SIGGRAPH THINK BEYOND

2020 S2020.SIGGRAPH.ORG

Spatiotemporal Reservoir Resampling for Real-Time Ray Tracing with Dynamic Direct Lighting

Benedikt Bitterli[1] Chris Wyman[2] Matt Pharr[2]
Peter Shirley[2] Aaron Lefohn[2] Wojciech Jarosz[1]

[1]Dartmouth College [2]NVIDIA

Goals

Goals

- Focus on (but not limited to) *direct lighting*

Goals

- Focus on (but not limited to) *direct lighting*
- Goal:
  - 1’000’000+ dynamic area light sources
  - Correct shadowing and BRDF
  - Large scenes with complex visibility

Direct Lighting is hard

Bayesian online regression for adaptive direct illumination sampling,
Vévoda et al. *2018*

The Design and Evolution of Disney's Hyperion Renderer,
Burley et al. *2018*

Importance Sampling of Many Lights with Adaptive Tree Splitting,
Estevez and Kulla *2018*

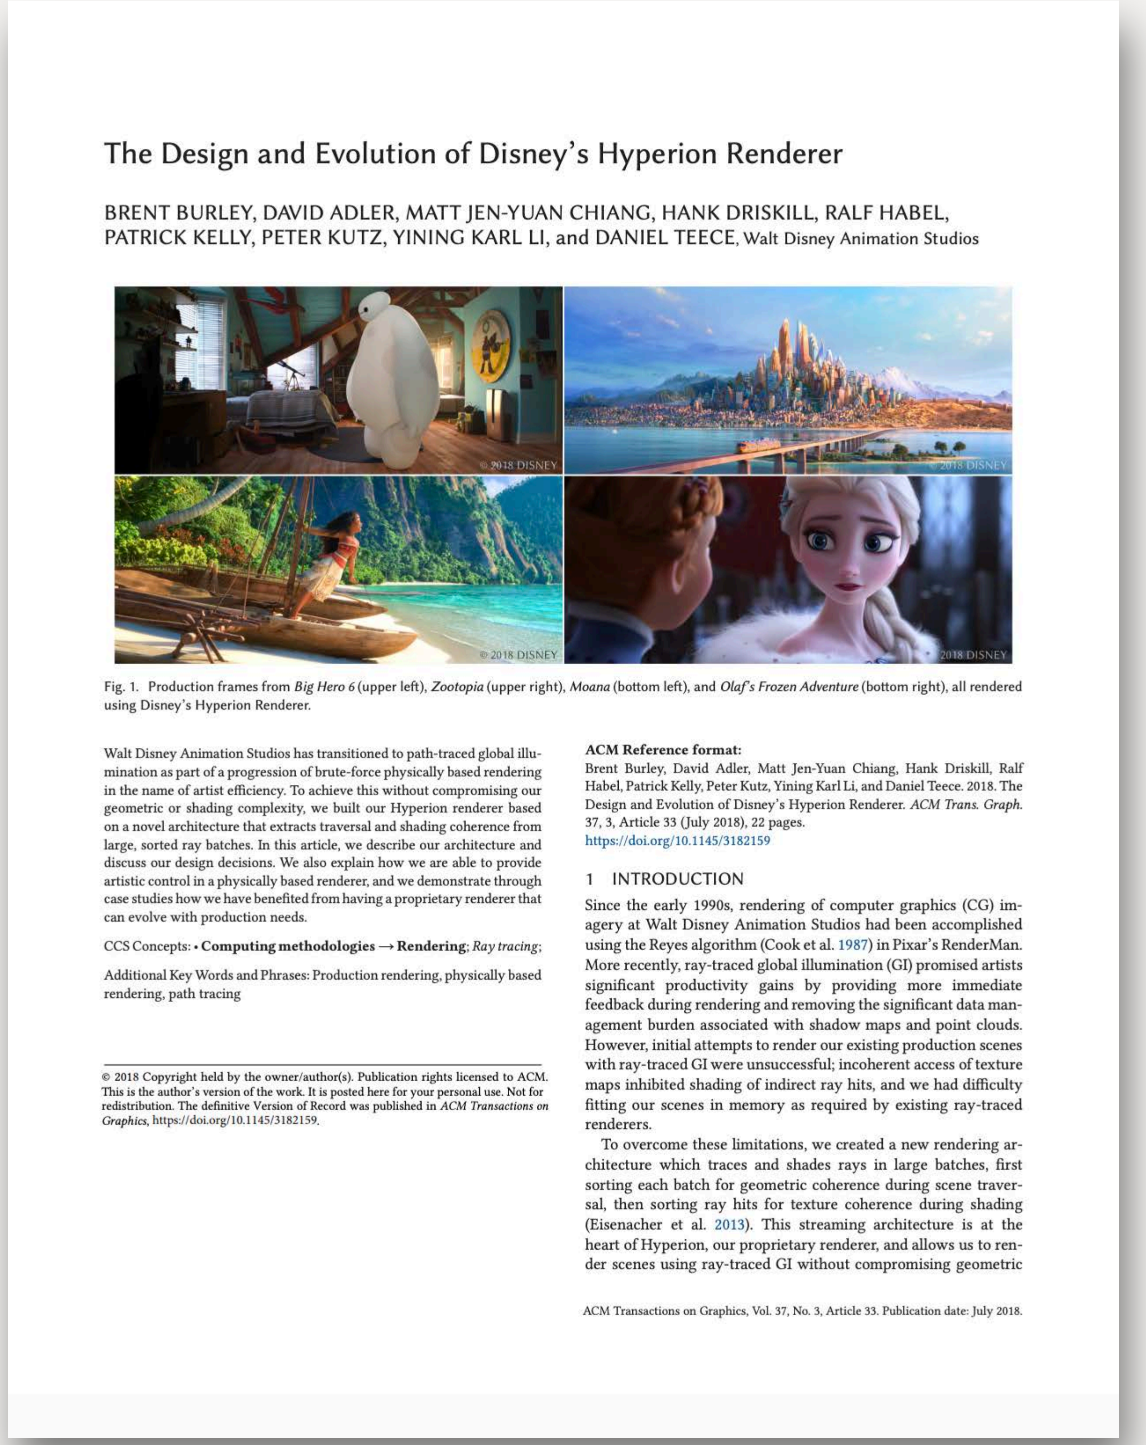

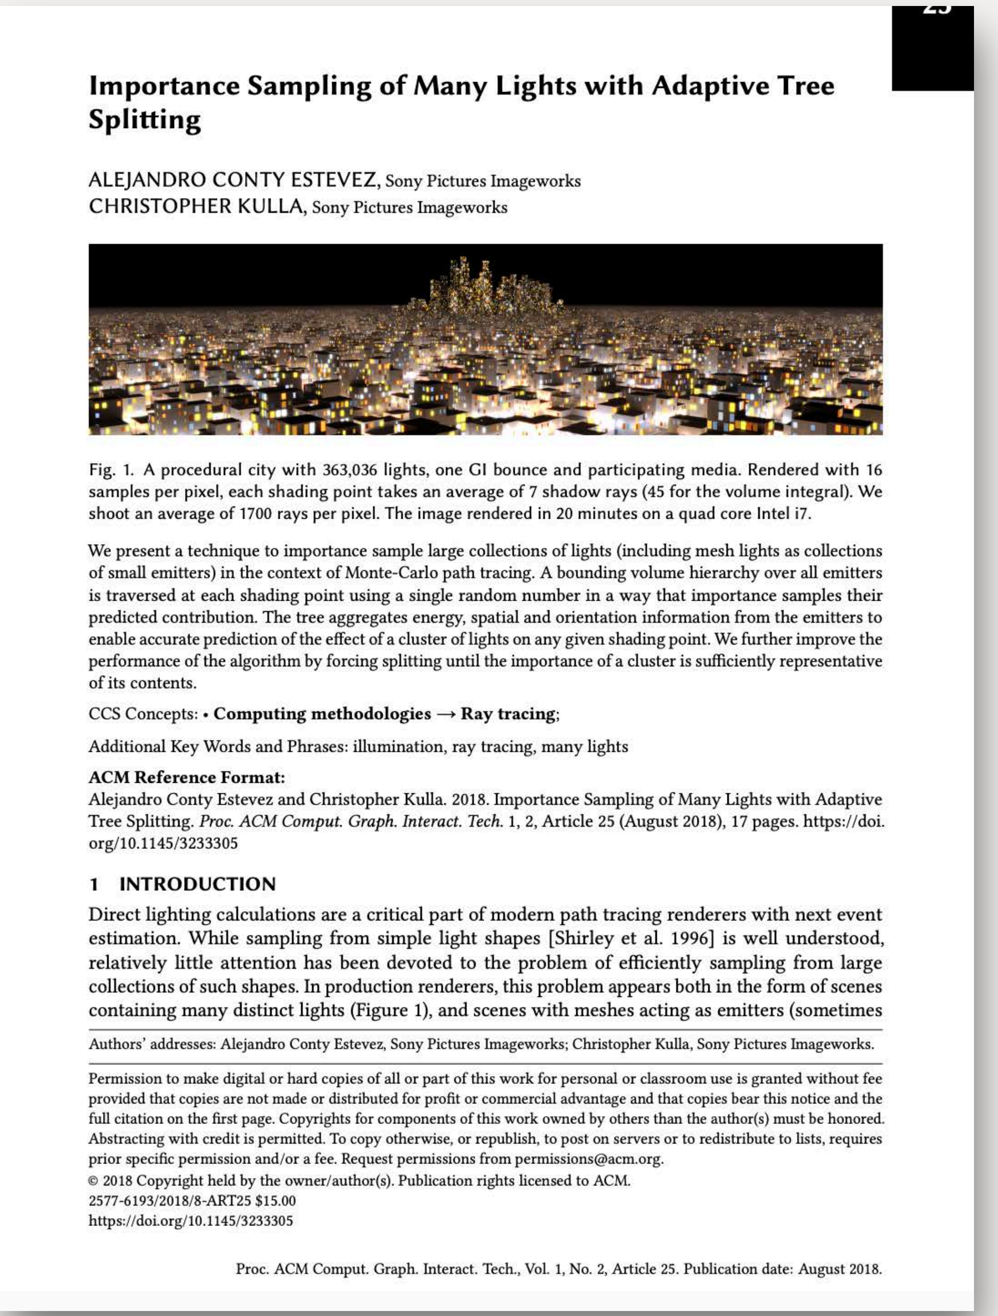

Real-Time is harder

Goals

- Good quality at low sample count ($\leq$4 spp)
- *Everything* is dynamic
- Massively parallel

Problem Description

Problem Description

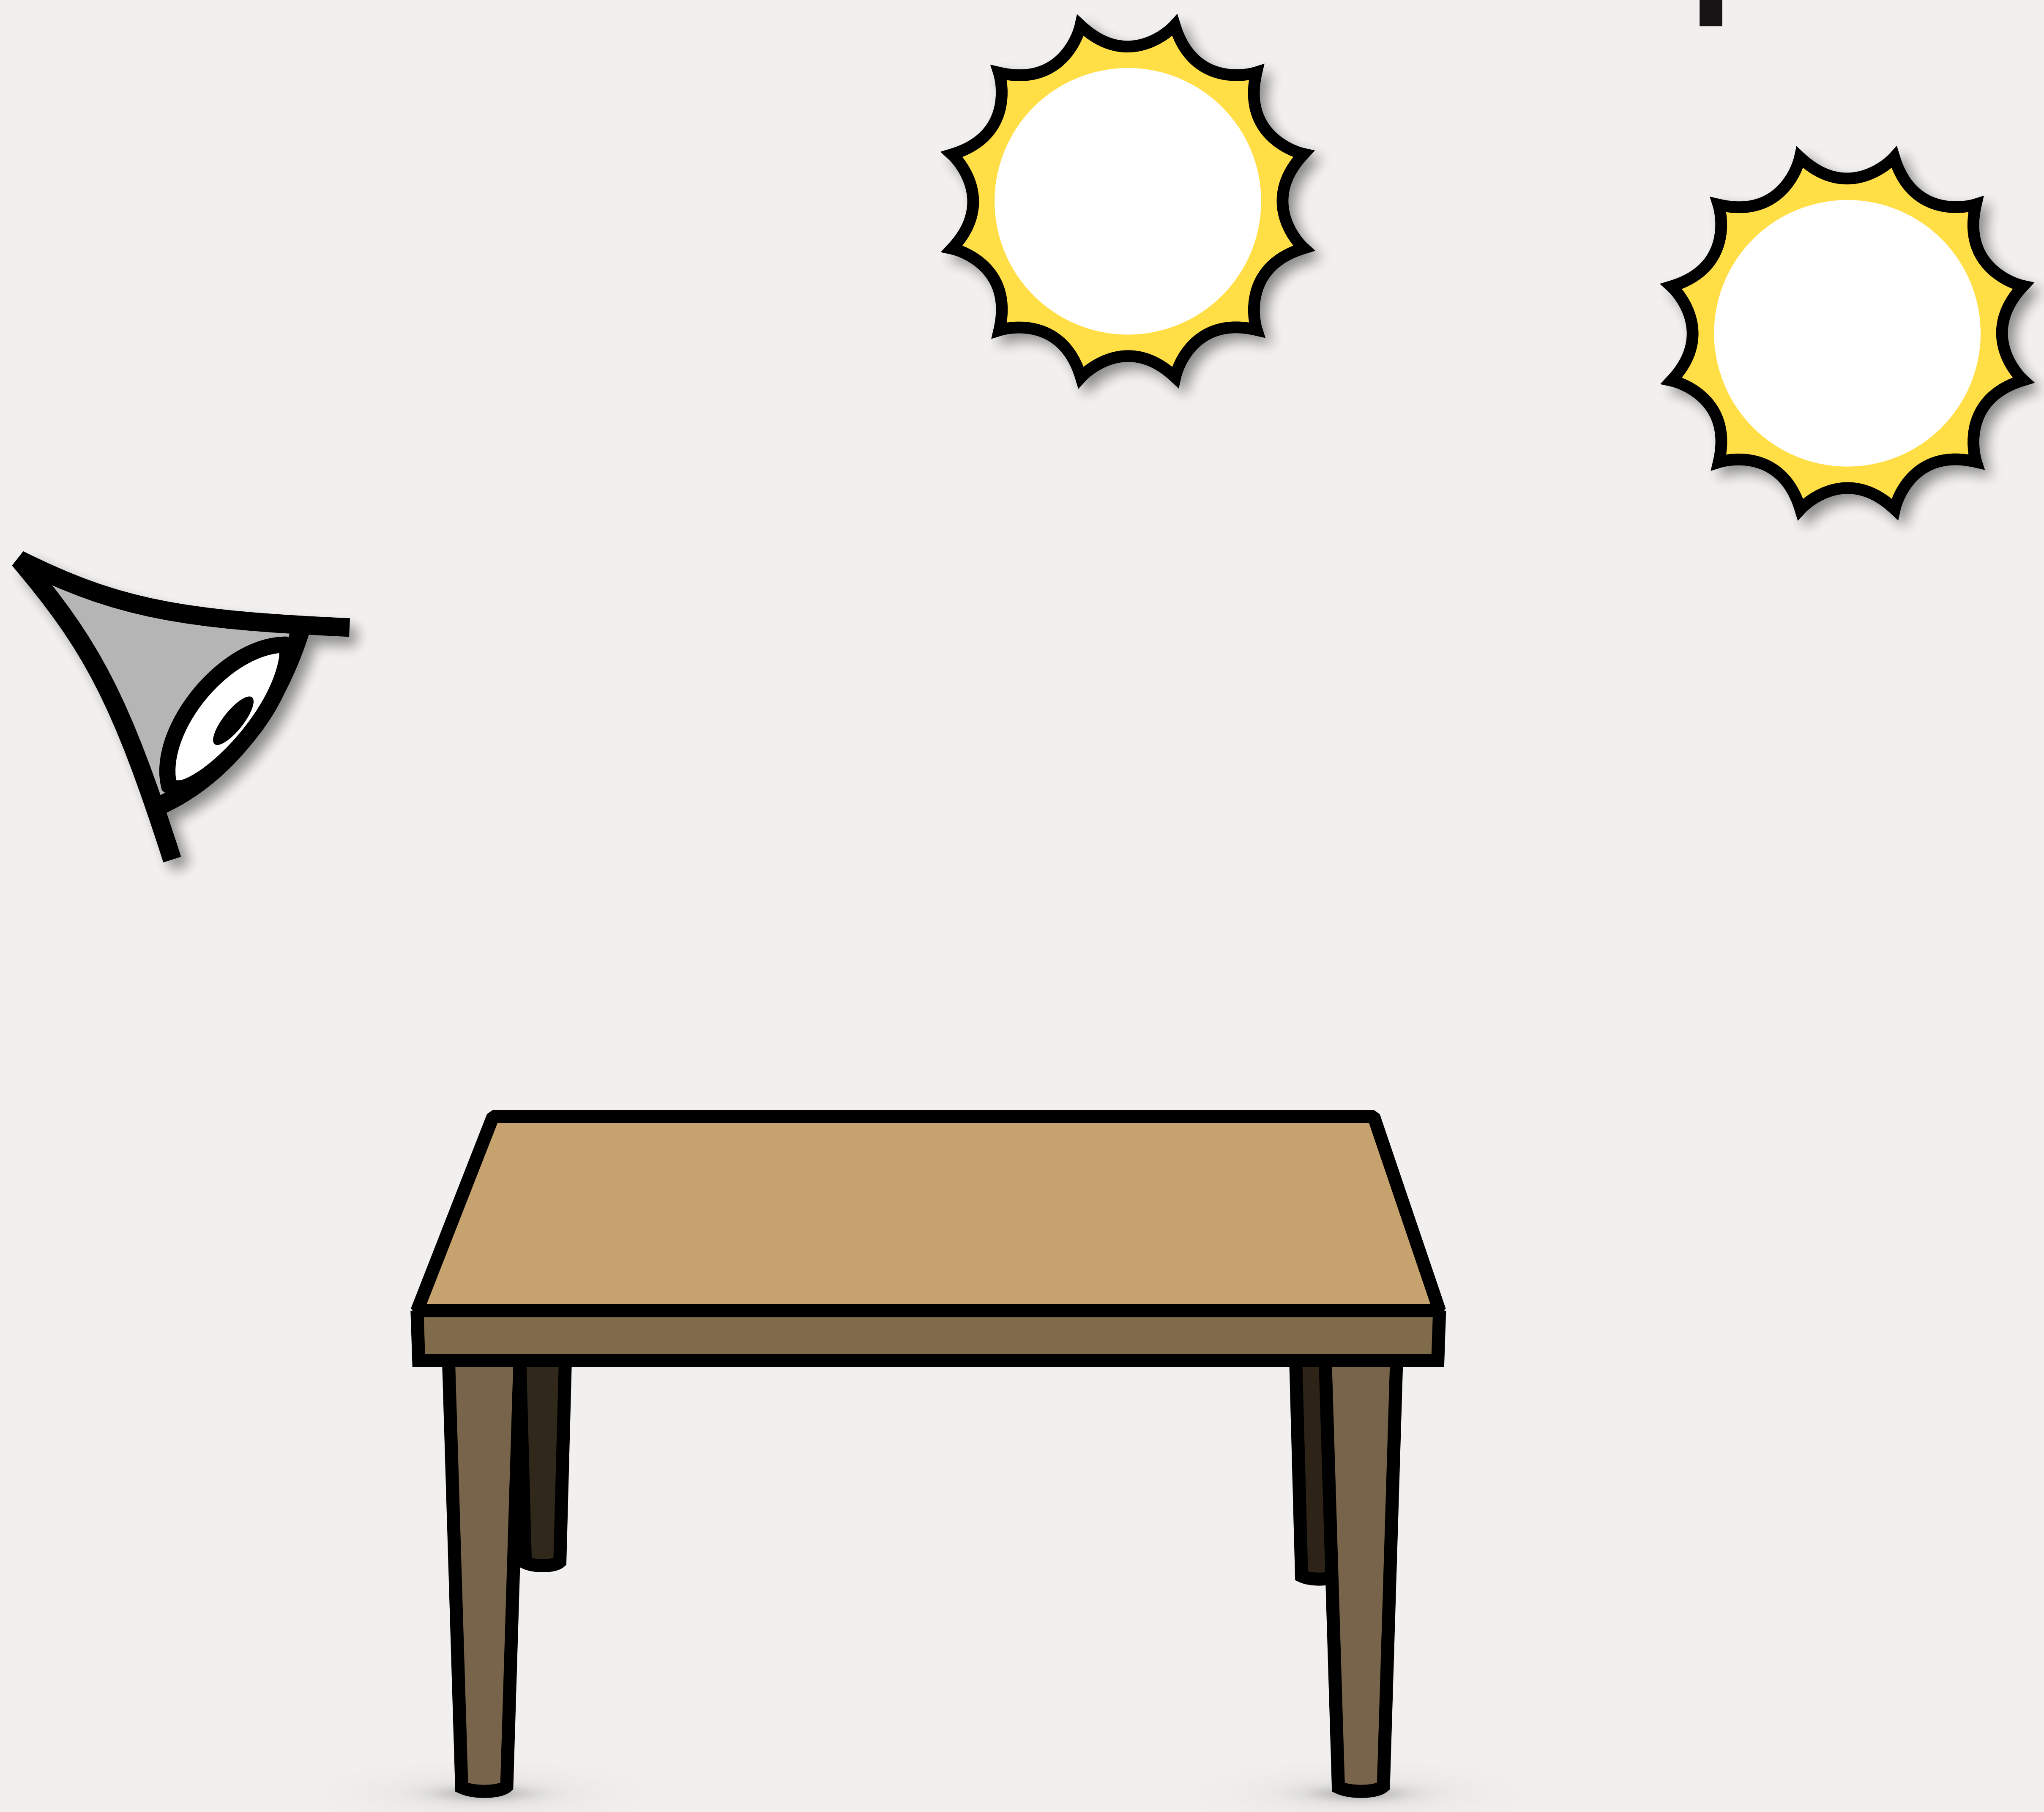

Problem Description

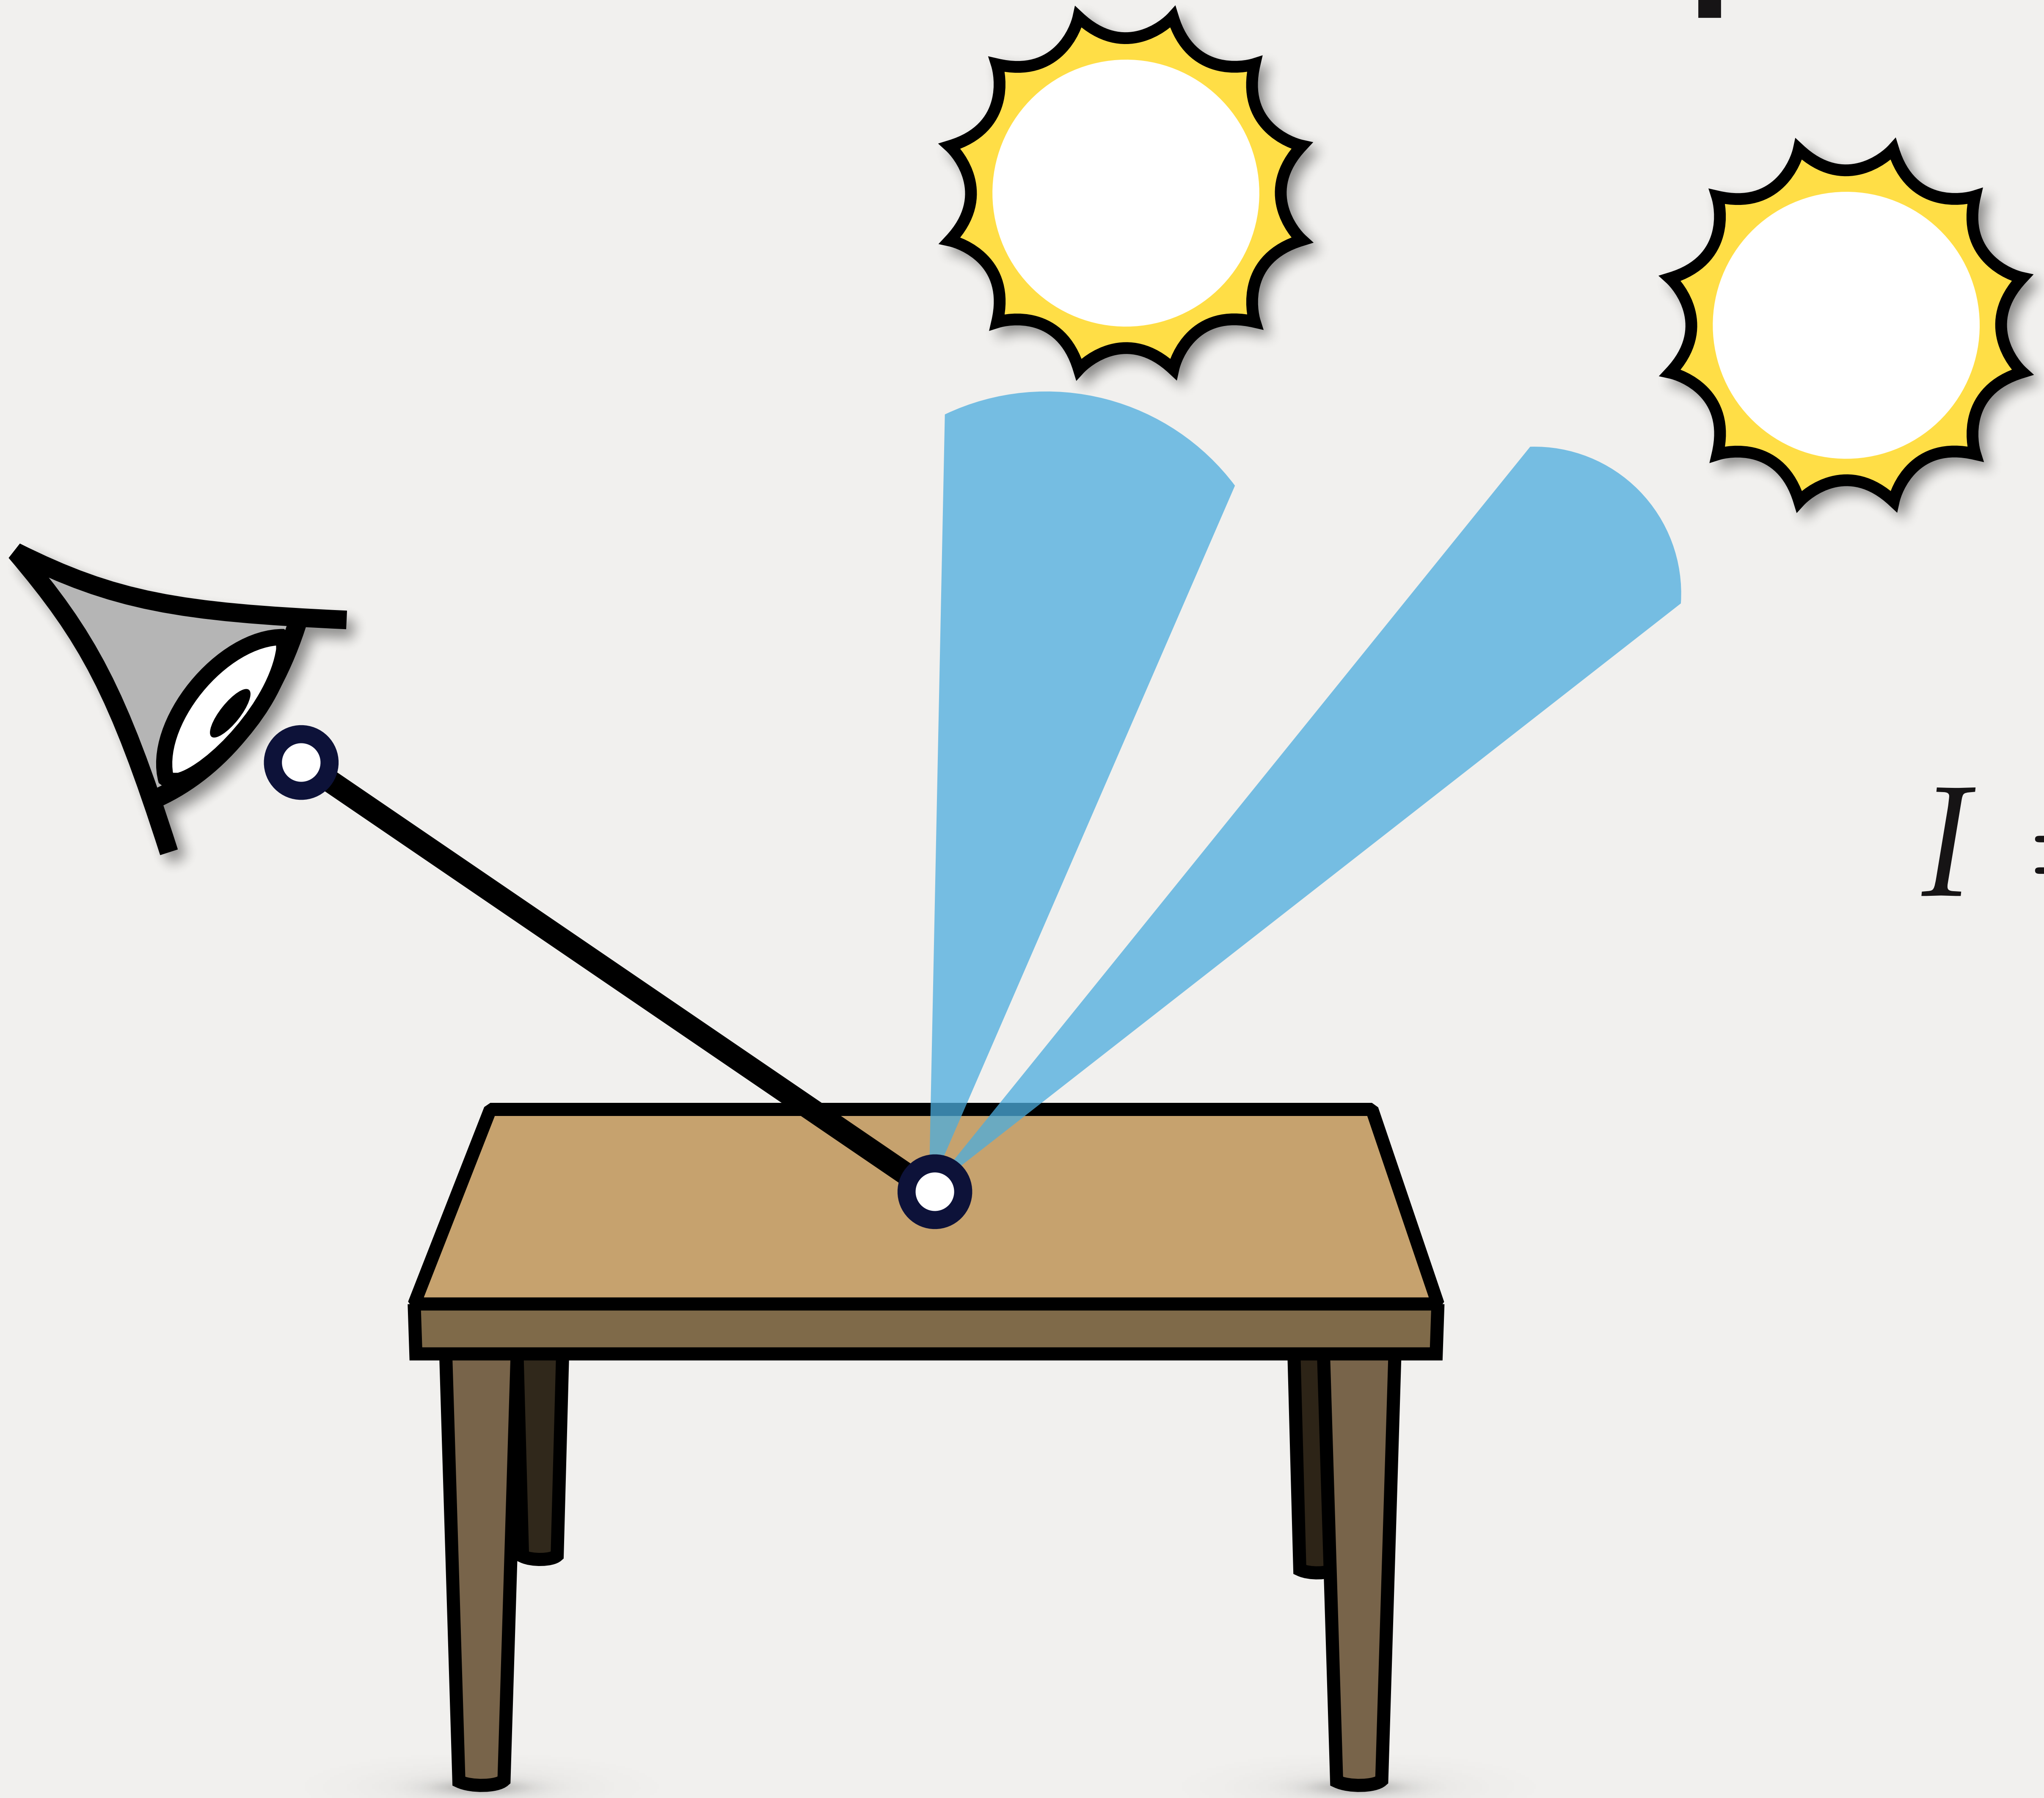

$$I = \int_{\Omega} f(x)\,\mathrm{d}x$$

Problem Description

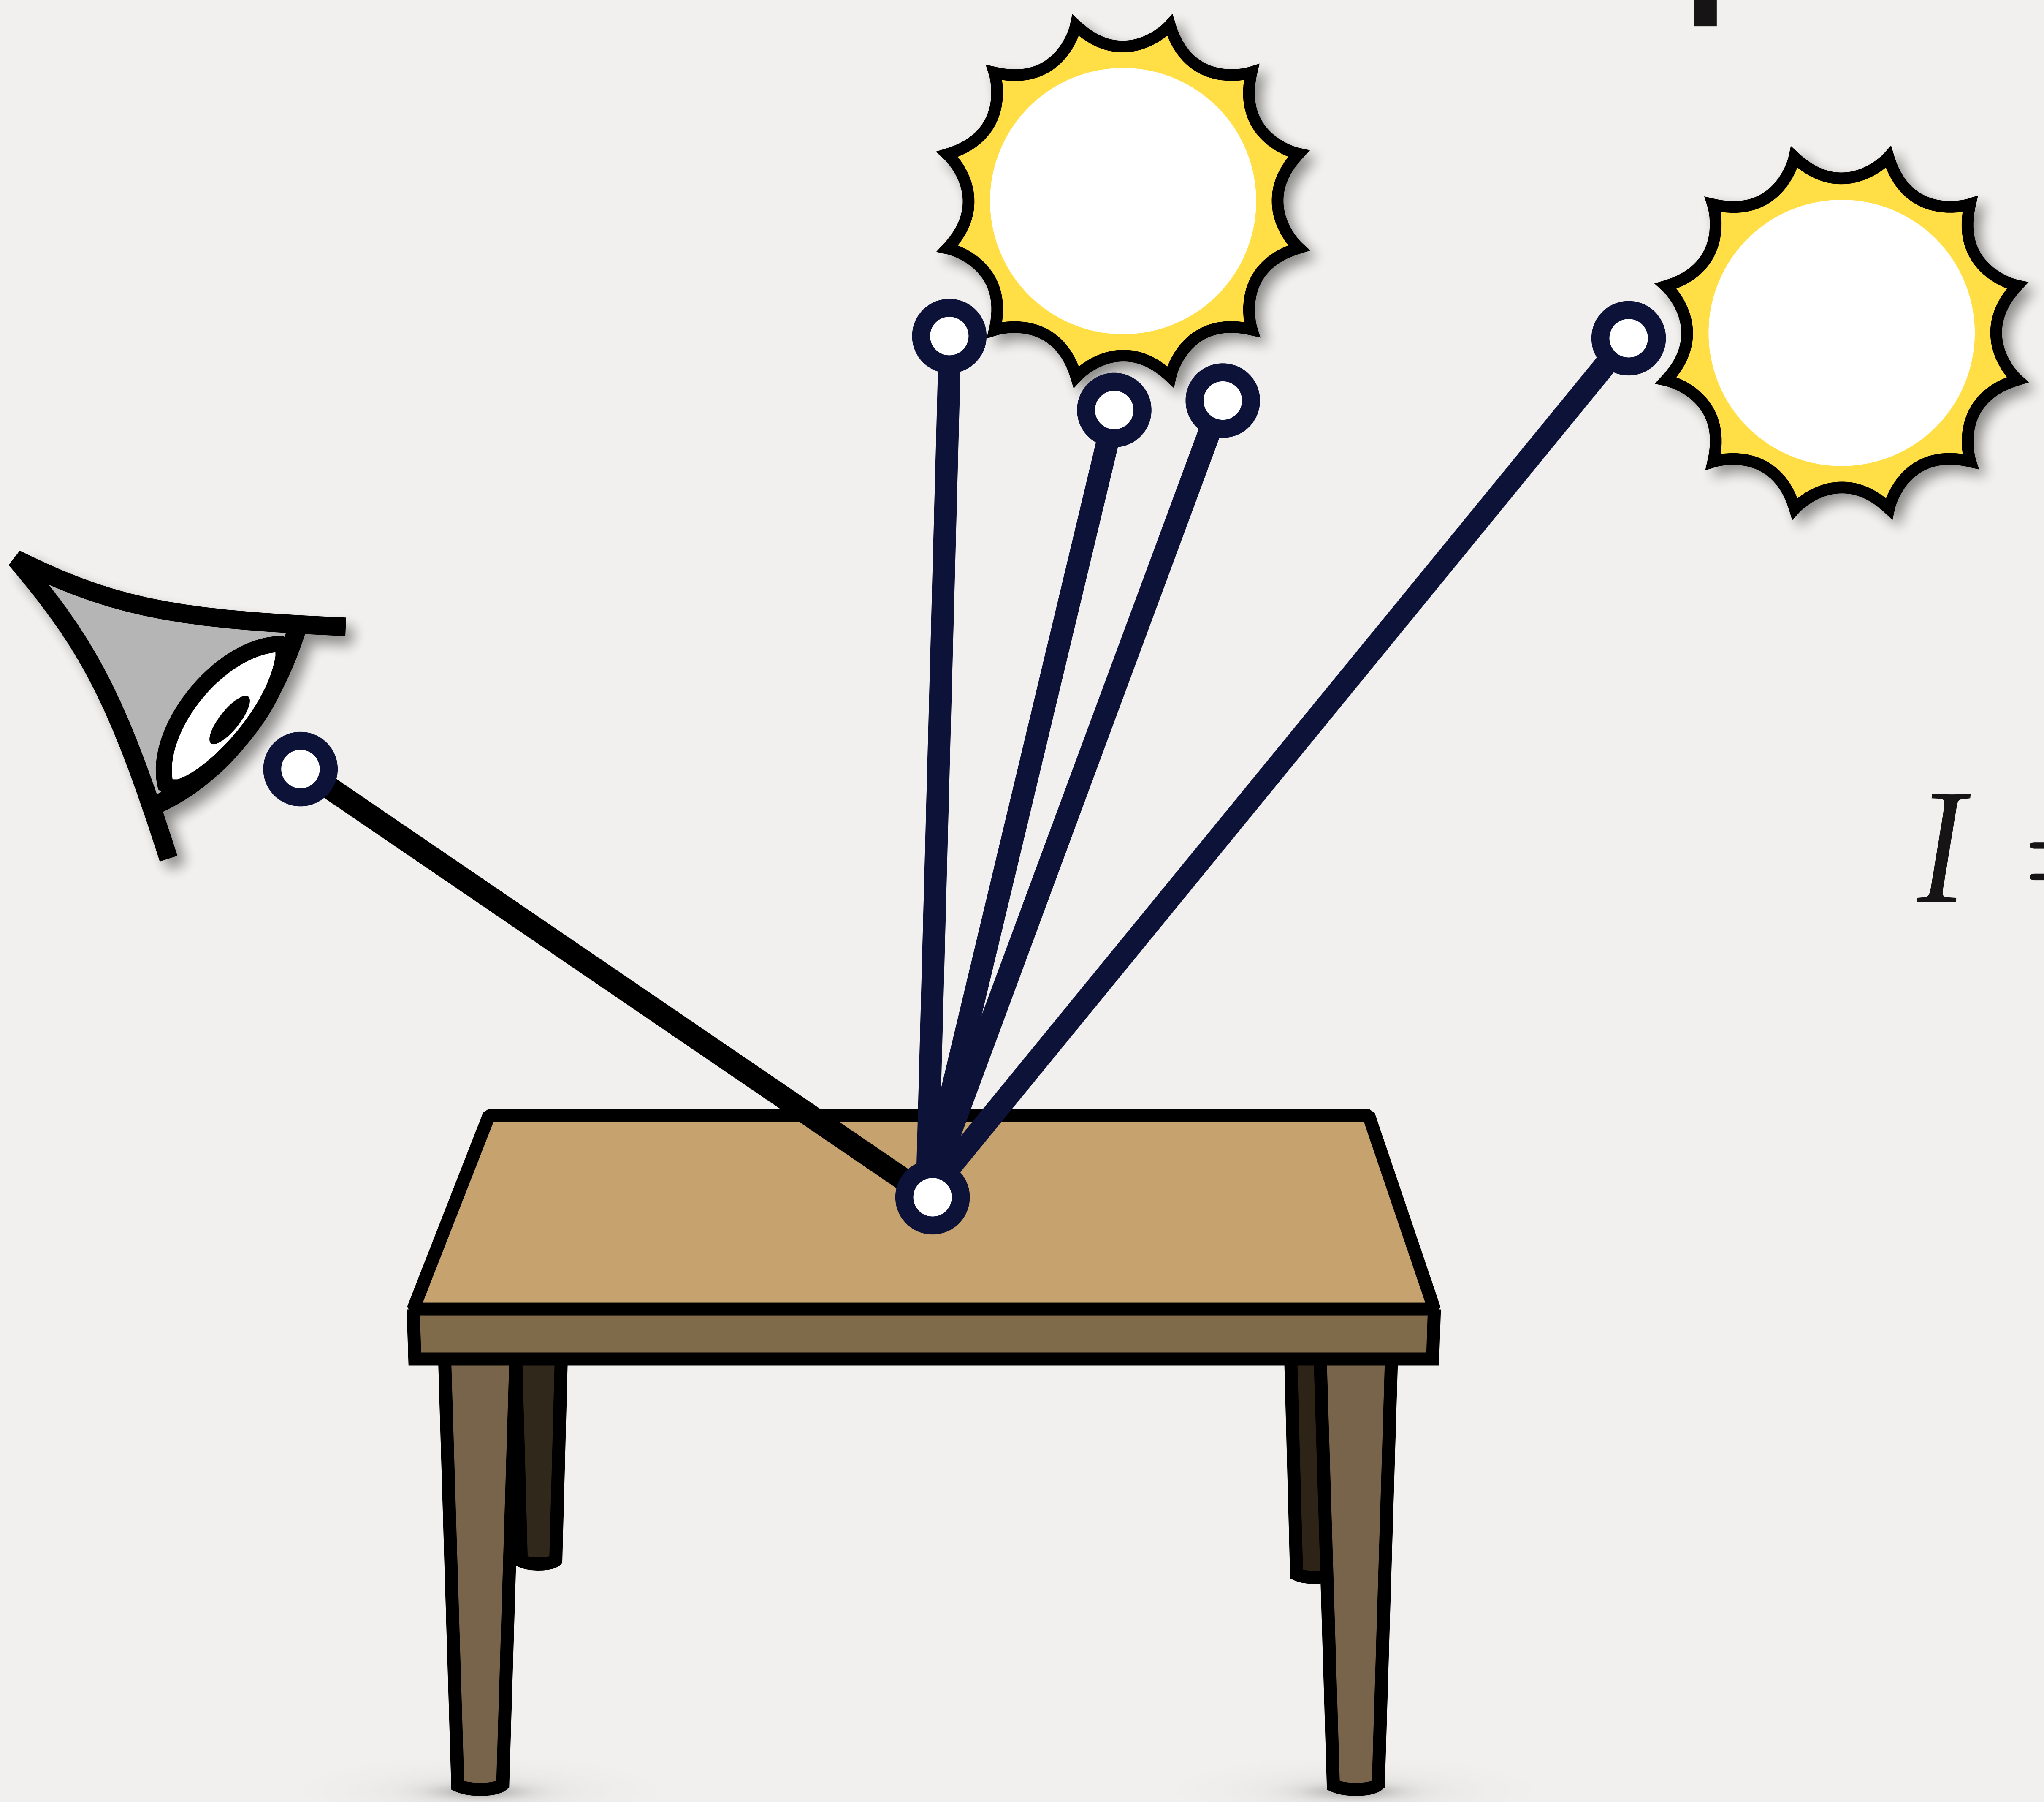

$$I = \int_\Omega f(x)\,\mathrm{d}x \approx \frac{1}{N}\sum_i^N \frac{f(x_i)}{p(x_i)}$$

Problem Description

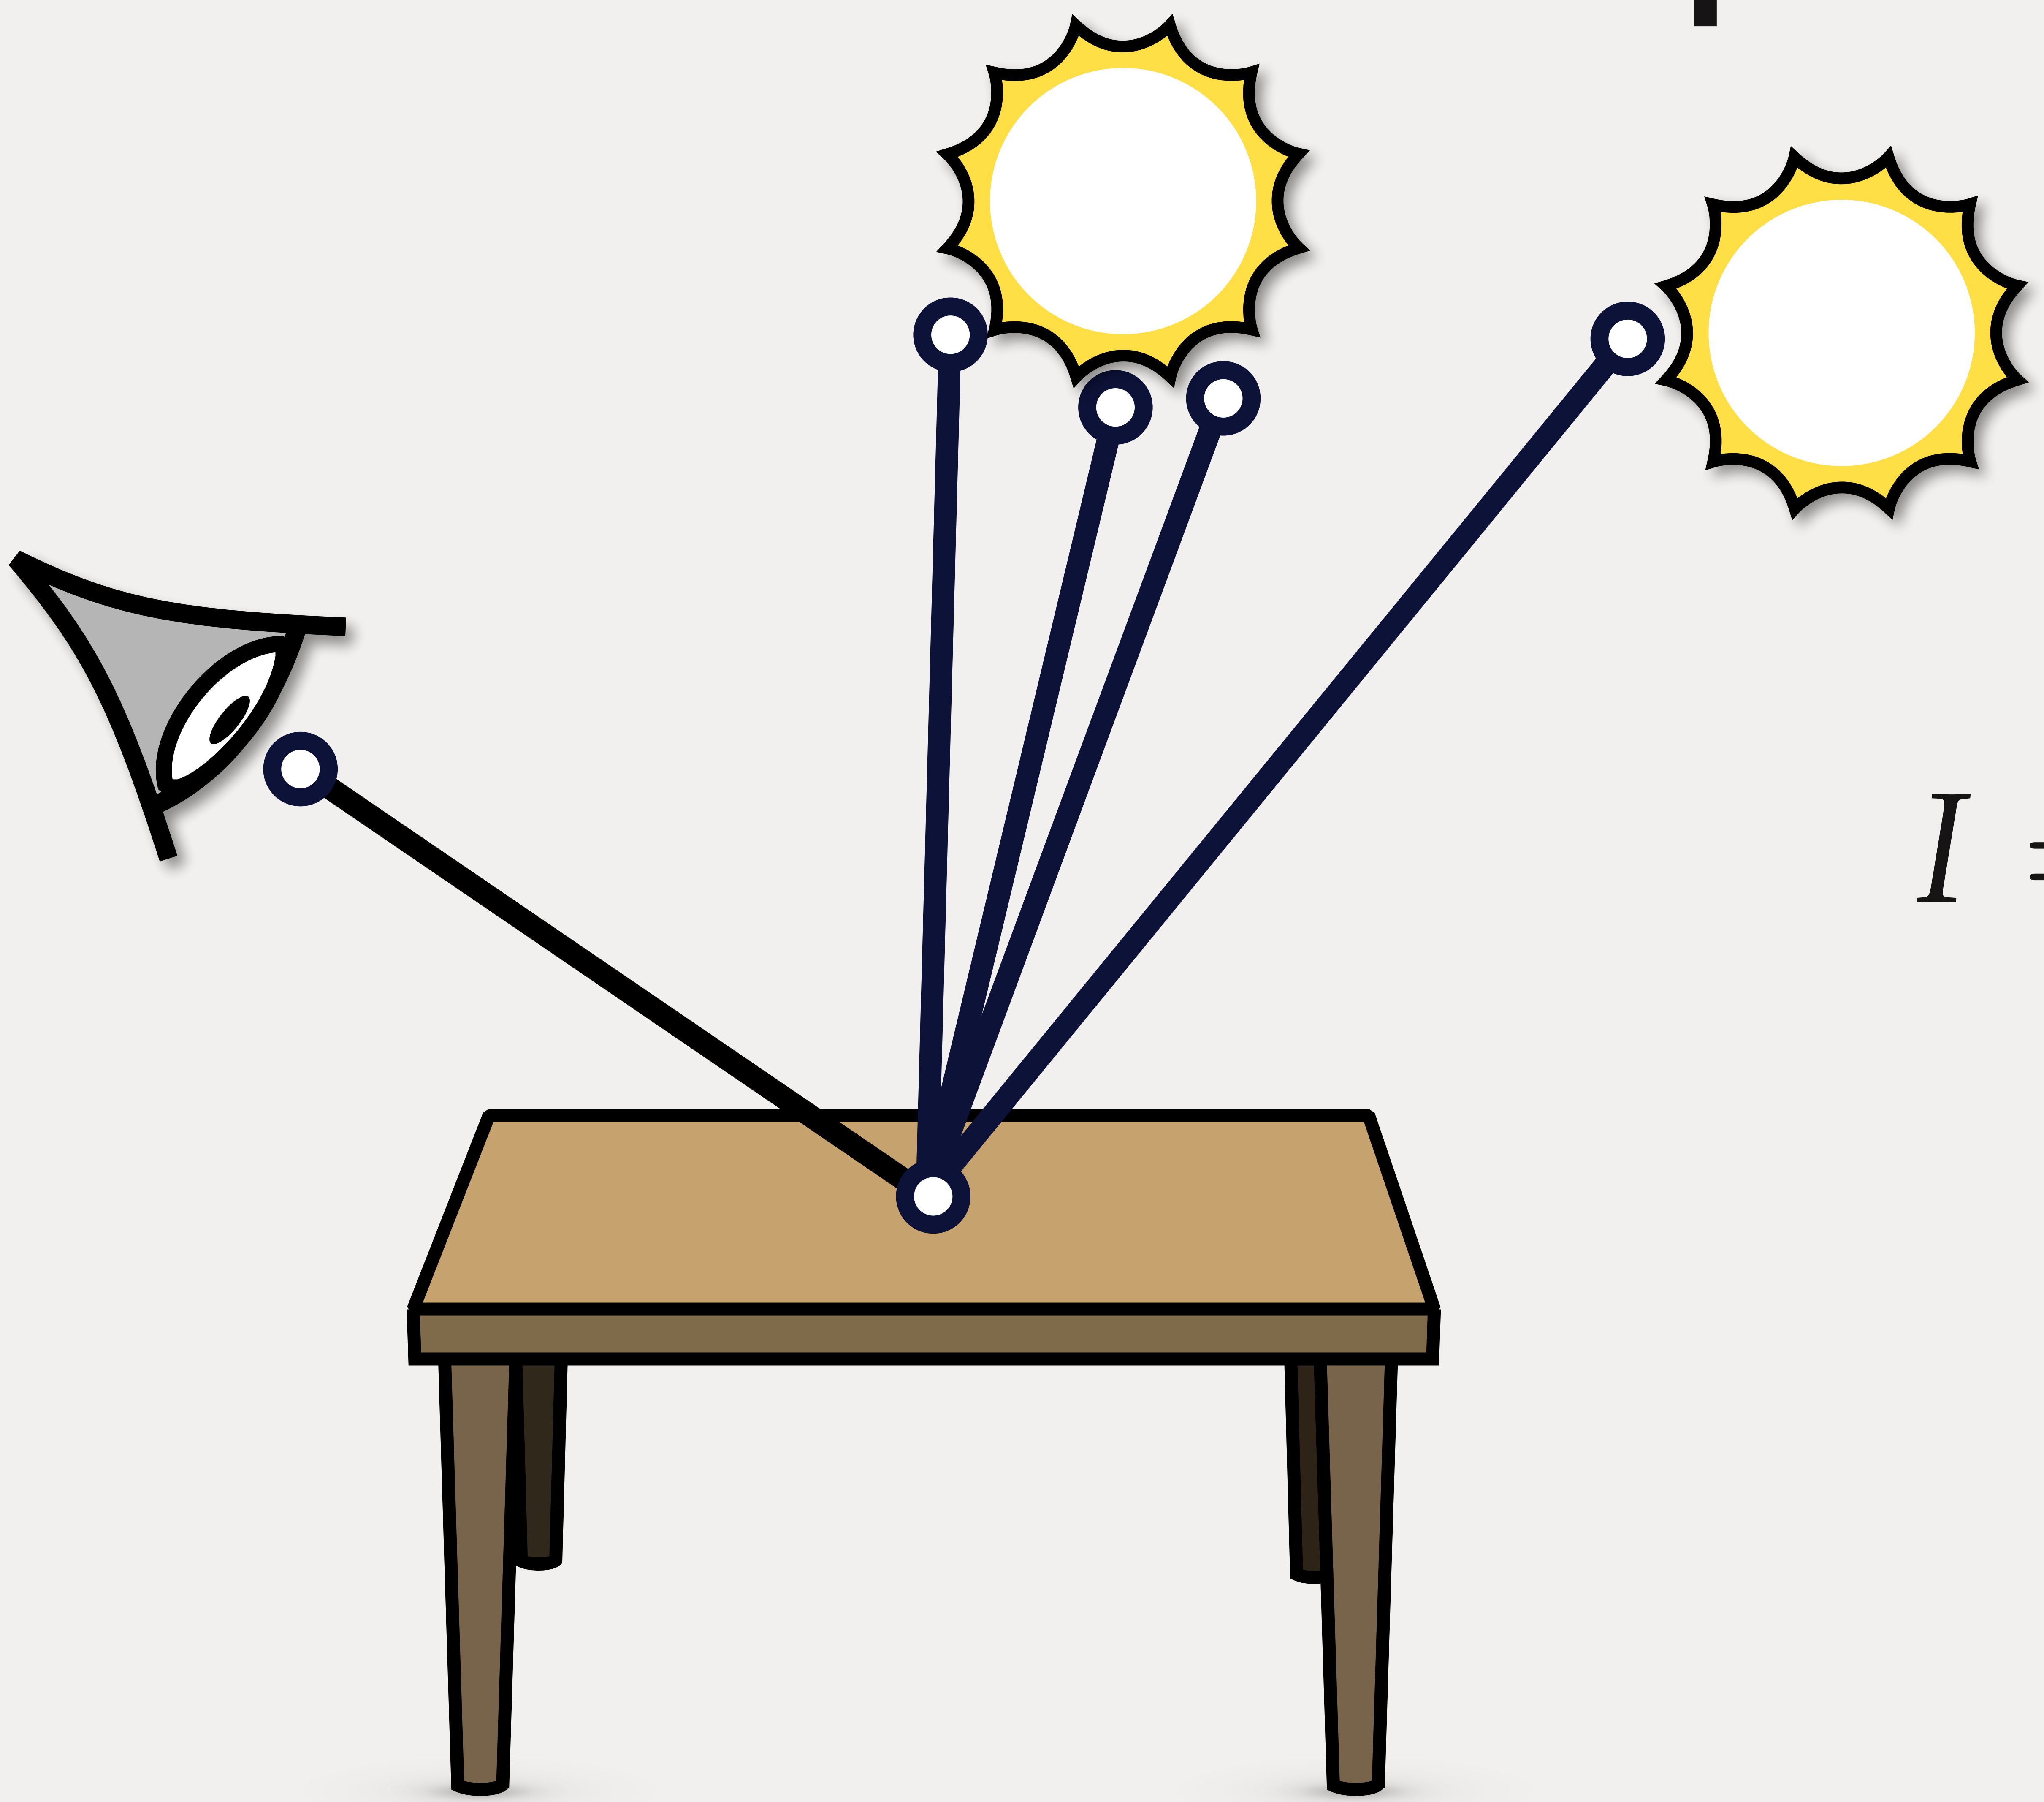

$$I = \int_\Omega f(x)\mathrm{d}x \approx \frac{1}{\boxed{N}} \sum_i^N \frac{f(x_i)}{p(x_i)}$$

Problem Description

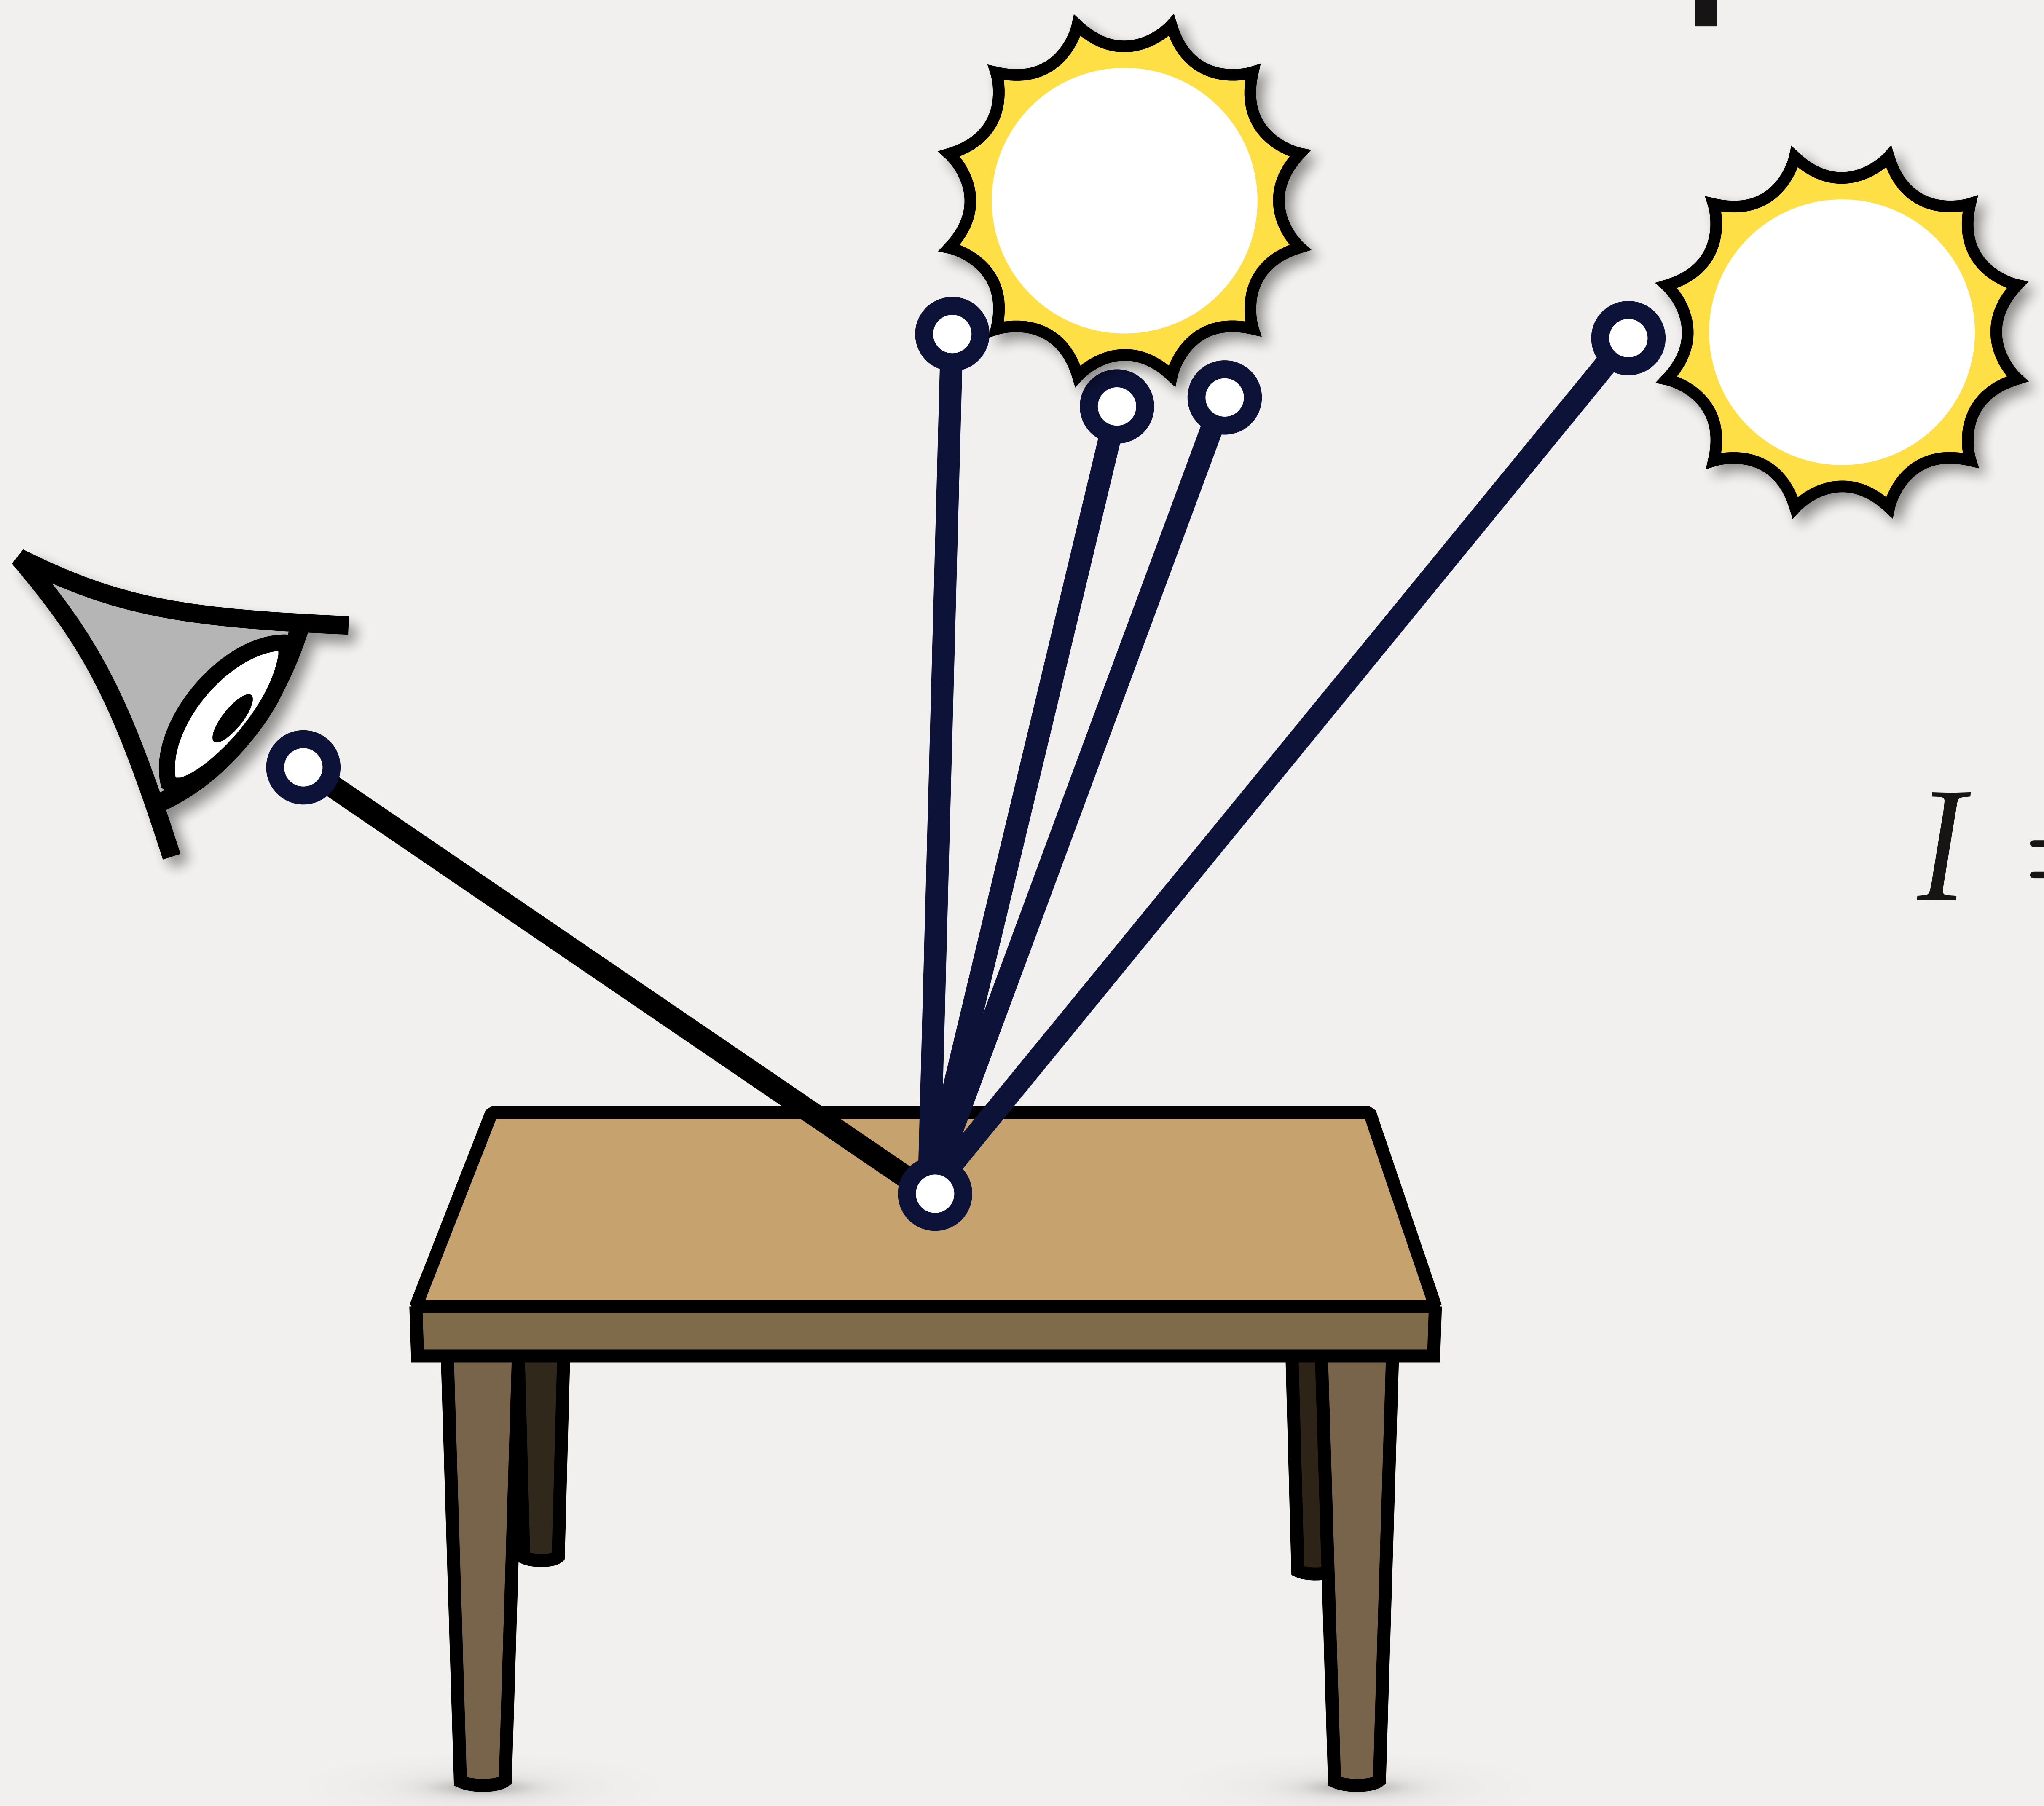

$$I = \int_\Omega f(x)\,\mathrm{d}x \approx \frac{1}{N}\sum_i^N \frac{f(x_i)}{\boxed{p(x_i)}}$$

Problem Description

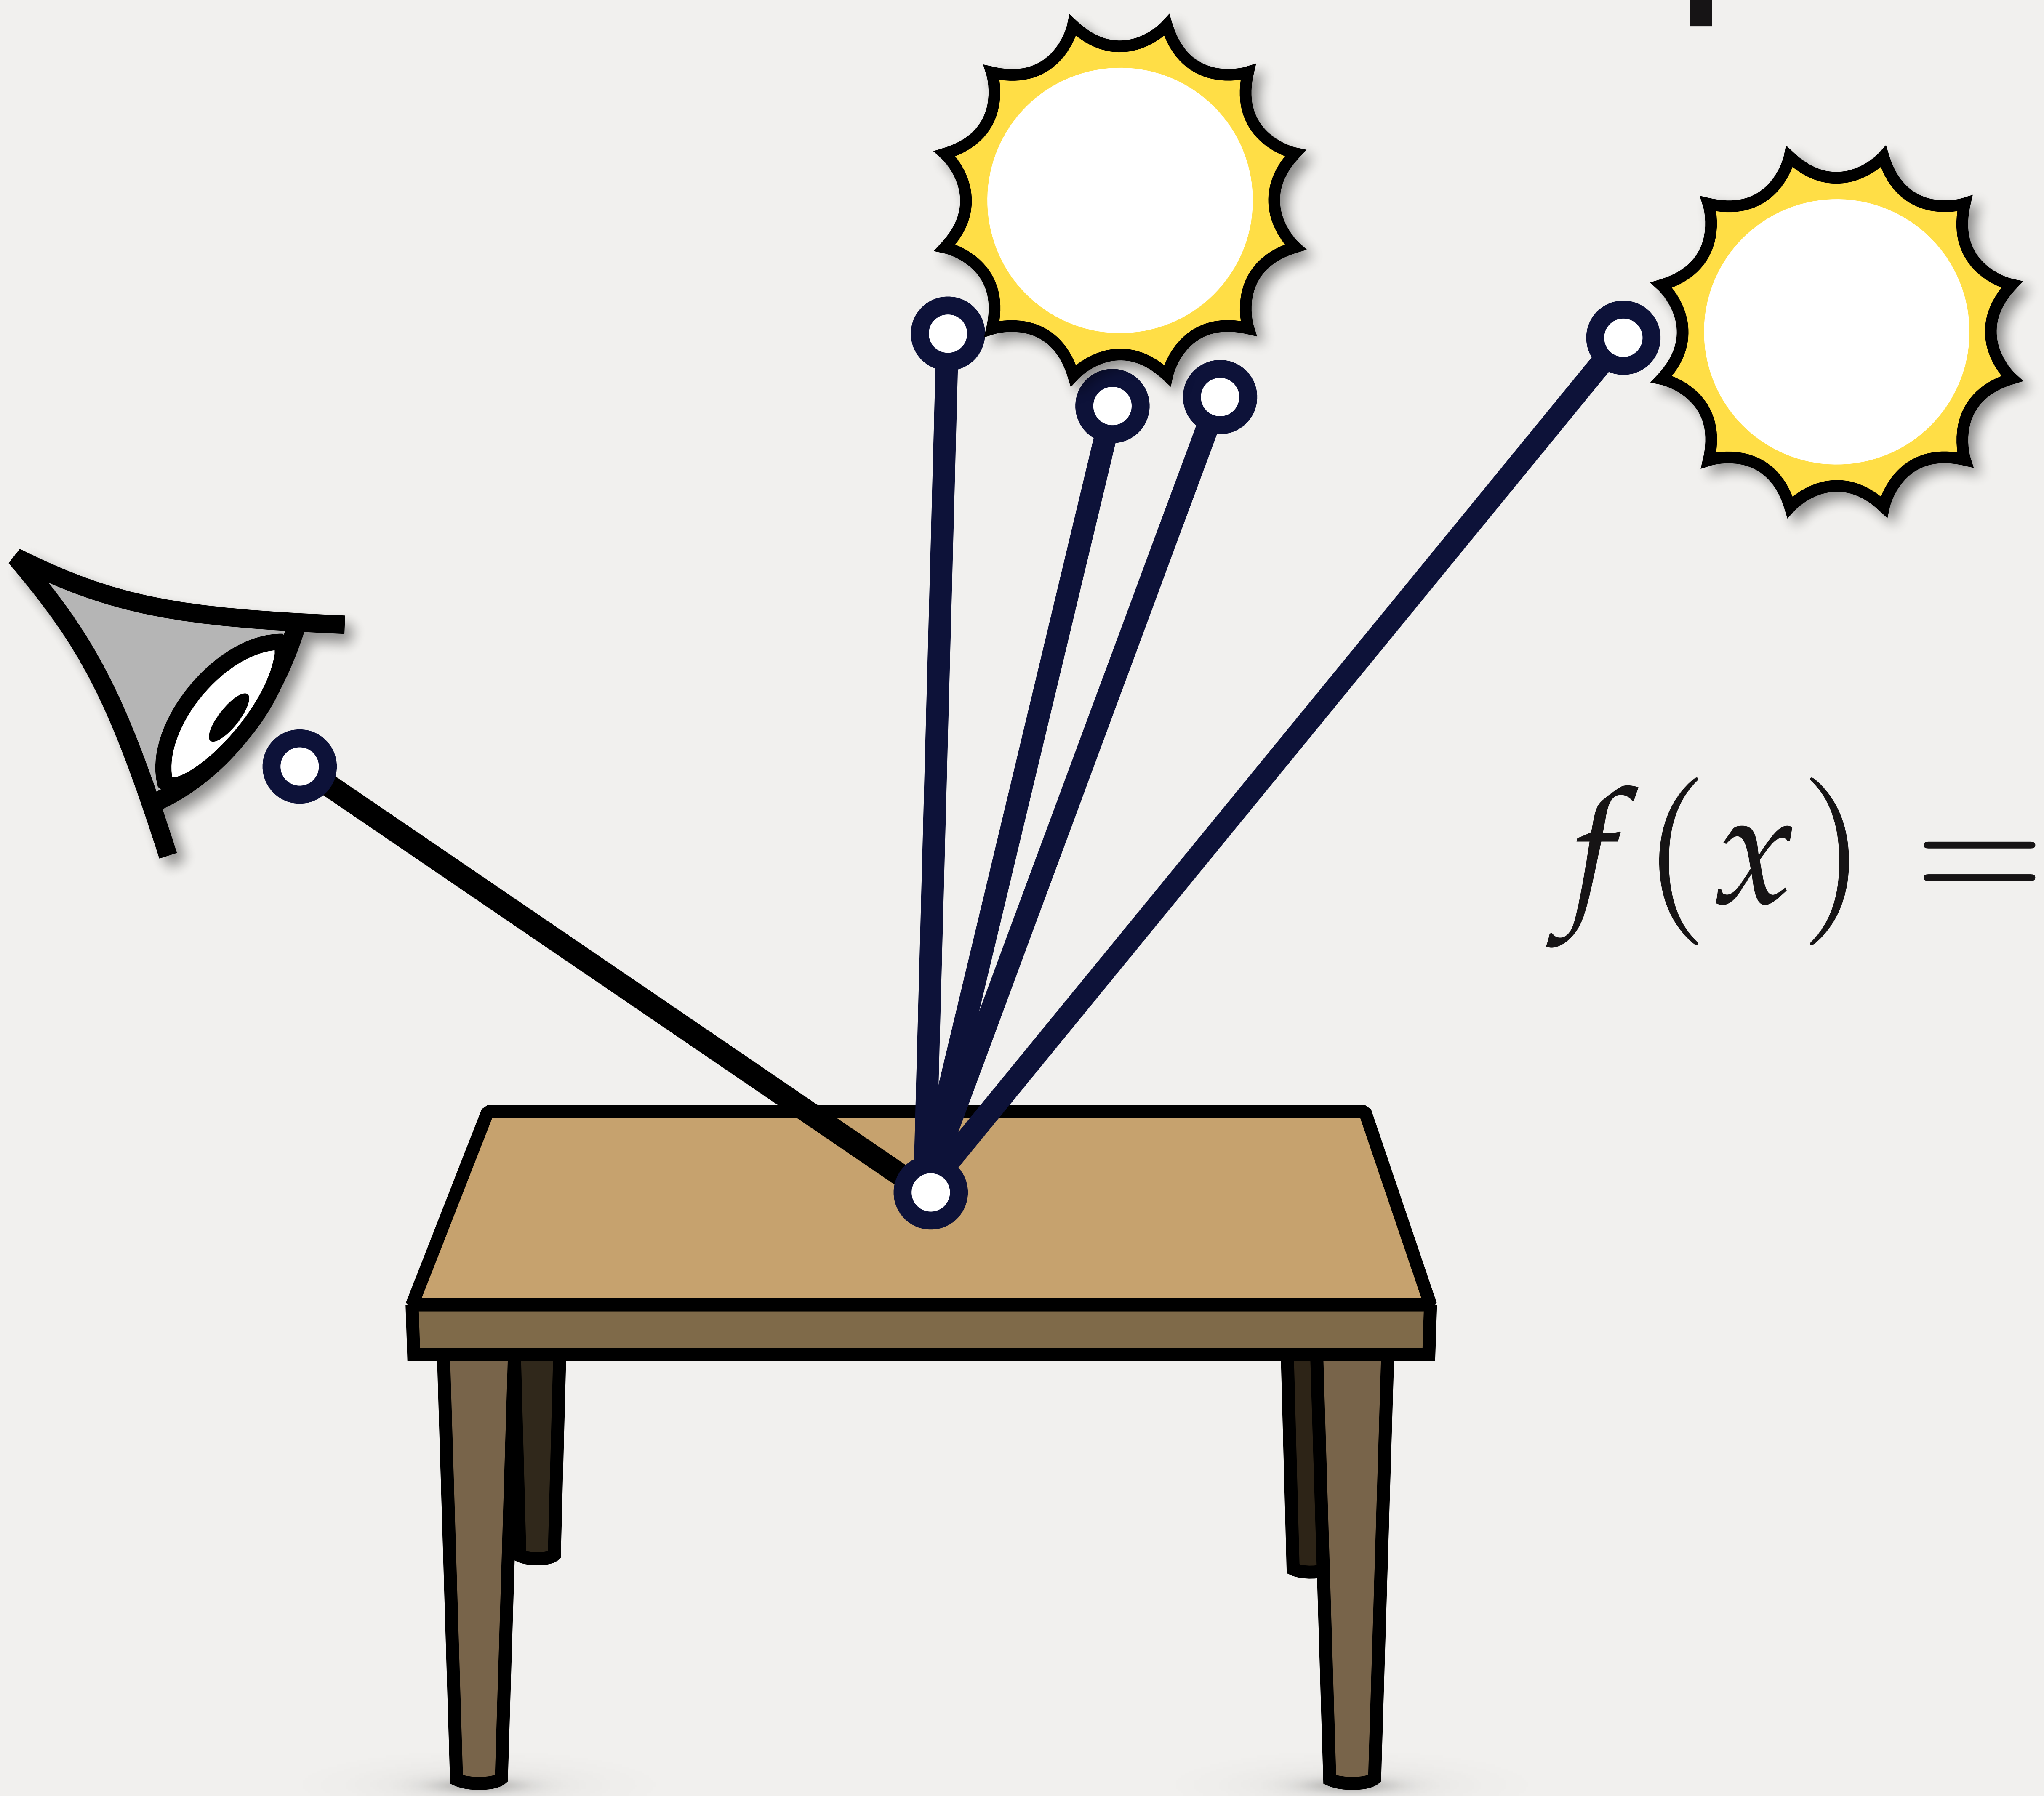

$$f(x) = L_e(x)\rho(x)V(x)\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

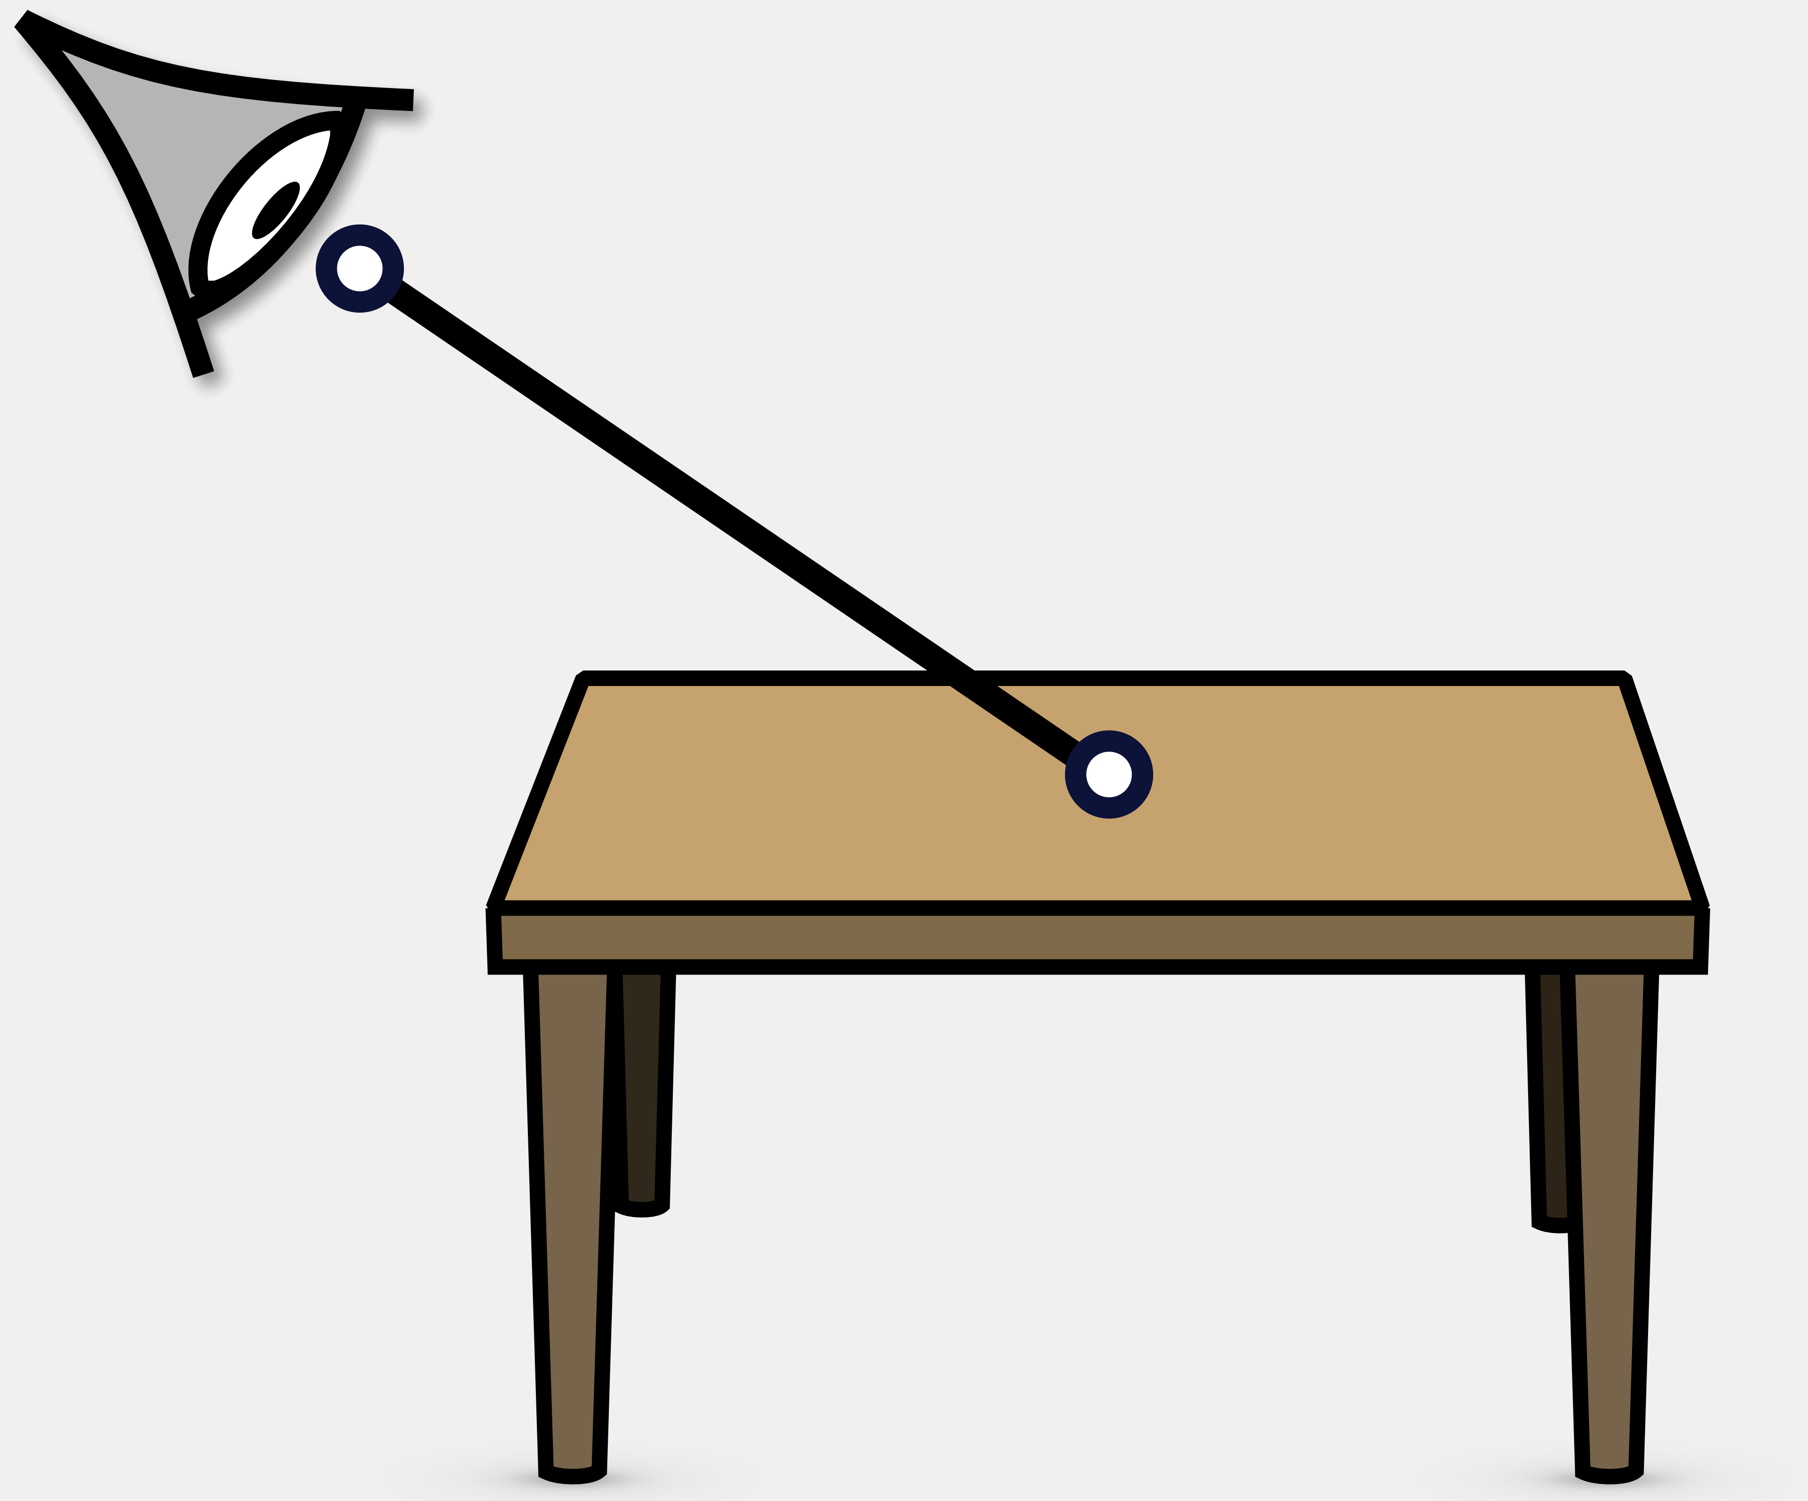

$$f(x) = L_e(x)\rho(x)V(x)\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

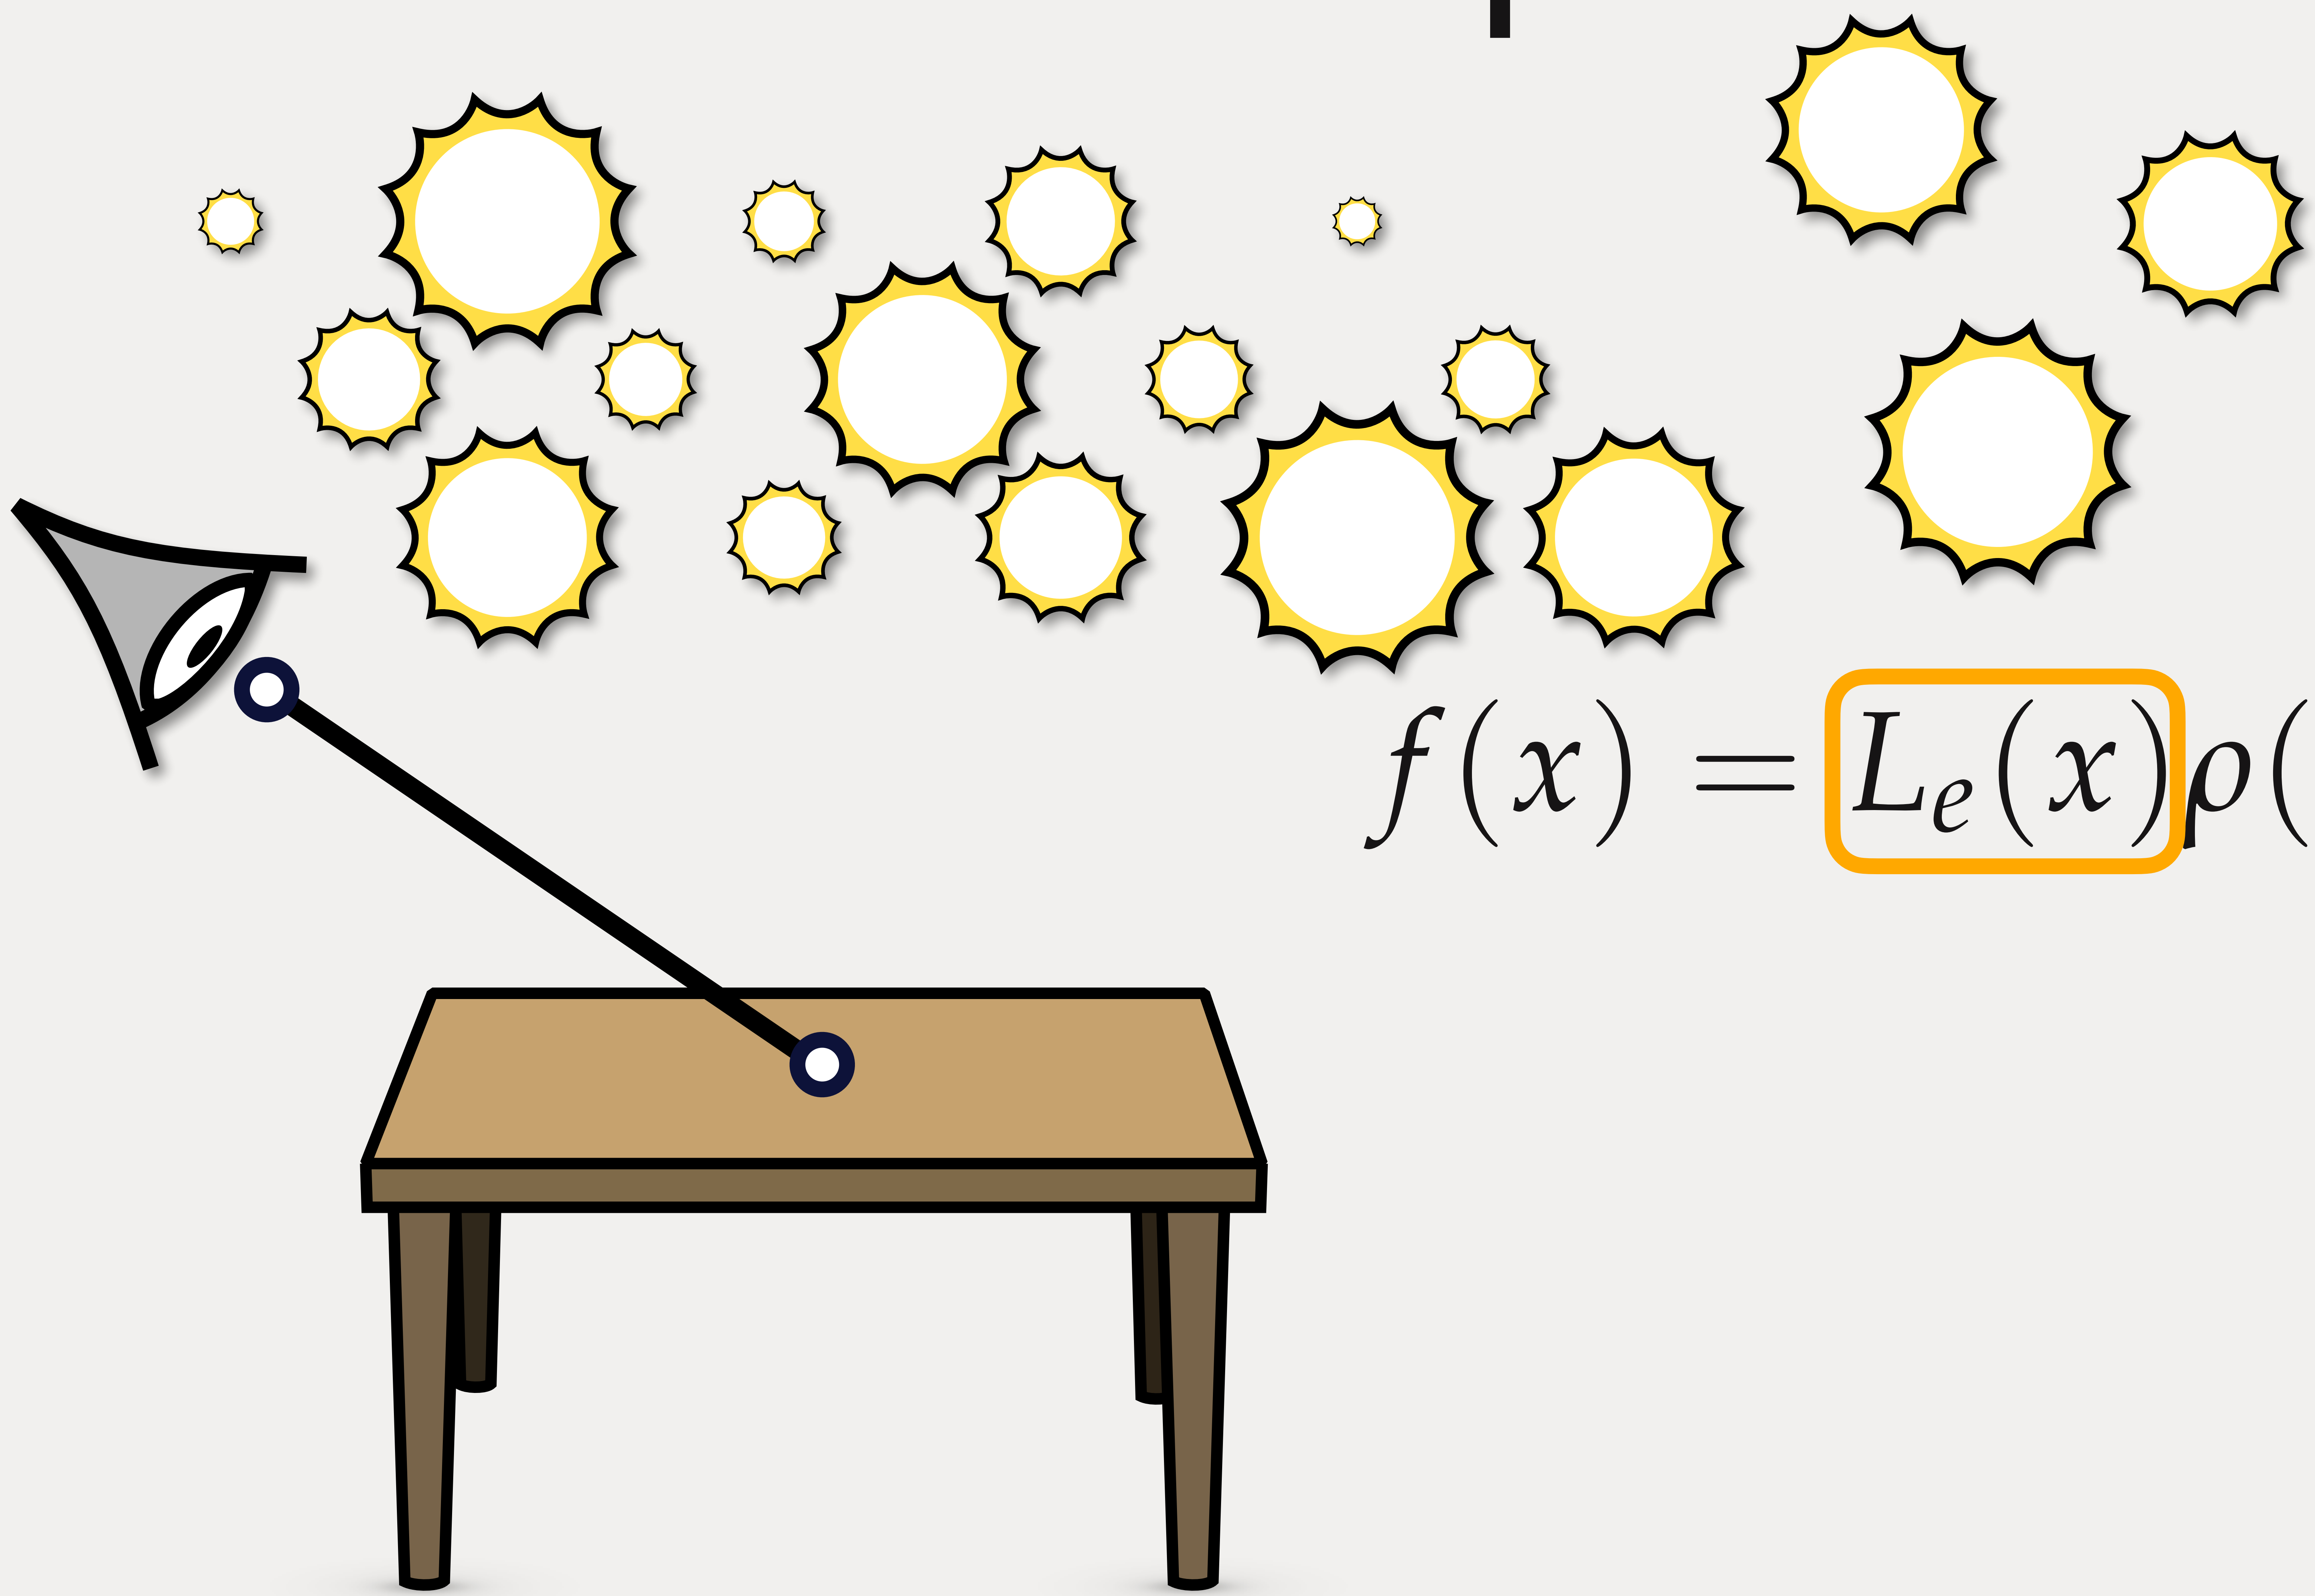

$$f(x) = \boxed{L_e(x)}\rho(x)V(x)\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

