Rendering: Multiple Importance Sampling

Adam Celarek

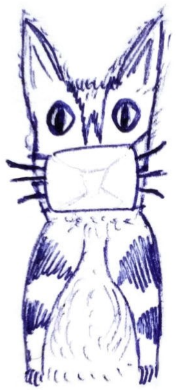

Research Division of Computer Graphics

Institute of Visual Computing & Human-Centered Technology

TU Wien, Austria

Overview

- Bad sampling
- Multiple Importance Sampling

source: own work

Sampling the light sources (128 samples)

glossy material

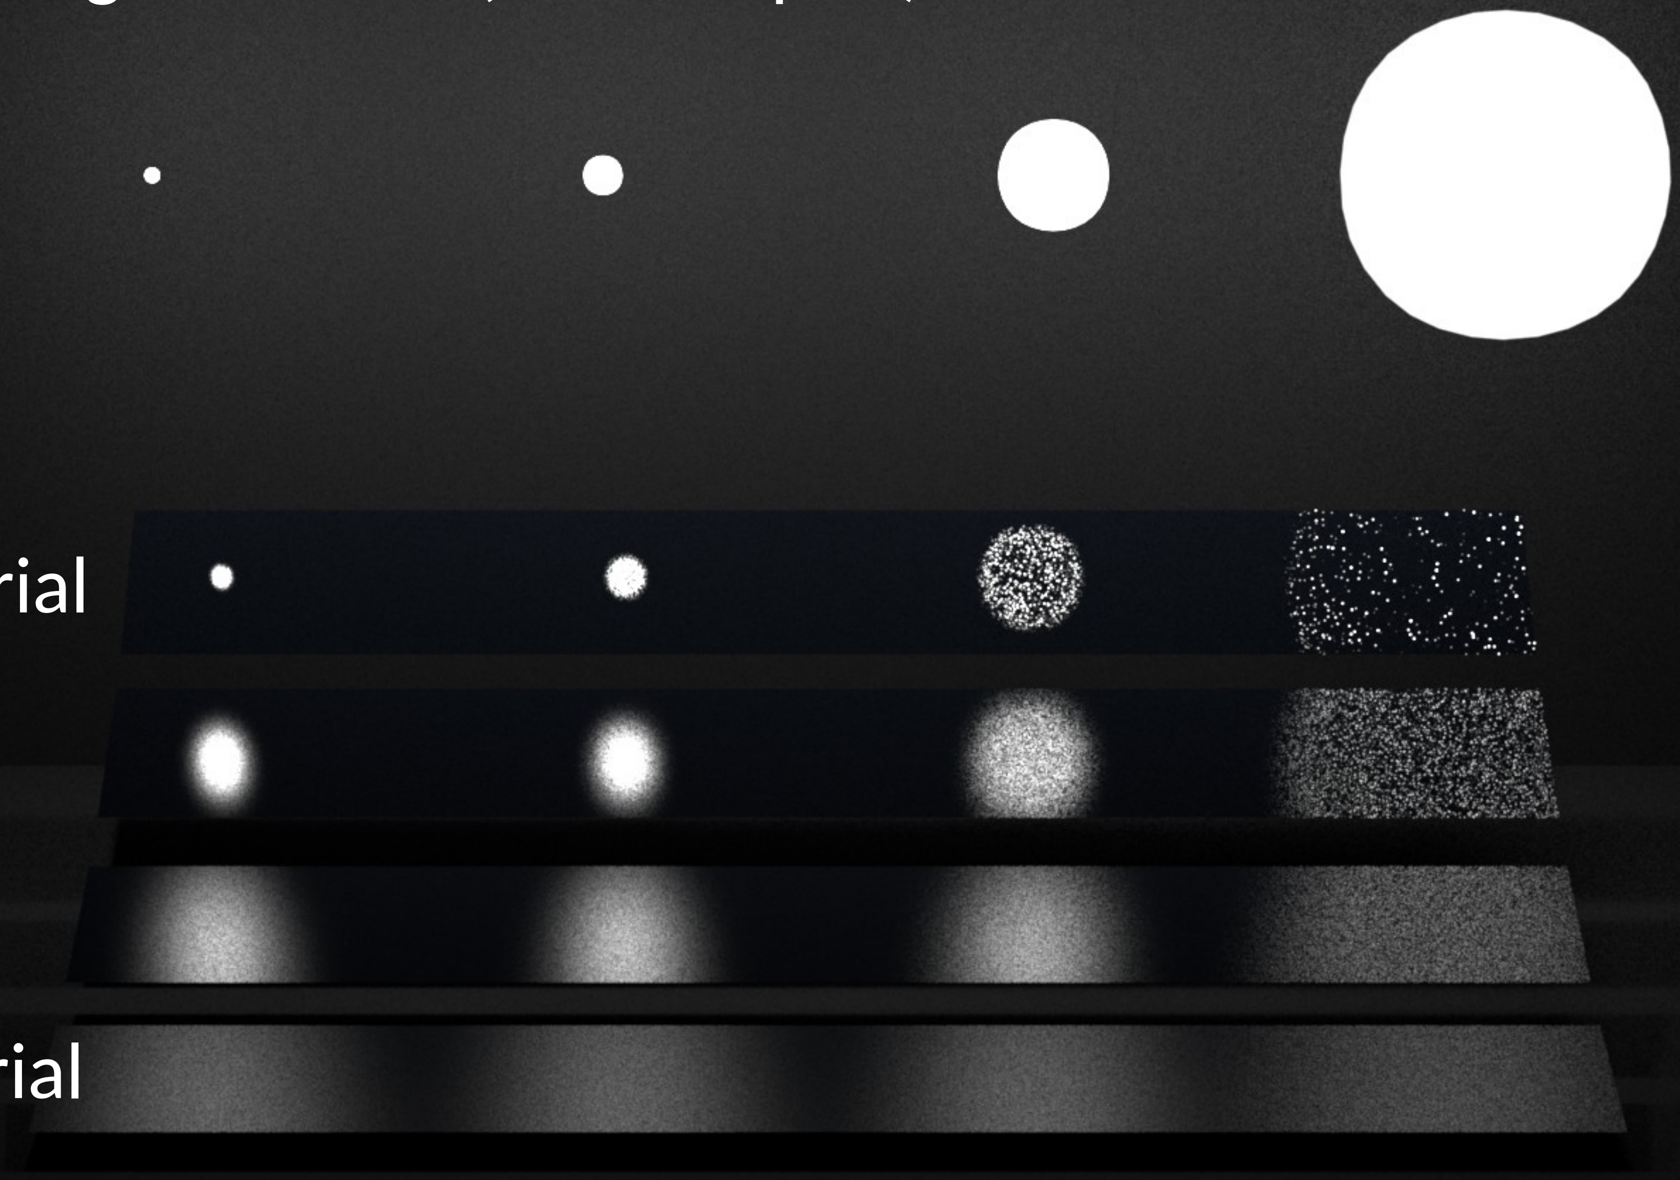

rough material

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Sampling the light sources (4096 samples)

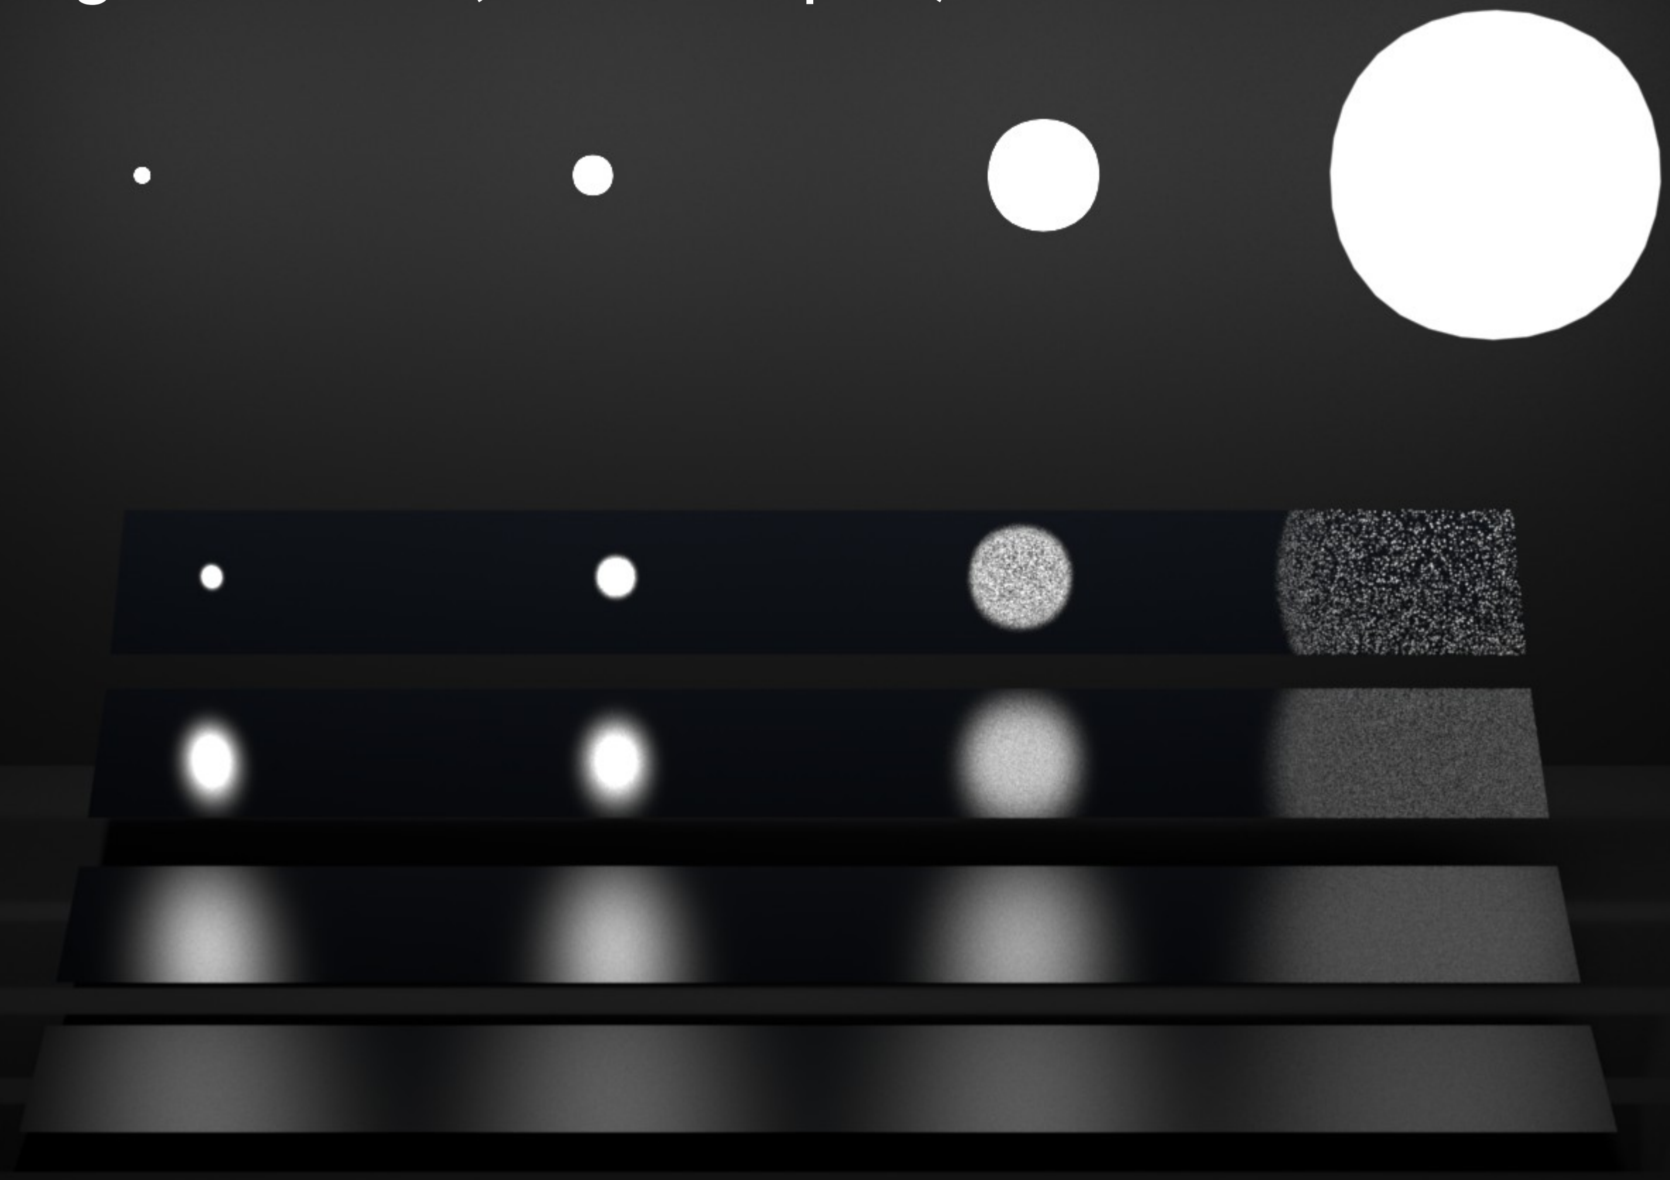

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Sampling the light sources (16384 samples)

