Learning Light Transport the Reinforced Way

Ken Dahm and Alexander Keller

Abstract We show that the equations of reinforcement learning and light transport simulation are related integral equations. Based on this correspondence, a scheme to learn importance while sampling path space is derived. The new approach is demonstrated in a consistent light transport simulation algorithm that uses reinforcement learning to progressively learn where light comes from. As using this information for importance sampling includes information about visibility, too, the number of light transport paths with zero contribution is dramatically reduced, resulting in much less noisy images within a fixed time budget.

1 Introduction

One application of light transport simulation is the computational synthesis of images that cannot be distinguished from real photographs. In such simulation algorithms [25], light transport is modeled by a Fredholm integral equation of the second kind and pixel colors are determined by estimating functionals of the solution of the Fredholm integral equation. The estimators are averages of contributions of sampled light transport paths that connect light sources and camera sensors.

Compared to reality, where photons and their trajectories are abundant, a computer may only consider a tiny fraction of path space, which is one of the dominant reasons for noisy images. It is therefore crucial to efficiently find light transport paths that have an important contribution to the image. While a lot of research in computer graphics has been focussing on importance sampling [19, 4, 3, 1, 24], for long there has not been a simple and efficient online method that can substantially reduce the number of light transport paths with zero contribution [33].

Ken Dahm · Alexander Keller

NVIDIA, Fasanenstr. 81, 10623 Berlin, Germany e-mail: ken.dahm@gmail.com, e-mail: keller.alexander@gmail.com

$$p \sim L_e f_s \cos \theta$$
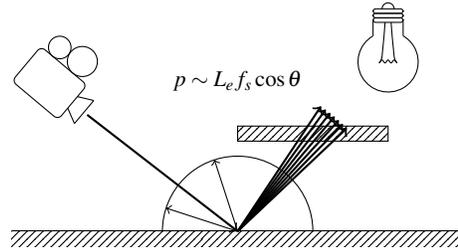
Fig. 1 In the illustration, radiance is integrated by sampling proportional to the product of emitted radiance L_e and the bidirectional scattering distribution function f_s representing the physical surface properties taking into account the fraction of radiance that is incident perpendicular to the surface, which is the cosine of the angle θ between the surface normal and the direction of incidence. As such importance sampling does not consider blockers, light transport paths with zero contributions cannot be avoided unless visibility is considered.

The majority of zero contributions are caused by unsuitable local importance sampling using only a factor instead of the complete integrand (see Fig. 1) or by trying to connect vertices of light transport path segments that are occluded, for example shooting shadow rays to light sources or connecting path segments starting both from the light sources and the camera sensors. An example for this inefficiency has been investigated early on in computer graphics [30, 31]: The visible part of the synthetic scene shown in Fig. 4 is lit through a door. By closing the door more and more the problem can be made arbitrarily more difficult to solve.

We therefore propose a method that is based on reinforcement learning [28] and allows one to sample light transport paths that are much more likely to connect lights and sensors. Complementary to first approaches of applying machine learning to image synthesis [33], in Sec. 2 we show that light transport and reinforcement learning can be modeled by the same integral equation. As a consequence, importance in light transport can be learned using any light transport algorithm.

Deriving a relationship between reinforcement learning and light transport simulation, we establish an automatic importance sampling scheme as introduced in Sec. 3. Our approach allows for controlling the memory footprint, for suitable representations of importance does not require preprocessing, and can be applied during image synthesis and/or across frames, because it is able to track distributions over time. A second parallel between temporal difference learning and next event estimation is pointed out in Sec. 4.

As demonstrated in Sec. 5 and shown in Fig. 8, already a simple implementation can dramatically improve light transport simulation. The efficiency of the scheme is based on two facts: Instead of shooting towards the light sources, we are guiding light transport paths to where the light comes from, which effectively shortens path length, and we learn importance from a smoothed approximation instead from higher variance path space samples [19, 10, 23].
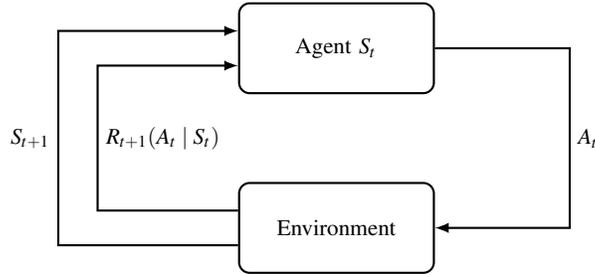
Fig. 2 The setting for reinforcement learning: At time t, an agent is in state S_t and takes an action A_t, which after interaction with the environment brings him to the next state S_{t+1} with a scalar reward R_{t+1}.

2 Identifying Q-Learning and Light Transport

The setting of reinforcement learning [28] is depicted in Fig. 2: An agent takes an action thereby transitioning to the resulting next state and receiving a reward. In order to maximize the reward, the agent has to learn which action to choose in what state. This process very much resembles how humans learn.

Q-learning [39] is a model free reinforcement learning technique. Given a set of states S and a set of actions A, it determines a function $Q(s,a)$ that for any $s \in S$ values taking the action $a \in A$. Thus given a state s, the action a with the highest value may be selected next and

$$Q(s,a) = (1-\alpha) \cdot Q(s,a) + \alpha \cdot \left(r(s,a) + \gamma \cdot \max_{a' \in A} Q(s',a') \right) \tag{1}$$

may be updated by a fraction of $\alpha \in [0,1]$, where $r(s,a)$ is the reward for taking the action resulting in a transition to a state s'. In addition, the maximum Q-value of possible actions in s' is considered and discounted by a factor of $\gamma \in [0,1)$.