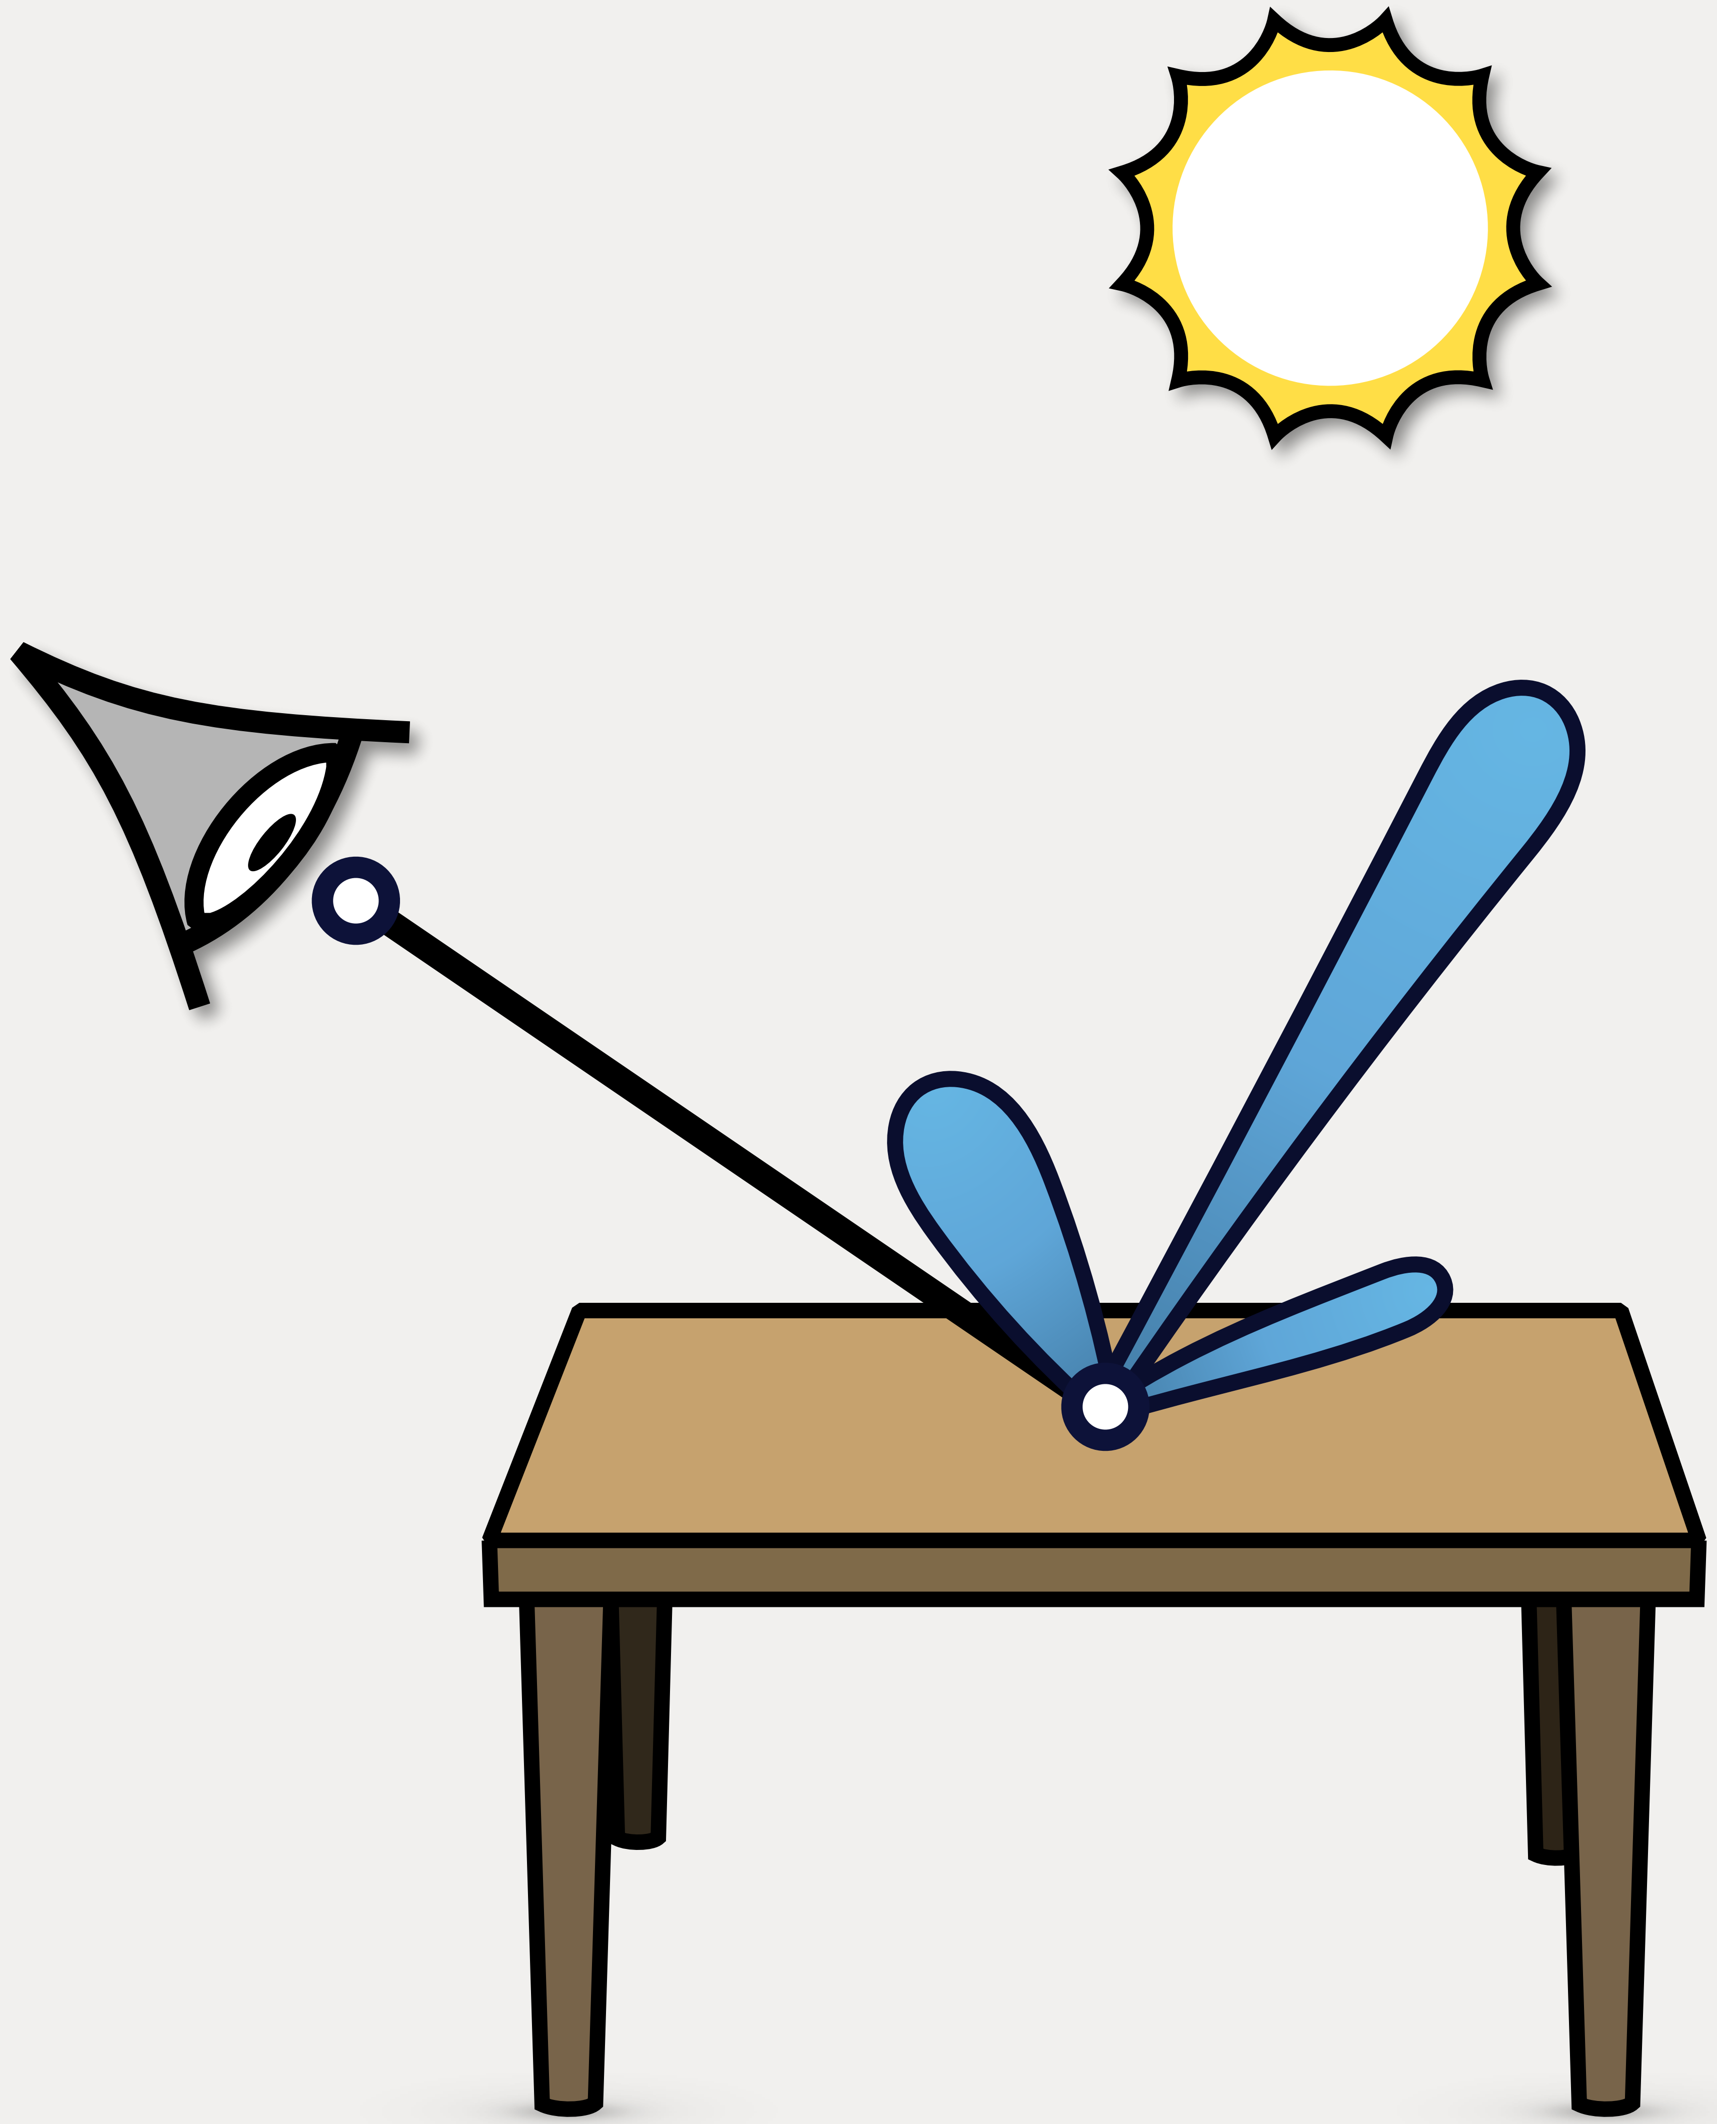

$$f(x) = L_e(x)\boxed{\rho(x)}V(x)\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

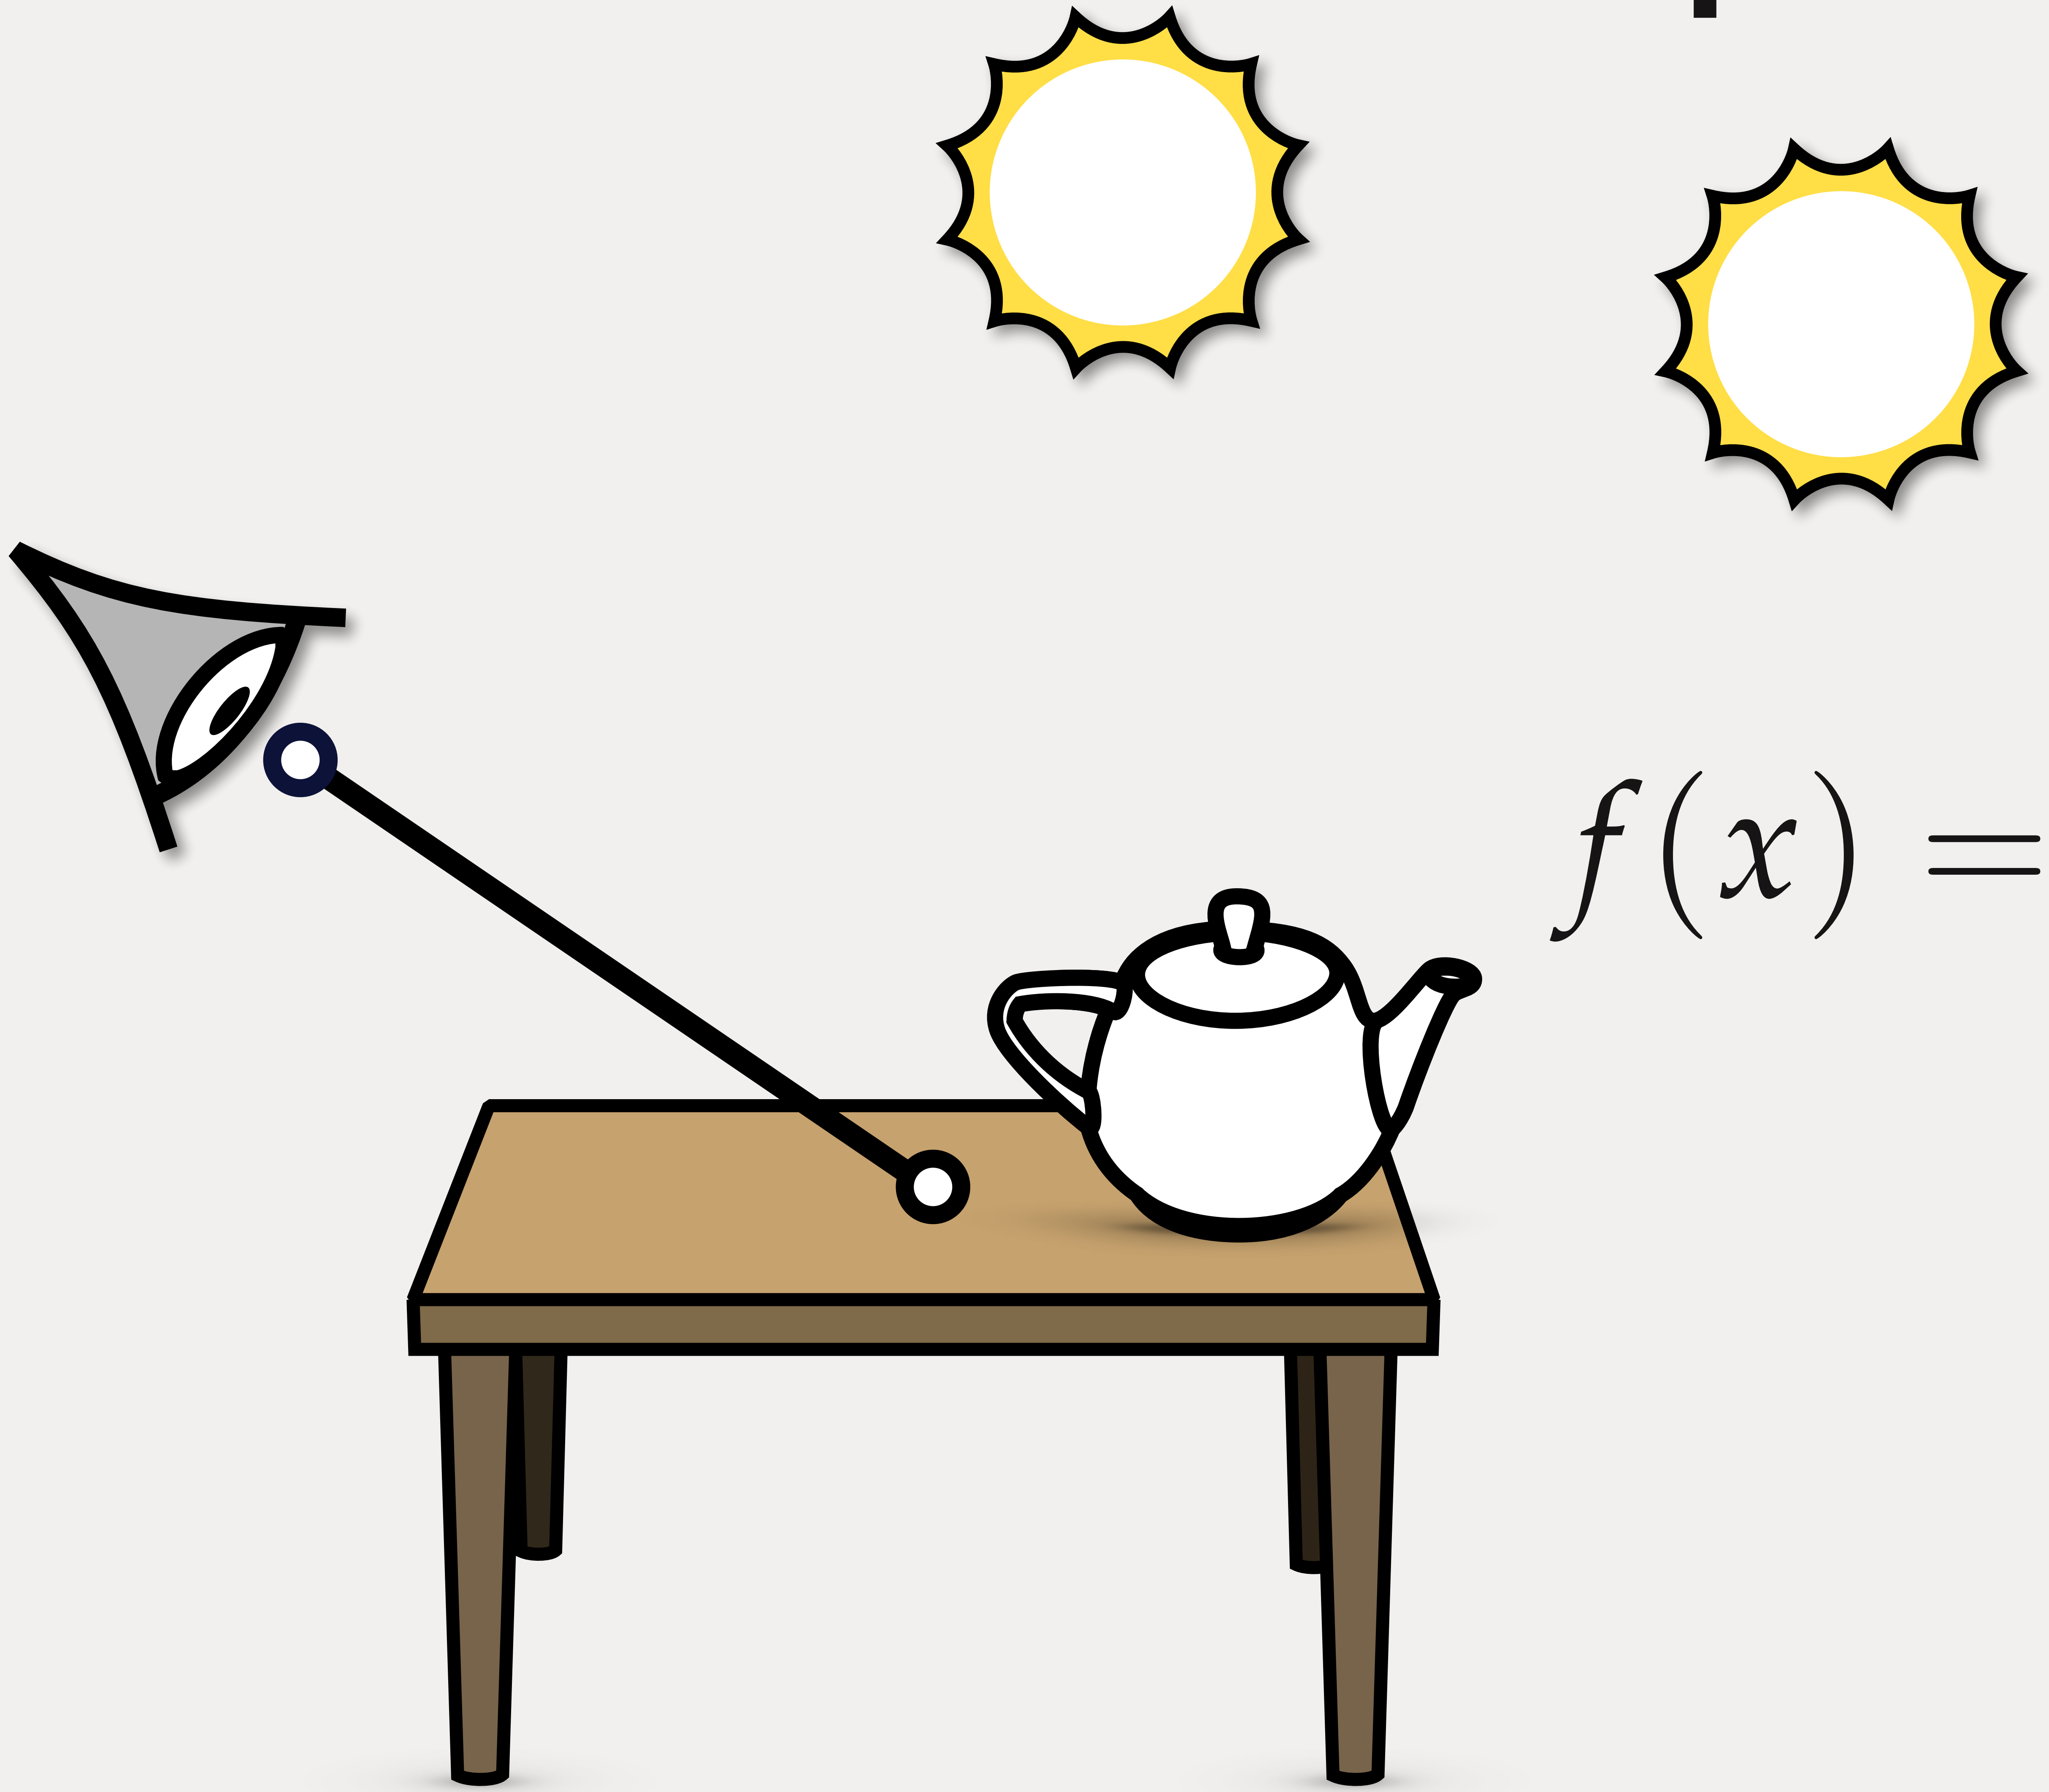

$$f(x) = L_e(x)\rho(x)\boxed{V(x)}\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

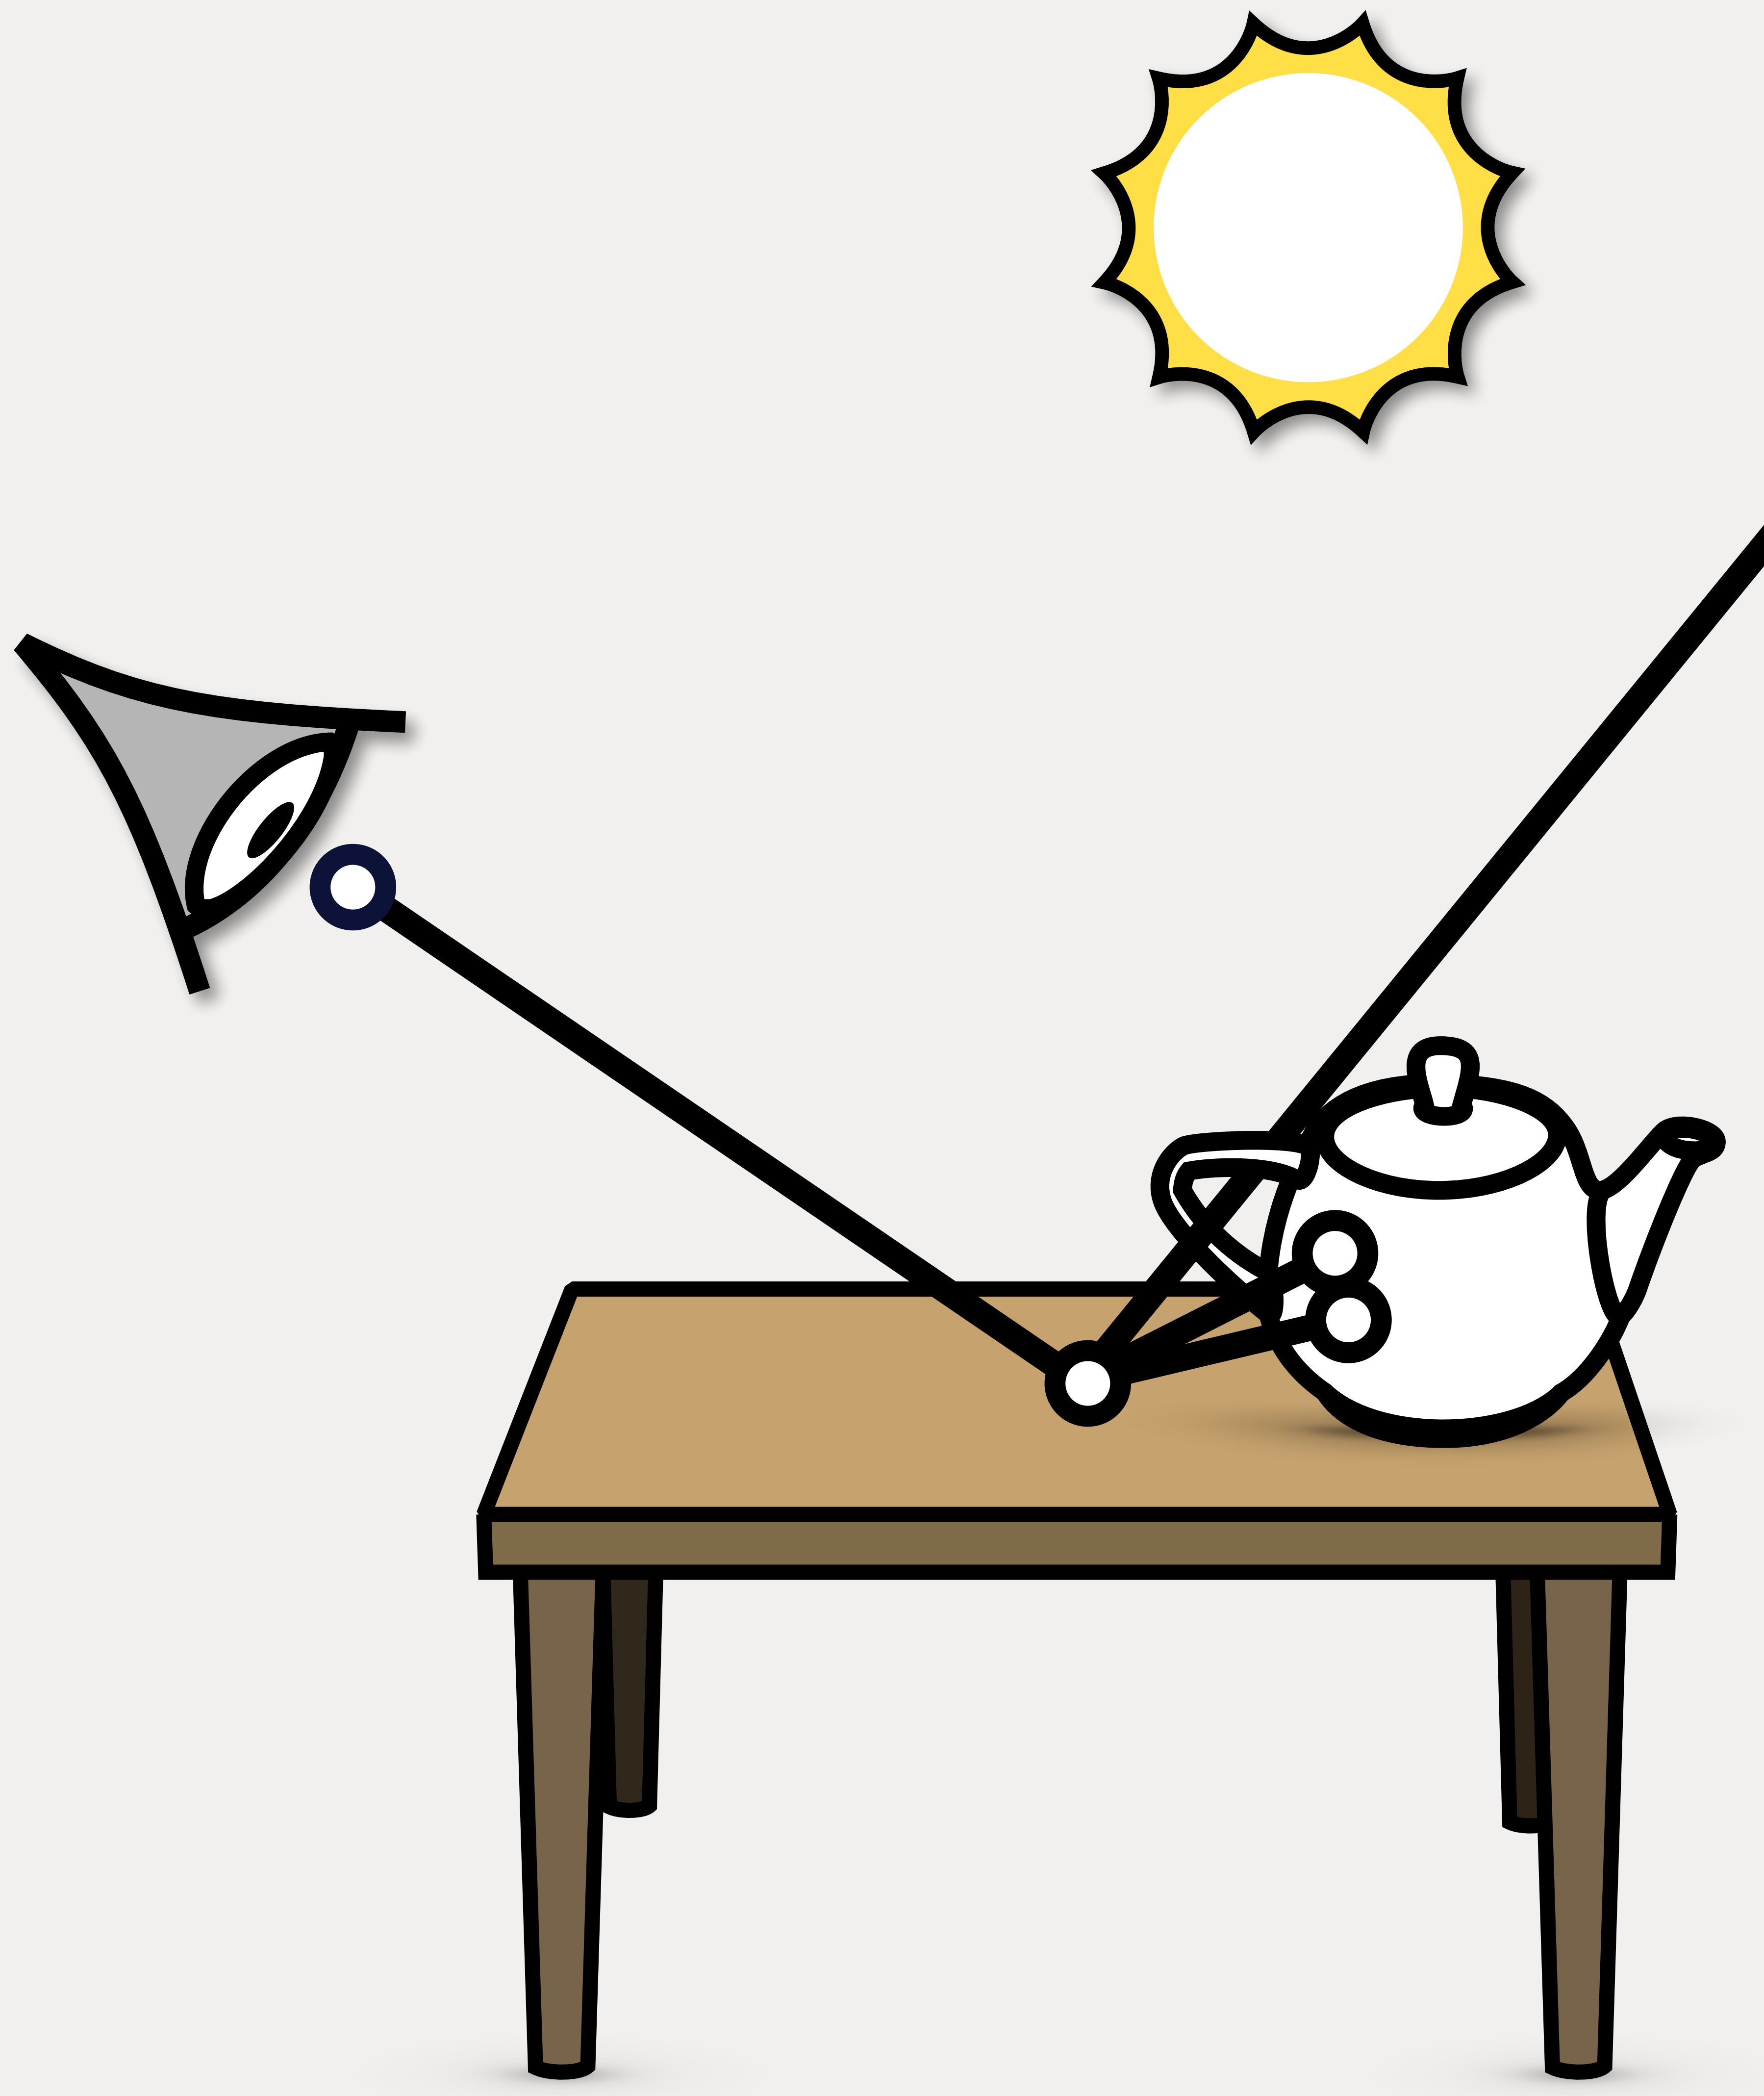

$$f(x) = L_e(x)\rho(x)\boxed{V(x)}\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Problem Description

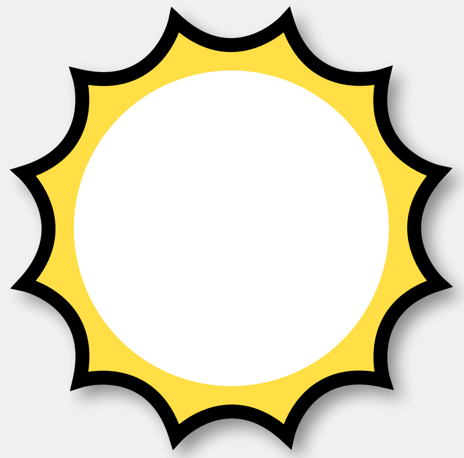

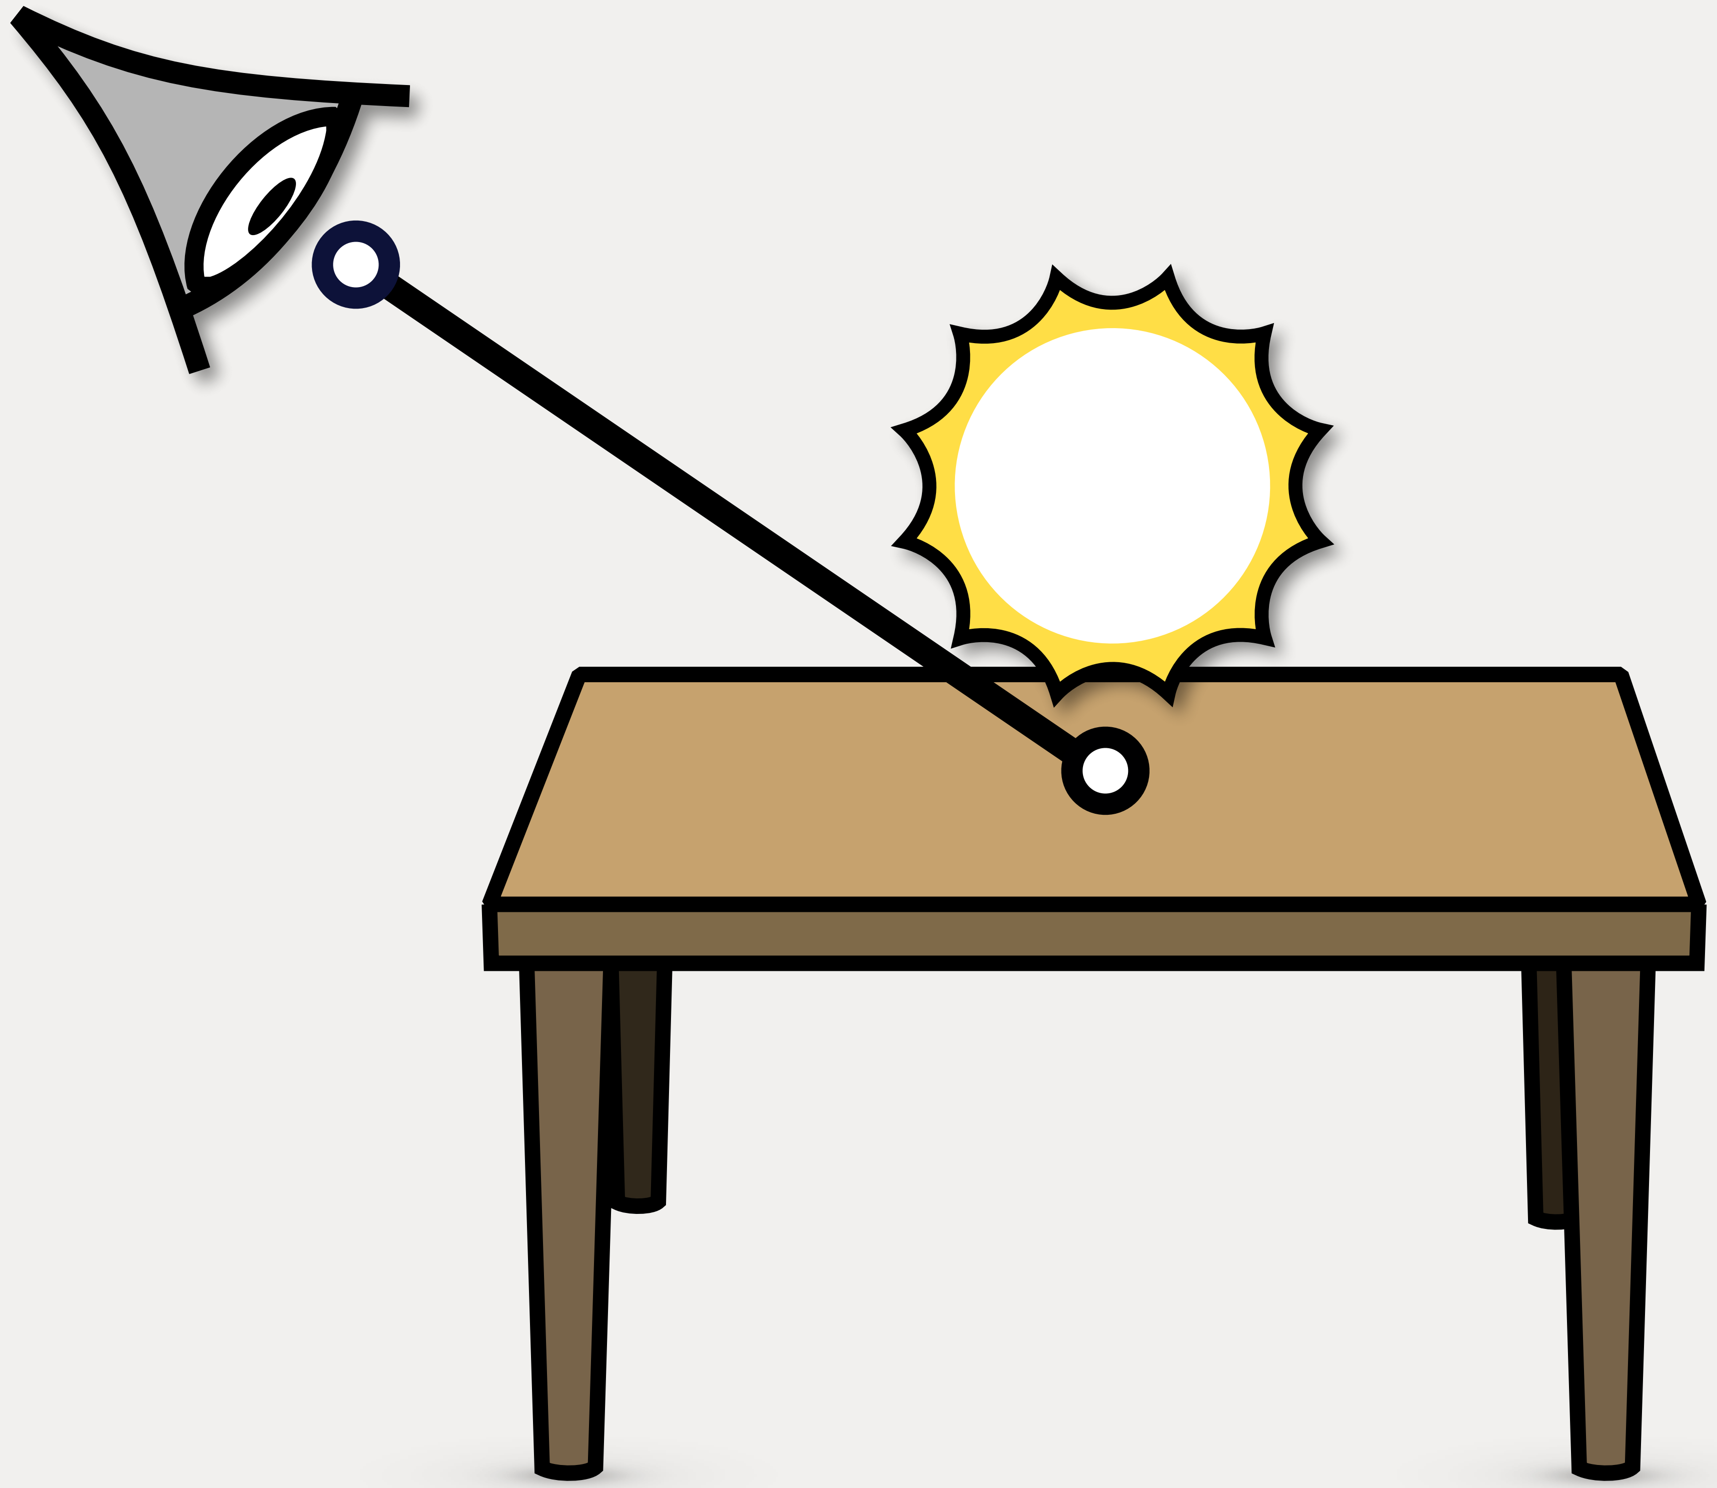

$$f(x) = L_e(x)\rho(x)V(x)\boxed{\frac{\cos\theta_i \cos\theta_o}{r^2}}$$

Problem Description

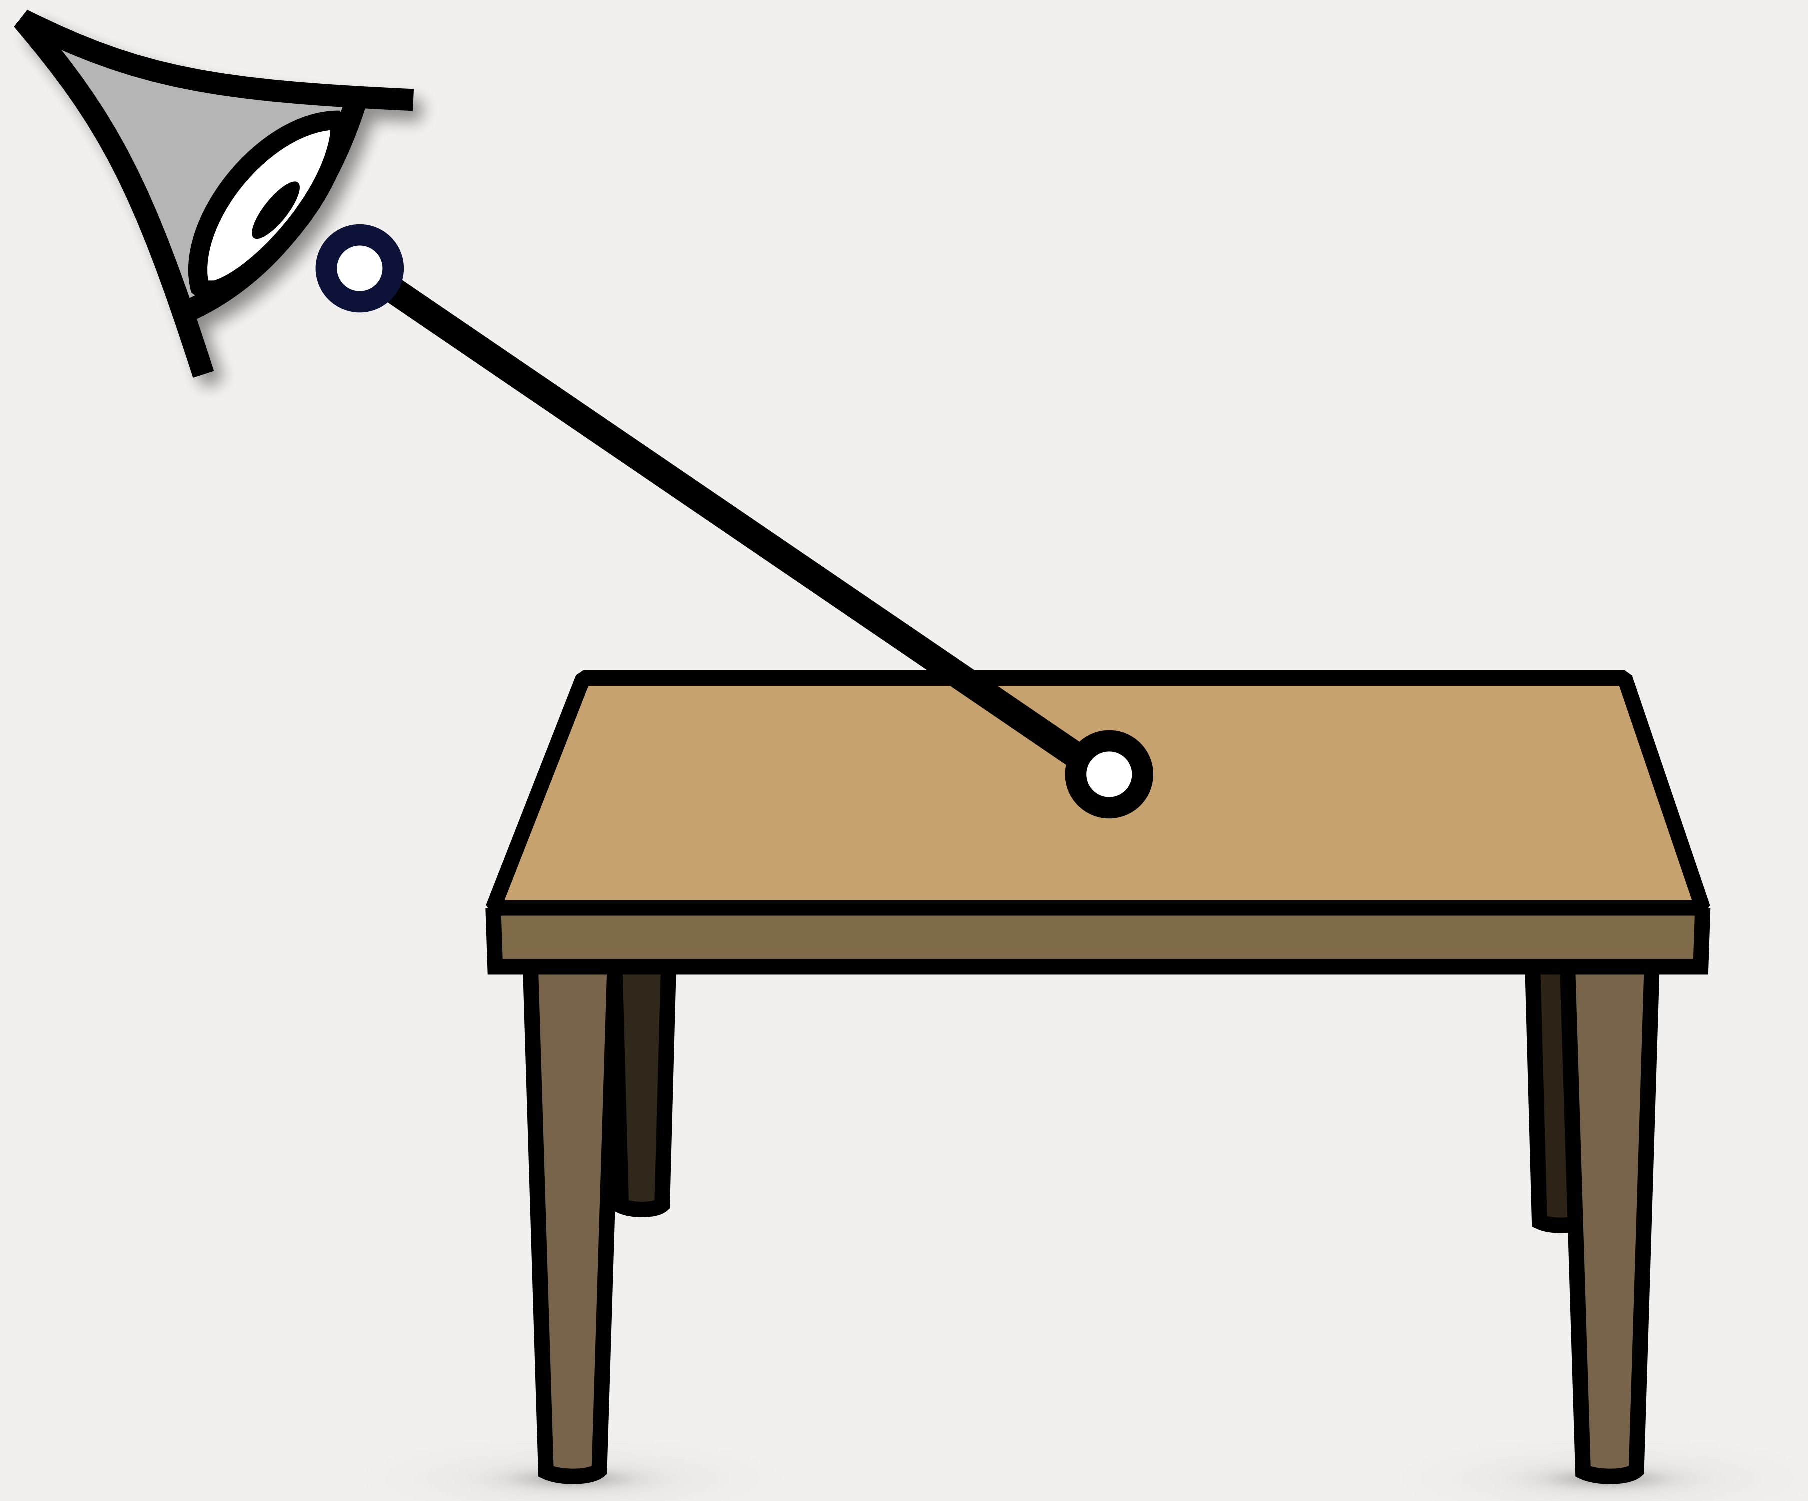

$$f(x) = \boxed{L_e(x)\rho(x)V(x)\frac{\cos\theta_i \cos\theta_o}{r^2}}$$

- Want to sample *product*
- This is hard!

Related Work

BRDF

Visibility

Geometry

Related Work

Donikian et al. 2006

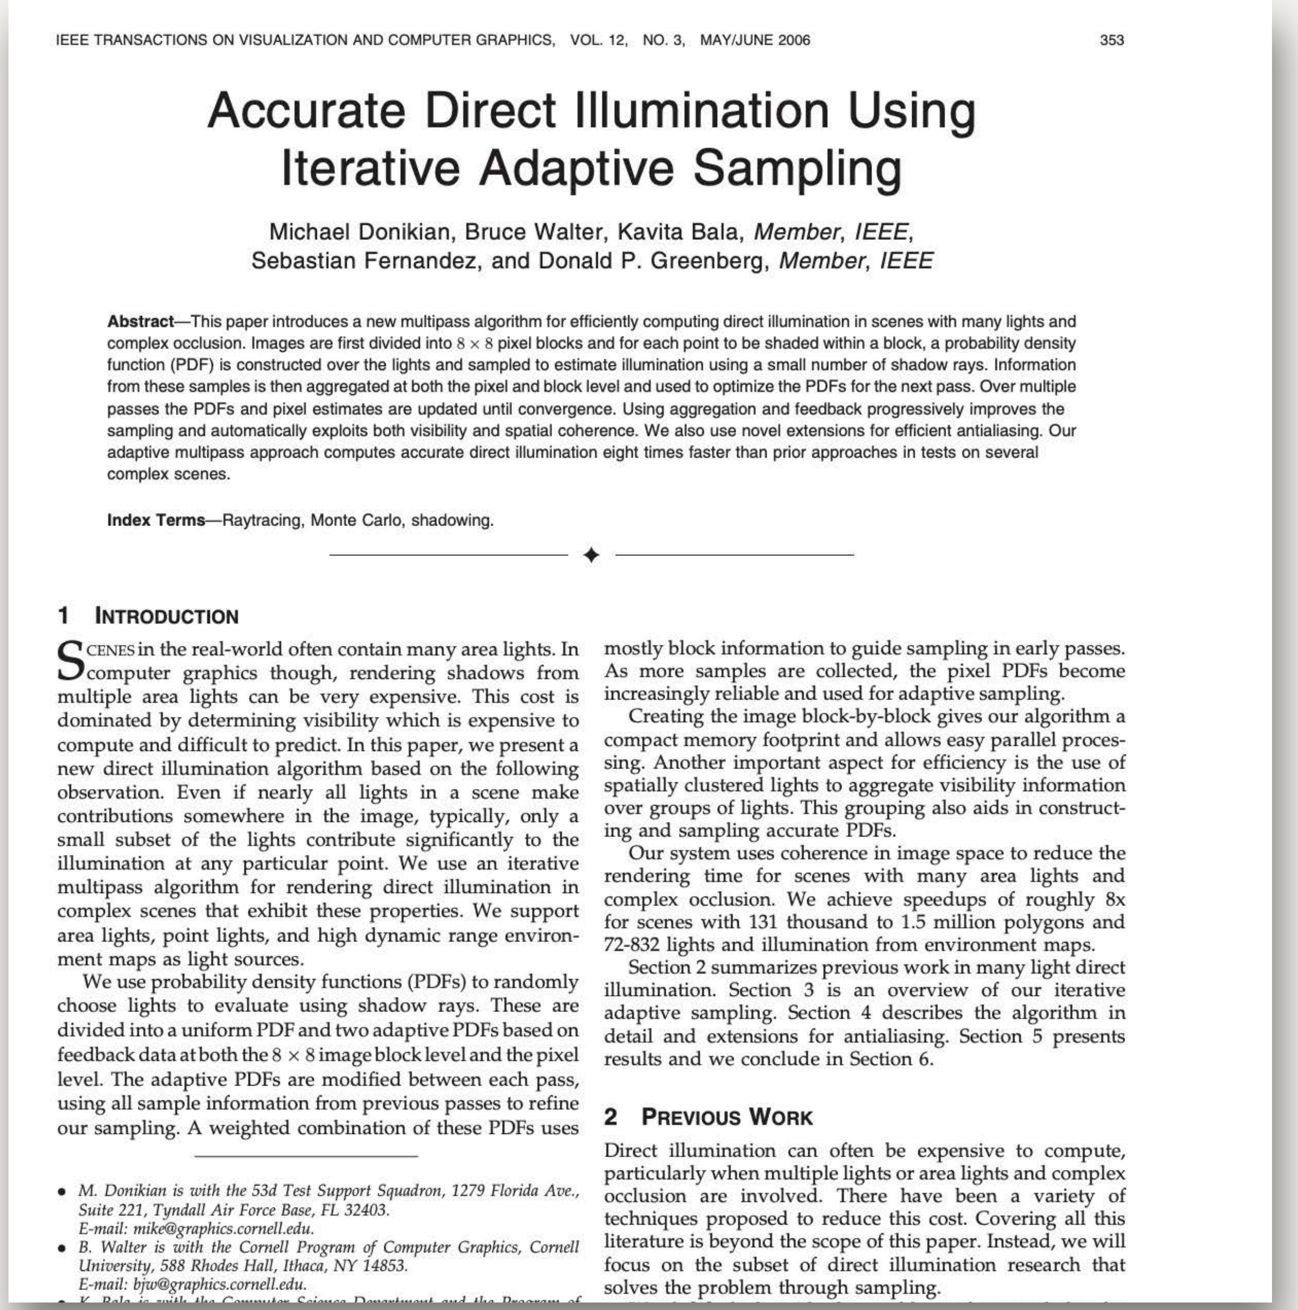

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. 12, NO. 3, MAY/JUNE 2006 353

Accurate Direct Illumination Using Iterative Adaptive Sampling

Michael Donikian, Bruce Walter, Kavita Bala, *Member, IEEE*, Sebastian Fernandez, and Donald P. Greenberg, *Member, IEEE*

Abstract—This paper introduces a new multipass algorithm for efficiently computing direct illumination in scenes with many lights and complex occlusion. Images are first divided into 8×8 pixel blocks and for each point to be shaded within a block, a probability density function (PDF) is constructed over the lights and sampled to estimate illumination using a small number of shadow rays. Information from these samples is then aggregated at both the pixel and block level and used to optimize the PDFs for the next pass. Over multiple passes the PDFs and pixel estimates are updated until convergence. Using aggregation and feedback progressively improves the sampling and automatically exploits both visibility and spatial coherence. We also use novel extensions for efficient antialiasing. Our adaptive multipass approach computes accurate direct illumination eight times faster than prior approaches in tests on several complex scenes.

Index Terms—Raytracing, Monte Carlo, shadowing.

1 INTRODUCTION

Scenes in the real-world often contain many area lights. In computer graphics though, rendering shadows from multiple area lights can be very expensive. This cost is dominated by determining visibility which is expensive to compute and difficult to predict. In this paper, we present a new direct illumination algorithm based on the following observation. Even if nearly all lights in a scene make contributions somewhere in the image, typically, only a small subset of the lights contribute significantly to the illumination at any particular point. We use an iterative multipass algorithm for rendering direct illumination in complex scenes that exhibit these properties. We support area lights, point lights, and high dynamic range environment maps as light sources.

We use probability density functions (PDFs) to randomly choose lights to evaluate using shadow rays. These are divided into a uniform PDF and two adaptive PDFs based on feedback data at both the 8×8 image block level and the pixel level. The adaptive PDFs are modified between each pass, using all sample information from previous passes to refine our sampling. A weighted combination of these PDFs uses mostly block information to guide sampling in early passes. As more samples are collected, the pixel PDFs become increasingly reliable and used for adaptive sampling.

Creating the image block-by-block gives our algorithm a compact memory footprint and allows easy parallel processing. Another important aspect for efficiency is the use of spatially clustered lights to aggregate visibility information over groups of lights. This grouping also aids in constructing and sampling accurate PDFs.

Our system uses coherence in image space to reduce the rendering time for scenes with many area lights and complex occlusion. We achieve speedups of roughly 8x for scenes with 131 thousand to 1.5 million polygons and 72-832 lights and illumination from environment maps.

Section 2 summarizes previous work in many light direct illumination. Section 3 is an overview of our iterative adaptive sampling. Section 4 describes the algorithm in detail and extensions for antialiasing. Section 5 presents results and we conclude in Section 6.

2 PREVIOUS WORK

Direct illumination can often be expensive to compute, particularly when multiple lights or area lights and complex occlusion are involved. There have been a variety of techniques proposed to reduce this cost. Covering all this literature is beyond the scope of this paper. Instead, we will focus on the subset of direct illumination research that solves the problem through sampling.

- *M. Donikian is with the 53d Test Support Squadron, 1279 Florida Ave., Suite 221, Tyndall Air Force Base, FL 32403. E-mail: mike@graphics.cornell.edu.*
- *B. Walter is with the Cornell Program of Computer Graphics, Cornell University, 588 Rhodes Hall, Ithaca, NY 14853. E-mail: bjw@graphics.cornell.edu.*

BRDF	✗
Visibility	✗
Geometry	✓

Related Work

Donikian et al. 2006

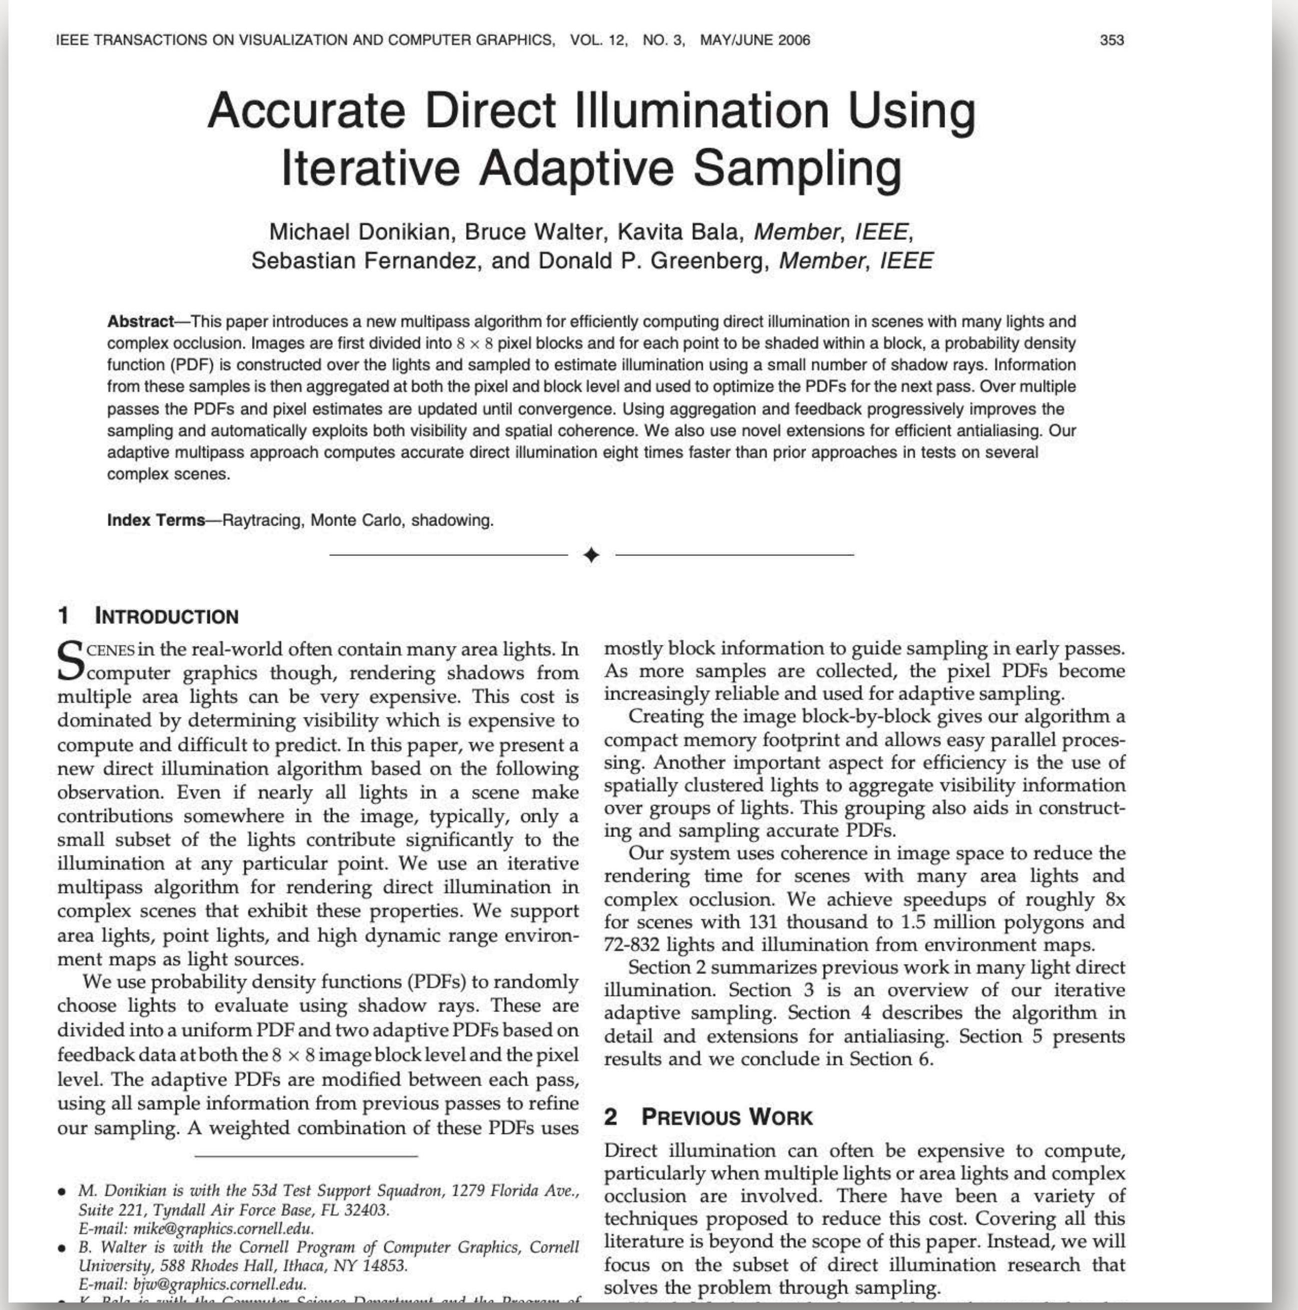

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. 12, NO. 3, MAY/JUNE 2006 353

Accurate Direct Illumination Using Iterative Adaptive Sampling

Michael Donikian, Bruce Walter, Kavita Bala, *Member, IEEE*, Sebastian Fernandez, and Donald P. Greenberg, *Member, IEEE*

Abstract—This paper introduces a new multipass algorithm for efficiently computing direct illumination in scenes with many lights and complex occlusion. Images are first divided into 8×8 pixel blocks and for each point to be shaded within a block, a probability density function (PDF) is constructed over the lights and sampled to estimate illumination using a small number of shadow rays. Information from these samples is then aggregated at both the pixel and block level and used to optimize the PDFs for the next pass. Over multiple passes the PDFs and pixel estimates are updated until convergence. Using aggregation and feedback progressively improves the sampling and automatically exploits both visibility and spatial coherence. We also use novel extensions for efficient antialiasing. Our adaptive multipass approach computes accurate direct illumination eight times faster than prior approaches in tests on several complex scenes.

Index Terms—Raytracing, Monte Carlo, shadowing.

1 Introduction

Scenes in the real-world often contain many area lights. In computer graphics though, rendering shadows from multiple area lights can be very expensive. This cost is dominated by determining visibility which is expensive to compute and difficult to predict. In this paper, we present a new direct illumination algorithm based on the following observation. Even if nearly all lights in a scene make contributions somewhere in the image, typically, only a small subset of the lights contribute significantly to the illumination at any particular point. We use an iterative multipass algorithm for rendering direct illumination in complex scenes that exhibit these properties. We support area lights, point lights, and high dynamic range environment maps as light sources.

We use probability density functions (PDFs) to randomly choose lights to evaluate using shadow rays. These are divided into a uniform PDF and two adaptive PDFs based on feedback data at both the 8×8 image block level and the pixel level. The adaptive PDFs are modified between each pass, using all sample information from previous passes to refine our sampling. A weighted combination of these PDFs uses mostly block information to guide sampling in early passes. As more samples are collected, the pixel PDFs become increasingly reliable and used for adaptive sampling.

Creating the image block-by-block gives our algorithm a compact memory footprint and allows easy parallel processing. Another important aspect for efficiency is the use of spatially clustered lights to aggregate visibility information over groups of lights. This grouping also aids in constructing and sampling accurate PDFs.

Our system uses coherence in image space to reduce the rendering time for scenes with many area lights and complex occlusion. We achieve speedups of roughly 8x for scenes with 131 thousand to 1.5 million polygons and 72-832 lights and illumination from environment maps.

Section 2 summarizes previous work in many light direct illumination. Section 3 is an overview of our iterative adaptive sampling. Section 4 describes the algorithm in detail and extensions for antialiasing. Section 5 presents results and we conclude in Section 6.

2 Previous Work

Direct illumination can often be expensive to compute, particularly when multiple lights or area lights and complex occlusion are involved. There have been a variety of techniques proposed to reduce this cost. Covering all this literature is beyond the scope of this paper. Instead, we will focus on the subset of direct illumination research that solves the problem through sampling.

• *M. Donikian is with the 53d Test Support Squadron, 1279 Florida Ave., Suite 221, Tyndall Air Force Base, FL 32403. E-mail: mike@graphics.cornell.edu.*
• *B. Walter is with the Cornell Program of Computer Graphics, Cornell University, 588 Rhodes Hall, Ithaca, NY 14853. E-mail: bjw@graphics.cornell.edu.*

Burley et al. 2018

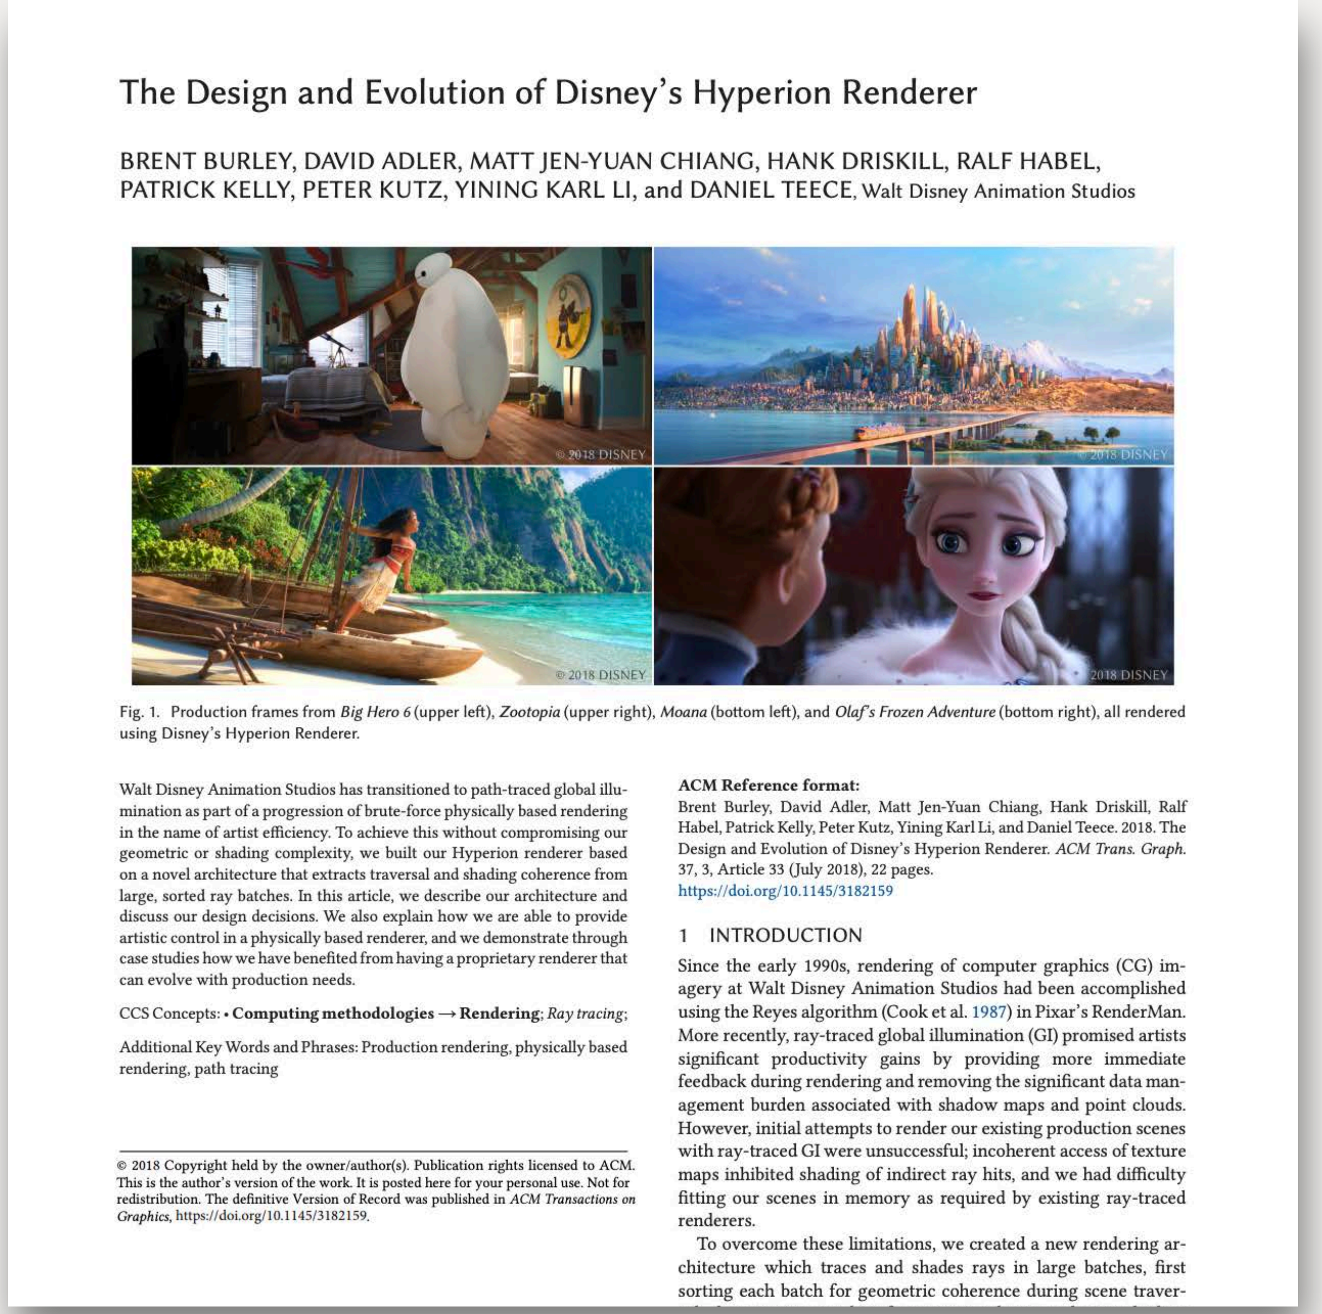

The Design and Evolution of Disney's Hyperion Renderer

BRENT BURLEY, DAVID ADLER, MATT JEN-YUAN CHIANG, HANK DRISKILL, RALF HABEL, PATRICK KELLY, PETER KUTZ, YINING KARL LI, and DANIEL TEECE, Walt Disney Animation Studios

© 2018 DISNEY

Fig. 1. Production frames from *Big Hero 6* (upper left), *Zootopia* (upper right), *Moana* (bottom left), and *Olaf's Frozen Adventure* (bottom right), all rendered using Disney's Hyperion Renderer.

Walt Disney Animation Studios has transitioned to path-traced global illumination as part of a progression of brute-force physically based rendering in the name of artist efficiency. To achieve this without compromising our geometric or shading complexity, we built our Hyperion renderer based on a novel architecture that extracts traversal and shading coherence from large, sorted ray batches. In this article, we describe our architecture and discuss our design decisions. We also explain how we are able to provide artistic control in a physically based renderer, and we demonstrate through case studies how we have benefited from having a proprietary renderer that can evolve with production needs.

CCS Concepts: • **Computing methodologies → Rendering**; *Ray tracing*;

Additional Key Words and Phrases: Production rendering, physically based rendering, path tracing

ACM Reference format:
Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The Design and Evolution of Disney's Hyperion Renderer. *ACM Trans. Graph.* 37, 3, Article 33 (July 2018), 22 pages.
https://doi.org/10.1145/3182159

1 INTRODUCTION

Since the early 1990s, rendering of computer graphics (CG) imagery at Walt Disney Animation Studios had been accomplished using the Reyes algorithm (Cook et al. 1987) in Pixar's RenderMan. More recently, ray-traced global illumination (GI) promised artists significant productivity gains by providing more immediate feedback during rendering and removing the significant data management burden associated with shadow maps and point clouds. However, initial attempts to render our existing production scenes with ray-traced GI were unsuccessful; incoherent access of texture maps inhibited shading of indirect ray hits, and we had difficulty fitting our scenes in memory as required by existing ray-traced renderers.

To overcome these limitations, we created a new rendering architecture which traces and shades rays in large batches, first sorting each batch for geometric coherence during scene traver-

	Donikian et al. 2006	Burley et al. 2018
BRDF	✗	✗
Visibility	✗	✓
Geometry	✓	✓

Related Work

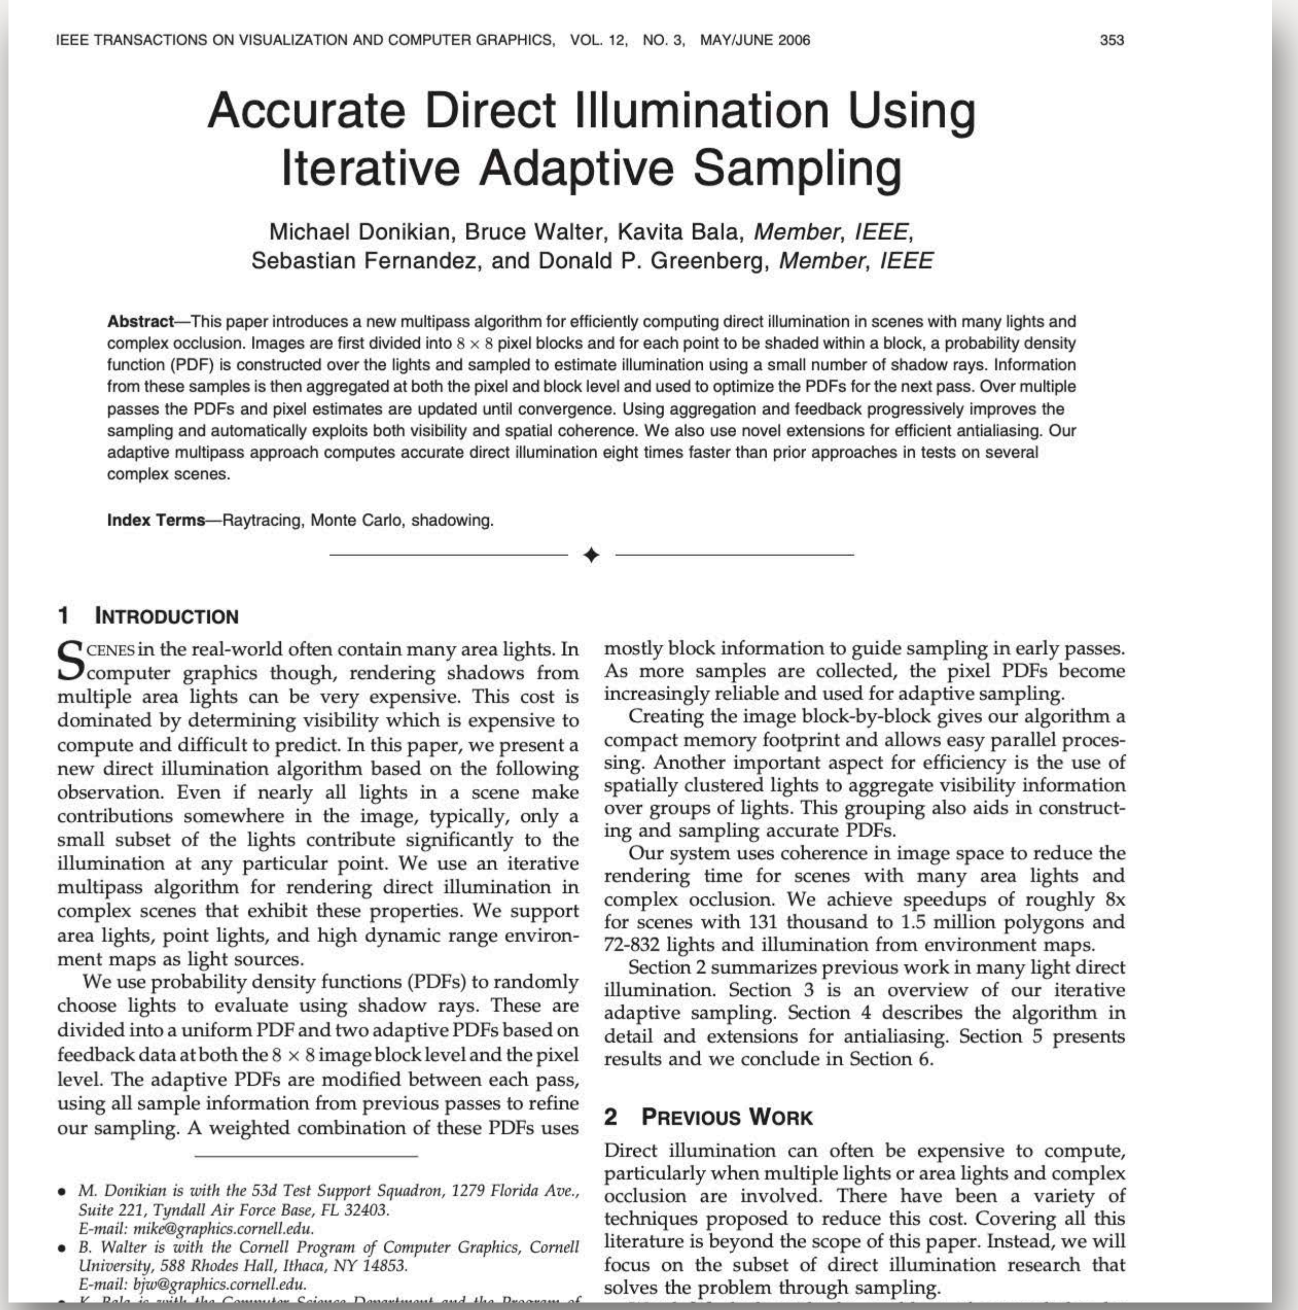

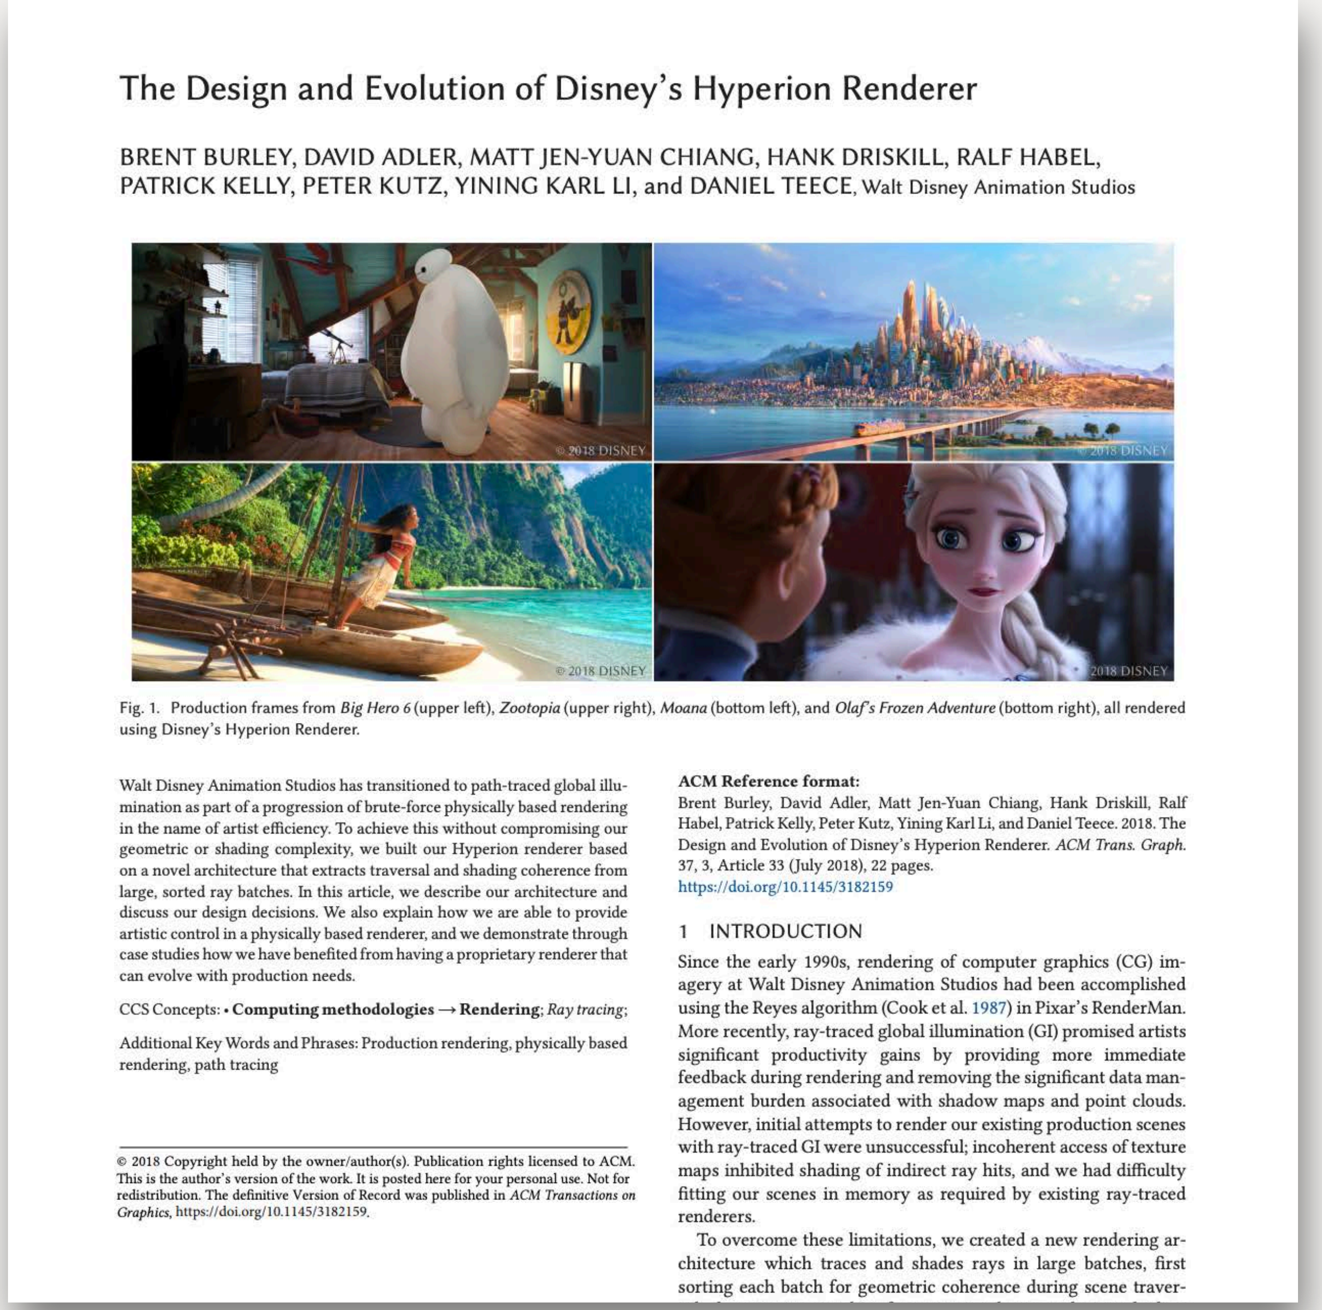

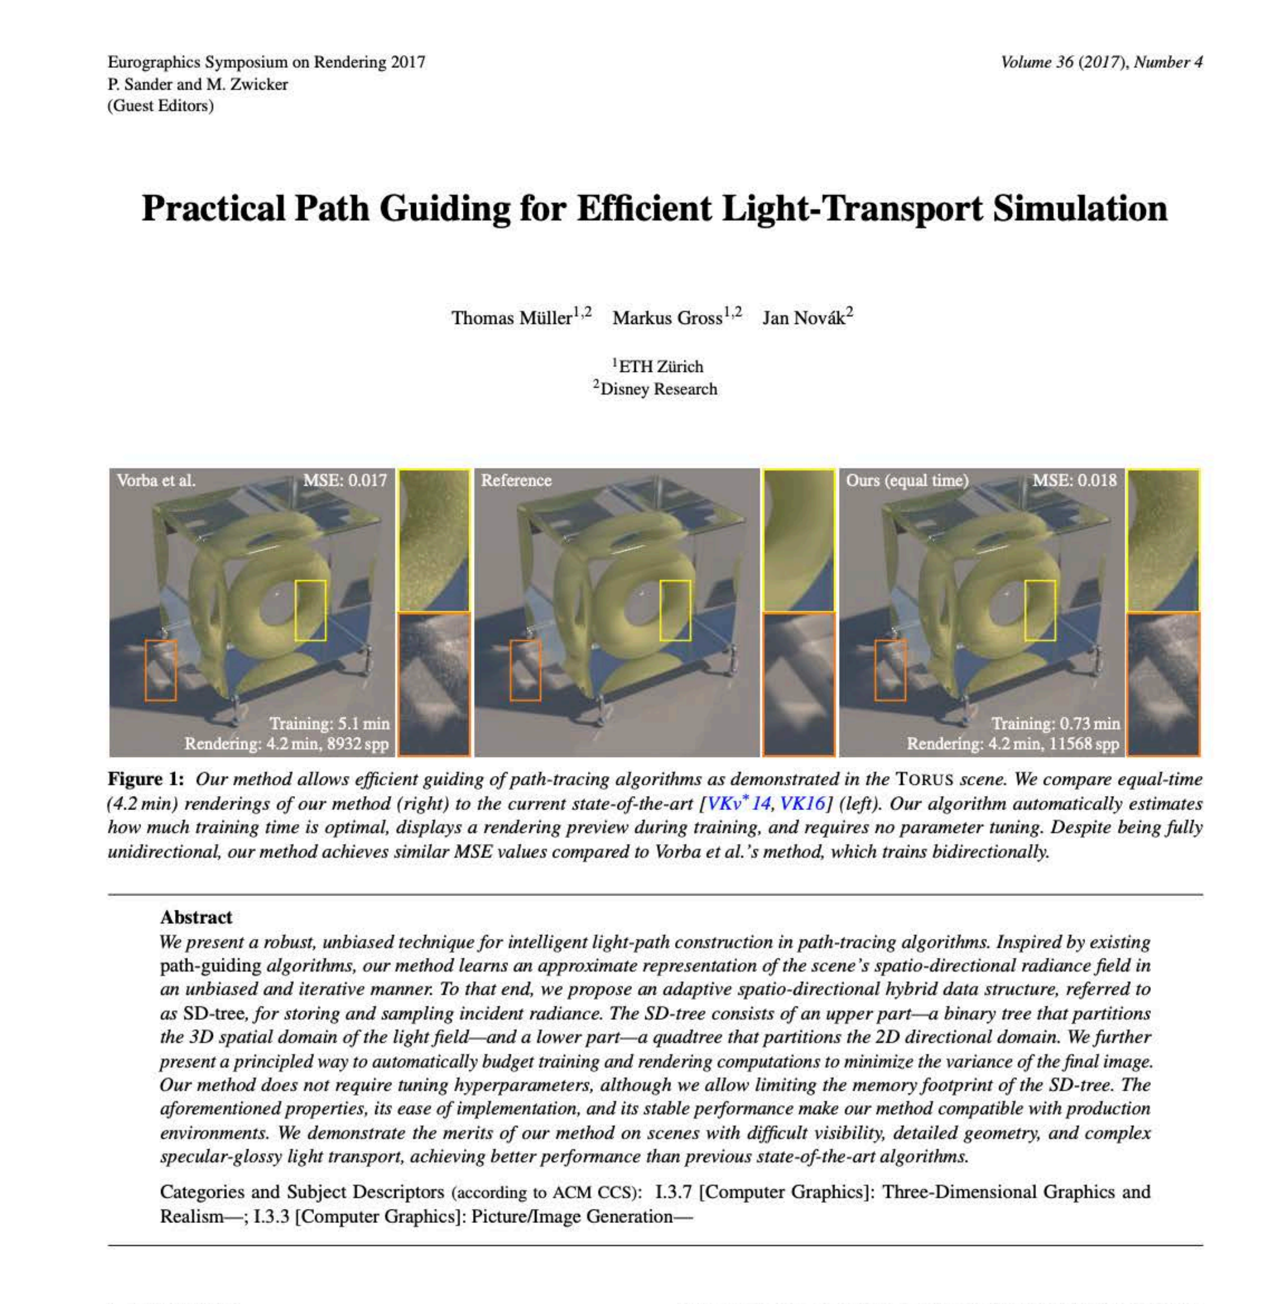

	Donikian et al. 2006	**Burley et al.** 2018	**Müller et al.** 2017
BRDF	✗	✗	✓
Visibility	✗	✓	✓
Geometry	✓	✓	✓

Related Work

BRDF

Visibility

Geometry

Related Work

Vévoda et al. 2018

Bayesian online regression for adaptive direct illumination sampling

PETR VÉVODA*, Charles University, Prague and Render Legion, a. s.
IVO KONDAPANENI*, Charles University, Prague
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

No learning
RMSE 0.0031 (0.0014 for DI only)

Our method
RMSE 0.0012, *6.7x speedup* (0.000062, *510x* for DI only)

No learning

Our method

Fig. 1. Equal-time comparison (60 s) of path-traced global illumination solutions computed using our learning-based direct illumination sampling method (right) and a baseline sampling method without learning (left). While both methods start off by sampling lights proportionally to rough estimates of their unoccluded contribution, our method progressively incorporates information about their actual contributions, including visibility, dramatically reducing image variance.

Direct illumination calculation is an important component of any physically-based renderer with a substantial impact on the overall performance. We present a novel adaptive solution for unbiased Monte Carlo direct illumination sampling, based on online learning of the light selection probability distributions. Our main contribution is a formulation of the learning process as Bayesian regression, based on a new, specifically designed statistical model of direct illumination. The net result is a set of regularization strategies to prevent over-fitting and ensure robustness even in early stages of calculation, when the observed information is sparse. The regression model captures spatial variation of illumination, which enables aggregating statistics over relatively large scene regions and, in turn, ensures a fast learning rate. We make the method scalable by adopting a light clustering strategy from the Lightcuts method, and further reduce variance through the use of control variates. As a main design feature, the resulting algorithm is virtually free of any preprocessing, which enables its use for interactive progressive rendering, while the online learning still enables super-linear convergence.

CCS Concepts: • **Computing methodologies → Rendering**; *Visibility*; *Machine learning*;

Authors' addresses: Petr Vévoda, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, Render Legion, a. s. Karlovo náměstí 288/17, Prague, 12000, petrvevoda@seznam.cz; Ivo Kondapaneni, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, ivo.kondapaneni@gmail.com; Jaroslav Křivánek, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, Render Legion, a. s. Karlovo náměstí 288/17, Prague, 12000, jaroslav.krivanek@mff.cuni.cz.

*Petr Vévoda and Ivo Kondapaneni share the first authorship of this work.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a

Additional Key Words and Phrases: direct illumination, adaptive sampling, visibility, learning.

ACM Reference Format:
Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Trans. Graph.* 37, 4, Article 125 (August 2018), 12 pages. https://doi.org/10.1145/3197517.3201340

1 INTRODUCTION

Realistic rendering today is almost entirely based on Monte Carlo (MC) methods. The *indirect* illumination component has traditionally been held responsible for the undesirable image noise produced by such algorithms, which is probably why the *direct* illumination has received disproportionately less attention in research. However, many scenes in digital production feature complex lighting setups, and practical experience shows that it is often direct illumination that is responsible for the majority of image noise.

In this paper, we aim at unbiased direct illumination estimation for MC renderers. Specifically, we address the problem of randomly choosing an appropriate light source for a given scene location, so that variance of the direct illumination estimator is minimized. This could be achieved by choosing lights with probability proportional to their respective contributions, but these are unknown at the outset, they are costly to evaluate and difficult to predict. This is true especially due to the visibility, since it can be discontinuous and its evaluation involves expensive ray casting.

One possible solution would involve constructing the light sampling distributions in a preprocessing step [Georgiev et al. 2012]. However, long preprocessing disqualifies any form of *interactive*

BRDF	✗
Visibility	✓
Geometry	✓

Related Work

Vévoda et al. 2018 **Estevez & Kulla** 2018

Bayesian online regression for adaptive direct illumination sampling

PETR VÉVODA*, Charles University, Prague and Render Legion, a. s.
IVO KONDAPANENI*, Charles University, Prague
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

Fig. 1. Equal-time comparison (60 s) of path-traced global illumination solutions computed using our learning-based direct illumination sampling method (right) and a baseline sampling method without learning (left). While both methods start off by sampling lights proportionally to rough estimates of their unoccluded contribution, our method progressively incorporates information about their actual contributions, including visibility, dramatically reducing image variance.

Direct illumination calculation is an important component of any physically-based renderer with a substantial impact on the overall performance. We present a novel adaptive solution for unbiased Monte Carlo direct illumination sampling, based on online learning of the light selection probability distributions. Our main contribution is a formulation of the learning process as Bayesian regression, based on a new, specifically designed statistical model of direct illumination. The net result is a set of regularization strategies to prevent over-fitting and ensure robustness even in early stages of calculation, when the observed information is sparse. The regression model captures spatial variation of illumination, which enables aggregating statistics over relatively large scene regions and, in turn, ensures a fast learning rate. We make the method scalable by adopting a light clustering strategy from the Lightcuts method, and further reduce variance through the use of control variates. As a main design feature, the resulting algorithm is virtually free of any preprocessing, which enables its use for interactive progressive rendering, while the online learning still enables super-linear convergence.

CCS Concepts: • **Computing methodologies** → **Rendering**; *Visibility*; *Machine learning*;

Authors' addresses: Petr Vévoda, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, Render Legion, a. s. Karlovo náměstí 288/17, Prague, 12000, petrvevoda@seznam.cz; Ivo Kondapaneni, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, ivo.kondapaneni@gmail.com; Jaroslav Křivánek, Charles University, Prague, Malostranské náměstí 25, Prague, 11800, Render Legion, a. s. Karlovo náměstí 288/17, Prague, 12000, jaroslav.krivanek@mff.cuni.cz.

*Petr Vévoda and Ivo Kondapaneni share the first authorship of this work.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a

Additional Key Words and Phrases: direct illumination, adaptive sampling, visibility, learning.

ACM Reference Format:
Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek. 2018. Bayesian online regression for adaptive direct illumination sampling. *ACM Trans. Graph.* 37, 4, Article 125 (August 2018), 12 pages. https://doi.org/10.1145/3197517.3201340

1 INTRODUCTION

Realistic rendering today is almost entirely based on Monte Carlo (MC) methods. The *indirect* illumination component has traditionally been held responsible for the undesirable image noise produced by such algorithms, which is probably why the *direct* illumination has received disproportionately less attention in research. However, many scenes in digital production feature complex lighting setups, and practical experience shows that it is often direct illumination that is responsible for the majority of image noise.

In this paper, we aim at unbiased direct illumination estimation for MC renderers. Specifically, we address the problem of randomly choosing an appropriate light source for a given scene location, so that variance of the direct illumination estimator is minimized. This could be achieved by choosing lights with probability proportional to their respective contributions, but these are unknown at the outset, they are costly to evaluate and difficult to predict. This is true especially due to the visibility, since it can be discontinuous and its evaluation involves expensive ray casting.

One possible solution would involve constructing the light sampling distributions in a preprocessing step [Georgiev et al. 2012]. However, long preprocessing disqualifies any form of *interactive*

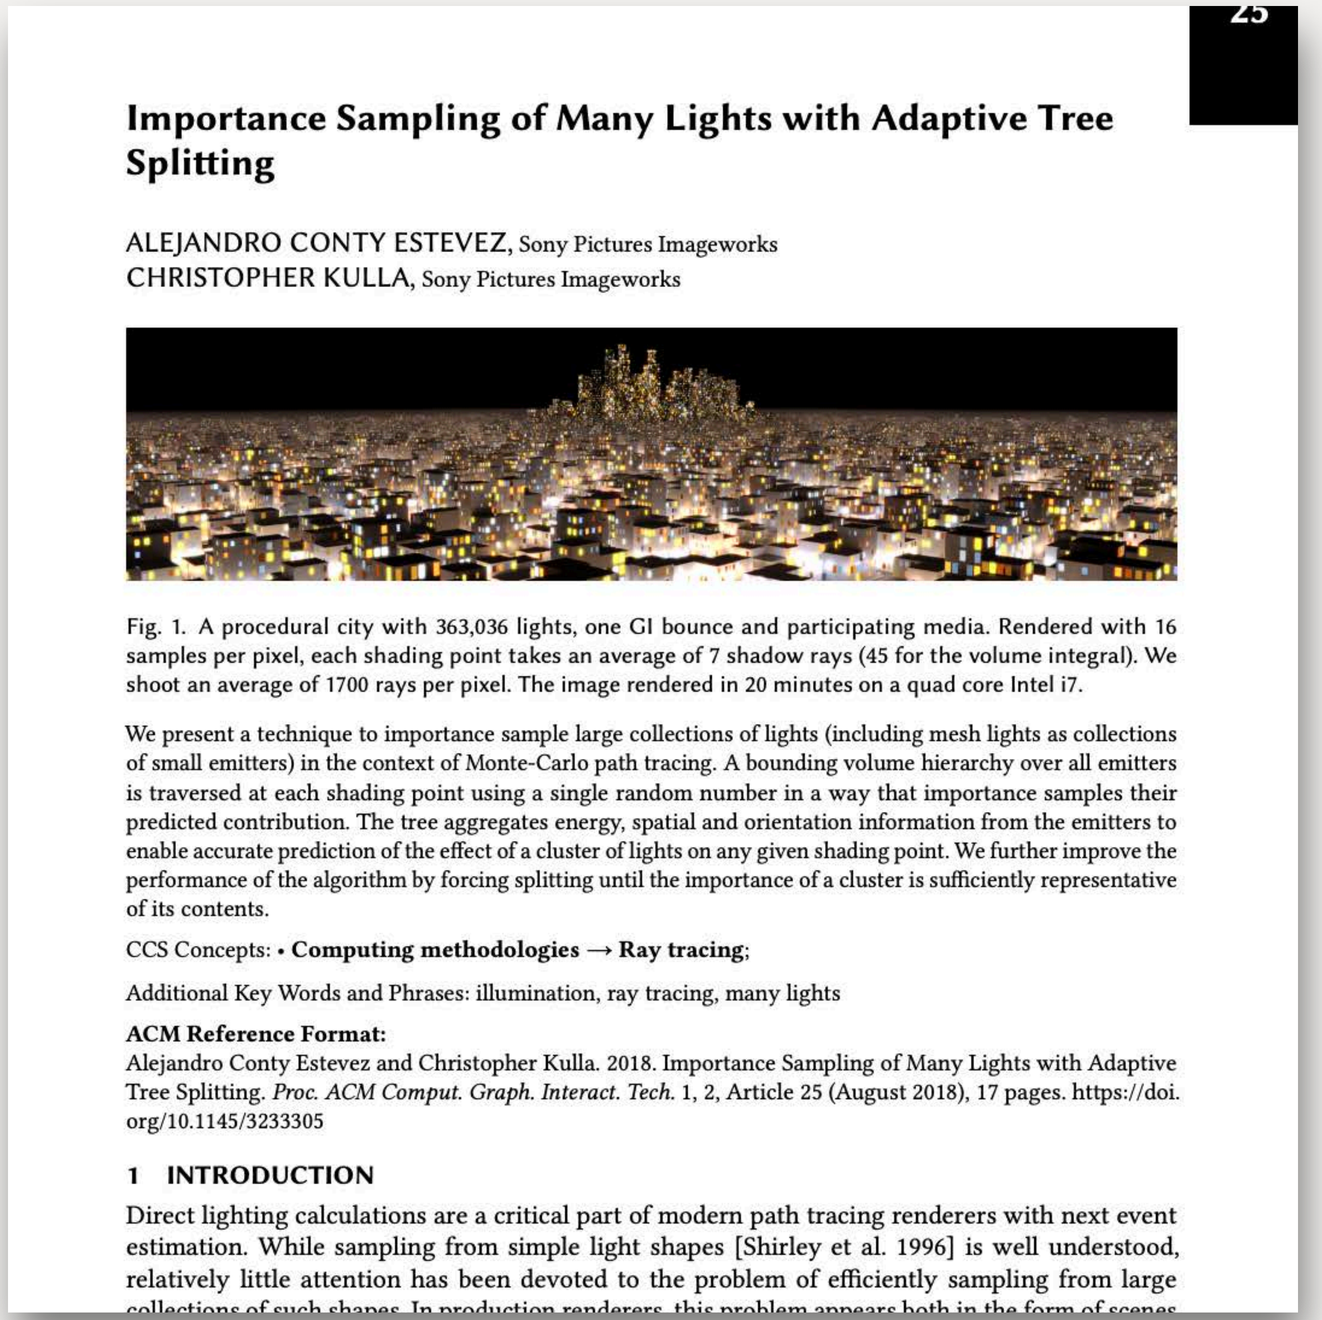

25

Importance Sampling of Many Lights with Adaptive Tree Splitting

ALEJANDRO CONTY ESTEVEZ, Sony Pictures Imageworks
CHRISTOPHER KULLA, Sony Pictures Imageworks

Fig. 1. A procedural city with 363,036 lights, one GI bounce and participating media. Rendered with 16 samples per pixel, each shading point takes an average of 7 shadow rays (45 for the volume integral). We shoot an average of 1700 rays per pixel. The image rendered in 20 minutes on a quad core Intel i7.

We present a technique to importance sample large collections of lights (including mesh lights as collections of small emitters) in the context of Monte-Carlo path tracing. A bounding volume hierarchy over all emitters is traversed at each shading point using a single random number in a way that importance samples their predicted contribution. The tree aggregates energy, spatial and orientation information from the emitters to enable accurate prediction of the effect of a cluster of lights on any given shading point. We further improve the performance of the algorithm by forcing splitting until the importance of a cluster is sufficiently representative of its contents.

CCS Concepts: • **Computing methodologies** → **Ray tracing**;

Additional Key Words and Phrases: illumination, ray tracing, many lights

ACM Reference Format:
Alejandro Conty Estevez and Christopher Kulla. 2018. Importance Sampling of Many Lights with Adaptive Tree Splitting. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 2, Article 25 (August 2018), 17 pages. https://doi.org/10.1145/3233305

1 INTRODUCTION

Direct lighting calculations are a critical part of modern path tracing renderers with next event estimation. While sampling from simple light shapes [Shirley et al. 1996] is well understood, relatively little attention has been devoted to the problem of efficiently sampling from large

	Vévoda et al. 2018	Estevez & Kulla 2018
BRDF	✗	✗
Visibility	✓	✗
Geometry	✓	✓

Related Work

Bayesian online regression for adaptive direct illumination sampling

PETR VÉVODA*, Charles University, Prague and Render Legion, a. s.
IVO KONDAPANENI*, Charles University, Prague
JAROSLAV KŘIVÁNEK, Charles University, Prague and Render Legion, a. s.