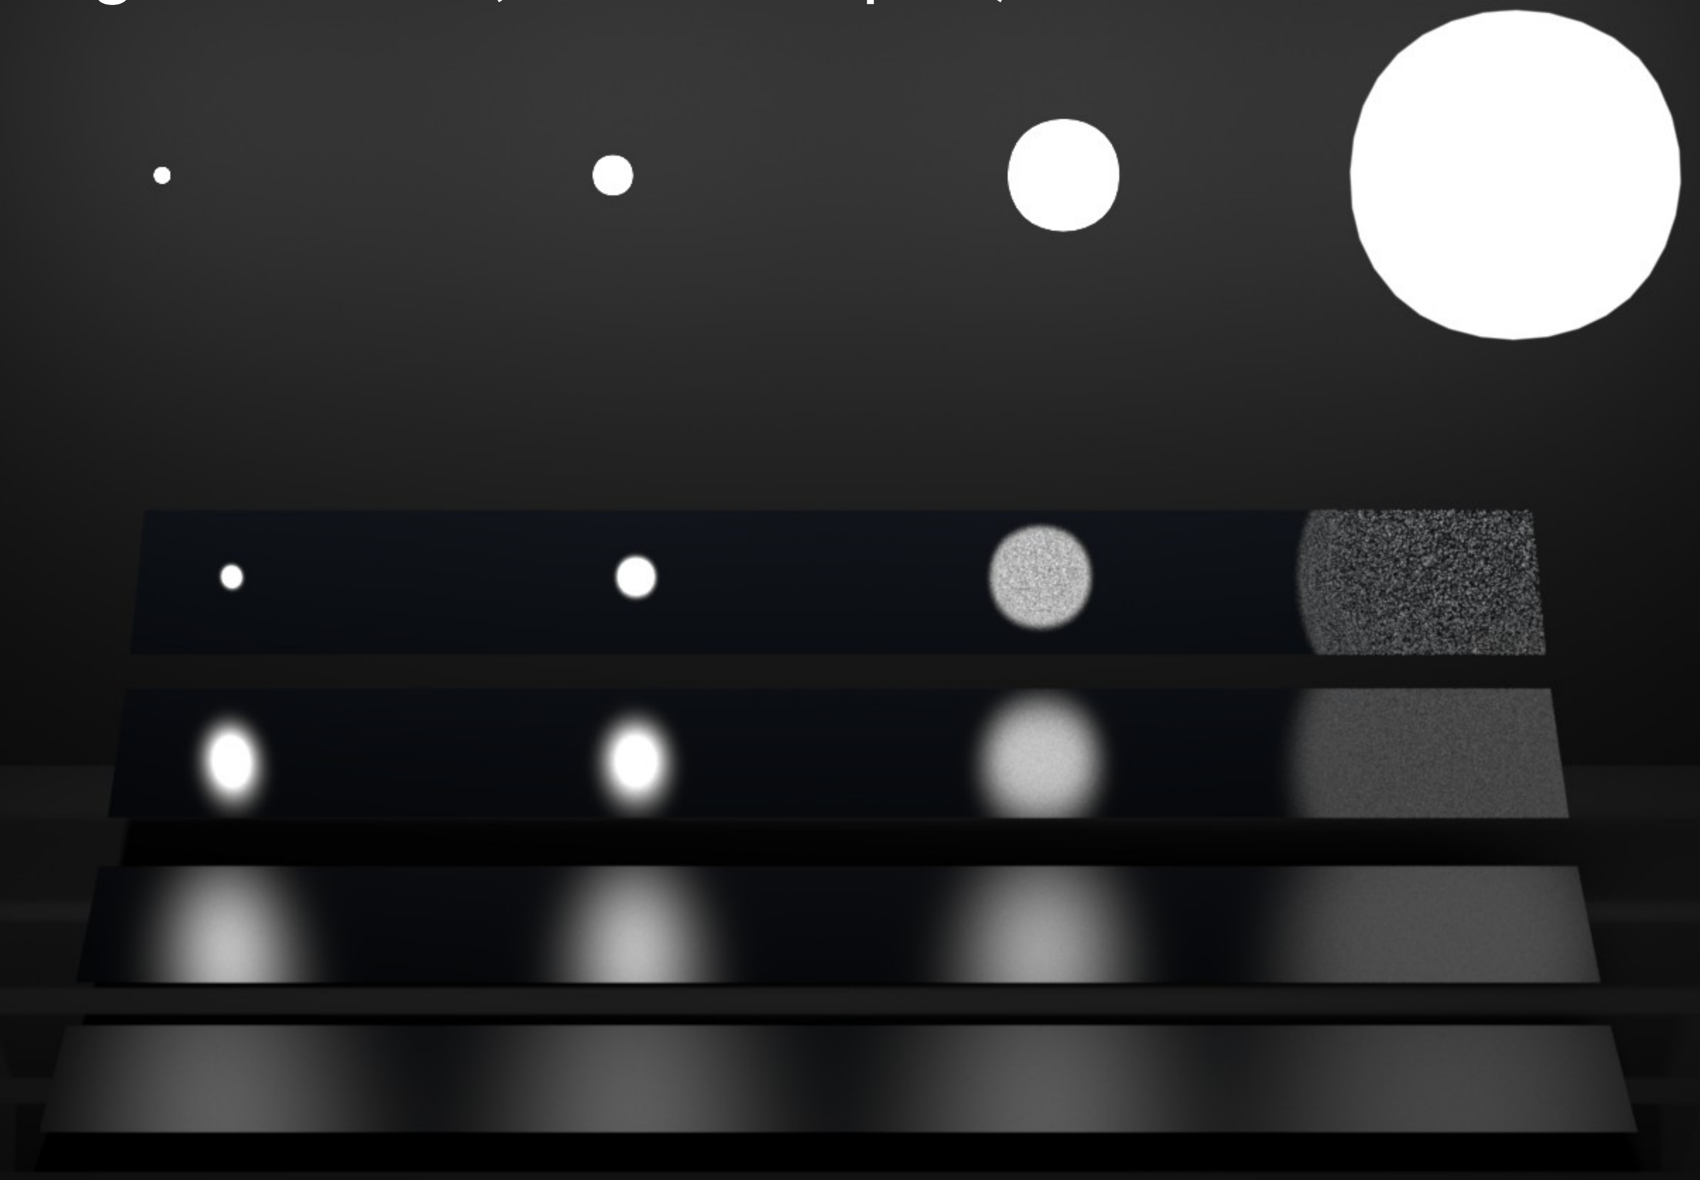

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Bad Sampling

- Clearly, we have a problem here. The sampling strategy has a hard time with the situation.

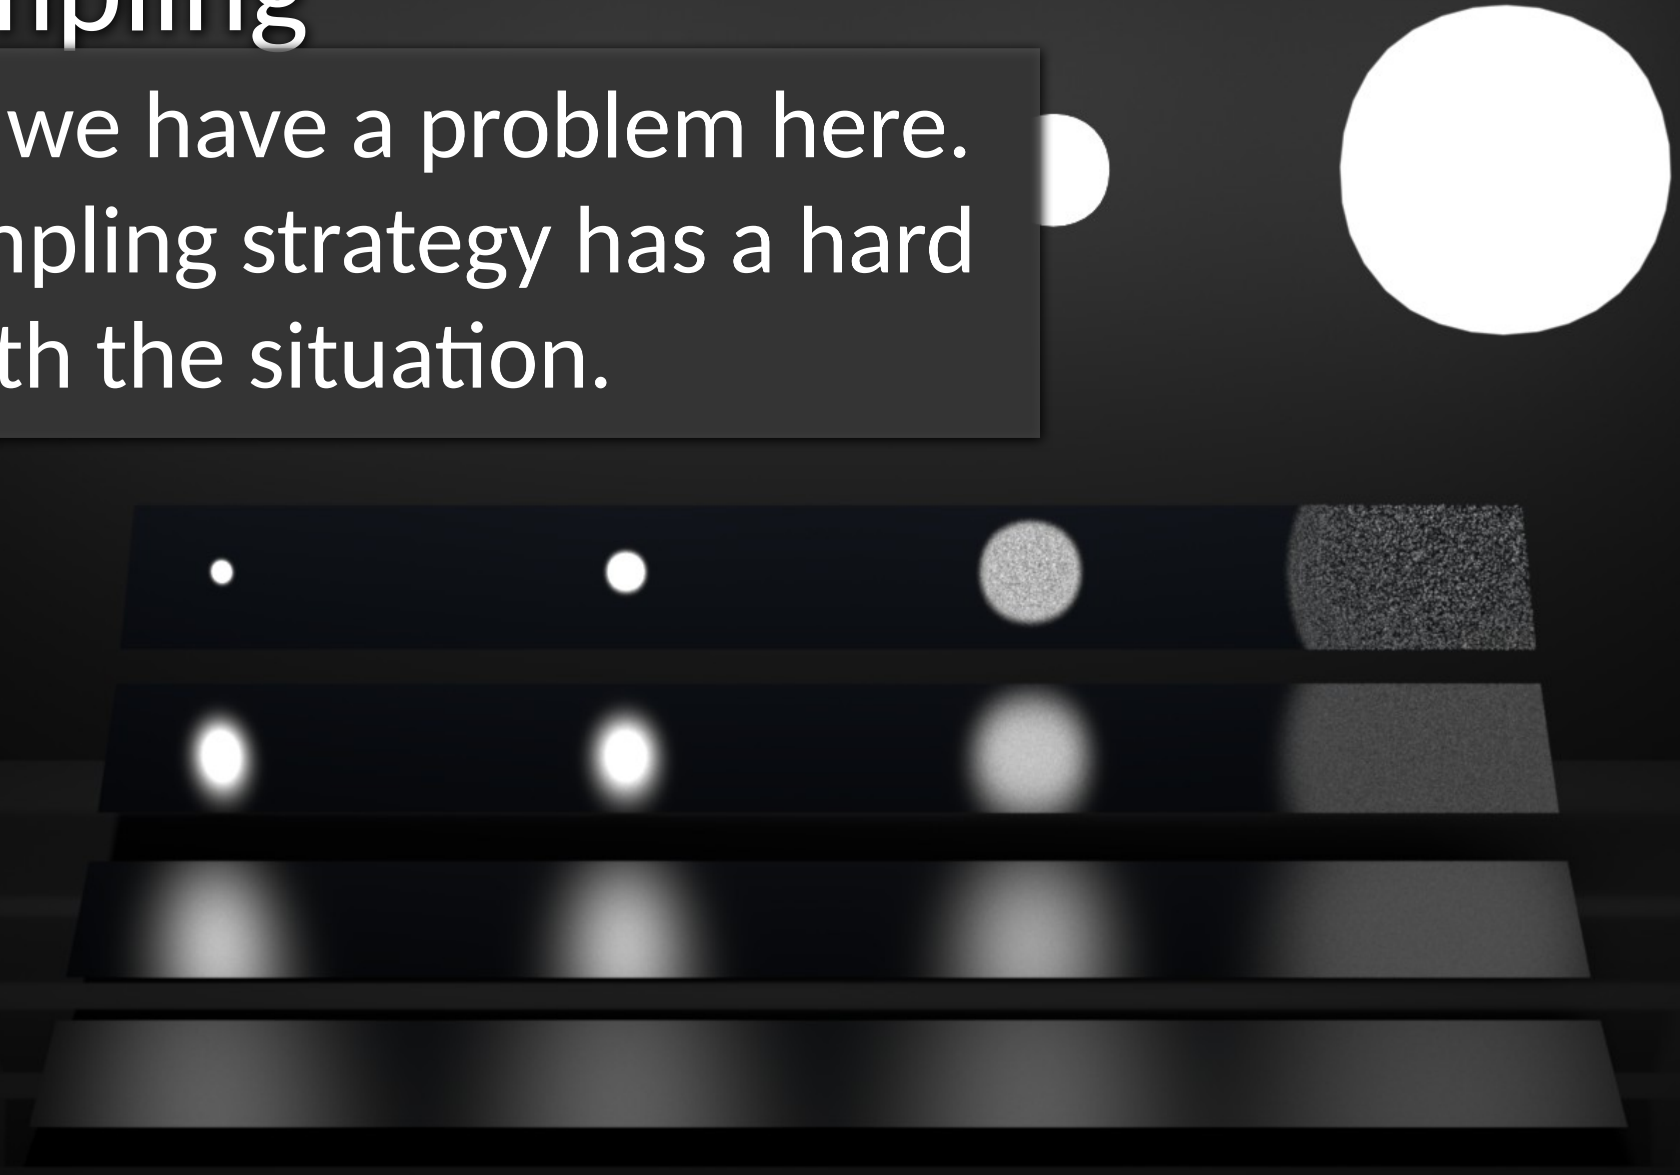

source: modified assignment scene rendered with Nori, based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Bad Sampling: What Happens