Instead of taking into account only the best valued action,

$$Q(s,a) = (1-\alpha) \cdot Q(s,a) + \alpha \cdot \left(r(s,a) + \gamma \cdot \sum_{a' \in A} \pi(s',a') Q(s',a') \right)$$

averages all possible actions in s' and weighs their values $Q(s',a')$ by a transition kernel $\pi(s',a')$, which is a strategy called *expected SARSA* [28, Sec.6.6]. This is especially interesting, as later it will turn out that always selecting the "best" action does not perform as well as considering all options (see Fig. 4). For a continuous space A of actions, we then have

$$Q(s,a) = (1-\alpha) \cdot Q(s,a) + \alpha \cdot \left(r(s,a) + \gamma \cdot \int_A \pi(s',a') Q(s',a') da' \right). \tag{2}$$

On the other hand, the radiance

$$L(x,\omega) = L_e(x,\omega) + \int_{\mathscr{S}^+(x)} L(h(x,\omega_i), -\omega_i) f_s(\omega_i, x, \omega) \cos\theta_i d\omega_i \qquad (3)$$

in a point x on a surface into direction ω is modeled by a Fredholm integral equation of the second kind. L_e is the source radiance and the integral accounts for all radiance that is incident over the hemisphere $\mathscr{S}^+(x)$ aligned by the surface normal in x and transported into direction ω. The hitpoint function $h(x, \omega_i)$ traces a ray from x into direction ω_i and returns the first surface point intersected. The radiance from this point is attenuated by the bidirectional scattering distribution function f_s, where the cosine term of the angle θ_i between surface normal and ω_i accounts for only the fraction that is perpendicular to the surface.

A comparison of Eqn. 2 for $\alpha = 1$ and Eqn. 3 reveals structural similarities of the formulation of reinforcement learning and the light transport integral equation, respectively, which lend themselves to matching terms: Interpreting the state s as a location $x \in \mathbb{R}^3$ and an action a as tracing a ray from location x into direction ω resulting in the point $y := h(x, \omega)$ corresponding to the state s', the reward term $r(s, a)$ can be linked to the emitted radiance $L_e(y, -\omega) = L_e(h(x, \omega), -\omega)$ as observed from x. Similarly, the integral operator can be applied to the value Q, yielding

$$Q(x,\omega) = L_e(y, -\omega) + \int_{\mathscr{S}^+(y)} Q(y, \omega_i) f_s(\omega_i, y, -\omega) \cos\theta_i d\omega_i, \qquad (4)$$

where we identified the discount factor γ multiplied by the policy π and the bidirectional scattering distribution function f_s. Taking a look at the geometry and the physical meaning of the terms, it becomes obvious that Q in fact must be the radiance $L_i(x, \omega)$ incident in x from direction ω and in fact is described by a Fredholm integral equation of the second kind - like the light transport equation 3.
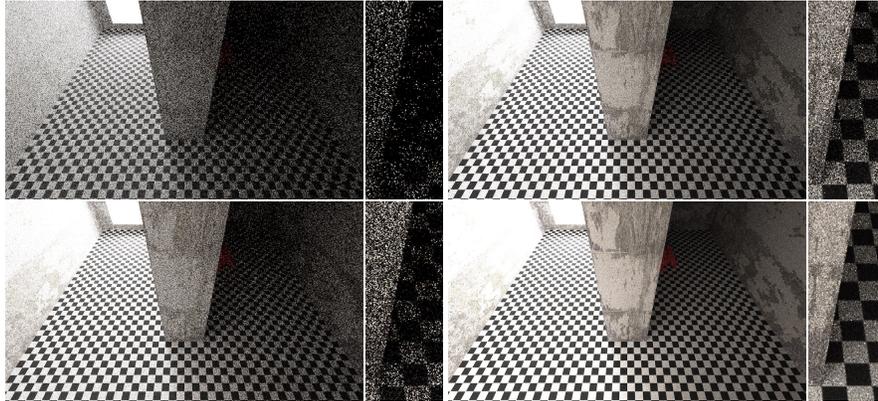
Fig. 3 Comparison of simple path tracing without (left) and with (right) reinforcement learning importance sampling. The top row is using 8 paths per pixel, while 32 are used for the bottom row. The challenge of the scene is the area light source on the left indirectly illuminating the right part of the scene. The enlarged insets illustrate the reduction of noise level.

