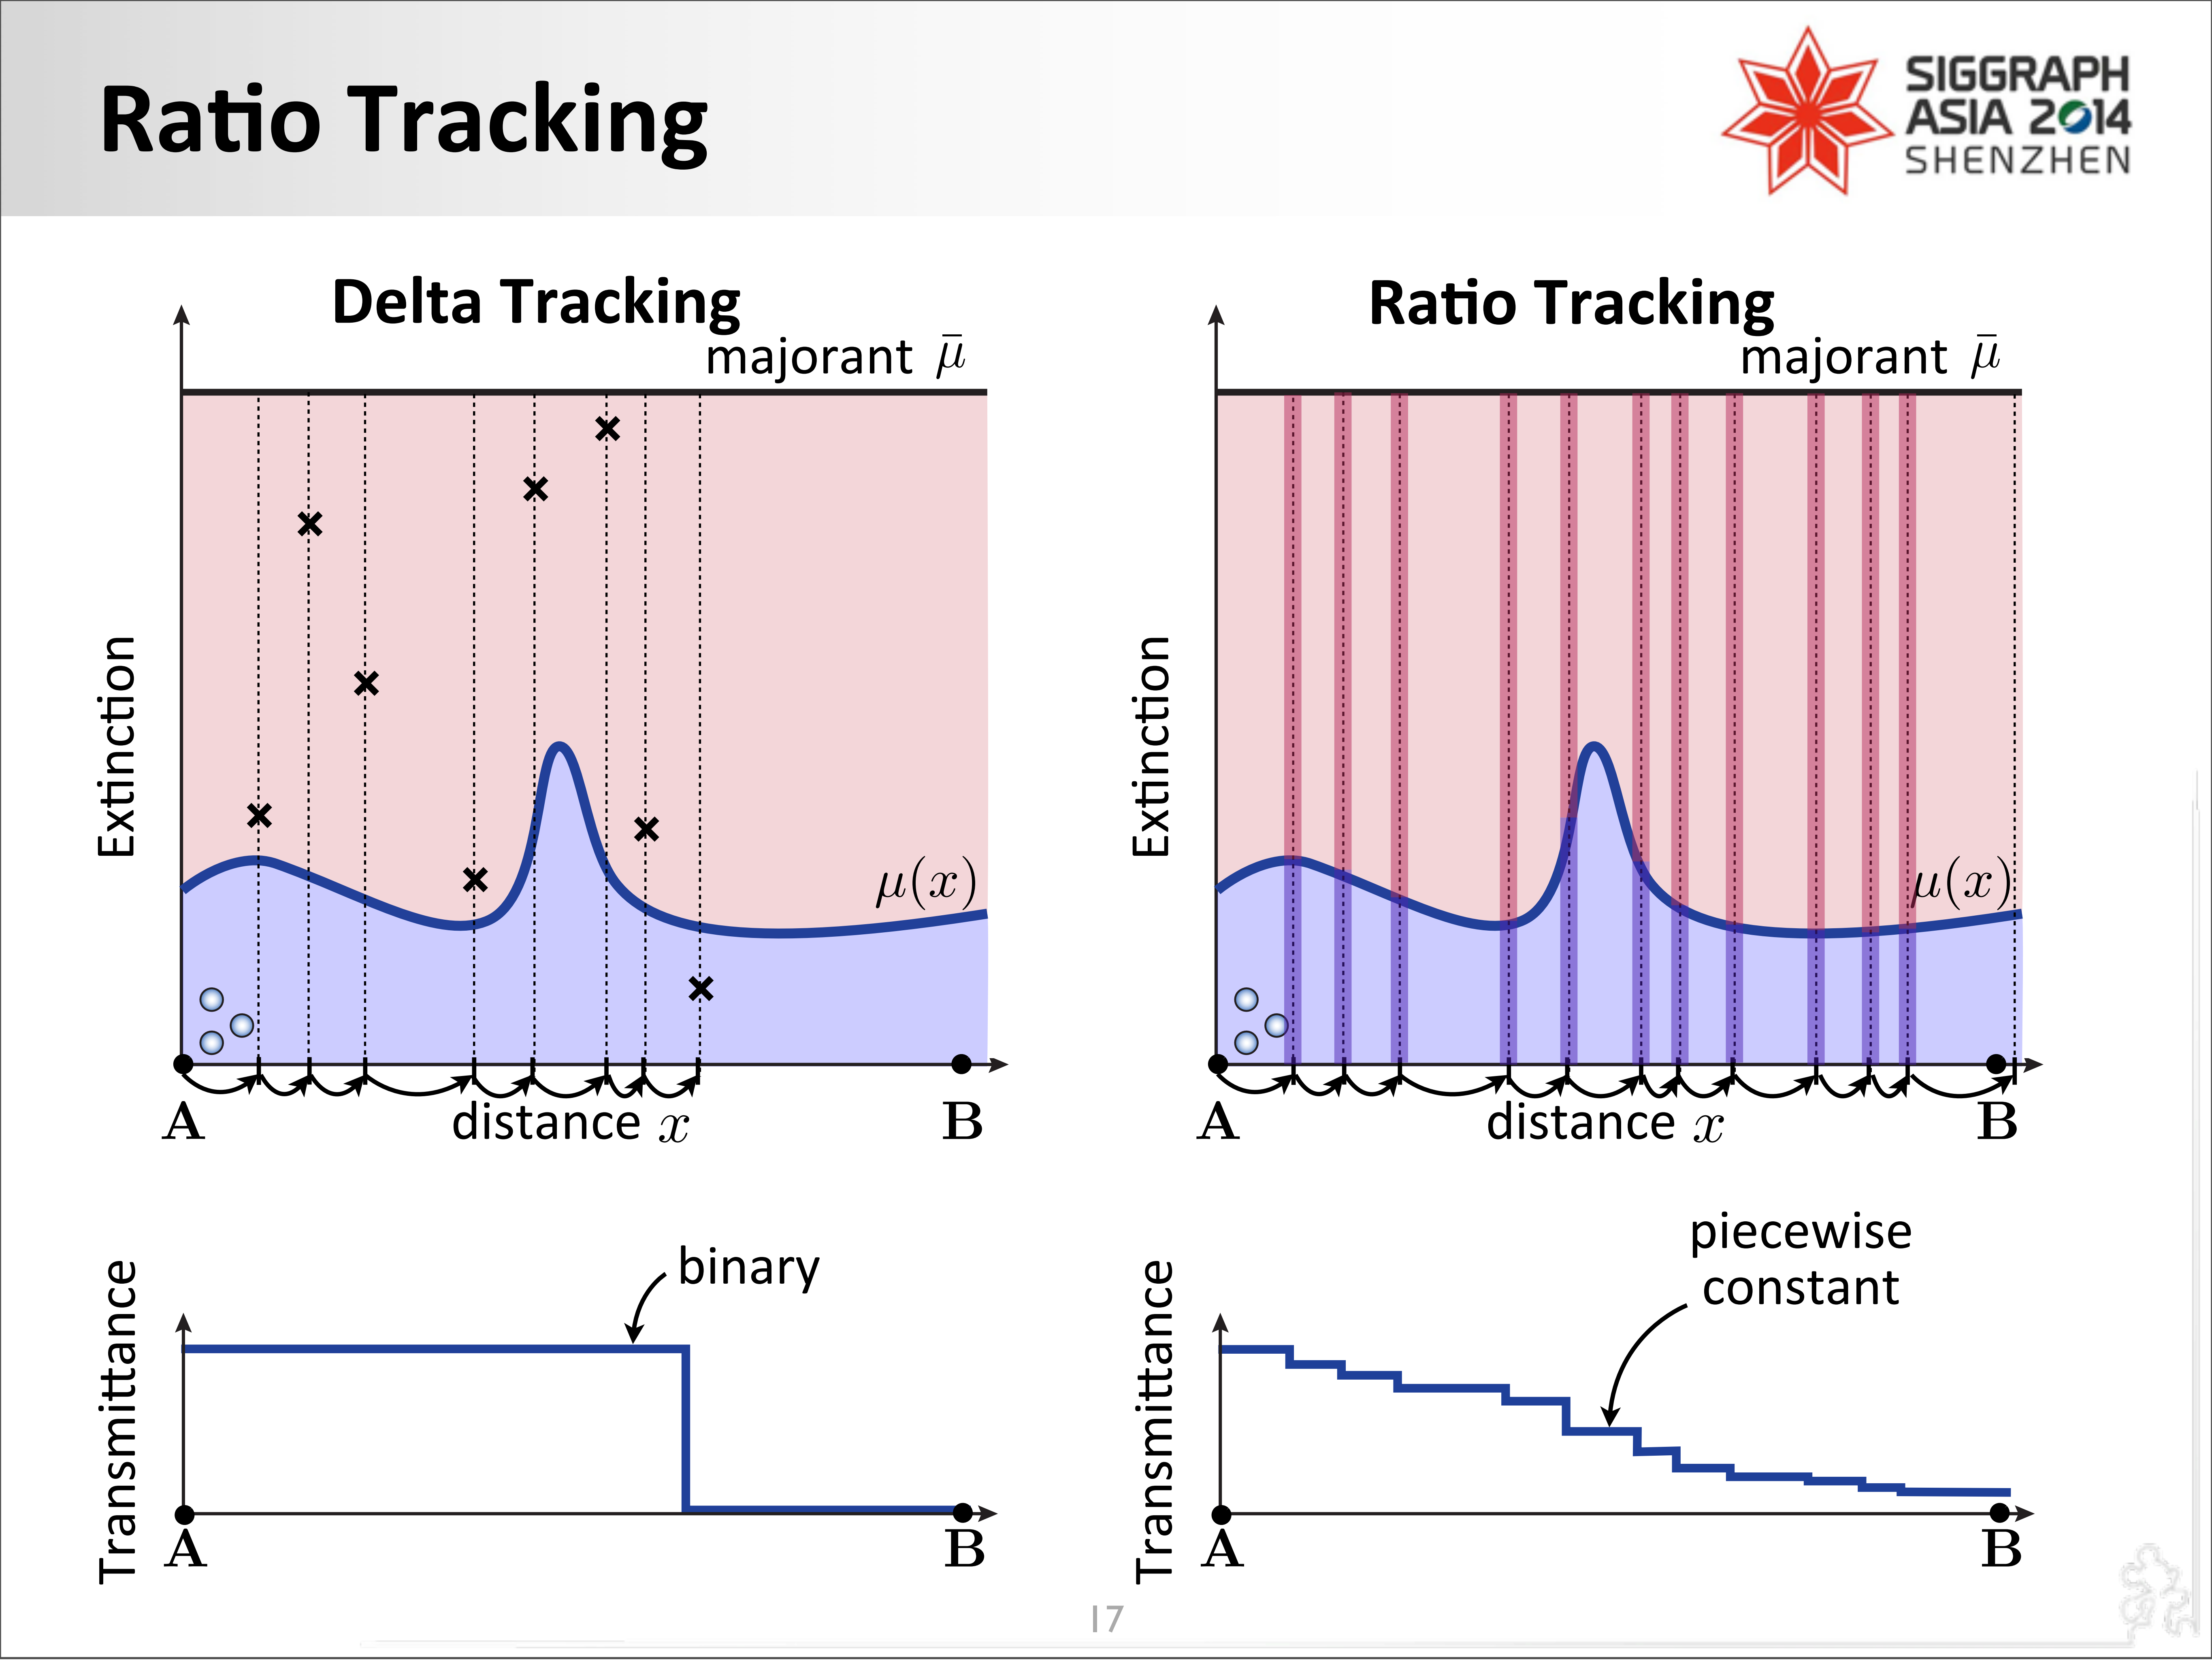

Let's quickly recap the Delta tracking

It builds a random walk. The walk is probabilistically terminated... at the first collision classified as real, and depending whether the walk reached the desired distance, the estimator scores 0 or 1. So we can think of Delta tracking as a random walk terminated by Russian roulette.

Our goal here is to refine the binary estimation. Instead of probabilistically terminating the walk, we always continue, but at each step, we compute a weight, which equals the ratio of fictitious particles to all the particles (to the majorant), the product of these weights then becomes the score of the estimator.

In a nutshell, we could say that the Ratio tracking disables the Russian roulette that was there previously, and factors its probabilities into the score.

Please see the paper for an exact definition of the algorithm.

The main advantage of Ratio tracking is that it produces a piecewise-constant approximation.

Ratio Tracking

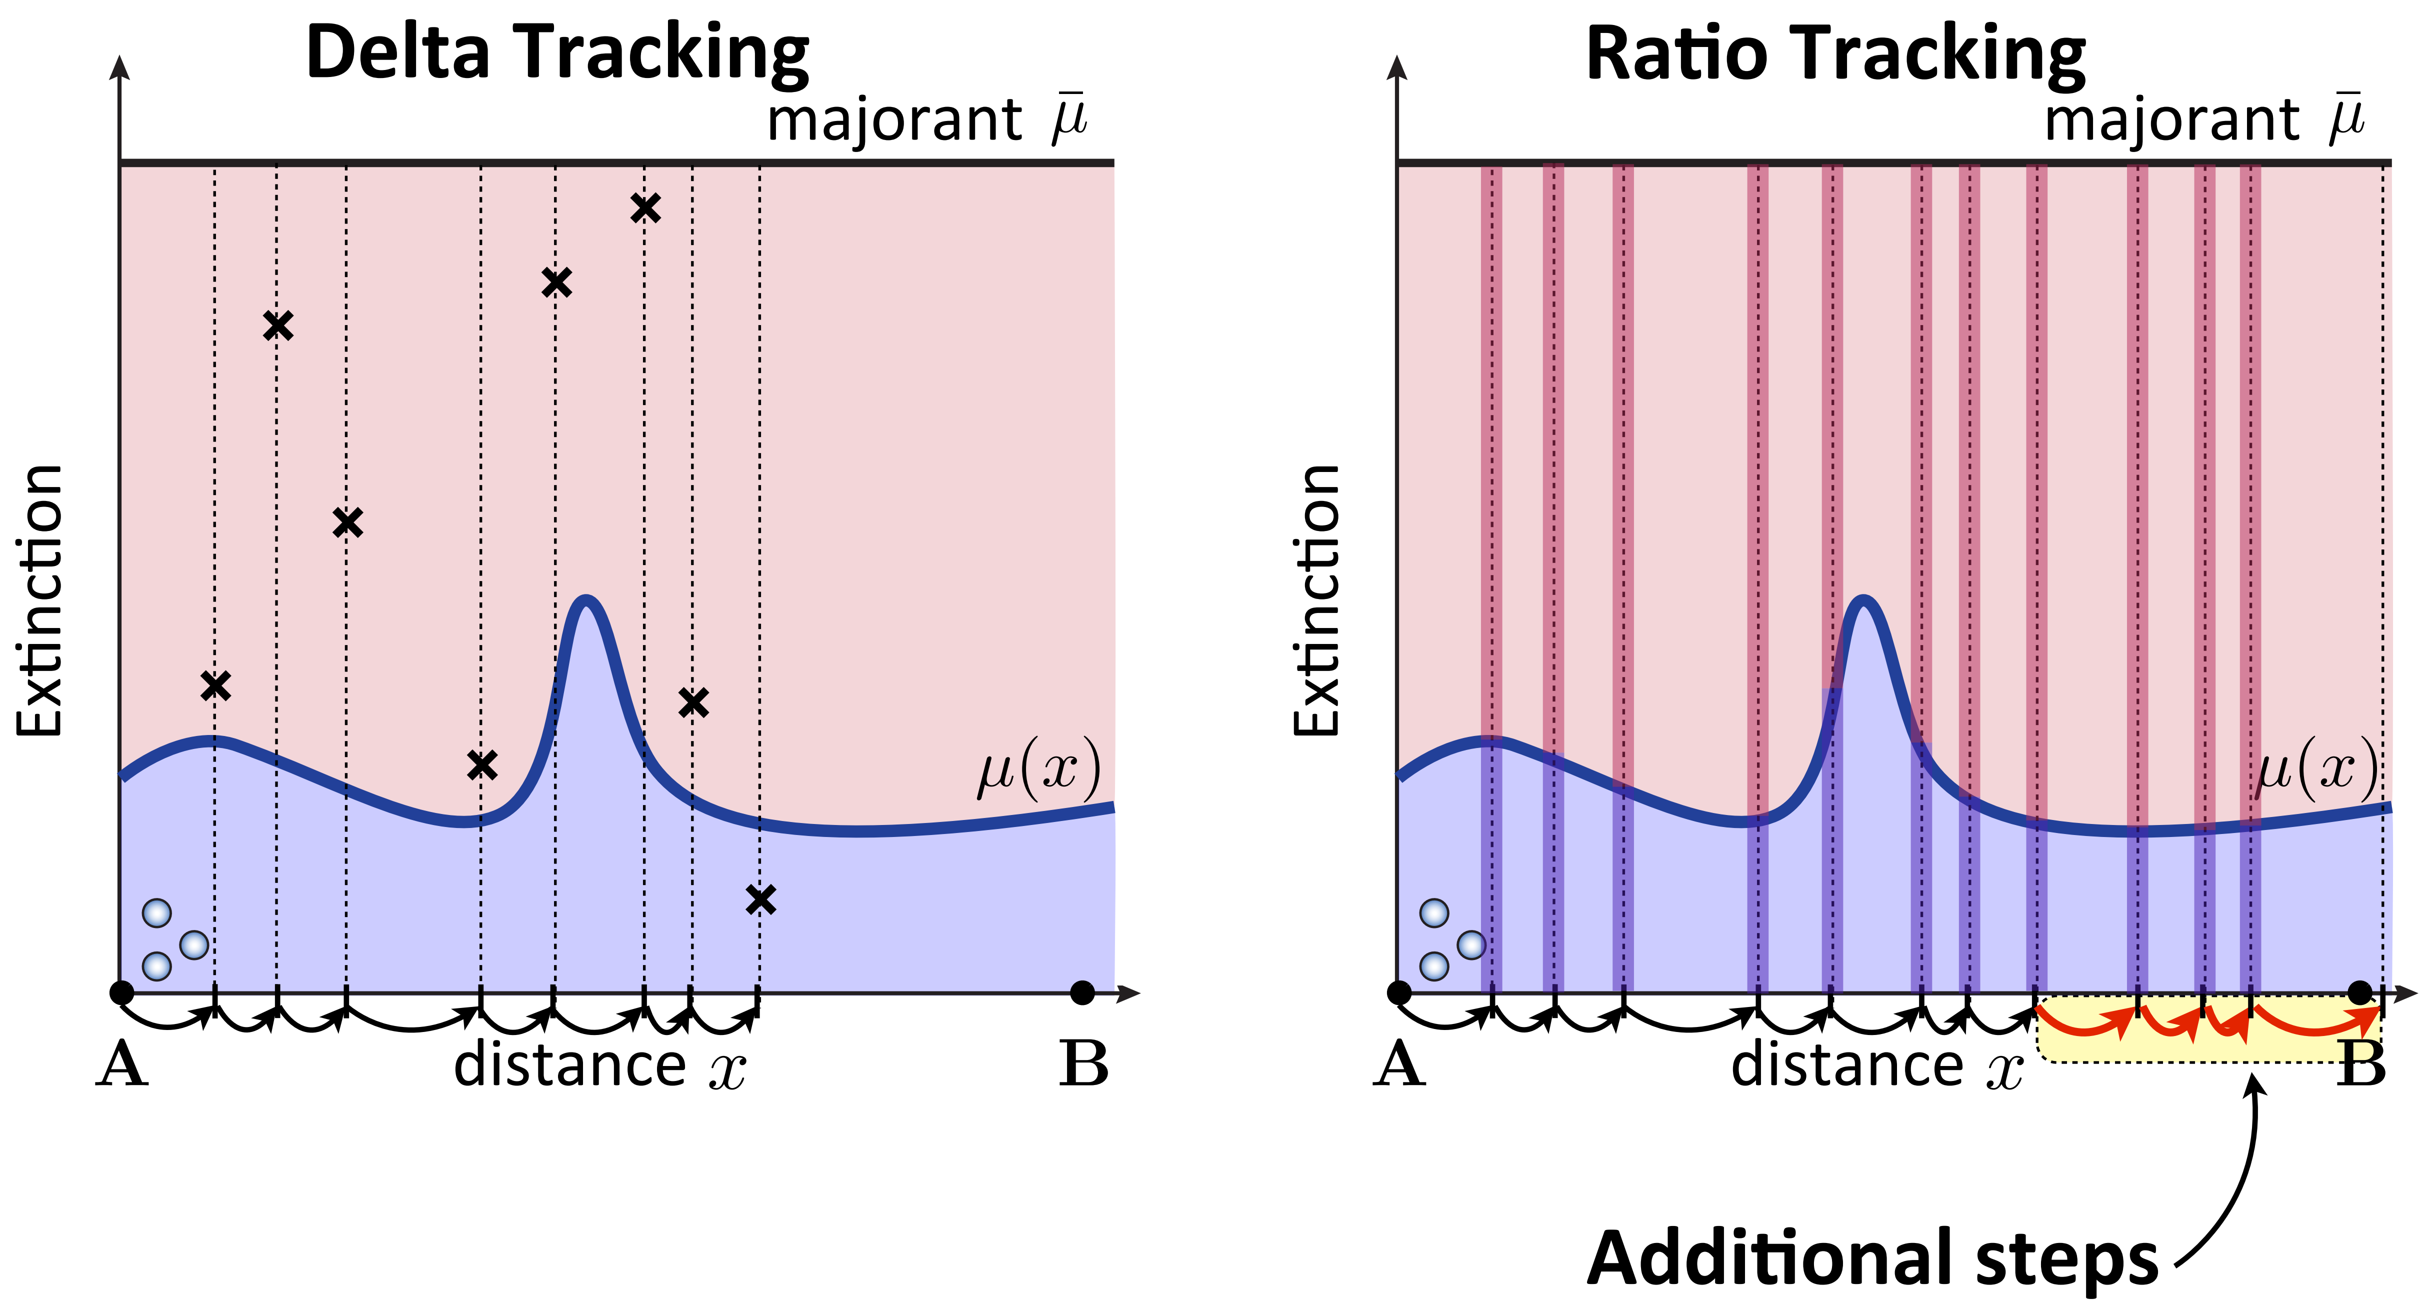

The only disadvantage of our approach is that it requires these additional steps. These can make it less efficient than Delta tracking in certain situations.

Delta Tracking vs. Ratio Tracking

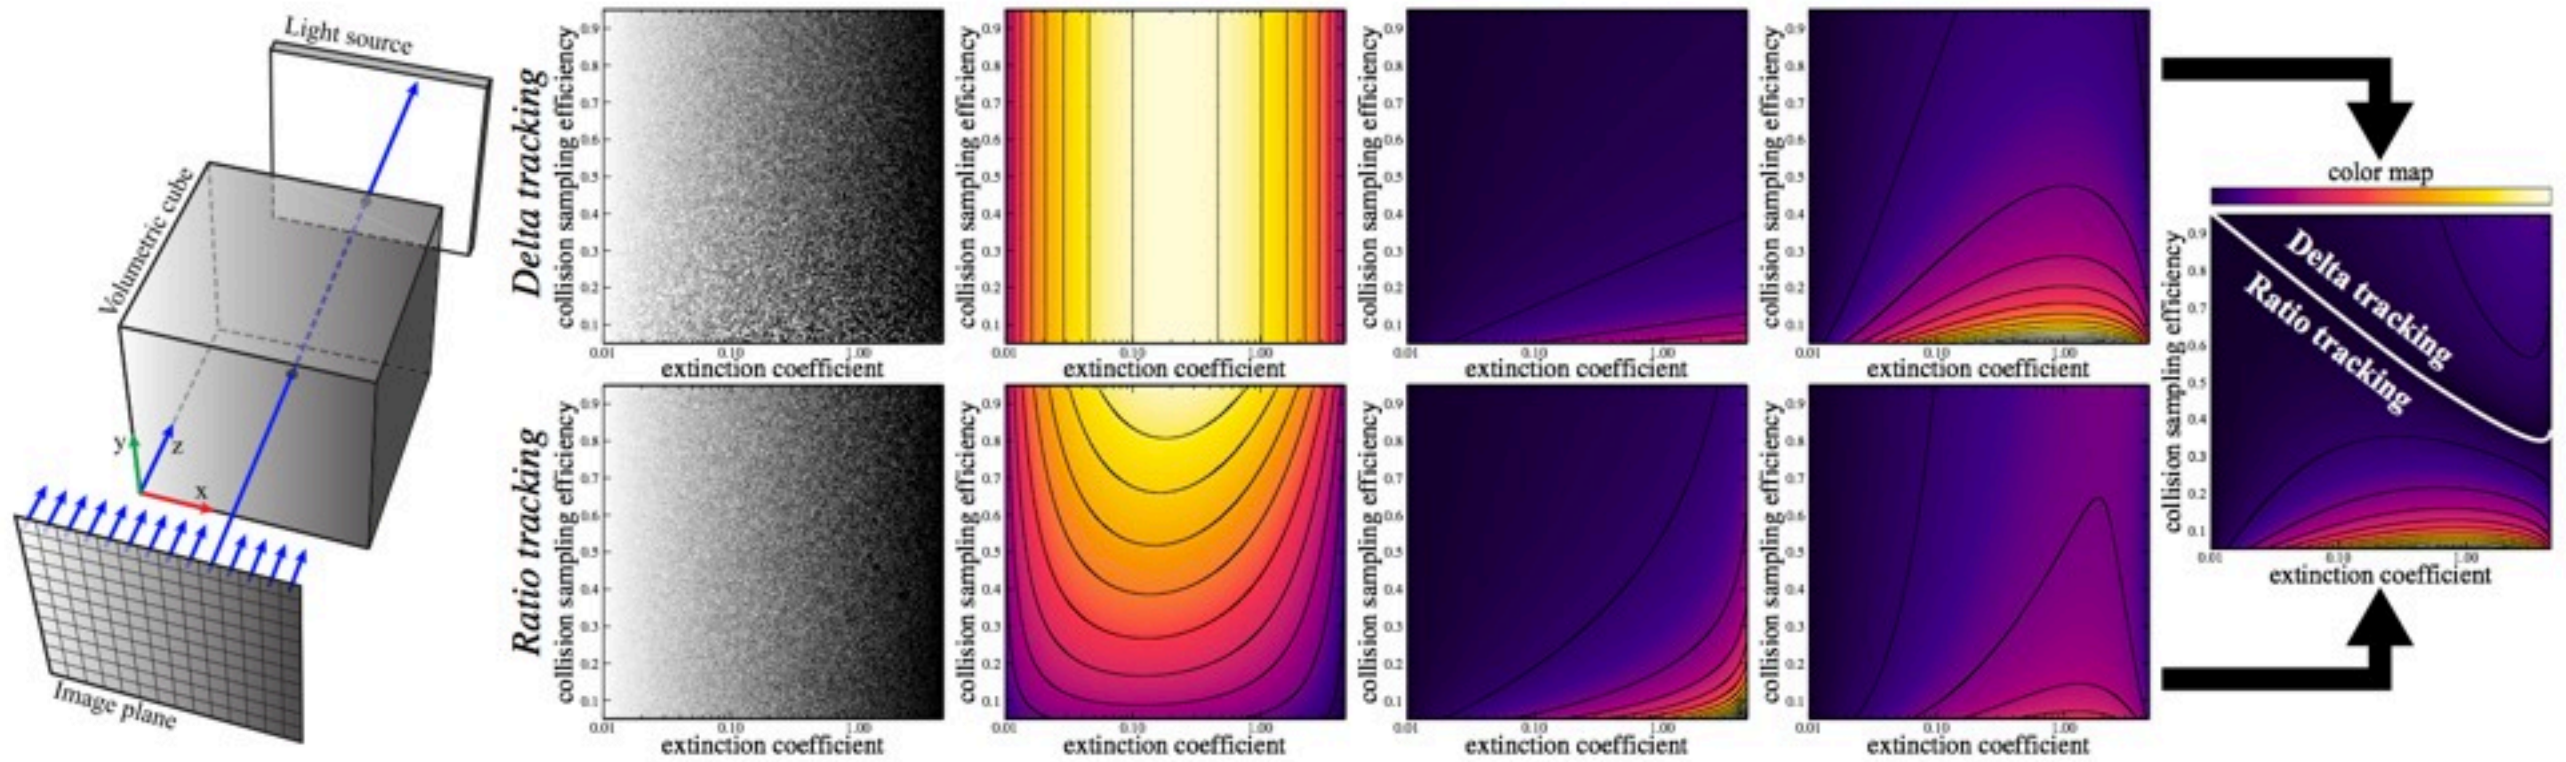

(a) Canonical scene (b) Equal cost render (c) Variance (per instance) (d) Cost (per instance) (e) Product (c) × (d) (f) Difference

We have a thorough analysis of the performance of the two algorithms in the paper,
but I will skip it here and summarize it by saying that:

In practical applications... Ratio tracking can be significantly better, but never significantly worse than Delta tracking.

Delta Tracking vs. Ratio Tracking

SIGGRAPH ASIA 2014 SHENZHEN

Ratio tracking **can be significantly better**
but **never significantly worse** than Delta tracking

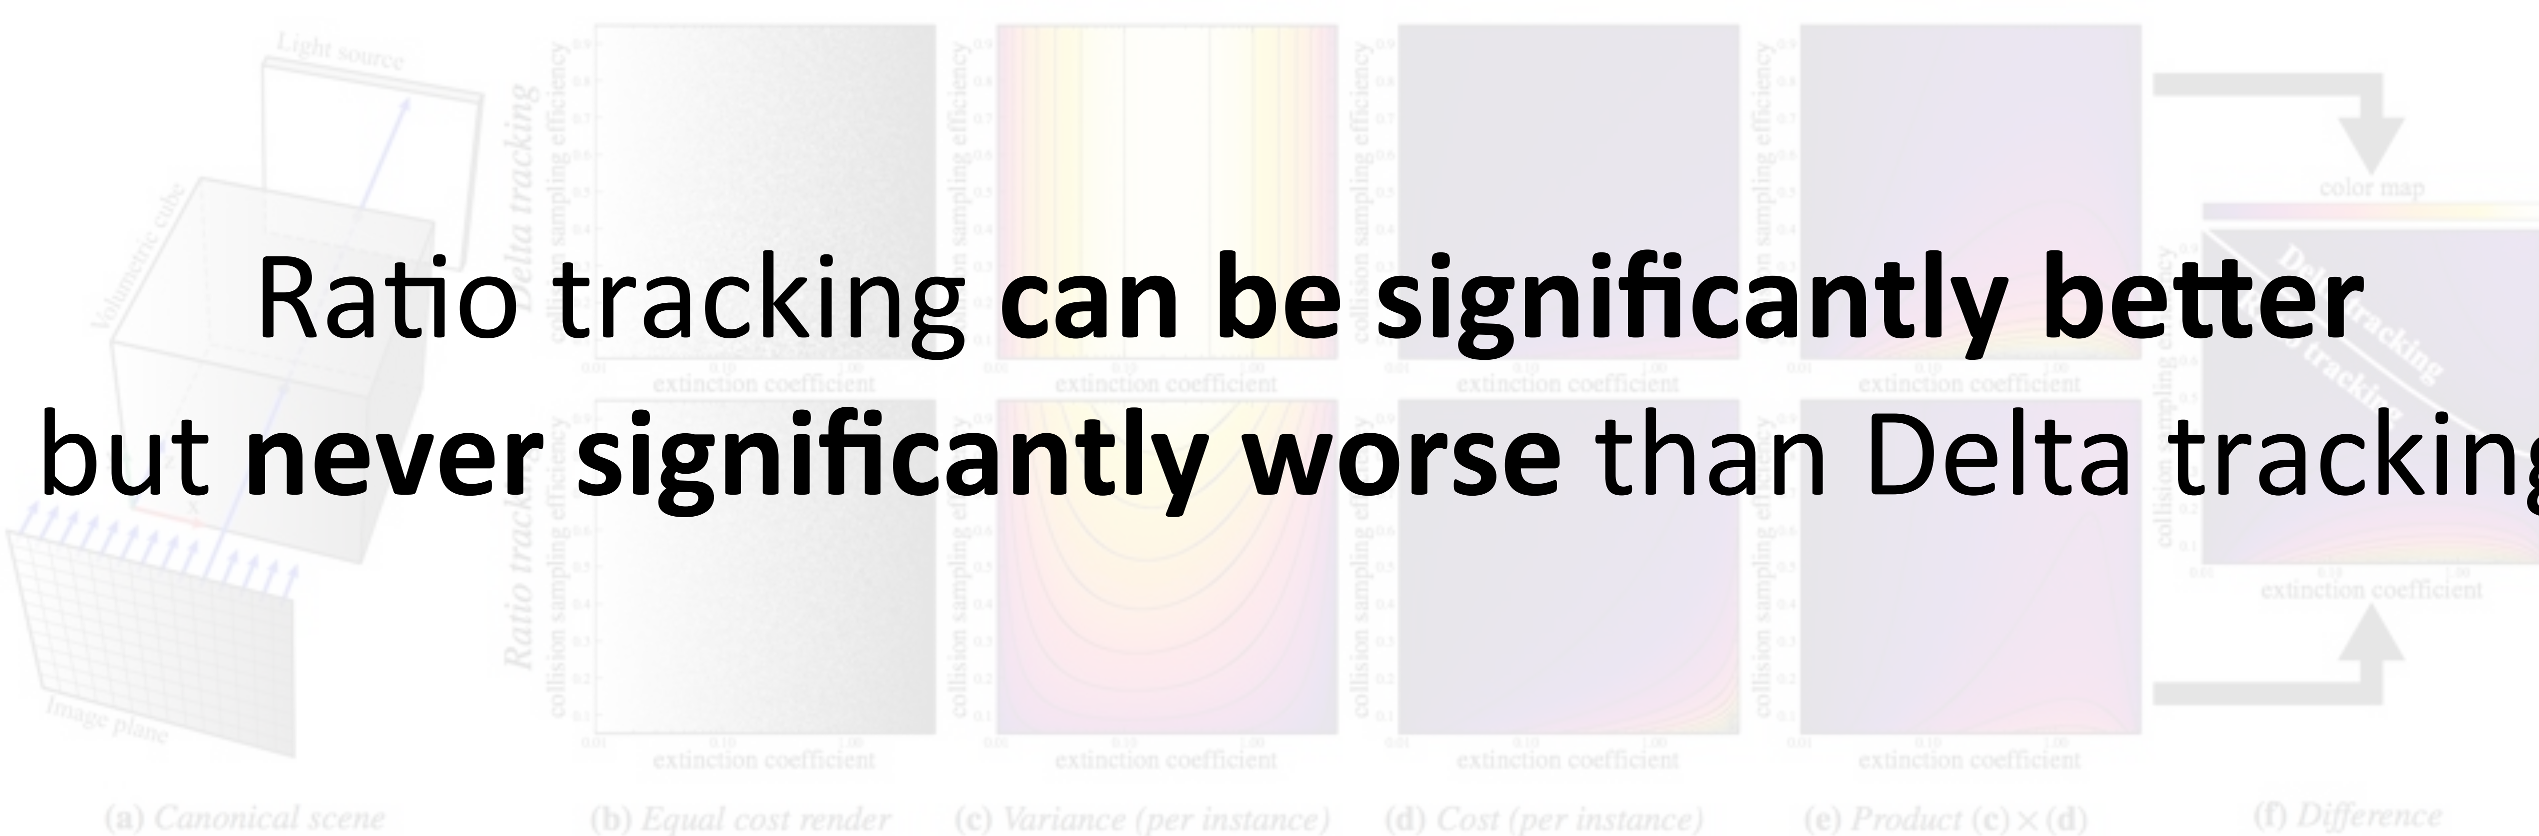

We have a thorough analysis of the performance of the two algorithms in the paper,
but I will skip it here and summarize it by saying that:

In practical applications... Ratio tracking can be significantly better, but never significantly worse than Delta tracking.

Delta Tracking vs. Ratio Tracking

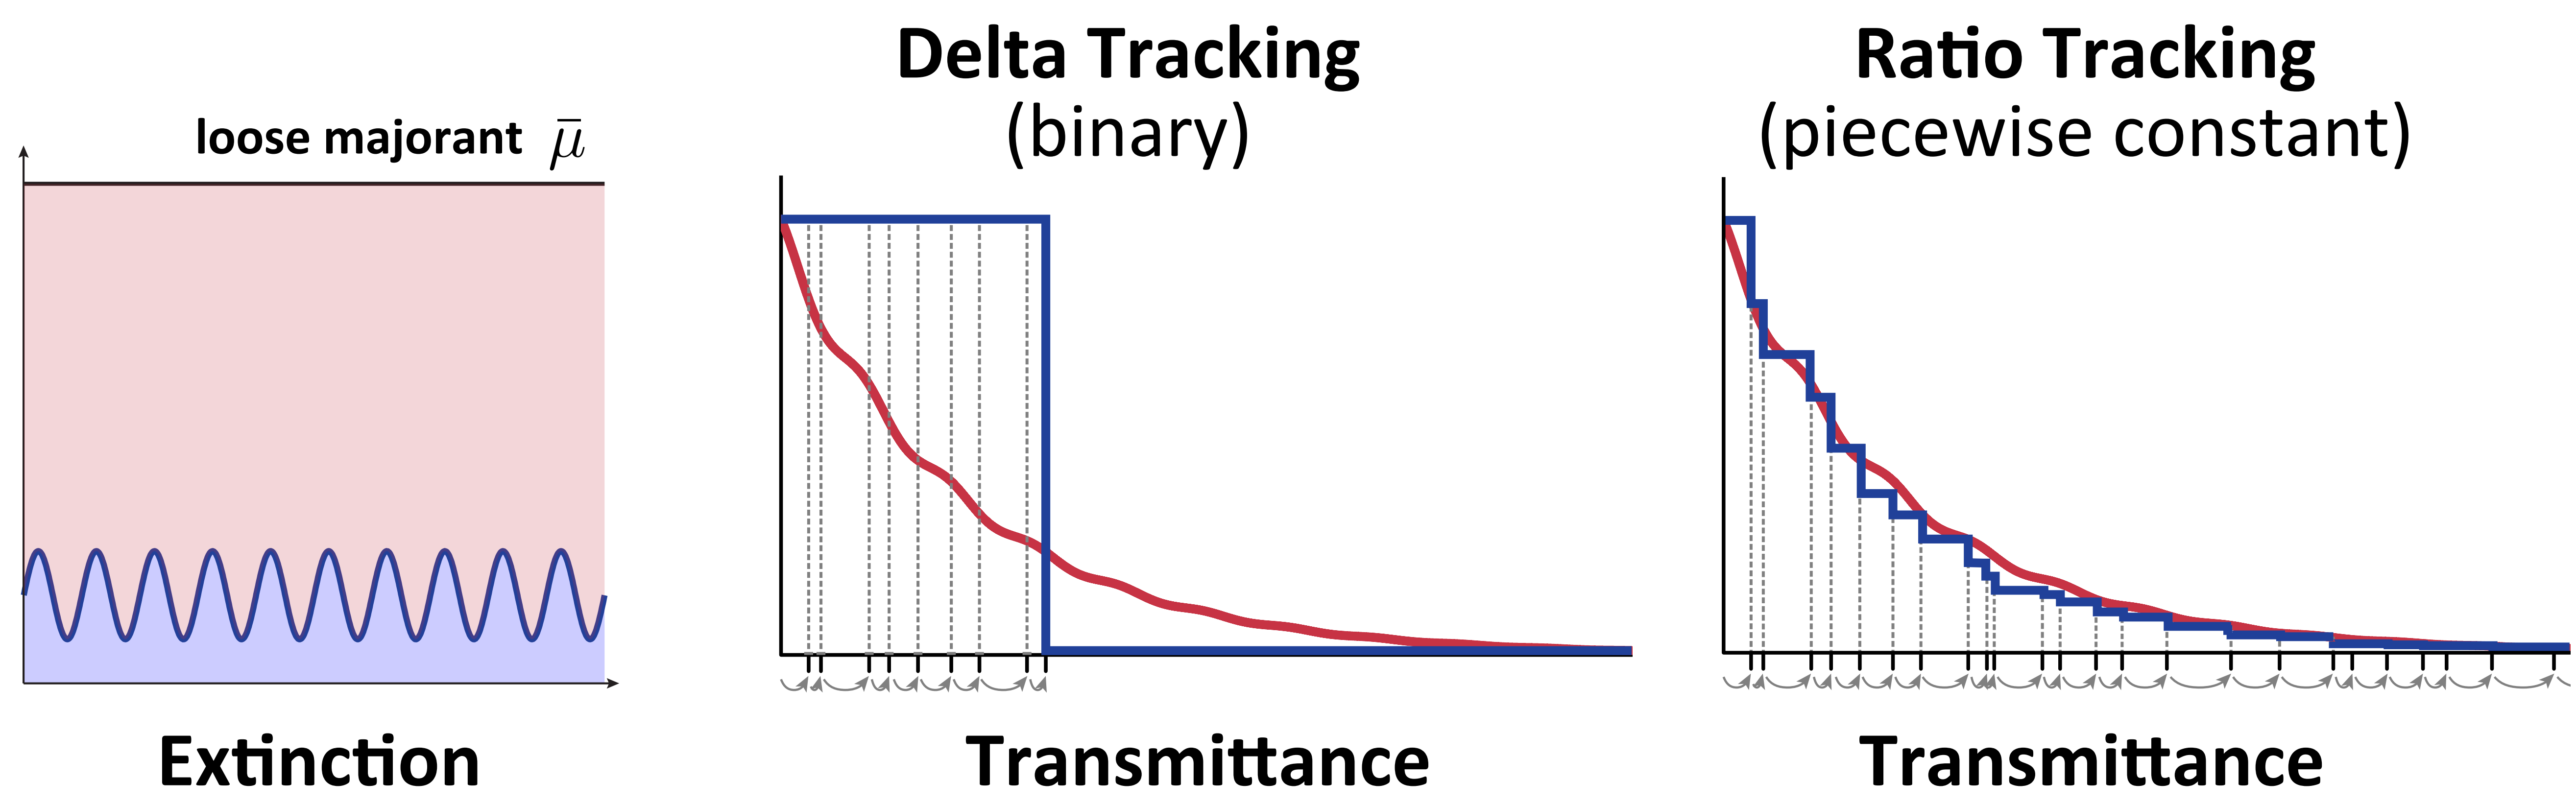

Let’s look at a few illustrations to better understand why. When the majorant is loose, the Ratio tracking provides a much finer, piecewise constant approximation. You can see that it better matches the red reference. This is because it leverages the tentative collisions more efficiently than Delta tracking.

Delta Tracking vs. Ratio Tracking

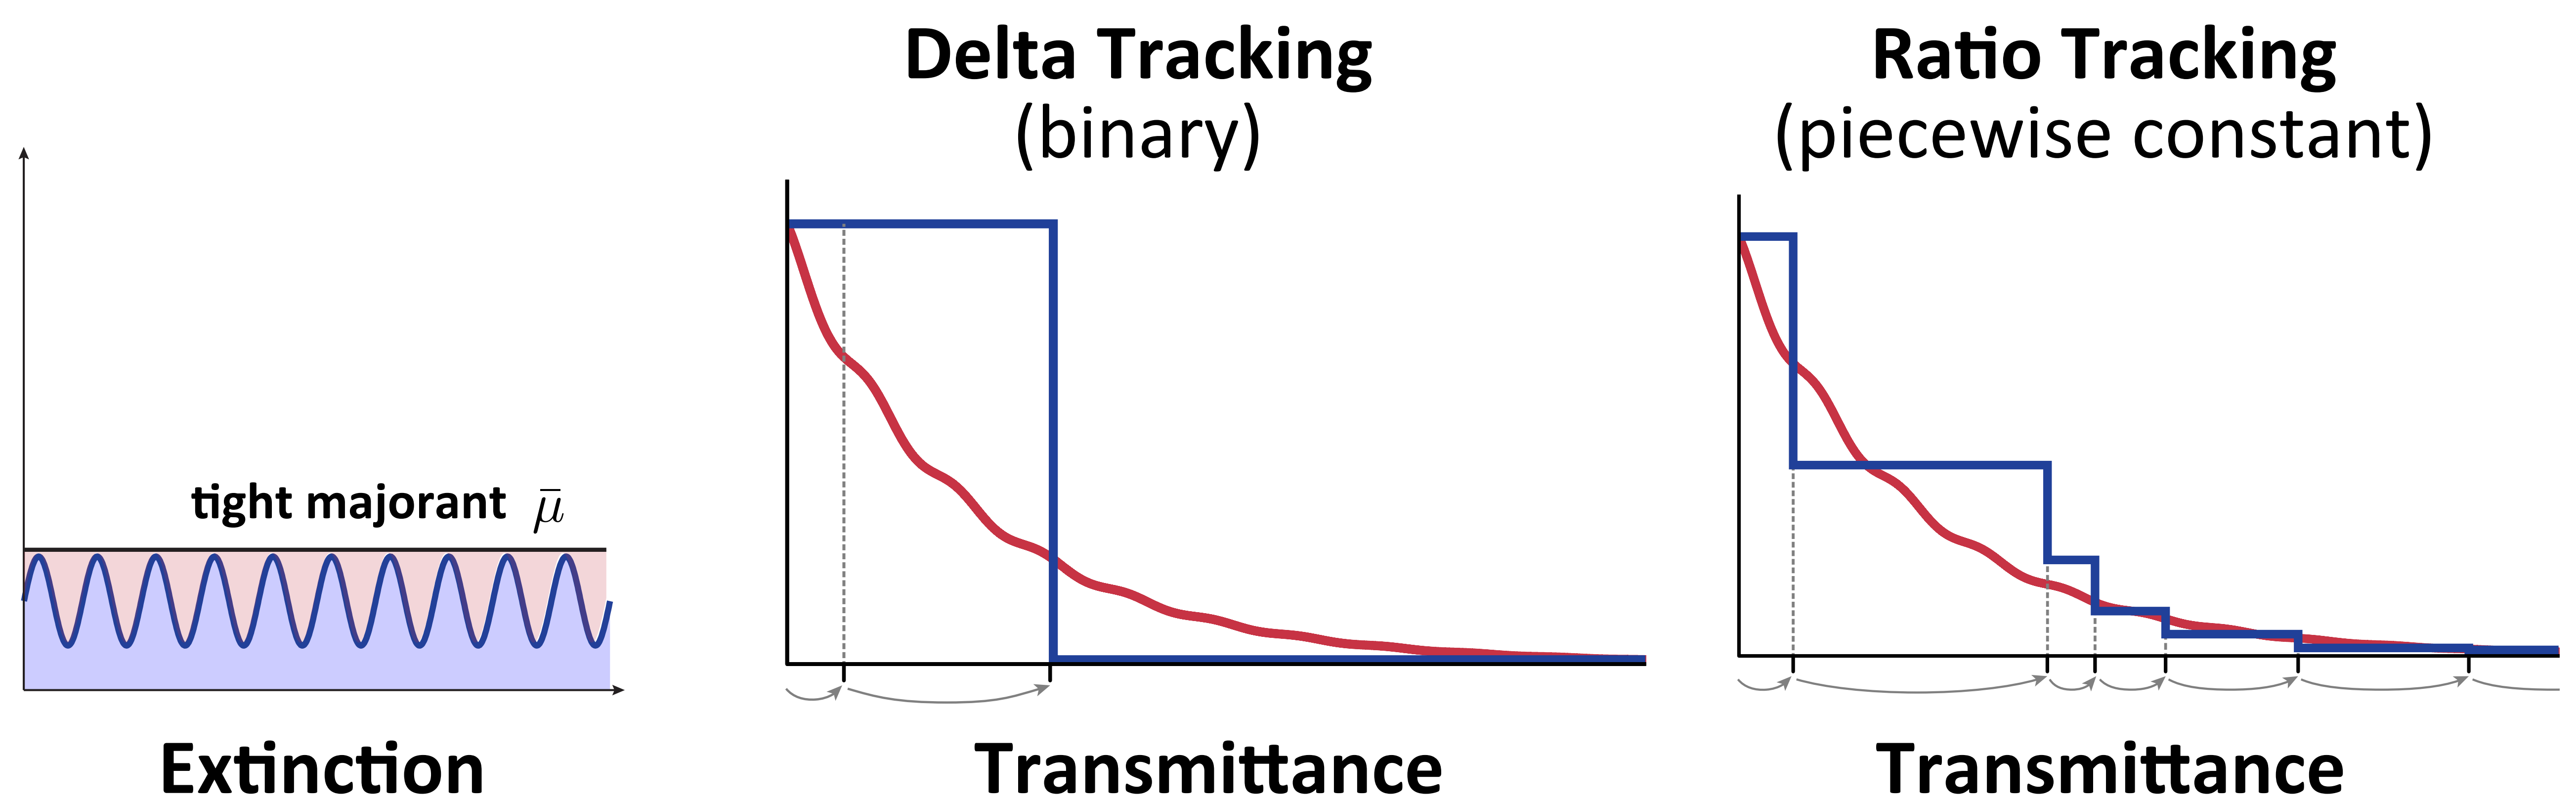

If the majorant tightly bounds the extinction, the Ratio tracking is still better, but the advantages over Delta tracking may not be as significant.