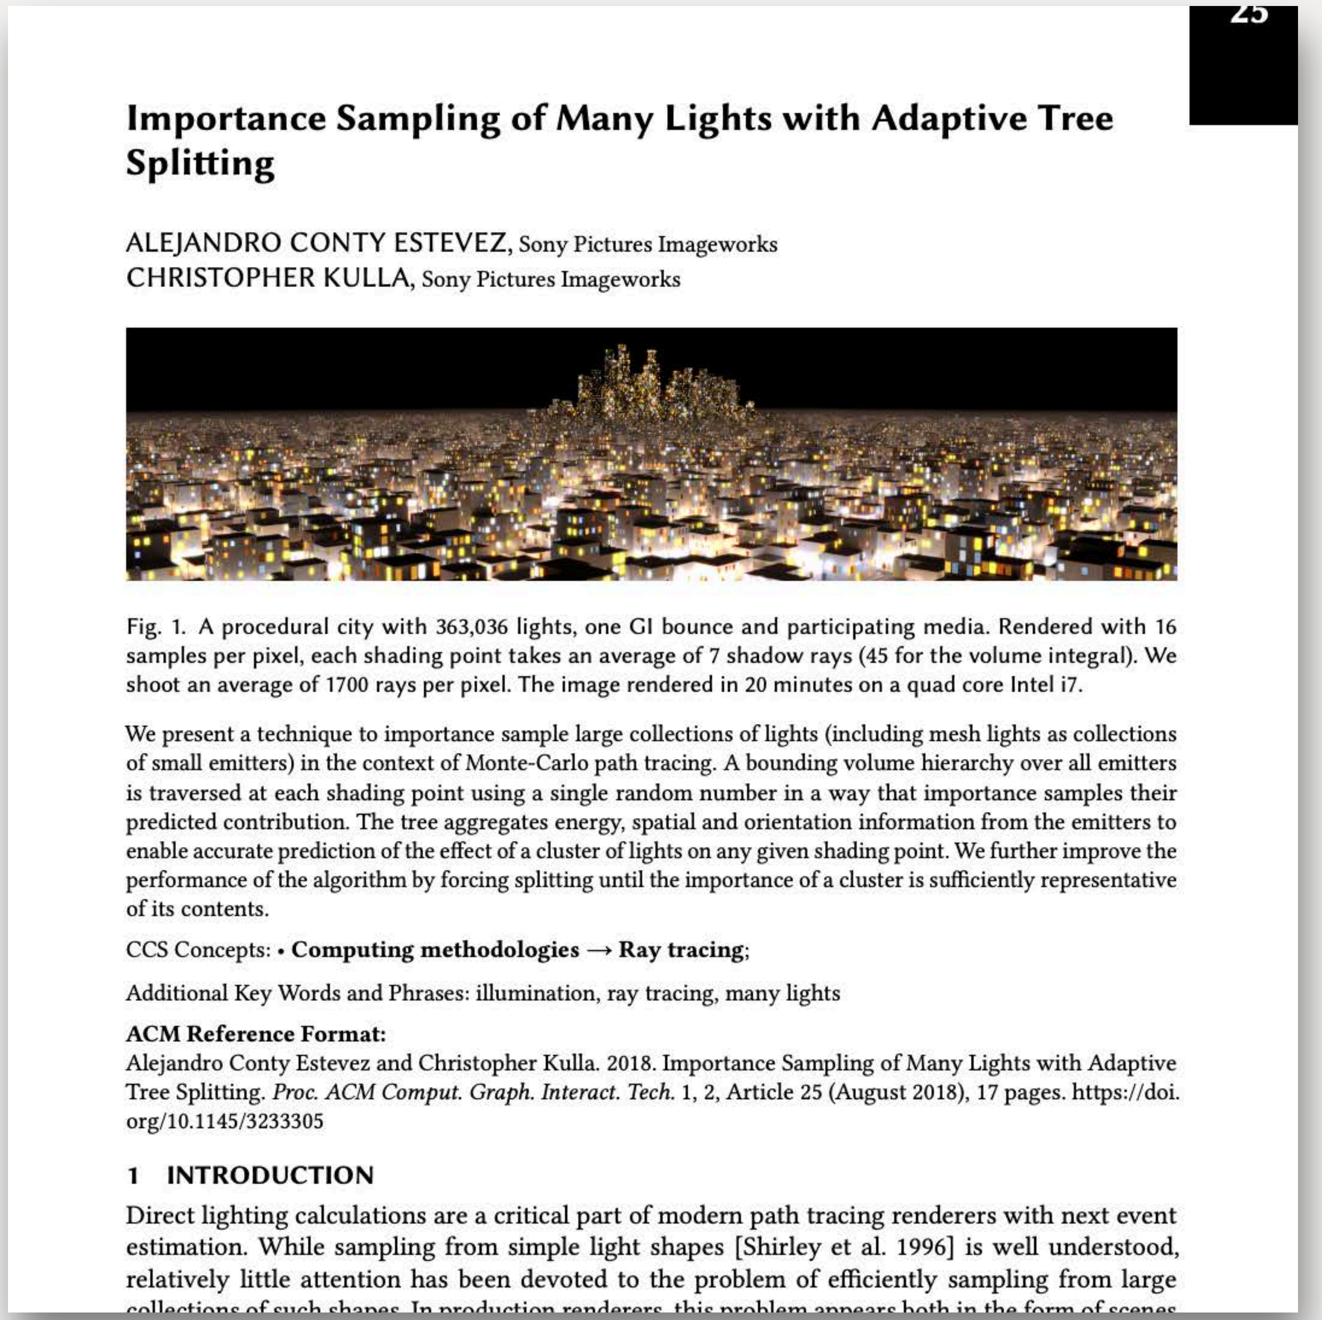

Importance Sampling of Many Lights with Adaptive Tree Splitting

ALEJANDRO CONTY ESTEVEZ, Sony Pictures Imageworks
CHRISTOPHER KULLA, Sony Pictures Imageworks

Fig. 1. A procedural city with 363,036 lights, one GI bounce and participating media. Rendered with 16 samples per pixel, each shading point takes an average of 7 shadow rays (45 for the volume integral). We shoot an average of 1700 rays per pixel. The image rendered in 20 minutes on a quad core Intel i7.

We present a technique to importance sample large collections of lights (including mesh lights as collections of small emitters) in the context of Monte-Carlo path tracing. A bounding volume hierarchy over all emitters is traversed at each shading point using a single random number in a way that importance samples their predicted contribution. The tree aggregates energy, spatial and orientation information from the emitters to enable accurate prediction of the effect of a cluster of lights on any given shading point. We further improve the performance of the algorithm by forcing splitting until the importance of a cluster is sufficiently representative of its contents.

CCS Concepts: • **Computing methodologies → Ray tracing**;

Additional Key Words and Phrases: illumination, ray tracing, many lights

ACM Reference Format:
Alejandro Conty Estevez and Christopher Kulla. 2018. Importance Sampling of Many Lights with Adaptive Tree Splitting. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 2, Article 25 (August 2018), 17 pages. https://doi.org/10.1145/3233305

1 INTRODUCTION

Direct lighting calculations are a critical part of modern path tracing renderers with next event estimation. While sampling from simple light shapes [Shirley et al. 1996] is well understood, relatively little attention has been devoted to the problem of efficiently sampling from large

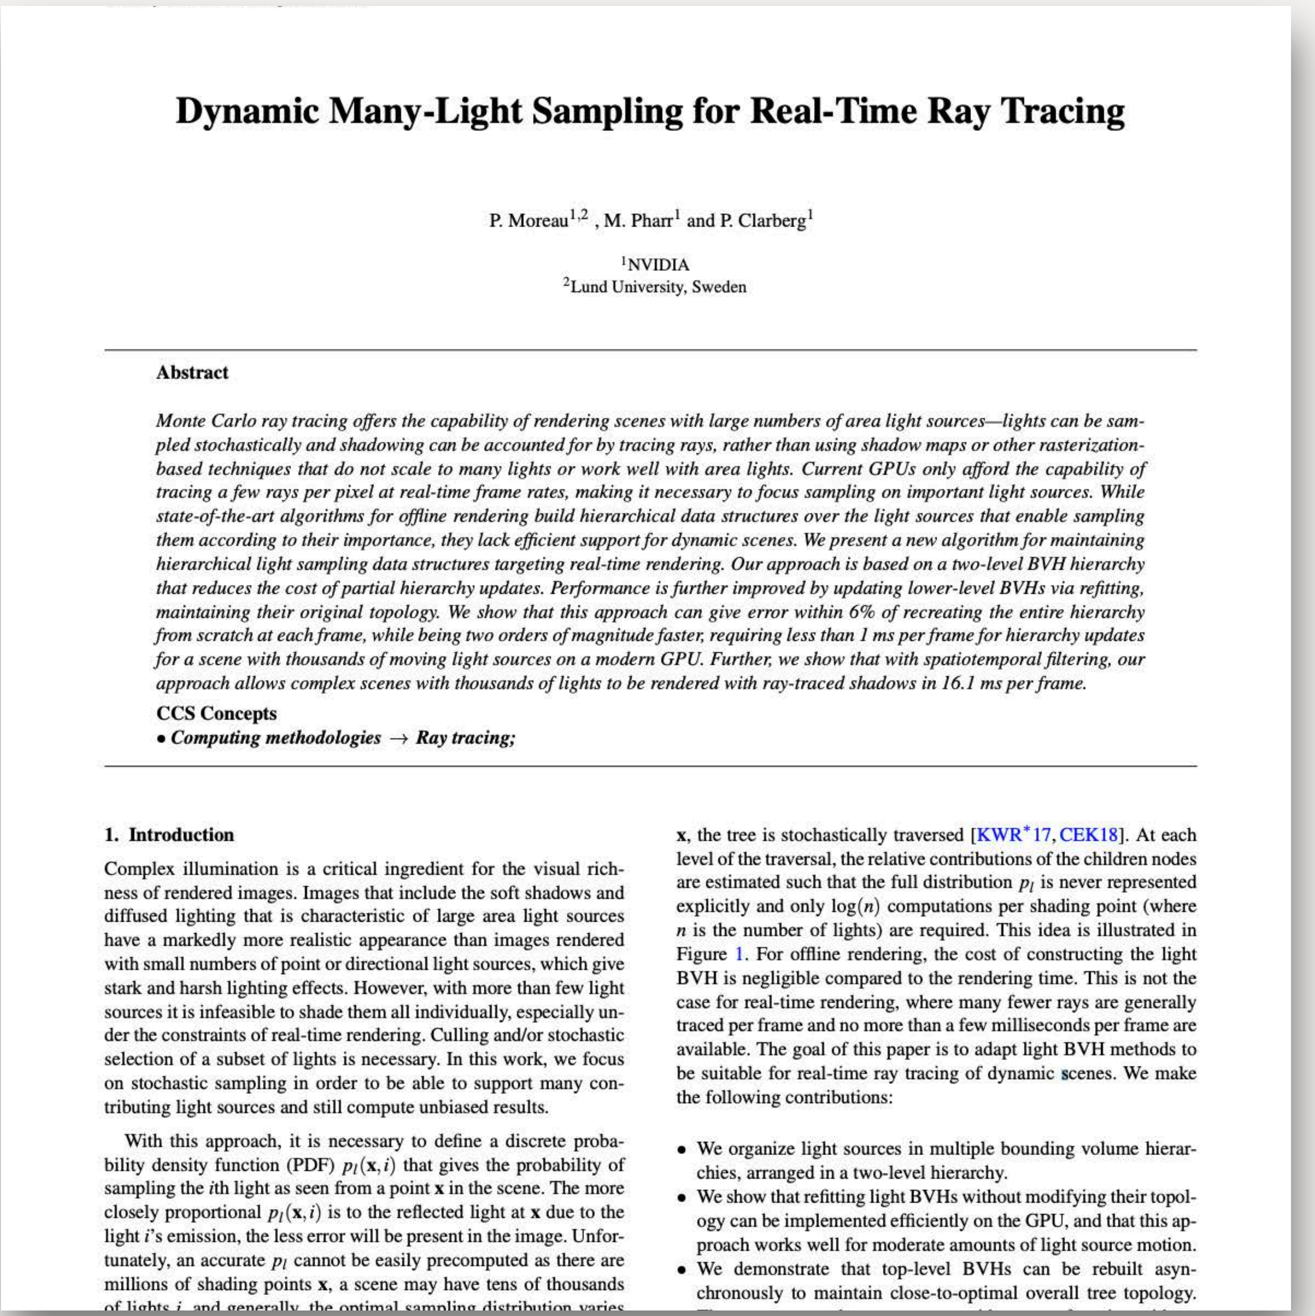

Dynamic Many-Light Sampling for Real-Time Ray Tracing

P. Moreau[1,2], M. Pharr[1] and P. Clarberg[1]

[1]NVIDIA
[2]Lund University, Sweden

	Vévoda et al. 2018	Estevez & Kulla 2018	Moreau et al. 2019
BRDF	✗	✗	✗
Visibility	✓	✗	✗
Geometry	✓	✓	✓

Related Work

- Overhead at low sample counts
- Sample only parts of product
- Maintain complex data structures

ReSTIR Overview

ReSTIR Overview

- Approximately sample full product

ReSTIR Overview

- Approximately sample full product
- No complex data structures

ReSTIR Overview

- Approximately sample full product
- No complex data structures
- Constant memory

ReSTIR Overview

- Approximately sample full product
- No complex data structures
- Constant memory
- Fixed computation per frame

ReSTIR Overview

- Approximately sample full product
- No complex data structures
- Constant memory
- Fixed computation per frame
- *Simple* algorithm

Method

Importance Resampling (RIS)

[Talbot et al. '05]

- Approximately sample *arbitrary* target distribution $\hat{p}(x)$
- Can be unnormalized

Importance Resampling (RIS)

[Talbot et al. '05]

- Start with *candidate* distribution $p(x)$…

Importance Resampling (RIS)

[Talbot et al. '05]

- Start with *candidate* distribution $p(x)$…

M "candidates" $x_i \propto p(x_i)$

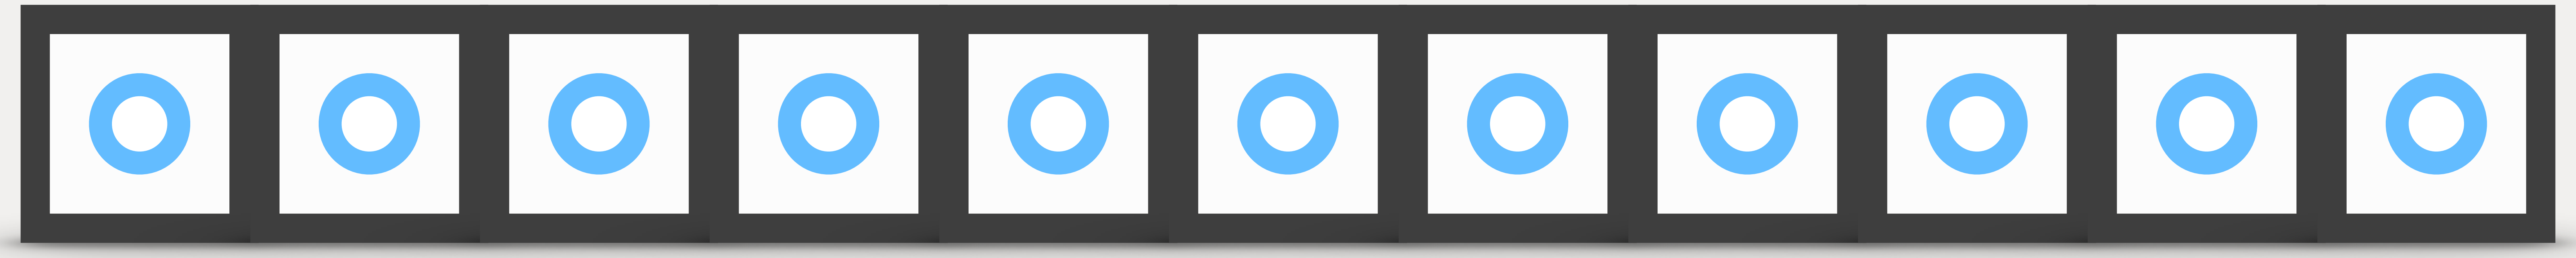

Importance Resampling (RIS)

[Talbot et al. '05]

- Start with *candidate* distribution $p(x)$...

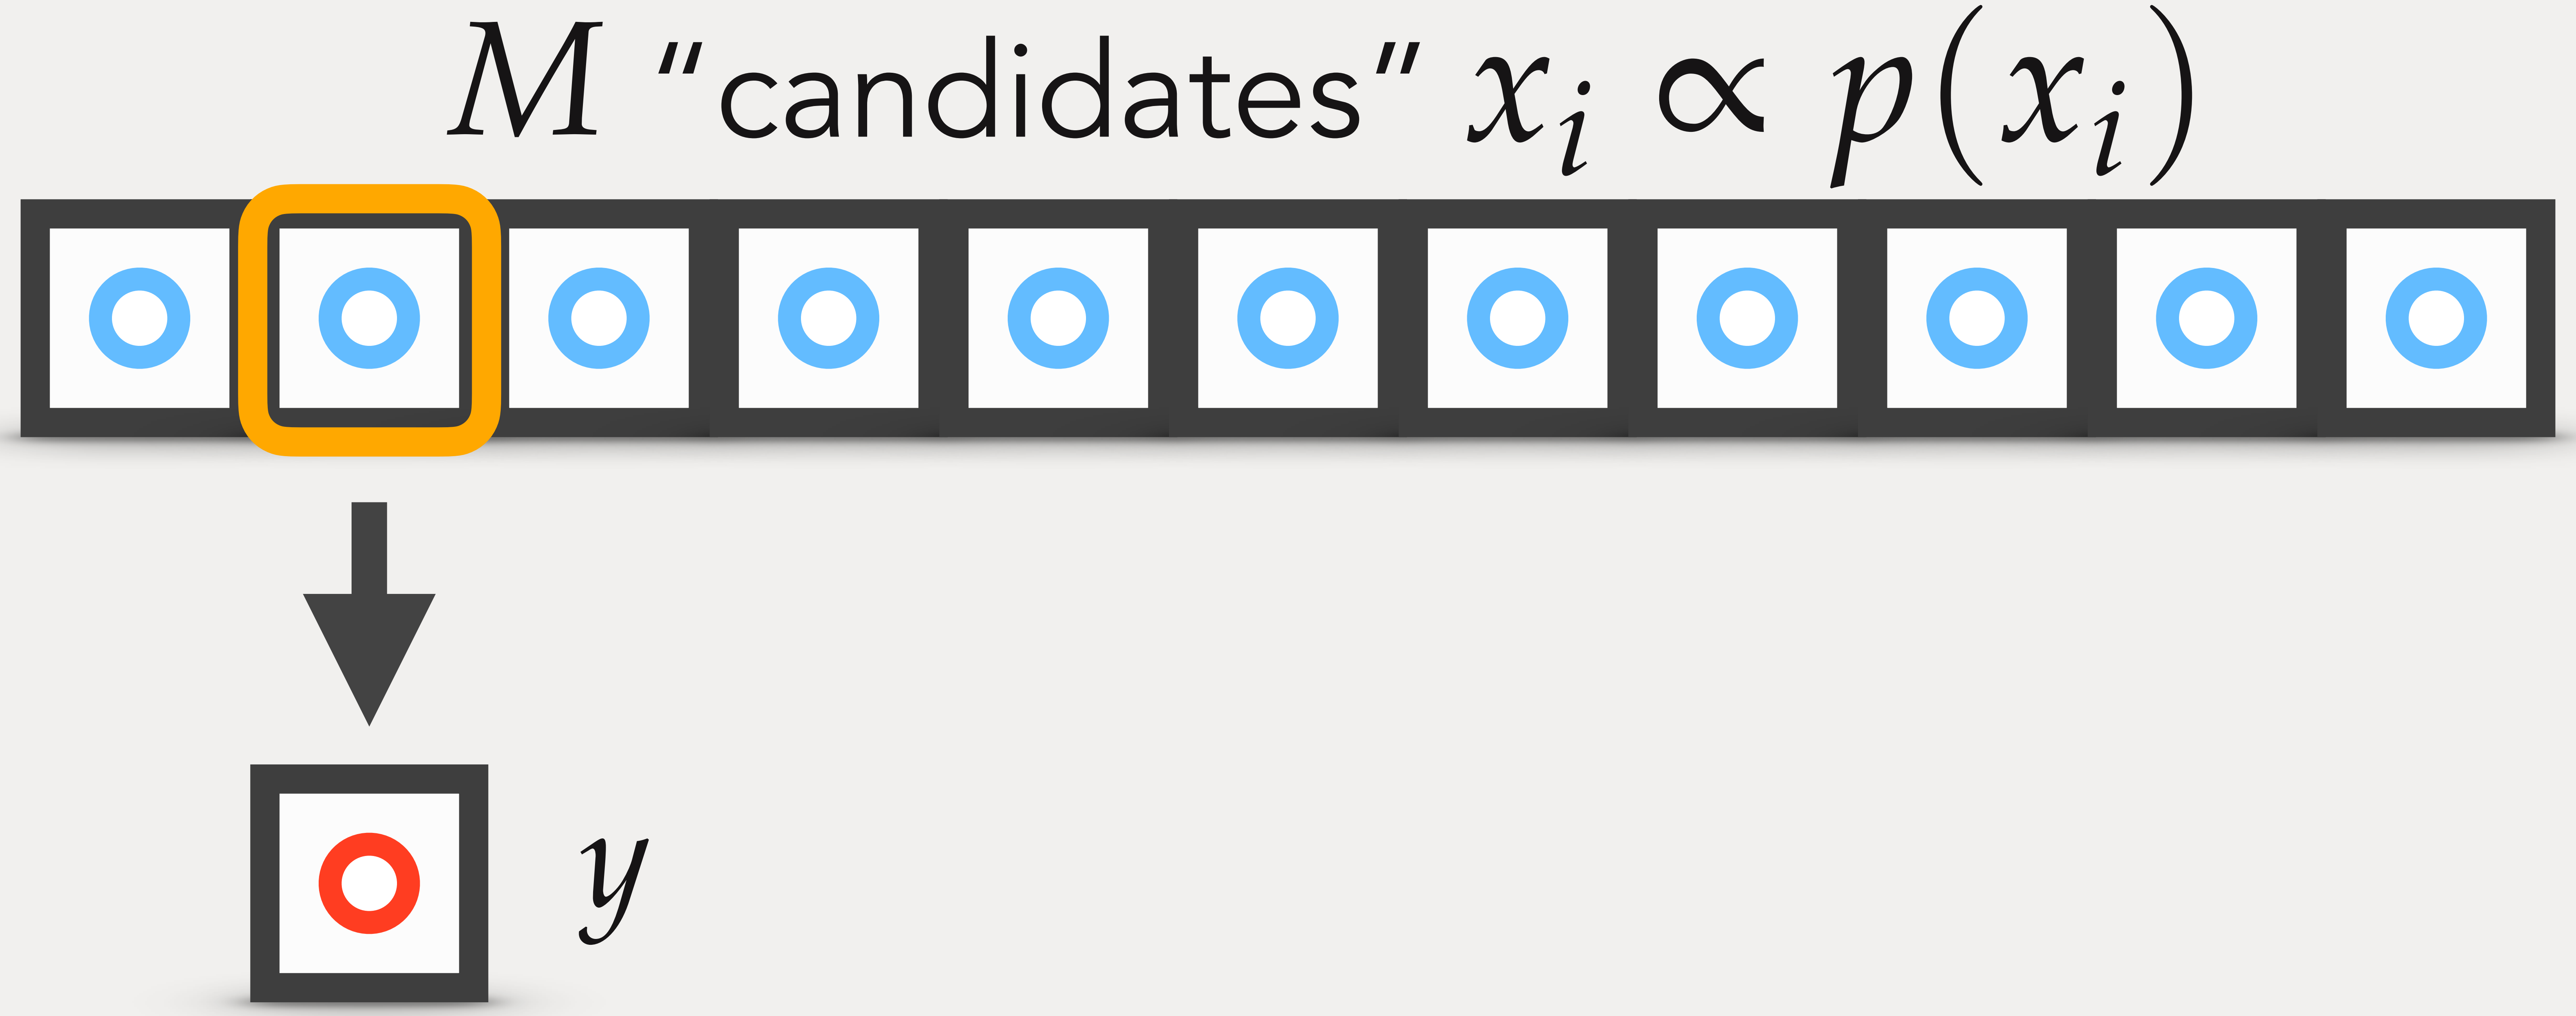

Importance Resampling (RIS)

[Talbot et al. '05]

- Start with *candidate* distribution $p(x)$…

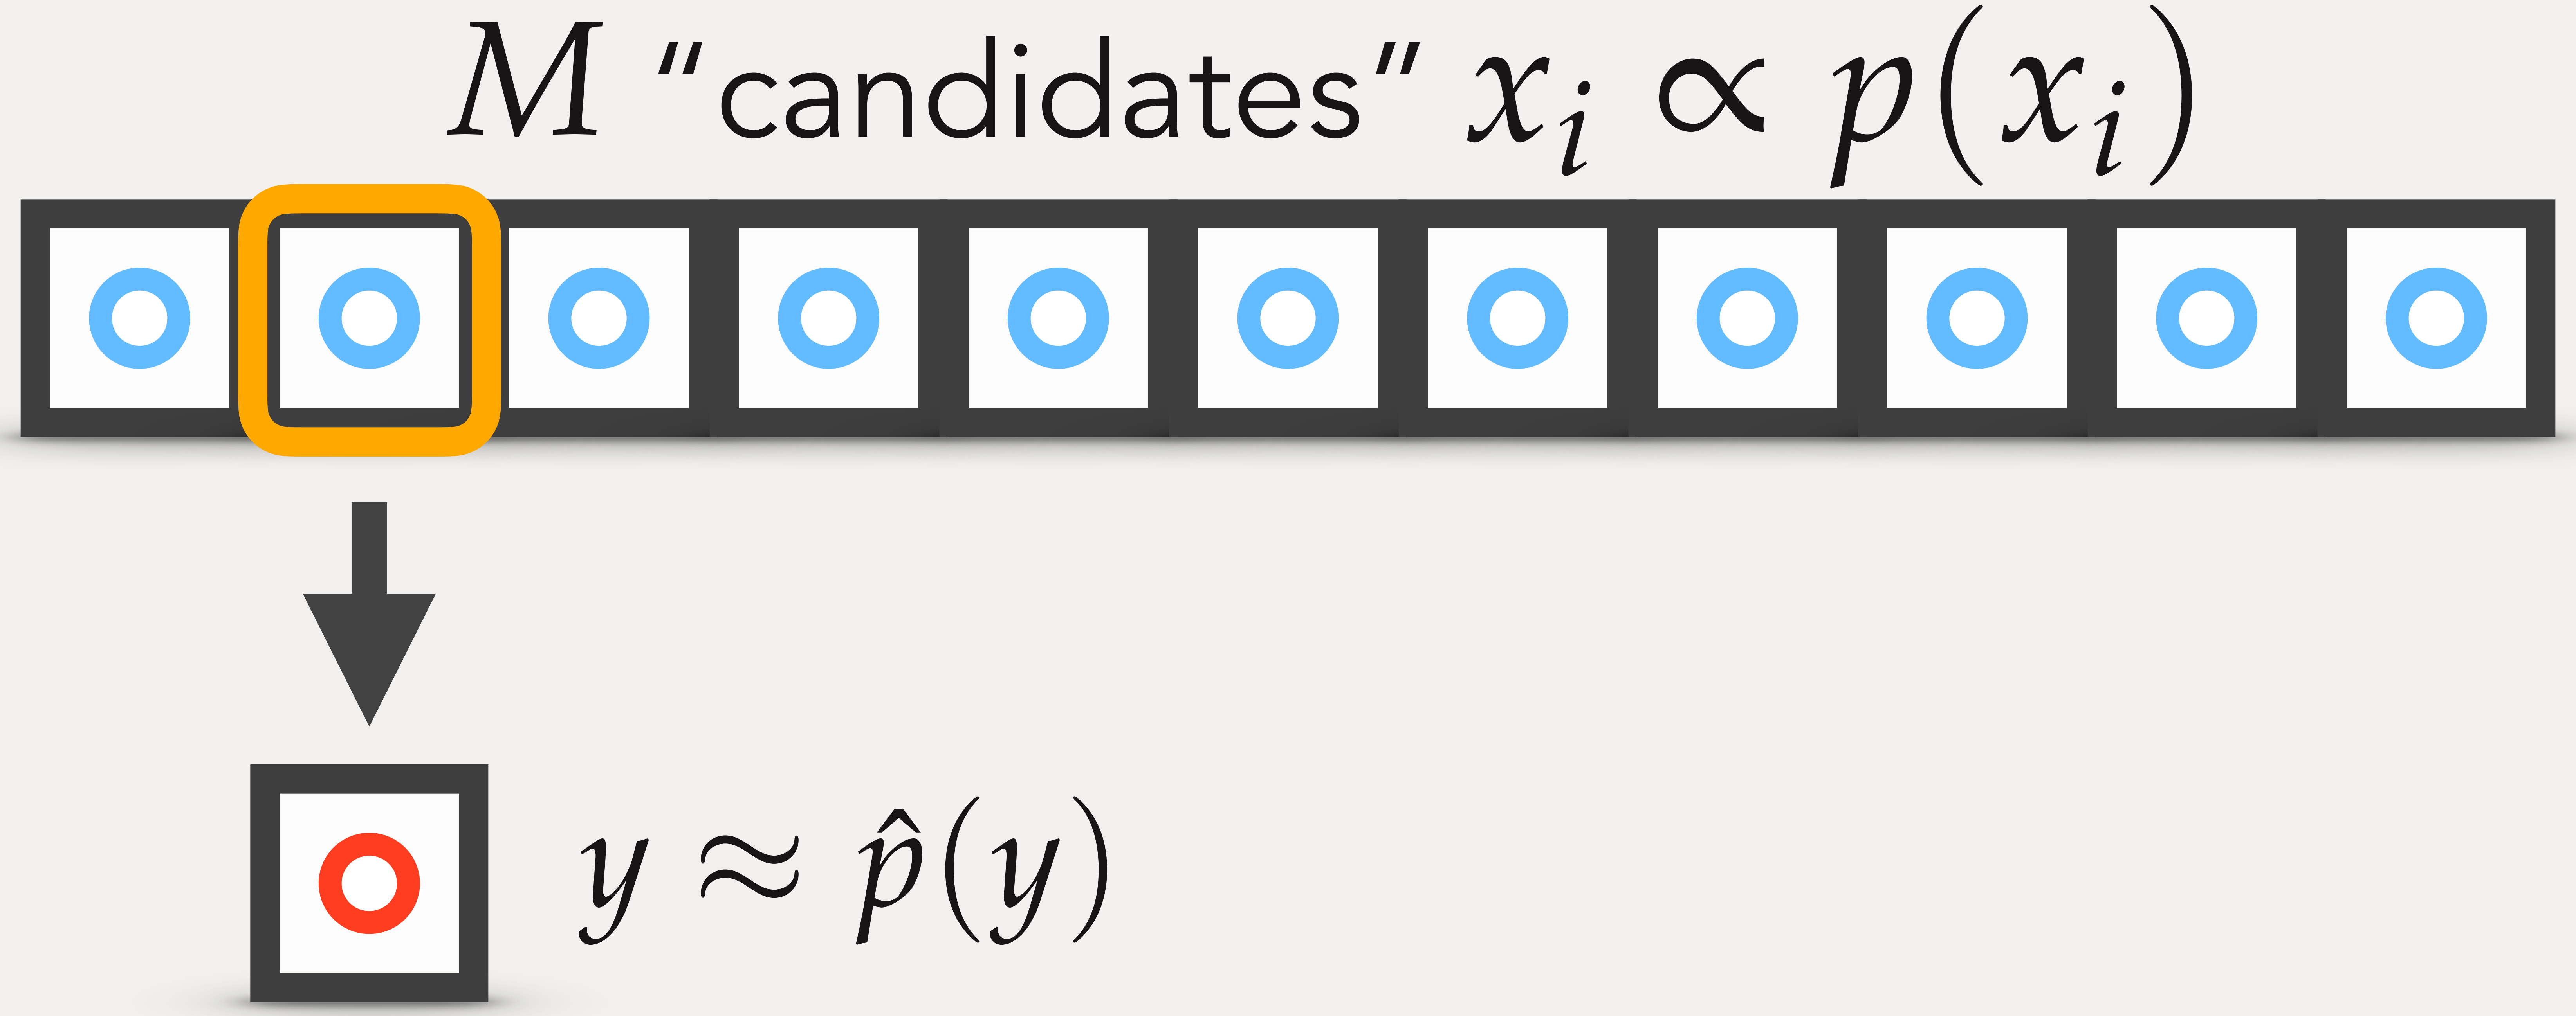

Importance Resampling (RIS)

[Talbot et al. '05]

- Start with *candidate* distribution $p(x)$…

M "candidates" $x_i \propto p(x_i)$

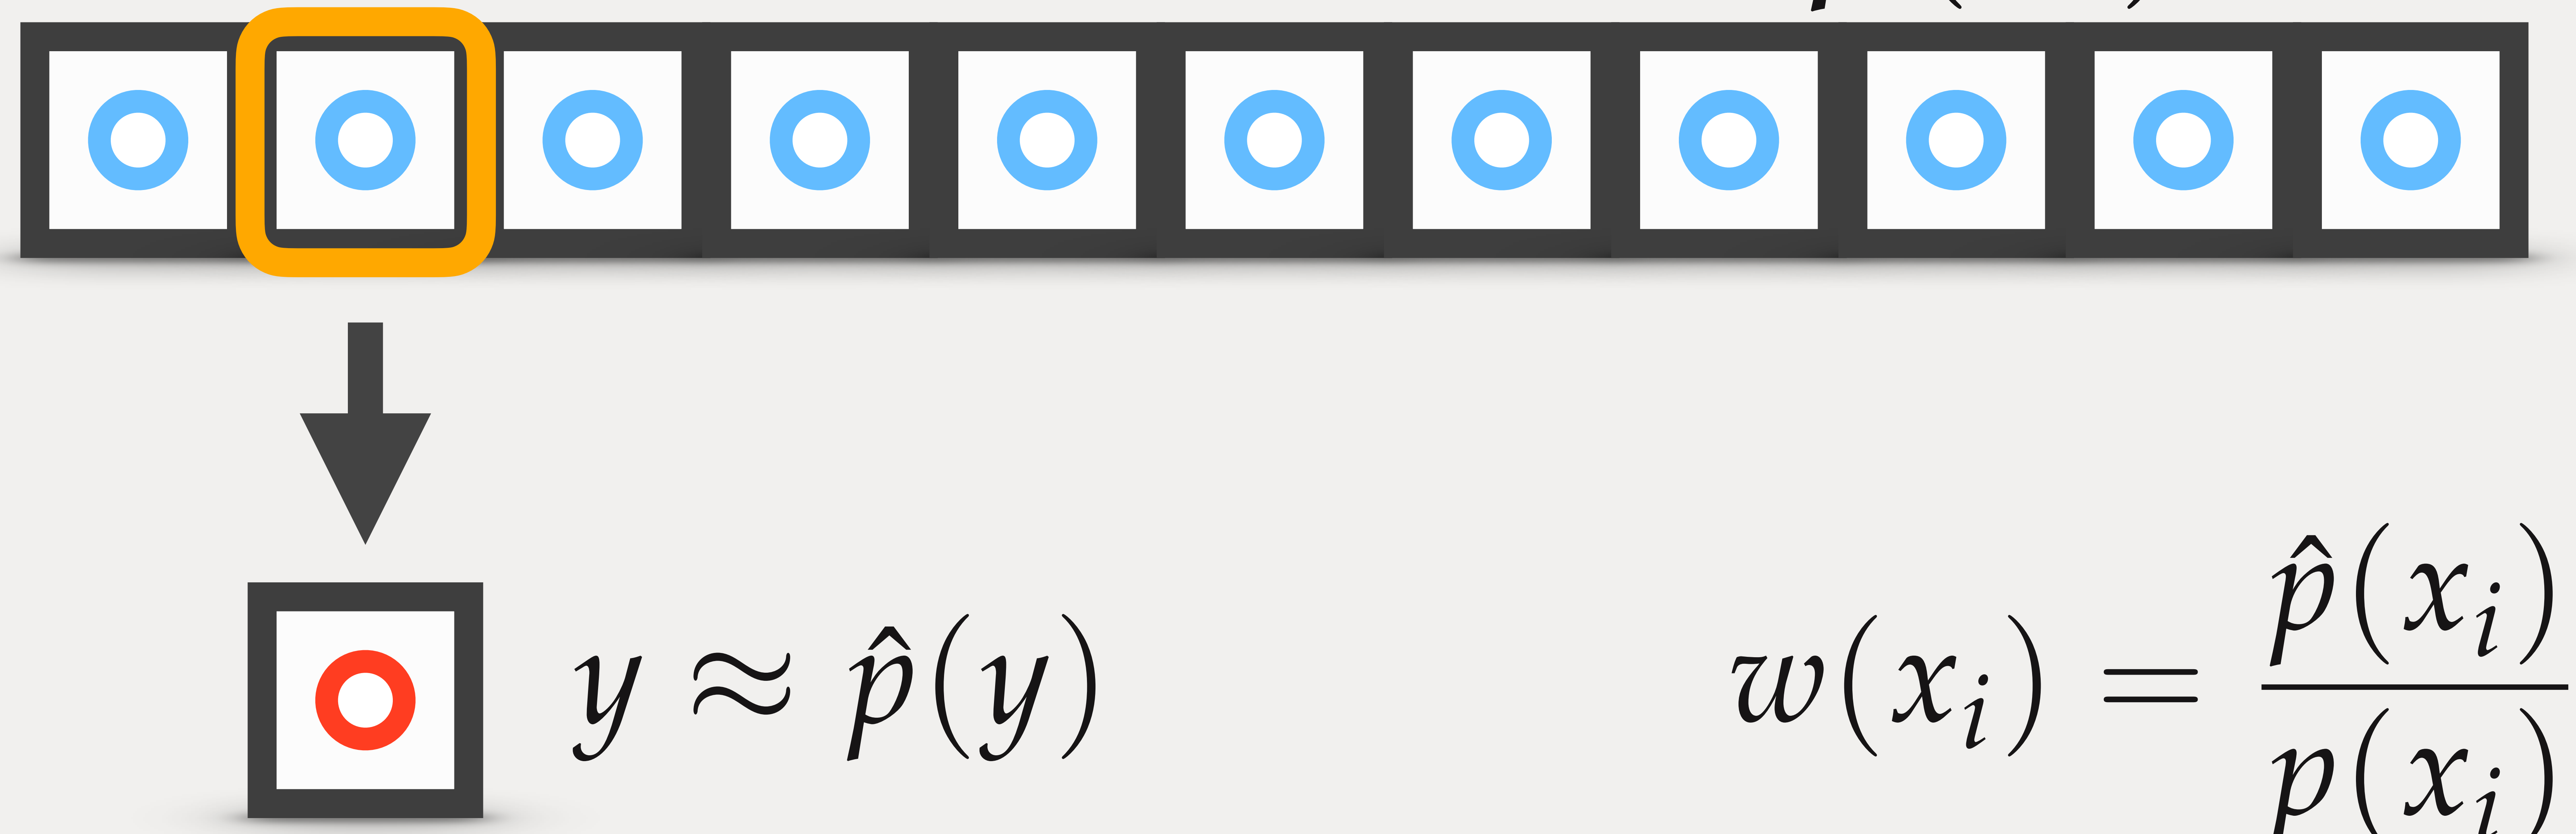

$y \approx \hat{p}(y)$

$$w(x_i) = \frac{\hat{p}(x_i)}{p(x_i)}$$

Importance Resampling (RIS)

[Talbot et al. '05]

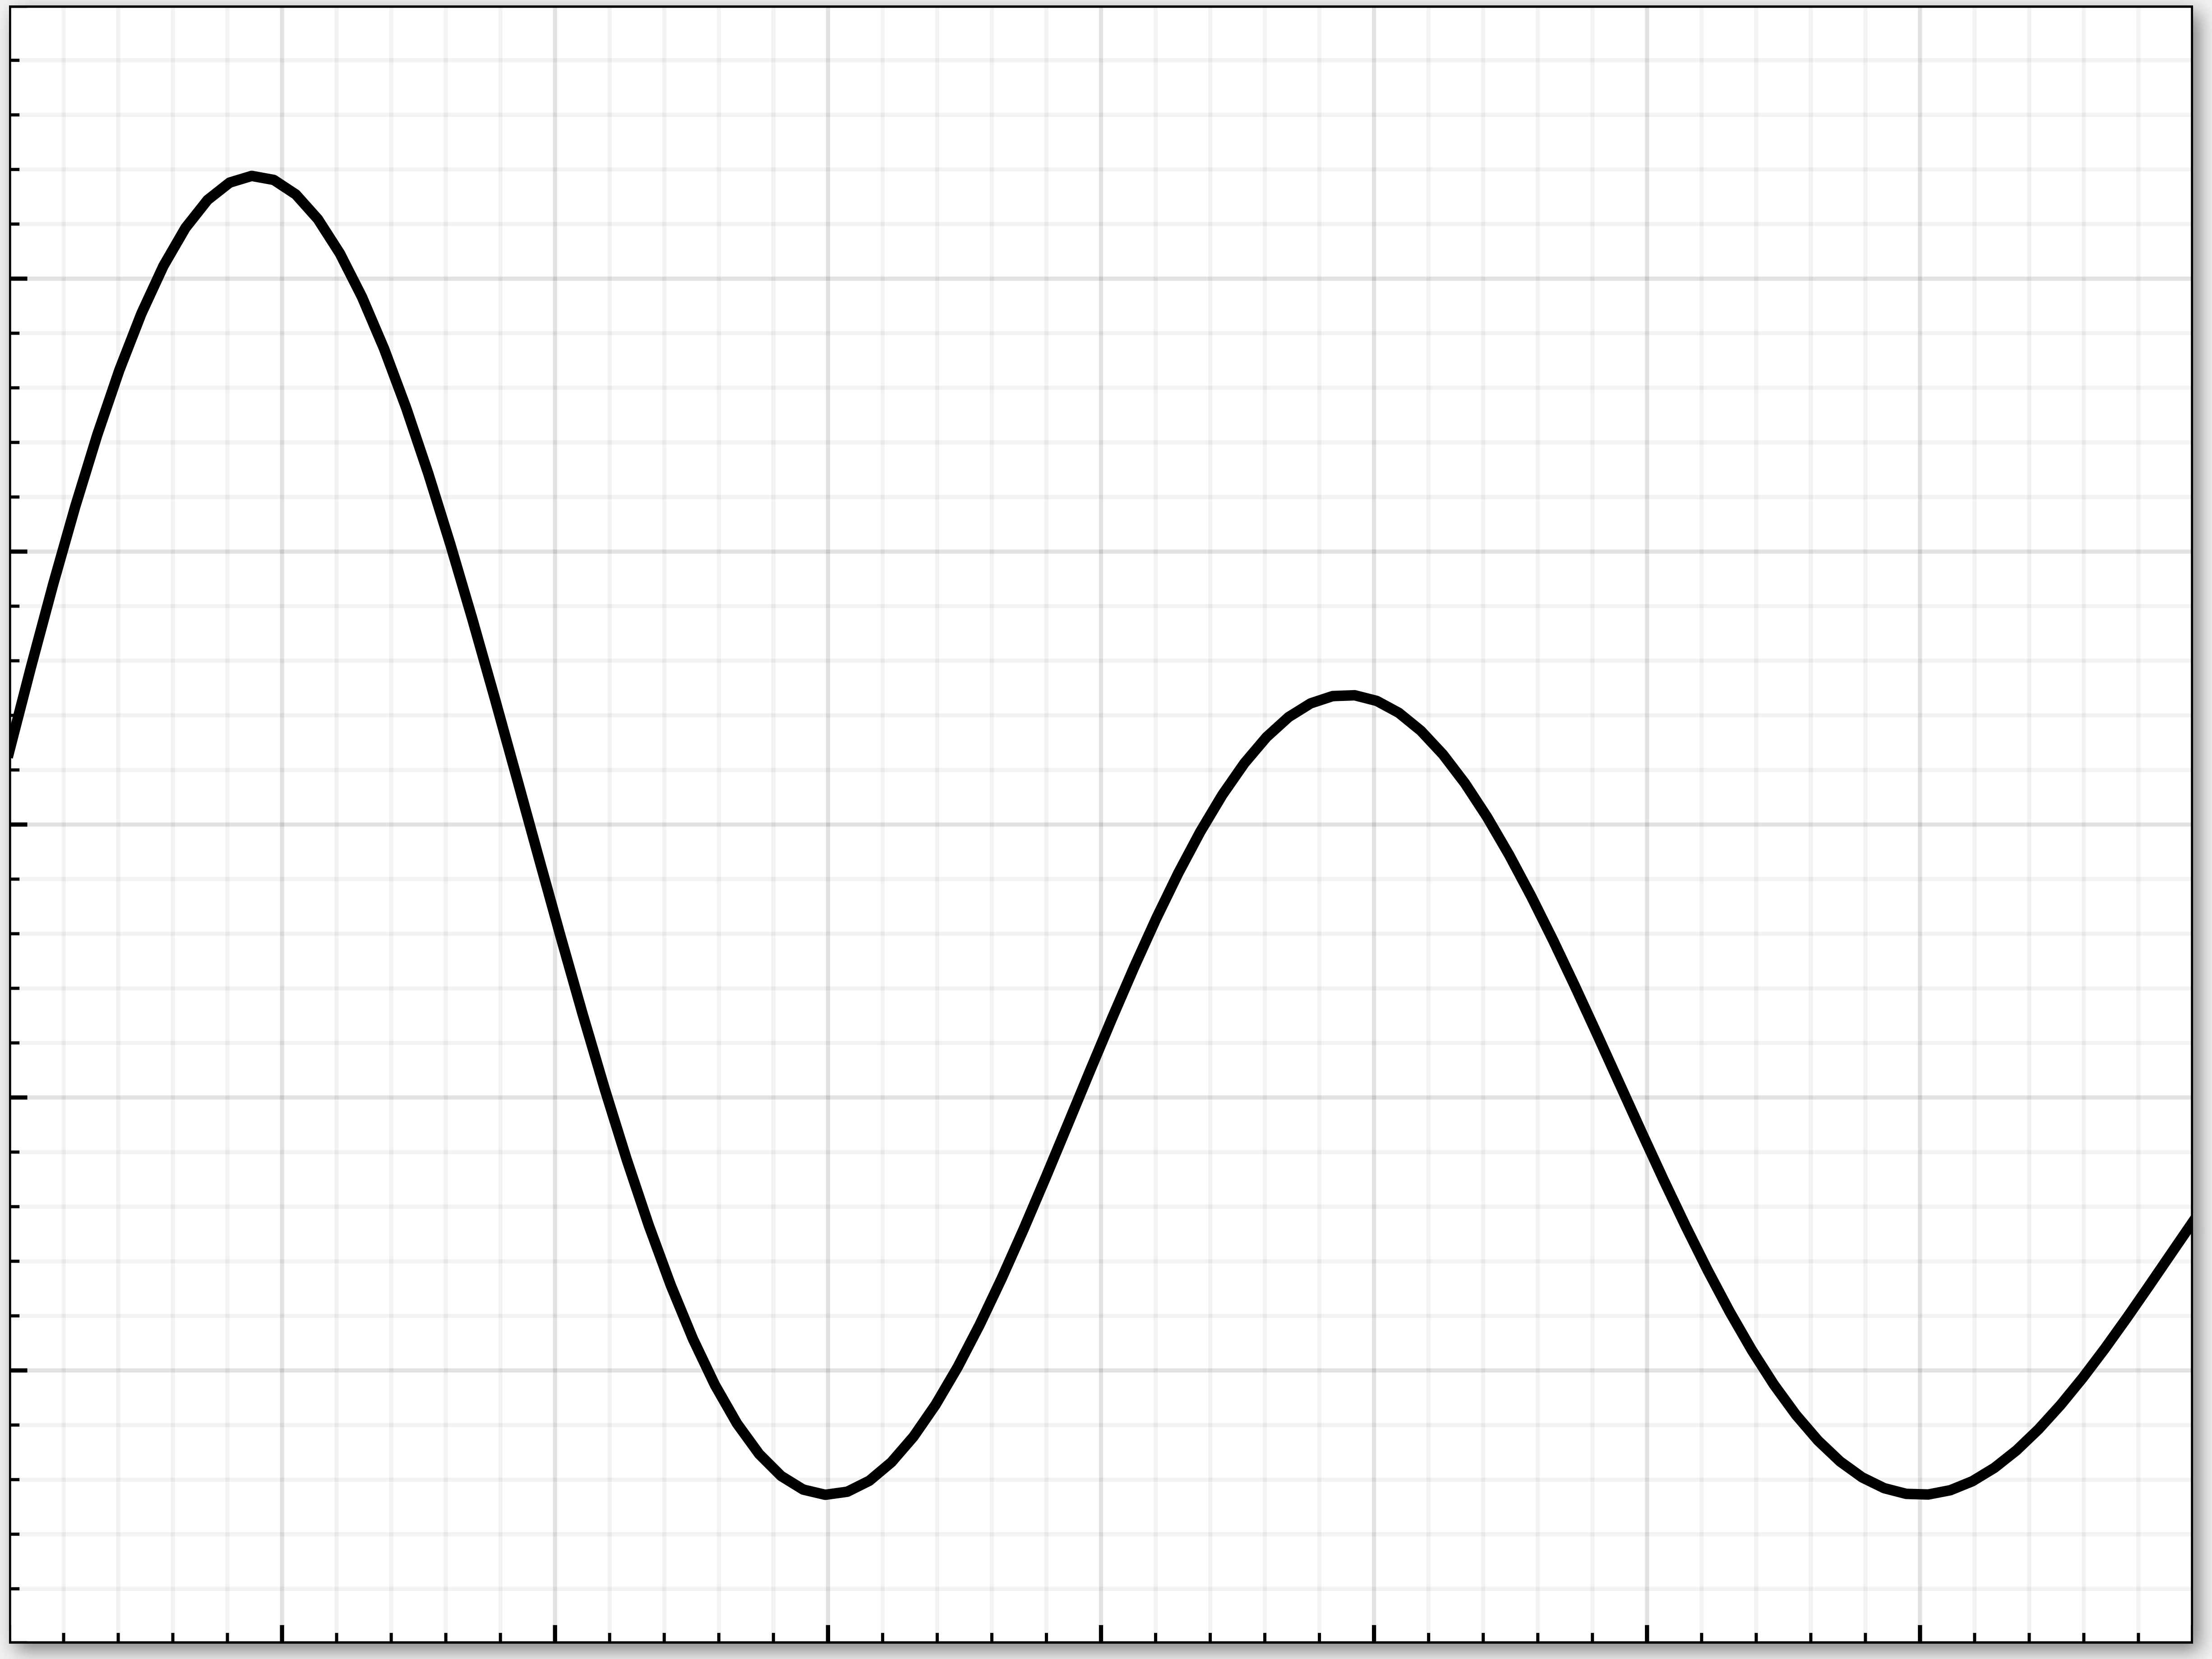

$\hat{p}(x)$

Importance Resampling (RIS)

[Talbot et al. '05]

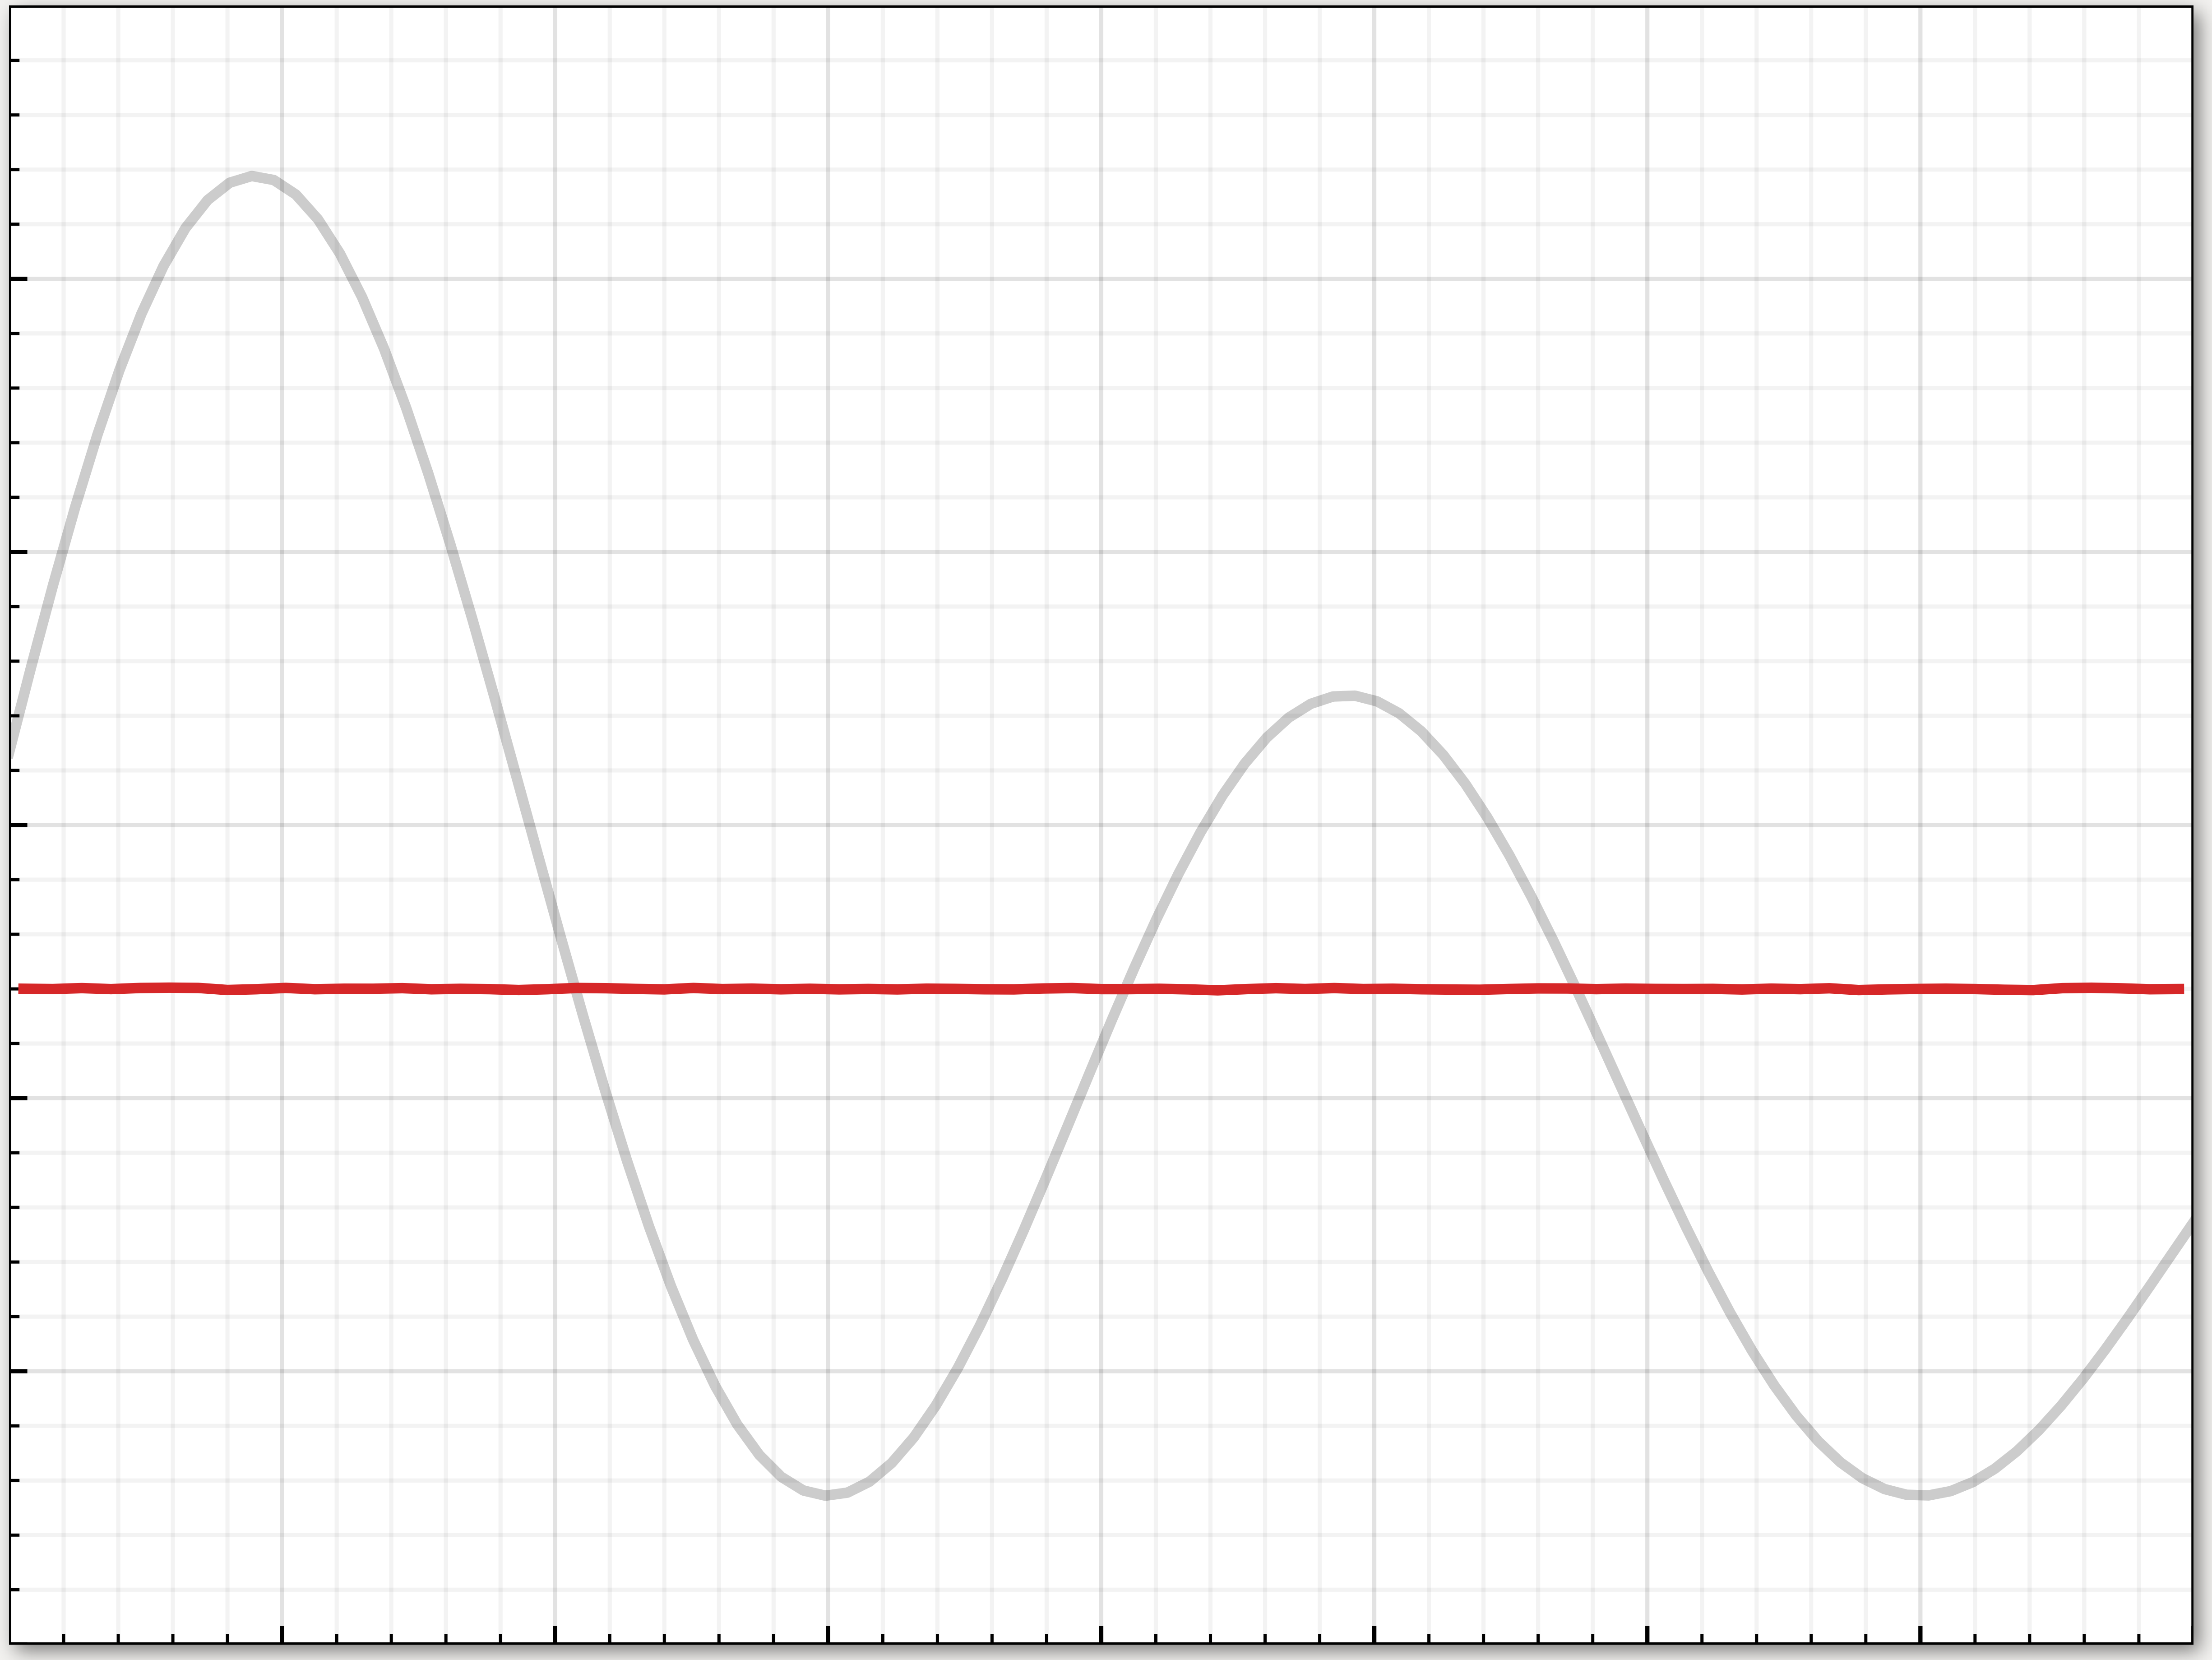

$p(x)$

Importance Resampling (RIS)

[Talbot et al. ’05]

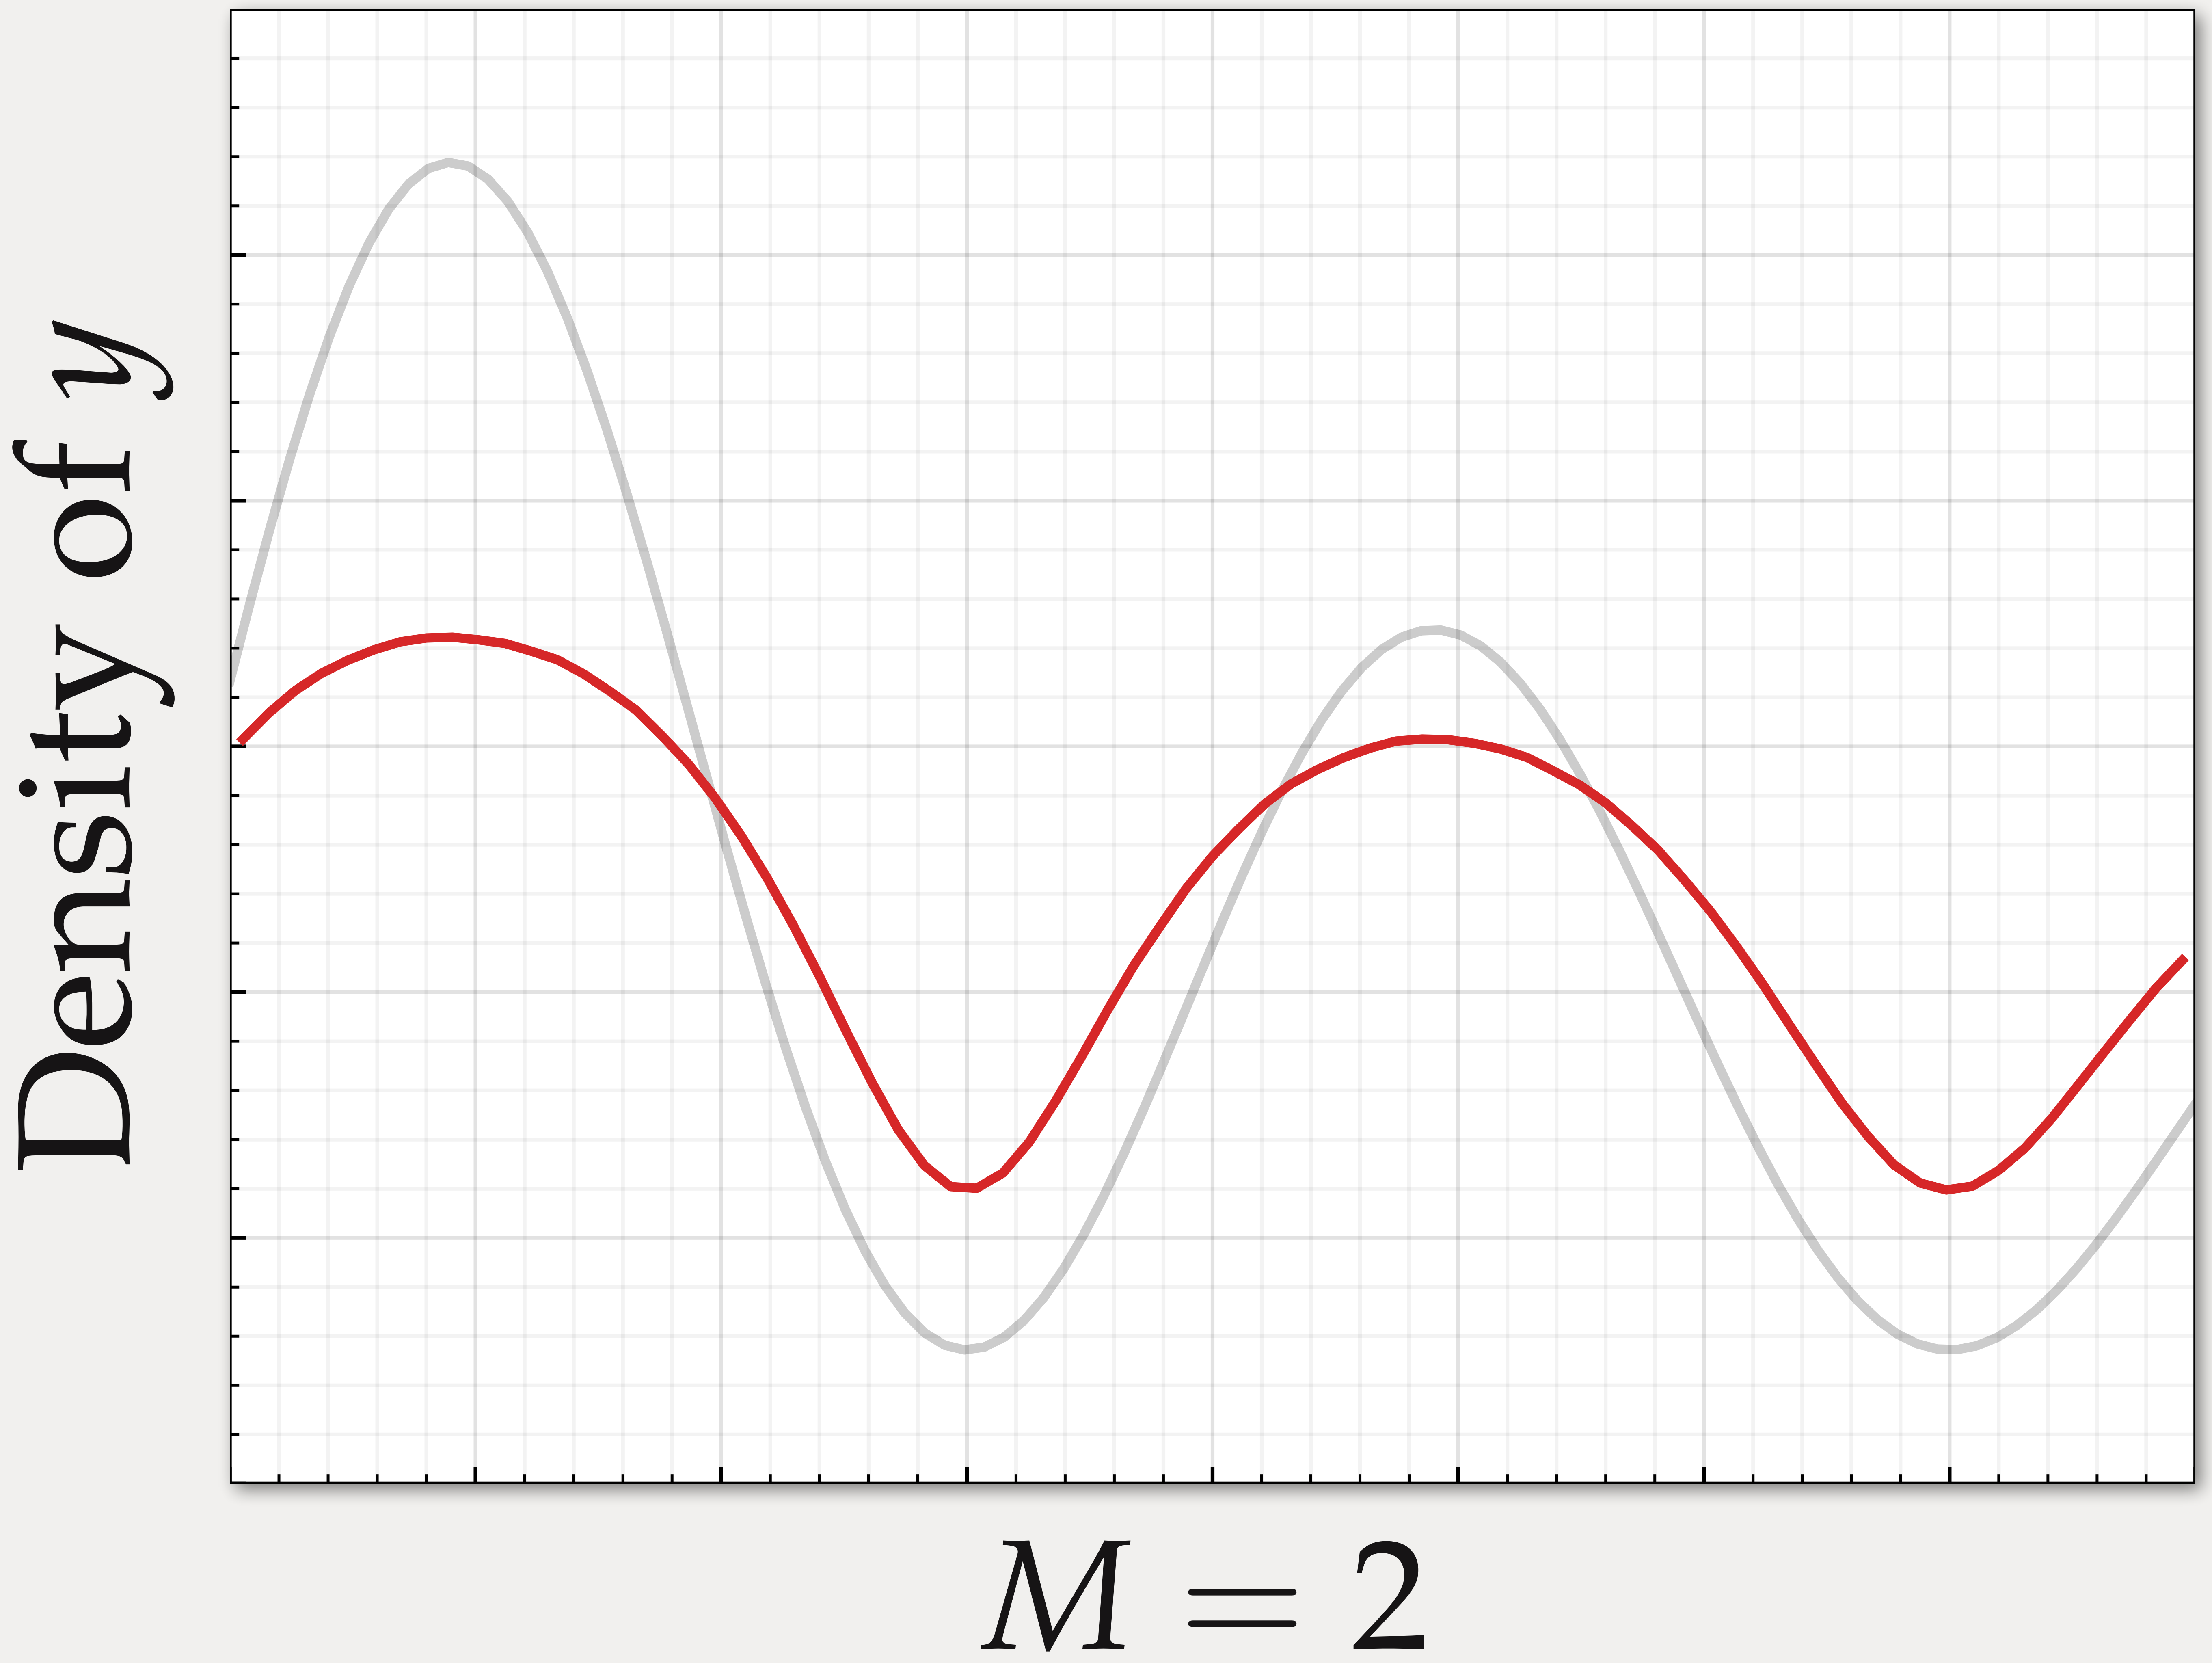

Importance Resampling (RIS)

[Talbot et al. '05]

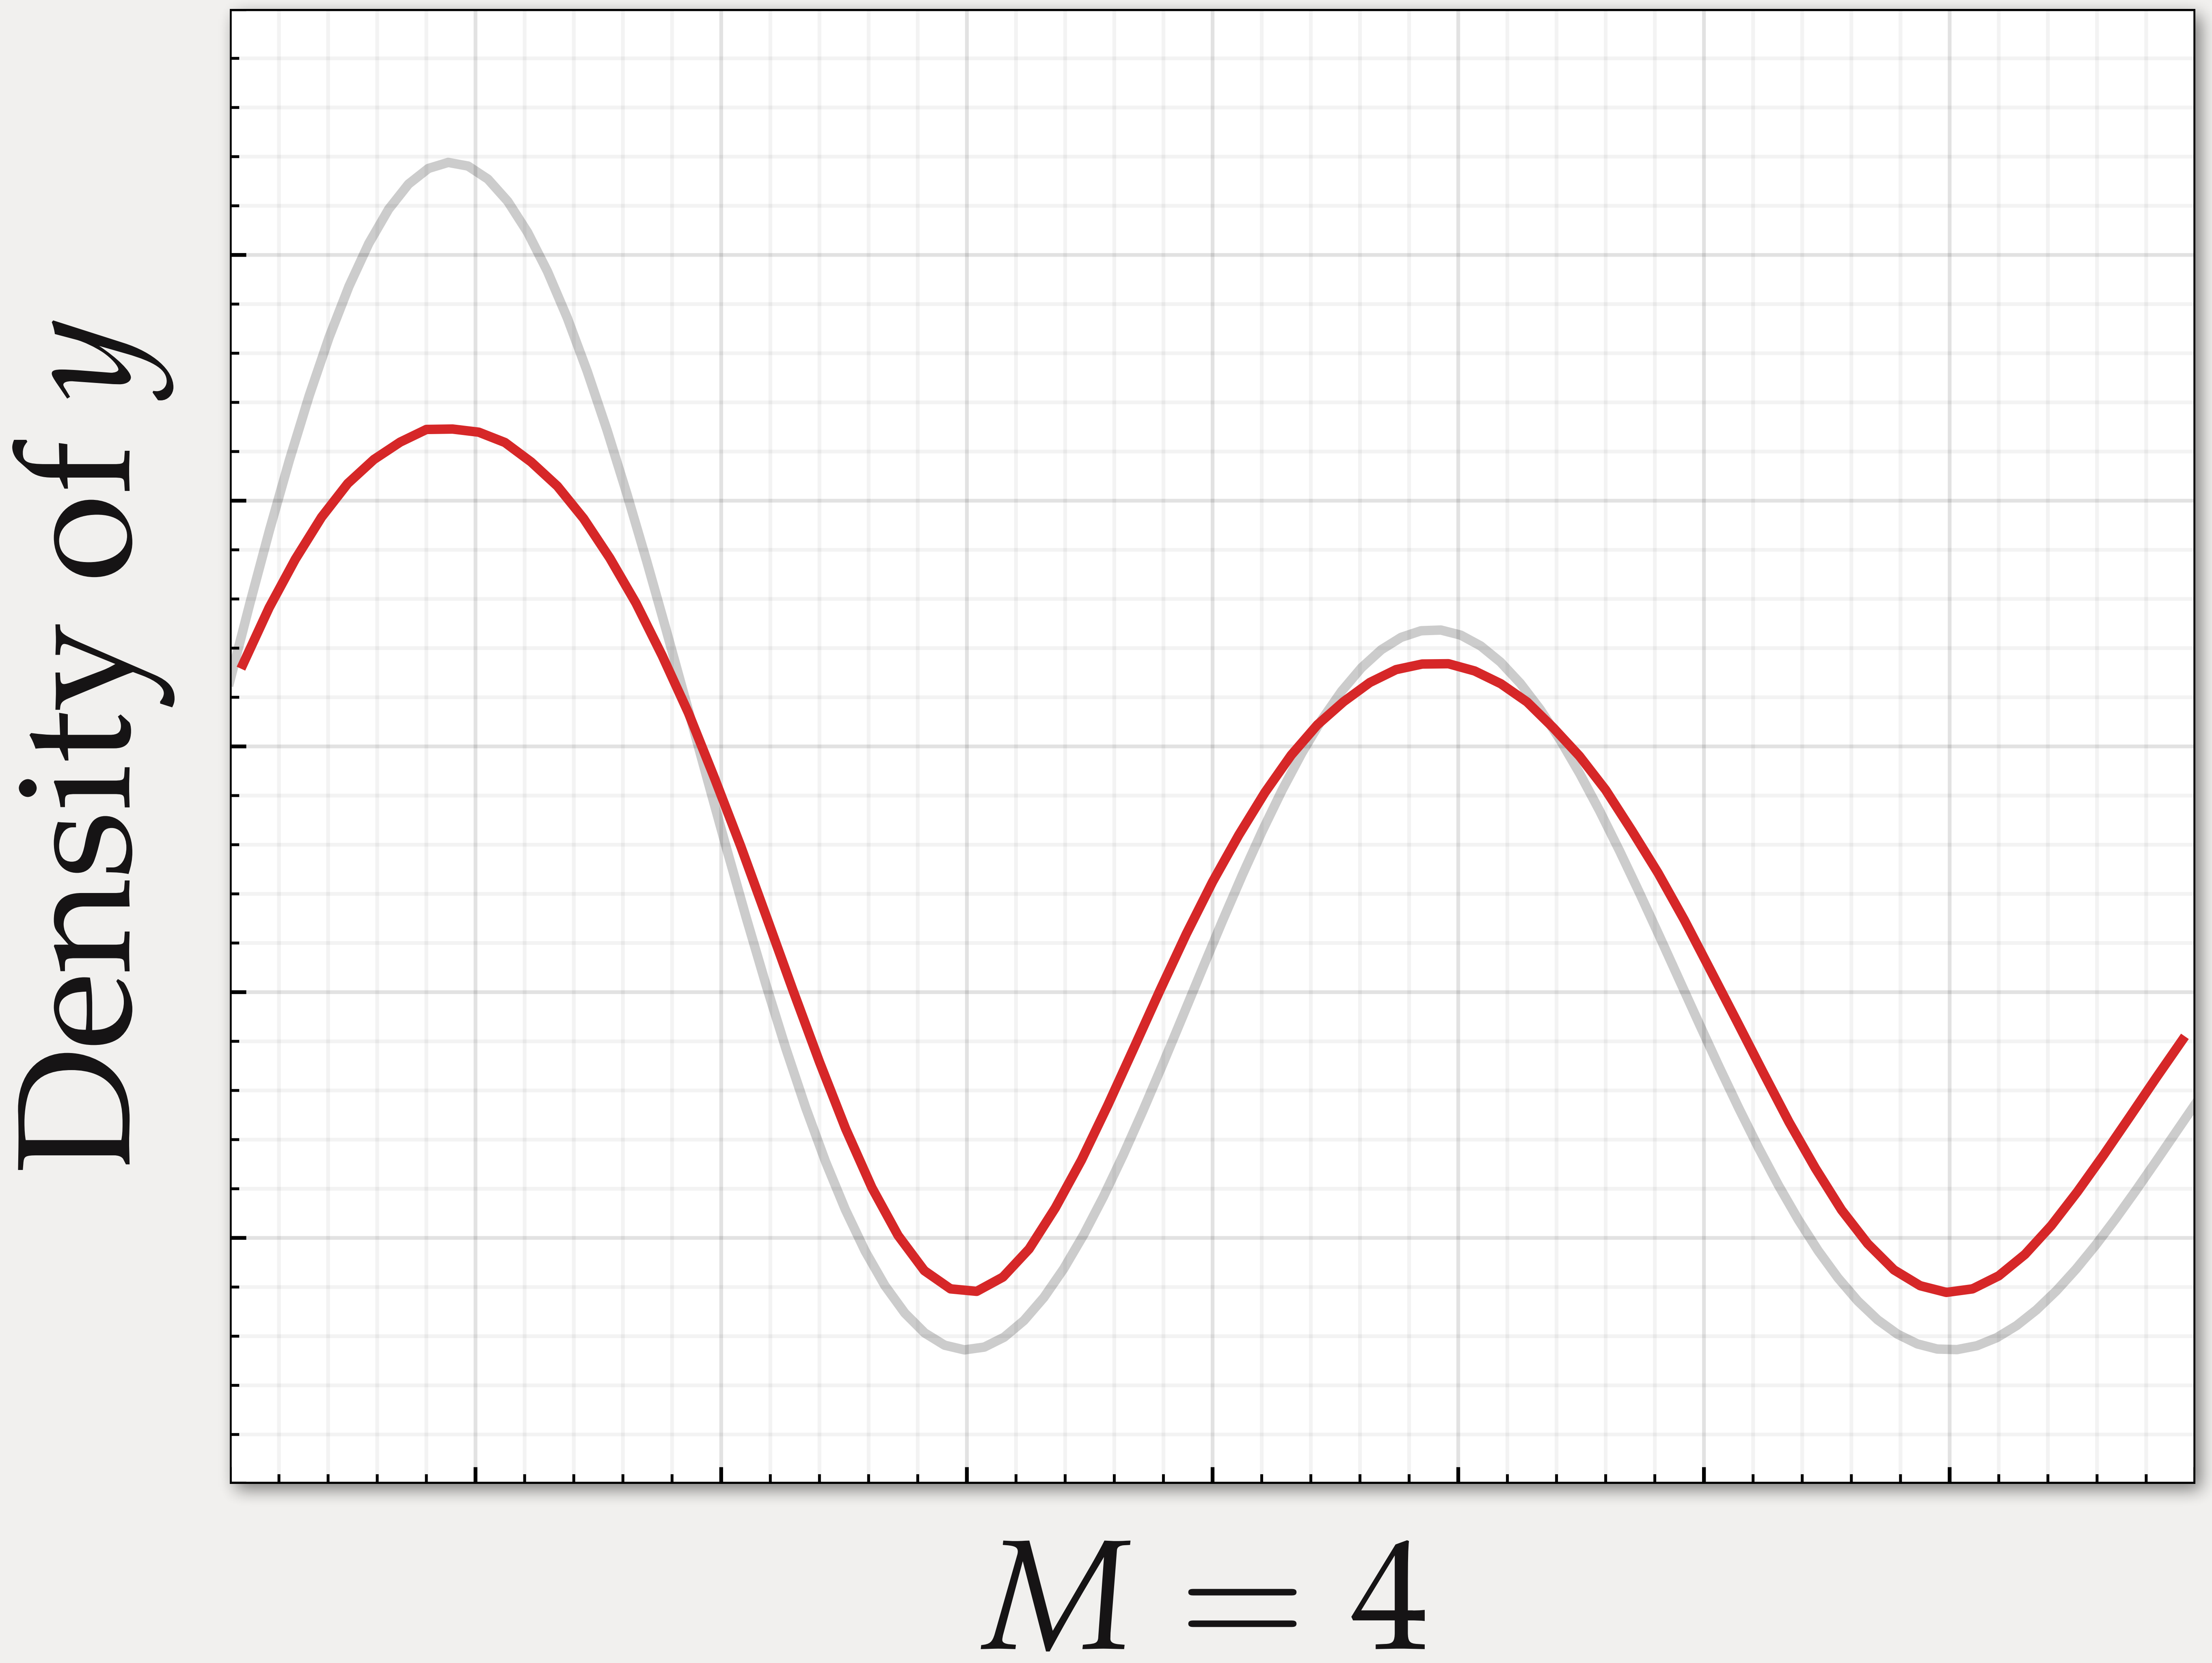

Importance Resampling (RIS)

[Talbot et al. '05]

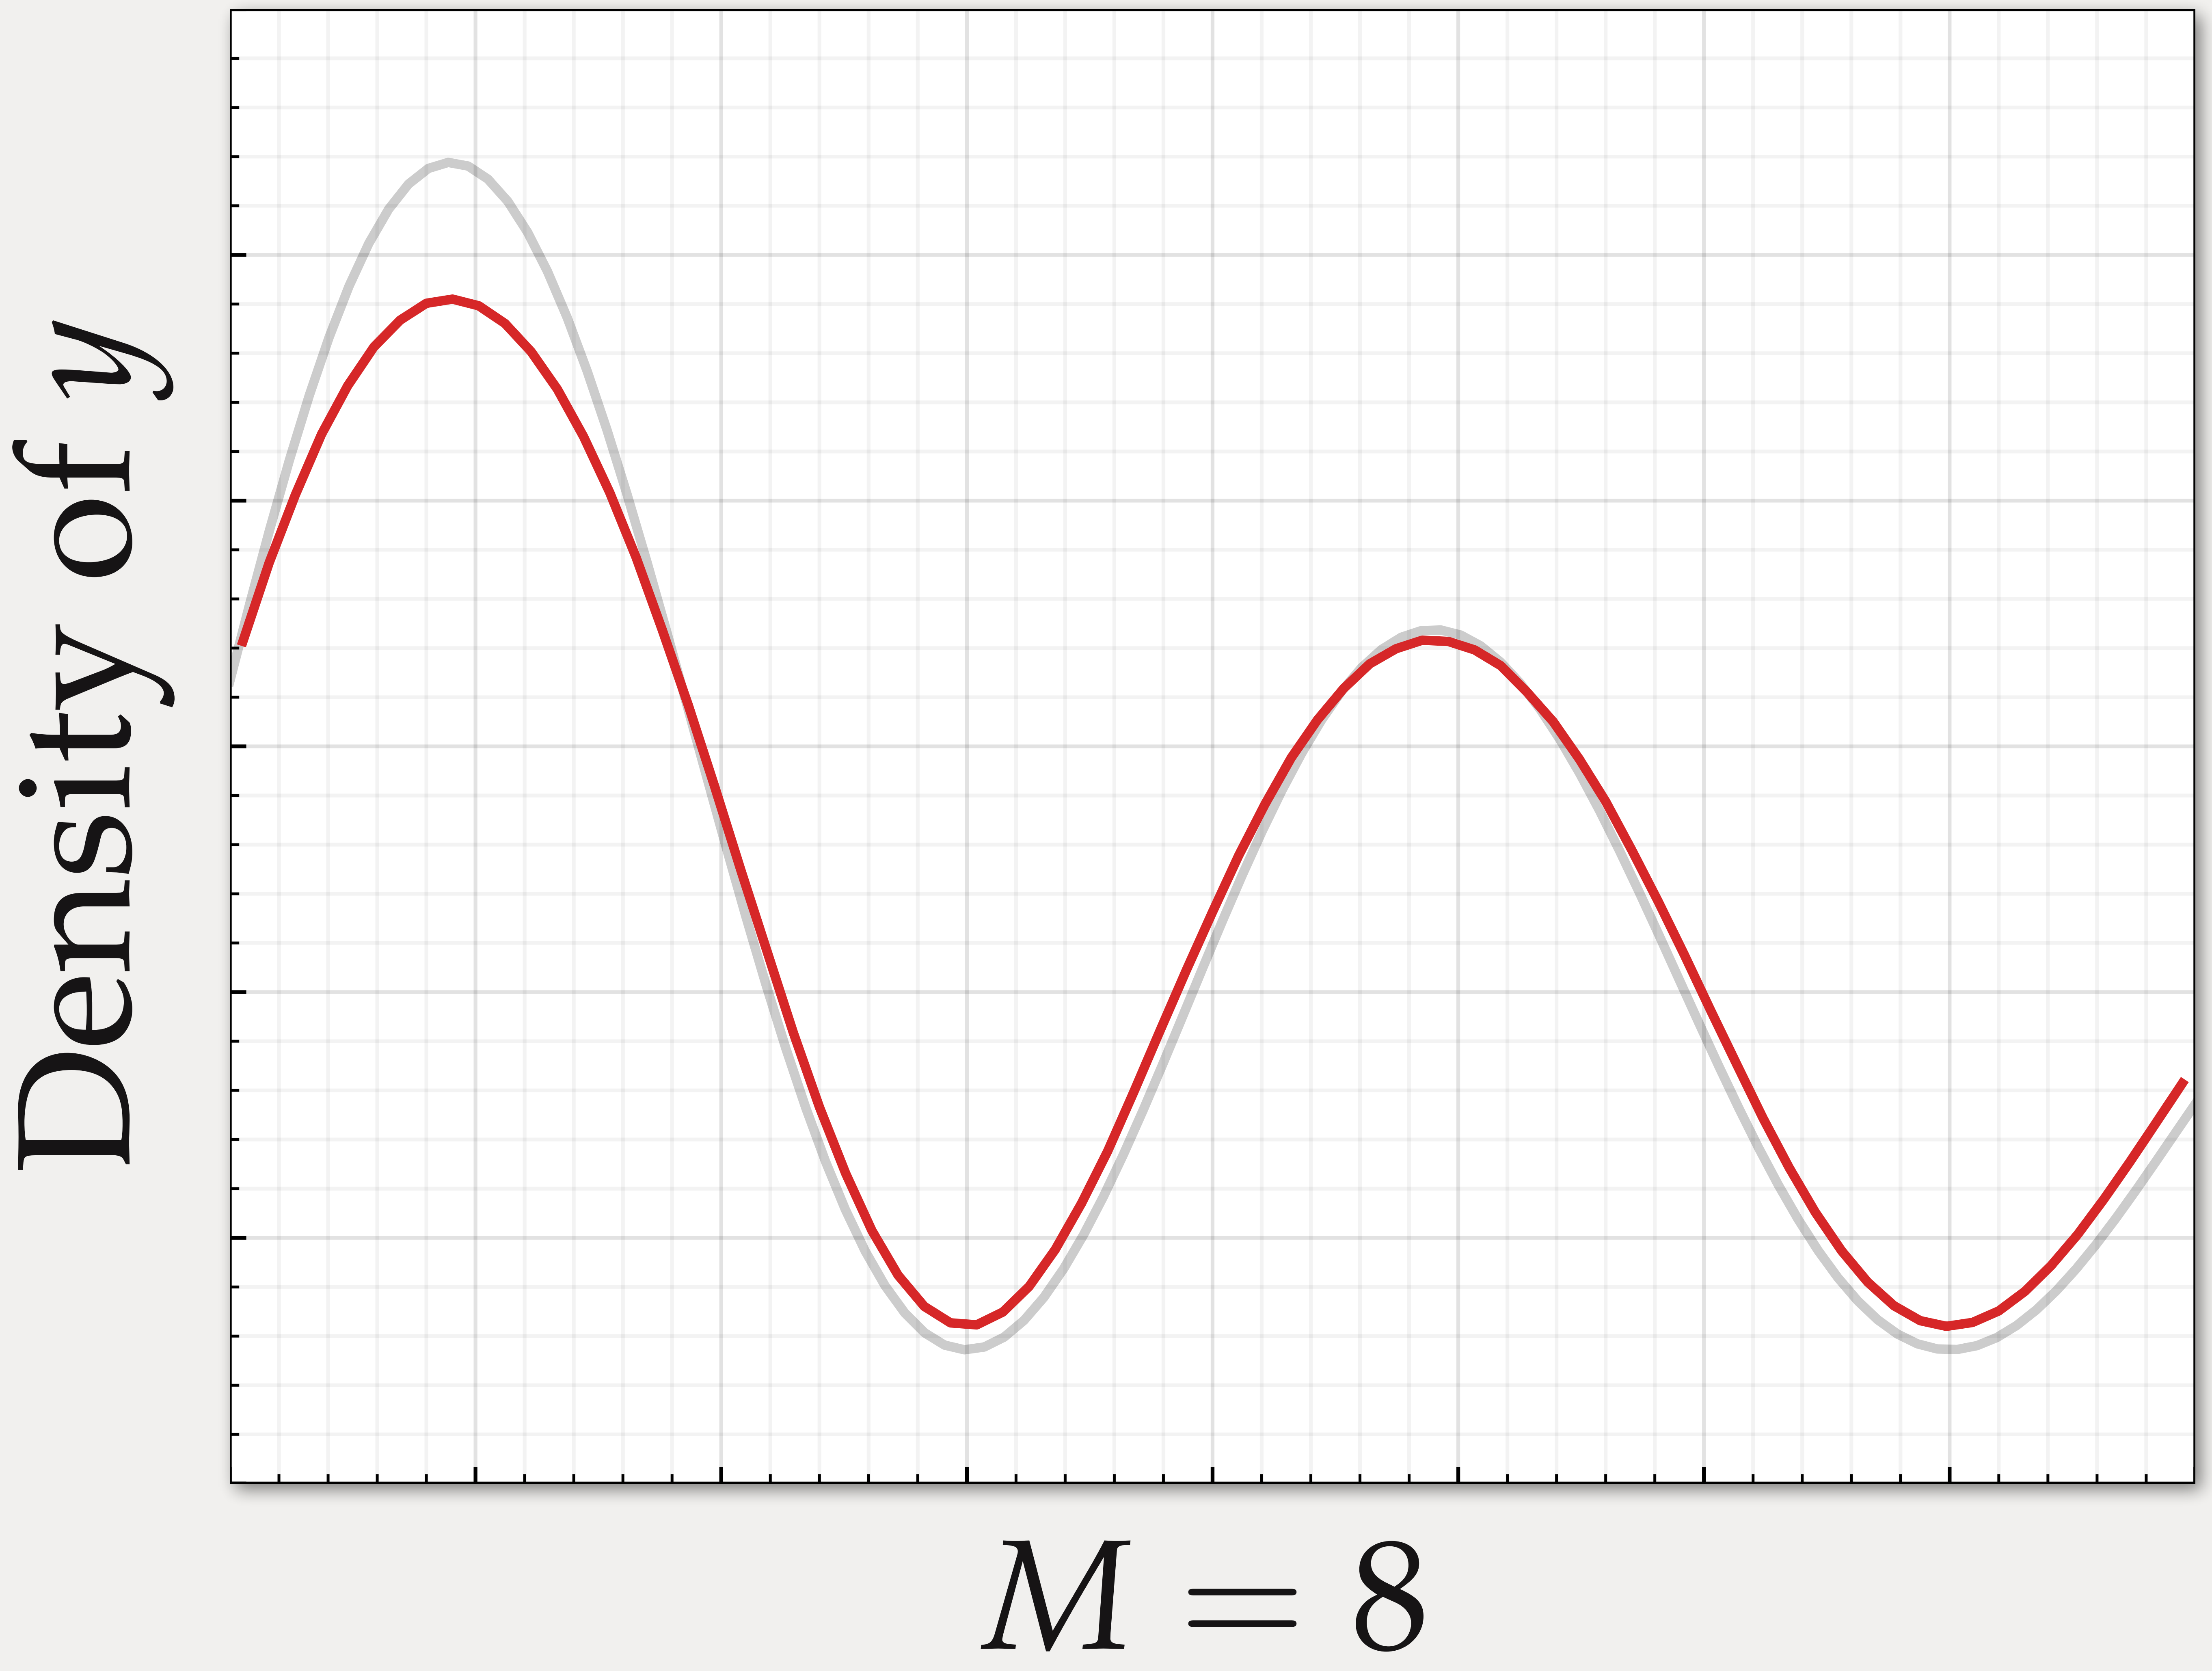

Importance Resampling (RIS)

[Talbot et al. '05]

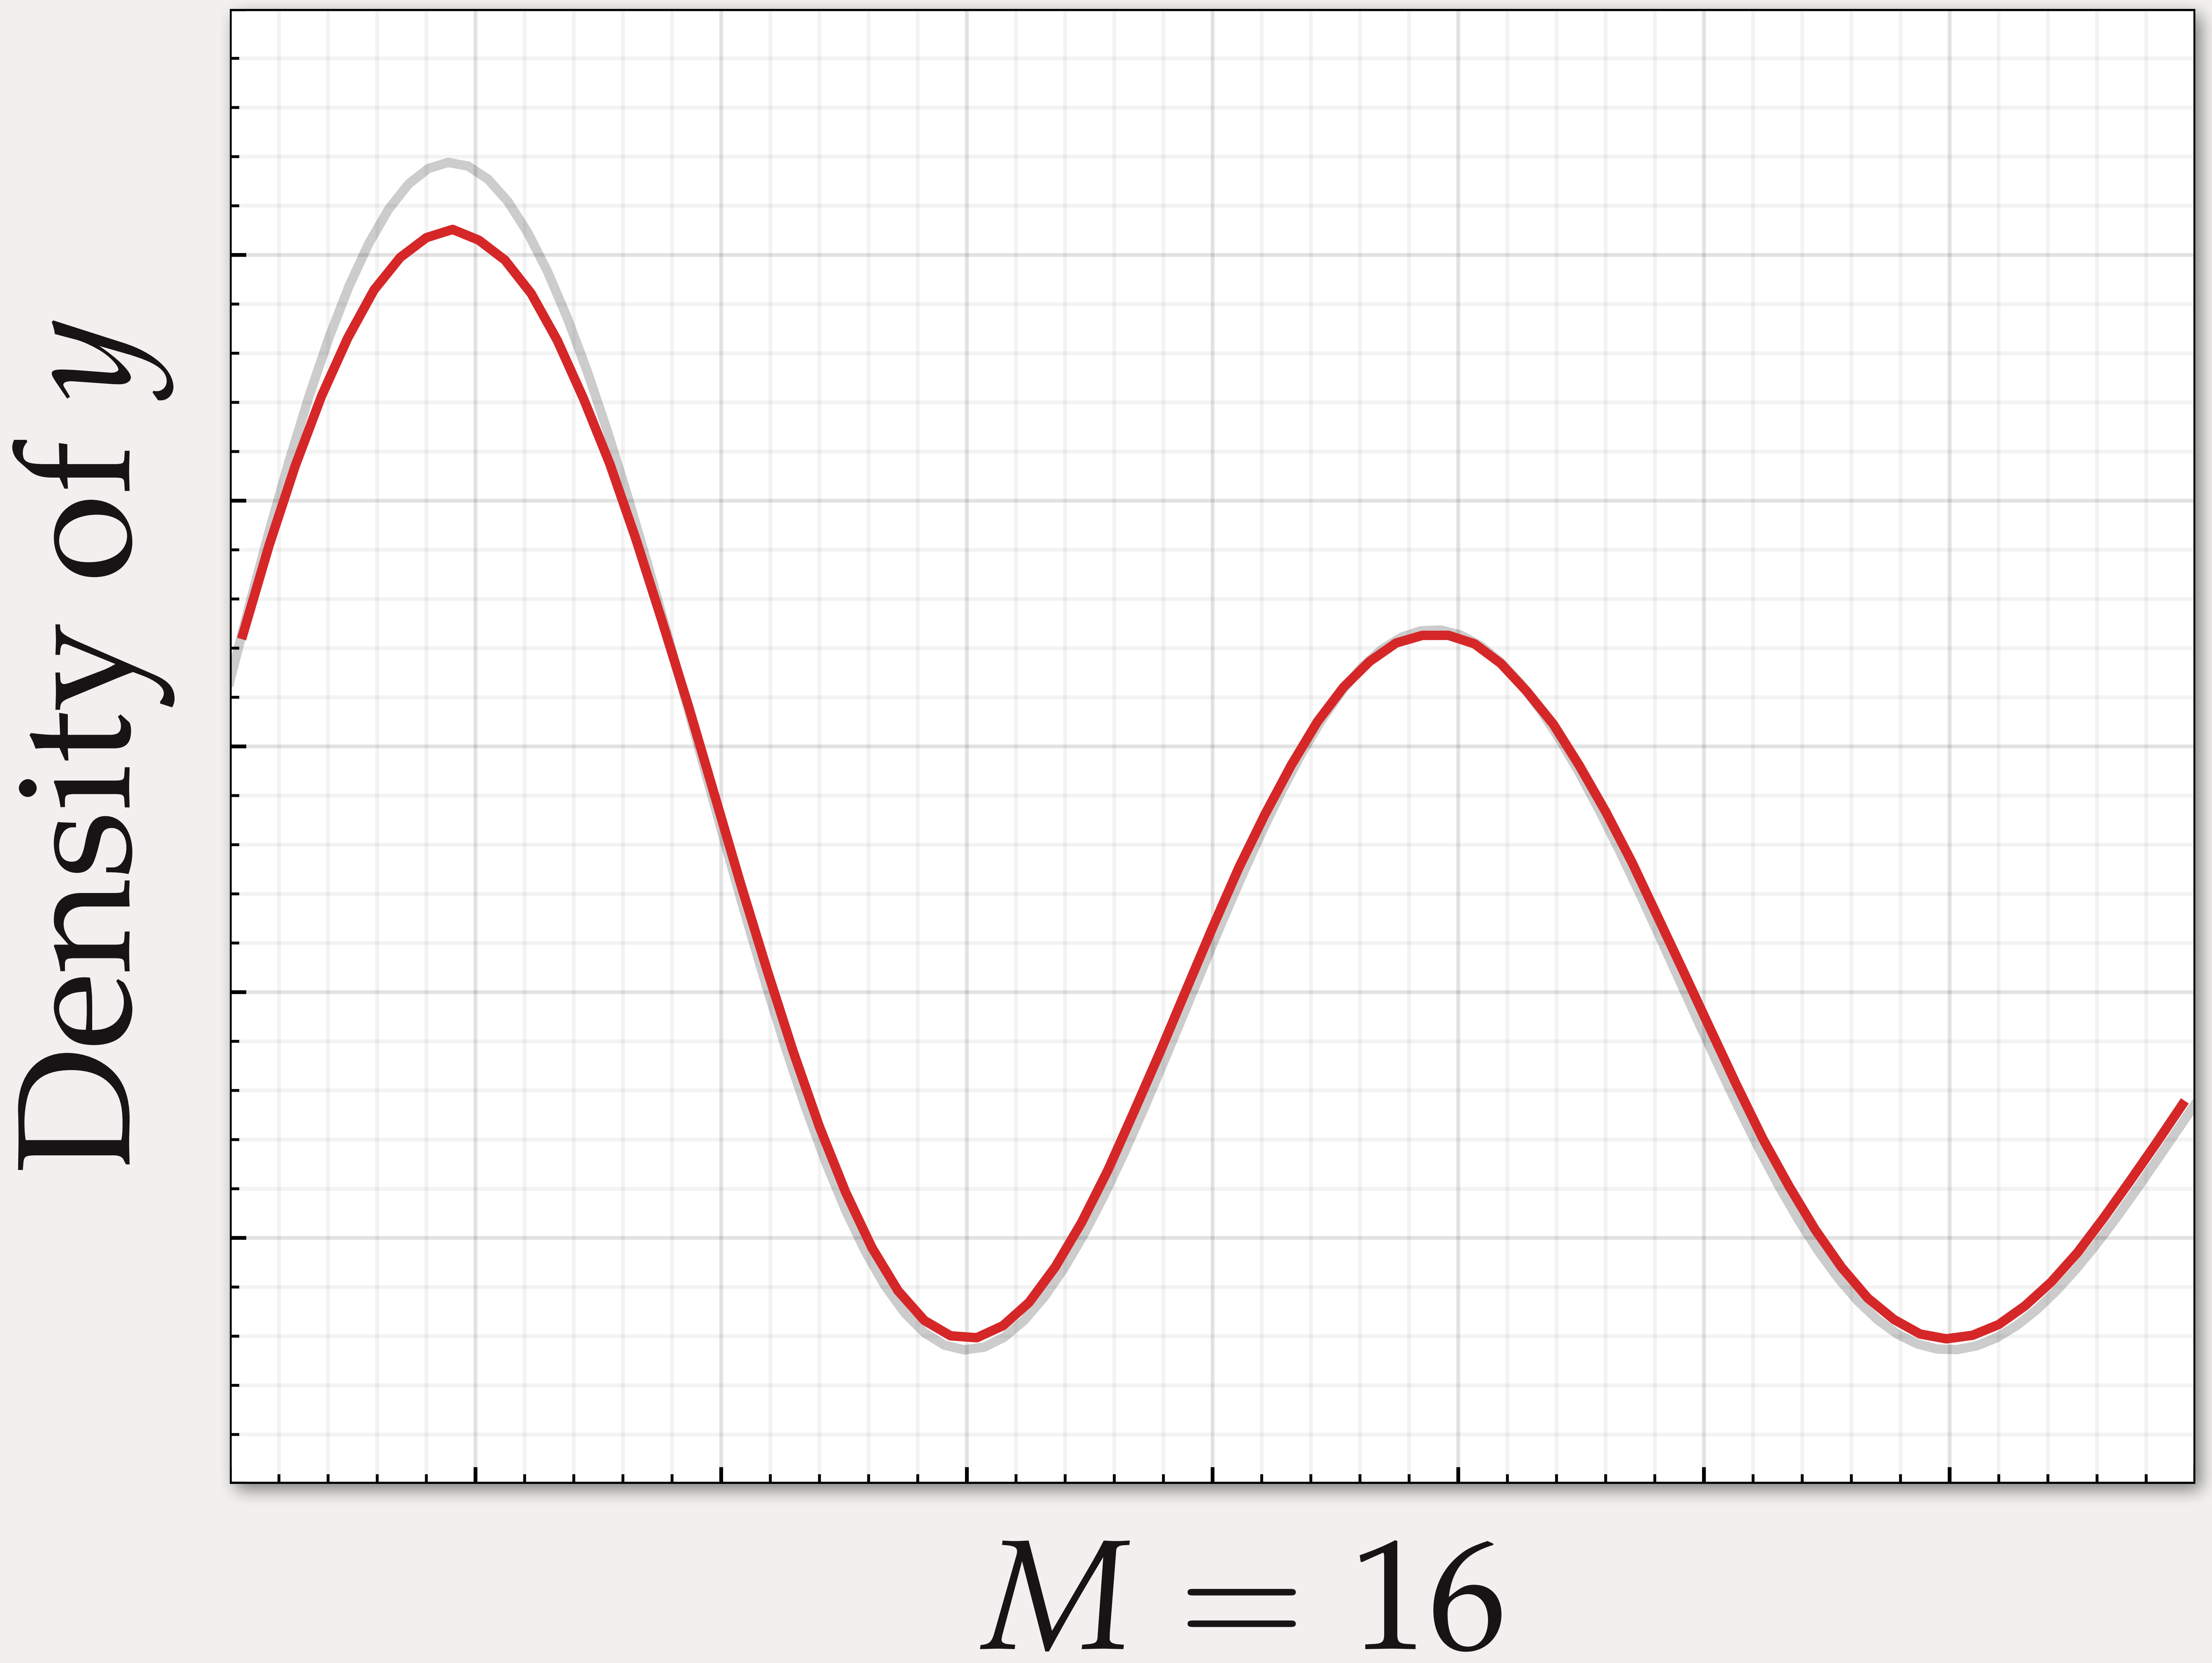

Importance Resampling (RIS)

[Talbot et al. '05]

$$\int_{\Omega} f(x)\mathrm{d}x$$

Importance Resampling (RIS)

[Talbot et al. '05]

$$\int_{\Omega} f(x)\mathrm{d}x \approx f(y) \cdot W$$

Importance Resampling (RIS)

[Talbot et al. '05]

$$\int_{\Omega} f(x)\mathrm{d}x \approx f(y) \cdot W$$

$$W = \frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M} \sum_{i=1}^{M} w(x_i) \right)$$

Importance Resampling (RIS)

[Talbot et al. '05]

$$\int_\Omega f(x)\mathrm{d}x \approx f(y) \cdot W$$

$$W = \frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M} \sum_{i=1}^{M} w(x_i) \right)$$

Unbiased if $p(x) > 0$ where $\hat{p}(x) > 0$

Resampled Direct Lighting

$$p(x) \propto L_e(x)$$

Resampled Direct Lighting

$$p(x) \propto L_e(x)$$

$$\hat{p}(x) = L_e(x)\rho(x)\frac{\cos\theta_i \cos\theta_o}{r^2}$$

Importance Resampling (RIS)

[Talbot et al. ’05]

Importance Resampling (RIS)
[Talbot et al. '05]

- Sampling candidates requires *storing* them

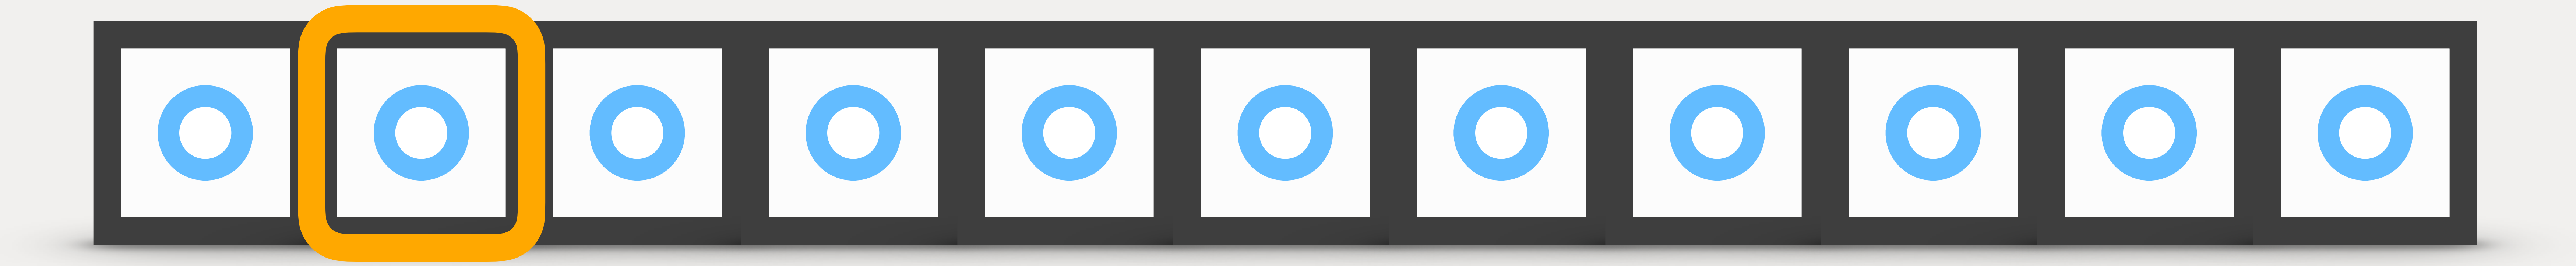

Importance Resampling (RIS)

[Talbot et al. '05]

- Sampling candidates requires *storing* them

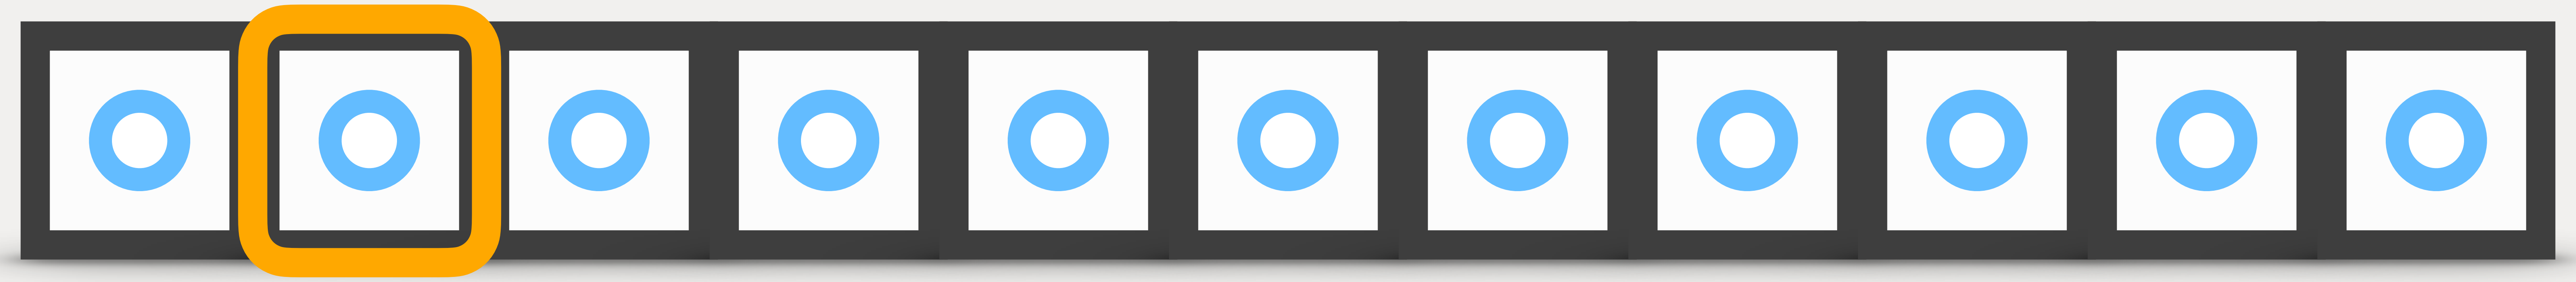

- Not GPU friendly

Importance Resampling (RIS)

[Talbot et al. '05]

- Sampling candidates requires *storing* them

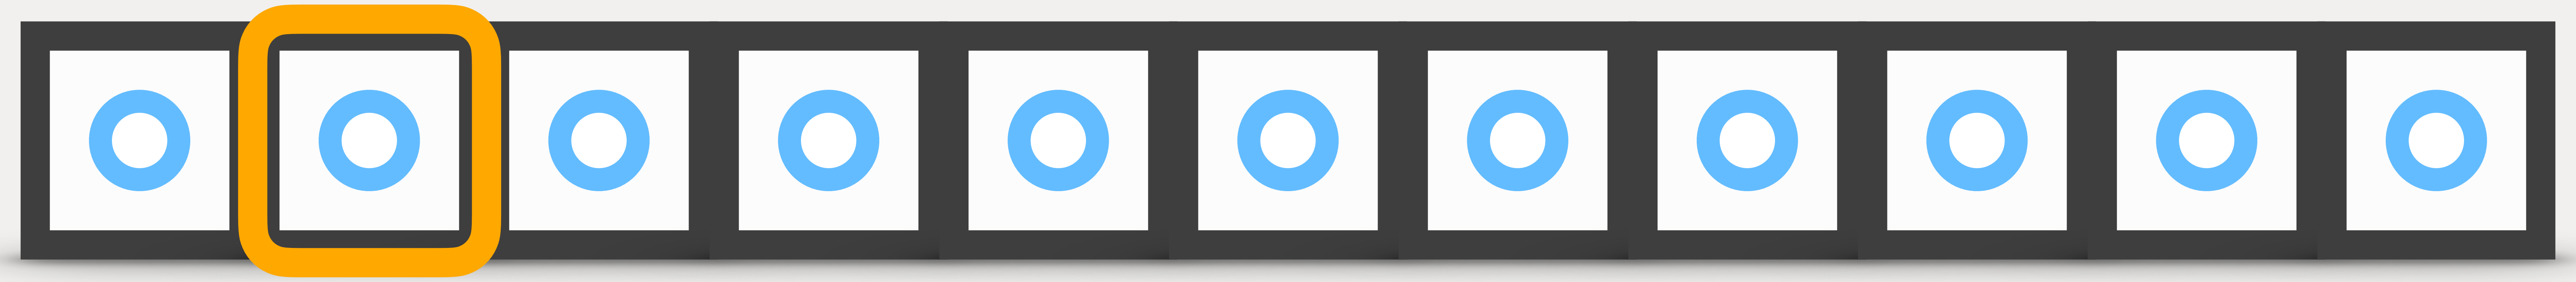

- Not GPU friendly

- Can we do better?