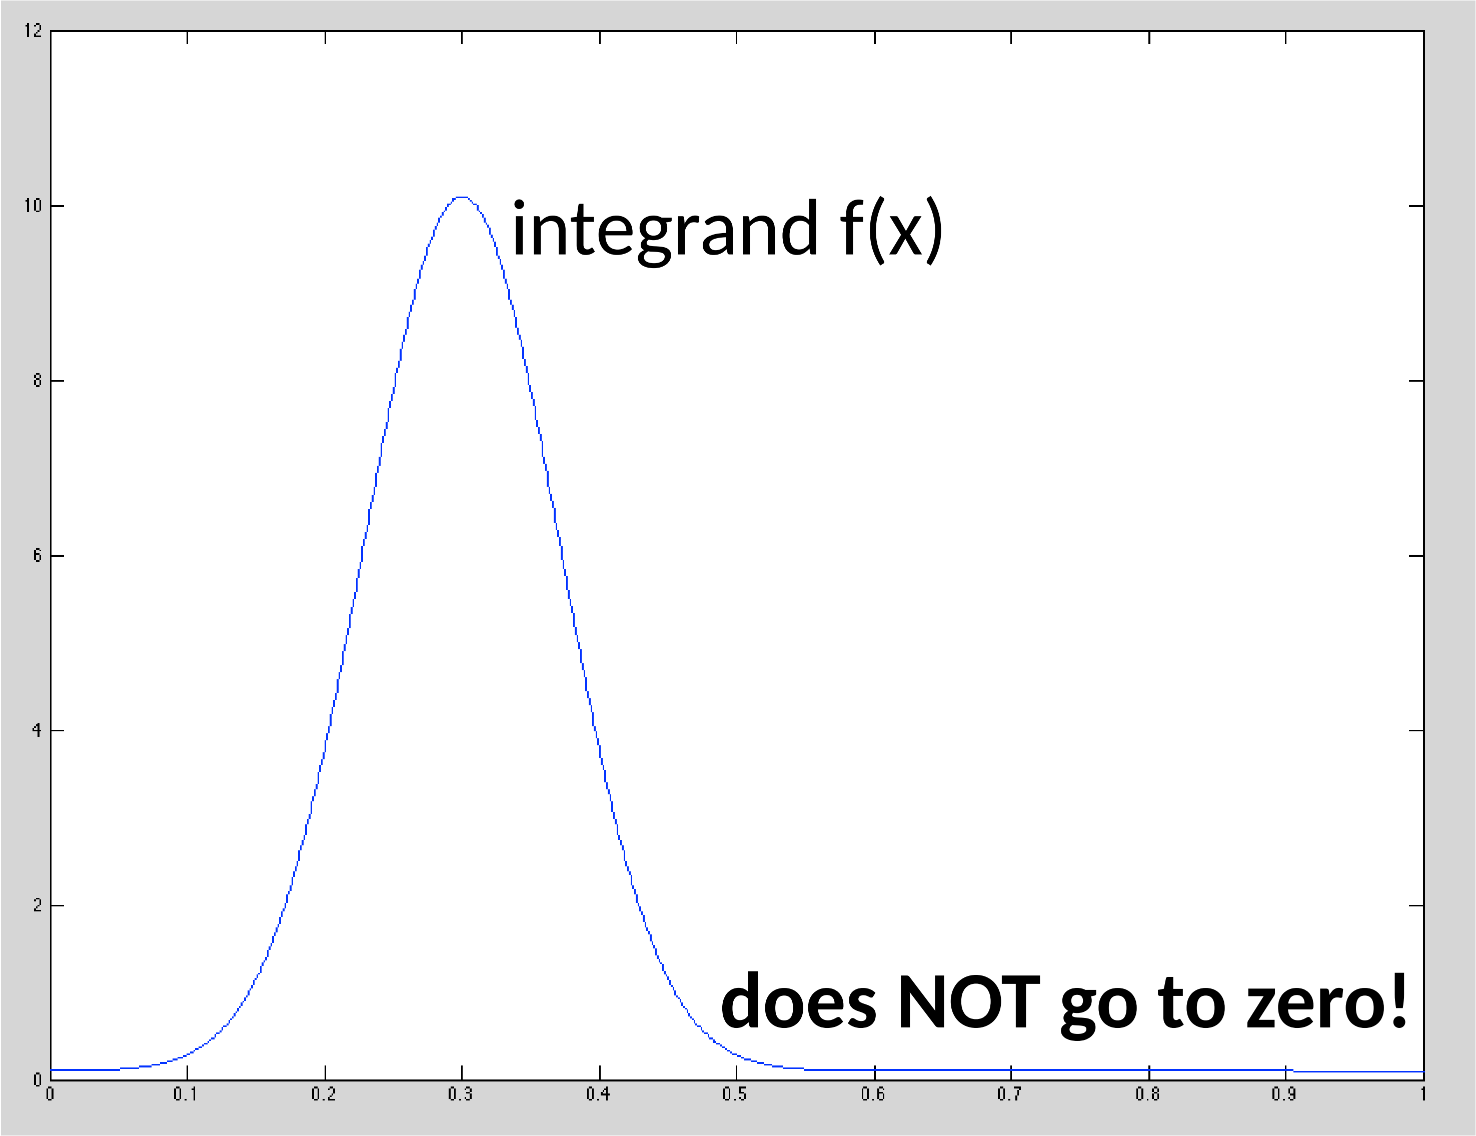

Slide modified from Jaakko Lehtinen, with permission

Bad Sampling: What Happens

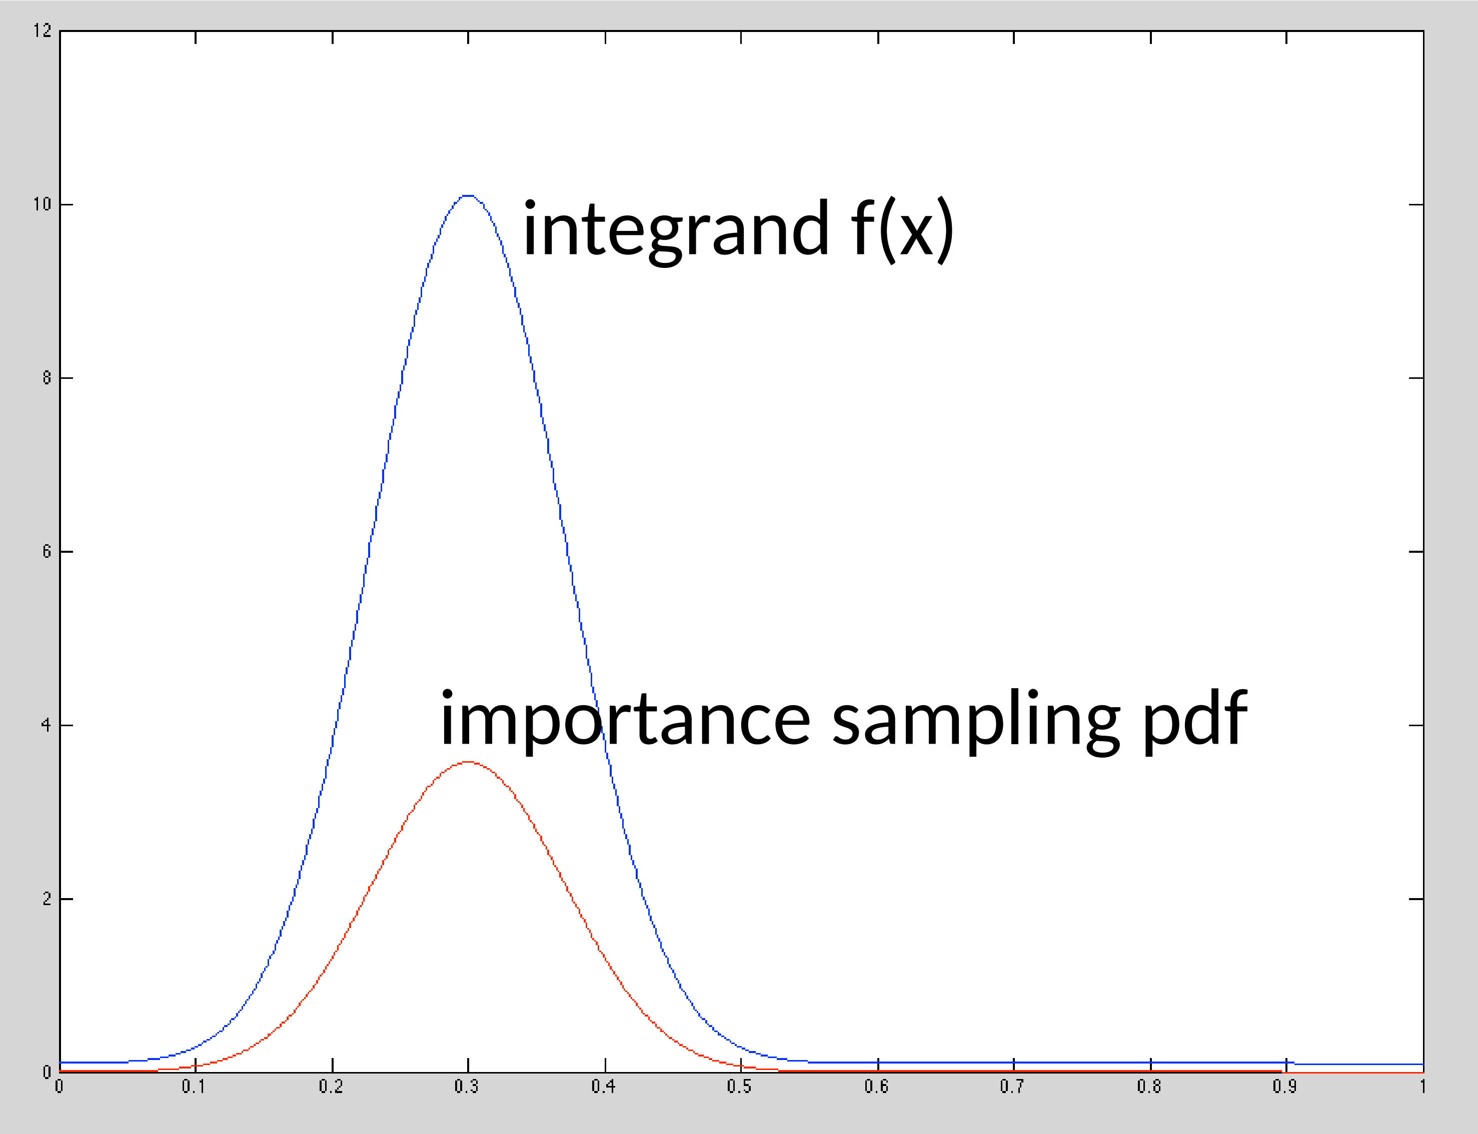

Slide modified from Jaakko Lehtinen, with permission

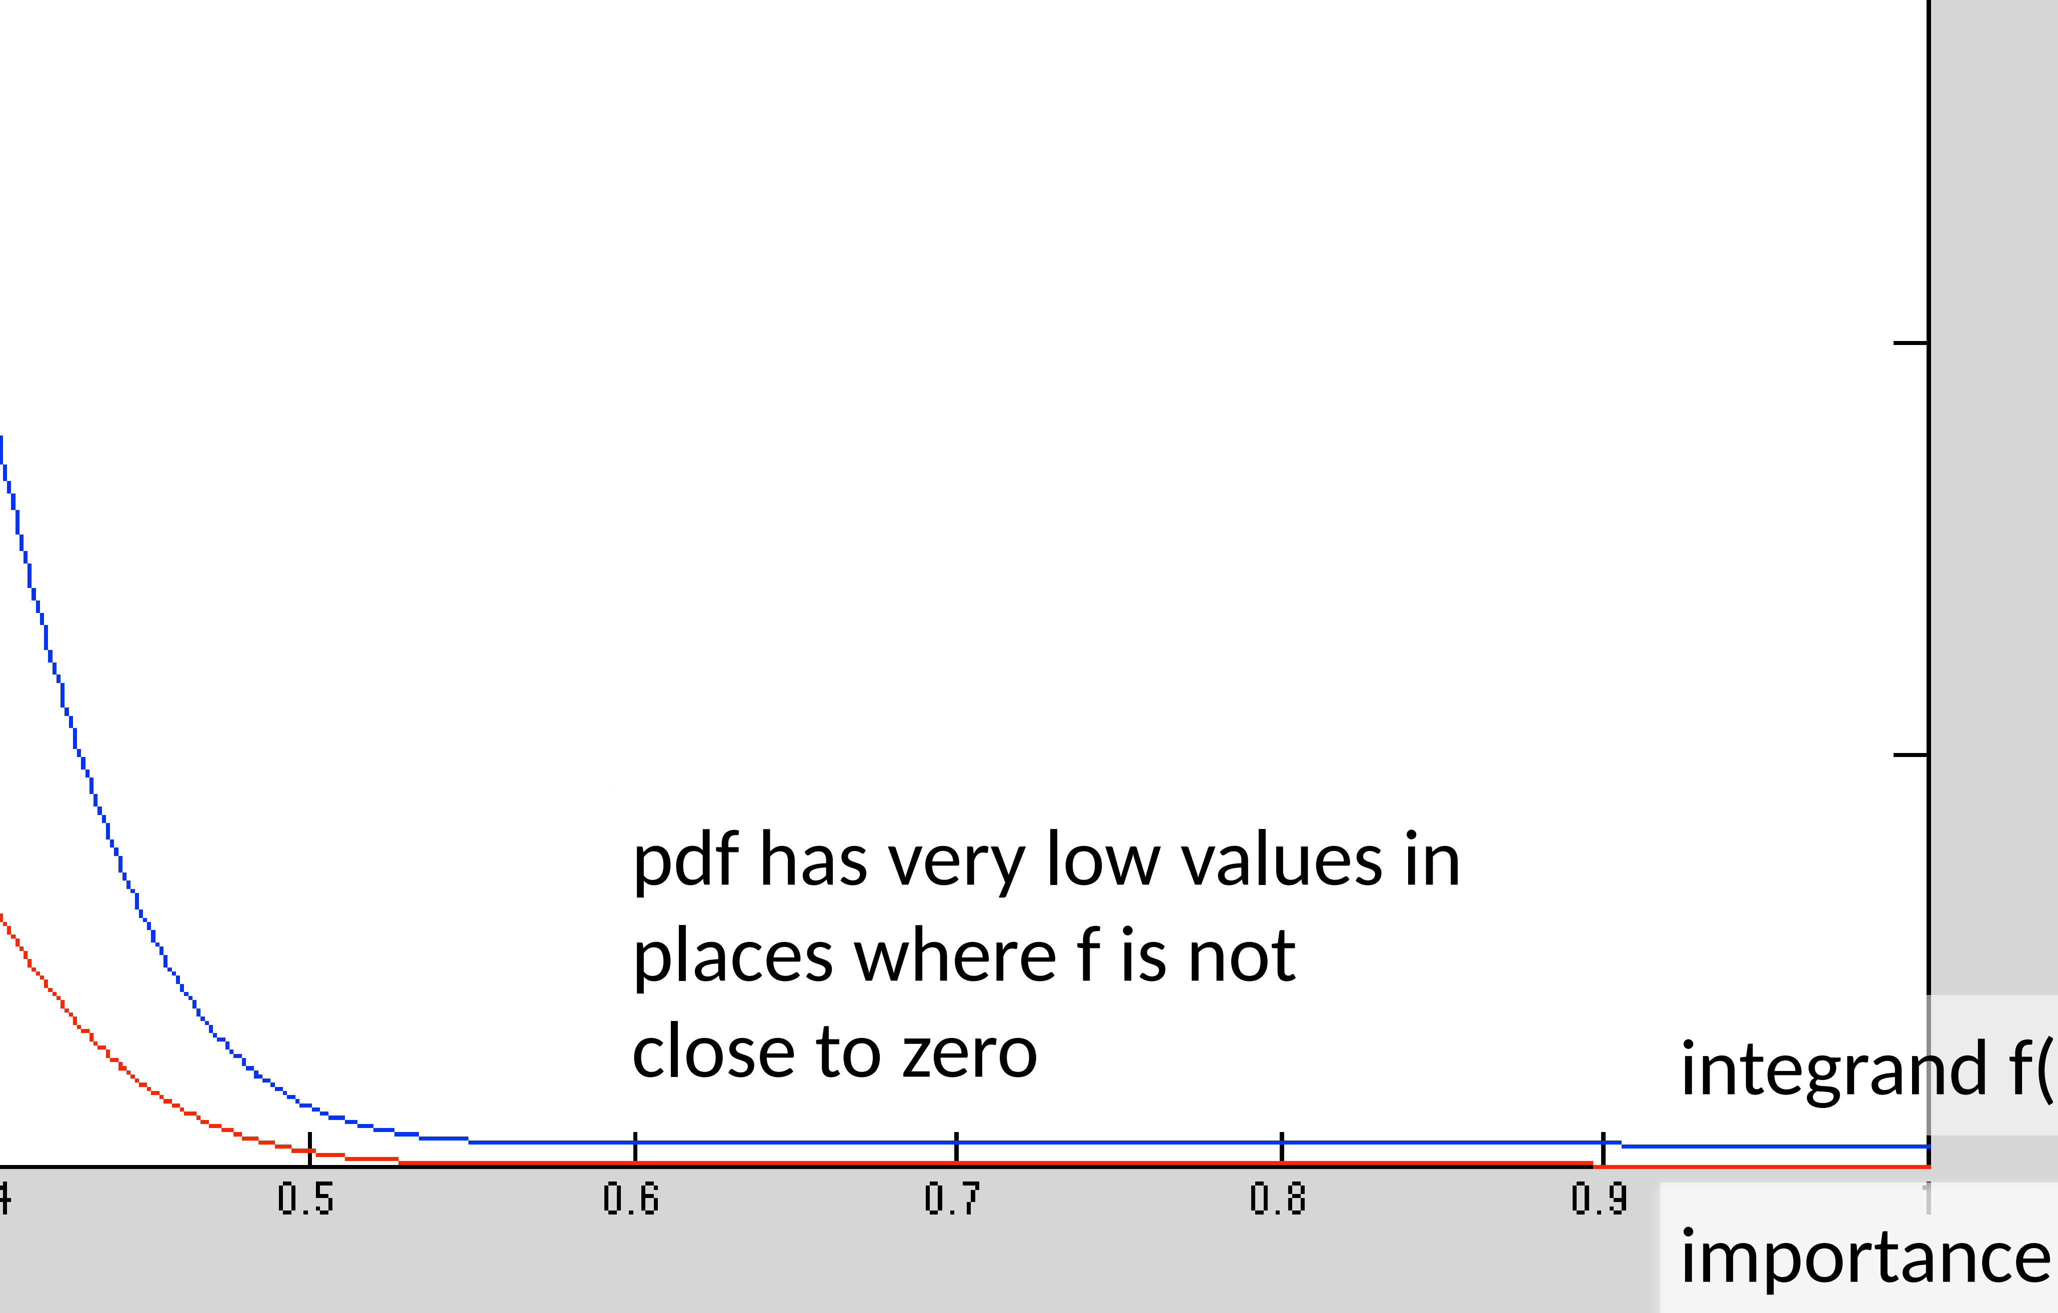

pdf has very low values in places where f is not close to zero