Delta Tracking vs. Ratio Tracking

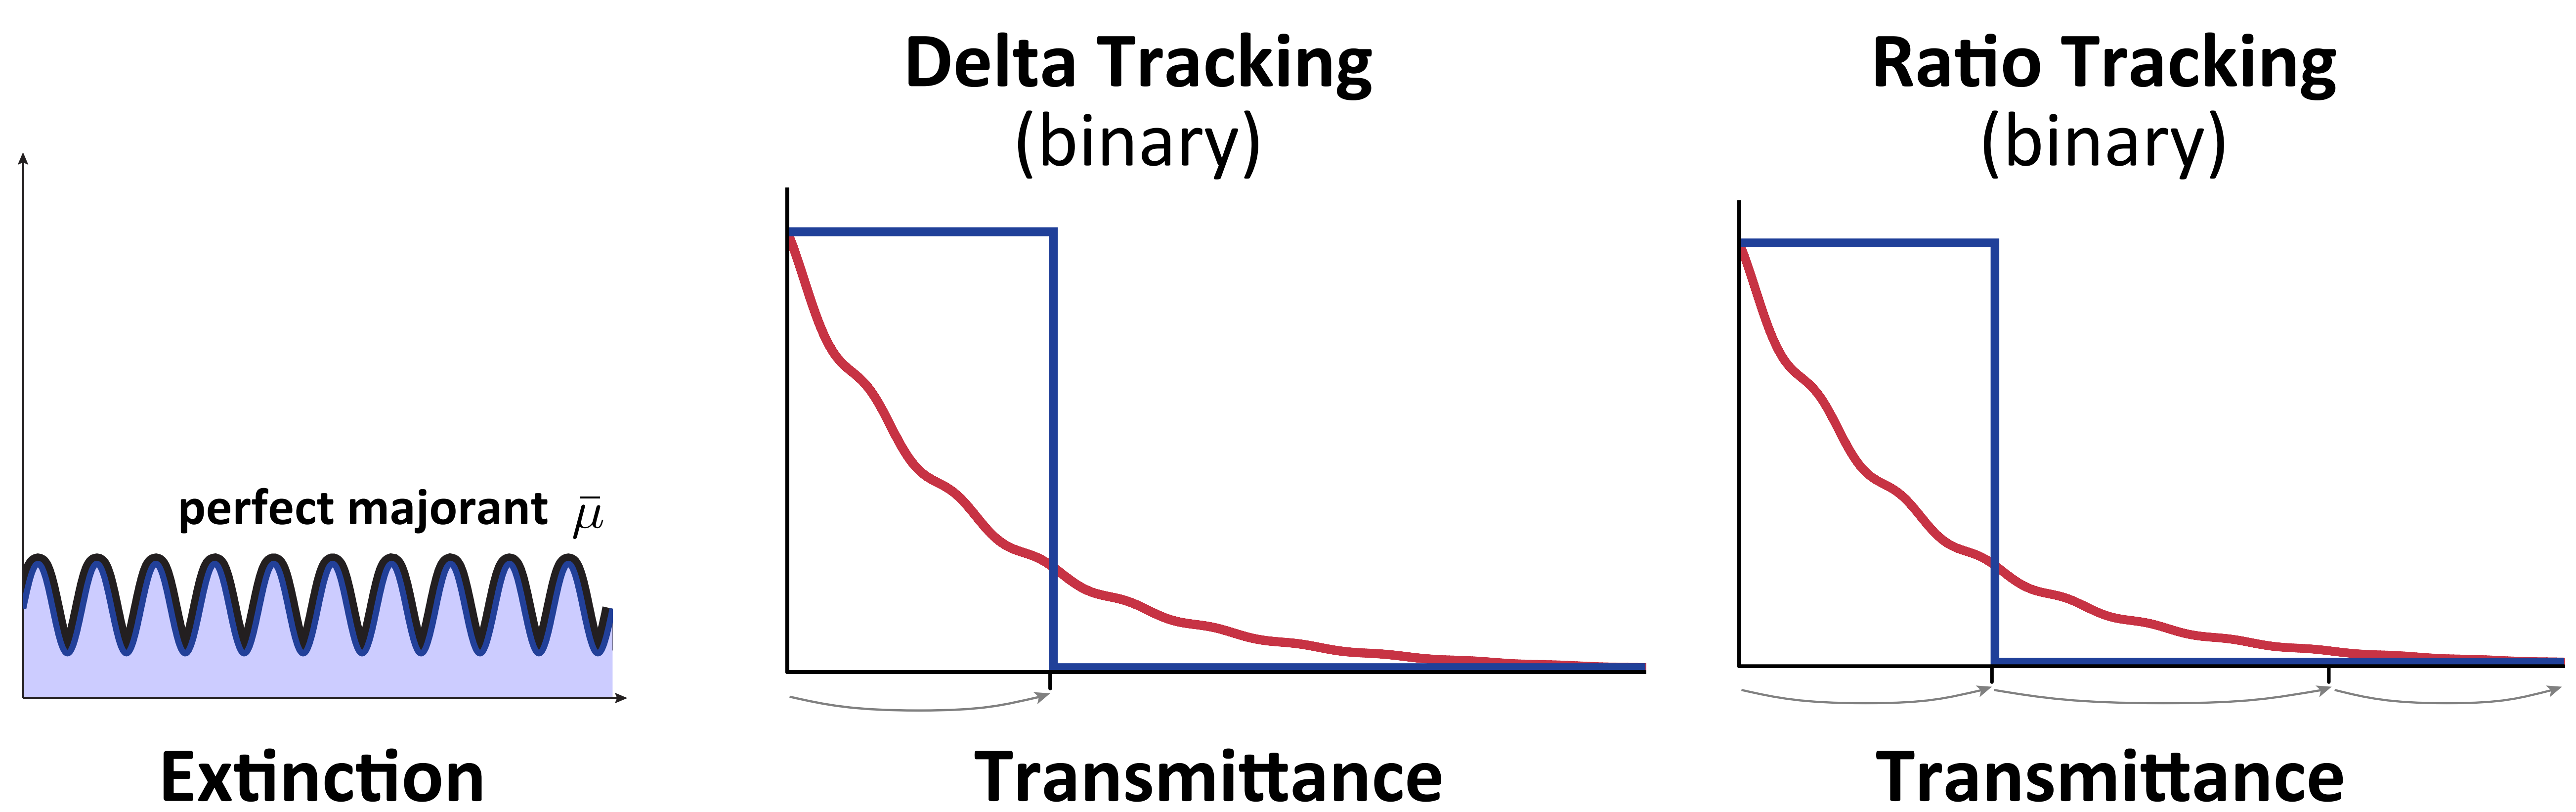

And in cases when the majorant is perfect, meaning that we don't need to add any fictitious particles to sample free-paths, both trackings provide just a binary transmittance estimate. So Ratio tracking helps when there are many fictitious particles, but it is less useful when there are only a few, or none.

Simple Scene

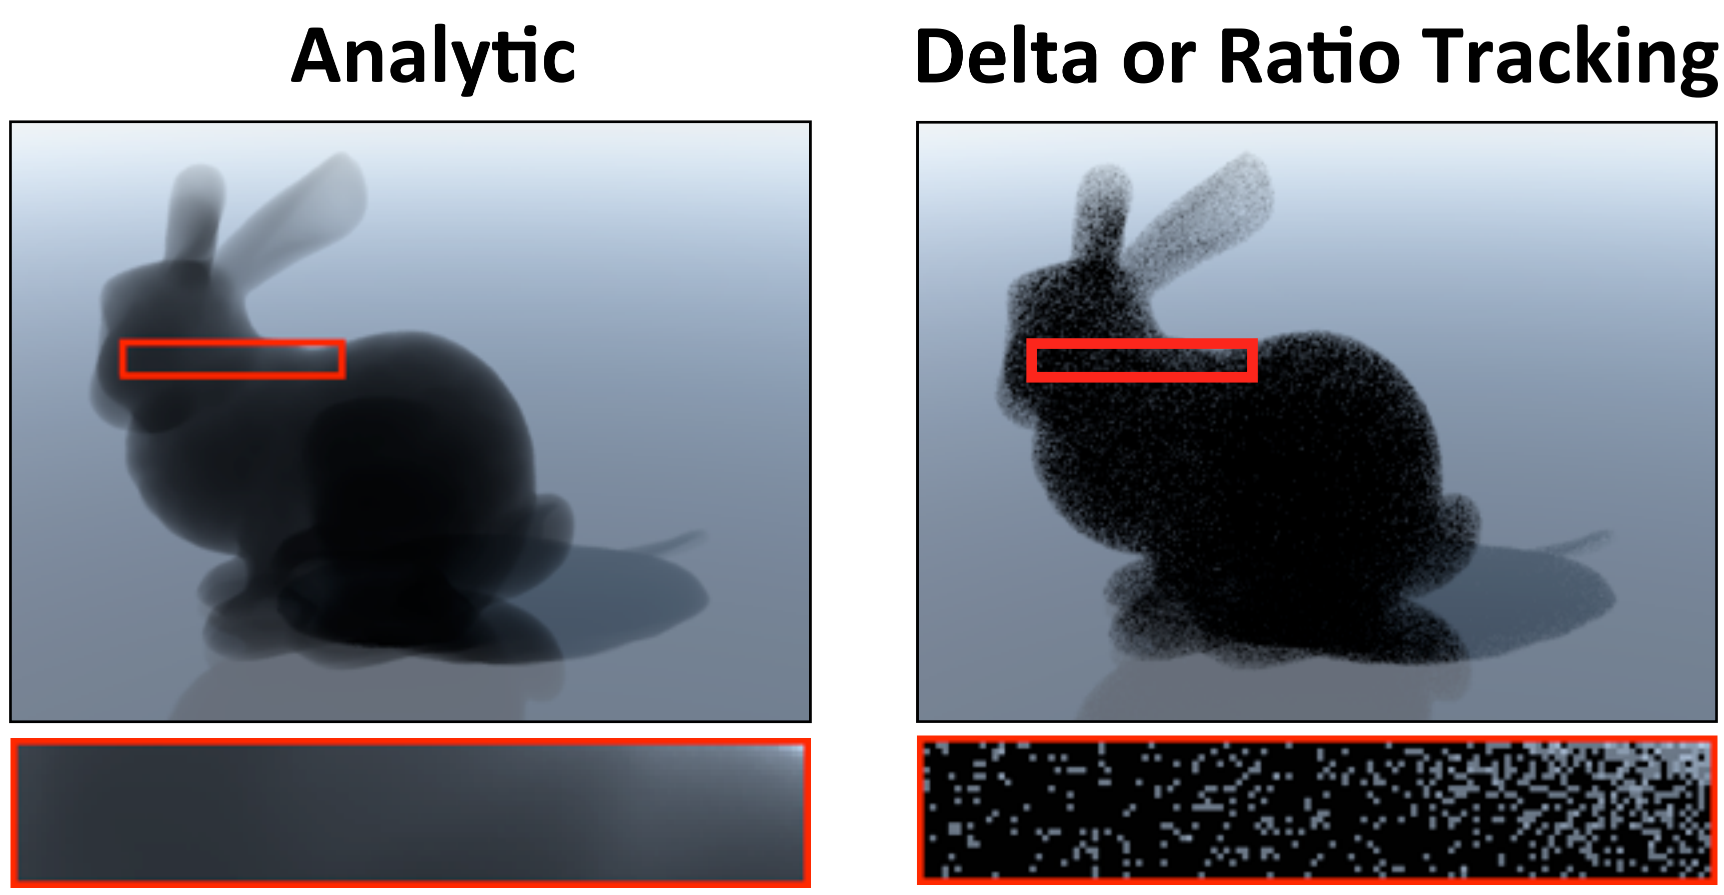

Here is one such example... a homogeneous medium bounded by an indexed-matched interface. Since the medium is homogeneous, we could evaluate the transmittance analytically.

But let's say we don't know this a-priori, and we use the Delta or Ratio tracking. Both techniques provide binary transmittance estimates yielding high variance.

So our next goal is to reduce the noise also in these cases, when the medium is homogeneous, or has a small degree of heterogeneity.

Simple Scene

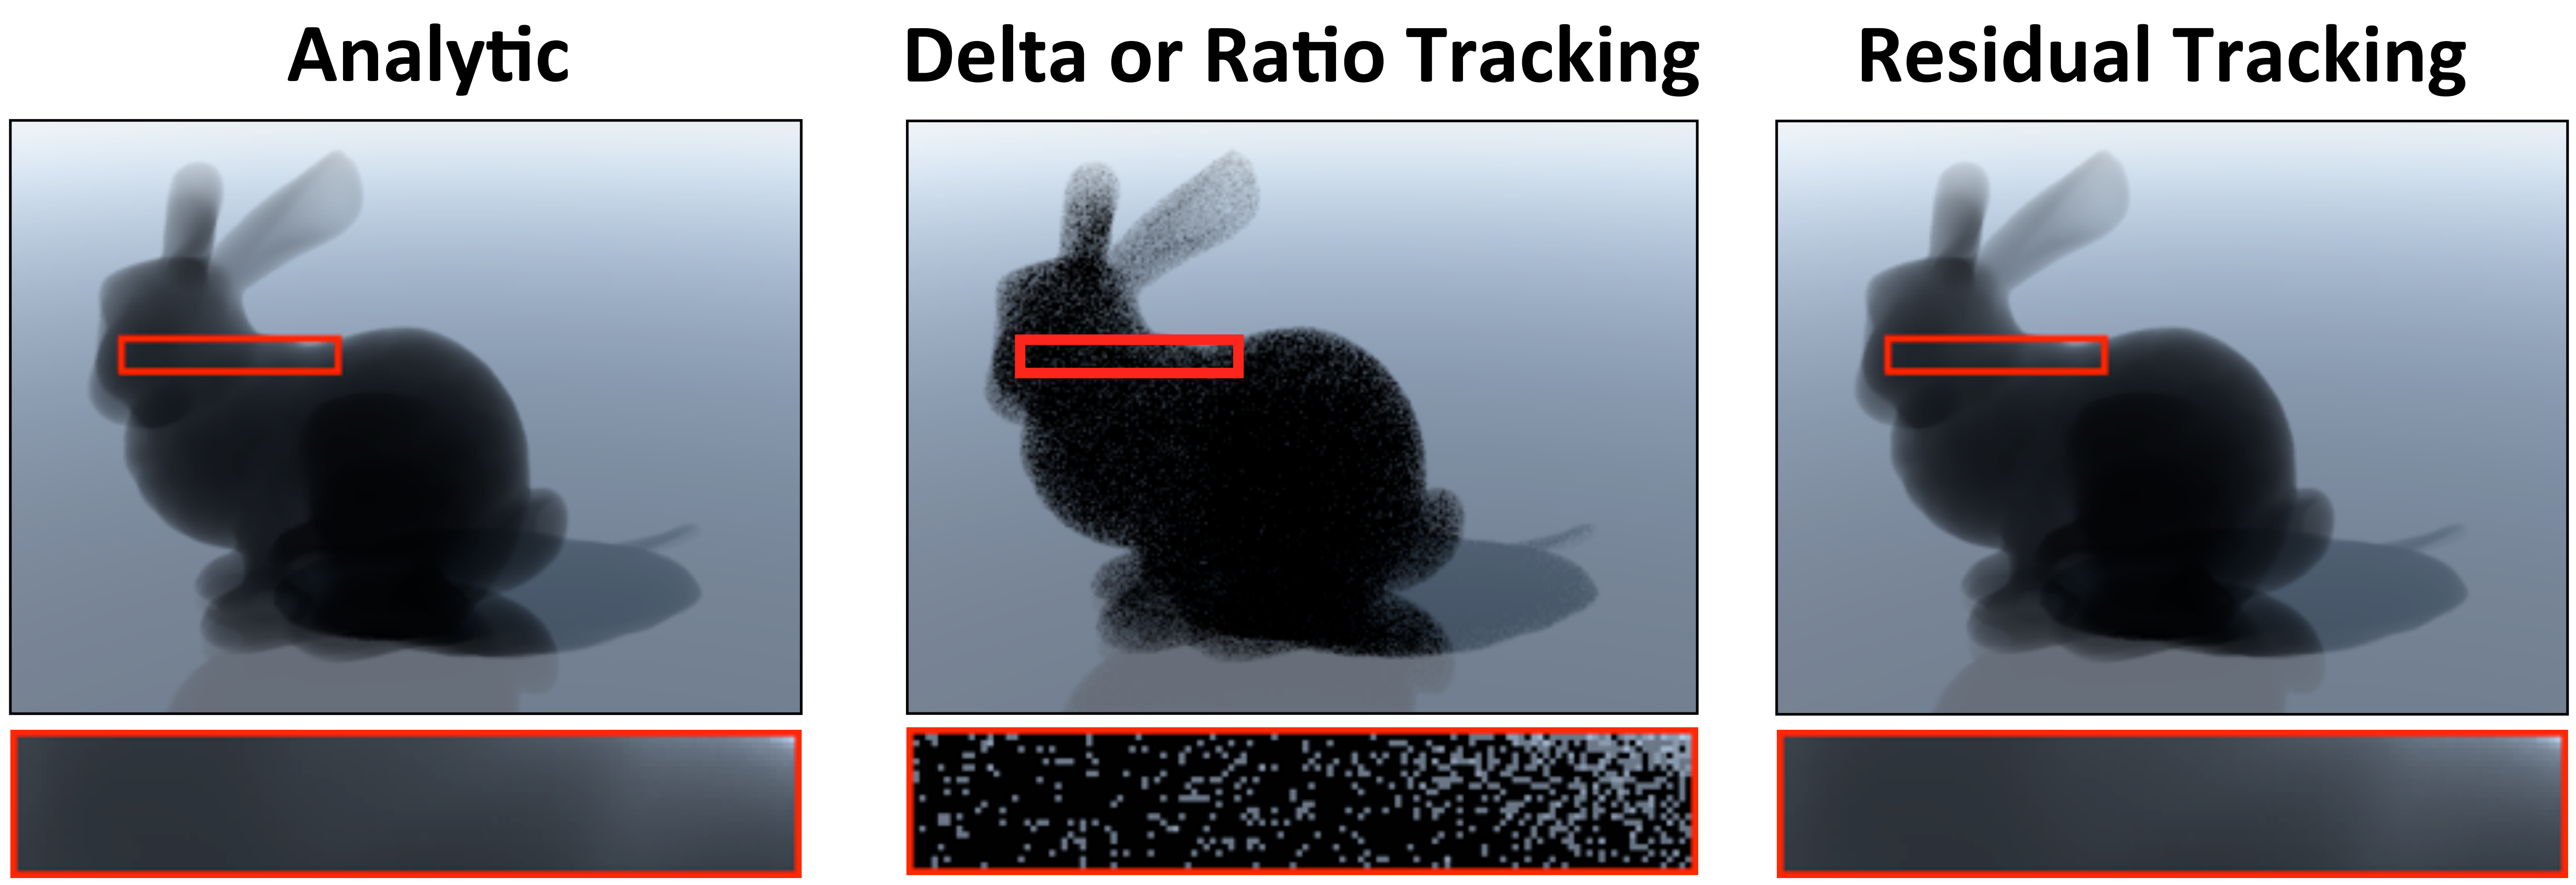

Here is one such example... a homogeneous medium bounded by an indexed-matched interface. Since the medium is homogeneous, we could evaluate the transmittance analytically.

But let's say we don't know this a-priori, and we use the Delta or Ratio tracking. Both techniques provide binary transmittance estimates yielding high variance.

So our next goal is to reduce the noise also in these cases, when the medium is homogeneous, or has a small degree of heterogeneity.

Residual Tracking

“from piecewise constant to piecewise exponential”

I will now talk about our second variance reduction approach called Residual tracking, which is inspired by control variates.

Residual Tracking

We attempt to reduce the variance by evaluating part of the transmittance analytically, and numerically estimate only the remainder. For this, we decompose the original medium into a sum of a control medium and a residual medium.

The transmittance can be then computed as the product of transmittances evaluated independently through each of the two media. As long as the control extinction matches the original one well,
the error of the numerical estimation will be reduced.

Residual Tracking

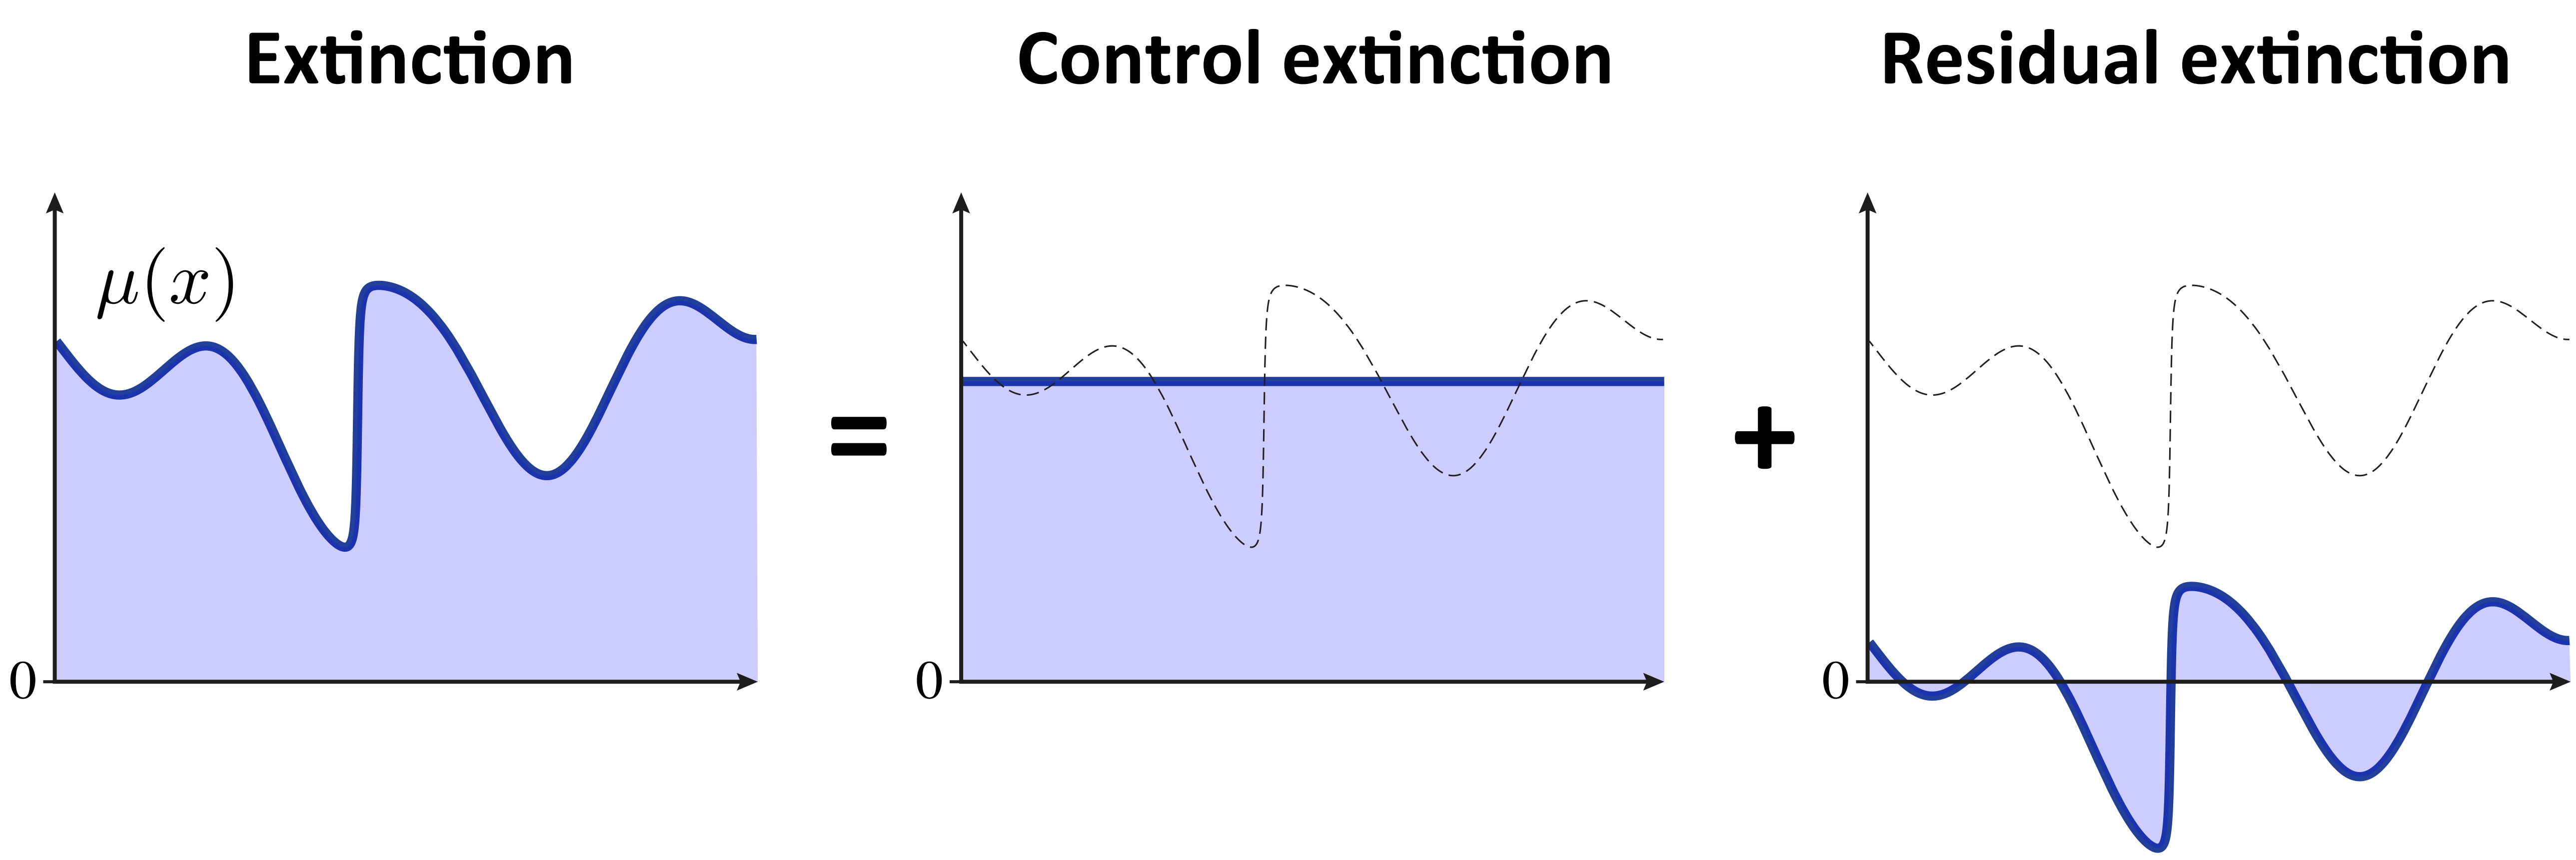

We attempt to reduce the variance by evaluating part of the transmittance analytically, and numerically estimate only the remainder. For this, we decompose the original medium into a sum of a control medium and a residual medium.

The transmittance can be then computed as the product of transmittances evaluated independently through each of the two media. As long as the control extinction matches the original one well,
the error of the numerical estimation will be reduced.

Residual Tracking

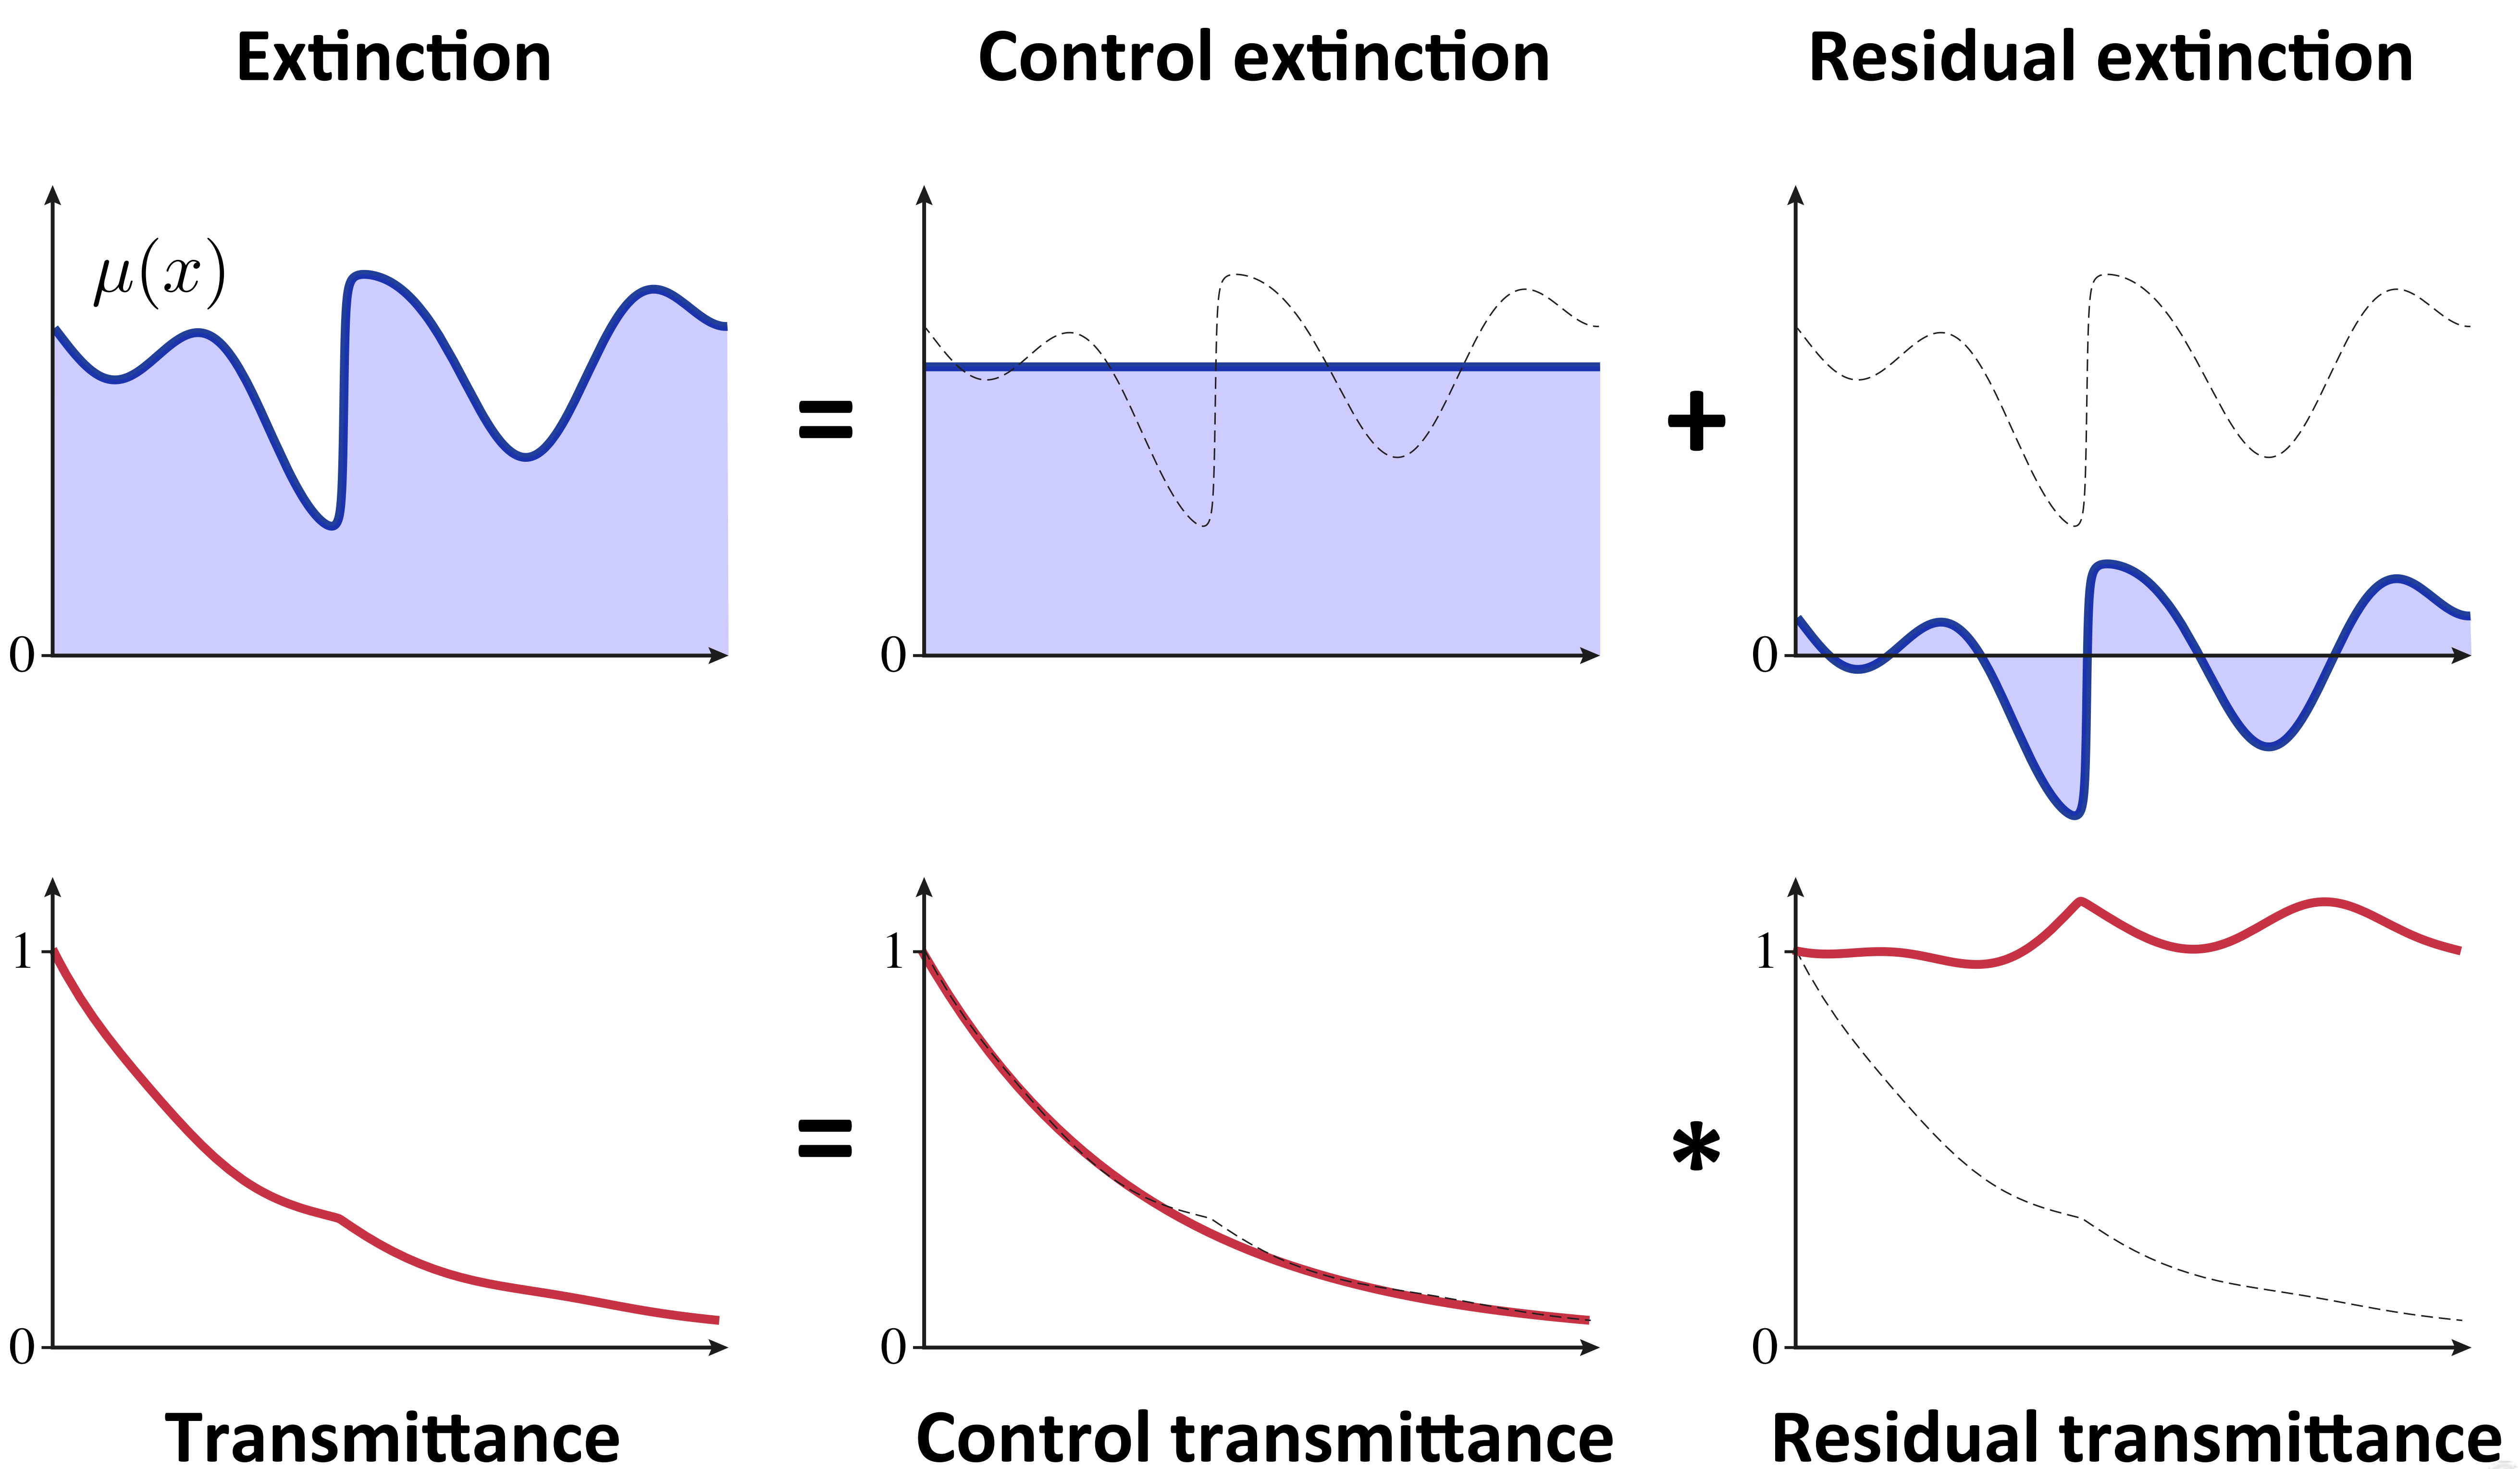

We attempt to reduce the variance by evaluating part of the transmittance analytically, and numerically estimate only the remainder. For this, we decompose the original medium into a sum of a control medium and a residual medium.

The transmittance can be then computed as the product of transmittances evaluated independently through each of the two media. As long as the control extinction matches the original one well,
the error of the numerical estimation will be reduced.

Residual Tracking

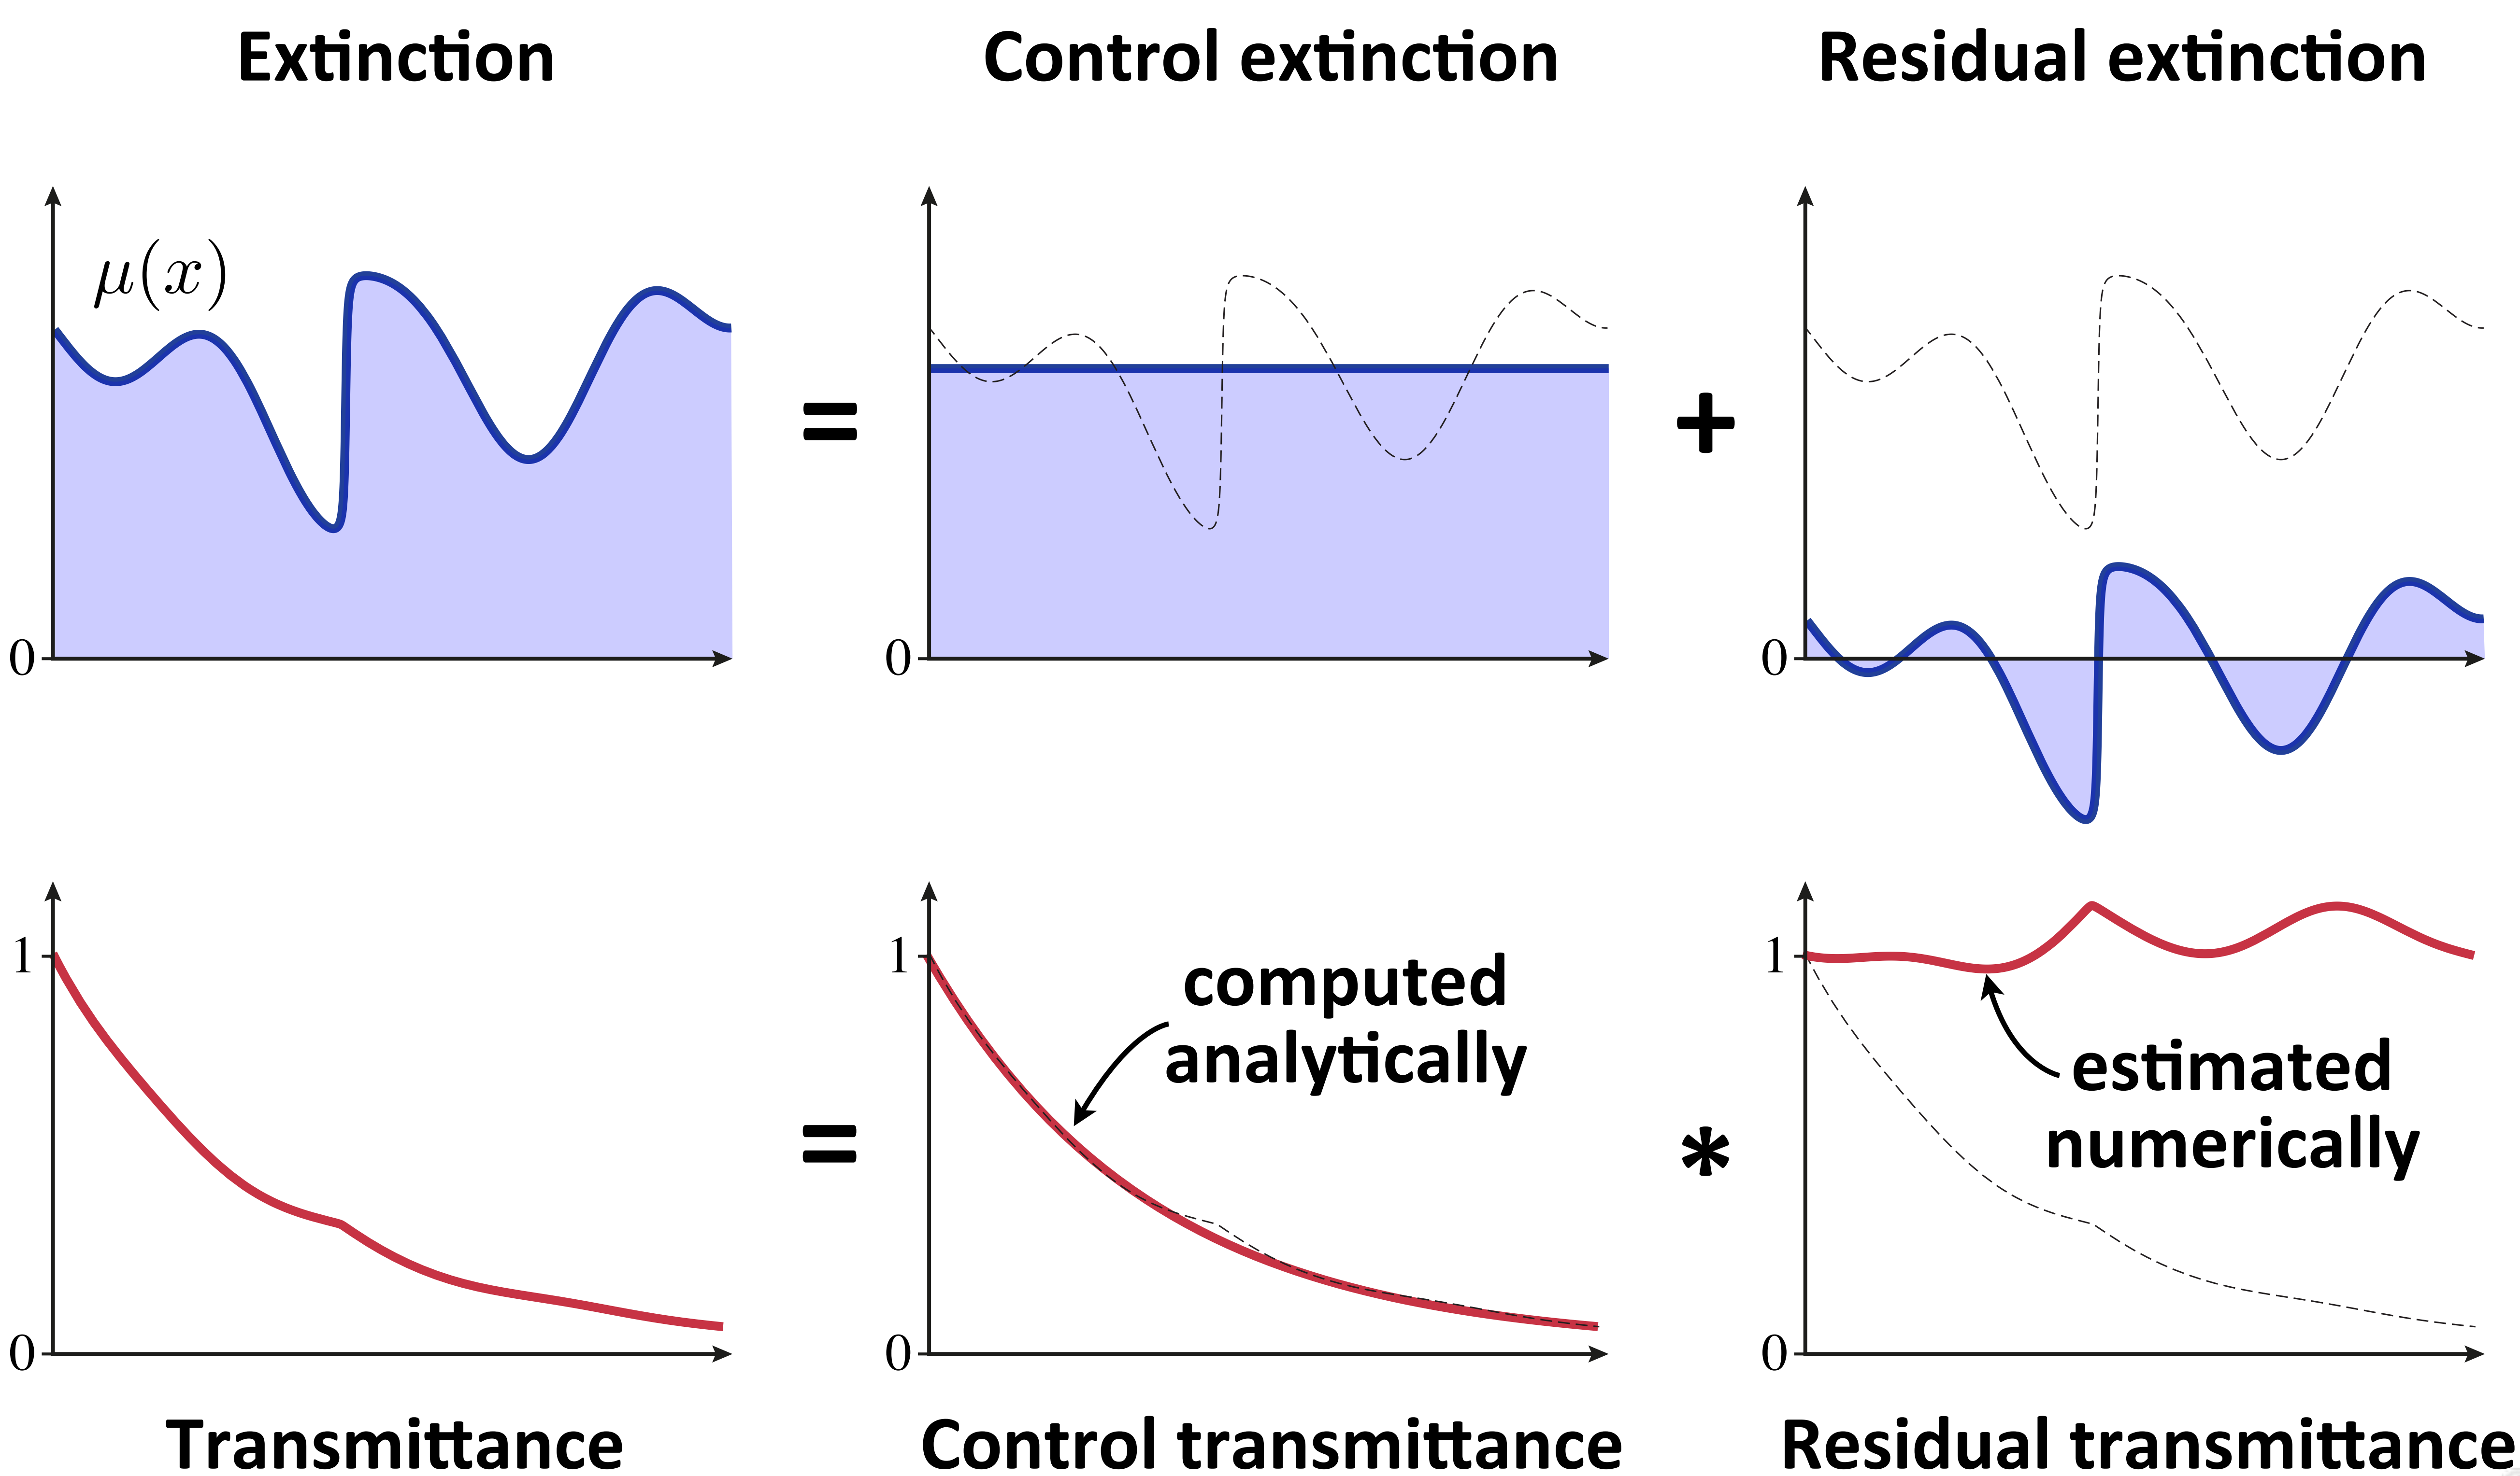

We attempt to reduce the variance by evaluating part of the transmittance analytically, and numerically estimate only the remainder. For this, we decompose the original medium into a sum of a control medium and a residual medium.

The transmittance can be then computed as the product of transmittances evaluated independently through each of the two media. As long as the control extinction matches the original one well,
the error of the numerical estimation will be reduced.

