

Bayesian Learning for Guided Direct Illumination Sampling



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



Computer
Graphics
Charles
University

- In this part of the course, I will focus on guiding in direct illumination sampling.

- Guiding needs **radiance approximations**
- How to learn them **reliably**?
- Our proposition:
 (Online, Bayesian) **Machine learning**
 [*Vorba et al. 2014, Vévoda et al. 2018*]



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- However, I would like to go beyond the specific problem of direct illumination and really focus on the fundamental methodology we used to address this problem, since it is applicable in a wider context in guiding and adaptive sampling.
- As we have seen in the course introduction, any guided sampling requires some information about the radiance distribution in the scene.
- The radiance is not available upfront, so we must resort to approximations.
- The approximations can be obtained by learning from the Monte Carlo samples themselves – but these provide noisy and unreliable estimates, especially early in the computation.
- So the general question we are striving to address is how to obtain reliable information from unreliable, noisy samples.
- Our proposition is that machine learning, and specifically Bayesian statistics, provide some excellent tools to tackle exactly this problem.
- This was the underlying idea both in our 2014 direction guiding work [Vorba et al. 2014] mentioned earlier, and also in the direct illumination guiding work by Petr Vévoda and Ivo Kondapaneni that I will present now [Vévoda et al. 2018].

Take home message

Machine Learning | Bayesian modeling

=

**Excellent framework for
guided/adaptive Monte Carlo**



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



Bayesian online regression for adaptive direct illumination sampling


Petr Vévoda, Ivo Kondapaneni, and Jaroslav Křivánek

Chaos Czech a.s.
Charles University, Prague




Computer
Graphics
Charles
University


- I will now describe our guided direct illumination solver build around these ideas.
- It is the default direct illumination solver used in Corona Renderer since version 3.




- Let us first motivate our work.
- MC rendering algorithms are currently getting more and more popular, but they suffer from noise.




- Traditionally, the indirect illumination component has been considered as the main source of noise, and it's been subject to lot of research.
- But in this, as well as in many other production scenes, it is the *direct component* that causes the trouble.



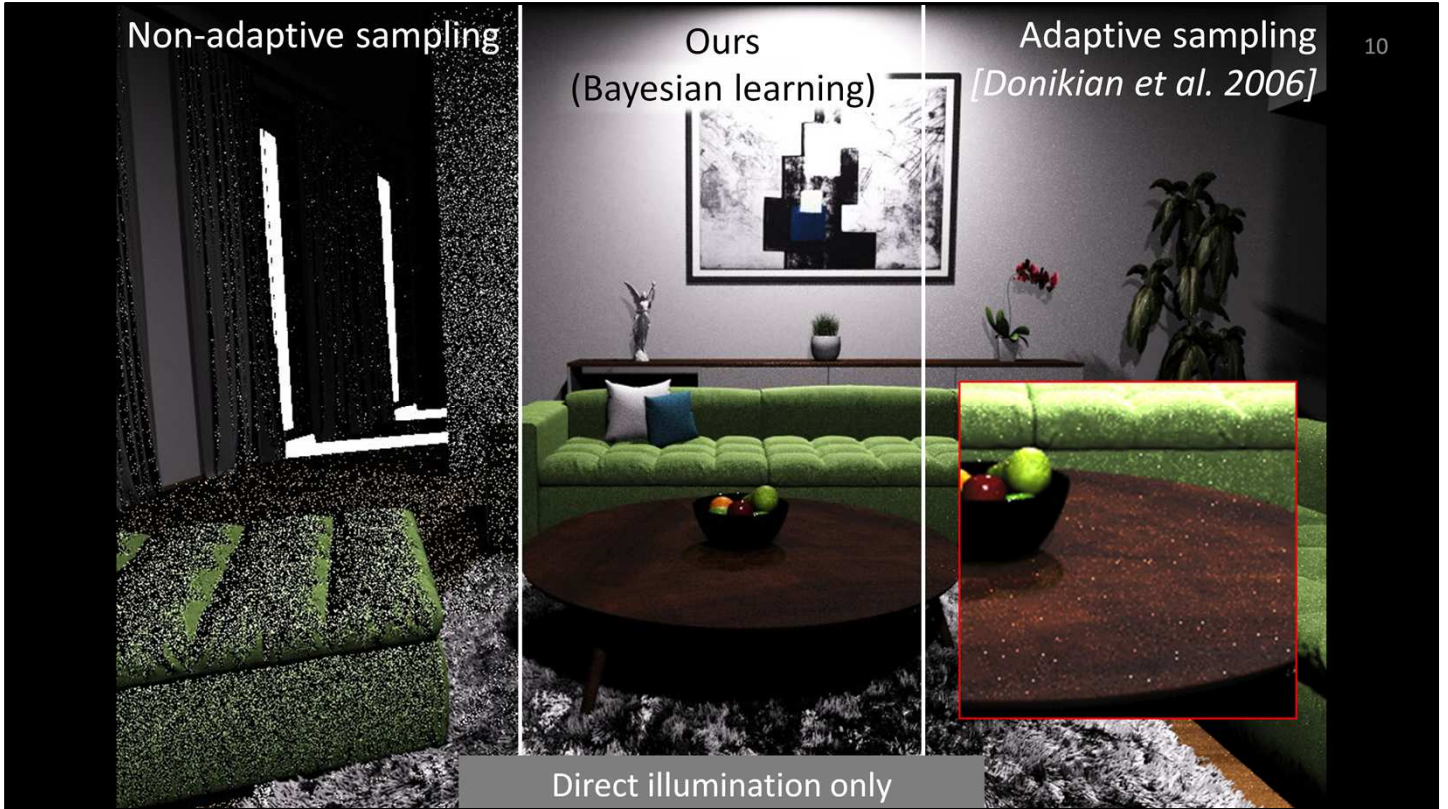
- In this image with only direct illumination we can clearly see the problem.
- A non-adaptive direct illumination sampling method that samples lights proportionately to the a conservative estimate of their unoccluded contribution is shown on the left.
- It struggles to work efficiently, because it wastes lot of samples on the strong but completely occluded sun.