a

b

c

d

e

f
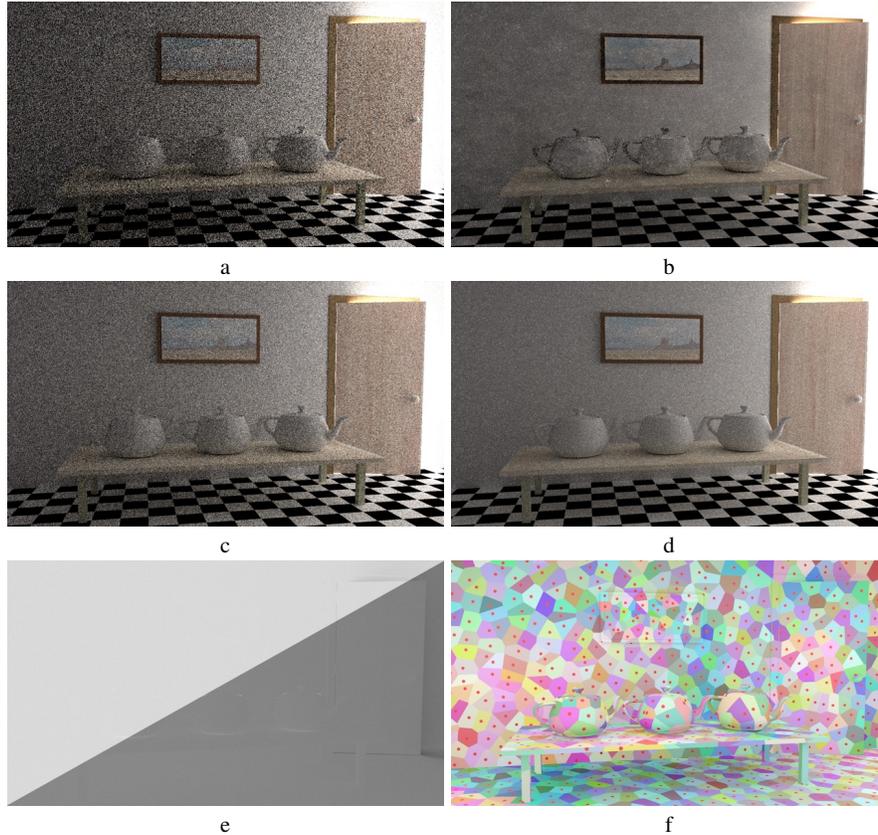
Fig. 4 Comparison at 1024 paths per pixel (the room behind the door is shown in Fig. 5): a) A simple path tracer with cosine importance sampling, b) the Kelemen variant of the Metropolis light transport algorithm, c) scattering proportional to Q, while updating Q with the maximum as in Eqn. 1 and d) scattering proportional to Q weighted by the bidirectional scattering distribution function and updating accordingly by Eqn. 5. The predominant reinforcement approach of always taking the best next action is inferior to selecting the next action proportional to Q, i.e. considering all alternatives. A comparison to the Metropolis algorithm reveals much more uniform lighting, especially much more uniform noise and the lack of the typical splotches. e) The average path length of path tracing (above image diagonal) is about 215, while with reinforcement learning it amounts to an average of 134. The average path length thus is reduced by 40% in this scene. f) Discretized hemispheres to approximate Q are stored in points on the scene surfaces determined by samples of the Hammersley low discrepancy point set. Retrieving Q for a query point results in searching for the nearest sample of Q that has a similar normal to the one in the query point (see especially the teapot handles). The red points indicate where in the scene hemispheres to hold the Q_k are stored. The colored areas indicate their corresponding Voronoi cells. Storing the Q_k in this example requires about 2 MBytes of memory. Scene courtesy (cc) 2013 Miika Aittala, Samuli Laine, and Jaakko Lehtinen (https://mediatech.aalto.fi/publications/graphics/GMLT/).

3 Q-Learning while Path Tracing

In order to synthesize images, we need to compute functionals of the radiance equation 3, i.e. project the radiance onto the image plane. For the purpose of this article, we start with a simple forward path tracer [25, 14]: From a virtual camera, rays are traced through the pixels of the screen. Upon their first intersection with the scene geometry, the light transport path is continued into a scattering direction determined according to the optical surface properties. Scattering and ray tracing are repeated until a light source is hit. The contribution of this complete light transport path is added to the pixel pierced by the initial ray of this light transport path when started at the camera.

In this simple form, the algorithm exposes quite some variance as can be seen in the images on the left in Fig. 3. This noise may be reduced by importance sampling. We therefore progressively approximate Eqn. 4 using reinforcement learning: Once a direction has been selected and a ray has been traced by the path tracer,

$$
\begin{aligned}
Q'(x, \omega) = (1 - \alpha) \cdot Q(x, \omega) \\
+ \alpha \cdot \left(L_e(y, -\omega) + \int_{\mathscr{S}^+(y)} Q(y, \omega_i) f_s(\omega_i, y, -\omega) \cos \theta_i d\omega_i \right)
\end{aligned}
\tag{5}
$$

is updated using a learning rate α. The probability density function resulting from normalizing Q in turn is used for importance sampling a direction to continue the path. As a consequence more and more light transport paths are sampled that contribute to the image. Computing a global solution to Q in a preprocess would not allow for focussing computations on light transport paths that contribute to the image.

3.1 Implementation

Often, approximations to Q are tabulated for each pair of state and action. In computer graphics, there are multiple choices to represent radiance and for the purpose of this article, we chose the data structure as used for irradiance volumes [6] to approximate Q. Fig. 5 shows an exemplary visualization of such a discretization during rendering: For selected points y in space, the hemisphere is stratified and one value $Q_k(y)$ is stored per sector, i.e. stratum k. Fig. 4f illustrates the placement of probe centers y, which results from mapping a two-dimensional low discrepancy sequence onto the scene surface.