Residual Tracking

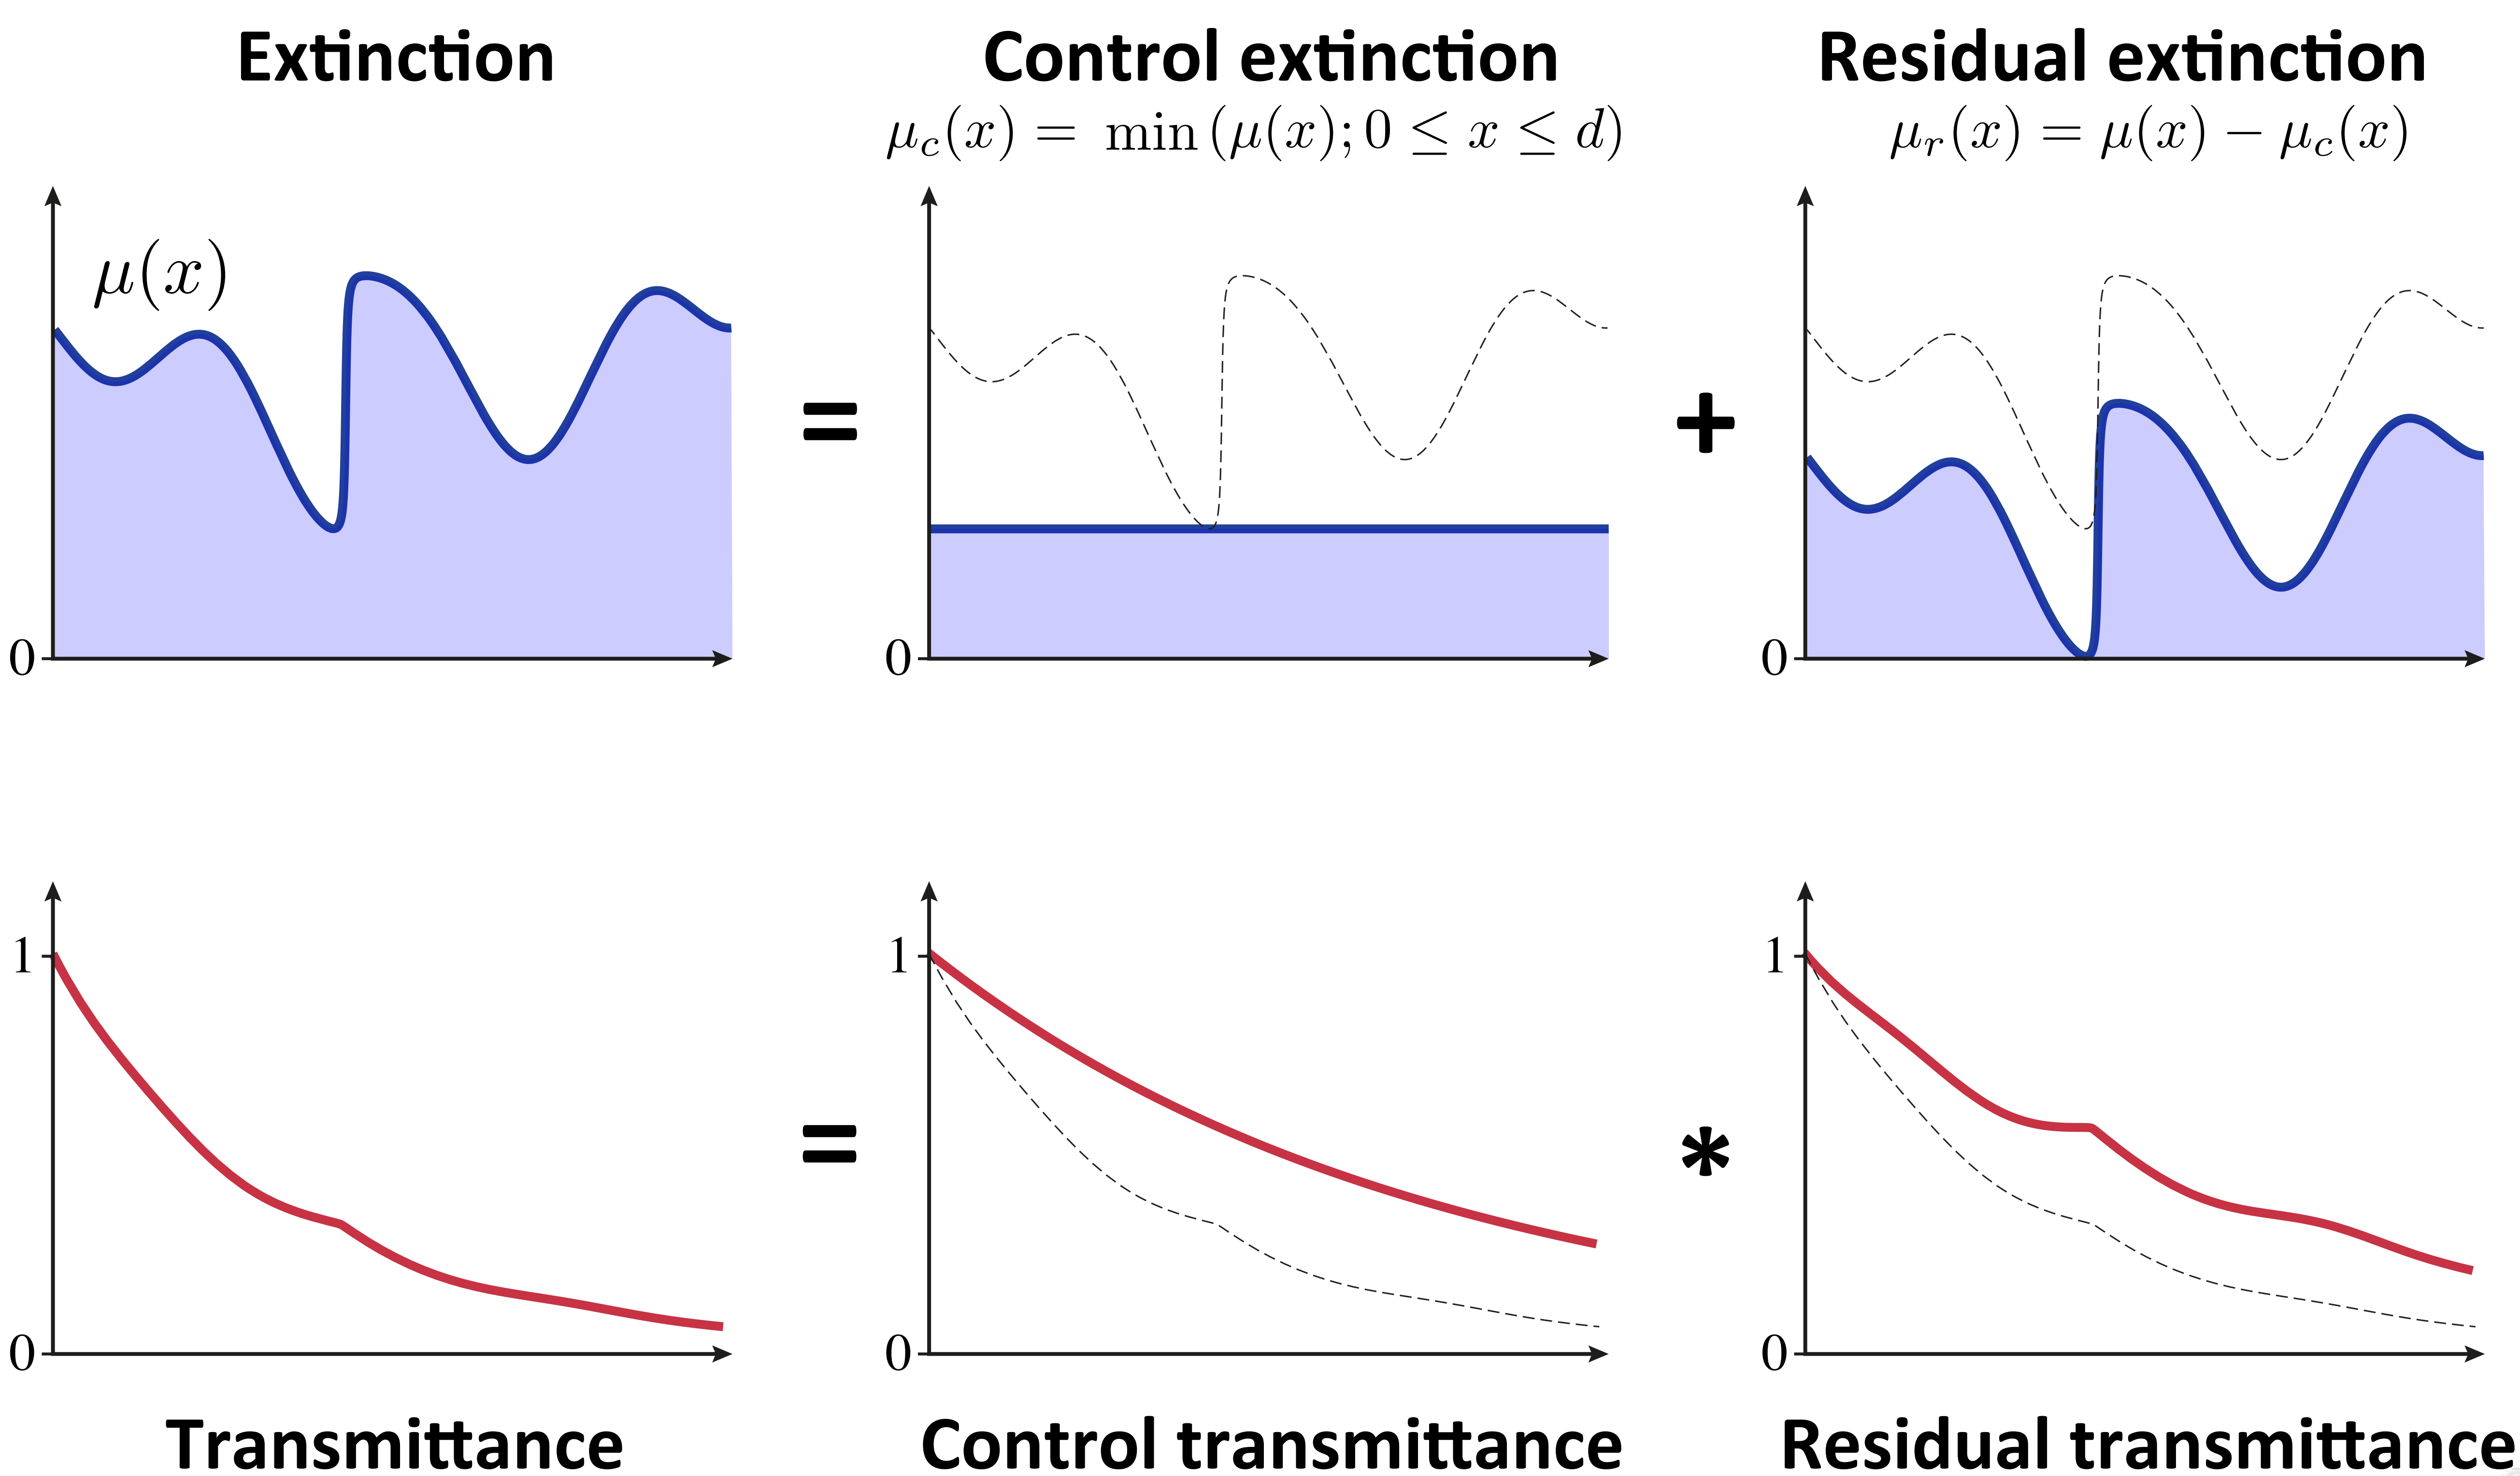

We have many options, how to choose the control extinction. We can for instance use a single global value, set to the minimum extinction in the volume, or to the average, or the maximum extinction over the entire volume, or even some arbitrary value.

You can see how it impacts the values of the control and residual transmittance.

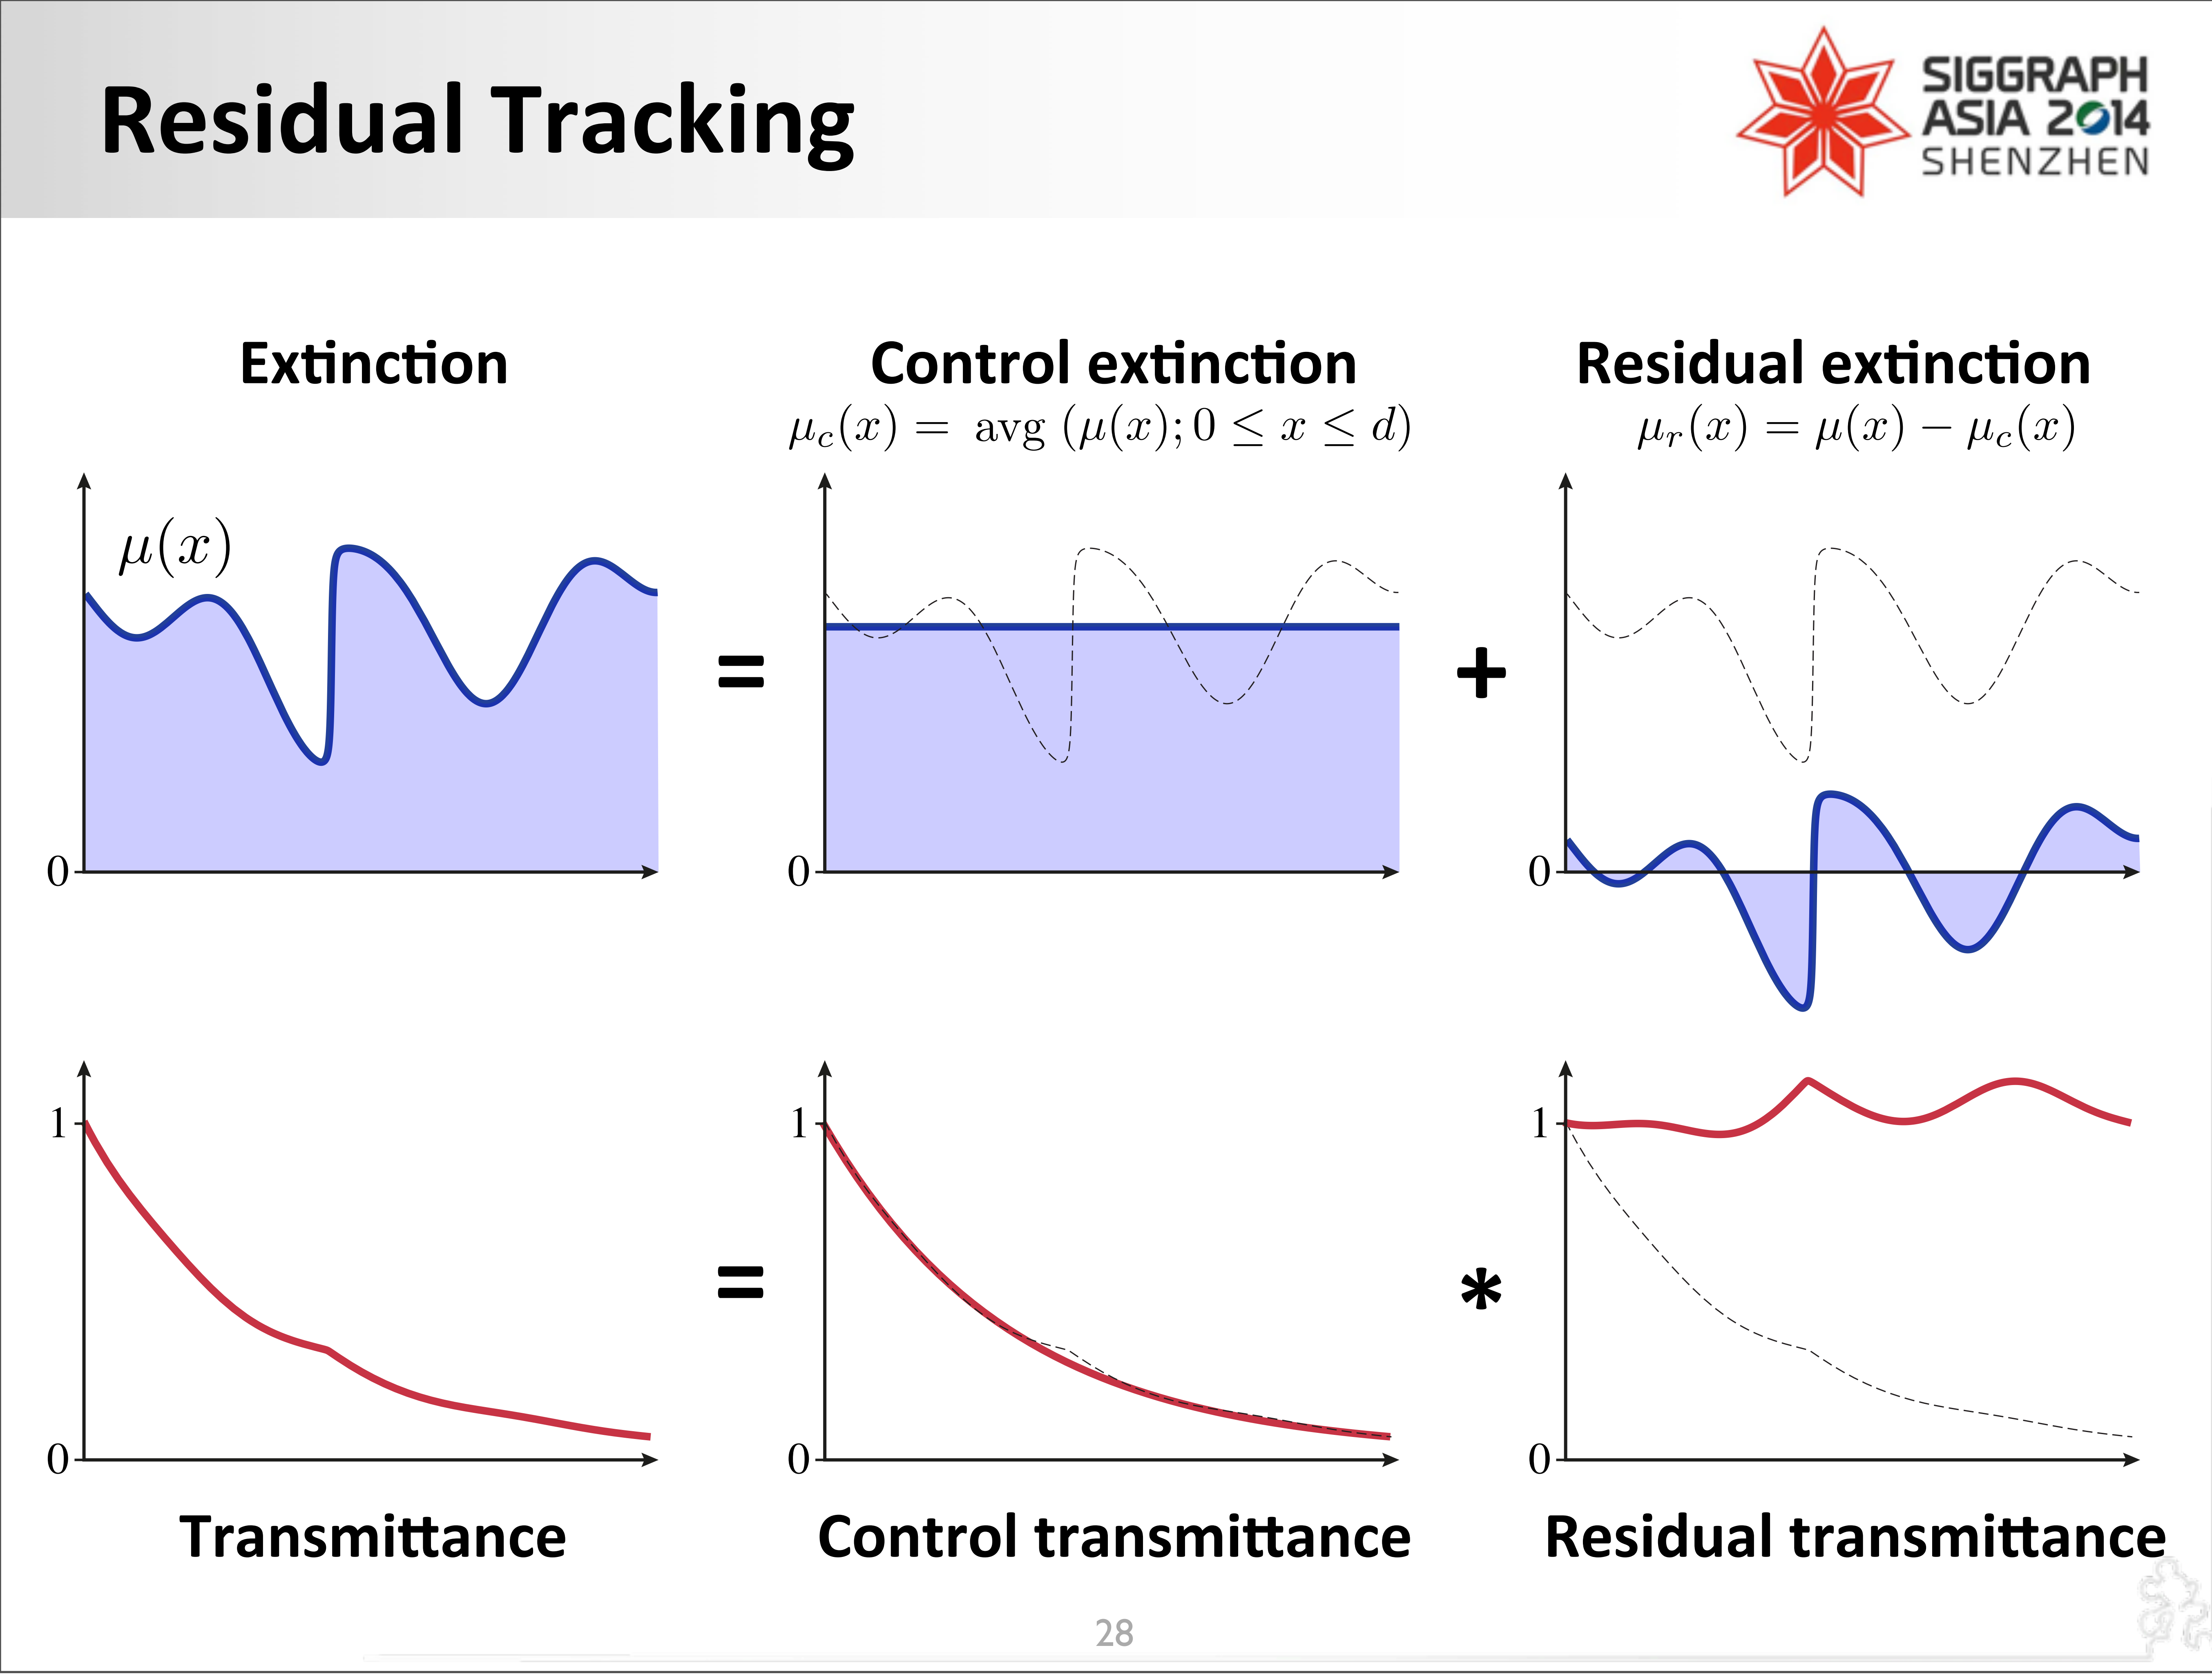

We have many options, how to choose the control extinction. We can for instance use a single global value, set to the minimum extinction in the volume, or to the average, or the maximum extinction over the entire volume, or even some arbitrary value.

You can see how it impacts the values of the control and residual transmittance.

Residual Tracking

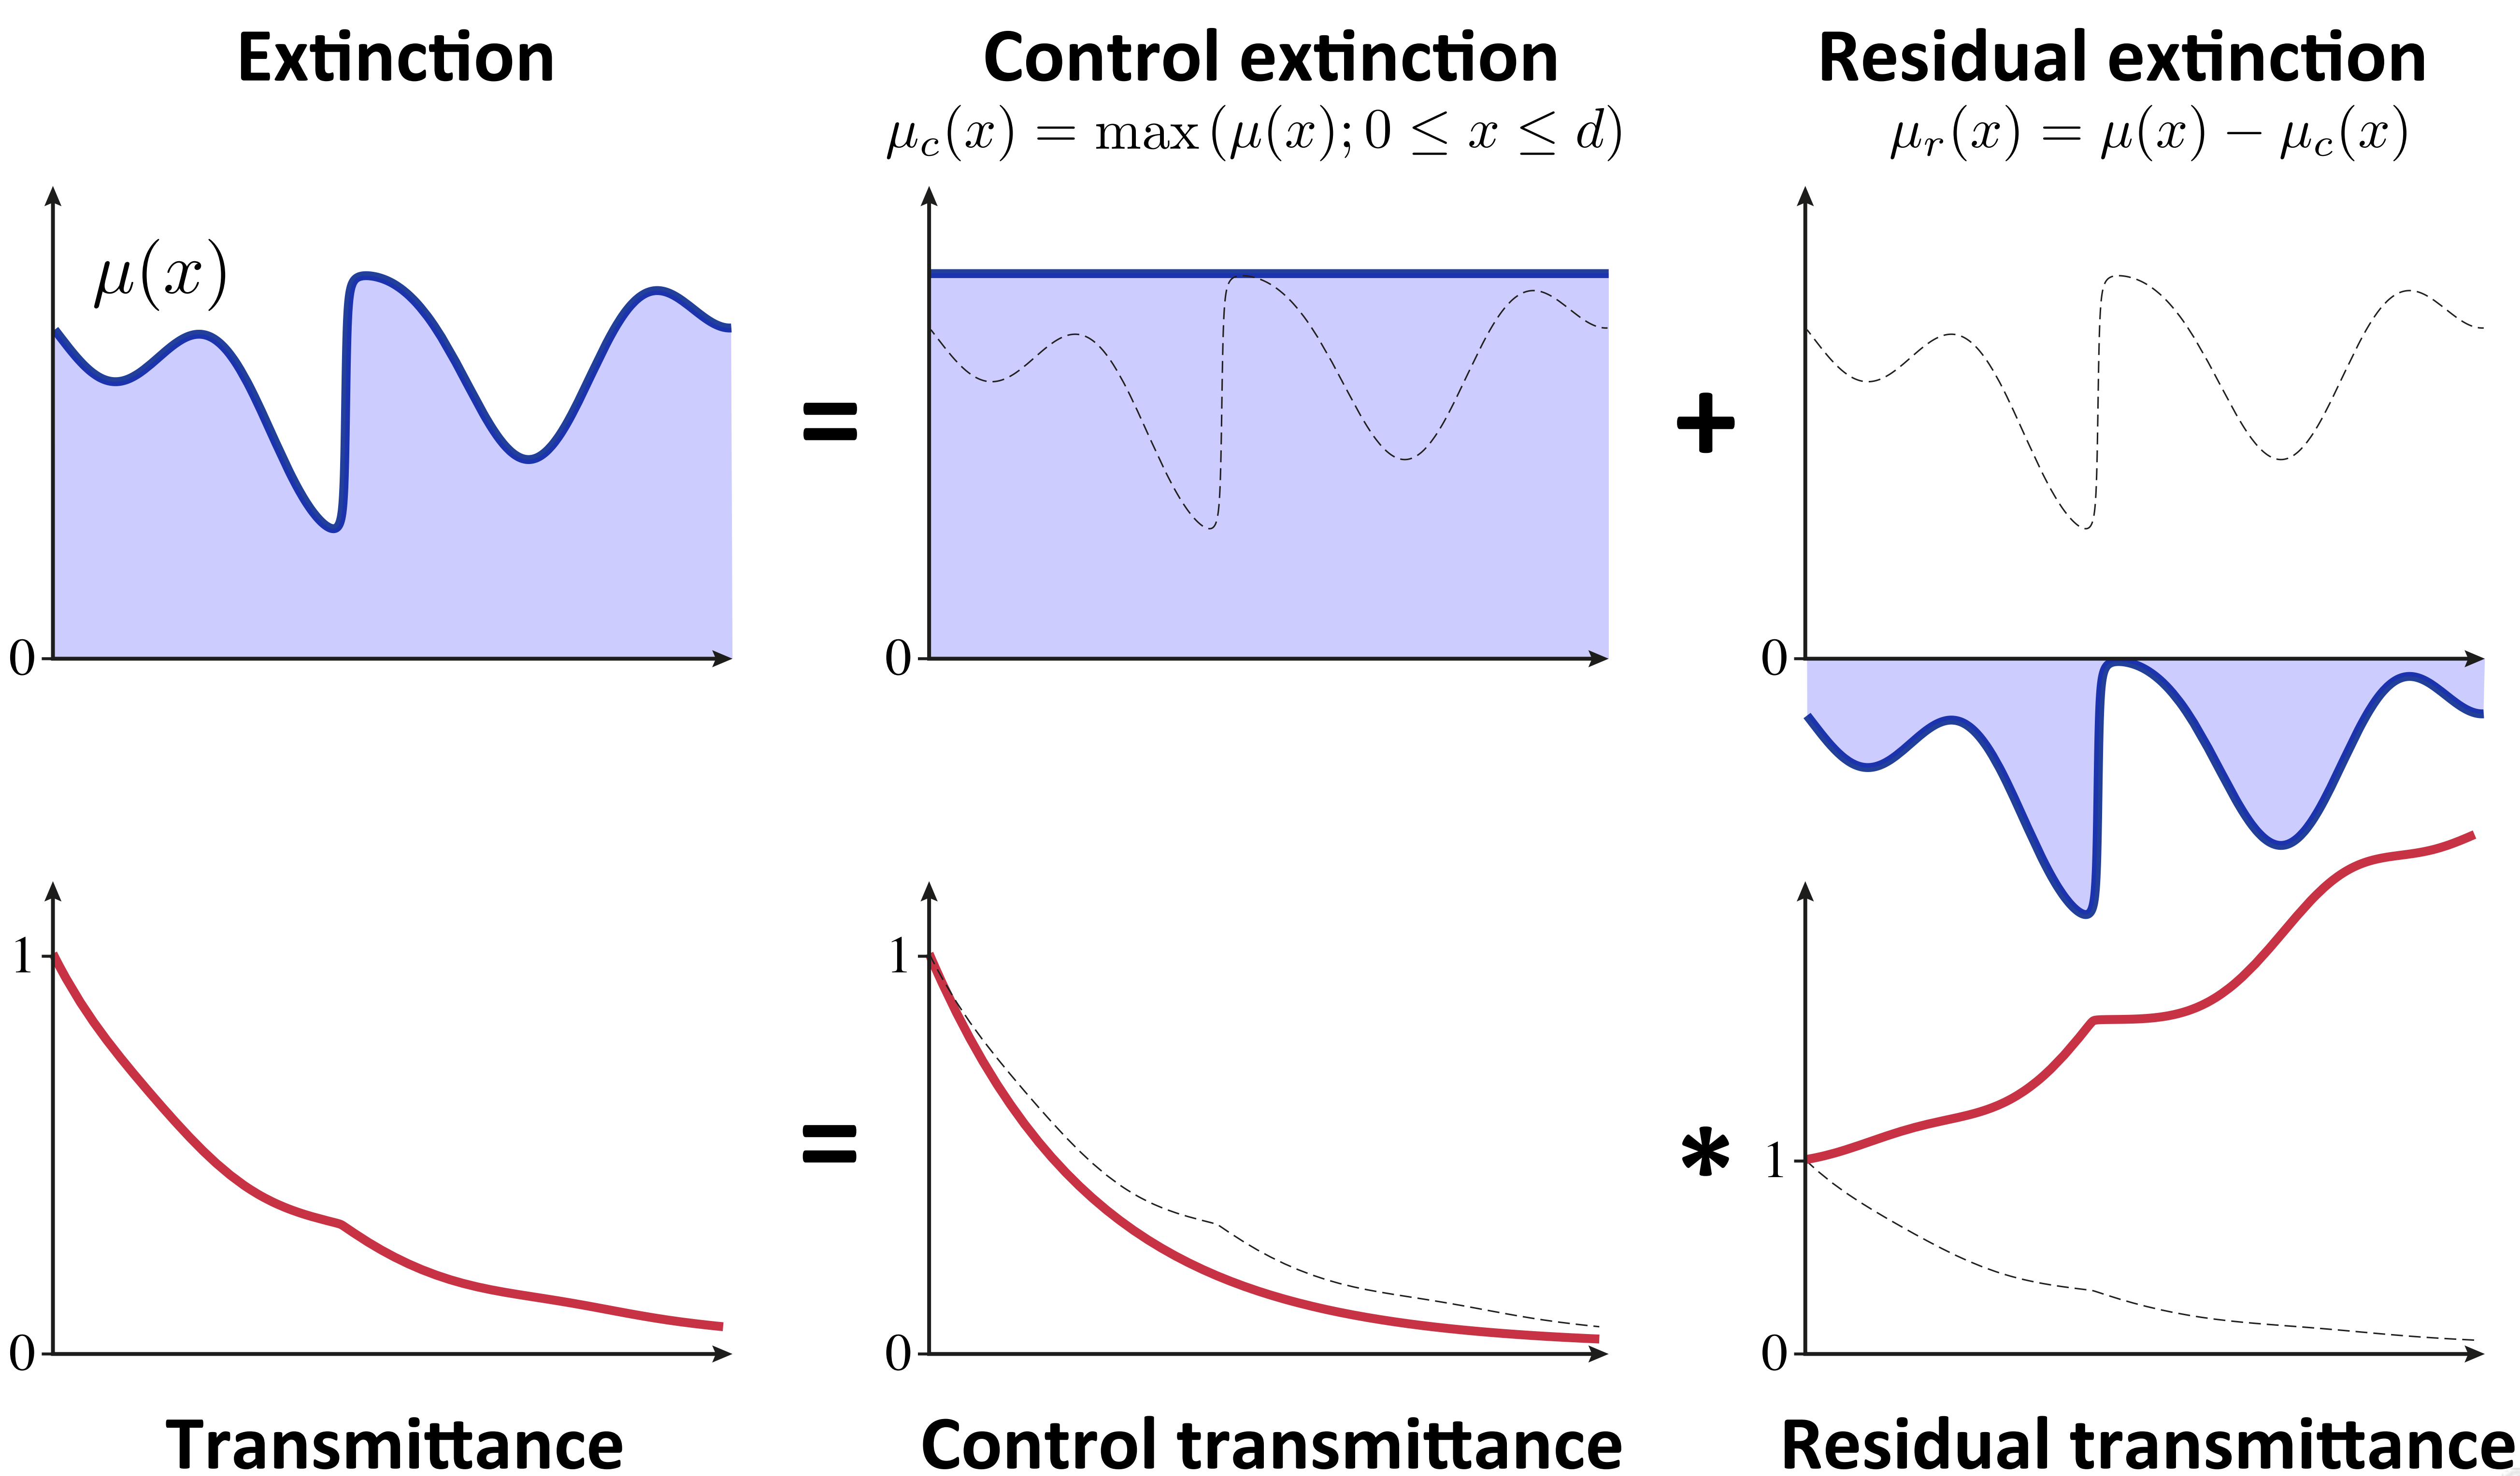

We have many options, how to choose the control extinction. We can for instance use a single global value, set to the minimum extinction in the volume, or to the average, or the maximum extinction over the entire volume, or even some arbitrary value.

You can see how it impacts the values of the control and residual transmittance.

Residual Tracking

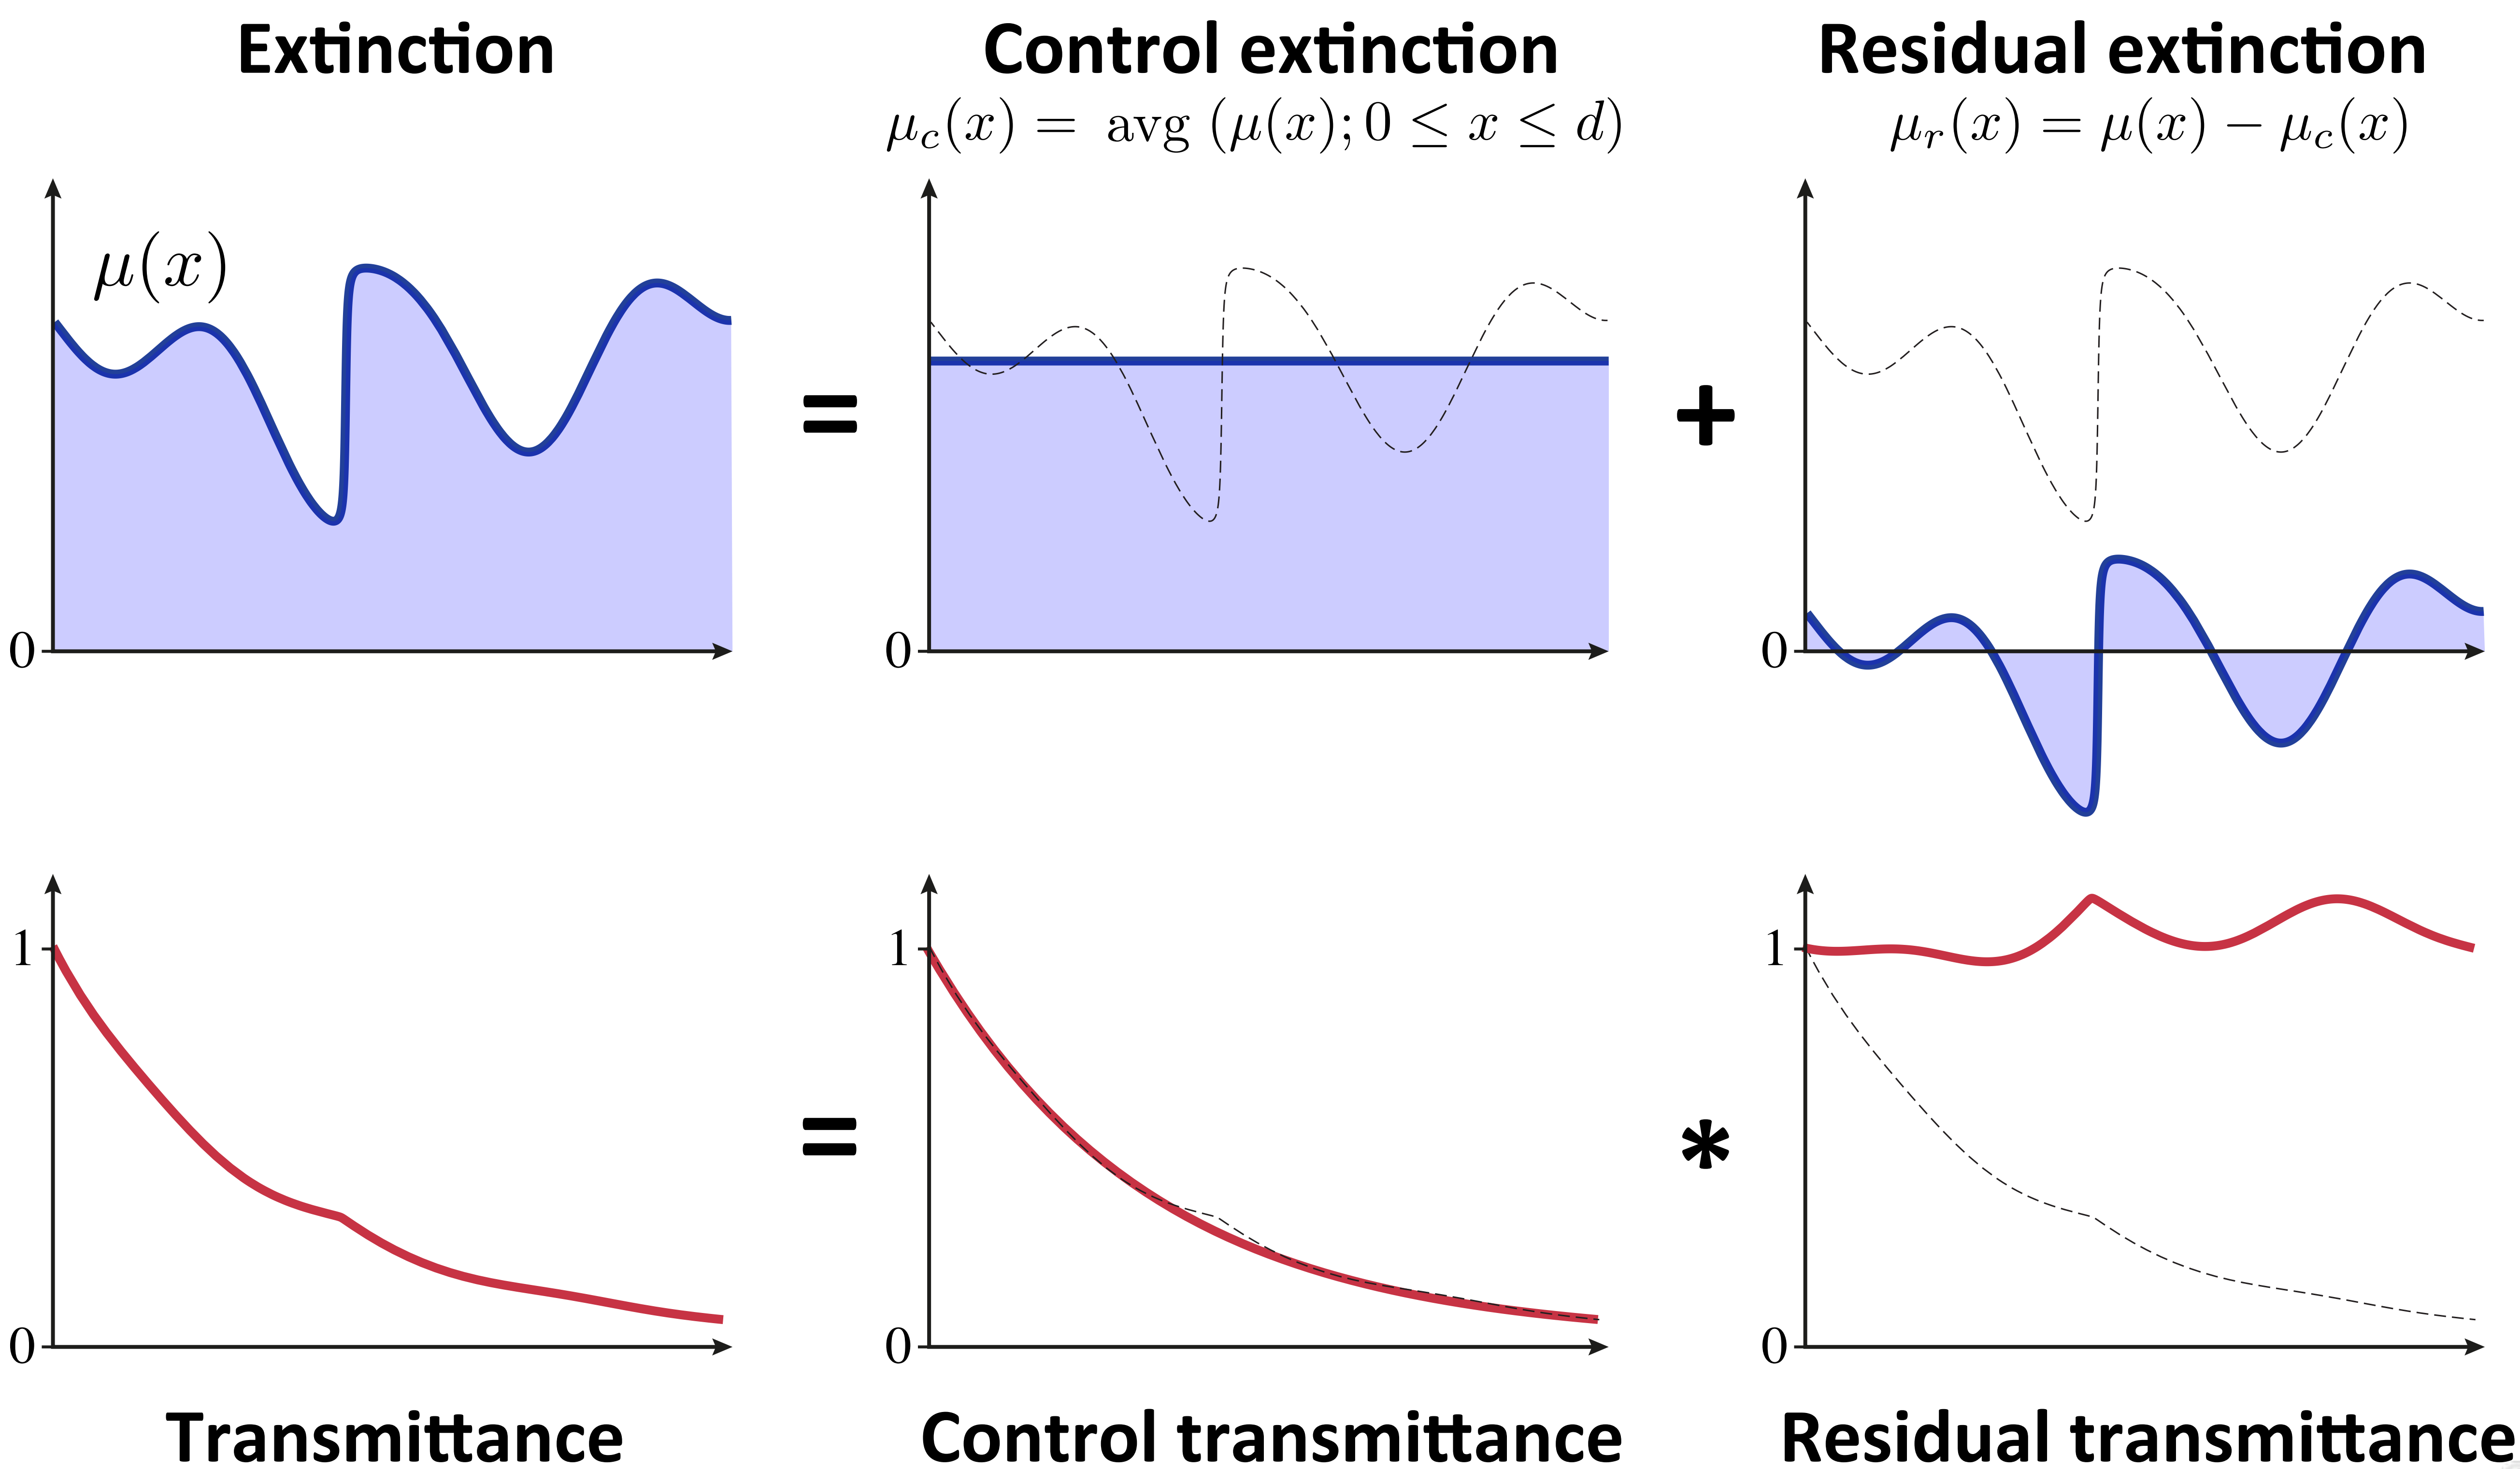

We have many options, how to choose the control extinction. We can for instance use a single global value, set to the minimum extinction in the volume, or to the average, or the maximum extinction over the entire volume, or even some arbitrary value.

You can see how it impacts the values of the control and residual transmittance.

Residual Tracking

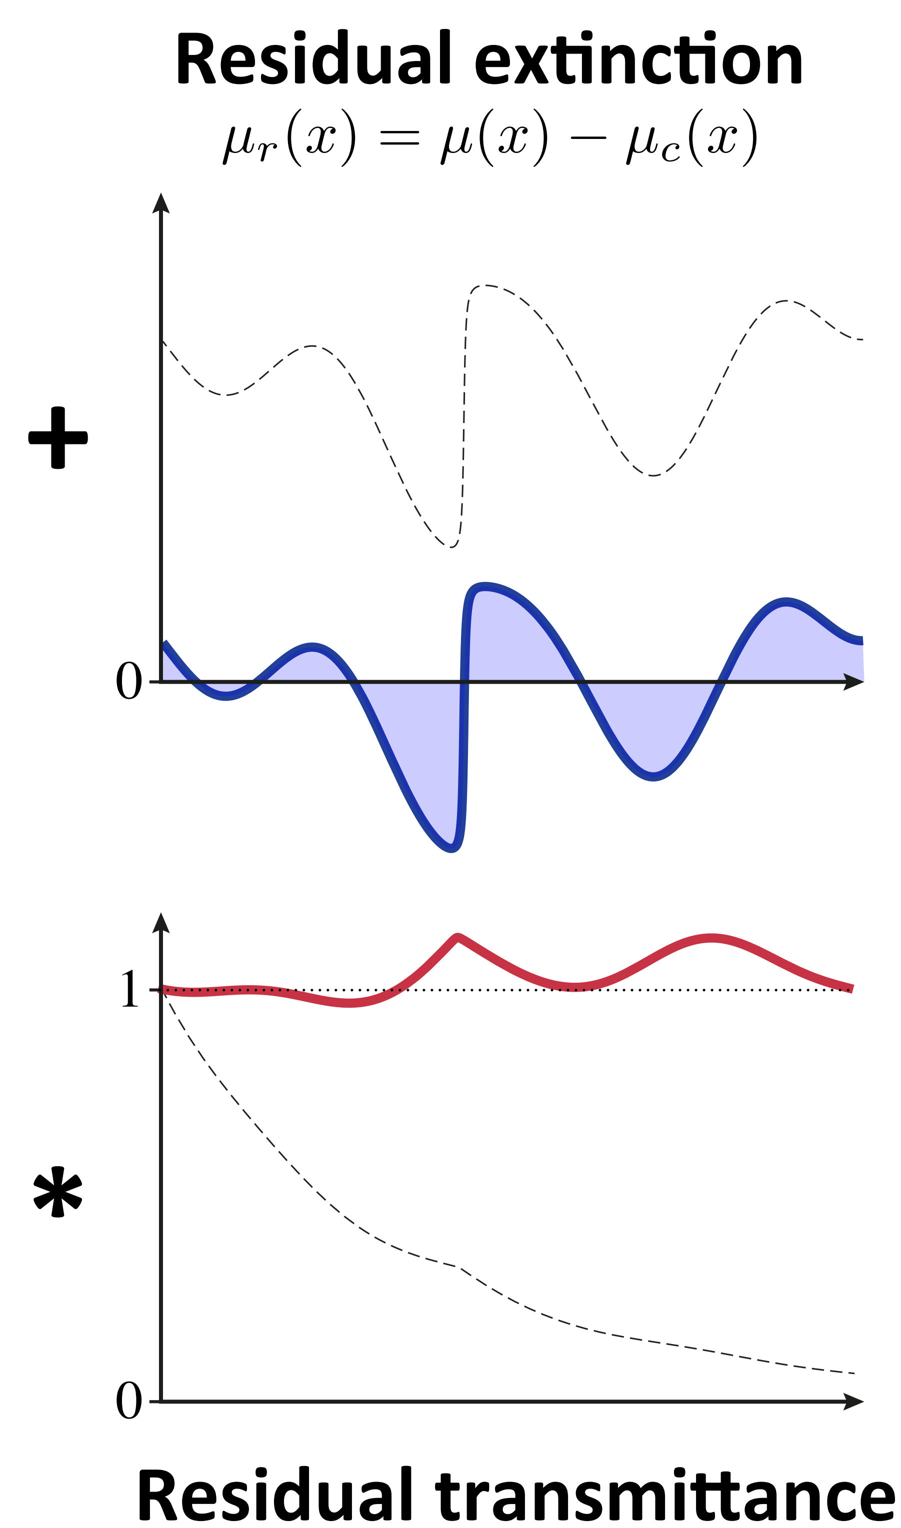

It is worth noting that the residual extinction can be negative. When this happens, it means that the transport should be amplified, instead of being attenuated.

This amplification cannot be handled by all algorithms, for instance we cannot use the Delta tracking as the amplification requires an additional weight.

But we can use the previously mentioned Ratio tracking, which handles negative extinctions natively without any problem.

Residual Tracking

Residual extinction

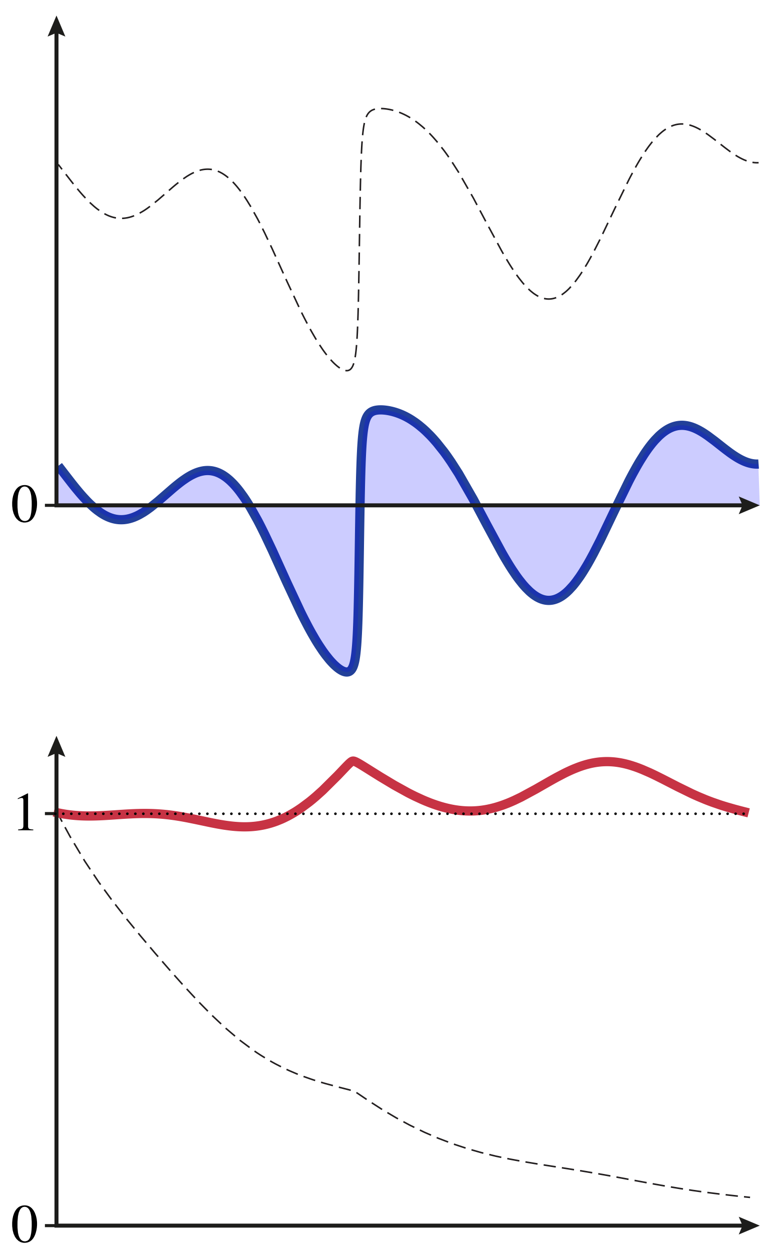

Residual transmittance

It is worth noting that the residual extinction can be negative. When this happens, it means that the transport should be amplified, instead of being attenuated.

This amplification cannot be handled by all algorithms, for instance we cannot use the Delta tracking as the amplification requires an additional weight.

But we can use the previously mentioned Ratio tracking, which handles negative extinctions natively without any problem.