- One possible solution is to use previous *adaptive methods* which try to improve sampling based on past samples.
- But while they can decrease the amount of noise significantly, ...




- ... they can also introduce various artifacts and spiky noise because they are based on adhoc solutions and they tend to overfit to the input noisy data.
- This lack of *robustness* is a consequence of adhoc solutions to crucial questions in adaptive sampling:
 - When is it safe to rely on the noisy samples?
 - How should the noisy samples be combined with any previous, a priori knowledge?



- We propose a solid theoretical framework based on Bayesian learning that addresses the above questions, and thanks to that enables robust adaptive sampling in rendering.



- In this scene, our solution is more than 500 times faster than the non-adaptive solution...



- ... and we can achieve much better robustness than the previous adaptive sampling method by Donikian et al. [2006].
- The Bayesian framework we introduce here is not limited to guided direct illumination. Other guiding / adaptive sampling methods can benefit from it as well.

Previous work



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Much existing work is related to ours.

Adaptive sampling

- General Monte Carlo
 - Vegas algorithm
 - [Lepage 1980]
 - Population MC
 - [Cappé et al. 2004, ...]
- Rendering
 - Image sampling
 - [Mitchell 1987, ...]
 - Indirect illumination (path guiding)
 - [Dutre and Willems 1995, Jensen 1995, Lafortune et al. 1995, ...]
 - [Vorba et al. 2014, Muller et al. 2017]
 - Direct illumination
 - [Shirley et al. 1996, Donikian et al. 2006, Wang et al. 2009]



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Adaptivity in Monte Carlo simulation is not a new concept.
- There is much work in general Monte Carlo as well as in rendering: for example works dealing with image sampling, indirect illumination sampling, and direct illumination sampling as well.
- A classic adaptive algorithm for general Monte Carlo estimation is the Vegas algorithm by Lepage [1980], which works by histogramming (stratifying) the integrand and using these histograms for better sampling in subsequent steps.
- Another example is Population Monte Carlo [Cappé et al. 2004], which uses a population of simulation particles and tracks how well they sample the integrand. Based on that, they keep the best individuals for subsequent sampling rounds.
- A lot of adaptive sampling work exists in rendering.
- Of the many works, we mention Mitchell [1987] which deals with allocation of samples into image parts with high-frequency content.
- In adaptive indirect illumination computation, some of the early works were done by Dutre and Willems [1995], where the authors adaptively trace particles from lights; Jensen [1995] who used photon maps to construct path sampling distributions (this is probably the first work in graphics on direction guiding); and Lafortune et al. [1995] who applied a Vegas-like approach in Monte Carlo light transport simulation.
- The ongoing renewed interest in path guiding / adaptive path space sampling could be attributed to our work [Vorba et al. 2014], which proposes online learning of Gaussian mixtures for direction guiding. The direction sampling algorithm by Muller

et al. [2017] is a version of the Vegas-style algorithm by Lafortune from 1995.

- As for direct illumination, we mention the pioneering work by Shirley et al. [1996] who adaptively classified lights into important and unimportant ones.
- The most closely related is the work by Donikian et al. [2006] that we describe in more detail later.
- Wang et al. [2009] sample lights adaptively based on surface reflectance and estimates of lights' contributions.
- None of these works deal with a problem of determining when and how to incorporate all the information into a robust, reliable 'trained' sampling distributions.

Bayesian methods in rendering

- Filtering
 - NonLocal Bayes [*Boughida and Boubekeur 2017*]
- Global illumination
 - Bayesian Monte Carlo [*Brouilat et al. 2009, Marques et al. 2013*]
 - Path guiding [*Vorba et al. 2014*]



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Bayesian methodology has so far not been widely applied in rendering.
- The few Bayesian methods in image filtering and global illumination are listed on the slide.

Background




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling




- Let me now provide some background related to the direct illumination problem.

Direct illumination problem



- Consider a scene with several light sources (yellow stars) and geometry (blue blocks).
- The goal is to calculate the *direct* contribution of all lights toward any given point in the scene.
- That is a complex task due the potentially high number of light sources, their highly uneven contributions, and due to occlusion
- We assume that the sampling problem is broken down to two stages:
 - A) pick a light at random
 - B) randomly sample a position on the chosen light
- While much previous work has addressed optimal choices in the step B, we focus solely on step A: the random choice of a light source.


Non-adaptive, un-occluded light sampling



- One common strategy for light sampling considers conservative bounds on the un-occluded light contribution to the surface point.
- We can use the bounds to build a sampling distribution over the lights and then use it for choosing the lights at random (step A on the previous slide).
- But this non-adaptive approach cannot approximate the true light contributions well especially when the individual lights are large and have complex emission distributions, and when of them many are (partially) occluded.

Adaptive light sampling

[Donikian et al. 2006]




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Donikian et al. [2006] propose an adaptive direct illumination sampling algorithm, which however suffers from limited robustness.
- They gather statistics about true light contributions in screen space for each pixel.
- But the per-pixel statistics are too noisy, so they additionally gather average stats over entire blocks of pixels.
- Then they mix both the per-pixel and per-block statistics to obtain the final sampling distributions: early on, the blocks are given more weight and as more samples arrive, the per-pixel statistics are given more weight.
- But the mixing is done in an ad-hoc way, which often results in overfitting of the sampling distribution to the samples, and consequently produces image artifacts, especially in the complex illuminations cases that we wish to be handled by the adaptive sampling in the first place.

Problem summary



- Let us now summarize the problem at hand.
- Our goal is to compute direct illumination by means of Monte Carlo sampling, and for that we need to find **optimal** discrete sampling distribution over individual lights.
- We choose to use **adaptive sampling (i.e. guiding)** so that we can deal with (partial) light occlusion, but we strive for a **robust, 100% reliable solution**, even in complex edge cases.
- We have two kinds of information at our disposal:
 - First, the bounds of light contribution towards a point, which are conservative and noise-free. These can be computed on the fly, so they are available for the start. As such, they can serve as our prior belief about the lights' real contribution to any given scene point.
 - Second, the Monte Carlo estimates of direct illumination (i.e. lights' real contribution). These are noisy at the start, but as more samples are taken, their average eventually converges to the correct answer - over time, they become more accurate than the conservative bounds.
- We propose a well-founded approach to combining these two sources of information in a robust way based on Bayesian inference, as we explain shortly.

