

Fig. 2. Illustration of delayed rejection-based algorithm. Mutating the primary sample  $\bar{u}$  using a first transition kernel  $Q_1$  yields a point  $\tilde{v}$  that maps to light path  $\bar{y} = S(\tilde{v})$  in red. In this example, the first stage proposal is rejected due to a significant (and unpredictable) change in the structure of the light path. A new vector  $\bar{w}$  is then proposed from  $\bar{u}$  using a second Markov kernel  $Q_2$ , producing path  $\bar{z}$  (in green) that is more likely to be accepted.

Metropolis et al. 1953] to the path integral [Veach 1997]—an integral over *path space*, the infinite-dimensional space of light transport paths. The statistical dependence intrinsic to Markov chains allows MLT-based techniques to locally explore high-contribution regions in path space, as they are encountered, by sequentially perturbing the path structure according to *mutation strategies*.

Despite advances in MCMC light transport [Bitterli and Jarosz 2019; Šik and Krivánek 2018], designing efficient, robust and general-purpose mutation strategies remains an open problem. Smaller perturbations can improve acceptance rates, but not without often slowing the exploration of the salient regions in state space. Alternatively, larger perturbations may reduce correlation between samples, but these proposals are often rejected as they tend to escape high density regions, leading to chains getting “trapped” and variance increase in the resulting estimator [Ashikhmin et al. 2001]. Another form of this trade-off appears when considering the runtime cost of transitions: intricate mutation strategies can be more effective at equal sample count while typically requiring more computation; applying such costly transitions everywhere is wasteful in simple scenarios, however completely abandoning them could lead to under-exploration of important regions. Our work focuses on designing efficient, flexible proposal mechanisms in such cases.

We present a two-stage proposal mechanism that automatically balances local exploration and computational efficiency. Our approach extends *delayed rejection* [Green and Mira 2001; Tierney and Mira 1999], a method for sequentially combining many transition kernels in order to propose more robust perturbations. Delayed rejection exploits prioritization by proposing bolder or less costly transitions at a first stage before falling-back to more timid or expensive kernels upon failure (Figure 2). By augmenting proposal mechanisms with such a “safety net”, our method can increase acceptance rates without compromising exploration of the state space,

especially in complex path sampling scenarios. A naive implementation of delayed rejection can lead to zero-acceptance in the second stage with a Markov kernel lacking global support (e.g., *Kelemen-style* mutations [Kelemen et al. 2002]). We thus introduce a novel combination of proposals to address this problem: our *pairwise orbital* mutation conditions on both the current and first rejected states to ensure that all states remain close to each other, thus eliminating degenerate cases that can drag acceptance rates down. This solution is general and straightforward to implement atop *any* primary sample space method.

We demonstrate the benefits of our method by integrating it into three algorithms: PSSMLT [Kelemen et al. 2002], MMLT [Hachisuka et al. 2014] and H2MC [Li et al. 2015] (Figure 1). Here, our two-stage variants outperform each of their original one-stage counterparts, resulting in smoother results on a set of challenging scenes with different lighting, geometry and material configurations.

Concretely, we present the following contributions:

- a two-stage proposal mechanism that adjusts to local structure in target densities,
- a novel transition kernel that alleviates vanishing acceptance in the original delayed rejection method, and
- a benchmark of the versatility of two-stage kernels on several MCMC applications in rendering, improving convergence at equal time and with minimal implementation effort.

## 2 RELATED WORK

We review the most relevant MCMC light transport works. Šik and Krivánek [2018] provide a comprehensive survey. Despite its name, our approach is neither related to two-stage MCMC resampling [Gelfand and Sahu 1994] nor to two-stage MLT [Veach 1997].

*State Spaces.* The MLT formulation [Veach and Guibas 1997] mutates paths by directly modifying their geometry (e.g., removing a path vertex). This approach was simplified by Kelemen et al. [2002] with primary sample space MLT (PSSMLT), a reparameterization using an abstract sampling domain over the space of random numbers consumed during path generation. With this simplification, paths are represented as vectors of random variates embedded in a unit hypercube referred to as the *primary sample space* (PSS).

When combined with bidirectional path tracing [Lafortune and Willems 1993; Veach and Guibas 1995], PSSMLT generates a family of paths of which many may contribute little to no energy. To address this limitation, multiplexed MLT (MMLT) [Hachisuka et al. 2014] augments PSS with an extra dimension indexing across bidirectional sampling techniques so that chains can also explore regions where certain path sampling methods outperform others.

Concurrent works aiming to unify the two state spaces, path space and primary sample space, apply inverse mappings from the paths to the random numbers that produced them. Otsu et al. [2017] and Pantaleoni [2017] convert path space mutation strategies to their PSS representations in a manner that is agnostic to the underlying MCMC framework. Reversible jump MLT (RJMLT) [Bitterli et al. 2018] instead applies reversible jump methods from statistics [Green 1995] to MMLT to map strategies without altering the sample path’s geometry. Our method similarly draws from the MCMC literature, adapting and extending delayed rejection to MCMC light transport.

*Parallel Tempering.* Otsu et al. [2013] apply parallel tempering—or *replica exchange*—to MLT to reduce the likelihood of cycling within high energy peaks in the target distribution. At its core, replica exchange tempers many target distributions and mutation strategies to encourage jumps between modes. This requires tracking parallel chains, carefully tuning temperatures parameters and/or judiciously choosing the subspaces over which to spread chains. Kaplanyan and Dachsbacher [2013] modify the state space by regularizing degenerate densities and delta emission profiles, increasing the likeliness of sampling such narrow interactions. Šik and Křivánek [2016] extend this idea to improve global exploration while applying parallel tempering to vertex connection and merging. Our method does not rely on parallel tempering, and so is orthogonal to these methods; however, the generality of our algorithm would allow its independent application to any of the tempered chains. It is worth noting that Hachisuka and Jensen [2011]’s approach, albeit formulated within the replica exchange framework, can be recasted as a special case of delayed rejection where the first transition kernel is uniform and the target distribution is the binary visibility.

*Specialized Perturbations.* Another way of tackling suboptimal exploration is by means of specialized mutations. Li et al. [2015] propose a mutation strategy inspired by Hamiltonian Monte Carlo [Duane et al. 1987], enabling long traversals in state space and mitigating the chances of getting caught in high energy peaks. Their H2MC approach leverages automatic differentiation to perform anisotropic Gaussian mutations, but the added cost of computing gradients and Hessians of the log-target can be significant. Jakob and Marschner [2012] model the local differential geometry of path space to design mutations that can explore low-dimensional manifolds of (near-)specular chains, with Kaplanyan et al. [2014] and Hanika et al. [2015] later extending this approach to a more natural half-vector reparameterization. Segovia et al. [2007] first suggested multiple-try MCMC [Liu et al. 2000] to sample paths from several candidates and, more recently, Otsu et al. [2018] employ mutations that rely on visibility-aware estimates of optimal cone angles at path vertices.

While capable of modeling local state space landscapes, specialized transitions all rely on additional technical machinery that tends to increase both the algorithmic complexity of their resulting rendering algorithms and the computational cost per perturbation. Moreover, these mutations are often applicable only to very specific forms of light transport effects. In contrast, our delayed rejection method is inexpensive and is capable of effectively performing local exploration only in challenging regions in state space. Most importantly, our method is agnostic to the underlying MCMC perturbation schemes and can be implemented atop any existing Metropolis–Hastings based algorithms with minimal code modifications.

*Adaptive Methods.* Adaptive MCMC [Haario et al. 1998] was first introduced to the rendering community by Hachisuka and Jensen [2011], where a global mutation size is adapted to light transport complexity during Metropolis-based photon tracing. Zsolnai and Szirmay-Kalos [2013] automatically adjust large step probabilities in PSSMLT by gathering acceptance statistics early on during the rendering process. Lai et al. [2007] apply population MC to improve sampling based on information collected during initial iterations of an energy-redistribution path tracer [Cline et al. 2005].

These methods are limited to either adjusting algorithm parameters or necessitating ad-hoc heuristics to maintain ergodicity as the chains no longer satisfy the Markov property [Roberts and Rosenthal 2009]. Here, it is often unclear whether the adjusted parameters will consistently perform well, or rather help only on a select set of specialized cases. Our two-stage proposal mechanism is similar in spirit, albeit not strictly adaptive, as it locally changes its behavior on an as-needed basis. Moreover, it maintains detailed balance by construction and does not require any parameter annealing over time. In this regard, our approach is not related to adaptive MCMC but is compatible with such schemes.

### 3 PRELIMINARIES

*Path Integral Formulation.* We briefly review the path integral formulation of light transport. In its most general form, a light transport algorithm seeks solutions to the *measurement equation* [Veach and Guibas 1997], an integral of the form

$$I_j = \int_{\mathcal{P}} h_j(\bar{x}) f(\bar{x}) d\mu(\bar{x}), \quad (1)$$

where  $I_j$  is a measurement (typically of the  $j$ -th pixel) expressed as the integral over all possible light paths  $\bar{x}$  of the product of a (pixel) reconstruction filter  $h_j$  and the path throughput  $f$ , with respect to the area-product (Lebesgue) measure  $\mu$ . Paths are defined as a sequence  $\bar{x} = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k)$  with endpoints  $\mathbf{x}_0$  and  $\mathbf{x}_k$  and intermediate scattering vertices  $\mathbf{x}_1, \dots, \mathbf{x}_{k-1}$ . The path throughput is a product of reflection operators across all vertices  $\mathbf{x}_i$ . The space of all possible paths, called *path space*  $\mathcal{P}$ , is the union over light paths of every fixed length ( $k \geq 2$  vertices):  $\mathcal{P} \triangleq \bigcup_{k=2}^{\infty} \mathcal{P}^k$ .