Now the integral

$$
\int_{\mathscr{S}^+(y)} Q(y, \omega_i) f_s(\omega_i, y, -\omega) \cos \theta_i d\omega_i \approx \frac{2\pi}{n} \sum_{k=0}^{n-1} Q_k(y) f_s(\omega_k, y, -\omega) \cos \theta_k
$$

in Eqn. 5 can be estimated by using each one uniform random direction ω_k in each stratum k, where θ_k is the angle between the surface normal in y and ω_k.
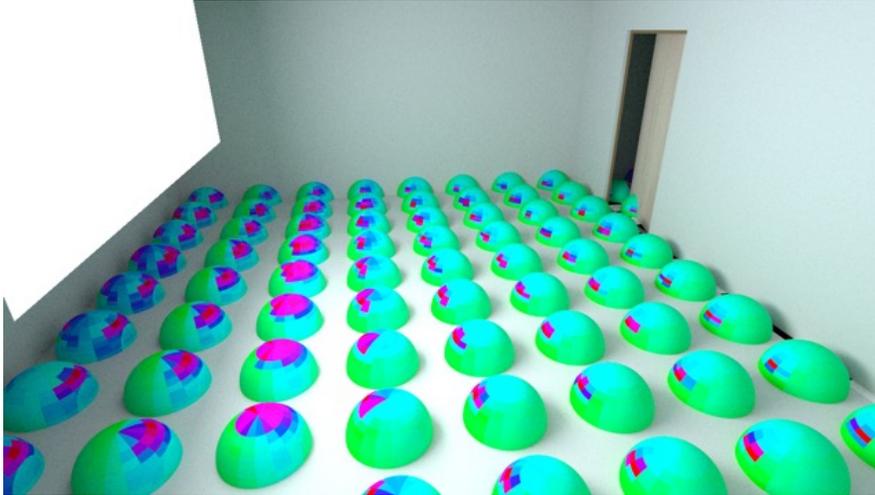
Fig. 5 The image shows parts of an example discretization of Q by a grid, where hemispheres are uniformly distributed across the ground plane. The false colors indicate magnitude, where small values are green and large values are red. The large values on each hemisphere point towards the part of the scene, where the light is coming from. For example, under the big area light source on the left, most radiance is incident as reflected radiance from the wall opposite to the light source.

The method has been implemented in an importance driven forward path tracer as shown in Alg. 1: Only two routines for updating Q and selecting a scattering direction proportional to Q need to be added. Normalizing the Q in a point y then results in a probability density that is used for importance sampling during scattering by inverting the cumulative distribution function. In order to guarantee ergodicity, meaning that every light transport path remains possible, all $Q(y)$ are initialized to be positive, for example by a uniform probability density or proportional to a factor of the integrand (see Fig. 1). When building the cumulative distribution functions in parallel every accumulated frame, values below a small positive threshold are replaced by the threshold.

The parameters exposed by our implementation are the resolution of the discretization and the learning rate α.

3.2 Consistency

It is desirable to craft consistent rendering algorithms [14], because then all renderer introduced artifacts, like for example noise, are guaranteed to vanish over time. This requires the $Q_k(y)$ to converge, which may be accomplished by a vanishing learning rate α.

In reinforcement learning [28], a typical approach is to count the number of visits to each pair of state s and action a and using

---

Algorithm 1: Augmenting a path tracer by reinforcement learning for importance sampling requires only two additions: The importance Q needs to be updated along the path and scattering directions are sampled proportional to Q as learned so far.

Function *pathTrace(camera, scene)*

    throughput $\leftarrow 1$
    ray $\leftarrow$ setupPrimaryRay(*camera*)
    for $i \leftarrow 0$ **to** ∞ **do**
        $(y, n) \leftarrow$ intersect(*scene, ray*)`//` $y := h(x, \omega)$
        `// addition 1: update `Q
        if $i > 0$ **then**
            $Q'(x, \omega) =$
            $(1 - \alpha)Q(x, \omega) + \alpha \left(L_e(y, -\omega) + \int_{\mathscr{S}_+^2(y)} f_s(\omega_i, y, -\omega) \cos \theta_i Q(y, \omega_i) d\omega_i \right)$
        if *isEnvironment(y)* **then**
            return *throughput* $\cdot$ getRadianceFromEnvironment(*ray, y*)
        else if *isAreaLight(y)*
            return *throughput* $\cdot$ getRadianceFromAreaLight(*ray, y*)
        `// addition 2: scatter proportional to `Q
        $(\omega, p_\omega, f_s) \leftarrow$ **sampleScatteringDirectionProportionalToQ**(*y*)
        throughput $\leftarrow$ *throughput* $\cdot f_s \cdot \cos(n, \omega) / p_\omega$
        ray $\leftarrow (y, \omega)$

$$\alpha(s, a) = \frac{1}{1 + \text{visits}(s, a)}.$$

The method resembles the one used to make progressive photon mapping consistent [7], where consistency has been achieved by decreasing the search radius around a query point every time a photon hits sufficiently close. Similarly, the learning rate may also depend on the total number of visits to a state s alone, or even may be chosen to vanish independently of state and action. Again, such approaches have been explored in consistent photon mapping [15].