Residual Tracking

Residual extinction

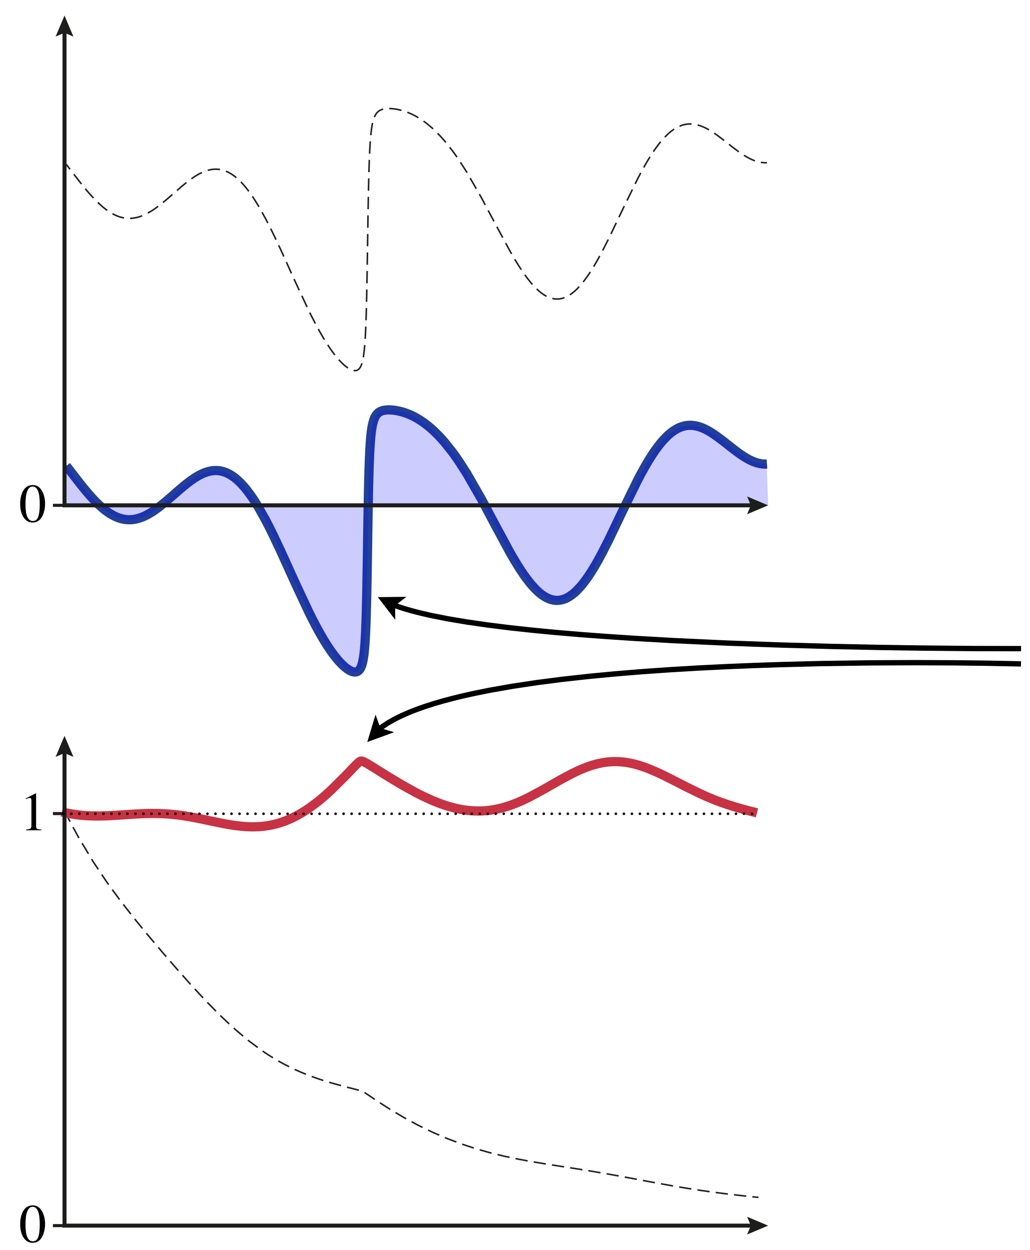

Residual transmittance

Transport "amplification"

- requires an additional weight
- cannot use **Delta tracking**
- **Ratio tracking** works well

It is worth noting that the residual extinction can be negative. When this happens, it means that the transport should be amplified, instead of being attenuated.

This amplification cannot be handled by all algorithms, for instance we cannot use the Delta tracking as the amplification requires an additional weight.

But we can use the previously mentioned Ratio tracking, which handles negative extinctions natively without any problem.

Residual Ratio Tracking

Example:

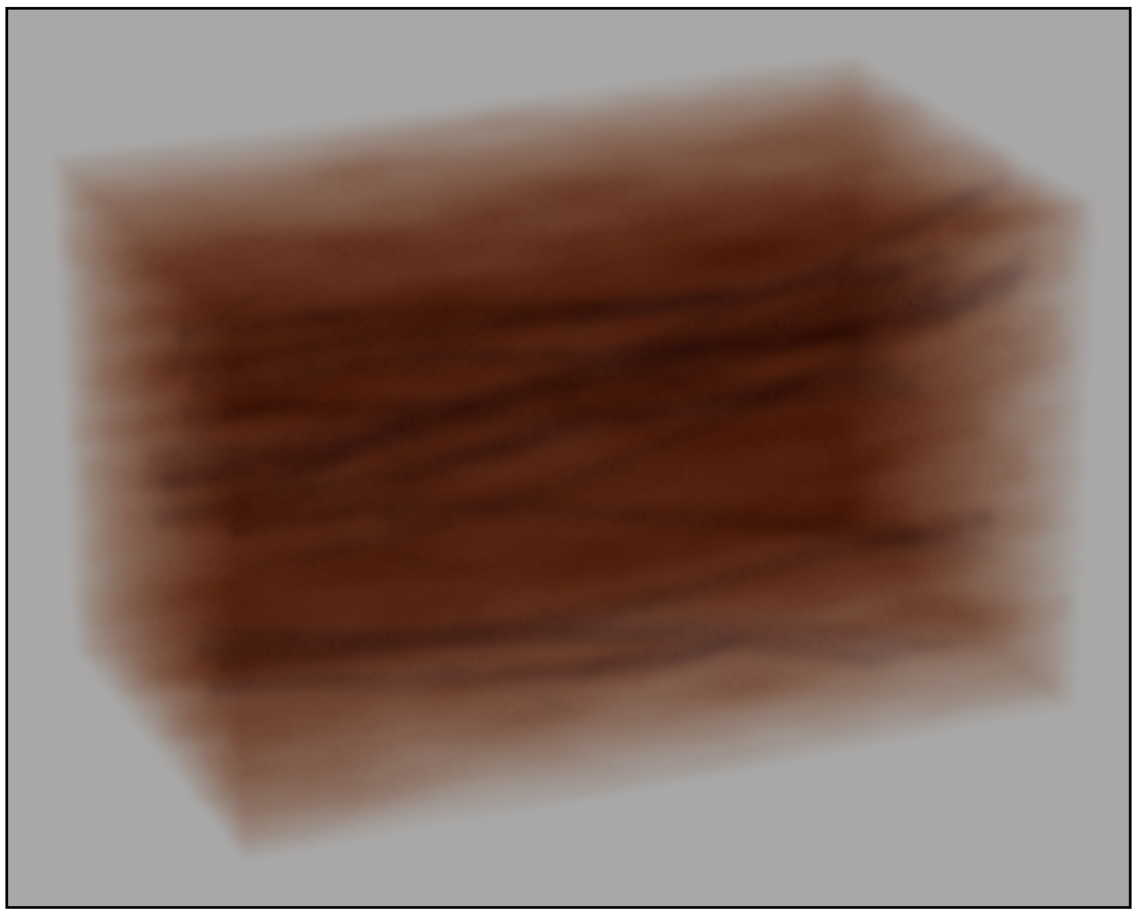

Simple, moderately heterogeneous volume

Let’s look at an example rendered with the residual ratio tracking, here we have a volume with a relatively low degree of heterogeneity.

Residual Ratio Tracking

SIGGRAPH ASIA 2014 SHENZHEN

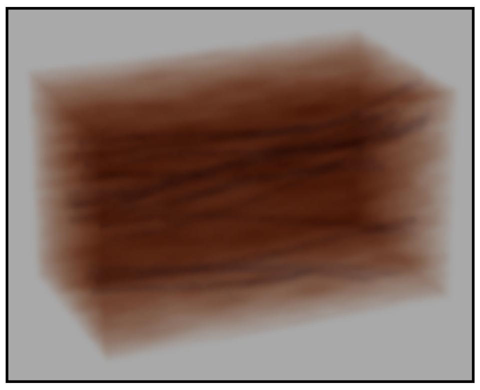

† over the entire volume

In the top row, we use the minimum extinction in the volume as the control extinction. In the middle we use the mean extinction as the control, and at the bottom we have the maximum as the control.

Note how the residual transmittance... in the middle... always corrects the control transmittance. and in all cases, the resulting product is an unbiased estimate of the transmittance.

But it seems that the average extinction works better than the other values in this simple example. However, using a single control for the entire volume leads a very non-uniform noise.

Residual Ratio Tracking

SIGGRAPH ASIA 2014 SHENZHEN

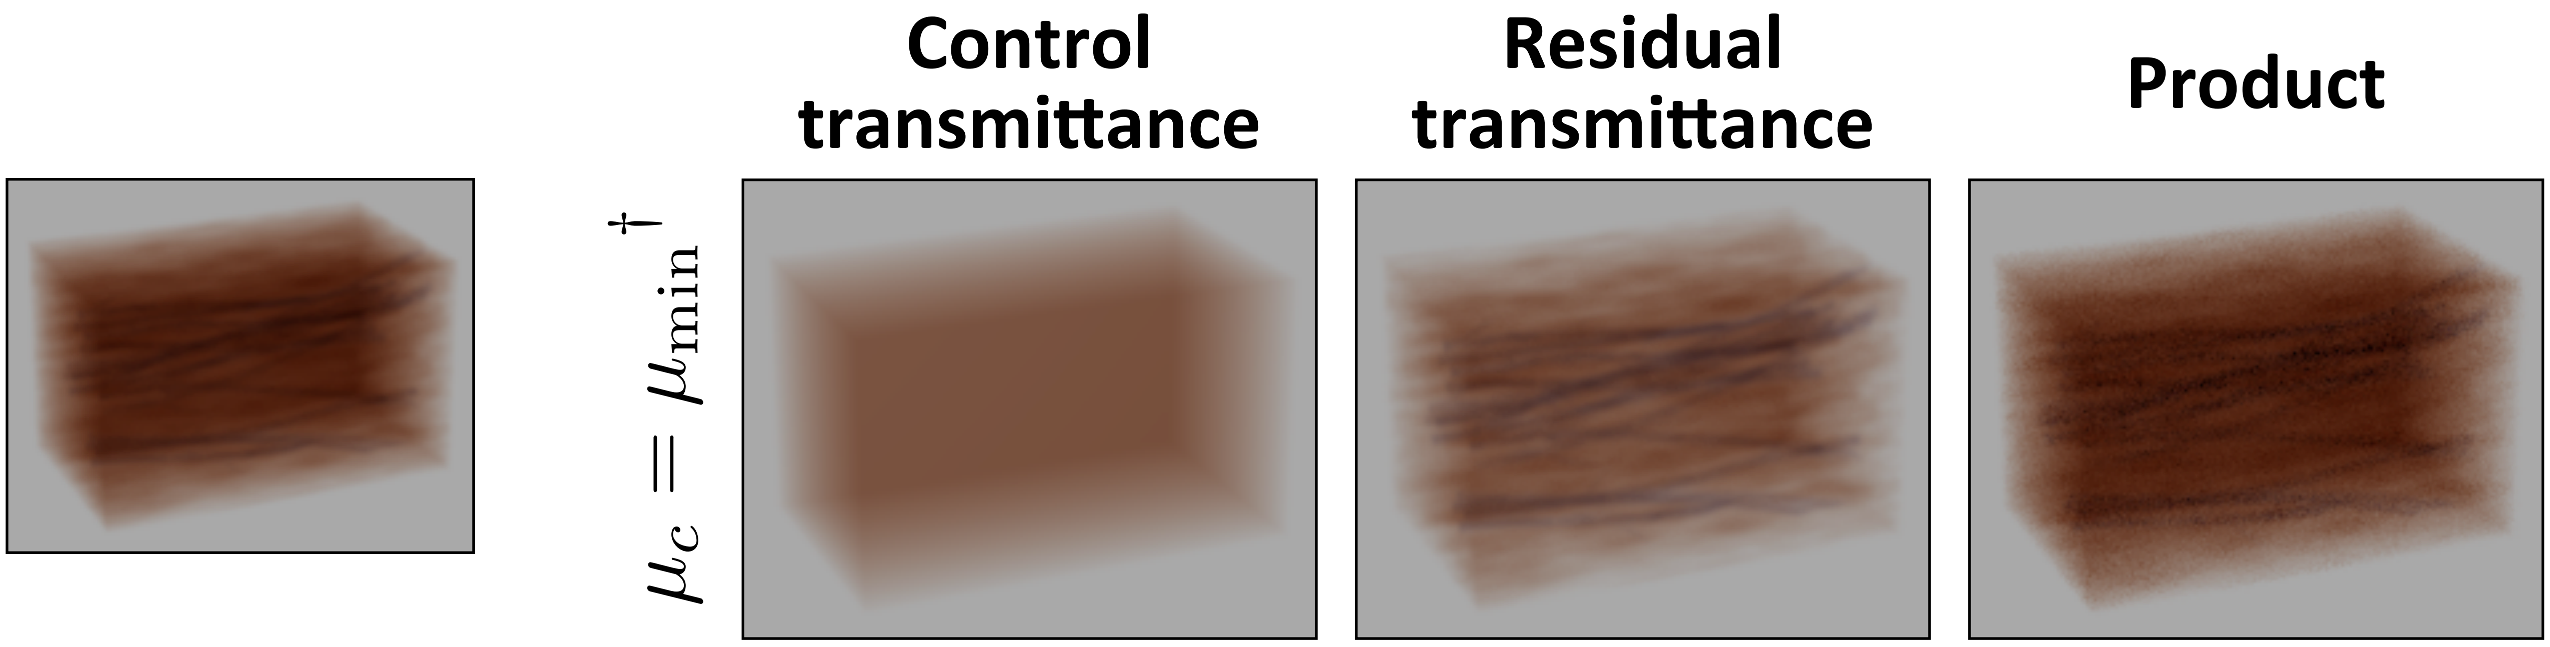

† over the entire volume

In the top row, we use the minimum extinction in the volume as the control extinction. In the middle we use the mean extinction as the control, and at the bottom we have the maximum as the control.

Note how the residual transmittance... in the middle... always corrects the control transmittance. and in all cases, the resulting product is an unbiased estimate of the transmittance.

But it seems that the average extinction works better than the other values in this simple example. However, using a single control for the entire volume leads a very non-uniform noise.

Residual Ratio Tracking

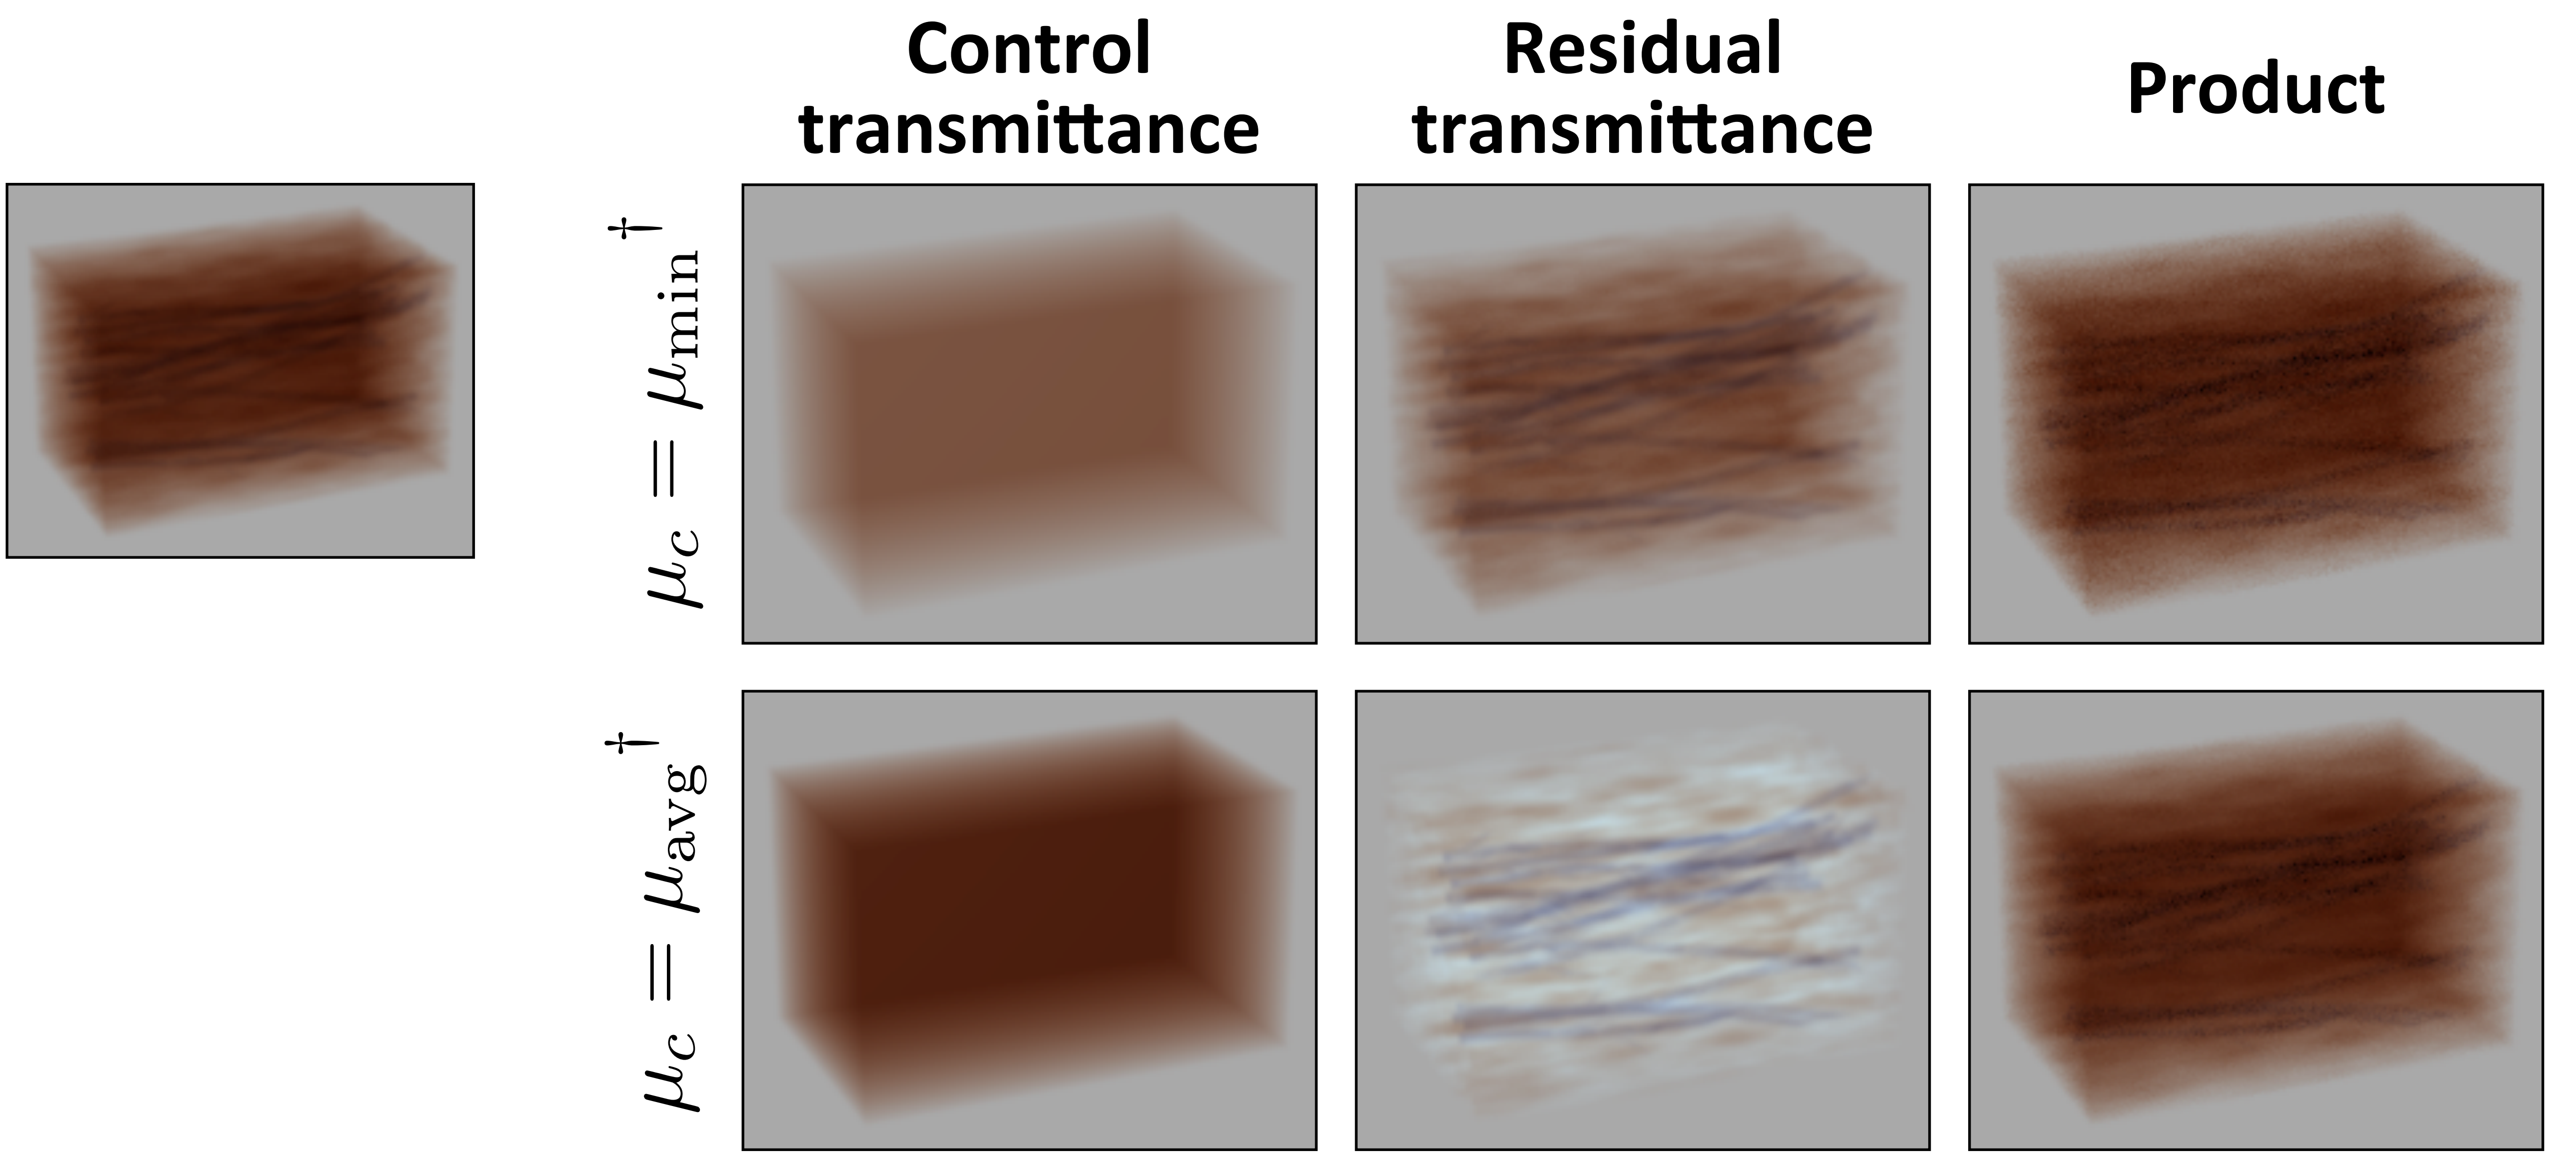

† over the entire volume

In the top row, we use the minimum extinction in the volume as the control extinction. In the middle we use the mean extinction as the control, and at the bottom we have the maximum as the control.

Note how the residual transmittance... in the middle... always corrects the control transmittance. and in all cases, the resulting product is an unbiased estimate of the transmittance.

But it seems that the average extinction works better than the other values in this simple example. However, using a single control for the entire volume leads a very non-uniform noise.

Residual Ratio Tracking

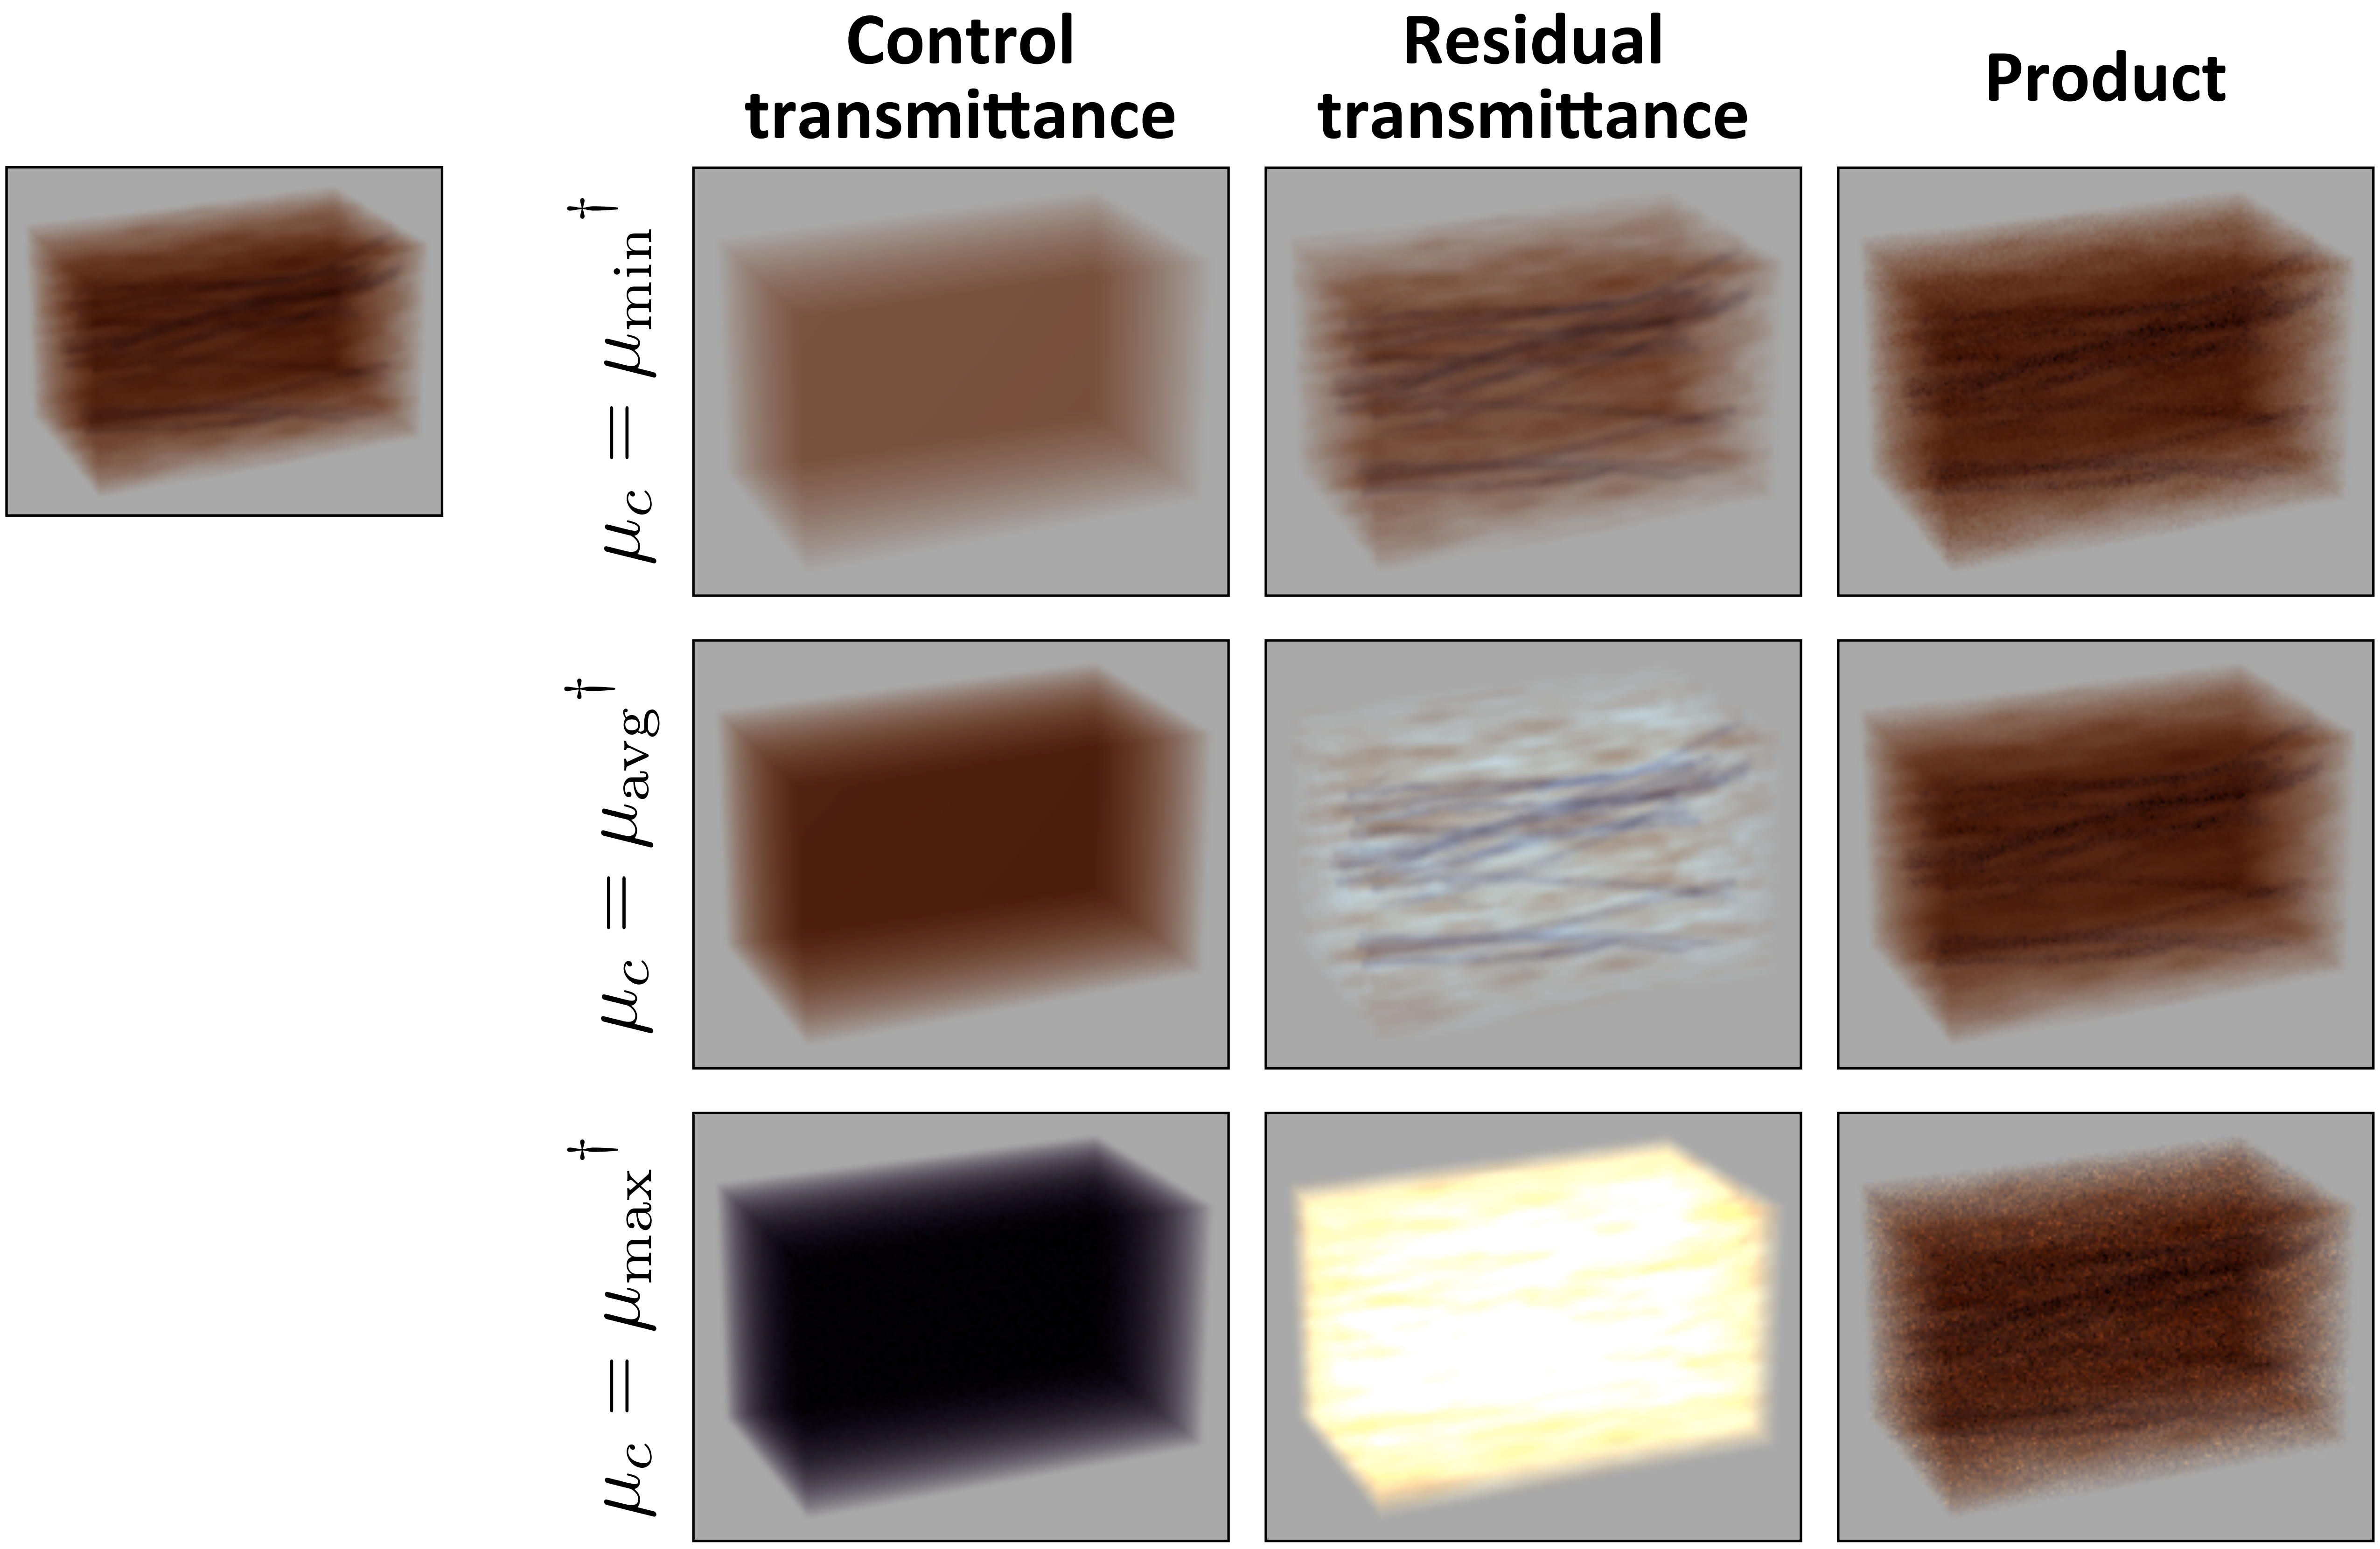

† over the entire volume

In the top row, we use the minimum extinction in the volume as the control extinction. In the middle we use the mean extinction as the control, and at the bottom we have the maximum as the control.

Note how the residual transmittance... in the middle... always corrects the control transmittance. and in all cases, the resulting product is an unbiased estimate of the transmittance.

But it seems that the average extinction works better than the other values in this simple example. However, using a single control for the entire volume leads a very non-uniform noise.

Localized Control Extinction

Optimally, we would want the control to be somewhat localized, so that it better matches the extinction function. There has been some research on localizing majorant coefficients for Delta tracking. We use the approach by Szirmay-Kalos and construct a grid of super-voxel, each storing a local control extinction function. We also experimented with constant and linearly interpolated control extinctions, I'll show an example in a moment.

Localized Control Extinction

- Localized majorants for Delta tracking
  - kd-tree [Yue et al. 2010]
  - super-voxels [Szirmay-Kalos et al. 2011]

Optimally, we would want the control to be somewhat localized, so that it better matches the extinction function. There has been some research on localizing majorant coefficients for Delta tracking. We use the approach by Szirmay-Kalos and construct a grid of super-voxel, each storing a local control extinction function. We also experimented with constant and linearly interpolated control extinctions, I'll show an example in a moment.

Localized Control Extinction

- Localized majorants for Delta tracking
  - kd-tree [Yue et al. 2010]
  - super-voxels [Szirmay-Kalos et al. 2011]
- Super-voxels
  - each voxel stores local control extinction (constant or interpolated)

Optimally, we would want the control to be somewhat localized, so that it better matches the extinction function. There has been some research on localizing majorant coefficients for Delta tracking. We use the approach by Szirmay-Kalos and construct a grid of super-voxel, each storing a local control extinction function. We also experimented with constant and linearly interpolated control extinctions, I'll show an example in a moment.

Localized Control Extinction

- Localized majorants for Delta tracking
  - kd-tree [Yue et al. 2010]
  - super-voxels [Szirmay-Kalos et al. 2011]
- Super-voxels
  - each voxel stores local control extinction (constant or interpolated)

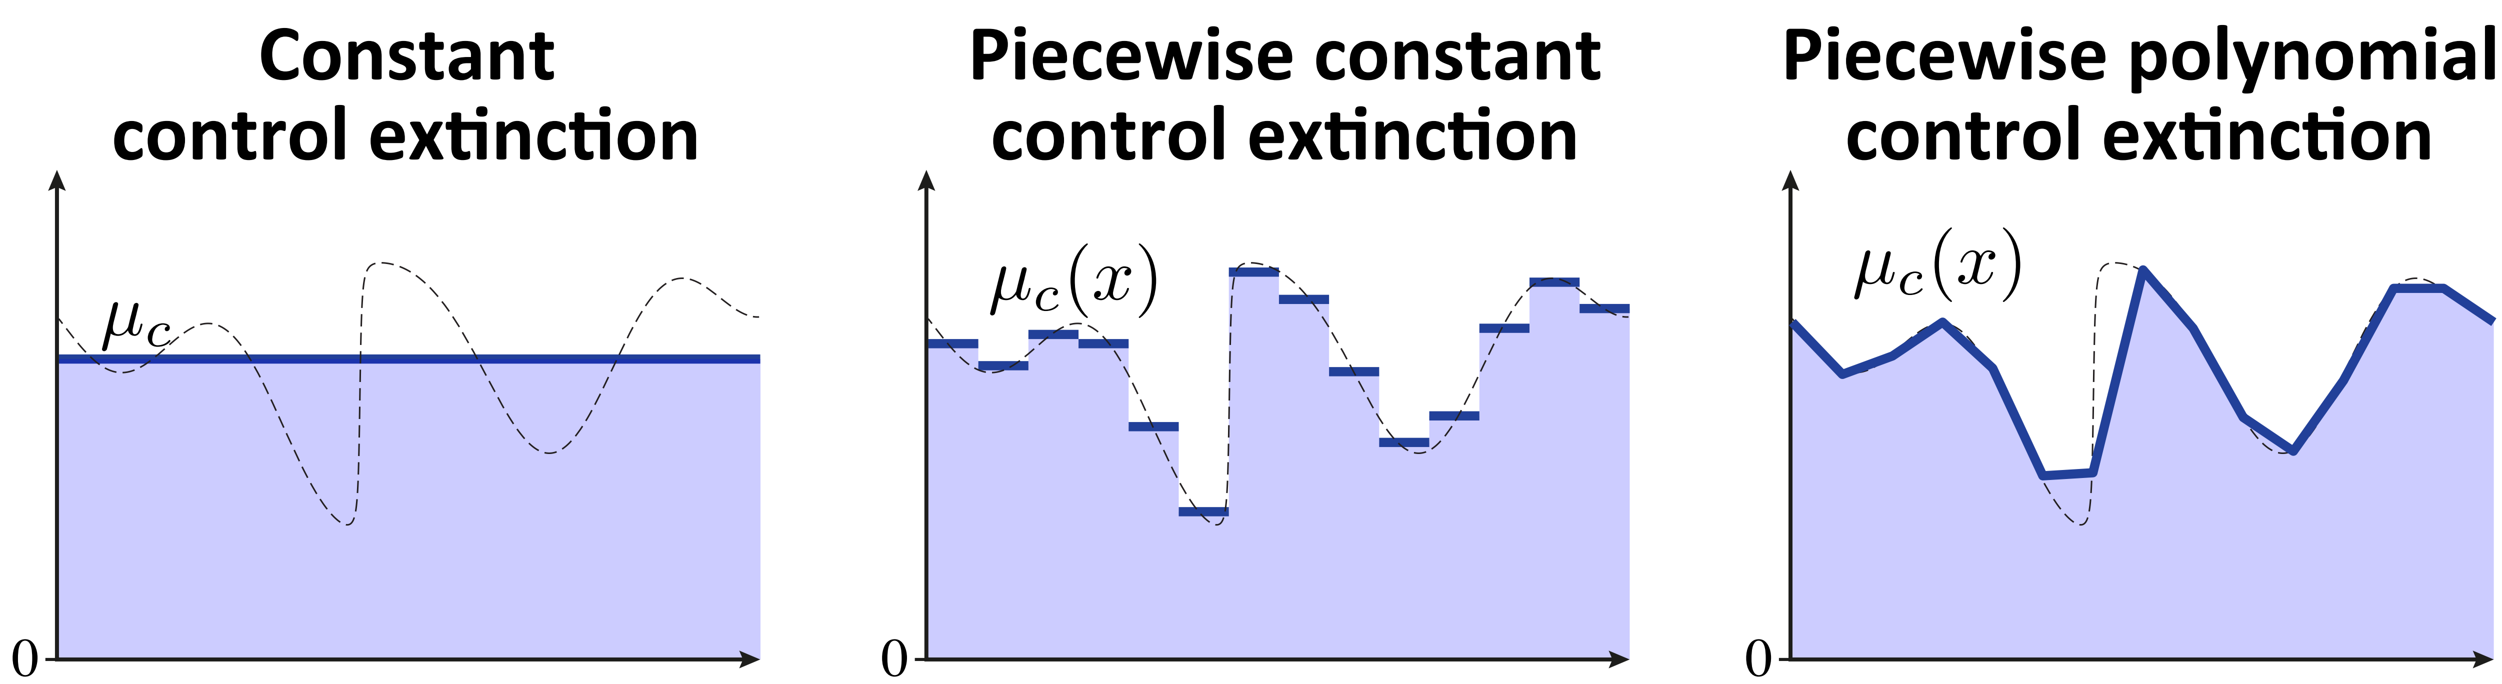

Optimally, we would want the control to be somewhat localized, so that it better matches the extinction function. There has been some research on localizing majorant coefficients for Delta tracking. We use the approach by Szirmay-Kalos and construct a grid of super-voxel, each storing a local control extinction function. We also experimented with constant and linearly interpolated control extinctions, I'll show an example in a moment.

Residual Ratio Tracking with Super-Voxels

Let’s look at some results, this is an absorbing heterogeneous medium.