Slide modified from Jaakko Lehtinen, with permission

Bad Sampling: What Happens

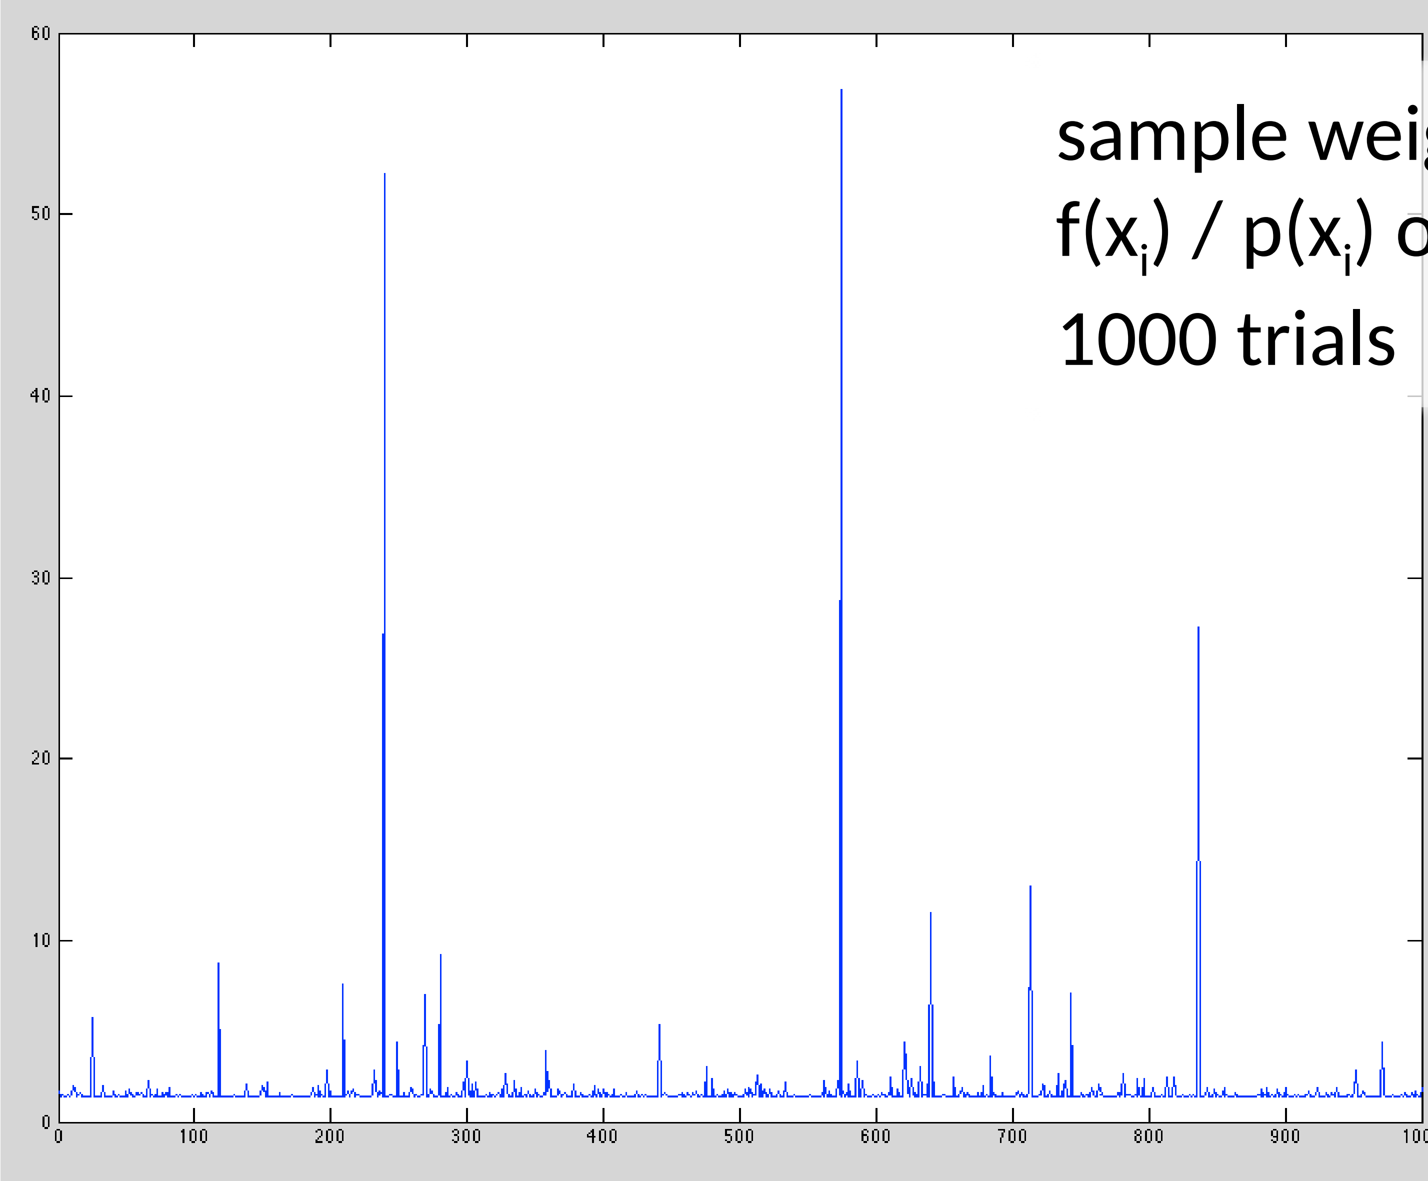

sample weight $f(x_i) / p(x_i)$ over 1000 trials

Slide modified from Jaakko Lehtinen, with permission

Bad Sampling: What Happens

spikes in cases where p(x) is very low, yet f(x) is never very low

in our example:
p(0.5) = 0.0027,
p(0.9) = 10^{-31}!