Our approach



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Let us now introduce our approach ...

Contributions

- What distribution should we learn?
- Learning the distribution through Bayesian inference




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- We make two main contributions:
 - First, we show what should be the optimal discrete probability distribution for choosing the lights, given the statistics of light contributions.
 - Our second and main contribution is the use of Bayesian inference to learn these sampling distributions. This gives us a robust solution and allows us to combine Monte Carlo samples with light contribution bounds in a principled manner.
- Let us start with the first point.

Optimal light sampling distribution

$$P(L) \propto \sqrt{\text{mean}^2 + \text{variance}}$$




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling




- Given a point in the scene, we want to sum up contributions from all lights.
- When the lights are too many, summing all the contributions would take too much time, so we randomly pick one or just a few. So far nothing new, this is a standard way to go about computing direct illumination in a Monte Carlo renderer.
- Choosing a few lights at random - instead of summing the contribution of all of them – can vastly increase efficiency, but it also introduces variance.
- The question is: What is the discrete probability mass function (PMF) for picking the lights at random that provides the lowest possible variance?
- The usual approach in Monte Carlo sampling would be to make the lights' sampling probabilities proportional to their expected (or **mean**) contribution to a given point.
- We show in the paper that this traditional choice may in fact be far from optimal, and we derive the optimal solution.
- To arrive at the optimal solution, it is important to realize that once we pick a light, another random decision follows, which selects a particular sample location on the light. The contributions from these sample locations will generally vary: usually the larger the light and/or the sharper the surface BRDF, the more variance in the individual sample contributions. Since this variance eventually creeps into the overall direct illumination estimator over all lights, we need to reduce it - by allocating more samples to those lights that yields samples with highly varying contributions.
- The final optimal sampling should therefore take into account **both the mean** contribution but **also the variance** of the contributions for each light.
- The specific formula is given on the slide. It is equivalent to making the light selection probability proportional to the **square root of the second moment** of the samples

from that given light.



- Let me show a practical example



- On the left we see an inset showing how sampling according to a mean performs. It under-samples some tricky light which leads to spiky noise.
- And on the right we see that sampling according to both the mean and variance eliminates this issue.

Contributions

- Optimal sampling distribution
- Learning the distribution through Bayesian inference

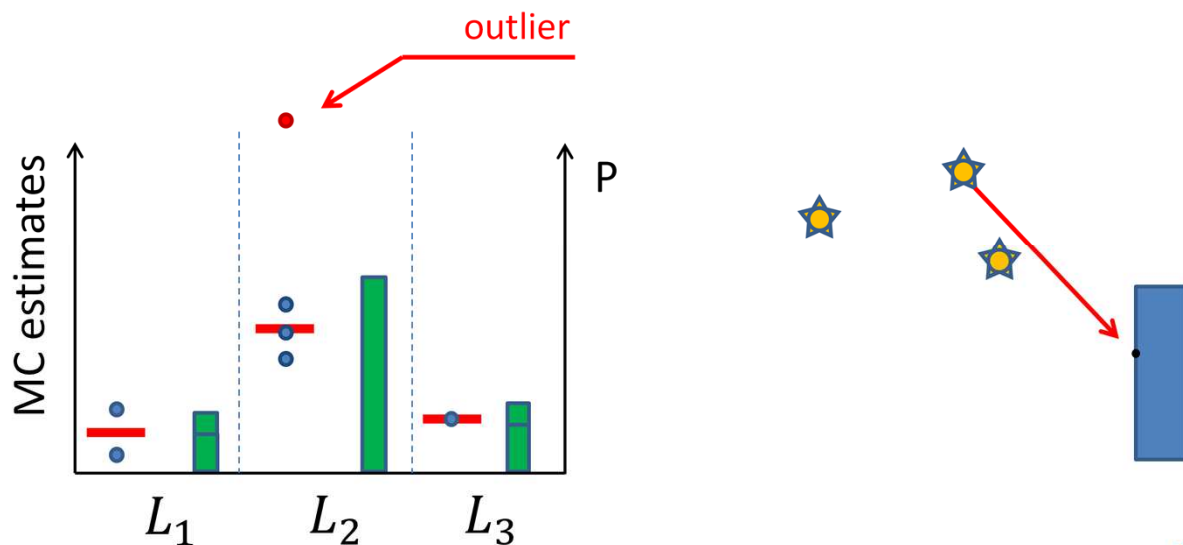


Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Having explained the target optimal sampling distribution that we strive to achieve, we now show how to learn this desired distribution using Bayesian regression.
- The issue is that the mean and the variance needed for the optimal sampling are not known upfront and need to be learned during rendering.

Naive adaptive light sampling




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



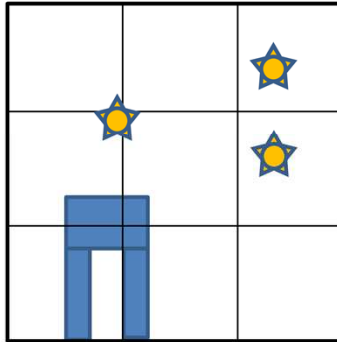
- To motivate our approach, let us first see how a naive adaptive approach would behave.
- Suppose we have already taken some samples from lights (blue dots) and obtained some light sampling probabilities (green bars).
- Now suppose we have taken a direct illumination sample which happens to be an outlier. For instance, the sample may lie on the light extremely close to the illuminated point.
- If we estimated the sample means directly, our estimates would change abruptly due to the outlier sample.
- That would have a disproportionately strong effect on further light sampling: one light would get sampled very often at the expense of other lights, leading to increased image noise or even strong fireflies.