Residual Ratio Tracking with Super-Voxels

Control	Residuum	Product	Variance

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

Residual Ratio Tracking with Super-Voxels

Control | Residuum | Product | Variance

$\mu_c = \mu_{\min}$

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

Residual Ratio Tracking with Super-Voxels

Control	Residuum	Product	Variance

$\mu_c = \mu_{\min}$

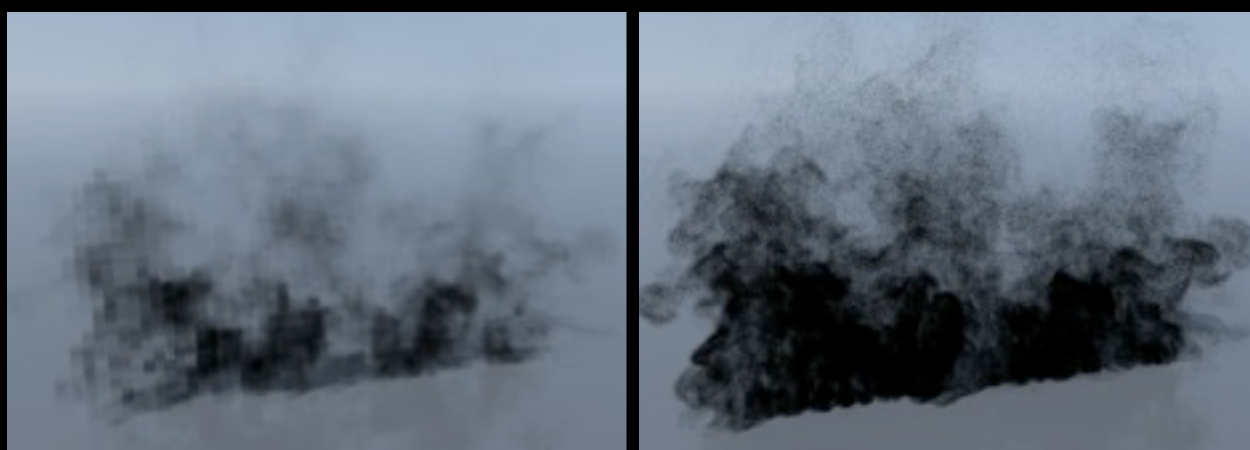

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

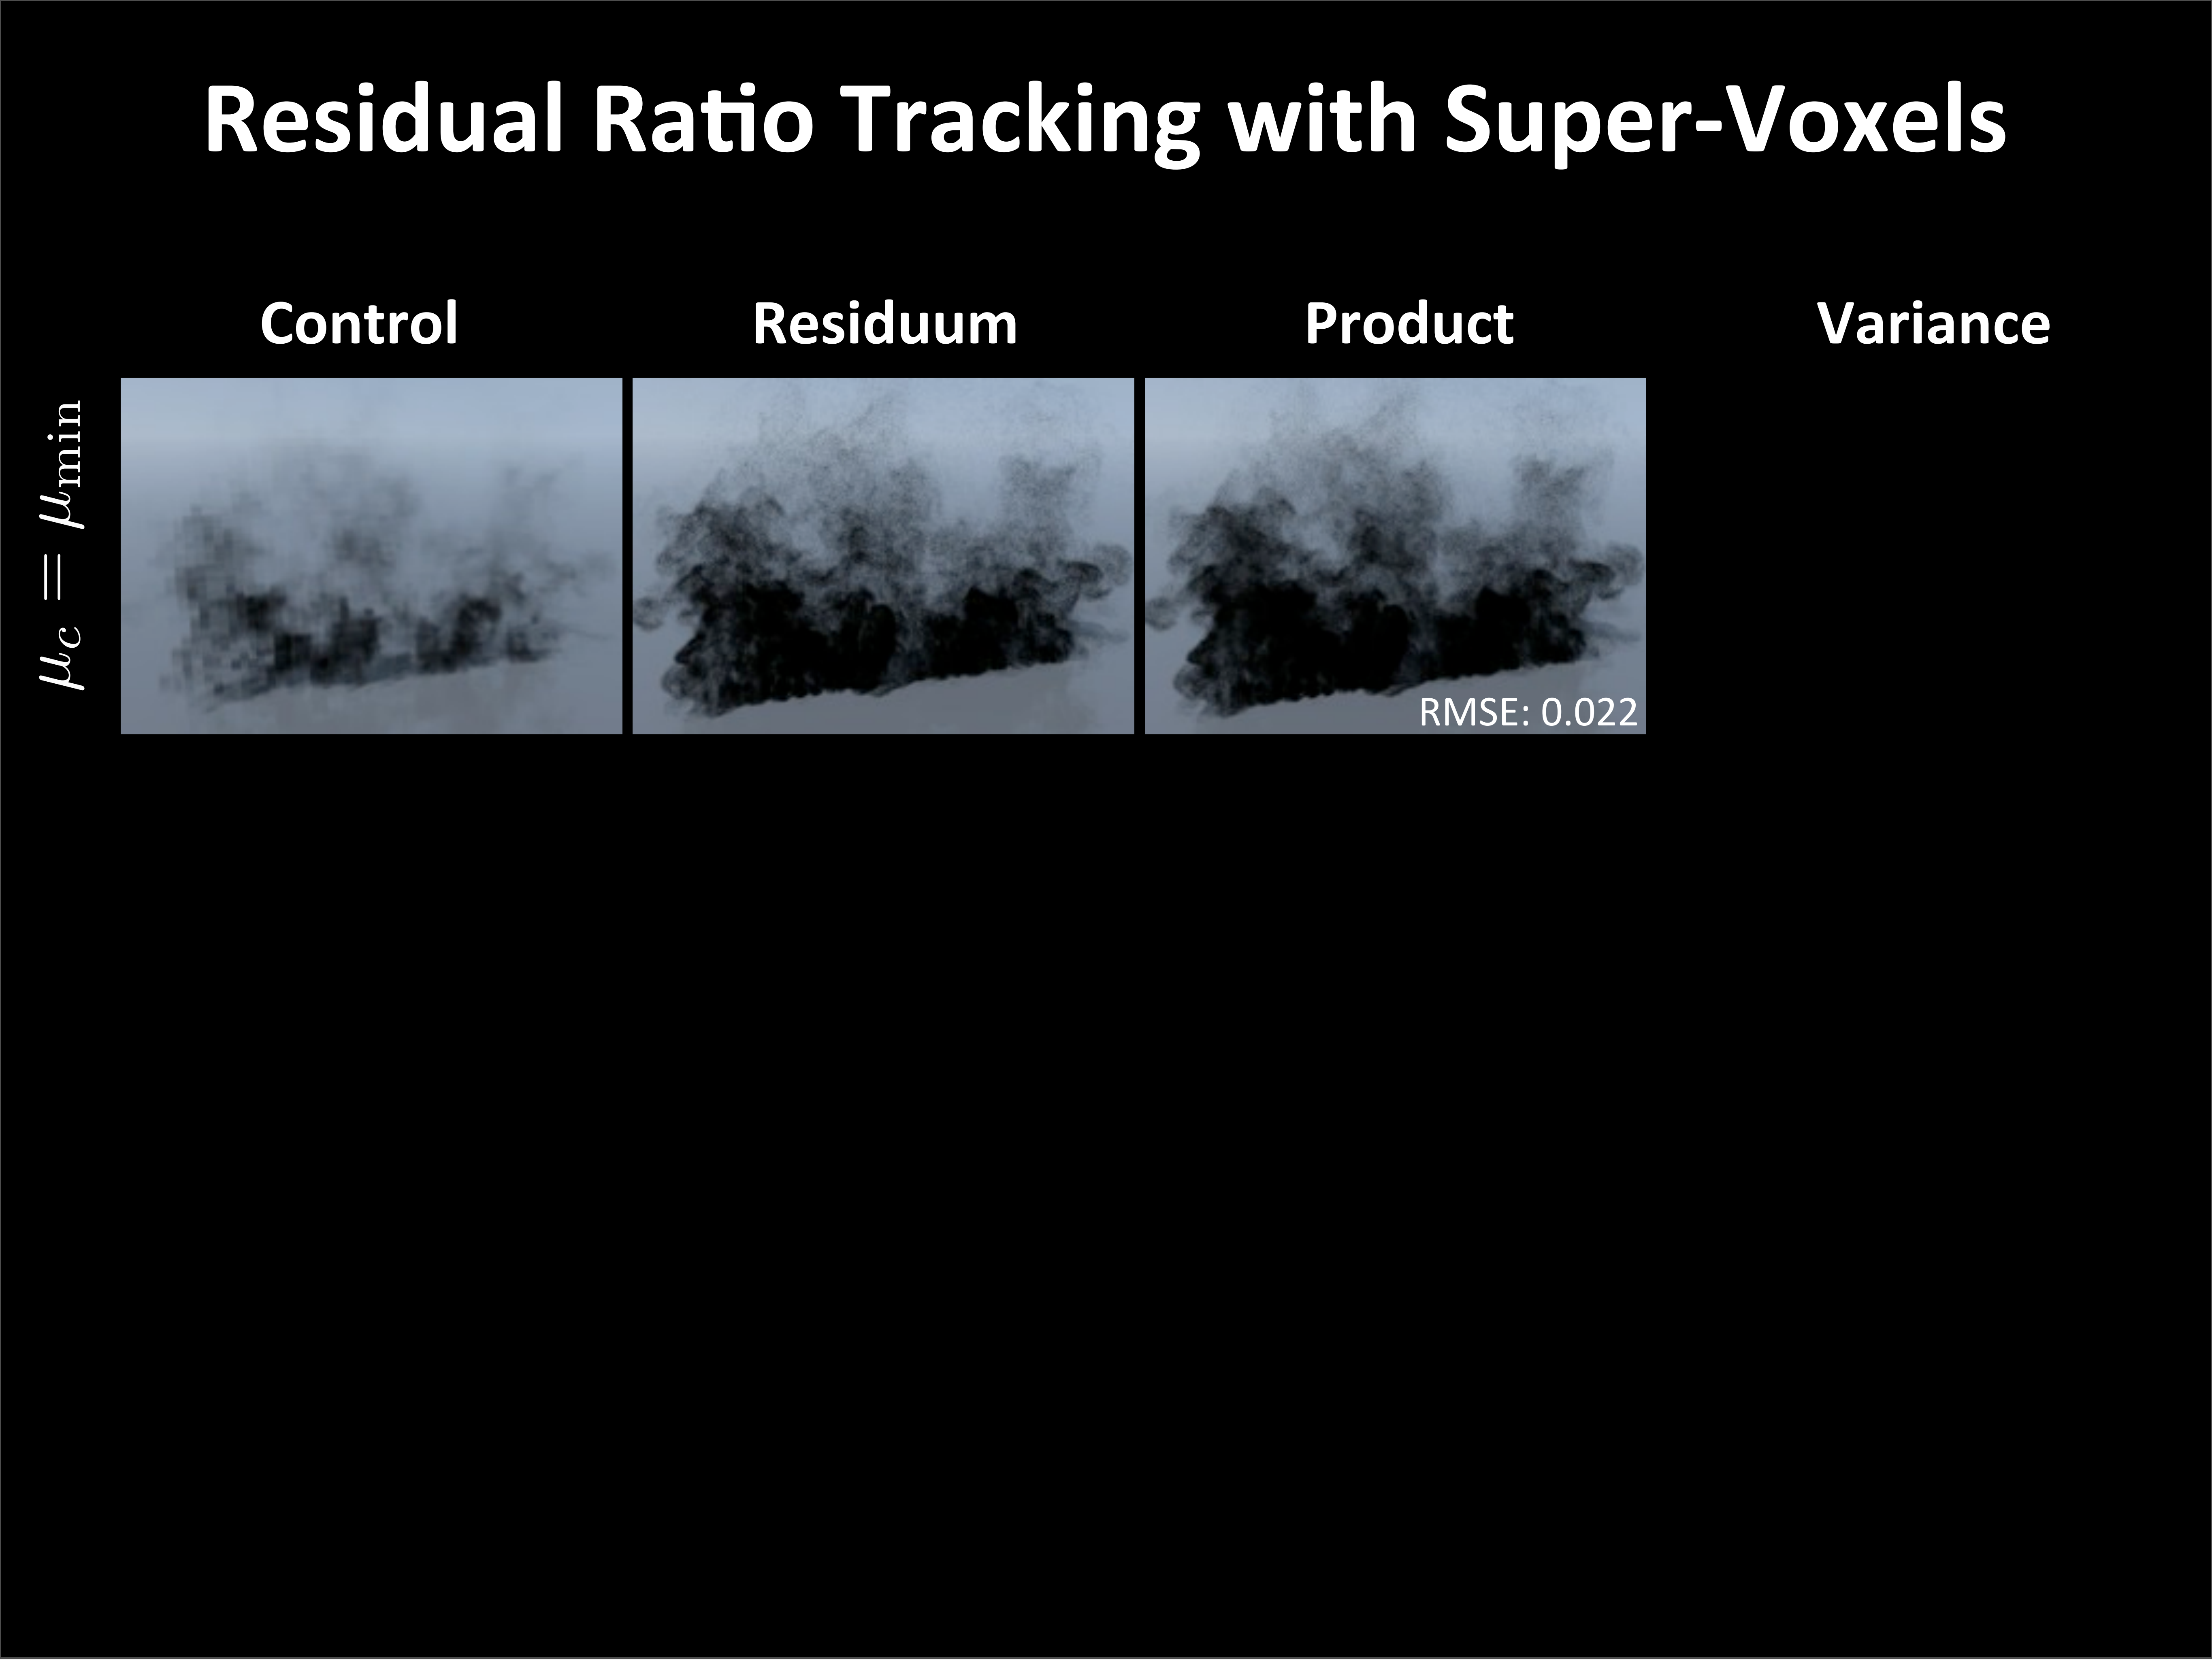

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

Residual Ratio Tracking with Super-Voxels

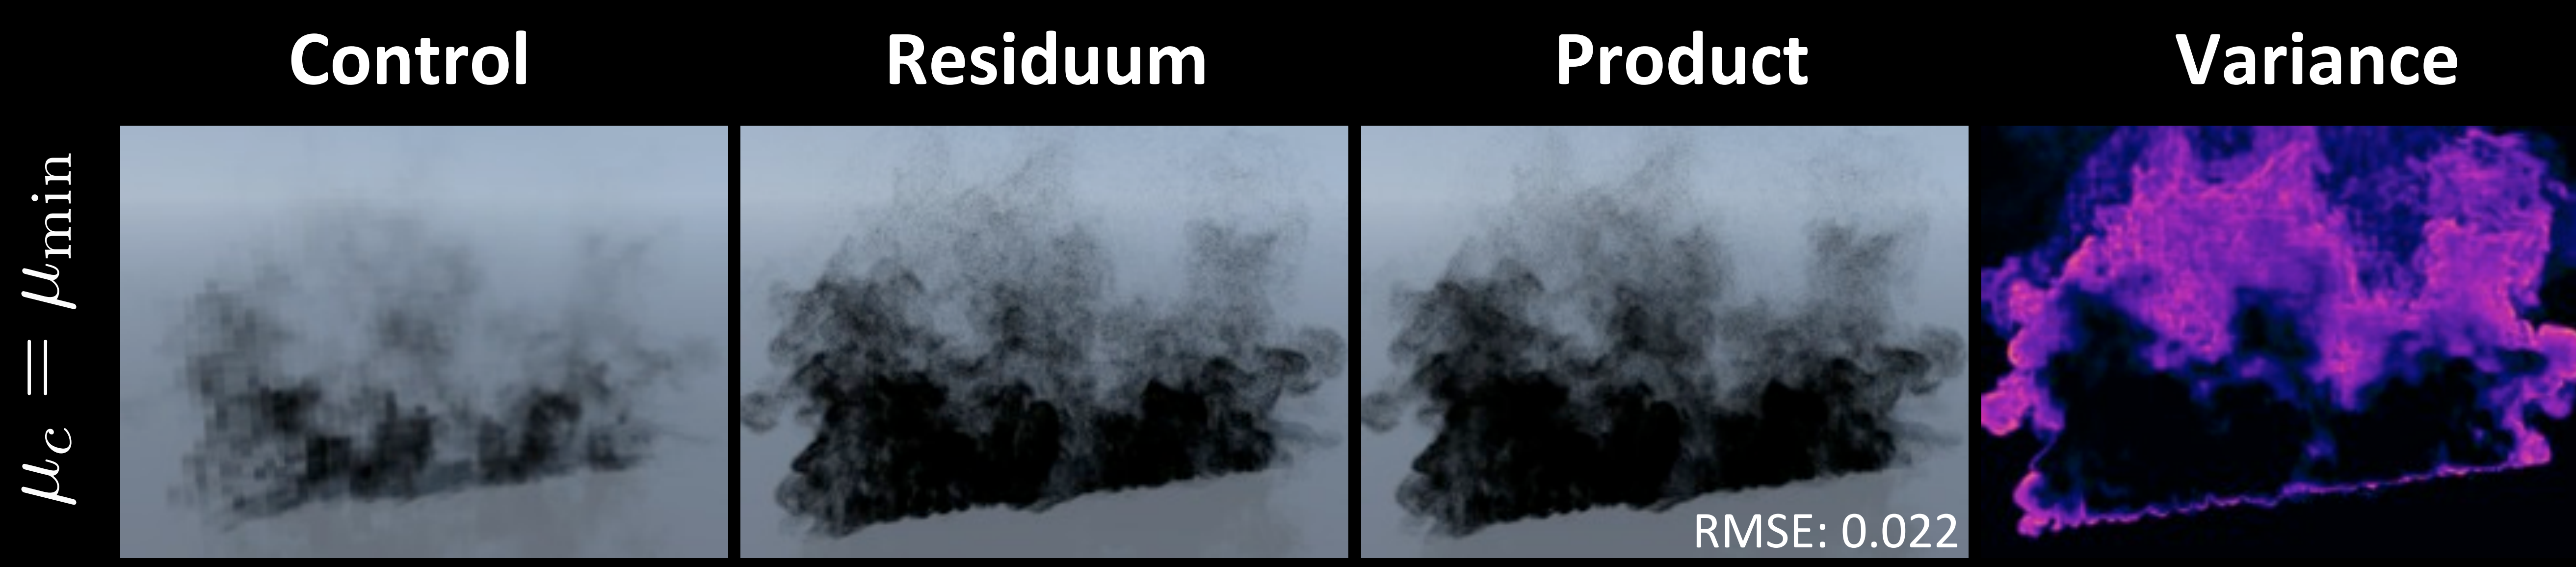

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

Residual Ratio Tracking with Super-Voxels

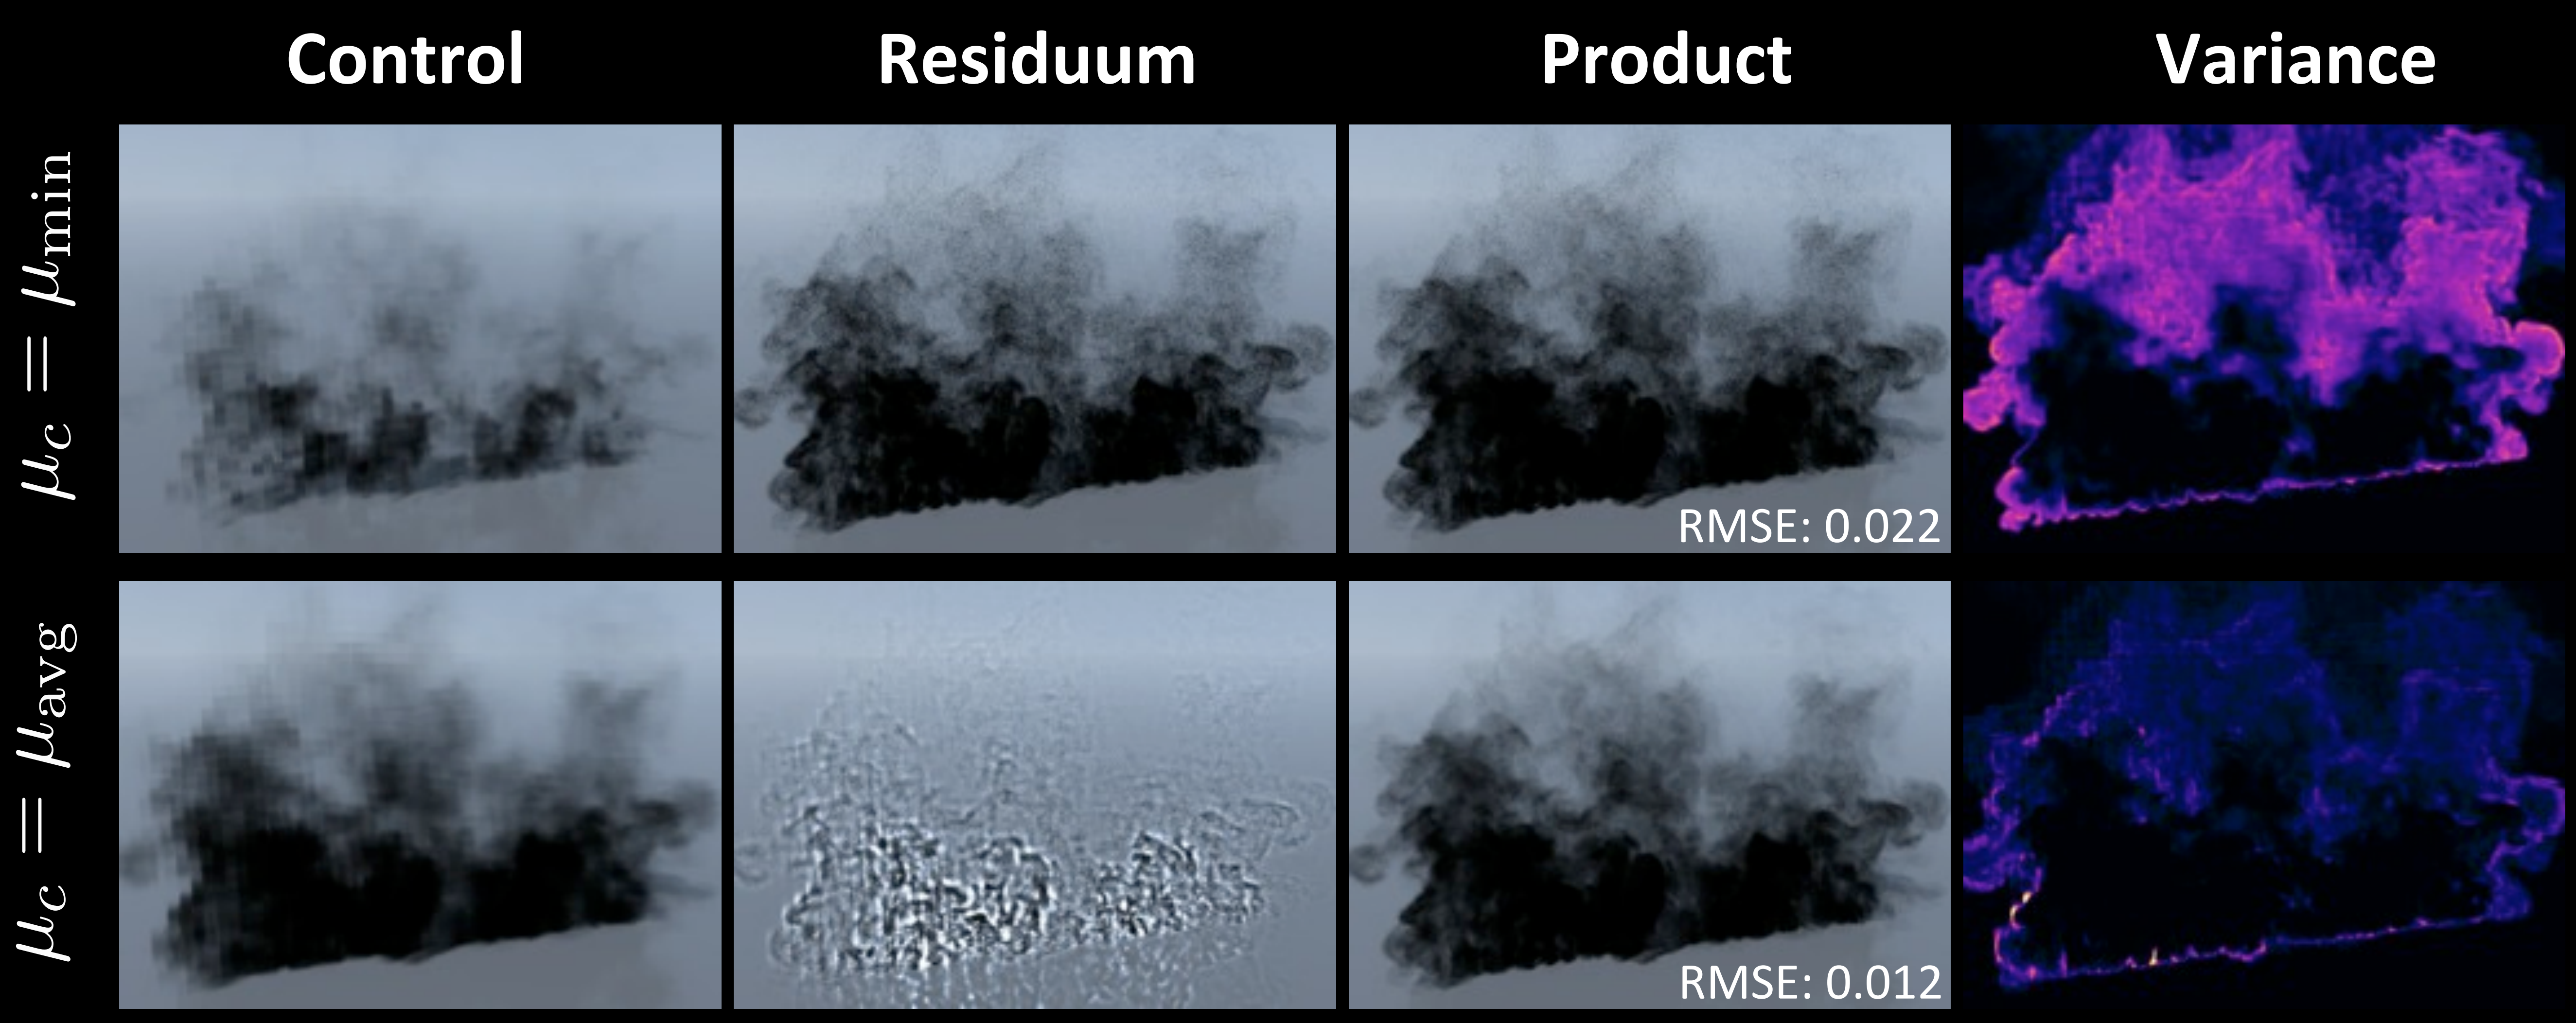

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

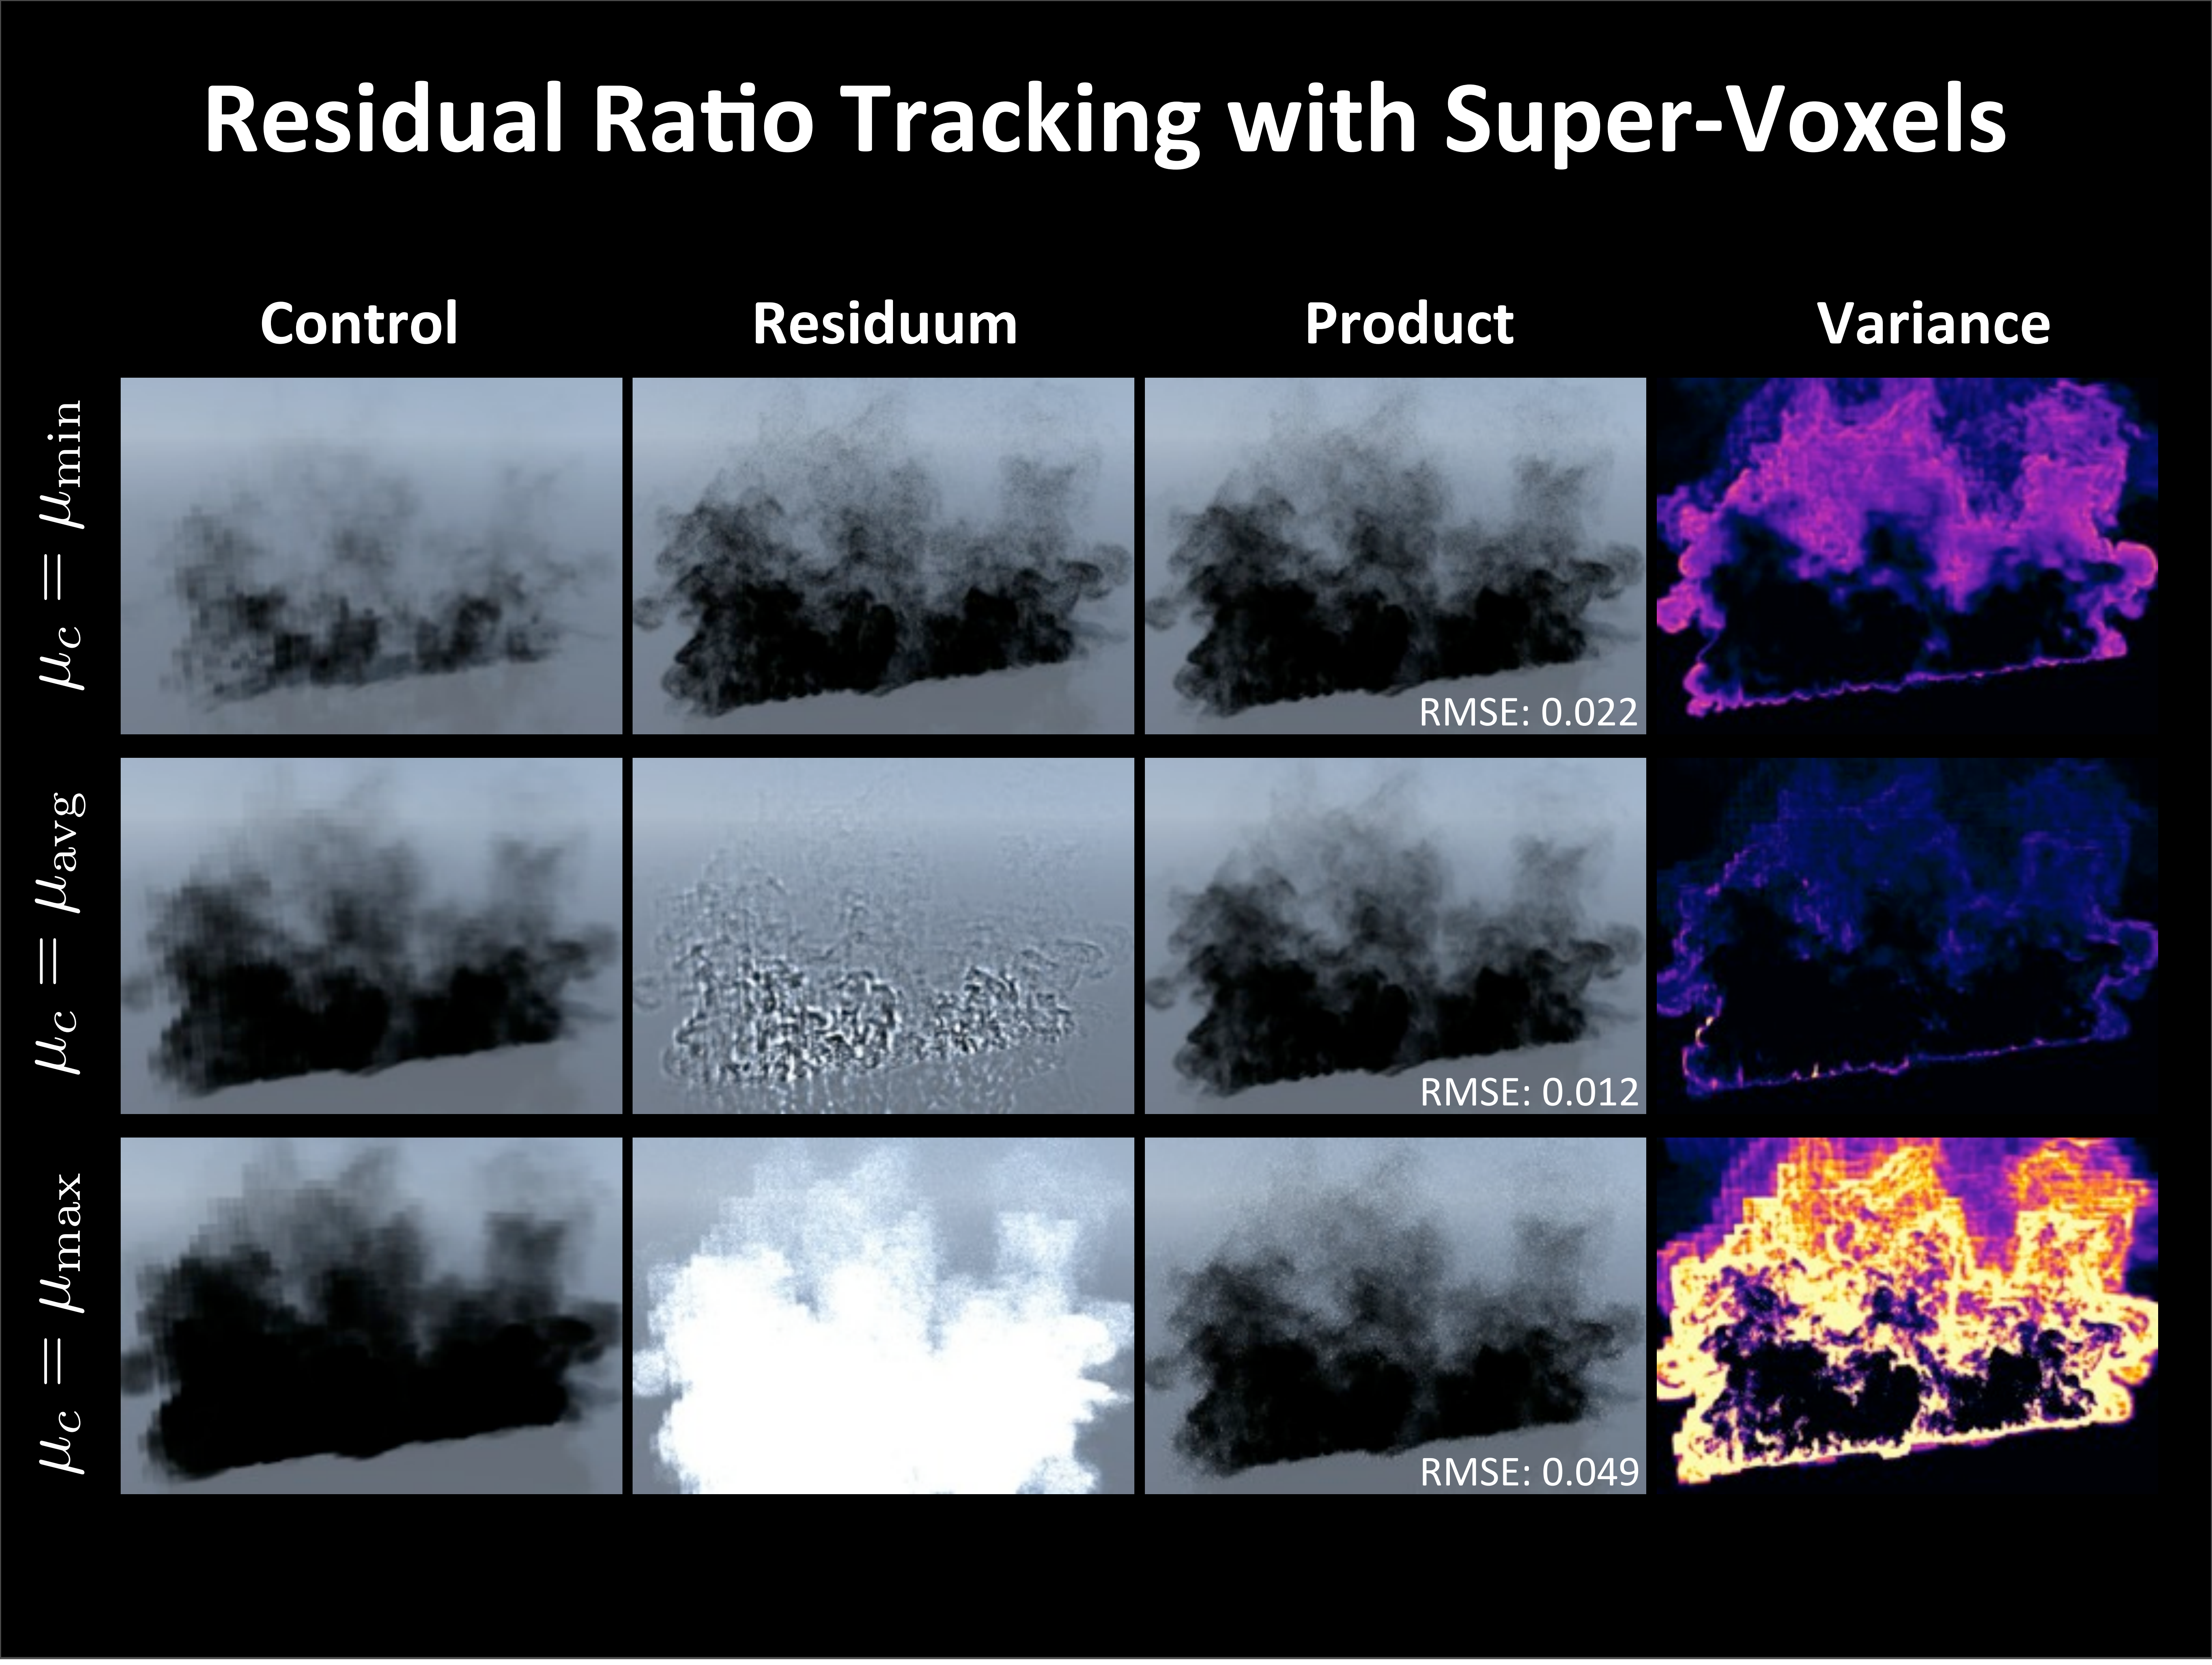

In the top row, we have the control medium that uses the super-voxels. In this case, each super-voxel stores the minimum extinction value, which is used as the control. This is the residual medium. And here is a product of the control and residual transmittance. The right-most false-color rendering shows the variance of the estimator.

In the second row, the super-voxels store the average extinction. In the last row, we used the maximum extinction.

You can see that the lowest overall variance is obtained with the average extinction.
In the paper, we propose an additional heuristics to make the control more robust.

Residual Ratio Tracking with Super-Voxels

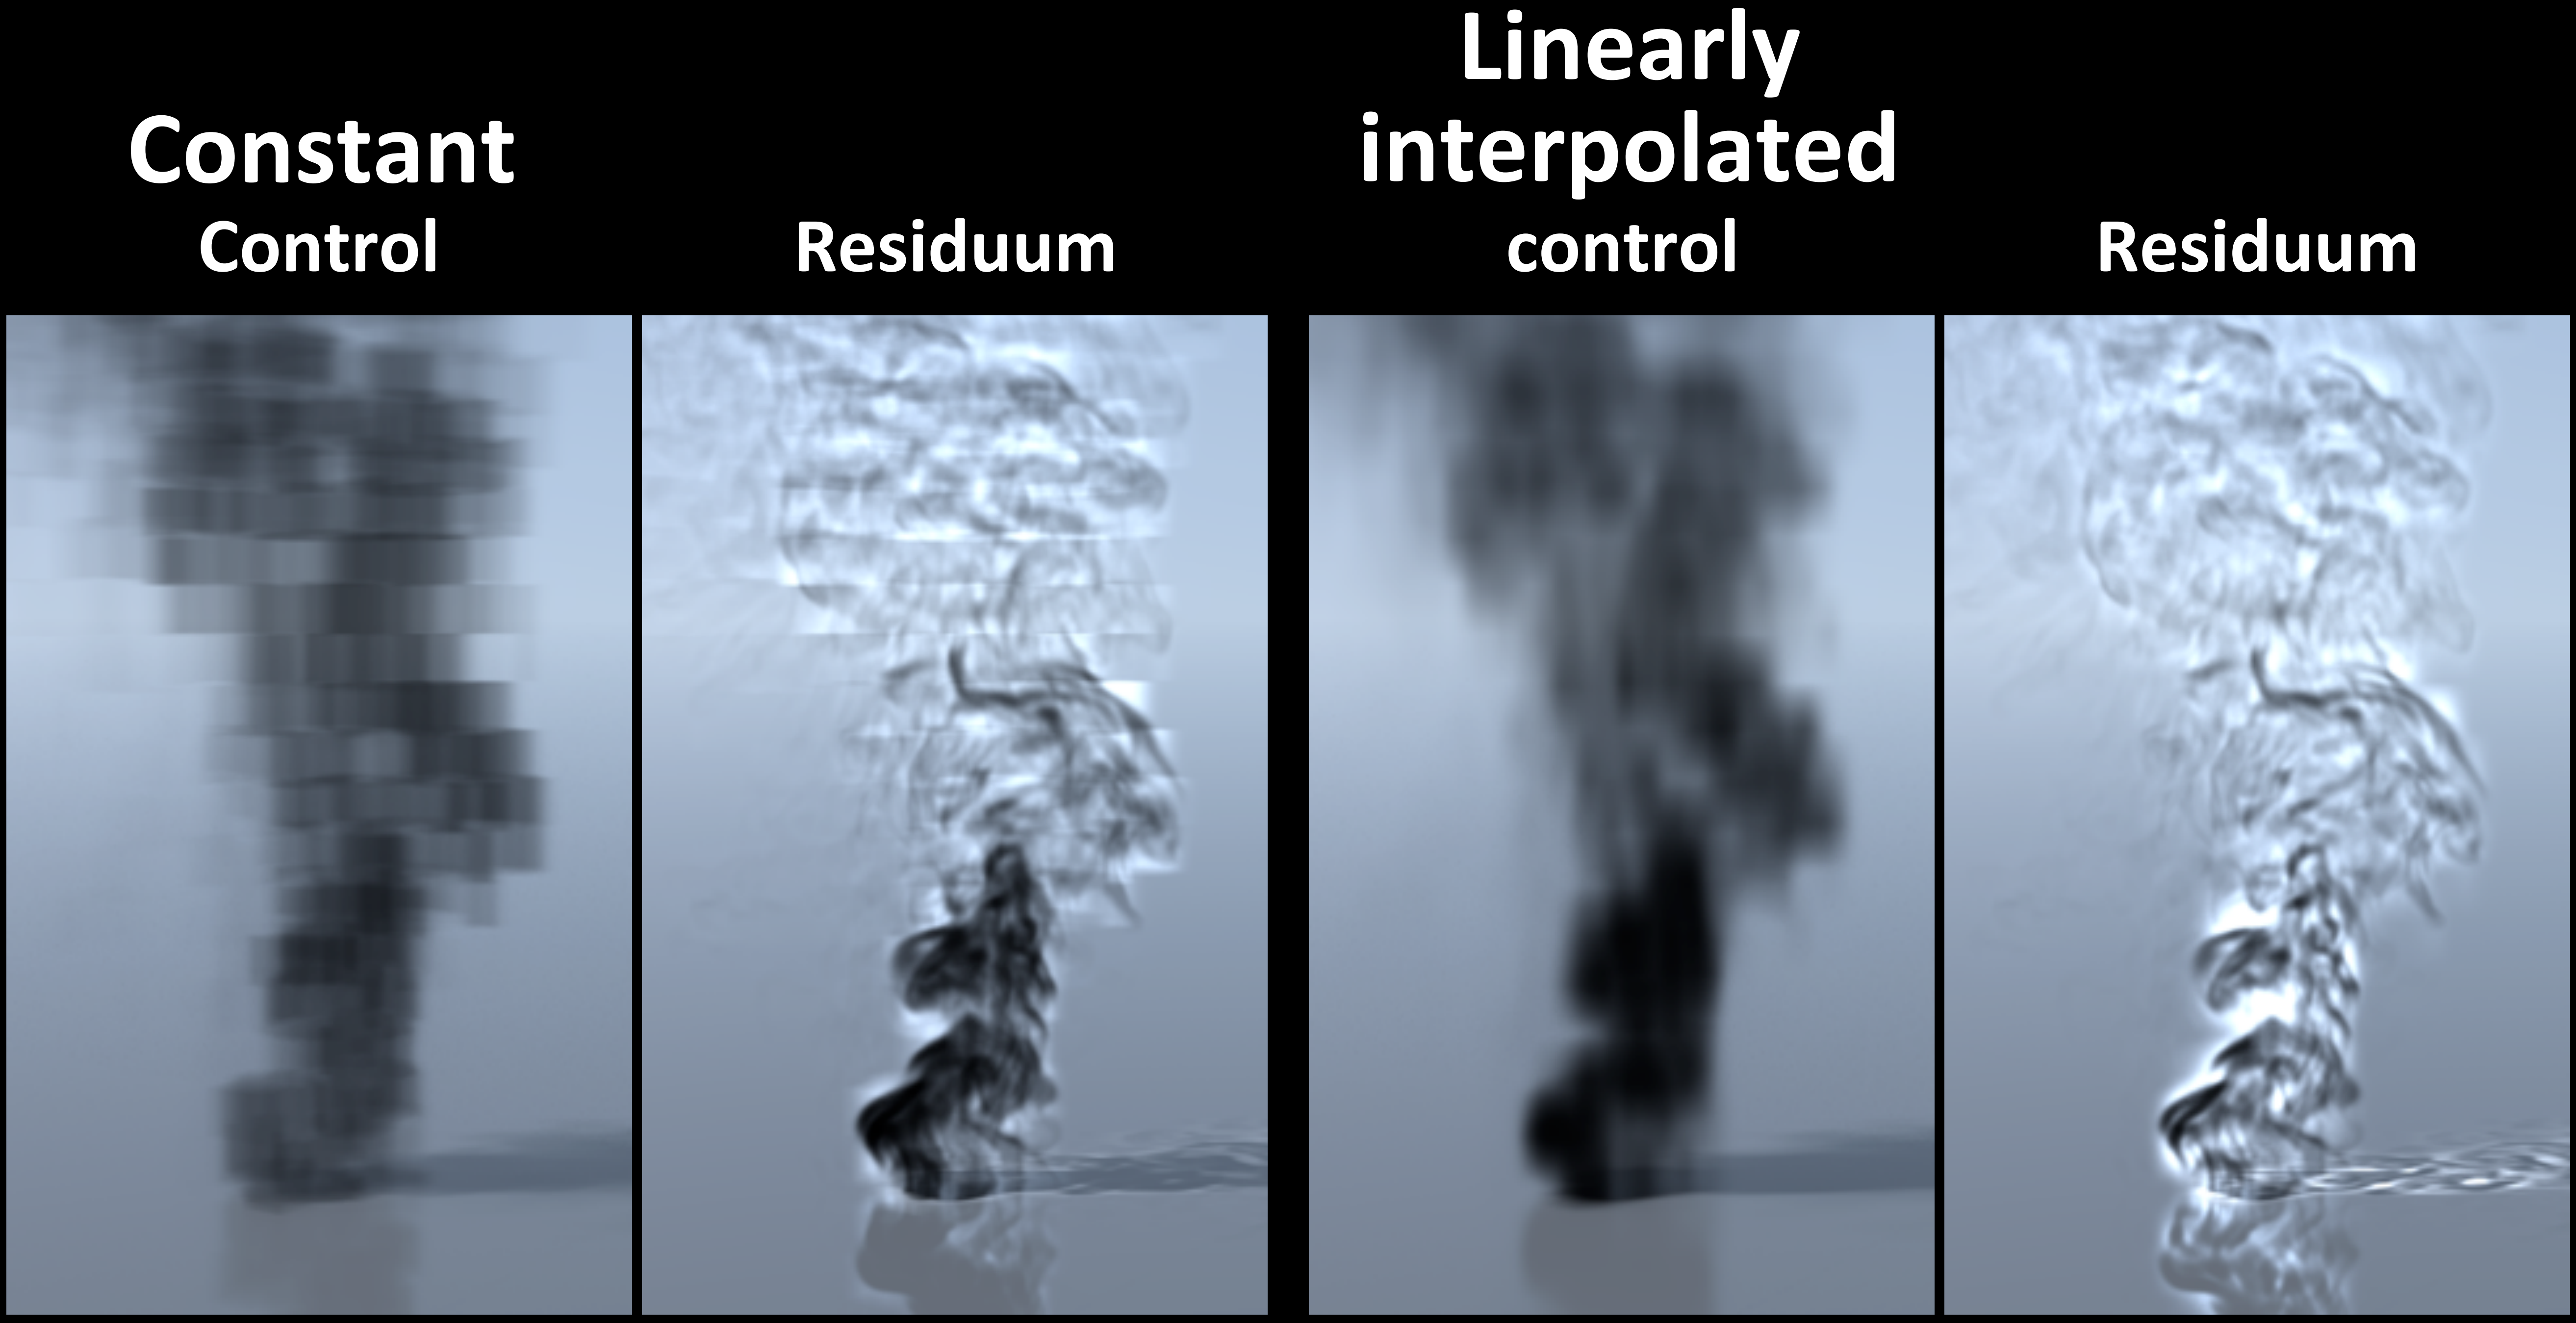

It is possible to also linearly interpolate the control extinction in each super-voxel. This makes the control smoother and it can further reduce the noise. It may also serve well for level-of-detail rendering, where we simply omit the residual tracking and use just the control transmittance.

Colored Extinction

Clouds

In this example, we have clouds with colored extinction. Let’s first look at the transmittance along the primary rays. All images were rendered at equal cost.

Delta tracking produces a lot of color noise, as it needs to handle the transmittance through each color channel independently.... in order to be efficient.

Ratio tracking can efficiently handle all color channels at once, removing the color noise and if we combine it with residual tracking, we get additional significant reduction in variance.

Colored Extinction

Transmittance along primary rays

Clouds

In this example, we have clouds with colored extinction. Let’s first look at the transmittance along the primary rays. All images were rendered at equal cost.

Delta tracking produces a lot of color noise, as it needs to handle the transmittance through each color channel independently.... in order to be efficient.

Ratio tracking can efficiently handle all color channels at once, removing the color noise and if we combine it with residual tracking, we get additional significant reduction in variance.

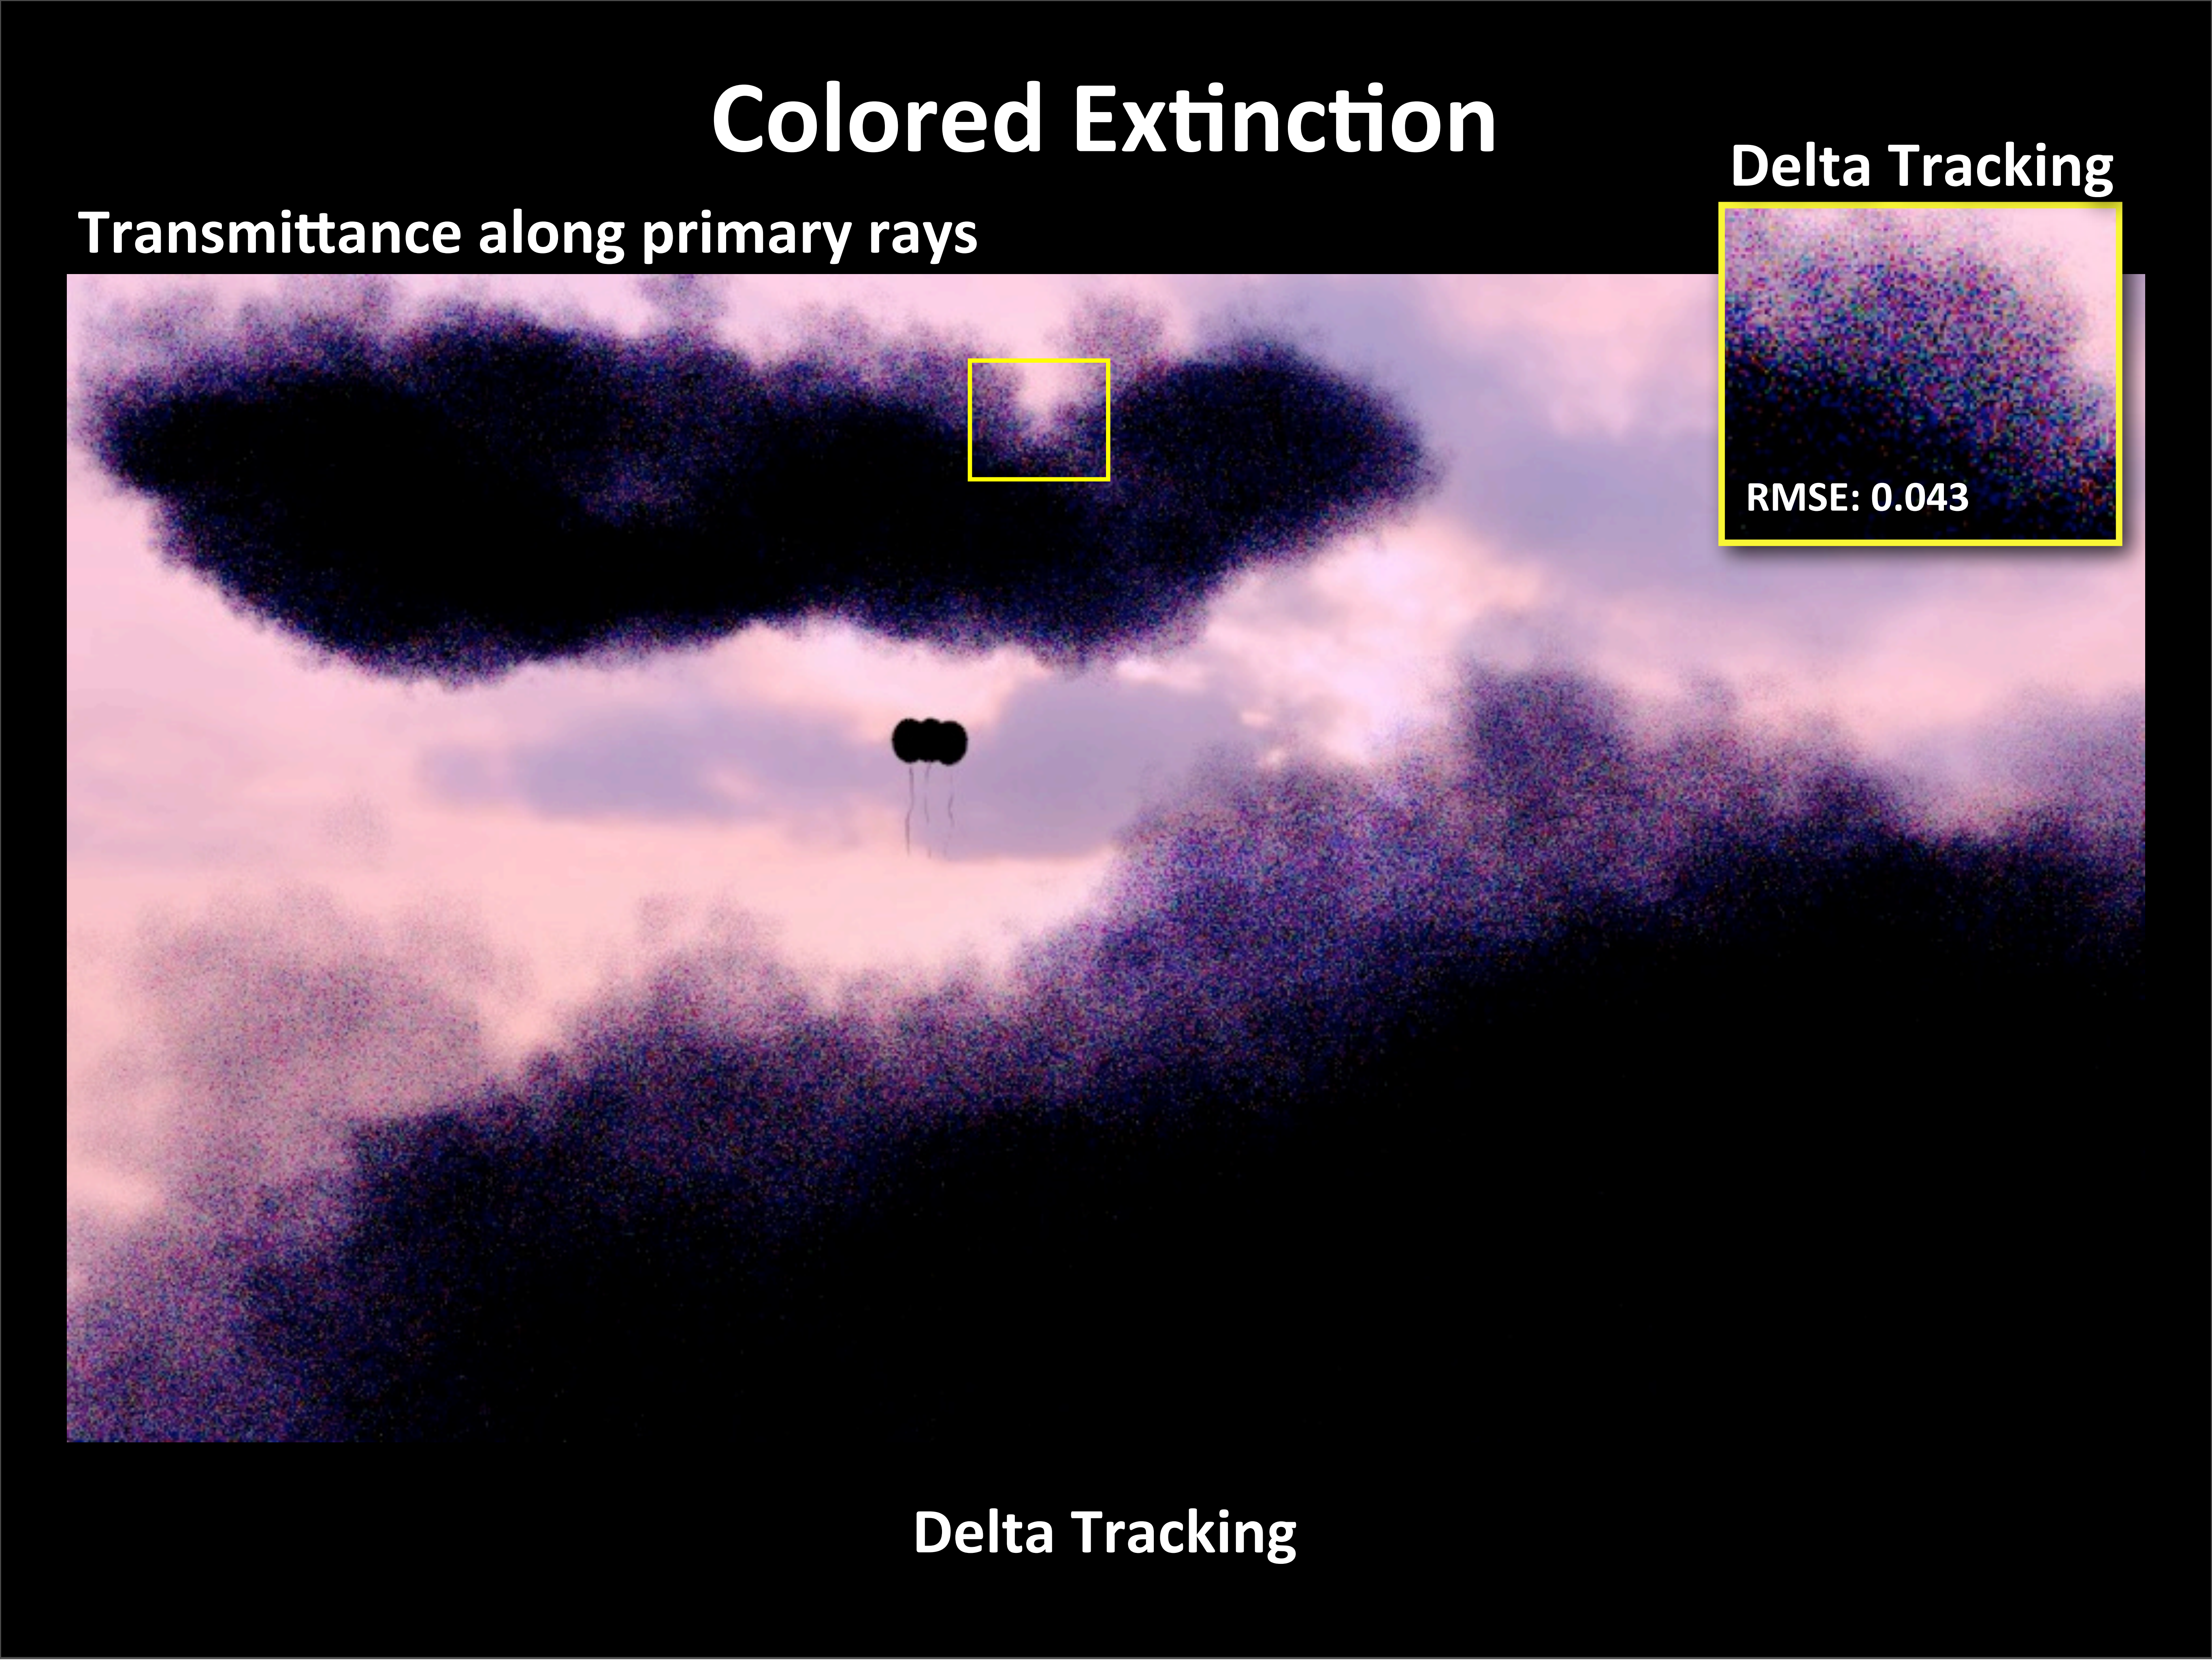

In this example, we have clouds with colored extinction. Let's first look at the transmittance along the primary rays. All images were rendered at equal cost.

Delta tracking produces a lot of color noise, as it needs to handle the transmittance through each color channel independently.... in order to be efficient.