While the $Q_k(y)$ converge, they do not necessarily converge to the incident radiance in Eqn. 4. First, as they are projections onto a basis, the $Q_k(y)$ at best only are an approximation of Q in realistic settings. Second, as the coefficients $Q_k(y)$ are learned during path tracing, i.e. image synthesis, and used for importance sampling, it may well happen that they are not updated everywhere at the same rate. Nevertheless, since all operators are linear, the number of visits will be proportional to the number of light transport paths [15] and consequently as long as $Q_k(y) > 0$ whenever $L_i(y, \omega_i) > 0$ all $Q_k(y)$ will be updated eventually.

3.3 Learning while Light Tracing

For guiding light transport paths from the light sources towards the camera, the transported weight W of a measurement (see [29]), i.e. the characteristic function of the image plane, has to be learned instead of the incident radiance Q. As W is the adjoint of Q, the same data structures may be used for its storage. Learning both Q and W allows one to implement bidirectional path tracing [29] with reinforcement learning for importance sampling to guide both light and camera path segments including visibility information for the first time. Note that guiding light transport paths this way may reach efficiency levels that even can make bidirectional path tracing and multiple importance sampling redundant [33] in many common cases.
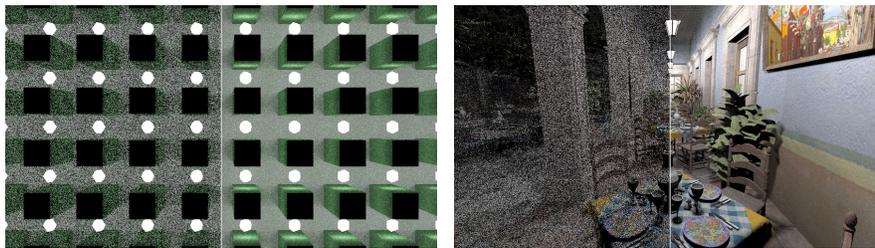
Fig. 6 Two split-image comparisons of uniformly selecting area light sources and selection using temporal difference learning, both at 16 paths per pixel. The scene on the left has 5000 area light sources, whereas the scene on the right has about 15000 (San Miguel scene courtesy Guillermo M. Leal Llaguno (http://www.evvisual.com/)).

4 Temporal Difference Learning and Next Event Estimation

Besides the known shortcomings of (bidirectional) path tracing [18, Sec.2.4 Problem of insufficient techniques], the efficiency may be restricted by the approximation quality of Q: For example, the smaller the light sources, the finer the required resolution of Q to reliably guide rays to hit a light source. This is where next event estimation may help [32, 16, 5].

Already in [38] the contribution of light sources has been "learned": A probability per light source has been determined by the number of successful shadow rays divided by the total number of shadow rays shot. This idea has been refined subsequently [17, 37, 2, 36].

For reinforcement learning, the state space may be chosen as a regular grid over the scene, where in each grid cell c for each light source l a value $V_{c,l}$ is stored that is initialized with zero. Whenever a sample on a light source l is visible to a point x to be illuminated in the cell c upon next event estimation, its value

$$V'_{c,l} = (1-\alpha)V_{c,l} + \alpha \cdot \|C_l(x)\|_\infty \tag{6}$$

is updated using the norm of the contribution $C_l(x)$. Building a cumulative distribution function from all values $V_{c,l}$ within a cell c, light may be selected by importance sampling. Fig. 6 shows the efficiency gain of this reinforcement learning method over uniform light source selection for 16 paths per pixel.

It is interesting to see that this is another relation to reinforcement learning: While the Q-learning equation 5 takes into account the values of the next, non-terminal state, the next state in event estimation is always a terminal state and Q-learning coincides with plain temporal difference learning [27] as in equation 6.

4.1 Learning Virtual Point Light Sources

The vertices generated by tracing photon trajectories (see Sec. 3.3) can be considered a photon map [11] and may be used in the same way. Furthermore, they may be used as a set of virtual point light sources for example the instant radiosity [13] algorithm.

Continuously updating and learning the measurement contribution function W [29] across frames and using the same seed for the pseudo- or quasi-random sequences allows for generating virtual point light sources that expose a certain coherency over time, which reduces temporal artifacts when rendering animations with global illumination.

4.2 Learning Environment Lighting

Rendering sun and sky is usually done by distributing samples proportional to the brightness of pixels in the environment texture. More samples should end up in brighter regions, which is achieved by constructing and sampling from a cumulative distribution function, for example using the alias method [35]. Furthermore, the sun may be separated from the sky and simulated separately. The efficiency of such importance sampling is highly dependent on occlusion, i.e. what part of the environment can be seen from the point to be shaded (see Fig. 1).