Bayesian adaptive light sampling



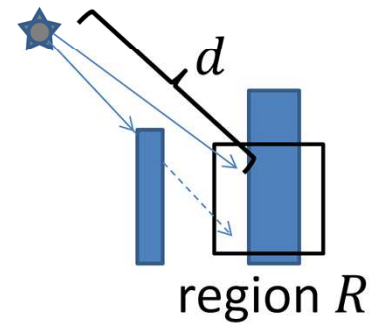

- We propose to estimate the means (and variances) in a Bayesian manner: **We model the distributions of MC estimates** seen so far, while we also have some **prior information about parameters of that distribution**.
- As a result, when we get a new – possibly outlier – sample, our distribution is affected less abruptly and so do the cluster sampling probabilities derived from it.
- **Resiliency to outlier samples without compromising the ability to learn from the new samples** is the basis of our robust solution.

Scene subdivided in regions




- Before explaining how we model the data (i.e. lights' direct illumination contributions), let us explain the basic set up.
- We split the scene into fixed spatial **Regions** (we use a regular grid, but that is not so important – any space subdivision could be used).

Light-region statistics



- Now we focus on samples collected from **one light** in a **single scene region**, and we keep track of the distance d in the geometry factor in each estimate.
- We can now then plot the data (i.e. MC direct illumination estimates) for a light in a region. This reveals the nature of the relation of illumination estimates to the distance.
- One may observe the inverse squared falloff with the distance, and a number of zero-valued (occluded) samples.

Regression data model



Parameters:

k, h - normal distr. parameters

p_0 - probability of occlusion


$$(1 - p_0) \times N(\text{est.} \mid \frac{k}{d^2}, \frac{h}{d^4})$$

$$p_0 \times \delta(\text{est.})$$

- Now that we understand the general characteristics of the gathered direct illumination samples (i.e. data), we proceed to establish a **model for the data**.
- Our model is a **parametric regression model**, which - for a given distance d - yields a distribution of direct illumination sample values. We design our data model as follows:
 - Non-zero samples are modeled by a normal distribution (i.e. a Gaussian) with mean and variance being a function of the distance associated with direct illumination samples. This part of the model has two parameters, k and h .
 - The zero valued samples are incorporated by mixing the inverse-square distance falloff model with a delta function, and it is controlled by the parameter p_0 , which has the meaning of occlusion probability.
- Having designed our data model, we proceed to define the prior distribution for all three model parameters, k , h , and p_0 .

Conjugate prior

$$\text{posterior} \propto \text{likelihood} \times \text{prior}$$



- The model we have just defined has parameters k , h and p_0 . In the Bayesian treatment, these are viewed as random variables governed by a certain probability distribution.
- The “prior” is a specific probability distribution that we believe these parameters follow, without observing any data at all.
- A “conjugate prior” is a convenience construction: we choose the prior to have a functional form that will be preserved when multiplied by the likelihood of observed data.

Our (conjugate) priors

$$p_0 \sim \text{Beta}(p_0 | \dots)$$

$$k, h \sim \text{Normal inverse gamma}(k, h | \mu_0, \dots)$$

Hyperparameters

Light contrib. estimate



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- The mathematical form of the likelihood function is that of our data model (recall that the data likelihood is calculated simply by plugging our data samples into the data model).
- We show in the paper that the conjugate prior in our case takes form of the Beta distribution for p_0 and the normal-inverse-gamma distribution for the parameters k and h .
- There are various **hyperparameters** in the equations, but one parameter which stands out is μ_0 , **for which we use the conservative light contribution bound**.
- This hyperparameter expresses our a priori belief about the mean of our data: That is, the light contribution bounds provide a prior information about the expected contribution each light will make to each scene region. This belief is then continually refined as we observe the actual direct illumination contributions made by sampling the lights.

Algorithmic summary

- During each Next event estimation (in a region)
 - Compute data distributions for each light (mean, variance).
 - Build sampling PMF over lights
 - Choose lights from the PMF & samples on lights at random
 - Update light-region stats




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- To wrap it up, our algorithm is as described on the slide.

Scalability – Light clustering

Technical detail – not essential for our take-home message



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- We see from the algorithm described on the previous slide that the overall complexity of calculating direct illumination at a given scene point is still linear with respect to the total number of lights.
- This is because our probabilistic model has to be evaluated for each light, the sampling distribution has to be constructed, turned to a cumulative distribution to facilitate the sampling, and properly normalized.
- This linear complexity can be a bottleneck in scenes with many lights.
- For this reason, we apply a classic solution: the lights are clustered such that all lights in a cluster have similar contribution to a given scene region. The light clustering is generally different for each scene region. These clusterings are constructed up front, before the rendering starts.
- To apply the Bayesian approach described so far to this more general case, simply replace “a light” by a “cluster of lights” on all previous slides: We are learning discrete distributions to randomly select a light cluster. A light from the cluster is then selected in a non-adaptive way (e.g. proportionately to the total emitted light power), and finally a sample is selected in the light.
- With respect to our take-home idea (i.e. the use of Bayesian learning in adaptive sampling), the light clustering is really just a technical detail. (Though most of the code actually deals with handling the clusterings.)

Results




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Let us now demonstrate our solution in practice.

Tests




Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- This is a list of all comparisons you will see.
- Let me start with performance testing in a scene with simple occlusion in direct illumination only setting.



- This is the living room scene from the beginning of our presentation.
- It is lit mostly by a few small area lights on the ceiling, only in the left part sunlight can enter the room through the windows.




- As you could already see, the non-adaptive sampling of Wang et al. does not perform well in this scene. Why is that? The sun is much stronger than the ceiling lights and is therefore sampled much more often even though it is actually occluded - and so most of the samples are wasted.
- Donikian et al.'s algorithm improves the result significantly, as it quickly learns the sun occlusion. On the other hand, it struggles with the ceiling lights. They are covered by shades which block some of the samples. The method believes these lights are actually completely occluded, and consequently undersamples these lights and introduces spiky noise.
- Our method can also quickly learn the sun occlusion and converges more than 500x faster than the non-adaptive method of Wang et al.
- Interestingly, we can even observe a higher empirical convergence rate.
- At the same time, thanks to the Bayesian treatment, our method is robust, does not get confused by the occluded samples and avoids the spiky noise.