Ratio tracking can efficiently handle all color channels at once, removing the color noise and if we combine it with residual tracking, we get additional significant reduction in variance.

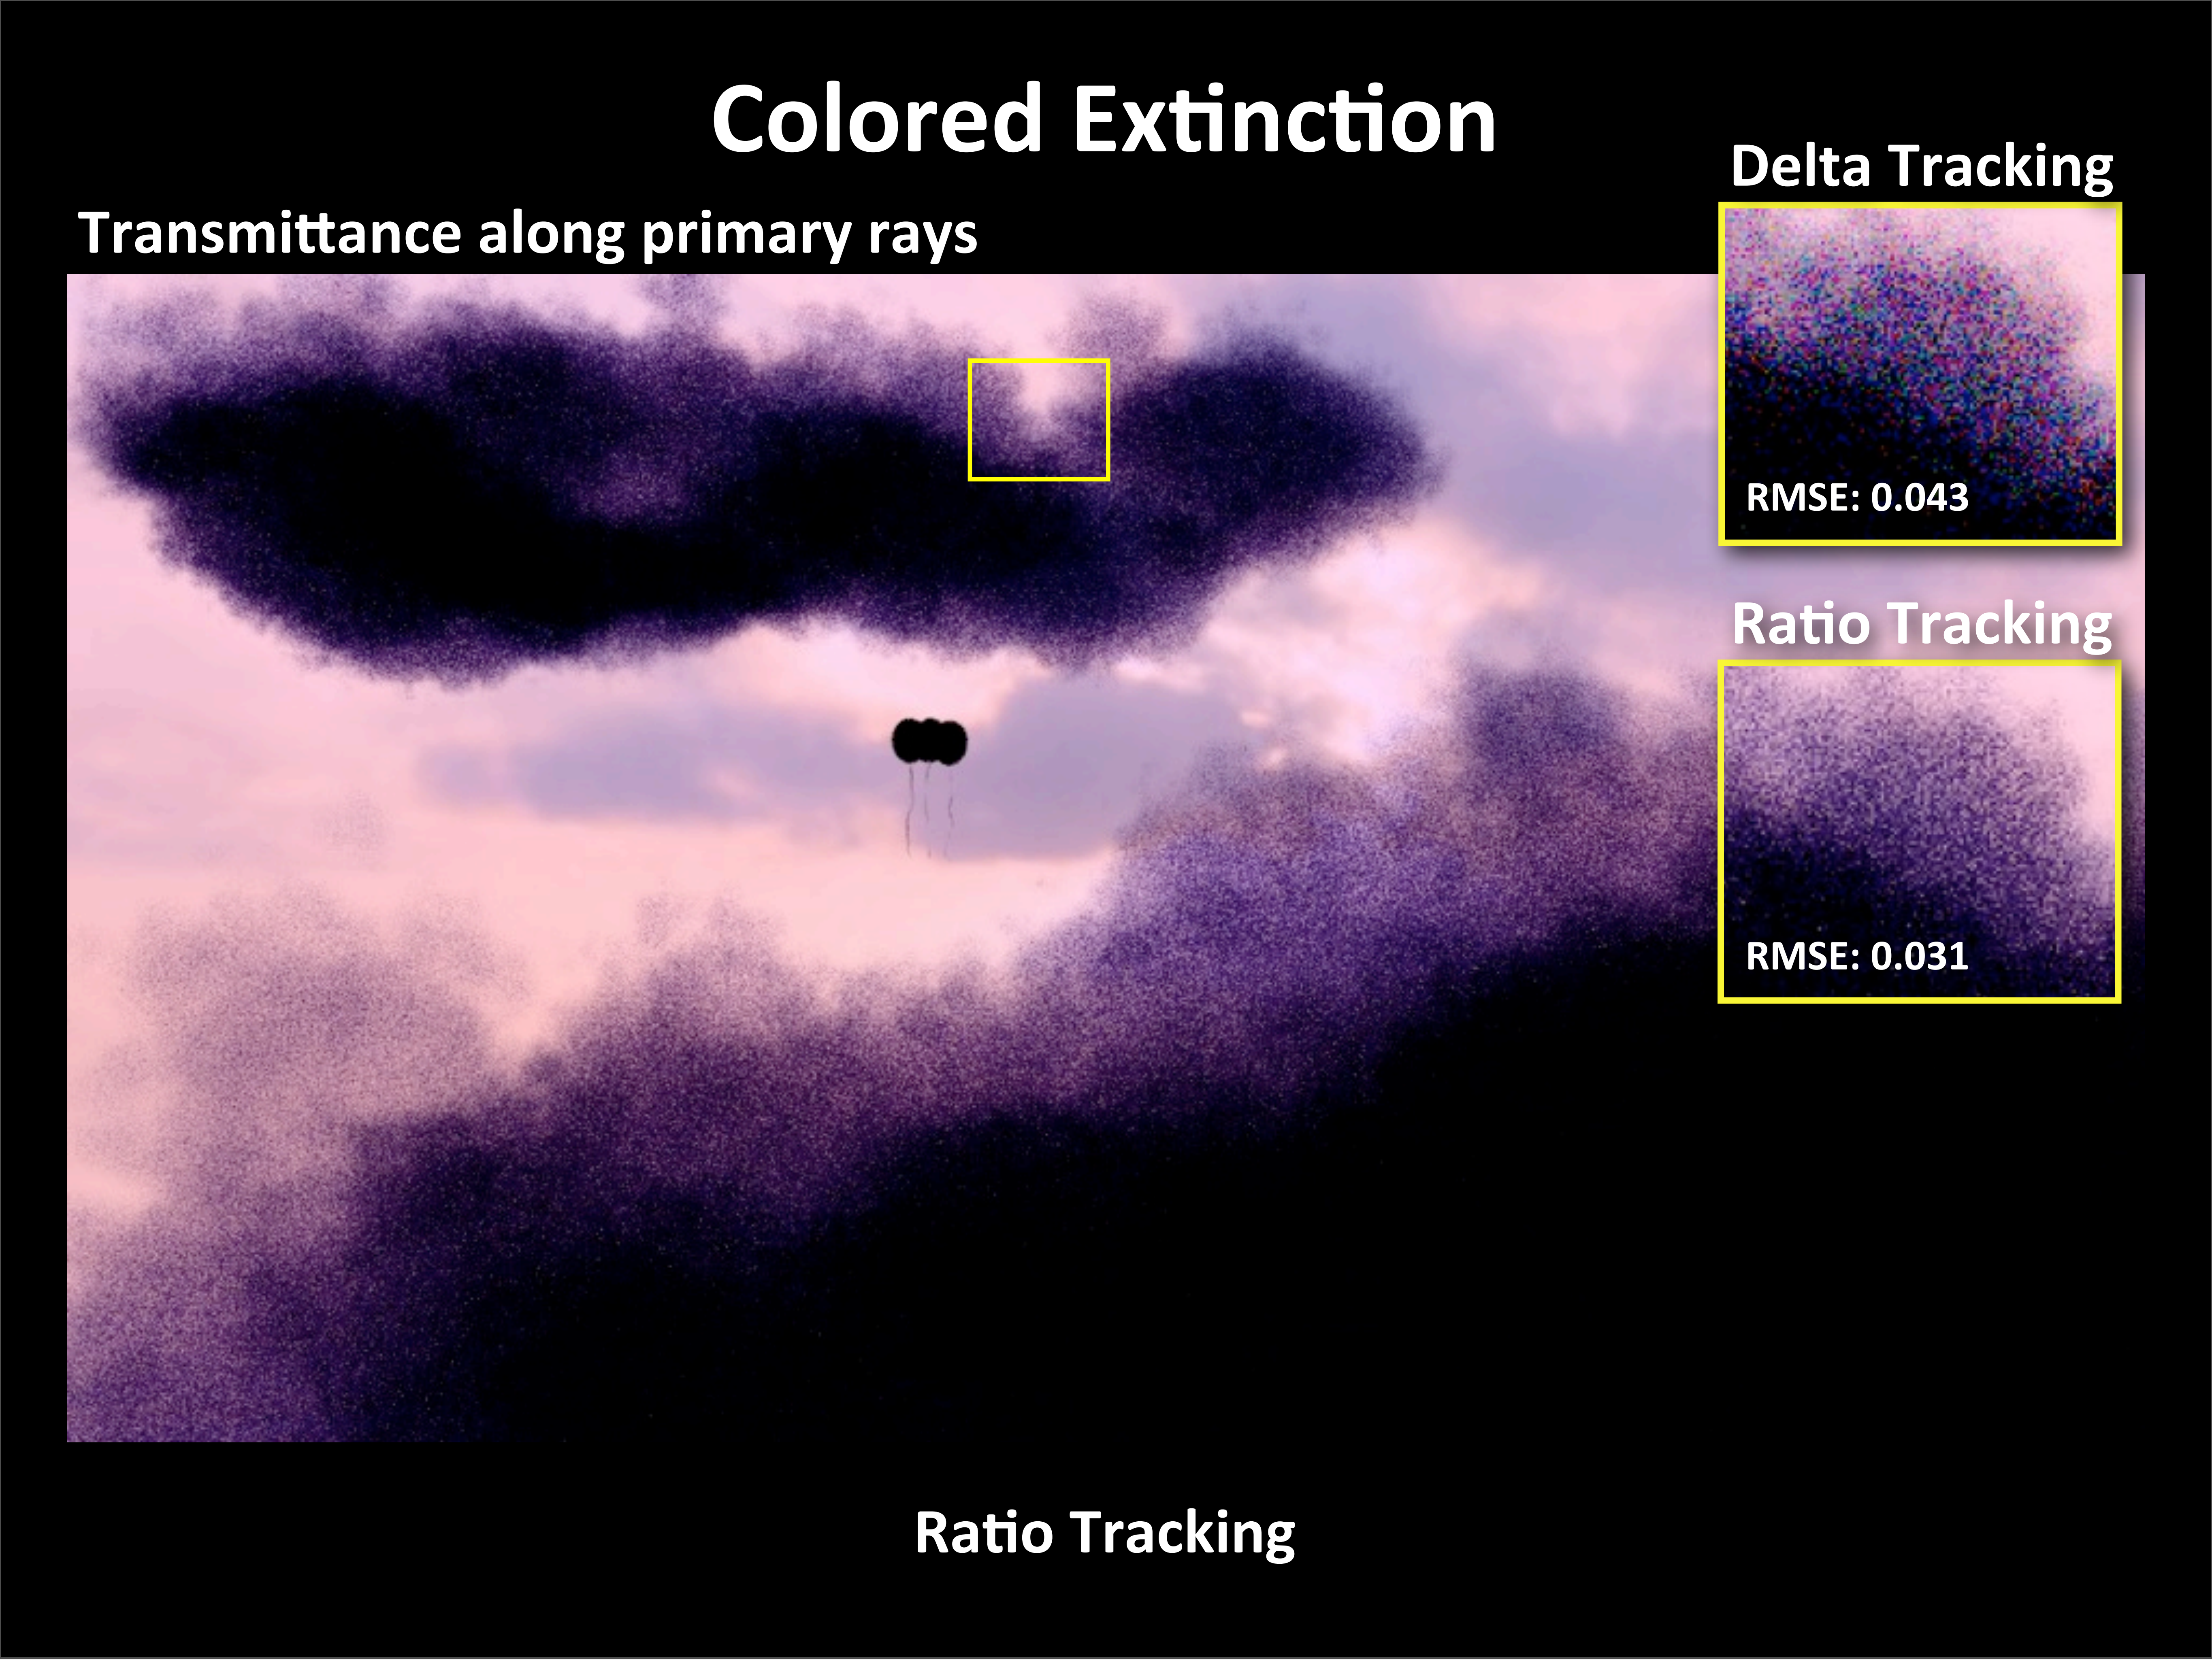

In this example, we have clouds with colored extinction. Let's first look at the transmittance along the primary rays. All images were rendered at equal cost.

Delta tracking produces a lot of color noise, as it needs to handle the transmittance through each color channel independently.... in order to be efficient.

Ratio tracking can efficiently handle all color channels at once, removing the color noise and if we combine it with residual tracking, we get additional significant reduction in variance.

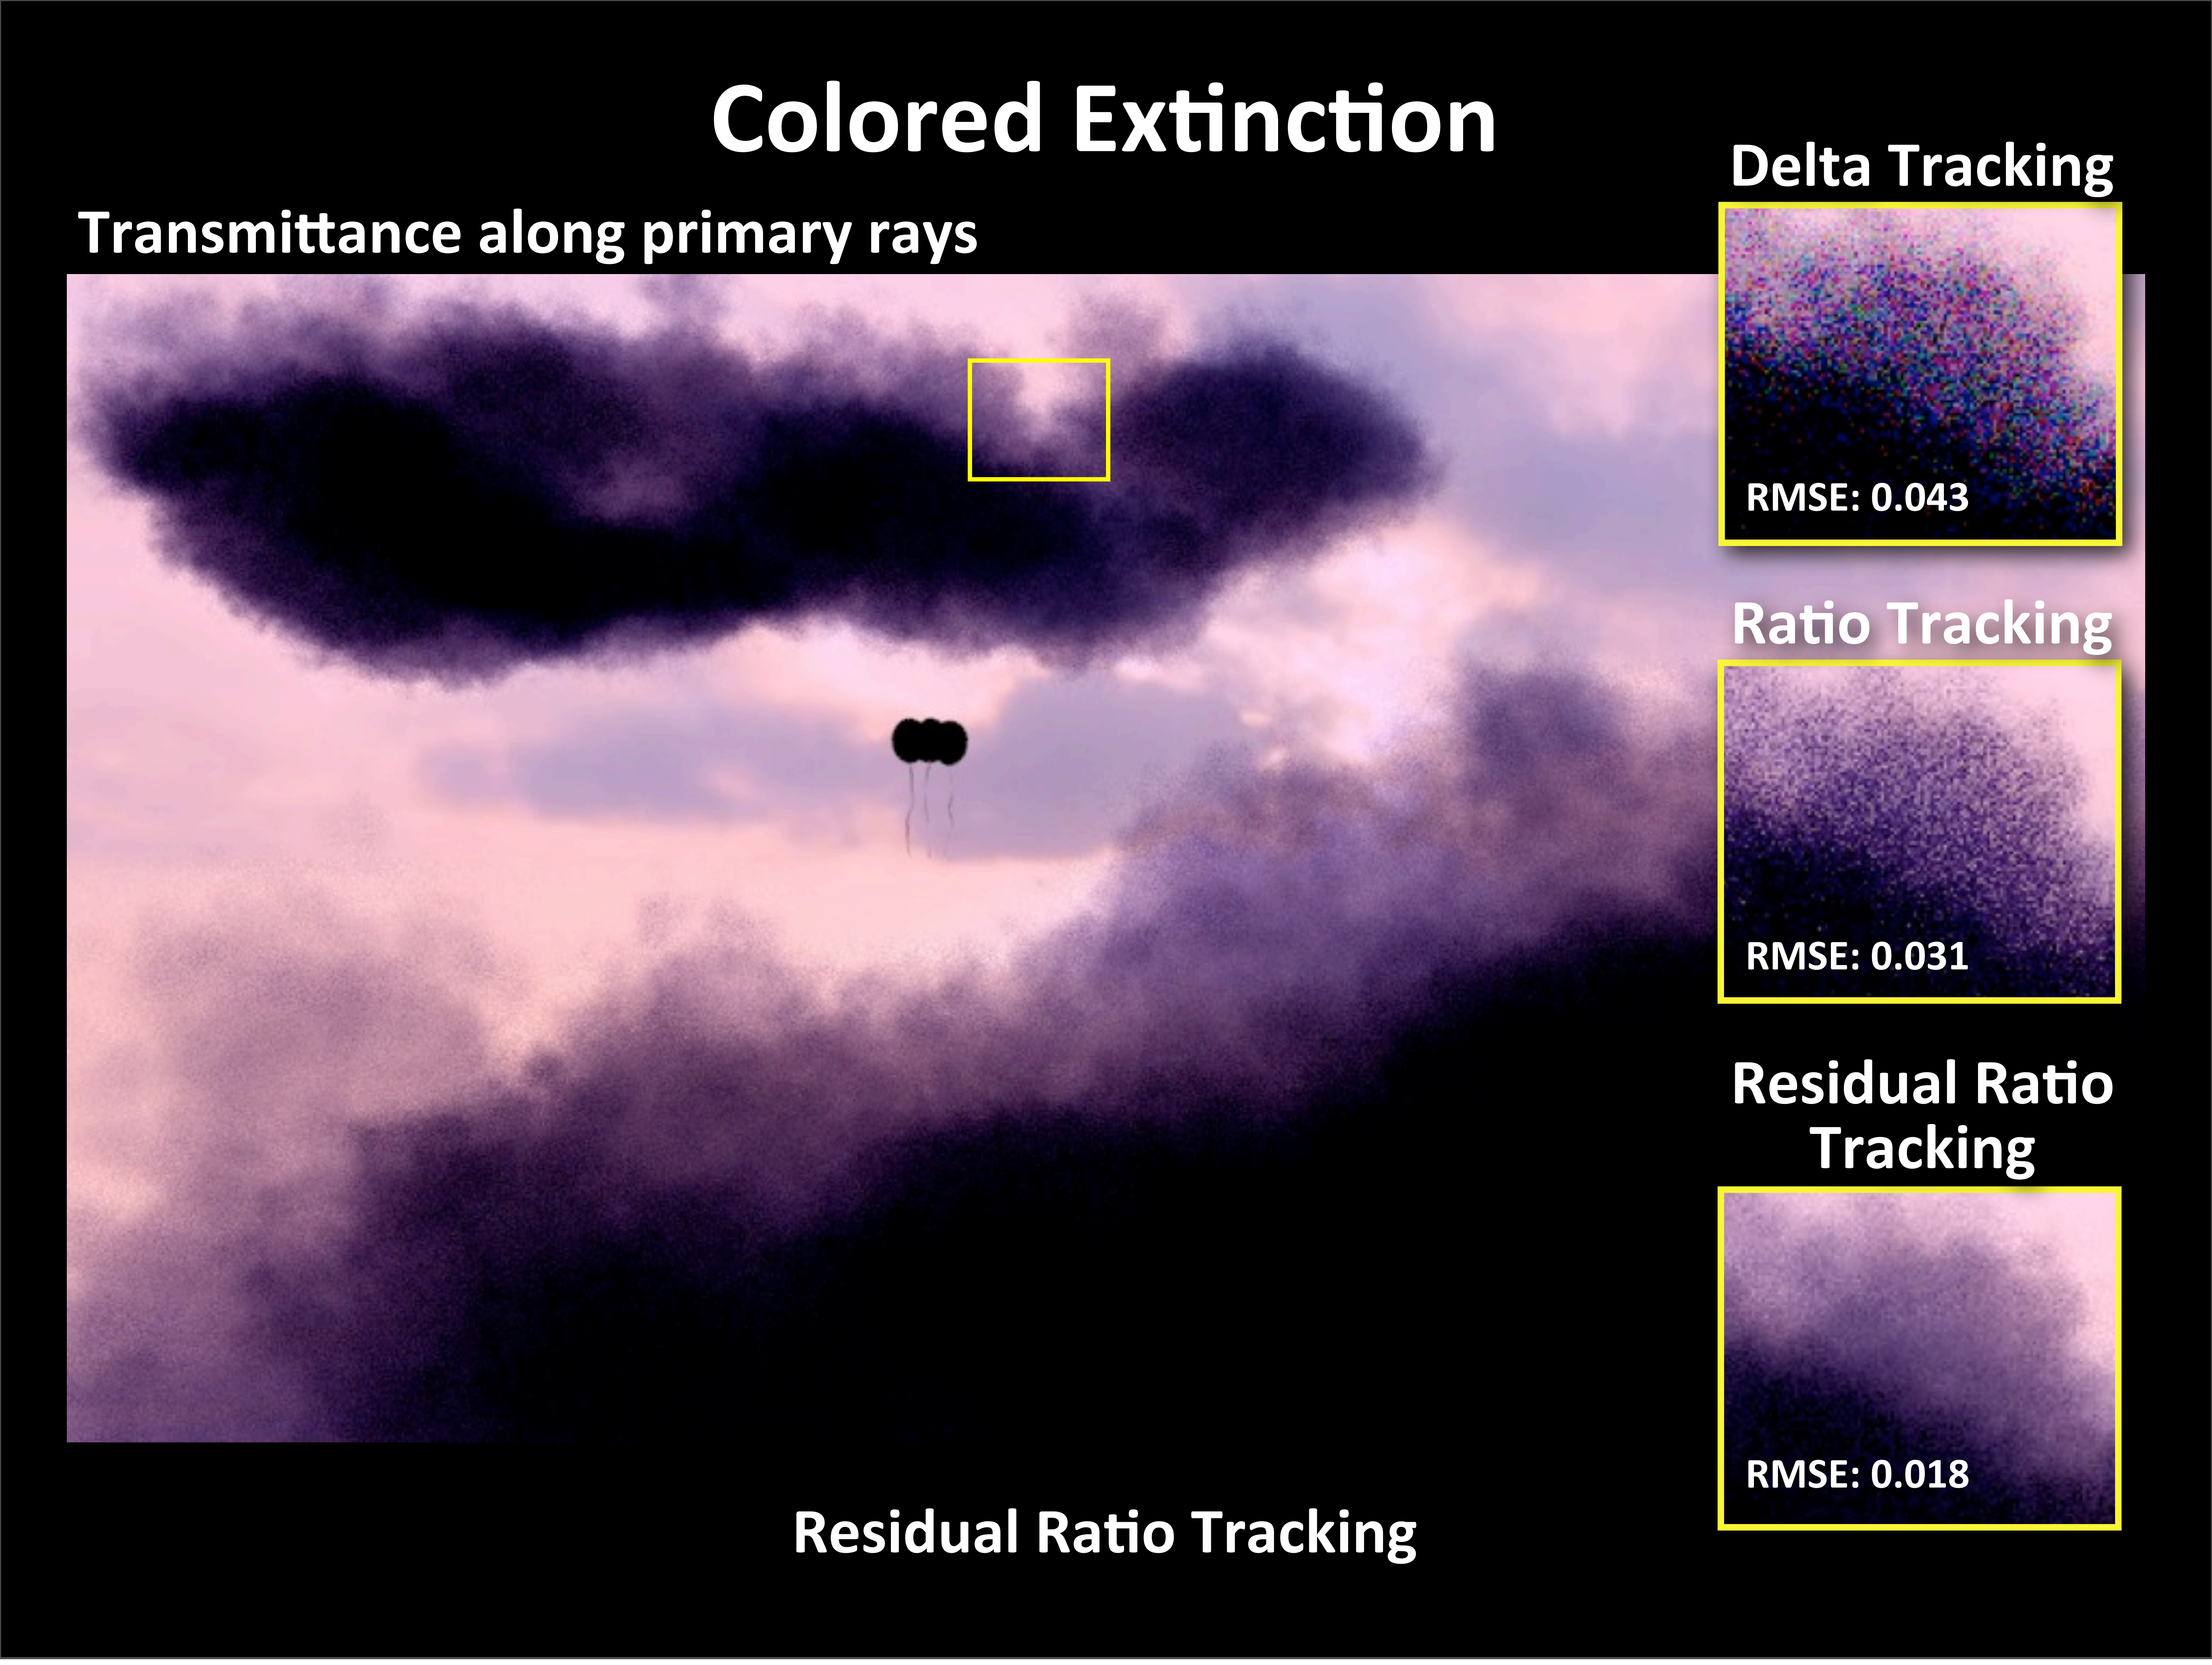

In this example, we have clouds with colored extinction. Let's first look at the transmittance along the primary rays. All images were rendered at equal cost.

Delta tracking produces a lot of color noise, as it needs to handle the transmittance through each color channel independently.... in order to be efficient.

Ratio tracking can efficiently handle all color channels at once, removing the color noise and if we combine it with residual tracking, we get additional significant reduction in variance.

Colored Extinction

Dual scattering

The next set of images shows the same scene, but this time rendered with two bounces in the medium.

You can see that even in this case, the reduction in variance that residual ratio tracking provides leads to a much lower overall amount of noise.

In the interest of time, I will refer you to the paper for additional results.

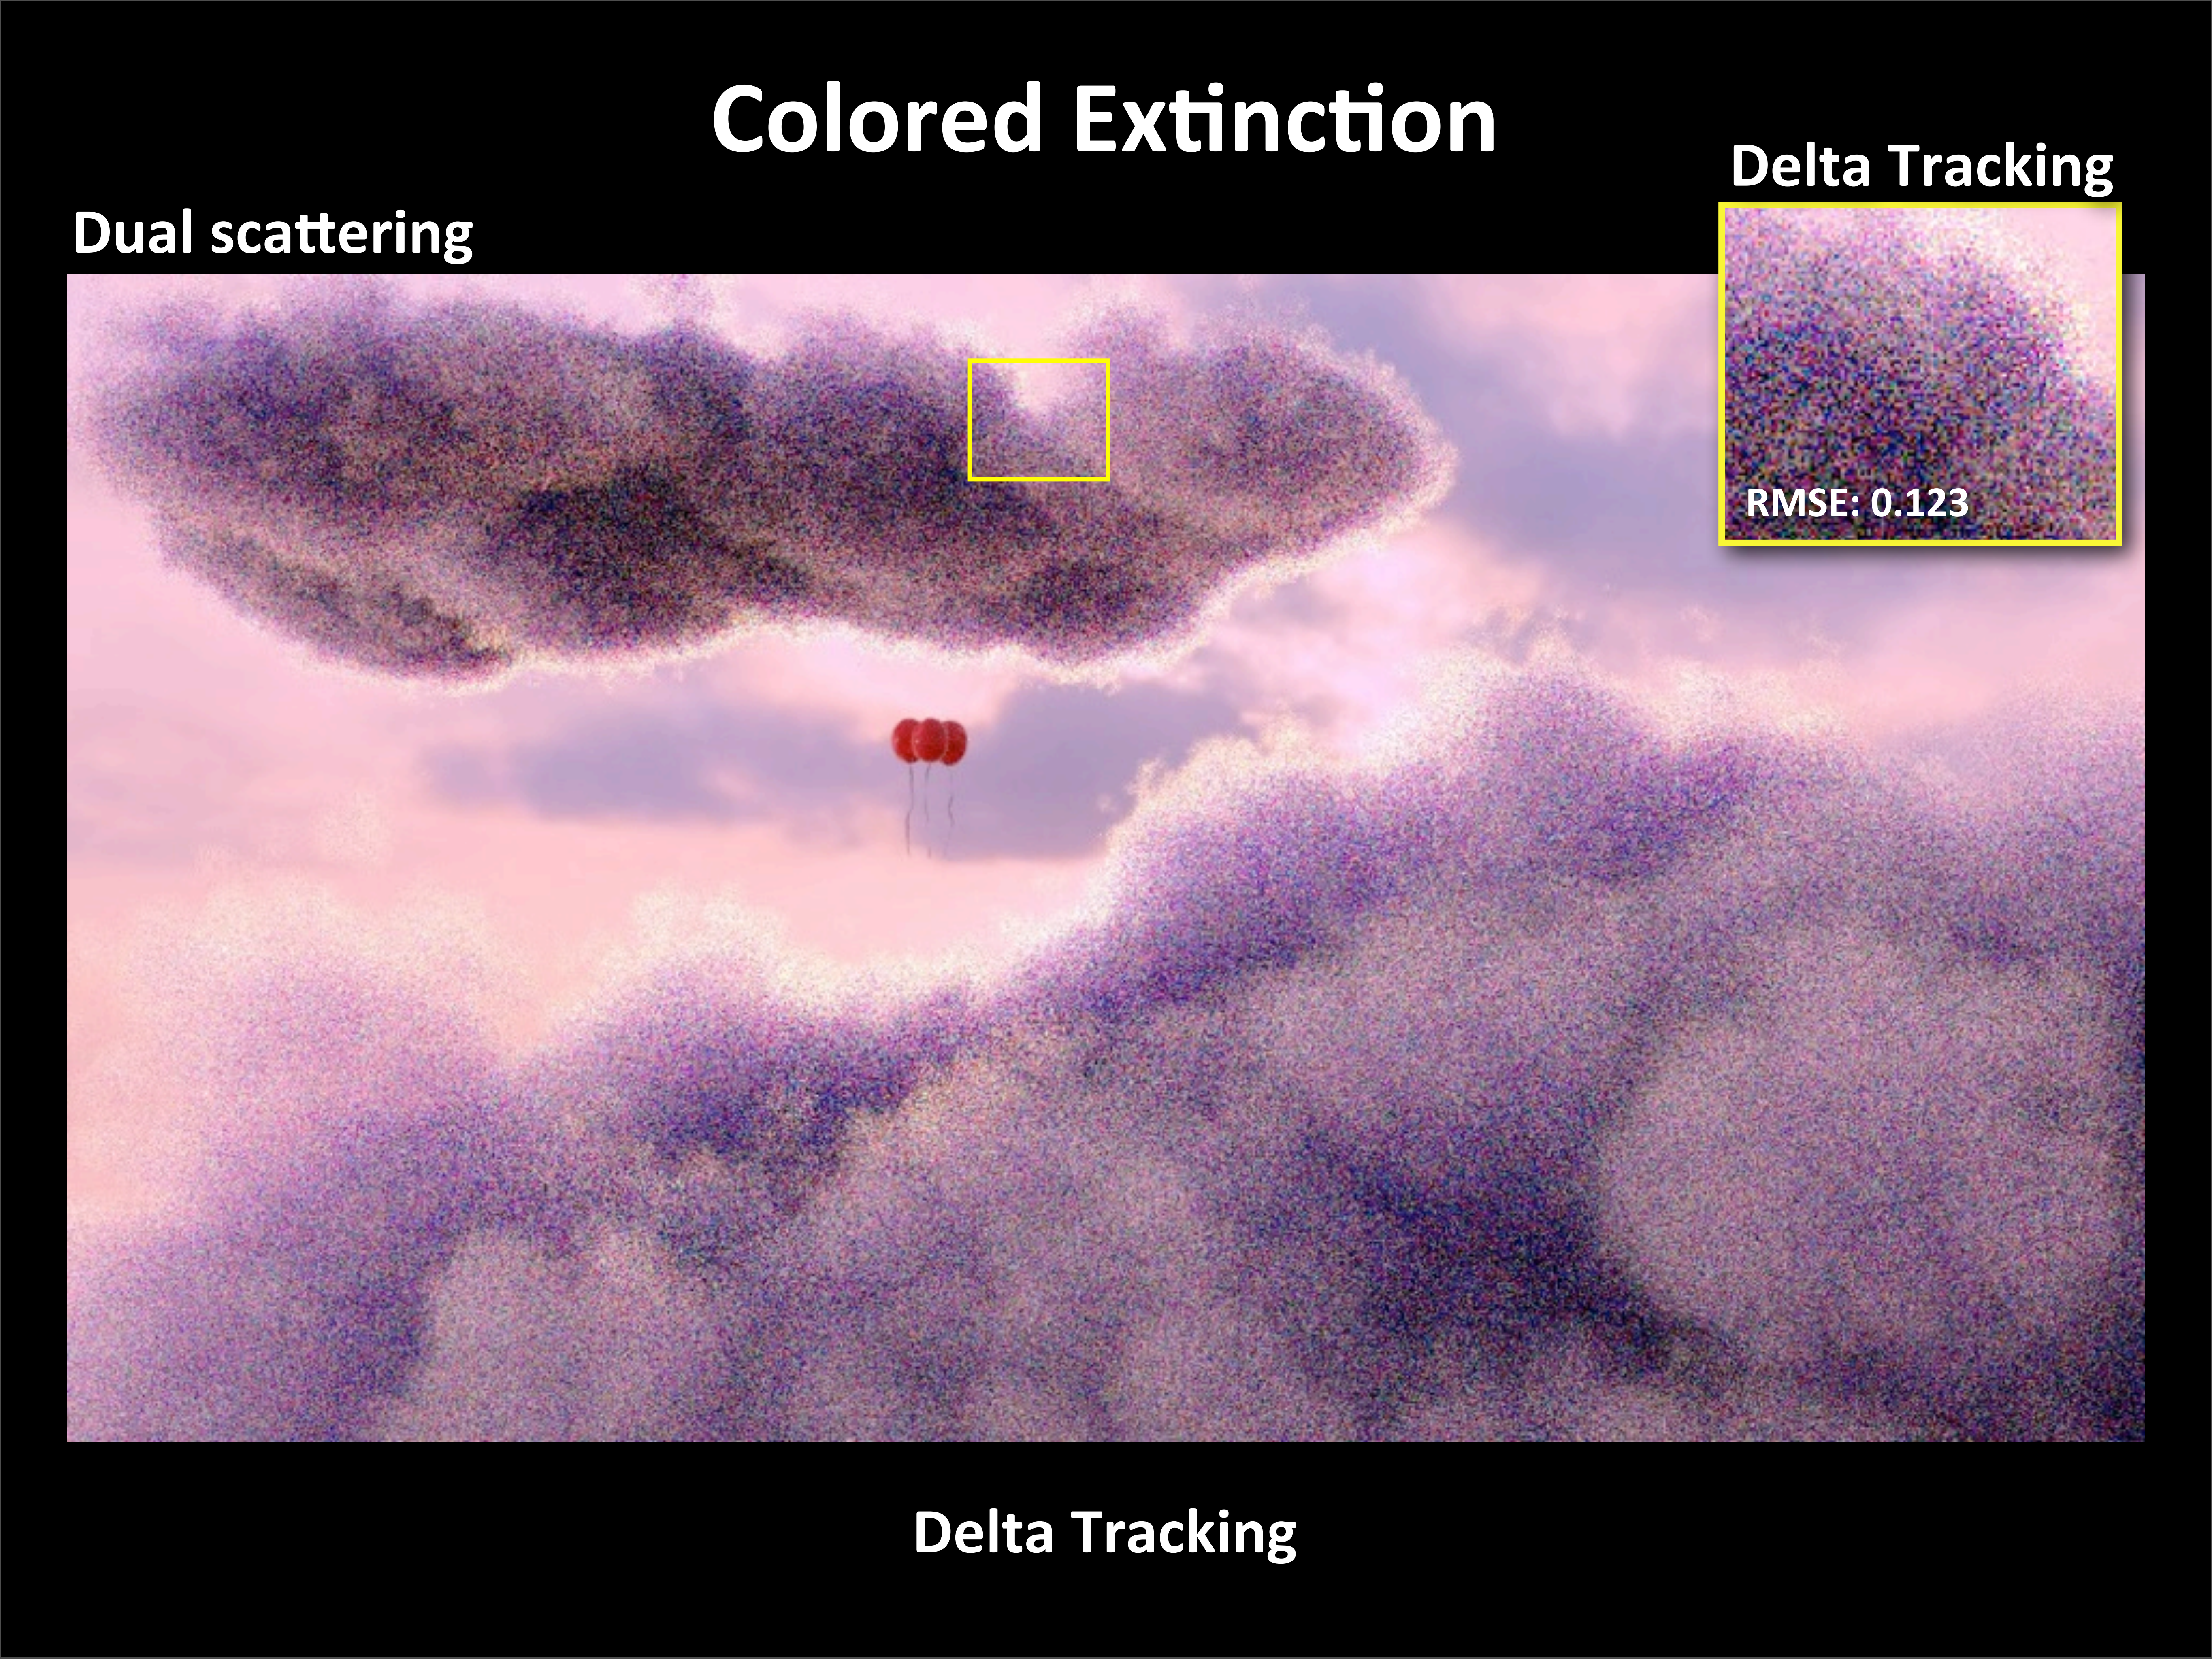

The next set of images shows the same scene, but this time rendered with two bounces in the medium.

You can see that even in this case, the reduction in variance that residual ratio tracking provides leads to a much lower overall amount of noise.

In the interest of time, I will refer you to the paper for additional results.

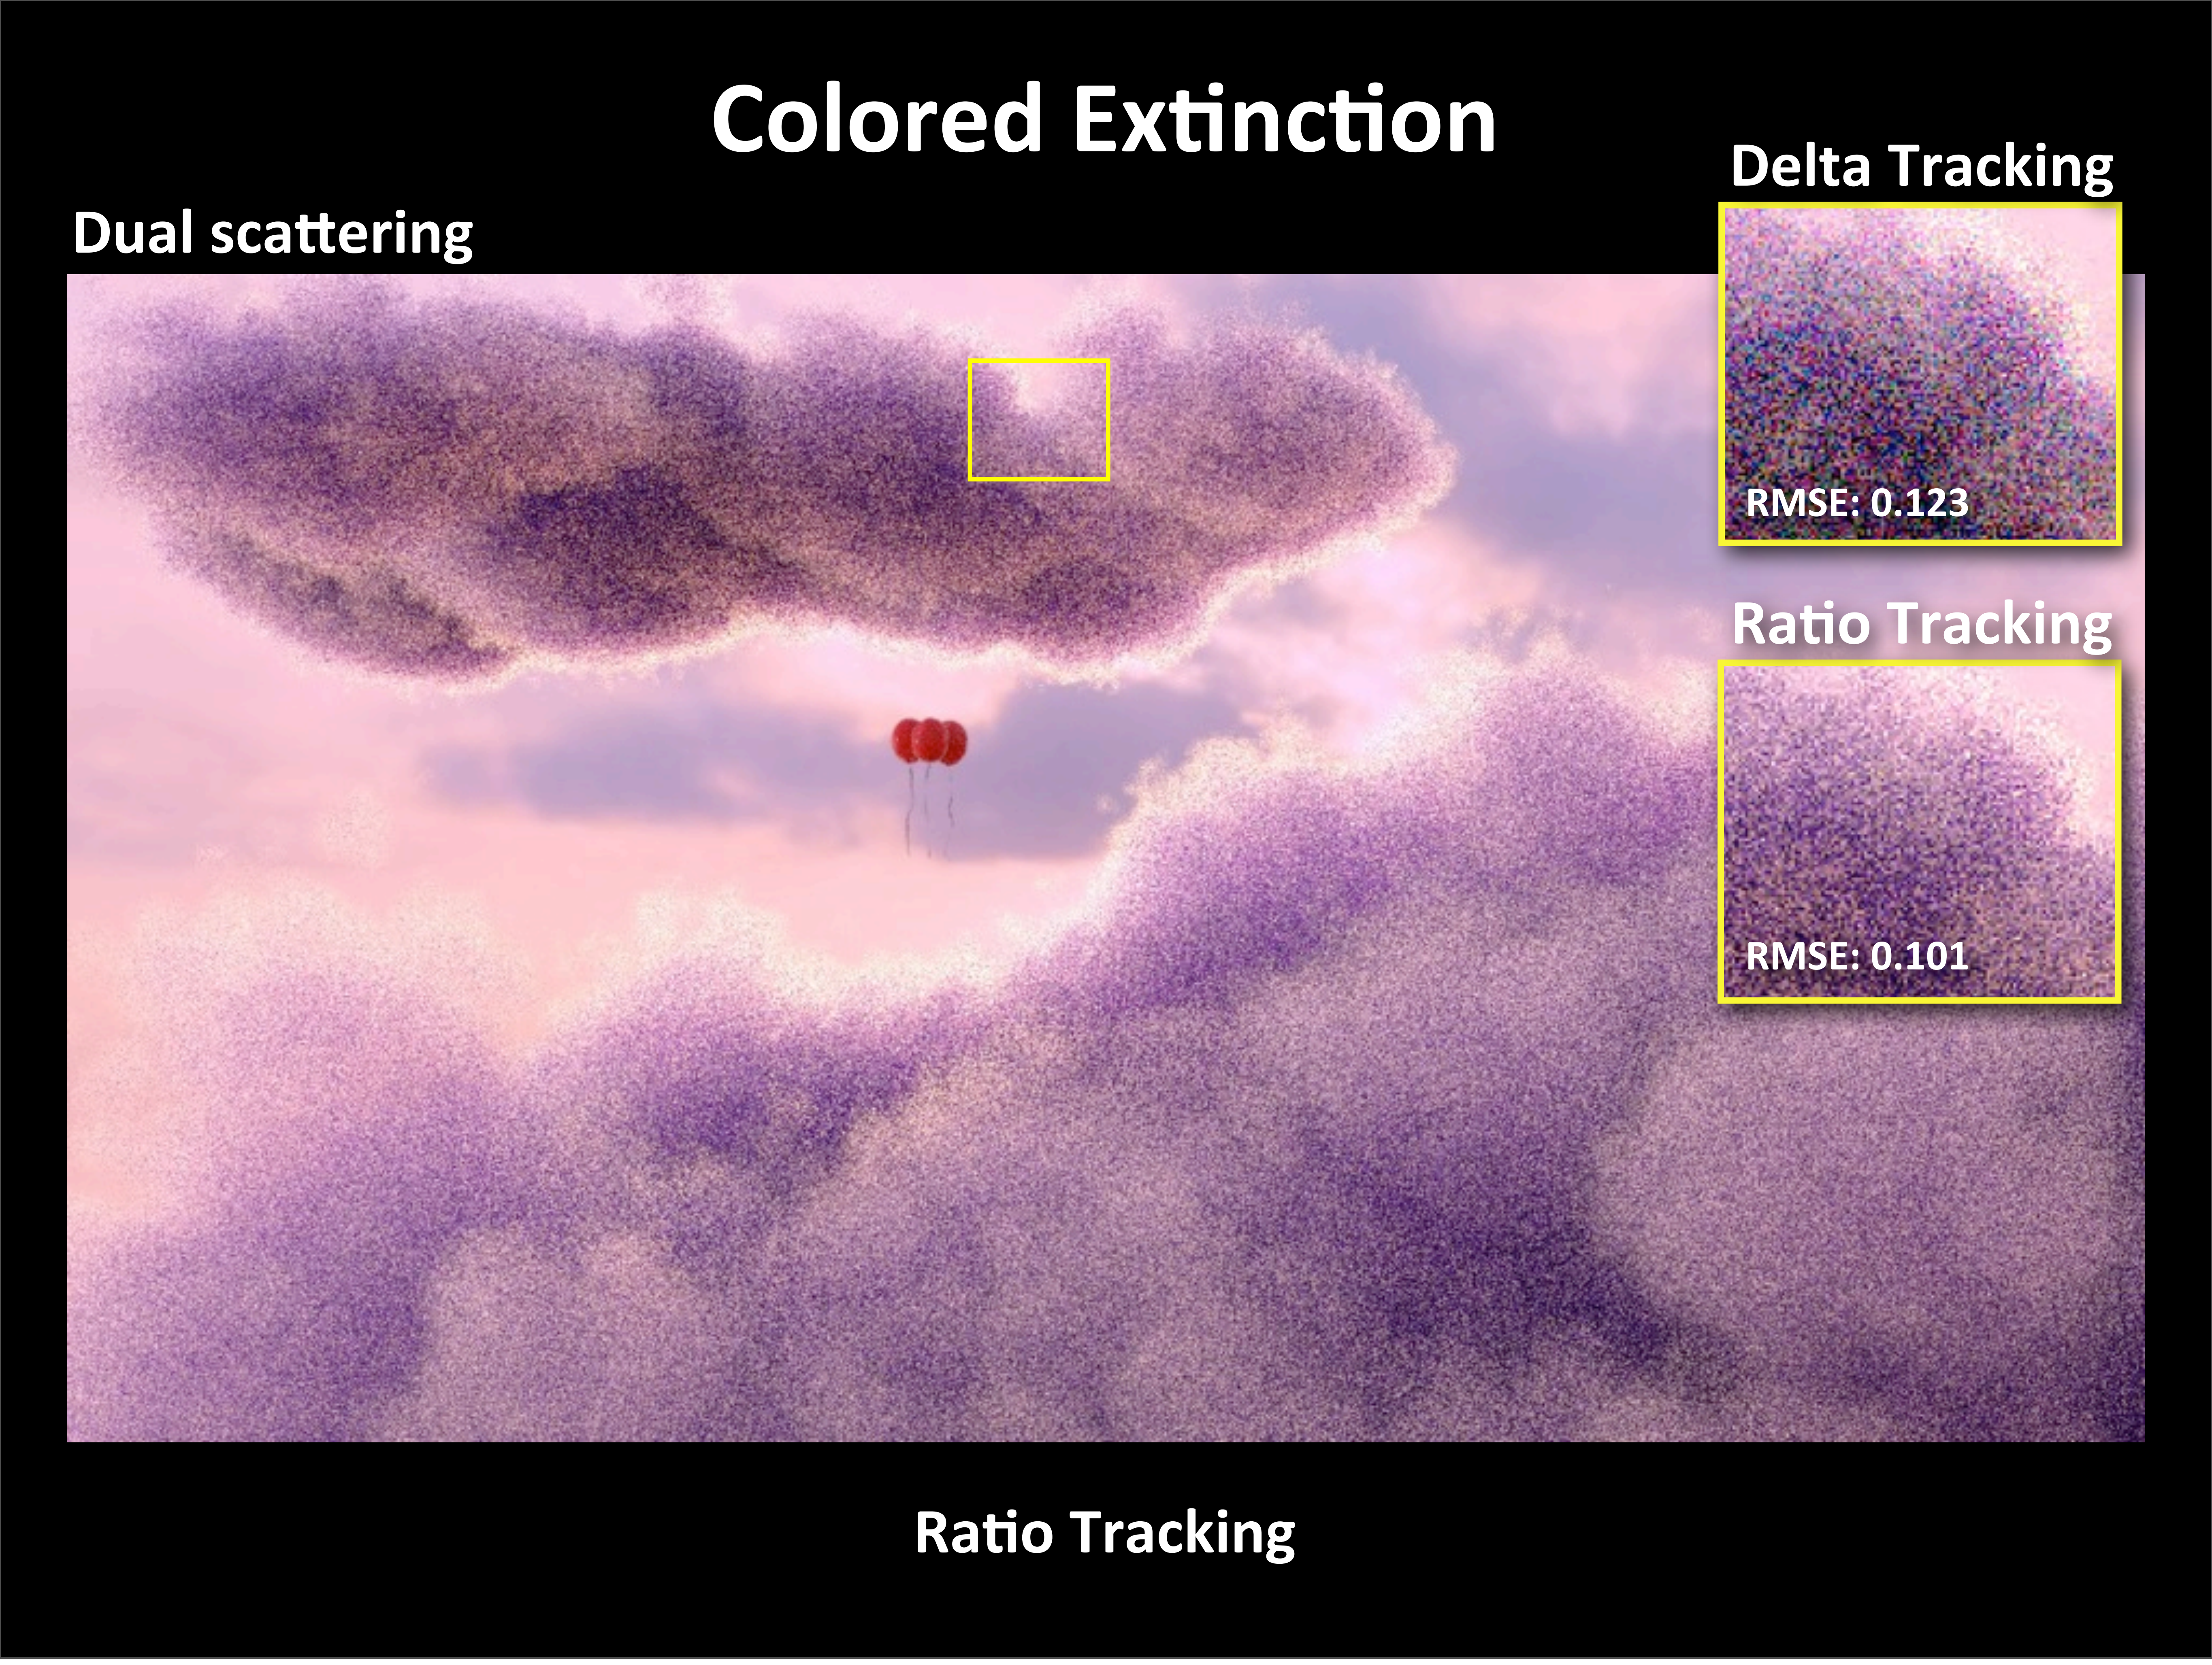

The next set of images shows the same scene, but this time rendered with two bounces in the medium.

You can see that even in this case, the reduction in variance that residual ratio tracking provides leads to a much lower overall amount of noise.

In the interest of time, I will refer you to the paper for additional results.