*Monte Carlo Integration.* Monte Carlo integration was first applied to light transport to estimate pixel intensities  $I_j$  by sampling and averaging  $N$  independent light paths  $\bar{x}_i \in \mathcal{P}$  drawn from a probability density function  $p$  [Kajiya 1986]:

$$I_j \approx \langle I_j \rangle^N = \frac{1}{N} \sum_{i=1}^N \frac{h_j(\bar{x}_i) f(\bar{x}_i)}{p(\bar{x}_i)} = \frac{1}{N} \sum_{i=1}^N h_j(\bar{x}_i) C(\bar{x}_i), \quad (2)$$

where  $C(\bar{x}_i) \triangleq f(\bar{x}_i)/p(\bar{x}_i)$  is the path contribution function. The choice of path sampling density implicitly controls the estimator’s variance, where having  $p \propto f$  increases sample efficiency and so reduces variance. Formulating effective density functions remains an open problem as only some factors of  $f$  are known prior to sampling [Vorba et al. 2019].

*Metropolis–Hastings Algorithm.* MCMC methods apply a different approach to the problem, drawing a series of correlated samples (i.e., paths) that converge to an arbitrary (not necessarily normalized) target distribution  $\pi$ . The Metropolis–Hastings (MH) algorithm [Hastings 1970; Metropolis et al. 1953] provides one such process: given a current state  $x$  in state space  $\mathcal{X} \subseteq \mathbb{R}^d$ , we draw a proposal state  $y$  from a pre-specified conditional density  $Q(y|x)$ —the *Markov transition kernel*—and accept it with probability

$$\alpha(x, y) = 1 \wedge \frac{\pi(y) Q(x|y)}{\pi(x) Q(y|x)}. \quad (3)$$

Otherwise, it is rejected and the state  $x$  is repeated. Here,  $1 \wedge \eta \triangleq \min(1, \eta)$  corresponds to the MH ratio.



Equation (3) ensures that the Markov chain is reversible with respect to its unique invariant density  $\pi$  and thus satisfies detailed balance. Under mild regularity conditions [Robert and Casella 2005], the chain will reach its stationary regime where the sequence of samples obtained from this accept/reject criterion will be distributed according to the *normalized* density  $\tilde{\pi} = \pi/b$ , with  $b = \int_{\mathcal{X}} \pi(x) dx$ . MH requires only that we evaluate  $\tilde{\pi}$  up to a proportionality constant, making it attractive for sampling complex, high-dimensional distributions. Indeed, MH can be used to estimate intractable integrals such as the path integral Equation (1) without any specialized knowledge of the target density.

*Metropolis Light Transport.* MLT [Veach and Guibas 1997] applies MH to the measurement equation to estimate pixel intensities, with the target  $\pi$  set to the scalar luminance  $\ell$  of the path throughput. Here, the Markov chain operates on the location of path vertices in  $\bar{x} \in \mathcal{P} = \mathcal{X}$ . Using specialized mutations, MLT mutates chains to locally explore regions where  $f(\bar{x})$  (and so  $\ell(\bar{x})$ ) is large, avoiding paths that contribute little energy to the image. While the normalization constant  $b$  is not required when running the MCMC sampler, it is needed to accumulate and finalize the contributions onto the image plane. This factor is commonly estimated using an independent path sampling-based MC estimator, which can also be used to seed initial Markov states. Unlike standard MC methods, MCMC *simultaneously* estimates *all* pixel intensities, as any chain can explore many locations over the image plane.

We build upon PSSMLT [Kelemen et al. 2002], operating directly on the space of uniform random numbers used to generate sampled paths. *Primary sample space* can similarly be factored according to path length as  $\mathcal{U} \triangleq \bigcup_{k=2}^{\infty} \mathcal{U}^k$ , where  $\mathcal{U}^k = [0, 1]^{O(k)}$  is the  $O(k)$ -dimensional unit hypercube. In this new domain, the target function is the luminance of

$$\hat{C}(\bar{\mathbf{u}}) = (f \circ S)(\bar{\mathbf{u}}) / (p \circ S)(\bar{\mathbf{u}}), \quad (4)$$

where  $S : \mathcal{U} \rightarrow \mathcal{P}$  is a path sampling strategy that maps primary samples  $\bar{\mathbf{u}} \in \mathcal{U} = \mathcal{X}$  to light paths, that is,  $S(\bar{\mathbf{u}}) = \bar{x}$ . This reparameterization abstracts the geometry of paths with a more compact representation. The new importance function  $\hat{C}(\bar{\mathbf{u}})$ , however, becomes flatter which facilitates sampling in high density regions when the path sampling strategy is well-fitted to the integrand. PSS also naturally allows for symmetric proposal densities, eliminating the need to compute transition densities  $Q$  in Equation (3).

*Problem Statement.* Most MCMC light transport algorithms inherit the same problems as classical MCMC: it is difficult to strike a balance between local exploitation and computational cost for a fixed number of iterations. The path contribution—which is typically part of the target function—is multi-modal and discontinuous, which makes reaching this compromise difficult in practice, especially in radiometrically-complex scenes. As regions of interest in state space can vary widely in scale, a proposal with constant footprint size can be suboptimal. In cases where proposals are not well suited to the target density’s profile, Markov chains can undergo several consecutive rejections and their mixing times suffer. In rendering, this erratic convergence manifests itself visually as fireflies, high-luminance “splotches” and structured artifacts. This unpredictable

behavior inherent to sample correlation arguably represents the most important practical limitations of MCMC samplers.

## 4 TWO-STAGE DELAYED REJECTION

We propose an extension of MLT-based algorithms to address the mutation strategy-selection problem. Our variant of delayed rejection (DR) [Green and Mira 2001; Tierney and Mira 1999], tailored to the form of integrands we face in light transport, allows us to *conservatively* apply a sequence of strategies, i.e., only when necessary. At its core, DR modifies MH to admit different types of transitions at different stages. Suppose a first candidate mutation is generated and then rejected by MH. Rejection suggests that the proposal is not well suited and so should be altered. Instead of retaining the current state and proceeding to the next iteration, a *different* transition kernel can be used to propose a new state within the *same* iteration. This subtlety here is crucial as the application of state-dependent kernels would invalidate reversibility [Andrieu and Thoms 2008]. Under the DR formalism, the newly-introduced kernel is allowed to depend on the previously rejected sample, and acceptance probabilities are computed in a manner that preserves reversibility.

This approach allows for the design of more versatile kernels that can be combined to propose transitions according to the underlying complexity of each mutation scenario, leading to broader exploration of regions in state space at scales that fit them best. Our two-stage framework readily applies to several MCMC applications (Section 5) which we efficiently combine atop many existing PSS algorithms. In general, DR-variants of these methods improve local exploration for difficult light paths without sacrificing performance in regions where simpler strategies suffice.

### 4.1 Delayed Rejection Metropolis–Hastings

We formally introduce the mathematical framework of delayed rejection [Tierney and Mira 1999]. Suppose we wish to draw samples from a target distribution  $\pi$  with density  $\pi(x)$  over an arbitrary state space  $\mathcal{X} \subseteq \mathbb{R}^d$ . A two-stage delayed rejection algorithm starts with a standard MH step—that is, proposing a first state  $y$  from  $x$  with density  $Q_1(y|x)$  and accepting it with probability

$$\alpha_1(x, y) = 1 \wedge \frac{\pi(y) Q_1(x|y)}{\pi(x) Q_1(y|x)}, \quad (5)$$

where the subscript on acceptance probability indicates the stage. If this first proposal is rejected, DR generates a new candidate  $z$  sampled from a different density  $Q_2(z|y, x)$ , accepting it with probability

$$\alpha_2(x, z) = 1 \wedge \frac{\pi(z) Q_1(y|z) Q_2(x|y, z) [1 - \alpha_1(z, y)]}{\pi(x) Q_1(y|x) Q_2(z|y, x) [1 - \alpha_1(x, y)]}. \quad (6)$$

Here, we read conditionals *right-to-left*:  $Q_2(z|y, x)$  is the density of sampling  $z$  conditioned on first proposing  $y$  from the current state  $x$ . The acceptance probability in Equation (6) greedily imposes detailed balance at both stages: the probability of proposing and rejecting a first candidate  $y$  is  $Q_1(y|x) [1 - \alpha_1(x, y)]$ , and so the probability of moving to the second stage is  $Q_2(z|y, x) \times Q_1(y|x) [1 - \alpha_1(x, y)]$ . This formulation is sufficient for preserving the chain’s reversibility.



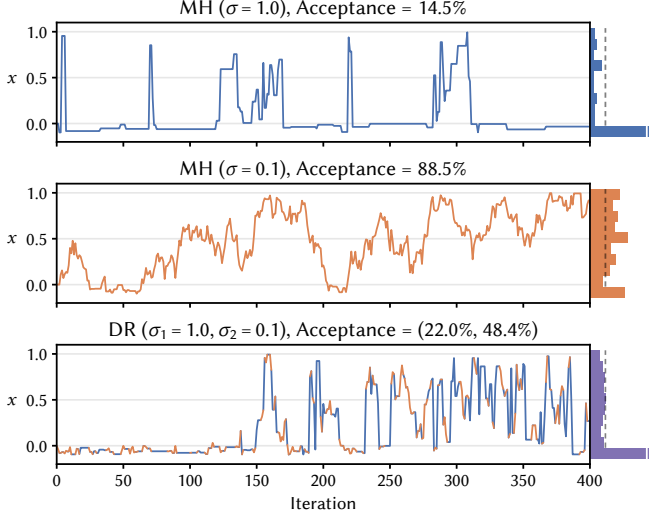


Fig. 3. 1D trace plots of the state trajectory for MH and DR. Some bars exceed the display range; the vertical dashed line represents the target  $\pi|_{[0,1]}$ . **Top:** A large proposal leads to long plateaus of repeated states near zero and low acceptance. **Middle:** When the proposal is too small, the chain moves slowly from one region to another, resulting in high correlation between successive states as shown by mountain-like patterns. This results in under-exploration of the high probability region  $[-1/10, 0)$ . **Bottom:** DR provides a more expressive proposal mechanism that captures both parts of the target density. Colors correspond to which stage was accepted.