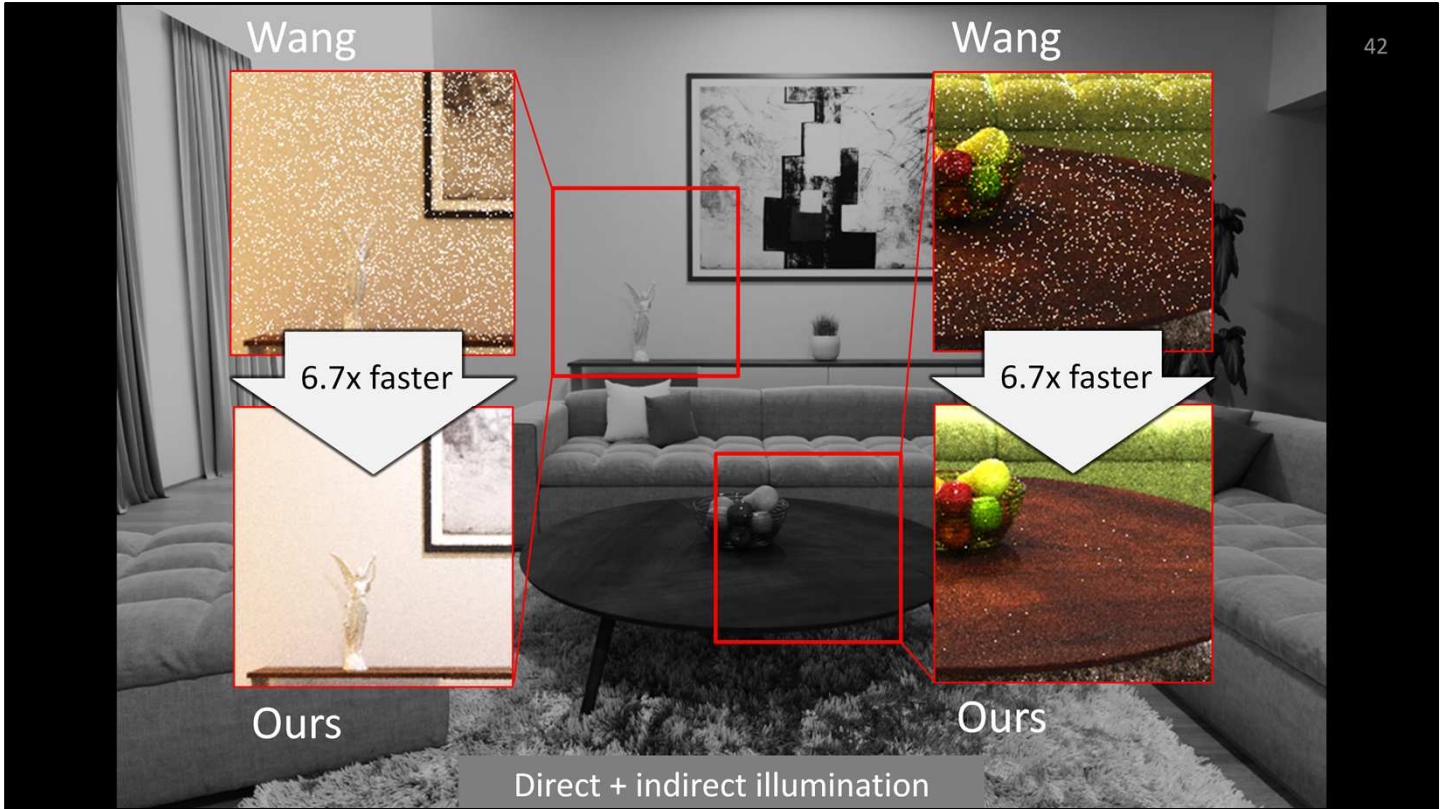
Tests

	Direct only	Direct + indirect
Simple occlusion	✓	
Complex occlusion		

- So, that was the direct illumination.
- But in practice one is usually more interested in images containing both direct and indirect components....



- ... like this one.




- We can see that the strong direct illumination noise of Wang et al. dominates also in the complete image. The direct component is definitely the main source of noise in this scene.
- By using our method in the next event estimation in path tracing, we are able to improve the light sampling on every path vertex and get more than 6x overall speedup.
- Note that the remaining noise at the bottom right of the image is caused solely by the indirect component and cannot be remedied by our method.