Weighted Reservoir Sampling

[M. T. Chao 1982]

- Stream-based sampling algorithm

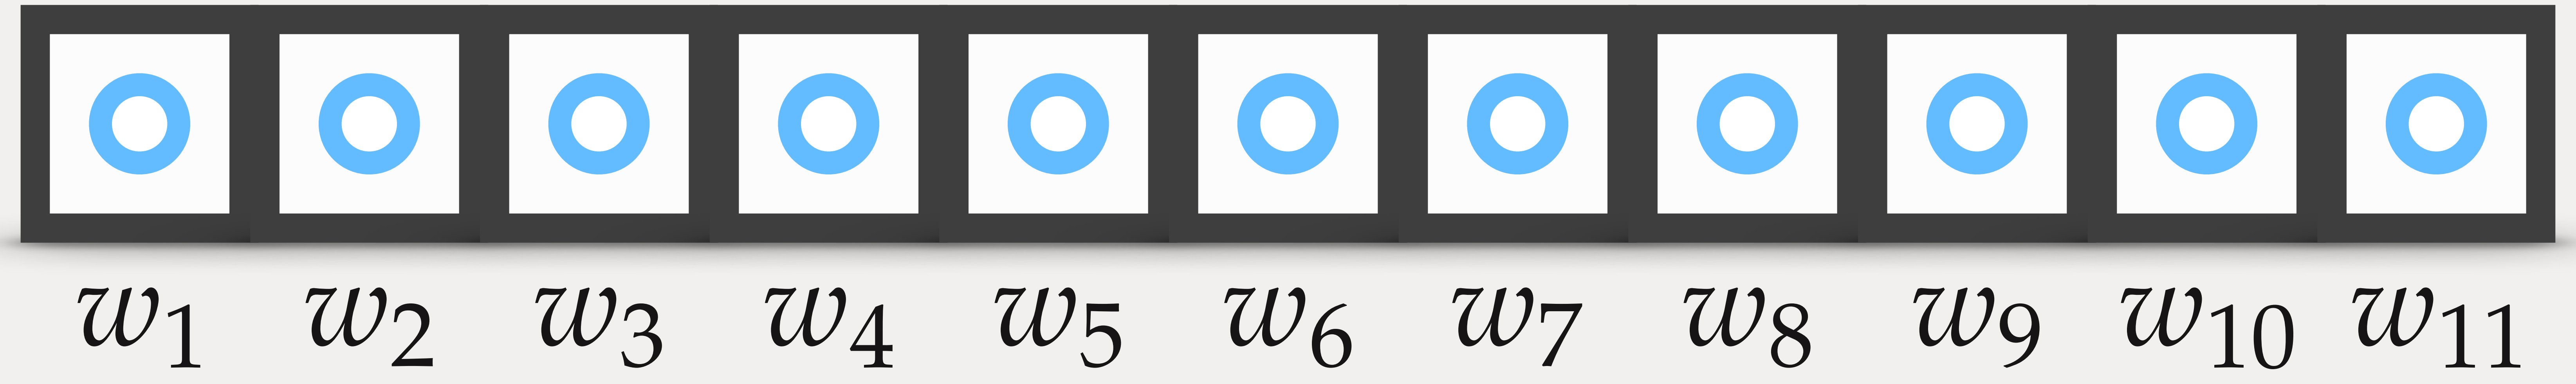

Weighted Reservoir Sampling

[M. T. Chao 1982]

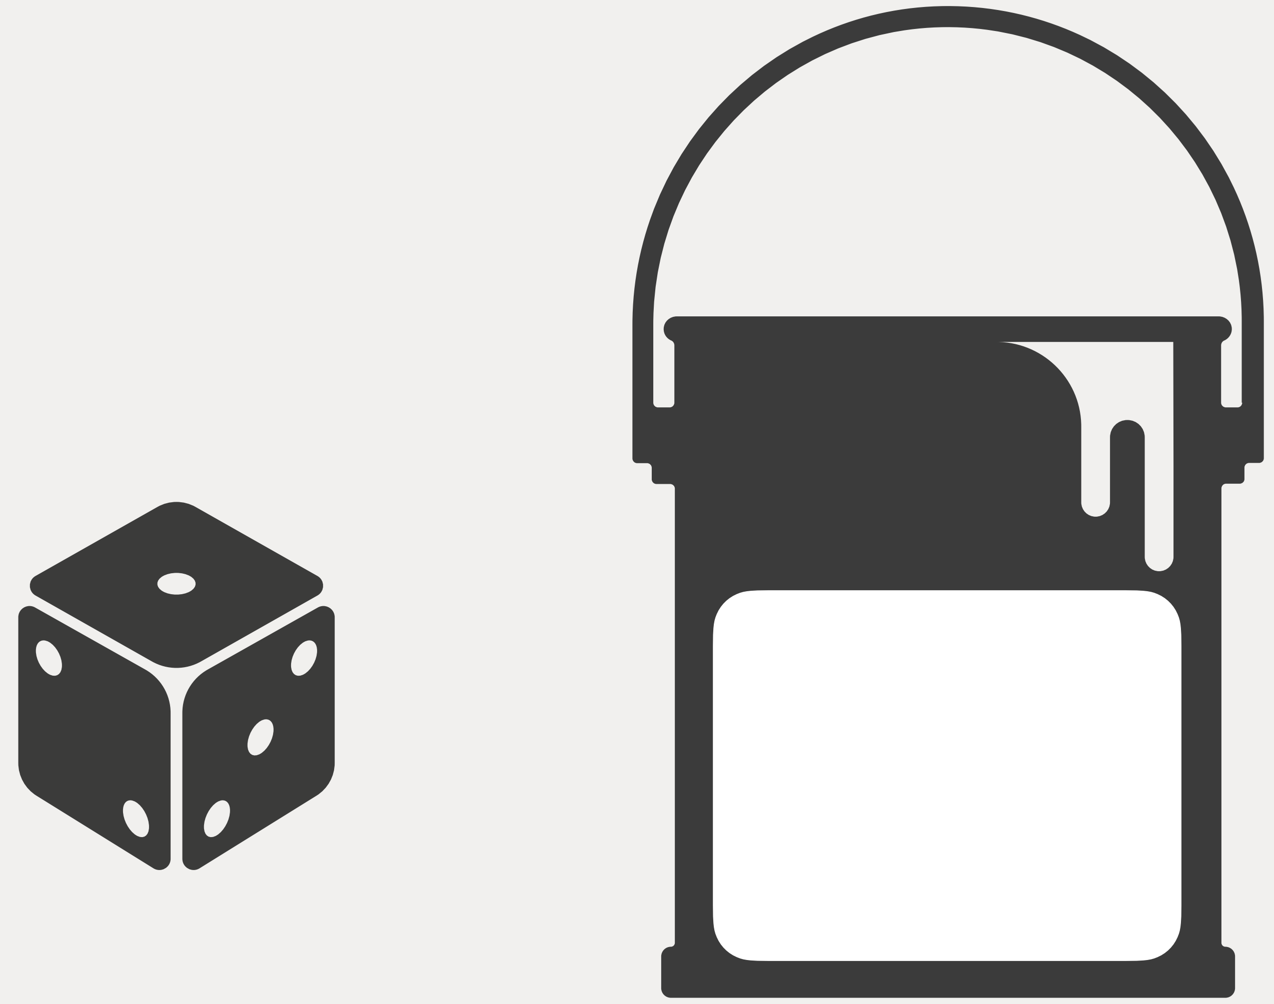

Weighted Reservoir Sampling

[M. T. Chao 1982]

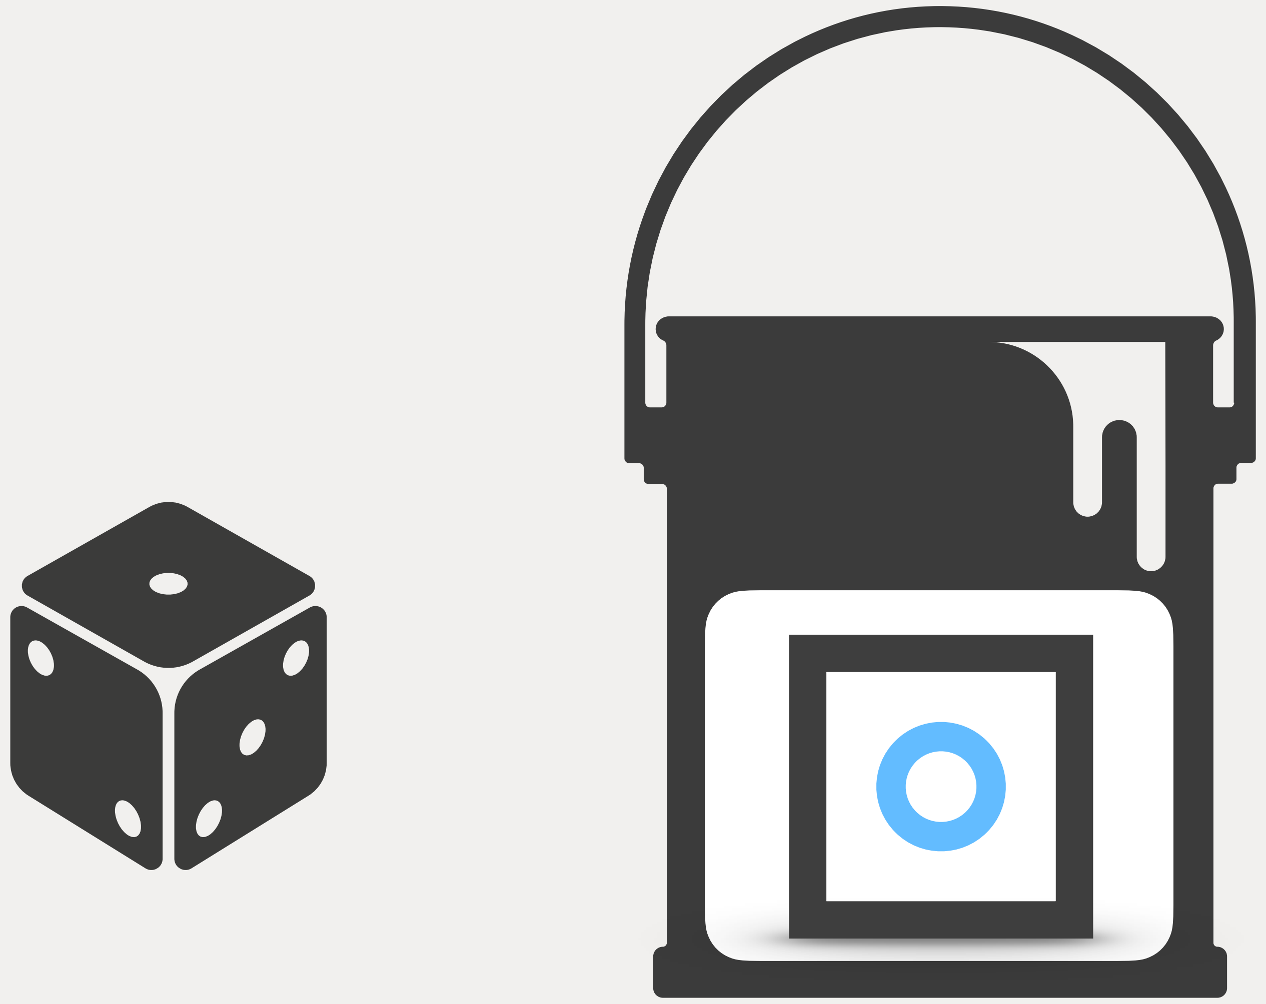

w_1

Weighted Reservoir Sampling

[M. T. Chao 1982]

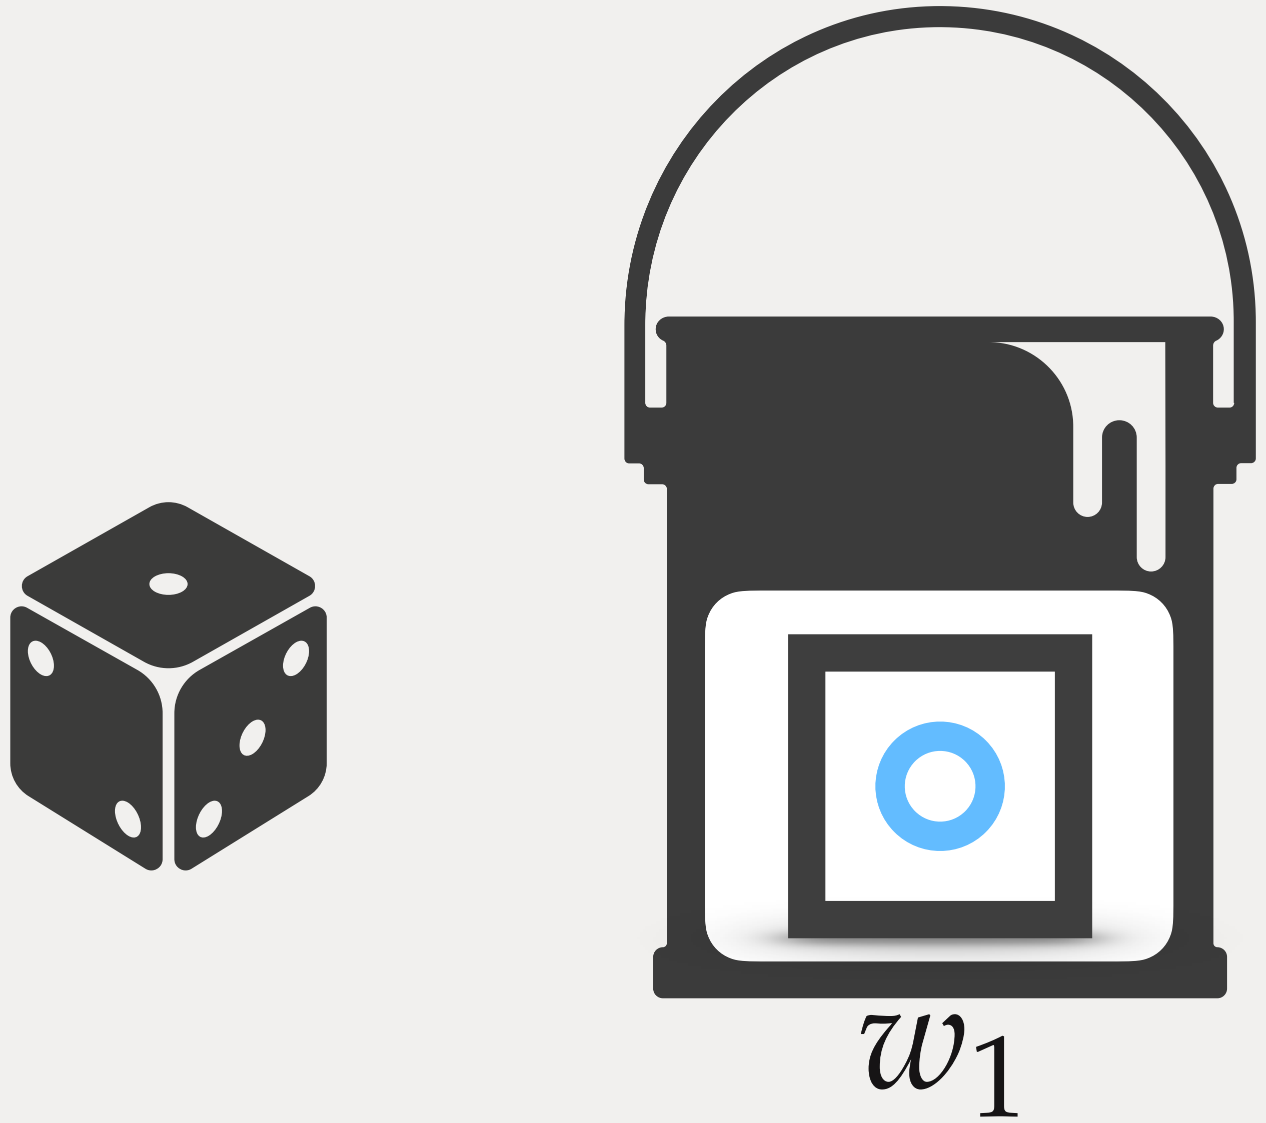

Weighted Reservoir Sampling

[M. T. Chao 1982]

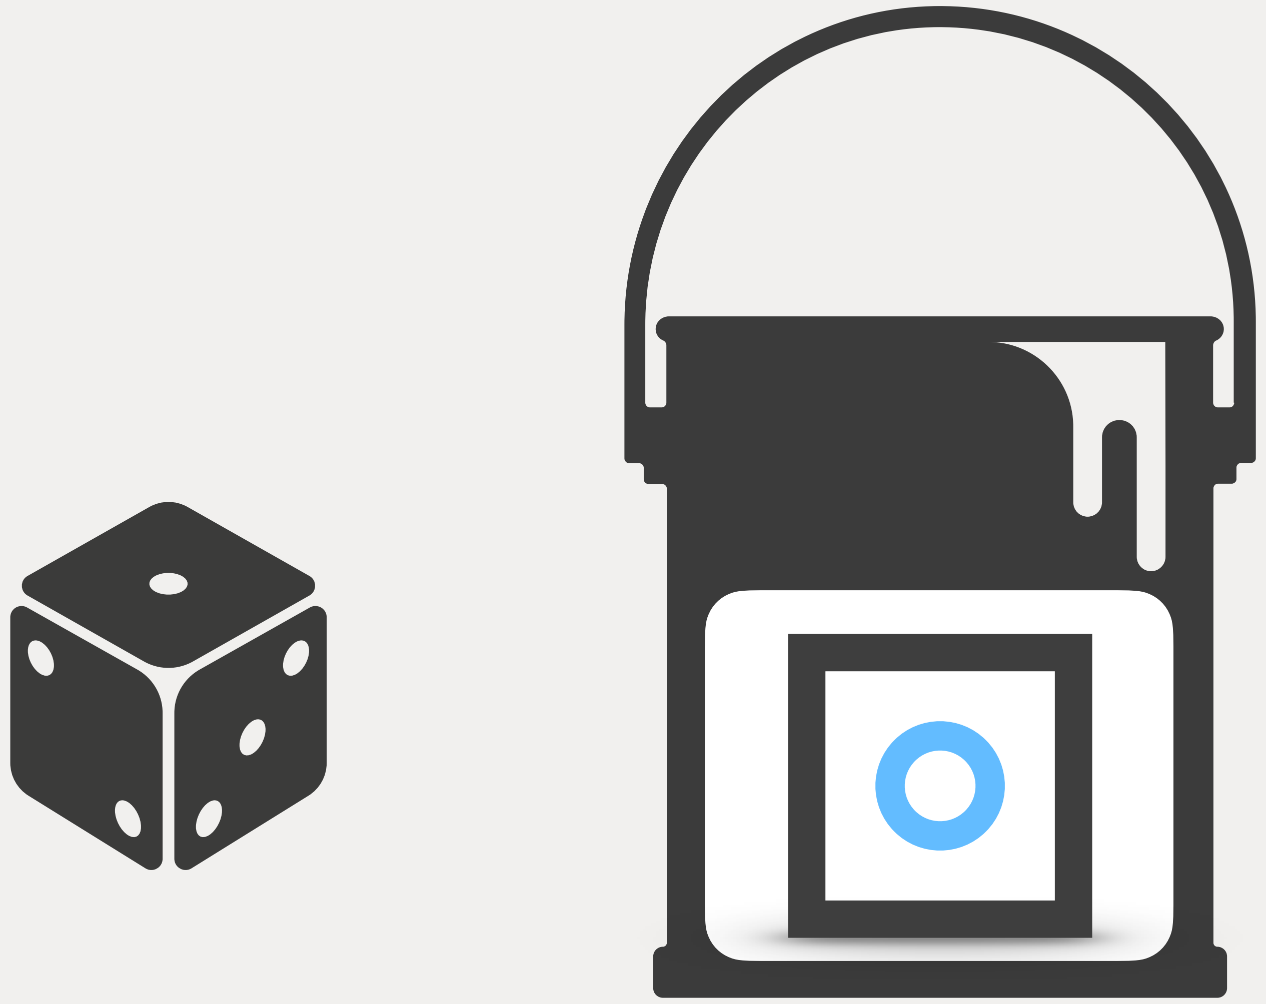

w_1

Weighted Reservoir Sampling

[M. T. Chao 1982]

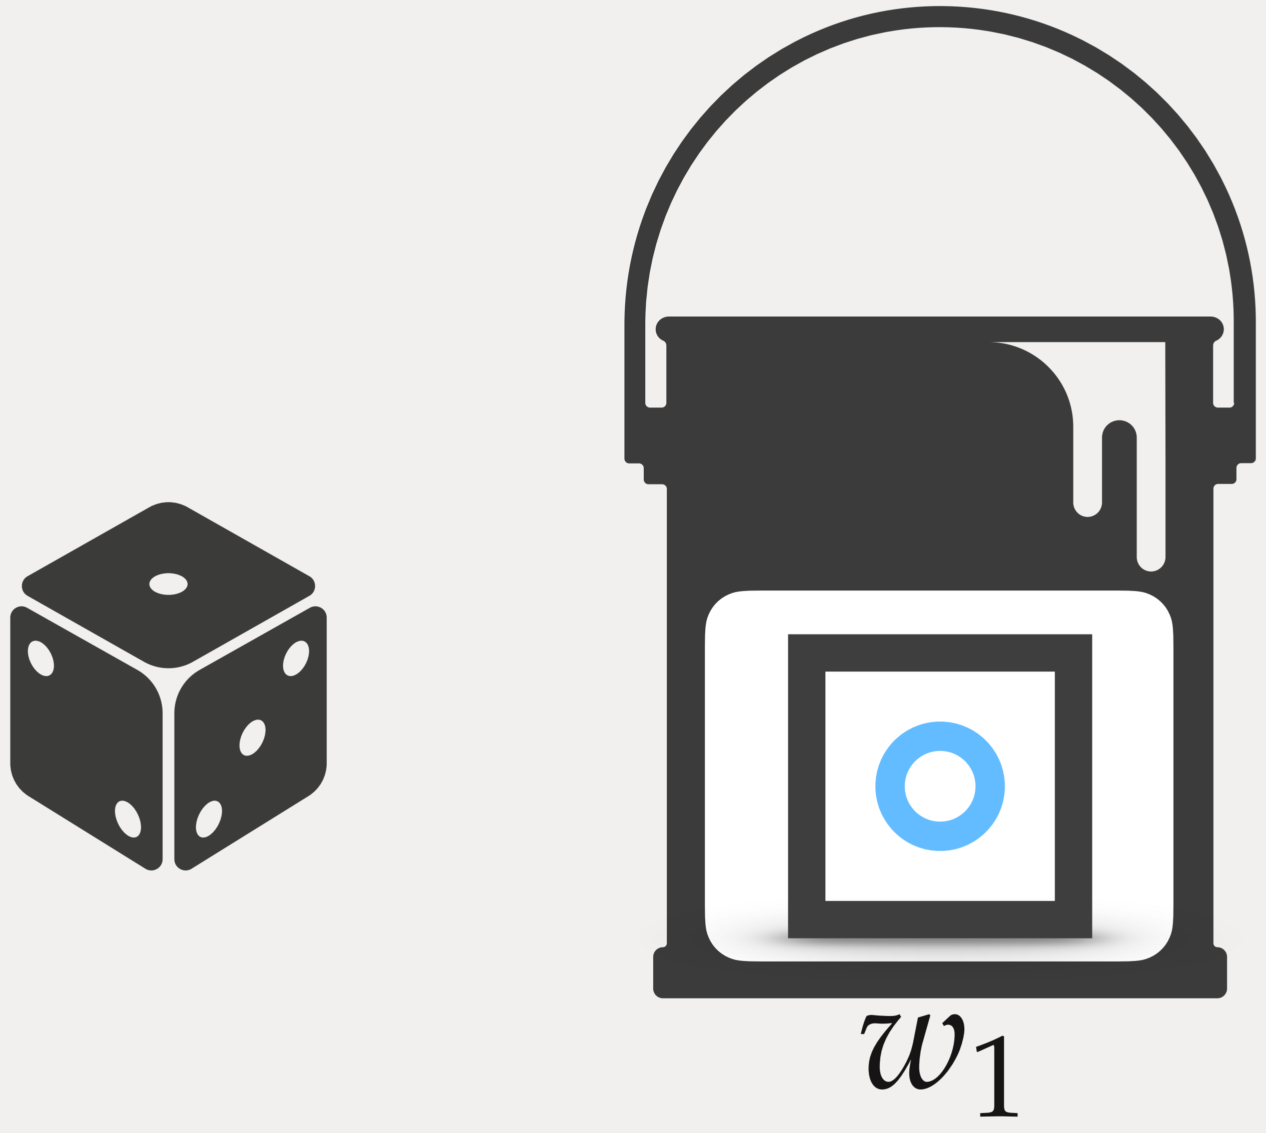

Weighted Reservoir Sampling

[M. T. Chao 1982]

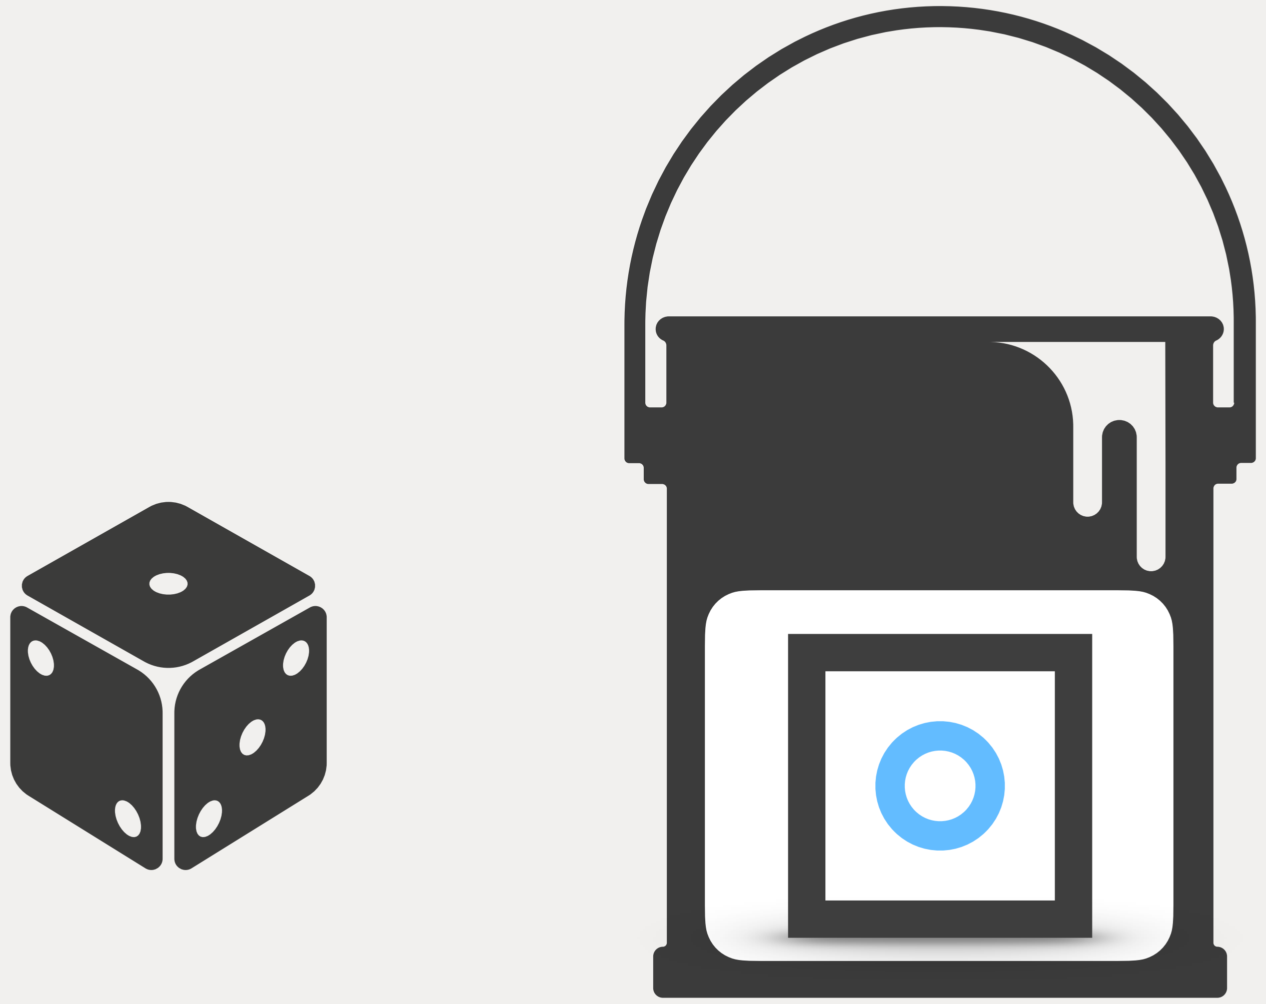

w_1

Weighted Reservoir Sampling

[M. T. Chao 1982]

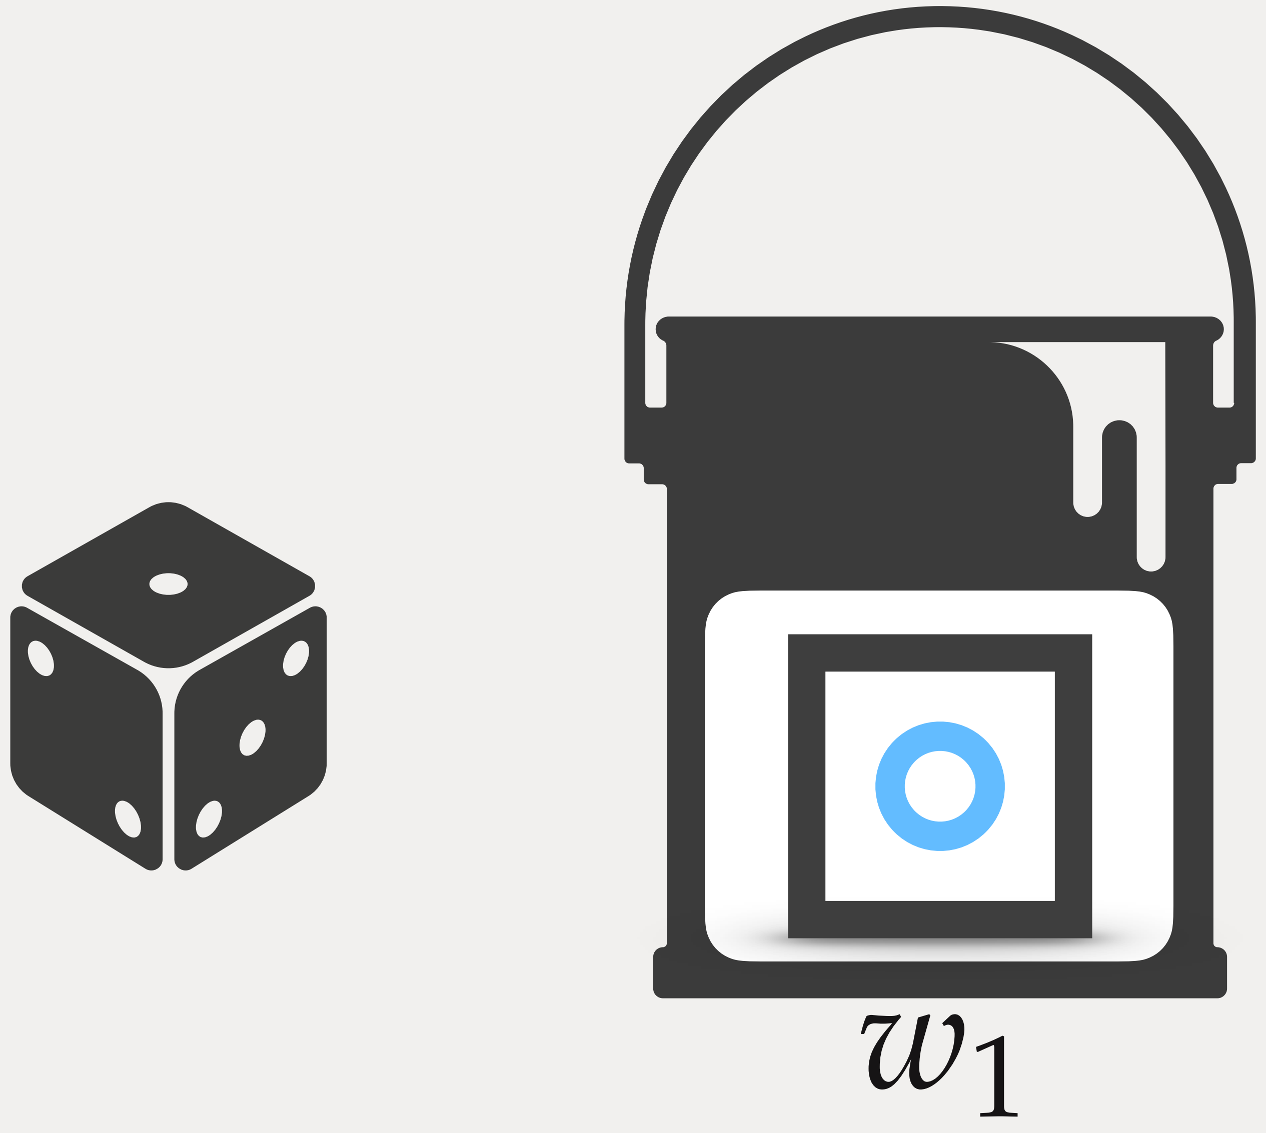

Weighted Reservoir Sampling

[M. T. Chao 1982]

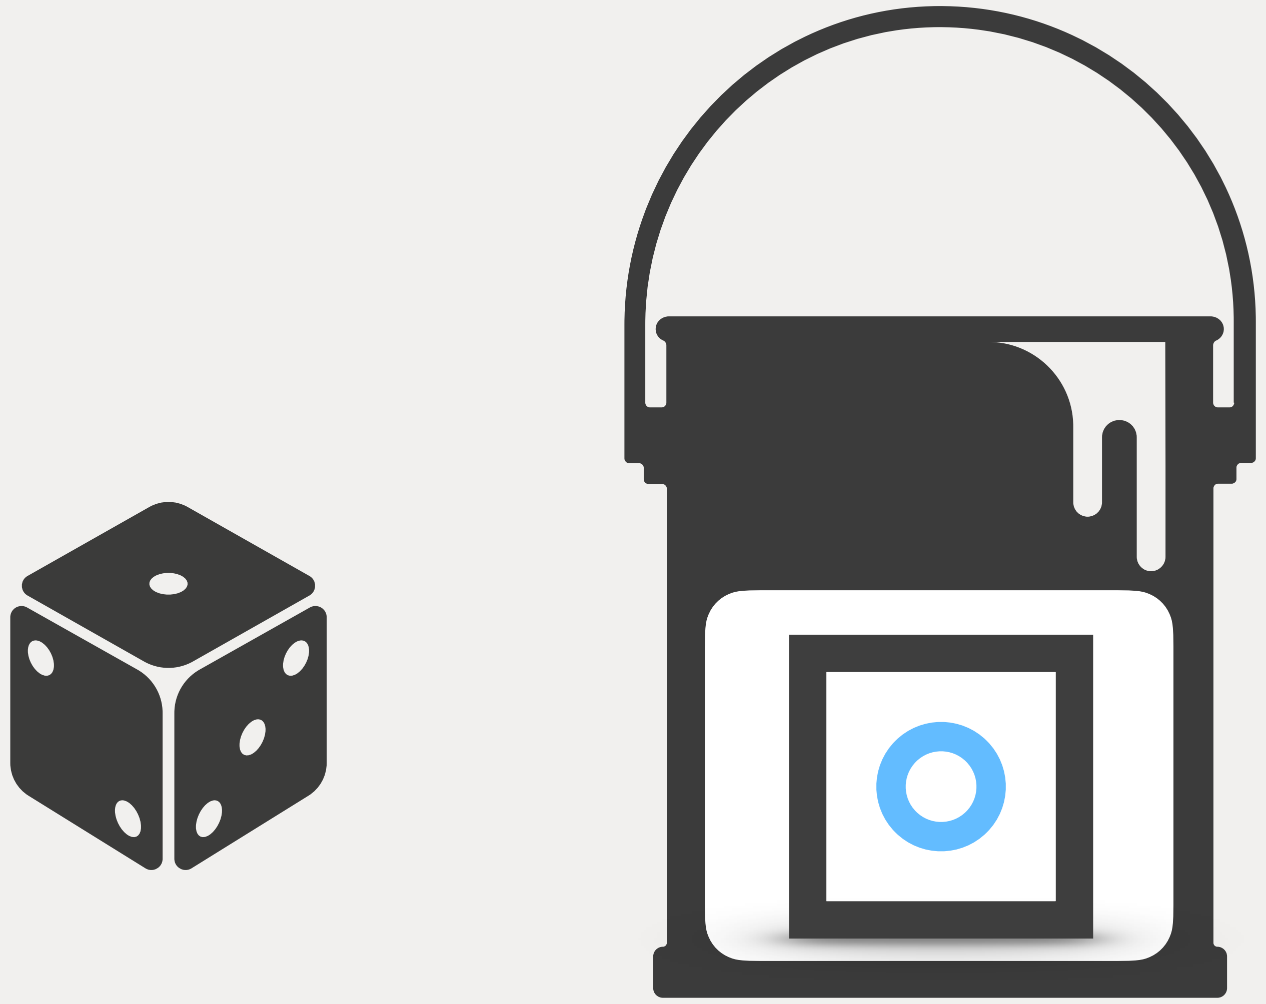

w_4

Weighted Reservoir Sampling

[M. T. Chao 1982]

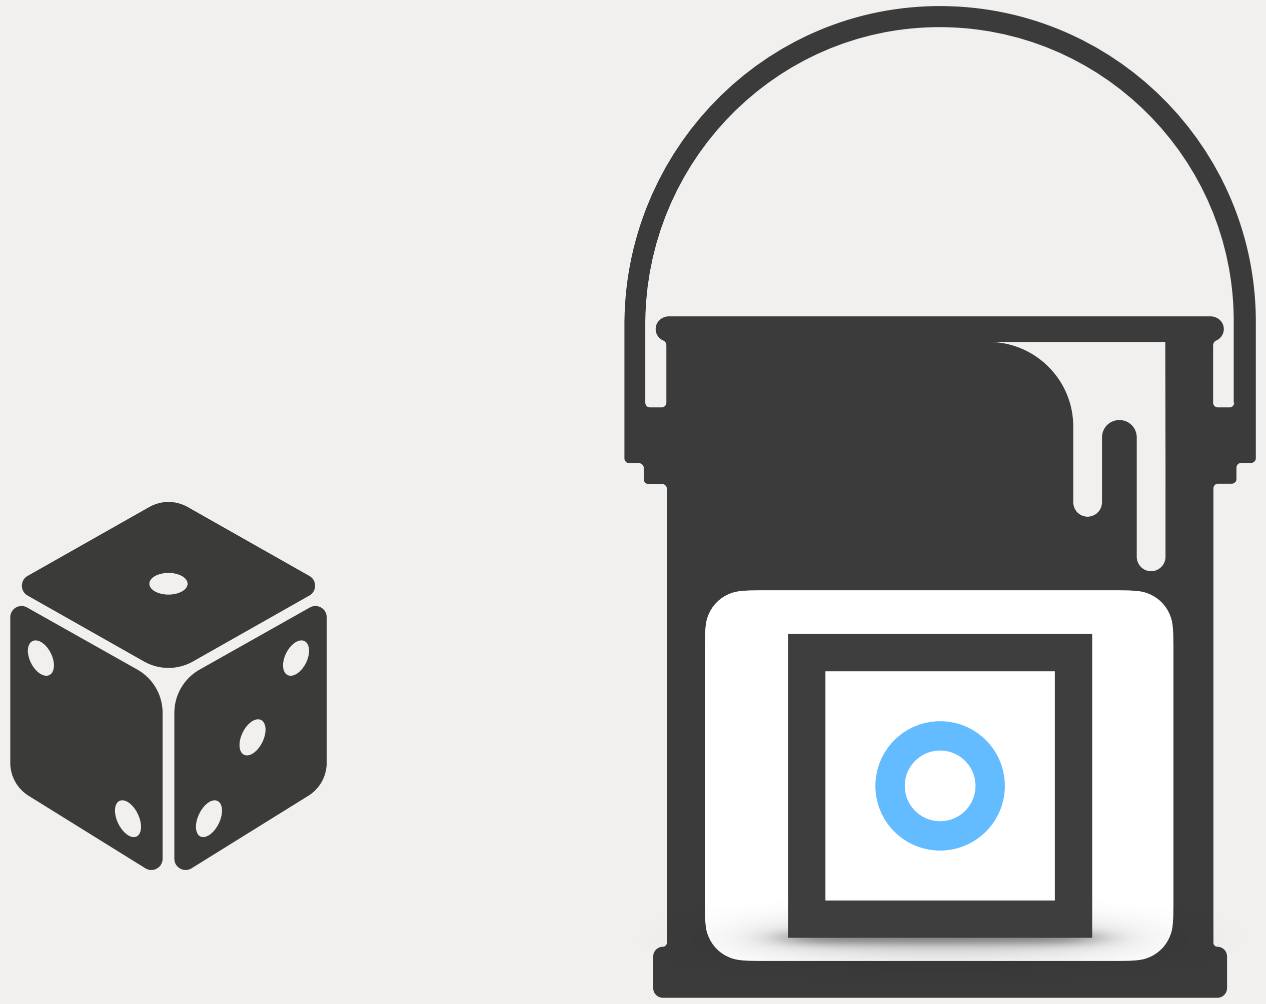

w_4

Weighted Reservoir Sampling

[M. T. Chao 1982]

w_4

$$\mathbb{P}[\text{Accept } i] = \frac{w_i}{\sum_j w_j}$$

Weighted Reservoir Sampling

[M. T. Chao 1982]

w_4

$$\mathbb{P}[\text{Accept } i] = \frac{w_i}{\sum_j w_j}$$

Weighted Reservoir Sampling

[M. T. Chao 1982]

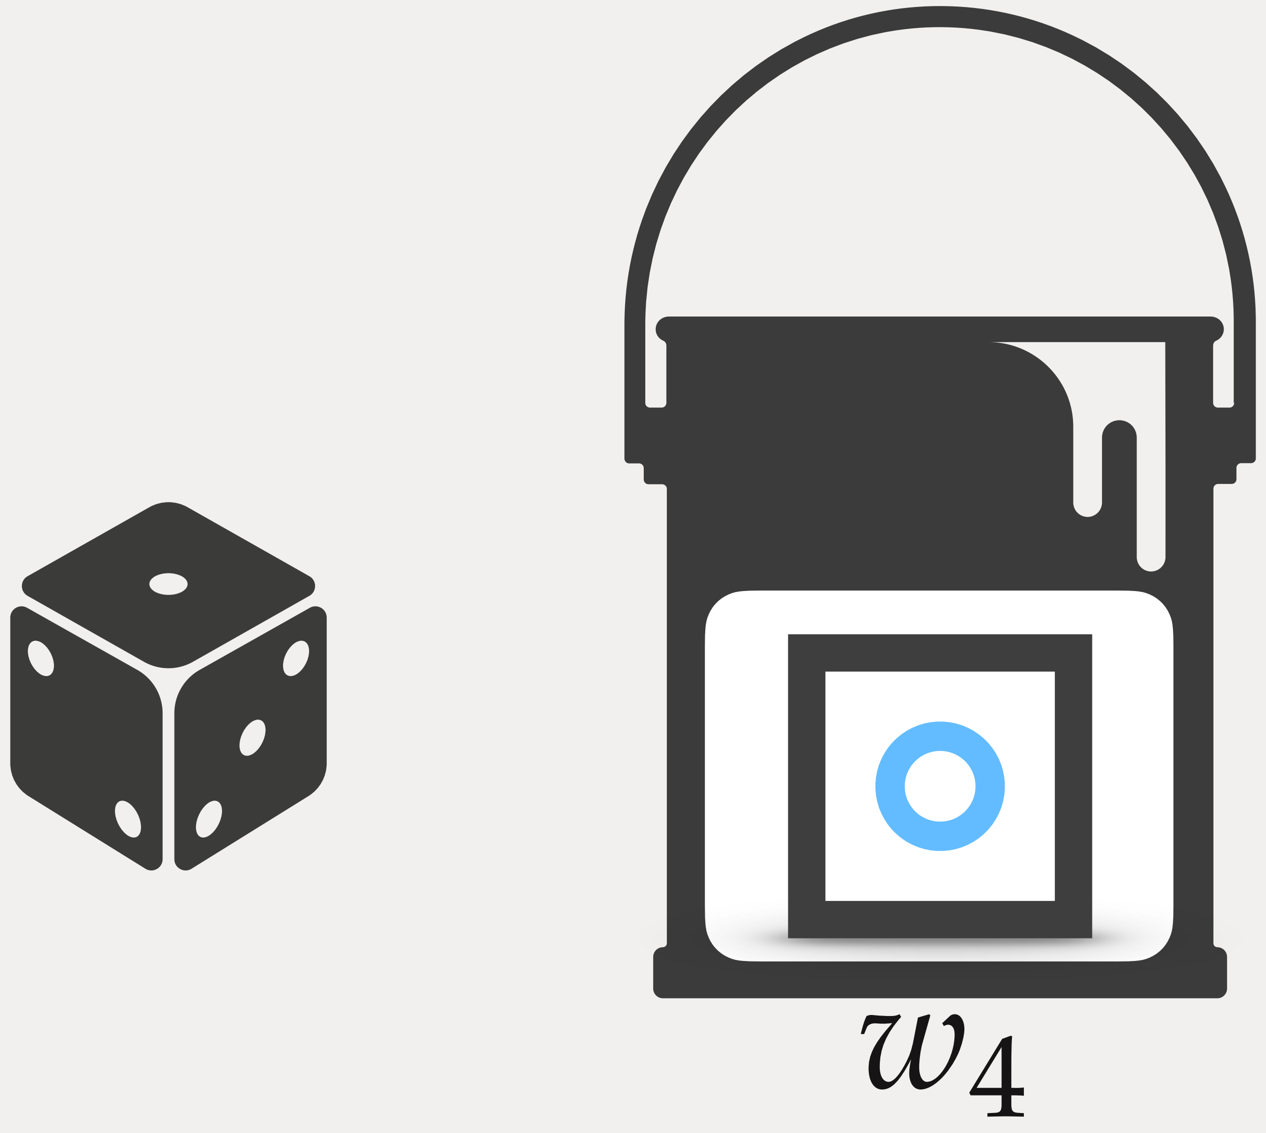

$$\mathbb{P}[\text{Accept } i] = \frac{w_i}{\sum_j w_j}$$

Weighted Reservoir Sampling

[M. T. Chao 1982]

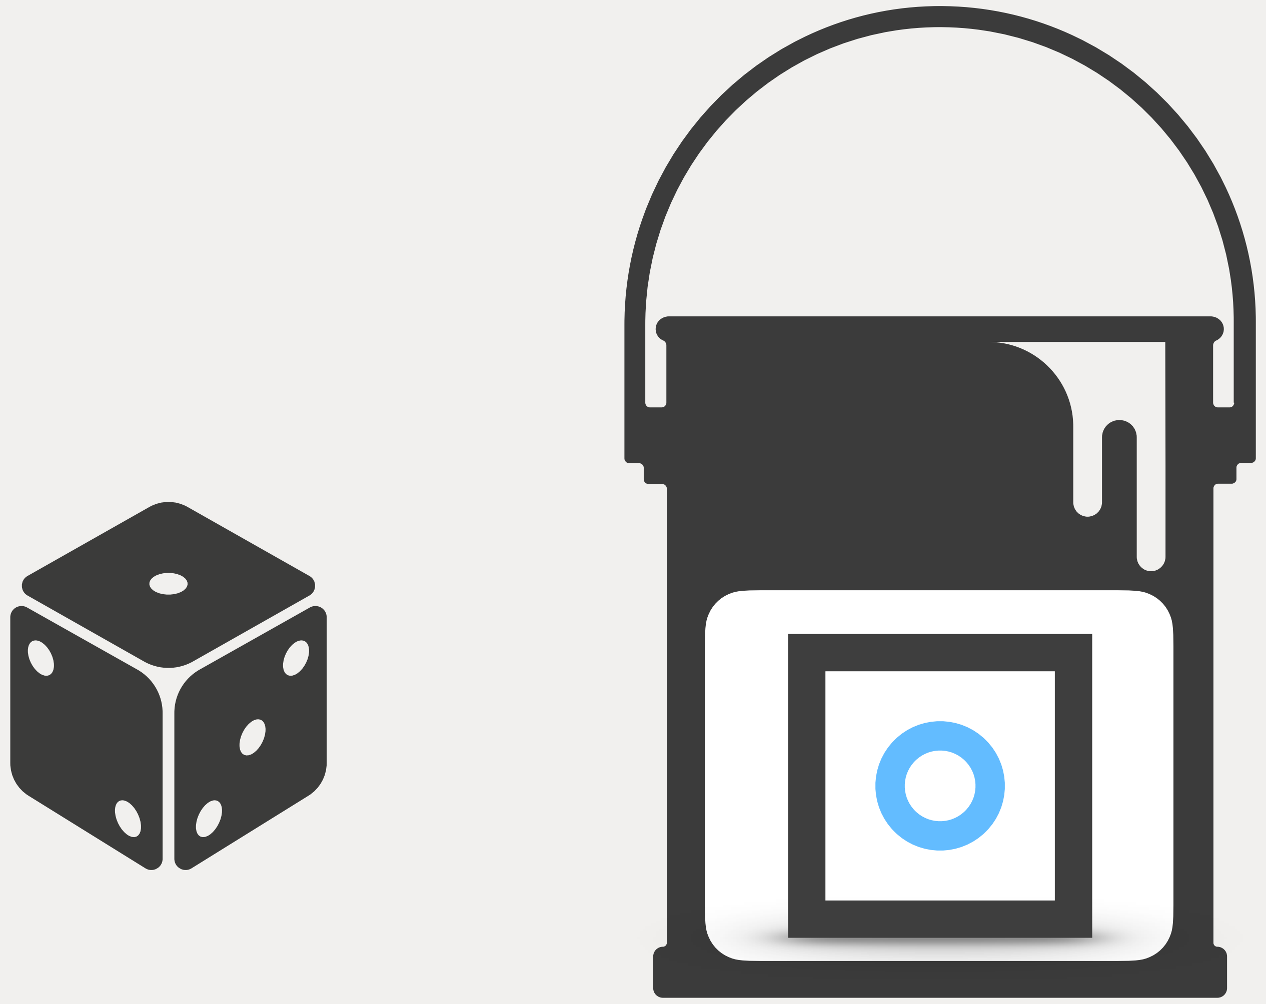

w_4

$$\mathbb{P}[\text{Accept } i] = \frac{w_i}{\sum_j w_j}$$

Weighted Reservoir Sampling

[M. T. Chao 1982]

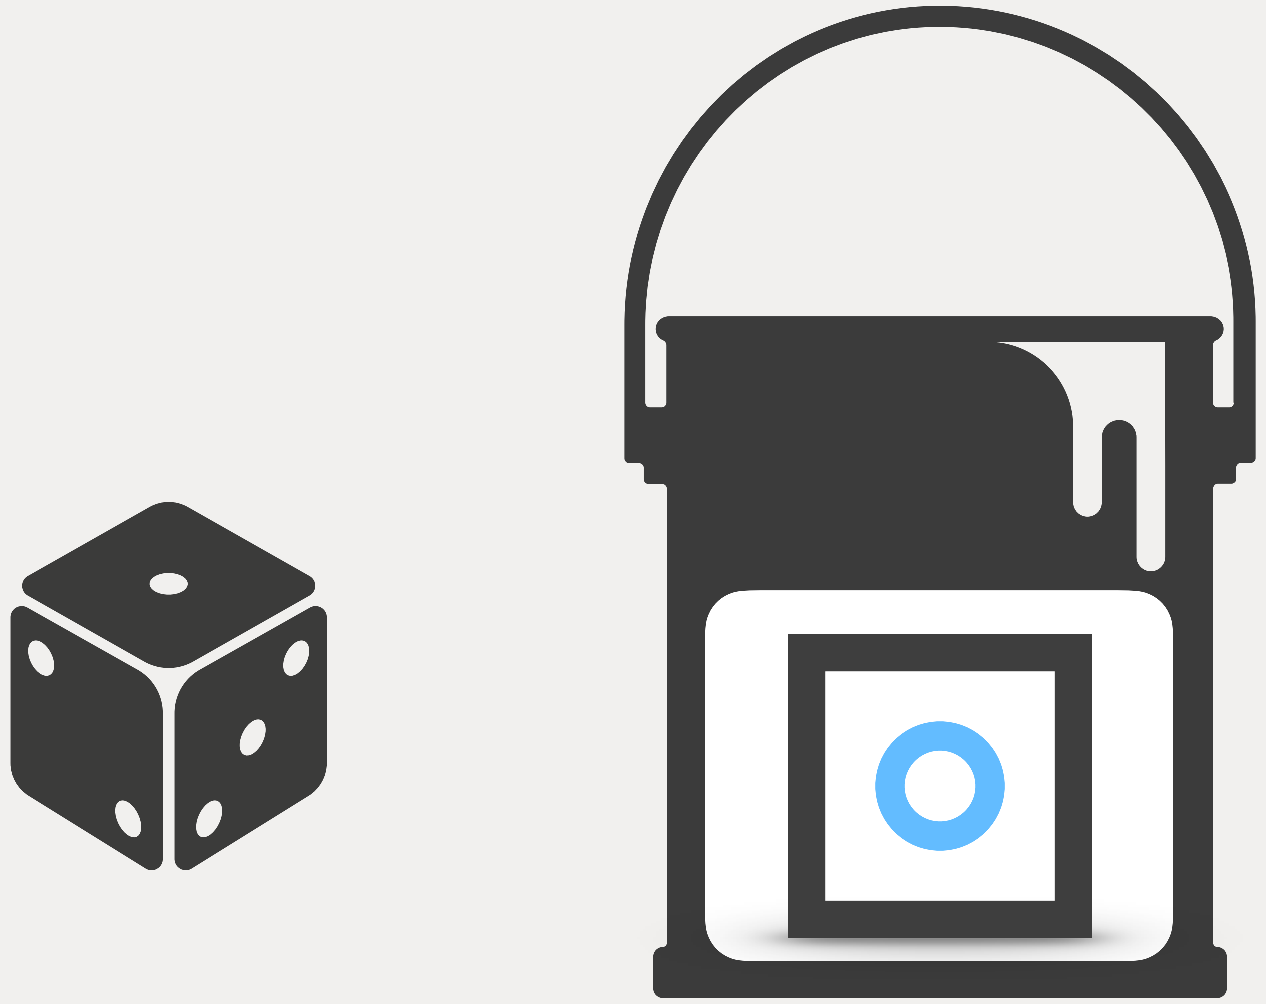

w_4

$$\mathbb{P}[\text{Accept } i] = \frac{w_i}{\sum_j w_j}$$

Reservoir Resampling

Reservoir Resampling

- Constant memory

Reservoir Resampling

- Constant memory
- Computation: $O(M)$

Reservoir Resampling

- Constant memory
- Computation: $O(M)$
- Can we do better?