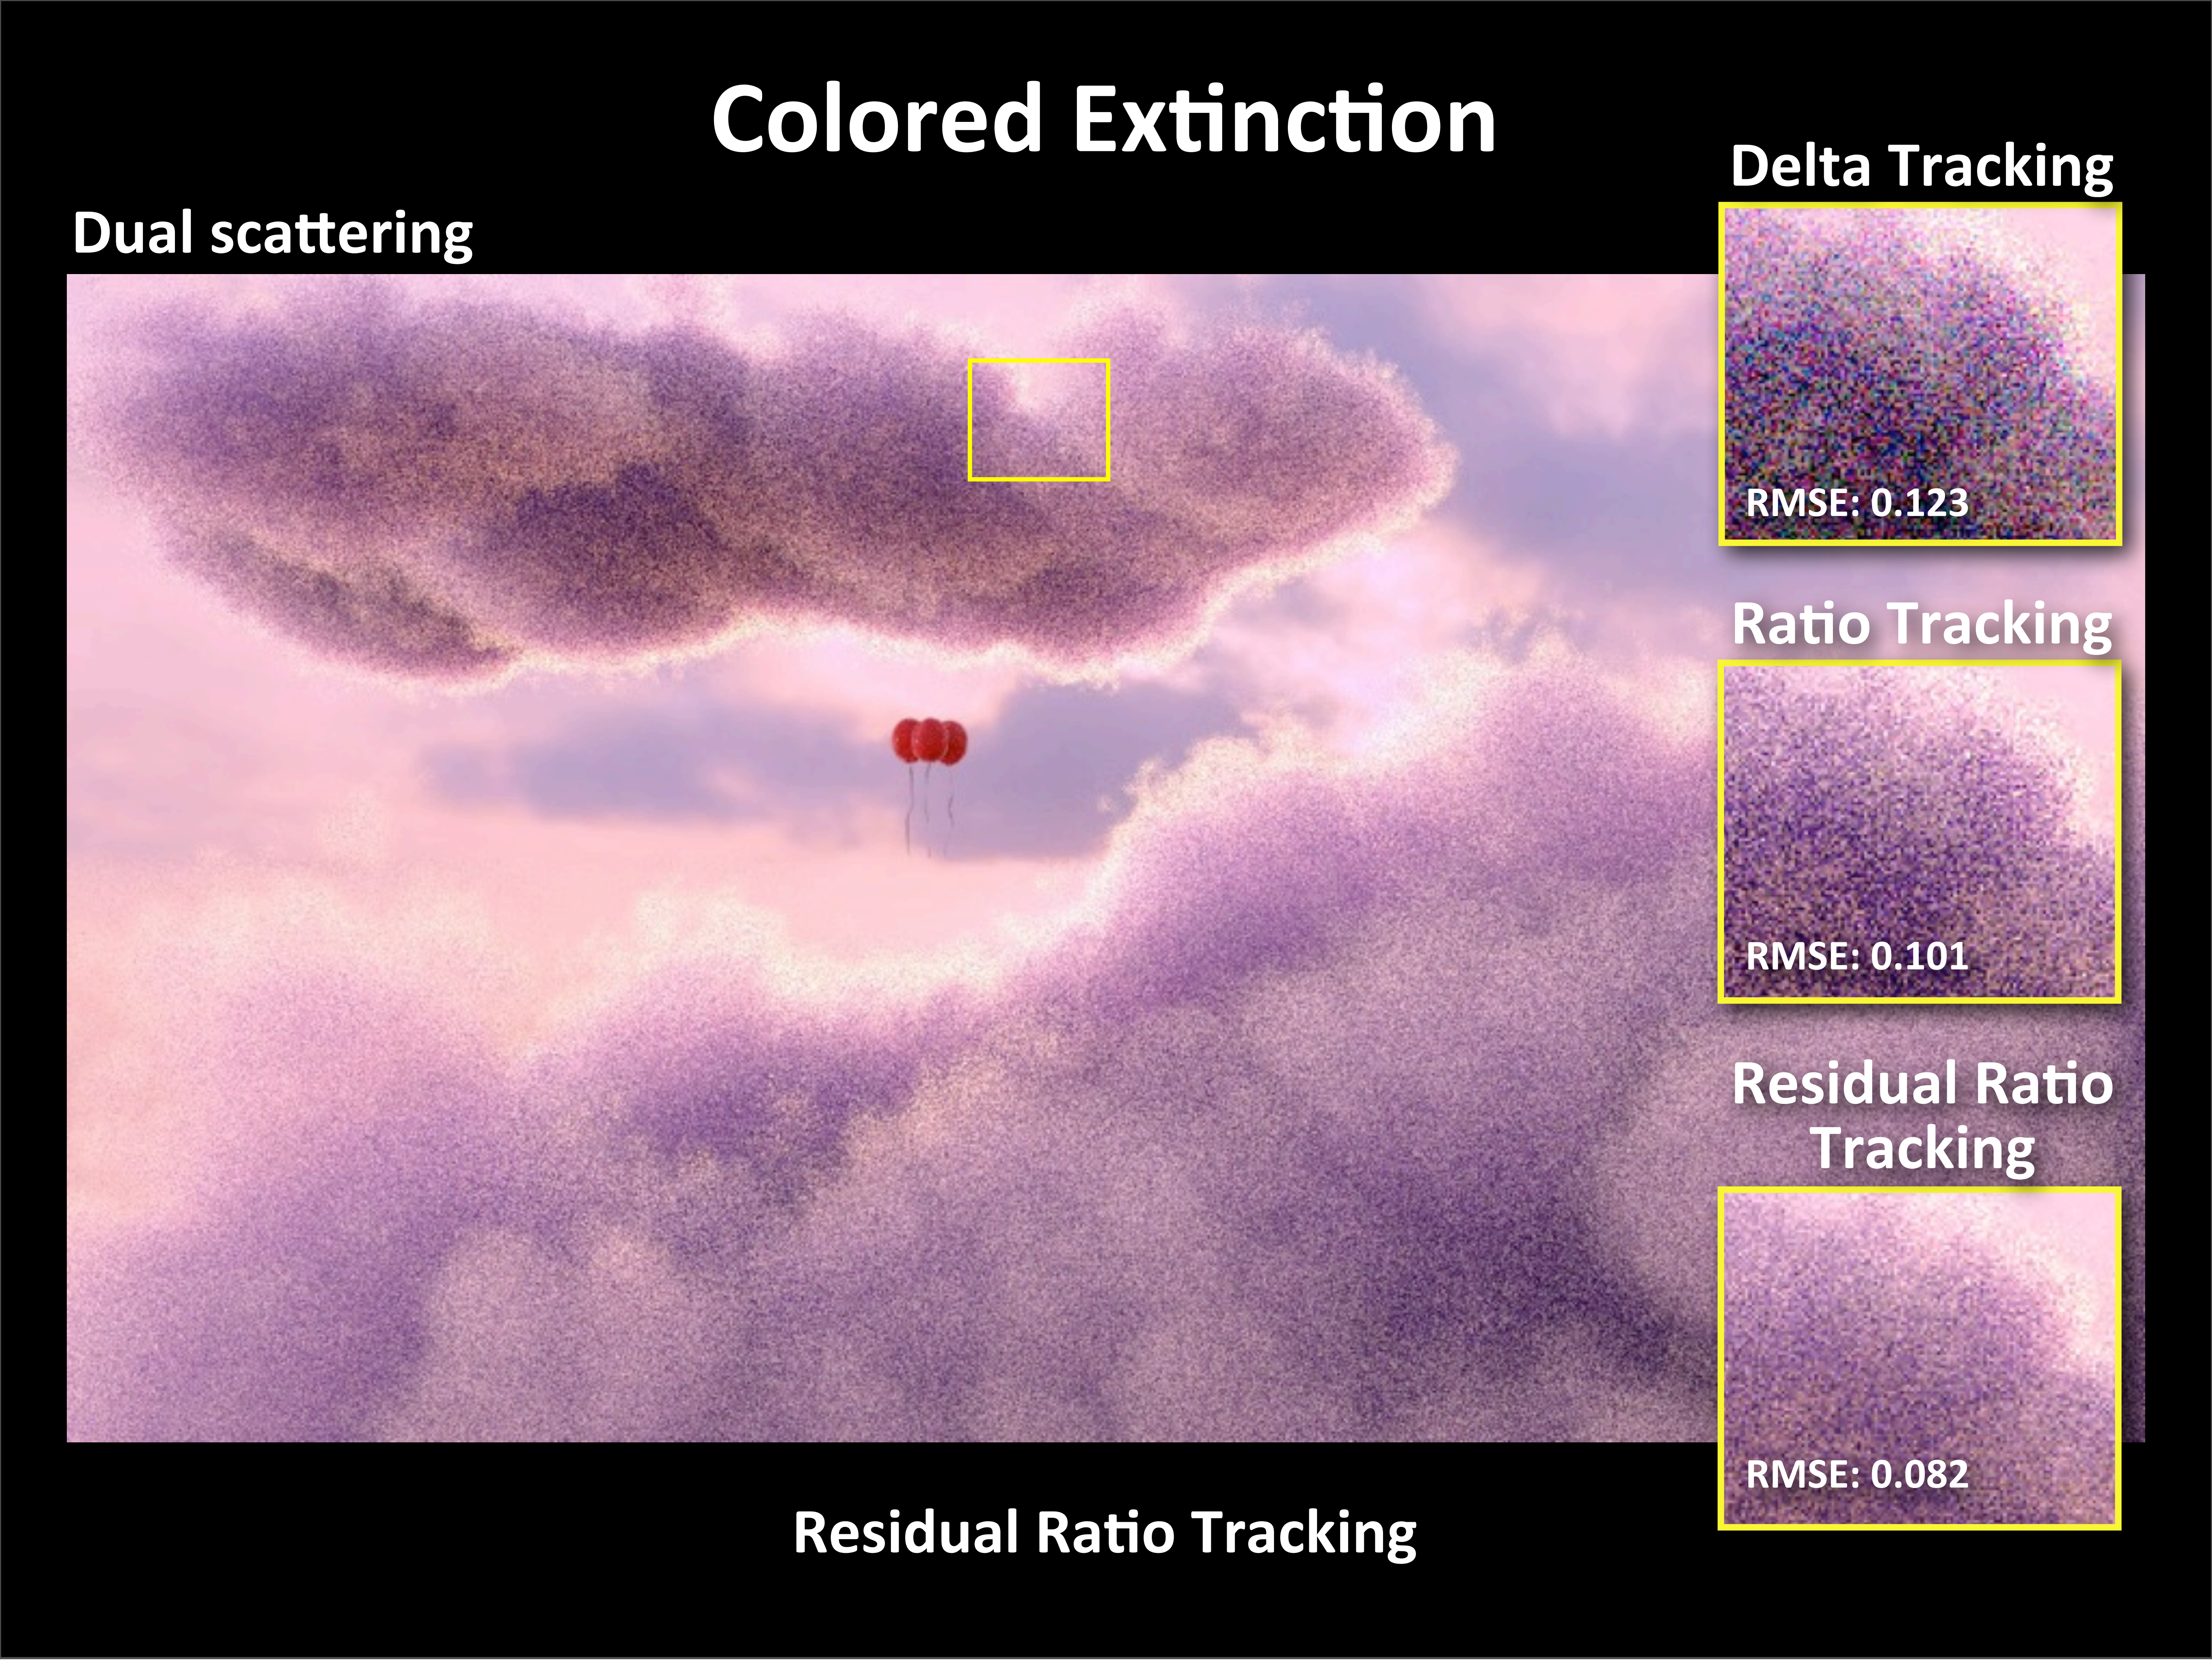

The next set of images shows the same scene, but this time rendered with two bounces in the medium.

You can see that even in this case, the reduction in variance that residual ratio tracking provides leads to a much lower overall amount of noise.

In the interest of time, I will refer you to the paper for additional results.

Conclusion

46

The techniques that I presented are weighted random walks
that provide a piecewise-constant or piecewise-exponential unbiased approximations

The Ratio tracking handles well media with high degree of heterogeneity.
Residual tracking then improves cases with low degree of heterogeneity.

In the future, we would like to investigate higher order basis functions for the control extinction
to make it better match the actual extinction.

On a more theoretical level, it would be interesting to further explore other variants of weighted (or non-analog) estimator.
For those interested I recommend looking at the integral formulation of tracking algorithms by Galtier et al., which provides a nice theoretical framework for expressing different variants of the tracking.

Conclusion

SIGGRAPH ASIA 2014 SHENZHEN

- **Ratio Tracking**
  - weighted random walk
  - piecewise-constant approximation (instead of binary)
  - efficient with loose majorants and colored extinctions
  - handles negative extinction coefficients

The techniques that I presented are weighted random walks
that provide a piecewise-constant or piecewise-exponential unbiased approximations

The Ratio tracking handles well media with high degree of heterogeneity.
Residual tracking then improves cases with low degree of heterogeneity.

In the future, we would like to investigate higher order basis functions for the control extinction
to make it better match the actual extinction.

On a more theoretical level, it would be interesting to further explore other variants of weighted (or non-analog) estimator.
For those interested I recommend looking at the integral formulation of tracking algorithms by Galtier et al., which provides a nice theoretical framework for expressing different variants of the tracking.

Conclusion

SIGGRAPH ASIA 2014 SHENZHEN

- **Ratio Tracking**
  - weighted random walk
  - piecewise-constant approximation (instead of binary)
  - efficient with loose majorants and colored extinctions
  - handles negative extinction coefficients
- **Residual Tracking**
  - decomposition into control and residual parts
  - piecewise-exponential approximation (analytic + numerical computation)
  - efficiently handles nearly homogeneous volumes
  - residual majorant is smaller -> cost reduction

The techniques that I presented are weighted random walks
that provide a piecewise-constant or piecewise-exponential unbiased approximations

The Ratio tracking handles well media with high degree of heterogeneity.
Residual tracking then improves cases with low degree of heterogeneity.

In the future, we would like to investigate higher order basis functions for the control extinction
to make it better match the actual extinction.

On a more theoretical level, it would be interesting to further explore other variants of weighted (or non-analog) estimator.
For those interested I recommend looking at the integral formulation of tracking algorithms by Galtier et al., which provides a nice theoretical framework for expressing different variants of the tracking.

Conclusion

SIGGRAPH ASIA 2014 SHENZHEN

Ratio Tracking

- weighted random walk
- piecewise-constant approximation (instead of binary)
- efficient with loose majorants and colored extinctions
- handles negative extinction coefficients

Residual Tracking

- decomposition into control and residual parts
- piecewise-exponential approximation (analytic + numerical computation)
- efficiently handles nearly homogeneous volumes
- residual majorant is smaller -> cost reduction

Future Work

- better basis functions for the control extinction
- integral formulation of tracking algorithms by Galtier et al. [2013]

The techniques that I presented are weighted random walks that provide a piecewise-constant or piecewise-exponential unbiased approximations

The Ratio tracking handles well media with high degree of heterogeneity. Residual tracking then improves cases with low degree of heterogeneity.

In the future, we would like to investigate higher order basis functions for the control extinction to make it better match the actual extinction.

On a more theoretical level, it would be interesting to further explore other variants of weighted (or non-analog) estimator.
For those interested I recommend looking at the integral formulation of tracking algorithms by Galtier et al., which provides a nice theoretical framework for expressing different variants of the tracking.

Summary

SIGGRAPH ASIA 2014 SHENZHEN

Disney BIG HERO 6

PIXAR's RenderMan

Acknowledgements
Simon Kallweit
Ralf Habel
Hyperion team @ WDAS

The tracking is very easy to implement, if you already have Delta tracking with some accelerating structure, it will take you less then an hour.... including debugging. The technique was used in all volumetric shots in the Big Hero 6 and parts of it are expected to appear in the next release of Renderman.

Residual Ratio ~~Delta~~ tracking

$T_c = \exp(-\mu_c * d)$

$T = 1$

$t = 0$

do:

$t = t - \frac{\log(1-\text{rand}())}{\bar{\mu}}$

if $t \geq d$: **break**

$\cancel{\xi = \text{rand}()}$ $T = T * \left(1 - \frac{\mu(o+t*\omega) - \mu_c}{\bar{\mu}}\right)$

while $\cancel{\xi > \frac{\mu(o+t*\omega)}{\bar{\mu}}}$ 1

return $\cancel{t \geq d}$ $T * T$

Acknowledgements
Simon Kallweit
Ralf Habel
Hyperion team @ WDAS

The tracking is very easy to implement, if you already have Delta tracking with some accelerating structure, it will take you less then an hour.... including debugging. The technique was used in all volumetric shots in the Big Hero 6 and parts of it are expected to appear in the next release of Renderman.