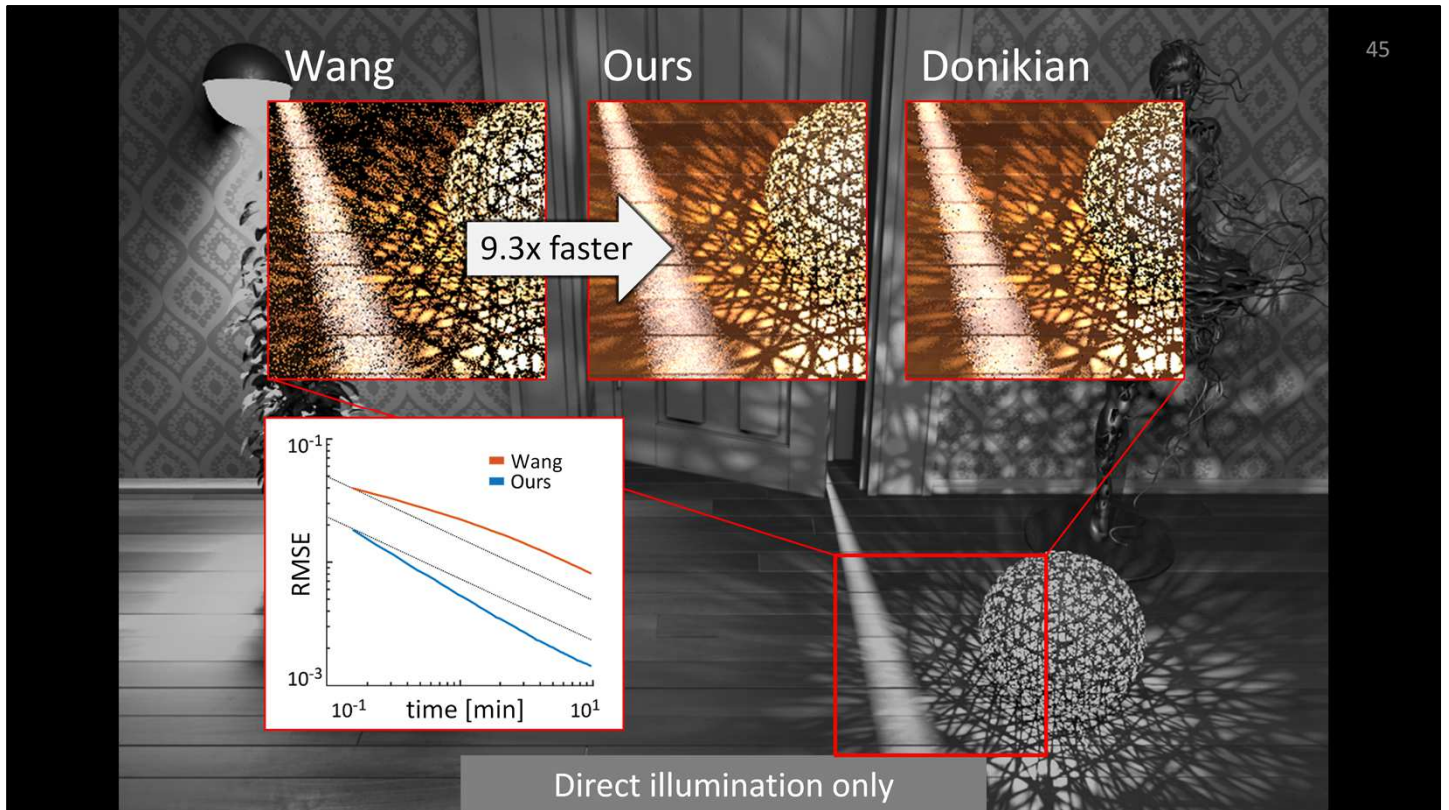
Tests

	Direct only	Direct + indirect
Simple occlusion	✓	✓
Complex occlusion		


- Now it is time to stress-test robustness of our method in a scene with complex occlusion.



- This scene presents a real challenge due to its highly structured illumination plus there are high-intensity lights in the other room behind the door.



- In this part Wang et al.'s method produces a lot of noise again, as it wastes samples on the lights behind the door.
- On the other hand, our method performs well.
- It is more than 9 times faster, and again we can observe higher empirical convergence rate. And all that without introducing any artifacts in such a complicated illumination setting.
- Donikian et al.'s method at first also seems to perform well but further inspection would discover small blocky artifacts in the shadows.




- In this inset, the non-adaptive sampling of Wang et al. again does not perform well.
- But this time also the Donikian et al.'s method fails badly: The illumination coming through the plant leaves is too complex for the ad-hoc learning to handle it well, the method overfits and produces square-shaped artifacts.
- This is exactly the problem of essentially all previous adaptive methods: While they can sometime provide substantial speedup, **they do not fail gracefully**, and one cannot rely on them.
- Our Bayesian learning makes our method much more robust and artifact-free.