Reservoir Resampling

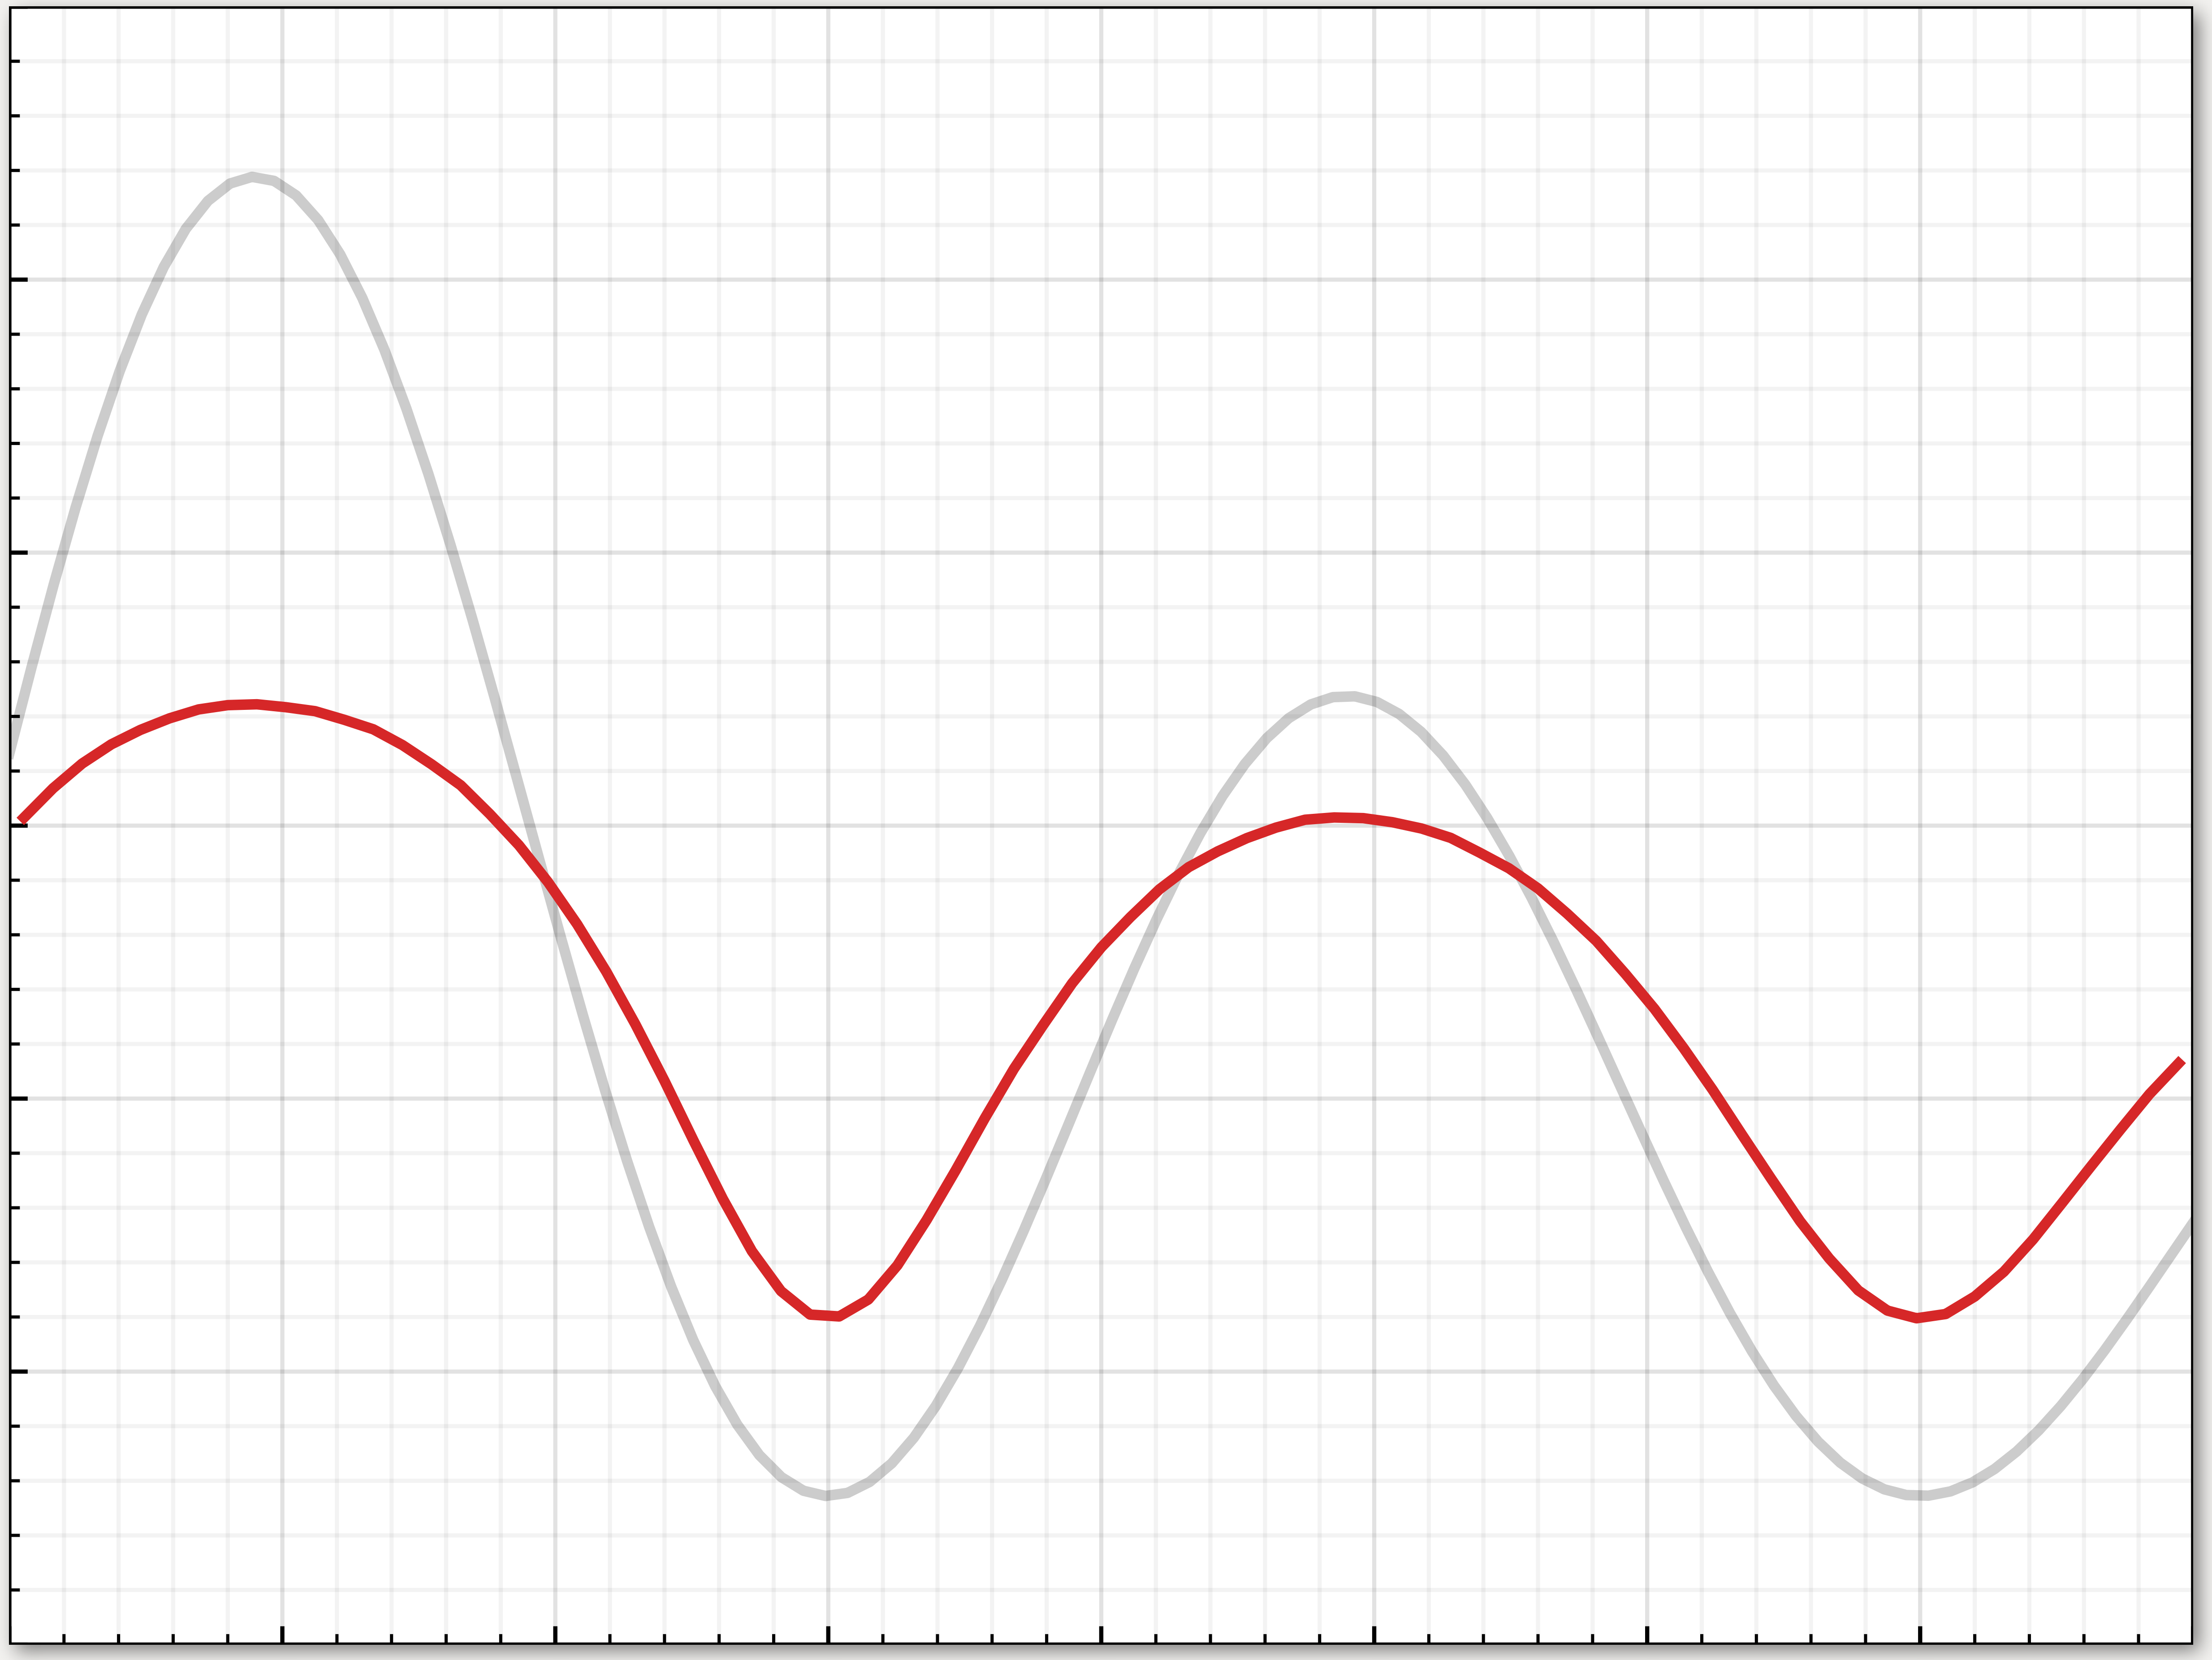

Improved Reservoir Resampling

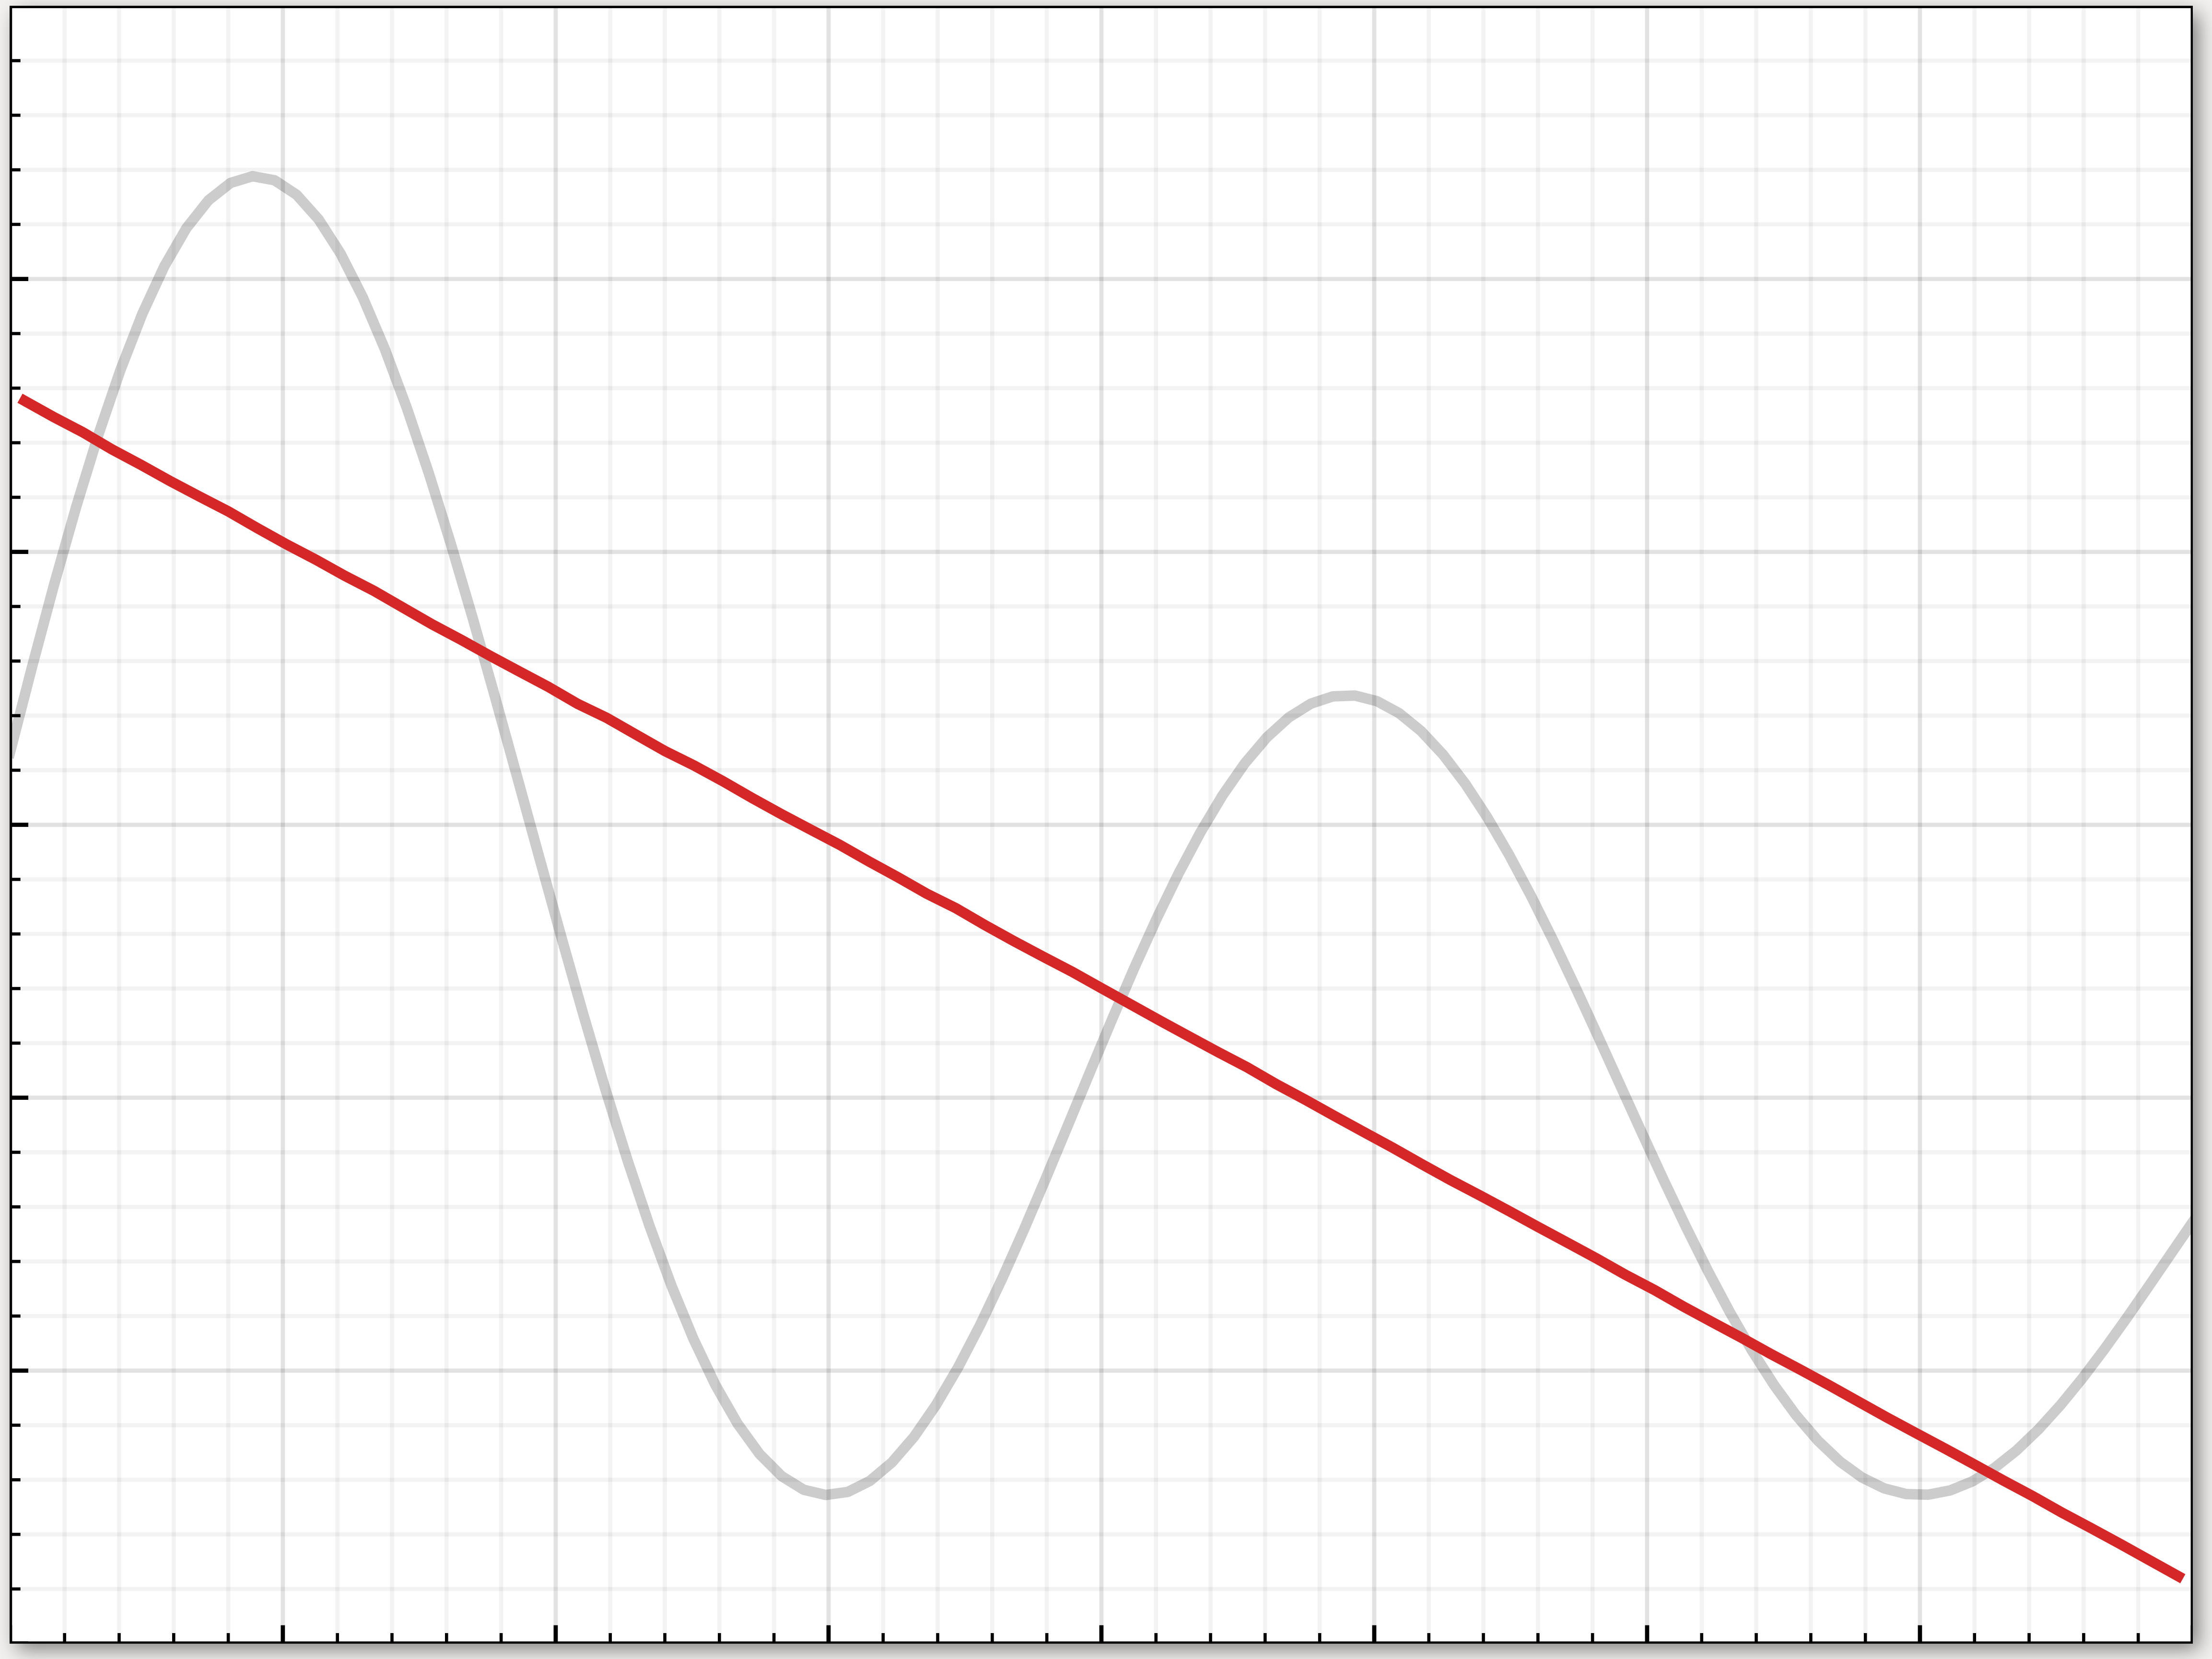

$p_1(x)$

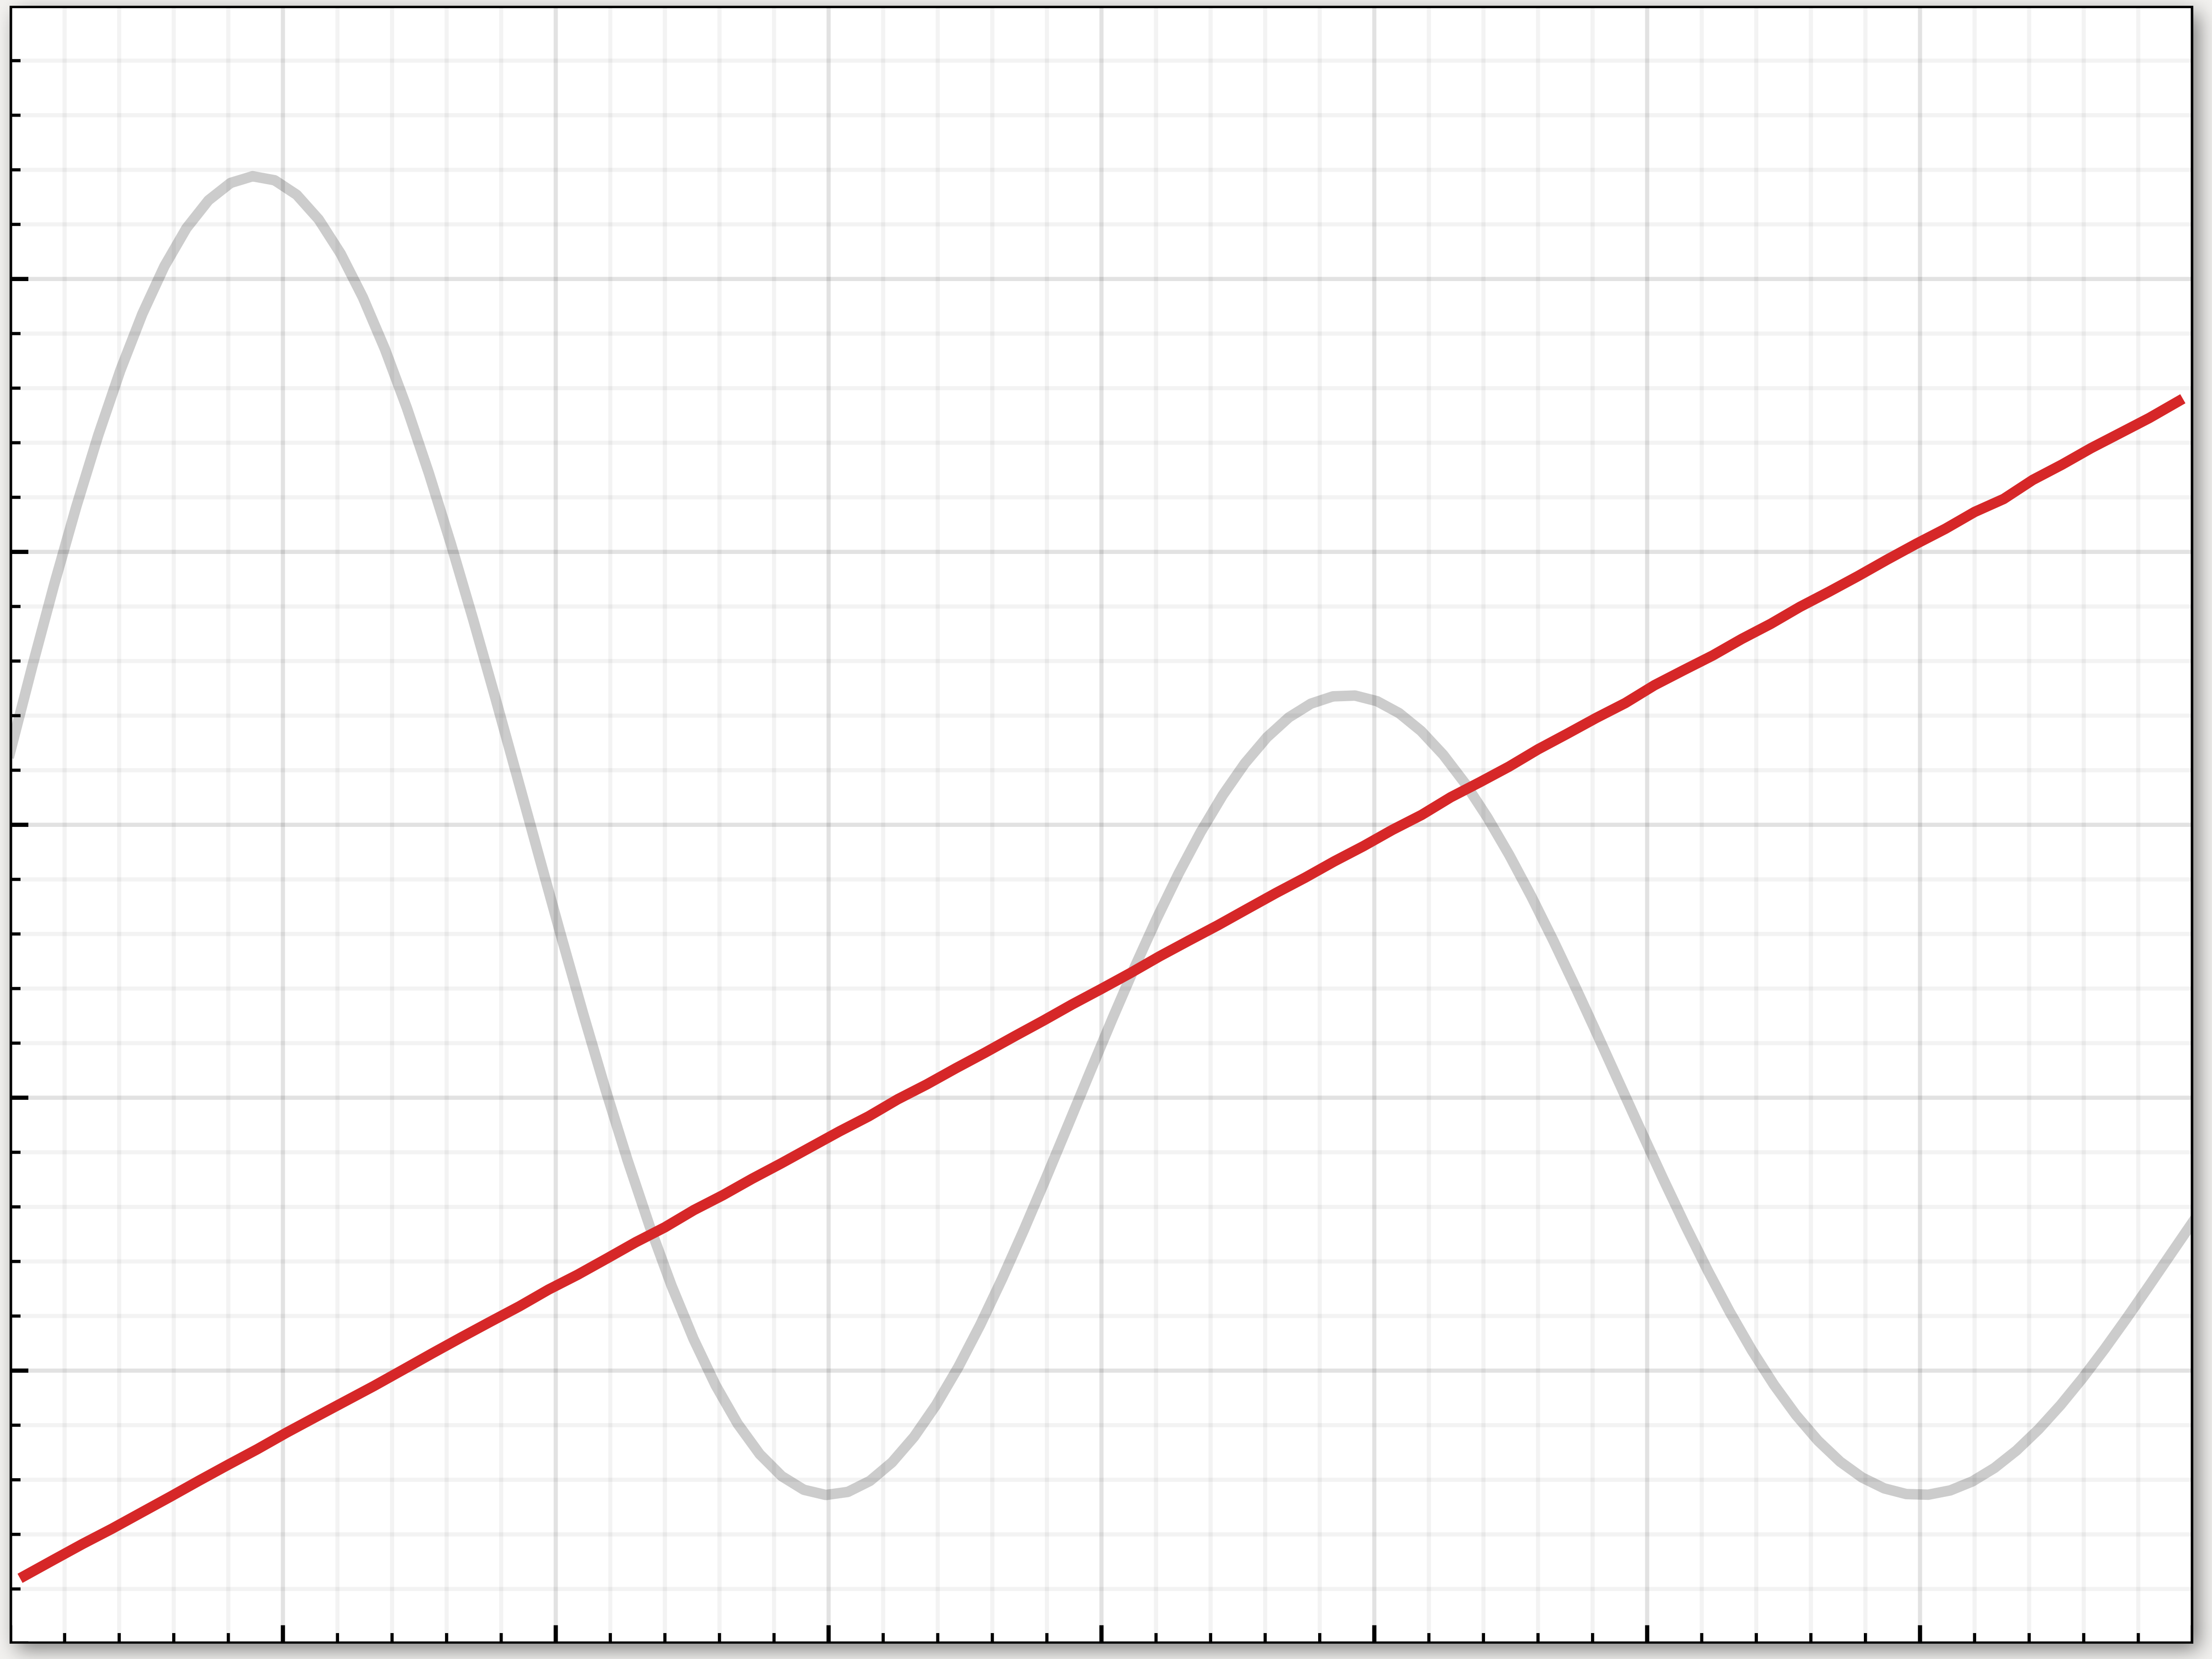

$p_2(x)$

Improved Reservoir Resampling

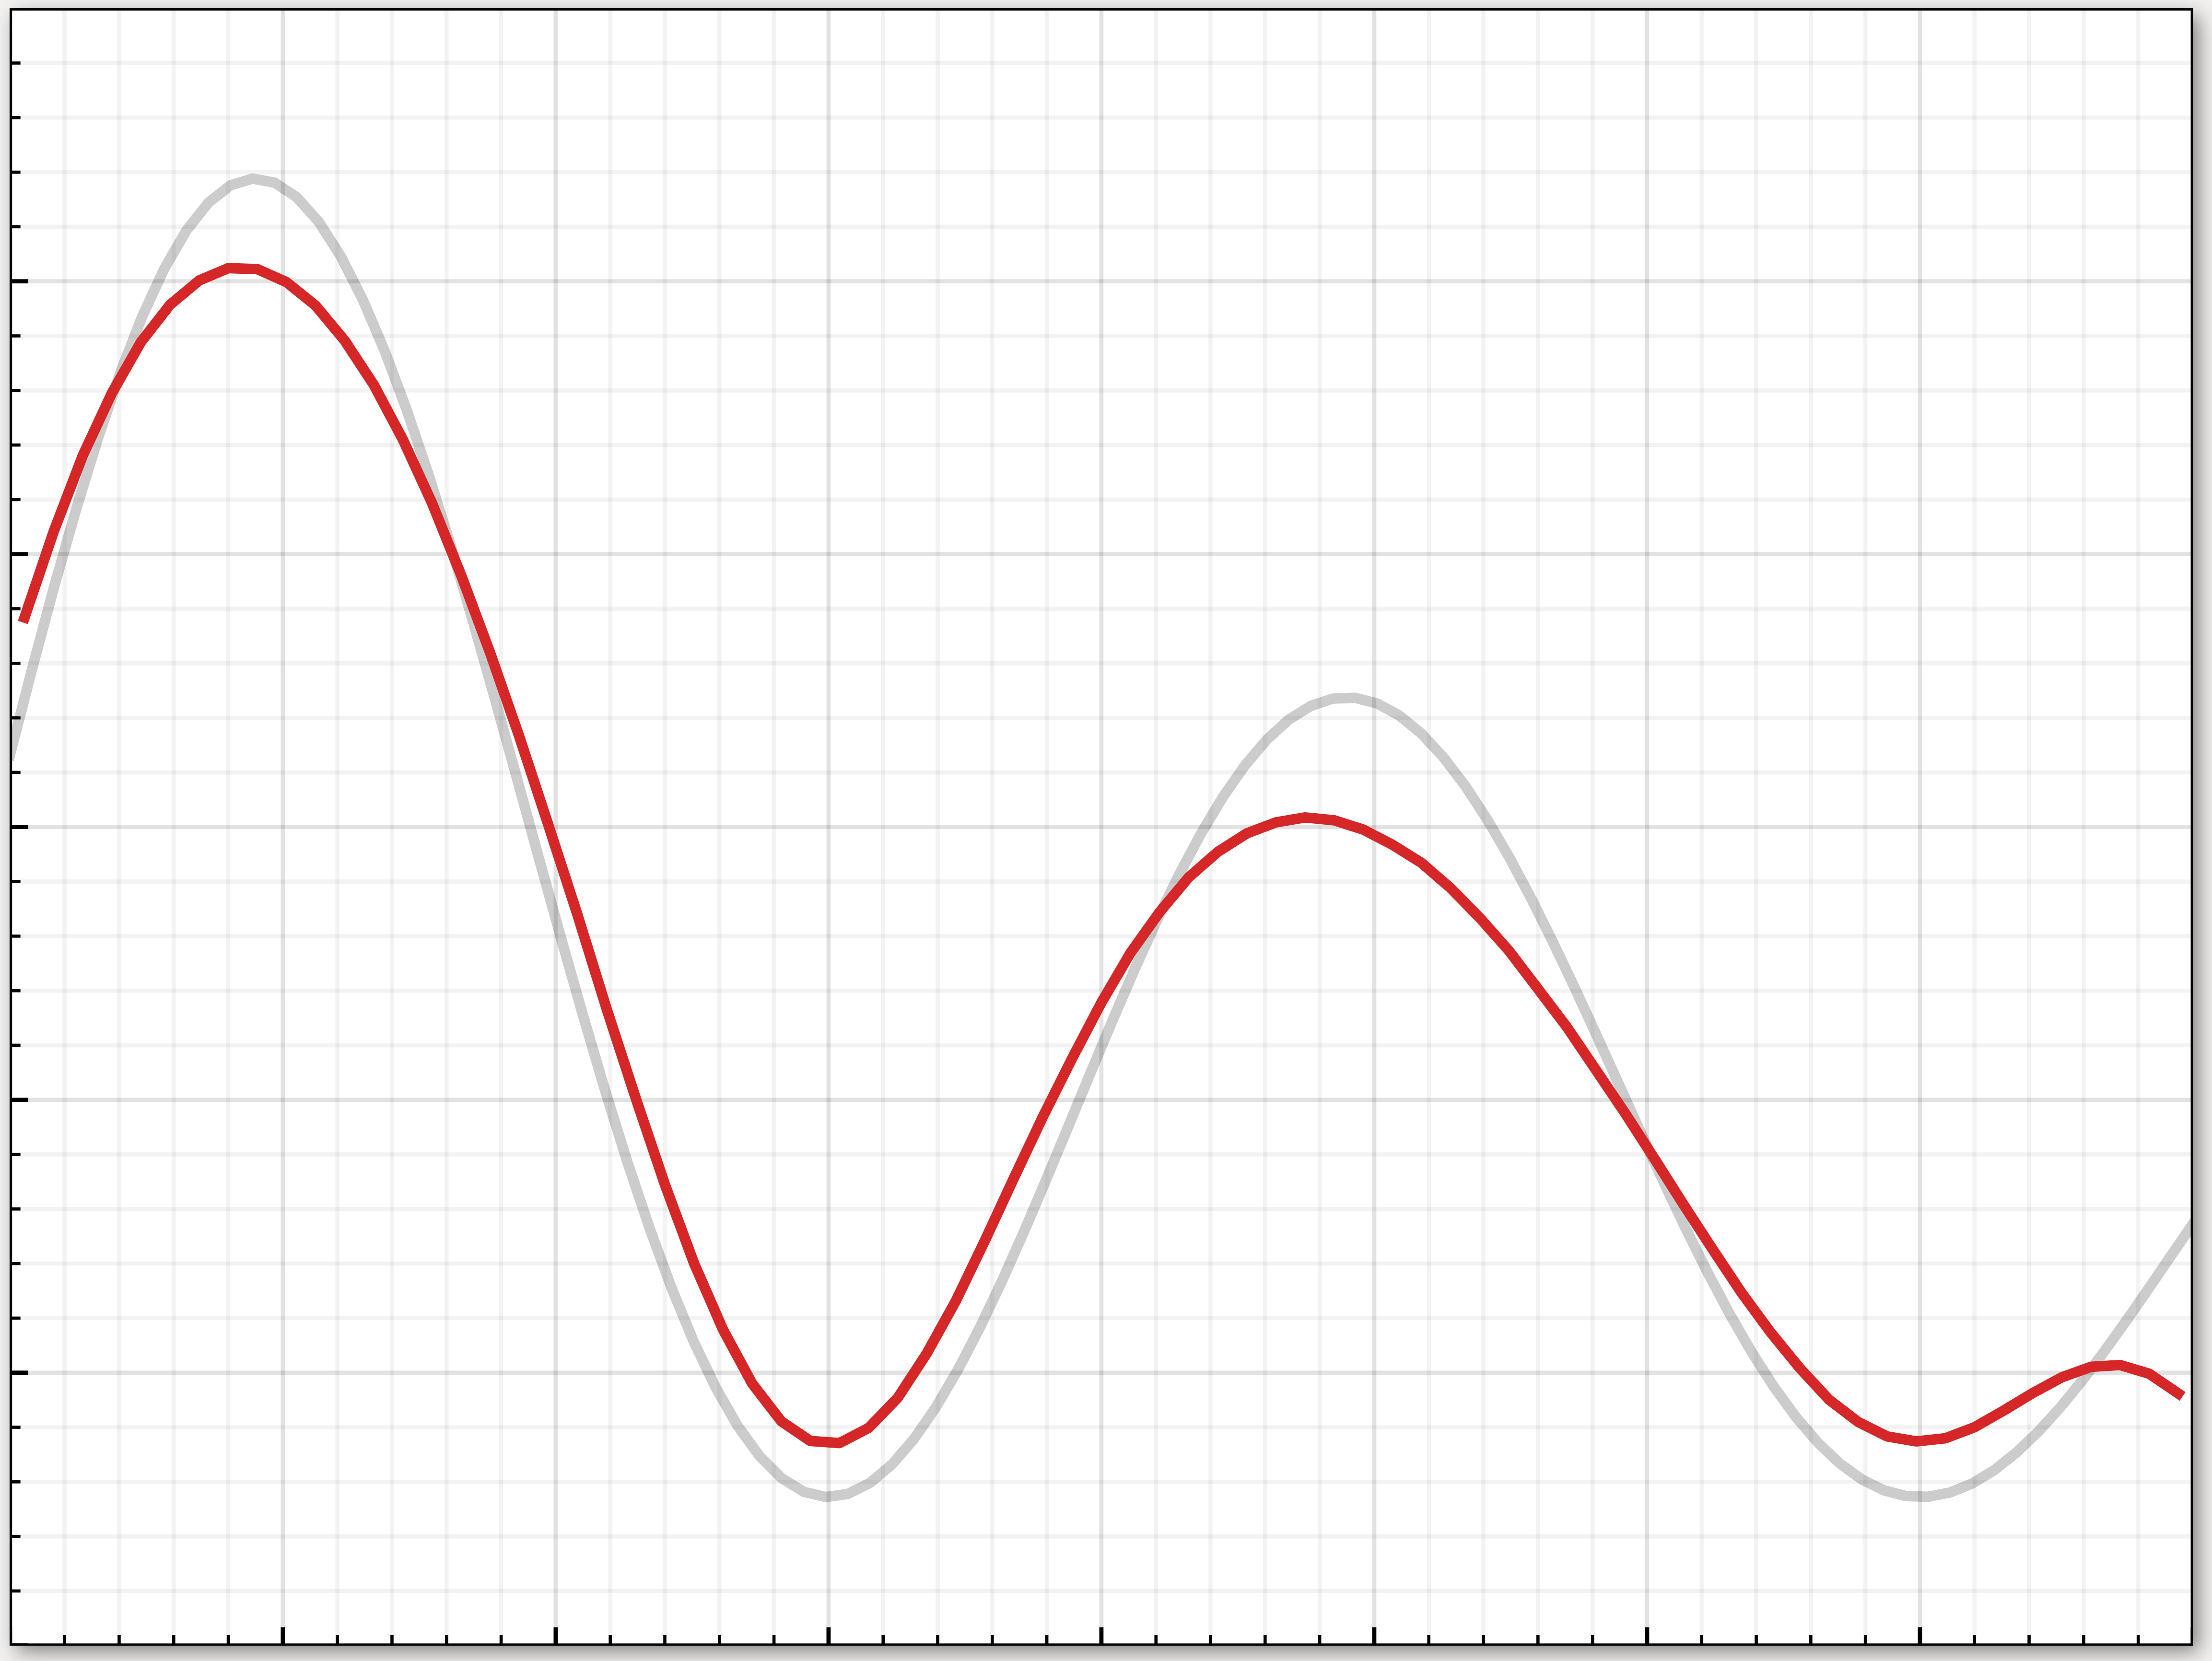

$p_1(x)$

$M = 4$

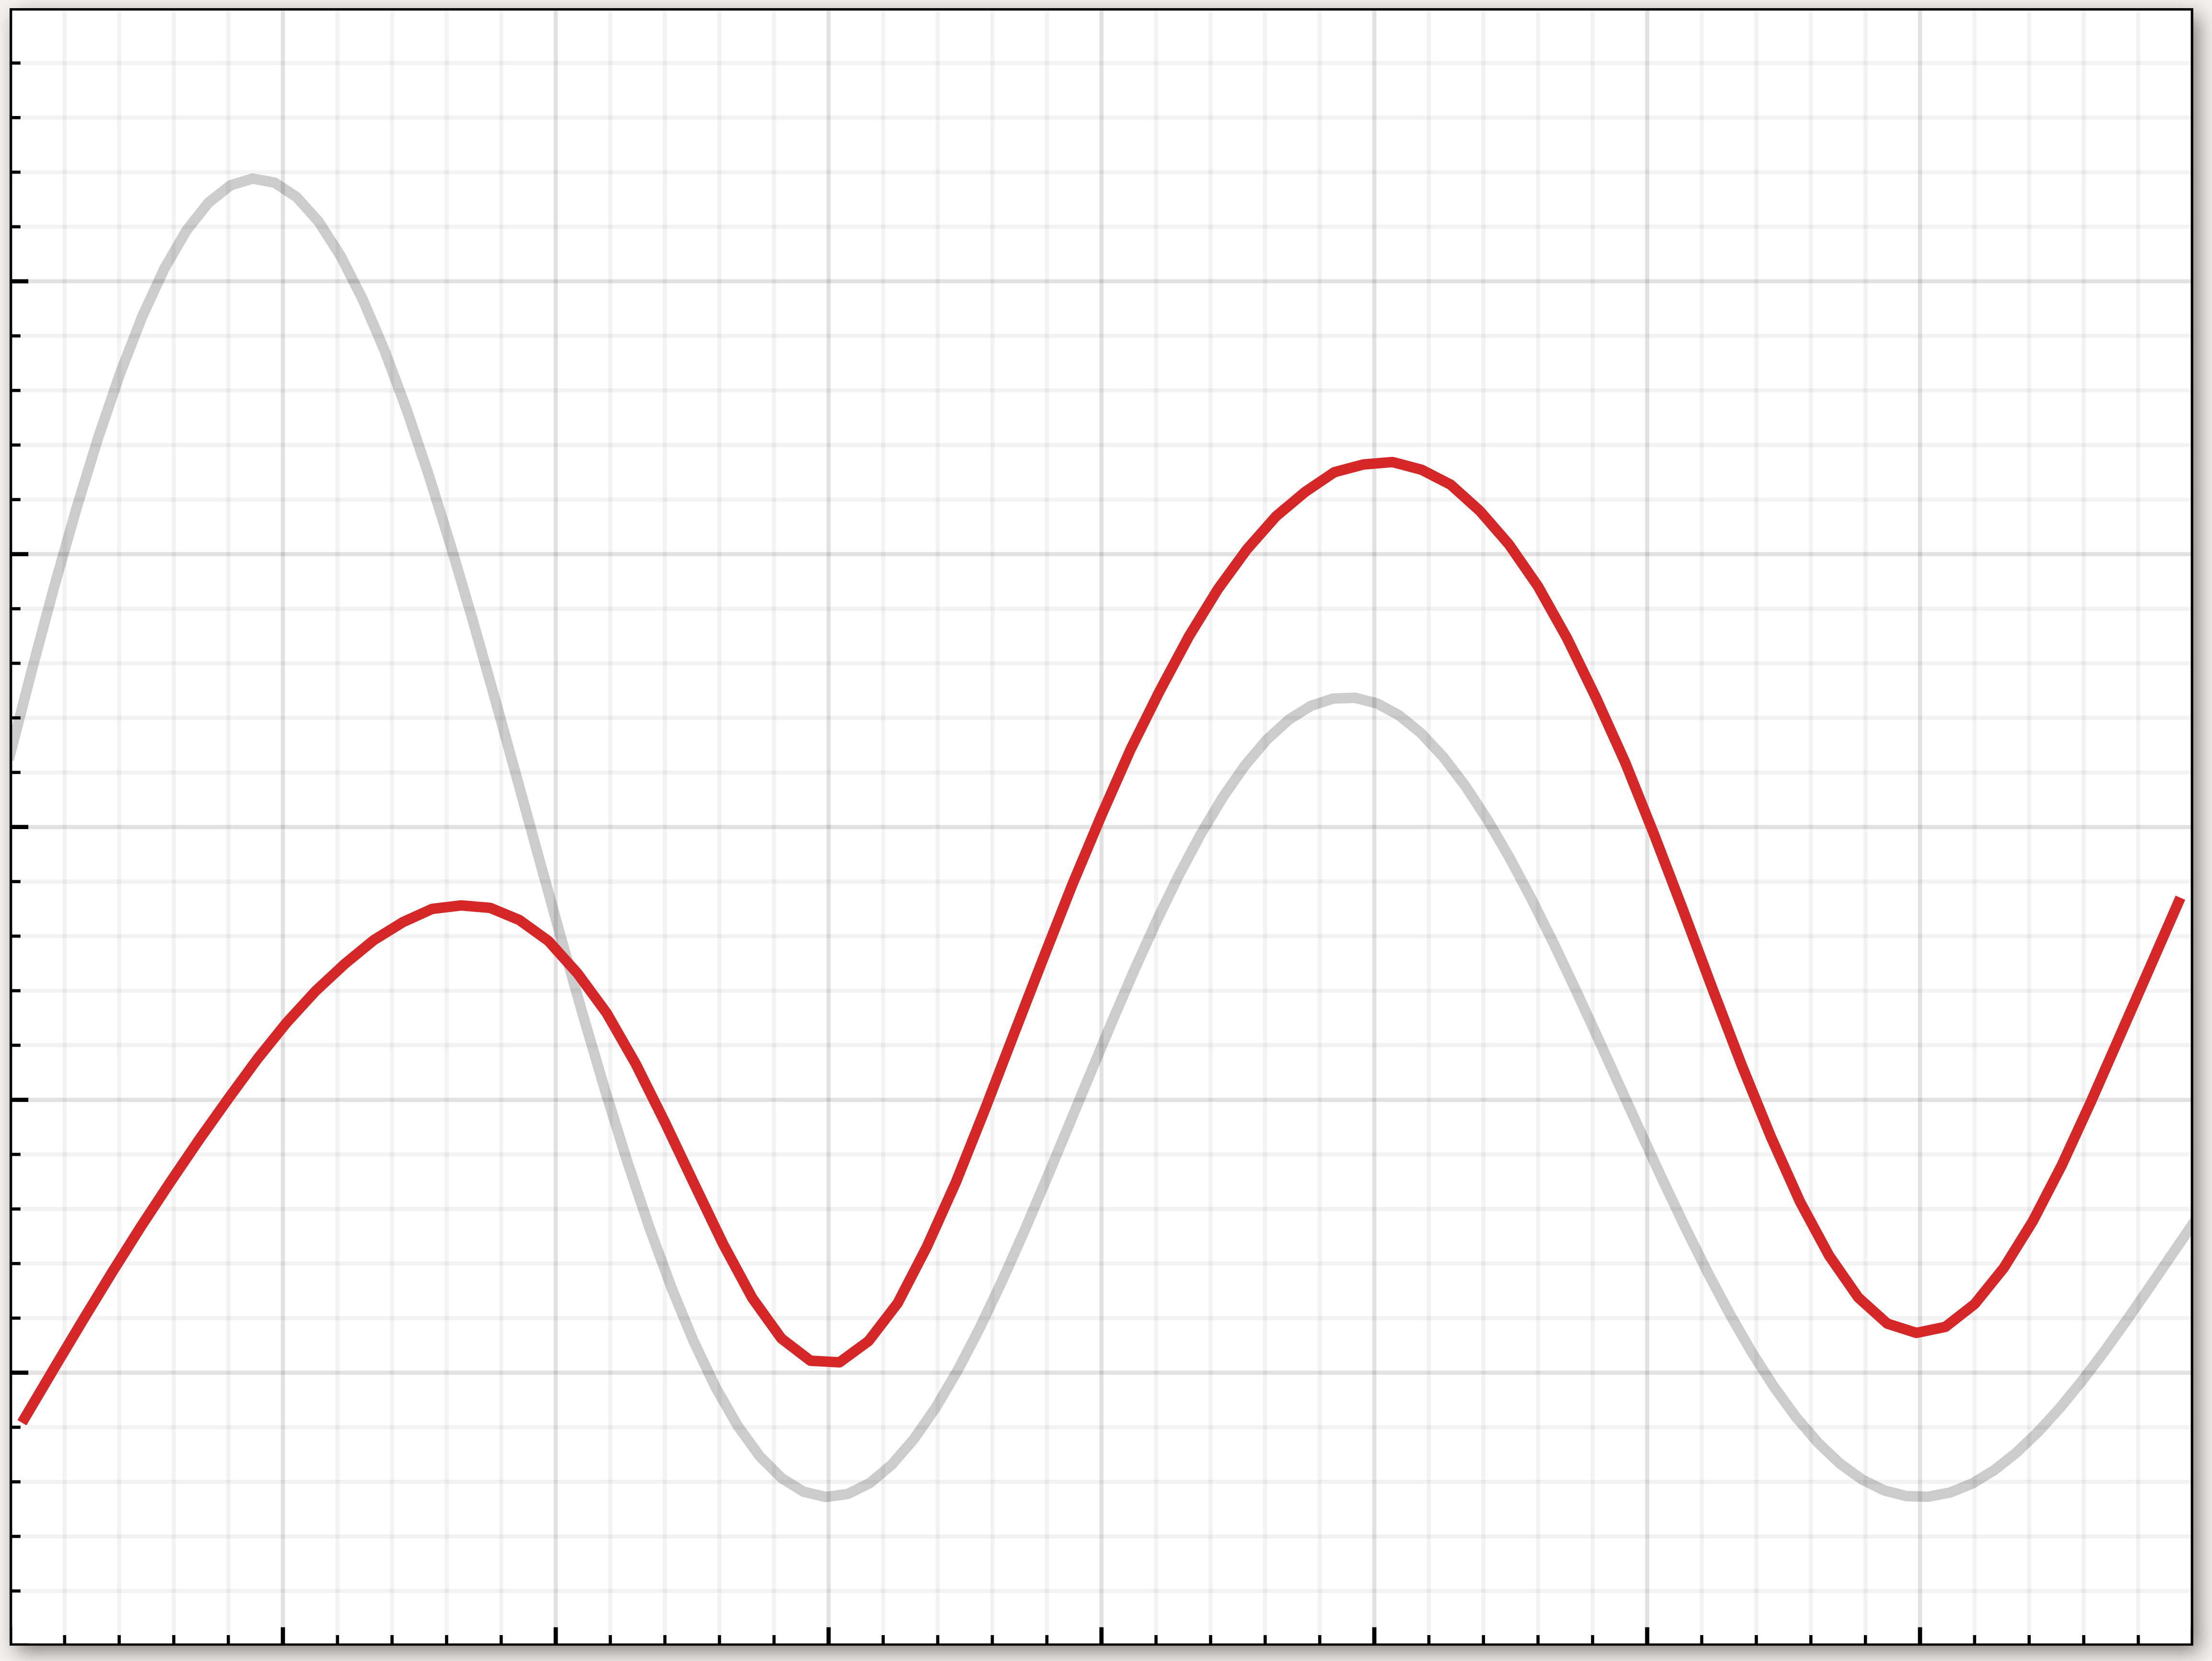

$p_2(x)$

Reservoir Reuse

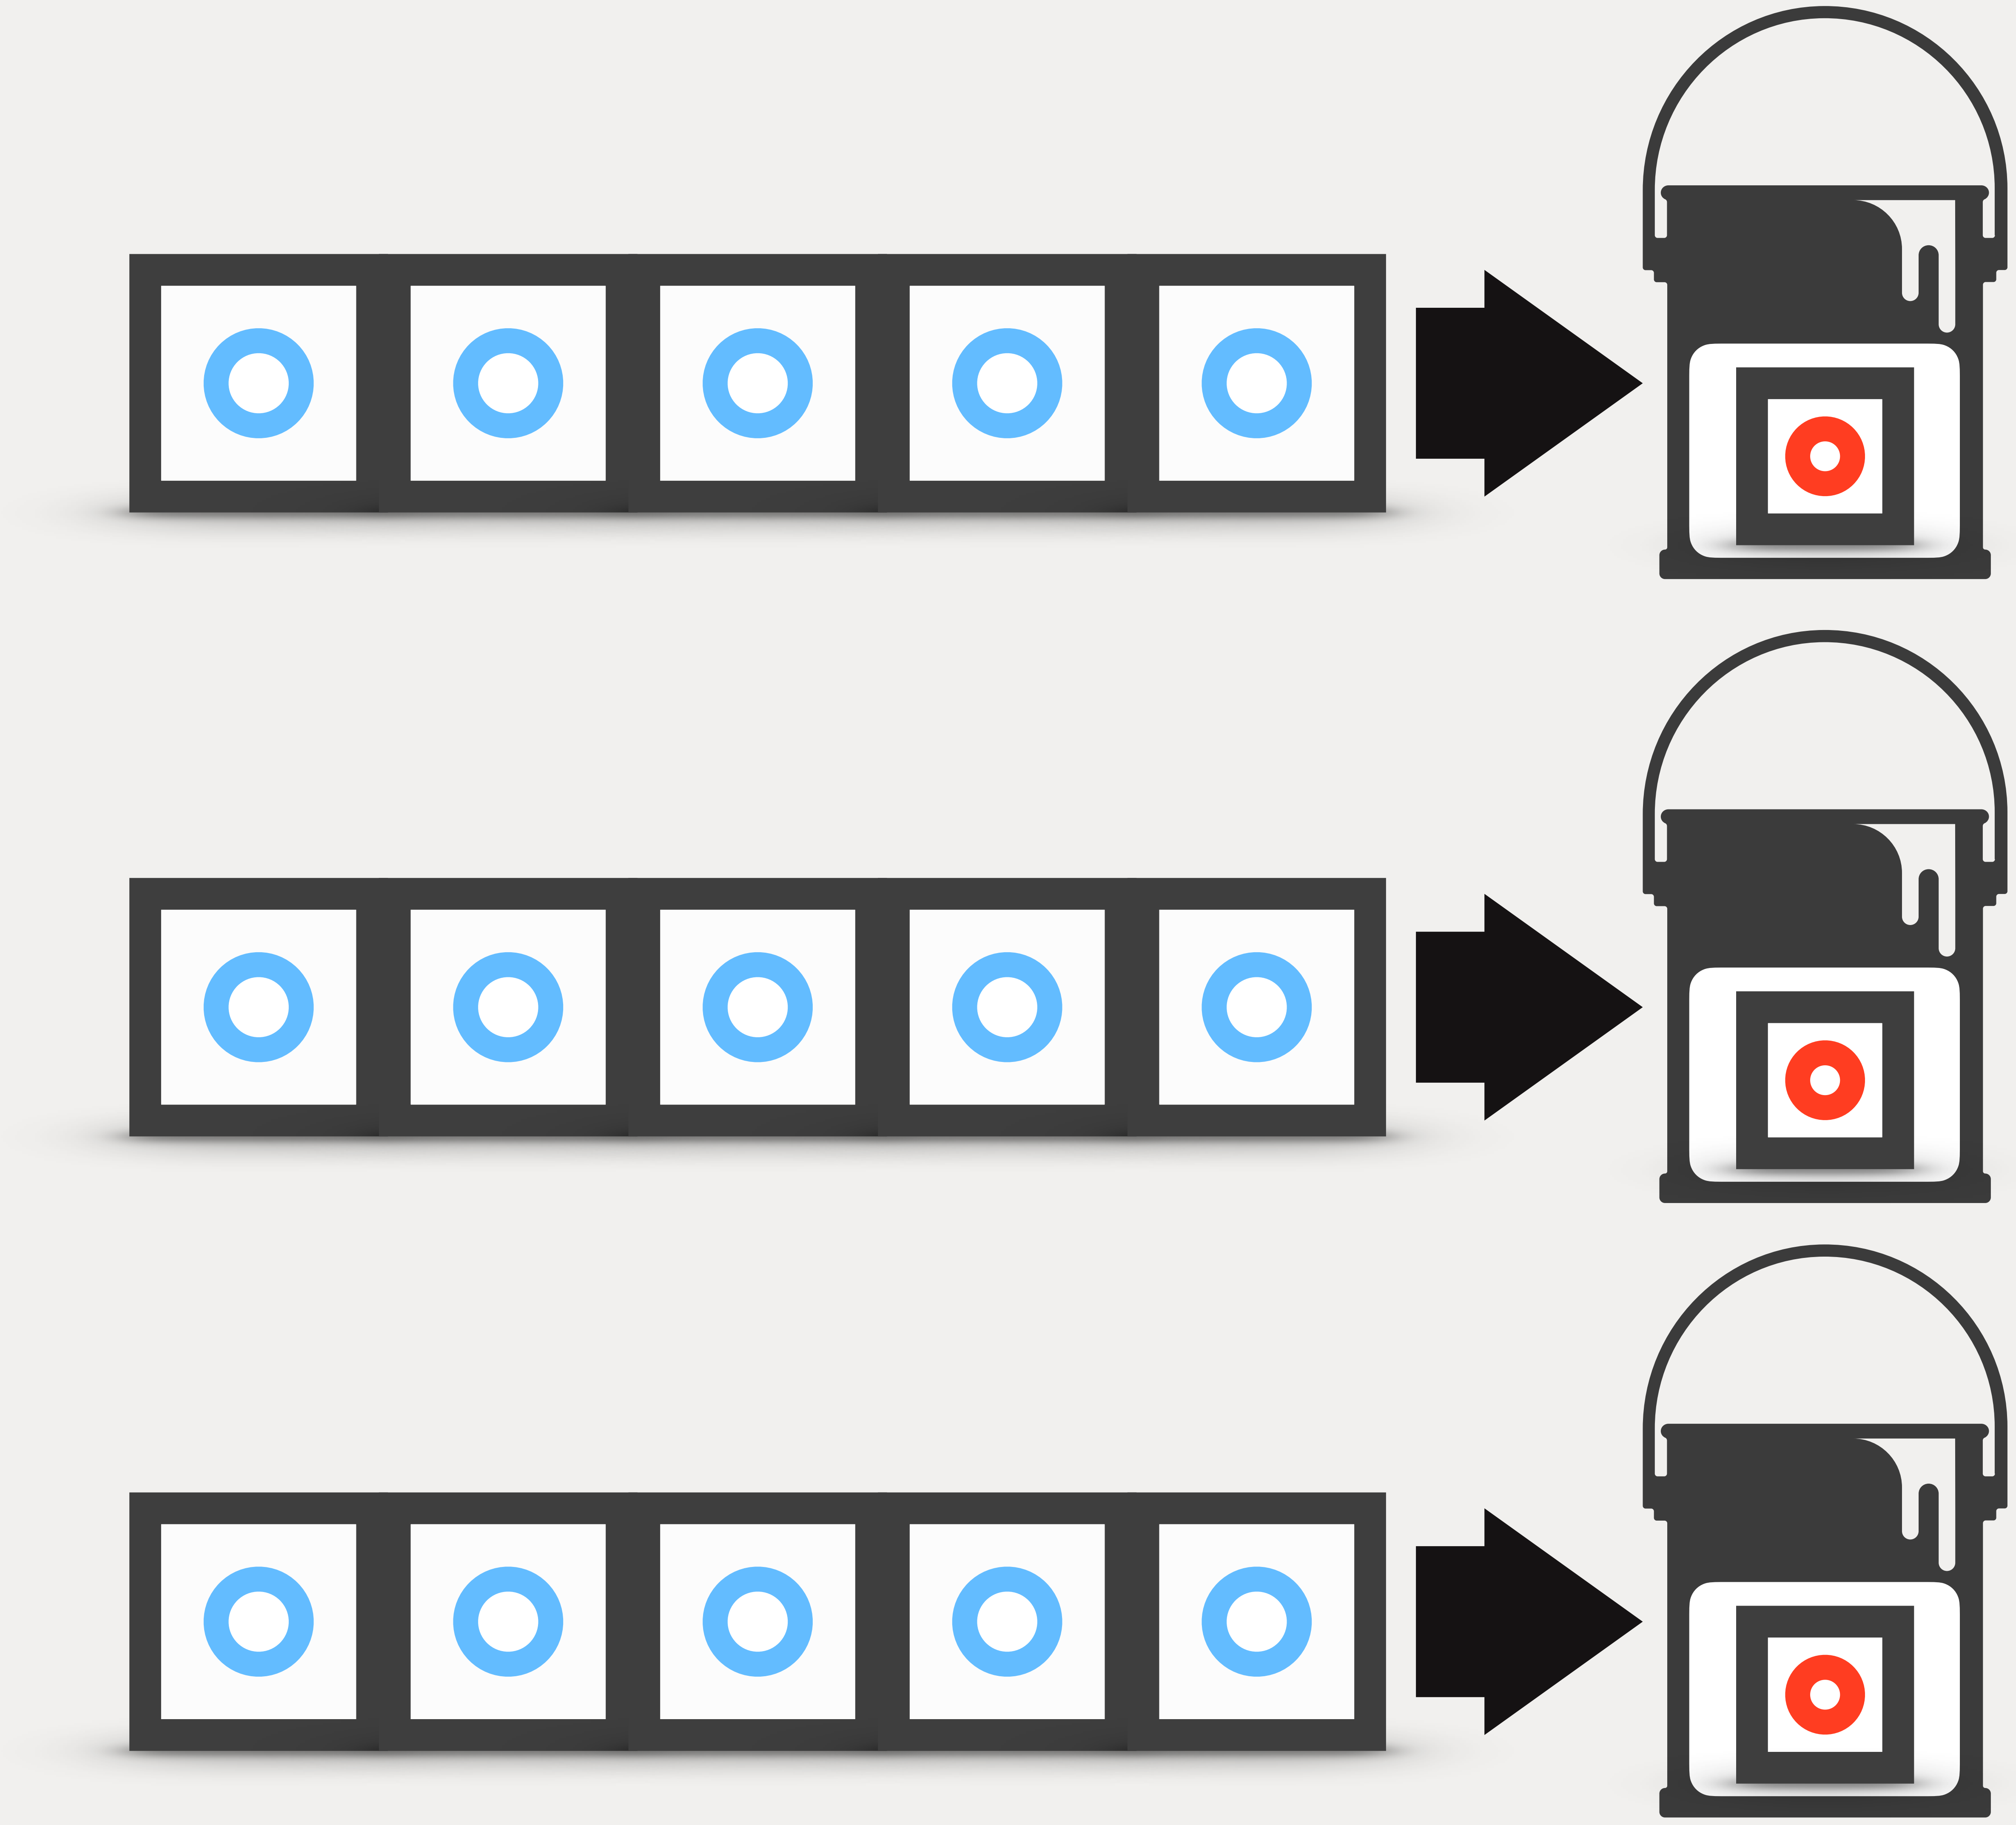

Reservoir Reuse

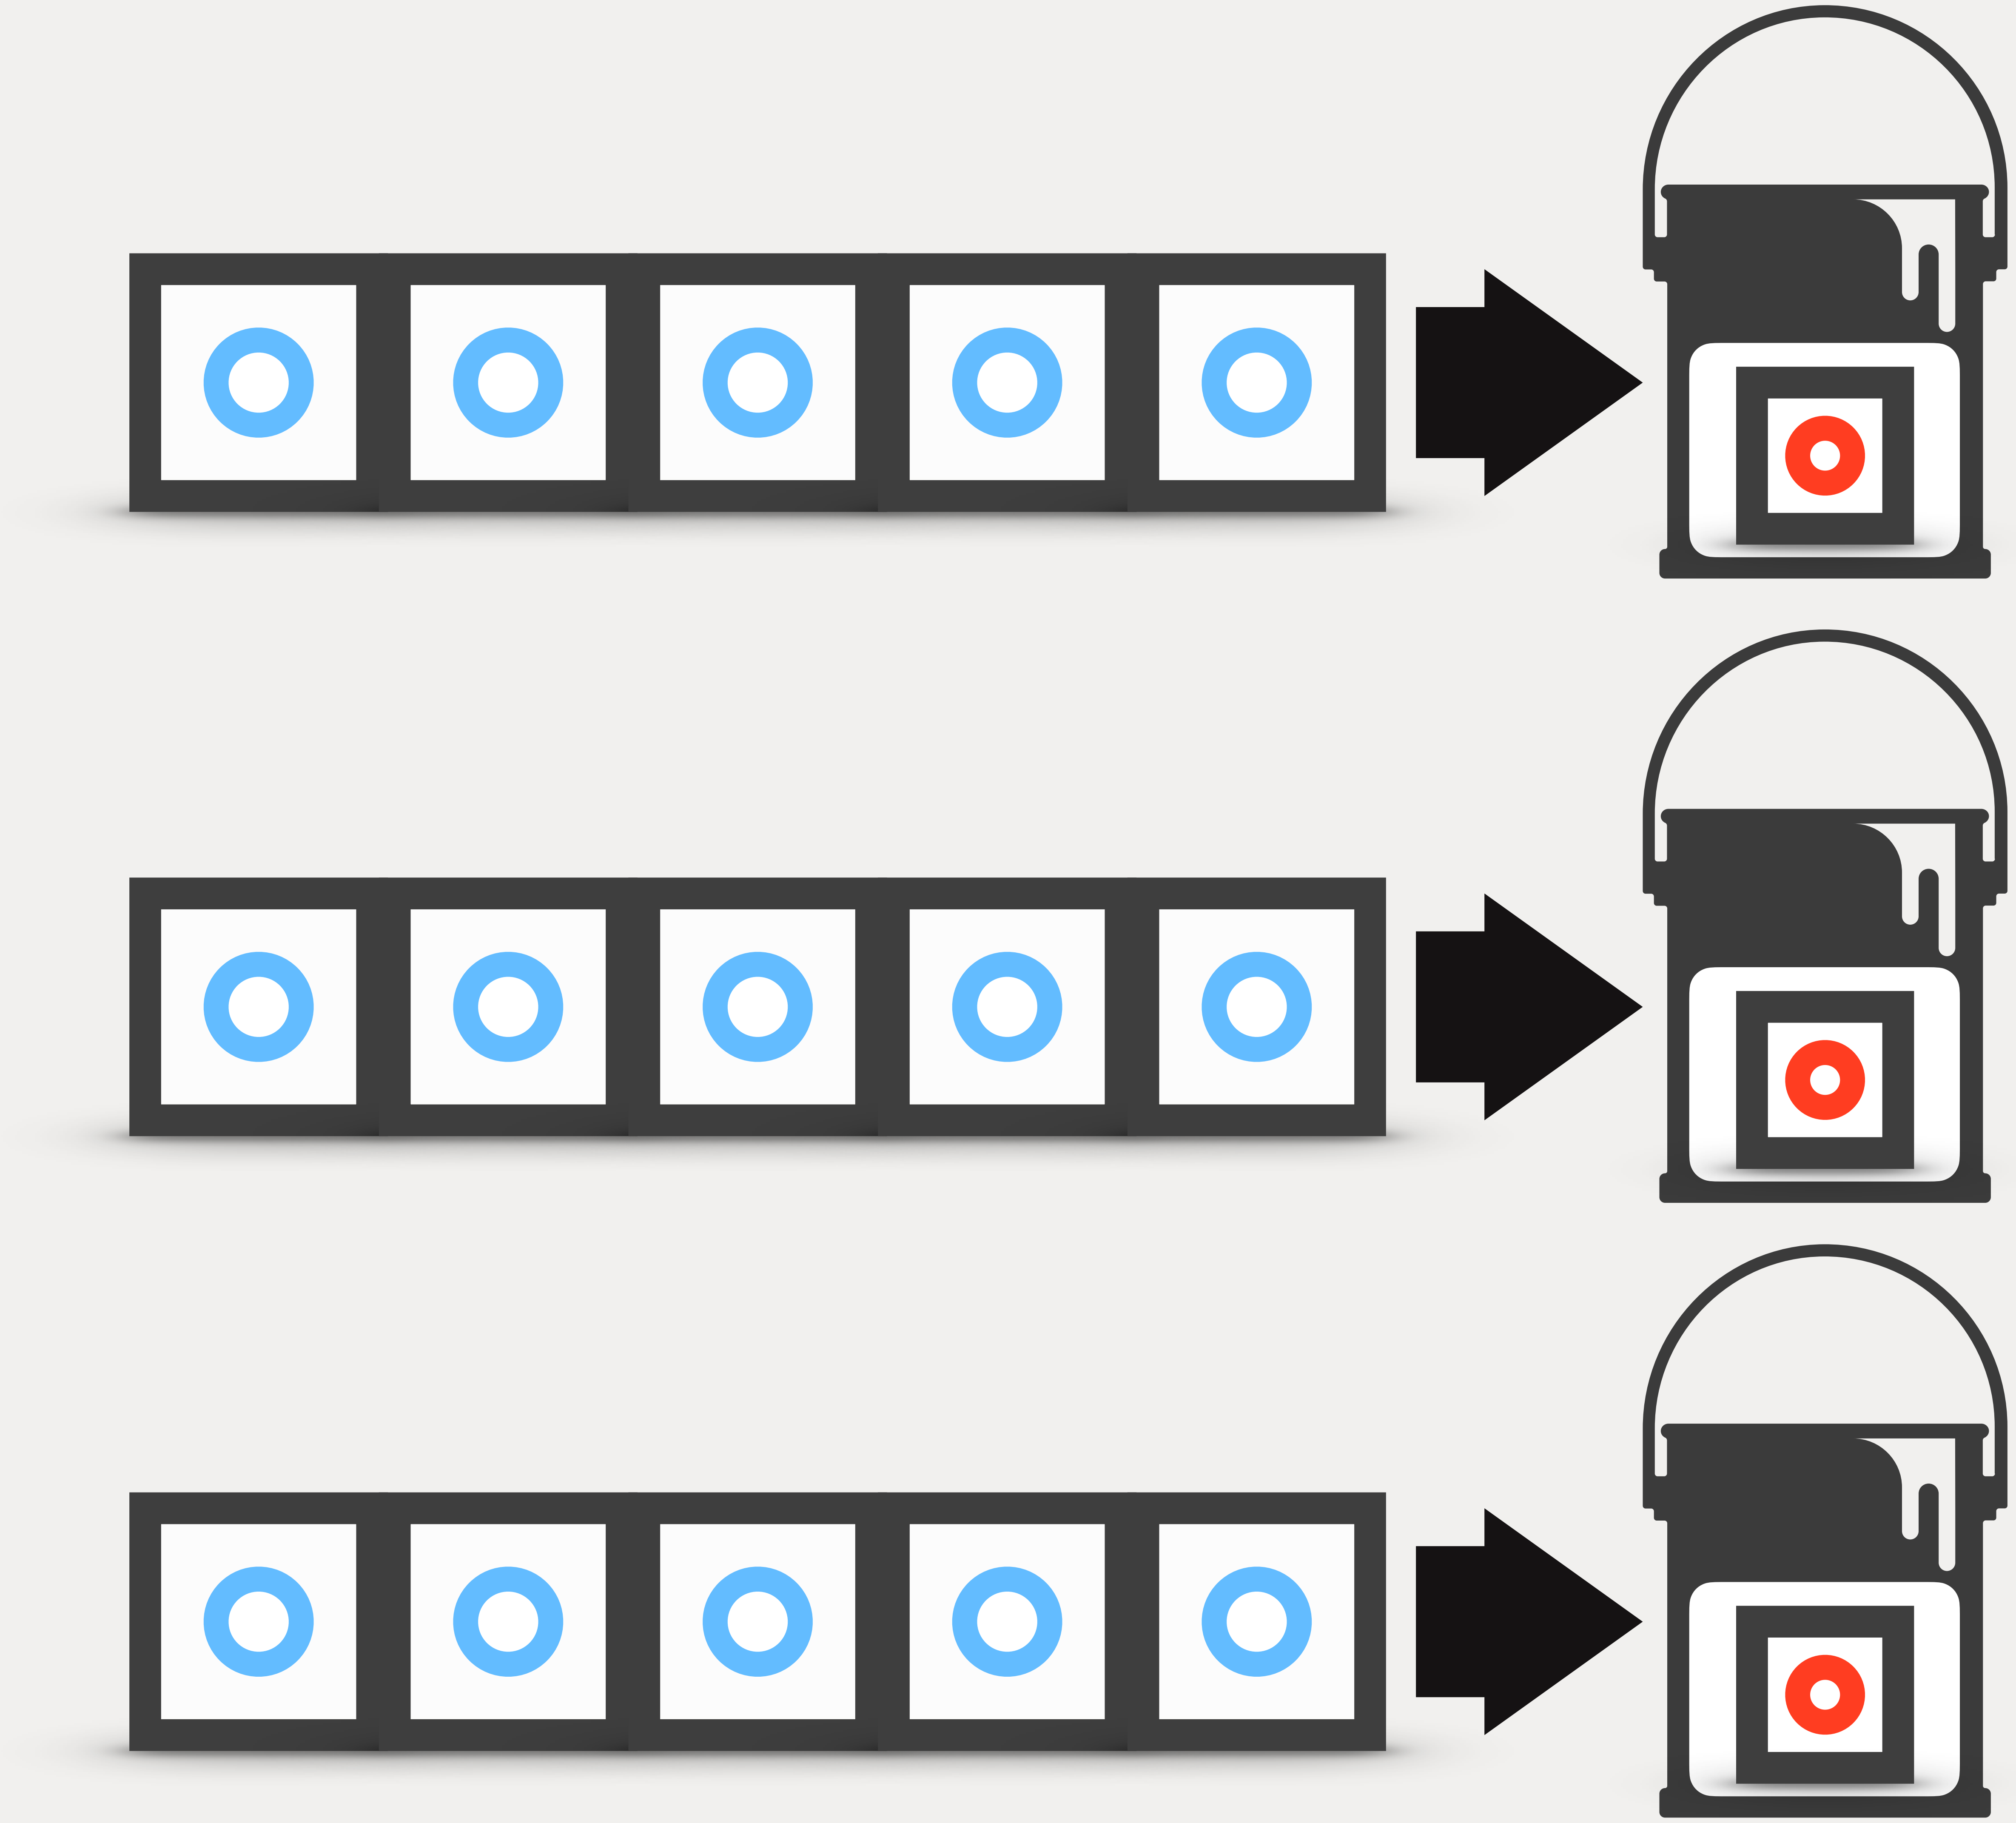

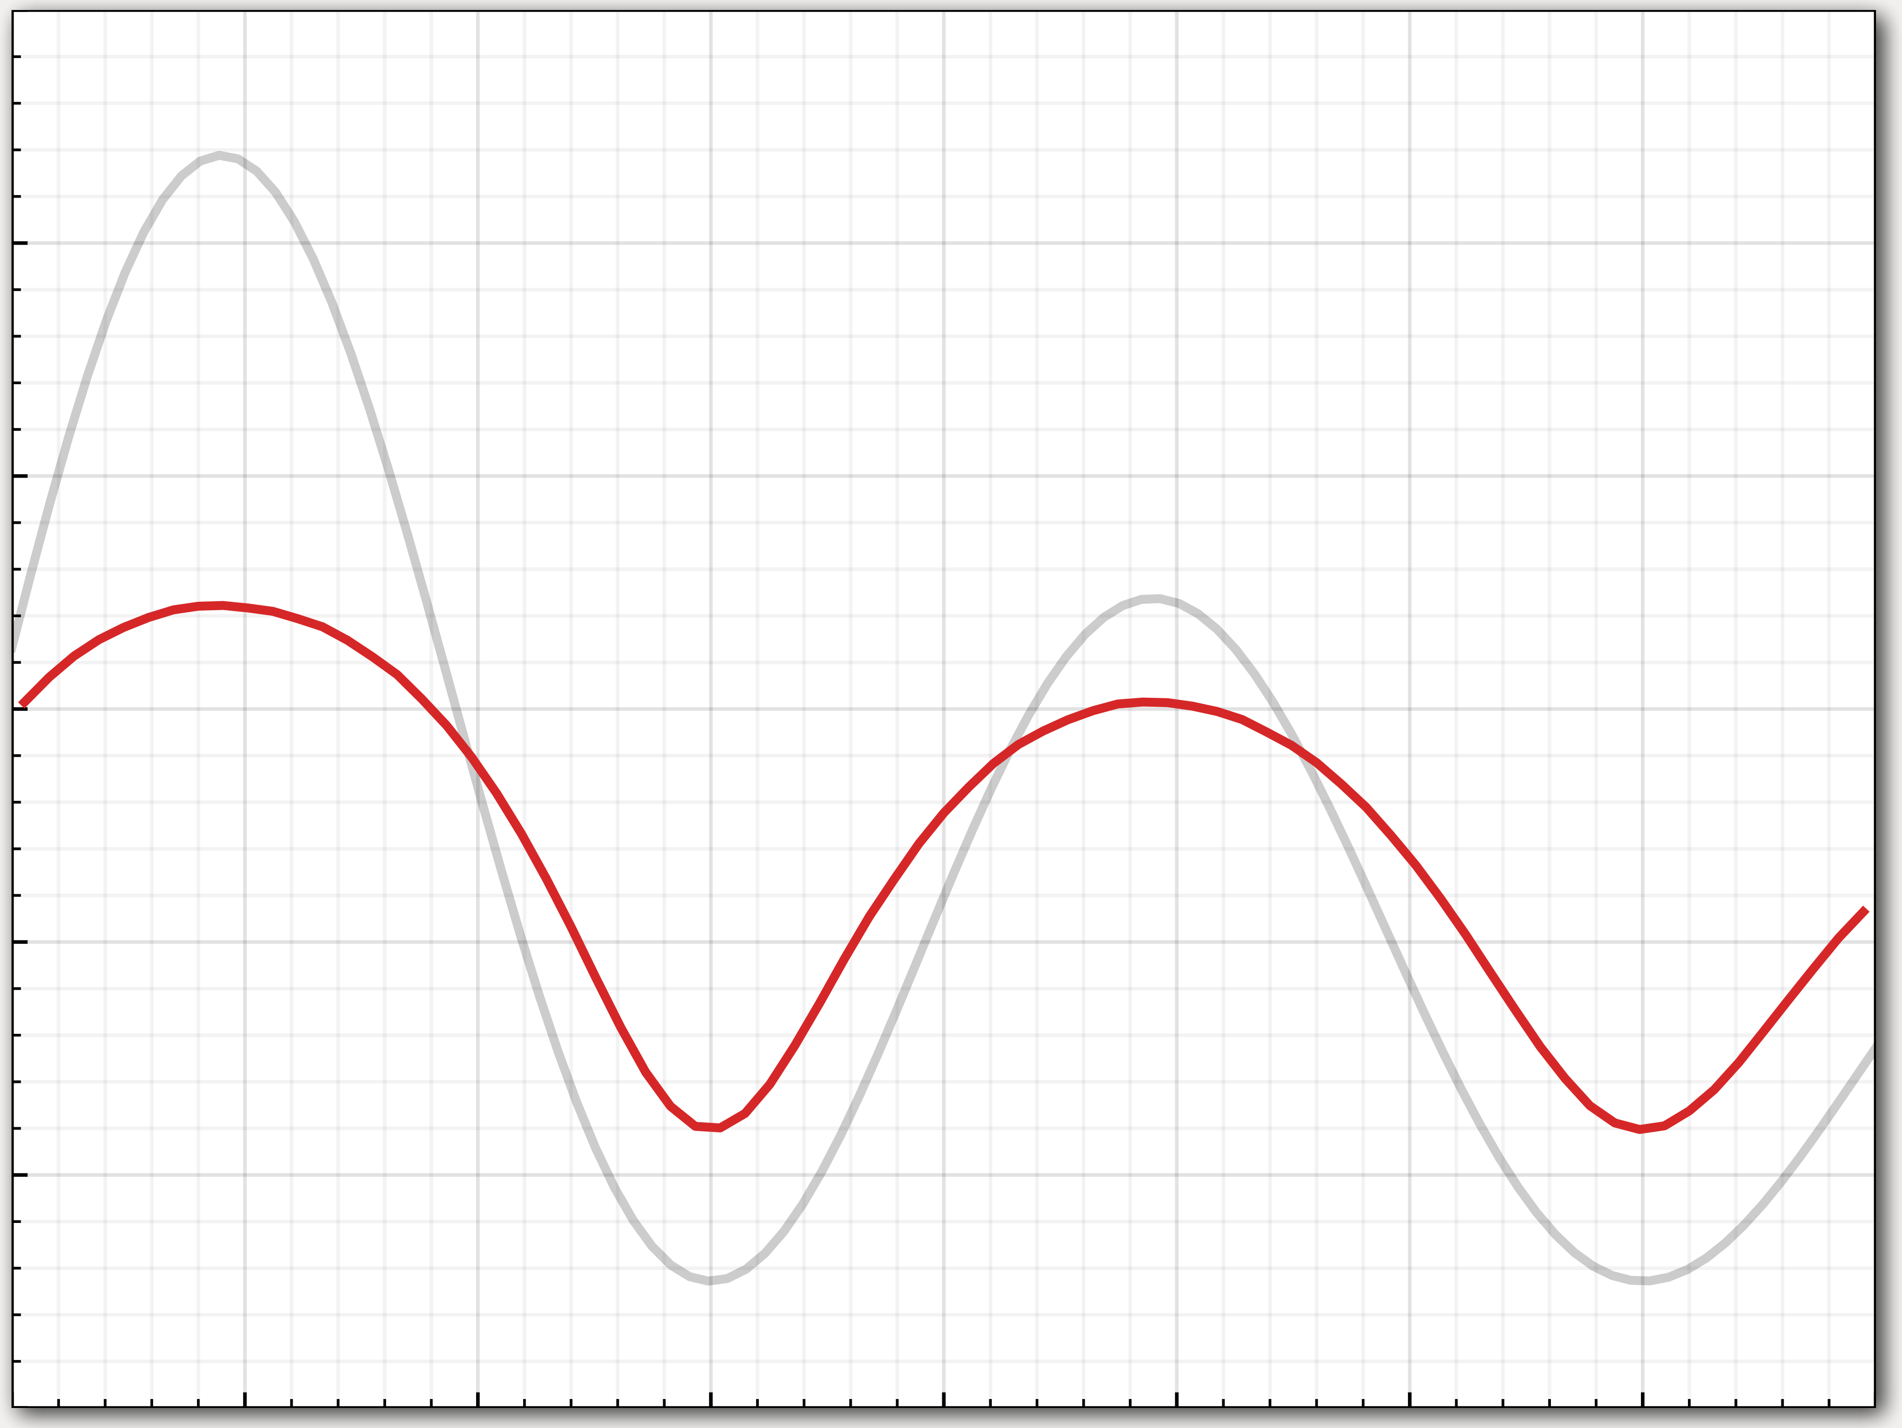

Reservoir Reuse

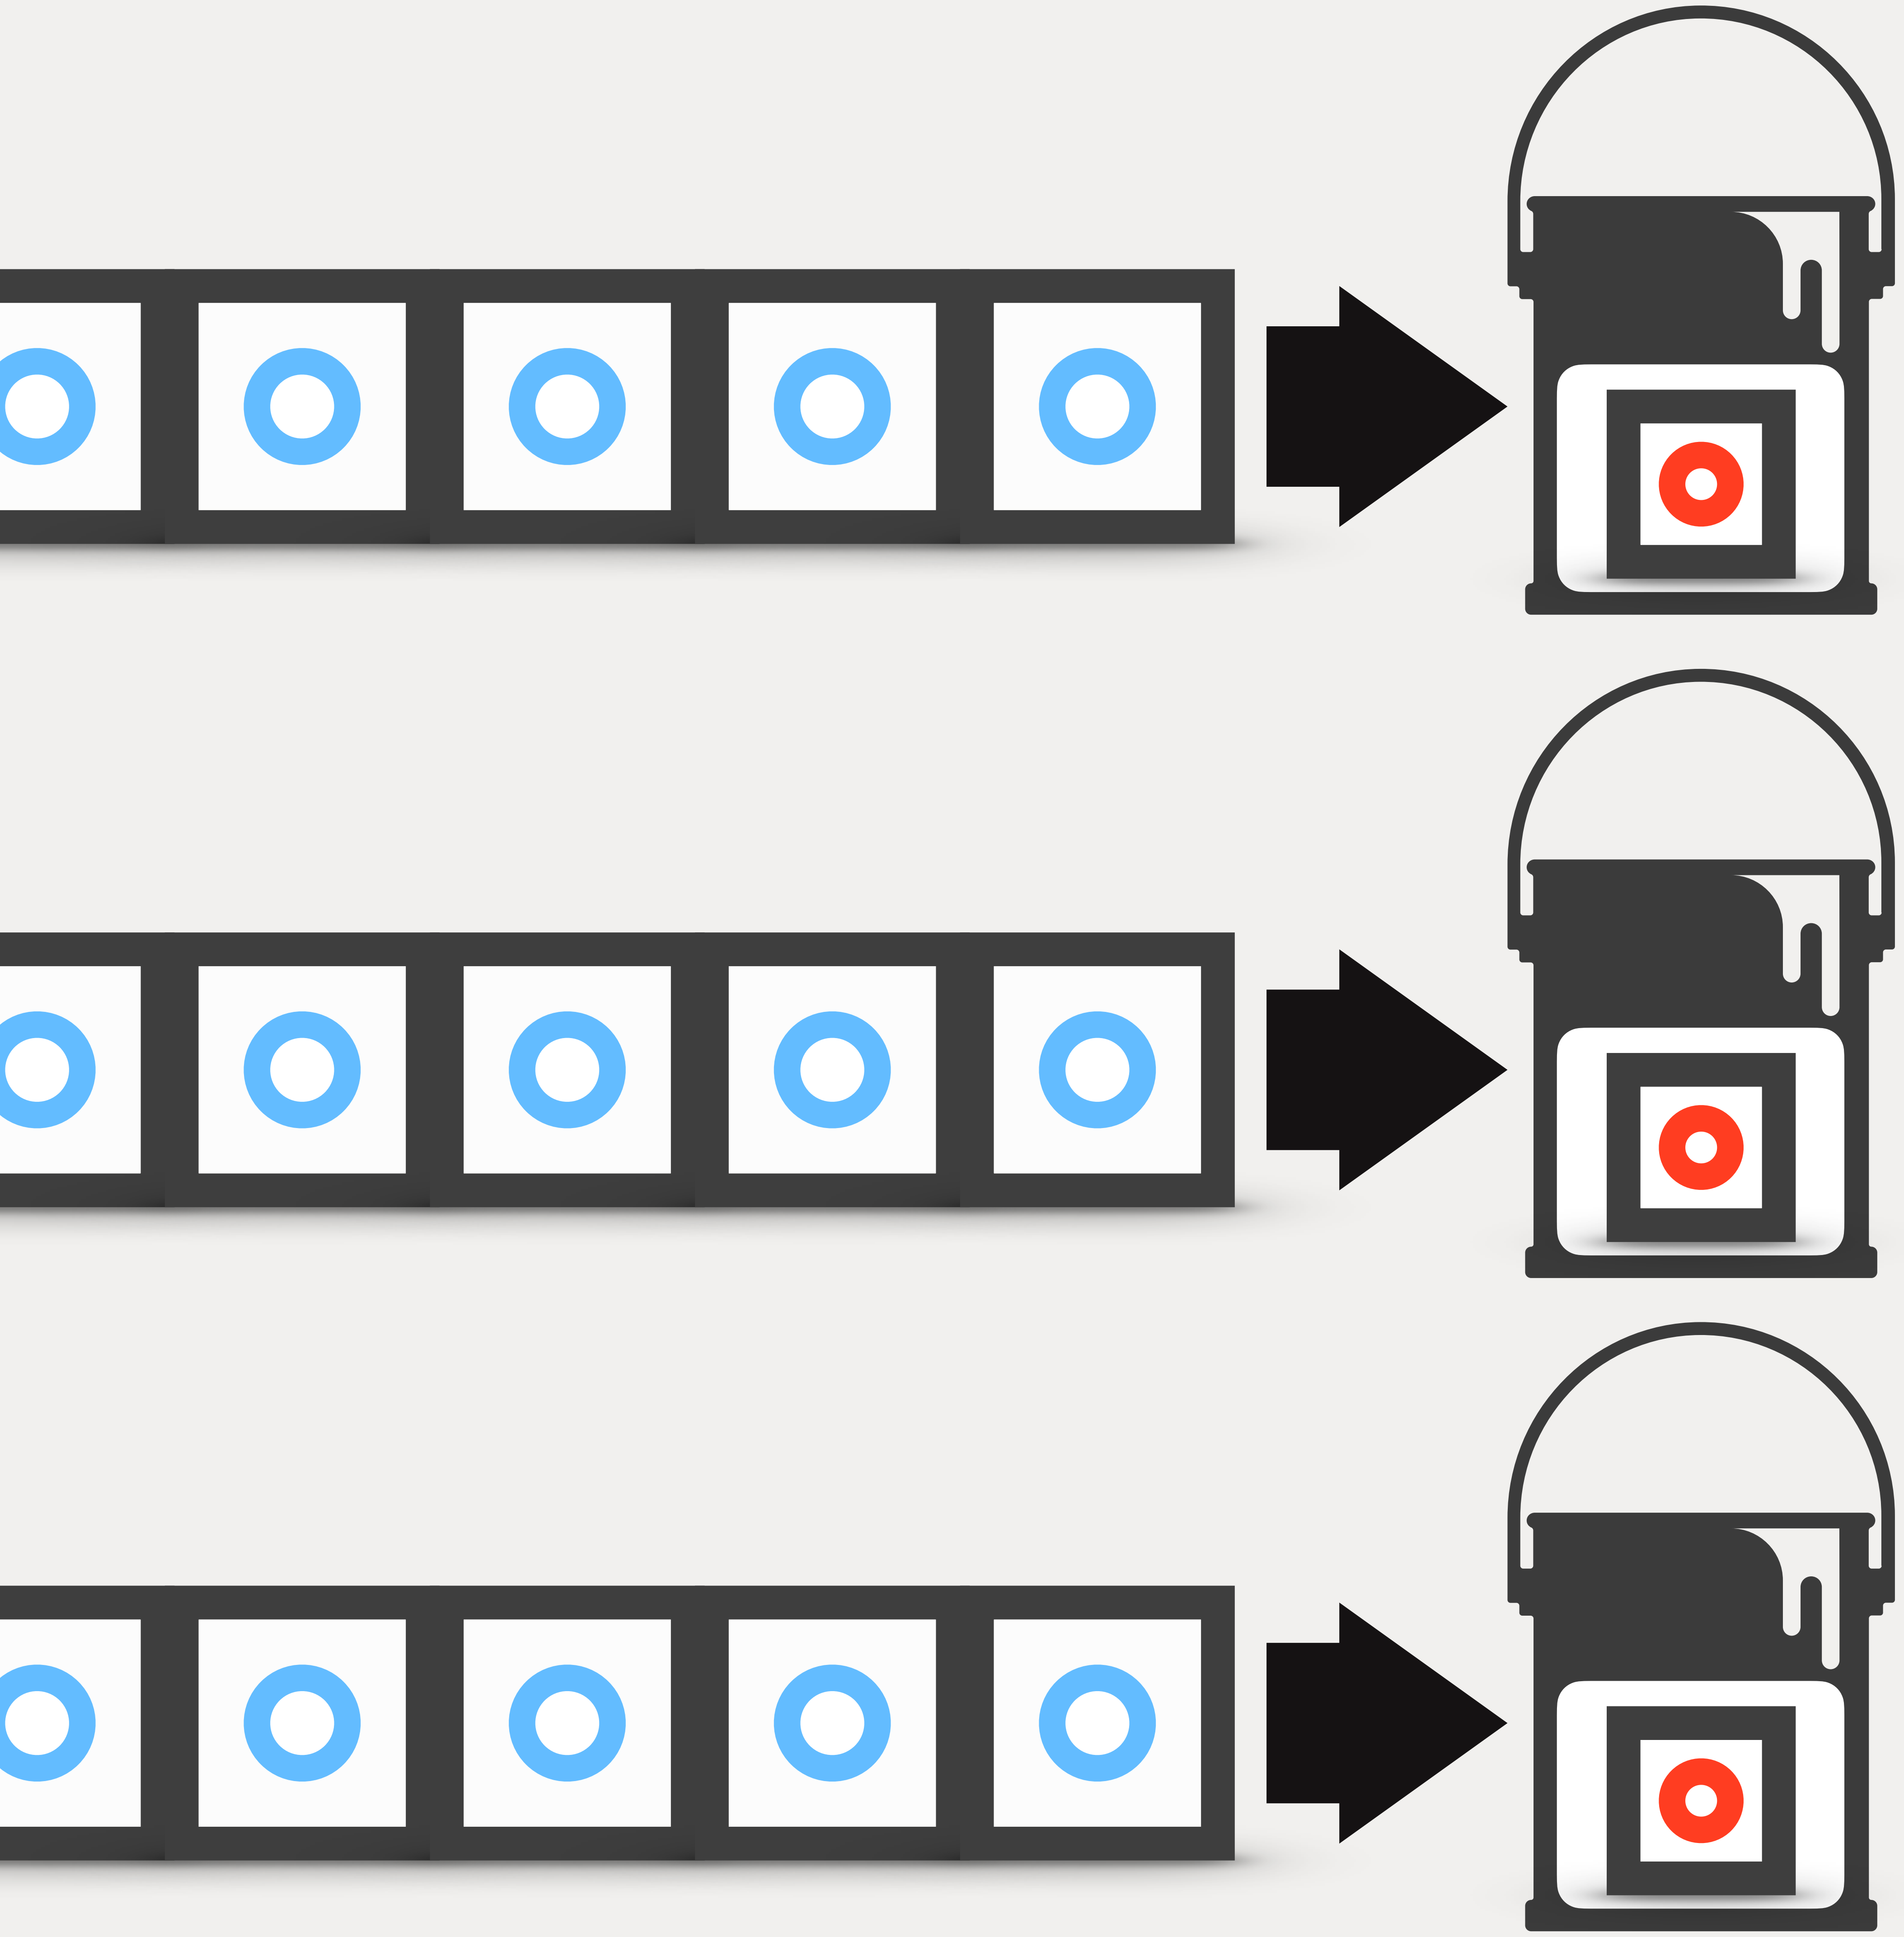

Reservoir Reuse

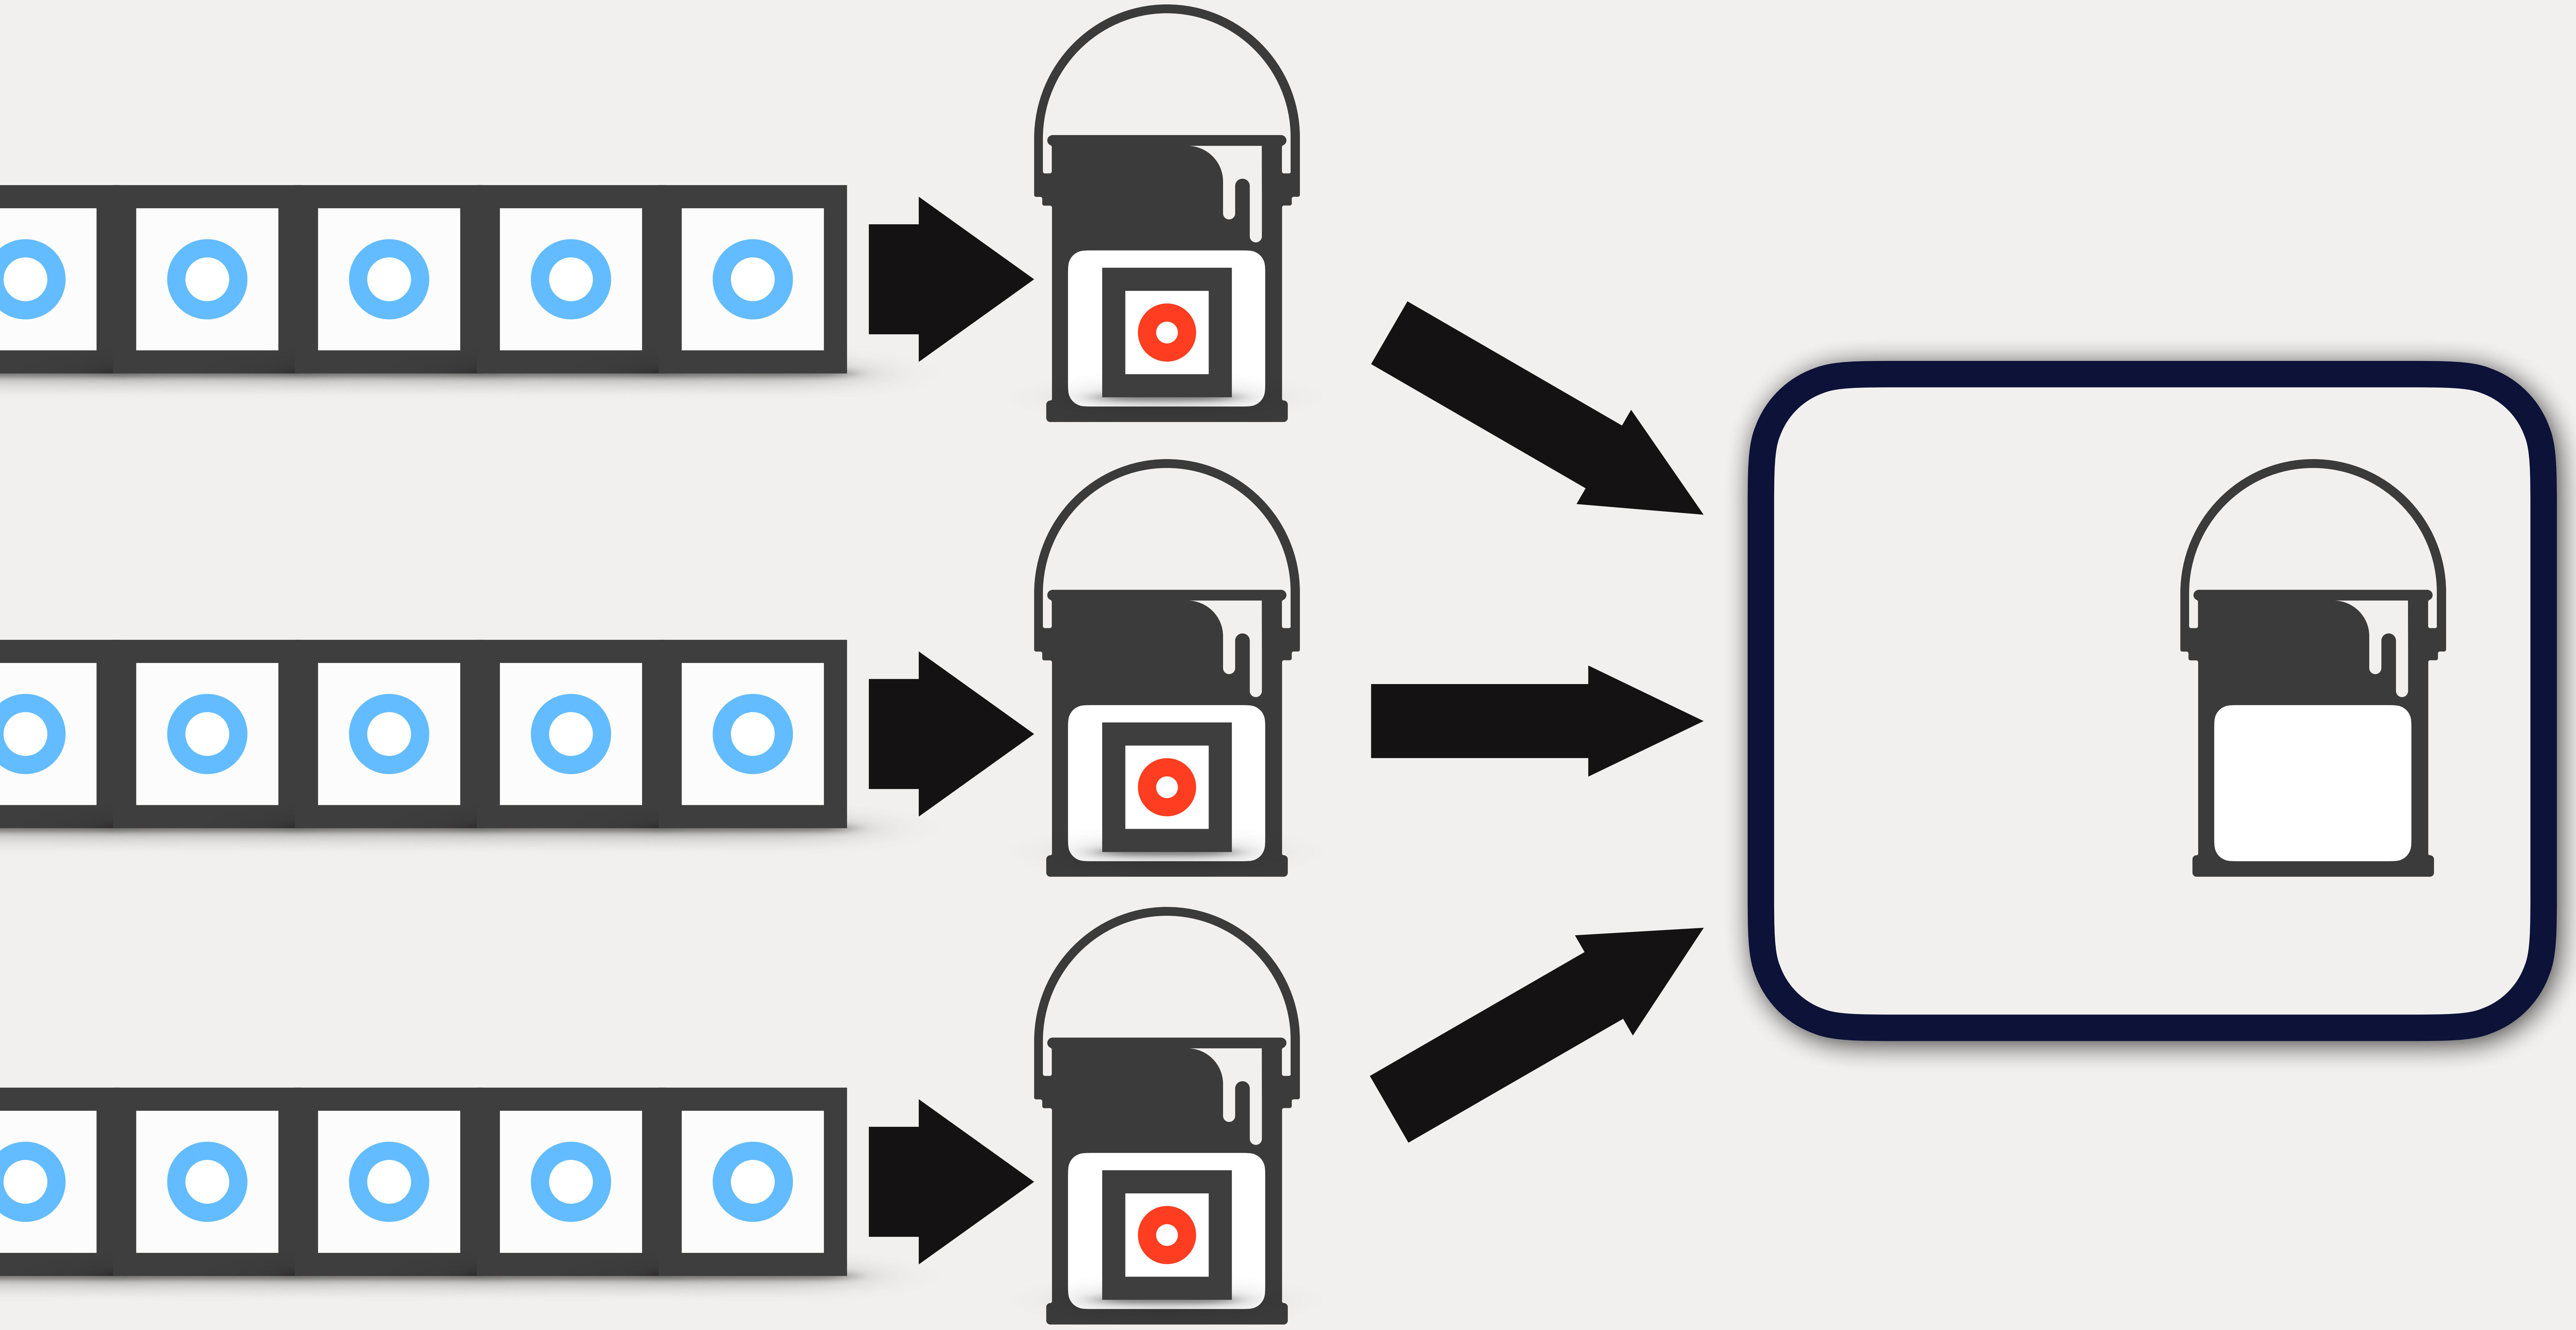

Reservoir Reuse

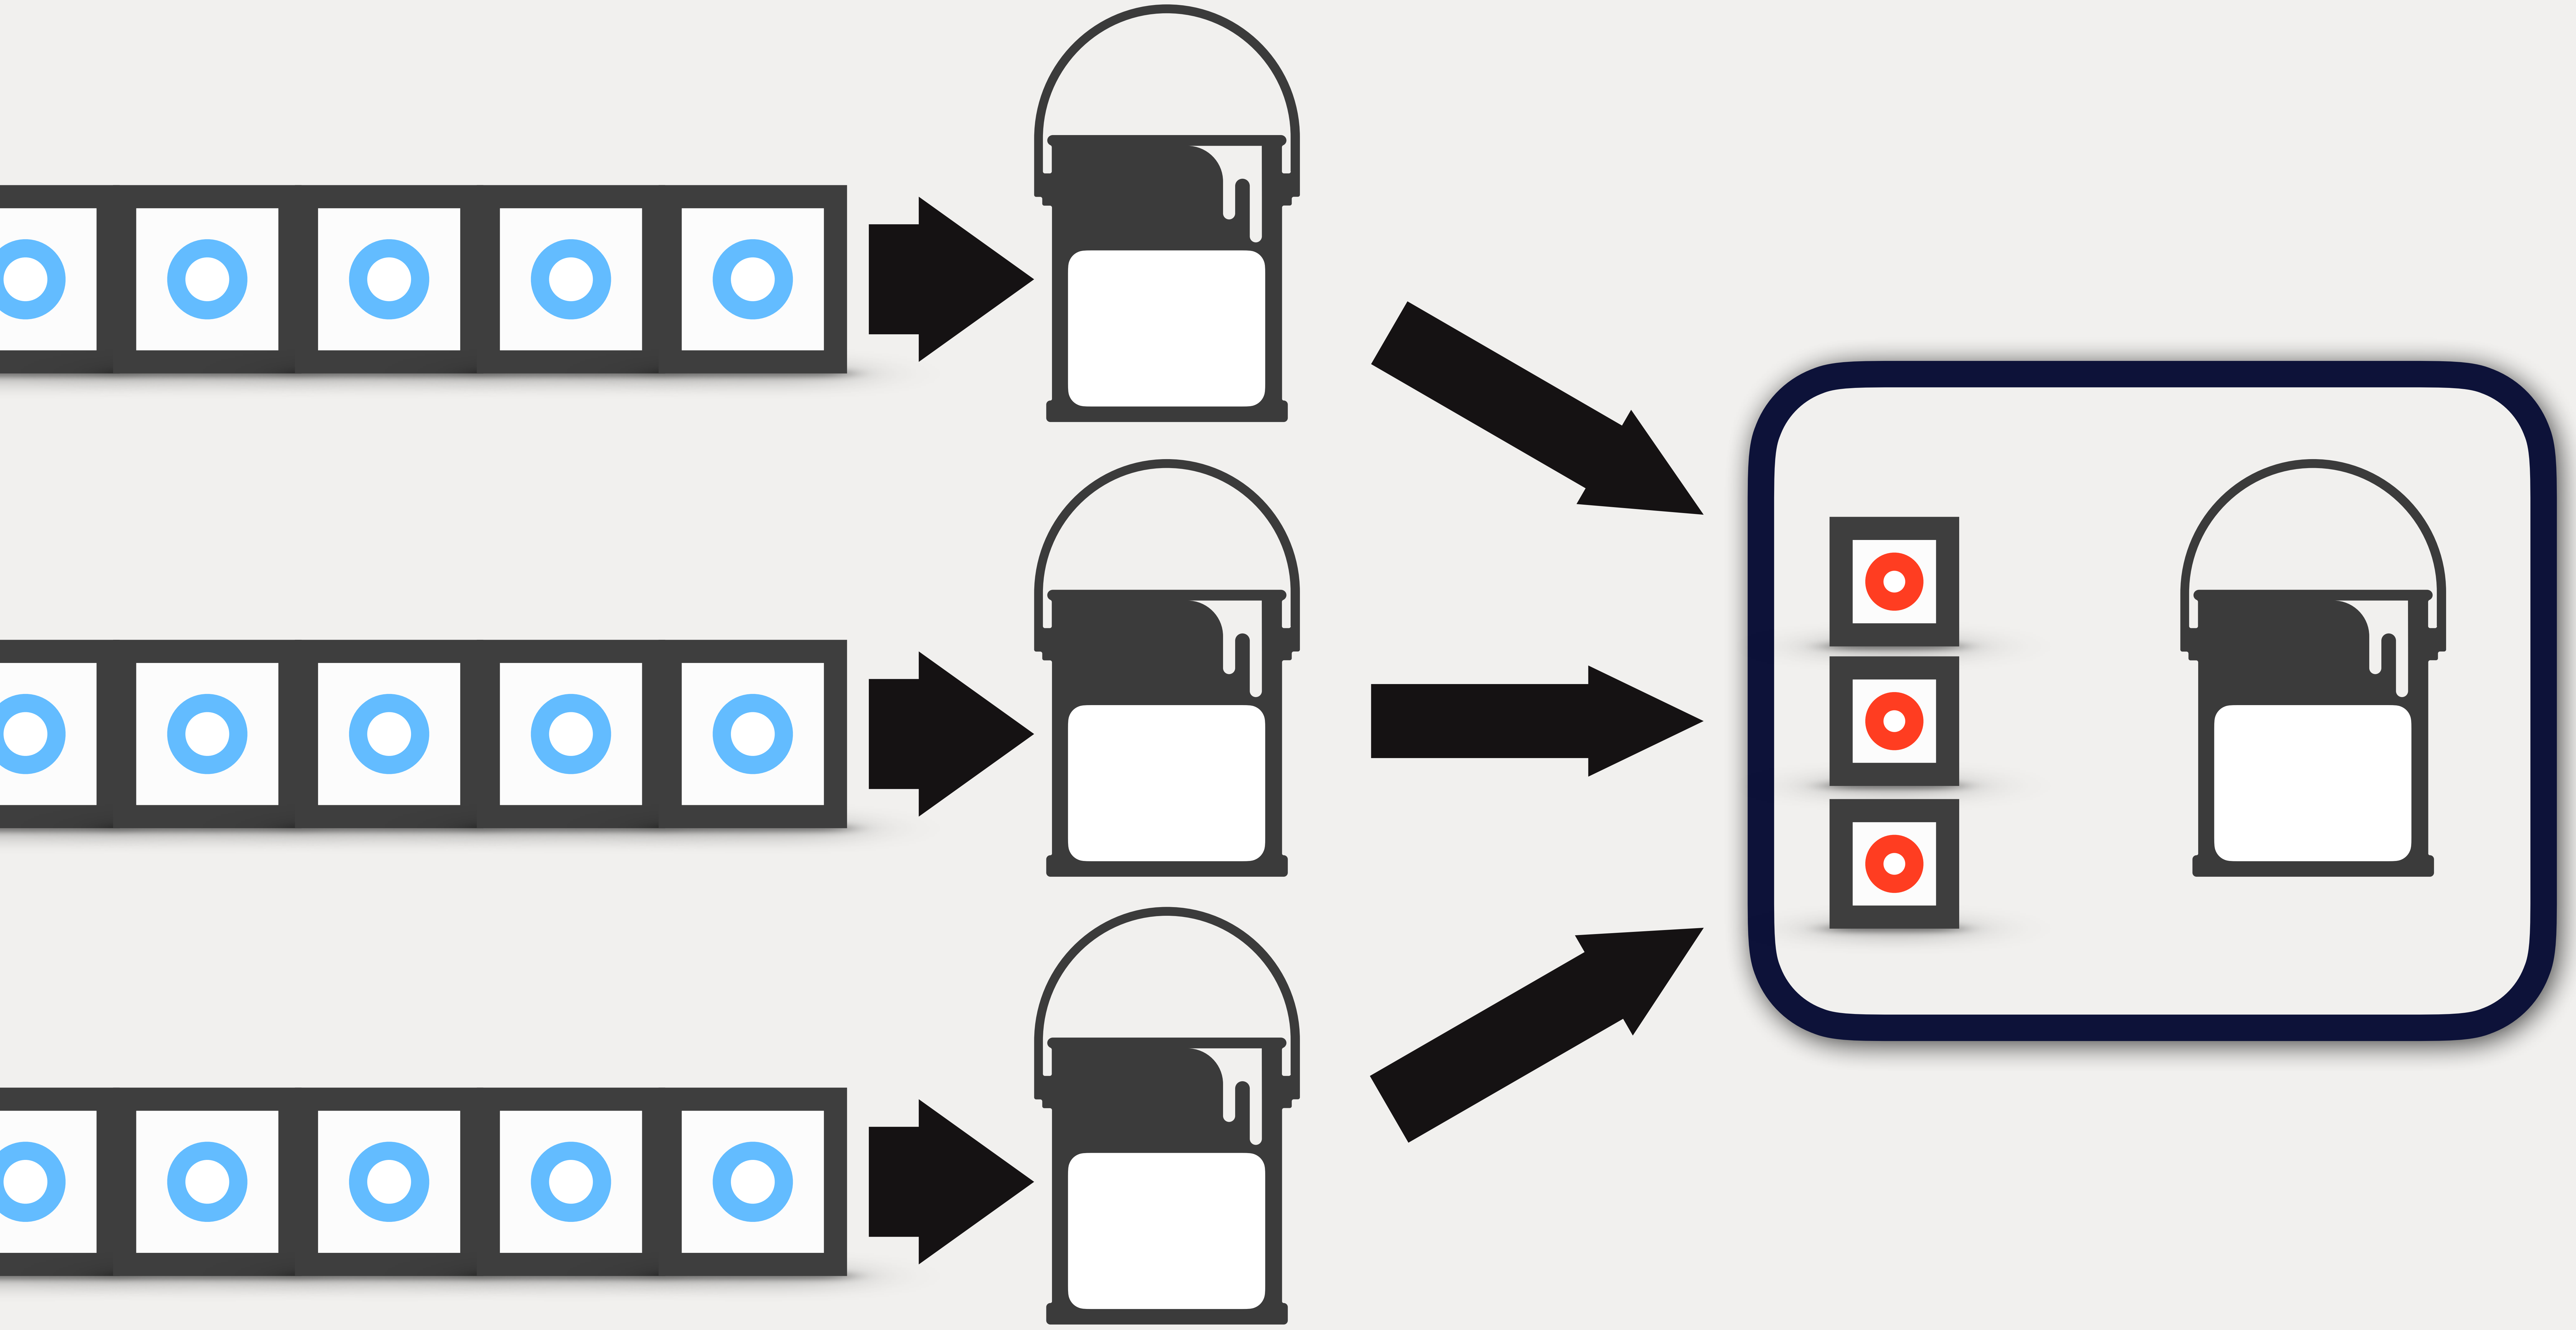

Reservoir Reuse

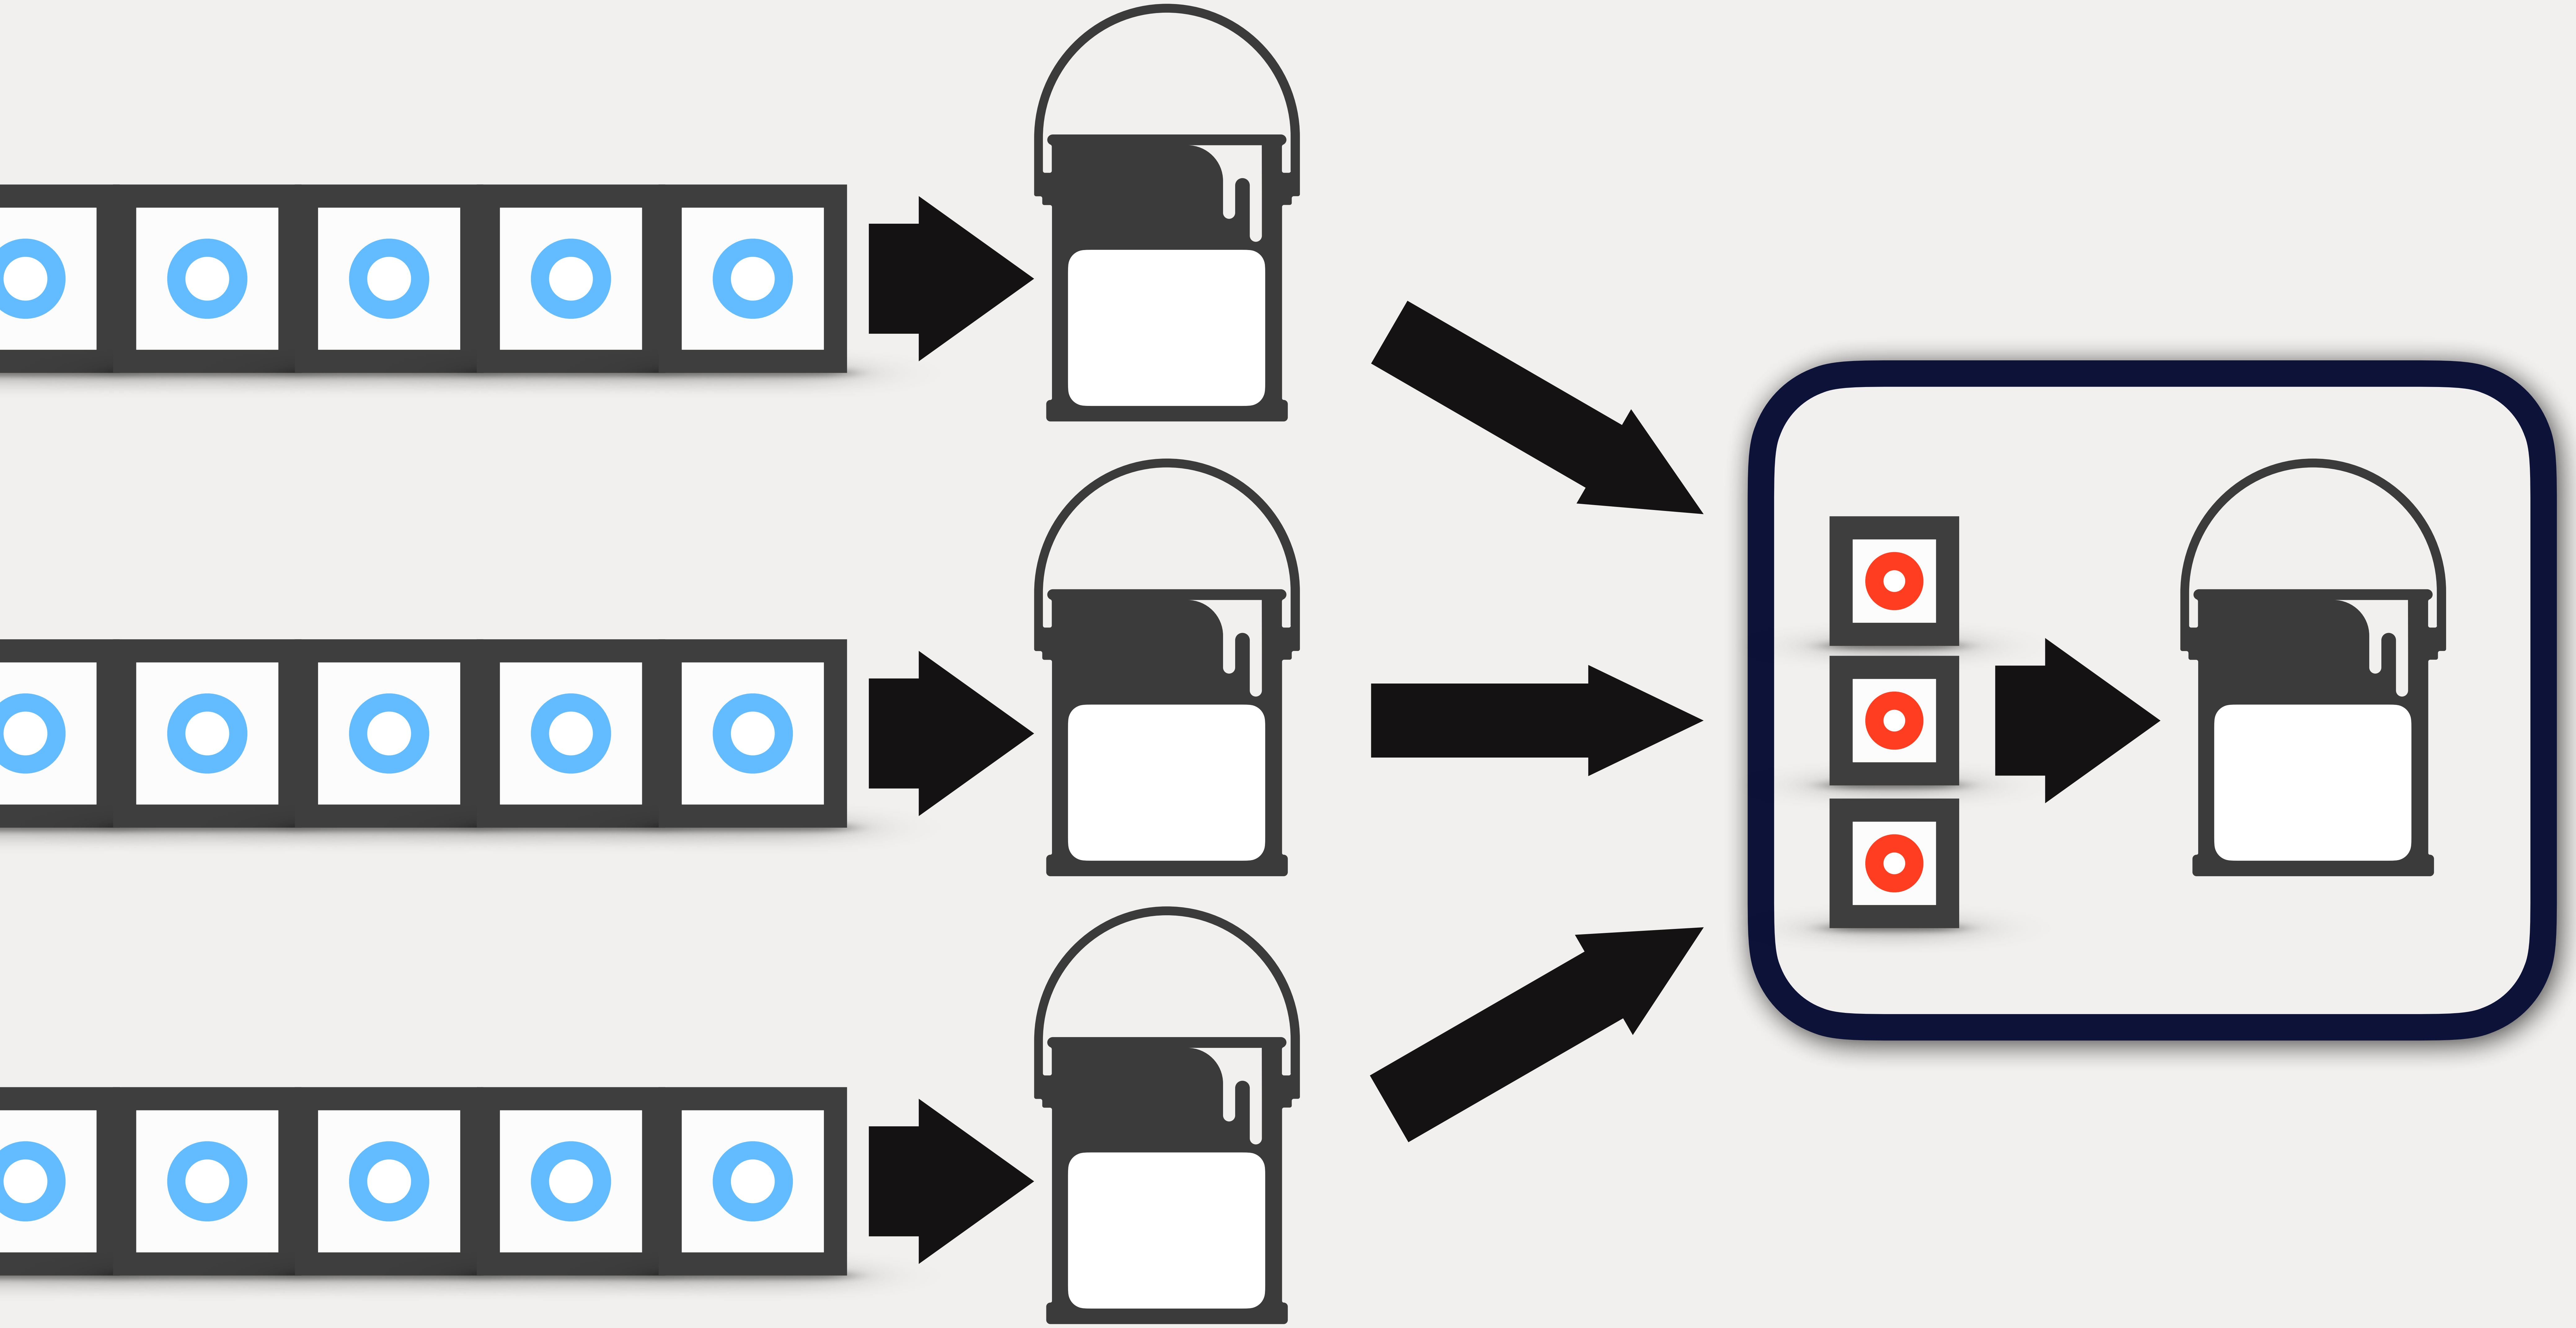

Reservoir Reuse

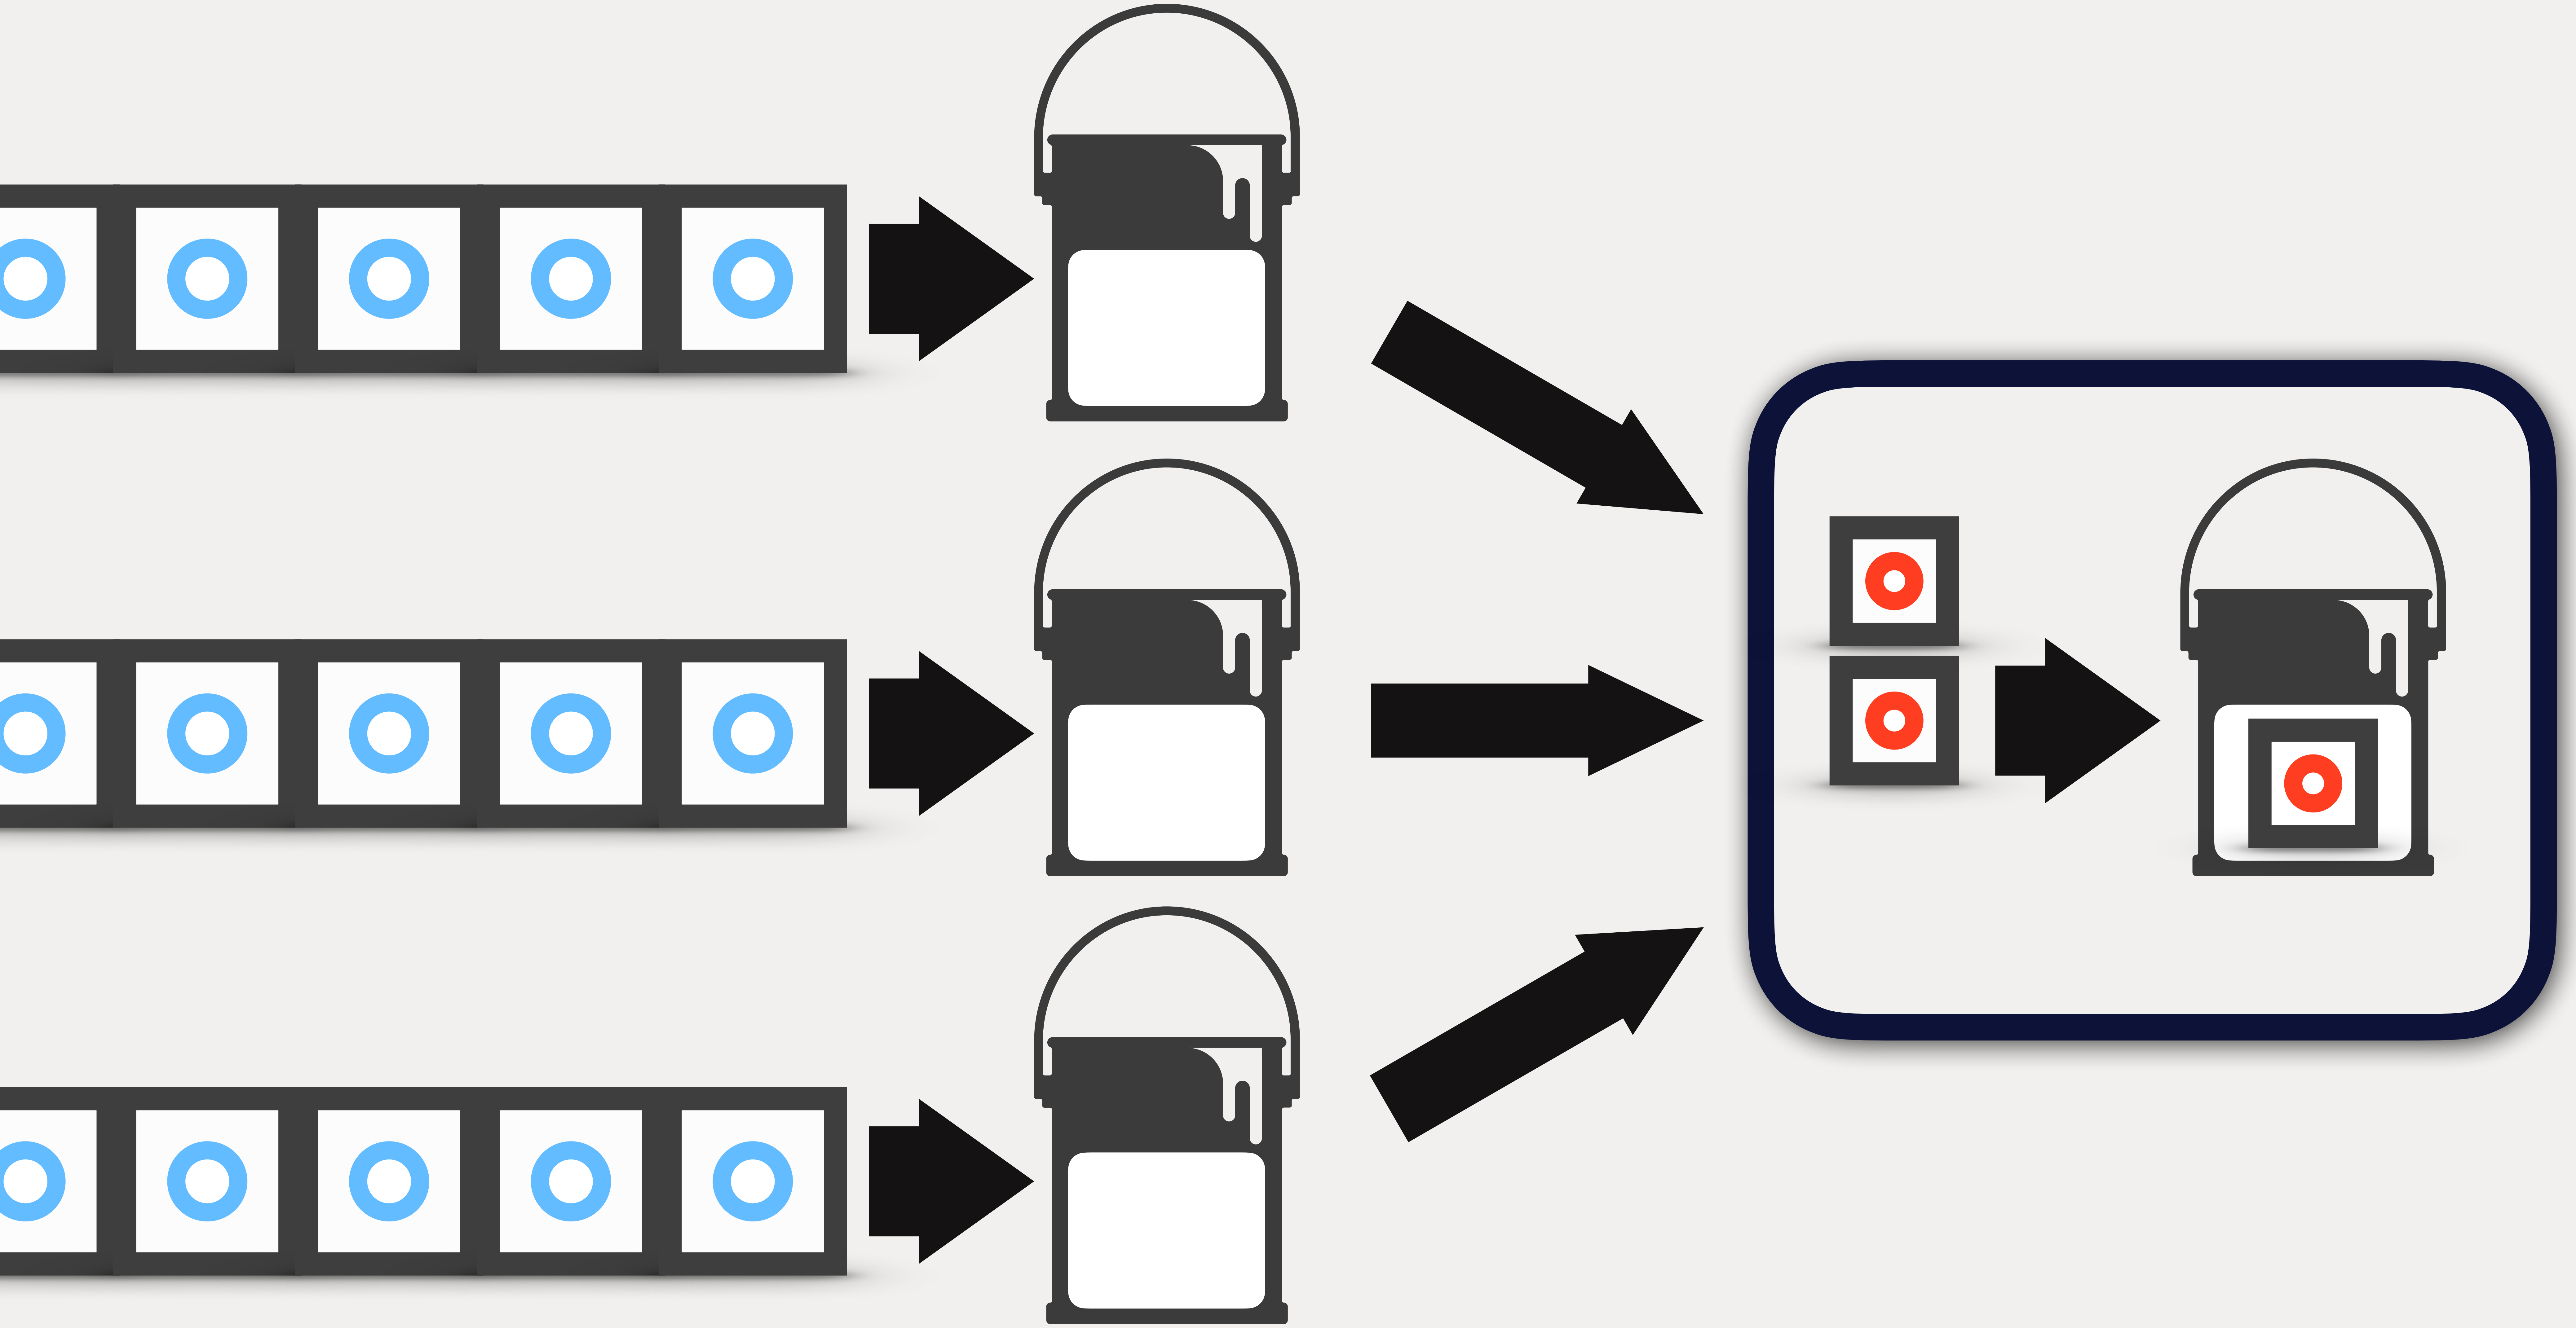

Reservoir Reuse