When  $Q_2$  is symmetric with respect to  $x$  and  $z$  (i.e.  $Q_2(z|y, x) \equiv Q_2(x|y, z)$ ), Equation (6) simplifies to

$$\alpha_2(x, z) = 1 \wedge \frac{\pi(z) Q_1(y|z) [1 - \alpha_1(z, y)]}{\pi(x) Q_1(y|x) [1 - \alpha_1(x, y)]}. \quad (7)$$

Note that symmetry in  $Q_1$  is not sufficient to eliminate it from the expression since, generally speaking,  $Q_1(y|z) \neq Q_1(y|x)$ .

## 4.2 Illustrative 1D Example

We demonstrate the behavior and benefits of DR on a simple example by Green and Mira [2001]. Consider a one-dimensional target distribution  $\pi$  with a density prescribed by a mixture of two uniform distributions:

$$\pi(x) = \frac{1}{2\alpha} \mathbf{1}_{[-\alpha, 0)}(x) + \frac{1}{2\beta} \mathbf{1}_{[0, \beta)}(x), \quad (8)$$

where  $\mathbf{1}_{[a, b]}$  is the indicator function on the interval  $[a, b]$ . Suppose  $\alpha \ll \beta$ , then when  $x \in [-\alpha, 0)$  the transition kernel should prefer smaller transitions, otherwise proposals are likely to be rejected as they often fall outside the support of  $\pi$ ; however, when  $x \in [0, \beta)$  larger transition kernels are preferable, otherwise the chain will mix slowly. In this case, the acceptance rate will be high but global exploration will be poor.

One option to address this scenario is to hand-tune the effective spread region of the transition kernel to balance local exploitation and global exploration. Unfortunately, this tuning becomes difficult when the target density is complex and irregular, which is almost always the case in light transport simulations. Alternatively, we can seek such a compromise directly through delayed rejection.

Figure 3 illustrates trace plots for a Markov chain evolving with different transition mechanisms, both with the standard Metropolis–Hastings accept/reject step and the delayed rejection technique. Each plot graphs the current state of each chain over time. We also plot the resulting histogram of the estimated distribution for each approach. Intuitively,  $\pi(x)$  exhibits a large spike for  $[-\alpha, 0)$  followed by a low plateau for  $[0, \beta)$ . This example uses  $\alpha = 1/10$ ,  $\beta = 1$  and a Gaussian proposal with variance  $\sigma^2$ . As expected, the chain gets stuck for larger  $\sigma$  and exhibits high autocorrelation for smaller  $\sigma$ . Delayed rejection mitigates the drawbacks of both proposals and yields better exploration, as evidenced in its resulting histograms whose bins should match the dashed line over  $[0, \beta)$ .

*Discussion.* In this 1D example, the optimal spread of the kernel lies somewhere between  $\alpha$  and  $\beta$  and hand-tuning  $\sigma$  would therefore not be too tedious: indeed, adaptive MCMC [Haario et al. 1998] could be employed to automatically tune  $\sigma$ . This scale, however, is global and might not be optimal for all regions. Relying on DR allows us to use multiple transition kernels and automatically revert to the one that is most suitable locally, improving chain mixing over (much) more of the domain.

## 4.3 Limitations of Original Framework

In general, Tierney and Mira’s original formulation in Section 4.1 can be quite restrictive. The simplicity of Equation (6) hides one important issue when computing the reverse probabilities: it implies that the backward path from  $z$  to  $x$  has to follow the forward path from  $x$  to  $z$  with time reversed. In other words, it has to pass through the intermediate state  $y$  and get rejected. This can be problematic when  $y$  is far from  $z$ , in which case the term  $Q_1(y|z) \rightarrow 0$  as does  $\alpha_2(x, z)$ . This holds even when  $Q_1$  is symmetric since it is not simplified by cancellation from a term in the denominator. A similar situation occurs when  $Q_1$  does not have support near its center and  $y$  is too close to  $z$  (Figure 4). This vanishing acceptance behavior is more pronounced when the Markov transition kernels used at both stages have limited overlapping densities. This configuration increases the likelihood of sampling a second point in the tails of the first proposal density, which may occur frequently when dealing with narrow, i.e., exponential distributions. The same issue was indirectly observed by Green and Mira [2001] and later rediscovered by Trias et al. [2009]. To the best of our knowledge, we are the first to confront this issue directly within the DR formalism.

To see why doing so is necessary for light transport simulation, consider the exponential transition kernel recommended by Kelemen et al. [2002]. This so-called Kelemen-style mutation strategy




Fig. 4. Failure of Tierney and Mira [1999]. Given a first stage kernel  $Q_1$  without global support over  $\mathcal{X}$  (e.g., Kelemen-style mutation), it is possible to sample  $z$  in a zero-density region at the second stage, zeroing out  $\alpha_2(x, z)$ . We refer to this problem as the vanishing acceptance at the second stage.




Fig. 5. *Our pairwise orbital mutation.* When a first state  $y$  is rejected, a second state  $z$  is greedily sampled from a Wrapped Cauchy distribution  $Q_2$  such that it lies on the orbit  $\mathcal{O}(x, y)$  for each pair of primary samples. This removes the need to account for transition ratios in the acceptance ratio.

proposes new states  $y = (y_1, \dots, y_d)$  from  $x$  as

$$y_i = x_i + \text{sgn}\left(\xi_1 - \frac{1}{2}\right) \epsilon_{\max} \exp\left(-\log \frac{\epsilon_{\max}}{\epsilon_{\min}} \xi_2\right), \quad (9)$$

where  $\xi_{1,2} \sim U[0, 1]$  are uniform variates and  $0 < \epsilon_{\min} < \epsilon_{\max}$  are parameters controlling the kernel size with  $|y_i - x_i| \in [\epsilon_{\min}, \epsilon_{\max}]$ . Kelemen et al. originally suggested setting  $(\epsilon_{\min}, \epsilon_{\max}) = (1/1024, 1/64)$ . This kernel has zero support around zero, forcing minimal-distance moves from the current state, which encourages image plane stratification [Sziromay-Kalos and Szécsi 2017] by driving newly proposed paths towards different pixel locations. The resulting density is symmetric and is the product of  $k$  independent PDFs at each primary sample, allowing for lazy path construction. One consequence of this zero-probability “hole” is that we cannot directly apply delayed rejection when first stage samples are drawn from a Kelemen-style kernel, as the mutation does not have global support. We propose two approaches to address this *vanishing acceptance* (Figure 4), each of which motivates a different application to light transport simulations (in Section 5).

#### 4.4 Pairwise Orbital Mutations

We first propose a novel perturbation technique designed to remove the need for computing the ratio  $\Gamma_1 \triangleq Q_1(y|z) / Q_1(y|x)$  in Equation (7). Suppose that  $Q_1$  is a product of circularly symmetric (i.e.  $Q_1(y|z) \equiv Q_1(\|y - z\|)$ ), independent and lower-dimensional densities that partition the state space. Our key insight is that sampling  $z \sim Q_2(z|y)$  such that  $\|z - y\| = \|y - x\|$  for each element of the partition is sufficient for the ratio  $\Gamma_1$  to cancel out. Since surface light transport with PSS typically uses pairwise vertex relationships, we group coordinates pairwise and sequentially, e.g.,  $(y_1, y_2)$ ,  $(y_3, y_4)$ , and so on. This corresponds to perturbing directional samples at each vertex. Visually, such a second stage proposal amounts to moving along a circle that is centered at  $y$  and that passes through  $x$  (Figure 5).

Any  $z$  located on the orbit  $\mathcal{O}(x, y) \triangleq \{o : \|o - x\| = \|y - x\|\}$  will satisfy our needs. Ideally, we want  $z$  to be close to  $x$  on the orbit  $\mathcal{O}$  to increase the likelihood of acceptance. Denote the direction between the current and the proposed state as  $w = x - y$ , the radius of the orbit as  $r = \|w\|$  and the angle between this direction and the first coordinate axis as  $\mu$ . In order to sample a point close to  $x$  on  $\mathcal{O}$ , we wish to sample a small azimuthal  $\theta$  such that the angle between  $z - y$  and the first coordinate axis is  $\mu' = \mu + \theta$ .

We construct circular distributions [Fisher 1995], over a domain with interval  $[0, 2\pi)$ , by warping 1D PDFs  $p$  periodically over the

unit circle. We choose the wrapped Cauchy distribution, with PDF

$$Q_{\text{WCauchy}}(\theta; \rho) = \frac{1}{2\pi} \frac{1 - \rho^2}{1 + \rho^2 - 2\rho \cos \theta}, \quad (10)$$

where the scale  $\rho \in [0, 1]$  plays a role analogous to the standard deviation  $\sigma$  in a 1D normal distribution. This specific choice is motivated by the distribution admitting a closed-form expression for its cumulative distribution function—a requirement for inversion sampling. The distribution approaches a uniform circular density as  $\rho \rightarrow 0$ , whereas when  $\rho \rightarrow 1$  the distribution degenerates into a Dirac comb (namely, a  $2\pi$ -periodic tempered Dirac). We finally recover the second stage sample components by projecting back onto the axis  $z = y + r(\cos \mu', \sin \mu') \in \mathcal{O}$ . Note that this is similar in spirit to elliptical slice sampling [Murray et al. 2010], except we do not have to iterate the sampler to obtain a new point on the orbit.