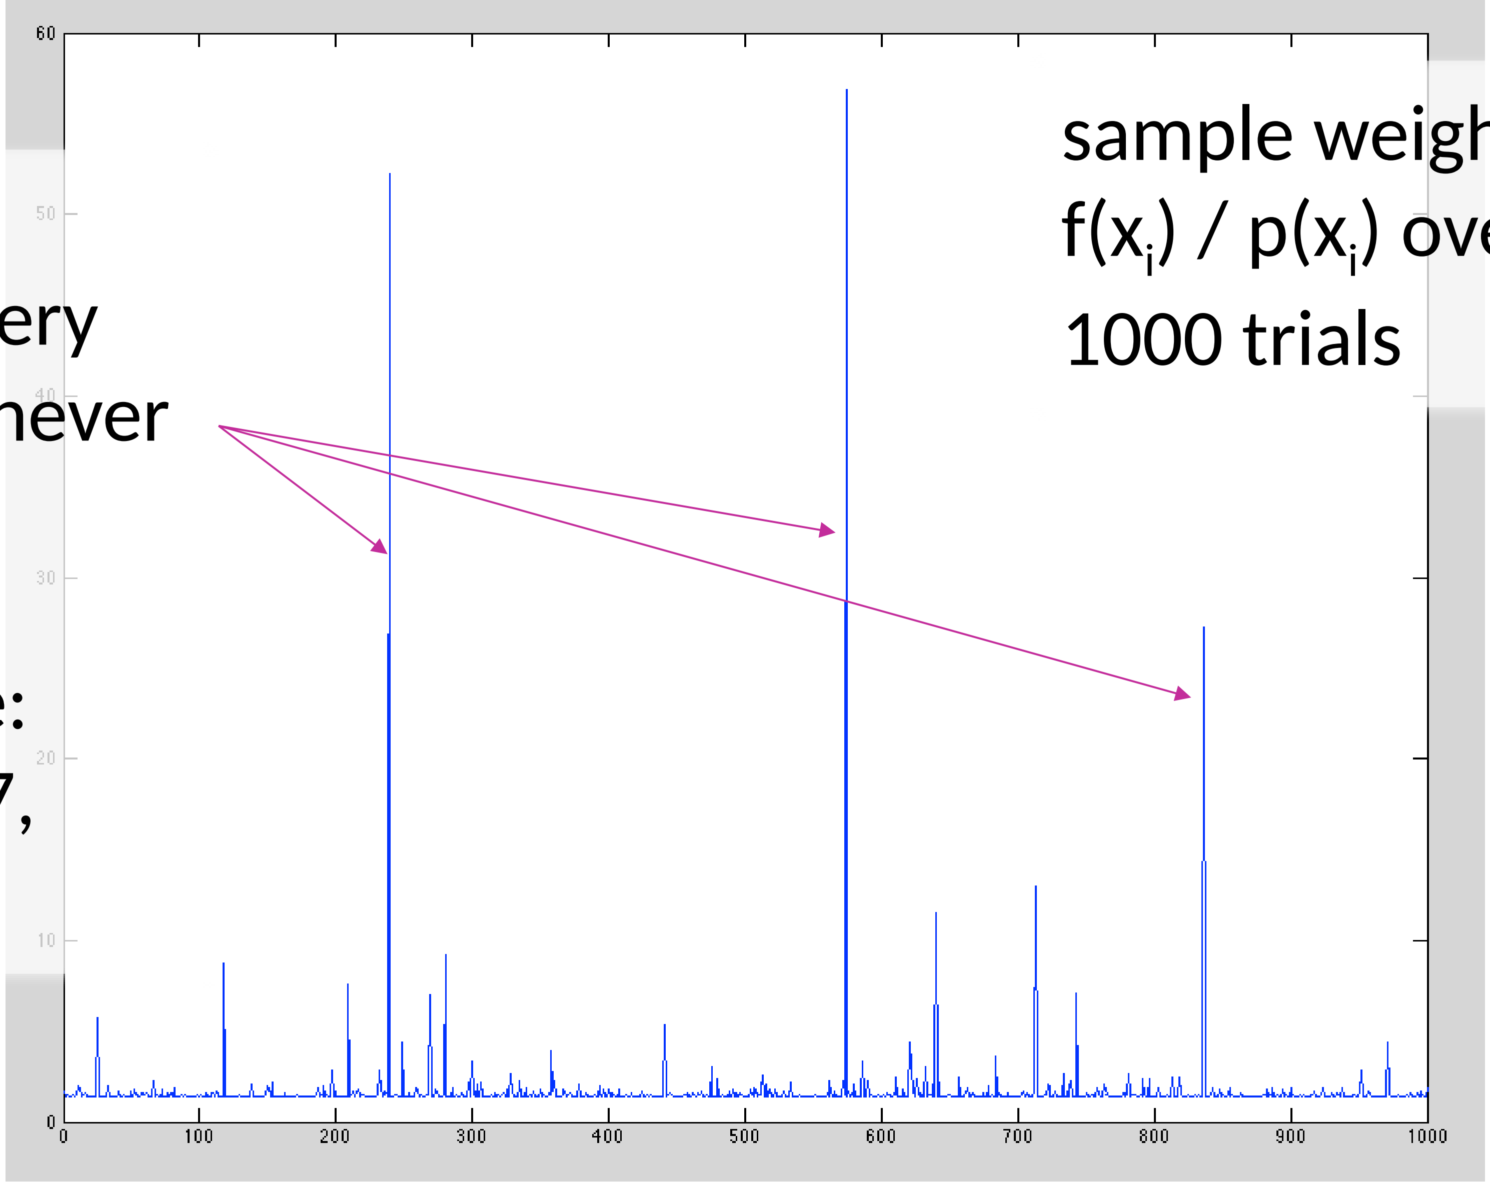

sample weight $f(x_i) / p(x_i)$ over 1000 trials

Slide modified from Jaakko Lehtinen, with permission

Bad Sampling: What Happens

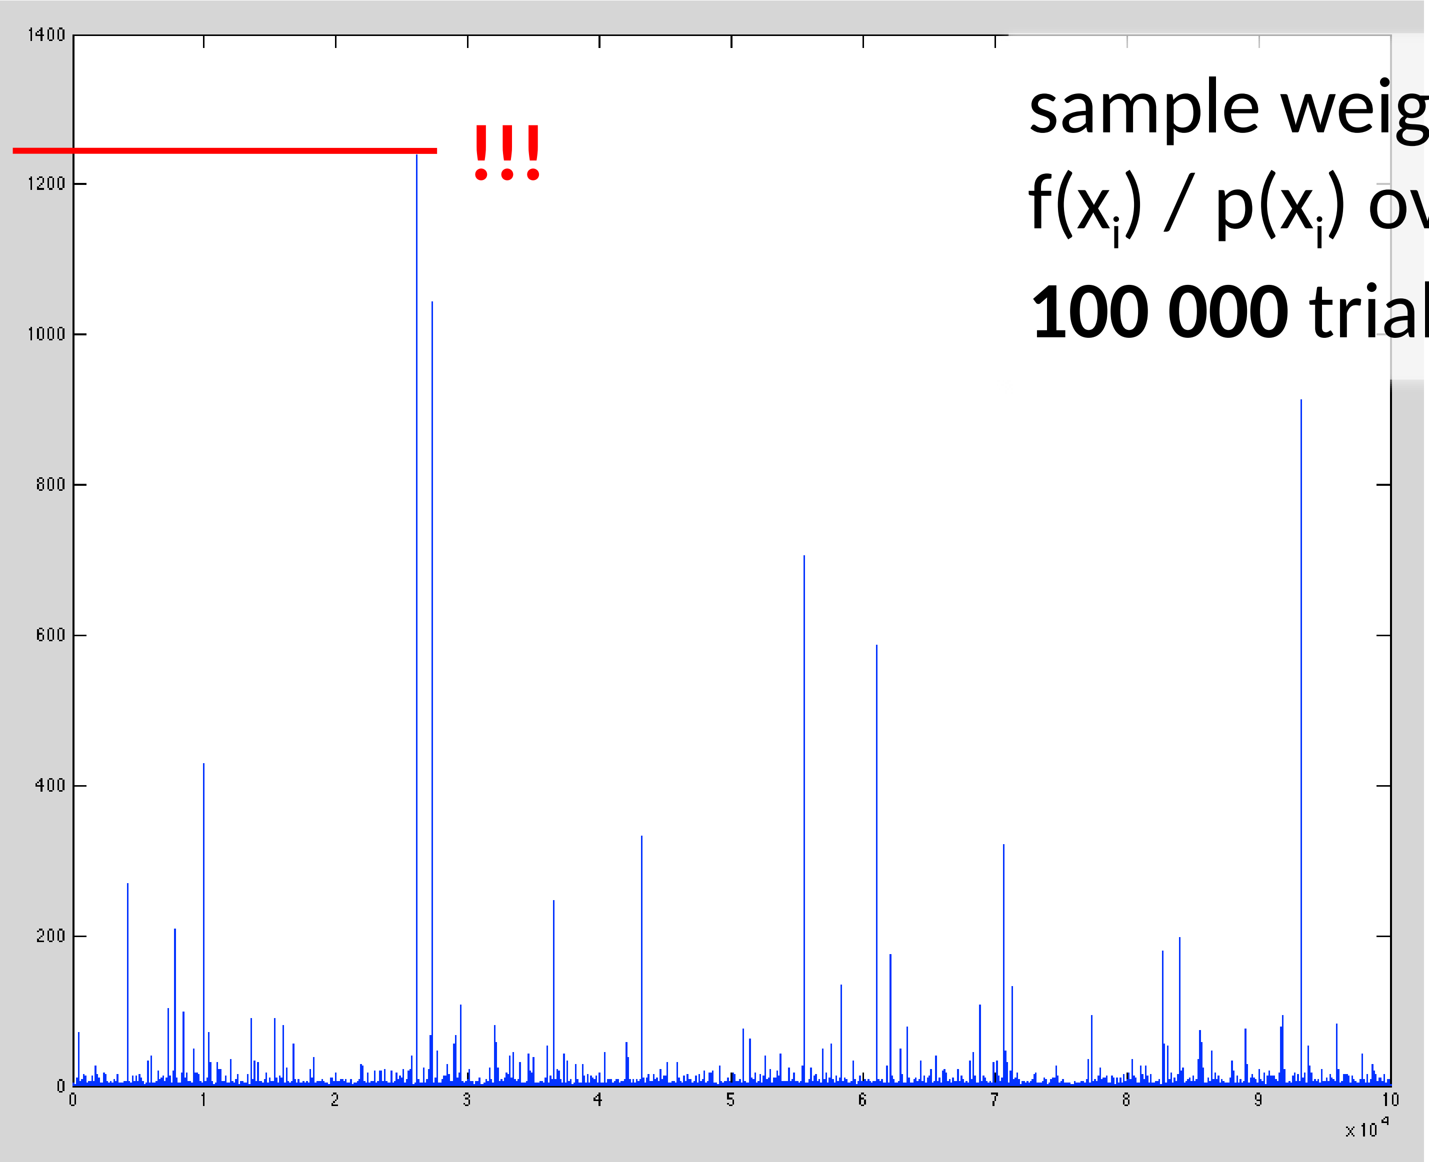

sample weight

$f(x_i) / p(x_i)$ over

100 000 trials

Slide modified from Jaakko Lehtinen, with permission

Bad Sampling: What Happens

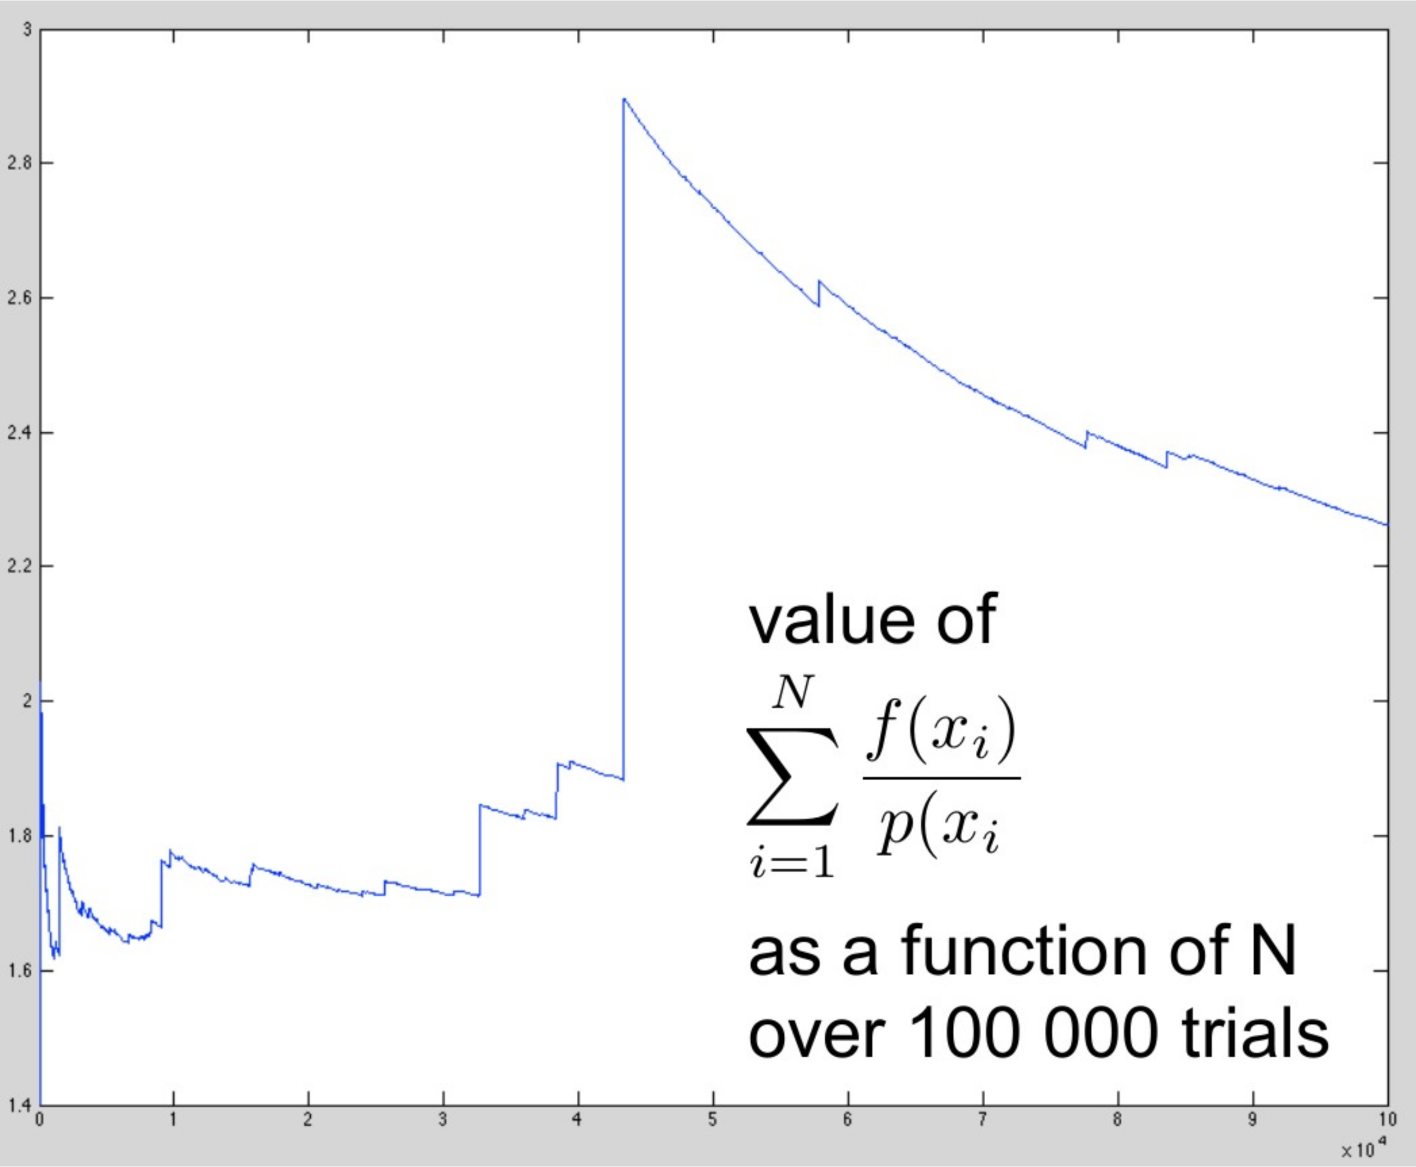

Slide modified from Jaakko Lehtinen, with permission

Bad sampling

When f(x) is large and p(x) small.

source: own work

Next: MIS

Multiple Importance Sampling

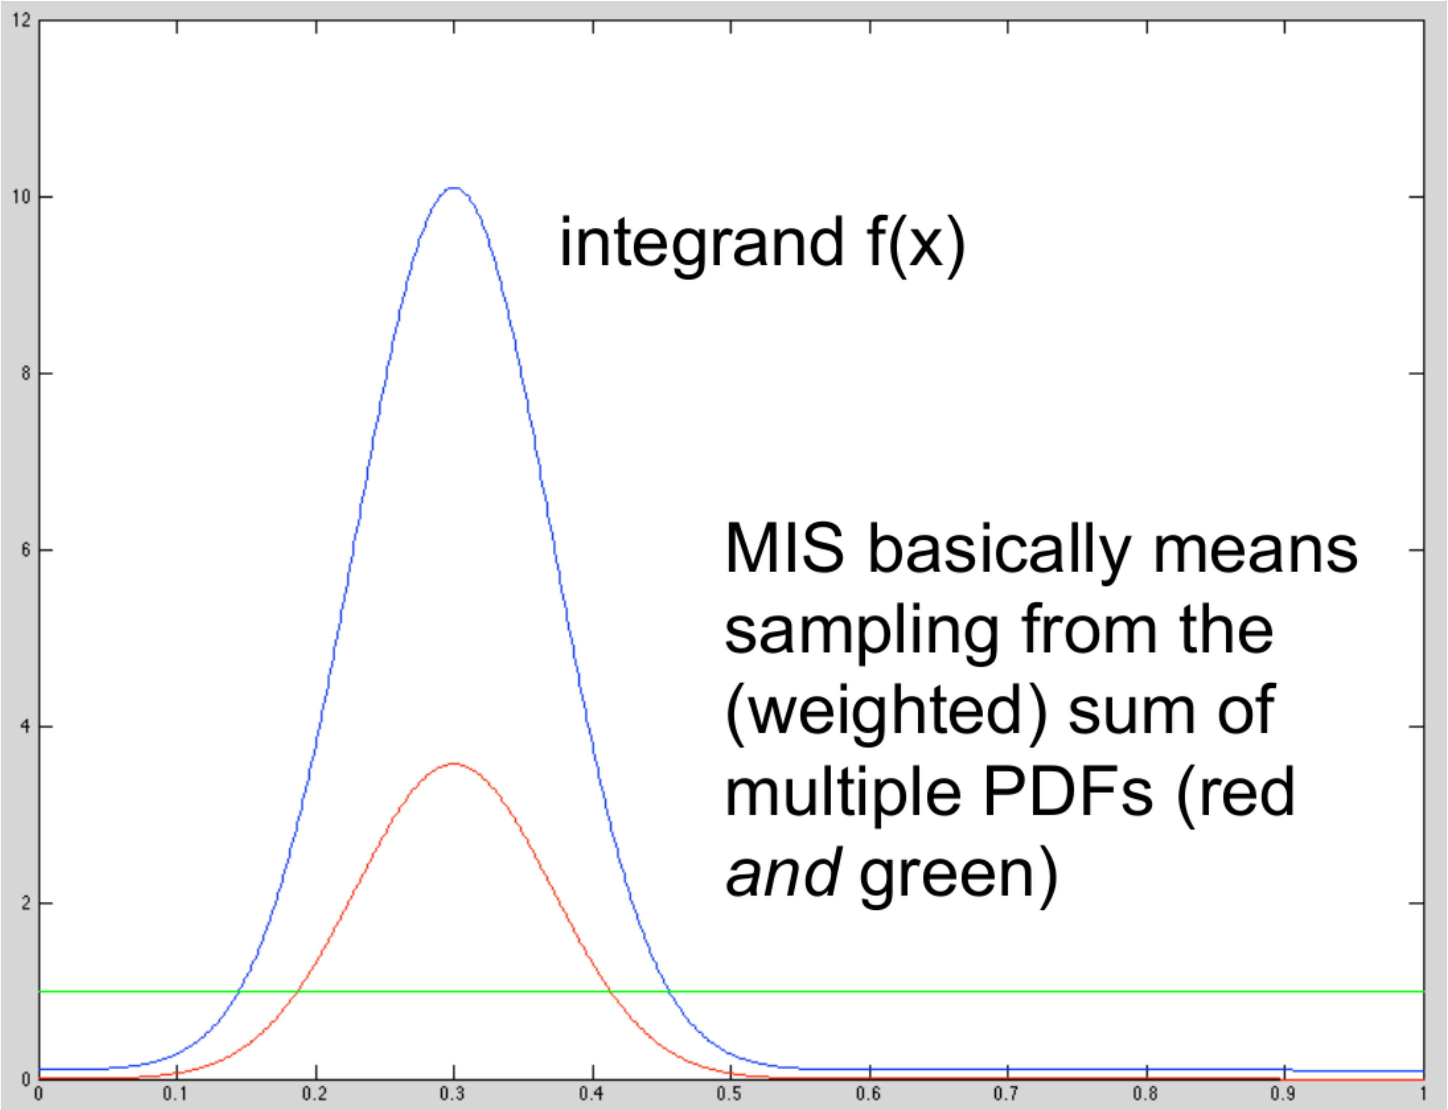

Slide modified from Jaakko Lehtinen, with permission

Sampling the light sources (128 samples)