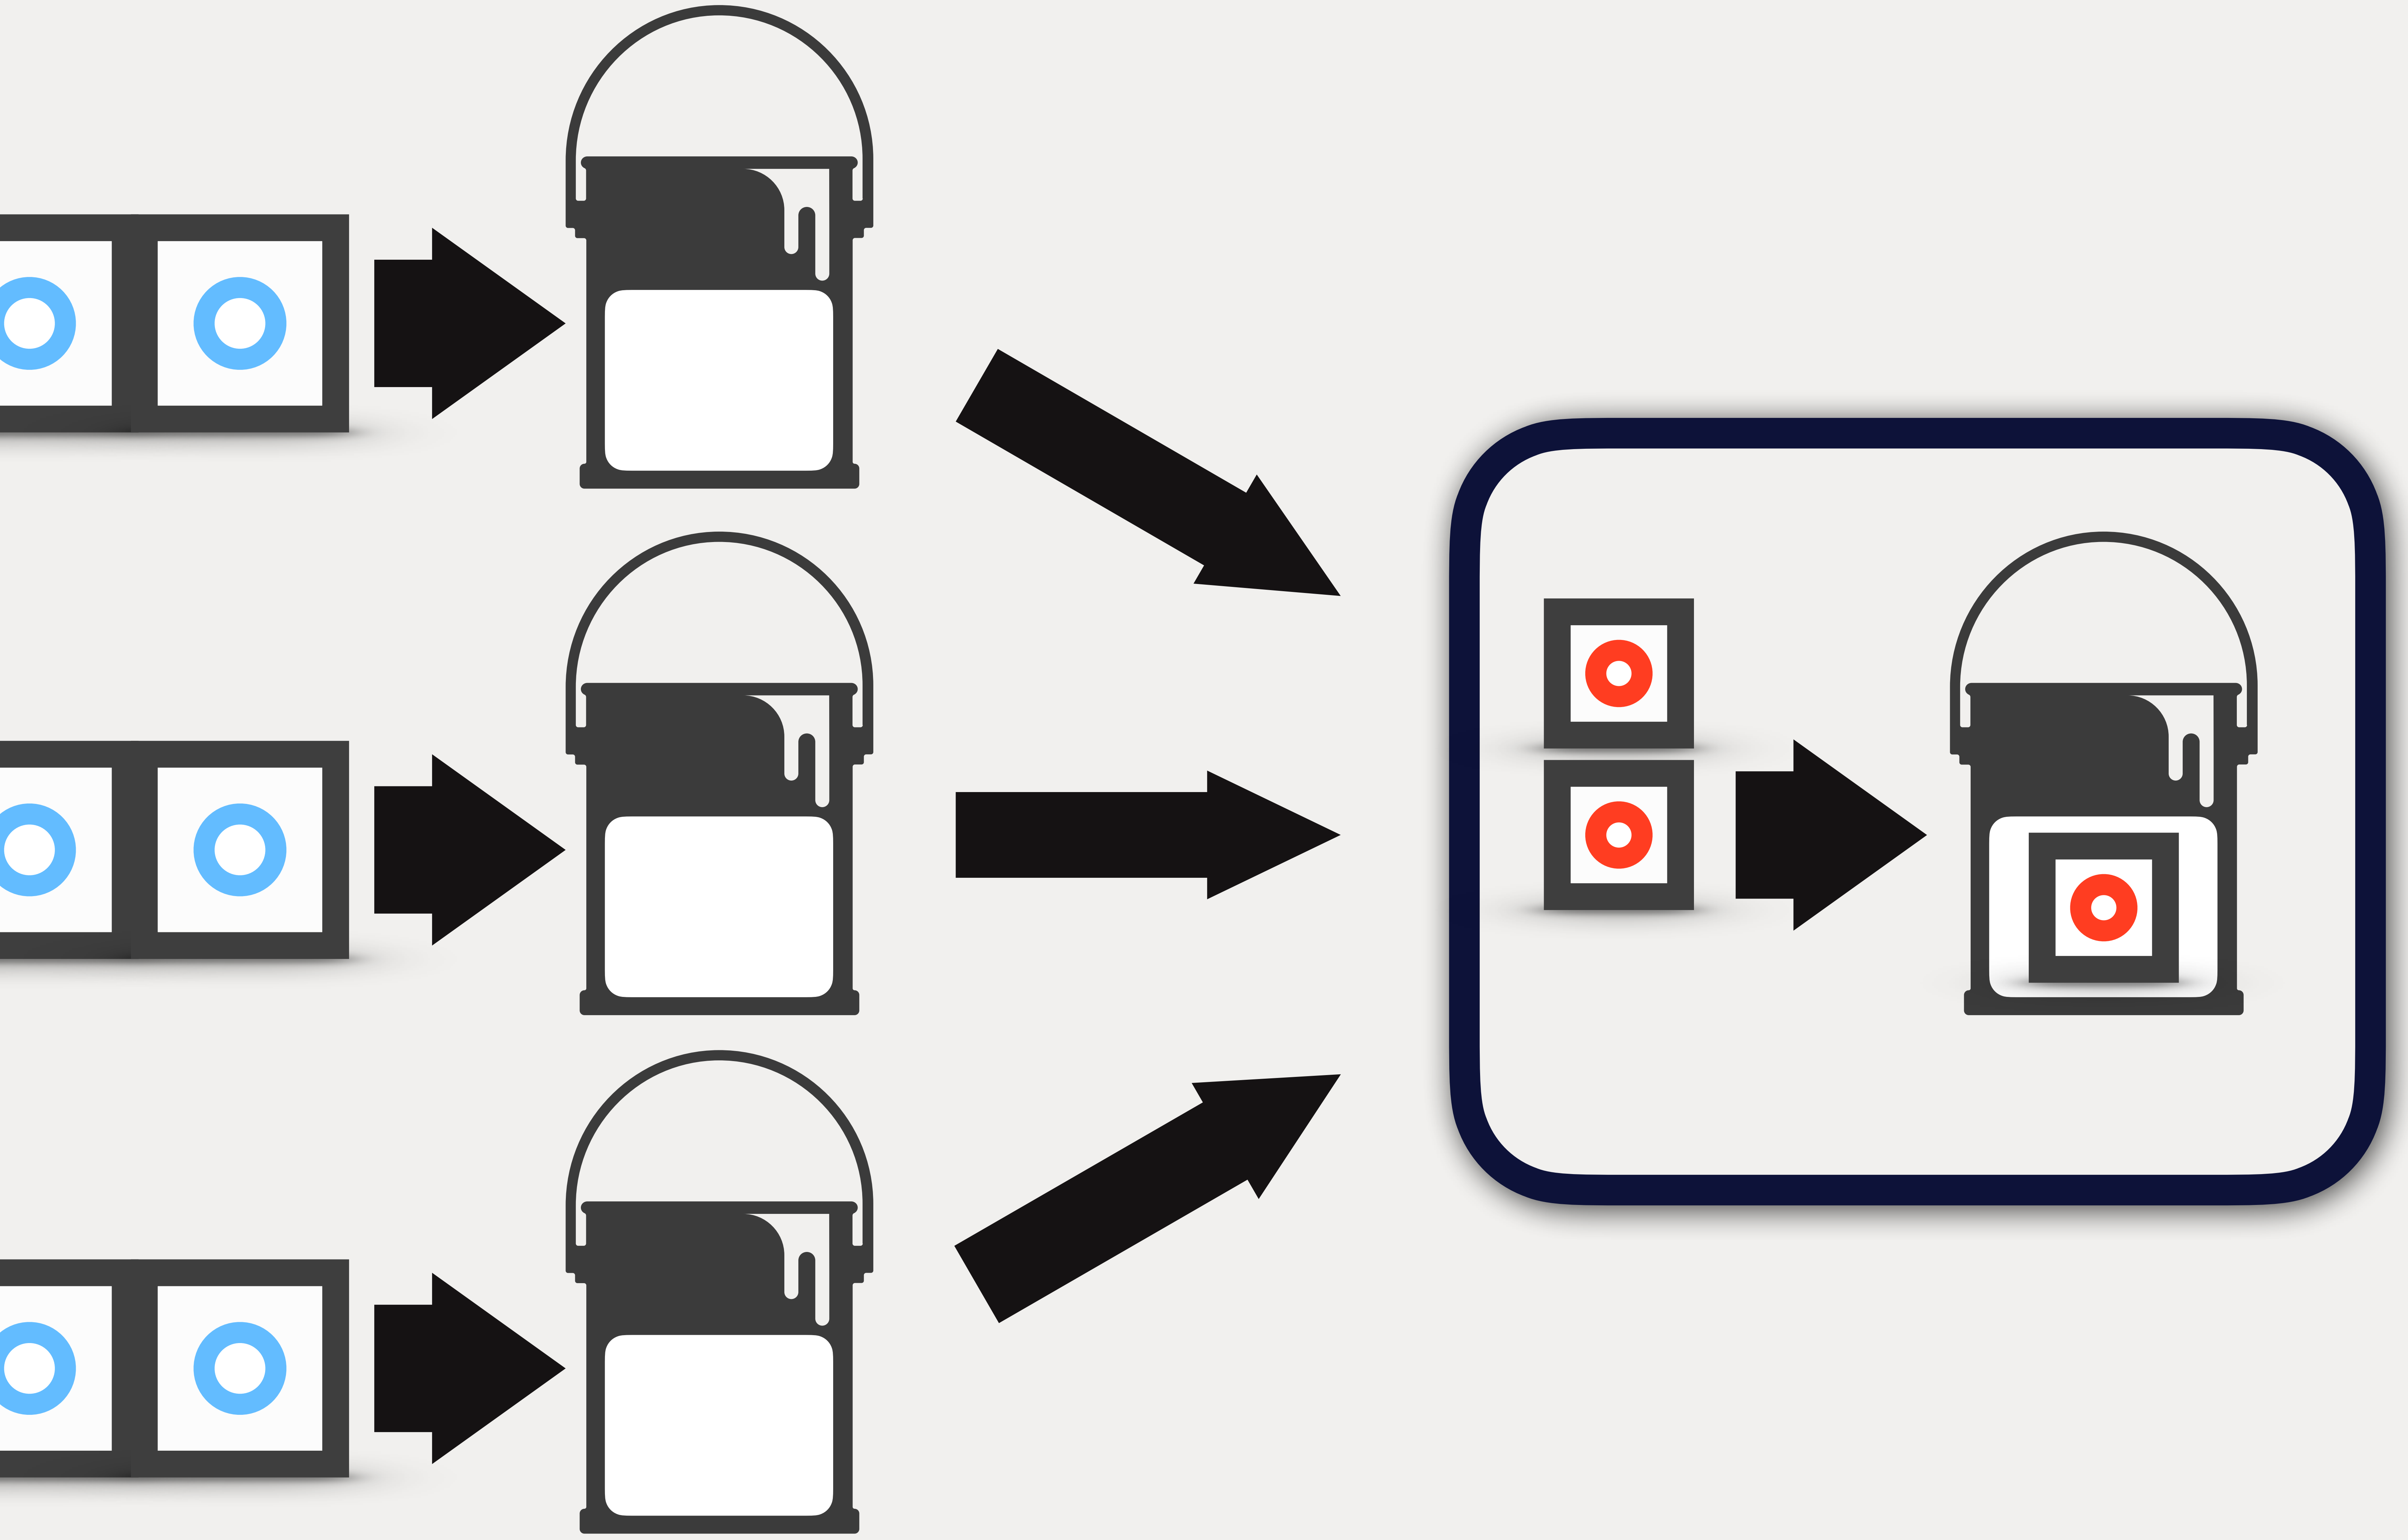

Reservoir Reuse

Spatial Reuse

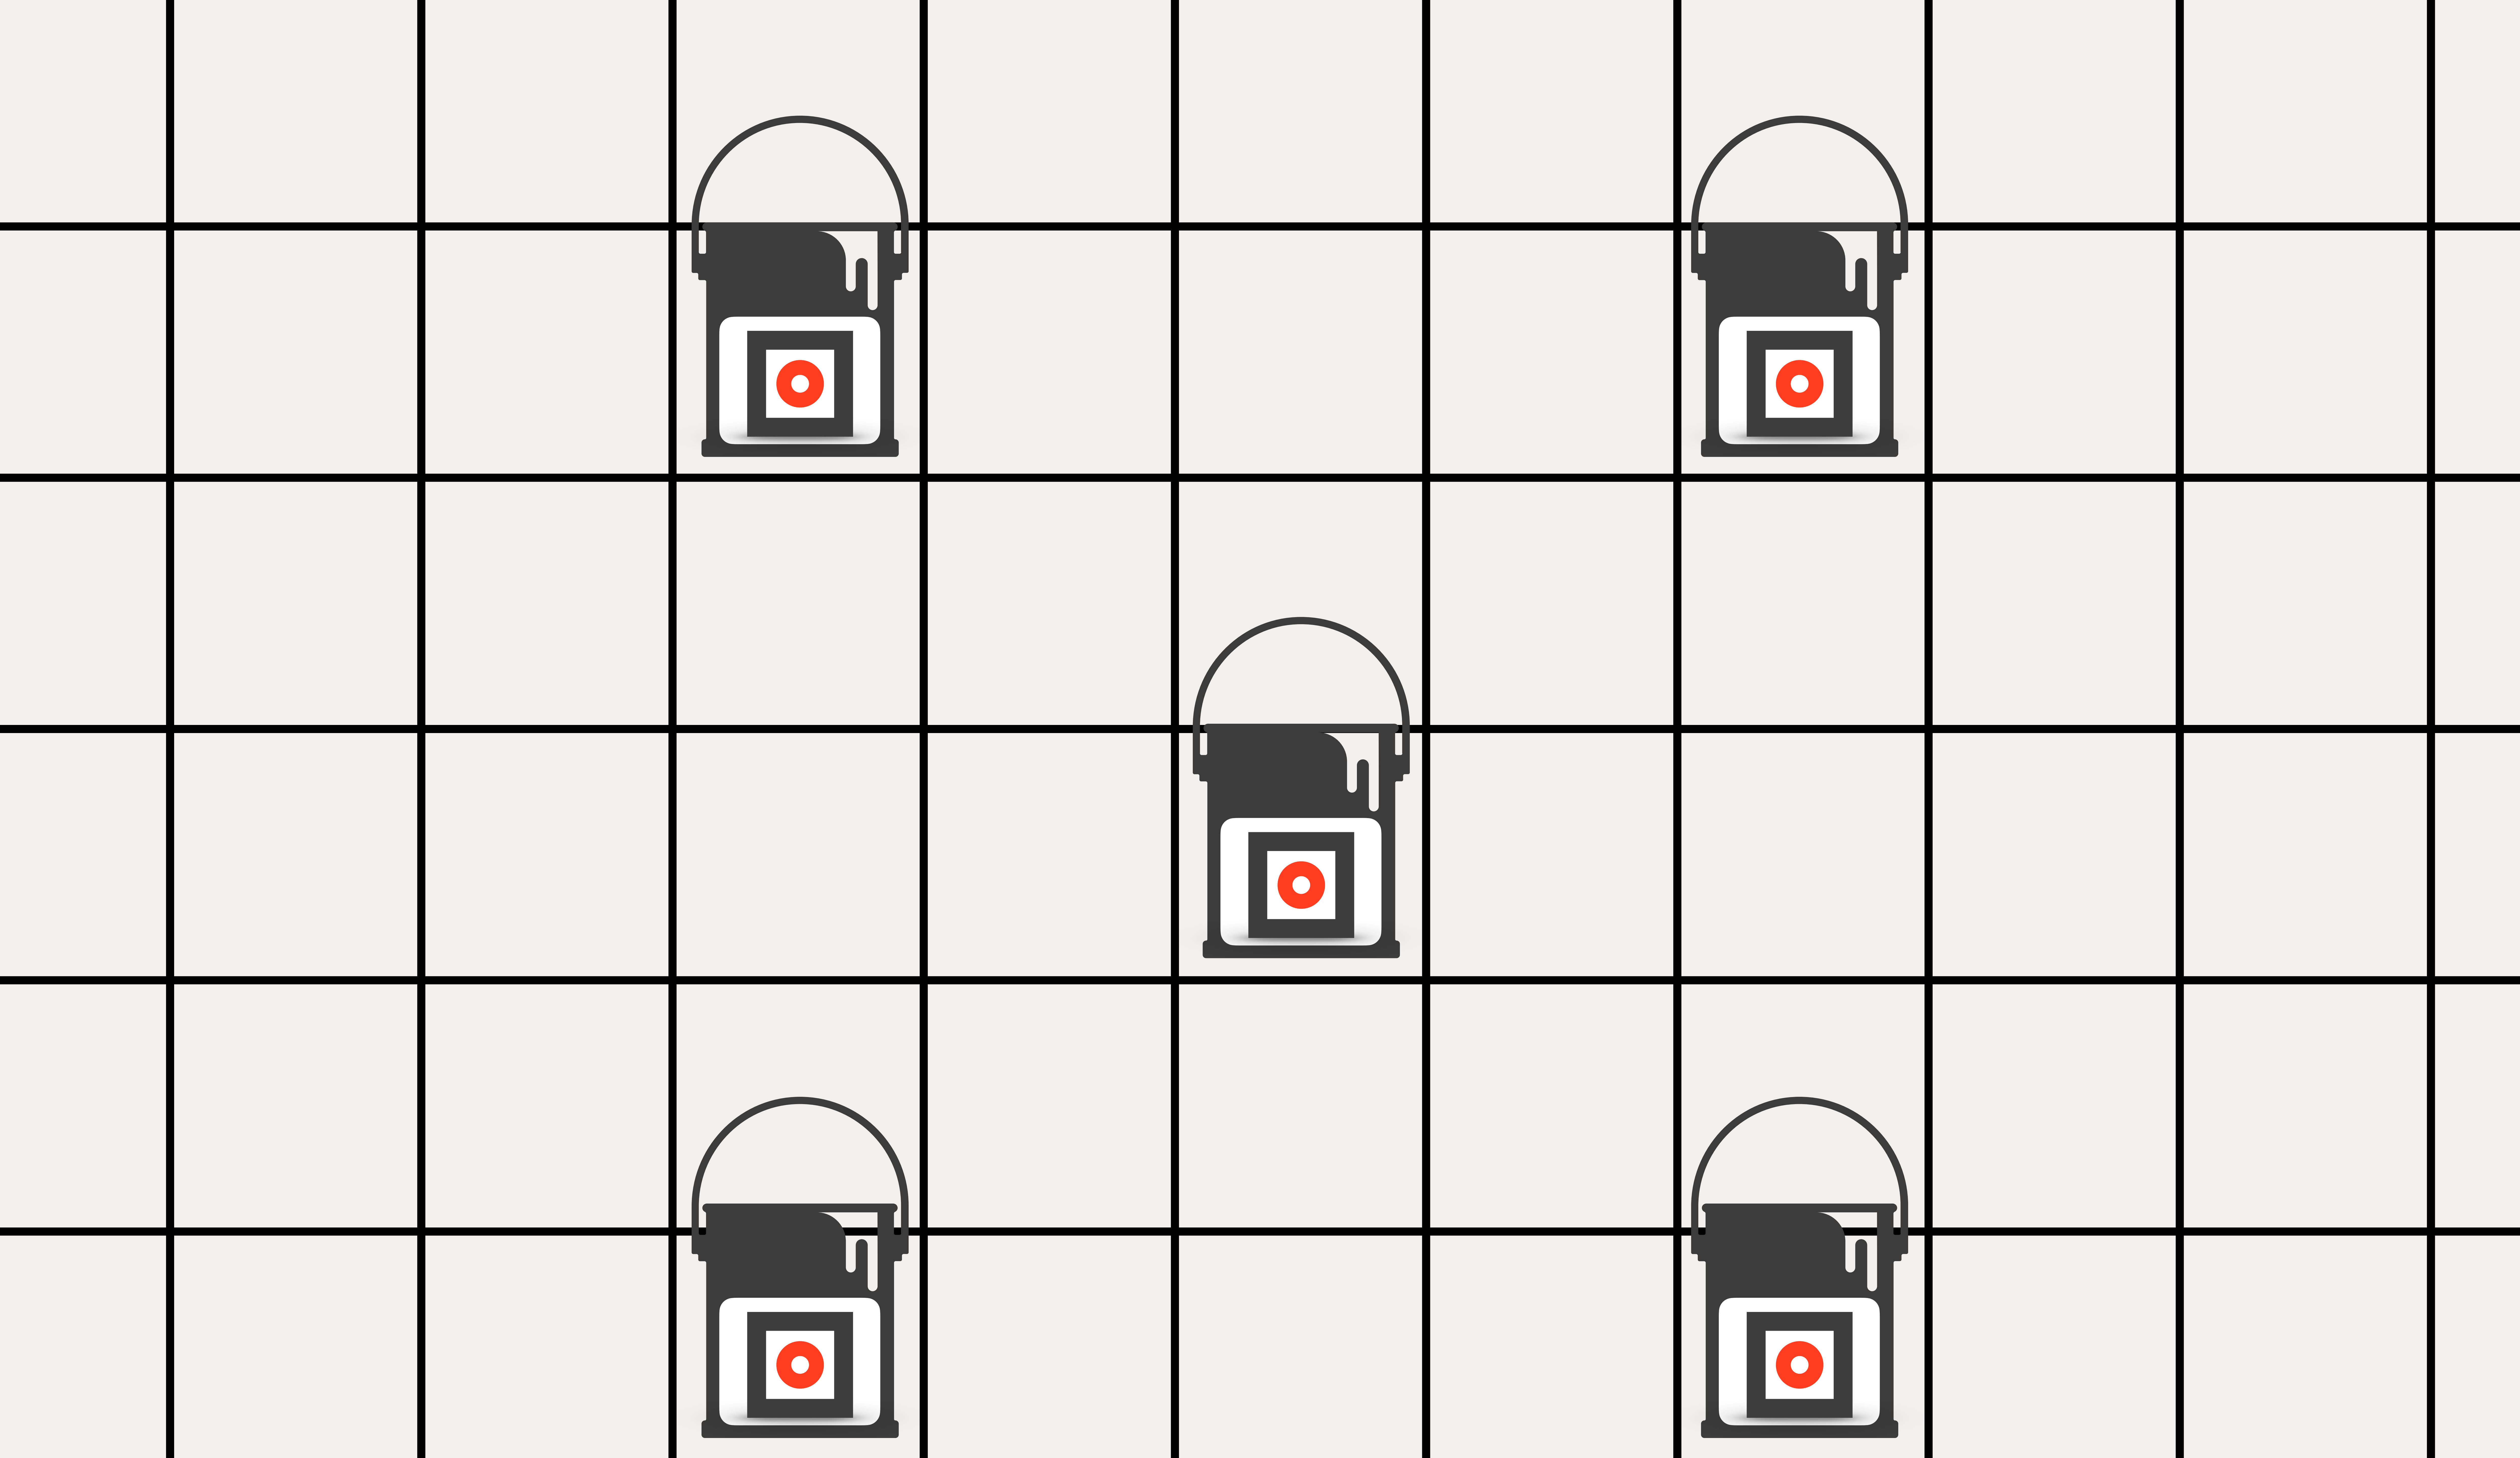

Temporal Reuse

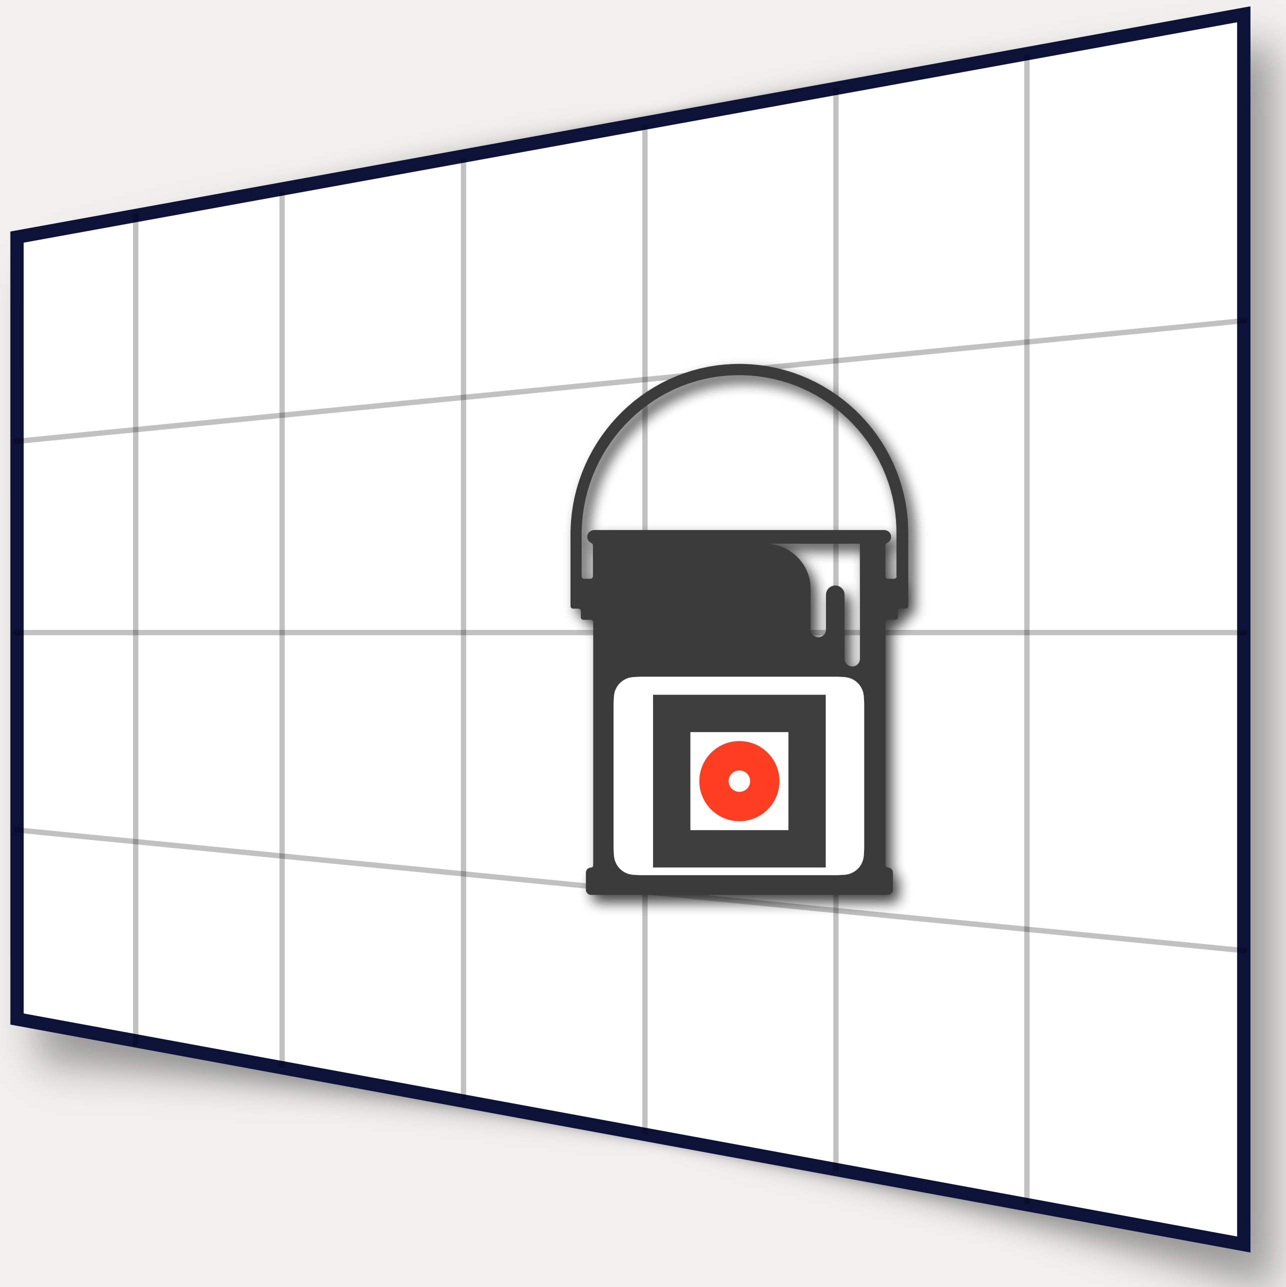

Temporal Reuse

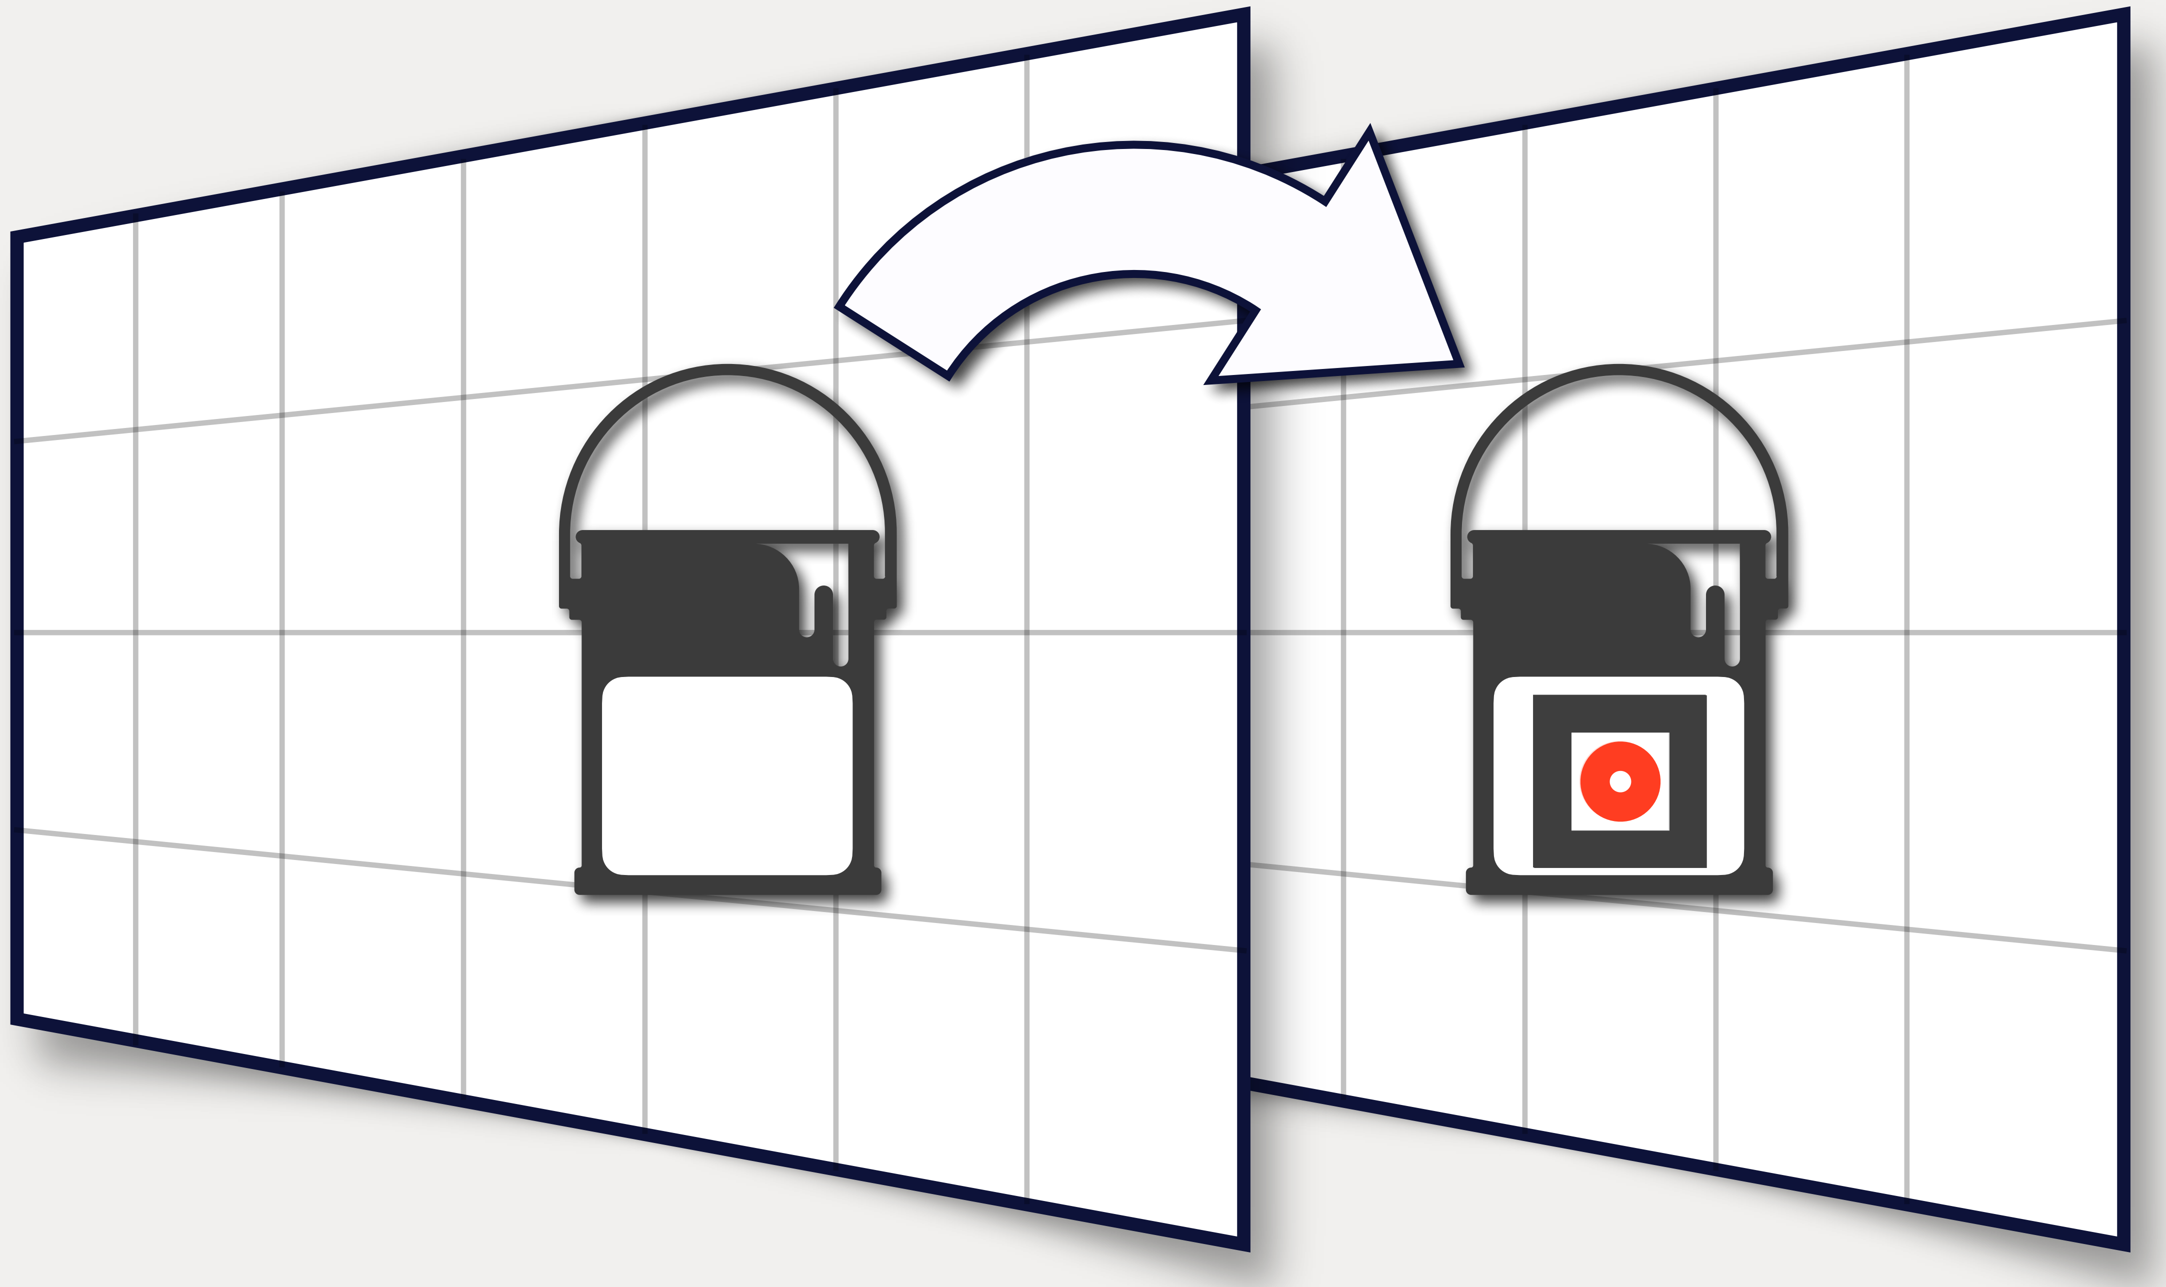

Temporal Reuse

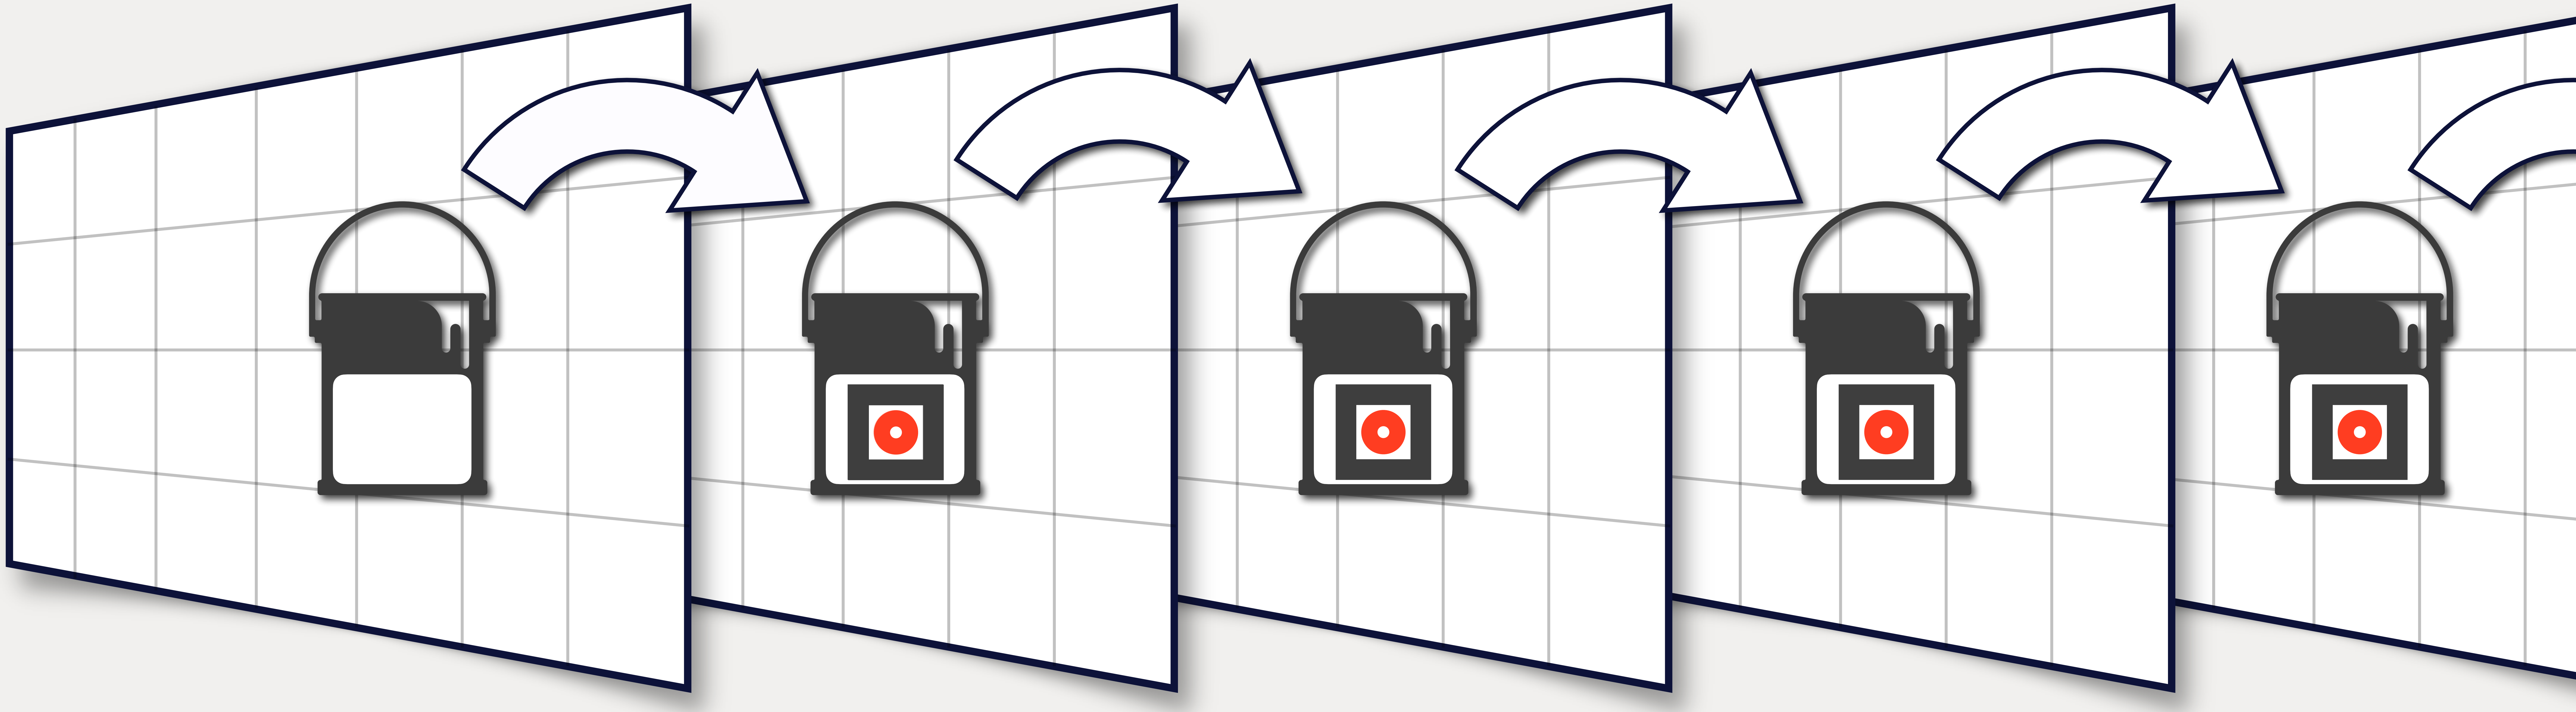

Visibility Reuse

Visibility Reuse

$$w(x_i) = \frac{\hat{p}(x_i)}{p(x_i)}$$

Visibility Reuse

Visibility Reuse

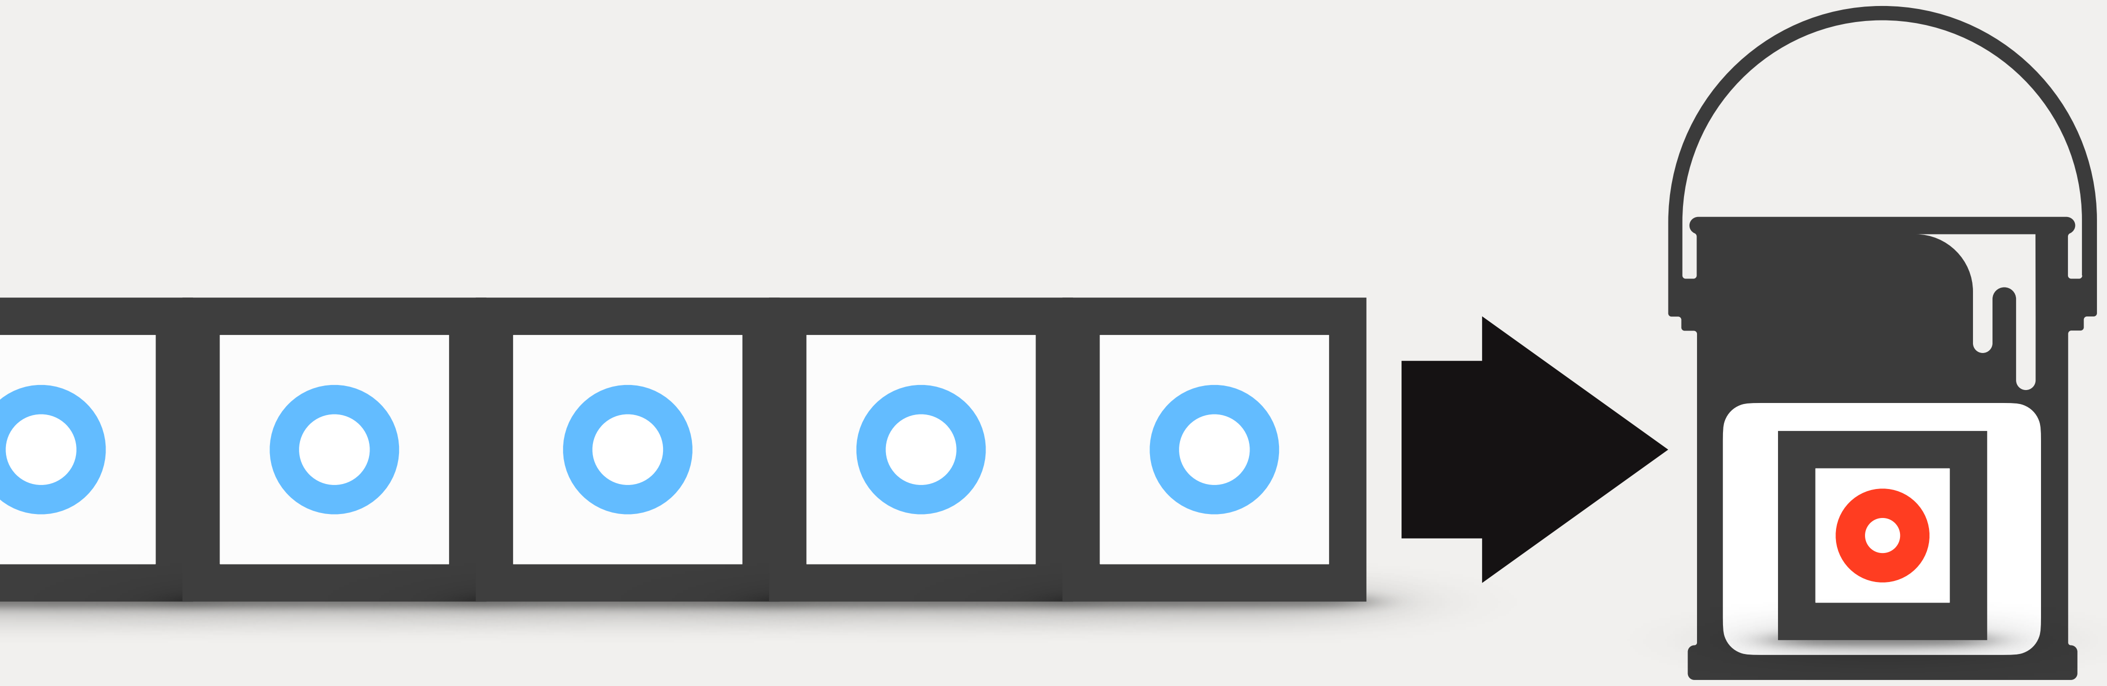

Visibility Reuse

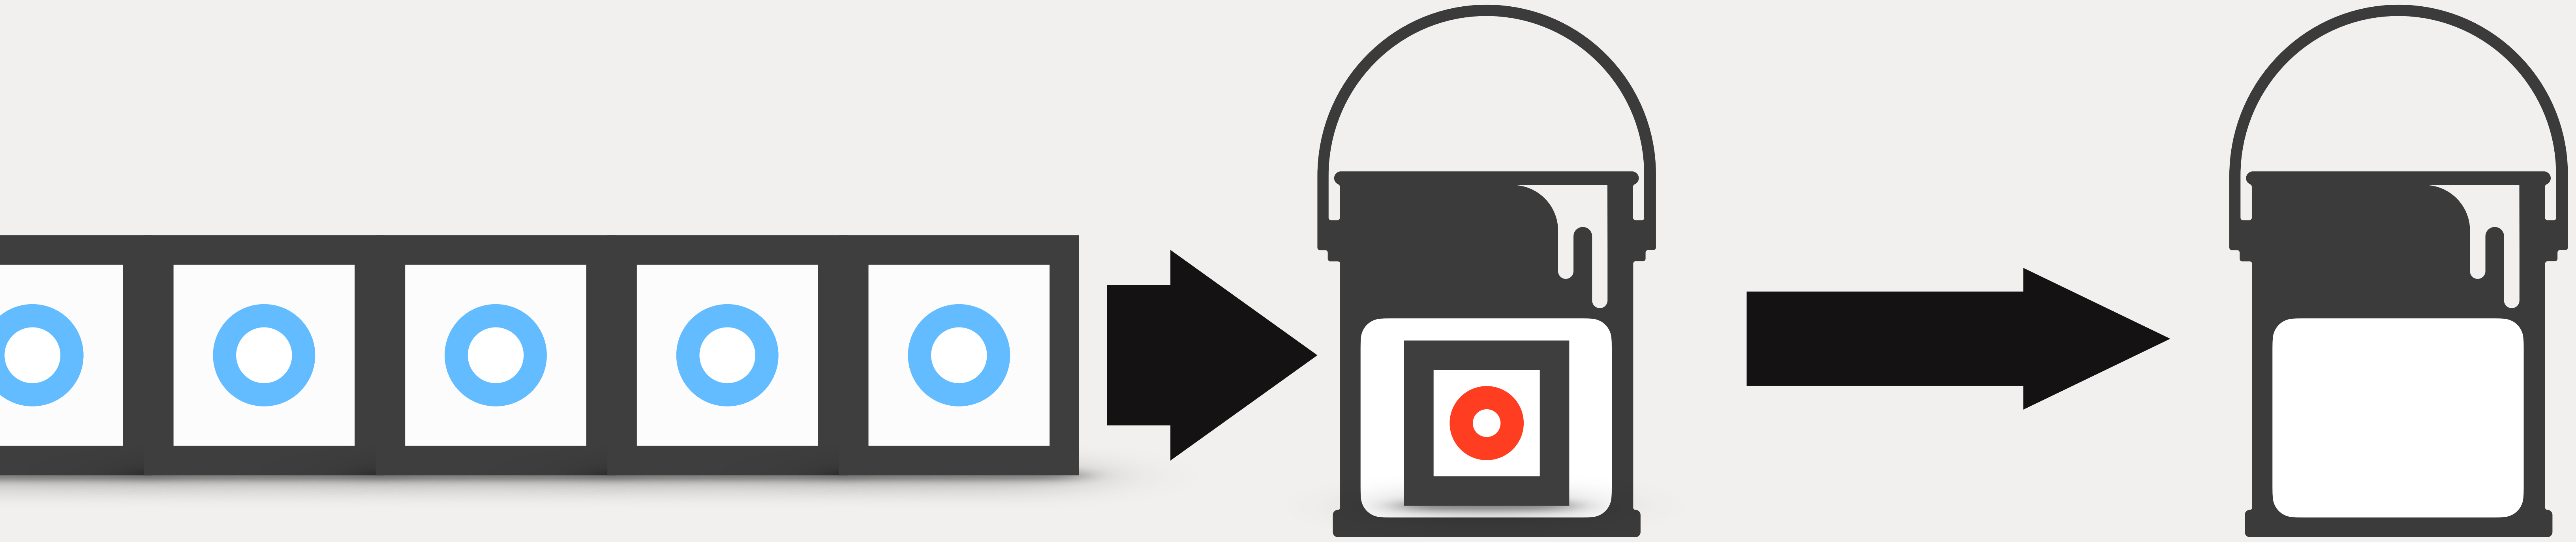

Visibility Reuse

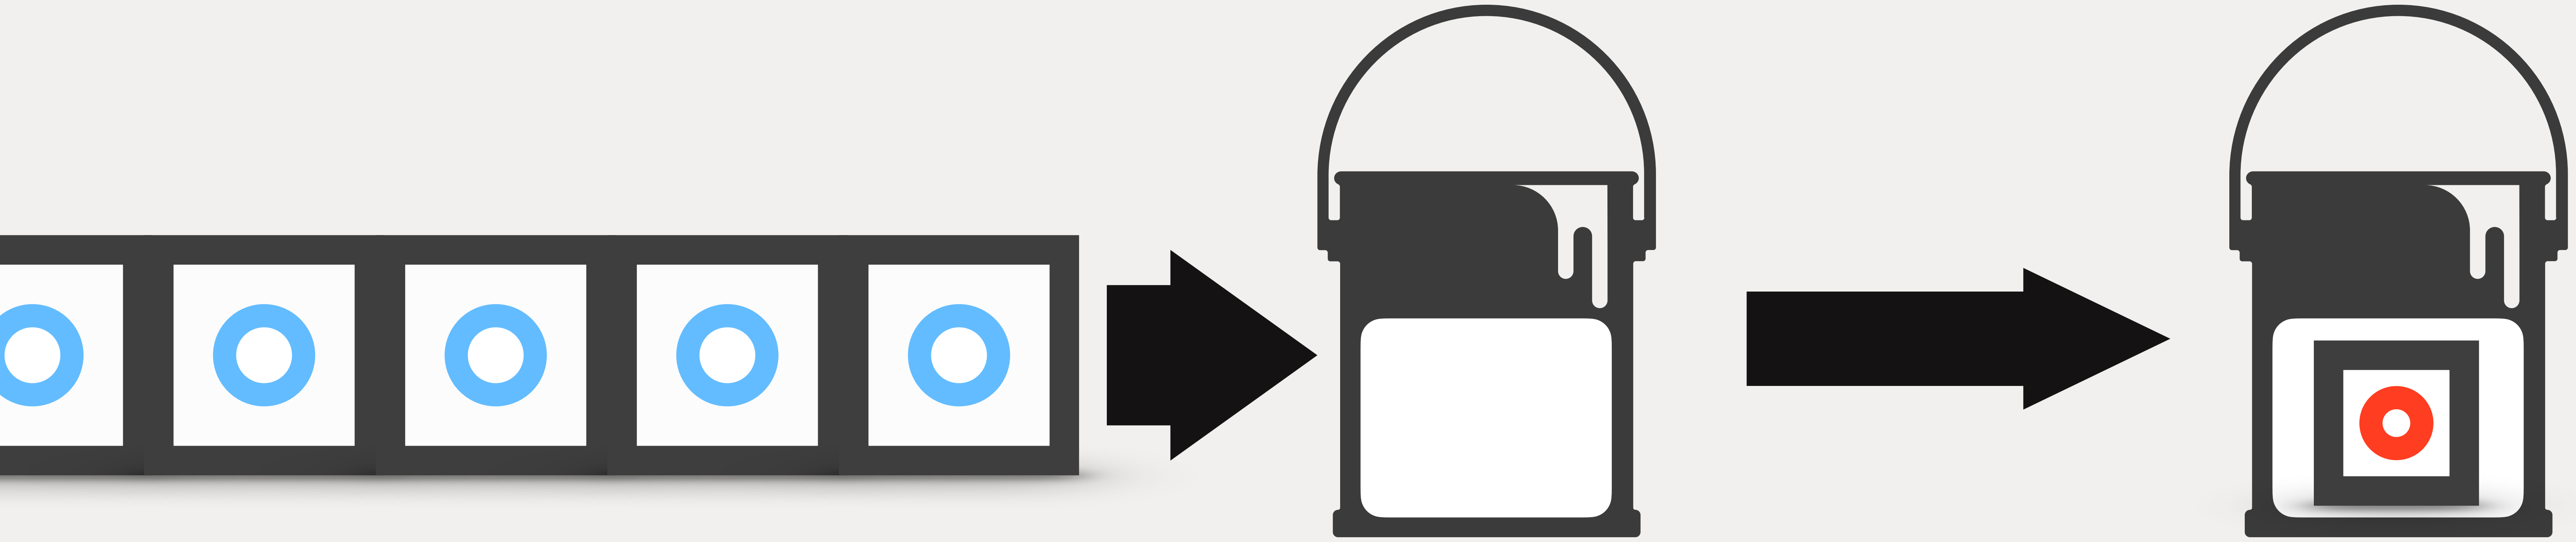

Visibility Reuse

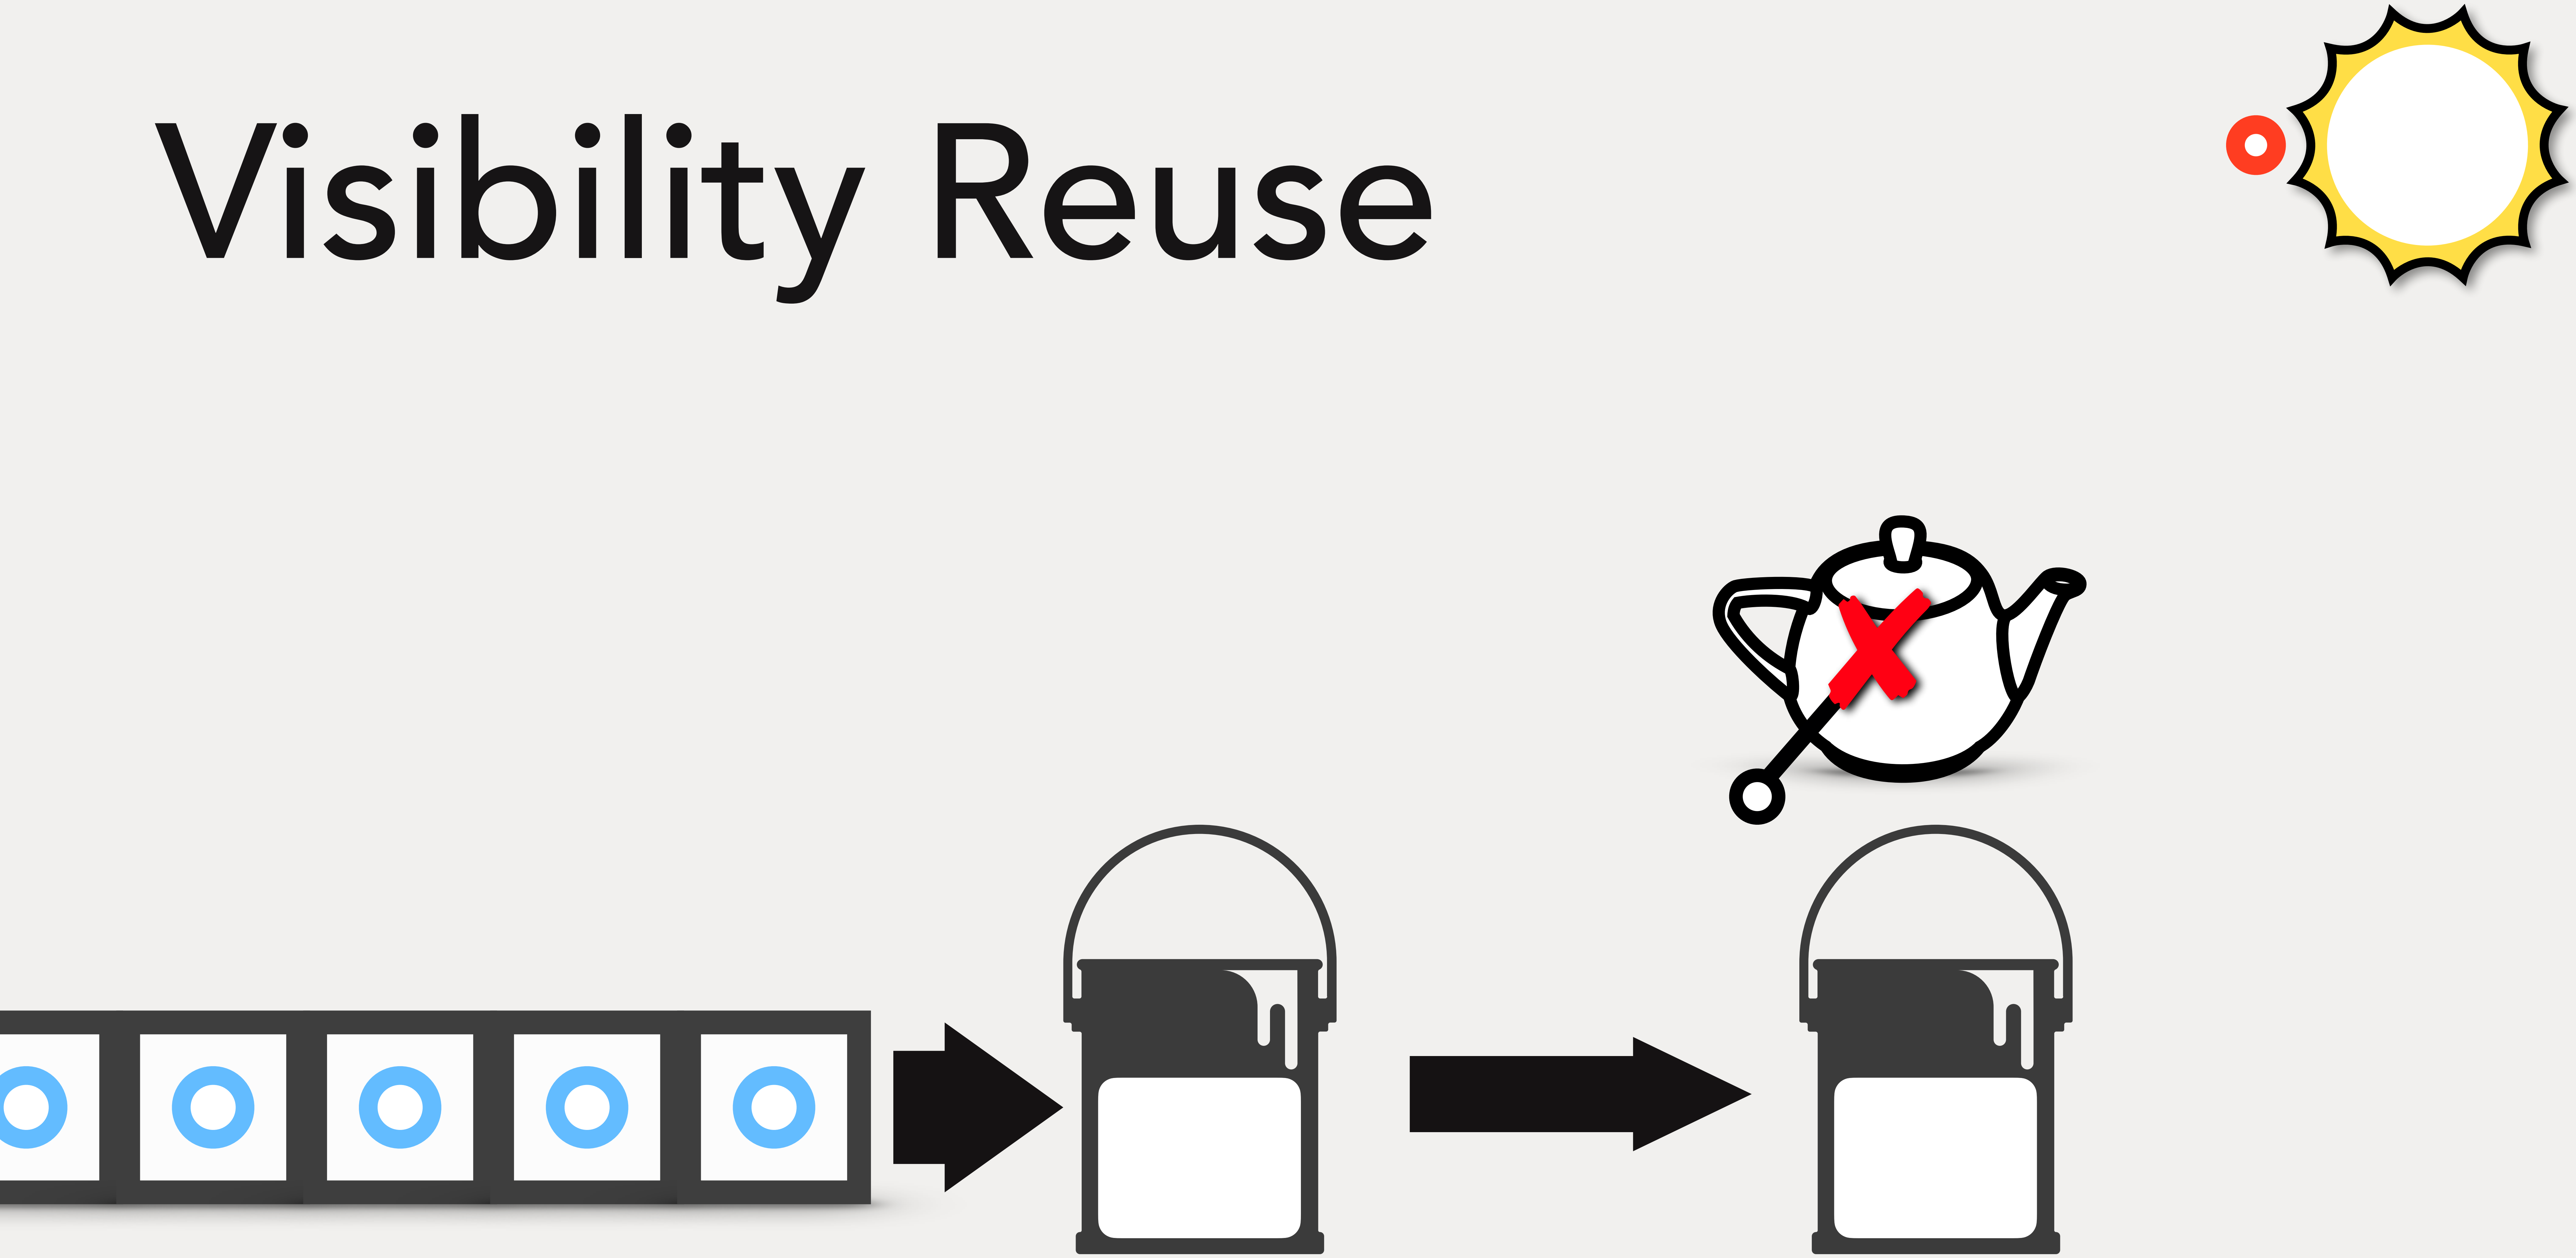

Visibility Reuse

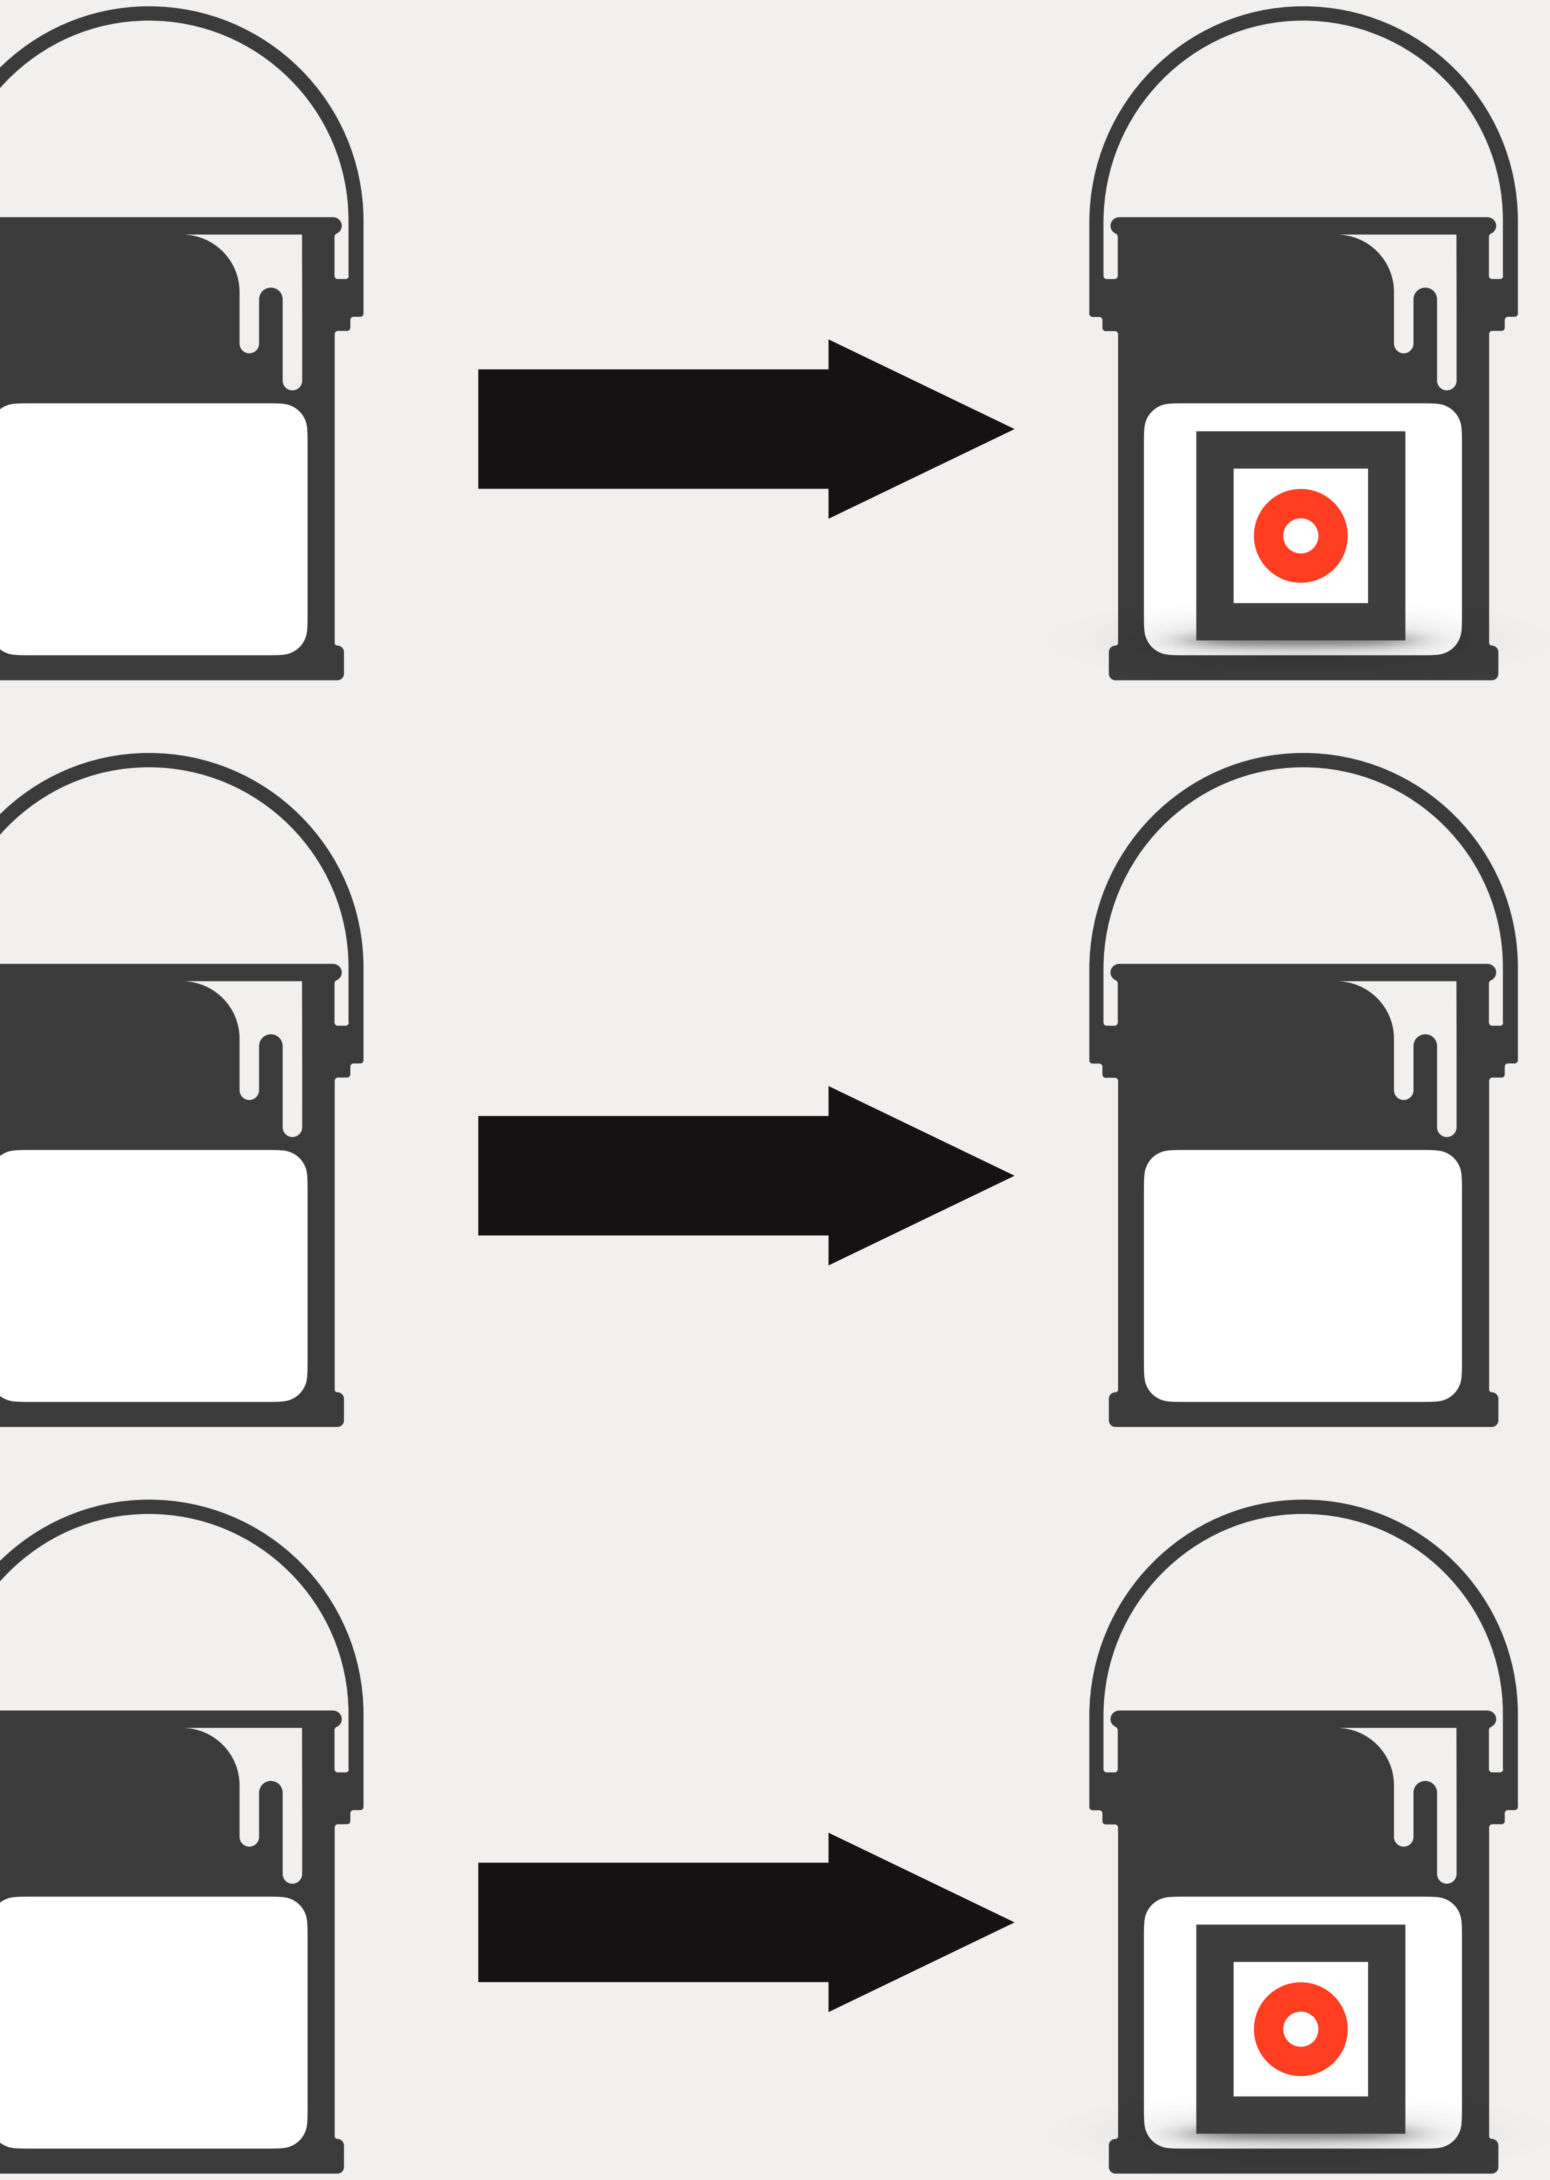

Visibility Reuse

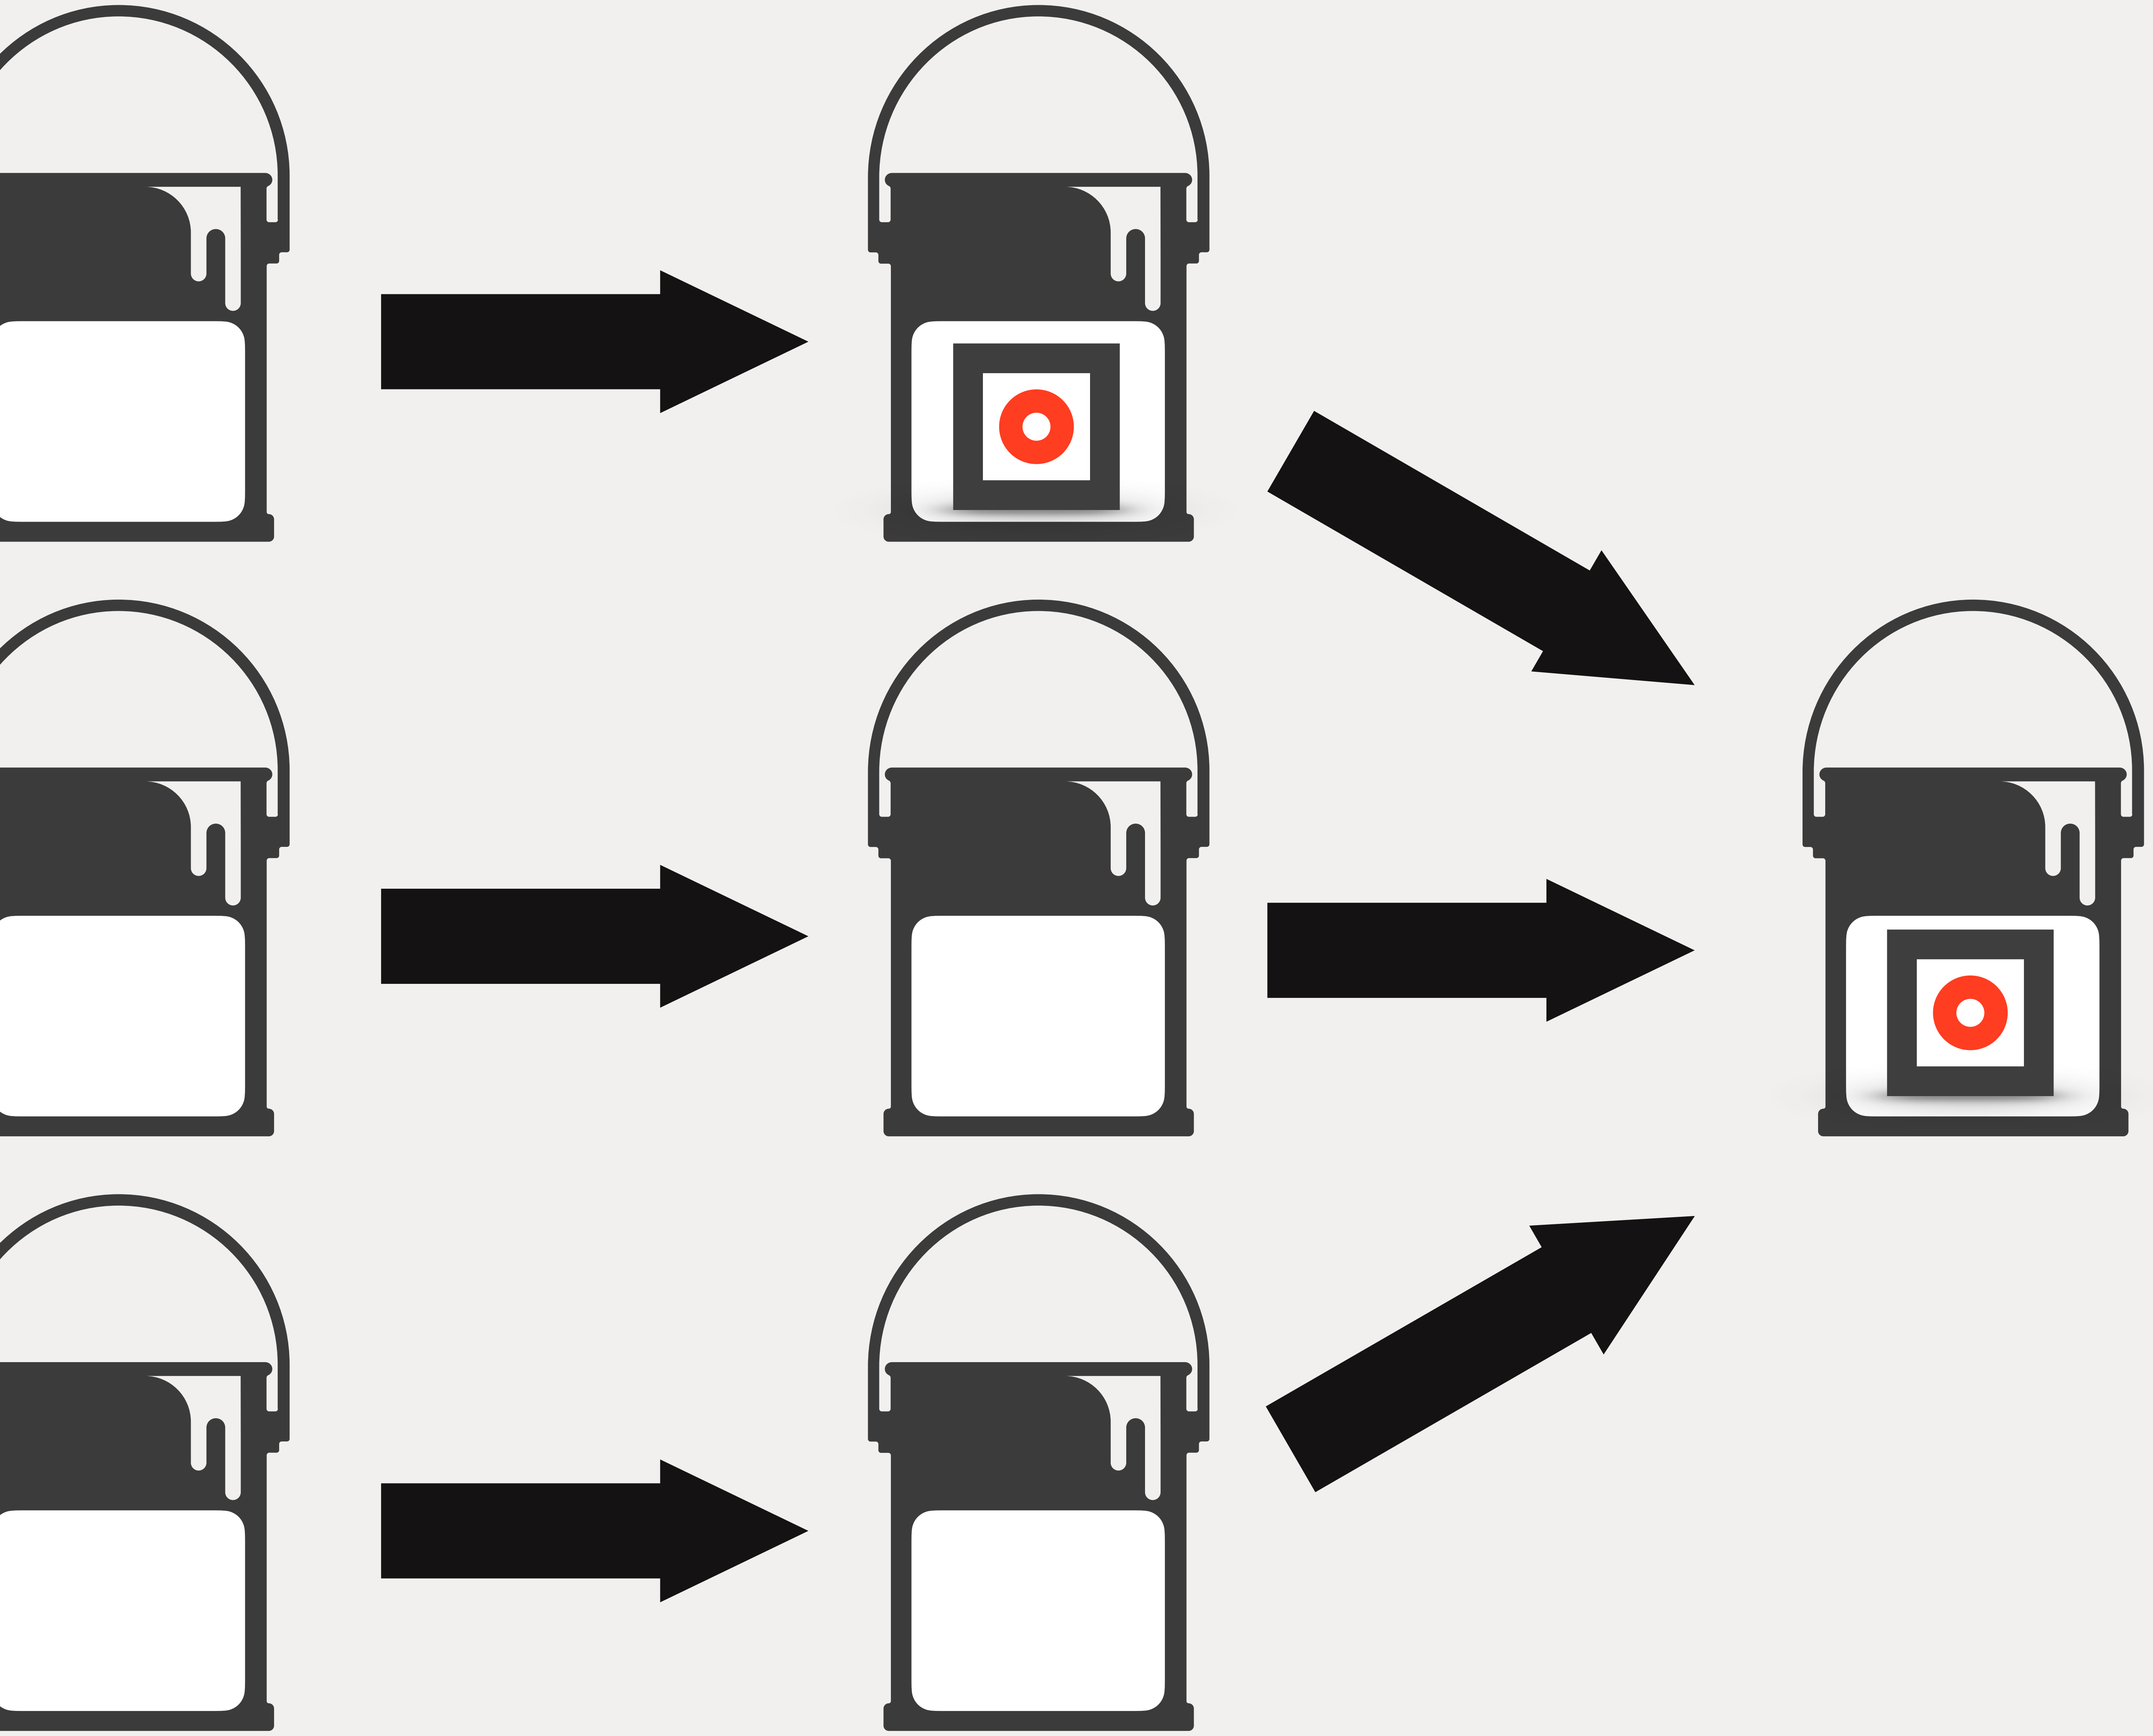

Visibility Reuse

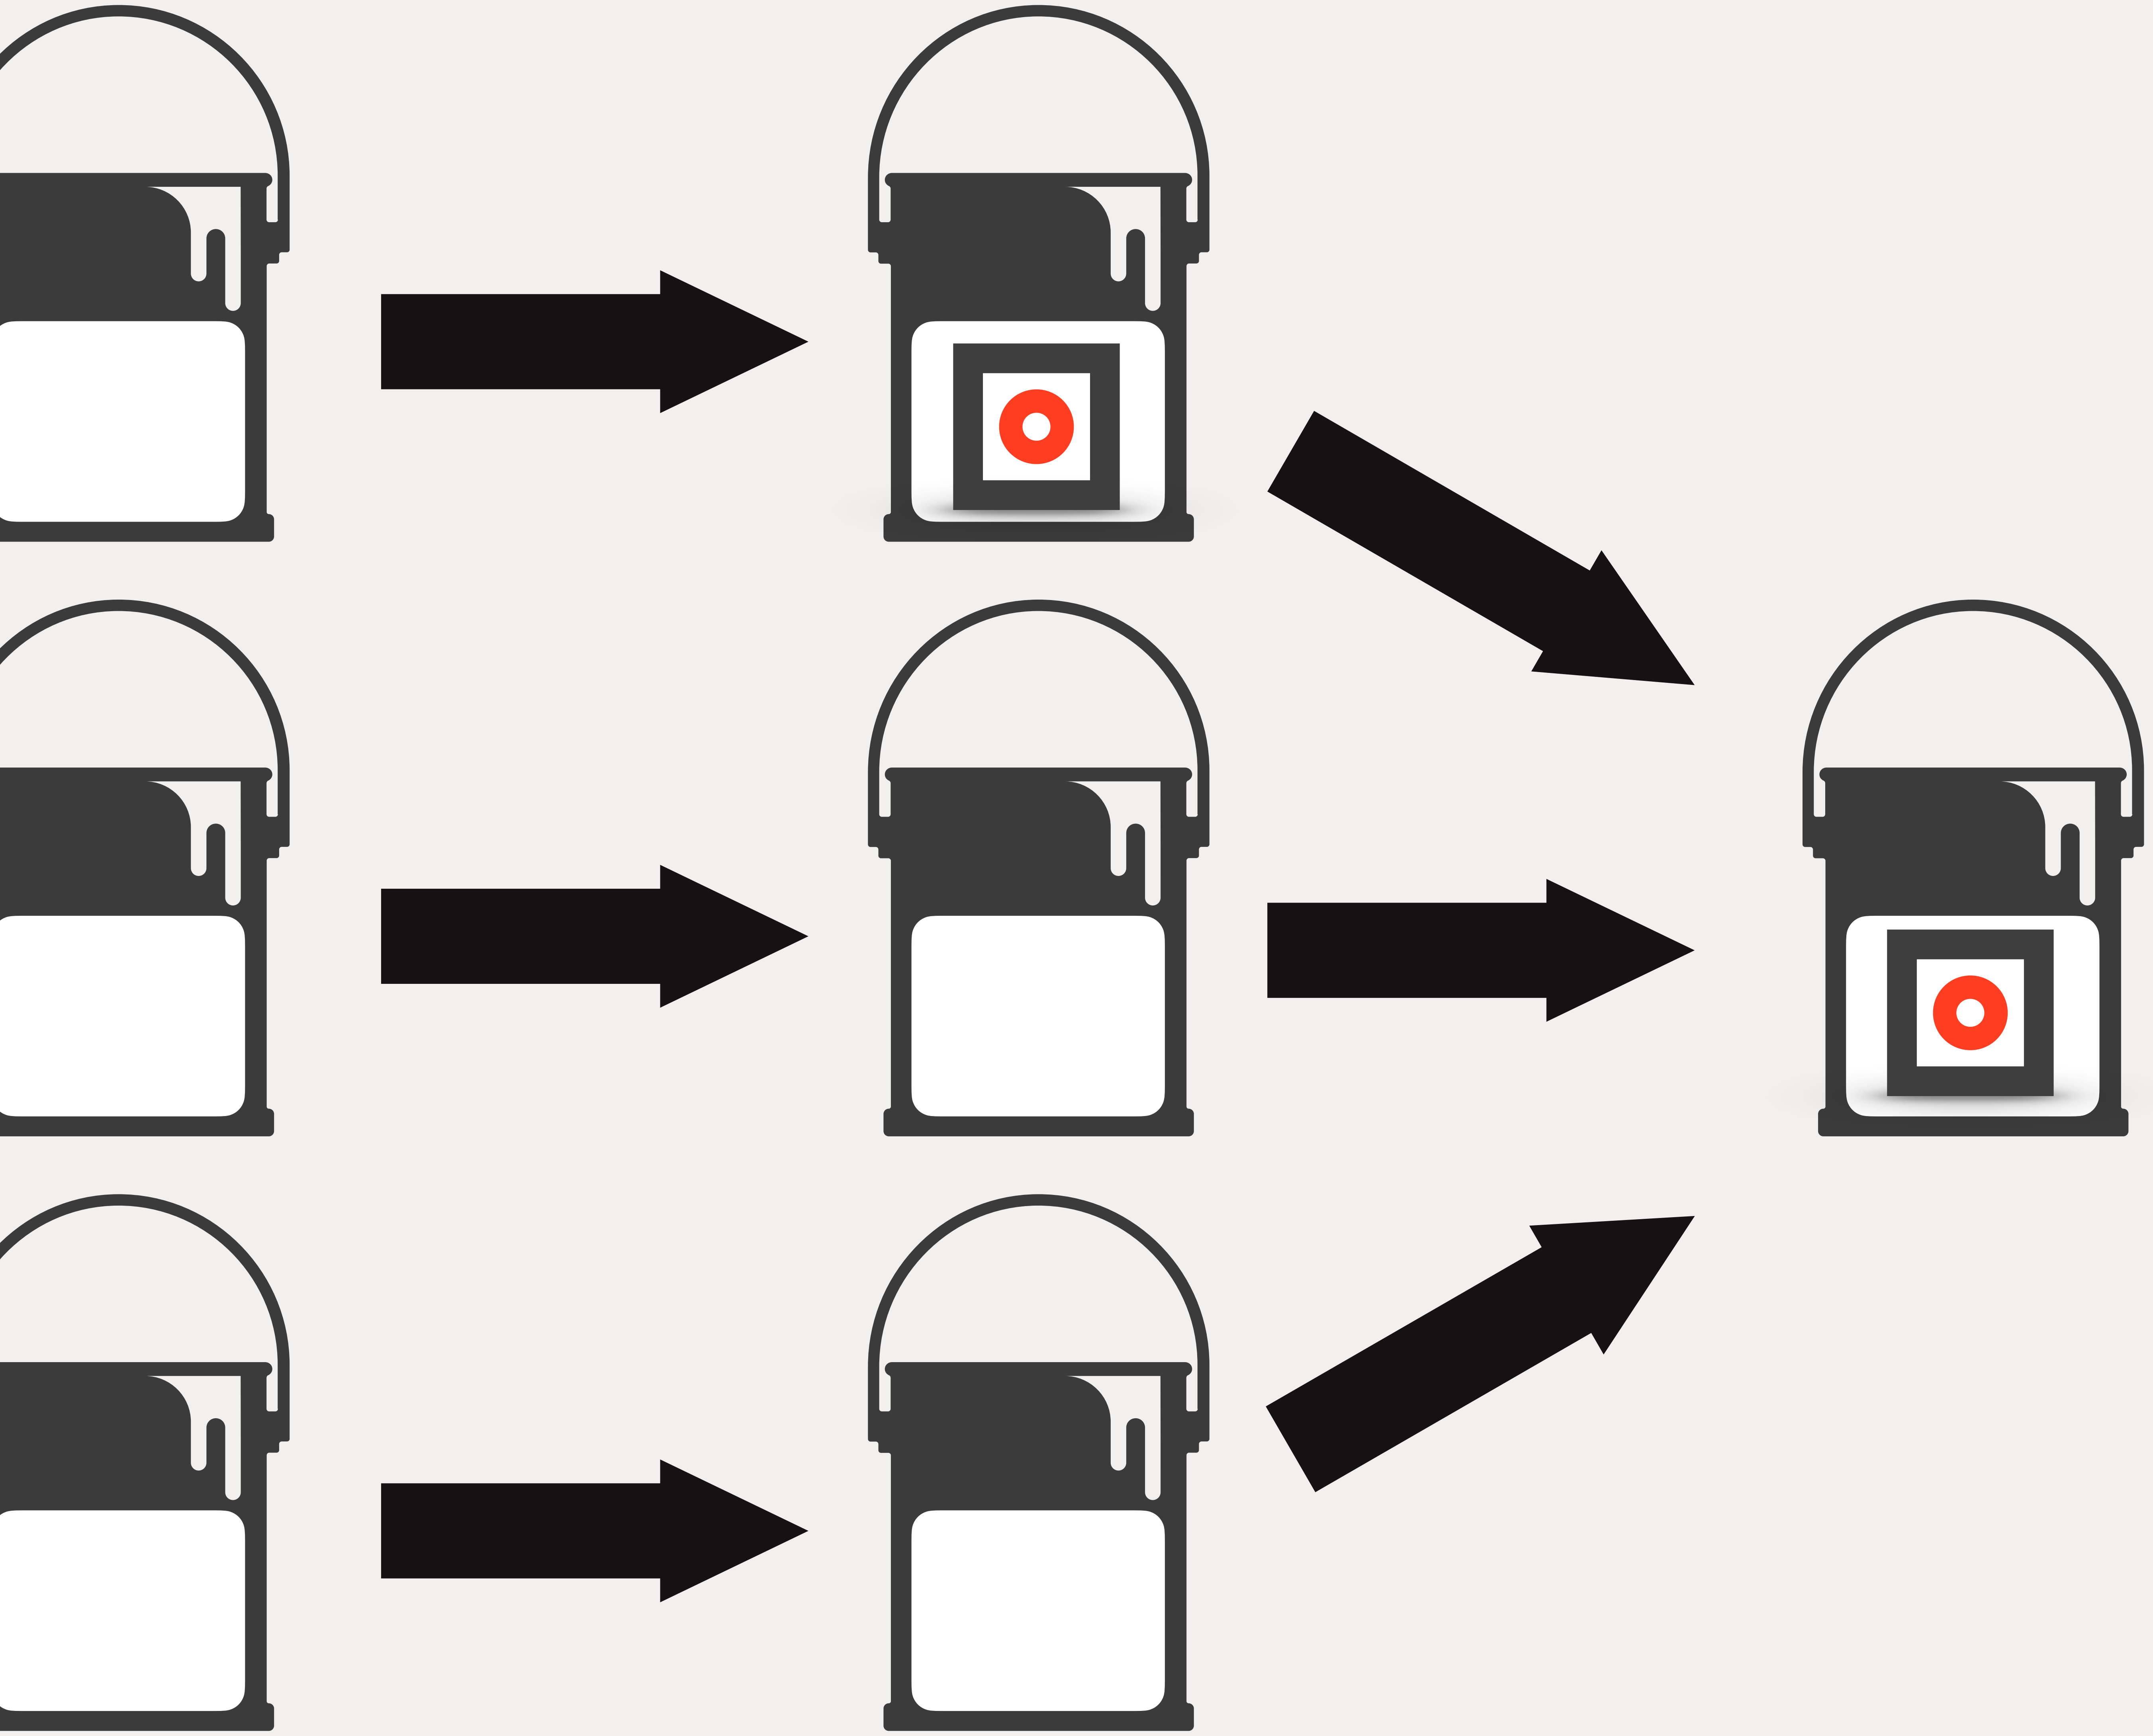

Includes visibility!