By construction, the resulting kernel satisfies the  $Q_2(z|y, x) \equiv Q_2(x|y, z)$  condition and the new state  $z$ , lying on  $\mathcal{O}$ , implies that  $Q_1(y|z) \equiv Q_1(y|x)$ . This results in the simplified acceptance ratio

$$\alpha_2(x, z) = 1 \wedge \frac{\pi(z)[1 - \alpha_1(z, y)]}{\pi(x)[1 - \alpha_1(x, y)]} = 1 \wedge \frac{0 \vee [\pi(z) - \pi(y)]}{\pi(x) - \pi(y)}, \quad (11)$$

where  $0 \vee \eta \triangleq \max(0, \eta)$ . This acceptance probability can be evaluated efficiently in a Metropolis sampler. We use orbital mutations and this ratio in our experiments when the transition kernel of the baseline MCMC sampler factors into a product of independent components.

It is worth noting that this particular combination of pairwise mutation and the wrapped Cauchy distribution is not the only approach one could use to obtain a simplified acceptance ratio. When Russian roulette is used in PSS or in the presence of a participating media, correlating triplets of samples (e.g., to account for free-flight distance sampling) would be better suited than pairs. This could be achieved, for instance, by sampling directions on a sphere according to a von Mises–Fisher distribution [Han et al. 2007].

*Limitations.* Generalizing our construction to higher dimensions with hyper-spherical distributions is feasible but more involved. Such a generalization would also require correlated multi-dimensional samples, and so would induce significant changes to underlying path sampling routines. We chose the pairwise approach for its simplicity and ease of integration atop traditional PSSMLT and lazy path evaluation, leaving these alternatives to future work.

While our orbital mutation can solve the vanishing acceptance problem, it reduces the second proposal’s choice to the choice of circular PDF. In fact, restricting the second states to lie on an orbit prevents our current method from using correlated proposals, such as the anisotropic Gaussian mutations proposed by Li et al. [2015]. Using such mutation strategies without the orbital constraints would reintroduce the vanishing acceptance problem as was discussed in Section 4.3. To remedy to this limitation, we propose an extension of our approach based on the work of Green and Mira [2001].

#### 4.5 Generalized Delayed Rejection

With Reversible Jump MCMC, Green and Mira [2001] relax the  $x \leftrightarrow y \leftrightarrow z$  reversibility constraint by employing an additional fictional state  $y^*$  that effectively generalizes the original DR algorithm to transdimensional moves. Below, we restrict ourselves to fixed




Fig. 6. Different states involved in Green’s two stage DR algorithm using a fictional state  $y^* = z - \delta y$ . Even though  $Q_1(y | z) \neq Q_1(y | x)$  in general, we do have that  $Q_1(y^* | z) = Q_1(y | x)$  for a symmetric  $Q_1$ .

dimensional mappings and present only a case of specific interest. For the full theoretical exposition we refer to [Green and Mira \[2001\]](#).

Similar to the original formulation, the first move consists of generating a sample  $y = x + \delta y$  in a standard MH step. If the move is rejected, we generate a new state  $z = x + \delta z$  sampled according to a different density  $Q_2(z | y, x)$  and then consider an intermediate, fictional state  $y^* = z - (y - x) = z - \delta y$  instead of  $y$  (Figure 6). The new state  $z$  is then accepted with probability

$$\alpha_2^*(x, z) = 1 \wedge \frac{\pi(z) Q_1(y^* | z) Q_2(x | y^*, z) [1 - \alpha_1(z, y^*)]}{\pi(x) Q_1(y | x) Q_2(z | y, x) [1 - \alpha_1(x, y)]}, \quad (12)$$

where

$$\alpha_1(z, y^*) = 1 \wedge \frac{\pi(y^*) Q_1(z | y^*)}{\pi(z) Q_1(y^* | z)}. \quad (13)$$

Note that the presence of  $y^*$  in Equation (12) induces symmetry. Contrary to the initial formulation, we observe that its contribution in every stage disappears when  $Q_1$  is symmetric:

$$\alpha_2^*(x, z) = 1 \wedge \frac{\pi(z) Q_2(x | y^*, z) [1 - \alpha_1(z, y^*)]}{\pi(x) Q_2(z | y, x) [1 - \alpha_1(x, y)]}, \quad (14)$$

solving the acceptance problem we discussed earlier. The special case  $y^* = y$  is precisely Equation (6) from [Tierney and Mira \[1999\]](#). [Green and Mira \[2001\]](#) provide the full derivation of Equation (14).

While [Green and Mira](#)’s approach generally gives comparable acceptance to our orbital mutation—and thus higher than [Tierney and Mira](#)’s—it involves an extra evaluation of the target density at  $y^*$ . In the context of light transport simulation, this requires tracing an additional light path. As shown in Figure 7, this overhead often results in lower performance than our orbital mutation when used in simpler applications. In such cases, the orbital approach should be preferred. However, when one wishes to use mutation strategies that are not compatible with orbital mutations (such as H2MC [[Li et al. 2015](#)]), [Green and Mira](#)’s generalization should be used.




Fig. 7. Equal time comparison of [Tierney and Mira \[1999\]](#), [Green and Mira \[2001\]](#) and our pairwise orbital approach on the GLASS OF WATER scene.




Fig. 8. Decision tree of our DR algorithm.

#### 4.6 Waste-recycling for Proposed States

By design, DR generates more states than a standard Metropolis–Hastings step. While accumulating only accepted states is viable, we want a scheme that involves *all* states to avoid dismissing important information. To this end, we adapt the *use of the expected values* [[Veach and Guibas 1997](#)] to our two-stage proposal mechanism. Consider the following observations:

- moving from state  $x$  to the first proposal  $y$  has probability  $\omega_y = \alpha_1(x, y)$ ,
- getting rejected at the first stage has probability  $1 - \alpha_1(x, y)$ ,
- once at  $y$ , moving to a second proposal  $z$  has probability  $\alpha_2(y, z)$ ,
- the probability of the full sequence  $x \rightarrow y \rightarrow z$  occurring and getting accepted at  $z$  is  $\omega_z = (1 - \alpha_1)\alpha_2$ , and
- similarly, the probability of both proposed stages failing and thus staying at  $x$  is  $\omega_x = 1 - \omega_y - \omega_z = (1 - \alpha_1)(1 - \alpha_2)$ .

We accumulate sample contributions at all three locations, weighted by their corresponding probabilities  $\omega_{\square}$  (Figure 8). We observe that, when using Equation (12), since  $y^*$  is not a state that is directly accessible from  $x$ , we cannot splat its contribution using the method of expected values. Indeed, this technique amounts to accumulating  $E[\pi(X_i) | X_{i-1} = x]$ , but  $y^*$  was sampled from  $z$ , not from  $x$ .

#### 4.7 Discussion on Mixture Models

While mixture models can simulate different transition kernels, they remain agnostic to the current state of the chain and rely on careful mixture weight settings. In contrast, delaying the rejection of samples provides an opportunity to adjust proposals along the chain, automatically tempering the negative consequences of local mismatches between transition kernel and target density. Our DR approach applies to MCMC light transport in order to afford chains a “second chance” to reach valuable regions in path space.

## 5 APPLICATIONS AND RESULTS

We validate our theory by employing our two-stage kernel to augment various MCMC light transport algorithms with more flexible transitions. We demonstrate that this simple—yet effective—addition increases the robustness of several techniques operating in PSS. We group our applications into two categories:

- (i) **Bold-then-Timid (Section 5.1)**. The first stage attempts a more adventurous transition; if it fails, we propose a more conservative move. Larger perturbations can improve exploration but are more likely to be rejected; smaller perturbations provide a “safety net” to avoid repeating the same state.
- (ii) **Cheap-then-Expensive (Section 5.2)**. The first stage uses a simple, efficient mutation; if it fails, a more intricate kernel is employed. This amortizes the cost of specialized mutations, limiting their use to challenging regions in state space.



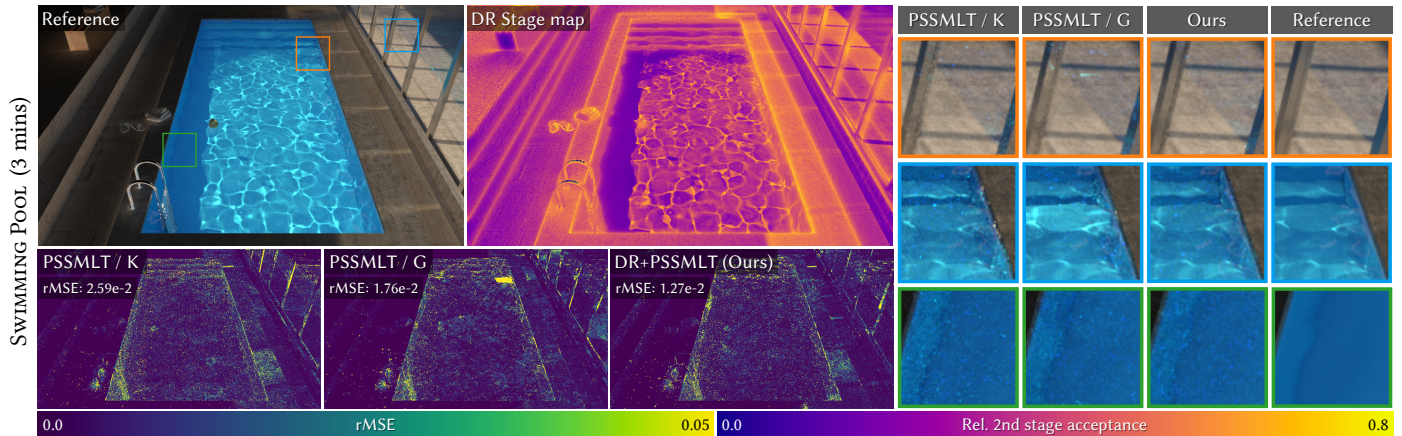


Fig. 9. Equal time comparison of the SWIMMING POOL scene for delayed rejection applied to PSSMLT. Compared to PSSMLT with the standard Kelemen-style kernel (PSSMLT / K) and with a Gaussian kernel (PSSMLT / G), DRMLT produces smoother results, thanks to our orbital mutation.