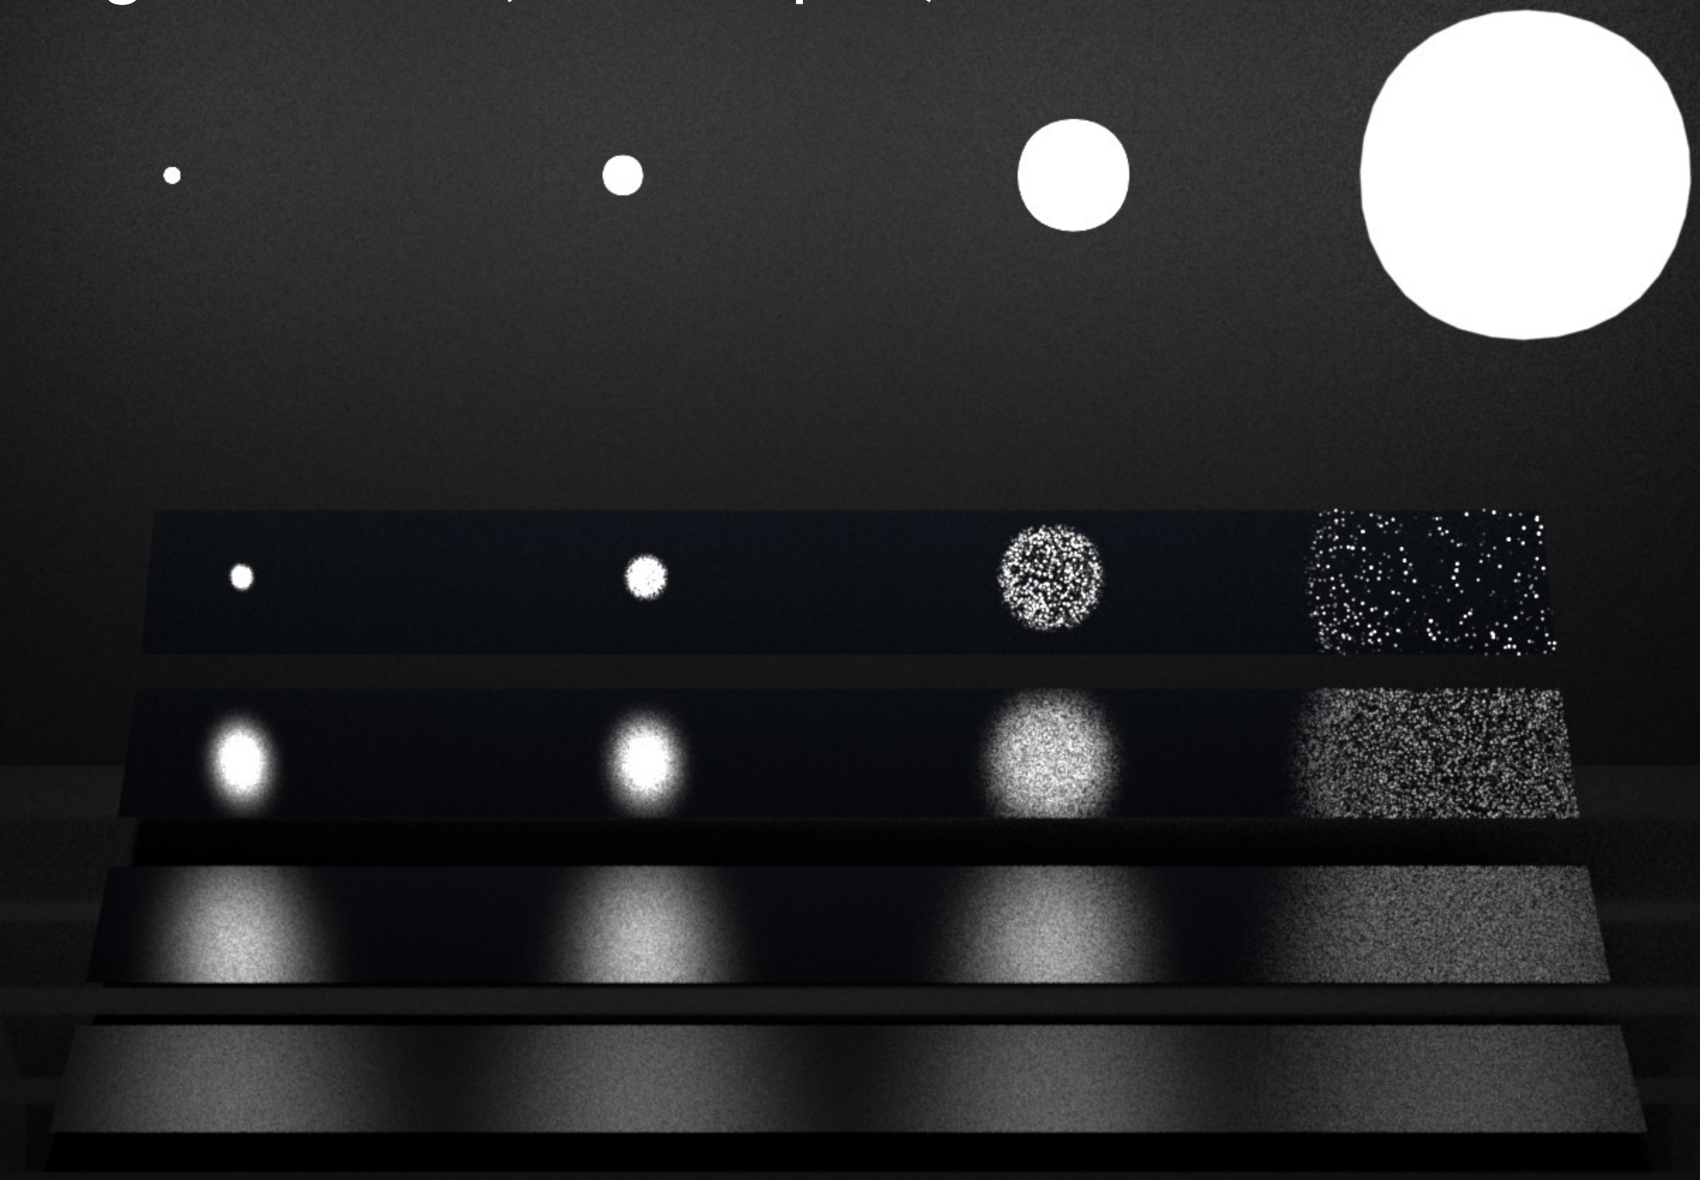

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Sampling the material (128 samples)

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Multiple Importance Sampling (128 samples)

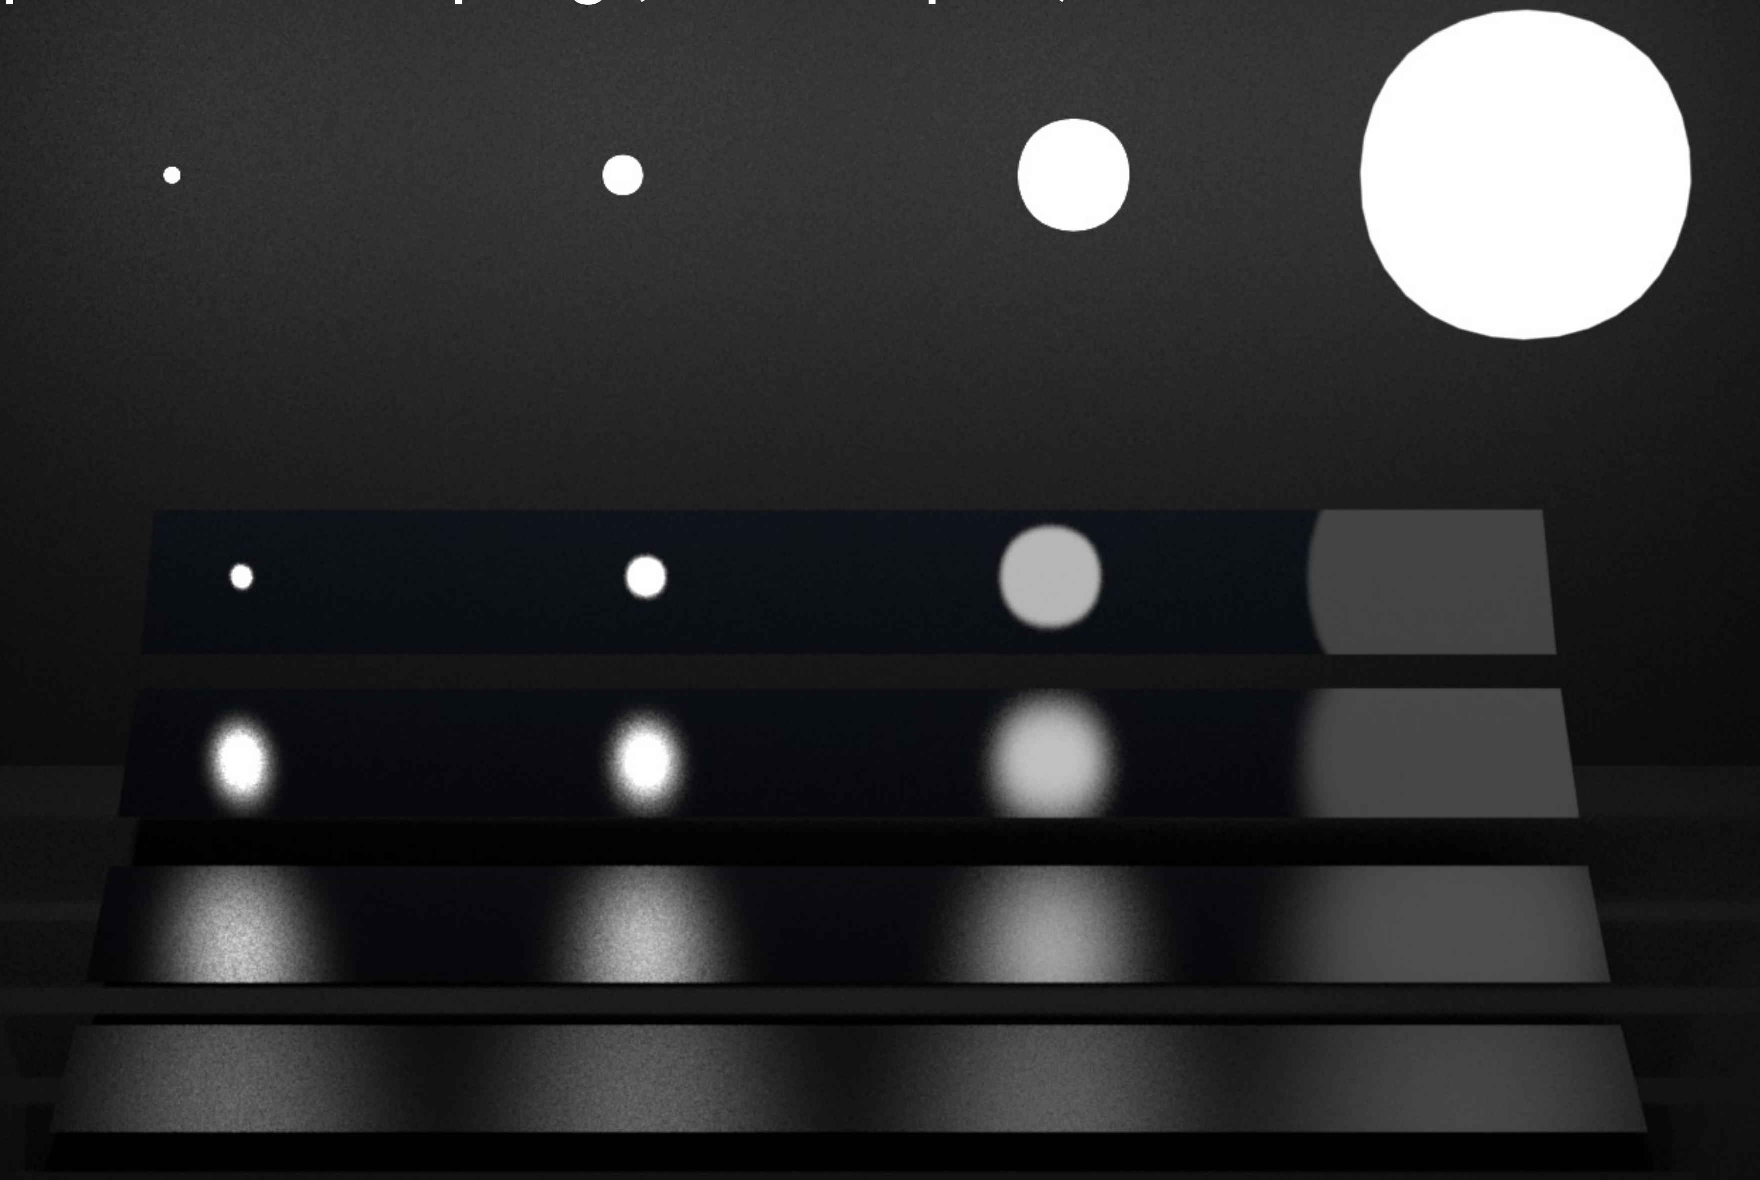

source: modified assignment scene rendered with Nori,
based on the MIS test scene by Eric Veach, modeled after a file by Steve Marschner

Multiple Importance Sampling: Math

- Let's start with plain Monte Carlo (what we already know)
- We have n estimators F_i and n_i samples each

$$F_i = \frac{1}{n_i} \sum_{j=0}^{n_i} \frac{f(X_j)}{p(X_j)}$$

- The expectation of all estimators is the integral

$$E[F_i] = \int_\Omega f(x)\,\mathrm{d}x$$

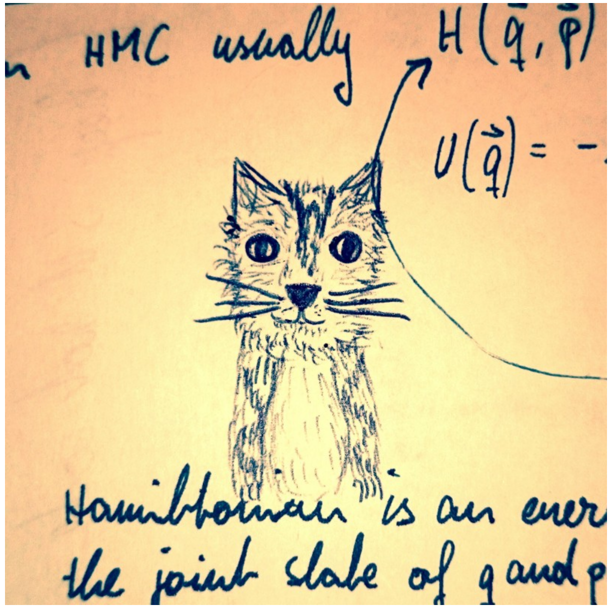

Multiple Importance Sampling: Math

Now, when we take the average, of these estimators

$$F = \frac{1}{n}\sum_{i=0}^{n} F_i$$

we again get an unbiased estimator

$$E[F] = \frac{1}{n}\sum_{i=0}^{n} E[F_i] = \int_{\Omega} f(x)\,\mathrm{d}x$$

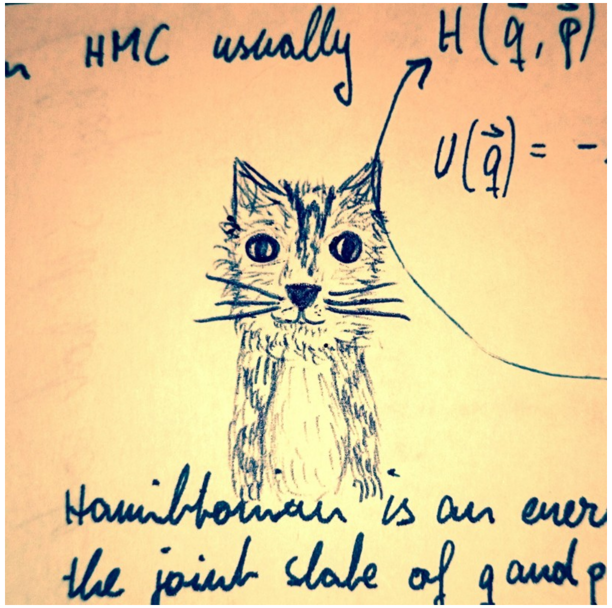

Multiple Importance Sampling: Math

Instead of a simple average, we can also take a weighted sum

$$E[F] = \sum_{i=0}^{n} w_i E[F_i] = \sum_{i=0}^{n} \frac{1}{n_i} \sum_{j=0}^{n_i} w_i E\left[\frac{f(X_{i,j})}{p(X_{i,j})}\right] = \int_{\Omega} f(x)\,\mathrm{d}x \text{ with } \sum w_i = 1$$

and move the weight into the estimators F_i

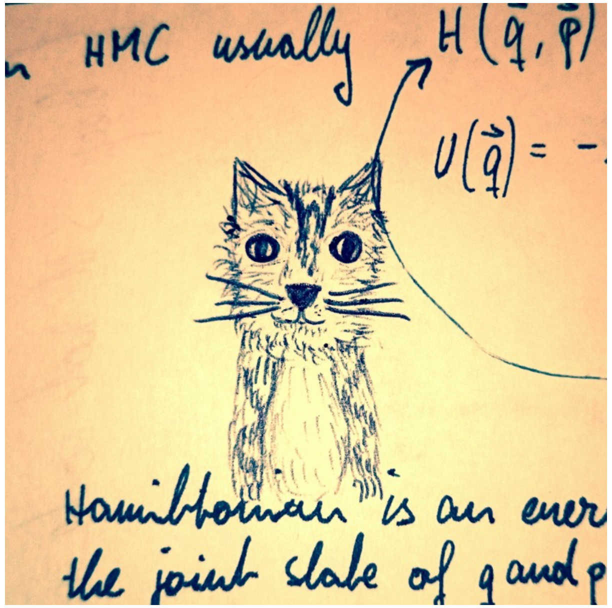

Multiple Importance Sampling: Math

And the weight can even depend on the sample.

$$E[F] = \sum_{i=0}^{n} \frac{1}{n_i} \sum_{j=0}^{n_i} E\left[w_i(X_{i,j}) \frac{f(X_{i,j})}{p(X_{i,j})}\right] = \int_\Omega f(x)\,\mathrm{d}x \text{ with } \sum w_i(X_{i,j}) = 1$$

Think about it that way:
We have our n strategies, but we draw only one sample each. By pure luck all samples $X_{i,0}$ are the same. In that case our weighting is clearly valid. But it's also valid when the samples are different. And this is the gist of MIS.