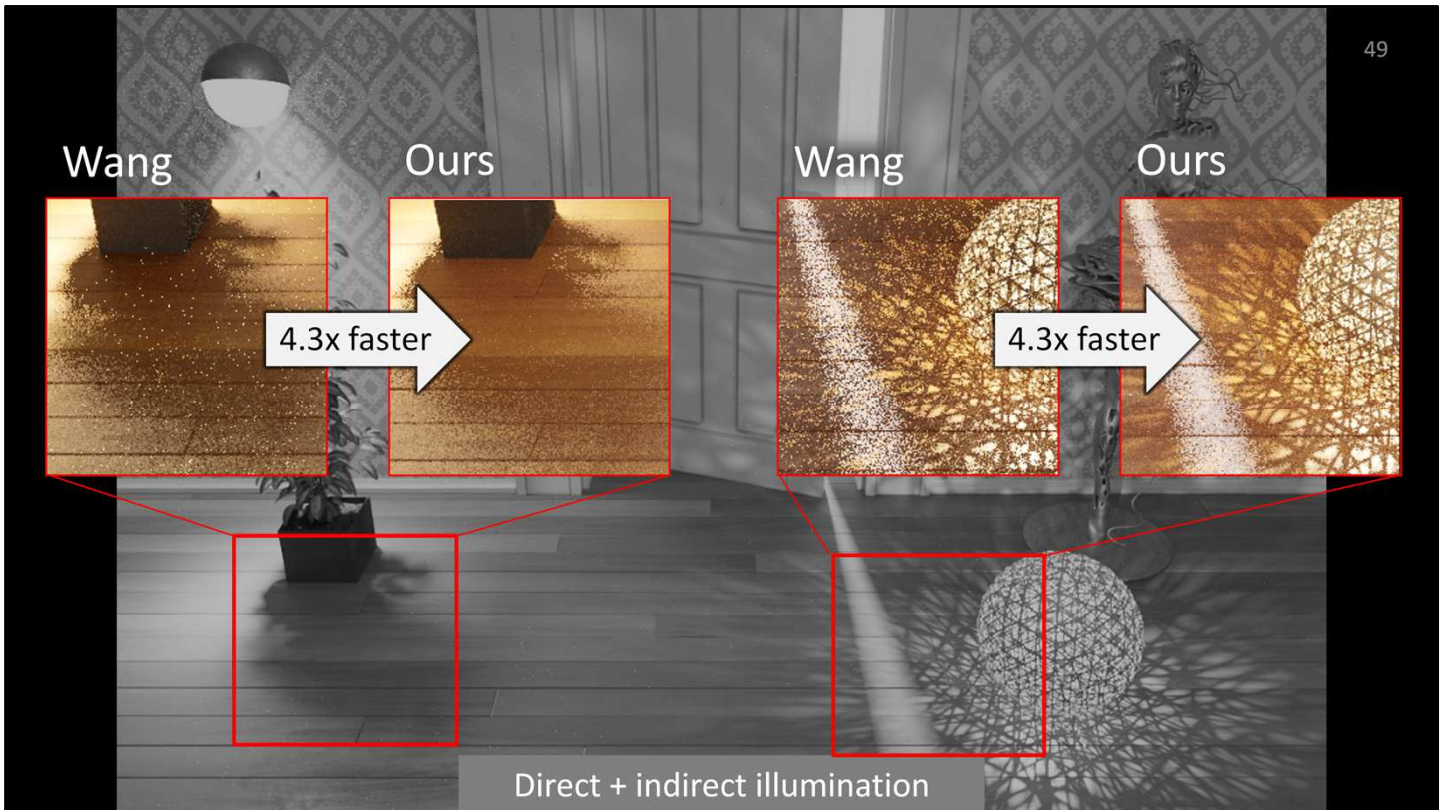
Tests

	Direct only	Direct + indirect
Simple occlusion	✓	✓
Complex occlusion	✓	

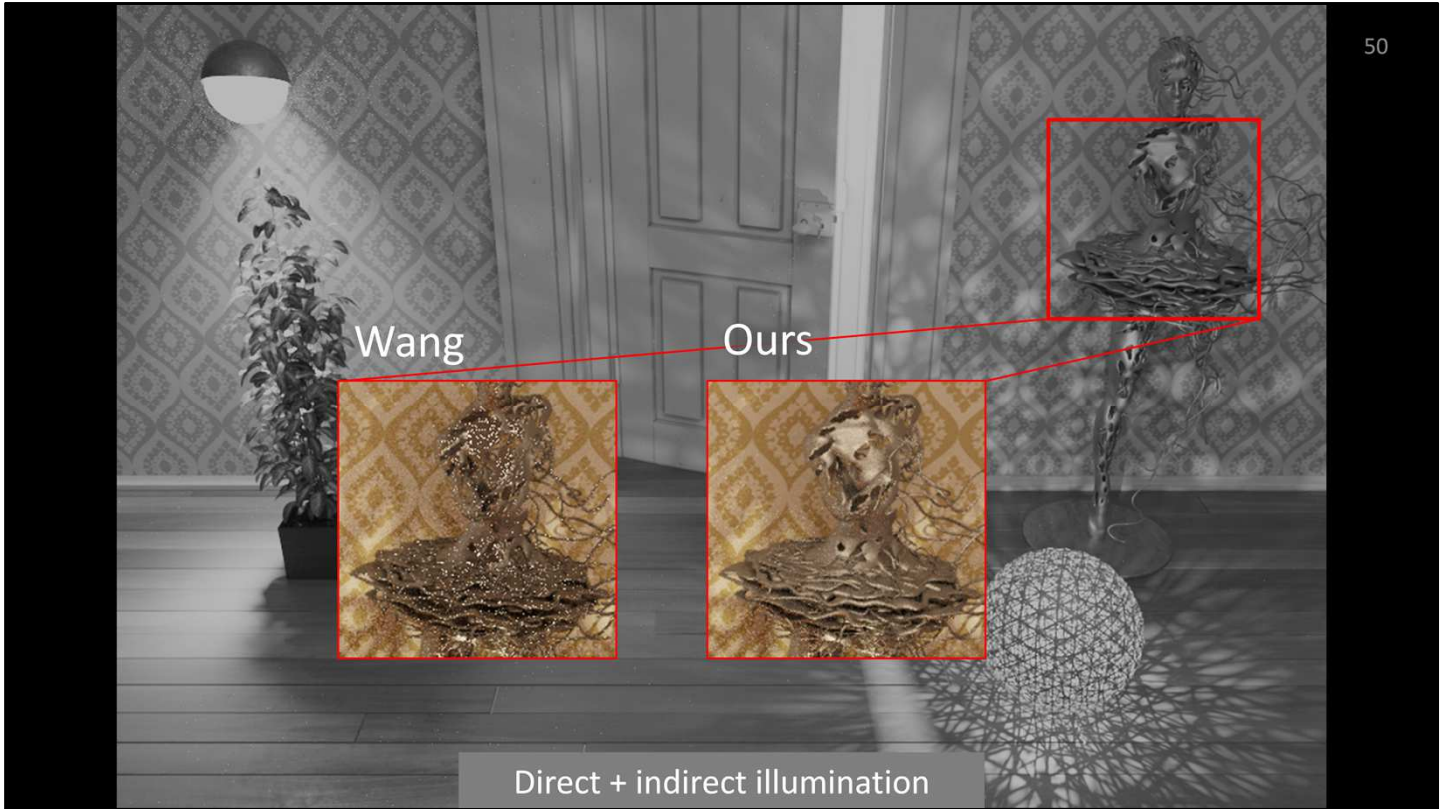
- Finally, let's test the complex occlusion also with the indirect component.



- If we take a look at the same scene...



- ...we see the direct illumination noise again dominates the complete image.
- Our method eliminates it and renders the complete image more than four times faster and without any artifacts.



- There is one more interesting area in this scene: The statue is made of glossy metal, and even though our method does not take the surface BRDF into account, it performs significantly better even there.