We show the pseudocode of our algorithm in Appendix A. We analyze the performance of our algorithm by tracking the relative Mean Squared Error (rMSE) compared to a reference image, reporting the median over several runs. Additional comparison metrics and an interactive viewer are provided in our supplemental material. All reference images are generated with several days of render time in their respective baseline renderers. All comparisons are equal time renderings running on an Intel Platinum 8160F Skylake CPU at 2.1 GHz with 48 threads. We only include a few results per application in our exposition and our supplemental material includes the comprehensive superset of results. Our implementation is publicly available online ([github.com/joeylitalien/drmlt](https://github.com/joeylitalien/drmlt)).

## 5.1 Bold-then-Timid

**5.1.1 Kelemen Then Pairwise Orbital.** The Kelemen kernel is the de facto mutation choice for most PSS-based methods. It can, however, be detrimental to the chain when the density is concentrated in thin regions of the space. In these scenarios, off-centered perturbations can lead to poor local exploration of thin highlights, which would be better served by a more localized transition kernel. As these small regions of the state space contribute significant energy to the final image, locally-bad fits introduce a disproportional amount of noise




Fig. 10. *Failure of Kelemen-style mutation.* This scene is explicitly designed to fail bold mutations as they cannot efficiently explore the thin highlight. Here, the Gaussian kernel (PSSMLT / G) clearly wins over Kelemen (PSSMLT / K). Even if our two-stage mechanism starts with a Kelemen mutation, it is able to recover and efficiently explore the reflection on the plate.

in these regions (Figure 10). Identifying occurrences that lead to these problems directly within a Metropolis sampler is thus difficult.

We use our pairwise orbital perturbation (Section 4.4) to reach a good compromise by mutating paths with a more *timid* transition kernel. In resorting to this more conservative proposal, chains can suggest paths that are closer to the current one—while maintaining a sufficiently high acceptance rate—without staying at the same state for many iterations. This extra stage has the desired effect of slightly moving the current path and provides an effective mechanism to better explore local modes.

**Results.** We integrated DRMLT atop a PSSMLT path tracer in Mitsuba v0.6 [Jakob 2013]. We use the kernel bounds recommended by Kelemen and our orbital perturbation with  $\rho = \exp(-1/4)$  for the second stage. We keep this value fixed across all the experiments that use our orbital mutation. We also show results with a Gaussian density with variance chosen to match the acceptance rate of our DR application. Doing so illustrates the versatility of our two-stage approach compared to a single-stage default kernel.

Figure 9 provides an equal time comparison on the SWIMMING POOL scene with complex transport from caustic and specular-diffuse-specular paths. Standard Kelemen-style mutations (PSSMLT / K) are too sensitive and have difficulties staying on the specular manifold, yielding noisy images. On the other hand, Gaussian mutations (PSSMLT / G) are less likely to fall outside a high density region, but they are too localized and generate luminosity spikes in the image. Our two-stage algorithm (DR+PSSMLT) work well overall, accommodating both types of mutations. Our method also produces a smoother output, thanks to the two-stage pixel splatting that makes use of all intermediate states. The DR stage map visualizes per-pixel relative acceptance at the second stage: as expected, a higher proportion of states are accepted at the second stage in complex regions, i.e., the caustics at the bottom of the pool and the reflected light around the pool.

**5.1.2 Multiplexed Then Subpath.** Our bold-then-timid approach can further be applied to multiplexed MLT (MMLT) [Hachisuka et al. 2014] to improve its efficiency. MMLT allows Markov chains to adaptively select the  $(s, t)$ -bidirectional sampling strategy, where

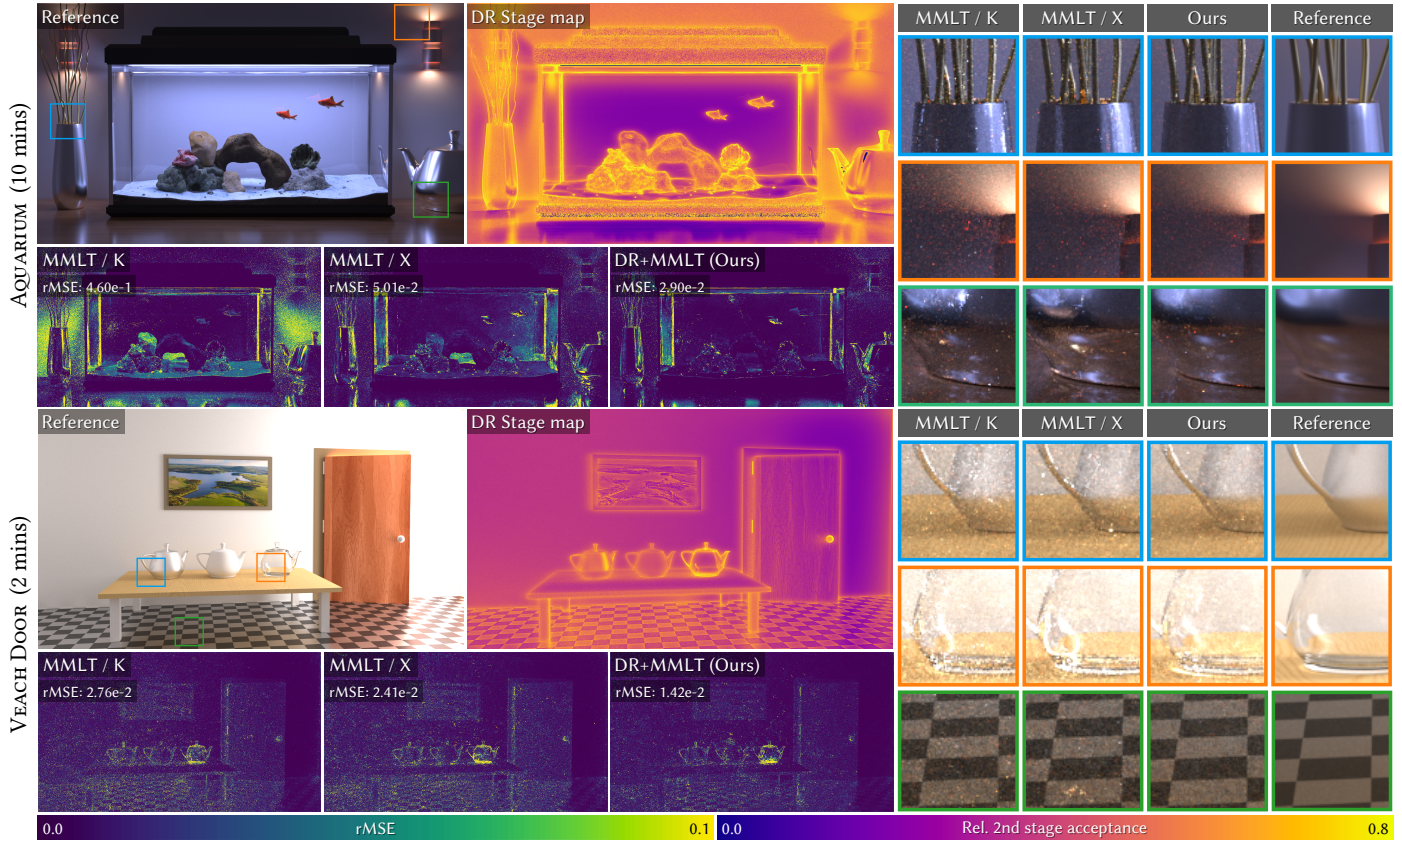


Fig. 11. Equal time comparisons of the scenes (AQUARIUM and VEACH DOOR) for delayed rejection applied to multiplexed MLT. By using a more flexible multiplexed kernel, our method significantly reduces variance compared to MMLT with a Kelemen-style kernel (MMLT / K) and a mixture model (MMLT / X).

$s, t$  are the length of the emitter and sensor subpaths. Here, only one path is computed at each iteration, avoiding the computation of all-pairs MIS-weighted vertex contributions.

We only slightly modify MMLT when performing local perturbations, *fixing* the bidirectional strategy instead of allowing it to vary along the chain. Since strategy changes are rarely proposed and accepted [Bitterli et al. 2018], this modification does not harm performance as most of MMLT’s gains come from evaluating a single path instead of an entire family of paths. We perturb all primary samples with our *bold* mutation; for our *timid* mutation, we *fix* the emitter subpath and only perturb the sensor subpath samples using our pairwise orbital mutation. This choice allows us to reuse parts of the computations from current and rejected states as only the sensor subpath needs to be re-traced.