Similar to Sec. 3.1 and in order to consider the actual contribution including occlusion, an action space is defined by partitioning the environment map into tiles and learning the importance per tile. Fig. 7 shows the improvement for an example setting.
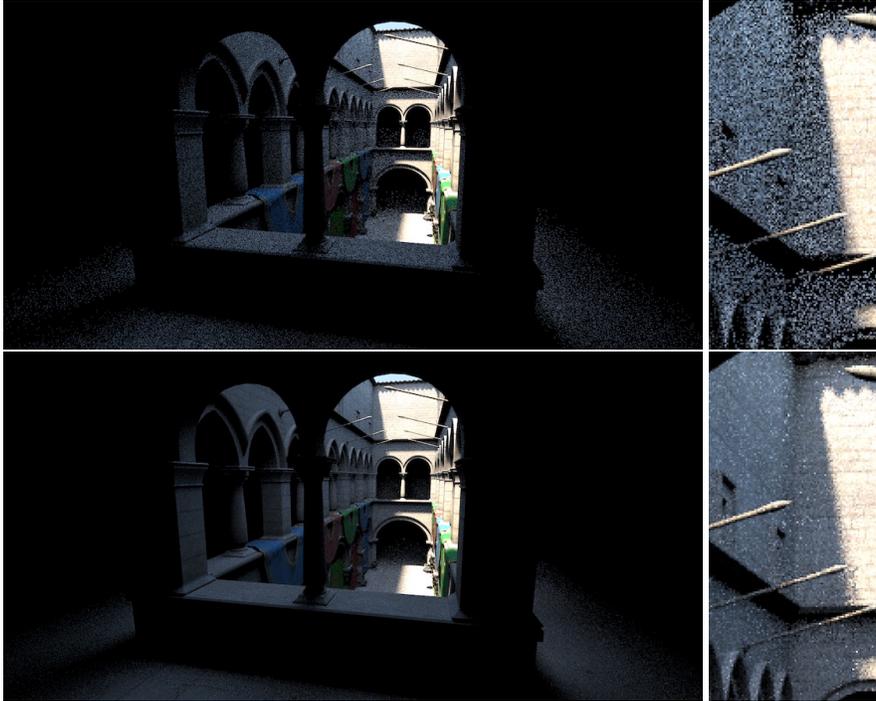
Fig. 7 Sun and sky illumination at 32 paths per pixel. Top: simple importance sampling considering only the environment map as a light source. Bottom: Importance sampling with reinforcement learned importance. The enlargements on the right illustrate the improved noise reduction. Scene courtesy Frank Meinl, Crytek (http://graphics.cs.williams.edu/data/meshes/crytek-sponza-copyright.html).

5 Results and Discussion

Fig. 4 compares the new reinforcement learning algorithm to common algorithms: For the same budget of light transport paths, the superiority over path tracing with importance sampling according to the reflection properties is obvious. A comparison with the Metropolis algorithm for importance sampling [31, 12] reveals much more uniform noise lacking the typical splotchy structure inherent with the local space exploration of Metropolis samplers. Note, however, that the new reinforcement learning importance sampling scheme could as well be combined with Metropolis sampling. Finally, updating Q by Eqn. 1, i.e. the "best possible action" strategy is inferior to using the weighted average of all possible next actions according to Eqn. 5. In light transport simulation this is not surprising, as the deviation of the integrand from its estimated maximum very often is much larger than from a piecewise constant approximation.

The big gain in quality is due to the dramatic reduction of zero contribution light transport paths (see Fig. 8), even under complex lighting. In Figs. 4a-d, the
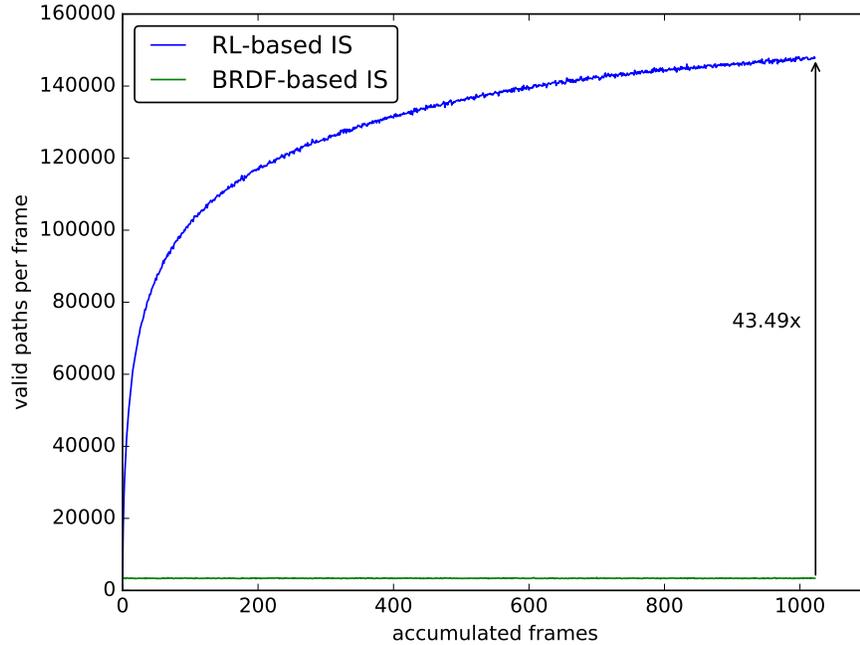
Fig. 8 Using reinforcement learning (RL), the number of paths actually connecting to a light source is dramatically improved over classic importance sampling (IS) using only the bidirectional scattering distribution function for importance sampling. As a result, more non-zero contributions are accumulated for the same number of paths, see also Fig. 4.

same number of paths has been used. In each iteration, for path tracing with and without reinforcement learning one path has been started per pixel, while for the Metropolis variant the number of Markov chains equals the number of pixels of the image. Rendering the image at 1280x720 pixels, each iteration takes 41ms for path tracing, 49ms for Metropolis light transport [31, 12], and 51ms for the algorithm with reinforcement learned importance sampling. Hence the 20% overhead is well paid off by the level of noise reduction.