Tests

	Direct only	Direct + indirect
Simple occlusion	✓	✓
Complex occlusion	✓	✓


- Grid resolution



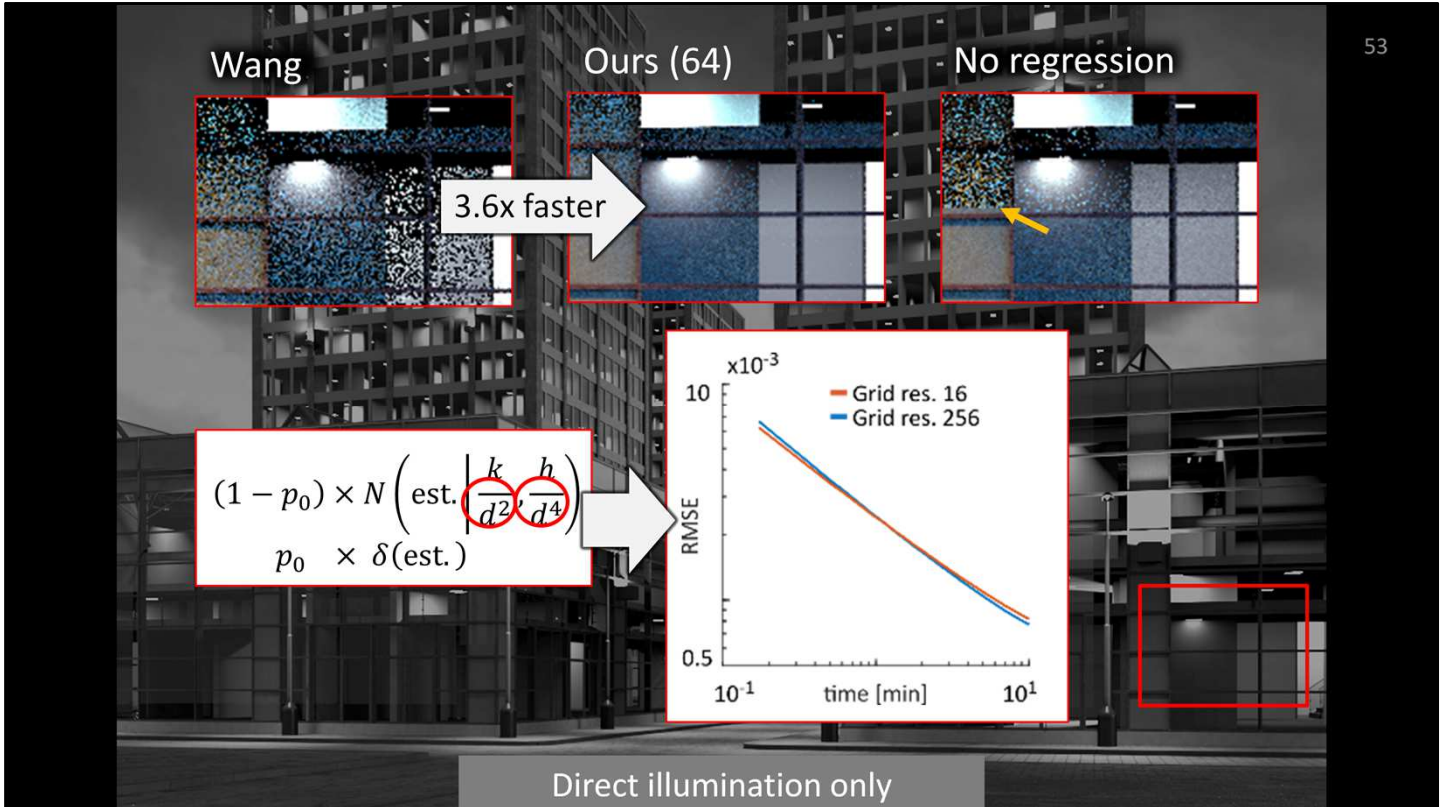
Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Recall that we divide the scene into regions by a uniform grid of a fixed resolution. We now test how this resolution affects the algorithm's performance.



- For that purpose we use this relatively large scene containing a many lights.



- With our default choice of 64 regions per the shortest grid dimension, our method performs more than three times faster than the method of Wang et al.
- So what about other grid resolutions?
- As it turns out, our method rather insensitive to the actual grid resolution. And so even much smaller as well as much higher resolutions all perform roughly the same.
- We hypothesize that this is due to the regression modeling of the distance falloff.
- Without the regression model we would have to use a much higher grid resolution, otherwise we would see sudden noise transitions between regions.

Tests

	Direct only	Direct + indirect
Simple occlusion	✓	✓
Complex occlusion	✓	✓

- Grid resolution ✓
- Temporal coherence



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Ok, so our method works well in all our static scenes.
- Now to our final test. It will verify temporal coherence in animations.



- We now show a short animation of this scene. The only thing that changes is the position of the sun. It starts behind the scene and will go towards the camera.
- Notice on the right that the non-adaptive method of Wang et al. works well on the back wall because the sun is behind it, so it is culled and no samples are wasted. As the sun moves towards the camera, it starts contributing to the back wall, cannot be culled anymore, numerous samples are allocated to it because of its high intensity, resulting in severe undersampling of the interior lights
- On the other hand, our method works well all the time without any artifacts or blinking.

Contribution

- **Bayesian framework for robust adaptivity/guiding**
- Optimal sampling distribution
- Algorithm for direct illumination
 - Unbiased, adaptive, robust
 - Easy to integrate



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- To conclude, the main contribution of our work is a Bayesian framework for adaptive/guided Monte Carlo quadrature.
- It enables exploiting the large potential of the adaptive sampling approach in Monte Carlo methods, while avoiding the biggest weakness of previous attempts – the lack of robustness.
- We applied this framework on the problem of direct illumination sampling.
- In the process we derived the optimal sampling distribution, taking variance into account, and developed an unbiased adaptive direct illumination algorithm with online learning of light sampling distributions. It is easily to integrate into a path tracer and suitable for interactive rendering (the up front cost is minimized as all learning happens on the fly during rendering).
- Our new framework is not limited to the direct illumination though and we are certain that other applications of adaptive sampling will benefit from it as well and it opens the path for many other tools of statistical machine learning (such as full Bayes or variational Bayes).

Acknowledgments

- Ludvík Koutný (a.k.a. rawalanche)
- Funding
 - Charles University: GAUK 1172416, SVV-2017-260452
 - Czech Science Foundation: 16-18964S, 19-07626S.



Vévoda, Kondapaneni, Křivánek - Bayesian online regression for adaptive illumination sampling



- Finally, We would like to thank Ludvík Koutný for modeling the test scenes and Charles University and Czech Science Foundation for supporting our work.

Take home message

Machine Learning | Bayesian modeling

=

**Excellent framework for
guided/adaptive Monte Carlo**



Computer
Graphics
Charles
University

And that's all. Thank you for your attention.