**Results.** We integrated MMLT atop Mitsuba v0.6. Figure 11 compares our DR+MMLT approach to standard MMLT on scenes with complex visibility. We also measure our performance against MMLT using an equally weighted mixture of our two stages (MMLT / X). Our DR+MMLT generates smoother images, has lower rMSE (on both scenes) and performs better than its mixture counterpart. This further supports our claim that, unlike mixture variants, our method is able to more appropriately (and automatically) select transition kernels based on the target density. AQUARIUM exhibits many light transport phenomena, evidenced in its stage acceptance map. The

interior of the aquarium is lit by an emitter inside its casing, and so is fairly easy to explore compared to the rest of the scene. In contrast, the first stage for VEACH DOOR behaves well except for a few localized regions. The contrast in the map—e.g., around the golden fishes and the wooden door—corresponds to abrupt changes in luminosity which tend to be initially rejected by our sampler.

**Handling Light Tracing.** Extra care is needed when a full emitter path forms the initial state of the chain, i.e., when  $t = 1, s = k - 1$  ( $k$  is the path length). In this scenario, we mutate the entire path twice using our orbital mutation at the second stage to avoid proposing a state identical to the previous one. This case alone motivates our fixing of the bidirectional strategy across the chain: if a path gets mutated to a full emitter one, extra considerations would be needed to account for the sudden change of proposal distribution, without violating reversibility. Therefore, we opt for a more transparent approach that treats pure light tracing as a special case without modifying the mathematical formulation of the acceptance ratios.

## 5.2 Cheap-then-Expensive

Simpler mutations, such as the one proposed by Kelemen, are less expensive and can be applied more often than complex ones (in fixed time). When working with limited resources and a target density that is nontrivial to evaluate, choosing when to apply a more






Fig. 12. Equal time comparison of the ROUGH TORUS scene for delayed rejection applied to H2MC. Our two-stage method explores difficult regions with strong directional transport and alleviates fireflies, compared to one-stage PSSMLT with an isotropic Gaussian kernel (PSSMLT / G) and fully anisotropic H2MC.

expensive/complicated transition kernel requires domain expertise. We propose to use of DR to automate this choice.

The H2MC proposal [Li et al. 2015] uses automatic differentiation to compute first- and second-order derivatives for every possible pair of materials. This is costly (and so performed as a preprocess) but enables more expressive transitions. Under the H2MC framework, a  $d$ -dimensional state  $y$  is proposed from  $x$  as  $y = x + \gamma$ , where  $\gamma \sim \mathcal{N}(\mu_x, \Sigma_x)$  is sampled from a multivariate (correlated) normal density. Here,  $\Sigma_x$  is an  $d \times d$  covariance matrix that depends on the throughput Hessian and  $\mu_x$  is a mean variable; both depend on the *full* vector of random numbers  $x$ . Most importantly, this correlation prevents the matrix from factoring into a product of joint densities. As such, sampling from the H2MC kernel requires a costly eigendecomposition of the inverse Hessian.

To amortize this cost, Li et al. propose to branch to anisotropic mutations based on a thresholded value of the Hessian  $L^2$ -norm, conservatively falling back to simple isotropic Gaussian transitions. The main issue with this approach is that it cannot capture anisotropy *across scales*, which limits exploration potential of the Hessian-based transition kernel.

We leverage the generalized DR framework (Section 4.5) to determine this branching more efficiently. We always perform a cheaper, isotropic Gaussian transition first; if this fails, we resort to the more costly H2MC transition. Note that by doing so, we need only compute the extra Hessian for the second stage at  $z$ . To demonstrate DRMLT’s flexibility, we adopt a hybrid approach with a second stage transition set to an equally-weighted mixture of anisotropic and isotropic Gaussians, the latter with smaller variance. This amortizes the cost of computing  $\Sigma_x$  every time a first proposal is rejected.

**Results.** We demonstrate the robustness of our approach by implementing DR in the differential DPT renderer [Li et al. 2015]. To ensure a fair comparison between DR+H2MC and H2MC, we disable path space lens perturbations in DPT and use the authors’ state space reparameterization for both techniques. We apply the same (original) parameters and a maximum path length of  $k = 10$ , allowing derivatives to be precomputed in a reasonable amount of time.

We compare to PSSMLT with a Gaussian kernel matching our first stage’s proposal and to H2MC (Figures 1 and 12).

We chose to use PSSMLT with a Gaussian transition kernel as an example of a cheap mutation strategy; this doubles as a baseline and the first stage in our DRMLT method. CHES has many modes of transport, and so is challenging for an isotropic- or anisotropic-only proposal. ROUGH TORUS exhibits significant state space anisotropies well-suited to the Hessian-Hamiltonian dynamics-based proposals.

Our method outperforms H2MC in both test scenes, as evidenced by lower errors and less variance in the rendered images. On CHES, isotropic Gaussian proposals are more effective in regions where the target is smooth, but struggle with directional transport such as on rough glass. In contrast, H2MC fares better on brushed metals and dielectrics, i.e., the chess pieces, but worse overall due to its overhead. ROUGH TORUS is dominated by transport concentrated in narrow bands of the space, as shown by the heavily skewed DR stage map, and our algorithm effectively balances the two stages and removes bright spikes in the images. Our method automatically finds an equilibrium between the two scenarios. Since first-stage mutations are an order of magnitude faster to evaluate, we are able to achieve more mutations in equal time, leading to a consistent improvements across scenes.

### 5.3 Empirical Convergence Analysis

Due to correlation between MCMC samples (as opposed to independent samples), a single run of an MCMC integrator is not representative of its ability to explore the relevant space. Therefore, we run 10 instances of each algorithm initialized at different random states to better capture the average behavior. Since the normalization factor can vary from one initialization to another, we integrate the luminance of the reference images to compute global scaling factors and supply these values to every algorithm. We also compute the standard deviation of the metrics and plot continuous error bars to visualize the variation, to assess the stability of our method.

Figure 13 shows one result per application. We refer to our supplementary material for more results. Our DRMLT approach **consistently outperforms** its one-stage PSS counterparts, both in  $L^1$  and



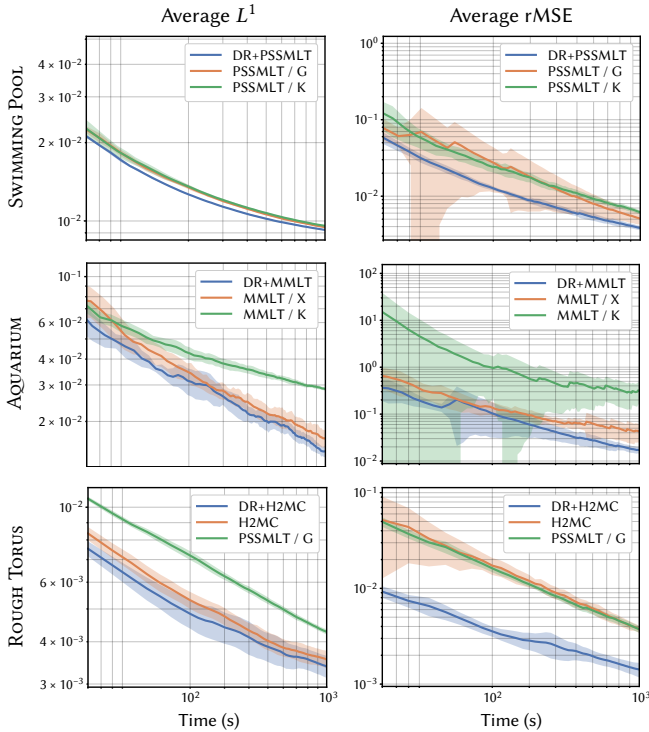


Fig. 13. *Convergence plots.* We visualize both the mean (thick curve) and standard deviation (shaded regions) of the error over 10 independent runs. Plots for all scenes and other metrics can be found in our additional material.

relative MSE. Our  $L^1$ -norm performance hints to a potential smoothing effect in low variance regions, due to our double-splating. In contrast, rMSE penalizes large outliers, such as fireflies. Our method is robust to both metrics, as shown by the lower curves.

## 6 LIMITATIONS AND FUTURE WORK

While DRMLT works well in many different scenarios, it has a few limitations in its current state.

*Portability to Path Space MLT.* There are several challenges in porting the original DR algorithm from Tierney and Mira [1999] to Veach and Guibas’ MLT. As path space perturbation strategies are mostly complementary and designed for paths with specific structures, choosing perturbation strategies for the two stages that are both *general* and *suitable* is rarely feasible. As such, a complex decision tree based on the possible path types should be established, further obfuscating the mutation routine. We can apply our orbital mutation directly to simple perturbations (e.g., lens and caustic) but more intricate mutation strategies (e.g., manifold exploration [Jakob and Marschner 2012]) need to be treated separately. This could potentially be accomplished through inverse mappings [Bitterli et al. 2018; Otsu et al. 2017; Pantaleoni 2017], transforming paths to their PSS representations. We leave this as future work.

Extra care is also needed when adapting Green and Mira’s algorithm to path space MLT. Indeed, this generalized version of DR requires tracing an intermediate light path  $y^* = z - \delta y$ . Naïvely adding or transforming vertices will almost always result in paths

that escape the scene manifold  $\mathcal{M}$ . One way of alleviating this issue is to project each vertex back onto the scene to obtain a new intermediate path  $\tilde{y} = \text{proj}_{\mathcal{M}}(y^*)$ . Accounting for this projection in Equations (12–13) to retain reversibility, however, may be difficult.

*Number of Stages.* Our approach only attempts two stages before conceding rejection and advancing time. While DRMLT can be generalized to  $m$  stages by enforcing detailed balance separately at each stage  $i \leq m$ , some problems remain. This construction gives rise to a product of rejection probabilities  $\prod_i Q_i [1 - \alpha_i]$  that can magnify the vanishing acceptance phenomenon in subsequent stages [Mira 2001]. Moreover, the number of reversibility constraints for these states grows geometrically with the number of stages, incurring a non-negligible additional cost. As noted by Green and Mira [2001], moving to three or more stages is typically not worthwhile, motivating our two-stage mechanism.