Shooting towards where the radiance comes from naturally shortens the average path length as can be seen in Fig. 4e. Based on the approach to guide light paths using a pre-trained Gaussian mixture model [33] to represent probabilities, in [34] in addition the density of light transport paths is controlled across the scene using splitting and Russian roulette. These ideas have the potential to further improve the efficiency of our approach.

While the memory requirements for storing our data structure for Q are small, the data structure is not adaptive. An alternative is an adaptive hierarchical approximation to Q as used in [19, 23]. Yet, another variant would be learning parameters for lobes to guide light transport paths [1]. In principle any data structure that has been used in graphics to approximate irradiance or radiance is a candidate. Which data structure

and what parameters are best, may depend on the scene to be rendered. For example, using discretized hemispheres limits the resolution with respect to solid angle. If the resolution is chosen too fine, learning is slow, if the resolution is to coarse, convergence is slow.

Given that Q asymptotically approximates the incident radiance L_i, it is worthwhile to investigate how it can be used for the separation of the main part as explored in [19, 24] to further speed up light transport simulation or even as an alternative to importance sampling.

Beyond what we explore, path guiding has been extended to consider product importance sampling [9] and reinforcement learning [28] offers more policy evaluation strategies to consider.

6 Conclusion

Guiding light transport paths has been explored in [19, 10, 4, 24, 1, 33, 23]. However, key to our approach is that by using a representation of Q in Eqn. 5 instead of solving the equation by recursion, i.e. a Neumann series, Q may be learned much faster and in fact during sampling light transport paths without any preprocess. This results in a new algorithm to increase the efficiency of path tracing by approximating importance using reinforcement learning during image synthesis. Identifying Q-learning and light transport, heuristics have been replaced by physically based functions, and the only parameters that the user may control are the learning rate and the discretization of Q.

The combination of reinforcement learning and deep neural networks [22, 8, 20, 21] is an obvious avenue of future research: Representing the radiance on hemispheres already has been successfully explored [26] and the interesting question is how well Q can be represented by neural networks.

References

1. Bashford-Rogers, T., Debattista, K., Chalmers, A.: A significance cache for accelerating global illumination. Computer Graphics Forum **31**(6), 1837–1851 (2012) 1, 5, 6
2. Benthin, C., Wald, I., Slusallek, P.: A scalable approach to interactive global illumination. Computer Graphics Forum (Proc. Eurographics 2003) **22**(3), 621–629 (2003) 4
3. Cline, D., Adams, D., Egbert, P.: Table-driven adaptive importance sampling. Computer Graphics Forum **27**(4), 1115–1123 (2008) 1
4. Dutré, P., Willems, Y.: Potential-driven Monte Carlo particle tracing for diffuse environments with adaptive probability functions. In: Rendering Techniques 1995 (Proc. 6th Eurographics Workshop on Rendering), pp. 306–315. Springer (1995) 1, 6