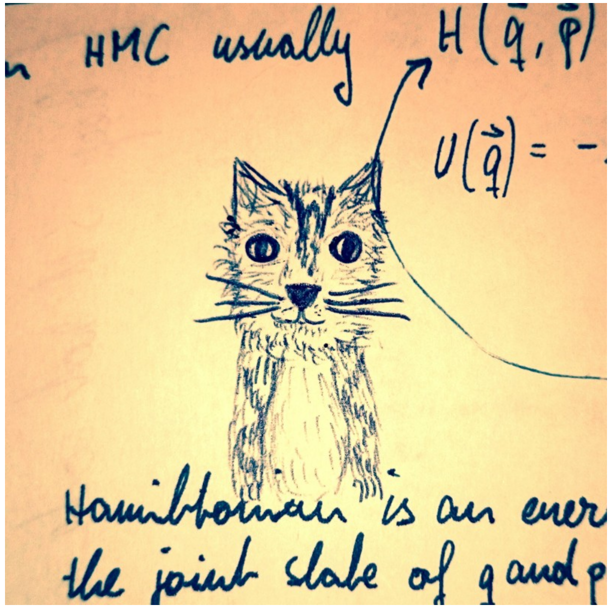

Multi-sample estimator is given by

$$F = \sum_{i=1}^{n} \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(X_{i,j}) \frac{f(X_{i,j})}{p_i(X_{i,j})}$$

It's unbiased when

$$\textbf{(W1)} \quad \sum_{i=1}^{n} w_i(x) = 1 \text{ whenever } f(x) \neq 0 \text{, and}$$

$$\textbf{(W2)} \quad w_i(x) = 0 \text{ whenever } p_i(x) = 0\,.$$

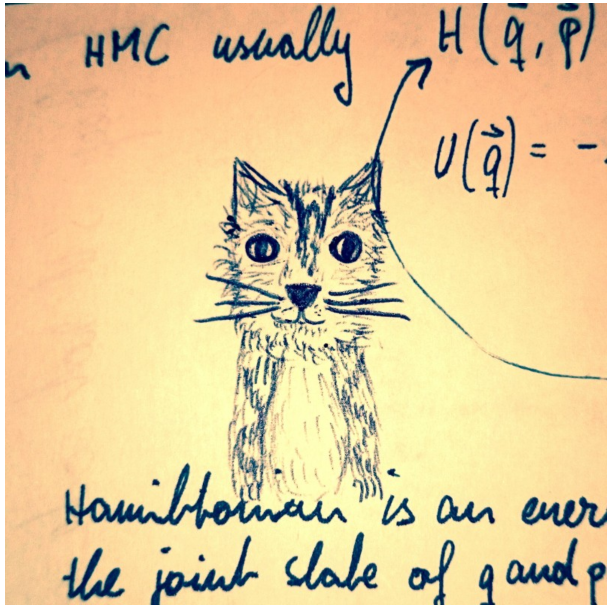

Robust Monte Carlo Methods for Light Transport Simulation (Veach 1997)

Multiple Importance Sampling: Math

Some examples of w_i

- Constant 1/n (from before, bad in practice because it doesn't kill variance effectively, see Veach 1997 PhD Thesis Chapter 9)
- 1 or 0 depending on $X_{i,j}$ (example 1d: use strategy A if x <0 otherwise B; You'll see examples of that in the path tracing lecture)
- Balance heuristic (You can't do much better than that, i.e. it's always within a bound of the best strategy Veach 1997, 9.2.2)

$$w_i(x) = \frac{p_i(x)}{\sum_{k=0}^{n} p_k(x)}$$

- Power heuristic (better if there is one strategy with very low variance)

$$w_i(x) = \frac{p_i(x)^\beta}{\sum_{k=0}^{n} p_k(x)^\beta}$$

Robust Monte Carlo Methods for Light Transport Simulation (Veach 1997)

Multiple Importance Sampling: Practical Example

Ok cat, my head is all mushy, can't you give me a practical example?

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Practical Example

Ok cat, my head is all mushy, can't you give me a practical example?

- Integrand f(x), estimator F
- Balance heuristic
- M sampling strategies (j=0..M)
- N samples (i=0..N)

Multiple Importance Sampling: Practical Example

- For each sample *i*
  - Pick a distribution using probabilities p(j)
  - Draw a sample x_i from it
  - Compute

$$F_i = \frac{f(x_i)}{\sum_{j=1}^{M} p(j) p_j(x_i)}$$

  - F += F_i (like you did before in MC)
- F /= N
- Done!

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Practical Example

- For each sample *i*
  - Pick a distribution using probabilities p(j)
  - Draw a sample x_i from it
  - Compute

$$F_i = \frac{f(x_i)}{\sum_{j=1}^{M} p(j) p_j(x_i)}$$

  - F += F_i (like you did before in MC)
- F /= N
- Done!

The p terms from page 24 are p(j)*pj(xi) here. Some terms cancel each other out, we had

$$F = \sum_{i=1}^{n} \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(X_{i,j}) \frac{f(X_{i,j})}{p_i(X_{i,j})}$$

and $w_i(x) = \frac{p_i(x)}{\sum_{k=0}^{n} p_k(x)}$.

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Practical Example

- For each sample *i*
  - Pick a distribution using probabilities p(j)
  - Draw a sample x_i from it
  - Compute

$$F_i = \frac{f(x_i)}{\sum_{j=1}^{M} p(j) p_j(x_i)}$$

  - F += F_i (like you did before in MC)
- F /= N
- Done!

The p terms from page 24 are p(j)*pj(xi) here. Some terms cancel each other out, we had

$$F = \sum_{i=1}^{n} \frac{1}{n_i} \sum_{j=1}^{n_i} w_i(X_{i,j}) \frac{f(X_{i,j})}{\cancel{p_i(X_{i,j})}}$$

and $w_i(x) = \frac{\cancel{p_i(x)}}{\sum_{k=0}^{n} p_k(x)}$.

On page 24 and before we had a fixed number of samples for each strategy, now we choose the strategy probabilistically and hence the additional p(j).

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Practical Example

- For each sample *i*
  - Pick a distribution using probabilities p(j)
  - Draw a sample x_i from it
  - Compute

$$F_i = \frac{f(x_i)}{\sum_{j=1}^{M} p(j) p_j(x_i)}$$

  - F += F_i (like you did before in MC)
- F /= N
- Done!

You can't do much better than equal chances, i.e., using probability p(j) = 1/M for all j (Veach 1995, Sec. 5.2)

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: What's Going On?

- The above process generates samples with the joint distribution

$$\bar{p}(x) = \sum_{j=1}^{M} p(j) p_j(x)$$

- Hence, we're just computing f/p with this new PDF.
  Note that the p(j)'s are a discrete distribution, their sum must be 1!
- *This is an unbiased estimate, just like regular MC.*

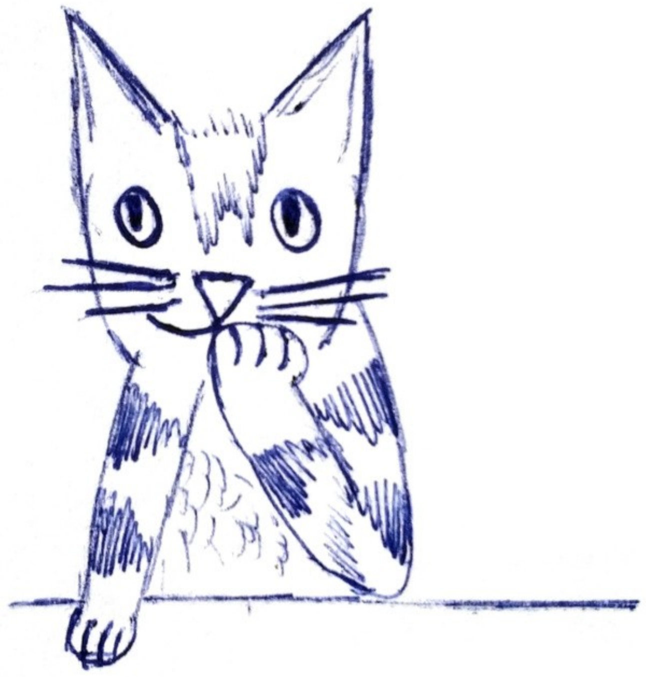

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Ha!

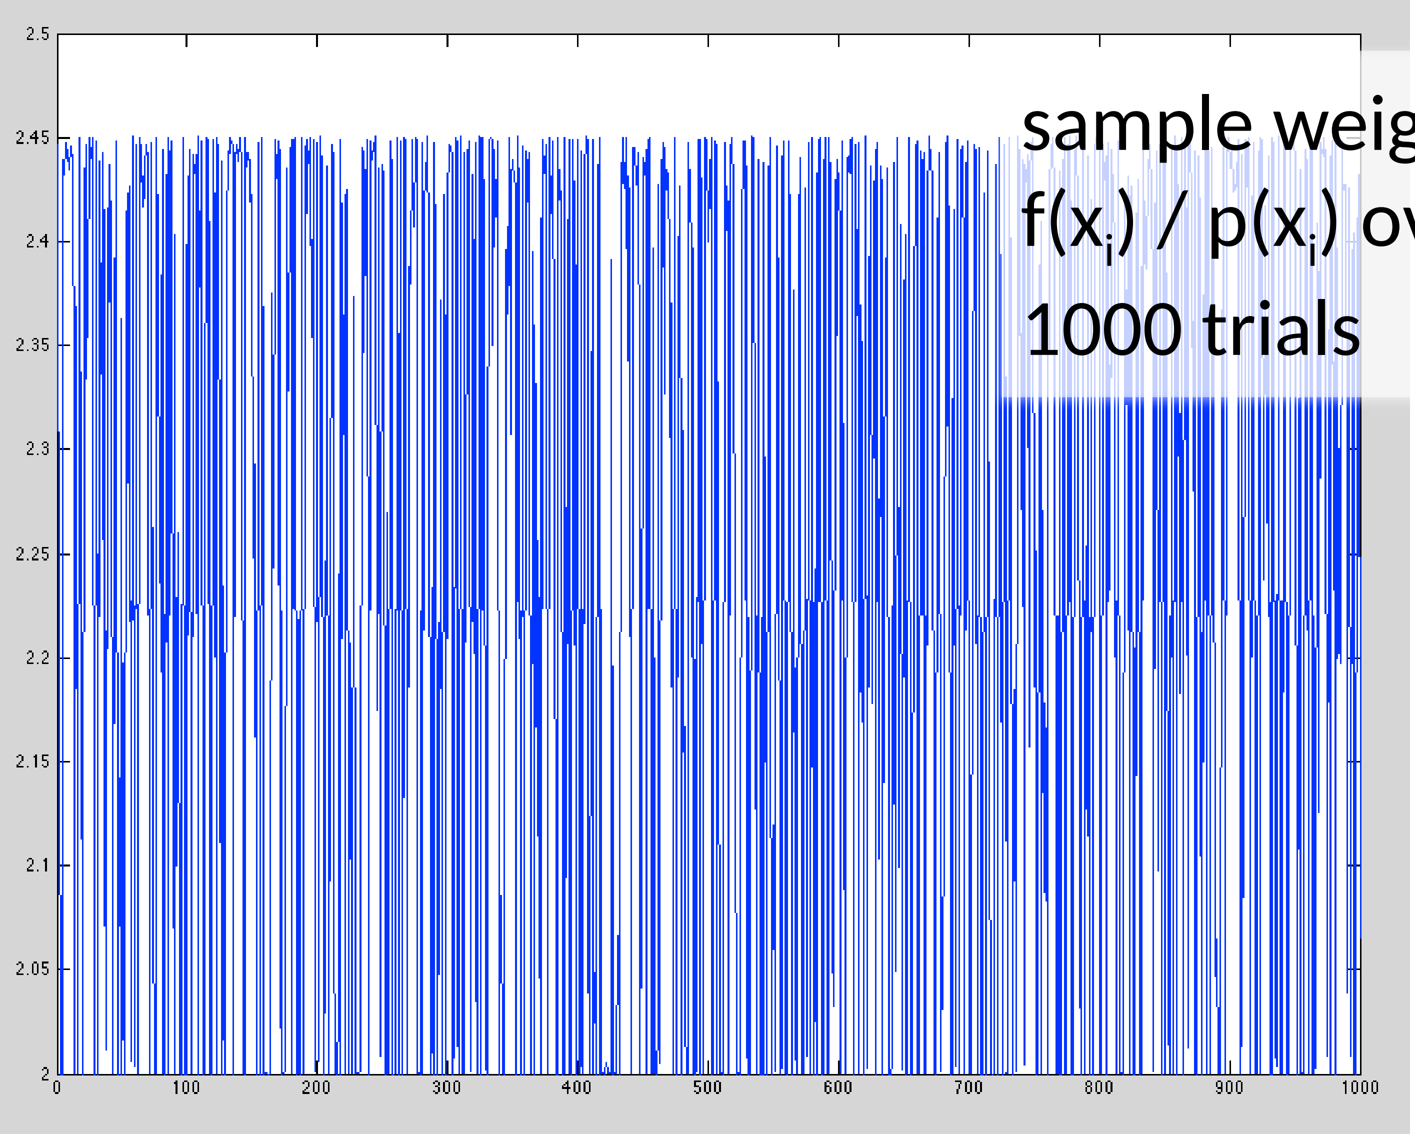

sample weight $f(x_i) / p(x_i)$ over 1000 trials

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Ha!