Unbiased Reuse

Unbiased Reuse

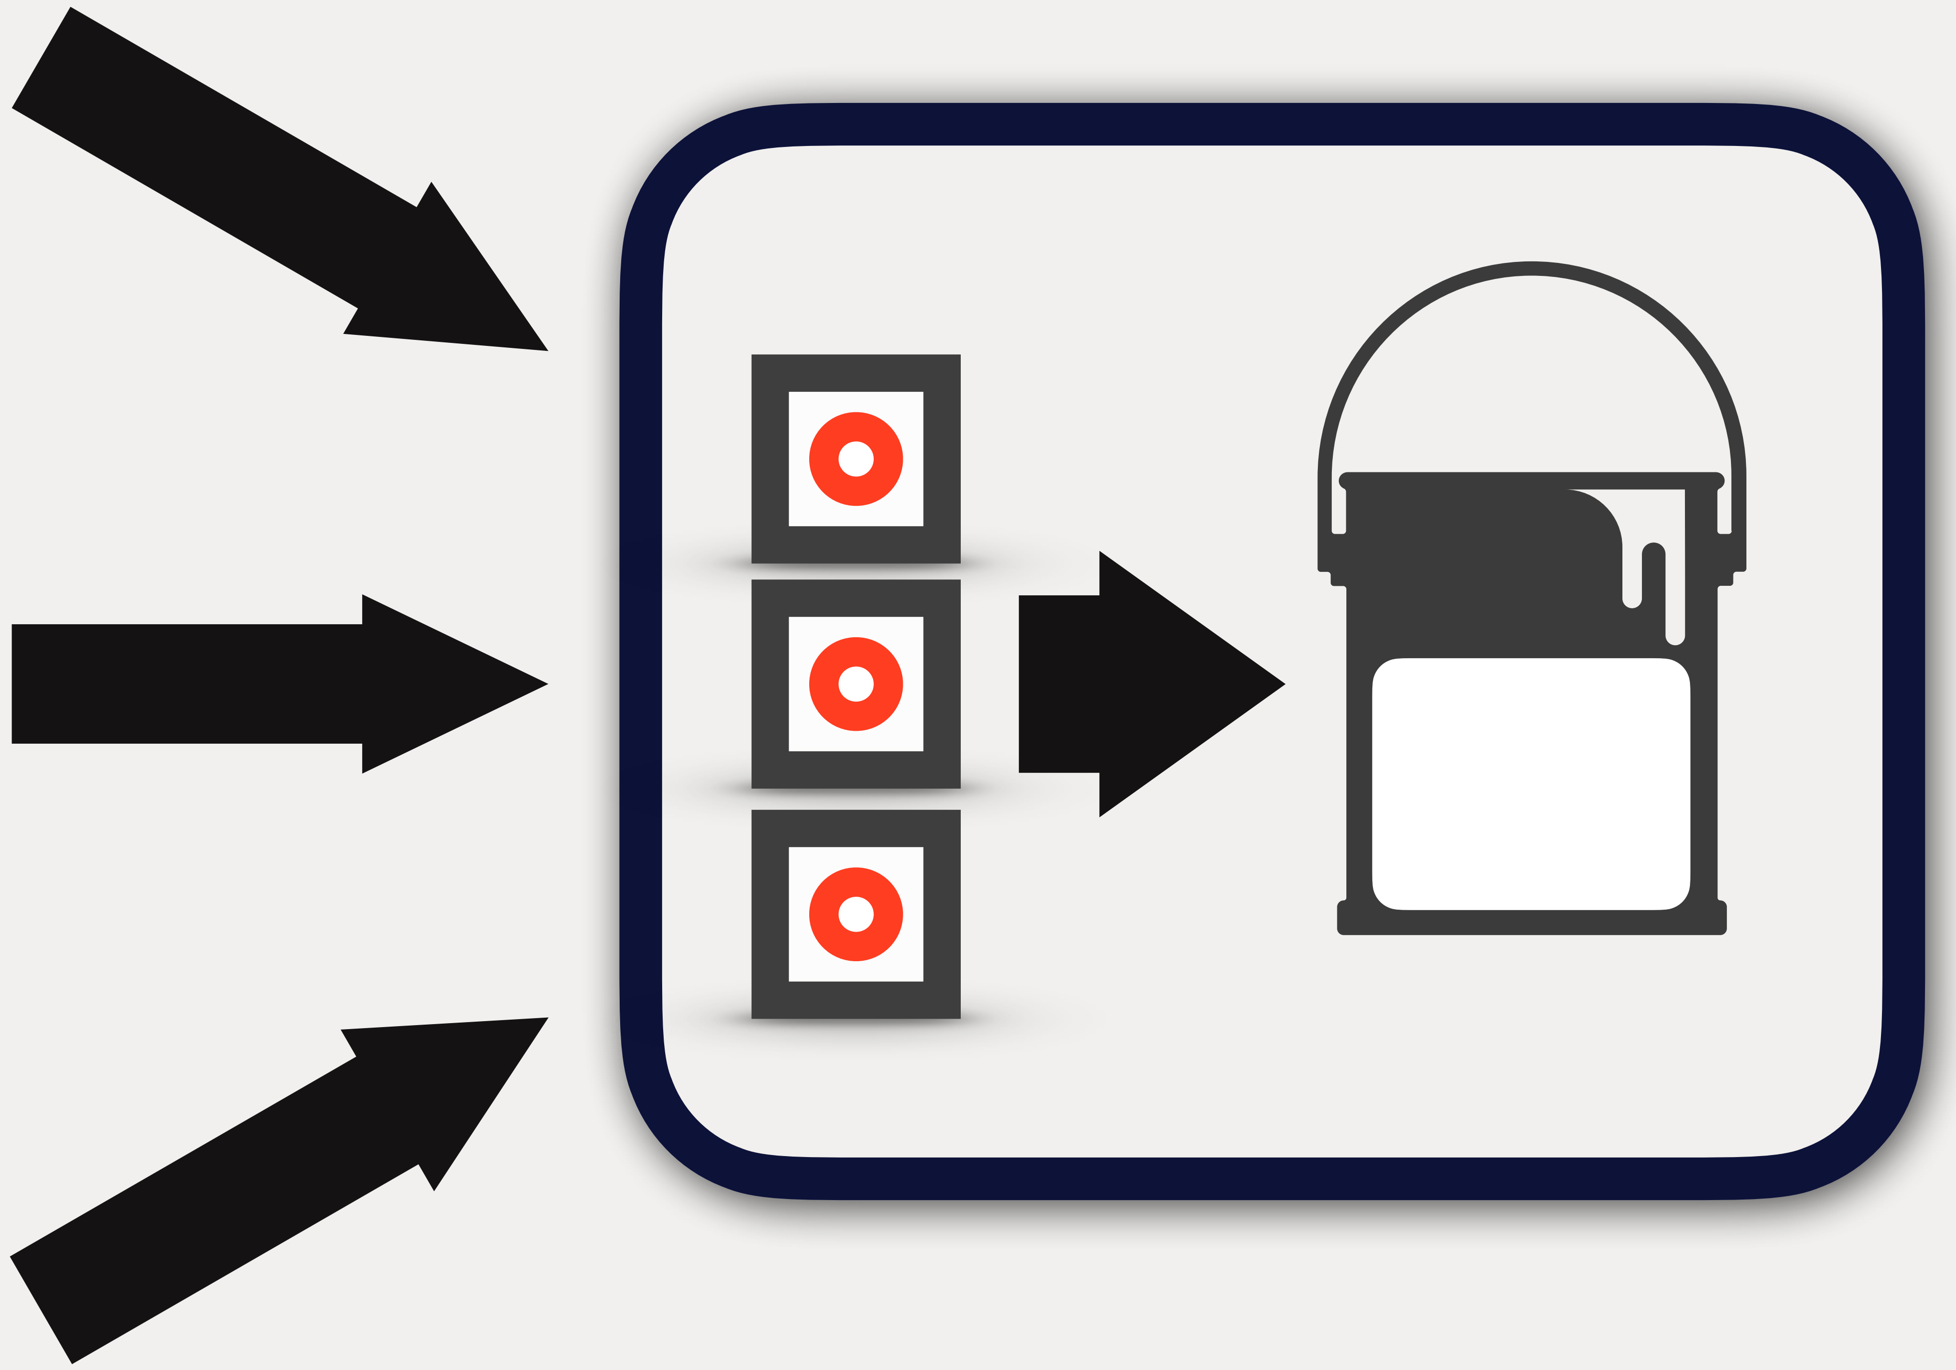

Unbiased Reuse

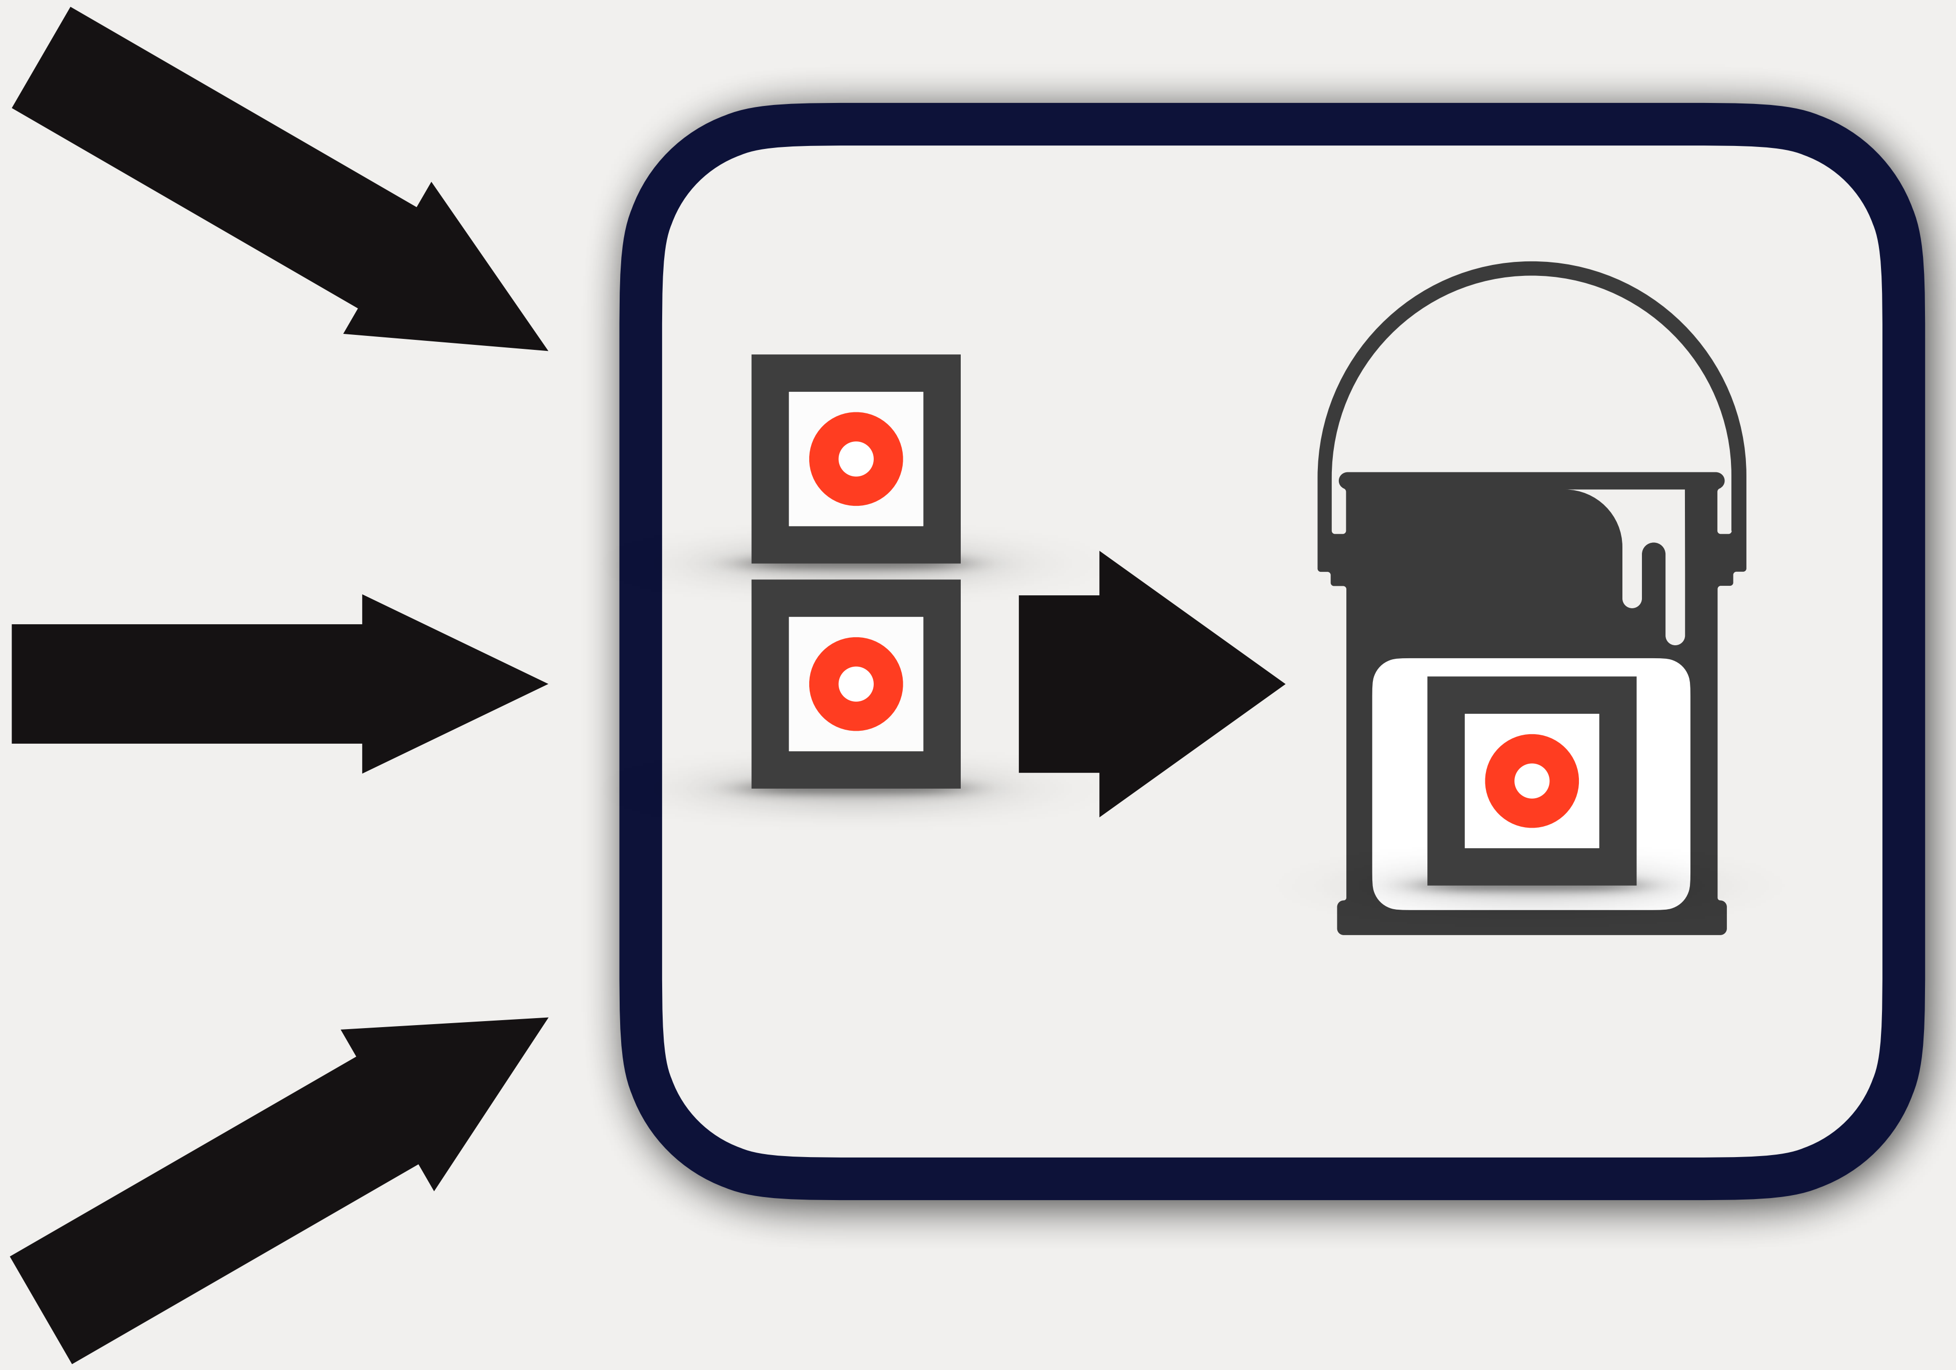

Unbiased Reuse

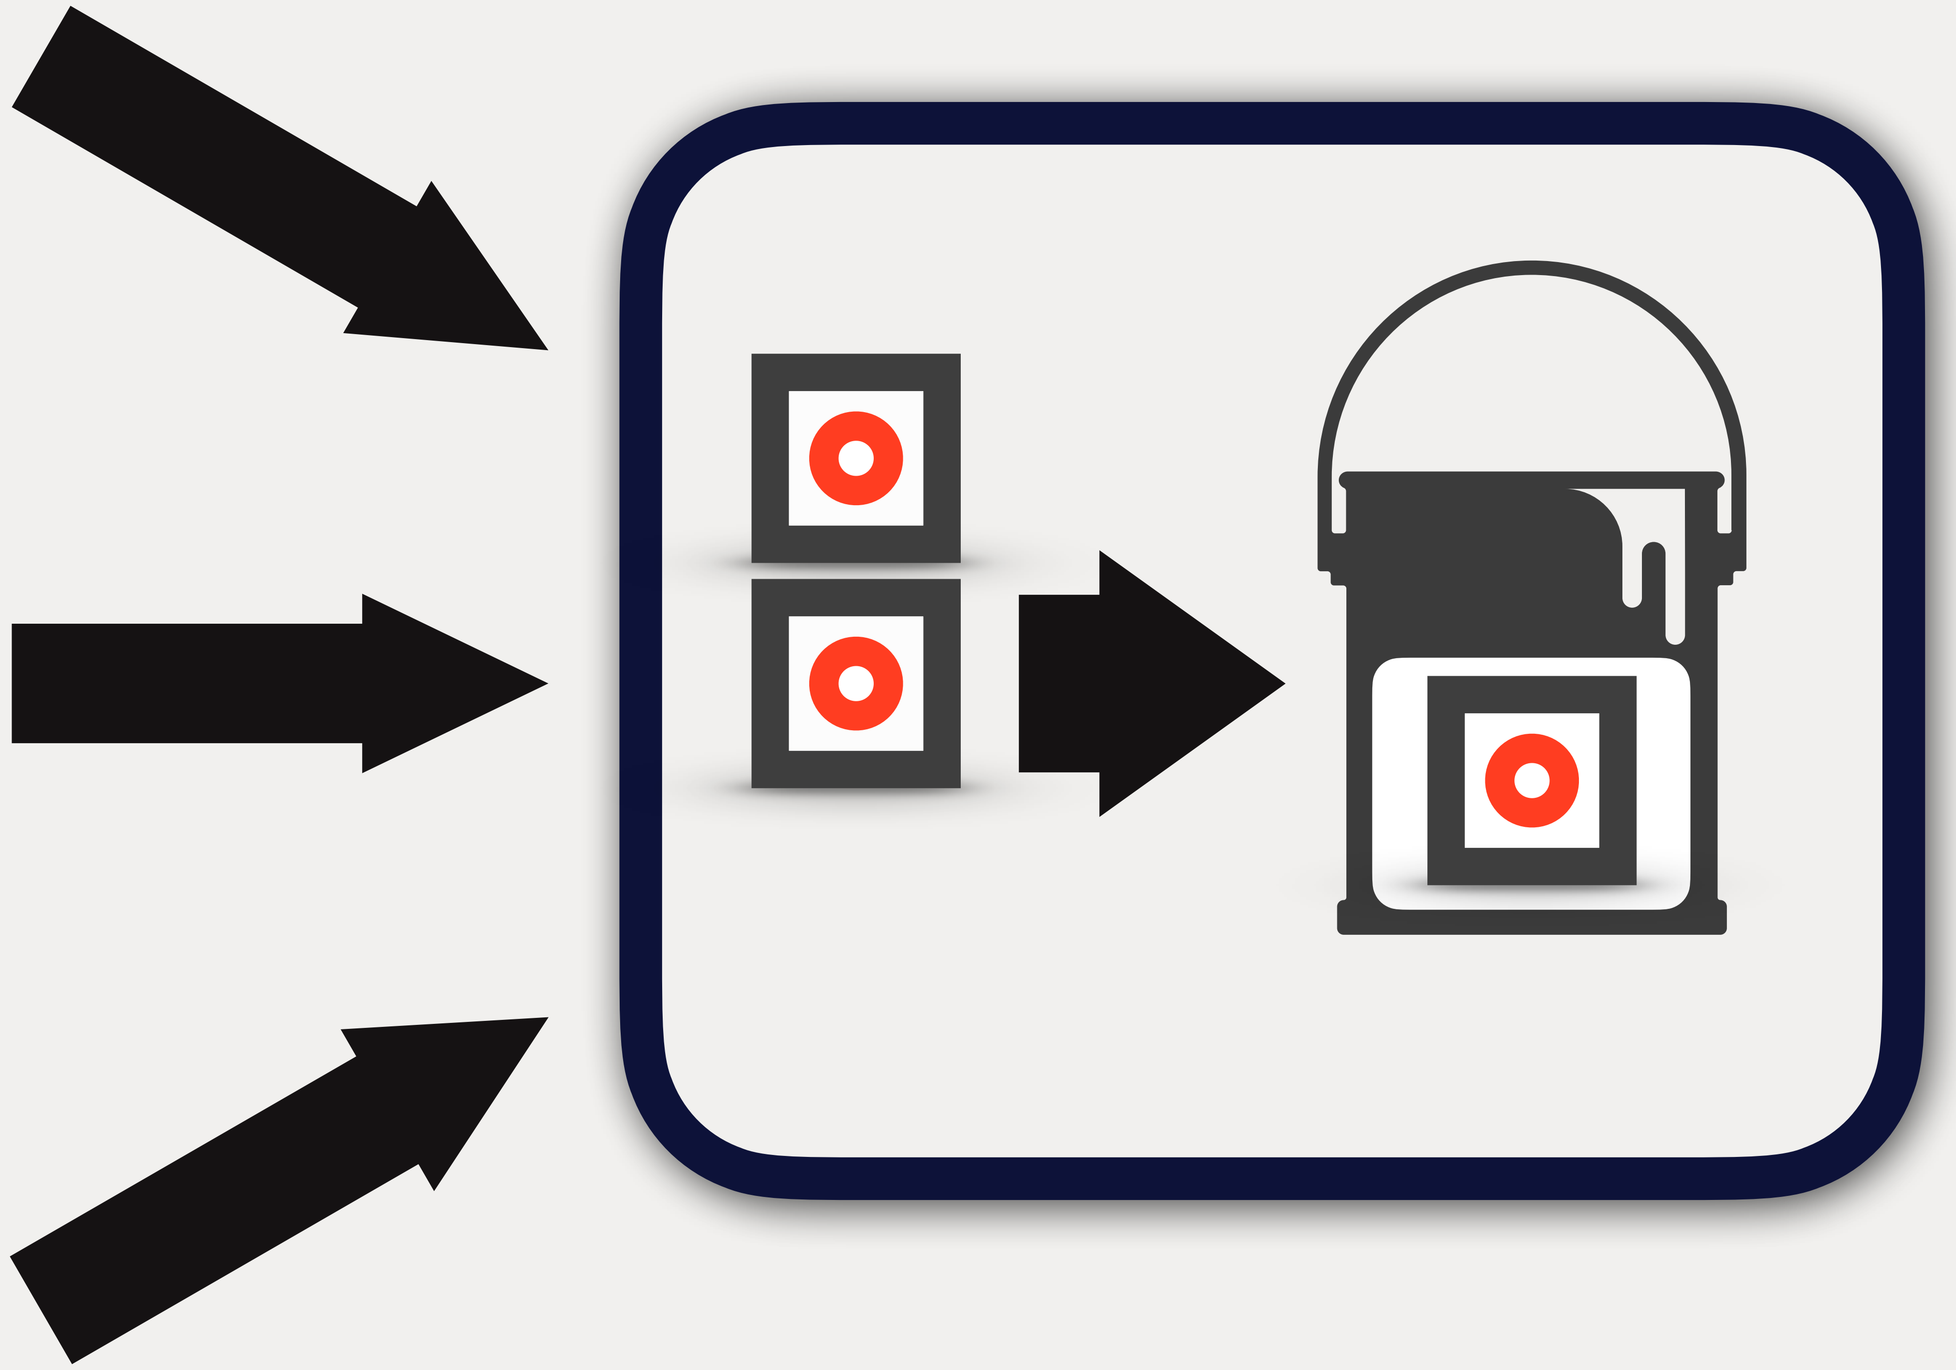

$$w(x_i) = \frac{\hat{p}(x_i)}{p(x_i)}$$

Unbiased Reuse

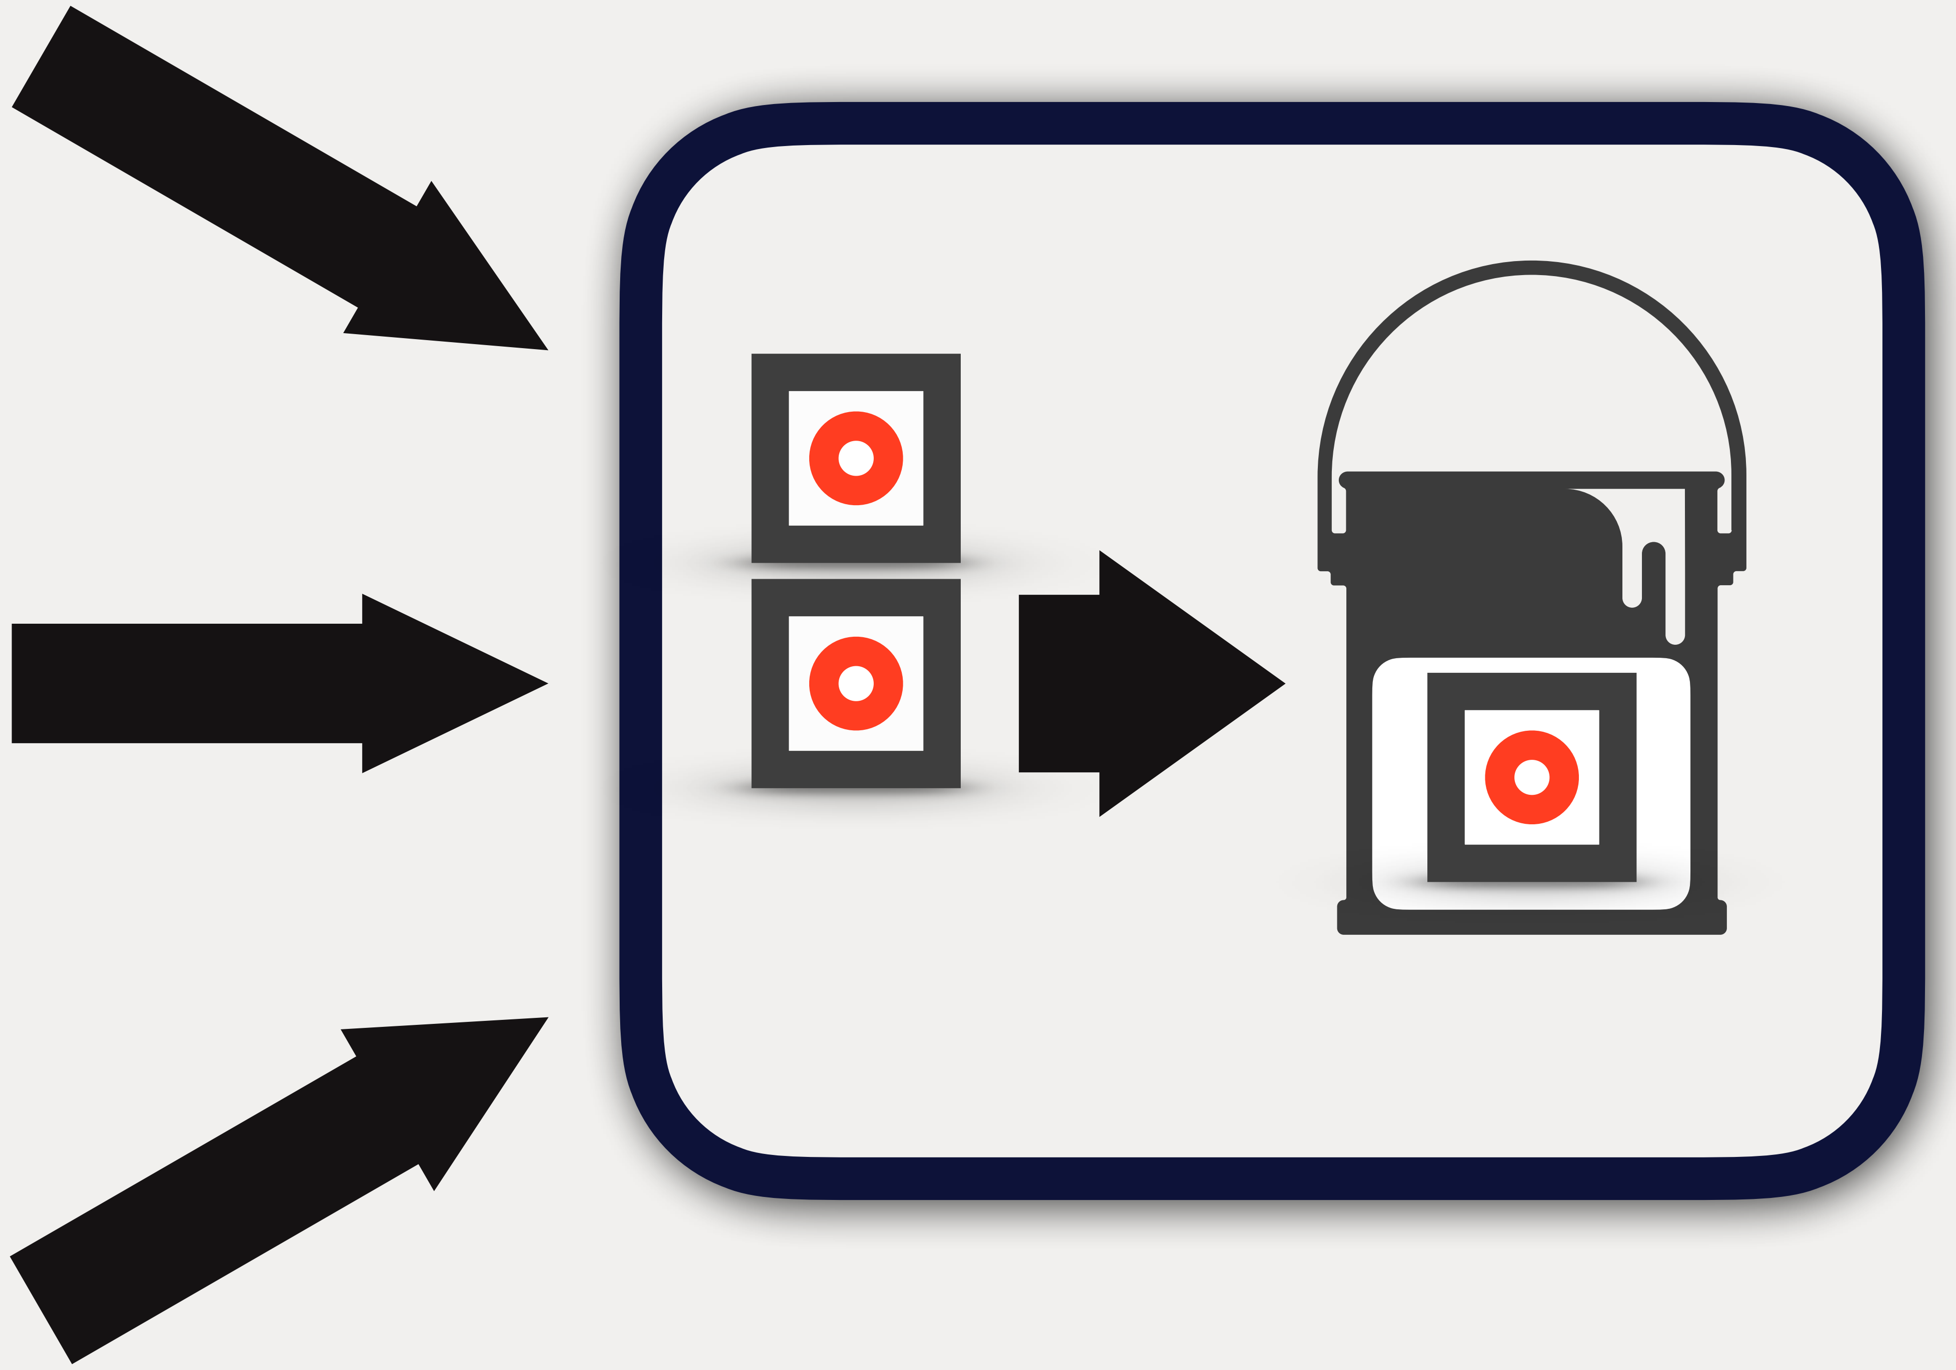

$$w(x_i) = \frac{\hat{p}(x_i)}{p_{\mathrm{RIS}}(x_i)}$$

Unbiased Reuse

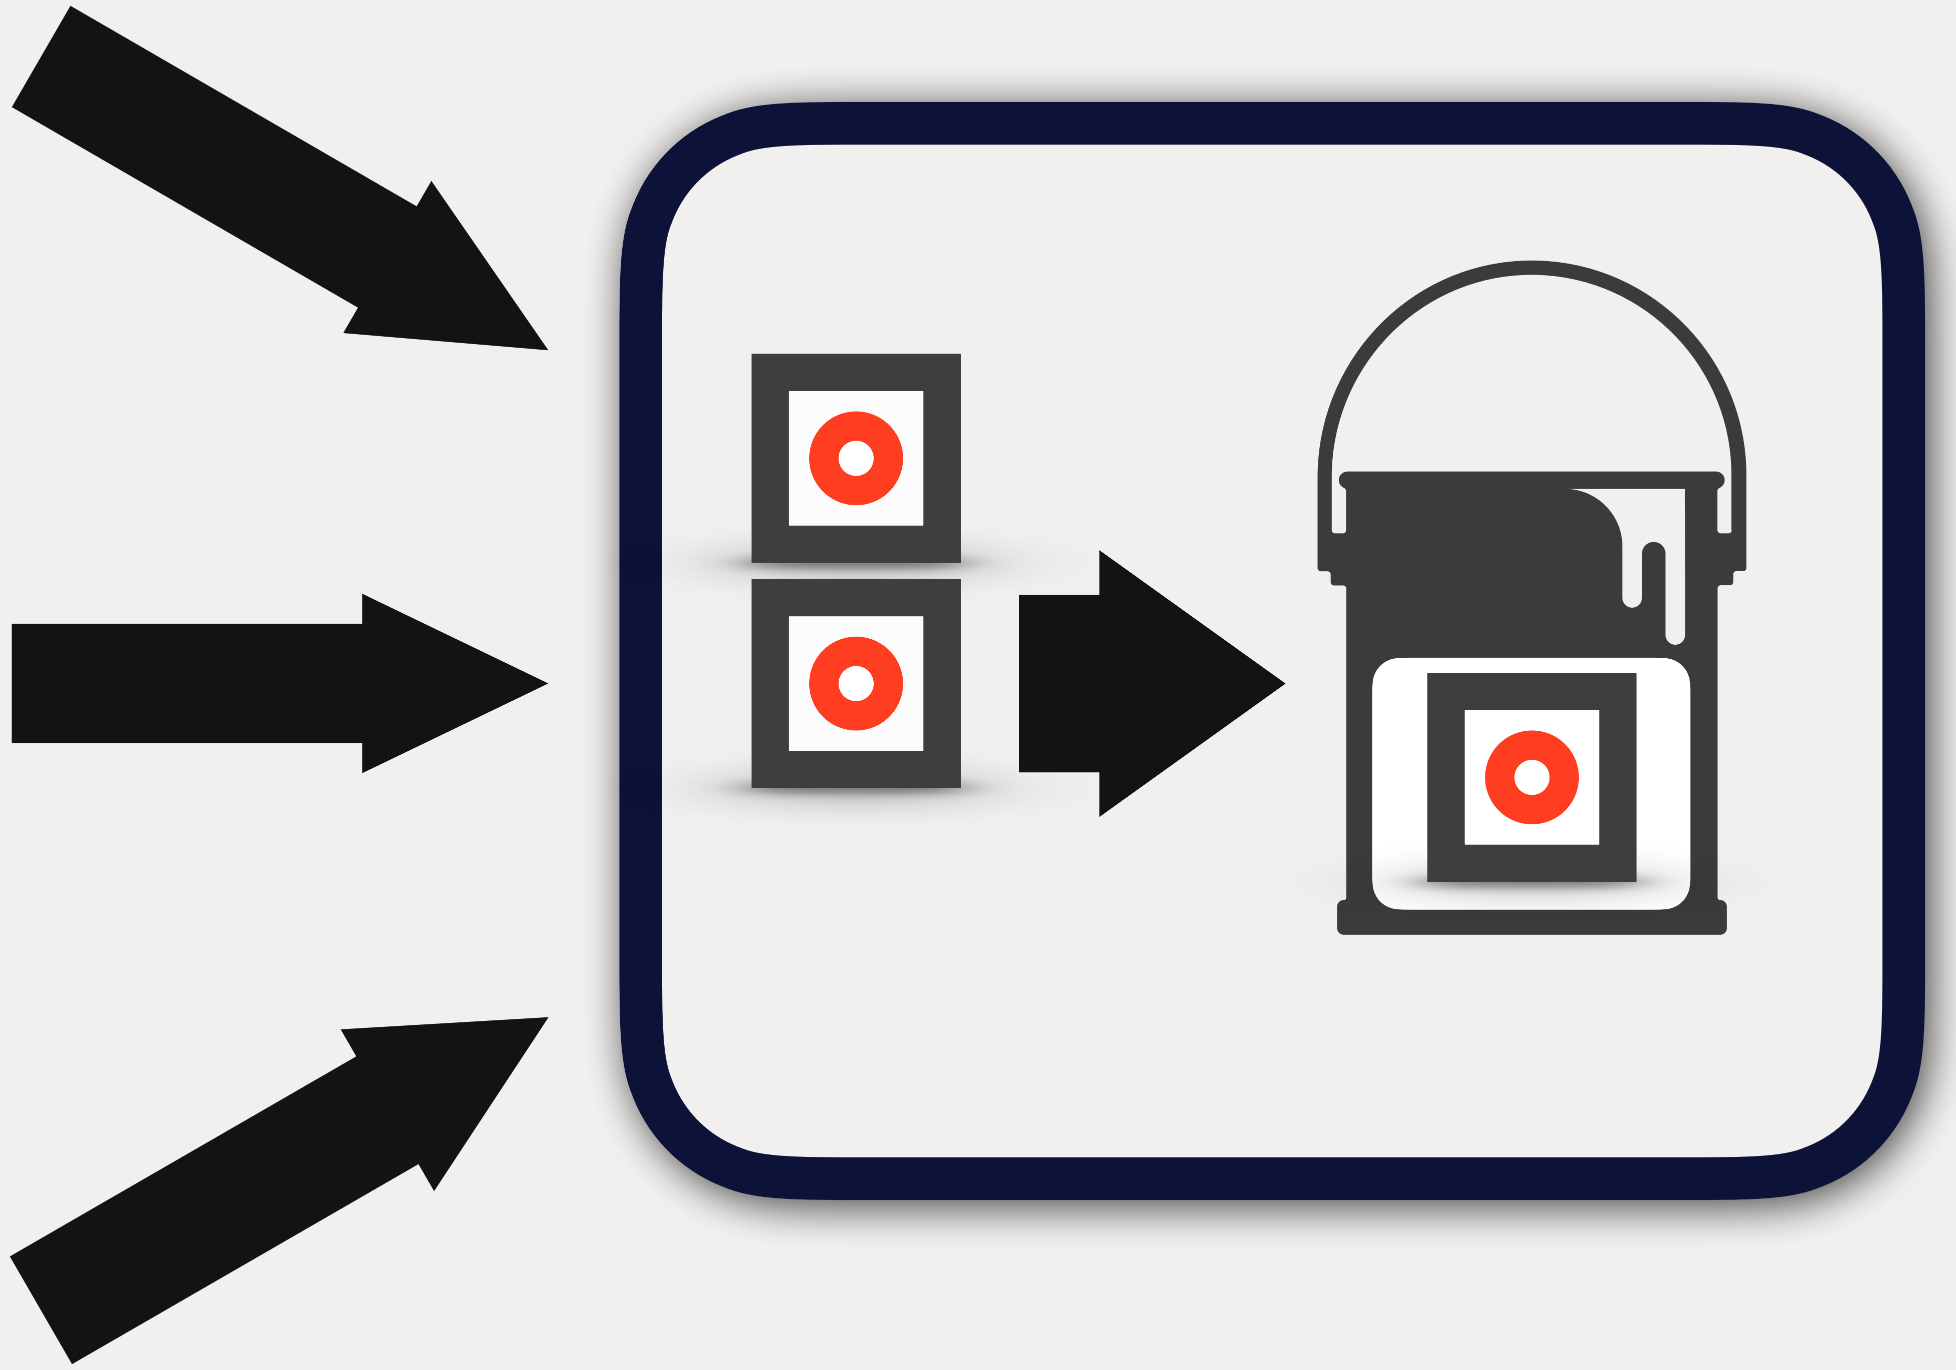

$$w(x_i) = \frac{\hat{p}(x_i)}{\boxed{p_{\mathrm{RIS}}(x_i)}}$$

Unbiased Reuse

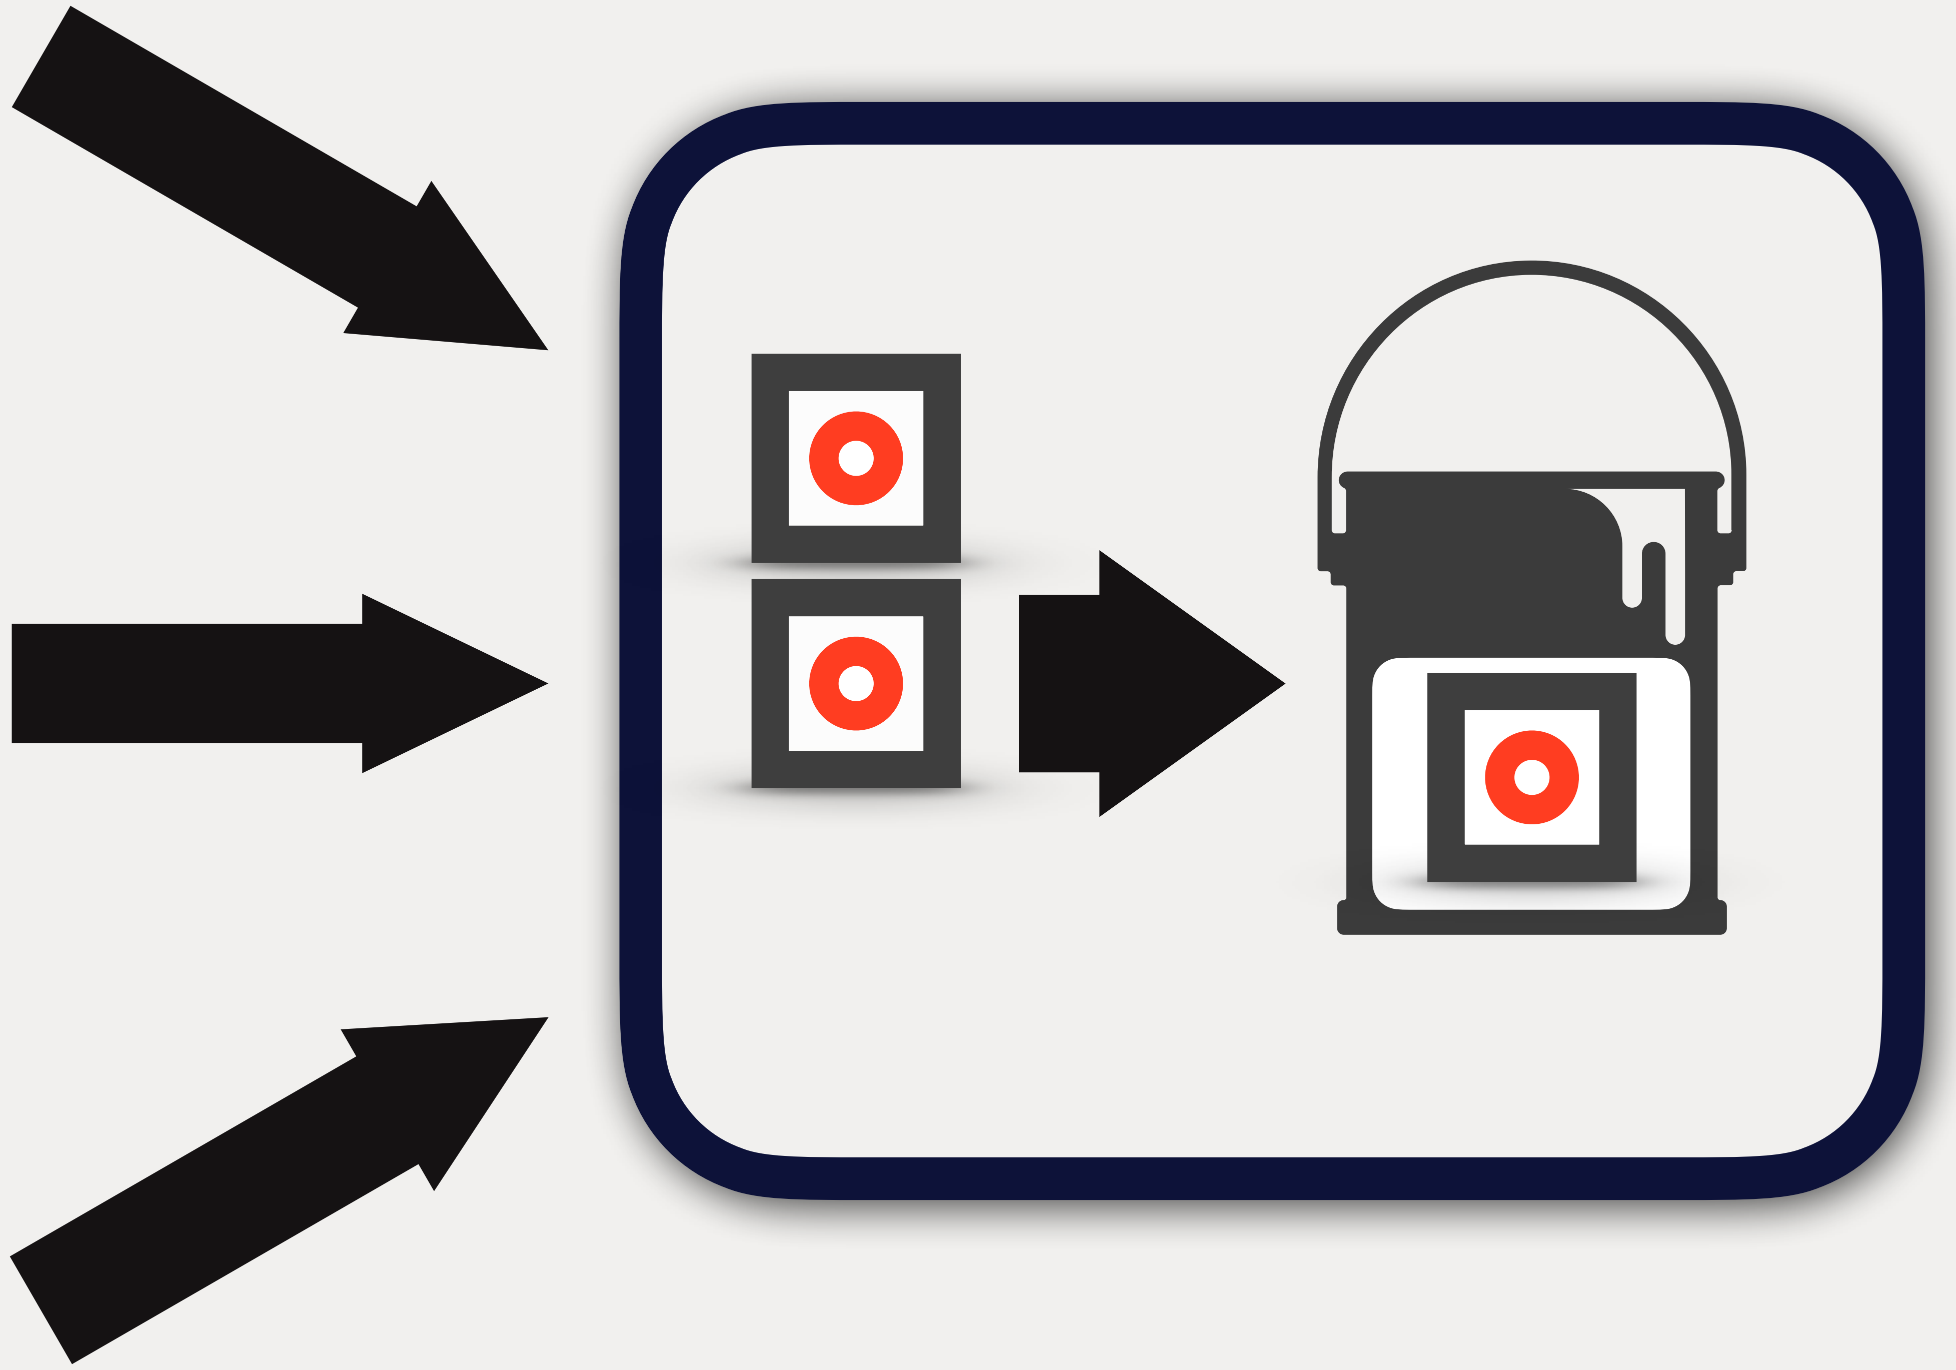

$$w(x_i) = \frac{\hat{p}(x_i)}{\boxed{p_{\mathrm{RIS}}(x_i)}}$$

Unbiased Reuse

$$p_{\mathrm{RIS}}(x)$$

Unbiased Reuse

$$p_{\mathrm{RIS}}(x) = \textbf{?}$$

Unbiased Reuse

$$p_{\mathrm{RIS}}(x) = \sum_{i=1}^{M} \underbrace{\int \cdots \int}_{M-1 \text{ times}} \left[\prod_{j=1}^{M} p(x_j) \right] \frac{w(x_i)}{\sum_{j=1}^{M} w(x_j)} \underbrace{\mathrm{d}x_1 \ldots \mathrm{d}x_M}_{M-1 \text{ times}}$$

Importance Resampling

[Talbot et al. '05]

$$\int_{\Omega} f(x)\mathrm{d}x \approx f(y) \cdot W$$

$$W = \frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M} \sum_{i=1}^{M} w(x_i) \right)$$

Importance Resampling

[Talbot et al. '05]

$$\int_{\Omega} f(x)\mathrm{d}x \approx f(y) \cdot W$$

$$\boxed{W = \frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M} \sum_{i=1}^{M} w(x_i) \right)}$$

Importance Resampling

[Talbot et al. ’05]

$$\int_\Omega f(x)\mathrm{d}x \approx f(y) \cdot W$$

$$\mathbb{E}[W] = \frac{1}{p_{\mathrm{RIS}}(y)}$$

Unbiased Reuse

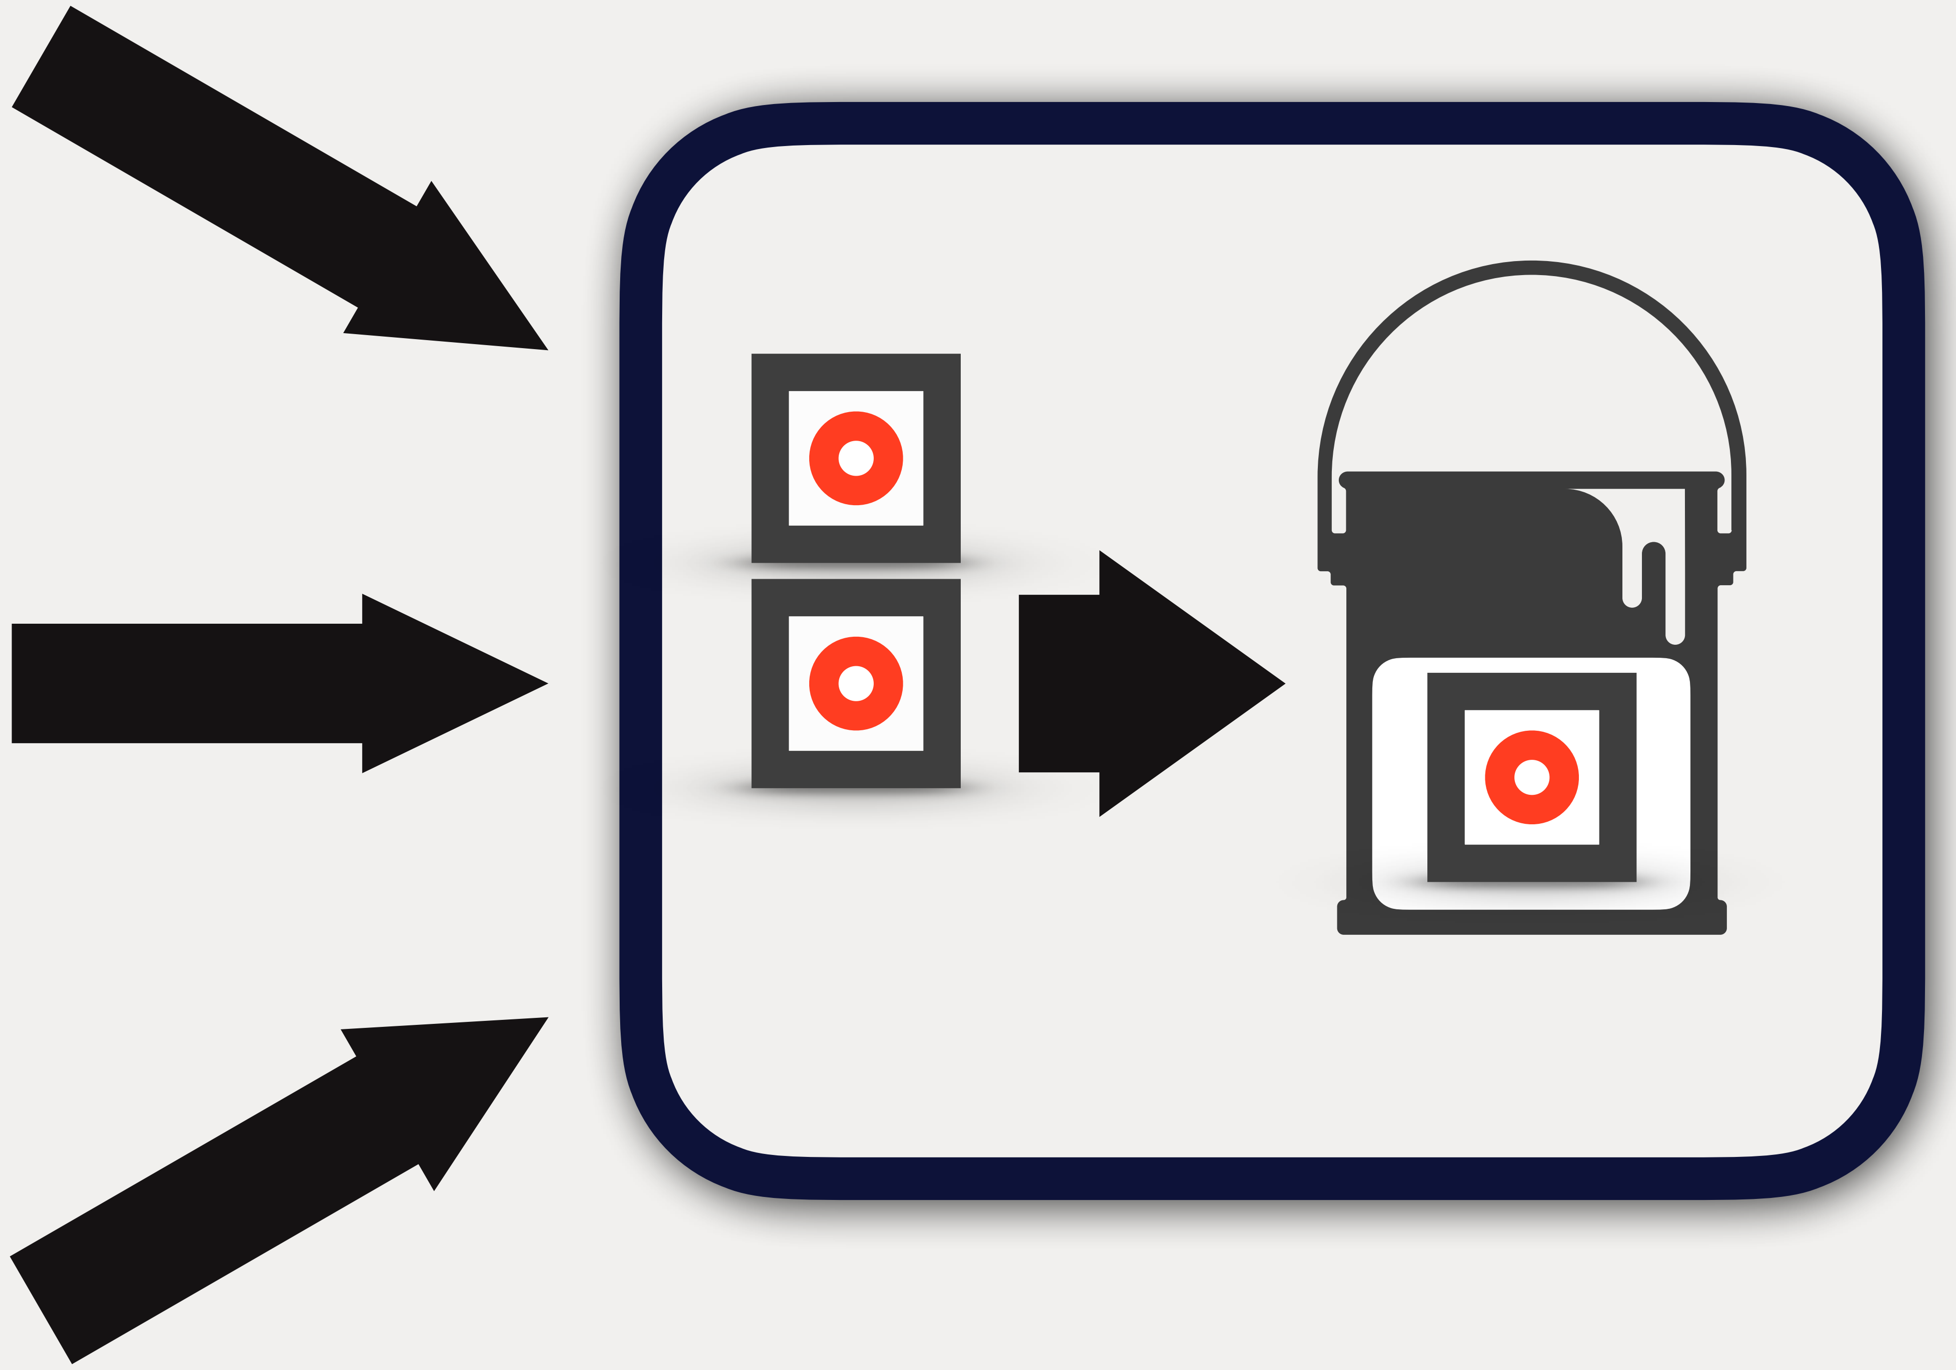

$$w(x_i) = \frac{\hat{p}(x_i)}{p_{\mathrm{RIS}}(x_i)}$$

Unbiased Reuse

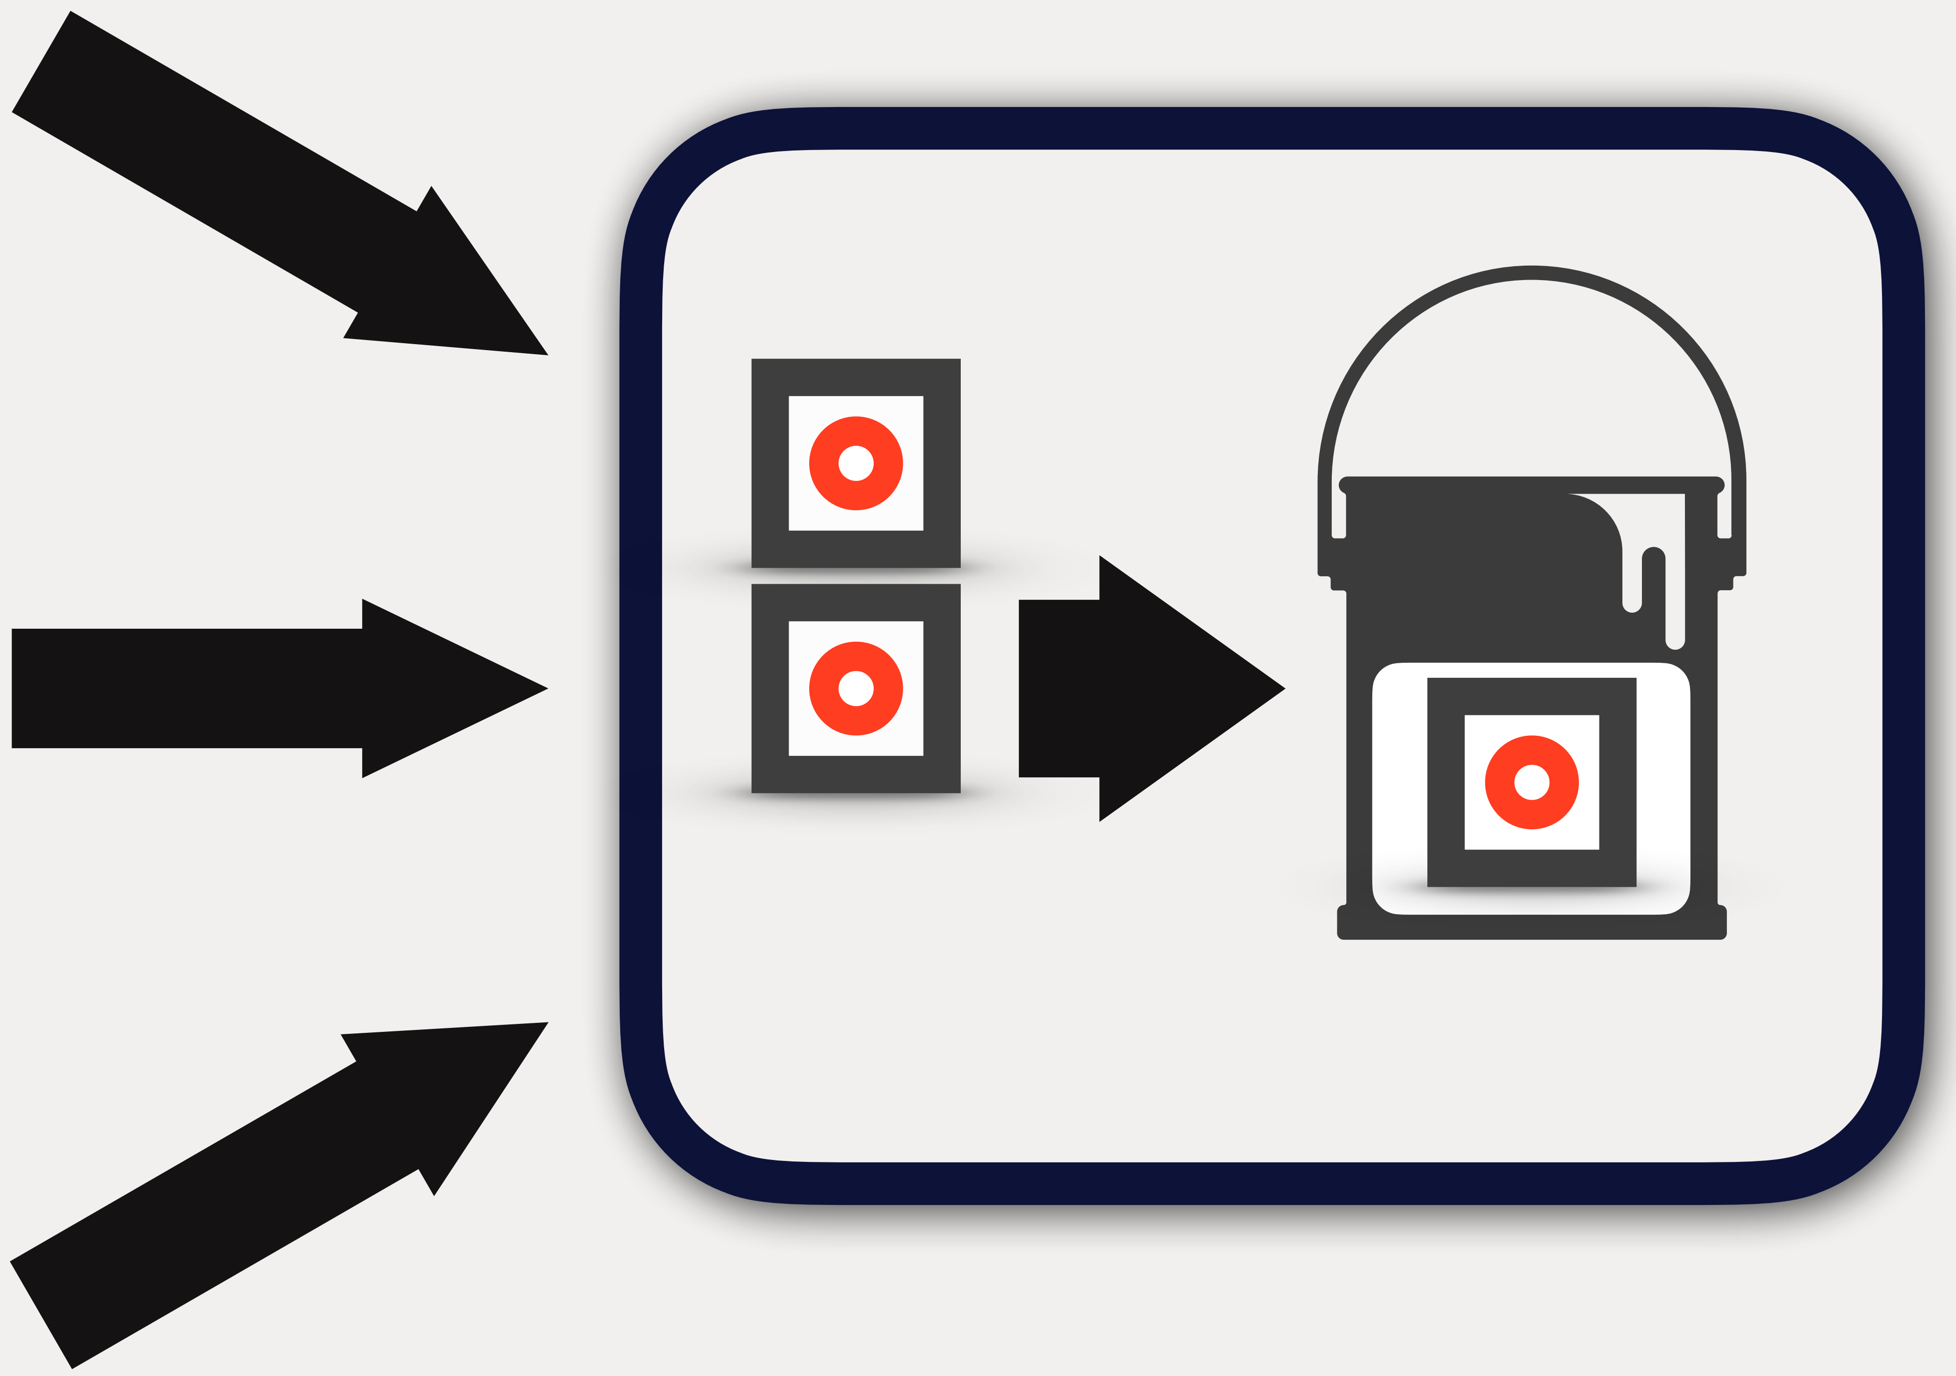

$$w(x_i) = \hat{p}(x_i) \cdot W_i$$

Unbiased Reuse

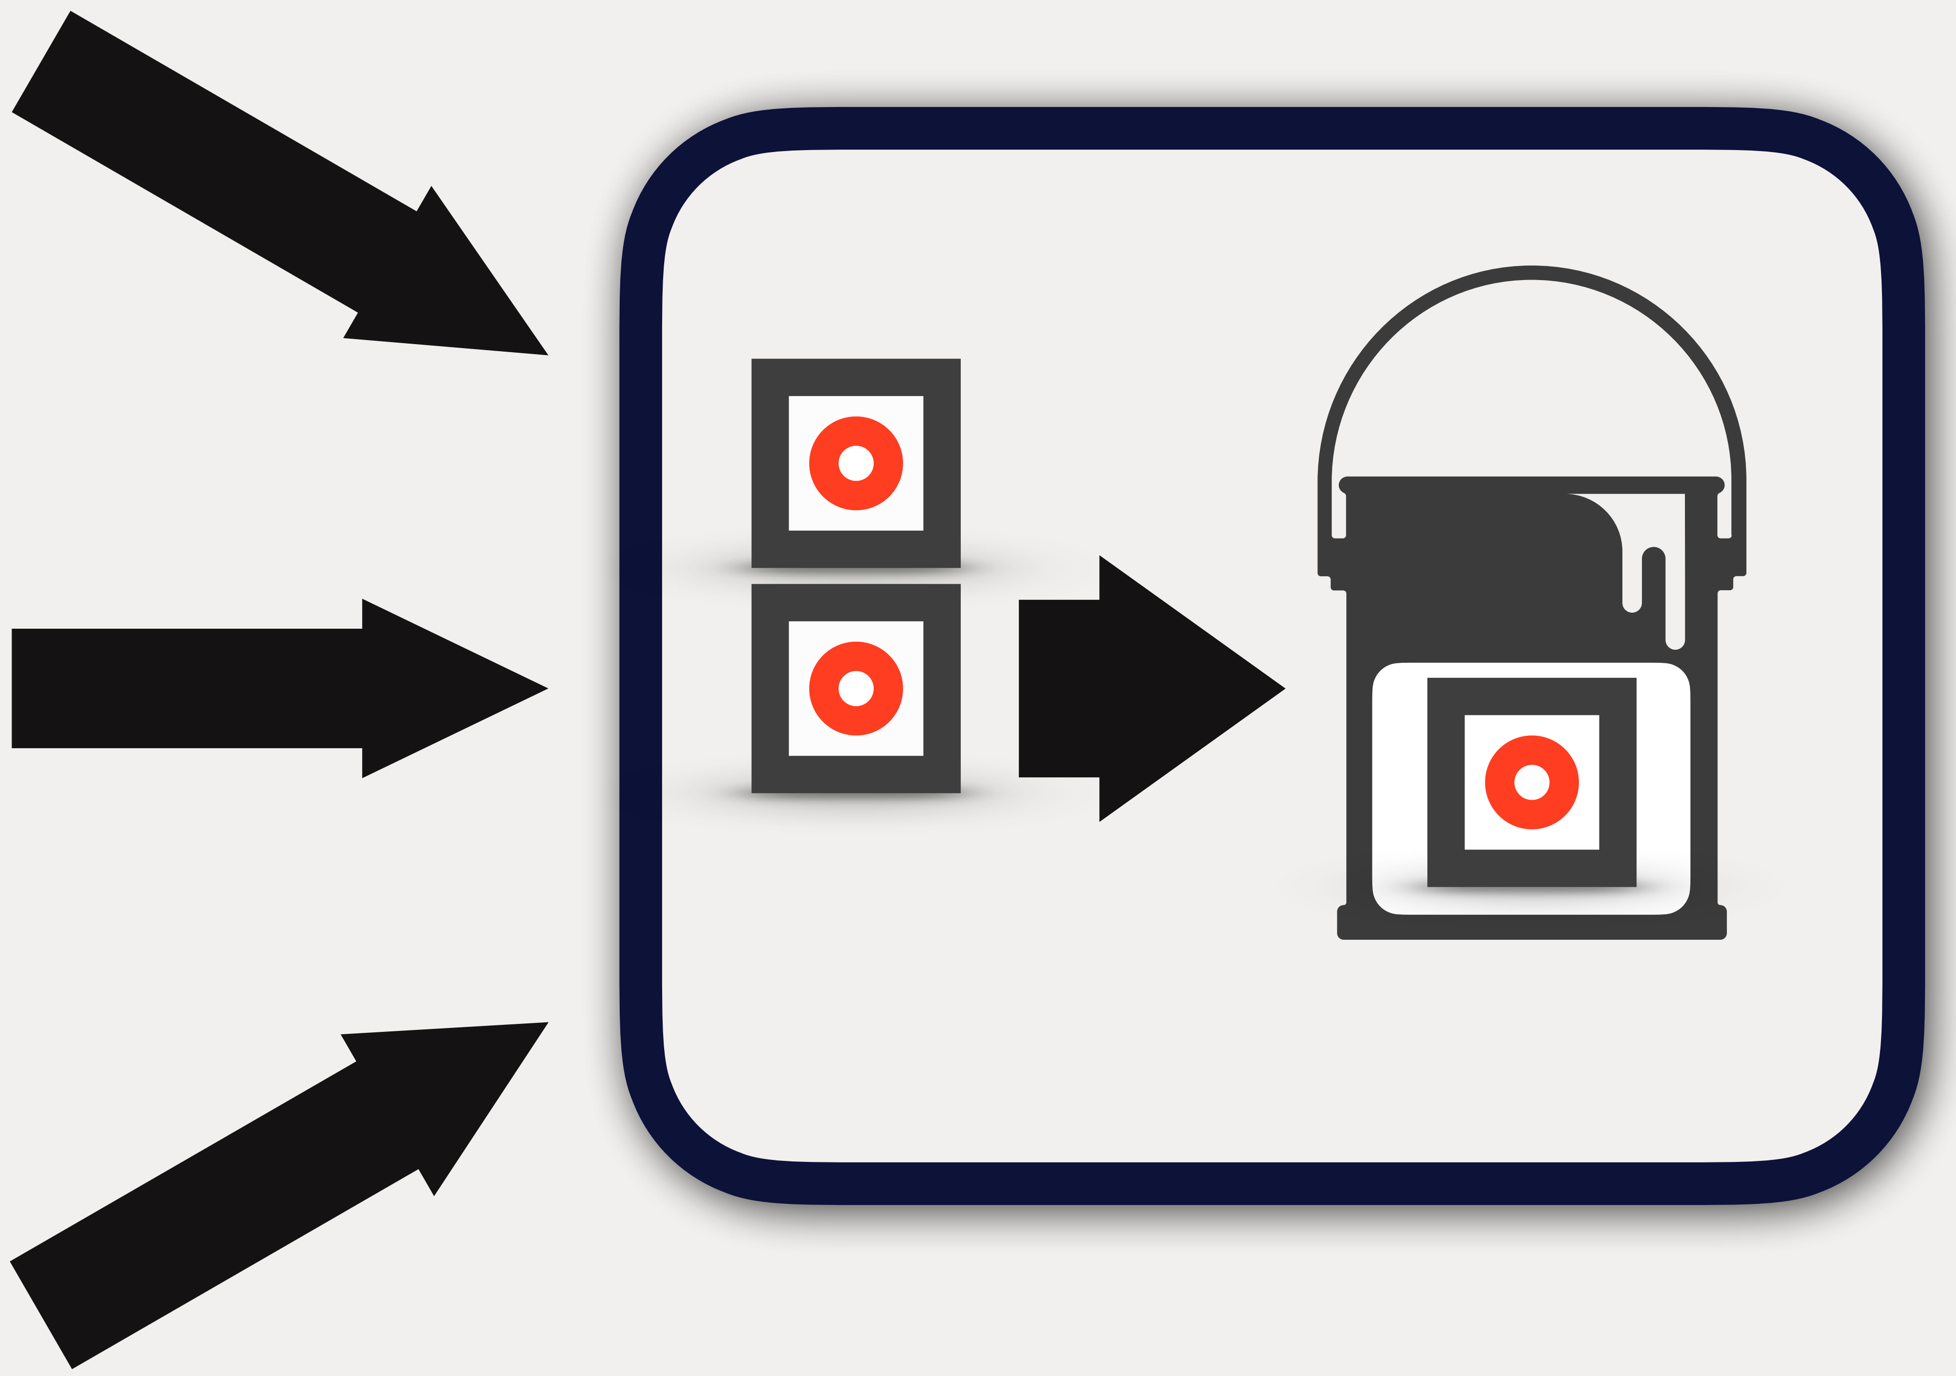

Unbiased Reuse

Resampling unbiased if $p(x) > 0$ where $\hat{p}(x) > 0$

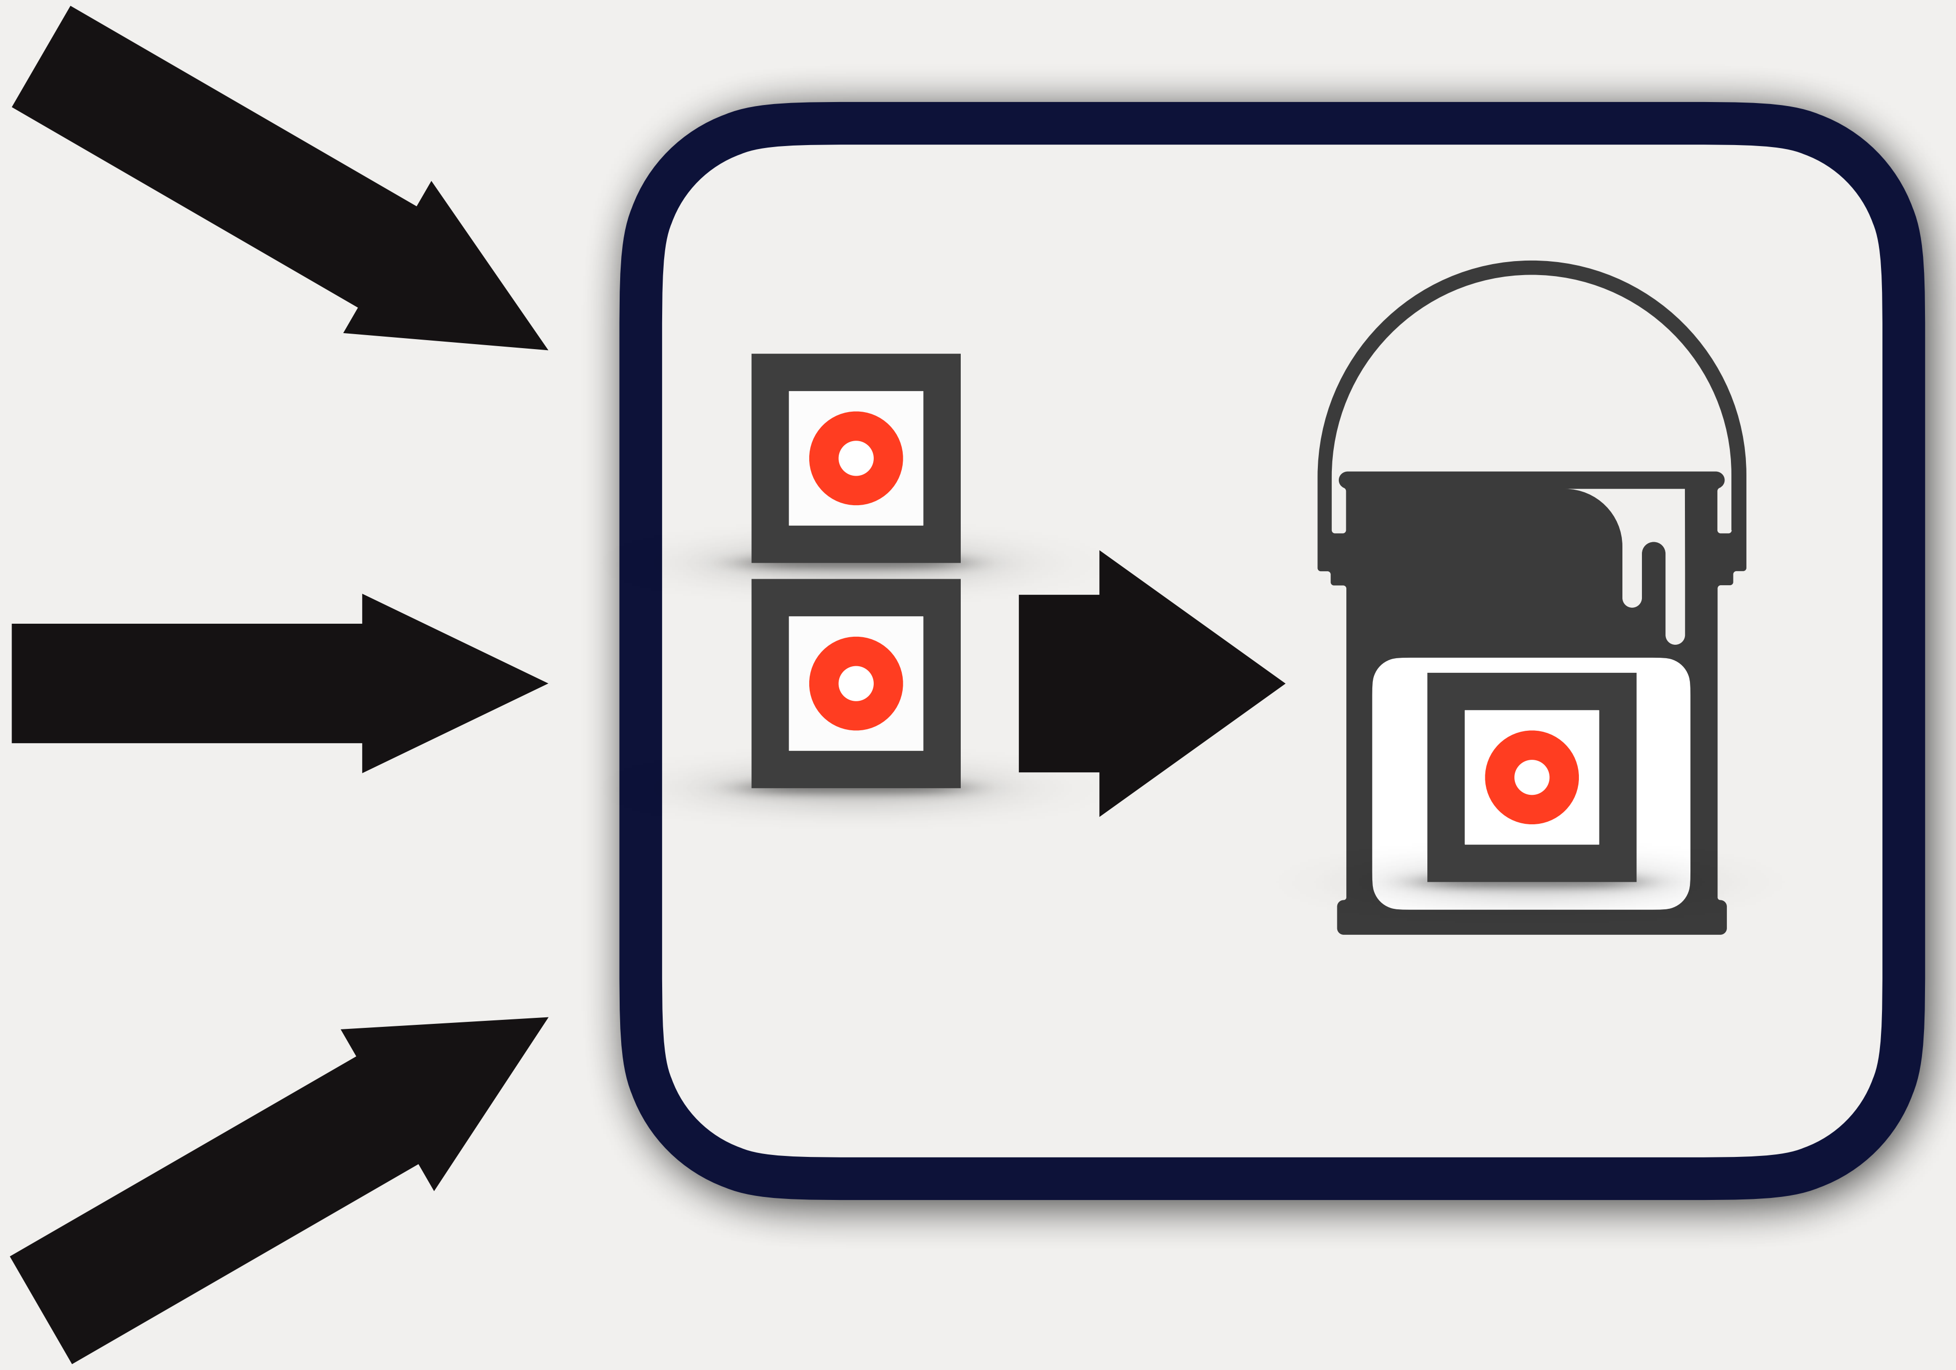

Unbiased Reuse

Resampling unbiased if $p(x) > 0$ where $\hat{p}(x) > 0$

Unbiased Reuse

Resampling unbiased if $p(x) > 0$ where $\hat{p}(x) > 0$

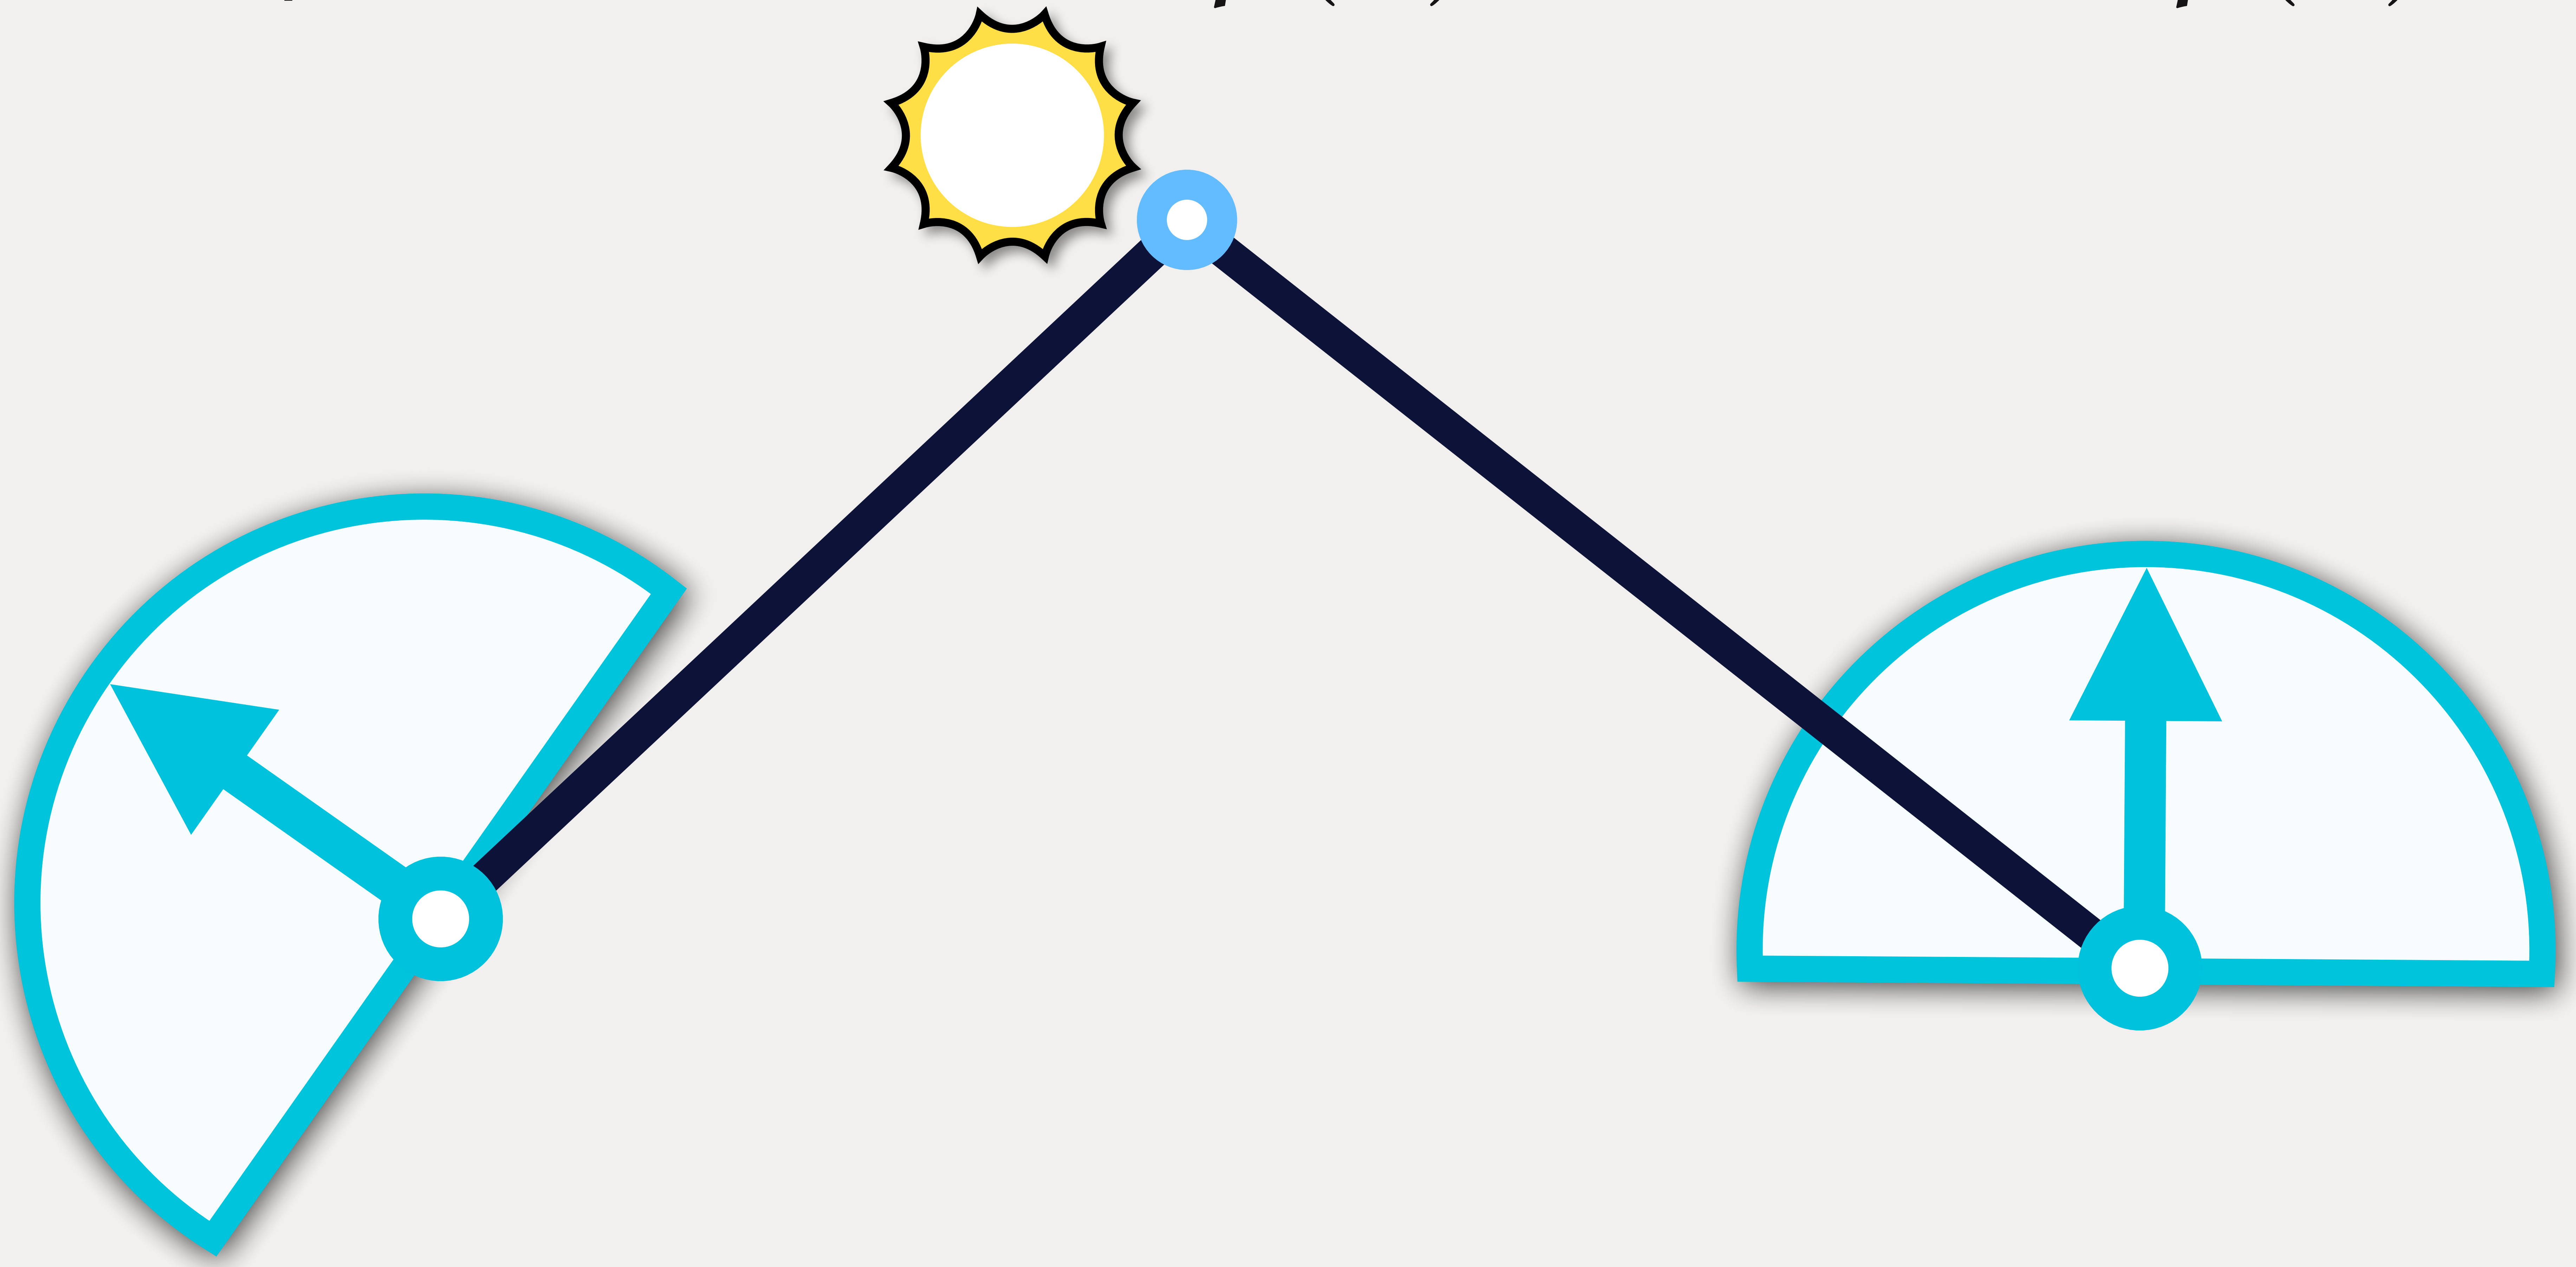

Unbiased Reuse

Resampling unbiased if $p(x) > 0$ where $\hat{p}(x) > 0$

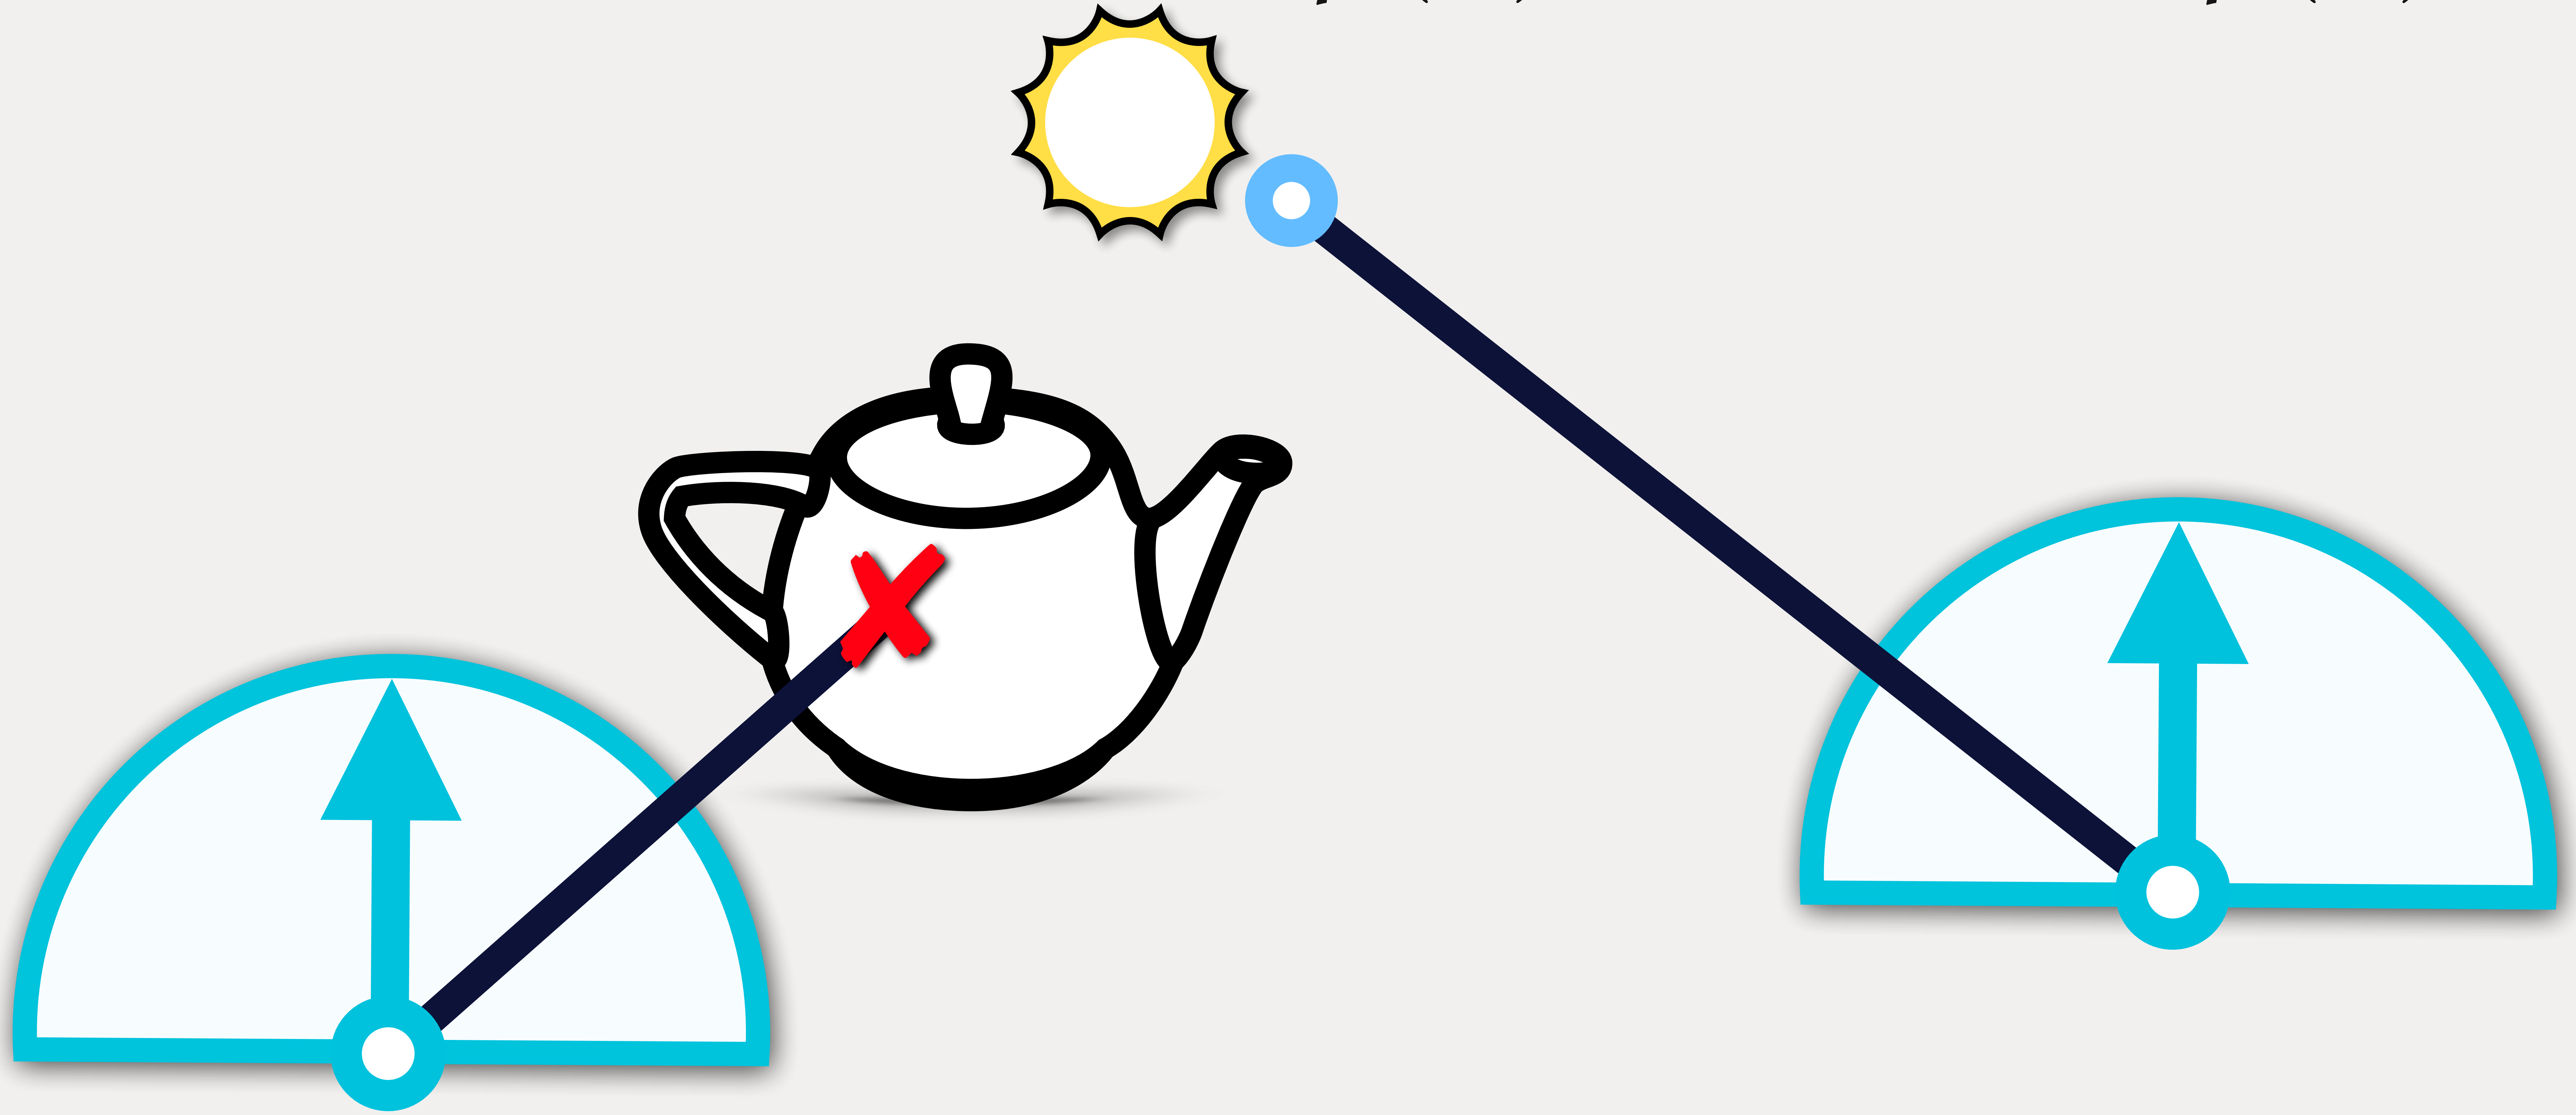

Unbiased Reuse

$$\mathbb{E}[W] \neq \frac{1}{p_{\mathrm{RIS}}(y)}$$

Unbiased Reuse

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M}\sum_{i=1}^{M} w(x_i)\right)\right] \neq \frac{1}{p_{\mathrm{RIS}}(y)}$$

Unbiased Reuse

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\boxed{\frac{1}{M}}\sum_{i=1}^{M} w(x_i)\right)\right] \neq \frac{1}{p_{\mathrm{RIS}}(y)}$$

Unbiased Reuse

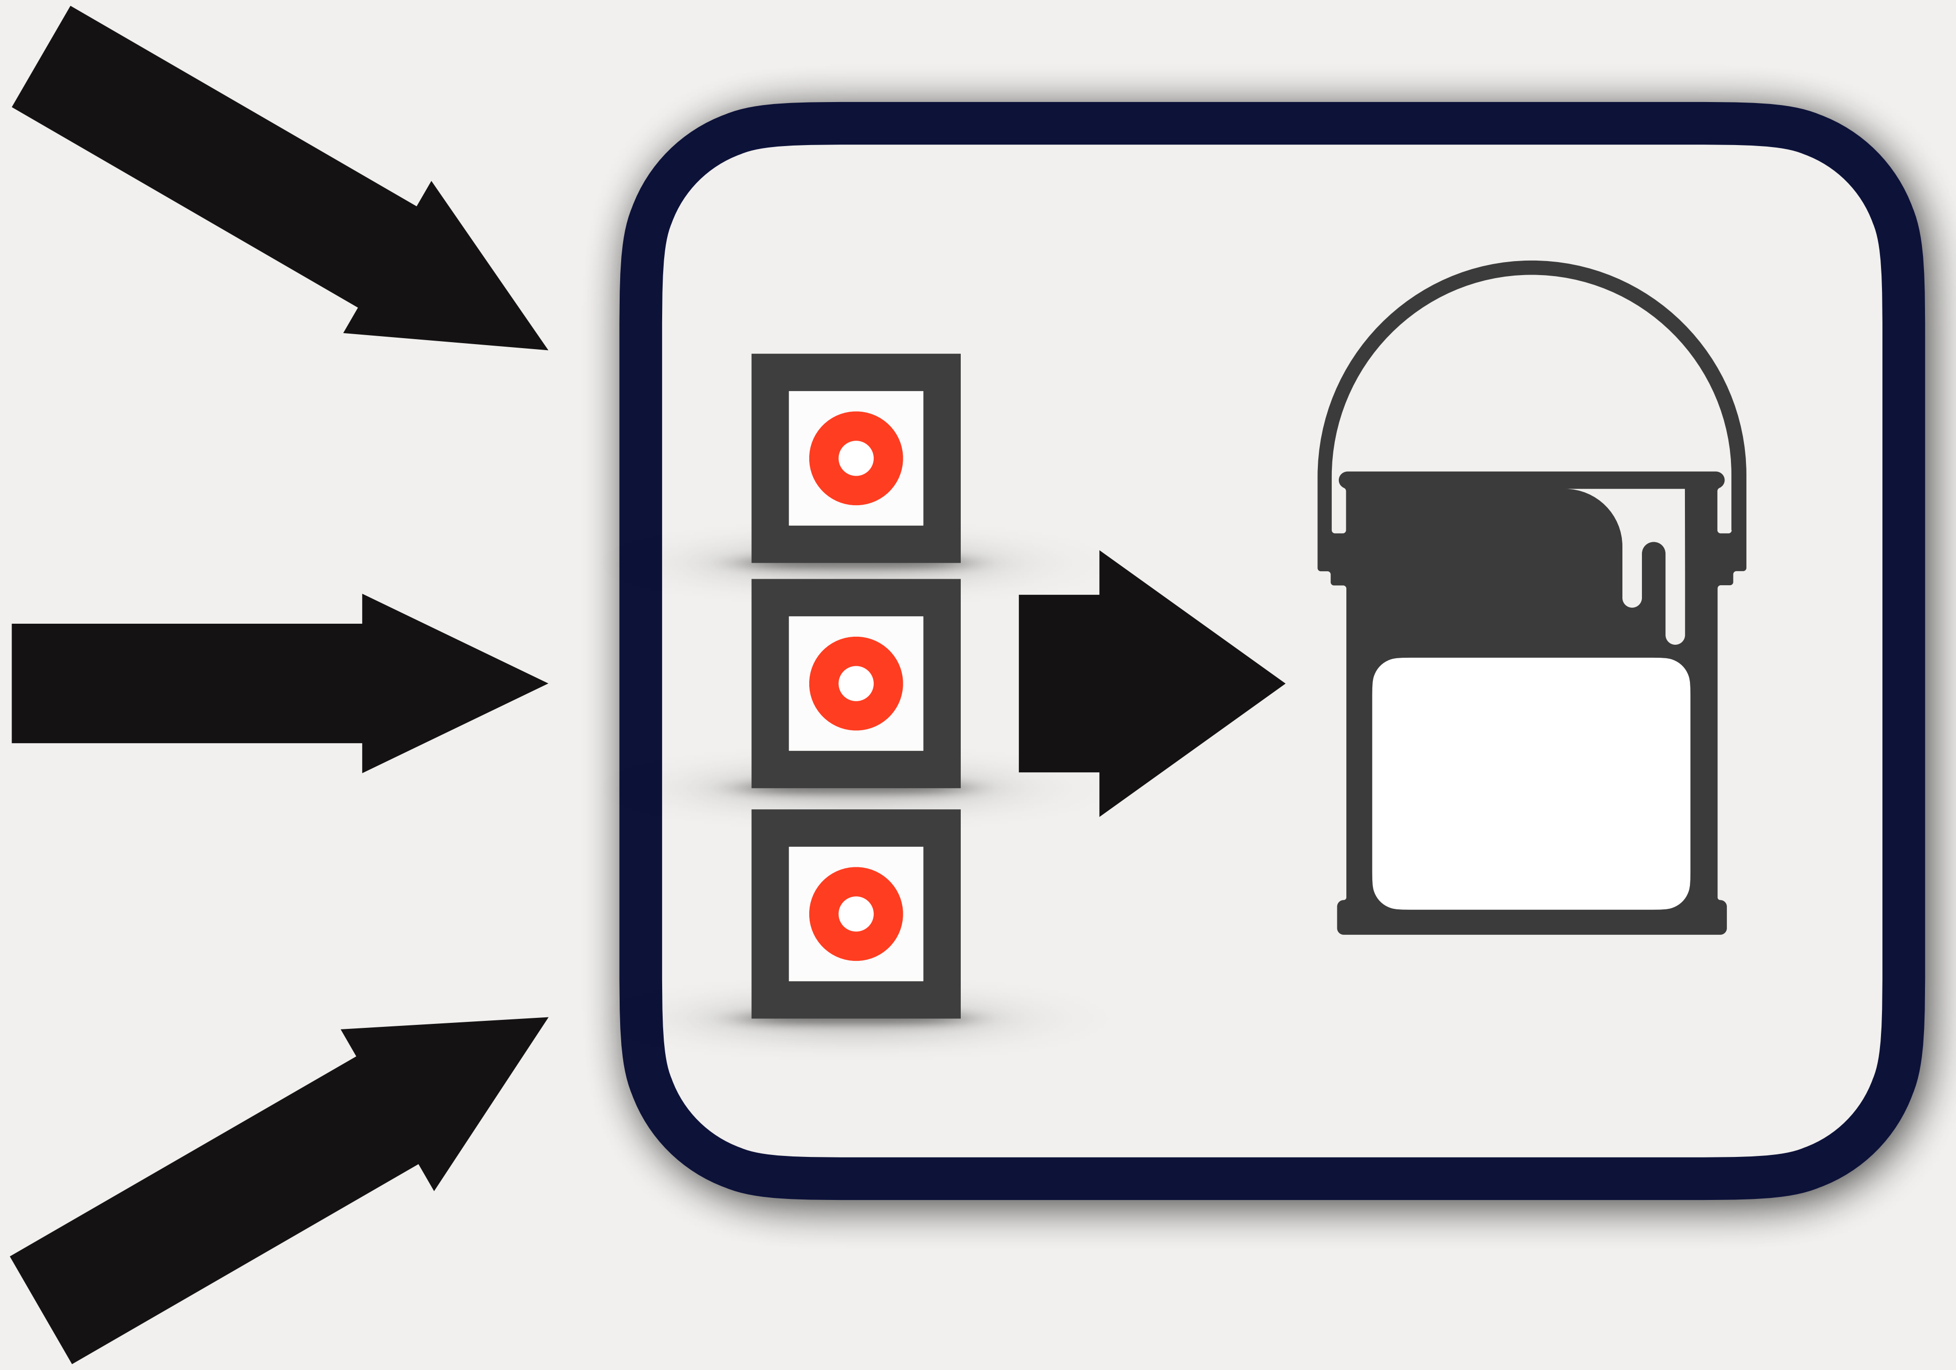

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)}\cdot\left(\frac{1}{M}\sum_{i=1}^{M} w(x_i)\right)\right]$$