5. Estevez, C., Kulla, C.: Importance sampling of many lights with adaptive tree splitting. In: ACM SIGGRAPH 2017 Talks, SIGGRAPH '17, pp. 33:1–33:2. ACM (2017) 4
6. Greger, G., Shirley, P., Hubbard, P., Greenberg, D.: The irradiance volume. IEEE Computer Graphics and Applications **18**(2), 32–43 (1998) 3.1
7. Hachisuka, T., Ogaki, S., Jensen, H.: Progressive photon mapping. ACM Transactions on Graphics **27**(5), 130:1–130:8 (2008) 3.2
8. van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double Q-learning. CoRR **abs/1509.06461** (2015). URL http://arxiv.org/abs/1509.06461 6
9. Herholz, S., Elek, O., Vorba, J., Lensch, H., Křivánek, J.: Product importance sampling for light transport path guiding. In: Proceedings of the Eurographics Symposium on Rendering, EGSR '16, pp. 67–77. Eurographics Association (2016). DOI 10.1111/cgf.12950. URL https://doi.org/10.1111/cgf.12950 5
10. Jensen, H.: Importance driven path tracing using the photon map. In: P. Hanrahan, W. Purgathofer (eds.) Rendering Techniques 1995 (Proc. 6th Eurographics Workshop on Rendering), pp. 326–335. Springer (1995) 1, 6
11. Jensen, H.: Realistic Image Synthesis Using Photon Mapping. AK Peters (2001) 4.1
12. Kelemen, C., Szirmay-Kalos, L., Antal, G., Csonka, F.: A simple and robust mutation strategy for the Metropolis light transport algorithm. Computer Graphics Forum **21**(3), 531–540 (2002) 5
13. Keller, A.: Instant radiosity. In: SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques, pp. 49–56 (1997) 4.1
14. Keller, A.: Quasi-Monte Carlo image synthesis in a nutshell. In: J. Dick, F. Kuo, G. Peters, I. Sloan (eds.) Monte Carlo and Quasi-Monte Carlo Methods 2012, pp. 203–238. Springer (2013) 3, 3.2
15. Keller, A., Binder, N.: Deterministic consistent density estimation for light transport simulation. In: J. Dick, F. Kuo, G. Peters, I. Sloan (eds.) Monte Carlo and Quasi-Monte Carlo Methods 2012, pp. 467–480. Springer (2013) 3.2
16. Keller, A., Wächter, C., Raab, M., Seibert, D., Antwerpen, D., Korndörfer, J., Kettner, L.: The Iray light transport simulation and rendering system. CoRR **abs/1705.01263** (2017). URL http://arxiv.org/abs/1705.01263 4
17. Keller, A., Wald, I.: Efficient importance sampling techniques for the photon map. In: Proc. VISION, MODELING, AND VISUALIZATION, pp. 271–279. IOS Press (2000) 4
18. Kollig, T., Keller, A.: Efficient bidirectional path tracing by randomized quasi-Monte Carlo integration. In: H. Niederreiter, K. Fang, F. Hickernell (eds.) Monte Carlo and Quasi-Monte Carlo Methods 2000, pp. 290–305. Springer (2002) 4
19. Lafortune, E., Willems, Y.: A 5D tree to reduce the variance of Monte Carlo ray tracing. In: P. Hanrahan, W. Purgathofer (eds.) Rendering Techniques 1995 (Proc. 6th Eurographics Workshop on Rendering), pp. 11–20. Springer (1995) 1, 1, 5, 6
20. Lillicrap, T., Hunt, J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. CoRR **abs/1509.02971** (2015). URL http://arxiv.org/abs/1509.02971 6
21. Mnih, V., Badia, A., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., Kavukcuoglu, K.: Asynchronous methods for deep reinforcement learning. CoRR **abs/1602.01783** (2016). URL http://arxiv.org/abs/1602.01783 6
22. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., Riedmiller, M.: Playing Atari with deep reinforcement learning. CoRR **abs/1312.5602** (2013). URL http://arxiv.org/abs/1312.5602 6
23. Müller, T., Gross, M., Novák, J.: Practical path guiding for efficient light-transport simulation. In: Proceedings of the Eurographics Symposium on Rendering (2017) 1, 5, 6
24. Pegoraro, V., Brownlee, C., Shirley, P., Parker, S.: Towards interactive global illumination effects via sequential Monte Carlo adaptation. In: Proceedings of the 3rd IEEE Symposium on Interactive Ray Tracing, pp. 107–114 (2008) 1, 5, 6
25. Pharr, M., Jacob, W., Humphreys, G.: Physically Based Rendering - From Theory to Implementation. Morgan Kaufmann, Third Edition (2016) 1, 3

26. Satilmis, P., Bashford-Rogers, T., Chalmers, A., Debattista, K.: A machine learning driven sky model. IEEE Computer Graphics and Applications pp. 1–9 (2016) 6

27. Sutton, R.: Learning to predict by the methods of temporal differences. Machine Learning **3**(1), 9–44 (1988) 4

28. Sutton, R., Barto, A.: Introduction to Reinforcement Learning, 2nd edn. MIT Press, Cambridge, MA, USA (2017) 1, 2, 2, 3.2, 5

29. Veach, E., Guibas, L.: Bidirectional estimators for light transport. In: Proc. 5th Eurographics Workshop on Rendering, pp. 147 – 161. Darmstadt, Germany (1994) 3.3, 4.1

30. Veach, E., Guibas, L.: Optimally combining sampling techniques for Monte Carlo rendering. In: SIGGRAPH '95 Proceedings of the 22nd annual conference on Computer graphics and interactive techniques, pp. 419–428 (1995) 1

31. Veach, E., Guibas, L.: Metropolis light transport. In: T. Whitted (ed.) Proc. SIGGRAPH 1997, Annual Conference Series, pp. 65–76. ACM SIGGRAPH, Addison Wesley (1997) 1, 5

32. Vévoda, P., Křivánek, J.: Adaptive direct illumination sampling. In: SIGGRAPH ASIA 2016 Posters, pp. 43:1–43:2. ACM, New York, NY, USA (2016) 4

33. Vorba, J., Karlík, O., Šik, M., Ritschel, T., Křivánek, J.: On-line learning of parametric mixture models for light transport simulation. ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014) **33**(4) (2014) 1, 1, 3.3, 5, 6

34. Vorba, J., Křivánek, J.: Adjoint-driven Russian roulette and splitting in light transport simulation. ACM Transactions on Graphics (Proceedings of SIGGRAPH 2016) **35**(4), 1–11 (2016) 5

35. Vose, M.: A linear algorithm for generating random numbers with a given distribution. IEEE Trans. on Software Engineering **17**(9), 972–975 (1991) 4.2

36. Wald, I., Benthin, C., Slusallek, P.: Interactive global illumination in complex and highly occluded environments. In: P. Christensen, D. Cohen-Or (eds.) Rendering Techniques 2003 (Proc. 14th Eurographics Workshop on Rendering), pp. 74–81 (2003) 4

37. Wald, I., Kollig, T., Benthin, C., Keller, A., Slusallek, P.: Interactive global illumination using fast ray tracing. In: P. Debevec, S. Gibson (eds.) Rendering Techniques 2002 (Proc. 13th Eurographics Workshop on Rendering), pp. 15–24 (2002) 4

38. Ward, G.: Adaptive shadow testing for ray tracing. In: 2nd Eurographics Workshop on Rendering. Barcelona, Spain (1991) 4

39. Watkins, C., Dayan, P.: Q-learning. Machine learning **8**(3), 279–292 (1992) 2