*Choice of Proposals.* One challenge when applying DRMLT lies in devising a sufficiently diverse sequence of proposals. This requires understanding the scale at which mutations operate and crafting a second stage mutation strategy that can capitalize on the shortcomings of the first. Finding innovative combinations that naturally identify these flaws is nontrivial. Also, in some cases, a grid search for “optimal” kernel spread and/or mixture weights could yield performance similar to our bold-then-timid approach. By design, DRMLT precisely avoids such tedious trial-and-error experiments and circumvents scene-dependent parameter optimization.

### 6.1 Future Work

Several interesting and promising directions result from our work. We highlight a few possibilities, below.

*Differential Geometric Mutations.* Given the growing ubiquity of differentiable rendering [Li et al. 2018; Nimier-David et al. 2019], our approach offers an attractive tool for amortizing the per-sample cost incurred by automatic differentiation. There are other MCMC samplers (e.g., Metropolis-adjusted Langevin [Livingstone and Girolami 2014] and its Riemannian manifold variant [Girolami et al. 2011]) that exploit the geometry of the target distribution to improve exploration. It is worth revisiting alternatives that were promptly dismissed for light transport simulation due to their seemingly high cost. Applying DR to H2MC was the first step in this direction, and investigating the performance of other differentiable mutations at the second stage is a promising avenue of future work.

*Improved Large Steps.* DR can be applied to the large step mutation, similar to Trias et al. [2009]. When a large step falls in the neighborhood of a target mode but is rejected by Metropolis–Hastings, DR could be applied to reach this mode with successively smaller steps to climb up the density hill. This effectively explores the new region of parameter space until a new state is finally accepted or a stopping condition is met. Identifying when an initial large step must be followed by additional small steps remains unclear, but doing so on a stochastic basis may foster global exploration.

*Delayed Rejection Adaptive Metropolis (DRAM).* Combining adaptive MCMC with DR, as suggested by Haario et al. [2006], is a natural direction to explore. This method relies on constructing proposals

by fitting the covariance of the distribution at different scales. One may (perhaps optimistically) expect to achieve similar performance as Li et al. [2015] here, however without the need for automatic differentiation. DRAM gives rise to a correlated transition kernel and is built atop the original DR algorithm, and so more investigation is required to address any potentially negative effects on the acceptance rates.

*Delayed Acceptance (DA)*. Similar to DR, delayed acceptance [Christen and Fox 2005; Sherlock et al. 2017] could be used to combine a cheap approximation of the target density with a two-stage version of MH. DA tests samples against the surrogate target, promptly dismissing those that would be unlikely to be accepted by the true target. A second stage then uses an acceptance probability to resolve the discrepancy between the approximation and the target, ensuring convergence. Since evaluations of the path throughput are generally the bottleneck of MLT-based methods, a successful application of this method could result in considerable speedups.

## 7 CONCLUSION

We introduced a two-stage mutation strategy based on the delayed rejection framework [Green and Mira 2001; Tierney and Mira 1999]. By sequentially combining two transition kernels, our method provides a form of “safety net” in order to improve local state space exploitation without comprising global exploration. We first showed that a naïve application of the original DR algorithm to light transport simulation could lead to poor acceptance rates. To address this problem, we developed a novel conditional mutation strategy at the second stage. Our resulting algorithm, *delayed rejection Metropolis light transport* (DRMLT), is simple to implement in existing rendering systems and can be applied to any PSS-based algorithm. We demonstrated various applications of DRMLT, namely *bold-then-timid* and *cheap-then-expensive* strategies, that automatically balance local exploration and computational efficiency across many complex light transport scenarios. We believe that DRMLT is a powerful approach to allow for specialized mutation strategies under tight computational budgets.

## ACKNOWLEDGMENTS

The scenes are courtesy of Bitterli [2016] (VEACH DOOR and GLASS OF WATER) and Ondřej Karlík (SWIMMING POOL). Computing resources were provided by the National Systems of Compute Canada. This research was partially funded by NSERC (RGPIN-2018-05669) and JSPS KAKENHI (18KK0309).

## REFERENCES

Christophe Andrieu and Johannes Thoms. 2008. A Tutorial on Adaptive MCMC. *Statistics and Computing* 18, 4 (2008). <https://doi.org/10.1007/s11222-008-9110-y>

Michael Ashikhmin, Simon Premoze, Peter Shirley, and Brian Smits. 2001. A Variance Analysis of the Metropolis Light Transport Algorithm. *Computers & Graphics* 25, 2 (2001), 287–294. [https://doi.org/10.1016/S0097-8493\(00\)00131-X](https://doi.org/10.1016/S0097-8493(00)00131-X)

Benedikt Bitterli. 2016. Rendering Resources. (2016). <https://benedikt-bitterli.me/resources>

Benedikt Bitterli, Wenzel Jakob, Jan Novák, and Wojciech Jarosz. 2018. Reversible Jump Metropolis Light Transport Using Inverse Mappings. *ACM Transactions on Graphics* 37, 1 (Jan. 2018), 1:1–1:12. <https://doi.org/10/gd52ph>

Benedikt Bitterli and Wojciech Jarosz. 2019. Selectively Metropolized Monte Carlo light transport simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Nov. 2019). <https://doi.org/10.1145/3355089.3356578>

J. Andrés Christen and Colin Fox. 2005. Markov Chain Monte Carlo Using an Approximation. *Journal of Computational and Graphical Statistics* 14, 4 (2005), 795–810. <https://doi.org/10.1198/106186005X76983>

David Cline, Justin Talbot, and Parris Egbert. 2005. Energy Redistribution Path Tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24, 3 (July 2005), 1186–1195. <https://doi.org/10/b3xtrn>

Simon Duane, A. D. Kennedy, Brian J. Pendleton, and Duncan Roweth. 1987. Hybrid Monte Carlo. *Physics Letters B* 195, 2 (1987), 216–222. [https://doi.org/10.1016/0370-2693\(87\)91197-X](https://doi.org/10.1016/0370-2693(87)91197-X)

Nicholas I. Fisher. 1995. *Statistical Analysis of Circular Data*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511564345>

Alan E. Gelfand and Sujit K. Sahu. 1994. On Markov Chain Monte Carlo Acceleration. *Journal of Computational and Graphical Statistics* 3, 3 (1994), 261–276. <https://doi.org/10.2307/1390911>

Mark Girolami, Ben Calderhead, and Siu A. Chin. 2011. Riemann Manifold Langevin and Hamiltonian Monte Carlo Methods. *J. of the Royal Statistical Society, Series B (Methodological)* (2011). <https://doi.org/10.1111/j.1467-9868.2010.00765.x>

Peter J. Green. 1995. Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika* 82, 4 (Dec. 1995), 711–732. <https://doi.org/10/bt2s2t>

Peter J. Green and Antonietta Mira. 2001. Delayed Rejection in Reversible Jump Metropolis–Hastings. *Biometrika* 88, 4 (2001), 1035–1053. <https://doi.org/10.1093/biomet/88.4.1035>

Heikki Haario, Marko Laine, Antonietta Mira, and Eero Saksman. 2006. DRAM: Efficient Adaptive MCMC. *Statistics and Computing* 16, 4 (01 Dec 2006), 339–354. <https://doi.org/10.1007/s11222-006-9438-0>

Heikki Haario, Eero Saksman, and Johanna Tamminen. 1998. An Adaptive Metropolis Algorithm. *Bernoulli* 7 (1998), 223–242. <https://doi.org/10.2307/3318737>

Toshiya Hachisuka and Henrik Wann Jensen. 2011. Robust Adaptive Photon Tracing Using Photon Path Visibility. *ACM Transactions on Graphics* 30, 5 (Oct. 2011), 114:1–114:11. <https://doi.org/10/fpwzq9>

Toshiya Hachisuka, Anton S. Kaplanyan, and Carsten Dachsbacher. 2014. Multiplexed Metropolis Light Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 100:1–100:10. <https://doi.org/10/f6cswv>

Charles Han, Bo Sun, Ravi Ramamoorthi, and Eitan Grinspun. 2007. Frequency Domain Normal Map Filtering. *ACM Transactions on Graphics* 26, 3 (2007), 28. <https://doi.org/10.1145/1276377.1276412>

Johannes Hanika, Anton Kaplanyan, and Carsten Dachsbacher. 2015. Improved Half Vector Space Light Transport. *Computer Graphics Forum* 34, 4 (July 2015), 65–74. <https://doi.org/10/gfzv83>

Wilfred K. Hastings. 1970. Monte Carlo Sampling Methods Using Markov Chains and Their Applications. *Biometrika* 57, 1 (April 1970), 97–109. <https://doi.org/10/dkbnmf>

Wenzel Jakob. 2013. Mitsuba Renderer. <http://www.mitsuba-renderer.org>

Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (July 2012), 58:1–58:13. <https://doi.org/10/gfzq4p>

James T. Kajiya. 1986. The Rendering Equation. *Computer Graphics (Proceedings of SIGGRAPH)* 20, 4 (Aug. 1986), 143–150. <https://doi.org/10/cvf53j>

Anton S. Kaplanyan and Carsten Dachsbacher. 2013. Path Space Regularization for Holistic and Robust Light Transport. *Computer Graphics Forum (Proceedings of Eurographics)* 32, 2 (2013), 63–72. <https://doi.org/10/gbc3p8>

Anton S. Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. 2014. The Natural-Constraint Representation of the Path Space for Efficient Light Transport Simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33, 4 (July 2014), 102:1–102:13. <https://doi.org/10/f6cz85>

Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A Simple and Robust Mutation Strategy for the Metropolis Light Transport Algorithm. *Computer Graphics Forum* 21, 3 (Sept. 2002), 531–540. <https://doi.org/10/bfrsqn>

Eric P. Lafortune and Yves D. Willems. 1993. Bi-Directional Path Tracing. In *Proceedings of the International Conference on Computational Graphics and Visualization Techniques (Compugraphics)*, Vol. 93. Alvor, Portugal, 145–153.