Unbiased Reuse

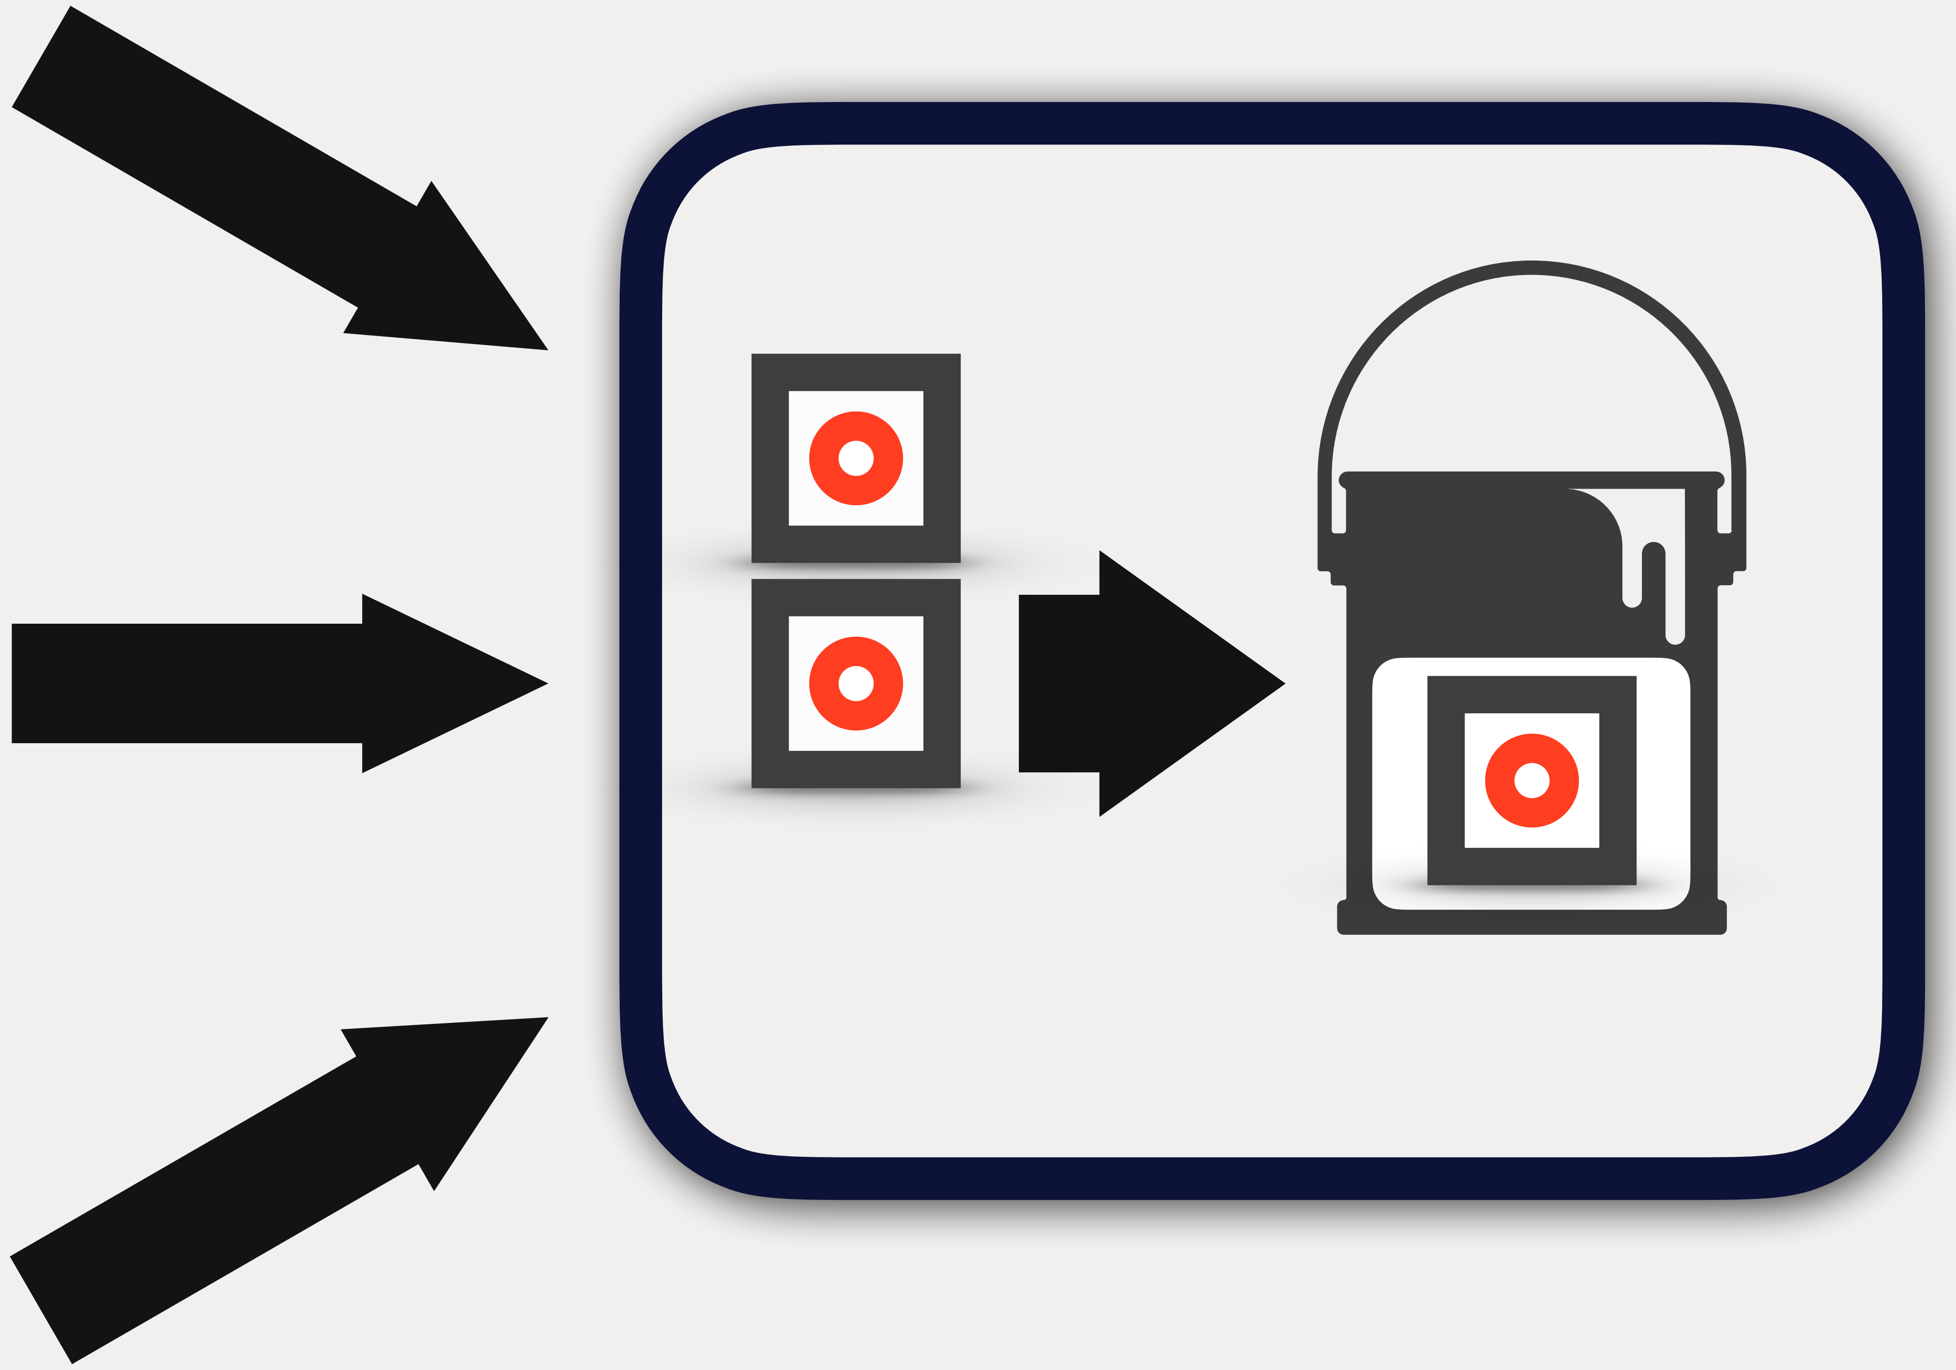

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M}\sum_{i=1}^{M} w(x_i)\right)\right]$$

Unbiased Reuse

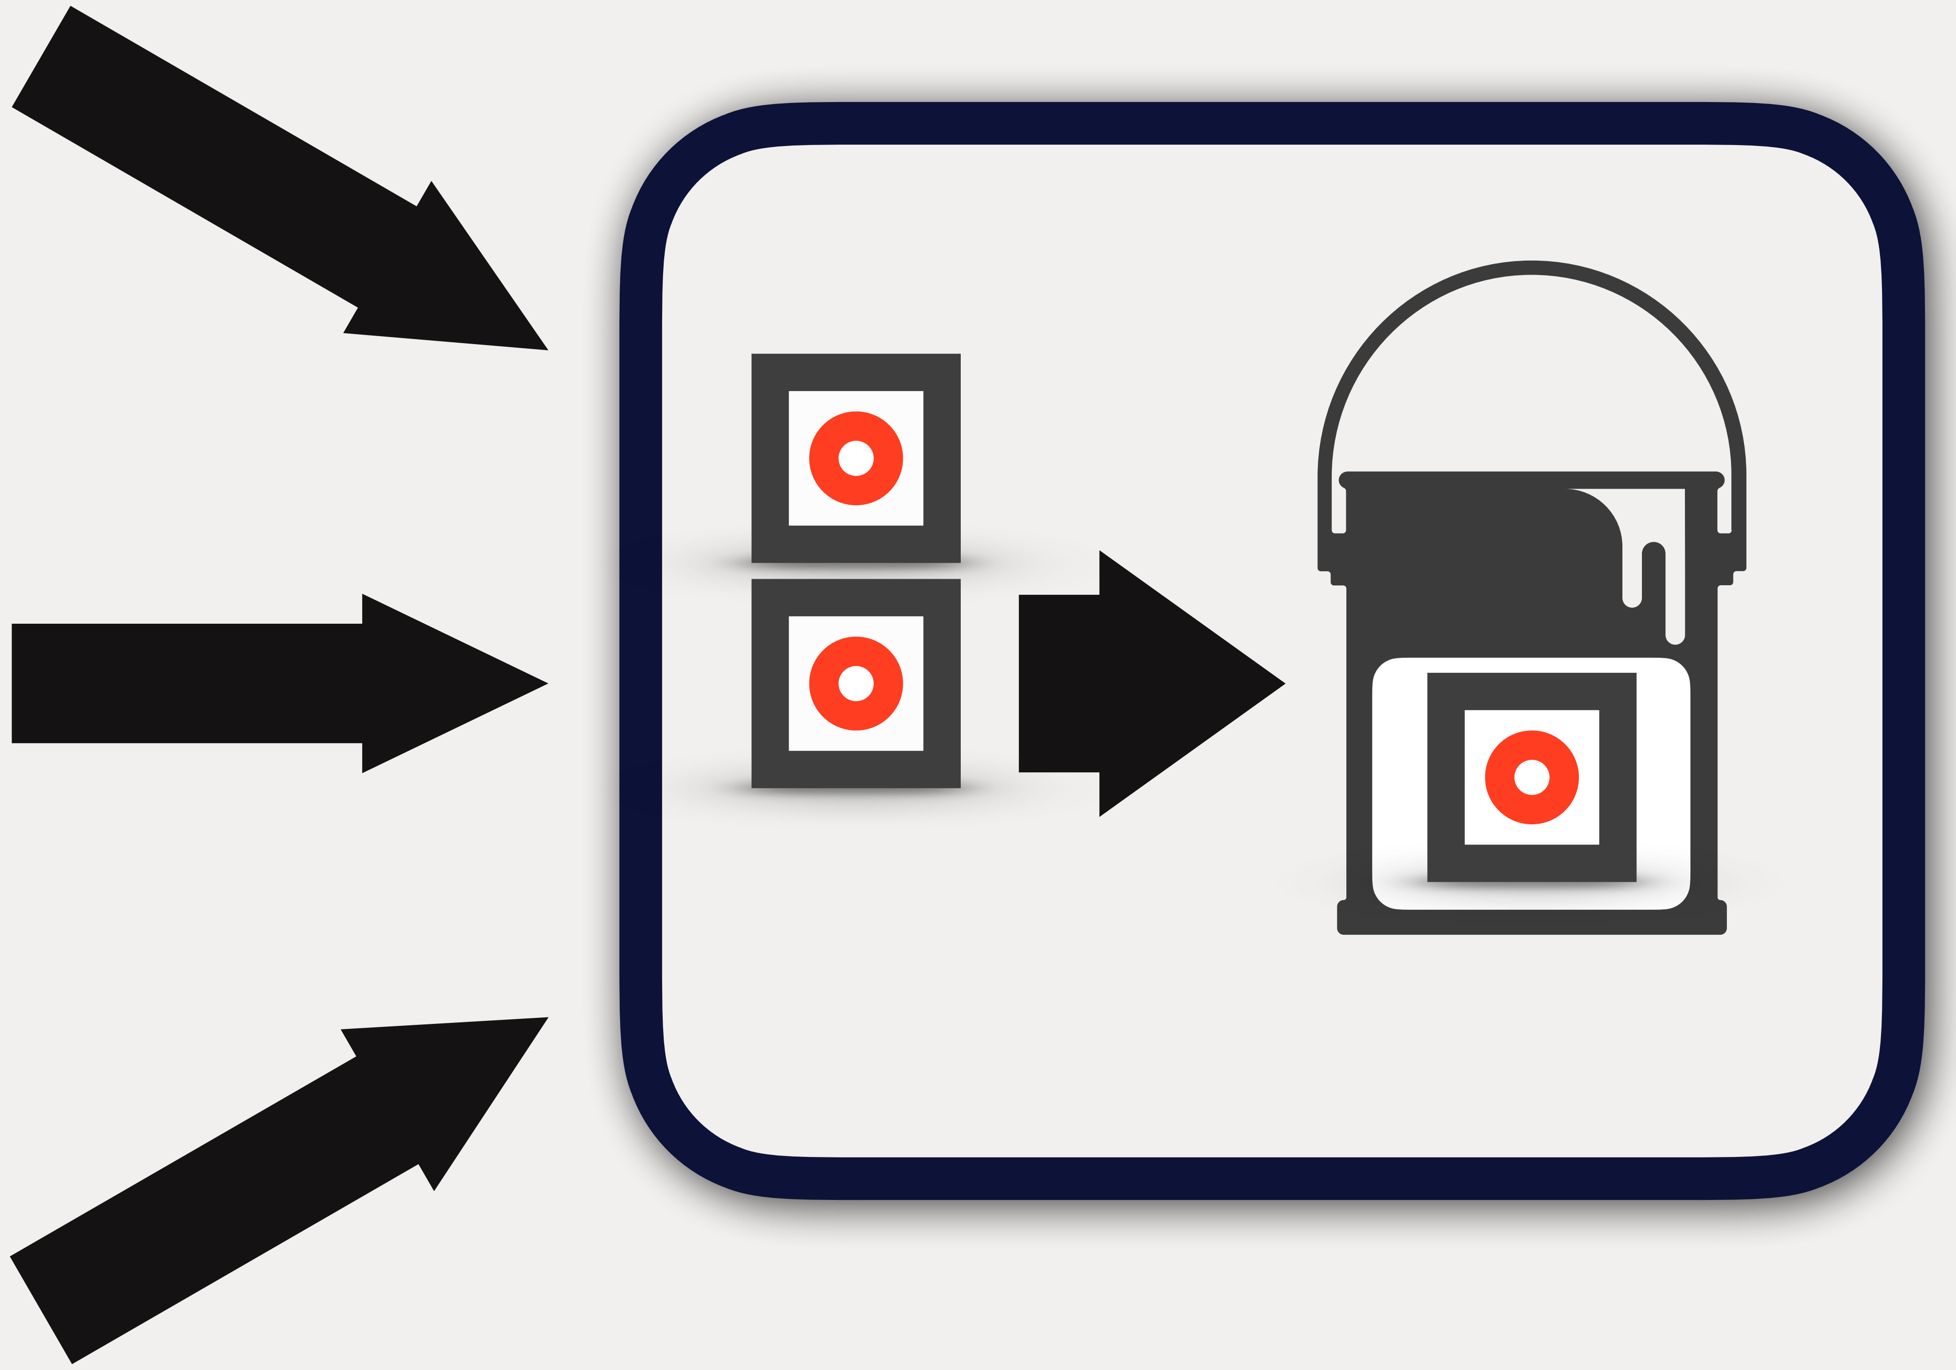

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M}\sum_{i=1}^{M} w(x_i)\right)\right]$$

Unbiased Reuse

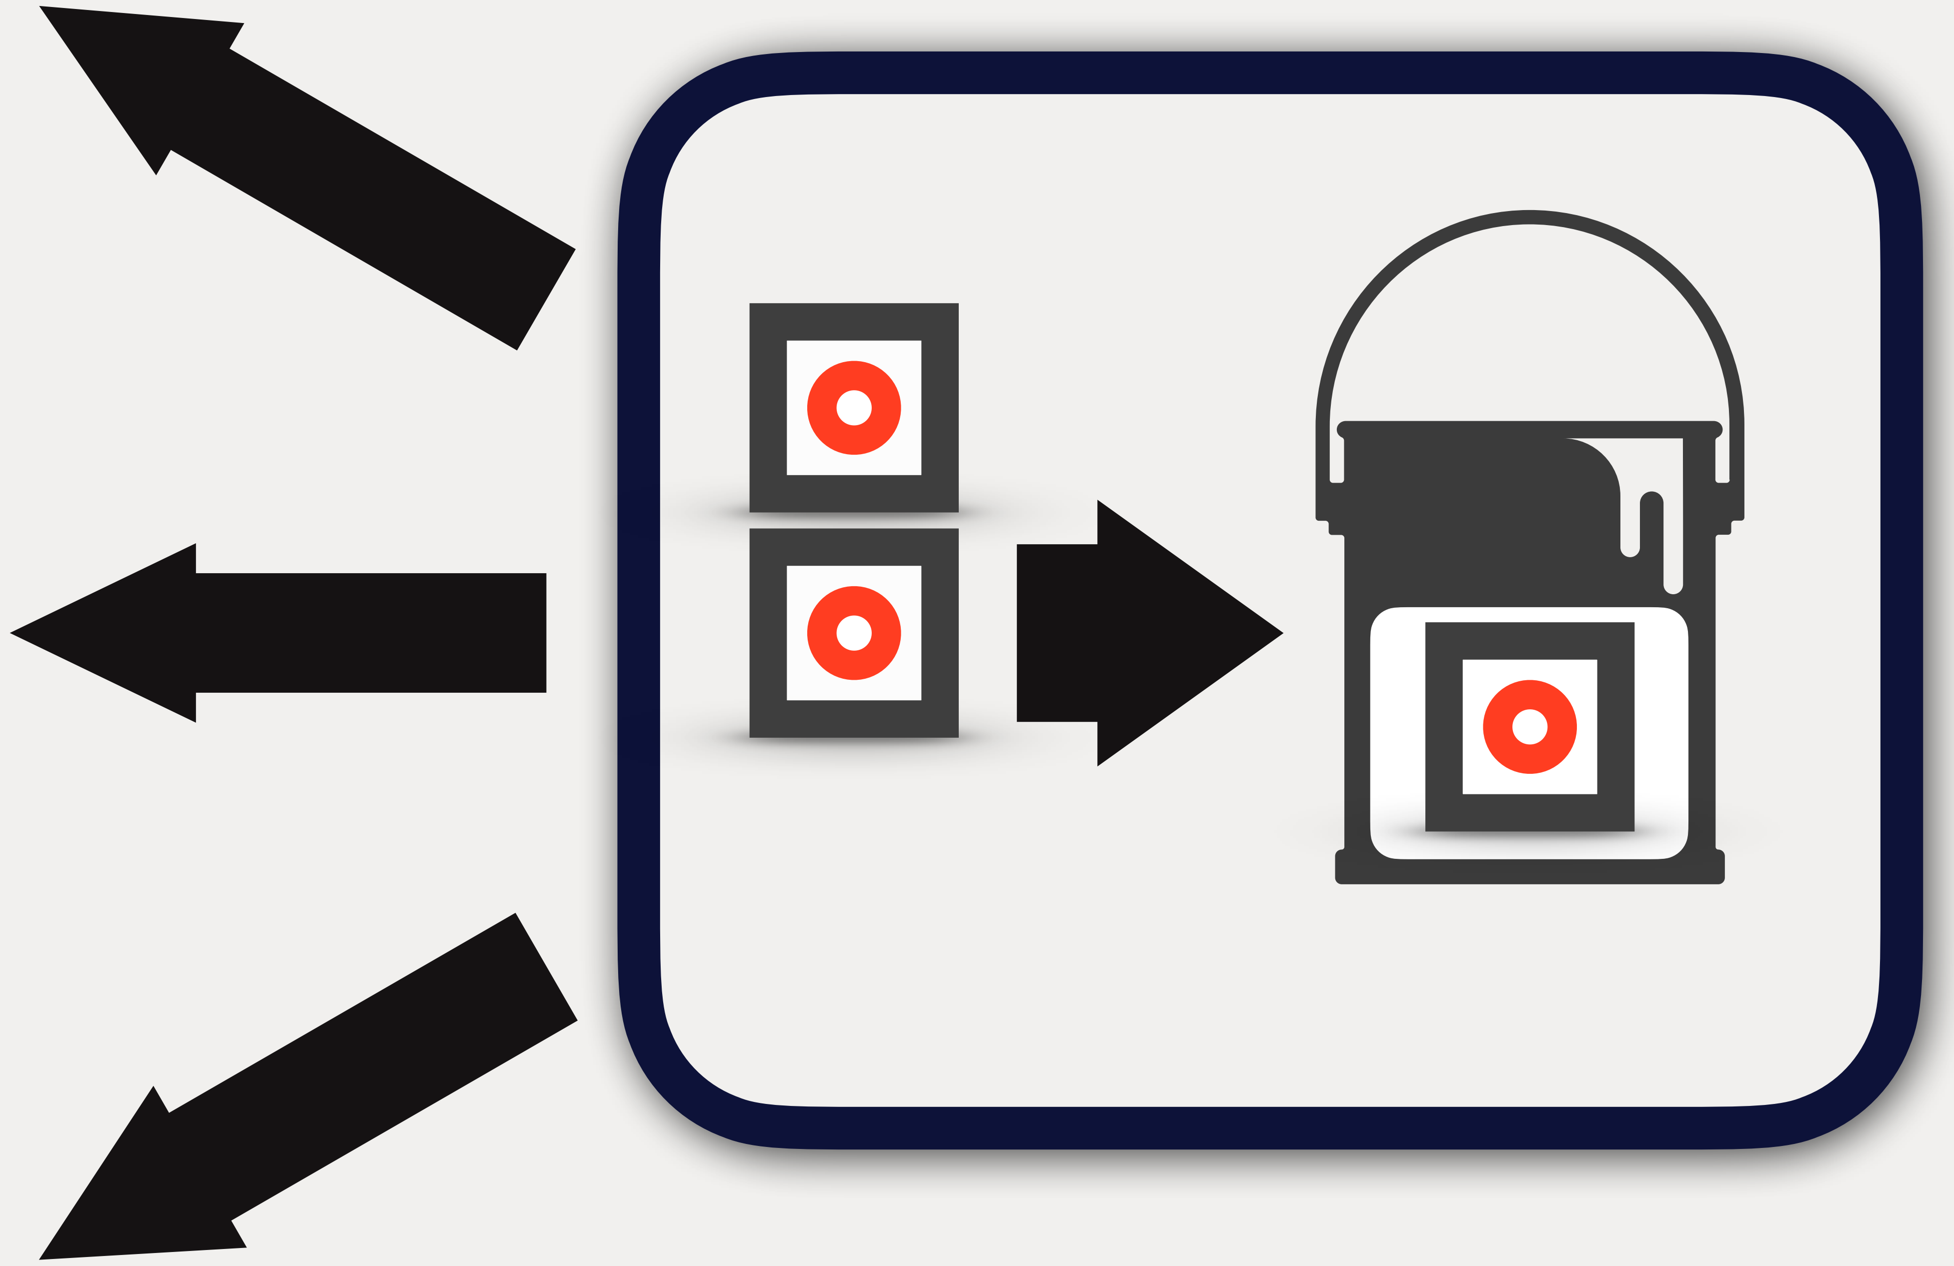

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{M}\sum_{i=1}^{M} w(x_i)\right)\right]$$

Unbiased Reuse

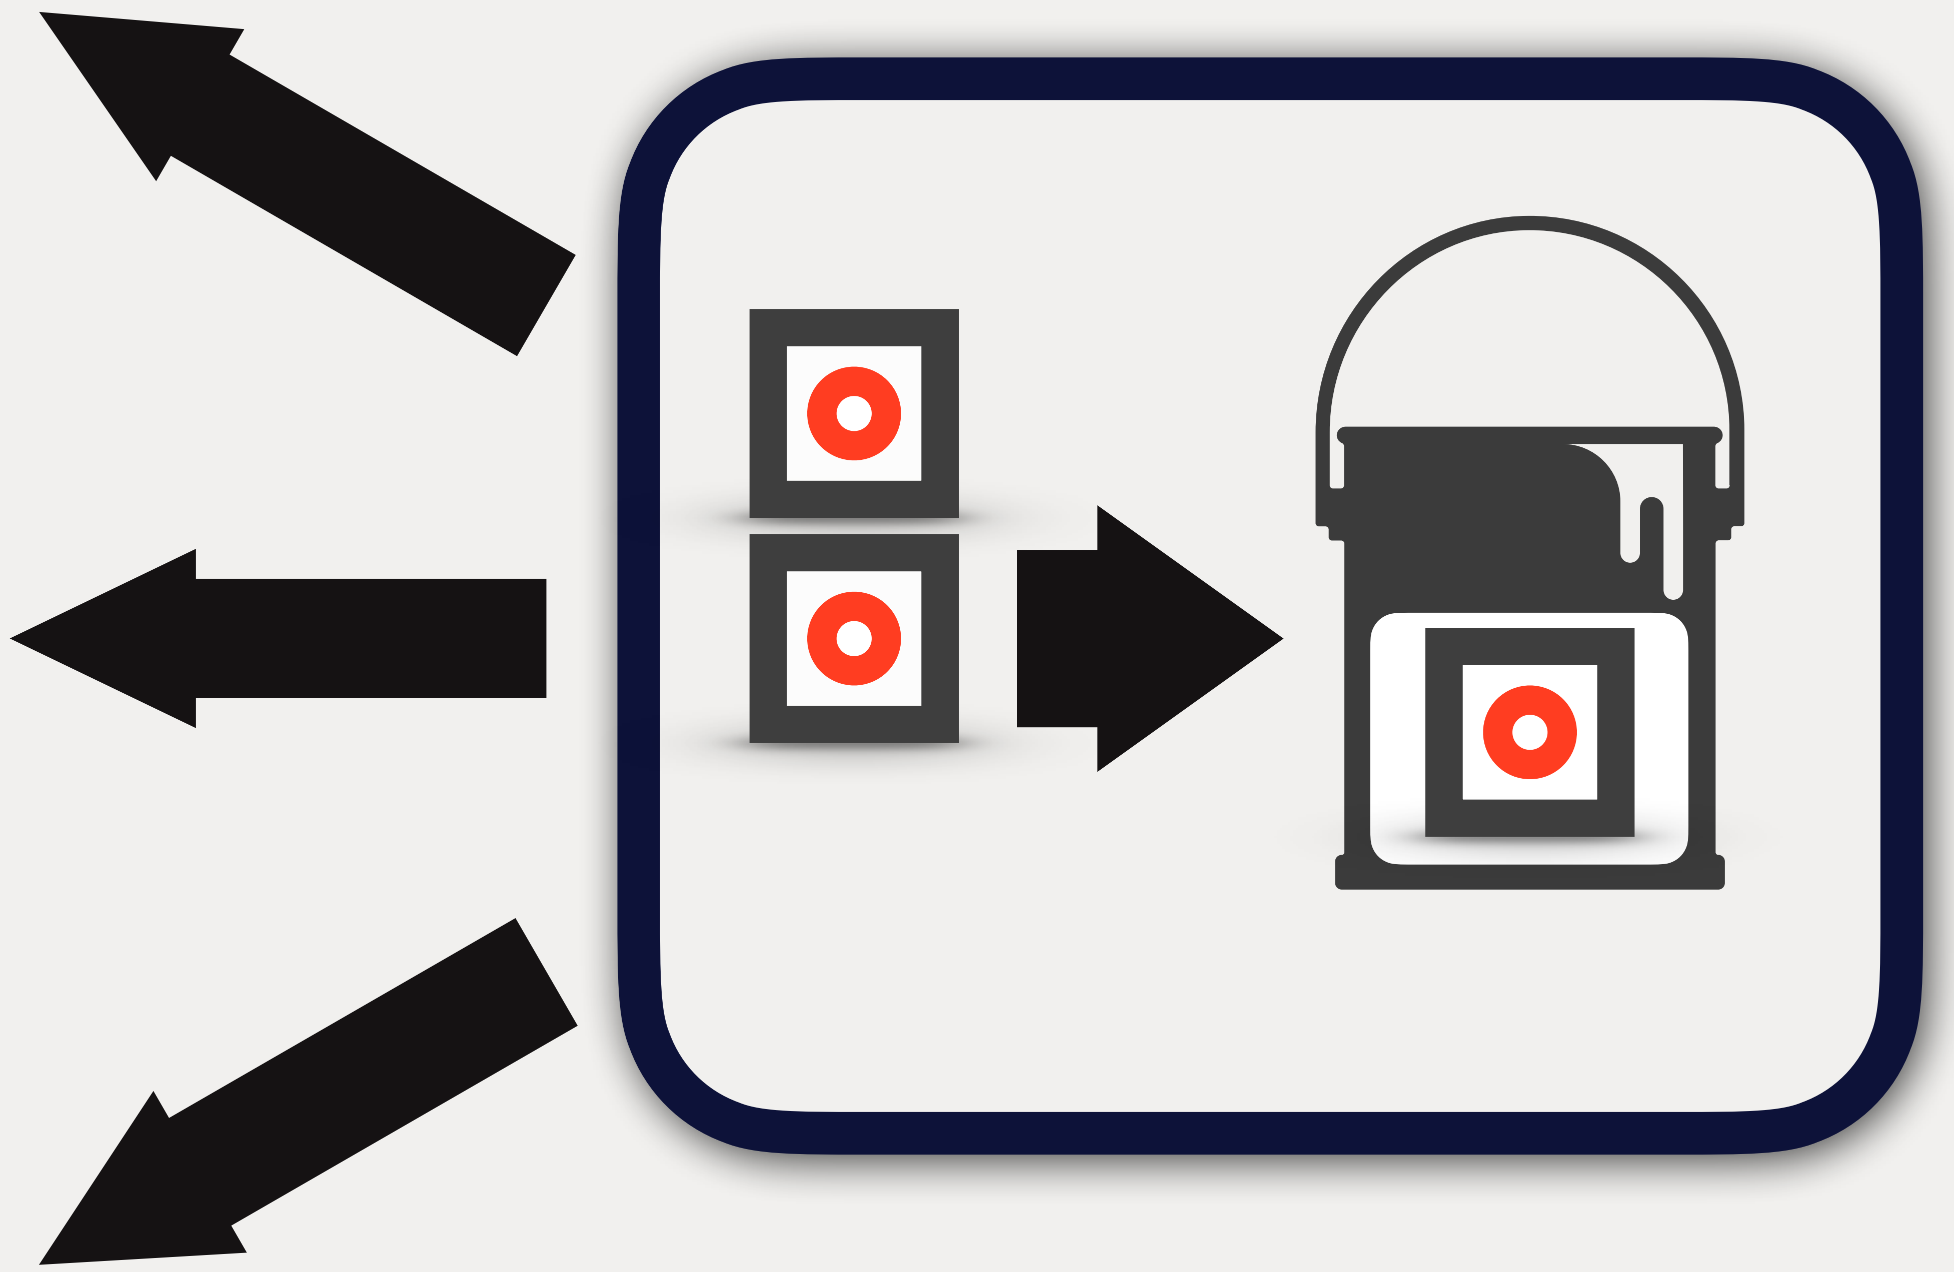

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(\frac{1}{Z}\sum_{i=1}^{M} w(x_i)\right)\right]$$

Unbiased Reuse

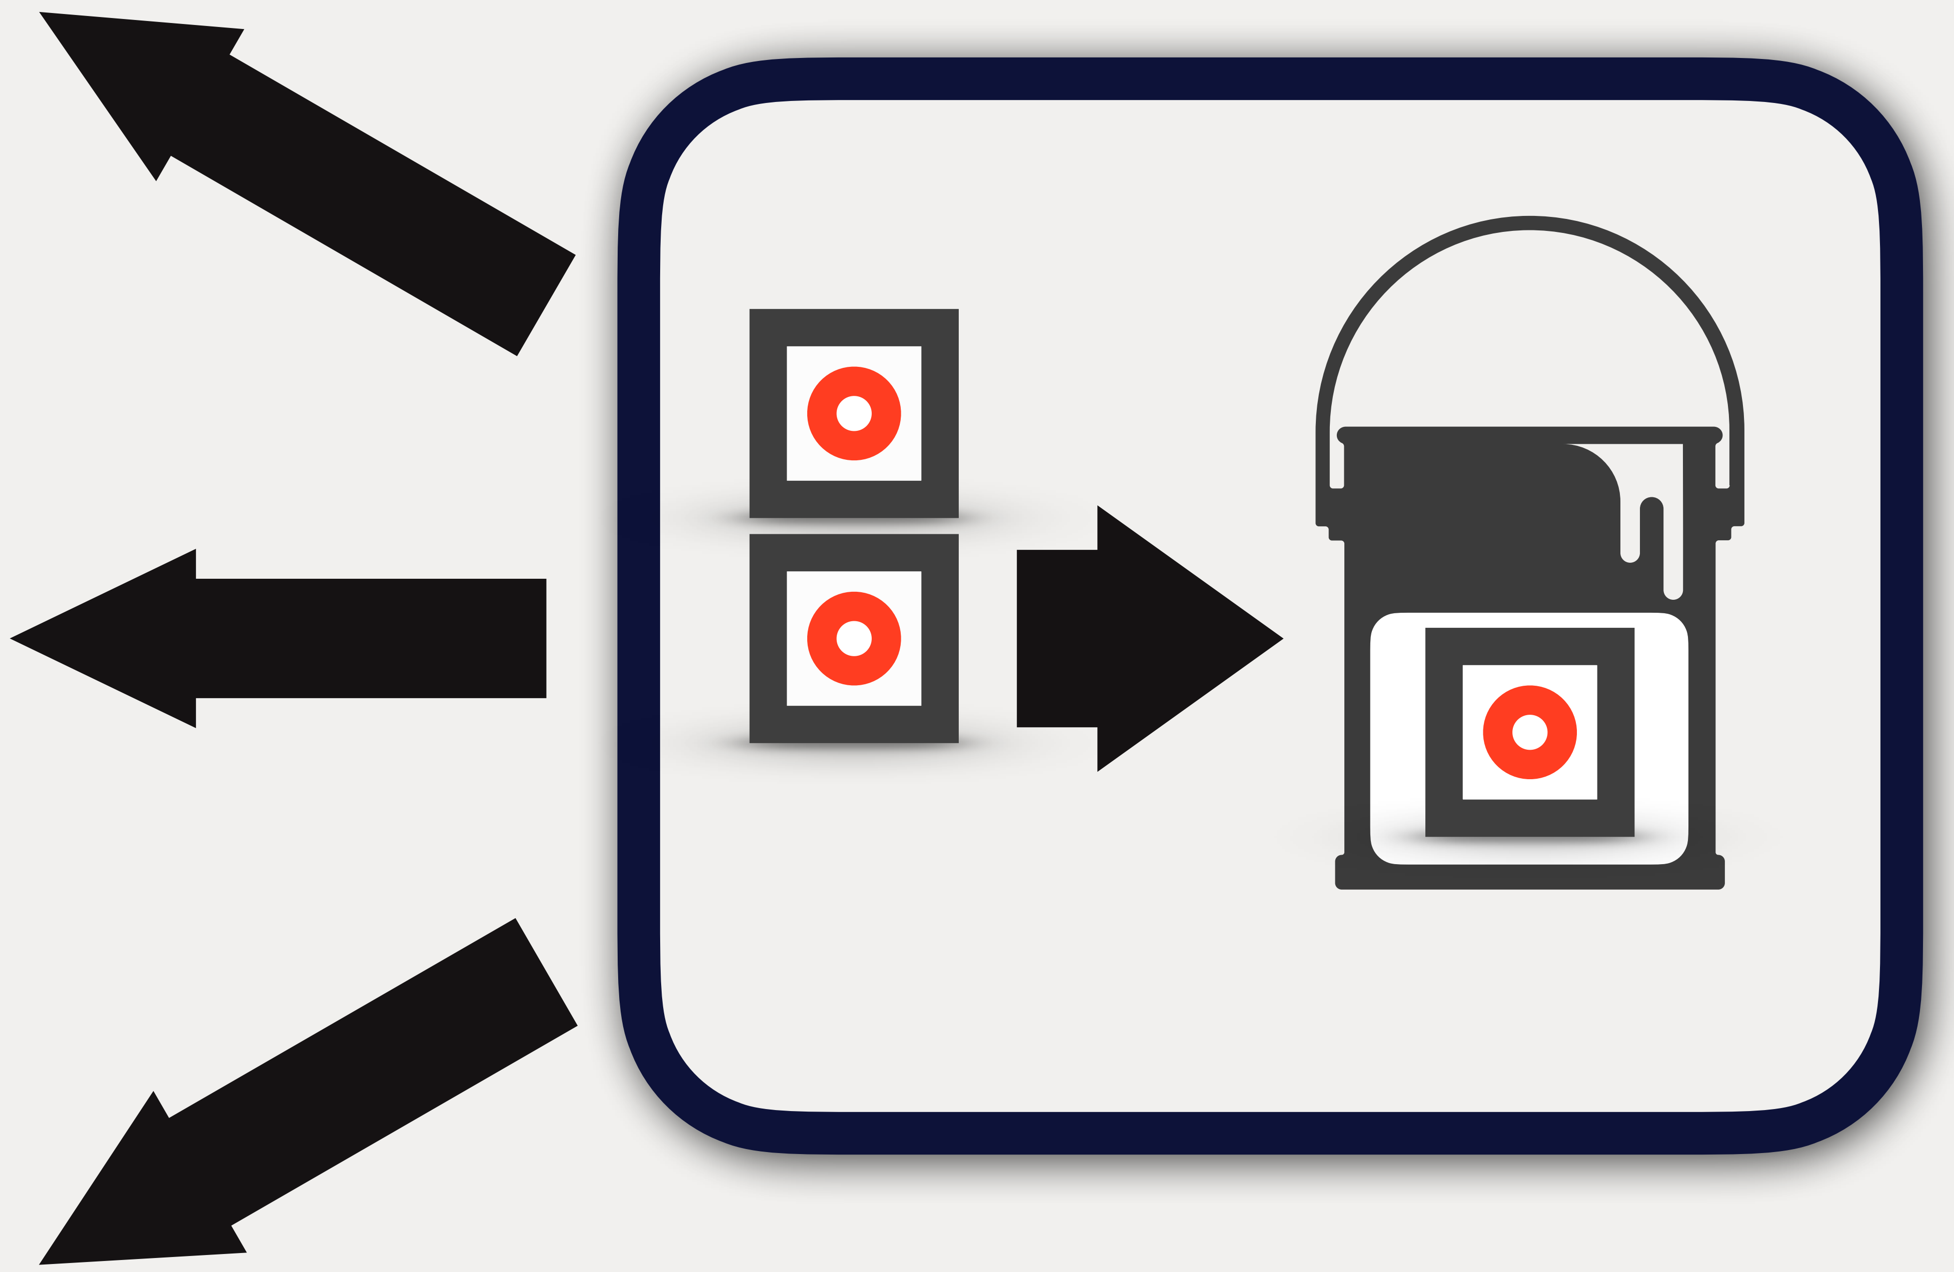

$$\mathbb{E}\left[\frac{1}{\hat{p}(y)} \cdot \left(m_i \sum_{i=1}^{M} w(x_i) \right)\right]$$

Biased Reuse

Biased Reuse

- Correcting bias is *expensive*

Biased Reuse

- Correcting bias is *expensive*
- Can we minimize it instead?

Biased Reuse

- Correcting bias is *expensive*
- Can we minimize it instead?

  $\implies$ Reject potentially troublesome neighbors

Biased Reuse

- Correcting bias is *expensive*
- Can we minimize it instead?

  $\Longrightarrow$ Reject potentially troublesome neighbors

- Compare normals/depth

Results

Algorithm Details

Algorithm Details

- Implemented in Falcor [Benty et al. 2019]

Algorithm Details

- Implemented in Falcor [Benty et al. 2019]
- Initial Resampling: M=32

Algorithm Details

- Implemented in Falcor [Benty et al. 2019]
- Initial Resampling: M=32
- Temporal Reuse: Reproject with motion vectors

Algorithm Details

- Implemented in Falcor [Benty et al. 2019]
- Initial Resampling: M=32
- Temporal Reuse: Reproject with motion vectors
- Spatial Reuse: Pick 5 random neighbors in 30 pixel disk
  - Repeated twice


Reference
20'000 Emitters
Restaurant
Le Petit Coin


[Moreau et al., 2019], 34ms
20’000 Emitters


ReSTIR (unbiased), 30ms
20'000 Emitters


ReSTIR (biased), 24ms
20’000 Emitters


Reference
20’000 Emitters
Le Petit Coin
Le Petit Coin
MENU


[Moreau et al., 2019], 30ms
20’000 Emitters


ReSTIR (unbiased), 26ms
20’000 Emitters
MENU


ReSTIR (biased), 21ms
20’000 Emitters
MENU


Reference
23'000 Emitters
S City Subway


[Moreau et al., 2019], 29ms
23’000 Emitters


ReSTIR (unbiased), 17ms
23’000 Emitters


ReSTIR (biased), 16ms
23’000 Emitters
City
Subway

Convergence Graphs

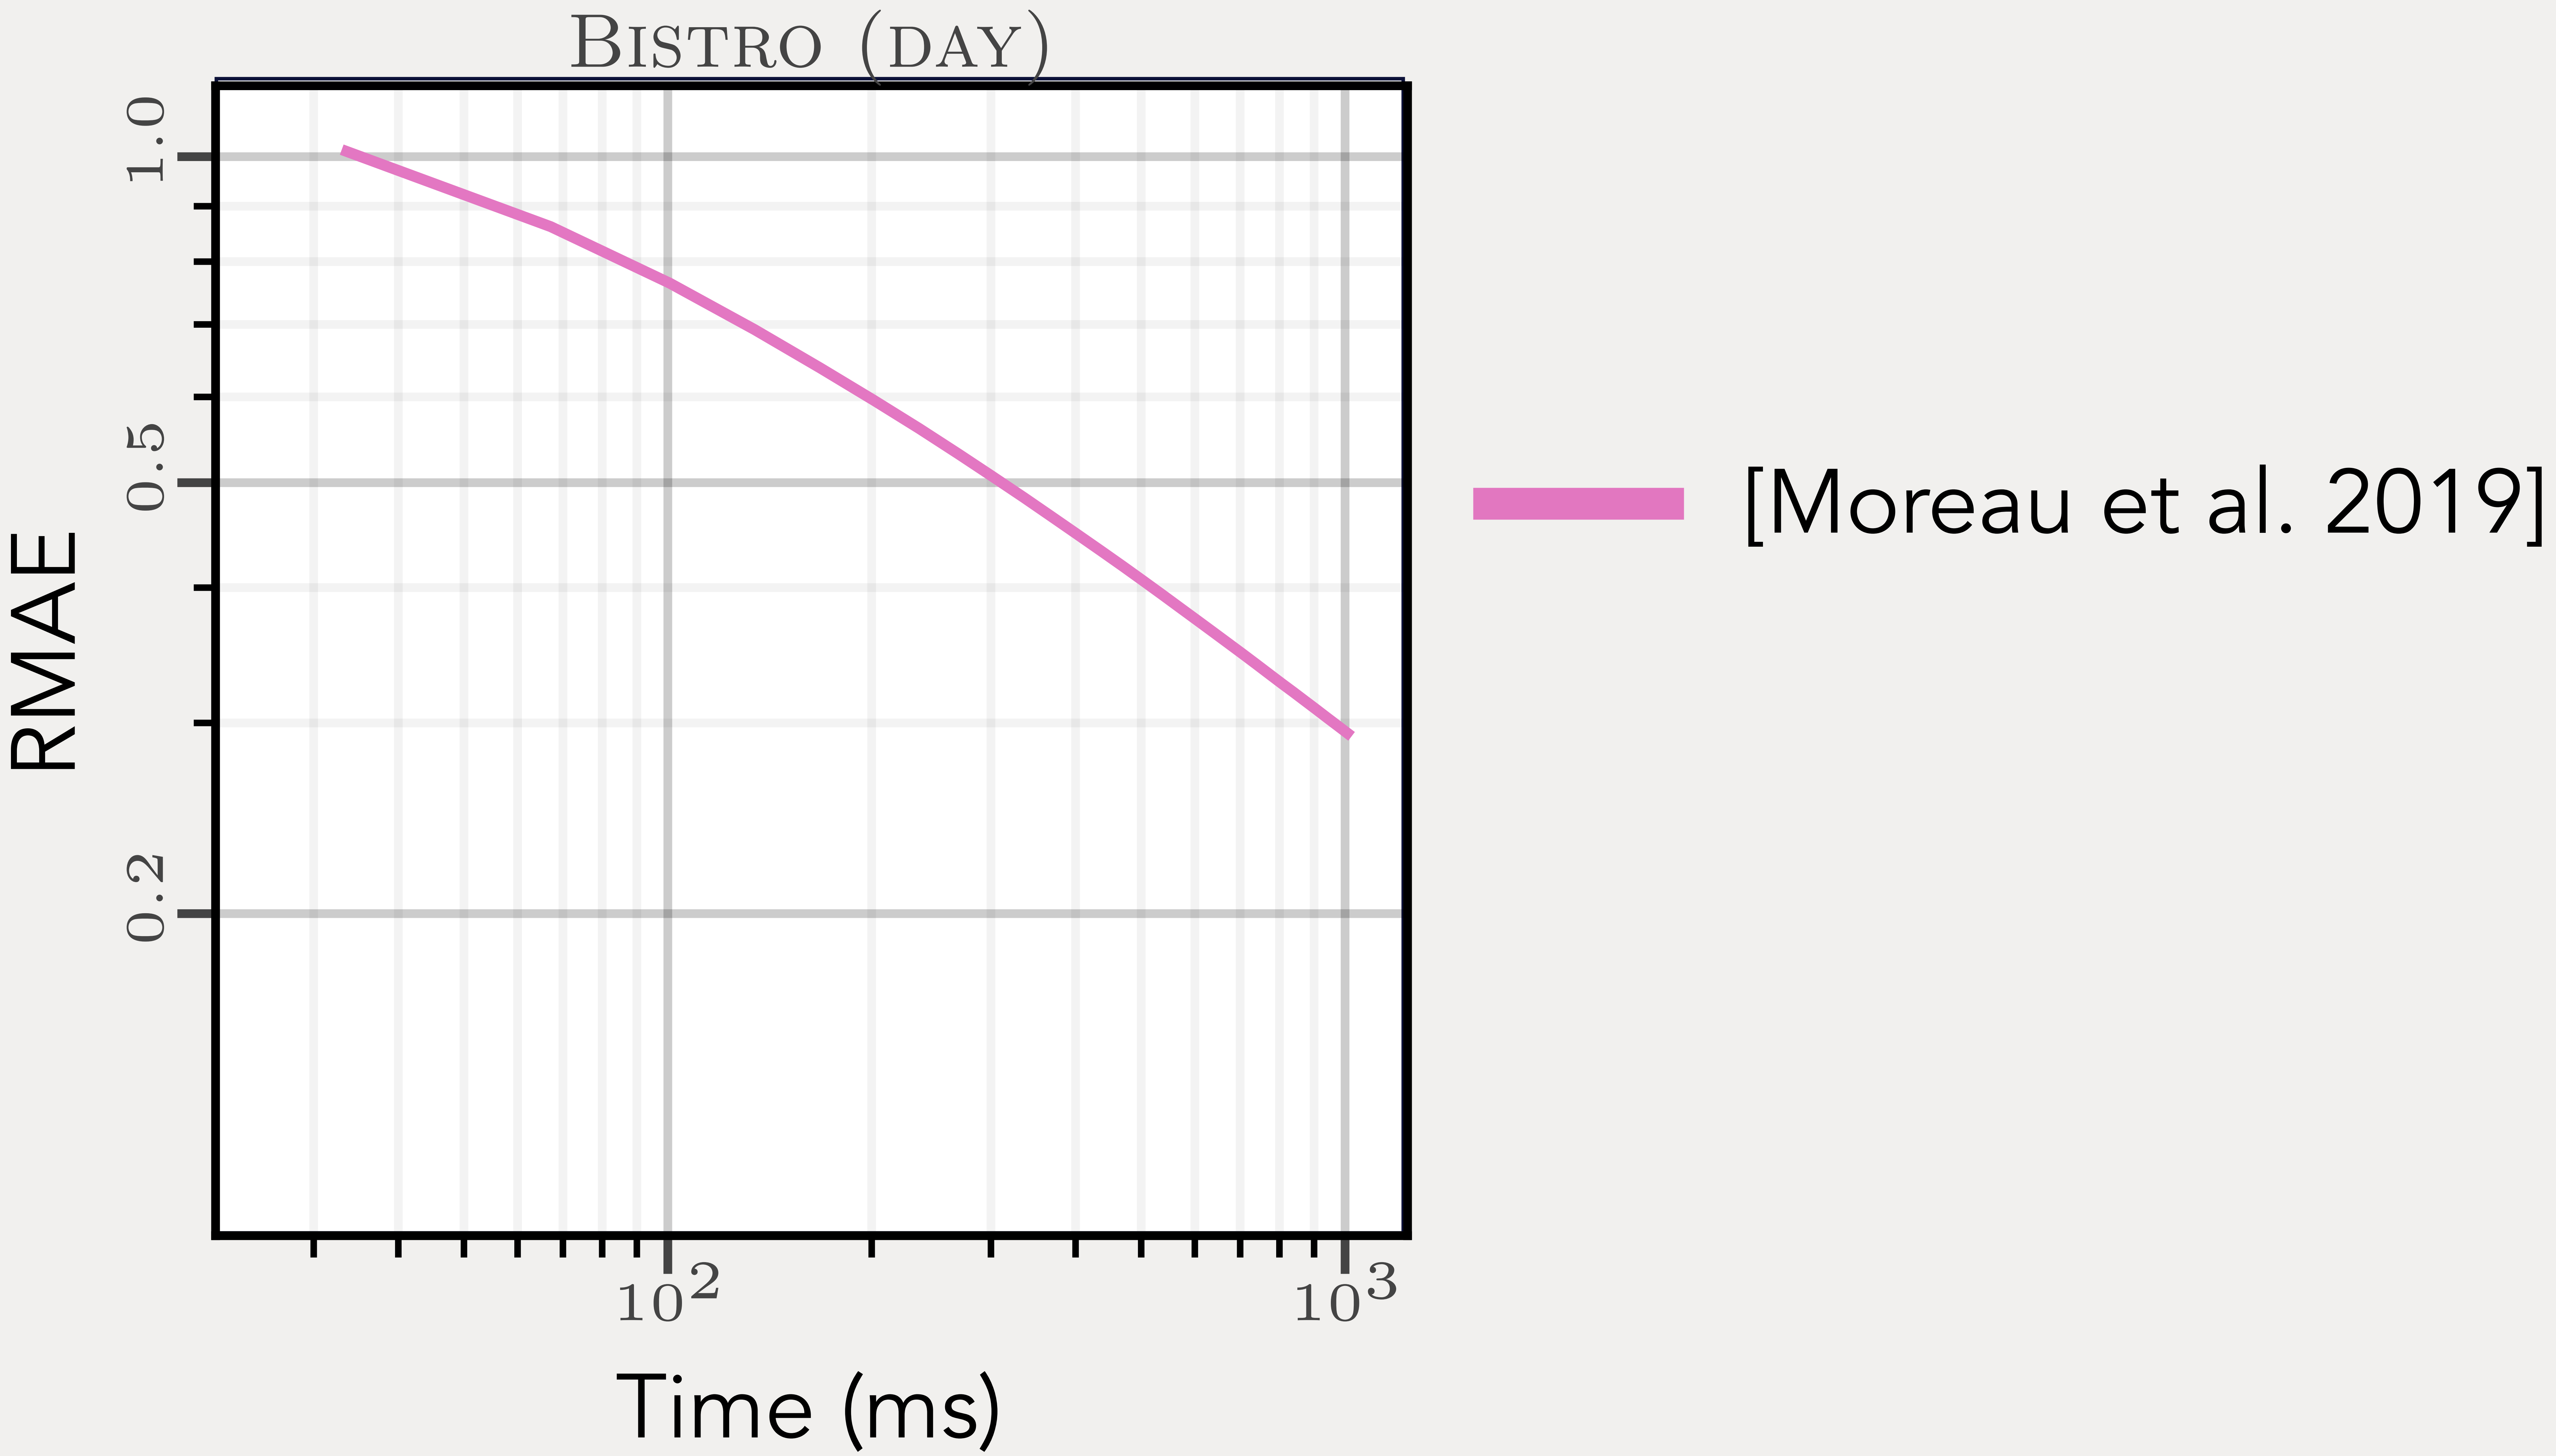

Convergence Graphs

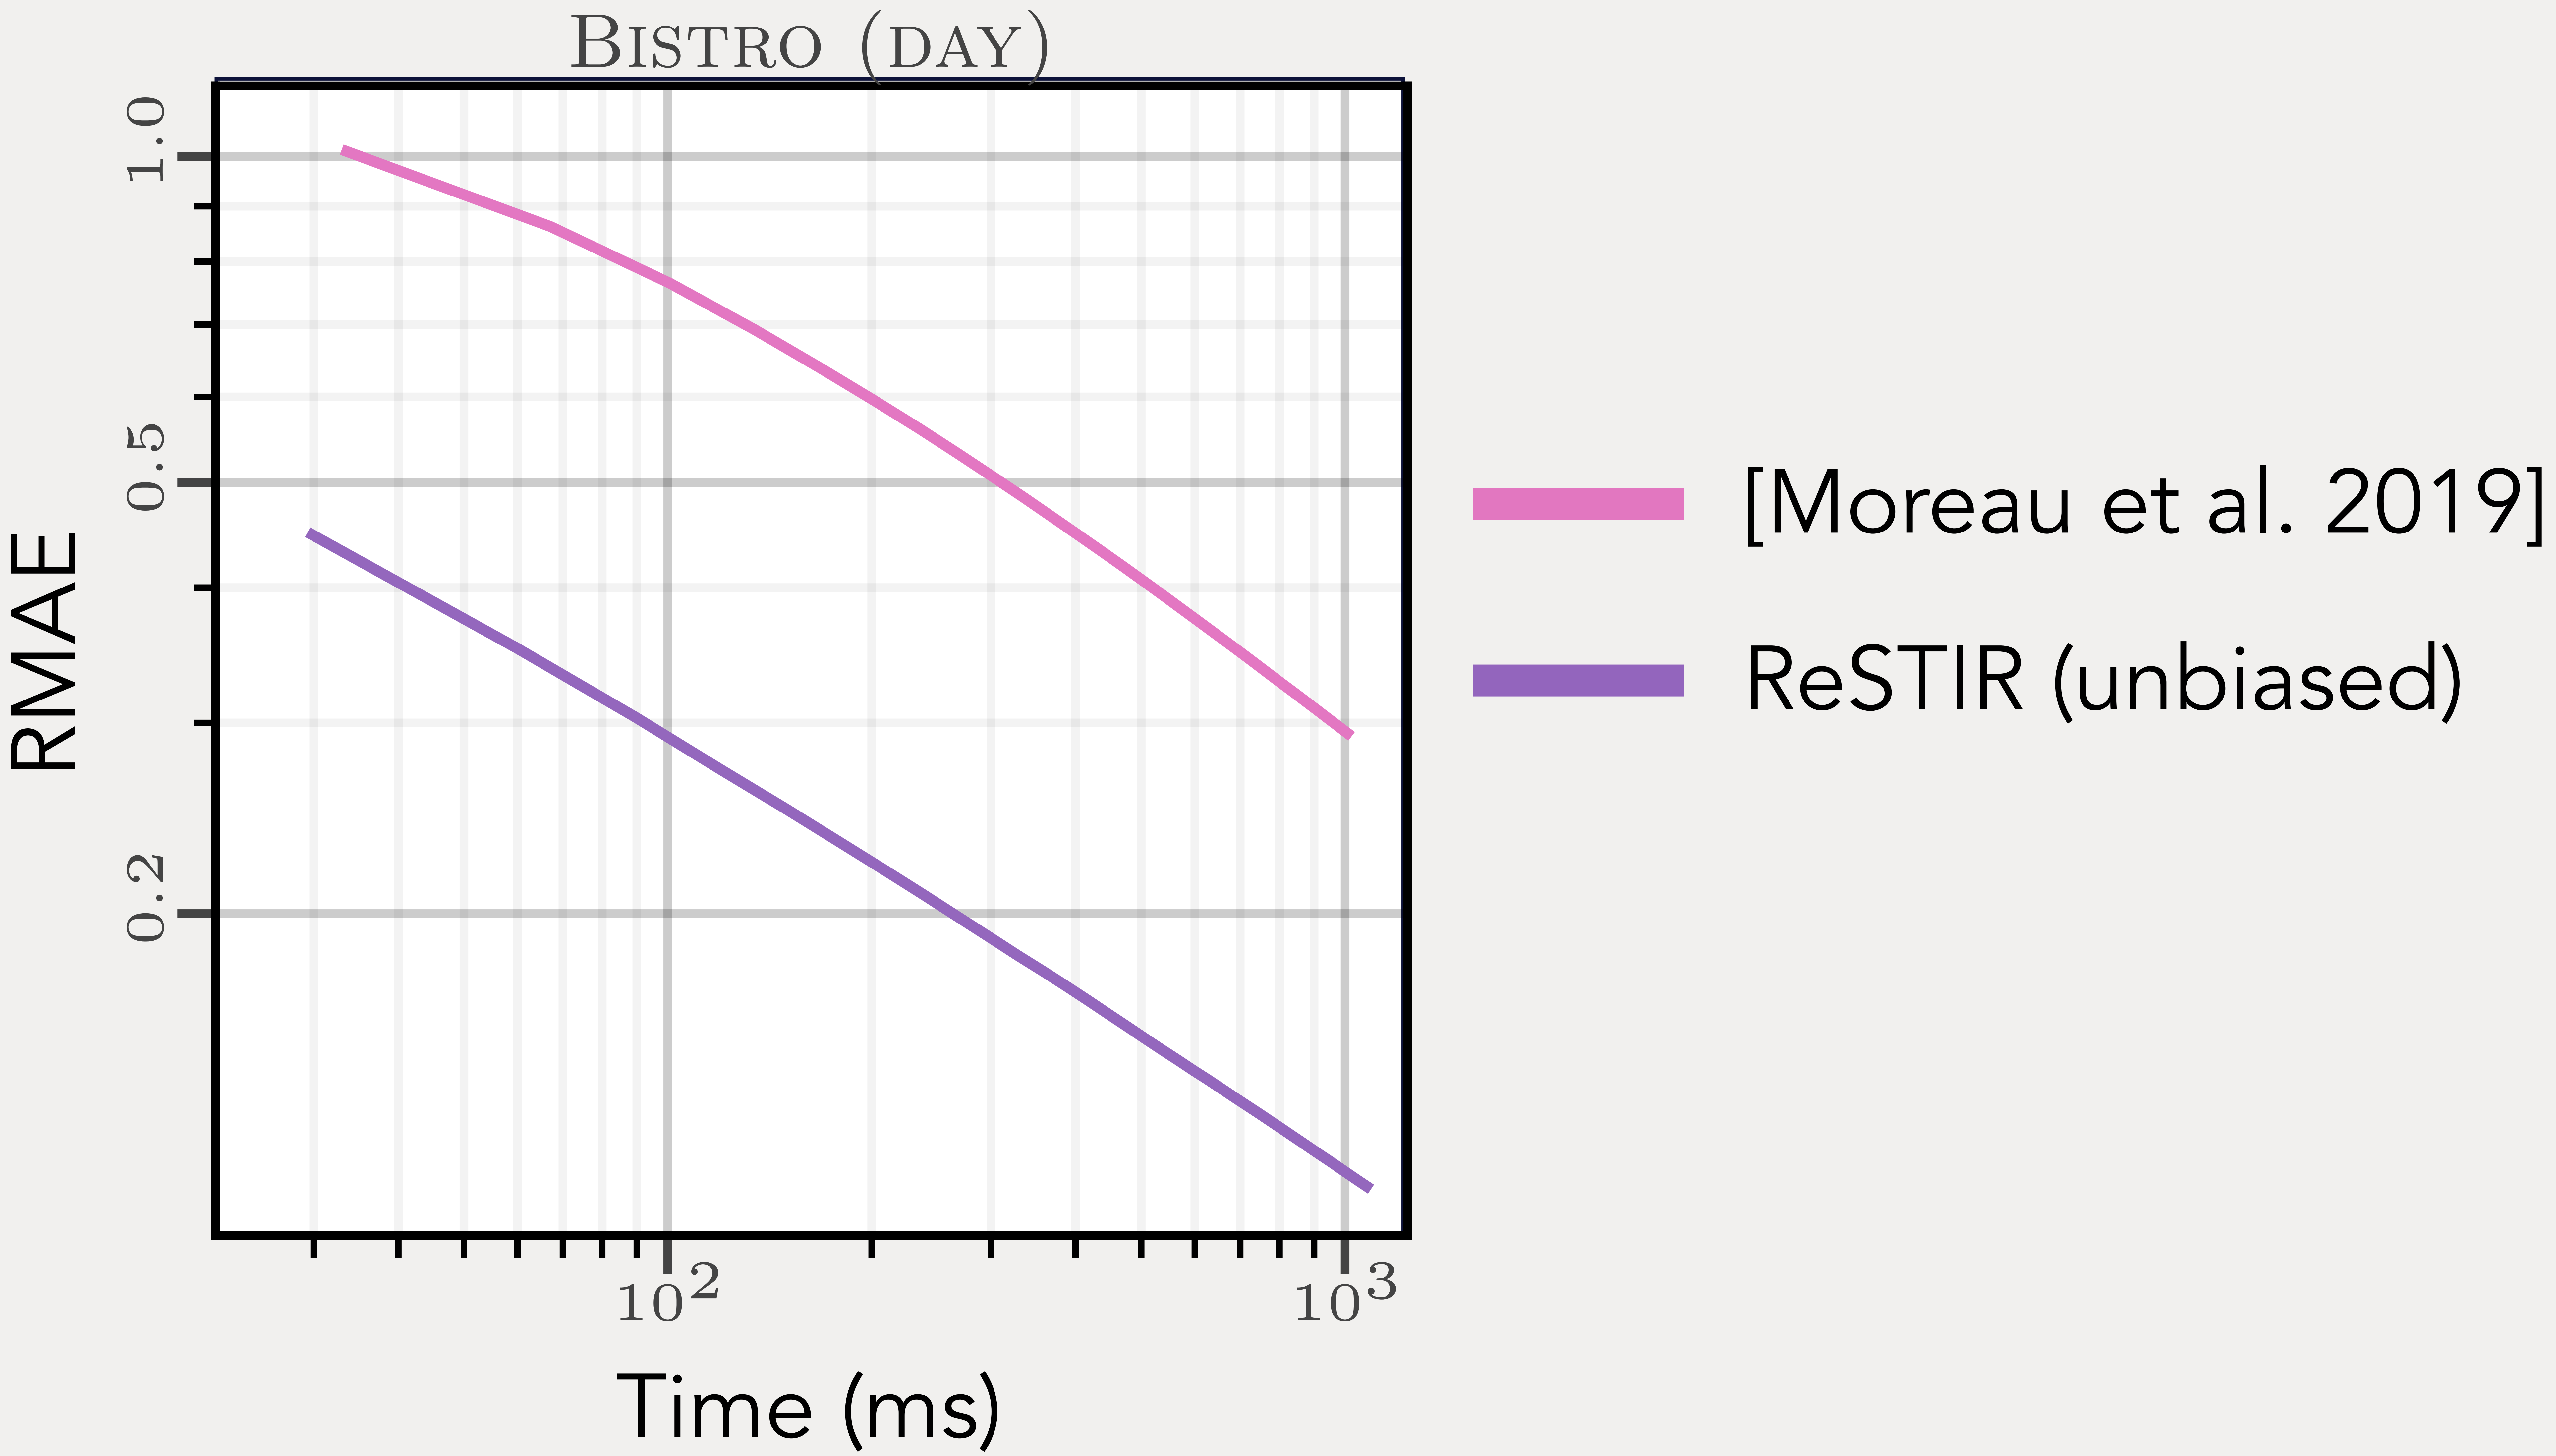

Convergence Graphs

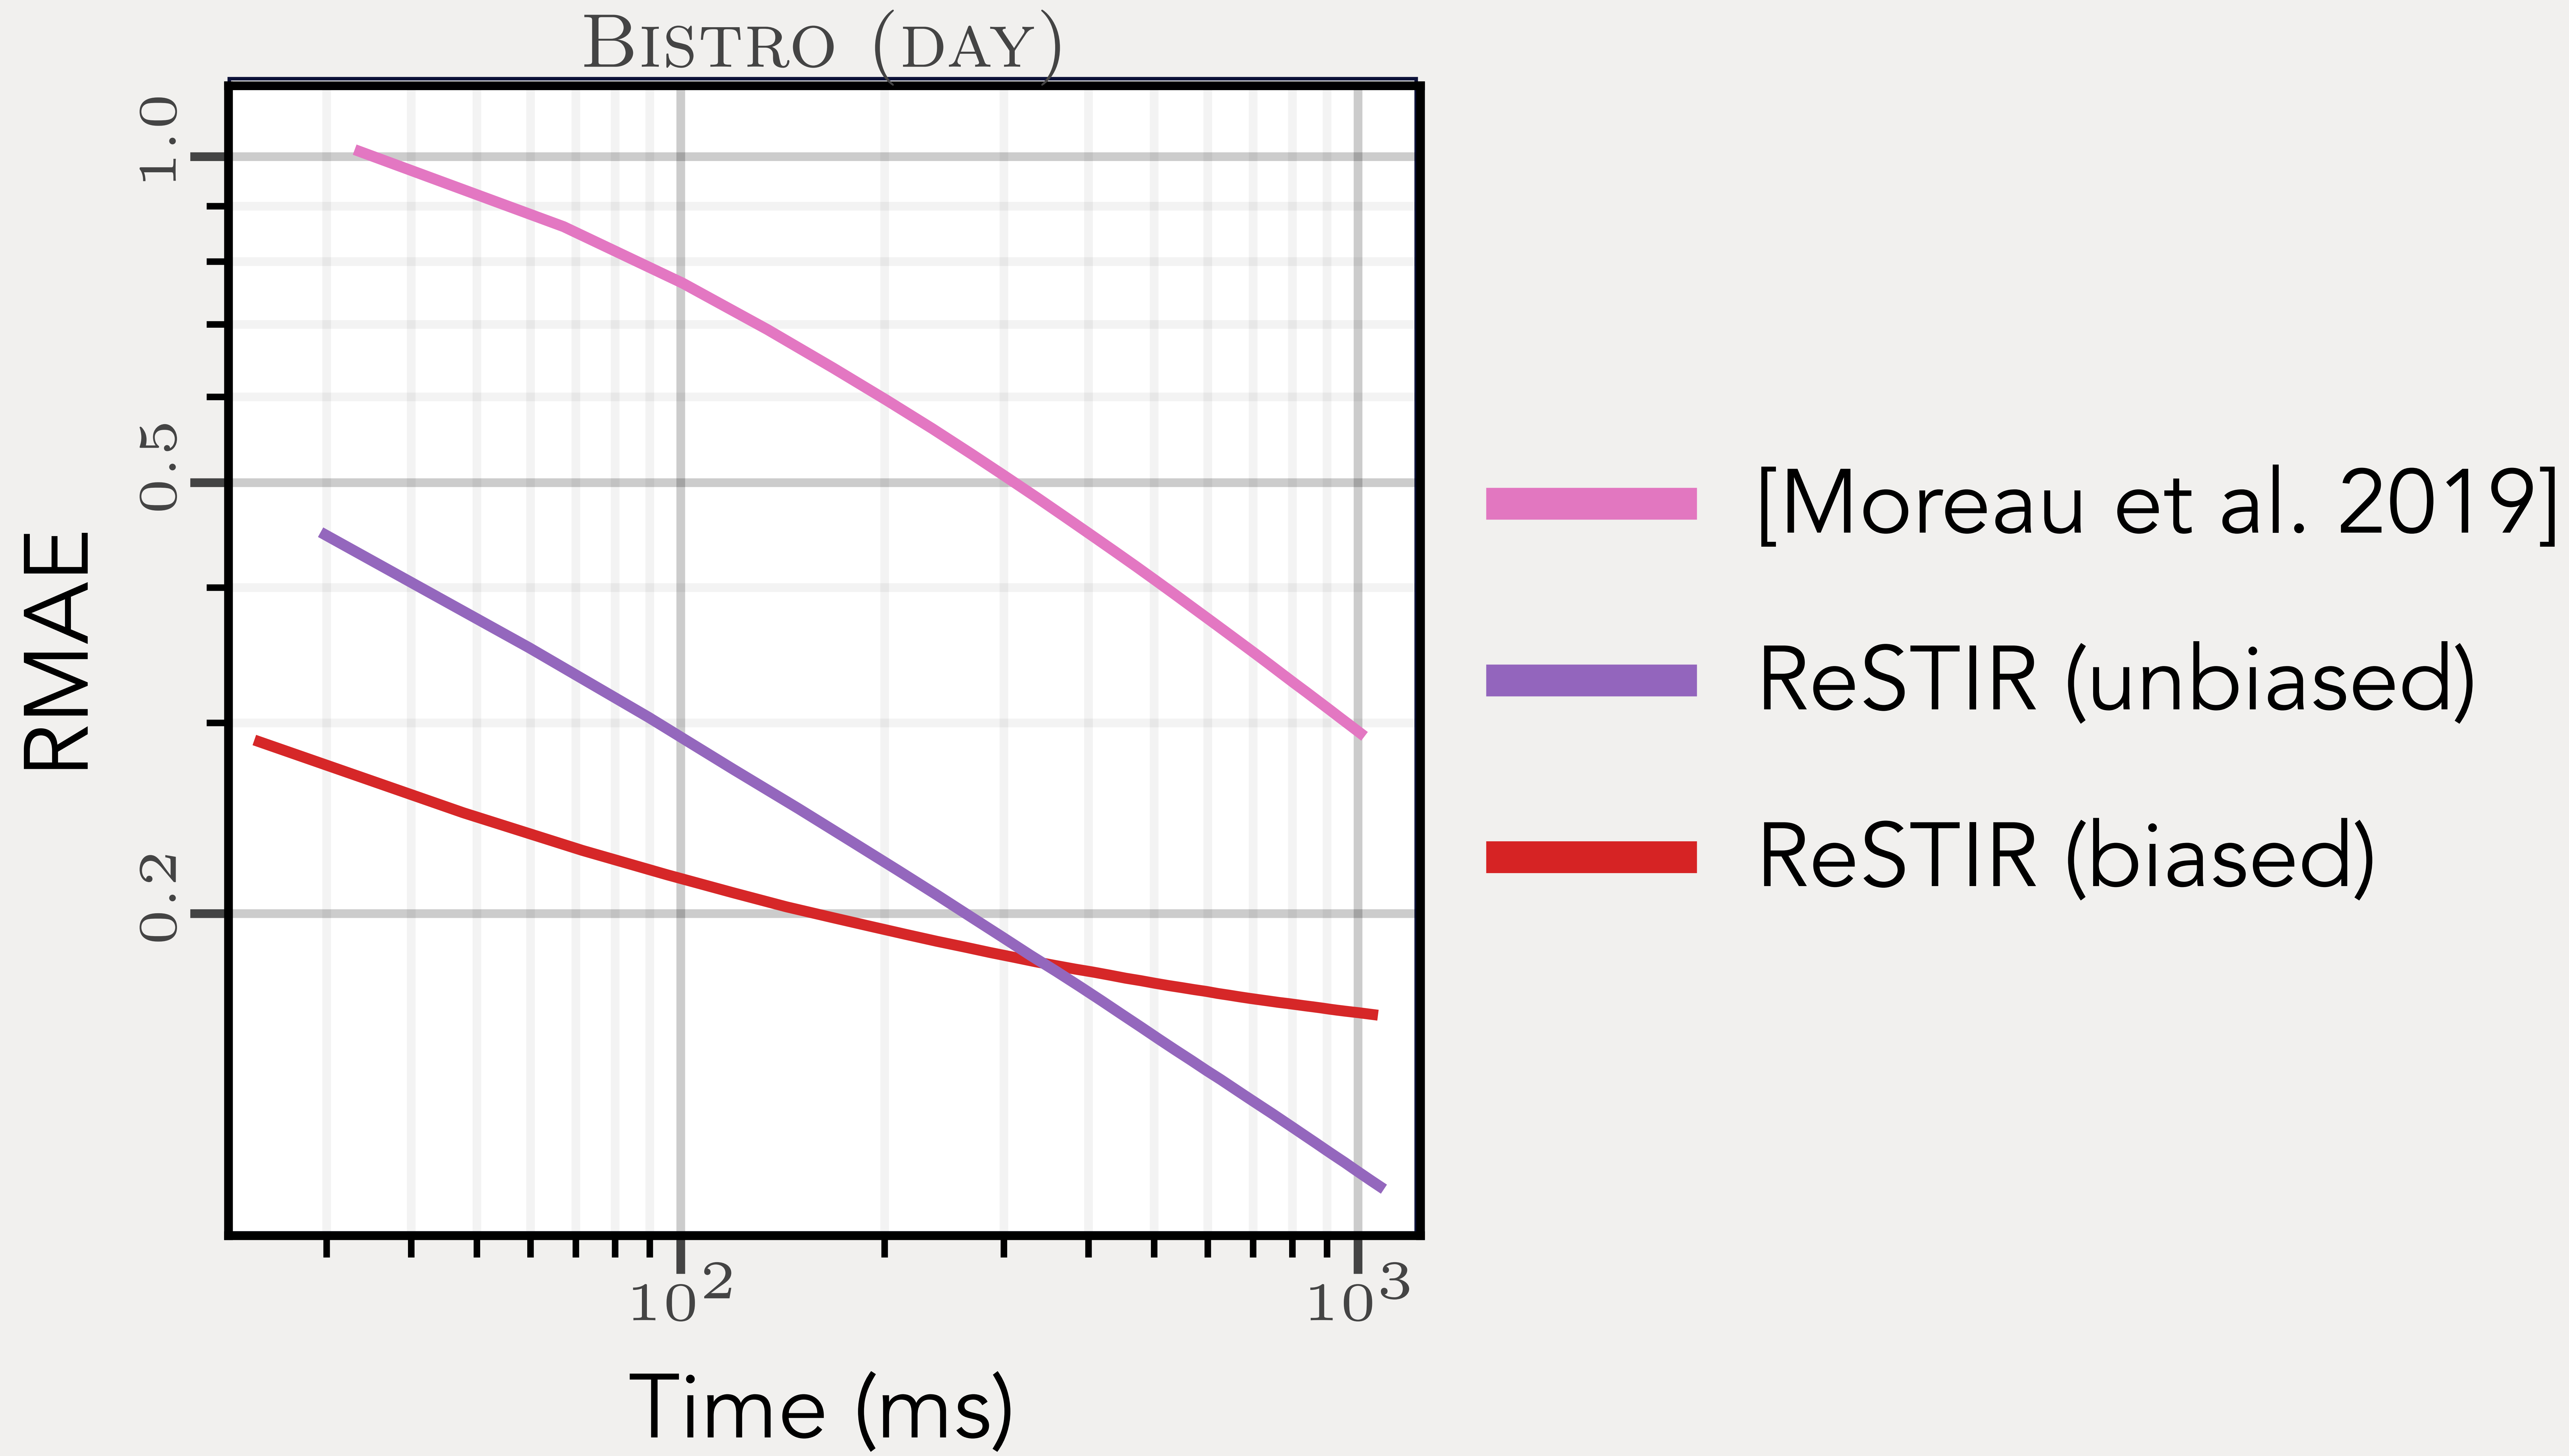

Convergence Graphs

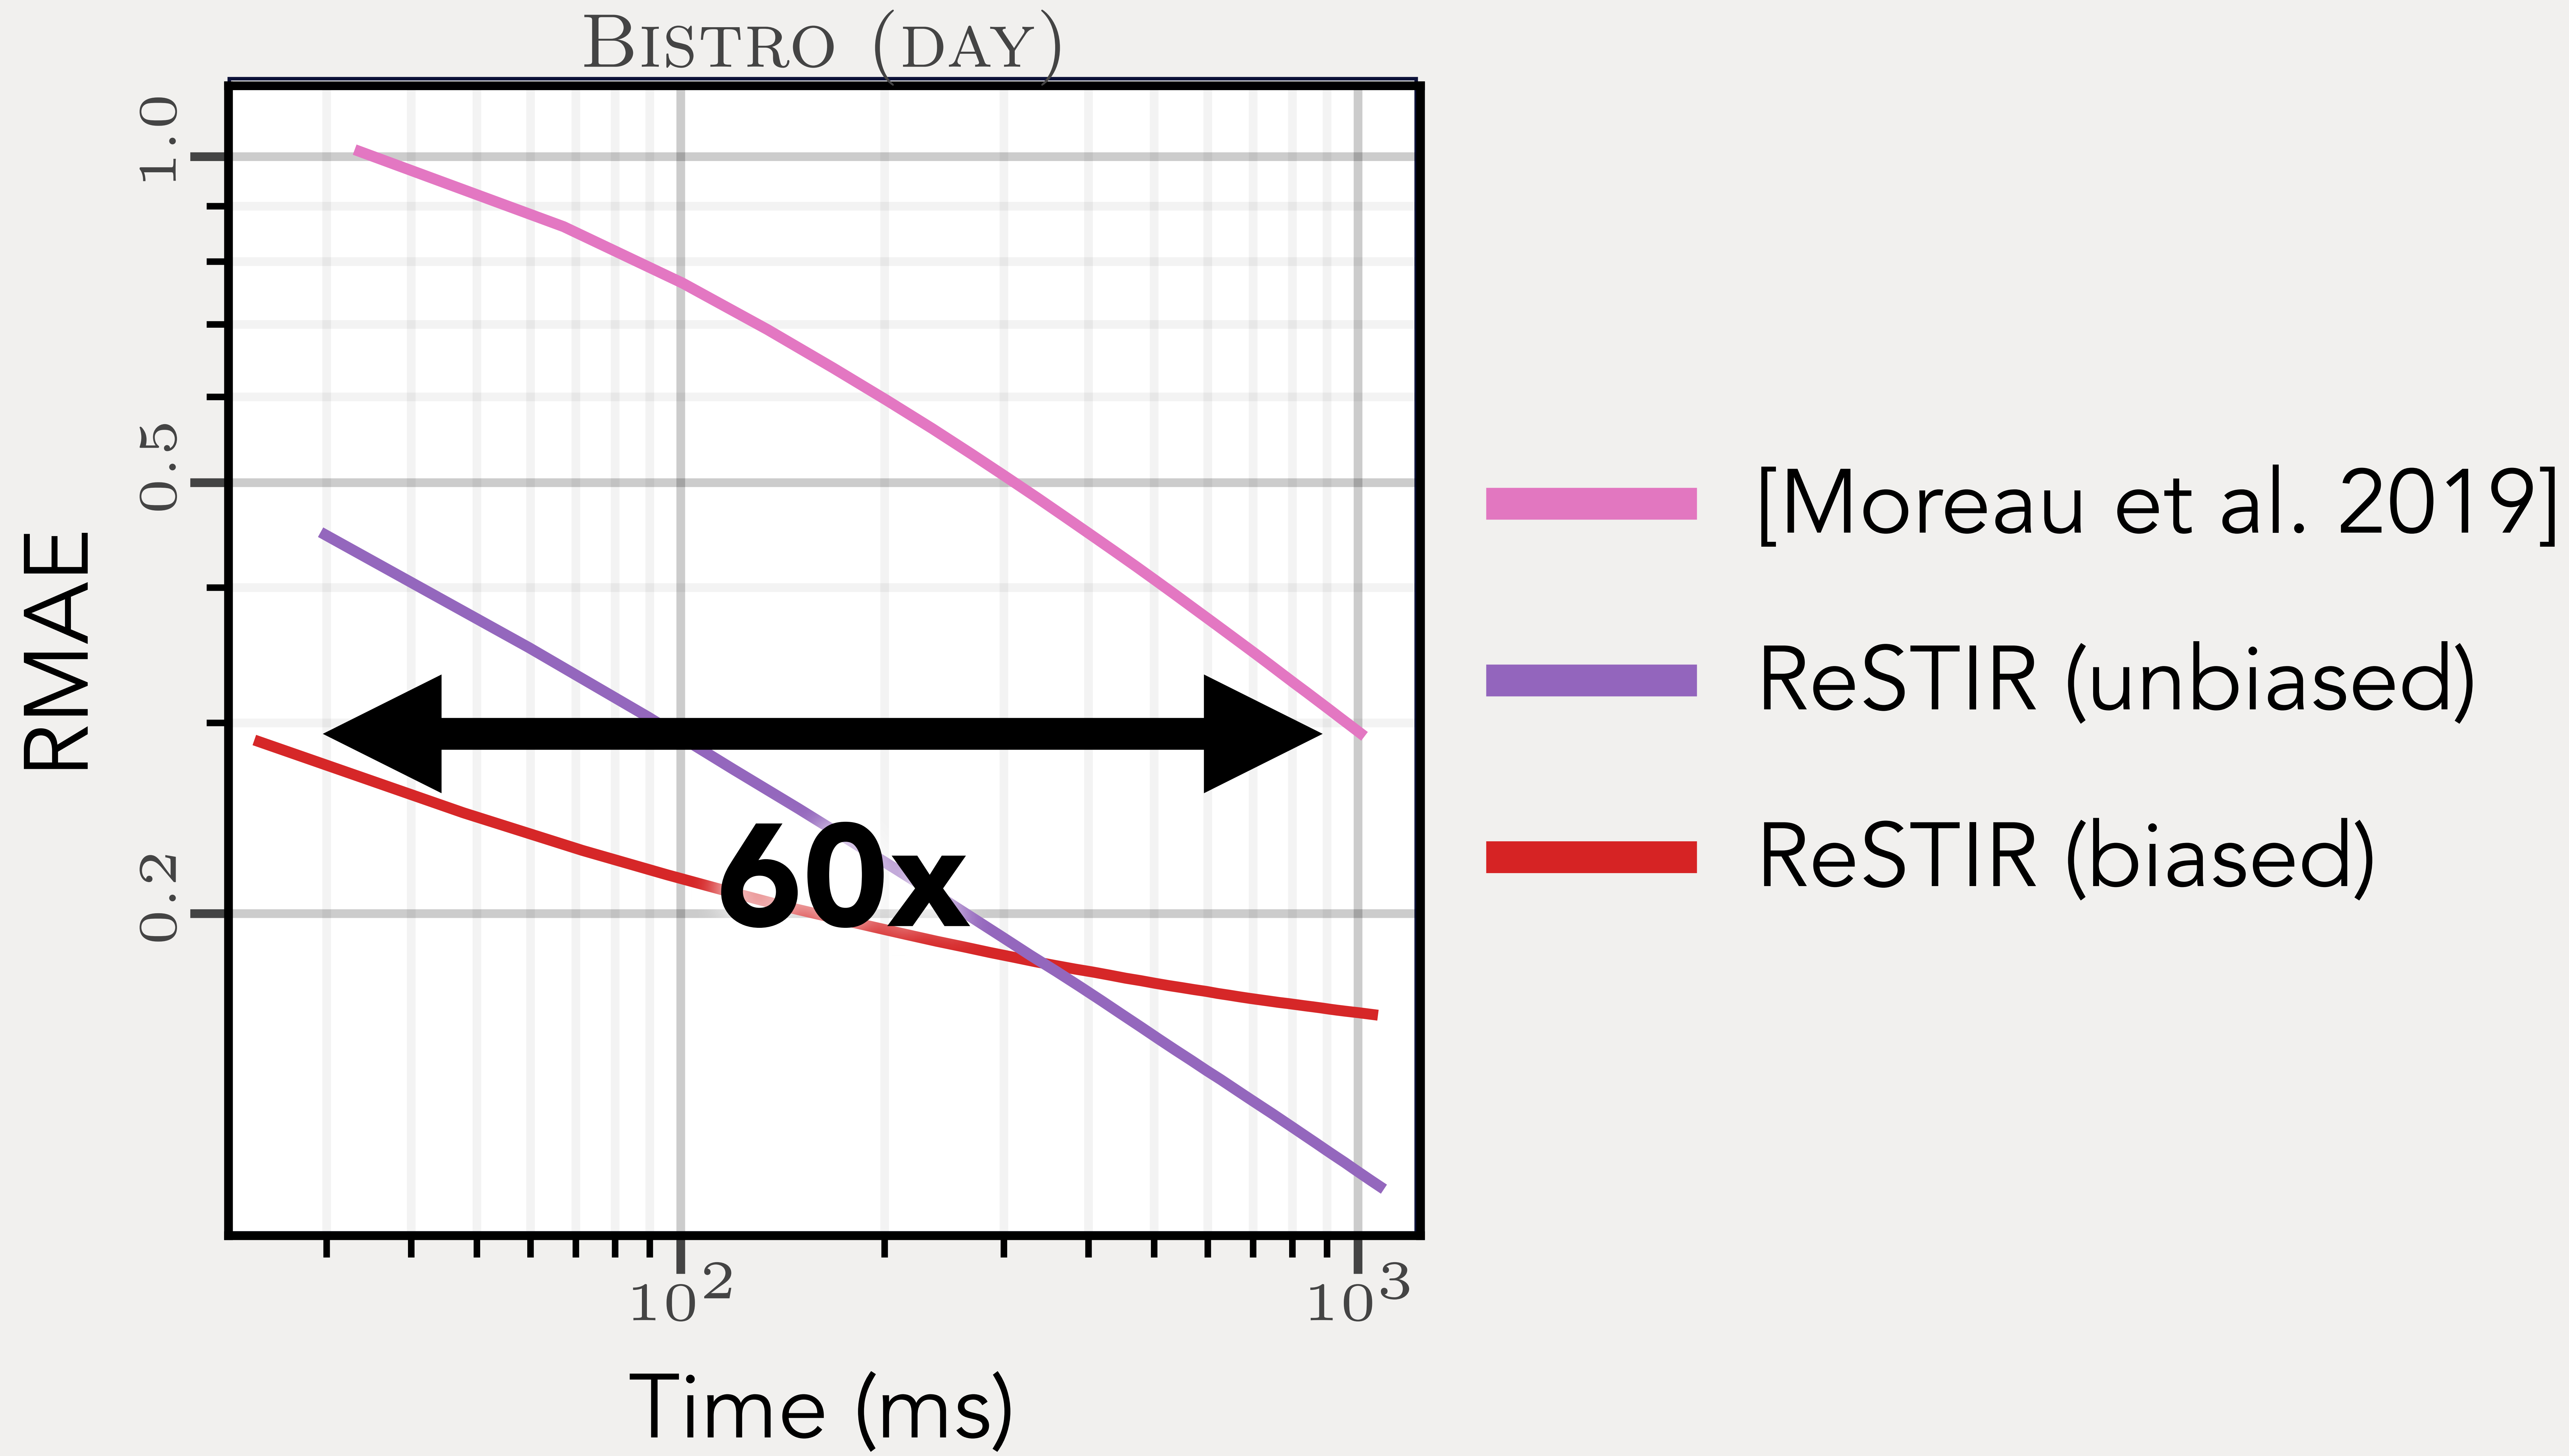

Convergence Graphs

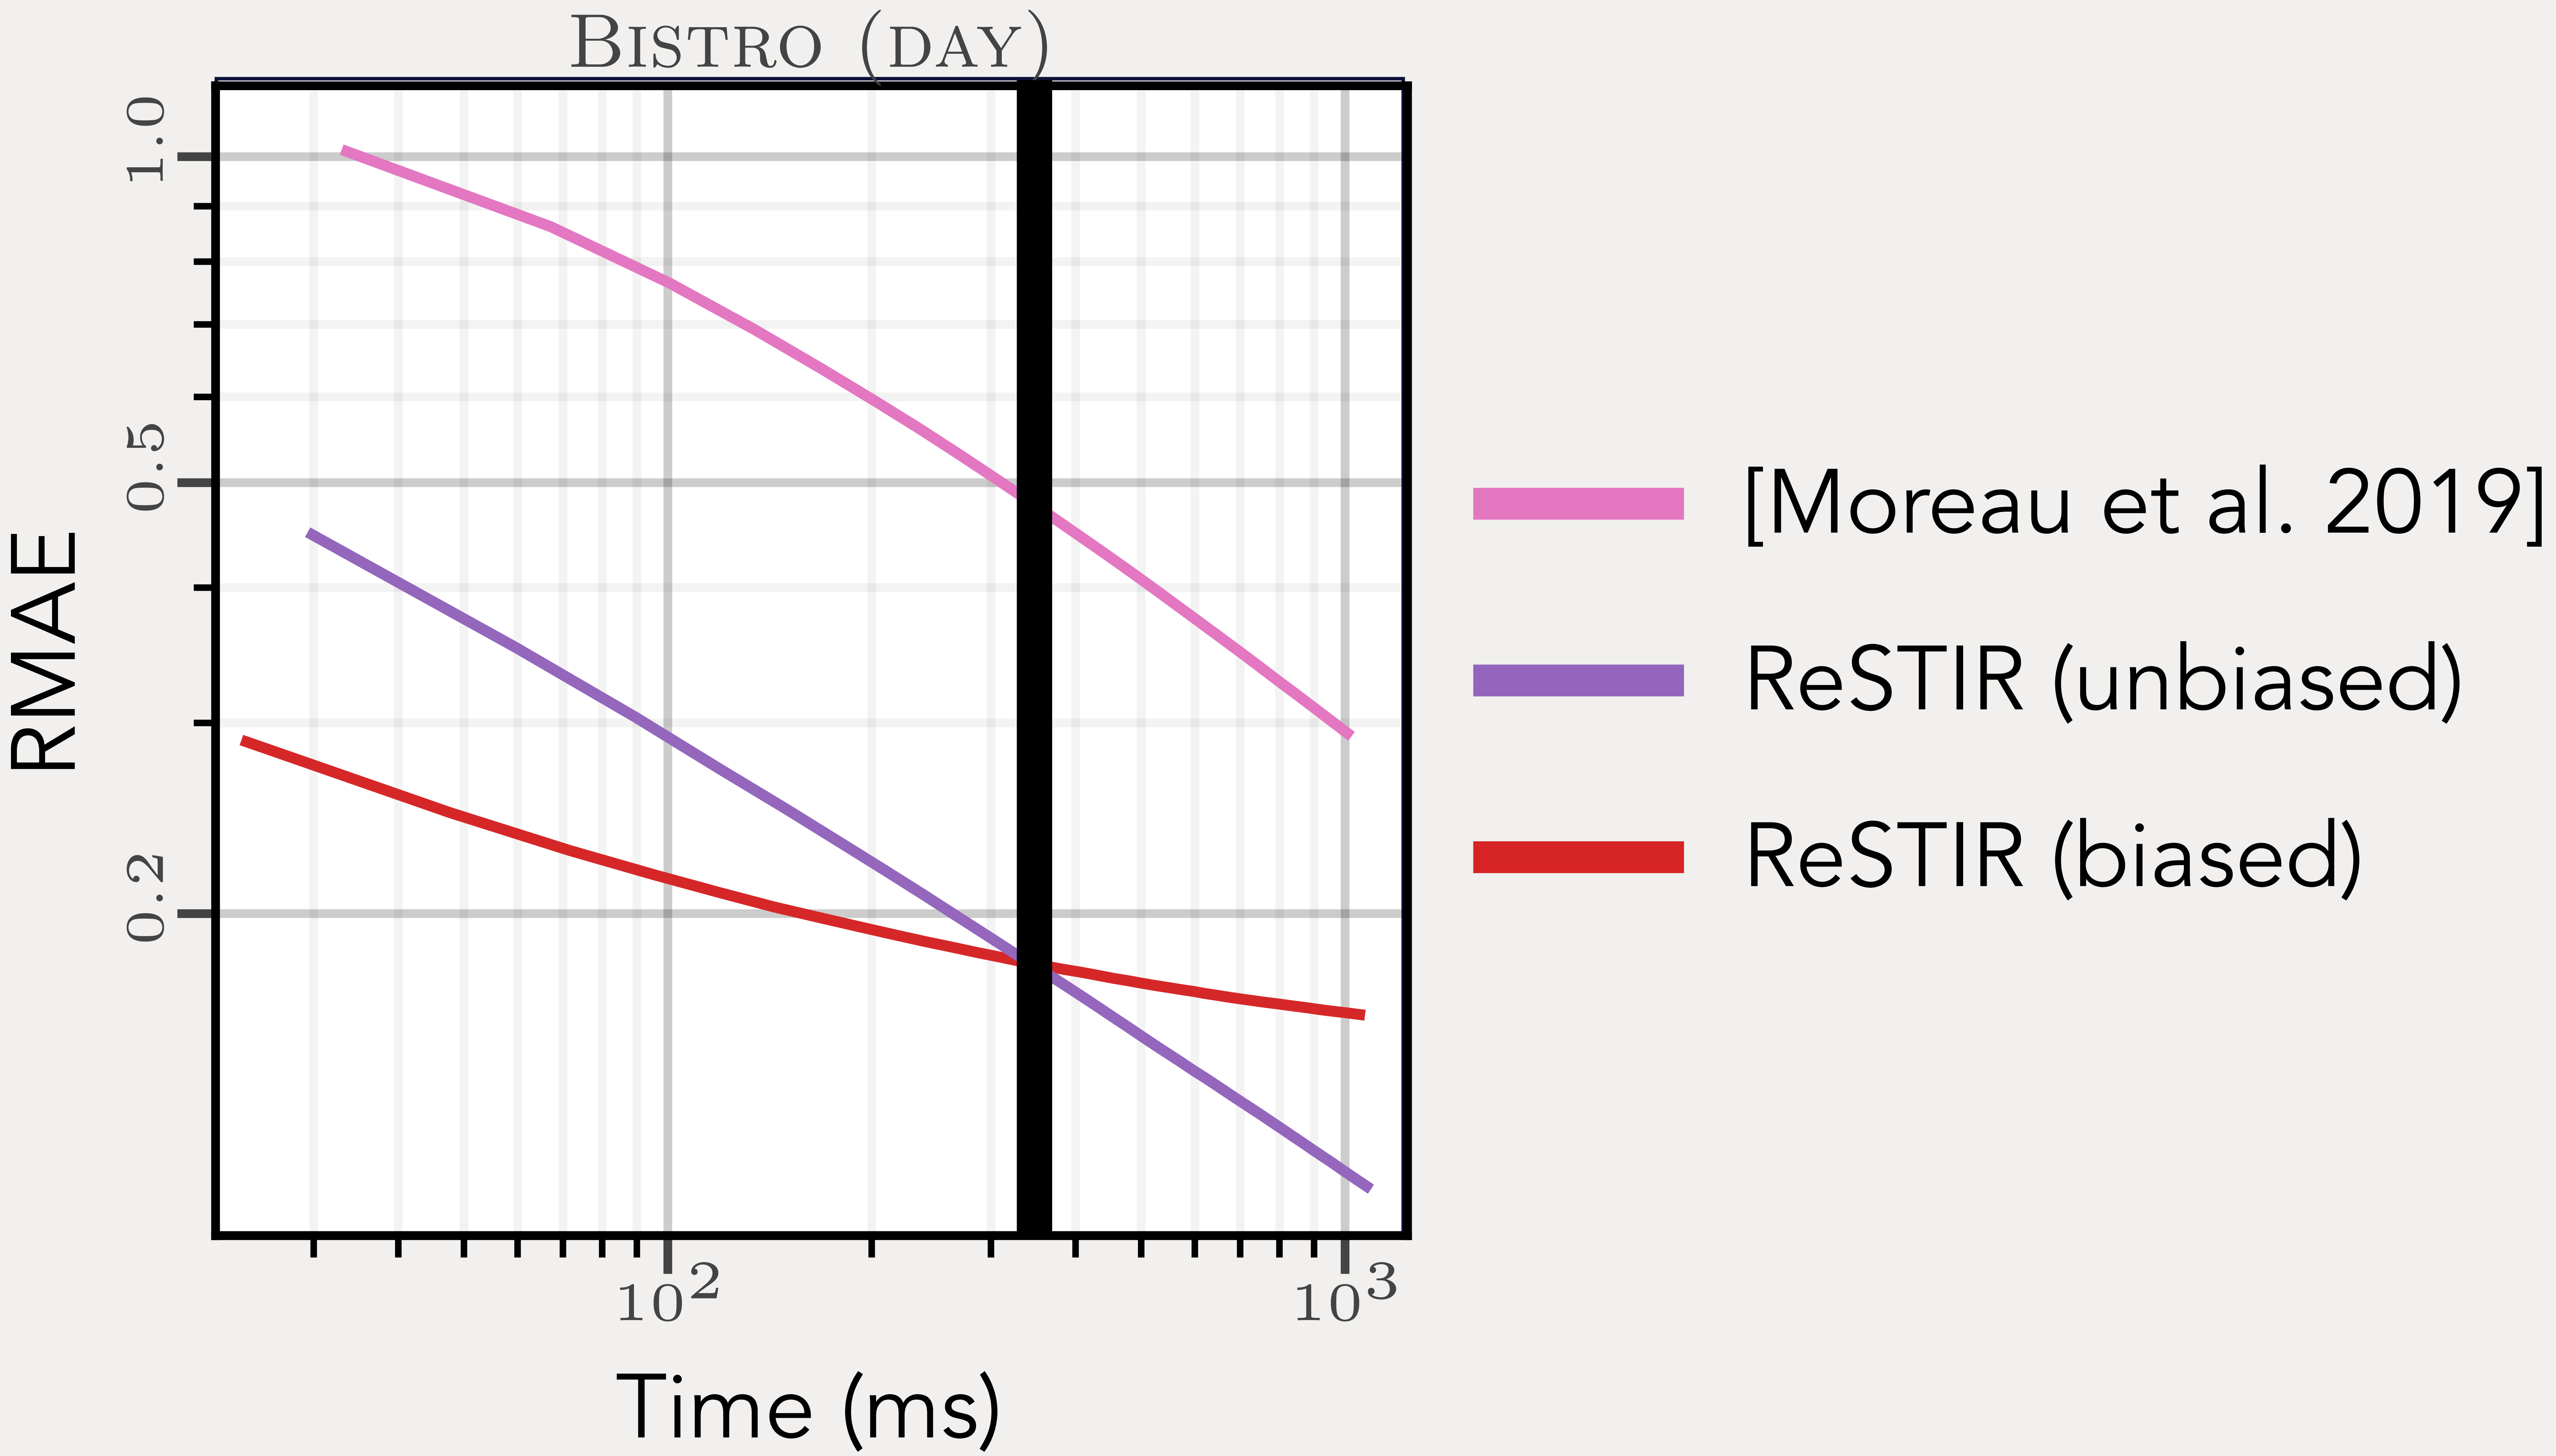

Conclusion

Conclusion

- Algorithm for sampling arbitrary distributions

Conclusion

- Algorithm for sampling arbitrary distributions
- Start with “bad” samples, improve them

Conclusion

- Algorithm for sampling arbitrary distributions
- Start with “bad” samples, improve them
- Reuse bad samples over multiple sample points

Limitations & Future Work

Limitations & Future Work

- Higher order bounces?

Limitations & Future Work

- Higher order bounces?
- World space?

Limitations & Future Work

- Higher order bounces?
- World space?
- How to best combine candidates from different distributions?

Limitations & Future Work

- Higher order bounces?
- World space?
- How to best combine candidates from different distributions?
- Disocclusions, visibility boundaries: Reuse limited

Limitations & Future Work

- Higher order bounces?
- World space?
- How to best combine candidates from different distributions?
- Disocclusions, visibility boundaries: Reuse limited
- BRDF etc. too expensive: Proxy needs to be used


Le Petit Coin
Le Petit Coin
Restaurant
Le Petit Coin
MENU
MENU


Burger Hit
Hit Cola


LIBRAIRIE


LIBRAIRIE