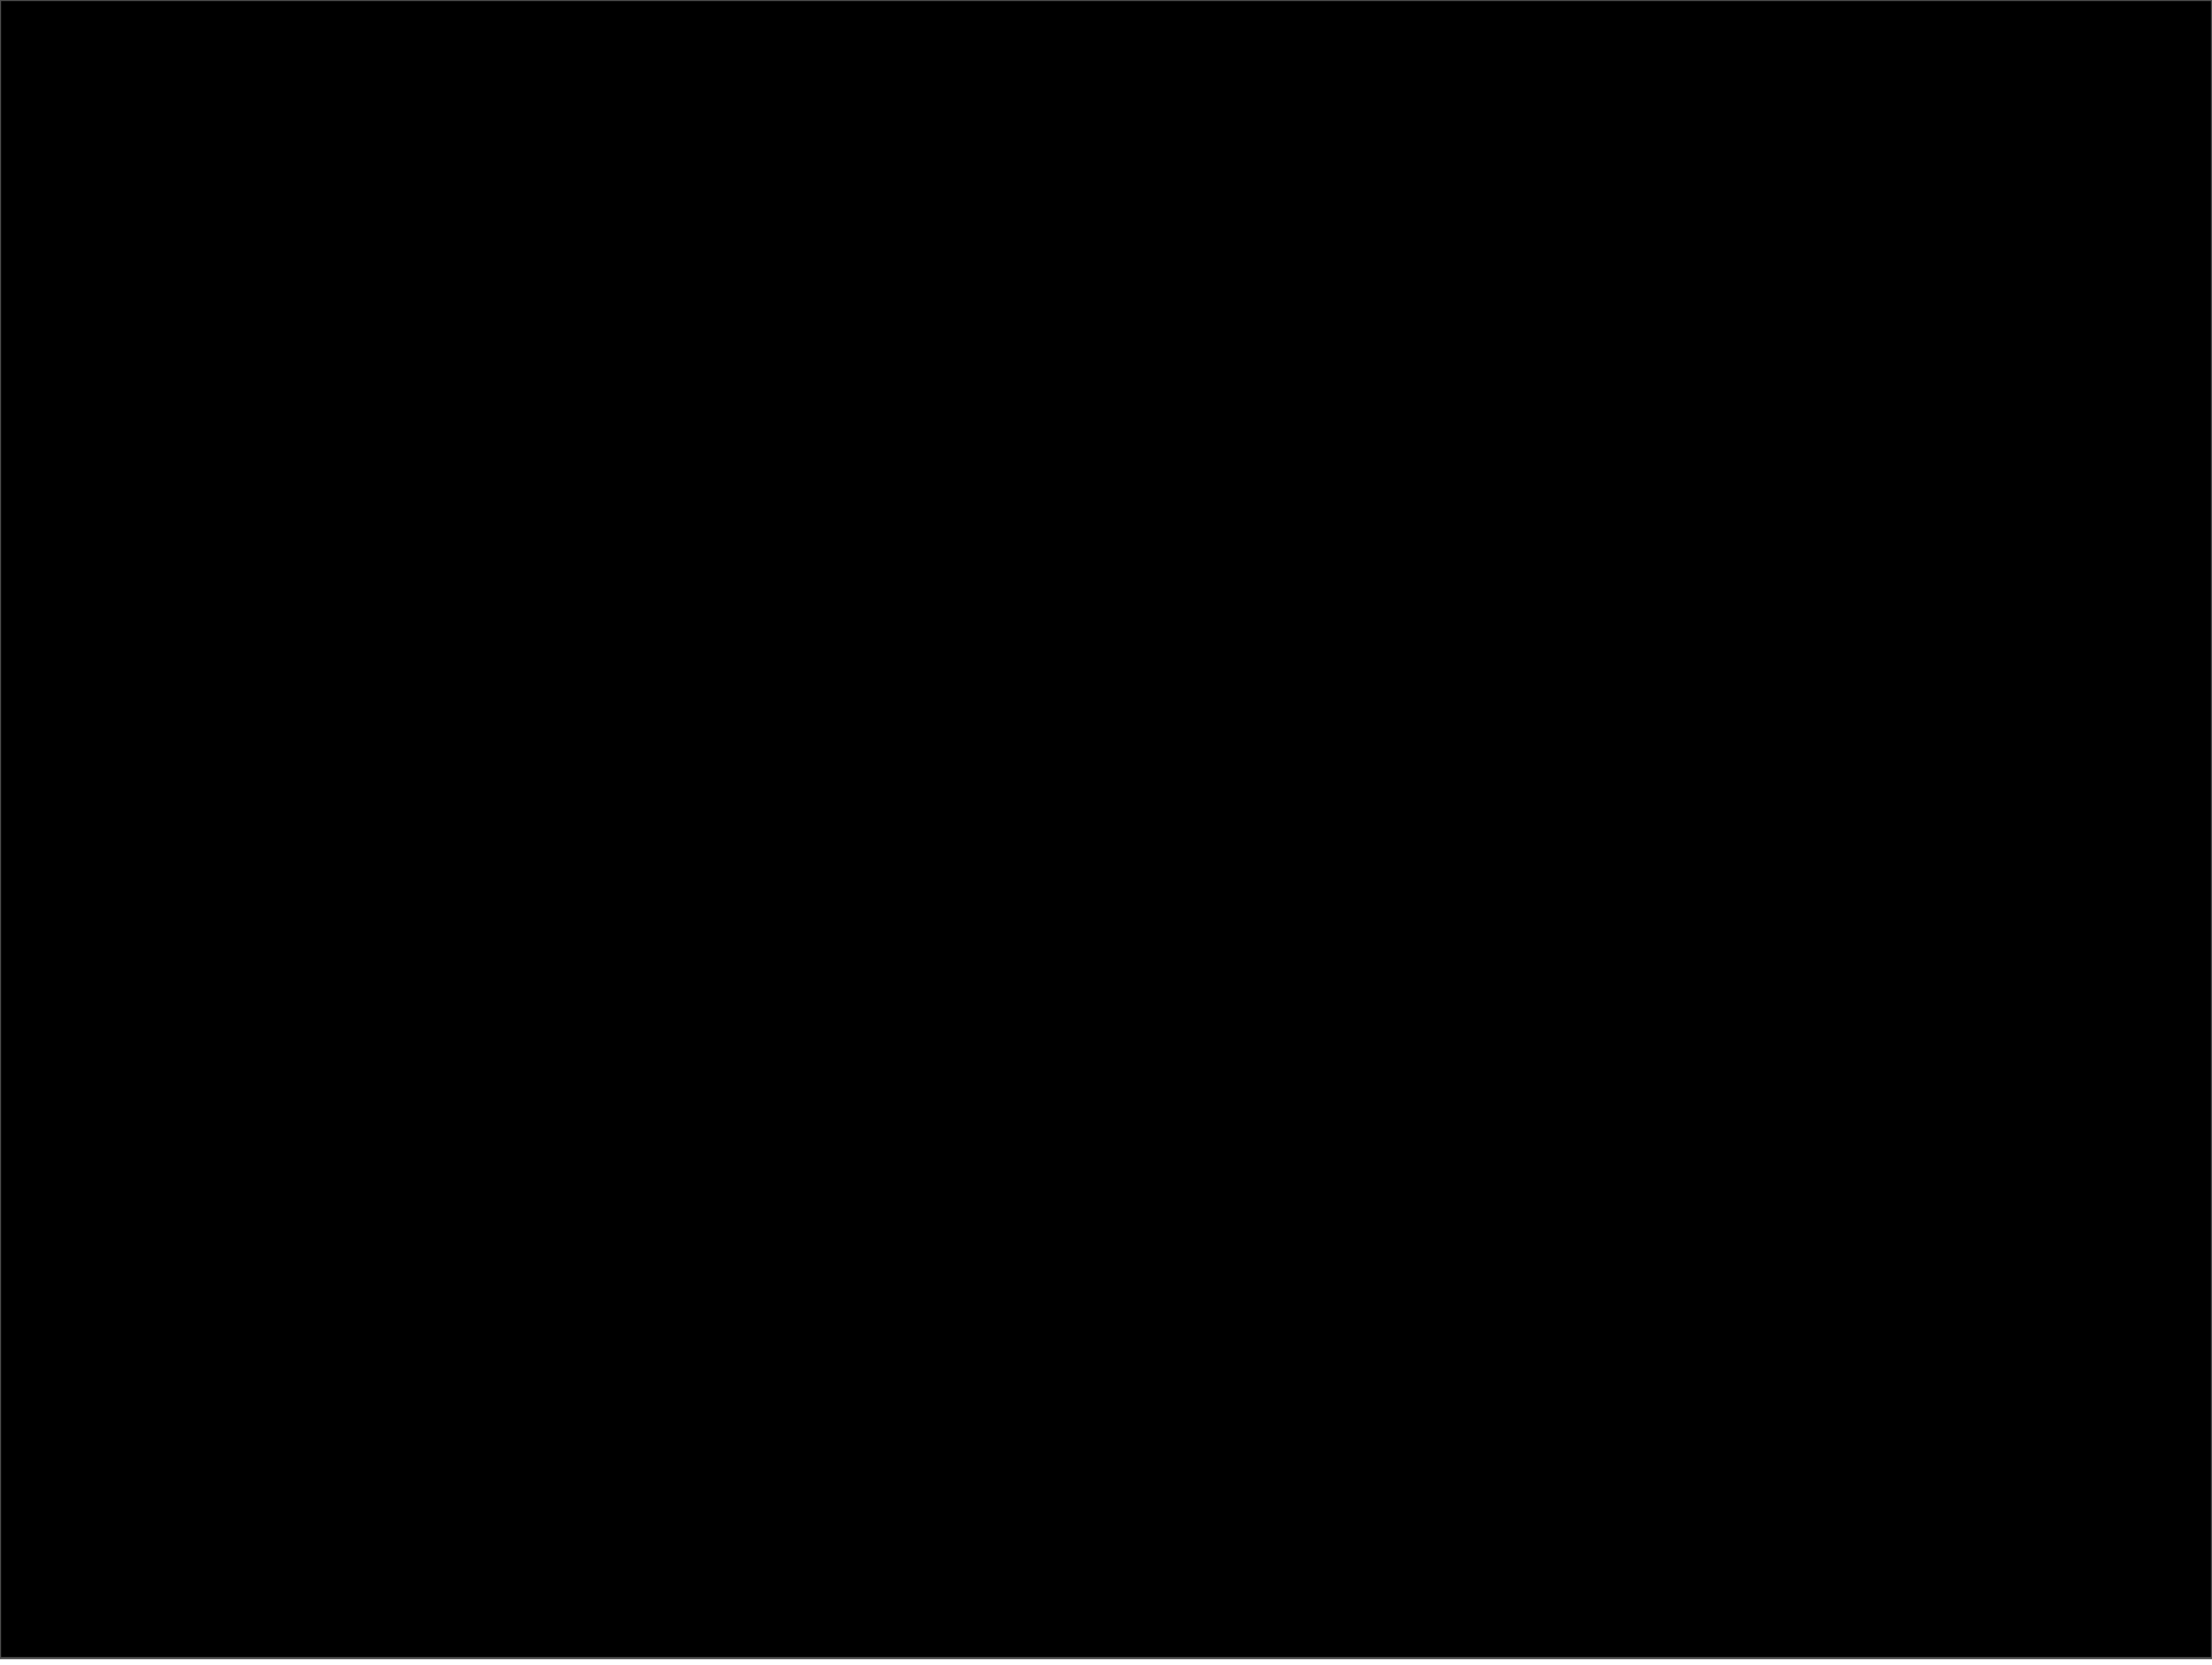

Delta Tracking vs. Ratio Tracking

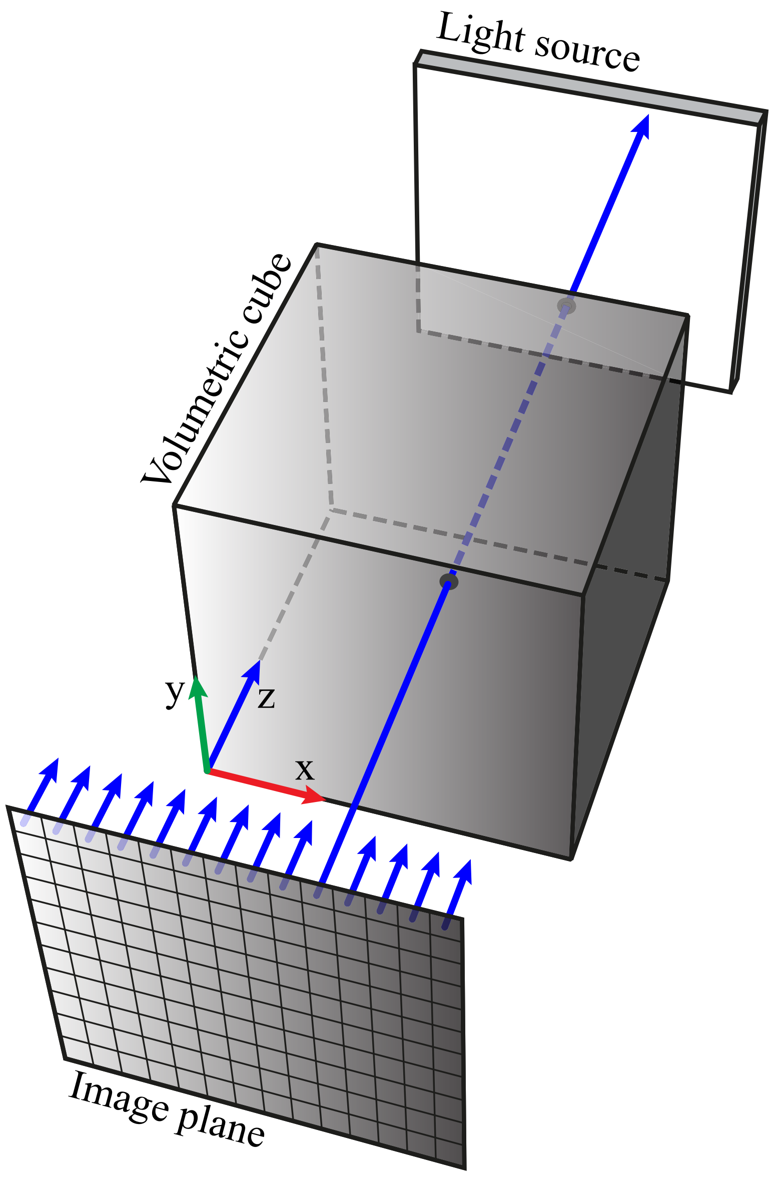

Delta Tracking vs. Ratio Tracking

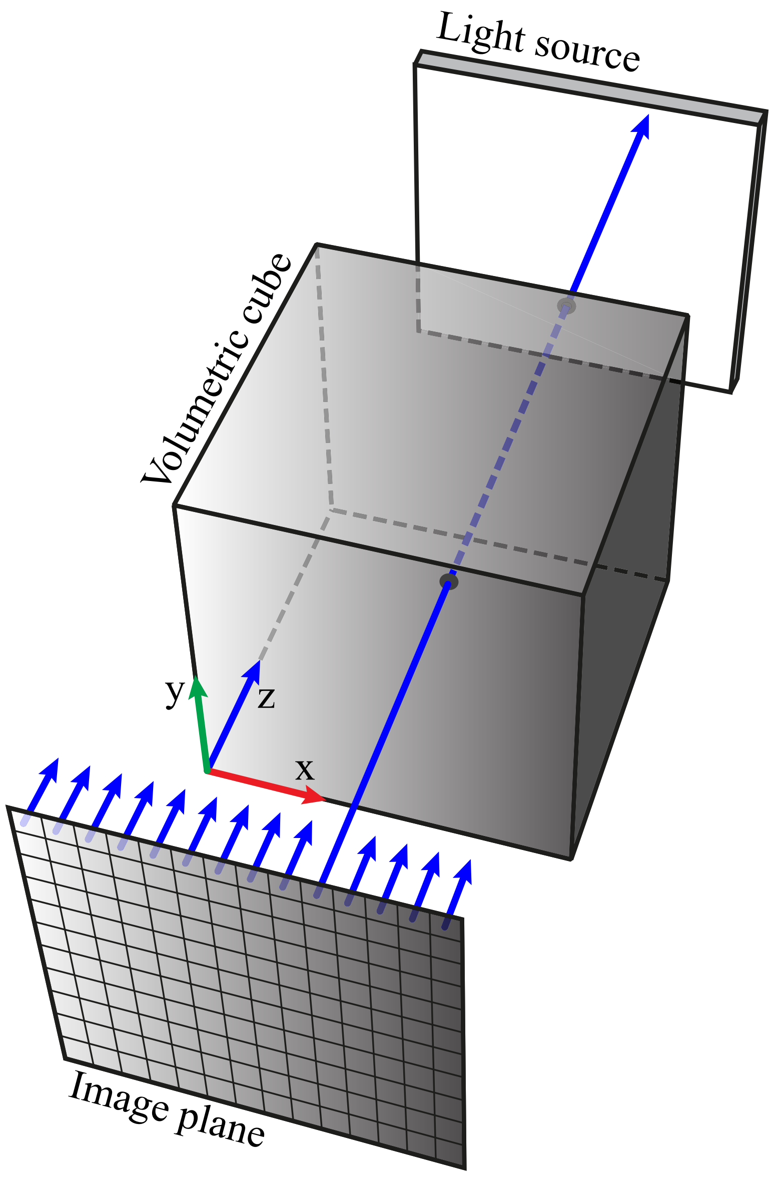

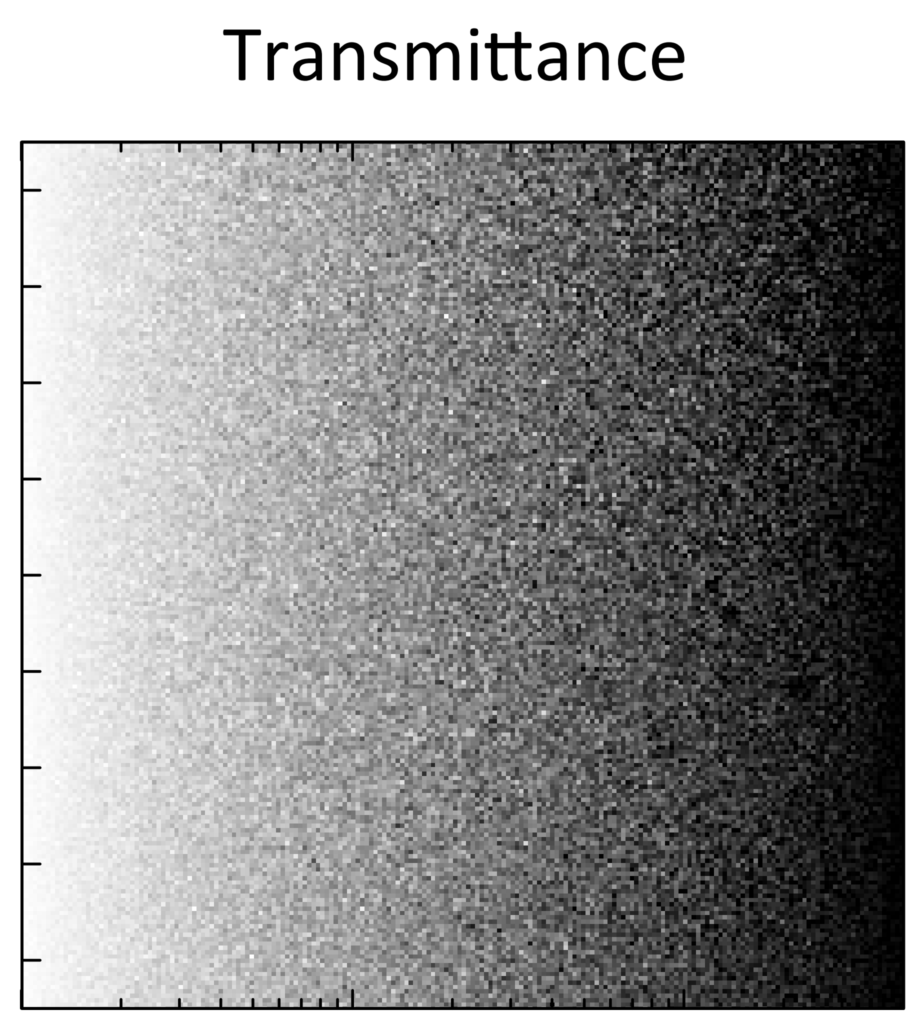

Delta Tracking vs. Ratio Tracking

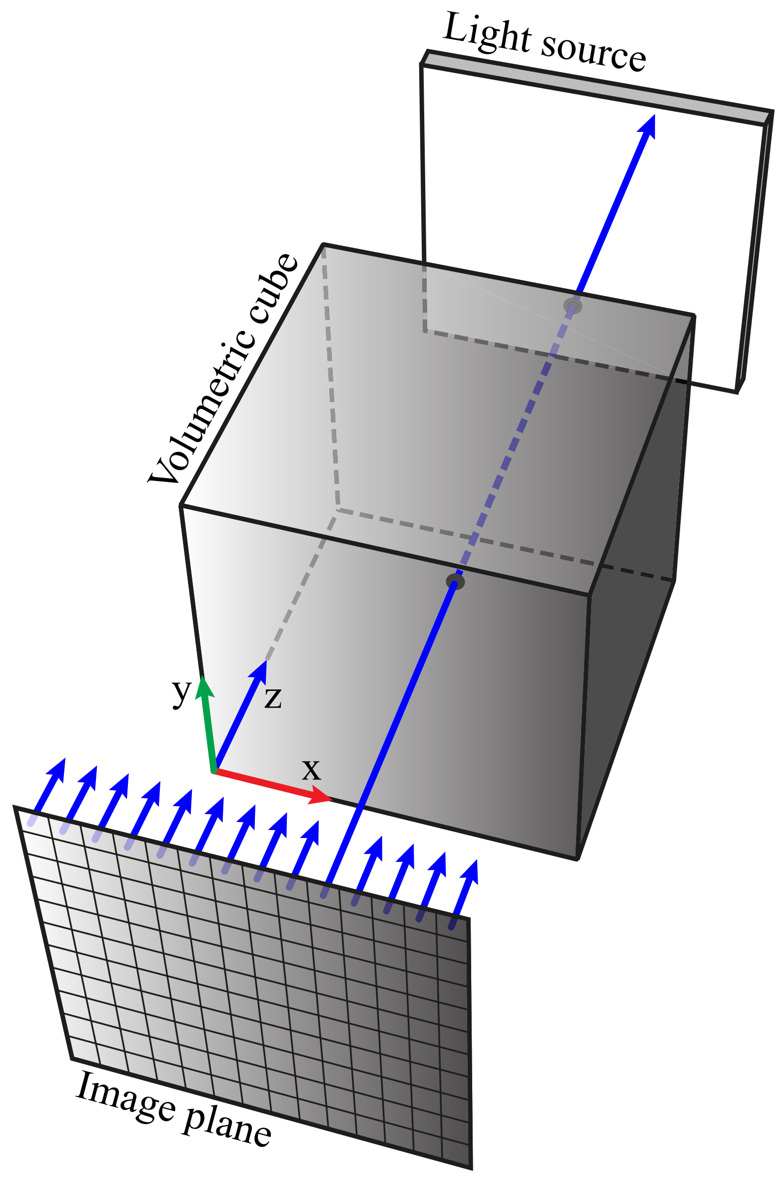

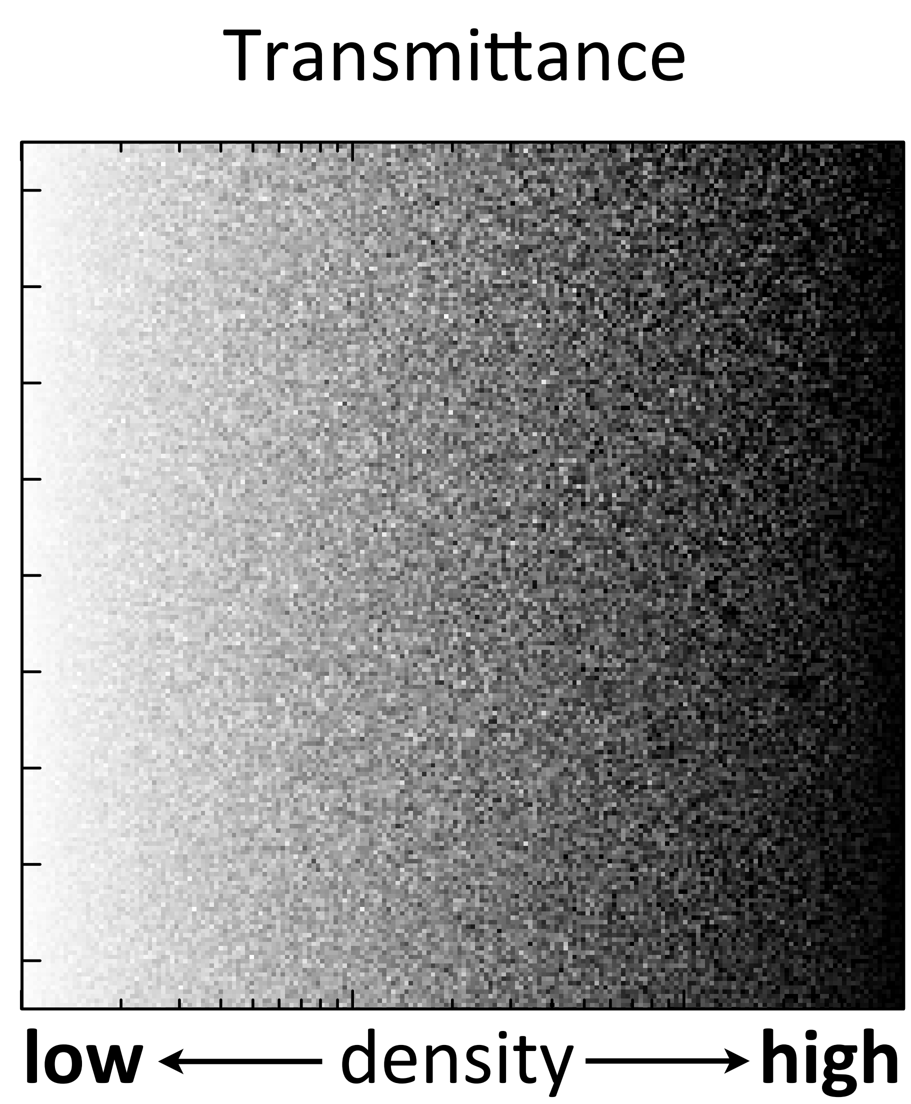

Delta Tracking vs. Ratio Tracking

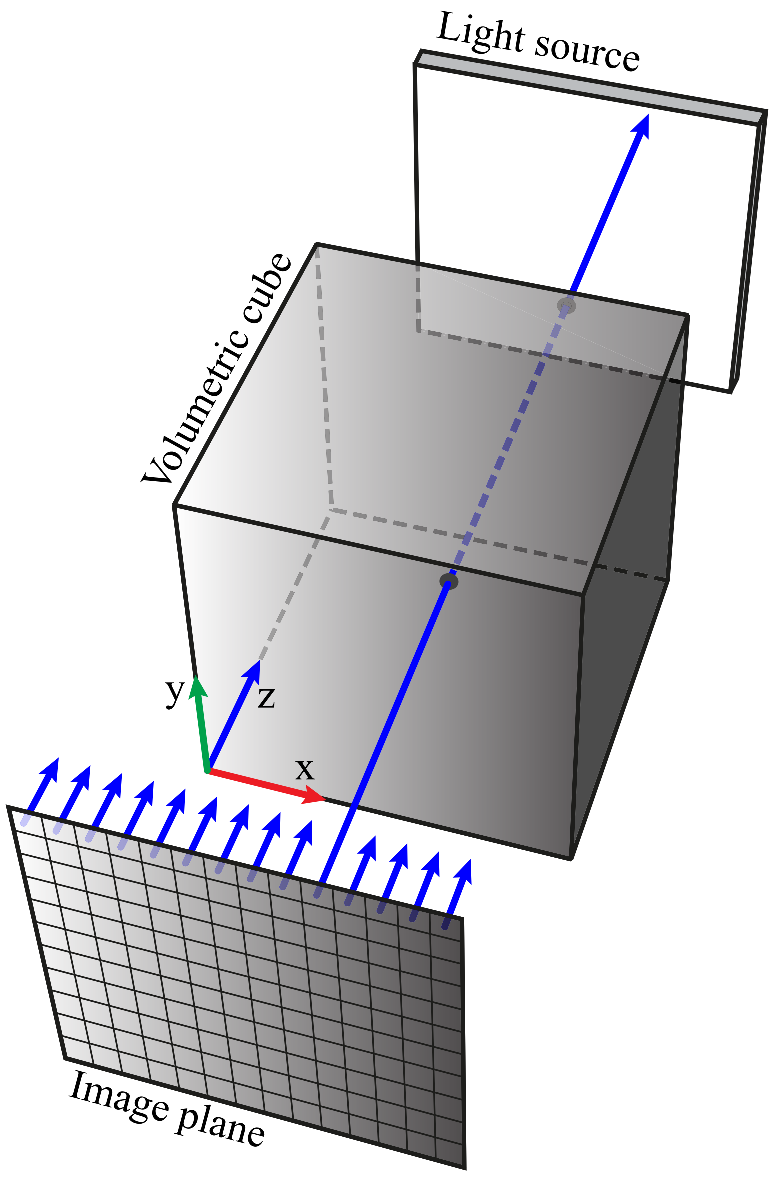

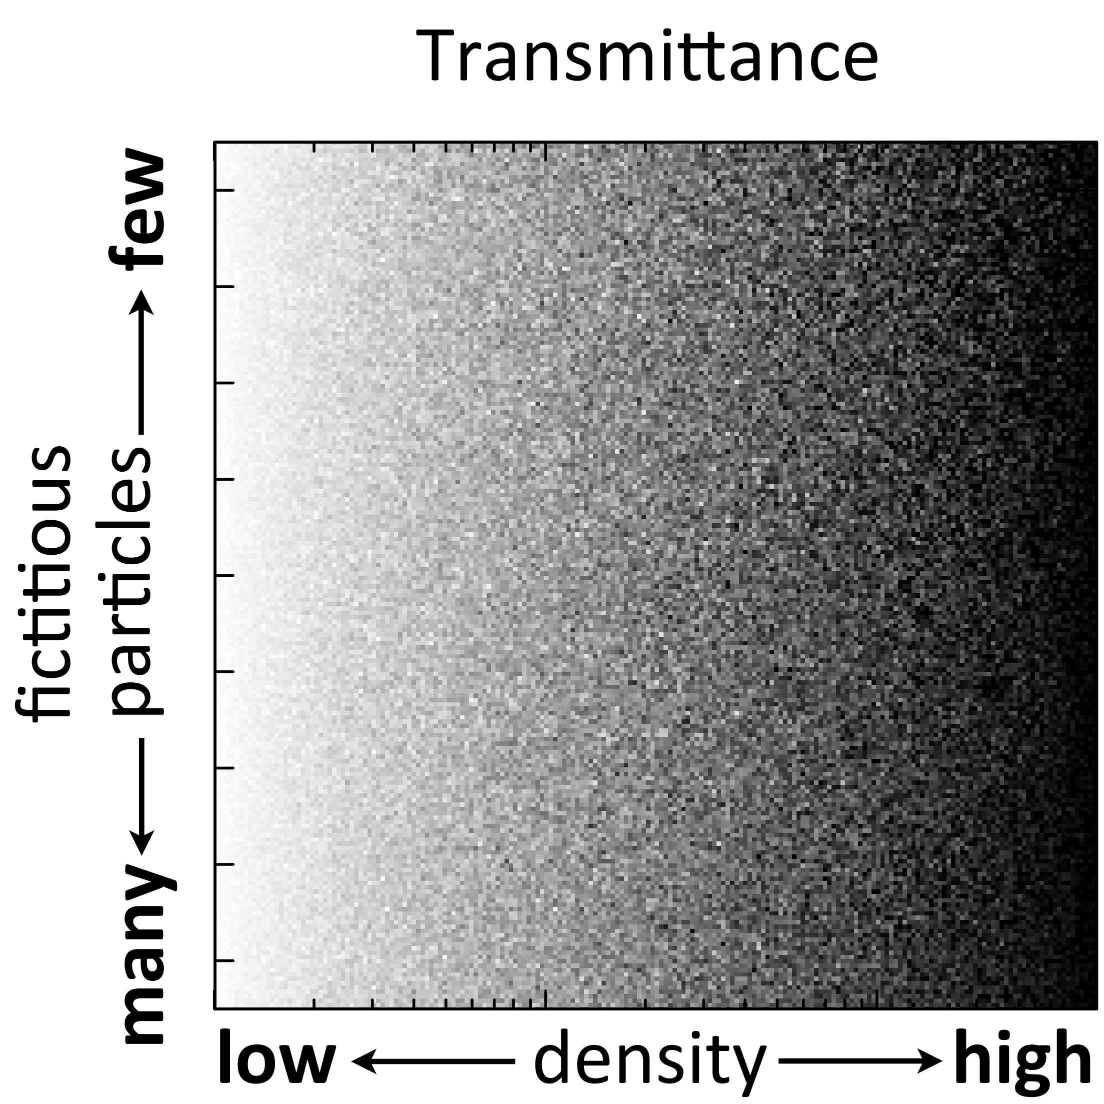

Delta Tracking vs. Ratio Tracking

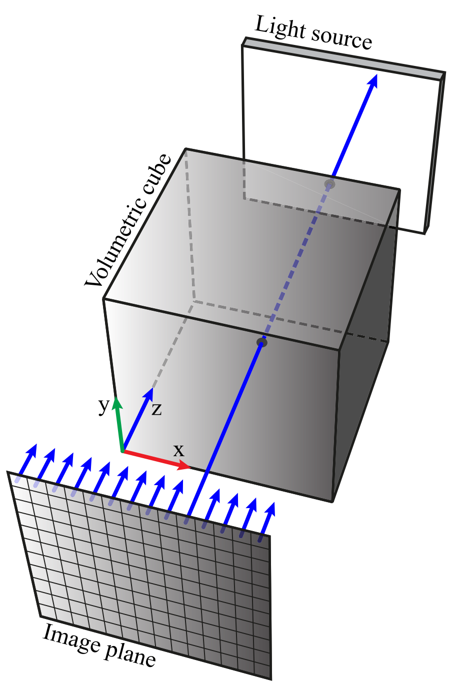

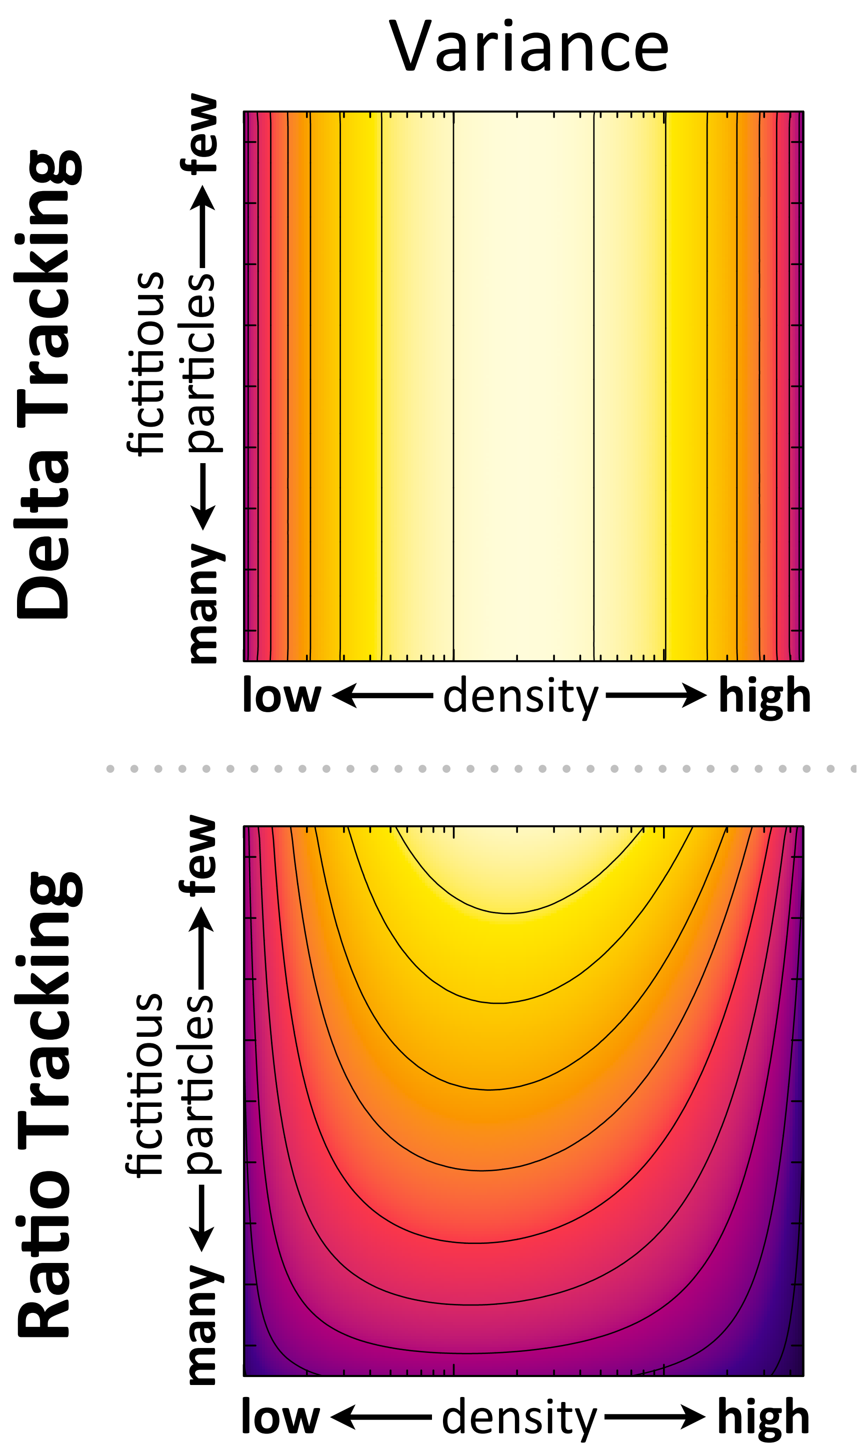

Delta Tracking vs. Ratio Tracking

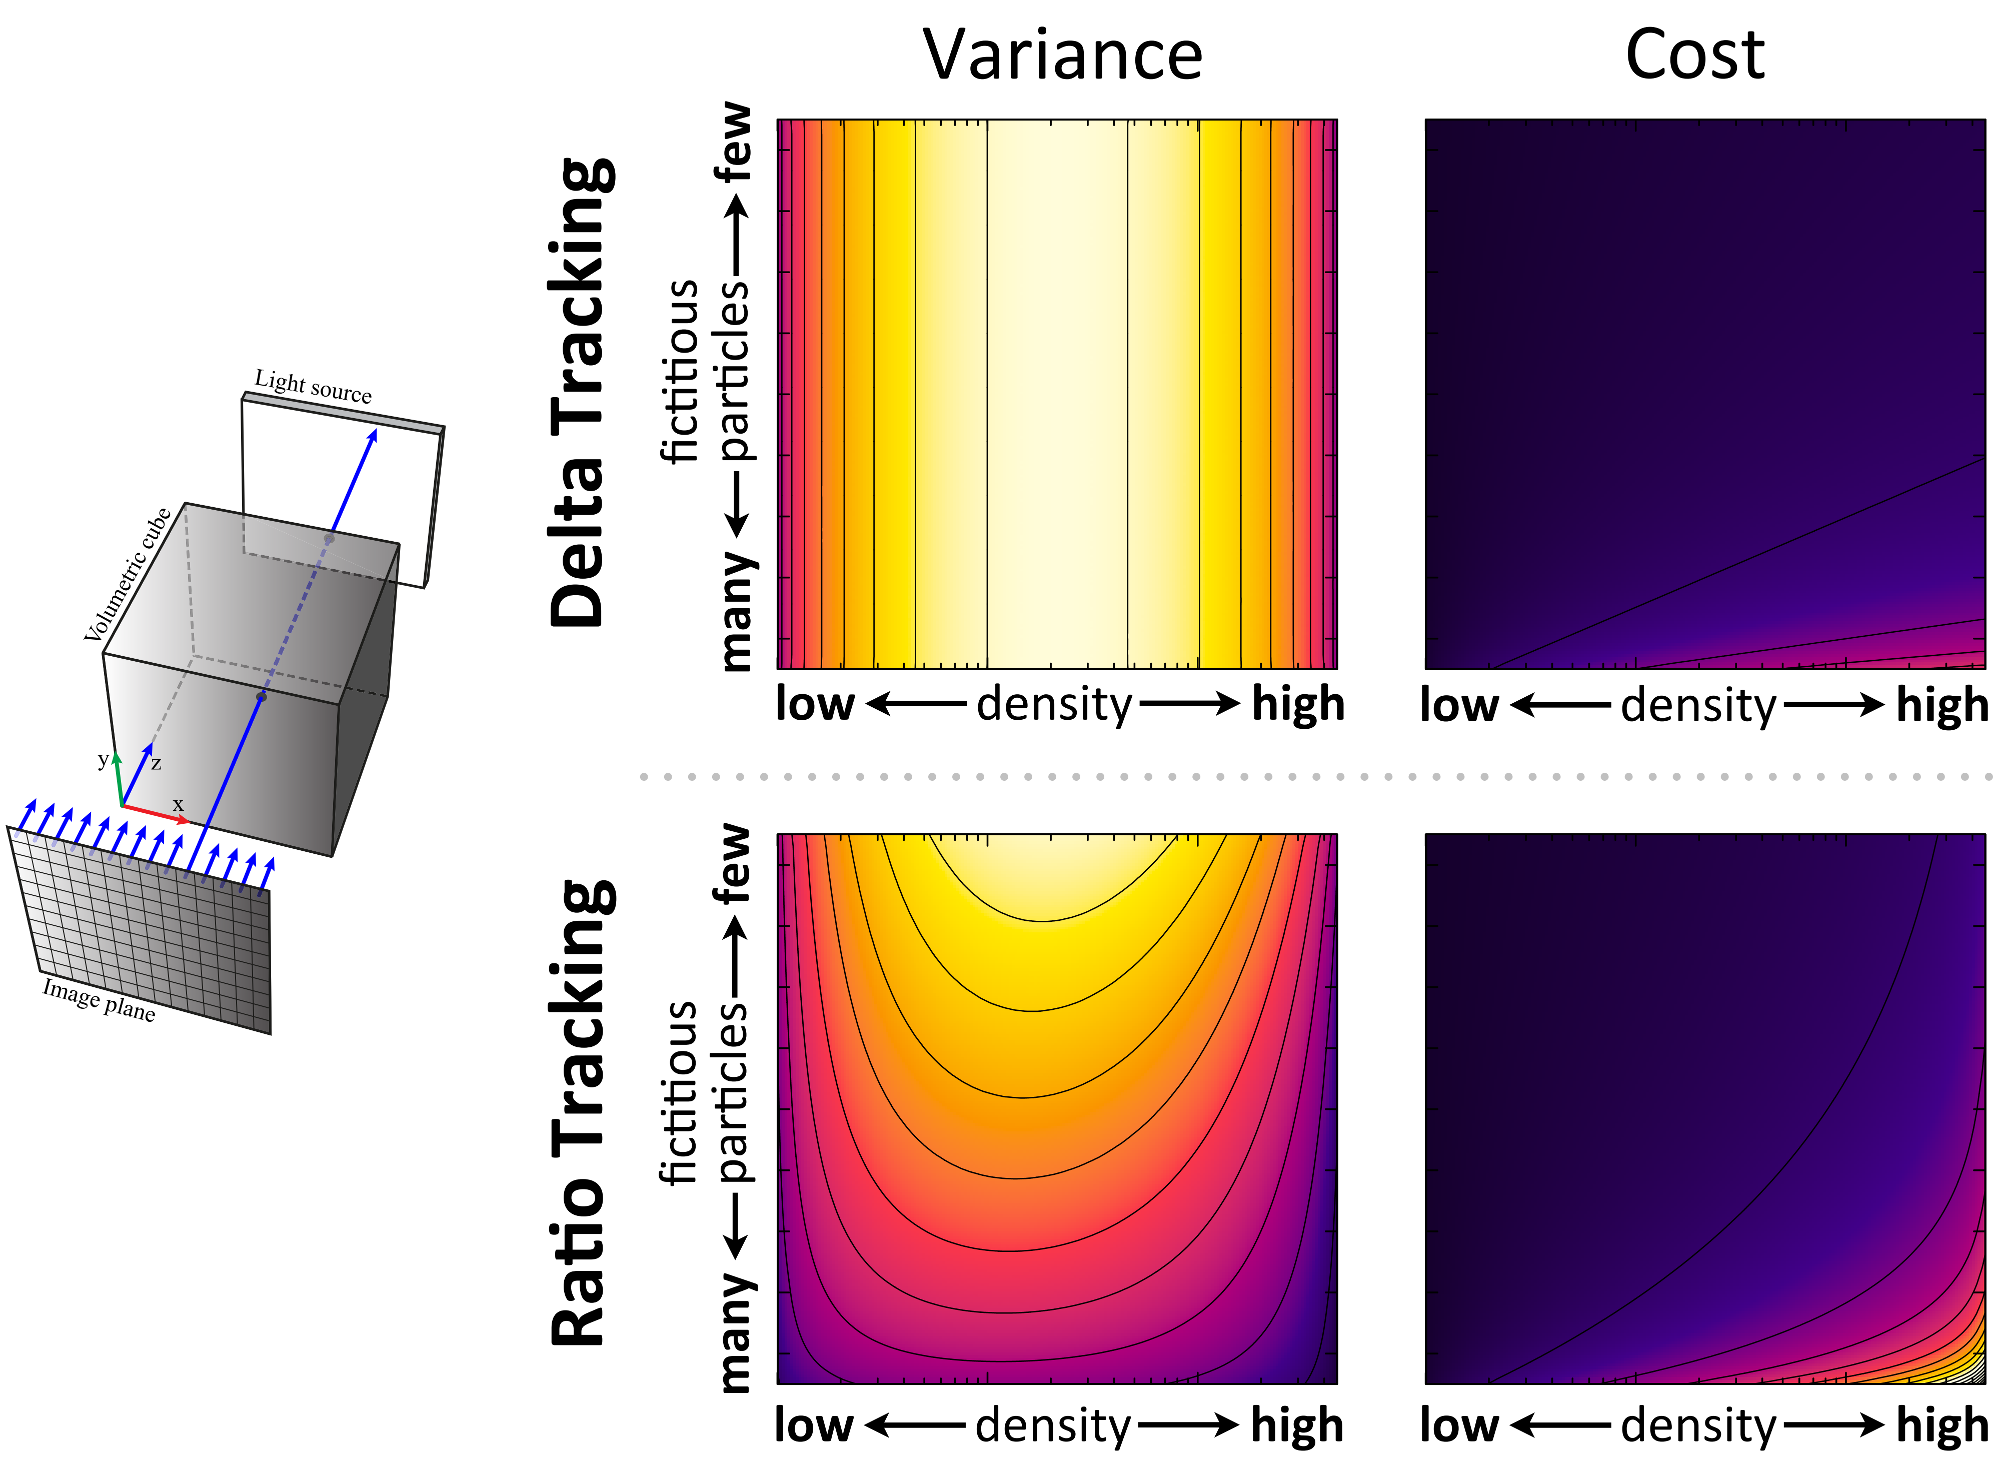

Delta Tracking vs. Ratio Tracking

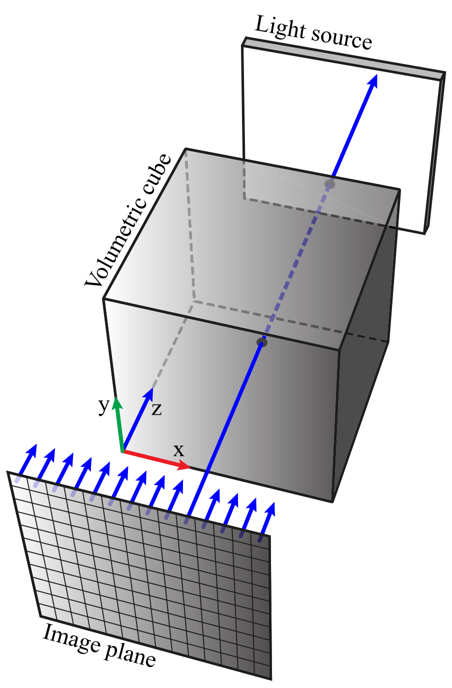

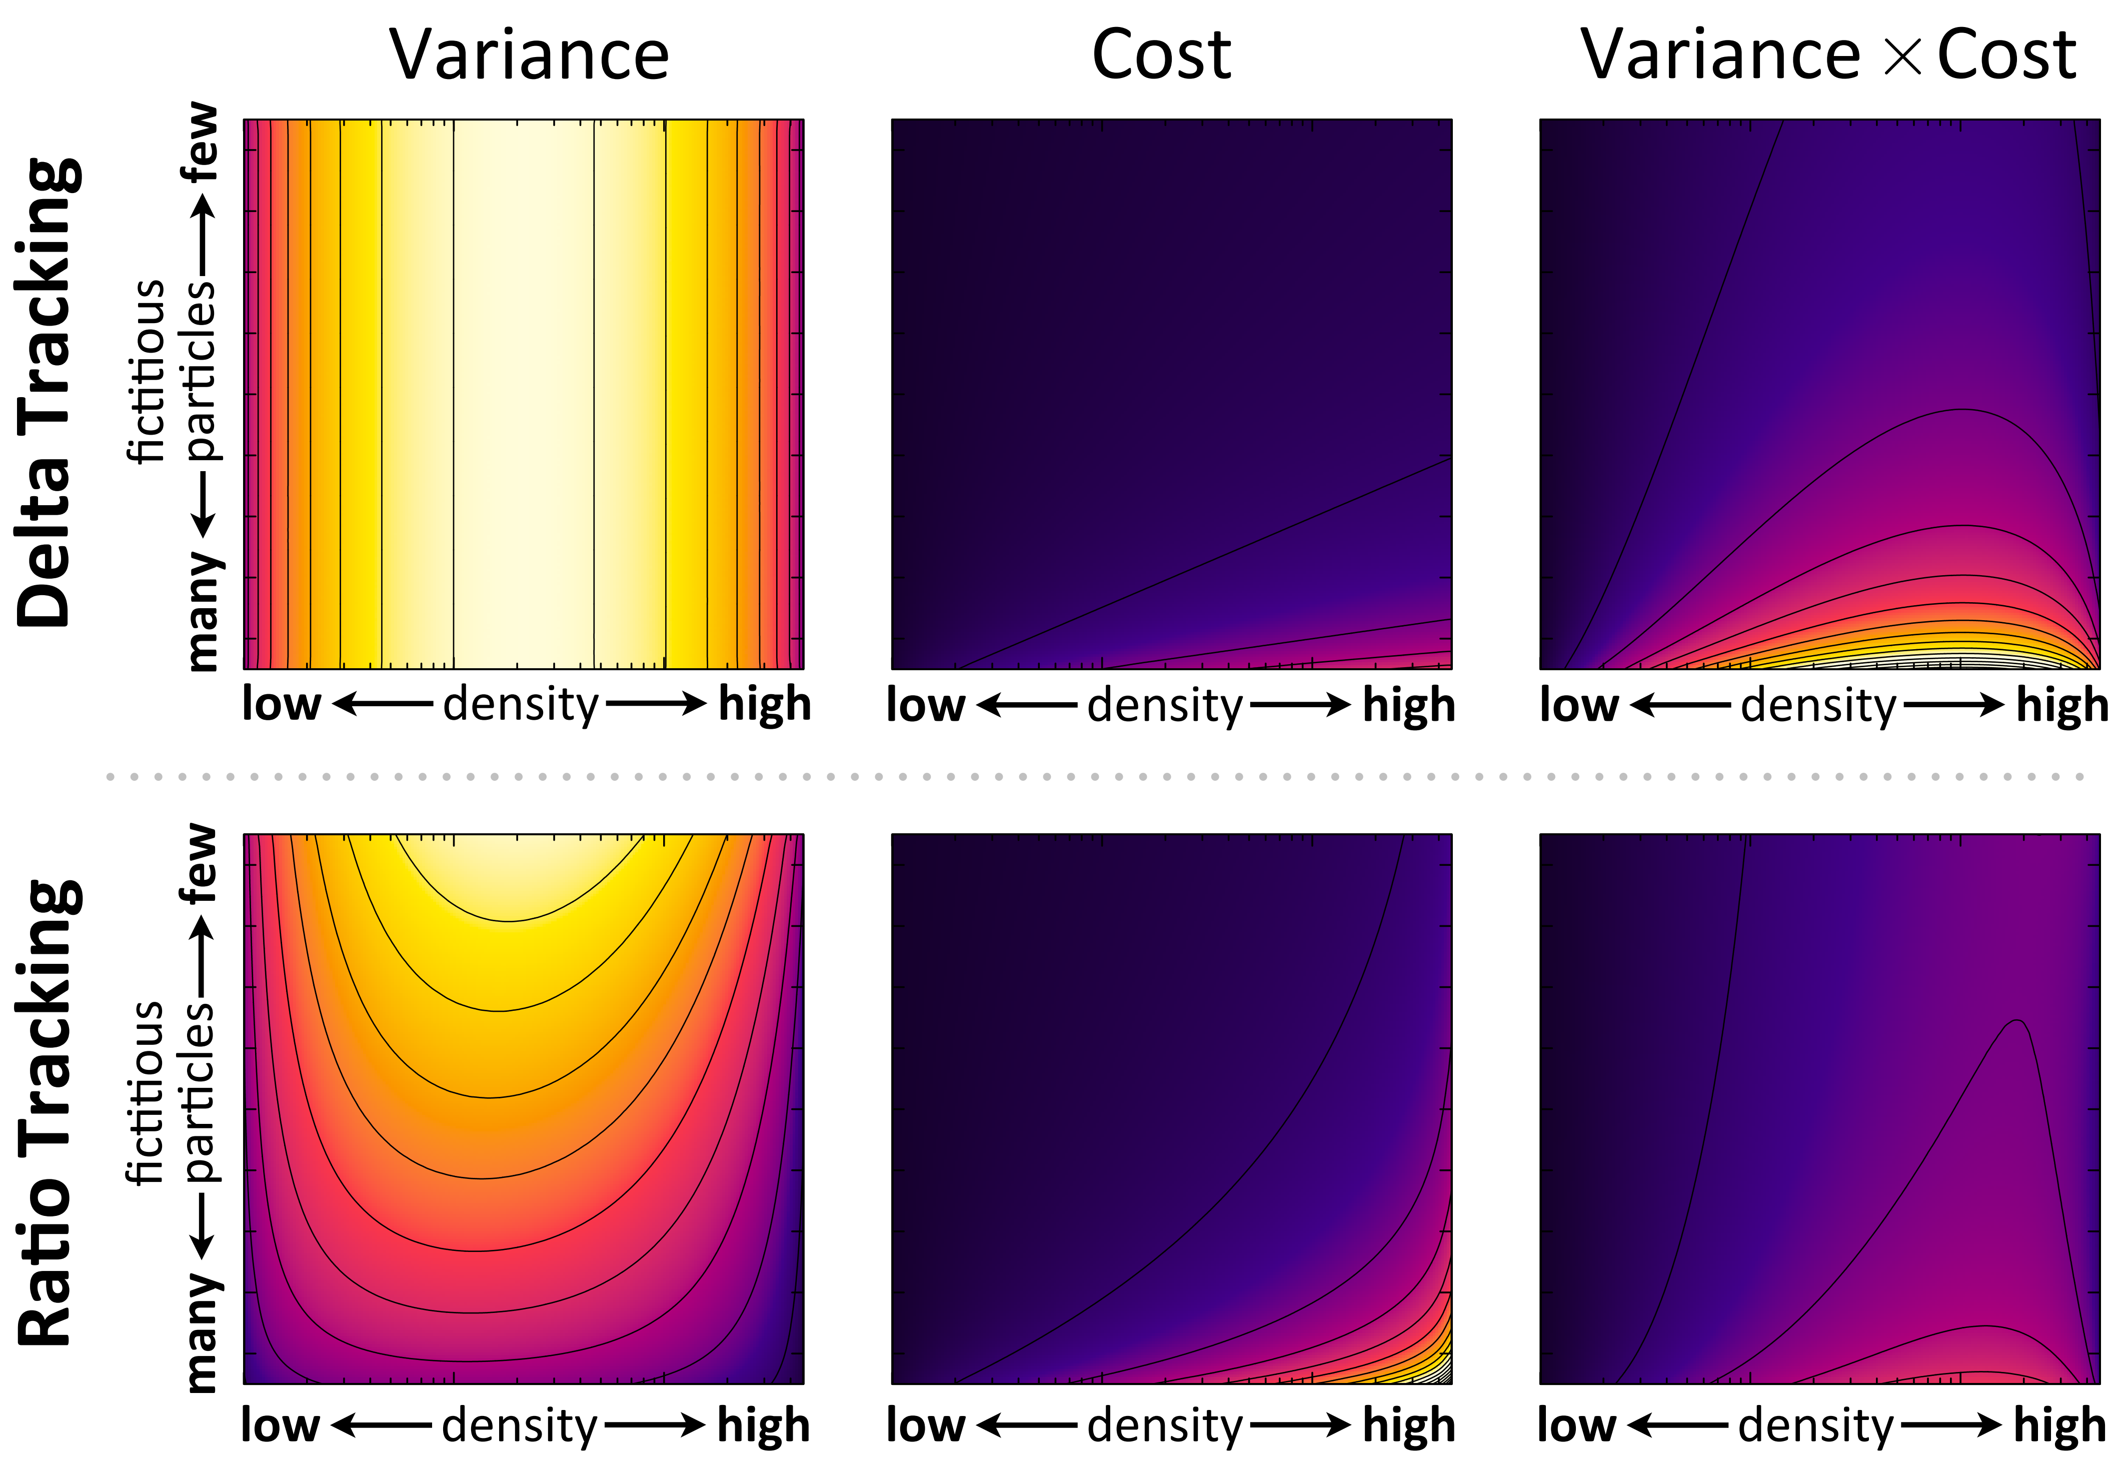

Delta Tracking vs. Ratio Tracking

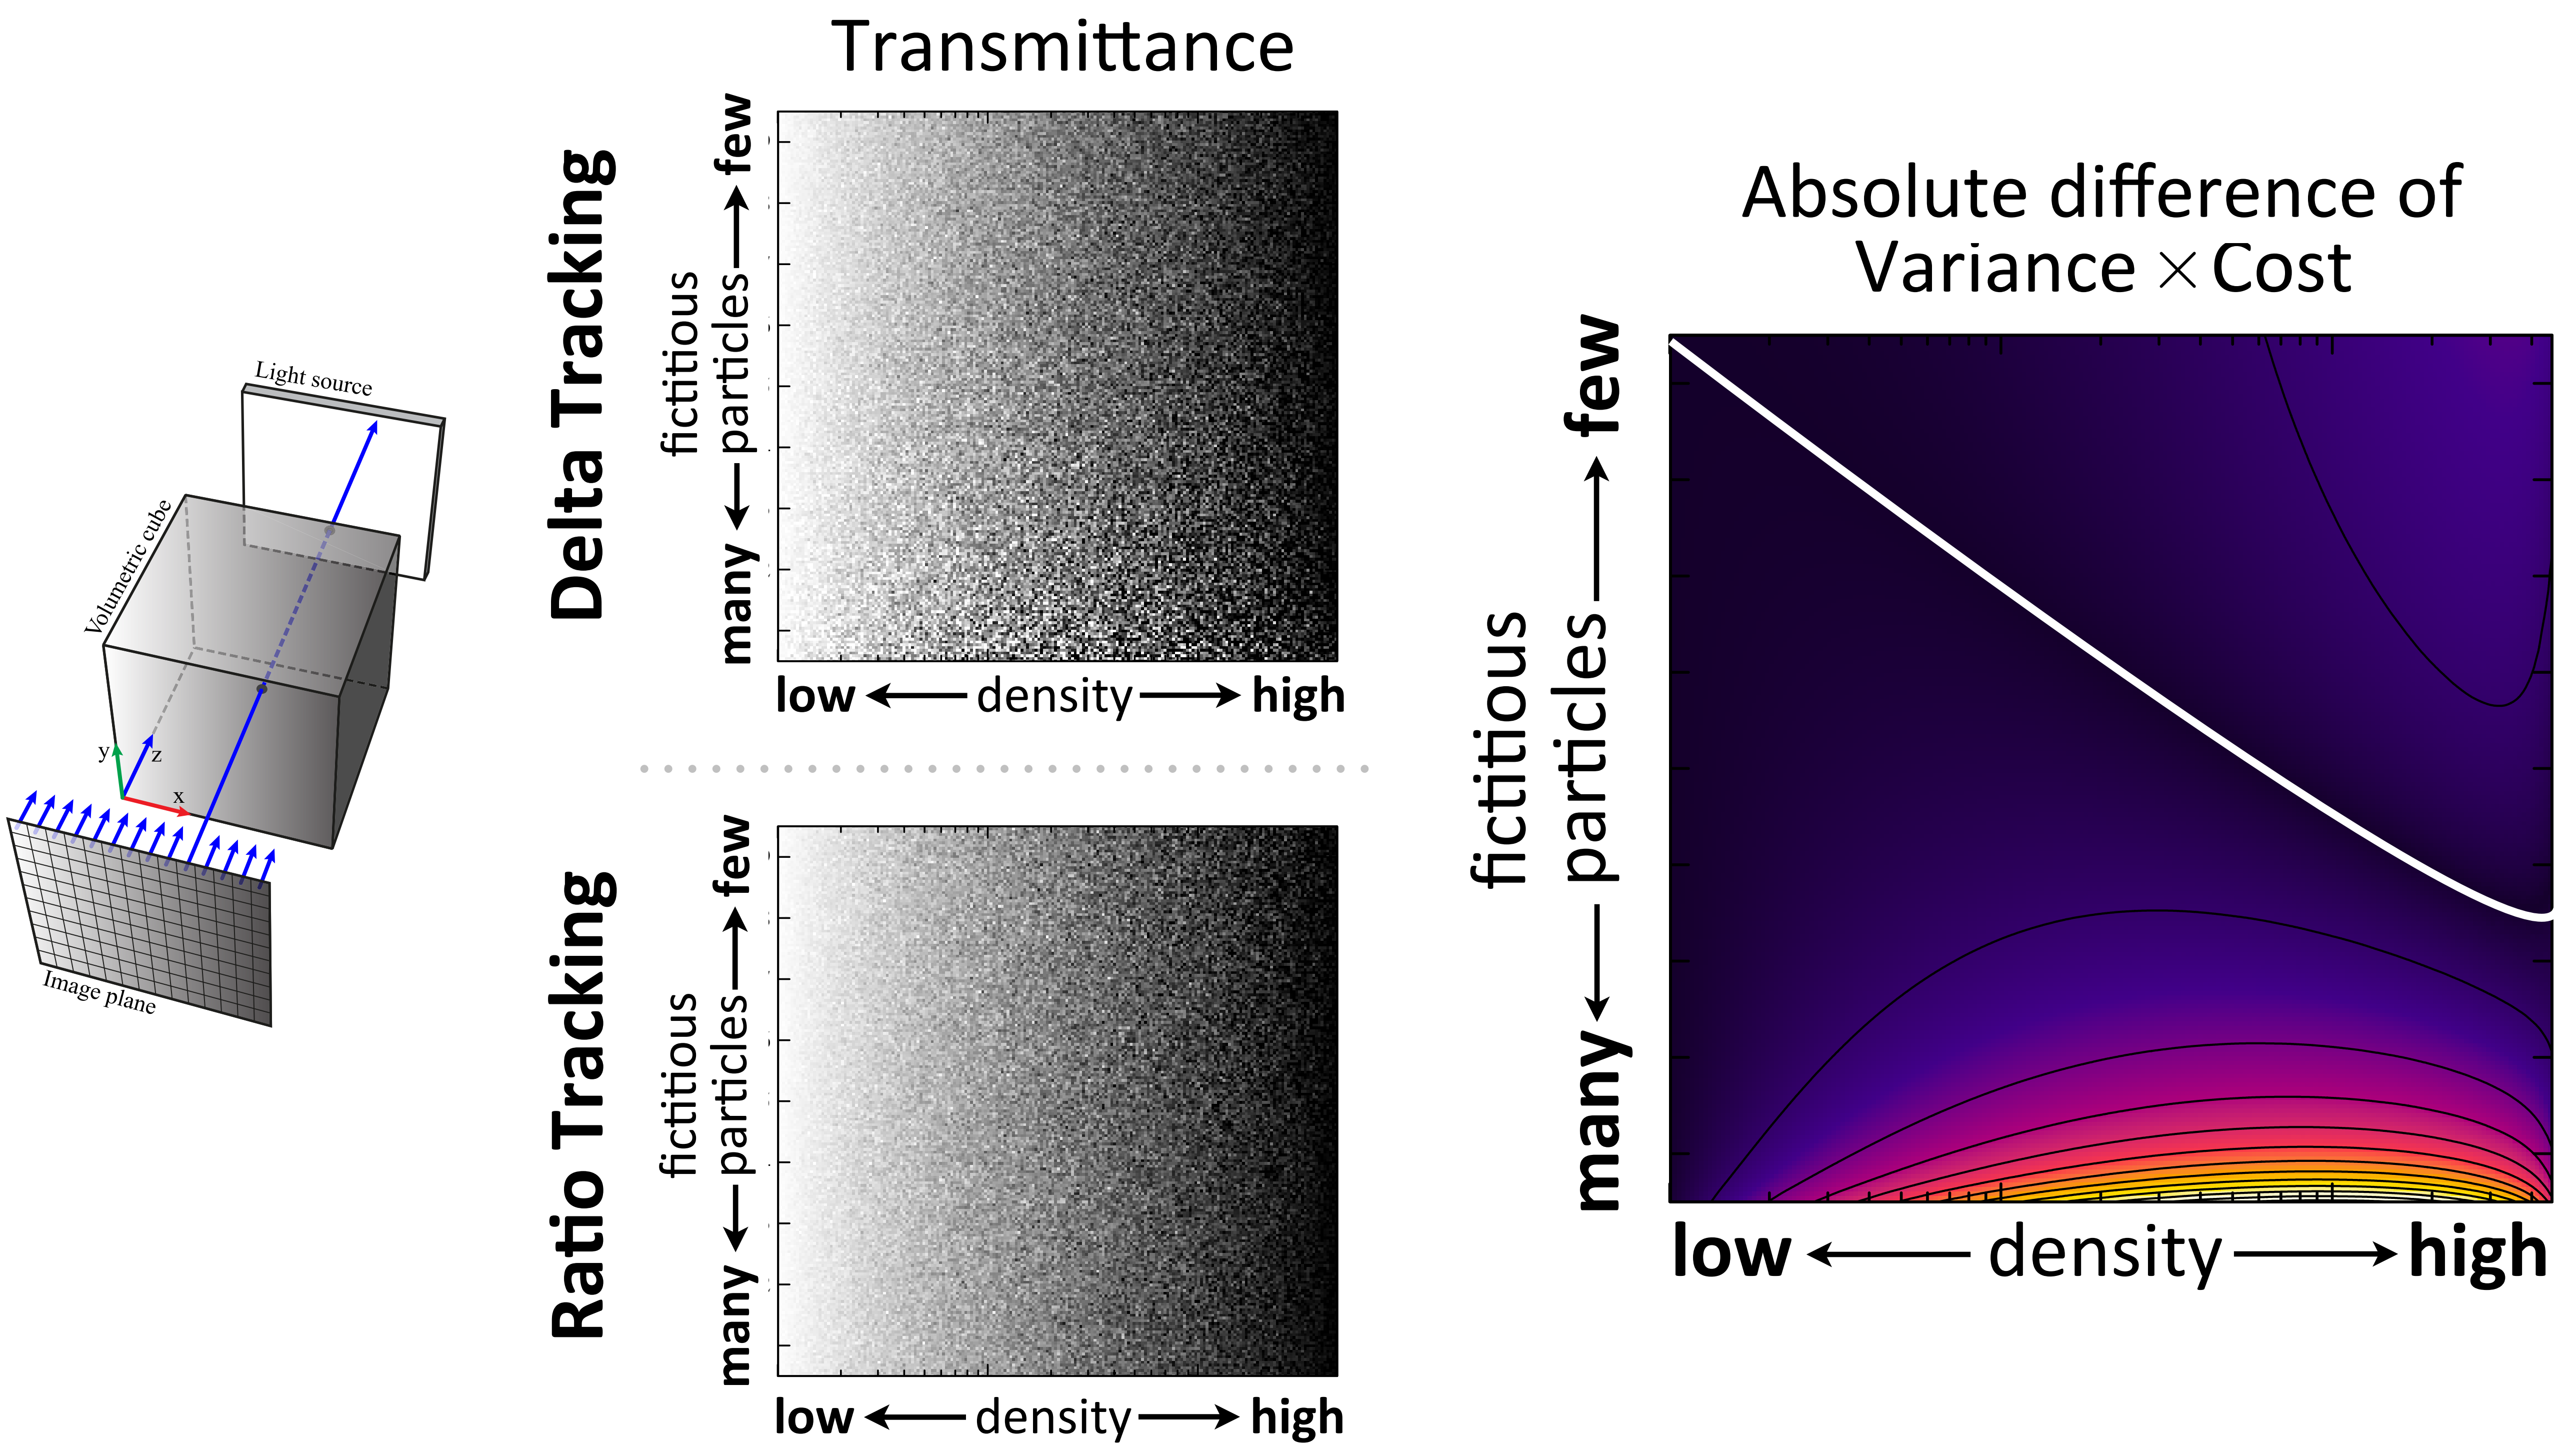

Delta Tracking vs. Ratio Tracking

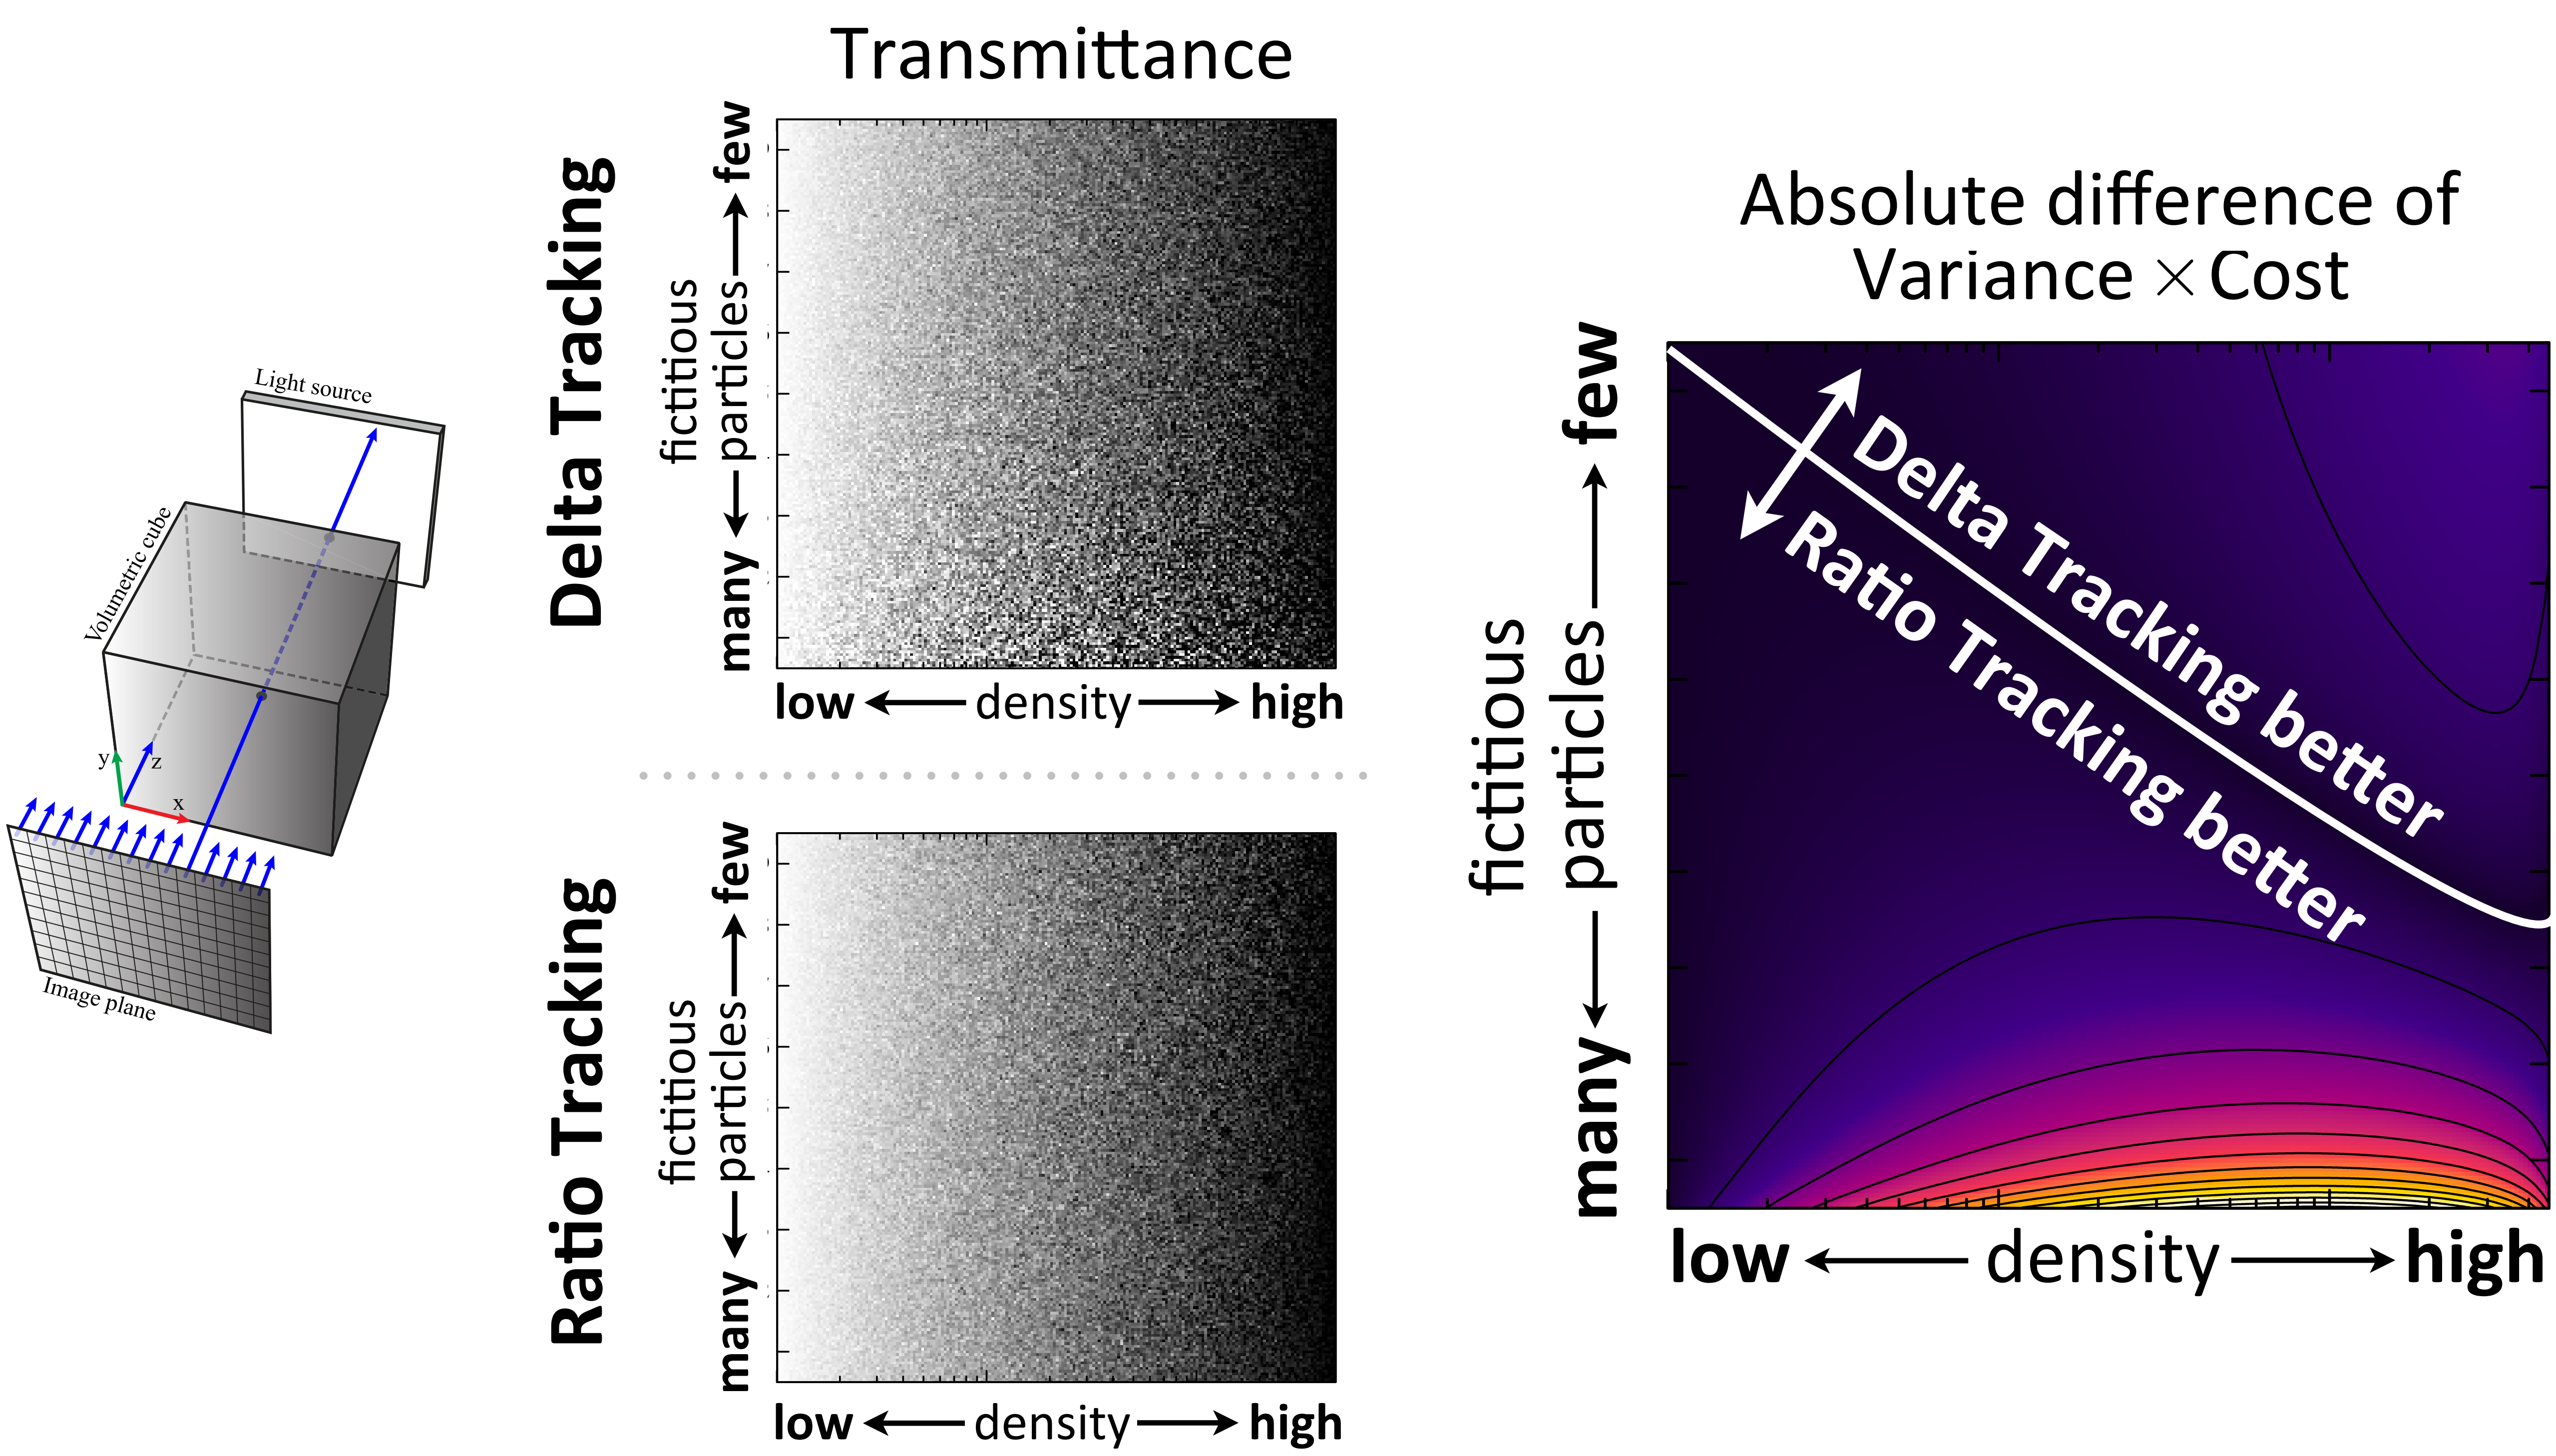

Delta Tracking vs. Ratio Tracking

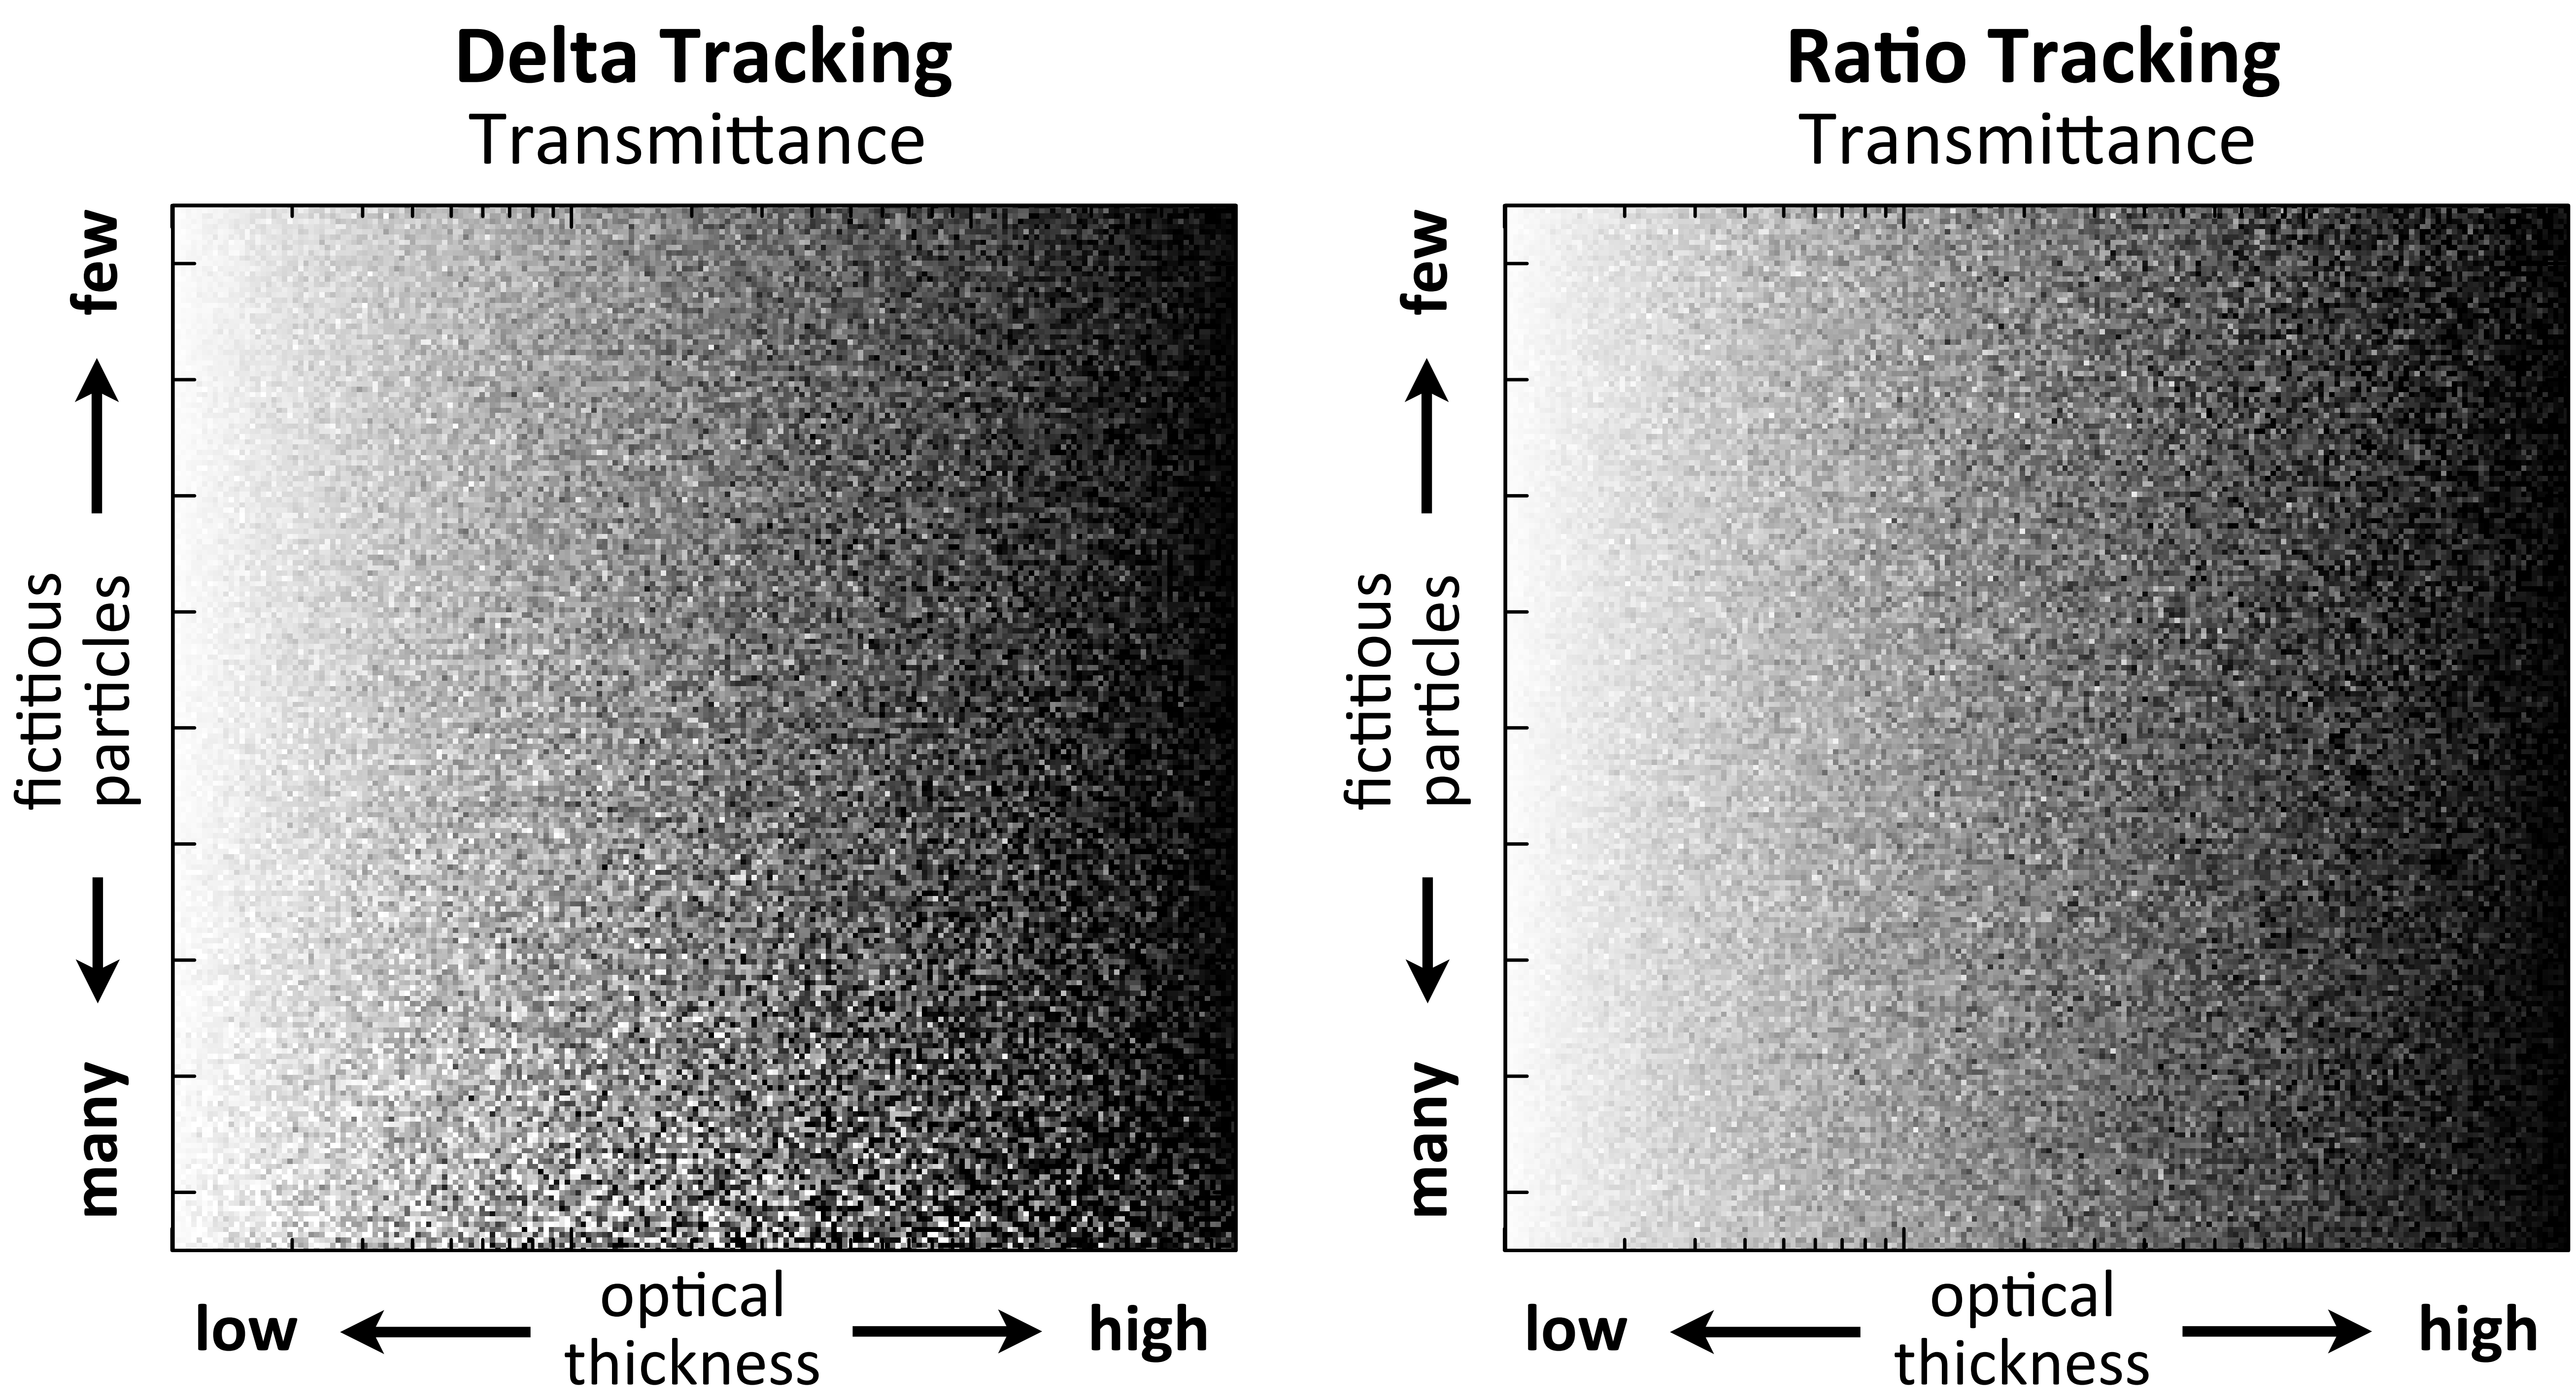

Residual Ratio Tracking with Super-Voxels