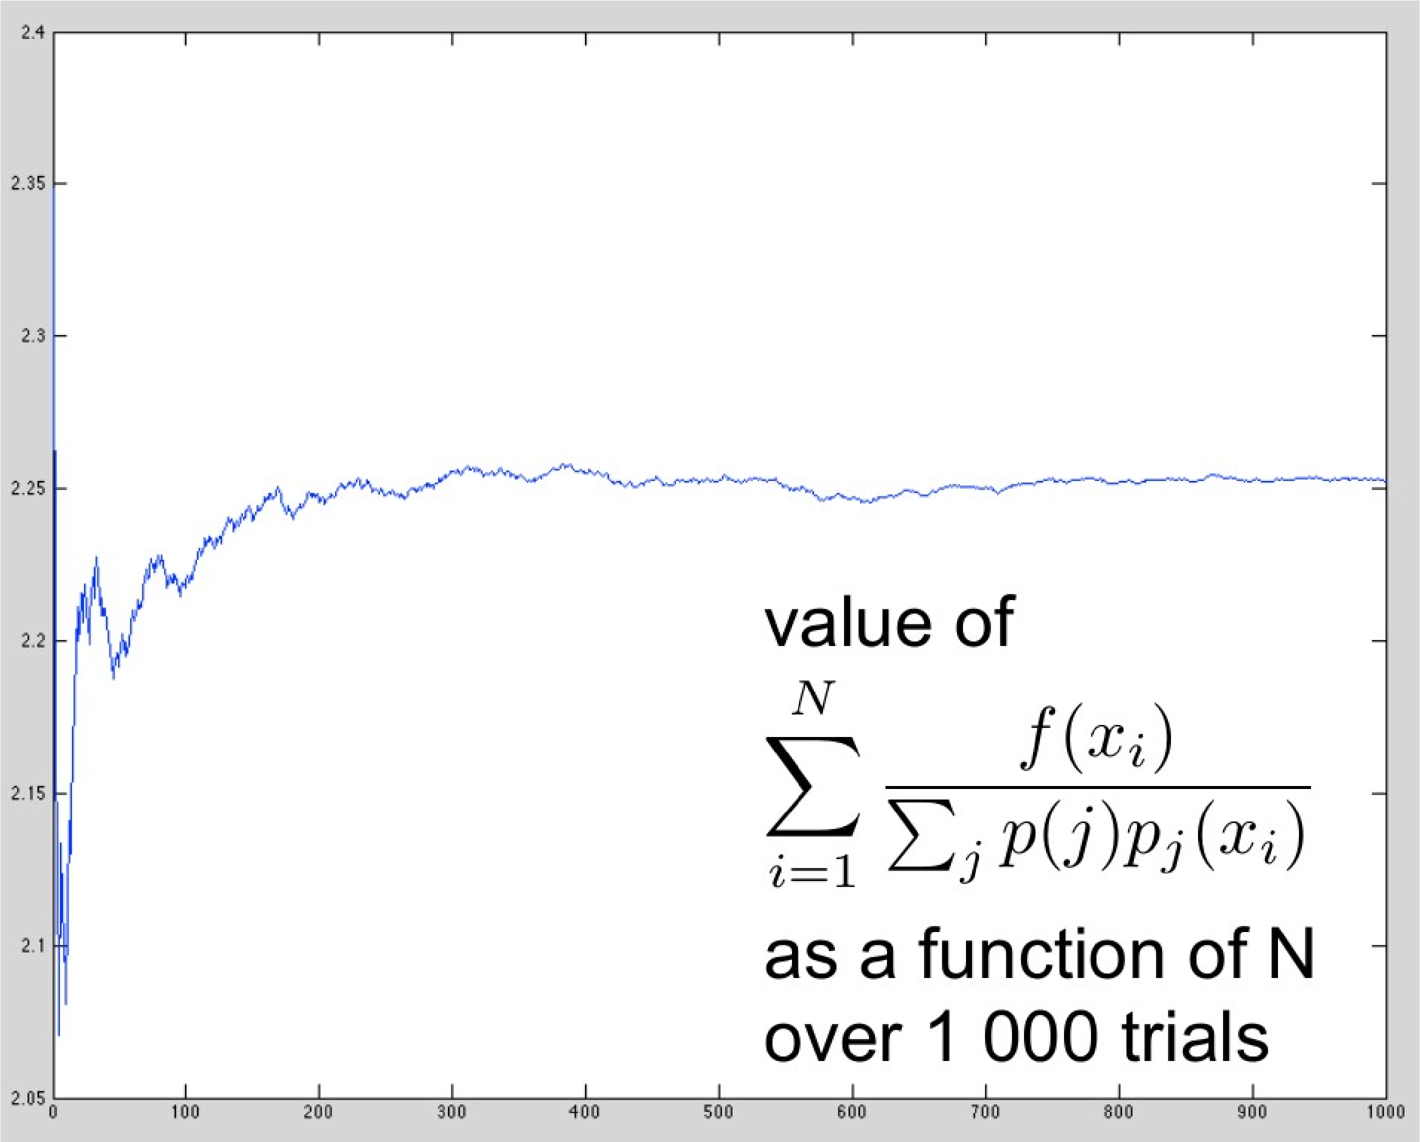

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Bah!

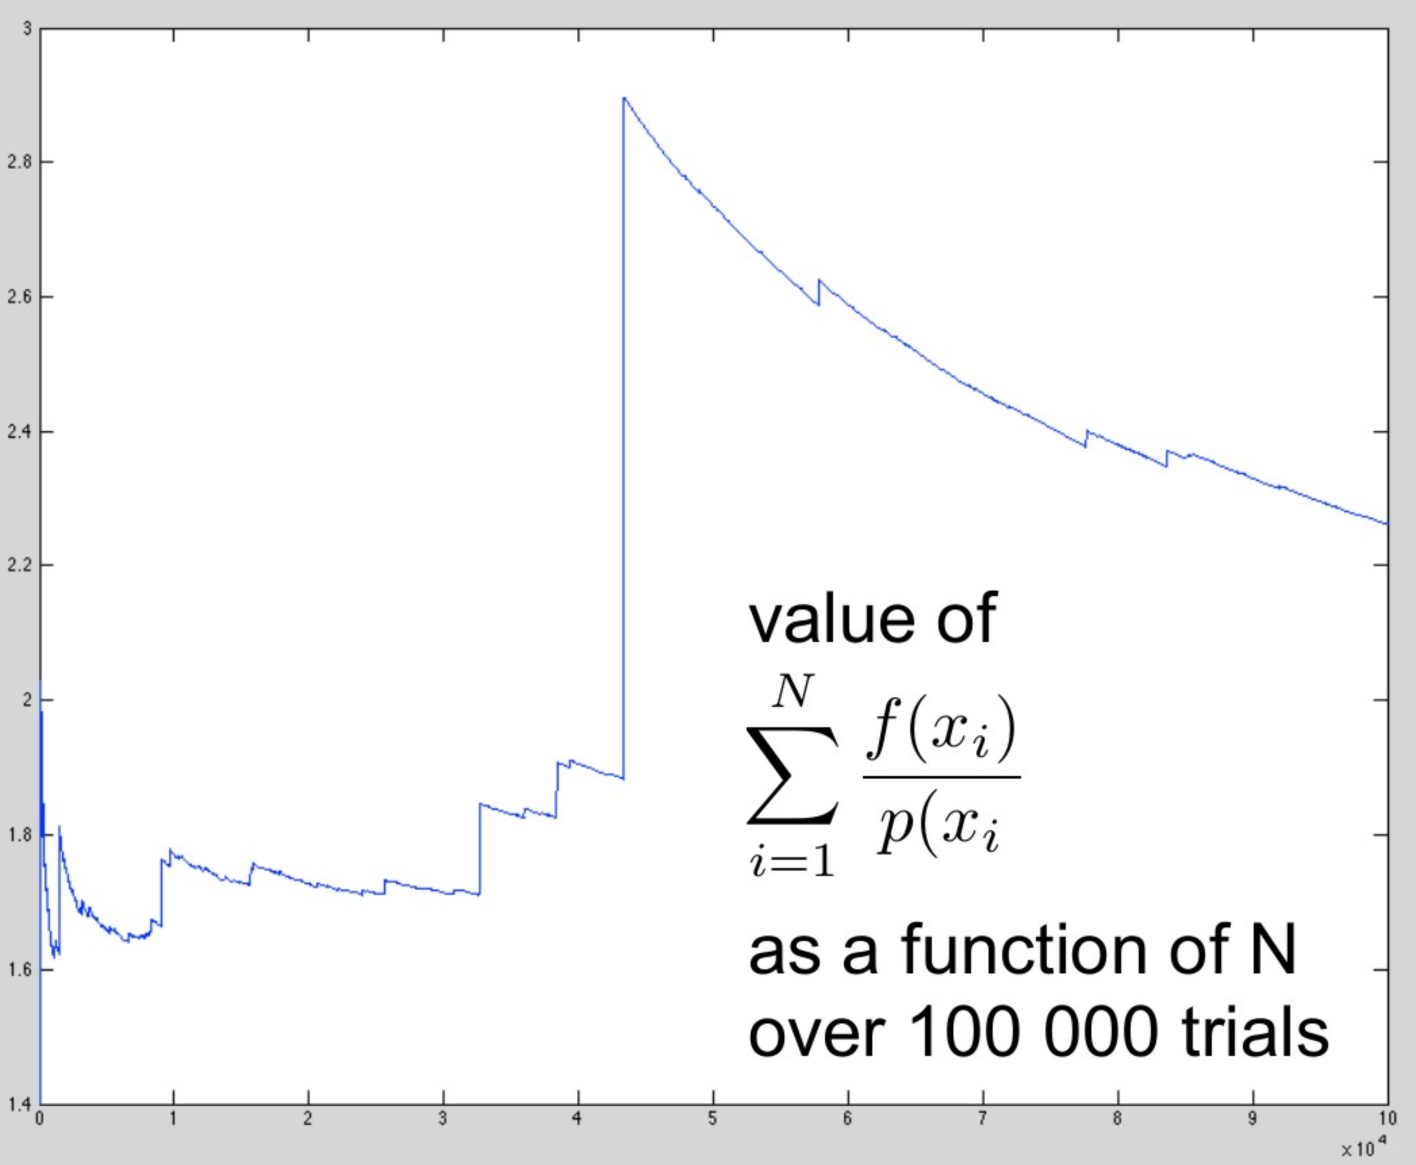

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Ha!

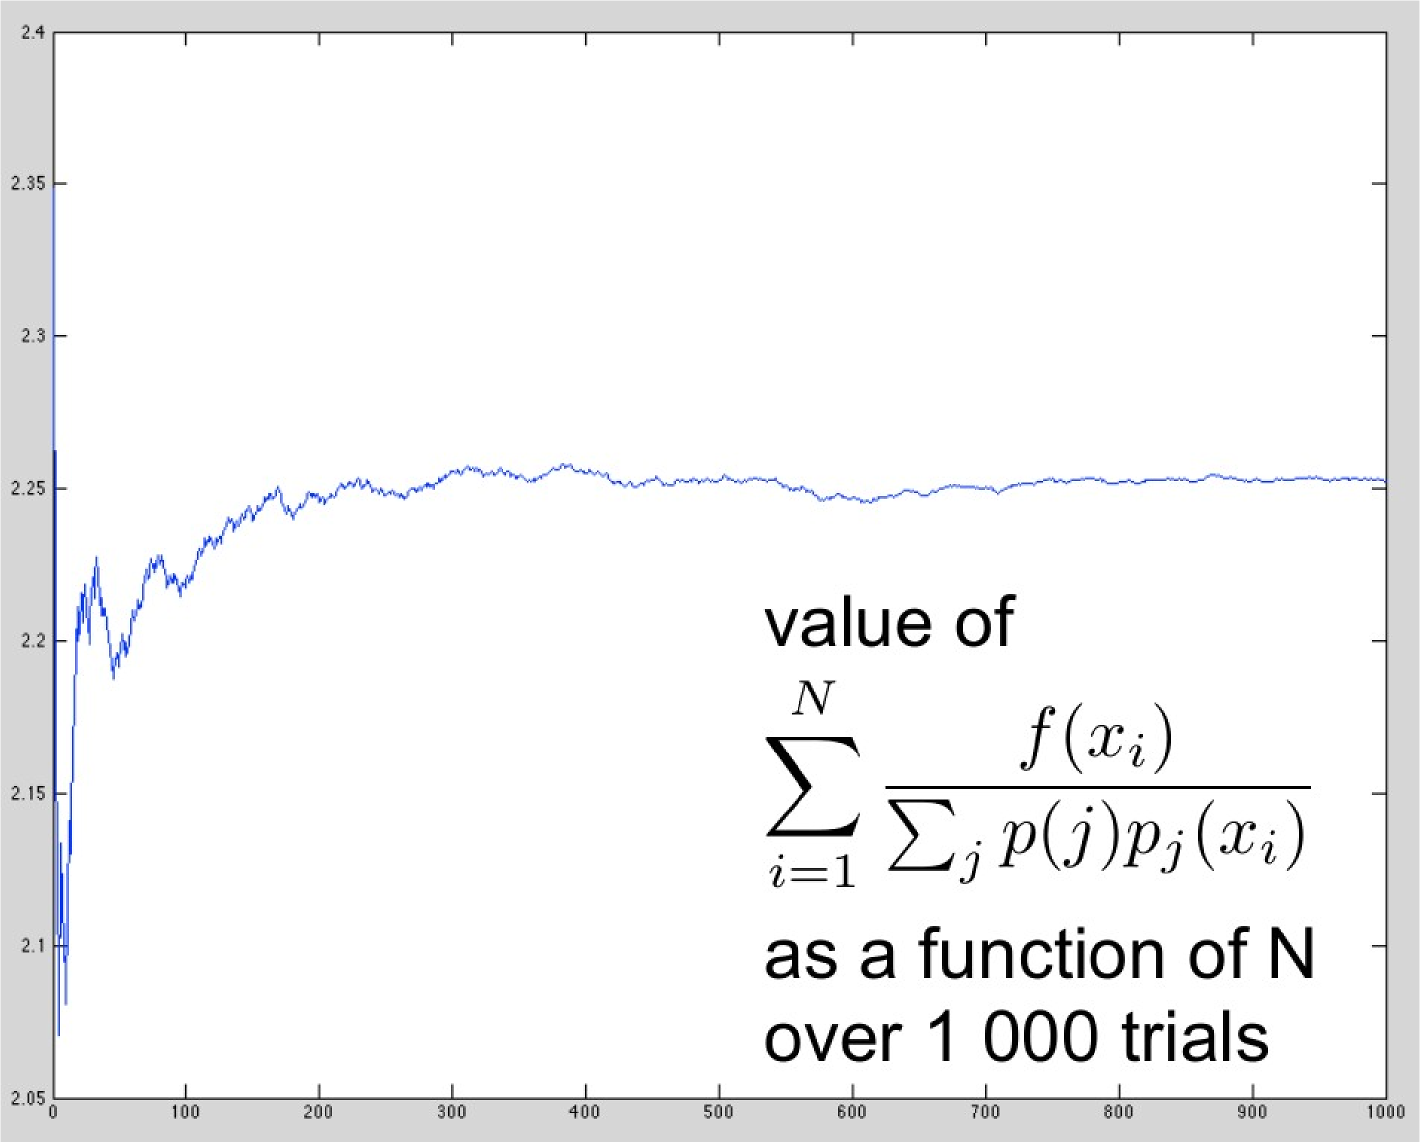

Slide modified from Jaakko Lehtinen, with permission

Multiple Importance Sampling: Bells and Whistles

- This is the basic intuition and approach
- Veach's 1995 paper and 1997 thesis contain a long treatment on how to choose the relative weighting between the PDFs and more general ways of constructing $\bar{p}(x)$ based on the individual distributions.
- Feel free to experiment with different strategies in your assignments :)

Slide modified from Jaakko Lehtinen, with permission

That's it..

There are some reading links on the next page, in case you feel bored :)

Useful reading (links)

- Jaakko Lehtinen’s slides (I borrowed a lot from lecture 4)
- My DA thesis, Section 2.3 (very brief write up of Monte Carlo Integration + MIS, but maybe you’ll like it)
- Last years lecture (recordings)
- Veach’s PhD Thesis (contains a lot of information, I liked it better than the papers)
- Veach’s 1995 paper