Yu-Chi Lai, Shao Hua Fan, Stephen Chenney, and Charle Dyer. 2007. Photorealistic Image Rendering with Population Monte Carlo Energy Redistribution. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (Grenoble, France) (EGSR'07)*. Eurographics Association, 287–295. <https://doi.org/10.2312/EGWR/EGSR07/287-295>

Tzu-Mao Li, Miika Aittala, Frédo Durand, and Jaakko Lehtinen. 2018. Differentiable Monte Carlo Ray Tracing through Edge Sampling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (2018). <https://doi.org/10.1145/3272127.3275109>

Tzu-Mao Li, Jaakko Lehtinen, Ravi Ramamoorthi, Wenzel Jakob, and Frédo Durand. 2015. Anisotropic Gaussian Mutations for Metropolis Light Transport Through Hessian-Hamiltonian Dynamics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 34, 6 (Oct. 2015), 209:1–209:13. <https://doi.org/10/f7wrcc>

- Jun S. Liu, Faming Liang, and Wing Hung Wong. 2000. The Multiple-Try Method and Local Optimization in Metropolis Sampling. *J. Amer. Statist. Assoc.* 95, 449 (March 2000), 121–134. <https://doi.org/10/gfkjaj>
- Samuel Livingstone and Mark Girolami. 2014. Information-geometric Markov chain Monte Carlo methods using diffusions. *Entropy* 16, 6 (2014), 3074–3102. <https://doi.org/10.3390/e16063074>
- Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. 1953. Equation of State Calculations by Fast Computing Machines. *Journal of Chemical Physics* 21, 6 (June 1953), 1087–1092. <https://doi.org/10/ds736f>
- Antonietta Mira. 2001. On Metropolis–Hastings Algorithms with Delayed Rejection. *Metron* 59, 3-4 (2001), 231–241.
- Iain Murray, Ryan Adams, and David MacKay. 2010. Elliptical slice sampling. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research)*, Vol. 9, 541–548.
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (2019). <https://doi.org/10.1145/3355089.3356498>
- Hisanari Otsu, Johannes Hanika, Toshiya Hachisuka, and Carsten Dachsbacher. 2018. Geometry-Aware Metropolis Light Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 37, 6 (2018), 278:1–278:11. <https://doi.org/10/gf2r3t>
- Hisanari Otsu, Anton S. Kaplanyan, Johannes Hanika, Carsten Dachsbacher, and Toshiya Hachisuka. 2017. Fusing State Spaces for Markov Chain Monte Carlo Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (July 2017), 74:1–74:10. <https://doi.org/10/gbxjs9>
- Hisanari Otsu, Yonghao Yue, Qiming Hou, Kei Iwasaki, Yoshinori Dobashi, and Tomoyuki Nishita. 2013. Replica Exchange Light Transport on Relaxed Distributions. In *ACM SIGGRAPH 2013 Posters (Anaheim, California) (SIGGRAPH '13)*. ACM, Article 106, 1 pages. <https://doi.org/10.1145/2503385.2503501>
- Jacopo Pantaleoni. 2017. Charted Metropolis Light Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (2017), 75:1–75:14. <https://doi.org/10/gfzq78>
- Christian Robert and George Casella. 2005. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag. <https://doi.org/10.1007/978-1-4757-4145-2>
- Gareth O. Roberts and Jeffrey S. Rosenthal. 2009. Examples of Adaptive MCMC. *Journal of Computational and Graphical Statistics* 18, 2 (2009), 349–367. <https://doi.org/10.1198/jcgs.2009.06134>
- Benjamin Segovia, Jean-Claude Iehl, and Bernard Peroche. 2007. *Coherent Metropolis Light Transport with Multiple-Try Mutations*. Technical Report RR-LIRIS-2007-015. Université Lyon, Lyon, France.
- Chris Sherlock, Andrew Golightly, and Daniel A. Henderson. 2017. Adaptive, Delayed-Acceptance MCMC for Targets with Expensive Likelihoods. *Journal of Computational and Graphical Statistics* 26, 2 (2017), 434–444. <https://doi.org/10.1080/10618600.2016.1231064>
- Martin Šik and Jaroslav Krivánek. 2018. Survey of Markov Chain Monte Carlo Methods in Light Transport Simulation. *IEEE Transactions on Visualization and Computer Graphics* (2018). <https://doi.org/10.1109/TVCG.2018.2880455>
- László Szirmay-Kalos and László Szécsi. 2017. Improved Stratification for Metropolis Light Transport. *Computers & Graphics* 68 (2017), 11–20. <https://doi.org/10.1016/j.cag.2017.07.032>
- Luke Tierney and Antonietta Mira. 1999. Some Adaptive Monte Carlo Methods for Bayesian Inference. *Statistics in Medicine* 18, 17-18 (1999), 2507–2515. [https://doi.org/10.1002/\(sici\)1097-0258\(19990915/30\)18:17<2507::aid-sim272>3.0.co;2-j](https://doi.org/10.1002/(sici)1097-0258(19990915/30)18:17<2507::aid-sim272>3.0.co;2-j)
- Miquel Trias, Alberto Vecchio, and John Veitch. 2009. Delayed Rejection Schemes for Efficient Markov-Chain Monte-Carlo Sampling of Multimodal Distributions. *arXiv preprint arXiv:0904.2207* (2009).
- Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Thesis. Stanford University.
- Eric Veach and Leonidas J. Guibas. 1995. Bidirectional Estimators for Light Transport. In *Photorealistic Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering)*. Springer-Verlag, 145–167. <https://doi.org/10/gfznbh>
- Eric Veach and Leonidas J. Guibas. 1997. Metropolis Light Transport. In *Annual Conference Series (Proceedings of SIGGRAPH)*, Vol. 31. ACM Press, 65–76. <https://doi.org/10/bkqj4>
- Jiří Vorba, Johannes Hanika, Sebastian Herholz, Thomas Müller, Jaroslav Krivánek, and Alexander Keller. 2019. Path Guiding in Production. In *ACM SIGGRAPH Course Notes*. ACM, Article 18, 77 pages. <https://doi.org/10.1145/3305366.3328091>
- Martin Šik and Jaroslav Krivánek. 2016. Improving Global Exploration of MCMC Light Transport Simulation. In *ACM SIGGRAPH 2016 Posters (Anaheim, California) (SIGGRAPH '16)*. Article 50, 2 pages. <https://doi.org/10.1145/2945078.2945128>
- Károly Zsolnai and László Szirmay-Kalos. 2013. Automatic Parameter Control for Metropolis Light Transport. In *Proceedings of Eurographics Short Papers*, M.-A. Otaduy and O. Sorkine (Eds.). The Eurographics Association. <https://doi.org/10/gf2r3s>

## A PSEUDOCODE FOR THE ALGORITHM

We summarize our DRMLT approach in Algorithm 1. Here, the function `MUTATE()` depends on the application, as shown in Section 5. `SAMPLEPATH()` maps a random number vector to a light path,  $\ell$  denotes the path luminance, and  $\xi \sim U[0, 1]$ . The `MUTATE()` function for our pairwise orbital mutation is described in Algorithm 2.

---

### Algorithm 1: Delayed Rejection MLT (DRMLT)

---

```

y ← MUTATE(x, _, 1), ℓy ← SAMPLEPATH(y)
α1 ← 1 ∧ [ℓy / ℓx]
if ξ1 < α1 then x ← y, ℓx ← ℓy
else
  z ← MUTATE(x, y, 2), ℓz ← SAMPLEPATH(z)
  η ← ℓz / ℓx
  if [Green and Mira 2001] then
    Γ2 ← Q2(x | y*, z) / Q2(z | y, x)
    y* ← z - (y - x), ℓy* ← SAMPLEPATH(y*)
    α* ← 1 ∧ [ℓy* / ℓz]
    α2 ← 1 ∧ [η Γ2 (1 - α*)] / (1 - α1)           ▷ Eq. (12)
  else if [Tierney and Mira 1999] then
    Γ1 ← Q1(y | z) / Q1(y | x)
    α2 ← 1 ∧ [η Γ1 (1 - α1)] / (1 - α1)           ▷ Eq. (6)
  else if [Pairwise Orbital Mutation – Sec. 4.4] then
    α2 ← 1 ∧ [0 ∨ (ℓz - ℓy)] / (ℓx - ℓy)           ▷ Eq. (11)
  if ξ2 < α2 then x ← z, ℓx ← ℓz
SPLATCONTRIBUTION(α1, α2, x, y, z)           ▷ Fig. 8

```

---



---

### Algorithm 2: MUTATE(x, y, stage)

---

```

while i < DIM(x) do
  if stage = 1 then
    d ~ QKelemen(εmin, εmax)           ▷ Eq. (9)
    φ ← 2πξ
    (yi, yi+1) ← (xi, xi+1) + d (cos φ, sin φ)
  else
    ▷ Pairwise Orbital Mutation – Sec. 4.4
    (u1, u2) ← (yi, yi+1) - (xi, xi+1)
    μ ← arccos(-u1 / ||u||)
    if u2 > 0 then μ ← 2π - μ
    θ ~ QWCauchy(ρ)           ▷ Eq. (10)
    μ' ← μ + θ
    (zi, zi+1) ← (yi, yi+1) + ||u|| (cos μ', sin μ')
  i ← i + 2
if stage = 1 return y else return z

```

---

Received October 2019; final version March 2020; accepted January 2020